

Collaboratively Building Reusable Job Configurations for HPC

Jeremy Cohen

London e-Science Centre, Department of Computing, Imperial College London
jeremy.cohen@imperial.ac.uk

With thanks to:

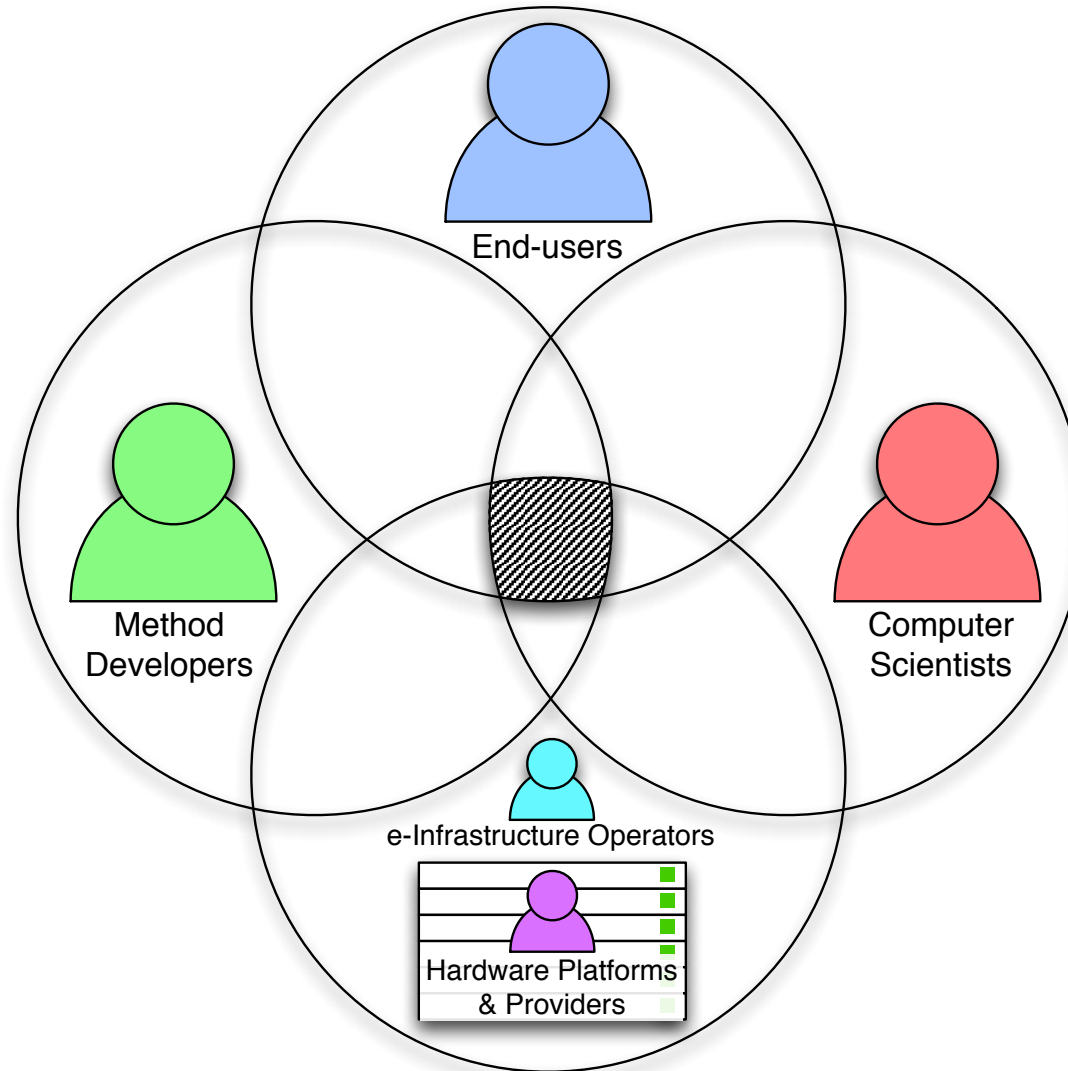
John Darlington, Chris Cantwell, David Moxey, Spencer Sherwin & Jeremy Nowell

