Exploiting International e-Infrastructures for Large Scale Computational Science

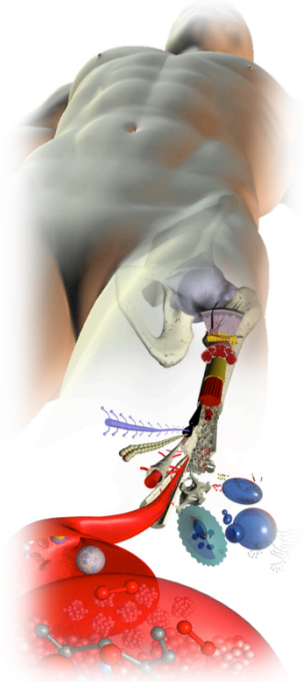

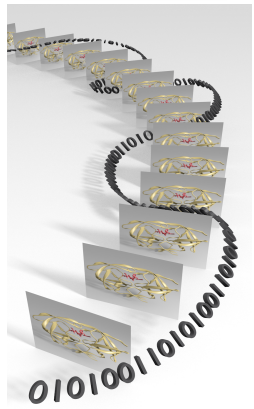

Peter Coveney

Centre for Computational Science

University College London

United Kingdom

Acknowledgements

UCL

- Agastya P. Bhati
- Shunzhou Wan
- David W. Wright
- Stefan J. Zasada

PL GRID

COMPAT
Computing Patterns for High Performance Multiscale Computing

EUDAT

PCSS

MRC Medical Research Council

EPSRC
Engineering and Physical Sciences Research Council

gsk

- Ian Wall
- Darren Green
- Paul Bamborough

RUTGERS

archer

epcc

Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities

PRACE

Hartree Centre
Science & Technology Facilities Council

Pfizer

- Sarah Skerratt
- Kiyoyuki Omoto
- Veerabahu Shanmugasundaram
- Sharan K. Bagal

الصندوق القطري لرعاية البحث العلمي
Qatar National Research Fund
Member of Qatar Foundation

CompBioMed

Overview

- AHE: A Lightweight Science Gateway
- Application 1: Calculating Drug Binding Affinities
  - Ensemble Methods
  - Binding Affinity Calculator
  - User Friendly Binding Affinity Calculator
  - High Performance & Cloud Computing
- Application 2: Multiscale simulation of nano materials
- Conclusions

AHE Motivation

UCL

- Problems for individual users installing/compiling/ optimizing application
- Complexities of using computational infrastructures
  - Job workflows, from staging data to launching jobs
- Security is an obstacle for users
  - Applying for certificates
  - Generating MyProxy
- A solution is needed to simplify usage for scientific end users, e.g. to allow clinicians to run simulations at the click of a button.

AHEApplication Hosting Environment

- **Application Hosting Environment**
  - Simplifying Access to the Grid
  - Community Model.
- **Simplifies security**
  - End-User avoids grid security and MyProxy configuration and generation.
- **Simplifies application setup**
  - End-User does not have to compile, optimise, install and configure applications.
- **Simplifies basic workflow**
  - AHE stages the data, runs and polls the job and fetches the results automatically
- **Simplifies compute access through RESTful web-services**
  - Provides a RESTful interface
  - Clients and services access infrastructure and apps with ‘Software as a Service’

AHEApplication Hosting Environment

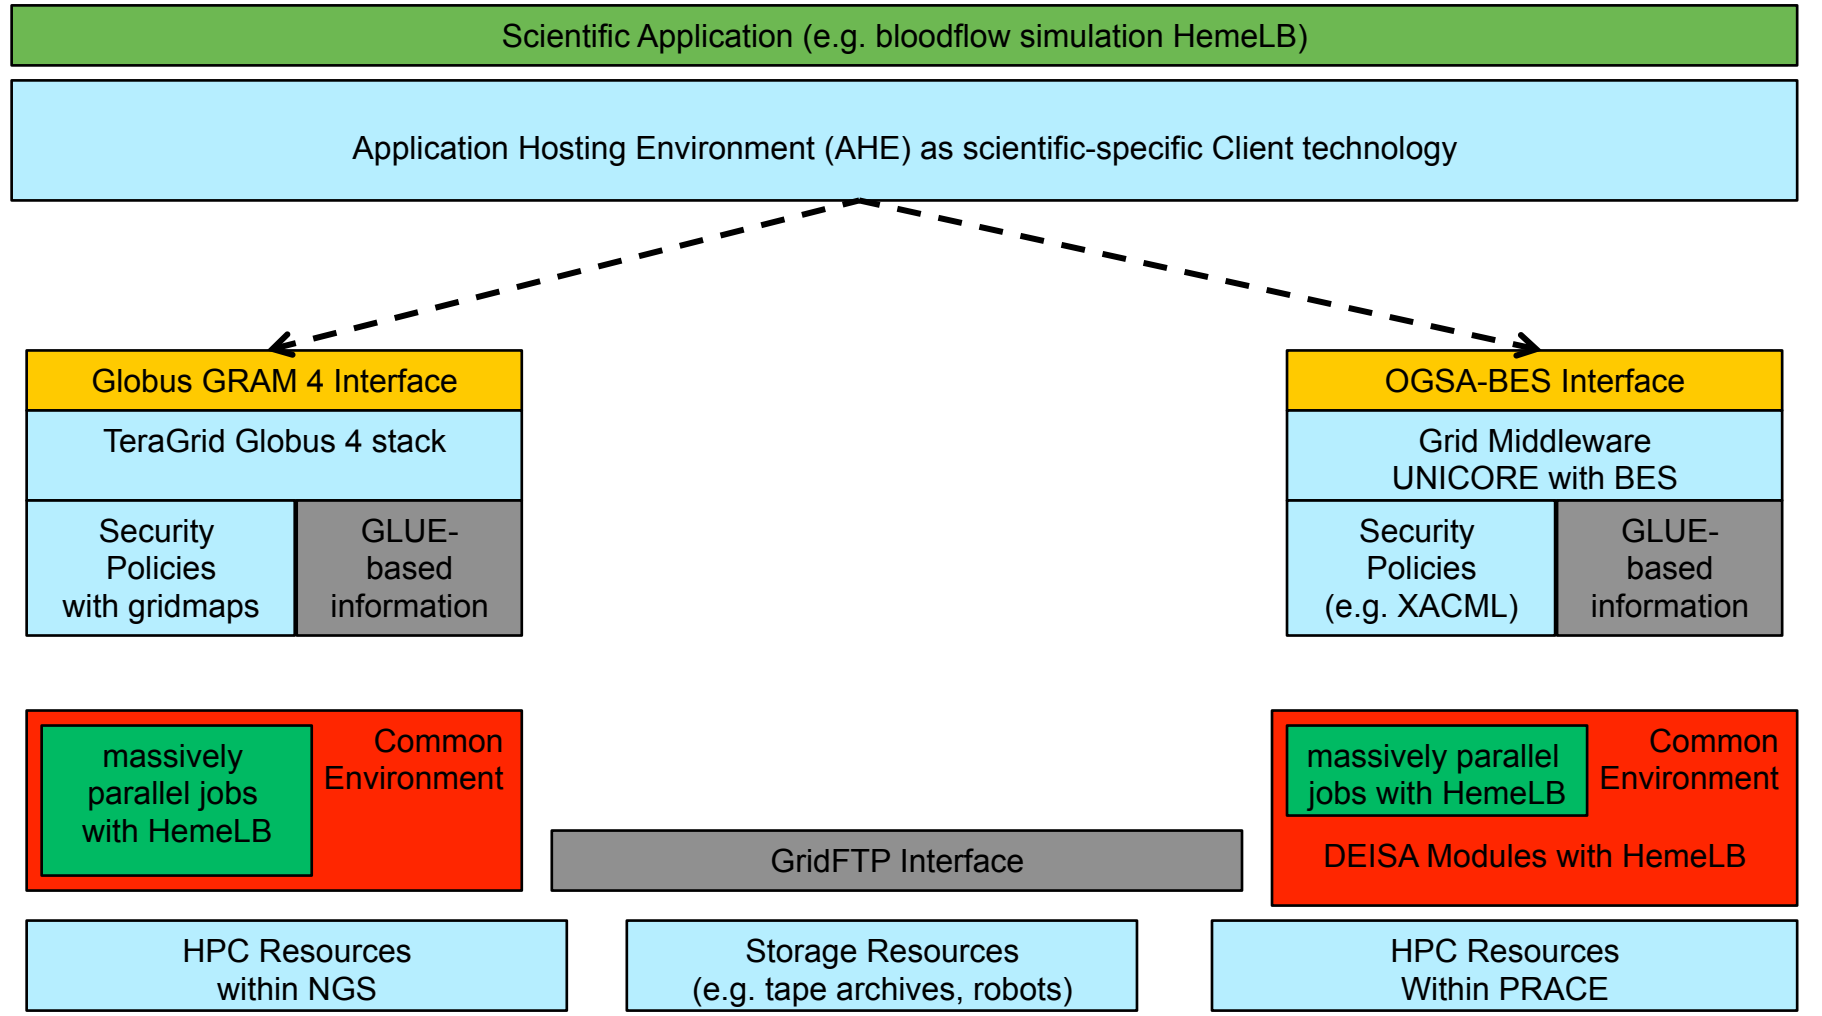

Stefan Zasada, Steven Manos, Morris Riedel, Johannes Reetz, Michael Rambadt et al.,
For the Virtual Physiological Human (VPH) projects that require interoperability of numerous Grids

AHEApplication Hosting Environment

UCL

Federating HPC resources

Globus

Local resources

GridSAM

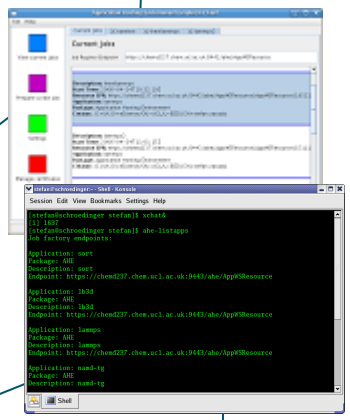

UK

ARCHER

Leeds

Manchester

Oxford

RAL

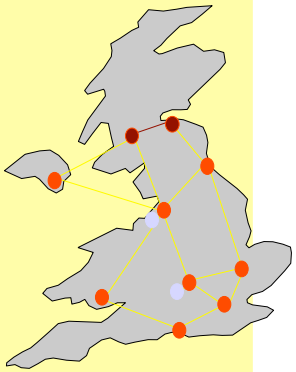

XSEDE

Globus

UNICORE

PRACE

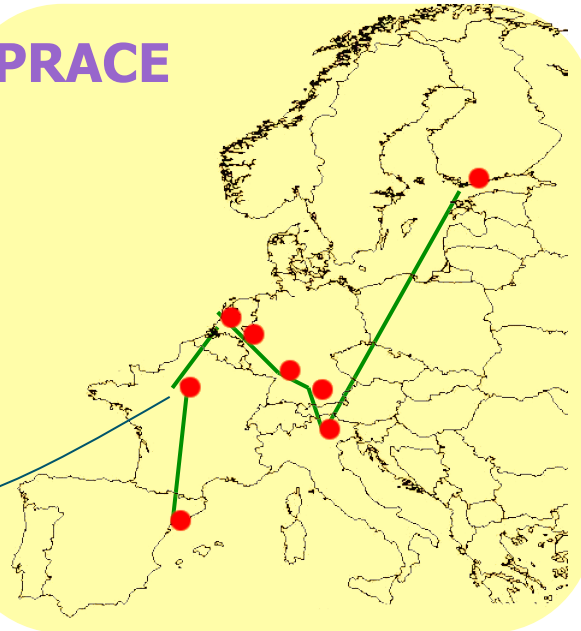

App 1: Calculating Drug Binding Affinities

- Ligand binding driven by changes in the Gibbs free energy
- The more negative the ΔG the stronger the binding

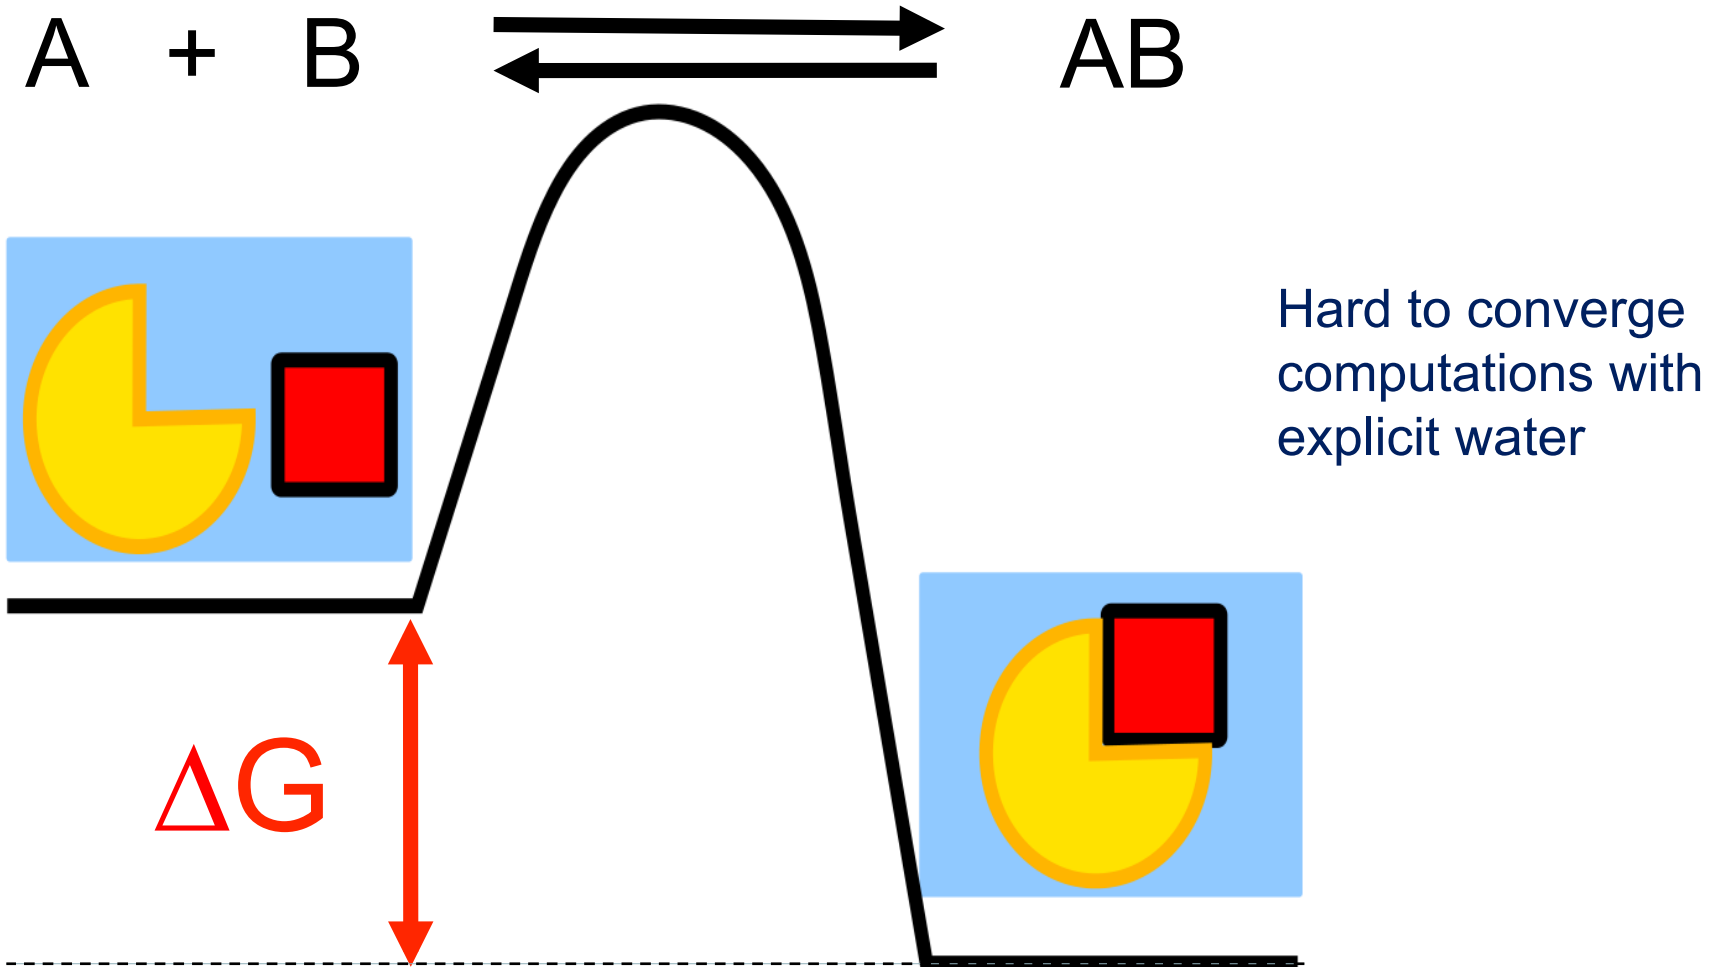

Hard to converge computations with explicit water

Binding Free Energy Calculation

Absolute binding free energy with end-point methodologies

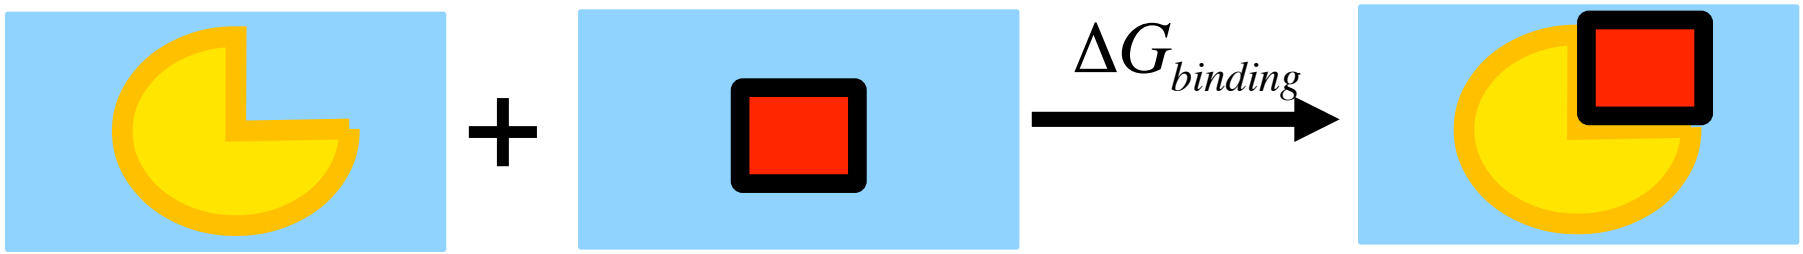

$$\Delta G_{binding} = G_{complex} - G_{protein} - G_{ligand}$$

The end-point free energy calculation methodologies **MMPB(GB)SA** are used to estimate the free energies of protein, ligands and their complexes using conformations generated via molecular dynamics simulation.

