

Efficient performance of the Met Office Unified Model v8.2 on Intel Xeon partially used nodes

I. Bermous, P. Steinle

Centre for Australian Weather and Climate Research, the Australian Bureau of Meteorology, Melbourne, Australia

Correspondence to: I. Bermous (i.bermous@bom.gov.au)

Abstract

The atmospheric Unified Model (UM) developed at the UK Met Office is used for weather and climate prediction by forecast teams at a number of international meteorological centres and research institutes on a wide variety of hardware and software environments. Over its 25 year history the UM sources have been optimised for better application performance on a number of High Performance Computing (HPC) systems including NEC SX vector architecture systems and recently the IBM Power6/Power7 platforms. Understanding the influence of the compiler flags, Message Passing Interface (MPI) libraries and run configurations is crucial to achieving the shortest elapsed times for a UM application on any particular HPC system. These aspects are very important for applications that must run within operational time frames. Driving the current study is the HPC industry trend since 1980 for processor arithmetic performance to increase at a faster rate than memory bandwidth. This gap has been growing especially fast for multicore processors in the past 10 years and it can have significant implication for the performance and performance scaling of memory bandwidth intensive applications, such as the UM. Analysis of partially used nodes on Intel Xeon clusters is provided in this paper for short and medium range weather forecasting systems using global and limited-area configurations. It is shown that on the Intel Xeon based clusters the fastest elapsed times and the most efficient system usage can be achieved using partially committed nodes.

1 **1 Introduction**

2 The Unified Model (UM) numerical modelling system (Brown et al., 2012) is used for short
3 and medium range weather forecasting, for both high resolution weather modelling and for
4 relatively coarser climate modelling. Such modelling software requires relatively powerful
5 High Performance Computing (HPC) systems to support operational forecast production.
6 Typically the computing systems have a peak performance comparable to the computer
7 systems included in the TOP500 list released every 6 months. Since September 2009 the UM
8 has been used in the Numerical Weather Prediction (NWP) component of the Australian
9 Community Climate and Earth System Simulator (ACCESS; Puri et al., 2010) at the
10 Australian Bureau of Meteorology (BoM).

11 Current operational systems at BoM are based on the UM version 7.5 (vn7.5) and the next
12 operational systems upgrade will be based on UM vn8.2. The latter version therefore was
13 used for the work described here.

14 UM versions include both science and performance upgrades, and extensive evaluation of
15 both types of changes in development mode is required prior to operational implementation.
16 In addition, changes to these systems have major consequences for many downstream
17 applications. For these reasons changing UM versions for operational systems is only done
18 every one to two years at the BoM.

19 Leading HPC systems have from tens of thousands to several million very powerful cores.
20 Since 1980 the trend in HPC development has been for the available processor performance to
21 increase at a greater rate than the available memory bandwidth (Graham et al., 2005, pp.106–
22 108). The authors concluded that a growing gap between processor and memory performance
23 could become a serious constraint in performance scaling for memory bound applications.
24 The gap between processor performance and memory bandwidth has been growing especially
25 quickly for multicore processors in the past 10 years (Wellein et al., 2012). This gap forces
26 the cores on a node to compete for the same memory causing resource contention, which can
27 become a major problem for memory intensive applications such as the UM.

28 Increasing the resolution of numerical models is one of the key approaches to improving
29 forecast accuracy. However, in an operational setting these models are constrained to run
30 within a fixed elapsed time on available computing resources. Increasing resolution requires
31 increased computation and therefore the performance efficiency (from herein just referred to
32 as efficiency) as measured by the run time on a given number of cores.

1 Finding the most efficient usage of the system for a particular application and the shortest
2 elapsed times varies depending on whether the application is run on all node cores (fully
3 committed case) or on a subset of the cores available on each node (partially committed
4 case). The placement of the threads and/or Message Passing Interface (MPI) processes across
5 partially committed nodes (sockets) also needs to be done carefully taking into consideration
6 all shared resources available to these cores.

7 Another practical aspect of the performance analysis discussed in the paper is to estimate
8 whether the coming upgrade for the BoM's operational models will be feasible given the
9 operational time windows and available HPC resources.

10 The performance analysis described here shows that on some modern HPC systems the
11 shortest run times can be achieved with the usage of partially committed nodes. The concept
12 of using partial nodes for UM applications allows reduced resource contention and improves
13 application performance (Bermous et al., 2013).

14

15 **2 Description of the models**

16 Regular upgrades to operational NWP systems are driven by the improvements made in both
17 the NWP software and the underlying science. The BoM is currently planning for the next
18 APS2 (Australian Parallel Suite 2) upgrade. The operational suite combines a number of short
19 and medium range weather forecasting systems based on the UM software. These systems
20 include the global NWP system (ACCESS-G), the regional NWP system (ACCESS-R,
21 12km), the tropical-cyclone forecast system (ACCESS-TC) and several city forecast-only
22 systems (ACCESS-C).

23 The current APS1 weather forecasting operational systems are based on UM vn7.5 for the
24 global (NMOC, 2012) and vn7.6 for the city systems (BoM, 2013). At this stage it is planned
25 that the operational weather forecasting software in APS2 will be upgraded to at least UM
26 vn8.2. This software includes improvements to physical parameterisations, computational
27 performance and performance scaling. Most of the scaling improvement is due to the
28 introduction of asynchronous I/O. With this upgrade the model resolutions for the Global and
29 City models will be increased. An increase in the model resolutions presents a challenge to fit
30 the model runs into the required operational time windows. As a result an initial analysis of
31 the performance measurements of the models is needed. In the current paper we will consider

1 two types of the weather forecasting models: a medium range N512L70 Global model and a
2 short range, limited area “City” model.

3 **2.1 Global N512L70 model**

4 The resolution of the currently operational Global N320 (40km) model with 70 vertical levels
5 in APS1 will be upgraded to N512 (25km), 70 levels in APS2. With a finite difference
6 discretisation of the UM mathematical model the latest Global model has a horizontal grid of
7 West-East x South-North of 1024 x 769. The existing "New Dynamics" dynamical core with
8 semi-implicit and semi-Lagrangian time integration (Davies et al., 2005) was used.

9 The operational systems run 4 times daily with two runs for 3 and two runs for 10 days. In
10 this paper performance scaling analysis for a 3 model day simulation with a 10min time step
11 was used. With the operational model settings this system produces 137GB of output data.
12 With relatively large amount of I/O, performance of the N512 global model and especially its
13 performance scalability is significantly affected by the I/O cost at high core counts.
14 Improvements in the scalability can be achieved with the usage of the asynchronous I/O
15 feature (Selwood, 2012) introduced into the model sources from UM release vn7.8. The I/O
16 servers functionality has been continually improving since then.

17 **2.2 UKV city model**

18 APS1 ACCESS-C operational systems nested in the ACCESS-R cover 5 Australian major
19 city domains: Adelaide, Brisbane (South-East Queensland), Perth, Sydney and
20 Victoria/Tasmania (Melbourne, Hobart). Each domain in the APS1 ACCESS-C has a
21 horizontal resolution of approximately 4km with 70 vertical levels. The corresponding short
22 range city models are set to run 36 hour forecasts 4 times daily. A significant horizontal
23 resolution increase is planned for the APS2 upgrade by reducing the horizontal resolution
24 from 4km to either 1.5km or 2.2km. As a result some initial performance analysis for the city
25 models is required to find the most efficient run configurations and arrangement within the
26 operational run-time schedule.

27 In this paper an example of a short-range limited-area forecasting is taken from a 1.5 km
28 high-resolution system for Sydney domain (Steinle et al., 2012). The experimental system
29 was created in 2012 and it is currently running 24 times per day. The range of the forecasts
30 provided by the system varies between 12 and 39 hours. The corresponding atmospheric

1 model is nested within the N512L70 global model and based on the UM vn8.2 sources using
2 its variable resolution version (the UKV) with the "New Dynamics" dynamical core.

3 The UKV modelling concept includes a high resolution of 1.5km in the inner domain, a
4 relatively coarse resolution of 4km near the boundary of the main domain and a transition
5 zone of using a variable grid size connecting the inner domain of 1.5 km with the "outer"
6 domain of 4 km.

7 The Sydney UKV model had a horizontal grid of E-W x N-S of 648x720 with 70 vertical
8 levels. The related forecast job was set to run a 25-hour simulation with a time step of 50 s
9 giving in total 1800 time steps for a run. The I/O in the job producing only 18 GB per run of
10 the output data is relatively small in comparison to the size of I/O in the global model job.
11 Therefore the usage of I/O servers does not have any major impact on the job performance,
12 even when a large number of cores are utilised.

