

Extending the MPI Backend of X10 by Elasticity

Marco Bungart and Claudia Fohry

Research Group Programming Languages / Methodologies

University of Kassel, Germany

{marco.bungart|fohry}@uni-kassel.de

U N I K A S S E L
V E R S I T Ä T

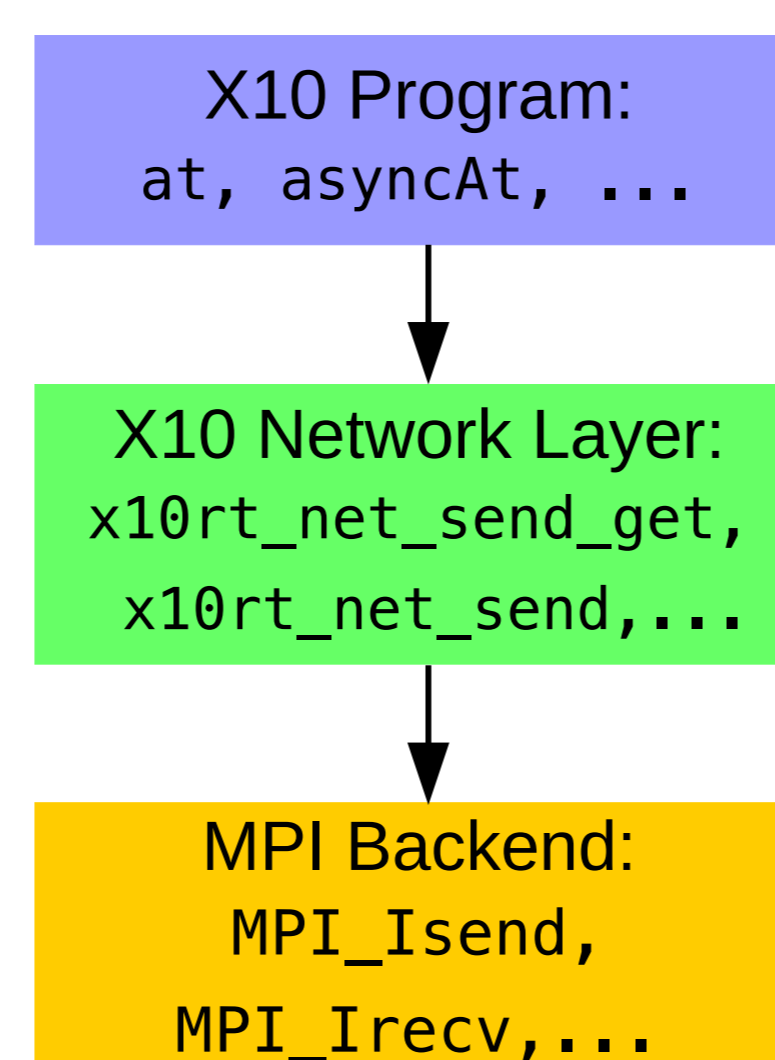


Overview

- Elastic X10 allows dynamic addition of places
- We implemented elasticity in MPI backend
- Switched on via environment variables: `X10_ELASTIC=1` and `X10_JOIN_EXISTING=<hostname>[:<port>]`
- Initial connection establishment through network socket
- Before replacement of `MPI_COMM_WORLD`, finish all communication (silent state)
- MPI connection establishment via `MPI_Comm_accept` and `MPI_Comm_connect`
- `MPI_Intercomm_merge` and `MPI_Group` help to construct new `MPI_COMM_WORLD`

MPI Backend of X10

- Implemented with Open MPI ULFM, thus `MPI_THREAD_SERIALIZED`
- Ranking must include all processes, including failed ones
- We use `MPI_Issend`



Acknowledgements

Special thanks go to Sara Salem Hamouda for providing an introduction to the design of X10s MPI backend. The ideas for this work have been developed during the cooperation within the project "Fehlertoleranz und Ressourcenelastizität für globale Taskpools" (FO 1035/5-1) funded by the Deutsche Forschungsgemeinschaft (DFG).

References

- [1] Wesley Bland. User level failure mitigation in MPI. In *Proc. Euro-Par*, pages pp. 499–504. Springer LNCS 7640, 2012.
- [2] David Cunningham, David Grove, Benjamin Herta, et al. Resilient X10: Efficient failure-aware programming. In *Proc. ACM SIGPLAN Symp. on Principles and Practice of Parallel Programming*, pages pp. 67–80, 2014.
- [3] Sara S. Hamouda, Benjamin Herta, Josh Milthorpe, David Grove, and Olivier Tardieu. Resilient X10 over MPI user level failure mitigation. In *Proc. ACM SIGPLAN X10 Workshop*, pages pp. 19–23, 2016.
- [4] Claudia Leopold, Michael Süß, and Jens Breitbart. Programming for malleability with hybrid MPI-2 and OpenMP - experiences with a simulation program for global water prognosis. In *Proc. European Conference on Modelling and Simulation*, pages pp. 665–670, 2006.

