OASIS - THE SOUTH FLORIDA WATER MANAGEMENT ADVISOR

Gary F. Goforth
Jeffrey S. Tongue

South Florida Water Management District
West Palm Beach, Florida

1. INTRODUCTION

The South Florida Water Management District (SFWMD) operates more than 200 hydrologic control structures along almost 2000 miles of canals in southern Florida to both control water levels and water quality for an area of nearly 18,000 square miles. Seasonal demands by municipal and agricultural interests, as well as perennial flood control, create a complex decision making arena for the operations staff at SFWMD. Additional factors of major concern for operators are the environment, which includes both enhancement and protection, and meteorologic parameters which center on rainfall. The basic operating strategy is to: 1) provide flood protection during the wet season of June through October by placing water in storage and discharging excess amounts to the Ocean and 2) to draw from storage areas as needed for water supply during the dry season.

To support the water control operations, SFWMD developed an extensive hydrologic data collection network. Real-time data are continuously collected automatically via a microwave telemetry system as well as manually by telephone. In total, there can be as many as 100,000 records of information that are generated, collected, and recorded each day. Decisions in the operation of the hydrologic control structures are based on the analysis of portions of these records. In addition, established water level and device criteria, as well as expected meteorologic conditions enter into the decisions. Thus, because of the enormous amount of information involved in making a decision and the fact that the operation of a control structure can have a "domino" effect on other structures, operator experience plays a most important role in the decision making process.

An artificial intelligence application was selected for evaluation to do the job because of the tremendous amount of information that requires assimilation as well as the experience that is needed to operate the SFWMD's water control facilities. After evaluation of available technology and operational needs, the decision was made to develop a prototype comprehensive real-time operations advisory system. The system, referred to as the Operational Assistant and Simulated System (OASIS), was prototyped "in-house" by a water resources engineer over a two year period. Now that a prototype has been developed, OASIS is undergoing a thorough evaluation before the development of an operational system.

2. SIZING THE PROJECT

A conceptual design document was developed following numerous conferences with users, experts and engineers. Included were capability objectives, integration into the existing operations of the SFWMD's water control facilities, the extent of the expert system aspect, and the conceptual man-machine interface. Of the 12 major hydrologic basins located with the area controlled by the SFWMD, the Everglades Agricultural Area (EAA) was selected. The EAA was selected because it contains the widest variety of water control devices and has a complex operation. Over 30 locations within the EAA are operated by the SFWMD. In all, there are over 80 sensors, more than 50 control gates, and 8 major pump stations.

3. DEVELOPMENT OF THE EXPERT SYSTEM

3.1 Knowledge Acquisition

The knowledge base for the District's water control operations comprises both declarative and procedural information. The declarative knowledge describes static and dynamic facts and relations critical to the operation of the District's water management facilities. Static domain knowledge encompasses physical system components, hydraulic characteristics, and other features that are relatively invariant over time. The dynamic component of the knowledge base includes hydrologic data collected in the field and transferred to the control room, policy decisions which influence daily operations, the meteorological forecast, antecedent rainfall, and notices from the public. Procedural knowledge includes heuristics, empirical relationships, the hydraulic operating guidelines, operating constraints brought about by physical limits of the system, as well as fundamental domain principles; e.g., hydraulic relationships and the analytical tools used by the experts to reach decisions.

Knowledge acquisition is the process of eliciting and encoding domain expertise from human experts and other knowledge sources, typically through a combination of observation, structured interviews, and research. A major challenge during the development of OASIS was to express the decision making algorithms, procedures that may have taken years of training and experience to acquire and utilize, in a way that was lucid enough for a non-expert to grasp and could be encoded in an appropriate computer representation, and yet retain the accuracy, efficiency, consistency, and comprehensive characteristic of the domain experts. This challenge and other components of the knowledge acquisition process in the development of OASIS are addressed below.

The goals of the knowledge acquisition process were to address the following issues:

1. How do the experts assess the situation?
2. How do the experts define and investigate a potential problem?
3. What is the appropriate strategy to solve the problem?
4. What are the appropriate decision factors, and how are their values determined?
5. What are the data and rules used by the experts?
6. What are the tools used by the experts?
7. Identify other people or agencies involved in operating decision.
8. What is the time frame of the decision process?

Four techniques were utilized in the knowledge acquisition process which addressed these goals:

1. Familiarization of the domain.
2. Observation of the experts in action.
3. Detailed problem analysis (interviews).
4. Knowledge base verification.

Familiarization. A preliminary component of the knowledge acquisition process was for the system developer to gain as much insight into the domain as practicable. This was accomplished in part through observation of the experts as they evaluated situations, accessed data and relevant operating manuals, utilized appropriate analysis tools, and conferred with other individuals and agencies involved in the decision-making process. In addition, a significant amount of procedural knowledge was obtained from the hydraulic operating guidelines developed during the design and subsequent operation of the District's water control facilities. Daily operations logs, structure status information and hourly water levels are maintained by control room personnel, providing convenient access to historic field conditions and operational activities. Additional perspectives on the decision process were obtained through discussion of the operations with individuals other than the domain experts; e.g., District management and staff of projects affected by daily operating decisions.

