

High Performance Computing solutions



Benchmarks Report



June 2007

Damien Déclat
damien.declat@bull.net
+33 (0) 4 76 29 71 93

Carole Vassard
carole.vassard@bull.net
+33 (0) 4 76 29 74 43

BULL SA — Centre HPC
1 rue de Provence
F-38130 Échirolles, France

Introduction

The aim of this document is to summarize the performances obtained for different benchmarks, such as GROMACS, HPCC, TauBench, ParPac and DDFEM.

Following the proposed configurations, we will give the results obtained for the different types of Xeon nodes (4-cores nodes and 8-cores nodes) and the results that will be obtained on the target configuration (**72 nodes of 8-cores nodes made up of 2 sockets of 4 cores – 2.66 GHz Cloverton processors**)

Note that for the GROMACS benchmarks, the version that was used is version 3.3.1. The test cases were done for single (float) and double precision with the FFTW 3.1.2 version library.

Description of the benchmarks environment

Two types of architectures have been used for the benchmarks. Following the specifications from LNCC, the clusters are equal equipment than the ones that are proposed.

The first type of architecture (A1) is a cluster of 4-cores Xeon nodes composed of:

- **72 compute nodes, R421 Bull Servers,**
- **For each nodes, 2 Xeon Sockets with 2 cores (3 GHz Woodcrest processor),**
- **4 Mo for the L2 cache size per socket,**
- **8 Go of memory per nodes,**
- **an InfiniBand interconnect (DDR 4x).**

The second kind of nodes (A2) used for the benchmarks is a **R421 Bull Server node composed of 2 Xeon Sockets with 4 cores (2,4GHz Cloverton processor).**

These two kinds of nodes will thus give information about the performances of the applications on a single 8-cores node but also the scalability of the codes. Using these two kinds of information, we will give the attended results on the target configuration.

The **software environment** used for the benchmarks is based on the Intel compilers and libraries:

- **Fortran/C/C++ Intel version 9.1/10.0 compilers,**
- **MKL version 9.0 numerical libraries,**

Whereas the MPI paradigm is managed by the **MPI-Bull2** library.

These environments for the benchmarks are thus reproducing the environment proposed for the target configurations.

SUMMARY OF THE RESULTS

This first part gives a summary of the different results for the 4 applications (Gromacs, ParPac, TauBench and DDFEM) and the different test case on the target configuration. The HPCC results follow this part.

1 Gromacs

Serial Single (Float) (ns/day)	
Villin	18,784
Lys/Cut	5,74
Lys/PME	3,7
DPPC	0,443
Poly-CH2	8,06

Paralelo Single (Float) (ns/day)	
Villin	73,662
Lys/Cut	25,97
Lys/PME	17,22
DPPC	1,748
Poly-CH2	34,2

Serial Double (Float) (ns/day)	
Villin	12,191
Lys/Cut	3,79
Lys/PME	2,62
DPPC	0,287
Poly-CH2	6,493

Paralelo Double (Float) (ns/day)	
Villin	39,9
Lys/Cut	17,41
Lys/PME	11,52
DPPC	0,937
Poly-CH2	20,2

2 ParPac

Numbers of cores	1	2	4	8	16	32	64
Total Simulation Time (sec)	133,57	159,18	185,93	194,18	214,97	196,08	202,88
Number of point /sec (millions)	1,12	2,06	3,16	6,18	11,92	23,93	47,33
GFLOPS	0,67	1,23	1,88	3,69	7,12	14,29	28,26

3 TauBench

Numbers of cores	1	2	4	8	16	32	64
MFLOPS	1069,325	2026,699	3233,845	6394,331	12870,511	34617,448	51230,017
Computing time	13,943	14,517	17,958	18,174	18,008	18,020	18,008
Communication Time	0,020	0,217	0,511	0,506	0,645	0,535	0,645
CommRatio	0,009	0,013	0,025	0,025	0,032	0,027	0,032

HPCC

HPCC is a test suite that gathers different benchmarks such as HPL, PTRANS, Star DGEMM, Single DGEMM, Star TREAM and Single STREAM. The results for these benchmarks on the target configuration are the following:

