

HPC for all

Porting *in silico* medicine applications to an HPC environment

Alessandro Melis

a.melis@sheffield.ac.uk

Department of Mechanical Engineering

The University of Sheffield

INSIGNEO Institute for *in silico* Medicine

HPC@Sheffield

Monday 26 March 2018

09.00 - 16.00

Firth Hall

Sheffield Teaching Hospitals 
NHS Foundation Trust

INSIGNEO
Institute for *in silico* Medicine



The
University
Of
Sheffield.

Outline

- HPC global race
- *In silico* medicine
- Case study, BoneDVC
- Benchmark and profiling
- Optimisation
- Lesson learned

Traditional HPC applications

- Dedicated to applications where the computational cost could not be reduced on normal systems
- Carrying enough scientific, social, and/or economic impact to justify the high cost of porting and optimising the codes for HPC (*very expensive computer systems*)
 - Weather forecasting
 - Fluid dynamics
 - Nuclear test simulations
 - Molecular dynamics

The global race

- H2020 HPC:
 - HPC should play a more pervasive role in science
 - All scientific disciplines are becoming computational
 - Industry is increasingly relying on the power of supercomputers to invent innovative solutions
- European HPC strategy:
 - Achieving excellence in HPC application delivery and use
 - Establishment of Centres of Excellence (CoEs)
 - Focus and coordinate support to the application of HPC in scientific or industrial domains that are most important for Europe
 - CompBioMed is the CoE dedicated to *in silico* Medicine
 - <http://www.compbioMed.eu>

In silico medicine

- Direct use of computer simulation in clinical practice for diagnosis, treatment, and prevention of disease in subject-specific cases
- Representation of biological processes in a virtual environment by means of
 - Modelling
 - Simulation
 - Visualisation
- Cardiovascular, Musculoskeletal, Biomolecular dynamics

Traditional HPC vs subject-specific modelling


Traditional HPC applications	Subject-specific modelling
Complex mathematics but simple structure	Complex workflows including multiple codes with specialised tasks
Developed by researchers and normally available in source code format: limited dependencies and no usage restrictions	Commercial software to shorten development cycle
Model cannot be reduced, long execution time on conventional PC, vast memory requirements	Combinatorial explosion in multiscale models
Dominated by fundamental research	Need to interface with clinical and industrial users
Run in batch	Run as soon as possible (urgent computing)