Observation. Documented observation of the experts in action was a principal step in the knowledge acquisition process. While passive; i.e., uninterrupted, observation is suggested by some expert system developers, often times decisions were based on subtle factors which required more thorough explanations by the experts. Primary facets of the decision-making process observed included the type and priority of different data, the experts' working environment, what data are used most, who does the expert confer with, what are the steps in the many types of operational decisions, what data format is most convenient for the expert, what type of interface capability is needed, what is the time frame for decisions, and how do the experts deal with data uncertainty; e.g., inaccurate, inconsistent, or incomplete data.

Interviews. After several weeks of observation, the domain experts were interrogated during structured interviews. The interviews focused on decision factors, operating modes, data sources, and expected outcomes of operations for a prototype region of the District. Selected combinations of hydrologic and meteorologic conditions were examined. These combinations were specific enough for the experts to focus their attention, yet diverse enough to cover all practicable ranges of conditions. The main instrument used during the interviews was a matrix of decision factors, with each unique permutation specifying an appropriate set of water control operations. A challenging variety of conditions was evaluated during each interview, resulting in intense one to three hour sessions. During the interviews the experts collaborated for accuracy and clarity of their response. As a supplement to the experts' interrogation, the daily operations logs and hourly records provided examples and further clarification of the operation rules.

Some general operating guidelines were enumerated; however, a large percentage of the water control operations depend on specific conditions and could not be generalized. The knowledge base expanded more rapidly as the general guidelines were extrapolated to undocumented operations, in contrast to incorporating specific rules with unique decision factors.

Verification. After the information obtained from these interviews was encoded in the expert system, the experts evaluated the accuracy, consistency, and completeness of the OASIS knowledge base. Discrepancies and other conclusions of the evaluation were investigated and revisions were incorporated and subsequently re-evaluated through this iterative process.

3.2 KNOWLEDGE REPRESENTATION WITHIN OASIS

The goal of knowledge representation during the development of OASIS was to incorporate the variety and degree of expertise which compose domain knowledge into the data structures and procedures of OASIS. The knowledge base comprises both declarative and procedural information. The variety and degree of expertise which compose this domain knowledge was incorporated into the data structures and procedures of the OASIS prototype through multiple knowledge representations.

Declarative knowledge was represented in the OASIS prototype by means of a semantic net data structure, consisting of nodes representing components in the taxonomy of the District's facilities, and links describing the nodes and their relationships to one another. Stations were defined as instances of particular station types according to their functions and capabilities. Their spatial relationship to canals, lakes, water conservation areas, and other stations were explicitly documented as slots in the station and water body schemata. Specific structures and sensors at a specific station are referenced, and each has a scheme to document relevant characteristics and their current status. The dynamic information; e.g., real-time data and changes in structure status, were explicitly contained in slots of the appropriate facility scheme.

The major portion of the OASIS prototype operations advisor procedural knowledge was represented by modules of condition(s)-action(s) expressions, referred to as production rules. The following is a simple example:

IF - all of the following conditions are met:

1. It has rained more than one inch in the preceding 24 hours;
2. The weather forecast calls for more rain;
3. The canal level has risen more than 0.25 feet in the last hour;
4. The average canal level is above 12.0 feet; and
5. A high risk of damage exists for residences in the area if the canal level exceeds 13.0 feet.

THEN - perform the following action:

Investigate ways to lower the canal level.

This action places another fact in the data base, indicating that the operator is looking for ways to lower the canal level, which is subsequently compared to conditions in other rules; e.g.,

IF:

1. The operator is investigating ways to lower the canal;
2. Personnel are standing by at the local pump station;
3. The pumps are functional at the local pump station; and
4. There are no environmental, legal or other restrictions on pumping.

THEN:

Conduct pumping operations at the local pump station in accordance with standard operating procedures.

Conflicts arose in the order that rules were slated for firing. However, as the rules were analyzed, often one or more conditions of one rule would mandate a higher priority; e.g., the risk of immediate flooding generally always carries a greater sense of urgency than the risk of reduced water supply six months from now by temporarily lowering the water level.

The initial efforts to program the operating rules focused on specifying unique combinations of decision factors; e.g., weather, water levels, trends in water levels, and water supply demands, as conditions on the left hand side of the production rules, and specifying the appropriate action on the right hand side. This technique required an enormous amount of code to represent every possible combination of decision factors as a single specific rule and was very cumbersome to work with. Review of the operating rules indicated that the majority of possible operating conditions could be covered by writing general production rules which spanned a range of magnitudes for each decision factor. Specific rules were written to cover the remaining unique cases. This change in knowledge representation greatly reduced the number of rules in the knowledge base, the rules for one pump station alone were reduced in number by an order of magnitude, although the rules did become more complex. In addition to creating a more compact rule base, the generalized representation of operating rules was easier to maintain and expand.

