Optimizing Group Utility in the Collaborative Decision Making Process

SYST 698 – Dr. Schum

Andrew J. Ryan

Introduction

Since 1960, air passenger traffic has increased at an average yearly rate of 9% [Donohue]. Because of this steady increase in demand, the airways have become increasingly more congested. While congestion in the airways is easy to deduce form this scenario. Another entity, the airport, is also affected. In its purest form, it is easiest to envision a flight in three stages: take-off, en-route, and landing. The airport dominates two of these three components and for that reason plays a large role in how traffic flows across the National Airspace System (NAS). Accordingly, the Federal Aviation Administration (FAA) has two separate organizations to monitor these functions. Air traffic control, which ensures safe separation between aircraft, and traffic-flow management (TFM), which balances demand and capacity to maintain safe and efficient traffic flow [Chang et al.]. TFM is tasked with the responsibility of minimizing the interruptions to the NAS, so that the available capacity is utilized by existing demand. When demand exceeds capacity, they must make decisions and take actions to create an optimal situation for NAS users. This paper describes this process and makes suggestions on how decision-making can be improved to provide better results.

Background

Unlike EuroControl (the FAA's European counterpart), flight scheduling in the United States operates under a sunny-day paradigm. Meaning, they assume there will be no environmental changes to the system to disrupt service. Meaning all aspects of the system from weather, to facilities, to human resources (i.e. controllers), will function as advertised and allow the total system capacity to be utilized. This utopian outlook works on most days, however when one or more of these resources is not at peak performance, the schedule is disrupted causing utilitarian decisions to be made across they system in deciding how to best cope with the situation.

To better deal with such conditions, in the early 1980s, the FAA began issuing Ground Delay Programs (GDP). The main purpose of a GDP is to allow planes to take their delay on the ground as opposed to in the air. By doing so, airlines would be able to save fuel (otherwise used doing airborne holding) and also increase safety (too many planes near the terminal area at an airport). Based on an airlines schedule, when demand at an airport surpasses capacity, flights inbound to that airport, which had not yet taken off, would be issued a delay to decrease the demand on the airport. In simplest terms, they were displacing traffic to allow the demand to equal capacity.

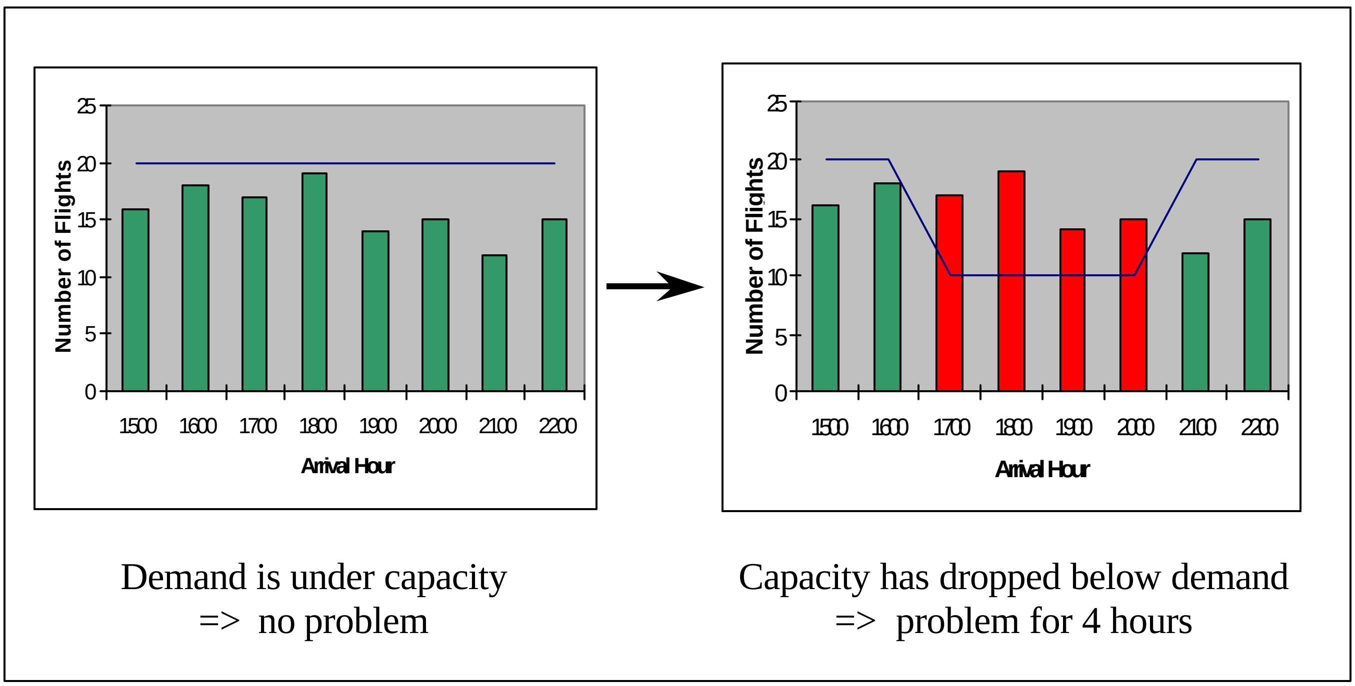

Figure 1: Demand surpassing Capacity brings about the need for a GDP

The Collaborative Decision Making Initiative

At a landmark meeting in 1993, stakeholders of the NAS sat down and aired their grievances. The result of this meeting was the birth of the Collaborative Decision Making Ground Delay Program – Enhancements project. The most important aspect of this program was the agreement by the airlines and FAA facilities to work together (in collaboration) to make decisions based on up to the minute information. “Rules of Engagement” were also developed which outlined the roles each entity would play in the process and the limitations of the influence they carry in the game.[1] An estimate by CDM participants calculated the CDM project saved the airlines $39,574,000[2] from January 20, 1998 to July 15, 1999 [Ball et al.]. From its inception, the plan was to look toward new ways to increase savings, “As the prototype stabilizes, our collective efforts will focus more and more on capturing benefits [Wambsganss].

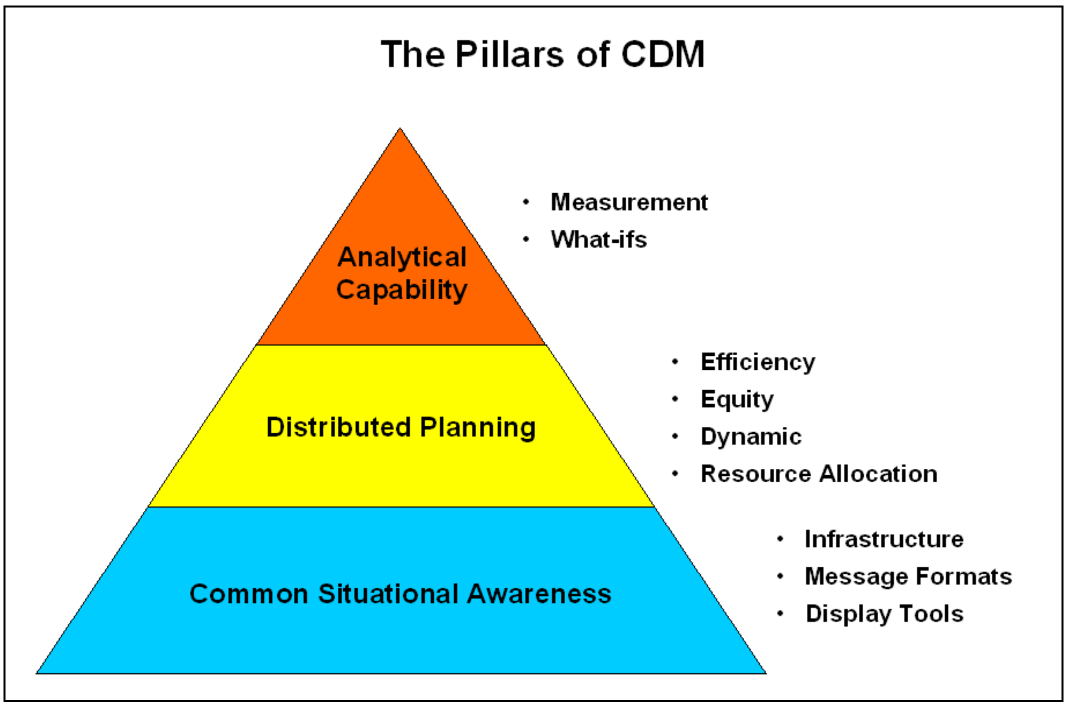

Figure 2: Functional Areas of CDM [Wambsganss]

As shown in Figure 2, distributed planning is one of the main areas of CDM. The collaboration among different stakeholders is central to arriving at a conclusion that is amenable to all parties involved. As described in “Collaborative Decision Making in Aviation Transportation,” “The overall objective for collaboration is to improve the efficient usage of scarce NAS resources [Kollman et al.]. While

[1] It is interesting that game theory was used in the early stages of CDM, as this paper reveals, there are more applications of game theory that can be applied to this area.
[2] This figure was calculated using the conservative estimate the each minute of delay mitigated by CDM saved the airline $25.

acceptable as a high level goal, ignored is the fact that each stakeholder has a unique set of low level goals which they tend to before this high level (system) goal.