PTRANS (GB/s)	27
HPL (Tflops)	3,02
StarDGEMM (Gflops/s)	7,92
SingleDEGEMM (Gflops/s)	9,47
STREAM - COPY (GB/s)	0,741
STREAM - SCALE (GB/s)	0,74
STREAM - ADD (GB/s)	0,737
STREAM - TRIAD (GB/s)	0,742
SingleStream - COPY (GB/s)	4,05
SingleStream - SCALE (GB/s)	4,05
SingleStream - ADD (GB/s)	4,05
SingleStream - TRIAD (GB/s)	4,05
MPIRandom (GUP/s)	0,15
StarRandom (GUP/s)	0,006
SingleRandom (GUP/s)	0,019
MPIFFT (Gflops/s)	26
StarFFT (Gflops/s)	0,47
SingleFFT (Gflops/s)	1,145
Ping Pong - Latência (msecs)	0,006
Ping Pong - Largura de banda (MB/s)	1000
Ring - Latência - Natural (msecs)	0,007
Ring - Largura de banda – Natural (MB/s)	340
Ring - Latência - Random (msecs)	0,02
Ring - Largura de banda – Random (MB/s)	101

GROMACS

GROMACS is a versatile package to perform molecular dynamics, i.e. simulate the Newtonian equations of motion for systems with hundreds to millions of particles.

It is primarily designed for biochemical molecules like proteins and lipids that have a lot of complicated bonded interactions, but since GROMACS is extremely fast at calculating the non bonded interactions (that usually dominate simulations) many groups are also using it for research on non-biological systems, e.g. polymers.

The benchmarks are divided in two parts; the first concerns performance on single node machines (i.e. no parallelization outside the machine, but possibly using several CPUs), while the second one shows the scaling capabilities for a large membrane system when running in parallel over an external network.

The following paragraphs give the results obtained for single and double precision executions on the different benchmarks environments and finally the results attended on the target cluster. It is important to notice that using mkl library instead of the FFT 3.1.2 version library increases the performances and the efficiency of gromacs especially up to 4 processes (and for double precision).

1 Single Precision Results for Serial (1 core) and Parallel (8 cores)

The following tables describe the results for the GROMACS test cases on the target architecture (8-cores nodes):

For the Serial Single test case:

Serial Single (Float) (ns/day)	
Villin	18,784
Lys/Cut	5,74
Lys/PME	3,7
DPPC	0,443
Poly-CH2	8,06

For the Parallel Single test case:

Paralelo Single (Float) (ns/day)	
Villin	73,662
Lys/Cut	25,97
Lys/PME	17,22
DPPC	1,748
Poly-CH2	34,2

2 Double Precision Results for Serial (1 core) and Parallel (8 cores)

The following tables describe the results for the GROMACS test cases on the target architecture (8-cores nodes):

For the Serial Double test case:

Serial Double (Float) (ns/day)	
Villin	12,191
Lys/Cut	3,79
Lys/PME	2,62
DPPC	0,287
Poly-CH2	6,493

For the Parallel Double test case:

Paralelo Double (Float) (ns/day)	
Villin	39,9
Lys/Cut	17,41
Lys/PME	11,52
DPPC	0,937
Poly-CH2	20,2

ParPac

The ParPac Application Benchmark is a lattice Boltzmann code, which simulates the dynamics of a fluid streaming through a porous structure. It calculates the structure's permeability and returns the calculation performance in Gflops.

The results on the ParPac benchmarks on architecture A1 (4-cores nodes 3GHz) are the following :

Numbers of cores	1	2	4	8	16	32	64
Total Simulation Time (sec)	119,36	141,49	169,56	177,09	196,05	178,82	185,02
Number of point /sec (millions)	1,257	2,314	3,461	6,779	13,076	26,25	51,909
GFLOPS	7,51E-01	1,382	2,066	4,047	7,806	15,67	30,99

When using 8-cores nodes (2,4 GHz) – Architecture A2 - , we obtain :

Numbers of cores	1	2	4	8
Total Simulation Time (sec)	148,04	176,42	206,07	372,47
Number of point /sec (millions)	1,014	1,856	2,847	3,214
GFLOPS	6,05E-01	1,108	1,7	1,919

Considering the single core test case, the ratio of the GFLOPS (7,51/6,05) corresponds to the ratio of the frequencies of the processors (3/2,4). This is also the case for 2 and 4 cores runs.

