SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures

INFRA-2007-2.2.2.1 - Preparatory phase for 'Computer and Data Treatment' research infrastructures in the 2006 ESFRI Roadmap

PRACE

Partnership for Advanced Computing in Europe

Grant Agreement Number: RI-211528

D5.4
Report on the Application Benchmarking Results of Prototype Systems

Final (Revised)

Version: 3.0
Author(s): Mark Bull, EPCC
Date: 18.06.2010

Project and Deliverable Information Sheet

<table>
<tr><td rowspan="8">PRACE Project</td><td colspan="2">Project Ref. №: RI-211528</td></tr>
<tr><td colspan="2">Project Title: Partnership for Advanced Computing in Europe</td></tr>
<tr><td colspan="2">Project Web Site: http://www.prace-project.eu</td></tr>
<tr><td colspan="2">Deliverable ID: D5.4</td></tr>
<tr><td colspan="2">Deliverable Nature: DOC_TYPE: Report</td></tr>
<tr><td rowspan="2">Deliverable Level:
RE *</td><td>Contractual Date of Delivery:
30 / June / 2010</td></tr>
<tr><td>Actual Date of Delivery:
30 / June / 2010</td></tr>
<tr><td colspan="2">EC Project Officer: Bernhard Fabianek</td></tr>
</table>

* - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	**Title:** Report on the Application Benchmarking Results of Prototype Systems	
	ID: **D5.4**	
	Version: 3.0	**Status:** Final
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2003	
	File(s): D5.4_v3.0	
Authorship	**Written by:**	Mark Bull, EPCC
	Contributors:	Lukas Arnold, FZJ Joachim Hein, EPCC Albert Farres, BSC Orlando Rivera, LRZ Andrew Sunderland, STFC Xavi Saez, BSC John Donners, SARA Pekka Manninen, CSC Sebastian von Alfthan, CSC Harald Klimach, HLRS Bertrand Cirou, CINES Xu Guo, EPCC Jussi Enkovaara, CSC Raúl de la Cruz, BSC Mauricio Araya, BSC Martin Polak, GUP David Henty, EPCC Paschalis Korosoglou, GRNET Andrew Porter, STFC Carlo Cavazzoni, CINECA Andrea Babuke, HLRS Fernando Nogueira, UC-LCA Andriy Chut, HLRS
	Reviewed by:	Dietmar Erwin, FZJ; Manuel Fiolhais, UC-LCA
	Approved by:	Technical Board

Document Status Sheet

Version	Date	Status	Comments
0.1	22/May/2009	Draft	
0.2	03/June/2009	Revised draft	
0.3	09/June/2009	For internal review	
1.0	26/June/2009	For TB approval	
1.1	26/Nov/2009	First revised draft	
1.2	30/Nov/2009	Second revised draft	
1.3	04/Dec/2009	For internal review	
2.0	16/Dec/2009	For TB approval	
2.1	11/Jun/2010	For internal review	
3.0	18/Jun/2010	For TB approval	

Document Keywords and Abstract

Keywords:	PRACE, HPC, Research Infrastructure, Prototype, Assessment, Benchmark, Applications
Abstract:	This deliverable describes the assessment of seven PRACE prototypes selected previously as candidates of likely Petaflop/s systems in 2009/10 by running applications benchmarks. The applications used are those identified by WP6 as being representative of the current and future usage of major European HPC systems. This deliverable contains assessment results for the IBM BlueGene/P at FZJ, the Cray XT5 at CSC, the IBM Power 6 at SARA, the IBM Cell at BSC, the NEC SX-9/x86 system at HLRS, the Sun x86 cluster at FZJ and the Bull x86 cluster at CEA.

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

[1] http://www.prace-project.eu.
[2] Planning document for the installation of the prototype systems, PRACE Deliverable D5.1.1, June 2008.
[3] Final benchmark suite, PRACE Deliverable D6.3.2, December 2009.
[4] Report on available performance analysis and benchmark tools, representative benchmark, PRACE Deliverable D6.3.1, November 2008.
[5] Final report on application requirements, PRACE Deliverable D6.2.2, March 2008.
[6] Installation report prototype systems, PRACE Deliverable D5.1.2, December 2008.
[7] Identification and categorisation of applications and initial benchmarks suite, PRACE Deliverable D6.1, June 2008.
[8] Preliminary report on application requirements, PRACE Deliverable D6.2.1, March 2008.
[9] Comparison and recommendations of Petaflop/s systems, PRACE Deliverable D5.5, December 2009.
[10] Final assessment of Petaflop/s systems to be installed in 2009/2010, PRACE Deliverable D7.1.3, June 2009.
[11] Technical requirement for the second Petaflop/s systems, PRACE Deliverable D7.5.2, November 2009.
[12] Jülich Benchmarking Environment (JuBE): http://www.fz-juelich.de/jsc/jube/
[13] Top500 Supercomputing sites: http://www.top500.org

List of Acronyms and Abbreviations

BSC	Barcelona Supercomputing Centre (Spain)
CEA	Commissariat à l'Energie Atomique (France)
CPU	Central Processor Unit
CSC	Center for Scientific Computing (Finland)
Flop	Floating –point operations (usually 64-bit, i.e. Double Precision, DP)
Flop/s	Floating –point operations per second (usually 64-bit, i.e. DP), also F/s
FZJ	Forschungszentrum Juelich GmbH (Germany)
GFlop/s	Giga (= 10^9) Floating–point operations per second (usually 64-bit), also GF/s
HLRS	Höchstleistungsrechenzentrum Stuttgart, High Performance Computing Centre Stuttgart (Germany)
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
JuBE	Jülich Benchmarking Environment
MPI	Message Passing Interface
MPP	Massively Parallel Processing
PFlop/s	Peta (= 10^{15}) Floating–point operations per second (usually 64-bit), also PF/s
PPE	Power Processing Element: component of the Cell Broadband Engine Architecture
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
SARA	Stichting Academisch Rekencentrum Amsterdam (Netherlans)
SMT	Simultaneous Multi threading. Issue multiple instructions from multiple threads in one cycle of a chip
SPE	Synergistic Processing Element: component of the Cell Broadband Engine Architecture
SVN	Subversion. Version control system
TFlop/s	Tera (= 10^{12}) Floating–point operations per second (usually 64-bit), also TF/s
WP	PRACE Work Package
WP5	PRACE Work Package 5 – Deployment of prototype systems
WP6	PRACE Work Package 6 – Software enabling for Petaflop/s systems
WP7	PRACE Work Package 7 – Petaflop/s systems for 2009/2010

Executive Summary

This report is the result of Task 5.4, "Evaluation and benchmarking of user applications" within Work Package 5 "Deployment of prototype systems".

The approach taken is to provide quantitative data measured with application benchmarks, provided by Work Package 6, to characterise the performance of the prototype systems as a whole.

In order to make meaningful comparisons between prototype systems with very distinct hardware architectures, a key feature of the approach is to measure the performance of applications on partitions of the various systems which have the same nominal peak floating point performance rather than, say, the same number of cores or processing elements.

1 Introduction

This report is a further updated version of D5.4, first submitted to the European Commission on June 30, 2009 and later revised on December 17, 2009. In the extension phase of the project, additional data has been collected from most of the prototype systems.

This chapter provides background information and defines the objectives, scope and structure of the document.

1.1 Background and Purpose

The PRACE project [1] has the overall objective of preparing for the creation of a persistent pan-European HPC service. PRACE is divided into eight inter-linked work packages (WP).

Work package 5 (WP5) within PRACE is titled "Deployment of prototype systems" and is responsible for installing and assessing operational prototypes in close to production conditions.

Based on the recommendations from WP7 and the PRACE Technical Board (TB), the PRACE Management Board (MB) has selected six systems that will be used as prototypes for production systems delivering about one Petaflop/s in 2009/2010. These prototype systems [2] are:

- IBM BlueGene/P at FZJ,
- Cray XT5 at CSC,
- IBM Power6 at SARA,
- IBM Cell/Power6 at BSC,
- NEC SX-9/x86 system at HLRS,
- Sun x86 cluster at FZJ
- Bull x86 cluster at CEA.

The assessment of the operational prototype systems in a close to production environment will result in recommendations for the selection of the first Petaflop/s production systems in 2009/2010.

Once the prototypes have been installed and integrated into a production-like environment, Task 5.4 in WP5 is responsible for application benchmarking on the prototype systems. The applications used are those identified in D6.3.2 [3]. The requirements for most of these applications were characterised in D6.2.2 [5].

Full descriptions of the prototypes are available in D5.1.2 [7].

1.2 Objectives

The objectives of this document are to:

1. Define the application benchmarking activities.
2. Record the results of the application benchmarking for each prototype system.
3. Analyse the results for each prototype system.
4. Provide information for the requirements for the first Petaflop/s systems.

5. Inform the decision on selecting production Petaflop/s systems.
6. Provide data to support the creation of a representative and scalable set of benchmarking codes.
7. Demonstrate that milestone M5.4.1 has been met.

1.3 Dependencies

Input from other PRACE tasks:

- D5.1.2 , reporting on the installation of prototype systems.
- D6.1 [7], identifying representative applications,
- D6.2.1[8] and D6.2.2 [5], reporting on applications requirements.
- D6.3.2 [3], describing the final benchmark suite

Output for other PRACE tasks:

- D5.5 [9], comparison and recommendations of Petaflop/s systems
- D6.3.1 [3], defining representative benchmarks,
- D7.1.3 [10], final assessment of Petaflop/s systems,
- D7.5.2 [11], technical requirements for procurement.

1.4 Structure of the Document

The document is divided into chapters.

Chapter 1 provides this introduction.

Chapter 2 explains the assessment methodology, including the approach taken.

Chapter 3 provides a brief description of each application benchmark, details of the test case(s) used, the results of the assessments, and an analysis of these results.

Chapter 4 provides a summary of the Chapter 3 results for each prototype.

Chapter 5 provides conclusions from the assessments and outlines future work.

Chapter 6 is an Appendix containing tables of raw data.

1.5 Audience

The intended audience for this document will have a technical background in HPC.

1.6 Changes between Version 2.0 and Version 3.0

The following is a summary of the changes in this document between Version 2.0 (submitted in M24) and Version 3.0 (submitted at the end of the extension period):

- **CODE_SATURNE**: added new larger test case, with data for IBM BlueGene/P, Cray XT5, IBM Power6, NEC x86 cluster and Sun x86 cluster.
- **CP2K**: added data for Cray XT5
- **EUTERPE**: added data for Cray XT5
- **NS3D**: added data for Sun x86 cluster

- **QCD**: added data for NEC x86 cluster
- **QUANTUM_ESPRESSO**: added data for NEC x86 cluster.
- **WRF**: added data for NEC x86 cluster
- **ELMER**: revised data for IBM Power6 and Cray XT5, added new data for Sun x86 cluster.
- **OCTOPUS**: additional data for Cray XT5

2 Assessment Methodology

2.1 Background

The process of application benchmarking consists of running full-scale, realistic application codes on the systems of interest. A benchmark is defined by the application code, together with one or more test cases. Each test case consists of a well-defined set of input parameters and data which define a calculation to be performed. The measurement made for each test case consists of a single number: the wall clock execution time. Typically this is not the execution time for the entire run of the application. For practical reasons, in order to keep the run time of benchmarks acceptably short, the test cases normally do not represent an entire scientifically meaningful calculation as would be expected to be run under production conditions. Instead, the run time is kept short by performing a short calculation which is nevertheless representative of production runs. Where an application is simulating the time evolution of a physical system, for example, the period of time simulated may be considerably reduced from that typically simulated under production conditions.

Applications normally include some start-up and shut-down phases which would, under production conditions, only consume an insignificant fraction of the total execution time. For benchmark test cases, however, these phases may no longer be insignificant. It is necessary to exclude them from the measured wall clock time so that the benchmark measurement is a fair reflection of the performance of the application in a production scenario. Precisely what is included and excluded from the measured time is determined by each application.

A benchmark framework called JuBE [12] has been selected by WP6 to manage the execution of both synthetic and application benchmark codes. Further information is available in [3].

2.2 Benchmark run rules

For the purposes of this assessment exercise, the rules governing the benchmark runs are kept deliberately open. When running the same test case on different systems, it is permissible to use system specific optimisations, including compiler flags, conditional compilation options, run-time parameter settings and code modifications, so long as the test case results are still identical to within scientifically acceptable tolerances on different systems. This is because we wish to assess the ability of the systems to deliver the science results rather than their ability to execute a particular set of lines of source code.

2.3 Approach taken

The PRACE prototype systems are quite varied in terms of system architecture. This poses a particular problem when making comparisons between them using application benchmarks. Traditionally, comparisons like this have been made by running the same test case on the same number of CPUs on different systems. There are two problems with this approach:

1. CPUs on different systems may have widely different performance characteristics. For example, an NEC SX-9 vector CPU has approximately 30 times the peak performance of an IBM BlueGene/P CPU.
2. With the advent of heterogeneous multicore architectures, such as the IBM Cell system at BSC, defining precisely what constitutes a CPU is becoming increasingly difficult.

To address these problems we have taken a different approach: we run the same test case on partitions of each system which have the same nominal peak floating-point performance (this metric is the R_{peak} value as reported in the Top500 list of supercomputers [13]). We have defined a series of partition sizes on which the test cases are run, where each partition size has twice the peak performance of the next smaller one. The partition sizes we use are 2.5, 5, 10, 20, 40, 80 and 160 TFlop/s. Unfortunately, not all the PRACE prototype installations are large enough to contain all these partitions: the smallest systems have peak performances between 10 and 20 TFlop/s, so the largest partition size we can use on these systems is 10 TFlop/s.

For each prototype system, the number of CPUs corresponding to each partition size can straightforwardly be calculated by dividing the required partition performance by the peak floating-point performance of a CPU. The results of this calculation are shown in Table 1.

A value of N/A is shown where the prototype installation is not large enough to deliver the required partition performance. A CPU is taken to be single core for all systems except the IBM Cell system, where a CPU is a whole Cell BE (one PPE core and eight SPE cores). For the IBM Cell system, the PPE is included in the peak flop/s rating, but the auxiliary Power6 processors are not. For the NEC SX-9 system, the scalar processors are not included in the peak flop/s rating. The x86 cluster part of the NEC system is evaluated as a separate prototype. Note that there are two columns for the Cray XT5, as the system contains two different models of AMD Opteron CPUs: Barcelona (denoted by XT5b) and Shanghai (denoted by XT5s). These two parts of the Cray XT5 have been treated as separate systems for the purposes of this study.

System	IBM BlueGene/P	Cray XT5b	Cray XT5s	IBM Power6	Sun x86 cluster	Bull x86 cluster	IBM Cell	NEC SX-9	NEC x86 cluster
Total # of CPUs	294912	5376	1440	3328	17664	1024	144	192	5600
Peak Gflop/s per CPU	3.4	9.2	10.8	18.8	11.72	11.72	108.8	102.4	11.2
# of CPUs = 2.5TF	735	272	231	133	213	213	23	24	223
# of CPUs = 5TF	1471	543	463	266	427	427	46	49	446
# of CPUs = 10TF	2941	1087	926	532	853	853	92	98	893
# of CPUs = 20TF	5882	2174	N/A	1064	1706	N/A	N/A	195	1786
# of CPUs = 40TF	11765	4348	N/A	2128	3413	N/A	N/A	N/A	3571
# of CPUs = 80TF	23529	N/A	N/A	N/A	6826	N/A	N/A	N/A	N/A
# of CPUs = 160TF	47059	N/A	N/A	N/A	13652	N/A	N/A	N/A	N/A

Table 1: Partition sizes for prototypes systems.

For some applications it is not possible to run a single test case across all partition sizes, for example because of memory constraints. Some applications therefore have multiple test cases, but for each test case the same physical problem is being simulated on different architectures and different partitions sizes. Thus (with a small number of exceptions) each individual test case is being run in a strong scaling mode, though different test cases may be used to span the full range of partition sizes.

Of the prototype systems, the IBM Power6 and the x86 clusters contain CPUs with simultaneous multithreading (SMT) support. For some applications, performance can be improved by running two MPI tasks (or two OpenMP threads) per core. In the results reported in this document, this feature has been utilised only for the IBM Power6 system. The execution time reported is the always the lowest for the given number of CPUs, regardless or whether SMT was used or not.

One problem we need to overcome in this approach is that applications may not run on the exact number of CPUs required for a given partition size, for example because the number of MPI tasks must be divisible by certain values. Furthermore, using a number of CPUs which does not correspond to a whole number of nodes may result in anomalous performance behaviour. The solution to this problem is to run on two different numbers of CPUs, one smaller than the target number and one larger. These results are then interpolated onto the target number of CPUs for the partition size. We use the following terminology:

$\mathbf{P_t}$ is the target number of CPUs for the partition size.

$\mathbf{P_l}$ is the number of CPUs for the smaller run for this partition size.

$\mathbf{P_u}$ is the number of CPUs for the larger run for this partition size.

$\mathbf{T_l}$ is the execution time for the smaller run for this partition size

$\mathbf{T_u}$ is the execution time for the larger run for this partition size.

$\mathbf{T_t}$ is the interpolated execution time for the target number of CPUs.

Since the naive expectation is that the execution time will be approximately inversely proportional to the number of CPUs, we interpolate linearly in the reciprocal of execution time, rather than in execution time itself. Thus T_t is given by:

$$T_t^{-1} = T_l^{-1} + (T_u^{-1} - T_l^{-1}) \times (P_t - P_l) \; / \; (P_u - P_l)$$

2.4 Reporting of results

To show comparisons between systems, we plot the data in two different ways. Firstly, we simply plot execution time against number of CPUs (on a log-log scale). Secondly we compute an efficiency metric by dividing performance (T^{-1}) by the partition size in peak TFlop/s. Since the units here are somewhat meaningless, we normalise this against the performance per peak TFlop/s on a given partition of a reference machine: for this we have chosen a 10 TFlop/s partition of the IBM Power6. Let $T^{X,A}$ be the execution time on an X TFlop/s partition of machine A. This dimensionless metric which we term *relative efficiency*, $E^{X,A}$, is then given by the expression:

$$E^{X,A} = (10\ T^{10,Power6})/(X\ T^{X,A})$$

This metric is then plotted against partition size. These plots have the following properties:

- They allow comparison of performance of an application across different machines relative to the peak performance of the machine, rather than relative to the number of CPUs.
- Scalability can be easily assessed, since perfect scaling is represented by a constant value of this metric.
- Small differences in relative performance can be easily seen, since no log scale is required on the y-axis.
- Due to the normalisation, it is possible to compare the value of this metric for different applications on the same machine.

$E^{X,A}$ is related to the percentage of peak performance attained by the application as follows: if, for example, the value of $E^{X,A}$ is twice the value of $E^{X,B}$, and we assume that the benchmark executes the same number flops on the two machines, this means that the percentage of peak flop rate for this application on machine A is twice that on machine B for a partition size X. In general, therefore, a higher relative efficiency indicates that an application is well suited to the architecture in question, since it achieves a higher percentage of the peak floating point performance, and therefore makes better use of the capabilities of the system.

In a few cases, benchmarks can only be run in a weak scaling mode, where the amount of computation is proportional to the number of CPUs. In these cases, we relative efficiency is calculated as:

$$E^{X,A} = T^{10,Power6}/T^{X,A}$$

Note that not all of the applications have been run on all of the prototype systems: to do so would have required a level of effort beyond the scope of this project.. In some cases, considerable effort would be required to port codes to, and optimise them for, certain of the prototypes, especially the IBM Cell and NEC SX-9 systems. Some prototype systems (the Intel Nehalem-based clusters) were not available until quite late in the project.

3 Application Benchmarking Results

This Chapter contains the results of the assessments for each application benchmark. For each test case we present two figures, The first shows the execution time as a function of the number of CPUs. The second shows the metric $E^{X,A}$ as a function of the partition size in TFlop/s. In this second figure, higher values represent better performance, and flatter curves show better scaling. The corresponding tables of raw data are presented in Chapter 6.

3.1 CODE_SATURNE

3.1.1 Summary

Code_Saturne is a general purpose CFD code, used for nuclear thermalhydraulics processes, coal and gas combustion, aeraulics, etc.

3.1.2 Test cases

Test Case A consists of an isothermal Large Eddy Simulation in a T-junction, containing 10 million cells. Only the dynamics of the flow is investigated. The runs are carried out for 100 time steps, starting from an already developed flow.

Test Case B consists of a flow around the body of a submarine, consisting of 107 million cells.

3.1.3 Results

Figure 1 and Figure 2 show the results for Test Case A. Figure 3 and Figure 4 show the results for Test Case B.

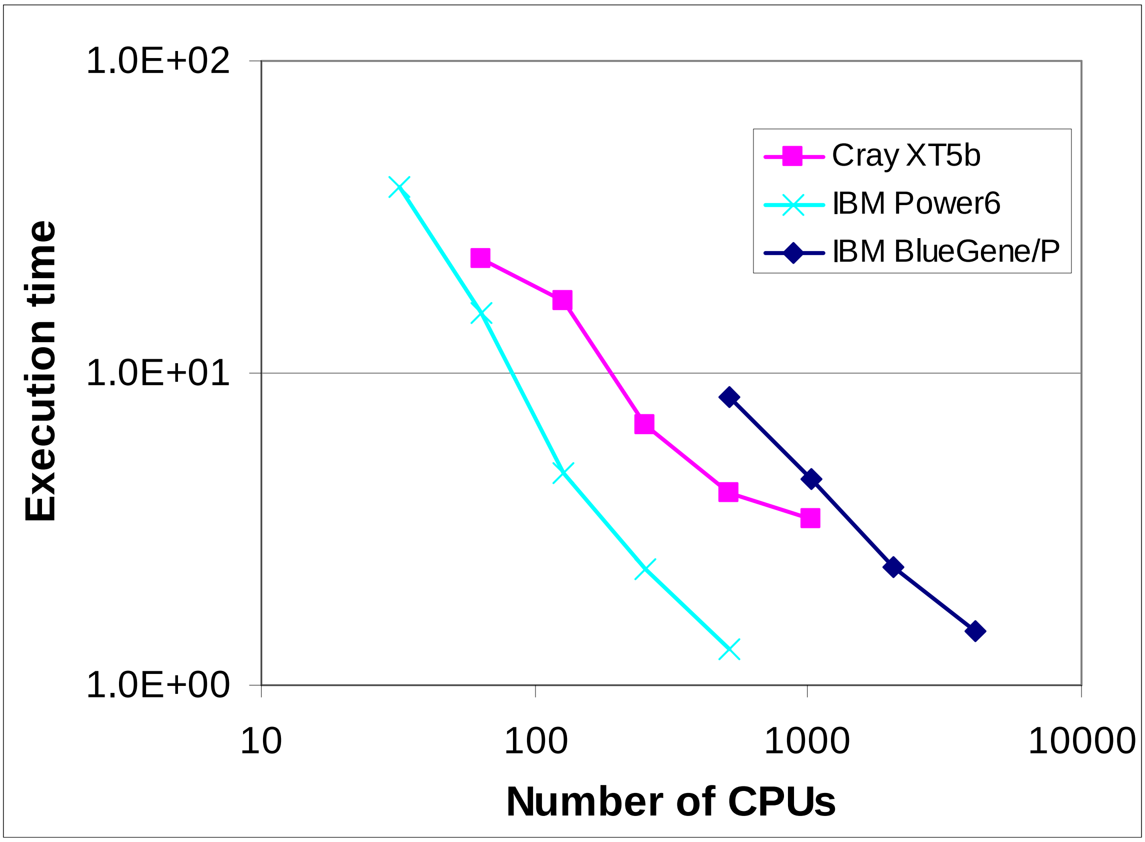

Figure 1: Execution time for Code_Saturne, Test Case A

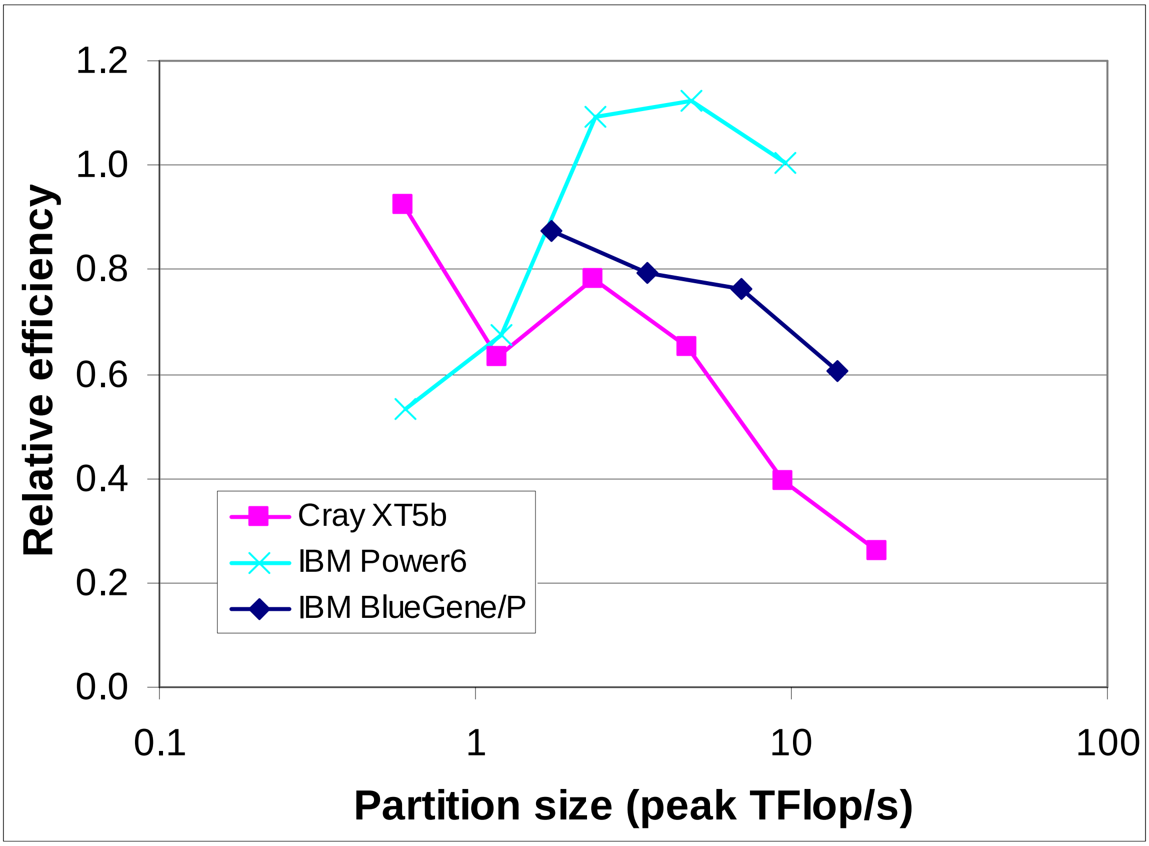

Figure 2: Relative efficiency of Code_Saturne, Test Case A

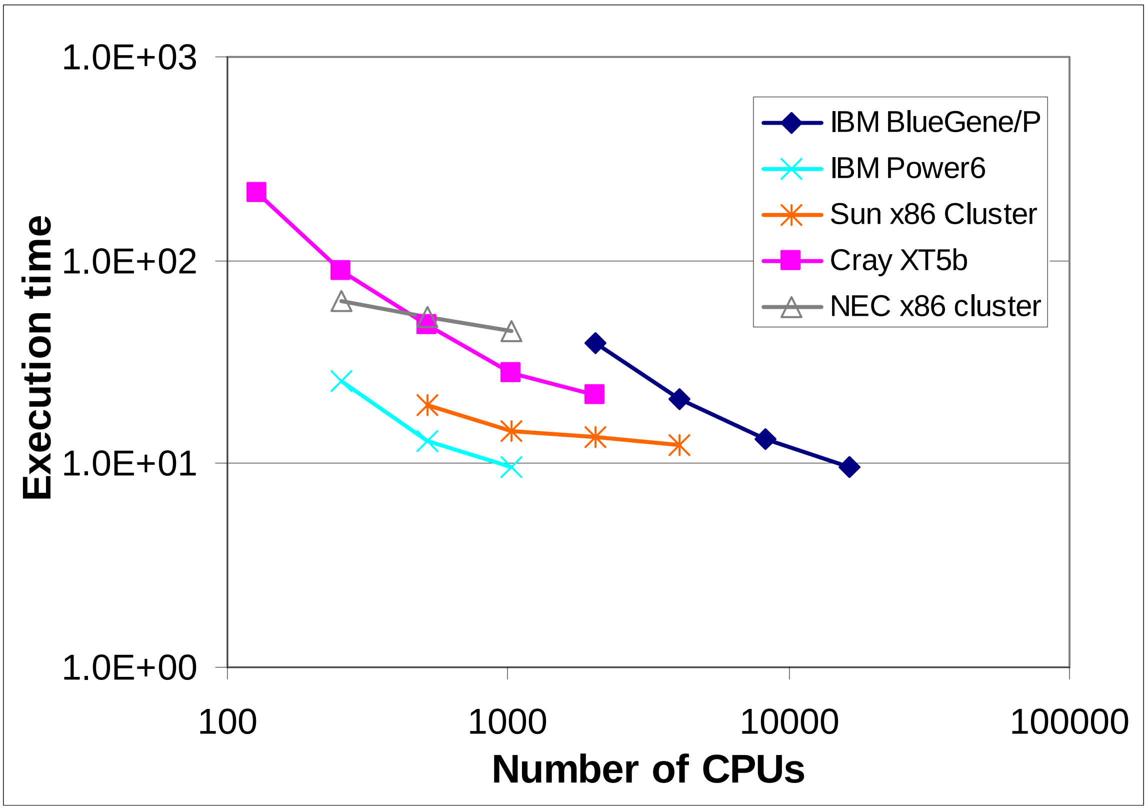

Figure 3: Execution time for Code_Saturne, Test Case B

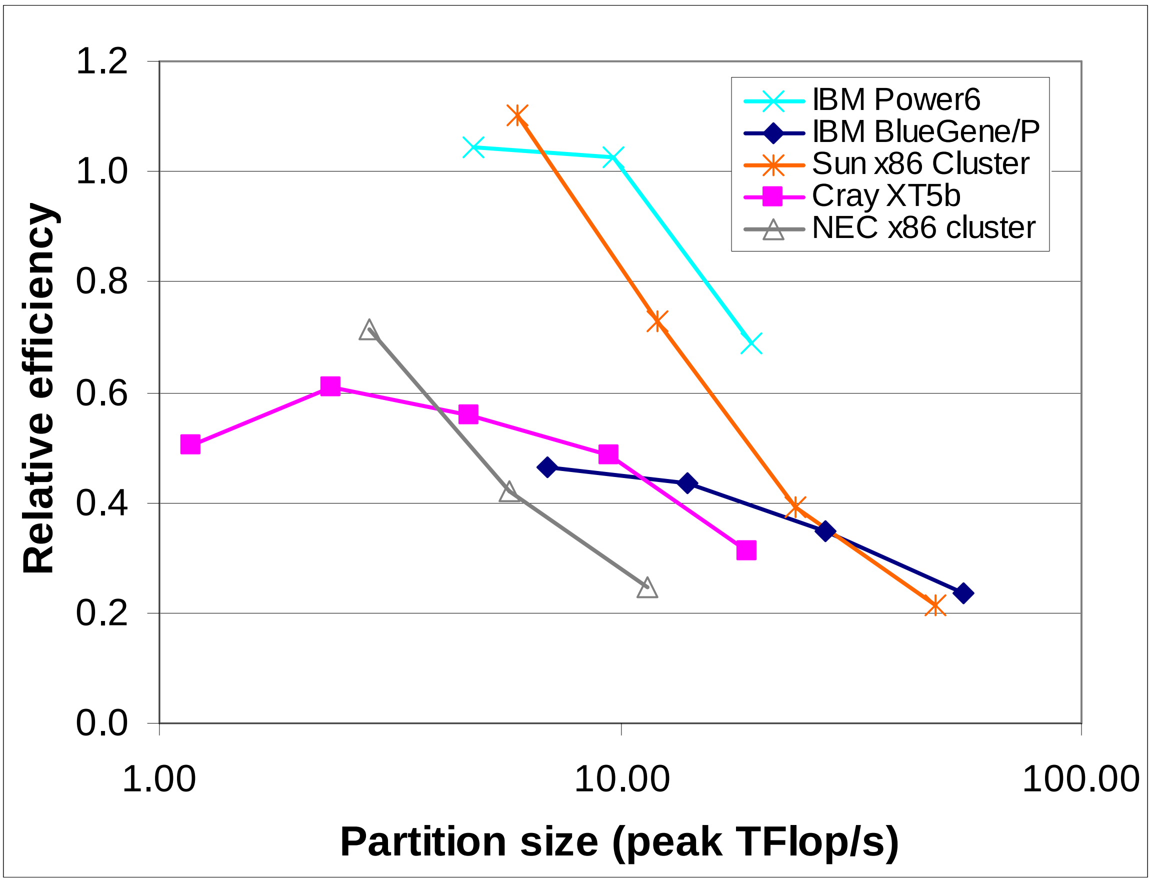

Figure 4: Relative efficiency of Code_Saturne, Test Case B

3.1.4 Analysis

Of the prototypes tested, the IBM Power6 shows the best relative efficiency for larger partition sizes, and also displays the best scaling. The Cray XT5 shows poorer relative efficiency and also poorer scaling behaviour. The IBM BlueGene/P shows reasonable scaling, despite the much higher numbers of CPUs required. On Test Case B, neither of the x86 clusters shows good scaling behaviour, though the Sun x86 cluster shows a much higher relative efficiency than the NEC x86 cluster: the reason for this is not clearly understood.

3.2 CP2K

3.2.1 Summary

CP2K is a community code to perform atomistic and molecular simulations of solid state, liquid, molecular and biological systems. It consists of several components for classical molecular dynamics, ab-initio density functional theory. etc.

3.2.2 Test cases

Test Case A consists of a simulation of 1024 water molecules.

3.2.3 Results

Figure 5 and Figure 6 show the results for Test Case A.

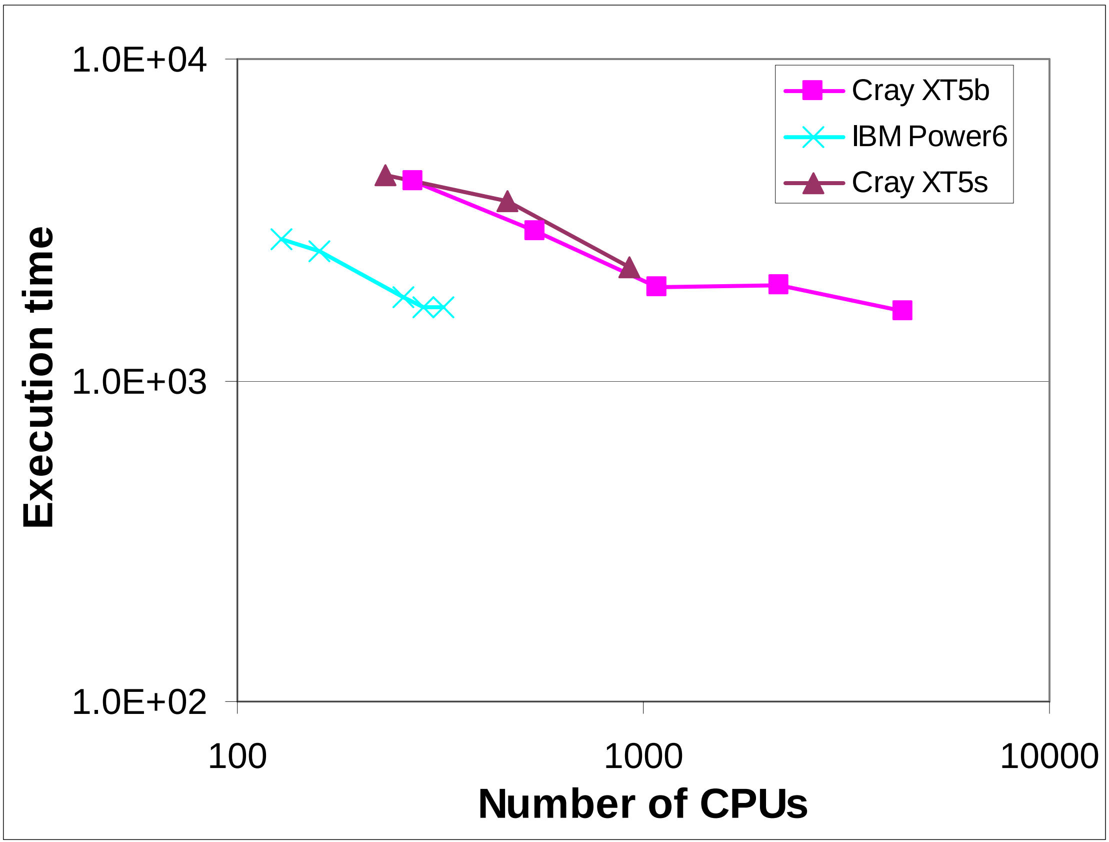

Figure 5: Execution time for CP2K, Test Case A

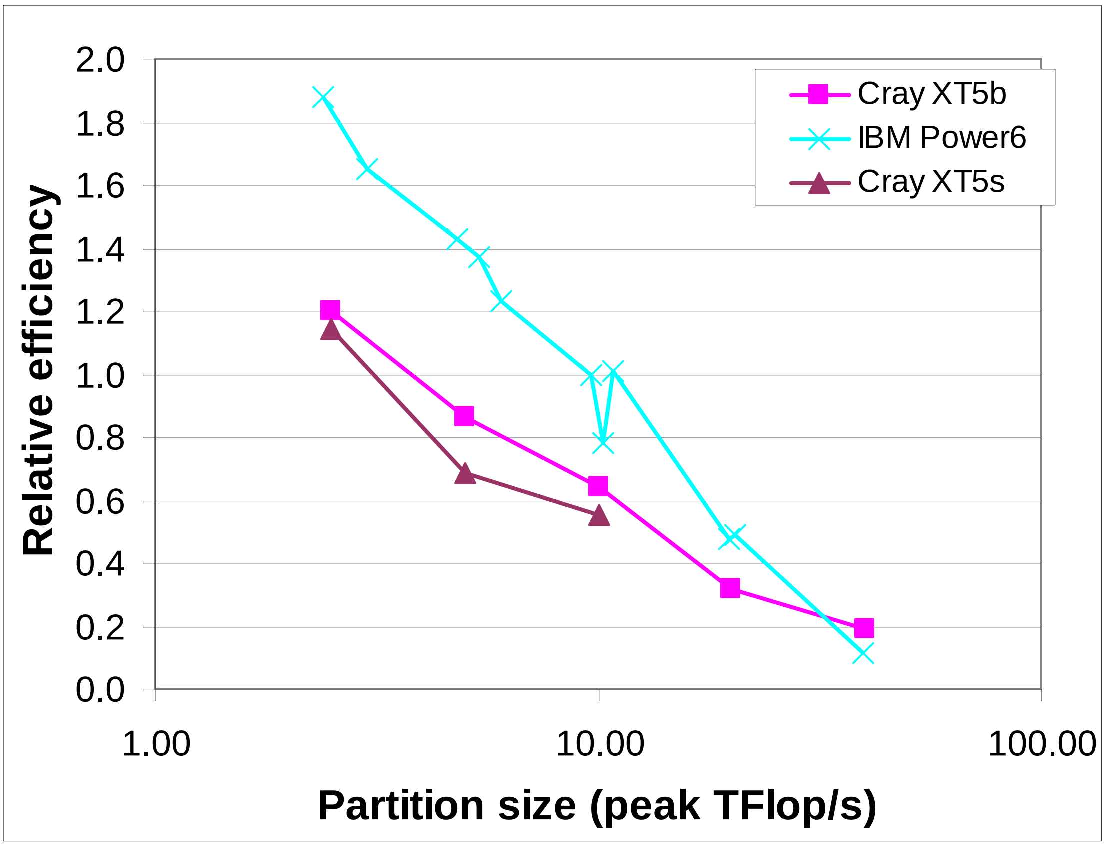

Figure 6: Relative efficiency of CP2K, Test Case A

3.2.4 Analysis

This test case does not scale well on either the IBM Power6 or the Cray XT5. For smaller problem sizes, the IBM Power6 shows higher relative efficiency, but the two systems have similar relative efficiency at 40 TFlop/s peak. The Cray XT5s shows similar scaling behaviour as the Cray XT5b, but slightly lower relative efficiency.

3.3 CPMD

3.3.1 Summary

CPMD is a parallelized plane wave/pseudopotential implementation of Density Functional Theory, particularly designed for ab-initio molecular dynamics.

3.3.2 Test cases

Test Case A is a simulation of 64 Ionic Liquids.

