# Running a climate model on workstations and Linux clusters

Andy Heaps
Centre for Global
Atmospheric Modelling
andy@met.rdg.ac.uk

## Outline

- Acknowledgements
- Original inspiration
- A little history
- The model
- Using the model
- Porting problems and results on various platforms
- Climateprediction.com
- Conclusions

# Acknowledgements

Jeff Cole (CGAM)

Anette Van Der Wal, Paul Burton and Bill Roseblade (Met Office)

Robin Harker and Paul Ingram (Workstations UK)

Pete Oliver and Alan Iwi (RAL)

Paul Valdes (Reading university)

Dave Stainforth (Oxford university)

The original inspiration for the work.
A Sun 712MP workstation.

# Some history

- 1643 Invention of the barometer by Torricelli.
- 1820 First synoptic map made (using 50 year old data!)
- Early in the 20th century Vilhelm Bjerknes developed a set of seven equations to predict large scale atmospheric motions.
- In 1922 Lewis Fry Richardson developed the first numerical weather prediction system.

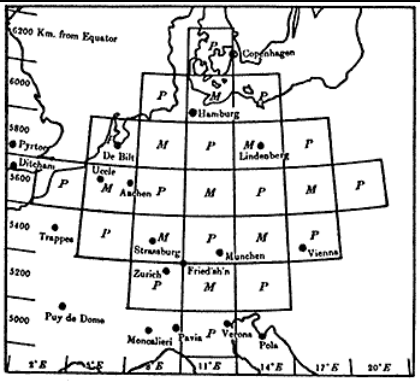


'P' provided atmospheric pressure.

'M' gave atmospheric momentum.

Six hour forecast took six weeks with a slide rule.

Forecast pressure change was 100 times greater than had actually been observed….

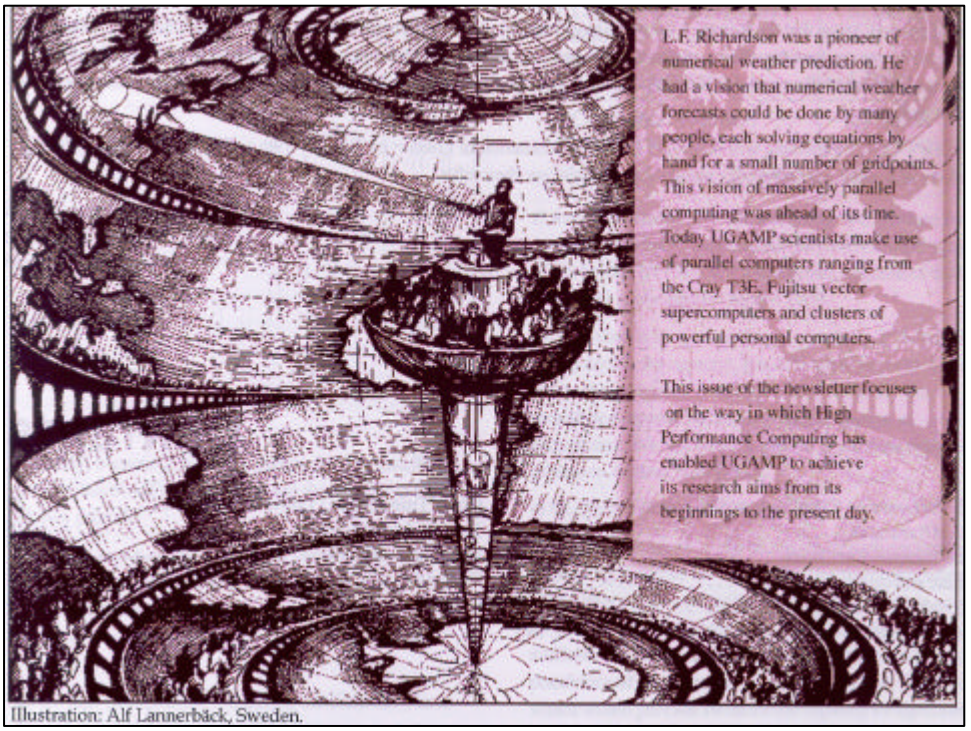

Illustration: Alf Lannerbäck, Sweden.

# The model

The model is based on the hydrostatic, primitive equations. The equations are discretised onto a Arakawa staggered 'B' grid in the horizontal, and onto a number of vertical levels or gridboxes of variable spacing. All variables except velocity are calculated at the gridbox centre; velocity is calculated at the gridbox verticies.

