Applying Cooperative Distributed Problem Solving Methods to Trumpeter Swan Management

Richard S. Sojda
USDI – Geological Survey
Biological Resources Division
101 Lewis Hall – Biology Department
Montana State University
Bozeman, Montana 59717-3460
sojda@montana.edu
http://www.mesc.usgs.gov/staff/sojda.html

Adele E. Howe
Department of Computer Science
Colorado State University
Fort Collins, Colorado 80521
howe@cs.colostate.edu
http://www.cs.colostate.edu/~howe

Abstract

We are developing a decision support system in an effort to assist biologists who are managing habitats for the Rocky Mountain population of trumpeter swans. Swan management is a domain that is ecologically complex, and this complexity is compounded by spatial and temporal issues. We are focused on providing decision support that allows managers to develop habitat management plans for local sites while recognizing that such decisions have ramifications not only at other sites but in the flyway as a whole. Because swan management is an inherently distributed problem, our system utilizes artificial intelligence methods including cooperative distributed problem solving, blackboards, and expert systems. The system will be made available to swan managers through the world wide web, using commercially available software that provides a common gateway interface between the web server software and an inference engine.

The Inherently Distributed Nature of Trumpeter Swan Management

We are developing a decision support system to assist biologists with the management of the Rocky Mountain population of trumpeter swans (Cygnus buccinator). The number of swans breeding in the Tri-State area where Montana, Idaho, and Wyoming come together has declined to just a few hundred pairs. They have abandoned, to large degree, what were thought to be traditional migratory pathways. Swans, like most migratory birds in North America, travel along migration corridors that link northern breeding areas with more southern wintering grounds. National wildlife refuges such as Grays Lake, Red Rock Lakes, National Elk Refuge, and Bear River Migratory Bird

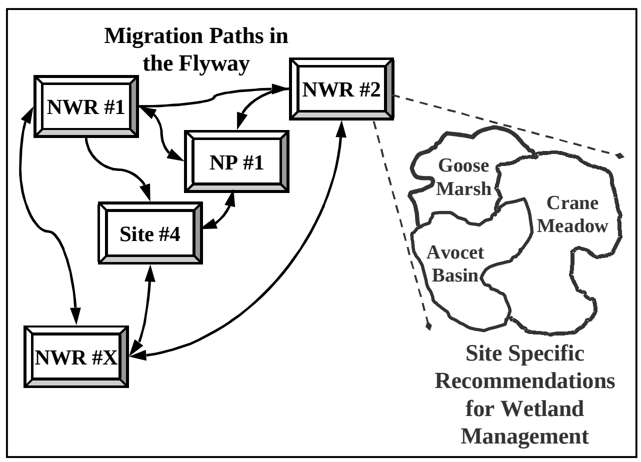

Figure 1. Spatial complexity shown as hypothetical migration paths among national wildlife refuges (NWR), national parks (NP), and other sites combined with recommendations needed for management of local wetlands.

Refuge, share swans at different times of the year with national parks such as Yellowstone and Grand Teton, and other areas such as Harrimann State Park. Swan management is a complex domain requiring reasoning across time and space among geographically dispersed managers (Figure 1.) Spatial interactions are inevitably intertwined with a temporal component as swans migrate. This complexity is further compounded by ecological issues that exist at specific wetlands in each location. Another component of complexity arises from local decisions having ramifications not only at other sites but in the flyway as a whole.

Swan managers have requested a decision support tool that will simulate and test management options for trumpeter swans throughout such corridors. The ultimate objective is to contribute to both population recovery and migration path development. They recognize that their decision making is cyclic, and they wish to iteratively plan, implement, evaluate, and improve their management strategies. Biologists are also concerned by their lack of ability to objectively assess critical information gaps, identifying those that contribute the most uncertainty to the selection of management options.

Unfortunately, optimizing any management of migratory birds throughout a flyway with cyclic planning is so complex that it is often all but impossible to implement without computerized decision support (Sojda, Dean, and Howe 1994). And, past conditions and future needs are ecological constraints to current decisions. Distributed decision making approaches suit problems where the complexity prevents an individual decision maker from conceptualizing, or otherwise dealing with the entire problem (Boland et al. 1992; Brehmer 1991). Our system is focused, then, on providing support for realistic and ecologically-based management of migratory birds at multiple geographic and temporal scales.

