Proposed RAPID UK model THC intercomparison project

1 Background and motivation

Human activities may significantly increase the risk of rapid climate change in the coming decades by inducing a possible slowdown of the Atlantic Ocean thermohaline circulation (THC), a fundamental controlling factor on the European climate owing to its crucial role in transporting northward the excess of heat imparted in the equatorial regions. The ultimate objective of the NERC funded RAPID programme is to quantify the probability and magnitude of this potential future rapid climate change, and the uncertainties in these estimates.

The main tools for making projections of anthropogenic climate change in the coming century are coupled atmosphere-ocean general circulation models (AOGCMs), as well as Earth Models of Intermediate Complexity (EMICs). Such models, however, exhibit widely divergent responses to a CO2 increase, which may range from no response of the THC to an almost complete disappearance. Such differences in the THC response imply large uncertainties in the character and magnitude of climate change over coming decades, especially in the north Atlantic and Europe.

To understand how to quantify and possibly reduce uncertainties in model predictions, one may usefully distinguish between three distinct categories: 1) Sensitivity to poorly known initial conditions, due to the inherently chaotic nature of the climate system; 2) Uncertainties in the forcing scenarios; 3) Sensitivity to the parameter values and functional form of parametrisations of subgridscale processes, as well as structural uncertainties associated with model bias, resulting from models being an imperfect representation of reality.

Among these, only parametric and structural uncertainties (3) can be reduced over time by continuous physical refinements, while the best that can be achieved for (1) and (2) is to quantify their importance. RAPID round 2 has funded a project on "understanding uncertainty in simulations of THC-related rapid climate change" (NE/C509366/1) with the aim of addressing (3), by comparing climate models of a range of complexity and design, run under similar scenarios, to identify and quantify the physical reasons for different predictions among them, and thereby to get insights into how to reduce parametric and structural uncertainties on predictions of climate change.

Part of the funding is to coordinate a UK model THC intercomparison project. The present document sets out a plan for this and outlines the resources requested. This proposed UK project is complementary and provides added value to the existing CMIP coordinated THC experiments, which have already been carried out by several international groups. This CMIP programme not being specially funded, however, it has been short on effort, only allowing for a cursory analysis of the results. The purpose of the UK project will be to construct its own hierarchy of coupled ocean/atmosphere models, and to conduct a much more thorough analysis of their results with the above objectives.

To achieve its objectives, the present project brings together the expertise of a critical mass of UK experts in ocean modelling, ocean theory, and statistics, which will be combined with the existing complementary expertise of the international participants in the CMIP project.

2 Status of the UK intercomparison project

As a first step toward setting up the UK intercomparison project, Rémi Tailleux was appointed project manager on June 20, 2005, and a workshop was held on June 28, 2005, as part of the annual Rapid meeting in Swansea, to identify potential participants, clarify the scientific objectives of the UK project, and assess the resources needed to run it. Jonathan Gregory, the PI of the project, first described the general context and objectives of the project as set up so far within the context of CMIP, and described some initial results reported in two papers (Gregory et al, GRL 2005, and Stouffer et al, to appear in J. Climate). Interested potential participants then described how their models would be of interest to the project. The models presented were HadCM3 and HadGEM1 (Anne Pardaens and Michael Vellinga, from the Met Office), HiGEM (Len Shaffrey, Reading), CHIME (Alex Megann, NOCS), FAMOUS (Jonathan Gregory, Reading), FORTE (Bablu Sinha, NOCS), and GENIE models, i.e., C-GOLDSTEIN and IGCM-GOLDSTEIN (Neil Edwards). After the workshop, Grant Bigg (Sheffield) expressed interest in the model FRUGAL being included in the UK intercomparison project. It was also decided, in consultation with the Rapid modelling subcommittee, that the model ECBilt-CLIO (KNMI) should be included in the comparison owing to its participation in other RAPID funded projects. The presentations were followed by a general discussion about the experiments to be carried out, the general scientific issues, practicalities and resources required.

From the discussions that resulted from and followed the workshop, while many groups expressed interest in participating in the project on their own time and funding, it appears that there would be need for the following:

- Support for a 2-year postdoc and 3 months' permanent staff support for carrying out experiments for CHIME and FORTE.
- 2/3 months postdoc support to run FRUGAL.
- 1 year staff support for HiGEM from the end of year 2 of the project.
- Support of expert time to investigate traceability within a rigorous statistical framework (which will include performing the runs with the two GENIE models, see details below).
- Support for holding semi-annual workshops to report progress, exchange results, and discuss papers to be written.

More information and justification for support is given in section 5.

3 Scientific issues and strategy of the UK project

3.1 Objectives

The objectives of the project, as stated in the case for support of the UK project, can be refined as follows:

“Physical” Approach Relate differences in model projections of THC-related climate change to the different formulations of the model, and in particular to discover any qualitative differences in behaviour which depend critically on complexity or resolution, through their effects on control state or feedbacks. For instance, irreversible collapse of the THC in response to CO_2 forcing has not been shown in any AOGCM, only in simpler models, raising the question of whether some aspect of the simplification permits this behaviour.

“Statistical” Approach Establish a “traceable” hierarchy of models for simulations of the THC and related climate changes by tuning simpler (faster) models to emulate more detailed (slower) ones. Faster models in such a hierarchy can then be used instead of slower models when the latter cannot be afforded, in the expectation that the conclusions will be consistent. Explore the probability of rapid THC change as a function of model formulation by using the emulators. This could contribute to a more widely based assessment of the probability using results from RAPID and Hadley Centre projects.

Note that these two main objectives are of a different nature, and hence require two separate approaches. Specifically, the first objective requires a physically-based approach, whereas the second objective is best addressed within a statistically-based framework. The integration of the physically and statistically based approaches will rely on good communication between the oceanographic and statistical communities. This will give a central place to the project workshops and other meetings among participants. We will try to find a common language. For further discussion of analysis methodology for the physical and statistical approaches, see Appendix A.

“Reduction” approach Finally, it is of interest to explore a third approach closely tied to the physical and statistical approaches, namely the construction of an emulator of the THC via purely physical means. This approach is similar but different to the statistical approach in that the form of the emulator is not specified a priori, but rather determined in an ad-hoc fashion to take into account as much as possible of expert knowledge on all the feedbacks known to be important for the THC. In other words, the physical approach seeks to establish the functional form of the emulator, whereas the statistical approach seeks to calibrate the coefficients of an emulator whose functional form is specified *a priori*. In the present case, the ‘physical’ emulator is intended to take the form of a box-model of the coupled ocean/atmosphere system, able to reproduce the behaviour of a more complex AOGCM. Work in this direction is being carried out at the Hadley Centre and would obviously benefit from close interactions with the groups involved with the physical and statistical approaches. If successful, the box model so constructed could also serve as the starting point of a rigorous statistical analysis along the lines developed in the statistical group.

