The Intel Parallel Computing Center at the University of Bristol

Simon McIntosh-Smith
Department of Computer Science

University of BRISTOL

UOB HPC

Bristol's rich heritage in HPC

The University of Bristol is one of the top HPC institutes in the UK:

- It has a vibrant HPC community of >500 researchers, >10% of all staff
- Invests over £2.5m p.a. in local HPC
- Trains over 100 HPC computer scientists each year (Bristol CS ranked #4 in UK)

HPC resources in Bristol

- Blue Crystal supercomputer:
  - £12m invested since 2006
  - Amongst the fastest in the UK
  - ~10,000 processor cores
  - ~250 TFLOPS
  - >1 PetaByte of data storage
  - Bristol is a leader in the use of many-core accelerators:
    - Intel Xeon Phi
    - Nvidia, AMD GPUs

University of BRISTOL

UOB HPC

Intel Parallel Computing Center

Intel chose to invest in the University of Bristol to establish its first "Intel Parallel Computing Center (IPCC)" in the UK (Feb 2014):

“The University of Bristol combines both a demonstrated ability to innovate and optimize parallel applications using open, industry-standard techniques with a focus on practical education of the next generation of application developers,”

- Joe Curley, Intel

Bristol IPCC activities

Optimising various HPC codes for Xeon Phi:

- BUDE
  - Molecular docking code
- ROTORSIM
  - Multi-block, multi-grid CFD code
- CloverLeaf/TeaLeaf
  - Hydrodynamics benchmark
- Lattice Boltzmann and more

