

The job of the airline dispatcher can be greatly enhanced with a new tool for making diversion decisions that satisfy safety and schedule needs.

A SUPERIOR TOOL FOR Airline Operations

BY MICHAEL C. DORNEICH,
STEPHEN D. WHITLOW,
CHRISTOPHER A. MILLER,
& JOHN A. ALLEN

AN AIRLINE DISPATCHER IS RESPONSIBLE for determining how and where to divert aircraft from its destination if the destination airport is unable to accommodate the aircraft because of certain circumstances, such as adverse weather. The dispatcher needs to choose which of the arriving aircraft should be diverted and to which airports. These decisions have a significant impact on downstream airline operations as well as the aircraft, crew, maintenance, and passenger schedules.

Key problems in making diversion decisions are maintaining an awareness of the extensive – and frequently contradictory – set of concerns that various stakeholders have and assessing the impact of a choice on those concerns. The Diversion Off-Gate Management Assistant (DOGMA) is a decision support tool that mitigates this problem by maintaining, integrating, and assessing the impact on the various concerns, or policies, of the affected stakeholders. This article describes the DOGMA tool and the human-centered process used to develop it.

Airline Operations Disruption Management

The dispatcher is the link between the airline, air traffic control, and the aircraft. Dispatchers have the following primary responsibilities:

- a. Determine the route the aircraft will take between the source and destination airports.
- b. Ensure that the aircraft is airworthy to make the flight.
- c. Track the aircraft's location and status.



In fulfilling these responsibilities, the dispatcher needs to interact with the crew, air traffic control, aircraft maintenance, aircraft weight and balance, baggage loaders, refuelers, customer service personnel, airline scheduling, crew scheduling and tracking, airport gate managers, meteorology, and ground traffic controllers. During normal operations, this process moves smoothly and is one of the central mechanisms that ensures the airline operates in a safe and timely manner. However, the airline schedule is frequently disrupted by external events, such as unscheduled maintenance or weather.

Weather at and around the airport can be particularly disruptive in that it often affects a number of flights, causing airlines to delay takeoffs and/or force the diversion of in-flight aircraft. Diversion decisions consist of two parts: which of the in-flight aircraft are to be diverted, and to which airports they will be diverted. These two decisions can have dramatic consequences in the disruption of an airline's four inter-linked schedules:

1. *Aircraft fleet schedule*, which specifies the movement of each aircraft as it travels from city to city. The greater the aircraft utilization (e.g., amount of time it is ferrying passengers vs. not), the better the return.

2. *Crew schedule*, which specifies which flight crew and cabin crew staff each flight. The schedule must satisfy all Federal Aviation Administration rules and minimize the number of labor union contract violations.
3. *Maintenance schedule*, which ensures aircraft are at suitable maintenance bases for scheduled maintenance checks when (but not too long before) necessary.
4. *Passenger schedule*, which is intended to avoid passenger delays and missed connections.

There are other stakeholders in the diversion decisions. However, the diversion decision is made by only one decision maker – the dispatcher – and there is very little time available to produce a diversion plan (one dispatcher characterized it as “0–10 minutes”). The relevant information about how a candidate plan will affect various schedules and their stakeholders is distributed across multiple systems and departments.

Consequently, in current practice the dispatcher’s decision is almost solely based on fuel limits and other aspects of aircraft safety. Although safety should always be the *primary* concern, in many cases, fuel limits are the only criterion on which diversion decisions are based. As a consequence, there are typically several different diversion plans possible that will maintain safe flight and landing profiles, but these plans differ widely in their impact on airline operations, profits, crew and staff convenience, and customer satisfaction.

DOGMA, a Policy-Based System

The Diversion Off-Gate Management Assistant (DOGMA) system is a critiquing tool that, while retaining safe operating practices, goes further toward providing decision makers with information about the broad and diverse set of concerns from the various stakeholders affected by diversion decisions.

