



An Update on HPC at the Met Office

Paul Selwood & Chris Maynard



- Modelling system overview
- Single precision solver
- HPC procurement




Met Office





Modelling System Overview





Two model Strategy: Global and UK



N x Global predictions at ~20km with lead times of days to years:
Synoptic drivers



<N x Regional predictions at ~1km:
Local meteorology



PDF of local hazard:
Impacts



Deterministic NWP model suite

Global


- 17km 70 Levels
- 48 hour forecast twice/day
- 6 day forecast twice/day

Euro4

- 4.4km 70 Levels
- 60 hour forecast twice/day
- 5 day forecast twice/day

UKV

- 1.5km 70 Levels
- 36 hour forecast eight times/day





Ensemble NWP model suite

MOGREPS-15


- 60km 70 Levels
- 15 day forecast 2 times/day
- 24 members

MOGREPS-G

- 33km 70 Levels
- 7 day forecast 4 times/day
- 12 members
- 24 member lagged products

MOGREPS-UK

- 2.2km 70 Levels
- 36 hour forecast 4 times/day
- 12 members





Ensemble NWP model suite

Changes

MOGREPS-15


will be retired
by March 2015
Week 2 from ECMWF

MOGREPS-G

- 33km 70 Levels
- 7 day forecast 4 times/day
- 12 members
- 24 member lagged products

MOGREPS-UK

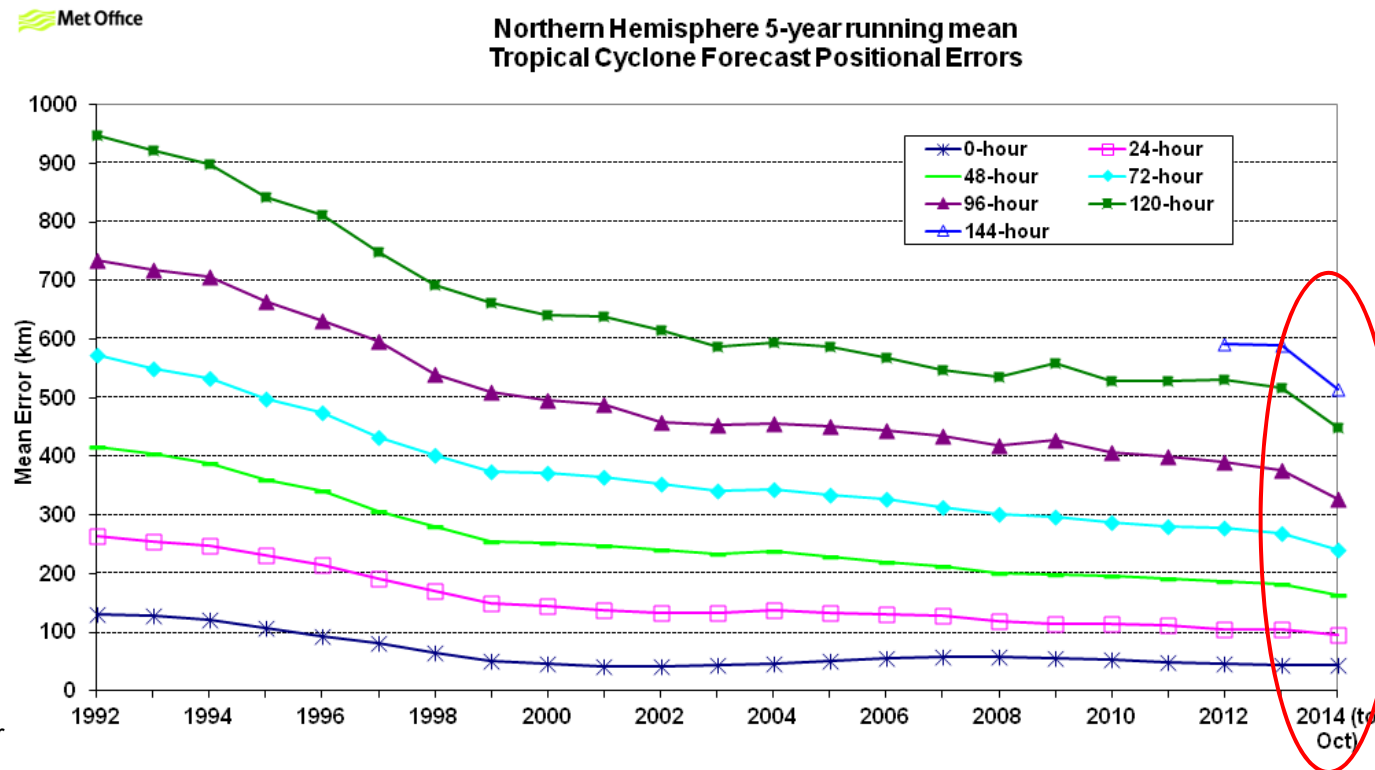
- 2.2km 70 Levels
- 36 hour forecast 4 times/day
- 12 members