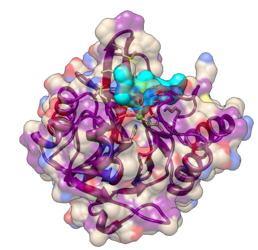

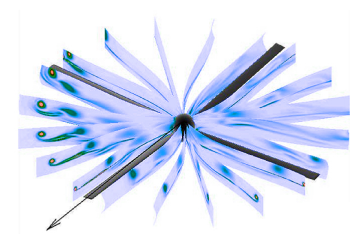

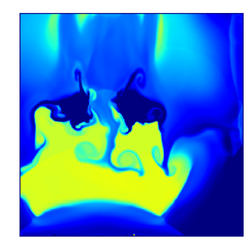

University of BRISTOL

UOB HPC

Molecular Docking in Bristol

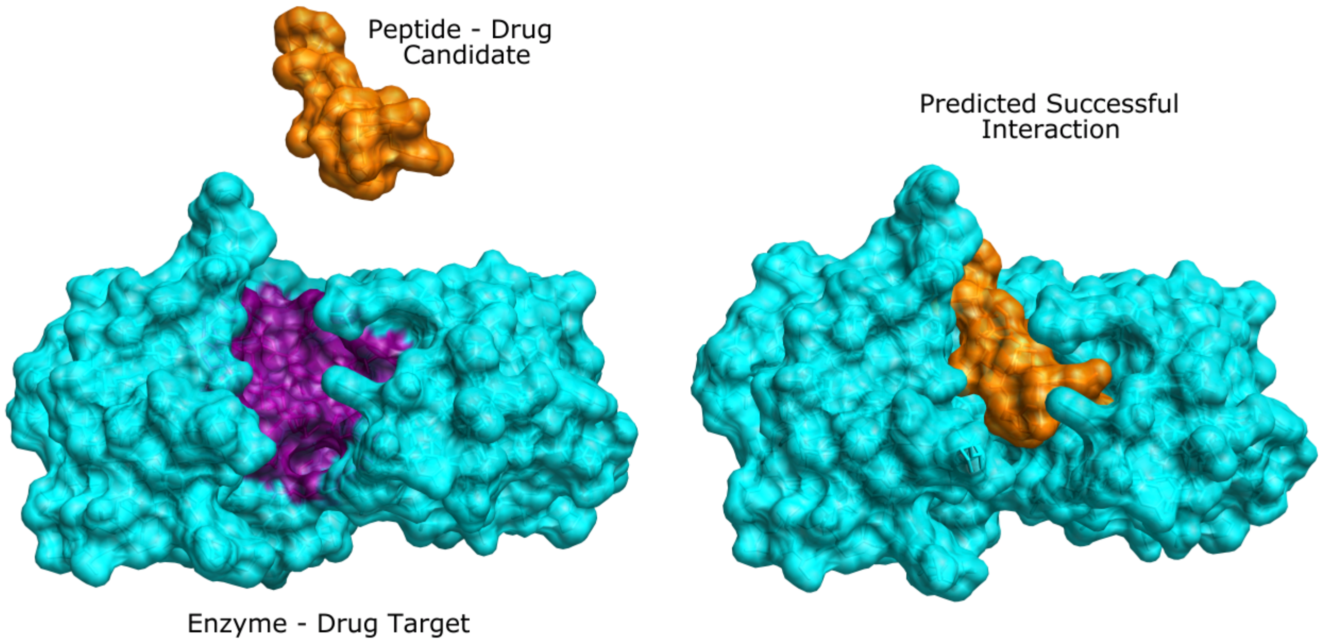

BUDE (Bristol University Docking Engine) is one of the fastest and most accurate molecular docking codes in the world.

BUDE is being used to find new drug targets for influenza, malaria, Alzheimer's, Emphysema, Insulin signalling and more

University of BRISTOL

"High Performance *in silico* Virtual Drug Screening on Many-Core Processors", S. McIntosh-Smith, J. Price, R.B. Sessions, A.A. Ibarra, IJHPCA 2014

UOB HPC

BUDE's algorithm

```
function DOCK(protein, ligand)
    Generate initial population poses at random
    energies = COMPUTE_ENERGIES(protein, ligand, poses)
    for each iteration of EMC do
        Select poses with lowest energies as parents
        Generate new population poses from parents
        energies = COMPUTE_ENERGIES(protein, ligand, poses)
    end for
    Output best poses
end function

function COMPUTE_ENERGIES(protein, ligand, poses)
    for i = 0 upto size(poses) − 1 do
        Transform ligand by poses[i]
        energies[i] = 0
        for each atom l_atom in ligand do
            for each atom p_atom in protein do
                energies[i] = energies[i] + INTERACTION(p_atom, l_atom)
            end for
        end for
    end for
    return energies
end function
```

University of BRISTOL

UOB HPC

BUDE's conditional behaviour

$$E_{complex} = E_{steric} + E_{electrostatic} + E_{desolvation} + E_{config\_entropy}$$

