UNIVERSITY OF Southampton
School of Engineering Sciences

Multi-objective Optimisation of GENIE Earth System Models

Andrew R. Price[1], Richard J. Myerscough[2], Ivan I. Voutchkov[1], Robert Marsh[2] & Simon J. Cox[1]

[1] School of Engineering Sciences, University of Southampton, Southampton SO17 1BJ, UK
[2] National Oceanography Centre, University of Southampton, Southampton SO14 3ZH, UK

9 September 2008

Overview

- GENIE Project
- Multi-objective Optimisation
- Surrogate Modelling
- Grid Computing Infrastructure
- Parameter Estimation for a new Ocean Mixing Scheme
- Conclusions

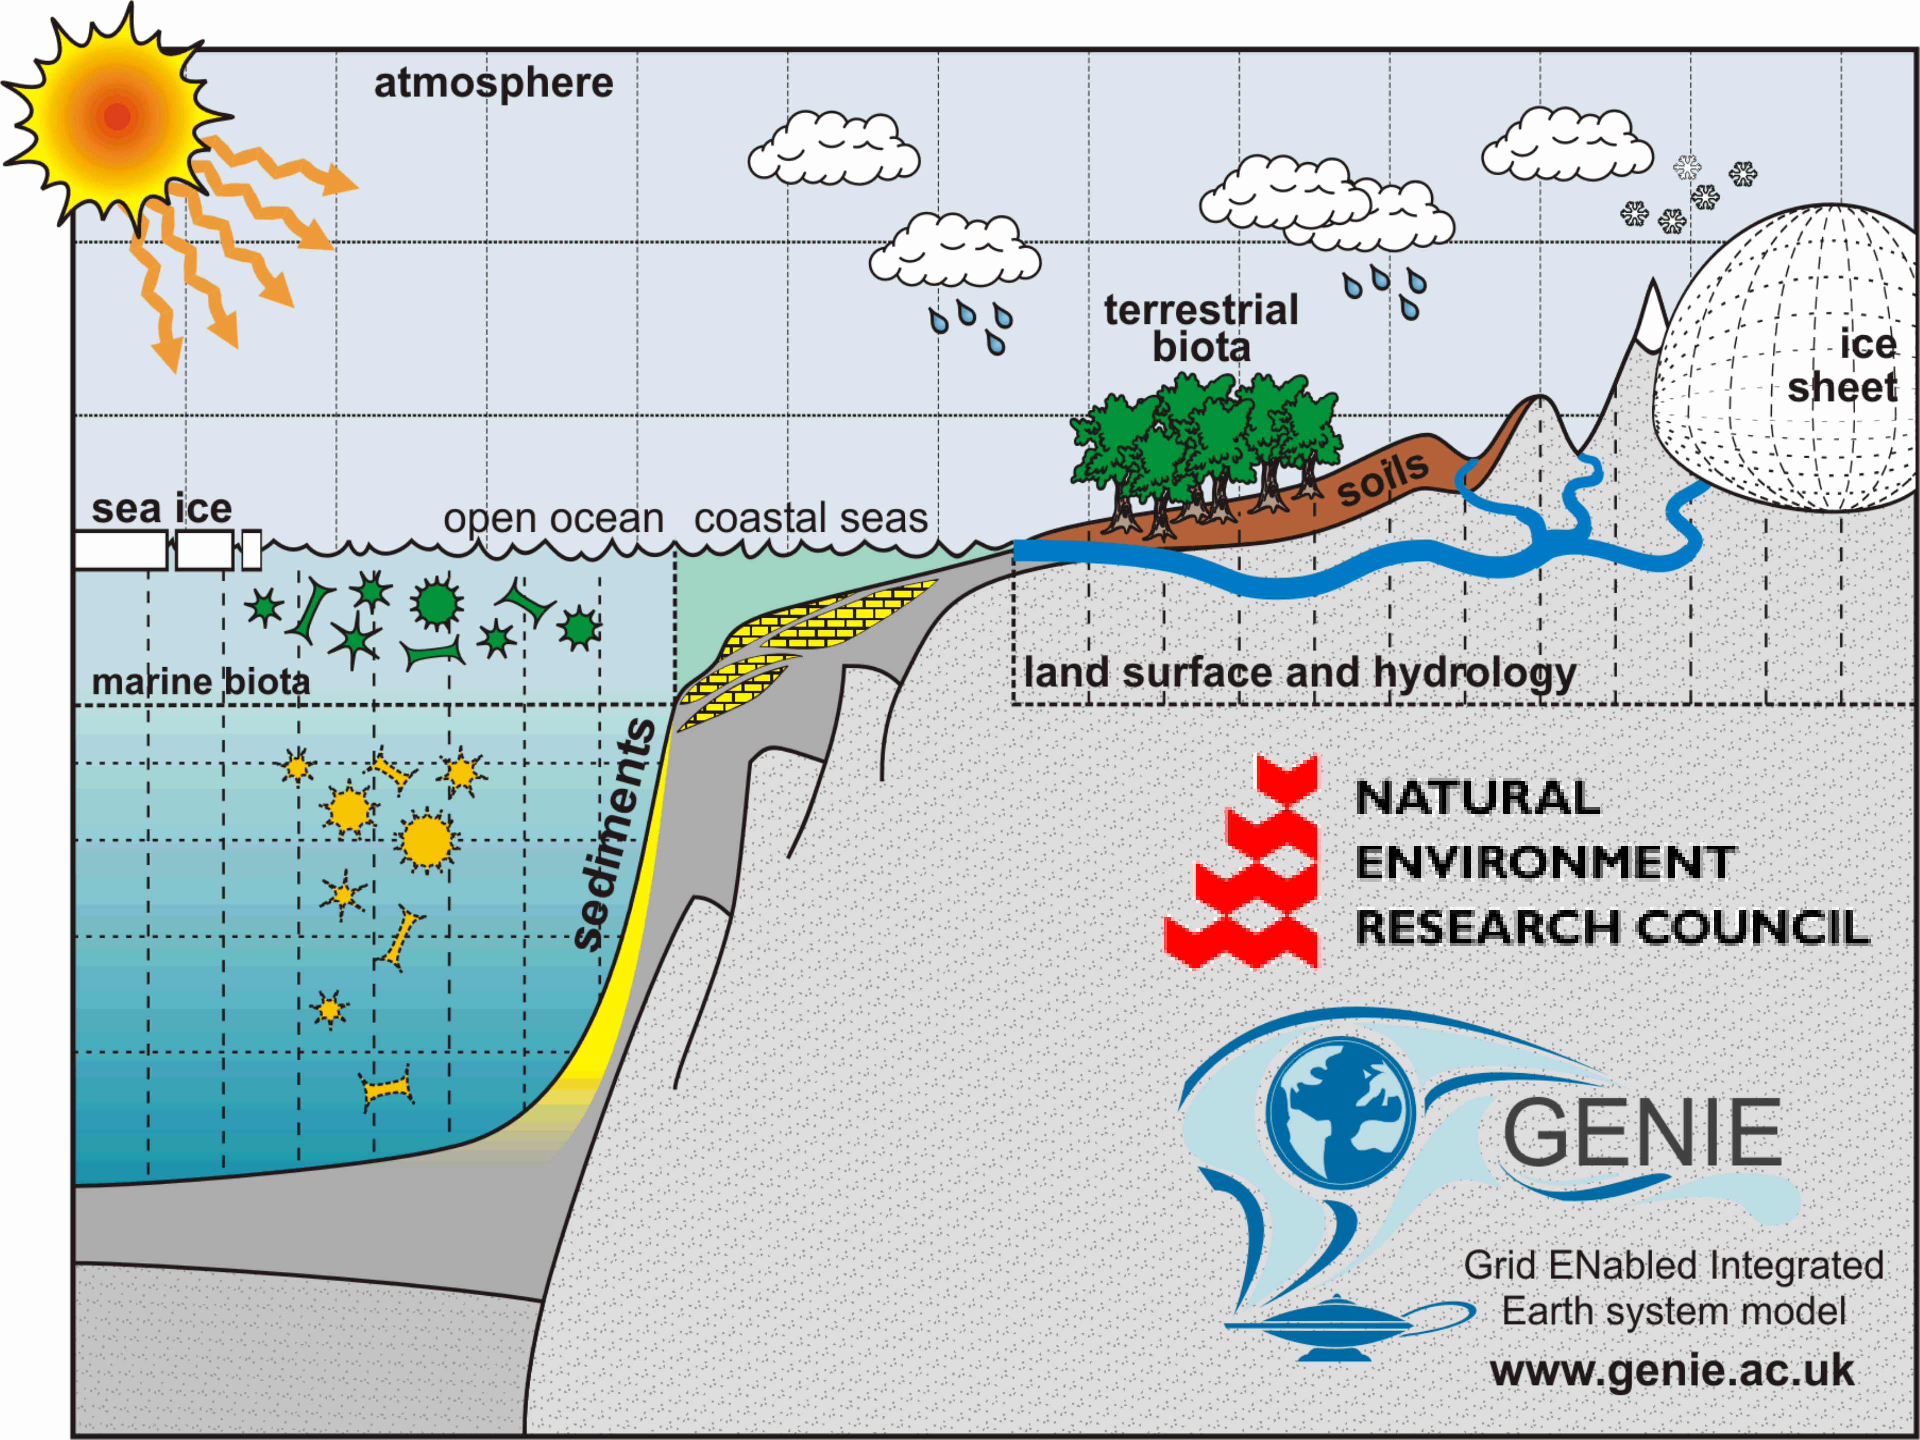
atmosphere
terrestrial biota
ice sheet
soils
sea ice
open ocean
coastal seas
marine biota
land surface and hydrology
sediments
NATURAL ENVIRONMENT RESEARCH COUNCIL
GENIE
Grid ENabled Integrated Earth system model
www.genie.ac.uk

