Scaling DBSCAN towards exascale computing for clustering of big data sets

Ernir Erlingsson[1], Gabriele Cavallaro[2], Morris Riedel[1,2], Helmut Neukirchen[1]

[1]University of Iceland
{ere29, helmut}@hi.is

[2]Jülich Supercomputing Centre, Germany
{g.cavallaro, m.riedel}@fz-juelich.

This project has received funding from the European Commission's Horizon 2020 Programme for research, technological development and demonstration under Grant Agreement n°754304

UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING, MECHANICAL ENGINEERING AND COMPUTER SCIENCE

JÜLICH FORSCHUNGSZENTRUM

DEEP Extreme Scale Technologies

Talk Overview

- **DBSCAN** Clustering
- Highly Parallel DBSCAN (**HPDBSCAN**)
- Large Point-cloud datasets
- **DEEP-EST** Modular Supercomputing Architecture (**MSA**)
- HPDBSCAN MSA mapping

DBSCAN Clustering Algorithm

- **D**ensity-**b**ased **s**patial **c**lustering of **a**pplications with **n**oise (**DBSCAN**).
  - Detects **arbitrarily shaped** clusters,
  - **Detects** and can filter **noise,**
  - **No need to know number of clusters**.

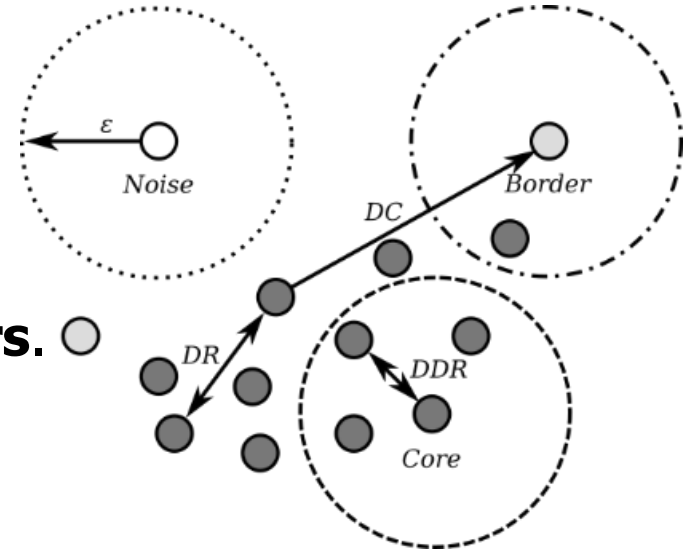

Götz et al. 2015

- Two parameters:
  - Spatial search radius *ε,*
  - Point density *minPts.*
    - At least *minPts* elements needed within *ε* radius to form a cluster.
    - Otherwise considered noise.

Ester, Kriegel, Sander, Xu, "Density-based spatial clustering of applications with noise" Proc. Second Int. Conf. on Knowledge Discovery and Data Mining. AAAI Press, 1996.

HPDBSCAN Specification

Highly **P**arallel **DBSCAN:**

- Data-space hypergrid overlay with load balancing.
- Uses MPI & OpenMP
- HDF5 I/O
- Strong scaling