$$\begin{aligned} G^i &= G^i_{MMPB(GB)SA} - TS^i_{conf} \\ &= E^i_{MM} + G^i_{solv} - TS^i_{conf} \\ &= E^i_{MM} + G^i_{PB/GB} + G^i_{SA} - TS^i_{conf} \end{aligned}$$

1-trajectory *vs* 3-trajectory

Computing Binding Free Energy Difference

Relative binding free energy with alchemical mutation: make use of thermodynamic cycle to calculate binding free energy difference

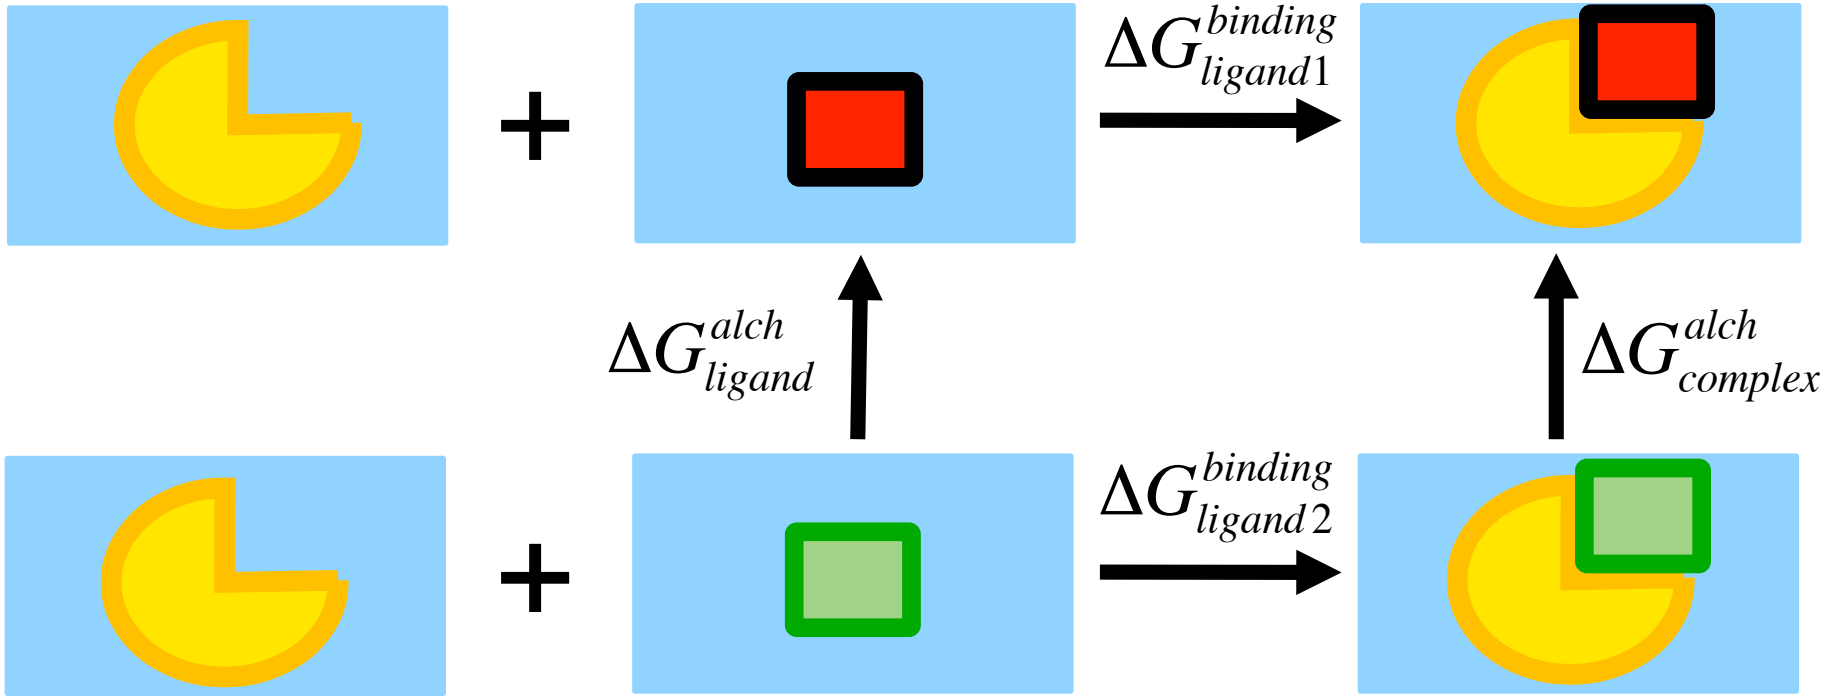

$$\Delta\Delta G^{binding} = \Delta G^{binding}_{ligand2} - \Delta G^{binding}_{ligand1} = \Delta G^{alch}_{ligand} - \Delta G^{alch}_{complex}$$

Ensemble MD Simulations

- The MM/PBSA results follow well defined Gaussian distributions.
- Configurational entropies, obtained from normal mode estimates, closely resemble normal distributions.

Drug – HIV-1 protease

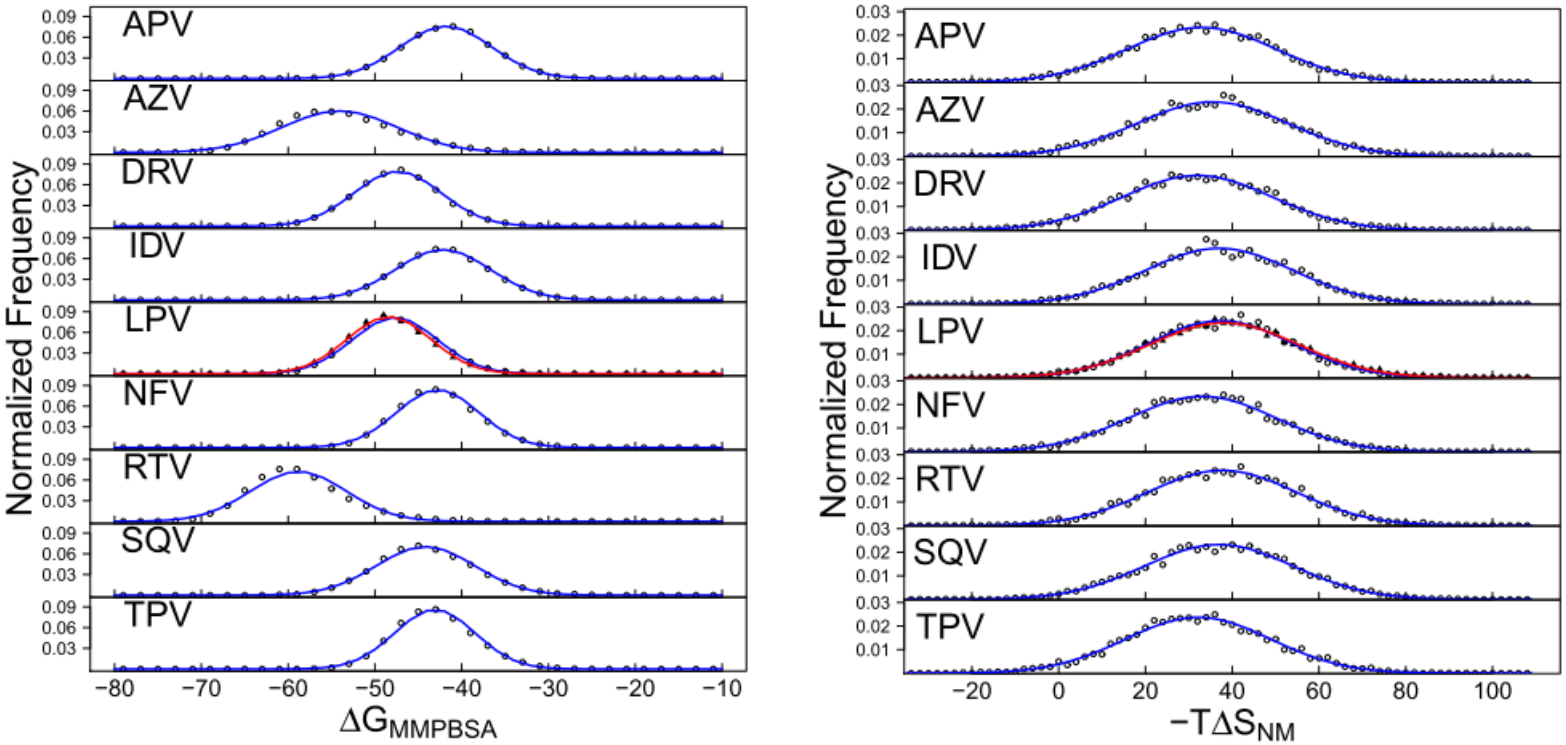

Wright, Hall, Kenway, Jha & Coveney, JCTC, (2014), DOI: 10.1021/ct4007037.

Computational Application to Drug Affinity Ranking – Single MD simulation

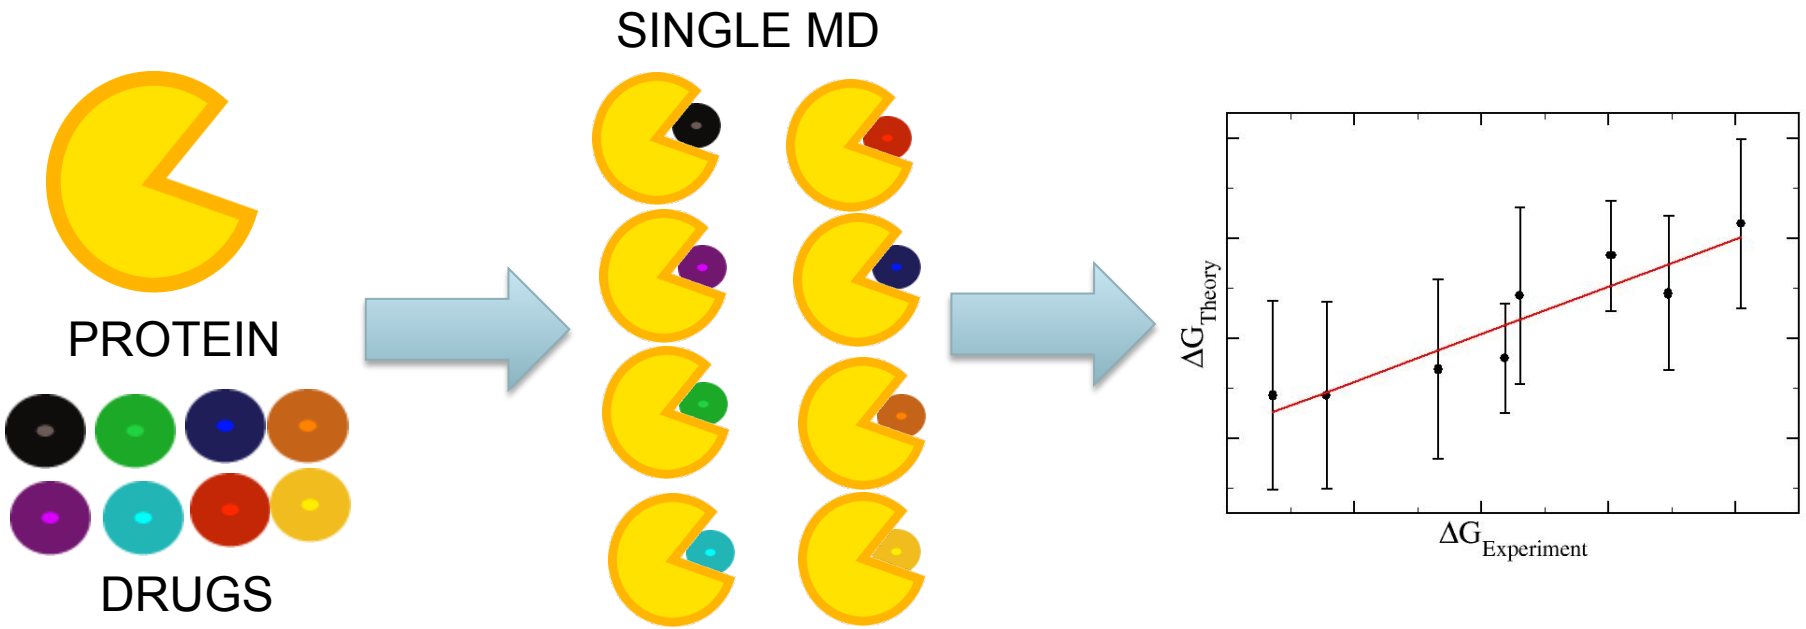

Errors uncontrolled
Results unreproducible

Computational Application to Drug Affinity Ranking – Ensemble Simulations

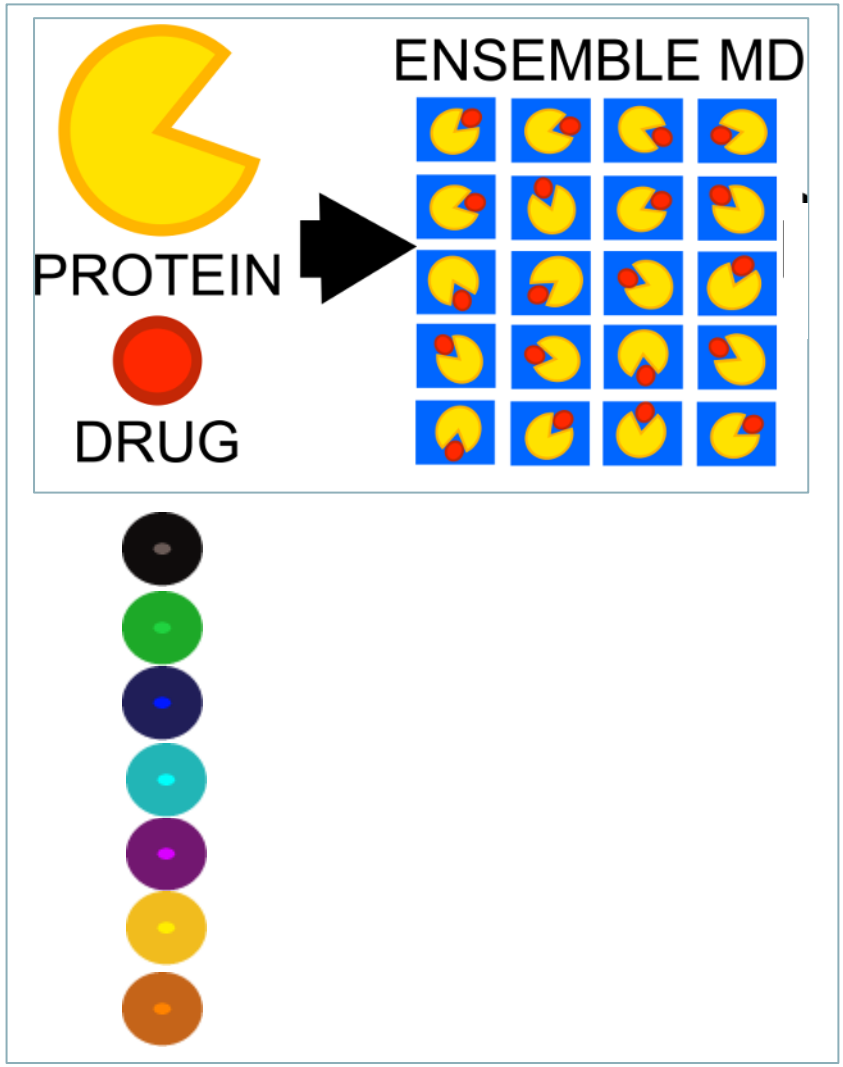

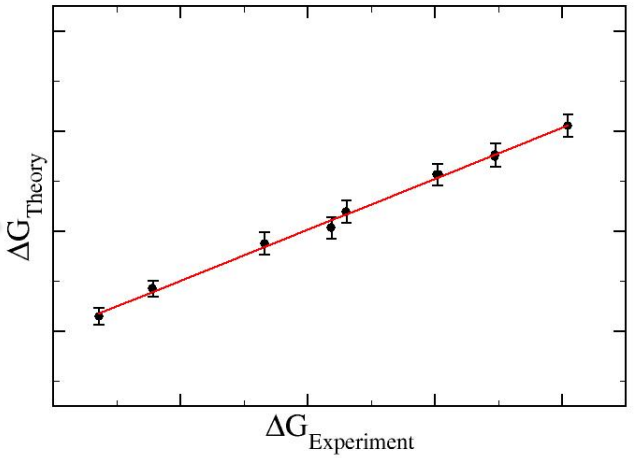

Errors fully under control
Results reproducible.

