Optimizations to compute large correlation matrix onto GPU system of hybrid HPC clusters

Dany Tello1, Fouad Boumezbeur2, Victor Arslan1, Vincent Ducrot1, Pierre Léonard2, Bouziane Moumen2, Sébastien Monot1, Nicolas Pons2, Tarik Saidani1, Pierre Renault2, Sean Kennedy2, Mathieu Almeida2, S.Dusko Ehrlich2 and J.M. Batto2

1**AS Plus, 22 rue René Coche 92170 Vanves, France**

2**Institut MICALIS, INRA CRJ, Domaine de Vilvert, 78352 Jouy-en-Josas, France**

Contact : s.monot@asplus.fr, jean-michel.batto@jouy.inra.fr

In the MetaHIT project, large computations are performed to correlate quantitative information related to metagenomic samples. This program, named MetaProf, helps MetaQuant (INRA Micalis) in processing metagenomics analysis.

In order to achieve this computation within an acceptable time frame, we have developed a version of this program targeted for execution on GPU based clusters. Optimizations have been done to use in the most efficient way the capability of the GPU processor. The optimized program has been benchmarked on the hybrid part (GPU/nVidia M2090) of the GENCI supercomputer named Curie installed at TGCC (Très Grand Centre de Calcul du CEA - http://www-hpc.cea.fr/en/complexe/tgcc-curie.htm).

How to explore, the Metagenomics Human Intestinal Tract (MetaHIT)?

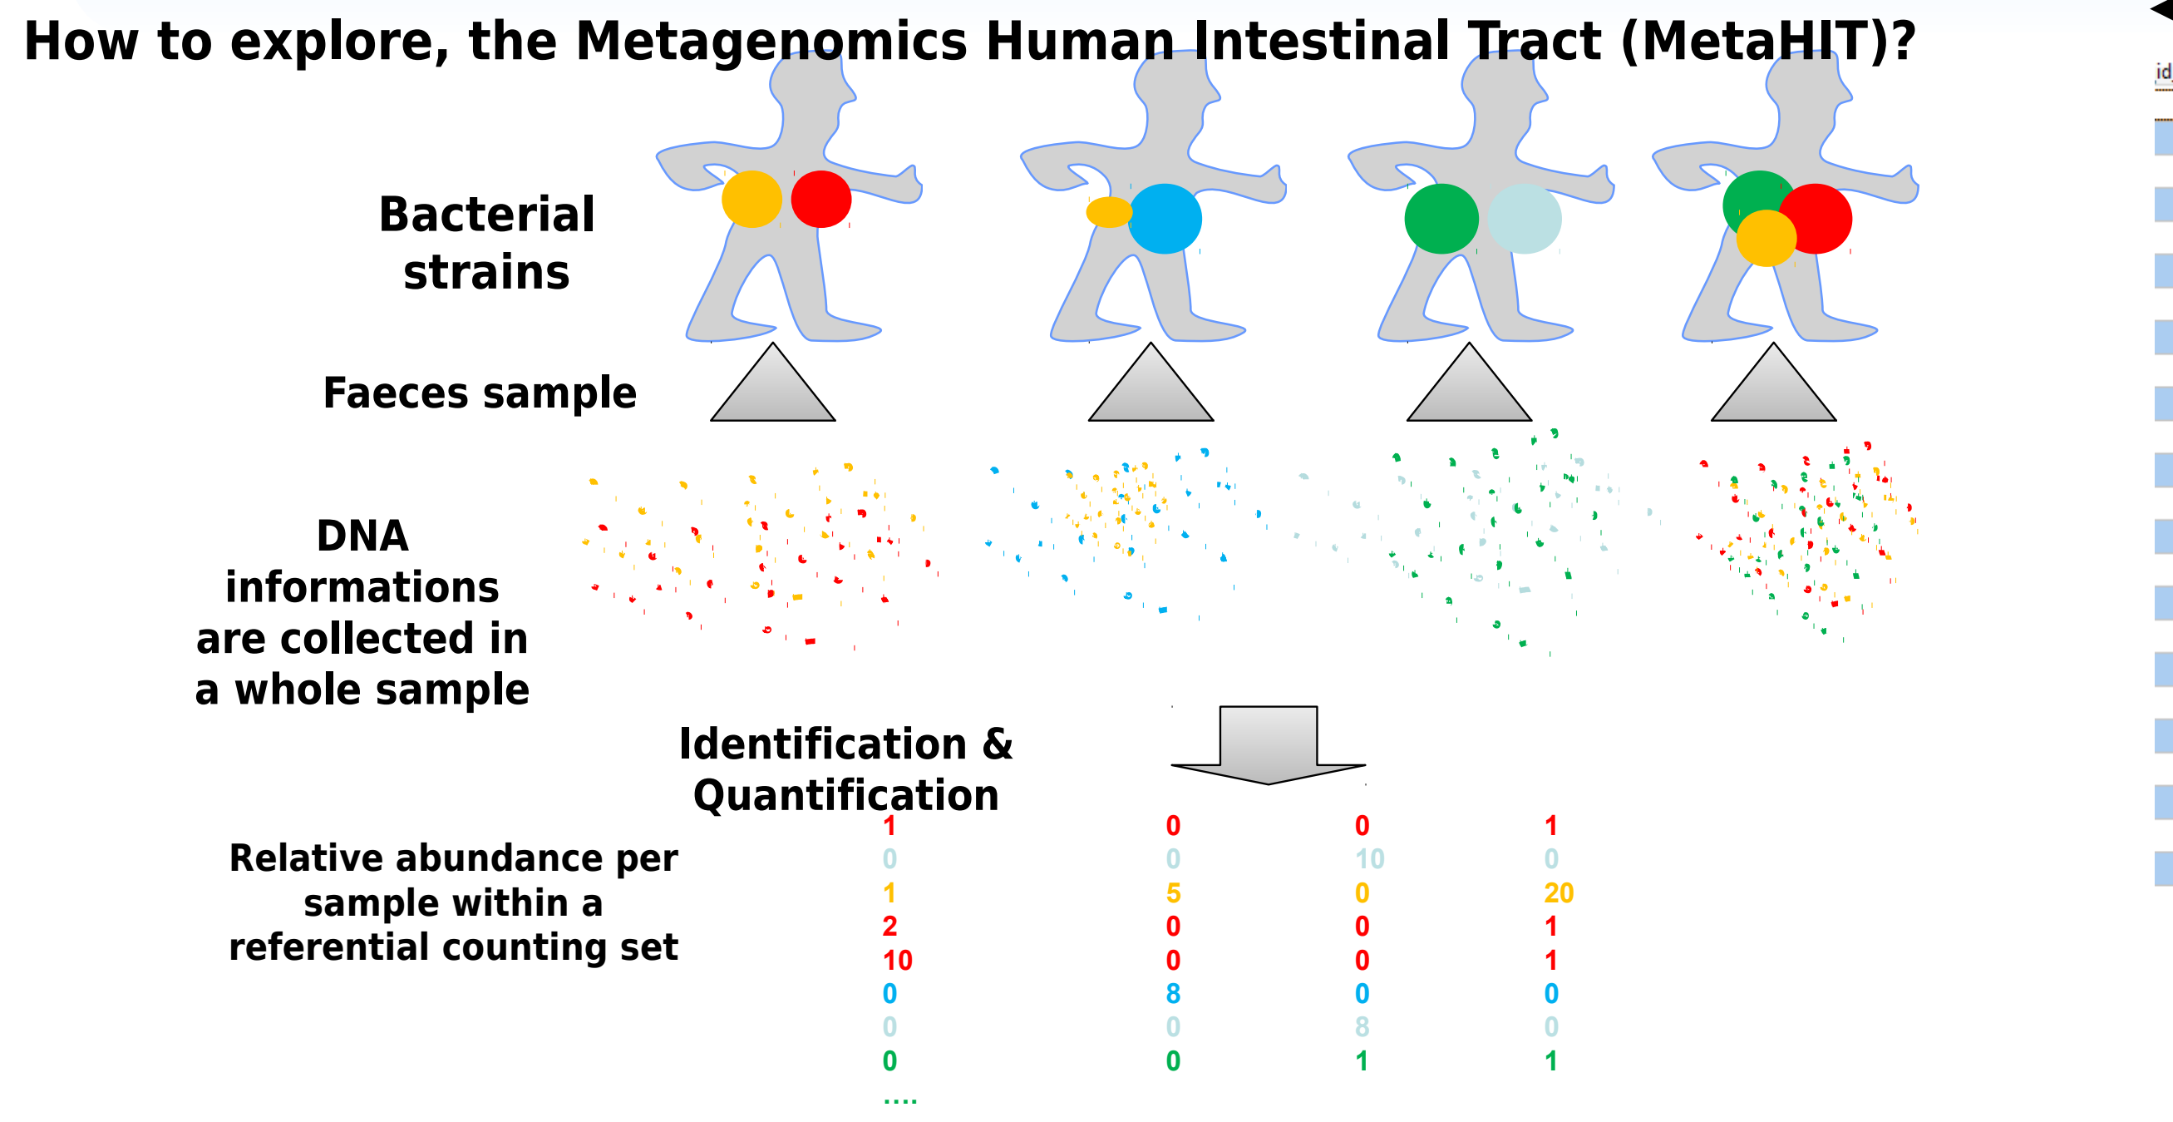
Bacterial strains
Faeces sample
DNA informations are collected in a whole sample
Identification & Quantification
Relative abundance per sample within a referential counting set
1 0 1 2 10 0 0 0 ....
0 0 5 0 0 8 0 0
0 10 0 0 0 0 8 1
1 0 20 1 1 0 0 1