# Modelling Global Climate

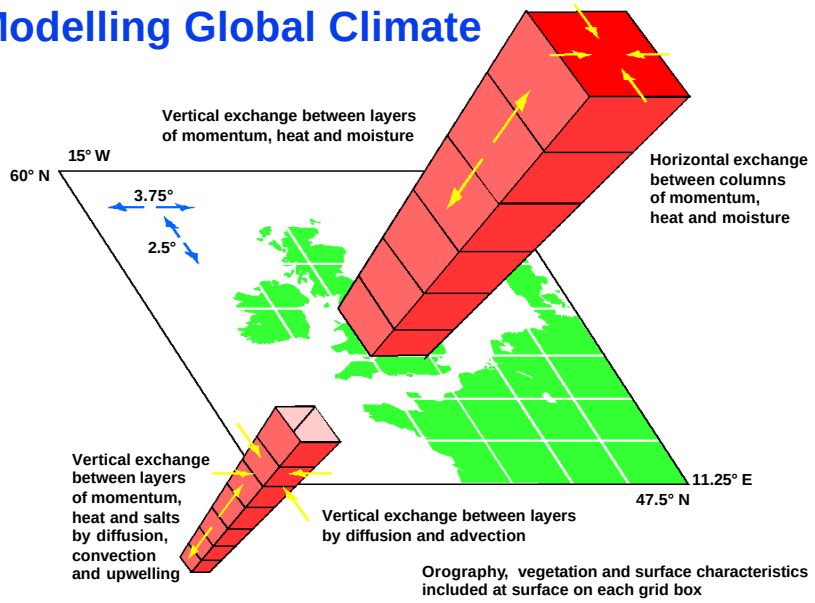




# The climate system

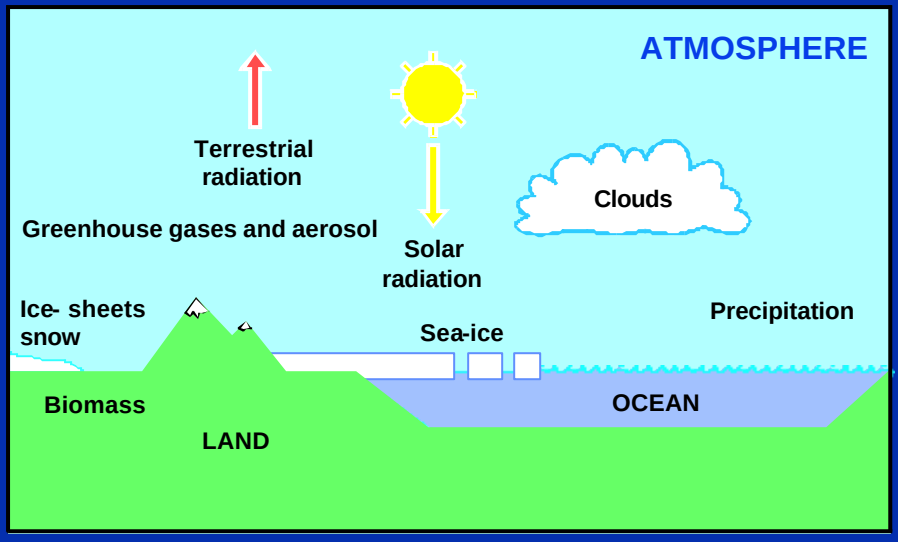

## Use of the Model – example 1

- The eruption of the Krakatau volcano (6 °S 105 °E) on August 27th 1883 led people to believe that the stratospheric wind above the equator blew in a westward direction.

  Dust from the eruption took 13 days to circle the equator and this upper air wind became known as the Krakatau easterlies.

- In 1908 Berson launched observational balloons above Lake Victoria in Africa and found westerly winds at about 15km (120mb). These westerly winds are called Berson's westerlies.

