


HPC vs. Cloud Benchmarking

An empirical evaluation of the performance and cost metrics


Kashif Iqbal* and Eoin Brazil
Kashif.iqbal@ichec.ie
ICHEC, NUI Galway, Ireland



ICHEC - in a nutshell



- **Irish Centre for High-End Computing**
 - National Tier-1 Centre
 - Run Irish National HPC service for Academics
 - PRACE partner
- Interest in understanding the competitive costs
 - Understanding various infrastructures & workloads
 - HPC, HTC, [HPC Cloud](#), [HTC Cloud](#)
- What is the most effective means to address our customers (Academics) needs?



Outline


- Benchmarking - Why, which benchmark?
- NAS Parallel Benchmark (NPB)
- Environment Setup
- Results
- Next Steps



Sitting in a 3.8-metre sea
kayak and watching
a four-metre great
white approach you is
a fairly tense experience


MOTIVATION

If there is a better reason to paddle, I don't know what it is.



Overview

- Diversity
 - Diverse computing infrastructures (HPC, HTC, Cloud)
 - Diverse workloads for various academic communities
- Cost analysis and performance metrics
 - Performance and configuration overhead as indirect costs
- System benchmarking for:
 - Comparison of HPC and HTC systems vs. Cloud offerings
 - Comparison of parallelism techniques (e.g. MPI/OMP)




HPC/HTC Benchmarks

- LINPACK – Top 500
- SPEC06 – CPU intensive benchmark
 - HEP-SPEC06
- HPC Challenge (HPCC)
- Graph 500
- STREAM – for memory bandwidth
- MPPtest – MPI performance
- NAS Parallel Benchmark (NPB)
- ...


NAS Parallel Benchmark

- Open-source and free CFD benchmark
- Performance evaluation of commonly used parallelism techniques
 - Serial, [MPI](#), [OpenMP](#), OpenMP+MPI, Java, HPF
- Customisable for different problem sizes
 - Classes S: small for quick tests
 - Class W: workstation size
 - Classes A, B, C: [standard test problems](#)
 - Classes D, E, F: large test problems



NPB Kernels


Kernel	Description	Problem Size	Memory (MW)
EP	Monte Carlo kernel to compute the solution of an integral – Embarrassingly parallel	2^{30}	18
MG	Multi-grid kernel to compute the solution of the 3D Poisson equation	256^3	59
CG	Kernel to compute the smallest eigenvalue of a symmetric positive definite matrix	75000	97
FT	Kernel to solve a 3D partial difference equation using an FFT based method	512x256x256	162
IS	Parallel sort kernel based on bucket sort	2^{25}	114
LU	Computational Fluid Dynamics (CFD) application using symmetric successive over relaxation	102^3	122
SP	CFD application using the Beam-Warming approximate factorisation method	102^3	22
BT	CFD application using an implicit solution method	102^3	96



Cloud Cluster Setup

- EC2 instance management
 - StarCluster Toolkit
 - <http://web.mit.edu/star/cluster/>
 - StarCluster AMIs – Amazon Machine Image
 - Resource manager plugin
- Login vs. compute instances
 - EC2 small instance as login node
 - File system shared via NFS across nodes

