Development of strategy and procedure for assessing portability of industrial codes to GPUs hardware: the TITAN example

Principal Investigator: Dr Luigi Capone

(Physical modelling group, CFD methods, Rolls Royce plc, UK)

Key Contributors: Dr Istvan Reguly

(Computer science department, University of Oxford, UK)

Matthew Street

(Computational sciences group, Rolls Royce plc, UK)

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Rolls-Royce

© 2016 Rolls Royce plc

Presentation Outline

- ARD100 project scope overview
- Current parallel implementation in Hydra
- OP2 and OPS libraries
- Hydra mini-application
- OP2 and OPS further extension and development
- Test cases
- Results

© 2016 Rolls Royce plc

Rolls-Royce

ARD100 project scope and overview

- The main objective for the present director discretionary allocation project was to develop an industrial strategy and procedure for testing portability to GPUs based hardware of the in-house Rolls Royce CFD code

- Dissemination of the GPUs technology in industry is based on an accurate evaluation of cost vs benefit

- The aim of the project was developing a careful procedure for assessing the potential benefit of the new computing architecture as well as evaluating the technical complexity, timescales and costs of the process itself

- This project provided a preliminary platform for developing a common framework for general studies on parallel optimization on new computing architectures for Rolls Royce

© 2016 Rolls Royce plc

Rolls-Royce

Current Parallel Strategy in Hydra

- Hydra is a compressible Navier-Stokes equation solver used for turbomachinery applications
- Currently the parallelization is based on MPI protocol
- Scaling analysis of the current Hydra parallel implementation was developed on TITAN as starting point
- Parallel operations are handled by the library OPlus originally developed by University of Oxford
- OPlus was implemented with the aim of systematize the coding of parallel operation and release the developer from direct knowledge of complex MPI operations

© 2016 Rolls Royce plc

Rolls-Royce

OPlus (1)

Mapping of Nodes vs Faces

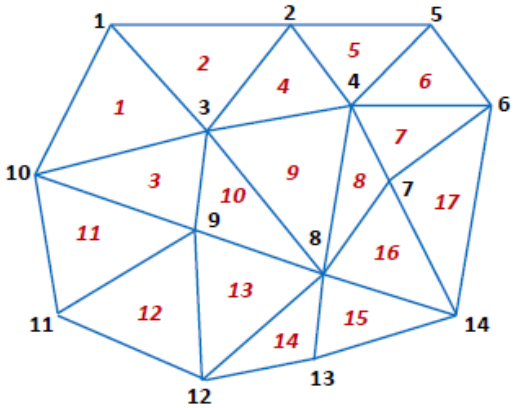

(1)

- There are very minor differences between a do loop in serial Fortran and in parallel OPlus
- One of the key aspect is the handling of the computational geometry in a purely unstructured fashion: MAPPING
- OPlus splits the base geometrical information in SETS of data
- Each of the SETS can be characterized by different ATTRIBUTES quantities

Standard Fortran loop commands

FORTRAN loop

```
DO IC = 1, NCELLS
  I1 = NCELL(1,IC)
  I2 = NCELL(2,IC)
  I3 = NCELL(3,IC)
  AREAN(I1) = AREAN(I1) + AREAC(IC)/3.0
  AREAN(I2) = AREAN(I2) + AREAC(IC)/3.0
  AREAN(I3) = AREAN(I3) + AREAC(IC)/3.0
ENDDO
```

OPlus Fortran loop commands

FORTRAN OPlus loop

```
DO WHILE(OP_PAR_LOOP(NCELLS,ISTART,IFINISH))
  CALL OP_ACCESS_R8('read'  ,AREAC,1,NCELLS,NULL ,0,0,1,1)
  CALL OP_ACCESS_R8('update',AREAN,1,NNODES,NCELL,1,1,1,3)
  DO IC = ISTART, IFINISH
    I1 = NCELL(1,IC)
    I2 = NCELL(2,IC)
    I3 = NCELL(3,IC)
    AREAN(I1) = AREAN(I1) + AREAC(IC)/3.0
    AREAN(I2) = AREAN(I2) + AREAC(IC)/3.0
    AREAN(I3) = AREAN(I3) + AREAC(IC)/3.0
  ENDDO
END WHILE
```

(2)

© 2016 Rolls Royce plc

Rolls-Royce

OPlus (2)

- Different SETS ATTRIBUTES may be interconnected to other SETS through POINTERS
- Any operation can be then described as a FUNCTION on the members of a set or on the connected members of a different SETS using a POINTER
- OPlus handles partitioning, halo cells construction and local parallel sets renumbering
- The main underlying hypothesis is that SETS operations need to be order independent (Gauss-Siedel relaxation or global implicit methods not applicable)

© 2016 Rolls Royce plc

Rolls-Royce

New concepts OP2 and OPS (1)

- Starting from the same background two new parallel libraries were created
- OP2 expands the OPlus approach on unstructured meshes creating source to source translation layer that can adapt the application to a wide range of hardware
- Once the application code has been rewritten using the OP2 API, it can be translated to CUDA, OpenMP or MPI
- OP2 initially did not have OpenACC capability
- The unstructured capability is maintained by the SETS mapping strategy that was already implemented in OPlus

Rolls-Royce

© 2016 Rolls Royce plc

New concepts OP2 and OPS (2)

- OPS is an API library source to source translator for block-structured meshes
- The Oplus/OP2 mapping abstraction is deprecated as the connectivity for structured SETS is implicit depending on the numerical method and order of accuracy of the stencil
- OPS provide source code translation to CUDA, OpenMP, OpenCL, MPI and OpenACC
- OPS initially did not have multigrid capability