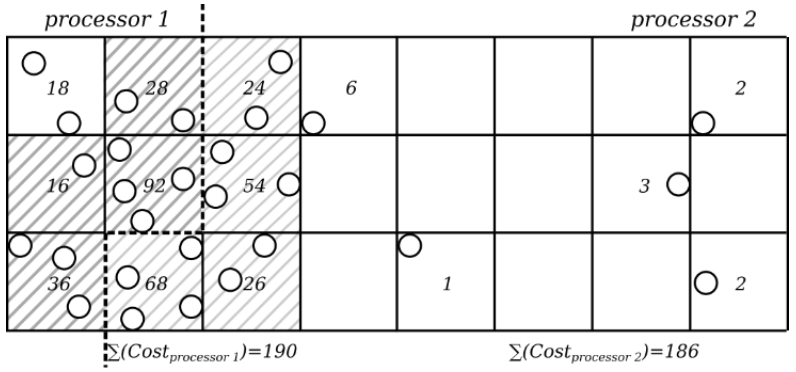

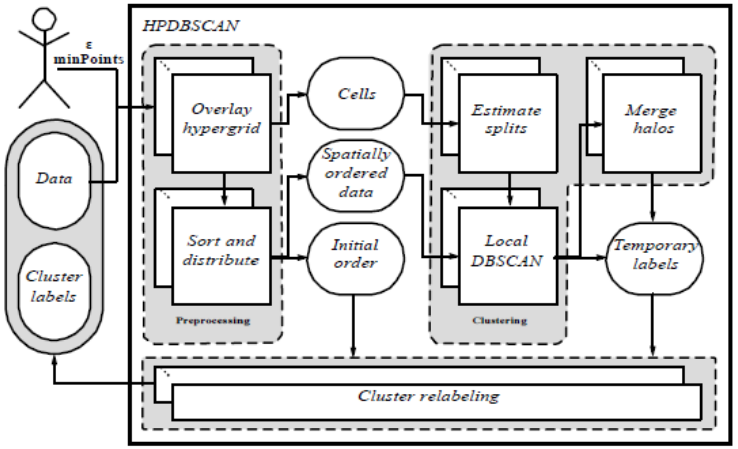

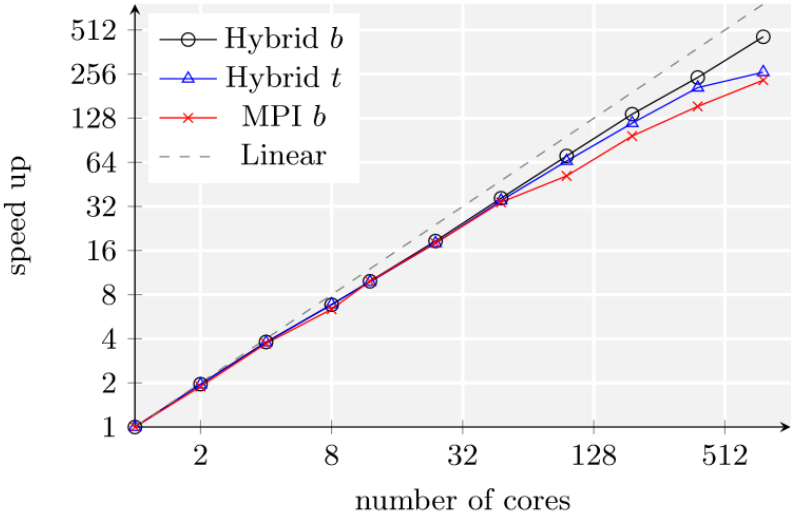

Götz, Bodenstein, Riedel, "HPDBSCAN: highly parallel DBSCAN," Proc. Workshop on Machine Learning in High-Performance Computing Environments, in conjunction with Super Computing 2015, ACM.

Point-cloud data acquisition

3D laser scans of buildings, landsmarks using ground or airborne sensors.

Examples include:

