Scalable Solvers and Applications

C.J. Ribbens

October 27, 2000

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

UCRL-CR-141281

Calvin J. Ribbens
Department of Computer Science
Virginia Polytechnic Institute & State University

Final Report for “Scalable Solvers and Applications Final Report”

Prepared for:

University of California
Lawrence Livermore National Laboratory
Attn: Jill Dunaway
P.O. Box 808, L-561
Livermore, CA 94551

For the period ending September 30, 2000

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes. Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract

UCRL-CR-141281

W-7405-ENG-48.

Scalable Solvers and Applications Final Report

Calvin J. Ribbens
Department of Computer Science
Virginia Polytechnic Institute & State University

The purpose of this report is to summarize research activities carried out under Lawrence Livermore National Laboratory (LLNL) research subcontract B501073. This contract supported the principal investigator (P1), Dr. Calvin Ribbens, during his sabbatical visit to LLNL from August 1999 through June 2000. Results and conclusions from the work are summarized below in two major sections. The first section covers contributions to the Scalable Linear Solvers and *hypre* projects in the Center for Applied Scientific Computing (CASC). The second section describes results from collaboration with Patrice Turchi of LLNL's Chemistry and Materials Science Directorate (CMS). A list of publications supported by this subcontract appears at the end of the report.

1 Scalable Linear Solvers

The purpose of the scalable linear solvers and *hypre* projects is to develop scalable algorithms and software for the solution of large, sparse linear systems of equations using massively parallel computers. In order to meet the needs of a variety of current and future users, an important goal of the *hypre* package and framework is interoperability and extensibility. Toward that end, the P1 worked with Andrew Cleary and Jeff Painter to formalize and extend an object model for *hypre*. From the beginning, *hypre* has been designed in an object-oriented style. The existing library includes several solvers and two user interfaces for defining problem components, all of which reflect a well-defined object model. Hence, the starting point in discussing an object model for *hypre* was the existing one. The goal was to extend this model, making it more precise and consistent, while attempting to reflect the ongoing work of the Equation Solvers Interface (ESI) Forum. The ESI Forum is a lead by representatives of several Department of Energy laboratories. Their goal is to develop a set of standards for scalable equation-solver services and components. The PT attended a meeting of the ESI Forum in September 1999 in order to gain insight into the relationship between their efforts and the *hypre* project.

Previously, the *hypre* model was defined only by informal description, convention, and example. Since the primary implementation language for *hypre* is C, this approach was necessary. The PT worked to formalize and extend the *hypre* object model using the tools and notation defined by the Babel language interoperability project. He also contributed to the initial implementation of several of the classes in this model. A significant advantage of using Babel in the context of *hypre* is that Babel's Scientific Interface Definition Language (SIDL) provides an unambiguous, compilable language for defining an object model, without dictating the implementation language. It proved extremely helpful to have a precise language in which to describe the object model, letting Babel provide (and enforce) the mechanism for implementing the object model in a particular language. As an important side-benefit, Babel provides convenient language-interoperability, e.g., application codes written in Fortran, C, or C++ can use *hypre* easily, and solver modules written in any of those languages can now be incorporated in the library in a much more straightforward fashion.

As a natural outgrowth of the *Babel/hypre* work, the PT also participated in the early activities of a group seeking to design a multi-level interface (MLI) for using and building multi-level solvers. This effort, lead by Charles Tong, is also using Babel and is leveraging much of the work that went into the *hypre* object model design and initial implementation.

Future Plans

The *Babel/hypre* work is ongoing. While the object model is relatively stable, there is still implementation work to be done, and there are other important issues to resolve. Two of these issues are merging “Babelized” *hypre* with the existing package and responding to the ongoing evolution of the EST. The PT will continue to collaborate on these issues. A draft of a report [1] describing the *Babel/hypre* work has been written. We plan to add at least one more good example of the power of the object model in doing algorithm development, at which point the report will be suitable for submission to a journal. The *Babel/hypre* work is also being reported in a paper [3] to be presented at the Tenth SIAM Conference on Parallel Processing for Scientific Computing. In this connection, the PI is also organizing a minisymposium at that conference on “Component Architectures for High Performance Scientific Computing.”