In the “Decide, Announce, Defend” model which best describes the FAA’s current thinking, a one-tiered decision goal is perfectly acceptable. However, with the varying objectives of stakeholders, a new outlook is needed. Clearly, such a complex system needs a richer decision making process. As described in *Wharton on Making Decisions*, [there are] three different levels [of decision making] – what should be done based on rational theories of choice (normative models), what is actually done by individuals and groups in practice (descriptive behavior), and how we can improve decision making based on our understanding about differences between normative models and descriptive behavior (prescriptive recommendations). This paper attempts to make prescriptive recommendations on the CDM process, using classical decision and game theory to develop a dedicated decision support aide to compliment existing CDM tools.

Group Decision Making Explored

Group decision-making is best defined as a collection of individuals having conflicting interests that must be resolved where each decision maker has a unitary interest motivating its decisions. Robert Scholl of the University of Rhode Island further categorizes different methods of group decision-making. The CDM process closely resembles the Individual Consultative Style. Here, the leader defines the problem and shares this definition with individual members of the group. Their advice is solicited and alternative solutions are generated. The leader then makes the choice of which alternative solution to employ[3][6]. Attempting to fuse an IT solution into this process is a necessary step in the maturity of CDM.

[3] It must be noted that in the CDM process, when a GDP is being planned, the ACSCC traffic flow specialist has 51% of the power – making them the leaders.

Insufficient information technology for decision support can lead to disastrous outcomes. In *The Collapse of Barings*, Stephen Fay comments on one of the causes that led to Barings Future Singapore (a financial management company) to lose over one billion dollars: " . . . the management at Barings failed to look at how technological advancements, such as computer-assisted programs could have prevented [this]" [Fay]. While the FAA has employed automation and computer-aided technologies, their scope has been limited. That is, they have not gone beyond the capability of integrating and displaying information. Technology is available which can aide in the decision-making process without imposing an increase in cognitive workload. Figure 3 shows the levels and stages of automation. Currently, TFM technology is on the low end of all three stages.

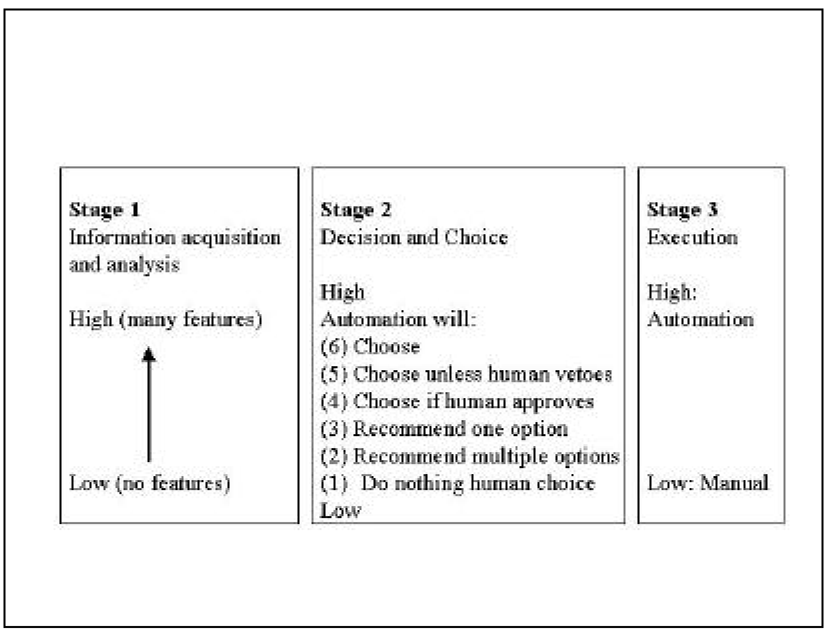

Figure 3: Stages and Levels of Automation [Wickens and Hollands]

Stage one refers to the amount of "work" that automation takes on to replace or assist the human in integrating/comprehending information. Stage two encompasses the degree of constraints that automation imposes on choices (i.e.

filtering out decisions that will lead to poor outcomes) by the human operator. Finally, stage three measures the amount of “work” needed to execute the decision. While a fragile domain such as air traffic control should not be operating at the highest end of this scale, there is definitely opportunity to make better use of the middle vicinity, by assisting the decision-maker where possible.

Adding Intelligence to CDM

While the Barings collapse occurred in the early 1990s (pre-Information Age), it is inexcusable for an information-laden society to ignore the potential uses of technology in the decision making process. Yet, what we propose in this paper is not a wholly technological solution. Technology must be used to support, not replace, the human in the decision making process. Both the human expert and the automated model have their strengths and weaknesses. These differences are highlighted in Table 1

Table 1: Comparing Experts and Models in Decision Making Tasks

[Hooch and Kunreuther]

Where Models Excel and Experts Fail	Where Experts Excel and Models Fail
Bias	*Subjectivity*
➢ Experts are subject to bias of perception and evaluation while models are deterministic in their interpretation of a problem	➢ Experts are proficient at attribute-valuation and provide evaluations of variables that are difficult to measure objectively. Models can only operate on the data provided.
Overconfidence	*Granularity*
➢ Experts suffer from overconfidence and may be influenced by organizational politics that encourage strategic responses; models are immune to social pressures for consensus	➢ Models know only what the expert tells the model builders; experts know what questions to ask and can garner further information if needed.
Vigilance	*Flexibility*
➢ Experts get tired, bored, and emotional, models do not	➢ Experts have highly organized, domain specific knowledge. They may have superior pattern matching skills which allow them to recognize and interpret outliers, while models tend to omit these cases.

Sociological issues play a larger role in group decision-making than individual decision-making. Issues like “conformity,” “availability,” “representativeness,” and “anchoring and adjustment” tend to skew the manner in which an individual performs in a group environment. Agarwal and Prasad point out that the size of the group, its cohesiveness, the particular task that confronts the group, the personal utility and preference functions are also contingencies that can affect the performance of the group. In taking this a step further, it is important to realize the level of experience of the individual members in the group. While an expert is able to pattern match from years of experience, a novice user may only have a few weeks experience to draw from.

When placed in a decision-making environment, the differences between these two actors are magnified. As researched by Clawson et al. and Niederman and Volkema, experts tend to focus on flexibility; they select some previous case from their personal collections and adapt it during the decision process. Conversely, novices tend to spend more time gaining a conceptual view about the process (through gathering more information) before responding to it. While both are worthwhile, efforts, why not combine the two? This is precisely the reason why meteorologists do well in their decision making process. They benefit from networked databases that contain forecasts from their peers as well as historical information on what happened in the past. Moreover, they receive daily feedback and can make adjustments in their calculations based on the strength/weakness of the prior days’ outcome [Raiffa].

The proposed method accomplishes this two-pronged approach while also adding intelligence to the process. CDM has been in existence for nearly eight years. In that time hundreds of people have been involved in the process (as operators, researchers, management) and have somehow remained an untapped source for knowledge. That is, the collective experience of those most intimate with CDM should be used as a means to advance the concept. The field of knowledge management has boomed in the last five years, primarily focusing on that subject. However, for our needs, leveraging the existing knowledge (lessons learned, heuristics, ‘best practices’) and embedding them into CDM tools, creating an element of “perfect recall”, is an essential step to creating the next generation CDM model.

Rationality

The proposed tool will be able to provide organizational learning, perfect recall, merged utilities, Bayesian inference, and closed-loop feedback. Emphasis will be placed on suggesting alternatives, not giving a final answer. The issue of rationality, which has been a very elusive term to define, is also important. Herbert A. Simon is probably the most significant contributor in the area of rationality. However, giving way to a context specific definition, Watson and Buede comment, “ . . . rationality has a lot to do with consistency . . . we are rational when, having adopted rules which our statements or actions should conform to, we act in a manner that is consistent with them.” A stakeholder is deemed rational when their decisions reflect an attempt to increase their position (i.e. meet their goals). Some decision support systems take this idea literally and implement models which are too rational. Due to the computational complexity of some problems, there are limitations on the quality of the decision a group can deduce. It is here that heuristics come into play.

Heuristics and Satisficing

The use of heuristics in decision-making has been around since the beginning of time. Heuristics are unscientific methods invoked in an attempt to formulate a response. We tend to gravitate towards heuristics as they minimize the cognitive resources needed to come up with a suitable answer. Instead of plugging through a full set of alternatives, heuristics allow us to identify a feasible solution. Economist Herbert Simon defines this as satisficing – ‘getting a result that is good enough.’ This often leads less than optimal outcomes as each satsificed decision decreases the expected value of its descendant. A final drawback of satsificing comes in the area of tradeoff evaluation. In an attempt to save cognitive effort decision makers will forego the process of weighing different decision attributes against one another. This is particularly damaging [as] explicit tradeoffs between attributes are generally acknowledged as the gold standard for accurate decision processes [Hoch and Kunreuther].

The third area where humans tend to satisfice (closely related to trade-offs) is integrating information. Humans process information sequentially, taking one piece at a time and until all pieces have been processed. This is the reason why in algebra, we are always told to solve for one variable at a time. Unlike the multi-process, high memory computer, we can only maintain between five and nine pieces of information at any point in time [Miller]. When dealing with a multi-attribute problem, it is clear to see why we tend to be selective in the information we choose to process as opposed to being comprehensive in our perception of the problem.