Case study

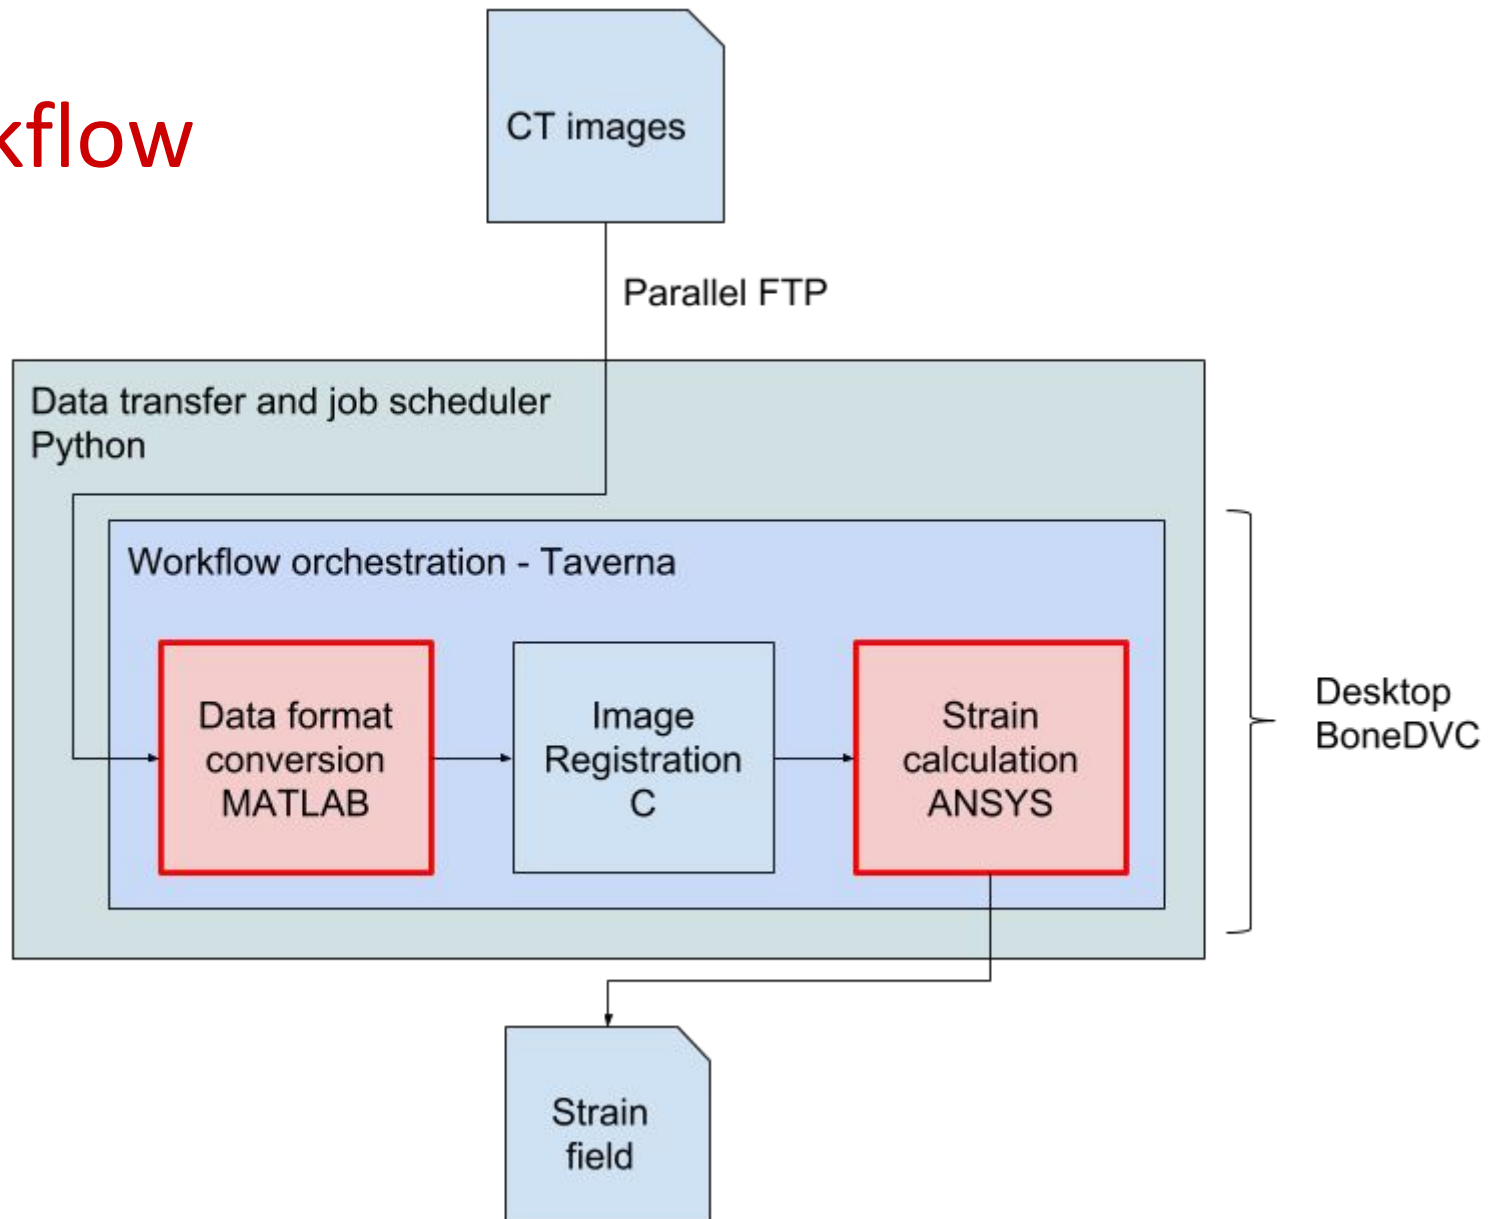
- Porting of typical *in silico* solution from desktop PC to Tier-1 national system (ARCHER EPCC) and Tier-3 local cluster (ShARC USFD)
- Document specific barrier that prevented a straightforward deployment
- Optimisation work needed to exploit HPC capabilities
- Musculoskeletal application, BoneDVC
 - Dall'Ara, E., Barber, D., & Viceconti, M. (2014). About the inevitable compromise between spatial resolution and accuracy of strain measurement for bone tissue: A 3D zero-strain study. *Journal of biomechanics*, 47(12), 2956-2963.
 - Palanca, M., Tozzi, G., Cristofolini, L., Viceconti, M., & Dall'Ara, E. (2015). Three-dimensional local measurements of bone strain and displacement: comparison of three digital volume correlation approaches. *Journal of biomechanical engineering*, 137(7), 071006.
 - Barber, D. C., & Hose, D. R. (2005). Automatic segmentation of medical images using image registration: diagnostic and simulation applications. *Journal of medical engineering & technology*, 29(2), 53-63.
 - Barber, D. C., Oubel, E., Frangi, A. F., & Hose, D. R. (2007). Efficient computational fluid dynamics mesh generation by image registration. *Medical image analysis*, 11(6), 648-662.

