Evaluation of Productivity and Performance Characteristics of CCE CAF and UPC Compilers

Sadaf Alam, William Sawyer, Tim Stitt,
Neil Stringfellow, and Adrian Tineo,
Swiss National Supercomputing Center (CSCS)

CSCS
Swiss National Supercomputing Centre

Motivation

- Upcoming CSCS development platform—Baker system with GEMINI interconnect
  - Availability of PGAS compilers on XT5
- HP2C projects
- PRACE WP8 evaluation

HP2C Projects (www.hp2c.ch)

- Effort to prepare applications for the next-gen platform

- **BigDFT - Large scale Density Functional Electronic Structure Calculations in a Systematic Wavelet Basis Set; Stefan Goedecker, Uni Basel**
- **Cardiovascular - HPC for Cardiovascular System Simulations; Prof. Alfio Quarteroni, EPF Lausanne**
- **CCLM - Regional Climate and Weather Modeling on the Next Generations High-Performance Computers: Towards Cloud-Resolving Simulations; Dr. Isabelle Bey, ETH Zurich**
- **Cosmology - Computational Cosmology on the Petascale; Prof. Dr. George Lake, Uni Zürich**
- **CP2K - New Frontiers in ab initio Molecular Dynamics; Prof. Dr. Juerg Hutter, Uni Zürich**
- **Gyrokinetic - Advanced Gyrokinetic Numerical Simulations of Turbulence in Fusion Plasmas; Prof. Laurent Villard, EPF Lausanne**
- **MAQUIS - Modern Algorithms for Quantum Interacting Systems; Prof. Thierry Giamarchi, University of Geneva**
- **Petaquake - Large-Scale Parallel Nonlinear Optimization for High Resolution 3D-Seismic Imaging; Dr. Olaf Schenk, Uni Basel**
- **Selectome - Selectome, looking for Darwinian Evolution in the Tree of Life; Prof. Dr. Marc Robinson-Rechavi, Uni Lausanne**
- **Supernova - Productive 3D Models of Stellar Explosions; Dr. Matthias Liebendörfer, Uni Basel**

PRACE Work Package 8

- Evaluation of hardware and software prototypes
  - CSCS focused on CCE PGAS compilers
  - "Technical Report on the Evaluation of Promising Architectures for Future Multi-Petaflop/s Systems" *www.**prace**-project.eu/documents/d8-3-2.pdf*

1-min introduction to PGAS

- PGAS—Partitioned Global Address Space
  - Not MPI message-passing API approach
  - Not a single, shared memory OpenMP approach
- Memory model with local and remote accesses
  - Access to local data—fast
  - Access to remote data—slow
- Language extensions
  - CAF (Co Array Fortran)
  - UPC (Unified Parallel C)

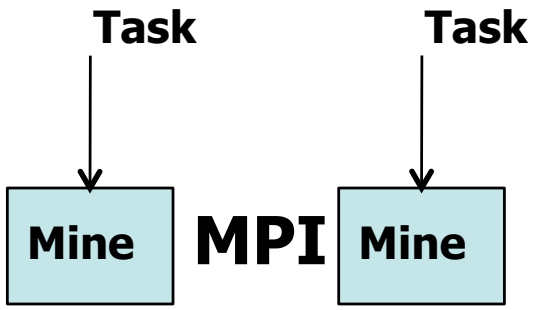

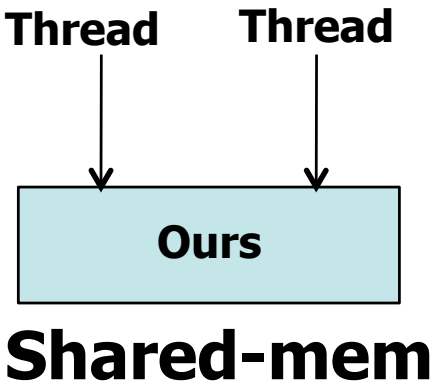

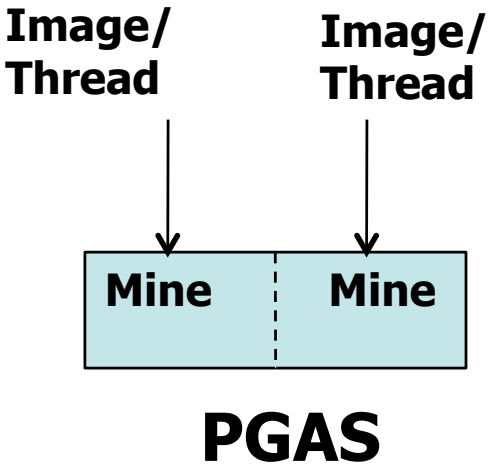

Yet another prog. Model?