- Inner city of Bremen dataset
- The national LiDAR datasets

http://www.igi-systems.com

http://www.igi-systems.com

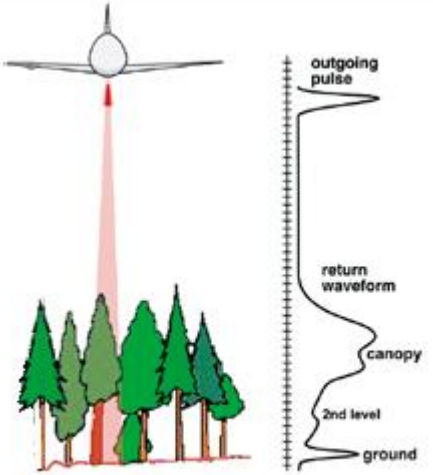

http://www.igi-systems.com

Bremen Dataset

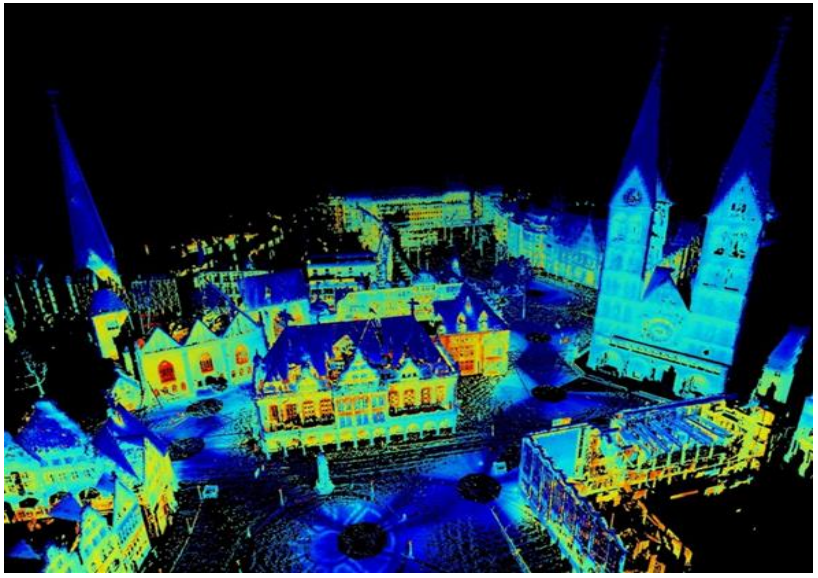

http://kos.informatik.uni-osnabrueck.de/3Dscans/

- 3D laser scans made in the inner city of the German city Bremen.
- Full Bremen dataset several GBs.

Not enough for exascale!

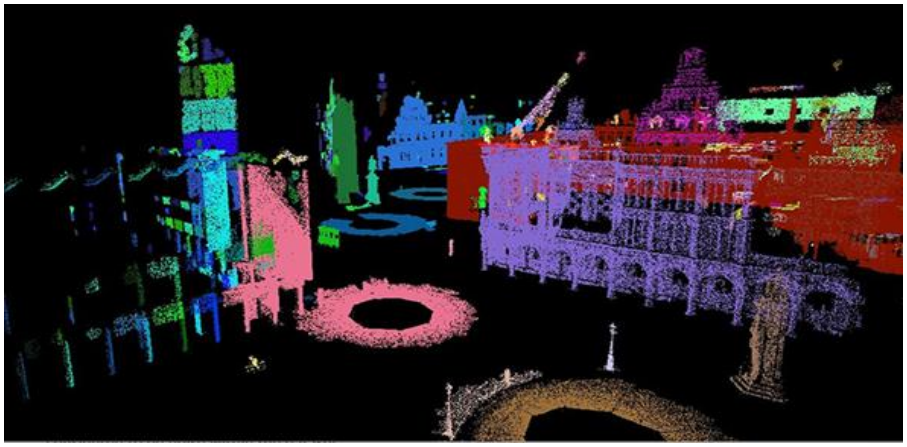

HPDBSCAN clustering

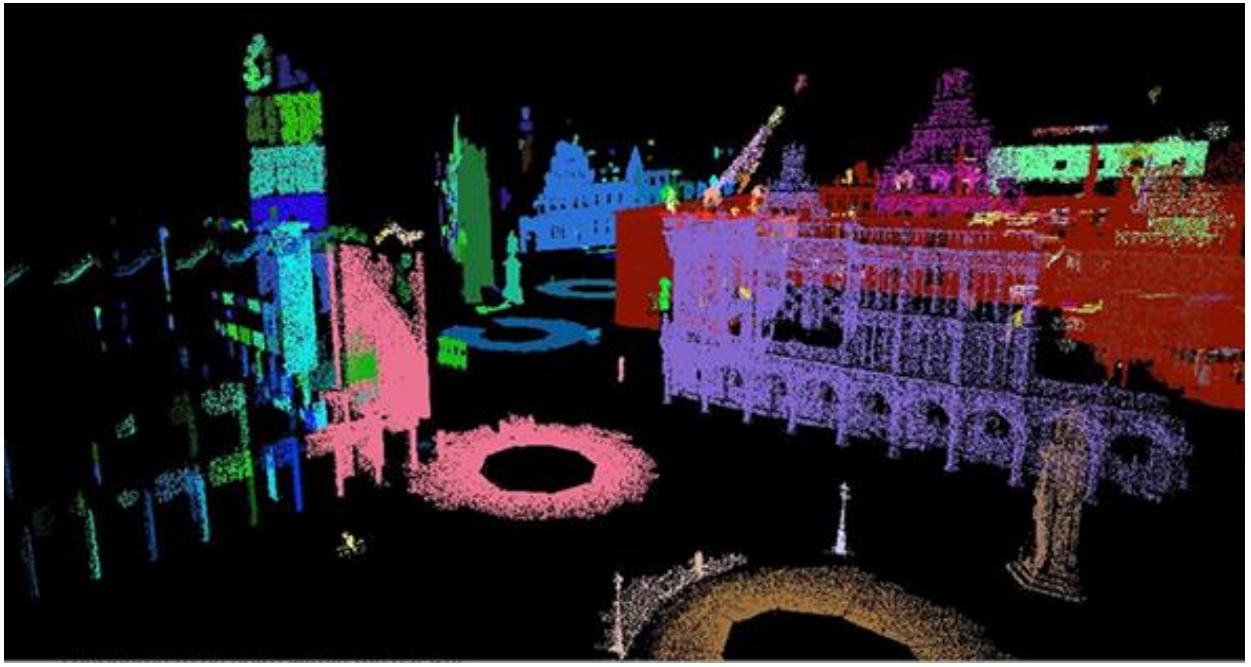

HPDBSCAN applied to the Bremen dataset where the colors represent different cluster labels.

