

Parallel Simulation Engines for Whole-Cell Models

Markus Schwehm, Nils Gehlenborg, Hendrik Post, Amelie Stein, Kirstin Weber

University of Tübingen

ZBIT - Zentrum für Bioinformatik Tübingen

Sand 1, 72076 Tübingen, Germany

schwehm@informatik.uni-tuebingen.de

ABSTRACT

The enormous complexity of whole-cell models needs new algorithmic approaches and high performant simulators. Our goal is the integration of deterministic and stochastic simulation algorithms in a highly scalable parallel simulator. Our hardware platform is a cluster of workstations as they are available in most research institutions. Therefore our parallel simulation engines are based on the *Message Passing Interface* (MPI) [1] standard, which has been implemented for many hard- and software environments. MPI is available as a library and represents a reliable platform for our work. We have implemented two parallel simulation engines, the *Parallel Deterministic Simulation Engine* (PDSE) and the *Parallel Stochastic Simulation Engine* (PSSE). Both use MPI for interprocess communication and share the same interface. Sharing a single library for parallel processing (MPI) reduces the amount of external software our simulator depends on and the system becomes easier to install and maintain.

The PDSE is written in C++ and based on the *Portable Extensible Toolkit for Scientific Computing* (PETSc) [2], which is a software library providing routines and data structures for the numerical solution of partial differential equations and related problems on parallel (and serial) computers. PETSc uses MPI (as described above) for interprocess communication. The PDSE performs simulations on a distributed (regular) grid, that can be either 2- or 3-dimensional. The size of the grid may be customized according to the reaction volume to be simulated. Using the forward or the backward Euler method, PETSc solves the ordinary differential equations arising from the discretization of time-dependent partial differential equations that describe the biochemical model under investigation. Concerning the solution of ordinary differential equations we can rely on the stable numerical methods provided by the PETSc library. However, we are also able to use our own routines, since PETSc offers an interface to external ODE solvers as well.

The PSSE is also written in C++ and uses an efficient implementation of the Next Reaction Method introduced by Gibson [3]. The reaction volume is subdivided into a cubic grid and each sub-

volume is mapped onto a MPI process instance. Diffusion between the sub-volumes is implemented by distributing 10 percent of the molecules of each sub-volume to its 26 neighbouring sub-volumes. The inter-process messaging is handled by native MPI/C++ routines. The internal task management is of a master-slaves process where the master process distributes substances equally to the sub-volumes and collects the results after completion of the slave tasks.

Besides the simulation engines implemented in C++ the user interface is implemented in pure Java. To bring both parts together in an efficient way, we have used the *Java Native Interface* [3] (JNI) which offers a flexible solution and good performance. JNI enables us to link native libraries dynamically to the Java project and thus call native routines wrapped by Java classes. Following this approach, the PDSE and the PSSE are developed as libraries, not as programs. That makes the simulator virtually independent of the simulation engines, allowing us to exchange them with small effort. For example we also have implemented a sequential stochastic simulation engine in Java. Another benefit of the JNI is, that a Java configuration module enables us to link both simulation engines together. Thus individual sub-volumes of the reaction volume can be simulated either by the PDSE or by the PSSE. Using this approach, it is possible to assign cell compartments to sets of sub-volumes which can then be executed on different (deterministic or stochastic) simulation engines.

REFERENCES

- [1] W. Gropp, E. Lusk and A. Skjellum. *Using MPI*. 2nd Edition, MIT Press, Cambridge, MA (1999)
- [2] S. Balay, W. D. Gropp, L. C. McInnes and B. F. Smith. *PETSc Users Manual*. Technical Report ANL-95/11 - Revision 2.1.0, Argonne National Laboratory (2001)
- [3] M. A. Gibson, *Computational Methods for Stochastic Biological Systems*. PhD. Thesis, Calif. Inst. Tech. (2000)
- [4] S. Liang. *The Java Native Interface*. Addison-Wesley, Reading, MA (1999)