3.2 Models and experiments in the UK project

The UK project will benefit from the possibility of constructing a hierarchy of models of a greater range of complexity than available to the CMIP programme, as in Table 1. Some further notes on the various models and their components follow:

- Hadley Centre models: HadCM3 and HadGEM1. These two models are expected to be the most realistic climate models within the UK model hierarchy of the present project,

Model	Resolution	CPU cost	storage
HadCM3	$3.75°\text{Lon} \times 2.5°\text{Lat} \times 19\text{L}$ $1.25°\text{Lon} \times 1.25°\text{Lat} \times 20\text{L}$	2.2 yr/day 1xSX6	0.7 Gbyte/yr 350 Gbyte/500yr
HadGEM1	$1.88 \times 1.25 \times 38$ $1.0 \times 1.0 \times 40$	1.3yr/day 8xSX6	5 Gbyte/yr 2500 Gbyte/500yr
HiGEM	HadGEM1 at $1.25 \times 0.83 \times 38$ HadGEM1 at $0.33 \times 0.33 \times 40$	0.3yr/day 128xHPCx	10 Gbyte/yr 5000 Gbyte/500yr
CHIME	HadCM3 HYCOM at $1.25 \times 1.25 \times 25$	2.0 yr/day 24xGreen	1.7 Gbyte/yr 850 Gbyte/yr
FAMOUS	HadCM3 at $7.5 \times 3.75 \times 11$ HadCM3 at 3.75×2.5	50yr/day 8xNewton	0.1 Gbyte/yr 50Gbyte/500yr
FRUGAL	enhanced UVic 2D MOM with variable grid allowing high resolution in the Arctic	$>$100 yr/day cluster	similar to FORTE
FORTE	IGCM3 $T21 \times 11$ or $T42 \times 11$ MOM $2 \times 2 \times 15$ or $4 \times 4 \times 20$	8 yr/day workstation	0.15 Gbyte/yr 75Gbyte/500yr
GENIE IGCM-GOLDSTEIN	IGCM3 at $T21 \times 5$ GOLDSTEIN $5.6 \times 2.8 \times 16$	480 yr/day workstation	0.1 Gbyte/yr 50 Gbyte/500yr
GENIE C-GOLDSTEIN	UVic 2D GOLDSTEIN $10 \times 5 \times 8$	$>$10,000yr/day workstation	0.01 Gbyte/yr 5Gbyte/500yr
ECBilt-CLIO	T21 3 layer QG model MOM-like 3x3x20	n.a.	similar to FORTE

Table 1: List of models and their associated computer and storage cost estimates. In the "Resolution" column, the atmosphere is given first, then the ocean, each in the form: degrees of longitude $\times$ degrees of latitude $\times$ number of vertical levels, except for IGCM3, which has a spectral atmosphere model with triangular truncation given as Tn. In the "CPU cost" column, the $n\times$ entries refer to the number of CPUs used. SX6 is the NEC supercomputer at the Met Office. HPCx is the IBM supercomputer at Edinburgh administered by the EPSRC. Green and Newton are the supercomputers run by the Computer Services for Academic Research at Manchester. Estimates of storage are based on about 500 equivalent years per model, since this is the requirement of the CMIP standard experiments (Appendix B).

along with HiGEM. They have been extensively validated and used for policy-relevant climate projections assessed by the IPCC, and are therefore to be regarded as references. Details of HadCM3 are found in Gordon et al (2000), whereas HadGEM1 is described in Johns et al (2005).

- HiGEM: This is the highest-resolution model of the project, and hence the most computationally expensive to run. Owing to the limited funding available for the UK project, the participation of HiGEM in the UK project will depend on the availability of funding from elsewhere. It is anticipated, however, that this will be successful, and that HiGEM will be able to run the 1% CO2 increase experiment by the end of 2007 (year 2) of the project. A drawback of this model, however, is that in spite of the *a priori* advantage of a very high resolution, it is still in a development phase, and therefore has not yet been extensively validated. HiGEM is a project of NERC and Hadley Centre.

- CHIME: The advantage of CHIME is its being the only one to use a (hybrid) isopycnal-coordinate model, while its atmosphere is the same as HadCM3, permitting a particularly useful comparison. It is relatively computationally expensive to run. Details of CHIME are found in the COAPEC report NER/T/S/2001/00187.

- GENIE IGCM-GOLDSTEIN and C-GOLDSTEIN: The acronym GENIE refers to a metamodelling framework allowing to create an earth system model by coupling various subcomponents. In this project, we will make use of two particular frameworks. These models both make use of GOLDSTEIN, a fast low-resolution 3D frictional-geostrophic ocean model. Details are given in Edwards and Marsh (2005) and the e-science report NER/T/S/2002/00217, but will use different atmospheres. The atmosphere of C-GOLDSTEIN is a simple 2D energy balance model developed at the University of Victoria, whereas that of IGCM-GOLDSTEIN makes use of a spectral atmosphere GCM will simple physics. The latter is currently being set-up and will be used extensively in 2006 in the two Rapid projects of Chellenor et al (Probability) and Bamber et al (Cryosphere).

- IGCM3: The IGCM is a spectral atmosphere GCM with relatively simple physics developed at the University of Reading.

- FRUGAL: The model is similar to that of the Earth System model of the University of Victoria, with the possibility of increased resolution in the Arctic ocean, including resolution of straits, and an improved version of the EBM for the atmosphere. At the coarsest resolution, FRUGAL is $7lon \times 5lat \times 19L$ decreasing to 1 degree at its finest around Greenland. Ressources permitting, FRUGAL-HIGH can also be run; its resolution is $1.25lon \times 1lat \times 19L$ decreasing to 0.25 degrees at its finest. Details of FRUGAL are given in Bigg and Wadley (2001), Wadley et al. (2002), Wadley and Bigg (2002), Wadley and Bigg (2000).

- FAMOUS: FAMOUS is a lower resolution version of HadCM3, which is tuned to reproduce its behaviour as much as possible, but which is computationally about 10 times cheaper to run. Details on FAMOUS are given in Jones (2003), Jones et al (2005), and the RAPID report NER/T/S/2002/00462.

