

Algorithms For Scaling Particle Simulations Across Ultra-High Core Count Machines

Kyle Thomas

The Problem

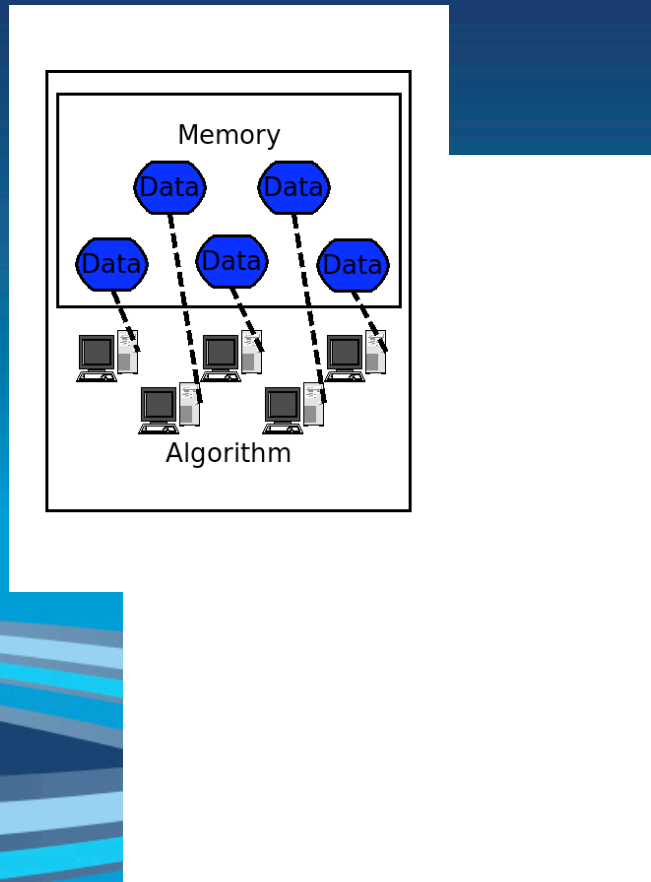
- Memory management in parallel computing is falling behind the ability to distribute large problems over many cores
- Particle problems can have very large domains that can be logically divided among memory in different cores

Proposed Solution

- Separation of the memory management from the computational aspects of the simulation
- Computational engine makes calls to the “Aura” for data

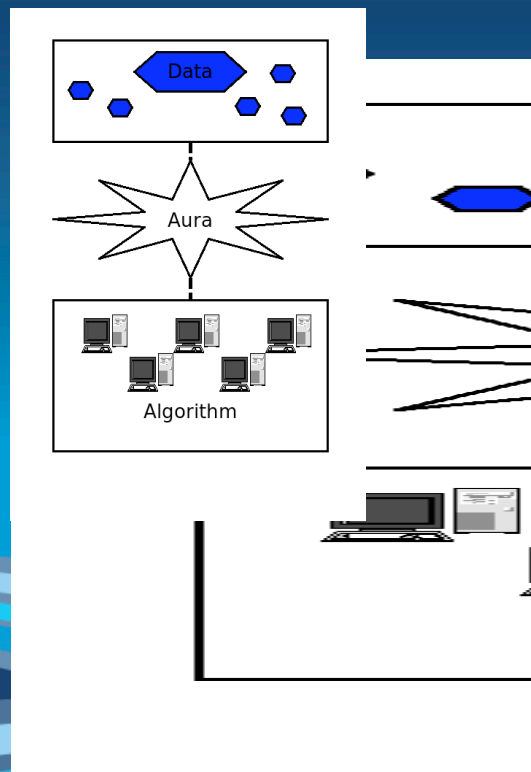
Visualizing the “Aura”

Currently, each node requires a copy of the domain



Visualizing the “Aura”


All nodes would request the necessary data from the “Aura” which distributes pieces of the domain among the nodes




BlobFlow

- Open source simulation for incompressible flows
- Testbed code for “Aura” memory management
- Well tested and understood


BlobFlow



BlobFlow Current Scaling



BlobFlow Current Scaling



Current Work

- Implementation of the “Aura” in BlobFlow using MPI-IO
- Finding bottlenecks in the code and their reasons
- Ensuring similar performance and results

MPI-IO

- The “Aura” will be based on MPI-IO
- MPI-IO has the ability to store large memory files across multiple nodes

Performance Metrics

- Reduction in the need for memory
- Scaling improvement
- Speed improvement

Future Work

- Completion of the conversion of BlobFlow to use a simple version of the “Aura” with MPI-IO
- Generalization of “Aura” method to allow for computations to be separated from memory management