


Estimating the Performance of BLAST runs on the EGEE Grid

Abel Carrión

Ignacio Blanquer

Vicente Hernández


- **The problem.**
- **Factors affecting the performance.**
- **Experiments on the grid.**
- **The Performance.**
- **Execution model.**
- **Conclusions and further work.**



- **Sequence alignment is an key operation in Bioinformatics**
 - It involves computing the comparison of proteomic and genomic samples with respect to annotated databases.
- **It is a part of many Bioinformatics pipelines**
 - Used to search for homologous in the study the functionality of different genes and regions.
 - Used in the phylogenetic taxonomy.
- **There are many tools developed in the literature**
 - Based on the Smith-Waterman transform (e.g. BLAST).
 - Based on Hash Tables (e.g. SSAHA, BLAT).
 - Based on Burroughs-Wheeler Transform (e.g. BWA, Bowtie).
 - And combinations of them (e.g. SSAHA2).

- **BLAST (Basic Local Alignment Search Tool) is the most widely used tool for performing the alignment of any length novel sequences against the ones contained in a determined database.**
 - Although it could be inefficient for many cases, It has a proven reputation.
- **Because a normal use case entails the alignment of millions of sequences, this kind of experiments are very computationally intensive (it demands years of CPU computation).**

- Problem is massive parallel and fits the requirements for using efficiently Grid infrastructures.
- The parallelization process follows the High-Throughput paradigm:
 - Segmenting the input data file into several chunks which are aligned in independent computation nodes.



- **Two factors are key for the general performance**
 - A good selection of the resources for computing and storage.
 - A good partition strategy.
- **This not only affects the response time, but also the failure ratio**
 - Queues have a limitation in the maximum job executing time.
 - Fault-tolerance automatic resubmission also needs to know if a job is executing slowly or simply it is blocked.
- **Thus, a key issue when performing the scheduling of thousands of jobs is estimating the response time of the tasks.**
 - Create a model to estimate the execution time of a job.

- **Quasi-Deterministic Factors**

- Application dependant.
 - Input data file size.
 - DB file size.
 - Similarity.
 - Number of hits.
- Resource dependant.
 - SPECint (SI00).
 - SPECfp (SF00).
 - Memory.

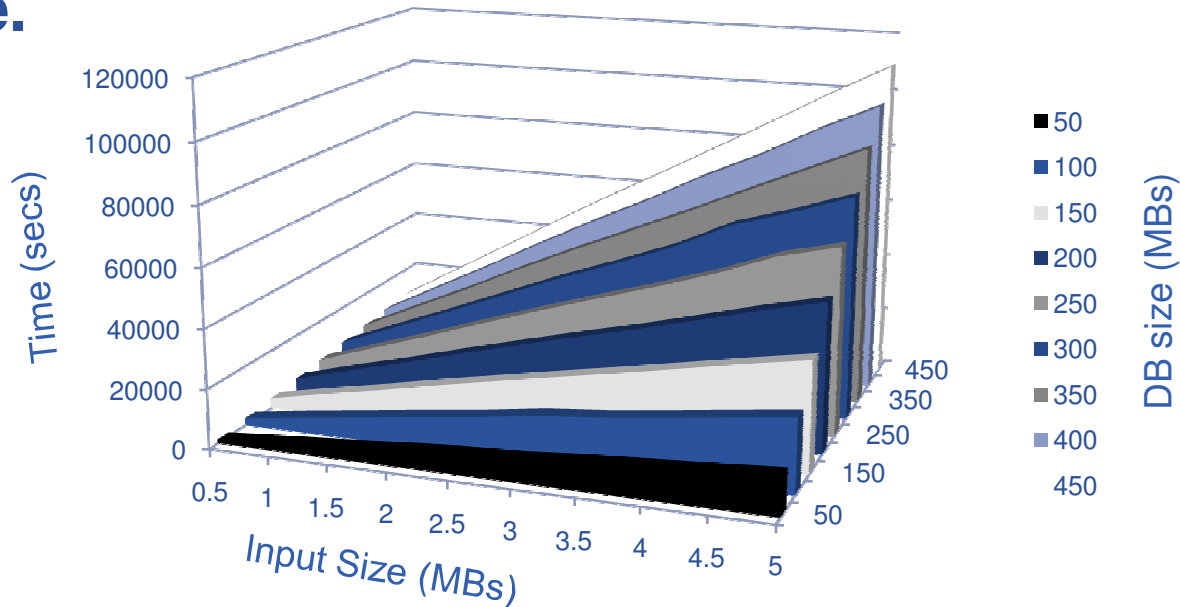
Obtained through experiments

- **Undeterministic Factors**

- Load dependant
 - Queue size
 - Average waiting time
- Site dependant
 - Site availability.
 - Specific job failure rates.


