

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

Received: 27 July 2006 / Accepted: 6 March 2007
© Springer-Verlag 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). Feedbacks involving both ocean and atmospheric dynamics are found to promote THC bi-stability. THC switch-off leads to increased import of freshwater at the southern boundary of the Atlantic associated with meridional overturning circulation. This is counteracted by decreased freshwater import associated with gyre and diffusive transports. However, these are localised such that the density gradient between North and South is reduced tending to maintain the THC off state. THC switch-off can also generate net atmospheric freshwater input to the Atlantic that tends to maintain the off state. The ocean feedbacks are 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 net oceanic freshwater import due to the

Electronic supplementary material The online version of this article (doi:10.1007/s00382-007-0254-9) contains supplementary material, which is available to authorized users.

T. M. Lenton (✉) · S. Goswami · V. N. Livina
School of Environmental Sciences,
University of East Anglia, Norwich, UK
e-mail: t.lenton@uea.ac.uk

T. M. Lenton · J. G. Shepherd · M. S. Williamson
Tyndall Centre, UK,
URL: <http://www.tyndall.ac.uk>

R. Marsh · Y. Aksenov · J. G. Shepherd · A. Yool
National Oceanography Centre,
University of Southampton, Southampton, UK

A. R. Price · T. Cooper-Chadwick · S. J. Cox · Z. Jiao
Southampton e-Science Centre,
University of Southampton, Southampton, UK

D. J. Lunt · A. J. Payne · I. C. Rutt · P. J. Valdes · G. Williams
School of Geographical Sciences,
University of Bristol, Bristol, UK

J. D. Annan · J. C. Hargreaves
Frontier Research Centre for Global Change,
Yokohama, Japan

N. R. Edwards
Centre for Earth, Planetary,
Space and Astronomical Research (CEPSAR), Earth Sciences,
Open University, Milton Keynes, UK

P. P. Harris
Centre for Ecology and Hydrology, Wallingford, UK

M. S. Williamson
School of Physics and Astronomy,
University of Leeds, Leeds, UK

overturning mode weakens, promoting THC switch-on, but the atmosphere counteracts this by increasing net freshwater input. This increases the extent of THC bi-stability in this version of the model.


1 Introduction

There currently exists a spectrum of Earth system models (ESMs) and climate models (Fig. 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 Fichefet 1999)

and its further development LOVECLIM (Roche et al. 2006), which achieve computational efficiency by using a 3-layer quasi-geostrophic 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 (Fig. 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 year 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 (Fig. 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

Fig. 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



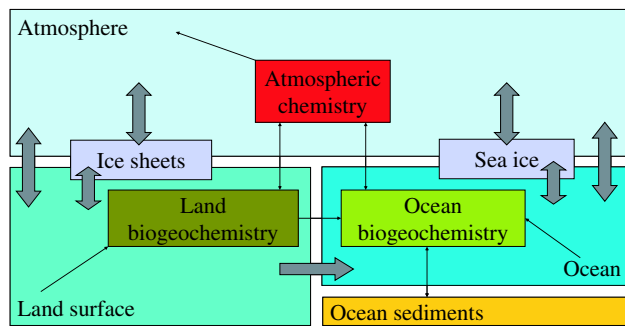


Fig. 2 Major components of an Earth system model. The GENIE framework currently offers at least one option for each of these components, except atmospheric chemistry, which is currently just an atmospheric tracer module without chemical reactions

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 models, 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 CO_2 suggests that this change in boundary condition shifts the model into a region of bi-stability (Stouffer and Manabe 2003). In contrast, under doubling CO_2 there is a reversible weakening, and under quadrupling CO_2 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 Niño events (Schmittner et al. 2003), 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 Sect. 2 we provide a description of the model. Section 3 introduces the novel software and techniques which made the experiment possible. In Sect. 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 Sect. 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 (Fig. 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 can solve 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

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°	–	–

of GOLDSTEIN (Table 1): (1) The initial “baseline” (standard) resolution ($36 \times 36 \times 8$), on a longitude versus $\sin(\text{latitude})$ horizontal grid, with quasi-logarithmic depth intervals (Edwards and Marsh 2005). (2) A higher resolution ($72 \times 72 \times 16$) version, on the same longitude versus $\sin(\text{latitude})$ horizontal grid. (3) An IGCM grid-matched ($64 \times 32 \times 8$) 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 2 m, 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 50 W m^{-2} .

Fig. 3 Alternative surface grids and corresponding bathymetry in the GENIE framework:

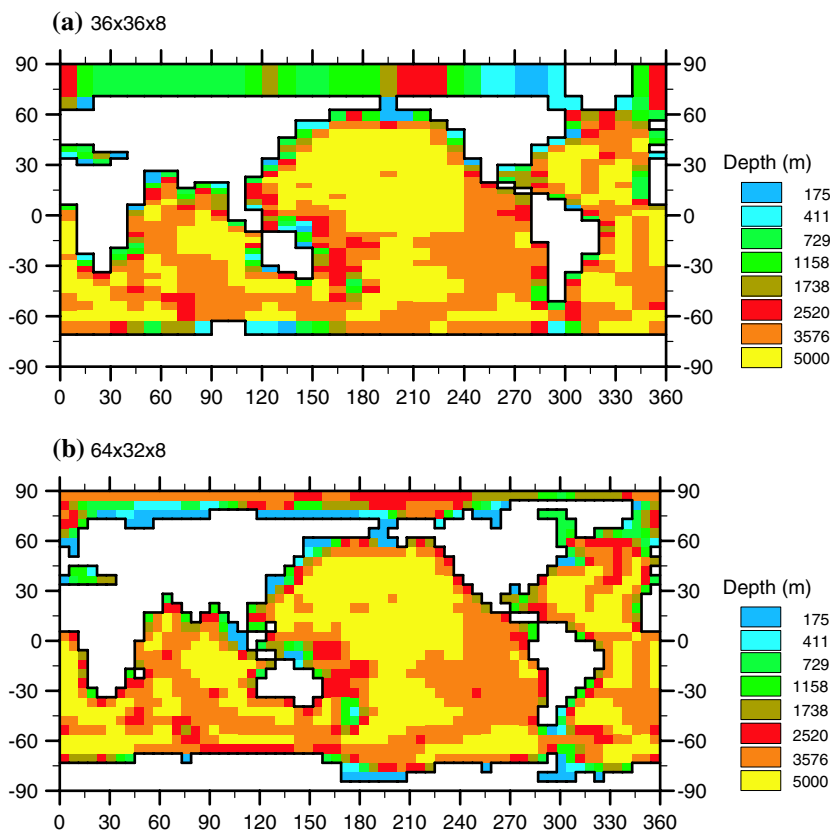
a $36 \times 36 \times 8$ longitude-sine(latitude) i.e. equal area,
b $64 \times 32 \times 8$ longitude-latitude. Note that the $72 \times 72 \times 16$ model used herein shares the $36 \times 36 \times 8$ surface grid and bathymetry, in order to isolate the effects of simply increasing the number of grid cells