PRISM Seminar, Thursday 26th March 2015

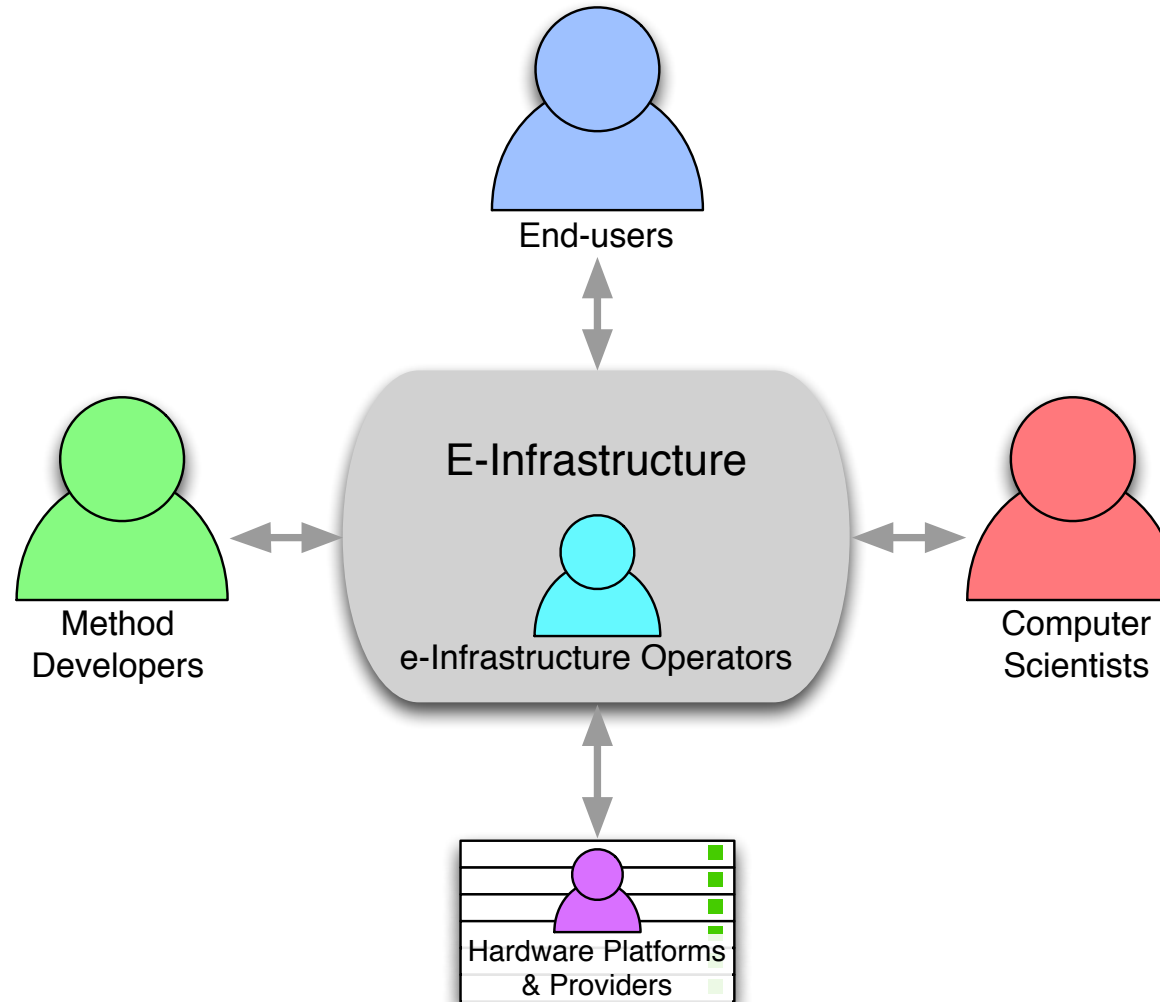
Reusable Job Configurations

- Provide a **high-level** approach for end-users to **configure** their **jobs**
- Address **complexity** of **configuration files** that stems from complex **software methods** and **heterogeneous hardware**
- Improve **usability** of **software methods** and their availability to **scientists/researchers** not from a computationally focused background

Integrated e-Infrastructure




A Decoupled e-Infrastructure



Abstraction versus Efficiency A Fundamental and Long-standing Problem

Simplicity and Ease of Use



Abstractions

Coordination Forms = Control Abstraction

Higher-order, functions as arguments

Abstract Components

=

Data Processing Abstractions

First-order, data as arguments

- Allows automated selection of optimal implementations
- Metadata plays a key role in enabling abstract => concrete mapping – maintains information

