SymGrid: a Framework for Symbolic Computations on the Grid

Kevin Hammond

As part of a major EU infrastructure project (SCIEnce: Symbolic Computing Infrastructure in Europe), we are developing SymGrid, a new framework that will, for the first time, allow multiple invocations of symbolic computing applications to interact via the Grid. SymGrid is designed to support the specific needs of symbolic computation, including computational steering (greater interactivity), complex data structures, and domain-specific computational patterns (for irregular parallelism). A key issue for the SymGrid design is heterogeneity: SymGrid is designed to orchestrate legacy components from different symbolic systems into a single coherent (possibly parallel) Grid application, building on the OpenMath standard for data exchange between mathematically-oriented applications.

SymGrid comprises two main parts: SymGrid-Services and SymGrid-Par. SymGrid-Service provides a set of WSRF-compliant interfaces from symbolic computations to both Grid and Web services, and allows straightforward encapsulation of symbolic computations as Grid service components, including automatic client generation. Complementing this, SymGrid-Par provides high-level orchestratation of symbolic components into (parallel) Grid-enabled applications. Each component executes within an instance of a Grid-enabled engine, which can be geographically distributed to form a wide-area computational Grid, built as a loosely-coupled collection of Grid-enabled clusters.

We have constructed implementations of both SymGrid-Services and SymGrid-Par. While we do not yet have performance results for a wide-area computational Grid, our initial results show that good performance can be achieved both on modern multicore systems and on clusters of workstations.

We demonstrate that for a number of simple, but representative testbed applications, good parallel speedup is possible over a sequential, optimised computer algebra system running on a single processor/core. In particular, we are able to achieve relative speedups of between 12.5 and 31.6 on 28 processors, and up to 8.3 on the eight-core machine, with slightly better absolute speedups for at least one legacy system.