

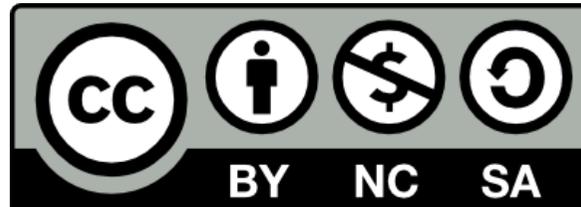
# Parallel Software usage on UK National HPC Facilities 2009-2015

---

Andy Turner, EPCC  
a.turner@epcc.ed.ac.uk



# Reusing this material



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

[http://creativecommons.org/licenses/by-nc-sa/4.0/deed.en\\_US](http://creativecommons.org/licenses/by-nc-sa/4.0/deed.en_US)

This means you are free to copy and redistribute the material and adapt and build on the material under the following terms: You must give appropriate credit, provide a link to the license and indicate if changes were made. If you adapt or build on the material you must distribute your work under the same license as the original.

Note that this presentation contains images owned by others. Please seek their permission before reusing these images.





EPSRC

NERC SCIENCE OF THE ENVIRONMENT

CRAY  
THE SUPERCOMPUTER COMPANY

| epcc |

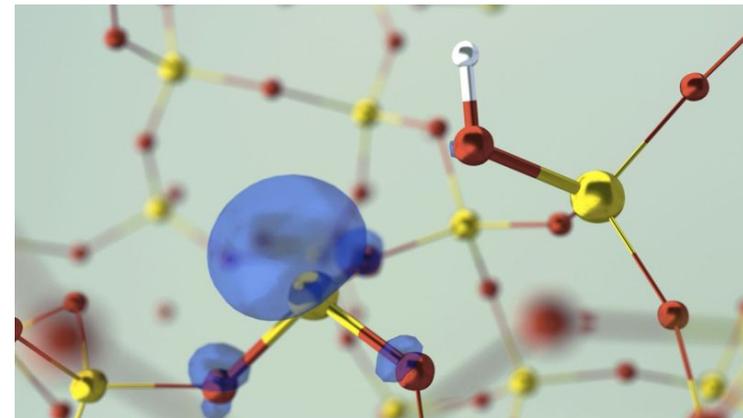
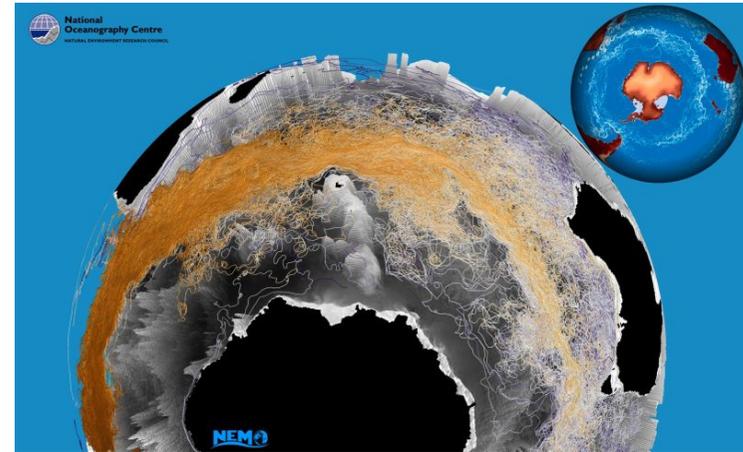


[www.epcc.ed.ac.uk](http://www.epcc.ed.ac.uk)  
[www.archer.ac.uk](http://www.archer.ac.uk)



# Outline

- Motivation
- Systems Included
- Analysis Tool
- Overall Comparisons
- Specific Research Areas
- Future Look
- Summary



# Acknowledgements

- Thanks to the EPCC User Support and Liaison Team for interesting discussions and comments around this analysis
  - Particularly Alan Simpson and Stephen Booth
- Thanks also to the Cray Centre of Excellence staff at the University of Edinburgh
  - Particularly Jason Beech-Brandt



# Motivation

- Understand the application profile on UK National Supercomputing Services so we can provide a better service
  - Different research areas will have different requirements
- Which application types are able to scale and which not?
  - Can we say anything about what is limiting scaling based on this?
- Which areas are growing in usage and which are decreasing?
  - Same question for applications: which new applications have appeared and which old ones have disappeared?
- What does this mean for future systems?



# Systems Included



# Systems Included

System (Type)	Processor Arch. (Clock Speed)	Cores per Node (Sockets)	Memory/Node (Bandwidth/core)	Nodes (Cores)	$R_{peak}$ / Tflop/s
HECToR Phase 2a (Cray XT4)	AMD Barcelona (2.3 GHz)	4 (1)	8 GB (3.2 GB/s)	5664 (22656)	63.4
HECToR Phase 2b (Cray XE6)	AMD Magny-Cours (2.1 GHz)	24 (2)	32 GB (3.6 GB/s)	1856 (44544)	372.8
HECToR Phase 3 (Cray XE6)	AMD Interlagos (2.3 GHz)	32 (2)	32 GB (2.7 GB/s)	2816 (90112)	829.0
ARCHER (Cray XC30)	Intel Ivy Bridge (2.6 GHz)	24 (2)	64 GB (4.9 GB/s)	4920 (118080)	2550.5



# Analysis Tool



# Analysis Tool

- Poll ALPS on an hourly basis and store:
  - User, size of job, executable name, ALPS ID
- Analyse logs using Python program
  - Extensible descriptions of applications include: programming language, parallel model, code type, research area, license
  - Executable names matched via regexp
- Can limit to specific period or project
- Text, graphical and/or CSV output
- Hints to help identify further applications



