

# NEMO Performance assessment

---

ISC BoF

June 21, 2016 - ISC

1


Background

# Background/Context

## NEMO presentation



- ▶ NEMO is an oceanographical application developed by an European consortium.
- ▶ ATOS is working with NEMO developers to improve its performances.
- ▶ We asked to BSC and POP to work on NEMO profiling to:
  - Evaluate the BSC tools
  - find some hints to improve the NEMO behavior
- ▶ A Priori
  - NEMO is a memory bandwidth limited application
  - not optimal scalability.



# Background/Context

## BSC deliverable

---



- ▶ From traces obtained by ATOS, BSC provided
  - a very clear and useful report on the NEMO analysis.
  - several hints on the work which could be done to improve the scalability, but also the vectorization.
- ▶ BSC expert are very reactive and propose an excellent service
- ▶ The next step will be to work with BSC on the memory pattern in order to decrease the memory bandwidth pressure

# 2 Analysis

---

# Application structure

- ▶ The spatio-temporal structure of the behavior of the run with for the whole run




Figure 1: Useful duration view for whole trace at 240 processes

- ▶ In the yellow box, the region of interest (ROI) with iterative computations

# ROI (Region of interest)

- ▶ Zoom on only 3 iterations




Figure 2: structure of ROI for 240 processes

- ▶ each lasts longer than 1ms

# Load balance

---

- ▶ extremely low inbalance

	<b>24</b>	<b>48</b>	<b>72</b>	<b>96</b>	<b>240</b>	<b>360</b>	<b>480</b>
<b>Instr. LB</b>	0.9995	0.9992	0.9992	0.9974	0.9991	0.9971	0.9914
<b>Cycles LB</b>	0.9934	0.9992	0.9840	0.9712	0.9811	0.9720	0.9777



# Computational cost of RIO regions

- ▶ Extrae shows that the major part of the code has a poor performance (**I**nstruction **P**er **C**ycle is low)




Figure 6. Evolution of #instructions vs. IPC for the main clusters when scaling from 24 to 480 processes.

3

Conclusion

# Summary observations/recommendations

---

- ▶ BSC gave us several hints to improve NEMO performances :
  - poor IPC which suggest that some part of the code should be restructured
  - some operation could be easely overlap thanks to OpenMP instructions.
- ▶ Services provided by POP are very precise and focused immediately on the point of interest that could allows to improve the performance of your code.
- ▶ “The services are **free of charge** to organizations in the EU!”

# Future work

---

- ▶ BSC have analyzed burst traces with sampling information of runs between 24 and 480 processes of the NEMO code. Traces should be re create with more information about the source code and memory pattern
- ▶ The future work will be focused on :
  - determine which part of the source code that could be optimized
  - make an analysis of the memory pattern



Funded by the  
European Union

Co-ordinated by  **ECMWF**



Thank You!

[www.hpc-escape.eu](http://www.hpc-escape.eu)



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671627