Tropical Cyclone Track Forecasts

- Global Model northern hemisphere tropical cyclone track forecast errors 2014 so far 20-25% lower than mean for previous 5 years
- Biggest drop in 5-year running mean for 15 years






Tropical Cyclone Intensity Forecasts

- Northern hemisphere 2014 to late October big reductions in tropical cyclone intensity errors (particularly at longer lead times)



Met Office Global Model Northern Hemisphere
Tropical Cyclone Central Pressure Mean Error





Met Office

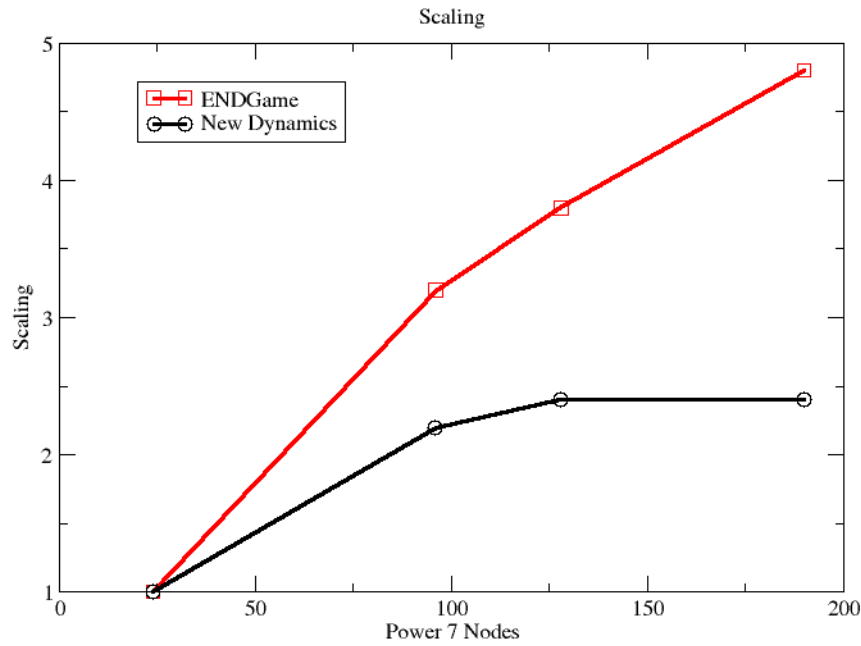


Single Precision Solver




ENDGame performance

N768 - New Dynamics vs ENDGame



N768 - New Dynamics vs ENDGame





Using single precision

Understanding the error

Accuracy of Krylov subspace solver – BiCGStab
Iterative solver, it improves the answer each iteration

In our model $\epsilon < 10^{-3}$
answer is *good enough*
What precision is needed to satisfy this?
Single precision *is good enough*

$$\|\vec{r}_k\| = \|\vec{b} - A\vec{v}_k\| < \epsilon$$

Modern FPU single prec operation is not significantly faster than double

Single prec words are **half the size** of double prec words

Compute values of A in **double precision**


Store them in **single precision**

Doubles the effectiveness of cache




Preconditioner scaling

Max total time for TRI_SOR
diagonal projection





Accuracy of answers



After 5 time-steps,
level 10 biggest
differences

$$f = (\xi_{64} - \xi_{32}) \times 10^{-4}$$

$$\xi = \text{exner} = \left(\frac{p}{p_0} \right)^k$$

24x32 EWxNS Proc
1024x769x70 grid

N512 6 time-steps (s)	64 bit	32 bit	Speed-up
EG_SL_HELMHOLZ	3.884	2.836	1.4
EG_BICGSTAB	2.876	1.809	1.6
TRI_SOR	2.075	1.124	1.8

