



Tobias Hilbrich
Technische Universität Dresden, Germany

Bronis R. de Supinski
Lawrence Livermore National Laboratory, USA

Andreas Knüpfer
Technisch Universität Dresden, Germany

Robert Dietrich
Technisch Universität Dresden, Germany

Christian Terboven
RWTH Aachen, Germany

Felix Münchhalphen
RWTH Aachen, Germany

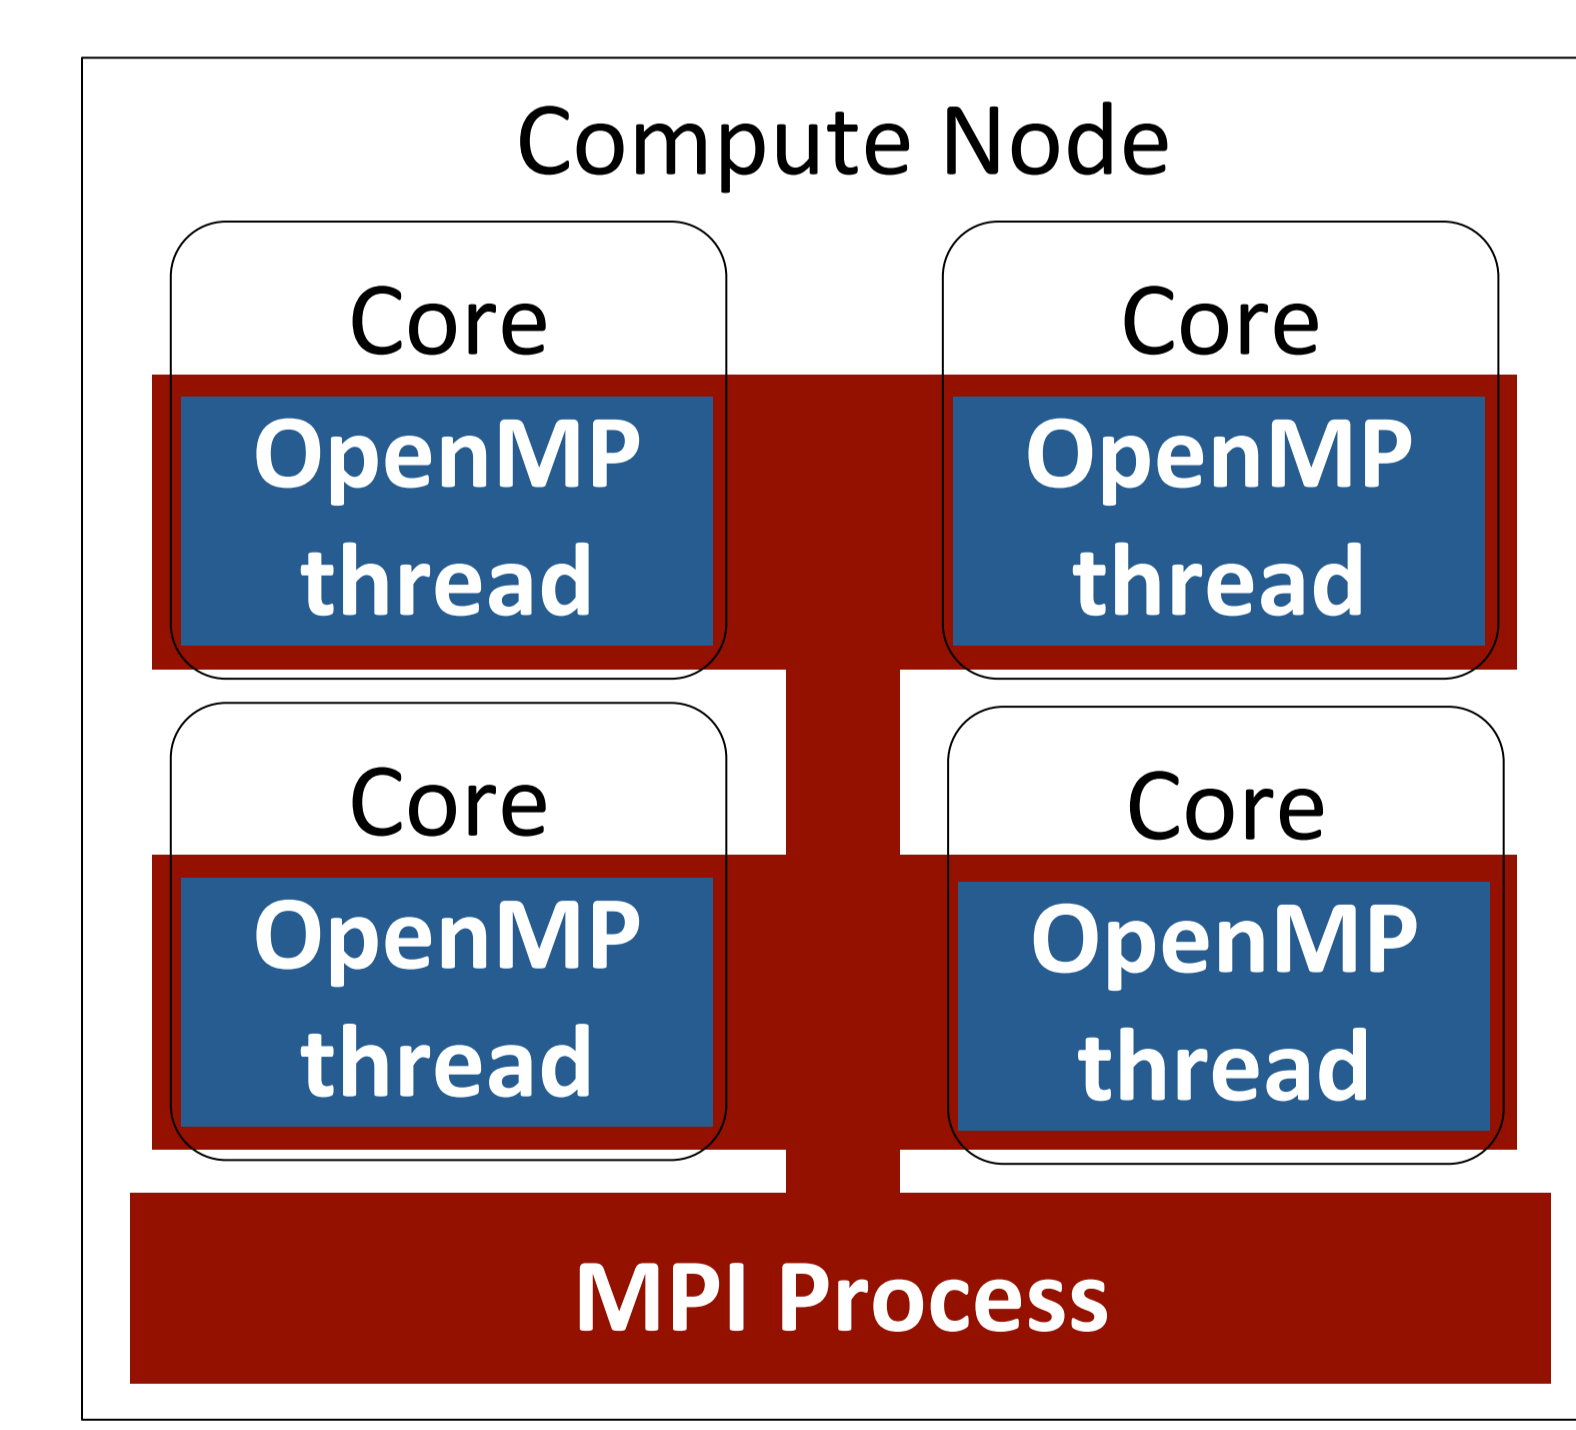
Wolfgang E. Nagel
Technische Universität Dresden, Germany