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 neighbour. 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 1,000 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 advection, diffusion and thermodynamic sea-ice) (Hargreaves et al. 2004; Edwards and Marsh 2005). Winds from the IGCM are scaled by a factor of 2 (as with prescribed winds from the EMBM) to counteract the excessive 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:

3.1.1 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.

3.1.2 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.

3.1.3 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

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—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 (h)
Minimum precipitation for which there is cloud	ia_pcloudmin	3.013 (mm day^{-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 (h)
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_zrdef	3.092 (μm)
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_b , go_14	$2,000$ (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

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 ($36 \times 36 \times 8$) 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 instead of removing freshwater from the Atlantic and adding it to the Pacific, the atmosphere model generally transports it in the opposite direction. 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

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 (year)	Resolution [binary]	Parameter	Range	Members	Restart state
Spin-ups	1,000	$36 \times 36 \times 8$	f	0–2	21	–
	1,000	$64 \times 32 \times 8$	f	0–2	21	–
	1,000	$72 \times 72 \times 16$	f	0–2	21	–
Restarts	1,000	$36 \times 36 \times 8$	f	0–1	21	$f = 0$
					21	$f = 1$
	1,000	$64 \times 32 \times 8$	f	0–1	21	$f = 0$
					21	$f = 1$
	1,000	$72 \times 72 \times 16$	f	0–1	21	$f = 0$
					21	$f = 1$
1,000	$72 \times 72 \times 16$	f	0–1	21	$f = 0$	
		[Linux]		21	$f = 1$	
2- parameter	2,000	$36 \times 36 \times 8$	f	0–2	11 × 11 = 121	
			log(ia_enhancestable)	–1–1		–

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 here on 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 IGCM 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 ($72 \times 72 \times 16$) 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 of the evaporation (hence latent heat) and sensible heat terms in the IGCM. Increasing it is expected to increase freshwater sources to the atmosphere, which are greatest in 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 2,000 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 and 1.0 Sv freshwater hosing experiments for 100 year as in Stouffer et al. (2006), to compare with their model inter-comparison.

5 Results

The THC exhibits a fairly high degree of variability under the dynamical atmosphere model, with a range in the maximum Atlantic 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 $64 \times 32 \times 8$, ≈ 35 Sv in $72 \times 72 \times 16$ and ≈ 32 Sv in $36 \times 36 \times 8$. In the annual average, the maximum of the Atlantic MOC is located at around 50°N in all resolutions (Fig. 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 $72 \times 72 \times 16$ (Fig. 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 as a function of Atlantic–Pacific freshwater flux correction, for all three ocean resolutions (Fig. 5). The amount of flux correction required to get the Atlantic MOC ‘on’ is noticeably less for the longitude-latitude surface grid $64 \times 32 \times 8$ relative to the other two. The transition is in a similar place for the $36 \times 36 \times 8$ and $72 \times 72 \times 16$ 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

Fig. 4 End-of-run annual average (a, c, e) Atlantic MOC and (b, d, f) Pacific MOC (both in Sv) after 2,000 years at the default flux correction, for (a, b) $36 \times 36 \times 8$ (c, d) $72 \times 72 \times 16$ (Linux binary) (e, f) $64 \times 32 \times 8$






Fig. 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” $36 \times 36 \times 8$, (dotted line) higher resolution $72 \times 72 \times 16$, (dashed line) IGCM-grid matched $64 \times 32 \times 8$. The IGCM atmosphere resolution (T21) is the same in all cases. Points are averages over the last 50 years of 1,000 year runs because in these initial ensembles output was restricted to 50 of each 100 years

the 2 parameters (Fig. 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




Fig. 6 Maximum Atlantic Meridional Overturning Circulation (MOC) strength (Sv) in the $36 \times 36 \times 8$ 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 2,000 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

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 here on 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 (Fig. 7). The region of bi-stability is noticeably wider for the longitude–latitude grid $64 \times 32 \times 8$ and shifted to lower values of flux correction. The bi-stable region is narrowest for $72 \times 72 \times 16$ and the transitions occur at the largest values of flux correction. The baseline $36 \times 36 \times 8$ 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 $64 \times 32 \times 8$, 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 ($f = 0.2–0.3$) (Fig. 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 Fig. 7). Illustrative cases (Fig. 8) are $f = 0.8$ of the default flux correction in resolution $36 \times 36 \times 8$,





Fig. 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” $36 \times 36 \times 8$, (dotted line) higher resolution $72 \times 72 \times 16$, (dashed line) IGCM-grid matched $64 \times 32 \times 8$. 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 1,000 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 $36 \times 36 \times 8$

Fig. 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) $36 \times 36 \times 8$ at $f = 0.8$ of default flux correction (c, d) $72 \times 72 \times 16$ (Linux binary) at $f = 0.85$ (e, f) $64 \times 32 \times 8$ at $f = 0.55$



$f = 0.85$ 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 $36 \times 36 \times 8$, which is consistent with higher resolution in straits allowing stronger through-flow, e.g. of the ACC. 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 the overturning mode diagnosed in the model. We diagnose the transport of freshwater (or conversely, salt) by the Atlantic MOC, and also by gyres and diffusion. The sum of these three transport modes is the total net oceanic freshwater transport.