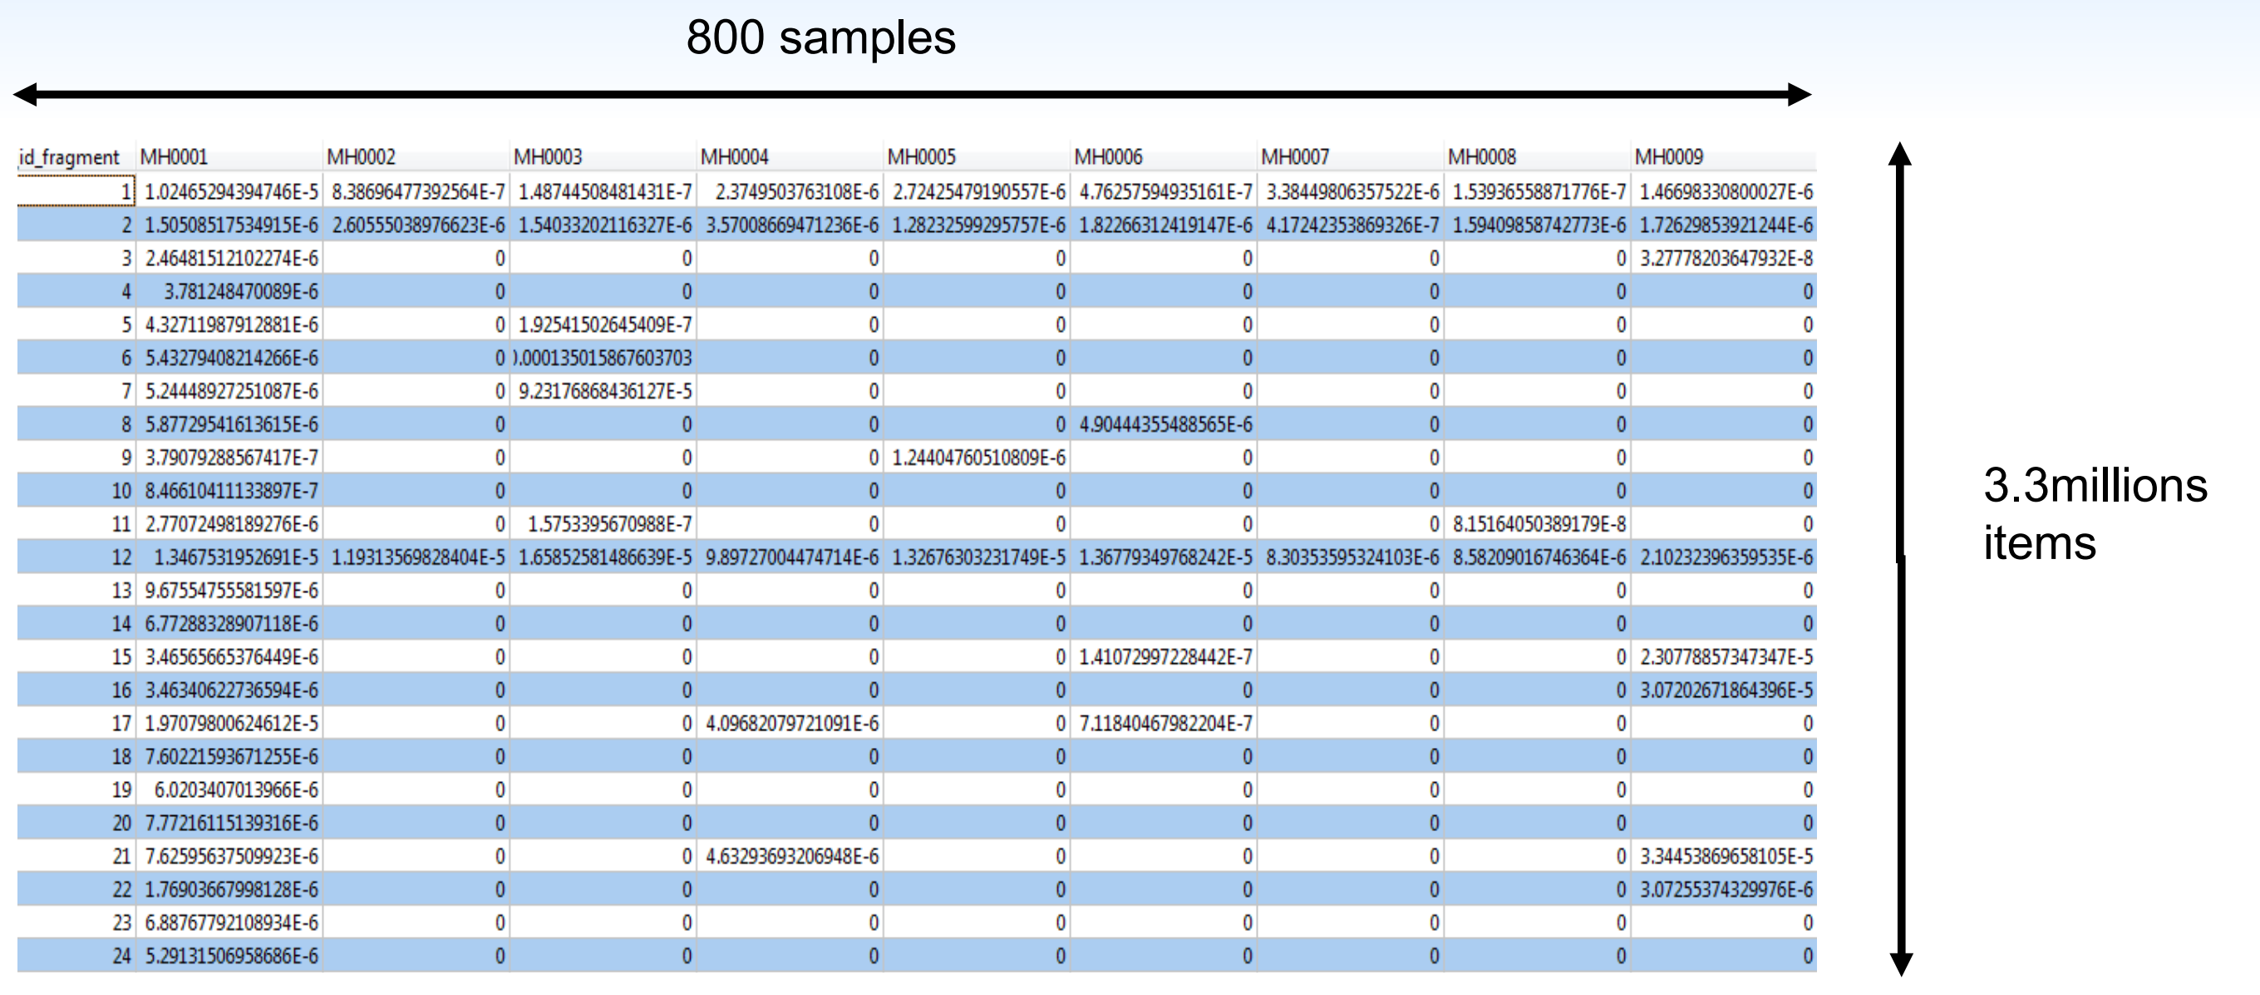
800 samples
3.3millions items

id_fragment	MH0001	MH0002	MH0003	MH0004	MH0005	MH0006	MH0007	MH0008	MH0009
1	1.02465294394746E-5	8.38696477392564E-7	1.48744508481431E-7	2.3749503763108E-6	2.72425479190557E-6	4.76257594935161E-7	3.38449806357522E-6	1.53936558871776E-7	1.46698330800027E-6
2	1.50508517534915E-6	2.60555038976623E-6	1.54033202116327E-6	3.57008669471236E-6	1.28232599295757E-6	1.82266312419147E-6	4.17242353869326E-7	1.59409858742773E-6	1.72629853921244E-6
3	2.46481512102274E-6	0	0	0	0	0	0	0	3.27778203647932E-8
4	3.781248470089E-6	0	0	0	0	0	0	0	0
5	4.32711987912881E-6	0	1.92541502645409E-7	0	0	0	0	0	0
6	5.43279408214266E-6	0	).000135015867603703	0	0	0	0	0	0
7	5.24448927251087E-6	0	9.23176868436127E-5	0	0	0	0	0	0
8	5.87729541613615E-6	0	0	0	0	4.90444355488565E-6	0	0	0
9	3.79079288567417E-7	0	0	0	1.24404760510809E-6	0	0	0	0
10	8.46610411133897E-7	0	0	0	0	0	0	0	0
11	2.77072498189276E-6	0	1.5753395670988E-7	0	0	0	0	8.15164050389179E-8	0