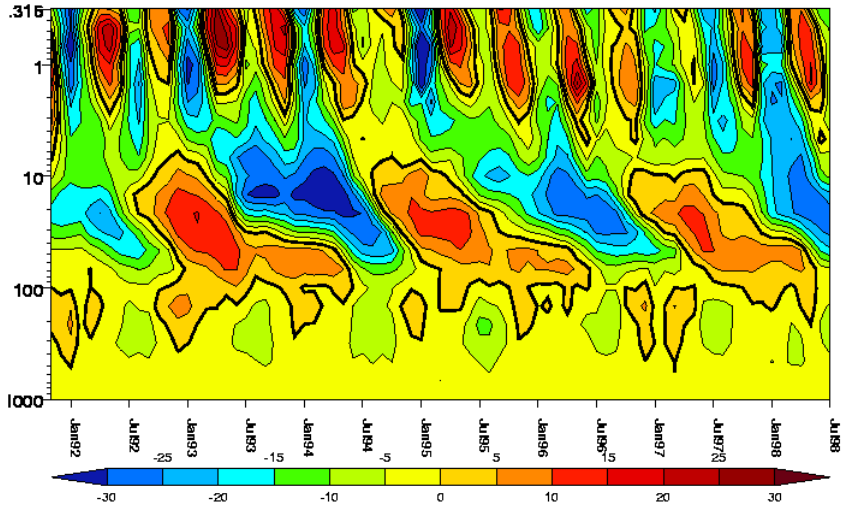


**The QBO as seen in the Met Office Assimilated dataset**

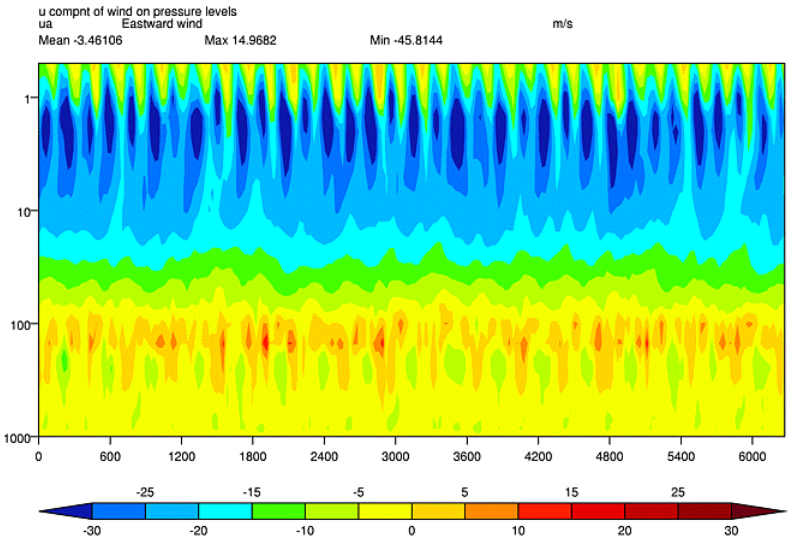


Lack of QBO in the Unified Model

# QBO theory

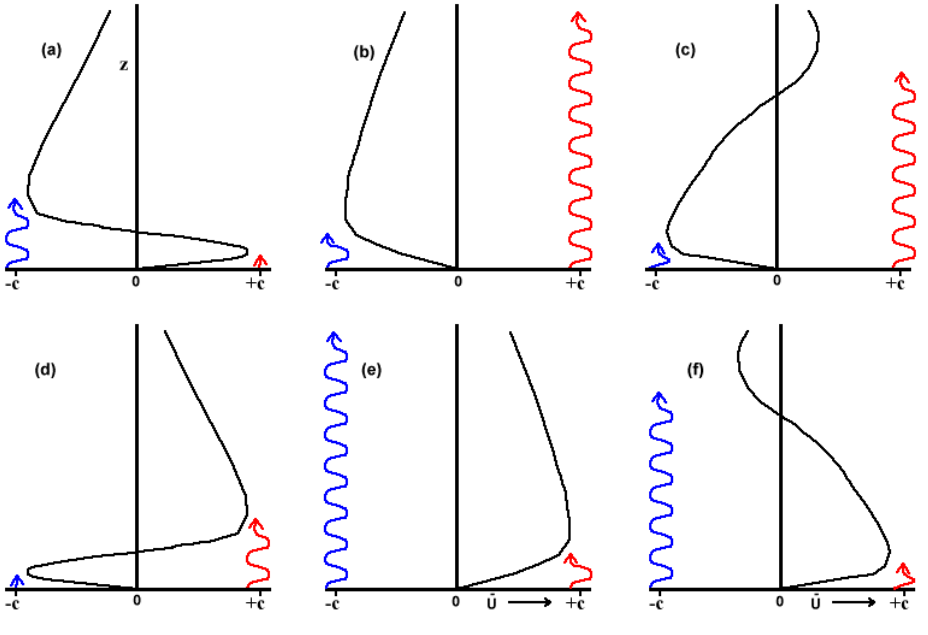

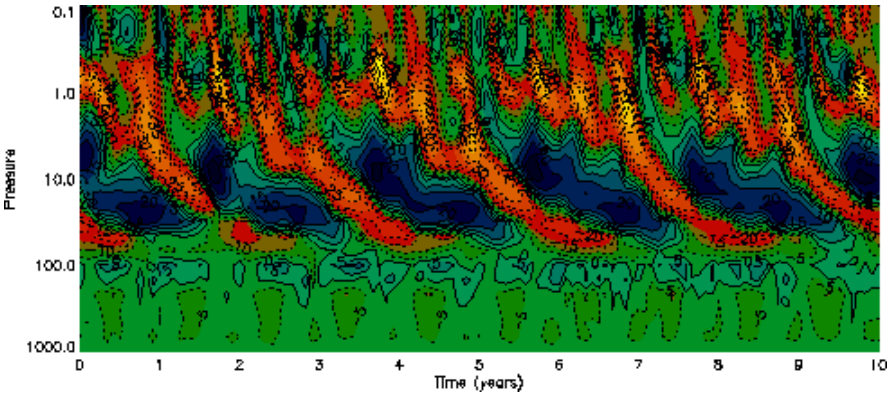


**QBO produced in the Unified Model by the Warner and McIntyre parameterization.**

## Use of the model – example 2

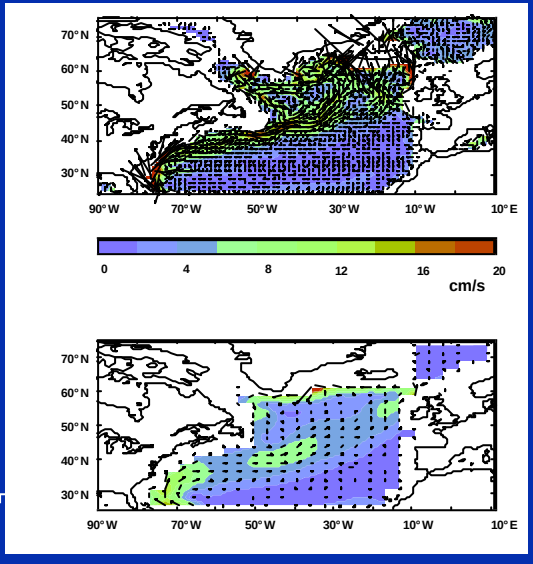


**Simulation of the Gulf Stream improves at higher ocean resolution**

**Above: Currents at 50m with 1.25 x 1.25° ocean (HadCM3)**

**Below: Currents at 50m with 2.5 x 3.75° ocean (HadCM3L)**

Hadley Centre for Climate Prediction and Research

# Results

Permission to quote these results must first be obtained from the Met Office.

# Porting problems

- 64 bit compiler – need reals, integers, logicals and variables to be the same size.
- MPP – wakeup message to an unopen file – a legacy from the days of Cyber205.