Scientific Aims

- Orbital parameters affect incident radiation and climate
- Biological and geological processes interact with, and feedback upon, the climate (via, for instance, CO_2)

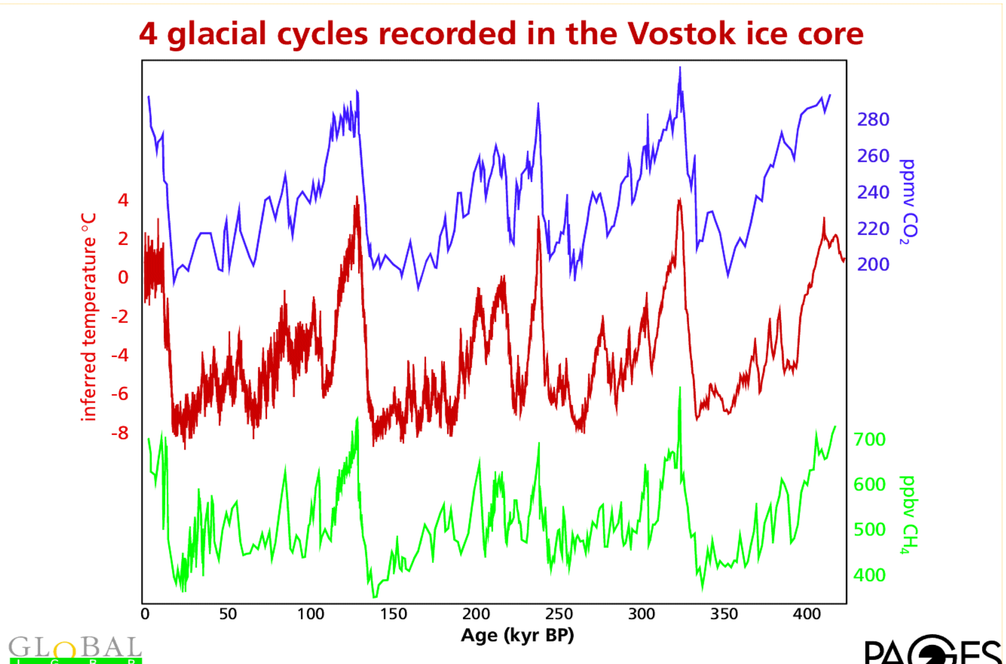

J.R. Petit et al., *Nature*, **399**, 429–36, 1999.

The mechanisms that have driven the most fundamental changes of planet Earth are not yet fully understood.

UNIVERSITY OF Southampton
School of Engineering Sciences

Parameter Estimation Problem

- Default parameters almost always sub-optimal for newly coupled models or existing models of increased resolution
- Non-linear response of a model to its parameters makes “tuning” a difficult task
- Often find conflicting design objectives (improvements in atmospheric representation can compromise ocean properties)
- Multi-objective design search and optimisation methods to Earth system models found to be effective

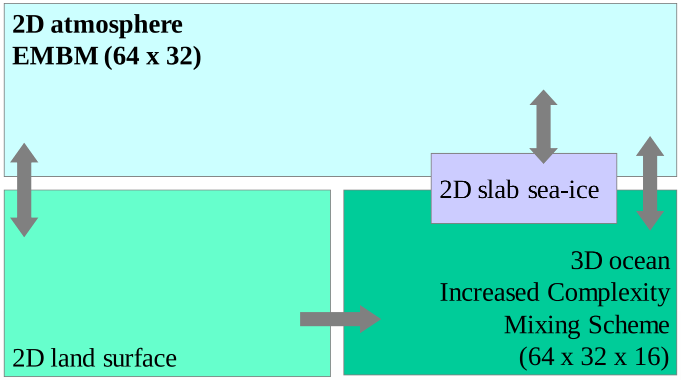

Composition of the GENIE model used in this study.
Ocean features an increased complexity mixing scheme over the default GOLDSTEIN code.

