

Article (refereed)

**Lenton, T. M.; Marsh, R.; Price, A. R.; Lunt, D. J.; Aksenov, Y.; Annan, J. D.; Cooper-Chadwick, T.; Cox, S. J.; Edwards, N. R.; Goswami, S.; Hargreaves, J. C.; Harris, P. P.; Jiao, S.; Livina, V. N.; Payne, A. J.; Rutt, I. C.; Shepherd, J. G.; Valdes, P. J.; Williams, G.; Williamson, M. S.; Yool, A. 2007 Effects of atmospheric dynamics and ocean resolution on bi-stability of the thermohaline circulation examined using the Grid ENabled Integrated Earth system modelling (GENIE) framework. *Climate Dynamics*, 29 (6). 591-613.
doi:10.1007/s00382-007-0254-9**

Copyright © Springer

This version available at <http://nora.nerc.ac.uk/1376/>

NERC has developed NORA to enable users to access research outputs wholly or partially funded by NERC. Copyright and other rights for material on this site are retained by the authors and/or other rights owners. Users should read the terms and conditions of use of this material at <http://nora.nerc.ac.uk/policies.html#access>

This document is the author's final manuscript version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this and the publisher's version remain. You are advised to consult the publisher's version if you wish to cite from this article.

www.springer.com

Effects of atmospheric dynamics and ocean resolution on bi-stability of the thermohaline circulation examined using the Grid ENabled Integrated Earth system modelling (GENIE) framework

T. M. Lenton* R. Marsh† A. R. Price‡ D. J. Lunt§
Y. Aksenov† J. D. Annan¶ T. Cooper-Chadwick‡
S. J. Cox‡ N. R. Edwards|| S. Goswami*
J. C. Hargreaves¶ P. P. Harris** Z. Jiao‡ V. N. Livina*
A. J. Payne§ I. C. Rutt§ J. G. Shepherd† P. J. Valdes§
G. Williams§ M. S. Williamson†† A. Yool†

January 31, 2007

Abstract

We have used the Grid ENabled Integrated Earth system modelling (GENIE) framework to undertake a systematic search for bi-stability of the ocean thermohaline circulation (THC) for different surface grids and resolutions of 3-D ocean (GOLDSTEIN) under a 3-D dynamical atmosphere model (IGCM). A total of 407,000 years were simulated over a three month period using Grid computing. We find bi-stability of the THC despite significant, quasi-periodic variability in its strength driven by variability in the dynamical atmosphere. The position and width of the hysteresis loop depends on the choice of surface grid (longitude-latitude or equal area), but is less sensitive to changes in ocean resolution. For the same ocean resolution, the region of bi-stability is broader with the IGCM than with a simple energy-moisture balance atmosphere model (EMBM).

*School of Environmental Sciences, University of East Anglia, Norwich, UK and Tyndall Centre, UK. E-mail: t.lenton@uea.ac.uk Tel: +44-1603-591414 Fax: +44-1603-591327

†National Oceanography Centre, University of Southampton, Southampton, UK and Tyndall Centre, UK

‡Southampton e-Science Centre, University of Southampton, Southampton, UK

§School of Geographical Sciences, University of Bristol, Bristol, UK

¶Frontier Research Centre for Global Change, Yokohama, Japan

||Centre for Earth, Planetary, Space and Astronomical Research (CEPSAR), Earth Sciences, Open University, Milton Keynes, UK

**Centre for Ecology and Hydrology, Wallingford, UK

††Tyndall Centre, UK and School of Physics and Astronomy, University of Leeds, Leeds, UK

Feedbacks involving both ocean and atmospheric dynamics are found to promote THC bi-stability. THC switch-off leads to increased export of salt at the S. boundary of the Atlantic that tends to maintain the off state. It can also generate net atmospheric freshwater input to the Atlantic that tends to maintain the off state. The ocean feedback is present in all resolutions, across most of the bi-stable region, whereas the atmosphere feedback is strongest in the longitude-latitude grid and around the transition where the THC off state is disappearing. Here the ocean response reverses, promoting THC switch-on by reducing Atlantic salt export, but the atmosphere counteracts this by increasing net freshwater input. This maintains some bi-stability even when the THC does switch on.

1 Introduction

There currently exist a spectrum of Earth system models (ESMs) and climate models (Figure 1) from the simplest box models to fully coupled atmosphere-ocean-land general circulation models (AOL-GCMs). State-of-the-art models such as the Hadley Centre model, are computationally too demanding for long-term integrations and ensembles are restricted in size, e.g. (Murphy *et al.*, 2004). Conversely, many existing efficient models of the coupled system, often termed Earth system models of intermediate complexity (EMICs) (Claussen *et al.*, 2002), e.g. CLIMBER-2 (Petoukhov *et al.*, 2000), employ highly idealized models of the individual components, with simplified physics, reduced dimensionality and low spatial resolution. Examples are the 2.5-D ocean model of Wright and Stocker (1991), the Statistical Dynamical Atmosphere model of Petoukhov *et al.* (2000), or the 2-D energy-moisture balance atmosphere model of Fanning and Weaver (1996). Alternatively, the Planet Simulator (Fraedrich *et al.*, 2005a) uses a fully dynamical 3-D atmosphere based on PUMA (Fraedrich *et al.*, 2005b) but only a 2-D mixed layer ocean. An exception is the fully 3-D ocean-atmosphere EMIC EcBilt/CLIO (Opsteegh *et al.*, 1998, Goosse and Fichet, 1999) and its further development LOVECLIM (Roche *et al.*, 2006), which achieve computational efficiency by using only 3 layers in the atmosphere.

In developing the Grid-ENabled Integrated Earth system modelling (GENIE) framework our aim is to populate the region of ‘model space’ between existing intermediate and full complexity models (Figure 1). Rather than develop a single model we are building a modular framework that allows different complexities of Earth system model to be created, by selecting different options for the various components (see Supplementary Information). The behavior one is trying to understand or predict, and its timescale, determine the components that need to be included. Our scientific drivers are to better understand changes in climate and the carbon cycle, on $10^3 - 10^6$ yr timescales, including the recent glacial-interglacial cycles (especially the last deglaciation), and in response to human activities. Our working hypothesis is that a realistic modelling framework for this purpose must include, as a minimum, component models of the atmosphere, ocean, sea-ice, marine biogeochemistry, marine sediments, land surface, vegetation and soil, and ice sheets (Figure 2). The framework must handle the exchange of energy, water, carbon and other biogeochemical tracers between components, ensuring rigorous conservation. Earth system models created from the framework must be fast enough to be integrated over multi-millennial time-scales, and to undertake large ensembles for sensitivity studies

and systematic model tuning. The design of the framework should allow other components, such as atmospheric chemistry, to be added at a later stage. It should also allow the continents to be reconfigured (i.e. change in the land surface mask, bathymetry and orography) in order to study events of interest in Earth history.

Computer power continues to increase exponentially, as described by Moore's Law; the number of transistors on an integrated circuit doubles approximately every 2 years. CPU speed, network bandwidth and physical storage have also all witnessed an exponential increase in capacity. Despite this, some compromise in model complexity is required to achieve our scientific goals now and in the next 5–10 years. Hence we have sought to enable traceability, meaning the ability to relate the simplified process representations and/or reduced resolutions used in our long-term and ensemble studies to more complete representations and higher resolutions used in state-of-the-art models. Support for modularity (i.e. interchangeable components) and scalability (i.e. variable resolution of the components) in the GENIE framework should help achieve this. It also allows us to couple in more complex and/or higher resolution component models in future as computer power increases.

In order to maximize access to available computer power we are developing software to run Earth system models on a "Grid" of distributed, non-specialist computing resources. "Grid computing" (Foster *et al.*, 2001) refers to systems that bring together people, resources (compute, data, network, sensors, devices, etc.) and services from across multiple administrative domains for a common purpose. The internet provides the infrastructure for such large scale computing across distributed domains. Grid technology supports the creation of dynamic "virtual organizations" to enable groups of individuals to collaborate. The work described herein brings together such a distributed group of Earth system modellers and computational scientists with the common purpose of building new models, executing them on distributed computing resources and sharing and recycling the data that they generate. Grid technology is used to ease the construction of new instances of Earth system model, automate the process of model tuning, enable large ensembles to be run, speed up the execution of long integrations, and recycle data back into model development. To make this a reality, it is particularly important to ensure that the Grid is useable directly from the Earth system modellers' working environment.

Here to demonstrate our modular framework and the use of Grid computing we undertake a new experiment to explore the stability of the ocean thermohaline circulation (THC) in different resolutions of a fully 3-D ESM. Changes in the THC are thought to have played a key role in past rapid climate changes, and the potential for shutdown of the THC is a key uncertainty in future climate projections. Theory and early models (Stommel, 1961) suggest that the THC exhibits bi-stability. Furthermore, all models systematically inter-compared by Rahmstorf *et al.* (2005) show bi-stability. However, the models differ in the position of the present climate state with respect to the region of bi-stability. Under sufficient freshwater forcing of the North Atlantic, the THC will collapse in all models, but in those starting in the bi-stable regime the collapse will be irreversible, whereas in those starting in the mono-stable regime the THC should recover. In the case of reversible THC collapse there may still be some hysteresis i.e. the recovery may be delayed. The failure (thus far) to find irreversible THC collapse in some AOGCMs suggests their initial climate state is in

a mono-stable THC regime, for example, HadCM3 bounces back from a THC collapse induced by freshwater hosing (Thorpe *et al.*, 2001) as does the AOGCM of Yin *et al.* (2006). In the GFDL model, the persistent collapse of the THC under a halving of pre-industrial CO₂ suggests that this change in boundary condition shifts the model into a region of bi-stability (Stouffer and Manabe, 2003). In contrast, under doubling CO₂ there is a reversible weakening, and under quadrupling CO₂ there is a collapse and later recovery indicating some hysteresis but not bi-stability and suggesting that these warmer model climates have a mono-stable THC.

Bi-stability of the THC can occur because a positive salt-advection feedback within the Atlantic acts to stabilize both the ‘on’ and ‘off’ states (Stommel, 1961). In more complex models it has been suggested that a larger scale salinity-overturning feedback, which can be positive or negative, is also critical in determining the stability regime (de Vries and Weber, 2005). Whether the Atlantic meridional overturning circulation (MOC) imports or exports salt at its Southern boundary is thought to determine the stability regime, the argument being that salt import (freshwater export) is required for bi-stability to occur (de Vries and Weber, 2005). Feedbacks from the atmosphere may also influence the stability regime of the THC, although de Vries and Weber (2005) find these to be of secondary importance in their model (EcBilt/CLIO). An atmospheric feedback involving net evaporation from the North Atlantic, in a pattern similar to that observed in El Nino events (Schmittner *et al.*, 2000), stabilizes the present ‘on’ state of the THC in some models (Latif *et al.*, 2000, Schmittner *et al.*, 2000). This does not deny bi-stability, indeed it may enhance it if there are also atmosphere feedbacks stabilizing the ‘off’ state of the THC.

It has been argued that increased variability in the THC induced by coupling to a fully dynamical atmosphere model would be expected to blur the boundaries of any bi-stable regime, and could potentially remove it (Schlesinger *et al.*, 2006). Thus far, studies using full primitive equation atmosphere models coupled to 3-D ocean models, have been restricted in their searching of model parameter space. Here we undertake a systematic search for THC bi-stability in a fully 3D ocean-atmosphere-sea-ice model. We employ the methodology of Marsh *et al.* (2004), following Wang and Birchfield (1992), of running the model to equilibrium under different atmospheric freshwater transport boundary conditions. To understand what determines the stability regime of the THC, we undertake an analysis of the freshwater budget of the Atlantic, distinguishing the atmosphere and ocean components, as suggested by de Vries and Weber (2005).

The paper is laid out as follows. In section 2 we provide a description of the model. Section 3 introduces the novel software and techniques which made the experiment possible. In section 4 we outline the experimental design. Section 5 reports the results with an emphasis on the mechanisms for THC bi-stability in different versions of our model. In section 6 we discuss these results in the context of previous model studies.

2 Model description

Here we briefly introduce the component models and the different ocean grids and resolutions (Table 1) used in the experiments described below. In the Supplementary Information, for each of the major components of the Earth

system (Figure 2) we discuss the appropriate modelling approaches for our goals and introduce all the corresponding model(s) that have been included thus far within the GENIE framework.

2.1 IGCM atmosphere

The Reading Intermediate General Circulation Model (IGCM3.1) has as its adiabatic core the 3-D spectral primitive equation model of Hoskins and Simmons (1975), with a default horizontal resolution of T21. In order to speed up the model for use in multi-millennial timescale integrations, the vertical resolution has been reduced from 22 levels to 7. Physical parameterizations include simplified versions of the turbulent flux scheme of Louis (1986), the convective adjustment scheme of Betts (1986), and the cloud scheme of Slingo (1987). The radiation scheme (which by default has no diurnal cycle) is based on a lookup table of transmittances in the longwave, and the two-band scheme of Morcrette (1990) in the shortwave. An earlier version of this model is described in some detail by de Forster *et al.* (2000); since their work, the surface scheme has been improved so that different vegetation types are associated with different roughness lengths and snow-covered and snow-free albedos, based on the lookup table used in the HadAM3 model (Pope *et al.*, 2000). The hydrological budget has been closed, and some restructuring of the original code has been carried out to increase modularity, including the integration of NetCDF input and output.

2.2 GOLDSTEIN ocean

The Global Ocean Linear Drag Salt & Temperature Equation Integrator (GOLDSTEIN) is a fast, intermediate complexity, 3-D frictional geostrophic model with linear drag. It incorporates eddy-induced and isopycnal mixing following Griffies (1998) and solves correctly for the flow around islands (Edwards and Marsh, 2005). The linear momentum balance of GOLDSTEIN allows for a local inversion of the baroclinic velocity field. At baseline resolution the model explicit timestep is around 3.5 days. Here we use two different grid types and three resolutions of GOLDSTEIN (Table 1): (1) The initial “baseline” (standard) resolution (36x36x8), on a longitude versus sin(latitude) horizontal grid, with quasi-logarithmic depth intervals (Edwards and Marsh, 2005). (2) A higher resolution (72x72x16) version, on the same longitude versus sin(latitude) horizontal grid. (3) An IGCM grid-matched (64x32x8) version, on a longitude versus latitude horizontal grid, which biases resolution towards the poles. Figure 3 shows the two grids and their realistic ocean bottom topography. In case (2), the higher horizontal and vertical resolution is implemented within the same surface grid and topography as case (1).

2.3 Slab sea ice

A very simple slab sea ice model has been separated from the IGCM. This is equivalent to a slab ocean component (see Supplementary Information), but with a slab thickness of 2m, and changes to the heat-capacity and albedo. In addition, there is the option to limit the implied ocean heat-flux term. For paleo simulations, tuning has indicated that in order to get a reasonable (defined here as being similar to that predicted by the HadSM3 model) simulation of modern,

pre-industrial, and LGM sea-ice area, it is necessary to limit the implied ocean heat flux for sea-ice to have an absolute value less than 50Wm^{-2} .

2.4 IGCM-land

The IGCM-land module was originally an integral part of the IGCM, as described in de Forster *et al.* (2000). In brief, it parameterises the surface energy, moisture, and momentum fluxes as a function of, respectively, temperature, moisture, and velocity vertical gradient. It also uses a bucket representation of soil-moisture, and parameterises changes in albedo related to snow cover. Some minor changes to the original scheme have been made so that it rigorously conserves energy and water.

2.5 Coupling

We couple the components described above in a modular fashion using the GENIE framework (for more details see Supplementary Information). Different resolutions of a particular model component are treated as different instances of that component. Hence the different ocean grids and resolutions can be readily interchanged at compilation time. To interpolate between different atmosphere and surface grids, we use a simple bilinear interpolation routine. In regions where there is a mismatch in land-sea mask, we extrapolate from the nearest neighbor. This process is in general non-conservative, so after interpolation we apply a correction factor in the tropics (30°S to 30°N) to ensure conservation of energy and moisture.

2.6 GENIE-2

By coupling the IGCM and GOLDSTEIN we have created the dynamical core of a fully 3-D Earth System Model (ESM). We use “GENIE-2” to describe the family of ESMs that use the IGCM dynamical atmosphere, because the choice of atmosphere model is the main determinant of overall model speed. The IGCM at T21 resolution integrates circa 10 years per CPU hour on an AMD 64 bit processor. In contrast the GENIE-1 model (Lenton *et al.*, 2006), which uses a single-layer Energy-Moisture Balance Model (EMBM) for the atmosphere, achieves over 1000 years per CPU hour at baseline resolution. Below we undertake a first large ensemble experiment using different grids and resolutions of GENIE-2, and contrast their behaviour with that of GENIE-1.

2.7 Parameter settings

Our versions of the fully coupled GENIE-2 have yet to be tuned. However, the IGCM with fixed ocean, fixed sea-ice, and IGCM-land has been tuned using a genetic algorithm (Price *et al.*, 2006) and we take 28 IGCM and IGCM-land parameter settings, ocean albedo and sea-ice albedo from that exercise (Table 2). We use default, untuned settings for the other IGCM, GOLDSTEIN and sea-ice parameters. The key GOLDSTEIN default parameters are shown in Table 2. Values for both isopycnal and diapycnal diffusivity and friction are relatively low compared to previous tunings of C-GOLDSTEIN (the version of GENIE-1 comprising GOLDSTEIN coupled to the EMBM and free-drift sea-ice)

(Hargreaves *et al.*, 2004, Edwards and Marsh, 2005). Winds from the IGCN are scaled by a factor of 2 (as with prescribed winds from the EMBM) to counteract the excessive surface drag in the frictional geostrophic ocean model.

3 Grid-enabled Problem Solving Environment

The collaborative Grid-enabled Problem Solving Environment we use for composing our model studies, accessing distributed computing resources, archiving, sharing and visualizing the results is built upon products of the first phase of the UK e-Science core programme (Hey and Trefethen, 2002), in particular the Geodise project (<http://www.geodise.org>). Their primary focus has been to provide solutions for design search and optimisation in the domain of aerospace engineering, resulting in a set of generic toolboxes for the Matlab and Jython environments (Eres *et al.*, 2005), a number of which we use.

3.1 Compute Toolbox

The Geodise Compute Toolbox provides intuitive high-level functions in the style of the hosting environment to allow users to easily manage the execution of a compute job on Grid resource. Functions are provided for three key activities:

Authentication In the UK e-Science community users are issued with a X.509 certificate by a trusted Certificate Authority. The toolbox enables the user with such a certificate to create a further time limited proxy certificate which effectively provides a single sign-on to the UK Grid. All subsequent activity on Grid resource (specifically resource implementing the Grid Security Infrastructure GSI (Welch *et al.*, 2003)) is authorised based upon the local rights of the account belonging to the certificate owner. Functionality is provided to instantiate, query and destroy proxy certificates.

File transfer The GridFTP (Allcock *et al.*, 2002) data movement service of the Globus Toolkit (2.4) (Foster and Kesselman, 1999) is exposed to the Matlab client through a set of functions that wrap the Java Commodity Grid (CoG) kit (von Laszewski *et al.*, 2001). Methods for transferring files to and from a GridFTP enabled resource are provided.

Job submission The user is enabled to execute work on resource managed by either the Globus Toolkit (implementing the Grid Resource Allocation and Management service (GRAM) (Czajkowski *et al.*, 1998)) or Condor (Thain *et al.*, 2005). By providing information describing the compute task (executable, input files, environment variables) the interface allows submission of jobs to the resource broker of the remote resource. Functions are provided to monitor the status of the job handles returned after submission and to kill those jobs if necessary.

