



Energy-Efficient Heterogeneous COmputing at exaSCALE

Iakovos Mavroidis

email: iakovosmavro@gmail.com

Telecommunication Systems Institute, Technical University of Crete- Greece

EXDCI Workshop, 10 May 2016, Prague



Outline

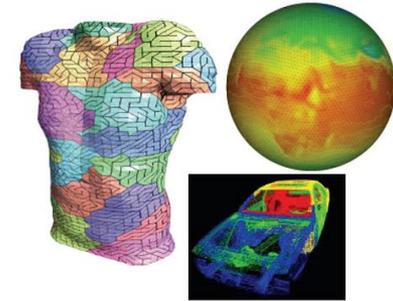
- ▶ **Challenges**
- ▶ **Approach**
 - Architecture
 - Sharing of Reconfigurable Resources
 - Runtime System
- ▶ **Collaborations**
- ▶ **Conclusion**

Challenges and Objectives

- ▶ **Scalability:** 100 millions cores, billions of tasks
 - 👉 1000x more performance
 - 👉 Hybrid MPI+OpenCL
 - 👉 UNIMEM+UNILOGIC (efficient resource sharing)
- ▶ **Energy–efficiency:** projected 1 GW.To reach 20MW...
 - 👉 50x more energy efficient
 - 👉 Reconfigurable Computing
- ▶ **Resilience:** projected MTBF less than 1 hour
 - 👉 1000x more reliable
 - 👉 Reconfiguration
 - 👉 Checkpointing

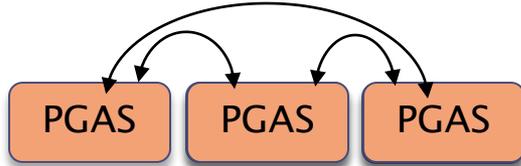
Holistic Approach

Massive Parallelism
Partition \Rightarrow Data Locality

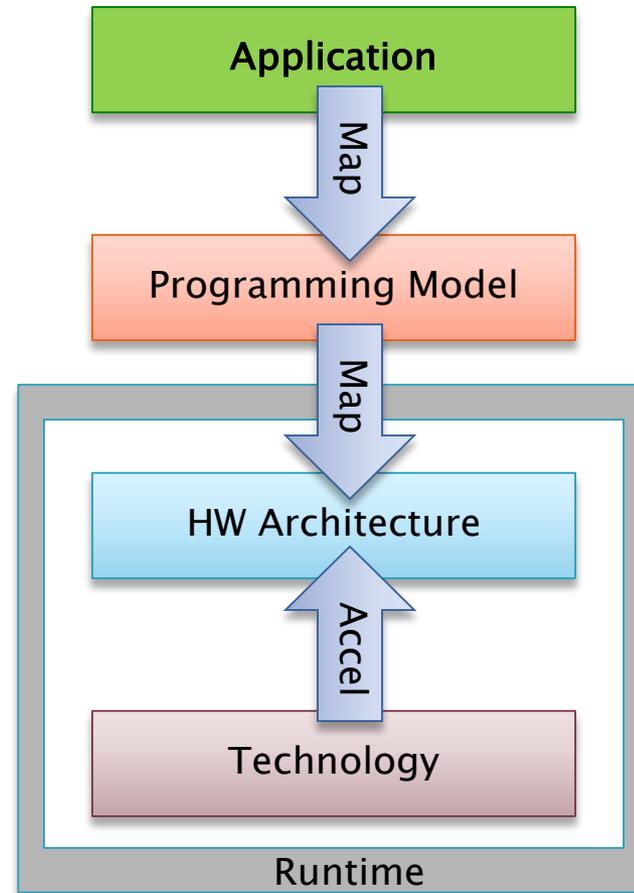
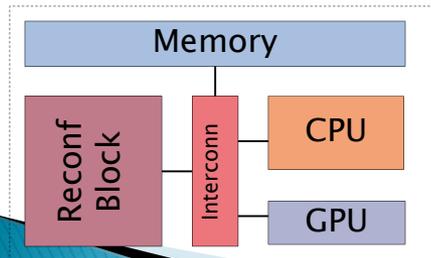