13 **2.2.1 UKV decomposition constraints**

14 The MPI decomposition in the Unified Model is based on horizontal domain decomposition
15 where each subdomain (MPI process) includes a full set of vertical levels. Within the
16 computational core of the UM, OpenMP is generally used to parallelise loops over vertical
17 dimension.

18 Due to semi-Lagrangian dynamics implementation, the halo size for each sub-domain limits
19 the maximum number of sub-domains in each direction. With the halo size of 10 grid points
20 used in this study and the horizontal grid size of 648x720, the corresponding limits for the
21 MPI decomposition sizes were 42 in the West-East direction and 48 in the South-North
22 direction. Another constraint in the UM model implementation is that the decomposition size
23 in the West-East direction must be an even number.

24

25 **3 Description of HPC clusters and software used**

26 This section includes hardware specifications and details of the software environment for the
27 HPC clusters used.

1 **3.1 Hardware: specifications of HPC clusters**

2 Numerical results have been obtained on three HPC clusters with Intel® Xeon® processors.
3 The first Constellation Cluster (Solar) with Intel Nehalem processors was installed by Oracle
4 (Sun) in 2009. This system was upgraded by Oracle to a new cluster (Ngamai) with Sandy
5 Bridge processors in 2013.

6 In addition to the Ngamai system, the BoM also has access to a Fujitsu system (Raijin)
7 installed at the National Computational Infrastructure (NCI) at the Australian National
8 University (ANU) in Canberra, and it is also based on Intel Sandy Bridge processors. The
9 NCI system at 1.2 Pflops was the fastest HPC system in Australia in 2013. Technical
10 characteristics of these three systems are provided in Table 1.

11 A node Byte/Flop value in the table was calculated as a ratio of the node maximum memory
12 bandwidth and the node peak performance. For the newer Ngamai and Raijin systems this
13 ratio is less than half of that for Solar. All three clusters have Lustre file systems.

14 It should be noted that turbo boost was enabled in Basic Input/Output System (BIOS) on
15 Raijin only. As per Intel turbo boost 2.0 technology processor can run at above its base
16 operating frequency which is provided in Table 1 (“Node processor cores” line). Having idle
17 cores on processor, power that would have consumed by these idle cores can be redirected to
18 the utilised cores allowing them to run at higher frequencies. Based on the turbo boost
19 additional multipliers a Base Clock Rate (BCLK) can be calculated. For example, we have

$$20 \quad \text{BCLK} = (3.3\text{GHz} - 2.6\text{GHz}) / 7 = 100\text{MHz}$$

21 and utilised cores operate at $2.6\text{GHz} + 5 * \text{BCLK} = 3.1\text{GHz}$ if an application is run on 6 cores
22 from 8 cores available on each Raijin processor.

23 **3.2 Software: compiler, MPI library**

24 With the UM vn8.2 sources separate executables were required for each model type: global
25 and limited-area. The Intel compiler version 12.1.8.273 was used to produce UKV
26 executables. In order for the global N512L70 system to use the UM async I/O features the
27 Intel 14.0.1.106 compiler release was needed to avoid crashes due to interaction between the
28 older compiler and the new I/O code.

29 On the BoM HPC systems an open source implementation of MPI (OpenMPI) was the only
30 library available. The Intel MPI library was available on Raijin; however, testing showed a

1 10-20% degradation in performance in comparison with OpenMPI. For this reason and to
2 maintain compatibility between the NCI and BoM systems, OpenMPI was used for all
3 comparisons presented below.

4 The UKV executable was built with OpenMPI 1.6.5. The usage of a UM async I/O feature
5 requires at least SERIALIZED thread safety level. Therefore the OpenMPI 1.7.4 library was
6 needed to enable the async I/O feature of UM vn8.2.

7 **3.3 Intel compiler options**

8 The Fortran and C sources of the UM were compiled on Sandy Bridge systems Raijin and
9 Ngamai with the following Intel compiler options

```
10 -g -traceback -xavx -O3 -fp-model precise (1)
```

11 Option `-O3` specifies the highest optimisation level with the Intel compiler. Combination of
12 `“-g -traceback”` options was required in order to get information on a failed subroutine
13 call sequence in case of a run time crash problem. The usage of these two options had no
14 impact on the application performance. Bit reproducibility of the numerical results on a rerun
15 is a critical requirement for the BoM operational systems. For this purpose compilation flag
16 `“-fp-model precise”` (Corden et al., 2012) was used in spite of causing between 5% to
17 10% penalty in the UM model performance. An additional pair of compilation options `“-i8`
18 `-r8”` was used to compile Fortran sources of the UM. These options make integer, logical,
19 real and complex variables 8 bytes long. Option `-openmp` was specified to compile all model
20 sources and to link the corresponding object files to produce executables used for
21 MPI/OpenMP hybrid parallelism.

22 Due to a very limited capacity for non-operational jobs on the BoM operational system,
23 Ngamai, the testing and evaluation of the forecast systems prior their operational
24 implementation is predominantly performed on the relatively larger Raijin system. The BoM
25 share on Raijin is 18.9%.

26 Code generated by the compiler option `-xHost` targets the processor type on which it is
27 compiled. On the old Solar Nehalem chip based system, this was equivalent to compiling with
28 `-xsse4.2`. During the porting stage of our executables to the new Sandy Bridge Ngamai
29 and Raijin systems, in order to ensure compatibility of executables across these machines, the
30 `-xHost` compiler option used on Solar was replaced with the compilation flag `-xavx` for

1 advanced vector extensions supporting up to 256-bit vector data and available for Intel Xeon
2 Processor E5 family. It was confirmed, as expected, that adding the `-xHost` compiler option
3 to the set of compilation options (1) did not have any impact on either the numerical results or
4 the model performance.

5 Compatibility of binaries across systems was achieved by having the same Intel compiler
6 revisions and OpenMPI library versions on the both Ngamai and Raijin systems as well as
7 system libraries dynamically linked to the executables at run time. Note that the usage of Intel
8 compilers and MPI libraries on all systems is via the environment modules package
9 (<http://modules.sourceforge.net>). It was found empirically that using Intel compiler options
10 (1) as described above provided both compatibility of the executables between the systems
11 and reproducibility of the numerical results between Ngamai and Raijin.

12 General COMmunications (GCOM) library which provides the interface to MPI
13 communication libraries is supplied with the UM sources. GCOM version 4.2 used with UM
14 vn8.2 was compiled with a lower `-O2` optimisation level.

15

16 **4 Description of the performance results**

17 Jobs were submitted from login nodes to compute nodes, and job scheduling was done via
18 Sun Grid Engine (SGE) software on Ngamai (earlier on Solar) and Portable Batch Systems
19 Professional Edition (PBS Pro) job scheduler on Raijin.

20 All model runs were made on very busy multi-user systems using the standard normal
21 queue and a shared Lustre file system for I/O with a potential impact of the I/O contention on
22 the performance results. At the same time each model run used exclusive nodes without being
23 effected by suspend and resume functionality on the systems.

24 Fluctuations in the elapsed times were usually around 3-5%, but they were up to 50% in 3-
25 5% of the runs. This was particularly noticeable on the Raijin system which had consistent
26 utilisation above 95%. Support staff at NCI (D. Roberts) investigated this problem. It was
27 found that if cached memory is not being freed when a Non-Uniform Memory Access
28 (NUMA) node has run out of memory, any new memory used by a program is allocated on
29 the incorrect NUMA node and as a result slowing down access. The UM is particularly
30 sensitive to this issue. An environment setting of

31 `OMPI_MCA_hwloc_base_mem_alloc_policy=local_only` (2)

1 was recommended to include in the batch jobs running UM applications. Setting (2) forces all
2 MPI processes to allocate memory on the correct NUMA node, but if the NUMA node is
3 filled, the page file will be used. As a result the usage of (2) greatly improved stability of the
4 run times on Rajjin. Addressing these findings on Ngamai, it appeared that (2) was a default
5 setting on that system. The best run times were initially taken from 3 or 4 runs. If this initial
6 estimate appeared to be an outlier from the estimated performance scaling curve, further runs
7 were made. It is noteworthy that the fluctuations in elapsed times were much higher on all
8 systems when more than 2000 cores were used. The cause of these large fluctuations was not
9 investigated.