- Yes and no
  - Been around for 10+ years
  - Limited success stories

- What is different now?
  - GEMINI provides NW support for PGAS access patterns
  - Compiler can potentially overlap comm./comp.

CUG 2010
EDINBURGH
SIMULATION COMES OF AGE

Target Platforms

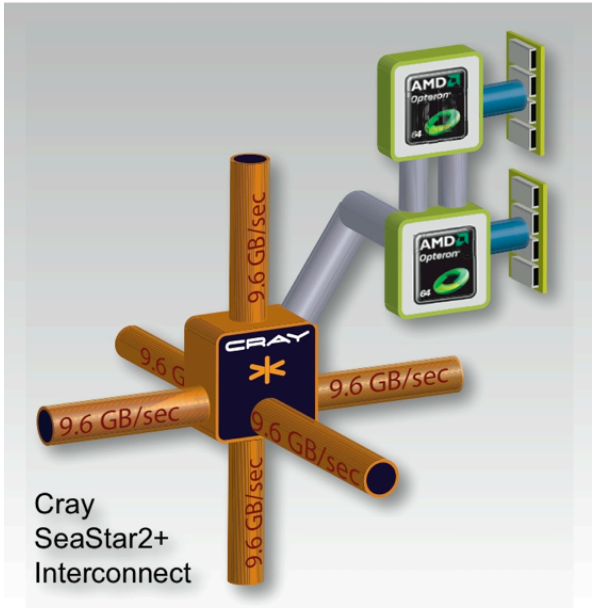

XT5 with commodity uProc and custom interconnect

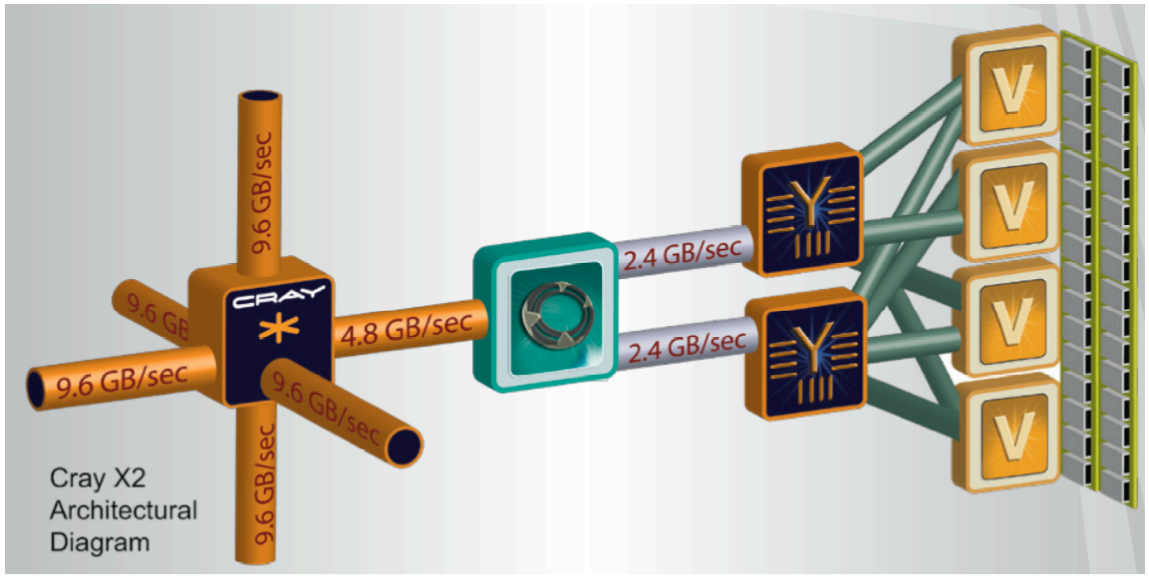

X2 with proprietary vector proc. and custom interconnect

CSCS
Swiss National Supercomputing Centre

Building Blocks of CCE PGAS Compilers

- Front end (C/C++/Fortran plus CAF and UPC)
- X86 back-end

- GASNet communication interface
  - Expected to change on GEMINI based systems

Test Cases

- Remote access STREAM
- Matrix Multiply
- Stencil based filter

X2

```
791. 1    Vr------<     DO  j = 1,n
792. 1    Vr             b(j) = scalar*c(j)[2]
793. 1    Vr------>      end DO
```

