Quantum Espresso benchmark

SGI-IZO

December 15, 2008

Abstract

In this report a brief performance analysis of Quantum Espresso is made by using pw.x and ph.x in itanium2, opteron, core2duo and pentium architectures.

1 Hardware description

1.1 Itanium2 1.6 (dc)

Itanium2 dual core processors at 1.6 GHz. 8 core nodes connected with infiniband network. HPSFS shared file system. Quantum Espresso was compiled with intel v10.0.023 compilers, fftw, mkl 10.0 libraries and mpiBull2.

1.2 Itanium2 1.6

Itanium2 single core processors at 1.6 GHz. 8 core nodes connected with infiniband network. HPSFS shared file system. Quantum Espresso was compiled with intel v10.0.023 compilers, fftw, mkl 10.0 libraries and hpmpi.

1.3 Itanium2 1.3

Itanium2 single core processors at 1.3 GHz. 4 core nodes connected with infiniband network. HPSFS shared file system. Quantum Espresso was compiled with intel v10.0.023 compilers, fftw, mkl 10.0 libraries and hpmpi.

1.4 Opteron 2.4

Opteron dual core processors at 2.4 GHz. 8 core nodes connected with infiniband network. HPSFS shared file system. Quantum Espresso was compiled with intel v10.1.017 compilers, fftw, mkl 10.3 libraries and hpmpi. PGI compiler was not used because our old version has a bug that avoids compile espresso.

1.5 Core2duo 2.1

Core2duo dual core processors at 2.1 GHz. 2 core nodes connected with eth100 network. NFS shared file system. Quantum Espresso was compiled with intel v10.0.023 compilers, fftw, mkl 10.0 libraries and Openmpi. These nodes belong to our Péndulo grid that use classrooms PCs at night for computing and this is the reason of the commodity 100eth network. The different labels corresponds to the following situations:

2.1

Local disk is used to store temporary files.

NFS

A NFS file system is used to store temporary files.

noIO

disk_io='none' is used to reduce the Input/Output.

1.6 Pentium 2.4

Core2duo dual core processors at 2.1 GHz. 2 core nodes connected with eth100 network. NFS shared file system. Quantum Espresso was compiled with intel v10.0.023 compilers, fftw, mkl 10.0 libraries and Openmpi. These nodes belong to our Péndulo grid that use classrooms PCs at night for computing and this is the reason of the commodity 100eth network.

2 Benchmark description

2.1 pw.x

FeAl system with 12 Fe atoms and 4 Al atoms. 120 irreducible k-points.

2.2 ph.x

AlAs system with 1 atom per specie. 60 irreducible k-points.

3 Results

We have measured the walltime used by the calculations in different configurations. For the pw.x case the results are shown in the table 1. Table 2 is a variation of the first one in which the results are shown in terms of the speed up of the calculation, the ideal case is a speed up equal to the number of cores. Nevertheless is better to pay more attention to table 1.

First we observe that the 1 core calculation in itanium 1.3 system is very slow. It improves a lot, or shows more normal results, when using more cores and this is the reason of the high speed up numbers.

Opteron nodes also seems to work better with calculations of more than 8 cores.

In the Péndulo grid (core2duo and pentium) the network is a clear bottleneck but through the NFS file system, not the interprocessor communication. Therefore we recover a perfect scaling when we reduce the IO traffic by using *disk_io='none'* in the input file.

Similar results are obtained for ph.x in table 3 and table 4. Just remark that opteron node runs better ph.x than pw.x in this architecture, compared with the others.

	1 core	2 cores	4 cores	8 cores	16 cores	32 cores
Itanium2 1.6(dc)	5700		1592	744	483	289
Itanium2 1.6	6540		1795	831	473	
Itanium2 1.3	9240		1850	968	601	371
Opteron 2.4	9420		2755	1047	579	400
Core2duo 2.1	5460	3000				
Core2duo (nfs)	7200	4260	2987			
Core2duo (noIO)		3457	1771	893	557	351
Pentium 2.4	7440	5040	3064	2232		

Table 1: PW.x performace in seconds.

	1 core	2 cores	4 cores	8 cores	16 cores	32 cores
Itanium2 1.6 (dc)	1		3.58	7.66	11.8	19.72
Itanium2 1.6GHz	1		3.64	7.87	13.83	
Itanium2 1.3	1		4.99	9.55	15.37	24.91
Opteron 2.4	1		3.42	9	16.27	23.55
Core2duo 2.1	1	1.82				
Core2duo (nfs)	1	1.69	2.41			
Core2duo (noIO)	1	1.74	3.39	6.72	10.77	17.09
Pentium 2.4	1	1.48	2.43	3.33		

Table 2: PW.x performace in terms of speed up.

	1 core	2 cores	4 cores	8 cores	16 cores
Itanium2 1.6 (dc)	6060		1548	796	452
Itanium2 1.6GHz	7680		1963	1018	540
Itanium2 1.3	9240		2378		
Opteron 2.4	6480		1861	775	412
Core2duo 2.1	4980	2603			
Core2duo (nfs)	6120	3900	2129	2850	2852
Core2duo (noIO)	6060	4560	2376	1954	2311
Pentium 2.4	8520	4500	2466		

Table 3: pw.x performace in seconds.

	1 core	2 cores	4 cores	8 cores	16 cores
Itanium2 1.6 (dc)	1		3.91	7.61	13.41
Itanium2 1.6GHz	1		3.91	7.54	14.22
Itanium2 1.3	1		3.89		
Opteron 2.4	1		3.48	8.36	15.73
Core2duo 2.1	1	1.91			
Core2duo (nfs)	1	1.57	2.87		
Core2duo (noIO)	1	1.33	2.55	3.1	2.62
Pentium 2.4	1	1.89	3.45		

Table 4: ph.x performace in terms of speed up.