APPLAUSE: Application & Assessment of Parallel Programming Using Logic

Liang-Liang Li, Mike Reeve, Kees Schuerman & André Véron[1]
Jacques Bellone & Claudine Pradelles[2]
Angelos Kolokouris[3]
Takis Stamatopoulos[4]
Dominic Clark, Chris Rawlings & Jack Shirazi[5]
Giuseppe Sardu[6]

[1] European Computer-Industry Research Centre
Arabellastrasse 17, 8000 Munich 19, Germany
[2] Dassault Aviation, Toulouse, France
[3] Expert Systems International, Athens, Greece
[4] University of Athens, Greece
[5] Imperial Cancer Research Fund, London, UK
[6] Systems & Management, Rome, Italy

Abstract. The APPLAUSE ESPRIT Project is building major applications using the ElipSys parallel constraint logic programming system developed at ECRC. Two major aims of the project are to advance the state of the art in four commercially significant application areas and to promote the use of ElipSys-like languages among applications developers. This brief paper gives an outline of ElipSys and an overview of the applications being developed within the APPLAUSE Project.

1 Project Overview

The APPLAUSE ESPRIT III Project (6708), the Application and Assessment of Parallel Programming Using Logic, is building major applications using the ElipSys parallel constraint logic programming system developed at ECRC (European Computer-Industry Research Centre). It brings together end-users, applications developers and technology providers. APPLAUSE is a three year project and it began in May 1992. The aims of the project are manifold, the major ones being:

- to advance the state of the art in four commercially significant application areas;
- to build a corpus of expertise in using ElipSys-like languages amongst applications developers;
- to generate training material to introduce applications developers to ElipSys-like languages;
- to assess the advantages (and disadvantages) of ElipSys-like languages.

2 ElipSys

ElipSys is a programming system designed to allow the performance potential of large-scale parallel machines to be exploited in the field of search-based applications, such as planning & scheduling and decision support.

ElipSys is based on constraint logic programming. Pure logic programming is recognized as a convenient paradigm for stating and describing search applications and combinatorial problems. It is only relatively recently that it has been complemented with constraint handling, a feature that makes realistic search-based applications tractable. Sequential constraint logic programming systems have been successfully applied to a number of operations research and combinatorial applications. Performances comparable to, and sometimes better than, traditionally coded applications have been exhibited. The paradigm also significantly enhances the flexibility of programming and programmer productivity.

ElipSys enables a search-based application to be parallelized by enabling different processors to be used for exploring the different branches of the search tree, which are normally traversed sequentially a branch at a time.

The ElipSys internal execution model has been designed for maximum portability between different parallel platforms. It uses a hybrid scheme based on shared memory and message passing.

- Shared memory is used exclusively to implement internal static and dynamic data areas. On distributed memory machines, the shared memory concept is supported by virtual shared memory. Thanks to the well defined structure of an ElipSys computation (tree search), the memory reference patterns are such that the usual impediments to the use of virtual shared memory can be minimised. For instance non-strongly coherent memory can be used, or the internal load balancing mechanism of ElipSys can adopt more or less strong policies for improving the locality of reference.
- Message passing is used for synchronization and control purposes between the processors. Depending on the platform, message passing can be implemented using either shared memory, OS primitives or hardware supported mechanisms.

ElipSys yields performance improvements ranging from linear to super-linear. They depend on the type of application and the kind of tree search the applications performs. In all cases the search trees must be bushy enough to let the effect of sharing work overwhelm the overheads of parallelization and load balancing. Linear speed-ups are to be encountered when search trees have to be entirely explored, for example when the optimal solution is sought. Super-linear speed-ups might appear when search trees need only be partially explored, as for instance during branch-and-bound search, where one is satisfied with any solution.

ElipSys is currently available on a wide range of machines. It runs on Sun workstations, in sequential and pseudo-parallel mode. The shared memory machines it has been ported to are the Sequent Symmetry and SPARC-based multi-processors from Sun and ICL. ElipSys has also been ported to the KSR-1 distributed memory machine, on which it makes effective use of the hardware supported shared virtual memory.

