

Networked Information Retrieval as Distributed Problem Solving

Tim Oates, M. V. Nagendra Prasad and Victor R. Lesser*

Computer Science Department, LGRC

University of Massachusetts

Box 34610

Amherst, MA 01003-4610

{oates,nagendra,lesser}@cs.umass.edu

Recent years have seen an explosion in both the amount of information available in electronic form and the complexity of networked environments, forcing the developers of information retrieval (IR) systems to re-evaluate their model of the world. Early IR systems assumed that users would supply both a query and corpus (data source) against which a query is to run. Users may have access to multiple corpora managed by a set of information servers that are heterogeneous in terms of their level of sophistication and the types of data they manage. In spite of this complexity, the user, rather than the IR system, is tasked with knowing which corpora are most likely to contain the correct answer. When a user has access to an information carrying environment as rich and complex as the Internet, it is no longer possible for them to know where the various bits of information that comprise a complete and coherent response to a query may be found, or to effectively manage the search for information so as to minimize resource usage while maximizing the quality of retrieved information. Clearly, something more is required of the IR system. The problem, as described, seems amenable to an agent based solution. In response to a query, multiple autonomous agents can be released to search this distributed “information space” in a cooperative manner for relevant items. Top level queries drive the creation of partially elaborated information gathering plans, resulting in the employment of multiple *semi-autonomous*, *cooperative* agents for the purpose of achieving goals and subgoals within those plans¹.

The task of information retrieval in a distributed setting can be viewed in general terms as either distributed processing or distributed problem solving (DPS). We argue that distributed information retrieval (DIR) in complex networked environments is by its very nature an instance of DPS. Distributed processing is characterized by complete independence of subproblems. Agents need nothing other than local information to arrive at a subprob-

*This material is based upon work supported by the NSF Center for Intelligent Information Retrieval, ONR contract N00014-92-J-1450, ARPA under ONR contract N00014-92-J-1698. The content of the information does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

¹For a more formal model that characterizes the environment and problem solving, please refer to [13].

lem solution of the required quality which can be synthesized with other agent subproblem solutions to arrive at a global solution. Distributed problem solving, on the other hand, is characterized by the existence of interdependencies between subproblems assigned to the individual agents, leading to a need for them to cooperate extensively during problem solving. They rely on communication to detect and exploit these interdependencies between subproblems. Subproblem interdependencies occur naturally in DIR tasks, where the identification of relevant corpora and the composition of retrieved data into a coherent response to a query are both difficult problems. Agents must cooperatively exchange information about relevant corpora to focus the retrieval process, and must communicate partial results of sub-queries so that a consistent response to the top level query can be pieced together from disparate sources. Initially, agents have only partial and incomplete views of global solution requirements. In spite of this deficiency in information, they may arrive at partial and tentative results which may be exchanged by the agents working on subproblems that are interdependent, to reduce the uncertainty that surrounds local problem solving. That is, agents may exploit the interdependencies between subproblems to their benefit. This is the essence of the functionally accurate, cooperative (FA/C) paradigm presented by Lesser et. al. [11, 12] as an approach to distributed problem solving.

Below, we will present the FA/C paradigm in more detail, using examples from distributed information retrieval to highlight its features. Imagine a user deciding to go on a vacation. She gives a travel planner program a few of her preferences for the vacation. Say she gives her specifications as a vacation for 3 or 4 days around July 20th, preferably in Massachusetts. In addition, either through user specifications or through user modeling, the travel planner knows that this particular user prefers historical sites or nature spots. The travel planner has to plan for at least four different aspects of the vacation — places/sites to visit, weather situation, accommodation and conveyance. So it sends off four agents to deal with the corresponding types of databases. However, the agents have to interact during the process of information retrieval and planning, both among themselves and with the travel planner.

