



From Exascale Supercomputing to FAIR data

Why (almost) everyone uses GPUs
and how to get a DOI for your dataset

Helmut Neukirchen
<https://uni.hi.is/helmut>

DOI: [10.23728/b2share.a6a4682fe1f74b32b8b67948f7ce6965](https://doi.org/10.23728/b2share.a6a4682fe1f74b32b8b67948f7ce6965)



UNIVERSITY OF ICELAND
FACULTY OF INDUSTRIAL ENGINEERING,
MECHANICAL ENGINEERING AND COMPUTER SCIENCE

Text on these slides:



Where sources are provided as reference, images
are copyright according to these source and
reproduced here under fair use principles.

The research leading to these results has received funding from
the European Union's Horizon 2020 research and innovation
programme under grant agreements n° 754304 and 857652



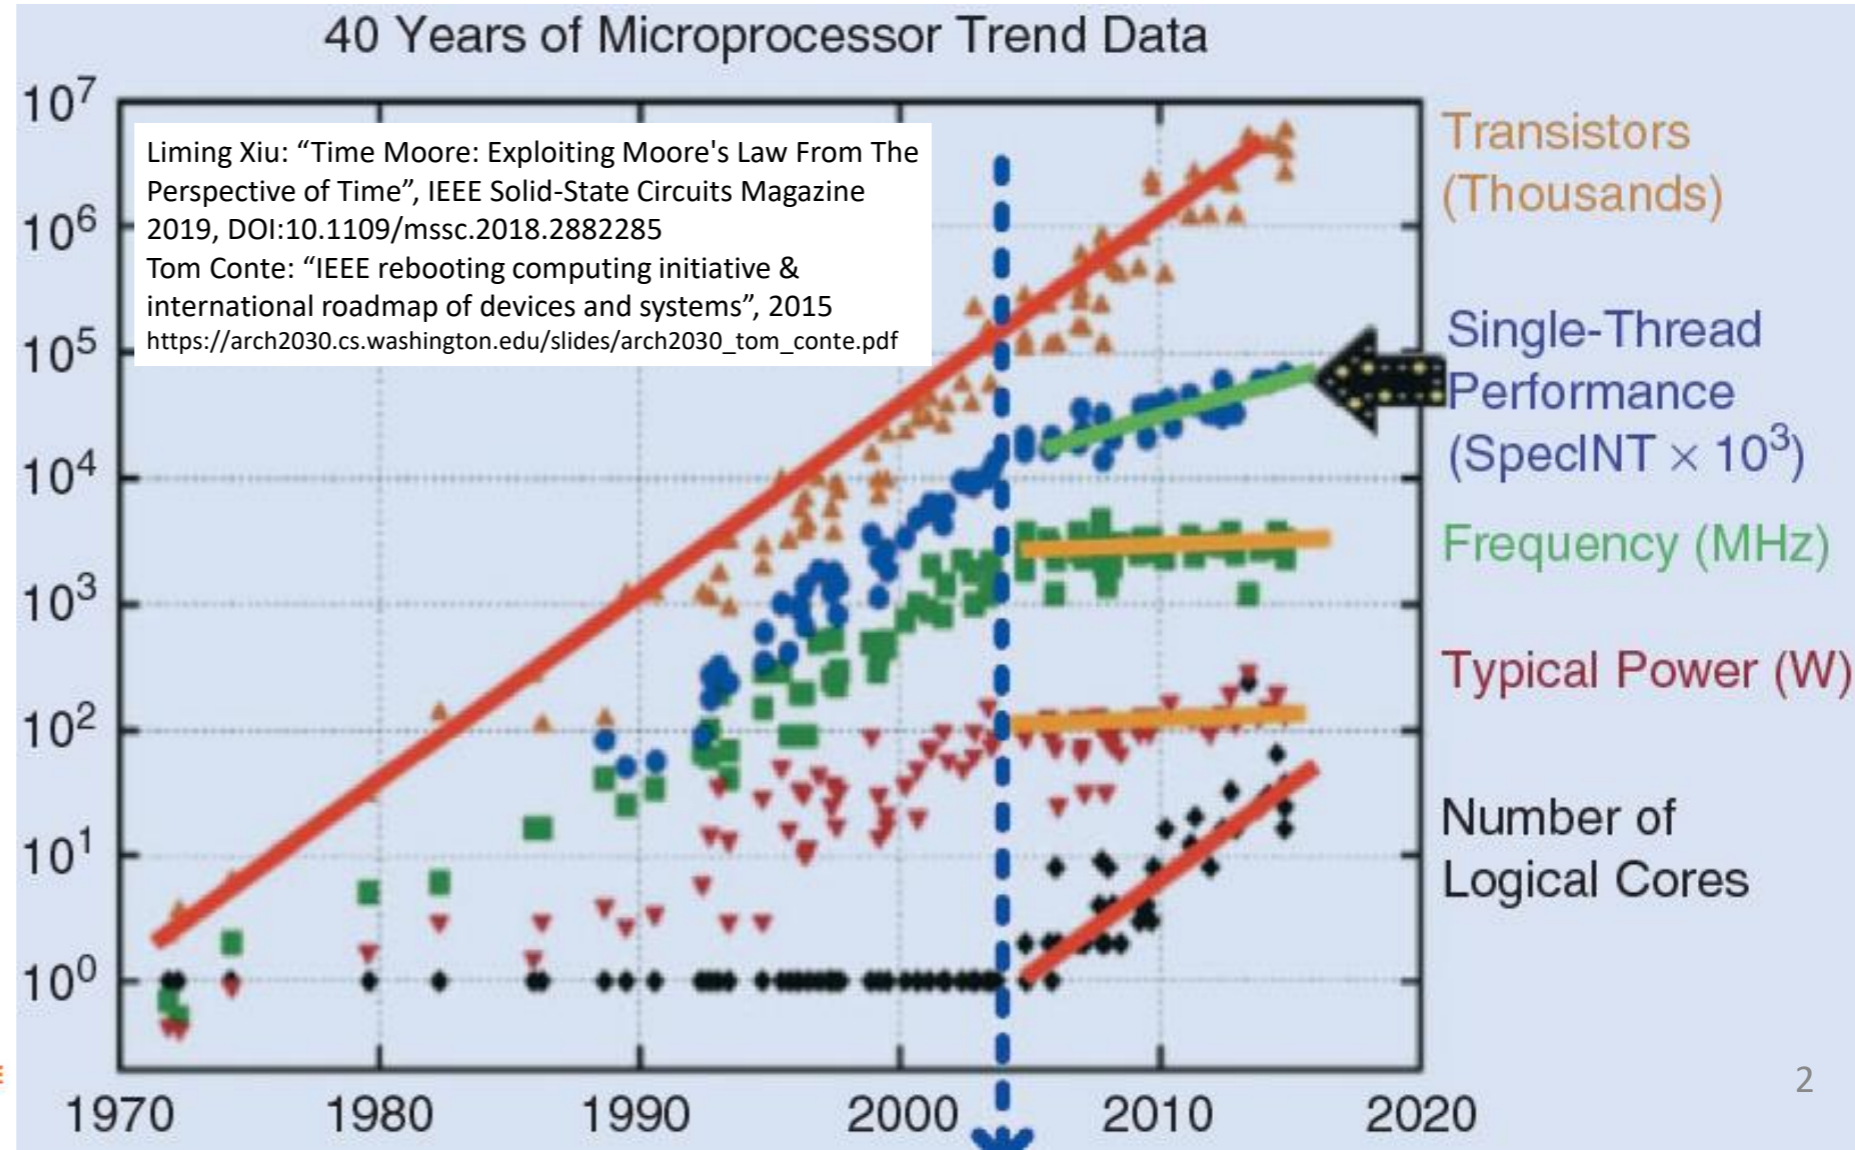


Moore's law

- Complexity (=number of transistors) in an integrated circuit doubles every two years.

