


01
02 **Performance of**
03 **computation-intensive,**
04 **parameter sweep applications on**
05 **Internet-based Grids of**
06 **computers: the mapping of**
07 **molecular potential energy**
08 **hypersurfaces**
09
10
11
12
13



14
15 S. Reyes¹, C. Muñoz-Caro¹, A. Niño^{1,*},[†], R. M. Badia² and J. M. Cela²
16

17 ¹*QCyCAR. E. S. Informática. Universidad de Castilla-La Mancha,*
18 *Paseo de la Universidad, 4. 13071 Ciudad Real, Spain*

19 ²*Barcelona Supercomputing Center, Campus Nord UPC, Mòdul D6, Jordi Girona,*
20 *1-3. 08034 Barcelona, Spain*

21
22
23 **SUMMARY**

24 **This work focuses on the use of computational Grids for processing the large set of jobs arising in parameter**
25 **sweep applications. In particular, we tackle the mapping of molecular potential energy hypersurfaces.**
26 **For computationally intensive, parameter sweep problems, performance models are developed to compare**
27 **the parallel computation in a multiprocessor system with the computation on an Internet-based Grid of**
28 **computers. We find that the relative performance of the Grid approach increases with the number of**
29 **processors, being independent of the number of jobs. The experimental data, obtained using electronic**
30 **structure calculations, fits the proposed performance expressions accurately. To automate the mapping of**
31 **potential energy hypersurfaces, an application based on GRID superscalar is developed. It is tested on the**
32 **prototypical case of the internal dynamics of acetone. Copyright © 2006 John Wiley & Sons, Ltd.**

33 *Received 4 April 2006; Revised 30 June 2006; Accepted 5 July 2006*

34 **KEY WORDS:** Grid computing; parameter sweep; performance modeling; GRID superscalar; potential energy
35 hypersurfaces; electronic structure calculations
36
37

38 *Correspondence to: A. Niño, QCyCAR. E. S. Informática. Universidad de Castilla-La Mancha, Paseo de la Universidad, 4,
39 13071 Ciudad Real, Spain.

40 [†]E-mail: alfonso.nino@uclm.es
41

42 Contract/grant sponsor: Consejería de Educación y Ciencia de la Junta de Comunidades de Castilla-La Mancha; contract/grant
number: PBI05-009

43 Contract/grant sponsor: Ministerio de Educación y Ciencia; contract/grant number: FIS2005-00293

44 Contract/grant sponsor: Universidad de Castilla-La Mancha

Contract/grant sponsor: Ministry of Science and Technology of Spain; contract/grant number: TIN2004-07739-C02-01



01 1. INTRODUCTION

02
03 Grid computing is a distributed computing model that permits the integration of computational
04 resources from different administrative domains, in order to define a new virtual system. This approach
05 permits to take advantage of individual computational resources, which are usually geographically
06 distributed [1–3]. The Grid is considered the next leap in computer interconnectivity. In the same
07 way as we now share files and facts over the Internet, the Grid lets us share other resources, such as
08 processing power, software or storage space. The connection between the distributed computational
09 resources and the Grid applications is performed by using a layer of middleware software. At present,
10 different middleware tools are available permitting to create and execute software applications on a
11 Grid environment in a transparent way. As examples we have GRID superscalar [4–6], GridWay [7]
12 or P-Grade [8]. In particular, GRID superscalar is a Grid programming framework for Grid unaware
13 applications that allows users to develop applications in a computational Grid. The underlying runtime
14 library is able to automatically parallelize the application. Applications that benefit from GRID
15 superscalar are those composed of tasks, of a certain granularity, which are called several times. GRID
16 superscalar automatically detects the data-dependencies between the tasks of the application, building
17 a task graph on run-time. This allows the user to exploit the inherent parallelism of the application,
18 and concurrent task submission is performed when resources are available. GRID superscalar can use
19 different underlying Grid middlewares as Globus Toolkit 2, Globus Toolkit 4, Ninf-G or ssh/scp.

20 Many scientific disciplines are considering Grid computing as a new tool. For instance, the
21 astrophysical [9], high-energy physics [10] or computational chemistry [11,12] communities are
22 resorting to Grid computing to perform complex simulations or to process large amounts of data.
23 In particular, Grid computing can be applied to parameter sweep problems, which arise naturally
24 in several scientific and engineering computing fields. Here, we have a system defined by a set of
25 independent parameters. The behavior of the system is analyzed by varying the parameters. This
26 produces a set of jobs that can be computed independently. Depending on the size of the parameter
27 space, the number of jobs can be very large. The problem of the scheduling of these jobs has been
28 considered in heterogeneous environments [13] and, specifically, on computational Grids [14,15].
29 In this context, a parameter sweep problem of great interest in the physical-molecular field is the
30 evaluation of molecular potential energy hypersurfaces.

31 Any kind of physical phenomenon related to the internal motion of a molecule depends on the
32 potential energy binding the nuclei of the system. Examples of these phenomena are the conformational
33 distribution of molecules in a vacuum or solution, or the characteristic pattern of energy levels
34 defining the rotovibrational spectrum of a molecule. In particular, the ability to compute an arbitrary
35 potential energy hypersurface will be an invaluable aid in molecular spectroscopic studies. This is so
36 because the potential hypersurface, and its associated molecular structures, permit the construction
37 of rotational and vibrational models. The models serve to interpret and predict the characteristics of
38 rotovibrational molecular spectra. These spectra are chiefly of interest in atmospheric and astrophysical
39 studies to identify molecular species, as well as to determine the conditions under which they are
40 observed [16,17].

41 For a system of N (nonlinear) nuclei, its relative position is defined by a total of $3N - 6$ parameters.
42 Therefore, the potential energy for the motion of the N nuclei defines a $3N - 6$ dimensional
43 hypersurface. In the framework of the Born–Oppenheimer, adiabatic, approximation, the sum of the
44 quantum mechanical electronic energy plus the electrostatic repulsion energy (for a fixed nuclei)



01 defines the potential for the nuclear motion [18]. Being able to perform quantum-mechanical, electronic Q1
02 structure calculations, the adiabatic approximation provides a parameter sweep approach for mapping
03 a potential energy hypersurface. The technique involves performing a series of electronic structure
04 calculations for different values of the position parameters of interest. From these results, it is possible
05 to write and solve a Hamiltonian for the motion. In this form, the corresponding energy levels
06 (spectroscopically observable) can be obtained, see, for instance, [19]. Formally, it would be possible
07 to deal with the full set of $3N - 6$ parameters, or vibration modes. However, in practice, until now only
08 a few internal motions can be considered simultaneously.

09 Nowadays, electronic structure methods are available in software packages, and some of them are
10 freely available [20–22]. The quantum-mechanical molecular models implemented in such software
11 tools are very accurate. The predictions they provide for molecular and electronic structure are
12 in very good agreement with the experimental measurements, see, for instance, [23]. However,
13 the computational cost grows with the sophistication of the methodology employed. From such
14 calculations, we can in principle derive a potential energy hypersurface as accurate as needed.
15 The problem is how, and where, to perform the large number of point computations needed for this
16 kind of study. In other words, we are facing a problem halfway between high-performance and high-
17 throughput computing.

