Hardware environment for CSLabGrid: Reaching maximum efficacy of computations in structural biology and bioinformatics

Nikolay Pydiura, Karpov Pavel, Blume Yaroslav

Institute of Food Biotechnology and Genomics NAS of Ukraine, 2A Osipovskogo str., Kyiv, Ukraine

pydiura@gmail.com, karpov.p.a@gmail.com, cellbio@cellbio.freenet.viaduk.net

Keywords

Virtual organisation, computational biology, scientific computing, molecular dynamics, bioinformatics, performance, grid.

1 Introduction

The number and complexity of computational biology tasks is constantly growing. Under such conditions the scientists are forced to take the best efforts to reach better performance with limited available resources. One of the factors that influence the performance of biological tasks is the software simulation parameter sets, and first of all the ones responsible for parallelisation. The other factor is the chosen hardware configuration. Computational biology tasks vary greatly in specificity and hardware requirements. Biological tasks include not only computations that require high parallelisation like molecular dynamics (MD) simulations of large systems but also tasks that are usually executed in a single thread. For instance, like sequence analysis, homology modelling, model evaluation, molecular docking, virtual screening and also MD simulations of the smaller systems in microsecond intervals.

Two common different resource intensive biological tasks are homology modelling and molecular dynamics simulation. Long MD simulations in hundreds of nanoseconds and microsecond intervals can greatly contribute to our knowledge of the dynamics of the proteins and protein complexes. One of such important objects is cytoskeleton and cytoskeleton-related proteins which are the targets of investigation within the CSLabGrid virtual organisation.

Due to the variety and heterogeneity of biological tasks and hardware configurations available on the clusters of the Ukrainian National Grid (http://ung.in.ua/) it is essential prior to each new series of experiments to make regular tests for the estimation of the optimal hardware configuration and software parameters. The performance of the latest server systems available on the market is sufficient for the solution of these tasks. However, the maximum efficacy of the calculations requires hardware and software parameters adjustment. Definitely, looking on the market for the best value hardware for a scientific cluster one should take into account the specificity of the tasks intended to be solved. We have also analysed the performance value for biological tasks of the Intel server CPUs available on the market.

2 Results

2.1 The test system Gromacs MD performance estimation

To evaluate the performance we used a previously constructed homology model of Arabidopsis thaliana alpha beta tubulin dimer. Using Gromacs software package the model was solvated. Na+ and Cl- ions were added to simulate the physiological ionic strength (0.15 M NaCl). After that, an energy minimization and a short position restrained run were performed. The final system configuration contains 146,101 atoms in 45,016 residues. The volume of the system box is 1444 nm^3 and the number of the SOL molecules is 44,119. The system is a “good” target for MD simulations as it is quit symmetrical (figure 1).

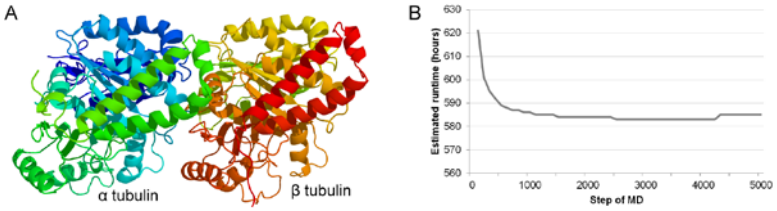

Fig. 1. (A) – the homology model of *Arabidopsis thaliana* alpha beta tubulin dimer. (B) – the plot of the Gromacs computation time estimation for the 100ns MD simulation of test system on the initial steps of MD simulation. The estimation is calculated basing on the current computational speed. The first plateau of the estimation time value is reached on average in 2000 to 3000 steps and the final computation time may only slightly differ from this estimate (data not shown). The preliminary estimate for the presented run is around 585 hours

To analyse the performance of the hardware configurations we used the Gromacs mdrun program runtime estimation data written to the standard output with the mdrun –v option. The figure 1A shows that in first 5,000 steps of MD simulation the estimate value reaches the plateau. For presented experiments we collected the estimate values at the step 5,000 from five test runs. In all considered configurations the execution of 5,000 steps took less than 5 minutes.