Gordon Moore: "Cramming more components onto integrated circuits", Electronics, McGraw-Hill, 1955

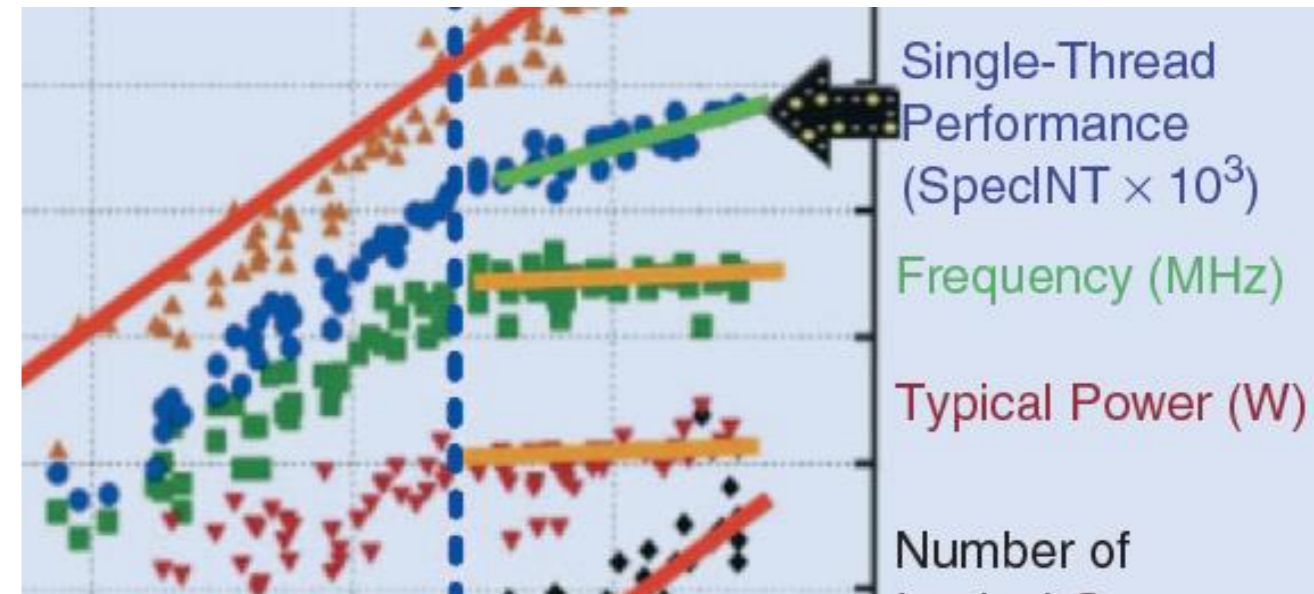
- CPU Clock speed & single-thread performance doubled as well...
 - ...until 2005.





Why Single Thread Performance Stalled

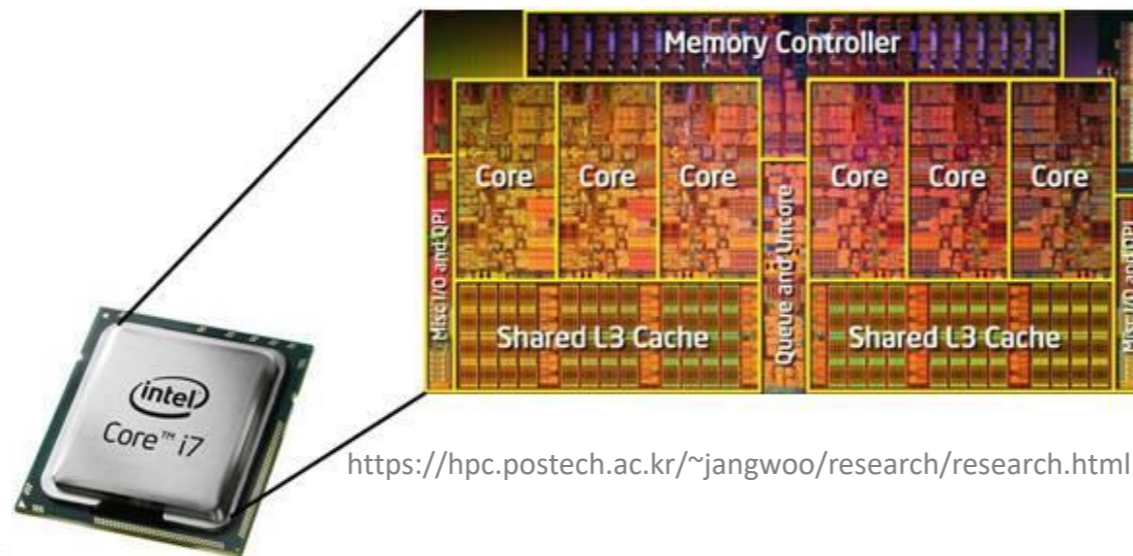
- Single-thread performance \approx Clock speed \sim Power
 \Rightarrow Silicon gets too hot!
- Still increase in single-thread performance, but limited.
 - No doubling every two years.



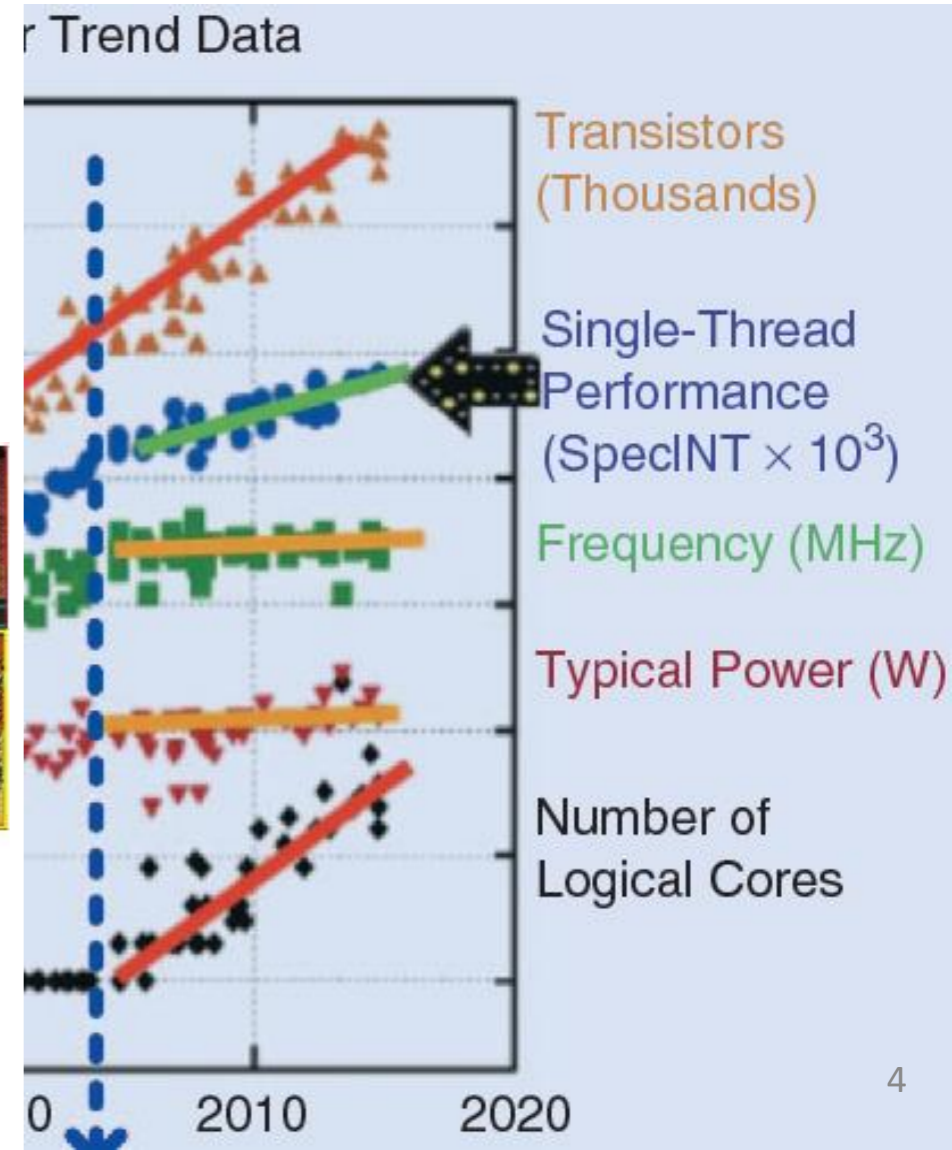


Doubling of Number Transistors Still True

- For what is the doubling of numbers of transistors used?
 - Multicore processors!



