

StarPU : A Runtime System for Hybrid Computing

Runtime team-project

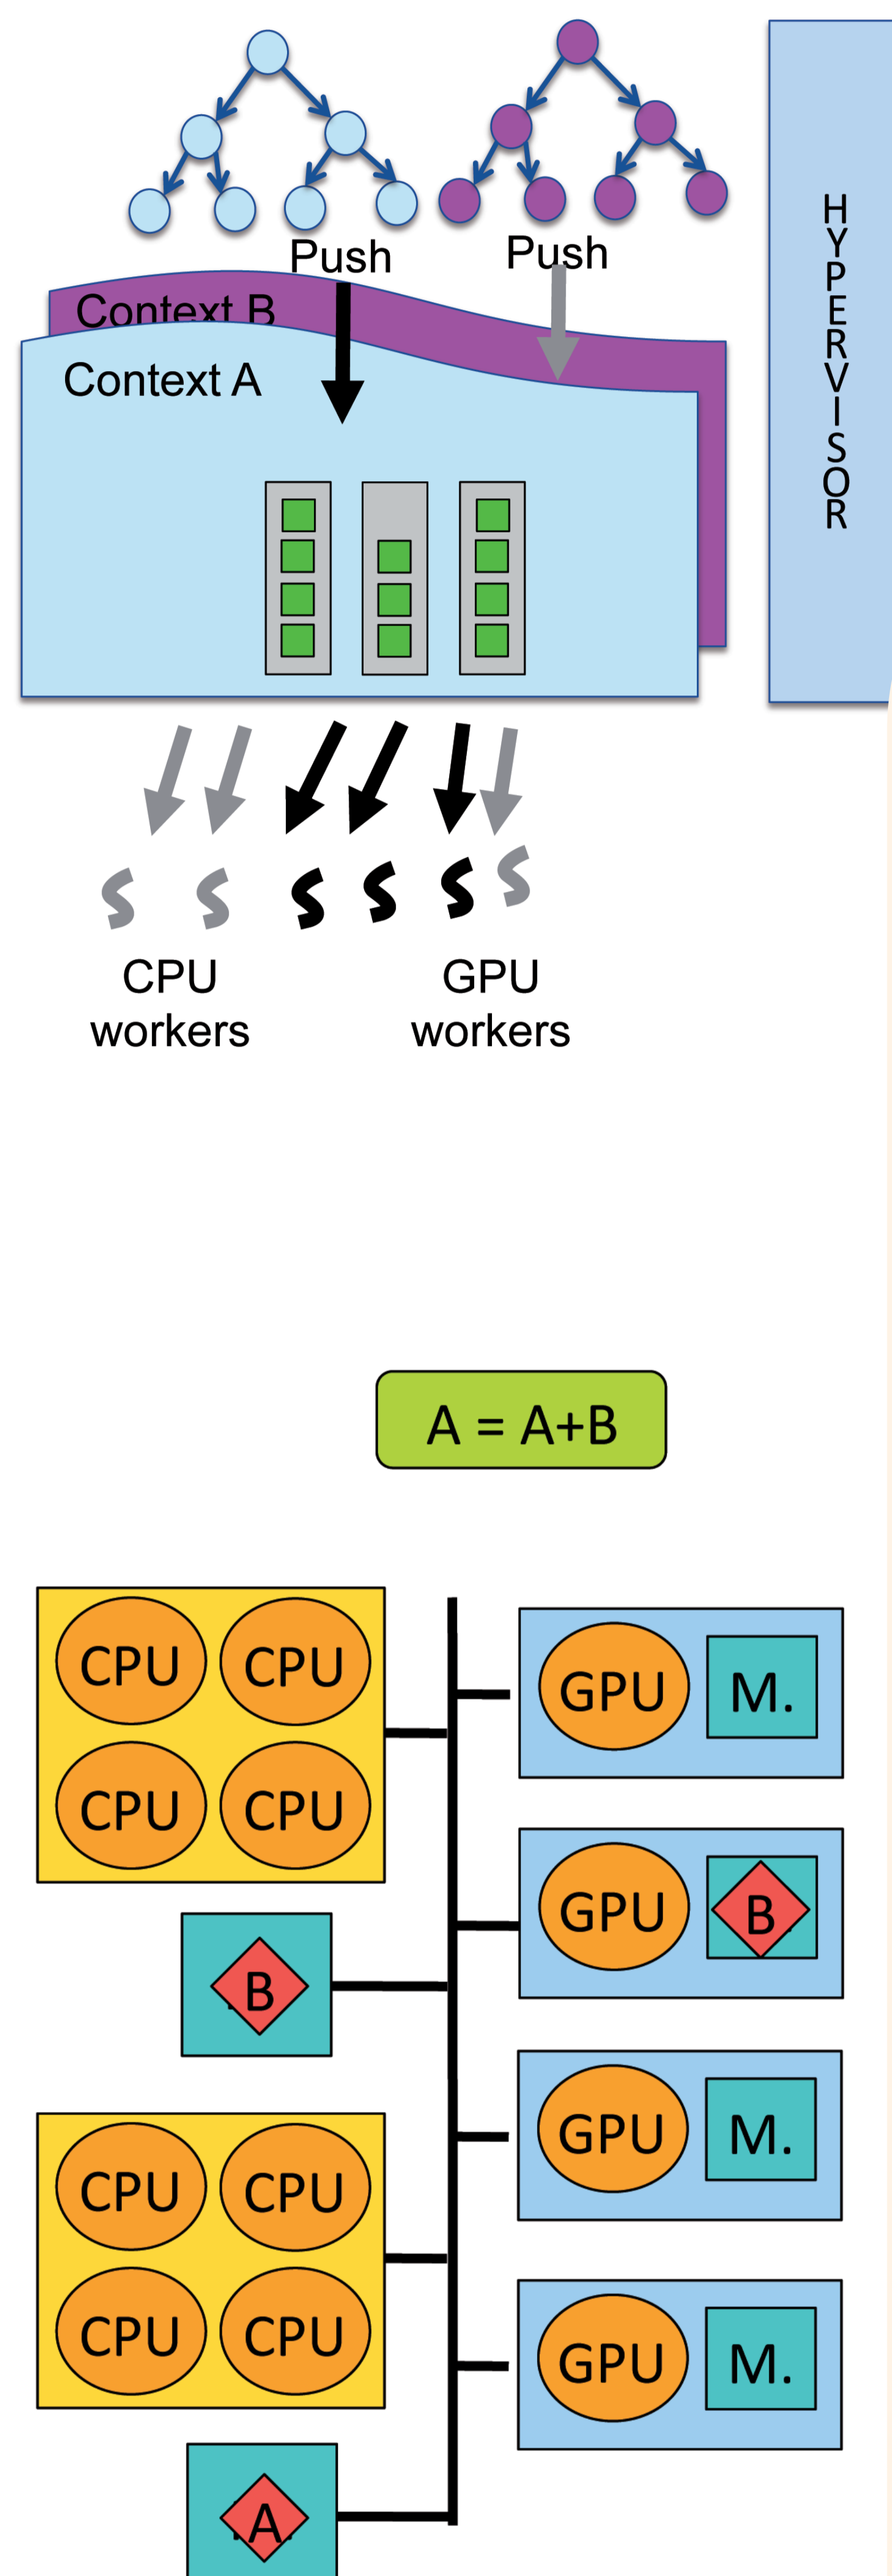
Contact: starpu-devel@list.gforge.inria.fr
http://runtime.bordeaux.inria.fr/StarPU