2.2 The influence of the domain decomposition on the performance of MD computations

Domain decomposition (DD) is a procedure implemented by Gromacs mdrun program. DD allows effectively dividing the volume of the system into cells. Each such cell is computed on a separate CPU thread. It is crucial ti divide system into cells as equally as it is possible taking into account the system constraints together with inhomogeneity provided by the shape and distribution of the biological molecules in the system. The load balance and, thus, the performance depend greatly on the correct domain decomposition. For our test system, we studied different 3D domain decomposition grids for 24 threads and have identified that automatic 3D decomposition grid (X Y Z) 6 4 1 is not the optimal one due to 1% performance loss (Fig 3A,B).

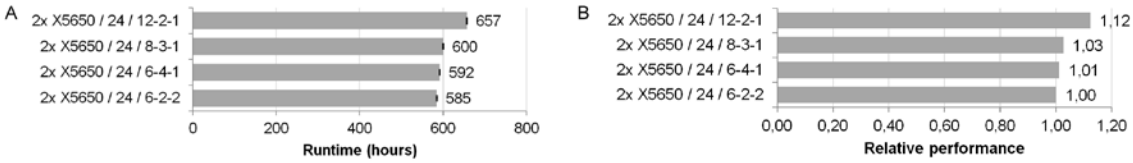

Fig. 2. MD performance results obtained with different domain decomposition grids for 24 threads. (A) – absolute execution time (hours). (B) – relative execution time

In overall, for our system an incorrectly selected decomposition grid may lead to up to 12 % longer computation time for domain decomposition into 24 threads (figure 2 A,B). If a system is more inhomogeneous than the presented test system the loss of performance could be much higher. This suggests the need to adjust the domain decomposition in each concrete case.

2.3 Domain decomposition depending on the different numbers of available threads

Electrostatic and covalent bonds and other restraints influence domain decomposition as for each system they predetermine a minimum further indivisible size of the cell. Due to these constraints, with an increase in number of parts into which the system is split, considerably increase the influence of the inhomogeneity provided by the shape and distribution of the biological molecules in the system. Heterogeneity of the cells volumes can be observed as “vol” value of Gromacs mdrun log file. The “vol” value shows the smallest volume of a domain decomposition cell in relation to the average volume. This value can be from 0 to 1.

In our experiments, active option of dynamic load balancing (every 10-th step) used 24 threads of the maximum performance (the loss of time due to load imbalance was – 0%). It was obtained with the “vol” values 0.48 and 0.72 for domain decomposition grids 6 4 1 and 6 2 2 correspondingly.

We compared the efficacy of the domain decomposition and MD computation in 24 and 32 threads. The DD grids were 6 4 1 and 8 4 1 correspondingly. Using 32 threads (DD grid 8 4 1) the heterogeneity of the load on the threads was 4.6%, and the time lost due to load imbalance was 1.5%. At the same time, on 24 threads (DD grid 6 4 1) the load imbalance and loss of performance were negligible. During the computations of the smaller or more inhomogeneous systems on 32 threads the performance losses due to load imbalance could be higher and level the performance gain from the greater number of threads used.

Thus, for effective and high-productive molecular dynamics calculations of the small (less than 50,000 atoms) and medium systems (50,000 – 150,000 atoms) with the number of threads used 24, 32 or more, individual core productivity becomes more important than the number of cores.

According to PassMark performance test – CPU Mark (http://www.cpubenchmark.net/) productivity of the Intel Xeon X5650 / 6 cores @ 2.67GHz /12 threads and Intel Xeon E5-2650 / 8 cores @ 2.00GHz / 16 threads is 7,419 and 10,966 marks respectively. Namely, the theoretical performance of the Intel Xeon E5-2650 CPU is 48% higher. In current MD simulation tasks we observed the 22% and 14% difference for runs on single CPUs and on nodes with two identical Intel Xeon X5650 or Intel Xeon E5-2650 CPUs respectively (figure 3A).

