

# Charm++

## Productive Parallel Programming

## Charm++ Features

### About Charm++

Charm++ is a mature, highly scalable parallel programming system, and is an alternative to current shared and distributed methods. Though written in a C++ skeleton, it is compatible with Fortran, C, and C++. Code written in MPI can call Charm++, and Charm++ can call MPI, OpenMP, CUDA, and more.

### Where Charm++ Helps

Charm++ helps creators of compute intensive software deliver better shared and distributed memory parallel application performance, improves overall throughput and reduces development and maintenance effort by enabling object oriented methods, dynamic load balancing, and fault tolerance with high level C++ style programming.



#### Processor Virtualization

The problem is broken down into logical units, which are automatically mapped to processors.



#### Load Balancing

Rather than make every program solve load balancing issues on their own, Charm++ provides automatic load balancing for all applications.



#### Automatic Communication and Computation Overlap

Charm++ exploits logical decomposition to enable dynamic overlap of communication and computation as the application executes.



#### Checkpointing & Resilience

Applications written in Charm++ can automatically checkpoint and restart with no extra code or special OS support. They can also run through node failures!



#### Asynchronous Execution

The dynamic scheduler in the Charm++ runtime system executes work as its data becomes available. This minimizes time spent idle, waiting for communication and synchronization. It also enables concurrent composition of multiple modules.