Cloud vs. HPC

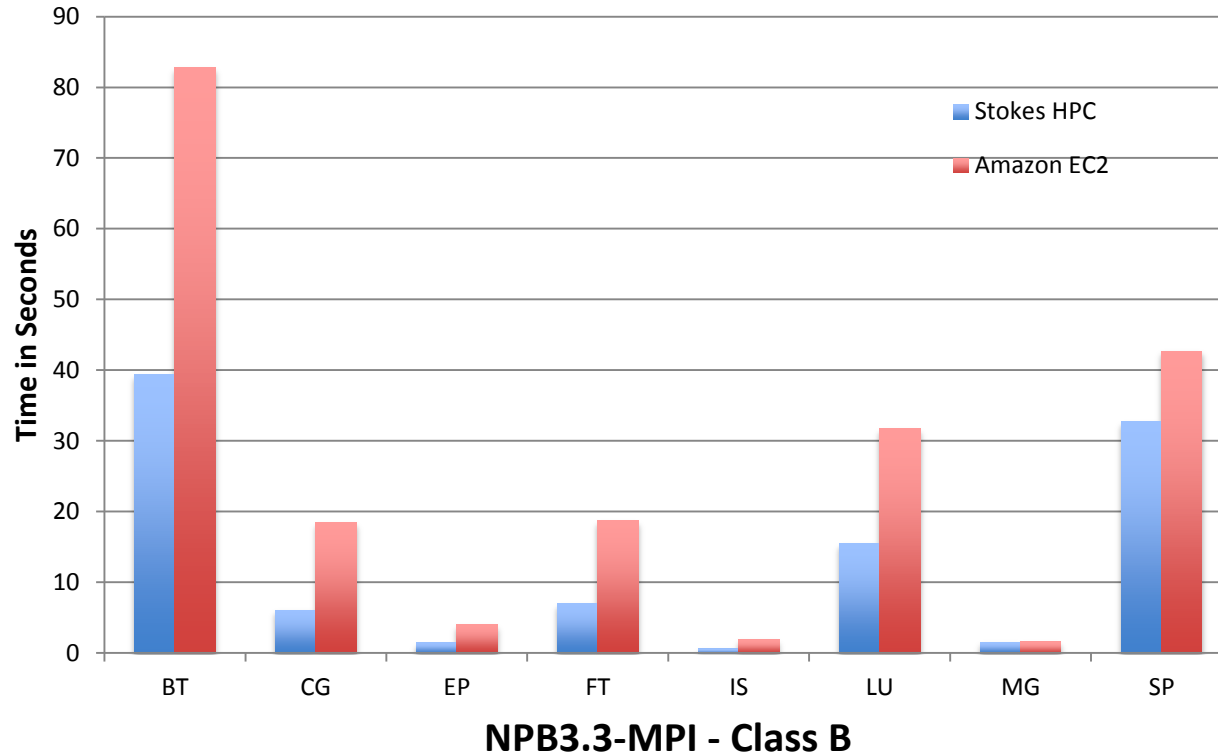


	Amazon EC2	Stokes HPC
Compute Node	23 GB of memory, 2 x Intel Xeon X5570, quad-core "Nehalem"	24 GB memory, 2 x Intel Xeon E5650, hex-core "Westmere"
Connectivity	10 Gigabit Ethernet	ConnectX Infiniband (DDR)
OS	Ubuntu, 64-bit platform	Open-SUSE, 64-bit platform
Resource manager	Sun Grid Engine	Torque
Compilers & libraries	Intel C, Intel Fortran, Intel MKL, Intel MVAPICH2	Intel C, Intel Fortran, Intel MKL, Intel MVAPICH2


- Non-trivial to replicate runtime environments
- Large variations in performance possible
- Logical vs. Physical cores
 - HT/SMT – Hyper or Simultaneous Multi-Threading

NPB – MPI

BT and SP using 16 cores, rest using 32 cores (22 runs)




The average performance loss was **48.42%**
(ranging from 1.02% to 67.76%).



NPB - OpenMP

8 cores with 8 OMP Threads (22 runs)



The average performance loss was **37.26%**
(ranging from 16.18 - 58.93%)

Cost



- 720 hours @ **99.29 USD** 😊
 - ~100 % utilisation
 - Compute cluster instance @ \$1.300 per Hour
 - Small instance @ \$0.080 per Hour
- Other useful insights:
 - Spot instances
 - Overheads (performance, I/O, setup)
 - Data transfer costs and time

Conclusions




- As expected a purpose built HPC cluster outperforms EC2 cluster for same number of OMP threads
 - Average performance loss over all NPB tests: **~37%**
- Similarly so for when comparing 10GigE versus Infiniband networking fabrics
 - Average performance loss over all NPB test: **~48%**
- Even at a modest problem size the differences in performances between systems is highlighted.

Next steps

- HTC vs. Cloud Benchmarking
 - HEP-SPEC on the virtualised EGI resources
 - and EC2 instances (small, medium, large)
- Benchmarking results in the cost model
 - As an extra weight in addition to monetary costs
- Publications





Thank you for your attention!

Questions??

kashif.iqbal@ichec.ie