IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS, VOL. SMC-11, NO. 1, JANUARY 1981

Natural and Social System Metaphors for Distributed Problem Solving: Introduction to the Issue

B. CHANDRASEKARAN, SENIOR MEMBER, IEEE

Abstract—Naturally occurring information systems provide a number of useful metaphors for distributed problem solving. An introduction to some aspects of these metaphors is given which simultaneously serves as a guest editorial for a special issue devoted to the topic. The ubiquity of a distributed mode of computation in information processing in natural phenomena in general and in human societies in particular is observed and related to the evolutionary and complexity-reducing advantages of this mode. The forms of communication media available to coordinate the problem solving activities of the individual processors are examined. Some general remarks are made on how problem solving is distributed and coordinated in some human organizations, and the potential usefulness of the "society of specialists" notion in explicating cognitive activity is pointed out. Along the way, the contents of the papers in the special issue are considered in relation to various points raised in the discussion.

I. UBIQUITY OF DISTRIBUTED INFORMATION PROCESSING

ISSUES about *distributed* computing, as about any other aspect of computing, can be formulated at various levels of abstraction. Each level has a different conceptual content, and raises a correspondingly different set of issues. In distributed computing most of the recent emphasis has been at a level that is closely related to physical connection of different processors, secure transmission of data among them, and the corresponding operating system problems of scheduling different processors. These issues have dominated discussion so much that the term distributed processing has come to mean almost exclusively that set of issues. The papers in this special issue deal with distributed processing at a different, "higher," level of abstraction. The questions of interest at this level concern the strategies by which the decomposition and coordination of computation in a distributed system are matched to the structural demands of the task domain. Distributed problem solving (DPS) is an appropriate term for the phenomena at this level of abstraction.

As the theme of the issue implies, a motivating belief is that information processing phenomena that occur in the natural world are a source of a number of useful metaphors for distributed processing in general and distributed problem solving in particular. It is clear that distribution of processing or computation is an intrinsic characteristic of most natural phenomena which can be captured within a computational or symbol processing framework. Social organizations from honeybee colonies to a modern corporation, from bureaucracies to medical communities, from committees to representative democracies are living examples of distributed information processing embodying a variety of strategies of decomposition and coordination. Computation in biological brains, especially in their sensory processors such as vision systems, displays a high degree of distribution. There is substantial evidence that higher cortical functions are also computed (and controlled) in the brain in an essentially distributed mode: regions have been identified in the cortex whose activities are highly correlated with specific higher cortical functions such as language processing. Geschwind [1] indicates that some regions in the brain are extremely specialized: there is an identifiable processor which specializes in human face recognition!

Control of movements in biological systems is also accomplished by distributed computation [2]. Evidence is available that control of normal walking movements resides in the spinal cord [3]. Volitional movement can be viewed as being generated by low-level programs coordinated and regulated by higher level controllers. The task of generating all the impulses for all muscle fibers for each movement is surely beyond the resources of any centralized biological movement processor.

In all these examples—from social organizations to brains and motor systems—the overall computational task is distributed among a collection of separate processors. These separate processors coordinate their computations by means of exchanging appropriate symbolic information.

II. WHY (AND WHY NOT) DISTRIBUTION

A. *Advantages of Distribution*

Why should distributed computation be such a ubiquitious mode in naturally evolved information processing systems? The following advantages of distribution may be relevant here.

1) *Decomposition* of processing is an absolutely basic strategy for controlling the *complexity* of computation. Central computation is just too costly in both memory and time. Distributing the computation among different

---

Manuscript received September 20, 1980. This work was supported in part by the Air Force Office of Scientific Research under Contract F 49620-79c-0152.

The author is with the Artificial Intelligence Group, Department of Computer and Information Science, The Ohio State University, Columbus, OH 43210.

0018-9472/81/0100-0001$00.75 ©1981 IEEE

processors generates possibilities for parallel activities by different processors which may be able to work on essentially nonoverlapping segments of the data most of the time. Also the scope of each processor is limited, i.e., the size of the input domain is much smaller. Complexity of computation is often an exponential function of input space size.

2) Appropriately distributed computing increases the prospects for graceful degradation of response when there is degradation of input data or failure of portions of the system.

3) Distribution is a natural attribute of *evolutionary* systems. As the system grows and increases in complexity, a distributed mode provides for replacing a processor with several processors and making *mostly* local changes in linkage among processors; or, as the external environment changes, distributed information processing makes *adaptation to change* easier, since again, as long as the rate of external change is not large, changes to the system can be mostly local, if the original decomposition reflected the structure of the task environment correctly. As programmers well know, these are really advantages of *modularity*, but a distributed architecture provides a natural means of implementing this modularity.

