Н

BACKGROUND

<**Project>** provides an abstraction of the **Parti**tioned Global Address Space (PGAS) programming paradigm based on C++11 by providing distributed data structures.

5

The PGAS paradigm decouples data transfer and synchronization and thus requires explicit synchronization primitives. So far, <Project> relied on collective operations for synchronization.

Array<int> arr(size()); // fill local part arr.local[0] = myid(); // wait for everyone to finish barrier(); // read value from unit 0 auto root_id = arr[0];

R

Figure 1: Use of collective synchronization in <**Project**>

DISTRIBUTED TASK DATA DEPENDENCIES



Figure 3: Communication scheme between schedulers on two processes. Dependencies are handled by the process that "owns" the referenced memory location.

Local scheduler handle locally created tasks as well as the dependencies on the local portion of the global memory space ((i) in Figure 3). Upon encountering a remote dependency, the local scheduler informs the respective scheduler, which **coordinates** the global execution of tasks with dependencies on this memory location. The scheduler also handles data hazards by observing WAR and WAW dependencies.

Regardless of the number of tasks only two global synchronizations are required to ensure that 1) all user-specified dependencies and 2) the resulting data hazards have been communicated. The task execution happens fully asynchronous.

GLOBAL TASK DATA DEPENDENCIES FOR PGAS APPLICATIONS <Author Names Hidden for Double-blind Review>





Figure 4: Blocked Cholesky factorization implemented using task data dependencies in <Project> (simplified).

REFERENCES



METHODOLOGY



Above: distributed task graph of a regular 1D stencil application running on 2 processes. Local dependencies are created in serial order and are thus trivial to match (depicted as black lines). Dependencies across process boundaries (red lines) cannot be matched reliably without additional information from the user as there may be **multi**ple tasks with output dependencies on the same memory location.

Our proposal: Task Phases Task phases serve as a **logical clock** for tasks and task dependencies that provide the necessary information to reliably match dependencies across process boundaries. Similar to barriers, task phases restore happens-before relations between tasks, i.e., any task that produces data in memory location L_k in phase ph_{i-1} has to be executed before tasks in phase ph_i that depend on L_k .



Email



Distributed data dependencies fulfilling the requirements depicted in Figure 2, supporting:

• Seamless global in/out dependencies • Flexible **copyin dependency** (implemented using get or send/recv) • Task scheduling **priorities** • Global/local task cancellation • Rescheduling task-yield • Intermediate dependency matching • MPI-based **active messages** for inter-scheduler communication • Poll-free waiting for data transfers

• Support for advanced dependencies types (commutative, concurrent) • Performance optimizations • Tool support • Porting target applications

CONTACT INFORMATION

<hidden> <hidden> <hidden> <hidden>