Using the CSIRO Mk3L climate system model
Part 1: Getting started

Steven J. Phipps
ARC Centre of Excellence for Climate System Science
Climate Change Research Centre
University of New South Wales

CLIM3001
7 May 2015

1. What is CSIRO Mk3L?

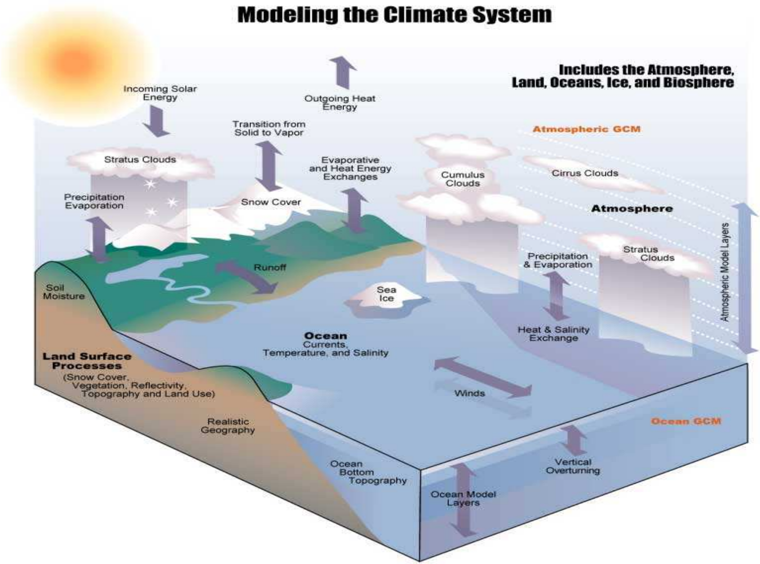
Modeling the Climate System
Includes the Atmosphere, Land, Oceans, Ice, and Biosphere
Incoming Solar Energy
Outgoing Heat Energy
Transition from Solid to Vapor
Atmospheric GCM
Stratus Clouds
Evaporative and Heat Energy Exchanges
Cumulus Clouds
Cirrus Clouds
Precipitation Evaporation
Snow Cover
Atmosphere
Precipitation & Evaporation
Stratus Clouds
Runoff
Atmospheric Model Layers
Soil Moisture
Sea Ice
Heat & Salinity Exchange
Ocean
Currents, Temperature, and Salinity
Land Surface Processes
(Snow Cover, Vegetation, Reflectivity, Topography and Land Use)
Winds
Ocean GCM
Realistic Geography
Vertical Overturning
Ocean Bottom Topography
Ocean Model Layers

Choosing the right model for you

- A model is a *tool* – the type that you use depends upon the question that you want to answer.
- Which components of the climate system do you need to model?
- Which processes do you need to model?
- Which quantities do you need to model?
- Do you need a regional or a global model?
- How much spatial resolution do you need?
- How long do you need to run the model for?
- These questions are inter-related – for example, it isn't feasible to run a high-resolution global model for 10,000 years!
- *No* model is a perfect representation of the real world.

Different types of models

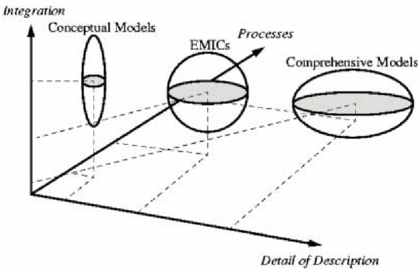

Fig. 1. Pictorial definition of EMICs. Adapted from Claussen (2000)

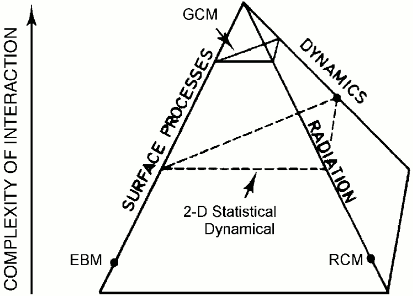

Fig. 2. The climate modeling pyramid. Adapted from Henderson-Sellers and McGuffie (1987)