- Array bounds not set correctly on one routine.
- Upper/lower bounds set too high for 32 bits.

  i.e. 1e+99 doesn't exist at 32 bit.

## Single processor timings – 32 bit mode

| Processor | Compiler | Options | Minutes/ climate day |
|---|---|---|---|
| Alpha 667MHz | Compaq | None (equivalent to -O4) | 3.30 |
| Sun Blade 1000 750 MHz | Forte 6.0 | -fast -xarch=v8plusb | 3.71 |
| Pentium 4 1.4GHz | Compaq 6.6 (windows only) | Maximum optimisations and P4 code | 3.82 |
| Athlon 1.4GHz | Portland 3.2 | -Mvect=prefetch -O2 -pc 32 | 3.95 |

## Comparison of 32 bit to 64 bit model timings

| Processor | 32/64 bit | Minutes per climate day | Ratio to previous 32 bit run |
|---|---|---|---|
| Athlon 850 MHz | 32 bit | 7.07 | 1.0 |
| | 64 bit | 14.14 | 0.5 |
| Sun 750 MHz | 32 bit | 3.71 | 1.0 |
| | 64 bit | 5.30 | 0.7 |

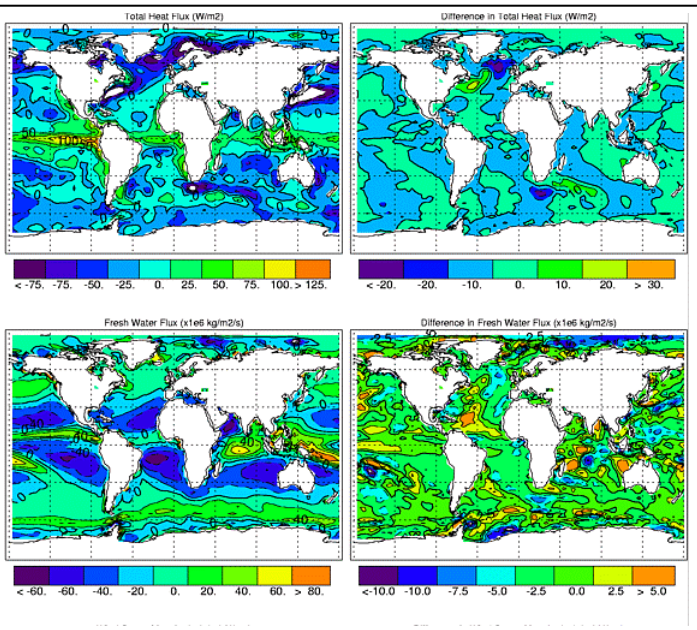


Decadal mean fields – left: 64 bit Cray run, right:difference between Cray and 32 bit PC run