Dutch National LiDAR dataset

- AHN-1 (2003) 4-5 m point spacing
- AHN-2 (2014) 6-10 points per m^2
- AHN-3 (2019) 0,5 m point spacing

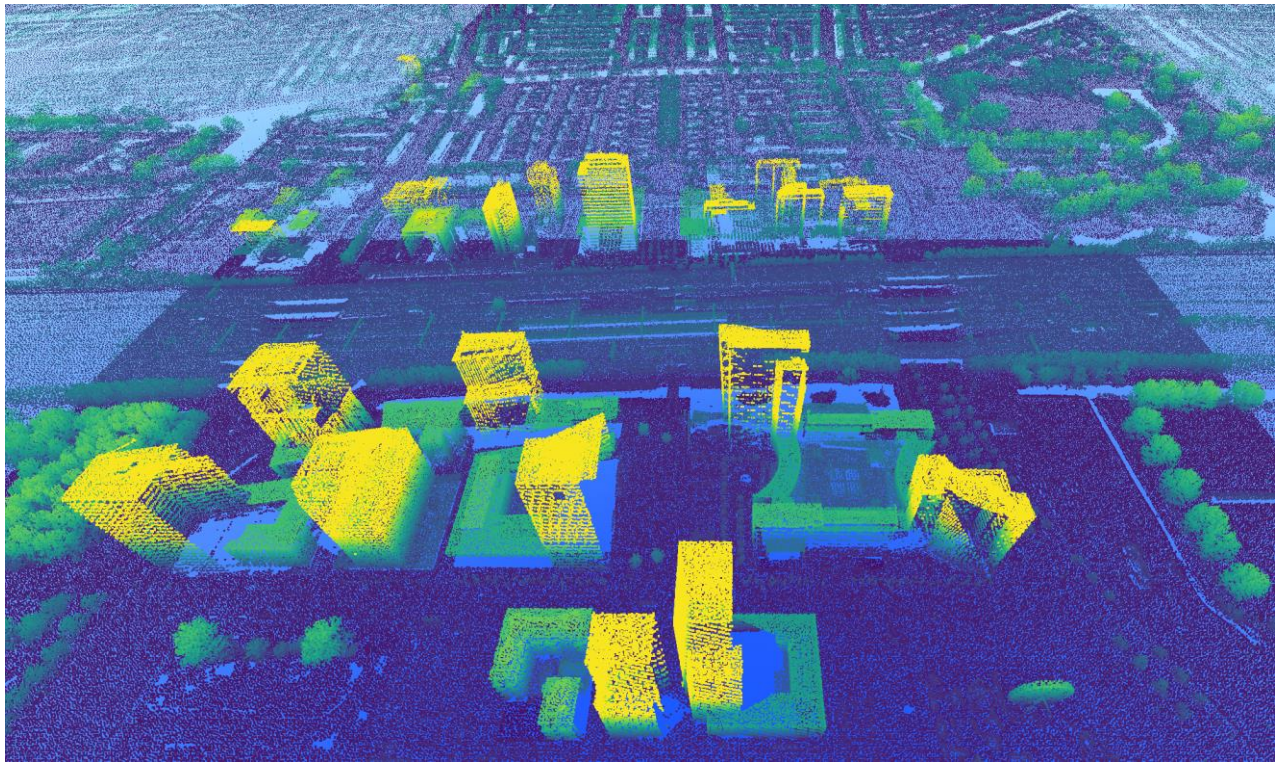

https://www.pdok.nl/en/pdok-products

DEEP-EST MSA Architecture

Heterogenous supercomputer with numerous modules which can be tailored to enhance an application.