Unstructured

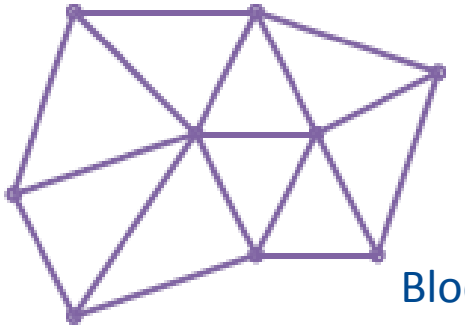

Block-structured

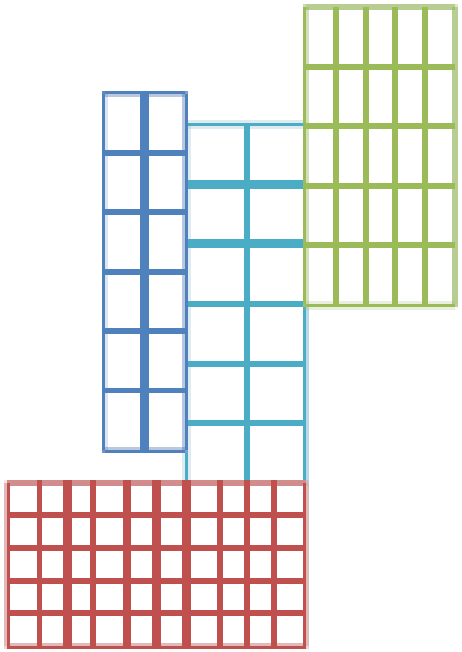

(3)

© 2016 Rolls Royce plc

Rolls-Royce

Geometric Multigrid

- Improve convergence and robustness
- Geometric multigrid requires creation of coarser grid levels through edge collapsing technique
- Restriction operator transfer to coarser levels the solution
- Prolongation operator maps the solution back to the finer levels
- Different error frequencies are damped by the smoother at different grid levels

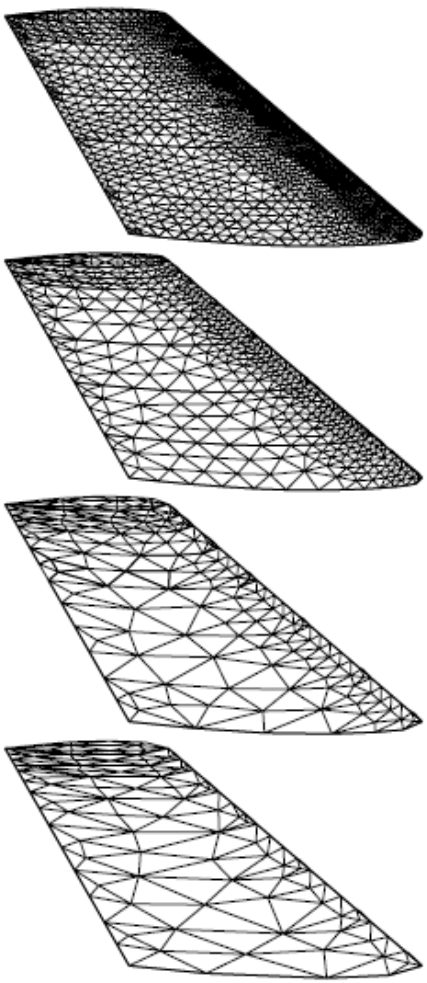

(2)

Rolls-Royce

© 2016 Rolls Royce plc

Hydra OPlus MPI TITAN scaling

- Two test cases 150M and 300M
- Comparison between TITAN and ARCHER supercomputing facility UK
- The main outcome is the limitation in terms of CPU and network performance if TITAN is used inefficiently with MPI