<https://hpc.postech.ac.kr/~jangwoo/research/research.html>

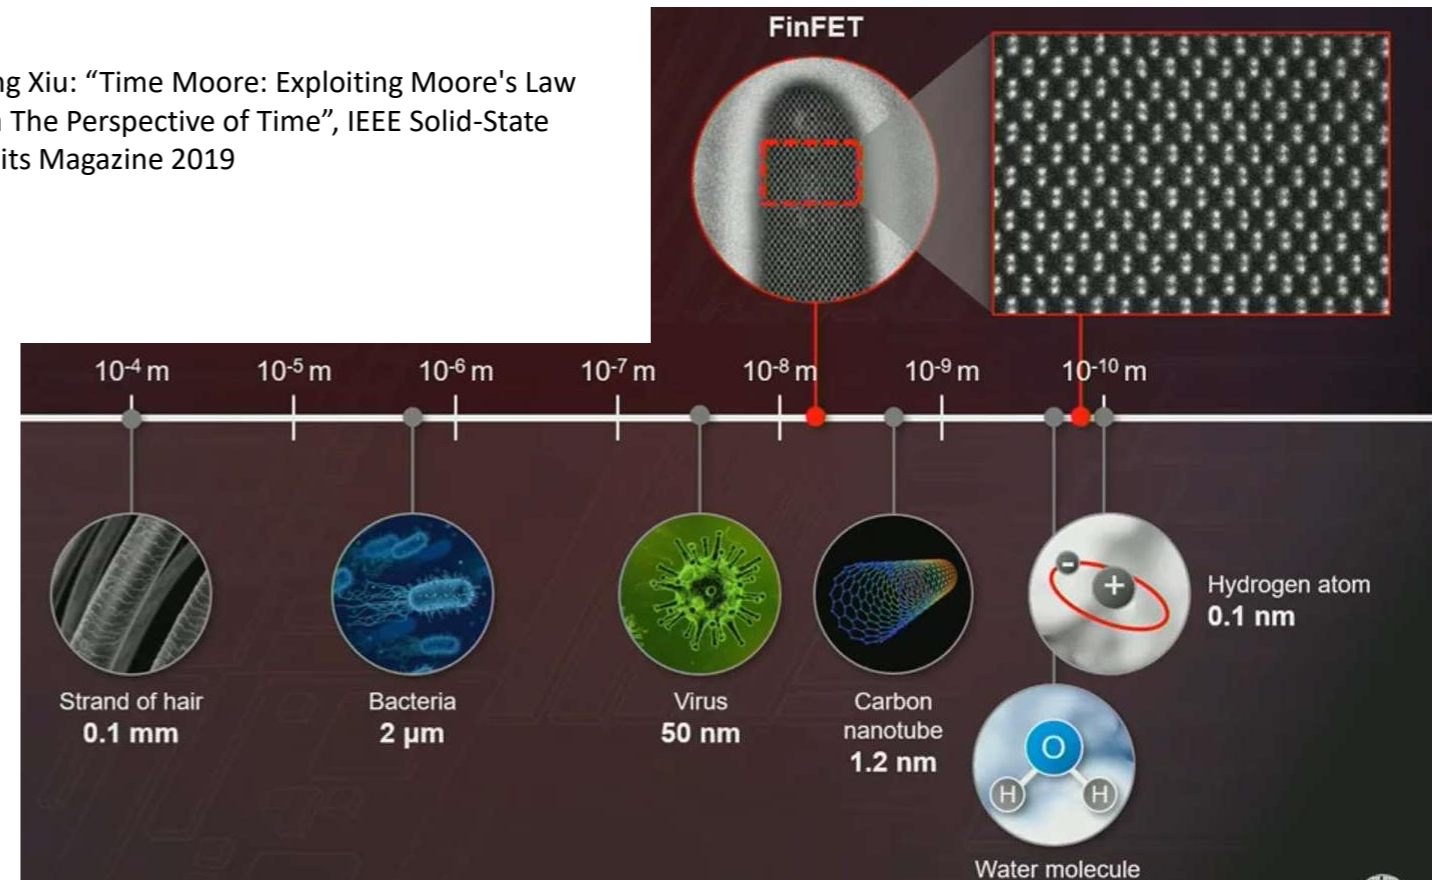




The End of Moore's law

- Currently: 7 nm semiconductor manufacturing process.
- 5 nm to start in 2020,
 - 5 nm = nine silicon crystal unit cells.
- 3 nm promised for 2022.
- Expected to reach atom level limits in 2025-2030.
 - The end of Moore's law?!
 - No need to buy new computers after 2025? (New CPU generations will not be faster anymore!?)

Liming Xiu: "Time Moore: Exploiting Moore's Law From The Perspective of Time", IEEE Solid-State Circuits Magazine 2019



<https://www.cnet.com/news/processor-progress-is-alive-and-well-tsmc-builder-of-apple-iphone-chips-says>





Beyond Moore's law

- From 2D to 3D silicon to continue increasing density (in particular memory).
- Smaller structures based on carbon nanotubes instead of silicon.
- Increase frequency: superconducting cryogenic computing.
- New paradigms:
 - Neuromorphic computing (e.g. human brain project),
 - Quantum computing.
 - (More in talk by Morris Riedel: “Demystifying Quantum Computing”, UTmessen 2020)
- Accelerators (e.g. GPUs):
 - Standard silicon manufacturing, but specialised hardware that is faster (and more energy efficient) for certain problems than generic CPUs.
 - Still always CPU needed for parts/problems not supported by accelerator.






Accelerators

- Graphics Processing Unit (GPU):
 - Do the same numerical operation in parallel to many pixels stored in high-bandwidth graphic memory.
 - The same approach fits well for many other parallel numerical problems, e.g. training a neural network.
- Google's Tensor Processing Unit (TPU):
 - Application-Specific Integrated Circuit (ASIC) for neural network machine learning operations.
- Field-Programmable Gate Array (FPGA):
 - Array of logic blocks that can be configured in the field.
- Downside: Need to rewrite code for accelerators!

Peak Floating Point Operations Per Second (FLOPS)





