

Adaptive MPI

Productive Parallel Programming

AMPI Features

About Adaptive MPI

Adaptive MPI (AMPI) is an implementation of MPI on top of Charm++'s adaptive runtime system. It provides automated support for virtualization, load balancing and fault tolerance—all with the familiar API of MPI. Users can recompile their existing applications, and start profiting from overdecomposition, load balancing, and more.

Where AMPI Helps

AMPI provides dynamic runtime support for pre-existing MPI applications. It can help alleviate scalability bottlenecks that are difficult to address directly in applications, such as static and dynamic load balances and difficult-to-hide communication latencies. AMPI also makes it possible to run through node failures and can run on dynamically shrinking and expanding job allocations.



Processor Virtualization

AMPI implements MPI ranks as lightweight, migratable user-level threads, rather than operating system processes.



Dynamic Load Balancing

The runtime system monitors load and supports migrating ranks at runtime across the cores and nodes of a system.



Latency Tolerance

Ranks are co-scheduled in a message-driven manner, automatically overlapping communication of one rank with computation of others.



Online Fault Tolerance

Users can checkpoint their application's state with a single call. The runtime system then detects and recovers from faults within the same job.



Job Shrink/Expand

Since AMPI ranks are virtualized and migratable, the number of physical processors that a job runs on can vary over time.



Communication Locality

AMPI optimizes communication between ranks in the same address space, providing low latency and high bandwidth for neighboring ranks.