- ECBilt-CLIO: ECBilt-CLIO is an intermediate complexity model similar in design to FORTE. Given that this model will not be funded, its degree of participation will depend on the effort that can be invested in it by the KNMI.

3.3 Experiments

Standard experiments The participating groups have agreed on running the standard CMIP THC 1% per year CO2 increase experiment, the hosing experiment, and the partially-coupled experiments (a brief description is provided in Appendix B). Owing to computational costs, with CHIME only the first two of these experiments will be carried out, whereas HiGEM will probably be able to run only the CO_2 increase experiment.

Other experiments Ensembles of model runs with different parameter settings will also need to be assembled for the construction of the emulators using a strategy to be defined, for GENIE (C-GOLDSTEIN and IGCM-GOLDSTEIN) and FAMOUS. Experiments with atmosphere-only or mixed-layer ocean models could be useful in diagnosing climate feedbacks on, or caused by, changes in the THC (see Appendix A). Further experiments may be carried out with the faster coupled models in order to investigate hypotheses suggested by the analysis of the standard experiments.

Tracers In order to get additional information on the behaviour of the models, it appears of interest to look at the behaviour of passive tracers for each kind of experiments. An appropriate formulation of spiciness (see Appendix A) could give insight into the distinction between passive advection and the dynamical influence of temperature and salinity. This is also the aim of the passive anomaly tracers implemented in HadCM3 and HadGEM1, which track the propagation of the anomalous heat and freshwater added at the ocean surface.

Storage The international CMIP project collected only a few diagnostics, and only for decadal means. AOGCMs typically archive most quantities for averaging periods no shorter than a month. In this project we will consider storing daily data for a few months and/or five-daily data for a few years, in order to investigate the variability of the THC on different time scales (see Appendix A).

3.4 Significance and links to other projects

The UK intercomparison project is part of a wider array of studies that seek to quantify the uncertainty and probability of rapid climate change associated with a change in the ocean thermohaline circulation:

- Rowan Sutton's RAPID round 2 project examines the predictability of European climate with respect to uncertainties related to the initial conditions.

- Peter Challenor's RAPID round 1 project aims to assess the probability of rapid climate change under future climate scenarios using statistical emulators of C-GOLDSTEIN.

- THCQUMP: This Hadley Centre project aims at quantifying uncertainty in model predictions associated with the THC in the case of HadCM3, by performing an ensemble of about 20 parameter sensitivity experiments.

- climateprediction.net is a NERC-funded project with public participation aimed at performing parameter-sensitivity ensembles of simulations using an atmospheric model and a slab ocean, and is intended to be extended to use an AOGCM.

- Observational monitoring programs of the THC, such as the monitoring along $26°N$ (Cunningham et al), as well as bottom pressure recording along the western boundary current (Hughes et al).

- Helen Jonhson's project is to investigate some effects of the arctic ocean on the THC.

- Bamber et al's project aims to investigate the role of the cryosphere in modulating the thermohaline circulation of the Atlantic.

4 Management plan

4.1 Management of the project

Rather than a centrally organised comparison project, the project is intended as a collaboration among modelling groups, to be conducted by exchange of results and conclusions, with shared participation in the writing of papers. To that end, the project manager will ensure the sharing of information by visiting each group regularly, and by maintaining a website with up-to-date information on the availability of data and scientific progress, as well as by ensuring the availability of manuscripts in the writing by the UK partners as well as by the international participants. The project manager will also take the responsibility of leading the analysis of the model results in concert with the participating groups, takes a central role in the collaborative writing up of progress reports and papers in the refereed literature, in the dissemination of the results at conference, and ensures the timely delivery of the project deliverables.

4.2 Work plan

The project will have three main components: I) The production of model experiments, and the assembling of a database of model results that will made available to the RAPID community via the British Atmospheric Data Centre (BADC); II) The development of physically-based diagnostics aimed at understanding the physical reasons for different predictions among the models; III) The development of a rigorous statistical framework to understand how to systematically reduce uncertainty on model predictions by relating the complex models to the simpler ones. The success of the project would be increased considerably by a fruitful interaction between II) and III).

4.3 Participating groups

The participating groups are at:

- Centre for Global Atmospheric Modelling at the Department of Meteorology of the University of Reading
- James Rennell Division of the National Oceanography Centre, Southampton
- Met Office Hadley Centre, Exeter
- Department of Earth Sciences of the Open University, Department of Mathematical Sciences of the University of Durham, Centre for Ecology and Hydrology at Edinburgh (these participants constitute the statistical group)
- Department of Geography of the University of Sheffield
- Climate Modelling group at the Dutch Royal Meteorological Institute (KNMI)

The participants are listed in Section 6.

4.4 Timeline and deliverables

We presume the project can start by late in 2005 or early in 2006, depending on recruitment.

Year 1

- The project manager and PI will organise a startup meeting. The preparation for the meeting will consist in making a synthesis of available theoretical and observational results that will be useful to constrain the physical analyses of the models.
- The startup meeting will convene the participants once the new staff have been appointed. The meeting will focus on reviewing the practicalities pertaining to the production of the experiments, and agree on a timeline. We will discuss the physics, diagnostics, and results obtained so far. We will identify the range of papers that can be written, and discuss how to coordinate the analysis.
- The Southampton group will run the experiments for CHIME and FORTE. To that end, an important step will be the porting of CHIME to the Bull system. This will consist of the following stages:
    1. Install the UMUI (Unified Model User Interface) to be able to generate run scripts.
    2. Port the HYCOM ocean model (including ice), test and optimise its performance.
    3. Port and test the PUM (Portable Unified Model), vn 4.5 (this provides the atmosphere model).
    4. Install and test OASIS (using MPI).
    5. Connect the atmosphere model to the ocean-ice model.
    6. Test and tune the coupled model

The critical stage will be step 3 (installing the PUM). This is the only step the group has not undertaken directly itself. However, a suitable version of the PUM has already been obtained through Jeff Cole in Reading. This has already been run on "Newton" at the CSAR service in Manchester, which has an architecture similar to that of the Bull at Southampton. Further advice will be sought from Jeff Cole if unexpected problems arise. All other stages of the port are similar to the NOC installation of CHIME at CSAR, and are not expected to involve additional problems.

We therefore feel that a clear breakpoint will be provided by the successful completion of Step 3. We anticipate this will be accomplished either within 2 months of the start of the project (with Dr. Sinha working full time on this), or, from a practical point of view, by January 2006. The Southampton group therefore propose to report back to the modelling sub-committee at this time, detailing progress to date, and, if necessary, available options.