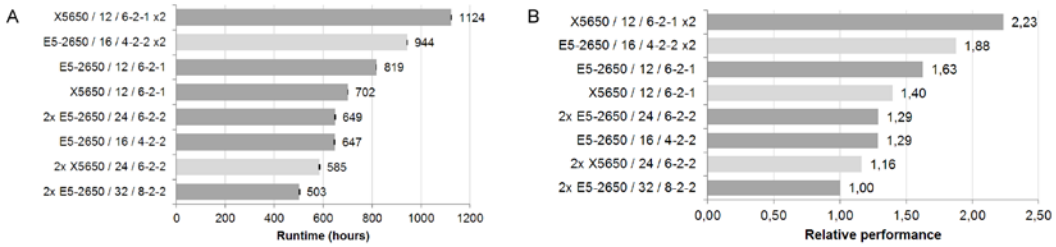

Fig. 3. (A) – estimated computation time in hours of 100 ns MD simulations of the test systems using different hardware configurations: E5-2650 and X5650 CPUs nodes. (B) – configuration performance relatively to the performance of fastest tested configuration of 2xE5-2650 / 32 threads with 3D domain decomposition 8 2 2

Two improve the overall performance of the MD calculations a strategy of utilizing two separate CPUs of the two-CPU node for two different independent tasks may be applied. However, the performance loss as compared to the case of single task per node is inevitable. Nevertheless, the simultaneous execution of a one task per each CPU of the two-CPU server gives a 4% and 7% benefit on X5650 and E5-2650 CPU nodes in compare to consequent execution of two same tasks using all threads of these nodes (E5-2650 / 32 / 8-2-2 versus E5-2650 / 16 / 4-2-2 x2 and X5650 / 24 / 6-2-2 versus X5650 / 12 / 6-2-1 x2 on figure 3A). Taking into account that in the case of consequent execution of the tasks the (re)submission of the new task to the node may require some time, the overall benefit of the strategy of simultaneous execution of the tasks becomes significant. But still, utilisation of such strategy causes additional load on the other system components of the server which will become a bottleneck under such conditions. The same tasks demonstrate 46% to 60% (figure 3A) increase in performance while running on a single CPU node than on each of the same CPUs of the dual CPU node. But taking into account the additional 56 % (figure 5) cost of two separate one E5-2650 CPU servers comparing to one dual CPU node this concern is negligible.

We have also estimated the theoretical performance of the single core of X5650 and E5-2650 CPUs in MD simulations. We have executed the same MD task using 12 threads of each CPU – X5650 / 12 / 6-2-1 and E5-2650 / 12 / 6-2-1 on figure 3A. As it can be seen from the figure 3A the computational performance of each core of Intel Xeon X5650 CPU running at 2.67GHz in this concrete task was about 10% to 14% higher than that of Intel Xeon E5-2650 CPU with 2.00GHz core (2.3 – 2.5GHz at Turbo Boost).

So, MD calculations of the medium and smaller MD systems may be effectively executed even on a single top level CPU. For such calculations the performance of a separate CPU thread is more important than the number of CPU cores. In turn, for larger MD systems (more than 200,000 atoms), such as for example the model of tubulin tetramer in water (400,000 – 500,000 atoms), the benefits from such parallelization are evident. Calculations of such MD systems will revealthe full potential of the latest CPUs with 8 and more cores.

2.4 Estimation of the value of the modern Intel Xeon CPUs for concrete examples of biological computational tasks

Biological computational tasks include not only tasks that require high parallelisation (like MD simulations of large systems) but also tasks that are usually executed in a single thread like bioinformatics sequence analysis, homology modelling, model evaluation, molecular docking, virtual screening and also MD simulations of smaller systems in microsecond intervals. To test the efficacy of the available processors we have performed homology modelling in Modeller 9v8 software. The calculation of a single homology model of the ensemble in this application cannot be run in parallel. So, we used the model generation time as an estimate of such single threaded task performance. The test modelling system is a human fibrinogen multi chain protein molecule consisting of 1981 amino acid residues (15,900 atoms). The model is subjected to an extensive modelling procedure (a.library_schedule = autosched.slow) and a very thorough MD optimisation (a.md_level = refine.slow). Ten models were generated in each run and average model generation time is presented on the figure 4.