3.2 Database Toolbox

The Geodise data management model allows local data (files, scripts, binaries, workspace variables, logical data aggregations) to be archived in a shared central

repository and for rich descriptive metadata to be associated with that data. The data can be archived to, queried in, and retrieved from the repository. The interface to the database is exposed using Web Services allowing users access to the same repository from distributed locations using standard web protocols. Files are stored on a GridFTP server hosted by the UK National Grid Service (<http://www.ngs.ac.uk/>). The Geodise Database Server has been augmented by mapping an XML Schema into the database to restrict the permissible metadata describing entities in the database. This significantly improves query and retrieval performance in the underlying database. Both programmatic and GUI interfaces are provided to the data repository allowing easy navigation of the data and enabling the database to be an integral part of scripted workflows.

3.3 OptionsMatlab

The GENIE client for Matlab also includes an interface to a third party design search and optimisation package, OPTIONS (Keane, 2003), that has been developed in the Computational Engineering and Design Centre at the University of Southampton. This software provides a suite of sophisticated multidimensional optimisation algorithms developed primarily for engineering design optimisation. The package has been made available to Matlab via the OptionsMatlab interface and has been exploited in conjunction with the Geodise Toolboxes to tune the IGCM model parameters (Table 2) (Price *et al.*, 2006).

3.4 GENIE Toolbox

A higher level abstraction of the Geodise functionality has been developed to provide intuitive management of time-stepping codes on the Grid. Scripted workflows wrapping the Geodise functions have been written to provide a uniform interface for the execution of GENIE Earth system models on local, Globus and Condor resource. The configuration and execution of a simulation is enabled through a single function call which accepts as input data structures describing the model instance (parameter settings, input files, etc.), the local runtime environment and the remote resource on which to execute the model. Further functionality is provided to coordinate the execution of ensemble studies mediated by the database. The toolbox methods enable users to expose models as tuneable functions or include the database as an integral part of a large ensemble study.

Ensemble studies are defined by creating a data structure within the database. This consists of a parent entity describing the experiment and contains a set of sub-entities describing the individual simulations comprising the ensemble. A simple template script is provided to aid the user in specifying the logical entities in the database. The subsequent execution of the ensemble is performed by autonomous client “worker” scripts that interrogate the database for work units, submit that work to available resource and post-process completed tasks. Similar systems such as Nimrod/G (Abramson *et al.*, 2002) and GridSolve (YarKhan *et al.*, 2006) provide bespoke definition languages for task farming studies but rely on a central agent to manage the execution of the study. Our system devolves responsibility for task farming to the client where informed decisions can be made based upon point-in-time queries on the contents of the

database. This enables more flexible integration of heterogeneous resource but means that no central control can be imposed on the study. A guarantee of the completion time for an ensemble study cannot be made but resource can be dynamically introduced or retracted by users contributing to the experiment.

3.5 Fault tolerance

Simulations are mediated through the database and are only progressed through the successful upload of completed work units into the repository. If the configuration, submission, execution, post-processing or data upload fail for a work unit then the client system will detect a problem and attempt corrective action. This may include a repeated attempt to perform the failed stage of progression, removal of the job from the database allowing the entire work unit to be attempted again or marking the simulation as failed (in cases where repeated failures have been detected). The means to tag a simulation as ‘failed’ is provided because studies may span areas of parameter space where models become numerically unstable. This fault tolerance also holds for client-side problems or network outages. If completed work units are not successfully uploaded to the database the study is unchanged and the work can be attempted again.

3.6 Brokering strategy

In the absence of a brokering service on the UK National Grid our data management system provides a means to maximize our responsible use of the available resource. A user provides metadata about the resources they want to use including the maximum number of jobs that they would like to submit to the job manager (PBS, SGE, condor) for any experiment. The database maintains a record of where a users jobs are active and the client will respect the usage limits and refrain from submitting work to a resource if it has already reached its maximum number of allowed jobs. A user can then set up client invocations to automatically attempt to push work to a list of resources. Once each resource reaches its limit the client will move on to the next and the system therefore keeps all available compute power busy (assuming sufficient work exists) without exceeding the users specified usage limits.

4 Experiment design

For the baseline resolution (36x36x8) we find that the resulting model climate spins up to a stable state. However, the ocean thermohaline circulation is reversed with sinking in the North Pacific and upwelling in the North Atlantic. This occurs largely because of insufficient transport of freshwater from the Atlantic to the Pacific by the atmosphere model. We diagnose the differences between IGCM freshwater transports and NCEP reanalysis data, and group them into three latitudinal sectors that correspond to the nearest on the IGCM grid to those used by Marsh *et al.* (2004): 0.41 Sv (28.125–90°N), 0.30 Sv (28.125°N–28.125°S), 0.08 Sv (28.125–50.625°S). We then apply flux corrections of these sizes in the three sectors, removing freshwater from the Atlantic and adding it to the Pacific in each case. The flux corrected model spins up to a qualitatively correct thermohaline circulation with sinking in the North Atlantic and upwelling

in the Pacific. We refer to this from hereon as the default flux correction.

Whilst the use of an untuned model with the need for this type of flux correction would be undesirable for predictive studies, we can make positive use of it for the idealized experiments conducted here. Following Marsh *et al.* (2004), we undertook 1-parameter experiments varying the imposed Atlantic to Pacific freshwater flux transported by the atmosphere. In a separate 2-parameter experiment, we identified and also varied an IGC parameter somewhat analogous to atmospheric moisture diffusivity in the EMBM, which primarily determines equator to pole moisture transport. In order to search for bi-stability of the THC we used restarts from chosen end-members of the initial ensembles, and to keep this computationally tractable we concentrated on restarts of the 1-parameter experiments.

The same default Atlantic-Pacific freshwater transport was applied in all 3 model resolutions, and in all the experiments it was varied by simply multiplying all three components of it by the same scaling factor. In initial 1-parameter ensembles for each of the 3 model resolutions, the scaling factor was varied between 0 (no flux correction) and 2 (twice default) in 21 steps. In all three resolutions, the switch between THC on and reversed states was found to occur in the range 0-1 for the scaling factor. This helped us determine 6 ensembles of restart experiments. For each of the 3 resolutions we ran ensembles from 2 different restart states – the end of the zero and default flux correction runs of the original ensemble. In the restart experiments, the scaling factor was varied between 0 and 1 in 21 steps, in order to better resolve switches between THC states. To examine the effect of compiling for different operating systems, we replicated our restart ensembles for one model resolution (72x72x16) with Linux and Win32 binaries.

For the 2-parameter experiment we used the baseline ocean resolution and chose as a second parameter `ia_enhancestable`, which is a constant multiplier at the front of the evaporation (hence latent heat) and sensible heat terms in the IGC. Increasing it is expected to increase freshwater sources to the atmosphere, which are greatest from the equatorial oceans. The Atlantic-Pacific freshwater flux correction was varied in 11-steps from 0 to 2 times default, and `ia_enhancestable` was varied in 11-steps on a logarithmic scale from 0.1 to 10 times its default value of 0.3731 (Table 2). This gave a 121 member ensemble, with each run for 2000 years in this case.

Table 3 summarises the 12 ensemble studies conducted. These were defined in the database and executed across the computational Grid by distributed members of the project team (see Appendix for further details). From the results of our 1-parameter experiments we selected some model versions and undertook 0.1 Sv and 1.0 Sv freshwater hosing experiments for 100 yr as in Stouffer *et al.* (2006), to compare with their model inter-comparison.

5 Results

The thermohaline circulation (THC) exhibits a fairly high degree of variability under the dynamical atmosphere model, with a range in the maximum Atlantic Meridional Overturning Circulation (MOC) of order ± 5 Sv. This THC variability is driven by the internal variability of the dynamical atmosphere model. Thus if atmospherically driven variability and associated feedbacks blur the

bi-stability of the THC we expect the effect to be fairly strong in our model.

5.1 Initial spin-ups

For the default flux correction, the maximum Atlantic MOC is rather strong in all GENIE-2 versions, being ≈ 42 Sv in 64x32x8, ≈ 35 Sv in 72x72x16 and ≈ 32 Sv in 36x36x8. In the annual average, the maximum of the Atlantic MOC is located at around 50°N in all resolutions (Figure 4), although in winter (not shown) it tends to shift to lower latitudes around 20°N. Both high and low latitude cells in the North Atlantic are apparent in the annual average, especially for 72x72x16 (Figure 4e). A strong low latitude cell in winter may be expected due to seasonal meridional shifts and changes in strength of low-latitude winds driving changes in surface Ekman transport, compensated by deep “sloshing” motions (Jayne and Marotzke, 2001). In the annual-average, these motions tend to cancel out, so that sinking is largely confined to higher latitudes.

Our initial 1-parameter experiment spin-ups revealed non-linear transitions in the strength of the maximum Atlantic MOC strength as a function of Atlantic-Pacific freshwater flux correction, for all three ocean resolutions (Figure 5). The amount of flux correction required to get the Atlantic MOC ‘on’ is noticeably less for the longitude-latitude surface grid 64x32x8 relative to the other two. The transition is in a similar place for the 36x36x8 and 72x72x16 resolutions but sharper for the latter. In all resolutions, removing flux correction gives an Atlantic MOC ‘off’ state and default flux correction gives an Atlantic MOC ‘on’ state, providing start points for our restart experiments.

In the 2-parameter study only 88 of 121 of the defined ensemble members completed successfully. The other 33 members were found to be in unstable regions of parameter space and failed quickly (at both high and low values of `ia_enhancestable`). Maximum Atlantic MOC varies with the 2 parameters (Figure 6) in a qualitatively similar way to results with GENIE-1 presented in Marsh *et al.* (2004). Increasing `ia_enhancestable` in the IGCM destabilises the Atlantic MOC, just as increasing atmospheric moisture diffusivity does in the EMBM. We infer that the increased evaporative flux under increased `ia_enhancestable` has a disproportionately larger effect in the tropics promoting an increased equator to pole atmospheric moisture transport, and thus tending to destabilise the THC. When `ia_enhancestable` is reduced to 0.15–0.2 of its default value, no Atlantic-to-Pacific freshwater flux correction is required to get the Atlantic MOC ‘on’ at ≈ 20 Sv. These results will be examined further in future work. From hereon we focus on the search for bi-stability varying only the Atlantic-to-Pacific freshwater flux correction.

5.2 Bi-stability

From the restart experiments, we find bi-stability of the Atlantic MOC in all 3 resolutions of GENIE-2 (Figure 7). The region of bi-stability is noticeably wider for the longitude-latitude grid 64x32x8 and shifted to lower values of flux correction. The bi-stable region is narrowest for 72x72x16 and the transitions occur at the largest values of flux correction. The baseline 36x36x8 model has the least defined transitions and there is some indication of bi-stability remaining at the default flux correction (although the Atlantic MOC is ‘on’ in both cases). Interestingly, for 64x32x8, the restarts from Atlantic MOC ‘on’

produce a transition in a different place ($f = 0.1$ – 0.15 of default flux correction) to the original spin-up experiments ($f = 0.2$ – 0.3) (Figure 5). This is not the case in the other two resolutions.

For each ocean resolution, we have examined the overturning streamfunction at a range of values of default flux correction that give rise to bi-stability of the Atlantic MOC (based on Figure 7). Illustrative cases (Figure 8) are $f = 0.8$ of the default flux correction in resolution $36\times 36\times 8$, $f = 0.85$ in $72\times 72\times 16$, and $f = 0.55$ in $64\times 32\times 8$. These are chosen to be close to the transition from bi-stability toward a mono-stable ‘on’ state, and the striking reversal of the Atlantic MOC is accompanied by a simultaneous reversal of the Pacific MOC (not shown). Thus, the overall THC reverses. At the lower end of the bi-stable region, on the upper branch, the positive cell of the Atlantic MOC often fails to reach the Southern Ocean and the Pacific MOC tends to have already reversed. This tendency for the Pacific circulation to reverse is partly due to the low (un-tuned) value used for the diapycnal diffusivity ($10^{-5} \text{ m}^2\text{s}^{-1}$).

Surface currents (not shown) vary considerably between the different model resolutions and change noticeably when switching between bi-stable states. Inertial currents such as the Antarctic Circumpolar Current (ACC) and the Gulf Stream are expected to be too weak due to missing dynamics in the frictional geostrophic ocean. In these un-tuned model versions, surface currents are noticeably stronger in $72\times 72\times 16$ than in either $64\times 32\times 8$ or $36\times 36\times 8$. The main North and South Pacific gyres are present in all model resolutions. The Atlantic gyres are less clear, but on going to higher resolution ($72\times 72\times 16$), the Gulf Stream becomes clearer and stronger. When the THC reverses, the Gulf Stream weakens and is deflected southwards and the Kuroshio current strengthens and moves northwards.

Equivalent states of the $72\times 72\times 16$ model using the Linux and Win32 binaries typically have small differences in spatial patterns. Restarts from the same state of the $72\times 72\times 16$ model conducted with Linux and Win32 binaries diverge rapidly, indicating a sensitivity to initial conditions, which is to be expected with a dynamical atmosphere model. However, the average behavior of the THC is similar – for example, in restart runs where the flux correction is such that the THC will wind down and then collapse, the timing of changes in the THC is the same in the Win32 and Linux binaries despite their differing inter-annual variability (results not shown).

5.3 Responsible mechanisms

To understand the mechanisms responsible for maintaining bi-stability and to elucidate the difference between different model grids and resolutions, we have examined the feedbacks from the ocean and atmosphere under changes in the THC state. Once again, for each resolution, we consider our illustrative examples of bi-stable states. Then we examine the responses of aggregate variables as a function of default flux correction. In the following sections, we use ‘‘THC’’ to refer to the overall circulation and ‘‘MOC’’ to refer specifically to a mode of transport diagnosed in the model.

5.3.1 Ocean feedback

Freshwater transport by the Atlantic MOC varies significantly as a function of latitude, and changes significantly when the THC switches state (Figure 9). For our illustrative runs, in the ‘on’ state, in all resolutions the self-sustaining nature of the THC can be seen in the N. Atlantic, where it transports freshwater southward. In 36x36x8 and 72x72x16, there is little net import or export of freshwater at the Southern boundary of the Atlantic, and when the THC switches off, there is a pronounced import of freshwater at the Atlantic S. boundary. By contrast, in 64x32x8, in the ‘on’ state, the MOC is importing freshwater at the Atlantic southern boundary, and there is little change when the THC switches off. A shift to increased import of freshwater should tend to maintain the ‘off’ state in the illustrative 36x36x8 and 72x72x16 runs, but is not apparent in the illustrative 64x32x8 run.

Examining the net MOC import or export of freshwater at the Atlantic southern boundary as a function of default flux correction (Figure 10) indicates that over much of the bi-stable region, in each resolution, changes in MOC freshwater import/export act to maintain bi-stability. In 36x36x8 for $f = 0.6$ – 0.9 , THC switch-off results in an increase in freshwater import, in 72x72x16 this occurs for $f = 0.65$ – 0.85 , and in 64x32x8 for $f = 0.15$ – 0.5 . This can go part way toward explaining the bi-stability of each resolution and the wider region of bi-stability in the 64x32x8 model (Figure 7). However, it cannot explain what is going on in the transition from bi-stability toward a mono-stable ‘on’ state, which also differs somewhat between resolutions. In 36x36x8 at e.g. $f=0.85$, although the ‘off’-initialized state has recovered to a strength of circa 17 Sv, greater freshwater import is still acting to maintain bi-stability. At $f = 0.95$ and above, this situation is reversed, yet there is still an indication of bi-stability in the overturning strength (Figure 7). In 72x72x16 at $f = 0.9$ and above the ‘off’-initialized state actually exports freshwater whilst the ‘on’ state imports it and yet the ‘off’-initialized state has a weaker overturning strength. In 64x32x8, at $f = 0.6$ and above, the recovering ‘off’-initialized state imports less freshwater than the ‘on’-initialized state but still has a weaker overturning strength.

5.3.2 Atmosphere feedback

To explain what is going on, especially in the transition from bi-stability to a mono-stable ‘on’ state, we must also consider atmospheric feedbacks. The change in net freshwater input (precipitation minus evaporation plus runoff) to the Atlantic from the atmosphere when the THC switches off shows a broadly similar pattern in our illustrative runs for all three resolutions (Figure 11). There is a general wetting up the storm track in the North Atlantic, a drying in the tropical North Atlantic, a wetting in the tropical South Atlantic, and a drying in the Southwest Atlantic. The changes in the tropics are indicative of a Southward shift in the Inter-Tropical Convergence Zone (ITCZ), which is typical of THC shut-down experiments in coupled models (Vellinga and Wood, 2002, Yin *et al.*, 2006). The amplitude of change is more pronounced in the higher resolutions than in 36x36x8.

Considering the net freshwater flux to the Atlantic as a function of latitude (Figure 12), the changes due to THC switch-off are small compared to the absolute flux in 36x36x8 and 72x72x16, but more pronounced in 64x32x8. A

wetting of the extra-tropical N. Atlantic is apparent in all resolutions, and a Southward shift of the ITCZ is most apparent in 64x32x8.

When integrating the net freshwater flux over the entire Atlantic basin, the illustrative cases all show an increase due to THC switch-off. However, there is considerable variation in this quantity across different runs as a function of Atlantic-Pacific flux correction (Figure 13). If there is multi-decadal variability in the total flux (not examined) then part of this variability may be due to averaging over 10 years. The atmospheric feedback is most pronounced in 64x32x8, with a substantial increase in freshwater input in going from the on-initialized to off-initialized states for $f = 0.45$ – 0.8 . This corresponds with the righthand section of the bi-stable region and the transition toward a mono-stable ‘on’ state (Figure 8) where the freshwater input will tend to inhibit recovery of the THC ‘off’ state. In 72x72x16 the atmosphere feedback appears substantial for $f = 0.75$ – 0.9 , which again corresponds with the righthand section of the bi-stable region and the transition toward a mono-stable ‘on’ state, although similar variability at $f = 0.4$ is well within the bi-stable off regime. In 36x36x8 it is hard to establish any qualitative relationship.

5.4 Freshwater hosing experiments

The maximum Atlantic MOC has been estimated at 18 Sv (Talley *et al.*, 2003), in reasonable agreement with estimates of North Atlantic Deep Water (NADW) formation of 16 ± 2 Sv (at 48°N) (Ganachaud, 2003) or 17.2 Sv (Smethie and Fine, 2001), and of the transport across 24°N as 18.5 ± 2 Sv (Ganachaud, 2003). In all three of our model variants, such values are clearly within the region of bi-stability (Figure 7) and rather close to the minimum sustainable value for the ‘on’ state of maximum Atlantic MOC. Hence they are close to the ‘cliff edge’ in the 2 parameter experiment with 36x36x8 (Figure 6). A value of ≈ 18 Sv corresponds to $f=0.65$ of the default flux correction in 72x72x16, $f=0.6$ – 0.65 in 36x36x8, and $f=0.2$ in 64x32x8. The real maximum Atlantic MOC may be somewhat larger than 18 Sv with estimates up to 26 Sv at 59°N (Talley *et al.*, 2003) and 23 ± 3 Sv at 30°S (Ganachaud, 2003). A value of ≈ 23 Sv corresponds to $f=0.7$ – 0.75 in 36x36x8, $f=0.7$ in 72x72x16, and $f=0.25$ in 64x32x8.