Bayes and Decision Making

Thus far, the focus has been on the qualitative shortcoming of human decision-making. Bayes offers a more concrete analysis. Humans are usually slow to react when evidence is introduced which changes the probability of an event occurring. A classic example of this phenomenon can be found in Kahnerman and Tversky.

A CDM related example of Bayes theorem could be easily constructed:

> *Stakeholders have decided to implement either a 12-West or a 10 West ground delay program over SFO. The specialist advises that in either case, there is an 80% chance that the program will be extended (increased in length) to cover the afternoon rush. Naturally, the airlines would not like to see this, as their schedule will be disrupted in a greater way if a program extends into the afternoon rush. However, empirical data, from programs run at that airport under similar conditions reveal, in a 12 West program, the*

chances of an extension adding more than 15 minutes of delay to all flights (caught in the GDP) is 40% while it is 15% for a 10 West.

According to Bayes, the probability of an event E, assuming some evidence F is symbolized by P(E|F). So, if E is the 10 West Program being extended and F is the evidence of the extension actually occurring, we can determine the probability of an event given some evidence using the following equation:

Equation 1: Calculating New Odds Using Bayes Theorem

Posterior Odds = Base Rate * Likelihood Ratio

OR

$$(F|E) = [(E)/(E')] * [(E|F)/(E'|F)]$$

Here, the base rate is the initial rate given by the specialist and the likelihood ratio is the rate given the evidence. It has been proven [Kahnerman and Tversky] that most people stay too grounded to the base rate even when presented with new evidence. In calculating the posterior odds (the new ratio given the evidence) new probabilities[4] are established for both the 10 West and 12 West scenarios. The new probability for a 10 West program is 41%, while in a 12 West it is 72%. Armed with this new information, airlines should choose the 10 West program. Currently, there is no mechanism in place to relay new probabilities to stakeholders.

[4] Once the posterior odds are calculated, we convert from ratio form to a decimal to establish the probability.

Why a Group Decision Support System?

Group Decision Support Systems (GDSS) have been researched as possible aides in group decision-making since the early 1980s. Due to technological limitations, much of the work was contingent upon technological advancements, "though our understanding of group processes is fairly sophisticated, the technology to support these processes tends to be basic in nature" [Agarwal and Prasad]. This is no longer the case. Today, we enjoy the fruits of the artificial intelligence research done in the last thirty years. Providing intelligent support via some electronic medium is well within reach.

Oftentimes, the term 'intelligent' when used in tandem with technology invokes the connotation that we are trying to replace people with computers. This is wholly inaccurate. Our goal is to create a system to support and compliment human decision-making in an attempt to increase the total utility of the outcome. Although definitions differ, very few researches claim that DSS can replace the human decision maker:

- Keen suggests, "[DSS] supports, rather than replaces their (manager's) judgment. The overall aim is to improve the effectiveness of their decision-making."
- Morton purports, "The terms decision support systems' DSS refers to the use of computer systems, often interactive, to support humans as they make certain types of partially structured decisions."
- Adelman offers, "The basic assumption is that, in its most general sense, DSS refers to any computer software designed primarily to support the decision making process by assisting decision makers in thinking about the various aspects of the decision problem(s) facing them. "

Implementing a decision support system will aide in the formulation of decisions as they help facilitate thought. Psychological research on decision making shows that unaided, people use simple decision making heuristics that typically violate theoretic axioms and often result in sub-optimal decision making [Adelman].

In keeping with the above guidelines, our system will have the following properties:

1. Understand the cognitive limitations of its users and not attempt to exceed it (bounded rationality).
2. Handle uncertainty exquisitely and prompt users when they are acting against Bayes' Rule.
3. Improve interaction within the group while building a profile of system users.
4. Generate 'canned' responses to situations from a database of past scenarios.
5. Provide feedback on the decision making process to keep user suggestions in line with their initial preferences.

Based on the above objectives, an agent-oriented approach would seem optimal. However, implementation is beyond the scope of this paper and will not be explicitly discussed. In the next section, we present a high level architecture for the system.

The Technology Aspect

Agent programming is an emerging paradigm within the software development community. Researchers have been plugging away for the last ten years in an attempt to create a paradigm which allows for learning, proactive response and autonomy in a programming language. A major benefit of agent-oriented programming (AOP) is that each agent can operate in a defined environment and assume full responsibility for the monitoring of its domain. Unlike the event-driven methodologies (i.e. object-oriented or procedural) that are currently popular, agents can react to a situation in real-time and 'learn' from its past responses to act in an efficient manner and issue an optimal solution.

It is important to pause and define what is meant by 'learning' and 'optimality'. One of the flaws of intelligent systems built to date is that they do not behave like humans. That is, they process inordinate amounts of information and ultimately issue an answer that while optimal, is incomprehensible to a human user. A glaring example of this is IBM's supercomputer Big Blue that beat World Chess champion Gary Kasparov. Big Blue was fed in hundreds of strategies and was able to use this unlimited data set to defeat the world's best chess player. While great for exhibition, in a real world context, automation must be within cognitive reach of the human user. Therefore, our agents will have the following characteristics:

- Do not have perfect information about the environment
- Do not have a perfect model of the environment
- Have limited computational power
- Have other resource limitations (e.g. memory)

In imposing these limitations, we attempt to adhere to Simon's idea of bounded rationality[5]. While our model searches for as much veracity as is possible, limitations such as time, cost, and technique will be taken into account when agents encounter a problem [Edmonds]. While the idea of bounded rationality and agents would be new to aviation research, Ljunberg and Curmi et al. have both yielded positive results in using agent technology to model aviation processes.

[5] Simon also defines procedural rationality as rationality without bounds. In order to prove this concept to the user community, both bounded and unbounded agents will be deployed to illustrate the performance gap between the two. Once the problem is fully understood, a medium can be decided upon where both ideals are represented.

Learning

In terms of learning, we again turn to the Simon's bounded rationality. IBM's Big Blue was fed information, thus making it a top class chess player overnight. When modeling human behavior, such an approach is unsuitable. Learning must take place over time and adapt to new situations. Such an evolutionary model mirrors the way humans learn and allows for greater acceptance by the user. A model that develops with the user is much more effective than one which is already beyond the cognitive capabilities of its user.

Figure 4 describes the system graphically. The agents begin with a given a priori body of knowledge (i.e. FAA rules, guidelines). Based on the input, a set of scenarios will be chosen with the closest matching the input selected for further consideration. Elaborated further, each agent has a finite memory which holds past data (scenarios) and the effects of its actions. According to the specificity of the model (how well the model and current situation match), the best model is chosen. Next, it uses that model and its goals/preferences to determine its action. Finally it takes that action and tracks the results in the environment for future use.

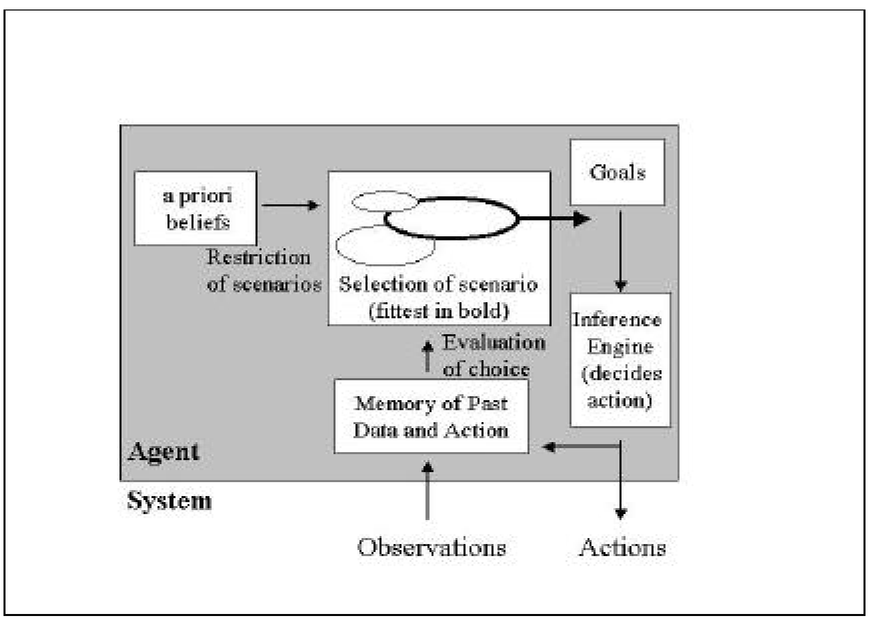

Figure 4: High Level Architecture of Proposed System

An exciting feature of this model is the evolutionary process in which it matures. With limitations placed on the amount of contact an agent can have with the environment and the amount of time it takes to reach a decision (no unlimited processing), the model grows with the user.[6]

About the Problem

On any given day, there are 40,000 flights operational in the National Airspace System that are under the auspices of the FAA. While most are familiar with Air Traffic Management (ATM), the cornerpost to cornerpost operations that control flights en-route, equally important is the Traffic Flow Management (TFM) arm of the FAA. While both organizations main concern is safety, ATM is responsible for separation of flights in the air, while TFM monitors the flow of traffics across the system. This includes making sure demand at an airport does not exceed capacity, and also monitoring the rate of traffic as it move across a region. Their work is heavily impacted by environmental disturbances such as inclement weather, runway closures, and facility abnormality (i.e. loss of radar, shortage of staff etc). Whenever one of these issues impacts traffic flow, a traffic initiative must be enacted. There are three initiatives which are most often employed:

1. **Expanded Miles in Trail (MIT)** – While en-route, the minimal distance that can separate two airplanes is five nautical miles. By increasing this number, flights are fed into the airport at a slower rate, thus decreasing the hourly rate without explicitly impacting an airlines schedule.
2. **Ground Stop (GS)** – In severe instances (thunderstorms, disabled plane on a runway, high airborne inventory), a ground stop is imposed which freezes all outgoing traffic from an airport. While airborne traffic may still land, any flights that are yet to depart are displaced to a later hour. Ground stops usually last about an hour.