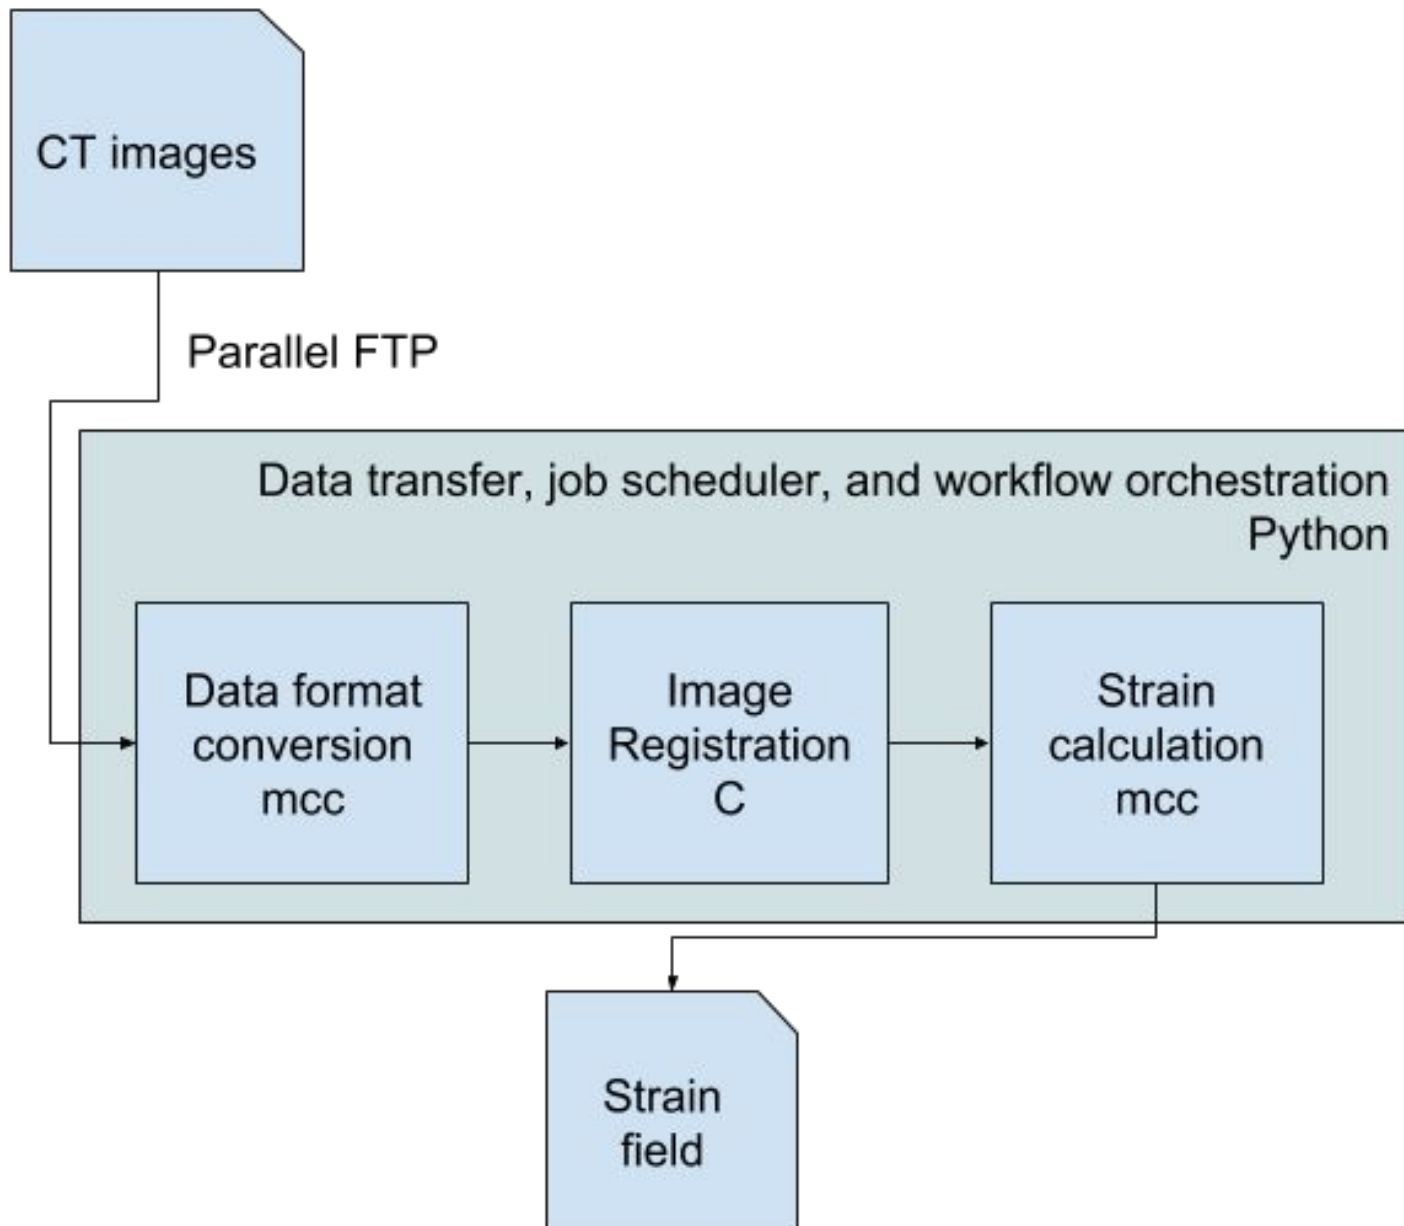
BoneDVC

- Analyse two stacks of CT images of the same bone tissue specimen
- Reference and displaced (axial compression) configurations
- Linear displacement field mapping computed by means of the digital volume correlation (DVC) method
- Strain field obtained by differentiation of the displacement field