XT

```
791. 1    1-------<     DO  j = 1,n
792. 1    1              b(j) = scalar*c(j)[2]
793. 1    1-------->     end DO
```

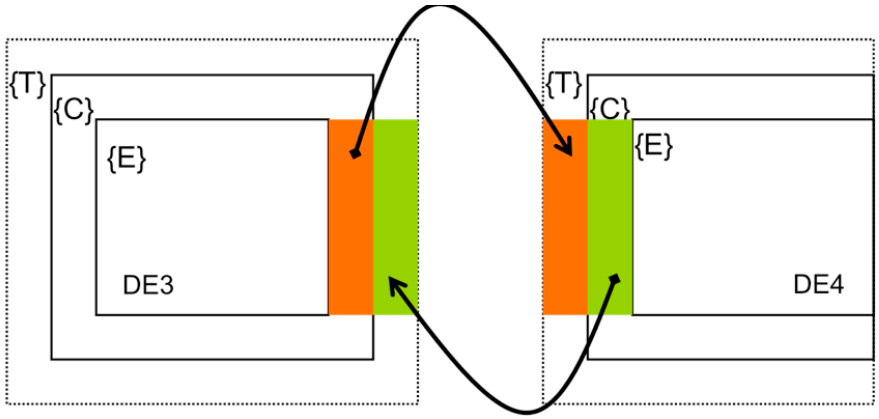

Compiler Listing

X2

```
1------<   upc_forall (i=0; i<N; i++; &c[i][0]) {
1 V----<     for (j=0; j<M; j++) {
1 V            c[i][j]=0;
1 V r--<       for (l=0; l<P; l++)
1 V r-->         c[i][j]+=a[i][l]*b[l][j];
1 V---->     }
```

XT5

```
1------<  upc_forall (i=0; i<N; i++; &c[i][0]) {
1 i----<    for (j=0; j<M; j++) {
1 i            c[i][j]=0;
1 i 3--<       for (l=0; l<P; l++)
1 i 3-->         c[i][j]+=a[i][l]*b[l][j];
1 i---->       }
1------>    }
```

CSCS
Swiss National Supercomputing Centre

X2 Results

	Single image (GB/s)	Two images (GB/s)
Copy	81.25	37.57
Scale	85.63	37.48
Add	57.54	34.95
Triad	60.37	34.95

Vectorization

Local memory copies

Remote memory copies

XT5 Results

	Single image (MB/s)	Two images (MB/s)
Copy	8524.85	3372.67
Scale	8450.93	1.42
Add	8792.65	1.50
Triad	8716.84	1.50

Vectorization

No vectorization

Local memory copies

Remote memory copies—one element at a time

CSCS
Swiss National Supercomputing Centre

Code Rewrite—Reducing Remote Accesses

Original matrix multiply

```
shared [N*P/THREADS] int a[N][P],c[N][M];
shared [M/THREADS] int b[P][M];
[…]
upc_forall (i=0; i<N; i++; &c[i][0]) {
   for (j=0; j<M; j++) {
      c[i][j]=0;
      for (l=0; l<P; l++)
         c[i][j]+=a[i][l]*b[l][j];
   }
}
```

Alternative matrix multiply

```
shared [N*P/THREADS] int a[N][P],c[N][M];
shared [M/THREADS] int b[P][M];
[…]
for(j=0;j<M;j++){
   for(l=0;l<P;l++){
      b_val = b[l][j];
      upc_forall(i=0;i<N;i++;&c[i][0])
         c[i][j]+=a[i][l]*b_val;
   }
}
```