<http://bit.ly/sc15poster>

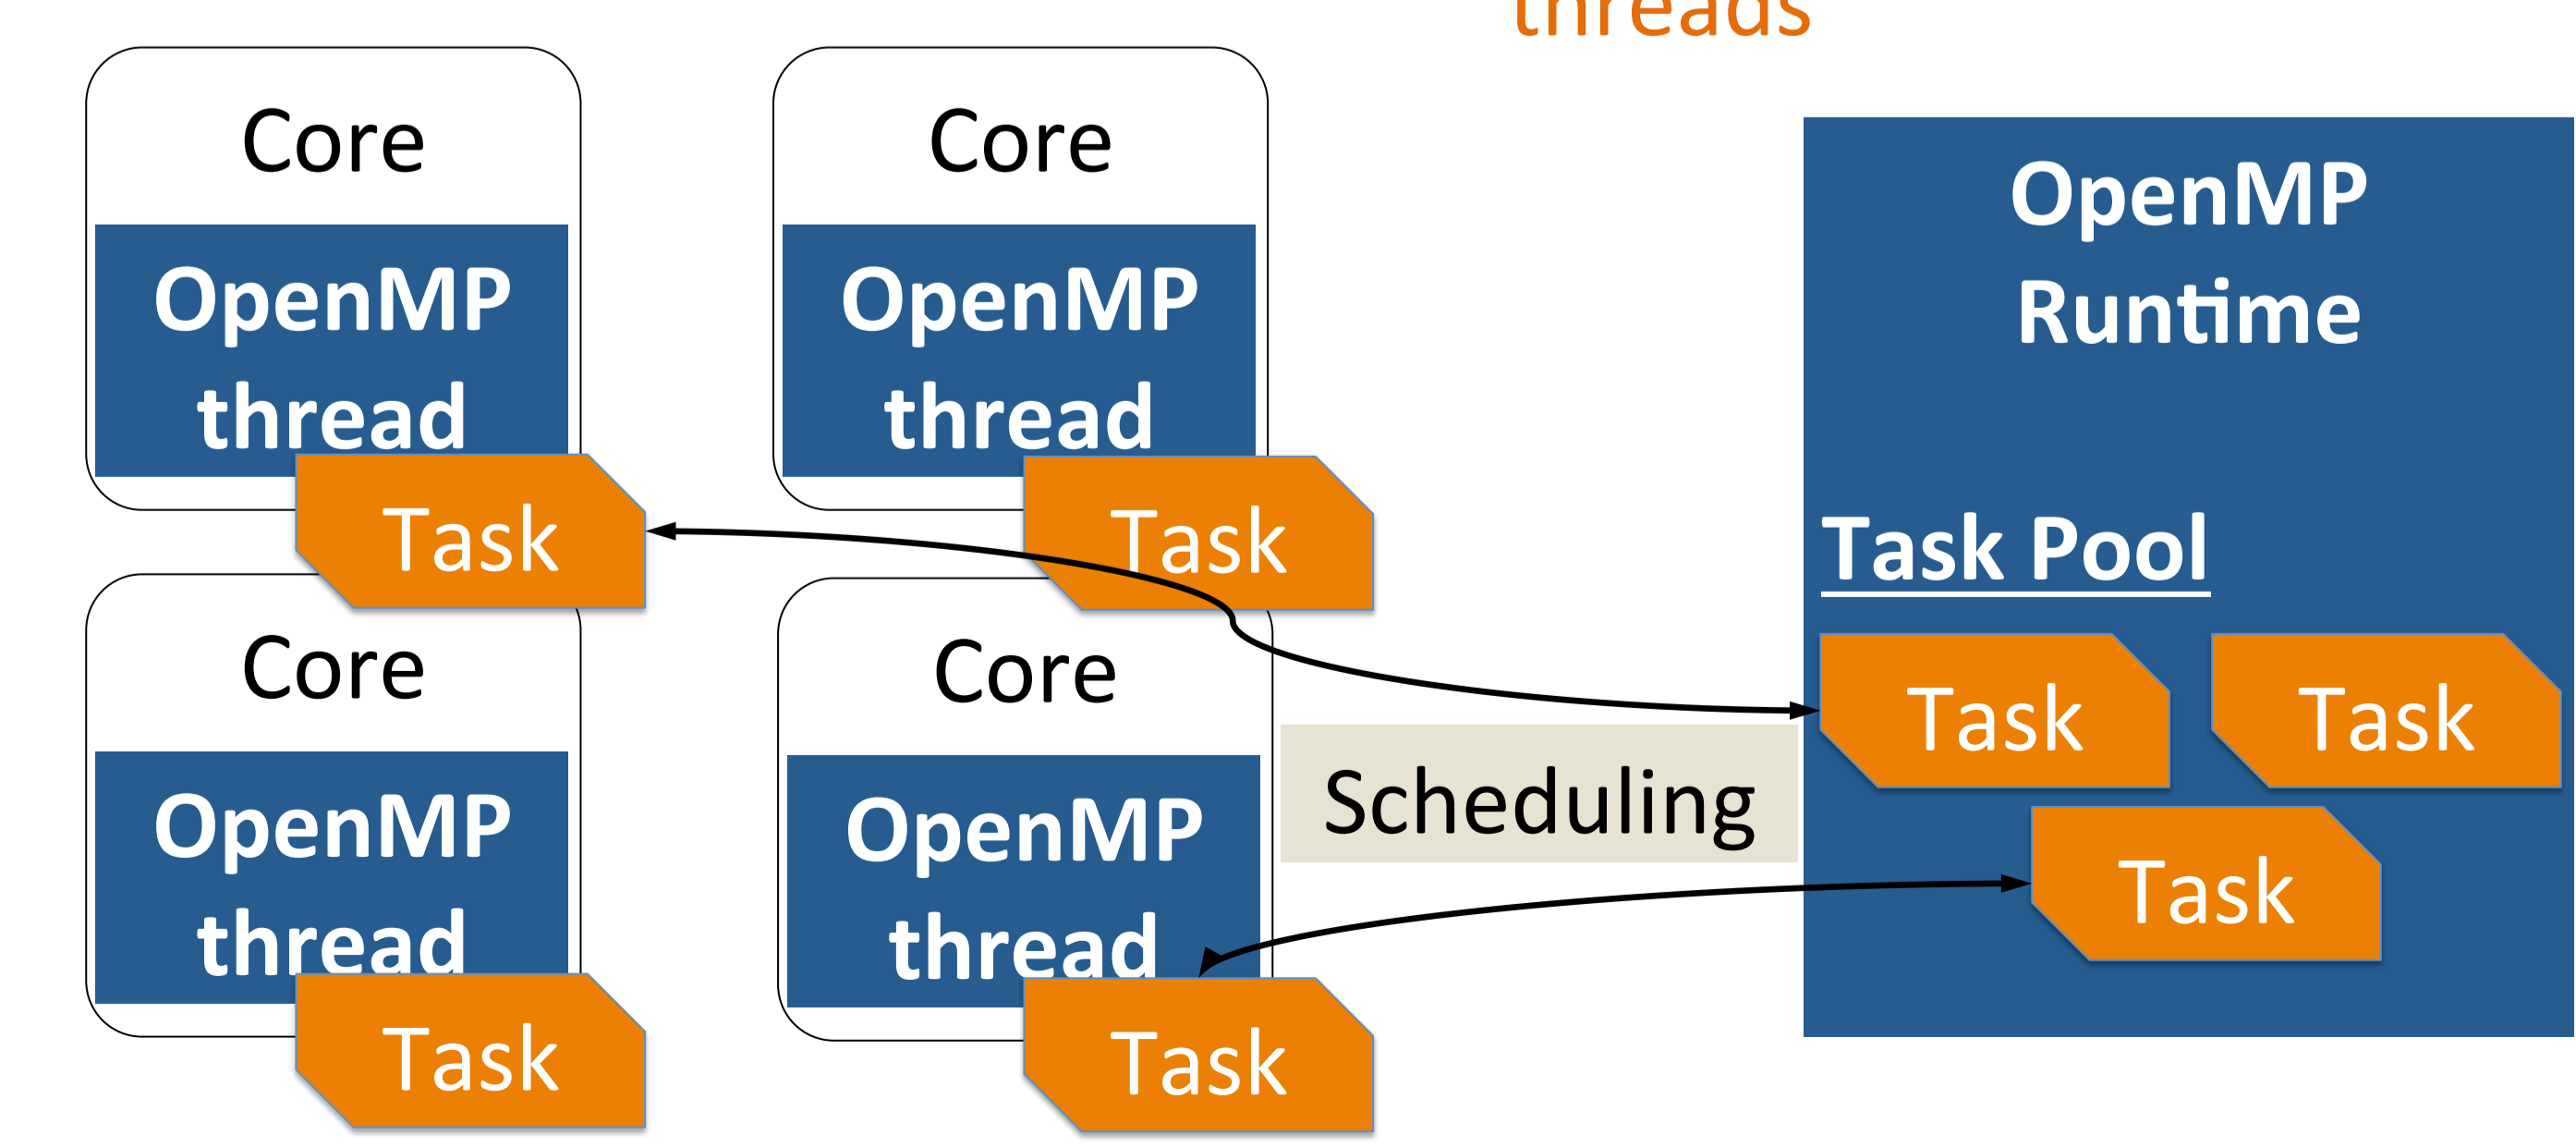
Current high performance computing applications often combine the Message Passing Interface (MPI) with threaded parallel programming paradigms, e.g., OpenMP. MPI allows fully hybrid applications in which multiple threads of a process issue MPI operations concurrently. Little study on deadlock conditions for this combined use exists. We propose a wait-for graph approach to understand and detect deadlock for such fully hybrid applications. It specifically considers OpenMP 3.0 tasking support to incorporate OpenMP's task-based execution model. Our model creates dependencies with deadlock criteria that can be visualized to support comprehensive deadlock reports. We use a model checking approach to investigate wide ranges of valid execution states of example programs to verify the soundness of our wait-for graph construction.

