Comparison of Parallel Simulation Techniques

mosis

Parsytec Supercluster / Parsytec Explorer

During a stay at the University of Glasgow there was the possibility to implement the simulation system mosis on two parallel computer systems installed there and to try the parallel comparison.

The parallel simulation system mosis was designed at the Technical University of Vienna. The results of two parallel implementations (Cogent XTM, RS6000 cluster under PVM) for this particular comparison have already been presented in SNE 11 and SNE 12.

mosis is a continuous simulation system especially designed for parallel computer architectures, but it can be used with single processor computers as well; the underlying parallelization concept is the *model interconnection concept* which allows a complex simulation model to be described as a connection of several smaller individual models that are executed autonomously as processes / threads and communication with each other on the network. A main advantage is that all issues concerning parallelism are hidden to the user: the simulation code (and even the experimentation commands) are exactly the same when the model is simulated on a multiprocessor network as when it is executed on a simple PC. The user also does not have to insert explicit communication commands into the simulation model; this is done automatically by the mosis translator.

mosis offers a powerful CSSL-type modelling language based on "C" and a "C"-like experimentation interpreter with programming features. Before models can be simulated, instances (corresponding to processes in the network) containing the actual data and connection information of those have to be created and eventually connected to other instances (for parallel models).

mosis is distributed as *freeware* and is available for a number of different computer systems (including PC/Windows).

At the University of Glasgow there are three parallel computer systems by the Parsytec, specialists for parallel computing; two of them were used for the implementation of mosis:

The "older" one is the Parsytec Supercluster consisting of 64 transputers T800, equipped with 4 MB RAM each.

The newer one (installed 1995) is a Parsytec Xplorer consisting of 4 PowerPC processors, 8 MB RAM each.

Both systems are connected to a SUN workstation for user interaction and run under the PARIX 1.x operating system (but different sub-versions that are unfortunately quite different for programming).

mosis was designed in a very modular way, so it should have been very easy to adapt it to any new computer architecture: three "C"-modules contain all operating system dependent code, one for graphics, one for user interaction (the same for all "raw" UNIX implementations) and one for process handling and message passing. Nevertheless there were some problems for the mosis implementations on the two machines: the first was that the PVM implementation on the Xplorer did not work correctly (the PVM implementation of mosis would have been easy, as it already existed), so the Parix system routines had to be used. The second problem was that the Parix implementations were quite different on the two systems so that many parts of the code had to be written twice.

Results of the Parallel comparison - General Issues

Although both computer systems were made by the same manufacturer, they showed an absolutely different behaviour for these comparisons: the Supercluster implementation works quite slow, but the relative speed-up factors are very good, while the Xplorer implementation works very fast in absolute terms, but the relative speed-up factors are much worse than with other systems (the relation between processor speed and communication speed is quite different, not only on the hardware level). For better comparability 4 processors were used in each system, in the Supercluster also 8 processors were used. All simulations were done with the demanded integration algorithms (all RK4) and system parameters.

Monte Carlo Study

With this problem, all systems had a relatively linear speed-up: the Xplorer version had $f = 3.1$; the Supercluster results were $f = 3.92$ (4 processors) and $f = 7.86$ (8 processors). The reason for the relative low speed-up factor for the Xplorer could be the relatively high communication overhead of mosis (polling for messages like simulation stop, view variables etc.)

Coupled Predator-Prey System

This part of the problem showed the differences between both systems in a dramatic way. The simulations were carried out with 5 processors (4 processors on the Xplorer) with communication either after each simulation step ($cint = dt$) or each 10th simulation step ($cint = 10.dt$):

	Xplorer 4 x PowerPC	Supercluster **5 x T800**
cint = dt	f = 0.0032	*f* = 0.24
cint = 10.*dt*	f = 0.03	*f* = 0.72

In this table one can see that the parallelization of this problem on the Xplorer is extremely ineffective: The parallel version takes around 300 times longer than the serial one! On the contrary the parallelization on the Supercluster is not so bad: When communication is done only each 10th step, the parallel version is not much slower than the serial one (note: this problem is intended to yield negative results, i.e. speed-up factors less than one). The problem is of too fine grain for systems with high computation speed.

Partial Differential Equation

The results for this example confirm the numbers described above. Simulation was done on 4 (Xplorer) or 8 (Supercluster) processors with eight tasks (800 lines split into 8 processes = 100 lines or 200 differential equations, resp. for each process). Experimentations were done with communication every step and each 4th step.

	Xplorer 4 processors	**Supercluster 8 processors**
cint = dt	f = 0.6	*f* = 3.6
cint = 4. *dt*	f = 1.5	f = 7.2

While the Xplorer does not gain a reasonable speed-up factor - even with sparse communication, the Supercluster gets a quite good factor even with regular communication which is even improved with less communication (nearly linear speed-up).

Conclusion

Although both computer systems are produced by the same manufacturer, they have a totally different behaviour in terms of parallelization: The Xplorer consists of very fast processors with a (relatively) very slow (with respect to the processor's speed) communication system; on the other side the Supercluster consists of many slow processors, but has a - relative to that - efficient communication system.

G. Schuster, W. Zeller, ARGESIM, Dept. of Simulation Techniques, Technical University Vienna, Wiedner Hauptstr. 8-10, A-1040 Wien, Email: argesim@argesim.tuwien.ac.at