Free Energy Methodologies

ESMACS: Enhanced Sampling of Molecular dynamics with the approximation of Continuum Solvent

- "Absolute" free energies
- Compare diverse ligands
- Statistical error analysis
- MM/PBSA
- Conformational entropy

TIES: Thermodynamic Integration with Enhanced Sampling

- Relative binding affinities
- "Exact"
- Limited range of application
- Ensemble needed for each λ

Ensemble Molecular Dynamics Protocol

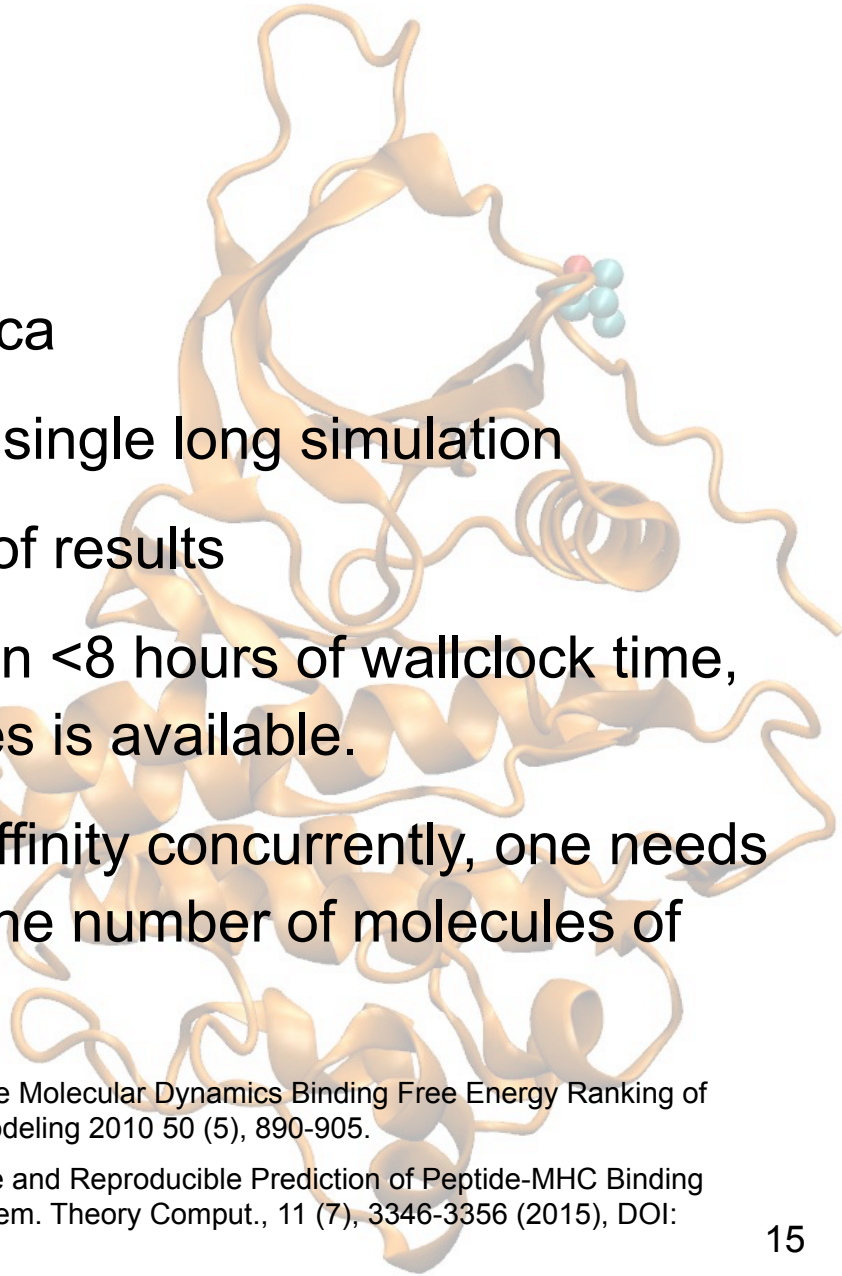

- Run <25 ‘replica’ simulations
- Vary only initial velocities
- 4 ns of production trajectory per replica
- More efficient sampling compared to single long simulation
- Allows us to examine reproducibility of results
- The workflow can be completed within <8 hours of wallclock time, provided the required number of cores is available.
- To compute more than one binding affinity concurrently, one needs to multiply the node requirement by the number of molecules of interest.

Sadiq, S.K, Wright, D.W., Kenway, O.A. and Coveney, P.V. “Accurate Ensemble Molecular Dynamics Binding Free Energy Ranking of Multidrug-Resistant HIV-1 Proteases.” Journal of Chemical Information and Modeling 2010 50 (5), 890-905.

Wan, S., Knapp, B., Wright, D.W., Deane, C.M., Coveney, P.V., "Rapid, Precise and Reproducible Prediction of Peptide-MHC Binding Affinities from Molecular Dynamics that Correlate Well with Experiment", J. Chem. Theory Comput., 11 (7), 3346-3356 (2015), DOI: 10.1021/acs.jctc.5b00179

Binding affinity calculator (BAC)

BAC can reliably predict binding affinities of compounds with their target proteins, and be used potentially as a drug ranking tool in clinical application or a virtual screening tool in pharmaceutical lead discovery.

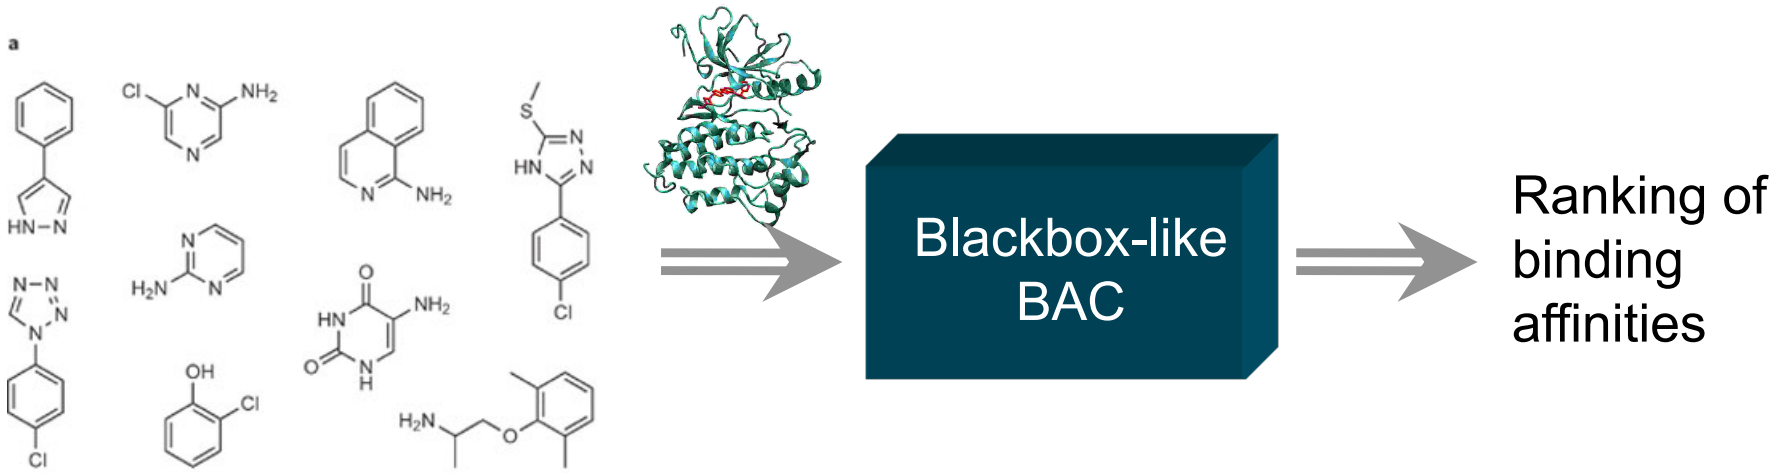

*S. K. Sadiq, D. Wright, S. J. Watson, S. J. Zasada, I. Stoica, Ileana, and P. V. Coveney, "Automated Molecular Simulation-Based Binding Affinity Calculator for Ligand-Bound HIV-1 Proteases", Journal of Chemical Information and Modeling, **48**, (9), 1909-1919, (2008), DOI: 10.1021/ci8000937.*

BAC Workflow: ESMACS Method

BAC: rapid and accurate binding affinity calculation on timescales relating to clinical decision making on drug selection and to pharmaceutical lead discovery.

Total of 1000s cores on HPC/cloud resources required per study

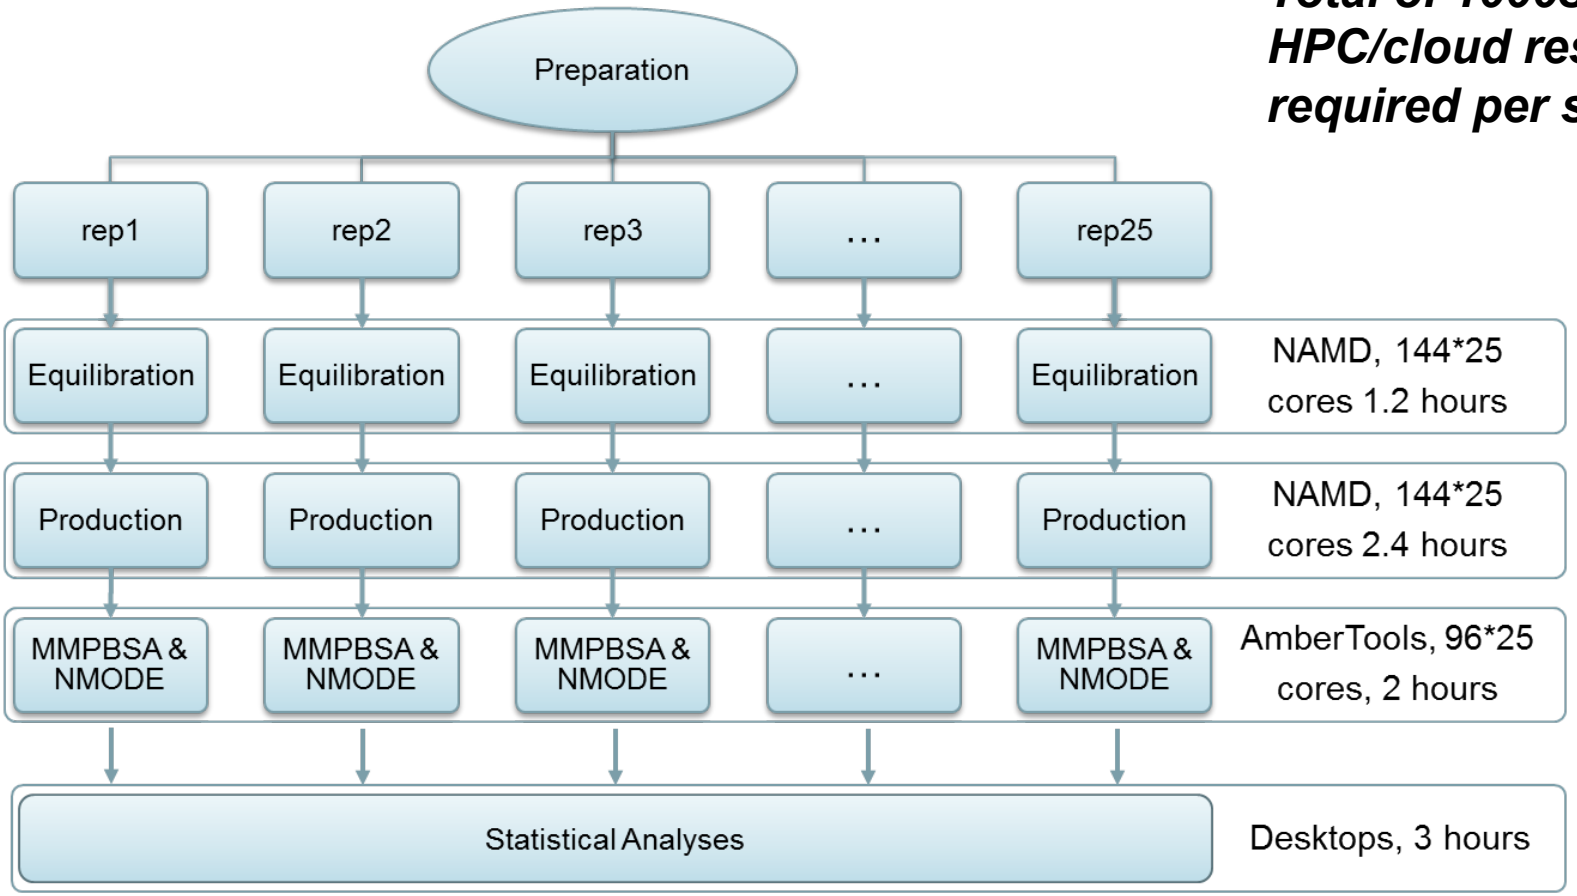

A. Bhati, S. Wan, D. Wright, P. V. Coveney, "Rapid, accurate, precise and reliable relative free energy prediction using ensemble based thermodynamic integration", Journal of Chemical Theory and Computation, DOI: 10.1021/acs.jctc.6b00979.

BAC Workflow: TIES Method

UCL

Binding Affinity Calculator (BAC) is a software toolkit which automates the implementation of TIES (and ESMACS) methods for binding affinity calculations

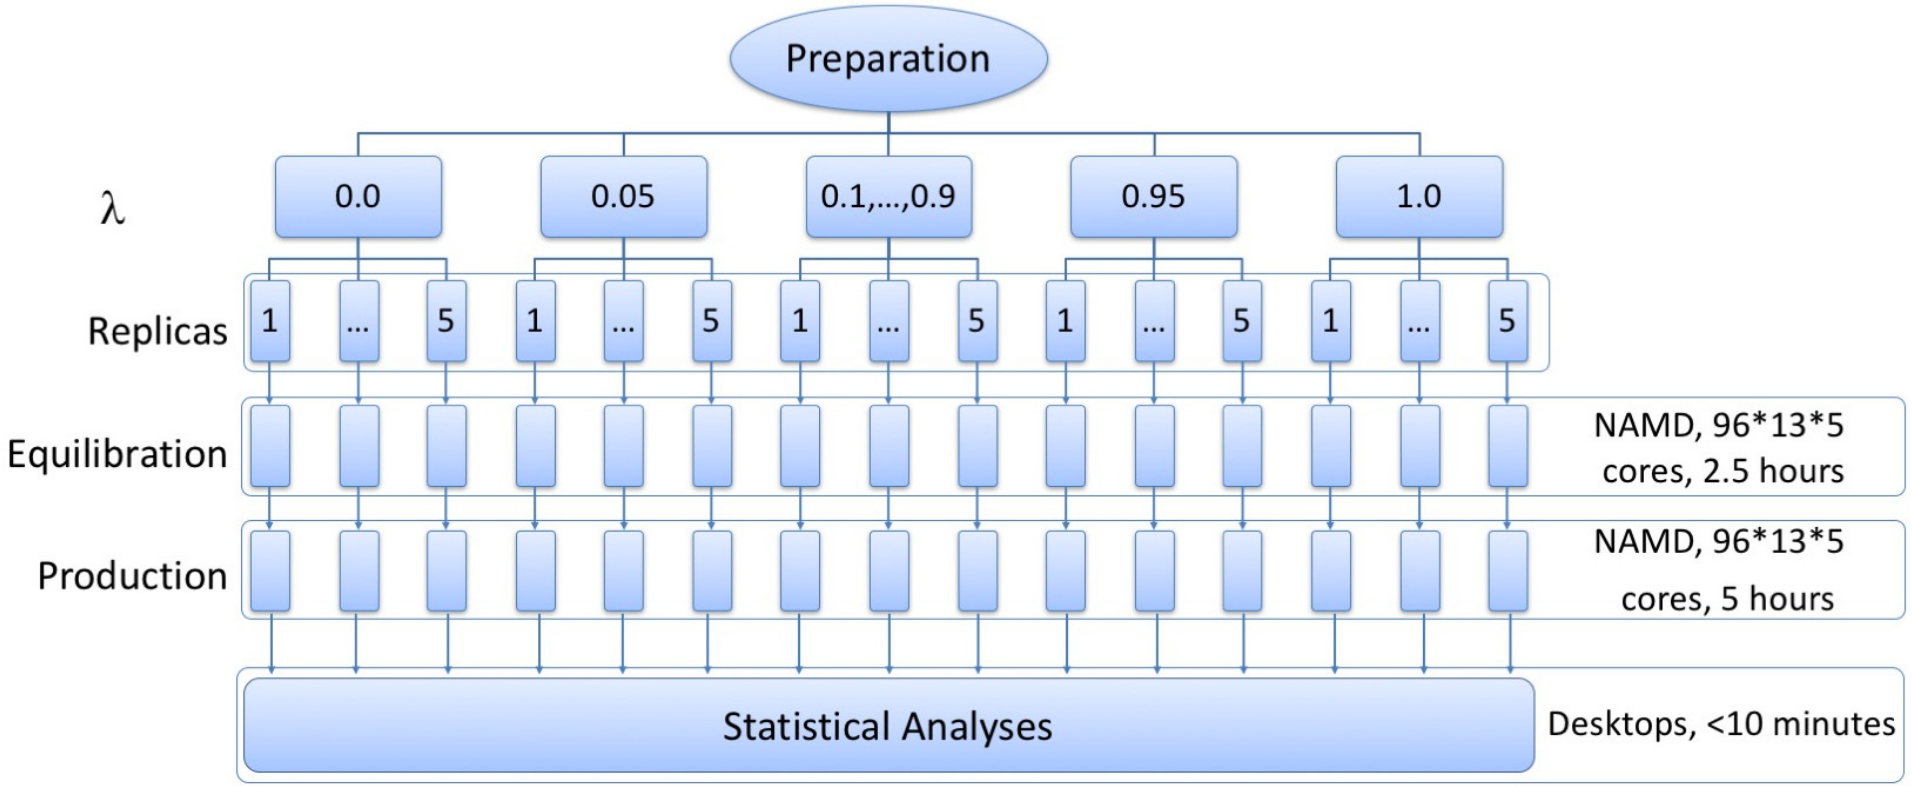

S.K. Sadiq, D. Wright, S.J. Watson, S.J Zasada, I. Stoica, P.V. Coveney, J. Chem. Inf. Model., **2008**, 48, 1909-1919.
A.P. Bhati, S. Wan, D.W. Wright & P.V. Coveney, *J. Chem. Theory Comput.*, **2017**, 13, 210–222.