18 Two extreme solutions are possible for a parameter sweep problem such as the mapping of
19 potential energy hypersurfaces. The classical solution is to compute each data point in parallel on a
20 multiprocessor system. Here, the parallelization requires the management of data dependence between
21 the tasks running on independent processors and belonging to the same calculation. We will term
22 this the fine-grained approach. On the other hand, it is possible to compute each point in a single
23 processor, but computing simultaneously as many points as processors. In this case, the calculations
24 are data independent. This last solution will be termed the coarse-grained approach. This approach
25 is well suited for implementation on an Internet-based Grid of computers. For instance, it has been
26 applied to the development of an effective pseudoatom potential for hybrid quantum mechanical–
27 molecular mechanical studies [24]. This potential was derived from a large number of electronic
28 structure calculations obtained by coupling the Gamess electronic structure code [20] to the Nimrod/G
29 Grid distribution tool [25]. However, no performance considerations were taken into account in this
30 work. Due to the difference in communication costs between a dedicated multiprocessor system and an
31 Internet distributed system, the adoption of the fine or coarse-grained approach depends on the relative
32 performance for the problem at hand.

33 For a given process in a multiprocessor system, the performance depends not only on the execution
34 time, but also on several communication factors. Different performance models have been developed
35 to explicitly account for these communication costs. Thus, the classical LogP approach [26] was based
36 on a Parallel Random Access Machine model. LogP includes the overhead in the transmission and
37 reception of each message, the time gap between consecutive message transmissions or receptions,
38 and the latency in communicating small messages. Variants of LogP have been proposed, reflecting
39 the evolution of the computing architectures. Thus, LogGP accounts for long messages [27]. LoPC
40 and LoGPC are extensions to active messages [28,29], and LoGPS accounts for synchronization
41 costs [30]. Performance models tailored for distributed systems have also been proposed [31–33]. Also,
42 an empirical study has been presented recently on the performance behavior of Grid connected clusters
43 using the NPB set of applications [34].

44 In this work, we analyze and test the efficiency of the computational Grid paradigm for
computationally intensive parameter sweep problems. In particular, we consider the automatic



01 exploration of potential energy hypersurfaces. Performance models are built to compare a (coarse-
 02 grained) Internet-based Grid approach with a parallel (fine-grained) approach on a multiprocessor
 03 system. Experimental measures are used to test the validity of the proposed mathematical models.
 04 In addition, a prototype of an application able to automatically generate a set of molecular structures,
 05 run the electronic structure calculations and organize the potential energy results on a computational
 06 Grid is presented. The application is tested by considering the torsion of the two methyl groups plus
 07 the coupling with the two CCO angles, which defines the molecular frame, in the acetone molecule.

08

09

10 2. PERFORMANCE MODELS

11

12 2.1. Fine-grained approach

13

14 Let us consider a parallel system (for instance a computer cluster) where one of the processors acts as
 15 the master and several processors act as servers. The total time, t_n , used by a parallel program running
 16 on n server processors is composed of the time associated to the communication costs plus the time
 17 corresponding to the execution of the code. The different factors affecting performance in currently
 18 distributed systems have been included by Le and Rejeb [32] in a modified LogGP performance model.
 19 Using this model, t_n can be expressed as,

20

21

$$t_n = o_m + L + G \cdot S + o_s + h_m + h_s + t_{\text{Seq}} + t_p \quad (1)$$

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

where o is the master (o_m) or server (o_s) overhead arising from the communication activities of the
 processors, L is the transmission latency, G is the message gap per byte, which is the transmission
 time per unit message, and S is the message length. Both, L and G can be considered independent of
 the number of processors, n . The term h is the function introduced by Le and Rejeb [32] to describe
 the overhead cost associated with the use of middleware in the parallelization process in the master
 (h_m) and server (h_s). Finally, t_{Seq} and t_p represent the time used by the sequential and parallelizable
 parts of the code, respectively. In Equation (1) we have magnitudes that increase with n , usually those
 associated to the master. However, other magnitudes decrease with n , such as S , t_p , or even those
 associated to the servers. Therefore, from a general standpoint, Equation (1) can be expressed as

31

32

$$t_n = A + B(n) + C(1/n) \quad (2)$$

33

34

35

36

37

38

39

40

41

42

43

44

where A is a constant, and B and C are functions that increase or decrease with n , respectively.
 When we observe the behavior of actual systems for increasing values of n , we find that the execution
 time reaches a minimum, and then begins to rise, as a consequence of the increasing role played by
 communication costs, term $B(n)$ in Equation (2). Therefore, for small or moderate values of n the
 $C(1/n)$ term dominates. In these conditions we have,

38

39

$$(dB/dn) \ll (dC/dn) \quad (3)$$

40

41

42

43

44

and Equation (2) can be approximated by

$$t_n \cong A' + C(1/n) \quad (4)$$

43

44

where A' is a constant. Thus, we have that the total execution time in n processors is composed
 by an approximately constant term plus a function of $1/n$. In this approach the computation time



01 will be always reduced when increasing n . This represents an overestimation of the performance
02 of parallelization for large n . In other words, we will have an upper (ideal) performance limit for
03 comparison with the Grid approach. This model is simple to use and it will be shown to be useful in
04 real cases, when a limited number of processors are considered.

05 Since in the conditions defining Equation (4) we have $t_1 > t_n$ for all n , we can establish an analogy
06 with Amdahl's law [35]. Thus, we can assume that t_n can be obtained from t_1 by considering that a
07 proportion, f , of the process is parallelizable. Applying Amdahl's law [35] we obtain

$$08 \quad t_n = t_1 \left(1 + \frac{f \cdot (1 - n)}{n} \right) \quad (5)$$

11 Comparing Equation (5) with Equation (4) we have

$$12 \quad f = \left(\frac{A' + C[1/n]}{t_1} \right) \frac{n}{(1 - n)} \quad (6)$$

13 Equation (6) shows that f is not a constant, and can be considered as a function of $1/n(f[1/n])$. In
14 other words, $f[1/n]$ does not depend exclusively on the proportion of parallelizable code because
15 it also incorporates communication costs. From Equation (6) we can derive a performance index
16 appropriate for the conditions defining Equation (3), i.e. not very large n values.

17 The usual performance index is the speedup, defined as $S = t_1/t_n$. Using Equation (6), we can define
18 a computational effort index, E , as the inverse of the speedup:

$$19 \quad E = S^{-1} = 1 + \frac{f[1/n](1 - n)}{n} \quad (7)$$

20 E , as the speedup, is a non-dimensional magnitude. However, E is normalized in the interval $[0,1]$ for
21 any parallel program. For one server processor the effort is maximum, $E = 1$, decreasing as the number
22 of processors increases. Equation (7) indicates that the highest attainable performance assuming no
23 overheads costs, i.e. the limit for $n \rightarrow \infty$, depends on the form of $f[1/n]$.

