Technoweek 2023
Cloud use case 4: ALBA Synchrotron

Speaker: Sergio Vicente
Authors: Adara Carrión and Sergio Vicente

Agenda

- The ALBA Synchrotron
- IT infrastructure for science
  - Description of the infrastructure on-premises
  - Ongoing upgrades
- The use of the cloud at ALBA
  - The use of OCRE IaaS Framework – AWS
  - The OCRE grant for science – GCP
  - Particular use cases
- Conclusions

The ALBA Synchrotron

ALBA is the national synchrotron **LIGHT SOURCE,** located in Cerdanyola del Vallès, Barcelona.
It is used for analyzing the structure and properties of the matter.

- **10 beamlines + 4 under construction**
- **6.000 hours** of synchrotron light per year
- **2.000 academic and industrial users** per year

GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA E INNOVACIÓN

Generalitat de Catalunya
Departament de Recerca
i Universitats

Synchrotrons in the World

The ALBA synchrotron is like a giant microscope

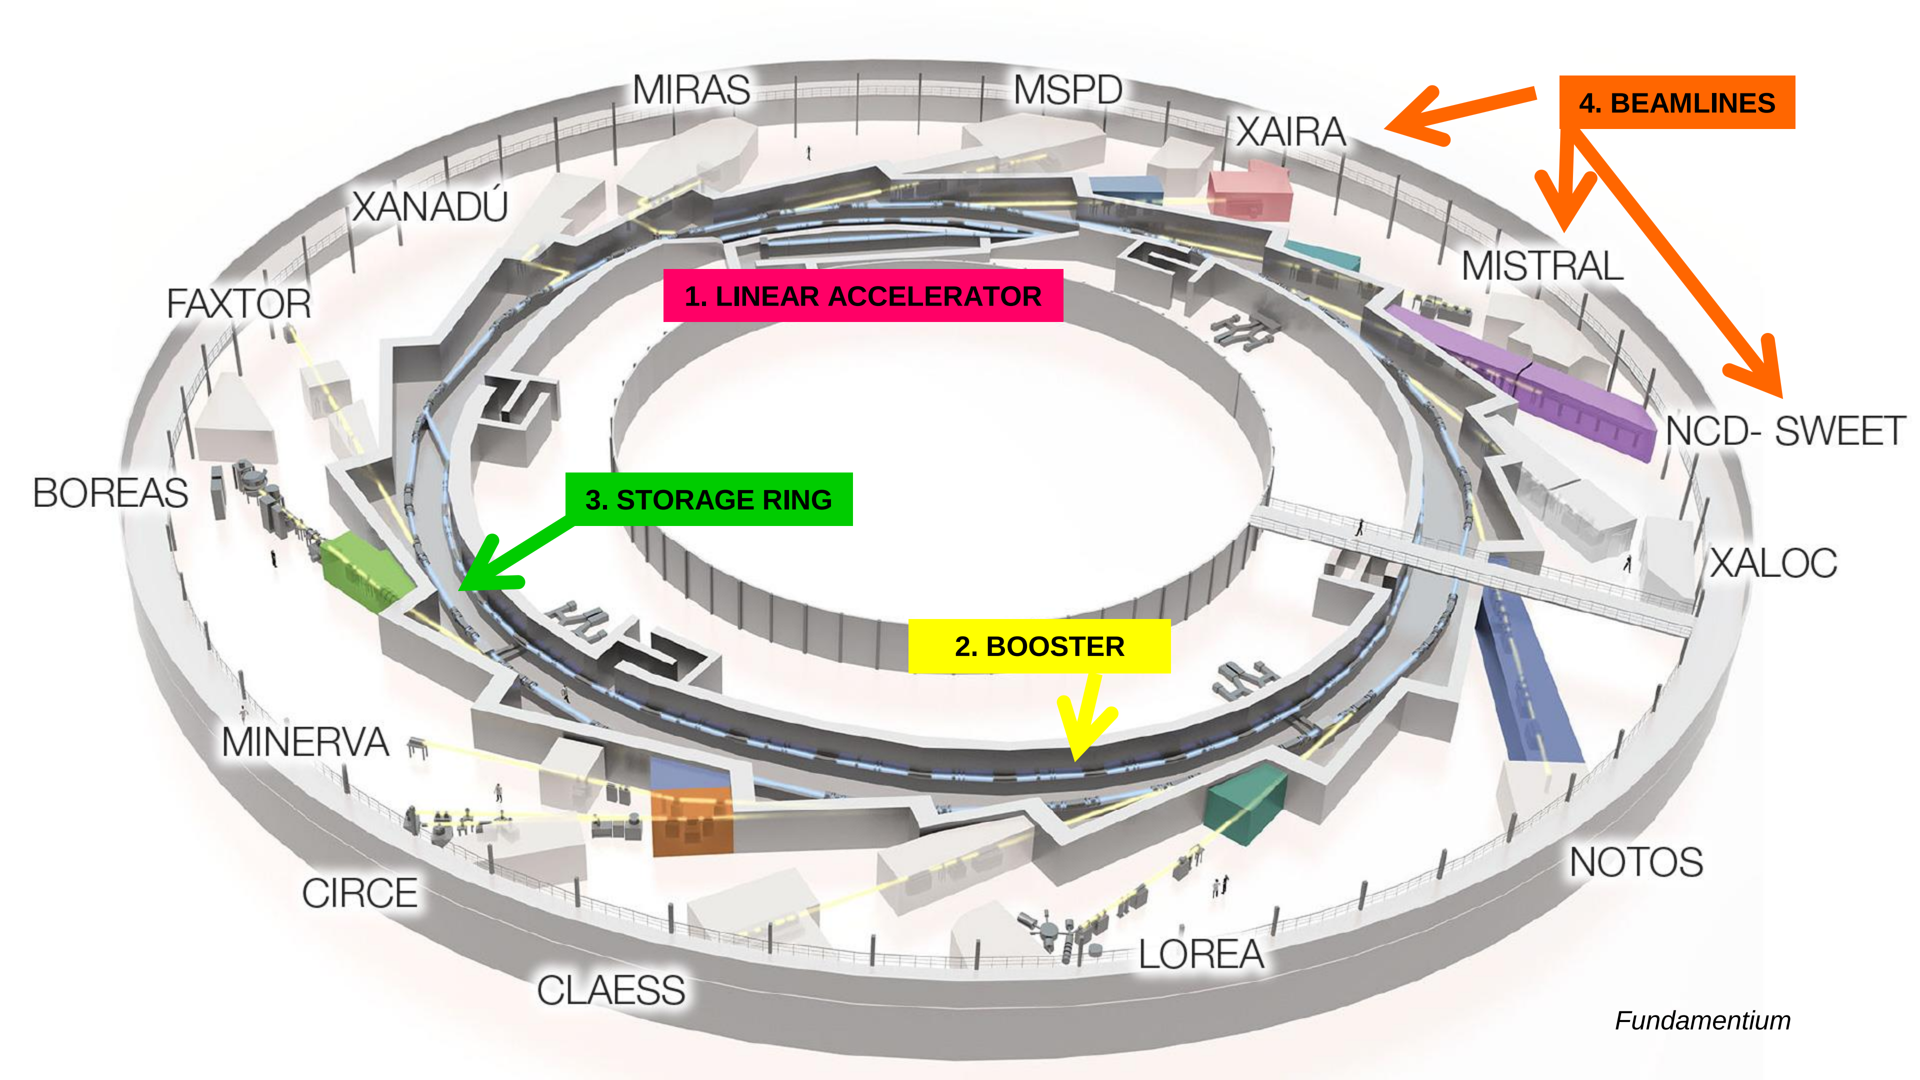
MIRAS
MSPD
XAIRA
4. BEAMLINES
XANADÚ
MISTRAL
FAXTOR
1. LINEAR ACCELERATOR
NCD- SWEET
BOREAS
3. STORAGE RING
XALOC
2. BOOSTER
MINERVA
NOTOS
CIRCE
LOREA
CLAESS
Fundamentium

We accelerate electrons

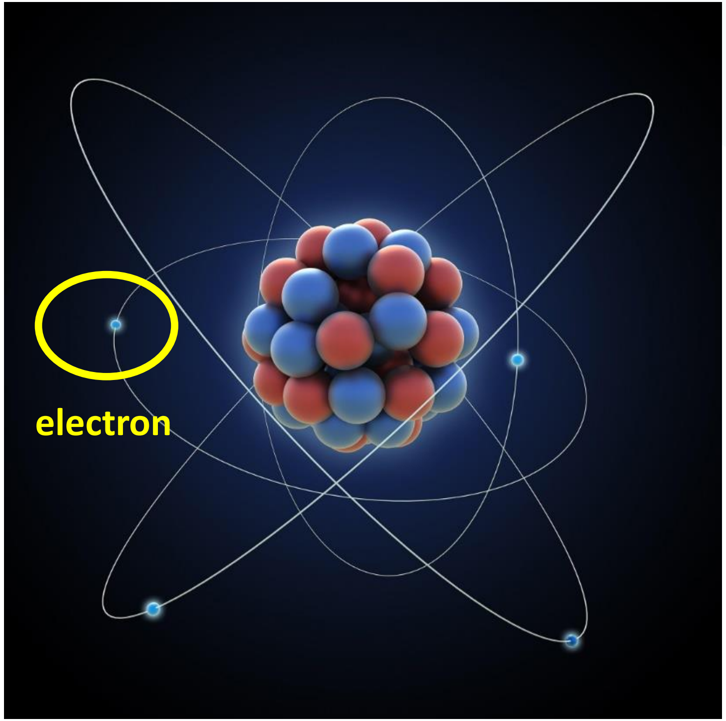

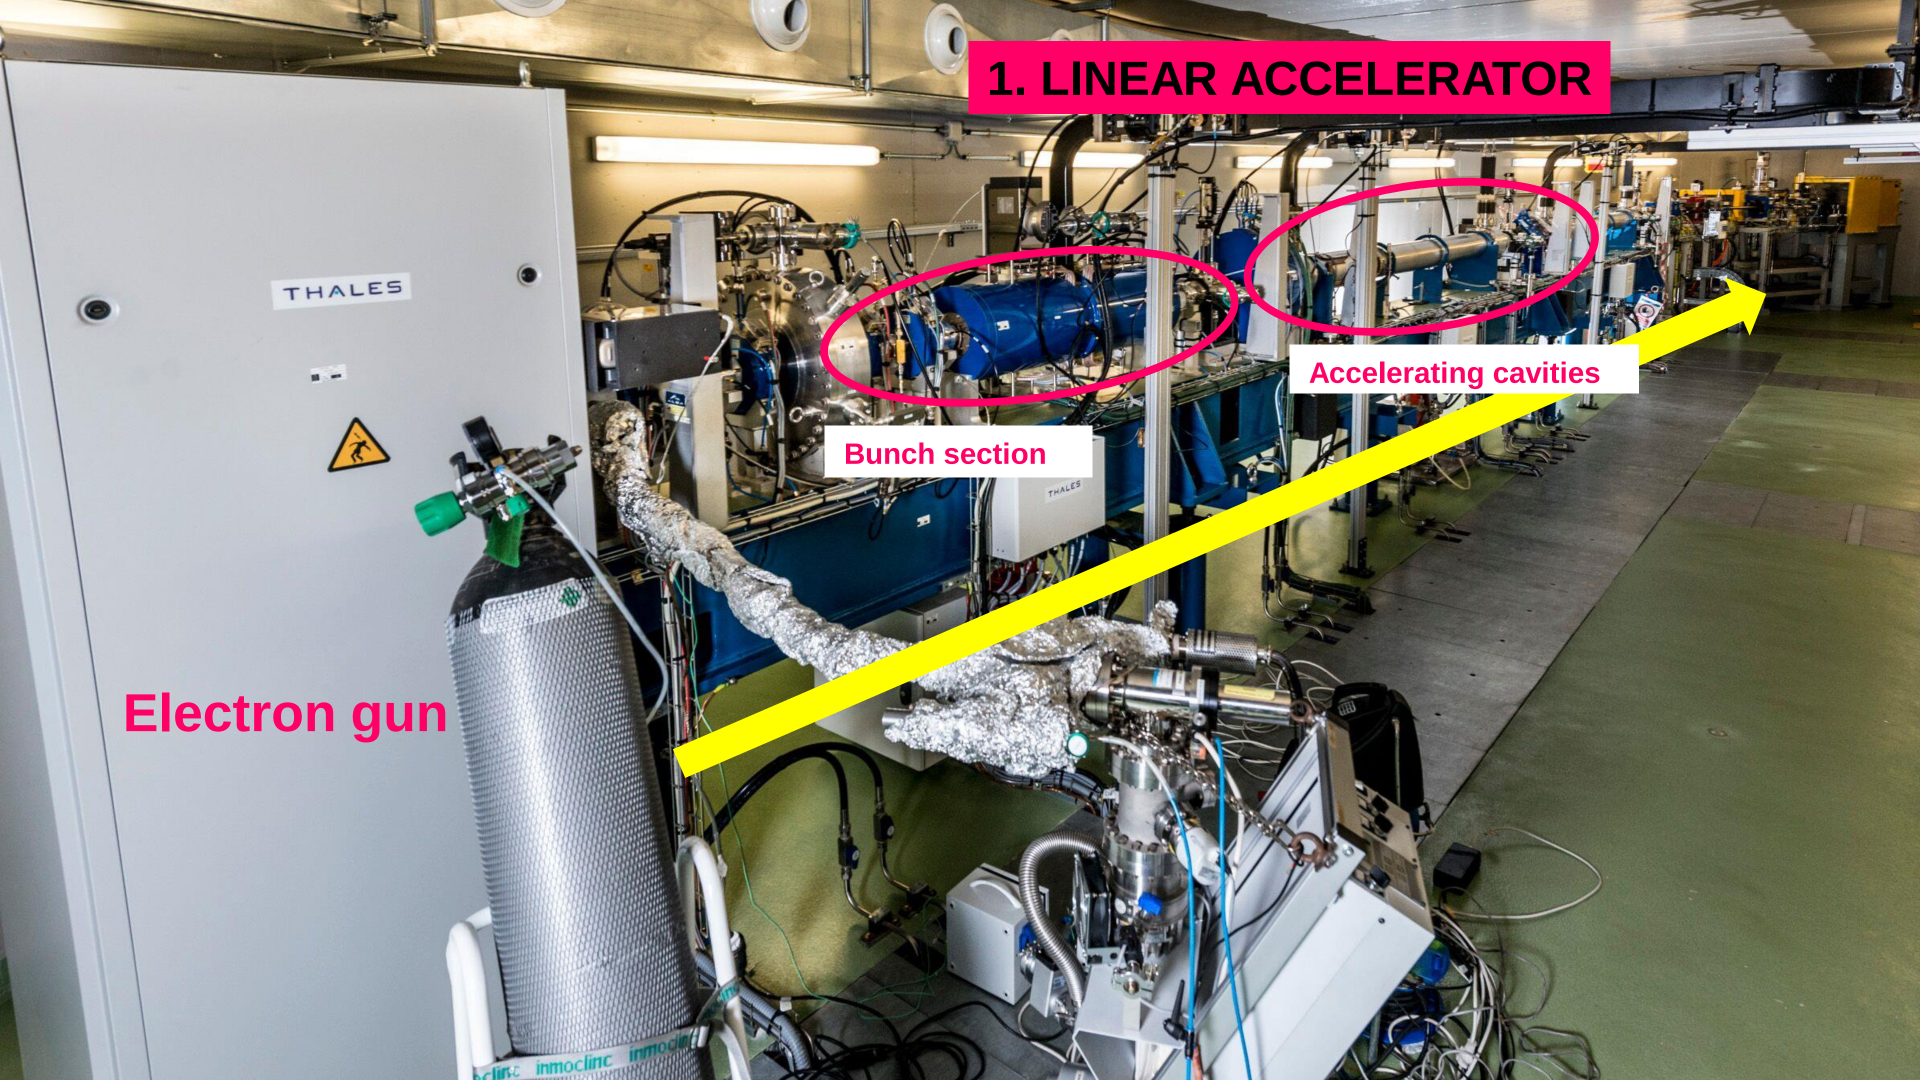
1. LINEAR ACCELERATOR
Accelerating cavities
Bunch section
Electron gun
THALES
THALES
inmoclinc