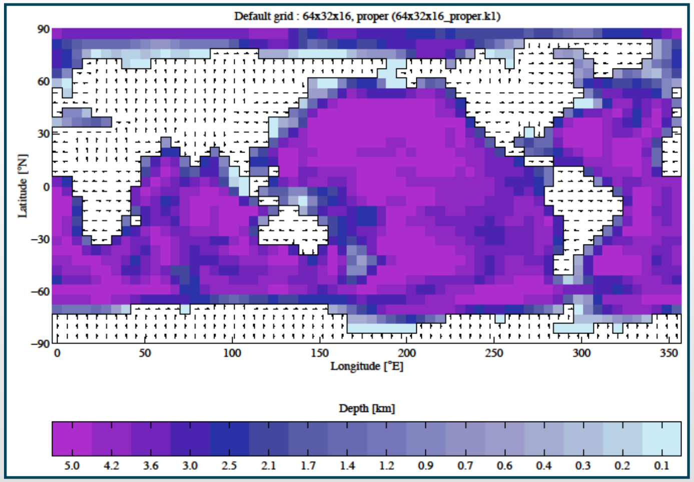

GENIE model bathymetry (depth profile) used in the 16-level model. The grid resolution is 64 x 32 x16.

UNIVERSITY OF Southampton
School of Engineering Sciences

Multi-objective Optimisation

- To compose a single objective function a decision maker must provide weighting factors for the individual targets
- The optimal or best choice for these weightings is often not known *a priori*
- Multi-objective method seeks Pareto optimal solutions
- The C-GOLDSTEIN function is easily split into its N constituents

$$f_i(x) = \sqrt{\frac{\left(s_i(x) - S_i\right)^2}{\hat{\sigma}_i^2}}, \qquad i = 1, \ldots, N$$

Multi-objective Optimisation

- Evolutionary programming and Genetic Algorithms are ideal for multi-objective methods
  - Maintain a population of solutions which "evolve" over generations of the algorithm
  - Such methods can capture Pareto optimal solutions
- Seek designs of high quality that are evenly distributed and widely spread in the objective space
- The NSGA-II algorithm is popular in the literature
  - The goal function used to drive the GA is based on relative ranking and spacing of the designs

UNIVERSITY OF Southampton
School of Engineering Sciences

Multiobjective Optimisation + Surrogates

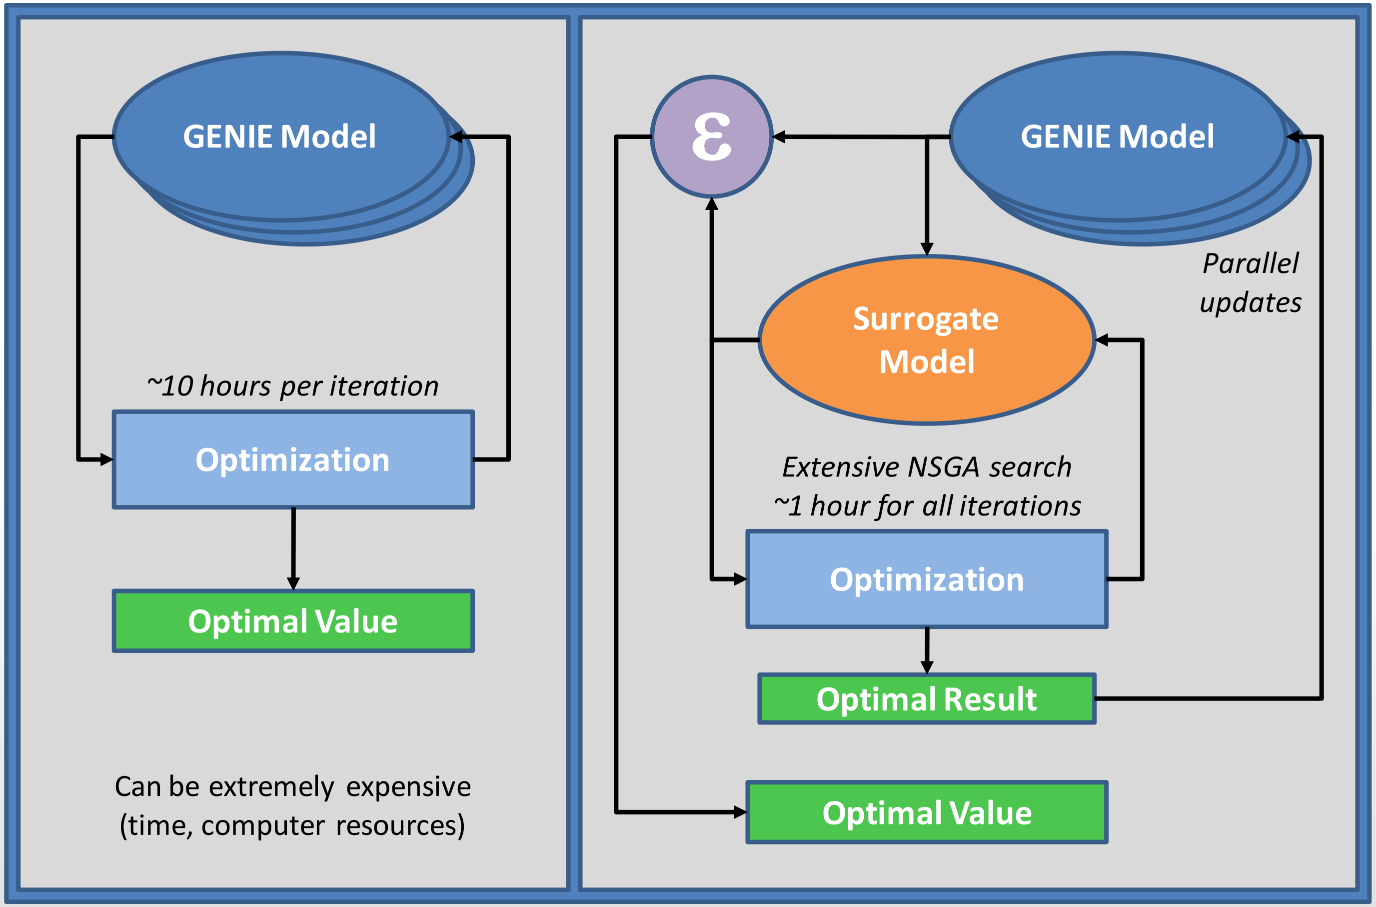

The use of surrogate models with the OptionsNSGA2 algorithm can reduce, by an order of magnitude, the total number of simulation years required for a high quality result in the calibration of a GENIE model. This approach provides surrogate models of the underlying problem which can be extensively searched at significantly less cost than the true expensive functions.