5.3.1 Ocean feedback

Freshwater transport by the Atlantic MOC, gyres and diffusion varies significantly as a function of latitude, and changes significantly when the THC switches state (Fig. 9). For our illustrative runs, in the ‘on’ state, in all resolutions the self-sustaining nature of the THC can be seen in the North Atlantic, where the MOC transports freshwater




Fig. 9 Meridional profiles of freshwater transport versus latitude (positive northwards), due to the Meridional Overturning Circulation, gyre circulation, diffusion and their total, for THC on and THC off restarts, and difference (off–on). **a** $36 \times 36 \times 8$ at $f = 0.8$ of default

flux correction, **b** $72 \times 72 \times 16$ at $f = 0.85$, **c** $64 \times 32 \times 8$ at $f = 0.55$. Individual annual averages are shown for a decade at the end of the restarted simulations

southwards. At the southern boundary of the Atlantic, in $36 \times 36 \times 8$ and $72 \times 72 \times 16$, there is little MOC import or export of freshwater in the ‘on’ state, whereas with the THC switched off, there is a pronounced MOC import of freshwater. By contrast, in $64 \times 32 \times 8$, in the ‘on’ state, the MOC imports freshwater at the Atlantic southern boundary, and there is little change when the THC switches off. A shift to increased MOC import of freshwater should tend to maintain the ‘off’ state in the illustrative $36 \times 36 \times 8$ and $72 \times 72 \times 16$ runs, but is not apparent in the illustrative $64 \times 32 \times 8$ run.

For the Atlantic basin to maintain freshwater balance at steady state, any changes in the MOC freshwater transport at the southern boundary (between ‘on’ and ‘off’ states) must be balanced by counteracting changes in other components of the freshwater balance. These could be in gyre transport, diffusive transport, net precipitation minus evaporation plus runoff, transport across the Bering Strait, or some combination of these. In all our simulations we find that changes in freshwater transport through Bering Strait are very small (maximum change ~ 0.01 Sv) and can be neglected. Changes in gyre and diffusive transport (Fig. 9) generally over-compensate for changes in MOC transport, with the net result of somewhat reduced total freshwater import in the ‘off’ state in all three resolutions. If only the net freshwater import were relevant, this should amount to a weak negative feedback tending to counteract THC collapse. Importantly, however, when the THC switches off, the net effect is to remove more freshwater from the South Atlantic than the North Atlantic, thus reducing the density gradient between North and South and inhibiting recovery of the THC, which instead suggests a positive feedback.

Examining the net import or export of freshwater due to various components at the Atlantic southern boundary as a function of default flux correction (Fig. 10) indicates that over much of the bi-stable region, in each resolution, changes in freshwater import/export due to the MOC act to import relatively more freshwater to the Atlantic in the ‘off’ state, while gyre and diffusive freshwaters transports over-compensate this, leading to a net reduction in import. Thus the positive feedback on the MOC required to explain the bi-stable regimes in our model cannot be explained by net changes in freshwater exchange at the southern boundary of the Atlantic. The effect on the density gradient appears to be more important, while an alternative explanation, which cannot be ruled out, is that changes in heat transport have a compensating effect.

5.3.2 Atmosphere feedback

To fully explain what is going on, especially in the transition from bi-stability to a mono-stable ‘on’ state, we must

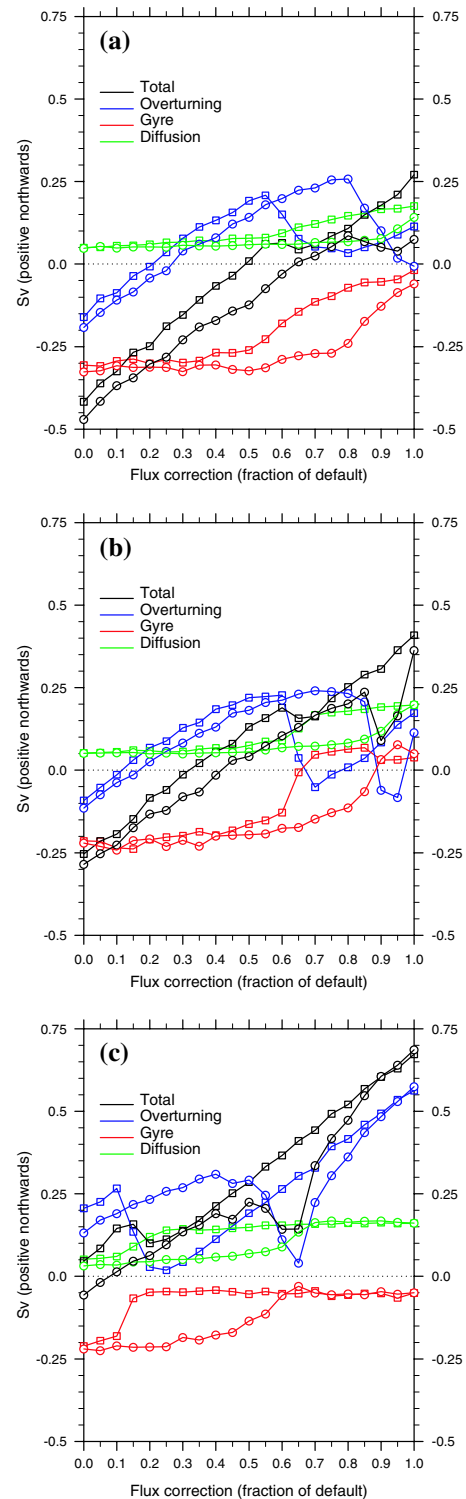


Fig. 10 Freshwater transport by the Meridional Overturning Circulation, gyre circulation, diffusion and their total at the southern boundary of the Atlantic (32°S) as a function of default flux correction, for restarts from THC on (squares) and THC off (circles) states, in different resolutions: **a** $36 \times 36 \times 8$, **b** $72 \times 72 \times 16$, **c** $64 \times 32 \times 8$. Calculated as a decadal average at the end of the restarted simulations