10 Choosing the best timing from a number of runs has been shown to provide reliable
11 estimates of timings under operational conditions – use of the highest priority queue and
12 dedicated file systems to avoid I/O contention and reserved and exclusive use of sufficient
13 nodes to run operational systems. These arrangements result in variations in elapsed times of
14 a few percent.

15 Starting from the old Solar system it was found that for I/O performance improvement
16 especially for applications with relatively heavy I/O in order of at least tens of gigabytes
17 Lustre striping had to be used. Based on the experimentation done on all three systems, Lustre
18 striping with a stripe count of 8 and a stripe size of 4M in a form of

```
19     lfs setstripe -s 4M -c 8 <run_directory>
```

20 was used to optimize I/O performance. Here <run_directory> is a directory where all output
21 files are produced during a model run. One of the criteria on whether striping parameters were
22 set to near optimal values was based on the consistency of the produced elapsed times for an
23 application running on a relatively busy system.

24 **4.1 UKV performance results**

25 Due to an earlier development of the Unified Model system at the end of 1980's - beginning
26 1990's on a massively parallel (MPP) system on which each CPU had its own memory,
27 initially the model had only a single level of parallelism using MPI. With the appearance of
28 symmetric multi-processor (SMP) systems in the 1990's and Open Multiprocessing (OpenMP)
29 software the hybrid MPI/OpenMP parallel programming concept which combines MPI across
30 the system nodes and multithreading with OpenMP within a single node was introduced. This
31 concept uses the shared address space within a single node. From the mid 2000's starting from

1 release 7.0 the hybrid parallel programming paradigm was introduced in the UM code and
2 since then the OpenMP implementation has been consistently improving in the model. Most
3 recent studies (Sivalingam, 2014) showed that even with UM vn8.6 the OpenMP code
4 coverage is limited. Furthermore the efficiency of pure MPI versus the hybrid parallelism
5 depends on the implementation, the nature of a given problem, the hardware components of
6 the cluster, the network and the available software (compilers, libraries) and the number of
7 cores used. As a result, there is no guarantee hybrid parallelism will improve performance for
8 every model configuration.

9 **4.1.1 Pure MPI vs MPI/OpenMP hybrid**

10 Comparison of the best elapsed times produced by running the UKV model with the usage of
11 pure MPI and MPI/OpenMP hybrid parallelism on Raijin and Ngamai is given in Figs 1 and 2
12 correspondingly. For simplicity the elapsed times are provided for 4 different decompositions
13 starting from the usage of 384 cores with a stride of 384. The run decompositions were 16x24,
14 24x32, 32x36 and 32x48 with pure MPI usage. In case of the hybrid parallelism, two
15 OpenMP threads were used and the related run configurations using the same number of cores
16 as in the pure MPI case were 2x6x32, 2x12x32, 2x16x36 and 2x16x48, where the first value
17 is the number of threads used. Fig.1 and Fig.2 include results for two cases: fully and partially
18 committed nodes. With partially committed nodes the application was running with the same
19 decomposition as in the fully committed node case, but only a part of each node was used: 8
20 cores from 12 on Ngamai and 12 cores from 16 cores on Raijin.

21 With the use of partially committed nodes, the placement/binding of cores to the
22 nodes/sockets should be done in a symmetrical way to give the same number of free cores on
23 each socket. This allows for a better usage of the shared L3 cache on each socket.

24 Based on the performance results plotted in Fig.2, the usage of pure MPI gives shorter
25 elapsed times than with the usage of the hybrid parallelism for all decompositions on Ngamai.
26 Fig.1 shows that a similar conclusion can be made for the elapsed times obtained on Raijin,
27 excluding the last point with the usage of 1536 cores. At the same time the shortest elapsed
28 times on Raijin are achieved using pure MPI on partially committed nodes (12 cores-per-
29 node) for the whole range of the used cores. Due to the limited proportion of the UM code
30 that can exploit OpenMP, the use of more than 2 threads (3 or 4) showed no improvement in
31 the performance efficiency.

1 Comparing the actual set of the obtained elapsed times between the two systems with the
2 usage of pure MPI on fully committed nodes of (2604; 1480; 1103; 942) s on Ngamai and
3 (2587; 1484; 1131; 1010) s on Raijin shows that the model performance on Raijin is slightly
4 worse than on Ngamai. At the same time comparing the corresponding elapsed times of
5 (2282; 1288; 947; 844) s on Ngamai and (2125; 1167; 857; 720) s on Raijin with the usage of
6 partially committed nodes for pure MPI, performance and especially performance scaling is
7 better on Raijin. For the same decomposition of 32x48 the elapsed time of 720 s on 2048
8 reserved cores on Raijin is 14.7% better than the corresponding elapsed time of 844 s
9 obtained on Ngamai on 2304 reserved cores. This improvement reduces with the number of
10 cores used, with only a 6.9% faster time for a decomposition of 16x24. Contributing factors to
11 this include the Raijin cores being slightly faster than Ngamai cores and turbo boost was not
12 enabled in BIOS on Ngamai. With the use of partially committed nodes, memory contention
13 between the processes/threads running on the same node is reduced, improving performance
14 for memory intensive applications. At the same time enabling turbo boost can increase
15 processor performance substantially, reaching peak speeds of up to 3.1 GHz on Raijin using
16 12 cores-per-node.

17 On Raijin the usage of 8 out of 16 cores with hybrid parallelism and 2 OpenMP threads
18 showed between 16.6% (for high core counts) to 31.0% (for low core counts) slower run
19 times in the 768-3072 reserved core range in comparison with the corresponding results
20 obtained using pure MPI. On Ngamai the usage of half committed (6 cores from 12 cores)
21 nodes with the hybrid parallelism gives two possible configurations for running an
22 application. Firstly there is the symmetrical case with 1 MPI process and 3 OpenMP threads
23 running on each socket. Another option is the non-symmetrical case with 2 MPI processes
24 running on one socket and 1 MPI process running on another socket with 2 threads per each
25 MPI process. Taking into account the limited OpenMP coverage in the UM vn8.2 sources,
26 tests using hybrid parallelism with a symmetrical case of 3 threads or an asymmetrical case
27 with 2 threads on half committed nodes were not performed on Ngamai.

28 **4.1.2 Fully committed nodes vs partially committed nodes**

29 Elapsed times for the UKV model with pure MPI usage on partially committed nodes on all 3
30 systems are provided in Fig.3- Fig.8. Each pair of figures (Fig.3-4 for Raijin; Fig.5-6 for Solar
31 and Fig.7-8 for Ngamai) shows speedup as a function of the number of cores actually used as
32 well as a function of the reserved cores (i.e. total number of cores allocated to the run, both

1 used and unused). The performance relative to the number of reserved cores is the most
2 important metric, however performance relative to the “Number of used cores“ provides
3 additional information on the value of reducing the number of active cores per node. This
4 extra information is particularly relevant to circumstances where the elapsed time is more
5 important than using nodes as efficiently as possible. Examples include climate runs and cases
6 where other restrictions mean that the number of nodes available is not a significant constraint
7 on an application's run time. The related performance information cannot be easily seen on
8 the graph using the "Number of reserved cores" metric.

9 For example, a 12 cores-per-node case on Raijin and a 6 cores-per-node case on Solar each
10 reserved full nodes (16 and 8 cores respectively), but left a quarter of cores unused. This
11 indicates a requirement of specifying by 1/3 of more cores using `-npersocket` or `-`
12 `npersocket` option of the `mpirun` command in comparison with the fully committed case
13 using the same run configuration. In the example of running the model with pure MPI on
14 Raijin and using 12 cores per node the following options

```
15     mpirun --npersocket 6 -mca orte_num_sockets 2 \  
16           -mca orte_num_cores 8 . . .
```

17 were used in the `mpirun` command with OpenMPI 1.6.5. The last two options specify the
18 number of sockets on a node and the number of cores on each socket. These options were
19 required to avoid bugs found in the OpenMPI 1.6.5 software.

20 Fig.3 shows the value of partially committed nodes on Raijin where using 12 cores-per-
21 node significantly improves the model scaling. This improvement generally increases as the
22 number of active cores increases. The improvement reaches 28.5% the performance of using
23 1728 fully committed nodes. The usage of 8 cores-per-node on Raijin gives an additional
24 performance improvement in comparison with the 12 cores-per-node case varying from
25 16.8% at 96 cores to 9.6% at 1728 cores. Examining the same performance results on the
26 reserved cores basis as in Fig.4 shows that it is more efficient to use 12 cores-per-node than
27 fully committed nodes just with over 768 cores. On 768 reserved cores the 12 cores-per-node
28 case has value of 1495 s for a 24x24 decomposition and the fully committed node case has
29 value of 1484 s for a 24x32 decomposition.

