

Improving Decision Making via Knowledge Management

Andrew J. Ryan

EMSE 270 -- Dr. Rudy Garrity

INTRODUCTION.....	3
THE COLLABORATIVE DECISION MAKING INITIATIVE.....	3
LEVELS OF DECISION MAKING.....	5
LEARNING WITH THE SYSTEM.....	5
DOUBLE LOOP LEARNING	6
KM ELEMENTS IN THE SYSTEM.....	9
PERFECT RECALL.....	10
EPISODIC MEMORY	10
KNOWLEDGE TRANSFER – DECISION MAKING IN KM	11
REGRET THEORY.....	12
DISAPPOINTMENT THEORY	12
PROSPECT THEORY	12
TECHNOLOGY ASPECT.....	14
CONCLUSION	16
FUTURE WORK.....	16
WORKS CITED.....	18
BIBLIOGRAPHY.....	21
APPENDIX A: CDM STAKEHOLDER PREFERENCES.....	24
APPENDIX B – VARIABLES INVOLVED IN THE GDP PROCESS.....	25

Introduction

Since 1960, air passenger traffic has increased at an average yearly rate of 9% [Donohue and Laska]. Because of this steady increase in demand, the airways have become increasingly more congested. While congestion in the airways is easy to deduce from 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 offers knowledge management as a way to improve the decision making in TFM. By way of an agent-based system, knowledge can be gathered and disseminated, allowing FAA personnel to make better decisions.

The Collaborative Decision Making Initiative

At a landmark meeting in 1993, stakeholders of the NAS (see Appendix A for list) 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. An estimate by Collaborative Decision Making (CDM) participants calculated the CDM project

saved the airlines \$39,574,000¹ 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].”

As shown in Figure 1, 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 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.

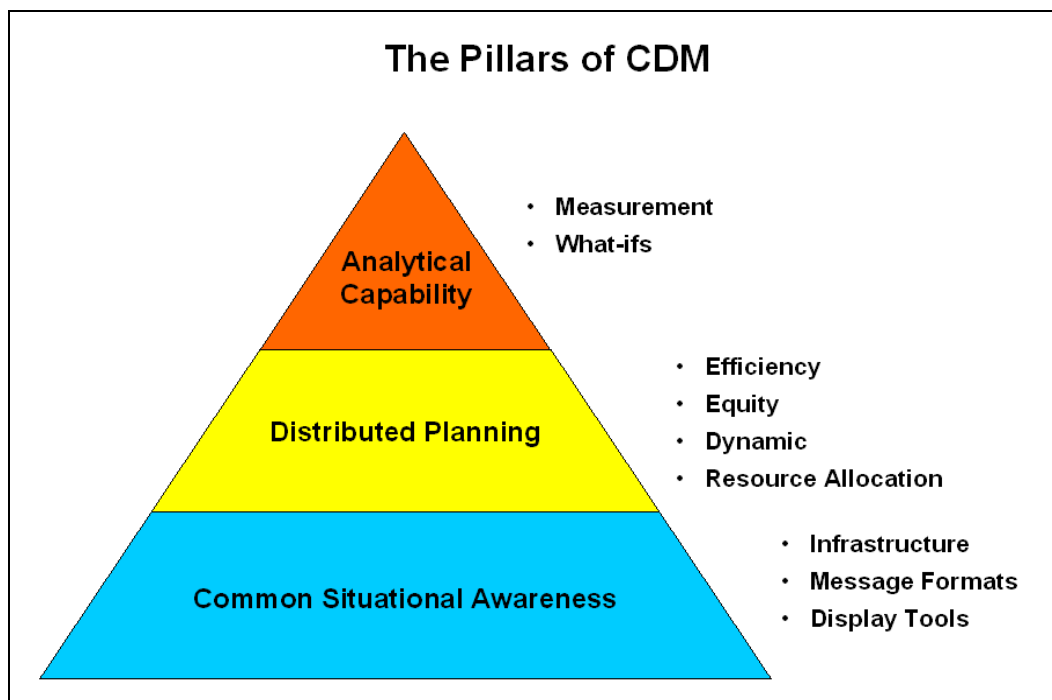


Figure 1: Functional Areas of CDM [Wambsganss]

¹ This figure was calculated using the conservative estimate each minute of delay mitigated by CDM saved the airline \$25.