The longitude-latitude 64x32x8 grid needs the least flux correction to get a THC ‘on’ state and has the widest region of bi-stability. Hence in order to examine the effect of freshwater hosing, we selected 8 ensemble members that reside on the upper branch of the bistable regime (restarted from the THC ‘on’) for 0.1 Sv and 1.0 Sv freshwater hosing experiments as in Stouffer *et al.* (2006). These range over $f=0.15$ – 0.50 of the default flux correction in steps of 0.05 and have a maximum Atlantic MOC ranging over ≈ 15 – 33 Sv. After 1000 years all have a stable Atlantic MOC with the possible exception of $f=0.15$, which is weakening slightly. Applying 0.1 Sv freshwater hosing for 100 yrs, the Atlantic MOC collapses in 2 ensemble members ($f=0.15$, 0.20) and weakens modestly in the others. When the forcing is removed, the collapsed MOC runs do not recover, indicating a switch to the other bi-stable state (although the run with $f=0.20$ shows brief resumptions of large-scale convection, e.g. around year 160). Applying 1.0 Sv freshwater hosing over 100 years causes all ensemble members to collapse. On removing the forcing, none have recovered after a further 140 yrs, indicating a switch to the other bi-stable state in all cases.

6 Discussion

We have described a search for bi-stability of the thermohaline circulation with the 3-D ocean–atmosphere–sea-ice core of the new GENIE-2 Earth system model. The study should be seen as a conceptual one because the fully coupled model has yet to be tuned. Such tuning would improve the fit of our model results to observations and quantitatively alter our predictions but would be unlikely to qualitatively alter the presence of bi-stability in the various configurations. The simulations presented here are also subject to a small scaling error in the ocean equation of state, which results in an under-prediction of ocean velocities by $\approx 10\%$. Despite these limitations, the results provide one of the first systematic, qualitative demonstrations of bi-stability of the thermohaline circulation in a 3-D ocean–atmosphere–sea-ice model. Although coupling to a fully dynamical atmosphere model clearly increases variability in the THC, it does not remove bi-stability or obviously blur the boundaries of the hysteresis loop, in contrast to recent suggestions (Schlesinger *et al.*, 2006, Yin *et al.*, 2006).

The imposed default flux correction in GENIE-2 totals 0.79 Sv from Atlantic to Pacific and is about 2.5 times the 0.32 Sv used by default in GENIE-1 and C-GOLDSTEIN (Marsh *et al.*, 2004). Figure 7 shows in grey a typical hysteresis loop for GENIE-1 obtained by varying the scaling of the Atlantic to Pacific freshwater flux adjustment, which gives very similar results to the established method used in Rahmstorf *et al.* (2005). With the baseline ocean resolution 36x36x8, the IGCM atmosphere demands more flux correction than the EMBM to get a THC ‘on’ state (Figure 7). This may be partly explained by the use of prescribed, NCEP-derived wind fields in the EMBM. Furthermore, the IGCM resolves vertical structure, simulates cloud cover and associated radiative properties, and is coupled in a more sophisticated way with the land surface. These differences may substantially alter atmospheric humidity, and hence large-scale moisture transports. The width of the region of THC bi-stability appears to be somewhat larger with the IGCM atmosphere (≈ 0.25 Sv) than with the EMBM (≈ 0.15 Sv) (Marsh *et al.*, 2004). With the 64x32x8 longitude-latitude grid and the IGCM, it is wider still (up to ≈ 0.5 Sv). Thus, including a 3-D dynamical atmosphere can actually broaden the region of THC bi-stability.

We conjectured in the introduction that feedbacks from the atmosphere (Latif *et al.*, 2000, Schmittner *et al.*, 2000) and the ocean (de Vries and Weber, 2005) may stabilize both the present ‘on’ state of the THC and the ‘reversed’ state. Our results show that both ocean and atmosphere feedbacks do indeed contribute to extensive bi-stability of the THC. From the bifurcation between mono-stable ‘off’ state and bi-stability through much of the region of bi-stability, THC switch-off leads to increased import of fresh water (export of salt) at the southern boundary of the Atlantic, which tends to maintain the ‘off’ state (Figure 10). However, the MOC imports freshwater (exports salt) in all experiments with the 64x32x8 model and in most experiments with the 36x36x8 model. Only in experiments with the 72x72x16 model does the MOC sometimes export fresh water at the southern boundary, favouring the THC ‘on’ state, while importing fresh water in the corresponding ‘off’ state. This behaviour is broadly consistent with findings of (de Vries and Weber, 2005), although the sign of MOC freshwater transport is clearly not the single determinant of THC bi-stability in GENIE-2.

This positive ocean feedback is also evident in all resolutions where the

bi-stable state that is initialized ‘off’ starts to recover. In those cases, the recovering ‘off’ state imports less fresh water (and even exports fresh water in experiments with the 72x72x16 resolution) as the resuming THC tends to draw salt into the Atlantic. Remarkably, less salt is exported than in those states that are initialized ‘on’, yet the MOC strength remains weaker in states initialised ‘off’. To explain this, atmosphere feedbacks must also be considered. Towards the point where the bi-stable ‘off’ state starts to break down, and through the region where it is recovering, positive feedback from the atmosphere tends to maintain bi-stability by increasing Atlantic freshwater input to the state that is initialized ‘off’ relative to the state that is initialized ‘on’ (Figure 13). The ocean and atmosphere mechanisms can be seen as counteracting one another in this regime, with the stronger atmospheric feedback extending the region of bi-stability towards higher values of Atlantic-Pacific freshwater flux correction. Even when the THC state that is initialized ‘off’ does switch on, it remains weaker than the state that is initialized ‘on’. Thus weaker and stronger THC ‘on’ states can be distinguished under the same boundary conditions and different initializations of the model (Figure 7).

The differences in the position, width and sharpness of the bi-stable regimes as a function of Atlantic-to-Pacific freshwater flux correction can be interpreted in terms of the varying strengths of feedbacks in the different model resolutions. The smaller flux correction required in the 64x32x8 model is related to the greater northward transport of high salinity water in the North Atlantic, by the MOC itself (Figure 9). The 64x32x8 model has the widest region of bi-stability because it has the widest region over which the ocean provides positive feedback, and where the ocean switches to negative feedback this is counteracted by the strongest positive feedback from the atmosphere. The 36x36x8 model has a slightly wider region of bi-stability than the 72x72x16 model, corresponding to a wider range over which the ocean provides positive feedback (where the atmosphere provides positive feedback it is weaker and more overlapping with the regime of ocean positive feedback in these resolutions). The main effect of the increase in resolution from 36x36x8 to 72x72x16 is to sharpen up the boundaries of the bi-stable regime, due to a stronger self-sustaining ocean feedback for the ‘on’ and ‘off’ states (Figure 10).

It is harder to pin down the deeper reasons why the resolutions differ in these ways. We suggest that the absence of interpolation in the 64x32x8 model is the main factor reducing the required flux correction. In the other two resolutions, interpolation leads to an additional flux correction applied in the tropics which may be shifting the model away from the THC ‘on’ state. Increasing the resolution from 36x36x8 to 72x72x16 within the same grid should improve the resolution of zonal pressure gradients in the ocean, leading to better THC structure, and improved salt transport. Surprisingly, this does not alter the required flux correction, but it may contribute to the larger changes in salt import/export with changes in the THC state in 72x72x16.

To put our results in the context of other studies, it is important to recognize that in GENIE-2, as in all other models that have been systematically tested (Rahmstorf *et al.*, 2005), there is a region of THC bi-stability, a region with a mono-stable ‘on’ state, and a region with a mono-stable ‘off’ state. The main difference between other models is in the location of the present climate with respect to the region of bi-stability – i.e. whether the present climate resides on the mono-stable ‘on’ branch, or on the ‘on’ branch in the region of bi-stability.

If the former, then a collapse of the THC caused by e.g. a temporary freshwater hosing perturbation, will be reversible, if the latter, then it will be irreversible. This observation seems to have escaped the notice of some authors who, on the basis of observing a reversible THC collapse in limited AOGCM runs, suggest that there is no bi-stability in their model. In fact, all that can be inferred is that the initial state of their model is not in a region of bi-stability. Given that our coupled model is un-tuned we do not consider that any of our model versions preferentially represent the present climate. Consequently, we are not in a position to say whether the real THC is in a mono-stable or bi-stable regime.

Even in a region of bi-stability, the size of perturbation required to trigger THC collapse depends on the position in the bi-stable regime, as is apparent in our hosing experiments (Figure 14). Versions of the 64x32x8 model that are close to the present estimate of maximum Atlantic MOC are vulnerable to 0.1 Sv of extra freshwater addition over just 100 years. This rate is equivalent to the freshwater flux expected from the Greenland ice sheet if it melts over a timescale of ≈ 1000 yr. The decrease of 10–20 Sv seen in the ensemble members that collapse exceeds that in any of the models inter-compared by Stouffer *et al.* (2006). However, this result should be treated with caution, as in these ensemble members the Pacific circulation is already reversed and the Atlantic MOC fails to reach the Southern Ocean. Our other ensemble members show a similar or lesser weakening than the ≈ 5 Sv in the ensemble mean of Stouffer *et al.* (2006). Under an extreme 1.0 Sv freshwater hosing, the collapse of the Atlantic THC in all our ensemble members agrees with all the models inter-compared by Stouffer *et al.* (2006). However, whereas a number of the other models recover after the perturbation is removed (indicating that they are in a mono-stable ‘on’ regime), none of our ensemble members do, because they have switched into a different stable state.

The response of the atmosphere to THC collapse in GENIE-2 exhibits similarities and important differences to other AOGCMs, noting that the atmospheric response in GENIE-2 is clearly a function of where the model is in parameter space, so the comparison may not be a direct one. A southward shift of the ITCZ is apparent in GENIE-2 as in other AOGCMs (Vellinga and Wood, 2002, Yin *et al.*, 2006). Yin *et al.* (2006) find a net drying of the tropical North Atlantic, which they argue creates a salinity anomaly near the Gulf of Mexico that propagates up the North Atlantic and aids THC recovery. This is somewhat similar to the anomaly observed in El Nino events (Schmittner *et al.*, 2000) that acts to stabilize the THC in global warming simulations with a different model (Latif *et al.*, 2000). Vellinga and Wood (2002) (their figure 4c) show some similar drying as well as wetting of the North Atlantic, but the changes associated with the ITCZ shifting are dominant. However, given their THC recovery, these models appear to be in the mono-stable ‘on’ regime. In contrast, the illustrative runs with GENIE-2 are in a bi-stable regime, and there is no clear Gulf of Mexico feature. Instead there is a net freshwater input to the North Atlantic concentrated up the storm track, which is strongest in the 64x32x8 model. Clearly the response of the atmosphere may be important in determining whether the THC recovers or remains collapsed, but this needs to be systematically examined with reference to the model stability regime. Furthermore, our results suggest that atmospheric feedback only dominates over ocean feedbacks in a restricted region of parameter space.

Acknowledgement The UK Natural Environment Research Council funded GENIE (NER/T/S/2002/00217) and funds the ongoing GENIE_{fy} (NE/C515904) project through its e-Science programme. The development of alternative ocean grids was funded by the NERC Rapid Climate Change project “The role of the cryosphere on modulating the THC of the North Atlantic” (Ref. NER/T/S/2002/00462). The authors acknowledge the use of the UK National Grid Service in carrying out this work.

7 Appendix: Grid computing performance

To obtain the results presented here we undertook a total of 319 GENIE-2 simulations totalling 407,000 model years and comprising 3,736 compute tasks, taking 46,992 CPU hours (some timings estimated). This was achieved in around 3 months, including time spent considering the initial results and designing the restart experiments. The throughput is equivalent to having about 20 model runs going continuously throughout the time, highlighting the power of the Grid to enable ensemble studies with Earth system models.

Typically, a 1000 year simulation of GENIE-2 requires ≈ 4 days of continuous compute time on a desktop machine (P4 3.0 GHz). To perform large ensembles of such runs presents a challenge in the traditional academic environment where compute clusters are often configured to limit single CPU jobs to less than 4 days and an individual user is unlikely to be permitted exclusive use of the resource. Furthermore, the effort involved in configuring, managing and executing a study of this nature on a single cluster can present a significant overhead. The GENIE data management solution provided the means for the project to harness the shared pool of compute resource available to individual project members and undertake the study with relative ease.

To contribute computational resource to a model study defined in the database a user simply requires a single piece of information; the unique identifier of the experiment entity in the database. The identifier is typically distributed by the experiment coordinator to interested parties although a query on the database can also be performed to obtain this information. Armed with the identifier, the user instantiates a time-limited proxy certificate to authenticate their actions on the Grid. They then create a data structure describing the computational resource(s) that they wish to submit work units to. This descriptive metadata provides information including hostname, resource type (Globus, Condor) and the remote brokering system (Portable Batch System, Sun Grid Engine, Condor, etc.) which the system can then use to manage job submissions. With this information defined in the workspace the user simply invokes a single function call in order to contribute to the study. The scripted workflows of the GENIE Toolbox then perform queries on the database to find available work units and submit jobs to the specified resource. Post-processing scripts are also invoked on completed work units in order to archive output data to the database.

In total 5 client installations were used to progress the simulations to completion. Each client queried the shared database for work units and submitted those compute tasks to resource available to the user in each case (see Figure 15). Each client was configured by the user to exploit both local and national resource available to them using their X.509 e-Science certificate. The resource

used included 5 nodes of the UK National Grid Service (NGS), 3 institutional High Performance Computing (HPC) clusters and 1 large university Condor pool. The database system enabled users to specify how much work was submitted, for example, short 5–10 hour jobs were typically submitted on Condor as the chances of an interruption increase with time (we cannot exploit Condor’s checkpointing facility because our pool consists of only Windows nodes). For the NGS and HPC clusters we were able to submit jobs that run close to the maximum allowable wall time of the scheduling system. Throughput is in principle only limited by the amount of available resource.

The daily breakdown of resource usage is shown in Figure 16. Periods of relative inactivity reflect development work on the system and transitions between experiments where the science team were analysing results from the first phase of experiments and configuring the second phase of the study. The computational throughput was not limited by the system and the rate of progression was determined by the amount of available work. The 5 clients functioning in tandem reached a sustained peak in excess of 90 model runs per day and each made a significant contribution to the studies. The breakdown of the number of jobs completed by each client installation (Figure 17a) reflects the amount of resource available to the user on each of these clients. The client responsible for 50% of the work was the only submission node on the large Condor pool and was also used to submit jobs to the National Grid Service. The distribution of jobs across the computational resources (Figure 17b) illustrates that the studies made even use of the available platforms. The jobs were well spread across the National Grid Service, the large Condor pool and the 3 dedicated clusters. The usage of different clusters broadly reflects the relative numbers of jobs a user could reasonably hope to run concurrently on each of these systems.

References

- Abramson, D., Buyya, R., and Giddy, J. (2002). A computational economy for Grid computing and its implementation in the Nimrod-G resource broker. *Future Generation Computer Systems*, **18**(8), 1061–1074.
- Allcock, B., Bester, J., Bresnahan, J., Chervenak, A. L., Foster, I., Kesselman, C., Meder, S., Nefedova, V., Quesnal, D., and Tuecke, S. (2002). Data management and transfer in high performance computational Grid environments. *Parallel Computing Journal*, **28**(5), 749–771.
- Betts, A. (1986). A new convective adjustment scheme. 1. observational and theoretical basis. *Q J R Meteorol Soc*, **112**, 667–691.
- Claussen, M., Mysak, L., Weaver, A., Crucifix, M., Fichet, T., Loutre, M., Weber, S., Alcamo, J., Alexeev, V., Berger, A., Calov, R., Ganopolski, A., Goosse, H., Lohmann, G., Lunkeit, F., Mokhov, I., Petoukhov, V., Stone, P., and Wang, Z. (2002). Earth system models of intermediate complexity: Closing the gap in the spectrum of climate system models. *Climate Dynamics*, **18**, 579–586.
- Czajkowski, K., Foster, I., Karonis, N., Kesselman, C., Martin, S., Smith,

- W., and Tuecke, S. (1998). A resource management architecture for meta-computing systems. In *Proc. IPPS/SPDP '98 Workshop on Job Scheduling Strategies for Parallel Processing*, pages 62–82.
- de Forster, P. M., Blackburn, M., Glover, R., and Shine, K. P. (2000). An examination of climate sensitivity for idealised climate change experiments in an intermediate general circulation model. *Climate Dynamics*, **16**, 833–849.
- de Vries, P. and Weber, S. L. (2005). The atlantic freshwater budget as a diagnostic for the existence of a stable shut-down of the meridional overturning circulation. *Geophysical Research Letters*, **32**, L09606.
- Edwards, N. R. and Marsh, R. (2005). Uncertainties due to transport-parameter sensitivity in an efficient 3-D ocean-climate model. *Climate Dynamics*, **24**, 415–433.
- Eres, M. H., Pound, G. E., Jiao, Z., Wason, J. L., Xu, F., Keane, A. J., and Cox, S. J. (2005). Implementation and utilisation of a Grid-enabled problem solving environment in Matlab. *Future Generation Computer Systems*, **21**(6), 920–929.
- Fanning, A. G. and Weaver, A. J. (1996). An atmospheric energy-moisture model: Climatology, interpentadal climate change and coupling to an ocean general circulation model. *Journal of Geophysical Research*, **101**, 15111–15128.
- Foster, I. and Kesselman, C. (1999). Globus: A toolkit-based Grid architecture. In I. Foster and C. Kesselman, editors, *The Grid: Blueprint for a New Computing Infrastructure*, pages 259–278. Morgan Kaufman Publishers, Inc., San Francisco, USA.
- Foster, I., Kesselman, C., and Tuecke, S. (2001). The anatomy of the Grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications*, **15**(3), 200–222.
- Fraedrich, K., Jansen, H., Kirk, E., Luksch, U., and Lunkeit, F. (2005a). The planet simulator: Towards a user friendly model. *Meteorologische Zeitschrift*, **14**(3), 299–304.
- Fraedrich, K., Kirk, E., Luksch, U., and Lunkeit, F. (2005b). The portable university model of the atmosphere (puma): Storm track dynamics and low-frequency variability. *Meteorologische Zeitschrift*, **14**(6), 735–745.
- Ganachaud, A. (2003). Large-scale mass transports, water mass formation, and diffusivities estimated from world ocean circulation experiment (WOCE) hydrographic data. *Journal of Geophysical Research (Oceans)*, **108**(C7), 3213.
- Goosse, H. and Fichefet, T. (1999). Importance of ice-ocean interactions for the global ocean circulation: a model study. *Journal of Geophysical Research*

- (*Oceans*), **104**(C10), 23337–23355.
- Griffies, S. M. (1998). The Gent-McWilliams skew flux. *Journal of Physical Oceanography*, **28**, 831–841.
- Hargreaves, J. C., Annan, J. D., Edwards, N. R., and Marsh, R. (2004). An efficient climate forecasting method using an intermediate complexity Earth System Model and the ensemble Kalman filter. *Climate Dynamics*, **23**, 745–760.
- Hey, T. and Trefethen, A. E. (2002). The UK e-Science core programme and the Grid. *Future Generation Computing Systems*, **18**(8), 1017–1031.
- Hoskins, B. J. and Simmons, A. J. (1975). A multi-layer spectral model and the semi-implicit method. *Quarterly Journal of the Royal Meteorological Society*, **101**, 637–655.
- Jayne, S. R. and Marotzke, J. (2001). The dynamics of ocean heat transport variability. *Reviews of Geophysics*, **39**, 385–411.
- Keane, A. J. (2003). *The OPTIONS design exploration system: reference manual and user guide*. <http://www.soton.ac.uk/~ajk/options.ps>.
- Latif, M., Roeckner, E., Mikolajewicz, U., and Voss, R. (2000). Tropical stabilization of the thermohaline circulation in a greenhouse warming simulation. *Journal of Climate*, **13**, 1809–1813.
- Lenton, T. M., Williamson, M. S., Edwards, N. R., Marsh, R., Ridgwell, A. J., Price, A. R., Shepherd, J. G., Cox, S. J., and The GENIE team (2006). Millennial timescale carbon cycle and climate change in an efficient Earth system model. *Climate Dynamics*, **26**, 687–711.
- Louis, J.-F. (1986). Forecast model research manual 3: physical parameterisation. *European Centre for Medium Range Weather Forecasts*.
- Marsh, R., Yool, A., Lenton, T. M., Gulamali, M. Y., Edwards, N. R., Shepherd, J. G., Krznaric, M., Newhouse, S., and Cox, S. J. (2004). Bistability of the thermohaline circulation identified through comprehensive 2-parameter sweeps of an efficient climate model. *Climate Dynamics*, **23**, 761–777.
- Morcrette, J. (1990). Impact of changes to the radiative transfer parameterisations plus cloud optical properties in the ECMWF model. *Monthly Weather Review*, **118**, 847–873.
- Murphy, J. M., Sexton, D. M. H., Barnett, D. N., Jones, G. S., Webb, M. J., Collins, M., and Stainforth, D. A. (2004). Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature*, **430**, 768–772.
- Opsteegh, J. D., Haarsma, R. J., Selten, F. M., and Kattenberg, A. (1998).