Supercomputing & Accelerators

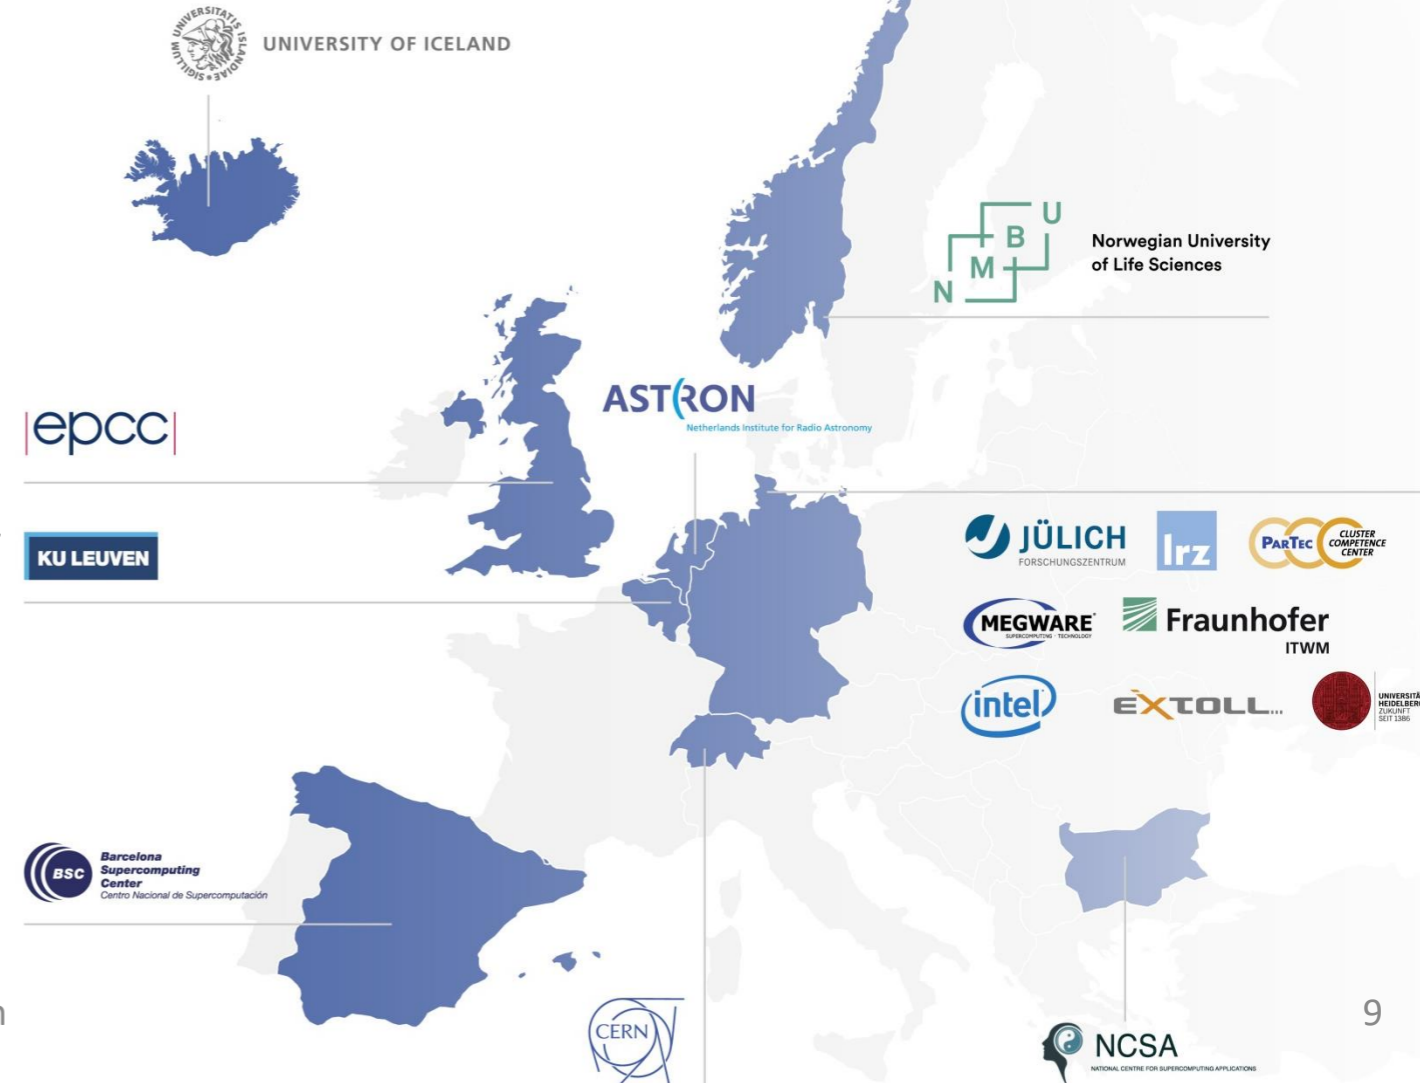
- Latest supercomputers: powerful CPUs each accompanied by GPUs.
 - More FLOPS,
 - Better energy efficiency.
 - 200 PetaFLOPS peak.
- Exascale expected 2023:
 - ExaFLOPS=10¹⁸ FLOPS

<https://www.top500.org/lists/2019/11/>

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States 2018	2,414,592	148,600.0	200,794.9	10,096
2	Sierra - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States 2018	1,572,480	94,640.0	125,712.0	7,438
3	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China 2016	10,649,600	93,014.6	125,435.9	15,371
4	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 , NUDT National Super Computer Center in Guangzhou China 2013	4,981,760	61,444.5	100,678.7	18,482
5	Frontera - Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR , Dell EMC Texas Advanced Computing Center/Univ. of Texas United States 2019	448,448	23,516.4	38,745.9	
6	Piz Daint - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray/HPE Swiss National Supercomputing Centre (CSCS) Switzerland 2016/2017	387,872	21,230.0	27,154.3	2,384

Dynamical Exascale Entry Platform - Extreme Scale Technologies (DEEP-EST)

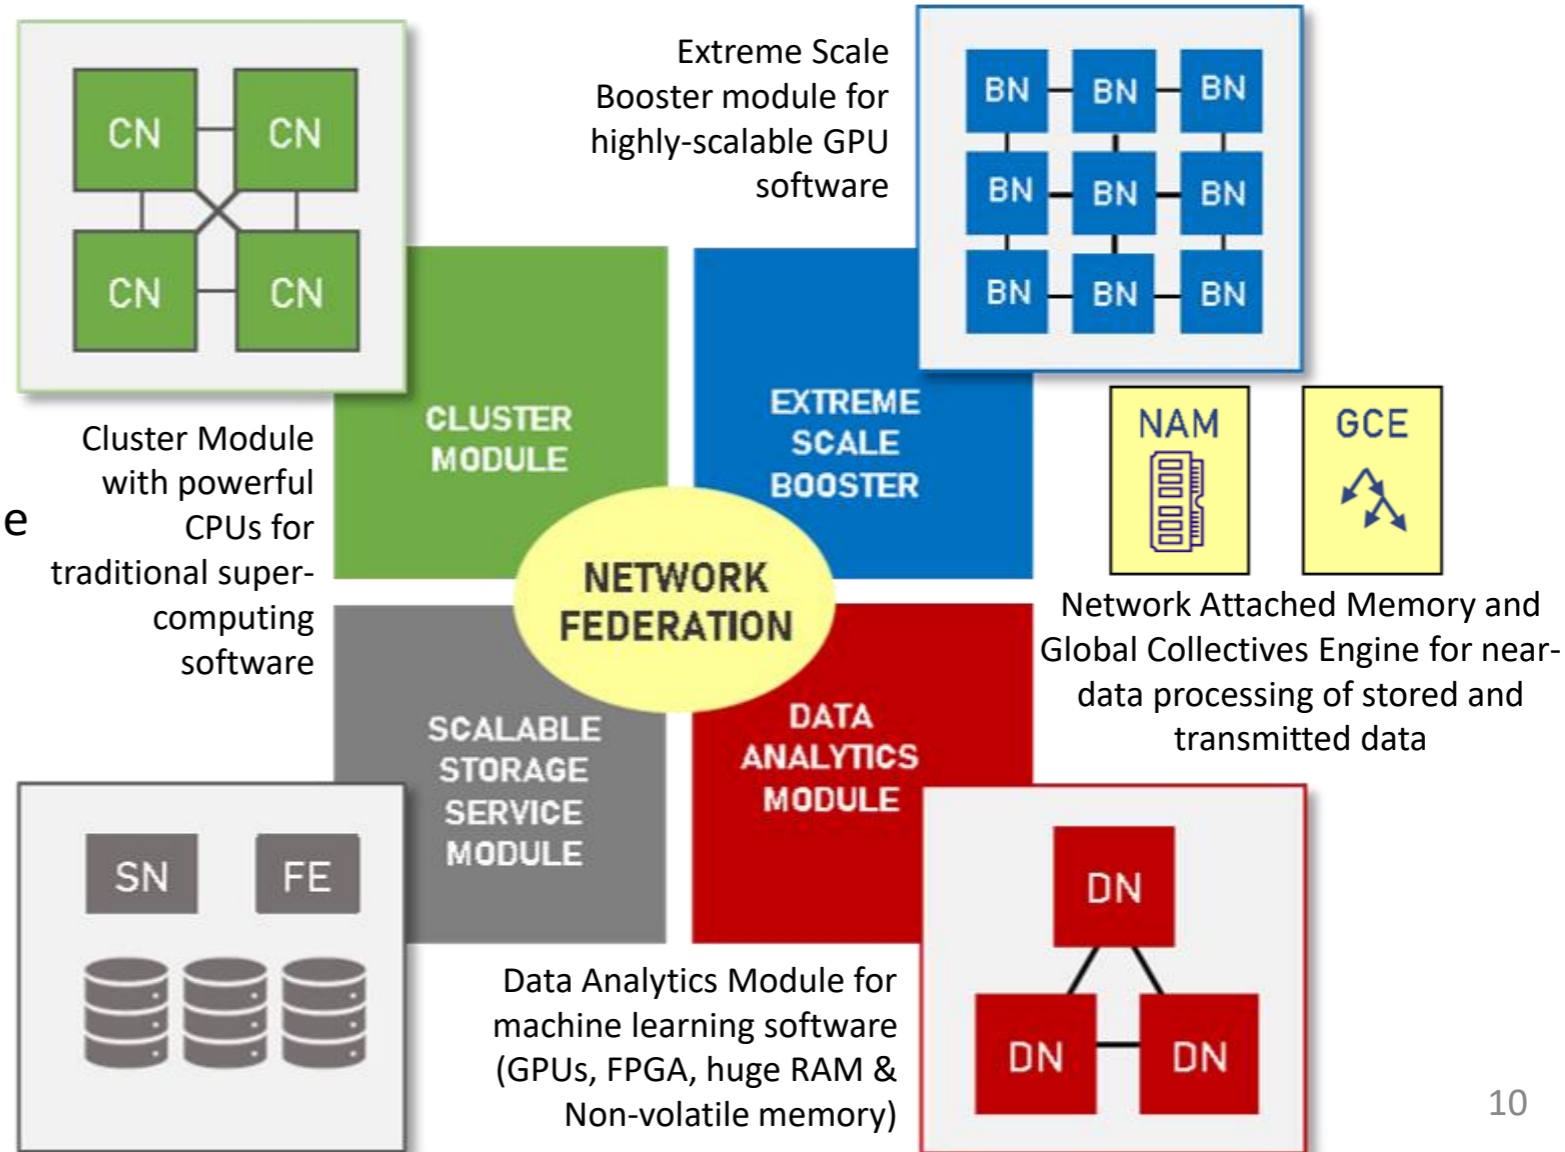
- H2020 project to build the blueprints of future Exascale supercomputers.
- Problems addressed, e.g.:
 - Current supercomputers have on each node a powerful CPU *and* powerful GPUs.
 - But: problems are often *either* CPU-intensive *or* GPU-intensive.
 - Waste of money and energy: all powered, but only one of these is productive.