At this time, there are no such decision support systems available for swan managers, nor any common databases for them to access. Furthermore, many managers are physically either located in relatively remote locations or simply distant from each other, making it difficult to meet frequently. They currently do get together once or twice a year to discuss and select broad management options for the flyway and specific recommendations for specific sites as deemed necessary. Additionally, on national wildlife refuges and some other areas, annual water management plans are prepared for individual wetlands. These are prepared manually, and often can not take into account conditions in other areas of the flyway except in a general sense. Plans are not usually updated during the course of the year. The past and current holistic situation for management of trumpeter swans, of which planning is only a part, has not yet resulted in population recovery for Tri-State swans. New planning approaches are welcome, and an approximately 80 percent increase in breeding pairs is still desired.

Requirements Analysis and Cooperative Distributed Problem Solving

Our research is pursuing three objectives. (1.) We intend to provide a decision support system that allows swan managers to examine management actions addressing population and migration objectives at a flyway scale, and allows them to evaluate management actions at a site specific scale. (2.) We will test the hypothesis that decision support technology which allows planning in multiple geographic and temporal scales results in an increased ability for managers to identify and capitalize on trumpeter swan management potentials. For our technology to give managers this capability, we must verify that the decision support system simulates future swan distributions that meet flyway goals; that habitat recommendations are satisfactory for supporting increasing populations; and, all recommendations remain reliable over a specified time period. Management potentials are those ecological conditions that can be exploited in pursuing trumpeter swan objectives. Included are habitat quality, quantity, distribution, and availability, as well as freedom from disturbance. (3.) We will test the hypothesis that our implementation of cooperative distributed problem solving among refuges, parks, other management areas, and the internal knowledge bases effectively integrates local management actions with small-scale landscapes. This integration will occur if information is shared among human and electronic nodes, if individual knowledge bases contribute to recommendations, and if principles of adaptive management (Holling 1978; Walters 1986) are incorporated.

Based on input from swan managers, we have identified four management questions to be addressed through decision support system simulations. Each of these is a relatively course-grained approach to extrapolate possible future scenarios, while retaining the need to address the practicality of the fine-grained needs of individual managers. This is being tackled by paying close attention to knowledge engineering efforts and the use of expert systems to connect the relatively qualitative knowledge of the domain experts with the heuristic guidance needed by managers.

Simulation #1. If a particular management action is implemented at a particular site and particular time, what are the consequences for that site and for other sites in the flyway?

Simulation #2. Given an objective for spatial and temporal distribution of swans, what is the best set of management actions across all sites to achieve this? The decision support system will also have the capability for the manager to provide an alternative objective.

Simulation #3. Given some subset of management action(s) across all sites, and given an objective for spatial and temporal distributions of swans, what is the best complementary subset of management actions at other sites to achieve this?

Simulation #4. Given a satisfactory set of management actions across all sites to achieve an objective for swan distribution, if an alternative management action were to be implemented at a particular site, what are the consequences for that site and for other sites in the flyway in terms of reaching their respective objectives?

To address Simulation #1, a blackboard approach will be taken. When a particular management action is proposed for a particular site, that information will be posted to the blackboard. Daemons residing there will fire as necessary to activate the use of appropriate rules and expert systems to simulate the effects of the proposed action for the current time at that site, as well as at other sites in the flyway. New and impending constraints that the proposed action will impose on future management will also be generated and presented.

To address Simulations #2-4, a more complex search of the solution space will be required, and cooperative distributed problem solving will be used. Each geographic node in the system will need to function both independently and collaboratively with themselves and with the knowledge bases, exchanging "tentative and partial results in order to converge on a solution" (Carver, Cvetanovic, and Lesser 1991). These more complex simulations will require the concurrent development and posting of partially completed plans and potential management options from all geographic sites. The goal is to find a satisfactory set of solutions for management at all sites. This will be done by sharing information among geographic nodes and with the knowledge bases and databases. Then, a recursive search will be made for a set of management options that satisfices the population level and distribution objectives, and that addresses the constraints in the system.

The swan decision support system will use a combination of artificial intelligence methods including expert systems, blackboards (Corkill 1991; Nii 1986a, 1986b), and cooperative distributed problem solving (Carver, Cvetanovic, and Lesser 1991; Durfee, Lesser, and Corkill 1989). Four basic modules form the system's framework (Figure 2): cooperative distributed problem solving, knowledge bases (expert systems), databases, and web interface. The decision space consists of knowledge and constraints, including population objectives, on-the-ground management capabilities, ecological principles, and implementations of adaptive management. In addition, an area's past management history, as well as its future needs, represent further temporal constraints to forming recommendations in the present, particularly related to wetland manipulations.