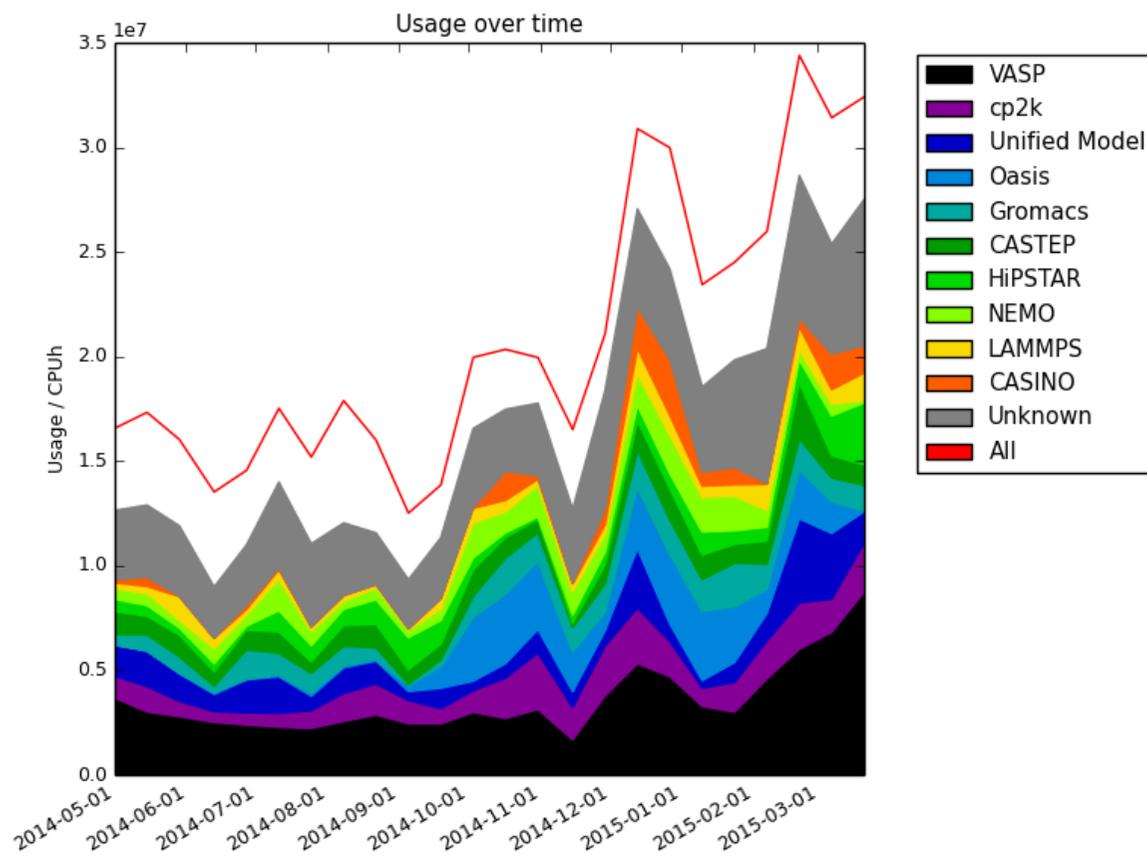
# Analysis Tool: Example Output

Total code usage (ordered by CPUh)

Code	CPUh	% Time	Jobs	% Jobs	Users	Mean	Median
====	=====	=====	====	=====	=====	=====	=====
VASP	90128736.0000	17.36	202742	40.78	177	2202.36	240.0000
cp2k	35669040.0000	6.87	21155	4.25	78	3702.55	672.0000
Unified Model	31979904.0000	6.16	18613	3.74	159	2087.09	1392.0000
Oasis	31491624.0000	6.07	2074	0.42	10	6249.98	5232.0000
Gromacs	26942616.0000	5.19	13759	2.77	69	941.73	432.0000
CASTEP	26060376.0000	5.02	74661	15.02	92	2434.20	360.0000
HiPSTAR	24043632.0000	4.63	1146	0.23	11	14439.75	10344.0000
NEMO	21366096.0000	4.12	19498	3.92	21	3589.34	1920.0000
LAMMPS	14493840.0000	2.79	9851	1.98	50	672.84	456.0000
CASINO	13404504.0000	2.58	861	0.17	4	5991.07	5016.0000
ONETEP	12339456.0000	2.38	3190	0.64	23	1079.92	864.0000
Hydra	10611552.0000	2.04	930	0.19	13	2069.26	1248.0000
NAMD	9475416.0000	1.83	8913	1.79	34	613.54	480.0000
CRYSTAL	8659632.0000	1.67	2237	0.45	26	6945.64	2808.0000
PDNS3D	7278528.0000	1.40	755	0.15	9	6721.09	6144.0000
WRF	6611784.0000	1.27	1170	0.24	20	2591.09	2064.0000
OpenFOAM	6474984.0000	1.25	2560	0.51	19	695.84	288.0000
MITgcm	5415576.0000	1.04	8284	1.67	23	571.12	384.0000