MPI, OpenMP, Tasks, and Hardware

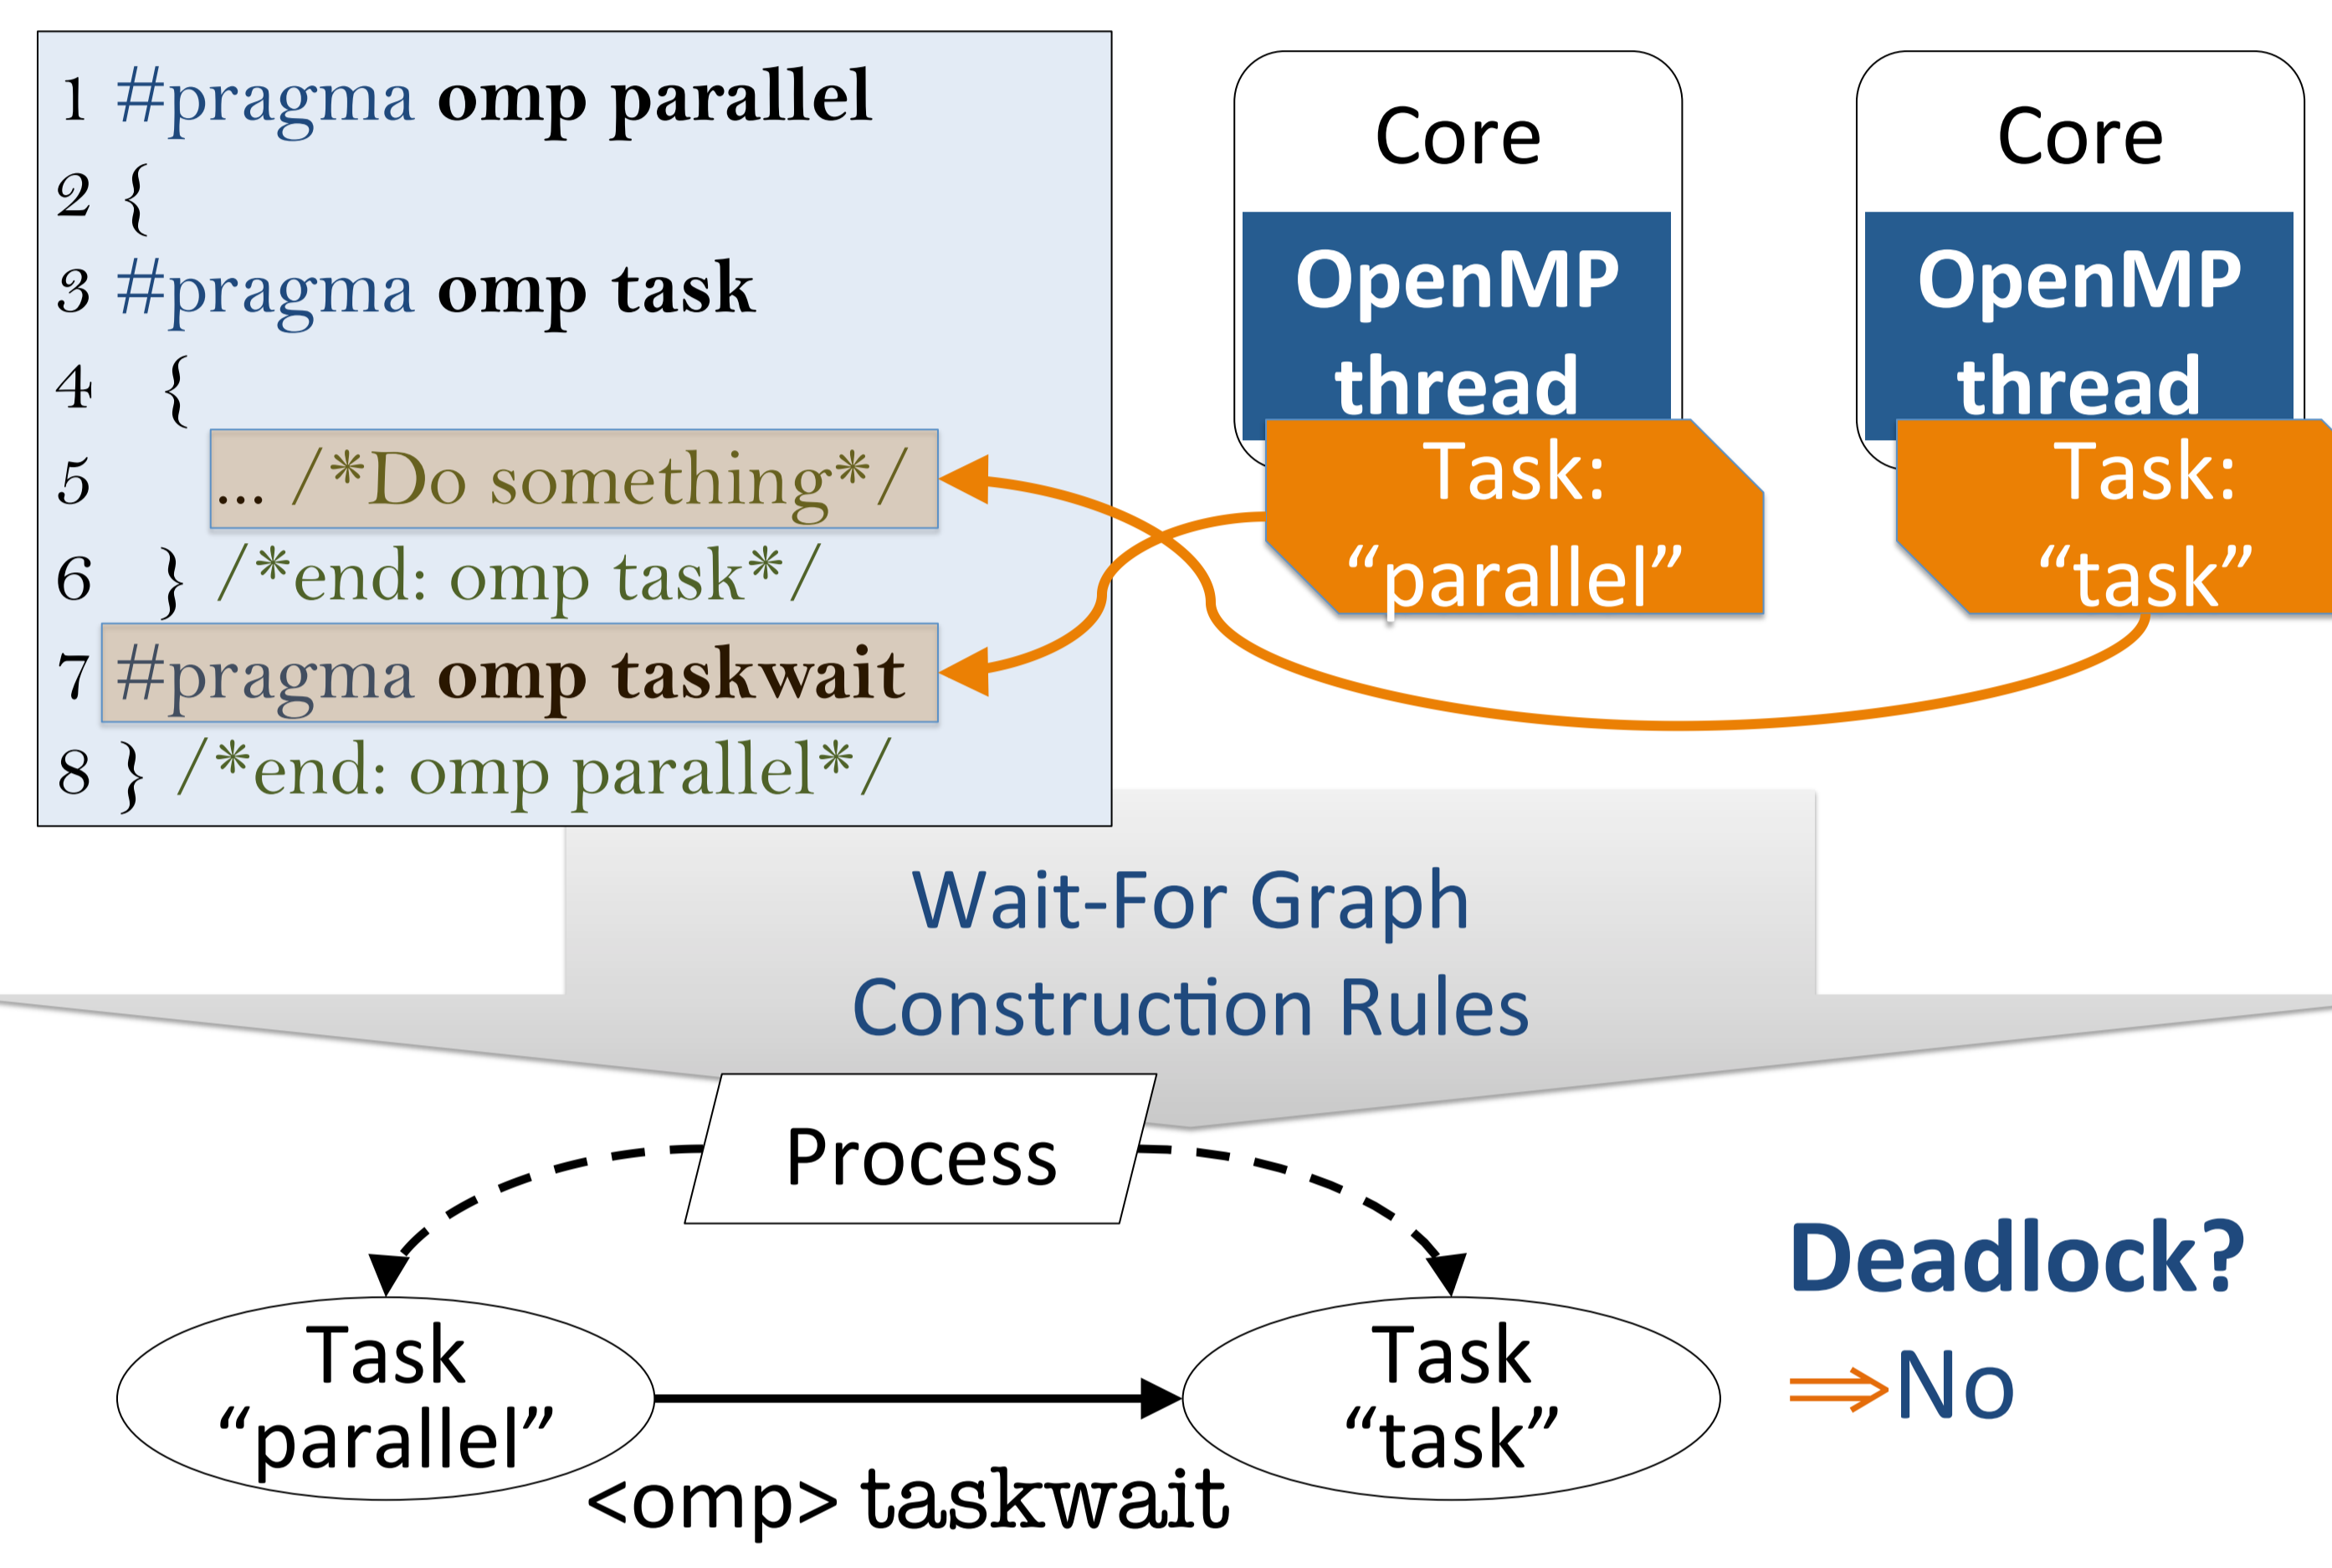
- Increasing core counts per compute node