Levels of Decision Making

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. *Wharton on Making Decisions* notes, 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).

Learning with the System

In terms of learning, one must consider Herbert Simon’s notion of bounded rationality which states that humans cannot be expected to act rationally at all times due to cognitive limitations as well as restrictions on memory, time (how long it will take to make a rational decision) and cost (how much will it cost the user to make a rational decision as opposed to a heuristic based decision). 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 for a knowledge-based system which requires human interaction. 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 2 describes the proposed system graphically. The system agents (discussed in detail later in the paper) 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.

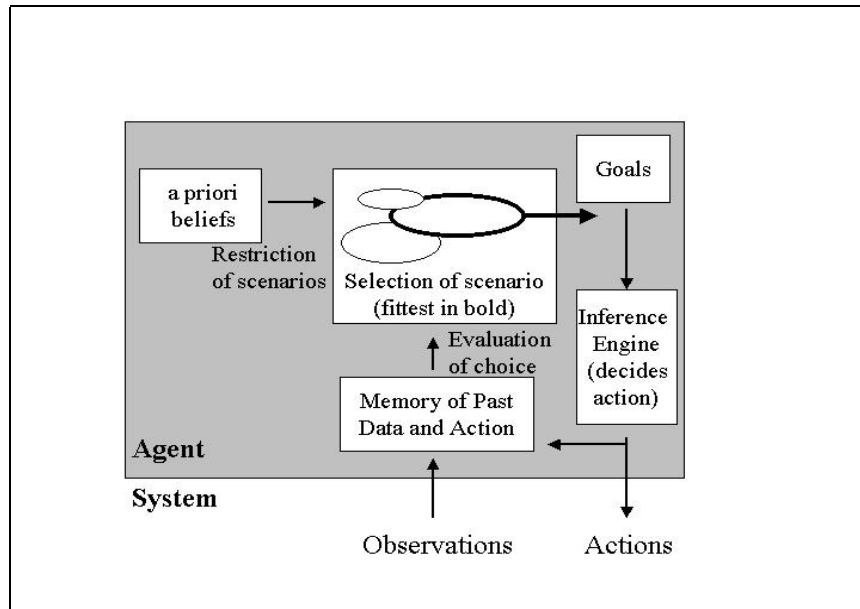


Figure 2: High Level Architecture of Proposed System

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.²

Double Loop Learning

In addition to the system growing with the user, it also allows for closed-loop (immediate) feedback. Examined from a knowledge management standpoint, the proper term would be “double loop learning.” As defined by Chris Argyris of the Harvard University Business School, “[double loop learning] allows decision makers to reflect critically on their own behavior, identify ways they often

² 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.

inadvertently contribute to the organization's problems, and then change how they react." Argyris makes the distinction between this and single loop learning where decision makers simply "sense and react" instead of adding the middle step to "asses and think," an essential step in double loop learning.

With the proposed system, the best possible scenarios will be presented to the user. From there, the user must decide which of those options are best for the situation. Additionally, the system will record these decisions, creating a "user-profile." Periodically, these decisions can be analyzed, using data mining and knowledge discovery techniques, to determine if the user is acting in accordance with their stated goals. Informal interviews with Air Traffic Control System Command Center (ATCSCC) traffic flow management specialists³ revealed that many of them could not articulate their desired goals when issuing a ground delay program. This result falls in line with the research of Argyris, "Ask people in an interview or questionnaire to articulate the rules they use to govern their actions, and they will give you what I call their 'espoused' theory of action. But observe these same people's behavior and you will quickly see it has very little to do with how they actually behave." Double loop learning will keep the user more in tune with their goals and even provide guidance if they begin to make decisions that are not harmony with their stated goals.