On 8 cores, we are running on 2 nodes for architecture A1 whereas we are running on 1 single node for Architecture A2.

The performance obtained for A1 is better than the one obtained for A2.

This phenomenon can be explained by 2 facts:

1. the communication pattern is well supported through the network. Therefore, there is no loss of performances when doing intra-nodes runs.
2. ParPac needs a high memory bandwidth. When using 4 processes per memory bus (for 8-cores nodes), the performances can much decrease than when using 2 processes per memory bus.

Therefore, in order to get the best performances, we will choose to run on the target configuration only 4 processes per nodes. To extrapolate these results we consider:

- The scalability obtained with Architecture A1,
- The fraction of frequencies to get the results from architecture A2 to the target configuration (from 1 to 4 processes runs).

Thus we will obtain on the target configuration :

Numbers of cores	1	2	4	8	16	32	64
Total Simulation Time (sec)	133,57	159,18	185,93	194,18	214,97	196,08	202,88
Number of point /sec (millions)	1,12	2,06	3,16	6,18	11,92	23,93	47,33
GFLOPS	0,67	1,23	1,88	3,69	7,12	14,29	28,26

TauBench

The TauBench is an unstructured grid benchmark. The respective kernels are derived from Tau -- a Navier Stokes solver which has been developed at the German aerospace research DLR in Germany.

The results on the TauBench benchmarks on architecture A1 (4-cores nodes 3GHz) are the following :

Numbers of cores	1	2	4	8	16	32	64
MFLOPS	1207,16	2291,13	3620,60	7229,84	14424,67	28779,68	57445,12
Computing time	5,23	5,44	6,76	6,82	6,78	6,77	6,80
Communication Time	0,01	0,09	0,24	0,19	0,25	0,30	0,26
Comm Ratio	0,02	0,02	0,04	0,03	0,04	0,05	0,04

When using 8-cores nodes (2,4 GHz) – Architecture A2 - , we obtain :

Numbers of cores	1	2	4	8
MFLOPS	964,486	1814,972	2965,240	3976,883
Computing time	19,657	20,609	25,134	37,033
Communication Time	0,034	0,319	0,485	1,171
Comm Ratio	0,020	0,015	0,019	0,032

As for ParPac, when considering the single core test case, the fraction of the GFLOPS (1,207/0,964) corresponds to the fraction of the frequencies of the processors (3/2,4). This is also the case for 2 and 4 cores runs.

On 8 cores, we are running on 2 nodes for architecture A1 whereas we are running on 1 single node for Architecture A2.

The performance obtained for A1 is better than the one obtained for A2.

This phenomenon can be explained by 2 facts:

1. the communication pattern is well supported through the network. Therefore, there is no loss of performances when doing intra-nodes runs.
2. TauBench needs a high memory bandwidth. When using 4 processes per memory bus (for 8-cores nodes), the performances can much decrease than when using 2 processes per memory bus.

Therefore, in order to get the best performances, we will choose to run on the target configuration only 4 processes per nodes. To extrapolate these results we consider:

- The scalability obtained with Architecture A1,
- The fraction of frequencies to get the results from architecture A2 to the target configuration (from 1 to 4 processes runs).

Thus we will obtain on the target configuration :

Numbers of cores	1	2	4	8	16	32	64
MFLOPS	1069,325	2026,699	3233,845	6394,331	12870,511	34617,448	51230,017
Computing time	13,943	14,517	17,958	18,174	18,008	18,020	18,008
Communication Time	0,020	0,217	0,511	0,506	0,645	0,535	0,645
Comm Ratio	0,009	0,013	0,025	0,025	0,032	0,027	0,032