Complex distributed search spaces, such as the one in the example above, are characterized by various soft and hard constraining *goal interrelationships* like facilitates, enables, overlaps and subsumes (see Figure 1). The ability to exploit these interrelationships to avoid negative interactions and take advantage of positive interactions can enhance the search quality by providing better solutions in possibly lesser time. In a distributed intelligent information retrieval (DIIR) task, potentially useful constraints may exist between different pieces of information, either via content or as a function of problem solving activity. The discovery and exploitation of such constraints is necessarily a dynamic and incremental process that occurs during problem-solving and entails communication of partial results among agents in a timely and selective manner. Such communication augments each agent's local view with a more global view of emerging solutions and of the current state of searches in terms of what has been examined and what is proposed to be examined. Figure 1 shows some of the subgoal interrelationships in the travel planning example. A facilitates interrelationship implies that the values of a solution output parameter of the facilitating goal can, in some way, determine optional goal specification parameters of the facilitated goal. The goal of planning flight connections and car rentals has weather report data as an optional goal specification parameter to achieve this goal. Availability of weather reports facilitates

the planning process for car rentals and flight connections by eliminating or attributing low importance to retrieval of flight reservation and car rental availability on those days when the weather is not conducive to travel. Even though Agent 2 can plan for car rentals and flights without weather data, it may make more efficient use of resources, if time pressure on when the answer is needed by the user is not tight, by delaying that planning until weather information is available. In the meanwhile, it can search for airlines offering cheaper deals. Along similar lines, retrieval of weather reports facilitates planning for outdoor spots. On the other hand, a plan of the places to visit will enable the accommodation agent to start its work on planning for places to stay during nights. An enables interrelationship implies that the enabling goal produces a solution output parameter value that determines a required goal specification parameter of the enabled goal. The place at which accommodations are to be secured is a required goal specification parameter for the "Setup Accommodation" goal. Note that abstractions of plans are all that is needed for the accommodation agent to start its work. Thus, while agent 2 is fleshing out the details of the abstract plan of the places to visit, agent 1 can, in parallel, start its work. Similarly, a favors interrelationship says that once you made the effort to design a plan to go to Concord, a plan for going to Lexington is obtained by minor modifications to the plan for Concord. An overlaps interrelationship says that the two agents involved may be doing similar work and can hence benefit by sharing their partial results. An overlaps interrelationship exists between two goals that share determinants of some of their solution output parameters.

Given the incomplete nature of the local views of the individual agents, another important aspect of FA/C systems is the explicit recognition of the role of solution and control uncertainty. Coupled with the fact that resources and time for conducting a search are limited in real-life problems, this leads to the notion of *satisficing search*. The environment in an information acquisition task is characterized by the fact that the supply of available data is almost limitless, whereas time, money and computational resources are not. Rather than being able to retrieve an exhaustively complete and accurate response to a query, intermediate results from disparate sources must be pieced together to form consistent islands of information that can be incrementally refined to form a more accurate solution depending on the extent of available resources and time.

Another aspect of the FA/C paradigm is the explicit recognition and exploitation or avoidance of *redundancy*, leading to increased robustness or decreased resource demands depending on the context and the structure of the domain. For example, the Internet may contain "mirror" sites for certain data repositories or it may contain redundant data from different sources for the same task. Data from different sources may be of different quality or may be organized differently. A particular task could possibly do with low quality data which perhaps could be locally acquired. Redundancy could be permitted if the control costs outweigh the benefits of avoiding it. Alternatively, if we are dealing with faulty systems or poor quality data, redundancy could help resolve the uncertainty in the retrieved data by providing additional constraints.

We believe that distributed information retrieval tasks characterized by complexity, heterogeneity and unstructured data environments, can be viewed as a distributed problem-solving task within the FA/C paradigm. The benefits of such a view not only stem from the fact that it provides a comprehensive conceptual model for the myriad of methods being proposed for intelligent information retrieval (IIR), but also from the fact that the view pro-

vides a direct map from the wealth of existing methods in distributed Artificial Intelligence (DAI) to the distributed information retrieval domain. These methods have evolved over more than a decade, since the time the FA/C paradigm was first proposed[11]. Below, we discuss various techniques and systems from DAI that may have direct bearing on DIIR viewed as a DPS task. These methods were originally proposed in contexts different from information retrieval, and most of them were developed as techniques to study, understand and exploit various aspects of the FA/C paradigm.

