

Runtime performance of IFS c36r4 on Beskow

Aims:

- Study the performance and scaling of IFS c36r4 EC-Earth configuration on the PDC Beskow CRAY XC40 system (<https://www.pdc.kth.se/resources/computers/beskow>).
- Compare the model compiled using the CRAY and Intel compilers.
- Use the Allinea Performance Reports tool (<http://www.allinea.com/products/allinea-performance-reports>) to extract additional information on the performance characteristics of the model.

Compile options:

1. CRAY:

- General F90 flags for compiling (Note a small number of subroutines compiled with -O0)
-sreal63 -em -hnoomp -O3
- General C flags for compiling
-O3
- Preprocessor defs for IFS source
linux LINUX LITTLE LITTLE_ENDIAN POINTER_64 BLAS

2. Intel:

- General F90 flags for compiling
-O2 -g -traceback -vec-report0 -r8
- General C flags for compiling
-O2 -g -traceback
- Preprocessor defs for IFS source
linux LINUX LITTLE LITTLE_ENDIAN POINTER_64 BLAS

Simulations:

- Standalone IFS, launched using the *run-atm.sh* EC-Earth 3.1 script.
- Two month simulations starting 1990-01-01.
- T255L91 and T511L91 resolutions.

Results:

- Regression/benchmark tests of the optimized codes were performed for both resolutions. These tests involve calculating the divergence of the spectral norms for the optimized and unoptimized (-O0) cases after 24 model timesteps. This is a standard benchmark used at ECMWF (http://cisl.ucar.edu/dir/CAS2K11/Presentations/mozdzynski/George-raps_v2.pdf). Divergence of 1% or greater is deemed to be a failure. Results from Table 1 show that all cases (CRAY/Intel and T255/T511) are well within this criteria on Beskow.

	T255L91	T511L91
CRAY	0.11757%	0.17031%
Intel	0.37547%	0.15466%

Table 1: Divergence of the spectral norms (optimized versus unoptimized compilations) after 24 model timesteps. Greater or equal to 1% divergence is considered a failure.