2. BOOSTER

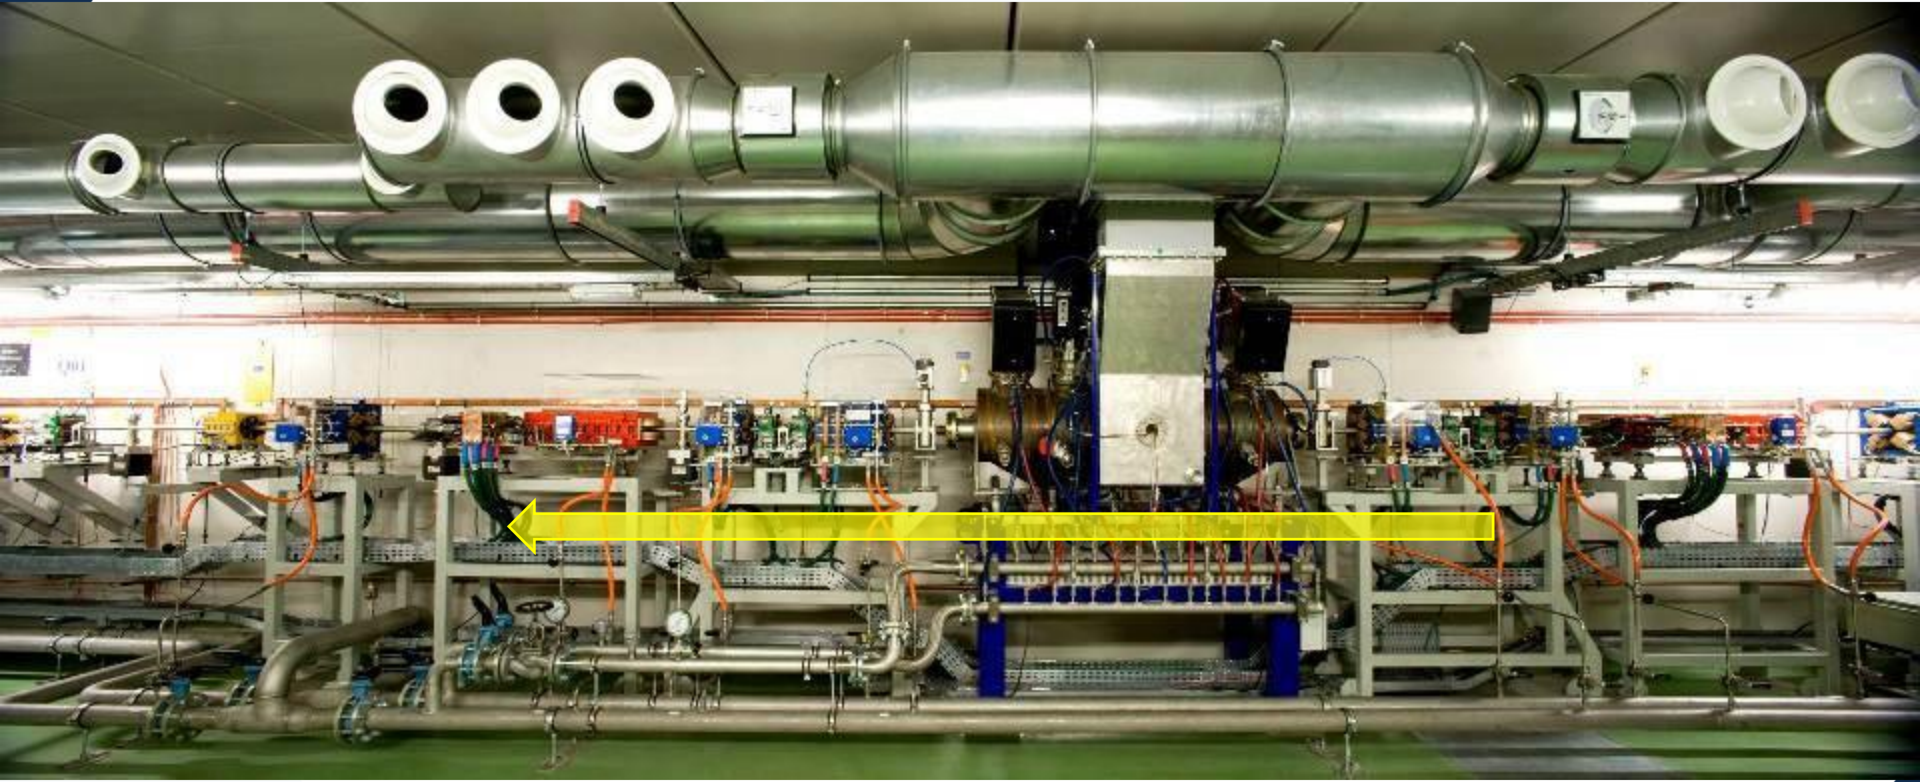

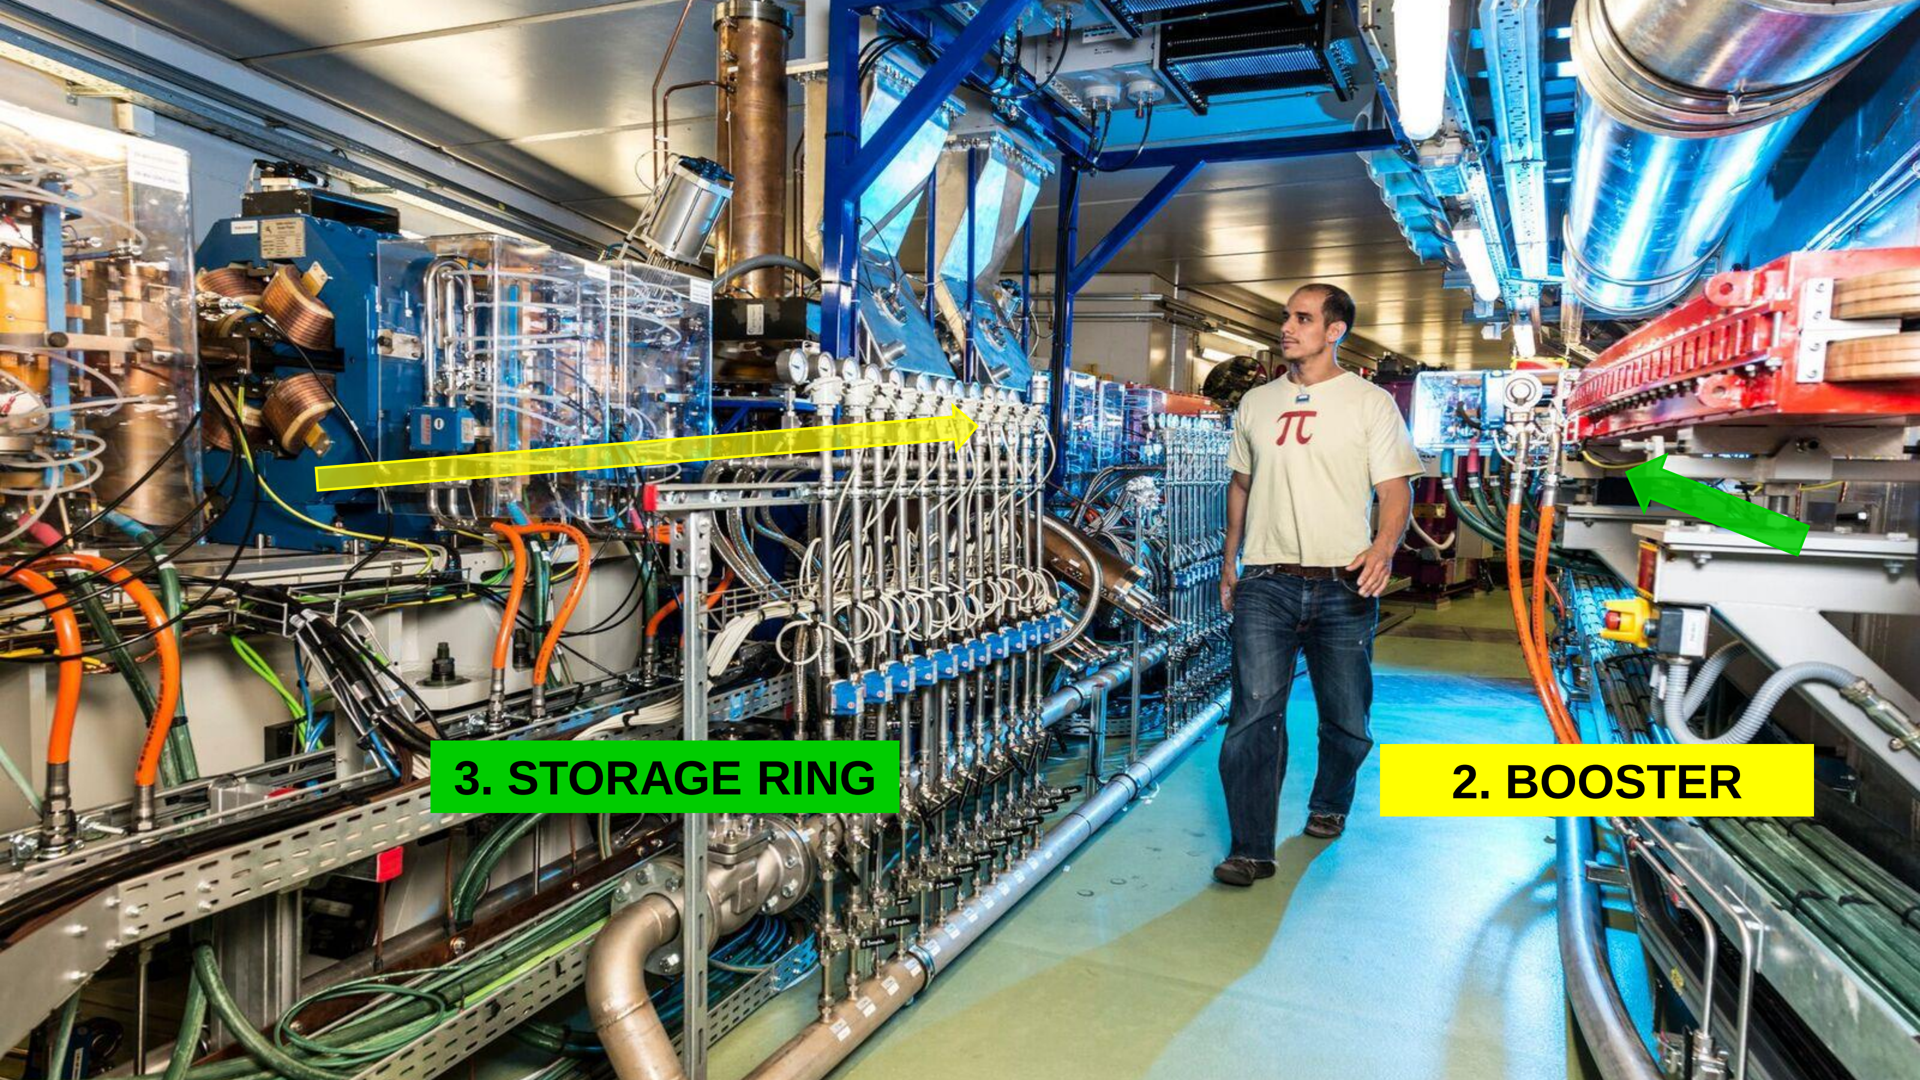
3. STORAGE RING
2. BOOSTER

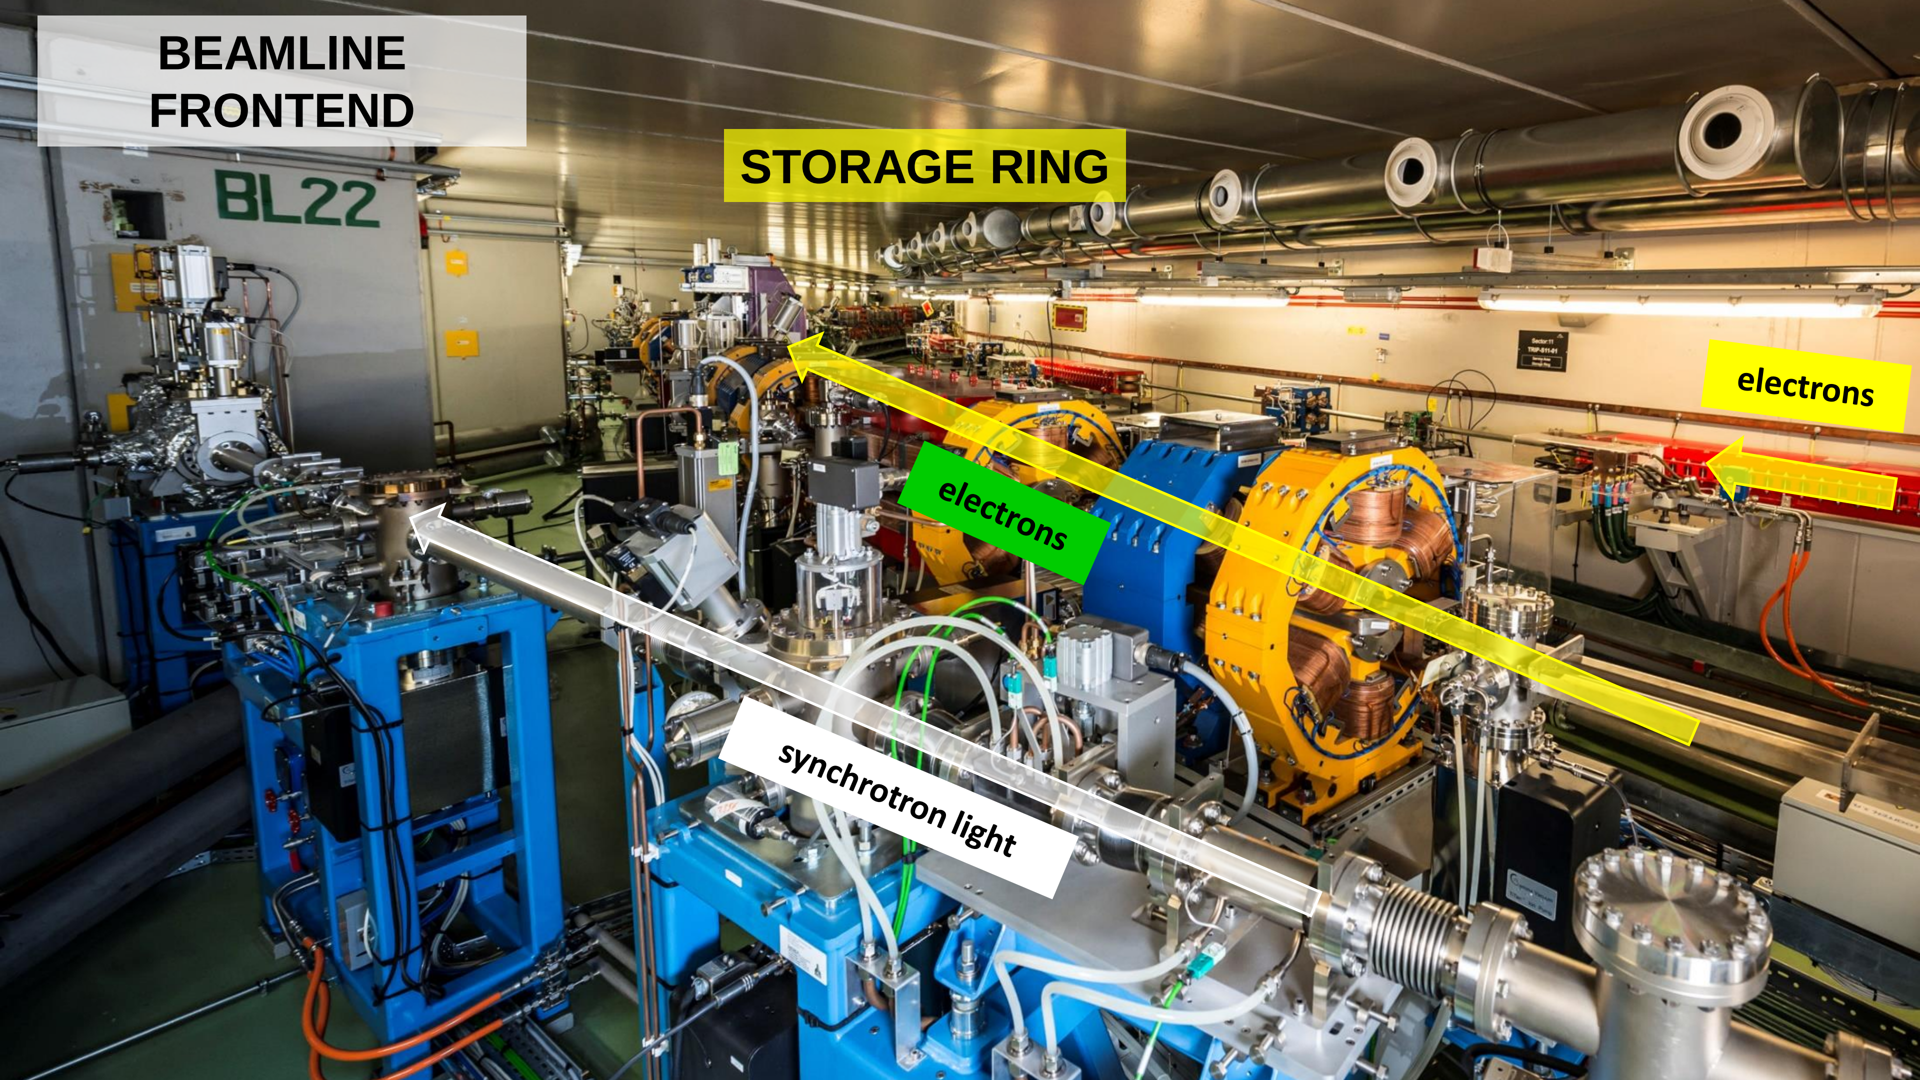
BEAMLINE
FRONTEND
STORAGE RING
BL22
electrons
electrons
synchrotron light