24 To give $f[1/n]$ an analytical form, it is useful to expand it in a Taylor series around the $n = 1$ point,

$$25 \quad f(1/n) = f_1 + \frac{\partial f(1/n)}{\partial(1/n)} \Big|_1 (1/n - 1) + \frac{1}{2} \frac{\partial^2 f(1/n)}{\partial(1/n)^2} \Big|_1 (1/n - 1)^2 + \dots \quad (8)$$

26 With Equations (7) and (8), we can give to E the general expression:

$$27 \quad E = 1 + a \frac{1 - n}{n} + b \frac{1 - n}{n^2} + c \frac{1 - n}{n^3} + \dots \quad (9)$$

28 where a, b, c, \dots , are constants. With respect to Equation (9) two interesting cases can be considered.
29 First, where $f[1/n]$ is a constant (see Equation (8)) we obtain

$$30 \quad E = 1 + a(1 - n)/n \quad (10)$$

31 Second, considering a linear variation of $f[1/n]$ on $1/n$ (see Equation (8)), we obtain

$$32 \quad E = 1 + a(n - 1)/n + b(n - 1)/n^2 \quad (11)$$

33 Thus, we can derive an analytical expression for E by making some calibration measures followed by
34 a fitting to Equations (10) or (11).



In the ideal case where there are no overhead costs, Equation (4) is valid for any n and the minimum possible effort can be obtained as

$$E_{\infty} = \lim_{n \rightarrow \infty} E = 1 - a \quad (12)$$

Equation (12) shows that $1 - a$ is the maximum attainable percentage of time reduction, $\max E$. The term 'a' can be interpreted as the maximum parallelizability proportion under ideal conditions.

2.2. Coarse-grained approach

Now, a total of m independent tasks will be run in n processors, with $n < m$. In principle, different scheduling schemas for allocating the m process in the n processors are possible [36–40]. We will assume, for the sake of simplicity, that the system is homogeneous and that the computation time of each task is the same, t_1 . In these conditions, we can apply a chunk self-scheduling scheme to minimize communication overheads [36]. Thus, we allocate a total of m/n tasks to each of the n processors. Since we are allocating jobs in different processors, the total computing time must include communication costs. Therefore,

$$t_{\text{total}} = \sum_i^{m/n} t_i = \frac{m}{n} t_1 + (o_m + L + G \cdot S + o_s + h_m + h_s) \quad (13)$$

where we have used the same nomenclature as in Equation (1). The experience shows that, when considering the same system, communication costs are smaller in Equation (13) than in Equation (1). However, the comparison is not so direct when the fine-grained approach is applied on a dedicated system, and the coarse-grained approach in an Internet-based Grid of computers, due to the differences in latencies, bandwidths and middlewares. In this last case, it is useful to define an F index measuring the ratio of communication to execution costs:

$$F = n(o_m + L + G \cdot S + o_s + h_m + h_s)/(m t_1) \quad (14)$$

The F index is formally defined in the interval $(0, \infty)$. Equation (14) shows increasing the number of jobs, m , we decrease the relative weight of the communication costs.

With respect to the electronic structure computations, we must consider that the input data size is small, usually one or two KB in length. On the other hand, and depending on the type of calculation, the results file can have a size of several hundred KB, reaching several hundred MB in the worst case scenario. However, only a limited amount of information is needed when mapping a hypersurface (total energy and vibrational coordinates, at least). Taking into account that a parameter sweep problem involves thousands of points, preselecting locally the return information will reduce communication costs dramatically. In addition, if we consider that typical calculations run at least for a few minutes, and it is not unusual for a calculation to run for a few hours, we can assume that the communication time can be neglected compared to the execution time, i.e. we have a small F factor. Therefore, we will have

$$t_{\text{total}} \cong \frac{m}{n} t_1 \quad (15)$$

Equation (15) shows that the total time depends linearly on the number of calculations and decreases with the inverse of the number of processors. Considering that for one processor $t_{\text{total}} = m t_1$, the



01 computational effort will be

$$02 \quad E = \frac{1}{n} \quad (16)$$

03
04 With the previous assumptions, let us quantify the relative performance of the fine versus the coarse-
05 grained approaches. In the fine-grained approach, for m simultaneous calculations each running on n
06 processors, the total time, t_f , is given by m times t_n . Using Equation (5), we have

$$07 \quad \lim_{n \rightarrow \infty} t_f = \lim_{n \rightarrow \infty} m t_1 \left(1 + \frac{f[1/n](1-n)}{n} \right) = m t_1 (1 - C) \quad (17)$$

08
09 where C is a constant equal to (see Equation (8))

$$10 \quad C = f_1 - \left. \frac{\partial f(1/n)}{\partial(1/n)} \right|_1 + \frac{1}{2} \left. \frac{\partial^2 f(1/n)}{\partial(1/n)^2} \right|_1 - \dots \quad (18)$$

11
12 Equation (17) is an underestimation, since the $B(n)$ term of Equation (2) is not taken into account.
13 This term will make the time increase for large n . On the other hand, from Equation (15), and taking
14 into account that the maximum number of processors in the coarse-grained approach equals the number
15 of tasks, we have for the optimal case

$$16 \quad \lim_{n \rightarrow m} t_c = \lim_{n \rightarrow m} \frac{m}{n} t_1 = t_1 \quad (19)$$

17
18 Comparing equations (17) and (19) we observe that the coarse-grained approach is more efficient as
19 the number of tasks, m , increases. Assuming m to be large (as in the case of the mapping of a potential
20 energy hypersurface, which can involve several thousand points), the coarse-grained approach seems
21 an appropriate solution.

22 We can determine the relative time gain (RTG) of the coarse versus the fine approach. Thus, dividing
23 $t_f - t_c$ by t_f and converting to percent, we obtain

$$24 \quad RTG(\%) = 100 \times \left(1 - \frac{1}{n + a(1-n) + b(1-n)/n + \dots} \right) \quad (20)$$

25
26 Equation (20) shows that the RTG is independent of the number of calculations, but depends on the
27 amount of parallelizable code and on the number of processors used. For the simplest case reflected in
28 Equation (11) we obtain

$$29 \quad RTG(\%) = 100 \times \left(1 - \frac{1}{n + a(1-n)} \right) \quad (21)$$

30
31 where 'a' represents the proportion of parallelizability. Equation (21) shows that even with a 90%
32 parallelizability a great gain is obtained with a small number of processors (as much as a 50% for
33 $n = 10$).

34 The present results show that a coarse-grained approach to a parameter sweep problem involving
35 several thousand time-consuming independent points, executed on individual processors of an Internet-
36 based Grid of computers, is a highly efficient process. The next step is to analyze experimentally the
37 behavior of this kind of computation on a real Grid system.




Figure 1. Structure of the acetone molecule. The numbering convention used in this work and the definitions of the θ_1 , θ_2 , α and β angles are included.