Measuring Outcomes

The final issue relating to learning is the measuring of outcomes. This involves creating a quantitative scale which can determine the effectiveness of a decision based on the outcome.⁴ Using multi-attribute utility theory (MAUT), also called multi-criteria decision making by some, weights can be given to attributes of the problem (attributes specific to the FAA problem can be found in Appendix B) and a final number can be calculated based on how much of each attribute

³ Interviews were conducted with various level employees at the ATCSCC facilities in Virginia. The ATCSCC is the center for all traffic flow operations in the nation.

⁴ It is important to note that a good outcome does not connote a good decision. In general, if the decision making process is well conceived and robust, the decision maker has done as much as they can to ensure a good outcome. However, if another user closes their eyes and chooses an option which yields a better outcome than the first decision maker, this does not mean that his *decision* was also better.

was “gained” by employing a particular option.⁵ The option which yields the highest utility (i.e. maximizes the most attributes) is deemed the best solution. All other decisions can be ranked in relation to this optimal solution. However, it must be noted that decision makers have varying preferences (documented in Appendix A) which may yield incongruent results. The next section offers a methodology to provide an optimal outcome when opposing preferences are involved.

Pareto Optimality

Pareto optimality describes the situation where all stakeholders, in a group decision, 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 stakeholder’s utility. When a pareto optimal solution has been found, all stakeholders have been guaranteed the best individual outcome while optimizing system resources.

To appreciate the concept of Pareto Optimality, one must first understand the notion of dominance. In Figure 3, 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.

⁵ Keeney and Raiffa’s *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, is considered the most complete reference on MAUT.

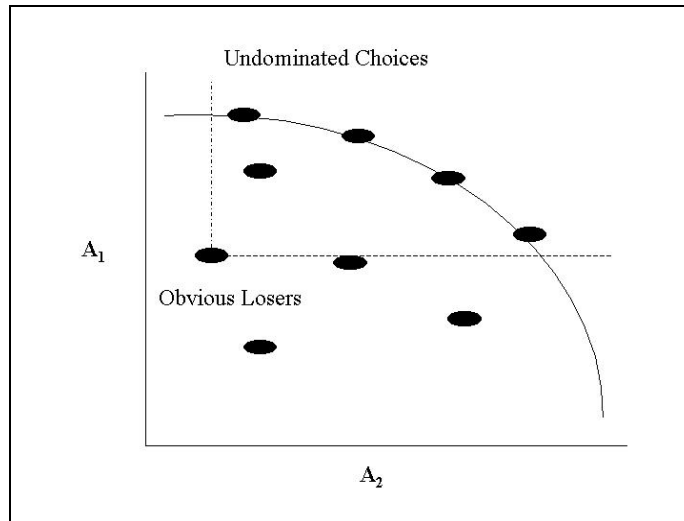


Figure 3: Pareto Optimality

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. With pareto optimality as the goal, let us look at knowledge management elements which could help us reach this goal.

KM Elements in the System

During the period starting November 1, 2000 and ending October 31, 2001, there were 166 ground delay programs (GDP) ⁶ run at Chicago's O'Hare airport. On each of these occasions, logs were created which captured important

⁶ 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.

events about the program including the duration of the GDP, the time it was put into effect, the reason for the GDP, and even user reactions and the effectiveness of the GDP. Violating the axioms of knowledge, none of this information is effectively used to take advantage of the potential benefits of the untapped knowledge stored in these files. Nonaka would find this to be a crack in the *spiral of knowledge*, “[where] an individuals’ personal knowledge is transformed into organizational knowledge valuable to the company as a whole.” While a noble objective, simply having the knowledge available does not guarantee improved performance. In the next session, we will examine some implications surrounding knowledge management and performance, focusing on decision-making.

Perfect Recall

In terms of decision-making, it has been noted that experts have the ability to respond ‘intuitively; and often very rapidly to situations – with a very high degree of accuracy and correctness [ChikGan]. Conversely, novices tend to spend more time analyzing a problem and do not have the same rate accuracy. One aspect of a knowledge management system is that it will have “perfect recall.” That is, it will be able to recall similar scenarios and present the specialist with important characteristics about the event which can aide in their development of a new GDP.