- ECBILT a dynamic alternative to mixed boundary conditions in ocean models. *Tellus A*, **50**, 348–367.
- Petoukhov, V., Ganopolski, A., Brovkin, V., Claussen, M., Eliseev, A., Kubatzki, C., and Rahmstorf, S. (2000). CLIMBER-2: A climate system model of intermediate complexity. Part I: Model description and performance for present climate. *Climate Dynamics*, **16**, 1–17.
- Pope, V., Gallani, M., Rowntree, P., and Stratton, R. (2000). The impact of new physical paramterisations in the Hadley Centre climate model: HadAM3. *Climate Dynamics*, **16**, 123–146.
- Price, A. R., Xue, G., Yool, A., Lunt, D. J., Valdes, P. J., Lenton, T. M., Wason, J. L., Pound, G. E., Cox, S. J., and the GENIE team (2006). Optimisation of integrated Earth System Model components using Grid-enabled data management and computation. *Concurrency and Computation: Practice and Experience*, page in press.
- Rahmstorf, S., Crucifix, M., Ganopolski, A., Goosse, H., Kamenkovich, I., Knutti, R., Lohmann, G., Marsh, R., Mysak, L. A., Wang, Z., and Weaver, A. J. (2005). Thermohaline circulation hysteresis: A model intercomparison. *Geophysical Research Letters*, **32**, L23605.
- Roche, D. M., Dokken, T. M., Goosse, H., Renssen, H., and Weber, S. L. (2006). Climate of the last glacial maximum: sensitivity studies and model-data comparison with the loveclim coupled model. *Climate of the Past Discussions*, **2**, 1105–1153.
- Schlesinger, M. E., Yin, J., Yohe, G. W., Andronova, N. G., Malyshev, S., and Li, B. (2006). Assessing the risk of a collapse of the Atlantic thermohaline circulation. In H. J. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, and G. Yohe, editors, *Avoiding Dangerous Climate Change*, pages 37–47. Cambridge University Press, Cambridge.
- Schmittner, A., Appenzeller, C., and Stocker, T. F. (2000). Enhanced Atlantic freshwater export during El Nino. *Geophysical Research Letters*, **27**, 1163–1166.
- Slingo, J. (1987). The development and verification of a cloud prediction scheme for the ECMWF model. *Q J R Meteorolo Soc*, **113**, 899–927.
- Smethie, W. M. and Fine, R. A. (2001). Rates of North Atlantic Deep Water formation calculated from chlorofluorocarbon inventories. *Deep Sea Research I*, **48**, 189–215.
- Stommel, H. (1961). Thermohaline convection with two stable regimes of flow. *Tellus*, **13**, 224–230.
- Stouffer, R. J. and Manabe, S. (2003). Equilibrium response of thermohaline circulation to large changes in atmospheric CO₂ concentration. *Climate*

Dynamics, **20**, 759–773.

- Stouffer, R. J., Yin, J., Gregory, J. M., Dixon, K. W., Spelman, M. J., Hurlin, W., Weaver, A. J., Eby, M., Flato, G. M., Hasumi, H., Hu, A., Jungclaus, J. H., Kamenkovich, I. V., Levermann, A., Montoya, M., Murakami, S., Nawrath, S., Oka, A., Peltier, W. R., Robitaille, D. Y., Sokolov, A., Vettoretti, G., and Weber, S. L. (2006). Investigating the causes of the response of the thermohaline circulation to past and future climate changes. *Journal of Climate*, **19**, 1365–1387.
- Talley, L. D., Reid, J. L., and Robbins, P. E. (2003). Data-based meridional overturning streamfunctions for the global ocean. *Journal of Climate*, **16**, 3213–3226.
- Thain, D., Tannenbaum, T., and Livny, M. (2005). Distributed computing in practice: the Condor experience. *Concurrency and Computation: Practice and Experience*, **17**(2–4), 323–356.
- Thorpe, R. B., Gregory, J. M., Johns, T. C., Wood, R. A., and Mitchell, J. F. B. (2001). Mechanisms determining the Atlantic thermohaline circulation response to greenhouse gas forcing in a non-flux-adjusted coupled climate model. *Journal of Climate*, **14**, 3102–3116.
- Vellinga, M. and Wood, R. A. (2002). Global climatic impacts of a collapse of the atlantic thermohaline circulation. *Climatic Change*, **54**, 251–267.
- von Laszewski, G., Foster, I., Gawor, J., and Lane, P. (2001). A Java Commodity Grid kit. *Concurrency and Computation: Practice and Experience*, **13**(8–9), 643–662.
- Wang, H. and Birchfield, G. E. (1992). An energy-salinity balance climate model: Water vapor transport as a cause of changes in the global thermohaline circulation. *Journal of Geophysical Research*, **97**(C2), 2335–2346.
- Welch, V., Siebenlist, F., Foster, I., Bresnahan, J., Czajkowski, K., Gawor, J., Kesselman, C., Meder, S., Pearlman, L., and Tuecke, S. (2003). Security for Grid services. In *Proceedings of the 12th IEEE International Symposium on High Performance Distributed Computing*, pages 48–57, Seattle, Washington, USA. The Institute of Electrical and Electronics Engineers, Inc.
- Wright, D. G. and Stocker, T. F. (1991). A zonally averaged ocean model for the thermohaline circulation. Part 1: Model development and flow dynamics. *Journal of Physical Oceanography*, **21**, 1713–1724.
- YarKhan, A., Seymour, K., Sagi, K., Shi, Z., and Dongarra, J. (2006). Recent developments in Gridsolve. *International Journal of High Performance Computing Applications*, **20**(1), 131–141.
- Yin, J., Schlesinger, M. E., Andronova, N. G., Malyshev, S., and Li, B. (2006). Is a shutdown of the thermohaline circulation irreversible? *Journal of*

Geophysical Research (Oceans), **111**, D12104.

Table 1: Components, models, grids and resolutions used in the GENIE-2 experiments described herein (many others are available in the GENIE framework – see Supplementary Information).

Component	Model	Resolution Variant	Longitude		Latitude		Vertical	
			Cells	Res.	Cells	Res.	Levels	Topog.
Ocean	GOLDSTEIN	Baseline	36	10°	36	$\sin(\pi/36)$	8	Yes
		High	72	5°	72	$\sin(\pi/72)$	16	Yes
		IGCM res.	64	5.625°	32	5.625°	8	Yes
Atmosphere	3-D IGCM	Baseline T21	64	5.625°	32	5.625°	7	Yes
Sea-ice	Slab	Baseline	64	5.625°	32	5.625°	1	–
Land	IGCM-land	Baseline	64	5.625°	32	5.625°	–	–

Table 2: Parameter settings used in the new experiments. Ocean albedo, sea-ice albedo and all IGCM and IGCM-land parameters listed are from genetic algorithm tuning of the IGCM with fixed ocean, fixed sea-ice, and IGCM-land (genie-ia-fo-fs-il) – they are the member of the final population with the lowest rms error. Other GOLDSTEIN parameters listed are un-tuned default values.

Parameter	Label	Value	Units
IGCM			
roughness length over ocean	ia_rough_ocean	3.711×10^{-4}	m
boundary layer gustiness	ia_blvad	1.132	ms^{-1}
enhancement in the stable case	ia_enhancestable	0.3731	–
factor limit for increases in speed due to gusts	ia_enhancelim	2.011	–
boundary layer relative humidity	ia_blrh	1.198×10^2	–
increase in flux	ia_fluxfactor	1.327	–
diffusion increase of u/v at 2nd lowest level	ia_increaseu	7.332	–
diffusion increase of t at 2nd lowest level	ia_increaset	0.6678	–
diffusion increase of q at 2nd lowest level	ia_increaseq	0.8888	–
vertical diffusion coefficient for q	ia_akqv	1.128	m^2s^{-1}
convection or not?	ia_lcr	true	–
timescale for convective adjustment	ia_cbadjt	4.911	hr
minimum precipitation for which there is cloud	ia_pcloudmin	3.013	mm d^{-1}
precipitation factor	ia_pcloudfact	0.09859	–
cloud fraction when there is shallow convection	ia_cloudshallow	0.2642	–
timescale for non-precipitation convection	ia_cubmt	1.371	hr
minimum RH for cloud fraction = 0	ia_humcloudmin	0.7263	–
maximum RH for cloud fraction = 1	ia_humcloudmax	0.9085	–
power for variation between min and max	ia_powercloud	1.002	–
cloud supersaturation for liquid water path	ia_supersatconv	0.01273	–
as above but for large-scale clouds	ia_supersatlarg	0.03299	–
cloud factor	ia_cloudconvnorm	0.4409	–
cloud factor at the base of the cloud	ia_cloudconvbase	0.8018	–
emissivity as f(liquid water path) constant a	ia_emissconsta	1.152	–
emissivity as f(liquid water path) constant b	ia_emissconstb	1.278×10^2	–
droplet radius	ia_zradef	3.092	microns
IGCM-land			
global albedo fix	ia_albedo_fix	-0.03963	–
height of snow for intermediate snow/veg albedo	ia_shsstar	0.1294	m
GOLDSTEIN			
isopycnal diffusivity	κ_h , go_14	2000	m^2s^{-1}
diapycnal diffusivity	κ_v , go_15	1×10^{-5}	m^2s^{-1}
friction	λ , go_16	2.5	days^{-1}
wind-scale	W , go_13	2	–
ocean albedo	go_albedo	0.07422	–
Slab sea-ice			
sea-ice albedo	ss_albedo_seaice	0.4652	–

Table 3: Ensemble studies of GENIE-2 performed on the Grid. The key parameter varied is the scaling factor of the default Atlantic-to-Pacific freshwater flux correction (f).

Experiment	Run duration (yr)	Resolution [binary]	Parameter	Range	Members	Restart state
Spin-ups	1000	36x36x8	f	0-2	21	-
	1000	64x32x8	f	0-2	21	-
	1000	72x72x16	f	0-2	21	-
Restarts	1000	36x36x8	f	0-1	21	$f=0$
					21	$f=1$
	1000	64x32x8	f	0-1	21	$f=0$
					21	$f=1$
	1000	72x72x16	f	0-1	21	$f=0$
					21	$f=1$
1000	72x72x16	f	0-1	21	$f=0$	
				21	$f=1$	
2-parameter	2000	36x36x8	f	0-2	11*11	-
			$\log(\text{ia_enhancestable})$	-1-1	= 121	-




Figure 1: The spectrum of Earth system models, after Claussen *et al.* (2002), defined in terms of number of grid cells (G), cumulative dimension (D), and number of interacting components (I). The positions of some well known and generic models are indicated, including a typical atmosphere-ocean-land (AOL) GCM. The GENIE framework spans a volume of model space in terms of resolution (G), dimensionality (D) and comprehensiveness (I), indicated by the black bars along the axes and the corresponding box.




Figure 2: Major components of an Earth system model. The GENIE framework currently offers at least one option for each of these components, except atmospheric chemistry, which is currently just an atmospheric tracer module without chemical reactions.

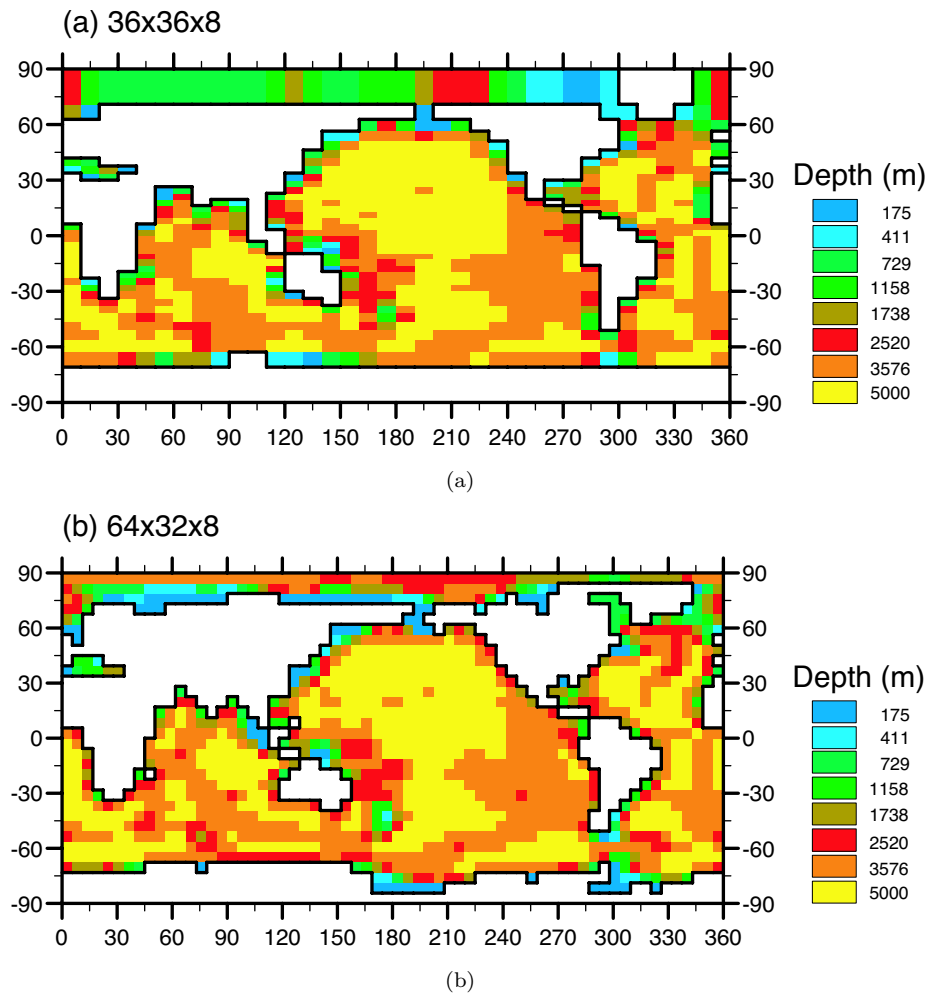


Figure 3: Alternative surface grids and corresponding bathymetry in the GENIE framework: (a) 36x36x8 longitude-sine(latitude) i.e. equal area, (b) 64x32x8 longitude-latitude. Note that the 72x72x16 model used herein shares the 36x36x8 surface grid and bathymetry, in order to isolate the effects of simply increasing the number of grid cells.

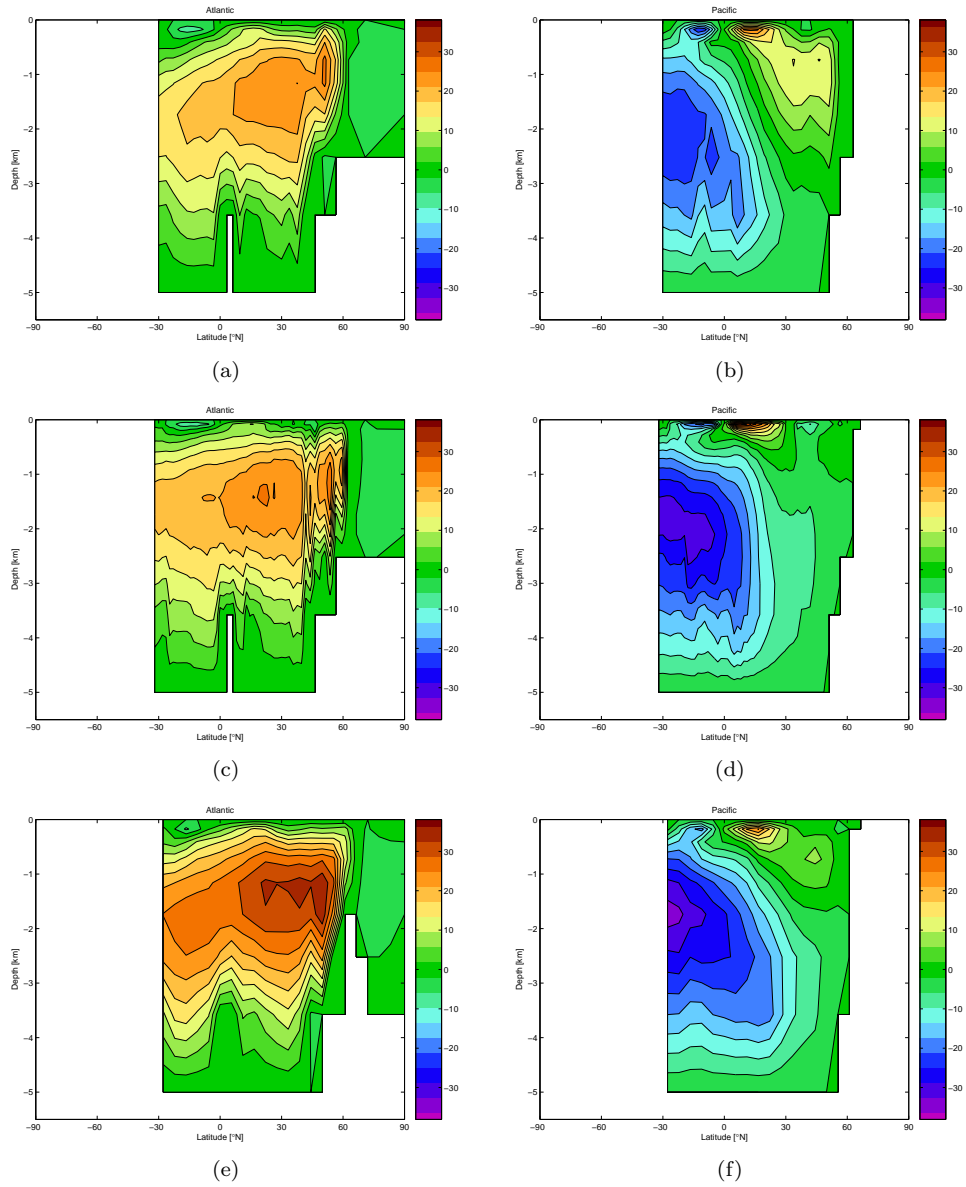


Figure 4: End-of-run annual average (a, c, e) Atlantic MOC and (b, d, f) Pacific MOC (both in Sv) after 2000 years at the default flux correction, for (a, b) 36x36x8 (c, d) 72x72x16 (Linux binary) (e, f) 64x32x8.