# Analysis Tool: Example Output



# Overall Comparisons

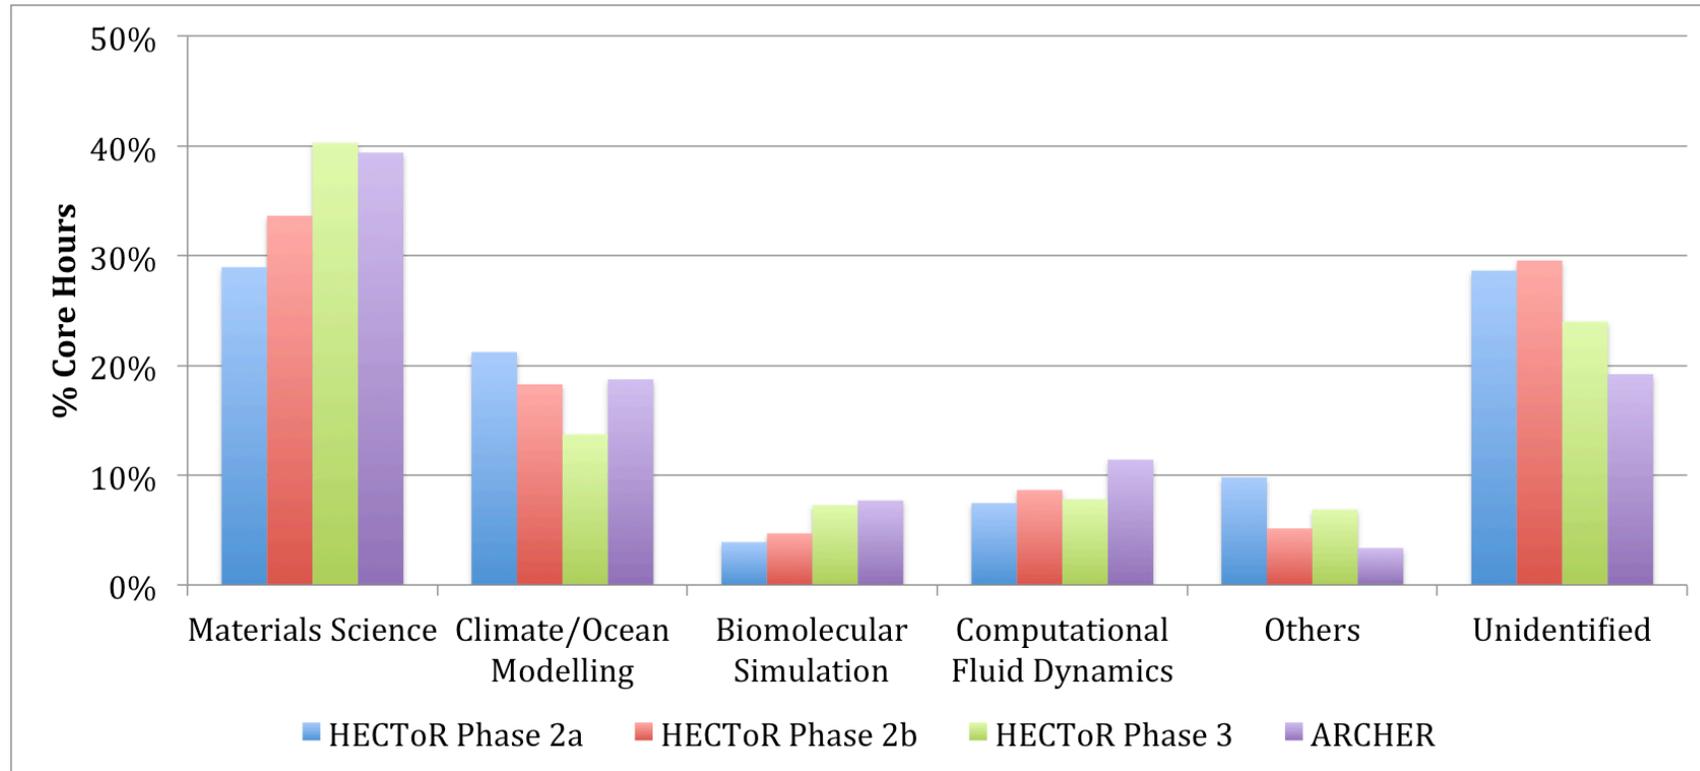


# Overall Comparisons: Top Ten Codes

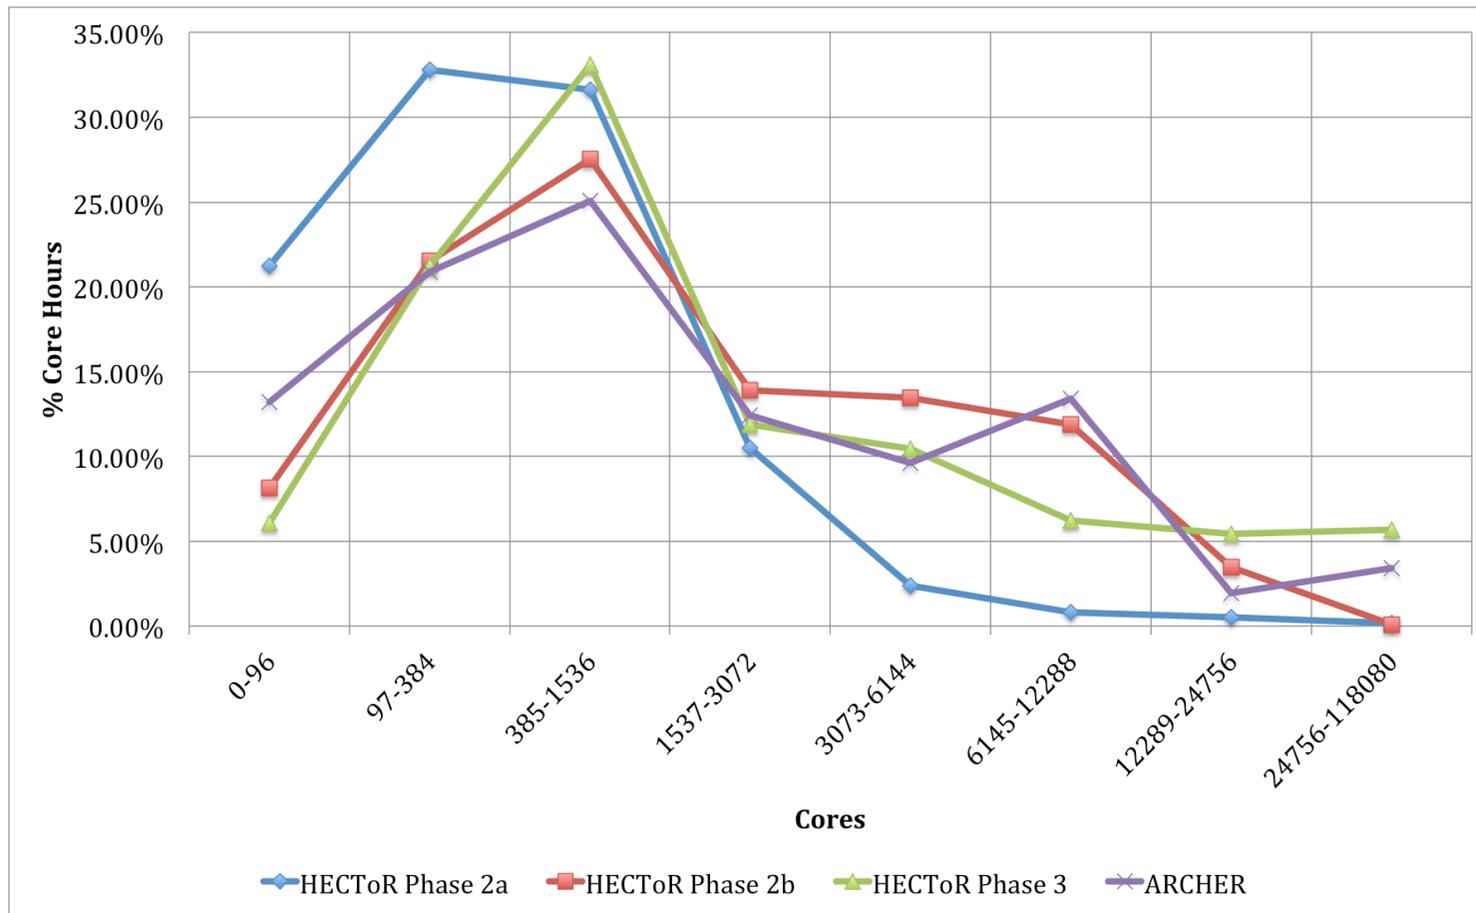
Rank	HECToR Phase 2a	HECToR Phase 2b	HECToR Phase 3	ARCHER
1	UM	VASP	VASP	VASP
2	VASP	UM	CP2K	CP2K
3	CASTEP	CASTEP	UM	UM
4	Hydra	CP2K	CASTEP	Oasis
5	CP2K	INCOMPACT3 D	Gromacs	Gromacs
6	Chroma	NEMO	DL_POLY	CASTEP
7	NAMD	Gromacs	PDNS3D	HiPSTAR
8	ChemShell	MITgcm	MITgcm	NEMO
9	WRF	ChemShell	NEMO	LAMMPS
10	DL_POLY	PDNS3D	CRYSTAL	CASINO