[6] While exciting from a research point of view, it will be left to the user community to decide if they prefer an evolutionary model or a finalized product in an operational setting.

3. **Ground Delay Program (GDP)** – Used when demand at an airport exceeds capacity. Enacted over a multi-hour period, a GDP sets an hourly rate at an airport that caps the number of flights which can land. Excess flights are given a delay that allows them to land at a time that the airport can accommodate them. A properly coordinated GDP can save airlines thousands of dollars as the delay is absorbed in the ground (as opposed to in the air) thus saving them fuel costs or even worse the possibility of a diversion.

While each initiative has a different impact on traffic flow, there are other repercussions involved. Airline schedule integrity is first and foremost. Airlines keep very tight schedules as they attempt to get maximum usage out of their planes. Moreover, connecting flights are based on their original (non-delayed) schedule. A hiccup in schedule means passengers will miss their next flight.

Another view is from the en route controller. A GDP generally means they will receive a decreased volume of traffic. Decreased volume means less workload.

The en route center also has a position when a GDP is enacted. The NAS is broken up into 20 contiguous operational centers. Depending on the airport the GDP is being issued, various sectors feel a backlash. For instance, if a GDP is issued at one of the New York area airports (EWR, JFK, EWR), while they are all located in ZNY (New York center), ZDC often feels a backlash as flights are slowed down and the delay propagates into its airspace. Conversely, ZNY sees a decreased flow of traffic yet is faced with congested sectors as the airport cannot accept all of its allotted traffic and airborne holding is utilized to keep the terminal area from becoming overcrowded.

The airport tower and approach is also involved. Their main concerns involve workload and the inability to handle flights as they arrive in the terminal area.

Once a plane is headed their way, they have no other choice but to control it, accepting more planes than they can control makes their job especially rigorous.[7]

Table 2 shows the basic objectives of each of the stakeholders:

Table 2: CDM Stakeholder Preferences

Stakeholder	Main Objective
Approach	Maintain workload limits based on radar thresholds (the number of planes they can handle within a given time frame)
ATCSCC	Keep flow evenly distributed in NAS
Major Airlines	Keep as close to original schedule as possible (economic)
Operational Center	Keep pressure on arrival fixes without saturating a sector
Tower	Maintain workload limits based on visibility

There are also a number of variables involved in the GDP process:

Scope – The scope of a program indicates its geographic spread. As previously mentioned, the NAS is split up into 20 centers and the wider the scope, the more centers (hence more flights) will be impacted by the GDP.

Duration – The average GDP lasts about four or five hours. The longer a GDP, the more flights involved and the more stable the airport demand will become.

Airport Arrival Rate – The most important variable. This is the agreed upon rate at which flights will be allowed into the airport (per hour). Each airport has a normal rate (the rate it usually operates at), and the AAR for a GDP will be below this level. A higher rate allows more planes to land and decreases overall delay.

[7] When this occurs, they will send planes into airborne holding to decrease the volume of planes coming into the airport.

Delay (minutes) – While not an explicit metric, there is a way to cap the maximum delay given to a particular flight. While airlines view delay minutes as a prime indication of the impact a GDP has had on their operations.

Equity – Measures fairness in a GDP. With several airlines operating flights into an airport, allotting delay unequally puts the more-delayed airlines at a competitive disadvantage.

Table 3 depicts the relationship among these variables. A zero (0) denotes there is no relationship between the variables, while plus (+) or minus (-) signifies a positive or negative association. Therefore, the first entry (++) for *scope* and *duration* indicates the larger the scope of the program, the longer the duration of the program. Understanding these relationships is an important first step to building a system that can generate intelligible resolutions when a GDP is needed.

Table 3: Relationships between GDP Variables[8]

	Scope	Duration	AAR	Delay	Equity
Scope	x	++	++	+-	++
Duration	++	x	+-	+-	0
AAR	0	+-	x	+-	++
Delay	+-	+-	+-	x	0
Equity	++	0	++	0	x

Table 4 explains the preferences of the actors in each of these variables. An 'X' denotes indifference.

[8] These relationships were derived from informal interviews held with FAA personnel.

Table 4: CDM Stakeholder Preferences

	Airlines	Approach	ATCSCC	Center	Tower
Scope	+	I	+	+	I
Duration	+	+	+	+	+
AAR	+	+	I	+	+
Delay	-	I	I	I	I
Equity	+	I	+	I	I

The Coordination Process

When issuing a GDP, the ATCSCC is in control of the process. They initiate a conference call which connects the other stakeholders. The constraints (adverse weather, runway closure etc) are discussed and the ATCSCC specialist solicits resolutions from stakeholders. From that point, the lines are open as the other stakeholders chime in on the proposed GDP. In group decision-making circles, this would loosely be termed the 'Delphi process.' Developed by the RAND Corporation in 1968, "the Delphi technique is a method for soliciting and collating group judgments on a particular topic through a set of carefully designed sequential questionnaires interspersed with summarized information and feedback of opinions delivered form earlier responses [Van de Ven]."

In our application we attempt to satisfy the following objectives:

1. To determine a range of possible alternatives based on selected parameters available when issuing a GDP
2. To seek out information which may generate a consensus
3. To correlate informed judgments on topics related to a GDP
4. To educate the group (provide feedback) as to the diverse and interrelated aspects of the topic.

Why not Intelligence Sooner?

Many in the user community will wonder why this type of research is needed. Without calling in the overused excuse of "the users never embrace change," this section is dedicated to pointing out some of the properties of this problem that contribute to an appearance of tranquility.

1. *The current system has no feedback loop* – without a baseline for what should have happened, users are never forced to evaluate their performance. Once users have experienced this cycle, reinforcement sets in and it becomes more difficult to break them of their poor habits.
2. *The system is not fully understood* – CDM has been around for less than ten years. Proper metrics and benchmarks have been slow to emerge. As a result, 'best practices' or optimal strategies[9] have not been developed.
3. *The task environment is forgiving of mistakes* – Very few industry leaders would shut down CDM and return to the old ways of traffic flow management. The benefits of CDM have been statistically proven, but the optimal solution is still out there. Learning by trial and error has been seen as the remedy for this. While not totally inconceivable, how long can the same model be picked, pulled, and prodded, before a paradigmatic change is needed?

Frames

The first step in the solution is to get all stakeholders to operate out of the same frame (or point of view). That is, currently each of the stakeholders sees the problem through their own lens. Each of the goals is defined in the metrics they are most familiar with (AAR for a tower, minutes lost for an airline), with no serious attempt made to take a global view of the situation. Daniel Kahnaman and Amos Tversky popularized the concept of decision frames by demonstrating

[9] While game theory seems to lend itself to this domain, it has been the airlines, and not the FAA who has used 'gaming' to selfishly satisfy their needs.

that changes in surface features of a decision (i.e. making the collaborative decision be seen as a gain for all parties as opposed to a loss for some) can alter choices. Implementing a multi-frame solution will be of great benefit in this area.

While a grand goal, there are many hurdles one must cross before making this a reality. While we have already defined the stakeholders and their preferences, we have not attempted to merge these desires to create a holistic view of the situation. Let us embark in this journey:

The first step in merging preferences is to have a full understanding of what each stakeholder actually wants. Assigning a numerical index to alternatives for each user allows inconsistencies in decision making to be easily pointed out. Written formally:

For any A_i and A_j, either $A_i \gtrsim A_j$, or $A_j \gtrsim A_i$, and if $A_i \gtrsim A_j$ and $A_j \gtrsim A_k$ then $A_j \gtrsim A_k$.

Going back to the definition of rationality, when one chooses against their previously stated goals (or acts intransitively), this should be pointed out, as they are not acting to increase their expected value. To indirectly cite the founding fathers of game theory, John von Neumann and Oskar Morgenstern, we are operating off of their presumption that if a person can express preferences between every possible pair of gambles, then one can introduce utility associations in such a manner that if the person is guided solely by the utility expected value, he is acting in accordance with his true tastes [Luce and Raiffa].