4) In complex information processing systems involving very large numbers of *sensors* and *effectors*, a central processor will require *very large bandwidths* for responding to sensors or activating effectors. Imagine an army whose commanding general alone is authorized to make *all* the field decisions!

5) Often a task decomposition will lead to the generation of a large number of identical subtasks, each, however, operating over different subsets of input data or regions of the environment. In this case, once a particular processor is optimized to be effective for such a subtask, in a distributed approach this processor can be *replicated* as often as needed. This provides for considerably increased efficiency due to parallelism, and is especially useful in systems that deal with large volumes of sensory information. However, it is useful in a variety of other situations also. An example is the optimization of the training of salespersons in a commercial organization and consequent production of large numbers of them.

B. *Possible Costs of Distribution*

Of course there could sometimes be a cost associated with distribution as opposed to a central computation. The computation of each processor is often a filtering operation, in that it communicates only the *result* of its computation. (Thus in a vision system a higher level processor may not have direct access to the image intensity data, but only to the outputs of edge detectors operating on these data.) Another processor may arrive at a different result with the same data if it did not depend on the computation by the other processor. Several strategies are used in naturally occurring distributed systems to deal with this problem. Marr [4] talks about the *principle of least commitment* one way by which the processors at a lower level in a visio system may be constrained from introducing too much of filtering. This principle suggests decomposing the proble in such a manner that at any level commitments are mad in a conservative fashion, i.e., to the least abstract enti that is necessary. If this principle is applied at each level abstraction carefully, the processors at the higher levels abstraction will have available to them generally reliab information from the lower level processors. The penal this extracts is a certain profligacy with respect to th number of processors, since typically this principle wou lead to an increase in the number of levels of abstractio Even with this principle, processors at a lower level abstraction may still be forced to make some commitmen which are not quite correct in specific instances, though the average it may be a reasonable thing to do. This sort thing explains certain kinds of visual illusions and visu effects, e.g., the "sun" effect in [4]. That is, visual illusio show dramatically the commitments made by lower lev processors that happen not to be warranted for speci classes of situations.

Another strategy is a sort of *local relaxation* by which th results of contiguous processors dealing with data in neighborhood are compared for consistency with ea other, and a processor's result would be ignored by higher level processor if it is substantially deviant fro those of its neighbors. This, of course, is a double-edg strategy, since in the occasional instances in which the deviant processor is correct, its result nevertheless does n get passed up to the higher level processors.

C. *From Committees to Hierarchies*

The fear of this filtering and consequent bias is at th heart of full participatory democracies like the city-sta democracy of ancient Athens, where all the citizens vot on almost all the issues. However, as the size and numb of issues grow large, this form of information processi begins to place great burdens on the available bandwidt and the democratic processes become more organized, a abstractions in the form of *representation of constituenc* begin to come about. These in turn produce the proble of filtering mentioned above. They also begin to manif another consequence of evolutionary distributed systems: tendency to swamp out changes which are local in space time. This is the other side of the coin of robustness th natural distributed systems often display. The changes the environment or input have to be sufficiently large overcome the filtering and the abstractions made by t various processors at levels of abstractions close to t sensory data.

The contribution of Wesson *et al.* for this special iss considers an experimental comparison of two distribut architectures for a message puzzle task, where a network human sensors, each of whom sees only a small portion a two-dimensional environment, attempts to interpret One arrangement was an "anarchic committee" archit ure— somewhat like the city-state democracy— where all he nodes were free to communicate with each other. The econd arrangement was a hierarchical one. For the particlar collection of experiments that were conducted, the ommittee architecture worked better than the hierarchies. everal aspects of the experiment are worth noting. The ize of the environment as well as the sensor network was mall. It would be interesting to see if the result would hold or larger size environments and large networks of nodes. f the observations in the preceding paragraph are correct, hen one would expect a gradual shift in favor of a more ierarchical arrangement with increasing size. Secondly, ierarchies may be appropriate only when the environment as a sufficient amount of structure. For example, when ifferent regions of the environment correspond to differnt identifiable configurations, groups of sensors will need o exchange information only within each group to identify e local configuration. The bias effect that we discussed arlier would be most pronounced when the architecture is ot matched to the structure of the environment.

