

High End Computing at Cardiff University


Focus on Campus Grids
James Osborne



Contents

- High End Computing Spectrum
- Facilities at Cardiff
- Condor at Cardiff
- Success Stories
- High End Computing Futures
- Questions

High End Computing Spectrum




The HEC Spectrum

HPC

Tightly Coupled

HTC

Loosely Coupled



Supercomputers

NUMA Machines

£ Million+

Large Clusters

SMP

£ H Thousand

£ Million

Small Clusters

Campus Grids


£ Thousand

£ H Thousand

The HPC End

HPC

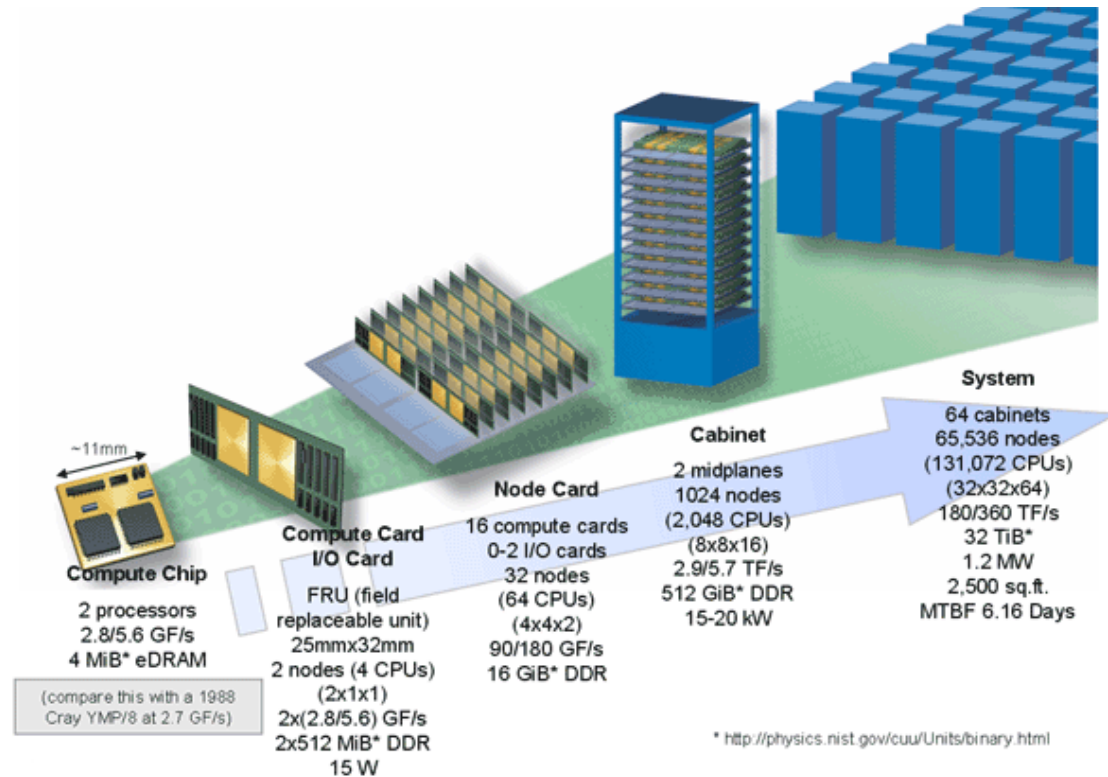
Tightly Coupled



Supercomputers

Bluegene L


131,072 CPUs



The HTC End



HTC
Loosely Coupled




Campus Grids
Condor@Cardiff
600+ CPUs

Facilities at Cardiff



Facilities at Cardiff - Helix

Large Clusters
Helix




200 CPUs



Facilities at Cardiff - SGI

Small Clusters
SGI Origin 300




32 CPUs



Facilities at Cardiff - Condor

