

EPIGRAM

Exascale ProGRAMming Models

Daniel Holmes
EPCC, University of Edinburgh

Message Passing and PGAS

Key Objectives

- Address the scalability (performance and memory consumption) problem for MP and PGAS models.
- Propose GPI as the European PGAS approach to exascale.
- Design a hybrid MP-PGAS programming model that combines the best features of the two approaches.
- Prepare two applications to exascale by redesigning and implementing their communications kernels.

State of the Art and Beyond

- Two programming models have demonstrated potential for exascale performance: MP and PGAS.
- Certain MPI collective operations are not scalable.
 - **Revise MP algorithms for scalable collectives.**
- PGAS approach has potential for exascale, but it requires considerable changes in codes to use it.
 - **Deploy a PGAS library (GPI) that does not require major restructuring in two real-world codes.**
 - **Design an MPI implementation based on the PGAS approach to combine advantages of MP and PGAS.**
- Several programming models for moving data between host and accelerators are being proposed.
 - **Investigate the potential of new approaches for diverse memory spaces.**

Rationale for EPIGRAM

- Message Passing programming model is the most used approach on peta-scale systems.
- PGAS approach has potential to be an exascale programming model.
- **Both approaches have limitations and they will not scale at exascale.**

EPIGRAM Vision

We will introduce new disruptive concepts in MP and PGAS programming models to fill the technological gap between petascale and exascale era in two ways:

- **Innovative and disruptive algorithms** will be used in both **MP** and **PGAS** approaches.
- We will **combine the best features of MP and PGAS** programming models, by developing and implementing an MP interface using a PGAS library as communication substrate.

A Window of Opportunity

- We have the chance to take a leading role in international MP programming model research.
- By extending and improving GPI to exascale we will consolidate the role of GPI and establish it as the European PGAS approach.
- EPiGRAM can complement the European CRESTA, DEEP, and Mont-Blanc exascale projects.
 - by exploring additional innovative PGAS approaches that go well beyond those considered in the current CRESTA project
 - by investigating efficient MP mechanisms that might be useful for hybrid Cluster-Booster architecture in DEEP
 - by studying and analyzing one-sided communication approaches for diverse memory spaces such as the one in hybrid ARM-GPU systems in Mont-Blanc.

Exascale Message Passing

- Objectives:
 - Investigate new, low-overhead, MP concepts and algorithms
 - initial focus on collectives
 - also threading, fault tolerance
 - Develop concrete specifications
 - go beyond the current standard where it limits scalability (e.g. persistent collectives).

Exascale PGAS

- Objectives:
 - To investigate current limitations of traditional PGAS.
 - To propose concrete solutions to current PGAS limitations.
 - To increase scalability of collective operations and synchronization in GPI.
 - To support fault-tolerance in GPI.
 - Exploitation of diverse and hierarchical memory spaces in PGAS.

PGAS-based MPI

- Objectives:
 - Implement and evaluate efficient message passing libraries on top of RDMA operations
 - Implement and evaluate collective operations on top of RDMA operations
 - Prototype implementation of MPI endpoints proposal
 - Develop recommendations for MPI to allow efficient implementation on top of RDMA
 - Develop recommendations for RDMA hardware

Applications

- Objectives:
 - Use of the exascale MP, PGAS and PGAS-based MPI software in two real-world applications: Nek5000, iPIC3D
 - Analyze the performance of newly developed communication kernels in Nek5000 and iPIC3D
 - Provide feed-back and guidance to the development of exascale programming models

Standardisation and Dissemination

- Objectives:
 - Participate and contribute to standardisation committees (especially MPI Forum).
 - Establish collaboration with EC projects and initiatives.
 - Disseminate knowledge from EPIGRAM.