Computers Elect. Engng Vol. 19, No. 6, pp. v–viii, 1993
Printed in Great Britain. All rights reserved
0045-7906/93 $6.00 + 0.00
Copyright © 1993 Pergamon Press Ltd

GUEST EDITORIAL

PARALLEL AND DISTRIBUTED COMPUTING FOR INTELLIGENT SYSTEMS

1. INTRODUCTION

Although it is not generally possible to precisely define an intelligent system or what constitutes such a system, there seems to be an agreement that many real-life practical tasks can only be achieved by intelligent systems that are complex in structure and functionality. For example, the problem of navigating a mobile robot in unstructural environments requires a number of sensors since a single sensor is inadequate to achieve all but some trivial navigational tasks [1,2]. Here the information from the sensor must be suitably fused and incorporated into path planning and obstacle avoidance algorithms.

Several advances in both hardware and software of computing systems have resulted in significant steps towards the development of such intelligent systems as evidenced by a number of special issues of journals devoted to topics such as architectures [4], distributed sensor systems [5], sensor fusion [6], autonomous machines [7] (to name a few).

The methods of parallel and distributed computing have become known to aid the development of such intelligent systems by providing several practical solutions to a number of problems [8]. In particular, these systems are able to overcome some of the difficulties of single processor systems in a number of applications. Typically in the area of neural networks, a large collection of rather simple processing elements is employed to collectively solve a complex problem. Our framework here is more general in that we allow concurrent interactions between disparate and complex modules of a system; this will be typically the case of a teleoperated robotic system that consists of sensors, computers and activators all operating in a coordinated and concurrent manner.

The design and development of intelligent systems can be conceptualized in two basic paradigms that very closely interact with each other in a given complex application:

1. **Task-Oriented Problems.** We are required to accomplish a task that requires a complex system as a part of the solution. Here the problem is to first identify the characteristics of a system that can perform a given task. Then a detailed design and development of various components has to be performed; then these components have to be integrated into a working system.
2. **System-Oriented Problems.** We are given a complex system that had been designed to perform a task, and are required to solve a particular problem for this system. The objective could be adapting a system to perform a task that is not identical but related to the task for which the system is first designed. In another scenario, the objective could be to use the elemental capabilities of the system in an algorithm that achieves a complex task.

Typically in the first paradigm, we are looking for a class of systems that can perform certain tasks, and in the second paradigm, we are given a system for which a specific task must be efficiently performed.

Advances have been made in several fronts of intelligent systems which are typically complex machines composed of disparate modules that interact with each other in a coordinated manner. Some of the important research topics include the following:

1. **Synthesis of Complex Systems.** The issues of combining a collection of modules into an integrated system are very critical. In general the constitutent modules could be similar or disparate; the former mode is generally used for fault tolerance, and the latter is to overcome the capabilities of individual modules.
2. **Coordination and Communication.** The problem of coordinating various components of the system is critical to the overall operation of the system. Also the information from

various modules has to be suitably routed and consolidated in order to achieve the overall objectives.

3. **Information and Computational Complexity.** Several of the tasks that are performed on computing devices have to be analyzed for their complexity in order to decide which solutions are practical and feasible. If some of the problems turn out to be computationally intractable or undecidable, it would be appropriate to design approximate methods that run in a reasonable time and compute suboptimal solutions.

4. **Pardigms and Algorithms.** Efficient algorithms and methods to handle several of the tasks that are required to achieve the overall objectives of the entire system.

5. **Modeling and Analysis.** Since complex systems consist of various modules, some simplified and abstract models can be employed to highlight the relevant aspects. Such models could alleviate the high complexity that results from a large number of unimportant parameters.

All of these issues are discussed, to various degrees, in the papers of this Special Issue.

2. SUMMARY OF PAPERS

In this Special Issue, we have a very diverse collection of papers written by researchers at universities and national laboratories. Several aspects of intelligent systems are addressed in the framework of parallel and distributed computing.