To be technically accurate in our description of the process, a social decision is what we are chasing. As prescribed by Luce and Raiffa, there are three entities that need to be defined and then incorporated into a social welfare function:

i. *Alternatives.* Let α = {x, y, , z} be a set of alternatives.

ii. *Individuals.* Let the individuals of the society be denoted by 1, 2, , i, , n.

iii. *Preferences.* For each individual I, and any alternatives u and v, one and only one of the following holds:

a. "i prefers *u* to *v*," which is written as uP_iv,

b. "i prefers *v* to *u*," which is written as vP_iu,

c. "i is indifferent between *u* and *v*," which is written as uI_iv,

iv. Design a social welfare function which will assign a rule which associates to each profile of preference orderings (n-tuple of orderings off all of the possible alternatives for each player) a preference for the society itself.

There are many challenges in establishing this social welfare function (swf). A few are listed below.

The Decision Maker as Amalgamator

Our model is based off the premise that the ATCSCC specialist is an amalgamator of preferences. That is, the final decision will attempt to maximize the total preferences of all individual decision makers. The following assumptions are made:

1. The decision maker's preferences (for a given scenario) are captured by the other individuals. If this case were not true, then there would be direct conflict and the space between an individual's preferences and any feasible solution would be too great to call 'compromise.'
2. The preference structure used by the supra-decision maker is completely specified by the individuals. Each individual must state either a preference or indifference for all categories. The addition/subtraction of categories is not allowed.
3. The supra-decision maker knows the value system of all individuals. This removes uncertainty from the problem. Doing so

allows this type of deterministic method to be properly employed. [Keeney and Raiffa].

Identifying a Social Welfare Function

In order to evaluate multiple approaches across a group of individual decision makers with varying preferences, a social welfare function must be formed. Luce and Raiffa state, " . . . the problem is to define 'fair' methods for amalgamating individual choices to yield a social decision." While there is no shortage on the numbers of pre-packaged welfare functions, the model which best fits this paradigm is Nanson's function. Nanson's functions is a Borda elimination function. An overview of Borda is presented as a precursor to Nanson.

Borda's function is the rank order function. Given a set of m alternatives, assign marks of m-1, m-1, ... 1, 0 to the first ranked, second ranked, .. . last ranked alternative for the each stakeholder. Then find the sum of all the individual rankings. Under Borda's function, the alternative with the highest score wins.[10] Nanson takes this concept a step further and adds the component of elimination. AT each stage in the sequential process, alternatives with the lowest Borda score at that stage are eliminated, unless all of the scores are equal. Fishburn has proven that Borda's function is Pareto optimal. It therefore satisfies our main criterion of providing a solution that is on the efficient frontier [Lin].

Defining the Social Welfare Function

Defined by Arrow, a social welfare function is a process or rule which, for each set of individual orderings, $R_1, \ldots R_n$, for alternative social states (one ordering for each individual), states a corresponding social ordering of alternative states, R. Alternatively defined by Hwang and Lin, a social welfare function can also be considered as a rule which maps the profile of individual preference orderings into one pf the possible preference orderings for the society itself. While Borda

[10] The USA Today college football poll is an example of the Borda function in use.

provides the qualification function; we must still determine a method to amalgamate individual preferences to create a group function.

By using multi-attribute utility analysis (MAUA), we are able to add the preferences of all stakeholders and calculate a group utility. Known as the Bentham-Edgeworth sum of individual utilities, there are many critics of this method especially when each individual uses an ordinal technique. Arrow points out one such case where there are three individuals and three alternatives. Two of the individuals have the utility 1 for alternative x, .9 for y and 0 for z; and the third has the utility .5 for x, 1 for 1 and 0 for z. Clearly z is the worse scenario for all individuals, however, based on ordinal ordering, x would be the preferred choice, while based on utility, y would be chosen. In this early stage of research, it is unclear which of these would best suit the stakeholders. However, it does underscore the importance of a social welfare function in the group decision-making process. Economist Kenneth Arrow concludes there is no procedure for combining individual rankings into a group ranking that does not explicitly address the question of interpersonal comparison of preferences. Termed Arrows Possibility Theorem [11] for Social Welfare Functions, it states there is no method one can employ to generate a group utility (via a social welfare function) which satisfactorily compares individual utilities in the process [Arrow]. A formal argument against Arrow is beyond the scope of this paper, however, due to the finite bound being imposed on the problem, the assumption shall be made that we can minimize the effect Arrow will have on the output on the social welfare function.

An Example from the Real World

An essential presumption in this argument is that there are a finite number of scenarios that mandate the creation of a ground delay program. We briefly address each below:

[11] Although Arrow defines it as his *possibility* function, because of its pessimistic forecast, it is sometimes referred to as the *impossibility* theorem.

1. **Weather** – primarily thunderstorms, planes do not perform very well in these situations. As a result, planes are rerouted around the weather (en-route), thus modifying the original schedule. When the storm is near the airport, the schedule impact is greater, sometimes forcing a complete stoppage of traffic until the front moves on.
2. **Runway closure** – During snowstorms, runways are often shut down for cleaning. Depending on the runway (i.e. direction), the impact must be calculated, resulting in a dip in AAR and the need for a revision of the GDP. Although not common, equipment failure (i.e. an immobile plane on the airport surface) may also cause the closure of a runway.
3. **Visibility** (low ceilings) – While technology has decreased the amount of human intervention needed to fly a plane, the landing procedure has managed to evade automation. Upon approach, both a pilot and the tower must be able to make visual contact with each other. When ceilings are low, due to fog, haze, or simply low lying clouds, airport operations must be decreased. When this occurs at peak periods, demand may surpass capacity and a GDP becomes necessary. San Francisco (SFO) faces this challenge nearly every day of the year.
4. **Wind** – Runways are built according to the wind patters of the airport. Varying runway configurations are utilized to allow planes to take off and land into the wind (when present). Periodically, a cross wind (i.e. going east to west across a north-south runway). This complicates the landing process forcing the pilot to compensate for the horizontal vector that is working against the plane. As stated earlier, the US air travel system is based on a sunny-day picture. Crosswinds shade this model and force airport operations to spread out (allowing more time for each landing) and thus decreasing the AAR.
5. **Airport congestion/volume** – The NAS can be thought of as a series of queues. From the time a flight pulls into a gate, it must queue to taxi-out of the gate, queue for take-off, queue for sector space (en-route), queue for clearance to land, and the queue to taxi-in to a gate. We assume each of these queues is moving its normal pace and serve is available. When

one or more of these queues provides degraded service, the delay back propagates creating a massive ripple effect in the system. At the airport level, there are times when there are more outbound flights than there are supposed to be. This is normally due to delays occurring in other parts of the system which are felt by a single airport. To alleviate this, a GDP[12] will be put into place which decreases the number of arrivals, thus giving precedence to departures to leave the airport.

Based ion this taxonomy, within bounds, similar scenarios call for repeatable solutions. While a confirmation bias often sets in, our model will suggest solutions based on past occurrences and user preferences. To demonstrate this point, on July 22-24 of 2001, Chicago O'Hare (ORD) was caught in thunderstorms. According to the Air Traffic Control System Command Center Severe Weather Area Log of July 24th: "Once again ORD was the main event, with a weather pattern that was a carbon copy of the last two nights." That summary goes on to state that thunderstorms started northwest of ORD and moved slowly in a southeasterly direction. The slow progression of the weather front caused various arrival and departure fixes to be shut down, thus impacting the capacity of the airport. Undoubtedly there are other variables to be considered (including time of day, day of week, time of year) which would have to be considered before drawing a conclusion. However, once a rich enough knowledge base is constructed, these can also be factored into the evaluation.

Putting the Method to Work

Earlier in this paper, profiles for each of the stakeholders were presented. In the absence of a full investigation these will be treated as the baseline for this paper. There is an obvious need for follow-up work to gain a broader understanding of the value system of each of the stakeholders, which will allow for a comprehensive study of the system. Before commencing on the task of simulating the system, a list of assumptions is presented:

[12] Ground stops are also used in these situations to completely stop inbound traffic and allow departures to leave the airport in a more expeditious manner.

1. The stakeholder preferences were honestly expressed and are representative of the entire group and not just the individual.[13]
2. The "supra" decision maker understands and values the preferences of the individual decision makers and their goal is to "maximize the well being" of those decision makers.
3. All individual decision makers have expressed their preferences for all variables – thus making the problem one of certainty.

In setting up the problem, the supra decision maker will be separated from the other stakeholders, the participatory group. The supra decision maker ultimately makes the final decision. She can be considered a "benevolent dictator" who is ultimately making decisions for large numbers of people. The participatory group, individuals I=1,2, . . . , N are the members of the decision group. Within this model, u is the group's utility function. While u_i specifies each individuals' utility function. They essentially form the basis for the supra decision makers final verdict.

While research in group decision-making has generated a plethora of assumptions and rules to be adhered to when employing an additive form, only the few most relevant ones will be mentioned here. First is the idea of universal agreement which states, if all members of a group have the same utility, then the group utility should be the common utility. Next, ordinal positive association necessitates if alternatives A and B are equally preferred by the group, and A is modified to alternative A' and some individual members prefer A' to A (but all others remain indifferent), then A' is preferred to the group [Keeney and Raiffa].

[13] A series of informal interviews were performed to gain insight into each stakeholder. While the questions asked did not ask for quantitative responses, the dialogue touched each relationship and established definite preference patterns.