1 Fig. 5 shows that the usage of partially committed nodes on Solar improves the runtimes
2 with 6 cores-per-node by 6.9-16.5% and a further reduction of 8.2-11.0% is achieved with the
3 usage of 4 cores-per-node.

4 The speedup curves as a function of used cores on Ngamai shown in Fig.7 indicate that the
5 model runs 10.4-14.6% faster with 8 cores-per-node. Unlike the other two systems (Raijin
6 and Solar), the use of half utilised nodes with 6 cores-per-node on Ngamai gives only a very
7 modest reduction of no more than 5.3%. These latter results indicate that a reduction in
8 memory contention with the 6 cores-per-node case has almost no impact over using 8 cores-
9 per-node.

10 Speedup curves as functions of the reserved cores for Solar (Fig. 6) and Ngamai (Fig. 8)
11 show that unlike Raijin, the efficiency gains on partial nodes were not achieved on up to 1152
12 reserved cores on Solar and 1728 on Ngamai. A relatively poor UKV performance on partial
13 nodes especially on Ngamai system in comparison with Raijin was due to unavailability of
14 turbo boost on that system. Turbo boost would allow active cores to run at up to 16% higher
15 clock speeds with the usage of 8 cores-per-node on Ngamai.

16 The shortest run times using fully committed nodes on Ngamai and Raijin with the usage of
17 up to 1728 cores were achieved on a decomposition of 36x48 with pure MPI. The largest
18 allowable decomposition under the constraints provided in Sect. 2.2.1 is 42x48 on 2016
19 cores. Increasing the number of cores from 1728 to 2016 provides further improvements in
20 the elapsed times of 1.7% on Ngamai and 2.2% on Raijin. This indicates that the
21 corresponding performance scaling have not begun to level off at the largest allowed
22 decomposition of 42x48. At the same time, with the use of partially committed nodes on
23 Raijin, an elapsed time of 950 s obtained on fully committed nodes for a 36x48
24 decomposition can be improved by: 28.5% (679 s) on 2304 reserved cores with 12 cores-per-
25 node or 35.4% (614 s) on 3456 reserved cores with 8 cores-per-node.

26 The above mentioned constraint of 2016 cores on fully committed nodes is applied for pure
27 MPI only. With the usage of hybrid parallelism and 2 OpenMP threads on fully committed
28 nodes performance of the model is still improving when the number of cores is increasing
29 from 1536 for decomposition of 16x48 to 3072 for a decomposition of 32x48 but the
30 corresponding run times are still greater than the run times obtained with pure MPI on partial
31 nodes using the same number of reserved cores. For example, the use of multithreading with 2
32 OpenMP threads and decomposition of 30x48 on 3072 cores gives an elapsed time of 669sec.

1 This run time is improved by 9.1% (608 s) with the use of pure MPI for a decomposition of
2 40x48 run on 10 cores-per-node with the same 3072 cores.

3 **4.2 N512L70 performance results**

4 As mentioned earlier, the UM async I/O feature (Selwood, 2012) was used to obtain good
5 performance and especially performance scaling when running the model on more than 1000
6 cores. Using the UM multithreaded I/O servers feature, all model runs were made with 2
7 OpenMP threads. The usage of 3 or even 4 threads showed no improvement in model
8 performance.

9 Elapsed times from runs without I/O were used as a target for the async I/O case. The run
10 times with full I/O in the fully committed node case were within 5%-10% of those without
11 I/O. Note that on a very busy such as Raijin system some improvement in the run times for a
12 few cases were achieved using different from (2) Lustre striping parameters, namely

```
13 lfs setstripe -s 8M -c <number_of_IO_servers> <run_directory>
```

14 As in the previous case (2) the values for the Lustre stripe count and the stripe size were
15 found experimentally. Performance results for the model were obtained with up to 3500 cores
16 on Ngamai. Due to large variations in the run times with over 4500 cores on Raijin, the
17 related model performance results are provided for only up to 3840 cores.

18 The best elapsed times obtained on Raijin are provided in Table 2, where the I/O servers
19 configuration of $(m \times n)$ included in the third column of the table have the following
20 meaning: m is a number of the I/O server groups, n is a number of I/O tasks per server. For up
21 to 2688 cores the best performance was achieved on fully committed nodes. For 3072 or more
22 cores the best performance results were achieved using partially committed nodes.

23 Performance scaling of the model as a function of a number of the reserved cores for fully
24 committed node case, 12 cores-per-node and 8 cores-per-node cases is shown in Fig.9. The
25 curves clearly show that the most efficient system usage with 3072 cores or higher is achieved
26 running the application on partially committed nodes with 12 cores on each node from 16
27 available. The curves corresponding to 12 cores-per-node and 8 cores-per-node cases show a
28 reasonably good scaling of the model with the usage of up to 4000 cores. Note that using
29 partially committed nodes the model performance is slightly worse when core usage is in the
30 range 384 to 2688.

1 The best elapsed times obtained on Ngamai are provided in Table 3. On this system in
2 contrast with Raijin, the most efficient usage is achieved using fully committed nodes.
3 Performance scaling of the model as a function of a number of the reserved cores for fully
4 committed node case and with the usage of 8 cores-per-node case is provided on Fig.10. For
5 the fully committed node case a relatively good performance scaling is achieved with the
6 usage of up to 1920 cores, after that performance scaling degrades slowly with the usage of
7 2304 and 2688 cores and levels out by 3072 cores. Based on the elapsed times produced with
8 the usage of up to 2688 cores, the most efficient usage of the system is with fully committed
9 nodes. At the same time the usage of 8 cores-per-node for up to 3456 reserved cores has
10 relatively good performance scaling and from the efficiency point of view runs with 3072
11 reserved cores and higher should use partially committed nodes. Our expectations are that this
12 efficiency in partial used nodes could be improved if turbo boost was enabled.

13 Performance results for a 6 cores-per-node case are not presented in Fig.10. As discussed at
14 the end of Sect. 4.1.2 there are two possible run configurations for this case. With
15 symmetrical usage of 3 threads per MPI process run on each socket, the model performance
16 was even worse in comparison with the fully committed node case. At the same time the non-
17 symmetrical usage with 3 MPI processes and 2 threads gave similar performance results as in
18 the 8 cores-per-node case.

19

20 **5 Conclusions**

21 With a trend in the HPC industry of decreasing the Byte/Flop ratio especially on multicore
22 processors and increasing the number of cores per CPU, the most efficient system usage by
23 memory intensive applications can be achieved with the usage of partially committed nodes.
24 In other words, the following factors such as increasing memory bandwidth per active core,
25 reduction in the communication time using less MPI processes and active cores running at
26 higher clock speeds with turbo boost can more than compensate for the reduced number of
27 cores in action. This approach can improve an application performance and most importantly
28 the application performance scaling. A conclusion on whether a specific application should be
29 running on fully or partially committed nodes depends on the application itself as well as on
30 the base operating frequency of the processor and memory bandwidth available per core.
31 Other factors such as availability of turbo boost, hyper-threading and type of node
32 interconnect on the system can also influence the best choice. This study showed that both the

1 regional and global models can run faster if partially committed nodes are used on Raijin. At
2 the same time taking into an account the similarities between Raijin and Ngamai systems,
3 there is a reasonable expectation that a similar affect would have been achieved on Ngamai if
4 turbo boost would be available on this system.

5 The usage of partially committed nodes can further reduce elapsed times for an application
6 when the corresponding performance scaling curve has flattened.

7 Another example when the use of partially committed nodes can reduce run times is when
8 the performance scaling has not flattened but the number of used cores cannot be increased
9 due to other constraints in the application. This case was illustrated by the UKV model
10 example in Sect. 4.1.2. This approach can be used for climate models based on the UM
11 sources and run at a relatively low horizontal resolution. As per the results of this Sect. 4.1.2
12 the usage of partial nodes can reduce elapsed times significantly. This has a very important
13 practical value for climate research experiments that require many months to complete.

14 The approach of using partially committed nodes for memory bandwidth-bound
15 applications can have a significant practical value for efficient HPC system usage. In addition,
16 this can also ensure the lowest elapsed times for production runs of time critical systems. This
17 approach is a very quick method for providing major performance improvements. In contrast,
18 achieving similar improvements through code related optimisation can be very time
19 consuming and may not even be as productive.