Hybrid: MPI+PGAS

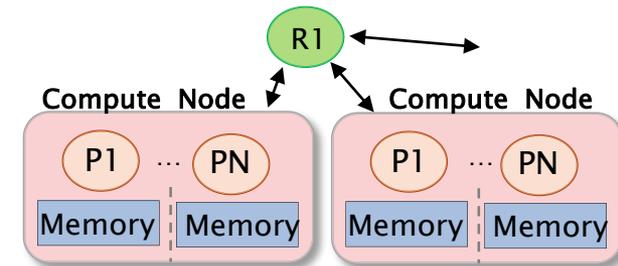
Hierarchical MPI



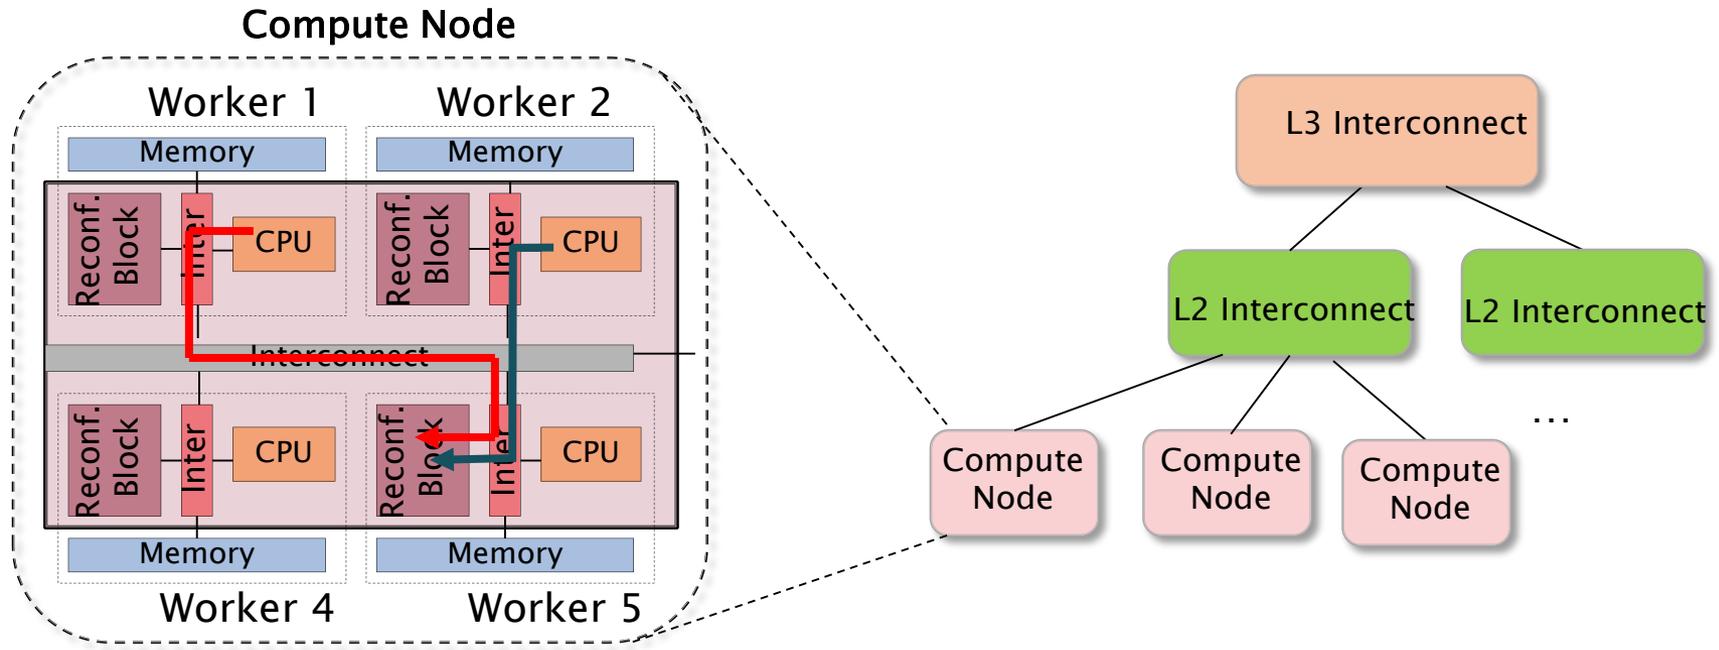
Heterogeneous Node:
Reconf. Accel.