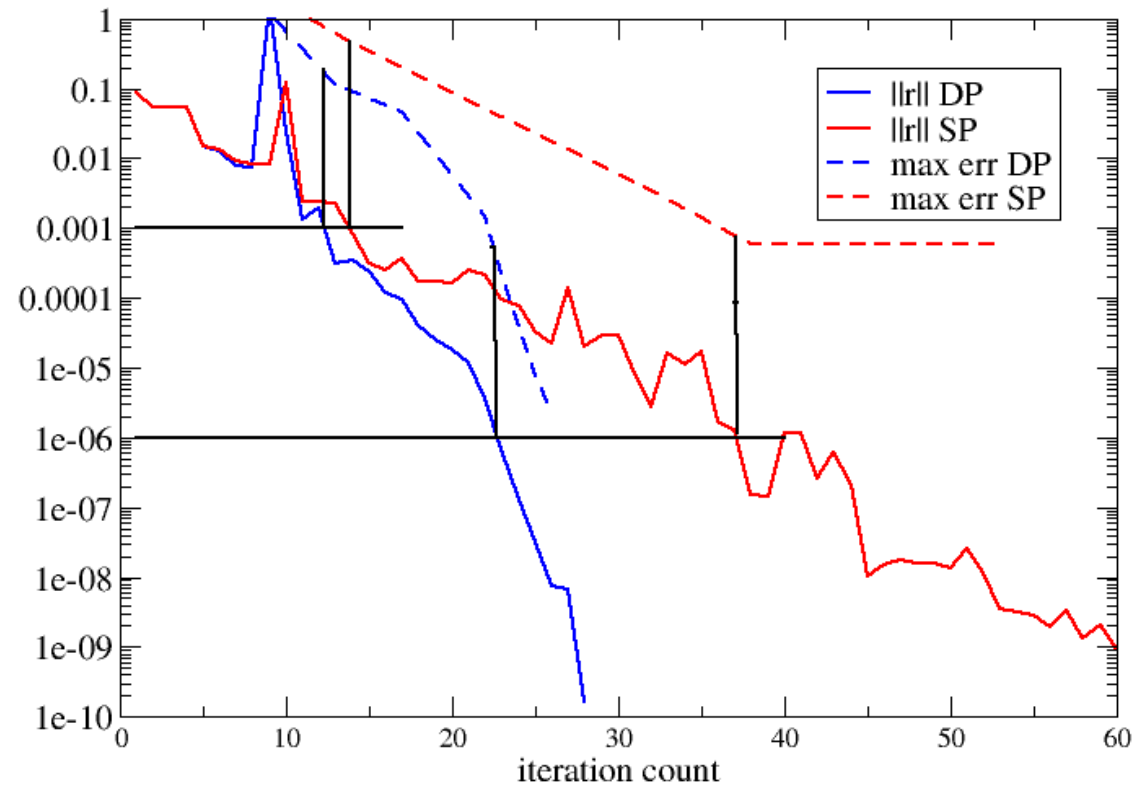


Convergence of Full SP Solver

Serial stand-alone solver (no comms)

Single precision except for correction vector

Single precision comms (halo-swap and global sum) are also faster in full UM





Results (N512, 768 cores)

	Atm_Step	Solver	EG_BiCGStab
Double Precision	1955	500	370
Single Precision Helmholtz Matrix	1812	355	227
Single Precision	1763	309	181

- Solver no longer most expensive part of timestep
- Further improvements still possible
 - Red-Black ordering of Helmholtz matrix
 - Communications improvements