20

21 **Code availability**

22 The Met Office Unified Model is available for use under licence. The Australian Bureau of
23 Meteorology and CSIRO are licensed to use the UM in collaboration with the Met Office to
24 undertake basic atmospheric process research, produce forecasts, develop the UM code and
25 build and evaluate Earth System models. For further information on how to apply for a
26 licence see <http://www.metoffice.gov.uk/research/collaboration/um-collaboration>.

27

28 **The Supplement related to this article is available online at**
29 **[doi:10.5194/gmdd-7-7395-2014-supplement](https://doi.org/10.5194/gmdd-7-7395-2014-supplement).**

30

1 **Acknowledgements**

2 The authors thank Yi Xiao (BoM) for providing the UM job settings used in the paper; Robert
3 Bell (CSIRO), Paul Selwood (UK Met Office), Gary Dietachmayer (BoM) and Martyn
4 Corden (Intel) for their useful comments; Jörg Henrichs (Oracle, Australia), Dale Roberts
5 (NCI, ANU) and Ben Menadue (NCI, ANU) for continuous application support provided on
6 all systems. A useful discussion with Dale Roberts on the Lustre striping parameter settings
7 helped improve elapsed times for a few configurations where the original settings gave small
8 anomalies (a few percent) in scaling results. We thank two anonymous reviewers for the
9 recommendations to improve the original version of the manuscript.

10

11 **References**

- 12 Bermous, I., Henrichs, J., and Naughton, M., 2013: Application performance improvement by
13 use of partial nodes to reduce memory contention. *CAWCR Research Letters*, 9, 19-22.
- 14 Brown, A.R., Milton, S., Cullen, M., Golding, B., Mitchell, J., and Shelly, A., 2012: Unified
15 modelling and prediction of weather and climate: a 25 year journey. *Bull. Am. Meteorol.*
16 *Soc.*, 93, 1865-1877.
- 17 Corden, M., Kreitzer, D.: Consistency of Floating-Point Results using the Intel® Compiler or
18 Why doesn't my application always give the same answer? available at:
19 [https://software.intel.com/en-us/articles/consistency-of-floating-point-results-using-the-](https://software.intel.com/en-us/articles/consistency-of-floating-point-results-using-the-intel-compiler)
20 [intel-compiler](https://software.intel.com/en-us/articles/consistency-of-floating-point-results-using-the-intel-compiler) (last access: 3 September 2014), 2012.
- 21 Davies, T., Cullen, M. J. P., Malcolm, A. J., Mawson, M.H., Staniforth, A., White, A. A. and
22 Wood, N. 2005: A new dynamical core for the Met Office's global and regional modelling
23 of the atmosphere. *Quarterly Journal of the Royal Meteorological Society*, 131, 1759–
24 1782.
- 25 Graham, S.L., Snir, M., and Patterson, C.A., 2005: *Getting Up To Speed: The Future Of*
26 *Supercomputing*. National Academies Press, 289pp.
- 27 NMOC Operations Bulletin Number 93, available at:
28 <http://www.bom.gov.au/australia/charts/bulletins/apob93.pdf>
29 (last access: 3 September 2014), 2012.

1 Puri, K., Xiao, Y., Sun, X., Lee, J., Engel, C., Steinle, P., Le, T., Bermous, I., Logan, L.,
2 Bowen, R., Sun, Z., Naughton, M., Roff, G., Dietachmayer, G., Sulaiman, A., Dix, M.,
3 Rikus, L., Zhu, H., Barras, V., Sims, H., Tingwell, C., Harris, B., Glowacki, T.,
4 Chattopadhyay, M., Deschamps, L., Le Marshall, J., 2010: Preliminary results from
5 Numerical Weather Prediction implementation of ACCESS. CAWCR Research Letters, 5,
6 15-22.

7 Selwood, P.: The Met Office Unified Model I/O Server, ENES Workshop: Scalable IO in
8 climate models, available at: [https://verc.enes.org/computing/hpc-collaborations/parallel-i-](https://verc.enes.org/computing/hpc-collaborations/parallel-io/workshop-scalable-io-in-climate-models/presentations/Unified_Model_IO_Server_Paul_Selwood.ppt/view)
9 [o/workshop-scalable-io-in-climate-](https://verc.enes.org/computing/hpc-collaborations/parallel-io/workshop-scalable-io-in-climate-models/presentations/Unified_Model_IO_Server_Paul_Selwood.ppt/view)
10 [models/presentations/Unified_Model_IO_Server_Paul_Selwood.ppt/view](https://verc.enes.org/computing/hpc-collaborations/parallel-io/workshop-scalable-io-in-climate-models/presentations/Unified_Model_IO_Server_Paul_Selwood.ppt/view) (last access: 3
11 September 2014), 2012.

12 Sivalingam, K.: Porting and optimisation of MetUM on ARCHER, The 16th ECMWF
13 Workshop on High Performance Computing in Meteorology, available at:
14 <http://www.ecmwf.int/sites/default/files/HPC-WS-Sivalingam.pdf> (last access: 28 January
15 2015), 2014

16 Staniforth, A., and Wood, N. 2008: Aspects of the dynamical core of a nonhydrostatic, deep-
17 atmosphere, unified weather and climate-prediction model. Journal of Computational
18 Physics, Volume 227, Issue 7, 3445-3464.

19 Steinle, P., Xiao, Y., Rennie, S., Wang, X.: SREP Overview: Meso-scale Modelling and Data
20 Assimilation, available at:
21 http://cawcr.gov.au/research/esm/Modelling_WS/10_MEW/Steinle_P.pdf (last access: 3
22 September 2014), 2012

23 The Australian Bureau of Meteorology (BoM): NMOC Operations Bulletin Number 99,
24 available at: <http://www.bom.gov.au/australia/charts/bulletins/apob99.pdf>
25 (last access: 3 September 2014), 2013.

26 Wellein, G., Hager, G., Kreutzer, M.: Programming Techniques for Supercomputers: Modern
27 Processors, available at:
28 [https://wiki.engr.illinois.edu/download/attachments/217842128/performance-](https://wiki.engr.illinois.edu/download/attachments/217842128/performance-bandwidth.pdf?version=1&modificationDate=1359049396000)
29 [bandwidth.pdf?version=1&modificationDate=1359049396000](https://wiki.engr.illinois.edu/download/attachments/217842128/performance-bandwidth.pdf?version=1&modificationDate=1359049396000)
30 (last access: 3 September 2014), 2012.

1 **Table 1.** Intel Xeon Compute System Comparison.

	Solar, BoM	Ngamai, BoM	Raijin, NCI (ANU)
Processor type	Intel Xeon X5570	Intel Xeon E5-2640	Intel Xeon E5-2670
Number of compute nodes	576	576	3592
Total number of cores	4608	6912	57472
InfiniBand interconnect	QDR	QDR	FDR
Memory size per node	24GB	64GB	32GB on 2395 nodes (67%) 64GB on 1125 nodes (31%) 128GB on 72 nodes (2%)
Node processor cores	2x(2.93GHz, 4-core)	2x(2.5GHz, 6-core)	2x(2.6GHz, 8-core)
Max turbo frequency	3.333GHz	3GHz	3.3GHz
Node cache size	2x8MB	2x15MB	2x20MB
Memory type	DDR3-1333MHz	DDR3-1333MHz	DDR3-1600MHz
Number of memory channels	3	4	4
Node peak performance (base)	85 GFlops	240 GFlops	332.8 GFlops
Node max memory bandwidth	64 GB/s	85.3 GB/s	102.4 GB/s
Byte/Flop	0.753	0.355	0.308

Turbo boost additional multipliers	2/2/3/3	3/3/4/4/5/5	4/4/5/5/6/6/7/7
Turbo boost	OFF	OFF	ON
Usage of hyper threading	No	No	No

1

1 **Table 2.** The best elapsed times for N512L70 on Raijin using 2 threads.