Response Surface Modelling - Kriging

- Kriging is a curve-fitting technique that originated in the field of geological surveying
- This method has been found to work very well for a wide range of multi-objective problems
- However, there is a computational cost to building the Krig models of the underlying functions
- The curvature of each Krig is controlled by a set of hyper-parameters that must themselves be tuned (optimised) to provide the best fit of the surface to the sampled data
- This is achieved by maximising a concentrated likelihood function (CLF) over a set of sampled data points
- The evaluations of the CLF involve the inversion of a matrix of correlation measures (an $O(N^3)$ operation) and consequently the tuning of the Krig can incur significant computational expense

UNIVERSITY OF Southampton
School of Engineering Sciences

Optimisation Workflow

1. Initial sampling of the underlying function (LPτ)
2. Tune the hyper-parameters of the Krig metamodel for each objective using the best training data available. High Performance Computing resource targeted for this

   High Performance Computing
3. Extensive NSGA-II searches of surrogate models
4. Select update points
   - Points from the Pareto front
   - Random points (escape from local minima)
   - Points from a small secondary NSGA-ii
   - Points of greatest Expected Improvement
   - Points of greatest RMS error in the Krigs
5. Evaluate the update points

   High Throughput Computing
6. Add the results to the existing data pool
7. Choose the best points in terms of closeness to the last Pareto front and separation in objective space
8. Rank the pool of function evaluations and extract the Pareto front
9. Return to 2

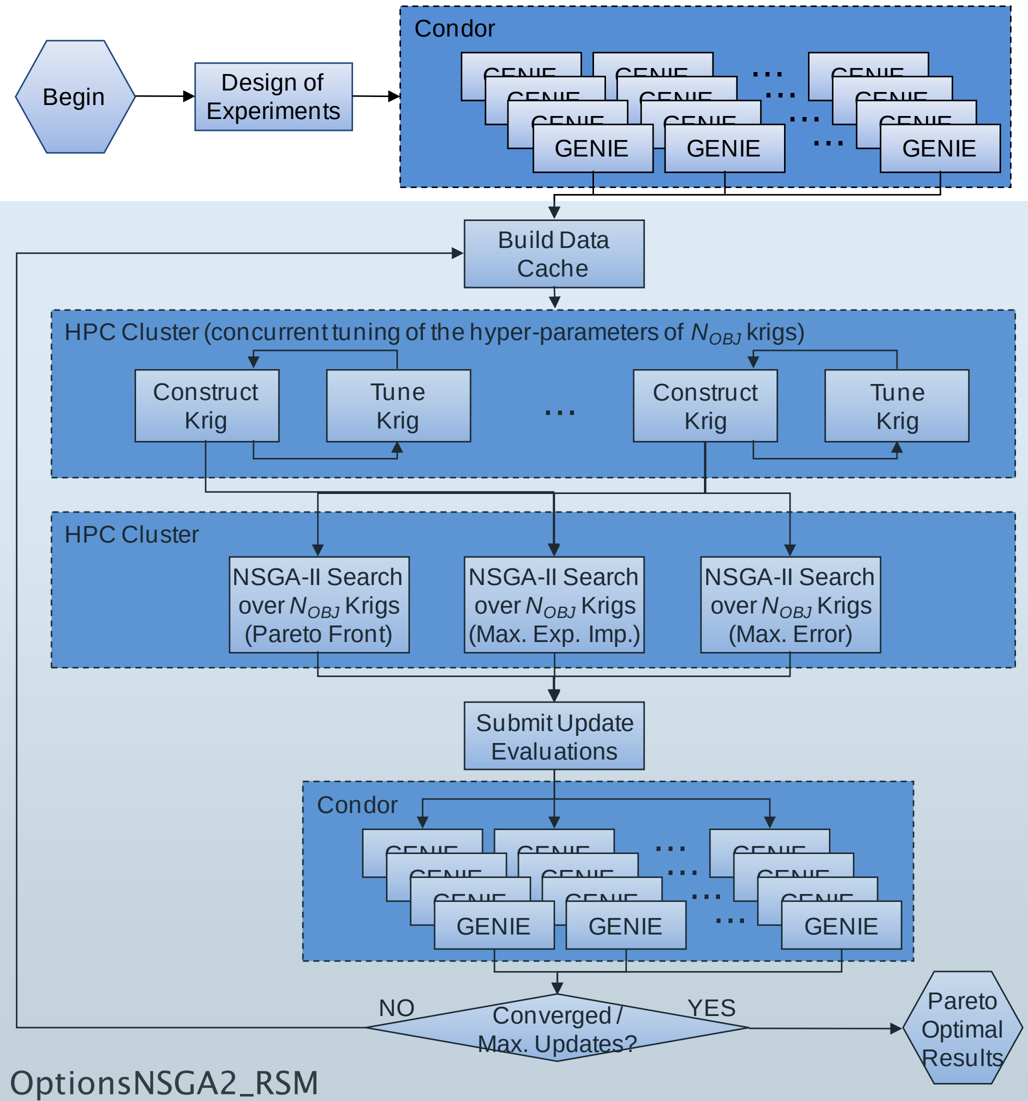

Software

- Geodise Compute Toolbox
  - Grid access from the Desktop
  - Matlab and Jython interfaces
  - Globus and Condor support
- Geodise Database Toolbox
  - Associate metadata with data
  - Programmatic and GUI access
- OptionsMatlab
  - Engineering Design Optimisation
  - Suite of multi-dimensional optimisation algorithms
- OptionsNSGA2
  - Multi-objective optimisation package
  - Augmented implementation of NSGA-II
  - Supplied courtesy of Rolls-Royce, PLC

Geodise Compute Toolbox	gd_createproxy.m	Creates a Globus proxy certificate for the user's credentials
	gd_destroyproxy.m	Destroys the local copy of the user's Globus proxy certificate
	gd_jobsubmit.m	Submits a compute job to a Globus GRAM job manager
	gd_jobstatus.m	Gets the status of a Globus GRAM job
	gd_putfile.m	Puts a remote file using GridFTP
	gd_getfile.m	Retrieves a remote file using GridFTP
	gd_rmfile.m	Deletes a remote file using GridFTP
	gd_makedir.m	Creates a remote directory using GridFTP
	gd_rmdir.m	Deletes a remote directory using GridFTP
Geodise Database Toolbox	gd_archive.m	Archives a file or data structure to the database
	gd_query.m	Query the database for data matching specified criteria.
	gd_retrieve.m	Retrieves a file or data structure from the database