Claussen et al. (2002), *Climate Dynamics*

The development of climate system models

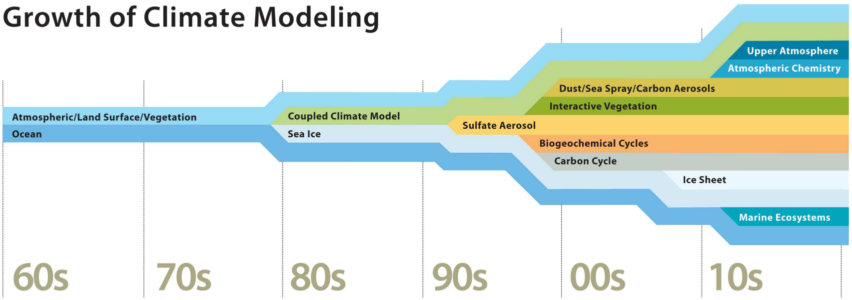

The first coupled atmosphere–ocean GCM

786 JOURNAL OF THE ATMOSPHERIC SCIENCES VOLUME 26

Climate Calculations with a Combined Ocean-Atmosphere Model

SYUKURO MANABE AND KIRK BRYAN

Geophysical Fluid Dynamics Laboratory, ESSA, Princeton University, Princeton, N. J.

13 March 1969 and 6 May 1969

Empirical evidence indicates that the poleward heat transport by ocean currents is of the same order of magnitude as the poleward transport of energy in the atmosphere (Sverdrup, 1957). A significant contribution to the heat exchange across latitude circles is also associated with polar pack ice. Thus, any serious attempt to calculate climate must take into account the entire fluid envelope of the earth, consisting of the atmosphere and the hydrosphere. Although the cryosphere, consisting of ice packs over the oceans and continental ice, is not a fluid in the usual sense, it must be included in a general climatic model because of its large reflectivity to the solar insolation and its ability to store and transport heat.

taken into consideration. Velocity, temperature, water vapor and surface pressure are calculated at each of the grid points which are spaced approximately 500 km apart. Calculations are carried out at 9 levels which are chosen so that they resolve the structure of the lower stratosphere and the Eckman boundary layer. The radiation model is essentially that described by Manabe and Strickler (1964). For the sake of simplicity, the seasonal and diurnal variation of solar insolation are not taken into consideration; instead, annual mean insolation is assumed for this study. The depletion of solar radiation and the transfer of terrestial radiation is computed taking into consideration cloud and gaseous absorbers such as water vapor, carbon dioxide and

The first coupled atmosphere–ocean GCM

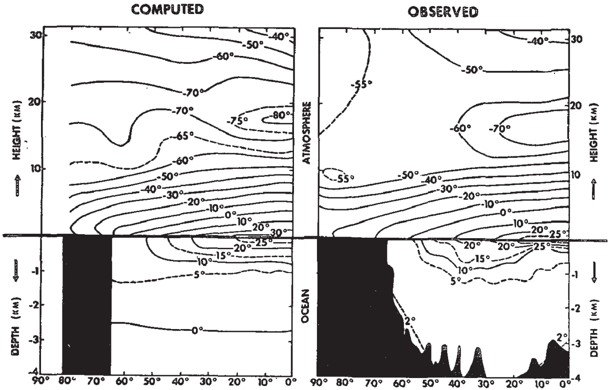

Fig. 2. Zonal mean temperature of the joint ocean-atmosphere system, left-hand side. This distribution, which is the average of two hemispheres, represents the time mean over two-sevenths of the period of the final stage of the time integration. The right-hand side shows the observed distribution in the Northern Hemisphere. The atmospheric part represents the zonally averaged, annual mean temperature. The oceanic part is based on a cross section for the western North Atlantic from Sverdrup *et al.* (1942).

Manabe and Bryan (1969), *Journal of the Atmospheric Sciences*

The CSIRO Mk3L climate system model

- Low-resolution version of the CSIRO climate system model (e.g. IPCC 1st, 2nd, 3rd, 4th and 5th Assessment Reports).
- Coupled atmosphere-land-sea ice-ocean general circulation model.
- Designed to enable millennial-scale simulations of climate variability and change e.g.
  - palaeoclimate simulations
  - projections of future climate
  - low-frequency climate variability
  - process studies
- Can simulate 1000 years in around a month.
- Community model.