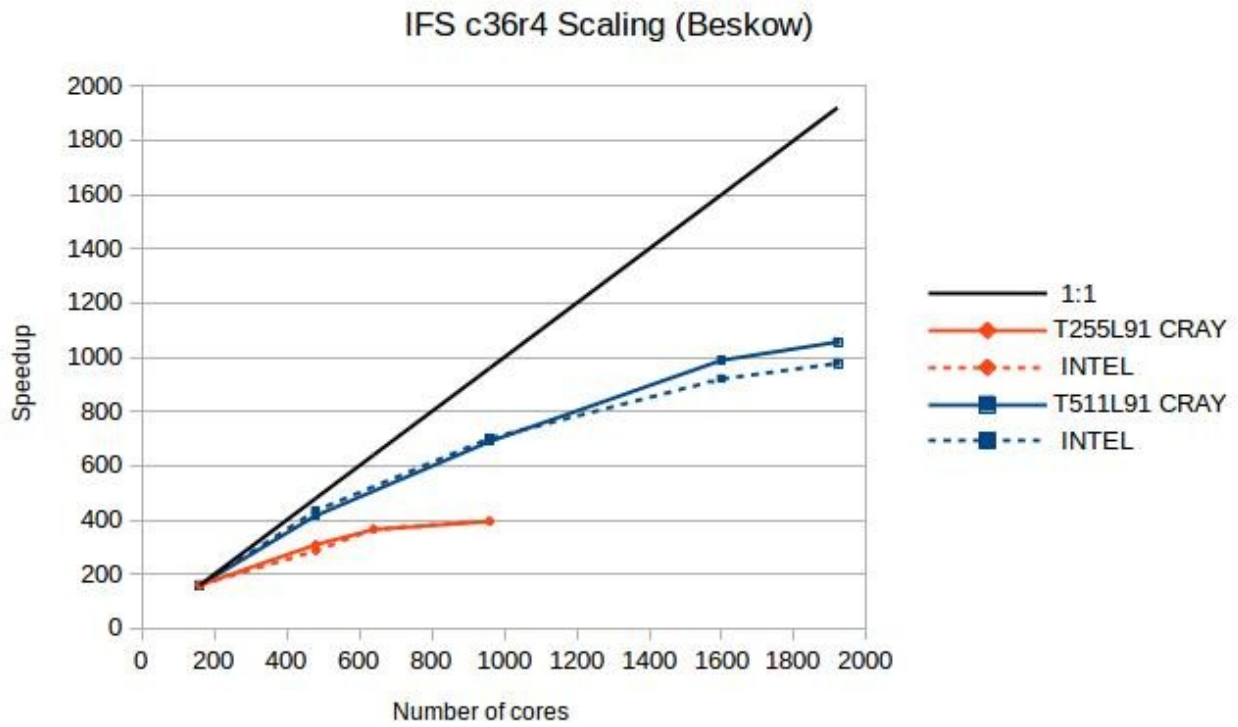


Figure 1: Scaling curves for T255 and T511, CRAY and Intel compilations. All results normalized to the 1:1 line for 160 cores (5 Beskow nodes)

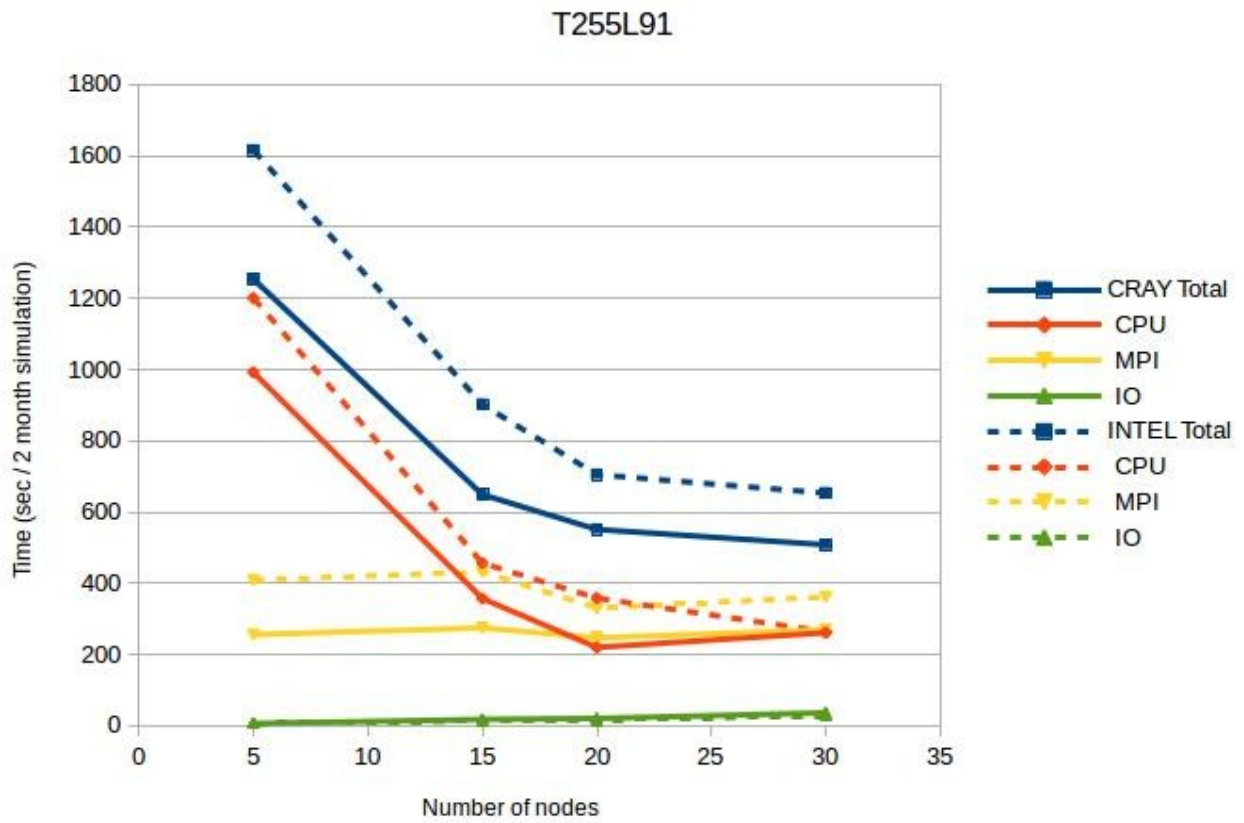


Figure 2: Allinea Performance Report output for T255L91 IFS (CRAY and Intel compilations).