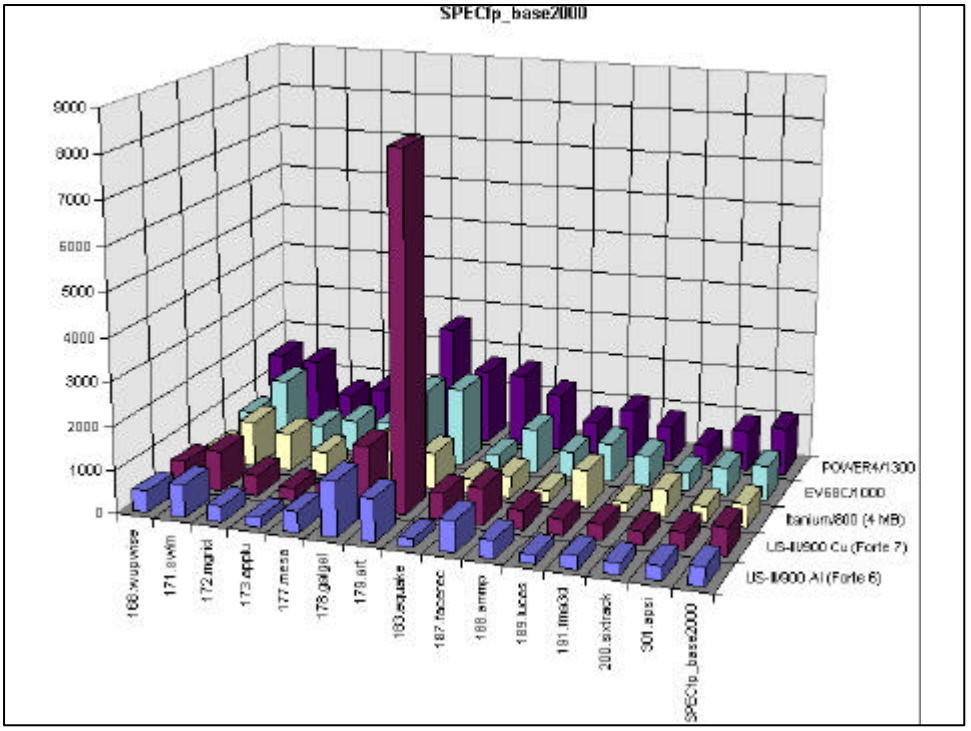

900MHz (Al) with Forte 6.0 vs. 1050MHz (Cu) with Forte 7.0

| Benchmark | Language | Category | Speedup |
|---|---|---|---|
| MHz | | | 1.17 |
| 168.wupwise | F77 | Physics / QCD | 1.59 |
| 171.swim | F77 | Shallow water modelling | 3.04 |
| 172.mgrid | F77 | Multi-grid solver: 3D potential field | 2.08 |
| 173.applu | F77 | Parabolic/Elliptic partial differential equations | 1.28 |
| 177.mesa | C | 3-D graphics library | 1.17 |
| 178.galgel | F90 | Computational fluid dynamics | 2.30 |
| 179.art | C | Image recognition/neural networks | 8.15 |
| 183.equake | C | Seismic wave propagation | 5.28 |
| 187.facerec | F90 | Image processing: face recognition | 1.45 |
| 188.ammp | C | Computational chemistry | 1.26 |
| 189.lucas | F90 | Number theory/Primality testing | 2.33 |
| 191.fma3d | F90 | Finite-element crash simulation | 1.42 |
| 200.sixtrack | F77 | High energy nuclear physics accelerator design | 1.29 |
| 301.apsi | F77 | Meteorology: pollution distribution | 1.55 |
| Ratio | | | 2.02 |

**Is galgel a good measure of this models performance?**

| **Processor** | **Galgel runtime** | **Unified Model (minutes/day)** | **Agree?** |
|---|---|---|---|
| Sun 750 MHz (Al) | 382 | 3.71 | - |
| P4 1.4 GHz | 383 | 3.82 | 3% out |
| Alpha 667 MHz (21264) | 349 | 3.30 | 3% out |
| AMD 1.4 GHz | 461 | 3.95 | 13% out |
| Sun 1050 MHz(Cu) | 139 | ? | ? |
| Power 4 690 Turbo 1.3GHz | 93 | ? | ? |

## MHz/Ghz – does it matter?

- 1.4 GHz P4 galgel = 383
- 2.0 GHz P4 galgel = 324
- Clock scaling 1.43
- Galgel scaling 1.18
- Power 4 1.3 GHz galgel = 93
- Equivalent P4 is >7 GHz

## MPP runs - Cray T3E at 64 bit

| Configuration | Minutes per climate day | Speedup |
|---|---|---|
| 1x1 | 14.88 | - |
| 2x1 | 7.86 | 1.89 |
| 2x2 | 4.33 | 3.44 |
| 3x3 | 2.25 | 6.61 |
| 4x4 | 1.67 | 8.91 |

# Initial MPP runs

| Number of longitude workstations | Number of latitude workstations | Minutes per climate day |
|---|---|---|
| 1 | 1 | 20 |
| 2 | 1 | 11 |
| 1 | 2 | 12 |
| 2 | 2 | 8 |
| 3 | 2 | 7 |
| 2 | 3 | 5 |

64 bit model runs.
333MHz Sun workstations on a 100Mb/sec switch

# Single 850 MHz Athlons with Myrinet 32 bit - maximum optimisations.

| Configuration | Minutes per climate day | Speedup |
|---|---|---|
| 1x1 | 7.07 | - |
| 2x1 | 4.04 | 1.75 |
| 2x2 | 2.33 | 3.03 |
| 3x3 | 1.09 | 6.49 |
| 4x4 | 0.72 | 9.82 |

# Dual 600 MHz processors using SCALI interconnect.

| Configuration | Minutes per climate day | Speedup |
|---|---|---|
| 1x1 | 8.37 | - |
| 1x1 | 10.62 | - |
| 2x2 | 3.89 | ? |

Preliminary results – better hardware now available.