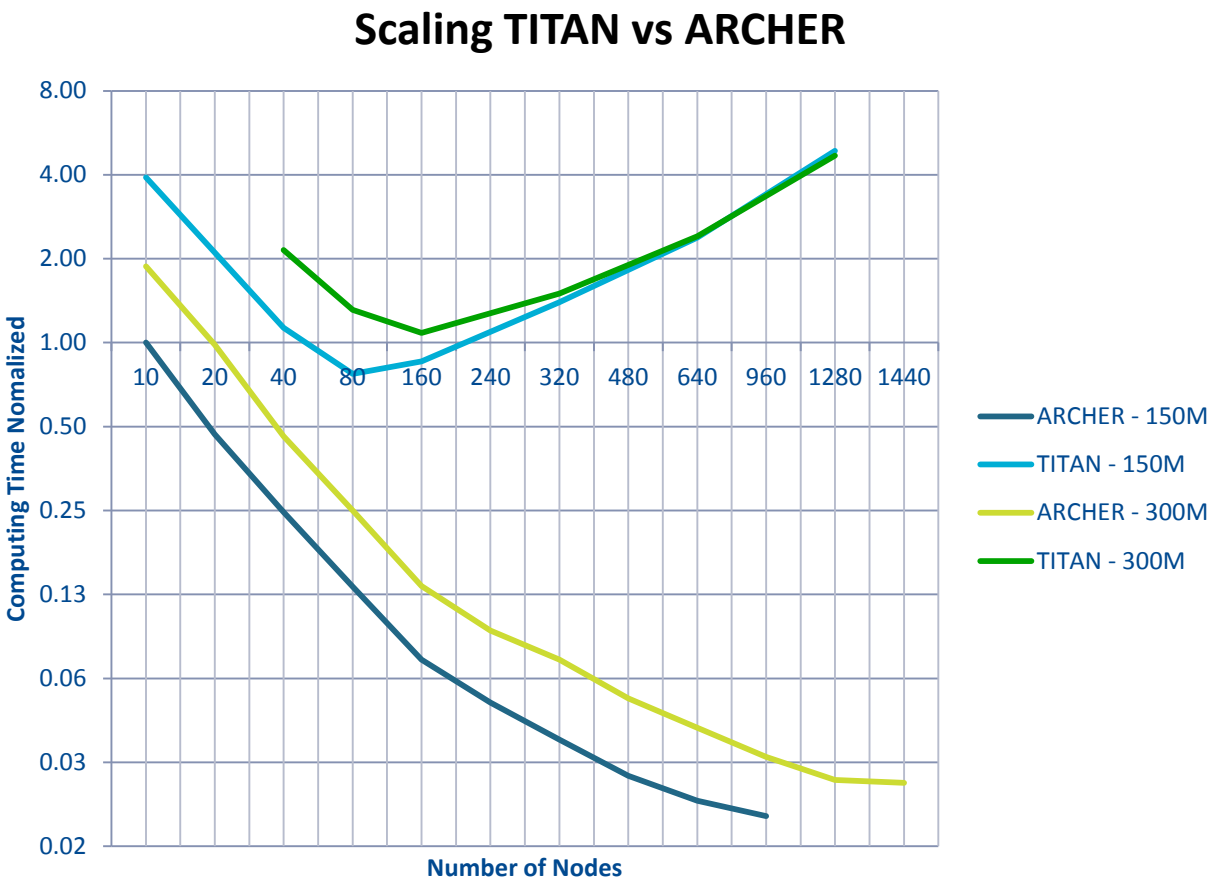

© 2016 Rolls Royce plc

Rolls-Royce

Hydra Mini-App

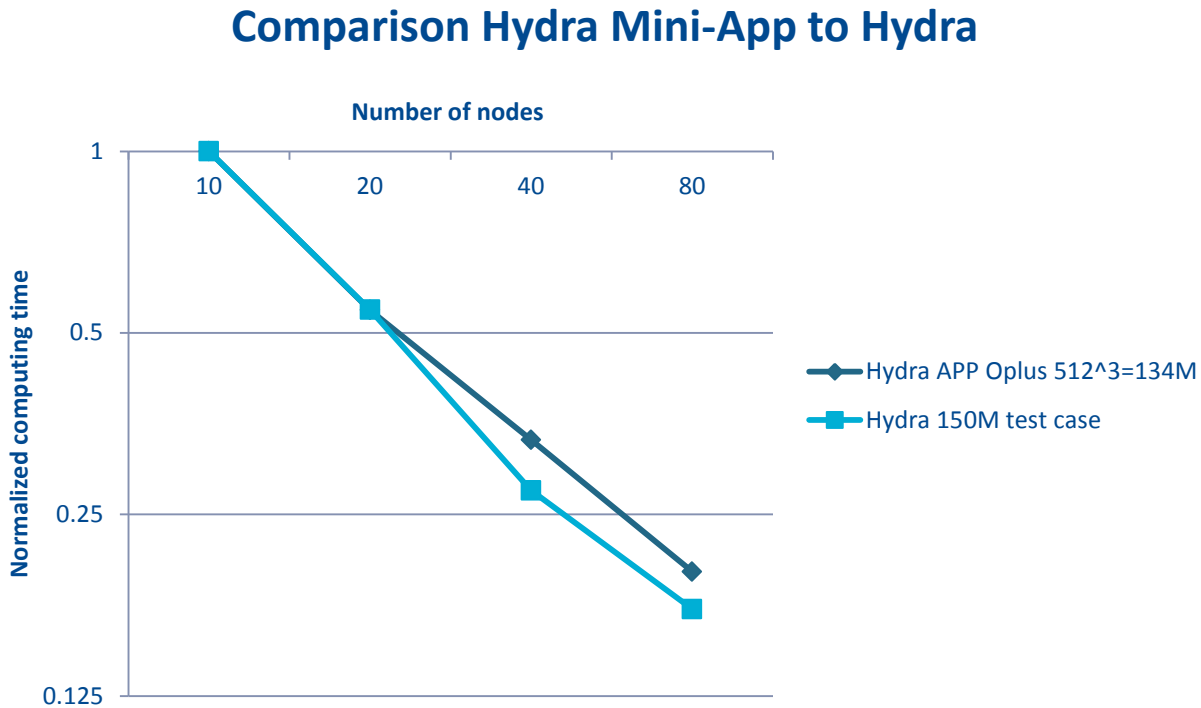


```
[lcapone@login02 src]$ wc -l *
   24 error.cpp
    8 error.h
   22 interp.cpp
   14 interp.h
  208 mg.cpp
   32 neighbours.cpp
   11 neighbours.h
   23 params.h
   21 setInitialGuess.cpp
   11 setInitialGuess.h
   45 smooth.cpp
   18 smooth.h
   39 transf.cpp
   14 transf.h
   10 update.cpp
   16 update.h
   48 utils.h
  564 total
```

- For overcoming the complexity of handling a complex code as Hydra a mini-app was created of only 564 lines in C++, original code order of 10^5 lines
- The idea of mini-apps for parallel performance was initially studied at Sandia in 2009
- The main feature to be reproduced was the geometric multigrid
- The mini app was then built using OPlus, OP2 and OPS API

Rolls-Royce

© 2016 Rolls Royce plc

Test Cases and Development od OP2 and OPS

- The mini-app solves a 3D cube structured mesh of 128, 256 and 512 elements per dimensions
- The OPlus and OP2 mapping is still full unstructured even if the mesh is structured
- OPlus mini-app version demonstrated similar performance to Hydra
- OP2 was further developed for this project for embedding OpenACC source to source translation capability
- Geometric multigrid algorithm was also built in OPS
- Obtaining versions of OPlus, OP2 and OPS with a the same set of core capability