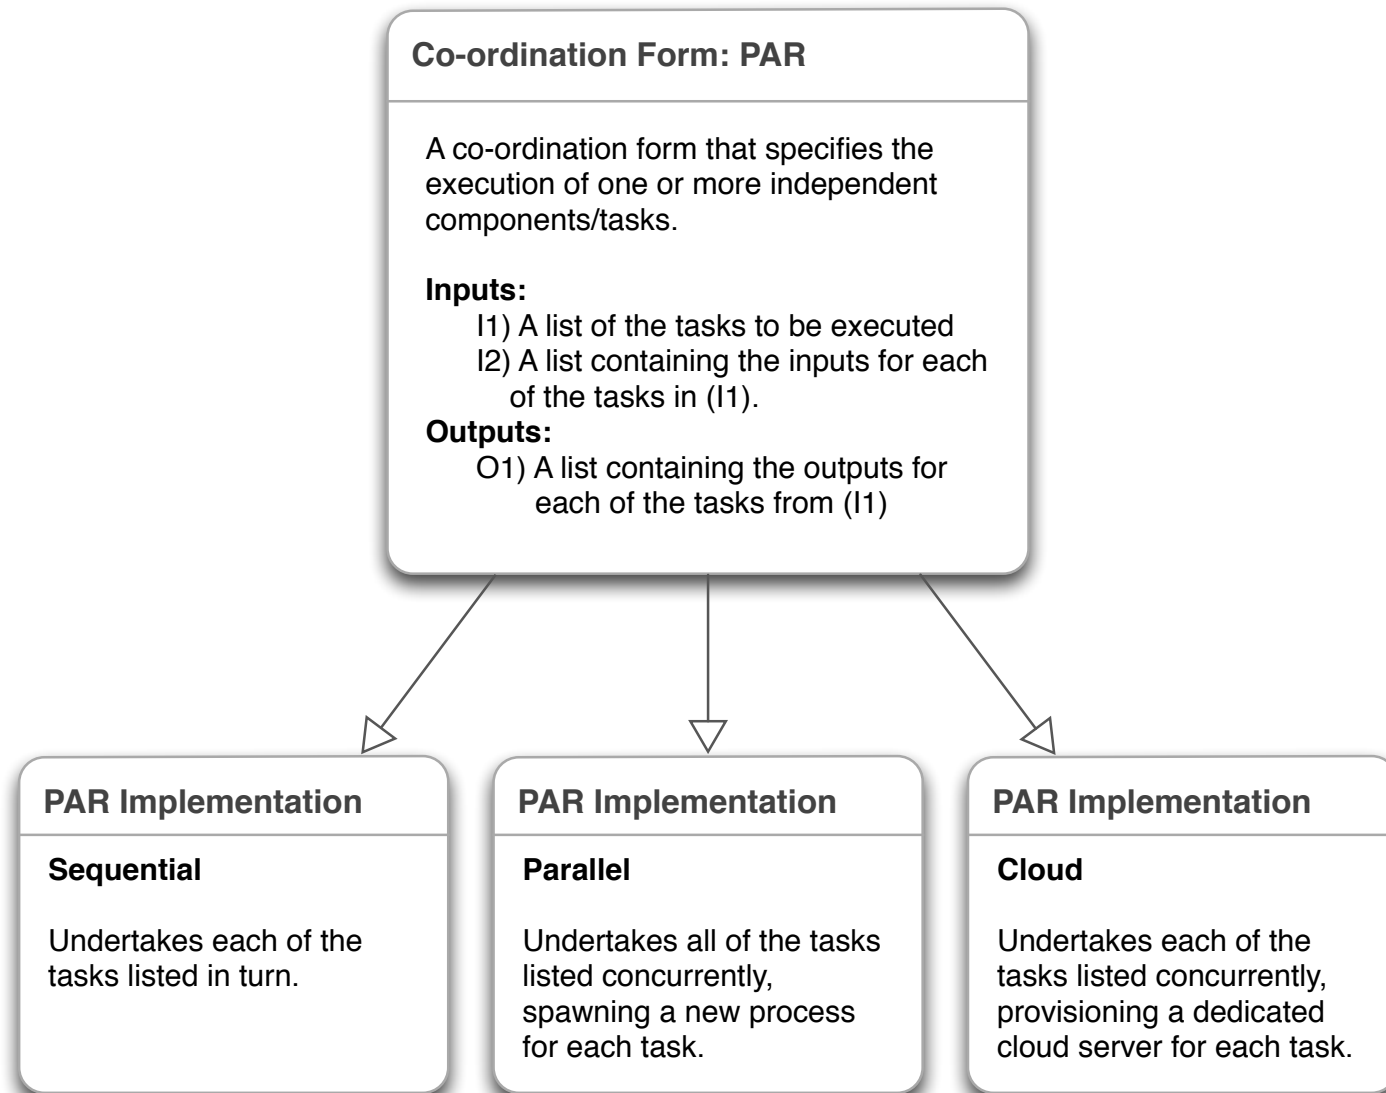
Coordination Forms

- A functional/mathematical approach to job specification
- Referentially transparent, Church-Rosser property
- Based on work by Darlington, et al.

J. Darlington, Y. Guo, H. W. To and J. Yang. Functional skeletons for parallel coordination. In proceedings of EURO-PAR '95 Parallel Processing, LNCS 966/1995, p. 55-66, 1995. Springer Berlin/Heidelberg

- Can have multiple implementations – e.g. sequential/parallel
- Compositions of coordination forms can be used to describe application flow