Not yet covered in the study

- The purpose of this experiment was to analyse the influence of the data size (input file and database).
- Using the UniProt database, several files with different sizes were generated and executed in the same machine.





- The results show that the input file size and database size have a direct linear impact on the response time.


- The heuristic nature of BLAST accelerates the comparison of two clearly unrelated sequences
- To check the influence of the similarity on the searches, three new versions of the UniProt database were produced, replacing 1%, 5% and 10% of their contents respectively.
- As it can be seen, the response time is independent from the similarity between sequences.
 - This factor will not be considered in the final performance model.



- An experiment has been executed in different computers for different values of the “BLAST Hits” argument
 - Values range from 5000 to 25000 (step of 2500) in 30 CEs
 - Soil NCBI Metagenome.
 - 270 Jobs with a duration from 130 to 200 minutes.
 - Although the output increases accordingly, no effect is observed on the response time nor the failure rate.



- The GlueSchema includes the values of Benchmarks SpecInt 2000 and SpecFloat 2000.
- 9 Different job types have been executed on 19 sites where this information was published.




- **The SpecInt and SpecFloat benchmarks do not seem to be correlated with actual performance for BLAST**
 - Correlation index ranges from 0.08 to 0.55 with poor values of significance (up to 77% of randomness).

	T1	T2	T3	T4	T5	T6	T7	T8	T9
Pearson	-0,433	-0,309	-0,533	-0,42	-0,521	-0,549	-0,413	0,079	0,55
p	0,052	0,175	0,014	0,06	0,017	0,011	0,081	0,772	0,089

- **This might be caused by different factors**
 - Unsuitability of the benchmark for BLAST.
 - Lack of accuracy of the benchmark, which is not computed in many sites but obtained from tables.
- **Moreover, the values of the benchmarks are not always published.**
- **A new estimator is needed.**

- **Ratio of average performance with respect to a fixed node**
 - Same tests are repeated in all nodes for different data sizes and speed-ups/downs are computed.
 - A single coefficient is obtained from the sum of all the computing times of all (the same number and type) of jobs.
- **Obviously, much more correlation is shown (high significance, always above 98%, but generally above 99,99%).**

	T1	T2	T3	T4	T5	T6	T7	T8	T9
Pearson	0,888	0,788	0,957	0,897	0,978	0,96	0,889	0,691	0,73
p	1,86E-07	3,33E-05	6,65E-11	8,72E-08	2,10E-13	3,12E-11	8,33E-07	0,004	0,015



- **Current conclusions**


- Direct dependence on the Input data size.
- Direct dependence on the database size.
- Direct dependence on the inverse of the process speed factor.
- No dependency on the similarity or number of hits.
- Low dependence on the memory, except for saturations.


- **Proposed model**

$$T = \frac{P_{Sp}}{Sp} P_{Tm} \left(\frac{T_{inp}}{T_{inp_basal}} \frac{T_{BD}}{T_{BD_basal}} \right) + K$$