ufBac

- User Friendly BAC makes reliable, repeatable binding affinity calculations available to anyone
- Web interface allows full BAC workflows via simple, user friendly client
- Manages execution of calculations on a range of resources, from HPC to commercial cloud platforms
- ufBac manages complete study execution and data archival on behalf of the user

BAC Requires Unified E-Infrastructures

The BAC workflow requires resources of different scales to execute

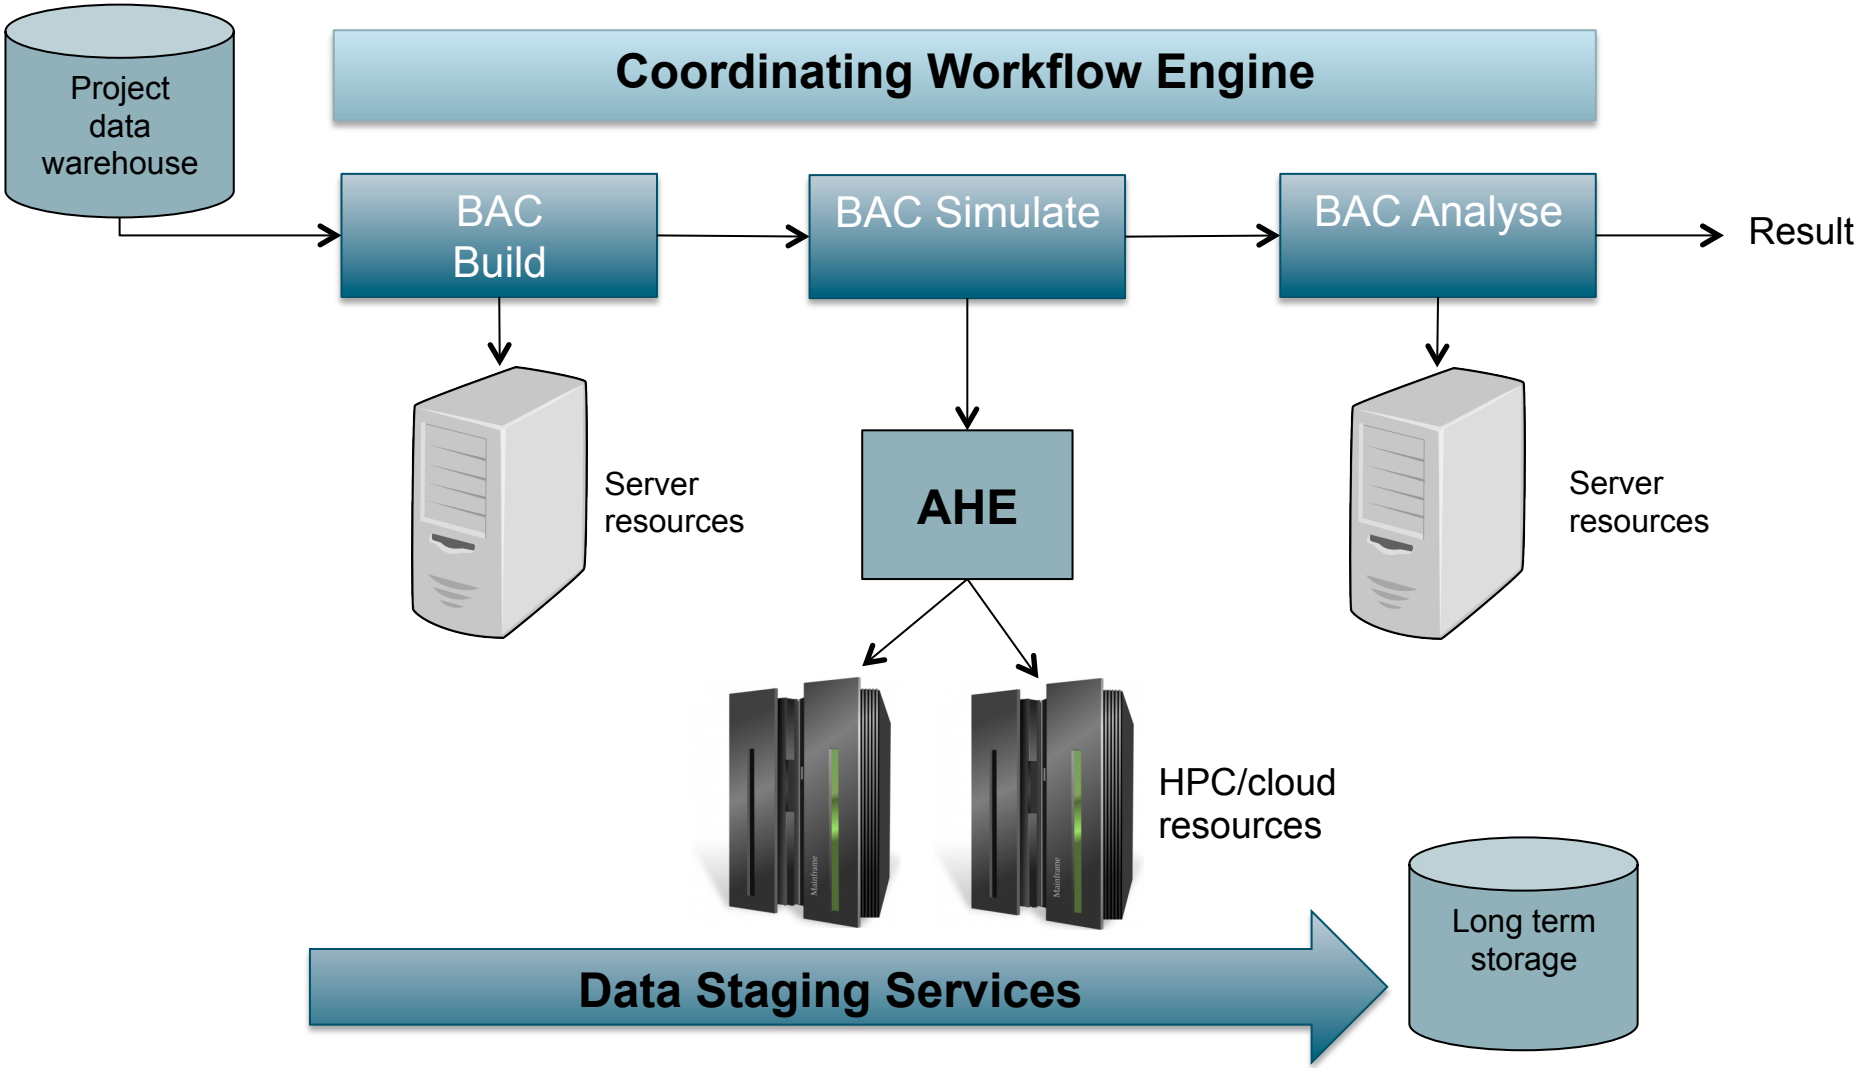

Petascale Computing Facilities Used by Us

UCL

XSEDE
Extreme Science and Engineering Discovery Environment

Piz Daint

SuperMUC

Cartesius

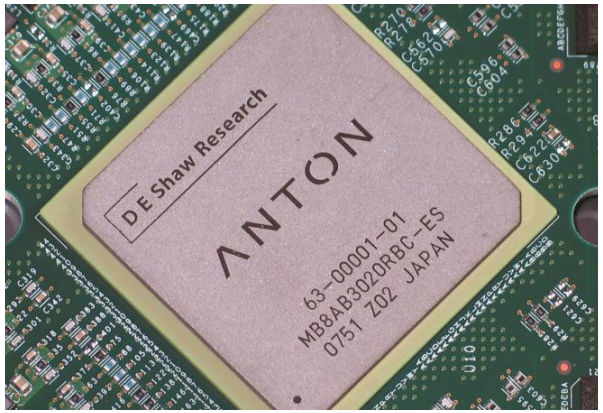

Anton

PRACE

HECToR

ARCHER

EMERALD

Blue Joule

Blue Wonder

Hartree Centre
Science & Technology Facilities Council

Giant Workflow on SuperMUC

UCL

Industrial Strength of BAC

Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities

"Scientists can work out the way that a candidate drug will act on a target in the body – a protein – and in a matter of a few hours."

11-13 June 2016
36 hours on active machine.

SuperMUC

LRZ Press Release:

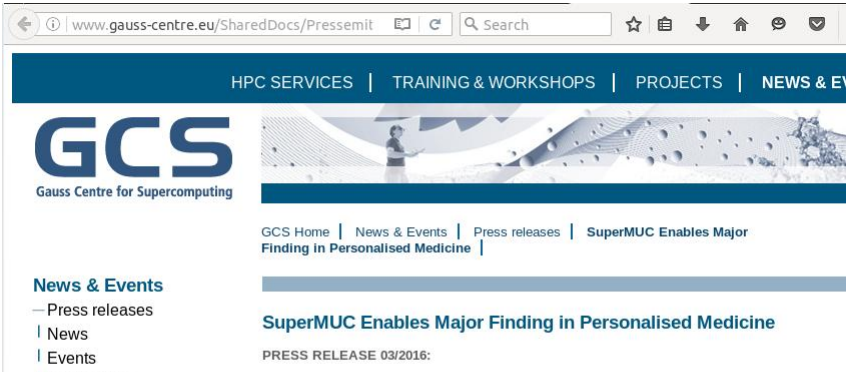

Post on the London Science Museum Blog:

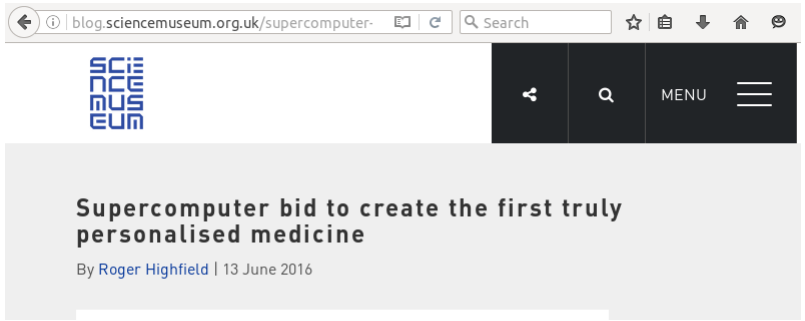

Giant Workflow

- 60 combinations of drug and protein investigated
  - 30 absolute free energy calculations of diverse ligands (ESMACS)
  - 30 relative free energy calculations of related ligands (TIES)
- Refine ESMACS protocols
  - Multiple simulations
  - Differences between chemical groups
- New insights into TIES methodology
  - Determinants of statistical error
- Hardened automation scripts
  - BAC
  - FabSim

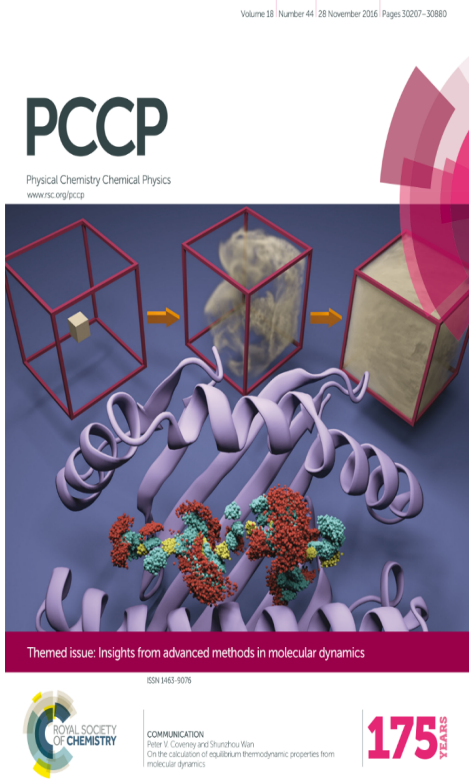

Rapid, accurate, precise and reliable relative free energy prediction using ensemble based thermodynamic integration, Agastya P. Bhati, Shunzhou Wan, David W. Wright, and Peter Vivian Coveney, *J. Chem. Theory Comput.*, 2016, **DOI:** 10.1021/acs.jctc.6b00979

CPUs *vs* GPUs

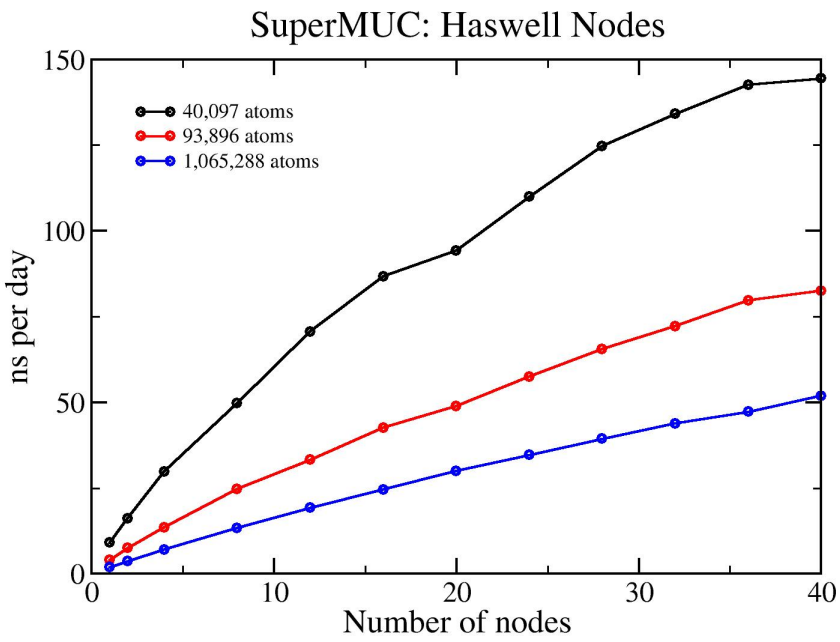

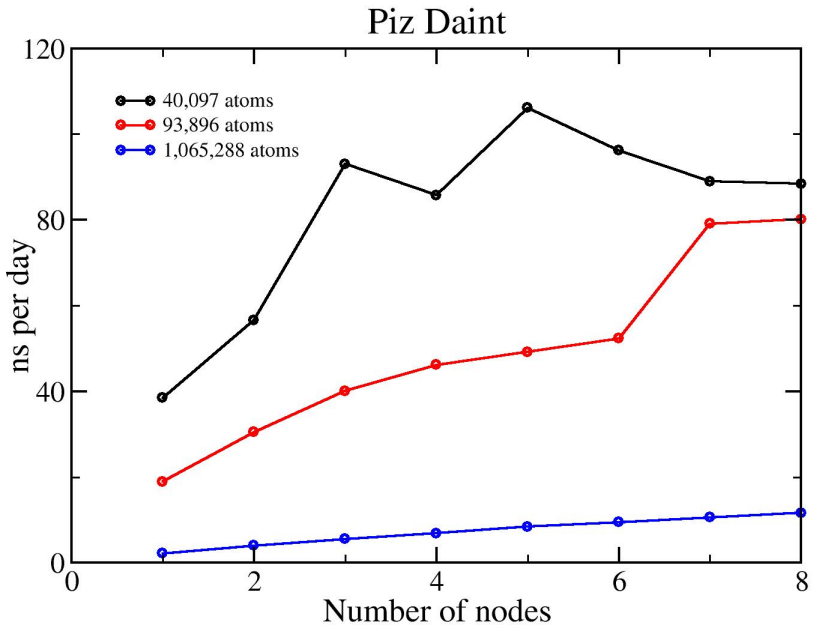

- Speed-up keeps growing on CPU machine, up to 40 nodes for a 40K atom system (36 atoms per core).
- Quicker speed is obtained on single GPU accelerated node, scaling up to 3 nodes on Piz Daint for the 40K system.