Campus Grids
Condor@Cardiff



600+ CPUs



Condor at Cardiff

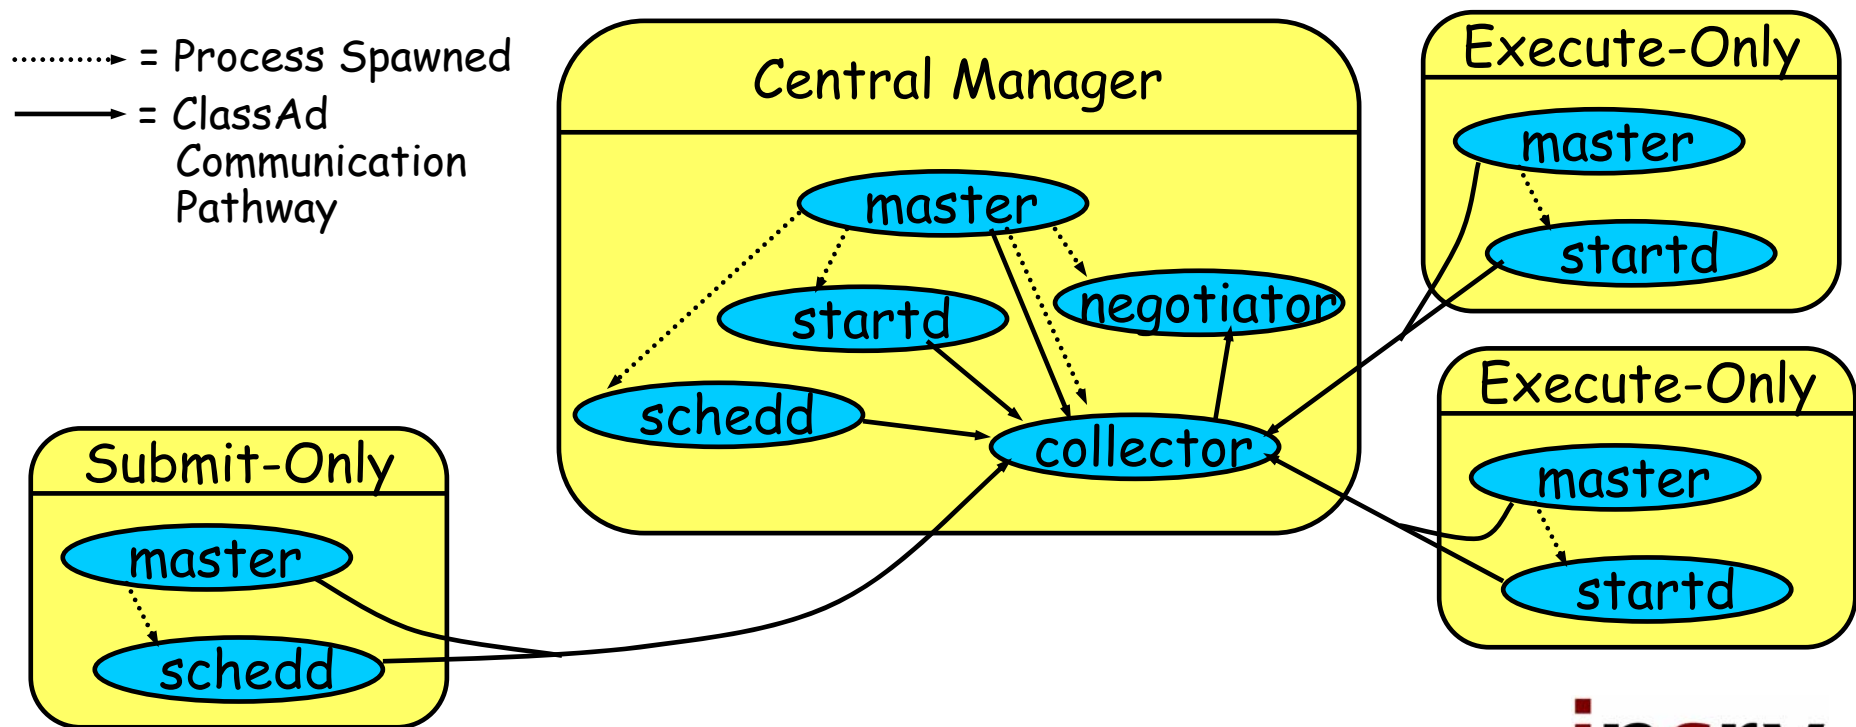


What is Condor ?

- Condor is a software system that creates a High-Throughput Computing (HTC) environment
- Condor effectively utilizes the computing power of workstations that communicate over a network
- Condor's power comes from the ability to effectively harness resources under distributed ownership

What is a Condor Pool ?

- A pool is a collection of workstations that communicate over a network



What is a Condor Job ?

- A command line windows executable
 - ▶ All files in a self-contained directory structure
 - ▶ Condor runs jobs in a sandbox *..lexecute\...*
 - ▶ Condor runs jobs as user *condor-reuse-vm1*
- One or more input files
- One or more output files
- A submit script
 - ▶ One or more logs – useful for debugging

What Goes In A Submit Script ?