We call the goals and priorities of interested parties in the diversion decision their *policies*. By capturing various stakeholders’ policies and showing the implications of those policies to the dispatcher, we make it easier for him or her to integrate those interests into the decision-making process. This broader awareness of the various concerns in the decision is learned gradually over time and is rarely possible to do completely, in real time, for any given diversion decision. Thus, one impact of DOGMA is expected to be superior diversion decisions from less experienced dispatchers. Another impact should be better and more consistent diversion decisions, which translates into minimizing the impact of time-critical diversion decisions and increasing the airline’s ability to recover from severe schedule disruptions.

A policy is an abstract, general, a priori statement expressing a value or goal and some notion of the relative importance of that goal. In its simplest form, a policy provides a method for human operators to mathematically define what constitutes “goodness.” Once defined, a policy statement can be treated as a rule, which is evaluated against a current or hypothetical context; if the rule is true in the context, then the context incurs the goodness (or “badness”) points stipulated by the rule.

Alternative contexts (which could be tied to the expected outcomes of alternative decisions) can then be evaluated against one another by examining the set of policy rules that are satisfied or violated and the resulting set of goodness/badness points accrued. Similarly, the different (and sometimes conflicting) value statements of various individuals, organizations, or perspectives can be examined separately or can be combined via various mathematical aggregation schemes (e.g., averaging, weighted averaging, min/max).

A set of individual policy statements can be bundled together to flexibly define the priorities that apply in a particular situation, given that priorities can change under various circumstances. The domain of airline flight and dispatch operations represents a highly constrained system wherein the specific situation greatly affects the optimal strategy. Thus, decision support systems in this domain would benefit greatly from a flexible definition of priorities that are context-dependent.

What constitutes good or bad can change from situation to situation, company to company, and season to season. Again, although safety is always the most important goal in airline operations, there may be times and situations when some secondary goals, such as passenger on-time arrival, are less important than other secondary goals, such as crew duty limits, and vice versa. This fact is too often ignored in the creation of decision-aiding and resource optimization systems, leaving the operator to either slavishly obey the limited set of considerations the aid reasons over or to go through extensive mental work to interpret the aid’s recommendations in light of what really counts *today*.

The diversion decision is made by only one decision maker – the dispatcher – and there is very little time available to produce a diversion plan.

It is important to be able to flexibly redefine and reapply the definition of *good*. This is accomplished by separating out the definition of goodness and collecting it into bundles corresponding to various major contexts or situations. There is evidence that dispatchers already do this with their evaluation of current priorities, but they do it entirely mentally and, as anecdotal evidence shows, with limited success.

One such major context shift is that between so-called normal operations and preholiday operations. For instance, dispatchers told us that it is much more important for passengers to reach their destination (even if late) during the holidays, whereas during normal operations it is more important for the majority of passengers to be on time (even at the cost of some passengers not reaching their destinations). The table on the following page illustrates two bundles of such policy statements for airline operations, one representing the badness of various circumstances under normal operations

POLICY BUNDLES FOR NORMAL AND HOLIDAY OPERATIONS

Policy Statements	Normal Operations Bundle	Holiday Operations Bundle
Do not exceed crew duty limits.	10	8
Do not divert international connecting passengers.	8	7
Do not delay flights greater than 15 minutes.	8	3
Do not cause passengers to fail to reach destination (even if late).	3	8
Do not divert a flight with an unaccompanied minor.	10	10
Do not divert to an airport that has its maximum capacity of aircraft.	8	8
Do not divert a flight in a protected market.	4	1

and one representing the badness of those same conditions during preholiday periods.

The implication of these alternative policy valuations is that a given diversion situation or plan might be better or worse depending on aspects of context (e.g., whether the airline is in a preholiday period) – that is, depending on the policy bundle under which it is evaluated. For example (and without going into detail about the mathematical computations to be discussed later), the table implies that a dispatch plan that causes some passengers not to get to their final destination but which avoids violating crew duty limits and avoids flight delays would generally be a good solution under normal operating situations, but it would be exactly the wrong thing to do in holiday operations.

We decided to implement a policy system via a critiquing approach in which the dispatcher's diversion plan would be reviewed by a computer partner.