The two of the most fundamental concepts of neurodynamics—irreversibility and creativity—are discussed by Zak. A very interesting architecture for neural systems is proposed here based on terminal noise which incorporates an element of irrationality into the dynamic behavior. In contrast to the conventional neural systems where the system "settles" to one of the "attractors", here the solution oscillates chaotically about some critical points in a manner qualitatively akin to the "classical" chaos. The phenomenological similarity between brain activity and the dynamics of the proposed system is to be emphasized: due to the terminal chaos, the dynamical system can be activated spontaneously driven by a global internal periodic rhythm.

Youssef studies the self-routing algorithms for Clos networks, which are three-level networks of elements that can connect the input lines to output lines in any permutation. The problem of realizing an arbitrary permutation of the input–output connections (to the entire network) is studied here. It has been illustrated that distributed algorithms that locally adjust the setting of each element can achieve very fast routing times with an extremely high probability. Both theoretical and empirical results are presented in this very interesting paper.

Several practical problems have been solved by employing suitable systems of neurons; most of the solutions are described in terms of a system of conceptual neurons. The actual solutions that employ these "conceptual" neural systems can be obtained by "simulating" these conceptual systems on suitable existing computer systems. Various issues of simulating neural networks on distributed memory architectures such as ring, hypercube, mesh and extended hypercube are studied by Patnaik and Mohan Kumar. They illustrate that communication-efficient networks of distributed memory systems perform better than several other topologies in simulating the artificial neural networks.

It has been established that many perceptual tasks require a system of sensors, and a single sensor is inadequate to accomplish reasonable tasks. A study of computational complexity of distributed detection problems with information constraints has been presented in the paper by Rao, Iyengar and Kashyap. They consider a simple detection system capable of identifying objects from a small finite set using a system of sensors. Even in this simple scenario many formulations, except for very simple ones, turn out to be computationally intractable, thereby prohibiting the exact and fast computational solutions for large systems.

In many of the existing neural systems the learning algorithms are implemented in software. Although there are several hardware implementations of neural systems, such implementations of learning mechanics are very few. Tawel describes a chip-in-the-hoop learning system assembled from custom analog building blocks hardware. Hardware implementations of learning algorithms based on gradient descent in feedforward networks and dynamically reconfigurable neural networks are described in detail in this paper.

The task-oriented problem of tracking multiple targets in dense threat environments is studied in the paper by Toomarian. This approach is based upon a continuum representation of a cluster of flying objects. Because of the impossibility of encounters in a high-density cluster, the velocities of flying objects are assumed to be embedded into a smooth velocity field. This hard tracking problem is then solved by reducing it to a problem in hydrodynamics.

In a complex system, the coordination of various disparate parts is very critical to the overall performance of the system. These isssues in the context of a system capable of shared/traded control of telerobots under time delay is described by Venkataraman and Hayati. A sharing strategy that combines the advantage of the shared and traded control capabilities is described; this work is geared towards a robotic system that will accept and execute commands from either a six-axis teleoperated device or an autonomous planner or a combination of the two. A two-tiered sheared control consisting of a task-level and servo-level is described.

Modeling and analyzing the information of a complex computational system is very important to judge the overall behavior of the system and also to compare the performance of various computational paradigms. An analytical technique based on queueing networks and Petri nets is introduced by Sundaram and Narahari for making a performance analysis of dataflow computations on multiprocessor systems. A four-parameter characterization, namely minimum parallelism, maximum parallelism, average parallelism and variance in parallelism, for parallelism in dataflow computations has been provided. A detailed investigation of these analytical models has been carried out to conclude that the average parallelism is a good characterization if the variance in parallelism is small, and significant differences in performance measures result otherwise.

A fundamental issue, namely chaotic manifestations due to the absence of proper conditioning, that directly impacts the scalability of theoretical neural networks to applicative embodiments is addressed by Jacob and Gulati. They introduce a mathematical framework for systematically reconditioning additive-type models, and derive a neuro-operator whose dynamics is neither concurrently synchronous nor sequentially asynchronous. Necessary and sufficient conditions guaranteeing concurrent asynchronous convergence are established in terms of contracting operators.