- Running myprog 100 times
universe = vanilla
executable = myprog.exe
input = myin.\$(PROCESS)
output = myout.\$(PROCESS)
error = myerr.\$(PROCESS)
queue 100

What Else Can Go In ?

root_dir = c:\mydirectory

transfer_files = ALWAYS

transfer_input_files = \$(ROOT_DIR)\afile.txt

transfer_output_files = \$(ROOT_DIR)\afile.txt

log = mylog.\$(PROCESS)

notification = NEVER | ERROR

arguments = -arg1 -arg2

What Else Can Go In ?

requirements =

OpSys == "WINNT51"

Machine == "hostname.cf.ac.uk"

How Do I Submit A Job ?

- In the first instance by sending all your files to condor@cardiff.ac.uk to allow us to tailor your jobs to our environment
- In time by seeking permission to submit your own jobs to condor@cardiff.ac.uk to allow us to enable your workstation as a submit host
 - ▶ Currently requires IP address change

How Do I Submit A Job ?

- Submitting your job

```
condor_submit myscript.sub
```

- Checking your job's progress

```
condor_q
```

- Checking the pool

```
condor_status
```

Terms of Use

- Any local researcher can use the campus grid on the proviso that they...
 - ▶ write a short summary of their research that we can use to publicise their use of the campus grid
 - ▶ provide references to journal articles and conference proceedings containing appropriate acknowledgements

Success Stories



Prof Tim Wess

- OPTOM
- X-Ray Diffraction
- Determine shape of molecules
- Time on a single workstation = 2-3 Days
- Time on the campus grid = 2-3 Hours
- Speed-up factor of ~20

Prof Tim Wess

- “This capability provides the final link in the chain that Cardiff has established to solve macromolecular structures”
- “Our involvement with synchrotron sources such as DIAMOND ... and the residence of CCP 13 ... ensures that we are well placed to be in the vanguard of structure determination”

Soyeon Lee

- CARBS
- Montecarlo Simulation
- 20,000 parameters for 90 different models
- Time on a single workstation = 42 Days
- Time on the campus grid = 2 Days
- Speed-up factor of ~20

Dr Kevin Ashelford

- BIOSI
- Distributed Search
- Identify corrupt records in a DNA database
- Time on a single workstation = 2.4 Years
- Time on the campus grid = 2.6 Weeks
- Speed-up factor of ~50