Figure 5: Initial spin-up ensembles of maximum Atlantic Meridional Overturning Circulation (MOC) as a function of Atlantic to Pacific freshwater flux correction (expressed as multiple of default 0.79 Sv) in different ocean resolutions of the 3-D Earth system model GENIE-2: (solid line) “baseline” 36x36x8, (dotted line) higher resolution 72x72x16, (dashed line) IGCM-grid matched 64x32x8. The IGCM atmosphere resolution (T21) is the same in all cases. Points are averages over the last 50 years of 1000 year runs because in these initial ensembles output was restricted to 50 of each 100 years.




Figure 6: Maximum Atlantic Meridional Overturning Circulation (MOC) strength (Sv) in the 36x36x8 ocean resolution version of GENIE-2, as a function of the IGCM parameter `ia_enhancestable` (which scales the latent heat and freshwater flux to the atmosphere) and Atlantic to Pacific freshwater flux correction (expressed as multiple of default 0.79 Sv). Results are after 2000 years of spin-up. 88 out of 121 runs completed, the white areas indicate runs that failed due to instability. Contour interval is 2.5 Sv.

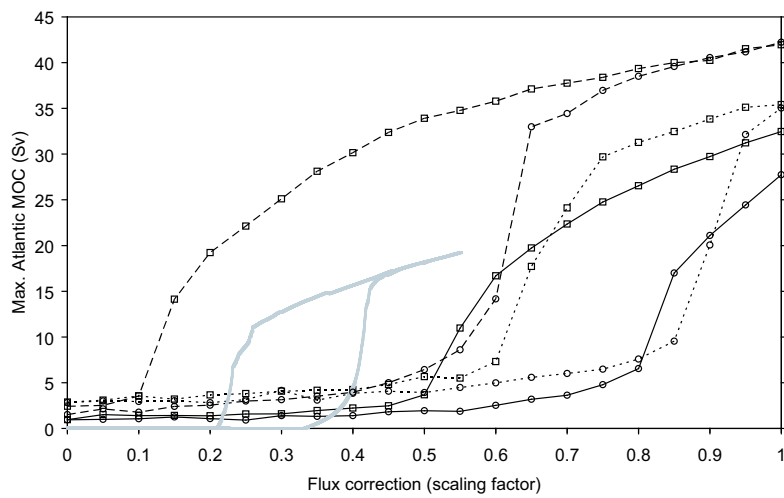


Figure 7: Bi-stability of the maximum Atlantic Meridional Overturning Circulation (MOC) as a function of Atlantic to Pacific freshwater flux correction (expressed as multiple of default 0.79 Sv) in different ocean resolutions of the 3-D Earth system model GENIE-2: (solid line) “baseline” 36x36x8, (dotted line) higher resolution 72x72x16, (dashed line) IGCM-grid matched 64x32x8. The IGCM atmosphere resolution (T21) is the same in all cases. For each resolution, ensembles were restarted from spin-ups with the default flux correction (upper branch, squares) or no flux correction (lower branch, circles). Points are averages of the last 100 years of a 1000 year run. Also shown in grey are the results of a hysteresis experiment with GENIE-1, which uses the 2-D EMBM atmosphere instead of the 3-D IGCM and has ocean resolution 36x36x8.

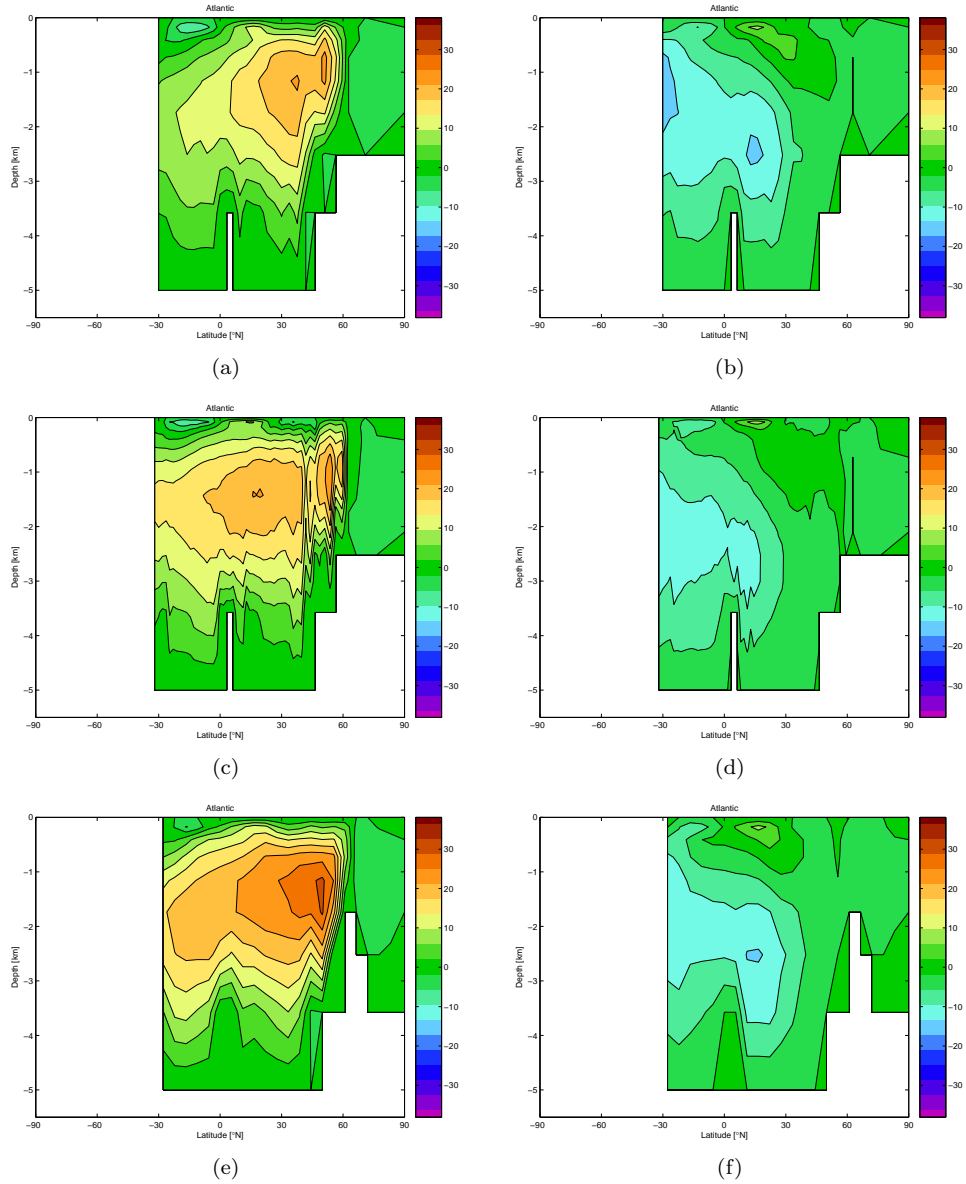





Figure 8: Bi-stability of the annual average Atlantic MOC for different ocean grids and resolutions, revealed by restarting from (a, c, e) THC on, or (b, d, f) THC reversed, for (a, b) 36x36x8 at $f = 0.8$ of default flux correction (c, d) 72x72x16 (Linux binary) at $f = 0.85$ (e, f) 64x32x8 at $f = 0.55$.



(a)




(b)




(c)


Figure 9: Meridional profiles of MOC freshwater transport versus latitude (positive northwards) for THC on and THC off restarts, and difference (off - on) (a) 36x36x8 at $f = 0.8$ of default flux correction (b) 72x72x16 at $f = 0.85$ (c) 64x32x8 at $f = 0.55$. Individual annual averages are shown for a decade at the end of the restarted simulations.



(a)



(b)



(c)

Figure 10: Freshwater transport at the Southern boundary of the Atlantic (32°S) as a function of default flux correction, for restarts from THC on (squares) and THC off (circles) states, in different resolutions: (a) $36 \times 36 \times 8$ (b) $72 \times 72 \times 16$ (c) $64 \times 32 \times 8$. Calculated as a decadal average at the end of the restarted simulations.

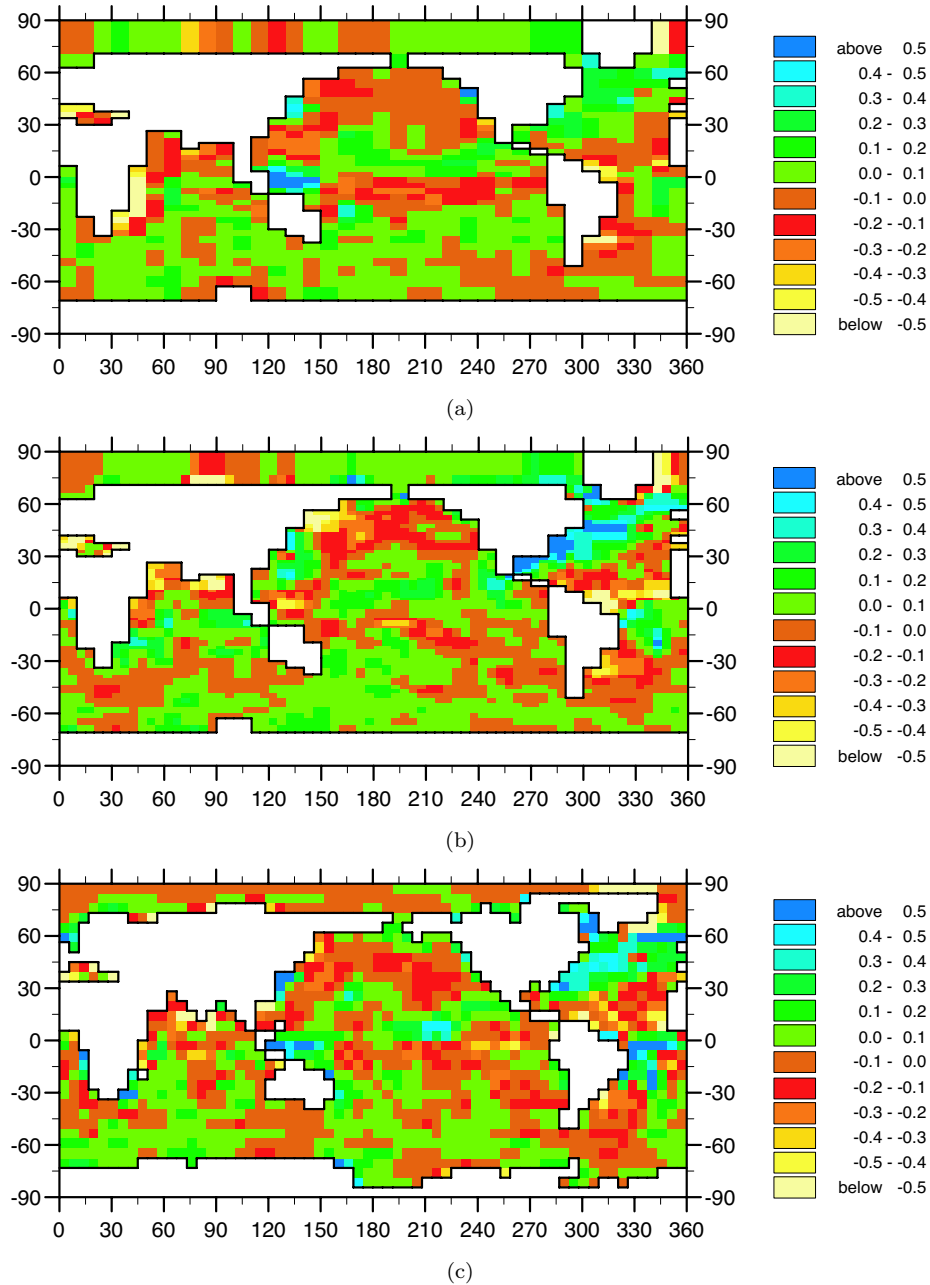





Figure 11: Maps of net surface freshwater flux difference due to THC switch off (off - on) in m/yr: (a) 36x36x8 at $f = 0.8$ of default flux correction (b) 72x72x16 at $f = 0.85$ (c) 64x32x8 at $f = 0.55$. Calculated as a decadal average at the end of the restarted simulations.



(a)



(b)



(c)

Figure 12: Meridional profiles of Atlantic surface freshwater flux per latitude strip in m/yr, for THC on and THC off restarts and difference (off - on) (a) 36x36x8 at $f = 0.8$ of default flux correction (b) 72x72x16 at $f = 0.85$ (c) 64x32x8 at $f = 0.55$. Calculated as a decadal average at the end of the restarted simulations.



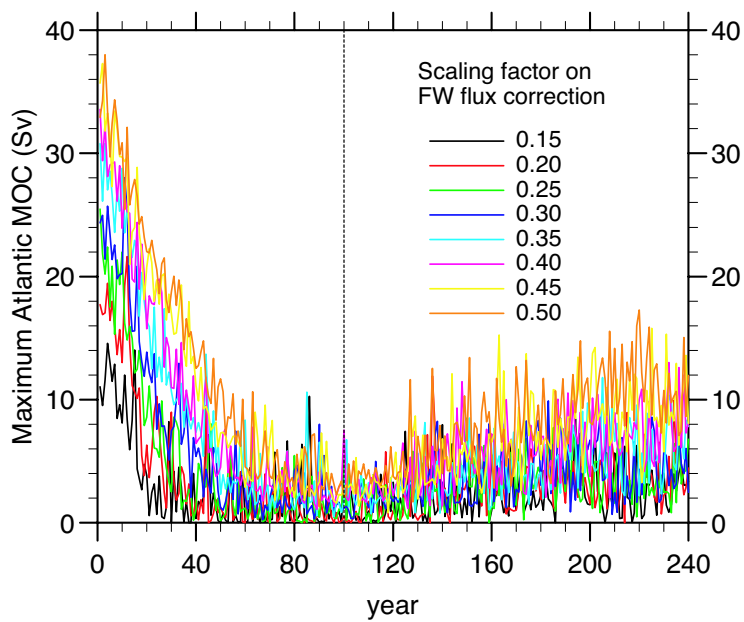


Figure 13: Net Atlantic surface freshwater flux difference (off - on) due to THC switch off, in different resolutions, as a function of default flux correction. Calculated as a decadal average at the end of the restarted simulations.



(a)



(b)

Figure 14: Freshwater hosing experiments (a) 0.1 Sv for 100 yr (b) 1.0 Sv for 100 yr, for variants of GENIE-2 with the 64x32x8 grid and differing fractions of the default flux correction. All ensemble members are within the bi-stable region for the THC.

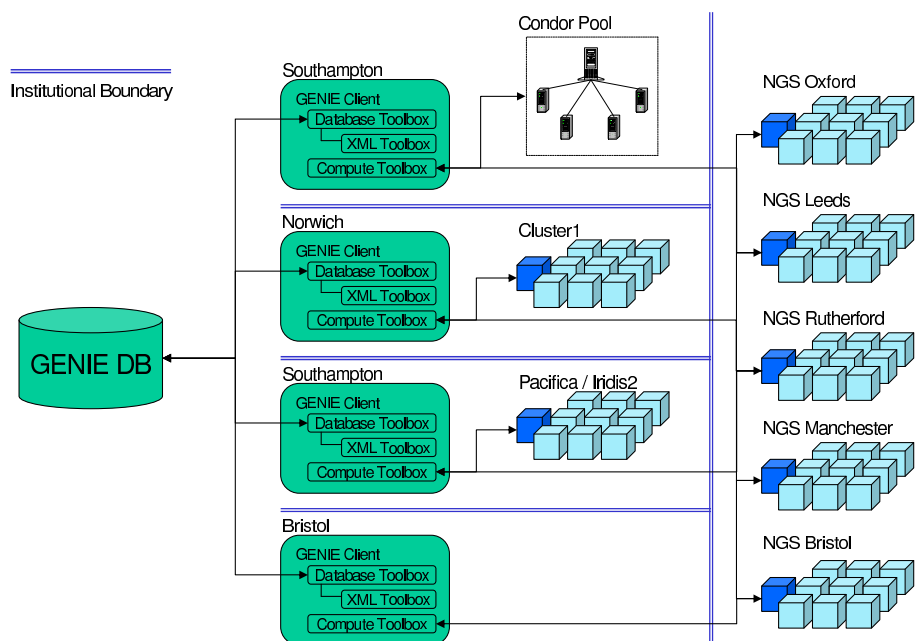


Figure 15: Collaborative study mediated by the GENIE database. Client installations at member institutions contribute to an ensemble study by retrieving work from the database and either submitting it to a resource within the institutional administrative domain or to a resource on the computational Grid

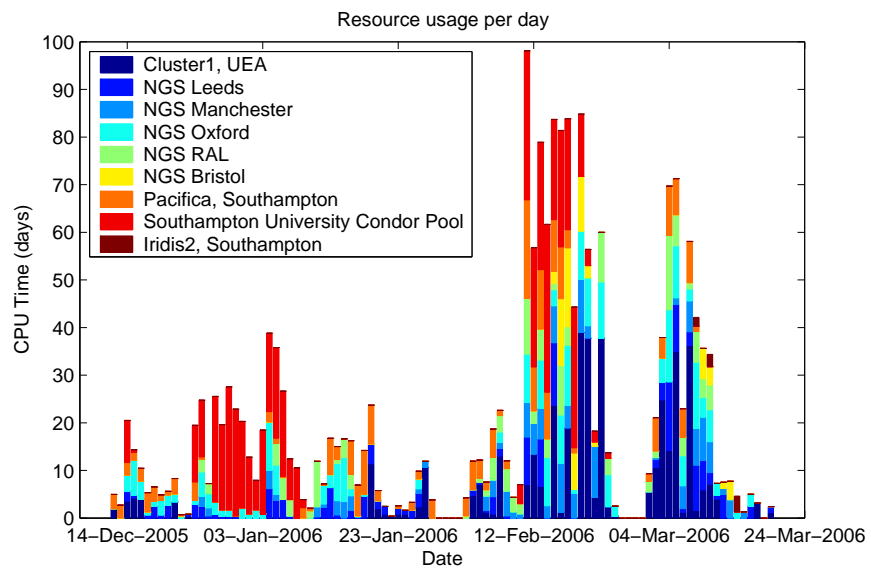


Figure 16: Daily resource usage for the twelve ensemble studies




Figure 17: Distributions of a) client job submissions and b) resource usage measured as total number of timesteps performed

Supplementary Information: The Grid ENabled Integrated Earth system modelling (GENIE) framework

T. M. Lenton* R. Marsh† A. R. Price‡ D. J. Lunt§
Y. Aksenov† J. D. Annan¶ T. Cooper-Chadwick‡
S. J. Cox‡ N. R. Edwards|| S. Goswami*
J. C. Hargreaves¶ P. P. Harris** Z. Jiao‡ V. N. Livina*
A. J. Payne§ I. C. Rutt§ J. G. Shepherd† P. J. Valdes§
G. Williams§ M. S. Williamson†† A. Yool†

January 31, 2007

Abstract

The Grid ENabled Integrated Earth system modelling (GENIE) framework supports modularity (i.e. interchangeable components) and scalability (i.e. variable resolution of the components), which aids traceability, meaning the ability to relate the process representation and results for different module choices and/or resolutions to one another. Here for each of the main components of the Earth system, we discuss the appropriate modelling approaches for our goals and introduce the component models adopted thus far. We describe their coupling to produce a range of computationally efficient Earth system models (ESMs) that span a spectrum from intermediate toward full complexity, and summarise the experiments undertaken with them thus far.