12	1.3467531952691E-5	1.19313569828404E-5	1.65852581486639E-5	9.89727004474714E-6	1.32676303231749E-5	1.36779349768242E-5	8.30353595324103E-6	8.58209016746364E-6	2.10232396359535E-6
13	9.67554755581597E-6	0	0	0	0	0	0	0	0
14	6.77288328907118E-6	0	0	0	0	0	0	0	0
15	3.46565665376449E-6	0	0	0	0	1.41072997228442E-7	0	0	2.30778857347347E-5
16	3.46340622736594E-6	0	0	0	0	0	0	0	3.07202671864396E-5
17	1.97079800624612E-5	0	0	4.09682079721091E-6	0	7.11840467982204E-7	0	0	0
18	7.60221593671255E-6	0	0	0	0	0	0	0	0
19	6.0203407013966E-6	0	0	0	0	0	0	0	0
20	7.77216115139316E-6	0	0	0	0	0	0	0	0
21	7.62595637509923E-6	0	0	4.63293693206948E-6	0	0	0	0	3.34453869658105E-5
22	1.76903667998128E-6	0	0	0	0	0	0	0	3.07255374329976E-6
23	6.88767792108934E-6	0	0	0	0	0	0	0	0
24	5.29131506958686E-6	0	0	0	0	0	0	0	0

Figure 1. The computation problem consists of a classical auto-correlation on the input genes x sample matrix. Using Spearman or Pearson correlation, main issues deal with the amount of floating point operations required. Computation complexity is O(N²). Hence, for a 3M genes input matrix, it requires to compute 3,5 1016 elementary operations.