DEEP-EST Modular Supercomputing Architecture

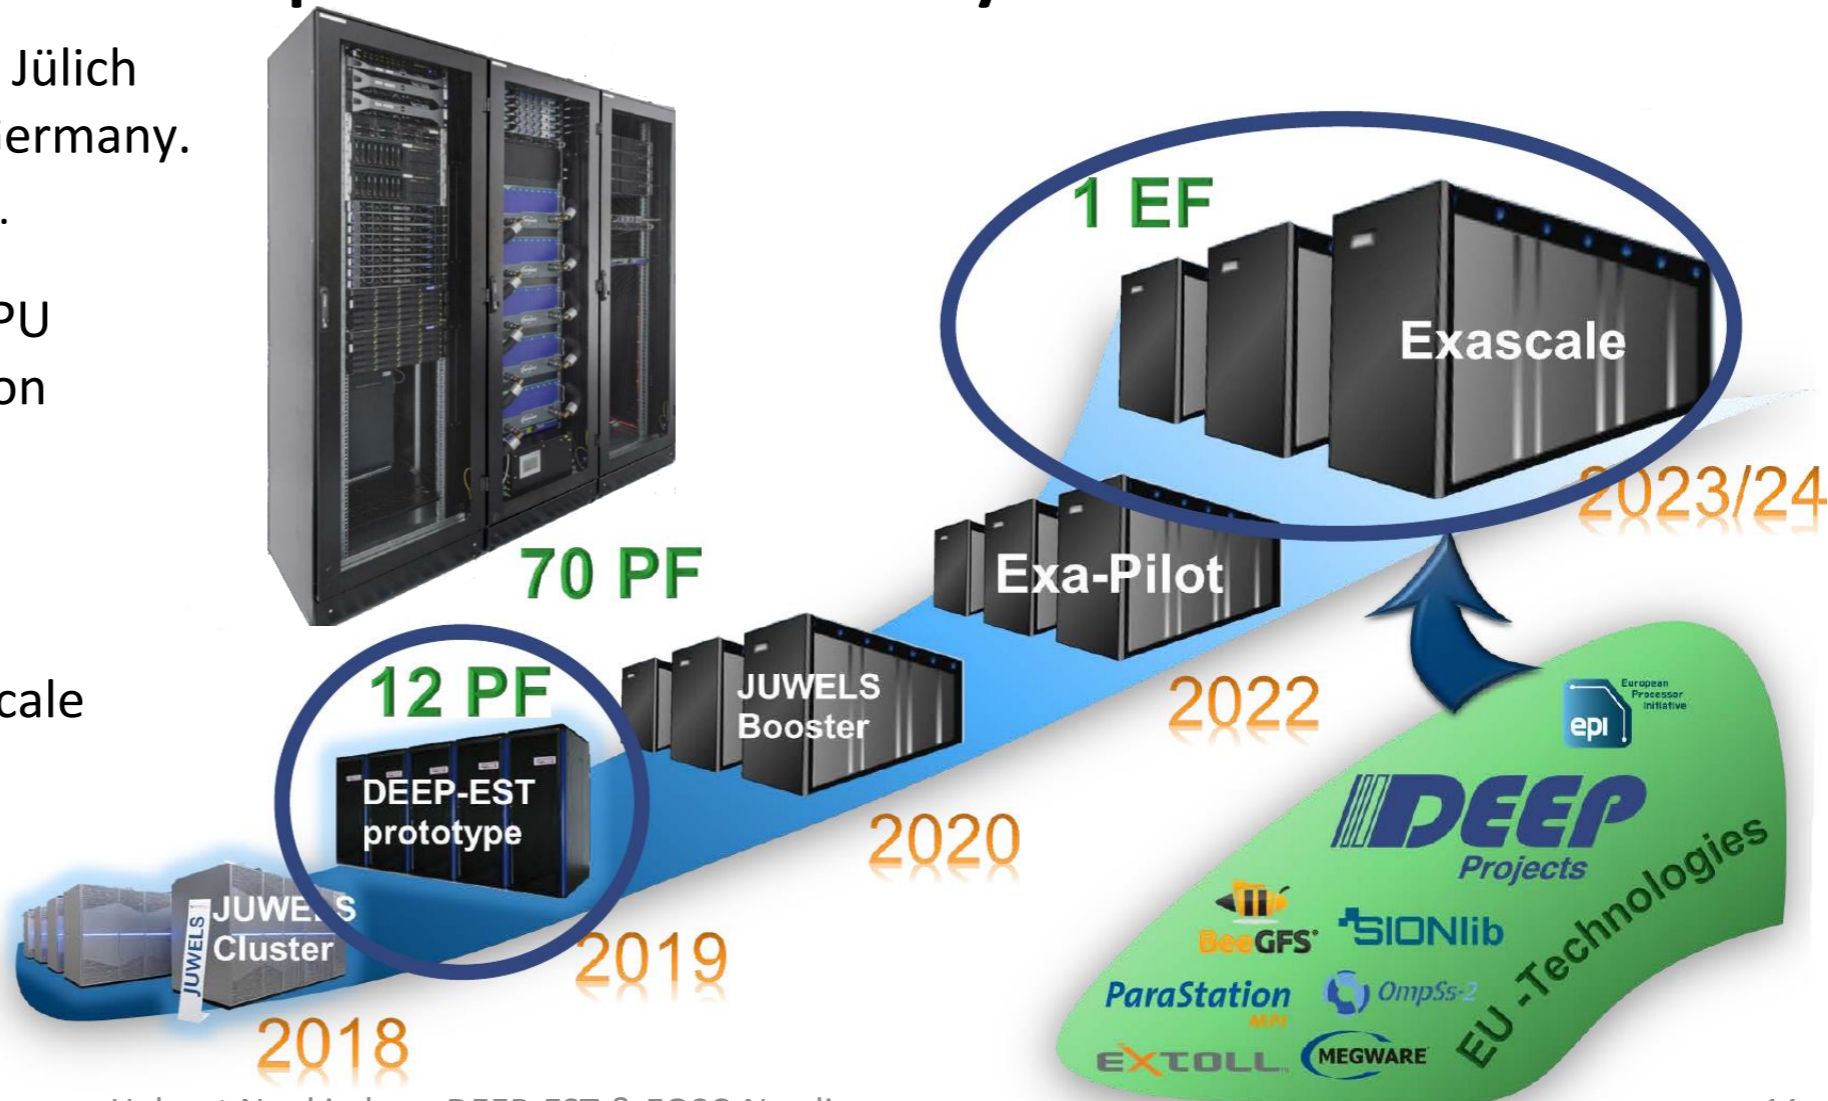
- Modules to support different supercomputing applications in the most (energy) efficient way:
 - Cluster Module with powerful CPUs,
 - Booster Module with powerful GPU (weak CPU only to drive GPU),
 - Data Analytics Module (DAM) for machine learning with huge memory,
 - future: quantum computing module?
 - all connected by a very fast network:
 - Same application can use different modules in parallel.
 - Near-data processing inside.