# Overall Comparisons



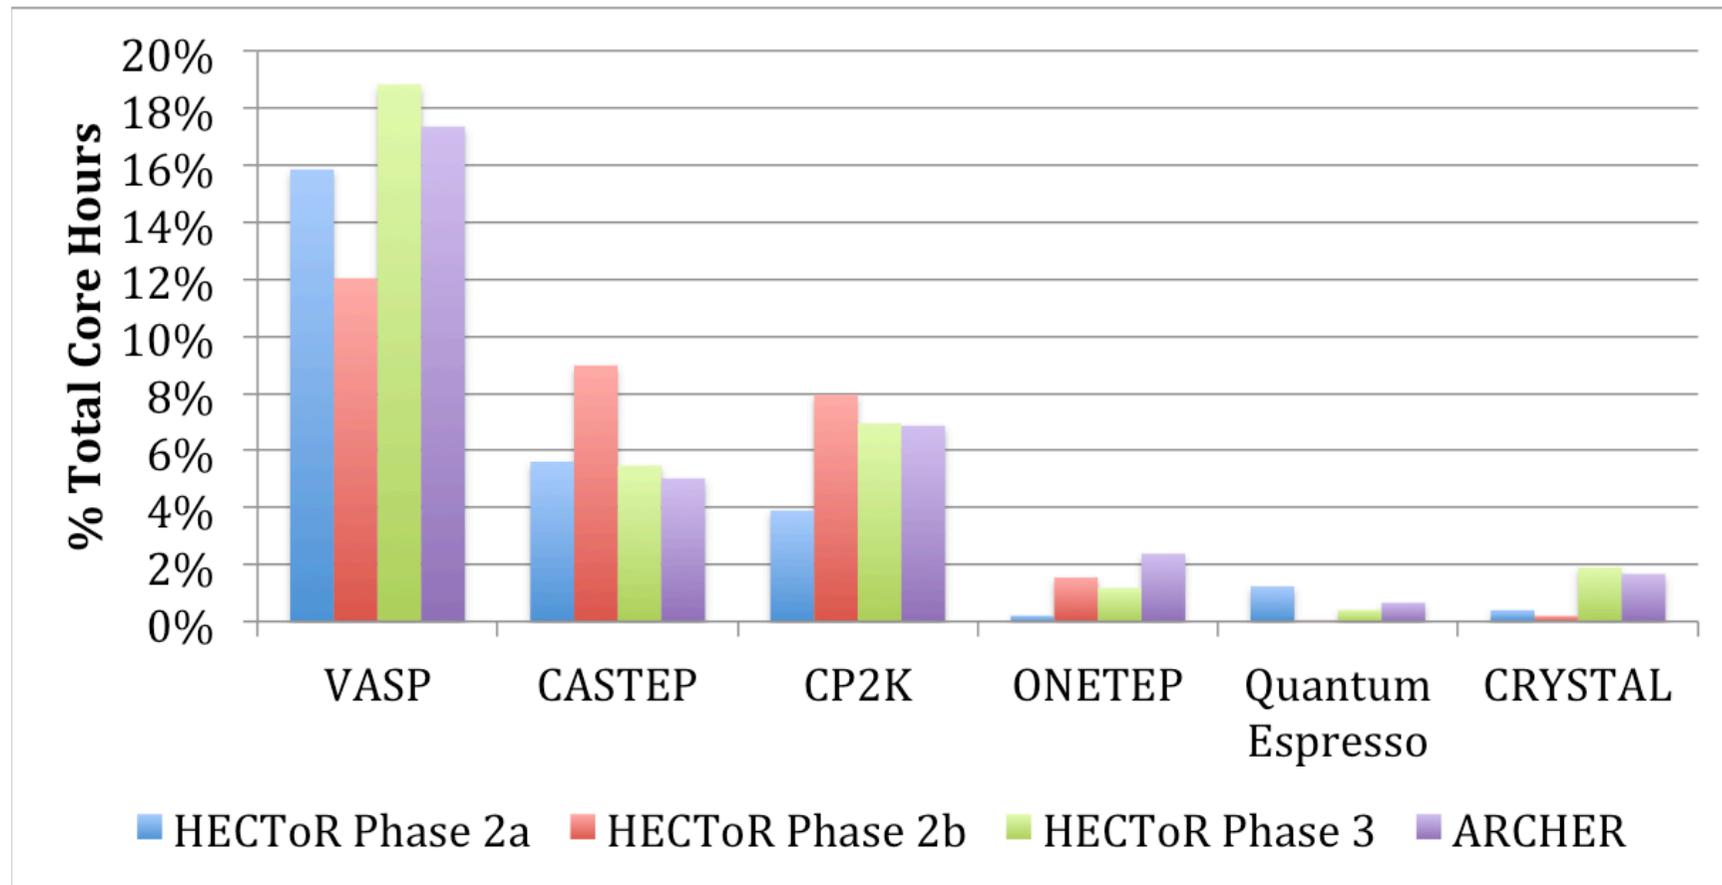
# Overall Comparisons



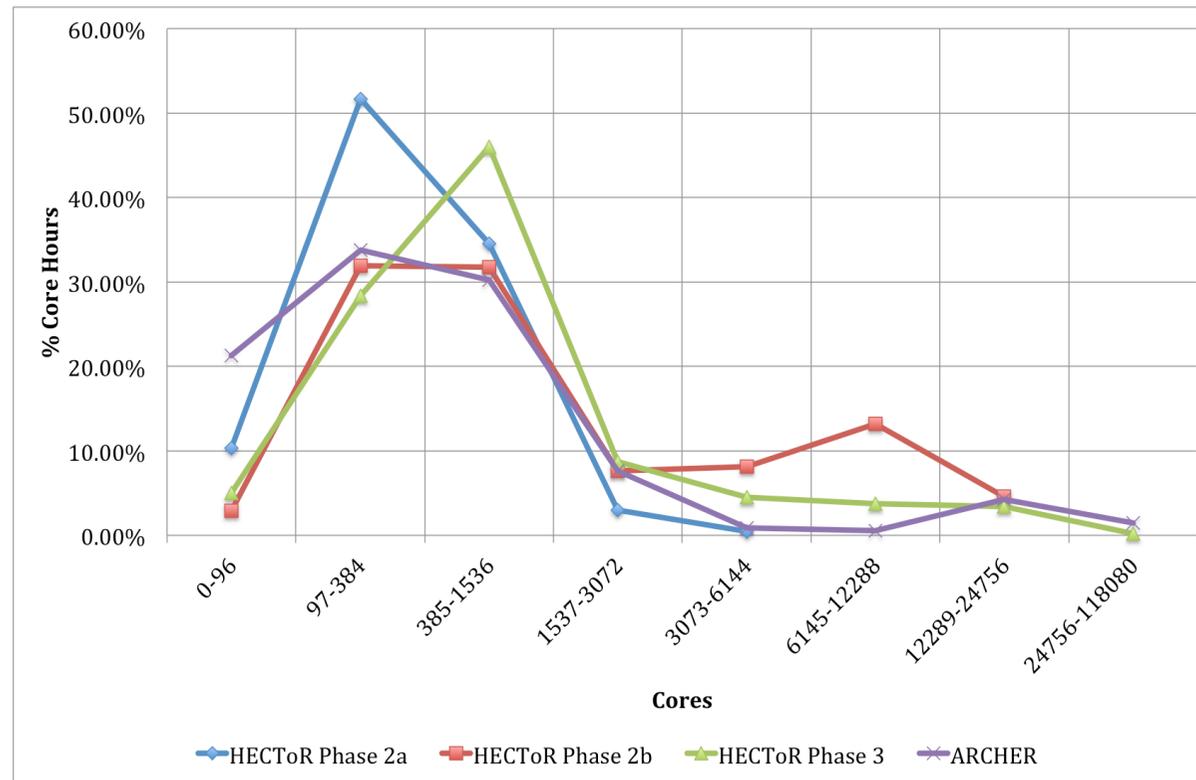
# Application Areas



# Periodic Electronic Structure



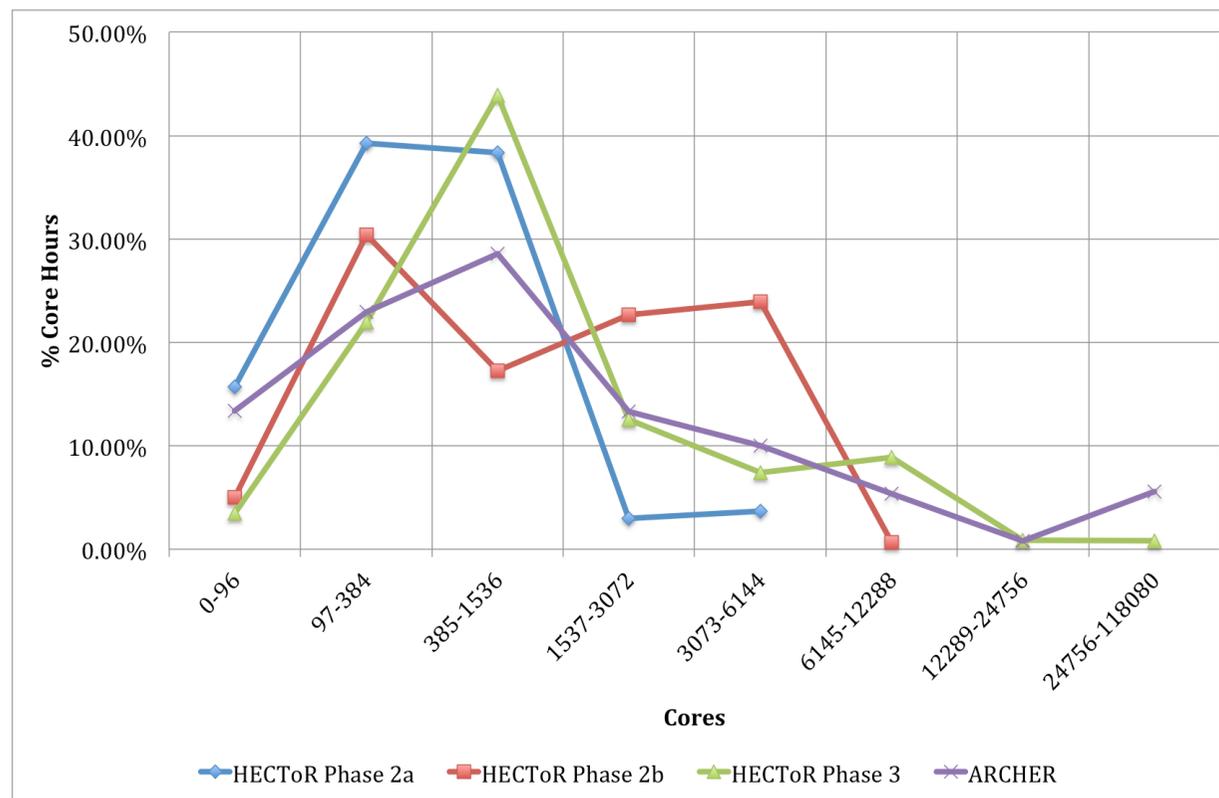
# Periodic electronic structure



CASTEP job size distribution



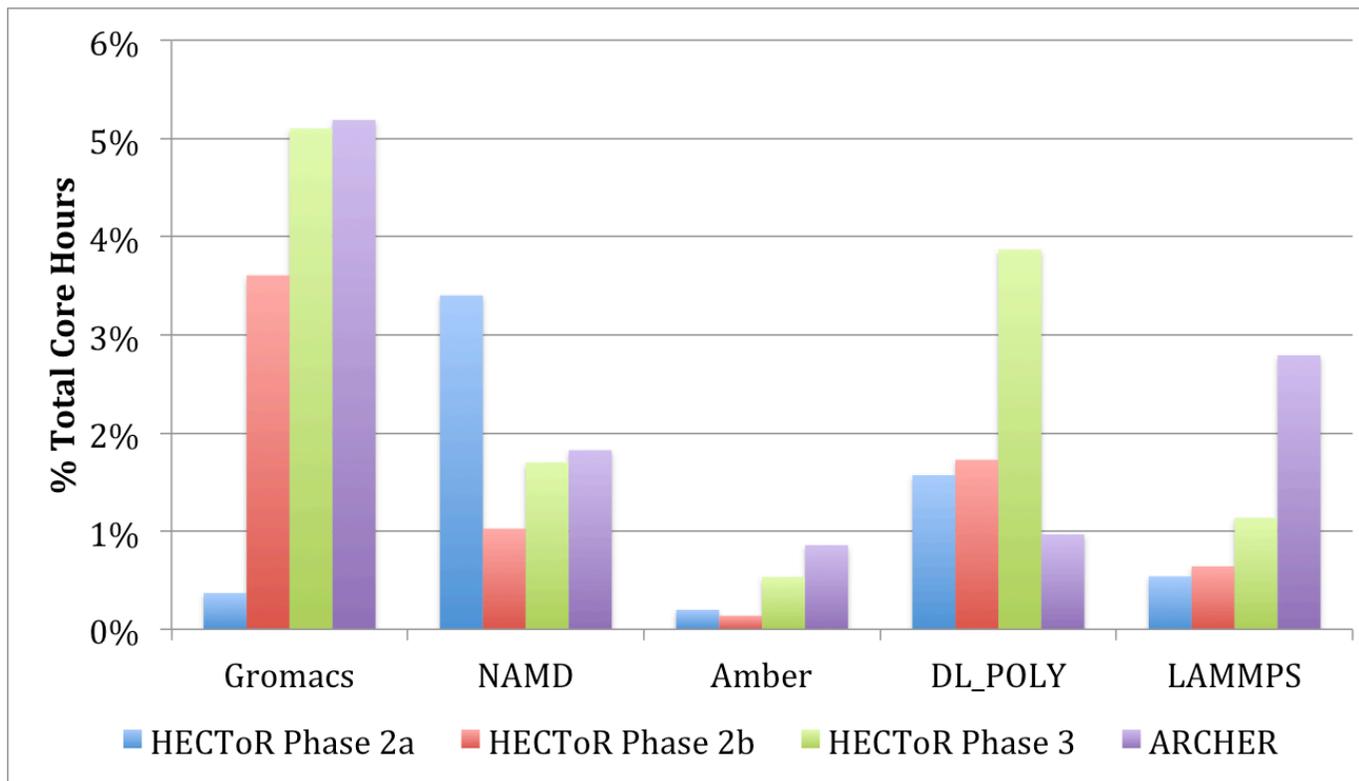
# Periodic electronic structure



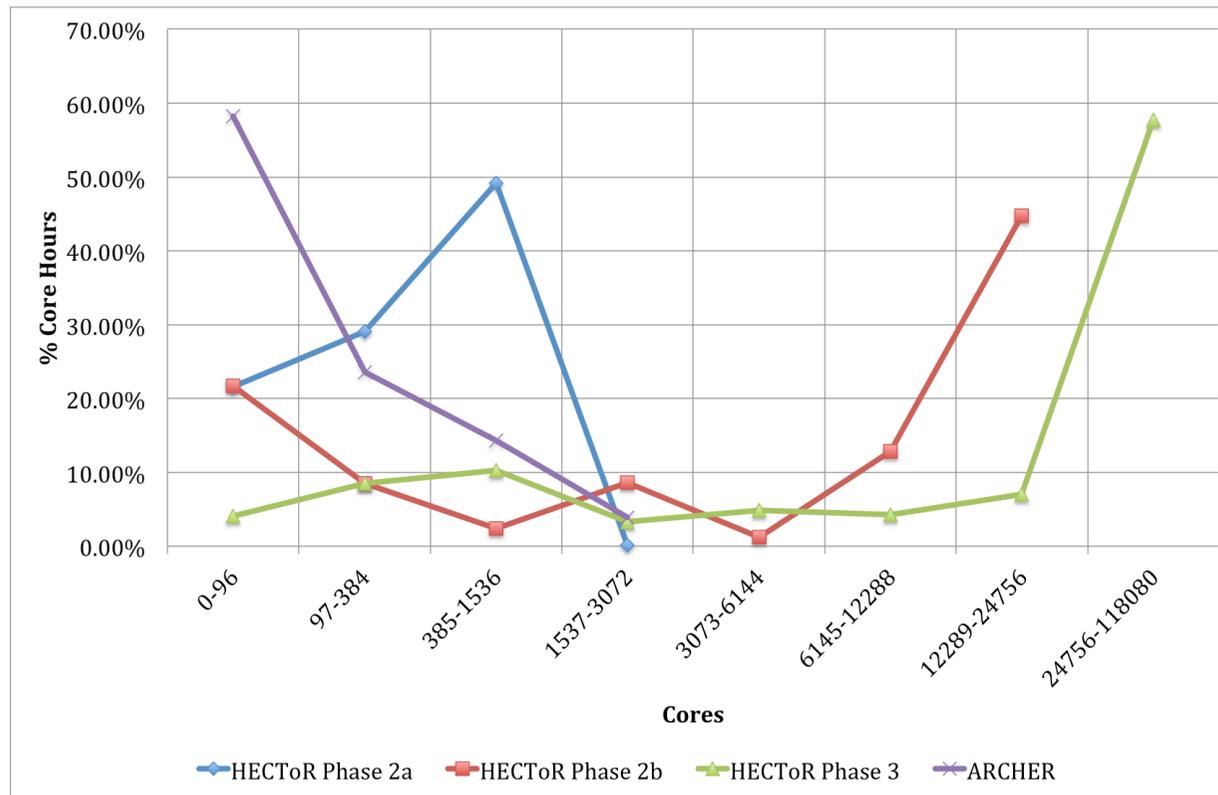
CP2K job size distribution