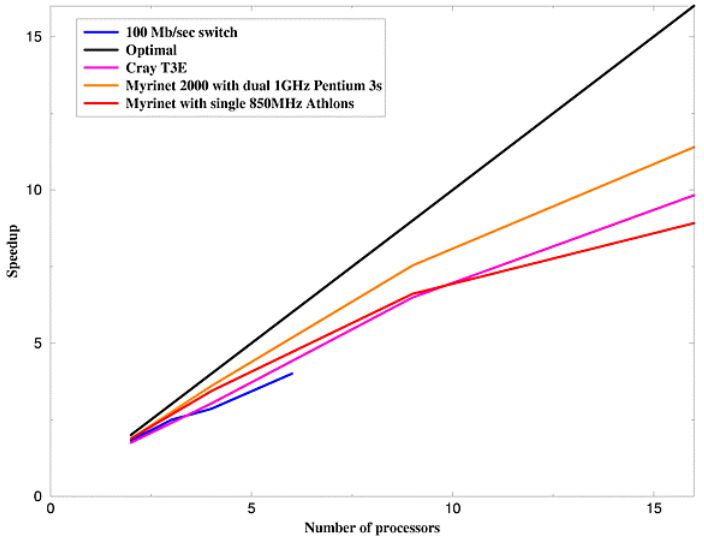

# Climateprediction.com

- Along the lines of the seti@home project.
- 100 years of coupled climate model starting in 1950 and running till 2050.
- Aiming for 2 million runs.

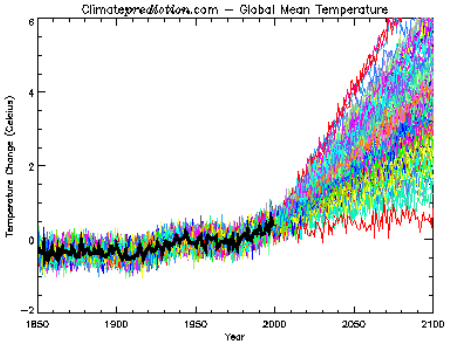


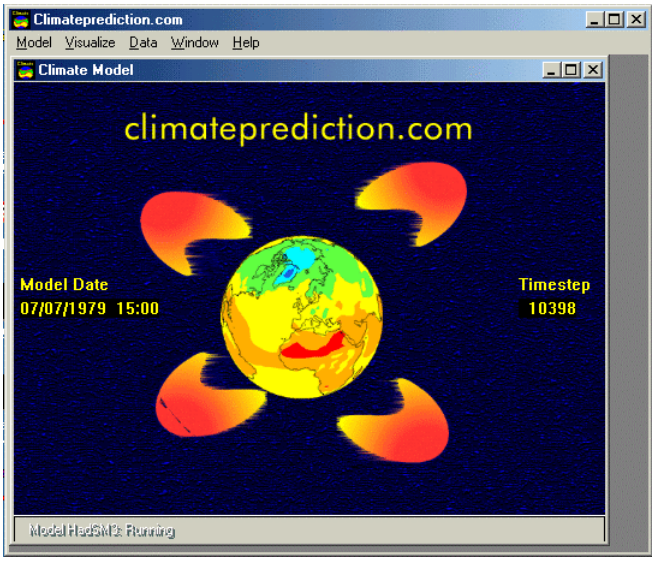

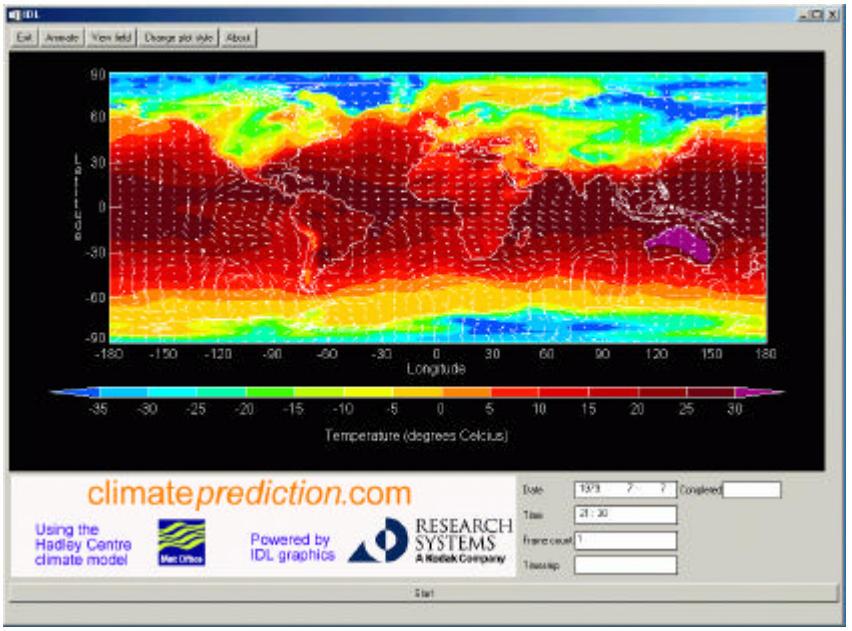
Exit
Animate
View field
Change plot style
About
90
60
30
0
-30
-60
-90
-180 -150 -120 -90 -60 -30 0 30 60 90 120 150 180
Longitude
-35 -30 -25 -20 -15 -10 -5 0 5 10 15 20 25 30
Temperature (degrees Celcius)
climateprediction.com
Using the Hadley Centre climate model
Met Office
Powered by IDL graphics
RESEARCH SYSTEMS
A Kodak Company
Date
Time
Frame count
Timestep
Completed
Start

