

Post-doc: Convergence of HPC and BigData infrastructures: case study on large graph drawing algorithms.

Description of the Cluster of Excellence SysNum (IdEx – université de Bordeaux)

The SysNum cluster aims to develop top level research on the next generation of digital interconnected systems, going from sensor design to the decision processes. The project is divided into 4 methodological work packages and three transversal packages focused on providing demonstrators and on engineering applications.

Each of the 4 methodological work packages contains several tasks at the interface of mathematics, computer science and engineering:

- Complex Systems and Interconnected Objects;
- Safety and Reliability;
- Modelling, Computing and Simulating;
- Massive and Heterogeneous Data.

The transverse work packages are aimed to develop various synergies between the methodological work packages in an engineering context of high interest at the regional, national or international level. These work packages are:

- Digital Ecological systems;
- Smart Campus;
- Robotics and drones.

They will provide experimental validation, proof of concept and demonstrators for the methodologies developed within the cluster.

More information on : <http://sysnum.labex.u-bordeaux.fr/en/> et <http://idex.u-bordeaux.fr/fr/>

Description of the project, activities and work context

Nowadays, networks (weighted graphs) are widely used for data analysis. The enthusiasm for this method of analysis is due to the fact that graph visualization facilitates the analysis of existing relationships between entities (persons, web pages, ...). These relationships can be automatically computed from the data, for instance, by computing a correlation matrix between entities and then thresholding it. These relationships can also evolve over time giving rise to dynamic-graphs or timestamped-graphs. By abstracting from the type of data and focusing only on relationships, graph visualization offers a remarkable tool for analysis of heterogeneous and dynamic Big Data.

In graph-based analysis methods, a first step is to lay out (or draw) the graph. This step algorithmically assigns a position to entities and relationships. The algorithmic complexity of the methods used [4] is in the general case of the order of $O(n^3)$. They cannot therefore be used on real data sets. For instance, in territorial surveillance tools to guarantee internal security the graphs have often hundred thousand of entities and millions of relationships. Recent work has shown that by combining "Barnes and Hut"[2] or "fast multipole"[3] (FMM) resolution methods and a multi-scale decomposition of graphs[5], it is possible to reduce this complexity to $O(n \log(n))$ while maintaining the properties of the original algorithm. Algorithms with $O(n \log(n))$ complexity can be efficiently launched on BigData infrastructures.

Therefore, it is theoretically possible to create methods that efficiently lay out graphs on these architectures [1].

At the moment, only heuristics [7] exist to lay out graphs on big data infrastructure and there is no method guaranteeing the same results as the original algorithms. Using of the resolution methods developed for HPC infrastructures on Big Data infrastructures remains an unresolved problem. The success the adaptation of these HPC resolution methods will allow to device graph layout algorithms which can be integrated into Big Data analysis ecosystems. This type of algorithm and study find their place in areas such as the observation and monitoring of social networks.

- The first objective of the post-doctoral fellowship will be to write an HPC version of the multi-scale mapping algorithm $O(n \log(n))$ based on an FMM. To do this, the candidate will implement the multi-scale algorithm above the FMM HPC library ScalFMMM developed in the team HiePACS. This algorithm will be evaluated on our HPC infrastructure PLAFRIM (www.plafrim.fr) and will lead to the writing of a report presenting this first distributed parallel implementation of this graph drawing algorithm. The data will then be transferred to the Big Data oriented infrastructure LSD to ensure visualization [6]. This first objective is mostly based on distributed algorithms and HPC.
- The second objective of the post-doctoral fellowship will be to port ScalFMM to make it work on the Big Data infrastructure LSD (<http://lsd.labri.fr>). This is mostly about making HPC - Big Data infrastructures converge, and it will allow us to compare traditional HPC software (StarPU, MPI) and Big Data ones (Spark, YARN) on a Big Data infrastructure. This comparison will result in the writing of a second report.
- The last objective will be to run the $O(n \log(n))$ distributed multi-scale parallel graph drawing algorithm, carried out for the first objective, directly on the Big Data oriented machine LSD. It will be done by taking advantage of the convergence effort achieved for the second objective. Finally, we will have a powerful graph drawing algorithm directly usable with the existing visualization on the LSD infrastructure. Having all the algorithms running on the Big Data infrastructure will prevent the costly transfer of a significant amount of data between an HPC machine and a Big Data machine. This work will lead to the writing of a third report.

References

[1] Alessio Arleo, Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. A distributed multilevel force-directed algorithm. In International Symposium on Graph Drawing and Network Visualization, pages 3–17. Springer, 2016.

[2] Josh Barnes and Piet Hut. A hierarchical $O(n \log n)$ force-calculation algorithm. Nature, 324 :446–449, 1986.

[3] Gereon Bartel, Carsten Gutwenger, Karsten Klein, and Petra Mutzel. An experimental evaluation of

multilevel layout methods. In International Symposium on Graph Drawing, pages 80–91. Springer, 2010.

[4] Tomihisa Kamada and Satoru Kawai. An algorithm for drawing general undirected graphs. Information processing letters , 31(1) :7–15, 1989.

[5] Yehuda Koren, Liran Carmel, and David Harel. Ace: A fast multiscale eigenvectors computation for drawing huge graphs. In Information Visualization, 2002. INFOVIS 2002. IEEE Symposium on , pages 137–144. IEEE, 2002

[6] Alexandre Perrot, David Auber. Cornac: Tackling Huge Graph Visualization with Big Data Infrastructure. IEEE transactions on big data, IEEE, 2018, 14, pp.1 - 1.

[7] Antoine Hinge, Gaëlle Richer, David Auber. MuGDAD: Multilevel graph drawing algorithm in a distributed architecture. *Conference on Computer Graphics, Visualization and Computer Vision, IADIS*, Jul 2017, Lisbon, Portugal. pp.189.

Candidate Profile

Skills (knowledge, know-how, know-how) :

MPI, parallel computing, distributed computing, graph drawing, graph theory, Hadoop ecosystem, Spark. Information visualization.

Degree required and / or level of qualification

PhD computer science

Experience required

Publications in at least one of the fields of research related to the project.

Job Information

Location : Bordeaux, LaBRI, university of Bordeaux and INRIA Bordeaux

Type of contract : Temporary

Job status : full-time

Duration: 2 years

Starting Date : as soon as possible

Remuneration: according to profile and experience

Contact(s)

David Auber: auber@labri.fr, Emmanuel Agullo : emmanuel.agullo@inria.fr

CV and cover letter should be sent to the contact(s).