

TERRESA A Task-Based Message-Driven Parallel Semantic Network System

Chain-Wu Lee

Faculty of the Graduate School of State

University of New York at Buffalo

Abstract

The ultimate goal of research in artificial intelligence is to build a system with human intelligence. Yet, despite considerable effort, artificial intelligence systems are still a long way from that goal. Various reasons attribute to why the goal cannot be achieved at current stage, one of the reasons is the vast amount of information to be processed. When presented with real world domains of millions of concepts and assertions, almost all existing systems are unable to deliver the results within a reasonable amount of time. This dissertation describes a parallel semantic network system, *TERRESA*, based on SNePS. *TERRESA* is a task-based, message-driven, multithreaded parallel semantic network system with the support of shared knowledge, load balancing and duplicate checking on MIMD distributed memory/message passing systems.

TERRESA can be split into two collaborating components: the host module and the slave module. The host module interacts with the user and processes the information for the slaves, while the slave module performs task execution. The current version of *TERRESA* focuses on SNePS path-based inference only.

TERRESA is designed to be portable among different parallel machines, and therefore most of the systemdependent functions are either avoided or re-implemented. *TERRESA* is written in ANSI C and MPI (Message Passing Interface) library with the help of *flex* lexical analyzer and *yacc* parser generator. Porting *TERRESA* to a new platform only requires recompiling in most cases. The performance studies of *TERRESA* have been done on SUN multiprocessor server.

References

- [1] ALASDIAR, R., BRUCE, A., MILLER, J., AND SMITH, G. *CHIMP/MPI User Guide*. Edinburgh Parallel Computing Centre, Edinburgh, UK, 1994.
- [2] ALVERSON, G., BRIGGS, P., COATNEY, S., KAHAN, S., AND KORRY, R. Tera Hardware-Software Cooperation. In *Proceedings of Supercomputing 1997* (San Jose, CA, Nov. 1997), ACM/IEEE.
- [3] BBN ADVANCED COMPUTERS INC. *TC2000 Technical Production Summary*. Cambridge, MA, Nov. 1989.
- [4] BHATT, S., GREENBERG, D., LEIGHTON, T., AND LIU, P. Efficient embeddings of trees in hypercubes. *SIAM Journal on Computing* (1991).
- [5] BRACHMAN, R. J., AND SCHMOLZE, J. G. An Overview of the KL-ONE Knowledge Representation System. *Cognitive Sci.* 9 (1985), 171–216.
- [6] CERCONI, N. The ECO Family. *Computers Math. Applic.* 23, 2-5 (1992), 95–131.
- [7] CHAN, C. M. Forward Path-Based Inference in SNePS. Tech. Rep. SNeRG Technical Note No. 17, Department of Computer Science at SUNY Buffalo, Feb. 1987.
- [8] CHUNG, M., AND MOLDOVAN, D. Parallel Memory-Based Parsing on SNAP. In 7th International Parallel Processing Symposium (Apr. 1993), pp. 680–684.
- [9] CHUNG, S., AND MOLDOVAN, D. I. Modeling Semantic Networks on the Connection Machine. *Journal of Parallel and Distributed Computing* 17 (1993), 152–163.
- [10] CLARK, K., AND GREGORY, S. Parlog: parallel programming in logic. *ACM Transactions on Programming Languages and Systems* 8, 1 (Jan. 1986), 1–49.

- [11] CRAY RESEARCH INC.. *Cray T3E Architecture Overview*.
<http://www.cray.com/products/systems/crayt3e/>.
- [12] CYBENKO, G. Dynamic Load Balancing for DistributedMemory Multiprocessors. *J. of Parallel Distrib. Computing* 7 (1989), 279–301.
- [13] DATE, C. J. *An Introduction to Database System*. Addison-Wesley Publishing Company, 1986.
- [14] R. F. DeMara, and D. I. Moldovan, “The SNAP-1 Parallel AI Prototype.” *IEEE Trans. on Parallel and Distributed Systems* 4, 8 (Aug. 1993), 841–854.