3.3.3 Results

Figure 7 and Figure 8 show the results for Test Case A

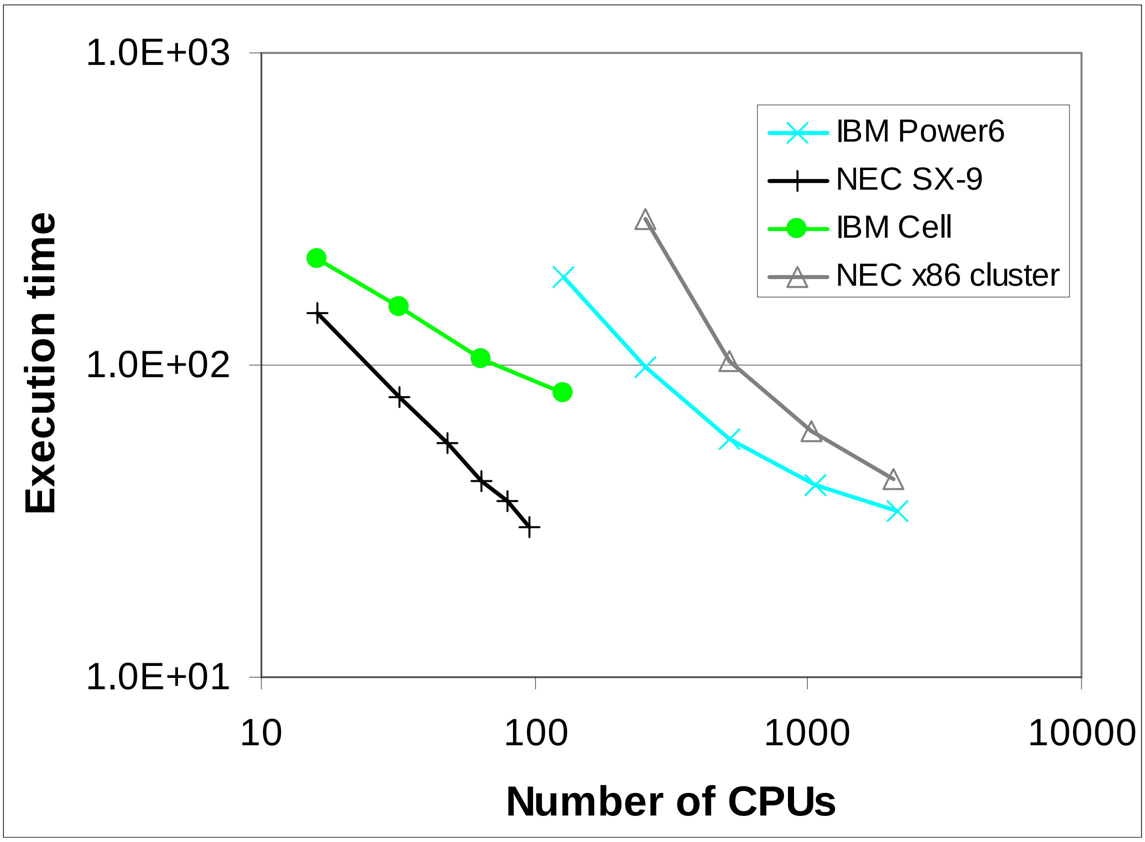

Figure 7: Execution time for CPMD, Test Case A

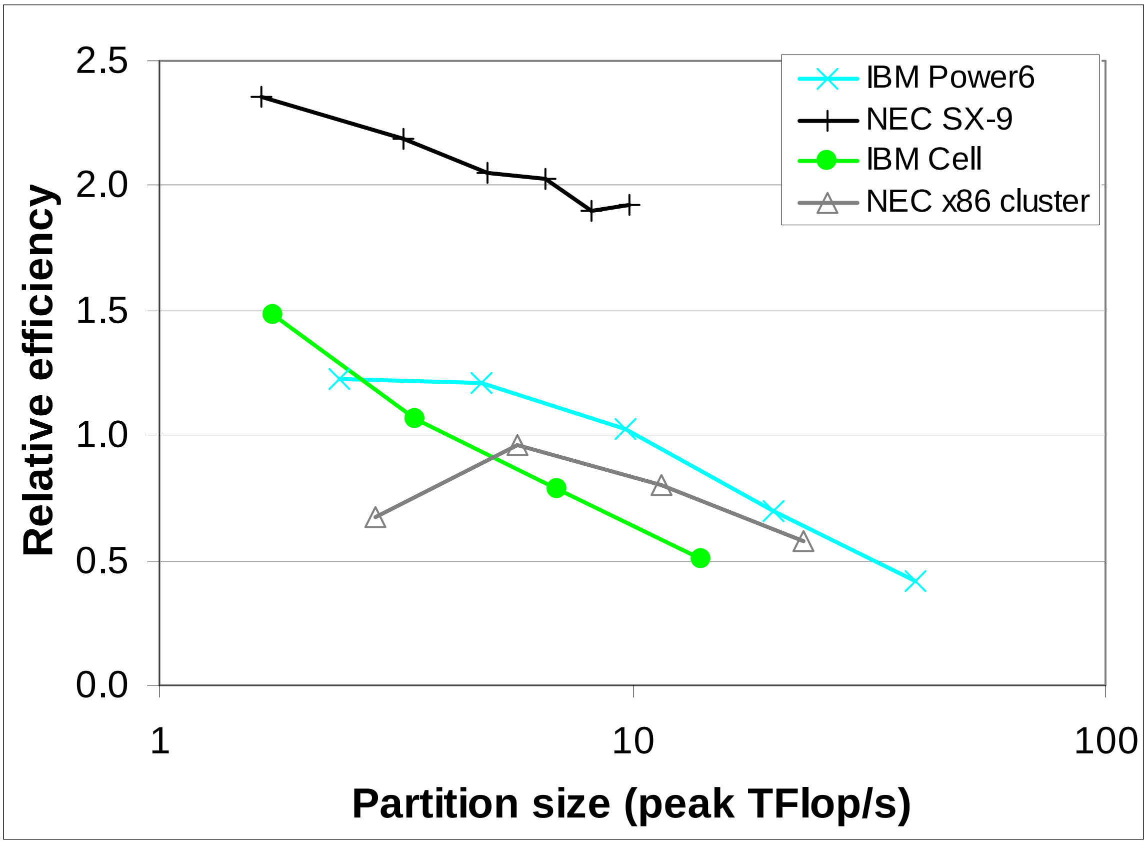

Figure 8: Relative efficiency of CPMD, Test Case A

3.3.4 Analysis

This test case scales well on the NEC SX-9, but not on the IBM Power6, NEC x86 cluster, or IBM Cell. At 10 TFlop/s peak , the NEC SX-9 is about twice as efficient as the IBM Power6. The NEC x86 cluster is a little less efficient than the IBM Power6. The IBM Cell shows poor relative efficiency as the partition size is increased: this can be attributed to an incomplete port of the application to the Cell SPEs.

3.4 EUTERPE

3.4.1 Summary

EUTERPE simulates a plasma in a cylindrical θ-pinch configuration. The code solves the coupled system of gyrokinetic equations for the ions, in the electrostatic approximation, and the quasi-neutrality equation, assuming adiabatically responding electrons.

3.4.2 Test cases

Test Case A consists of a simulation of 1000 million particles.

3.4.3 Results

Figure 9 and Figure 10 show the results for Test Case A.

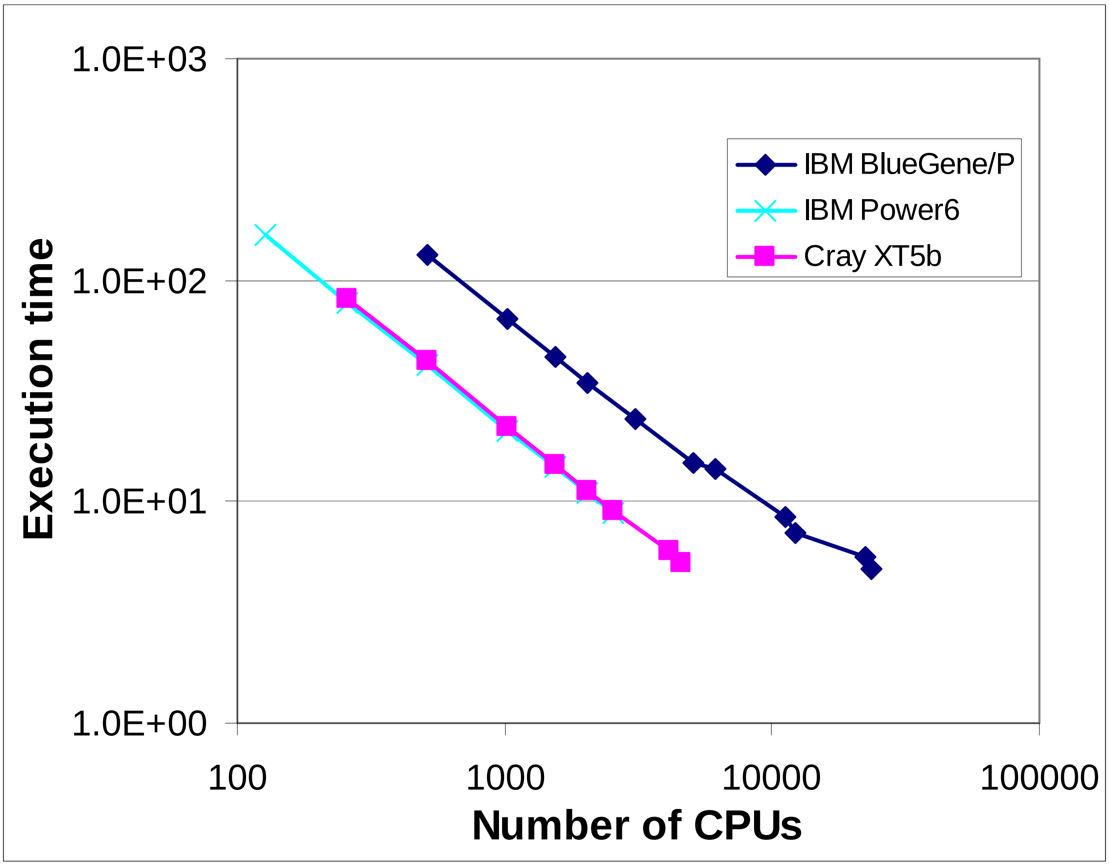

Figure 9: Execution time for EUTERPE, Test Case A

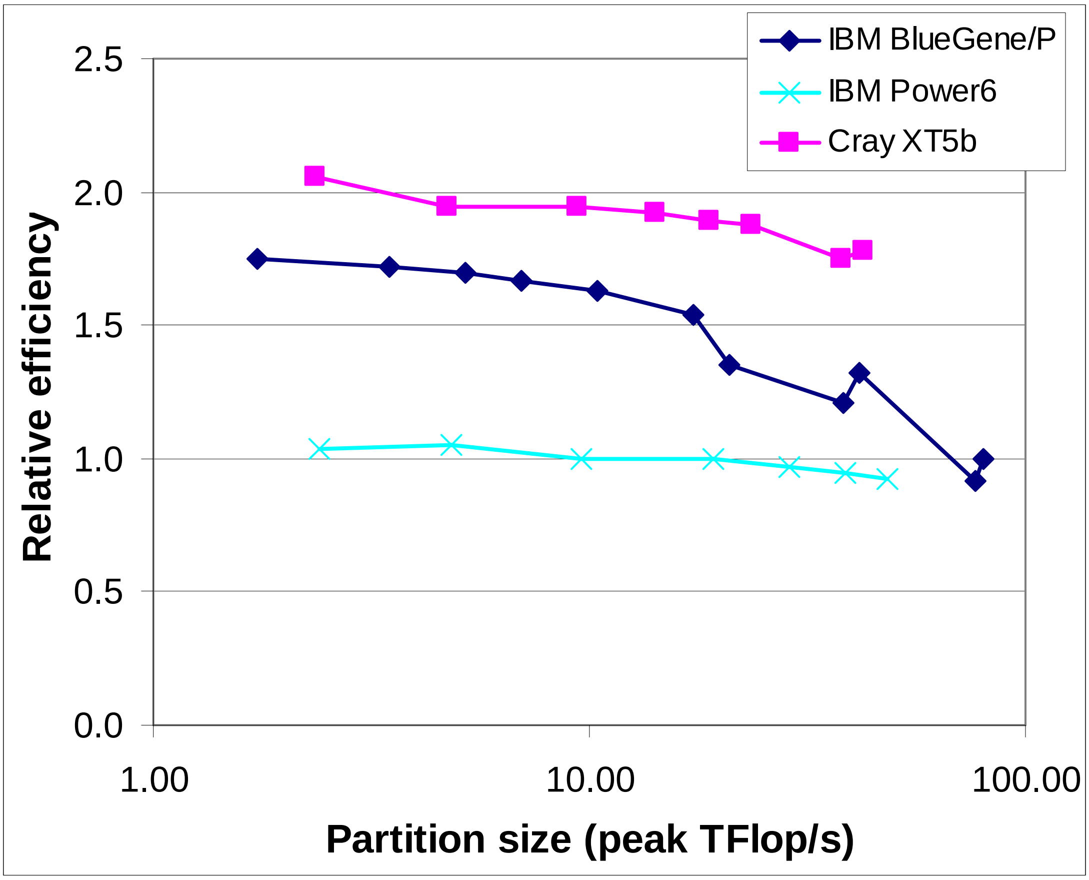

Figure 10: Relative efficiency of EUTERPE, Test Case A

3.4.4 Analysis

This test case scales well on all three systems up to 40-50 TFlop/s. For smaller partition sizes, the IBM BlueGene/P is significantly more efficient than the IBM Power6, but this advantage diminishes as the partition size increases. The Cray XT5b shows a relative efficiency approximately twice that of the IBM Power6.

3.5 GADGET

3.5.1 Summary

GADGET is a freely available code for cosmological N-body/SPH simulations on massively parallel computers with distributed memory.

3.5.2 Test cases

Test Case A is a simulation consisting of 32 million particles.

Test Case B is a simulation consisting of 64 million particles.

3.5.3 Results

Figure 11 and Figure 12 show the results for Test Case A; Figure 13 and Figure 14 show the results for Test Case B.

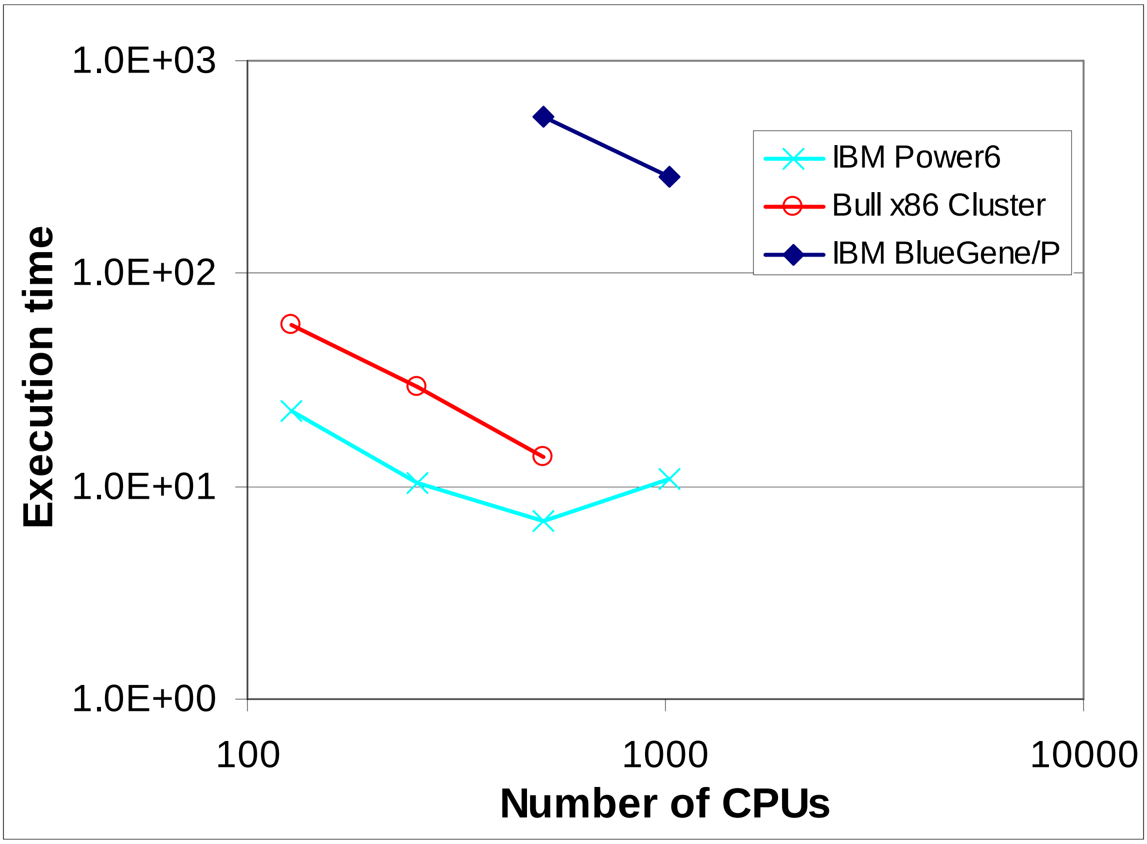

Figure 11: Execution time for GADGET, Test Case A

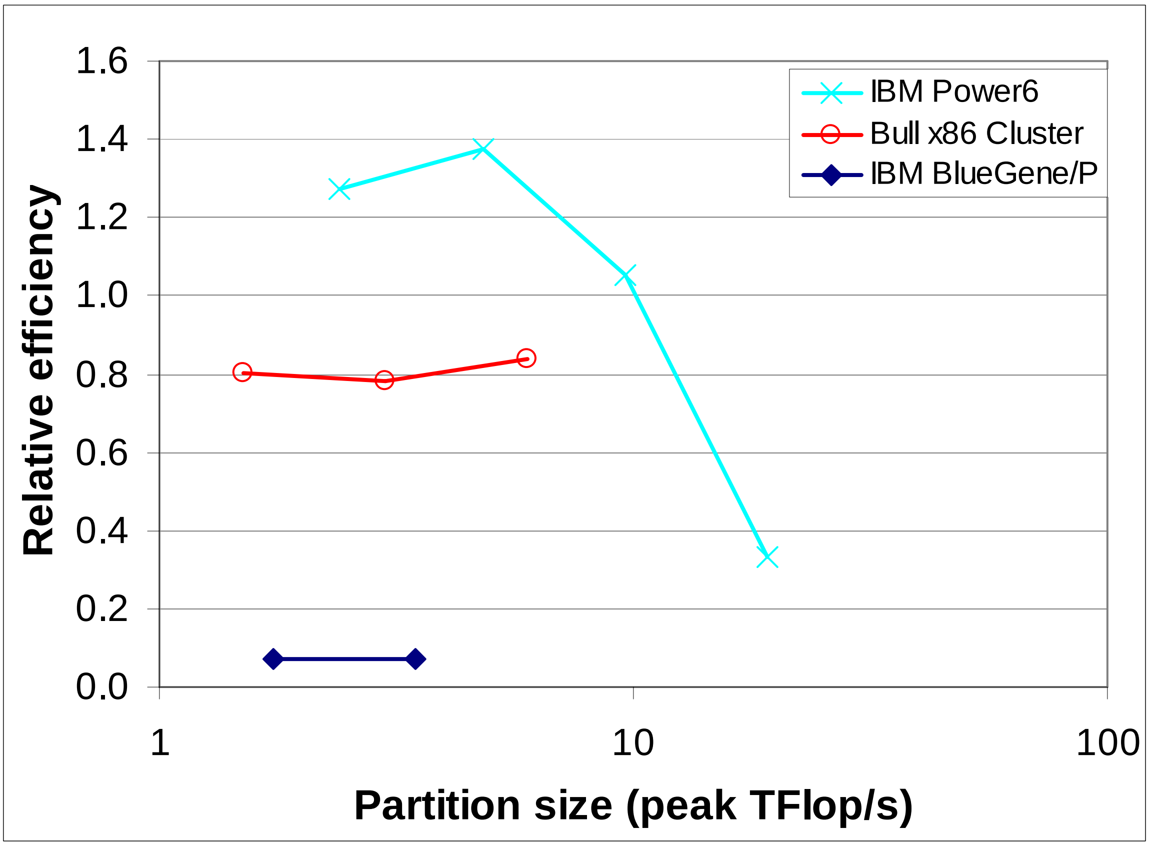

Figure 12: Relative efficiency of GADGET, Test Case A

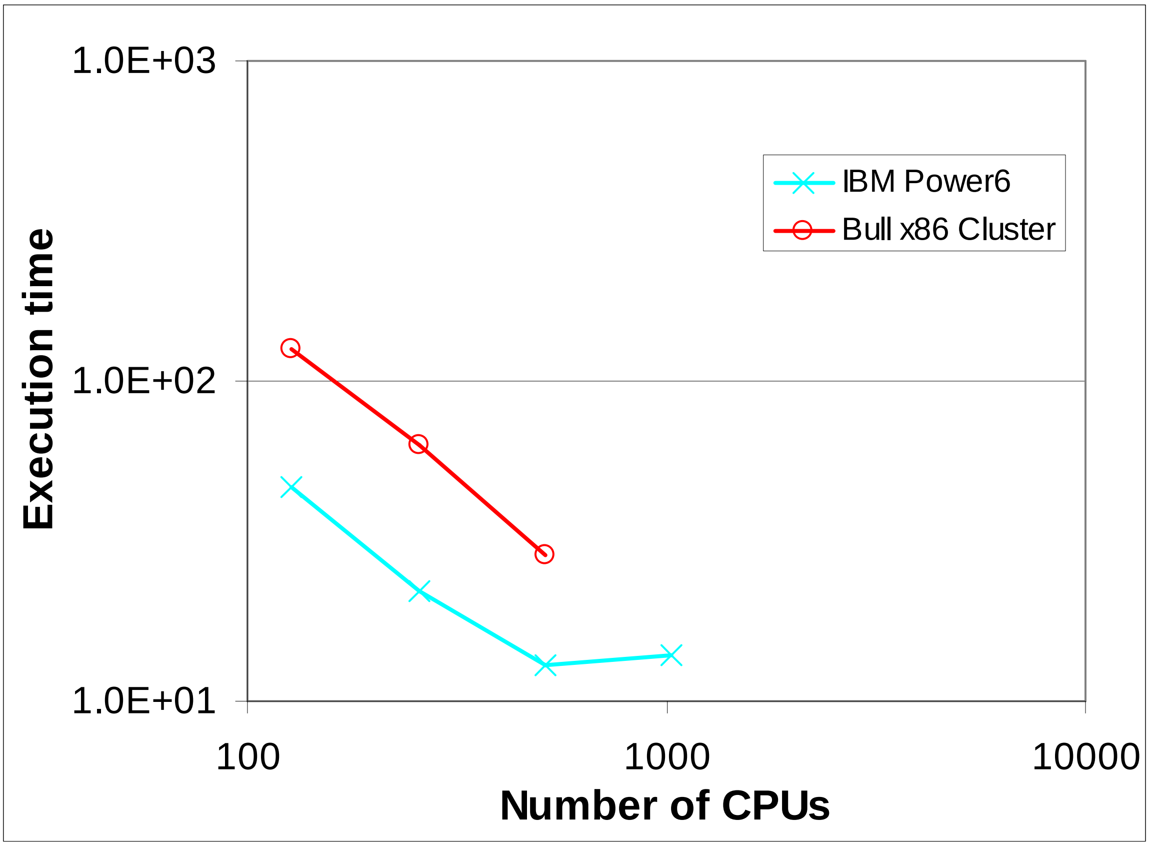

Figure 13: Execution time for GADGET, Test Case B

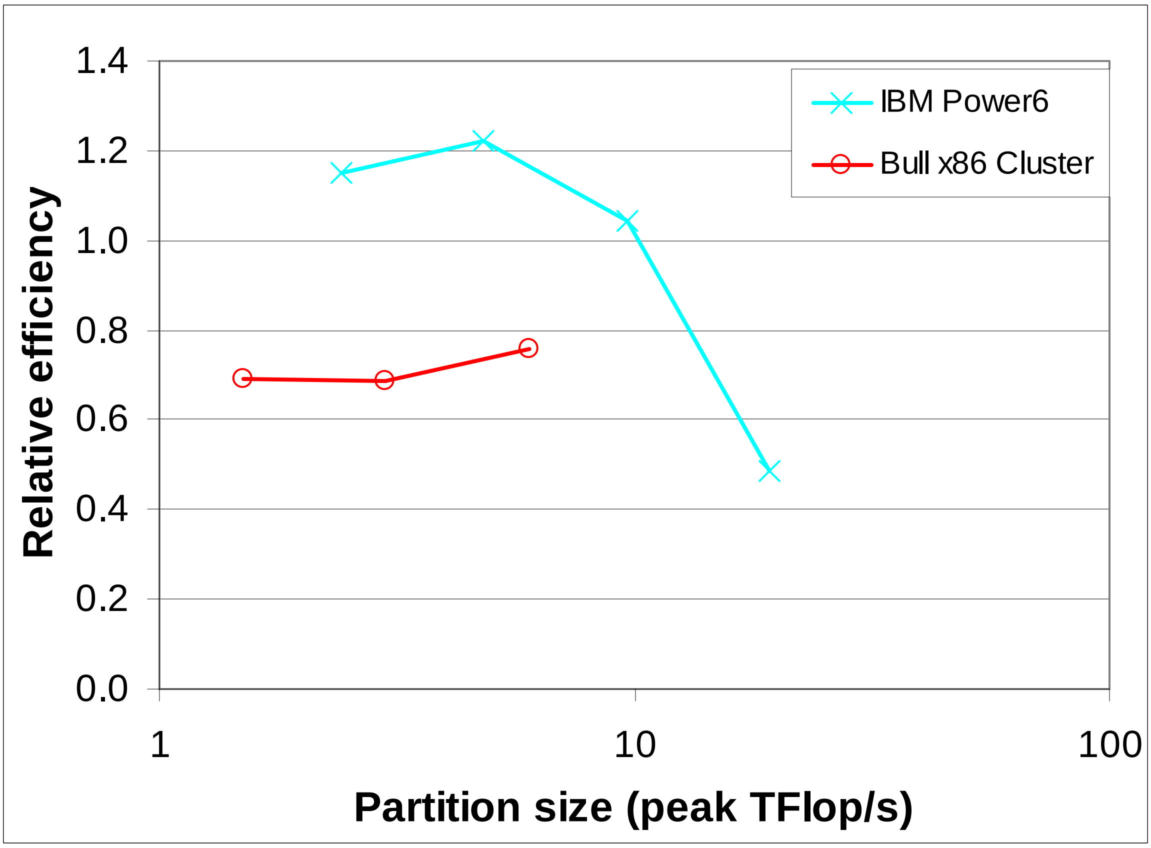

Figure 14: Relative efficiency of GADGET, Test Case B

3.5.4 Analysis

Both test cases scale well up to 5 TFlop/s peak on both the Bull x86 cluster and on the IBM Power6, but at 10 TFlop/s on the IBM Power6 all scalability has been exhausted. Efficiency is lower on the Bull x86 cluster, than on the IBM Power6, and the efficiency is very poor on the IBM BlueGene/P: due to memory restrictions this code is run using only one core per node.

3.6 GROMACS

3.6.1 Summary

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 such as proteins and lipids that have many 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.

3.6.2 Test cases

Test Case A comprises two vesicles embedded in water with 1752 POPC lipids and 334489 water molecules giving in total 1094681 atoms. The long ranged forces are calculated using the reaction-field (RF) approximation. Dynamic load balancing is turned on and thus the domain decomposition varies throughout the simulation to balance the load on each node.

3.6.3 Results

Figure 15 and Figure 16 show the results for Test Case A.

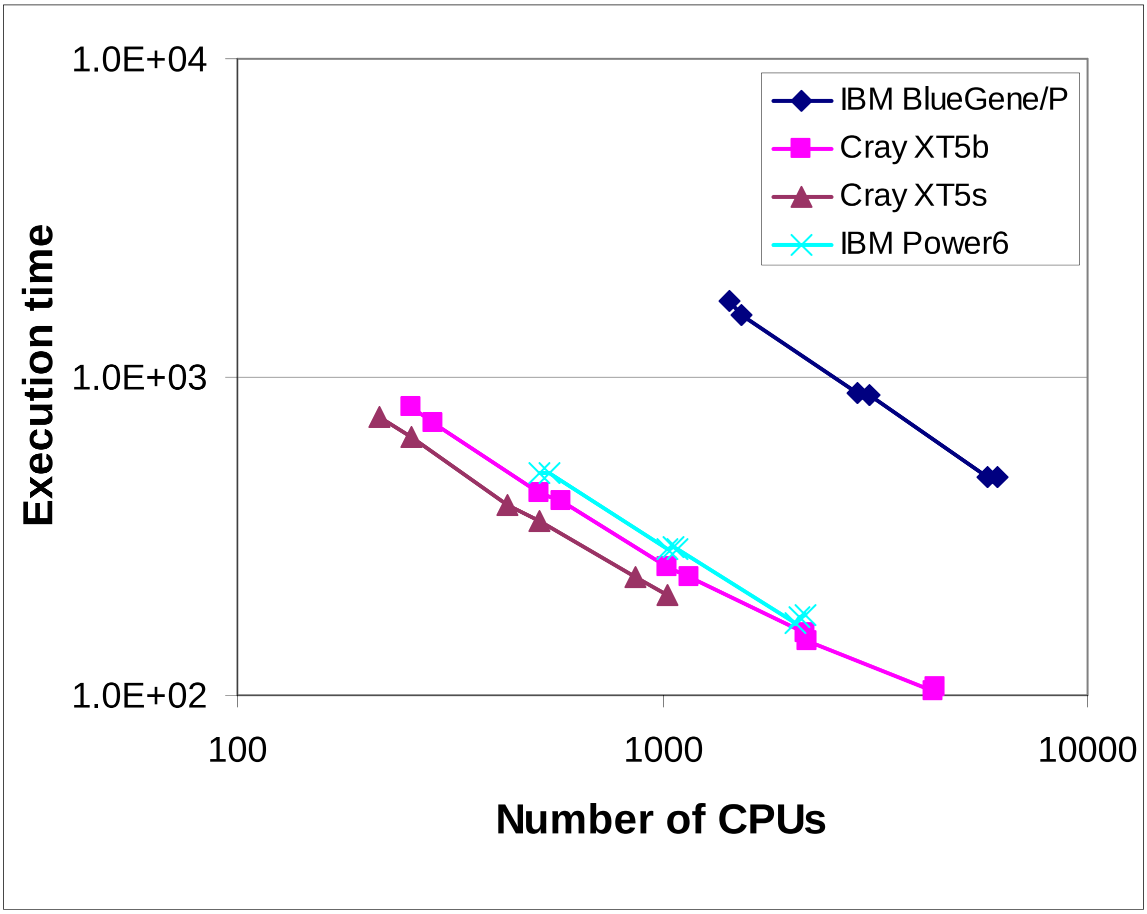

Figure 15: Execution time for GROMACS, Test Case A

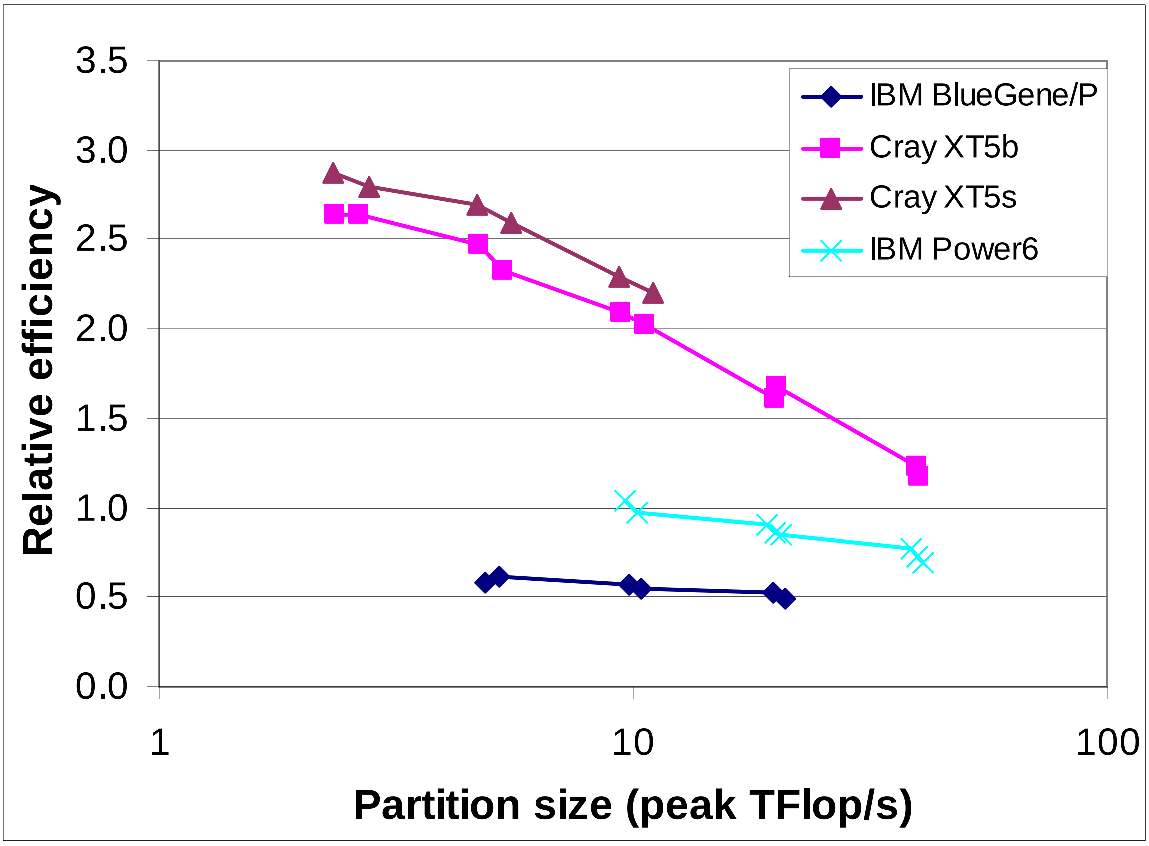

Figure 16: Relative efficiency of GROMACS, Test Case A

3.6.4 Analysis

The Cray XT5 shows a high relative efficiency on this application compared to the IBM Powe6, though this advantage is diminished as the partition size reaches 40 TFlop/s, while the IBM BlueGene/P is about half as efficient as the IBM Power6. On the Cray XT5, the Shanghai cores are slightly more efficient than the Barcelona cores.

3.7 NAMD

3.7.1 Summary

NAMD is a parallel molecular dynamics code for high-performance simulation of large biomolecular systems. For the assessments reported here, a beta version of NAMD 2.7 was used, which significantly reduces memory requirements compared to previous versions.

3.7.2 Test cases

Test Case A is a simulation consisting of 1 million atoms.

Test Case B is a simulation consisting of 9 million atoms.

3.7.3 Results

Figure 17 and Figure 18 show the results for Test Case A: Figure 19 and Figure 20 show the results for Test Case B. Note that SMT was utilised on the IBM Power6.