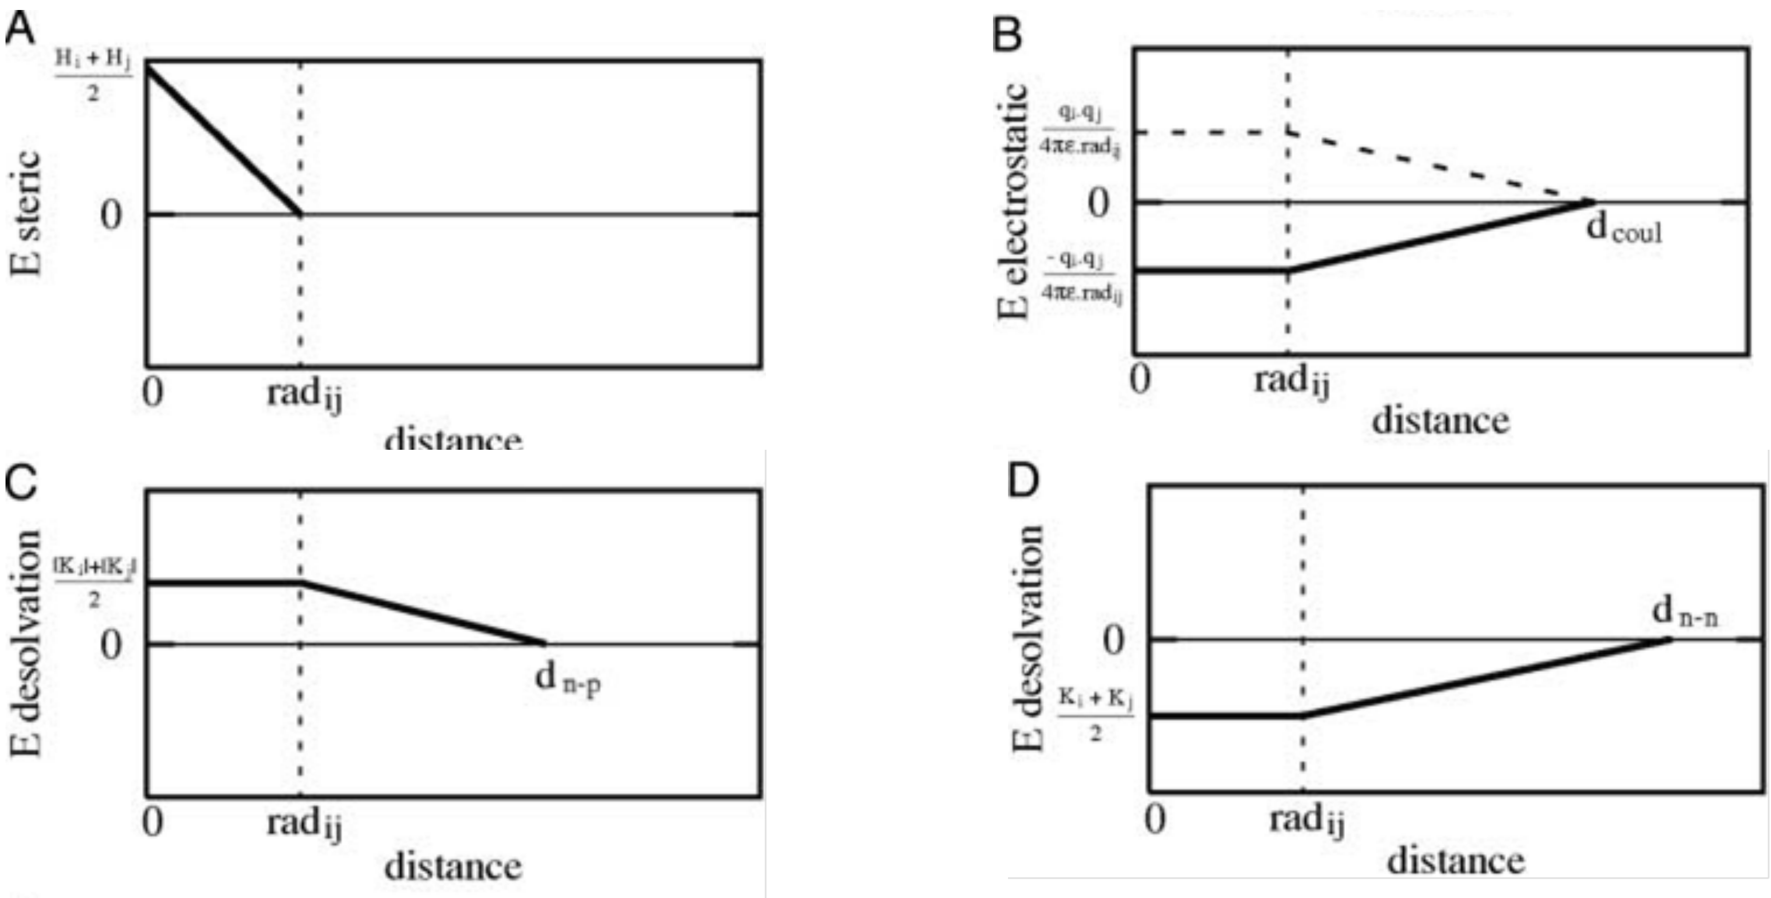

University of BRISTOL

McIntosh-Smith, S., et al., *Benchmarking Energy Efficiency, Power Costs and Carbon Emissions on Heterogeneous Systems.* Computer Journal, 2012. **55**(2): p. 192-205.

UOB HPC

BUDE optimisations for Xeon Phi

```
if (a > b)
{
  accumulator += (a - b*c);
}
```

```
setp.gt.f32 %pred, %a, %b
@!%pred bra $endif
mul.f32 %f0, %b, %c
sub.f32 %f1, %a, %f0
add.f32 %accumulator, %accumulator, %f1
$endif:
```

```
temp = (a - b*c);
mask = (a > b ? 1 : 0)
accumulator += (mask * temp);
```

```
mul.f32 %f0, %b, %c
sub.f32 %temp, %a, %f0
setp.gt.f32 %pred, %a, %b
selp.f32 %mask, %one, %zero, %pred
mad.f32 %accumulator, %mask, %temp, %accumulator
```

University of BRISTOL

UOB HPC

More BUDE optimisations for Phi

- Work-group and NDrange sizes are very important (multiples of 16, 240 etc.)
- Lots of input into Xeon Phi OpenCL driver
  - Have seen a 2-3X improvement in last year
- Vtune on Phi proving very useful