Because different bundles of policy statements can be created a priori (and at times when the dispatcher is not under high time pressure), they can be asserted at the time the decision is to be made (i.e., runtime). Earlier we illustrated policy bundling based on major operational contexts, but many other forms or dimensions of bundling are also possible and useful. For example, the set of concerns (i.e., policy statements) stemming from different operational units in the airline may be bundled and examined separately – allowing the dispatcher to see, for example, how good or bad a situation might be to marketing, to maintenance, and so on. Policy statements can even be used to reflect an individual operator's personal preferences.

In other words, which policy bundle is used to evaluate a situation can depend on the situation itself, on the user's

preference, or on what the user wants or needs to learn about the situation. The choice of which policy bundles to assert at what time may be left to the user or may be automatically triggered by situation monitoring and assessment. Finally, though it may be simplest to assert a single policy bundle at a time and avoid conflicting statements within the policy, we have also proposed techniques for resolving conflicting policies against each other and thereby allowing a diverse set of bundles to be examined simultaneously (Funk, Miller, Johnson, & Richardson, 2000).

Basis of Interaction Design

In this section we describe human-automation interaction principles that underpin the decision of the appropriate interaction design for a policy-based system in an airline dispatcher environment for the diversion management task.

After deciding to apply policy to the domain of airline operations diversion management, our next issue was determining how dispatchers should interact with a policy-based system. We considered having the system automatically generate diversion solutions that would subsequently be reviewed by dispatchers. However, previous research indicates that such traditional decision support systems suffer from brittleness (the system model does not account for all possible scenarios) and complacency (overreliance on system recommendations biases operators to not consider some factors and accept computer recommendation without adequate review; Parasuraman, Molloy, & Singh, 1993; Smith, McCoy, & Layton, 1997). Moreover, domain issues such as FAA regulations, the dispatchers' union, and organizational personality precluded this option as well. We also considered another common problem: Users of automated systems are reluctant to relinquish control to the automation (Miller & Hannen, 1999).

In light of these issues, we decided to implement a policy system via a critiquing approach in which the dispatcher's diversion plan would be reviewed by a computer partner that would offer feedback if there was a problem with the plan. This interactive critiquing approach has been shown experimentally to be an effective form of decision support

(Guerlain et al., 1999). Other research has indicated that interactive critiquing systems have less obtrusive interaction styles than traditional decision support, which results in greater user acceptance (Langlotz & Shortliffe, 1983).

The key need that has been identified by interviewees over and over is the ability to see the associated affects of diversion decisions.

Research by Kaber, Omal, and Endsley (1999) suggested that requiring dispatchers to construct their own solution enables them to maintain better situation awareness and minimizes complacency problems. The use of policy to critique the solutions makes available feedback that is relevant to their decisions only, thus minimizing the potential for information overload. Additionally, dispatchers are given feedback on the consequences of their decisions, which helps them assess the best strategies for minimizing downstream disruptions, allowing airlines to recover the schedule more quickly.

Rationale and Endorsement for the Policy Approach

Initial interviews with airline operations personnel identified the need for decision support tools to enhance the situation awareness of dispatchers. What was needed was a common data view across the airline (e.g., central operations, station operations), where typically data access is difficult and inconsistent across functional areas of operations. On the problem of diversion management, one airline supervisor said, "It's always been a problem. [In] 26 years, my entire career has been in dispatch, and there's never been a good way of managing [it]."

One typical example of the problems of situation awareness and lack of collaboration is coordination of alternatives by dispatchers, whereby during a major disruption, too many planes are diverted to a single station. One major airline's dispatch director summed up the needs as follows: "The biggest thing we can give the dispatchers is information: ...[so] they can make their choice based on better information, rather than just where is the flight coming from and where is it going." In addition to better information, dispatchers need tools at the operations end that enable the airline to recover from multiple diversions. A dispatcher said, "We know we're going to get hit, don't know when or where, but what you need are the tools to recover when you do."

The key need that has been identified by interviewees over and over is the ability to see the associated affects of diversion decisions. During initial formative evaluation studies of the prototype, we received near-unanimous endorsement for the application of policy to the diversion management domain as a way to quickly understand the effects of decisions on down-line operations and the ability of the airline to recover. The domain experts resonated to the idea of having stakeholders' voices present at the decision point. They felt that "the more people involved in a decision the better, both economically and for safety."