The PT intends to keep in touch with the MLI effort as well. The possibilities for new algorithm development here are very good, including work in the area of domain decomposition, a topic of long-time interest to the PT.

2 Computational Materials Science

Patrice Turchi is a senior research scientist in the Chemistry and Materials Science (CMS) Directorate at LLNL. He and his co-workers are developing a new tight-binding-based electronic structure scheme for the prediction of the thermodynamic and physical properties of complex multi-component alloys. The scheme combines molecular dynamics (MD) and monte carlo (MC) simulations so that issues related to different time-scales can be addressed. The energetics which describes the diffusion part of the scheme (MC) is obtained within a real-space electronic structure (ES) method by solving first for the inhomogeneous chemically random alloy in the framework of the coherent potential approximation with a recursion technique, and second for the ordering part of the internal energy with the embedded cluster method and an orbital-peeling technique. Once a new alloy configuration is obtained from MC simulations, the alloy topology is relaxed with MD in the next cycle. The MD-ES-MC loop is iterated until self-consistency is reached and chemical short-range order, energetics and relaxed topology can be predicted.

Stand-alone sequential codes exist for each of three main components of the MD-ES-MC loops. The MD and ES codes are computationally very demanding, requiring many hours of CPU time for a problem of modest size (0(1000) atoms). During his visit to LLNL, the PT became familiar with these codes, concentrating his efforts on the most expensive step (ES), as implemented in the (sequential) *tbepi* code. This code was first parallelized using OpenMP for the Compaq cluster shared-memory machines at LLNL. Then a distributed-memory MPI version of the code was developed, targeting the IBM ASCI Blue Pacific machine. To take advantage of this architecture a hybrid implementation was developed, with OpenMP-based threads at inner loops and MPJ based problem decomposition at outer loops. A special static load balancing scheme was developed in order to achieve good performance in the distributed-memory version of *tbepi*. Efficient and scalable performance was demonstrated on up to 2000 atoms on 250 nodes (1000 processors). In addition to the work on *tbepi*, the PT implemented the first MPI-based parallel version of Turchi’s MD code *tbmd*.

Future Plans

A paper [6] is in preparation describing the parallel *tbepi* code and reporting on some early scientific results derived using this code. Ribbens and Turchi are now looking at the details of combining the three codes into the global scheme. Besides the many applications questions that are being investigated, interesting computer science issues

UCRL-CR-141281

regarding code interoperability, data management, load balancing, and computational steering are arising. Collaboration will continue on this project. It is likely that a NSF or DOE proposal will be written to further this work. In addition, collaboration among the PT, Turchi, and Dr. Diana Farkas of Virginia Tech's Materials Science and Engineering Department will continue. This collaboration made good progress during a visit of Dr. Farkas to LLNL during April 2000, when the details of one publication [2~ were finalized and plans for new experiments and model and code integrations were made.

References

- [1] A. Cleary, S. Kohn, C. Kumfert, J. Painter, and C. Ribbens. Babel as a tool for language interoperability and object-oriented design in *hypre*. In preparation.
- [2] D. Farkas, M. Duranduru, W. A. Curtin, and C. J. Ribbens. Multiple dislocation emission from the crack tip in ductile fracture of Al. *Philosophical Magazine A*, 2000. Accepted.
- [3] S. Kohn, Gary Kumfert, Jeff Painter, and Cal Ribbens. Divorcing language dependencies from a scientific software library. To be presented at the Tenth SIAM Conference on Parallel Processing for Scientific Computing.
- [4] C. Mateescu, C. J. Ribbens, and L. T. Watson. A domain decomposition algorithm for hermite collocation problems. *Num. Meth. PDEs*, 2000. Submitted.
- [5] N. Ramakrishnan and C. J. Ribbens. Mining and visualizing recommendation spaces for elliptic PDEs with continuous attributes. *ACM Trans. Math. Softw.*, 2000. To appear.
- [6] P. Turchi and C. Ribbens. $O(N)$ tight-binding CPA approach to chemical order in alloys on massively parallel computers. In preparation.