The CSIRO Mk3L climate system model

- Atmosphere:
    - Three-dimensional general circulation model
    - Horizontal resolution of $5.6^\circ \times 3.2^\circ$ with 18 vertical levels
- Ocean:
    - Three-dimensional general circulation model
    - Horizontal resolution of $2.8^\circ \times 1.6^\circ$ with 21 vertical levels
- Sea ice:
    - Dynamic-thermodynamic sea ice model
    - Three layers (two ice, one snow)
- Land surface:
    - Soil-canopy scheme (13 land surface/vegetation types, 9 soil types)
    - Six soil layers, three snow layers

The topography of the land surface within CSIRO Mk3L

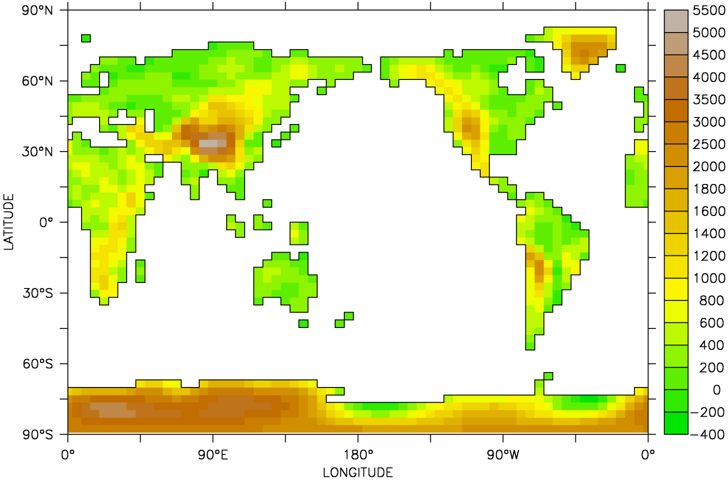

The bathymetry of the oceans within CSIRO Mk3L

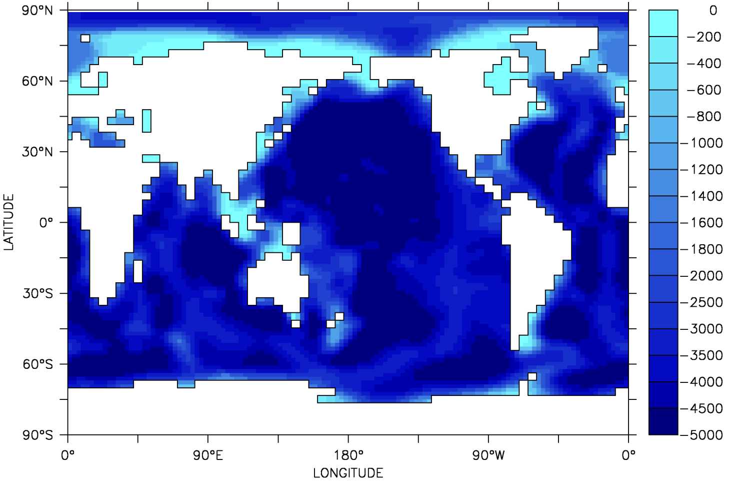

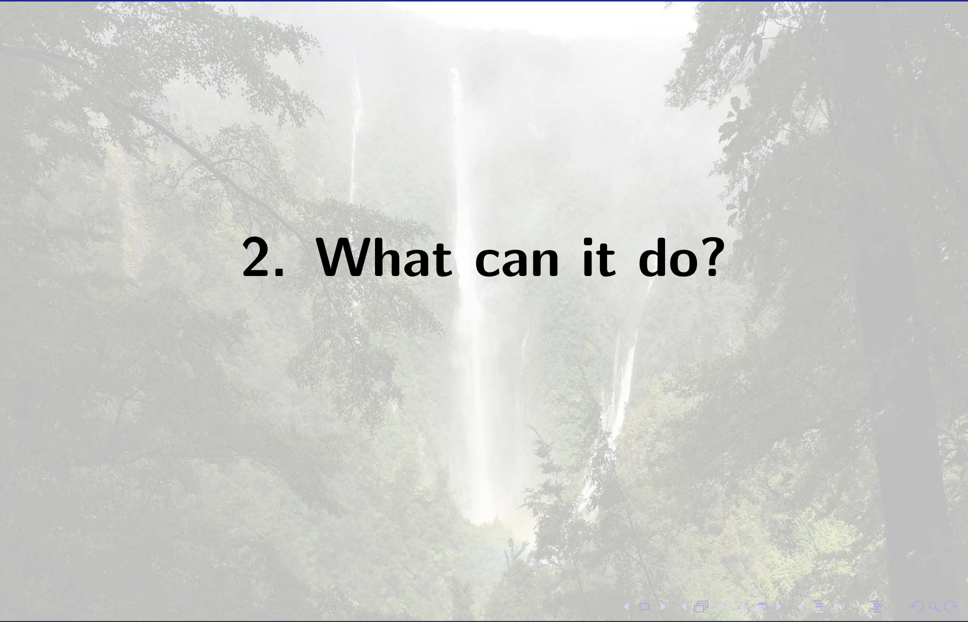

2. What can it do?

Simulated annual-mean surface air temperature (°C)

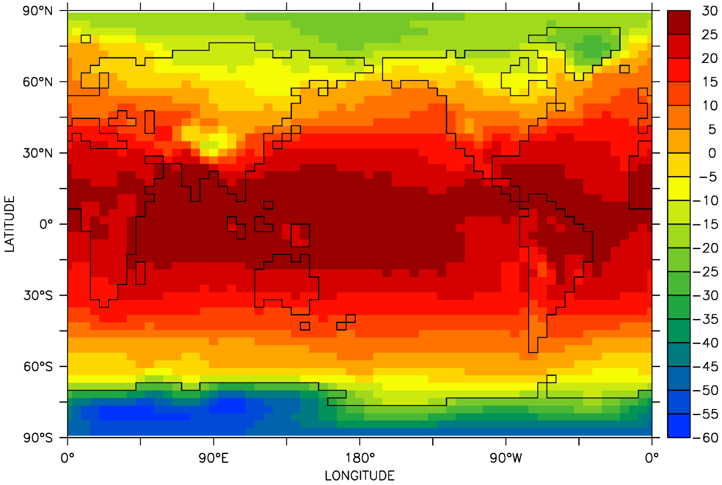

Simulated annual precipitation (mm)

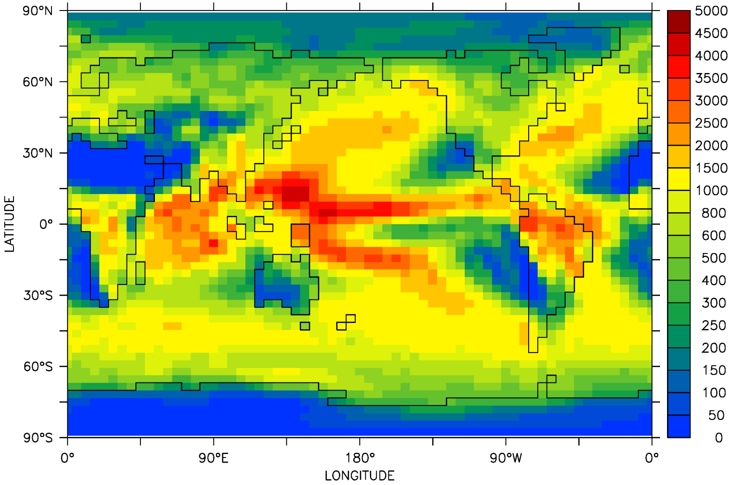

Simulated annual-mean sea surface temperature (°C)