3.3 ACCESS TO REAL-TIME DATA

A critical step in the development of the prototype was the interface with the extant real-time data base, comprised of both automated data acquisition and records of manual readings. This enormous existing data base serves as the source of field data for the OASIS advisor. The technique selected for accessing this information within the prototype was to write a simple communications protocol to pass real-time data from existing Modcomp equipment to the prototype OASIS Symbolics equipment through a direct RS-232 serial connection. The method has serious drawbacks for full-scale implementation, but sufficed for the prototype development. To reduce the CPU load on the Symbolics, preliminary data filtering and format transformation were carried out on the Modcomp side. A Lisp function within the OASIS prototype inserts the data directly into the expert system's fact base.

3.4 THE OASIS PROTOTYPE

The OASIS prototype incorporates the dual capability of providing decision-making support to the control room operators with the supplemental ability to complete the decision process by suggesting appropriate control structure operations. The OASIS prototype incorporates four functional elements:

1. Operations Status displays real-time hydrologic, rainfall and structural information;
2. Operations Assistant displays current and historic time series of data for trend analysis;

3. Operations Advisor is the control structure operations expert system; and
4. Alarm Status provides continuous background data analysis for detecting present and anticipated alarm conditions, complete with suggestions for ameliorating the identified alarm conditions.

The prototype was developed jointly by the District and Inference Corporation. District staff developed the conceptual design, conducted the knowledge acquisition interviews, formalized the knowledge representation, and established the communications links between the Symbolics and other computers. On a contractual basis, Inference staff encoded the functional framework; i.e., the interfaces between the different components of OASIS, prepared the initial District maps, and designed the preliminary data storage format.

The prototype is executed through a highly interactive color graphics interface which utilizes a combination of the Symbolics Lisp color system and Inference's color ARTIST capabilities. A series of maps serves as the means to locate stations of interest, and screen menus are available to select OASIS functions. Desired basins, stations and operations are accessed by moving the mouse-controlled cursor arrow to the appropriate location on the screen. Keyboard interaction is required to enter some manually-processed data and optional station information.

The background process within the OASIS prototype analyzes the incoming data for current or projected alarm conditions. Present trends are extrapolated for 24 hours and seven days to identify impending conditions which would require control operations. The presence of any level of alarm urgency is communicated to the control room operators through an alarm window which is present during all OASIS operations. The color of the alarm window reflects the urgency of the situation, flashing red represents conditions which require immediate attention, such as high water levels with a threat of health or economic damages, flashing yellow represents conditions which need attention soon to prevent alarm conditions that may arise within the next 24 hours, and flashing green signals the possibility that a problem may arise within the next seven days. Details of the alarm conditions are obtained by selecting the alarm window with the mouse. Suggested operations for ameliorating each alarm condition are provided upon request.

3.5 PROTOTYPE EVALUATION

During the 12 months following the completion of the OASIS prototype, the prototype will undergo extensive internal and external review. The District's operations staff will evaluate the accuracy, consistency and completeness of the expert system and other OASIS features. In addition, the performance speed of the prototype will be analyzed with respect to the time requirement for the operations decisions. The data format of the knowledge base will be evaluated with regard to storage and retrieval performance, primarily considering the interaction with the Modcomp system and the archive data storage computers. External consultants will analyze the prototype on the basis of efficient and effective utilization of hardware and software capabilities.

Following the OASIS prototype evaluation, results of the technical analyses will be reviewed. If the District staff and external consultants agree that the prototype can retain its basic framework, the design revisions will be incorporated. However, if the consensus is that the prototype structure is not appropriate for the OASIS objectives, major modifications will be implemented, up to discarding the prototype. Regardless of the final evaluation recommendations, the prototype has satisfied the fundamental requirement of a knowledge-base prototype in that it has allowed the District to test the concept of applying artificial intelligence techniques to the decision-making arena of water control structure operations.

4. SUMMARY

OASIS will be used in an advisory capacity in the Operations Control Room. It will be a dynamic advisor, incorporating new operating procedures as they are implemented. Anticipated benefits of OASIS fall into two categories, daily operations and operations support. Benefits expected to accrue during daily operational activities include:

1. Routine system-wide monitoring will be provided on a consistent and comprehensive basis.
2. Decision makers will be fully informed of changing conditions at all times.
3. Continuous analysis of data trends will identify impending alarm conditions.
4. The operations expertise of multiple experts are combined.
5. Operation advice during crisis conditions will not be influenced by emotions.
6. Consistency of operations should provide more defensibility of actions, as well as enable better coordination with local water management districts.

Benefits expected to accrue in support of operations include:

1. The knowledge extraction process has provided documentation of previously undocumented operating rules.
2. The rules have been further evaluated as they have been organized during the knowledge acquisition process.
3. Using OASIS to determine the influence of proposed operating strategies on existing operations.
4. OASIS can act as a training tool for the Operations staff.