3. METHODOLOGY

As an example of a parameter sweep problem we will use the mapping of molecular potential energy hypersurfaces. For the tests we have selected the acetone molecule. This molecule is considered a prototype for large amplitude vibrations due to the rotation of its two methyl groups. Figure 1 shows the structure and numbering convention of acetone used in this work. We will consider as the main parameters the two methyl torsional angles θ_1 , and θ_2 , defined by dihedrals O1-C2-C4-H5 and O1-C2-C3-H8, respectively (see Figure 1). As usual, an a-b-c-d dihedral angle corresponds to the angle subtended by planes a-b-c and b-c-d. In addition, we will consider the α and β valence angles (see Figure 1). These two angles define the structure of the CCCO main frame of acetone. Considering the four angles, we are describing the coupling of the two methyl torsions with the acetone molecular frame. Electronic structure calculations on acetone will be carried out at the MP2/6-311G (2d, p) theory level [23]. This level of theory has shown to provide excellent results for the potential energy function for two [41] and three [42] internal degrees of freedom in the molecule of acetaldehyde.

The starting point is a full geometry optimization of acetone at the selected theory level. This defines the zero of relative energy on our hypersurface. Starting from this optimized structure, we perform a scan of the considered angles. In this form, we generate the structures for the mapping of the potential energy hypersurface. For each point on the hypersurface, the values of the angles considered have been kept fixed. All other structural parameters are fully relaxed. All calculations have been performed with the Gamess electronic structure package [20].

The parallel, fine-grained calculations have been carried out on a cluster (named Tales) formed by 12 single-processor Pentium IV, 2.4–3.0 GHz, 1 MB RAM, using a fast Ethernet connection and local storage of intermediate results, as described in reference [43]. The system runs under the cluster configuration and management tool Rocks [44], 3.0.0 version. The details of the Gamess code parallelization can be found in reference [45].

The coarse-grained calculations have been performed both on Tales and on a computational Grid formed by five nodes. The first, Sofocles, is the Grid client. It is a single-processor machine AMD-2.4 GHz from the QCyCAR research group at Ciudad Real, Spain. The next two are two PC clusters of 12 nodes each, also from the QCyCAR research group. The first, Tales, is described above. The second (Hermes) is similar, but consists of Pentium IV, 2.6–2.8 GHz machines. The third node is formed



01 by the Kadesh supercomputing system at the Barcelona Supercomputing Center, Spain. Kadesh is a
02 cluster formed by 128 Power3/375 MHz + 32 Power4/1 GHz nodes with 512 MB of RAM each.
03 The fourth node is a cluster (Popocatepetl) of nine biprocessor 64 bits AMD machines, 1.4 GHz,
04 1 GB RAM, local storage and ethernet network, placed at the Laboratorio de Química Teórica in
05 the Universidad de Puebla, Mexico. Thus, the Grid implies a transatlantic connection. Hermes and
06 Popocatepetl run under Rocks, version 3.0.0 and 3.2.0, respectively. The Grid runs the Globus toolkit
07 system [46], 2.4.3 version. The heterogeneity of the system is also reflected in the different queuing
08 systems used: Hermes and Tales run OpenPBS [47], Popocatepetl uses Sun Grid Engine [48] and
09 Kadesh LoadLeveler [49]. The independent jobs are sent on the Grid using the 1.6.0 version of GRID
10 superscalar. This version is built on top of Globus Toolkit 2.4.

11 In order to maintain the consistency of the experimental conditions, all the electronic structure
12 calculations use (independently of the number of processors) the parallel algorithm implemented in
13 Gamess for optimization at the MP2/6-311G (2d, p) level of theory [45]. To just focus on performance
14 issues, the Grid has been used as a dedicated system in the tests.

15

16

17 4. RESULTS AND DISCUSSION

18

19 The performance models presented in Equations (10), (11) and (21) can be compared by making some **Q2**
20 experimental measurements. The tests are carried out considering a potential energy hypersurface for
21 the acetone molecule.

22

23 4.1. Determination of Gamess performance at the MP2/6-311G (2d, p) theory level

24

25 First, we determine the proportion of parallelizability when working with Gamess at the MP2/6-311G
26 (2d, p) theory level. Thus, we run a full geometry optimization of the acetone molecule varying the
27 number of processors from 1 to 11 in the Tales cluster. Figure 2 collects the results expressed as
28 computational effort, E , relative to a single-processor time of 1448 seconds. The data are fitted to
29 expressions of the type given by Equation (10), for a constant proportion of parallelizability, and
30 Equation (11), for a linear dependence with $(1 - n)/n$. The results, including squared correlation
31 coefficient (R^2) and residual sum of squares (RSS), are respectively:

32

$$33 E = 1 + 0.7725(n - 1)/n \quad \text{with } R^2 = 0.956, \text{ RSS} = 0.023 \quad (22)$$

34

$$35 E = 1 + 0.8665(n - 1)/n - 0.5558(n - 1)/n^2 \quad \text{with } R^2 = 0.990, \text{ RSS} = 0.005 \quad (23)$$

36

36 We observe that the fit to the linear case has a higher correlation and a much smaller residual
37 sum of squares. Thus, we will use this case as reference. Therefore, with Equation (23), and from
38 Equation (12), the proportion of effective parallelizability is 86.65%, or approximately 90%.

39

40 4.2. Automatic generation of molecular structures

41

42 To generate the battery of molecular structures as a function of the θ_1 , θ_2 , α , and β angles in
43 our parameter sweep problem, see Figure 1, we need the minimum energy structure of acetone.
44 Previous experimental [50], as well as theoretical [51–53] studies identified the $\theta_1 = \theta_2 = 0^\circ$ as the




Figure 2. Experimental variation of the E computational effort index with the number of processors, n , in the Tales cluster. The data are obtained for the full geometry optimization of acetone at the MP2/6-311G (2d, p) theory level. The diamonds mark the measured data. The dashed and continuous lines correspond to the fitted curves given by Equations (22) and (23), respectively.

most stable conformer. Thus, we perform a full molecular structure optimization, starting from this configuration, to determine all the structural parameters (distances and angles) in the minimum energy conformation. As expected, our results are $\theta_1 = \theta_2 = 0^\circ$ for this conformer. The other parameters relevant for this study, the α and β angles, are found to be 121.8° . This value is in very good agreement with the 121.9° derived from the experimental data [50]. The battery of points is generated from these starting values.

The results of Niño *et al.* [42], on the coupling of the methyl group and the CCCO molecular frame in acetaldehyde show a maximum variation of 1.5° for the CCO angle. Therefore, we have selected a 2° positive and negative variation around the $\approx 122^\circ$ α and β equilibrium value. Thus, the α and β values are generated starting from 120° to 124° in increments of 1° . On the other hand, considering the symmetry of acetone [51–53], the torsion of the two methyl groups is described using increments of 20° from 0° to 120° on the θ_1 , and θ_2 angles with the restriction $\theta_1 \geq \theta_2$ for $\alpha = \beta$. A total of 1120 different molecular structures are obtained in this way. Generation of the set of points is carried out using nested loops involving the θ_1 , θ_2 , α and β angles.