Basic elements of the accelerator

electricity

magnets

Radiofrequency cavities

Electromagnets

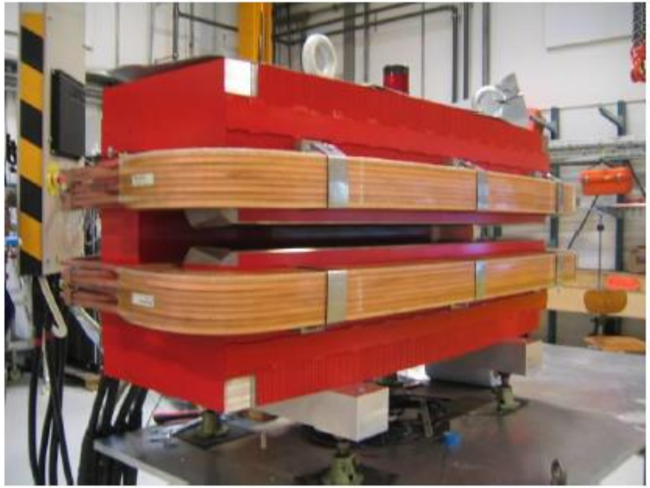

32 imants dipolars

112 quadrupols

120 sextupols

How do we change the direction of electrons?

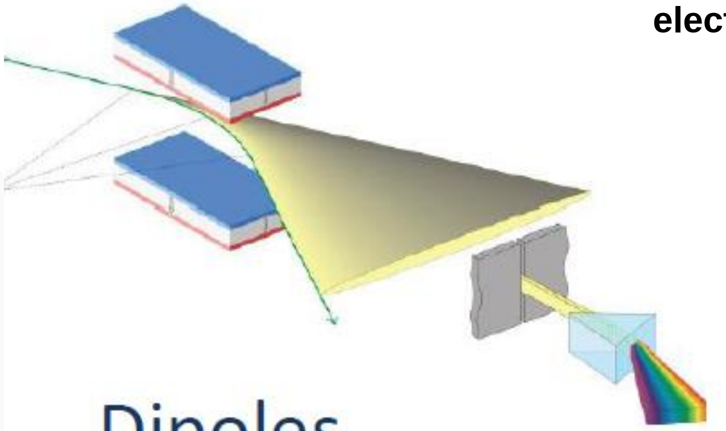

Dipoles

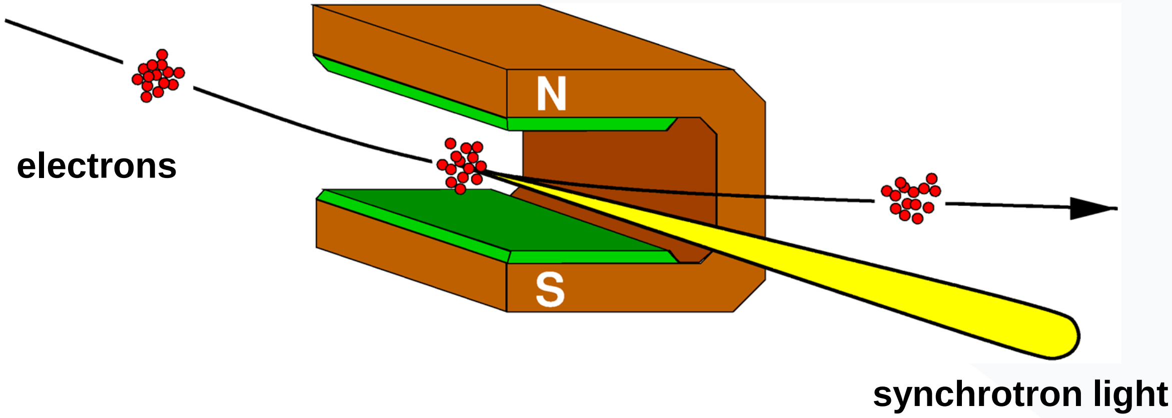

We use BIG MAGNETS

Insertion devices

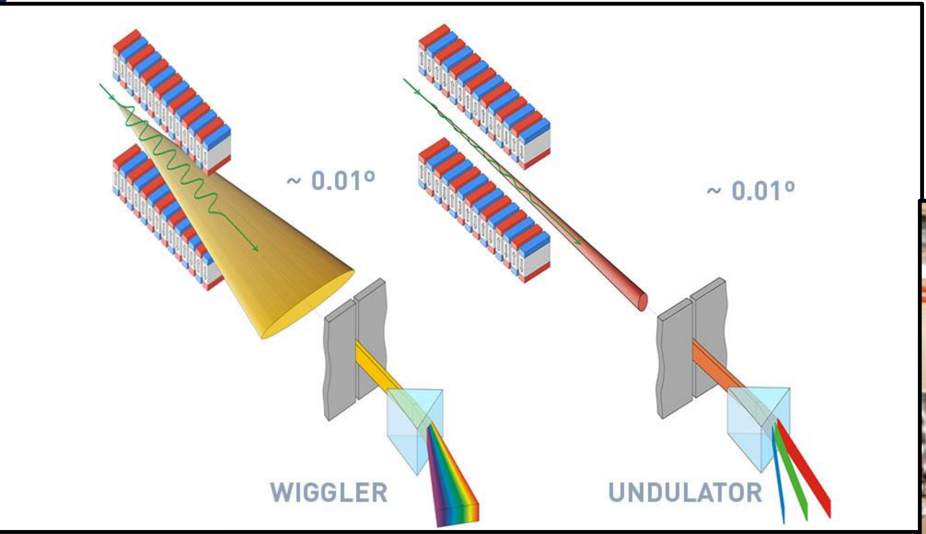

Characteristics of the synchrotron light

WIDE SPECTRUM from infrared to X rays	**BRILLIANT**
POLARIZED	**PULSED**

BEAMLINES

where the experiments happen

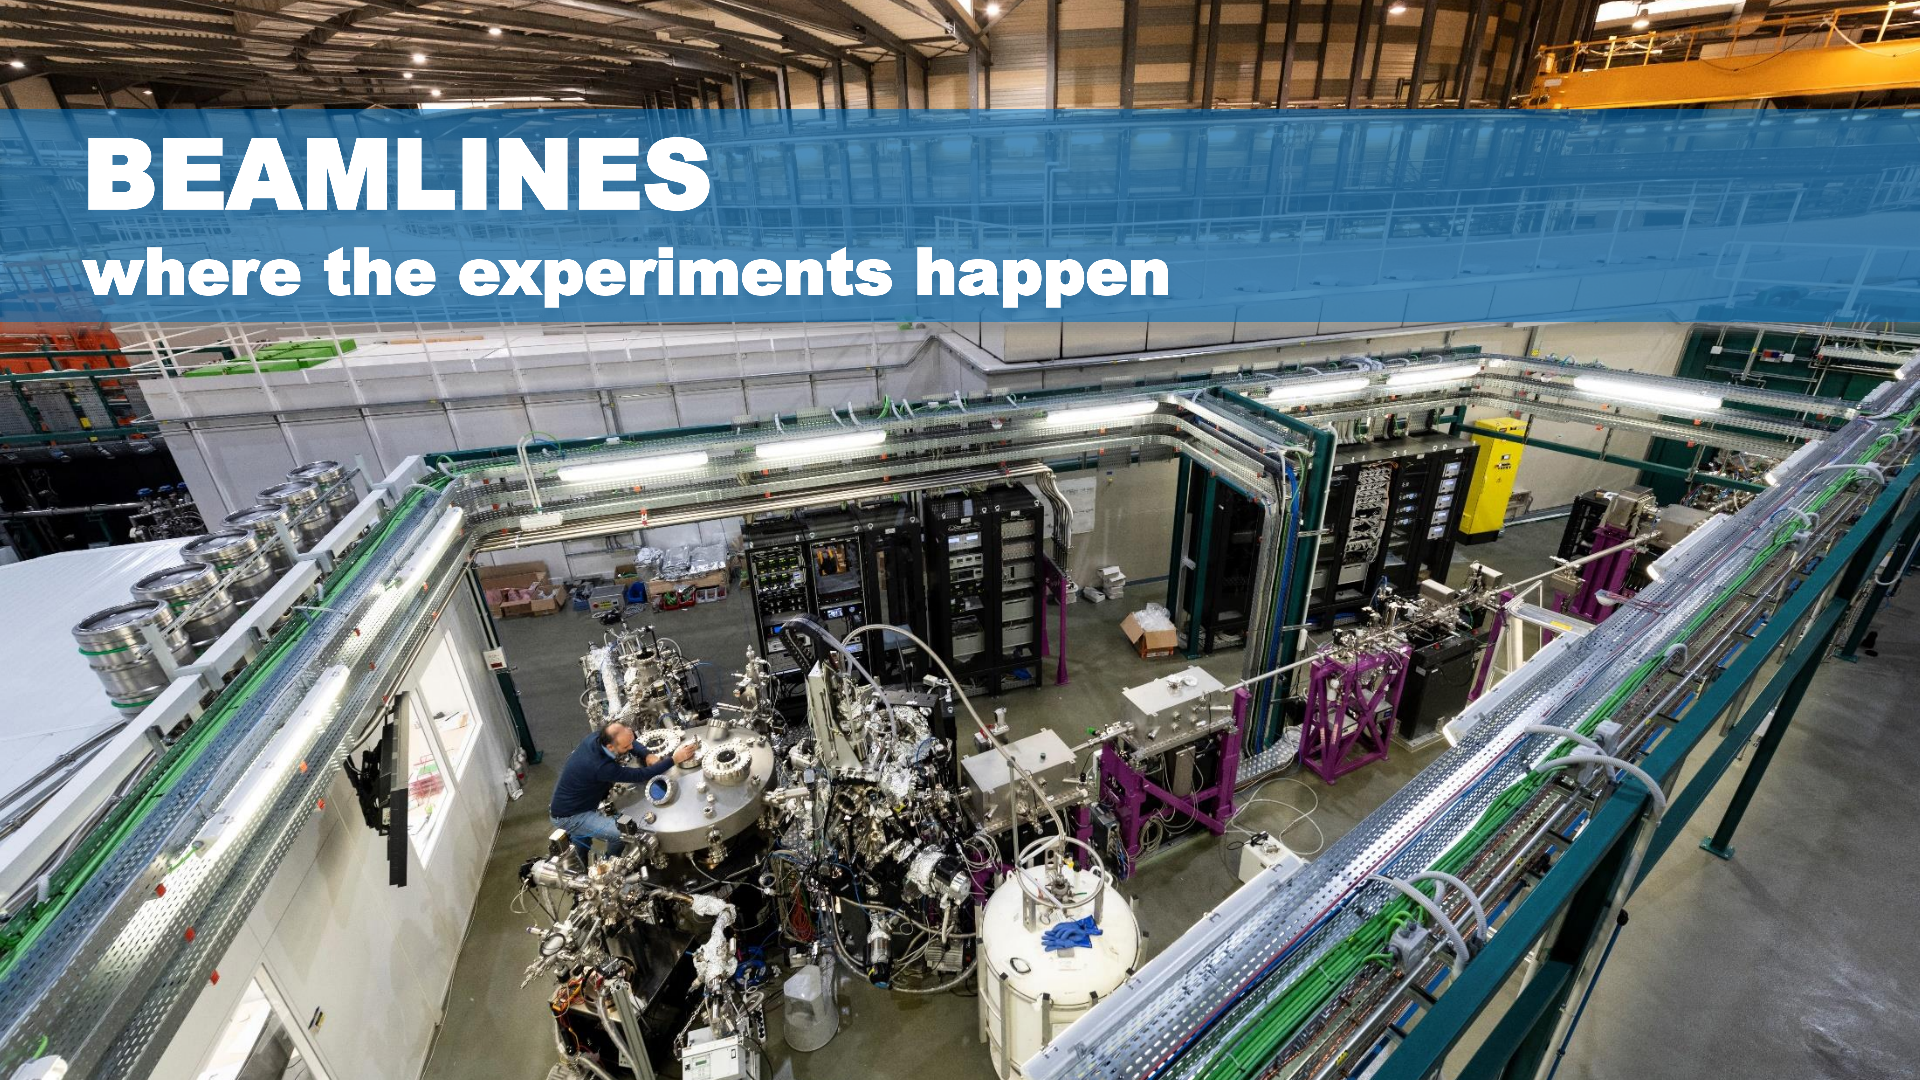

Parts of a beamline

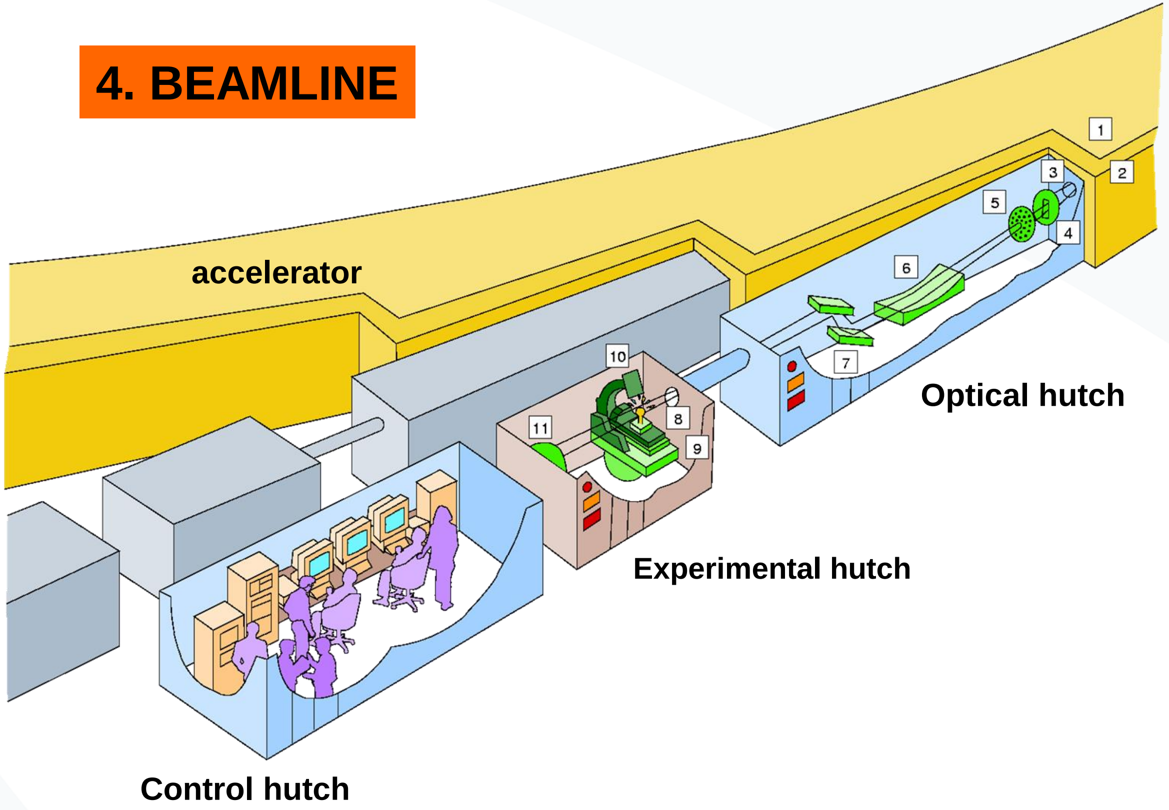

CLAESS & NOTOS

Chemistry, Batteries

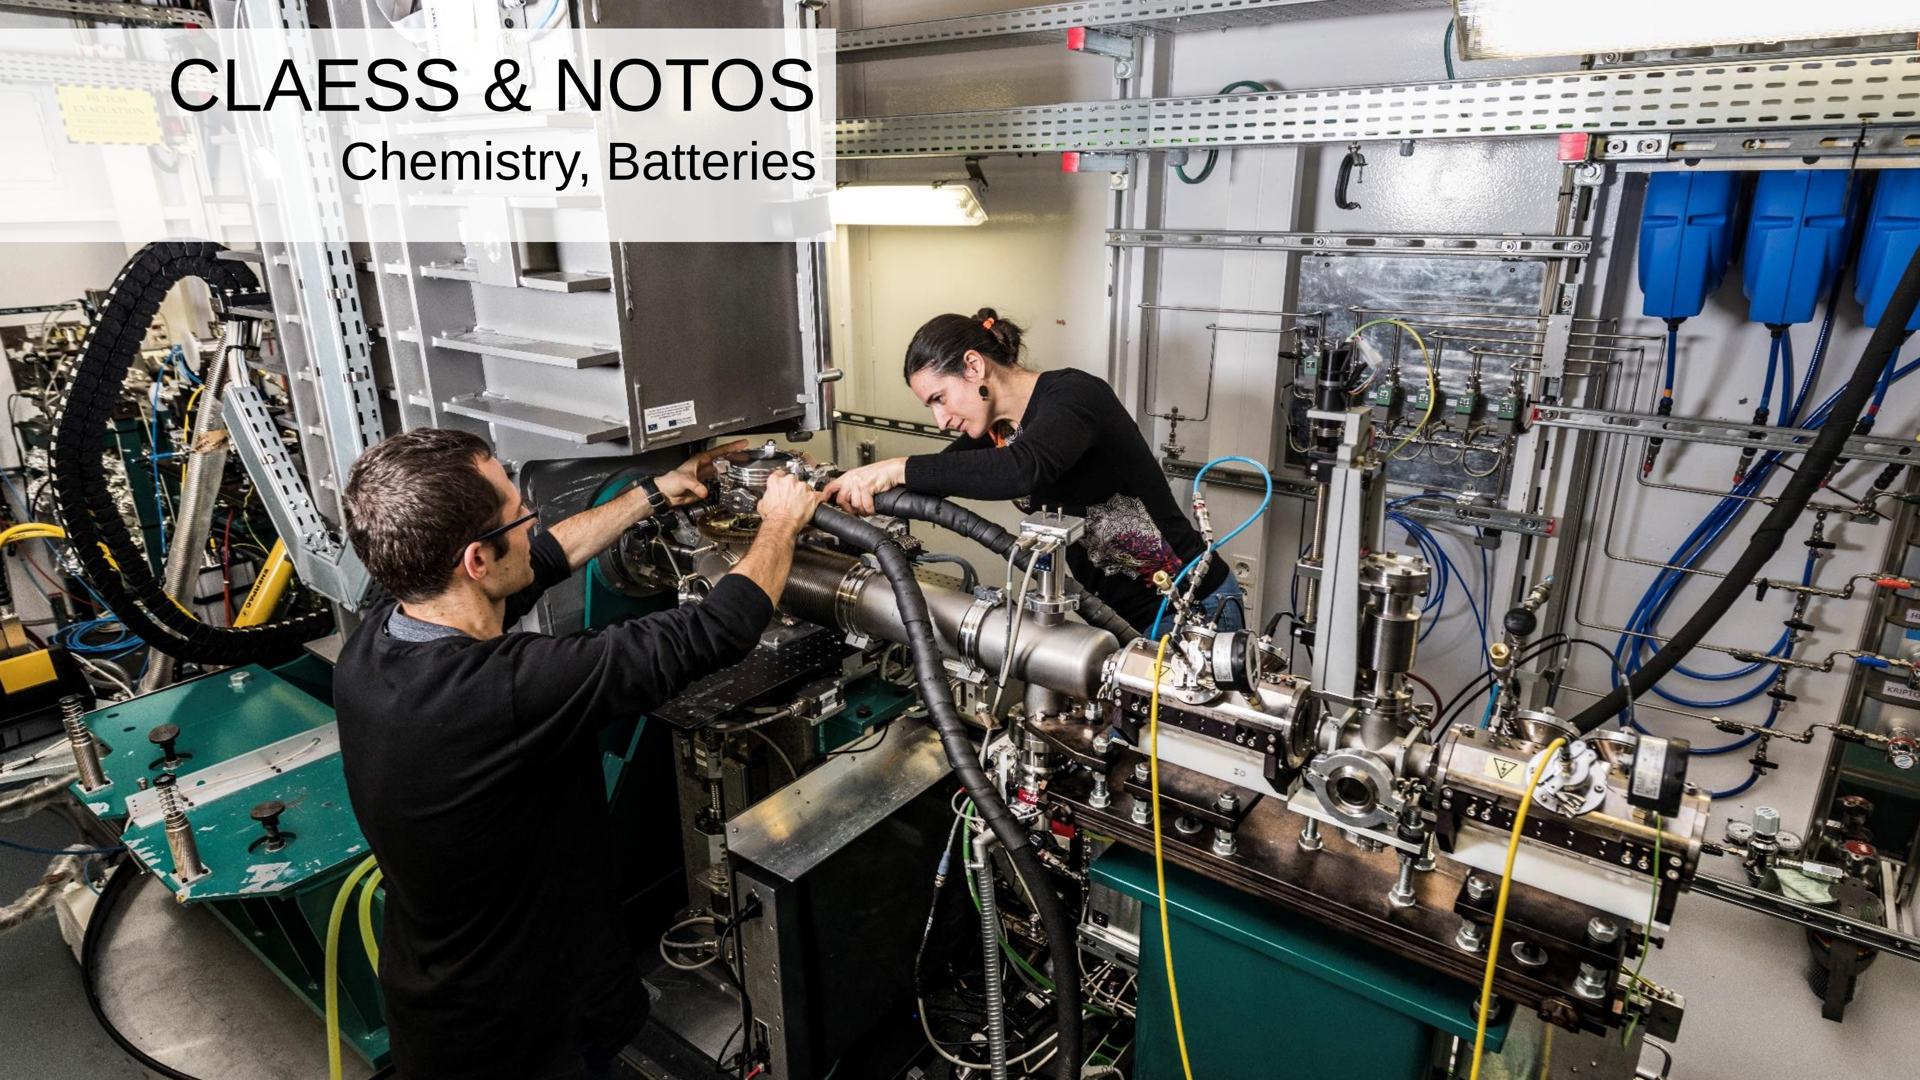

CARBON NANOSPHERES FOR IMPROVED SODIUM-SULFUR BATTERIES

Sodium-sulfur batteries are promising electrical energy storage technologies that can serve as a key solution to intermittency problems and can be integrated with renewable forms of energy generation. An international research team has reported the synthesis of micro-mesoporous carbon nanospheres with continuous pore distribution as an efficient sulfur host for sodium-sulfur batteries. The work sheds new light on the progress of the sulfur cathode in sodium-sulfur batteries and provides a promising strategy for the viable design of other metal–sulfur batteries. Experiments at the CLAESS beamline in ALBA allowed determining the sulfur species during charge/discharge processes.

Cerdanyola del Vallès, 25 May, 2022. Solar and wind power are useful resources for energy generation but they are intermittent (at night or on cloudy days solar panels do not work, for example). Electrical energy storage technologies serve as a key solution to these intermittency problems and can be integrated with **renewable forms of energy generation**. Among these technologies, **room-temperature sodium-sulfur (Na–S) batteries** are deemed to be one of the **most promising candidates**, owing to their **high theoretical energy density** – the amount of energy they can store - and **low cost**. Nonetheless, this battery system suffers from a slow reaction rate at room temperature, which radically limits battery performance and makes difficult its practical commercialization.

An efficient strategy to deal with this challenge is the **use of porous carbon material as a host to encapsulate molecular sulfur**, significantly enhancing its conductivity. This system acts as the cathode of the battery, which is the electrode where reduction occurs. To make the battery work, sodium ions have to migrate from the anode to the cathode. However, in these systems, it is a **challenge to provide fully accessible sodium ions that do not obstruct the sub-nanosized pores of the carbon host**.

