

Work-in-Progress Abstract

Compiler-Assisted Scientific Workflow Optimization

Hadia Ahmed¹, Peter Pirkelbauer²,
Purushotham Bangalore², Anthony Skjellum³



¹ Lawrence Berkeley National Laboratory

² University of Alabama at Birmingham

³ University of Tennessee at Chattanooga

Exascale Systems


Data analytics will face tremendous challenges on Exascale systems

- Many compute nodes communicate with analytics nodes
- Simulations produce vast amount of data
- In-situ (in-transit) analytics necessary to deal with limited bandwidth
- Simulation / analytics code need to be re-organized

Describe Re-organization

Users specify re-organization with an annotation language

Tool generates optimized version




- Move code from analytics node to simulation (or vice versa)
- Describe reductions
- ...

Approach


Compiler-based

Use ROSE to read, analyze, and re-organize source files



Restructure Bonds-CSym

On a single system, we achieved speedups between 4% and 12%.



- Restructured Bonds-CSym in a 1:1 configuration
- Re-organized code
 - Eliminates storage to file system
 - Eliminates data container conversion
 - Enables further compile-time optimizations
- Bonds-CSym is quadratic, smaller input sizes exhibit larger speedups
- Reduced need for network communication

Thank you

contact: Peter Pirkelbauer (UAB)
e-mail: pirkelbauer@uab.edu