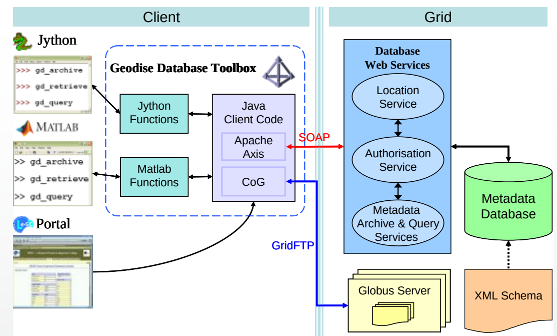

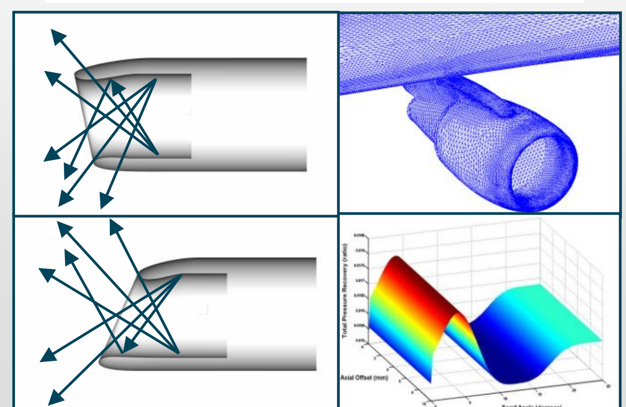

Grid Computation

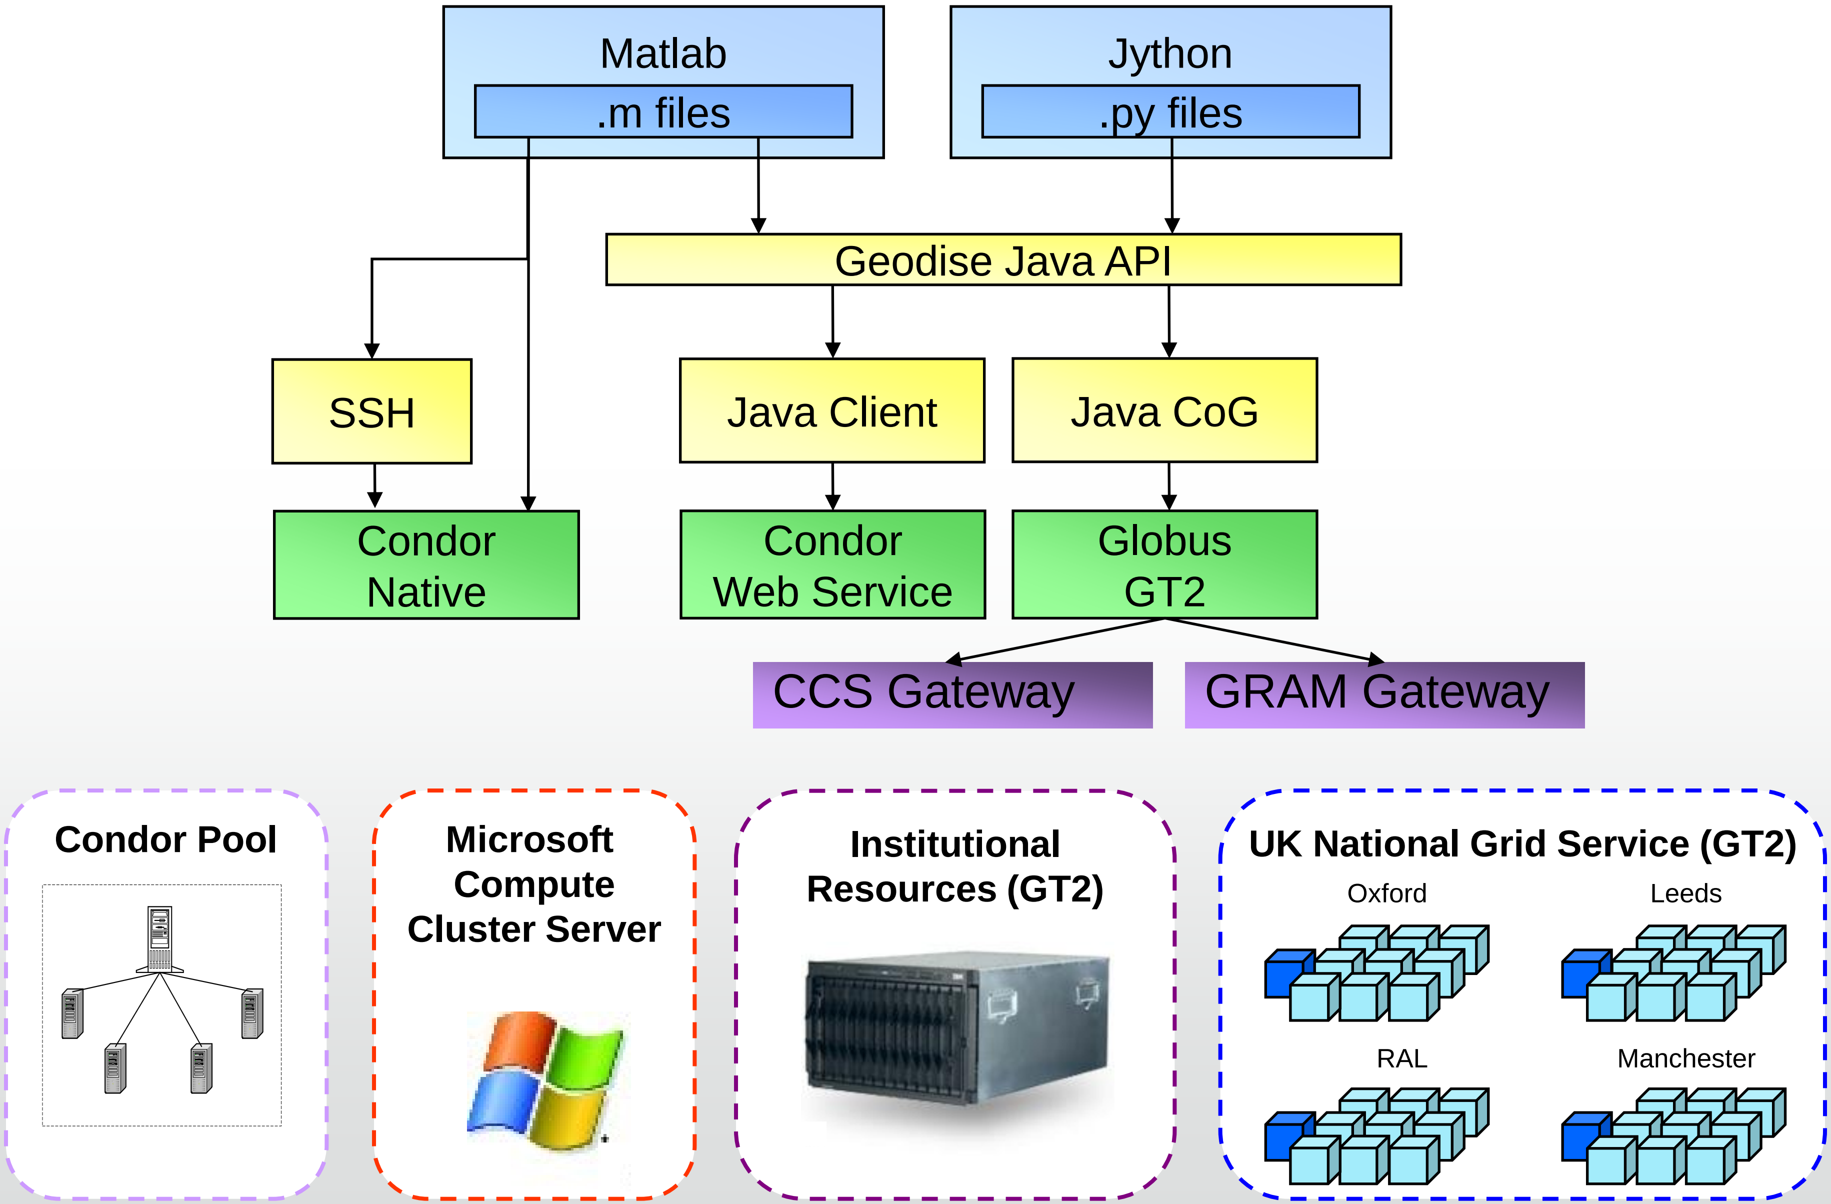

UNIVERSITY OF Southampton
School of Engineering Sciences

Matlab PSE Scripting

- OptionsNSGA2_RSM requires the user to provide two Matlab function pairs
  - Submission and post-processing functions for the hyper-parameter tuning process
  - Submission and post-processing functions for managing the GENIE simulations
- Users are free to target the most appropriate resource for their problem

```
function retrievalID=krigtune_ccs(i,USERDATA)
% Configure the location of the CCS cluster
GT2HOST='ccsglobusgateway.soton.ac.uk';
GT2DIR =['/home/andrew/tuning/',num2str(i)];
...
try
    % Create a remote directory and transfer files
    gd_makedir(GT2HOST,GT2DIR);
    gd_putfile(GT2HOST,'tune.zip',[GT2DIR,'/tune.zip'],'binary');
    ...
    % Write the input training structure to disk and transfer
    ...
    gd_putfile(GT2HOST,[jobid '/input.dat'],[GT2DIR,'/input.dat']);
    % Write the RSL string and submit the compute job
    rslstr=['&(executable=' GT2DIR '/tuneHP.bat' ')'...
        '(directory=' GT2DIR ')'...
        '(stdout=' GT2DIR '/gt2stdout.txt)'...
        '(stderr=' GT2DIR '/gt2stderr.txt)'...
        '(count=1)'...
        '(jobType=single)'...
        '(maxWallTime=' num2str(60) ')'];
    handle=gd_jobsubmit(rslstr,[GT2HOST,'/jobmanager-ccs']);
    % Return job handle
    retrievalID.handle=handle;
catch
    retrievalID.handle='failed to submit';
end
```

```
function eval=krigtune_ccs_parse2(rID)
while true,
    % Poll job status
    status=gd_jobstatus(rID.handle);
    % Handle failures
    ...
    % Process if job complete
    if status==3,
        gd_getfile(rID.GT2HOST,[rID.GT2DIR '/hyperDHC.dat'], ...
            [rID.jobid '/hyperDHC.dat']);
        % Load the tuned hyper parameters
        hyperDHC=dat2struct([rID.jobid '/hyperDHC.dat']);
        eval.OBJHYPER=hyperDHC.OBJHYPER;
    end;
    pause(checkfrequency);
end
% Clean up remote resource
gd_rmuniquedir(rID.GT2HOST,rID.GT2DIR);
```

UNIVERSITY OF Southampton
School of Engineering Sciences

Condor DAGMan

- 4,000 simulated model years required ~5 - 7 hours CPU time on the range of resource in the Condor pool
- Compute tasks of this duration at high risk of pre-emption, suspension and eviction
  - Throughput adversely affected
- University of Southampton pool exclusively Win32 machines
  - Native Condor check-pointing not available
- Use Condor DAGMan to manage simulations through a linear series of checkpoints and restarts

Condor DAGMan

- Matlab scripts auto-generate the Condor Directed Acyclic Graph (DAG) for a given number of checkpoints
- Condor DAGMan manages the submission of the DAG of compute tasks to Condor
- Pre-processing scripts manage the staging of the output files to the following task
- Some fault tolerance is provided through retries of failed tasks.

```
# Auto-generated DAG file
Job A stageA.condor
Job B stageB.condor
Job C stageC.condor
Job D stageD.condor
Job E stageE.condor
Job F stageF.condor
Job G stageG.condor
Job H stageH.condor
Job I stageI.condor
Job J stageJ.condor
Script PRE B PRE_mv_restart.sh B
Script PRE C PRE_mv_restart.sh C
Script PRE D PRE_mv_restart.sh D
Script PRE E PRE_mv_restart.sh E
Script PRE F PRE_mv_restart.sh F
Script PRE G PRE_mv_restart.sh G
Script PRE H PRE_mv_restart.sh H
Script PRE I PRE_mv_restart.sh I
Script PRE J PRE_mv_restart.sh J
PARENT A CHILD B
PARENT B CHILD C
PARENT C CHILD D
PARENT D CHILD E
PARENT E CHILD F
PARENT F CHILD G
PARENT G CHILD H
PARENT H CHILD I
PARENT I CHILD J
Retry A 3
Retry B 1
Retry C 1
Retry D 1
Retry E 1
Retry F 1
Retry G 1
Retry H 1
Retry I 1
Retry J 1
```

Compute Tasks

Task Pre-Processing

Directed Acyclic Graph (DAG)

Fault Tolerance

UNIVERSITY OF Southampton
School of Engineering Sciences

Condor Pool Usage

ISS Condor pool Condor Pool User Statistics for Apr - Windows Internet Explorer

http://www.soton.ac.uk/~condor/condor_view/UserApr.html

ISS Condor pool Condor Pool User Statistics for Apr

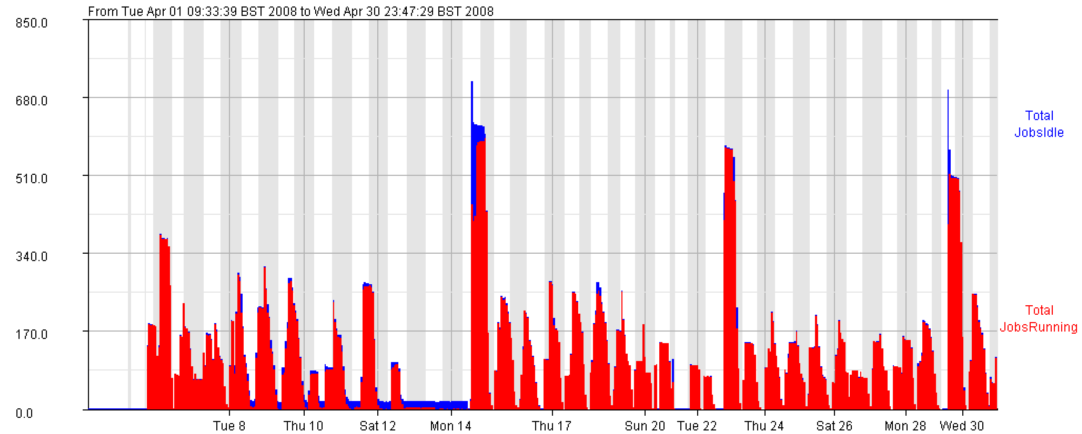

Configure... | Zoom In | Zoom Out | Reset | About

[Graph Hints: The Y-axis is number of jobs, the X-axis is time. When graph finishes updating, press "Configure.." to view different Architecture or State data. Also, you can use the mouse to draw a rectangle on the graph and then press "Zoom In". Press "Reset" to center/resize the data after Configure or when done zooming. Nighttime shows up on graph background as grey.]

User	Total Allocation Time (Hours)	JobsRunning Average	JobsIdle Average	JobsRunning Peak	JobsIdle Peak
Total	66138.8	114.9 (76.1%)	6.3 (23.9%)	590.0 (100.0%)	294.0 (100.0%)
andrew pacifica.iridis.soton.ac.uk	66117.2	128.1 (85.0%)	6.6 (15.0%)	590.0 (100.0%)	293.0 (100.0%)
gbd CC-0187.soton.ac.uk	0.0	0.0 (0.0%)	1.0 (100.0%)	0.0 (0.0%)	1.0 (100.0%)
wasadmin portal.e-science.soton.ac.uk	10.4	7.8 (90.0%)	0.2 (10.0%)	35.0 (100.0%)	1.0 (50.0%)

D.J.Baker@soton.ac.uk

Done | Internet | 100%

University of Southampton Condor pool usage in April 2008. Three OptionsNSGA2 multi-objective optimisation processes were running concurrently from the middle of the month. The management of the short individual compute tasks by Condor DAGMan keeps the optimisations in phase.

Multi-objective Optimisation Results

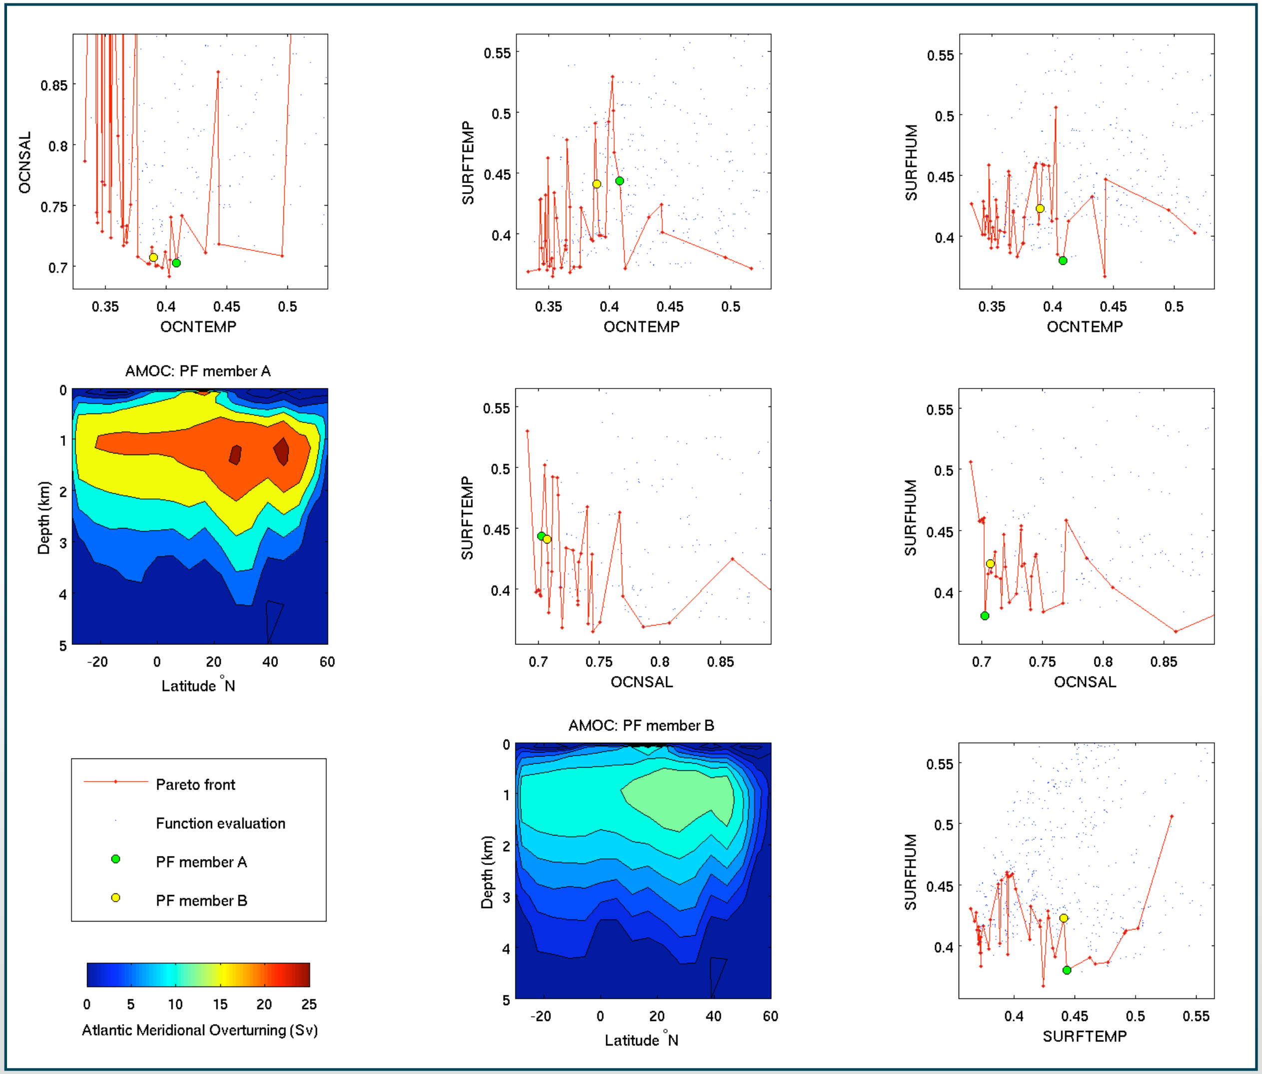

Results of the multi-objective optimisation. Top right triangle of plots shows the function evaluations projected onto 2D objective space for each pair of objectives. Two points from the Pareto front are highlighted (A,B) which have a similar score by a "traditional" single objective measure but exhibit significantly different behaviour in the Atlantic Meridional Overturning Circulation (AMOC).

Optimisation Results

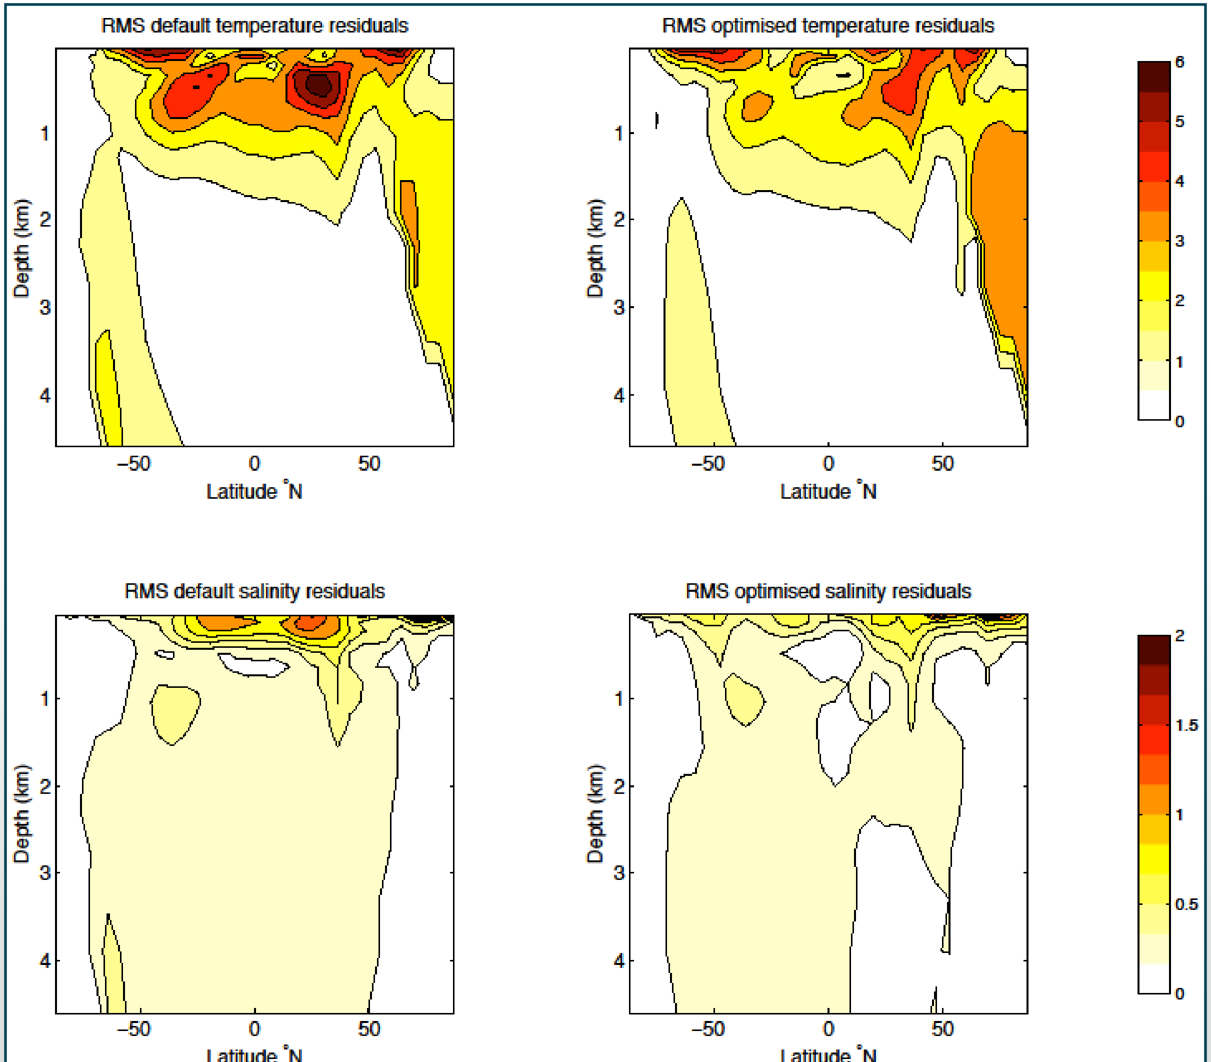

Plots of the latitudinally averaged RMS residuals for the ocean temperature and salinity profiles for the default and tuned parameters sets compared to the target observational data.

UNIVERSITY OF Southampton
School of Engineering Sciences

Conclusions

- Multi-objective optimisation
  - Avoids the need for a single weighted composite objective
  - Surrogate modelling significantly reduces number of expensive objective function evaluations
- Grid computing
  - OptionsNSGA2 implemented in Matlab Problem Solving Environment
  - Geodise software provides an interface to the Computational Grid
  - Tailor the demands of the calculation to the most appropriate resource
  - Concurrent executions of the expensive model code performed using High Throughput Computing
  - Condor DAGMan used to manage each simulation through a series of checkpoints and restarts
  - RSM hyper-parameter tuning process targeted at High Performance Computing resource

UNIVERSITY OF Southampton
School of Engineering Sciences

The GENIE Team

Principal Investigator – GENIEfy (2005-2008):

Tim Lenton – *UEA Norwich*

Research Team and Collaborators:

James Annan – *FRSGC, Japan*
Chris Armstrong – *Manchester*
Chris Brockwell – *UEA Norwich*
David Cameron – *CEH Edinburgh*
Peter Cox – *Hadley Centre (UKMO)*
Neil Edwards – *Open University*
Sudipta Goswami – *UEA Norwich*
Robin Hankin – *NOC Southampton*
Julia Hargreaves – *FRSGC, Japan*
Phil Harris – *CEH Wallingford*
Zhuoan Jiao – *Southampton e-Science Centre*
Martin Johnson – *UEA Norwich*
Eleftheria Katsiri – *London e-Science Centre*
Valerie Livina – *UEA Norwich*
Dan Lunt – *Bristol*
Richard Myerscough – *NOC Southampton*
Sofia Panagiotidi – *London e-Science Centre*
Andrew Price – *Southampton e-Science Centre*
Andy Ridgwell – *Bristol*
Ian Rutt – *Bristol*
Gethin Williams – *Bristol*
Mark Williamson – *UEA Norwich*
Gang Xue – *Southampton e-Science Centre*
Andrew Yool – *NOC Southampton*

Principal Investigator – GENIE (2002-2005):

Paul Valdes – *Bristol*

Co-Investigators / Management team:

Peter Challenor – *NOC Southampton*
Trevor Cooper-Chadwick – *Southampton e-Science Centre*
Simon Cox – *Southampton e-Science Centre*
John Darlington – *London e-Science Centre*
Rupert Ford – *Manchester*
Eric Guilyardi – *Reading*
John Gurd – *Manchester*
Richard Harding – *CEH Wallingford*
Robert Marsh – *NOC Southampton*
Tony Payne – *Bristol*
Graham Riley – *Manchester*
John Shepherd – *NOC Southampton*
Rachel Warren – *UEA Norwich*
Andrew Watson – *UEA Norwich*

UEA NORWICH
CEH
London e-Science Centre Le-SC
The Open University
UNIVERSITY OF Southampton
MANCHESTER 1824
Met Office