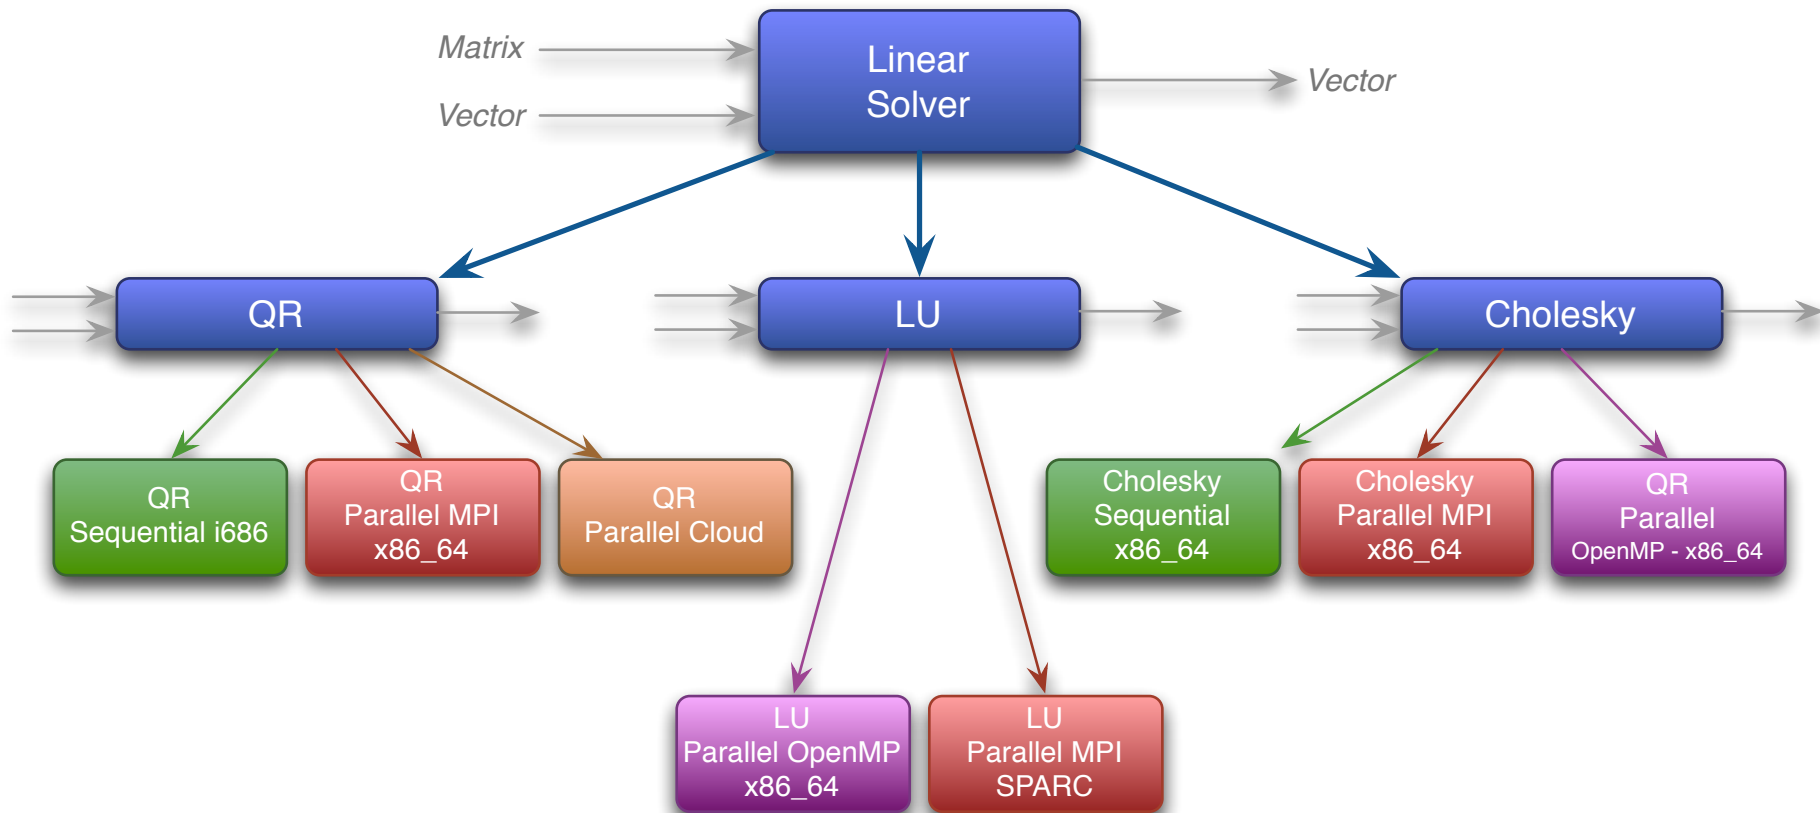
Alternative implementations: Coordination Forms



Software Components

- Granularity varies
 - **Fine-grained**: small libraries, individual functions, command-line tools
 - **Coarse-grained**: Whole application!
- **Abstract** – metadata wrapper, no implementation
- **Concrete** – Runnable component, metadata + implementation
- Components can have **multiple implementations**