Each of the variables identified in table 3 will be defined per airport. For example, at ORD, there are several scopes available. For example, an ALL program would affect all 20 centers in the NAS while a 1st Tier program would only control surrounding airports (Midway (MDW), St. Louis (STL), Indianapolis (IND)). Each of these scenarios will be given a utility based on stakeholder preference.

Equation 2: Calculating Individual Stakeholder Utility

$$u_i(x_1) = \lambda_1 u_1(s_1) + \lambda_1 u_1(l_1) + \lambda_1 u_1(a_1) + \lambda_1 u_1(d_1) + \lambda_1 u_1(e_1)$$

Where:

λ = the scaling constant (between 0 and 1)which measures the relative importance of a variable

u = the automated utility for the variable in question

s = Scope

l = Duration

a = AAR

d = Delay

e = Equity

From equation 2, the group utility function can be determined:

Equation 3: Calculating Group Utility

$$u(x) = \sum_{i=1}^{n} \lambda_1 u_1(x)$$

Therefore as the model weans out low scoring alternatives, the group will be left with fewer choices and less comparisons to be made. As stated earlier, the system will not identify a single solution so it will be up the user to make the final assessment. With this being the case, it is important to define a mechanism for comparing trade-offs between solutions. The next three sections describe various approaches for elucidating trade-off.

Comparing Trade-offs

Once preferences have been defined, and the group utility calculated, the true power of the method comes into play. When it comes to trade-offs, humans are poor preceptors of value. While the expected value (i.e. expected payoff) of a problem remains a salient comprehensible attribute, the function itself and changing quantities among the attributes, raise doubt in decision makers. Consider the following alternatives:

Act A earn $100,000 for sure
Act B earn $200,000 or $0, each with probability 0.5
Act C earn $1,000,000 with probability 0.1 or $0 with probability 0.9
Act D earn $200,000 with probability 0.9 or lose $800,000 with probability 0.1

While detractors may argue that Act A is the obvious choice, it should be noted that all four outcomes have the same expected utility of $100,000. Understandably, by introducing the element of uncertainty, we move away from the fundamentals of our initial method, however, the basic premise remains: When the components of the function change (in value), we are poor preceptors of how this affects the expected value.

In an attempt to correct this, two techniques are employed, indifference curves and aspiration levels. Both of these concepts assist the decision maker in moving away from a binary outcome, towards a manifold scenario. While all individual decision makers have their preferences, the ordinal rankings between each is unclear. Both indifference curves and aspiration levels attempt to disambiguate and more importantly quantify the gaps between outcomes.

Indifference Curves

Indifference curve quantify the gap between units for each individual. That is, once the group utility is calculated, when a change in the "calculated" group utility (i.e. some variables are increased/decreased) is made, how does that affect each user. Moreover, how much of a change can take place before a user actually *notices* an outcome that works against their preferences. This idea is closely related to just noticeable difference (JND)[14] which examines the amount of change that must take place before an individual realizes a change in their environment.

TO give a context specific example, we consider the oft-encountered scenario at SFO where a GDP is being discussed and the stakeholders are being polled for their preferences. Two popular scopes for SFO programs are 6 West and 12 West. The six West has a smaller scope and is often favored by the smaller airlines. While larger airlines bid for a 12 West so they are not left with airplanes that do not receive a delay and then have to wait for the delayed planes to catch up. While these two stakeholders have justification for their partiality, others do not have the same tangible reasoning. An indifference curve can be informative in helping the stakeholder formalize what the trade-off between two attributes means to them.

Figure 5 below depicts an indifference curve graphically. Consider two attributes A_1 and A_2 with increasing utilities -- meaning having more of either attribute is better than having less. Each of the points on the curves represents a consequence (or outcome). If we consider each curve a different scope, ("6 West," "10 West," "12 West"), the user should be indifferent to any changes in A_1 and A_2 which fall between the curves. That is, the stakeholder should always choose the lesser curve.

[14] Psychophysicists attempt to test this by taking a light bulb and slowly increasing its intensity, until the subject admits to realizing a difference.

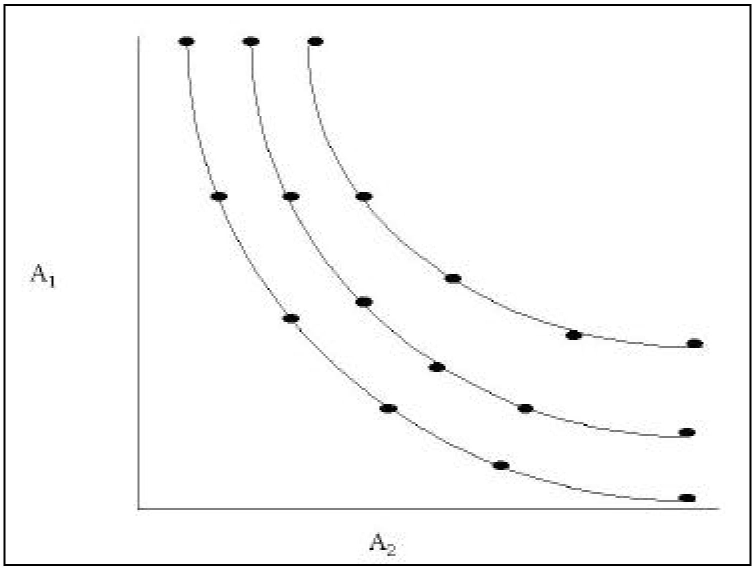

Figure 5: Indifference Curves Pictorially [Schum]

Aspiration Levels

Aspiration levels are another way to measure trade-off for users. Each user has an optimal value for each variable. An aspiration level sets a range for all variables where an additional quantity of one variable can be offset with a comparative decrease across one or more other variables while still attaining the user's goals. In the figure below, based on aspiration, option one is the better choice, even though option two exceeds option one in three of the four attributes. Commentators, such as Schum, have argued that using aspiration levels do not give consideration to the relative importance of the attributes. By using other, ordinal-oriented methods, in concert with aspiration levels, we can circumvent this shortcoming.

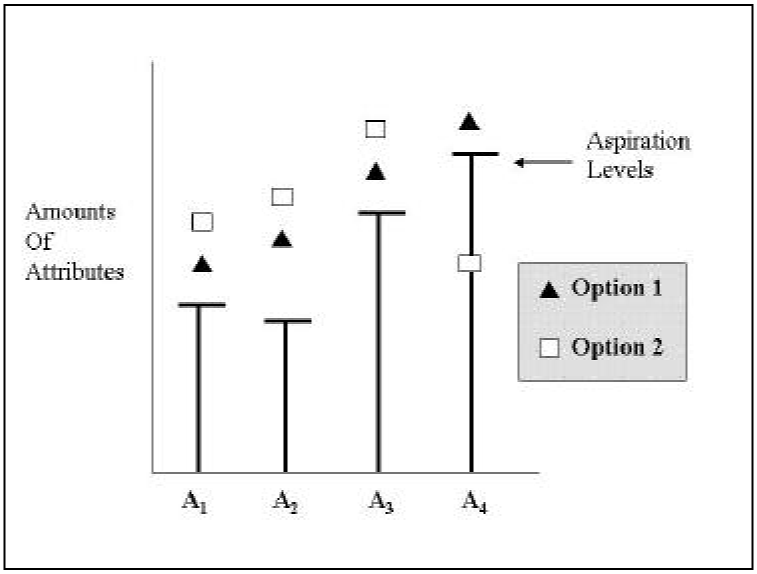

Figure 6: Aspiration Levels Across Four Attributes [Schum]

Pareto Optimality

Pareto optimality describes the situation where all stakeholders in a group decision makers have outcomes on an efficient frontier. The efficient frontier describes the region where each stakeholder's utility is optimized to the extent that an increase in any one's stakeholder's utility would result in a decrease in another player's utility. When a pareto optimal solution has been found, all stakeholders have been guaranteed the best individual outcome while get the most out of system resources. By using the group utility to create a pareto optimal scenario, it would essentially eliminate any 'gaming' of the system by individual stakeholders.

To appreciate the concept of Pareto Optimality, one must first understand the notion of dominance. In Figure 7, having more of attributes A_1 and A_2 is better than having less. In this scenario, any of the outcomes that produce an 'obvious loser' can be eliminated from consideration. As the dotted line demonstrates, any of the outcomes in that range dominate it. That is, they are at least as good, or better for the attributes in question.