Episodic Memory

Somewhat related to “perfect recall,” “episodic memory” was termed by John Sowa, “ [to] define the operational knowledge that is stored in a database, acquired from a user during dialog, or derived by inference in answer to a question.” While a full review of semantic networks is beyond the scope of this effort, it is necessary to point out the five kinds of information that constitute the semantic memory of a system and hence formulate its episodic memory. According to Sowa these five elements are: ontology, definitions, constraints,

defaults, and behavior. It is important for all of these areas to be covered in a knowledge management system.

Critics of knowledge management would argue that knowledge alone does not make one an expert. In *Applied Cognitive Task Analysis in Aviation*, Seamster et al. comment on the difference between knowledge and skills. They remark, “Knowledge is the information required to develop skills . . . skill is the ability to perform the action or activity.” Therefore, knowledge alone cannot improve ability.⁷ The next section will explore ways that KM does improve ability.

Knowledge transfer – Decision Making in KM

Knowledge Transfer = Transmission and Absorption (and Use). Davenport and Prusak remark that most people are incapable of changing without the threat of crisis. The events of September 11th have brought about a real enough crisis where aviation workers are beginning to see the need for change. As the case with most highly skilled workers, ATCSCC specialists have an extreme confidence in their work. ChikGan’s edited volume, *Progress in Decision Utility and Risk Theory*, examines this phenomenon by way of the calibration paradigm. This paradigm is used to answer the question: Do people know how much they know? The experiment⁸ reveals an overconfidence phenomenon where people overestimate their knowledge. While a seemingly subtle discovery, decision theory offers another layer of reasoning to this argument. The following sections introduce theories which provide obstacles to

⁷ Seamster et al. offer five key findings about skill development:

1. Mental structures and processes change in important ways over the course of the learning process
2. The limited processing capacity of working memory affects how mental resources are allocated during performance
3. Humans have a certain degree of control (meta-cognition) over their mental processes, implying that people can monitor, control, and optimize their cognitive processes
4. Skills develop during a three-stage process that begins with the mastery of item-specific knowledge and progressing to the compilation of that knowledge into a rational structure. The final stage is skill refinement and expertise.
5. Experts in a domain can be distinguished from novices by their cognitive structures and skills (7)

⁸ The Calibration Method asks a subject an A or B question. The subject is asked to give a subjective probability (between 50% and 100%). They found that subjects are overconfident as their relative frequency of correct results are much less than the associated subjective probability.

the human decision maker. These must be taken into account before making the proclamation that KM improves decision-making

Regret Theory

Regret theory is best explained as an individual decision maker taking account not only the consequences they might experience as a result of the action not chosen, but also how each consequence compares with what they would have experienced had they chosen the other action under the occurred state of the world referred to this action.⁹

Disappointment Theory

Disappointment theory occurs when an individual compares the experience of a chosen action with the expected basic utility of that action. In narrative form, if an individual chooses some action, he/she forms some prior expectation about that action. After the uncertainty is resolved, the individual experiences one particular consequence which maybe better/worse than the prior expectation. The basic utility of the consequence will be modified by some decrement/increment to disappointment or elation.

Prospect Theory

Prospect theory can only be applied to events with probabilistic outcomes. It is an attempt to eliminate discrepancy from real human behavior.¹⁰ Prospect theory allows three behavioral effects:

1. Certainty effect – the tendency to give greater weights to determinant outcomes
2. Reflection effect – the tendency to change preferences upon passing from gains to losses
3. Isolation effect – the tendency to simplify choice by eliminating the common components of decision variants.

⁹ This result is called modified utility.

¹⁰ The gap between theory and human behavior is also referred to as prescriptive theory.

Finally, it is important to consider how humans process information. KM proponents extol a KM system's ability to capture and organize information. Mistakenly, they ignore the limitations the system users face. A few of these limitations are described below as part of the *Qualitative Model of the Human Decision Making Process*:

1. The features of the human information processing system
 - Limited span of working memory (Miller's Magical Number Seven Plus or Minus 2) of the decision situation by replacing some of the criteria limitation by eliminating in some of the criteria, by grouping alternatives etc.
2. Limited exactness in quantitative measures
 - A human being is not an exact measurement device producing quantitative measurements [Kahneman and Tversky]. This is the reason for intransitive behavior some problem of choice.
3. Human errors and contradiction
 - People err when processing information. Some reasons for this include: weariness, lack of attention, and habitual heuristics.