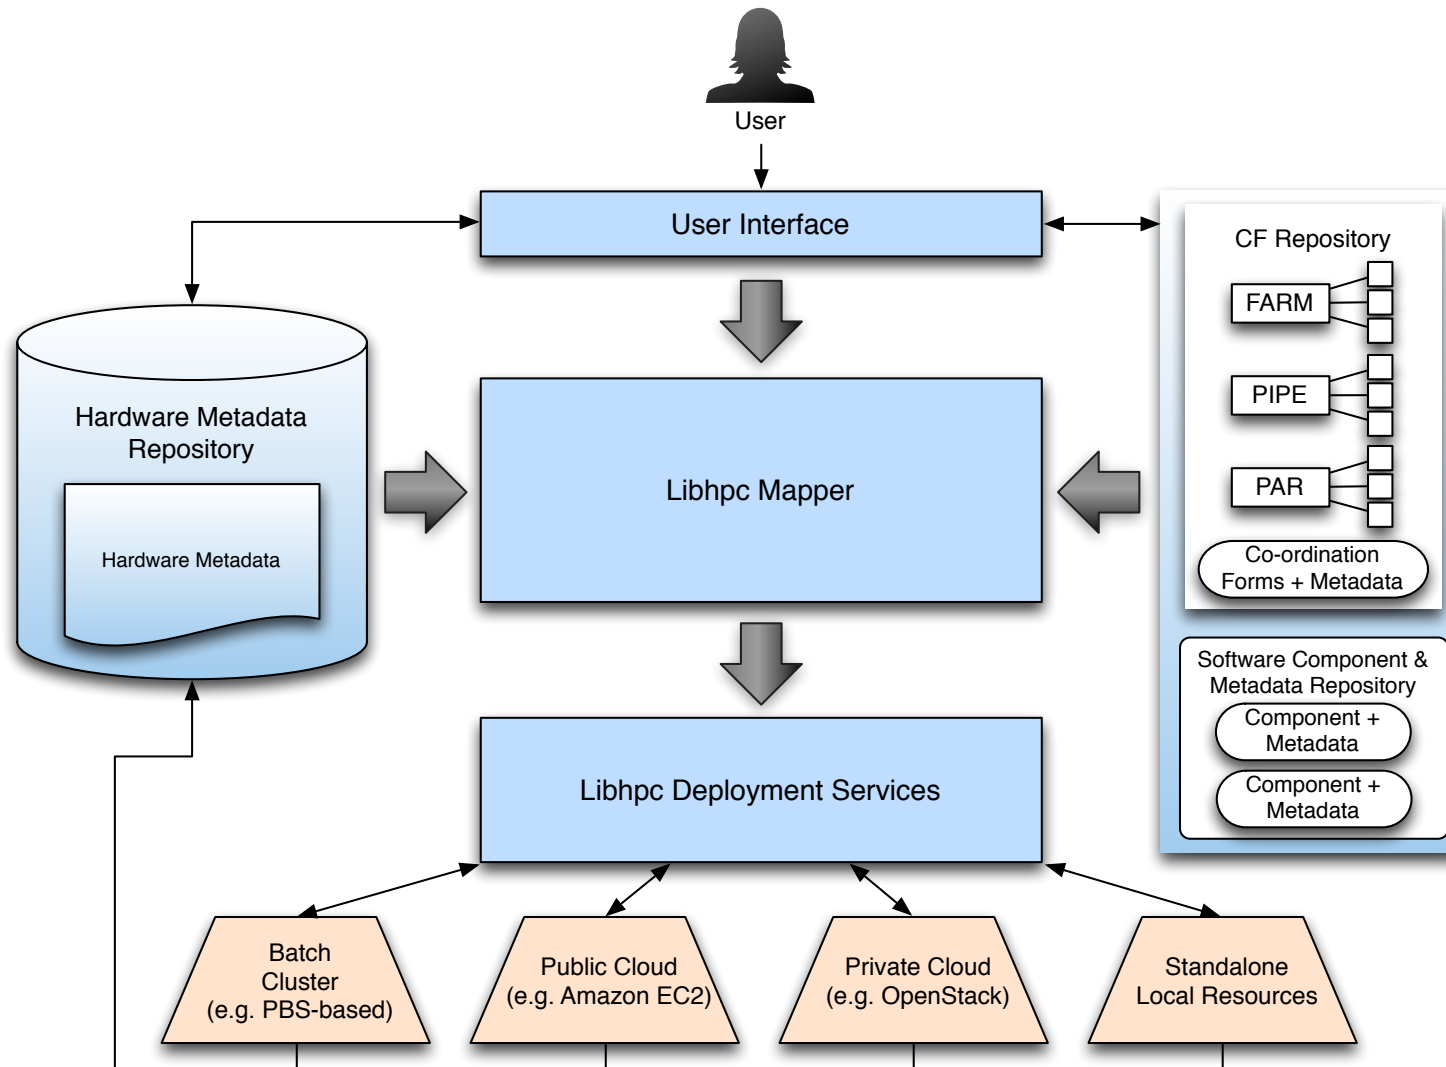
Alternative implementations: Components



Libhpc Projects

- Libhpc 2 runs to end October 2015
 - Builds on Libhpc 1 which ran from July 11 -> Jun 13
 - Developing framework model and a range of associated tools, services and demonstrators
-
- Imperial College London
 - Dept of Computing (LeSC/SCG)
 - Dept of Aeronautics
 - CISBIO / Bioinformatics Support Service
 - Epidemiology, School of Public Health
 - University of Edinburgh
 - Edinburgh Parallel Computing Centre (EPCC)