BUDE Xeon Phi results

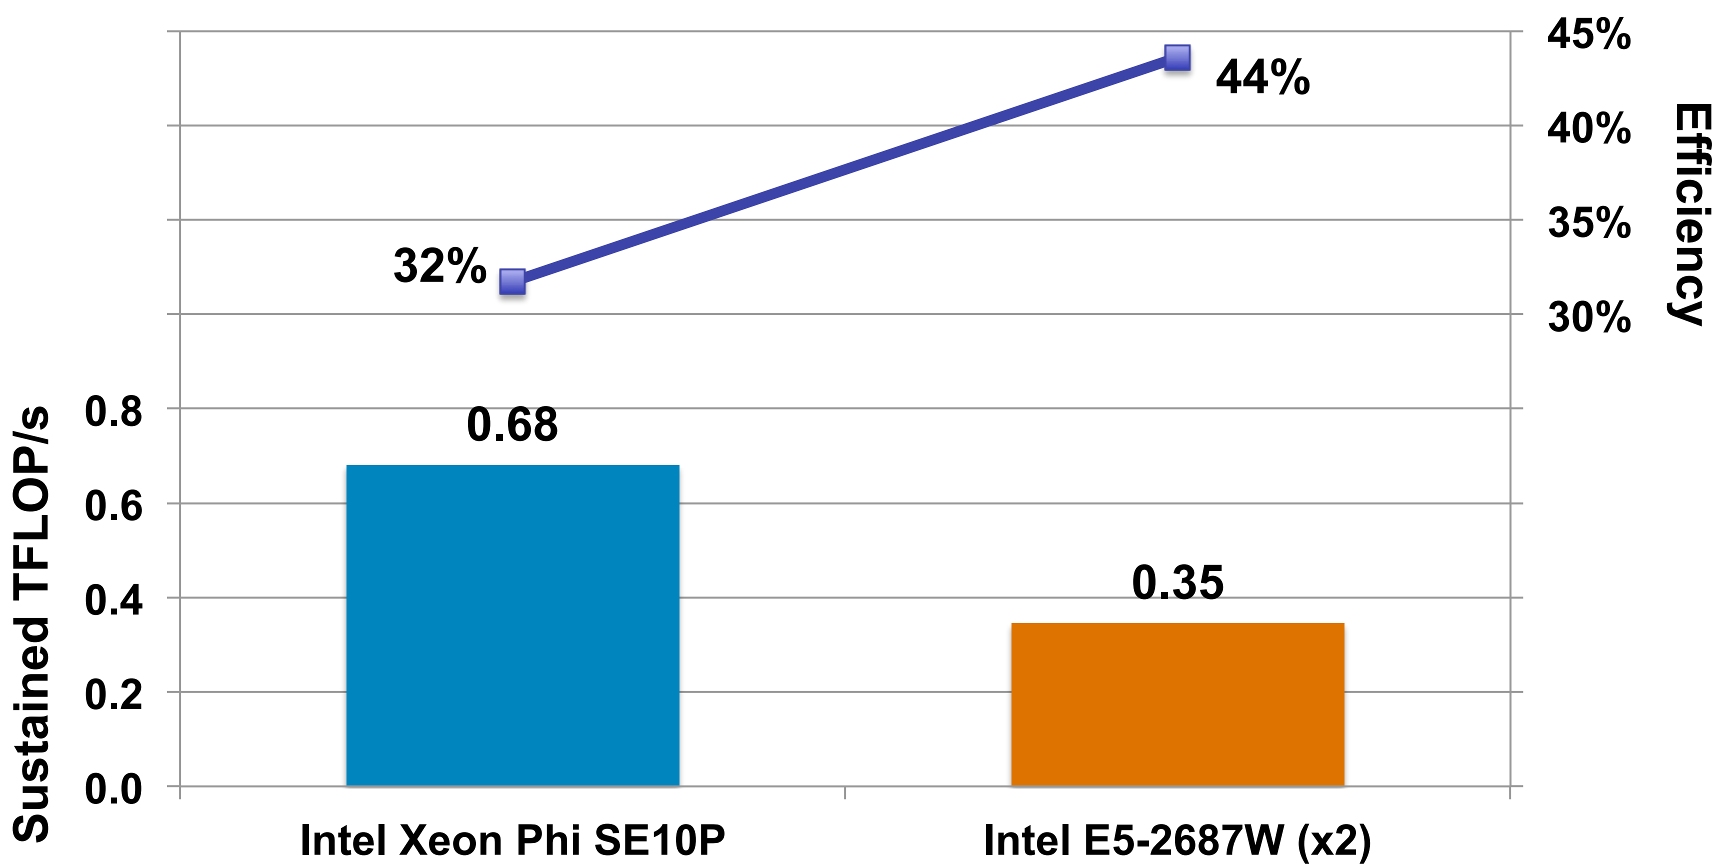

One Xeon Phi SE10P **1.94X** faster than 16 cores of Sandy Bridge at 3.1GHz

University of BRISTOL

UOB HPC

Cloverleaf: A Lagrangian- Eulerian hydrodynamics benchmark

- A bandwidth-limited structured grid code that is part of Sandia's "Mantevo" benchmark suite
- Solves the compressible Euler equations, which describe the conservation of energy, mass and momentum in a system.
- These equations are solved on a Cartesian grid in 2D with second-order accuracy, using an explicit finite-volume method.
- Optimised parallel versions exist in OpenMP, MPI, OpenCL, OpenACC, CUDA and Co-Array Fortran