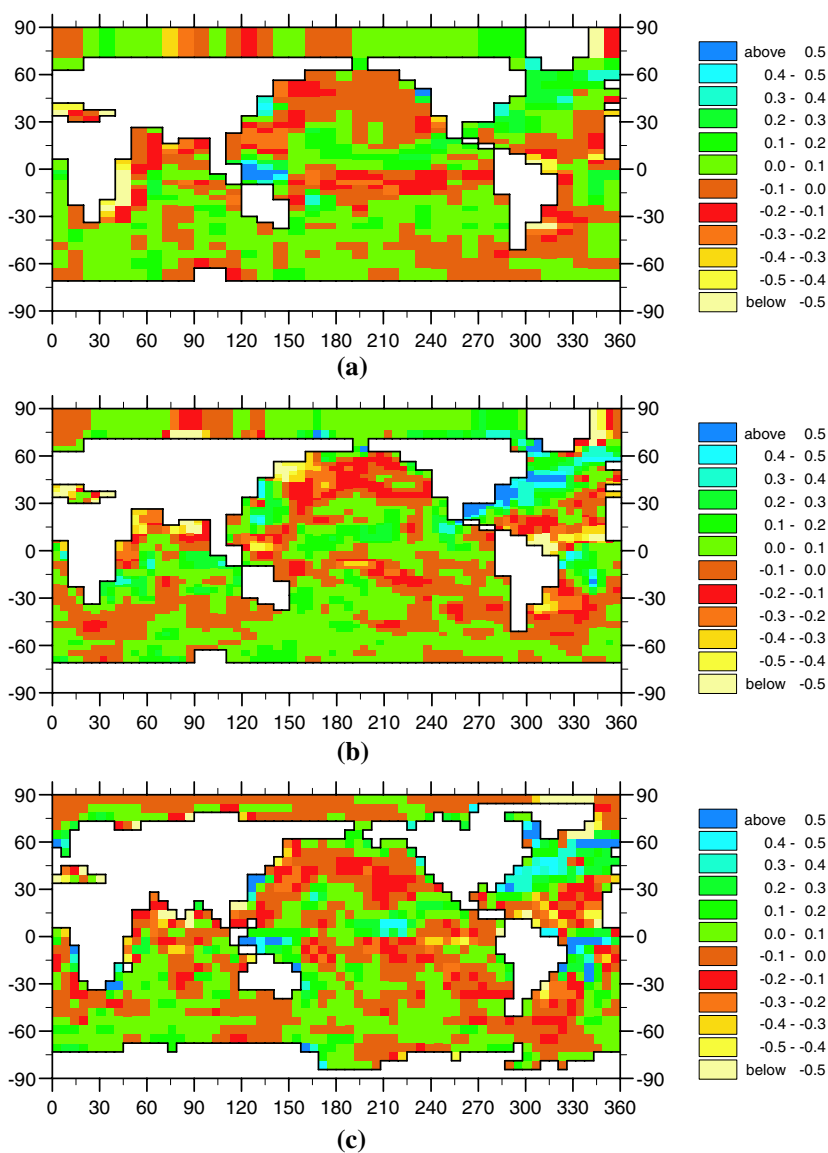
also consider atmospheric feedbacks. The change in net freshwater input (precipitation minus evaporation plus runoff) to the Atlantic from the atmosphere when the THC switches off shows a broadly similar pattern in our illustrative runs for all three resolutions (Fig. 11). There is a general wetting along 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 $36 \times 36 \times 8$.

Considering the net freshwater flux to the Atlantic as a function of latitude (Fig. 12), the changes due to THC

switch-off are small compared to the absolute flux in $36 \times 36 \times 8$ and $72 \times 72 \times 16$, but more pronounced in $64 \times 32 \times 8$. A wetting of the extra-tropical North Atlantic is apparent in all resolutions, and a southward shift of the ITCZ is most apparent in $64 \times 32 \times 8$.

When integrating the net surface freshwater flux over the entire Atlantic basin, the illustrative cases all show an increase due to THC switch-off. However, the change due to THC switch-off is modest compared to the absolute magnitude of the flux, and both the sign and magnitude of the change varies considerably as a function of Atlantic–Pacific flux correction (Fig. 13). If there is multi-decadal variability in the total flux (not examined) then part of the variability in the effect of THC switch-off may be due to averaging over 10 years. The atmospheric feedback (Fig. 13b) is most pronounced in $64 \times 32 \times 8$, with a sig-

Fig. 11 Maps of net surface freshwater flux ($P - E + R$) difference due to THC switch off (off–on) in $m/year$ (positive downwards): **a** $36 \times 36 \times 8$ at $f = 0.8$ of default flux correction, **b** $72 \times 72 \times 16$ at $f = 0.85$, **c** $64 \times 32 \times 8$ at $f = 0.55$. Calculated as a decadal average at the end of the restarted simulations



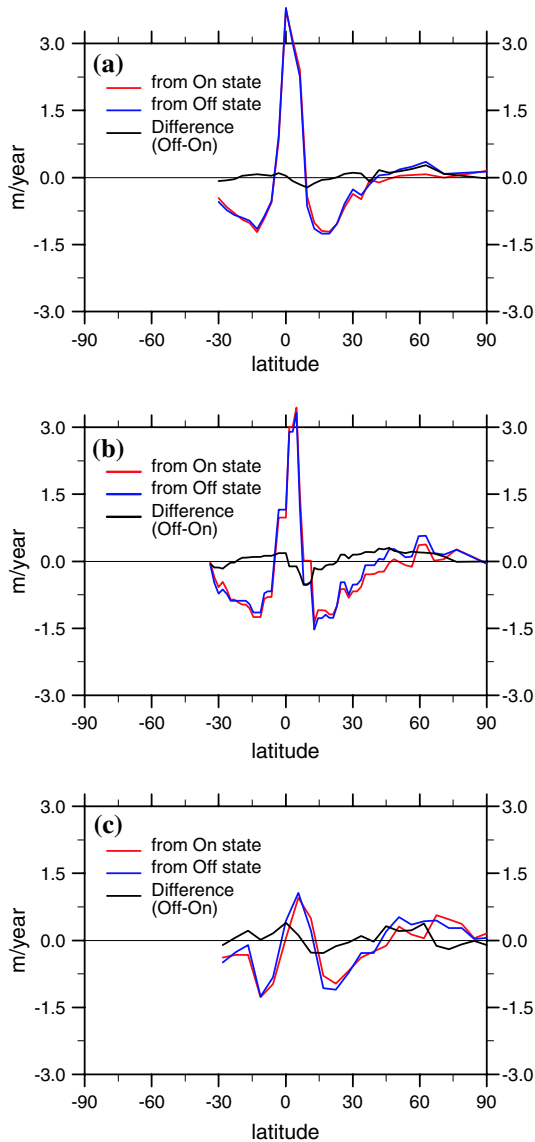


Fig. 12 Meridional profiles of net Atlantic surface freshwater flux ($P - E + R$) per latitude strip in m year^{-1} (positive downwards), for THC on and THC off restarts and difference (off-on). **a** $36 \times 36 \times 8$ at $f = 0.8$ of default flux correction, **b** $72 \times 72 \times 16$ at $f = 0.85$, **c** $64 \times 32 \times 8$ at $f = 0.55$. Calculated as a decadal average at the end of the restarted simulations

nificant increase in freshwater input in going from the on-initialized to off-initialized states for $f = 0.55-0.8$. This corresponds with the righthand section of the bi-stable region and the transition toward a mono-stable ‘on’ state (Fig. 7) where the freshwater input will tend to inhibit recovery of the THC ‘off’ state. In $72 \times 72 \times 16$ 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 $36 \times 36 \times 8$ it is hard to establish any qualitative relationship.