*School of Environmental Sciences, University of East Anglia, Norwich, UK and Tyndall Centre, UK. E-mail: t.lenton@uea.ac.uk Tel: +44-1603-591414 Fax: +44-1603-591327

†National Oceanography Centre, University of Southampton, Southampton, UK and Tyndall Centre, UK

‡Southampton e-Science Centre, University of Southampton, Southampton, UK

§School of Geographical Sciences, University of Bristol, Bristol, UK

¶Frontier Research Centre for Global Change, Yokohama, Japan

||Centre for Earth, Planetary, Space and Astronomical Research (CEPSAR), Earth Sciences, Open University, Milton Keynes, UK

**Centre for Ecology and Hydrology, Wallingford, UK

††Tyndall Centre, UK and School of Physics and Astronomy, University of Leeds, Leeds, UK

1 Components and models

Here, for each of the major components of the Earth system we discuss the appropriate modelling approaches for our goals and introduce the corresponding model(s) that have been included thus far within the GENIE framework. The components, models, and their notation are summarized in Table 1. As a rule, we have taken component models that already exist in the published literature and have developed them as necessary. Each of the physical components interfacing at the Earth’s surface – atmosphere, ocean, sea-ice, land, ice-sheets – can be substituted by a fixed boundary condition (static or time-varying). For completeness we introduce all currently available components here, thus there is some overlap with the main article.

1.1 Atmosphere

The atmosphere is the component of the Earth system with the shortest intrinsic timescale. Therefore, if the full primitive equations are to be solved, a relatively short time-step is required (about half an hour to an hour, depending on the model resolution). Hence fast Earth system models have traditionally made simplifications to these primitive equations, for example the quasi-geostrophic approximation employed in the EcBilt model (Opsteegh *et al.*, 1998). Alternatively they have done away with a dynamical atmosphere altogether in favour of an energy-moisture balance approach, as in the UVic model (Weaver *et al.*, 2001), or a statistical-dynamical approach as in CLIMBER-2 (Petoukhov *et al.*, 2000). Simplifications such as these mean that model predictions, for example of heat transport by eddies, may be inaccurate. Furthermore, aspects of the atmosphere’s structure have to be prescribed, including winds in the energy-moisture balance model. However, to achieve large ensemble integrations over multi-millennial timescales, the only practical approach is still to use a reduced-complexity atmosphere. Therefore, we currently have two different atmospheric models available in the GENIE framework – a primitive-equation model, and a fast energy-moisture-balance model. Even the more complex, fully dynamical model is capable of simulating 1,000 years in under a week on a single CPU, but the simpler model is two orders of magnitude faster. In the future we envisage using the two models in combination, with the dynamical atmosphere generating the wind field that is prescribed in the energy-moisture balance model.

1.1.1 IGCM

The Reading Intermediate General Circulation Model (IGCM3.1) has as its adiabatic core the 3-D spectral primitive equation model of Hoskins and Simmons (1975). We use a default horizontal resolution of T21 with 7 vertical levels. It is easy to change the model resolution, and we have begun exploring T10 and T42 horizontal resolutions and 5 and 22 vertical levels. The convection scheme of Tiedke (1993) has been included as an option, because of poor modeled convective precipitation in the original convective adjustment scheme of Betts (1986). For more details see the main article.

1.1.2 Energy-Moisture-Balance Model (EMBM)

The EMBM is a single-layer, 2-D model, comprising diffusive-advective transports of heat and moisture in the atmospheric boundary layer, similar to Weaver *et al.* (2001). Climatological annual-mean (Edwards and Marsh, 2005) or monthly-mean (Lenton *et al.*, 2006) wind fields are prescribed. Latitudinal profiles of heat diffusivity and moisture diffusivity are specified as external inputs. Albedo is treated as a zonal planetary mean (Edwards and Marsh, 2005) or separated into surface and atmosphere/cloud components, the latter treated as monthly means (Lenton *et al.*, 2006). Typically, the EMBM is implemented at an identical horizontal resolution to the underlying ocean and land surface and it takes five implicit timesteps per ocean timestep. See Edwards and Marsh (2005) and Lenton *et al.* (2006) for further details.

1.2 Atmospheric tracers

Atmospheric tracers are effectively a sub-component of the atmosphere. In future we plan to include an atmospheric chemistry module.

1.2.1 ATCHEM

A simple module ‘ATCHEM’ is currently used for storing the concentration of atmospheric gases – CO₂, CH₄, N₂, N₂O, O₂, H₂S, CFC-11 and CFC-12 as well as the relevant isotopic properties ($\delta^{13}\text{C}$, $\delta^{14}\text{C}$, $\delta^{15}\text{N}$, $\delta^{18}\text{O}$, $\delta^{34}\text{S}$). The storage array has the same resolution as the EMBM including a single vertical level. Currently, ATCHEM assumes the atmosphere is well mixed and homogenizes the chemical composition at each time-step. Future extensions will allow the atmosphere to modify the distribution of concentrations.

1.3 Ocean

Ocean models used for Earth system modelling should ideally be integrated long enough to reach an equilibrium state of the large-scale thermohaline circulation (THC). The THC is strongly affected by diapycnal mixing in the deep ocean, at low latitudes in the Indo-Pacific and at mid/high latitudes in the Southern Ocean. Diapycnal diffusivity in many parts of the deep ocean is small, and this mixing process occurs on a millennial timescale. In “traditional” general circulation models (GCMs) of the global ocean, based on the primitive equations, time-steps of typically 1 day at low resolution (3-4°) lead to integration times of around 100 years per day - e.g., the CLIO model (Opsteegh *et al.*, 1998). Hence several weeks of computing are required to reach THC equilibrium. Consequently, equilibrium studies of the THC are limited to just a few experiments in each instance and the sensitivity of simulated THC to model boundary conditions and/or key parameters cannot be easily established.

Alternative, faster approaches have been developed. For some time, a leading example has been the Hamburg “Large Scale Geostrophic” (LSG) model (Maier-Reimer *et al.*, 1993), which has longer but more computationally expensive time-steps of 1 month at 3.5° horizontal resolution (with 11 vertical levels). The LSG model was initially used for tracer studies (Maier-Reimer, 1993), but has since been coupled with atmospheric GCMs of varying complexity (ECHAM, PUMA) Prange *et al.* (2003).

Somewhat simpler, but several orders of magnitude more efficient, is a zonally-averaged 2.5-D ocean model (Wright and Stocker, 1991) which represents the dynamics of the THC in three basins (Pacific, Atlantic, Indian) interconnected by a circumpolar Southern Ocean. Crucially, the east-west pressure gradient (which supports a meridional overturning circulation) is parameterized in terms of the meridional pressure gradient. This invokes implicit assumptions regarding the circulation state which may be questioned if appropriate 3-D simulations are not available for calibration purposes. Despite this structural simplification, THC states are obtained which compare well with those simulated by 3-D models and large ensembles of multi-millennial integrations are feasible (Knutti *et al.*, 2002). The 2.5-D ocean model is used in the Bern 2.5-D model (Joos *et al.*, 1999) and in CLIMBER-2 (Petoukhov *et al.*, 2000).

1.3.1 GOLDSTEIN

A fast, intermediate complexity, 3-D frictional geostrophic model with linear drag, GOLDSTEIN (Global Ocean Linear Drag Salt & Temperature Equation Integrator) is used in GENIE. A coarse (18x18x8) resolution global version with idealised geometry (rectangular ocean basins), was first used for exploratory studies (Edwards and Shepherd, 2002) and for model development and testing (e.g. incorporation of marine biogeochemistry). This was extended to incorporate eddy-induced and isopycnal mixing following Griffies (1998) and to solve correctly for the flow around islands by Edwards and Marsh (2005). The linear momentum balance of GOLDSTEIN is marginally simpler than that of the LSG model, allowing for a local inversion of the baroclinic velocity field. Both models are very similar to GCMs in their nonlinear treatment of tracer dynamics.

Within the GENIE framework, two different grid types and a range of horizontal and vertical resolutions of GOLDSTEIN are currently available (Table 2). As well as the grids and resolutions used in the main paper, an “equal latitude” (36x60x8) version has been developed, on the longitude versus latitude horizontal grid, which retains similar equatorial resolution to the “baseline” resolution, and is intended for studies using more realistic sea-ice. It is easy to vary the number and spacing of depth intervals and the horizontal grid can be readily altered using Matlab-based tools that allow users to generate arbitrary resolution grids, and non-modern continental configurations. A suite of ocean topography files to represent changing continental and basin configurations through the Phanerozoic (based on PALEOMAP project reconstructions) is in preparation. At baseline resolution the model explicit timestep is around 3.5 days. A version with implicit timestep also exists.

1.3.2 Slab ocean

The IGCM3.1 originally included a thermodynamics-only ocean, which has been separated out as a “slab” ocean component. The default thickness of the slab is 25m. The temperature evolution of the surface is determined by the net latent, sensible, short-wave, and long-wave surface fluxes, and an implied ocean heat-flux term. The ocean heat-flux term is calculated in a prior calibration integration, as being that flux which results in SSTs equivalent to a given climatology.

1.4 Sea-ice

Many Earth system models in current use feature the simple sea-ice thermodynamics of Semtner (1976), following Hibler (1979), e.g. EcBilt/CLIO (Opsteegh *et al.*, 1998), CLIMBER-2 (Petoukhov *et al.*, 2000), and the McGill Paleoclimate Model (Wang and Mysak, 2000). An exception is the UVic Earth System Climate Model (ESCM), which includes several options for more sophisticated representation of sea ice thermodynamics and thickness distribution, including full rheology, ice thickness distribution and multi-layer thermodynamics (Weaver *et al.*, 2001). The computational expense of any 2-D sea ice scheme is small compared to 3-D ocean and atmosphere models, although the degree of complexity has implications for the realism of simulated sea ice distributions, especially in the Arctic where sea ice prevails throughout the year under present climate and rheology is important. The GENIE framework includes a recently modularised sea ice model allowing different levels of complexity in both the dynamics and thermodynamics, and an alternative slab sea-ice model.

1.4.1 Dynamics: EVP / Free drift / Ocean advection

There are three options for sea-ice dynamics. The most advanced is a model with Elastic-Viscous-Plastic (EVP) rheology (Hunke and Dukowicz, 1997) and a sub-cycling procedure for ice dynamics (Hunke, 2001). EVP dynamics is only appropriate if using higher resolution at high latitudes. A more computationally efficient alternative is the free-drift mode with the internal ice stress in the ice momentum balance equation set up to zero, as used in C-GOLDSTEIN (Edwards and Marsh, 2005). For both these options, the dynamical coupling between the ice and ocean occurs through the interfacial drag stress with McPhee's quadratic dependence on the shear velocity (Hunke and Zhang, 2001). Ice transport is calculated using either explicit or implicit advection schemes. The third and simplest option advects ice by the surface ocean currents with some horizontal ice diffusion. Here the wind stress is applied directly to the ocean surface. For all options, ice thickness and ice fractional coverage of each ocean cell are prognostic variables. If there is snow on the ice it is included in the momentum balance and transport calculations producing snow depth and snow fraction prognostic fields. There is also a simple scheme to account for land-fast ice appearance; if ice is thicker than the ocean depth available it is grounded. This should be useful in snowball Earth simulations.

1.4.2 Thermodynamics: Snow-ice / Ice-only

Thermal evolution of sea ice is simulated using either snow-ice thermodynamics or a simple ice-only thermal balance model. The former follows Semtner (1976), whereas the latter was used in C-GOLDSTEIN (Edwards and Marsh, 2005). In both models, ice slab grows vertically and can be formed in the open ocean. In the ice-only model, all precipitation is added directly to the ocean as freshwater. In the snow-ice model, snow cover is accumulated on top of the ice as the surplus of precipitation over evaporation. It melts if the surface temperature is above the melting point, or if the sea-ice is flooded by water above the melting point. If snow is too thick and ice is submerged under the snow load, then to reach hydrostatic balance the excess snow is converted into ice. During the conversion no salt or heat exchange between sea-ice and ocean occur. In both models sea ice

albedo is taken as a function of surface air temperature (Holland *et al.*, 1993). In the snow-ice model, the snow albedo has two different values for freezing and melting conditions at the surface. Neither snow nor ice has any thermal inertia and they are assumed to be fresh. The salt or fresh water is added to the upper layer of the ocean when ice volume increases or decreases or snow melts.

1.4.3 Slab sea ice

A very simple slab sea ice model has been separated from the IGCM. This is equivalent to the slab ocean component (see section 1.3.2), but with a slab thickness of 2m, and changes to the heat-capacity and albedo. In addition, there is the option to limit the implied ocean heat-flux term. For paleo simulations, tuning has indicated that in order to get a reasonable (defined here as being similar to that predicted by the HadSM3 model) simulation of modern, pre-industrial, and LGM sea-ice area, it is necessary to limit the implied ocean heat flux for sea-ice to have an absolute value less than 50Wm^{-2} .

1.5 Ocean biogeochemistry

Ocean biogeochemistry is effectively a sub-component of the ocean, because the dissolved chemical compounds created and transformed by marine organisms are subject to advection by ocean physics. A minimal set of tracers; dissolved inorganic carbon (DIC), alkalinity (Alk), plus one major nutrient - phosphate (PO_4) or nitrate (NO_3) - is required to capture the carbon cycle to a first approximation. EMICs that include marine biogeochemistry generally have this basic set of tracers, with the generic nutrient (usually PO_4) subject to Michaelis-Menten uptake in the surface ocean or restoration toward observations. Exceptions include the UVic model and CLIMBER-2, which now have more complex Nutrient Phytoplankton Zooplankton Detritus (NPZD) schemes, following Six and Maier-Reimer (1996). The GENIE framework has yet to include explicit marine biology, however we do include a very broad range of marine biogeochemical tracers. At run time the user can select a subset of these (without the need to recompile), with un-selected tracers excluded from calculation in order to limit the computational burden. The selection is automatically checked for self-consistency at run time.

1.5.1 BIOGEM

A marine biogeochemical model “BIOGEM” has been developed based on the scheme from an existing 2.5-D model (Ridgwell, 2001). BIOGEM handles the biological and chemical transformation (and physical sinking of solid materials) of a wide range of chemical species (tracers), whilst the ocean circulation model transports them in 3 dimensions by advection, diffusion, and convection, in the same way as temperature and salinity. The pre-defined and selectable tracers in BIOGEM include: DIC, PO_4 , Alk, O_2 , NO_3 , H_4SiO_4 (silicic acid), Fe, DOM C (carbon in dissolved organic matter), DOM N, DOM P, and DOM O_2 . For addressing the biogeochemistry of low oxygen environments, CH_4 , NH_3 , N_2O , and N_2 are defined as additional tracer options, while transformations between SO_4 , H_2S are already explicitly included (Ridgwell *et al.*, 2006a). Other selectable

tracers include; Ca, B, and F (for more explicit carbonate chemistry representations), an iron-binding ligand, trace metals of relevance to paleoceanography (e.g. Cd, Ge), and diagnostic tracers of ocean mixing (SF_6 , CFC-11, CFC-12) and particulate fluxes (^{230}Th , ^{231}Pa). The isotopes of various tracer species are tracked enabling the calculation of: $\delta^{13}\text{C}$, $\delta^{14}\text{C}$, $\delta^{15}\text{N}$, $\delta^{18}\text{O}$, $\delta^{30}\text{Si}$, and $\delta^{34}\text{S}$ isotopic properties. Associated with many of the dissolved tracers are particulate (solid) ones, for instance; particulate organic carbon (POC) which is associated with DIC, or calcium carbonate (CaCO_3) which is associated with both DIC and ALK.

Biological uptake in the surface layer of the ocean has thus far been based around a single nutrient (PO_4) limitation of biological production, although we now have the scope to incorporate co-limitation by a range of nutrients. Both a phosphate restoring scheme (either towards observations or zero) (Cameron *et al.*, 2005) and Michaelis-Menten type kinetics (Ridgwell *et al.*, 2006a) have been used. Air-sea gas exchange fluxes are calculated as explicit functions of wind speed using gas solubility coefficients and gas transfer velocities that depend on environmental conditions, following Wanninkhof (1992).

BIOGEM calculates the redistribution of tracer concentrations due to biological uptake of dissolved species in the surface ocean, sinking of particulate species, and their remineralisation (due to metabolic and dissolution processes) back to dissolved species at depth. This sequence of processes occurs rapidly relative to transport by the large-scale circulation of the ocean. An array stores the concentration of suspended and settling solids throughout the ocean. This allows the option of either parameterizing instantaneous redistribution of dissolved tracers due to prescribed remineralisation profiles (the default), or explicitly calculating sinking speeds and the transfer of solids between depth layers. The fractionation between carbon and nutrients and between different isotopes as particles sink through the water column can be represented (although by default we assume no fractionation during remineralisation). The residual flux of particulate matter that escapes remineralisation and reaches the sea floor either provides input to the sediment module, or if that is not selected, is completely remineralised and dissolved constituents returned to the overlying ocean.

1.6 Marine sediments

The circa 10,000 year adjustment time of marine sediment composition is too long for ocean-sediment processes to be included in most 3-D ocean carbon cycle models. Hence the majority simply return all settling material that reaches the ocean floor instantaneously back to the ocean. Some groups have undertaken off-line studies in which a pre-calculated, fixed ocean circulation field provides input to a sediment model, e.g. (Archer *et al.*, 2000). However, this does not allow feedbacks via atmospheric CO_2 and climate to be accounted for. A sediment model (Archer, 1991) has recently been coupled to the zonally-averaged (2.5D) ocean of CLIMBER-2. However, testing such a model against data is problematic because the substantial zonal heterogeneity in the observed sediment composition cannot be captured (Ridgwell, 2001). This problem is even more acute in box models. The GENIE framework interactively couples marine sediments to a 3-D ocean and marine biogeochemistry. The sediments are treated as an optional module, which resides under the ocean and exchanges matter of various types with the ocean biogeochemistry module (Ridgwell and

Hargreaves, 2006). A separate sediment module is assigned to each ocean grid point.

1.6.1 SEDGEM

The sediment module, SEDGEM (for; SEDimentGEOchemical Model), is based on the scheme of Ridgwell (2001). Sediments comprise of a single surface layer underlain by a stack of storage layers. The surface layer represents the zone in which bioturbation effectively homogenizes solid composition with depth and where the primary diagenetic processes (such as carbonate dissolution) take place. The storage layers provide a synthetic proxy record of biogeochemical and climate changes occurring in the ocean and can be contrasted directly to paleoceanographic cores (Ridgwell, 2001). A diffusive-like transfer is prescribed to represent the vertical mixing of solids due to biological activity at depths below the surface layer.

Currently, we consider the preservation of calcium carbonate (CaCO_3) (following Archer (1991)) and opal (following Ridgwell *et al.* (2002)) as well as a refractory (detrital phase). However, only the preservation of CaCO_3 in GENIE has been formally described and characterized to date (Ridgwell and Hargreaves, 2006). Work is underway incorporating as a user-selectable option the diagenetic model of Archer *et al.* (2002) (MUDS), which will allow the explicit prediction of organic carbon preservation (and burial). A variety of properties of solids are also tracked in the sediment module – currently, $\delta^{13}\text{C}$ and $\delta^{14}\text{C}$ as well as a numerical age of sedimenting carbonate (Ridgwell, 2001).