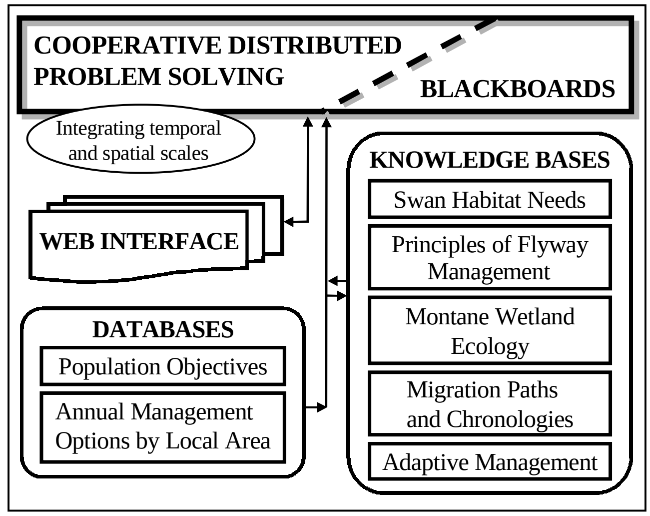

Figure 2. The framework for the swan management decision support system.

The essence of the distributed nature of swan ecology stems from birds moving among areas as seasons and other ecological conditions change, especially habitat availability. Migration stimuli are also related to annual life cycle events and physiological condition in individual swans. Meta-rules for handling the integration of all such spatial and temporal issues in the domain will be developed and integrated at a high level in the system.

It is clear to us that wetland ecology is a domain where the complexity of relationships, the interactions among ecological parameters, and the lack of empirical data makes the programming of rule-bases and decision trees complicated. By the same token, we are becoming increasingly convinced that the complexity of ecological systems is, in fact, what makes the application of expert systems, cooperative distributed problem solving, and other artificial intelligence methods so potentially useful. This domain has a nearly infinite number of ecological conditions, but the number of potential recommendations is more limited. Backward chaining approaches are allowing us to appropriately search the decision space in a goal-directed manner. Artificial intelligence based multi-agent methods are another approach that might be used for such a planning problem. However, their contribution often lies in searching exceptionally large and dispersed information sources, in providing real-time solutions, or utilizing the reasoning power of individual agents. None of these attributes exists in our domain. On the other hand, there are some similarities in our approach to the asynchronous backtracking algorithm presented by Armstrong and Durfee (1997), except that we are not using a complex agent implementation.

Status of System Implementation

We are developing the system on a personal computer using a commercially available expert system development shell that has blackboard capabilities. The system is deployed on a Unix workstation acting as a web server connected directly to the Internet through Montana State University. This is accomplished using software affiliated with the development shell that provides a common gateway interface between the web server and the inference engine, developing HTML web pages on the fly.

Our primary goal is to explore whether cooperative distributed problem solving can solve actual ecological problems characterized by geographically distributed issues that are compounded by temporal scales. There were several, general institutional concerns governing our selection of technologies. These included palatability to end users, availability of off-the-shelf software, probability of long-term software support, and cost. Following the scheme describing expert system use and research provided by Hollnagel (1991), our project is using known methods and addressing unknown problems. However, this categorization is not clear-cut because, to our knowledge, the application of cooperative distributed problem solving has not been implemented using our current software.

We have developed knowledge bases for swan habitat needs and management of montane wetlands. Each of our knowledge engineering sessions was approximately three days in length, and utilized one to two experts each. In some cases, one of the experts has been involved in previous expert system development, making knowledge acquisition efforts relatively easy. In particular, this individual was more apt to provide us detailed chains of logic in his reasoning without us needing to continually prompt him to do so. One technique that we used extensively was to provide the experts with detailed slide shows of actual field situations depicting wetland condition and management options. This seemed quite effective. It continues to be difficult to have our experts delineate their level of confidence in pieces of knowledge so that we might assign uncertainties within a knowledge base. Although our only evaluation to date has been qualitative, we have been pleased with the acceptance of the knowledge bases that we have demonstrated to swan managers.