Acknowledgements—We thank the Editor Professor Mo Jamshidi for giving us the opportunity to put together this Special Issue. We wish to thank all the authors for their contributions to this Special Issue. Bulk of the reviewing was performed by us; in particular each paper had been reviewed by at least one member (often 3 members) of the Guest Editorial Board. We thank all our external reviewers, in particular Venu M. Padakanti and Binyun Xie, for their timely and critical reviews. N. S. V. Rao was partially funded by the National Science Foundation under Grant No. IRI-9108610.

Department of Computer Science,
Old Dominion University,
Norfolk, VA 23529-0162
NAGESWARA S. V. RAO

Jet Propulsion Laboratory,
4800 Oak Grove Drive, MS 303-310,
Pasadena, CA 91109-8099
SANDEEP GULATI

Department of Computer Science,
Louisiana State University,
Baton Rouge, LA 70803
S. SITHARAMA IYENGAR

Division of Electronics,
Office of Naval Research,
Arlington, VA 22217-5000
RABINDER N. MADAN

REFERENCES

1. S. Chen, Multisensor fusion and navigation of mobile robots. *Int. J. Intell. System.* **2,** 227–251 (1987).
2. S. S. Iyengar and A. Elfes (Eds), *Autonomous Mobile Robots: Perception, Mapping and Navigation.* IEEE Computer Society Press (1991).
3. B. Wah and C. V. Ramamoorthy, *Computers for Artificial Intelligence Processing.* Wiley, New York (1990).
4. J. Herrath, Guest Editor's introduction: computer architectures for intelligent systems. *IEEE Comput.* **May,** 6–9 (1992).

5. S. S. Iyengar, R. L. Kashyap and R. N. Madan, Distributed sensor networks—introduction to the special section. *IEEE Trans. Syst. Man Cybernet.* **21,** 1027–1031 (1991).
6. M. Brady, Forward to the special issue on sensor data fusion, *Int. J. Robot. Res.* **7,** 2–4 (1988).
7. S. S. Iyengar and R. L. Kashyap, Guest editors' introduction: autonomous intelligent machines. *IEEE Comput.* **Jun,** 14–15 (1989).
8. J. A. Anderson and E. Rosenfeld (Eds), *Neurocomputing: Foundations of Research.* MIT Press, Cambridge, MA (1988).

AUTHORS' BIOGRAPHIES

Nageswara S. V. Rao—Nageswara Rao received B.Tech. from Regional Engineering College, Warangal, India in 1982, M.E. from Indian Institute of Science, Bangalore in 1984, and Ph.D. in Computer Science from Louisiana State University in 1988. He is an Assistant Professor in the Department of Computer Science at Old Dominion University. He was a visiting scientist at Oak Ridge National Laboratory during the summers of 1988–1992. His research interests include distributed sensor systems, robot navigation and sensing and fault diagnosis.

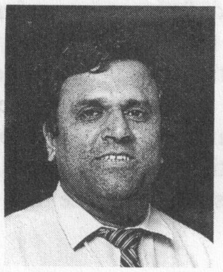

S. Sitharama Iyengar—Sitharama Iyengar received Masters degree in engineering from Indian Institute of Science, Bangalore, India in 1970 and Ph.D. in engineering from Mississippi State University in 1974. He is currently a Professor of Computer Science and has been Director of the Robotics Research Laboratory since its inception in 1986 at Louisiana State University. He edited two books and published more than 150 publications including 75 archival journal papers in areas of high-performance parallel and distributed algorithms and data structures for image processing and pattern recognition, autonomous navigation and distributed sensor networks.

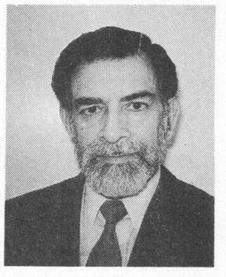

Rabinder N. Madan—Rabinder Madan received B.Sc and M.Sc. in physics from Delhi University in 1954 and 1956, and A. M. and Ph.D. in physics from Princeton University in 1964 and 1967, respectively. From 1967 to 1977 he was on the faculties of the Universities of California, Massachusetts and North Carolina. From 1977 to 1979 he was the manager of Systems Analysis at the Bunker-Ramo Corporation, and from 1979 to 1984 he was a senior scientist at Hughes Aircraft Company. Since 1984 he has been at the Office of Naval Research as program manager of the systems and electromagnetic theory.