Total reserved cores	Decomposition	I/O servers configuration	Cores used per node	Elapsed time (sec)
384	8x23	2x4	16	2881
768	12x31	2x6	16	1575
1152	14x40	2x8	16	1131
1536	16x47	2x8	16	927
1920	16x59	2x8	16	802
2304	18x63	2x8	16	701
2688	20x66	2x12	16	640
3072	18x63	2x8	12	603
3584	20x66	2x12	12	543

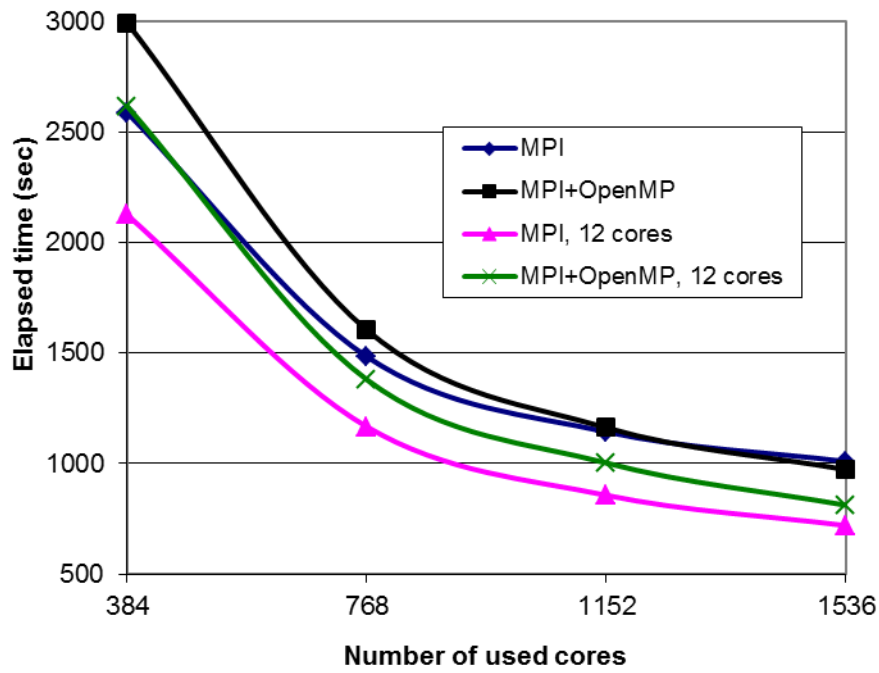
2

1 **Table 3.** The best elapsed times for N512L70 on Ngamai using 2 threads on fully committed
2 nodes.

Total reserved cores	Decomposition	I/O servers configuration	Elapsed time (sec)
384	8x23	2x4	3068
768	12x31	2x6	1639
1152	14x40	2x8	1217
1536	16x47	2x8	1026
1920	18x52	2x12	877
2304	18x63	2x9	805
2688	20x66	2x12	759
3072	22x69	2x9	756

3

UKV on Raijin



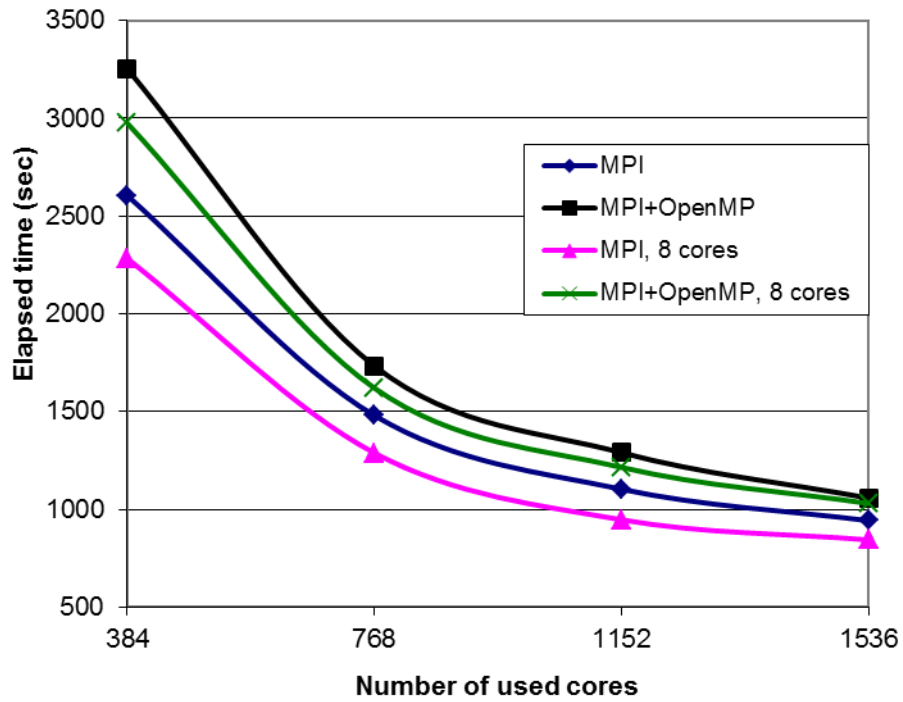
1

2

3 **Figure 1.** Elapsed times for the UKV model runs on Raijin versus the number cores actually

4 used in each run.

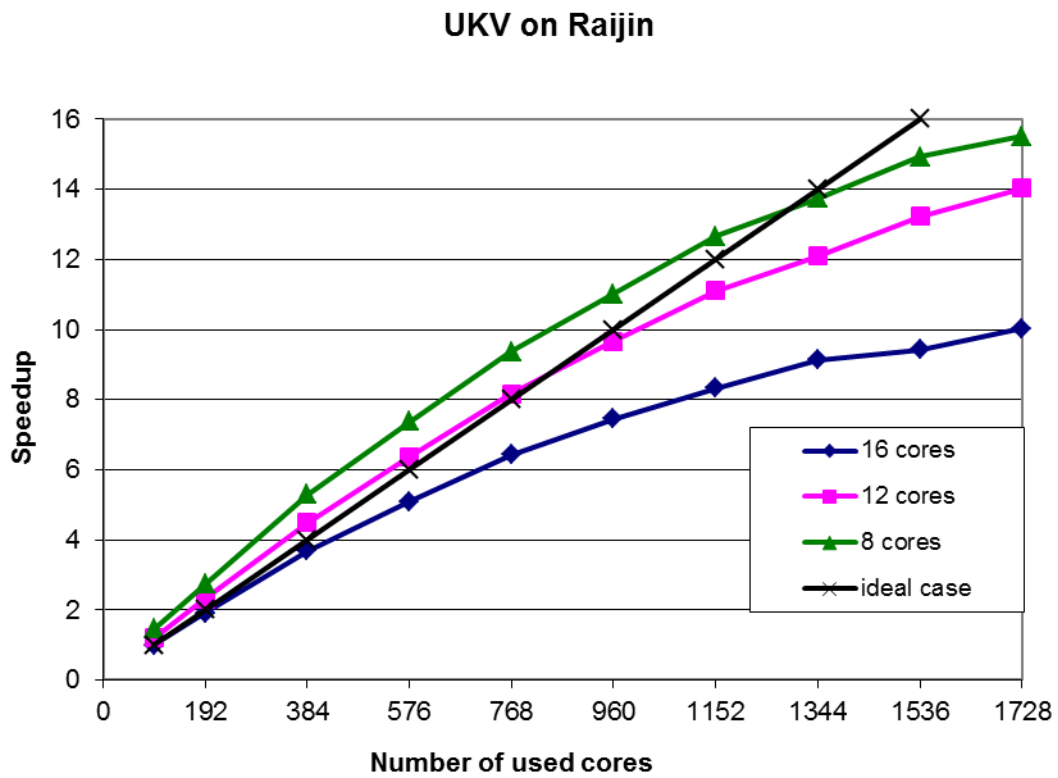
UKV on Ngamai



1

2

3 **Figure 2.** Elapsed times for the UKV model runs on Ngamai versus the number cores actually
4 used in each run.

