COHX: CHAPEL ON HSA+XTQ

Adventures of a PGAS Language in a Heterogeneous World

Deepak Majeti

deepak@rice.edu

Graduate Student, Rice University

2014 Summer co-op, AMD® Research

MOTIVATION

- Heterogeneous architectures are here to stay
  - Exist in current mobiles to super-computers

- PGAS language support for modern heterogeneous architectures is in its initial stages
  - Support necessary for wider adaptation of such languages

- Support Chapel on leading heterogeneous architectures
  - HSA (**H**eterogeneous **S**ystem **A**rchitecture)
  - XTQ (e**X**tended **T**ask **Q**ueuing)

CHAPEL ON HSA + XTQ

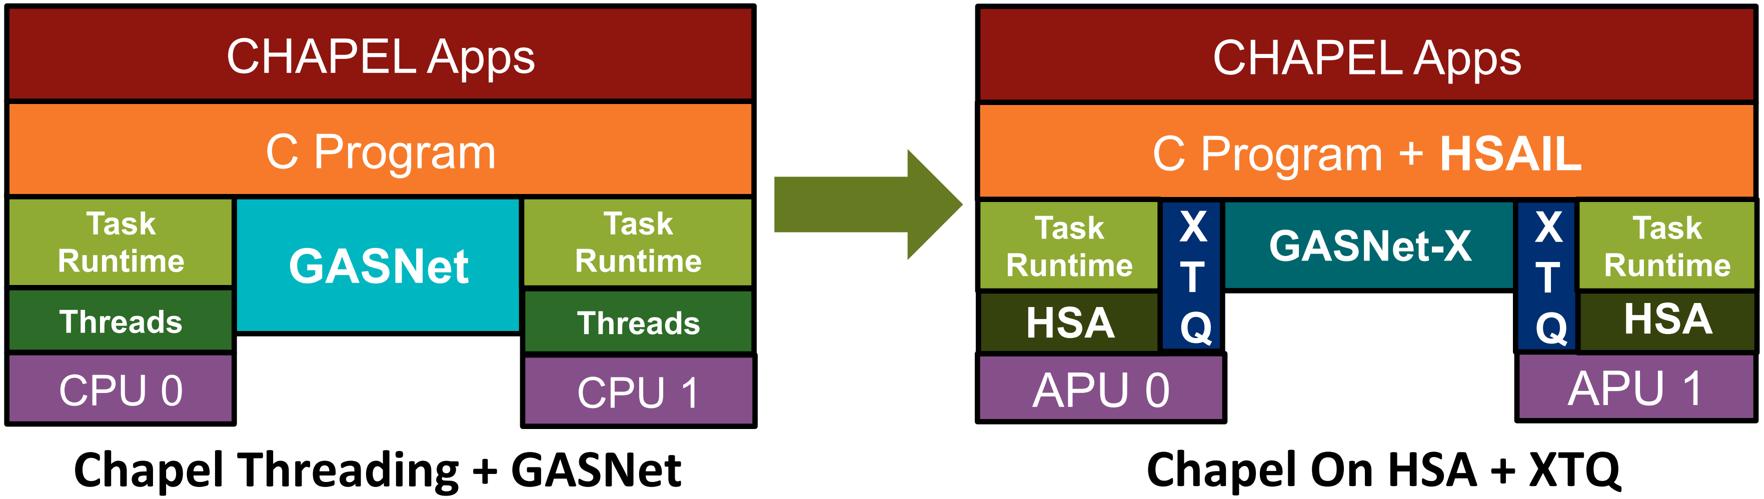

Chapel Threading + GASNet

Chapel On HSA + XTQ

- Current Chapel Framework
  - Local tasks via Threads
  - Remote tasks via GASNet Active Messages

- Proposed Chapel Framework
  - Local tasks via **HSA**
  - Remote tasks via **XTQ**

APPROACH

- Target Chapel tasks to generate HSAIL
  - Focus on “for loops” to target GPUs
- Extend the “on” construct of Chapel
  - Use HSA if execution is on a local locale
  - Use XTQ for execution on a remote locale
- Build an active message interface similar to GASNet AMs on top of XTQ
  - Use existing GASNet RDMA
- Unified memory support in HSA avoids data copies to/from GPU

FEW IMPLEMENTATION CHALLENGES

- Chapel threading layer assumes thread local storage
  - HSA has no such feature

- Improve generated C code
  - Need “clean loops” for better performance on GPU
  - Perform loop invariant code motion

- Need support for user annotations
  - Hints for the compiler to target a specific processor

CONCLUSION

- Need to support Chapel on current heterogeneous architectures for wider adaptation

- Newer heterogeneous architectures with HSA support require little compiler development effort

- Few implementation challenges which can be resolved with effort from Chapel community