This paper draws upon a long history of thought in DAI to present a model of cooperative retrieval of distributed information. We briefly touch upon various implemented aspects of FA/C that have direct relevance to the DIIR task. Decker and Lesser[6, 7] provide detailed studies of quantitative trade-offs involved in explicit recognition and exploitation of task interrelationships for use in multi-agent coordination. Garvey and Lesser[8] discuss design-to-time algorithms which basically endow the local problem solver with abilities to deal with real-time considerations and goal inter-dependencies. Such a scheduler is, perforce, satisficing in the solutions it provides and relies on the use of approximate processing techniques. Carver and Lesser[4, 5] present RESUN and its distributed derivative DRESUN as architectures that explicitly recognize and resolve uncertainties associated with the partial, evolving solutions in the interpretation domain. Interpretation is viewed as an incremental process of resolving sources of uncertainty (SOUs) through directed and intentional accrual of evidence. This seamless integration of data-driven bottom-up and goal-driven top-down processes opens up a huge set of opportunities for IIR. Information on hand can in turn serve to instantiate and actively direct further retrieval process to resolve the deficiencies in the partial data.

We hope that this paper encourages IR system designers to take a radically new view of distributed information retrieval as a distributed problem solving activity. Distributed information retrieval systems proposed in the literature typically either do not fully exploit the potential of knowledge-intensive methods for retrieval[1, 2, 3, 9] or tend toward distributed processing, failing to exploit the dependencies between agents working on different aspects of a retrieval task[10, 14]. Cooperating to enhance efficiency of a resource-limited retrieval process or negotiating to dynamically resolve conflicts and inconsistencies in the retrieved data, leading to further search or retrieval, may be important aspects of DIIR systems in the future. Our model is an initial step in this direction. In fact, this complex view may be more appropriately termed information acquisition - an incremental plan-based information seeking activity. We also suggest that existing methods in DAI can serve to leverage future implementations of DIR systems based on this view.

References

- [1] M. C. Bowman, P. B. Danzig, U. Manber, and M. F. Schwartz, "Scalable Internet Resource Discovery: Research Problems and Approaches", *Communications of the ACM*, 37(8), 1994, pp 98 - 107, cntd on 114.
- [2] J. P. Callan, W. B. Croft, and S. M. Harding, "The INQUERY Retrieval System", in *Proceedings of the 3rd International Conference on Database and Expert Systems Applications*, pp 78-83, 1992.




Figure 1: Goal Tree for the Vacation Planning Example

- [3] J. Callan, *Personal Communication*, July 1994.
- [4] N. Carver and V. Lesser., "A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty", in *Proceedings of AAAI-91*, 1991, pp 724-731.
- [5] N. Carver, Z. Cvetanovic, and V. Lesser., "Sophisticated Cooperation in FA/C Distributed Problem Solving Systems", in *Proceedings of AAAI-91*, 1991, pp 191-198.
- [6] K. S. Decker and V. R. Lesser, "Generalizing the Partial Global Planning Algorithm", *International Journal of Intelligent and Cooperative Information Systems* 1(2), June 1992, pp 319-346.
- [7] K. Decker and V. R. Lesser, "Quantitative Modeling of Complex Computational Task Environments", in *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pp 217-224, Washington, 1993.
- [8] A. Garvey and V. R. Lesser, "Design-to-time Real-Time Scheduling", *IEEE Transactions on Systems, Man, and Cybernetics: Special Issue on Scheduling, Planning, and Control*, 23(6), 1993.
- [9] M. Huhns, U. Mukhopadhyay, L. M. Stephens, and R. Bonnell, "DAI for Document Retrieval: The MINDS Project", in *Distributed Artificial Intelligence* ed. by M. Huhns, Pittman Publishing/Morgan Kauffmann Pub., pp. 249-284.
- [10] C. A. Knoblock, and Y. Arens, "An Architecture for Information Retrieval Agents", in *Working Notes of the AAAI Spring Symposium on Software Agents*, 1994, pp. 49-56.

- [11] V.R. Lesser and D. Corkill., "Functionally accurate, cooperative distributed systems", in *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1), 1981, pp. 81-96.
- [12] V.R. Lesser., "A retrospective view of FA/C distributed problem solving", *IEEE Systems, Man, and Cybernetics*, 21(6), 1991, pp. 1346-1363.
- [13] T. Oates, M. V. Nagendra Prasad, V. R. Lesser., "A Distributed Problem Solving View of Networked Information Retrieval", Forthcoming Technical Report, Department of Computer Science, University of Massachusetts at Amherst, 1994.
- [14] E. M. Vorhees, "Software Agents for Information Retrieval", in *Working Notes of the AAAI Spring Symposium on Software Agents*, 1994, pp. 126-129.