Compute Node (PGAS)

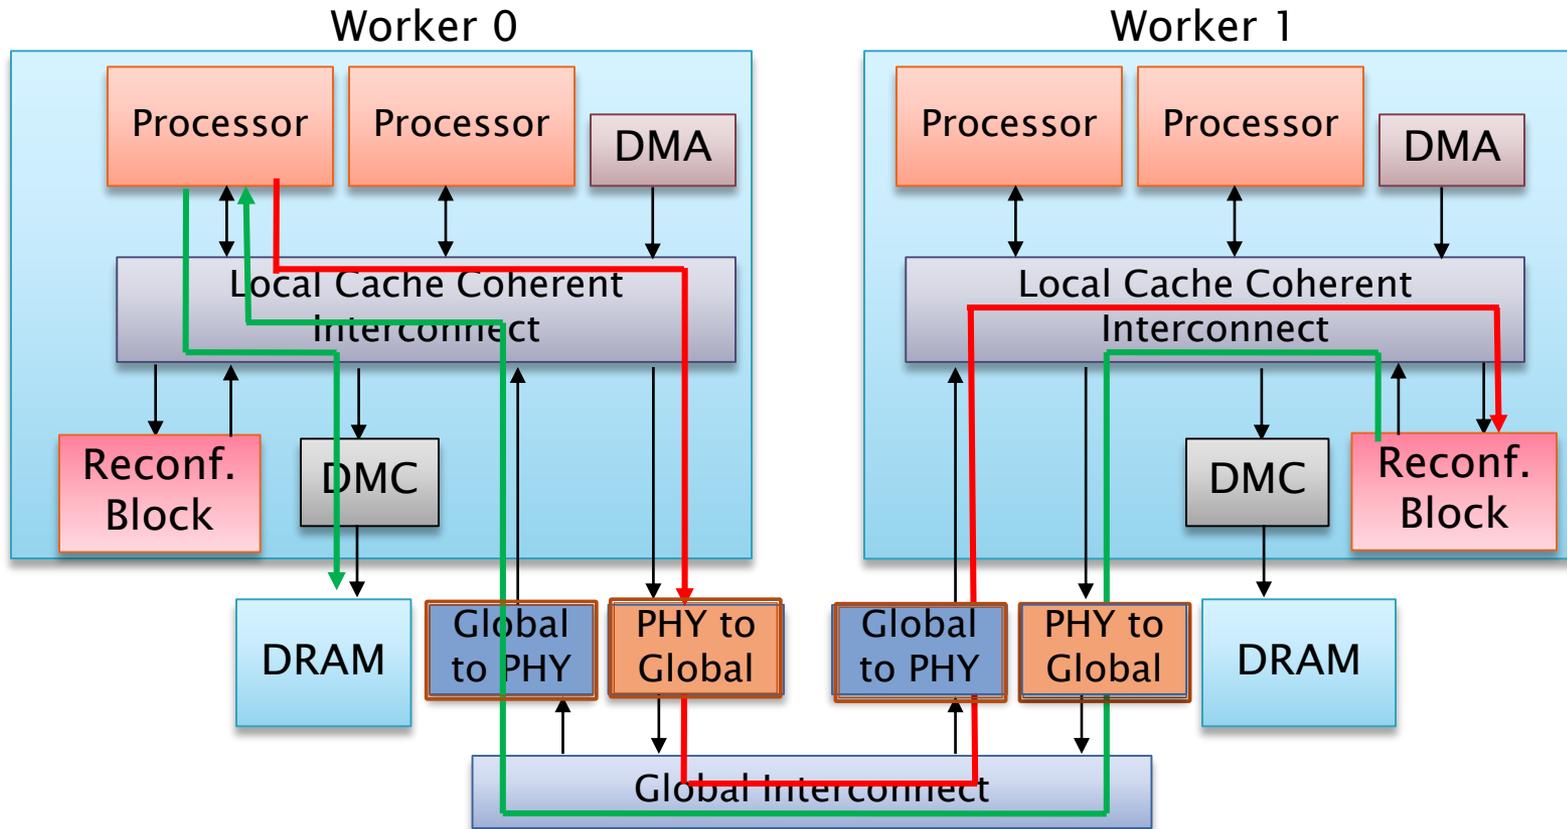


HW Architecture



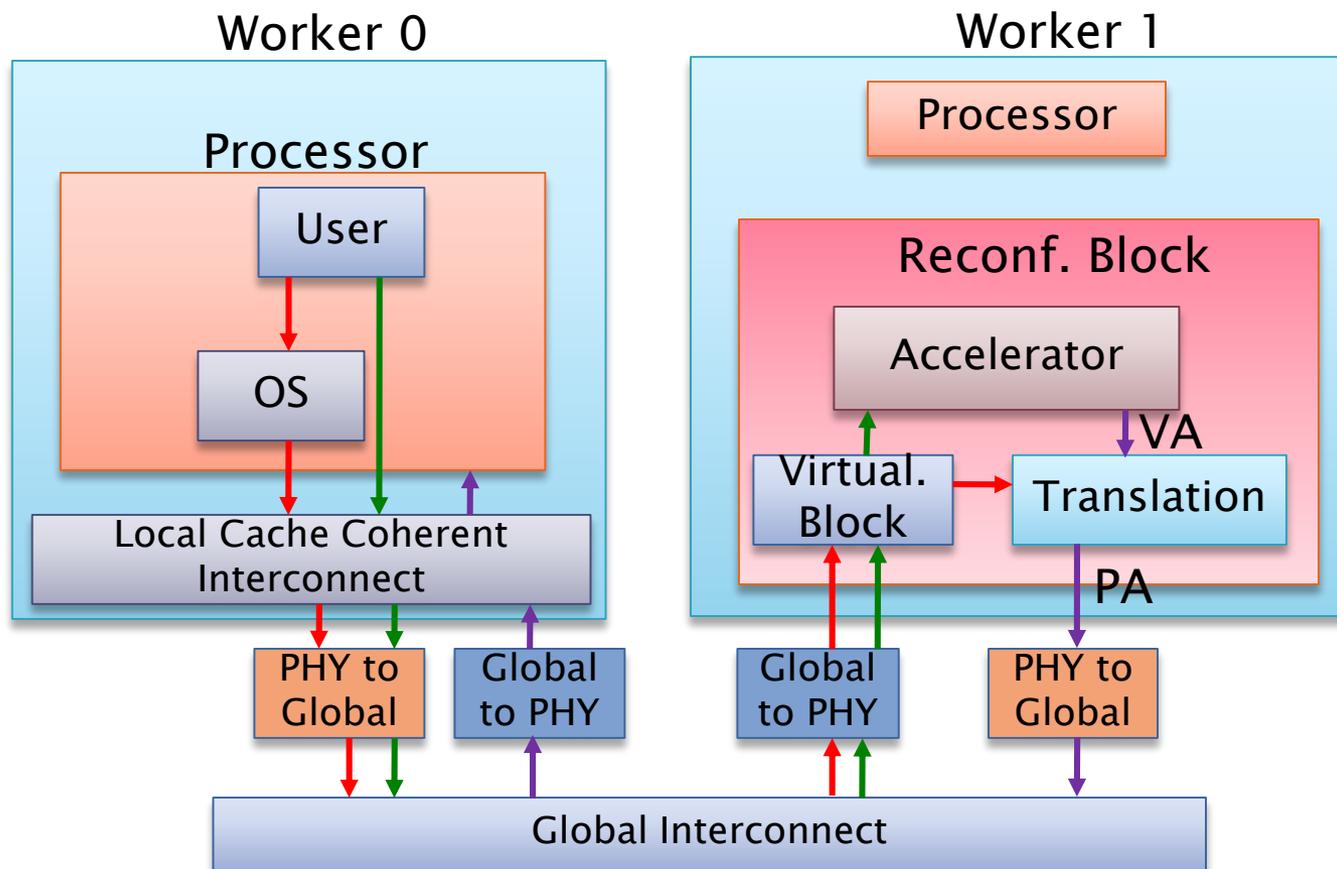
- ▶ Hierarchical Topology (ExaNeSt)
- ▶ Heterogeneous Computing (OpenCL tasks)
- ▶ Resource Sharing (UNIMEM+UNILOGIC)
 - Access to Remote Reconf. Resources
 - Sharing of Reconf. Resources