4.3. Validation of the relative performance model for the coarse versus the fine-grained parameter sweep approach

Using the eleven processors of the Tales cluster, we can validate Equation (21) considering a variable number of points on the potential energy hypersurface of acetone. These measures are carried out at the cluster level using the OpenPBS queuing system [47]. Thus, sets from 10 to 500 molecular structures, arbitrarily selected from the 1120 molecular structures, are considered. The coarse-grained



Table I. Experimental results of the fine versus the coarse-grained approaches for different number (m) of molecular structures in the acetone molecule. Each calculation implies a partial molecular structure optimization at the MP2/6-311G (2d, p) theory level. The FG column collects the total time, in seconds, for the calculation of each set of points in the fine-grained approach. The CG column collects equivalent results for the coarse-grained case.

m	FG	CG	RTG (%)
10	2775	1320	52.4
50	17 314	7080	59.1
100	35 772	14 520	59.4
200	75 297	31 080	58.7
300	123 752	49 140	60.3
400	166 350	62 640	62.3
500	209 231	79 380	62.0

approach is applied using the chunk self-scheduling schema [36] with a chunk size of one, i.e. a pure self-scheduling approach. This approach optimizes load balancing. This is specially useful in our case, where the computation time of each job will vary from case to case and where the communication overhead can be neglected compared to the computation time. Thus, we allocate single-processor calculations to free processors until the set of jobs is finished. Table I collects the results. We observe that the RTG value exhibits little variation. The value stabilizes in a 62% for large values of m , the number of data points. This stabilization can be considered as a consequence of the statistical average of the heterogeneity of the system. The observed RTG value of 62% is in good agreement with the 64.8% predicted by Equation (21) using the results of Equation (23).

The previous results validate the performance model proposed in Equation (21). The next step is to apply this approach to the Grid mapping of molecular potential energy hypersurfaces.

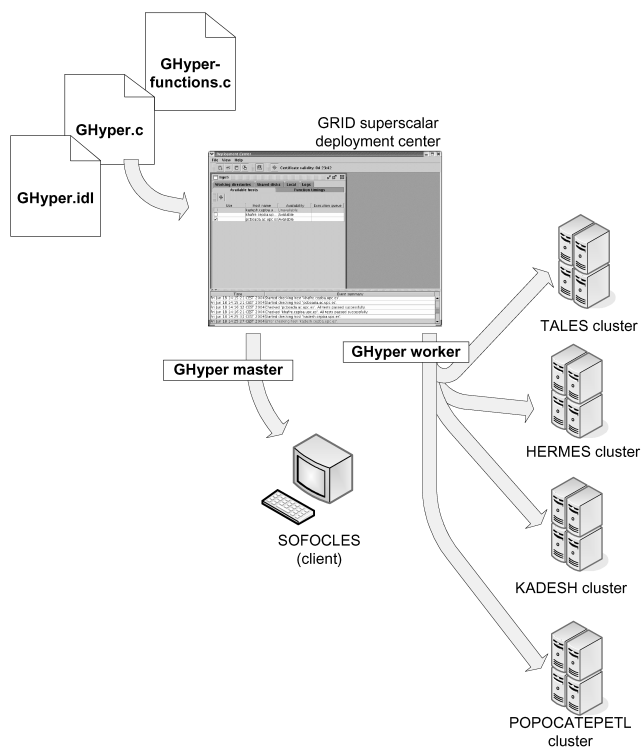
4.4. Automating the mapping of potential energy hypersurfaces

The computation of a potential energy hypersurface can be regarded as a single problem, rather than as a sum of individual electronic structure calculations. Thus, we need to implement a unique process able to generate the molecular structures, to perform the individual calculations and to integrate the results. Each calculation needs to run an instance of an electronic structure code with its own input and output file. Our aim is to implement this technique at the Grid level.

From this point of view, a general methodology can be postulated to solve the problem of the generation of molecular potential energy hypersurfaces on a computational Grid. The process is organized in three steps.



01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

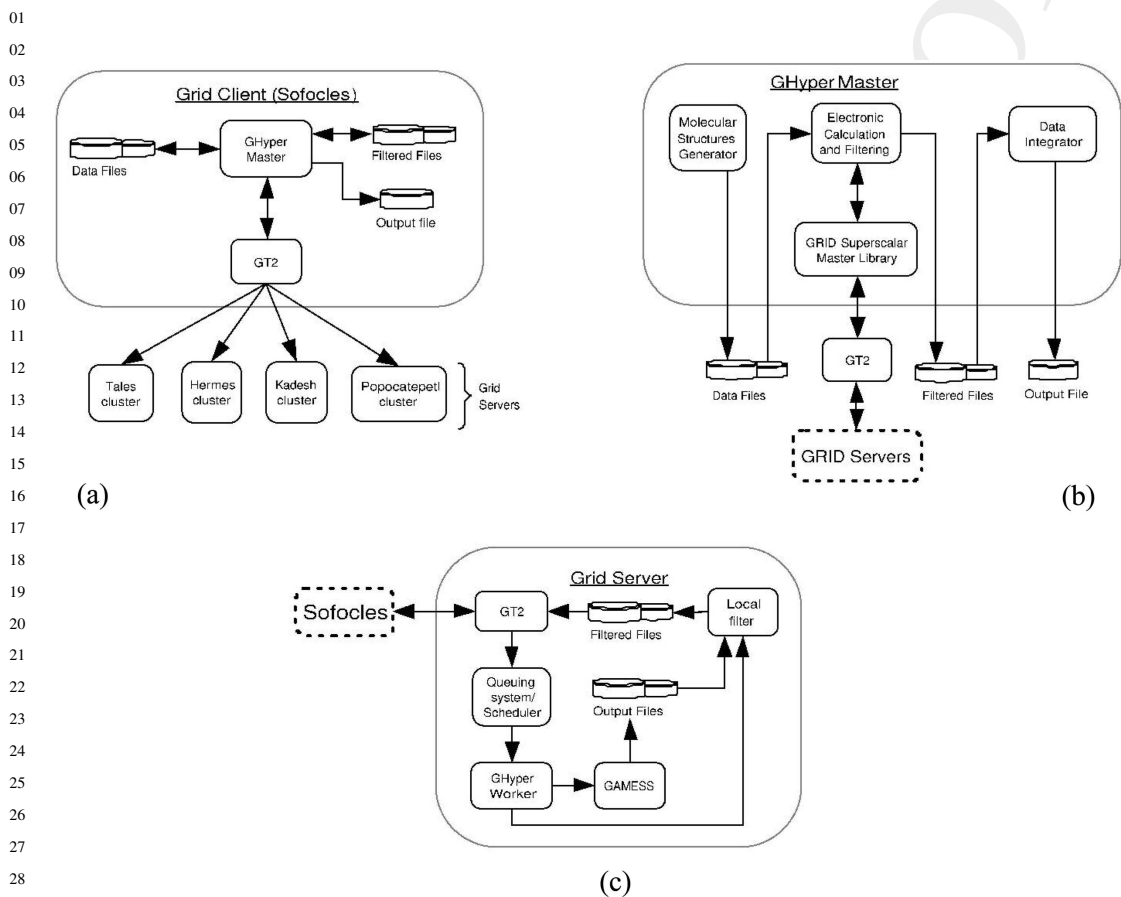


26 Figure 3. Representation of the files, and the installation process, needed to run the GHyper application with GRID
27 superscalar in a computational Grid.