Matrix Multiply Results on XT5

No difference on X2 platform—slowdown for the alternate implementation

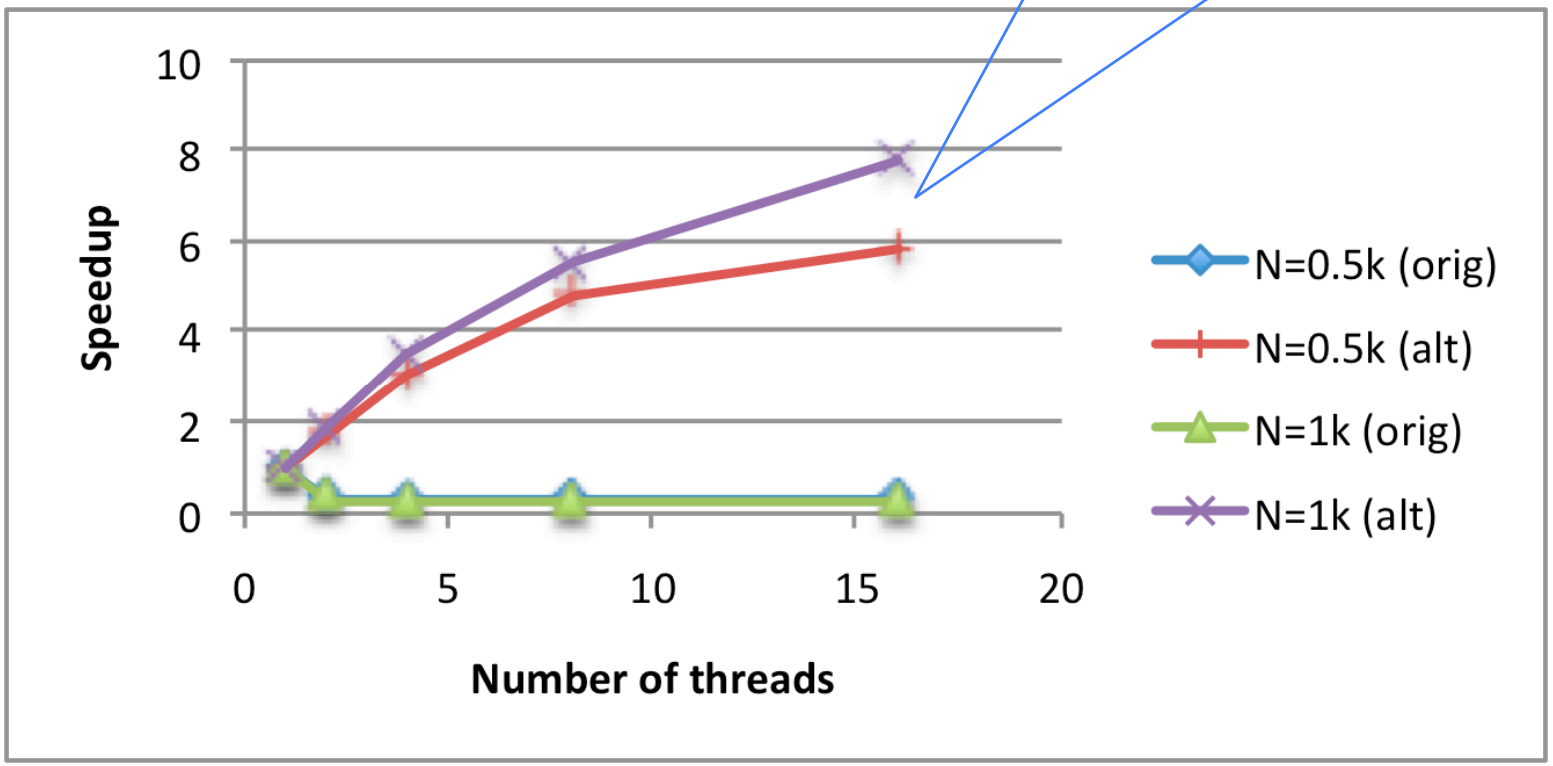

CSCS
Swiss National Supercomputing Centre

Productivity Evaluation

	CAF	UPC
Compiler interface	☺	☺
Runtime control	☺	☺
Debugging tools	😐	☹
Performance tools	☹	☹

Biggest Issue is availability of multi-platform compilers esp. for CAF

Conclusions

- Need to retain uProc level optimization
- Memory and comm. Hierarchy aware runtime
- CCE PGAS compilers for x86 and GASNet supported platforms
- PGAS aware debugging and performance tools

Looking forward to experimenting with GEMINI

CSCS
Swiss National Supercomputing Centre

Acknowledgements

The authors would like to thank Dr Jason Beech-Brandt from the Cray Centre of Excellence for HECToR in the UK for providing access to the X2 nodes of the system. We also appreciate the feedback from Bill Long, Cray for advice on the CAF development of the stencil application.

CSCS
Swiss National Supercomputing Centre

THANK YOU

CSCS
Swiss National Supercomputing Centre