- The Sheffield group will run the experiments for FRUGAL.
- The Met Office group will run the HadGEM1 experiments. (The HadCM3 experiments have already been carried out.)
- The Reading group will run the experiments for FAMOUS.
- The KNMI modelling group will provide the data for ECBilt-CLIO.
- The statistical group will construct the statistical framework to link the different models, and define the experimental design as to the optimal way to construct the ensemble of model runs required to construct the statistical emulators. This will include performing the standard experiments, as well as complementary experiments, with the GENIE models. More details on the steps required are provided at the end of appendix A.
- Richard Wood will pursue ongoing work on constructing a physically-based emulator of the THC using a box-model (not RAPID-funded)
- The project manager will perform diagnostics and analyses of already available model results in order to select those that will be useful for the model intercomparison. This will be done in concert with members of the other groups engaged in the analysis of the model results.
- Helen Johnson will participate in the analysis of the modelled water masses.
- The PI and project manager will communicate the results of the project at international conferences (AGU fall meeting and EGU conference), and meet with the international partners of the project at these occasions to maintain collaboration and discuss joint efforts.

The deliverables of year 1 will be

- The database of model output from GENIE, CHIME, FAMOUS, FORTE, FRUGAL, ECBilt-CLIO, HadCM3 and perhaps HadGEM1 (this effort is not RAPID-funded).
- Definition of a clear strategy for analysing the model outputs, for both the physical approach and the statistical approach.

- Determination of the need for additional experiments if deemed necessary in the light of results emerging from the study. For instance, it could be useful to know more about the sensitivity of the THC response to the hosing region. Taro Hosoe has already planned to investigate this issue with FAMOUS, and will be able to report on his work by the end of year 1, to determine whether it would be worthwhile to conduct that kind of sensitivity experiments with other models. Insights into useful additional experiments can also be obtained from the ensemble of experiments constructed as part of the statistical part of the project.

Year 2 and 3 The second and third year will largely depend on the progress achieved during year one. It will include:

- Holding regular meetings (every six months) to review progress, discuss model results, and refine strategies.
- Continue the analysis of the model results and collaborative writing of papers.
- Possibly carry out further experiments with the faster coupled models in order to investigate hypotheses suggested by the analysis of the standard experiments achieved so far.
- Carry out the CO_2 experiment with HiGEM during year 2.
- Organisation of a final meeting to which the CMIP participants will be invited to compare results and analyses. This meeting will be supported by the recently successful proposal by Peter Challenor to the NERC International Opportunities fund.
- Writing of the final report detailing the results achieved, and recommendation for computing probabilities of rapid climate change, and reducing uncertainties.

5 Request for support

5.1 Justification of resources

Running the standard experiments with CHIME will first require migrating the model from CSAR to a NOCS Bull NOVASCALE computer. The latter is composed of 1.3 Ghz Intel Itanium 2 processors, linked together in various clusters or "nodes", each with its own shared memory. The system presently has 56 processors. The notional cost for the Bull system is about £80k/year over its lifetime. It is estimated that the CHIME runs will take about 25% of the machine time for a year. £20k are therefore requested to contribute to the running cost of the system, which will serve to enhance the system, most likely by buying a cluster of 4 Itanium processors, with a shared memory of 8 Gbyte, and quadrix interconnect to the other nodes on the Bull. It is proposed that the migration be carried out by Bablu Sinha during the fall of 2005, who would be funded for three months. Subsequently, a Grade I Research Assistant level would be appointed for two years with the responsibility of running CHIME and FORTE, and contributing to the analysis of the results. Funding is also requested to contribute toward the storage costs of the model results. At present, archiving at NOCS costs £6k for 1Tb/3 years. Scratch disk cost £2.5k for 1Tb/3 years. Backed up raid disk with fast access cost £3k for 50Gb/4 years.

The statistical work would be carried out by a team consisting of Jonathan Rougier, David Cameron, Michael Goldstein, and Neil Edwards. Rougier and Cameron would be funded for a total of 18 months over 3 years (6 months and 12 months respectively). Rougier and Goldstein jointly developed the reification approach, and are leading experts in the design and analysis of computer experiments. Edwards is a principal developer of GOLDSTEIN-based models, on which most of the work will be focussed. Cameron has extensive experience in running large ensemble experiments using these models and will be in charge of running the experiments with the GENIE models.

The work of running FRUGAL would be carried out by Martin Wadley (University of East Anglia), in collaboration with Grant Bigg, at the University of Sheffield, on a local LINUX cluster system, for which some funding should be allocated toward upgrading the system to make it more efficient. We request funding for Martin Wadley on a 1/3 part time basis over the first six months of 2006.

One year funding is requested in 2007 to support Len Shaffrey for production and analysis of the HiGEM results at the University of Reading. Len Shaffrey has been contributing to the development of HiGEM, and has become highly experienced with the model.

The database will be stored long-term at BADC. The current cost of storage is about £1.5k per Terabyte. We estimate that a minimum of 4 Terabytes will be necessary. To allow for further experiments, and backup of the most important data, we ask for an additional 2 Terabytes. This amounts to 6 Terabytes, for an estimated cost of £9k.

We plan to hold one-day workshops every six months or so to exchange and present new results, and to discuss how to report them in peer-reviewed journals. We shall assume an average of 10 people per meeting, with an average of £100 per participant, which totals £5k for 5 meetings. In addition, travel money (£300) is requested to allow collaboration between Grant Bigg (Sheffield) and Martin Wadley (UEA, Norwich) to carry out the work with FRUGAL, as well as £700 to allow collaboration with KNMI. We also request additional travel money (£5,000) in order to be able to invite CMIP international partners to our workshops.