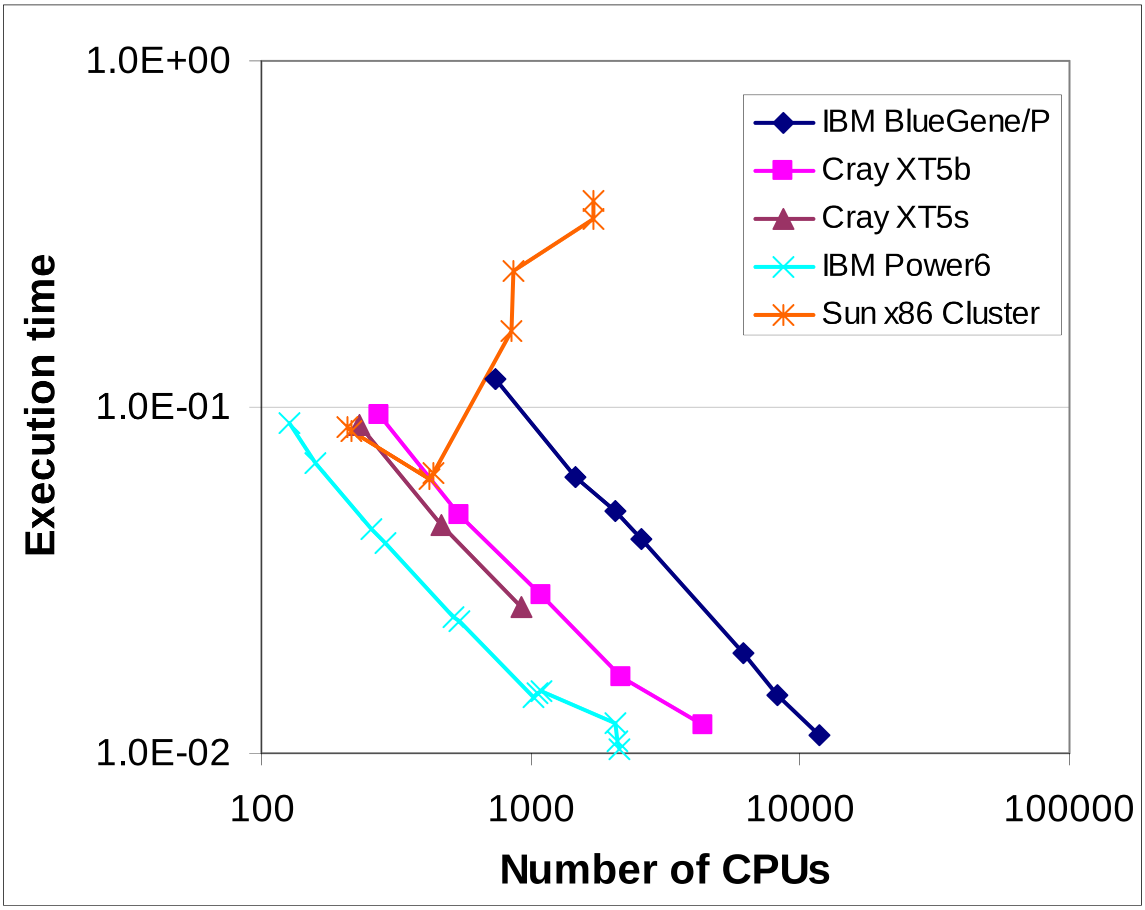

Figure 17: Execution time for NAMD, Test Case A

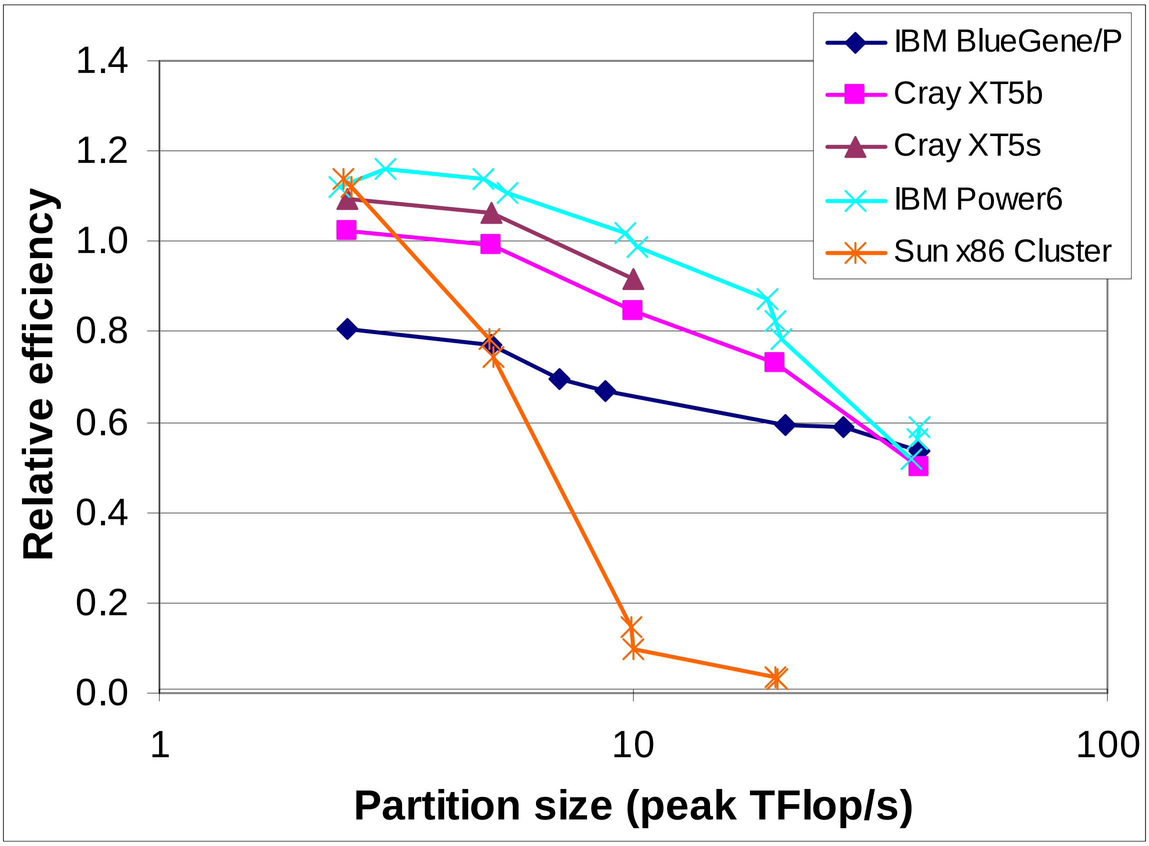

Figure 18: Relative efficiency of NAMD, Test Case A

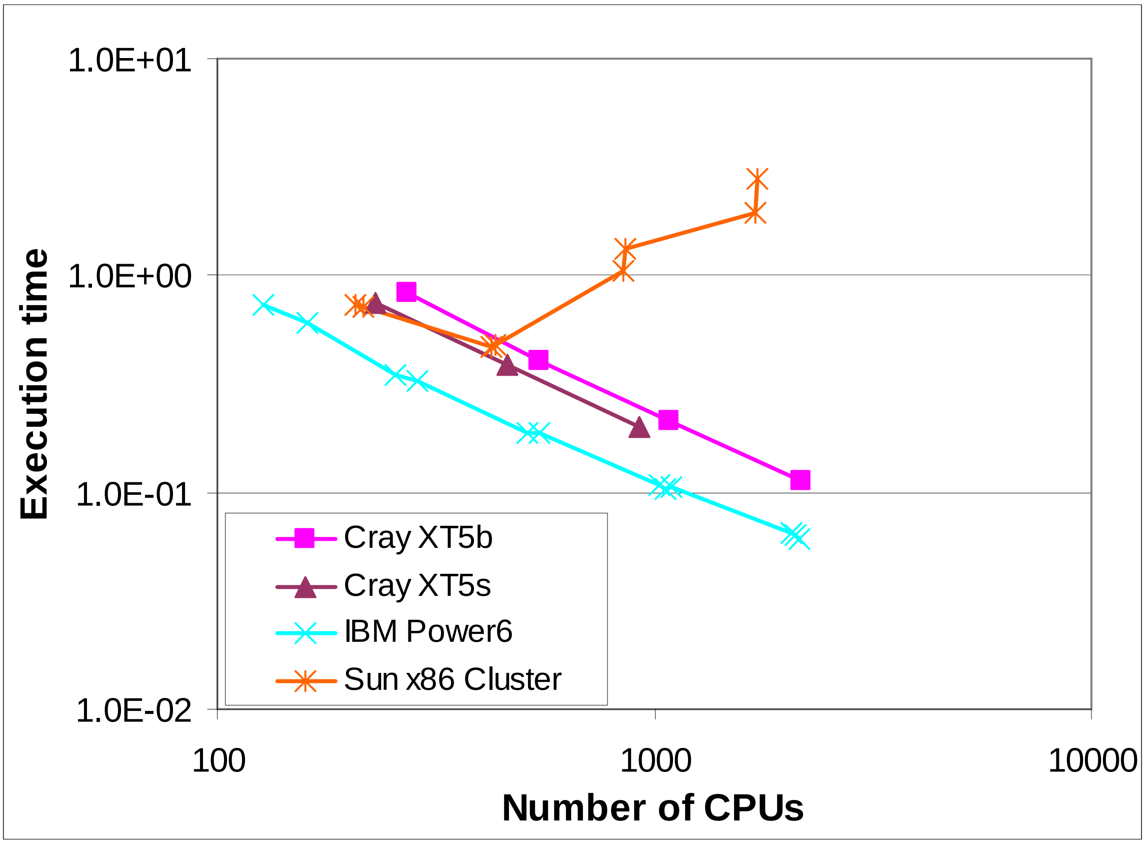

Figure 19: Execution time for NAMD, Test Case B

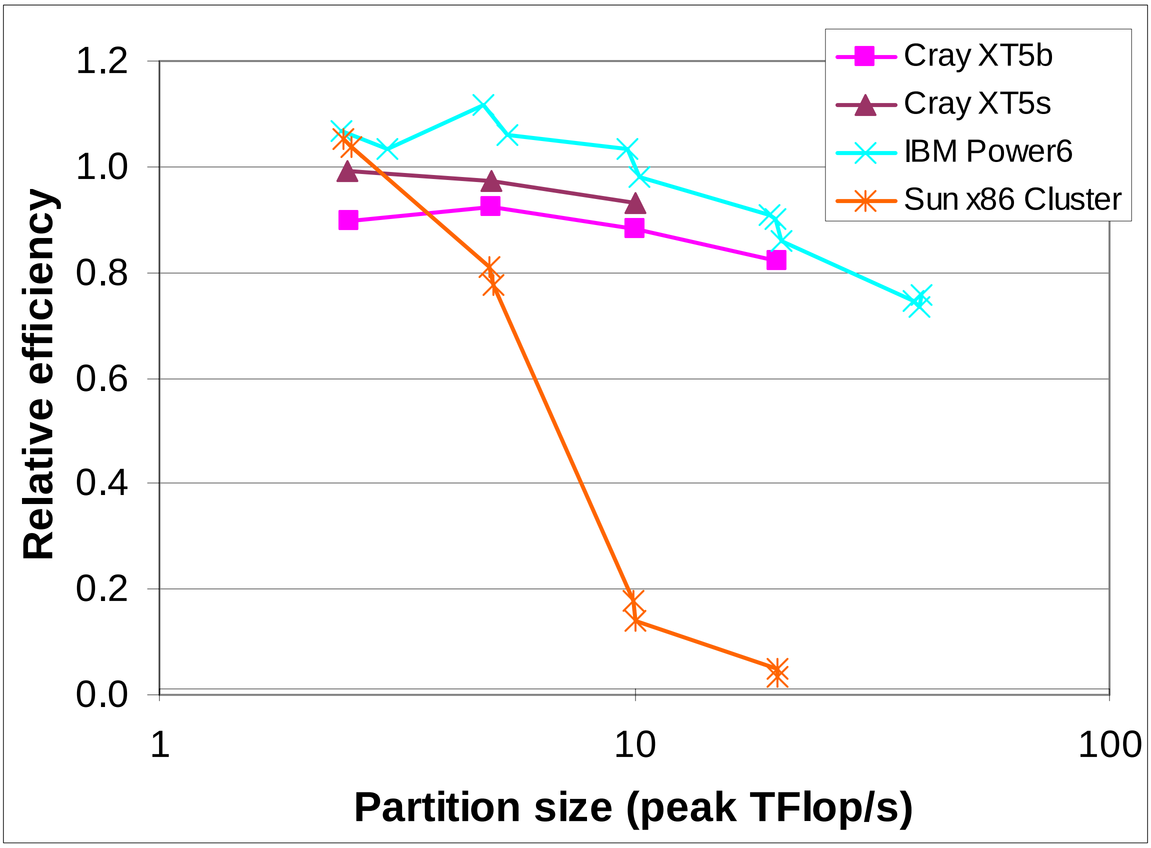

Figure 20: Relative efficiency of NAMD, Test Case B

3.7.4 Analysis

Differences in relative efficiency between the Cray XT5, IBM Power6 and IBM BlueGene/P are small, especially at large partition sizes. Test Case B could not be run on the IBM BlueGene/P, however, due to memory restrictions. The Sun x86 cluster performs very poorly on both test cases: this is under investigation but is possibly related to the Charm++ library deployed on the system implementing all communications using non-blocking point-to-point communications and using MPI_Iprobe to test whether data has arrived prior to posting any receives.

3.8 NEMO

3.8.1 Summary

NEMO is a numerical platform for simulating ocean dynamics and biochemistry, and sea-ice.

3.8.2 Test cases

Test Case A is a global, 0.25 degree resolution simulation with sea-ice. It uses a grid-size of 1442x1021, and 46 depth levels.

3.8.3 Results

Figure 21 and Figure 22 show the results for Test Case A.

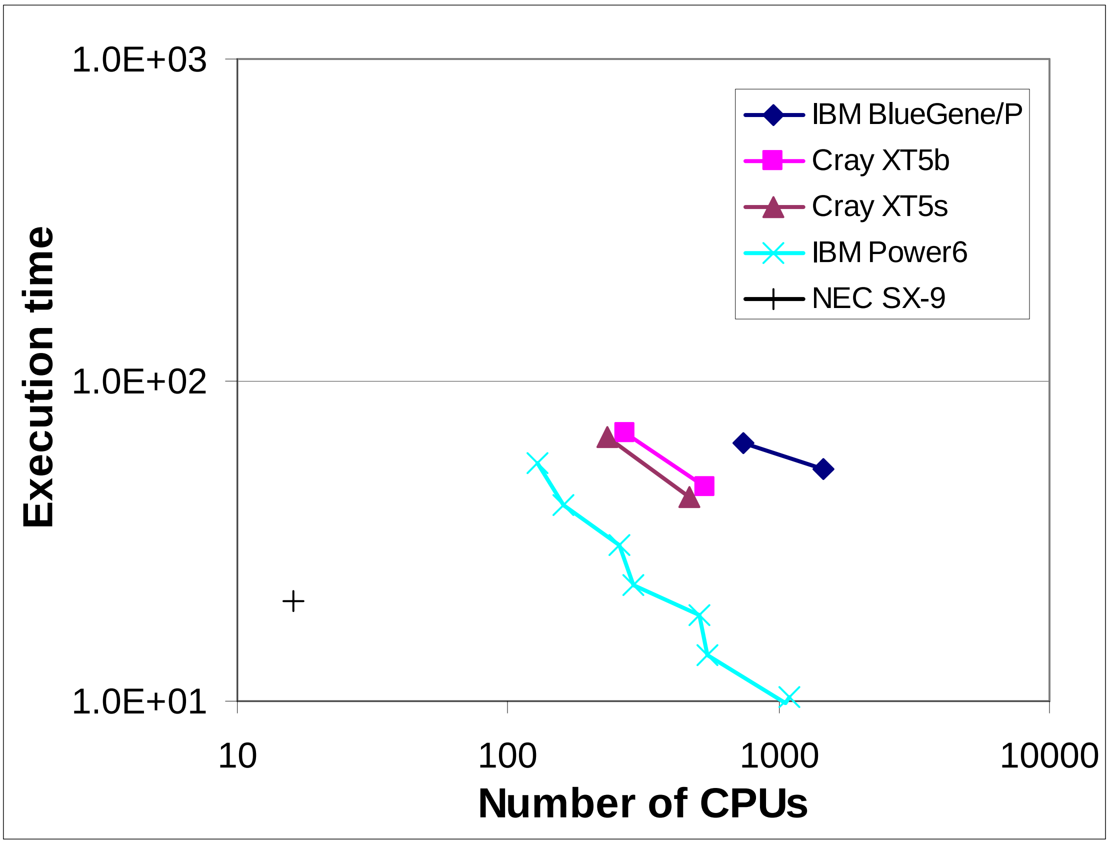

Figure 21: Execution time for NEMO, Test Case A

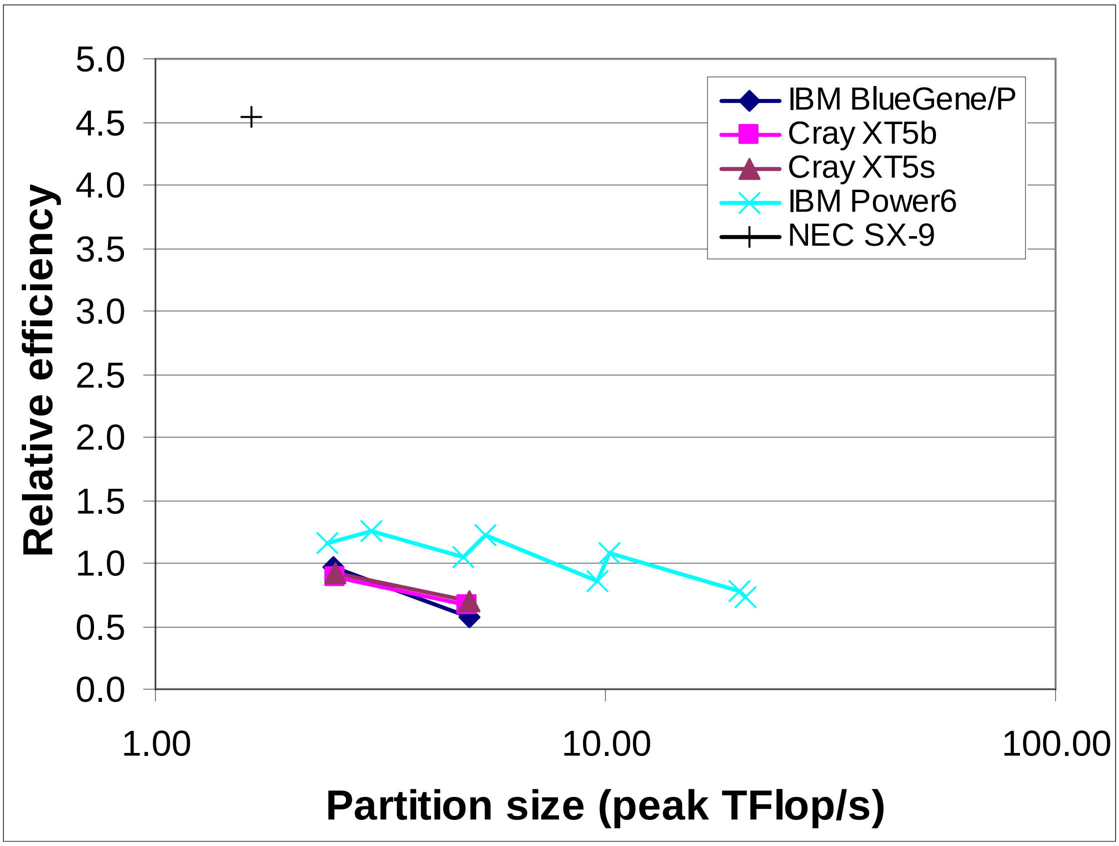

Figure 22: Relative efficiency of NEMO, Test Case A

3.8.4 Analysis

This test case does not scale well on any of the systems tested. Relative efficiency is similar on the Cray XT5 and IBM BlueGene/P, and is lower on these systems than on the IBM Power6. The NEC SX-9 shows very high relative efficiency on the small partition size tested.

3.9 NS3D

3.9.1 Summary

NS3D solves the incompressible Navier-Stokes equations by Direct Numerical Simulation (DNS).

3.9.2 Test cases

Test Case A is a shearwake with 1024 x 840 x 128 grid points.

Test Case B is a shearwake with 2048 x 840 x 128 grid points.

Test Case C is a shearwake with 5632 x 840 x 128 grid points.

3.9.3 Results

Figure 23 shows the results for Test Case A: note that relative efficiency cannot be computed for this test case as it has not been run on a 10 TFlop/s partition of the IBM Power6. Figure 24 and Figure 25 show the results for Test Case B, while Figure 26 and Figure 27 show the results for Test Case C.

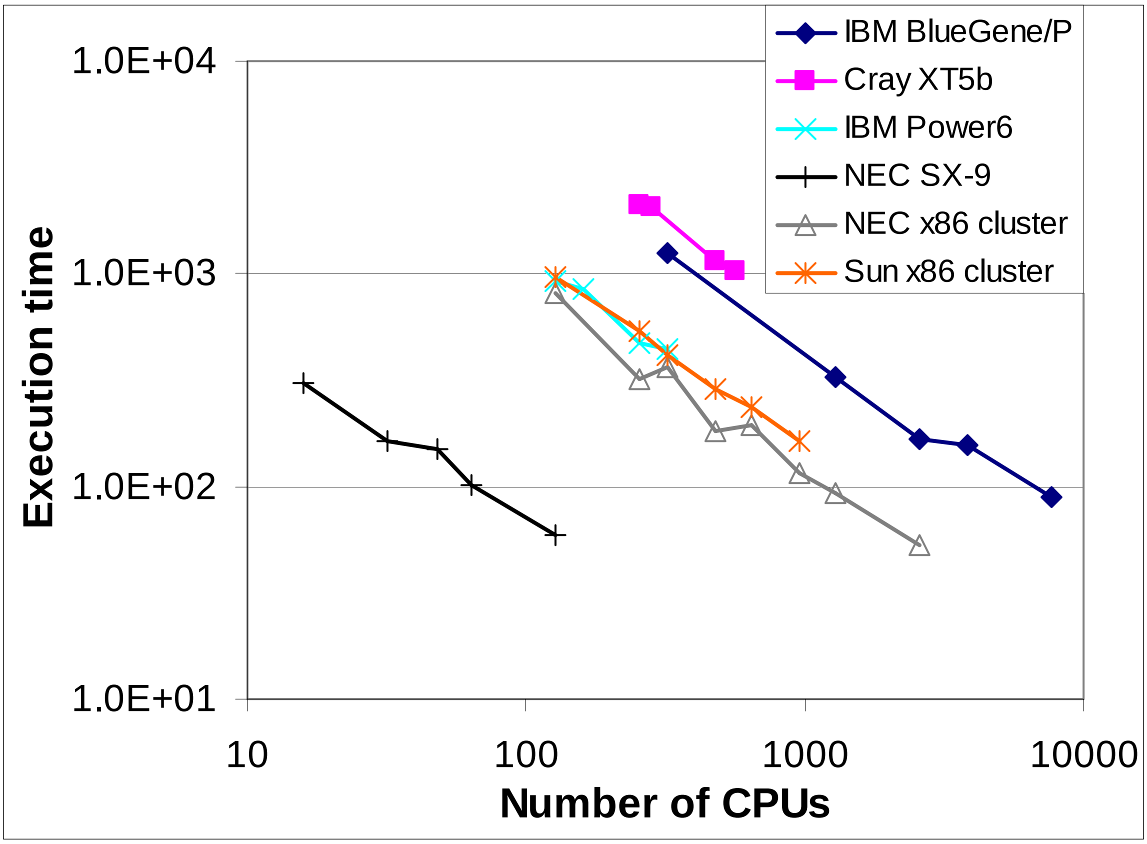

Figure 23: Execution time for NS3D, Test Case A

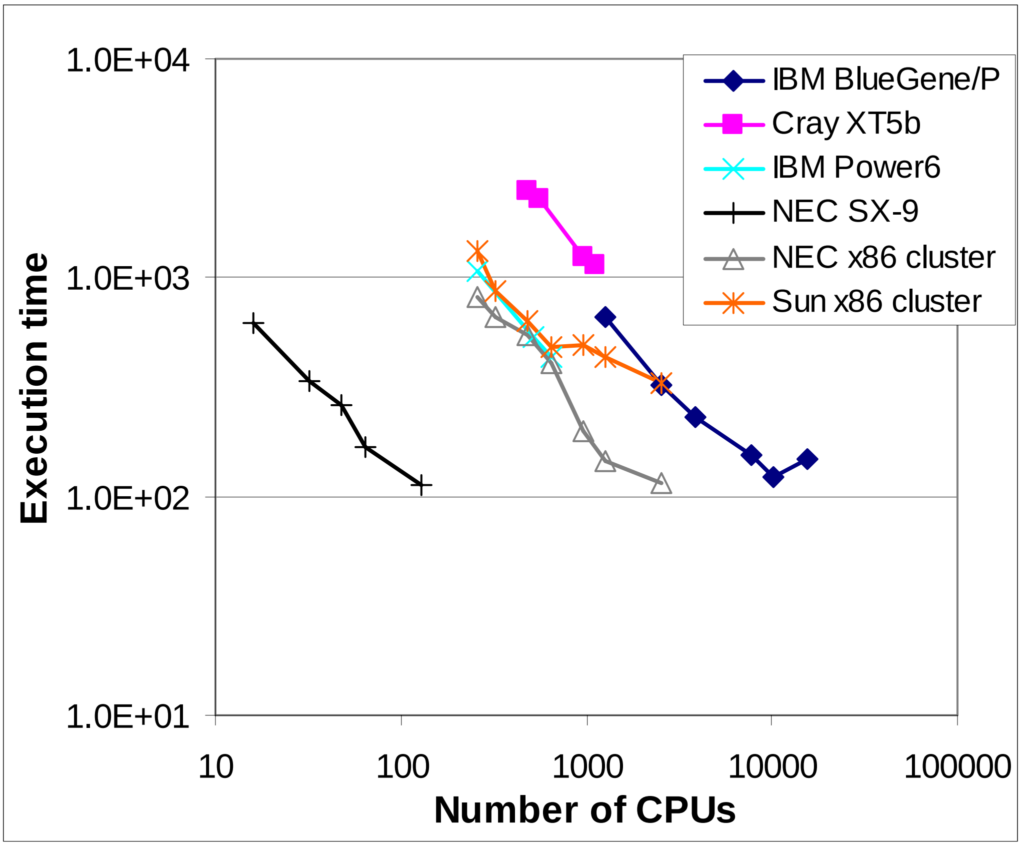

Figure 24: Execution time for NS3D, Test Case B

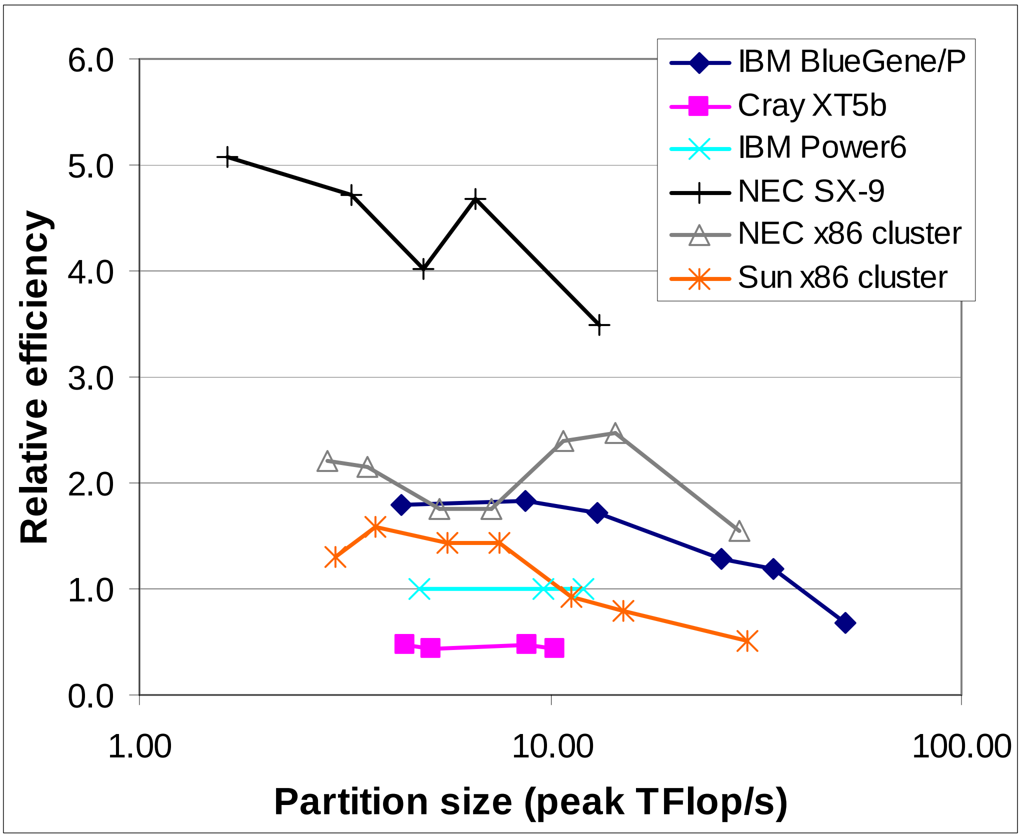

Figure 25: Relative efficiency of NS3D, Test Case B

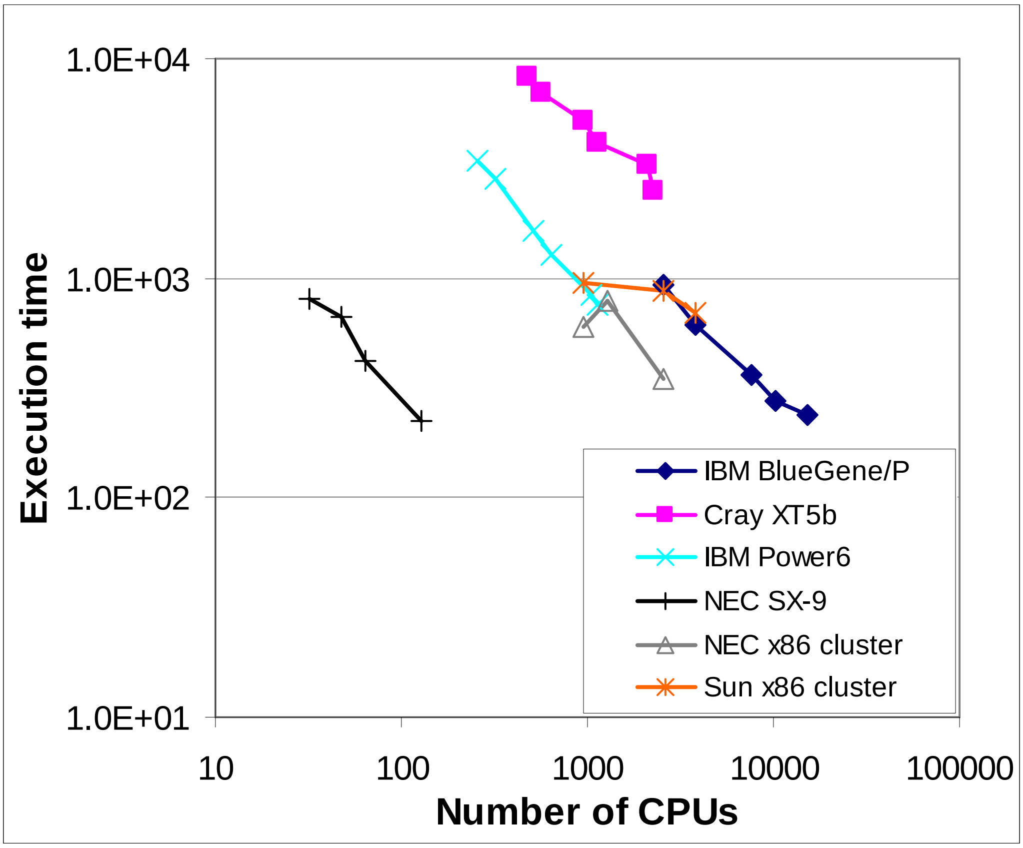

Figure 26: Execution time for NS3D, Test Case C

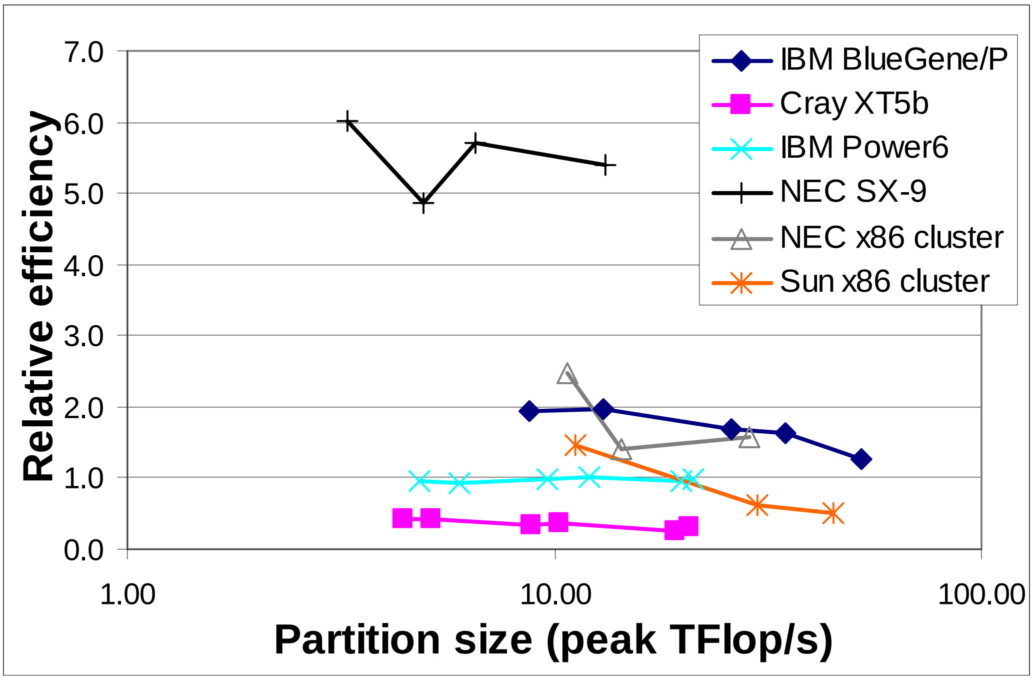

Figure 27: Relative efficiency of NS3D, Test Case C

3.9.4 Analysis

The results from all three test cases show similar trends. The test cases scale well on all the systems, except for on large core counts on the IBM BlueGene/P and on the Sun x86 cluster. Note that on the IBM BlueGene/P and the x86 clusters the code was run using explicit finite differences, instead of compact finite differences. The explicit finite difference method has poorer sequential performance but better scalability. This application shows a wide spread of relative efficiency on the different systems: the IBM Power6 and Sun x86 cluster are about twice as efficient as the Cray XT5, while the IBM BlueGene/P and NEC x86 cluster are about twice as efficient again. The NEC SX-9 is between 4 to 6 times as efficient as the IBM Power6: this reflects the fact that this application has been developed for, and extensively optimised for, vector systems.

3.10 QCD

3.10.1 Summary

The QCD benchmark is, unlike the other benchmarks in the PRACE application benchmark suite, not a full application but a set of 5 kernels which are representative of some of the most compute-intensive parts of QCD calculations.

3.10.2 Test cases

Each of the 5 kernels has one test case:

Kernel A is derived from BQCD (Berlin Quantum ChromoDynamics program), a hybrid Monte-Carlo code that simulates Quantum Chromodynamics with dynamical standard Wilson fermions. The computations take place on a four-dimensional regular grid with periodic boundary conditions. The kernel is a standard conjugate gradient solver with even/odd pre-conditioning. Lattice size is 32^2 x 64^2

Kernel B is derived from SU3_AHiggs, a lattice quantum chromodynamics (QCD) code intended for computing the conditions of the Early Universe. Instead of "full QCD", the code applies an effective field theory, which is valid at high temperatures. In the effective theory, the lattice is 3D. Lattice size is 256^3

Kernel C Lattice size is 8^4. Note that Kernel C can only be run in a weak scaling mode, where each CPU stores the same local lattice size, regardless of the number of CPUs. Ideal scaling for this kernel therefore corresponds to constant execution time, and relative efficiency is computed as $T^{10,Power6}/T^{X,A}$.

Kernel D consists of the core matrix-vector multiplication routine for standard Wilson fermions. The lattice size is 64^4

Kernel E consists of a full conjugate gradient solution using Wilson fermions. Lattice size is 64^3 x 3

3.10.3 Results

Figure 28 and Figure 29 show the results for Kernel A, Figure 30 and Figure 31 for Kernel B, Figure 32 and Figure 33 for Kernel C, Figure 34 and Figure 35 for Kernel D, and Figure 36 and Figure 37 for Kernel E. Figure 38 shows the (geometric) mean relative efficiency over all 5 kernels.