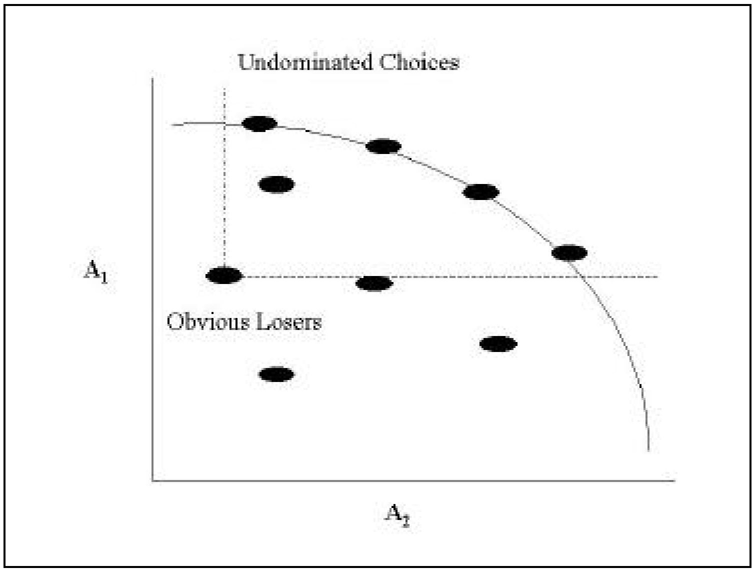

Figure 7: Pareto Optimality

Those points which lie on the curve reflect the efficient frontier. No dominant region can be drawn which includes another alternative. With this being the case, these solutions are pareto optimal, where no one stakeholder's utility can be bettered, without decreasing the utility of another stakeholder. Efficient frontiers have been used in the past by Decision and Designs, Inc. to assist clients in choosing designs for a car [Winterfeldt and Edwards] as well as an innovative application by Metron Inc. to help minimize the travel time by salespeople within a region [Lent]. It should be noted that outcomes rarely fall on the efficient frontier, but solutions that are near the efficient frontier are often available.

Summary

Figure 8 demonstrates how each of the aforementioned methods can be used in tandem to aide in the group decision-making process. Starting at the bottom of the figure, each of stakeholder will establish their utility, using multi-attribute utilities. These are then summed and a group utility is calculated. At that point, each of the three techniques can be employed with the results sent back to the stakeholder. Adjustments can be made and the cycle repeated if needed.

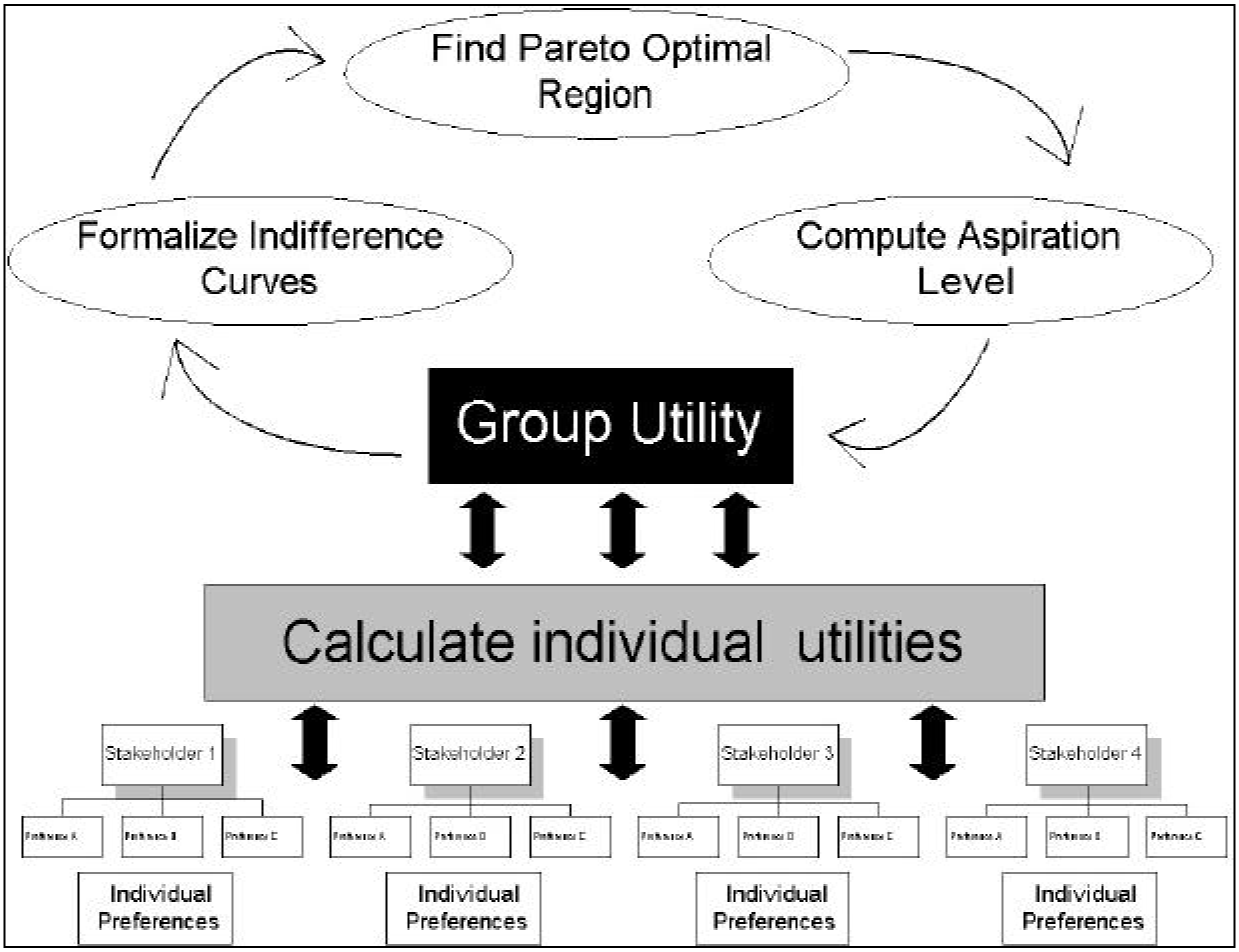

Figure 8: The System Cycle

Recall Figure 4 illustrated the high level architecture of the system. A fully automated system would be able to perform all of the above operations without human intervention. Once each stakeholder's preferences have been 'learned' by the system, the need for an iterative decision making process will become unnecessary. Although utopian in nature, it is possible that enough scenarios can be cataloged by the system where 'canned' decisions become possible. Buede and Watson define this as a programmed decision, "where it is possible to prescribe how a decision should be made prior to its appearance (like in a warehouse, supplies are ordered based on projected demand on quantity, no long process is needed beforehand.)

Discounting the implementation of a fully automated system, few would argue against the notion that the time has come to move forward in Wicken's stages of automation. The system described in this paper will go beyond presenting data (a set of discrete, objective facts about events) and provide information (data endowed with relevance and purpose) and knowledge (information with guidance for action) to the user to aid in the decision making process. Epistemologists have derived a chain which links each of these entities and shows their relationships with each other (shown in Figure 9).

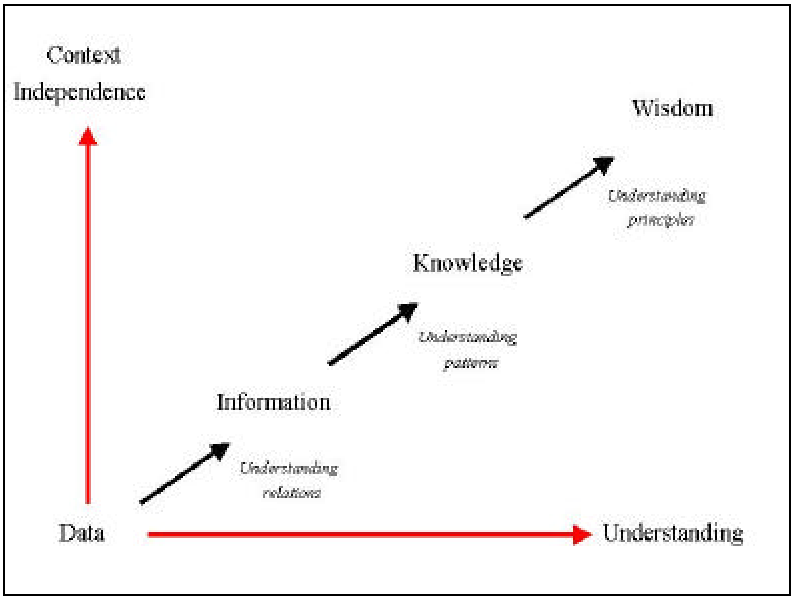

Figure 9: Knowledge Hierarchy

Infusing these elements into the system will provide the user with a more complete understanding of the solution space allowing them to make better-informed decisions, which will in effect benefit the system. Transparent to the user, the system will also provide a feedback loop as it informs the user of their best alternatives and warns the user when their choices go against their stated preferences. Finally, the system will also have "perfect recall," the ability to summon past events and offer them as archetypes (if the pattern fits).

Future Study

The importance of applicability of research to the real world (find the case study word for this) can never be understated. This paper relied on informal interviews with FAA employees to create stakeholder profiles. While each was intimately involved with the CDM process, there comments were not limited to one specific locale. It will be necessary to choose an airport and conduct a case study on their operations. This will include cataloging (at least) one year's worth of ground delay programs and analyzing stakeholder decisions during various scenarios. Although this paper focuses on the decisions made when formulating a GDP, many decision are made across time until the GDP is ended. Finally, examining sequential decision-making as it relates to these determinations could reveal new ways to approach these questions.

Works Cited

Adelman, Leonard. "Real-Time Computer Support for Decision Analysis in a Group Setting: Another Class of Decision Support Systems." *The Institute of Management Science*, 1984.

Air Traffic Control System Command Center Severe Weather Area Log, Federal Aviation Record Of Facility Operation, July 24, 2001.

Agarwal, Ritu, and Kislaya Prasad. "Enhancing the Group Decision Making Process: An Intelligent Systems Architecture." *Proceedings of the Twenty-Second Annual Hawaii International Conference on Systems Sciences*, 1989.