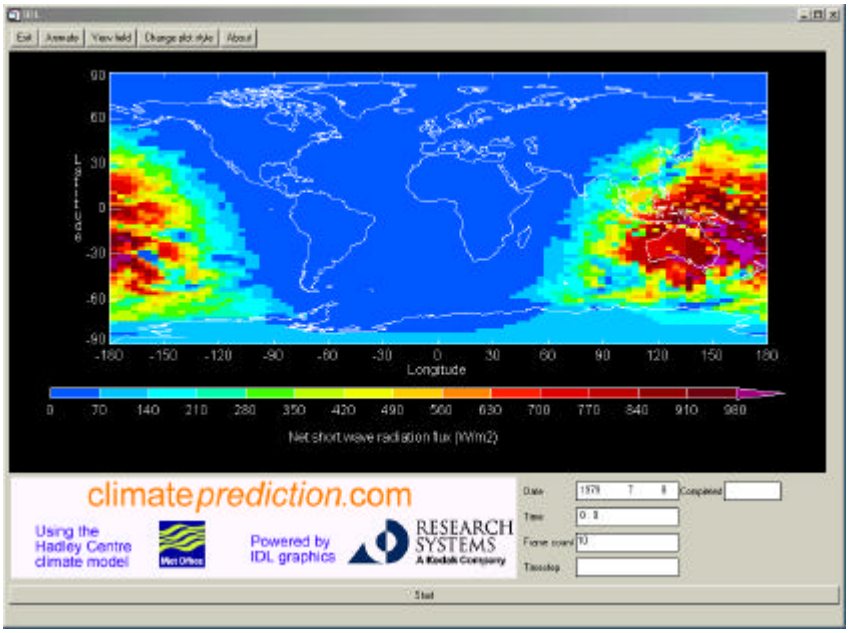
Exit
Animate
View field
Change plot style
About
90
60
30
0
-30
-60
-90
-180 -150 -120 -90 -60 -30 0 30 60 90 120 150 180
Longitude
0 70 140 210 280 350 420 490 560 630 700 770 840 910 980
Net short wave radiation flux (W/m2)
climateprediction.com
Using the Hadley Centre climate model
Met Office
Powered by IDL graphics
RESEARCH SYSTEMS
A Kodak Company
Date
Time
Frame count
Timestep
Completed
Start

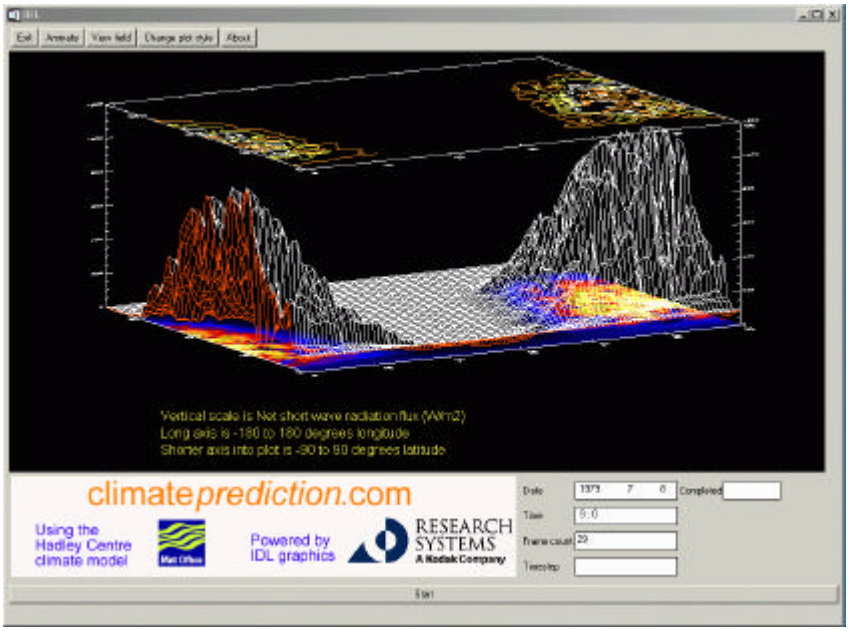
Exit
Animate
View field
Change plot style
About
Vertical scale is Net short wave radiation flux (W/m2)
Long axis is -180 to 180 degrees longitude
Shorter axis into plot is -90 to 90 degrees latitude
climateprediction.com
Using the Hadley Centre climate model
Powered by IDL graphics
RESEARCH SYSTEMS
A Kodak Company
Date
Time
Frame count
Timestep
Completed
Start