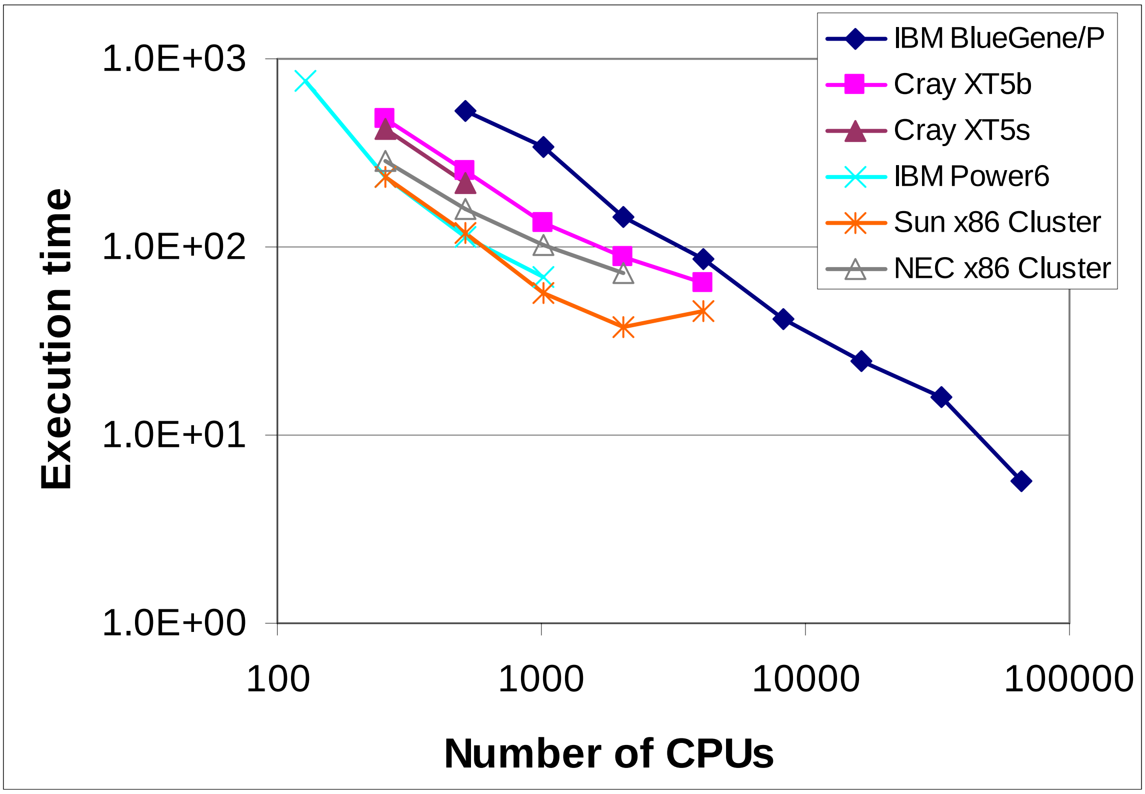

Figure 28: Execution time for QCD, Kernel A

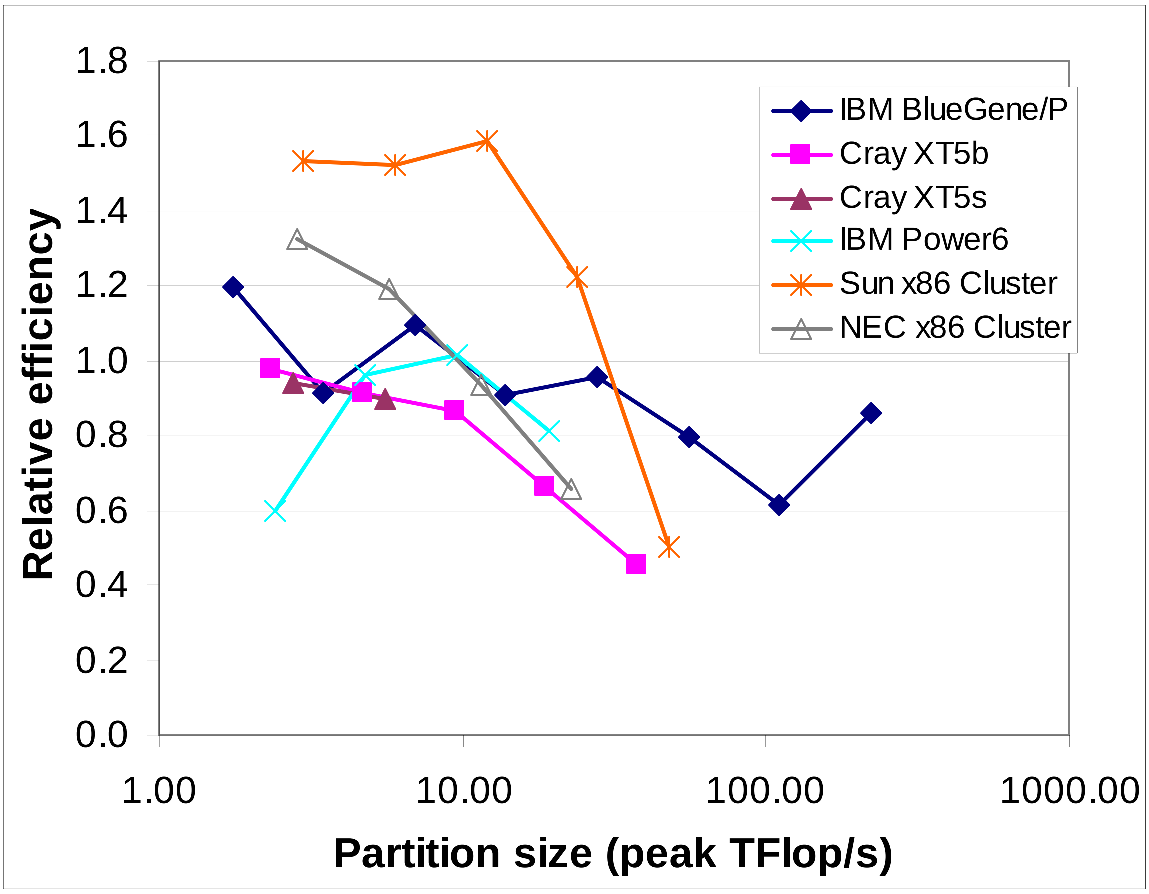

Figure 29: Relative efficiency of QCD, Kernel A

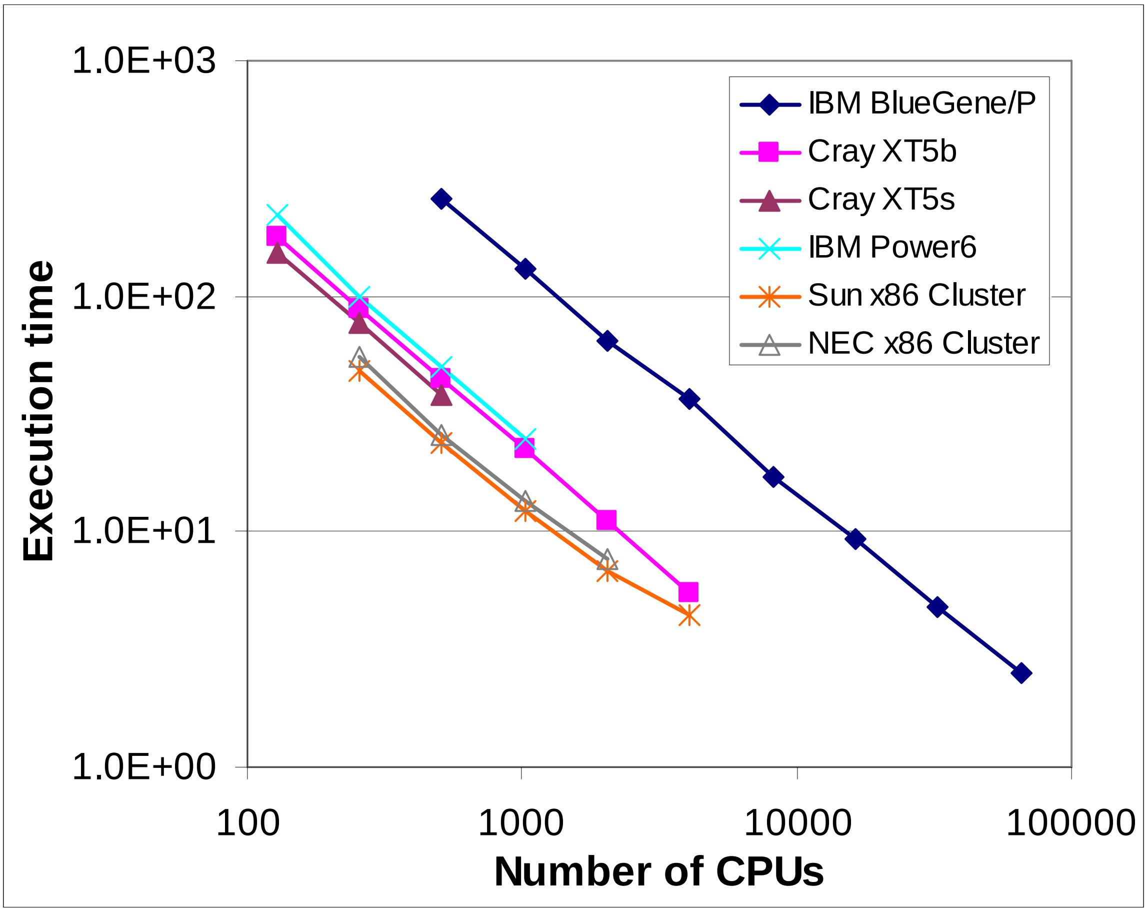

Figure 30: Execution time for QCD, Kernel B

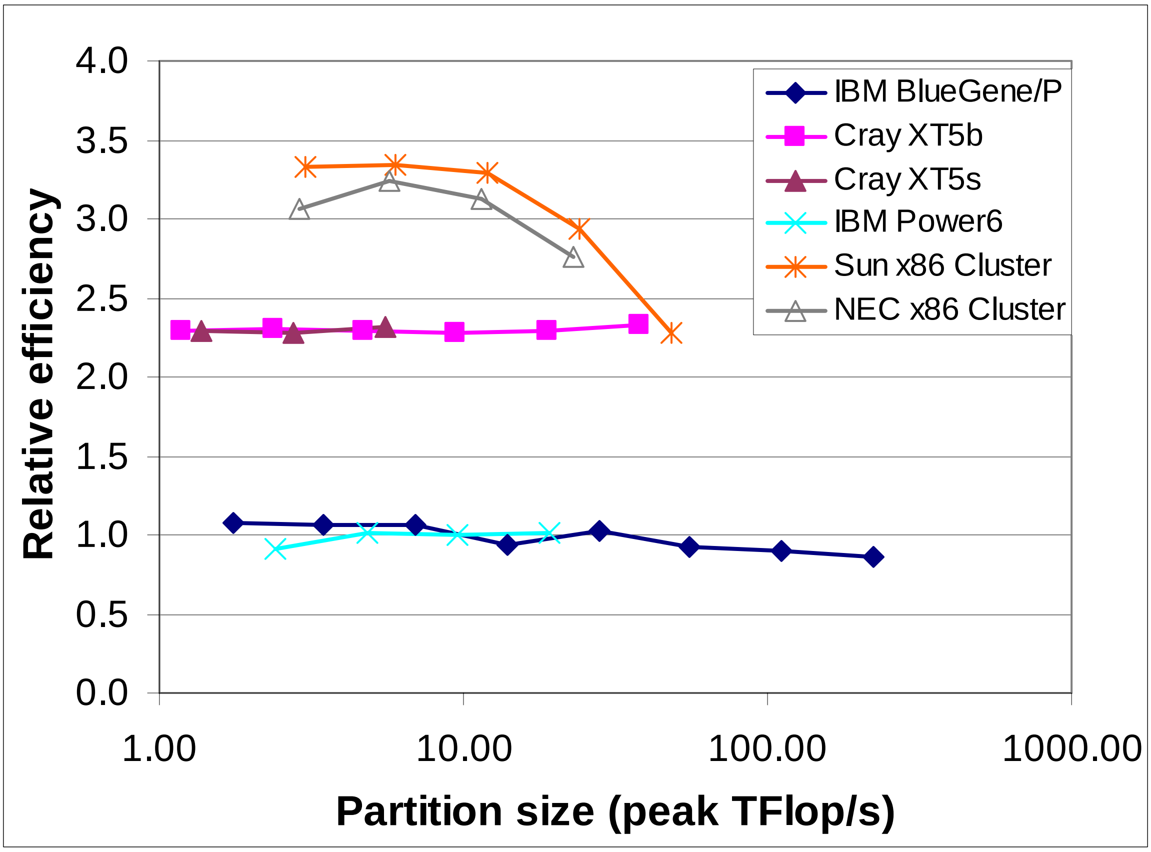

Figure 31: Relative efficiency of QCD, Kernel B

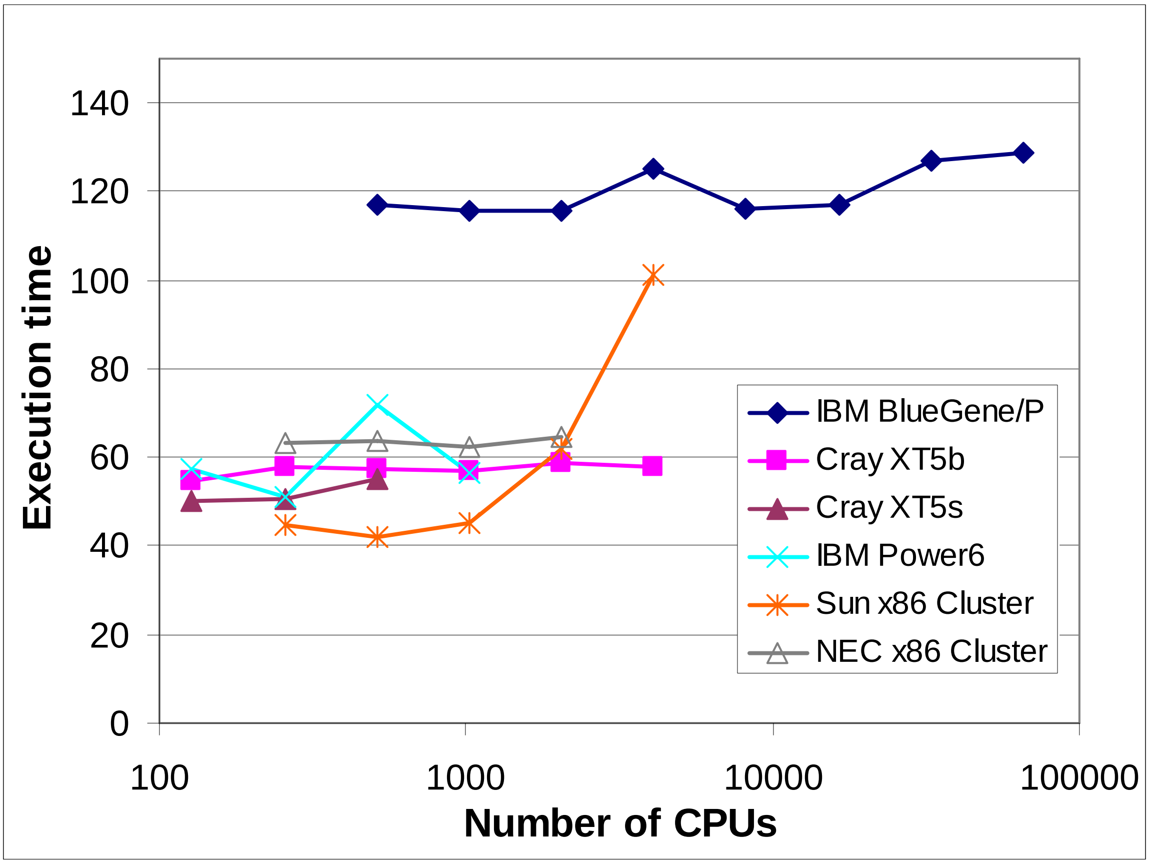

Figure 32: Execution time for QCD, Kernel C

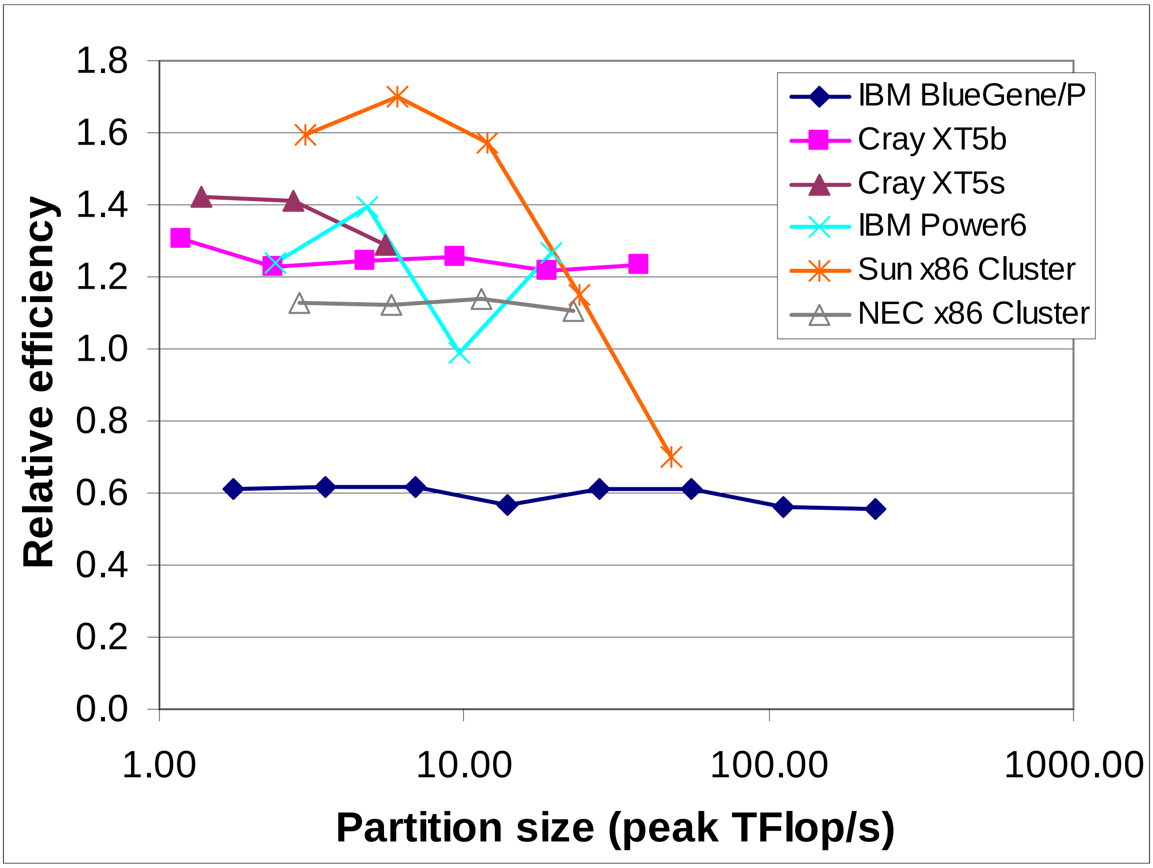

Figure 33: Relative efficiency of QCD, Kernel C

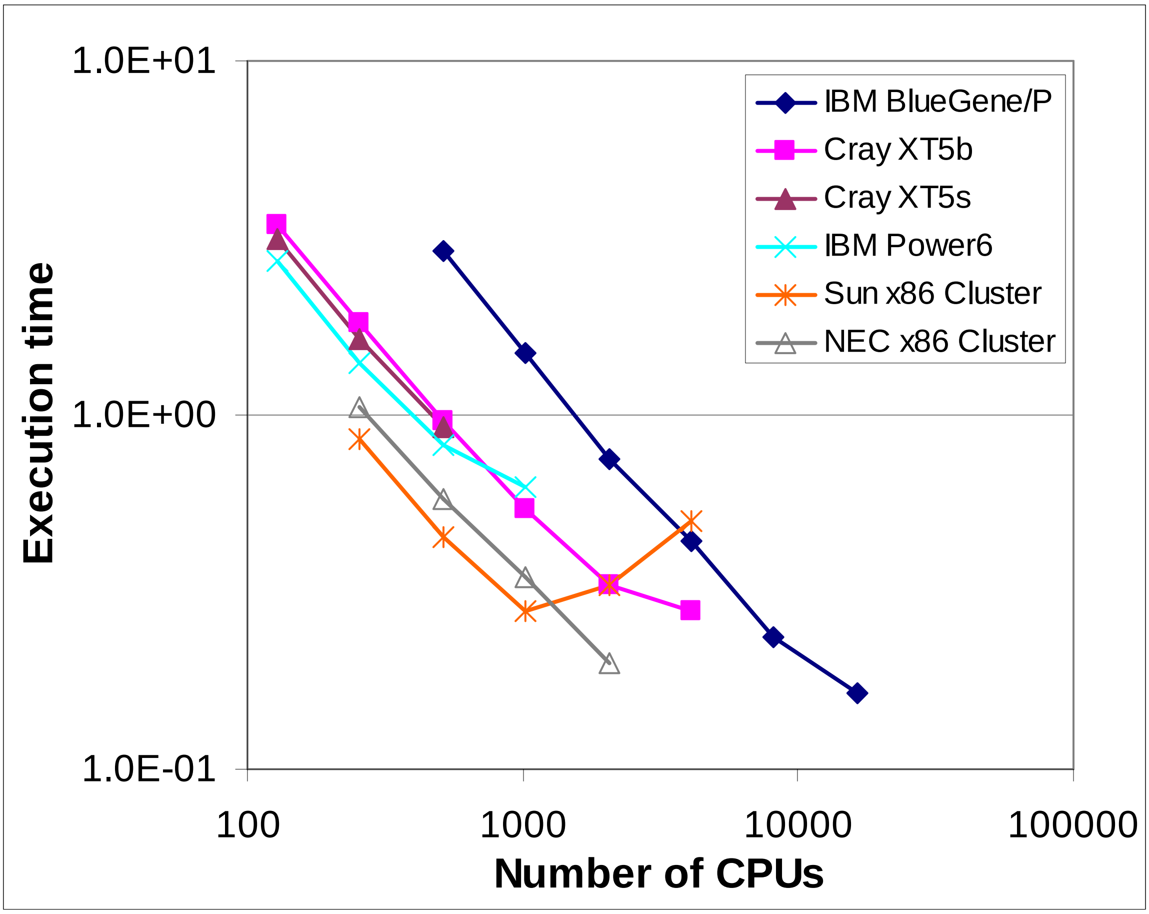

Figure 34: Execution time for QCD, Kernel D

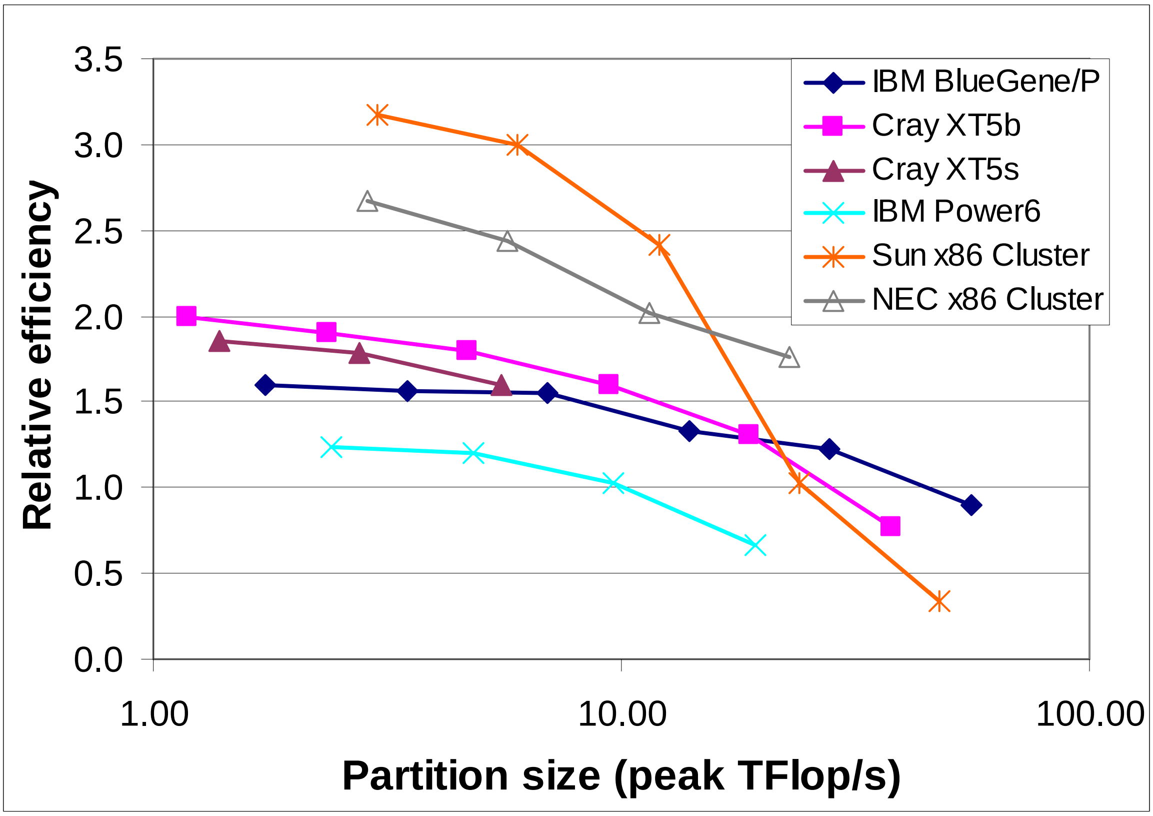

Figure 35: Relative efficiency of QCD, Kernel D

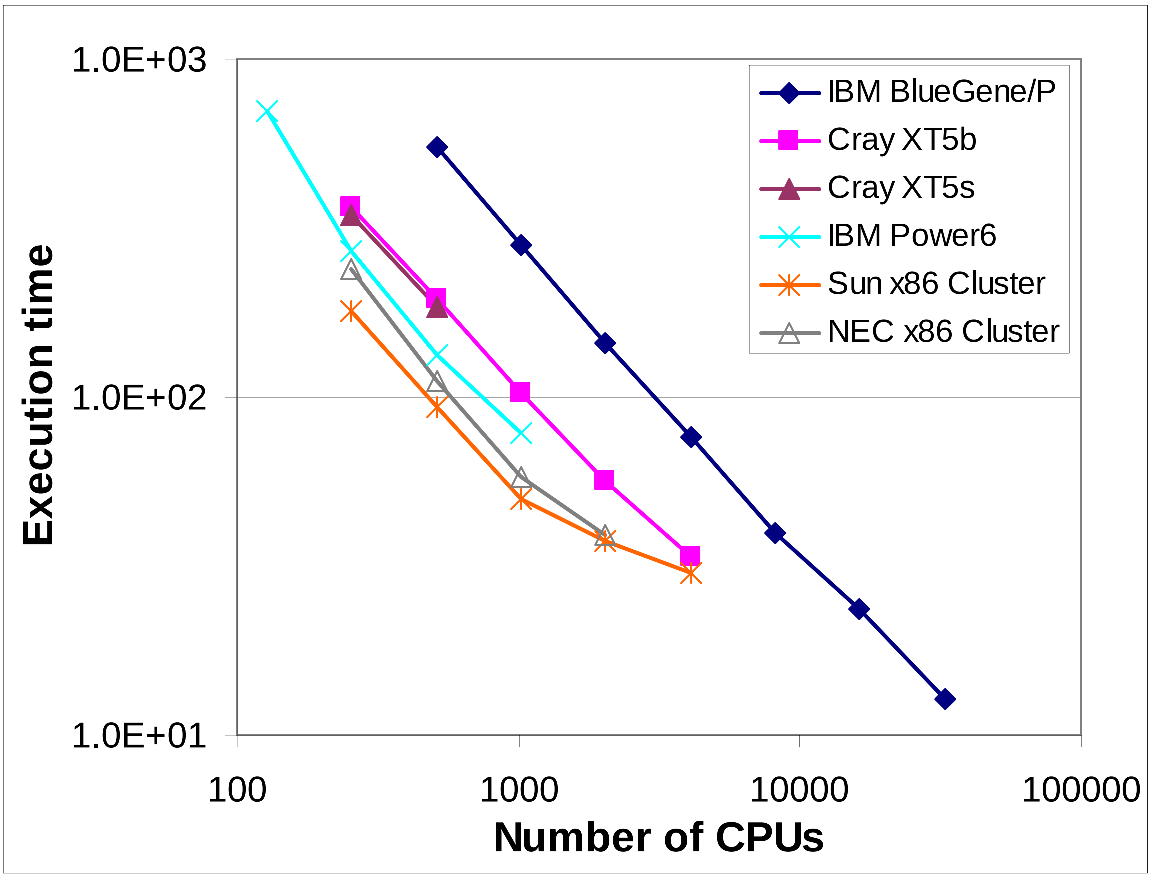

Figure 36: Execution time for QCD, Kernel E

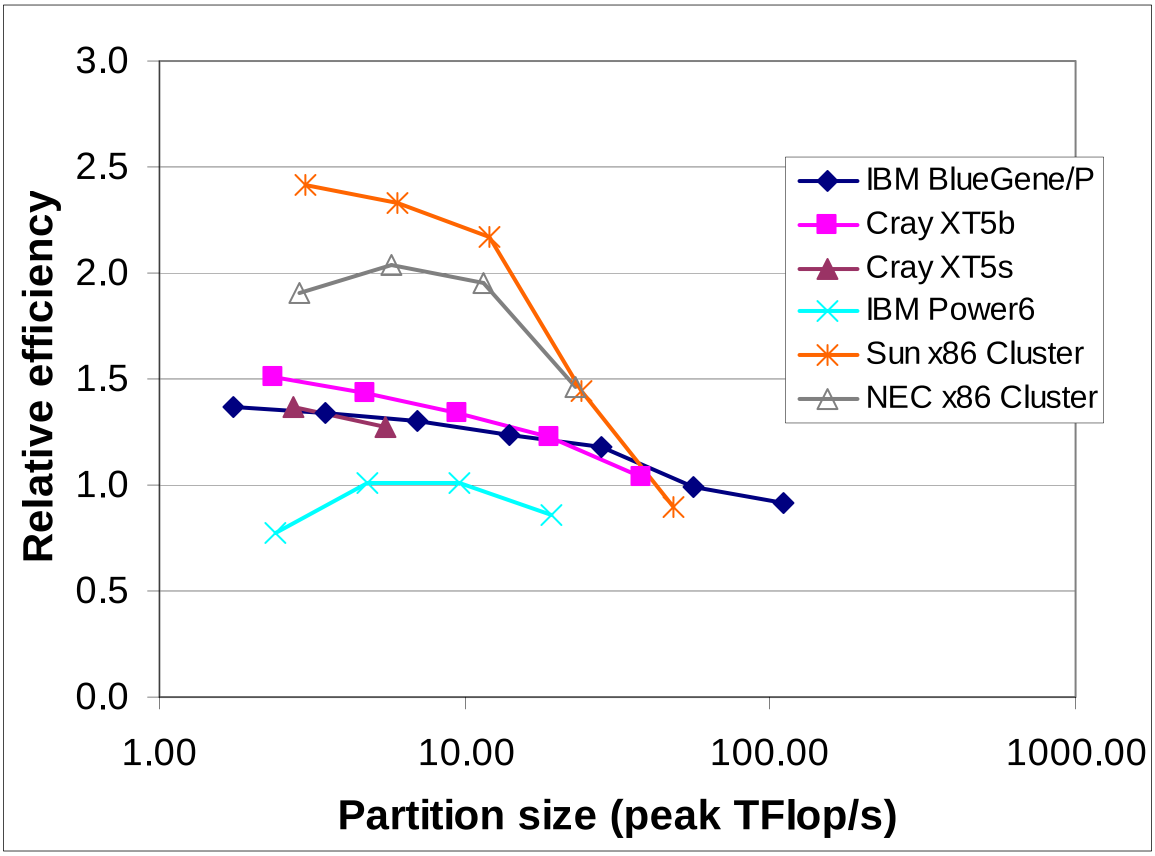

Figure 37: Relative efficiency of QCD, Kernel E

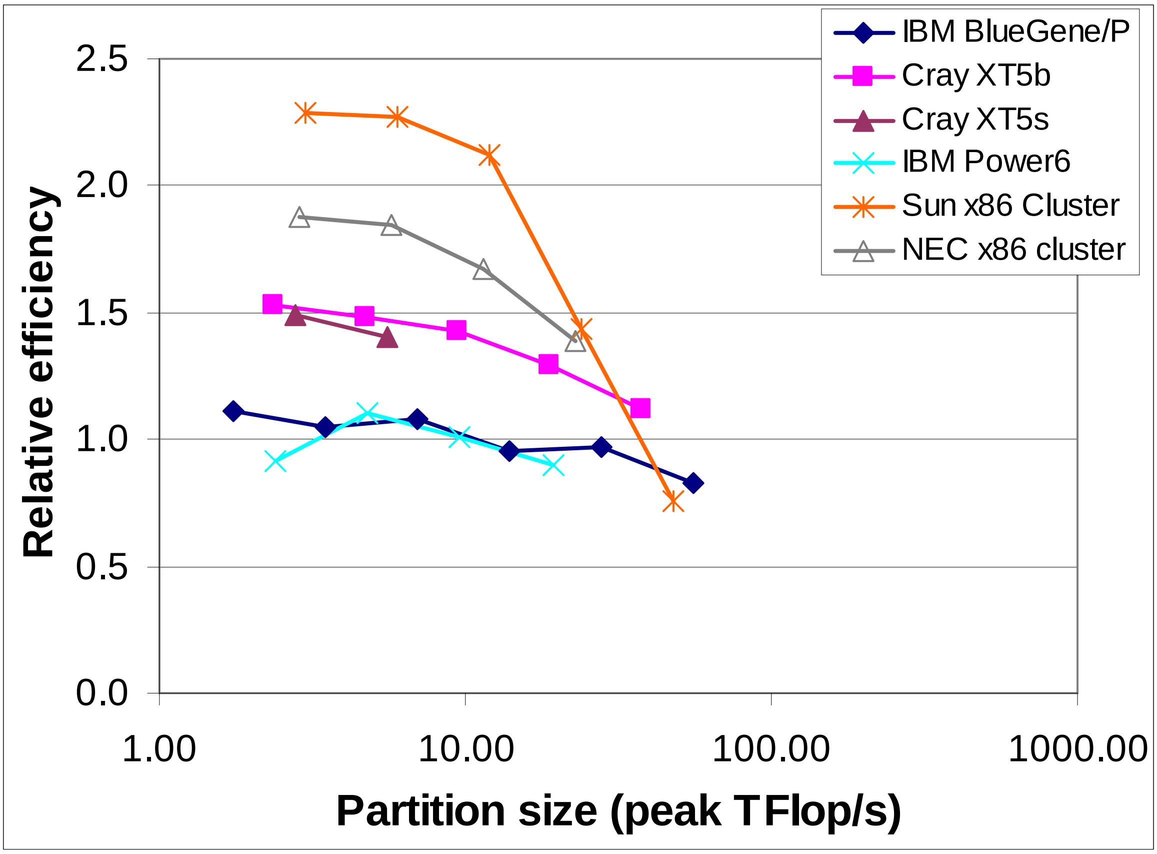

Figure 38: Mean relative efficiency of QCD, all kernels

3.10.4 Analysis

On average, over the 5 kernels, the IBM Power6 and IBM BlueGene/P have similar relative efficiency, though there are significant differences on individual kernels. On average, the Cray XT5 shows better relative efficiency than either of the IBM systems. The Sun x86 cluster is much more efficient than the other systems for small partition sizes, but shows very poor scalability above 10 TFlop/s peak. The NEC x86 cluster has lower efficiency than the Sun x86 cluster (though still higher than all the other systems), but does exhibit such poor scalability.

3.11 QUANTUM ESPRESSO

3.11.1 Summary

QUANTUM ESPRESSO is an integrated suite of computer codes for electronic-structure calculations and materials modelling at the nanoscale, based on density-functional theory, plane waves, and pseudopotentials (norm conserving, ultrasoft, and PAW).

3.11.2 Test cases

Test Case A is a simulation of an iridium surface with a graphene sheet monolayer on top of it, using 443 atoms.

Test Case B is a simulation of an iridium surface with a graphene sheet monolayer on top of it, using 686 atoms.

3.11.3 Results

Figure 39 and Figure 40 show the results for Test Case A. Figure 41 and Figure 42 show the results for Test Case B.

3.11.4 Analysis

Test Case A scales better on the Cray XT5 than on the IBM Power6, but the relative efficiency at 20 TFlop/s peak is very similar. The Sun x86 cluster shows high relative efficiency for small core counts, but scalability beyond 512 CPUs is poor. The NEC x86 system shows slightly poorer efficiency than the Sun x86 system on small partition sizes, but much better scalability. The NEC SX-9 shows modest efficiency for a vector system, and poor scalability: little effort has been expended on optimising this application specifically for the SX-9. The IBM BlueGene/P shows good scalability at very high core counts, but the relative efficiency is very low: due to memory restrictions, this application was run using mixed-mode MPI + OpenMP on this system.

Test Case B shows similar relative efficiencies, although there are too few data points to make clear statements about scalability.

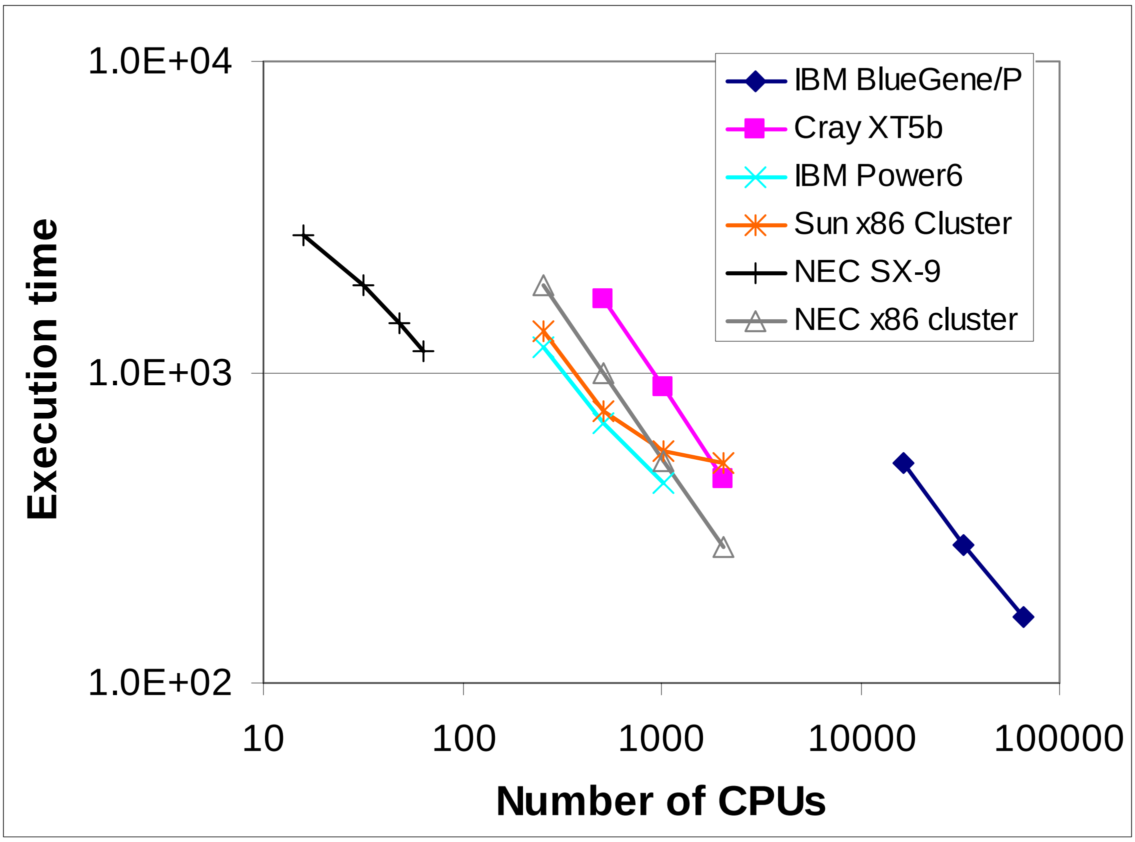

Figure 39: Execution time for QUANTUM ESPRESSO, Test Case A

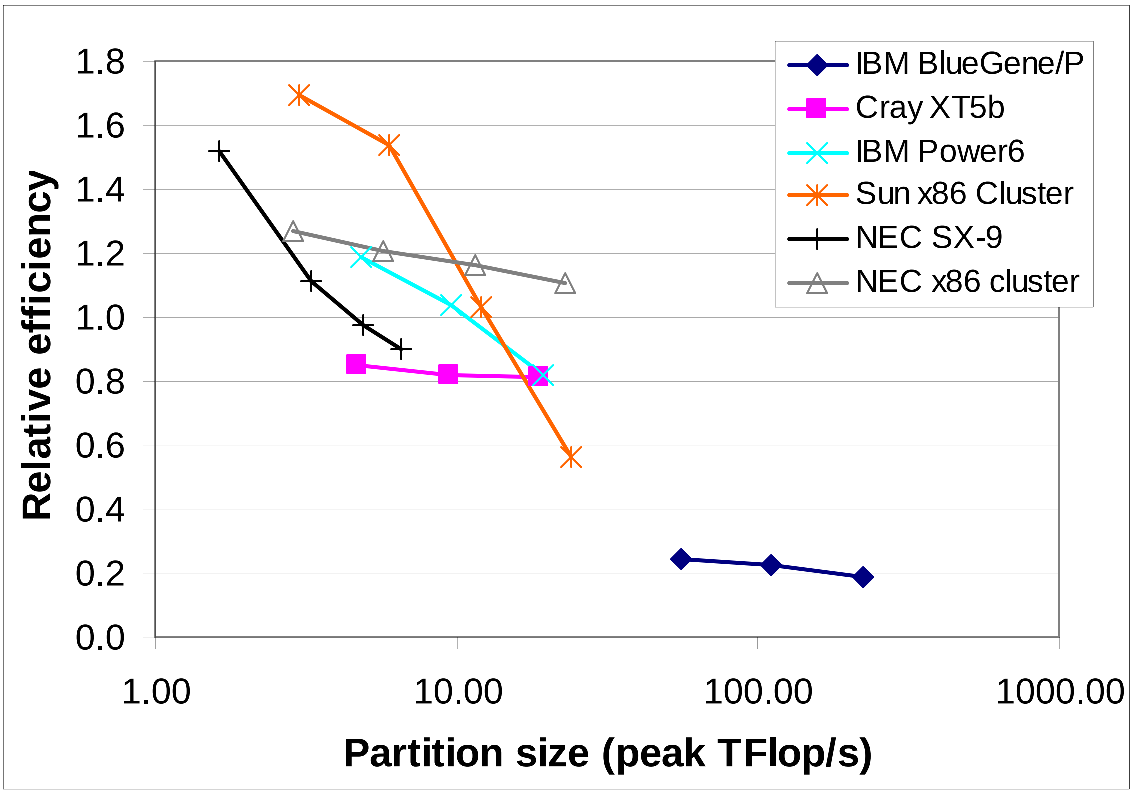

Figure 40: Relative efficiency of QUANTUM ESPRESSO, Test Case A

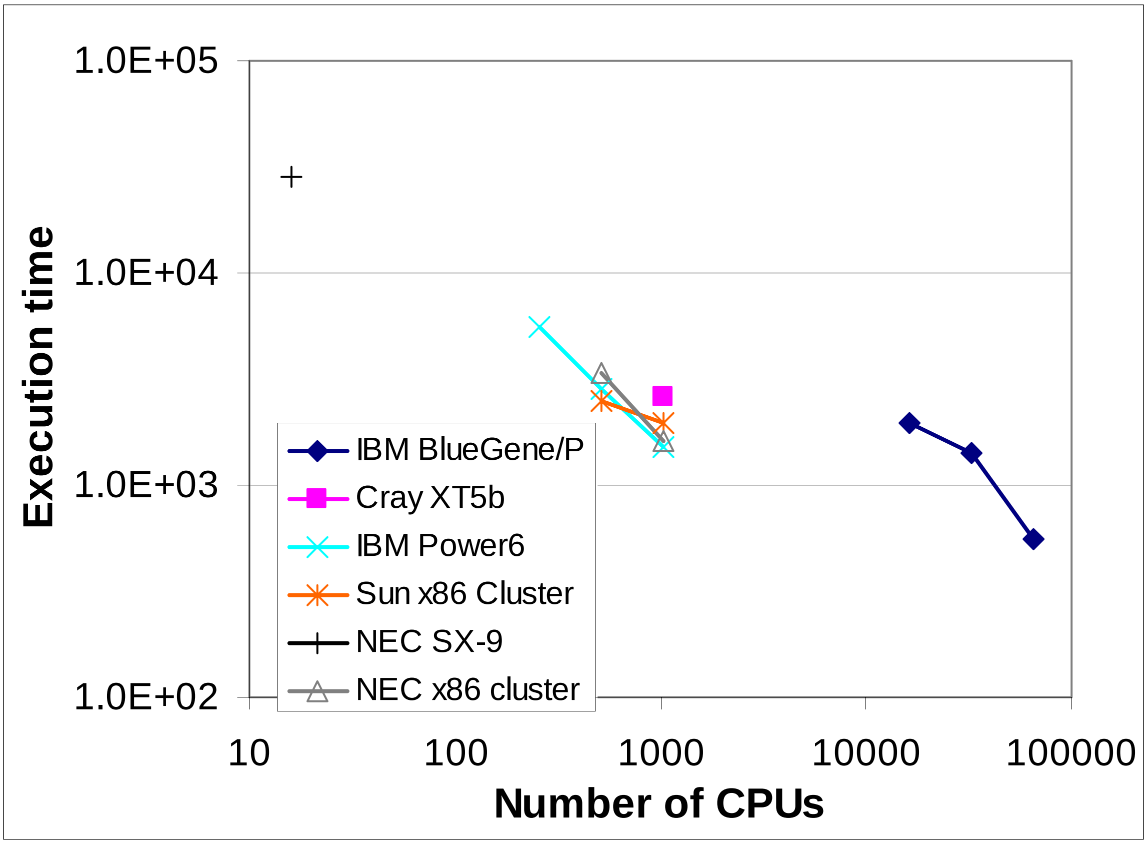

Figure 41: Execution time for QUANTUM ESPRESSO, Test Case B

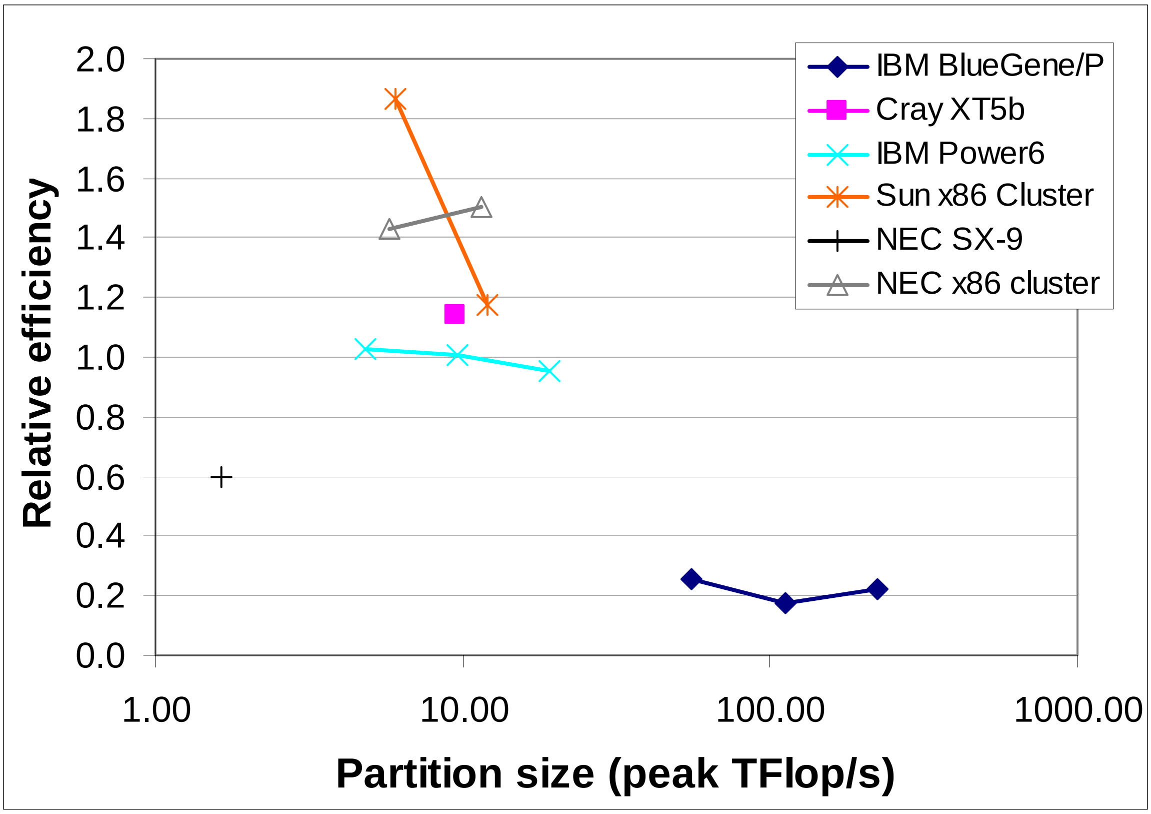

Figure 42: Relative efficiency of QUANTUM ESPRESSO, Test Case B

3.12 WRF

3.12.1 Summary

The Weather Research & Forecasting (WRF) model was developed at the National Center for Atmospheric Research (NCAR) in the United States as a regional- to global-scale model for both research applications and operational weather-forecast systems.

3.12.2 Test cases

Test Case A is a single domain configuration over the North Atlantic and UK consisting of approximately 1200 x 1200 grid points.

3.12.3 Results

Figure 43 and Figure 44 show the results for Test Case A.