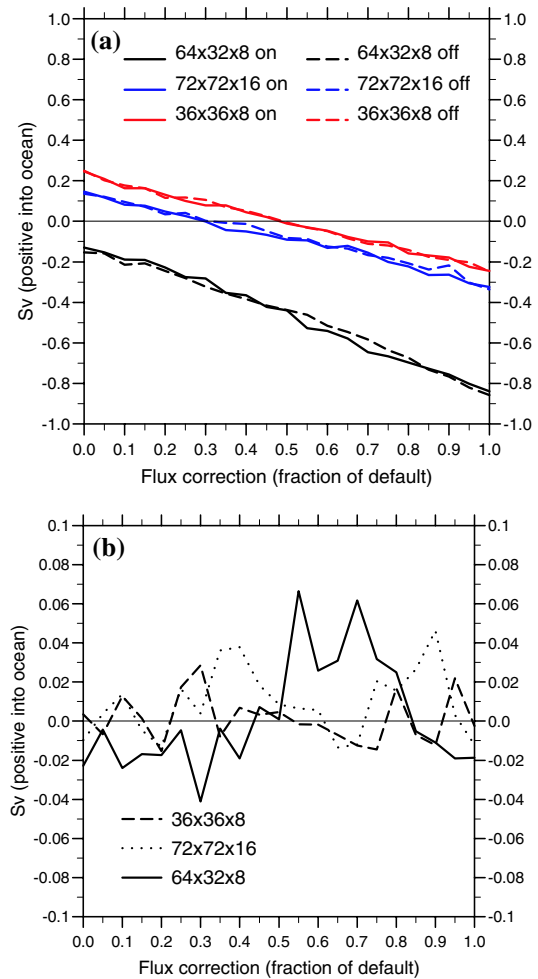


Fig. 13 a Net Atlantic surface freshwater flux ($P - E + R$), and **b** difference (off-on) due to THC switch off, in different resolutions, as a function of default flux correction. Calculated as a decadal average at the end of the restarted simulations

5.3.3 Atlantic freshwater balance

Combining the total freshwater transport by the ocean across the southern boundary of the Atlantic (Fig. 10) and the net surface freshwater flux integrated across the Atlantic basin (Fig. 13a) gives us the two major components of the Atlantic freshwater budget. For each model resolution, an approximately linear relationship can be seen (Fig. 10) between Atlantic-to-Pacific flux correction and total freshwater transport by the ocean across the southern boundary of the Atlantic, for the initial spun-up state of the model. This is to be expected as the flux correction represents a net removal of freshwater from the Atlantic (of default size 0.79 Sv) that must be counterbalanced.

When the Atlantic-to-Pacific flux correction is removed, large net exports of freshwater by the ocean are seen in $36 \times 36 \times 8$ and $72 \times 72 \times 16$ (Fig. 10). These are partly accounted for by net inputs of freshwater to the Atlantic

basin from the atmosphere (Fig. 13a), but there is a discrepancy which may indicate an additional source of freshwater from the flux correction associated with interpolation. In $36 \times 36 \times 8$ atmospheric input was diagnosed as 0.27 Sv, implying an additional source of ≈ 0.2 Sv. In $72 \times 72 \times 16$, atmospheric input is ≈ 0.15 Sv, implying an additional source of ≈ 0.1 Sv. In $64 \times 32 \times 8$, which requires no interpolation and has no corresponding flux correction, there is little net addition or removal of freshwater by the ocean, yet there is a net removal of ≈ 0.15 Sv by the atmosphere, implying a corresponding source that cannot be due to interpolation. Transport across the Bering Strait cannot explain the discrepancies because it is small and out of the Atlantic in all three resolutions; -0.030 Sv in $36 \times 36 \times 8$, -0.034 Sv in $72 \times 72 \times 16$, and -0.010 Sv in $64 \times 32 \times 8$.

According to NCEP reanalysis data, the atmosphere removes 0.52 Sv freshwater from the Atlantic basin, which was added to the diagnosed model flux to define the default flux correction of 0.79 Sv in $36 \times 36 \times 8$. When this default value is applied, the $36 \times 36 \times 8$ model then adjusts such that the atmosphere achieves only ≈ 0.25 Sv removal (Fig. 13a). The corresponding THC ‘on’ state imports a roughly counterbalancing flux (Fig. 10a). In $72 \times 72 \times 16$, under the default flux correction, the atmosphere removes ≈ 0.32 Sv and the ocean adds ≈ 0.4 Sv in the THC ‘on’ state (Fig. 10b). In $64 \times 32 \times 8$ at the default flux correction, the atmosphere removes ≈ 0.85 Sv and the ocean imports ≈ 0.7 Sv. There is often a deviation of total oceanic freshwater transport from a straight line (Fig. 10) for ‘off’ states switching ‘on’ and ‘on’ states switching ‘off’. This suggests that these runs are still adjusting towards freshwater balance after 1000 years, as the other components of the freshwater budget are too small to account for the discrepancy.

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 (Fig. 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 $36 \times 36 \times 8$ (Fig. 6). A value of ≈ 18 Sv corresponds to $f = 0.65$ of the default flux correction in $72 \times 72 \times 16$, $f = 0.6$ – 0.65 in $36 \times 36 \times 8$, and $f = 0.2$ in $64 \times 32 \times 8$. 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 $36 \times 36 \times 8$, $f = 0.7$ in $72 \times 72 \times 16$, and $f = 0.25$ in $64 \times 32 \times 8$.