Met Office

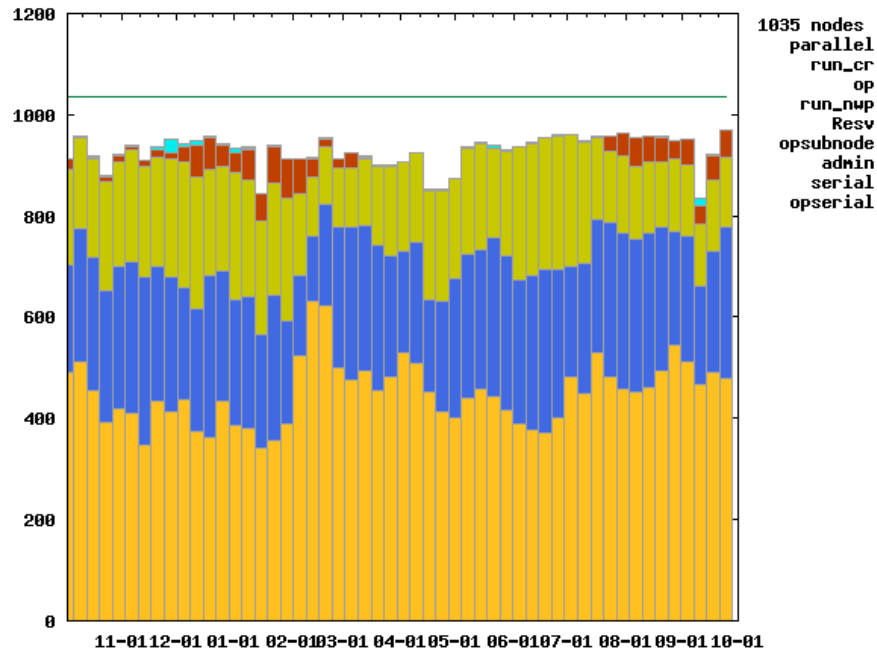


HPC Procurement

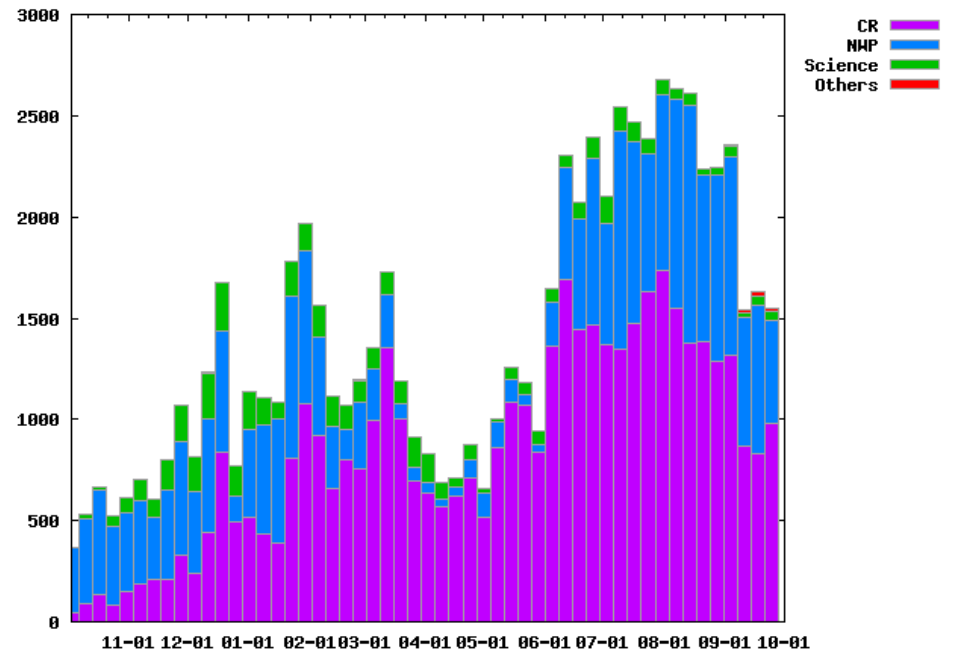


HPC Utilisation (1 year)

Last 52 weeks All Divisions/Accounts : + resvs: both Phase 2 clusters
Parallel jobs: wallclock


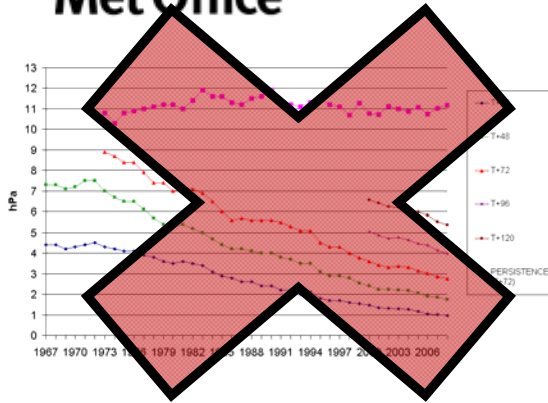


Last 52 weeks All Divisions/Accounts : both Phase 2 clusters
Parallel jobs: queueing - weighted by node





Making the case for more HPC



Socio-economic benefit case studies on

- Aviation
- Flooding
- Food security
- Renewables
- Winter travel
- Decadal – centennial advice for mitigation/adaptation



New Funding Model

- Previously had a loan from owning department
- Now applied for £97 million grant to cover,
 - New HPC
 - New off-site IT hall
 - Archive and other downstream impacts
- Still need to fund running costs out of our normal business



Phasing

- Initial Test and Development systems in Autumn 2014
- Phase 1a – two clusters to replace Power 7s by September 2015 – power is a problem!
- Phase 1b – extend both clusters to power limit by March 2016
- Phase 1c – 1 new cluster in new IT Hall by March 2017



Procurement Timelines

- First RAPS release – December 2012
- Wide ranging discussions with potential suppliers, September 2013 – March 2014
- Full and final RAPS release – October 2013
- September 2013 – Draft requirements released
- January 2014 – 2nd draft requirements

- ITT - February 2014
- Shortlist – April 2014
- Preferred Bidder – August 2014



Benchmarking


Model	Weight in Evaluation	No. of copies on IBM P7	Nodes per copy on IBM P7
UM-N1024	30%	5.5	192
UM-N144 + Chemistry	30%	33	32
NEMO 0.25 + CICE & Tracers	20%	33	32
4DVAR - N320	20%	22	48

The Benchmark Challenge :

1. *At least match Power 7 runtimes*
2. *Define Capacity (number of nodes) to match existing capacity as a weighted average of the 4 benchmarks, running sufficient copies to fill both clusters.*
3. *Scale up that capacity within the Affordability (and for current IT Halls the Power) Constraints*
4. *Optimisations allowed, but limited LOC changes only*



New IT Hall – Planning Application



- Exeter Science Park
- Modern IT Facility
- 5.5 MVA, upgradeable
- Collaboration space

Thank You!
Questions?

