Distributed Object Library (DOLIB) Emulates Shared Memory on Intel Paragon

Problem: Tracking particles is a key computational kernel in Lagrangian transport algorithms, molecular dynamics simulation and particle in cell methods. A common technique involves a spatial decomposition of the computational grid with padding (or extended "ghost cells") to contain flow field information from neighboring regions. The time-step is then chosen to ensure no particle can escape the extended region. This technique either imposes a very severe constraint on the time step allowed or requires high overhead in memory use and communication volume associated with a large ghost region. Load imbalance is usually a problem since the cost of tracking a particle may depend on the local characteristics of the flow field.

Approach: We take a more natural approach and consider the flow field stored in a globally accessible array. Processors "directly" access flow field information. As a result, particles need not be statically assigned to processors and dynamic redistribution of particles for load balancing is possible.

`DOLIB` provides a Fortran and C callable library for runtime dynamic creation and destruction of global arrays. Access to shared memory is through explicit calls to `gather` and `scatter` and requires no new language extensions, no compiler nor special operating systems support. *Read-only* data is cached within `DOLIB` for good performance. `DOLIB` provides atomic accumulate operations (without explicit locks) and generalized test-and-set operation for implementing dynamic load balancing. `DOLIB` uses the IPX message system developed at BNL and is available for PVM and Intel multiprocessors.

Progress and Results:

- Shared memory provided by `DOLIB` is exploited by `DONIO` (Distributed Object Network I/O Library) to speed up non-sequential parallel direct access I/O by 15-fold on the Intel Paragon.
- `DOLIB` has been used to parallelize the CHAMMP Semi-Lagrangian Transport (SLT) code for modeling advection transport of moisture. The shared memory code retains most of the serial code and was competitive with another version that uses explicit message passing.
- `DOLIB` has been used to produce a molecular dynamics code based on SOTON_PAR. The new code is more memory efficient and has runtime dynamic load rebalancing.

Future Work: We are exploring handling paging to disk to create hundred gigabytes global arrays.

Sponsor: `DOLIB` was supported by the PICS (Partnership in Computational Sciences) Computational Techniques and Software Tools projects.

For additional information contact:

Ed F. D'Azevedo or Charles H. Romine
Mathematical Sciences Section
Oak Ridge National Laboratory
P.O. Box 2008, Bldg. 6012
Oak Ridge, TN 37831-6367
E-mail: Ed: `efdazedo@msr.epm.ornl.gov`, Charles: `rominech@ornl.gov`
Phone: Ed: (615) 576-7925, Charles: (615) 574-3141
FAX: (615) 574-0680