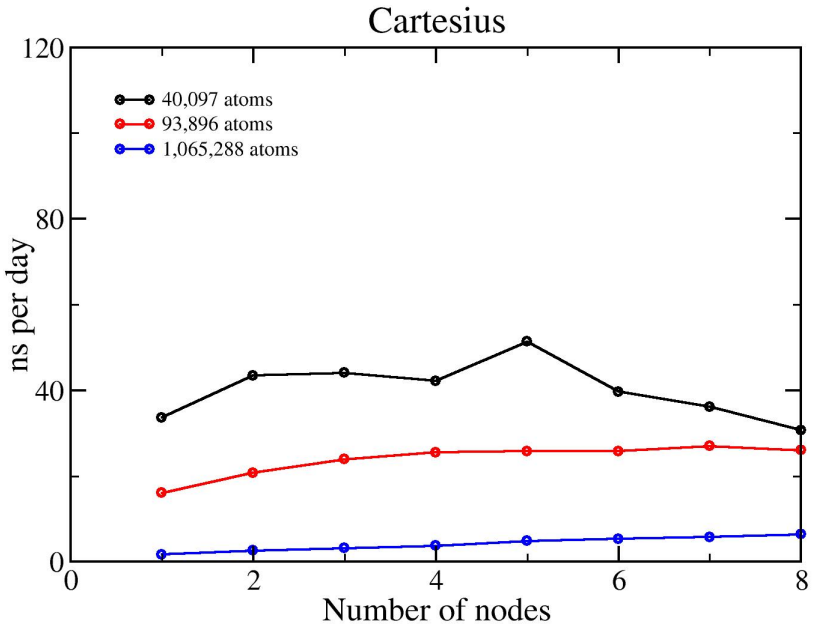

Commercial Clouds

- Cloud computing is an alternative schema for running applications on remote resources.
- Access to compute is provided in return for monetary payment.
- *Infrastructure as a Service* (IaaS) clouds provide access to CPU, memory and storage.

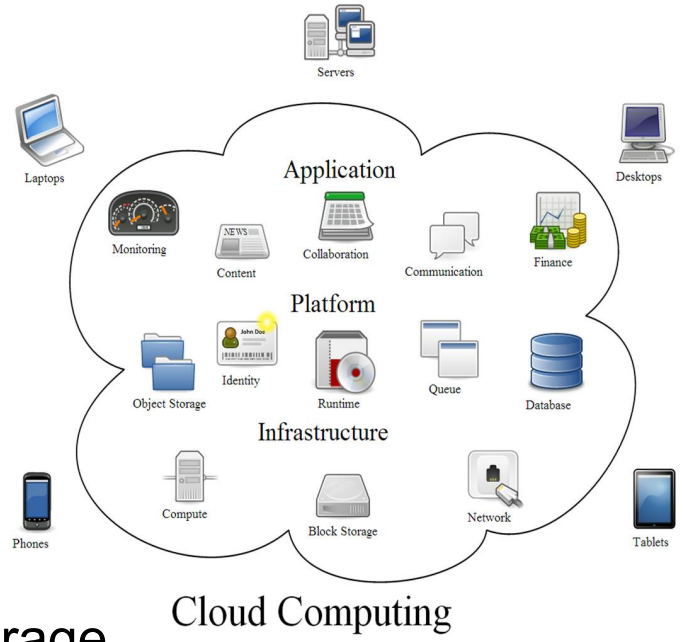

- *Software as a Service* (SaaS) clouds provide access to applications.

We've worked to deploy BAC on AWS and DNAnexus cloud platforms, and expect to publish it through the Amazon Market Place. We are currently evaluating deployment on Microsoft Azure.

Cloud Deployment: Amazon Web Services

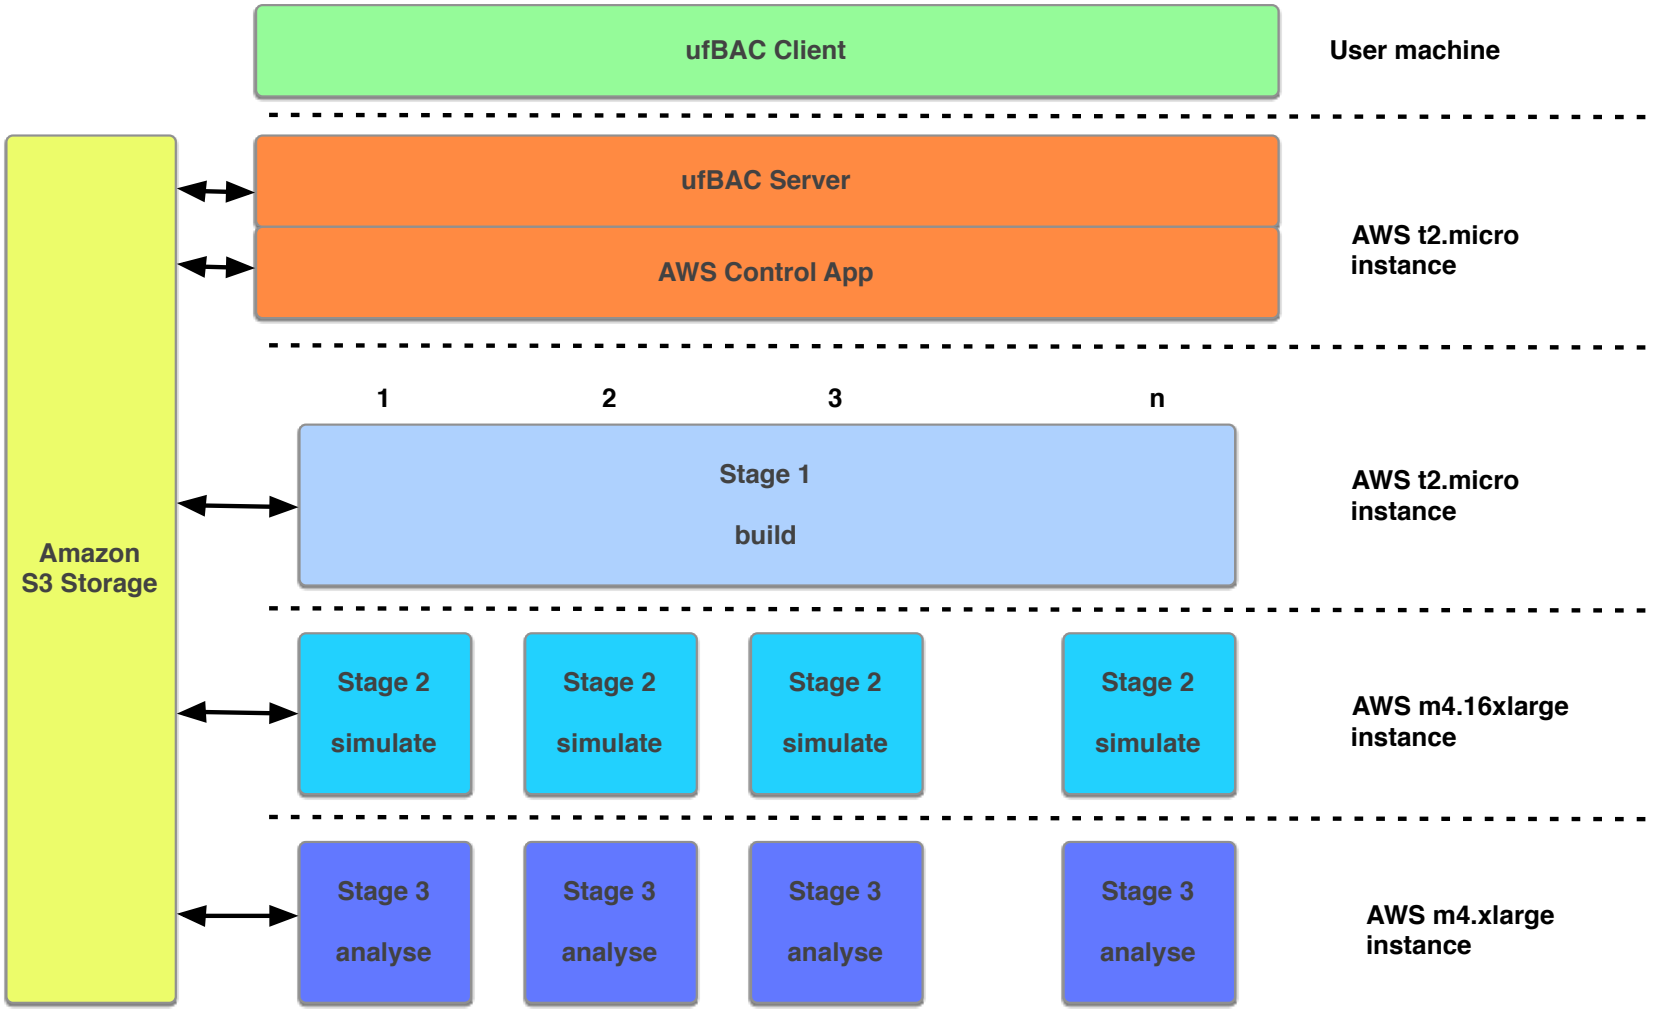

Cloud Scaling Performance

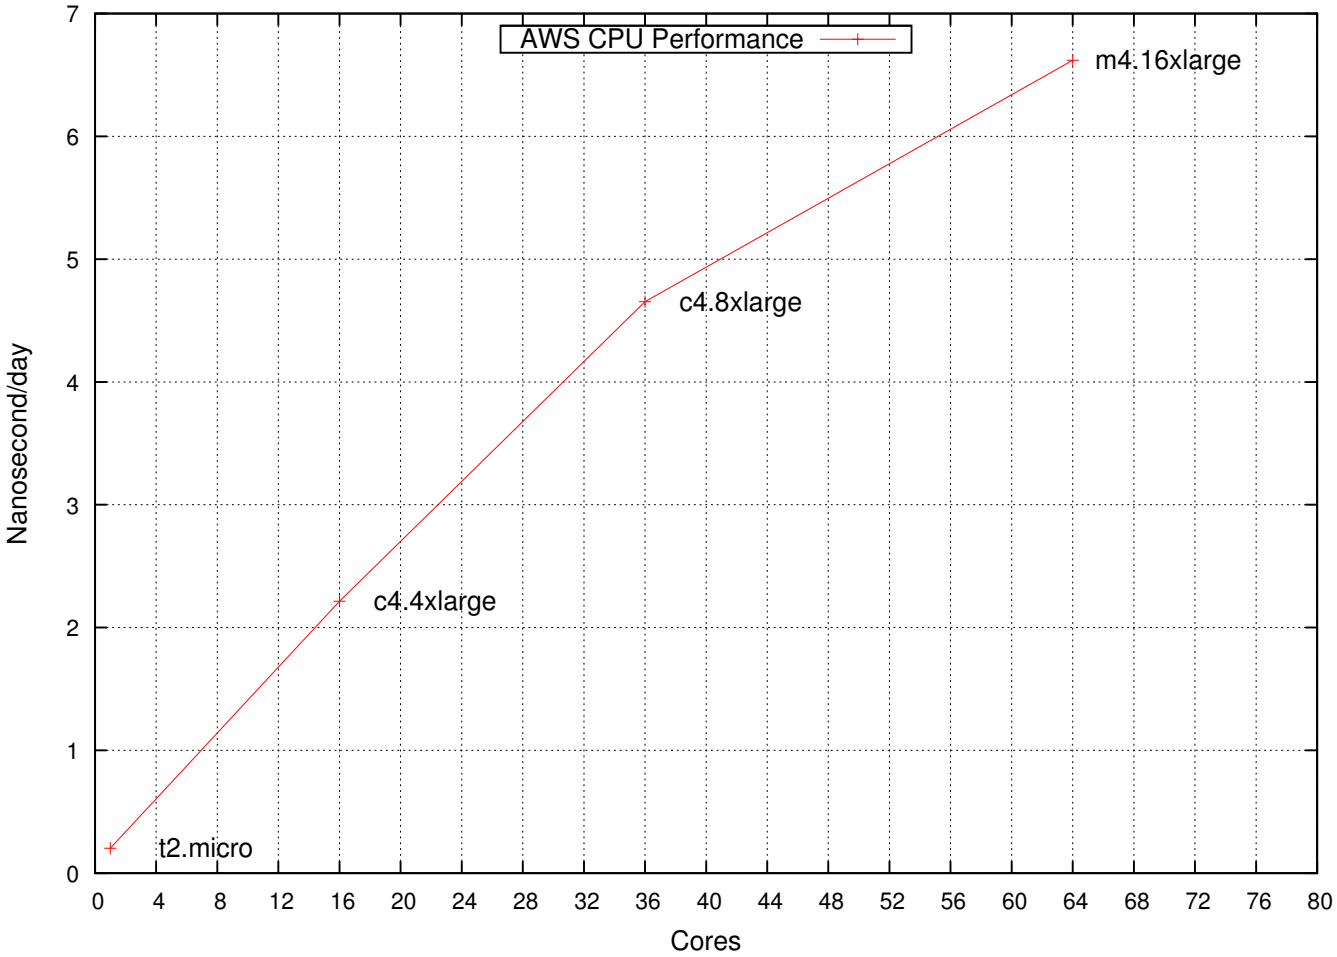

Instance	Size
t2.micro	1 CPU/1 GB RAM
c4.4xlarge	16 CPU/30 GB RAM
c4.8xlarge	36 CPU/60 GB RAM
m4.16xlarge	64 CPU/256 GB RAM

GPU Scaling Performance

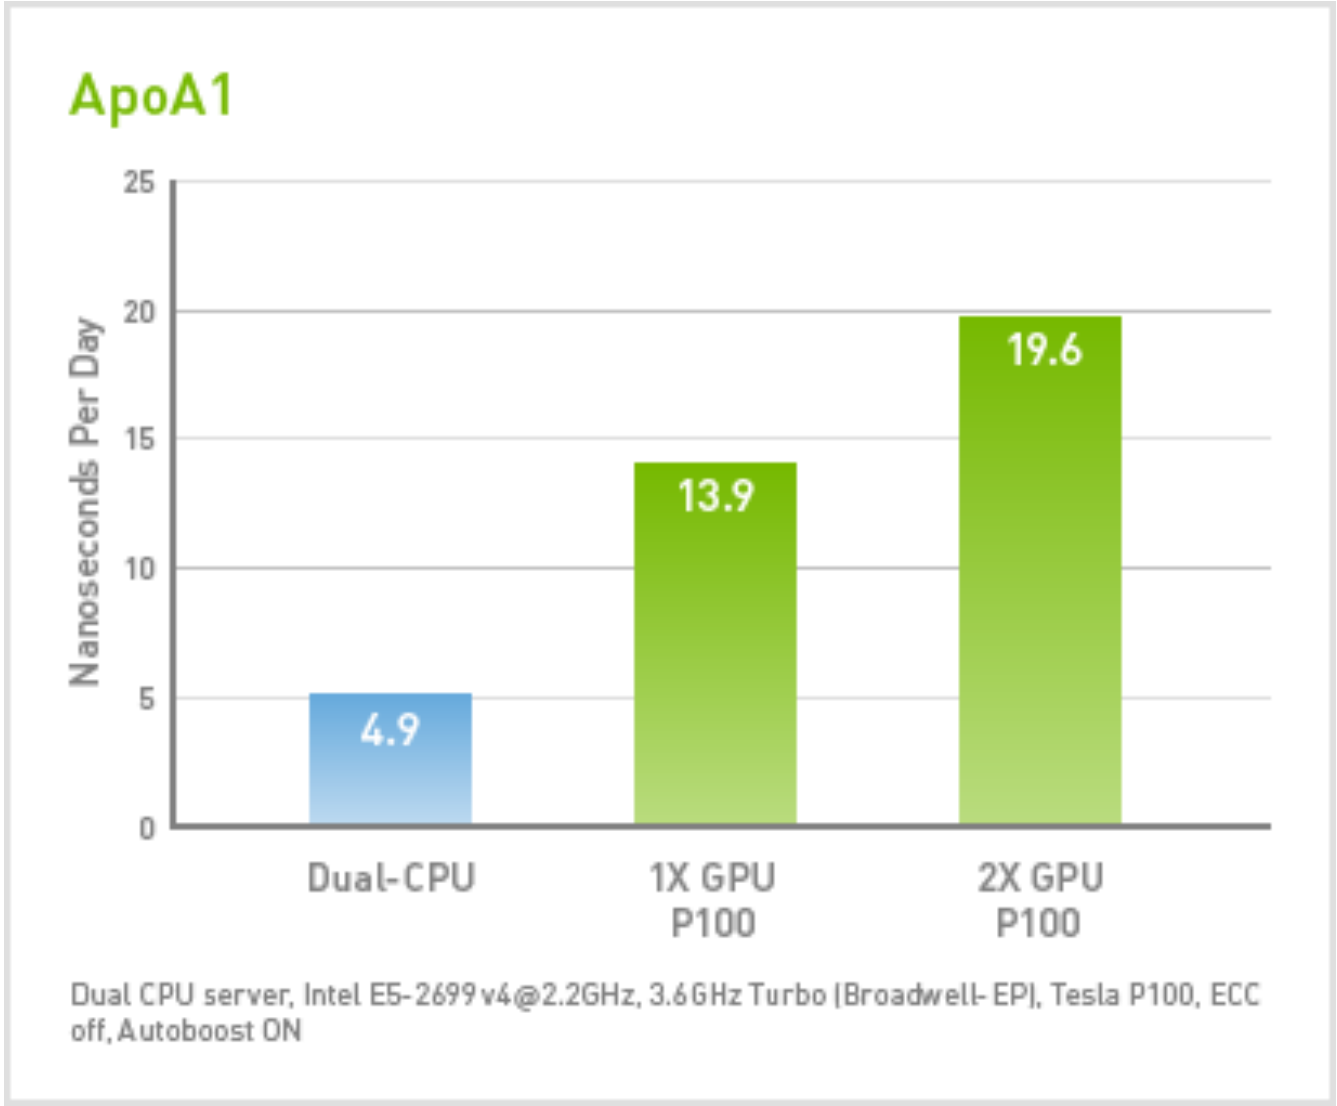

Microsoft Azure HPC in the Cloud

- Microsoft Azure Big Compute has **InfiniBand-interconnectec** HPC nodes (H-series) for Linux and Windows
- 2-3 microsec MPI latency
- nVidia K80 GPU nodes with InfiniBand available (K-series)
- Azure Batch Shipyard allows Dockerized deployment of MPI applications – including NAMD
- No queueing