Libhpc Architecture






Templates and Profiles

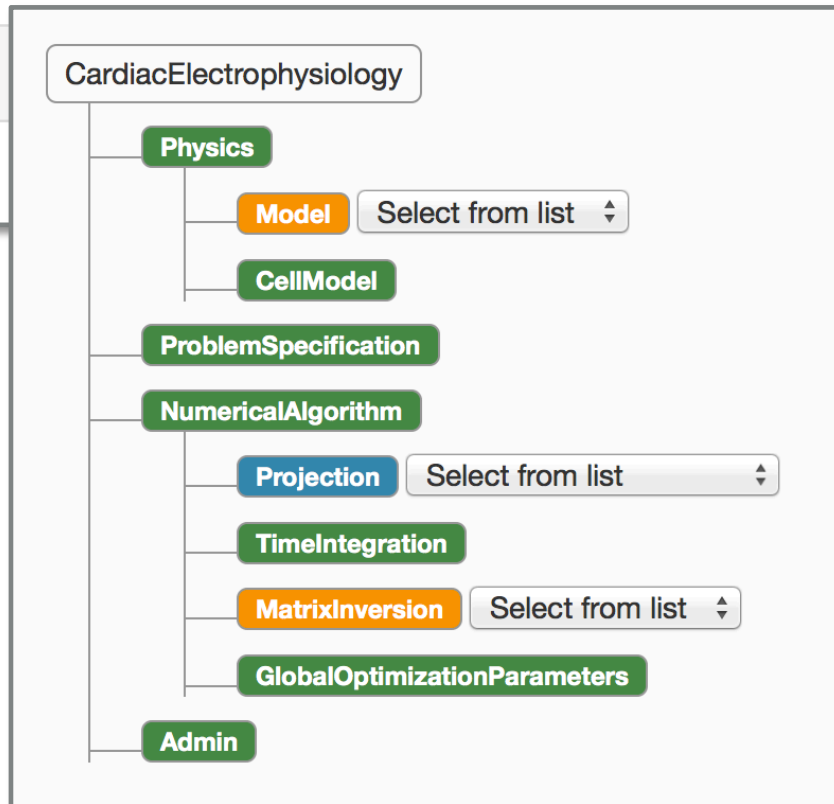
Templates & Profiles

- Libhpc software parameter templates
 - Represent an application's possible **configuration parameters/decisions**
 - **Tree structure** with semantic parameter grouping
 - Defined using **XML Schema**
 - Does **not** contain **values for** any of the **specified parameters**
 - Includes **validation** and **documentation** metadata

Templates & Profiles

Templates

Solver	Created by	Create profile
Caridac Electrophysiology	jhc02	
Incompressible Navier Stokes		
Compressible Flow Solver		



Templates & Profiles

- Libhpc profiles
 - Provides an **instantiation of a template's** parameters
 - **XML** document – profile structure can be **validated against template**
 - May be:
 - **Partial**: contains a subset of the required values from template
 - **Complete**: Contains a full set of required values and can be used to run a job

Templates & Profiles

Profiles

Name	Solver	Valid	Editable	Created by	Created	
Default profile - Cardiac Electrophysiology	CES	✓	✓	jhc02	25 Jul 2014, 4:30 p.m.	<input type="button" value="Edit"/>
Default profile - Incompressible Navier Stokes	INS	✓	✓	jhc02	12 Aug 2014, 25 p.m.	<input type="button" value="Edit"/>
test profile					0 Oct 2014, 11 p.m.	<input type="button" value="Edit"/>

```
<?xml version="1.0" encoding="utf-8"?>
<?xml-stylesheet type="text/xsl"
  href="/data/nekkcloud/src/main/
  resources/Transform/LibhpcNektarToTrueNektar.xsl"?>
<IncompressibleNavierStokes>
  <Physics>
    <KinematicViscosity>1</KinematicViscosity>
  </Physics>
  <ProblemSpecification>
    <SolutionMethod>VelocityCorrectionScheme</SolutionMethod>
    <EvolutionOperator>Adjoint</EvolutionOperator>
    <Geometry>CylinderGeometry.xml</Geometry>
    <InitialConditions>
      <Constant>-81</Constant>
    </InitialConditions>
    <Expansion>
      <PolynomialOrder>7</PolynomialOrder>
      <BasisType>MODIFIED</BasisType>
    </Expansion>
  </ProblemSpecification>
  <NumericalAlgorithm>
    <Projection>ContinuousGalerkin</Projection>
    <TimeIntegration>
```

Templates & Profiles

- Templates **defined** and built **by developers / domain experts**
- Partial **profiles** may be **saved; extended** by **different entities**
- **End-users** may be provided with an almost **complete profile** and then finalise this **to run** their **job(s)**
- Helps to **decouple interactions** required for configuration of **complex applications** for **heterogeneous resources**