- 28
29
- 30 1. *Molecular structure generation*: this step involves the generation of the set of molecular
31 structures defining the potential energy hypersurface. The result is a large number of input data
32 files.
 - 33 2. *Electronic structure calculation*: one evaluation with the electronic structure package is executed
34 for each of the data files generated in step 1. Since the output of each of these evaluations is a
35 large output file, a filtering process that obtains the required information (molecular coordinates
36 and total energy) is applied.
 - 37 3. *Data integration*: the data generated by each of the calculations in step 2 is integrated in a single
38 ASCII file.

39 While steps 1 and 3 are not computationally intensive, the evaluation of the different data of step 2
40 would take a long time in a sequential machine. Since the number of points for this kind of job is large
41 (in the current case 1120 electronic structure-independent computations), it is interesting to find a way
42 to parallelize and execute this step in the Grid.

43 For this *gridification* we use GRID superscalar. The process described by the three previous steps
44 is implemented in a sequential C program, shown in Figure 3 as file GHyper.c. Also, to use GRID



30 Figure 4. The structure of the application proposed for the automatic computation of molecular potential energy
31 hypersurfaces on a computational Grid. GT2 refers to Globus toolkit version 2.4.3. Case (a) is the high-level
32 description of the whole system. Case (b) is the structure of the GHYper Master process. Case (c) shows the
33 internal working of one of the Grid nodes.

34
35
36 superscalar, we provided the GHYper.idl and GHYper-functions.c files, containing the interfaces and
37 code of the functions to be run in the Grid [4–6]. Two functions are needed here: one for wrapping and
38 calling the electronic structure package, and a second function for the filtering process. The rest of the
39 functions in the application, for example the functions that execute steps 1 and 3, are not listed in the
40 GHYper.idl file, and therefore will be locally executed.

41 From these three files, and using the GRID superscalar deployment center, two binaries are
42 generated: the GHYper master, which is installed in the client machine; and the GHYper worker,
43 installed on each of the Grid servers machines. The behavior of the whole system is schematically
44 depicted in Figures 4(a)–(c).



Table II. Number of jobs (molecular structures) computed on each node of the Grid. The table also collects the total time (in seconds) used to compute the 1120 total jobs. Cases (a) and (b) corresponds to the 32 and 64 nodes distributions, respectively, as described in the text.

	Case (a)	Case (b)
Tales	310	223
Hermes	290	218
Popo	286	262
Kadesh	234	417
Total time	112 500	61 500

The GHYper master starts performing the generation of the input molecular structures files. The program loops over the considered angles and prepares the input files for the Gamess electronic structure package (step 1), see Figure 4(b). Next, the GHYper master continues with step 2. Since the GHYper master is dynamically linked with the GRID superscalar master library, for each electronic structure calculation and filtering process the GRID superscalar broker is queried to decide to which server in the Grid the processes are submitted, see Figure 4(b). Each server in the Grid receives the GRID superscalar requests and the local GHYper worker finally calls the electronic structure calculation (Gamess package) and the filtering function, see Figure 4(c).

All file transfers, scheduling and allocation decisions are performed by the GRID superscalar libraries, in a totally transparent way for the initial GHYper application. Finally, the information (files) generated at the different Grid nodes is sent back by GRID superscalar to the client machine, and integrated in a single output ASCII file (step 3), see Figure 4(a).

The described application is tested considering the four-dimensional problem of acetone. Two different Grid configurations have been used in the tests. The first one uses eight machines (processors) in each cluster of the Grid. Thus, 32 processors are used, distributed as follows: eight on Tales, eight on Hermes, eight on Kadesh and eight on Popocatepetl. The second configuration uses a total of 64 processors: 11 on Tales, 11 on Hermes, 14 on Popocatepetl and 28 processors on Kadesh. The results for both cases are collected in Table II. In absolute terms, we observe that even with 32 processors the time needed to compute the 1120 data points is more than reasonable (31 hours and 15 minutes). When the number of processors doubles, the total time is reduced almost to half (0.55), despite the physical and logical heterogeneity of the system, as a consequence of the linear scaling of the coarse approach, as described by Equations (13) and (15).

These results show that the Grid approach makes possible the computation of the thousands of points needed to map a molecular energy hypersurface in several dimensions. Finally, it is remarkable that the client machine (Sofocles, a 32-bit PC) has been able to handle all the connections on the Grid without trouble. This is a consequence of the asynchrony of the connections needed to send (or receive) data (or results). This asynchrony is favored by the heterogeneity of the system, which makes execution times different from node to node.




Figure 5. Two-dimensional potential energy hypersurface for acetone as a function of the θ_1 , and θ_2 angles. Data obtained in Grid at the MP2/6-311G (2d, p) theory level. Energy in cm^{-1} and angles in degrees. The interval between isocontour lines is 100 cm^{-1} .

4.5. Example of application: Grid evaluation of a potential energy hypersurface for acetone

For the two-dimensional analysis of the acetone hypersurface on the θ_1 , and θ_2 angles, our Grid application generates 91 molecular structure files. After submission to the Grid and the collection of the results, the energy data points are used to generate a potential energy hypersurface. This is performed by interpolating with the Kriging method [54] using the Surfer package [55]. Figure 5 shows the resulting hypersurface. The barrier to rotation is obtained as the difference between the minima (for instance, point A in Figure 5) and the first order saddle point connecting two minima (for instance, point B in Figure 5). The value obtained from the interpolated Kriging grid is 271 cm^{-1} , whereas the value obtained directly from the data is 270 cm^{-1} . Table III compares our value to previous theoretical and experimental data. We find that our result, $270\text{--}271 \text{ cm}^{-1}$, is within the interval of experimental results, $251\text{--}290 \text{ cm}^{-1}$, obtained from fitting experimental rotational and vibrational (torsional) data to different rotovibrational models.

On the other hand, our 1120 four-dimensional results define a database of energy information as a function of the θ_1 , θ_2 , α , and β angles. Thus, we can gain original information just by selecting the appropriate combination of angles. For instance, this database permits the user to analyze the dependence of the barrier to rotation on the α and β angles. Therefore, we select all the θ_1 and θ_2 data for $\alpha = \beta$ values. For the $\alpha = \beta = 120^\circ, 121^\circ, 122^\circ, 123^\circ$ and 124° values we obtain barriers of $231 \text{ cm}^{-1}, 260 \text{ cm}^{-1}, 290 \text{ cm}^{-1}, 322 \text{ cm}^{-1}$ and 357 cm^{-1} , respectively. Thus, the barrier, H , increases with the value of the angles. The variation fits extremely well to a quadratic form:

$$H(\text{cm}^{-1}) = 9517.89 - 182.74\alpha - 0.80\alpha^2 \quad \text{with } R^2 = 0.99998, \text{ RSS} = 0.019 \quad (24)$$

The increase of the barrier can be explained by an increase of the steric repulsion between the methyl groups, which approach each other when the α and β angles increase.



