

Interactive comment on “Efficient performance of the Met Office Unified Model v8.2 on Intel Xeon partially used nodes” by I. Bermous

Anonymous Referee #1

Received and published: 24 November 2014

The paper describes the performance of Met Unified model on 3 Intel machines and discusses the best parameters for efficient performance. The UKV city model shows improvement in performance when the nodes are only partially used. The N512L70 model scales up to 4000 cores. When more than 3072 cores are used, partially committed nodes shows relatively better performance compared to fully committed nodes. The author explains that the UM is memory intensive and using partially committed nodes improves the memory bandwidth of UM.

Page 7396 line 9 : The paper mentions about “Shortest elapsed times” and then again about finding the “most efficient run configurations” at line 2 , page 7397. Which is more important?

Page 7396 line 15: “memory bandwidth intensive applications” – Suggestion is that
C2396

Memory intensive or memory bandwidth limited applications will be better suited.

Page 7396 line 23: “relatively powerful HPC systems” - Comparison to other HPC systems will be helpful.

Page 7396 line 24: numerical weather prediction – The first letters should be capitalized.

Page 7397 line 9: “asynchronous computation and I/O “should be “computation and asynchronous I/O” as only the I/O is asynchronous

Page 7397 line 12: “these are not expected to significantly impact on the major results of this paper.” is an assumption which can be avoided.

Page 7399 line 11: It is not clear if it is 3 day simulation or 3 model day simulation.

Page 7400 line 1: First use of 3D-VAR and has not been described before.

Page 7400 line 11: It is not clear if 18 GB of data produced per run or per model day?

Page 7401 line 5: “reduced by over 2 times” should be reworded.

Page 7401 line 10: The more advanced optimizations are not stated.

Page 7403 line 25: Figure 1 and 2 shows only the elapsed times and not the efficiency.

Page 7404 line 1: The author describes using partially committed nodes by using 12 cores from 16 cores of Raijin (with OPENMP). Why is using only 8 cores from 16 cores not explored? And similarly for Ngamai why is 6 cores from 12 not explored.

Page 7405 line 19: The author should describe the options `orte_num_sockets` and `orte_num_cores` used.

Page 7406 line 1: The two efficiency measures are not described in the paper and are not shown in Fig 3 or Fig 4.

Page 7406 line 23: The author does not describe if and why the number of sub domains

in the East –West direction is limited to only 36.

Page 7408 line 15: The turbo boost or its usage is not described by the author.

Page 7409 line 26: “Couple of hundred cores” is vague and exact number should be mentioned.

Page 7415 table 3: Table 3 does not have information about the number of cores used per node.

The paper in its current form should not be published. Even though the findings about using partially committed nodes for UM are good, the paper’s presentation is not clear. The language is very difficult to understand at places. There are some technical details which are not described and some assumptions which are not clearly stated.

Interactive comment on Geosci. Model Dev. Discuss., 7, 7395, 2014.