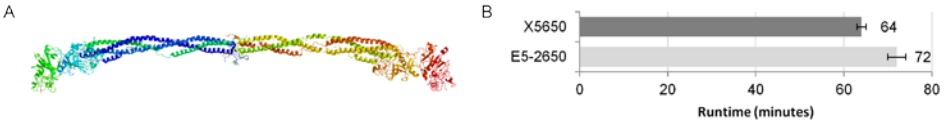

Fig. 4. (A) – homology model of Homo sapiens fibrinogen. (B) – time in minutes used by single threads of Intel Xeon E5-2650 @ 2.0 GHz and X56502.67 GHz CPUs to produce a single model of the test Modeller9v8 homology modelling task

Execution time estimation of a single thread test homology modelling task (figure 4) suggests that the X5650 CPU (2.67 GHz) of previous generation shows the performance 16% high as that of the last generation CPU E5-2650 (2.0 GHz). As most biological tasks are not yet optimised for specific CPU architectures, there is no need going deep into architectural differences between these two CPUs responsible for such results. By now, the key factor substantial for the performance in such small biological tasks remains the CPU frequency. However, in the nearest future it is possible that biological software packages will have better support of the latest CPUs architecture, especially the commercial ones.

Based on obtained data we estimated the value of modern Intel Xeon CPUs and servers for the performance of concrete biological computational tasks – figure 5.

CPUs				Servers		
Model	**Number of cores/threads**	**Main frequency**	**CPU price, USD**	**Number of CPUs**	**Price, USD**	**"Thread price", USD**
E3-1270 V2	4 / 8	3.50GHz	320	1	2800	350
E5-2620	6 / 12	2.00 GHz	575	2	3000	125
E5-2630	6 / 12	2.30 GHz	700	2	3800	160
E5-2640	6 / 12	2.50 GHz	1050	2	5500	230
E5-2650	8 / 16	2.00 GHz	1250	1	4600	289
E5-2650	**8 / 16**	**2.00 GHz**	**1250**	**2**	**5900**	**184**
E5-2660	8 / 16	2.20 GHz	1650	2	8400	263
E5-2670	8 / 16	2.60 GHz	2000	2	8600	269
E5-2690	8 / 16	2.90 GHz	2600	2	10600	331
X5650	**6 / 12**	**2.67 GHz**	**1350**	**2**	**5200**	**217**

Fig. 5. Characteristics and average prices of CPUs and servers in Kyiv, Ukraine. Tested systems are shown in bold. Prices are obtained using hotline.ua web-site. "Thread price" is calculated as node price divided by number of threads

It is often required to execute hundreds and thousands of small biological tasks for different objects (e.g. docking and screening of certain groups of tubulin microtubule (de)stabilising compounds versus tubulins of different species in our VO CSLabGrid – http://infrastructure.kiev.ua/en/monitoring/47/). In this case the higher single thread performance of the CPU is essential. However, the latest top higher frequency E5-26XX CPUs are pretty costly. Taking this fact into account, it might be practically useful and economically sound to have at the disposal of the biological scientific VOs at the same time both servers with the high frequency / high single thread performing CPUs and the two-CPU nodes with larger number of cores but lower frequency, the first ones for the computation of the single threaded tasks and the second ones – for large MD systems. Obviously, that for a virtual organisation users it is not important where exactly within the Ukrainian National Grid (UNG - http://ung.in.ua/) infrastructure these resources reside as long as grid users have access to them.

The server for small tasks could be equipped with 4 core / 8 thread Intel Xeon E3-12xx V2 CPUs which to date have the highest single thread performance among the CPUs available on the market according to PassMark – CPU Mark. Namely, its single thread performance may be 58% higher than that of the tested here Intel Xeon E5-2650. This may give molecular biological or any other computational tasks which utilize up to 8 threads a great performance boost.

In turn, two-CPU nodes could be equipped with Intel Xeon E5-26xx CPUs, e.g. E5-2650 which was tested here. At the same time, for almost double the "thread price" of E5-2650 one can get Intel Xeon E5-2670 with 8 cores / 16 threads running at 2.6 GHz. This could be a good universal solution suitable for bioinformatics tasks of all scales, though, slightly overpriced. According to PassMark – CPU Mark, E5-2690 is only 23% faster than E5-2650 but 60% more expensive.

3 Acknowledgments

All molecular computations were carried out with the assistance of The Ukrainian National Grid (UNG) – (http://ung.in.ua/), and we would like to acknowledge Targeted State Scientific and Technological Program on the Development and Implementation of Grid-Technologies for 2009-2013 (http://grid.nas.gov.ua/).