As described above, humans are not optimal decision makers. Figure 4 illustrates Simon's bounded rationality and the fact that actual judgment falls below normative judgment. Coupled with the deficiencies in prospect theory, regret theory and disappointment theory – it is obvious that knowledge management techniques can be utilized to bring the human decision maker closer to normative outcomes. Finally, we turn our attention to the technology that will drive the knowledge management system. For without technology, many of the improvements in decision-making will never be realized.




Figure 4: Bounded Rationality

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 systems¹¹ 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

¹¹ Bielawski and Leward give an excellent definition of an intelligent system, “Intelligent systems draw on knowledge and the power of association and inference to steer the direction of a running program towards useful results.”

human user. As stated by Quinn et al. an important part of knowledge management involves managing “human/professional intellect”¹² and converting it into useful products and services. It is self-defeating to allow the resulting system to have an intelligence that the *source* cannot comprehend.

While the aforementioned Big Blue example violates the idea of bounded rationality, Xerox heeded Simon’s advice as they developed an intelligent system to decrease the cost of copier maintenance [Brown]. They designed a system that was able to self-diagnose itself and tell the user (or service person) what was wrong with the machine. Furthermore, once the machine was fixed, they logged all of these events into a database, allowing other service workers to learn from the work of their peers. Taking lessons from both of the aforementioned examples and the decision theory axioms, 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)
- Yield explainable results

In imposing these limitations, we attempt to adhere to Simon’s idea of bounded rationality.¹³ 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]. Although the idea of bounded

¹² Human/professional intellect has three characteristics:

- a. Cognitive knowledge (know how) – translating book learning into effective execution
- b. Systems understanding (know why) – deep knowledge of the web of cause and effect relationships underlying a discipline
- c. Self-motivated creativity (know-why) – consists of will, motivation and adaptability for success

¹³ 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.

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.

Conclusion

This paper examines knowledge management as it relates to decision-making. While it is easy to conclude that the more information a decision maker has at their disposal, the better the decision will be. This is not always the case. Decision theory offers counterpoints to this thesis, forcing a more robust system to be used. Such a system will be able to filter options for the user and provide information relevant to the current situation; allowing ATCSCC personnel to have the most relevant information before making a decision.

Currently the ATCSCC employs no type of knowledge management. A data-laden organization, they take a myopic approach to analysis – rarely mining their repositories in an attempt to learn from the past. Utilizing a knowledge management system, which grows with the user, will allow ATCSCC specialists to learn from past experiences. As the NAS continues to become more congested ¹⁴ the increase in decision-making and the ‘lessons learned’ and ‘best practices’ that will evolve from a KM system will allow our airways to be operable well into the future.

Future Work

Examining knowledge management in use, there are opportunities to advance the total life cycle of the system. The goal will be to encourage users to add values by transforming data and information into knowledge. Currently this requires manual intervention. An exciting extension to this practice would be to setup an expansive ontology of the domain which captures the data objects, their relationships, and any transformation that represent the physical and cognitive entities necessary to accomplish some task. Once successfully formulated, the

¹⁴ While the events of September 11 had a profound effect on the NAS, the *number* of flights in the air has remained the same – with an increase on private aircraft offsetting the decrease in commercial traffic.

system can mature as new knowledge is captured by the system and entrenched in the knowledge base.

This type of maintenance is routine for all databases, but with the richness of a knowledge base, it is paramount to keep the knowledge current. According to Coenen and Bench-Capon, this type of perfective maintenance takes up 50-65 percent of the maintenance effort in a system. Through the use of agent technology and a well flushed out ontology, the system will be able to 'learn' with limited user intervention. This in turn allows for programmed selection, "[which] applies to classes of well defined repetitive problems and is carried out by a well-defined decision rule which determines which action will be taken in any problem of the class" [White].