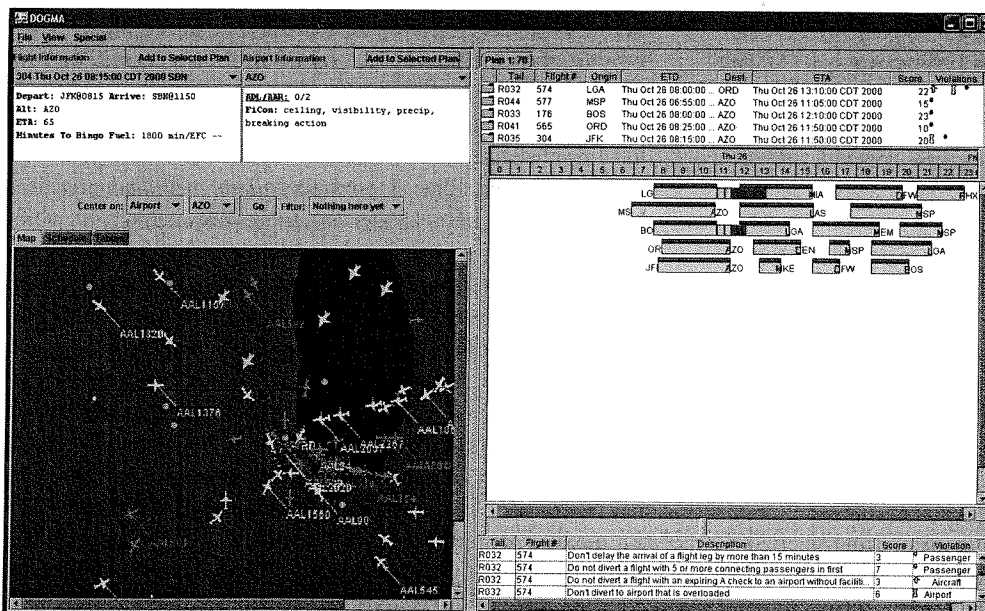
DOGMA integrates multiple information sources to improve dispatchers' situation awareness of the current state of flight, aircraft, maintenance, crew, and passenger schedules. Policy can capture the goals and priorities of all interested parties in the diversion decision, thereby integrating their interests into the decision-making process.

DOGMA Prototype

The DOGMA prototype, illustrated below, is divided into two principal spaces: (a) the information space and (b) the diversion plan workspace.

The information space, found on the left half of the interface, provides an integrated view of available information. The primary goal of the information space is to maintain the dispatcher's situation awareness across multiple information sources by allowing for rapid access to relevant information.

The lower half of the display provides a view of the available data, such as airline schedules, crew schedules, maintenance schedules, current aircraft position, and airport characteristics (e.g., current airport arrival rate). These data can be seen in multiple formats depending on task needs (map display, schedule view, tabular view). Users



Diversion Off-Gate Management Assistant.

can filter and sort information on various criteria using either predefined sorts (e.g., by tail number, time, arrival/departure airport, dispatcher, or fleet multiple criteria) or user-defined sorts (via a sort-query builder). The upper half displays detailed information about the user-selected aircraft and airport. Users can add selected aircraft or airports to the diversion plan workspace.

The right side of the interface contains the diversion plan workspace, where dispatchers construct diversion plans by selecting aircraft and deciding where they should be diverted. The upper third lists the candidate aircraft chosen for diversion. The user diverts the aircraft by changing its destination airport and its estimated time of arrival. Policy violation scores and categories are listed for each aircraft. Policy rules (like the holiday example in the table on page 20) are evaluated over instances in a context. For example, a plan might incur 10 badness points for every unaccompanied minor who gets delayed in a candidate situation.

Any other aircraft that suffer policy violations are also listed, to highlight dependencies that result in propagating policy violations from a single decision. The middle third of the diversion plan workspace displays the aircraft schedule of the relevant aircraft in order to highlight dependencies between aircraft, crew, and maintenance schedules. Thus, diverting Flight 123 may result in Flight 234 being delayed because of a lack of a crew if Flight 123's crew members were supposed to transition to Flight 234 when they arrived at the original destination.