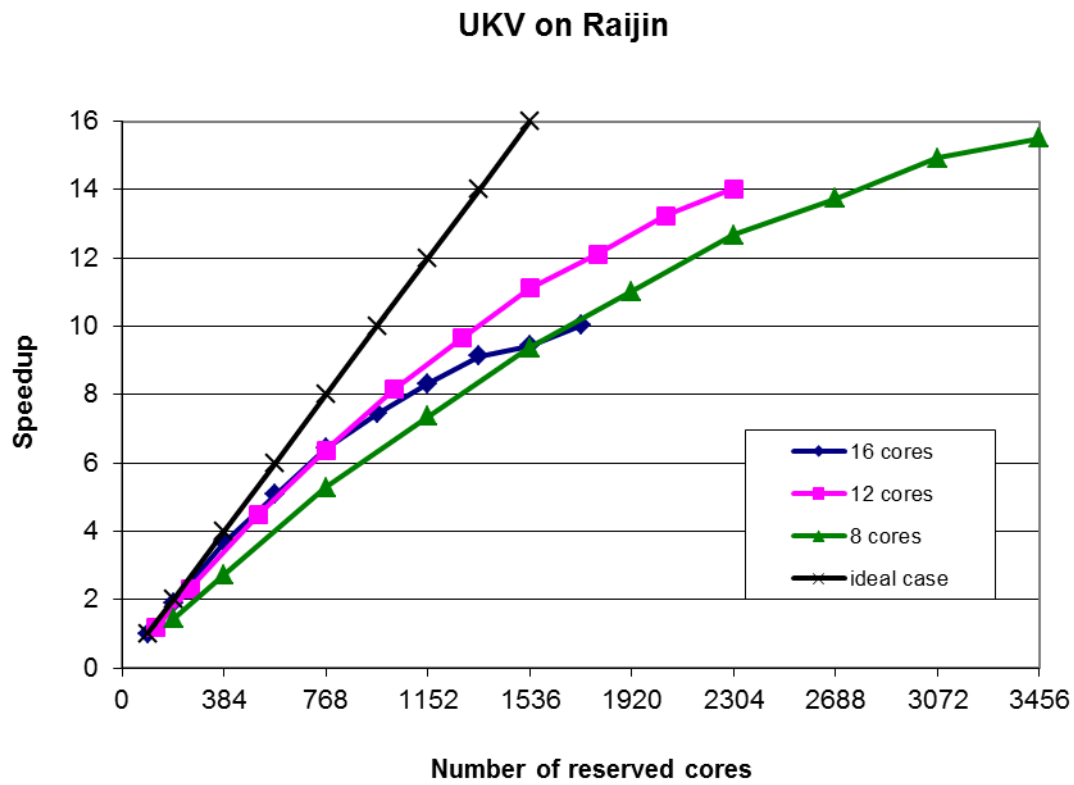


1

2

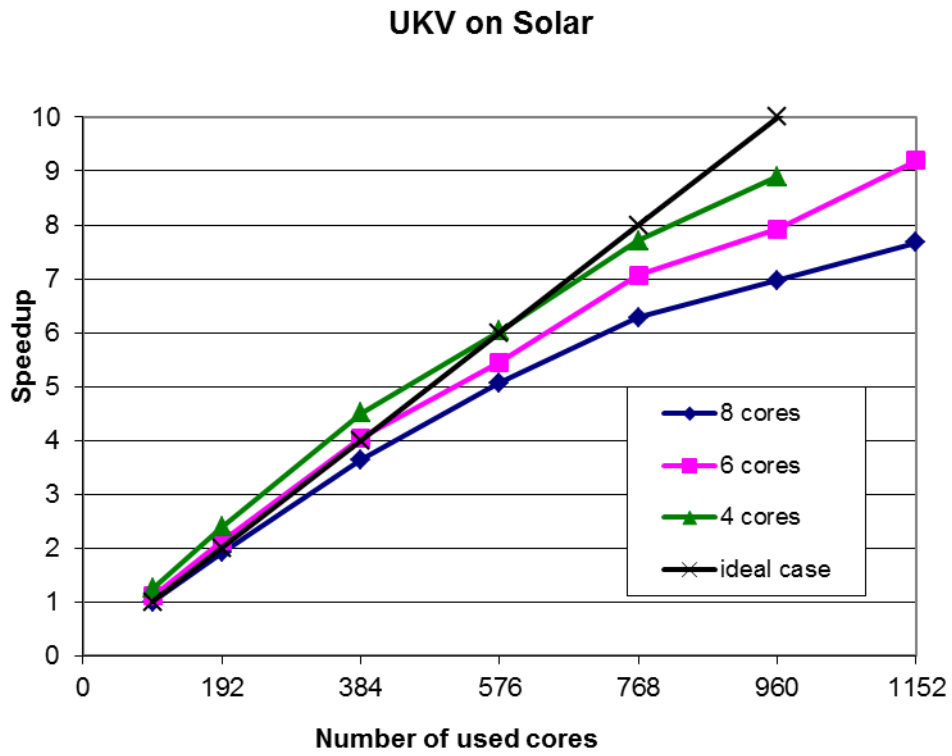
3 **Figure 3.** Speedup as a function of number of used cores on Raijin. Speedup was calculated

4 in relation to the elapsed time of 9523 s obtained for a 96 core run on fully committed nodes.



1
2
3
4
5

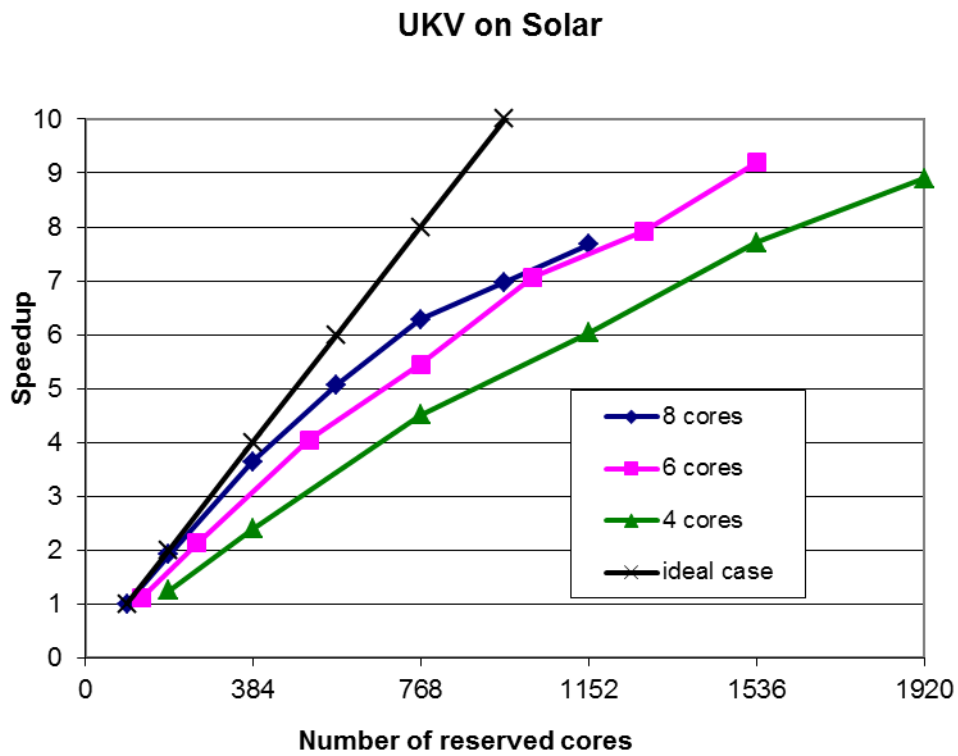
Figure 4. Speedup as a function of number of reserved cores on Raijin. Speedup was calculated in relation to the elapsed time of 9523 s obtained for a 96 core run on fully committed nodes.



1

2

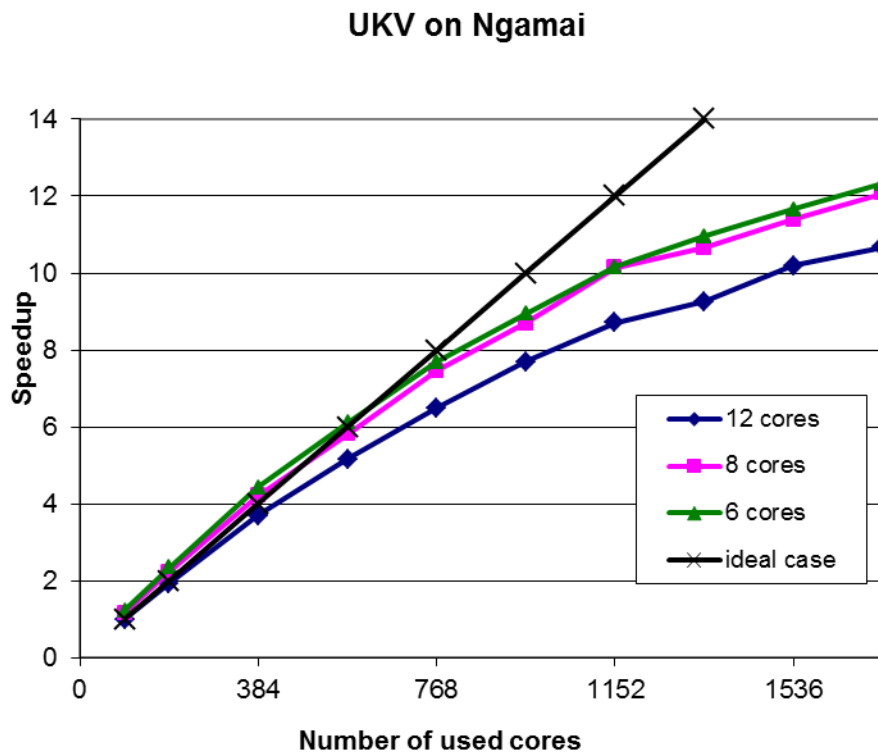
3 **Figure 5.** Speedup as a function of number of used cores on Solar. Speedup was calculated in
 4 relation to the elapsed time of 11488 s obtained for a 96 core run on fully committed nodes.



