

Elastic MPI Applications iMPI and Dynamic Resource Management

Swen Böhm, Christian Engelmann, Michael Gerndt,
Martin Schulz, Felix Wolf

Code Structure

```
MPI_Init_adapt (... , &status) ;  
for (...) {  
    MPI_Probe_adapt (&adapt, ...) ;  
    if (adapt) {  
        MPI_Comm_adapt_begin (...) ;  
        // redistribution code  
        MPI_Comm_adapt_commit (...) ;  
    }  
    // compute and MPI code  
}
```

Discussion

- MPI standardization
 - Integration of elastic MPI as part of the Sessions working group (2025)
- Extensions:
 - poll frequency, interrupt,
 - Interface extensions: constraints, hints
- Advantages
 - Power and energy management
 - Increased application efficiency
 - Job scheduling

Discussion

- Related work and applications for dynamic resource management
 - Charm++
 - Workflow applications
 - Task parallel languages
 - Heterogeneous programming (Deep)
 - Coupled applications
 - System maintenance
- Integration into resource manager
 - Slurm probably the most important one

Discussion

- Complexity of adapting applications
 - Currently done by the programmer
 - Libraries like LAIK might help.
 - Applications might have redistribution support by checkpointing
- Relationship to Spark@HPC
 - Integration with iRM at TUM
 - Potential extension of existing works (Spark on MPI)