5.2 Budgeting

5.2.1 Staff (current costing scheme, incl. 46% overhead)

Category	2005	2006	2007	2008	total
RA1 NOCS	0	42,060	43,342	0	85,402
Bablu Sinha (NOCS) 3 months	16,078	0	0	0	16,078
Martin Wadley (U. East Anglia)	0	9,981	0	0	9,981
Len Shaffrey (Reading)	0	0	56,895	0	56,895
D. Cameron (CEH, Edimburgh)	4,565	18,932	19,711	15,435	58,643
J. Rougier (Durham)	0	0	14,544	15,267	29,811
Total	20,643	70,973	134,492	30,702	**256,810**

Note: The estimate for the RA1 in NOCS is based on a bottom scale PDRA. The 2005 figure for Bablu Sinah, previously given with a starting date of 1st October, now assumes a starting date of 1st January 2006. This is given for the year 2005, as the work will start as soon as the project is approved. The figures for David Cameron are based on the OU financial year based on a starting date of 1st of January 2006. The number for 2005 are therefore for the financial year 2005/2006 and so on. The figure for Martin Wadley is based on 1/3 part time over the first six months of 2006. The figure for David Cameron is based on 1/3 part time basis over 3 years, i.e., 12 months in total, while that for Jonathan Rougier is based on 1/4 part time basis over the 2 last years of the project, i.e., 6 months in total. The figure for Len Shaffrey is based on full time for the whole year 2007. All figures include salary+ pension + NI + 46% overhead.

5.2.2 Staff (FEC scheme)

Category	Year 1	Year 2	Year 3	Total
RA1, NOCS	28,808	29,686	0	58,494
Bablu Sinha (NOCS)	11,012	0	0	11,012
PI Time (Adrian New, NOCS)	6,066	6,066	6,066	18,198
Co-I Time (Alex Megann, NOCS)	4,354	4,354	4,354	13,062
Co-I Time (Bob Marsh, NOCS)	4,741	4,741	4,741	14,223
Co-I Time (Bablu Sinah, NOCS)	0	4,405	4,405	8,810
Estates Cost (NOCS)	17,626	15,422	4,406	37,454
FEC indirect (NOCS)	57,317	50,152	14,329	121,798
Martin Wadley (Sheffield)	19,568	0	0	19,568
PI Time (Grant Bigg, Sheffield)				
Len Shaffrey (Reading)	0	0	91,394	91,394
PI Time (Warwick Norton)				
D. Cameron (CEH, Edimburgh)	12,756	39,797	27,552	80,105
N. Edwards (O. University)	5,267	5,359	5,430	16,056
J. Rougier (Durham)	0	18,046	18,540	36,586
M. Goldstein (Durham)	0	1,853	1,853	3,706
Total	167,515	179,881	183,070	**530,466**

(1) The above numbers assume 1/1/2006 as a starting date, except for the three months of Bablu Sinha expected to take place as soon as possible, and completed before 1st April 2006.
(2) Pi and Co-I time at NOC is based on 10 percent for three years in the case of Adrian New,

Bob Marsh, and Alex Megann, and on 10 percent for two years in the case of Bablu Sinah.
(3) The computation for David Cameron assumes 36 working days by 1st April 2006.
(4) PI time for Neil Edwards and Michael Goldstein is based on 2hours/week for three years for N. Edwards, and on 2hours/week for the last two years of the project for M. Goldstein.

5.2.3 Consumables/computing

Description	Laboratory	Cost
Storage and maintenance Contribution to Bull system costs	NOCS	15,000 20,000
BADC Storage costs	BADC	9,000
Consumables	Sheffield	3,000
Consumables	Open University	400
Publication/color figure costs		10,000
Total		57,000

5.2.4 Travel and related

Description	Laboratory	Cost
workshops/travel	Reading	11,000
Total		11,000

5.2.5 Summary of resources

Staff	Consumables/storage	Travel	Total
256,810 (current scheme)	57,400	11,000	325,210
530,466 (FEC)	57,400	11,000	598,866

6 List of UK project participants

Name	Institution	Role
J. Gregory	U. Reading	PI, FAMOUS+analysis
R. Tailleux	U. Reading	Project Manager, FAMOUS+analysis
H. L. Johnson	U. Reading	analysis
W. Norton	U. Reading	HiGEM
Len Shaffrey (1year during year 2)	U. Reading	HiGEM
A. New	NOCS	CHIME + analysis
A. Megann	NOCS	CHIME+analysis
B. Sinha	NOCS	CHIME+FORTE+analysis
R. Marsh	NOCS	GENIE
RA1 (2years)	NOCS	CHIME+FORTE
A. Pardaens	MetOffice	HadGEM+ analysis
M. Vellinga	MetOffice	HadCM3+analysis
R. Woods	MetOffice	Traceability studies
N. Edwards	Open University	GENIE, Traceability
J. Rougier	Durham	statistics, traceability
D. Cameron	Edimburgh	Statistics, traceability
M. Goldstein	Edimburgh	Statistics, traceability
G. Bigg	Sheffield	FRUGAL
M. Wadley	U. East Anglia	FRUGAL
S. Weber	KNMI	ECBilt-CLIO

A Analysis methodology and interpretative frameworks

A.1 Physical approach

Section 5 of the CMIP case for support proposed an analysis methodology in which the basis for understanding inter-model differences is linked to the existence of a conceptual picture capable of factorising influences on THC change into (a) the effect of climate on surface fluxes; (b) the effect of surface flux changes on the density field; (c) the effect of the density field on the ocean circulation. The extent to which this methodology can be successful depends critically on how far physical understanding can be achieved in the process of relating inter-model differences to well-identified physical processes, either explicitly resolved or parametrised. Achieving physical understanding will rely on the use of innovative diagnostics on model data and observations, in relation with existing theoretical frameworks or ones to be developed, for which some details are provided below. The physical approach does not *a priori* require additional experiments beyond the CMIP coordinated sets of experiments.

- *What are the driving mechanisms of the thermohaline circulation?* The large-scale ocean circulation is geostrophic at leading order, so that the strength of the MOC must be primarily controlled by the East-West pressure gradient, which itself is related to the differences in bottom pressure between the eastern and western parts of the basin in a realistic ocean with sloping bottom topography. Several studies indicate that the strength of the

MOC is also dictated to various extents by the north-south density gradient, vertical diffusion of heat and salt, and by deep water mass formation. Other studies also point to the possibility of significant remote control of winds in the Antarctic Circumpolar Current on the MOC. To make progress, therefore, it will be important to understand the respective importance of each of these effects in controlling the strength of the MOC among the various models, in order to determine whether these can account for the different behaviours observed. An important issue will be to understand the possible time dependence of the relationships between MOC, density gradients, mixing, and deep water formation, and to determine the range of time scales for which they are valid. To that end, it will be useful to investigate the dynamical processes controlling the variability of the THC, as well as its adjustment to changes in the forcing such as boundary, equatorial, and planetary waves. Understanding the latter processes will require model outputs saved at a higher-resolution time frequency than monthly.

