

## Task and Core-based Automated Fault Tolerance in High-Performance Computing Systems

High-performance computing systems require manual intervention if one or more computing cores fail.

This places a cost on the maintenance of computing tasks.

Intelligent approaches which can pro-actively detect computing core failures and take action to relocate the computing core's task onto reliable cores can make a significant step towards automating fault tolerance in high-performance computing systems.

This paper describes an experimental investigation into multi-agent approaches to bring in this intelligence to high-performance computing systems.

Three approaches are studied to incorporate agent intelligence with high-performance computing systems; the first at the task level, the second at the core level and the third both at the task and core level.

The approaches are investigated for single core failure scenarios that can occur in the execution of parallel reduction algorithms on computer clusters.

The key result is that a task can be relocated without manual intervention and with a time delay in the order of milliseconds.