Research on Parallel Algorithms and Applications
A Trip Report of My Visit to ICOT

Vipin Kumar
Department of Computer Science,
University of Minnesota
Minneapolis, MN 55455
kumar@cs.umn.edu

October 1990

1 Introduction

This is a report on my activities during my visit to ICOT from September 17, 1990 to October 4, 1990. I would like to thank Dr. Fuchi and Dr. Furukawa for the invitation, and Ichiyoshi-san and other members of the Seventh Research Laboratory for their wonderful hospitality during my stay. The purpose of my visit was to exchange views and discuss future collaboration with members of the Seventh Lab on the following topics: (i) development of new parallel algorithms and methods for mapping them to parallel architectures; (ii) scalability analysis of parallel programs; (iii) applications of parallel processing to specific practical problems.

As the Fith Generation Computing Project moves to its last stage, there is a lot of emphasis on developing applications to run on Multi-Psi. This seems to be the major part of the activities within the Seventh Lab. Since my current research is primarily concerned with the development and analysis of parallel programs for a variety of applications, I found a substantial overlap in research interests with that of the members of the Seventh Lab. During my

stay, I had extensive discussions with various members of the Seventh Lab on topics of common interests, and discussed (and even started) collaborations on various topics. In the following, I summarize some of this.

The Seventh Lab is primarily concerned with the development of parallel applications in KL-1 to be run on Multi-PSI, whereas my research group (now at Minnesota, and formerly at the University of Texas) mostly writes parallel programs C (and its variants) for commercially available parallel computers such as Ncube, Intel Hypercube, Symult, Sequent and Butterfly. Despite this, the underlying parallel algorithms and their scalability analysis are often very similar. This dichotomy of paradigms also allowed us to investigate as to which paradigm is better than the other for different problem domains. Clearly, the major attractive feature of Kl-1 is possible ease in programming the parallel machine, whereas the major attraction of the parallel variants of C is their flexibility in coding different parallel algorithms. One interesting question we are investigating is whether some of the parallel algorithms developed by my research group can be as easily (or more easily) programmed in KL-1. Wada-san had finished a preliminary implementation of our concurrent heap scheme during my stay. That implementation had some serious sequential bottlenecks that are not present in our C implementation of the same data structure. I have heard from Ichiyoshi-san at the time of writing this report that she has finished another more sophisticated version that does not have the sequential bottleneck.

A number of parallel algorithms investigated at ICOT as well as in my research group require balancing and ditribution of work dynamically among parallel processors. We have developed a variety of load balancing schemes. We are now investigating the relative superiority of these schemes. Researchers at ICOT are implementing the load balancing schemes developed by us for their parallel applications, and my students have now started implementing the single-level and multi-level load balancing developed at ICOT. It would be interesting to see as to which of these schemes performs better in various applications.

Since effective use of large scale parallel computers is essential to the fullfilment of ICOT's aim of achieving gigalip perfromance, it is important to figure out whether the parallel application running on a small number of processors (say 64) will also run and achieve good speedup on much larger (say 100,000) number of processors. The scalability analysis technique (called isoefficiency analysis) developed by us is very effective in analyzing

the scalability of parallel algorithm/architecture combinations, and in predicting the performance on large number of processors when experimental results are available on a small number of procerssors. We are now analyzing the scalability of various parallel algorithms developed at ICOT. In particular, Kimura-san and Ichiyoshi-san have completed isoefficiency analysis of the single-level and multi-level load-balancing scehems. They have used a probabilistic queuing model to compute the effects of load imbalance on the overhead of the overall parallel algorithm (that uses the load balancing scheme), anf from this they have computed the isoefficiency functions of these two load-balancing schemes.

Kimura-san and Ichiyoshi-san also plan to undertake the isoefficiency analysis of the parallel algorithm for the shortest path problem. Many parallel algorithms for this problem have been developed by Wada-san and Ichiyoshi-san. A distinguishing feature of these algorithms is that parallel algorithm may perform more computation than is done by the best sequential algorithm. This happnes due to the speculative nature of parallel search. At various points in time, a process(or) may have the choice of either waiting to determine whether a certain computation needs to be done or to go ahead and do it anyway. In the first case, processors will be idle unnecessarily (at least in some cases), and in the second cases, unnecessary (i.e., speculative) work will be done which may lead to more such speculative work. So isoefficiency analysis of such algorithms will involve modeling this speculative overhead.

Ichiyoshi-san discussed distributed hash table as a data structure to be used in parallel algorithms for set manipulation. One possible application of these data structures is in parallel best-first search algorithms used for solving combinatorial problems. In my research group, we have investigated various methods of parallelizing best-first search using distributed priority queues, and we look forward to more insights that will be gained by the development of distributed has tables in developing better parallel algorithms for solving combinatorial problems.

My visit to ICOT overlapped (in the first week) with the US/Japan workshop. I also served on the panel "Technology and Future of Parallel Symbolic Processing" at InfoJapan'90 on October 2. I thoroughly enjoyed both the workshop, and the panel discussion.

In summary, the trip turned out to be intellectually quite stimulating and rewarding for me. I was quite impressed with the range of applications

being pursued in the Seventh Lab, and am looking forward to a continuing collaboration. Here is a log of activities that I engaged in at ICOT.

2 Log of Activities ICOT

Monday, September 17:

Dr. Iwata gave me an introduction to the general organization of ICOT, and introduced me to several people including Mr. Hiroshigue, executive director of ICOT, Dr. Fuchi and Dr. Nitta. Nitta-san gave me an overview of the activities of the Seventh Lab. I was also introduced to the PSI machine, and learned how to connect from this machine to icot21 and icot 34, and from there to the internet. Later in the day, I had a discussion with Ichiyoshi-san and Taki-san on the possibile activities, and developed a tentative schedule of activities and discussions for my visit.