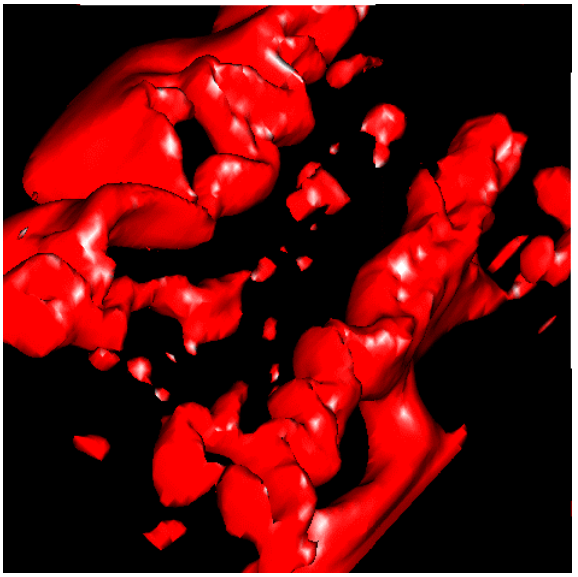

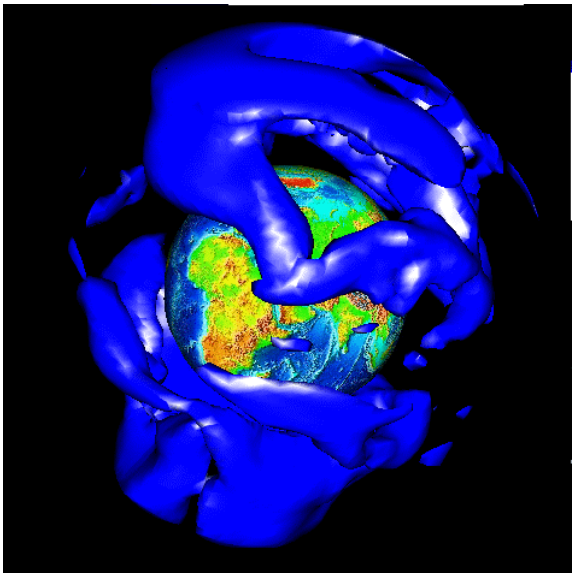

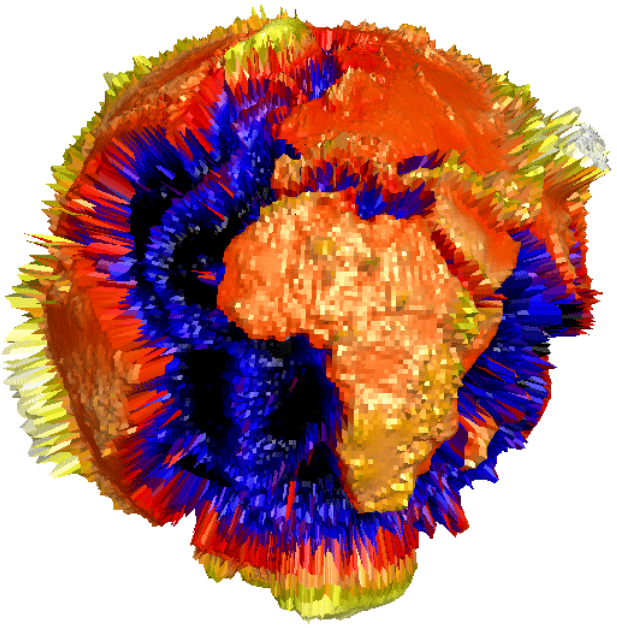

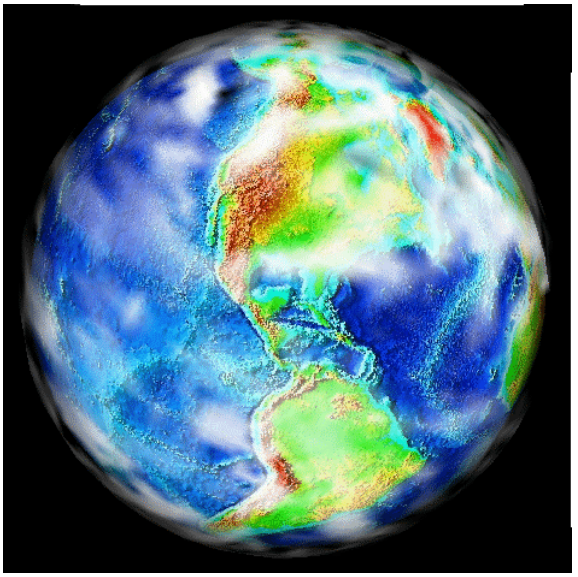

**http://www.climateprediction.com**

## Conclusions

- Always benchmark your code.
- Compilers are very important.
- Stream / Spec2000 are an indication of performance only.
- GHz aren't everything.
- Use whatever compute resources you have available to achieve your goal.