Remote Acceleration



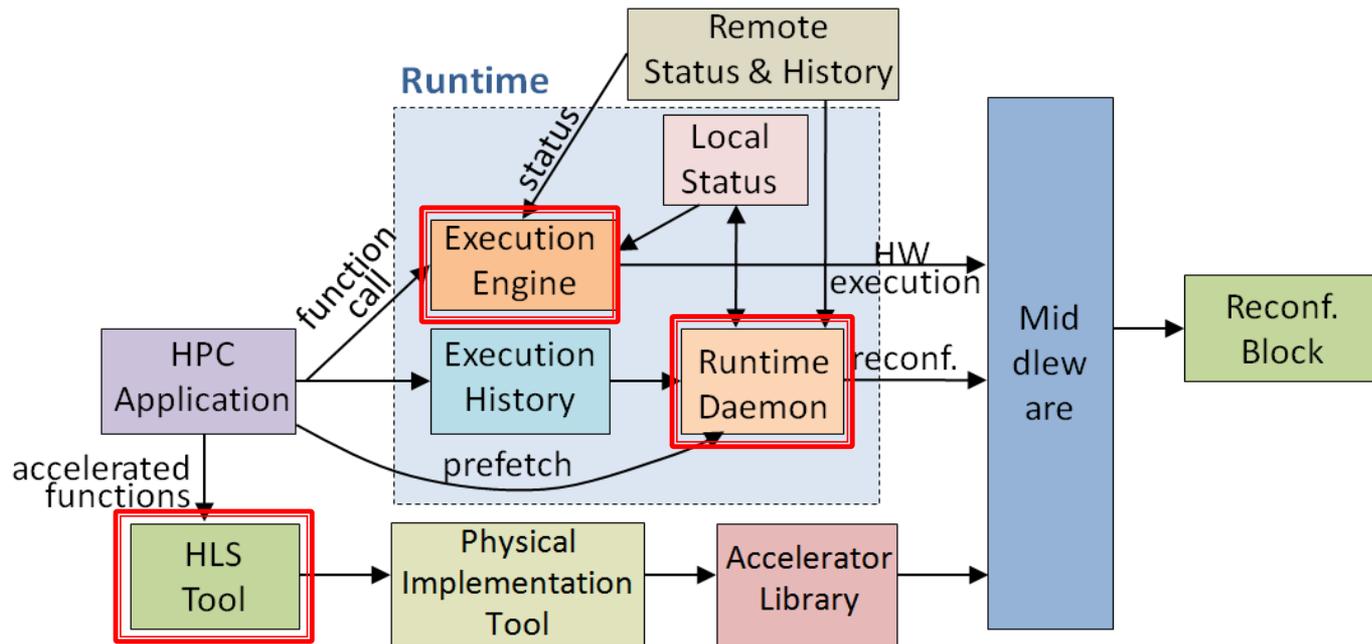
- ▶ UNIMEM provides global addr. space
- ▶ Remote Reconfigurable logic access
- ▶ Remote coherent memory access