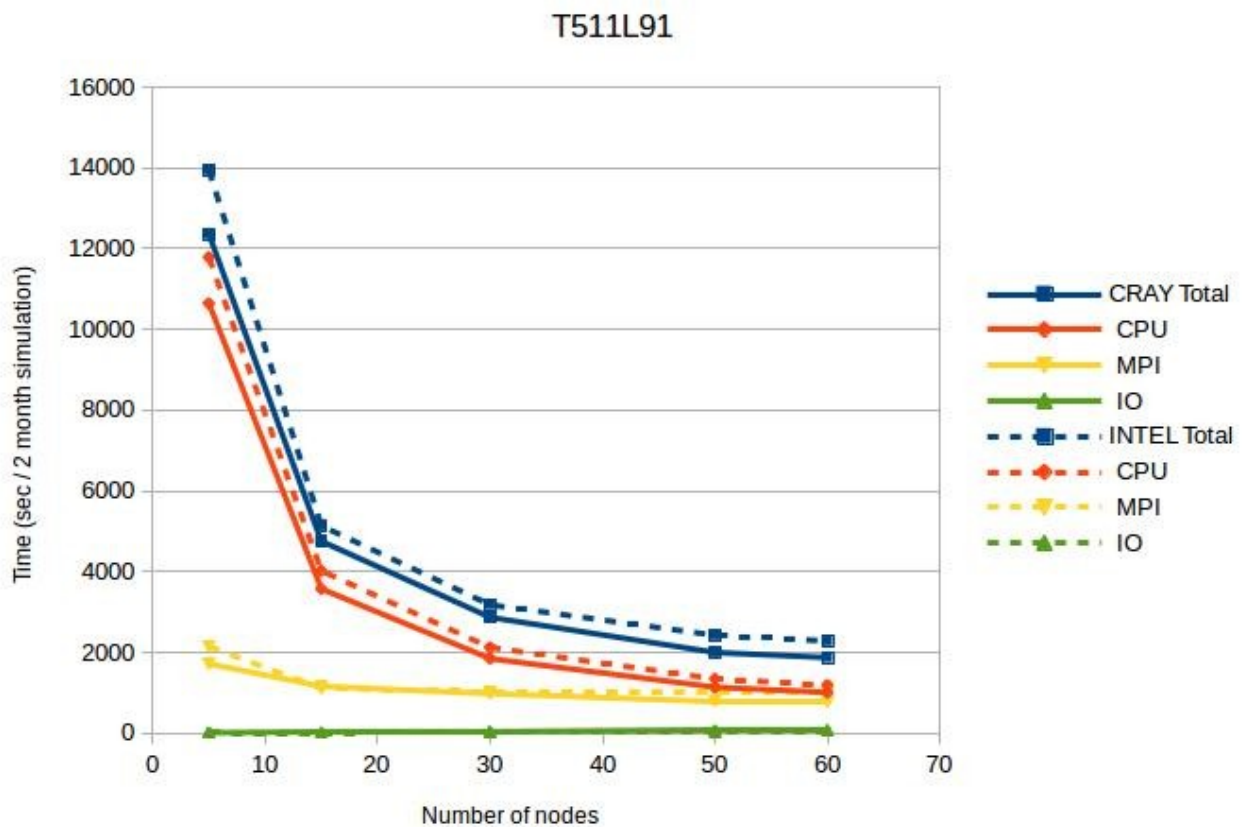


Figure 3: Allinea Performance Report output for T511L91 IFS (CRAY and Intel compilations).

Conclusions

- For IFS, the CRAY compilation shows better performance and scaling than Intel
- IO does not significantly contribute to simulation times
- T255: Beyond ~25 nodes the model becomes MPI bound (time spent in MPI communication greater than CPU and IO)
- T511: The model model is still CPU bound at 60 nodes, and fails to successfully launch at 70 nodes.