Scaling to thousands of cores with NAMD benchmark (preliminary results shown)

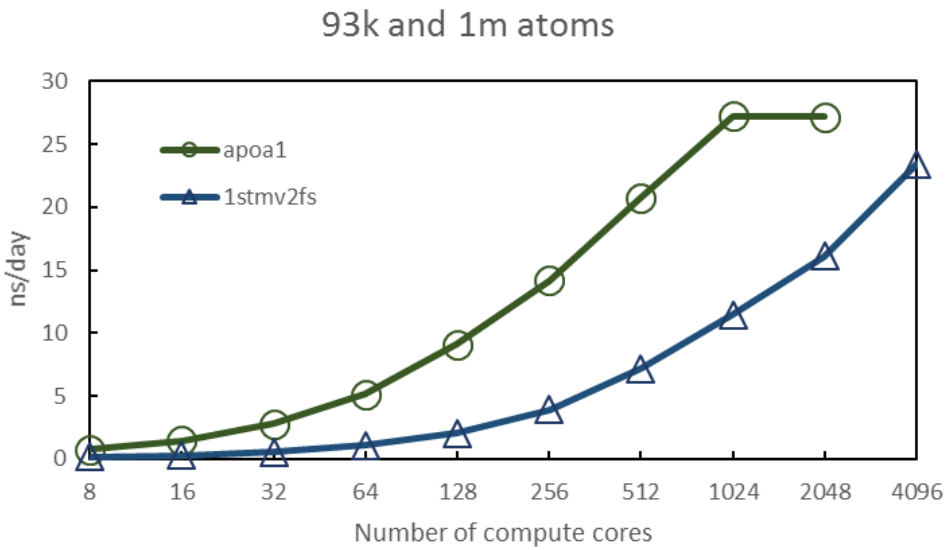

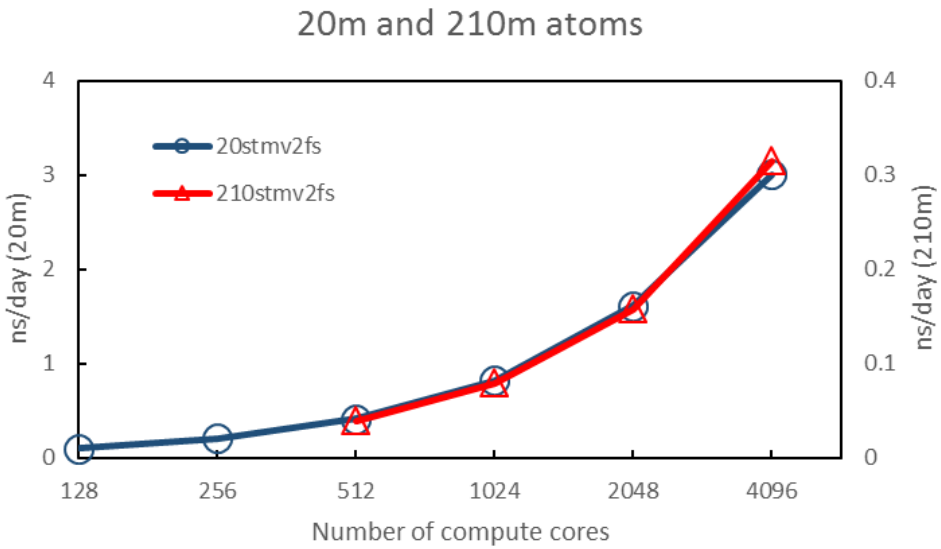

Cloud Deployment vs HPC Investment

- Cloud is ‘elastic’ − should always have enough resource to run your workflow.
- Public HPC runs on first come first served queue
- Clouds have security architectures beyond those provided by conventional HPC providers.
- Cloud gives access to resources at scale to ‘try things out’ but, at large scale, production cloud is expensive
- Learning curve is steep − need to intall most things

HPC Deployment: Radical/Ensemble Toolkit

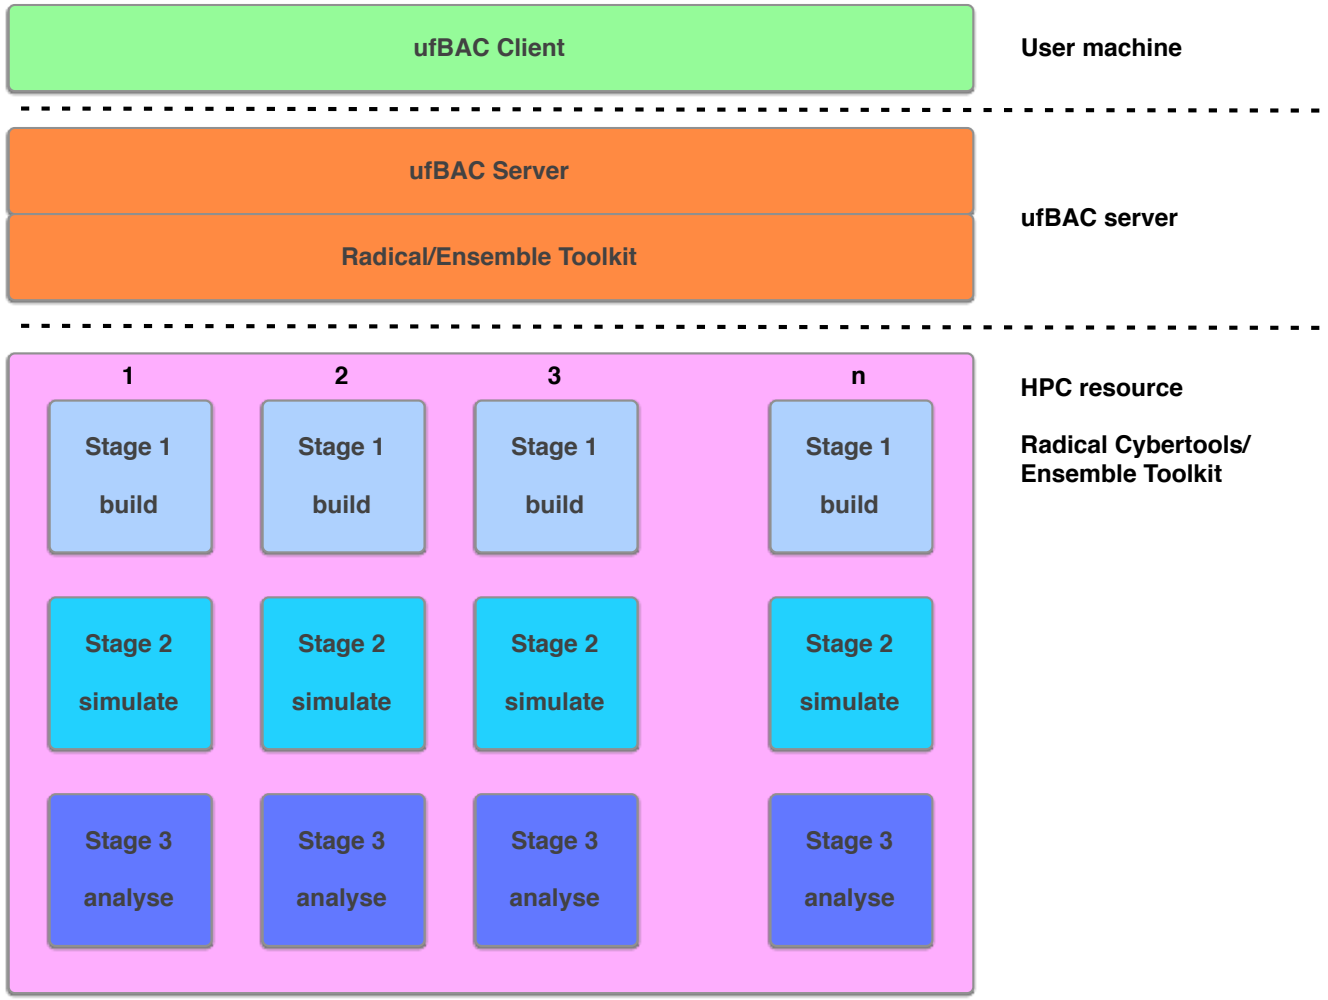

RADICAL-Cybertools

Four Layers:

- L4: Application
- L3: Workload Management (**WLMS**)
- L2: Task Run-time (**TRS**)
- L1: Resource Access Layer

Abstractions & Building Blocks:

- L1: **RADICAL-SAGA** Distributed job submission & standard interface
- L2: **RADICAL-Pilot (RP)** Abstraction for Resource Management
- L3: **RADICAL-WLMS, Ensemble Toolkit**

Cross-layer: **RADICAL-Analytics**

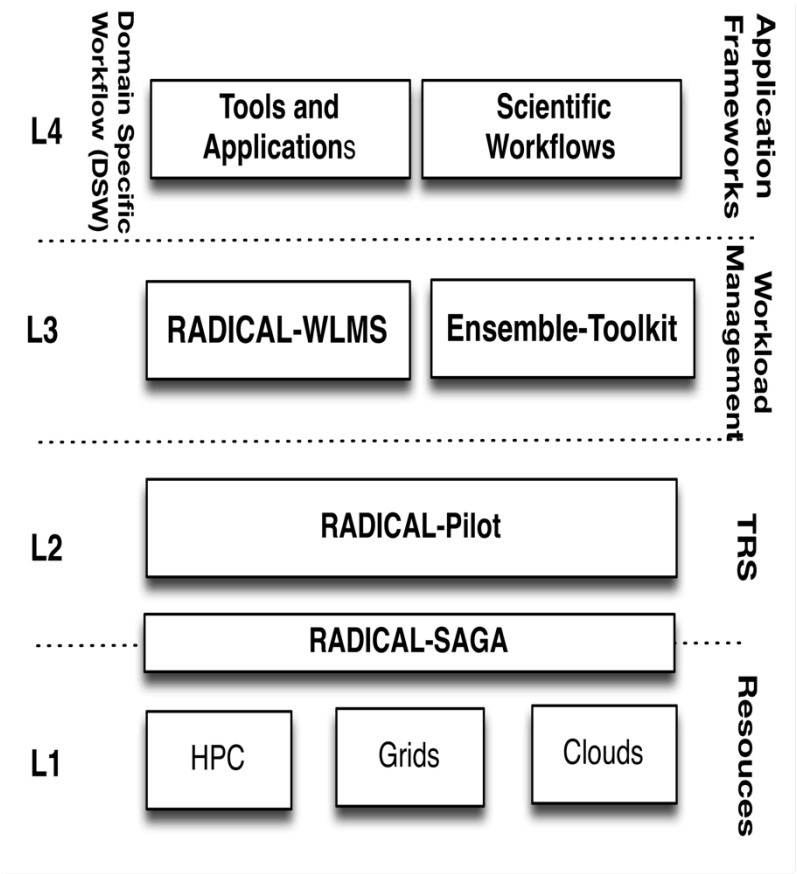

“Building Blocks” Approach to Workflow Systems

Workflows aren’t what they used to be!

Pervasive, sophisticated but no longer confined to “big science”

Diverse requirements, “design points”; unlikely “one size fits all”

Extend traditional focus from **end-users to workflow system/tool developers**!

Building Blocks (BB) permit workflow tools and applications

Two illustrative examples of building blocks:

WLMS (**AIMES Model**)

Pilot Job Systems (**P* Model**)

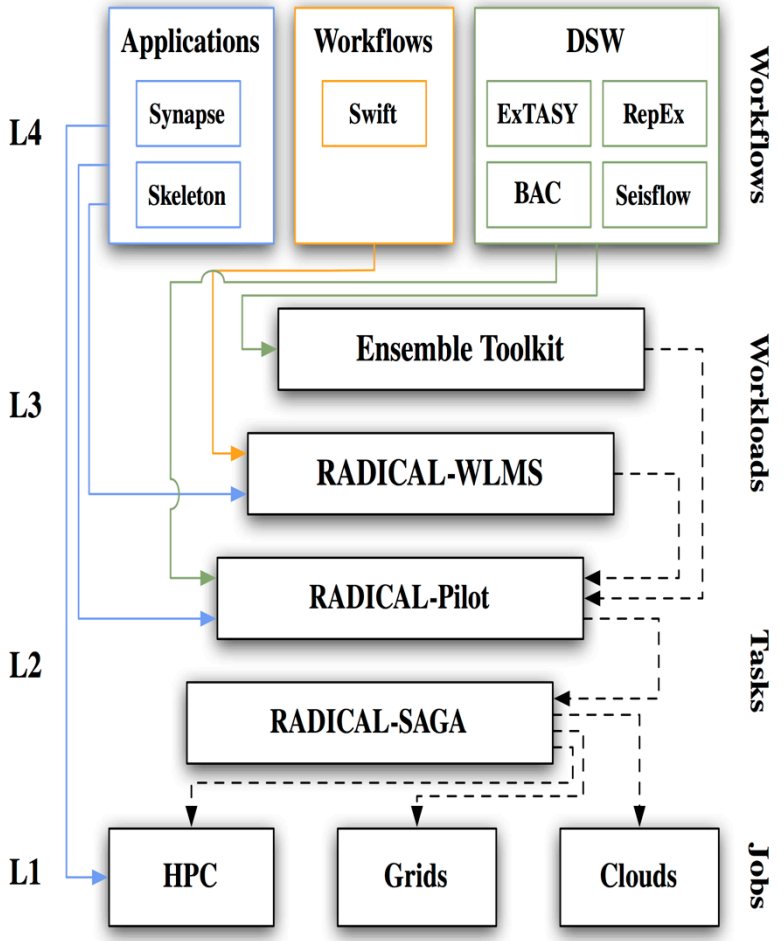

RUTGERS

RADICAL-Cybertools: Ensemble Toolkit (EnTK)

Ensemble-Member = task = **Execution Unit**

Multi-node, sub-node, MPI/non-MPI...

Simulation, Analysis, ..

AIMES Execution Model

Support for heterogeneous tasks

Adaptive Workload: Tasks and/or relations between tasks changes, or unknown *a priori*

Multiple dimensions of scalability:

Concurrency: O(10K) tasks

Task size: O(1)-O(1,000) cores

Launch: O(100+) tasks per second

Task duration: O(1)-O(10,000) seconds

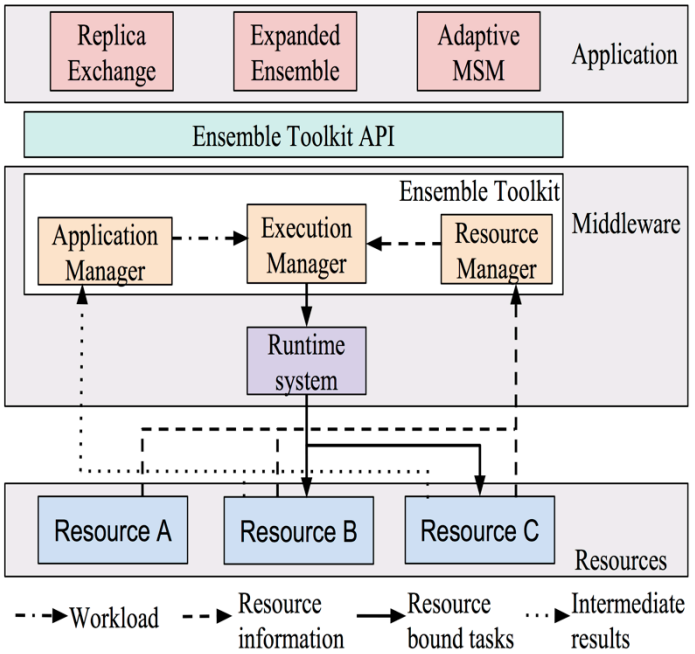

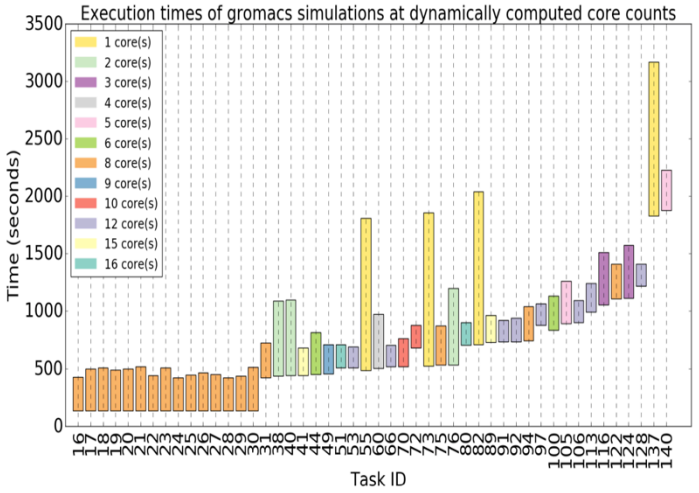

RUTGERS

HT-BAC: BAC + RADICAL-Cybertools

Design Challenge: How to (i) provide performance (ii) adequate functional extensibility, (iii) keep new software footprint.