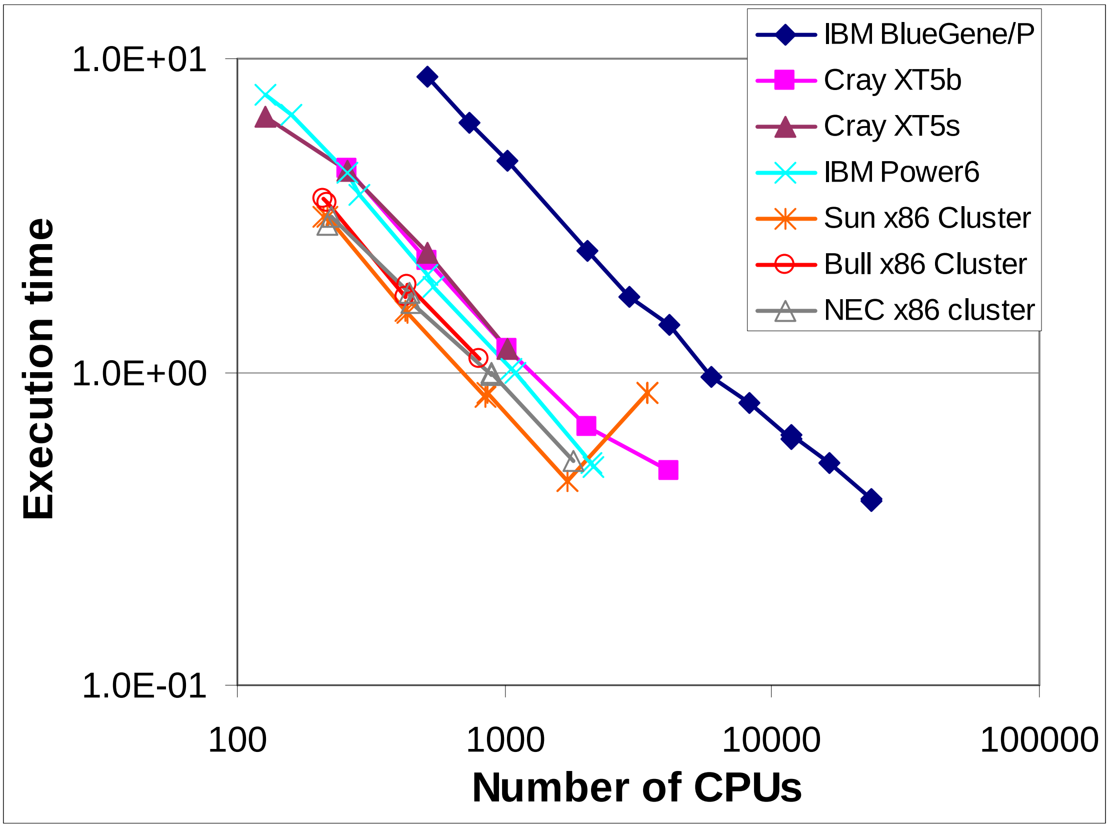

Figure 43: Execution time for WRF, Test Case A

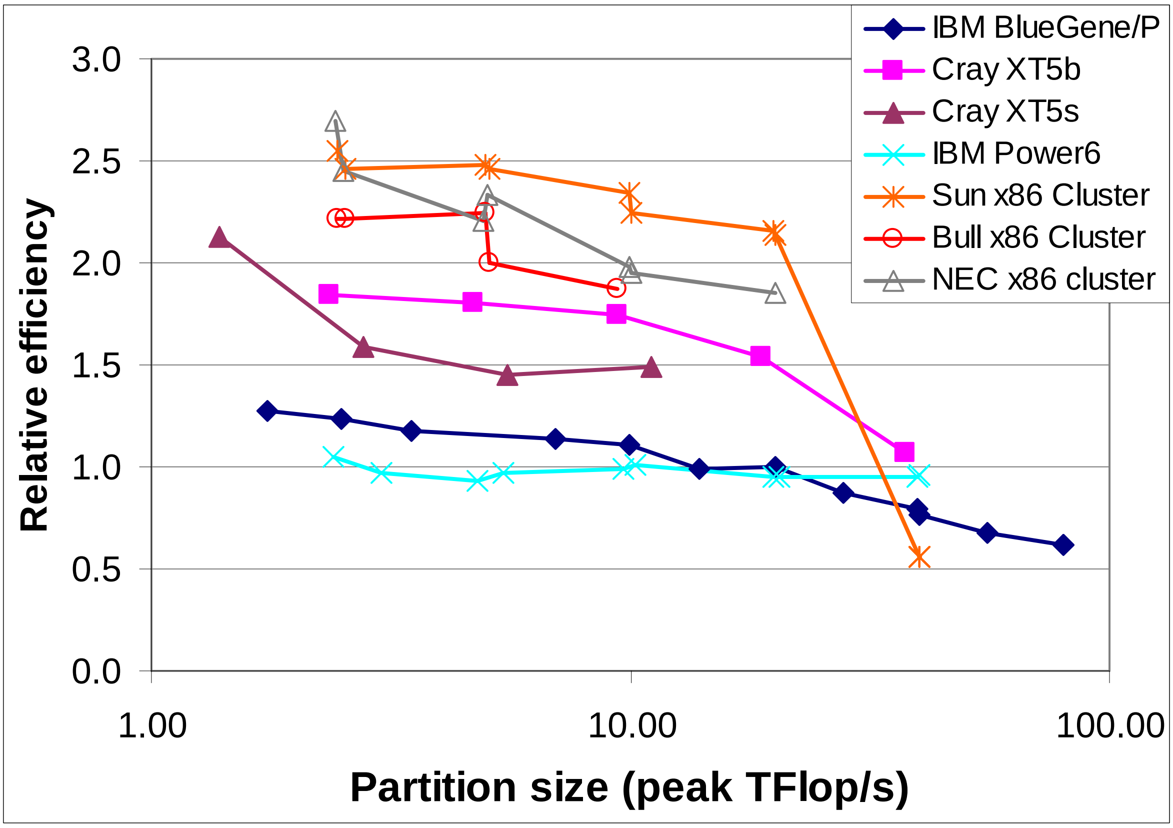

Figure 44: Relative efficiency of WRF, Test Case A

3.12.4 Analysis

For small partition sizes the IBM Power6 has the lowest relative efficiency, but scalability is good. The IBM BlueGene/P is slightly more efficient than the Power6 up to 20 TFlop/s peak, but slightly less efficient at larger partition sizes. On the Cray XT5, almost no performance gain is seen when the faster Shanghai cores are used instead of Barcelona cores. Efficiency on this system is good on small partition sizes, but the scalability is poor above 20 TFlop/s Peak.

The x86 clusters are over twice as efficient as the IBM Power6 for small partition sizes, but on the Sun x86 cluster, scalability is completely exhausted at 40 TFlop/s.

3.13 ALYA

3.13.1 Summary

ALYA is a Computational Mechanics (CM) code capable of solving different physics, each one with its own modelling characteristics, in a coupled way. Among the problems it solves are: Convection-Diffusion-Reaction, Incompressible Flows, Compressible Flows, Turbulence, Bi-Phasic Flows and free surface, Excitable Media, Acoustics, Thermal Flow, Quantum Mechanics (TDFT) and Solid Mechanics (Large strain).

3.13.2 Test cases

Test Case A is a 3D cavity with 256^3 mesh elements.

3.13.3 Results

Figure 45 and Figure 46 show the results for Test Case A.

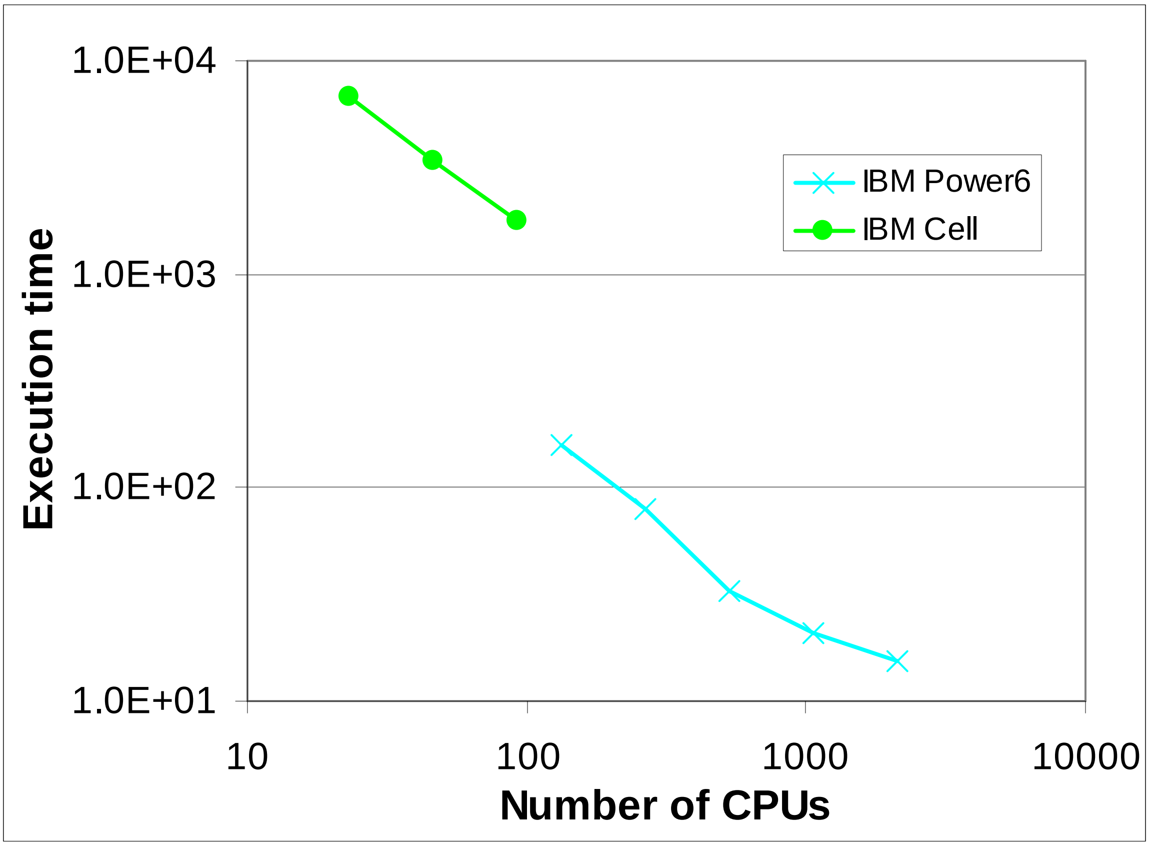

Figure 45: Execution time for ALYA, Test Case A

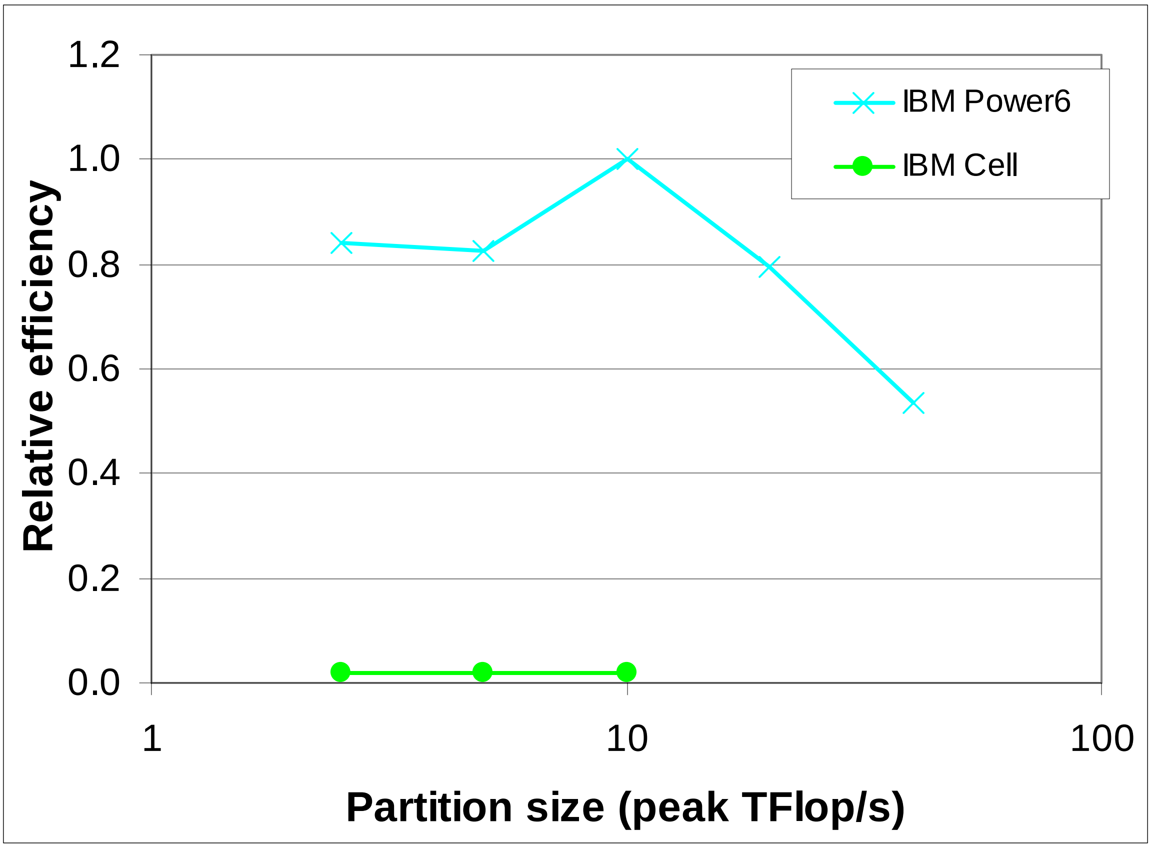

Figure 46: Relative efficiency of ALYA, Test Case A

3.13.4 Analysis

The IBM Cell system has very poor relative efficiency for this application: this is because the code is running on the PPE part of the Cell processor only: porting to the SPEs is still in progress.

3.14 AVBP

3.14.1 Summary

AVBP is a massively parallel CFD code that solves laminar and turbulent compressible reacting flows.

3.14.2 Test cases

Test Case A is a simulation containing 37 million cells.

3.14.3 Results

Figure 47 and Figure 48 show the results for Test Case A.

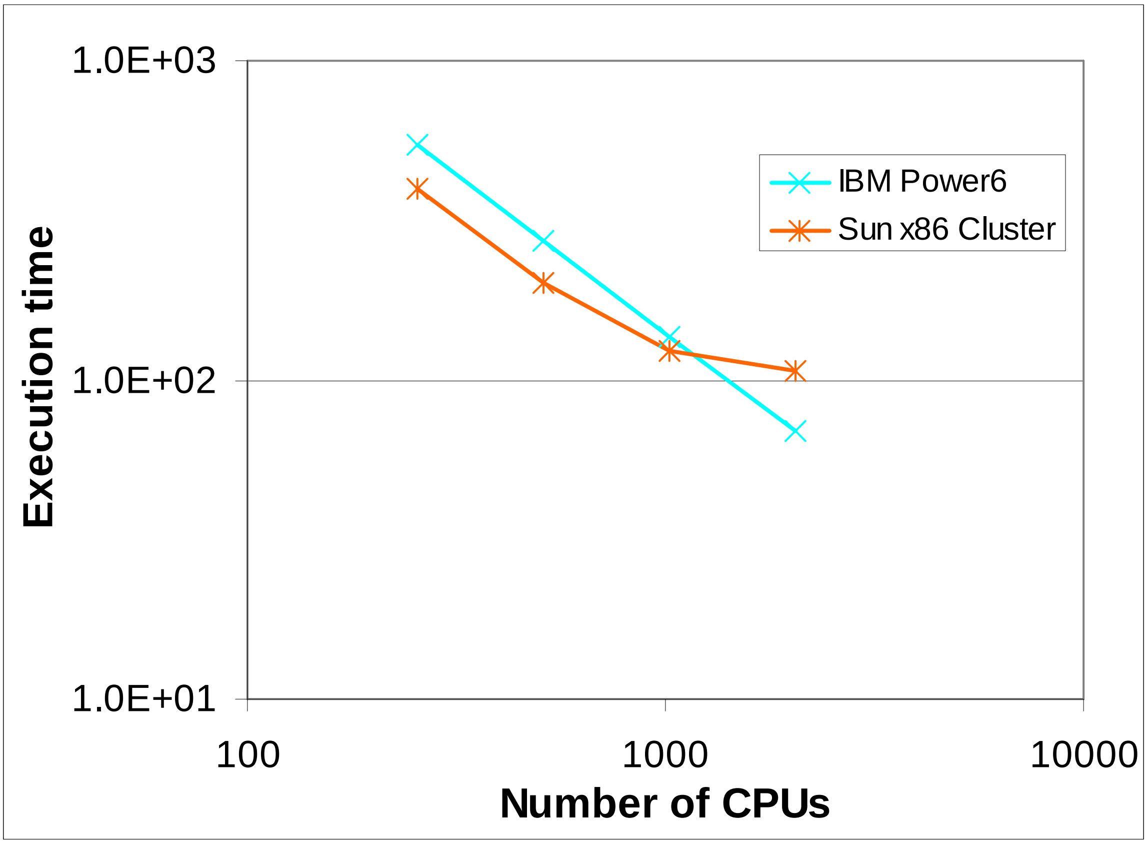

Figure 47: Execution time for AVBP, Test Case A

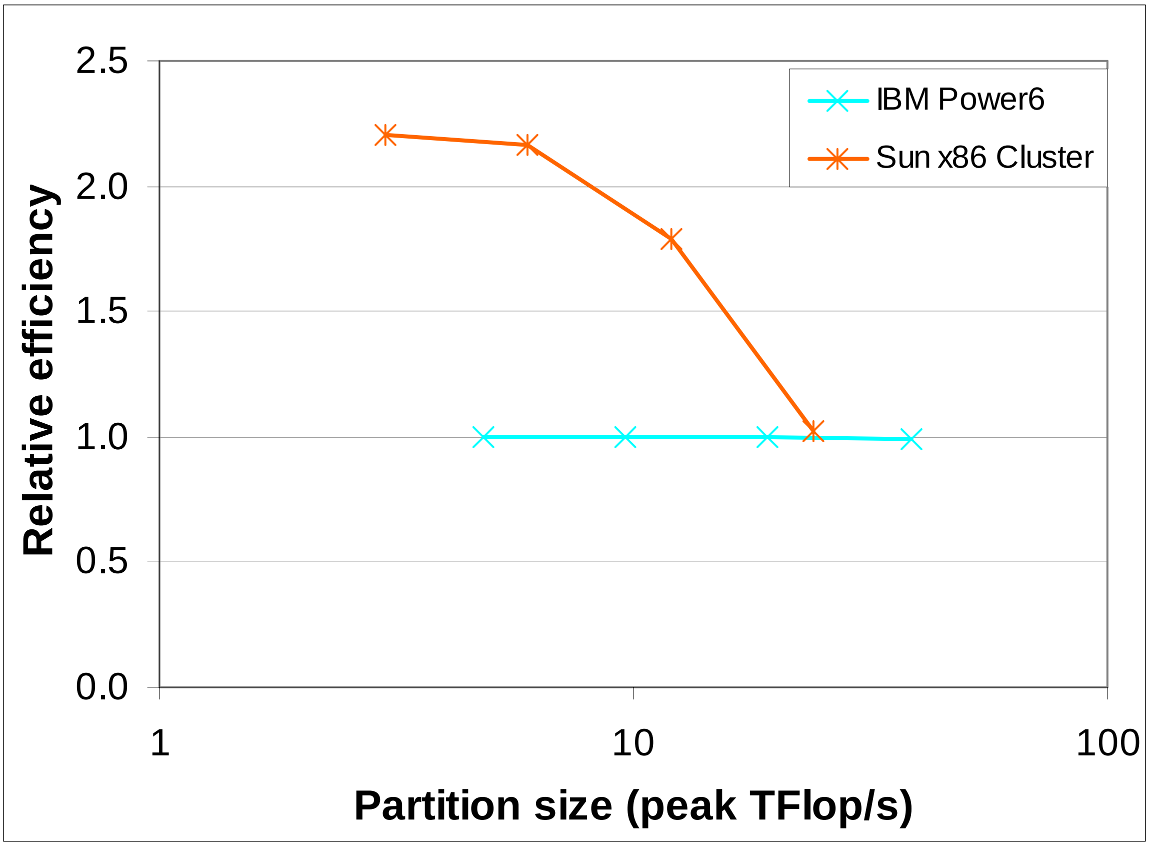

Figure 48: Relative efficiency of AVBP, Test Case A

3.14.4 Analysis

This test case scales very well on the IBM Power6 up to 40 TFlop/s. Although the relative efficiency on the Sun x86 cluster is over twice that on the IBM Power6 for small partition sizes, the scalability is much poorer.

3.15 BSIT

3.15.1 Summary

BSIT is a tool set for computational geophysics. The computational kernel which forms the benchmark application performs Reverse Time Migration.

3.15.2 Test cases

Test Case A is a weak scaling test where each node computes an input dataset of size 5443.

Ideal scaling for this application therefore corresponds to constant execution time, and relative efficiency is computed as T10,Power6/TX,A.

3.15.3 Results

Figure 49 and Figure 50 show the results for Test Case A.

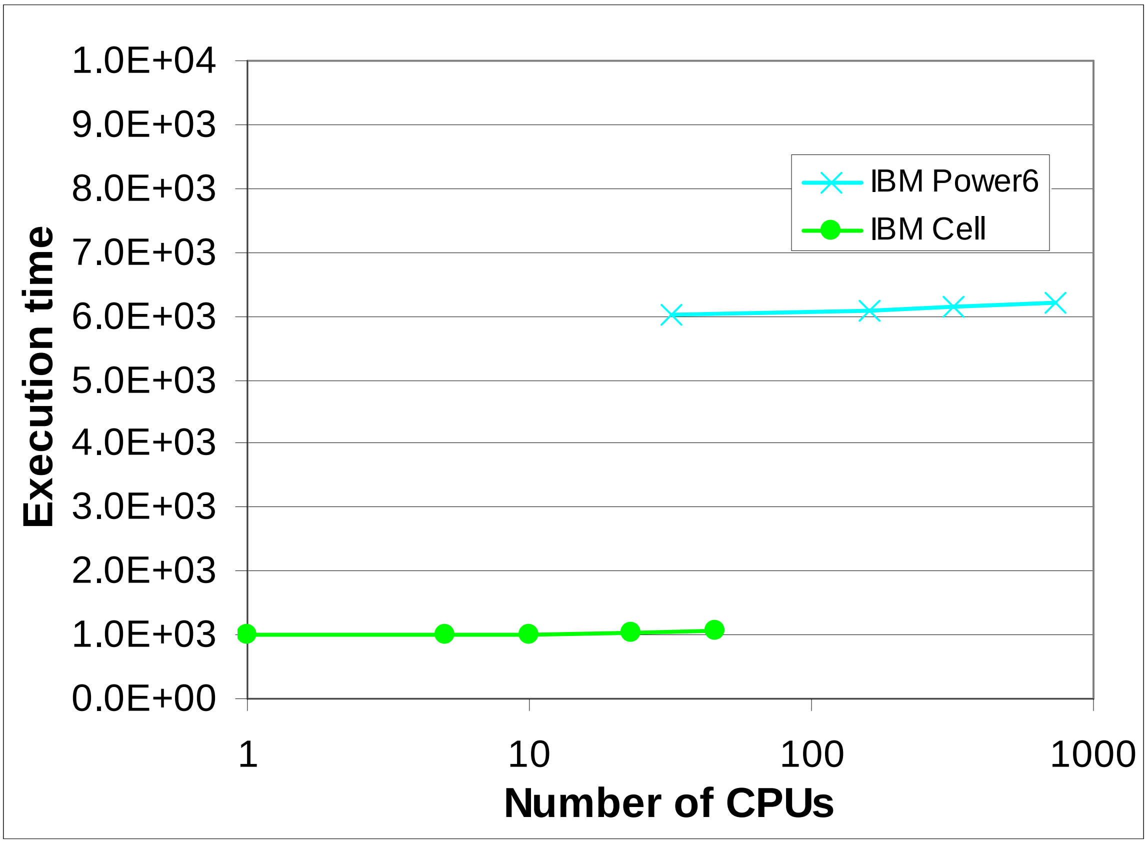

Figure 49: Execution time for BSIT, Test Case A

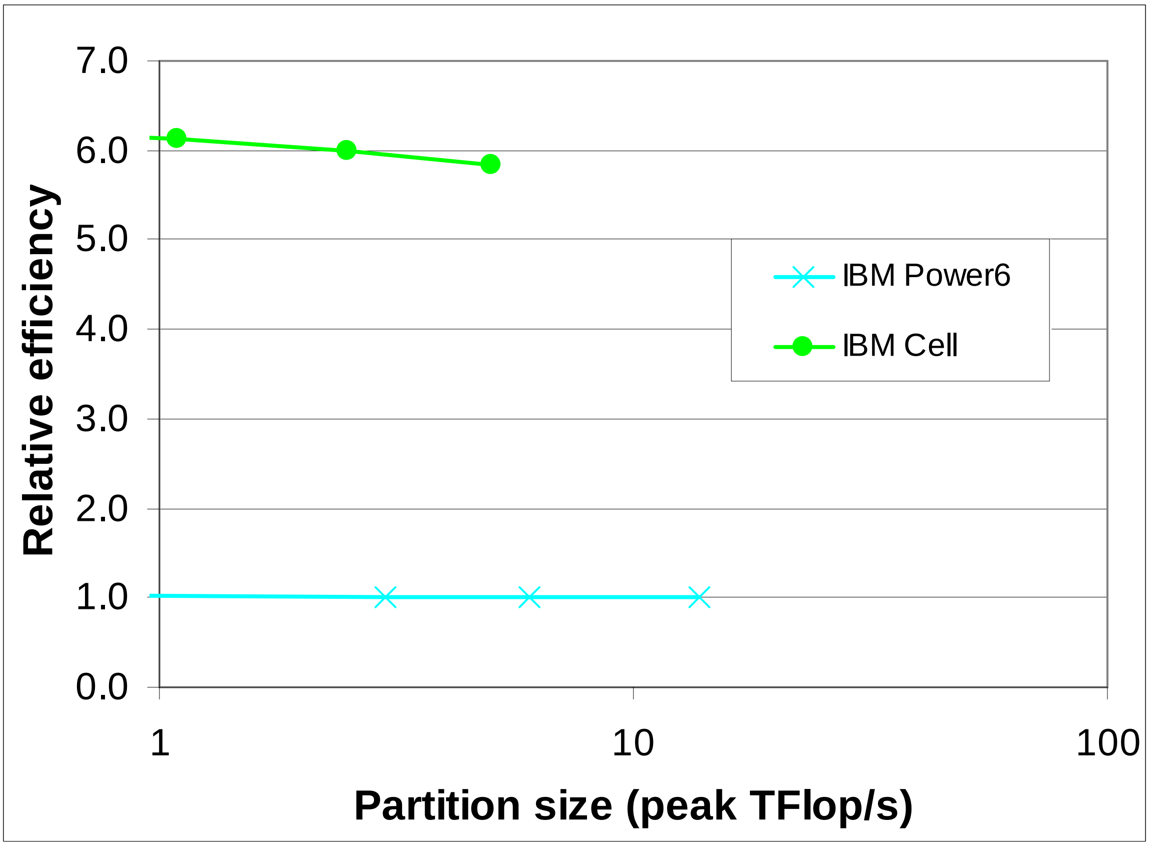

Figure 50: Relative efficiency of BSIT, Test Case A

3.15.4 Analysis

For this application, the IBM Cell system is around six times more efficient than the IBM Power6. This application has been fully ported to the Cell SPE cores.

3.16 ELMER

3.16.1 Summary

Elmer is a finite element software package for the solution of partial differential equations. Elmer can deal with a great number of different equations, which may be coupled in a generic manner making Elmer a versatile tool for multiphysical simulations.

3.16.2 Test cases

Test Case A is a finite element model for heat conduction in a three dimensional solid.

The problem is solved by trilinear hexagonal elements, incomplete LU-precoditioniong (ILU0), and iterative Krylov space methods (BiCGStab) for linear algebra.

3.16.3 Results

Figure 51 and Figure 52 show the results for Test Case A.

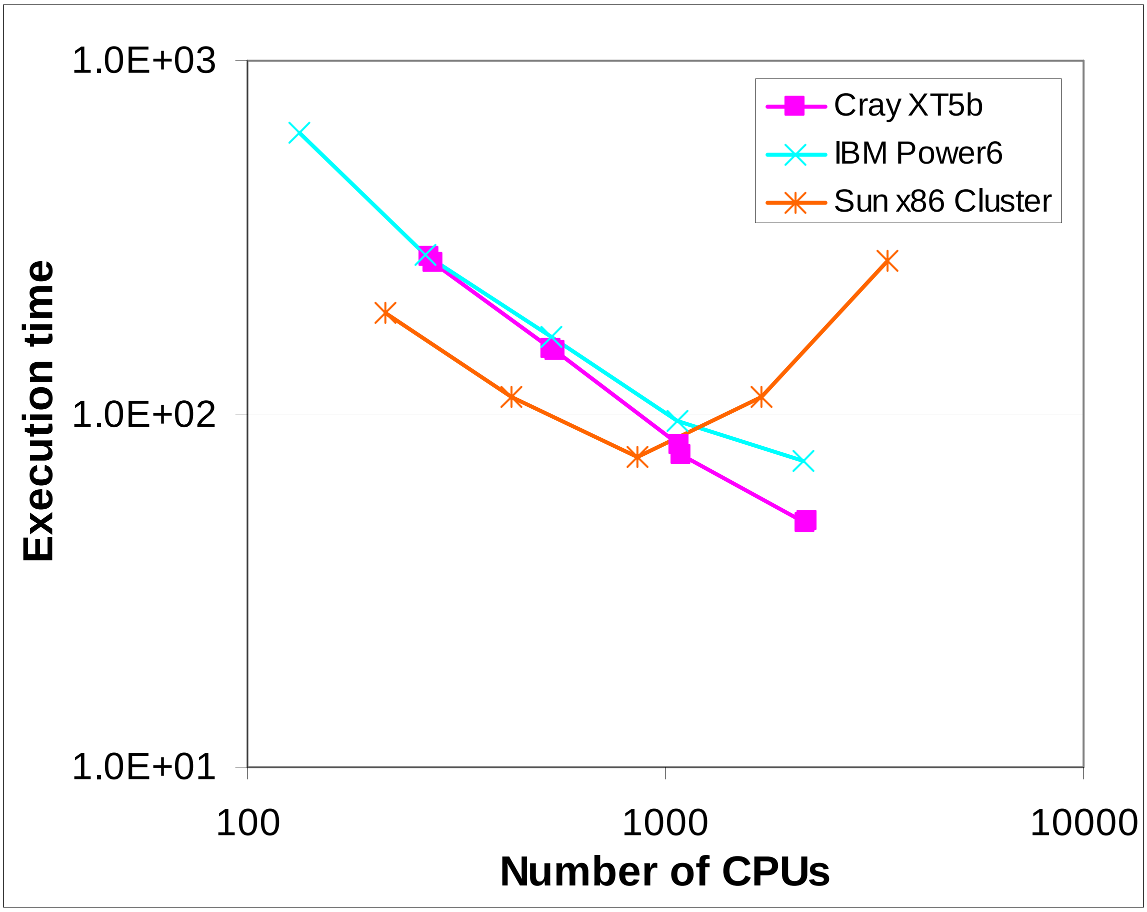

Figure 51: Execution time for ELMER, Test Case A

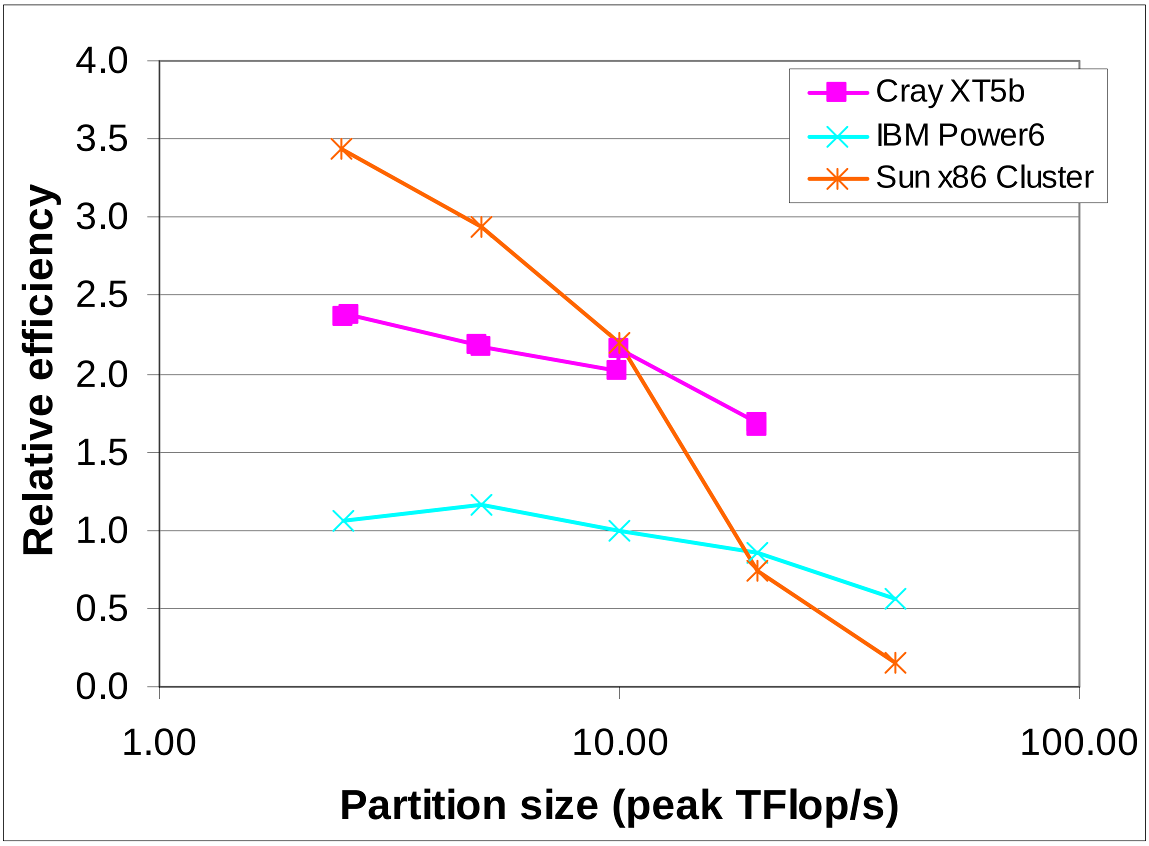

Figure 52: Relative efficiency of ELMER, Test Case A

3.16.4 Analysis

The test case scales well on both the Cray XT5 and IBM Power6, but the XT5 is between 2 and 2.5 times more efficient then the Power6. The Sun x86 cluster is over 3 times as efficient as the IBM Power6 on small partition sizes, but the scalability is very poor compared to the other two systems.

3.17 GPAW

3.17.1 Summary

GPAW is a density-functional theory (DFT) code based on the projector-augmented wave (PAW) method. It uses real-space uniform grids and multigrid methods.

3.17.2 Test cases

Test Case A simulates the ground state electronic structure of 256 water molecules. The number of atoms is 768, a real-space grid of 112 x 112 x 112 is used and the number of electronic states is 1056. Parallelisation is done with domain decomposition over the real-space grid.

Test Case B is a calculation of optical spectra within the time-dependent density functional theory for an Au55 cluster. The number of atoms is 55, a real-space grid of 96 x 96 x 96 is used and the number of electronic states is 320. Parallelisation is done with domain decomposition over the real-space grid and over the electronic states.

3.17.3 Results

Figure 53 and Figure 54 show the results for Test Case A. Figure 55 and Figure 56 show the results for Test Case B.

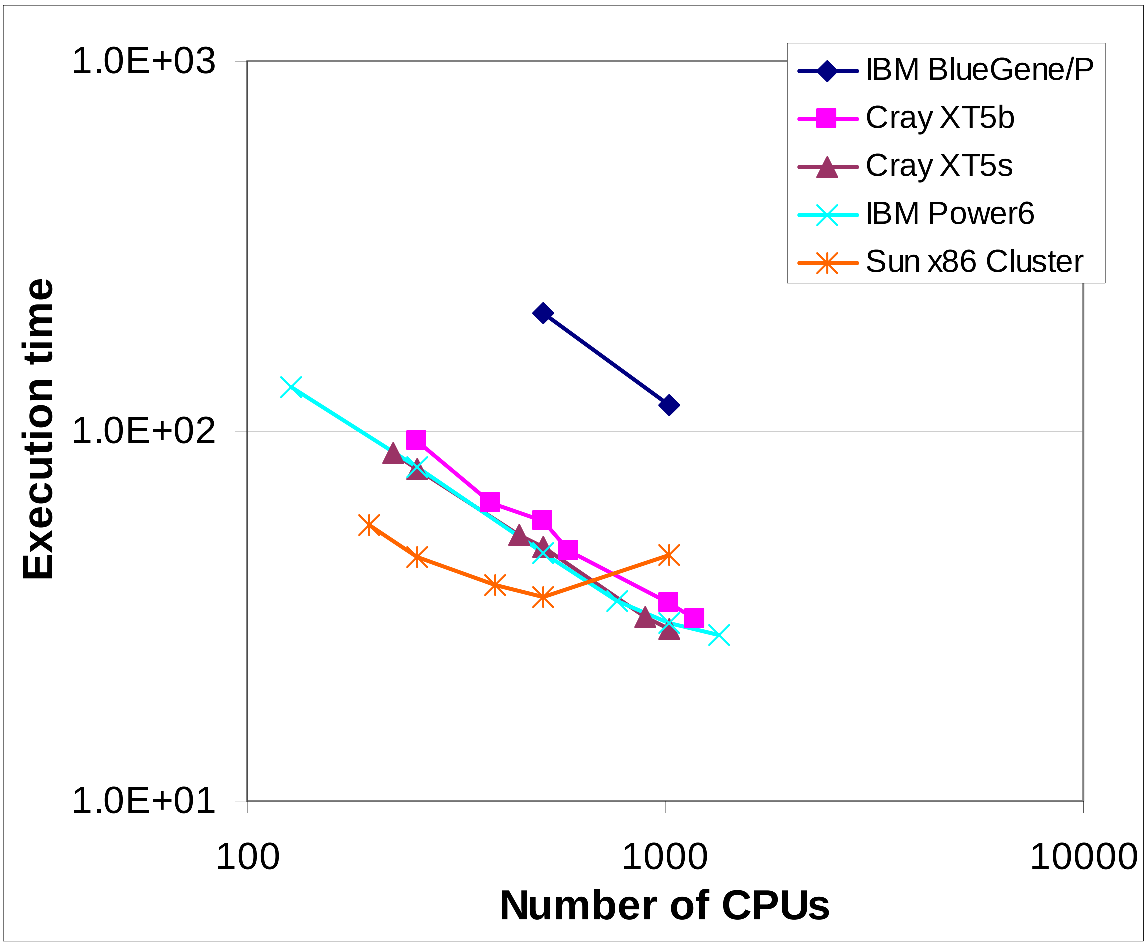

Figure 53: Execution time for GPAW, Test Case A

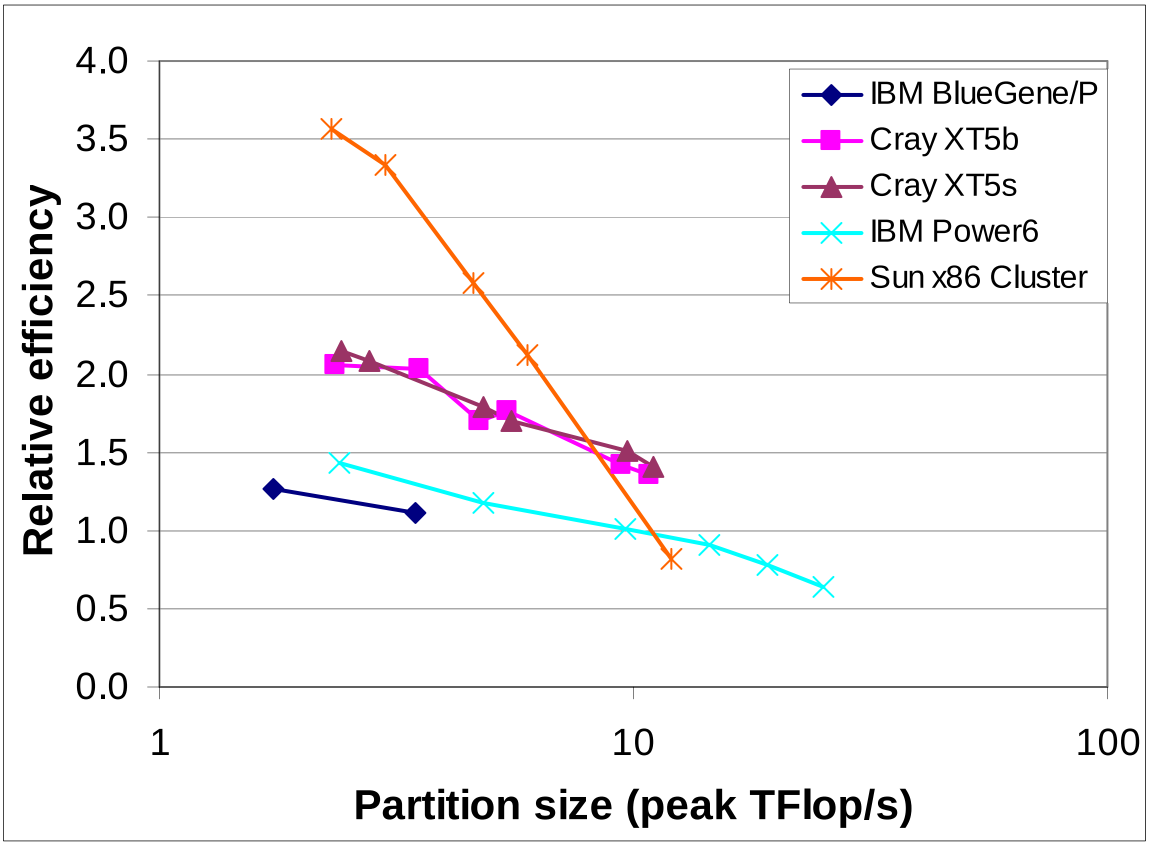

Figure 54: Relative efficiency of GPAW, Test Case A

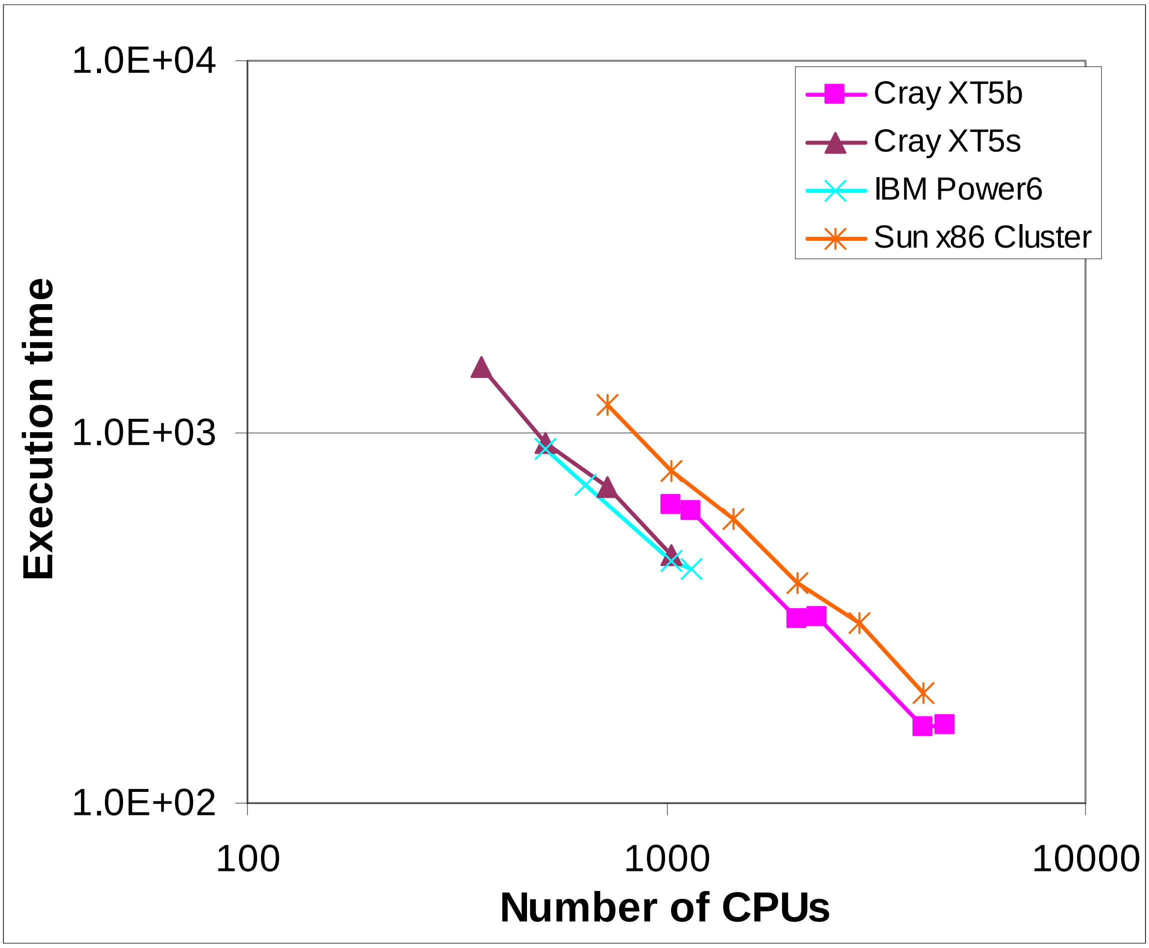

Figure 55: Execution time for GPAW, Test Case B

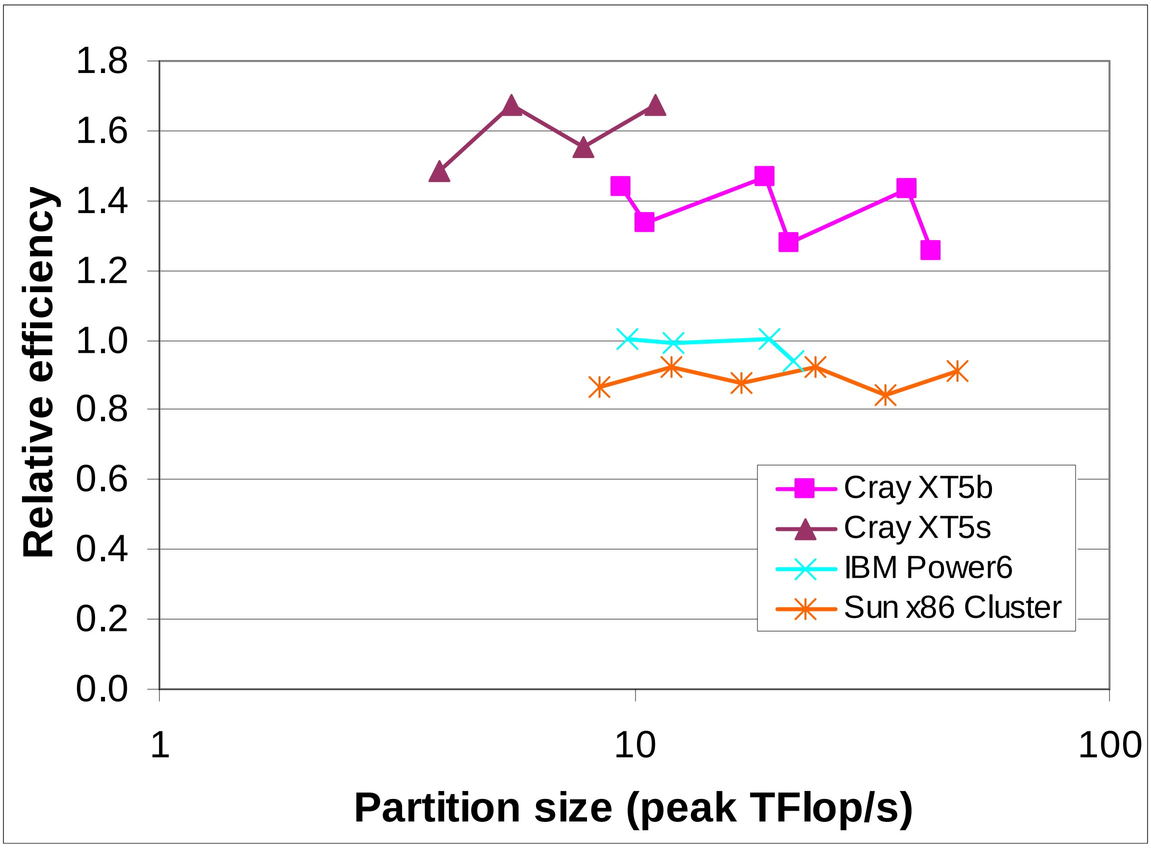

Figure 56: Relative efficiency of GPAW, Test Case B

3.17.4 Analysis

Test Case A does not scale well on any of the systems tested, but the Sun x86 cluster is especially poor in this respect. For Test Case B, scalability is good on all the systems tested, though the problem is clearly better load balanced on some core counts than on others. The Sun x86 cluster is slightly less efficient than the IBM Power6, while the Cray XT5 is significantly more efficient than the other two systems.