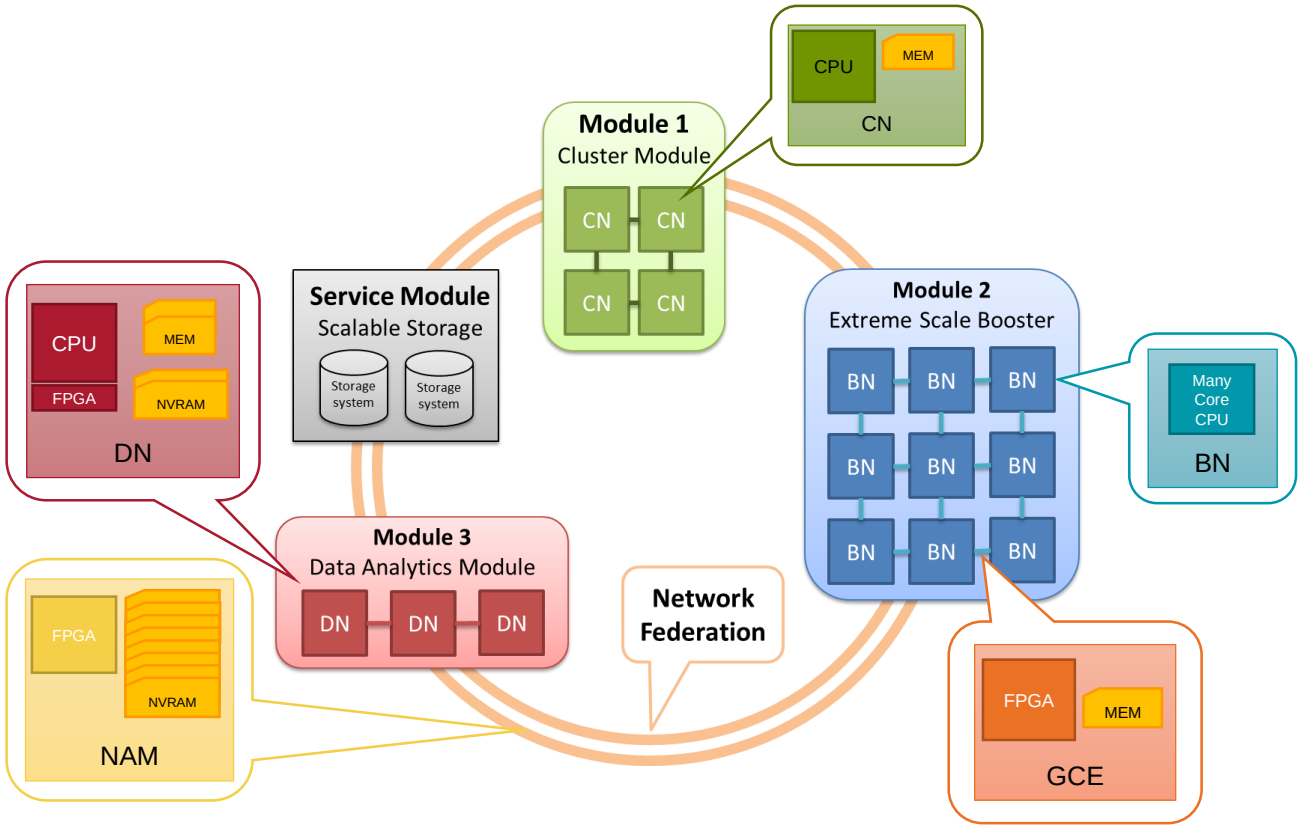

Network Access Memory (NAM)

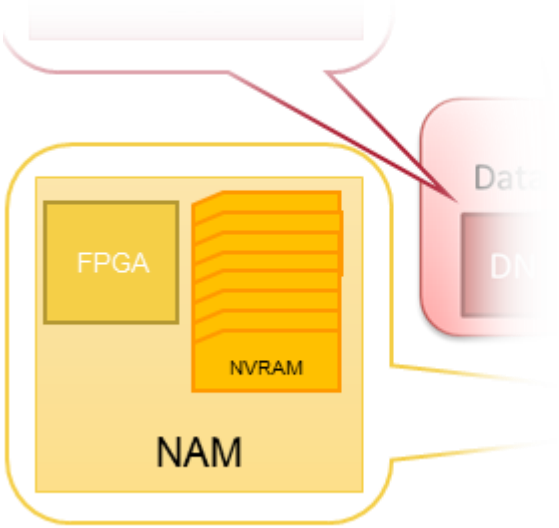

- Part of the fabric via a fast EXTOLL interconnect
- Useful as a fast general purpose memory with processing capabilities

- Dedicated device with **N**on **V**olatile **RAM** and a **F**ield **P**rogrammable **G**ate **A**rray