CloverLeaf optimisations for Xeon Phi

- Focused on OpenCL and OpenMP
- Task granularity is crucial
  - Important to get rid of bounds checking
- Memory alignment and access patterns are also very significant
  - Many simultaneous memory streams can cause the TLBs to thrash
- Barrier placements critical
  - Adding barriers can *improve* performance

Vtune can be really useful on Phi

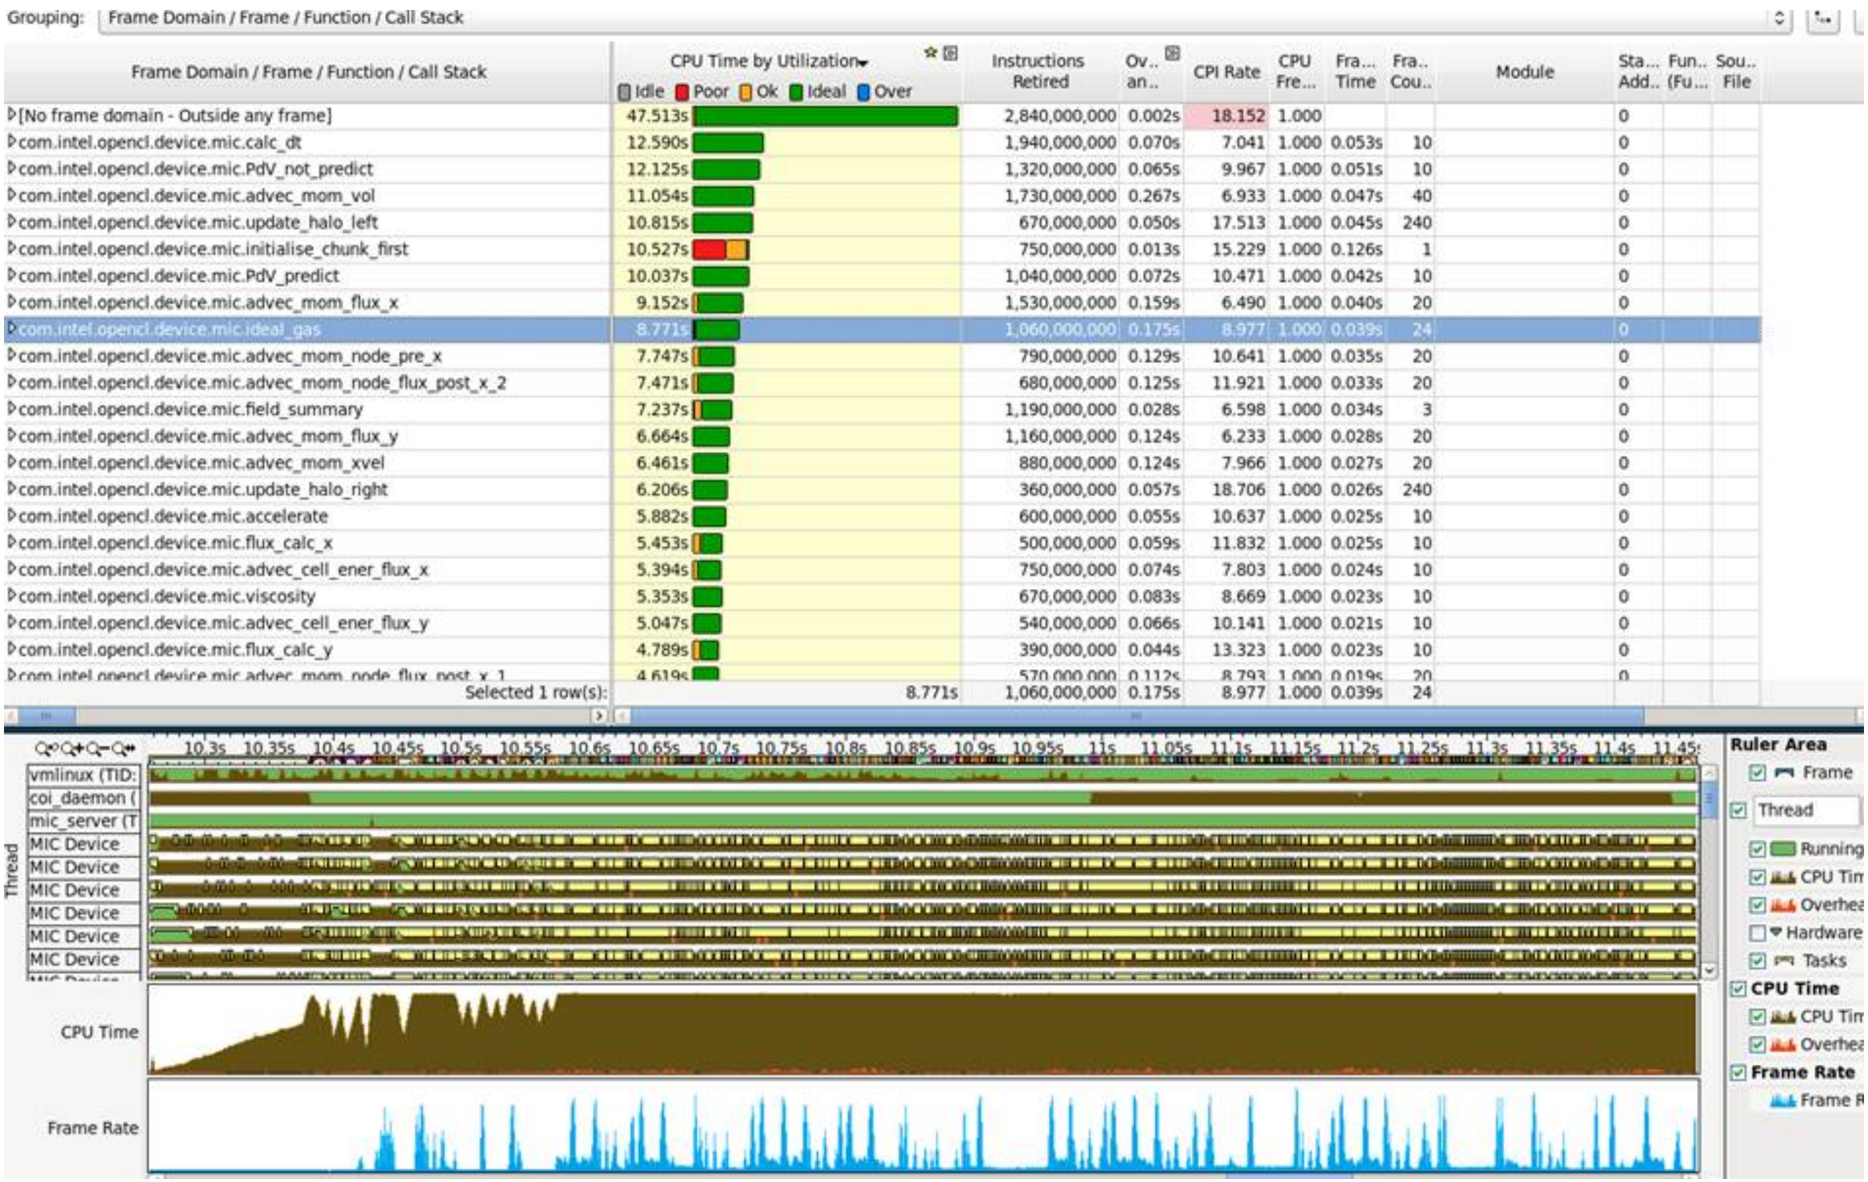

Grouping: Frame Domain / Frame / Function / Call Stack

Frame Domain / Frame / Function / Call Stack	CPU Time by Utilization (Idle, Poor, Ok, Ideal, Over)	Instructions Retired	Ov.. an..	CPI Rate	CPU Fre...	Fra... Time	Fra.. Cou..	Module	Sta... Add..	Fun.. (Fu...	Sou.. File
[No frame domain - Outside any frame]	47.513s	2,840,000,000	0.002s	18.152	1.000				0		
com.intel.opencl.device.mic.calc_dt	12.590s	1,940,000,000	0.070s	7.041	1.000	0.053s	10		0		
com.intel.opencl.device.mic.PdV_not_predict	12.125s	1,320,000,000	0.065s	9.967	1.000	0.051s	10		0		