# N-body codes



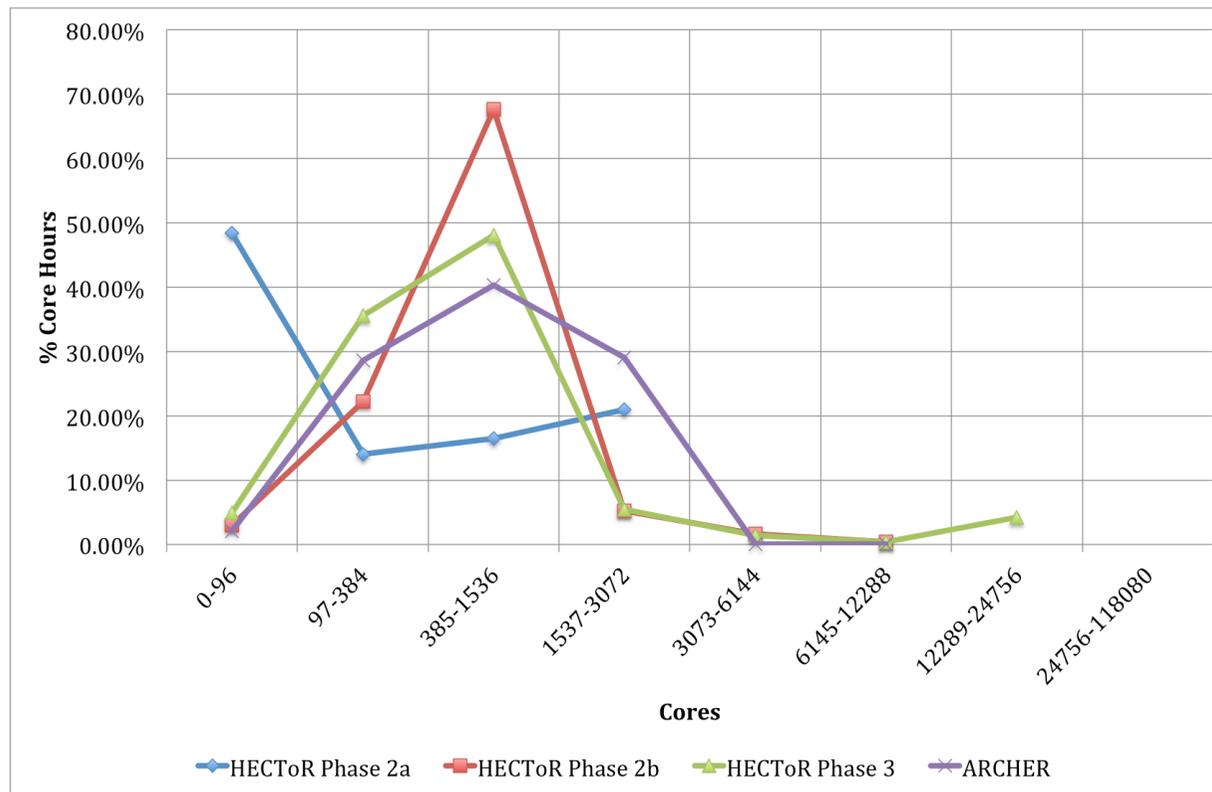
# N-body codes



DL\_POLY job size distribution



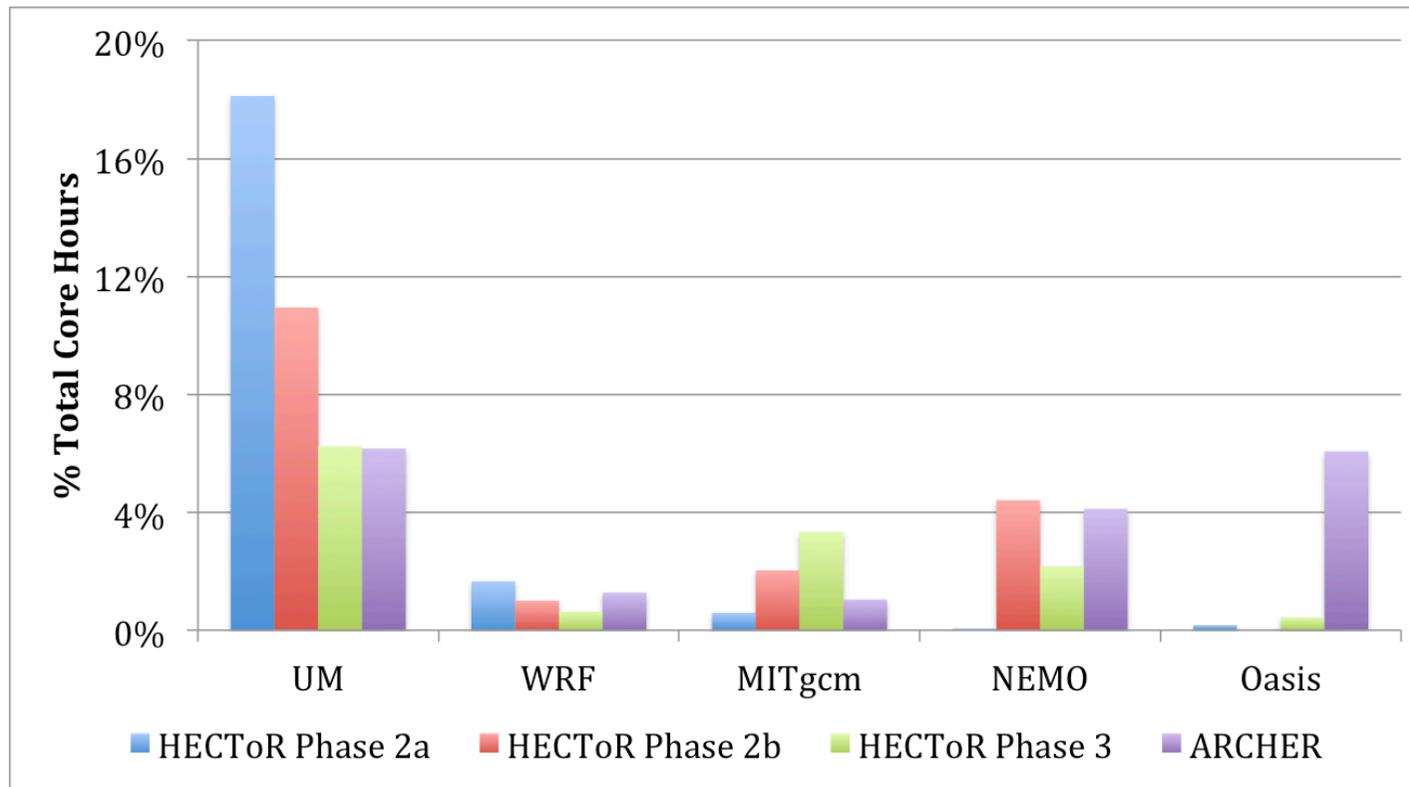
# N-body codes



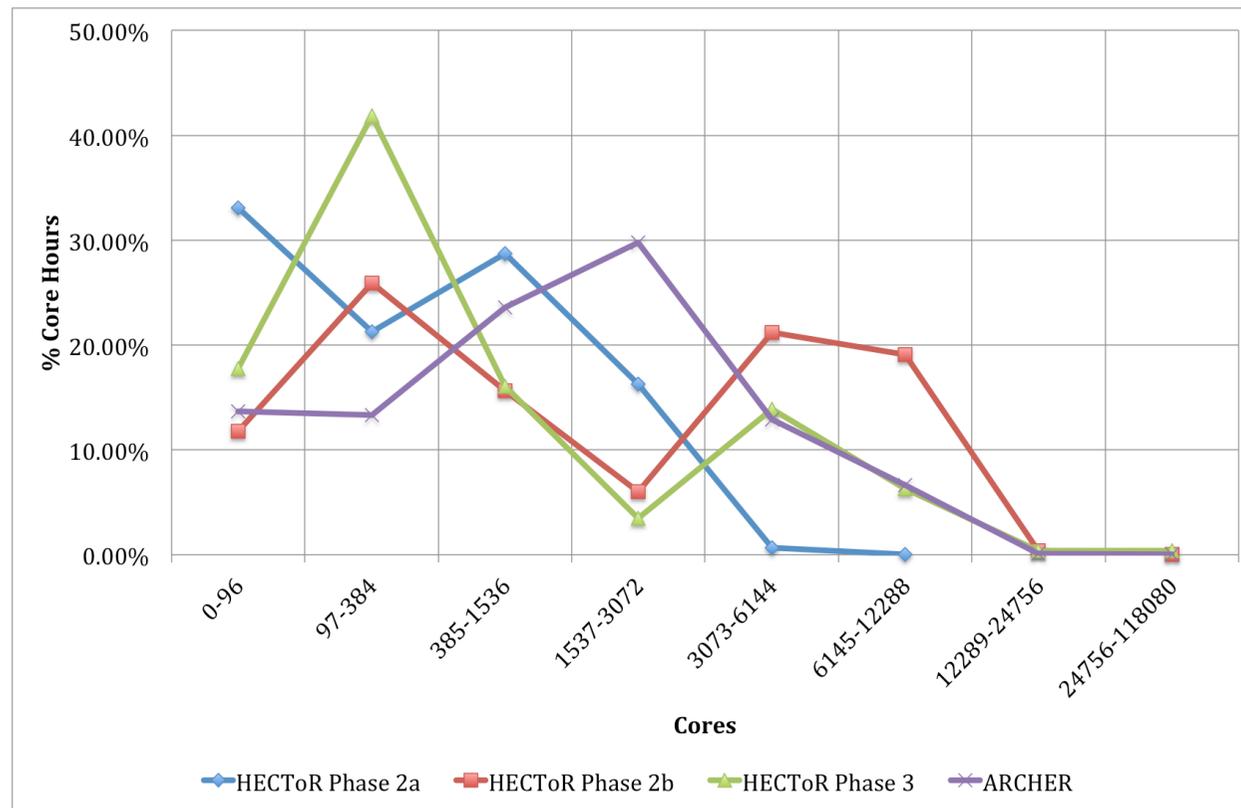
Gromacs job size distribution



# Grid-based Codes: Climate/Ocean



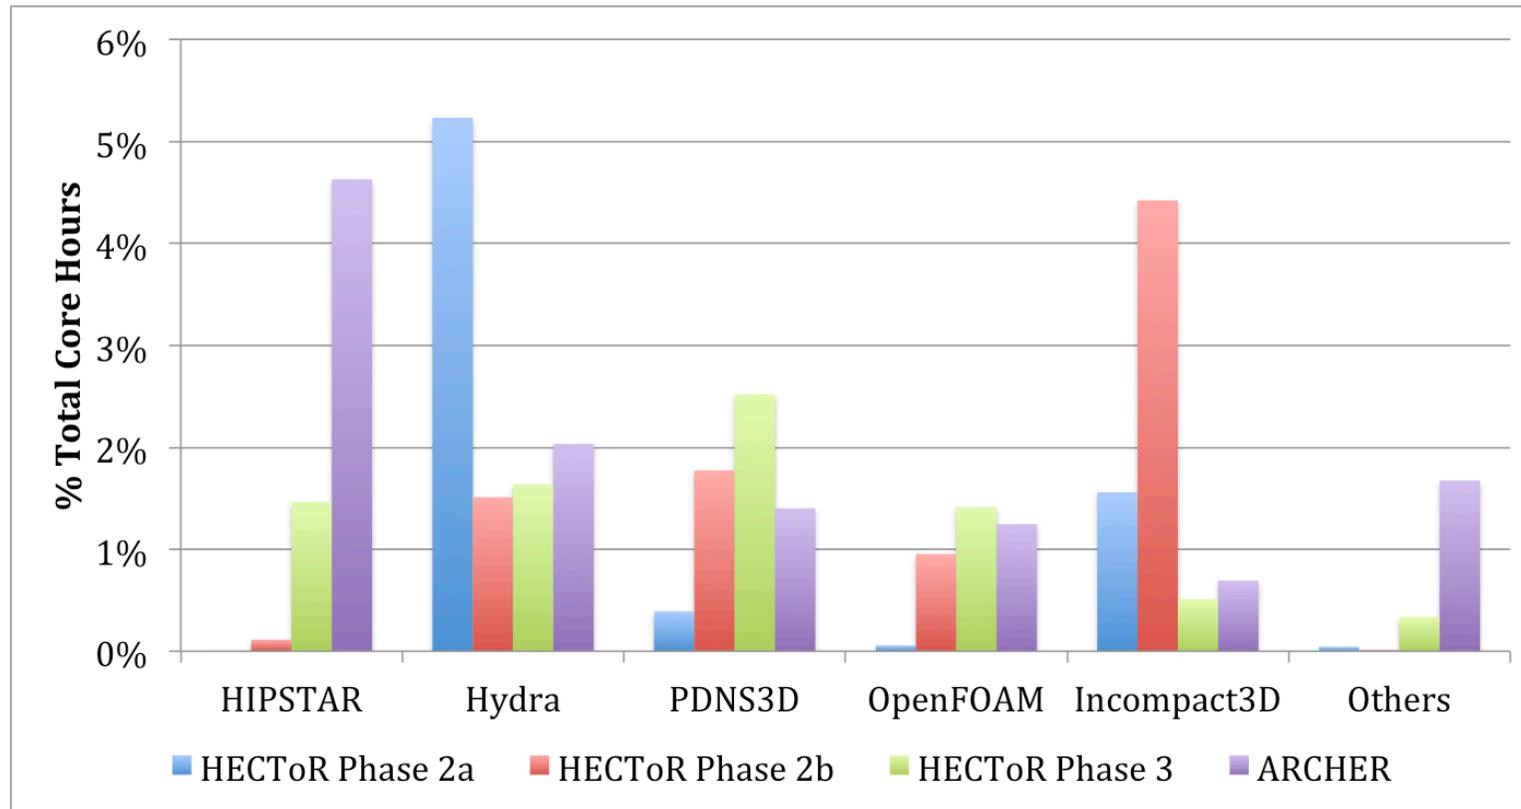
# Structured Grid: Climate Simulation



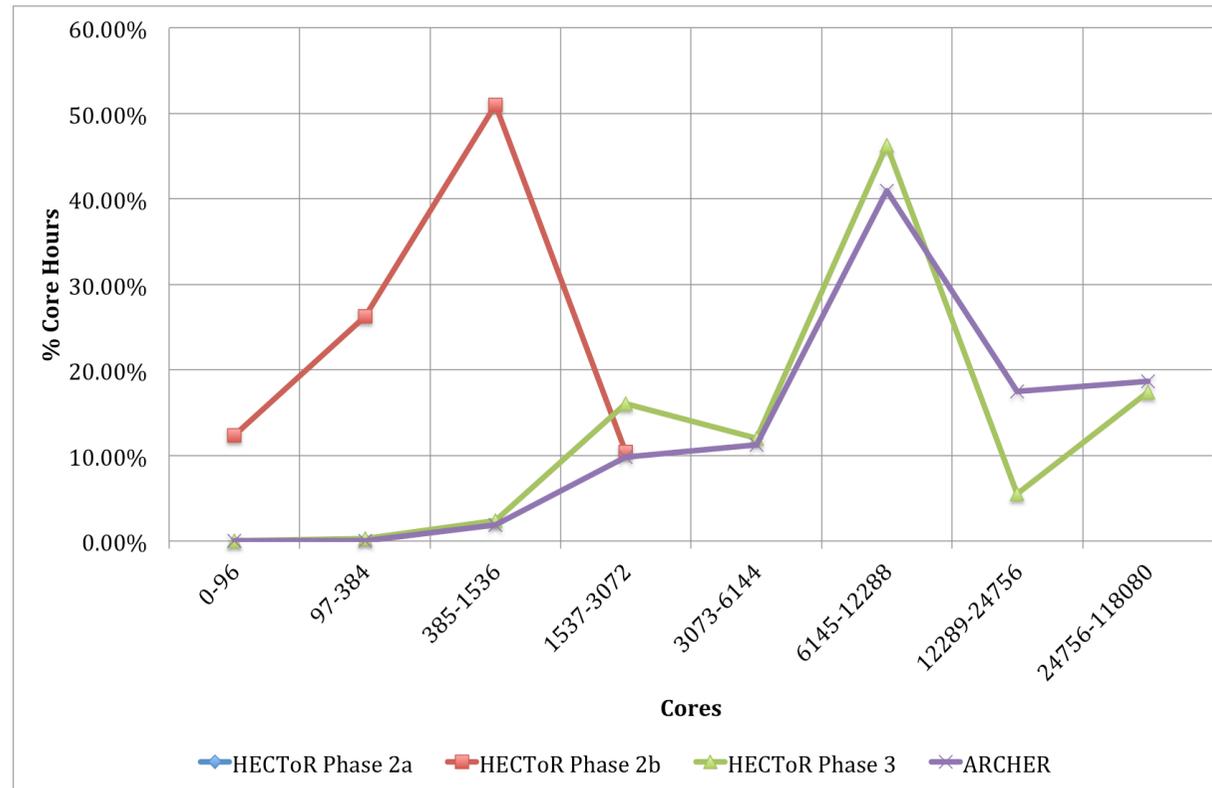
Met Office UM job size distribution



# Grid-based Codes: CFD



# Unstructured Grid: CFD



HiPSTAR job size distribution



# Future Look

1. Scaling limited by scientific problem
    - e.g. biomolecular simulation
    - Use additional throughput to access more sophisticated sampling
    - Via application or code agnostic frameworks (e.g. PLUMED)
  2. Scaling not limited by scientific problem
    - e.g. grid-based methods
    - Opportunity for single calculations to scale to large core counts
    - (As well as exploiting additional throughput)
- In both cases continued software development is key to exploiting future HPC architectures



# Summary

- Most applications able to increase scaling with switch from quad core to multicore
  - Generally by 2 times rather than 6 times
- For some areas scaling is generally limited by research problem rather than the application issues
  - Application scaling limits may never be reached for problems that are scientifically relevant
- In other areas application scaling is key to furthering research
- Future HPC systems offer opportunities for applications in both classes
  - In the UK we maybe need to be less focused on single application scaling and look at sampling and coupling frameworks more
- Lots more analysis of this data to be done!



# Questions?