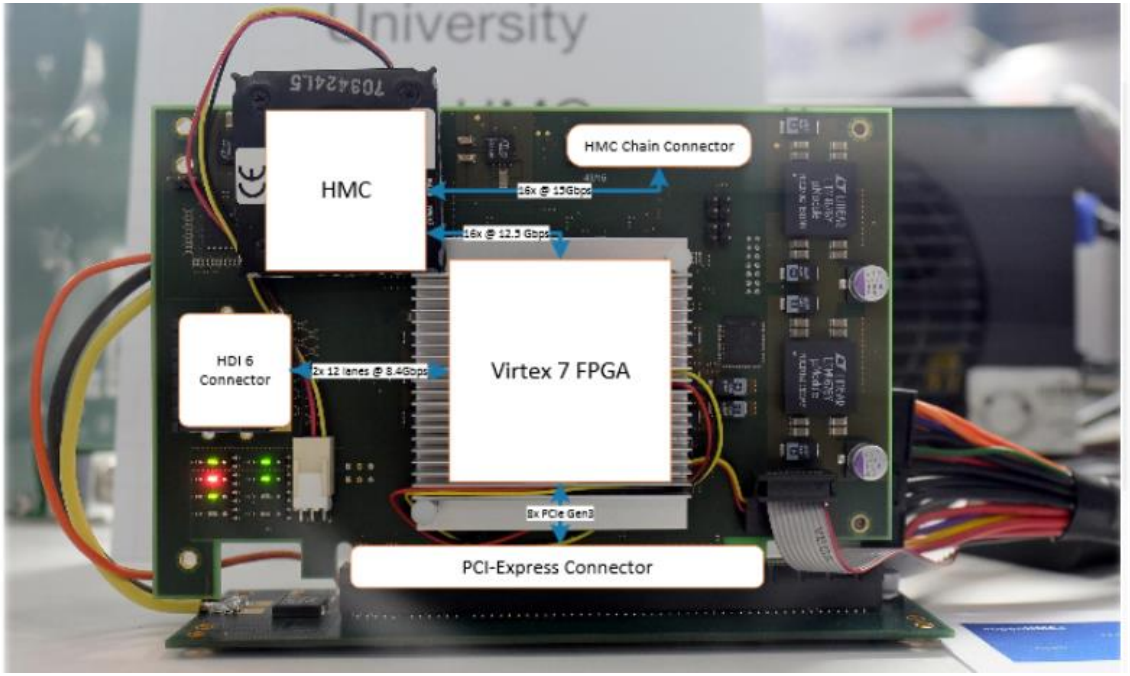

Extreme Scale Booster (ESB)

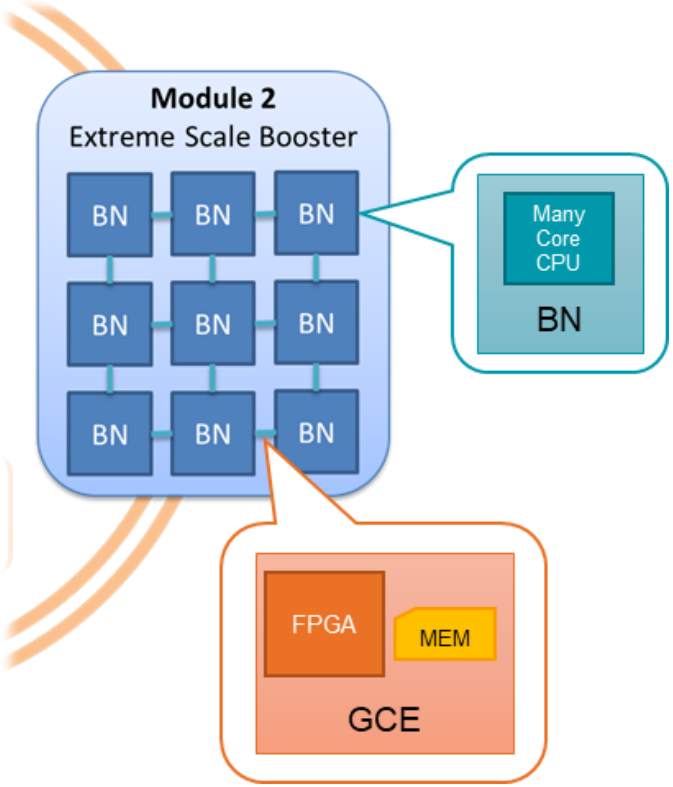

- ESB production module already implemented in JURECA supercomputer at the Jülich Supercomputing Centre (JSC)
- Global Collectives Engine (GCE) for speeding up MPI collectives.