com.intel.opencl.device.mic.advec_mom_vol	11.054s	1,730,000,000	0.267s	6.933	1.000	0.047s	40		0		
com.intel.opencl.device.mic.update_halo_left	10.815s	670,000,000	0.050s	17.513	1.000	0.045s	240		0		
com.intel.opencl.device.mic.initialise_chunk_first	10.527s	750,000,000	0.013s	15.229	1.000	0.126s	1		0		
com.intel.opencl.device.mic.PdV_predict	10.037s	1,040,000,000	0.072s	10.471	1.000	0.042s	10		0		
com.intel.opencl.device.mic.advec_mom_flux_x	9.152s	1,530,000,000	0.159s	6.490	1.000	0.040s	20		0		
com.intel.opencl.device.mic.ideal_gas	8.771s	1,060,000,000	0.175s	8.977	1.000	0.039s	24		0		
com.intel.opencl.device.mic.advec_mom_node_pre_x	7.747s	790,000,000	0.129s	10.641	1.000	0.035s	20		0		
com.intel.opencl.device.mic.advec_mom_node_flux_post_x_2	7.471s	680,000,000	0.125s	11.921	1.000	0.033s	20		0		
com.intel.opencl.device.mic.field_summary	7.237s	1,190,000,000	0.028s	6.598	1.000	0.034s	3		0		
com.intel.opencl.device.mic.advec_mom_flux_y	6.664s	1,160,000,000	0.124s	6.233	1.000	0.028s	20		0		