ACKNOWLEDGMENTS

- AMD Research
  - Jay Owen
  - Mauricio Breternitz
  - Brad Benton
  - Mike Chu
  - Steve Reinhardt

- Computer Science Department, Rice University
  - Vivek Sarkar
  - Habanero Extreme Scale Software Research Project

DISCLAIMER & ATTRIBUTION

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

BACKUP

HSA

HETEROGENEOUS SYSTEM ARCHITECTURE

- Improved processor design to efficiently program heterogeneous architectures
  - User level task queues provide lightweight task execution
  - **“True”** unified address space with the help of hardware support
  - Seamless use of different compute units present

- HSAIL: HSA Intermediate Language
  - Interface to expose underlying hardware capabilities to the software

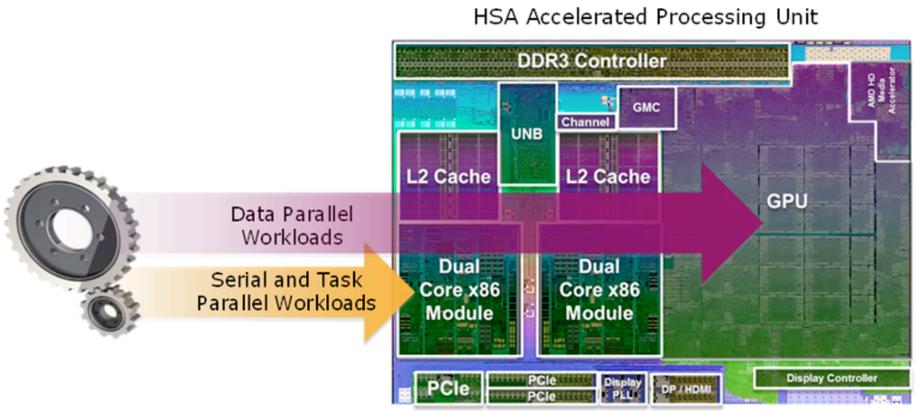

Heterogeneous System Architecture (HSA) with CPU + GPU

Img src: http://amd-dev.wpengine.netdna-cdn.com/wordpress/media/2012/10/HSAAcceleratedProcessingUnit.png

XTQ
EXTENDED TASK QUEUING

- Motivation: *Extend HSA to a distributed computing environment*

- XTQ (eXtended Task Queuing)
  - – A light weight remote task execution framework.

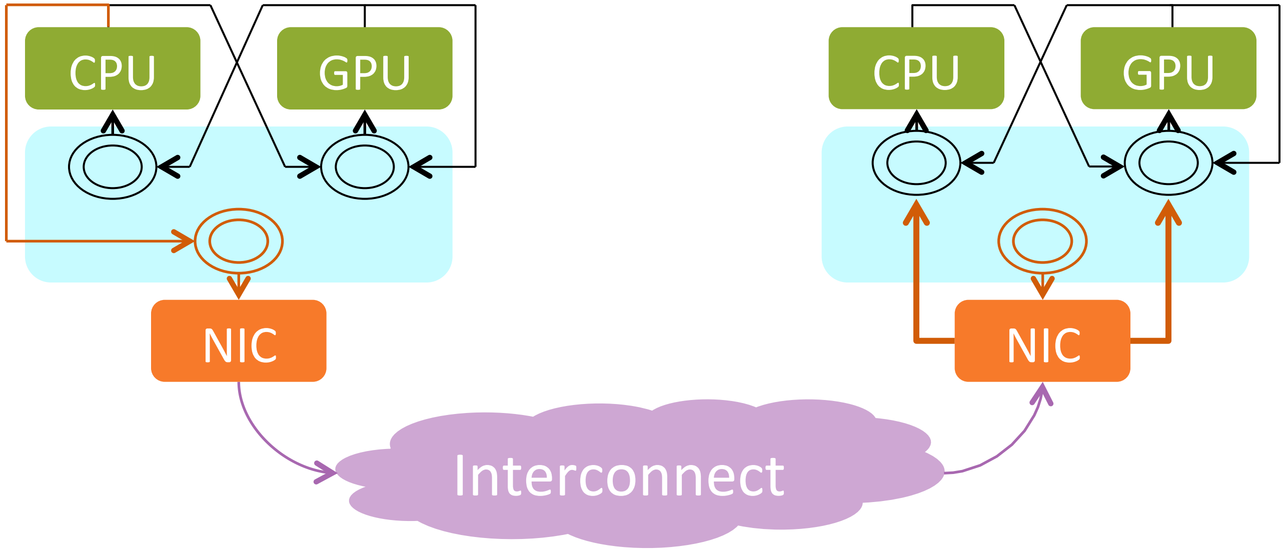
CPU
GPU
NIC
Interconnect
CPU
GPU
NIC