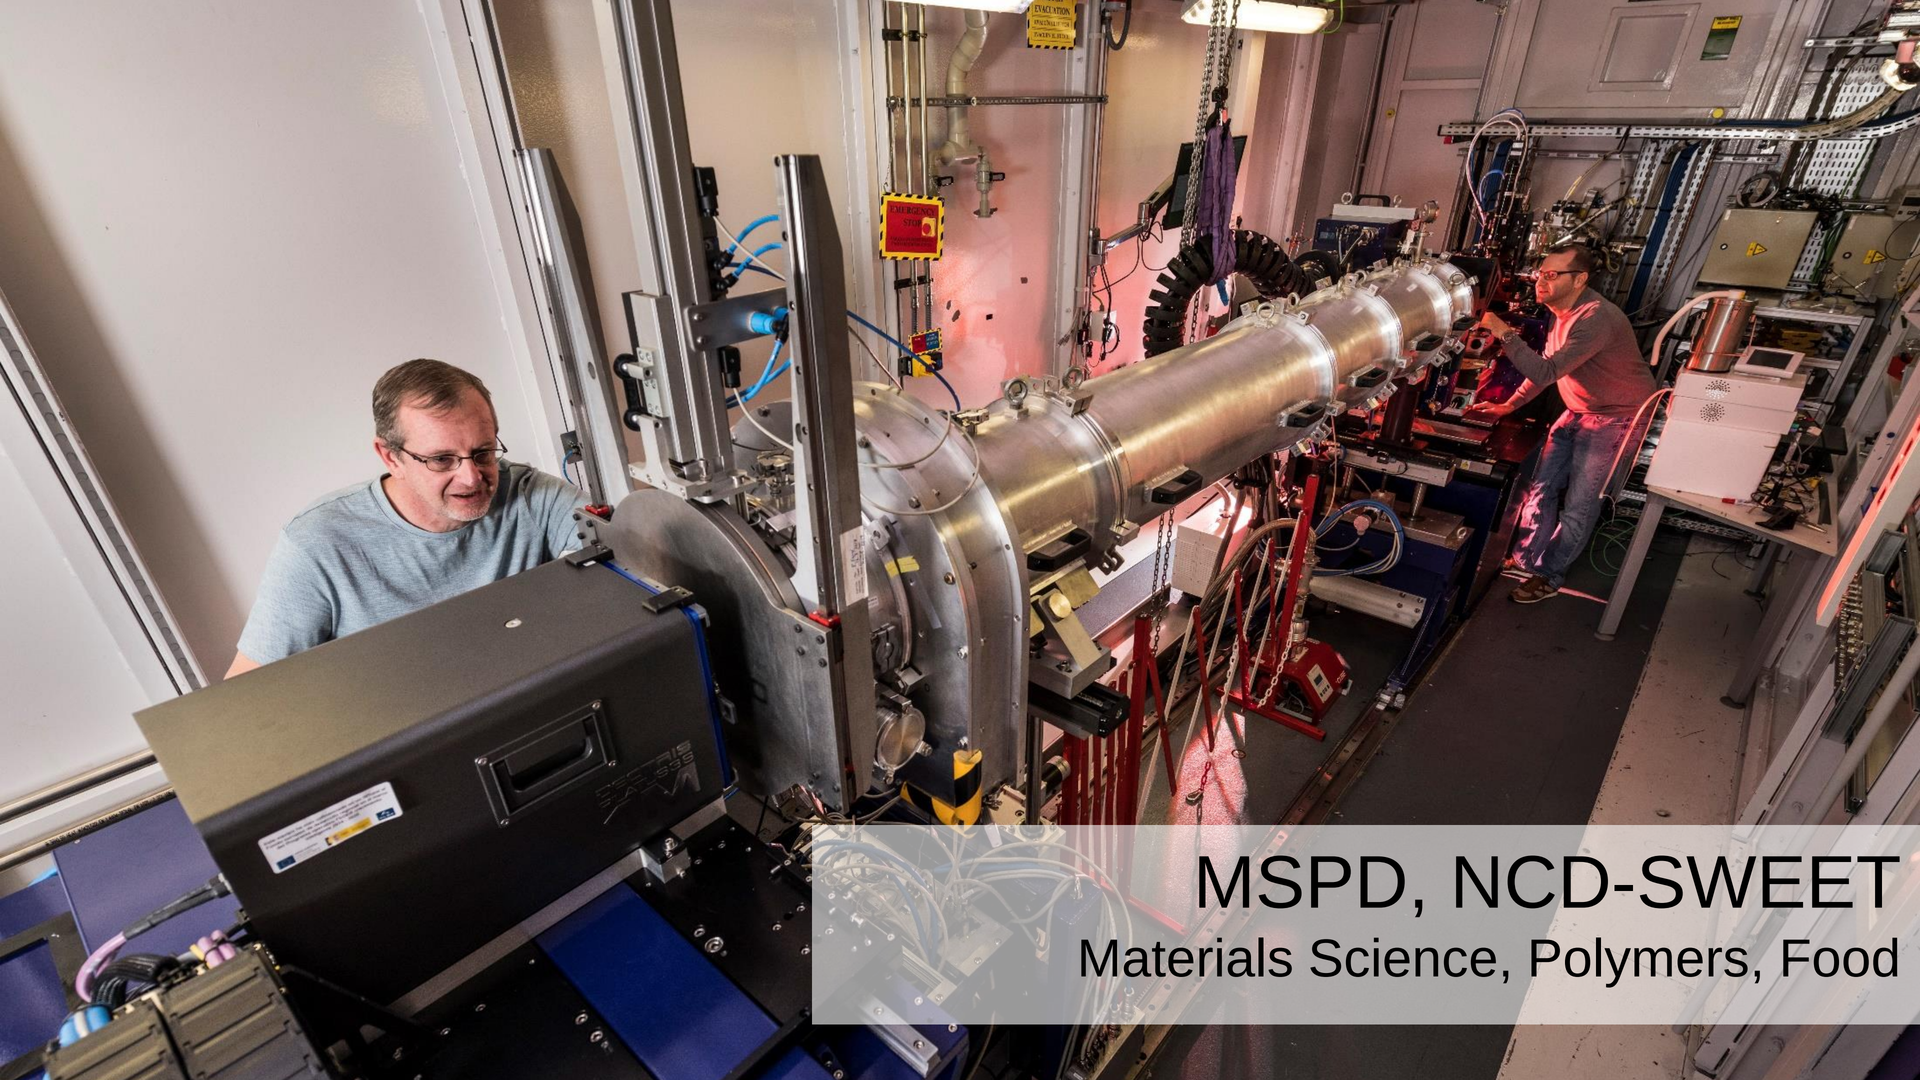

MSPD, NCD-SWEET

Materials Science, Polymers, Food

CYBORG PLANTS: ROOTS CAN STORE ENERGY

Researchers of the HyPhOE European Project have developed biohybrid plants with an electronic root system, which could be used to store energy or as electronic sensors. This study proved the integration of circuits and electrochemical devices into the plants without damaging them, so that they continued to grow and adapt to their new hybrid state. Experiments at the NCD-SWEET beamline of the ALBA Synchrotron were crucial to shed light on the plant-based technology field.

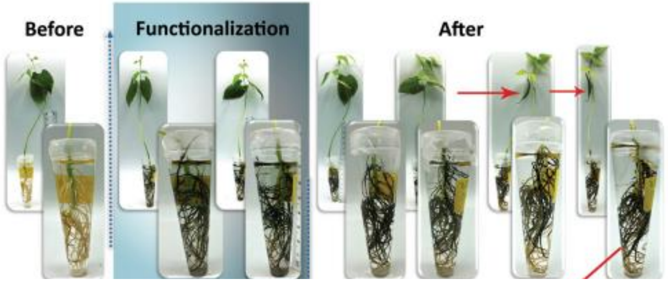

Cerdanyola del Vallès, 17th January 2021. **Plants are amazing machines**: not only they are solar-powered and convert carbon dioxide into chemical energy, but they are also capable of producing cellulose, the most abundant biopolymer on Earth, and can self-repair via tissue regeneration. All these factors make plants **the perfect candidates for developing biohybrid technological systems**, integrating smart materials and devices into their structure.

In a recent publication, the team led by researcher Eleni Stavrinidou from the Linköping University (Sweden) has presented a study about **biohybrid plants with an electronic root system**. They found out how to integrate circuits and electrochemical devices into the plants **without damaging them, so that they can continue to grow and develop**, and use them as supercapacitors or electronic sensors.

MISTRAL, XALOC, MIRAS, & XAIRA

Biology, Life Sciences

Structure of the proteins

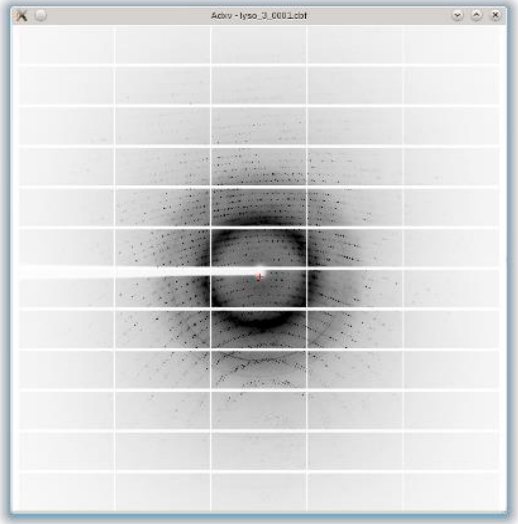

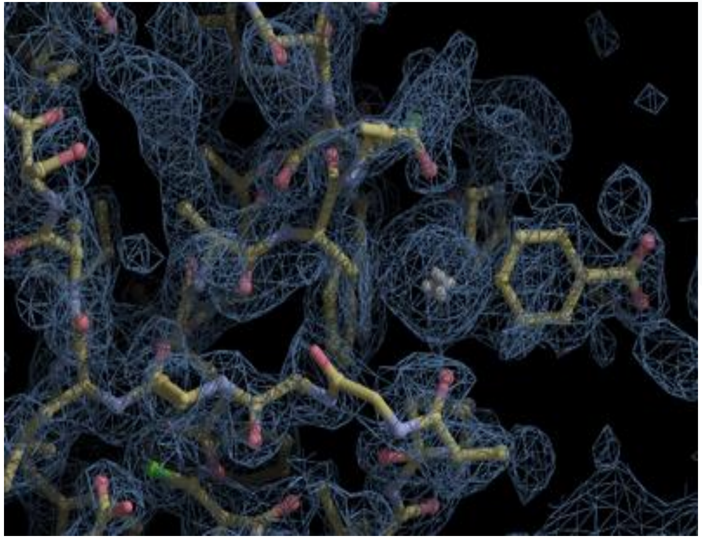

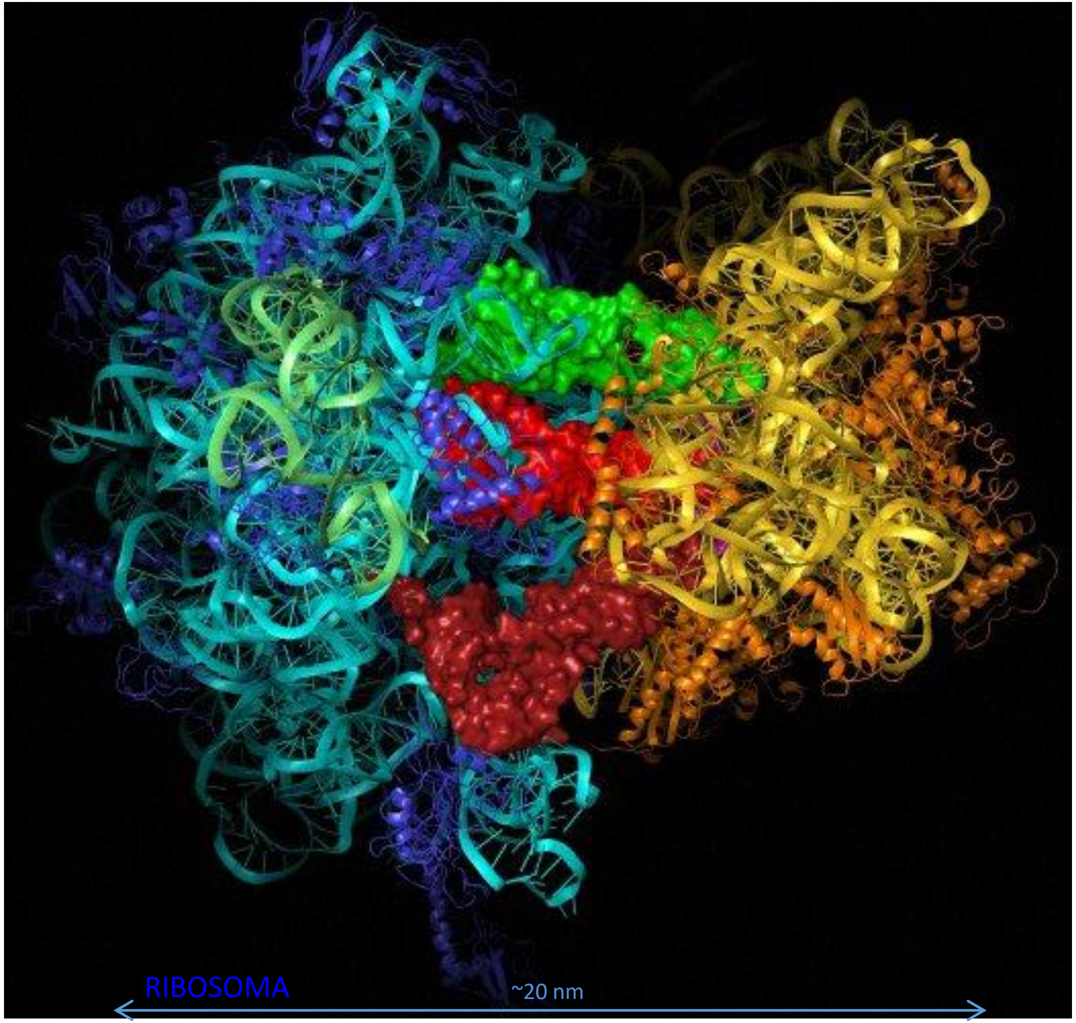
RIBOSOMA
~20 nm

SYNCHROTRON LIGHT UNVEILS THE EFFICACY OF A NEW NANOMATERIAL FOR CARDIAC FIBROSIS

ALBA

Experiments led by a PhD student at the ALBA Synchrotron revealed, for the first time, the location of a novel hybrid nanomaterial designed to inhibit the collagen overproduction after a myocardial fibrosis event and the induced morphological changes in the cells. Obtaining this high-resolution 3D information of cells is crucial for developing pre-clinical studies of novel therapeutic agents.

Cerdanyola del Vallès, 8th November, 2021. A collaboration between ALBA Synchrotron PhD student Johannes Groen and researcher Eva Pereiro and the groups of Aitziber L. Cortajarena (CIC BiomaGUNE, San Sebastián) and Ana V. Villar (IBBTEC, Santander) have visualised for the first time in 3D the exact location and how a nanomaterial drug behaves in whole cells. **This drug is a therapeutic protein-nanomaterial hybrid specifically designed to inhibit the collagen overproduction after a myocardial fibrosis event**.

Using a new approach, researchers combined **cryo-3D correlative fluorescent light and X-ray tomography** to locate, for the first time, the nanomaterial and to obtain high resolution information on the **cellular morphological changes after treatment**, proving the antifibrotic properties of the drug.

First, the intracellular location allows to understand **what possible routes are used to enter the cell**, secondly, **how the cell deals with a specific agent** and finally, **what are the effects induced by the treatment on the structural morphology of the cells**.

The results obtained in this study pave the way for the **introduction of nanomaterial-based drugs and nanomedicine into the clinic** as they reinforce the usefulness of imaging techniques for evaluating cellular structure after the application of specific treatments.

