The VELaSSCo framework: a software platform for end-user analytics and visualization of large simulation datasets

G. Filippone[1*], A. Janda[1*], P. Totoo[1], J.P. Morrissey[1], K.J. Hanley[1], S. Papanicolopulos[1], J.Y. Ooi[1], A. Coll[2], M.A. Pasenau de Riera[2], J. Mora[2], J. Pérez[2], B. Lange[3*], B. Raffin[3], I. Cores[3], T. Nguyen[3*], A. Dietrich[4], F. Michel[4], T. Gierlinger[4], H. Dahl[5], O. Barrowclough[5], I. Martinez Rodriguez[6], T. Pariente Lobo[6], M.A. Tinte Garcia[6], C. Pumar Garcia[6], J. Haenisch[7], K. Bengtsson[7], O. Liestøl[7]

[1] Institute for Infrastructure and Environment, School of Engineering, The University of Edinburgh, Edinburgh, UK

[2] Edificio C1 Campus Nord, UPC, calle Gran Capitán s/n, 08034 Barcelona, Spain

[3] Inria Grenoble Rhône-Alpes, 655 Avenue de l'Europe, 38330 Montbonnot-Saint-Martin, France

[4] Fraunhofer IGD, Fraunhoferstraße 5, 64283 Darmstadt, Germany

[5] SINTEF, Forskningsveien 1, 0314 Oslo, Norway

[6] Atos Spain SA, C/ Albarracin 25, 28037 Madrid, Spain

[7] Jotne EPM Technology AS, Grenseveien 107, 0663 Oslo, Norway

[*] indicates former affiliation

THE VELASSCO PROJECT

- VELaSSCo (Visual Analysis for Extremely Large-Scale Scientific Computing) is an EC FP7 project involving a consortium of seven European partners (Fig. 1).
- VELaSSCo aims to provide new visual analysis methods for large-scale simulations serving the petabyte era.
- The main output of the project is the VELaSSCo platform which has been designed and developed to perform distributed post-processing and visualisation of the results of very large engineering simulations.
- The VELaSSCo platform exists in two versions: a fully open-source version and a proprietary version which use Apache HBase and EXPRESS Data Manager™ (EDM) as database systems, respectively.