3.18 HELIUM

3.18.1 Summary

HELIUM simulates the behaviour of helium atoms using time-dependent solutions of the full-dimensional Schrödinger equation.

3.18.2 Test cases

Test Case A is a simulation using 1540 x 1540 blocks and L_{max}=16.

3.18.3 Results

Figure 57 and Figure 58 show the results for Test Case A

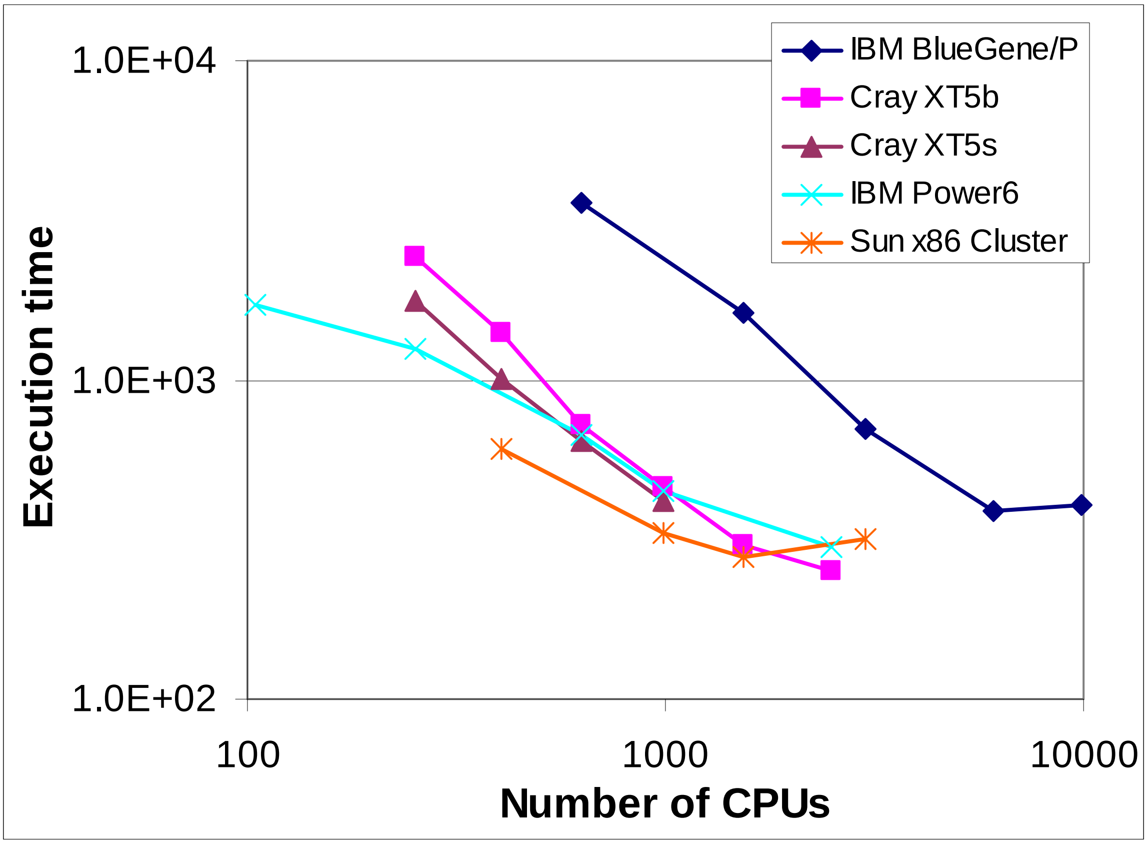

Figure 57: Execution time of HELIUM, Test Case A

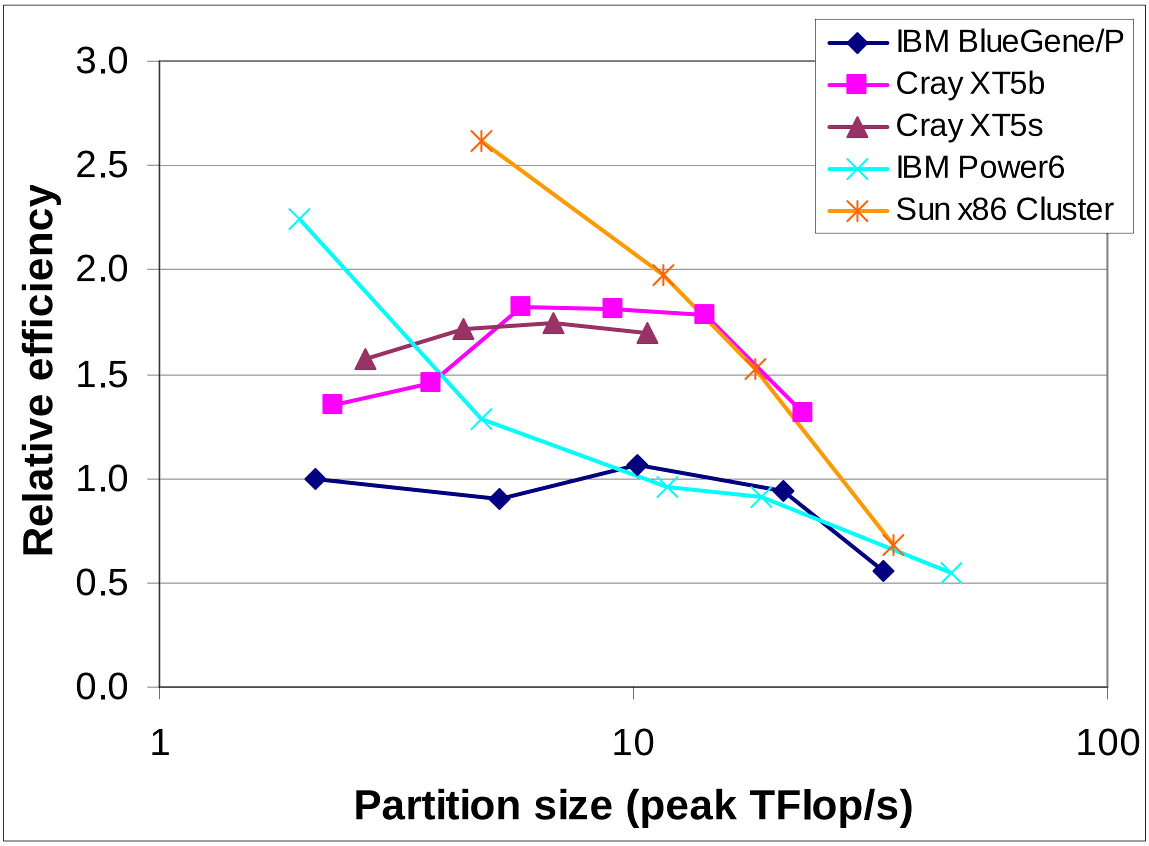

Figure 58: Relative efficiency of HELIUM, Test Case A

3.18.4 Analysis

On this test case, the IBM BlueGene/P scales well, but has low relative efficiency. The Cray XT5 also scales well, and is between 1.5 and 1.7 times as efficient as the IBM BlueGene/P. The IBM Power6 and Sun x86 cluster show poorer scalability.

3.19 OCTOPUS

3.19.1 Summary

Octopus simulates complex electronic processes in medium to large systems, using Density-Functional Theory (DFT) and in particular its time-dependent formulation.

3.19.2 Test cases

Test Case A is a simulation of C_{240}.

Test Case B is a simulation of chlorophyll.

3.19.3 Results

Figure 59 and Figure 60 show the results for Test Case A. Figure 61 and Figure 62 show the results for Test Case B.

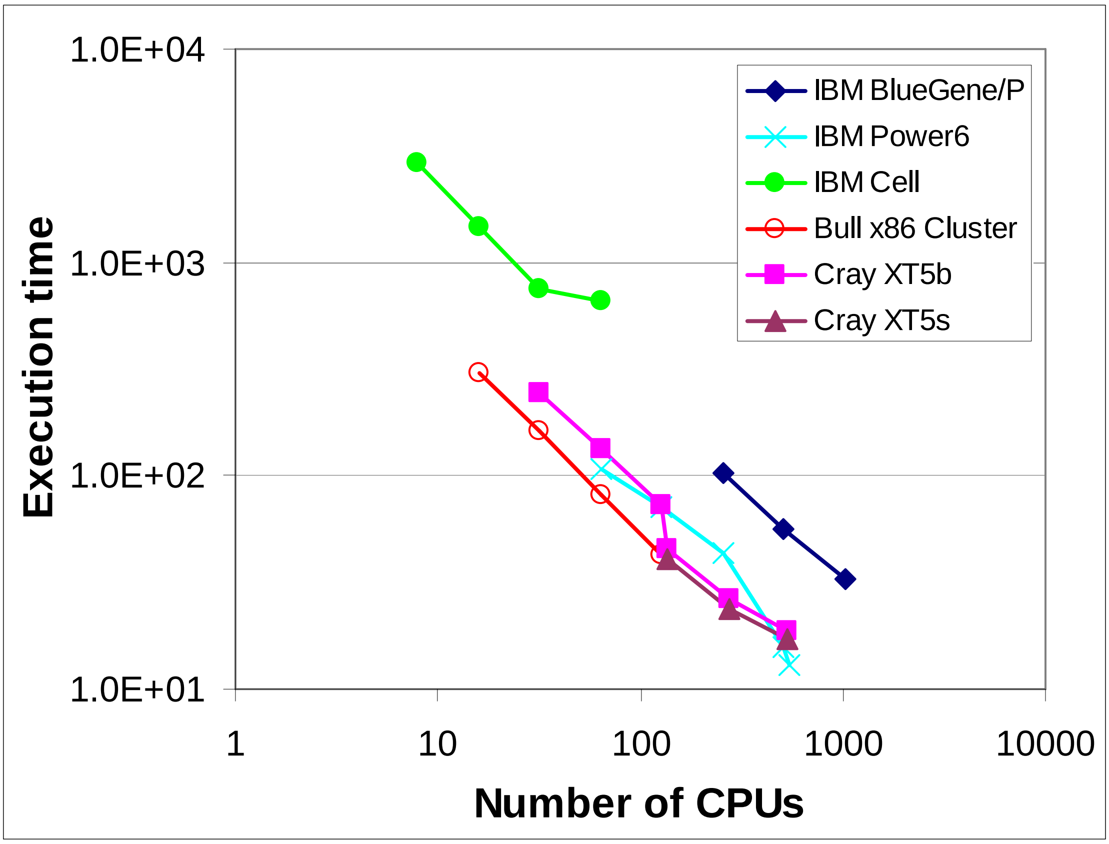

Figure 59: Execution time of OCTOPUS, Test Case A

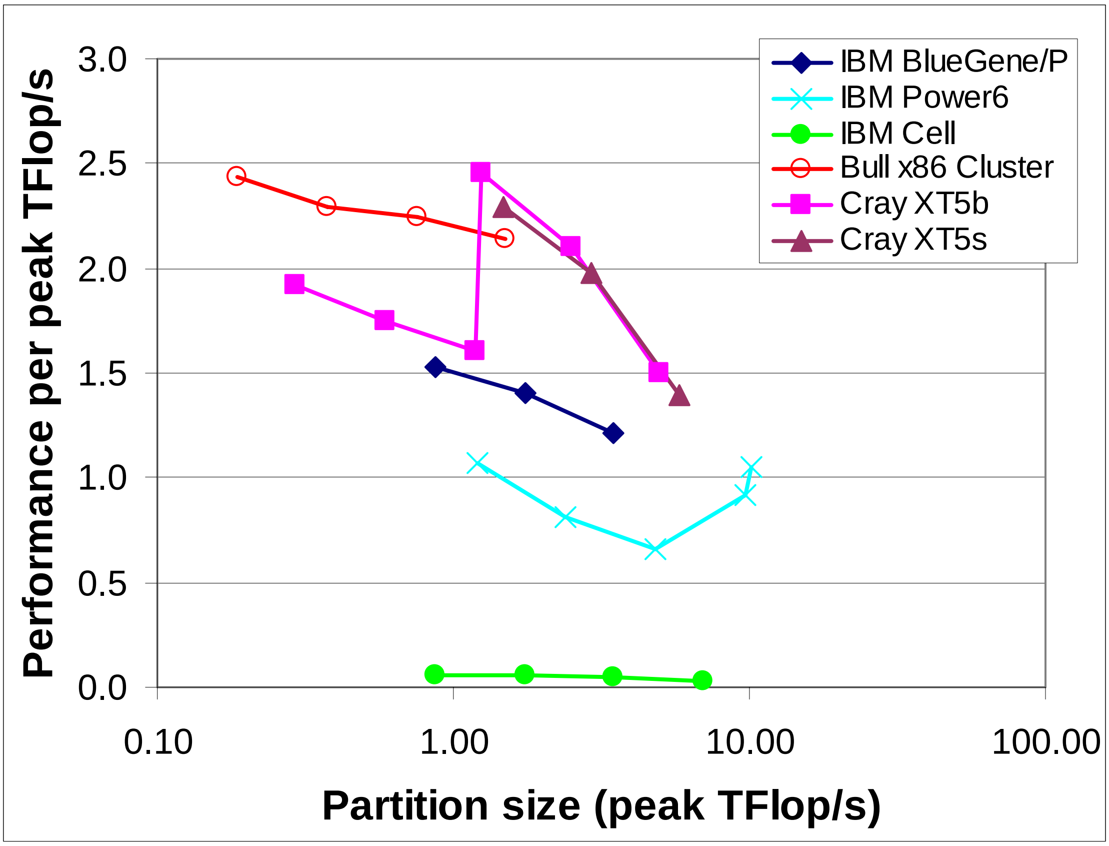

Figure 60: Relative efficiency of OCTOPUS, Test Case A

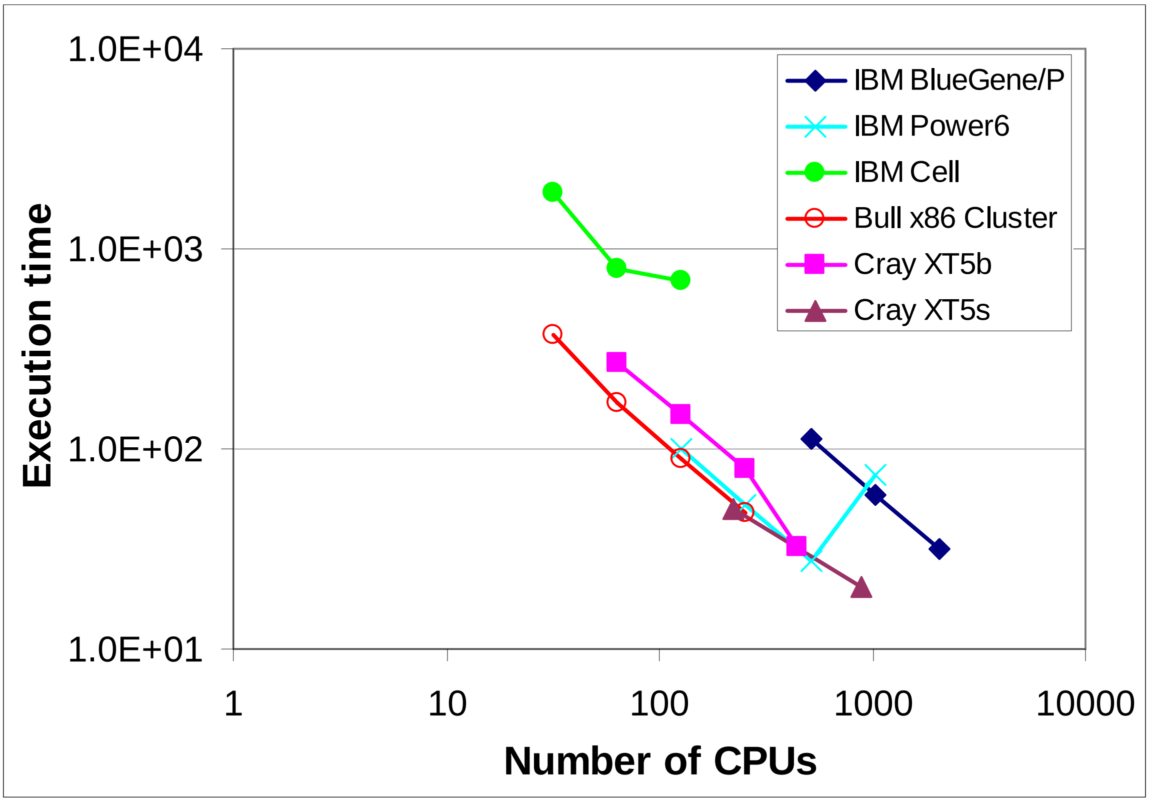

Figure 61: Execution time of OCTOPUS, Test Case B

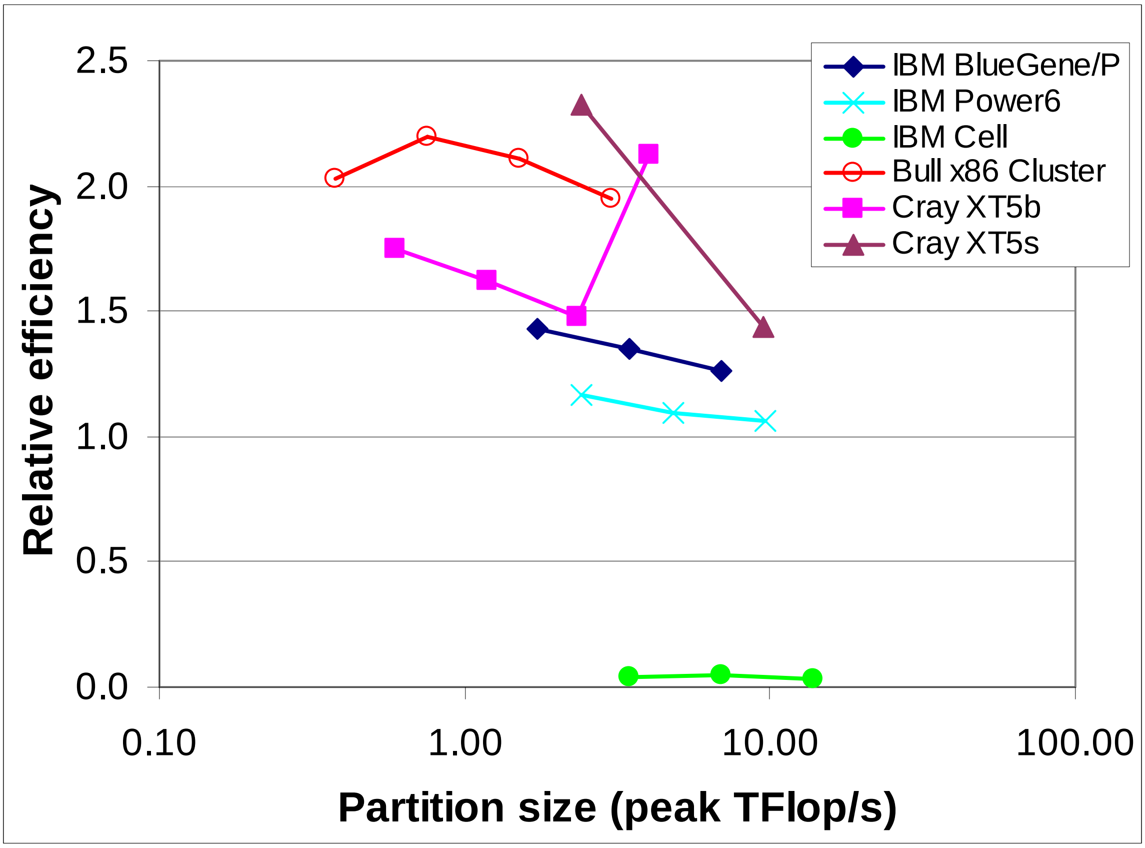

Figure 62: Relative efficiency of OCTOPUS, Test Case B

3.19.4 Analysis

Neither test case shows very good scalability on any of the systems. The relative efficiency of the IBM Cell system is very poor, reflecting the fact that this application makes little if any use of the SPE cores on the Cell processor. Of the other systems, the Bull x86 cluster is the most efficient, followed by the Cray XT5, the IBM BlueGene/P and then the IBM Power6.

3.20 PEPC

3.20.1 Summary

PEPC is a parallel tree-code for computation of long-range Coulomb forces. The forces are calculated based on the Barnes-Hut algorithm. The code takes advantage of multipole-groupings of distant particles to reduce the original O(N2) scaling of the calculation to an O(N log N) scaling.

3.20.2 Test cases

Test Case A consists of a simulation of a 3D-sphere homogeneously filled with 5 million ions, which are initially at rest.

Test Case B consists of a simulation of a 3D-sphere homogeneously filled with 50 million ions, which are initially at rest.

3.20.3 Results

Figure 63 and Figure 64 show the results for Test Case A. Figure 65 and Figure 66 show the results for Test Case B.

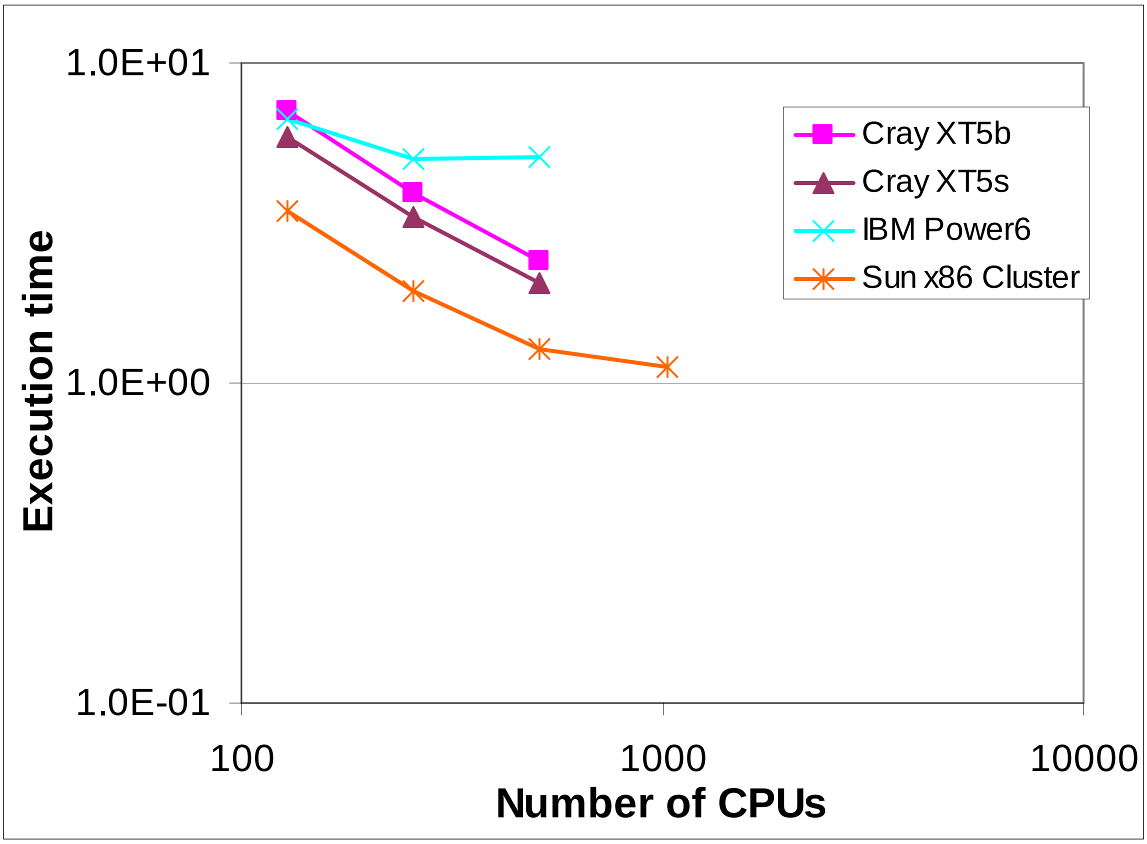

Figure 63: Execution time of PEPC, Test Case A

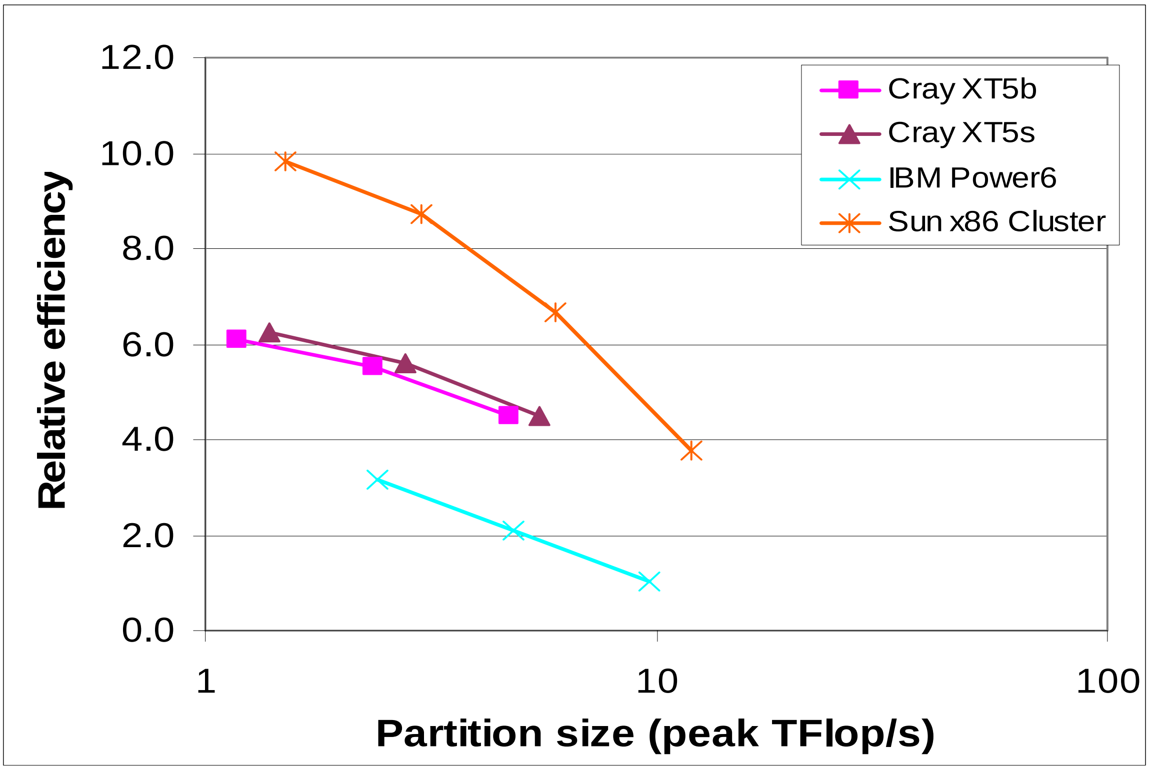

Figure 64: Relative efficiency of PEPC, Test Case A

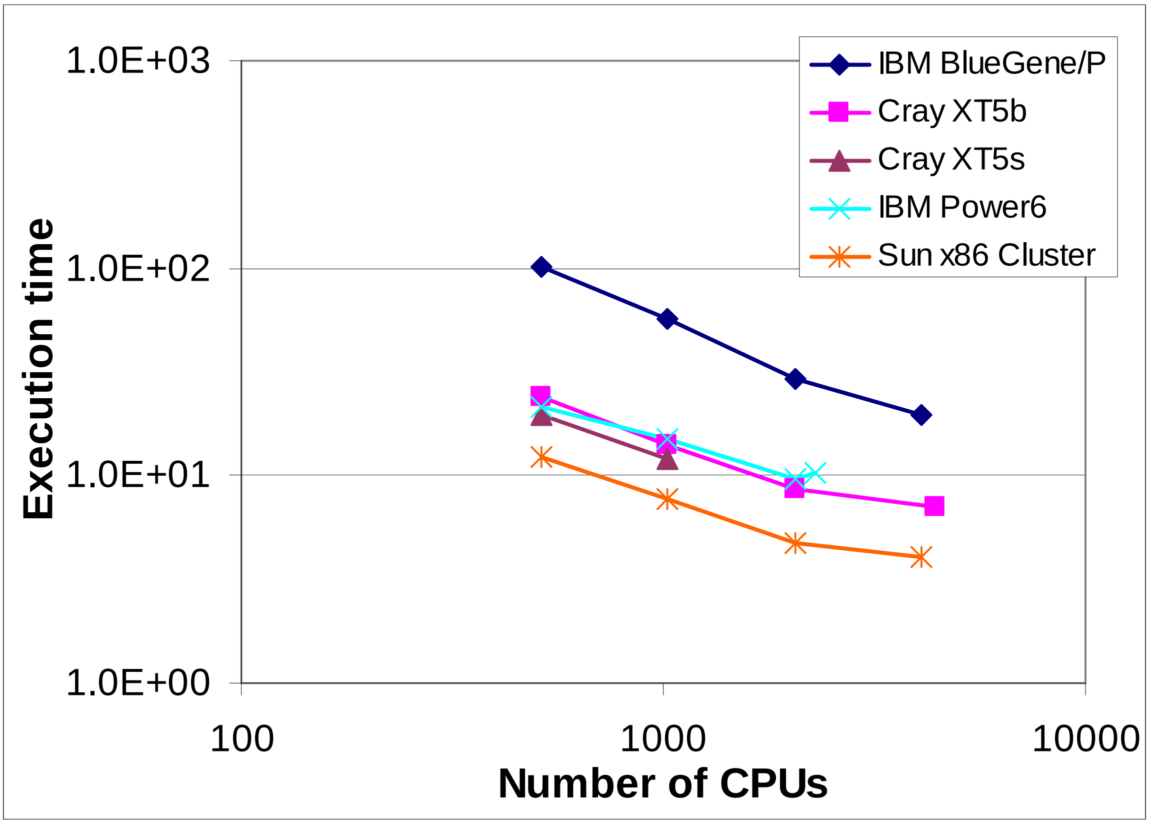

Figure 65: Execution time of PEPC, Test Case B

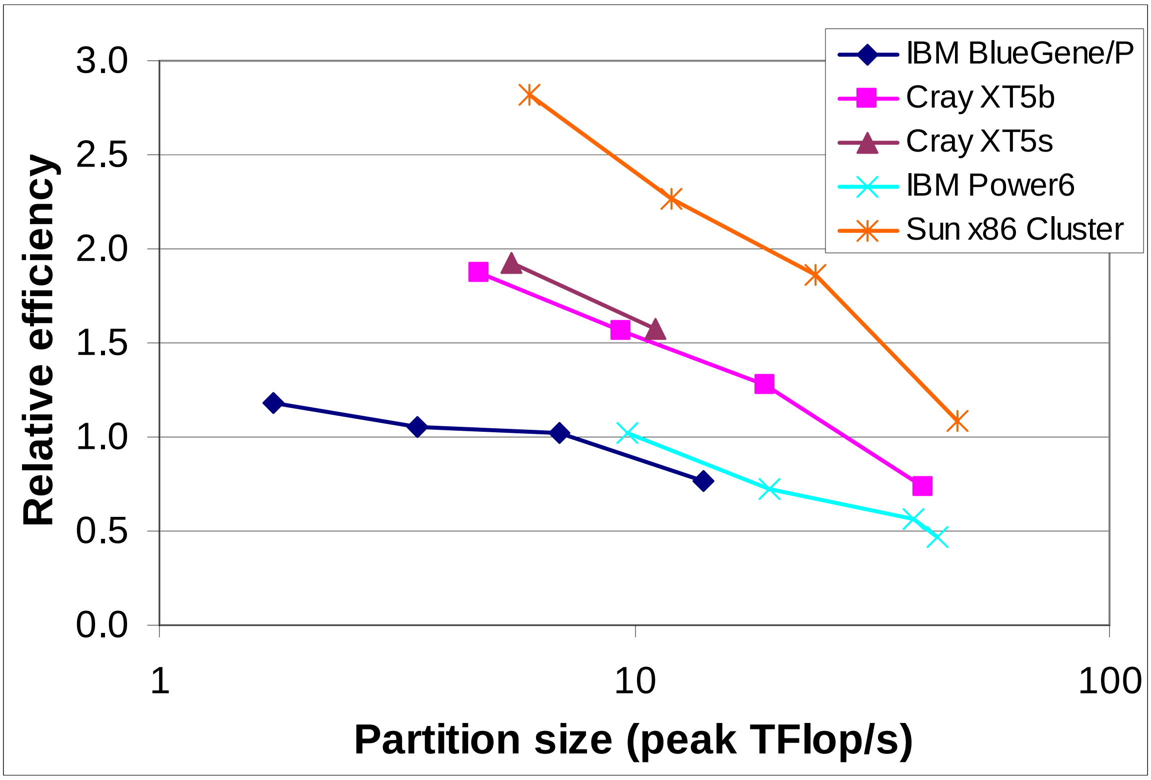

Figure 66: Relative efficiency of PEPC, Test Case B

3.20.4 Analysis

For Test Case B, the IBM Power6 and IBM BlueGene/P show similar relative efficiency. The Cray XT5 and Sun x86 cluster have higher relative efficiency, but do not scale as well. Test Case A shows similar trends in relative efficiency, but does not scale very well on any of the systems.

4 Summary by Prototype

In this Chapter, we summarise the results reported in Chapter 3 by prototype system.

Table 2 shows the relative efficiency of each application of each prototype for a 10 TFlop/s peak partition. These results are obtained by interpolating execution time data as described in Section 2.3. (in a few cases some modest extrapolation has been used). The value of this is exactly 1 on the IBM Power6, by definition. Note that this figure can only be computed where data exists for the given application on a 10 TFlop/s partition size of *both* the prototype in question *and* the IBM Power6.

	IBM Power6	IBM BlueGene/P	Cray XT5b	Cray XT5s	Sun x86 cluster	Bull x86 cluster	IBM Cell	NEC SX-9	NEC x86 cluster
Application									
Code_Saturne	**1.00**	**0.45**	**0.47**		**0.80**				**0.27**
CP2K	**1.00**		**0.64**	**0.55**					
CPMD	**1.00**						**0.61**	**1.92**	**0.83**
EUTERPE	**1.00**	**1.63**	**1.94**						
GADGET	**1.00**								
GROMACS	**1.00**	**0.56**	**2.06**	**2.25**					
NAMD	**1.00**	**0.65**	**0.85**	**0.92**	**0.10**				
NEMO	**1.00**								
NS3D	**1.00**	**1.96**	**0.36**		**1.64**			**5.48**	**2.57**
QCD	**1.00**	**1.01**	**1.38**		**2.16**				**1.83**
Quantum_Espresso	**1.00**		**0.82**		**1.14**				**1.17**
WRF	**1.00**	**1.11**	**1.72**	**1.48**	**2.28**	**1.87**			**1.96**
ALYA	**1.00**								
AVBP	**1.00**				**2.35**				
BSIT	**1.00**								
ELMER	**1.00**		**2.14**		**2.19**				
GPAW	**1.00**		**1.38**	**1.64**	**0.89**				
HELIUM	**1.00**	**1.06**	**1.80**	**1.71**	**2.04**				
OCTOPUS	**1.00**						**0.04**		
PEPC	**1.00**	**0.87**	**1.53**	**1.62**	**2.38**				

Table 2: Relative efficiency of applications on a 10 TFlop/s peak partition.

4.1 IBM Power6

Figure 67 and Figure 68 show the relative efficiency of applications on the IBM Power6 prototype system. For ease of reading, the applications are split into the core list (Figure 67) and the extended list (Figure 68).