1.7 Land surface physics and biogeochemistry

In principle, schemes for land-surface physics (dealing with energy and water) and biogeochemistry (dealing with carbon, nutrients and potentially vegetation dynamics) can be considered as separate components. In reality, however, they are intertwined because the fluxes of water, associated latent heat and carbon from the vegetated land-surface occur through the same plant-leaf stomata. Opening and closing of stomata is tied to photosynthesis with vegetation acting to maximize carbon uptake while minimizing water loss. A separation can be maintained by treating photosynthesis and stomatal behavior as part of the land surface physics and treating the corresponding photosynthetic carbon flux as an input to land surface biogeochemistry and vegetation dynamics. In practice this means treating the biogeochemistry and any vegetation dynamics as a sub-component of the land surface physics. This is the approach taken by e.g. the Hadley Centre in the separation of the MOSES land surface physics scheme and TRIFFID dynamic global vegetation model (DGVM). We follow this with a simplified version of MOSES, plus TRIFFID as a sub-component. We have also developed an Efficient Numerical Terrestrial Scheme (ENTS) in which land-surface physics and biogeochemistry are combined. Finally, we have separated the land surface physics scheme from the IGCM. Other EMICs typically have land surface physics based on simple bucket models, or the biosphere-atmosphere transfer scheme (Dickinson *et al.*, 1986). To this they often add VECODE (Brovkin *et al.*, 1997).

1.7.1 GENIE-land / Fast MOSES

The GENIE-land surface scheme is based on a simplified version of the MOSES land-surface scheme (Cox *et al.*, 1999) similar to that being used by the UVic group (Meissner *et al.*, 2003, Matthews *et al.*, 2004). It allows a longer time step than full MOSES (typically the same as the atmospheric time step, i.e. around 6 hours rather than 30 minutes), by excluding fast processes (e.g. canopy interception, separation of fast and slow runoff) and correcting for the local sunlit day length rather than resolving the diurnal cycle. Leaf-level photosynthesis is calculated for 5 plant functional types (PFTs; broadleaf and needle-leaf trees, C3 and C4 grasses, and shrubs) as a function of near-surface climate and vegetation state (Collatz *et al.*, 1991, Collatz *et al.*, 1992). Like MOSES, GENIE-land is a tiled scheme in which the surface energy budget is calculated separately for 7 tile types (5 PFTs, bare soil, and land-ice). The grid-box representative output fields required by other GENIE components are derived internally by GENIE-land from the tiled values. Land-surface hydrology, including snow-melt, soil moisture change and runoff, are calculated at the grid-box rather than the tile level. An energy conserving parameterisation of snow is based around the sub-surface temperature, with precipitation falling as snow below a threshold temperature, and snow-melt occurring when the ground heat flux would raise the sub-surface temperature above a threshold of 0°C. This provides a delay to the spring warming of high-latitudes by accounting for the energy used in water phase change. Tile albedos are a function of snow depth and account for masking of snow by tall vegetation. Runoff is routed to the ocean using pre-defined destinations, which are currently based on the TRIP routing data (Oki and Sud, 1998).

1.7.2 TRIFFID

The full TRIFFID model from the Hadley Centre (Cox, 1998) captures vegetation dynamics and carbon storage in vegetation and soil. It allows 5 PFTs to grow and compete for the available space based on the accumulated carbon fluxes calculated by GENIE-land and a Lotka-Volterra approach to competition between the PFT populations.

1.7.3 ENTS

The Efficient Numerical Terrestrial Scheme (ENTS) is a minimal representation of land surface physics, hydrology and carbon cycling for use in a spatial context (Williamson *et al.*, 2006). The model does not distinguish between plant functional types. Instead ENTS has a single pool of vegetation carbon at each land grid point which feeds a single pool of soil carbon through carbon fluxes of leaf litter. These land carbon pools are connected to atmospheric carbon dioxide through photosynthesis, plant respiration and soil respiration.

A simple bucket model is used to represent hydrology. In each grid cell the bucket is filled by precipitation and emptied by evaporation. If the bucket is filled beyond its capacity the excess is transported to the ocean as runoff. The physical properties in the model such as soil water and land temperature are tightly linked to the state of the biosphere and vice-versa. For example moisture and heat fluxes are determined by the roughness of the surface which is in turn determined by the amount of vegetation. Land surface albedo, affecting the

absorption of solar radiation, is a function of the vegetation and soil carbon densities, denser vegetation causing more radiation to be absorbed. The field capacity of the water buckets is controlled by the amount of soil carbon. Greater densities of soil carbon act like a sponge enabling the land to hold more water.

ENTS also has a very simple snow scheme to capture the dominant albedo effect. A grid cell is set to snow albedo (taking into account vegetation) when precipitation falls and the air and land temperatures are below critical values. The snow persists until one of the temperatures rises above the critical value.

In the GENIE framework solar insolation is prescribed as a diurnal average. This enables an efficient solution to the land radiation balance by calculating for the steady state. With this approximation ENTS can be coupled to the atmosphere at the same frequency as the ocean.

1.7.4 IGCM-land

The IGCM-land module was originally an integral part of the IGCM, as described in de Forster *et al.* (2000). In brief, it parameterises the surface energy, moisture, and momentum fluxes as a function of, respectively, temperature, moisture, and velocity vertical gradient. It also uses a bucket representation of soil-moisture, and parameterises changes in albedo related to snow cover. Some minor changes to the original scheme have been made so that it rigorously conserves energy and water.

1.8 Ice sheets

Ice sheet modelling poses a horizontal resolution challenge for Earth system modelling. Important physical processes that determine the evolution of the ice sheet tend to occur in small regions — examples of this are the melting of ice in a narrow margin at the edge of an ice sheet, and the formation of mountain glaciers, which can potentially lead to large-scale glacial inception. In both cases, the scales involved are in the tens of km. Traditionally, ice sheet modelling in an Earth system context has been done on a global latitude-longitude grid. However, with this approach, it is hard to generate the required resolution, as well as being potentially inefficient (since only limited regions are generally ice-covered), especially if a uniform resolution is used. To address this issue, we have coupled a regional ice model (GLIMMER) into the GENIE framework. Other ice sheet models have been coupled to EMICs, ranging from zonally-averaged models (Gallee *et al.*, 1992) to 3-D models (Calov *et al.*, 2005).

1.8.1 GLIMMER

GLIMMER is based on an existing ice sheet model (Payne, 1999). It uses a projected grid scalable over 20-100+ km grid cells. The time steps (years to decades) are longer than other components. Different levels of physics are available; mass balance calculations using energy balance and day-degree calculations; coupled ice flow, thermodynamics and ice-thickness evolution; and isostasy. Isostasy along with a simple parameterization of sliding is necessary for glacial-interglacial simulations and, in particular, work on Heinrich events. The code uses quasi-object-oriented techniques in Fortran 95 to allow the model

to be run as multiple instances over different regions of the globe. Thus the individual ice sheets of the world (Greenland, Antarctica, Laurentide, Scandinavian etc) can be modeled on separate, local grids and the need for high-resolution global coverage is avoided.

2 Coupling and Software

The GENIE framework couples the components described above in a modular fashion, such that they can be readily interchanged at compilation time. Different resolutions of a particular component model are treated as different instances of that component. Spatial grids and masks can differ between components, in particular the atmosphere, ocean, and sea-ice, because they have been derived from different sources. To interpolate between the different grids, we currently use a simple bilinear interpolation routine. In regions where there is a mismatch in land-sea mask, we extrapolate from the nearest neighbor. This process is in general non-conservative, so after interpolation we apply a correction factor in the tropics (30°S to 30°N) to ensure conservation of energy and moisture.

At present the default approach to component coupling is a master Fortran routine with a number of compilation time switches and a corresponding suite of configuration files to set up different Earth system models. We are currently exploring other software approaches to flexible coupling, which will be described in separate publications. Achieving full modularity poses considerable scientific and technical challenges – for example in reconciling the different time-stepping of atmosphere, ocean and sea-ice routines that is inherent in the parent models we are modularising. This represents ongoing work, in that not all possible combinations of the modules we have described are currently available and working. Furthermore, the modularization of ENTS is still in progress.

The GENIE framework is intended to be sufficiently generic that alternative models for any given component can be included. In practice, this means that we must define the quantities exchanged at each interface between components. We subdivide the interfaces between components of the Earth system into “physical” (energy, momentum, water, salt) and “biogeochemical” (carbon, nutrients, dust). Most existing studies have only considered physical couplings, furthermore, biogeochemical couplings generally depend on information from the physical couplings but not *vice versa*. Thus a purely physical coupling of various instances of Earth system model (with fixed biogeochemistry) should be achievable.

Proposals for standard physical interfaces for GCM modelling have been put forward for Atmosphere–Ocean–Sea-ice coupling by PRISM (the PRogram for Integrated Earth System Modelling) (Guilyardi *et al.*, 2002), and for Atmosphere–Land-surface coupling by PILPS (the Project for Inter-comparison of Land surface Parameterisation Schemes) (Polcher *et al.*, 1998). These provide useful overviews of the key scientific and technical issues, and are providing the basis for our ongoing modularisation work.

For the biogeochemical interfaces, no prototype standard exists hence we have started to develop one. To maximize overall computational speed, the exchange of biogeochemical tracers between surface sediments (SEDGEM) and overlying ocean (BIOGEM) is carried out asynchronously, with the composition of deep-sea sediments and the return dissolution flux to the ocean re-calculated

at a longer time-step (by default, 1 year) than ocean circulation and biogeochemistry. This coupling is handled by an interfacing module which integrates the sediment rain flux over the SEDGEM time-step before it is added to the surface sediment layer. In the opposite direction, the sediment dissolution flux supplied to the ocean is assumed constant for the duration of the SEDGEM time-step. The interfacing module also enables simple transformations between different grid resolutions, allowing us to use a finer grid for the sediments than the ocean circulation (Ridgwell and Hargreaves, 2006).

2.1 Software tools

The GENIE framework is developed by a team of researchers spread over many institutions. Several software development tools and approaches have been key to making this effort coordinated and efficient. A single, central source code repository is used, maintained using a version control system – the open source Concurrent Version System (CVS). This enables developers to keep their copy of the source code up to date with other developers changes and so avoid the pitfalls of inconsistent versions. CVS provides the ability to view differences between versions of the source code (and a log of the changes). When used in concert with frequent testing, this provides a means to rapidly track down and correct coding bugs.

In line with good software engineering practice, we have adopted a policy of frequently testing the source code during development. We test for water conservation, restarts, and the integrity of model runs relative to a suite of standard simulations. To make testing as easy as possible, we have created an integrated build and test system, stored alongside the source code in the repository. This is particularly challenging for a model which embodies sensitivity to initial conditions and is developed by researchers working with different compilers and hardware architectures. It is worthwhile because it facilitates rigorous quality control.

A Matlab graphical user interface (GUI) called “Aladdin” allows the novice user to generate multiple configuration files and then writes a master script to submit model ensembles to Grid computing resources.

2.2 Tuning methods

A substantial part of the workload in the creation of a model must be devoted to calibrating or tuning model parameters. Often, the physically correct parameter values are not known with much precision, and furthermore, minor missing processes in the model may lead to modest biases (eg in overall energy flux) which it can be desirable to tune out in order for the model to simulate a reasonable climatological state. The nonlinear and multivariate relationship between parameters and model climate makes model tuning a challenging task, and the implementation and development of tuning and estimation techniques has proceeded in step with the model development.

A multi-parameter Latin hypercube exploration of 12-dimensional parameter space is described by Edwards and Marsh (2005). A variety of alternative approaches using the Design Search and Optimization package OPTIONS are described by Price *et al.* (2006b). These include the application of a Kriging meta-model (Keane, 2003) and the use of a Genetic Algorithm (Michalewicz,

1992) to locate minima in 12- and 30- dimensional parameter space respectively. An implementation of the non-dominated sorting genetic algorithm (NSGA-II) (Deb *et al.*, 2002) has been used to perform a multi-objective tuning of GENIE models (Price *et al.*, 2006a). A novel, computationally cheap approach has also been developed by Hargreaves *et al.* (2004) and Annan *et al.* (2005b), based on the ensemble Kalman filter (Evensen, 2003). This relies on a linear approximation to the error covariances. Nevertheless, the method seems to work well in the presence of significant nonlinearity. Finally, a proximal analytic centre cutting plane method (proximal ACCPM), which is highly efficient for scalar processing, is described by Beltran *et al.* (2005). Using any or all of these methods, it is possible to both tune the model to a best fit and also explore the range of possible outcomes (Annan *et al.*, 2005a, Hargreaves and Annan, 2006).

A potentially powerful tool for the analysis of climate model stability and feedbacks is the adjoint model, which effectively calculates the derivative of the model response with respect to model parameters. To demonstrate the feasibility of constructing an adjoint for GENIE an adjoint has recently been constructed for the C-GOLDSTEIN configuration (described below) using the TAPENADE software, and tested by utilising the calculated gradients to reproduce the ACCPM parameter calibration exercise.

3 Earth system models and existing studies

A range of Earth System Models have been constructed (or reconstructed) from the GENIE framework and a range of experiments have been conducted with them (Table 3). The notation system used to specify complete ESMs in the GENIE framework is: `genie_aa_oo_ss_ll_ii_vv_gg_dd_cc`. This is based on the labels for components given in Table 1. Particular models for each component have a 2 letter identifier (Table 1), which is substituted for the relevant label. Where a component is not included at all, that part of the notation is removed.

3.1 GENIE-1: Simple atmosphere ESMs

Our starting model was the coupled 3-D ocean (36x36x8 resolution), 2-D atmosphere (EMBM), and 2-D sea-ice (free-drift) model “C-GOLDSTEIN” (Edwards and Marsh, 2005). We have modularised and recreated this using the GENIE framework and have reproduced the original results. The modular version is known as `genie_ea_go_gs`. C-GOLDSTEIN has been tuned using a variety of techniques (see Section 2.2). Distributed computing was used to undertake a systematic search for bistability of the thermohaline circulation (Marsh *et al.*, 2004). A large number of simulations were also undertaken for the IPCC Fourth Assessment Report. In a separate strand of work, C-GOLDSTEIN has been coupled to a highly simplified global economic growth model (Drouet *et al.*, 2005, Drouet *et al.*, 2006) and to MERGE: a much more detailed, sectorial macroeconomic and energy supply sector model (Bahn *et al.*, 2006). In both cases C-GOLDSTEIN replaced a highly idealised climate module, resulting in an integrated assessment model (IAM) with a much more detailed representation of climate dynamics. The consistent representation of climate-economy feedbacks was achieved within an oracle-based overall optimisation approach

(Drouet *et al.*, 2005) using ACCPM, which nevertheless demanded many climate integrations and thus relied heavily on the efficiency of the climate model.

ESMs using the simple EMBM atmosphere are grouped under the heading “GENIE-1”. Ocean biogeochemistry (BIOGEM), ocean sediments (SEDGEM), and the ENTS land surface and carbon cycle, have all been added to the physical components of C-GOLDSTEIN. A version with only BIOGEM added has been used to examine the effect of different ocean circulation states, marine carbon pumps, and climate feedback on atmospheric CO₂ (Cameron *et al.*, 2005). The biogeochemistry of this model has been tuned using an EnKF (Ridgwell *et al.*, 2006a), as has a version with sediments (SEDGEM) added (Ridgwell and Hargreaves, 2006). The coupling of the ENTS land model necessitated some improvements to the EMBM, including the addition of seasonality (Lenton *et al.*, 2006). With ENTS and BIOGEM together the model has a closed carbon cycle and has been forced with CO₂ emissions in future global change studies (Lenton *et al.*, 2006). The more complex GENIE-land surface plus TRIFFID have also been coupled under the EMBM, as in the UVic model (Meissner *et al.*, 2003). The physical core of GENIE-1 using ENTS with fixed vegetation cover, and prescribed changes in ice sheet cover, orography, orbital forcing and CO₂ has been used for transient simulations of the last deglaciation (Lunt *et al.*, 2006), and recently, the whole of the last glacial-interglacial cycle (Marsh *et al.*, 2006). GENIE-1 has also recently been coupled to MERGE to form a detailed, consistent IAM with a closed carbon cycle.

3.2 GENIE-2: Dynamical atmosphere ESMs

ESMs using the IGCM dynamical atmosphere are grouped under the heading “GENIE-2”, because the choice of atmosphere model is the main determinant of overall model speed. The IGCM at T21 resolution integrates circa 10 years per CPU hour on an AMD 64 bit processor, in contrast to GENIE-1 at baseline resolution achieving over 1000 years per CPU hour. The IGCM (T21) in isolation, i.e. coupled to a fixed ocean, fixed sea-ice and its original land surface routine is known as `genie.ia.fo.fi.il`. This has been tuned using an EnKF (Annan *et al.*, 2005a) and a Genetic Algorithm (Price *et al.*, 2006b). The IGCM has been asynchronously coupled to GLIMMER and used to study glacial inception (Rutt *et al.*, 2005). The IGCM has also been coupled to GOLDSTEIN to create a fully 3-D ocean-atmosphere model. This has been achieved for a range of surface grids and ocean resolutions (36x36x8, 72x72x16, 64x32x8). The different sea-ice options are available along with GENIE-land or the IGCM’s own land surface scheme. Different variants and resolutions of the fully coupled model are currently being tuned using different methods. In the main article we undertake a first large ensemble experiment using different resolutions of GENIE-2.

Acknowledgement The UK Natural Environment Research Council funded GENIE (NER/T/S/2002/00217) and funds the ongoing GENIE_{fy} (NE/C515904) project through its e-Science programme. The initial inclusion of BIOGEM and the development of ENTS were funded by the Tyndall Centre, UK. The inclusion of an EVP sea ice option, and the development of alternative ocean grids were funded by the NERC Rapid Climate Change project “The role of the cryosphere on modulating the THC of the North Atlantic” (Ref. NER/T/S/2002/00462).