Rolls-Royce

© 2016 Rolls Royce plc

Mini-App OPlus vs OP2 Scaling

128^3 Mesh OP2 vs OPlus

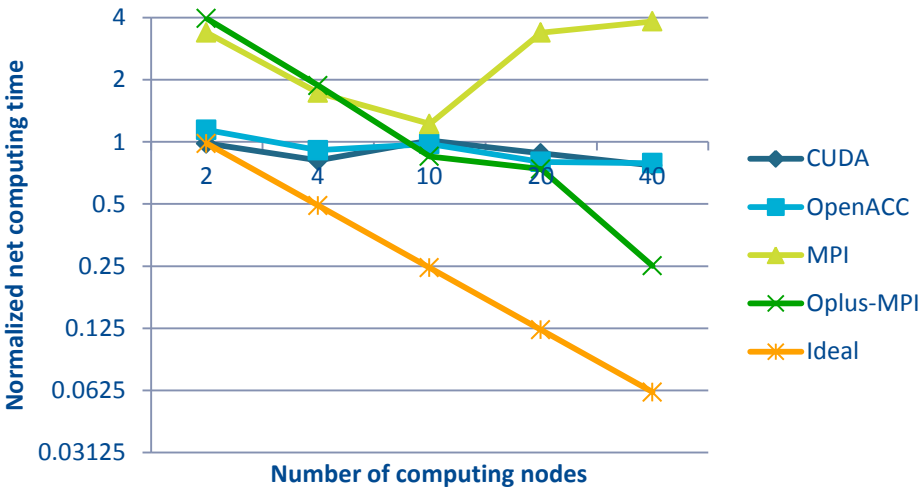

256^3 Mesh OP2 vs OPlus

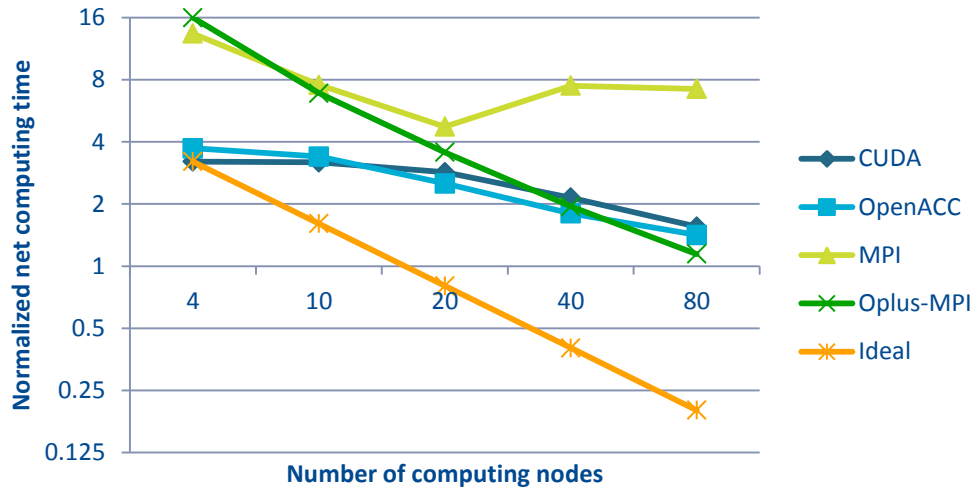

512^3 Mesh OP2 vs OPlus

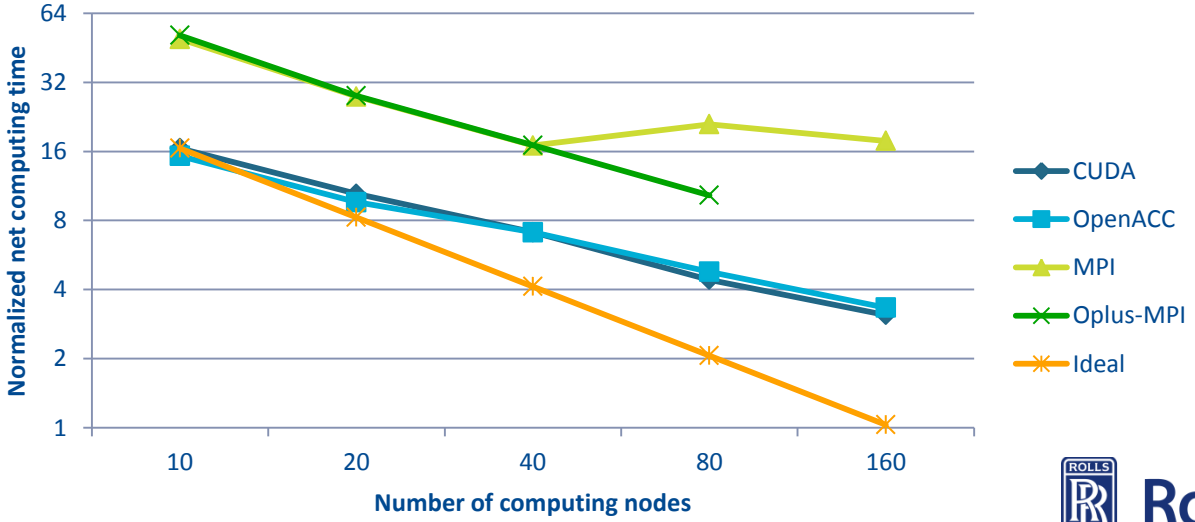

© 2016 Rolls Royce plc

Rolls-Royce

Mini-App OPS vs OPlus Scaling

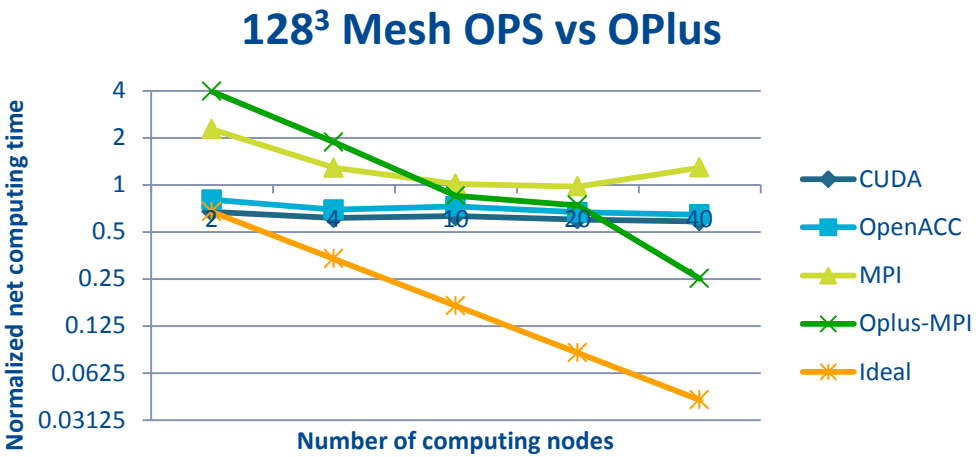

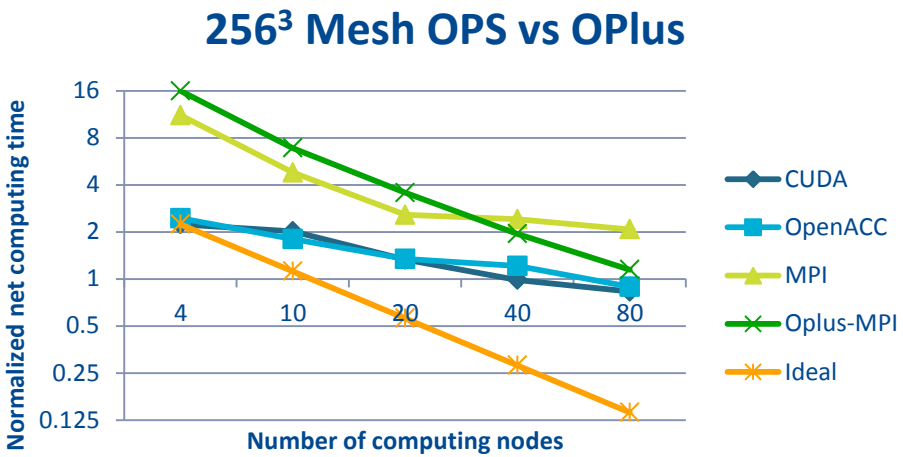

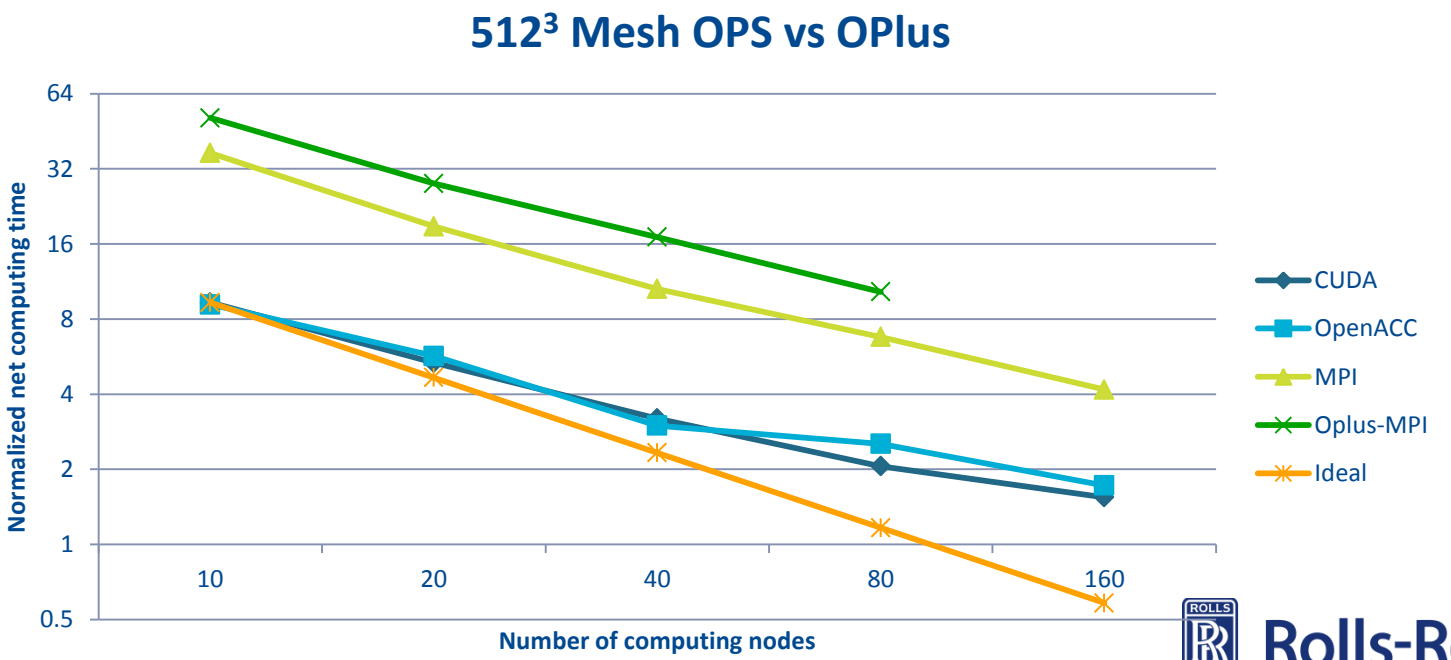

© 2016 Rolls Royce plc

Rolls-Royce

Overall Scaling

128^3 Mesh Global Comparison

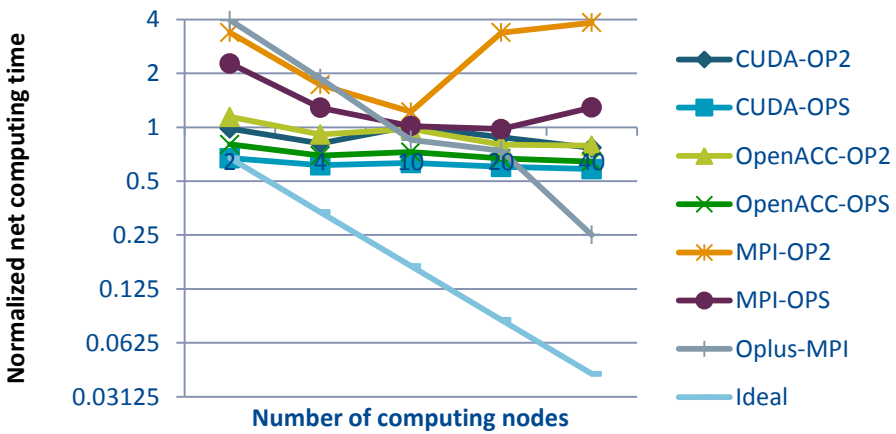