III. INTERPROCESSOR COMMUNICATION

It seems reasonable to suppose that different architecures of distribution would emerge depending upon the osts of communication among processors or, equivalently, pon available bandwidth. When a multiplicity of media ith differing bandwidths and accessibilities is available, e architecture of a naturally occurring evolutionary disibuted system would be organized so as to use the availale bandwidth most effectively. A modern corporation ith telephones, radio, and other media available has a ifferent architecture than one in times or regions with ore primitive communication structures. Increased bandidth availability would seem to decrease processor autonmy, i.e., place greater constraints on the amount of filterg allowed at the local processor level. This, however, ould typically be counteracted by the increased burdens n the top-level processors that an overload of information ill pose. So the degree of autonomy is a balance between ese contending tendencies in distributed natural systems.

The communication media available can be categorized to two broad classes: one which is used by senders and ceivers who know the identity of each other, and another hich has more of a *broadcast* character, and is more ssociative in nature, i.e., the receiver uses whatever information in the medium that it deems appropriate to its eeds. In large distributed systems, it is not practical, even the bandwidth were available, for the first class of ommunication media to be used without constraint. For ne thing, as the size of the system grows, the directory size r each processor would grow rapidly, burdening the formation processing capabilities of the processor. For nother, as processors are deleted or added in response to cal changes, the directories will have to be updated all ver the system. This collection of constraints typically ads to this class of media being used within a narrow and cal scope. Most often the communication is hierarchical, thus eliminating directory size and updating problems. Information meant for or needed from other processors is directed to the broadcast media, such as *blackboards* or *journals*. However, once again, this cannot also be done with abandon, since these media will then be clogged with the outputs of the large number of processors, making it useless for receivers, unless they are willing to invest a large portion of their limited processing resources to sort out the clutter. Extremely careful and powerful abstractions, closely matching the structure of the class of tasks for which the distributed system is designed, will need to be generated as appropriate inputs to the broadcast media. Understanding how this is done is a central theoretical enterprise. We shall later comment upon the use of blackboards by three papers in this issue: those by Gomez and Chandrasekaran, Rieger and Small, and Cullingford.

IV. DPS IN HUMAN COMMUNITIES

As we mentioned earlier decomposition is the basic weapon against complexity. Thus when a socially important task is too complex for individual humans, organizations with a number of humans evolve whose architecture matches the structure of the task, and whose total computational capacity is adequate for it. Task decomposition in human organizations often provides a great deal of parallelism, which is conceptually and operationally important for increased efficiency. Several hundred years ago, each competent physician possessed almost all of the medical knowledge then available. Today the complexity of medical knowledge has resulted in the creation of a complex organization of specialists, where no one knows more than a small part, but the community overall advances medical knowledge and provides care. The modern corporation is often a large distributed system with a large number of specialized subdivisions which, when successful, mesh together in a miracle of purposefulness, but when the overall structure strays too far from the changing environment, it resembles a maladaptive dinosaur (see comments on adaptation and distribution in Section II). The scientific community is another human organization whose architecture has evolved in a distributed fashion, the shaping forces in this case being the dictates of the scientific method, and the communication requirements for the creation and verification of new scientific knowledge. There are countless other examples of human organizations.

Gomez and Chandrasekaran's work in this issue concentrates on the epistemological structure of medical diagnosis, which is independent of whether the task is accomplished by a single human, a community of specialists, or a collection of microprocessors. They relate the identity and structure of specialists to the *conceptual content* of the domain. The distributed problem solving that they propose has a great deal of parallelism in it.

In their explication of the scientific community metaphor, Kornfeld and Hewitt emphasize its inherent parallelism. They develop some concurrent language primitives to emulate some of the problem solving behavior of scientific communities. Fox proposes that, among the criteria for distribution in human organizations are *complexity*, *uncertainty*, and *resource constraints*. He considers the organization of a distributed system such as Hearsay-II [5], and studies the extent to which it already incorporates the insights of organization theory.

Markets are an interesting kind of distributed system. They are not primarily information processing systems, but they use a distributed, mostly local information exchange to achieve a certain kind of global optimality in resource utilization. They are the sources of the metaphors of *prices* (a kind of abstraction) and *contract and bids* (a kind of mechanism) that enable a global optimality to be reached over a period of time. The paper by Smith and Davis in this issue deals with some aspects of these metaphors and their use in distributed systems.