References

- Annan, J. D., Lunt, D. J., Hargreaves, J. C., and Valdes, P. J. (2005a). Parameter estimation in an atmospheric GCM using the Ensemble Kalman Filter. *Nonlinear Processes in Geophysics*, **12**, 363–371.
- Annan, J. D., Hargreaves, J. C., Edwards, N. R., and Marsh, R. (2005b). Parameter estimation in an intermediate complexity earth system model using an ensemble Kalman filter. *Ocean Modelling*, **8**, 135–154.
- Archer, D. E. (1991). Modeling the calcite lysocline. *Journal of Geophysical Research (Oceans)*, **96**, 17037–17050.
- Archer, D. E., Winguth, A., Lea, D., and Mahowald, N. (2000). What caused the glacial/interglacial pCO₂ cycles. *Reviews of Geophysics*, **38**, 159–189.
- Archer, D. E., Morford, J. L., and Emerson, S. R. (2002). A model of suboxic sedimentary diagenesis suitable for automatic tuning and gridded global domains. *Global Biogeochemical Cycles*, **16**(1), GB001288.
- Bahn, O., Drouet, L., Edwards, N. R., Haurie, A., Knutti, R., Kypreos, S., Stocker, T. F., and Vial, J.-P. (2006). The coupling of optimal economic growth and climate dynamics. *Climatic Change*, page in press.
- Beltran, C., Edwards, N. R., Haurie, A., Vial, J.-P., and Zachary, D. S. (2005). Oracle-based optimisation applied to climate model calibration. *Environmental Modelling and Assessment*, **11**, 31–43.
- Betts, A. (1986). A new convective adjustment scheme. 1. observational and theoretical basis. *Q J R Meteorol Soc*, **112**, 667–691.
- Brovkin, V., Ganopolski, A., and Svirezhev, Y. (1997). A continuous climate-vegetation classification for use in climate biosphere studies. *Ecological Modelling*, **101**, 251–261.
- Calov, R., Ganopolski, A., Claussen, M., Petoukhov, V., and Greve, R. (2005). Transient simulation of the last glacial inception. Part I: glacial inception as a bifurcation in the climate system. *Climate Dynamics*, **24**(6), 545–561.
- Cameron, D. R., Lenton, T. M., Ridgwell, A. J., Shepherd, J. G., and Marsh, R. (2005). A factorial analysis of the marine carbon cycle and ocean circulation controls on atmospheric CO₂. *Global Biogeochemical Cycles*, **19**, GB4027.
- Challenor, P. G., Hankin, R. K. S., and Marsh, R. (2006). *Towards the Probability of Rapid Climate Change*. In “Avoiding Dangerous Climate Change”. Cambridge University Press, Schellnhuber, H J., Cramer, W., Nakicenovic, N., Wigley, T. and Yohe, G (Eds.), Cambridge.
- Collatz, G. J., Ball, J. T., Grivet, C., and Berry, J. A. (1991). Physiological and environmental regulation of stomatal conductance, photosynthesis and

- transpiration: a model that includes a laminar boundary layer. *Agricultural and Forest Meteorology*, **54**, 107–136.
- Collatz, G. J., Ribas-Carbo, M., and Berry, J. A. (1992). A coupled photosynthesis-stomatal conductance model for leaves for C4 plants. *Australian Journal of Plant Physiology*, **19**, 519–538.
- Cox, P. M. (1998). TRIFFID: A Top-Down Model of Interactive Foliage and Flora Including Dynamics. Climate Research Technical Note 86. Technical report, Hadley Centre, UK Met. Office.
- Cox, P. M., Betts, R. A., Bunton, C. B., Essery, R. L. H., Rowntree, P. R., and Smith, J. (1999). The impact of new land surface physics on the GCM simulation of climate and climate sensitivity. *Climate Dynamics*, **15**(3), 183–203.
- de Forster, P. M., Blackburn, M., Glover, R., and Shine, K. P. (2000). An examination of climate sensitivity for idealised climate change experiments in an intermediate general circulation model. *Climate Dynamics*, **16**, 833–849.
- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *Evolutionary Computation, IEEE Transactions on*, **6**(2), 182–197.
- Dickinson, R. E., Henderson-Sellers, A., Kennedy, P. J., and Wilson, M. F. (1986). Biosphere-atmosphere transfer scheme (BATS) for the NCAR community climate model. Technical Report 275, NCAR.
- Drouet, L., Edwards, N. R., Beltran, C., Haurie, A. B., Vial, J.-P., and Zachary, D. S. (2005). *An Oracle Method to Couple Climate and Economic Dynamics. In "The Coupling of Climate and Economic Dynamics"*. Kluwer, Haurie, A. and Viguier, L. (eds), Dordrecht.
- Drouet, L., Edwards, N. R., and Haurie, A. B. (2006). Coupling climate and economic models in a cost-benefit framework: A convex optimization approach. *Environmental Modelling and Assessment*, **11**, 101–114.
- Edwards, N. R. and Marsh, R. (2005). Uncertainties due to transport-parameter sensitivity in an efficient 3-D ocean-climate model. *Climate Dynamics*, **24**, 415–433.
- Edwards, N. R. and Shepherd, J. G. (2002). Bifurcations of the thermohaline circulation in a simplified three-dimensional model of the world ocean and the effects of inter-basin connectivity. *Climate Dynamics*, **19**(1), 31–42.
- Evensen, G. (2003). The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Climate Dynamics*, **53**, 343–367.
- Fanning, A. G. and Weaver, A. J. (1996). An atmospheric energy-moisture model: Climatology, interpentadal climate change and coupling to an ocean general circulation model. *Journal of Geophysical Research*, **101**,

15111–15128.

- Gallee, H., van Ypersele, J.-P., Fichefet, T., Marsiat, I., Tricot, C., and Berger, A. (1992). Simulation of the last glacial cycle by a coupled, sectorially averaged climate-ice sheet model. Part II: Response to insolation and CO₂ variation. *Journal of Geophysical Research*, **97**, 15713–15740.
- Griffies, S. M. (1998). The Gent-McWilliams skew flux. *Journal of Physical Oceanography*, **28**, 831–841.
- Guilyardi, E., Polcher, J., Fichefet, T., and Planton, S. (2002). PRISM discussion paper: a proposal for standard atmosphere/ocean/sea ice interfaces.
- Hargreaves, J. C. and Annan, J. D. (2006). Using ensemble prediction methods to examine regional climate variations under global warming scenarios. *Ocean Modelling*, **11**, 174–192.
- Hargreaves, J. C., Annan, J. D., Edwards, N. R., and Marsh, R. (2004). An efficient climate forecasting method using an intermediate complexity Earth System Model and the ensemble Kalman filter. *Climate Dynamics*, **23**, 745–760.
- Hibler, W. D. (1979). A dynamic thermodynamic sea ice model. *Journal of Physical Oceanography*, **9**, 815–846.
- Holland, D. M., Mysak, L. A., Manak, D. K., and Oberhuber, J. M. (1993). Sensitivity study of a dynamic thermodynamic sea ice model. *Journal of Geophysical Research (Oceans)*, **98**, 2561–2586.
- Hoskins, B. J. and Simmons, A. J. (1975). A multi-layer spectral model and the semi-implicit method. *Quarterly Journal of the Royal Meteorological Society*, **101**, 637–655.
- Hunke, E. (2001). Viscous plastic dynamics with the EVP model: Linearization issues. *Journal of Computational Physics*, **170**, 18–38.
- Hunke, E. and Zhang, Y. (2001). Recent Arctic change simulated with a coupled ice-ocean model. *Journal of Geophysical Research*, **106**(C3), 4369–4390.
- Hunke, E. C. and Dukowicz, J. K. (1997). An elastic-viscous-plastic model for sea-ice dynamics. *Journal of Physical Oceanography*, **27**(9), 1849–1867.
- Joos, F., Plattner, G.-K., Stocker, T., Marchal, O., and Schmittner, A. (1999). Global warming and marine carbon cycle feedbacks on future atmospheric CO₂. *Science*, **284**, 464–467.
- Keane, A. J. (2003). Wing optimization using design of experiment, response surface and data fusion methods. *Journal of Aircraft*, **40**(4), 741–750.

- Knutti, R., Stocker, T. F., Joos, F., and Plattner, G.-K. (2002). Constraints on radiative forcing and future climate change from observations and climate model ensembles. *Nature*, **416**, 719–723.
- Lenton, T. M., Williamson, M. S., Edwards, N. R., Marsh, R., Ridgwell, A. J., Price, A. R., Shepherd, J. G., Cox, S. J., and The GENIE team (2006). Millennial timescale carbon cycle and climate change in an efficient Earth system model. *Climate Dynamics*, **26**, 687–711.
- Lunt, D. J., Williamson, M. S., Valdes, P. J., and Lenton, T. M. (2006). Comparing transient, accelerated, and equilibrium simulations of the last 30,000 years with the genie-1 model. *Climate of the Past Discussions*, **2**, 267–283.
- Maier-Reimer, E. (1993). Geochemical cycles in an ocean general circulation model. Preindustrial tracer distributions. *Global Biogeochem. Cycles*, **7**, 645–677.
- Maier-Reimer, E., Mikolajewicz, U., and Hasselmann, K. (1993). Mean circulation of the Hamburg LSG OGCM and its sensitivity to the thermohaline surface forcing. *Journal of Physical Oceanography*, **23**, 731–757.
- Marsh, R., Yool, A., Lenton, T. M., Gulamali, M. Y., Edwards, N. R., Shepherd, J. G., Krznaric, M., Newhouse, S., and Cox, S. J. (2004). Bistability of the thermohaline circulation identified through comprehensive 2-parameter sweeps of an efficient climate model. *Climate Dynamics*, **23**, 761–777.
- Marsh, R., Smith, M. P. L. M., Rohling, E. J., Lunt, D. J., Lenton, T. M., Williamson, M. S., and Yool, A. (2006). Modelling ocean circulation, climate and oxygen isotopes in the ocean over the last 120,000 years. *Climate of the Past Discussions*, **2**, 657–709.
- Matthews, H. D., Weaver, A. J., Meissner, K. J., Gillett, N. P., and Eby, M. (2004). Natural and anthropogenic climate change: Incorporating historical land cover change, vegetation dynamics and the global carbon cycle. *Climate Dynamics*, **22**, 461–479.
- Meissner, K. J., Weaver, A. J., Matthews, H. D., and Cox, P. M. (2003). The role of land surface dynamics in glacial inception: a study with the UVic Earth system model. *Climate Dynamics*, **21**, 515–537.
- Michalewicz, Z. (1992). *Genetic Algorithms + Data Structures = Evolution Programs*. AI Series. Springer-Verlag, New York.
- Oki, T. and Sud, Y. C. (1998). Design of total runoff integrating pathways (TRIP) - A global river channel network. *Earth Interactions*, **2**.
- Opsteegh, J. D., Haarsma, R. J., Selten, F. M., and Kattenberg, A. (1998). ECBILT a dynamic alternative to mixed boundary conditions in ocean models. *Tellus A*, **50**, 348–367.

- Payne, A. J. (1999). A thermomechanical model of ice flow in West Antarctica. *Climate Dynamics*, **15**(2), 115–125.
- Petoukhov, V., Ganopolski, A., Brovkin, V., Claussen, M., Eliseev, A., Kubatzki, C., and Rahmstorf, S. (2000). CLIMBER-2: A climate system model of intermediate complexity. Part I: Model description and performance for present climate. *Climate Dynamics*, **16**, 1–17.
- Polcher, J., McAvaney, B., Viterbo, P., Gaertner, M.-A., Hahmann, A., Mahfouf, J.-F., Noilhan, J., Phillips, T., Pitman, A., Schlosser, C. A., Schulz, J.-P., Timbal, B., Verseghy, D., and Xue, Y. (1998). A proposal for a general interface between land surface schemes and general circulation models. *Global and Planetary Change*, **19**, 261–276.
- Prange, M., Lohmann, G., and Paul, A. (2003). Influence of vertical mixing on the thermohaline hysteresis: Analyses of an OGCM. *Journal of Physical Oceanography*, **33**(8), 1707–1721.
- Price, A. R., Jiao, Z., Voutchkov, I. I., Lenton, T. M., Williams, G., Lunt, D. J., Marsh, R., Valdes, P. J., Cox, S. J., and the GENIE team (2006a). Collaborative study of GENIEfy Earth system models using scripted database workflows in a Grid-enabled PSE. In S. J. Cox, editor, *Proceedings of the UK e-Science All Hands Meeting 2006*, page submitted, Nottingham, UK. EPSRC.
- Price, A. R., Xue, G., Yool, A., Lunt, D. J., Valdes, P. J., Lenton, T. M., Wason, J. L., Pound, G. E., Cox, S. J., and the GENIE team (2006b). Optimisation of integrated Earth System Model components using Grid-enabled data management and computation. *Concurrency and Computation: Practice and Experience*, page in press.
- Ridgwell, A. J. (2001). *Glacial-interglacial perturbations in the global carbon cycle*. Ph.D. thesis, University of East Anglia.
- Ridgwell, A. J. and Hargreaves, J. (2006). An efficient Earth System Model of Global biogeochemical cycling: Regulation of atmospheric CO₂ by deep-sea sediments. *Global Biogeochemical Cycles*, page submitted.
- Ridgwell, A. J., Watson, A. J., and Archer, D. E. (2002). Modelling the response of the oceanic Si inventory to perturbation and consequences for atmospheric CO₂. *Global Biogeochem. Cycles*, **16**(10).
- Ridgwell, A. J., Hargreaves, J. C., Edwards, N. R., Annan, J., Lenton, T. M., Marsh, R., Yool, A., and Watson, A. J. (2006a). Marine geochemical data assimilation in an efficient Earth System Model of global biogeochemical cycling. *Biogeosciences Discussions*, **3**, 1313–1354.
- Ridgwell, A. J., Zondervan, I., Hargreaves, J. C., Bijma, J., and Lenton, T. M. (2006b). Significant long-term increase of fossil fuel co₂ uptake from reduced

- marine calcification. *Biogeosciences Discussions*, **3**, 1763–1780.
- Rohling, E. J., Marsh, R., Wells, N. C., Siddall, M., and Edwards, N. R. (2004). Similar meltwater contributions to glacial sea level changes from Antarctic and northern ice sheets. *Nature*, **430**, 1016–1021.
- Rutt, I. C., Payne, A. J., Lunt, D. J., and Valdes, P. J. (2005). Simulation of seasonal snow-cover and glacial inception with GENIE and GLIMMER. *EOS*, **86**(52 (Fall Meeting Supplement)), Abstract C12A-01.
- Semtner, A. J. (1976). A model for the thermodynamic growth of sea ice in numerical investigations of climate. *Journal of Physical Oceanography*, **6**, 379–389.
- Six, K. D. and Maier-Reimer, E. (1996). Effects of plankton dynamics on seasonal carbon fluxes in an ocean general circulation model. *Global Biogeochem. Cycles*, **10**, 559–583.
- Tiedke, M. (1993). A comprehensive mass flux scheme for cumulus parameterization in large-scale models. *Monthly Weather Review*, **117**, 1779–1800.
- Wang, Z. and Mysak, L. A. (2000). A simple coupled atmosphere-ocean-sea ice-land surface model for climate and paleoclimate studies. *Journal of Climate*, **13**, 1150–1172.
- Wanninkhof, R. (1992). Relationship between wind speed and gas exchange over the ocean. *Journal of Geophysical Research (Oceans)*, **97**(C5), 7373–7382.
- Weaver, A. J., Eby, M., Weibe, E. C., Bitz, C. M., Duffy, P. B., Ewen, T. L., Fanning, A. F., Holland, M. M., MacFadyen, A., Matthews, H. D., Meissner, K. J., Saenko, O., Schmittner, A., Wang, H., and Yoshimori, M. (2001). The UVic earth system climate model: Model description, climatology, and applications to past, present and future climates. *Atmosphere-Ocean*, **39**(4), 361–428.
- Williamson, M. S., Lenton, T. M., Shepherd, J. G., and Edwards, N. R. (2006). An efficient numerical terrestrial scheme (ENTS) for Earth system modelling. *Ecological Modelling*, page in press.
- Wright, D. G. and Stocker, T. F. (1991). A zonally averaged ocean model for the thermohaline circulation. Part 1: Model development and flow dynamics. *Journal of Physical Oceanography*, **21**, 1713–1724.

Table 1: Components currently available in the GENIE framework and their notation

Component	Label	Model	Identifier	Reference
Atmosphere	aa	IGCM	ia	de Forster <i>et al.</i> (2000)
		EMBM	ea	Fanning and Weaver (1996)
		Fixed	fa	
Ocean	oo	GOLDSTEIN	go	Edwards and Marsh (2005)
		Slab	so	de Forster <i>et al.</i> (2000)
		Fixed	fo	
Sea-ice	ss	Multi-option	gs	Herein, Edwards and Marsh (2005)
		Slab	ss	de Forster <i>et al.</i> (2000)
		Fixed	fs	
Land	ll	GENIE-land	gl	Meissner <i>et al.</i> (2003)
		IGCM-land	il	de Forster <i>et al.</i> (2000)
		ENTS	el	Williamson <i>et al.</i> (2006)
		Fixed	fl	
Ice-sheets	ii	GLIMMER	gi	Payne (1999)
		Fixed	fi	
Vegetation	vv	TRIFFID	tv	Cox (1998)
		ENTS	ev	Williamson <i>et al.</i> (2006)
		Fixed	fv	
Biogeochemistry	gg	BIOGEM	bg	Herein, Ridgwell (2001)
Sediments	dd	SEDGEM	sd	Herein, Ridgwell (2001)
Atmospheric tracers	cc	Atchem	ac	Herein

Table 2: Grids and resolutions currently available in the GENIE framework.

Component Model	Resolution Variant	Longitude		Latitude		Vertical	
		Cells	Res.	Cells	Res.	Levels	Topography
Ocean GOLDSTEIN	Baseline	36	10°	36	$\sin(\pi/36)$	8	Yes
	High	72	5°	72	$\sin(\pi/72)$	16	Yes
	IGCM res.	64	5.625°	32	5.625°	8	Yes
	Equal lat.	36	10°	60	3°	8	Yes
Atmosphere 2-D EMBM	Baseline	36	10°	36	$\sin(\pi/36)$	1	Optional
	High	72	5°	72	$\sin(\pi/72)$	1	Optional
	IGCM res.	64	5.625°	32	5.625°	1	Optional
	Equal Lat.	36	10°	60	3°	1	Optional
Atmosphere 3-D IGCM	Baseline (T21)	64	5.625°	32	5.625°	7	Yes

Table 3: Large experiments undertaken with the GENIE framework. Where GOLDSTEIN has been used it was at baseline (36x36x8) resolution, except for some of the experiments described in the main article.

Experiment	Components included									Ensemble size (yr)	Reference
	aa	oo	ss	ll	ii	vv	gg	dd	cc		
Initial	ea	go	gs	-	-	-	-	-	-	2,400,000	Edwards and Marsh (2005)
Bistability 1	ea	go	gs	-	-	-	-	-	-	38,440,000	Marsh <i>et al.</i> (2004)
EnKF 1	ea	go	gs	-	-	-	-	-	-	1,188,000	Annan <i>et al.</i> (2005b)
EnKF 2	ea	go	gs	-	-	-	-	-	-	1,944,000	Hargreaves <i>et al.</i> (2004)
EnKF 3	ea	go	gs	-	-	-	-	-	-	972,000	Hargreaves and Annan (2006)
Krig tuning	ea	go	gs	-	-	-	-	-	-	246,000	Price <i>et al.</i> (2006b)
Meltwater	ea	go	gs	-	-	-	-	-	-	512,000	Rohling <i>et al.</i> (2004)
Emulator 1	ea	go	gs	-	-	-	-	-	-	10,000	Challenor <i>et al.</i> (2006)
Emulator 2	ea	go	gs	-	-	-	-	-	-	10,800	McNeall <i>et al.</i> , in prep.
IAM 1	ea	go	gs	-	-	-	-	-	-	84,600	Drouet <i>et al.</i> (2005)
IAM 2	ea	go	gs	-	-	-	-	-	-	3,955,200	Drouet <i>et al.</i> (2006)
IAM 3	ea	go	gs	-	-	-	-	-	-	106,400	Bahn <i>et al.</i> (2006)
IPCC AR4	ea	go	gs	-	-	-	-	-	-	22,897,890	in prep.
Factorial	ea	go	gs	-	-	-	bg	-	ac	4,895,000	Cameron <i>et al.</i> (2005)
Calcification	ea	go	gs	-	-	-	bg	-	ac	648,000	Ridgwell <i>et al.</i> (2006b)
Acidification	ea	go	gs	-	-	-	bg	-	ac	648,000	Ridgwell <i>et al.</i> (2006a)
Fossil C fate	ea	go	gs	-	-	-	bg	sd	ac	600,000	Ridgwell and Hargreaves (2006)
Millennial	ea	go	gs	el	-	ev	bg	-	ac	262,000	Lenton <i>et al.</i> (2006)
Deglaciation	ea	go	gs	el	fi	fv	-	-	-	109,000	Lunt <i>et al.</i> (2006)
Atl. salinity	ea	go	gs	el	fi	fv	-	-	-	100,000	Marsh <i>et al.</i> (2006)
Glacial cycle	ea	go	gs	el	fi	fv	-	-	-	480,000	Marsh <i>et al.</i> (2006)
EnKF 4	ia	fo	fs	il	-	-	-	-	-	38,400	Annan <i>et al.</i> (2005a)
GA tuning	ia	fo	fs	il	-	-	-	-	-	60,000	Price <i>et al.</i> (2006b)
Inception	ia	so	ss	-	gi	-	-	-	-	70,000	Rutt <i>et al.</i> (2005)
Bistability 2	ia	go	ss	il	-	-	-	-	-	407,000	Main article