NCAS CMS support for SWAMMA

William McGinty, Simon Wilson, Grenville Lister

NCAS-CMS, Dept of Meteorology, University of Reading

Introduction

SWAMMA (Sahara-West African Monsoon Multi-scale Analysis) is a four year project to investigate the influences of dust and convection on the timing of the West African Monsoon.

The project aims to study May to September in three different years 2006 (AMMA), 2011 and 2012 (Fennec). The following prognostic dust runs were planned.

Expt	Grid, km	Convection	Radiative Effects of Dust included
xlhua	12	Explicit	Yes
xlkkb	12	Parameterized	No
xlhua	12	Explicit	Yes
xlkkb	12	Parameterized	No
xlhpm	4	Explicit	Yes
xlhpn	4	Explicit	No

Domains

Both the 12km and 4km domains covered the same region

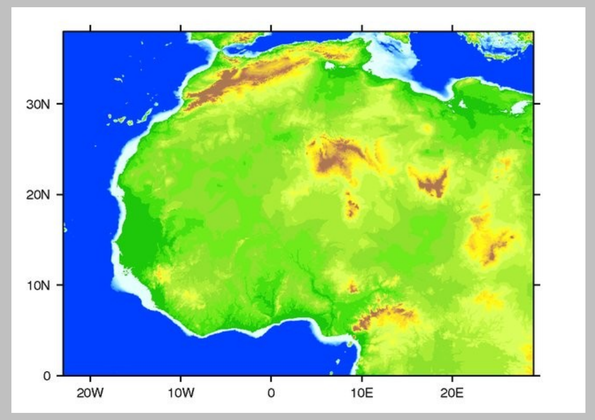

Grid, km	Columns	Rows	Levels
12	460	340	70
4	1380	1020	70

Model Development

CMS developed a 12km model based on a robust UM8.2 model designed for modelling in the Singapore region. CMS added a prognostic dust model.

To reduce data volumes and ease post processing the model included netCDF output data rather than the standard UM output. To increase the efficiency of the model output IO servers were used.

The model output 260 diagnostics, some on 70 levels. Some of the diagnostics revealed bugs in the code which had to be patched. Various test runs revealed problems with “out of memory” errors which were eventually traced to the Cray Compiler. These were overcome after two months work.

The 4km model was derived from the 12km model by changing the grid scale.

LBCs and CRUNs

LBC data was required at hourly intervals. Each of the 612 ECMWF analysis GRIB start dumps was therefore processed by an N216 global model to produce the LBC diagnostics and MAKEBC used to create the LBC files.

To allow MAKEBC to complete within the 24 hours allowed on the ARCHER serial queue, multiple LBC files 30 days long were created. An LBC switcher was developed to support the UM to enable continuous running.

CRUNs were therefore a sub multiple of the LBC length.

Data Transfer and Archiving

The project estimated that 300TB of data would be produced. Data transfer from Lustre (/work) to the RDF takes about 3 hours/TB. So the entire transfer would take about 37.5 days. Fortunately, data transfer can be overlapped with data production.

Data is also archived to the JASMIN data/archive processing facility. A similar transfer rate was achieved using the special JASMIN xfer2 node.

The large data volumes created during the 4km runs exceeded the quota on the ARCHER Lustre (/work) file system The LBC switcher was therefore augmented by a data pruner which backed up data to the RDF and then removed data leaving a few CRUNs worth.

Summary results

Grid, km	Expt	Cost MAU	Data Vol (TB)	Run duration (days)
12	xlhub	0.28	4.6	23
12	xlhua	0.27	4.6	24
12	xlkkb	0.26	4.6	15
12	xlkka	0.25	4.6	15
4	xlhpm	12.4	52	25
4	xlhpm	12.4	52	25
Total		25.9	122.4	

Highlights

After one year, CMS have delivered

- Initial GRIB data for AMMA 2006, Fennec 2011, Fennec 2012
- 12km and 4km UM8.2 models
- Ancillary data sets including SSTs
- Four 153 day 12km runs
- Two 153 day 4km runs

The models ran perfectly apart from ARCHER failures

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

Computational Modelling Services
NATIONAL CENTRE FOR ATMOSPHERIC SCIENCE
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk