

Performance analysis and Optimisation of the Met Unified Model on a Cray XC30

Karthee Sivalingam, Grenville Lister, Bryan Lawrence

Abstract—The Unified Model (UM) code supports simulation of weather, climate and earth system processes. It is primarily developed by the UK Met Office, but in recent years a wider community of users and developers have grown around the code. Here we present results from the optimisation work carried out by the UK National Centre for Atmospheric Science (NCAS) for a high resolution configuration (N512 \approx 25km) on the UK ARCHER supercomputer, a Cray XC-30. On ARCHER, we use Cray Performance Analysis Tools (CrayPAT) to analyse the performance of UM and then Cray Reveal to identify and parallelise serial loops using OpenMP directives. We compare performance of the optimised version at a range of scales, and with a range of optimisations, including altered MPI rank placement, and addition of OpenMP directives. It is seen that improvements in MPI configuration yield performance improvements of between 5 and 12%, and the added OpenMP directives yield an additional 5-16% speedup. We also identify further code optimisations which could yield yet greater improvement in performance. We note that speedup gained using addition of OpenMP directives does not result in improved performance on the IBM Power platform where much of the code has been developed. This suggests that performance gains on future heterogeneous architectures will be hard to port. Nonetheless, it is clear that the investment of months in analysis and optimisation has yielded performance gains that correspond to the saving of tens of millions of core-hours on current climate projects.

Index Terms—Unified Model, UM, Climate Modelling, Cray XC30, Performance analysis, Optimisation, ARCHER

1 INTRODUCTION

The Unified Model (UM) is a simulation code which has been designed to support both predicting weather and projecting and understanding climate. It represents more than two decades of development and evolution by the UK Met Office and collaborators. During this time regular upgrades added both improved science and better performance. The UM can be configured in a range of modes, from single-column through to global mode, and with a range of horizontal and vertical resolutions. The active use of this

wide range of configurations is termed seamless prediction; the history of the evolution of the UM in this context is described in [1]. While the UM is mainly used by the UK Met Office, it is increasingly used in other organisations, both in the UK, and elsewhere. The UK academic community are one set of such users of the UM, with use of the UM underpinning a significant proportion of weather, and particularly, climate science. Over the years the academic community have fed improvements in science back into the Met Office trunk, and occasionally they have provided performance improvements - generally those associated with migrating the code to new architectures. One such set of improvements is described here. We show how a combination of improving a range of MPI settings and the

-
- *K. Sivalingam, B. Lawrence and G. Lister are with the The National Centre for Atmospheric Science - Computational Modelling Services (NCAS-CMS), University of Reading, Reading, UK
E-mail: k.sivalingam@reading.ac.uk*

TABLE 1
Hardware specification of HPC machines

	HECTOR	ARCHER	MONSooN
Machine	Cray XE6	Cray XC30	IBM 775
Compute node	Opteron	Ivybridge	Power 7
Interconnect	Gemini	Aries	IBM
File system	Lustre	Lustre	GPFS
Compute Cores	90,112	118,080	5,120
Memory (in TB)	90	318.5	10

use of OpenMP directives can make speed ups of in excess of 20% on a Cray XC-30 for some configurations of the UM.

The version of the UM discussed here is an evolution (GA6, V8.6) of the HadGEM3 global atmosphere configuration which includes a sophisticated land surface sub-model, JULES [2]. Most of the integration time is taken up by the atmospheric dynamical core. The original UM hydrostatic dynamical core of the UM described in [3] was replaced in 2002 with the New Dynamics [4]. The New Dynamics, a non-hydrostatic, semi-implicit, semi-Lagrangian scheme was used until late 2014 in operational numerical weather prediction and climate science configurations. Over the last year, variants of the model have progressively seen their dynamical cores replaced with the Even New Dynamics(ENDgame) scheme [5], [6]. The ENDgame scheme itself is described in [7], with further scientific analysis of the performance and comparisons with the New Dynamics in [8]. The results presented here are the first analysis of the computational performance of an ENDgame climate configuration to appear in the open literature.

In the UM, processes are discretized on a horizontal latitude-longitudinal grid, and over the years the grid-spacing has decreased, resulting in more simulated grid points and larger computational domains. This march to higher resolution has enabled more complex process interactions, and significantly improved scientific outcomes. The code itself is mainly FORTRAN with a few calls to

ANSI C routines, and can be run on any platform, but at anything except low resolution, requires a supercomputer. The primary versions of the shared UM code are currently maintained on an IBM Power 755 [9] supercomputer at the Met Office (MONSooN, to be replaced by a Cray XC-40 in late 2015), and codes are typically ported from that environment onto target platforms. In November 2011, the UM was ported to the HERMIT (Cray XE6) supercomputer at HLRS in Germany and optimised for the N512 high-resolution configuration (≈ 25 km grid spacing) as part of a large simulation campaign (UPSCALE, [10]). Here we discuss a new port from the MONSooN environment onto the ARCHER platform, which is aimed at extending and improving the UPSCALE optimisations.

ARCHER is a Cray XC-30, deployed at the Edinburgh Parallel Computing Centre (EPCC) as the national computing facility for environmental and engineering science. ARCHER has Intel Ivy-Bridge compute nodes [11] and an Aries interconnect [12]. In what follows we examine performance of the UM on ARCHER at three different resolutions (N96, N216 and N512, corresponding to ≈ 130 , 60 and 25 km) and compare with the MONSooN performance. We begin by discussing the default performance on MONSooN and ARCHER for these three configurations, then examine the impact of MPI communications on performance, leading to recommendations as to the best way of organising the layout of MPI ranks. We then show that there is considerable scope for performance improvement if thread imbalance can be addressed, before exploiting the Cray Reveal tool to add OpenMP directives and get a significant speedup on ARCHER. We conclude by putting this work in context with similar work on previous versions of the UM as well as other similar codes, and making suggestions as to future optimisation potential.

Table 1 lists the hardware specifications of the HPC machines that are discussed in this paper. In this study, we use the IBM Power 775 machine as a baseline for com-

TABLE 2

Standard UM jobs at different resolution. The number of columns and rows describes the grid of the global model in North-South and East-West (horizontal) direction respectively. Land points refers to the number of simulated land points. Vertical levels describes the vertical grid of the atmosphere. Timesteps refers to the number of physics timesteps per simulated day. Resolution refers to resolution of the global grid.

Jobname	Columns	Rows	Land Points	Vertical levels	Timesteps	Resolution
N96	192	144	11271	85	20 min	130 km
N216	432	324	52614	85	15 min	60 km
N512	1024	768	280592	85	10 min	25 km

parison. For MONSoon, we assume that UM jobs are fully optimized by the MetOffice. For ARCHER, optimizations discussed in the UPSCALE project are applied by default unless stated otherwise. These jobs are used as a baseline for further performance analysis. The three resolutions of the UM models used are listed in table 2 with the context discussed in [13]. We assume that the UM standard jobs will be used for performance analysis unless otherwise stated.

2 PERFORMANCE MEASUREMENT

In measuring the performance of UM models, we will use the number of model years simulated per day (M_{year}) as a metric.

$$M_{year} = \frac{1200}{T_{model}} \quad (1)$$

where T_{model} is the time taken for modelling 5 model days. In this paper, a model year is assumed to be 360 days long and T_{model} is measured from the total wallclock time ($T_{wallclock}$ in seconds) and the initial setup time ($T_{initial}$ in seconds) of a 5 model day run.

$$T_{model} = T_{wallclock} - T_{initial} \quad (2)$$

For global climate modelling, M_{year} is the most useful metric and will be compared against the number of physical cores n_{core} used. The cost in core-hours of simulating a model year (C) can be evaluated as follows

$$C = \frac{1}{M_{year}} \times n_{core} \times 24 \quad (3)$$

In this paper, C will be scaled by 1/1000 for ease of representation and 1 C will represent 1000 core-hours or a kilo core-hours.

2.1 Performance tuning parameters

Bit reproducibility is a requirement for climate modelling and is strictly enforced in all the runs. The following Cray FORTRAN compiler flags are used by default

```
-e m -s real64 -s integer64 -h O2
-hflex_mp=intolerant -h omp
```

to enforce bit reproducibility. For all the UM jobs studied the default Lustre stripe count (4) and stripe size (1MB) are used.

For MPI parallelism, we can set the number of processes in the East-West (longitude) and North-South directions (latitude). The UM uses iterative solvers to solve a Helmholtz equation and each iteration requires halos to be communicated between the MPI processes. The interpolation order used in semi-lagrangian advection and the maximum wind speed allowed in the E-W direction determines the size of the halo. The UM uses an extended halo size of up to 8 and this restricts the maximum number of MPI processes in any direction. When the MPI processes are increased, the data columns/rows per MPI process reduces. This leads to overlapping halos that makes the UM model not to bit compare.

We can do an exhaustive search to find the optimal processor decomposition that can be used. This is very

expensive and on the Cray XC30 we find that the peak performance of high resolution jobs has weak dependence on the decomposition. In further studies, we try to use a decomposition that is proportional to the number of columns and rows of a job (as listed in table 2).

OpenMP and IO server [14] support (asynchronous file IO) have been added in the recent versions. An IO server (Listener) puts the UM data writes in a FIFO queue and an IO server(Writer) processes the queue in an asynchronous manner. Parallelism is achieved by having multiple IO servers and using a threaded implementation for Listener and Writer.

For the UPSCALE project, IO servers are configured in dedicated node islands that are under populated [10]. This requires many nodes to be dedicated to IO. For MONSooN, the IO servers perform efficiently when they are spaced across the nodes running the UM [15]. IO performance benchmarks on ARCHER show that the UM runs the fastest when all the IO servers are placed on a single dedicated node.

There are many other parameters that can be tuned and are dependent on the resolution and how the physics of the model is setup. Finding all the optimal setting for the UM jobs is beyond the scope of this paper.

2.2 Threading

In the UM, parallelisation is achieved through message passing and threads. Symmetric multi threading (SMT) is supported in MONSooN and achieves significant speedup for the UM. ARCHER (Intel Ivy bridge) supports hyper-threading(HT) which can be enabled by using the '-j' option in aprun. '-j 2' enables 2 hardware threads per PE as shown below.

```
aprun -n 24 -j 2 UM.exe
```

Enabling HT (or SMT) increases the number of processing elements (PEs) available per node. This enables the UM

TABLE 3
Performance of Symmetric Multi-Threading (SMT) and Hyper-Threading technology (HT) on MONSooN and ARCHER respectively. %Speedup refers to the relative performance speedup achieved by using SMT or HT.

ARCHER			
Model years per day			
Cores	HT OFF	HT ON	%Speedup
24	0.42	0.48	15.96
48	0.81	0.90	12.20
96	1.54	1.74	13.04
192	2.80	3.06	9.44
384	4.78	4.82	0.80
MONSooN			
Model years per day			
Cores	SMT OFF	SMT ON	%Speedup
32	0.69	0.89	29.30
64	1.31	1.70	30.26
128	2.38	3.09	29.90
256	4.10	5.13	25.21

to run with twice the number of MPI tasks or threads. Table 3 shows the performance of HT and SMT on ARCHER and MONSooN respectively for a N96 job. The performance is measured using only 2 threads per MPI task and the number of MPI tasks is doubled when SMT or HT is enabled.

While SMT shows a speedup of up to 30% and shows good scaling up to 256 cores, the speedup achieved from using HT is $\approx 16\%$ but the speedup vanishes as the number of cores are increased to 384. For high resolution jobs that run on thousands of cores, we see that HT slows the UM. So in all our model runs, HT is disabled and the UM is run with 12 MPI tasks and two threads per node by default on ARCHER. For MONSooN, SMT is enabled and 32 MPI tasks and two threads are used per node by default.

3 UM PERFORMANCE

Tables 4 and 5 show the performance scaling of the UM jobs at three different resolutions on ARCHER and MONSooN respectively. Performance is measured as number of model

TABLE 4

Performance scaling of UM jobs on ARCHER. EW refers to number of PEs in East-West direction, NS - number of PEs in North-South direction, T_{model} - Wallclock time taken to complete 5 model days, n_{core} - number of physical cores, M_{year} - Model years simulated in a day, C - cost in core-hours per model year.

EW	NS	Node	T_{model}	n_{core}	M_{year}	C
N96						
4	3	1	2858	24	0.42	1.38
4	6	2	1476	48	0.81	1.43
8	6	4	771	96	1.54	1.50
8	12	8	416	192	2.80	1.65
16	12	16	237	384	4.78	1.93
24	16	32	144	768	7.14	2.59
N216						
12	8	8	2395	192	0.50	9.20
12	16	17	1239	408	0.97	10.11
24	16	33	685	792	1.75	10.85
24	32	65	407	1560	2.95	12.70
36	32	97	305	2328	3.93	14.20
48	32	129	252	3096	4.76	15.60
48	40	161	224	3864	5.36	17.31
N512						
36	24	73	2218	1752	0.54	77.72
36	36	109	1617	2616	0.74	84.60
48	36	145	1263	3480	0.95	87.90
48	48	193	1025	4632	1.17	94.96
60	48	241	885	5784	1.36	102.38
72	50	301	773	7224	1.55	111.68

years simulated in a day (M_{year}) and the cost in core-hours per model year (C).

3.1 N96

Figure 1 shows the scaling of the N96 job on ARCHER (ARC) and MONSooN (MON). ARCHER perfect scaling (ARC PS) and MONSooN perfect scaling (MON PS) is plotted for reference and refer to the respective perfect scaling that can be expected based on UM performance on a single node (or lowest number of nodes).

The UM scales to $M_{year} \approx 3$ with 128 and 192 cores on

TABLE 5

Performance scaling of UM jobs on MONSooN. EW refers to number of PEs in East-West direction, NS - number of PEs in North-South direction, T_{model} - Wallclock time taken to complete 5 model days, n_{core} - number of physical cores, M_{year} - Model years simulated in a day, C - cost in core-hours per model year.

EW	NS	Node	T_{model}	n_{core}	M_{year}	C
N96						
4	8	1	1300	32	0.89	0.86
8	8	2	696	64	1.70	0.90
8	16	4	377	128	3.09	0.99
16	16	8	218	256	5.13	1.20
N216						
4	16	2	4010	64	0.30	5.13
8	16	4	2117	128	0.57	5.42
8	32	8	1140	256	1.05	5.84
16	32	16	625	512	1.92	6.40
32	32	32	392	1024	3.06	8.03
N512						
16	32	17	4270	544	0.28	46.46
32	32	33	2425	1056	0.49	51.22
32	48	49	1793	1568	0.67	56.23
50	40	64	1500	2032	0.8	60.96

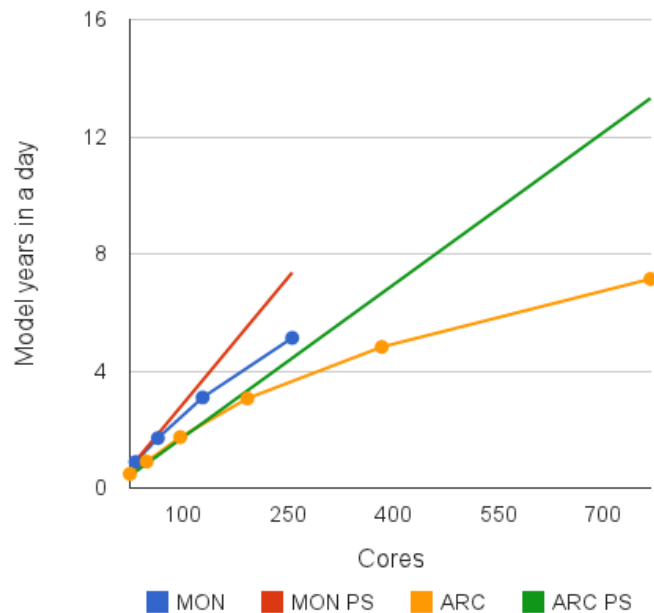


Fig. 1. Performance scaling of the N96 job on ARCHER (ARC) and MONSooN (MON). Cores refers to the actual number of physical cores used and performance is measured as number of model years simulated in a day (M_{year}). MON PS and ARC PS refers to perfect scaling that can be expected on MONSooN and ARCHER respectively.




Fig. 2. Performance scaling of the N216 job on ARCHER(ARC) and MONSooN(MON). Cores refers to the actual number of physical cores used and performance is measured as number of model years simulated in a day (M_{year}). MON PS and ARC PS refers to perfect scaling that can be expected on MONSooN and ARCHER respectively.

MONSooN and ARCHER respectively. We can infer from this that the IBM Power 7 cores are 1.5 times faster than the Intel Ivy bridge cores.

On MONSooN, the N96 job scales only up to 256 cores as extended halo size restricts the number of MPI PEs in any direction. Further scaling can be obtained by running the UM underpopulated (i.e. less than 32 MPI tasks per node). On ARCHER, the UM scales to 768 cores. Using 3 times the number of cores as on MONSooN, ARCHER has a peak performance that is 1.4 times than that on MONSooN.

3.2 N216

Figure 2 shows the scaling of the N216 job on ARCHER (ARC) and MONSooN (MON). N216 shows similar scaling performance to that of N96. Using ≈ 3.8 times the number of cores as on MONSooN, ARCHER has a peak performance that is 1.75 times than that on MONSooN. On ARCHER, IO servers are used to hide the overheads due to IO. On MONSooN N216 scales only up to 32 nodes (1024 cores) and IO servers do not improve the performance at this scale.




Fig. 3. Performance scaling of the N512 job on ARCHER(ARC) and MONSooN(MON). Cores refers to the actual number of physical cores used and performance is measured as number of model years simulated in a day (M_{year}). MON PS and ARC PS refers to perfect scaling that can be expected on MONSooN and ARCHER respectively.

Though the N216 scales to 3824 cores on ARCHER, the performance reduces almost by half compared to the perfect scaling. This implies that the cost of simulating a model year (C) doubles as the number of cores is increased from 192 to 3864.

3.3 N512

Figure 3 shows the scaling of the N512 job on ARCHER (ARC) and MONSooN (MON). IO servers are used both on ARCHER and MONSooN. Using ≈ 3.55 times the number of cores as on MONSooN, ARCHER has a peak performance that is ≈ 1.94 times than that on MONSooN. On ARCHER, the cost of simulating a model year (C) increases 30% as the number of cores is increased from 1752 to 7224. On MONSooN the cost of simulating a model year (C) increases 23% as the number of cores is increased from 544 to 2552.

Figure 4 compares the cost of simulating a model year (C) of the UM jobs on ARCHER(ARC) and MONSooN(MON) to the number of physical cores (n_{core}) and model years simulated in a day (M_{year}). In this figure, a

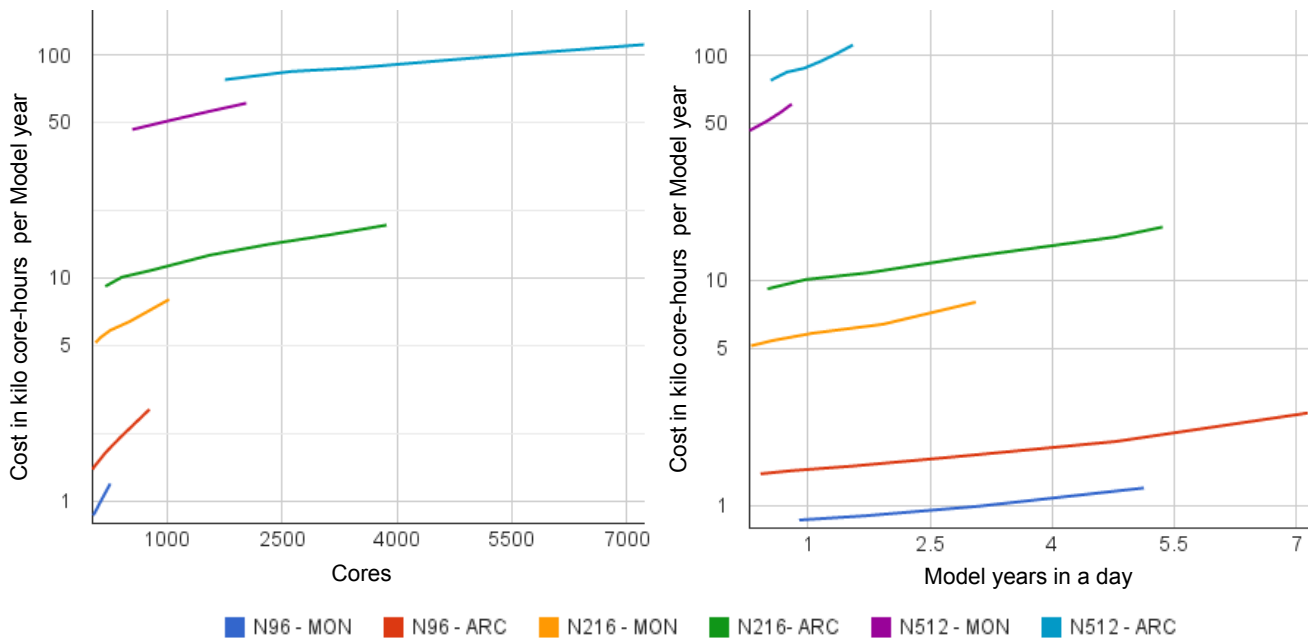


Fig. 4. Cost of simulating a model year (C) of the UM job on ARCHER(ARC) and MONSooN(MON) compared to the number of physical cores - n_{core} (left) and Model years in a day M_{year} (right).

horizontal line represents perfect scaling as this represents a constant cost of simulating a model year (C) as n_{core} or M_{year} is increased. Also a longer line represents the better performance scaling. n_{core} scales well as the resolution is increased. We are more interested in M_{year} and we can infer directly that the cost of jobs scales poorly with increase in resolution. We see a general trend where MONSooN is cost efficient when compared to ARCHER, but ARCHER scales better than MONSooN. This is based on the assumption that the usage of IBM power 7 core cost the same as the Ivy bridge core. On ARCHER N512 achieves only 1.55 model years in a day whereas we can model 5 and 7 model years of N216 and N96 respectively. Further comparing the peak performance of these jobs on ARCHER, N512 is 6.5 times costlier than N216 and 43 times costlier than N96 jobs. This clearly shows the need to analyze and optimize the high resolution N512 model.

4 PERFORMANCE ANALYSIS

Profiling the application is the first step in analyzing the performance and understanding the bottlenecks. Cray Per-

formance Analysis Tool (CrayPAT) [16] is used to profile the UM on ARCHER. Using CrayPAT, we can also profile a specific user defined function such as IO, SHMEM and more. For the UM which has a flat profile, we use the automated program analysis. This analyses the UM performance and identifies interesting areas/functions that should be instrumented.

In this analysis, we will use the N96 job running on a single node ($N96_1$) as a baseline and compare it with the profile of N512 job, running on 73 ($N512_{73}$) and 241 ($N512_{241}$) nodes. This will help us understand the bottlenecks of a high resolution job compared to a lower resolution job and also understand how the profile changes as N512 is scaled from 73 to 241 nodes.

Figure 5 shows the profile of the $N96_1$, $N512_{73}$ and $N512_{241}$ jobs in the form of pie charts. The inner pie charts show the summary of the profile and the outer pie charts reveal the finer details of the exact functions/procedures. In the profiles, UM refers to the profile of UM user code (like `um_main`, `tri_sor`) and ETC to that of all library calls (like `mpi_barrier`, `mpi_alltoall`). `UM_Others`/`ETC_Others`




Fig. 5. Pie chart showing profile of UM jobs. (a) N96 job running on a single node. (b) N512 job running on 73 nodes (c) N512 job running on 241 nodes. The inner pie chart shows the overall profile in which UM includes profile of all UM user code and ETC includes all other library calls.

includes all other UM/ETC instrumented functions that do not individually consume significant execution time.

The profiles are based on a 5 day model run. The initial setup and other related overheads are included in 'um_main'. Hence this does not include the actual simulation time. In long climate runs, this overhead ('um_main') becomes negligible.

We can infer from the figure 5 that the UM has a flat profile as even the most expensive functions (excluding

um_main) consume less than 10% of the profile. The profiles of UM jobs change with resolution as the most expensive function for $N96_1$, $N512_{73}$ and $N512_{241}$ is not the same.

4.1 MPI rank reorder

The $N96_1$ job runs on a single node and does not include any off node communication. The N512 jobs run on more than 72 nodes and require message passing between nodes. For ARCHER, the off node communications are significantly

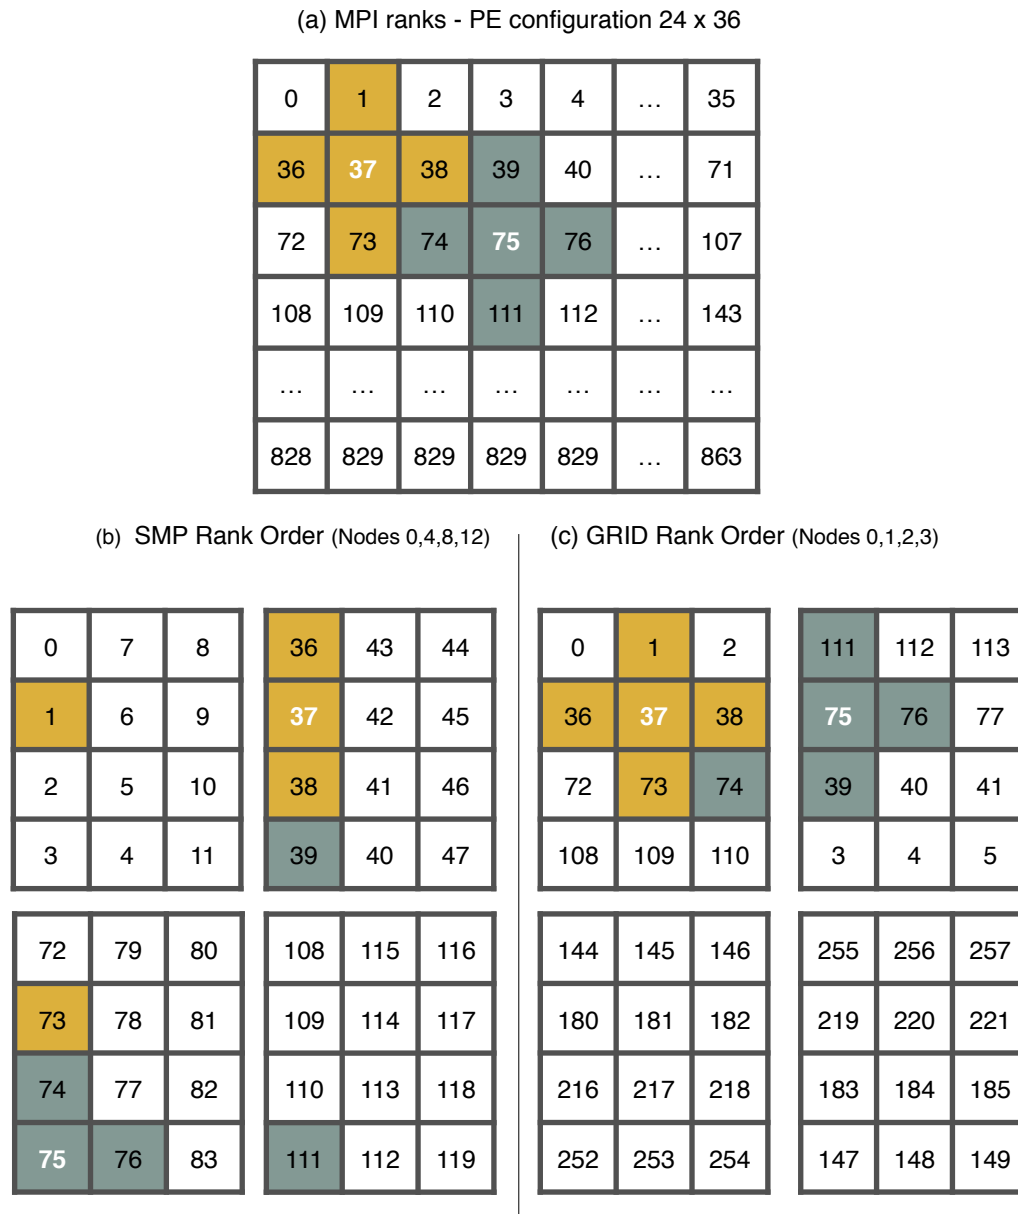


Fig. 6. (a) MPI ranks of a UM job with 24×36 PE configuration. (b) Placement of MPI ranks using the default SMP rank order. (c) Placement of MPI ranks using the GRID rank order. MPI rank 37 (and 75) along with the ranks involved in nearest neighbor communications are highlighted. Rank order is based on using 12 MPI ranks per node on ARCHER.

more expensive when compared to intra node communications.

The ETC profile represents the message passing overheads and increases from 10% on N_{96_1} to 42% on $N_{512_{73}}$. For N_{512} , ETC increases by 13% as the number of nodes is increased from 73 to 241. Also ETC consumes more the half of the total CPU time when scaled to 241 nodes. We can deduce that message passing between nodes is the most significant bottleneck for scaling of high resolution models.

CrayPAT can be configured to detect the MPI grid and the MPI communication pattern used. Based on this it suggests several MPI rank orders that will reduce off node traffic and increase the MPI bandwidth. On ARCHER the MPI ranks can be reordered by setting the environment variable `MPICH_RANK_REORDER_METHOD` to 3 and specifying the rank order in `MPICH_RANK_ORDER` file.

`MPICH_RANK_REORDER_METHOD` is set by default to 2, which is symmetric multiprocessing (SMP) style

placement. For the N512 job, CrayPAT suggests a different grid order that is based on nearest neighbor communications. This rank order is generated using a utility called `grid_order`. For example, a N512 job running with $24(EW) \times 36(NS)$ PE decomposition, CrayPAT recommends the rank order generated by the following command.

```
grid_order -R -P -c 4,1 -g 24,36
          -m 864 -n 12 -N 12 (4)
```

Here -R refers to row-major order, -m the maximum rank count, and -N the number of ranks per node. Refer to [16] for the details of the options used. We will refer to this rank reorder as GRID.

Figure 6 shows MPI ranks for UM jobs with $24(EW) \times 36(NS)$ PE configuration and how the MPI ranks can be placed on ARCHER nodes in SMP and GRID style. We assume that 12 MPI ranks will be placed per node as used in all our UM jobs. To illustrate the message passing, we can consider the nearest neighbor communications of ranks 37 and 75 (as highlighted).

For SMP, 4 nodes are involved in the communication whereas the number of nodes is reduced to 2 for GRID. We are more interested in the off-node communications as they are costlier than inter-node communications. In SMP, rank 37 has to communicate with ranks 1 and 73 which are off-node. In GRID, all the nearest neighbors of 37 reside on the same node. Similarly for rank 75, the number of MPI ranks that are off-node are reduced by half if GRID rank order is used instead of SMP.

Figure 7 compares the performance of a N512 job using a SMP and GRID rank order. GRID achieve a speedup up to 12% compared to the SMP rank order. GRID rank order results in almost a perfect speedup when scaled up to 192 nodes. In performing these measurements, IO is turned OFF. The IO performance is dependent on the Lustre file system

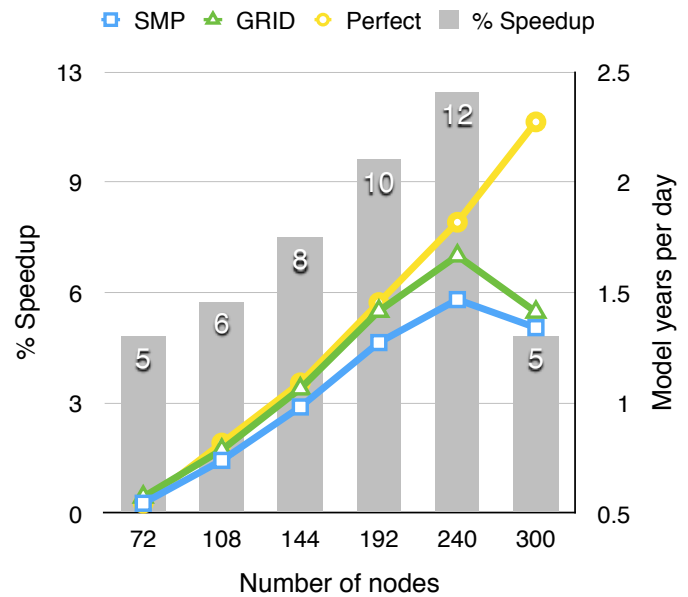


Fig. 7. Comparison of performance scaling of the N512 job with different MPI rank orders. %Speedup refers to the relative speedup achieved by using GRID rank order instead of SMP rank order.

which is a shared resource. In measuring speedup of the order of 10%, the measurements become unreliable as the shared file system performance is noisy.

4.2 Load imbalance

Another important metric is the load balance of the application. CrayPAT reports on the load balance information over all PEs and threads. It also provides finer details of load balance of different functions/routines. In CrayPAT, imbalance is measured as imbalance percentages which are relative to the set of threads or PEs. For example, if we consider the UM running with 2 OpenMP threads, an imbalance percentage of 50% implies that one thread is idle for 50% of the time when the other thread is busy.

Figure 8 shows the imbalance percentage of $N96_1$, $N512_{73}$ and $N512_{241}$. All these measurements are based on a 5 day model run using 12 MPI tasks and 2 OpenMP threads per node. For $N96_1$, 'mpi_allreduce', 'mpi_scatterv' and 'mpi_alltoall' have no imbalance but increases up to 64% for $N512_{241}$. This further emphasizes the need to reduce message passing overheads.




Fig. 8. Imbalance percentage of UM jobs : N96 job running on a single node; N512 job running on 73 nodes; N512 job running on 241 nodes; Thread imbalance percentage are relative to the set of threads used. Imbalance percentage of UM functions are relative to set of PEs.

In SMP based supercomputers, we need thread based parallelism to improve the scaling of the UM and reduce MPI overheads which increase to more than 50% of the profile as the N512 model is scaled to 241 nodes. For example, using n threads will allow the UM to scale to n times the number of cores as that of a MPI only version and also reduce the number of MPI packets communicated between nodes by half. The number of MPI tasks in the UM is limited by the extended halo size. Also the extended halos increase the actual memory consumed. . This clearly shows the need for better thread performance to improve the scaling of high resolution models.

On ARCHER, if we assume that the OpenMP implementation is 100% efficient and scales well, we can ideally set the number of threads to 12 and MPI tasks to 2. ARCHER nodes have two, 12-core NUMA regions and the 12 cores of a single NUMA region have fast access to the shared local memory. $N96_1$ has only 10% MPI overhead and the 46% thread imbalance of $N96_1$ can be attributed directly to

loops that are not thread parallelised. So based on Amdahl's law, we can expect poor scaling of the UM as the number of threads is increased. Functions 'glue_conv' and 'ls_ppnc' have higher imbalance compared to other UM functions and can be improved by parallelising loops.

5 OPENMP OPTIMISATION

OpenMP provides a standard and portable way of parallelising loops. This involves scoping the loop variable (as shared, private ...) and inserting OpenMP directives before a loop. The UM has a flat profile with thousand of serial loops that can be parallelised. Parallelising all the loops is expensive and careful consideration is required to ensure data consistency. Also race conditions involving parallel threads are hard to debug.

Cray Reveal is an integrated performance analysis and code optimisation tool. It provides loop analysis and scoping of serial loops and suggests OpenMP directives that can be inserted to a loop. Performance data collected during ex-

TABLE 6

Performance scaling of UM jobs using increasing number of OpenMP threads on ARCHER. Wallclock time refers to the time taken to complete 2 model days. $MPI \times THR$ refers to the number of MPI tasks \times number of OpenMP threads used per node. $UM_{8.6}$ refers to the original UM code and UM_{Reveal} to the code with new OpenMP directives. %Speedup is measured as a relative performance improvement achieved by adding new OpenMP directives.

<i>M_{year} with</i>				
Node	$MPI \times THR$	$UM_{8.6}$	UM_{Reveal}	%Speedup
N96 (PE config - 4 \times 3)				
1	12 \times 2	0.408	0.429	4.9
2	6 \times 4	0.568	0.625	9.1
3	4 \times 6	0.577	0.712	18.9
N512 (PE config - 36 \times 24)				
72	12 \times 2	0.494	0.518	4.6
144	6 \times 4	0.667	0.731	8.9
216	4 \times 6	0.734	0.872	15.9

ecution by CrayPAT can easily be attached to Cray Reveal to identify the profile of loops. This will help us in prioritizing the loops that consume more CPU time.

The tool also shows the compiler optimizations that have been applied. Even though this tool is user friendly, it requires knowledge of OpenMP to resolve conflicts, race conditions and scoping issues. It works only with the Cray compiling environment. This tool does not provide support for parallel regions, task based parallelism, barrier, critical or atomic regions.

5.1 UM - Reveal on ARCHER

UM has a flat profile and has thousands of loops that can be parallelised. Parallelising all these loops is beyond the scope of this paper. As a case study, 2389 serial loops are parallelised by adding OpenMP directives as suggested by Cray Reveal. default(none) option is used for all the newly added directives. For Fortran array notation expressions like 'for all' and 'where' statements, are parallelised using the following directives

```
!$OMP PARALLEL WORKSHARE
!$OMP END PARALLEL WORKSHARE
```

TABLE 7

Performance scaling of a N96 job (PE config - 4 \times 4) using increasing number of OpenMP threads on MONSooN. Refer to table 6 for definition of Wallclock time, $MPI \times THR$ and %Speedup.

<i>M_{year} using</i>				
Node	$MPI \times THR$	$UM_{8.6}$	UM_{Reveal}	%Speedup
1	16 \times 2	0.276	0.277	0.7
2	8 \times 4	0.370	0.373	0.9
4	4 \times 8	0.427	0.374	-14.3

Bit reproducibility is enforced to test for correctness of OpenMP directives.

Table 6 shows the performance scaling of UM jobs when the number of threads is increased on ARCHER. In all the performance measurements, each thread is assigned to a PE which means PEs are not oversubscribed. For example, in case of ARCHER nodes, when the number of OpenMP threads is increased from 2 to 4, the number of MPI tasks per node is reduced from 12 to 6.

In this section, the original UM code will be referred to as $UM_{8.6}$ (8.6 is the original UM version number used) and the UM code with Cray Reveal changes will be referred to as UM_{Reveal} . Hyperthreading and symmetric multi-threading are switched off for ARCHER and MONSooN respectively, to enable easier comparison.

On ARCHER, UM_{Reveal} achieves a speedup of 18.9% and 15.9% for N96 and N512 jobs respectively with 6 OpenMP threads. The fall in speedup for N512 job compared to N96, is due to additional MPI overheads. For both N96 and N512 jobs, the speedup increases by more than 3 times as the number of threads are trebled. This is a significant improvement and can be further improved by adding OpenMP directives to thousands of other serial loops.

5.2 UM - Reveal on MONSooN

UM_{Reveal} is based on the OpenMP standard and can be ported easily to most other supercomputers. The IBM compiler supports OpenMP and UM_{Reveal} is easily ported to




Fig. 9. %Speedup of UM functions on MONSooN and ARCHER. %Speedup is measured as the relative improvement achieved by adding OpenMP regions. Performance of UM is measured using 6 MPI and 4 OpenMP threads per node. In X-axis labels, the calltree of the function is specified using ':' delimiter. For example in AS:AP1, 'Atm Step' calls function 'Atmos Physics 1'.

MONSooN. Table 7 shows the performance scaling of UM jobs when the number of threads is increased on MONSooN. On MONSooN, UM_{Reveal} does not result in significant speedup and when 8 threads are used, UM_{Reveal} slows the performance by 14.3%.

Figure 9 shows a closer look at the relative speedup achieved for different functions on MONSooN and ARCHER using 4 OpenMP threads. %Speedup is measured as the relative improvement achieved by using UM_{Reveal} instead of $UM_{8.6}$. In ARCHER, all the UM functions show considerable speedup whereas on MONSooN, 'Atmos Physics 1' and 'Cloud Simulator' routines slow down considerably.

Even though OpenMP provides a standard and portable way for implementing thread based parallelism, the performance improvements are not portable and depend on the hardware architecture. One major difference we found is that ARCHER has a unified L3 cache whereas MONSooN does not. This may be the reason for the significant difference in performance. Since MONSooN is being replaced by a Cray XC40 machine, this is not investigated further.

5.3 Related work

The performance of UM is dependent not only on threads and MPI processes, but also on the resolution, the physics and the algorithms used. This requires careful modelling of UM performance to study the impact of all these parameters. Researchers have modelled the performance of the UM to study scalability with increased resolution and core counts [17], [18]. Historically the UM has been ported by Met Office to different architectures like Cray T3E, NEC SX-6, NEC SX-8, IBM Power 6 and Power 7 [19]. They studied components that affect the scalability and found IO and the Helmholtz solver to be the biggest obstacle to UM scaling. Also these components interact with each other to a different extent and this made model performance evaluation difficult.

Researchers at the Australian National University have developed efficient profiling methodology and scalability analysis of the UM at different resolutions [20]. They have identified that the high resolution N512 L70 job scales up to 2048 cores but is affected by load imbalance and MPI communications. On an Intel Sandy Bridge based cluster,

they identified that hybrid OpenMP/MPI will provide the best opportunity for optimisations. These findings are based on UM version 7.5 (New Dynamics) released in 2010 and we see similar trends as discussed in this paper.

Met Office researchers have studied the scaling of high resolution jobs on a IBM Power 6 and an Intel Nehalem cluster (Juropa). They found that the scaling can be improved by using IO servers, OpenMP and identified further scope for improvement [21]. This study reports $\approx 6\%$ speedup achieved using 2 OpenMP threads and SMT. Also the GungHo project [22] which is research collaboration between the Met Office, NERC funded researchers and STFC Daresbury strongly supports introducing thread based parallelism to improve scalability by reducing the cost of MPI communications. Also they emphasise the need to exploit thread parallelism that will be available in future exascale machines. The OpenMP optimisations we have discussed in this paper make good strides in improving the thread parallelism of UM.

As part of the UPSCALE project, Tom Edwards at Cray studied the performance of the UM (version 8.0, N512 L80, New Dynamics) on HERMIT, a Cray XE6 machine [23]. In his report, the UM was also optimised by reordering the ranks so that the IO servers are placed in separate dedicated nodes. This and other optimisations resulted in 39% improvement in runtime performance of the UM for 14% increase in number of cores. This reports also suggests the need to extend OpenMP based thread parallelism which severely affects the scaling of the UM.

In [25], the authors prescribe that the MPI ranks should be reordered to match the application communication pattern with the underlying hardware. Also [26] found MPI reordering to be a promising optimisation for unstructured CFD code. In [24], the researchers studied the performance of a simple shallow water model on a Cray XE6 machine (HECTOR). They suggest that MPI ranks to physical cores should be mapped such that the off-node data transfer vol-

ume in a nearest neighbour communication can be reduced. This is in agreement with our findings based on the performance of the UM on ARCHER, a Cray XC30 machine.

Similar efforts have been made to study the performance of other weather models. Weather and Research Forecasting model (WRF) is used all around the world and its performance has been studied in detail on Blue Gene/L [27], Blue Gene/P [28], Cray XE6 [29], Cray XT [30] and many other machines. ECMWF (European Centre for Medium-Range Weather Forecasts) have ported IFS (Integrated Forecasting System) from IBM Power 7 to Cray XC30 [31] and identified performance testing to be an important tool in meeting application challenges. As a part of ECMWF's scalability project, IFS has been ported to GPU, Intel Xeon Phi and Intel Xeon Haswell and its performance compared [32].

6 CONCLUSION

We have analysed of the performance of the UM at three different resolutions on ARCHER using CrayPAT tools and compared to that of MONSooN. Even though IBM Power 7 cores are more powerful compared to Intel Ivy bridge, ARCHER scales better compared to MONSooN. Performance analysis shows that MPI communication and thread imbalance affect the scaling of high resolution UM jobs. Reordering the MPI ranks using GRID rank order speed up the UM by up to 12% compared to default SMP rank order. Using Cray Reveal, new OpenMP directives are added to the UM that results in improved speedup of up to 16% on ARCHER whereas it slows the performance on MONSooN. These performance optimisations have resulted in savings of tens of millions of core-hours in current climate projects. We can further improve the UM performance by allowing both threads to make MPI calls and by adding OpenMP directives to loops that can be parallelised.

For high resolution (N512) jobs, thread imbalance increases from 66% to 79% as the number of nodes is increased from 73 to 241(as shown in figure 8). This is because only

thread0 is involved in message passing. The imbalance can be reduced by allowing both threads to make MPI calls. This requires not only using a thread-safe MPI implementation [33], but also careful coding to ensure data consistency and avoid race conditions. Exploring this beyond the scope of this paper. Even though we parallelised 2389 serial loops using Cray Reveal, there are thousands of loops that can still be parallelised. Also parallel regions can be added to reduce threading overheads. This will further improve the scaling and efficiency of the UM performance.

ACKNOWLEDGMENTS

The work done here was funded by a contract to the National Center for Atmospheric Science - Computational and Modelling Services group (NCAS-CMS) at the University of Reading by the Met Office via the Joint Weather and Climate Research Programme (JWCRP). We are grateful to colleagues at the Met Office and at CMS for assistance and advice throughout this project, in particular Matthew Mizielinski, Rosalyn Hatcher and Simon Wilson. We thank EPCC for their help in porting and supporting the UM jobs on ARCHER.

REFERENCES

- [1] A. Brown, S. Milton, M. Cullen, B. Golding, J. Mitchell, and A. Shelly, "Unified Modeling and Prediction of Weather and Climate: A 25-Year Journey," *Bulletin of the American Meteorological Society*, vol. 93, no. 12, pp. 1865–1877, May 2012. [Online]. Available: <http://journals.ametsoc.org/doi/abs/10.1175/BAMS-D-12-00018.1>
- [2] D. N. Walters, K. D. Williams, I. A. Boutle, A. C. Bushell, J. M. Edwards, P. R. Field, A. P. Lock, C. J. Morcrette, R. A. Stratton, J. M. Wilkinson, M. R. Willett, N. Bellouin, A. Bodas-Salcedo, M. E. Brooks, D. Copesey, P. D. Earnshaw, S. C. Hardiman, C. M. Harris, R. C. Levine, C. MacLachlan, J. C. Manners, G. M. Martin, S. F. Milton, M. D. Palmer, M. J. Roberts, J. M. Rodríguez, W. J. Tennant, and P. L. Vidale, "The Met Office Unified Model Global Atmosphere 4.0 and JULES Global Land 4.0 configurations," *Geosci. Model Dev.*, vol. 7, no. 1, pp. 361–386, Feb. 2014. [Online]. Available: <http://www.geosci-model-dev.net/7/361/2014/>
- [3] M. J. P. Cullen and T. Davies, "A conservative split-explicit integration scheme with fourth-order horizontal advection," *Quarterly Journal of the Royal Meteorological Society*, vol. 117, no. 501, pp. 993–1002, Jul. 1991. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1002/qj.49711750106/abstract>
- [4] T. Davies, M. J. P. Cullen, A. J. Malcolm, M. H. Mawson, A. Staniforth, A. A. White, and N. Wood, "A new dynamical core for the Met Office's global and regional modelling of the atmosphere," *Quarterly Journal of the Royal Meteorological Society*, vol. 131, no. 608, pp. 1759–1782, 2005. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1256/qj.04.101/abstract>
- [5] D. N. Walters *et al.*, "The met office unified model global atmosphere 6.0 and jules global land 6.0 configurations," *Geoscientific Model Development*, pp. 523, 524, 525, 529, 533, 535, 538, in preparation 2014.
- [6] T. Melvin and M. Gross, *ENDGame Dynamical Core*, Met Office, https://earthsystemcog.org/site_media/docs/DCMIP_19_mentor_ENDGame_Melvin_Gross.pdf, 2012.
- [7] N. Wood, A. Staniforth, A. White, T. Allen, M. Diamantakis, M. Gross, T. Melvin, C. Smith, S. Vosper, M. Zerroukat, and J. Thuburn, "An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the deep-atmosphere global non-hydrostatic equations," *Quarterly Journal of the Royal Meteorological Society*, vol. 140, no. 682, pp. 1505–1520, Jul. 2014. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1002/qj.2235/abstract>
- [8] N. J. Mayne, I. Baraffe, D. M. Acreman, C. Smith, N. Wood, D. S. Amundsen, J. Thuburn, and D. R. Jackson, "Using the UM dynamical cores to reproduce idealised 3-D flows," *Geosci. Model Dev.*, vol. 7, no. 6, pp. 3059–3087, Dec. 2014. [Online]. Available: <http://www.geosci-model-dev.net/7/3059/2014/>
- [9] D. Quintero *et al.*, *IBM Power Systems 775 for AIX and Linux HPC Solution*, ser. IBM redbooks. IBM Redbooks, 2014. [Online]. Available: http://books.google.co.uk/books?id=0_LHAgAAQBAJ
- [10] M. S. Mizielinski *et al.*, "High-resolution global climate modelling: the upscale project, a large-simulation campaign," *Geoscientific Model Development*, vol. 7, no. 4, pp. 1629–1640, 2014. [Online]. Available: <http://www.geosci-model-dev.net/7/1629/2014/>
- [11] *Intel 64 and IA-32 Architectures Optimization Reference Manual*, Intel, <http://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-optimization-manual.pdf>, September 2014.
- [12] B. Alverson *et al.*, "Cray xc series network," Cray, <http://www.cray.com/Assets/PDF/products/xc/CrayXC30Networking.pdf>, Tech. Rep., 2012.
- [13] S. Bell, "Strategy for operational nwp modelling sys-

- tems," Met Office, http://www.metoffice.gov.uk/media/pdf/c/4/MOSAC_17.5_SBell.pdf, Tech. Rep. PAPER 17.6, 2012.
- [14] J. M. Dennis *et al.*, "An application-level parallel i/o library for earth system models," *International Journal of High Performance Computing Applications*, vol. 26, no. 1, pp. 43–53, 2012. [Online]. Available: <http://hpc.sagepub.com/content/26/1/43.abstract>
- [15] P. Selwood. (2012) The met office unified model i/o server. https://verc.enes.org/computing/hpc-collaborations/parallel-i-o/workshop-scalable-io-in-climate-models/presentations/Unified_Model_IO_Server_Paul_Selwood.ppt.
- [16] *Using Cray Performance Measurement and Analysis Tools*. <http://docs.cray.com/books/S-2376-622/S-2376-622.pdf>: Cray, 2014, no. S-2376-622.
- [17] D. J. Kerbyson and P. W. Jones, "A performance model of the parallel ocean program," *Int. J. High Perform. Comput. Appl.*, vol. 19, no. 3, pp. 261–276, Aug. 2005. [Online]. Available: <http://dx.doi.org/10.1177/1094342005056114>
- [18] A. Osprey *et al.*, "Performance modelling for climate models," http://www.met.reading.ac.uk/~annette/performance/poster_ncas_5July2010.pdf, July 2010.
- [19] P. Selwood *et al.*, "Strategies for improving the scalability of the um in response to changing computer architectures." <http://research.metoffice.gov.uk/research/nwp/publications/mosac/doc-2009-10.pdf>, November 2009.
- [20] P. Strazdins *et al.*, "Profiling methodology and performance tuning of the met office unified model for weather and climate simulations," in *Parallel and Distributed Processing Workshops and Phd Forum (IPDPSW), 2011 IEEE International Symposium on*, May 2011, pp. 1322–1331.
- [21] A. Malcolm, M. Glover, and P. Selwood, "Scalability of the met office unified model," *14th ECMWF workshop*, 2010.
- [22] R. Ford *et al.*, "Gung ho: A code design for weather and climate prediction on exascale machines," *Advances in Engineering Software*, 2014.
- [23] T. Edwards, "Optimising upscale on hermit," Cray, Tech. Rep., 2012.
- [24] A. Osprey and other, "A benchmark-driven modelling approach for evaluating deployment choices on a multicore architecture," in *Proceedings of the International Conference on Parallel & Distributed Processing Techniques & Application (PDPTA'13)*, 2013. [Online]. Available: <http://centaur.reading.ac.uk/34184/>
- [25] G. Mercier and E. Jeannot, "Improving mpi applications performance on multicore clusters with rank reordering," in *Proceedings of the 18th European MPI Users' Group Conference on Recent Advances in the Message Passing Interface*, ser. EuroMPI'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 39–49. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2042476.2042483>
- [26] B. Brandfass, T. Alrutz, and T. Gerhold, "Rank reordering for {MPI} communication optimization," *Computers and Fluids*, vol. 80, no. 0, pp. 372 – 380, 2013, selected contributions of the 23rd International Conference on Parallel Fluid Dynamics ParCFD2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S004579301200028X>
- [27] D. J. Kerbyson, K. J. Barker, and K. Davis, "Analysis of the weather research and forecasting (wrf) model on large-scale systems," in *Parallel Computing: Architectures, Algorithms and Applications*, vol. 15. IOS Press, 2007, pp. 89–98.
- [28] P. Malakar *et al.*, "Performance evaluation and optimization of nested high resolution weather simulations," in *Proceedings of the 18th International Conference on Parallel Processing*, ser. Euro-Par'12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 805–817. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-32820-6_80
- [29] P. Johnsen *et al.*, "Petascale wrf simulation of hurricane sandy deployment of ncsa's cray xe6 blue waters," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, ser. SC '13. New York, NY, USA: ACM, 2013, pp. 63:1–63:7. [Online]. Available: <http://doi.acm.org/10.1145/2503210.2503231>
- [30] A. R. Porter and M. Ashworth, "Configuring and optimizing the weather research and forecast model on the cray xt," in *Cray User Group 2010 Proceedings*, 2010.
- [31] D. Salmond and P. Towers, "Ifs migrates from ibm to cray 'cpu, comms and i/o'," in *Workshop on High-Performance Computing*. ECMWF, October 2014.
- [32] S. Saarinen, D. Salmond, and R. Forbes, "Preparation of ifs physics for future architectures," in *ECMWF HPC workshop*. ECMWF, October 2014.
- [33] W. Gropp and R. Thakur, "Issues in developing a thread-safe mpi implementation," in *Proceedings of the 13th European PVM/MPI User's Group Conference on Recent Advances in Parallel Virtual Machine and Message Passing Interface*, ser. EuroPVM/MPI'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 12–21. [Online]. Available: http://dx.doi.org/10.1007/11846802_11

Karthee Sivalingam Karthee Sivalingam is a Computational Scientist and holds a joint NCAS / MetOffice post with the aim of supporting the JWCRP high resolution and earth system modelling programmes by improving the portability, scalability, data handling, and overall performance of the Unified Model.

Bryan Lawrence Bryan Lawrence is the Director of Models and Data at the UK National Centre for Atmospheric Science, Professor of Weather and Climate Computing at the University of Reading, and the Director of the STFC Centre for Environmental Data Archival (CEDA).

Grenville Lister Grenville Lister is the head of Computational Modelling Support - National Centre for Atmospheric Science (NCAS)