BOREAS, CIRCE, LOREA

Electronic and magnetic properties

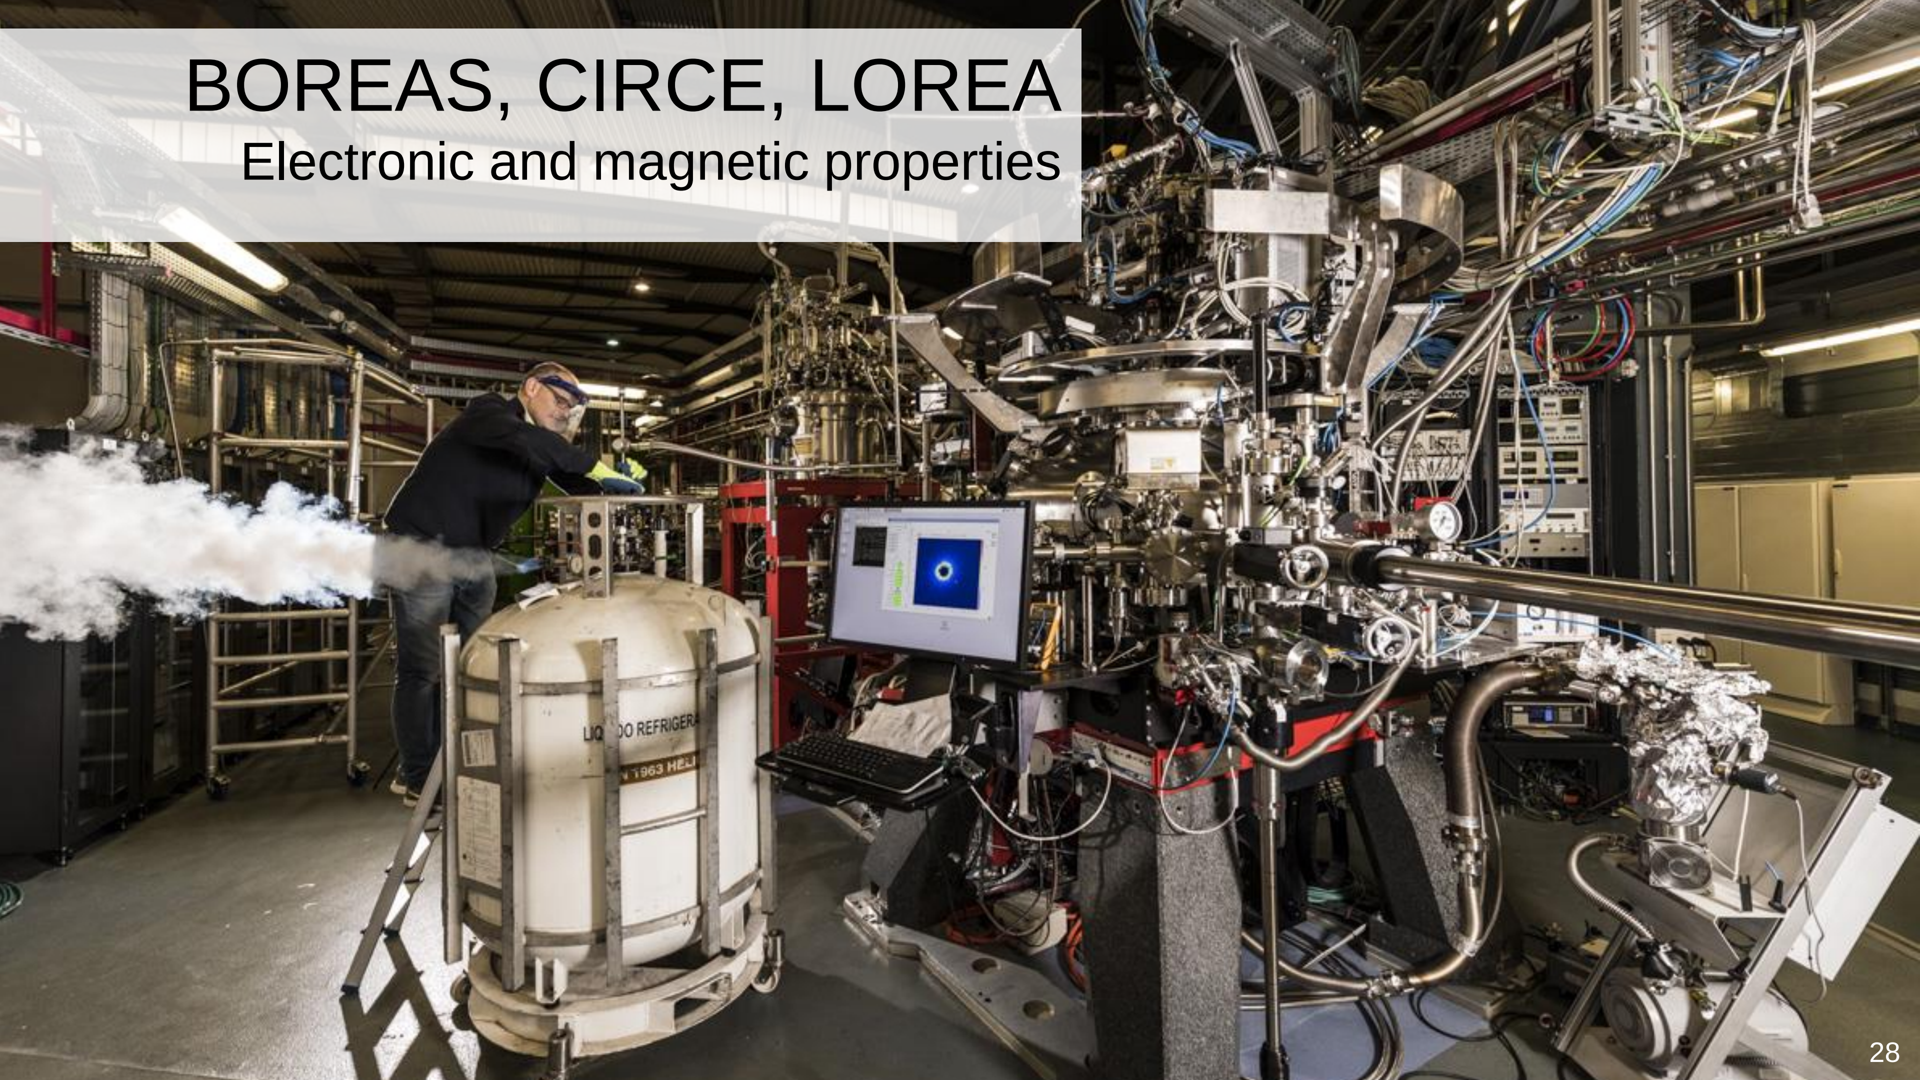

RESEARCHERS DISCOVER A NEW MECHANISM ENABLING ULTRAFAST MEMORY ARCHITECTURES

A study published in Advanced Functional Materials and led by researchers in the United Kingdom, in collaboration with BOREAS beamline at the ALBA Synchrotron, has proved the possibility to store magnetic information in picosecond timescales at a fullerene - oxide interface by using the photocurrent generated in the molecular layer. This holds great potential for the development of eco-friendly, ultra-fast hybrid information memories and magneto-optic sensors working via light or electrical irradiation at high frequencies.

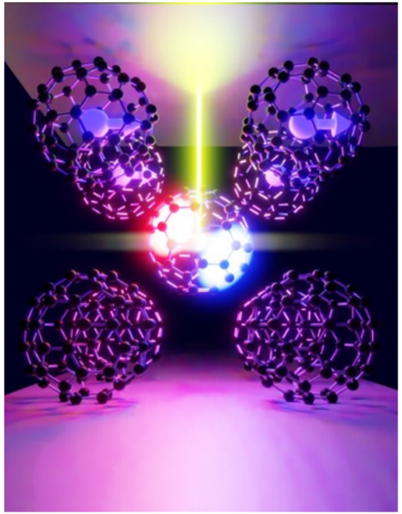

Cerdanyola del Vallès (Barcelona), 16th March 2023. The amount of information stored and circulated doubles approximately every 2 to 4 years in what is known as the zettabyte era (the period in mid 2010s when the amount of digital data in the world first exceeded a zettabyte -2012- and when the global IP traffic first exceeded one zettabyte -2016-).

The power and resources required to sustain the exchange and storage of information represents a continually increasing percentage of the global resources. To make this progress sustainable in the coming decades, **it is needed to find disruptive technologies** operating at higher frequencies than the ones we use today (from MHz to THz) whilst reducing the operational power and promoting the use of eco-friendly materials. However, current devices frequently need of heavy metals and face limitations due to the speed of magnetization reversal and the current densities needed to store or switch information.

Researchers from the **University of Leeds**, in collaboration with the **Science and Technology Facilities Council**, the universities of **Edinburgh** and **Exeter**, and the **ALBA Synchrotron**, have published in the journal *Advanced Functional Materials* a **new device architecture** where the information is written using **light instead of traditional methods** such as magnetic fields in hard disks or electrical currents in solid state drives. That is achieved using a **fullerene layer**, a form of carbon such as graphene or graphite, sandwiched between a ferromagnetic metal and a transition metal oxide. In that material, light irradiation generates bounded pairs of holes and electrons (i.e excitons) that can be then separated using dielectric fields at each interface. These electrons and holes are finally spin-filtered, which generates **a magnetic layer where the information is stored.**

The JEMCA / InCAEM

- Currently there are two big electron microscopes in operation (owners ICN2 & IBMB).
- A third one will be acquired in InCAEM project (owner ALBA).
  - That involves that we **must give support to Data Management, Data Analysis of this Electron Microscope**.
  - The project includes challenging correlative analysis with different beamlines and cooperation with the PIC.

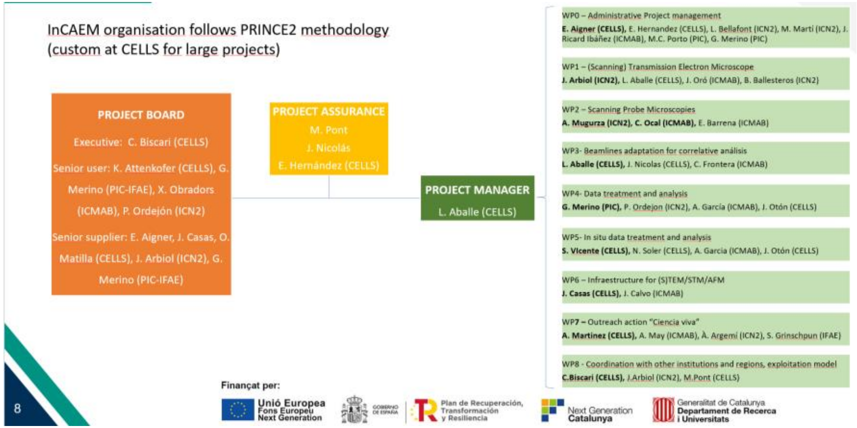

Around 260 people of permanent staff

ORGANIZATION/DIVISIONS

- ACCELERATORS
- ADMINISTRATION
- COMPUTING AND CONTROLS
- EXPERIMENTS
- ENGINEERING
- DIRECTOR'S OFFICE









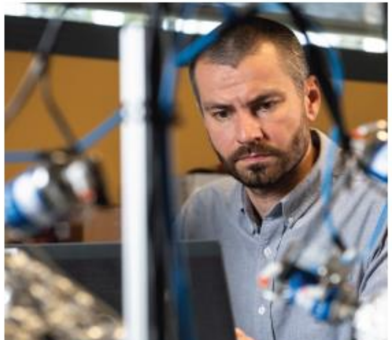


IT infrastructure for science

- Description of the infrastructure on-premises
- Ongoing upgrades

Storage

- 1 storage cabin for scientific data (1PB)
- 1 storage cabin for administrative data (190TB)
- 1 tape library with LTO5-7 drives (3.3PB)

Virtualization

- 29 virtualization servers
- 394 virtual machines
- 1 AFF NVME storage cabin (22TB)
- 1 AFF SSD storage cabin for SA VMs (17TB)
- 1 storage cabin for online backup (130TB)
- 10 to 25 GbE network ports
- Independent ethernet switches for the storage access (VMSAN)

Containers

- 1 Kubernetes cluster based on docker
- 3 master nodes
- 6 worker nodes (256 cores – 625 GB RAM)
- 76 persistent volumes (4 TB)
- External NL-SAS cabin for storage

High Performance Computing

- CPU cores: 524
- Total RAM: 4224 GB
- Connectivity: 4 x 10GbE ports
- GPUs:

Model	Units	CUDA cores	Total CUDA cores	Current use
Nvidia Tesla T4	2	2560	5120	General use HPC
Nvidia Tesla A100	4	6912	27648	Xaira
Nvidia Tesla P100	1	1126	1126	HPC interactive (Mistral)
Nvidia Tesla V100	2	5120	10240	1 x Amira (Mistral) + 1 x Dragonfly (Faxtor)
Nvidia Tesla L4	2	795	1590	Opera (Accelerators)
Nvidia Tesla A2	1	1280	1280	em01 (Glacios)
Nvidia Tesla A16	1	5120	5120	em01 (Glacios)
	13		**52124**	

Network

- 40 x 10GbE ports in High Availability
- 48 x 1GbE ports in High Availability
- Based in Extreme Networks and Nvidia-Mellanox technology

IT Infrastructure upgrade

- Future prospects RAW Data generated at ALBA

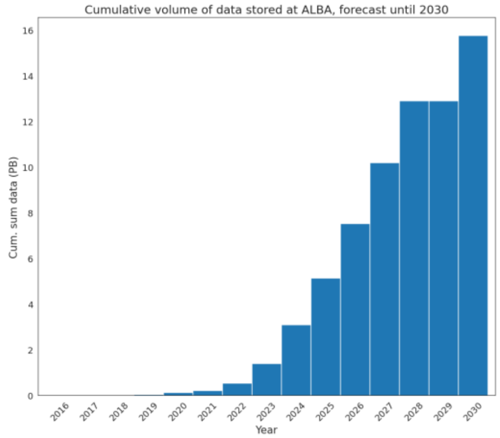

(TB)	BL06-XAIRA	BL13-XALOC	BL31-FAXTOR	BLXX-3Sbar
2022	0	200	0	0
2023	200	300	200	0
2024	500	300	700	0
2025	500	300	1000	0
2026	500	300	1000	0.3
2027	500	300	1000	0.5
2028	500	300	1000	0.5
2029	0	0	0	0
2030	500	300	1000	0.5

- Two years ago, we were confronted with a dangerous gradual ageing of our data center. Mainly acquired in the early years of ALBA.

- Fortunately, we have reversed the situation, and now there is an IT infrastructure investment plan that allowed us to invest 2.7 million euros in 2022.

IT Infrastructure upgrade

- Data Storage will be increased **from 1PB to 4.7 PB** with a drastically **improved performance**

Current storage for BLs and EMs
- Capacity: 1000 TB
- Write performance: 1.8 GBps

Current storage for Xaira commissioning
- Capacity: 15 TB
- Write performance: 6 Gbps

New Ultra Fast Storage
- Capacity: 3000 TB at >1 GBps
- Burst Buffer Xaira: > 35 TB at > 2.6 GBps
- Burst Buffer Faxtor: > 70 TB at > 12.5 GBps
- Currently under installation
- 900 k€ investment

New storage for EMs
- Capacity: >=1000 TB
- Write performance: 2.6 GBps
- Currently under installation
- 419k€ investment

IT Infrastructure upgrade

- Backup and archiving capacity will be updated, in a first phase, without increasing capacity

Current backup & archive tape library

- Maximum write transfer: 3930 MB/s
- Capacity*: 3360 TB
- Tape technology: LTP5 & LTO7

New tape library

- Maximum write transfer: 7200 MB/s
- Capacity*: 3600 TB
- Tape technology: LTO9 (+ 1 LTO7 drive for backwards compatibility)
- Currently under purchase process
- 50 k€ investment

* Raw maximum capacity without extracting tapes offline

IT Infrastructure upgrade

- The **CPU computation** capacity will be **multiplied x3** - Tender in awarding process
- **GPU computation** upgrade will be defined during this year.

Current HPC cluster

- CPU cores: 524
- Total RAM: 4224 GB
- Connectivity: 2 x 10GbE ports
- 13 GPUs of different characteristics

HPC cluster upgrade (additional to the current)

- CPU cores: + 1080
- Total RAM: + 5760 Gb
- Connectivity: 4 x 25GbE ports
- 10 nodes reserved to Xaira online analysis
- To be installed at the end of 2023
- 299k€ investment

UNIÓN EUROPEA
Fondo Europeo de Desarrollo Regional
"Una manera de hacer Europa"

GPU upgrade

- GPUs computation requirements for Faxtor and ML/AI are under definition
- GPU hardware will be acquired this year

IT Infrastructure upgrade

- **Network core** capacity will be increased in a **x10** factor

Current core switches

- 40 x 10GbE ports in High Availability
- 48 x 1GbE ports in High Availability
- Obsolete hardware of 2009

New core switches

- 48 x 100 GbE ports in High Availability
- 96 x 25 GbE ports in High Availability
- Required for the connectivity for new high demanding environments
  - Ultra Fast Storage
  - HPC cluster upgrade
- Purchase done. Waiting the hardware delivery
- 219k€ investment

ALBA
GOBIERNO DE ESPAÑA MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES
UNIÓN EUROPEA Fondo Europeo de Desarrollo Regional "Una manera de hacer Europa"

Future Challenges: Open Science

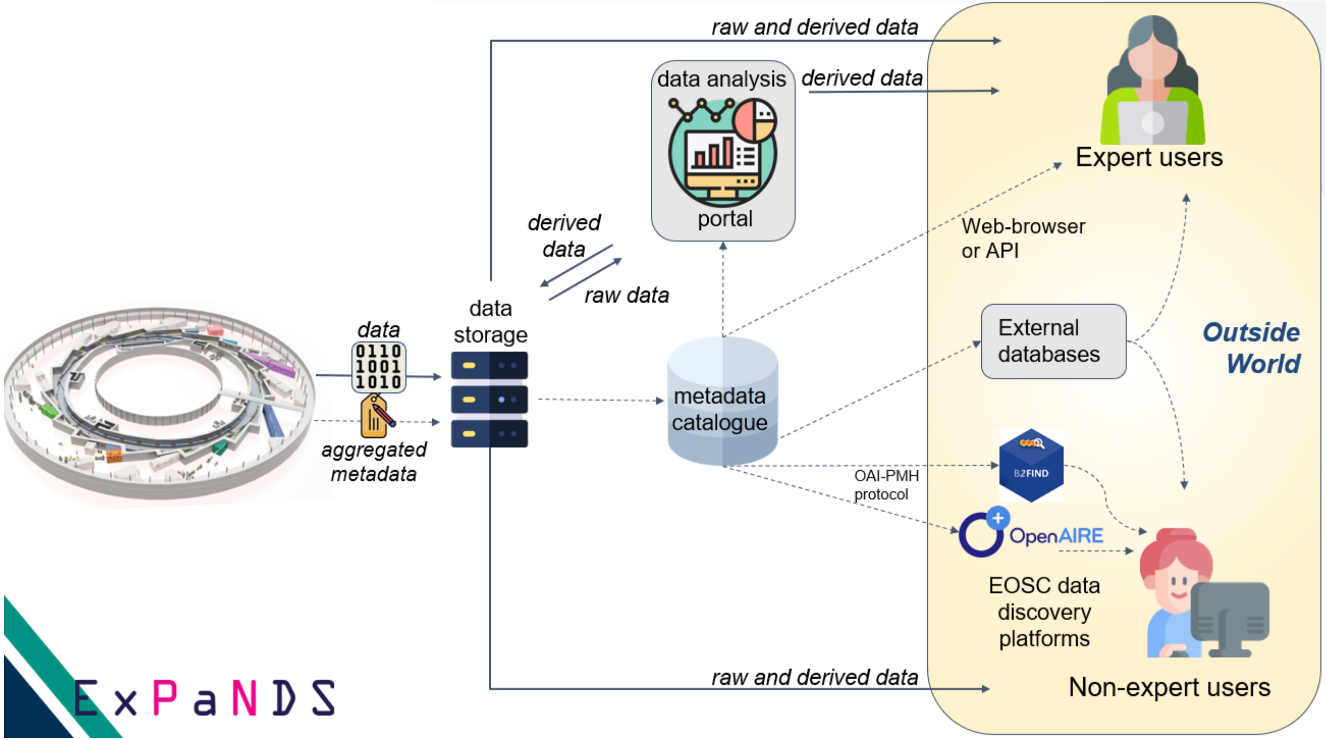

ALBA II, project for the update of the infrastructure

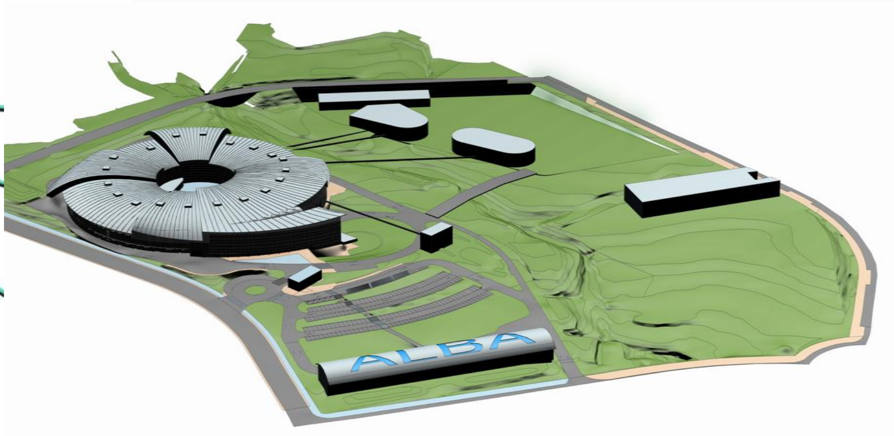

a)

b)