256^3 Mesh Global Comparison

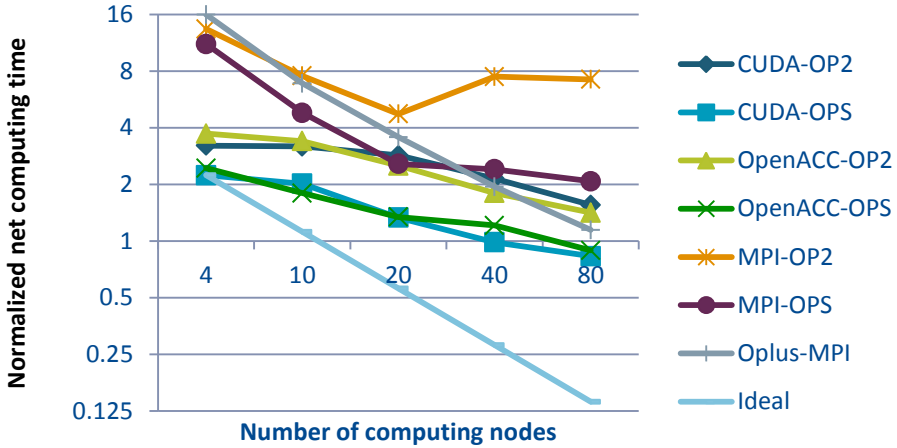

512^3 Mesh Global Comparison

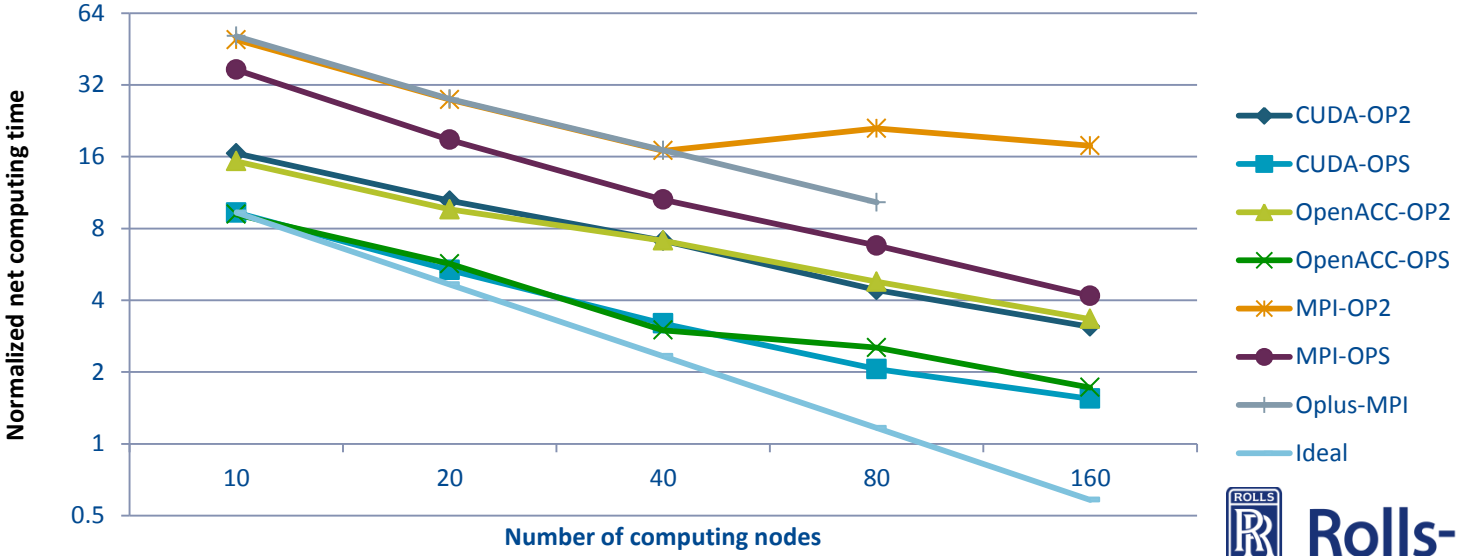

Rolls-Royce

© 2016 Rolls Royce plc

Weak Scaling

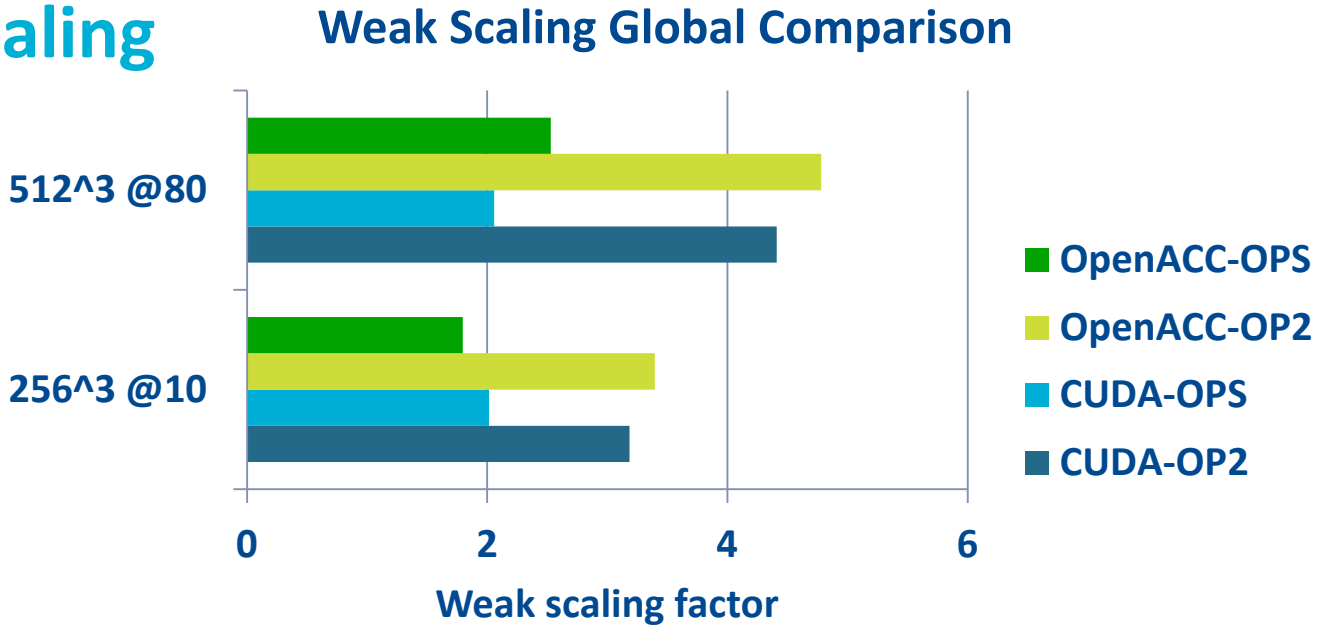

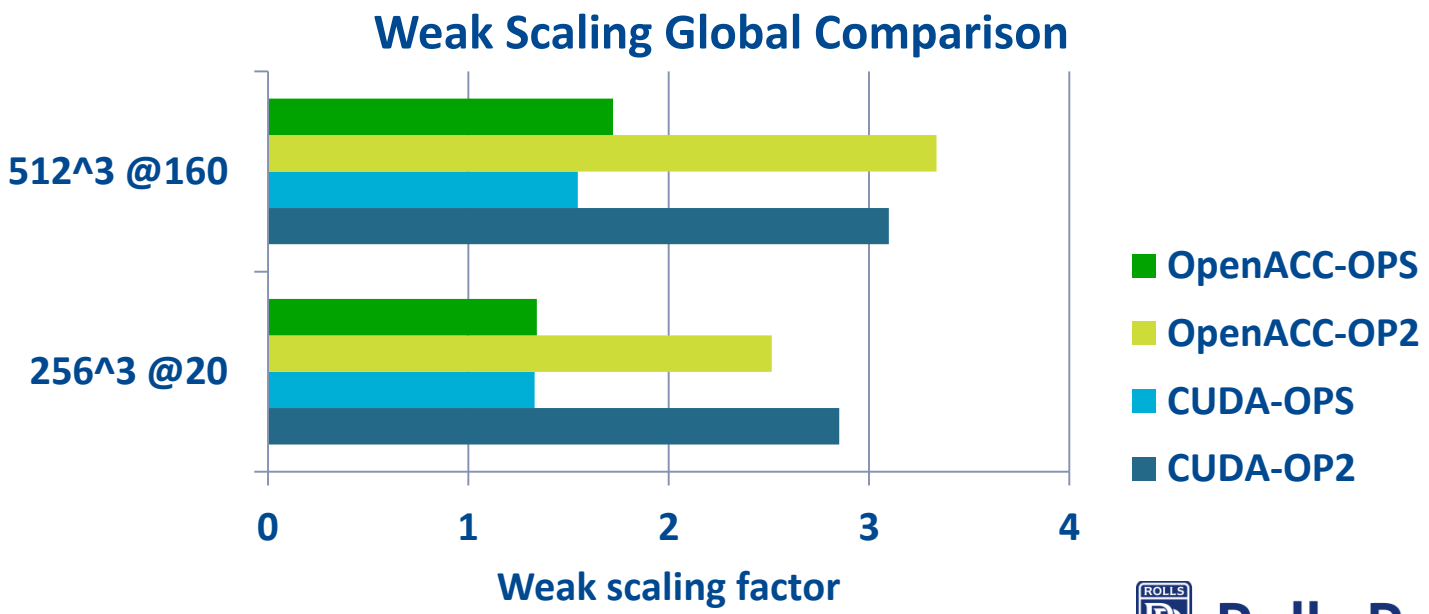

© 2016 Rolls Royce plc

Rolls-Royce

Speed Up

Speed up factors for 128^3 Mesh

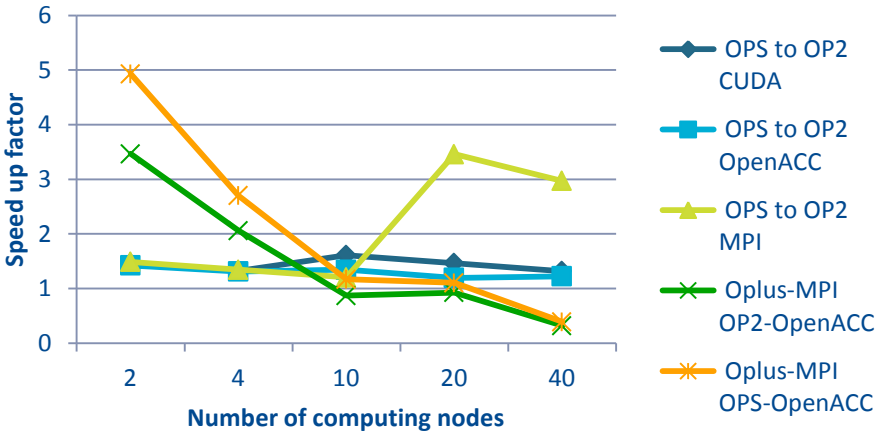

Speed up factors for 256^3 Mesh

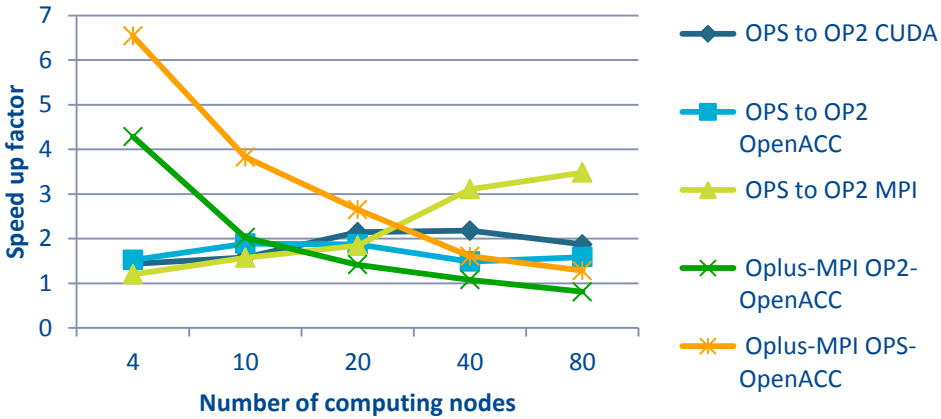

Speed up factors for 512^3 Mesh

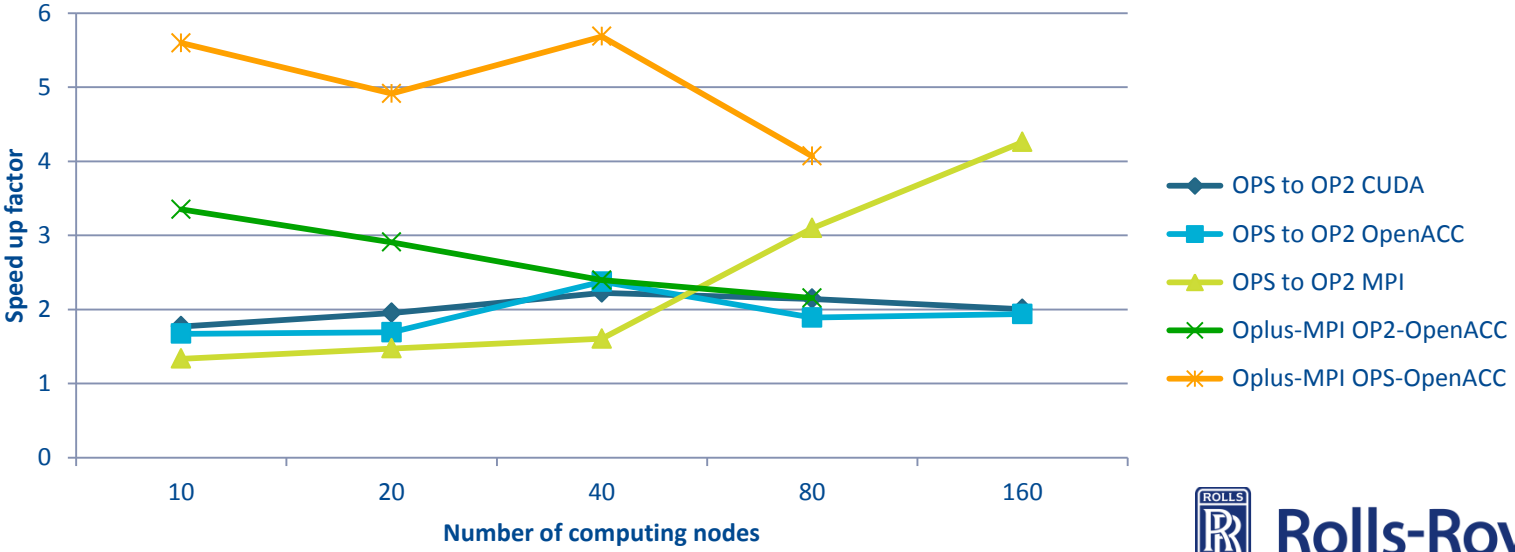