- *How important are mixing processes in controlling the MOC?* Many studies point to the crucial importance of mixing in controlling the strength of the MOC. In the past, simple relationships have often been found between the strength of the MOC and vertical diffusivity, when the latter is a constant. In most recent models, however, the formulation of mixing processes escapes such a simple description. Owing to the importance of mixing and water mass formation on the MOC, there is an obvious need for being able to discriminate the behaviours of the different models with respect to such quantities. To that end, we clearly need a way to attach each model with an 'effective bulk mixing efficiency'. Some possible ways to look at this issue is global energetics using the ideas of Winters et al (1995) or Toggweiler and Samuels (1998) for instance. Another possible approach would be by looking at the residual term in mass budgets of various density classes *à la* Speer and Tziperman (1992) for instance.

- *What is the link between a change in surface heat and freshwater fluxes and changes in the MOC?* Changes in surface heat and freshwater fluxes result in changes in buoyancy fluxes, which alter the overall mass balance of each density classes, and changes in spiciness fluxes, which do not. We expect that a change in surface buoyancy fluxes will make the ocean strive toward regaining balance of the sources and sinks of mass within each density classes by reorganising the density field (Tziperman 1986, Speer and Tziperman 1992). Provided that we have been successful in linking the strength of the MOC to some measure of the north-south density gradient, even for transient cases, we may be able to understand how changes in surface fluxes will affect changes in the MOC by investigating how the changes in surface fluxes affect the north-south density gradient for instance. Insights into this issue should benefit from systematically investigating the distribution of diapycnal fluxes along isopycnal surfaces, as well as from systematically analysing the MOC in density/latitude coordinates, in addition to the more classical depth/latitude representation.

- *What are the processes controlling the existence of multiple equilibria, hysteresis, and irreversible collapse, of the THC?* Multiple equilibria of the THC for a given surface forcing have long been invoked as a possible source of rapid climate change, yet their precise nature is still very poorly understood. Existence of multiple equilibria has been demonstrated in coarse resolution models, whereas they are much harder to find in higher-resolution models, raising the question as to the proximity of the present day THC to an instability threshold. The issue of the stability of the THC is a crucial issue for assessing the probability of rapid climate change, so that it appears important to understand why the THC can irreversibly collapse in some models, whereas it does not in other models. In order to tackle this issue, a possible starting point is to assume that there must in general be an energy barrier separating two stables regimes of the THC, so that work must be provided to the system for any transition to occur. We may try to understand how this source of energy is provided to the ocean in cases where a collapse of the THC is observed, and why the same does not occur in the other models, by investigating the time-dependent energy budget of the THC.

- *What is the three-dimensional structure of the MOC?* In most studies of the MOC, the strength of the latter is usually measured by the maximum of the meridional overturning streamfunction. This is of course a highly-aggregated measure of the MOC, which hides its three-dimensional character, and therefore may prevent us from fully understanding it. To get insights into this issue, it seems important to understand how the maximum of the meridional overturning streamfunction relates to various representation of its three-dimensional structure, as achieved for instance by looking at various zonal sections of the meridional velocities, or by investigating other ways to aggregate information on the MOC by means of EOFs for instance.

- *How does the strength of the MOC relate to that of the meridional heat transport?* The present paradigm of the climate impact of the MOC is that there is a strong correlation between the strength of the MOC and the meridional heat transport, so that a slowdown (or even collapse) of the MOC is expected to induce a cooling (warming) of the northern (southern hemisphere) hemisphere, e.g., Vellinga and Wood (2002). This would moderate the impact of global warming in the northern hemisphere, but probably exacerbate it in the southern hemisphere. The precise connection between meridional overturning and heat transport, however, is still uncertain, so that it will be of interest to investigate how such a such a relationship varies from model to model, which may possibly allow us to get a more physical understanding of it.

- *What are the climate feedbacks on the MOC?* The important issue here is to understand how the atmosphere system will react to a change in the MOC, and whether it will react to exacerbate or compensate for its effects? For instance, some studies indicate a stabilisation effect in the form of enhanced evaporation in the tropics which allow for an increase of salinity compensating for the freshening of the northern regions (Latif et al 2000). Other studies indicate a stabilising effect of climate, as the THC seems able to recover after a collapse is artificially induced, e.g., Vellinga et al (2002). It will be of interest to investigate the respective behaviour of each model in this respect, in order to quantify their destabilising, stabilising, or possibly neutral impact on the MOC.

- *What are the mechanisms for changes in total heat content in the ocean?* As said above, changes in surface heat and freshwater fluxes result both in changes in spiciness and buoyancy fluxes. Spiciness is an important quantity because it does not affect the density field. For instance, if changes in the surface heat and freshwater fluxes would exactly compensate as to induce only a change in spiciness fluxes but not in density fluxes, we would not expect changes in the density field, and hence in the MOC. Yet, it would still be possible for the heat and salinity fields to change over time; as a result, it would be in principle possible to observe a change in the meridional heat transport without observing any change in the MOC. This thought experiment therefore suggests that the precise link between the strength of the MOC and meridional heat transport depends to some unknown extent on spiciness fluxes and the transport of spiciness anomalies in the ocean interior. Spiciness anomalies are also important because they have been invoked as a possible way for midlatitudes to impact on equatorial regions via the so-called thermocline bridge, see Tailleux et al. (2005) for references. This way, spiciness makes it possible for a temperature and salinity anomalies subducted in the midlatitudes to resurface year later in the equatorial regions to impact on the local climate there, with possible stabilising or destabilising effects on the MOC as suggested by Latif et al. (2000). The above comments shows therefore that there is a link between understanding the ocean/atmosphere interactions, and understanding the behaviour of passive tracers in the ocean. In this project, we propose to look at the behaviour of spiciness. If needed, it might also be of interest to investigate the behaviour of more idealised tracers to look at more specific aspects of the problem.

- *Observational constraints.* The intercomparison of the models does not provide by itself insights into their respective shortcomings, so that there is obviously the need for observational constraints to assess the physical realism of each model. The issue is complicated by the fact that the project focusses on the MOC, a quantity that is not directly observable. Nevertheless, there is current observational effort within RAPID at measuring the strength of the MOC in the North Atlantic, by means of arrays along $26°N$, as well as by means of pressure recorders in the western boundary. Such observations could be compared with the results of each model. In the present context, it would also be of interest to determine whether the quantities investigated by Vellinga and Wood (2004) prove to be as informative as indicators of MOC change in each model as they are in HadCM3. Such a result would provide a firm basis for extending the observational strategy of the MOC beyond what is presently done.