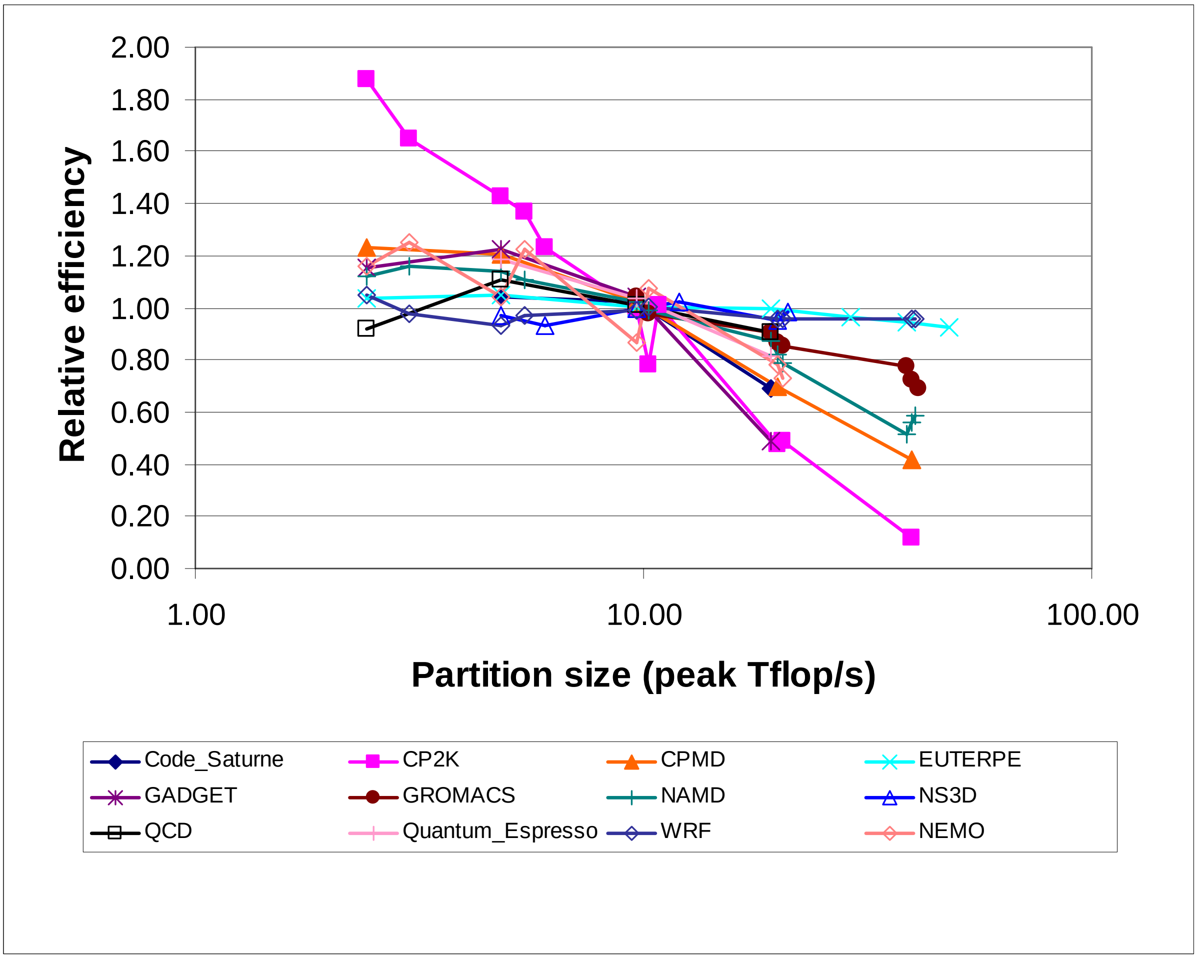

Figure 67: Relative efficiency of applications on IBM Power6 (core list)

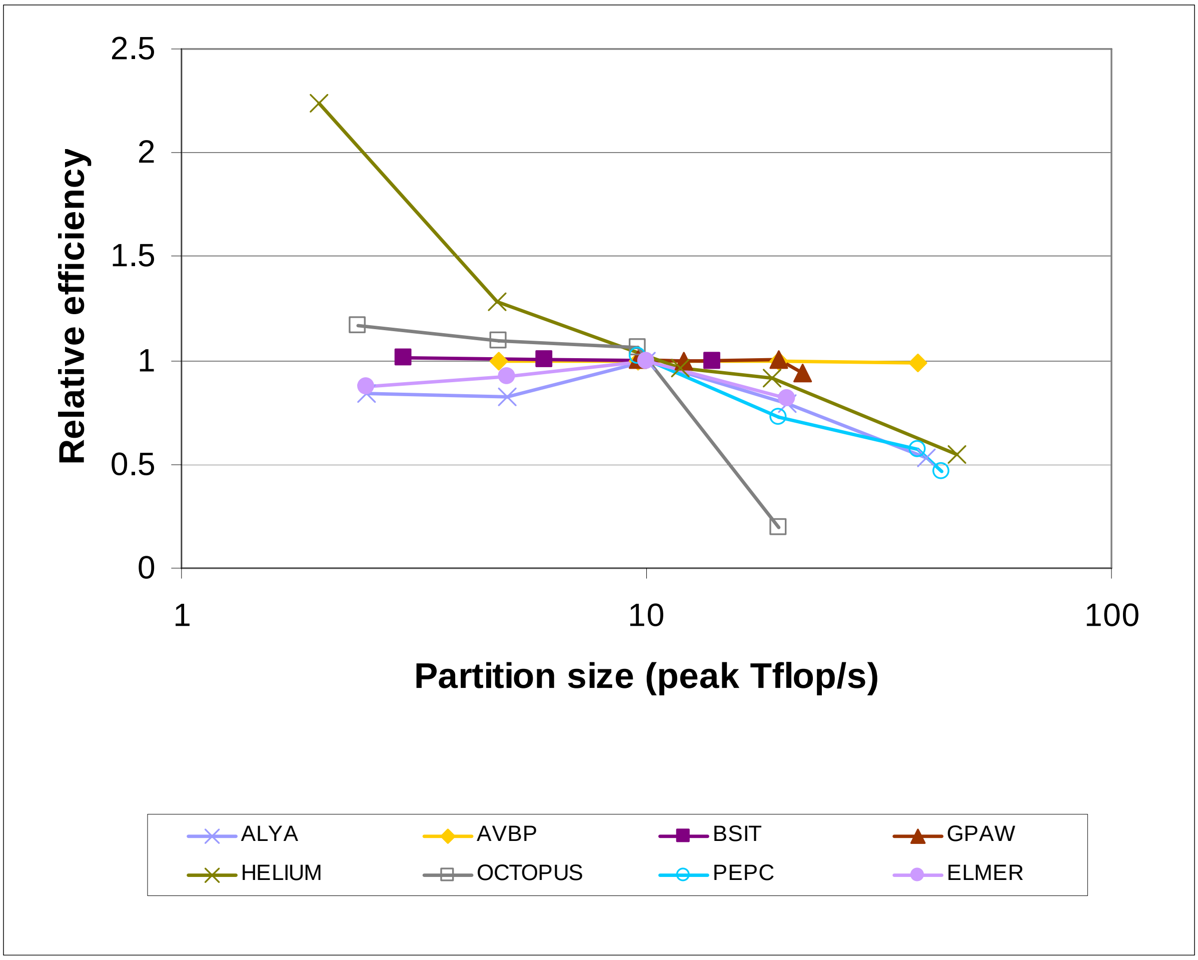

Figure 68: Relative efficiency of applications on IBM Power6 (extended list)

Since the IBM Power6 system is used as the reference system for making comparisons between prototypes, these figures serve only to illustrate the scalability of applications on the IBM Power6 system. Scalability is in general good, at least up to a partition size of 20 TFlop/s.

4.2 Cray XT5

Figure 69 and Figure 70 show the relative efficiency of applications on the Cray XT5, with Barcelona and Shanghai cores, respectively.

Compared to the IBM POWER6 system, the Cray XT5 with Barcelona cores shows broadly similar performance across the range of applications. At 10 TFlop/s peak, the ratio of performance to the IBM POWER6 lies typically between 0.5 and 2.0, with roughly half the applications performing better on the Cray XT5 and half performing worse. The two CFD applications (NS3D and Code_Saturne) show the poorest performance compared to the IBM POWER6.

The Cray XT5 also shows similar scaling behaviour to the IBM POWER6 as the partition size is increased above 10 TFlop/s. As would be expected, most applications perform better on the Cray XT5 with Barcelona cores than on the same number of Shanghai cores. When we compare partitions with the same peak performance, there is little difference between the two Cray systems.

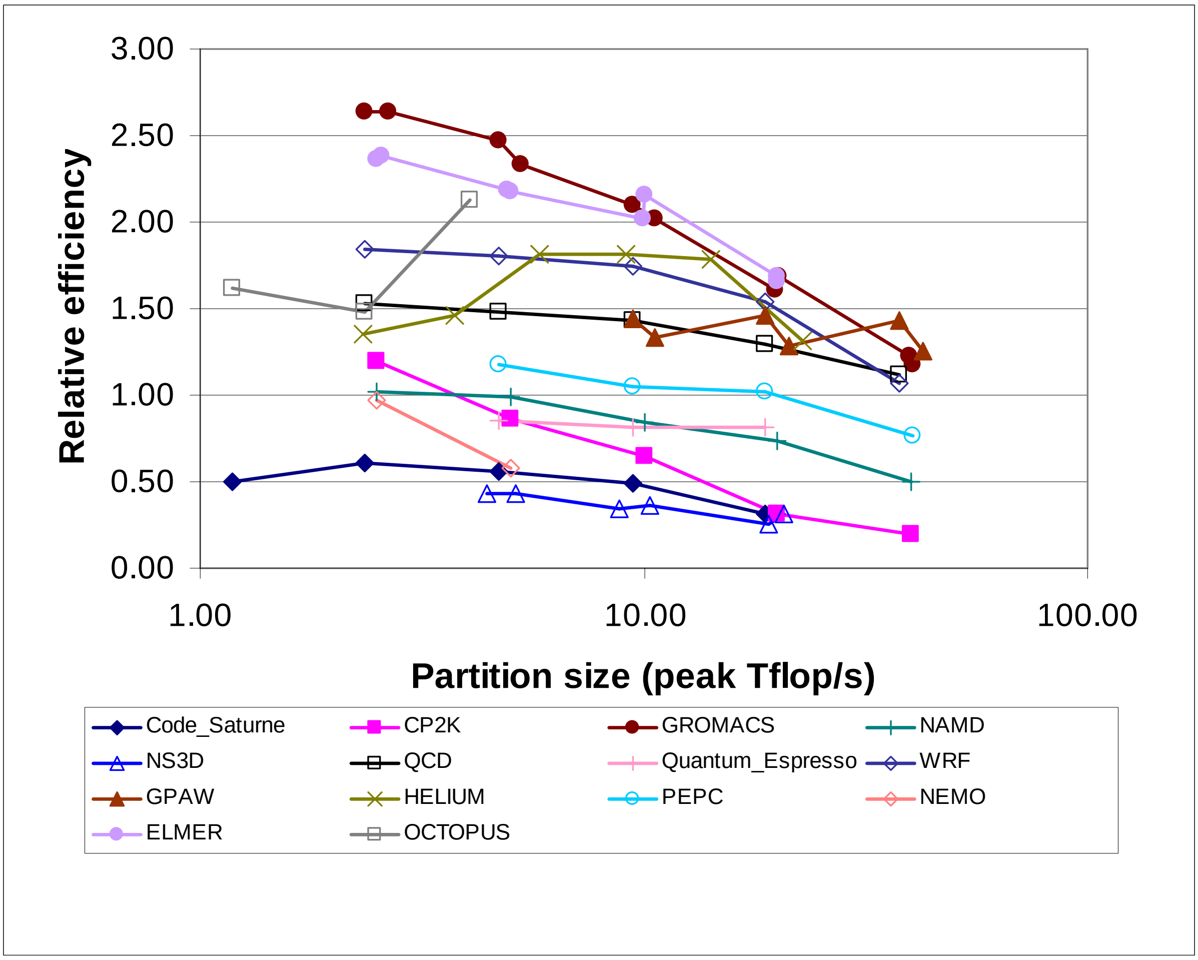

Figure 69: Relative efficiency of applications on Cray XT5, Barcelona cores

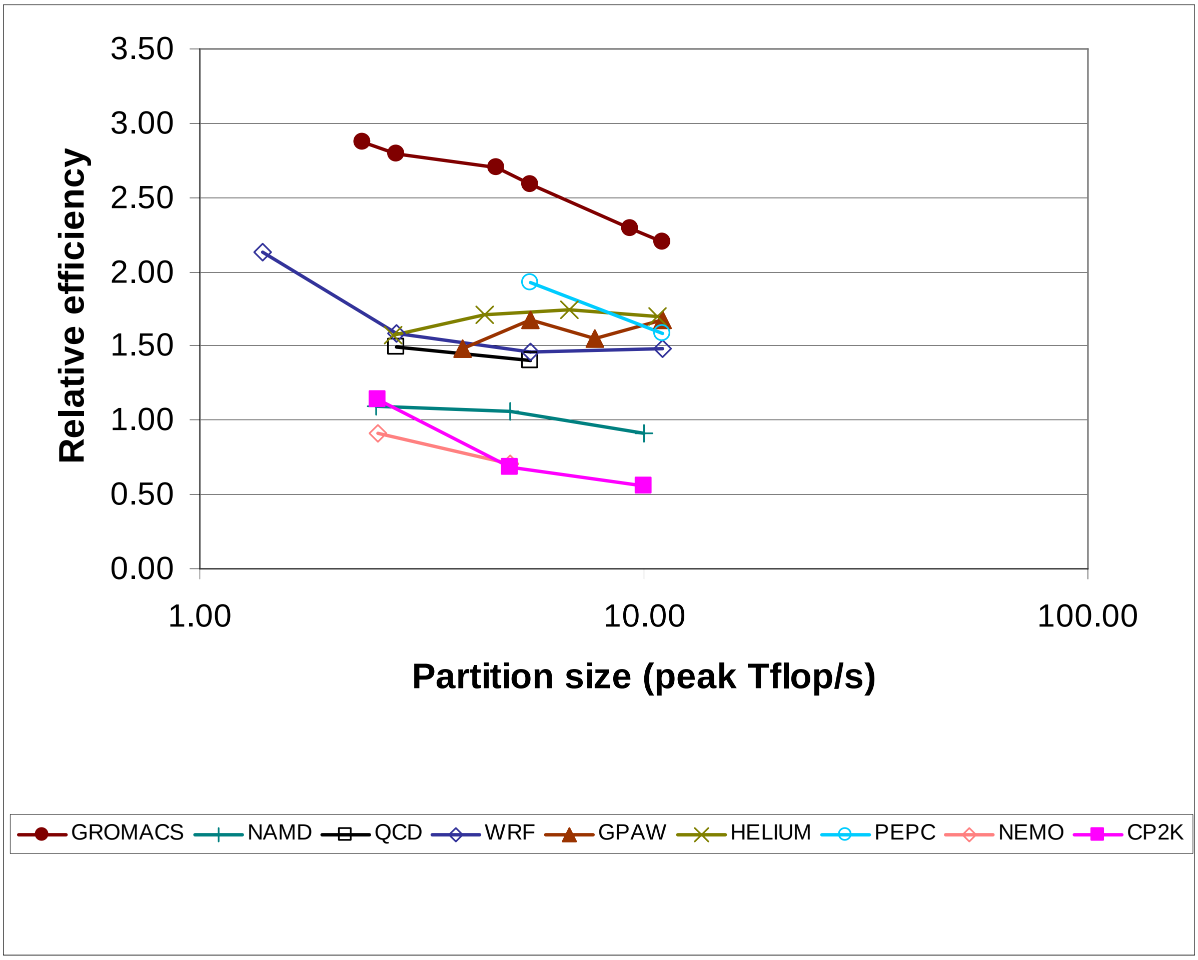

Figure 70: Relative efficiency of applications on Cray XT5, Shanghai cores

4.3 IBM BlueGene/P

Figure 71 shows the relative efficiency of applications on the IBM BlueGene/P.

Compared to the IBM POWER6 system, the IBM BlueGene/P also shows broadly similar performance across the range of applications at modest partition sizes, but scales better to larger partition sizes. At 10 TFlop/s peak, the ratio of performance to the IBM POWER6 again lies typically between 0.5 and 2.0, with roughly half the applications performing better on the BlueGene/P and half performing worse. The scalability of the BlueGene/P is markedly better than the IBM POWER6 system. Not only does the BlueGene/P show comparable performance to the POWER6 at the same peak performance (which the BlueGene/P requires six times as many cores to achieve) but the scalability is better than the POWER6 as the partition size is increased above 10 TFlop/s. Although the peak performance of the individual cores in the BlueGene/P is low, this is more than compensated for by the quality of the memory subsystem and of the interconnect.

A small number of application datasets do perform poorly on the BlueGene/P, and some cannot be run at all: this is due to the restricted amount of memory per core. In such cases, to allow the application to run at all, it is sometimes necessary to run fewer MPI tasks than cores on each node.

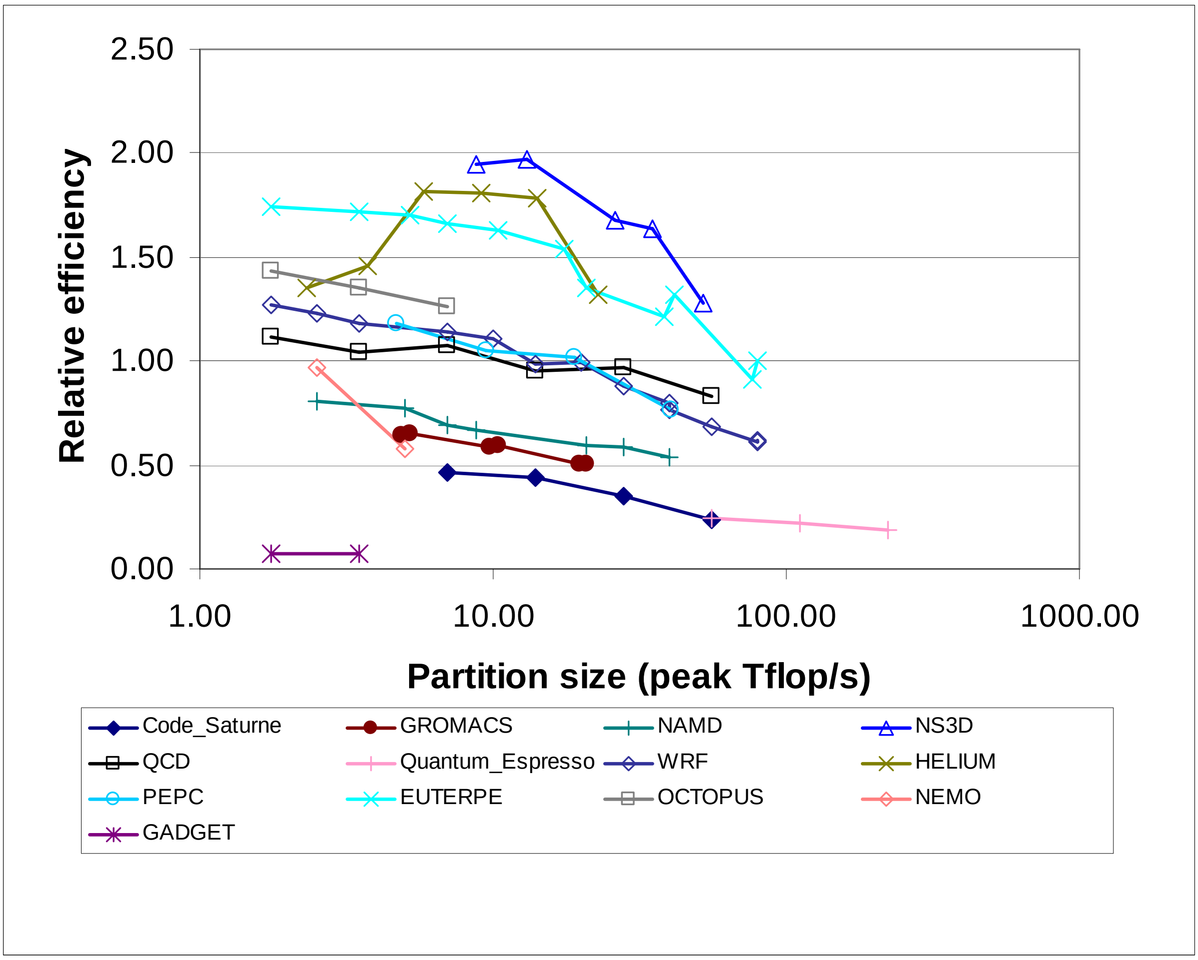

Figure 71: Relative efficiency of applications on IBM BlueGene/P

4.4 Sun x86 cluster

Figure 72 shows the relative efficiency of applications on the Sun x86 cluster.

At modest partition sizes, the Sun x86 cluster shows significantly better performance per peak TFlop/s that the IBM POWER6, but the scalability to large partition sizes is very poor. At 10 TFlop/s peak, the ratio of performance to the IBM POWER6 lies typically between 1.0 and 2.5. The large cache sizes and high memory bandwidth of the Intel Nehalem cores are the likely reason why applications can achieve a higher percentage of peak performance on this system than on the POWER6, Cray XT5 or BlueGene/P. However, as the partition size is increased above about 10 TFlop/s (≈ 850 cores), the relative performance declines steeply. This is likely to be related to either the Infiniband interconnect, or the MPI implementation, but the precise reasons are still not fully understood.

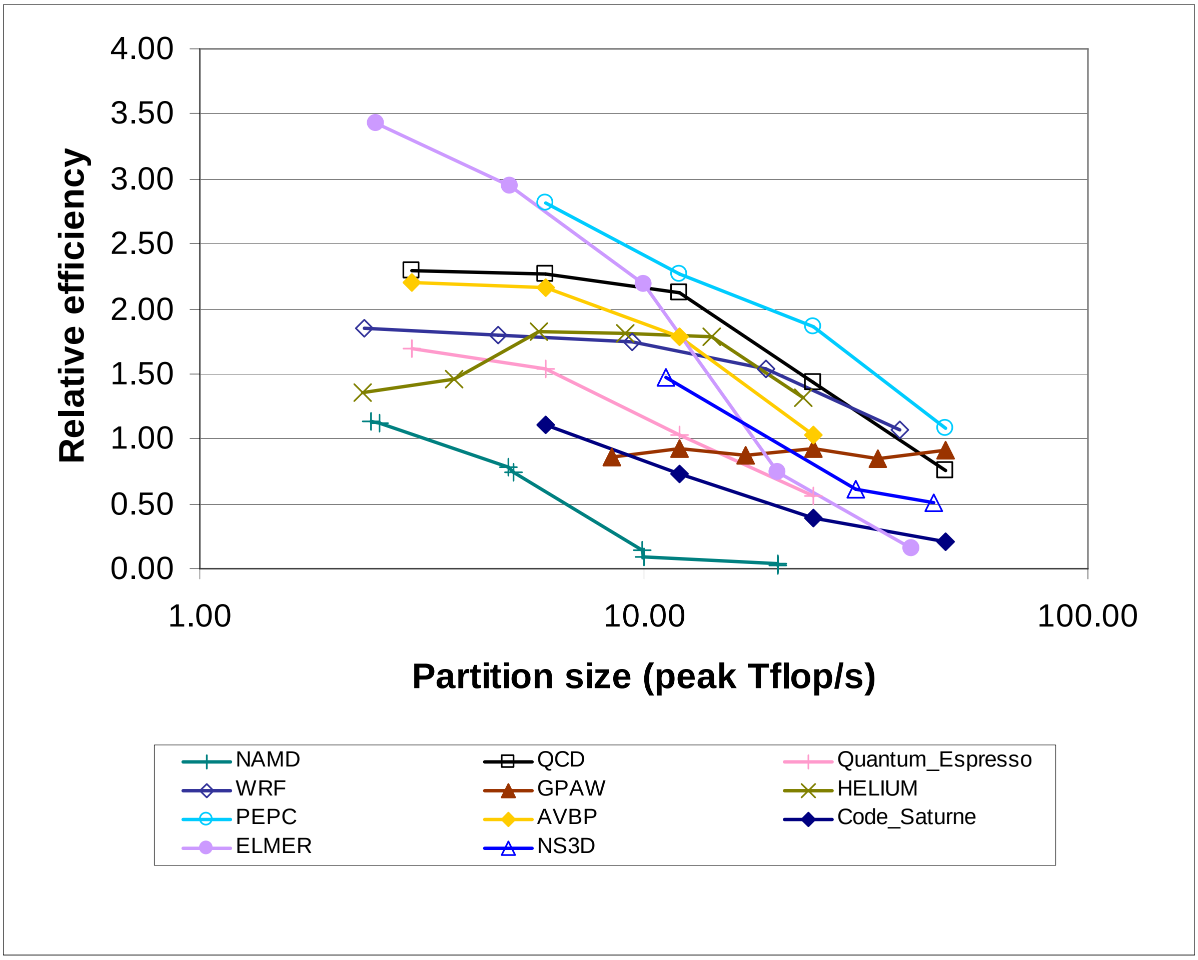

Figure 72: Relative efficiency of applications on Sun x86 cluster

4.5 Bull x86 cluster

Figure 73 shows the relative efficiency of applications on the Bull x86 cluster.

Only a small subset of the application benchmarks has been run on the Bull x86 cluster. This prototype system is similar to the Sun x86 cluster (both use Intel Nehalem processors and an Infiniband interconnect), but is smaller in scale. Applications which were run on both x86 cluster show similar performance, but due to the size of the Bull system the scalability behaviour above 10 TFlop/s peak could not be observed.

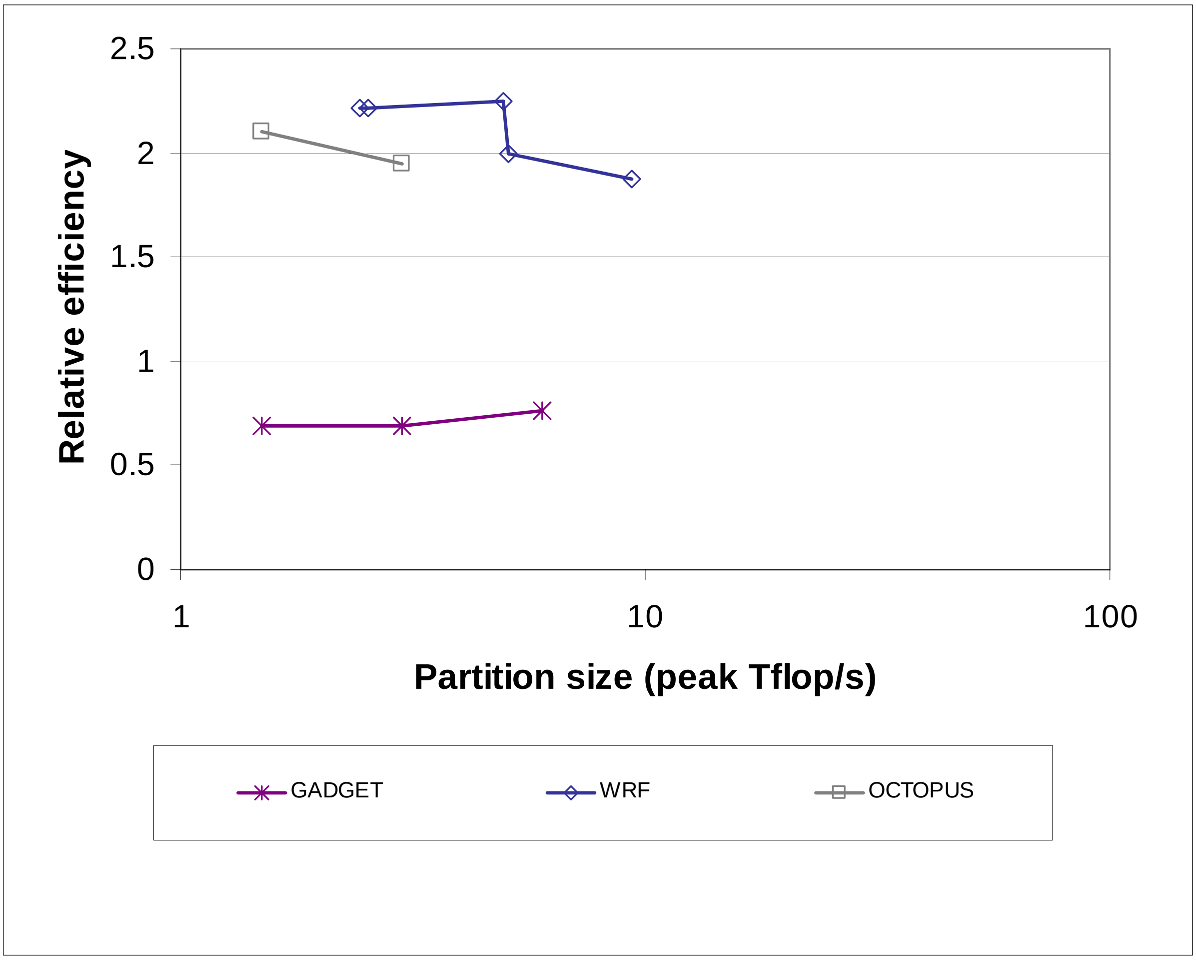

Figure 73: Relative efficiency of applications on Bull x86 cluster

4.6 IBM Cell

Figure 74 shows the relative efficiency of applications on the IBM Cell.

Only a small subset of the application benchmarks has been run on the IBM Cell system. This reflects the difficulty of porting applications to this prototype. Compared to the IBM POWER6, the performance of applications varies widely. On application (BSIT) which has been fully ported to the SPE cores, runs about six times faster per peak TFlop/s on the Cell, while others, which use only the PPE core, run more than 20 times slower. For certain applications, and given sufficient programmer effort, this system can be very effective, but for other applications, results are very poor. This is a modest scale system, so scalability above 10 TFlop/s peak could not be assessed.

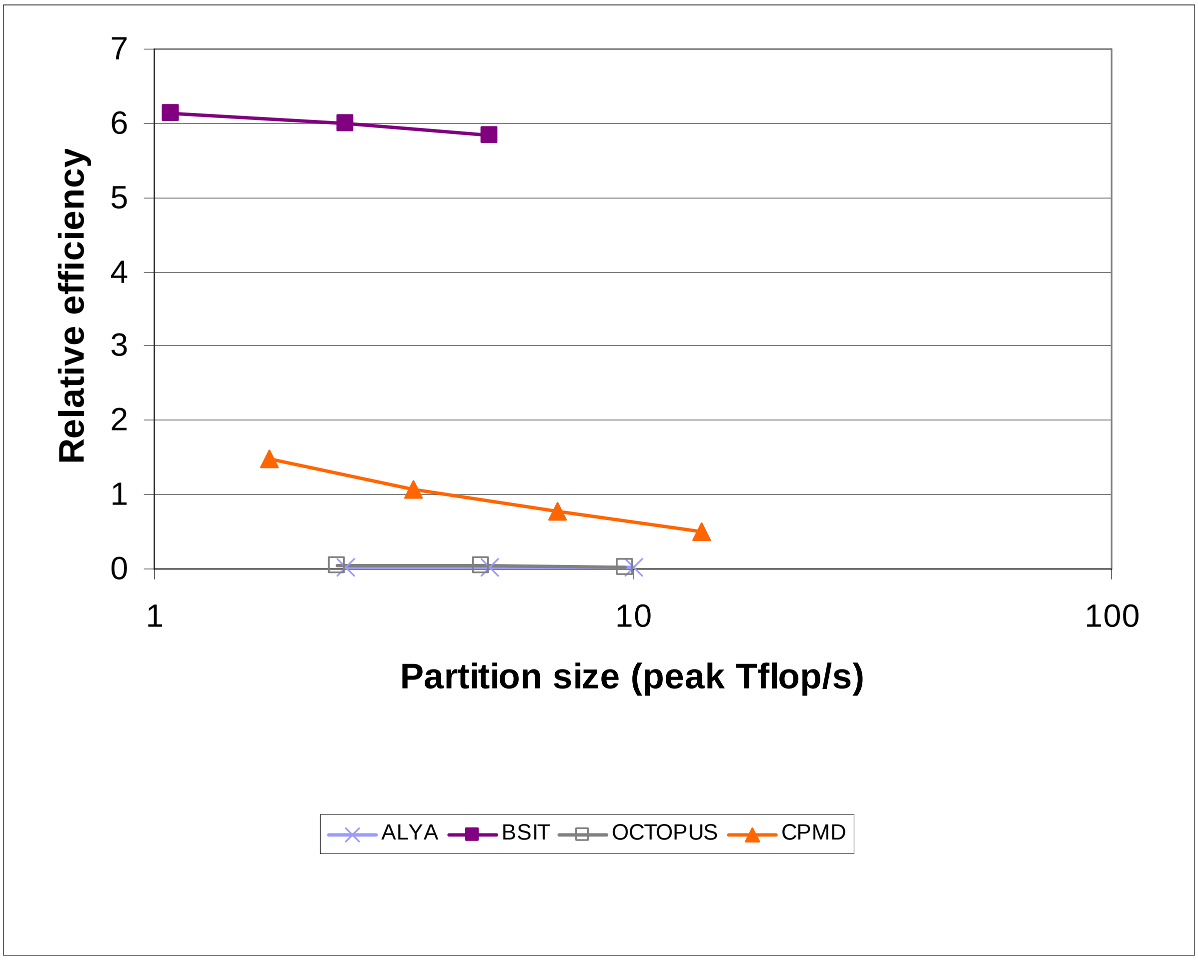

Figure 74: Relative efficiency of applications on IBM Cell

4.7 NEC SX-9

Figure 75 shows the relative efficiency of applications on the NEC SX-9.

Only a small subset of the application benchmarks has been run on the NEC SX-9. One application (NS3D) runs over five times faster per peak TFlop/s than on the IBM POWER6. NEMO also shows a high relative efficiency on the small partition size which was tested (one node). Although the porting effort for this system is not as great as for the IBM Cell, it is clear that straightforward porting to this system does not automatically result in significant performance gains over more conventional architectures with scalar processors. This is also a modest scale system, so scalability above 10 TFlop/s peak could not be assessed.

Figure 75: Relative efficiency of applications on NEC SX-9

4.8 NEC x86 cluster

Figure 76 shows the relative efficiency of applications on the NEC x86 cluster.

Despite the similarities in the hardware, there are significant differences between the NEC x86 cluster and the Sun x86. For partition sizes smaller than 10 Tflop/s, the NEC system typically has somewhat lower relative efficiency than the Sun system. This may be due to the fact that the two systems contain different variants of the Nehalem processor, or else the differences may be due to different software environments. Above 10 Tflop/s, the NEC system, however, does not appear to exhibit similar scalability problems to the Sun x86 cluster. This conclusion should be treated with caution, however, as few applications have been run on more than about one third of the full system size.

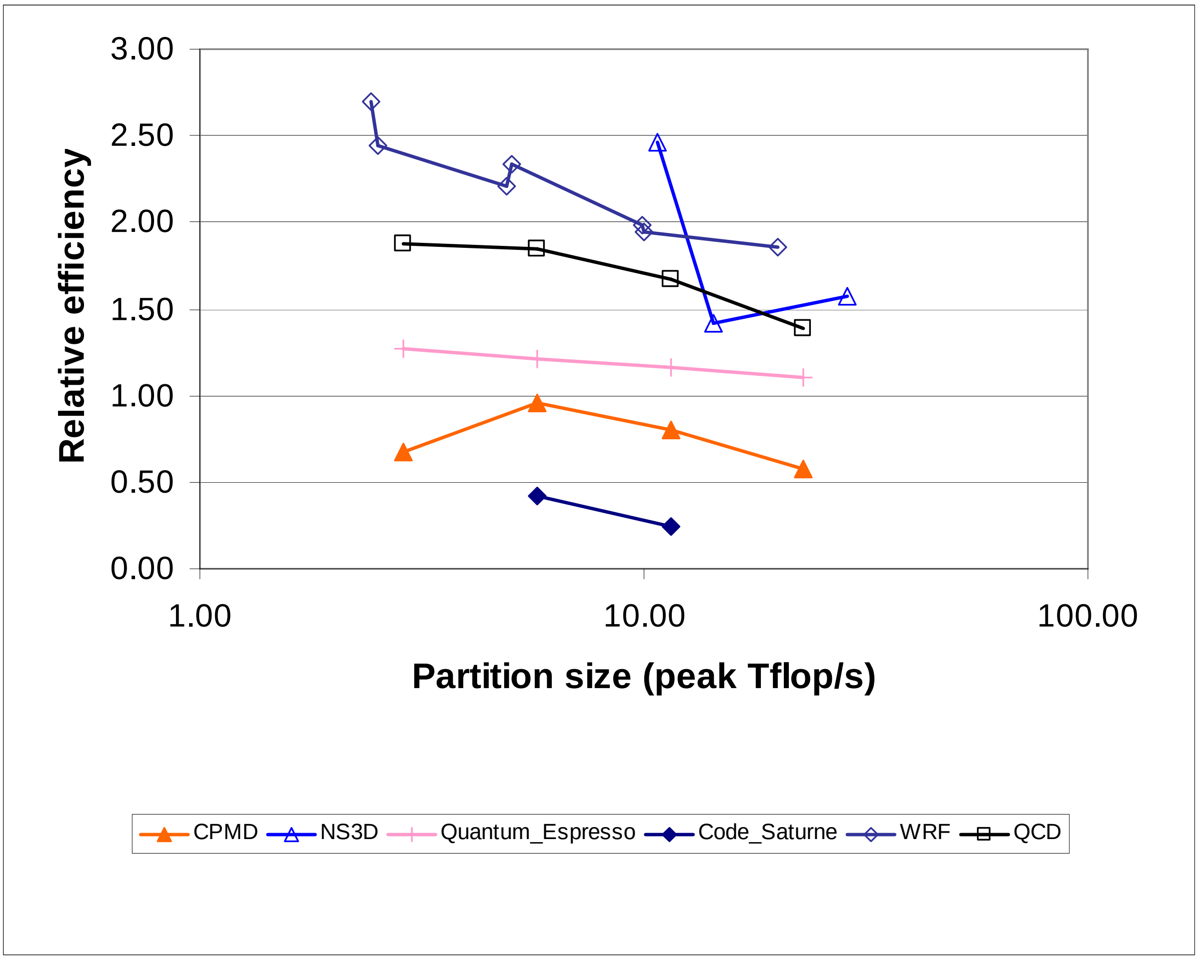

Figure 76: Relative efficiency of applications on NEC x86 cluster

5 Conclusions and Future Work

This document reports the results of running the PRACE application benchmark suite on the installed prototype systems. In presenting these results we have used both execution time and relative efficiency as the principal metrics. The advantage of the latter is that it allows meaningful comparisons to be made between systems containing processing elements of widely differing computational power.

Due to constraints on the available effort, and in part to the unavailability of some prototypes until quite late in the project, it has not been possible to run all the applications on all the prototypes. Porting and optimising applications for some of the prototypes (such as the IBM Cell and NEC SX-9) is not a trivial task.

A useful extension to this work would be to use hardware counters to measure the number of flops executed by the application benchmarks on a reference system (e.g. the IBM Power6). This would allow us to compute an absolute measure of efficiency (the percentage of peak floating point performance attained) as well as the relative efficiency reported here.

6 Appendix

This Chapter contains the raw data used in the figures in Chapters 3 and 4. For definitions of quantities, see Section 2.3.

6.1 CODE_SATURNE

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	8.33	0.87
1024	3.48	4.59	0.79
2048	6.96	2.39	0.76
4096	13.93	1.50	0.61

Table 3: Results for Code_Saturne, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
64	0.59	23.32	0.92
128	1.18	16.96	0.63
256	2.36	6.86	0.78
512	4.71	4.12	0.65
1024	9.42	3.41	0.39
2048	18.84	2.56	0.26

Table 4: Results for Code_Saturne, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
32	0.60	39.49	0.53
64	1.20	15.57	0.68
128	2.41	4.81	1.09
256	4.81	2.34	1.12
512	9.63	1.31	1.00

Table 5: Results for Code_Saturne, Test Case A on IBM Power6.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	2.39	1.50	1.90	0.67
Cray XT5b	1086.96	1024	2048	3.41	2.56	3.34	0.38
IBM Power6	531.91	256	512	2.34	1.31	1.27	1.00

Table 6: Results for 10 TFlop/s partitions: Code_Saturne, Test Case A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
2048	6.96	39.58	0.46
4096	13.93	21.01	0.44
8192	27.85	13.24	0.35
16384	55.71	9.7	0.24

Table 7: Results for Code_Saturne, Test Case B on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.18	215.17	0.50
256	2.36	89.08	0.61
512	4.71	48.43	0.56
1024	9.42	27.89	0.49
2048	18.84	21.86	0.31

Table 8: Results for Code_Saturne, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	4.81	25.39	1.05
512	9.63	12.94	1.03
1024	19.25	9.64	0.69

Table 9: Results for Code_Saturne, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	6.00	19.3	1.10
1024	12.00	14.63	0.73
2048	24.00	13.62	0.39
4096	48.01	12.42	0.21

Table 10: Results for Code_Saturne, Test Case B on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	62.41	0.71
512	5.73	53.04	0.42
1024	11.47	44.9	0.25

Table 11: Results for Code_Saturne, Test Case B on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512	1024	12.94	9.64	12.77	1.00
IBM BlueGene/P	2941.18	2048	4096	39.58	21.01	28.57	0.45
Cray XT5b	1086.96	1024	2048	27.89	21.86	27.42	0.47
Sun x86 Cluster	853.24	512	1024	19.3	14.63	15.91	0.80
NEC x86 cluster	892.86	512	1024	53.04	44.90	46.74	0.27

Table 12: Results for 10 TFlop/s partitions: Code_Saturne, Test Case B.