The lower third of the display lists the policy violations, sorted initially by severity of the violation (the list can be resorted by category or flight). Dispatchers are presented with a set of policy violations relevant to each diversion plan. The associated penalties are simply added up to give a total score for the plan – though alternative weighting schemes are possible.

By viewing the policies, dispatchers are informed of which airline priorities the plan violates and to what degree and can choose to modify the plan accordingly. For instance, a decision to divert Flight 123 may violate the policy of "Do not divert a flight with an unaccompanied minor on board." The dispatcher did not need to know an unaccompanied minor was on board until that fact affected his or her diversion decisions. In this way, policy is used to present relevant information to a dispatcher only when he or she needs it.

The total policy violation score for a particular diversion plan is displayed at the top, on the tab for the particular plan. Multiple plans can be created (each with its own page accessible via a tab) and compared.

Policies are generated by allowing various stakeholders to express their priorities in the language of the domain. These policies are then integrated onto a common scale in predefined bundles, allowing airlines to express their priorities in different contexts. Policies can be updated as needed or as new contexts are identified. For example, as a result of what was learned during the airspace shutdown in September 2001, new terrorism-related policy bundles could be developed that allow airlines to preplan emergency response so they will

be able to recover more quickly during large-scale, airspace-wide disruptions. Another example of a terrorism-related policy bundle concerning airline safety may include defining diversion responses of hijacked aircraft to airports that are prewarned and have appropriate response capabilities.

Environmental Challenges to Implementing DOGMA

The events of 9/11 and its impact on the airline industry caused us to suspend our work in airline operations until a later date. As a consequence, we have not yet had the opportunity to install DOGMA within an airline's operations. However, by 9/11 we had made significant progress toward that goal and had jumped many of the technical and organizational hurdles related to delivering software to an airline dispatcher.

It is not enough to make the dispatcher's overall workload management better; a system must satisfy each individual interaction between the dispatcher and pilot.

The dispatch organization within an airline is an unusual one. Its members are employed and paid by the airlines but are accredited and regulated by the FAA. In this way, the dispatcher serves two masters, both of which can cause a dispatcher to lose his or her job. Although the airlines are interested in the smooth operations that generate revenue, they also understand the importance of aircraft safety. The FAA's interest is almost exclusively on aircraft safety. As a consequence, safety is the overriding concern of dispatchers, and they will accept nothing that might compromise this mandate. DOGMA bridged this gap by demonstrating that there was a range of diversion plans that differed greatly in their impact on down-line schedule disruptions, but all maintained high levels of safety.

The second issue that makes the dispatcher unusual is that he or she is co-responsible, with the aircraft pilot, for the aircraft itself. This is a relatively new role and one that the pilots have been understandably slow to accept. This newly developing trust between pilots and dispatchers has had several software impediments that in the past have caused disagreements between the two. Perhaps the largest of these deals with route planning. Both parties have access to systems that calculate time of arrival and fuel usage given a particular flight path. The dispatcher has the ground flight planning system, and the pilot has the flight management system in the cockpit. Because these systems use different data, they often come up with different answers that can be the basis of disagreement between the dispatcher and the pilot. As a consequence, the dispatcher is very sensitive to the needs of supporting and nurturing this relationship, and

so any supporting procedure or software must also support the needs and requirements of the pilot.

It is not enough to make the dispatcher's overall workload management better; a system must satisfy each individual interaction between the dispatcher and pilot. DOGMA addresses dispatcher workload by restricting feedback to relevant information through diversion plan critiques. By enumerating the down-line consequences of competing diversion plans, DOGMA provides the data on which pilots and dispatchers can jointly understand the pros and cons of particular decisions (or plans), in the context of both individual and system response to diversion disruptions.