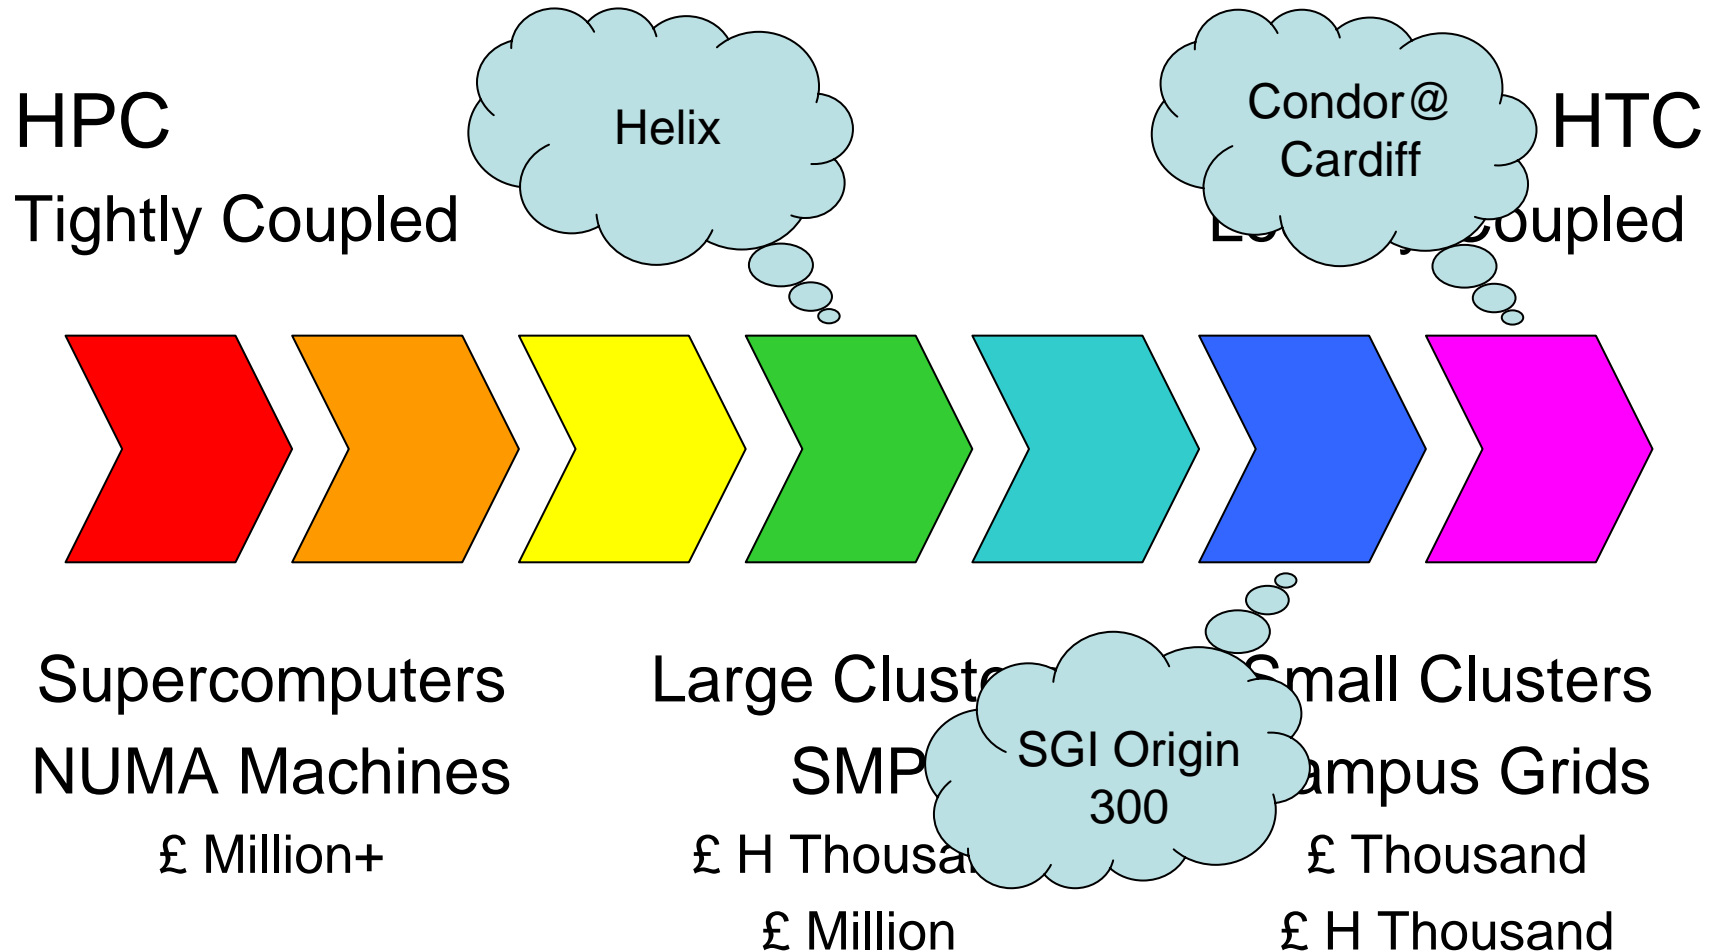
Dr Kevin Ashelford

- “This is a significant contribution to microbial research and will hopefully be the required impetus for the world-wide research community to improve current methods”