Examples and Demos

Bioinformatics: Genome Read Pre-Processing/Mapping

Input files –

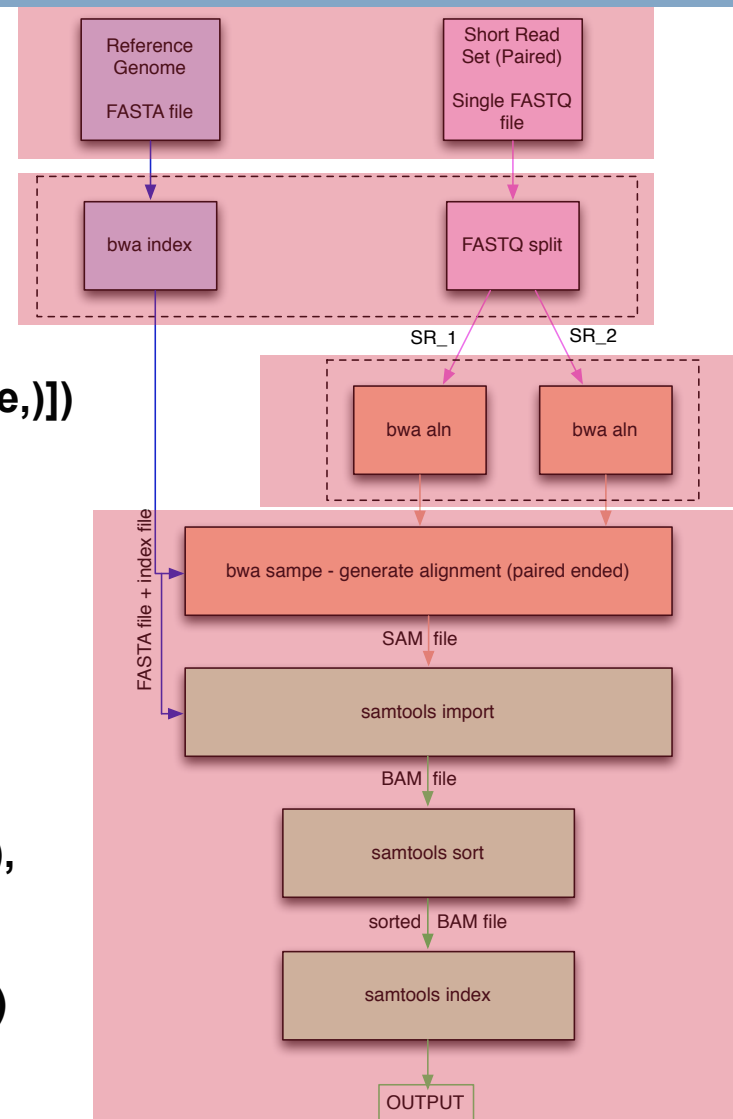
Reference Genome – FASTA file

Reads from sequencing machine - FASTQ

```
((sr1, sr2), u) = PAR([fastq_split, bwa_index],
  [(short_read_file, None, None), (ref_genome_file,)])
```

```
(v, w) = PAR([bwa_aln, bwa_aln],
  [(ref_genome_file, sr1, None),
  (ref_genome_file, sr2, None)])
```

```
result = PIPE([samtools_index, samtools_sort,
  (samtools_import, ref_genome_file),
  bwa_sampe],
  [ref_genome_file, [v,w], [sr1, sr2], None])
```



Nekkloud: Simplifying Access to Nektar++

The screenshot displays the Nekloud web interface. On the left, a configuration panel for 'IncompressibleNavierStokes' is visible, with sections for Physics, ProblemSpecification, NumericalAlgorithm, and Admin. The ProblemSpecification section includes dropdowns for SolutionMethod (VelocityCorrectionScheme) and AdvectionForm (Convective), and input fields for FinalTime (10) and TimeStep (0.01). The main content area features the Nekloud logo, a description of the service, and a 'Sign up' button. A 'Sign in' form is also present, with fields for Username and Password, and a 'Remember me' checkbox. On the right, a 'Job snapshots' table lists two snapshots with their IDs, checkpoint counts, and times. Below the table is an 'Output files' section listing files like 'job-8f7ad7f6_output.tar.gz' and 'job-8f7ad7f6_movie.avi' with their respective sizes.