Shared Reconfigurable Resources



- ▶ Configure Remote Translation Table
- ▶ Remote OpenCL call
- ▶ Remote Execution

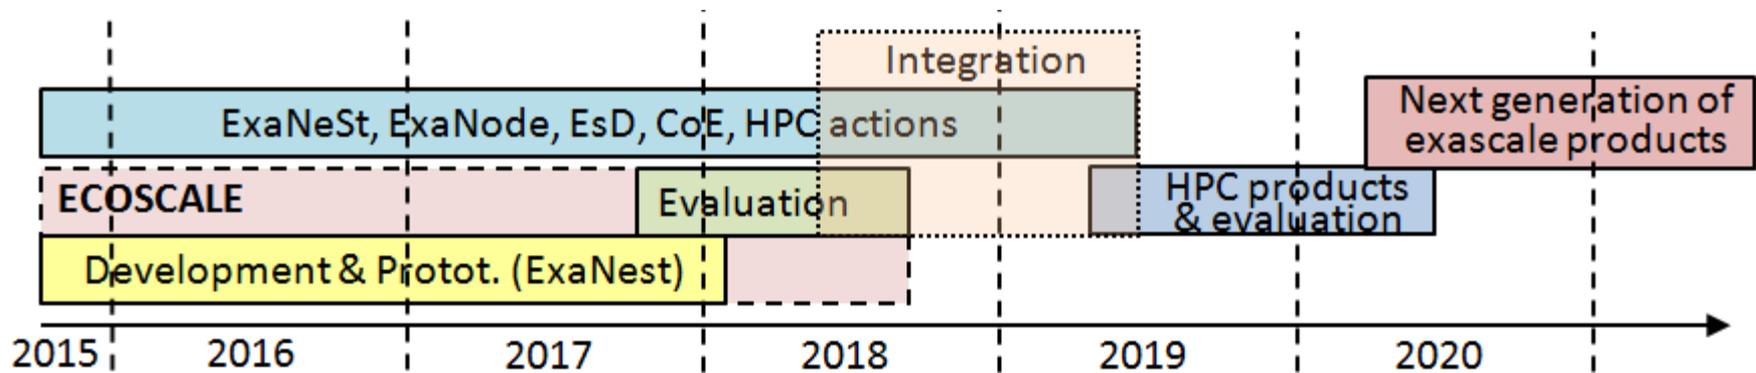
Runtime System and HLS



- ▶ Reconfiguration at runtime
- ▶ Task Execution Scheduling
- ▶ Task Execution Monitoring
- ▶ Locality Management

Collaborations / Common Goals

- ▶ **EuroExa group:** ECOSCALE, ExaNeSt, ExaNode
- ▶ **SRA roadmap:**
 - Stepping stone towards future energy-efficient reconfigurable computing
 - EsD's can provide first ECOSCALE-based platforms
 - Centers of Excellence (CoE) can provide ECOSCALE-based applications



Conclusion

- ▶ UNILOGIC offers
 - Distributed Reconfigurable Logic
 - Scalability
 - Logic Sharing
 - Direct Accesses
 - Cache Coherency (UNIMEM)
- ▶ ECOSCALE Runtime System
 - OpenCL Tasks
 - Reconfiguration at runtime
 - Task execution scheduling
 - Task Execution Monitoring

Stay Tuned!



www.ecoscale.eu



@ecoscale_H2020