com.intel.opencl.device.mic.advec_mom_xvel	6.461s	880,000,000	0.124s	7.966	1.000	0.027s	20		0		
com.intel.opencl.device.mic.update_halo_right	6.206s	360,000,000	0.057s	18.706	1.000	0.026s	240		0		
com.intel.opencl.device.mic.accelerate	5.882s	600,000,000	0.055s	10.637	1.000	0.025s	10		0		
com.intel.opencl.device.mic.flux_calc_x	5.453s	500,000,000	0.059s	11.832	1.000	0.025s	10		0		
com.intel.opencl.device.mic.advec_cell_ener_flux_x	5.394s	750,000,000	0.074s	7.803	1.000	0.024s	10		0		
com.intel.opencl.device.mic.viscosity	5.353s	670,000,000	0.083s	8.669	1.000	0.023s	10		0		
com.intel.opencl.device.mic.advec_cell_ener_flux_y	5.047s	540,000,000	0.066s	10.141	1.000	0.021s	10		0		
com.intel.opencl.device.mic.flux_calc_y	4.789s	390,000,000	0.044s	13.323	1.000	0.023s	10		0		
com.intel.opencl.device.mic.advec_mom_node_flux_post_x_1	4.619s	570,000,000	0.112s	8.793	1.000	0.019s	20		0		
Selected 1 row(s):	8.771s	1,060,000,000	0.175s	8.977	1.000	0.039s	24				