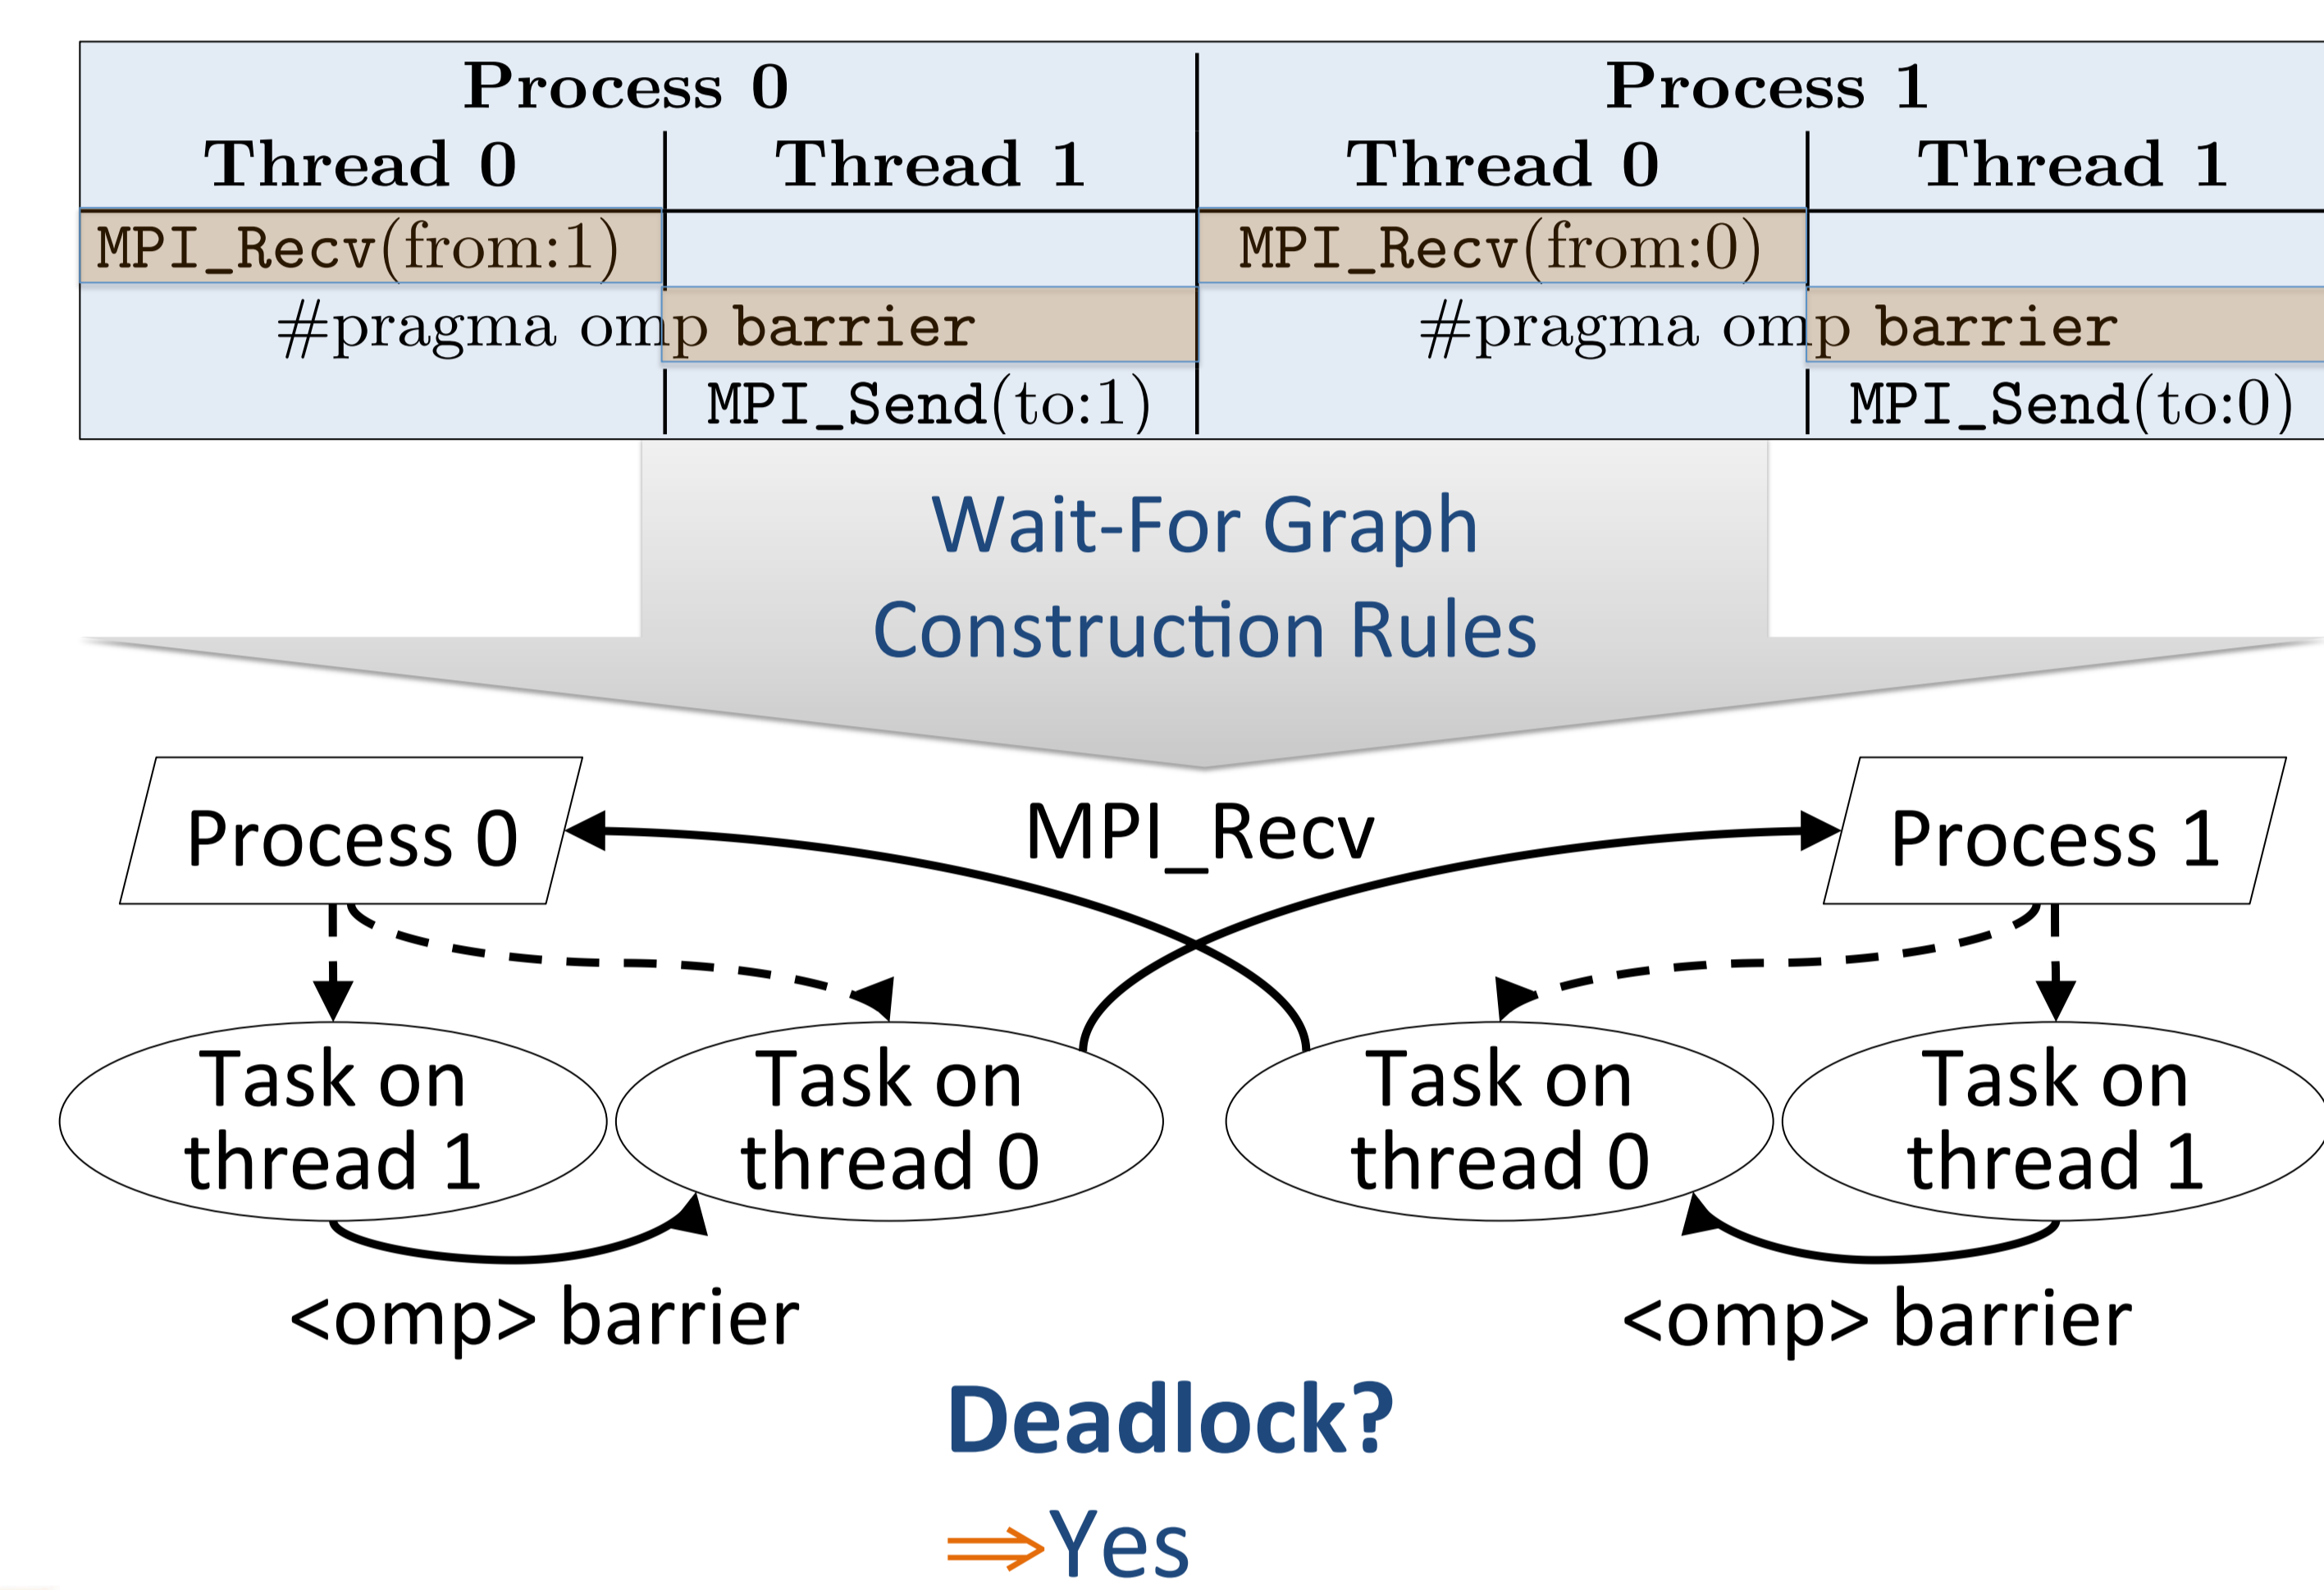
- ⇒ **Process:** From MPI, execute the same program, distinct virtual memory, message communication
- ⇒ **Thread:** From OpenMP, share memory within their process
- ⇒ **Task:** Managed by OpenMP runtime, executes on threads



OpenMP Tasks and Deadlocks



OpenMP-MPI Deadlock



Is my MPI-OpenMP application deadlocked?

Wait-For Graph (WFG) Construction Rules

- Each MPI process i get one node in the WFG
- Each task t_i of a process gets one node in the WFG
- A node for process i waits for all of its task nodes t_i (OR semantic)
- An unscheduled task t_i waits for its parent process i (with AND semantic)
- If any task t_i is at a task scheduling point, ignore rule IV. for all tasks that could replace t_i
- Model individual wait-for dependencies for MPI/OpenMP directives active on tasks
- MPI Dependencies target process nodes, not task nodes

Soundness

- Are the construction rules sound?
- Approach: Enumerate a large set of
 - ⇒ Evaluate large numbers of application execution states
 - ⇒ Apply construction rules to these states
 - ⇒ Test WFG for deadlock
 - ⇒ Compare result to expectation => Should this state have a deadlock or not?
 - ⇒ We demonstrate no false positives/negatives for 13 examples