3 Applications Using ElipSys

3.1 Manufacturing Planning: Dassault Aviation

Dassault Aviation is using ElipSys to enhance and extend the PLANE manufacturing planning system, currently implemented using the CHIP constraint programming

system also developed at ECRC. PLANE is an aid for the planning of long term (5 to 10 years) production schedules for a collection of assembly lines manufacturing a mix of different aircraft types. The aim of the system is to pace the assembly lines, under a given set of two-dimensional constraints, so as to minimize the combination of the stock costs and the production rate changes. The first dimension consists of precedence constraints between the assembly lines and the second of disjunctive and precedence constraints on the production rate changes. The combinatorics of the problem is the main computational issue and these arise when the system has to decide for which aircraft it is wise to change the production rate and how to change it, e.g. in the most difficult case to choose one aircraft from the 250 that may be involved and to choose one value for the production rate among 90 possible values.

A port and enhancement of PLANE on ElipSys, known as PSAP (Planning System for Aircraft Production), will use the features of ElipSys to explore disjunctive constraints, labelling strategies and search heuristics in parallel.

3.2 Tourist Advice: Expert Systems International & the University of Athens

MaTourA is a Multi-agent Tourist Advisor being built by Expert Systems International (ESI) and the University of Athens as a demonstrator for the Greek National Tourist Organization. The purpose of the system is to support travel agencies in constructing personalized tours, selecting predefined package tours and accessing basic tourist information.

MaTourA consists of a set of specialized autonomous agents. Some of these agents handle information related to activities, events, sites, accommodation, transportation, ticketing, and so on. Other agents are responsible for more complicated tasks, such as the construction of personalized tours that satisfy constraints imposed by the tourist. These tours are actually instantiated Daily Plan Templates (DPTs) which are given by the user. In order for an agent to solve a specific problem, it may require the services offered by other agents. In such a case, it send a requests to the appropriate "expert" agents, which, after the required computation is carried out, send back their responses (results). The agents of MaTourA work in a physically distributed environment and not all will be implemented in ElipSys.

The system provides an interface which helps the user to state his/her problem via a well designed dialog procedure. A training facility will allow even casual users to use MaTourA.

3.3 Molecular Biology: Imperial Cancer Research Fund

The Imperial Cancer Research Fund (ICRF) is addressing a number of challenging problems in protein structure analysis and molecular genetics. These can both be characterized as combinatorially complex and computationally intensive.

The rationale for the use of CLP methods is the observation that many aspects of scientific activity can be viewed as attempts to produce the most general consistent interpretation of a broad range of possibly heterogeneous but interrelated data.

Additionally it should be emphasized that since molecular biology is a broad field of scientific endeavour, encompassing many distinct types of problem-solving

requirements, an important aim of the project is the development of a generic schema for both identifying problems that are amenable to solutions that exploit parallel CLP and providing an implementation methodology.

The potential advantages arising from the use of ElipSys occur both: as a result of the existence of a priori constraints (e.g. rules that govern protein folding and prior knowledge of local gene order) which can be used to prune the search tree; and the remaining, possibly massive, search spaces which can be searched in parallel.

The specific systems initially being developed are a system for the prediction of protein topologies from secondary structure assignments and topological folding rules, and a system for generating a physical genetic map from hybridization fingerprinting data to support the human genome mapping work of Hans Lehrach's laboratory at the ICRF.

3.4 Environmental Monitoring and Control: Systems & Management

Systems & Management is developing a knowledge based decision support system (DSS) concerned with the monitoring and control of pollution in the Venice Lagoon as a demonstrator for the Venice Water Magistracy. It is aimed at assisting the authorities in two major aspects: the correct evaluation of the state of the pollution and the planning of technical interventions aimed at restoring an acceptable state at acceptable costs.

The DSS includes a relational database describing the emissions and a hydrodynamic model of the lagoon able to plot the distribution of polluting substances. These modules are integrated through a knowledge based core, seen as the interpretive model of the analysis and decision making activities, which is being developed in ElipSys. The constraints and parallelism features of ElipSys are seen as complementary means to attack the combinatorial aspects of the evaluation and decision-making processes.

A complete modelization must consider up to 20,000 (registered) pollution sources and tens of dangerous substances. The system should provide an early prediction of potentially dangerous combinations by considering the combinatorial explosion of possible effects. Finally, it should enumerate a suitable mix of technical instruments to be applied to pollution sources to restore acceptable states at affordable costs.

4 Conclusion

By bringing together end-users, applications developers and technology providers, the APPLAUSE Project is building novel sophisticated applications designed to advance the state of the art in four commercially significant application areas. By using the ElipSys parallel constraint logic programming system, developed at ECRC, the project will provide a basis for other applications developers to assess the suitability of ElipSys-like systems for their tasks.

Acknowledgements

The APPLAUSE Project is partially funded by the ESPRIT Programme of the Commission of the European Communities as ESPRIT Project 6708.