Exascale Roadmap Influenced by DEEP-EST



- DEEP-EST prototype installed at Jülich Supercomputing Centre (JSC), Germany.
 - Final installation: 12 PetaFLOPS.
- DEEP-EST blueprints used for GPU booster of current JSC production supercomputer JUWELS:
 - Will add 60 GPU PetaFLOPS to current 10 CPU PetaFLOPS.
- JSC is optimistic to host an Exascale system by 2023/24.





DEEP-EST – University of Iceland

- Scalable, parallel machine learning implementations.
- Case studies to evaluate benefits of Modular Supercomputing Architecture:
 - Classification of data using Deep Neural Networks.
 - Pointless to compete with Google, Facebook, Uber: Use their frameworks (Tensorflow, Horovod) on GPUs but with fast memory & storage hardware of DEEP-EST.
 - Classification of data using Support Vector Machine (SVM).
 - SVM useful if not enough training data for neural network available.
 - Clustering of data using Density-Based Spatial Clustering of Applications with Noise (DBSCAN).
 - Detect clusters in point data, e.g. LIDAR scans (point clouds).



NextDBSCAN (Ernir Erlingsson):

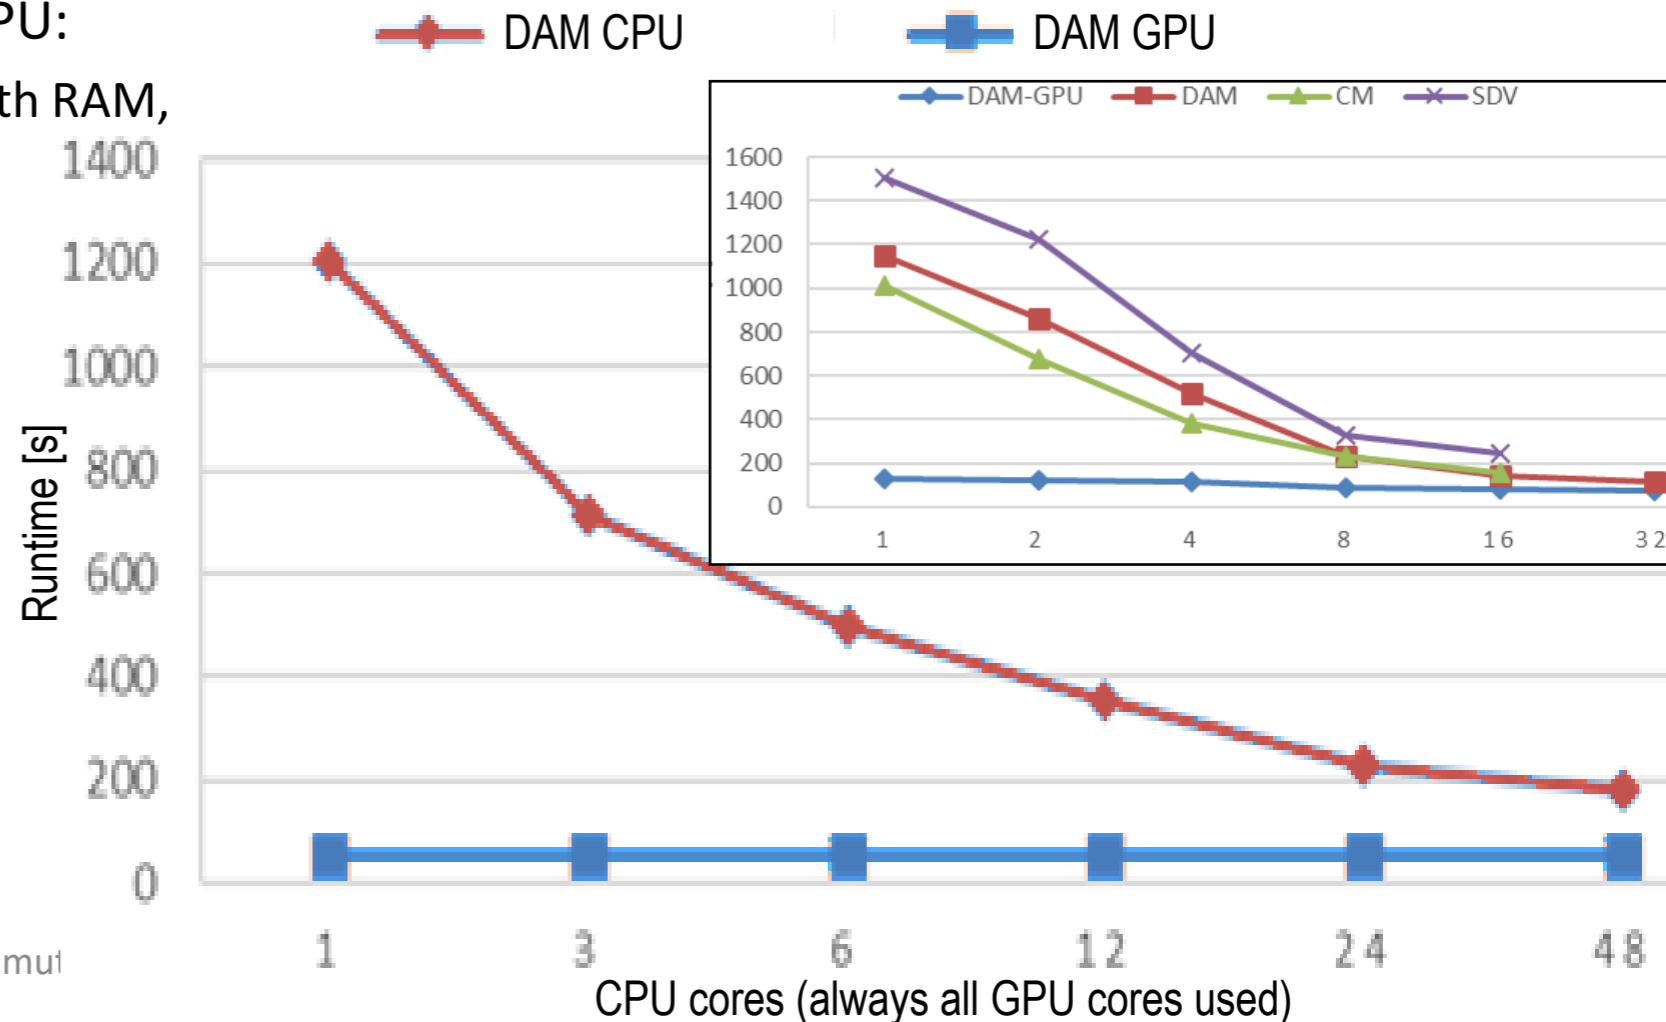
Scalable CPU & GPU Implementation of DBSCAN

- Depending on clustering parameters:
GPU sometimes faster, sometimes CPU:

- GPU limited to 32 GB high-bandwidth RAM,
- CPU: 384 GB RAM
+ 3 TB NVDIMM.
 - NVDIMM: slower than RAM
but faster than SSD;
by using RAM as cache:
only 10%-20% slower
than RAM.