CloverLeaf Xeon Phi results

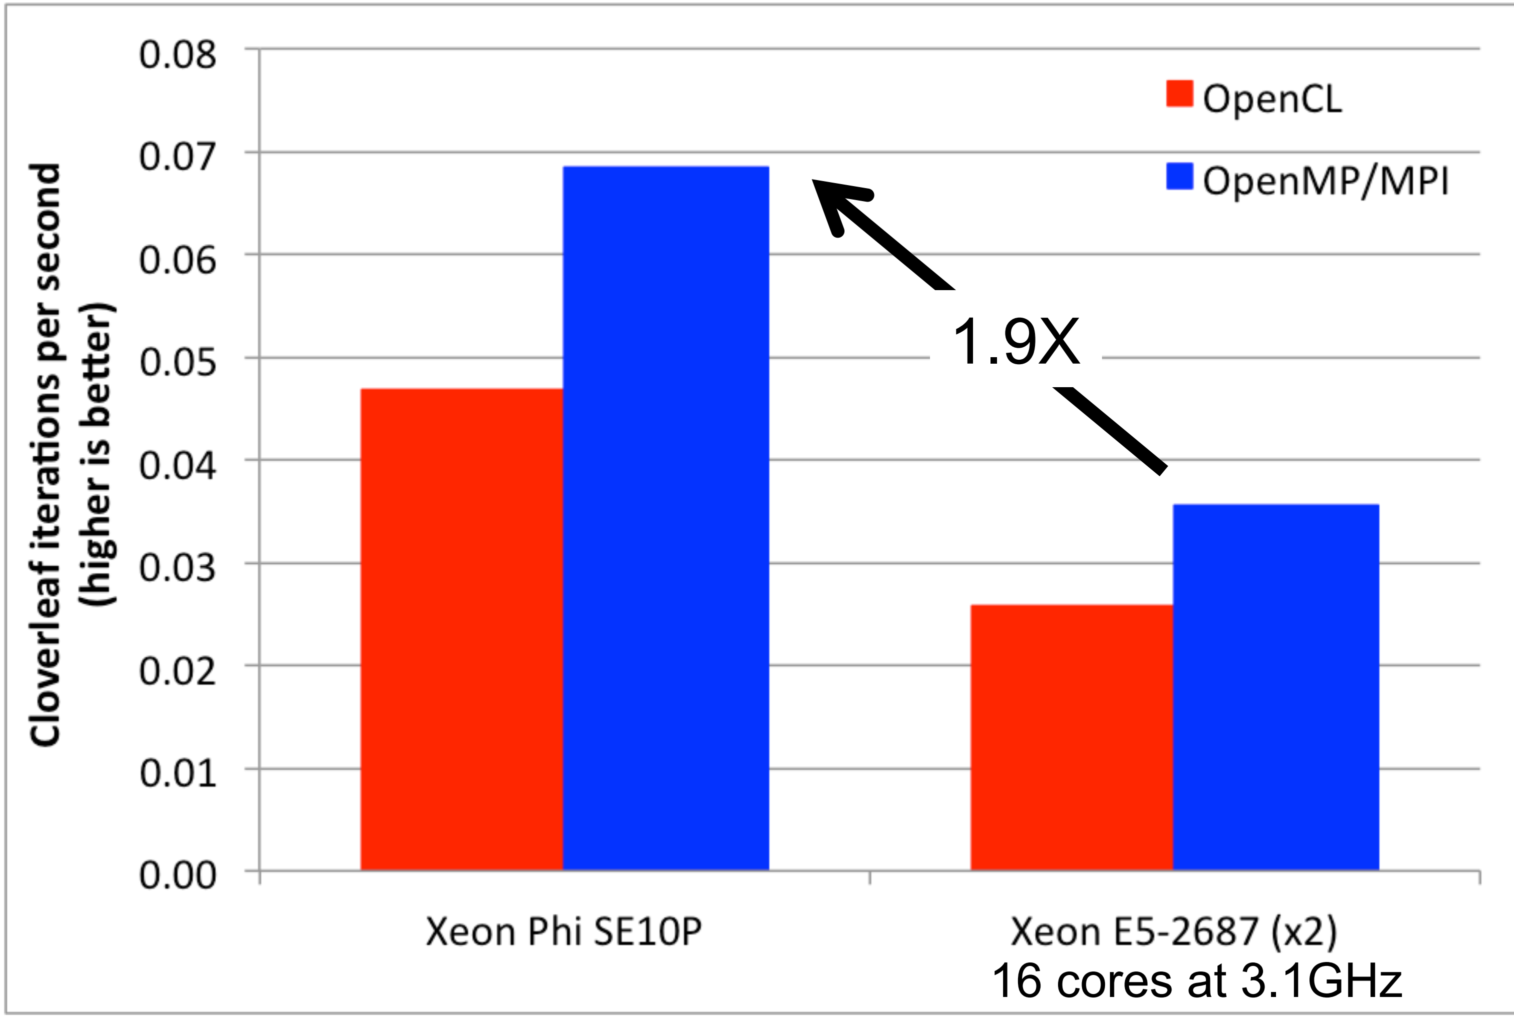

University of BRISTOL

S.N. McIntosh-Smith, M. Boulton, D. Curran, & J.R. Price, "On the performance portability of structured grid codes on many-core computer architectures", To appear, International Supercomputing, Leipzig, June 2014.

UOB HPC

Summary

- Bristol is a leader in exploiting many-core architectures to deliver cutting-edge HPC
- Xeon Phi can deliver acceleration of 1.5-2.0X for real HPC codes
- There's a very worrying trend that many HPC codes are not evolving fast enough to be ready for the many-core trend
  - Implication: these codes will fail to get good performance on Xeon, never mind Xeon Phi!

OpenCL conference in Bristol

IWOCL 2014
INTERNATIONAL WORKSHOP ON OPENCL
OFFICIAL DATE & VENUE:
12-13 MAY, 2014
BRISTOL UNIVERSITY, ENGLAND

http://iwocl.org

- IWOCL ("eye-wok-ul")
- May 12-13th 2014
- Bristol, UK
- In an award-winning science museum
- 2 days of technical talks and workshops

University of BRISTOL

UOB HPC