Fig. 1 →

The 7 European partners comprising the VELaSSCo consortium

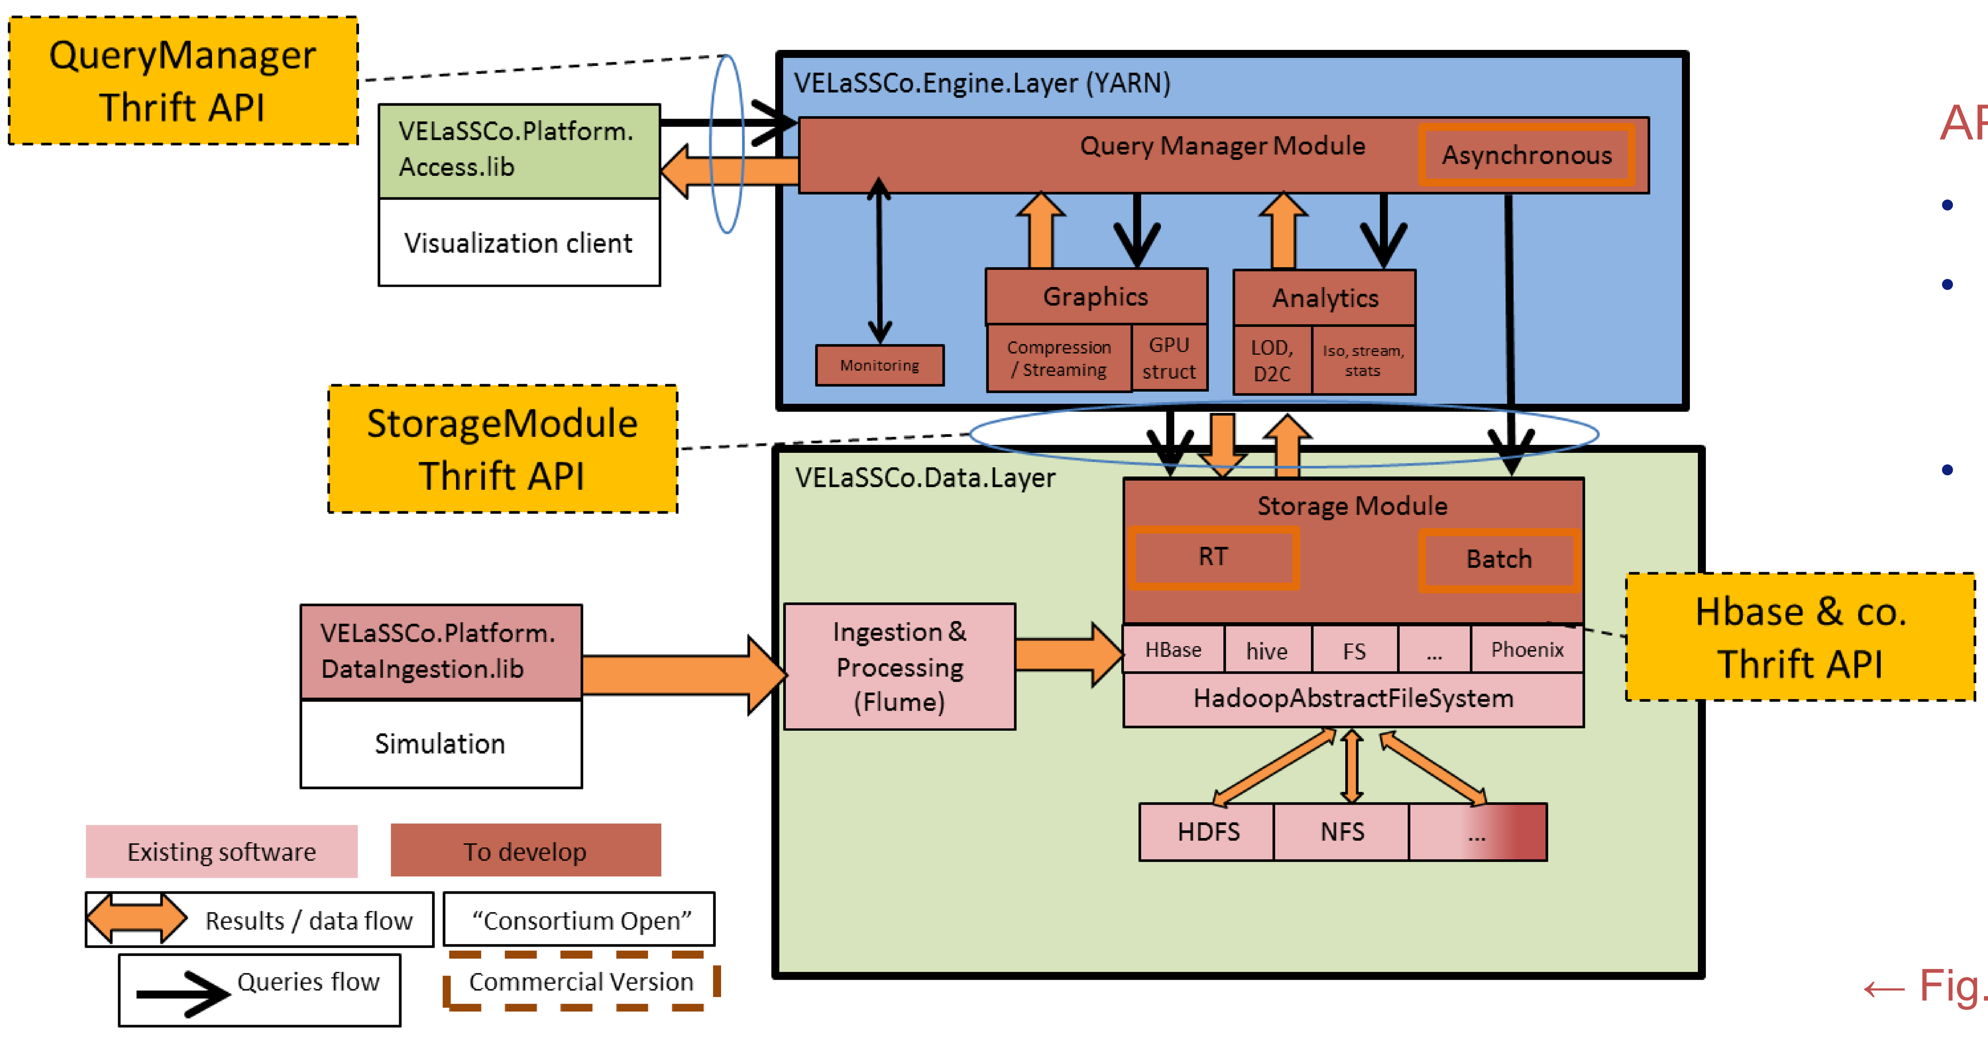

ARCHITECTURE OF THE VELASSCO PLATFORM

- The open-source version of the platform is shown in Fig. 2.
- The architecture of the VELaSSCo platform is based on the open-source Hadoop software stack, a Java-based framework for distributed storage and processing of Big Data.
- It is composed of two main layers:
  1. Data Layer: responsible for storing, accessing and translating the simulation data. It is composed of both standard tools such as Hadoop with HDFS, Apache Flume and HBase, and a bespoke Storage module which is based on a HBase Thrift server.
  2. Query Engine Layer: in charge of receiving the user queries from the visualisation client, extracting and/or analysing the simulation data and returning the results in a GPU-friendly format for fast visualisation.

← Fig. 2

The VELaSSCo architecture used for the open-source platform

PRELIMINARY EVALUATION OF THE PLATFORM

- A first evaluation of the platform was held in February 2016.
- Fig. 3 shows one of the evaluation examples: a fluidised bed simulated using the discrete element method (DEM) coupled with computational fluid dynamics (CFD).
- The model contained around 40,000 time-steps, 12,000 particles per time-step and more than 3,000 particle-particle and particle-wall contacts per time-step.
- Participants in the evaluation were able to compute the *discrete-to-continuum* transformation which applies temporal and spatial coarse graining methods (Goldhirsch, 2010; Labra et al., 2013) to DEM simulation data in order to compute bulk quantities that are projected onto an underlying continuum mesh.
- This first Hadoop implementation of the discrete-to-continuum transformation has shown excellent results in terms of its scalability and normalised speedup.