The longitude–latitude $64 \times 32 \times 8$ grid needs the least flux correction to get a THC ‘on’ state and has the widest region of bi-stability. Hence we focus on this version 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 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 to 33 Sv. After 1,000 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 years, 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 years, 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 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).

Figure 7 shows alongside the GENIE-2 results, 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). 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). This is partly because NCEP reanalysis implies a 0.2 Sv larger net atmospheric removal of freshwater from the Atlantic (0.52 Sv in total) and partly because with the baseline $36 \times 36 \times 8$ ocean resolution, the IGCM transports freshwater in the wrong direction adding 0.27 Sv to the Atlantic. The EMBM uses NCEP-derived wind fields minimising associated errors, whereas the IGCM generates its own winds. 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, all of which 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 $64 \times 32 \times 8$ 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 reveal the role of oceanic and atmospheric feedbacks in THC bi-stability, through changes in the Atlantic freshwater budget. 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 by the MOC, which de Vries and Weber (2005) have argued acts to maintain the ‘off’ state (Fig. 10). For their argument to hold, this positive feedback must somehow dominate over counterbalancing negative feedbacks from gyre and diffusive freshwater transport at the southern boundary of the Atlantic. We surmise this is either because of undiagnosed competing effects on heat transport, or because the net freshwater or buoyancy export is concentrated in the South Atlantic: when the THC switches off, the net effect of changes in the components of freshwater transport is to reduce the density gradient between North and South, inhibiting recovery of the THC (a positive feedback). In most of our model experiments, the MOC imports freshwater to the Atlantic (exports salt) because it is acting to counterbalance a net removal of freshwater by the imposed Atlantic-Pacific flux correction. Only in some experiments with the $72 \times 72 \times 16$ model does the MOC export fresh

water at the southern boundary, favouring the THC ‘on’ state, while importing fresh water in the corresponding ‘off’ state. Thus our model behaviour is broadly consistent with findings of de Vries and Weber (2005), but the sign of MOC freshwater transport at the southern boundary of the Atlantic is clearly not the single determinant of THC bi-stability in GENIE-2.

To explain the extensive bi-stability in our model, atmosphere feedbacks must also be considered. In particular, when increasing the Atlantic-to-Pacific freshwater flux towards the point where the bi-stable ‘off’ state starts to recover, 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’ (Fig. 13). The ocean and atmosphere mechanisms can be seen as counteracting one another in this regime, with the atmospheric feedback extending the region of bi-stability towards higher values of Atlantic-to-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 (Fig. 7).

The differences in the 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. In $64 \times 32 \times 8$ there is a greater northward transport of high salinity water in the North Atlantic, by the MOC itself (Fig. 9). The $64 \times 32 \times 8$ model has the widest region of bi-stability perhaps because it has the widest region over which freshwater transport by the MOC provides positive feedback, and where the MOC switches to negative feedback this is counteracted by the strongest positive feedback from the atmosphere. The $36 \times 36 \times 8$ model has a slightly wider region of bi-stability than the $72 \times 72 \times 16$ model, corresponding to a wider range over which the overturning 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 $36 \times 36 \times 8$ to $72 \times 72 \times 16$ is to sharpen up the boundaries of the bi-stable regime, due to a stronger self-sustaining MOC feedback for the ‘on’ and ‘off’ states (Fig. 10).

The differences in the position of the bi-stable regimes as a function of Atlantic-to-Pacific freshwater flux correction is due to differences in the net freshwater exchange between the Atlantic and the atmosphere in the different resolutions, with an additional unexplained factor that may be partly due to the presence or absence of interpolation,

which brings an additional flux correction. The $64 \times 32 \times 8$ model needs the least Atlantic-to-Pacific flux correction because it already has a net removal of ≈ 0.15 Sv freshwater by the atmosphere (although this magnitude is too small) and it avoids interpolation. In the $36 \times 36 \times 8$ model, the atmosphere adds 0.27 Sv and there is an additional ≈ 0.2 Sv source, potentially from interpolation. In the $72 \times 72 \times 16$ model, the atmosphere source of ≈ 0.15 Sv is smaller, but again interpolation may add freshwater. Increasing the resolution from $36 \times 36 \times 8$ to $72 \times 72 \times 16$ 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 greatly alter the required flux correction, but it may contribute to the larger changes in salt import/export with changes in the THC state in $72 \times 72 \times 16$.

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. Thus it cannot be claimed purely on the basis of observing a reversible THC collapse in limited AOGCM runs, that there is no bi-stability in a given model. In fact, all that can be inferred is that the initial state of the 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 (Fig. 14). Versions of the $64 \times 32 \times 8$ 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 $\approx 1,000$ years. The decrease of 10–20 Sv seen in the ensemble members that collapse exceeds that in any of the models compared by Stouffer et al. (2006). However, this result should be treated with caution, as in these ensemble members the North Atlantic deep water fails to reach the Southern Ocean. Our other

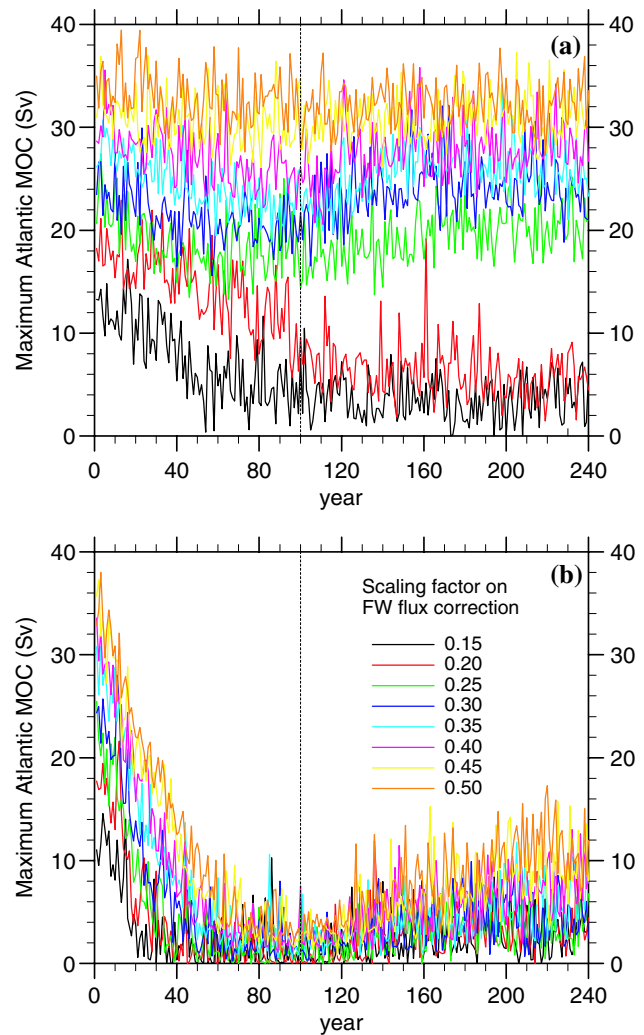


Fig. 14 Freshwater hosing experiments **a** 0.1 Sv for 100 year, **b** 1.0 Sv for 100 year, for variants of GENIE-2 with the $64 \times 32 \times 8$ grid and differing fractions of the default flux correction. All ensemble members are within the bi-stable region for the THC