V. DPS IN COGNITIVE ACTIVITY

It is easy to conceive of distributed computing in the case of what are evidently communities of individual processors, such as ants building hills, armies, corporations, or the scientific community. What is less obvious is the utility of this conception in understanding the information processing of an individual human being. The metaphor of a society of little minds–the homunculi—has come up repeatedly in psychology and philosophy of mind. Dennett [6] gives a brief but useful account of the history of this metaphor. These models have floundered on the apparent infinite regress involved in explaining a mind by postulating a collection of minds. It has been only recently that, due to work in artificial intelligence, we can begin to see how "mind-like" is not an all-or-none affair, that more complex mind-like behavior can be obtained by the coordination of less complex mind-like entities. The less complex entities are specialists, i.e., they have a narrower scope. This kind of decomposition can be applied recursively at quite a few levels: Minsky [7] has recently formulated a "society of minds" model dealing with epistemological issues all the way down at the "neurodevelopmental" level, and Marr's work is very suggestive of how vision can be conceptualized as a society of specialists: groups of specialists all the way from very low-level ones (edge specialists) through those at slowly increasing levels of abstraction to high-level conceptual specialists.

When top-level control in a society of specialists is weak or nonexistent, the subordinate specialists may speak up in different, possibly conflicting, voices. "Multiple control" substitutes unitary control. Jaynes [8] recalls the *Iliad* and its heroes' ascription of many of their actions to the demands of gods whose voices they hear. He relates this to the evidence that many schizophrenics—examples of non-unitary consciousnesses—report hearing "inner voices" during acute attacks. Hilgard [9] discusses hypnosis as a breach of this unitary control. Freud's theories, of course, were based on the view of the mind as an interacting society of many agents. In this context it is tempting to speculate that one of the roles of consciousness is in providing a blackboard (we earlier discussed it as one form of communication medium), for the agents at the top level of cognition.

For our purposes, the foregoing speculations can b given a concrete cast by asking: What scientific value doe the distributed processing metaphor—or the society o minds notion—have in explicating high-level cognitiv phenomena? Gomez and Chandrasekaran in their paper i this issue point out that there is really not much that is new in the notion of a society of specialists as a model o complex computations: almost all large programs are modular and the modules are specialists. Thus they suggest tha for this metaphor to be technically useful criteria ar needed for *decomposing* tasks into specialists. Their pape provides such criteria for one well-defined class of cognitive activity, viz., diagnosis. The specialists are conceptua specialists who are hierarchically organized. In addition t the hierarchical communication, they also use a *blackboard*, i.e., a broadcast form of communication. The author provide an explicit account of the structure of the blackboard for this particular task.

Similarly, Rieger and Small propose a particular criterion to organize the specialists in parsing natural languag utterances. They propose that the specialists be "wor experts," and suggest how their activities should be coord nated. Again it is instructive that the experts use a broadcast communication medium, viz. the "control workspace." Cullingford, in his paper in this issue, considers the problem of integrating and controlling the experts in a system to "understand" a class of newspaper stories. Again th blackboard notion finds expression in his work. For th particular class of tasks considered by him, he propose that the blackboard data contain not only results compute by specialists, but in addition some control informatio i.e., some indication of *how* the blackboard information i changed, e.g., *who* changed an item. In addition to th papers in this issue, I should like to draw the reader attention to the work of Sacerdoti [10] who has considere how a distributed architecture might be designed for natural language understanding.

VI. CONCLUDING REMARKS

Hearsay-II [5] was one of the first large artificial intell gence (AI) systems to use an essentially distributed architecture for a complex problem and was the first system t use a blackboard as an interprocessor data structure. Th paper by Lesser and Corkill in this issue attempts t concretize some of the lessons from Hearsay-II and othe knowledge-base systems for the design of distribute processing systems. They concentrate on how uncertaint in input data and inaccurate processing by individua specialists can be compensated for by the collective so tha the whole is more robust than the individual processors They call their approach functionally accurate (referring t robustness of performance), cooperative (indicating som form of relaxation procedure—see Section II-C— by whic each contributes to the reduction of uncertainity of other processors) distributed systems.

Distributed problem solving, or more generally distributed artificial intelligence, is important conceptually, strategically, and practically. Its conceptual importance lies in that a distributed paradigm eludicates the structure of the processes of intelligence, such as vision, speech processing, or language understanding, whether or not an AI system is in fact implemented in a distributed manner. Its strategic importance arises from the fact that it is a good research strategy to look for decompositions of a complex problem. Its practical significance arises both from applications that are essentially distributed in nature, such as distributed sensor networks [11], as well as from the technological breakthroughs in microprocessors that make a distributed implementation of AI systems an elegantly practical possibility. I hope that the papers in this special issue direct attention to this important scientific and engineering enterprise.