Arrow, Kenneth J. *Social Choice and Individual Values.* New Haven: Yale University Press, 1951.

Ball, Michael, et al. "Assessing the Benefits of Collaborative Decision Making in Air Traffic Management," *Proceedings of the 3rd USA/Europe Air Traffic Management R & D Conference*, 2000.

Chang, Kan et al. "Enhancements to the FAA Ground Delay Program Under Collaborative Decision Making." *International Journal of the Institute for Operations Research and the Management Sciences*, Vol. 31, 1, 2001.

Clawson, V et al. "The Role of Facilitator in Computer-Supported Meetings." *Small Group Research,* 1993.

Curmi, J. et al. “Modeling Pilot Behavior in Air Traffic Control Simulation” *In Proceedings of the Simulation Technology and Training (SimTexT) Conference*, 1996.

Donohue, George L and William D. Laska. “United States and European Airport Capacity Assessment using GMU Macroscopic Capacity Model (MCM).” *Proceedings of the 3rd USA/Europe Air Traffic Management R&D Seminar*, 2000.

Edmonds, Bruce. “Modeling Bounded Rationality In Agent-Based Simulations using the Evolution of Mental Models.” In Brenner, T. (ed.), *Computational Techniques for Modeling Learning in Economics*, Kluwer, 305-332

Fay, Stephen, *The Collapse of Barings*. New York: W.W. Norton, 1996.

Hoch, Stephen J. and Howard Kunreuther (eds.). *Wharton on Decision Making*. New York: John Wiley & Sons, 2001.

Hwang, Ching-Lai and Ming-Jeng Lin. *Group Decision Making Under Multiple Criteria: Methods and Applications*. Springer-Verlag: Berlin, 1987.

Kahneman, D. and A. Tversky. “Subjective Probability: A Judgment of Representativeness.” *Cognitive Psychology*, 3, 1972.

Keen, P. G. W., “Decision Support Systems: Translating Analytic Techniques into Useful Tools,” *Sloan Management Review,* Vol 21, No. 3, 1980.

Keeney, Ralph L. and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* Cambridge University Press: Canada, 1993.

Kollman, Kevin et al. “Collaborative Decision Making in Aviation Transportation: Improving decision making through shared information.” 1997.

Lent, Stephen. Electronic Mail Communication. 15 October 2001.

Lim, L. H., and K.K. Wei. “Does GDSS Promote More Democratic Decision Making? -- The Singapore Experiment.” *Proceedings of the Twenty-Third Annual Hawaii International Conference on Systems Sciences*, 1990.

Ljunberg M., and A. Lucas, A. “The Oasis Air Traffic Management System”. *Proceedings of the Second Pacific Rim International Conference on AI (PRICAI-92)*, 1992.

Luce, Duncan R. and Howard Raiffa. *Games and Decisions: Introduction and Critical Survey.* New York: Dover, 1957.

Miller, G.A. *The Magical Number Seven*, Plus or Minus Two: Some Limits on Our Capacity for Processing Information, Psychological Review, Vol. 63, 1956.

Morton, M.S. Scott Book. “Review of Decision Support Systems: Current Practice and Continuing Challenges.” *Sloan Management Review*, Vol. 21, No. 3, 1980.

Niederman, F. and R. Volkema. “The Influence of Agenda Creation and Use on Meeting Activities and Outcomes: Report on Initial Results.” *Proceedings of the Conference on ACM SIGCPR/SIGMIS Conference*, 1996.

Raiffa, Howard. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty. Addison Wesley:* London 1970.

Scholl, Richard “Methods of Group Decision Making.” 13 August, 2001 http://www.cba.uri.edu/Scholl/Notes/Group_Decision_Making.htm

Schum, David, *Class Notes*. Systems Engineering 563. Fall 1998.

Van de Ven, Andrew H. *Group Decision Making and Effectiveness*. Ohio: Kent State University Press, 1974.

Wambsganss, Michael. “Collaborative Decision Making through Dynamic Information Transfer.” *Air Traffic Control Quarterly*, 4, 1997.

Watson, Stephen R. and Dennis Buede. *Decision Synthesis: The Principles and Practice of Decision Analysis*. New York: Cambridge University Press, 1987.

Wickens, Christopher D. and Justin Hollands. *Engineering Psychology and Human Performance.* New York: Prentice Hall, 1999*.*

Winterfeldt, Detlof von and Ward Edwards. *Decision Analysis and Behavioral Research*. New York: Cambridge University Press, 1986.

Bibliography

Adelman, Leonard and Sharon Riedel. *Handbook for Evaluating Knowledge Based Systems: Conceptual Framework and Compendium of Methods.* Kluwer Academic Publishers: Boston, 1997.

Bard, Jonathan F, and Stephen F. Bousk. "A Tradeoff Analysis for Rough Terrain Handlers Using the AHP: An Example of Group Decision Making," *IEEE Transactions on Engineering Management*, August 1990.

Brandstatter, Hermann, et al. (eds.) *Group Decision Making*. London: Academic Press, 1982.

Cao, Patrick, and Frada Burstein. "An Empirical Study of Influences of the Coordination Models in Supporting Group Multiple-Criteria Decision Making" *Proceedings of the Australian Conference in Information Systems*, 2000.

Cao, Patrick P, and Frada Burstein. "An Asynchronous Group Decision Support System Study for Intelligent Multicriteria Decision Making" *Proceedings of the Thirty-Second Annual Hawaii International Conference on Systems Sciences*, 1999

Chung, Thomas H. "Approximate Methods for Sequential Decision Making Using Expert Advice" *Proceedings of the Seventh Annual ACM Conference on Computational Learning Theory,* 1994.

Davenport, Thomas H. and Lawrence Prusak. *Working Knowledge: How Organizations Manage What They Know.* Boston: Harvard Business School Press, 1998.

Gigernzer, Gard. “Bounded Rationality: Models of Fast and Frugal Inference.” *Swiss Journal of Economics and Statistics*, 133, 2, 2, 1997.

Gintis, Herbert. *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. New Jersey: Princeton University Press, 2000.

Goncalves, Nuno Pina, and Pedro Antunes. “An Information Systems Approach to Support Decision Makers Selecting Group Decision Processes.” *Proceedings of the 6th International Workshop on Groupware*, 2000

Hanappi-Egger, Edeltraud and Hardy Hanappi. “’Learning by Trying’: Simulating Strategic Behavior in Group Decision Making.” *Proceedings of the Thirtieth Annual Hawaii International Conference on Systems Sciences*, 1997

Hilmer, Kelly M., and Alan Dennis. “Stimulating Thinking in Group Decision Making.” *Proceedings of the Thirty-Third Annual Hawaii International Conference on Systems Sciences,* 2000.

Hoffman, R., Hall, W., Ball, M., Odoni, A. and Wambsganss, M., (1999), Collaborative Decision Making in Air Traffic Flow Management, NEXTOR Research Report # RR-99-2, UC Berkeley.

Iz, Peri H. “An Experimental Assessment of Preference Aggregation in a Group Decision Support System Based on Multiple Criteria Optimization.” *Proceedings of the Twenty-Fifth Annual Hawaii International Conference on Systems Sciences*, 1992

Jennings, Nicholas, R. and Michael Woolridge “Intelligent Agents: Theory and Practice.” *Submitted to Knowledge Review*, 1995.

Kraemer, Kenneth L. and John Leslie King. "Computer-Based Systems for Cooperative Work and Group Decision Making." *ACM Computing Surveys*, 1988.

Kraemer, Kenneth L. and John Leslie King. "Systems for Cooperative Work and Group Decision Making: Status of Use and Problems in Development." *ACM Computing Surveys*, 1988.

Limayen, Moez, and Chelbi Anis "Improving Multicriteria Group Decision Making with Automated Decision Guidance." *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, 1997.

McLeod, Poppy Lauretta. "A Comprehensive Model of Anonymity in Computer-Supported Group Decision Making." *Proceedings of the ACM International Conference on Information Systems*, 1997.

Parks, Craig D. and Lawrence Sanna. *Group Performance and Interaction*. Colorado: Westview Pres, 1999.

Rardin, Ronald L., *Optimization in Operations Research.* New Jersey: Prentice Hall, 1998.

Ryan, Andrew J. "An Approach to Quantitative Non-Functional Requirements in Software Development." *Proceedings of the 34th Annual Government Electronics and Information Association Conference*, 2000.

Sanford, Beverly D. "Decision-Aiding Automation for the En-Route Controller: A Human Factors Field Evaluation." *6th IFAC Symposium on Automated Systems Based on Human Skill*, 1997

Seidl, Christian, and Stefan Traub. “Testing Decision Rules for Multiattribute Decision Making.” *Discussion Paper Center of Economic Research at Tilburg University*, 1996.

Senge, Peter M. *The Fifth Discipline.* New York: Doubleday, 1990.

Simon, Herbert A. *Models of Man.* New York: John Wiley and Sons, 1957.

Takao, Shinji. “The Effects of Narrow Bandwidth Multipoint Videoconferencing on Group Decision Making and Turn Distribution.” *International Joint Conference on Work Activities Coordination and Collaboration,* 1999.