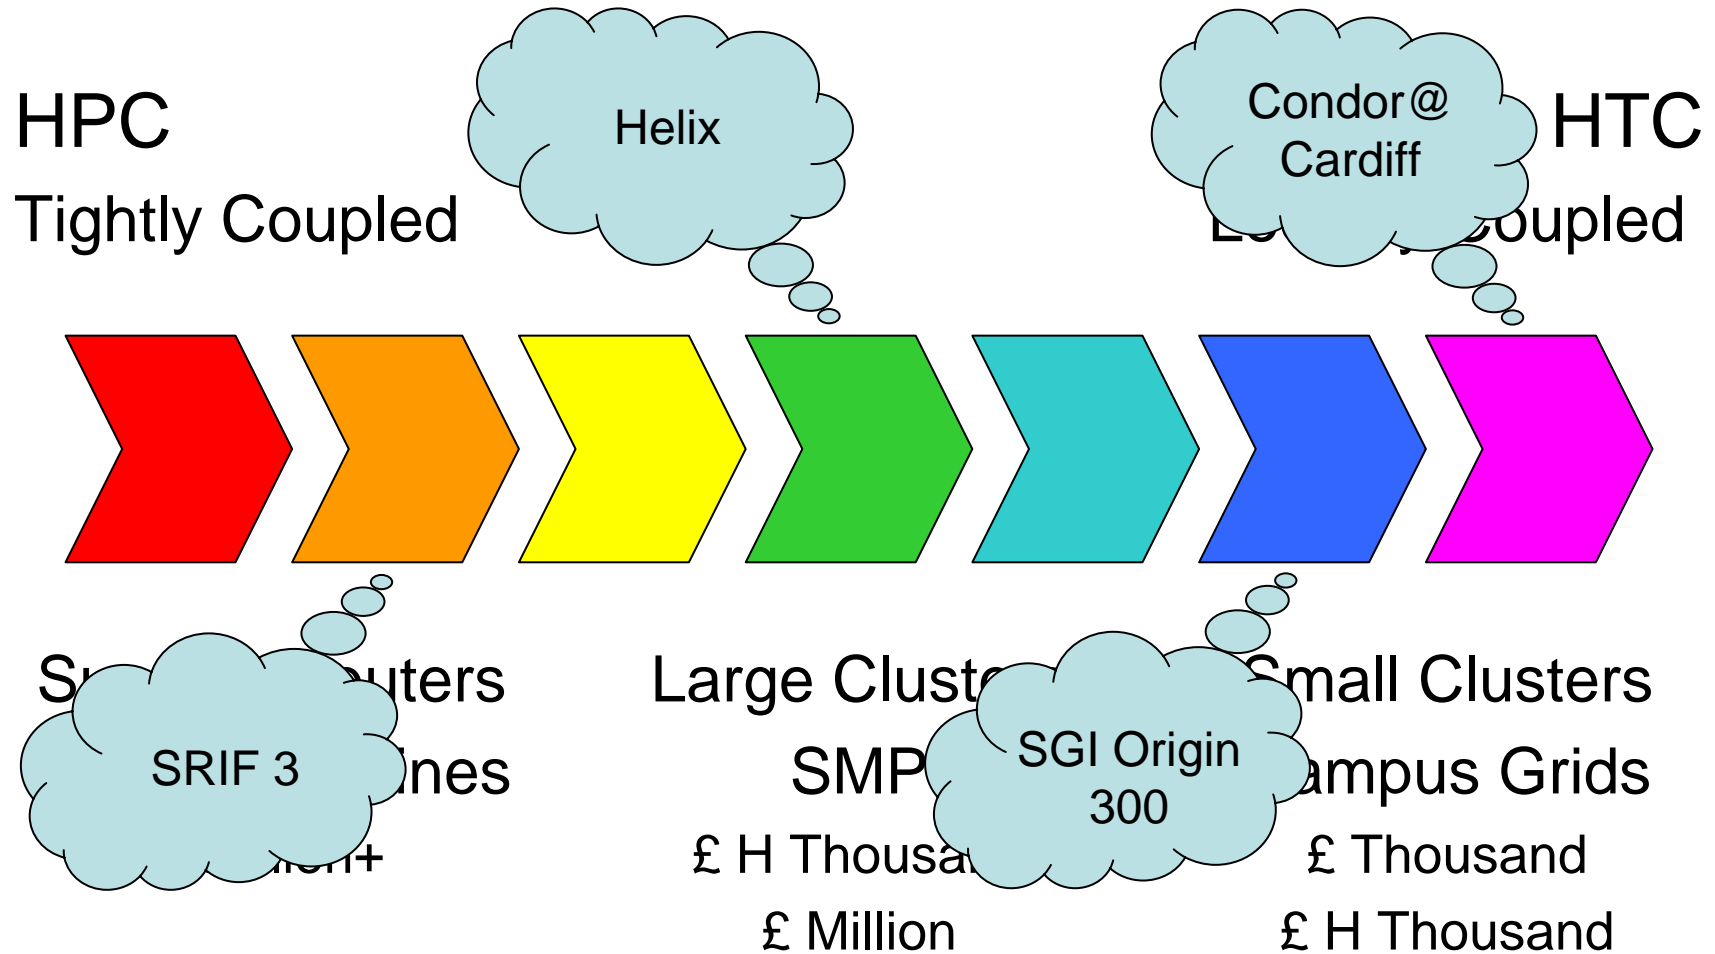
High End Computing Futures



The HEC Spectrum



The HEC Spectrum





Questions ?

condor@cardiff.ac.uk