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). In some other models, a net drying of the tropical North Atlantic creates a salinity anomaly near the Gulf of Mexico that propagates up the North Atlantic and aids THC recovery (Thorpe et al. 2001; Yin et al. 2006). This is somewhat similar to the anomaly observed in El Niño 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 Fig. 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 $64 \times 32 \times 8$ 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. In our model, atmospheric feedback only dominates over ocean feedbacks in a restricted region of parameter space.

Acknowledgments The UK Natural Environment Research Council funded GENIE (NER/T/S/2002/00217) and funds the ongoing GENIEfy (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.

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 single-processor experiments running continuously throughout the time, highlighting the power of the Grid to enable ensemble studies with Earth system models.

Typically, a 1,000 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 five 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 Fig. 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 five nodes of the UK National Grid Service (NGS), three institutional High Performance Computing (HPC) clusters and one large university Condor pool. The database system enabled users to specify how much work was submitted, for example, short 5–10 h 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 Fig. 16. Periods of relative inactivity reflect development work on the system and transitions between experiments

Fig. 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






Fig. 16 Daily resource usage for the 12 ensemble studies

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 five 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 (Fig. 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 (Fig. 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 three 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.

Fig. 17 Distributions of a client job submissions and b resource usage measured as total number of timesteps performed



References

- Abramson D, Buyya R, Giddy J (2002) A computational economy for Grid computing and its implementation in the Nimrod-G resource broker. *Future Generation Comput Syst* 18(8):1061–1074
- Allcock B, Bester J, Bresnahan J, Chervenak AL, Foster I, Kesselman C, Meder S, Nefedova V, Quesnal D, Tuecke S (2002) Data management and transfer in high performance computational Grid environments. *Parallel Comput J* 28(5):749–771
- Betts AK (1986) A new convective adjustment scheme. 1. observational and theoretical basis. *Q J R Meteorol Soc* 112:667–691
- Claussen M, Mysak LA, Weaver AJ, Crucifix M, Fichefet T, Loutre MF, Weber SL, Alcamo J, Alexeev VA, Berger A, Calov R, Ganopolski A, Goosse H, Lohmann G, Lunkeit F, Mokhov I, Petoukhov V, Stone P, Wang Z (2002) Earth system models of intermediate complexity: closing the gap in the spectrum of climate system models. *Clim Dyn* 18:579–586
- Czajkowski K, Foster I, Karonis N, Kesselman C, Martin S, Smith W, Tuecke S (1998) A resource management architecture for meta-computing systems. In: *Proceedings of IPPS/SPDP '98 workshop on job scheduling strategies for parallel processing*, pp 62–82
- de Vries P, Weber SL (2005) The atlantic freshwater budget as a diagnostic for the existence of a stable shut-down of the meridional overturning circulation. *Geophys Res Lett* 32:L09606
- Edwards NR, Marsh R (2005) Uncertainties due to transport-parameter sensitivity in an efficient 3-D ocean-climate model. *Clim Dyn* 24:415–433
- Eres MH, Pound GE, Jiao Z, Wason JL, Xu F, Keane AJ, Cox SJ (2005) Implementation and utilisation of a grid-enabled problem solving environment in Matlab. *Future Generation Comput Syst* 21(6):920–929
- Fanning AG, Weaver AJ (1996) An atmospheric energy-moisture model: climatology, interpentadal climate change and coupling to an ocean general circulation model. *J Geophys Res* 101:15111–15128
- de Forster PM, Blackburn M, Glover R, Shine KP (2000) An examination of climate sensitivity for idealised climate change experiments in an intermediate general circulation model. *Clim Dyn* 16:833–849
- Foster I, Kesselman C (1999) Globus: a toolkit-based Grid architecture. In: Foster I, Kesselman C (eds) *The grid: blueprint for a new computing infrastructure*. Morgan Kaufman Publishers, Inc., San Francisco, pp 259–278
- Foster I, Kesselman C, Tuecke S (2001) The anatomy of the grid: enabling scalable virtual organizations. *Int J High Perform Comput Appl* 15(3):200–222
- Fraedrich K, Jansen H, Kirk E, Luksch U, Lunkeit F (2005a) The planet simulator: towards a user friendly model. *Meteorol Z* 14(3):299–304
- Fraedrich K, Kirk E, Luksch U, Lunkeit F (2005b) The Portable University Model of the Atmosphere (PUMA): storm track dynamics and low-frequency variability. *Meteorol Z* 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. *J Geophys Res (Oceans)* 108(C7):3213
- Goosse H, Fichefet T (1999) Importance of ice-ocean interactions for the global ocean circulation: a model study. *J Geophys Res (Oceans)* 104(C10):23337–23355
- Griffies SM (1998) The Gent-McWilliams skew flux. *J Phys Oceanogr* 28:831–841
- Hargreaves JC, Annan JD, Edwards NR, Marsh R (2004) An efficient climate forecasting method using an intermediate complexity Earth System Model and the ensemble Kalman filter. *Clim Dyn* 23:745–760
- Hey T, Trefethen AE (2002) The UK e-Science core programme and the grid. *Future Generation Comput Syst* 18(8):1017–1031
- Hoskins BJ, Simmons AJ (1975) A multi-layer spectral model and the semi-implicit method. *Q J R Meteorol Soc* 101:637–655
- Jayne SR, Marotzke J (2001) The dynamics of ocean heat transport variability. *Rev Geophys* 39:385–411
- Keane AJ (2003) The OPTIONS design exploration system: reference manual and user guide. <http://www.soton.ac.uk/~ajk/options.ps>
- von Laszewski G, Foster I, Gawor J, Lane P (2001) A Java commodity grid kit. *Concurr Comput Pract Exp* 13(8–9):643–662
- Latif M, Roeckner E, Mikolajewicz U, Voss R (2000) Tropical stabilization of the thermohaline circulation in a greenhouse warming simulation. *J Clim* 13:1809–1813
- Lenton TM, Williamson MS, Edwards NR, Marsh R, Ridgwell AJ, Price AR, Shepherd JG, Cox SJ, The GENIE team (2006) Millennial timescale carbon cycle and climate change in an efficient Earth system model. *Clim Dyn* 26:687–711
- Louis JF (1986) Forecast model research manual 3: physical parameterisation. European Centre for Medium Range Weather Forecasts
- Marsh R, Yool A, Lenton TM, Gulamali MY, Edwards NR, Shepherd JG, Krznanic M, Newhouse S, Cox SJ (2004) Bistability of the thermohaline circulation identified through comprehensive 2-parameter sweeps of an efficient climate model. *Clim Dyn* 23:761–777
- Morcrette JJ (1990) Impact of changes to the radiative transfer parameterisations plus cloud optical properties in the ECMWF model. *Monthly Weather Rev* 118:847–873
- Murphy JM, Sexton DMH, Barnett DN, Jones GS, Webb MJ, Collins M, Stainforth DA (2004) Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature* 430:768–772
- Opsteegh JD, Haarsma RJ, Selten FM, 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, Rahmstorf S (2000) CLIMBER-2: a climate system model of intermediate complexity. Part I: Model description and performance for present climate. *Clim Dyn* 16:1–17
- Pope V, Gallani M, Rowntree P, Stratton R (2000) The impact of new physical parameterisations in the Hadley Centre climate model: HadAM3. *Clim Dyn* 16:123–146
- Price AR, Xue G, Yool A, Lunt DJ, Valdes PJ, Lenton TM, Wason JL, Pound GE, Cox SJ, the GENIE team (2006) Optimisation of integrated Earth system model components using grid-enabled data management and computation. *Concurr Comput Pract Exp* (in press)
- Rahmstorf S, Crucifix M, Ganopolski A, Goosse H, Kamenkovich I, Knutti R, Lohmann G, Marsh R, Mysak LA, Wang Z, Weaver AJ (2005) Thermohaline circulation hysteresis: a model inter-comparison. *Geophys Res Lett* 32:L23605
- Roche DM, Dokken TM, Goosse H, Renssen H, Weber SL (2006) Climate of the last glacial maximum: sensitivity studies and model-data comparison with the LOVECLIM coupled model. *Clim Past Discussions* 2:1105–1153
- Schlesinger ME, Yin J, Yohe GW, Andronova NG, Malyshev S, Li B (2006) Assessing the risk of a collapse of the Atlantic thermohaline circulation. In: Schellnhuber HJ, Cramer W, Nakicenovic N, Wigley T, Yohe G (eds) *Avoiding dangerous climate change*. Cambridge University Press, Cambridge, pp 37–47
- Schmittner A, Appenzeller C, Stocker TF (2000) Enhanced Atlantic freshwater export during El Niño. *Geophys Res Lett* 27:1163–1166