Table III. Barriers to methyl rotation in acetone as determined by different theoretical and experimental methods. All data in cm^{-1} .

Method	Barrier (cm^{-1})
MP2/6-311G (2d, p) ^a	270–271
Rotational data [50]	274
Rotational data [56]	265
Rotational data [57]	272
Torsional data [58]	290
MP2/6-31G (d, p) [53]	267
Rotational + Torsional data [59]	251

^aThis work.

5. CONCLUSIONS

This work considers the suitability of an Internet-based computational Grid for tackling computationally intensive parameter sweep problems. In particular, the mapping of molecular potential energy hypersurfaces is considered. Thus, we compare the computation of large sets of independent jobs in an Internet-based Grid of computers (coarse-grained approach), with the execution on a parallel multiprocessor system (fine-grained approach). The performance models developed show that the performance of the coarse versus the fine-grained approach increases proportionally to the number of jobs. On the other hand, we find that for time-consuming jobs the relative time gain between the coarse and the fine approach is independent of the number of calculations. Even for a 90% of parallelizability, and as few as ten processors, we have a gain of 50%. Thus, the implementation of a parameter sweep problem, such as the mapping of molecular potential energy hypersurface, at the Grid level results a very efficient process. The performance models are contrasted by considering a set of electronic structure calculations on the acetone molecule. The results fit closely the proposed expressions for the parallel performance at the fine-grained level and for the performance of the coarse versus the fine-grained approach.

The practical viability of the mapping of molecular potential energy hypersurfaces on an Internet-based Grid of computers is shown by developing and testing a new Grid application. The application automatically generates molecular structures as a function of the molecular coordinates considered. In addition, the structures are submitted to the Grid for computation in an electronic structure package. The results are filtered in the worker nodes and assembled in the client node. Considering the acetone molecule, a total of 1120 molecular structures are generated. The results show that it is possible to use the Grid to compute thousands of molecular structure data in a more than reasonable time.

As a practical example, the application to acetone permits the mapping map of potential energy hypersurfaces in two and four dimensions. The barrier to methyl rotation obtained from these results is in good agreement with the experimental data. The four-dimensional hypersurface shows that the barrier to methyl rotation increases with the CCO angles. This fact can be attributed to sterical hindrance between the methyl groups. The variation follows a quadratic form extremely well.

01 **ACKNOWLEDGEMENTS**

02

03 The authors wish to thank the Barcelona Supercomputing Center-Centro Nacional de Supercomputación (BSC-
04 CNS) and the Facultad de Ciencias Químicas (Laboratorio de Química Teórica) of the Universidad Autónoma de
05 Puebla (Mexico) for the use of the Kadesh and Popocatepetl clusters, respectively.

Q3

06

07 **REFERENCES**

08

09

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

1. Foster I, Kesselman C, Tuecke S. The anatomy of the Grid. Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications* 2001; **15**:200–222.
2. Foster I. THE GRID: Computing without bounds. *Scientific American* 2003; **288**:78–86.
3. Foster I, Kesselmann C (eds.). *The Grid2: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann: San Francisco, CA, 2004.
4. Badia RM, Labarta J, Sirvent R, Pérez JM, Cela JM, Grima R. Programming Grid applications with Grid SuperScalar. *Journal of Grid Computing* 2003; **1**:151–170.
5. GRID superscalar Website. <http://www.bsc.es/grid> [30 July 2006].
6. Sirvent R, Pérez JM, Badia RM, Labarta J. Automatic Grid workflow based on imperative programming languages. *Concurrency and Computation: Practice and Experience* 2006; **18**(10):1169–1186. DOI: 10.1002/cpe.1000.
7. Gridway. <http://www.gridway.org/> [30 June 2006].
8. Kacsuk P, Dózsa G, Kovács J, Lovas R, Podhorszki N, Balaton Z. P-GRADE: A Grid programming environment. *Journal of Grid Computing* 2003; **1**:171–197.
9. AstroGrid Project. <http://www.astroGrid.org/> [30 June 2006].
10. LHC computing Grid Project. <http://lcg.web.cern.ch/LCG/> [30 June 2006].
11. The Extensible Computational Chemistry Environment (ECCE). <http://ecce.emsl.pnl.gov/> [30 June 2006].
12. The Computational Chemistry Grid. <https://www.gridchem.org/> [30 June 2006].
13. Maheswaran M, Ali S, Siegel HJ, Hensgen D, Freund R. Dynamic matching and scheduling of a class of independent tasks onto heterogeneous computing systems. *Proceedings of the 8th Heterogeneous Computing Workshop (HCW'99)*, April 1999.
14. Casanova H, Legrand A, Zagorodnov D, Berman F. Heuristics for scheduling parameter sweep applications in Grid environments. *Proceedings of the 9th Heterogeneous Computing Workshop (HCW00)*, May 2000; 349–363.
15. Huedo E, Montero RS, Llorente IM. Experiences on adaptive Grid scheduling of parameter sweep applications. *Proceedings of the 12th Euromicro Conference on Parallel Distributed and Network-Based Processing (PDP'04)*, 2004; 28.
16. Hollas JM. *Modern Spectroscopy* (3rd edn). Wiley: New York, 1996.
17. *Vibrational-Rotational Spectroscopy and Molecular Structures and Dynamics (Advanced Series in Physical Chemistry, vol. 9)*, Papoušek D (ed.). World Scientific: Singapore, 1997.
18. Szabo A, Ostlund NS. *Modern Quantum Chemistry*. McGraw-Hill: New York, 1989.
19. Muñoz-Caro C, Niño A, Senent ML. Theoretical study of the effect of torsional anharmonicity on the thermodynamic properties of methanol. *Chemical Physics Letters* 1997; **273**:135–140.
20. Gamess. <http://www.msg.ameslab.gov/GAMESS/GAMESS.html> [30 June 2006].
21. NWChem. <http://www.emsl.pnl.gov/docs/nwchem/nwchem.html> [30 June 2006].
22. Dalton. <http://www.kjemi.uio.no/software/dalton/dalton.html> [30 June 2006].
23. Jensen F. *Introduction to Computational Chemistry*. Wiley: New York, 1998.
24. Sudholt W, Baldrige KK, Abramson D, Enticott C, Garic S. *Applying Grid Computing to the Parameter Sweep of a Group Difference Pseudopotential (Lecture Notes in Computer Science, vol. 3036)*, 2004; 148–155.
25. Nimrod/G. <http://www.csse.monash.edu.au/~david/nimrod/index.htm> [30 June 2006].
26. Culler DE, Karp RM, Patterson D, Sahay A, Santos EE, Schauser KE, Subramonian R, von Eicken T. LogP: A practical model of parallel computation. *Communications of the ACM* 1996; **39**(11):78–85.
27. Alexandrov A, Ionescu MF, Schauser KE, Scheiman C. LogGP: Incorporating long messages into the LogP model— one step closer towards a realistic model for parallel computation. *Proceedings of the 7th Annual Symposium on Parallel Algorithms and Architecture (SPAA'95)*, Santa Barbara, CA, July 1995; 95–105.
28. Frank MI, Agarwal A, Vernon MK. LoPC: Modeling contention in parallel algorithms. *Proceedings of the 6th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Las Vegas, NV 18–21 June 1997; 276–287.
29. Moritz CA, Frank MI, LoGPC: Modeling network contention in message-passing programs. *Proceedings of the 1998 ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*, Madison, WI, 1998; 254–263.