Figure 2. A snapshot of Curie Supercomputer – hosted by GENCI in TGCC/CEA, Bruyères-Le-Châtel, France. The Curie Supercomputer has 144 nodes of dual Intel® Westmere® 2.66 GHz per node and dual nVidia® Tesla™ 2090 GPU (Fermi - 512 cuda cores, picture on the right side) per node – 192 Teraflops peak

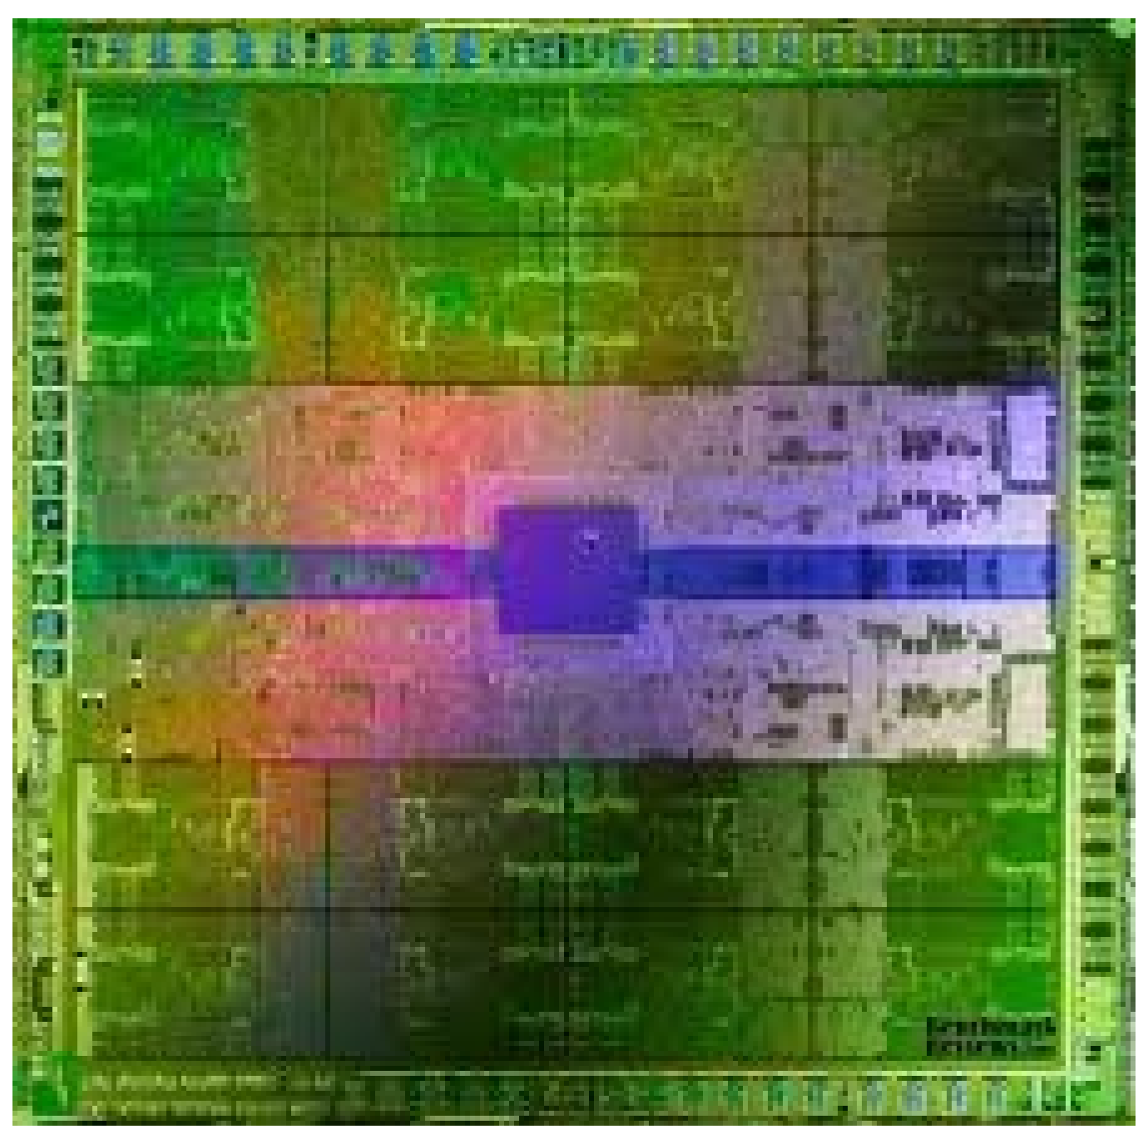

Figure 3. Our implementation targets large clusters consisting of GPU based nodes hence a hybrid programming model : MPI at highest level aimed at distributed computing between nodes and Cuda at the node level to benefit from the massively parallel architectures of the GPU.