Clustering and Indexing

(1) The point cloud dataset is loaded with parallel I/O using HDF5 in the DEEP-EST Scalable Storage Service Module (SSSM) and the DEEP-EST CLUSTER module
(2) The indexing through sorting and cost heuristic small computing elements takes advantage of the FPGA in the DEEP-EST Global Collective Engine (GCE) module in combination with MPI collectives to distribute the points equally among the DEEP-EST CLUSTER
(3) Clustering with HPDBSCAN is performed on the DEEP-EST CLUSTER locally (OpenMP) for shared memory elements and the load provided by (2)
(4) Merging the different computed clusters on chunk edges according to specific rules using halos across nodes is performed on the DEEP-EST CLUSTER globally (MPI)
(5) Cluster ID and noise ID are written to the HDF5 file but it can also be written to the Network Attached Memory (NAM) for further study, e.g. level of detail (LoD)

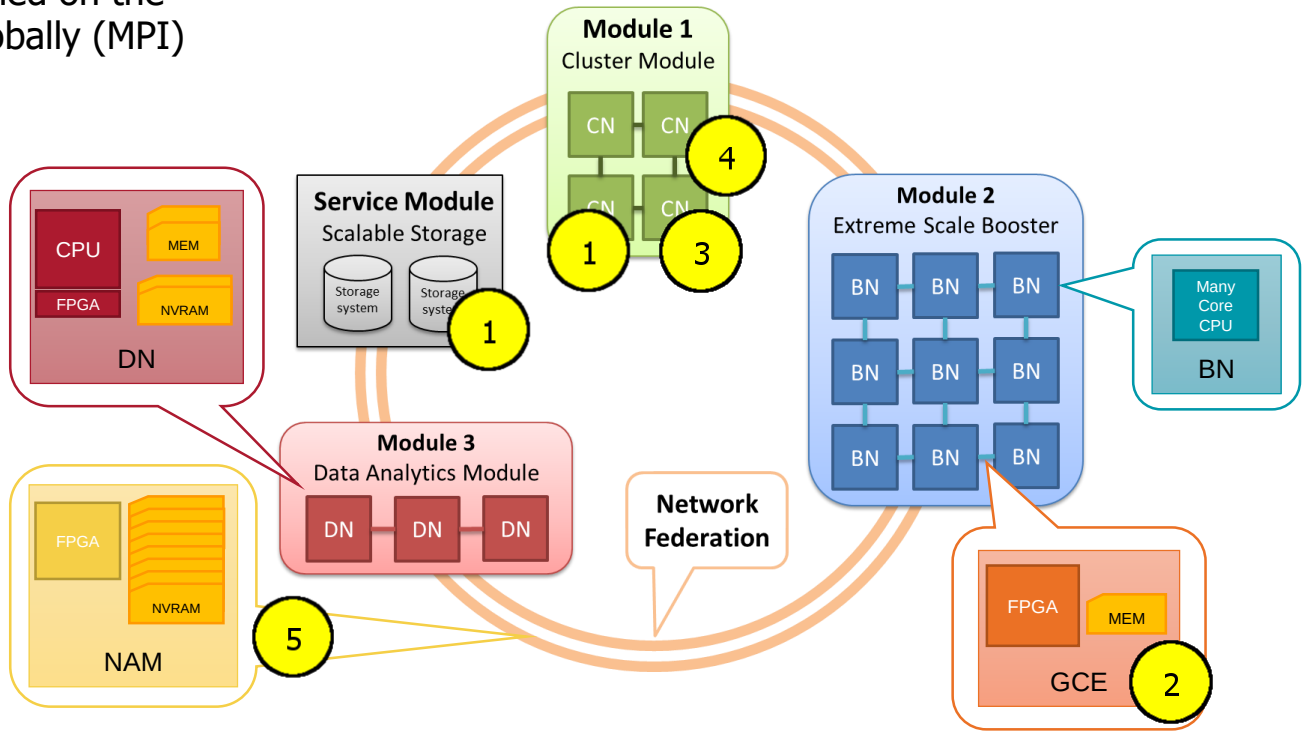

LoD and cLoI studies

(1) The point cloud dataset is loaded with parallel I/O using HDF5 in the DEEP-EST SSSM and the ESB

(2) After clustering following Workload A, the data reside in the DEEP-EST NAM as a fixed number of levels of importance (w.r.t. detail/scale)

(3) Selected point cloud LoD studies require continuous levels of importance using importance values of a point regarded as an added dimension to space and time using n-D space filling curves or tree structures whereby the latter may take advantage of MPI collectives using the DEEP-EST GCE enabling small computing modifications to the original clustered datasets in combinations with ESB

(4) The different data set results of the various modifications towards continuous levels of importance (cLoI) can be placed in different sections of the DEEP-EST NAM in order to take advantage of a variety of (semi-) continuous and spatio-temporal representations (e.g. zoom-in/out) for scientific studies.