StarPU in a nutshell

- ▶ **Task = data input/output + multiple implementations + scheduling hints**
 - Minimize data transfers
 - Maintain software caches
 - Lazy transfers
- ▶ **Target Software**
 - Parallel Languages : HMPP [Caps Entreprise]
 - Numerical Libraries :
 - MAGMA [University of Tennessee-Knoxville],
 - qr_mumps [INRIA-CNRS]
 - Pastix[INRIA],
 - ScalFMM[INRIA],
 - Algorithmic Skeletons: SkePU [University of Linköping]
- ▶ **Supported accelerators**
 - Nvidia GPU
 - Intel Xeon Phi
 - Hybrid CPU+GPU processors
 - AMD Fusion
 - Intel Ivy Bridge

Easy interfacing with MPI



Scheduling Contexts

- ▶ **StarPU Virtual machines**
 - Feature their own scheduler
 - Minimize interferences
 - Enforce data locality
- ▶ **Allocation of resources**
 - Explicit :
 - Programmer's input
 - Supervised :
 - Tips on the number of resources
 - Tips on the number of flops
 - Shared processing units

Hypervisor

- ▶ **Idea**
 - Dynamically resize scheduling contexts
 - Different resizing policies
- ▶ **Optimization criteria :**
 - Minimize resources' idle time
 - Maximize the instant speed of the resources/contexts
 - Minimize total execution of the application
 - Workload of the application provided
 - Linear programs to evaluate the best distribution of the resources