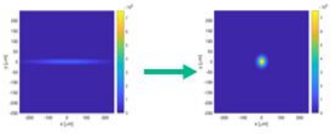

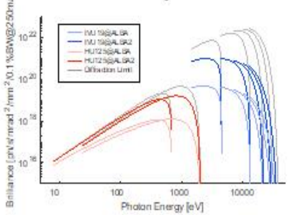

The use of the CLOUD at ALBA

- The use of OCRE IaaS Framework - AWS
- The OCRE grant for science - GCP
- Particular use cases

OCRE Open Clouds for Research Environments

Google Cloud

aws

Contracting

Public institution = “hard” procurement mechanisms

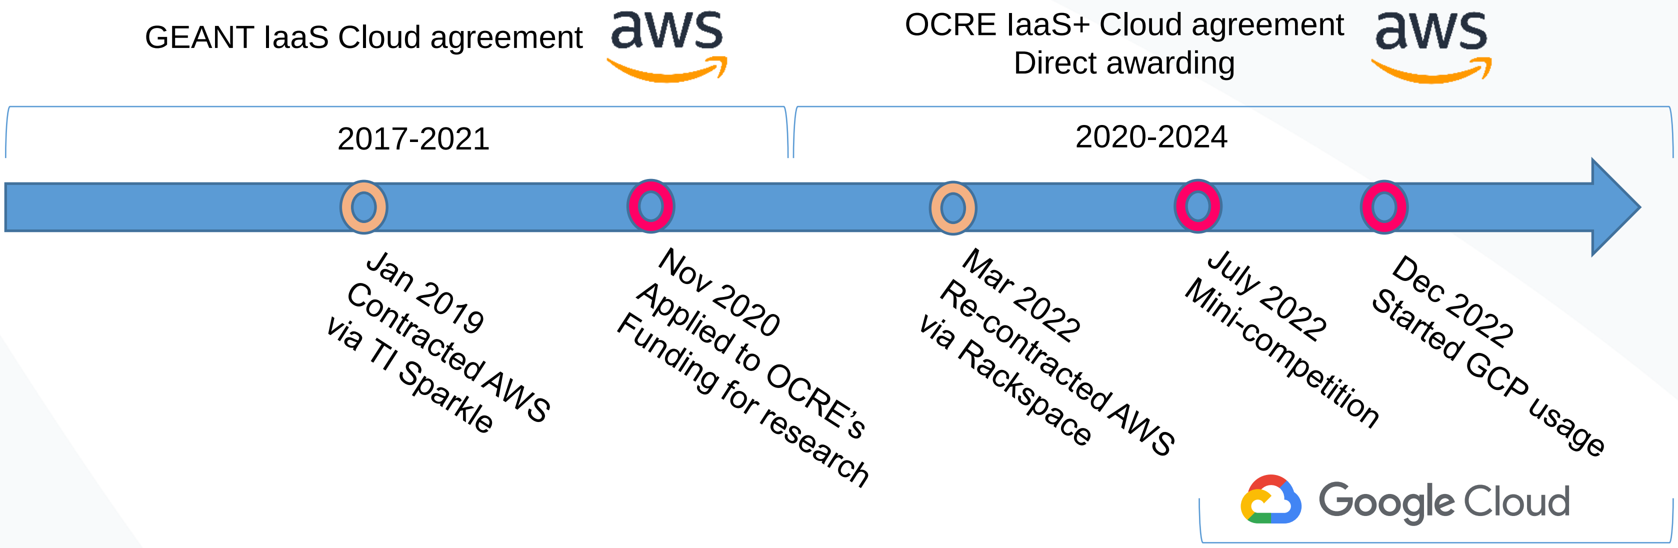

AWS – Archiving of scientific data

- 660 TBs (286.360.275 files) stored in AWS S3 (Glacier Deep Archive tier)

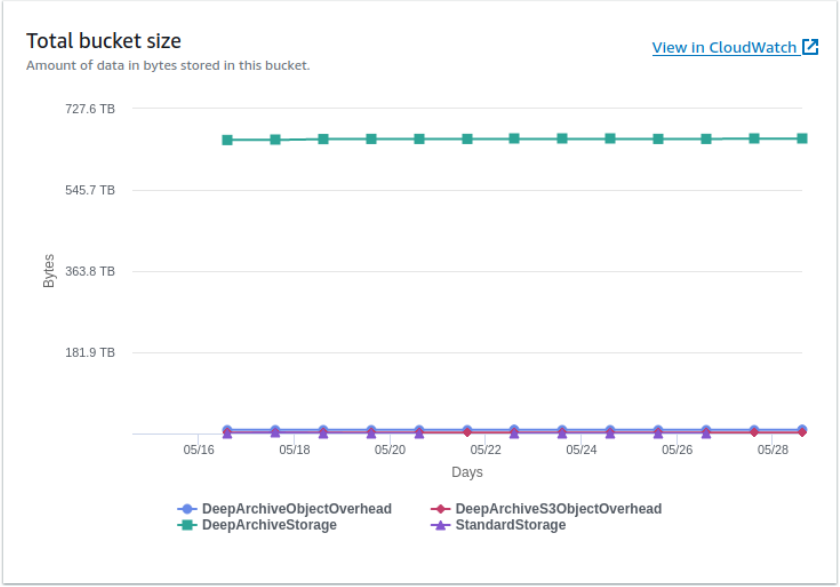

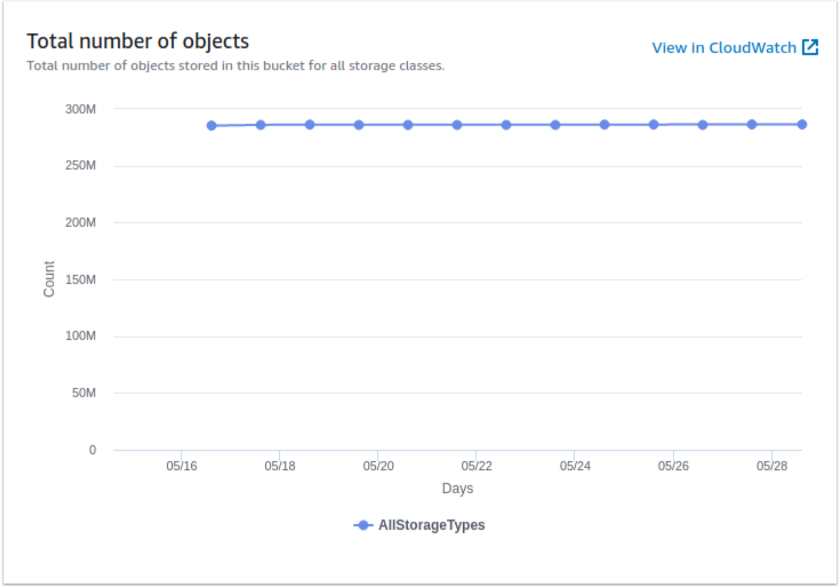

OCRE funding - Use cases

- Virtual machines
    - Prepare reusable base images (with needed drivers and software installed)
    - Upload our own datasets (25GB – 400GB) to persistent disks attached to the VMs
    - Run the same data pipelines in different types of VMs to find the best specs
    - Compare cloud processing times with on premise processing times
    - Integrate with our on premise data analysis workflows to start cloud instances on-demand* (need to have data uploaded, data transfer is slow)
- HPC cloud bursting
    - Extend the on-premise high performance computing capacities to the cloud to satisfy all the needs in moments of high workloads demands
    - Offload Slurm jobs to a cloud HPC alternative cluster
- Kubernetes
    - Deployment of a cloud Kubernetes cluster that can be managed from our on premises Rancher

GCP architecture

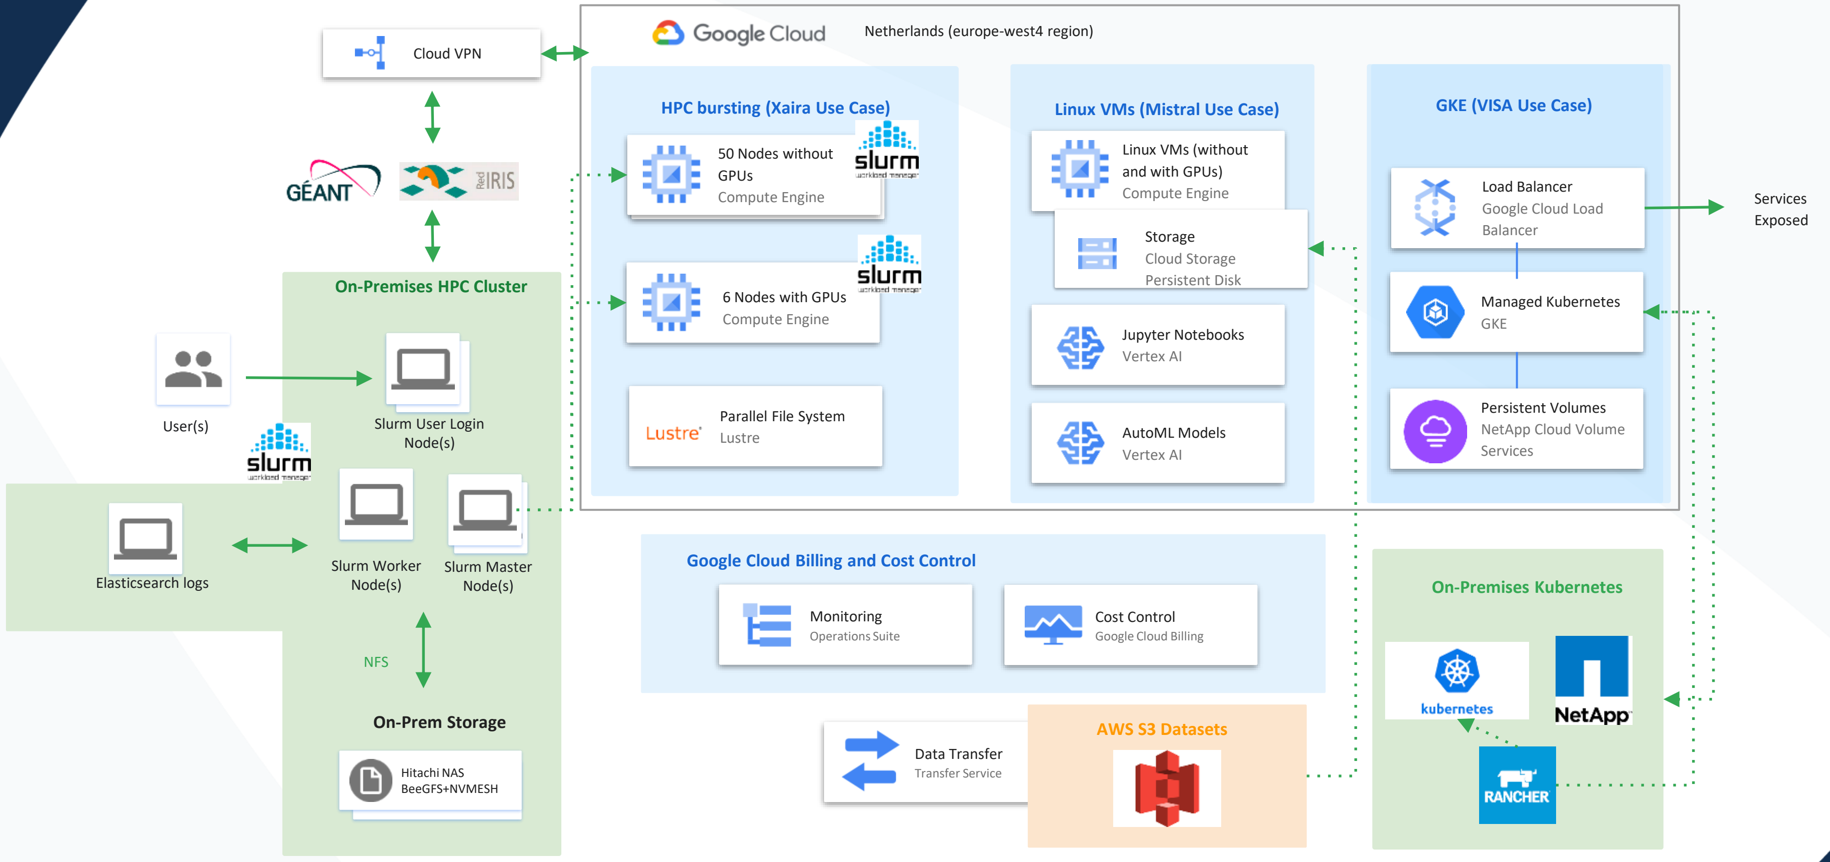

Image processing with Dragonfly

- FAXTOR is currently under construction and commissioning. The first user operation is expected by the beginning of 2024
- The possible fields of applications are multifold and include material sciences, paleontology, earth science, cultural heritage and industrial application
- Techniques: Absorption and phase contrast X-ray imaging (e.g., propagation-based and grating-based imaging), 4D X-Ray Computed Tomography
- Dragonfly is used to process and analyze the reconstructed volumes obtained from computed tomography scan. It can be used to study static and dynamic samples, following their evolution with time

Image processing with Dragonfly

- Service: Compute Engine
- Instance type: n1-highmem-96 (96 vCPUs, 624 GB memory)
- GPUs: 4 x Nvidia Tesla P100 (with Virtual Workstation, ex Nvidia GRID)
- OS: Windows Server 2022
- Scientists access: IAP Desktop (Identity Aware Proxy, using their own Google credentials) / RDP