Implementation: Use (i) functionally well-defined building blocks, (ii) provide well-defined interfaces and (iii) separate (performance and interoperability) from functionality.

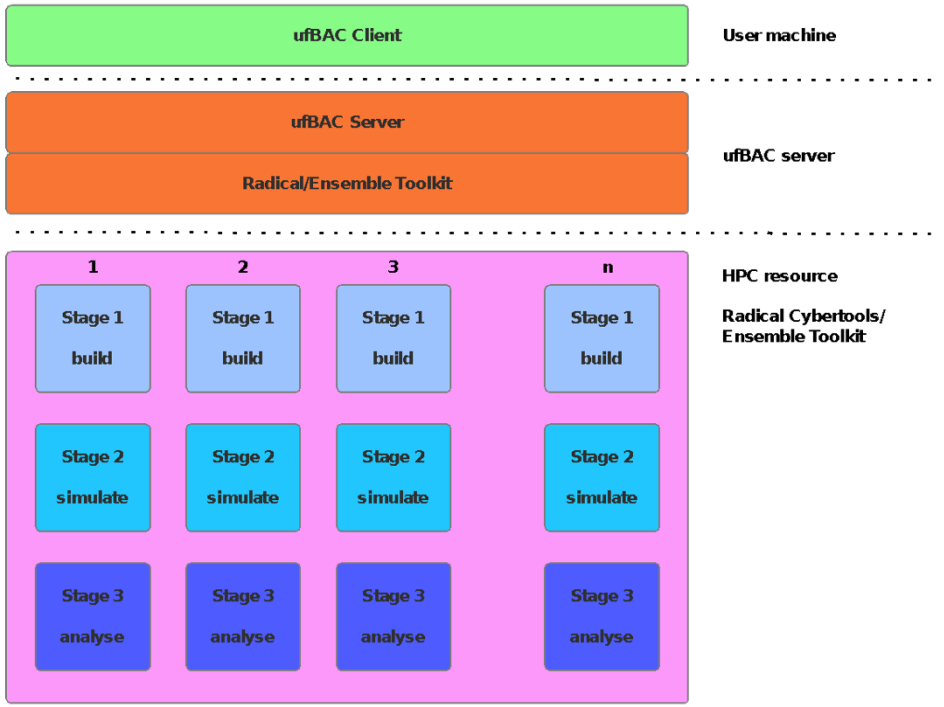

UCL

RUTGERS

App 2: Multiscale simulation of nanomaterials

From micro to meso ….

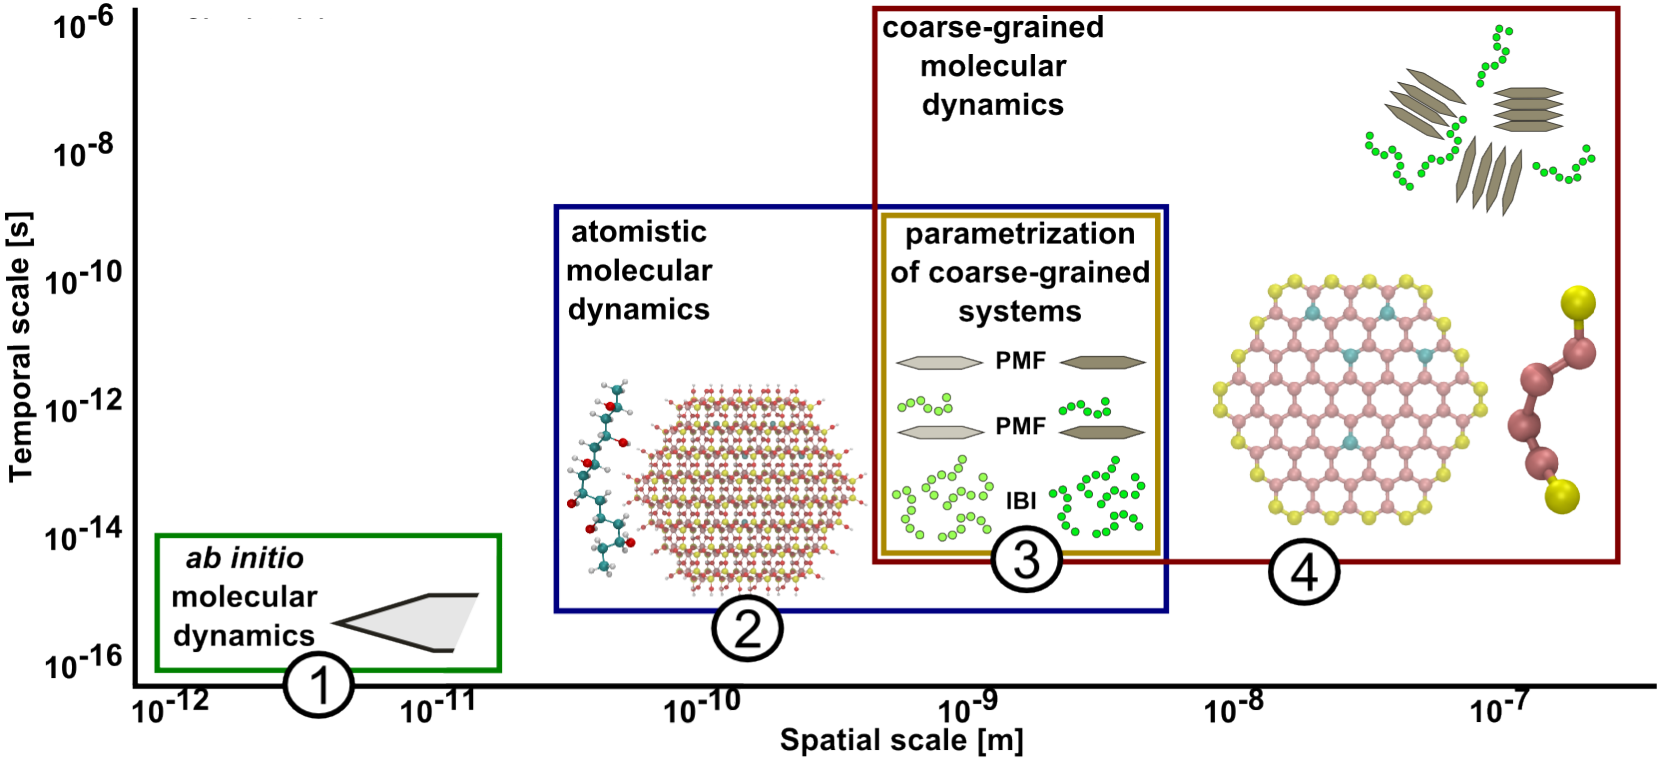

① CPMD, local cluster

② LAMMPS (AA), HECToR

③ LAMMPS (CG) BlueWonder, HECToR

④ LAMMPS (CG) BlueJoule, HECToR, ARCHER

Size of data generated: ~10TB . All data is stored on the RDF and on EUDAT (www.eudat.eu).

App 2: Multiscale simulation of nanomaterials
From micro to meso ….

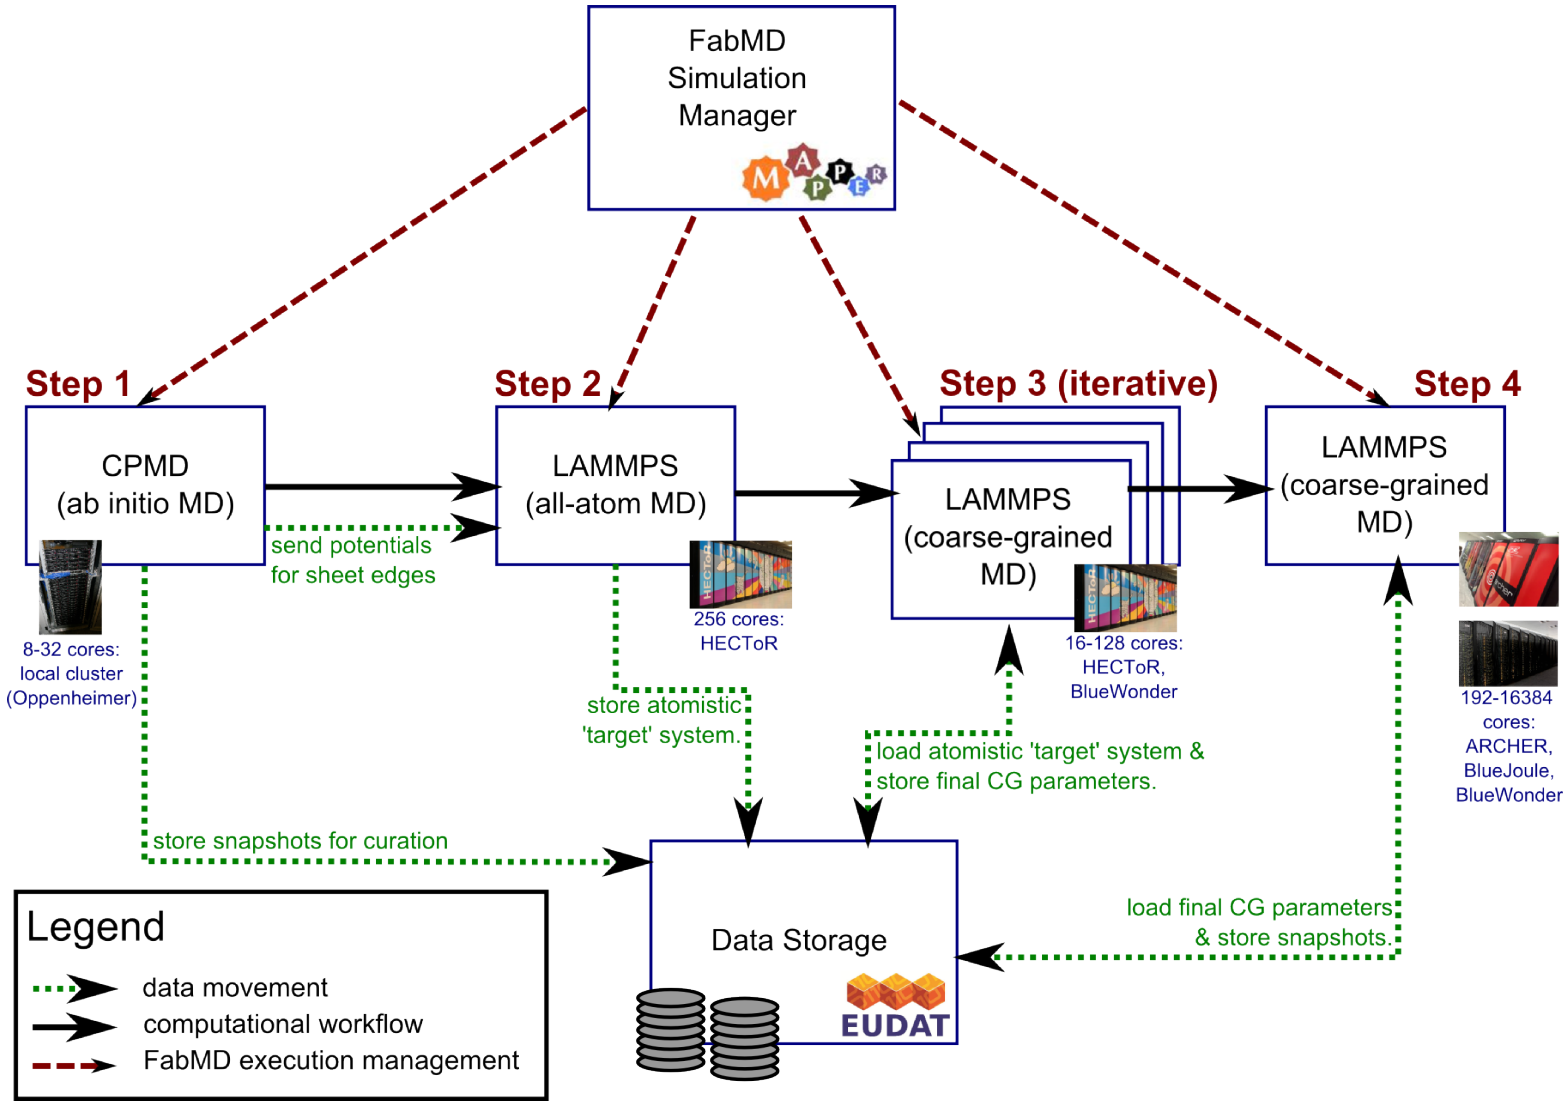

D. Groen, A. Bhati, J. Suter, J. Hetherington, S. Zasada, P. V. Coveney, ", *Comp Phy Commun*, **207**, 375–385 (2016)

App 2: Multiscale simulation of nanomaterials

From meso to macro …

- Material properties at the macroscopic scale **depend tightly** on the structure of the material and its evolution at **finer scales**
- Using **reduced models** to account for this dependency such as constitutive equations has **limited accuracy** and **stability** (e.g. when nonlinear, inelastic, anisotropic, or history-dependent behaviors)

↪ **Multi-scale simulations** replace reduced models with structural models of the material at finer scale

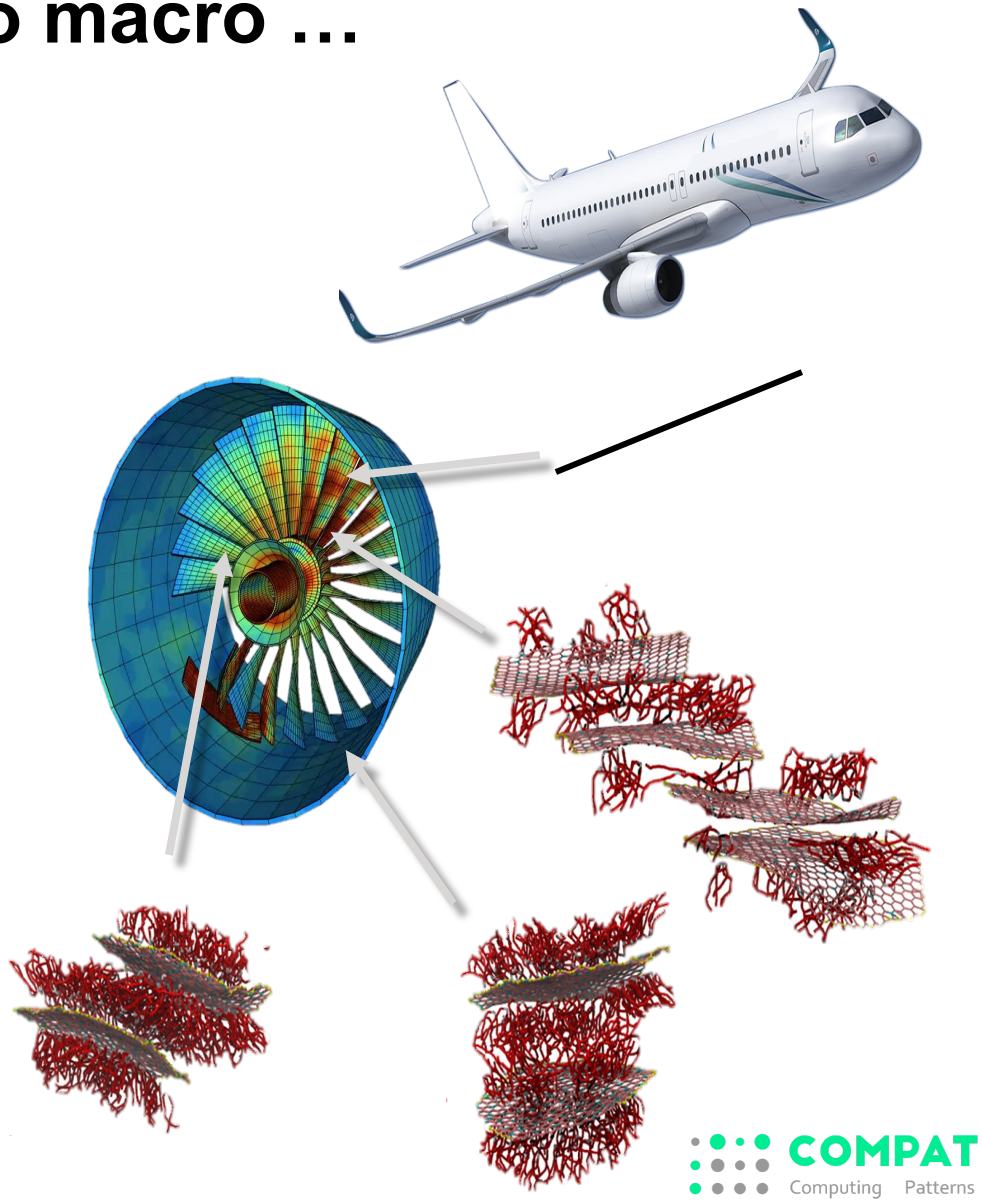

COMPAT
Computing Patterns for High Performance Multiscale Computing

Heterogeneous Multiscale Methods (HMM)

- Two or more structural models loosely coupled as part of a single simulation
  - partial information exchange in **chosen space-time locations**
  - information exchange **both ways** (bottom-up and top-down)
  - **discontinuity** of time and space discretization
    - faster **relaxation** times at lower scales $\rightarrow t_n \ll \Delta t_m$
    - **periodicity** in space $\rightarrow x_n \ll \Delta x_m$

- Implementation of a FE-MD HMM
  - Macroscale model based on continuum mechanics using the **finite element** library **Deal.II**
  - Nanoscale model based on **molecular dynamics** using **LAMMPS**
  - Exchange of **strains** (top-down) and **stresses** and **stiffness** (bottom-up)

