

Building Dependable Distributed Systems Using the AQuA Architecture

William H. Sanders
whs@crhc.uiuc.edu

Center for Reliable and High-Performance Computing
Coordinated Science Laboratory and Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

This talk provides an overview of the AQuA architecture, which is an infrastructure for building dependable distributed systems. In particular, the AQuA architecture provides dependable distributed objects through replication, where the type and degree of replication is managed dynamically, depending on the needs of applications. To do this, the AQuA architecture uses the Proteus dependability manager to configure the system in response to faults and availability requests. More specifically, Proteus uses gateways to provide fault-tolerant mechanisms and group communication services to standard CORBA applications transparently as directed by an application and the Proteus dependability manager. The AQuA architecture thus raises the level of abstraction at which a programmer thinks about fault tolerance much higher, relative to existing dependability middleware, allowing an application programmer high-level control over the type of faults that should be tolerated and the level of availability desired from a distributed object. In addition, AQuA dynamically adapts the configuration of a system at runtime, in response to faults that occur, to try to maintain a desired level of availability. Finally, it recognizes that applications will require different levels of availability during different phases of their executions, and supports system reconfiguration in response to changing application requirements.