Looking towards the future, there are some issues that we envision will be particularly challenging. First, developing the rules to implement cooperative distributed problem solving as a specific expert system, in essence a meta-system guiding the rest, has never been tried in this type of ecological venue. We are examining a number of ways to utilize blackboard algorithms (Carver and Lesser 1992) in domains such as ours. The multi-agent system of Pinson, Louca, and Moraitis (1997) which includes artificial agents, blackboards, and a constraint base may hold promise. Similar to their system, ours will be able both to make satisficing recommendations and to present incompatible management options through the use of subgoals. Our subgoals are represented by output from the knowledge bases as well as partially completed habitat plans for individual management areas (Figure 2.) Local control structures on the blackboard will critique and assemble partial plans from individual users, using rules to determine when knowledge base or database interaction is necessary. The scheduling of when such knowledge base or database output is necessary will be handled at a meta-level similar to that described by Maitre and Laasri (1990). Our constraint satisfaction approach will be implemented as knowledge bases invoked as part of the meta-level control structure.

Another challenge will be determining the best way to propagate uncertainties in the system. In the development of the current prototype, we intend to accept uncertainties provided as output from individual modules at face value. Then, the global propagation issues will be tackled within the cooperative distributed problem solving algorithm, allowing each knowledge base module to pass its own internal confidence assessment, essentially unchallenged, to the broader system. Future system development may address uncertainty issues within individual modules.

Finally, although empirical evaluation of the system is planned, we anticipate that it will not be straightforward. Gold standards for validation of various ecological components and models do not exist, plus the system will be predicting and guiding future scenarios as they, in fact, unfold. And, from a holistic perspective, we are developing decision support for issues that appear to beyond the capability of single persons to conceptualize and solve.

Acknowledgements

The ecological expertise and insights of L. Fredrickson, T. Grant, D. Hamilton, and J Kadlec have been invaluable during knowledge acquisition. They, along with J. Cornely, D. Dean, and J. Loomis have guided the development of this project from its inception. We are grateful to the Biological Resources Division of the U.S. Geological Survey and the U.S. Fish and Wildlife Service for funding, and accordingly note the encouragement of R. Stendell.

References

Armstrong, A., and E. Durfee. 1997. Dynamic prioritization of complex agents in distributed constraint satisfaction problems. *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 620-625.

Boland, R. J., A. K. Mahewshwari, D. Te'eni, D. G. Schwartz, and R. V. Tenkasi. 1992. Sharing Perspectives in Distributed Decision Making. In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 306-313. New York, NY: Association for Computing Machinery.

Brehmer, B. 1991. Distributed Decision Making: Some Notes on the Literature. In *Distributed Decision Making: Cognitive Models for Cooperative Work*, 3-14, J. Rasmussen, B. Brehmer, and J. Leplat, eds. Chichester, England: John Wiley and Sons.

Carver, N., Z. Cvetanovic, and V. Lesser. 1991. Sophisticated Cooperation in FA/C Distributed Problem Solving Systems. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 191-198, Menlo Park, Calif.: AAAI Press.

Carver, N., and Lesser, V. 1992. The Evolution of Blackboard Control Architectures CMPSCI Technical Report 92-71, Department of Computer Science, Univ. Massachusetts-Amherst.

Corkill, D. D. 1991. Blackboard Systems. *AI Expert* 6(9): 40-47.

Durfee, E. H., V. R. Lesser, and D. D. Corkill. 1989. Cooperative Distributed Problem Solving. In *The Handbook of Artificial Intelligence, Vol. IV*, 84-147, A. Barr, P. R. Cohen, and E. A. Feigenbaum, eds. Reading, Mass.: Addison-Wesley.

Hollnagel, E. 1991. The pragmatic and the academic view on expert systems. *Expert Systems With Applications 3:179-185*.

Holling, C. S. 1978. *Adaptive Environmental Assessment and Management*. New York, NY: John-Wiley & Sons.

Maitre, B., and Laasri, H., 1990. Cooperating Expert Problem-Solving in Blackboard Systems: ATOME Case Study. In *Decentralized A.I.: Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 251-263, Y. Demazeau, ed. Cambridge, England: Elsevier Science Publishers.

Nii, H. P. 1986a. Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures. *AI Magazine* 7(2):38-53.

Nii, H. P. 1986b. Blackboard Systems: Blackboard Application Systems, Blackboard Systems from a Knowledge Engineering Perspective. *AI Magazine* 7(3):82-106.

Pinson, S. D., J. A. Louca, and P. Moraitis. 1997. A Distributed Decision Support System for Strategic Planning. *Decision Support Systems* 20:35-51.

Sojda, R. S., D. J. Dean, and A. E. Howe. 1994. A Decision Support System for Wetland Management on National Wildlife Refuges. *AI Applications* 8(2):44-50.

Walters, C. 1986. *Adaptive Management of Renewable Resources*. New York, NY: Macmillan Publishing Co.