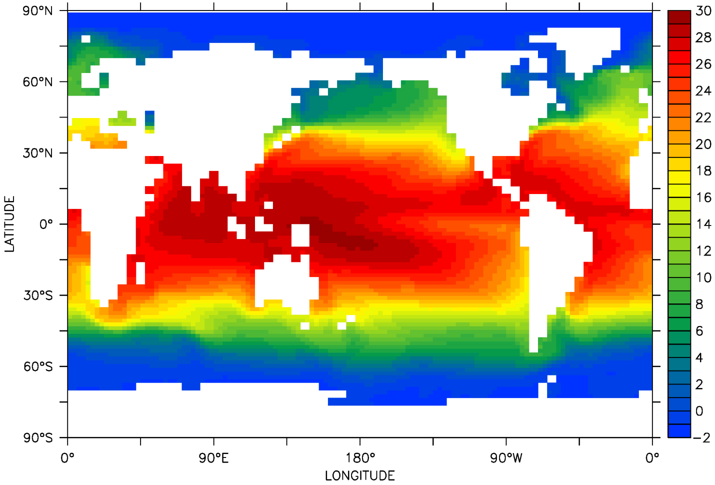

The evolution of SAM over the last millennium

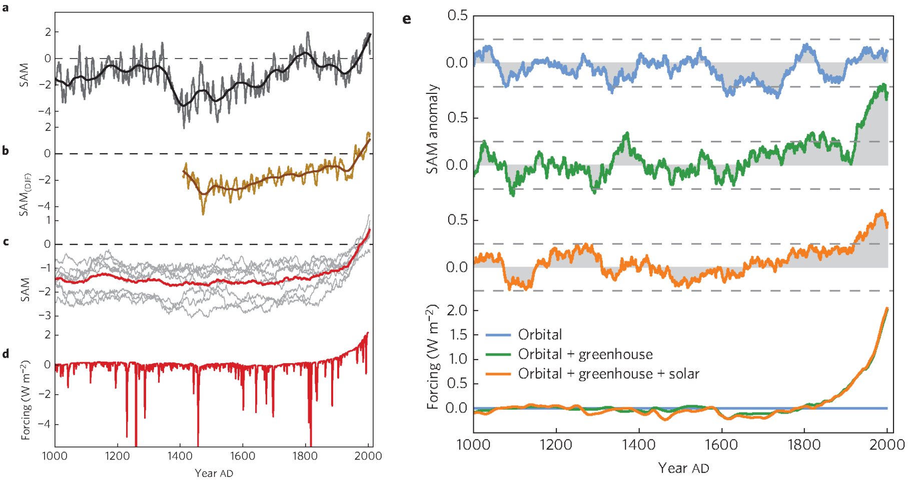

Abram et al. (2014), *Nature Climate Change*

The role of climate forcings over the past 1500 years

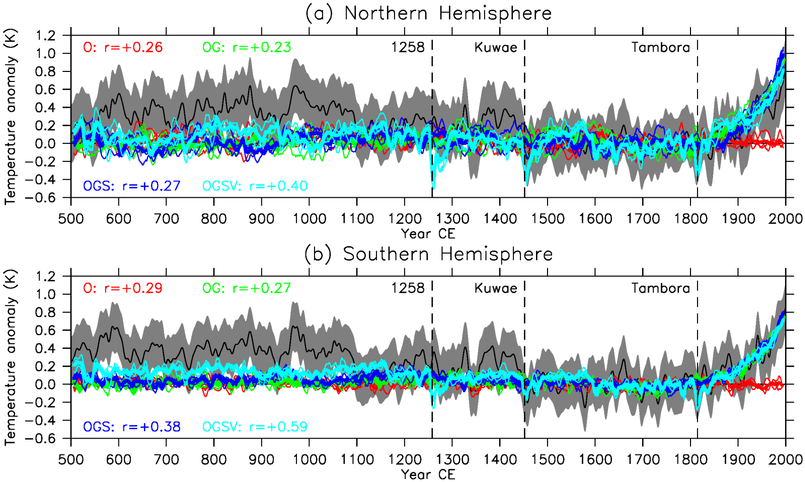

Phipps et al. (2013), *Journal of Climate*

The long-term response of the climate system to CO_2

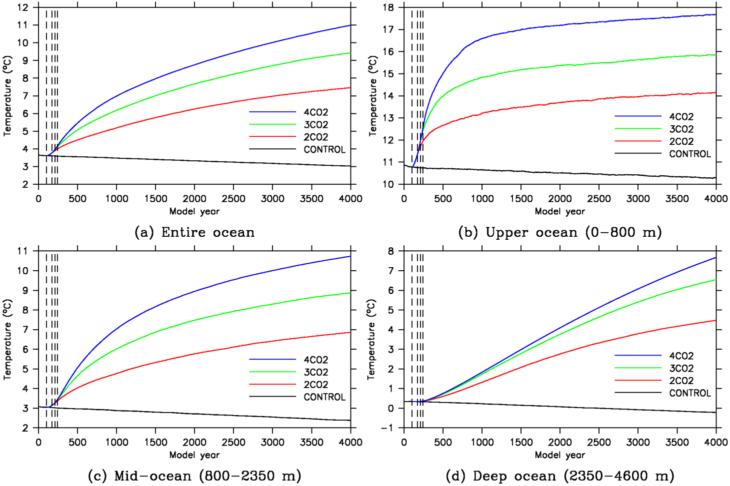

Phipps et al. (2012), *Geoscientific Model Development*

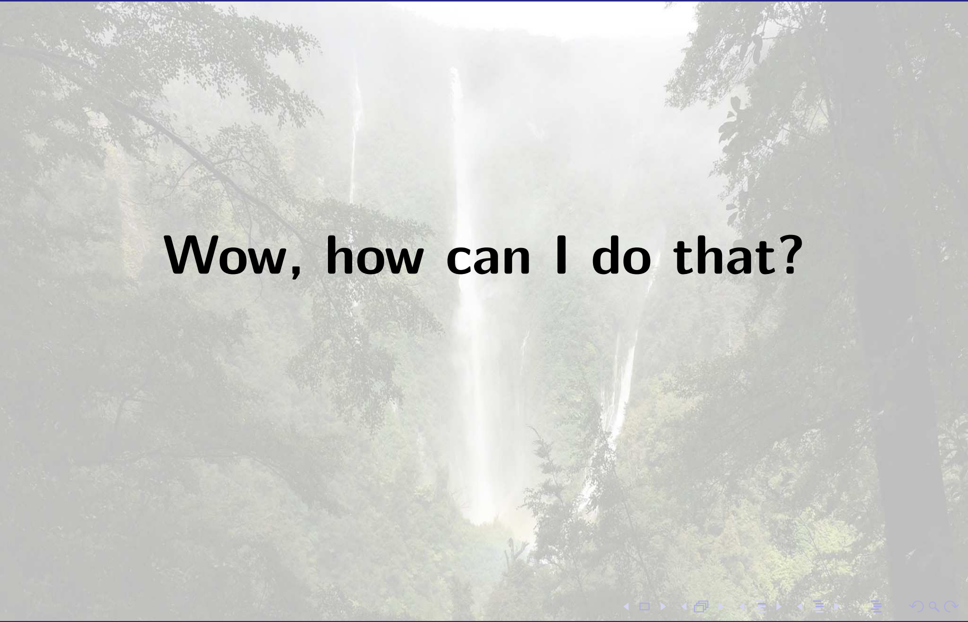
Wow, how can I do that?

3. Installing CSIRO Mk3L

High-performance computing, 1969-style

High-performance computing, 1969-style

Katana: A computational cluster

- Owned and used by UNSW Faculty of Science
- 140 x Dell blade compute nodes
- Total of 2,000 Intel CPU cores
- Linux operating system
- Portable Batch System (PBS) for running jobs
- Hostname is `katana.science.unsw.edu.au`
- For more information about Katana see:
  - `www.hpc.science.unsw.edu.au/cluster/katana`
- For more information about using the clusters see:
  - `www.hpc.science.unsw.edu.au/about/getting-started`

Exercise 1: Using Katana

- Launch Xming (`Programs > Xming > Xming`).
- Launch PuTTY (`Programs > PuTTY > PuTTY`).
- Using PuTTY, do the following:
  - Select `Connection > SSH > X11`
  - Check the `Enable X11 forwarding` box
  - Select `Session`
  - In the `Host Name` box, enter `katana.science.unsw.edu.au`
  - Click `Open`
  - Log in using your zNumber and zPass.
- Familiarise yourself with the basic Linux commands (see the next slide).