(5) Based on the CLoI data available in the DEEP-EST NAM the DEEP-EST CM could use this information to re-cluster the data but on a modified dataset and/or using different parameters

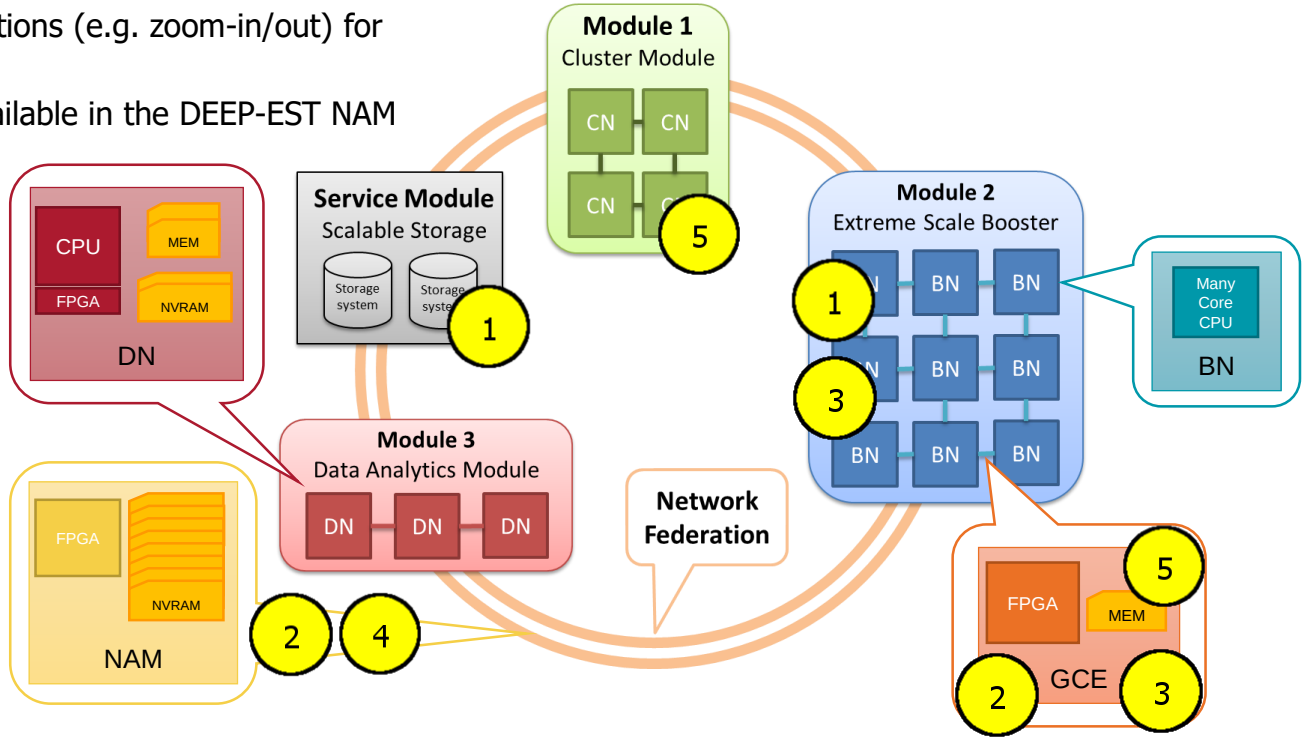

Scaling DBSCAN towards exascale computing for clustering of big data sets

Thank you for listening!

Any questions?

This project has received funding from the European Commission's Horizon 2020 Programme for research, technological development and demonstration under Grant Agreement n°754304

UNIVERSITY OF ICELAND
SCHOOL OF ENGINEERING AND NATURAL SCIENCES
FACULTY OF INDUSTRIAL ENGINEERING, MECHANICAL ENGINEERING AND COMPUTER SCIENCE

JÜLICH FORSCHUNGSZENTRUM

DEEP Extreme Scale Technologies