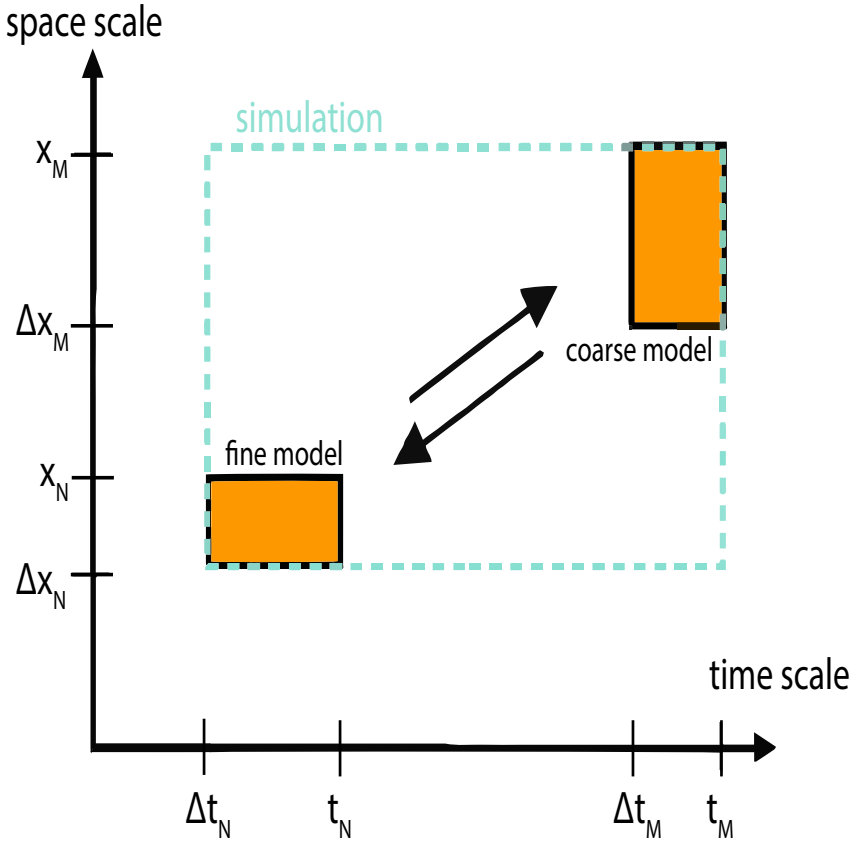

COMPAT
Computing Patterns for High Performance Multiscale Computing

Workflow of the Finite Element - Molecular Dynamics HMM Coupling Scheme

- Solving continuum mechanics **quasi-static equilibrium** using an **incremental iterative algorithm**
- Finite element computations are short and parallelised
- At a given time, nanoscale state updates can be run **independently**

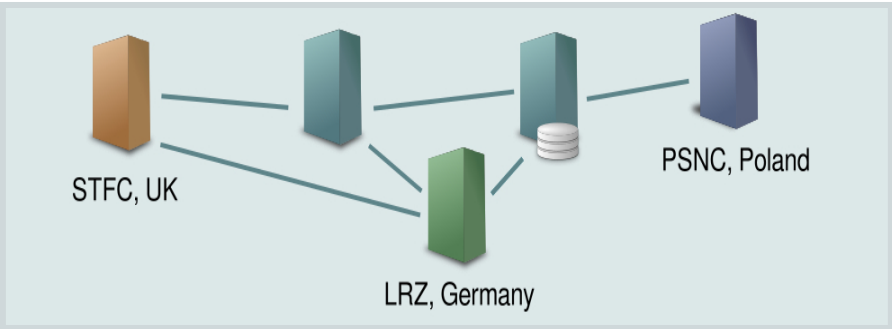

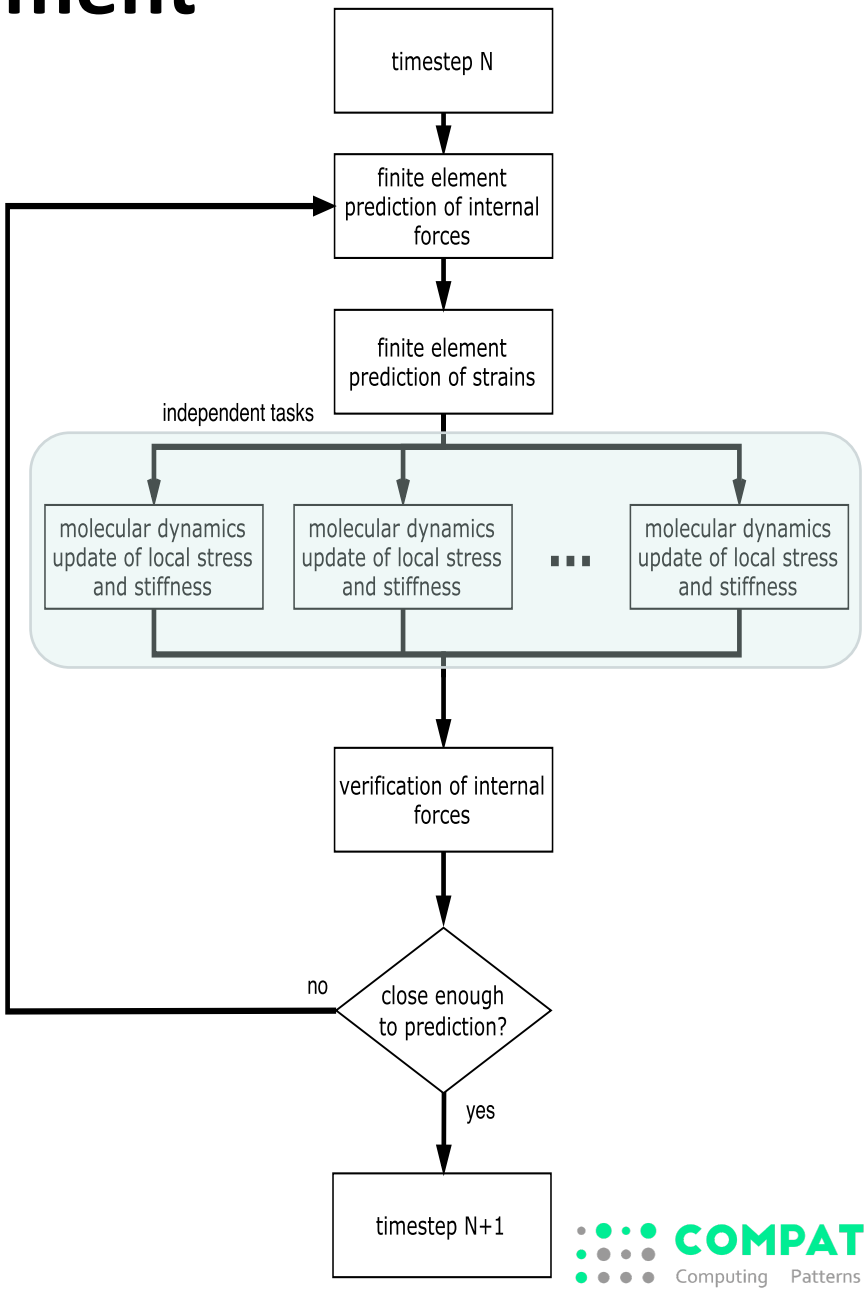

COMPAT
Computing Patterns for High Performance Multiscale Computing

The FE-MD HMM in the ComPat project

COMPAT
Computing Patterns for High Performance Multiscale Computing

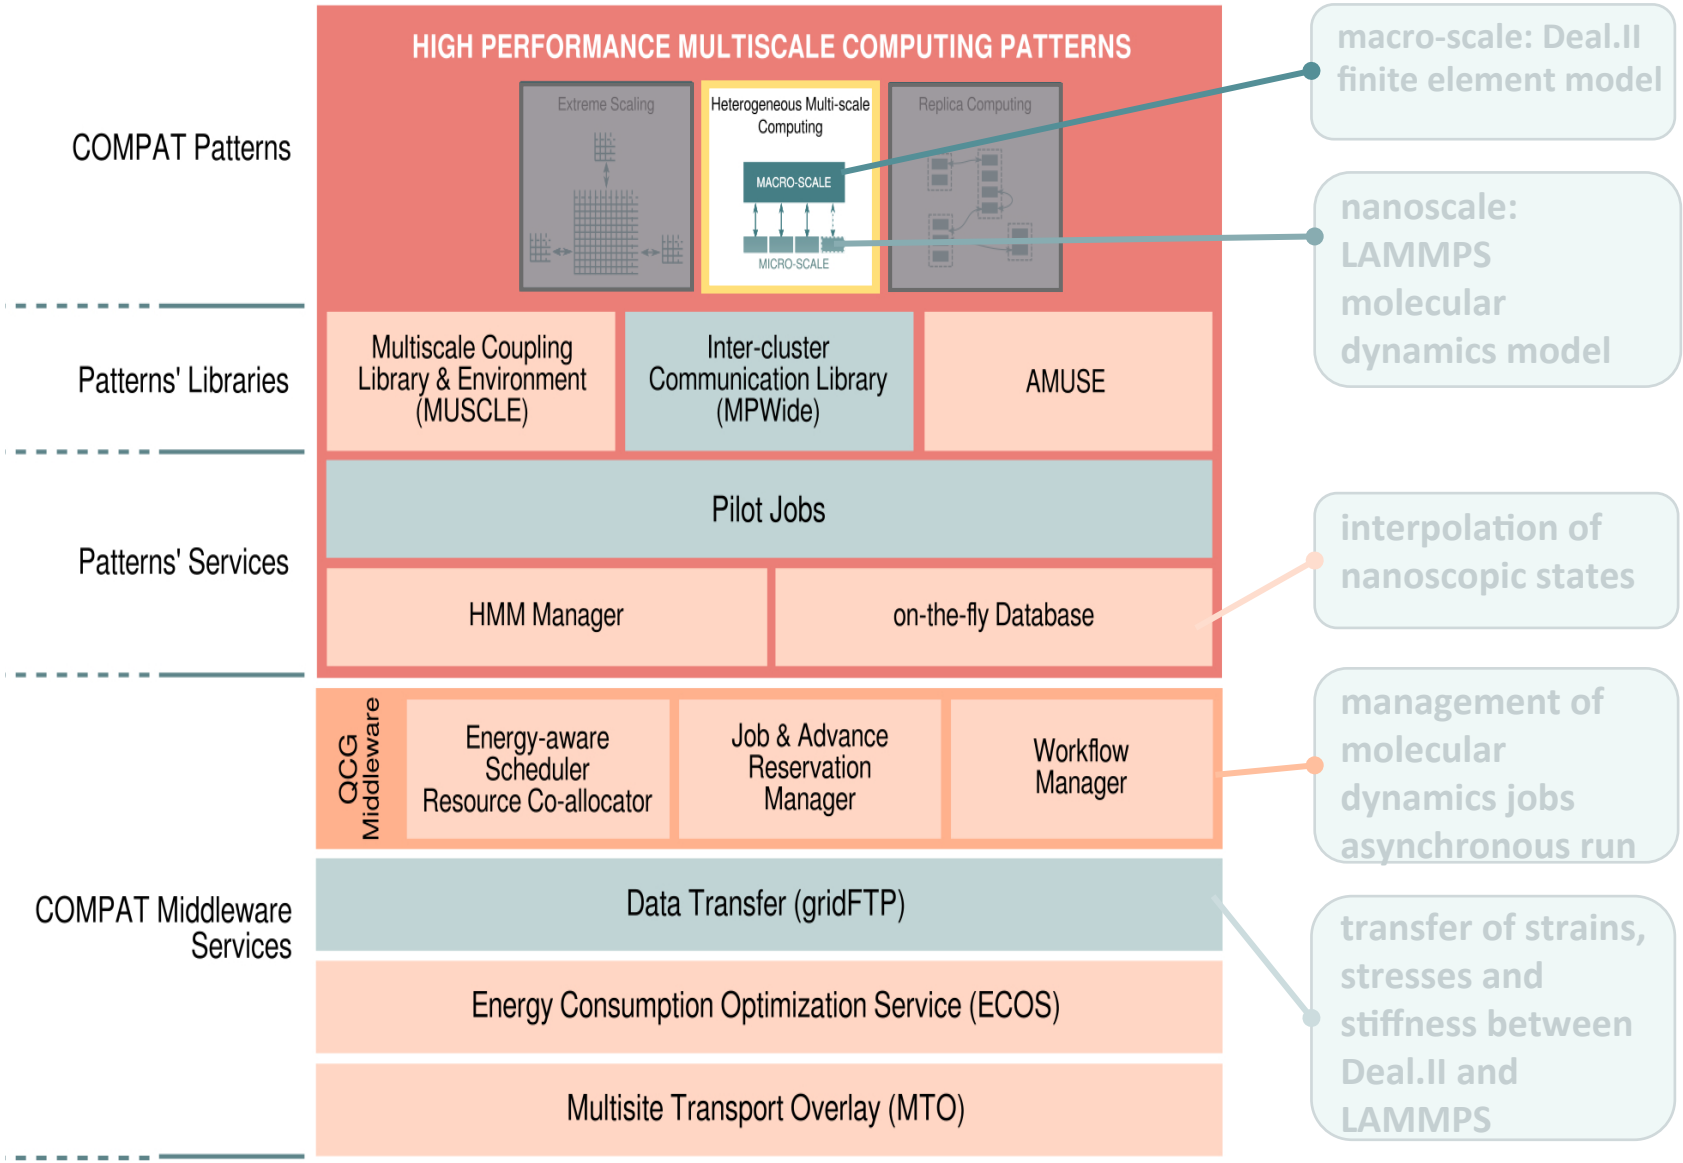

Conclusions

- Reliable, accurate and reproducible binding affinity ranking can be obtained only by ensemble simulations.
- **Scalable approach**: results in hours.

- ufBAC is a science gateway that hides the complexity of the workflow in a simple, web accessible application
- QCG Computing, RADICAL-Cybertools and Ensemble ToolKit (EnTK) are a realization of the Building Blocks approach to scalable workflows.
- Allows BAC developers to focus on implementing new ensemble based approaches independent of platform and performance.

- Heterogeneous multiscale materials simulations comprise complex workflows of operations.
- EU H2020 ComPat project e-infrastructure used to manage simulations and mutiscale modelling linkages, constituting a science gateway for multiscale computing.

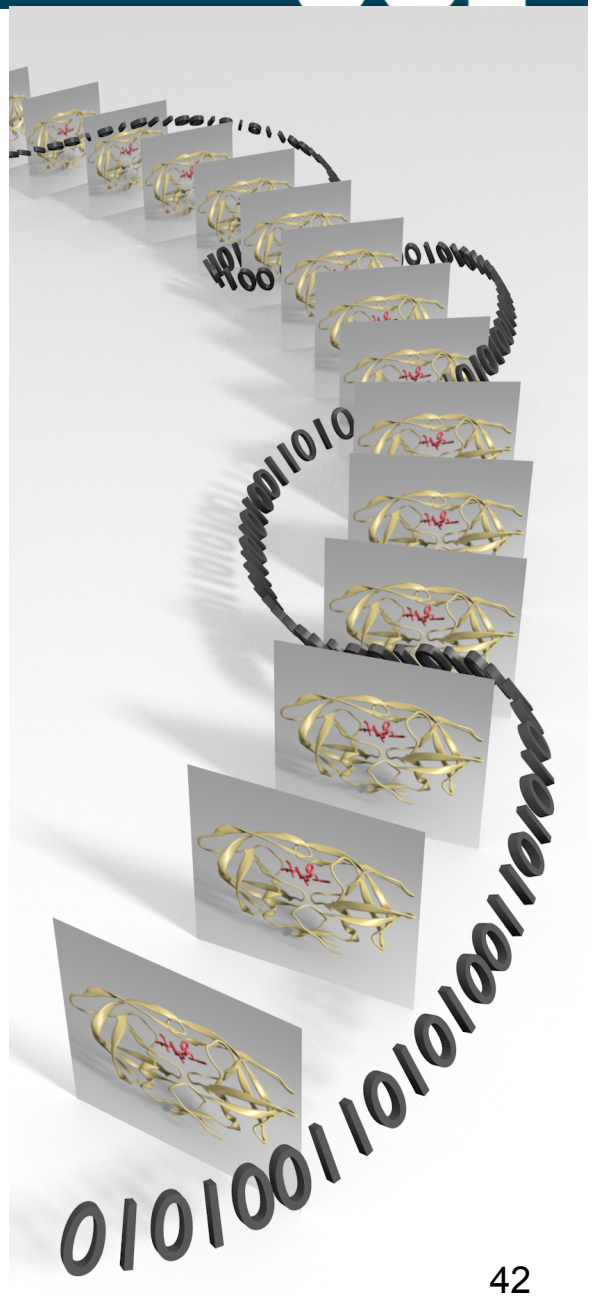

COMPAT
Computing Patterns for High Performance Multiscale Computing

Acknowledgements

UCL

- Agastya P. Bhati
- Shunzhou Wan
- David W. Wright
- Stefan J. Zasada

gsk

- Ian Wall
- Darren Green
- Paul Bamborough

Pfizer

- Sarah Skerratt
- Kiyoyuki Omoto
- Veerabahu Shanmugasundaram
- Sharan K. Bagal

COMPAT
Computing Patterns for High Performance Multiscale Computing

EUDAT

MRC Medical Research Council

EPSRC
Engineering and Physical Sciences Research Council

archer

epcc

Irz Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

PRACE

Hartree Centre
Science & Technology Facilities Council

الصندوق القطري لرعاية البحث العلمي
Qatar National Research Fund
Member of Qatar Foundation

CompBioMed