Benefits and Impact

The policy-based DOGMA system can capture the goals and priorities of all interested parties in the diversion decision, thereby integrating their interests into the decision-making process. This broader awareness of the various concerns in the decision is only gradually learned over time. Thus, one impact of DOGMA is expected to be superior decisions from less experienced dispatchers. The use of policy enables the efficient, seamless integration of enterprise-wide goals at the decision point, thus enabling organizational control over decisions. Furthermore, policy systems can facilitate the propagation of high-level enterprise goals, such as customer satisfaction, down to the operational level via relevant policy feedback (e.g., "don't delay a passenger twice on a trip," or, more accurately, "delaying a passenger twice on a trip incurs 8 badness points"). Such a system can improve visibility into other stakeholders' priorities, thus minimizing so-called bunker mentality within departments of an organization.

Another impact of this increased awareness and broader input into the diversion decision should be better and more consistent diversion decisions that minimize the negative impact of diversions and improve an airline's ability to recover from severe schedule disruptions. Moreover, critiquing reduces workload by providing feedback about only those policies that are violated by a user action. The system can further reduce workload by enabling a quick comparison of diversion options by providing a simple metric: total policy penalty points on a selected flight. This approach can also accommodate both novice and expert dispatchers. The interaction does not interfere with an expert's workflow unless a decision violates a policy; and policy feedback provides learning experiences for novice dispatchers. Consistent outcomes are also ensured by the systematic and automated evaluation of user actions.

References

- Funk, H., Miller, C., Johnson, C., & Richardson, J. (2000). Applying intent-sensitive policy to automated resource allocation: Command, communication and most importantly, control. In *Proceedings of the Conference on Human Interaction with Complex Systems* (pp. 179-183). Los Alamitos, CA: IEEE Computer Society Press.
- Guerlain, S., Smith, P. J., Obradovich, J. H., Rudmann, S., Strohm, P., Svrbely, J., & Sachs, L. (1999). Interactive critiquing as a form of decision support: An empirical evaluation. *Human Factors*, 41, 72-89.
- Kaber, D. B., Omal, E., & Endsley, M. R. (1999). Level of automation effects on telerobot performance and human operator situation awareness and subjective workload. In *Automation Technology and Human Performance: Current Research and Trends* (pp. 165-170). Mahwah, NJ: Erlbaum.
- Langlotz, C. P., & Shortliffe, E. H. (1983). Adapting a consultation system to critique user plans. *International Journal of Man-Machine Studies*, 19, 479-496.
- Miller, C., & Hannen, M. (1999). The rotorcraft pilot's associate: Design and evaluation of an intelligent user interface for a cockpit information manager. *Knowledge Based Systems*, 12, 443-456.
- Parasuraman, R., Molloy, R., & Singh, I. L. (1993). Performance consequences of automation-induced "complacency." *International Journal of Aviation Psychology*, 3, 1-23.
- Smith, P. J., McCoy, E. C., & Layton, C. (1997). Brittleness in the design of co-operative problem-solving systems: The effects of user performance. *IEEE Transactions on Systems, Man, and Cybernetics - Part A*, 27(3): 360-371.

Michael C. Dorneich is extensively involved in aviation-related projects at Honeywell Laboratories, 3660 Technology Dr., Minneapolis, MN 55418; michael.dorneich@honeywell.com, including developing flight deck alerting and notification systems, developing critiquing tools to help collaborative decision making in airline dispatch, and developing frameworks for classifying aircraft system failures. His expertise is in the application and development of cognitive systems engineering and human factors design techniques to collaborative system design, including interactive learning environments, adaptive information systems, and decision support systems. Stephen D. Whitlow, Honeywell Laboratories, works on human-automation interaction design issues, decision aid development, and the integration of neurophysiological feedback into adaptive systems. Christopher A. Miller, chief scientist, Smart Information Flow Technologies, has spent more than 15 years designing methods for humans to interact with complex automation and information systems. Smart Information Flow Technologies is a small research and consulting firm based in Minneapolis. John A. Allen, Honeywell Laboratories, conducts research on planning, scheduling, machine learning, and tactical control. The authors acknowledge Jim Richardson, Jim White, Jon Schewe, Ranjana Deshpande, Rahima Kaba, and Trent Reusser for their contributions to this project. ■■