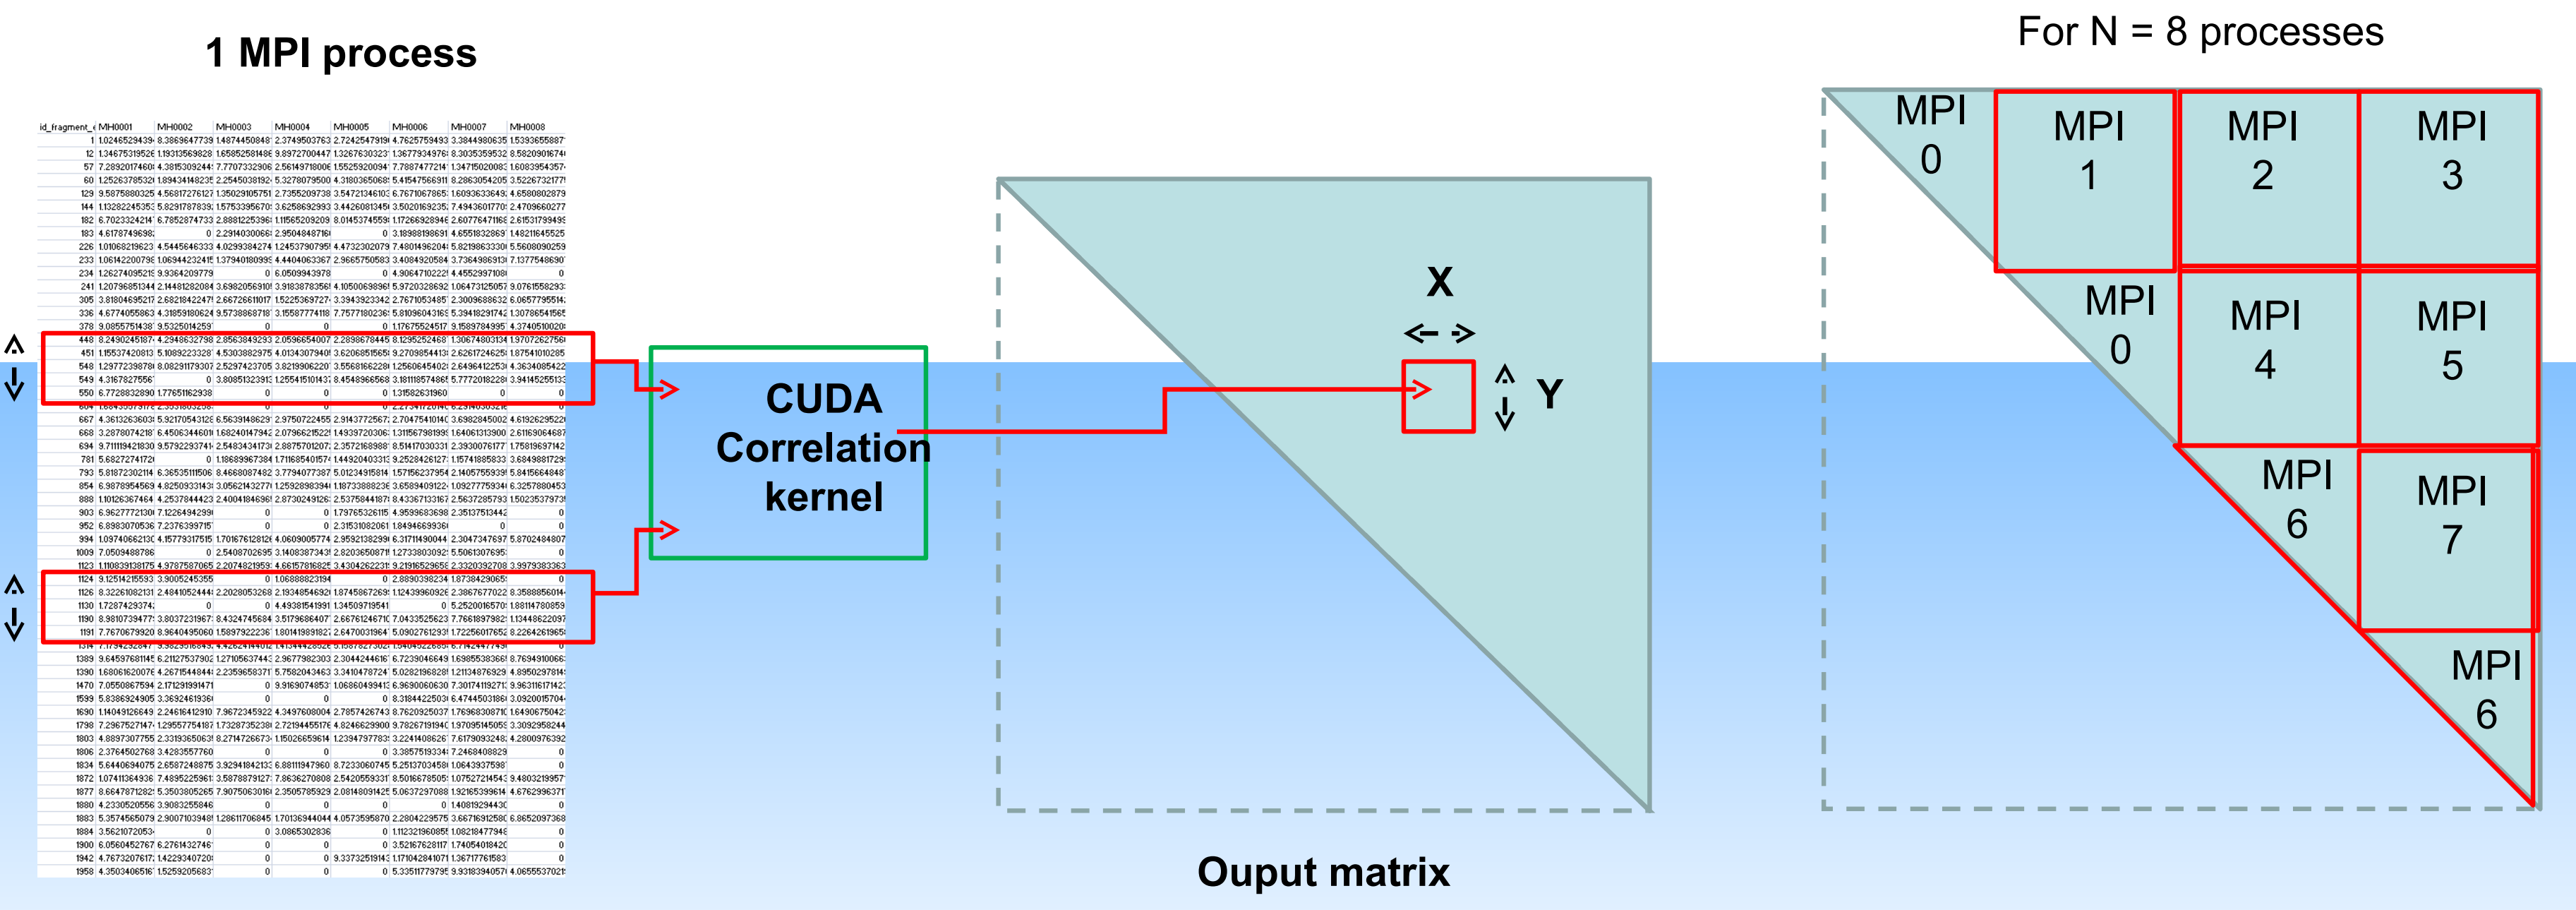

A divide and conquer strategy applied to the output matrix was used to get a fair load balance between the cluster nodes and limit the memory requirement per process. (illustrated here with eight processes)

The MetaProf code is released on request with a non disclosure license.

Quantitative Metagenomic and Big Computation

Figure 4. A universal pipeline

Metaprof is part of the ***Meteor*** processing pipeline designed by MetaQuant team to help in referencing data bank creation and in diagnostic evaluation. The MetaQuant Platform analyses 10 000 samples per year with a average size per project of 1000 samples. The MetaHIT project helps to size the pipeline and give a strong pressure to perform the computing analysis in a time below the week for a project.