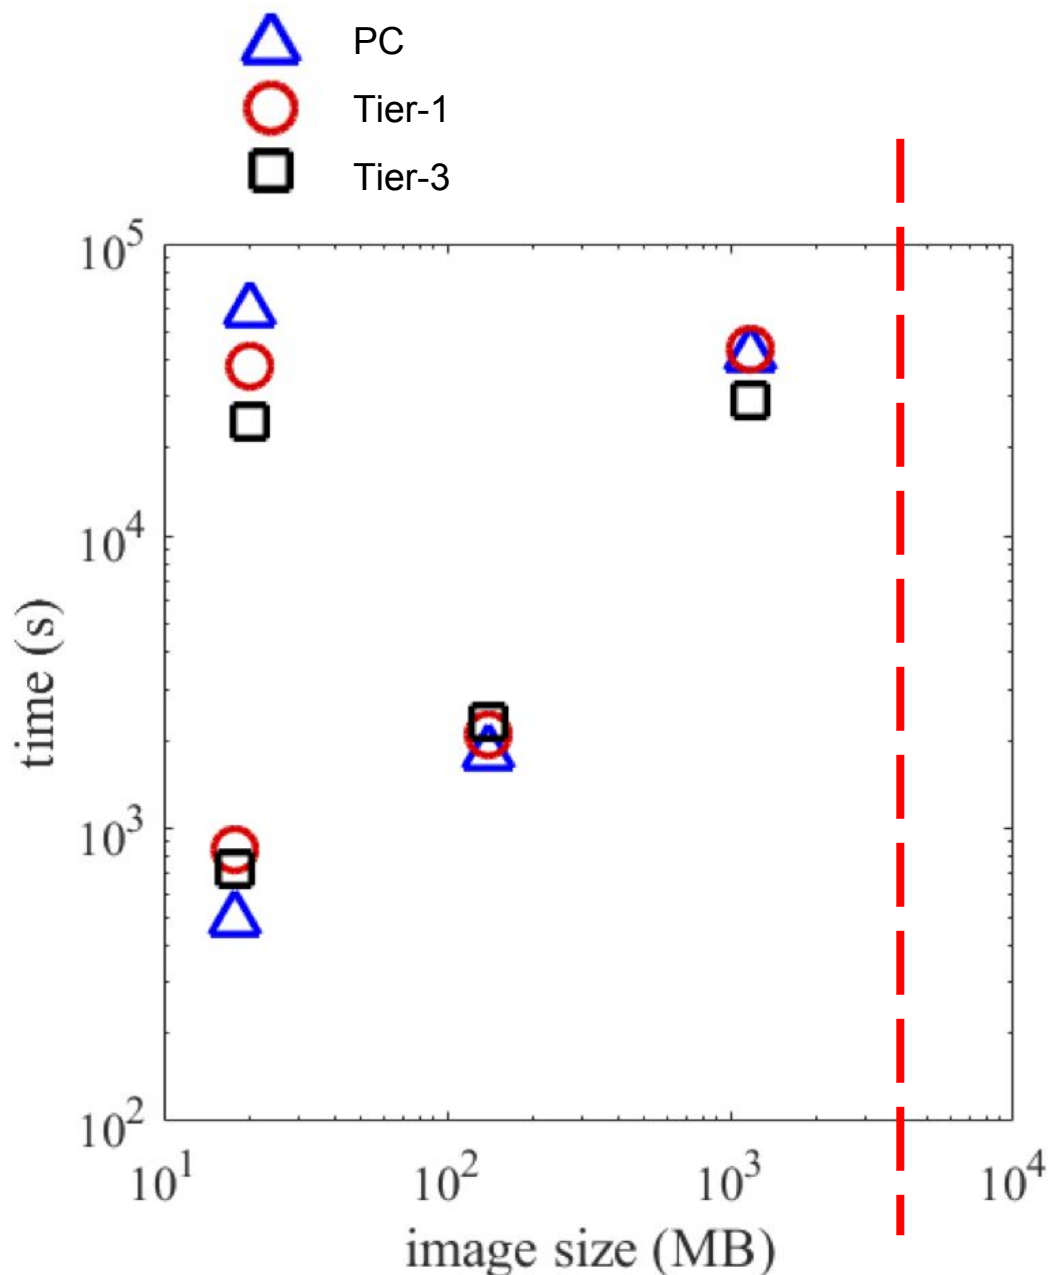
Workflow






Benchmark

- Serial execution: one single node was employed on both Tier-1 and Tier-3 systems
- Performance assessed on increasing input image size (3 sets)



Profiling



Bottleneck

- Image registration code developed for research
 - high accuracy
 - low performance
- Developed in 2000
 - Little documentation
 - Not designed to work on large 3D images (didn't exist at that time)
 - Not designed to use parallel architectures

Workflow optimisation

- Remove all the commercial dependencies
- Tier-1 systems are optimised for parallel execution
- *New* image registration code:
 - Clean room re-writing
 - Open-source with less-restrictive license to aid deployment on different platforms
 - Compute the 10,000s of voxel in parallel
 - Employ MPI to scale to higher number of nodes

Lesson learned

- Usually, in computational medicine, a single component of the workflow would benefit from executing on an HPC system: workflow should be flexible enough to allow the execution of single parts on different platforms.
- Tier-1 systems are not optimised for serial execution and their performance is dependent on the size of the simulation: best results are obtained when dealing with 100s if not 1000s of parallel tasks.
- Research software for computational medicine should be written with HPC in mind. Software we write today will continue to be used for years to mode, and the growing data footprint in our domain will pose scalability issues soon.
- **Plan for HPC** should become the motto of computational medicine software developers.

Acknowledgements

- Prof M. Viceconti - PI NMSK applications
- Dr. Shannon Li - NMSK applications
- Dr. A. Marzo - Deputy PI CVSR applications
- Dr. A. Narracott - IT infrastructures CVRS applications
- Dr. S. Wood - IT infrastructures Clinical data integration
- Mrs. L. Statham - Research Administration
- Dr. P. Tooley - Research Associate Development
- Dr. E. Dall'Ara - BoneDVC
- Dr. M. Giorgi - BoneDVC
- Dr. P. Bhattacharya - BoneDVC

Acknowledgements



Acknowledgements

CompBioMed Centre of Excellence (H2020-EINFRA-2015-1-675451)

EPSRC funded project MultiSim (Grant N. EP/K03877X/1)

UK NC3R funded project (Grant N. 113629)

FP7 European projects CHIC (Grant N. 600841)

Thank You!

Alessandro Melis

a.melis@sheffield.ac.uk

Department of Mechanical Engineering

The University of Sheffield

INSIGNEO Institute for *in silico* Medicine

HPC@Sheffield

Monday 26 March 2018

09.00 - 16.00

Firth Hall

Sheffield Teaching Hospitals 
NHS Foundation Trust

INSIGNEO
Institute for *in silico* Medicine



The
University
Of
Sheffield.