Image processing with Dragonfly

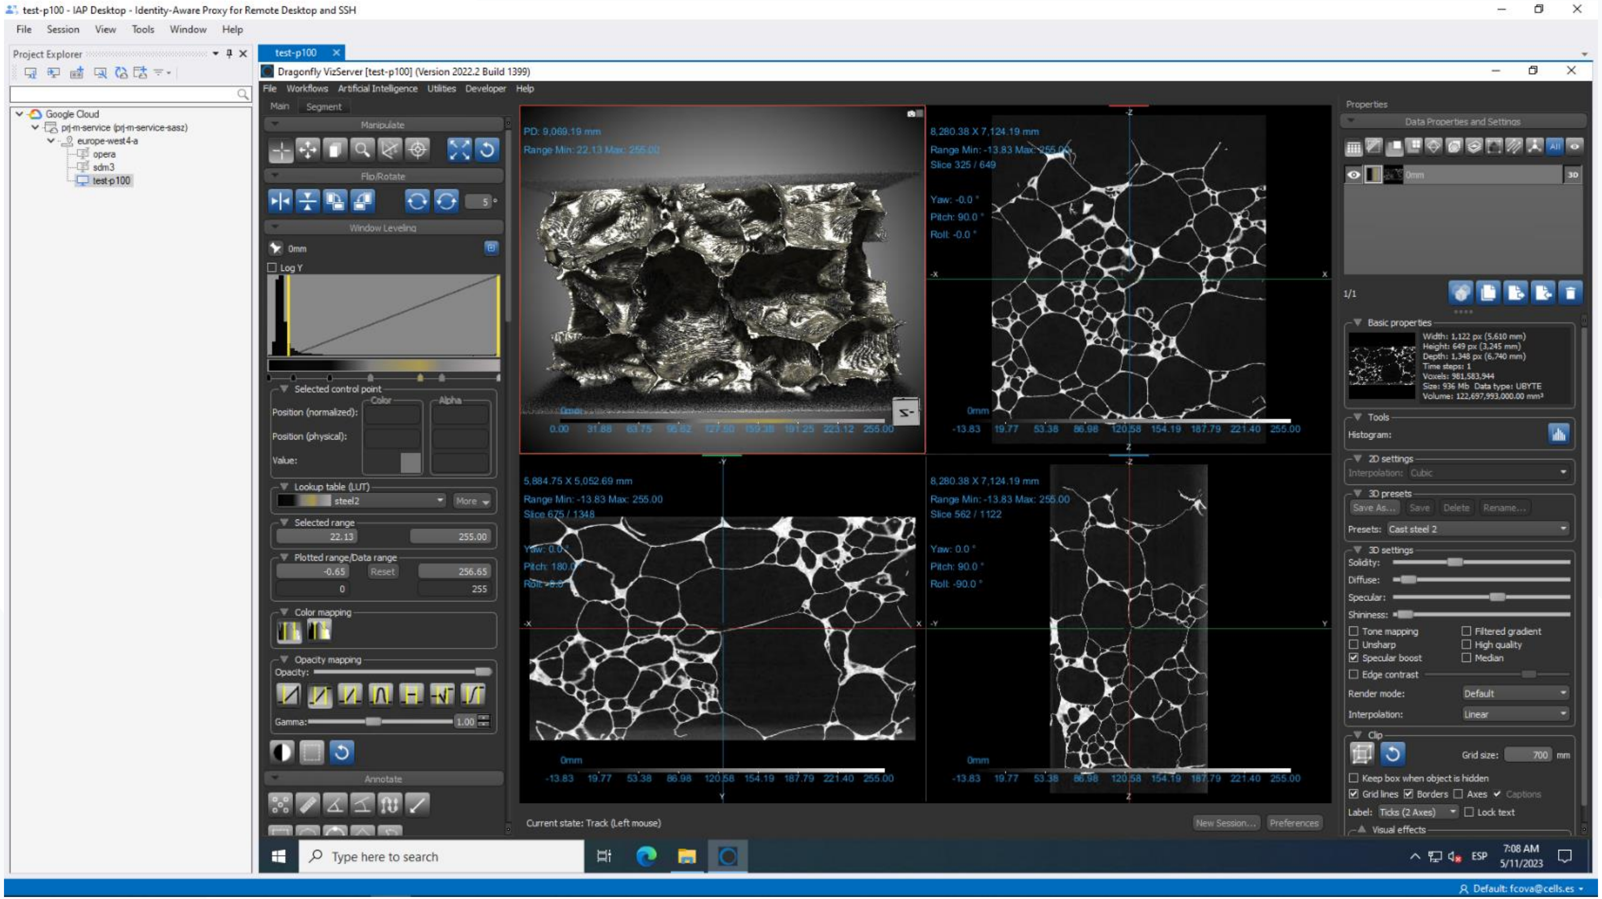

Image processing with Dragonfly

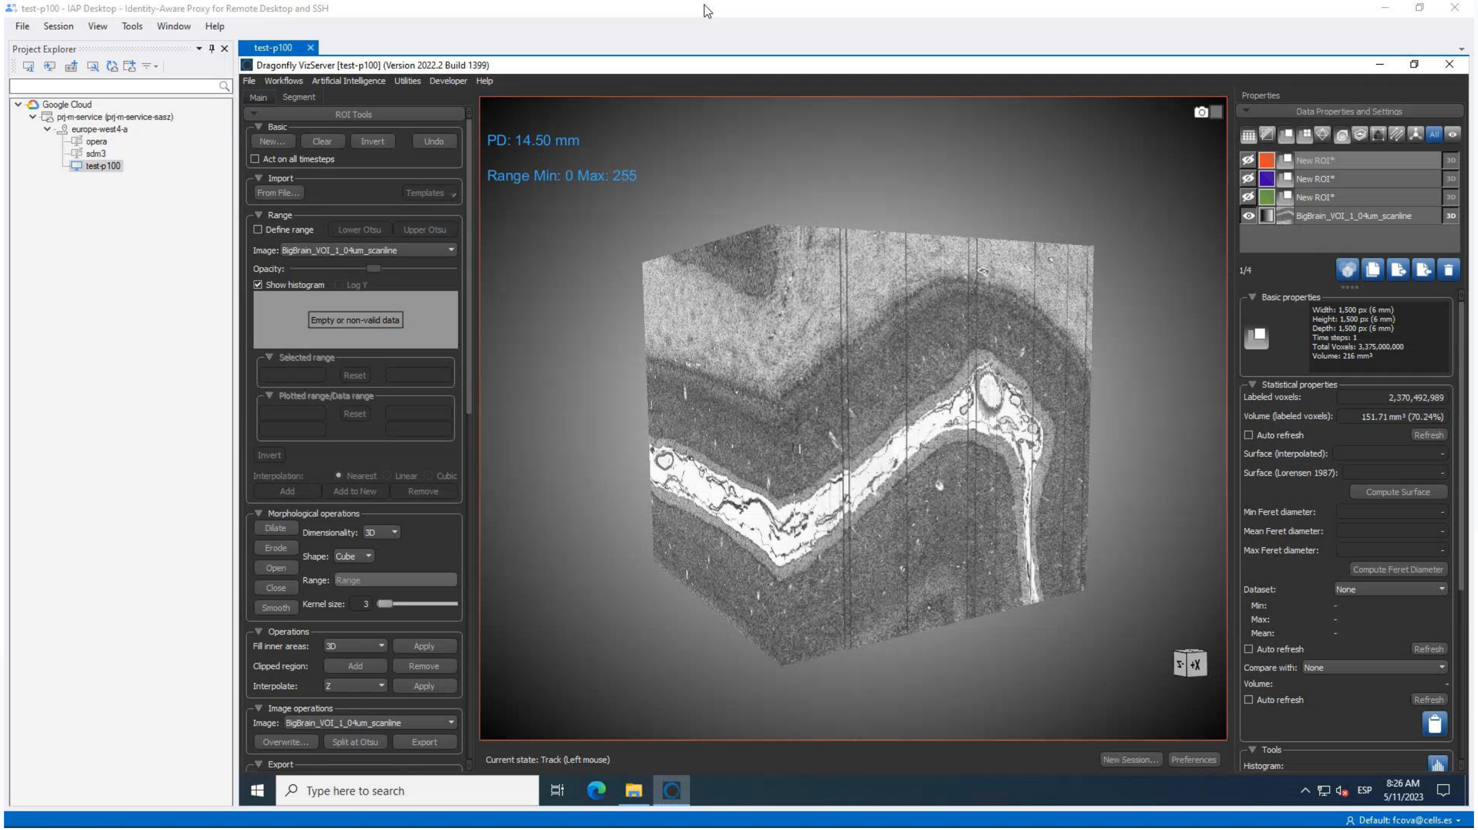

Opera

- Electromagnetic and electromechanical simulation software
- Calculation process with Opera 3D post-processing

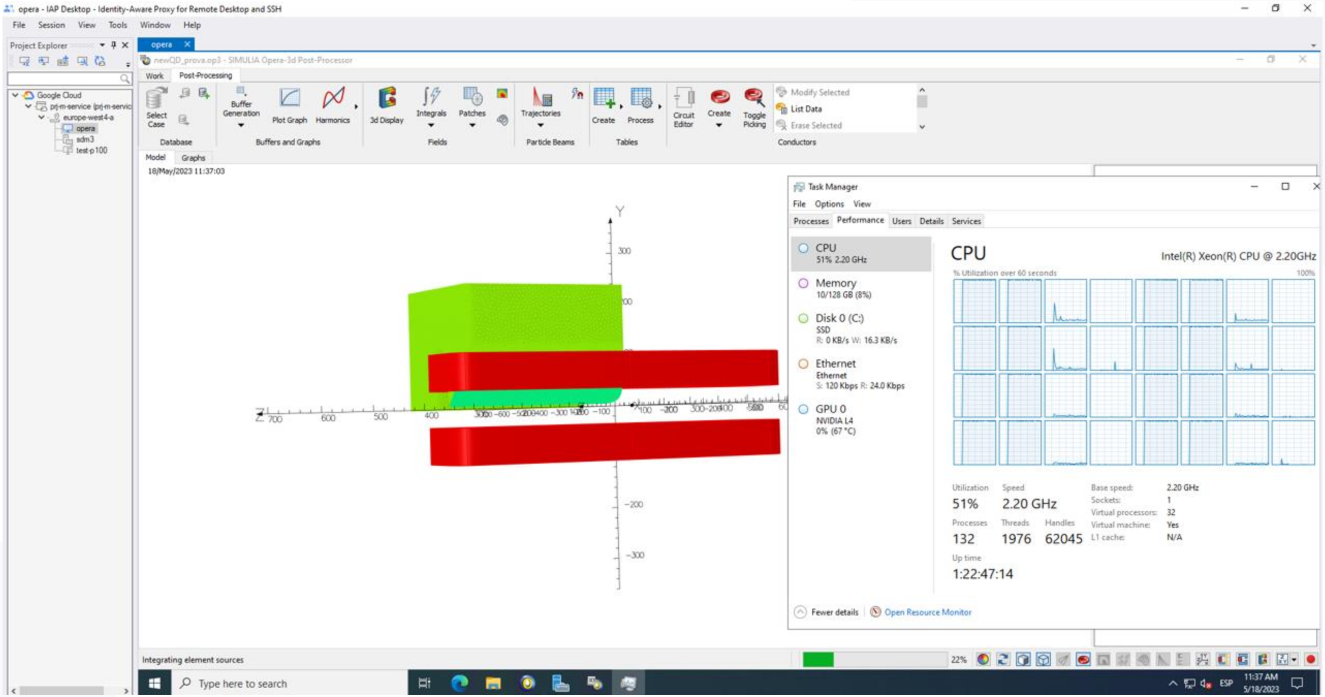

Opera

- Service: Compute Engine
- Instance type: g2-standard-32 (32 vCPUS, 128 GB memory)
- GPUs: 1 x Nvidia L4 (with Virtual Workstation, ex Nvidia GRID)
- OS: Windows Server 2022
- Engineers access: IAP Desktop (Identity Aware Proxy, using their own Google credentials) / RDP

Magnetism data processing with custom libraries

- Goal: Higher quality reconstructions could be generated in much less time
  - The goal is to reduce a sample processing of a 29 GB dataset from 12 hours to 1-2 hours
  - More reconstruction processes could be launched simultaneously
  - Scientists should be able to get better scientific results earlier
- Results:
  - The pre-processing time reduced from 2h26m at our workstations to 37m in c2-standard-16

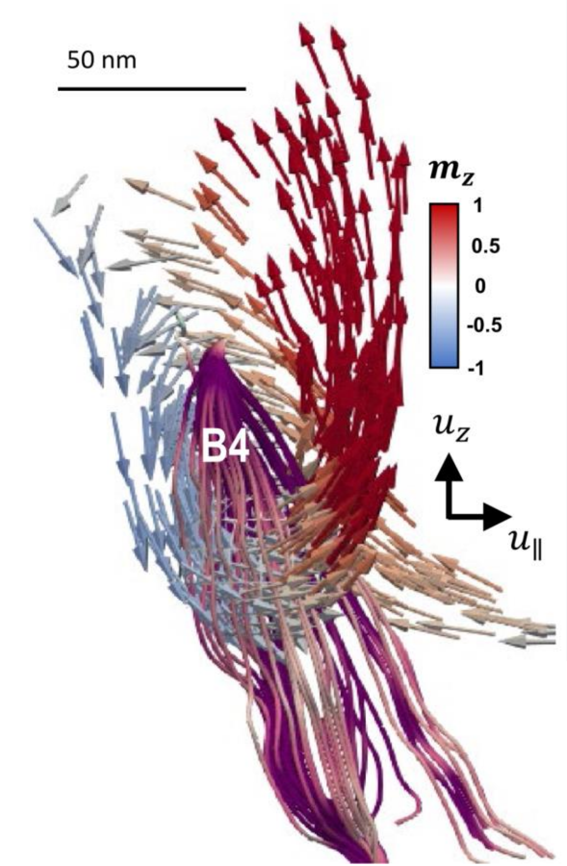

Magnetism data processing with custom libraries

- Service: Compute Engine
- Instance type: (preemptible instances)
  - c2-standard-16 (16 vCPUs, 64 GB memory)
  - c3-highcpu-176 (176 CPUs, 352 GB memory)
  - n2d-standard-224 (224 vCPUs, 896 GB memory)
- OS: Debian 11
- Engineers access:

  gcloud compute ssh [instance] –ssh-flag="-X" --tunnel-through-iap (Identity Aware Proxy, using their own Google credentials)

Conclusions

- Advantages of using the CLOUD over on premise
- When? For what?

Advantages of the CLOUD vs. ON-PREMISES

- Access to the services and data from anywhere with an Internet connection
- Scalability
- Flexibility for testing different hardware
- Quick tests with different VMs types
- Hardware and services quickly available (vs public tenders)

Usually mentioned advantages that were not in our case:

- Latency
- Cost efficiency (it depends, specially if you already have a DC)
- Disaster recovery (it depends)
- Automation (Infrastructure as Code can be also used on premises)

Challenges of the CLOUD

- Improve data transfer
- Integrate with existing workflows
- Extend use to scientists
- Difficult to estimate the costs in advance
- Security, data policy, what can we upload and what not?
- Vendor lock-in - Terraform
- Hybrid cloud might be the perfect solution

Acknowledgments

IT Systems section – Setting up the cloud infrastructure

- Adara Carrión
- Jordi Salabert
- Ramon Escribà

Scientific Data Management section – Creating use cases and tests

- Nicolas Soler
- Gabriel Jover
- Przemyslaw Karczmarczyk

Scientists of the beamlines - experiments division

- Federico Cova – BL31 FAXTOR
- Alessandra Patera – BL31 FAXTOR

Accelerators division

- Valentí Massana

THANK YOU