© 2016 Rolls Royce plc

Rolls-Royce

Conclusions

- A general framework was created for consistent comparison between the OPlus parallel library and the newer OP2 and OPS
- Structured vs unstructured meshes were evaluated for GPUs porting
- An Hydra mini app was created for flexible and general validation and parallel performance optimization
- An extensive evaluation of the OpenACC was achieved
- Results demonstrates how the OpenACC strategy could be a valid alternative for porting efficiently and cost effectively existing source codes to new hardware, but further investigation is needed.
- Complete test of the TITAN structure and capabilities was performed opening the opportunity for a future INCITE program proposal for Rolls Royce

Rolls-Royce

© 2016 Rolls Royce plc

Acknowledgements

- This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

- The project was possible also thanks to the valuable support of J. Wells, S. Tichenor, A. Barker, S. Parete-Koon, F. Foertter and all the OLCF users support team.

© 2016 Rolls Royce plc

Rolls-Royce

References

1) I. Reguly, G. R. Mudalige, C. Bertolli, M. B. Giles, A. Betts, P. H. J. Kelly and D. Radford, “Acceleration of a full-scale industrial CFD application with OP2”, ACM transaction on parallel computing, October 2013.
2) P. I. Crumpton and M. B. Giles, “Aircraft computations using multigrid and an unstructured parallel library”, AIAA meeting, January 1995, AIAA 95-0210.
3) G. R. Mudalige, I. Z. Reguly, M. B. Giles, W. Gaudin, A. Mallinson and O. Perks, ”high-level abstraction for performance, portability and continuity of scientific software on future computing systems”, On-line report.

© 2016 Rolls Royce plc

Rolls-Royce