Basic Linux commands

`ls`	list the contents of a directory
`ls -l`	create a long listing
`mkdir <directory>`	create the directory `<directory>`
`cd <directory>`	change to the directory `<directory>`
`cp <file1> <file2>`	copy the file `<file1>` to `<file2>`
`mv <file1> <file2>`	move the file `<file1>` to `<file2>`
`rm <file>`	delete the file `<file>`
`rmdir <directory>`	delete the directory `<directory>`
`man <command>`	display the manual page for `<command>`

- For some more Linux commands see:
  - `www.dummies.com/how-to/content/linux-for-dummies-cheat-sheet.html`

Subversion

- Subversion is a *version control system*.
- Used to manage current and historical versions of files.
- Operates via the internet, allowing a community of users and developers to seamlessly share a piece of software.
- CSIRO Mk3L is managed and distributed using subversion.
- The repository is located at the Tasmanian Partnership for Advanced Computing in Hobart (but could be anywhere).
- For further information see:
  - `http://subversion.apache.org` (includes free book!)

Exercise 2: Getting CSIRO Mk3L

- We're not going to use subversion today.
- To save time, I've put a copy of the model distribution on Katana.
- Get version 1.2 of CSIRO Mk3L by entering these commands:

```
cd
mkdir CSIRO_Mk3L
cd CSIRO_Mk3L
tar zxvf /srv/scratch/z3210932/mk3l-1.2.tar.gz
```

- You'll see a lot of text scroll by as each file is extracted.

Exercise 2: Compiling CSIRO Mk3L

- Before you can run CSIRO Mk3L, you need to *compile* it.
- This turns the source code into a program that you can run.
- Compile the model by entering these commands:

```
cd ~/CSIRO_Mk3L/version-1.2/core/scripts/
./compile
```

- Compilation will take around two minutes.
- You will see text scroll by as each source file is compiled.

Exercise 2: Testing CSIRO Mk3L

- Before you can run CSIRO Mk3L, you should also *test* it.
- You can do this by entering this command:

```
./test_cpl
```

- This runs the full climate system model for one day. (It is possible to run the model in other modes, such as atmosphere only or ocean only, but we're not going to cover those in this course.)
- The test will take around 20 seconds to run.
- The model will display diagnostic text as it runs.
- If the test is successful, then the final line of text will be:

```
Stopped after  1 days.
```

4. Running CSIRO Mk3L

Running CSIRO Mk3L

- The basic command which runs Mk3L is simply:

```
./model < input
```

- `model` is the *executable*. This is the "model".
- `input` is the *control file*. This contains the instructions which tell the model what to do.
- The above command *executes* the model and feeds it the information contained within the control file.
- You've already seen what happens when you do this: a lot of diagnostic information is written to the command line.

Running CSIRO Mk3L

- The model is usually run using the command:

```
./model < input > output
```

- This command takes the diagnostic information generated by the model, and *redirects* it to an output file.
- For short jobs, the model can be run interactively.
- However, production jobs can take weeks or months to complete.
- We therefore need to use a *queueing system*.
- Katana uses the Portable Batch System (PBS).

Exercise 3: Running CSIRO Mk3L

- Run the model by entering this command:

```
qsub qsub_test_cpl
```

- This runs the model for one day.
- The `qsub` command submits a job to the queueing system.
- The file `qsub_test_cpl` tells the queueing system what to do.
- Use the command `qstat` to check the progress of your job:

```
qstat
```

- Hint: `qstat` lists all the jobs running on the cluster!

Exercise 3: Running CSIRO Mk3L

- The file `qsub_test_cpl` is called a *script*. The instructions contained within this file describe how to run the model.
- Using the `less` command, examine the contents:

```
less qsub_test_cpl
```

- Hint: type `q` to exit `less`.
- What does the script do?
- The lines beginning with `#` are comments.
- The lines beginning with `#PBS -l` tell the queueing system which resources are required to run the job.

Requesting resources

- When using a queueing system, you need to request sufficient resources to run your job.
- The script that you just ran uses three different options to do this:

`nodes`	The number of nodes to run on
`vmem`	The total amount of memory required
`walltime`	The expected run time

- It's important to request sufficient resources, but not *too* much.
- For further information see:
  - `www.hpc.science.unsw.edu.au/about/resource-requirements`

5. Output files

Output files

- When the model runs, it generates output. This is what you want!
- The model generates two types of output:

output files	save the state of the model *during* a simulation
restart files	save the state of the model at the *end* of a simulation

- The output files contain the simulated climate.
- See Chapter 6 of the Users Guide for further information.
- In common with almost all climate models, CSIRO Mk3L saves its output in a format called netCDF.