⇒ CPU version can cache
calculations that GPU needs
to recalculate in different
phases of the DBSCAN algorithm.

527 162 888 3D data points, single node/shared memory (dual CPU, 1 GPU)





And Now For Something Completely Different: From Computing to Data






The European Open Science Cloud (EOSC)

- H2020 funded range of projects.
- Environment for researchers to store, analyse and re-use data.
- Open science – reproducibility, integrity, e.g.:

- Open methodology,
- Open source,
- Open access,
- Open/FAIR data.

Vicente-Saez, R.,
Martinez-Fuentes, C.
Open Science now:
A systematic literature review
for an integrated definition.
J Business Research 88 (2018),
DOI:10.1016/j.jbusres.2017.12.043

Woelfle, M., Olliaro, P., Todd, M.
Open science is a research accelerator.
Nature Chem 3, 745–748 (2011).
DOI:10.1038/nchem.1149





Open Data vs. FAIR Data

<http://opendatahandbook.org>

- Open data: data that can be used, reused, and redistributed freely by any person.
- FAIR data: Wilkinson, M. D. et al.. The FAIR Guiding Principles for scientific data management and stewardship. Sci. Data 3:160018 (2016) DOI:10.1038/sdata.2016.18
<https://www.openaire.eu/how-to-make-your-data-fair>
 - Findable via rich metadata including a global ID (e.g. DOI) to make it identifiable and locatable.
 - Accessible: always available and obtainable; even if the data is restricted, the metadata is open.
 - Interoperable via a defined syntax and semantics, allowing exchange between researchers.
 - Reusable: sufficiently described; least restrictive license, allowing the widest reuse possible.

- European Commission advocates Data Management Plan based on FAIR data.

- Example metadata using Dublin Core format (simplified):

```
<metadata
  <title>From Exascale Supercomputing to FAIR data</title>
  <description>Presentation about DEEP-EST and EOSC-Nordic</description>
  <creator id="https://orcid.org/0000-0001-8595-3748">Neukirchen, H.</creator>
  <date>2020</date>
  <rights>http://creativecommons.org/licenses/by/4.0/legalcode</rights>
  <type>Text; Image</type>
  <format>application/pdf</format>
  <identifier>https://dx.doi.org/10.23728/b2share.
    b6c1b824a8f84db2bbc1d08a77a4e187</identifier>
</metadata>
```