A programmed selection would allow the system to give recommendations (as stated earlier) but also yields the following advantages:

1. Makes the process quicker in terms of time needed to produce a solution
2. Decrease the effort needed to produce a solution
3. Freed time can then be devoted to other (less routinized) matters

The only glaring disadvantage lay in the fact that there must be a capability to ensure that the selections made are identical with the person who would have made them without formal aids. Even with this possible shortcoming in place, the ability to have a knowledge management system 'learn' and then feed its findings into a decision aid is a promising outlook. As companies continue to search for a *collective intelligence*, this seems like an excellent way to reach that goal.

Works Cited

- Adelman, Leonard and Sharon Riedel. *Handbook for Evaluating Knowledge Based Systems: Conceptual Framework and Compendium of Methods*. Boston: Kluwer Academic Publishers, 1997.
- Argyris, Chris. "Teaching Smart People How to Learn." *Harvard Business Review on Knowledge Management*. Boston: Harvard Business School Publishing, 1998.
- 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.
- Brown, John Seely. "Research that Reinvents the Corporation." *Harvard Business Review on Knowledge Management*. Boston: Harvard Business School Publishing, 1998.
- 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.
- ChikGan, Attila. *Progress in Decision, Utility and Risk Theory*. Boston: Kluwer Academic Publishers, 1991.
- Coenen, Frans and Trevor Bench-Capon. *Maintenance of Knowledge-Based Systems*. London: Academic Press, 1993.
- Curmi, J. et al. "Modeling Pilot Behavior in Air Traffic Control Simulation." *In Proceedings of the Simulation Technology and Training (SimTexT) Conference*, 1996.

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

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

Gal, Tomas, et al. (eds). *Multi Criteria Decision Making Advances in MCDM Models, Algorithms, Theory, and Applications*. Boston: Kluwer Academic Publishers, 1997.

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

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.

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.

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.

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

Nonaka, Ikujiro. "The Knowledge-Creating Company." *Harvard Business Review on Knowledge Management*. Boston: Harvard Business School Publishing, 1998.

Quinn, James Brian, et al. "Managing Professional Intellect: Making the Most of the Best." *Harvard Business Review on Knowledge Management*. Boston: Harvard Business School Publishing, 1998.

Seamster, Thomas L, et al. *Applied Cognitive Task Analysis in Aviation*. Aldershot: Avebury Aviation, 1997.

Sowa, John. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Pacific Grove: Brooks/Cole, 2000.

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

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

White D. J. *Decision Theory*. Chicago: Aldine Publishing Company, 1969.

Bibliography

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.

Benyon-Davies, Paul. *Knowledge Engineering for Information Systems*. London: McGraw-Hill Book Company, 1993.

Bielawski, Larry and Robert Lewand. *Integrating Expert Systems, Hypermedia, and Database Technologies*. New York: John Wiley & Sons, 1991.

Billings, Charles E. *Aviation Automation: The Search for a Human-Centered Approach*. New Jersey: Lawrence Erlbaum Associates, 1997.

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

Cutcher-Gershenfeld, et al. *Knowledge Driven Work*. Oxford: Oxford University Press, 1998.

Galotti, Vincent P. *The Future Air Navigation Systems: Communication Navigation Surveillance Air Traffic Management*. Cambridge: Ashgate Publishing Company, 1997.

Gigerzner, 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

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.

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

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

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.

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

Mollaghasemi, Mansooreh and Julia Pet-Edwards. *Multi-Objective Decisions*. Los Alamitos: IEEE Computer Society, 1997.

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

Myers, Paul S. *Knowledge Management and Organizational Design*. Boston: Butterworth-Heinemann, 1996.

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.

Ryan, Andrew J. "Improving Group Utility in the Collaborative Decision Making Process." <<http://mason.gmu.edu/~ajryan/research.html>> 10 October 2001.

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

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

Suijs, Jeroen. *Cooperative Decision-Making Under Risk*. Boston: Kluwer Academic Publishers, 2000.

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

Appendix A – 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

APPENDIX B – 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.

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.