Tuesday, September 18 - Friday, September 21

US/Japan Workshop on Parallel Processing and Logic Programming.

Tuesday, September 25:

In the morning, Ohtake-san, Matsumoto-san and Kimura-san discussed current progress on their parallel formulations (in KL-1) of a number of CAD problems including wire-routing and circuit simulation. I suggested to them to look at the work of Professor Sahni on a parallel maze router on the hypercube (the paper appears in the proceedings of 1987 International Conference of Parallel Processing), and of Seitz and Wen-King Su on circuit simulation on hypercube.

In the afternoon, I saw the demo of a tool called VISTA that is used for performance monitoring of KL-1 programs running on multi-PSI.

In the late afternoon, I had extensive discussions with Taki-san, Ichiyoshi-san and Furuichi-san on the merits of different load balancing schemes.

Wednesday, September 26:

In the morning, Chikayama-san presented details of the PIMOS operating system. In the afternoon, I had a discussion with Wada-san, Taki-san and Ichiyoshi-san on how to implement concurrent heap in KL-1.

Thursday, September 27:

I gave an introductory talk on the general need and utility of scalability analysis to the CPP working group.

Friday, September 28:

In the morning, I had a long discussion with Kimura-san and Ichiyoshi-san on the scalability analysis of the single-level and multi-level load balancing scheme. In the afternoon, I had extensive discussion with Ichiyoshi-san and Taki-san on possible collaboration between our groups.

Monday, October 1:

I gave a detailed and technical talk on isoefficiency analysis to the KL1-Task Group. I went into the details of computing isoefficiency functions for some simple load balancing schemes. I also explained as to how the concept of isoefficiency analysis may help in performance prediction as well as in choosing among various parallel algorithms for a given problem. In this talk, Ichiyoshi-san was kind enough to provide commentary in Japanese at various points. My impression was that this mode of presentation was very well received by the audience.

Tuesday, October 2:

In the morning, I prepared for the panel discussion on "Technology and Future of Parallel Symbolic Processing" to be held at InfoJapan'90 the same afternoon. I participated in the panel along with Dr. Uchida, Dr. Overbeek and Dr. Amamiya. The panel was chaired and moderated by Professor Tanaka.

Wednesday, October 3:

In the morning, I continued discussion with Kimura-san on the isoefficiency analysis of single-level and multi-level load balancing schemes. In the afternoon, I visited Professor Tanaka's research laboratory at the University of Tokyo. I was quite impressed by the hardware of the PIE machine.

Curriculum Vitae

Vipin Kumar

Computer Science Department
University of Minnesota
Minneapolis, MN 55455
Tel: 612-624-8023
Arpanet: kumar@cs.umn.edu
FAX: 612-625-0572

Date of Birth: 21 October 1956

EDUCATION:

Ph.D. (Computer Science)
Univ. of Maryland, College Park, December 1982
Dissertation Title: A Unified Approach to Problem Solving Search Procedures.

M.E. (Electronics Engineering)
Philips International Institute, Eindhoven, Netherlands, May 1979

B.E. (Electronics & Communication Engineering)
University of Roorkee,(U.P.),India, June 1977

PROFESSIONAL AFFILIATIONS:

Member of the ACM, SIGART, SIGARCH, IEEE, IEEE computer society, American Association for Artificial Intelligence, Association for Logic Programming

PROFESSIONAL ACTIVITIES:

Program Committee:
1990 National Conference on Artificial Intelligence (AAAI-90).
1988 National Conference on Artificial Intelligence (AAAI-88).
1987 National Conference on Artificial Intelligence (AAAI-87).
1990 National Conference on Logic Programming
1989 National Conference on Logic Programming
1989 AAAI Symposium on Search and Planning (sponsored by American Association for Artificial Intelligence).

Co-chair of Workshop on Parallel Algorithms for Machine Intelligence and Pattern Recognition, August 1988, St. Paul, Minnesotta. (sponsored by American Association for Artificial Intelligence).

Co-chair of IJCAI-89 Workshop on Parallel Algorithms for Machine Intelligence August 1989, Detroit, Michigan. (sponsored by American Association for Artificial Intelligence).

WORK EXPERIENCE:

8/89-present
Associate Professor
Department of Computer Sciences
University of Minnesota

1/83-7/89
Assistant Professor
Department of Computer Sciences
University of Texas at Austin

8/79-12/82
Graduate Research Assistant
Machine Intelligence and Pattern Analysis Lab.
Dept. of Computer Science, Univ. of Maryland
Worked on the development of a unified model for Problem Solving search algorithms, and their parallel implementations.

11/78-5/79

Philips Research Laboratories
Eindhoven, Netherlands
Designed a SPOOLING system for PHIDIAS - Philips Distributed Asynchronous System being developed at the Research Lab.

10/76-5/77
Department of Electronics Engineering
University of Roorkee
Designed and partially constructed a small educational computer (12 bit central processor and its ASR-33 teletype interface) using SSI and MSI chips.

RESEARCH INTERESTS:

Parallel Processing
Artificial Intelligence
Theory of Algorithms

COURSES TAUGHT:
Introduction to Parallel Computing
Introduction to Artificial Intelligence
Parallel Architectures for Artificial Intelligence
Topics in Parallel Symbolic Computing
Introduction to Expert Systems
Computer Systems Architecture