Initial steps towards implementing some aspects of the above will include:

1. Synthesis of theoretical knowledge on the links between the MOC and such quantities as the density field, mixing, and deep water formation to facilitate the quest for empirical links between the MOC and some measure of the north-south density gradient for instance. Theoretical investigation of the time-dependent aspects of the problem, as well as of the effects of realistic geometry and topography (as opposed to the use of a flat-bottom and vertical walls, as found in many theoretical studies).

2. Assessment of the respective merits of existing formulations of spiciness for use as a passive tracer in the ocean; most likely, construction of a new spiciness variable as a function of salinity, temperature, and pressure minimising the coupling with the density field (which arises as a result of the nonlinearity of the equation of state).

3. Assessment of the respective merits of various ways to investigate the global energetics of the coupled system, i.e., use of potential energy versus available potential energy (APE).

Estimate the errors associated with using approximate forms of APE, as well as the feasibility and cost of using more accurate forms of APE. Assessment of the practical details, effort, and additional computational costs required to implement global energetics budget in z-coordinate models as well as isopycnic ones.

4. Evaluate the feasibility and accuracy of computing mass budgets for various density classes (with the purpose of estimating a bulk amount of mixing and water mass transformations among other things) with models outputs saved at and averaged over various time intervals. Estimate the problems arising from the transience and drift of the solutions. Assess the minimum conditions for achieving sufficient accuracy with such budgets, and the possible need to save additional variables (such as the time-averaged depths of relevant isopycnal surfaces, as well as the time-averaged buoyancy fluxes at the bottom of the mixed layer for instance).

A.2 Statistical approach

A possible limitation of the above physically-based analysis methodology, however, may lie in its inability to clearly relate inter-model variability in the THC predictions to intra-model variability generated by parametric uncertainty, which is essential for quantification of real-world uncertainties, and which thus points to the need for a complementary statistically based approach. Given the cost of high-resolution model simulations, the only realistic way to assess parametric uncertainty in the very high-dimensional space of uncertain model parameters is to use simplified models which can be run very many times, and to relate the behaviour of the simple and complex models within a rigorous statistical framework. In the approach[1] of Goldstein and Rougier (2005) a hierarchy of physically-based models 'simulators' is embedded in a framework of closely-related statistical 'emulators' which reproduce the behaviour of the simulators. Using this structure, it is possible to transfer information about sensitivity to parameters from simple to complex models in a quantitative way, allowing us to make statements about parametric uncertainty in complex models that are informed by evaluations of simpler, faster models. The framework also makes it possible to quantify the reduction in structural error within the hierarchy, a particularly difficult task. More ad-hoc assimilation-based techniques have been successful in the choice of model parameters, but there is limited basis for the estimates of uncertainty that they produce, and limited repeatability in model calibration.

The statistical framework consists of a set of emulators (effectively models of models) typically taking the form of expansions in a polynomial basis, which link simulator inputs and outputs. Uncertainty is encapsulated by the variance of the random coefficients of the basis functions, and by the covariance of the discrepancy between model outputs and observations. Construction of the framework involves the specification and determination of all the relevant functions and coefficients by a combination of simulator evaluations and physical interpretation. The framework can be used to quantify the uncertainty in predictions, which should reduce as more simulations are incorporated at any level in the hierarchy.

A vital component of the work is the interaction between the rigorous statistical approach and the process-based physical attribution of cause and effect. This interaction is central to the analysis because physical understanding informs the specification of functional forms and error covariances. For example, a functional dependence on eddy diffusivity at low resolution

[1] baptised 'reified analysis', the process of reification roughly meaning making abstract concepts real

translates, further up the hierarchy, into a joint dependence on sub-grid scale diffusivity and resolved-scale eddy flux. Such hierarchical dependencies can be built in to the relationships between emulator coefficients. This can allow physically based causal interpretations of inter-model variability, e.g., via differences in convection or water-mass production, to be compared with intra-model parametric variability in the context of overall uncertainty.

Such a statistically-based approach was not evoked in the CMIP case for support, and is therefore a novelty of the UK intercomparison project. Since a crucial element of the statistical part of the project is the construction of emulators, for which model ensembles exploring part of their parameter spaces is necessary, it follows that additional experiments will be needed to construct such ensembles, which the standard coordinate sets of CMIP experiments does not provide.

Rather than attempting to relate all models in the intercomparison project, we will address a subset of three: C-GOLDSTEIN and IGCM-GOLDSTEIN, which use the same ocean model, and FAMOUS. This subset spans a wide range of complexity from a very efficient model (C-GOLDSTEIN) to one closely related to HadCM3 (FAMOUS). The interpolation of an intermediate model (IGCM-GOLDSTEIN) makes it possible to trace the physics across the hierarchy, and also simplifies the statistical framework. Depending on time availability, the study may also make use of FORTE, which has the same atmosphere component as IGCM-GOLDSTEIN.

To implement the statistical approach involves:

1. Construction of the statistical framework linking models that respects the hierarchy of model structure and the generalisations in the model physics (so-called 'reified analysis'). The framework is effectively a set of 'models of models' which include error terms.

2. Experimental design: to map the behaviour of the models across their high-dimensional parameter spaces efficiently using a limited number of simulations, it is essential to select parameter value intelligently. Initially we use a broad design across all models (favouring those that are cheaper) but this can be refined sequentially through the project to concentrate later evaluations on important and highly informative regions of the models' parameter spaces.

3. Model simulations and the assimilation of both model results and physical insight as they become available: new results and new insights into physical processes can be used to refine the statistical framework, and thus the estimates of uncertainty, while the framework can be used to search for explanations of differences in model behaviour.

B Description of the CMIP standard coordinated experiments

B.1 CO_2 experiments

This part studies the role of the surface fluxes in producing the response to time-dependent climate change on the century timescale. It comprises four experiments of 140 years each, which are described in detail in Gregory et al. (2005). The first is a CONTROL experiment with constant CO_2 and steady-state climate. The other experiments all begin from the initial state of the CONTROL. The TRANSIENT experiment adopts a scenario of CO_2 increasing at 1% per year compounded, bringing it after 140 years to four times its initial concentration (denoted $4 \times CO_2$).