1

2

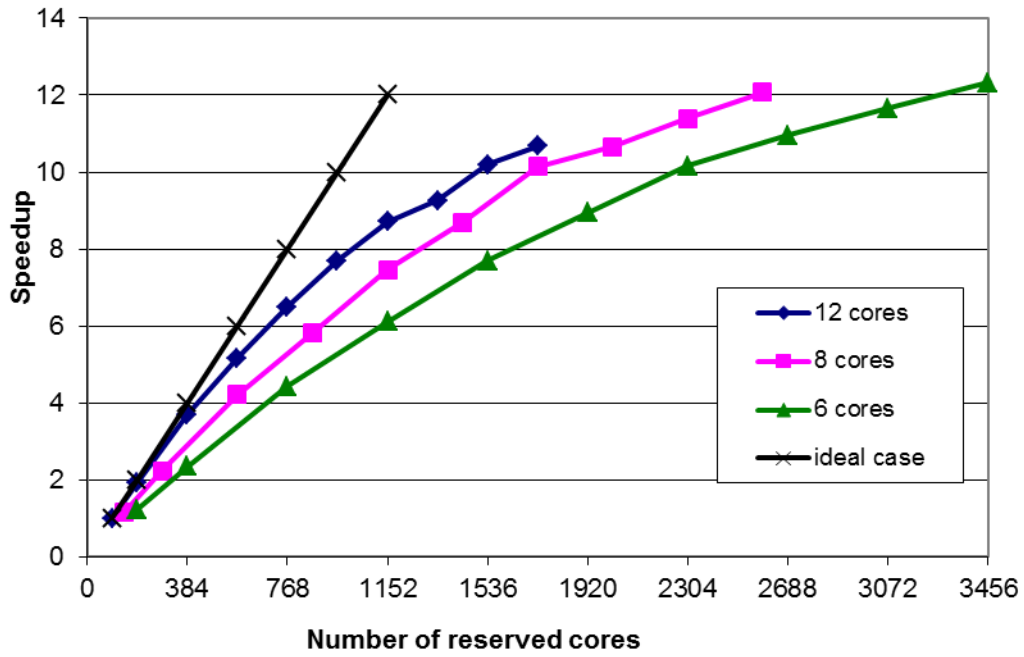
3 **Figure 6.** Speedup as a function of number of reserved cores on Solar. Speedup was
 4 calculated in relation to the elapsed time of 11488 s obtained for a 96 core run on fully
 5 committed nodes.



1
2
3
4

Figure 7. Speedup as a function of number of used cores on Ngamai. Speedup was calculated in relation to the elapsed time of 9608 s obtained for a 96 core run on fully committed nodes.

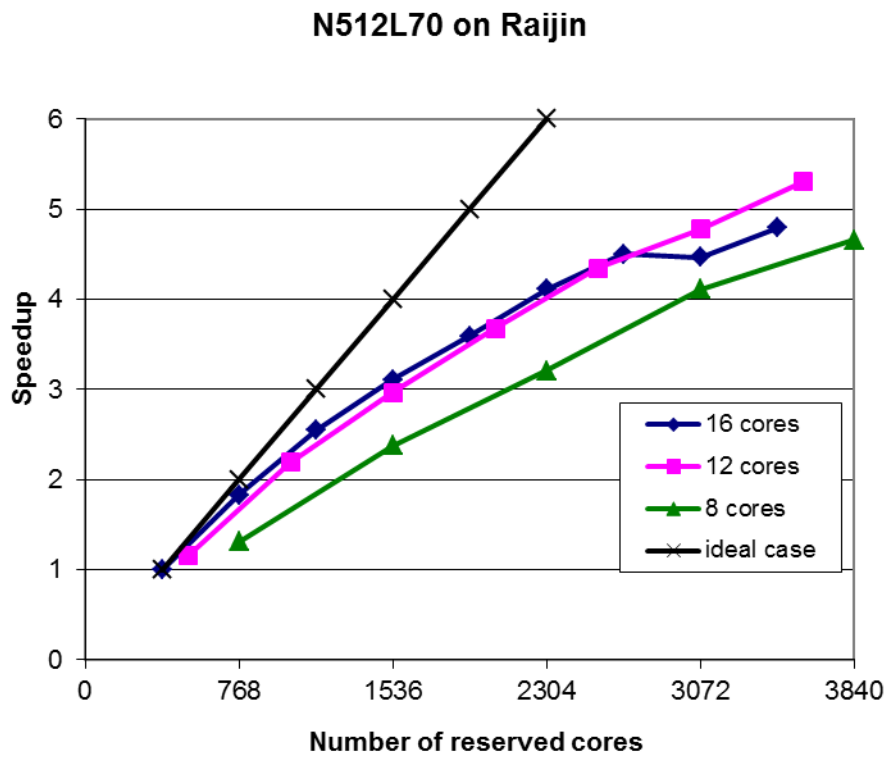
UKV on Ngamai



1

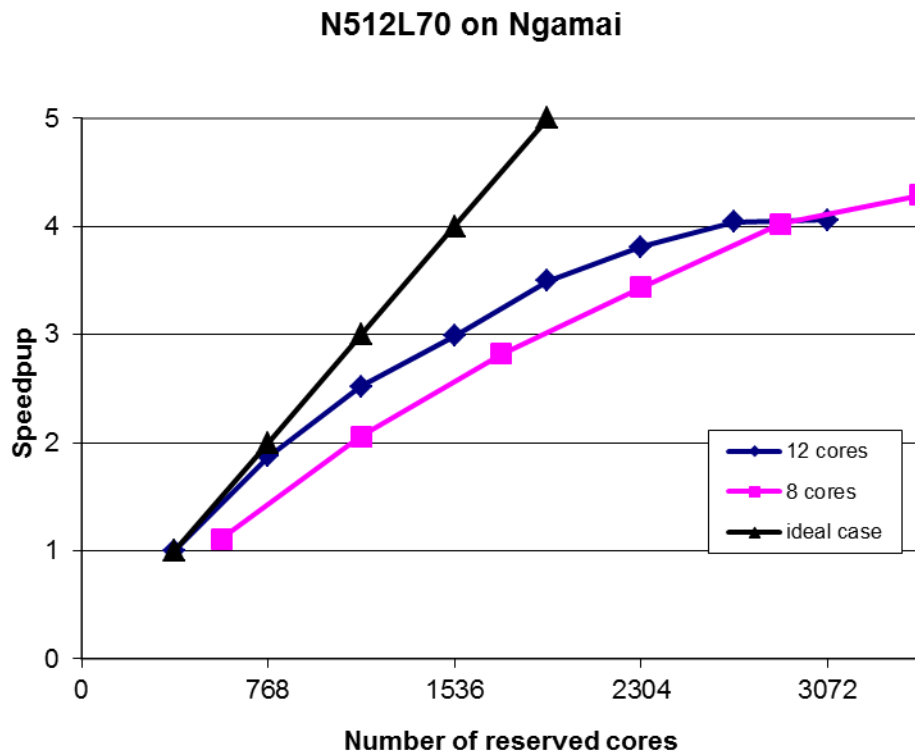
2

3 **Figure 8.** Speedup as a function of number of reserved cores on Ngamai. Speedup was
4 calculated in relation to the elapsed time of 9608 s obtained for a 96 core run on fully
5 committed nodes.



1
2
3
4
5

Figure 9. Speedup as a function of number of reserved cores on Raijin. Speedup was calculated in relation to the elapsed time of 2881 s obtained for a 384 core run on fully committed nodes.



1

2

3 **Figure 10.** Speedup as a function of number of reserved cores on Ngamai. Speedup was
 4 calculated in relation to the elapsed time of 3068 s obtained for a 384 core run on fully
 5 committed nodes.