

JENCI Intel Knights Landing: performance and usability assessment for scientific communities

E. Boyer, G. Hautreux (GENCI) - Technological watch group, GENCI-CEA-CNRS-INRIA and French Universities



EQUIP@MESO

 $(-j \vdash N)$

• Promote the use of supercomputing for the benefit of French scientific communities and for industries

- Represent France in the **PRACE** research infrastructure

Specific action to SMEs through the Simseo intiative

- Equip@meso, 15 partners
- Partnerships at the regional level
- More than 6,5 Pflops available on 3 national centers (CINES, IDRIS and TGCC)
- INES FPA Partie calcul intensi (depuis 2015 • In charge of national HPC strategy for civil research UPS (CNRS) Compétences et ressources (depuis 2013) GENC en calcul intensi MAITRISE **D'OUVRAGE** TGCC frastructure de calcul du CEA Exploitation de calculateurs (depuis 2012) pétaflopiques MAITRISES D'ŒUVRE











Frioul: KNL based prototype

- Atos Sequana cell @ CINES, Montpellier (France)
 - 3 partitions of 48 KNL 7250 nodes, 146 Tflops peak
 - 16 Quadrant cluster + Cache memory nodes
 - 16 Quadrant cluster + Flat memory nodes
 - 16 Quadrant cluster + Hybrid memory nodes
- Quadrant mode used for both partitions
 - Poor results in other processor modes
 - BIOS versions should be updated Work in progress with Atos
- High level support
 - Thanks to Atos and Intel teams





- The comparison is made between
 - Frioul (Xeon Phi, Knights Landing, EDR interconnect)
 - Occigen (Broadwell node: 2x14 cores @ 2.60GHz, FDR interconnect)

5,00

4,00

3,00

2,00

1.00

0.00

- The overall performance for those applications at the moment
 - Node to node comparison
 - Energy efficiency is an estimation using TDP ratio
 - → Results for 14 applications
 - Assesment of memory modes, vectorization impact, ...



KNL memory modes



- Impact of MCDRAM
 - Applications ran in « full flat » and « cache » mode
 - Results provided are for the best test case (favorable to flat mode)
 - No real impact of using MCDRAM for our applications
 - Testcases do not fit in MCDRAM, real applications are using memory
 - Lots of indirections leading to be latency bound
 - Currently testing the hybrid mode, it could be the solution



IXPUG BoF ISC17



First conclusions



- Ease of use
 - No trouble with portability
 - Good speed-up in a few days work
 - Energy efficiency
 - Still troubles with scalability
 - Performance of IB drivers ? IB parameters ?
 BIOS fixes ?
- A choice concerning the memory can be made
 - Cache for large infrastructures with a wide variety of users
 - A comparison Flat vs Cache should be done if you only have a few users
- Full opening of the platform in April
 - French community at first
 - Help us to have a better feedback on the platform
 - more applications and new focus on DL applications
 - A Knight Mill platform may be considered









Thank you for your attention

Questions?

eric.boyer@genci.fr

gabriel.hautreux@genci.fr