The remaining two experiments are CRAD_TH2O, which has constant CO_2 at the CONTROL values and imposes on the ocean the surface water fluxes saved from TRANSIENT, and TRAD_CH2O, which has CO_2 increasing at the same rate in TRANSIENT but with the surface water fluxes from CONTROL. These integrations are thus called "partially coupled" (Mikolajewicz and Voss, 2000) because the exchange of water has been disrupted; the water fluxes computed from the atmosphere and surface do not equal those applied to the ocean. The partially coupled integrations are intended to quantify the relative importance of changes in surface heat fluxes and surface water fluxes on the ocean circulation. The idea is that the changes in CRAD_TH2O will be caused by surface water flux changes and in TRAD_CH2O by surface heat flux changes. The design supposes that these will add linearly to produce the total change of TRANSIENT; this can be tested from the results of the experiments.

B.2 Hosing experiment

This part studies sensitivity to an abrupt change in surface water flux forcing. A (fresh) water flux of 0.1 Sv is applied for 100 years to the north Atlantic. The flux is then switched off and the experiment continues, in order to see whether the THC recovers its initial strength. A control integration, with no water flux forcing, runs in parallel. The size of the forcing ("hosing") flux is chosen to be of the order of magnitude predicted for a large CO_2-induced climate change.

References

Bigg, G. R., and M. R. Wadley, 2001: Millennial changes in the oceans: an ocean modeller's viewpoint. *J. Quater. Sci.*, **16**, 309-319.

Bryan, F., 1987: Parameter sensitivity of primitive equation ocean general circulation models. *J. Phys. Oceanogr.*, **17**, 970-985.

Edwards, N. E., and R. Marsh, 2005: Uncertainties due to parameter-transport sensitivity in an efficient 3-D ocean climate model. *Climate Dynamics*, **24**, 415-433.

Flament, 2002: A state variable for characterising water masses and their diffusive stability: spiciness. *Progress in oceanography*, **54**, 493-501.

Goldstein and Rougier (2005): Reified Bayesian modelling and inference for physical systems, acccepted in JSPI.

Goldstein, M., and J.C. Rougier (2005): Probabilistic formulations for transferring inferences from mathematical models to physical systems. *SIMA J. Sci. Comp.*, **26**, 467-487.

Gordon, C., C. Cooper, C.A. Senior, H. Banks, J.M. Gregory, T.C. Johns, J.F.B. Mitchell and R.A. Wood, 2000. The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre coupled model without flux adjustments. *Climate Dynamics*, **16**, 147-168.

Gregory et al. 2005: A model intercomparison of changes in the Atlantic thermohaline circulation in response to increasing atmospheric CO2 concentration. *Geophys. Res. Lett.*, **32**, L12703, doi:10.1029/2005GL023209.

Johns, T.C., C.F. Durman, H.T. Banks, M.J.Roberts, A.J.McLaren, J.K. Ridley, C.A. Senior, K.D. Williams, A. Jones, G.J. Richard, S. Cusack, W.J. Ingram, D.M.H. Sexton, M.M. Joshi, B.-W. Dong, H. Spencer, R.S.R. Hill, J.M. Gregory, A.B. Keen, A.K. Pardaens, J.A. Lowe, A. Bodas-Salcedo, S. Stark, and Y. Searl (2005): The new Hadley Centre climate model HadGEM1: Evaluation of coupled simulations in comparison to previous models. *J. Climate*, submitted.

Jones, C., 2003: A fast GCM without flux adjustments. *J. Atm. Ocean Tech*, **20**, 1857-1868.

Jones, C. D., J. M. Gregory, R. B. Thorpe, P. M. Cox, J. M. Murphy, D.H.M. Sexton, and P. Valdes, 2005: Systematic optimisation and climate simulation of FAMOUS, a fast version of HadCM3, *Clim. Dyn.*, doi:10.1007/s00382-005-0027.2.

Latif et al., 2000: Tropical stabilisation of the thermohaline circulation in a greenhouse warming simulation. *J. Climate*, **13**, 1809-1813.

Roberts, M.J., *et al*, 2004: Impact of an eddy-permitting ocean resolution on control and climate change simulations with a global coupled GCM. *J. Climate*, **17**, 3-20.

Rougier (2005): Probabilistic inference for future climate using an ensemble of simulator evaluations, submitted.

Schmidt, G., 1998: Oxygen-18 variations in a global ocean model. *Geophys. Res. Lett.*, 25, 1201-1204.

Sinha, B., and R.S. Smith, 2002: Development of a fast coupled general circulation model (FORTE) for climate studies, implemented using the OASIS coupler, Southampton Oceanography Centre Internal Document, No. 81, 67 pp.

Speer and Tziperman, 1992: Rates of deep water formation in the North Atlantic Ocean. *J. Phys. Oceanogr.*, **22**, 93-104.

Stouffer et al. 2005: Investigating the causes of the response of the thermohaline circulation to past and future climate changes. J. Climate, in press.

Toggweiler and Samuels (1998): On the ocean's large-scale circulation near the limit of no vertical mixing. *J. Phys. Oceanogr.*, **28**, 1832-1852.

Tailleux et al. (2005): Physics and dynamics of density-compensated temperature and salinity anomalies. Part I: Theory, *J. Phys. Oceanography*, **35**, 849-864.

Tziperman, 1986: On the role of interior mixing and air-sea fluxes in determining the stratification and circulation of the ocean. *J. Phys. Oceanogr.*, **16**, 680-693.

Vellinga, M., R.A. Wood, and J.M. Gregory, 2002: Processes governing the recovery of a perturbed thermohaline circulation in HadCM3. *J. Climate*, **15**, 764-780.

Vellinga, M., and Wood, R.A., 2002: Global climatic impacts of a collapse of the Atlantic thermohaline circulation. *Climatic Change*, **54**, 251-267.

Vellinga, M., and Wood, R.A., 2004: Timely detection of anthropogenic change in the Atlantic meridional overturning circulation. *Geophys. Res. Let.*, doi:10.1029/2004GL02306,2004.

Wadley et al., 2002: On modelling present-day and last glacial maximum oceanic $\delta^{18}O$ distributions. *Global Plan. Change*, **32**, 89-109.

Wadley, M. R. and G. R. Bigg, 2000: Implementation of variable time stepping in an ocean general circulation model. *Ocean Modelling*, **1**, 71-80.

Wadley, M. R. and G. R. Bigg, 2002: Impact of flow through the Canadian Archipelago on the North Atlantic and Arctic thermohaline circulation: an ocean modelling study. *Quart. J. Roy. Meteor. Soc.*, **128**, 2187-2203.

Winters et al. (1995): Available potential energy and mixing in density-stratified fluids. *J. Fluid Mech.*, **289**, 115-128.