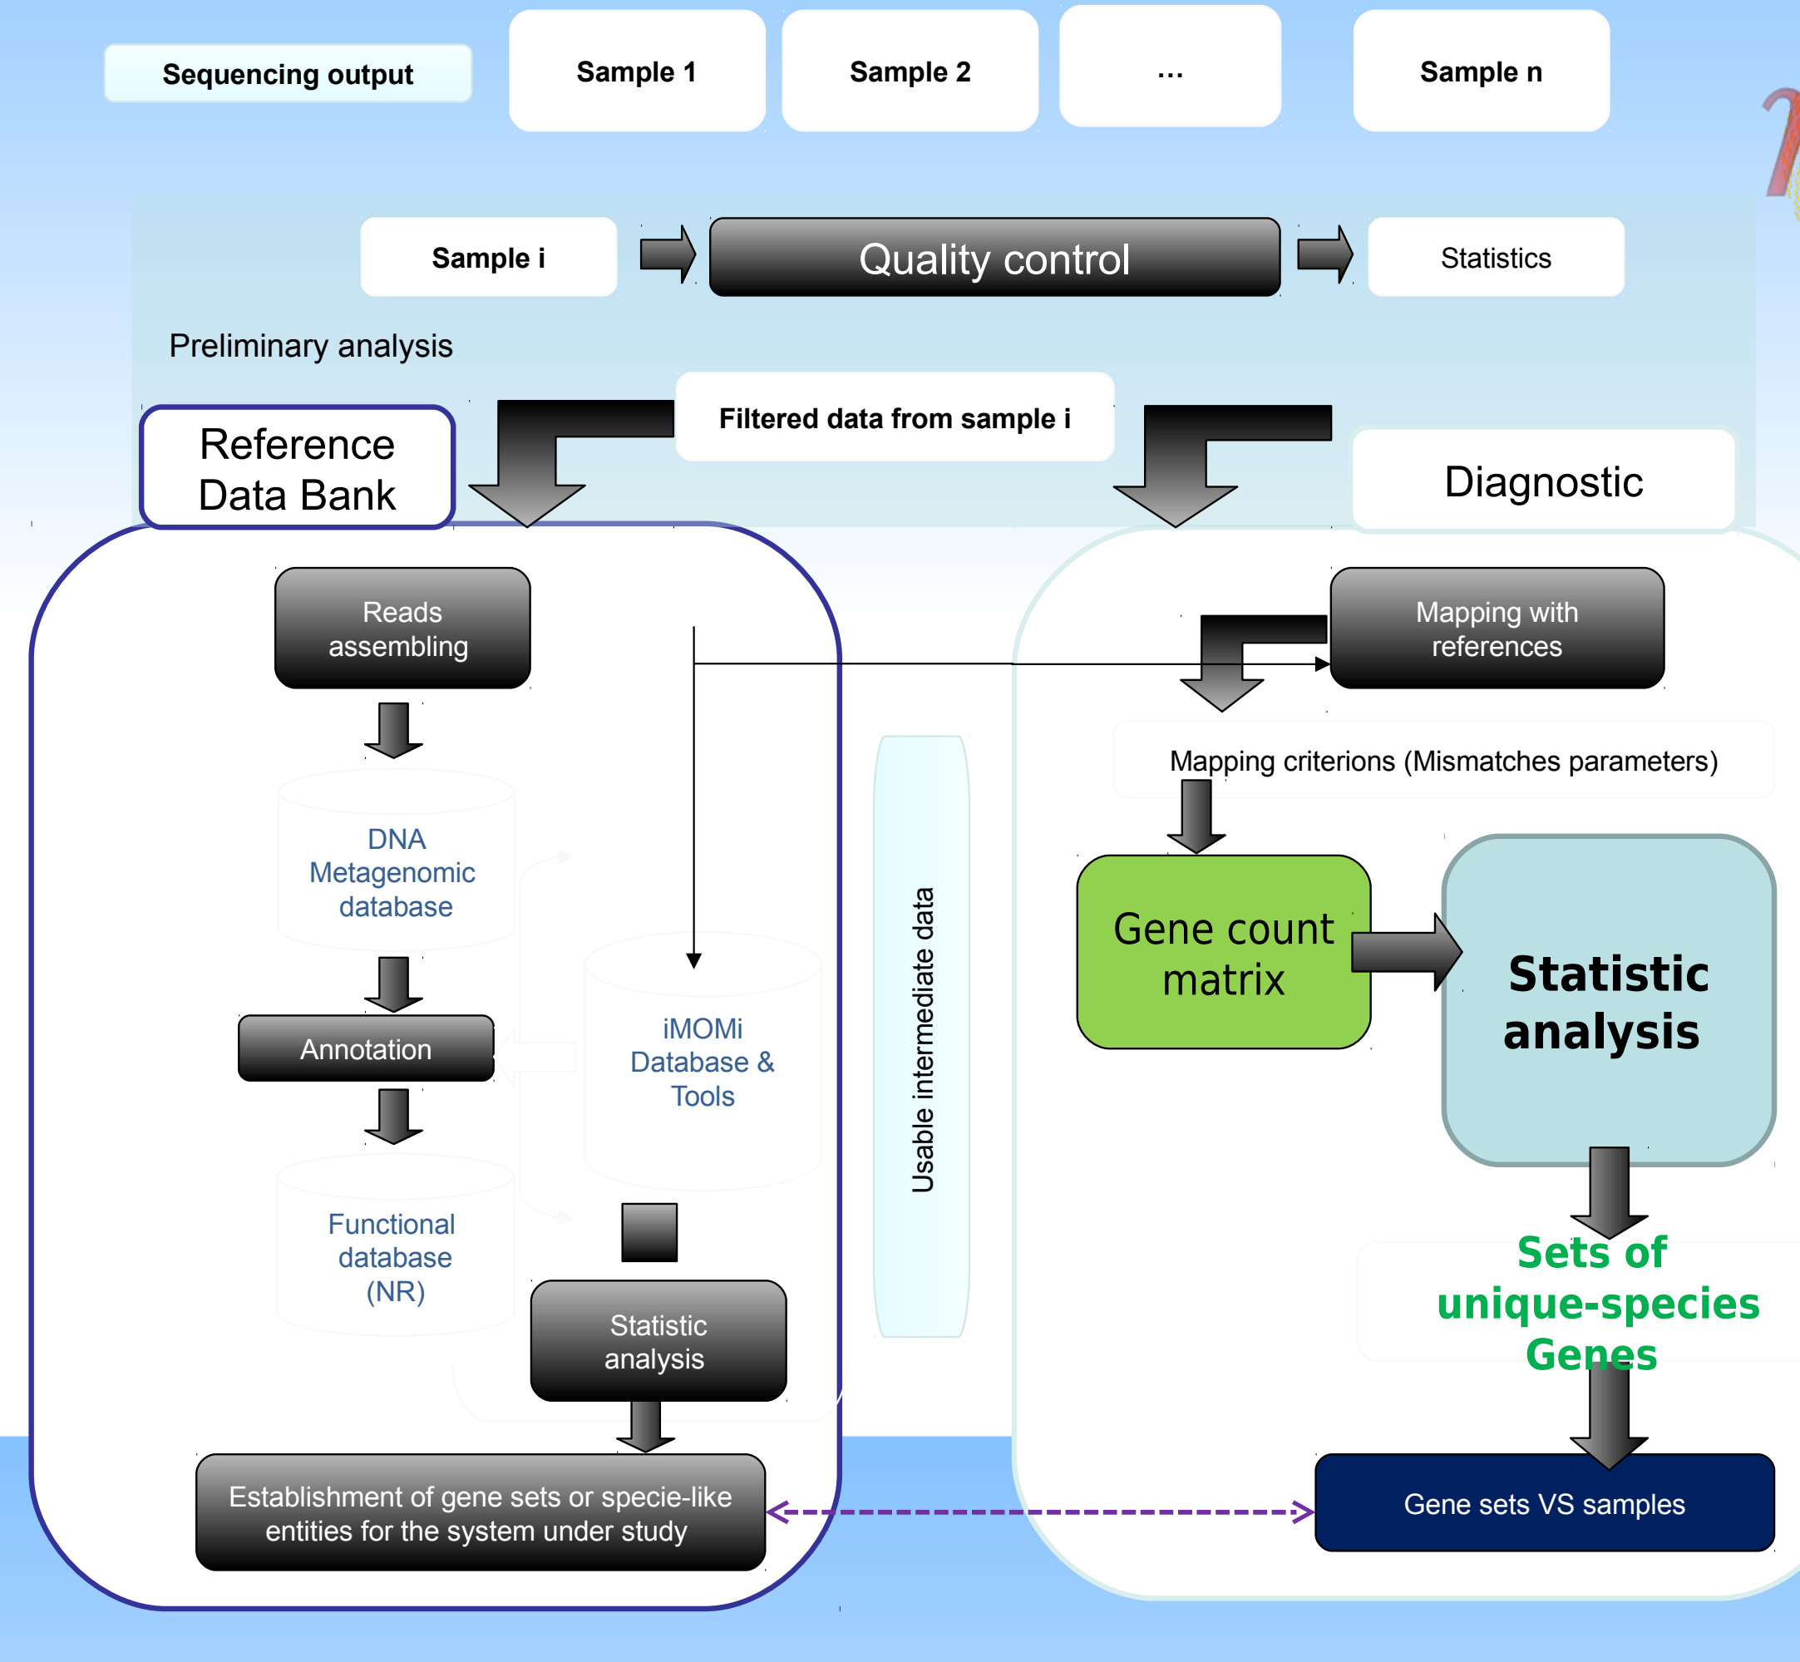

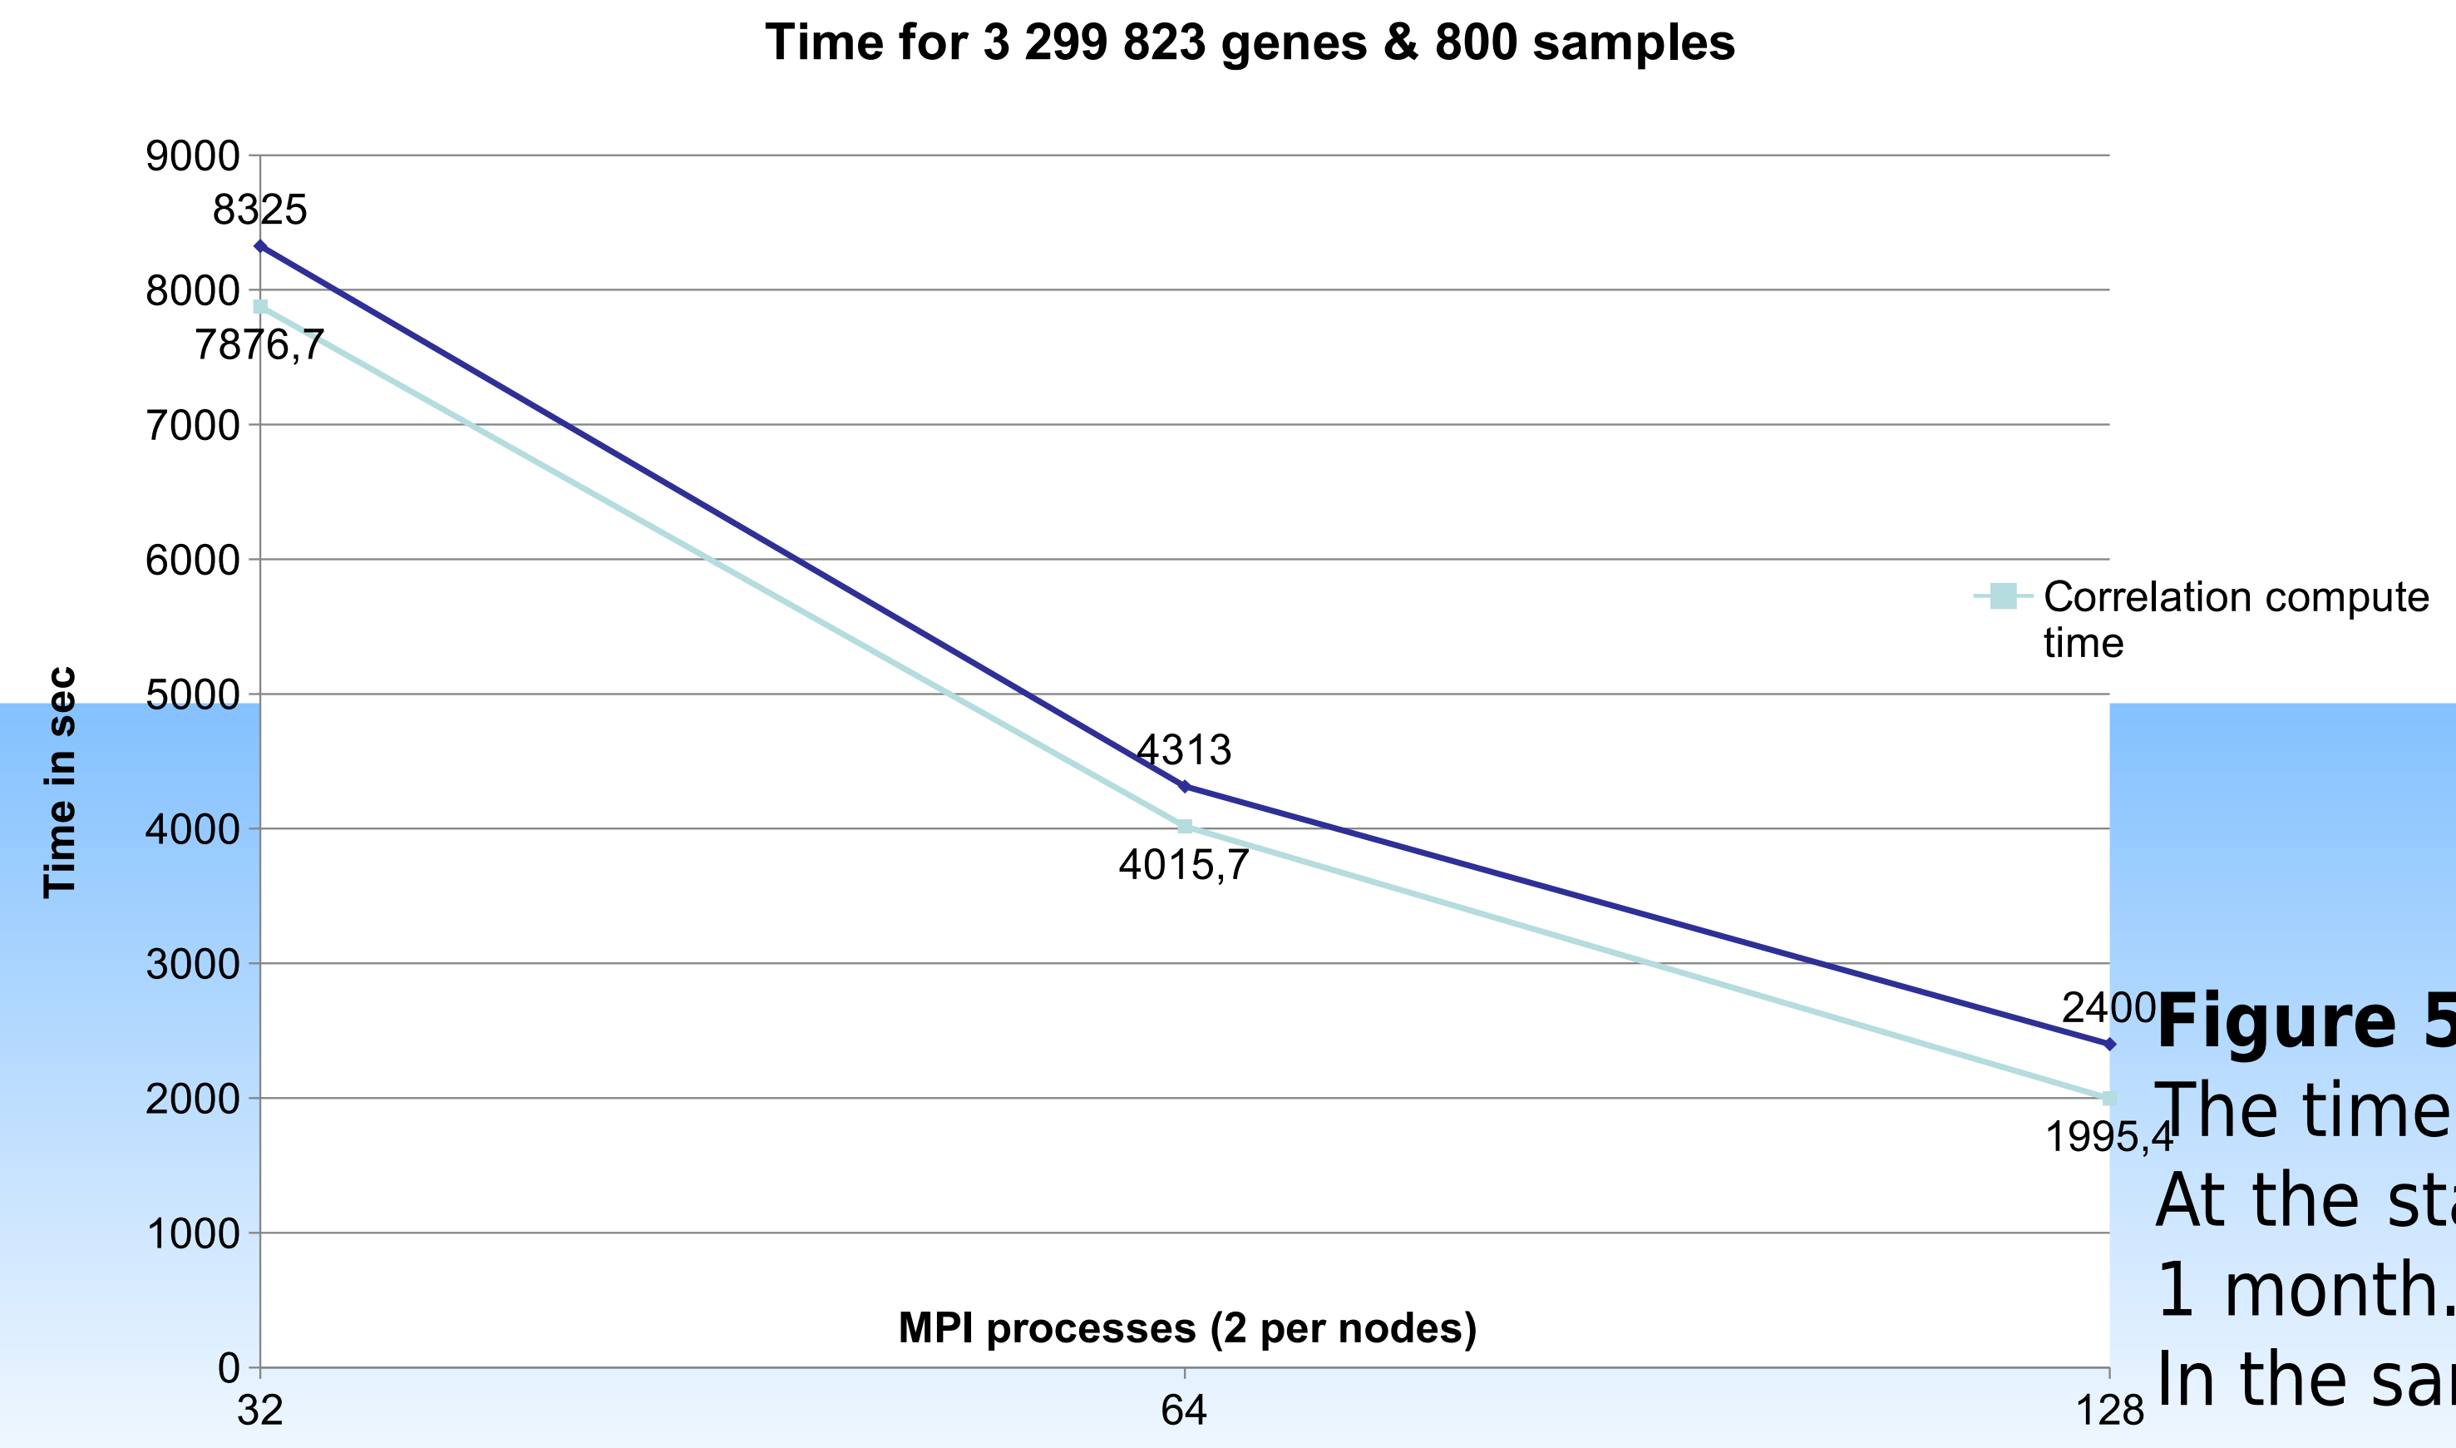

Figure 5. Time measure for doing the computation

The time measure benchmarked on a real use case – from the MetaHIT project.

At the starting point of the project, performing such a computation was expected to be done in 1 month. With the current optimization, the computation can be done quickly – less than 1 day.

In the same time, the throughput of the NGS device is increasing faster than the Moore law.

Expected results and potential impact

This optimization onto the Curie supercomputer proves that it is possible to use the GPU to perform huge computation and furthermore it brings GPU as a common component in the bioinformatics laboratory. It opens new opportunity : the possibility to perform more sophisticated computation on big data. Furthermore, as the benchmark proves that the time execution scales according to GPU number, it helps to design a small in house cluster. With a 16 GPU cluster, doing the MetaHIT use-case can be done in less than a day.

Links :

MetaQuant : www.metaquant.eu MetaHIT : www.metahit.eu Micalis : www.micalis.fr AS+ : www.asplus.fr OpenGPU : www.opengpu.net/EN