Exercise 4: Model output

- Get the course material for this week:

```
cd
tar zxvf /srv/scratch/z3210932/week1.tar.gz
```

- These commands create a new directory, `week1`, which contains some typical output from a coupled model simulation.
- Use the Linux command `ls` to examine the contents of this directory.
- You will see that the names of the files end with `.nc`.
- These are netCDF files.

What is netCDF?

- network Common Data Form.
- A self-describing, machine-independent data format.
- Probably the most common data format in the climate sciences.
- The names of netCDF files usually end with `.nc`.
- The command `ncdump` can be used to examine the contents of netCDF files.
- For further information see:
  - `http://www.unidata.ucar.edu/software/netcdf/`

Exercise 5: netCDF

- Load netCDF by entering the command:

```
module load netcdf
```

- Use `ncdump` to examine the contents of the sample atmosphere model output file, `stsc_spi62.nc`. Try commands such as:

```
ncdump -h stsc_spi62.nc
ncdump -c stsc_spi62.nc
ncdump stsc_spi62.nc | less
```

- What can you see?

Ferret

- A free visualisation and analysis package.
- Specifically designed for visualising climatic data.
- Makes it a breeze to visualise, analyse and manipulate the contents of netCDF files.
- Very powerful and easy-to-use averaging, interpolation and re-gridding capabilities.
- Your new best friend!
- For further information see:
    - `http://ferret.pmel.noaa.gov/Ferret/`

Basic Ferret commands

`use <file>`	Load the netCDF file `<file>`
`show data`	List the data which is available
`list <variable>`	List the values of `<variable>`
`plot <variable>`	Produce a line plot of `<variable>`
`shade <variable>`	Produce a shade plot of `<variable>`
`fill <variable>`	Produce a filled plot of `<variable>`
`contour <variable>`	Produce a contour plot of `<variable>`
`exit` or `q`	Exit

Basic Ferret transformations

- If the variable `tsc` contains surface air temperature as a function of longitude and latitude, then you can slice and dice the data using these expressions:

`tsc[i=10,j=8]`	Temperature at gridpoint (10, 8)
`tsc[x=140e,y=35s]`	Temperature at 140°E, 35°S
`tsc[x=90e:180e,y=45s:0]`	Temperature within 90–180°E, 45-0°S
`tsc[i=@ave]`	Zonal-mean temperature
`tsc[i=@ave,j=@ave]`	Global-mean temperature
`tsc[i=@max,j=@max]`	Global-maximum temperature
`tsc[i=@min,j=@min]`	Global-minimum temperature

Exercise 6: Ferret

- Now, load and run Ferret:

```
module load ferret
ferret
```

- Within Ferret, load the sample atmosphere model output:

```
yes? use stsc_spi62.nc
```

- This file contains data for surface air temperature.

Exercise 6: Ferret

- Try commands such as:

```
show data
fill tsc[k=1,l=1]
fill tsc[k=@ave,l=@ave]
fill tsc[i=@ave,k=@ave]
fill tsc[k=@max,l=@max]
plot tsc[i=@ave,j=@ave,k=@ave]
plot tsc[i=@ave,k=@ave,l=@ave]
plot tsc[x=140e,y=35s,l=@ave]
list tsc[i=@ave,j=@ave,k=@ave,l=@ave]
show transform
```

Exercise 7: Ocean model output

- A sample ocean model output file, `com.spi62.00001.nc`, is also provided. Examine the contents of this file using `ncdump` and Ferret.
- Within Ferret, try commands such as:

```
yes? use com.spi62.00001.nc
```

- Try commands such as:

```
shade/lev=1d temp[k=1,l=1]
fill/lev=1d temp[i=@ave,l=@ave]
fill/lev=2dc motg[l=@ave]
plot mota[y=30n:60n@max,k=@max]
```

- Table 6.1 of the Users Guide will be useful here.