- Slingo JM (1987) The development and verification of a cloud prediction scheme for the ECMWF model. *Q J R Meteorol Soc* 113:899–927
- Smethie WM, Fine RA (2001) Rates of North Atlantic Deep Water formation calculated from chlorofluorocarbon inventories. *Deep Sea Res I* 48:189–215
- Stommel H (1961) Thermohaline convection with two stable regimes of flow. *Tellus* 13:224–230
- Stouffer RJ, Manabe S (2003) Equilibrium response of thermohaline circulation to large changes in atmospheric CO₂ concentration. *Clim Dyn* 20:759–773
- Stouffer RJ, Yin J, Gregory JM, Dixon KW, Spelman MJ, Hurlin W, Weaver AJ, Eby M, Flato GM, Hasumi H, Hu A, Jungclaus JH, Kamenkovich IV, Levermann A, Montoya M, Murakami S, Nawrath S, Oka A, Peltier WR, Robitaille DY, Sokolov A, Vettoretti G, Weber SL (2006) Investigating the causes of the response of the thermohaline circulation to past and future climate changes. *J Clim* 19:1365–1387
- Talley LD, Reid JL, Robbins PE (2003) Data-based meridional overturning streamfunctions for the global ocean. *J Clim* 16:3213–3226
- Thain D, Tannenbaum T, Livny M (2005) Distributed computing in practice: the Condor experience. *Concurr Comput Pract Exp* 17(2–4):323–356
- Thorpe RB, Gregory JM, Johns TC, Wood RA, Mitchell JFB (2001) Mechanisms determining the Atlantic thermohaline circulation response to greenhouse gas forcing in a non-flux-adjusted coupled climate model. *J Clim* 14:3102–3116
- Vellinga M, Wood RA (2002) Global climatic impacts of a collapse of the Atlantic thermohaline circulation. *Clim Change* 54:251–267
- Wang H, Birchfield GE (1992) An energy-salinity balance climate model: Water vapor transport as a cause of changes in the global thermohaline circulation. *J Geophys Res* 97(C2):2335–2346
- Welch V, Siebenlist F, Foster I, Bresnahan J, Czajkowski K, Gawor J, Kesselman C, Meder S, Pearlman L, Tuecke S (2003) Security for Grid services. In: Proceedings of the 12th IEEE international symposium on high performance distributed computing. The Institute of Electrical and Electronics Engineers, Inc., Seattle, pp 48–57
- Wright DG, Stocker TF (1991) A zonally averaged ocean model for the thermohaline circulation. Part 1: Model development and flow dynamics. *J Phys Oceanogr* 21:1713–1724
- YarKhan A, Seymour K, Sagi K, Shi Z, Dongarra J (2006) Recent developments in Gridsolve. *Int J High Perform Comput Appl* 20(1):131–141
- Yin J, Schlesinger ME, Andronova NG, Malyshev S, Li B (2006) Is a shutdown of the thermohaline circulation irreversible? *J Geophys Res (Oceans)* 111:D12104