Acknowledgment

I thank Andy Sage for a very helpful and cooperative role in his capacity as Editor of this Transactions, the uthors of all submitted papers for their cooperation, and he reviewers R. Bhaskar, R. Cullingford, R. Davis, M. Fox, P. Kugel, J. Reggia, C. Rieger, R. Smith, N. S. Sridharan, K. Stevens, L. Svobodova, S. Tanimoto, and S. W. Zucker—for their invaluable help. R. Bhaskar, Bruce Flinchbaugh, and Fernando Gomez made useful comments on earlier drafts of this introduction. I am grateful to David Zeltzer for information on biological movement control.

References

[1] N. Geschwind, "Neurological knowledge and complex behaviors," *Cognitive Sci.*, vol. 4, no. 2, pp. 185–193, 1980.

[2] N. Bernstein, *The Coordination and Regulation of Movement*, New York: Pergamon, 1967.

[3] K. Pearson, "The control of walking," *Scientific American*, vol. 235 (6), pp. 72–86, Dec. 1976.

[4] D. Marr, "Early processing of visual information," *Phil. Trans. Roy. Soc.*, vol. 275, ser. B 942, pp. 489–519, 1976.

[5] L. D. Erman, and V. R. Lesser, "The Hearsay-II system : A tutorial," in *Trends in Speech Recognition*, W. A. Lea, Ed.. Englewood Cliffs, N. J., Prentice-Hall, 1979.

[6] D. C. Dennett, "Artificial intelligence as philosophy and as psychology," in *Brainstorms*. Bradford Books, pp. 109–126, 1978.

[7] M. Minksy, "K-Lines: A theory of memory," *Cognitive Sci.* vol. 4, no. 2, pp. 117–133, 1980.

[8] J. Jaynes, *The Origin of Consciousness in the Breakdown of the Bicameral Mind*. Boston: Houghton Mifflin, 1976.

[9] E. R. Hilgard, *Divided Consciousness: Multiple Controls in Human Thought and Action*. New York: Wiley, 1977.

[10] E. Sacerdoti, "What language understanding research suggests about distributed artificial intelligence," in *Proc. Distributed Sensor Nets Workshop*, Carnegie-Mellon University, pp. 8–11, 1978.

[11] *Proc. Distributed Sensor Nets Workshop*, Carnegie-Mellon University, 1978.

Network Structures for Distributed Situation Assessment

ROBERT WESSON, FREDERICK HAYES-ROTH, JOHN W. BURGE, CATHLEEN STASZ,
AND CARL A. SUNSHINE

***Abstract*—A new approach to situation assessment is an automated istributed sensor network (DSN) consisting of many "intelligent" sensor evices that can pool their knowledge to achieve an accurate overall ssessment of a situation. Laboratory experiments were conducted to nvestigate potential DSN organizations and to ascertain some general esign principles. These experiments have been performed with a network f "sensor nodes," each of whom sees only a small portion of the entire nvironment and attempts to identify the environmental mobile entities as uickly as possible. To do this, they must cooperatively communicate their ypotheses and data, using a limited number of messages. Two general DSN organizations were tested. The first was hierarchical. The second was an "anarchic committee" whose nodes could each send messages to one, some, or all other nodes. The performance of the committee organization consistently surpassed the hierarchical one. This lent support to the contention that DSN architectures need to emphasize cooperative aspects of problem-solving. A machine-based simulation of such a network that achieved performance levels comparable to that of the human committee DSN organization was also constructed and tested. Because most situation assessment communications concern hypothesis updating and revision, minimizing communication requirements through the concept of active "hypothesis processes," which are responsible for predicting their own evolution over time, is suggested.**

I. Introduction

SITUATION ASSESSMENT (SA) involves acquiring, organizing, and abstracting information about the environment which may either correlate well with our expec-

---

Manuscript received January 24, 1980; revised August 26, 1980.

R. Wesson, F. Hayes-Roth, and C. Stasz are with The Rand Corporation, 1700 Main Street, Santa Monica, CA 90406.

J. W. Burge was with The Rand Corporation, Santa Monica, CA. He is ow with Operating Systems, Inc., Woodland Hills, CA 91364.

C. A. Sunshine was with The Rand Corporation, Santa Monica, CA. He now with the Information Science Institute, University of Southern alifornia, Marina Del Rey, CA 90291.