

The TOP-C Parallel Model and Symbolic Algebra

Gene Cooperman¹

College of Computer and Information Science
Northeastern University
Boston, MA 02115
`gene@ccs.neu.edu`

Software name: TOP-C (and ParGAP, ParGCL, ParGeant4, Marshalgen)

Short description: Parallelization features of TOP-C:

- minimal modifications to sequential code (using ‘raw mode’)
- small footprint (30 KB stripped library under Linux)
- self-contained (including MPI subset)
- extensive manual (approximately 50 pages)
- compiles into distributed or shared memory architecture
- optionally compiles into sequential emulation for easy debugging

Public access: <http://www.ccs.neu.edu/home/gene/XXX> for XXX one of `topc.html`, `pargap.html`, `pargcl.html`, or `pargeant4.html`

Abstract

The TOP-C parallel model was specifically designed to be optimal for many symbolic algebra algorithms. Symbolic algebra problems are irregular and differ from typical problems of numerical analysis in two respects.

1. They have large amounts of intermediate swell.
2. They are usually not able to make good use of data parallelism.

The TOP-C model was developed on the basis of experience over ten years. The model was first implemented for distributed memory for GCL (GNU Common LISP), and is currently available as ParGCL [1]. That work was extended to ParGAP [3], a binding to the interactive language GAP (Groups, Algorithms and Programming). Finally, a C/C++ version, Task Oriented Parallel C/C++, lent its acronym to the TOP-C model [2].

A sign of the maturity of TOP-C is the parallelization of Geant4 (Geometry and Tracking) [4]. Geant4 is a million-line program for the simulation

¹This work was partially supported by the National Science Foundation under Grant CCR-0204113 and earlier grants.

of particle-matter interaction used in the design of collider experiments at CERN. Current issues facing parallelization include the slowness of RAM [5] and the need for semi-automated marshalling (serialization of objects) [6].

-
- A. *The Task*: tasks executed by each slave process, in parallel
 - B. *The Action*: control of parallel strategy for TOP-C
 - C. *The Shared Data*: shared data, read-only by all processes




Figure 1: TOP-C model (distributed memory): For Gröbner bases, the *shared data* is the current basis of the polynomial ideal, the *task input* is an S-pair, and the *task output* is a reduced polynomial. For a non-trivial reduction, the master declares an UPDATE action, and each node reduces its replicate of the current basis.

References

- [1] G. Cooperman. STAR/MPI: Binding a parallel library to interactive symbolic algebra systems. In *Proc. of International Symposium on Symbolic and Algebraic Computation (ISSAC '95)*, pages 126–132. ACM Press, 1995.
- [2] G. Cooperman. TOP-C: A Task-Oriented Parallel C interface. In *5th International Symposium on High Performance Distributed Computing (HPDC-5)*, pages 141–150. IEEE Press, 1996.
- [3] G. Cooperman. Parallel GAP: Mature interactive parallel computing. In *Groups and Computation III*, pages 123–138. DeGruyter Publishers, 2001.
- [4] G. Cooperman, H. Casanova, J. Hayes, and T. Witzel. Using TOP-C and AMPIC to port large parallel applications to the computational grid. *Future Generation Computer Systems (FGCS)*, 19:587–596, 2003.
- [5] G. Cooperman and V. Grinberg. Scalable parallel coset enumeration: Bulk definition and the memory wall. *J. Symbolic Computation*, 33:563–585, 2002.
- [6] G. Cooperman, N. Ke, and H. Wu. Marshalgen: A package for semi-automatic marshalling of objects. In *Proc. of The 2003 International Conference on Internet Computing (IC'03)*. CSREA Press, 2003. to appear.