- 01 30. Ino F, Fujimoto N, Hagihara K, LogGPS: A parallel computational model for synchronization analysis. *Proceedings of the*
02 *8th ACM SIGPLAN Symposium on Principles and Practices of Parallel Programming*, Snowbird, UT, 2001; 133–142.
- 03 31. Cameron KW, Ge R. Predicting and evaluating distributed communication performance. *Proceedings of the 2004*
04 *ACM/IEEE Conference of Supercomputing (SC2004)*, Pittsburgh, PA, 6–12 November 2004; 43–43.
- 05 32. Le TT, Rejeb J. A detailed MPI communication model for distributed systems. *Future Generation Computer Systems* 2006;
06 **22**:269–278.
- 07 33. Chronopoulos AT, Penmatsa S, Xu J, Ali S. Distributed loop-scheduling schemes for heterogeneous computer systems.
08 *Concurrency and Computation: Practice and Experience* 2006; **18**(7):771–785.
- 09 34. Sokolowski P, Grosu D, Xu C-Z. Analysis of performance behaviors of Grid connected clusters, *Performance Evaluation*
10 *of Parallel, Distributed and Emergent Systems*, Ould-Khaoua M, Min G (eds.). Nova Science: Hauppauge, NY, 2006.
- 11 35. Amdahl GH. Validity of the single processor approach to achieving large scale computing capabilities. *Proceedings of the*
12 *AFIPS Conference*, Spring Joint Computer Conference, Atlantic City, NJ, 30 April 1967; 483–485.
- 13 36. BBN Advanced Computers Inc. *Mach 1000 Fortran Compiler Reference* (revision 1.0 edn), November 1998. BBN
14 Advanced Computer Inc: Cambridge, MA, 1998.
- 15 37. Polychronopoulos CD, Kuck DJ. Guided self-scheduling: A practical scheduling scheme for parallel supercomputers. *IEEE*
16 *Transactions on Computers* 1987; **36**:1425–1439.
- 17 38. Hummel SF, Schonberg E, Flynn LE. Factoring: A method for scheduling parallel loops. *Communications of the ACM*
18 1992; **35**(8):90–101.
- 19 39. Tzen TH, Ni LM. Trapezoid self-scheduling: A practical scheduling scheme for parallel compilers. *IEEE Transactions on*
20 *Parallel and Distributed Systems* 1993; **4**(1):87–98.
- 21 40. Chronopoulos AT, Andonie R, Benche M, Grosu D. A class of loop self-scheduling for heterogeneous clusters. *Proceedings*
22 *of the 3rd IEEE International Conference on Cluster Computing (CLUSTER 2001)*, Newport Beach, CA, 8–11 October
23 2001; 282–291.
- 24 41. Niño A, Muñoz-Caro C, Moule DC. Wagging and torsion vibronic structure in the $T_1 \leftarrow S_0$ electronic spectrum of
25 acetaldehyde. *Journal of Physical Chemistry* 1994; **98**:1519–1524.
- 26 42. Niño A, Muñoz-Caro C, Moule DC. Three-dimensional vibrational study of the coupling between the methyl torsion and
27 the molecular frame in the S_0 state of acetaldehyde. *Journal of Physical Chemistry* 1995; **99**:8510–8515.
- 28 43. Reyes S, Niño A, Muñoz-Caro C. Customizing clustering computing for a computational chemistry environment. The case
29 of the DBO-83 nicotinic analgesic. *Journal of Molecular Structure: THEOCHEM* 2005; **727**:41–48.
- 30 44. <http://www.rocksclusters.org/wordpress/> [30 June 2006].
- 31 45. Schmidt MW *et al.* General atomic and molecular electronic structure system. *Journal of Computational Chemistry* 1993;
32 **14**:1347–1363.
- 33 46. <http://www.globus.org/> [30 June 2006].
- 34 47. <http://www.openpbs.org/> [30 June 2006].
- 35 48. <http://www.rocksclusters.org/roll-documentation/sge/4.1/> [30 June 2006].
- 36 49. Kannan S, Roberts M, Mayes P, Brelsford D, Skovira JF. *Workload Management with LoadLeveler*. IBM RedBooks, IBM,
37 2001.
- 38 50. Swalen JD, Costain CC. Internal rotation in molecules with two internal rotors: Microwave spectrum of acetone. *Journal*
39 *of Chemical Physics* 1959; **31**:1562–1574.
- 40 51. Smeyers YG, Bellido MN. Internal dynamics of nonrigid molecules. I. Application to acetone. *International Journal of*
41 *Quantum Chemistry* 1981; **19**:553–565.
- 42 52. Ozkabak AG, Goodman L. Skeletal flexing during methyl rotation in small dimethyl molecules. *Chemical Physics Letters*
43 1991; **176**:19–26.
- 44 53. Smeyers YG, Senent ML, Botella V, Moule DC. An ab initio structural and spectroscopic study of acetone. An analysis of
the far infrared torsional spectra of acetone- h_6 and d_6 . *Journal of Chemical Physics* 1993; **98**:2754–2767.
54. Armstrong M. *Basic Linear Geostatistics*. Springer: Berlin, 1998.
55. Golden Software Inc. *Surface Mapping System*, Surfer (V. 8.0), 11 February 2002. Golden Software Inc: Golden, CO, 2002.
56. Peter R, Dreizler H. Das mikrowellenspektrum von acetone im torsionsgrundzustan. *Zeitschrift für Naturforschung A* 1965;
20:301.
57. Nelson R, Pierce L. Microwave spectrum, structure, and barrier to internal rotation of acetone. *Journal of Molecular*
Spectroscopy 1965; **18**:344–352.
58. Groner P, Guirgis GA, Durig JR. Analysis of torsional spectra of molecules with 2 internal C_{3v} rotors. 24. High-resolution
far infrared-spectra of acetone- d_0 , acetone- d_3 , and acetone- d_6 . *Journal of Chemical Physics* 1987; **86**:565–568.
59. Groner P. Experimental two-dimensional torsional potential function for the methyl internal rotors in acetone. *Journal of*
Molecular Structure 2000; **550–551**:473–479.