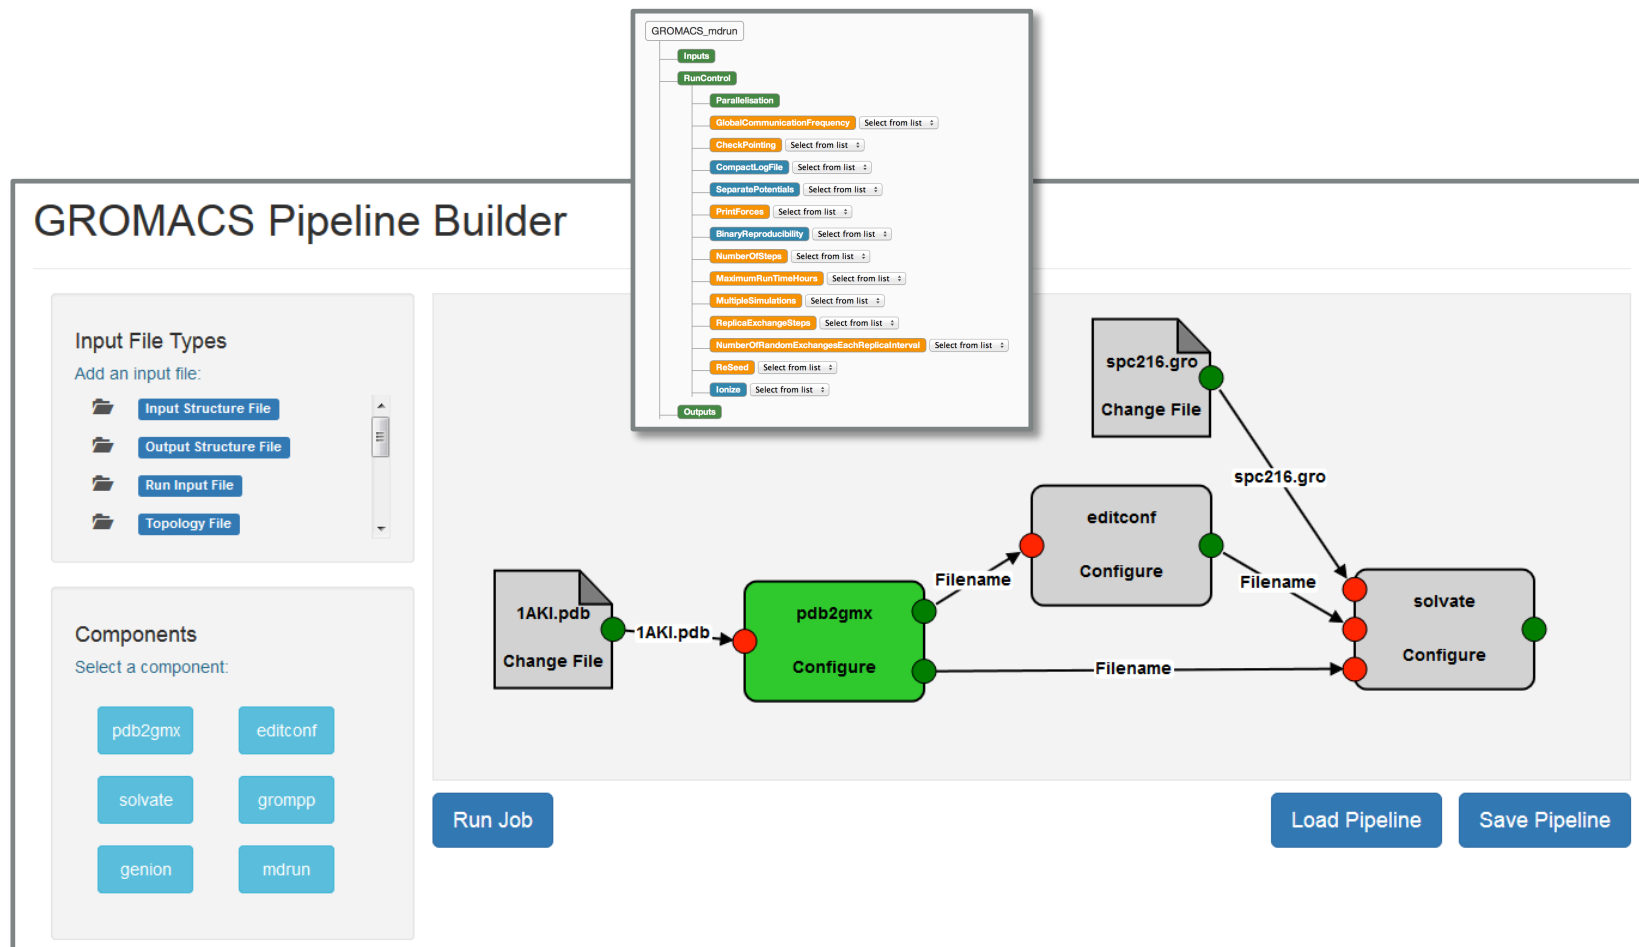
Preview	Snapshot ID	Job checkpoint	Snapshot time
	snap-ca6ff15c	3	14:01:39, 18 Mar 2015
	snap-9297def1	8	14:01:51, 18 Mar 2015

Name	Size
job-8f7ad7f6_output.tar.gz	15.7 MB
job-8f7ad7f6_output_combined.tar.gz	15.7 MB
job-8f7ad7f6_movie.avi	40.8 KB

For more info see: J. Cohen, D. Moxey, C. Cantwell, et al., "Nekkloud: A software environment for high-order finite element analysis on clusters and clouds," IEEE Cluster 2013, Sep 2013, Indianapolis, IN, USA. DOI: 10.1109/CLUSTER.2013.6702616

Molecular Dynamics: GROMACS

- GROMACS is a high performance molecular dynamics package providing a range of MD algorithms – <http://www.gromacs.org>
- Ideal example of an application that includes both tightly coupled parallel processes but also a higher-level pipeline of tools



Nekkloud Demo

Thanks & Acknowledgements

Thank You

Questions?

Acknowledgements:

- Nektar++ team – <http://www.nektar.info/wiki/Latest/Team> – including Chris Cantwell, David Moxey and Spencer Sherwin
- London e-Science Centre – John Darlington, Peter Austing
- EPCC – Jeremy Nowell, Xu Guo;
- Bioinformatics Support Service, Imperial College London – Sarah Butcher, James Abbott and Filippo Mortari
- Additional thanks to members of the above groups who have provided content for this presentation