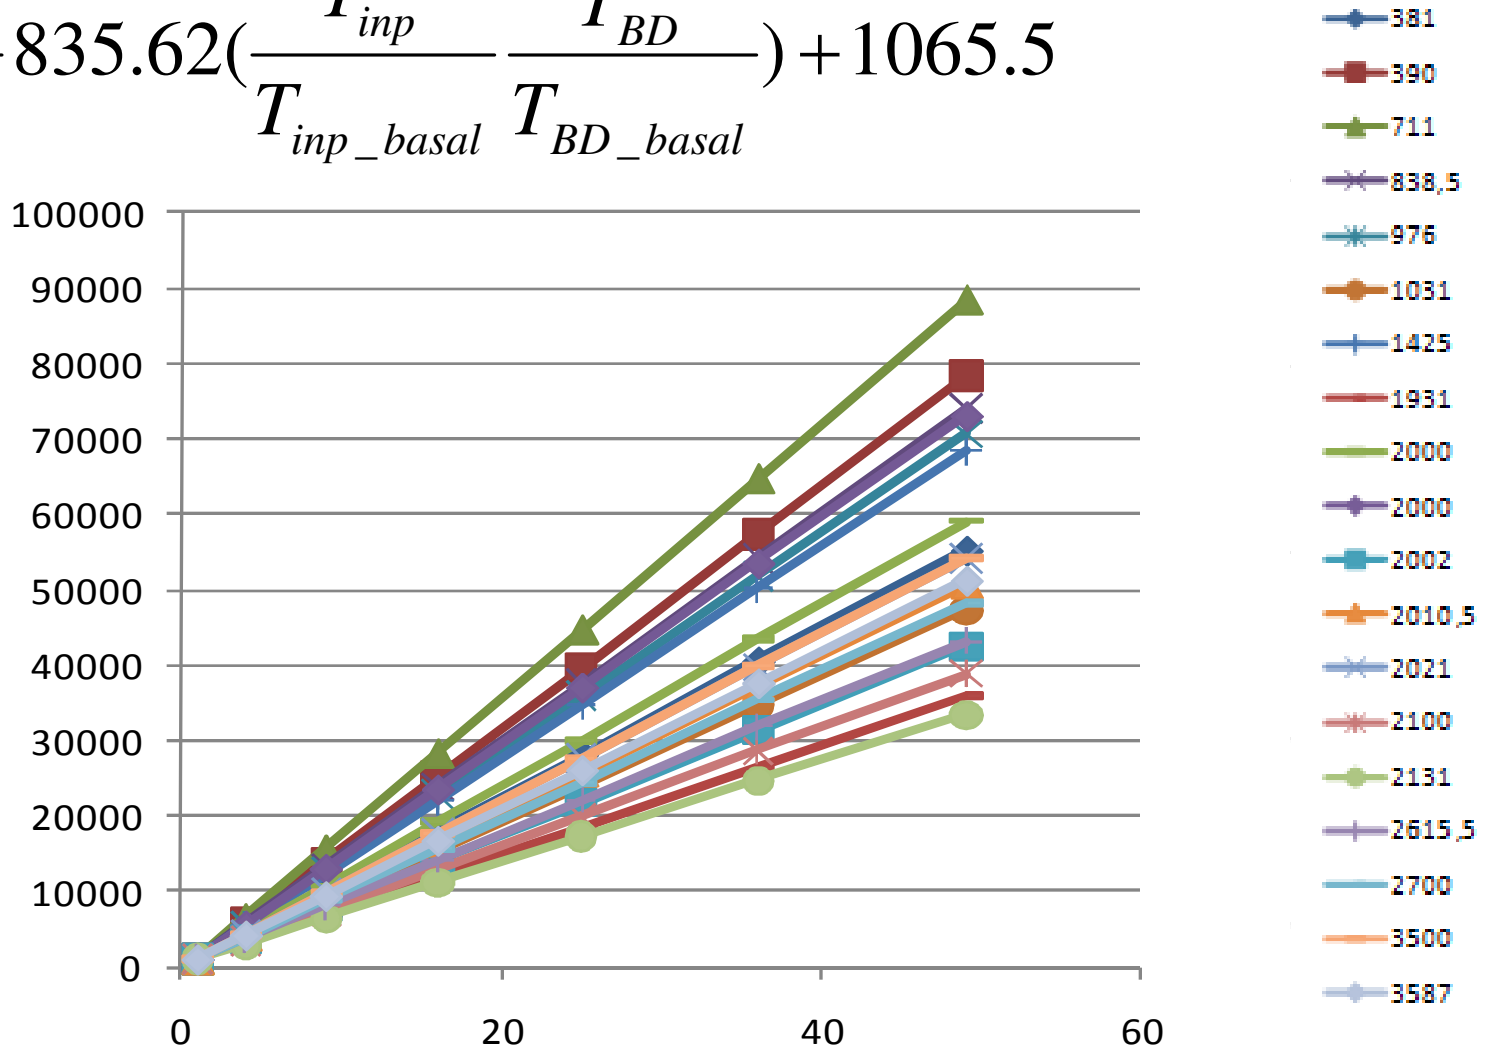
- **Where**

- P_{Sp} , P_{Tm} and K are the unknown parameters of the regression model.
- T_{inp} and T_{BD} are the fixed values for the size of data and database.
- Sp is the ratio between the response time in a reference site and each other site.
- T_{inp_basal} is 0,5 y T_{BD_basal} is 50.





$$T = \frac{1.35}{Sp} 835.62 \left(\frac{T_{inp}}{T_{inp_basal}} \frac{T_{BD}}{T_{BD_basal}} \right) + 1065.5$$



- **This work presents a performance model for estimating the response time of BLAST runs in the EGEE grid**
 - Direct dependence with the performance indicator and the data size of input and reference database.
 - No direct dependence with memory or blast hits.
- **Very important parameter for load balancing and pre-emptive resubmission.**
- **The work will be extended, introducing new parameters from other components**
 - Workload Manager System.
 - Local Resource Management Systems.
 - Bandwidth between CEs y SEs.