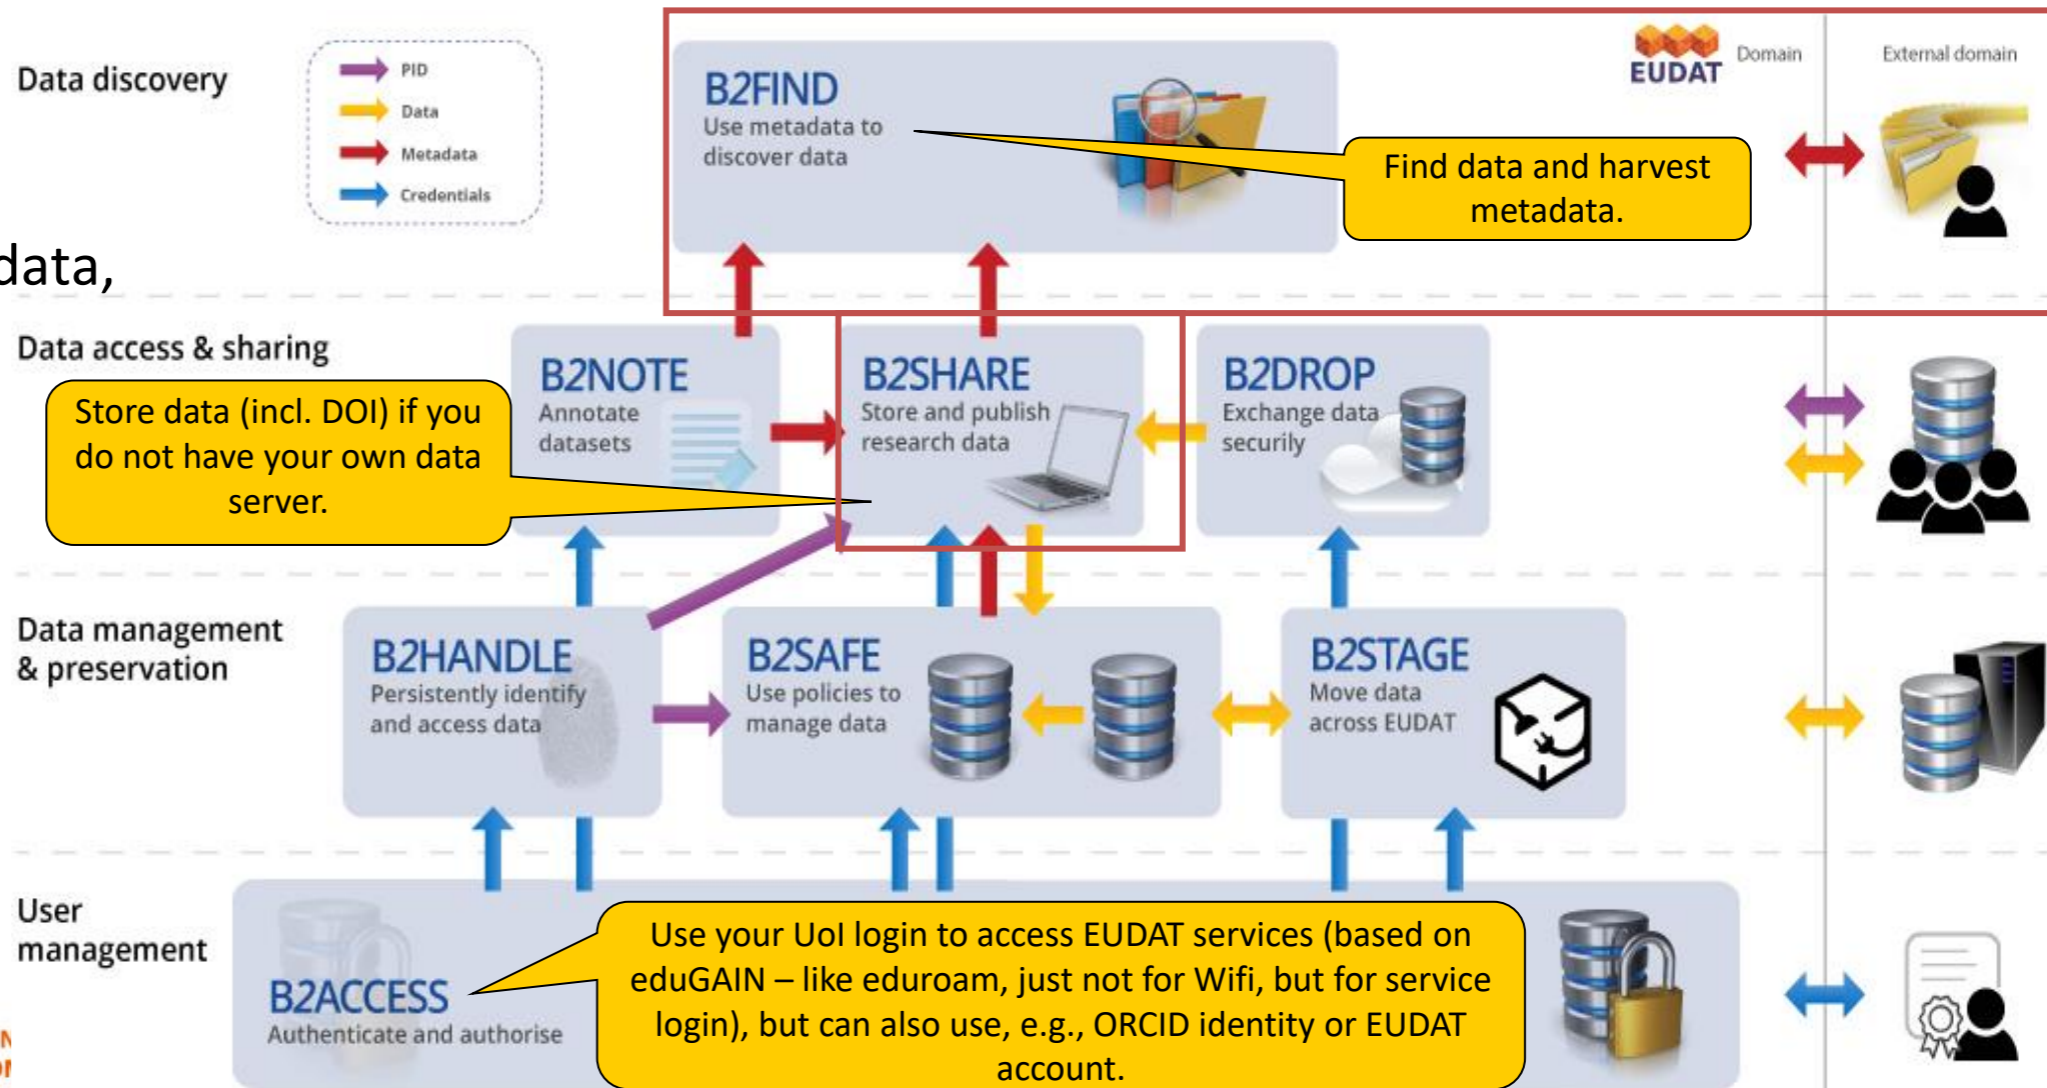
EOSC-Nordic

- H2020 project: part of European Open Science Cloud.
- Objectives, e.g.:
 - Create knowledge hub (e.g. how to create a data management plan).
 - Map Nordic research data repositories and services.
 - Make Nordic data FAIR by adding them to EOSC metadata catalogue.
 - Based on EUDAT Collaborative Data Infrastructure.



EUDAT Services

- H2020 project (part of EOSC).
- B2SHARE: long-term storage of research data,
- DOI for data.
- B2FIND: metadata harvesting & search.
- Where login required: use your existing account.





B2FIND: Finding Data

- B2FIND harvests metadata from various sources.
 - We need to add metadata for Icelandic datasets as sources!
- Example usage:
 - <https://eosc-portal.eu/>
 - ACCESS EOSC SERVICES & RESOURCES
 - Sharing & Discovery
 - B2FIND
 - Access the Service
 - You end up at B2FIND: <http://b2find.eudat.eu>
 - Enter your search keywords to search metadata and find DOI to access the data.



B2FIND: Finding Data

The screenshot shows the B2FIND website interface. At the top, there are navigation links for GUIDELINES, COMMUNITIES, FACETED SEARCH, and ABOUT. Below the navigation is a search bar containing the keyword 'geothermal'. A yellow callout bubble points to the search bar with the text 'Metadata search keyword(s)'. Below the search bar, the results are displayed as a list of datasets. The first result is titled 'Greenland geothermal heat flux distribution and estimated Curie Depths, links...' and includes a map of Iceland. A yellow callout bubble points to the map with the text 'Search filters, e.g. geo, date'. Below the first result, there are two more results: '(Supplement 1) Selected chemical composition of Icelandic river waters and be...' and 'Lead isotopes of feldspar in North Atlantic Ocean sediments, supplement to: G...'. On the left side of the search results, there are filters for 'Filter by location' (with a map of Iceland) and 'Filter by time' (with start and end date pickers).

The screenshot shows the detail page for the dataset 'Greenland geothermal heat flux distribution and estimated Curie Depths, links to gridded files, supplement to: Martos, Yasmina M; Jordan, Tom A; Catalan, Manuel; Jordan, Thomas M; Bamber, Jonathan L; Vaughan, David G (in press): Geothermal heat flux reveals the Iceland hotspot track underneath Greenland. Geophysical Research Letters'. The page includes a map of Greenland, a 'Dataset extent' section, and a 'Metadata' section. A yellow callout bubble points to the 'Metadata' section with the text 'Metadata'. Another yellow callout bubble points to the 'Identifier' field in the metadata table with the text 'Access to actual data via data identifier, e.g. DOI. Will lead to the server where actual data is stored.'.

Identifier	
Source	https://doi.pangaea.de/10.1594/PANGAEA.892973
Metadata Access	http://ws.pangaea.de/oai/provider?verb=GetRecord&metadataPrefix=datacite3&



B2SHARE: Storing data and sharing via DOI

- Access it via EOSC (just like B2FIND).
 - <https://b2share.eudat.eu>
 - If you need a playground for training, use: <https://trng-b2share.eudat.eu> (created DOIs are dummies).
- To create a new dataset, you need to log in.
 - Use your University of Iceland login as identity provider (eduGAIN).
 - If that does not work, you can use many other identity providers:
 - E.g. ORCID, Facebook, Google, or create a EUDAT account.
 - ⇒ Also Icelandic researchers not being part of eduGAIN can use EUDAT!
- Create a new record.
 - You need to select a community: → Select “EUDAT”
 - Add files, fill-in metadata, save and publish. ⇒ Voilà: data stored with DOI





B2SHARE



Search records for...

HELP COMMUNITIES UPLOAD CONTACT

Store and publish your research data

Search in public datasets or register as a user to upload and publish your data!

Please use <https://trng-b2share.eudat.eu> for testing or training.


[Login or Register](#)

Create record

[Forgotten password?](#)

Title

Community



You can also update the data in an existing record by creating a new version of that record. Search for the 'Create new version' button on the record's page.





B2SHARE

B2SHARE and B2FIND demo slides and videos for Un

by [Neukirchen, Helmut](#);

Jan 7, 2020

Last updated at Jan 10, 2020

Abstract: Demo slides (PDF and Microsoft Powerpoint format) and demo videos (MP4 container -- the B2FIND v B2SHARE is 20:36 long) used for demonstrating B2FIND (finding research data based on metadata) and B2SHA including getting a DOI for it). Both are services offered by the EUDAT project (<https://eudat.eu/>) that research Iceland, can use for finding and storing data.

University of Iceland is partner in the H2020 EOSC-Nordic project (<https://www.eosc-nordic.eu/>) that aims at data discoverable and usable within EOSC, the European Open Science Cloud (<https://eudat.eu/>).

Note that you can also watch the videos as stream in your web browser via: <https://hi.cloud.panopto.eu/Panopto/List.aspx?folderID=96a64bf1-fde9-480f-ace0-ab3b00e3c4bc>

Disciplines: [4.1.16.1](#) → [Information science](#) → [Data management](#);

Keywords: [B2SHARE](#); [B2FIND](#); [EUDAT](#); [EOSC](#); [EOSC-Nordic](#);

DOI: [10.23728/b2share.9e43ad2b67464350a594d693882b54b0](https://doi.org/10.23728/b2share.9e43ad2b67464350a594d693882b54b0)

PID: [11304/db6201d3-ef87-4939-b512-6db3a203fa51](https://purl.org/urn:nbn:de:hbz:5:1-63864-p0011-7)

Add files

Drop files here, or click to select files

Add B2DROP files

Uploaded files

Name	Size
> labeled_data.csv	30.71KB

Basic fields

Community *

Titles *

 Clear
 Add

Descriptions

Files

Name	Size
> b2find_demo-2020-01-07_11.22.10.mp4	12.92MB
> b2share_demo-2020-01-07_13.16.47.mp4	41.21MB
> eudat.pdf	470.74KB
> eudat.ppt	668.67KB

Basic metadata

Open Access

License

Contact Email



Questions?

-  deep-projects.eu

-  eosc-nordic.eu



– Do you have any dataset that is worth to be made FAIR?

- Contact me: helmut@hi.is