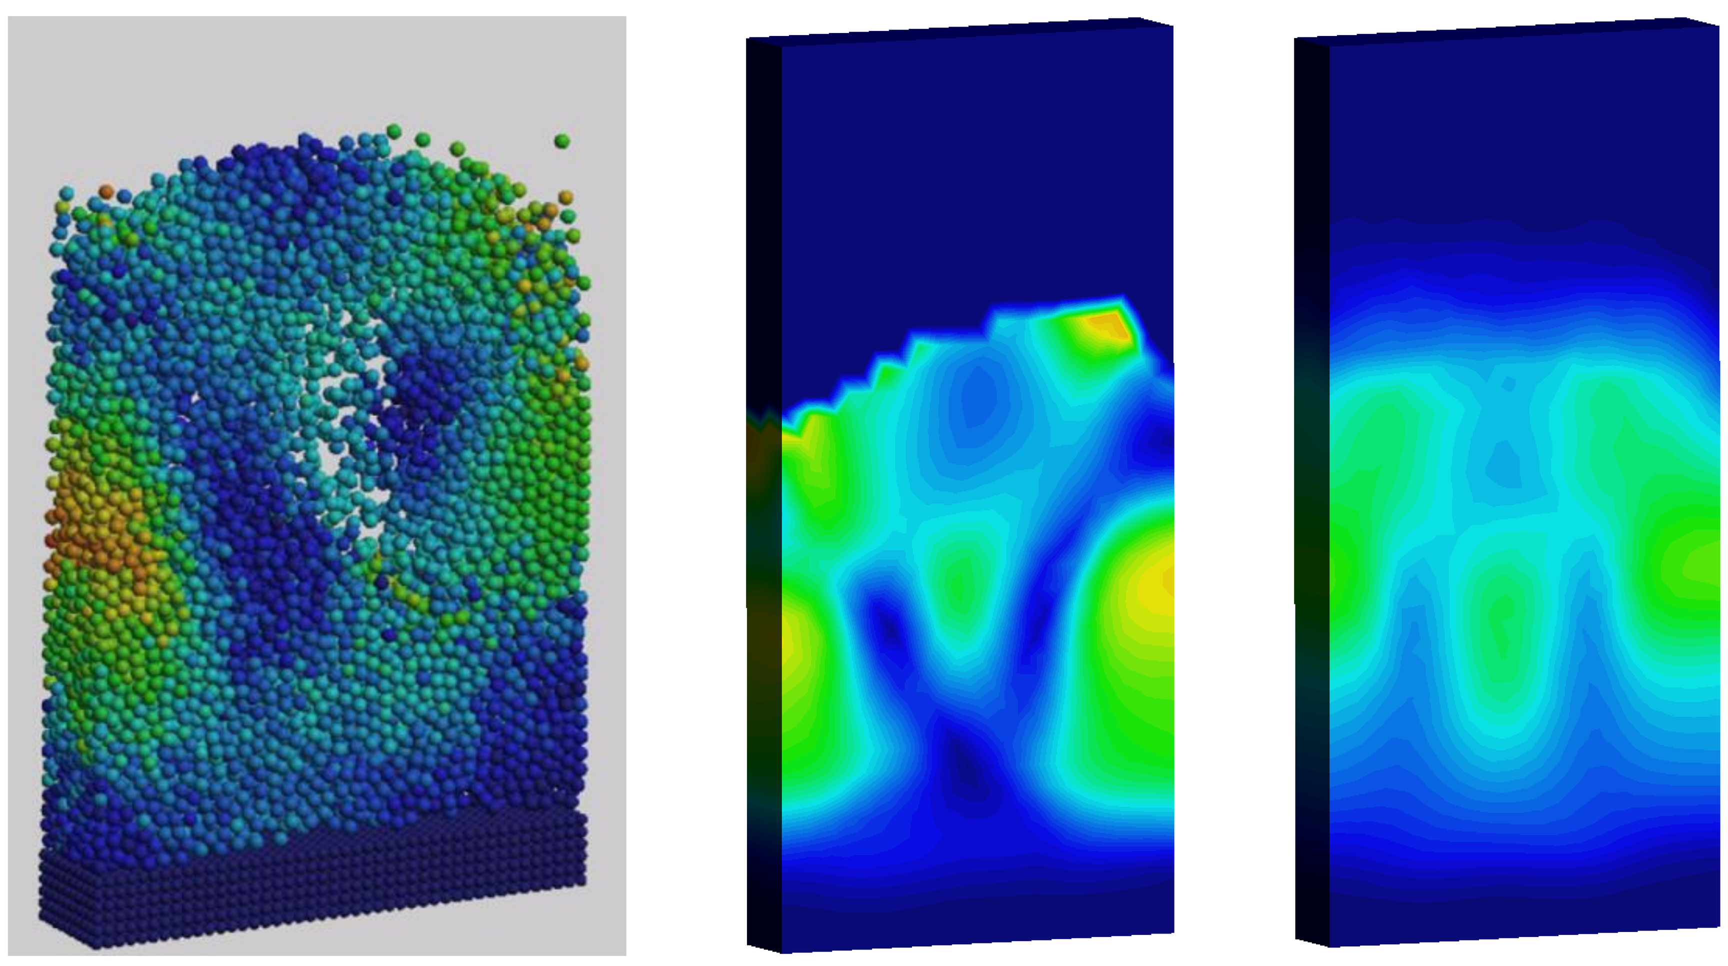

↑ Fig. 3

Example of the discrete-to-continuum transformation applied to fluidised bed data. The colours indicate modulus of velocity. The leftmost image shows discrete particle data at one time-step. The central image shows the result of spatial averaging only. The rightmost image shows spatial and temporal averaging over multiple time-steps.

REFERENCES

- Goldhirsch, I. (2010). Stress, stress asymmetry and couple stress: from discrete particles to continuum fields. Granular Matter, 12: 239–256
- Labra, C., Ooi, J.Y. & Sun, J. (2013). Spatial and temporal coarse-graining for DEM analysis. AIP Conference Proc., 1542: 1258–1261, Powders and Grains 2013, Sydney, Australia

The development of the VELaSSCo platform is funded by the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 619439.