6.2 CP2K

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
272	2.50	4164.0	1.20
543	5.00	2902.7	0.86
1087	10.00	1943.2	0.64
2174	20.00	1964.0	0.32
4348	40.00	1635.5	0.19

Table 13: Results for CP2K, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
232	2.51	4369	1.14
464	5.01	3637	0.69
928	10.02	2257	0.55

Table 14: Results for CP2K, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	2767.8	1.88
160	3.01	2520.9	1.65
256	4.81	1820.5	1.43
288	5.41	1686.7	1.37
320	6.02	1686.2	1.23
512	9.63	1304.5	1.00
544	10.23	1561.8	0.78
576	10.83	1147.1	1.01
1056	19.85	1322.6	0.48
1088	20.45	1252.6	0.49
2112	39.71	2750.9	0.11

Table 15: Results for CP2K, Test Case A on IBM Power6.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
Cray XT5b	1086.96	1087	1087	1943.23	1943.23	1943.23	0.64
Cray XT5s	925.93	464	928	3637.00	2257.00	2260.83	0.55
IBM Power6	531.91	512	576	1304.47	1147.10	1251.06	1.00

Table 16: Results for 10 TFlop/s partitions: CP2K, Test Case A.

6.3 CPMD

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	192.05	1.23
256	4.81	97.83	1.21
512	9.63	57.63	1.02
1056	19.85	41.23	0.69
2112	39.71	34.23	0.42

Table 17: Results for CPMD, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.64	147.15	2.36
32	3.28	79.18	2.19
48	4.92	56.25	2.05
64	6.55	42.78	2.03
80	8.19	36.59	1.90
96	9.83	30.09	1.92

Table 18: Results for CPMD, Test Case A on NEC SX-9.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.74	219.92	1.48
32	3.48	152.77	1.07
64	6.96	103.90	0.79
128	13.93	81.08	0.50

Table 19: Results for CPMD, Test Case A on IBM Cell.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	293.69	0.67
512	5.73	103.04	0.96
1024	11.47	61.59	0.80
2048	22.94	43.18	0.57

Table 20: Results for CPMD, Test Case A on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512	1056	57.63	41.23	56.80	1.00
NEC SX-9	97.66	80	96	36.59	30.09	29.55	1.92
IBM Cell	91.91	64	128	103.90	81.08	92.54	0.61
NEC x86 cluster	892.86	512	1024	103.04	61.59	68.67	0.83

Table 21: Results for 10 TFlop/s partitions: CPMD, Test Case A.

6.4 EUTERPE

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	130.93	1.75
1024	3.48	66.49	1.72
1536	5.22	44.83	1.70
2048	6.96	34.34	1.66
3072	10.44	23.42	1.63
5120	17.41	14.87	1.54
6144	20.89	14.12	1.35
11264	38.30	8.57	1.21
12288	41.78	7.22	1.32
22528	76.60	5.68	0.91
23552	80.08	4.96	1.00

Table 22: Results for EUTERPE, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	82.07	2.06
512	4.71	43.45	1.94
1024	9.42	21.7	1.95
1536	14.13	14.63	1.92
2048	18.84	11.15	1.89
2560	23.55	9	1.88
4096	37.68	6.03	1.75
4608	42.39	5.27	1.78

Table 23: Results for EUTERPE, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	159.71	1.04
256	4.81	78.77	1.05
512	9.63	41.32	1.00
1024	19.25	20.76	1.00
1536	28.88	14.27	0.97
2048	38.50	10.91	0.95
2560	48.13	8.94	0.92

Table 24: Results for EUTERPE, Test Case A on IBM Power6.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	3072	34.34	23.42	24.41	1.63
Cray XT5b	1086.96	1024	1536	21.70	14.63	20.48	1.94
IBM Power6	531.91	512	1024	41.32	20.76	39.79	1.00

Table 25: Results for 10 TFlop/s partitions: EUTERPE, Test Case A.

6.5 GADGET

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	549.4	0.07
1024	3.48	284.63	0.07

Table 26: Results for GADGET, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	22.625	1.27
256	4.81	10.46	1.37
512	9.63	6.82	1.05
1024	19.25	10.89	0.33

Table 27: Results for GADGET, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.50	57.34	0.80
256	3.00	29.56	0.78
512	6.00	13.76	0.84

Table 28: Results for GADGET, Test Case A on Bull x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512	1024	6.82	10.89	6.92	1.00

Table 29: Results for 10 TFlop/s partitions: GADGET, Test Case A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	46.85	1.15
256	4.81	22.09	1.22
512	9.63	12.96	1.04
1024	19.25	13.84	0.49

Table 30: Results for GADGET, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.50	125.17	0.69
256	3.00	62.91	0.69
512	6.00	28.51	0.76

Table 31: Results for GADGET, Test Case B on Bull x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512	1024	12.96	13.84	12.99	1.000

Table 32: Results for 10 TFlop/s partitions: GADGET, Test Case B.

6.6 GROMACS

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
1440	4.90	1745	0.59
1536	5.22	1565	0.61
2880	9.79	896	0.57
3072	10.44	872	0.55
5832	19.83	485	0.52
6144	20.89	485	0.49

Table 33: Results for GROMACS, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	807	2.64
288	2.65	718	2.63
512	4.71	431	2.47
576	5.30	406	2.33
1024	9.42	254	2.09
1152	10.60	234	2.02
2160	19.87	157	1.61
2188	20.13	148	1.68
4320	39.74	103	1.22
4374	40.24	106	1.18

Table 34: Results for GROMACS, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
216	2.33	747	2.88
256	2.76	649	2.79
432	4.67	398	2.70
512	5.53	350	2.59
864	9.33	234	2.30
1024	11.06	206	2.20

Table 35: Results for GROMACS, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	9.63	500	1.04
544	10.23	502	0.98
1024	19.25	287	0.91
1056	19.85	292	0.86
1088	20.45	288	0.85
2048	38.50	168	0.77
2112	39.71	175	0.72
2176	40.91	178	0.69

Table 36: Results for GROMACS, Test Case A on IBM Power6.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2880	3072	896.00	872.00	888.21	0.56
Cray XT5b	1086.96	1024	1152	254.00	234.00	243.75	2.06
Cray XT5s	925.93	864	1024	234.00	206.00	222.31	2.25
IBM Power6	531.91	512	544	500.00	502.00	501.24	1.00

Table 37: Results for 10 TFlop/s partitions: GROMACS, Test Case A.

6.7 NAMD

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
735	2.50	0.1206	0.80
1471	5.00	0.0630	0.77
2048	6.96	0.0501	0.70
2560	8.70	0.0416	0.67
6144	20.89	0.0195	0.59
8192	27.85	0.0148	0.59
11765	40.00	0.0113	0.53

Table 38: Results for NAMD, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
272	2.50	0.0947	1.02
543	5.00	0.0488	0.99
1087	10.00	0.0286	0.85
2174	20.00	0.0166	0.73
4348	40.00	0.0121	0.50

Table 39: Results for NAMD, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
231	2.49	0.0887	1.10
463	5.00	0.0456	1.06
926	10.00	0.0264	0.92

Table 40: Results for NAMD, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	0.0898	1.12
160	3.01	0.0693	1.16
256	4.81	0.0442	1.14
288	5.41	0.0405	1.11
512	9.63	0.0247	1.02
544	10.23	0.0240	0.99
1024	19.25	0.0145	0.87
1056	19.85	0.0149	0.82
1088	20.45	0.0151	0.79
2048	38.50	0.0122	0.52
2112	39.71	0.0109	0.56

Table 41: Results for NAMD, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
208	2.44	0.0872	1.14
216	2.53	0.0853	1.12
424	4.97	0.0622	0.78
432	5.06	0.0643	0.74
848	9.94	0.1649	0.15
856	10.03	0.2483	0.10
1704	19.97	0.3507	0.03
1712	20.06	0.3949	0.03

Table 42: Results for NAMD, Test Case A on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2560.0000	6144	0.0416	0.02	0.0371	0.65
Cray XT5b	1086.96	1087.0000	1087	0.0286	0.03	0.0286	0.85
Cray XT5s	925.93	926.0000	926	0.0264	0.03	0.0264	0.92
IBM Power6	531.91	512.0000	544	0.0247	0.02	0.0242	1.00
Sun x86 Cluster	853.24	848.0000	856	0.2483	0.25	0.2483	0.10

Table 43: Results for 10 TFlop/s partitions: NAMD, Test Case A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
272	2.50	0.8331	0.90
543	5.00	0.4057	0.92
1087	10.00	0.2118	0.88
2174	20.00	0.1136	0.82

Table 44: Results for NAMD, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
231	2.49	0.7547	0.99
463	5.00	0.3839	0.97
926	10.00	0.2005	0.93

Table 45: Results for NAMD, Test Case B on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	0.7275	1.07
160	3.01	0.6018	1.03
256	4.81	0.3472	1.12
288	5.41	0.3257	1.06
512	9.63	0.1878	1.03
544	10.23	0.1863	0.98
1024	19.25	0.1071	0.91
1056	19.85	0.1044	0.90
1088	20.45	0.1066	0.86
2048	38.50	0.0651	0.75
2112	39.71	0.0640	0.73
2144	40.31	0.0614	0.76

Table 46: Results for NAMD, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
208	2.44	0.7285	1.05
216	2.53	0.7116	1.04
424	4.97	0.4645	0.81
432	5.06	0.4759	0.78
848	9.94	1.0538	0.18
856	10.03	1.3225	0.14
1704	19.97	1.9268	0.05
1712	20.06	2.7655	0.03

Table 47: Results for NAMD, Test Case B on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
Cray XT5b	1086.96	1087.0000	1087	0.2118	0.21	0.2118	0.88
Cray XT5s	925.93	926.0000	926	0.2005	0.20	0.2005	0.93
IBM Power6	531.91	512.0000	544	0.1878	0.19	0.1869	1.00
Sun x86 Cluster	853.24	848.0000	856	1.0538	1.32	1.2157	0.15

Table 48: Results for 10 TFlop/s partitions: NAMD, Test Case B.

6.8 NEMO

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
735	2.50	63.4	0.97
1472	5.00	53.0	0.58

Table 49: Results for NEMO, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
272	2.50	68.3	0.90
536	4.93	46.7	0.67

Table 50: Results for NEMO, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
232	2.51	66.6	0.92
464	5.01	43.3	0.71

Table 51: Results for NEMO, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	54.9	1.16
160	3.01	40.6	1.25
256	4.81	30.6	1.04
288	5.41	23.1	1.22
512	9.63	18.4	0.86
544	10.23	13.9	1.08
1056	19.85	9.9	0.78
1088	20.45	10.3	0.73

Table 52: Results for NEMO, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.64	20.6	4.54

Table 53: Results for NEMO, Test Case A on NEC SX-9

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512	544	18.4000	13.90	15.3145	1.00

Table 54: Results for 10 TFlop/s partitions: NEMO, Test Case A.

6.9 NS3D

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
320	1.09	1252.5	-
1280	4.35	325.0	-
2560	8.70	166.2	-
3840	13.06	155.7	-
7680	26.11	88.6	-

Table 55: Results for NS3D, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	2090.6	-
280	2.58	2077.9	-
480	4.42	1147.6	-
560	5.15	1025.8	-

Table 56: Results for NS3D, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	922.5	-
160	3.01	844.9	-
256	4.81	472.3	-
320	6.02	441.4	-

Table 57: Results for NS3D, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.50	965.4	-
256	3.00	533.4	-
320	3.75	415.8	-
480	5.63	288.5	-
640	7.50	233.8	-
960	11.25	162.0	-

Table 58: Results for NS3D, Test Case A on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.64	304.7	-
32	3.28	164.8	-
48	4.92	149.3	-
64	6.55	101.1	-
128	13.11	59.3	-

Table 59: Results for NS3D, Test Case A on NEC x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.43	805.1	-
256	2.87	321.0	-
320	3.58	366.2	-
480	5.38	181.4	-
640	7.17	193.2	-
960	10.75	116.6	-
1280	14.34	93.0	-
2560	28.67	52.7	-

Table 60: Results for NS3D, Test Case A on NEC SX-9.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
1280	4.35	661.3	1.79
2560	8.70	324.3	1.83
3840	13.06	228.7	1.73
7680	26.11	154.9	1.27
10240	34.82	123.6	1.20
15360	52.22	147.0	0.67

Table 61: Results for NS3D, Test Case B on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
480	4.42	2494.4	0.47
560	5.15	2267.5	0.44
960	8.83	1244.1	0.47
1120	10.30	1143.4	0.44

Table 62: Results for NS3D, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	4.81	1071.2	1.00
512	9.63	534.7	1.00
640	12.03	431.1	0.99

Table 63: Results for NS3D, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	1311.7	1.31
320	3.75	872.1	1.58
480	5.63	637.3	1.44
640	7.50	477.1	1.44
960	11.25	492.2	0.93
1280	15.00	433.2	0.79
2560	30.00	332.3	0.52

Table 64: Results for NS3D, Test Case B on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.64	620.1	5.07
32	3.28	334.0	4.71
48	4.92	261.3	4.01
64	6.55	167.8	4.69
128	13.11	112.9	3.48

Table 65: Results for NS3D, Test Case B on NEC SX-9.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	816.1	2.20
320	3.58	665.9	2.16
480	5.38	545.0	1.76
640	7.17	409.2	1.76
960	10.75	200.5	2.39
1280	14.34	145.3	2.47
2560	28.67	116.0	1.55

Table 66: Results for NS3D, Test Case B on NEC x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
2560	8.70	933.0	1.95
3840	13.06	614.3	1.97
7680	26.11	361.2	1.68
10240	34.82	277.9	1.63
15360	52.22	237.4	1.27

Table 67: Results for NS3D, Test Case C on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
480	4.42	8293.2	0.43
560	5.15	7043.5	0.44
960	8.83	5262.0	0.34
1120	10.30	4183.5	0.37
2080	19.14	3275.1	0.25
2240	20.61	2478.9	0.31

Table 68: Results for NS3D, Test Case C on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	4.81	3393.6	0.97
320	6.02	2824.7	0.93
512	9.63	1651.4	0.99
640	12.03	1281.1	1.03
1056	19.85	837.5	0.95
1120	21.06	763.1	0.98

Table 69: Results for NS3D, Test Case C on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
960	11.25	955.2	1.47
2560	30.00	868.4	0.61
3840	45.00	694.8	0.51

Table 70: Results for NS3D, Test Case C on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
32	3.28	803.3	6.00
48	4.92	660.1	4.87
64	6.55	423.3	5.70
128	13.11	223.8	5.39

Table 71: Results for NS3D, Test Case C on NEC SX-9.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
960	10.75	595.8	2.47
1280	14.34	780.6	1.41
2560	28.67	349.9	1.58

Table 72: Results for NS3D, Test Case C on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2560.0	3840	932.95	614.34	808.14	1.96
Cray XT5b	1086.96	960.0	1120	5261.98	4183.45	4368.36	0.36
IBM Power6	531.91	512.0	640	1651.43	1281.10	1580.35	1.00
Sun x86 cluster	853.24	960.0	2560	955.2	868.40	961.65	1.64
NEC SX-9	97.66	64.0	128	423.28	223.77	288.17	5.48
NEC x86 cluster	892.86	960.0	2560	595.75	349.93	613.85	2.57

Table 73: Results for 10 TFlop/s partitions: NS3D, Test Case C.

6.10 QCD

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	525.00	1.20
1024	3.48	344.30	0.91
2048	6.96	143.60	1.09
4096	13.93	86.61	0.91
8192	27.85	41.00	0.96
16384	55.71	24.64	0.80
32768	111.41	16.05	0.61
65536	222.82	5.71	0.86

Table 74: Results for QCD, Kernel A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	475.40	0.98
512	4.71	253.61	0.92
1024	9.42	134.20	0.87
2048	18.84	88.03	0.66
4096	37.68	64.22	0.45

Table 75: Results for QCD, Kernel A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.76	421.3	0.94
512	5.53	220	0.90

Table 76: Results for QCD, Kernel A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	757.00	0.60
256	4.81	237.00	0.96
512	9.63	112.00	1.02
1024	19.25	69.80	0.81

Table 77: Results for QCD, Kernel A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	237.9	1.53
512	6.00	119.6	1.52
1024	12.00	57.45	1.59
2048	24.00	37.24	1.22
4096	48.01	45.47	0.50

Table 78: Results for QCD, Kernel A on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	288.4	1.32
512	5.73	160.1	1.19
1024	11.47	102.2	0.93
2048	22.94	72.83	0.66

Table 79: Results for QCD, Kernel A on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	143.60	86.61	111.58	0.98
Cray XT5b	1086.96	1024	2048	134.20	88.03	130.01	0.84
IBM Power6	531.91	512	1024	112.00	69.80	109.43	1.00
Sun x86 Cluster	853.24	512	1024	119.6	57.45	69.49	1.57
NEC x86 Cluster	892.86	512.00	1024	102.2	72.83	78.62	1.39

Table 80: Results for 10 TFlop/s partitions: QCD, Kernel A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	257.00	1.08
1024	3.48	129.90	1.07
2048	6.96	64.96	1.07
4096	13.93	36.88	0.94
8192	27.85	16.91	1.03
16384	55.71	9.37	0.92
32768	111.41	4.80	0.90
65536	222.82	2.51	0.86

Table 81: Results for QCD, Kernel B on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.18	178.70	2.29
256	2.36	89.01	2.30
512	4.71	44.66	2.29
1024	9.42	22.48	2.28
2048	18.84	11.17	2.29
4096	37.68	5.50	2.33

Table 82: Results for QCD, Kernel B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.38	152.80	2.29
256	2.76	76.53	2.28
512	5.53	37.72	2.31

Table 83: Results for QCD, Kernel B on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	220.00	0.91
256	4.81	99.00	1.01
512	9.63	50.20	1.00
1024	19.25	24.80	1.01

Table 84: Results for QCD, Kernel B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	48.34	3.33
512	6.00	24.03	3.35
1024	12.00	12.21	3.29
2048	24.00	6.84	2.94
4096	48.01	4.41	2.28

Table 85: Results for QCD, Kernel B on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	55.01	3.06
512	5.73	25.97	3.24
1024	11.47	13.49	3.12
2048	22.94	7.63	2.76

Table 86: Results for QCD, Kernel B on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	64.96	36.88	48.77	0.99
Cray XT5b	1086.96	1024	2048	22.48	11.17	21.16	2.28
IBM Power6	531.91	512	1024	50.20	24.80	48.28	1.00
Sun x86 Cluster	853.24	512	1024	24.03	12.21	14.61	3.31
NEC x86 Cluster	892.86	512.00	1024	25.97	13.49	15.38	3.14

Table 87: Results for 10 TFlop/s partitions: QCD, Kernel B.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	117.00	0.61
1024	3.48	115.80	0.62
2048	6.96	115.50	0.62
4096	13.93	125.20	0.57
8192	27.85	116.10	0.61
16384	55.71	117.00	0.61
32768	111.41	126.80	0.56
65536	222.82	128.60	0.55

Table 88: Results for QCD, Kernel C on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.18	54.45	1.31
256	2.36	58.02	1.23
512	4.71	57.19	1.25
1024	9.42	56.84	1.25
2048	18.84	58.63	1.21
4096	37.68	57.64	1.24

Table 89: Results for QCD, Kernel C on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.38	50.00	1.42
256	2.76	50.51	1.41
512	5.53	55.2	1.29

Table 90: Results for QCD, Kernel C on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	57.60	1.24
256	4.81	51.10	1.39
512	9.63	72.00	0.99
1024	19.25	56.30	1.27

Table 91: Results for QCD, Kernel C on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	44.6	1.60
512	6.00	41.86	1.70
1024	12.00	45.33	1.57
2048	24.00	61.83	1.15
4096	48.01	101.4	0.70

Table 92: Results for QCD, Kernel C on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	63.31	1.13
512	5.73	63.49	1.12
1024	11.47	62.42	1.14
2048	22.94	64.54	1.10

Table 93: Results for QCD, Kernel C on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	115.50	125.20	119.54	0.60
Cray XT5b	1086.96	1024	2048	56.84	58.63	56.95	1.25
IBM Power6	531.91	512	1024	72.00	56.30	71.23	1.00
Sun x86 Cluster	853.24	512	1024	41.86	45.33	44.11	1.61
NEC x86 Cluster	892.86	512.00	1024	63.49	62.42	62.69	1.14

Table 94: Results for 10 TFlop/s partitions: QCD, Kernel C.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	2.92	1.59
1024	3.48	1.49	1.56
2048	6.96	0.75	1.55
4096	13.93	0.44	1.33
8192	27.85	0.24	1.23
16384	55.71	0.16	0.89

Table 95: Results for QCD, Kernel D on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.18	3.44	2.00
256	2.36	1.81	1.90
512	4.71	0.96	1.79
1024	9.42	0.54	1.59
2048	18.84	0.33	1.30
4096	37.68	0.28	0.77

Table 96: Results for QCD, Kernel D on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.38	3.16	1.86
256	2.76	1.64	1.79
512	5.53	0.92	1.59

Table 97: Results for QCD, Kernel D on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	2.73	1.23
256	4.81	1.40	1.20
512	9.63	0.82	1.03
1024	19.25	0.63	0.67

Table 98: Results for QCD, Kernel D on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	0.85	3.18
512	6.00	0.45	3.00
1024	12.00	0.28	2.41
2048	24.00	0.33	1.02
4096	48.01	0.5	0.34

Table 99: Results for QCD, Kernel D on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	1.06	2.67
512	5.73	0.58	2.44
1024	11.47	0.35	2.02
2048	22.94	0.2	1.77

Table 100: Results for QCD, Kernel D on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	0.75	0.44	0.57	1.42
Cray XT5b	1086.96	1024	2048	0.54	0.33	0.52	1.56
IBM Power6	531.91	512	1024	0.82	0.63	0.81	1.00
Sun x86 Cluster	853.24	512	1024	0.45	0.28	0.32	2.53
NEC x86 Cluster	892.86	512.00	1024	0.58	0.35	0.39	2.08

Table 101: Results for 10 TFlop/s partitions: QCD, Kernel D.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	546.00	1.37
1024	3.48	279.90	1.34
2048	6.96	143.50	1.31
4096	13.93	75.91	1.23
8192	27.85	39.78	1.18
16384	55.71	23.55	0.99
32768	111.41	12.80	0.91

Table 102: Results for QCD, Kernel E on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	365.80	1.51
512	4.71	193.70	1.43
1024	9.42	103.10	1.34
2048	18.84	56.55	1.22
4096	37.68	33.40	1.04

Table 103: Results for QCD, Kernel E on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.76	343.80	1.37
512	5.53	185.1	1.27

Table 104: Results for QCD, Kernel E on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	698.00	0.78
256	4.81	269.00	1.01
512	9.63	134.00	1.01
1024	19.25	78.70	0.86

Table 105: Results for QCD, Kernel E on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	180.1	2.41
512	6.00	93.36	2.33
1024	12.00	50.04	2.17
2048	24.00	37.59	1.45
4096	48.01	30.21	0.90

Table 106: Results for QCD, Kernel E on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	239.1	1.90
512	5.73	111.4	2.04
1024	11.47	58.2	1.95
2048	22.94	38.9	1.46

Table 107: Results for QCD, Kernel E on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	143.50	75.91	103.36	1.26
Cray XT5b	1086.96	1024	2048	103.10	56.55	98.13	1.33
IBM Power6	531.91	512	1024	134.00	78.70	130.44	1.00
Sun x86 Cluster	853.24	512	1024	93.36	50.04	59.20	2.20
NEC x86 Cluster	892.86	512.00	1024	111.4	58.20	66.31	1.97

Table 108: Results for 10 TFlop/s partitions: QCD, Kernel E.

System	Mean $E^{10,A}$
IBM BlueGene/P	1.01
Cray XT5b	1.38
IBM Power6	1.00
Sun x86 Cluster	2.16
NEC x86 cluster	1.83

Table 109: Results for 10 TFlop/s partitions: mean relative efficiency for QCD, Kernels A-E.

6.11 QUANTUM_ESPRESSO

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16384	55.71	513.00	0.24
32768	111.41	279.00	0.22
65536	222.82	163.00	0.19

Table 110: Results for QUANTUM_ESPRESSO, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	4.71	1722.00	0.85
1024	9.42	896.00	0.82
2048	18.84	450.00	0.81

Table 111: Results for QUANTUM_ESPRESSO, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	4.81	1206.00	1.19
512	9.63	690.00	1.04
1024	19.25	439.00	0.82

Table 112: Results for QUANTUM_ESPRESSO, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	1360.00	1.69
512	6.00	747.00	1.54
1024	12.00	556.00	1.03
2048	24.00	510.00	0.56

Table 113: Results for QUANTUM_ESPRESSO, Test Case A on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.87	1900.00	1.27
512	5.73	996.00	1.21
1024	11.47	518.00	1.16
2048	22.94	272.00	1.11

Table 114: Results for QUANTUM_ESPRESSO, Test Case A on NEC x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.64	2775.00	1.52
32	3.28	1897.00	1.11
48	4.92	1440.00	0.97
64	6.55	1170.00	0.90

Table 115: Results for QUANTUM_ESPRESSO, Test Case A on NEC SX-9.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
Cray XT5b	1086.96	1024	2048	896.00	450.00	844.54	0.82
IBM Power6	531.91	512	1024	690.00	690.00	690.00	1.00
Sun x86 Cluster	853.24	512	1024	747.00	556.00	607.83	1.14
NEC x86 cluster	892.86	512.00	1024	996	518.00	590.60	1.17

Table 116: Results for 10 TFlop/s partitions: QUANTUM_ESPRESSO, Test Case A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16384	55.71	1955.00	0.25
32768	111.41	1419.00	0.17
65536	222.82	560.00	0.22

Table 117: Results for QUANTUM_ESPRESSO, Test Case B on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
1024	9.42	2573.00	1.14

Table 118: Results for QUANTUM_ESPRESSO, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	4.81	5580.00	1.03
512	9.63	2857.00	1.00
1024	19.25	1505.00	0.95

Table 119: Results for QUANTUM_ESPRESSO, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	6.00	2466.00	1.87
1024	12.00	1957.00	1.18

Table 120: Results for QUANTUM_ESPRESSO, Test Case B on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	1.64	28356	0.59

Table 121: Results for QUANTUM_ESPRESSO, Test Case B on NEC SX-9.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	5.73	3362.00	1.43
1024	11.47	1599.00	1.51

Table 122: Results for QUANTUM_ESPRESSO, Test Case B on NEC x86 cluster.

6.12 WRF

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	8.73	1.27
735	2.50	6.28	1.23
1024	3.48	4.72	1.18
2048	6.96	2.44	1.14
2941	10.00	1.74	1.11
4096	13.93	1.40	0.99
5882	20.00	0.97	1.00
8192	27.85	0.79	0.88
11712	39.82	0.61	0.80
11776	40.04	0.63	0.77
16384	55.71	0.51	0.68
23520	79.97	0.39	0.61
23552	80.08	0.39	0.62

Table 123: Results for WRF, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	4.44	1.85
512	4.71	2.28	1.80
1024	9.42	1.18	1.74
2048	18.84	0.67	1.54
4096	37.68	0.48	1.07

Table 124: Results for WRF, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.38	6.56	2.13
256	2.76	4.41	1.58
512	5.53	2.40	1.45
1024	11.06	1.18	1.49

Table 125: Results for WRF, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	7.67	1.05
160	3.01	6.60	0.97
256	4.81	4.32	0.93
288	5.41	3.68	0.97
512	9.63	2.03	0.99
544	10.23	1.88	1.01
1056	19.85	1.02	0.96
1088	20.45	0.99	0.96
2112	39.71	0.51	0.96
2144	40.31	0.50	0.96

Table 126: Results for WRF, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
208	2.44	3.11	2.55
216	2.53	3.11	2.46
424	4.97	1.57	2.48
432	5.06	1.55	2.46
848	9.94	0.83	2.34
856	10.03	0.86	2.24
1704	19.97	0.45	2.15
1712	20.06	0.45	2.14
3408	39.94	0.86	0.56
3416	40.04	0.86	0.56
3424	40.13	0.86	0.56

Table 127: Results for WRF, Test Case A on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
208	2.44	3.58	2.22
216	2.53	3.45	2.21
424	4.97	1.73	2.25
432	5.06	1.91	2.00
800	9.38	1.10	1.88

Table 128: Results for WRF, Test Case A on Bull x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
216	2.42	2.96	2.70
224	2.51	3.15	2.45
440	4.93	1.78	2.20
448	5.02	1.65	2.34
888	9.95	0.98	1.98
896	10.04	0.99	1.95
1792	20.07	0.52	1.85

Table 129: Results for WRF, Test Case A on NEC x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2941.00	4096	1.74	1.40	1.74	1.11
Cray XT5b	1086.96	1024.00	2048	1.18	0.67	1.13	1.72
Cray XT5s	925.93	512.00	1024	2.40	1.18	1.30	1.48
IBM Power6	531.91	512.00	544	2.03	1.88	1.93	1.00
Sun x86 Cluster	853.24	848.00	856	0.83	0.86	0.85	2.28
Bull x86 Cluster	853.24	432.00	800	1.91	1.10	1.04	1.87
NEC x86 cluster	892.86	888.00	896	0.98	0.99	0.99	1.96

Table 130: Results for 10 TFlop/s partitions: WRF, Test Case A.

6.13 ALYA

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
133	2.50	156.79	0.84
266	5.00	79.76	0.82
532	10.00	32.89	1.00
1064	20.00	20.65	0.80
2128	40.01	15.39	0.53

Table 131: Results for ALYA, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
23	2.50	6850.3	0.02
46	5.00	3416.4	0.02
92	10.01	1762.6	0.02

Table 132: Results for ALYA, Test Case A on IBM Cell.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	266	532	79.76	32.89	32.90	1.00
IBM Cell	91.91	46	92	3416.4	1762.60	1764.24	0.02

Table 133: Results for 10 TFlop/s partitions: ALYA, Test Case A.

6.14 AVBP

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	4.81	548.91	1.00
512	9.63	273.46	1.00
1024	19.25	136.75	1.00
2048	38.50	68.93	0.99

Table 134: Results for AVBP, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	3.00	398.55	2.20
512	6.00	202.59	2.17
1024	12.00	122.65	1.79
2048	24.00	107.22	1.02

Table 135: Results for AVBP, Test Case A on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512	1024	273.46	136.75	263.22	1.00
Sun x86 Cluster	853.24	512	1024	122.65	107.22	111.92	2.35

Table 136: Results for 10 TFlop/s partitions: AVBP, Test Case A.

6.15 BSIT

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
32	0.60	6027.11	1.03
160	3.01	6087.43	1.01
320	6.02	6147.64	1.00
736	13.84	6207.91	1.00

Table 137: Results for BSIT, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
1	0.11	988.04	6.25
5	0.54	997.91	6.19
10	1.09	1007.79	6.13
23	2.50	1031.43	5.99
46	5.00	1057.2	5.84

Table 138: Results for BSIT, Test Case A on IBM Cell.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	320	736	6147.64	6207.91	6178.20	1.00

Table 139: Results for 10 TFlop/s partitions: BSIT, Test Case A.

6.16 ELMER

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
272	2.50	259.76	2.52
280	2.58	255.1	2.50
536	4.93	135.8	2.45
544	5.00	136.67	2.40
1080	9.94	79.46	2.08
1088	10.01	72.01	2.28
2168	19.95	48.96	1.69
2176	20.02	49.7	1.66

Table 140: Results for ELMER, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
133	2.50	747	0.88
266	5.00	354	0.93
532	10.00	164	1.00
1064	20.00	100	0.82
2128	40.01	73.63	0.56

Table 141: Results for ELMER, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
213	2.50	192.77	3.43
427	5.00	112.09	2.94
853	10.00	75.34	2.19
1706	19.99	111.5	0.74
3413	40.00	272.13	0.15

Table 142: Results for ELMER, Test Case A on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
Cray XT5b	1086.96	1080	1088	79.46	72.01	72.90	2.25
IBM Power6	531.91	266	532	354	164.00	164.03	1.00
Sun x86 Cluster	853.24	853	1706	75.34	111.50	75.35	2.19

Table 143: Results for 10 TFlop/s partitions: ELMER, Test Case A.

6.17 GPAW

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	207.40	1.26
1024	3.48	117.00	1.12

Table 144: Results for GPAW, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	2.36	94	2.05
384	3.53	63.5	2.03
512	4.71	56.8	1.70
588	5.41	47.6	1.77
1024	9.42	34.1	1.42
1176	10.82	31.1	1.35

Table 145: Results for GPAW, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
224	2.42	87.7	2.14
256	2.76	78.9	2.09
448	4.84	52.6	1.79
512	5.53	48.5	1.70
896	9.68	31.3	1.50
1024	11.06	29.2	1.41

Table 146: Results for GPAW, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	131.8	1.43
256	4.81	80	1.18
512	9.63	46.7	1.01
768	14.44	34.8	0.91
1024	19.25	30.3	0.78
1344	25.27	28.2	0.64

Table 147: Results for GPAW, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
196	2.30	55.5	3.57
256	3.00	45.4	3.34
392	4.59	38.3	2.59
512	6.00	35.7	2.12
1024	12.00	46.3	0.82

Table 148: Results for GPAW, Test Case A on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
Cray XT5b	1086.96	1024	1176	34.10	31.10	32.79	1.39
Cray XT5s	925.93	896	1024	34.10	31.10	33.35	1.36
IBM Power6	531.91	512	768	46.70	34.80	45.49	1.00
Sun x86 Cluster	853.24	512	1024	35.7	46.30	42.13	1.08

Table 149: Results for 10 TFlop/s partitions: GPAW, Test Case A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
1024	9.42	638.6	1.44
1152	10.60	611.4	1.34
2048	18.84	313.4	1.46
2304	21.20	318.9	1.28
4096	37.68	160.4	1.43
4608	42.39	162.6	1.25

Table 150: Results for GPAW, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
360	3.89	1500.9	1.48
512	5.53	934.2	1.67
720	7.78	714.9	1.56
1024	11.06	467.5	1.67

Table 151: Results for GPAW, Test Case B on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	9.63	897.6	1.00
640	12.03	722.9	0.99
1024	19.25	448.9	1.00
1152	21.66	425.1	0.94

Table 152: Results for GPAW, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
720	8.44	1184.5	0.87
1024	12.00	782.8	0.92
1440	16.88	583.4	0.88
2048	24.00	390.1	0.92
2880	33.75	304.2	0.84
4096	48.01	197.3	0.91

Table 153: Results for GPAW, Test Case B on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
Cray XT5b	1086.96	1024	1152	638.6	611.40	624.93	1.38
Cray XT5s	925.93	720	1024	714.9	467.50	526.25	1.64
IBM Power6	531.91	512	640	897.60	722.90	865.07	1.00
Sun x86 Cluster	853.24	720	1024	1184.50	782.80	967.01	0.89

Table 154: Results for 10 TFlop/s partitions: GPAW, Test Case B.

6.18 HELIUM

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
630	2.14	3563	1.00
1540	5.24	1615	0.90
3003	10.21	700	1.07
6105	20.76	389	0.94
9870	33.56	405	0.56

Table 155: Results for HELIUM, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
253	2.33	2420	1.35
406	3.74	1400	1.46
630	5.80	723	1.82
990	9.11	462	1.81
1540	14.17	302	1.78
2485	22.86	253	1.32

Table 156: Results for HELIUM, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
253	2.73	1776	1.57
406	4.38	1013	1.72
630	6.80	642	1.74
990	10.69	419	1.70

Table 157: Results for HELIUM, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
105	1.97	1723	2.24
253	4.76	1248	1.28
630	11.84	670	0.96
990	18.61	449	0.91
2485	46.72	298	0.55

Table 158: Results for HELIUM, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
406	4.76	612	2.62
990	11.60	333	1.97
1540	18.05	277	1.52
3003	35.20	317	0.68

Table 159: Results for HELIUM, Test Case A on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	1540	3003	1615.00	700.00	717.17	1.06
Cray XT5b	1086.96	990	1540	462.00	302.00	422.54	1.80
Cray XT5s	925.93	630	990	642.00	419.00	446.61	1.71
IBM Power6	531.91	253	630	1248.00	670.00	761.79	1.00
Sun x86 Cluster	853.24	406	990	612.00	333.00	372.80	2.04

Table 160: Results for 10 TFlop/s partitions: HELIUM, Test Case A.

6.19 OCTOPUS

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
256	0.87	103.10	1.66
512	1.74	56.16	1.52
1024	3.48	32.65	1.31

Table 161: Results for OCTOPUS, Test Case A on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
32	0.29	243.66	2.07
64	0.59	133.66	1.89
128	1.18	72.83	1.73
136	1.25	44.853	2.45
272	2.50	26.1765	2.10
536	4.93	18.636	1.50

Table 162: Results for OCTOPUS, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
136	1.47	40.86	2.29
272	2.94	23.70	1.98
536	5.79	17.04	1.39

Table 163: Results for OCTOPUS, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
64	1.20	107.23	1.15
128	2.41	70.58	0.87
256	4.81	43.06	0.72
512	9.63	15.59	0.99

Table 164: Results for OCTOPUS, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
8	0.87	2949.58	0.06
16	1.74	1478.58	0.06
32	3.48	756.26	0.06
64	6.96	657.95	0.03

Table 165: Results for OCTOPUS, Test Case A on IBM Cell.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
16	0.19	301.10	2.63
32	0.38	160.28	2.47
64	0.75	81.77	2.42
128	1.50	42.86	2.31

Table 166: Results for OCTOPUS, Test Case A Bull x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	256.00	512	43.06	15.59	14.85	1.00

Table 167: Results for 10 TFlop/s partitions: OCTOPUS, Test Case A.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	112.18	1.43
1024	3.48	59.35	1.35
2048	6.96	31.84	1.26

Table 168: Results for OCTOPUS, Test Case B on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
64	0.59	270.96	1.75
128	1.18	146.37	1.62
256	2.36	80.16	1.48
440	4.05	32.5	2.12

Table 169: Results for OCTOPUS, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
224	2.42	49.74	2.32
880	9.50	20.44	1.44

Table 170: Results for OCTOPUS, Test Case B on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	99.72	1.16
256	4.81	53.08	1.09
512	9.63	27.26	1.07
1024	19.25	74.61	0.19

Table 171: Results for OCTOPUS, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
32	3.48	1897.79	0.04
64	6.96	797.39	0.05
128	13.93	697.90	0.03

Table 172: Results for OCTOPUS, Test Case B on IBM Cell.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
32	0.38	367.32	2.03
64	0.75	169.75	2.20
128	1.50	88.48	2.11
256	3.00	47.82	1.95

Table 173: Results for OCTOPUS, Test Case B Bull x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM Power6	531.91	512.00	1024	27.26	74.61	27.95	1.00
IBM Cell	91.91	64.00	128	797.39	697.90	750.72	0.04

Table 174: Results for 10 TFlop/s partitions: OCTOPUS, Test Case B.

6.20 PEPC

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.18	7.08	6.10
256	2.36	3.92	5.51
512	4.71	2.41	4.48

Table 175: Results for PEPC, Test Case A on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.38	5.89	6.25
256	2.76	3.28	5.61
512	5.53	2.05	4.49

Table 176: Results for PEPC, Test Case A on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	2.41	6.67	3.17
256	4.81	5.01	2.11
512	9.63	5.08	1.04

Table 177: Results for PEPC, Test Case A on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.50	3.45	9.83
256	3.00	1.94	8.74
512	6.00	1.27	6.67
1024	12.00	1.12	3.78

Table 178: Results for PEPC, Test Case A on Sun x86 cluster.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	1.74	102.00	1.18
1024	3.48	57.30	1.05
2048	6.96	29.50	1.02
4096	13.93	19.60	0.77

Table 179: Results for PEPC, Test Case B on IBM BlueGene/P.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	4.71	23.7	1.88
1024	9.42	14.2	1.57
2048	18.84	8.7	1.28
4400	40.48	7.03	0.74

Table 180: Results for PEPC, Test Case B on Cray XT5b.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
128	1.38	5.89	6.25
256	2.76	3.28	5.61
512	5.53	2.05	4.49

Table 181: Results for PEPC, Test Case B on Cray XT5s.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	9.63	21.3	1.02
1024	19.25	15	0.73
2048	38.50	9.57	0.57
2304	43.32	10.3	0.47

Table 182: Results for PEPC, Test Case B on IBM Power6.

No. of CPUs	Partition size (TFlop/s)	Time (sec)	Relative Efficiency
512	6.00	12.4	2.82
1024	12.00	7.69	2.27
2048	24.00	4.7	1.86
4096	48.01	4.03	1.08

Table 183: Results for PEPC, Test Case B on Sun x86 cluster.

System	P_t	P_l	P_u	T_l	T_u	T_t	$E^{10,A}$
IBM BlueGene/P	2941.18	2048	4096	29.50	19.60	24.17	0.87
Cray XT5b	1086.96	1024	2048	14.20	8.70	13.67	1.53
Cray XT5s	925.93	512	1024	19.70	12.00	12.97	1.62
IBM Power6	531.91	512	1024	21.3	15.00	20.96	1.00
Sun x86 Cluster	853.24	512	1024	12.4	7.69	8.81	2.38

Table 184: Results for 10 TFlop/s partitions: PEPC, Test Case B.