Referaggio Richieste Risorse di Calcolo Commissione Scientifica Nazionale 5

Anno 2023

Alessandro Lonardo

INFN Sezione di Roma

6-7 Settembre 2022

Il processo (1)

1. **Raccolta delle richieste** di risorse di calcolo da parte delle sigle (attive e nuove proposte) in fase di preventivo. Predisposto un form e interagito con i responsabili delle sigle:
   - Selezione delle richieste, filtrate:
     - Richieste “sotto soglia” (pc desktop, …)
     - Dispositivi da integrare in apparati sperimentali (host per schede DAQ, sistemi di sviluppo, …)
     - Richiesta per un server di calcolo da gestire localmente (SETIN2)
   - Classificazione della tipologia di risorse richieste (Cloud/HPC)
   - Individuazione della tempistica per l’uso delle risorse
     - Immediata: richiesta SJ per il 2022 (AI_MIGHT)
     - Inizio del 2023: sigle attive e nuove proposte
     - Legata alla disponibilità sul mercato della tecnologia (ML_INFN)

Il processo (2)

2. **Inserimento delle richieste** selezionate nella sigla CALC5_Tier1 (tot. 221k€)
   - Criterio di quantificazione del costo (<u>per eccesso, x3?)</u>
     - Preventivi presentati per l'<u>acquisto</u> delle risorse richieste
   - Risorse "HPC"
     - Core-hour e GB per core di memoria RAM (concordati con i richiedenti fissando una ragionevole percentuale di uso delle risorse richieste in origine come acquisti)
     - Numero e tipologia di acceleratori GPU
     - TB storage di rete
   - Risorse "Cloud"
     - Caratteristiche delle macchine richieste in origine
3. **Post chiusura preventivi**
   - Revisione delle richieste
   - Interazioni con i responsabili
   - Eliminazione/Rimodulazione/Accorpamento delle richieste

Stato delle richieste al termine del processo

Sigla CS	Tipologi	Descrizione delle risorse richieste	Nome e cogn
AI_MIGHT	CLOUD	a) Server con 64 core, 256 GB RAM, equipaggiato con 4 schede GPU NVIDIA A100 80 GB b) 50 TB con possibilita' di storare dati sensibili (parte su dischi NVMe da usare durante elaborazione)	Setareh Fatemi
nextAIM	CLOUD	Richieste accorpate con quelle di AI_MIGHT	Francesca Lizzi
FRIDA	HPC	a) CPU 1M core-hour, 2GB/core (TPS, codici PIC e idrodinamici, modelli biofisici) b) CPU 400k core-hour, 6GB/core per attivita' WP1 Roma3/TIFPA. c) 1 GPU A100 40 GB per 60gg/anno. d) 60 TB storage, protocolli NFS/ssh/ftp.	Alessio Sarti
FUSION*	HPC	a) CPU 200k core-hour, 4 GB/core. b) 1 TB Storage	Raffaella Testoni
ML_INFN**	CLOUD	a) 1 Server in grado di ospitare fino a 5 GPU NVIDIA Hopper. b) 1 GPU NVIDIA Hopper. c) 30 TB NVMe per caching locale dei dati per attività di training.	Lucio Anderlini
QUANTEP	CLOUD	2 workstation ognuna con 2 CPU Intel Xeon-Gold 6252N (2.3GHz/24-core/150W), 256 GB RAM, 4 TB storage, per progettazione con ANSYS/Lumerical.	Andrea Salamon
SETIN2*	HPC	a) 140k core hour, 10 GB/core. b) 3 TB	Gaetano Salina

nuove sigle 2023 (non ancora approvate)
****sigla in chiusura nel 2023 che richiede estensione (da approvare)***

AI_MIGHT

Artificial **I**ntelligence methods applied to **M**edical **I**ma**G**es to en**H**ance and personalize BNCT **T**reatment planning

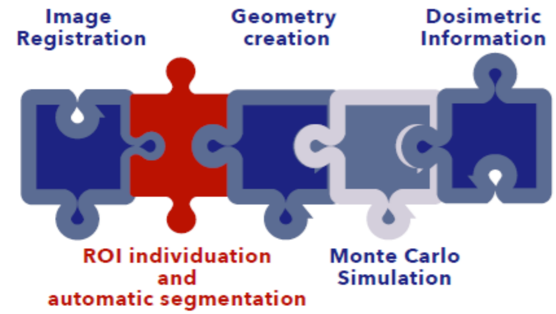

AI_MIGHT aims to use Deep Neural Networks **DNN** to **automatically segment** medical images used in **Boron Neutron Capture Therapy.**

Specifically the automatically segmented images will be used to create the geometrical input needed by the BNCT **Treatment Planning System**

HEAD & NECK CANCER

GLIOBLASTOMA MULTIFORME

AI_MIGHT will focus on two types of tumuor commonly treated in BNCT clinical centers:

Head & Neck cancer and Gliomas

Two different imaging modalities will be taken into consideration for the project. **CT images represent the gold standard** for the diagnosis and are used to evaluate the treatment plan combined with PET images.

MRI images have recently become object of study in the BNCT community to replace the CT+PET methodology commonly used so they represent **the future in this field**.

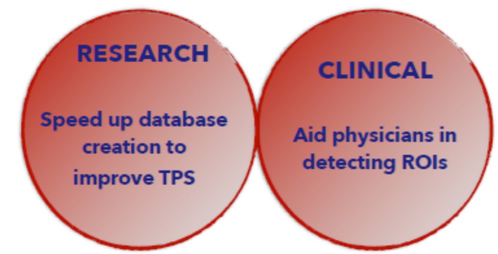

AI_MIGHT aims to create a software able to analyze large amounts of data in short time that can be helpful **both for clinicians** studying the best treatment plan for the patients **and for researchers** aiming to enhance their treatment planning systems.

Moreover, AI_MIGHT would investigate the use of MRI images for the dosimetric calculations and study the feasibility of a TPS based on MRI images.

Work packages and objectives

WP 1: Acquisition and standardisation of the medical images

First step will be acquiring the medical images for both tumor types and both imaging modalities. The images will undergo a standardization process, data augmentation and for MRI images data harmonization. For the realization of this part of the work, the necessary resources are **computational power and HD space.**

WP 2: Training and testing of the segmentation algorithms

Once the image database for each case has been set up, the collaboration will start **training the segmentation algorithm** for each case. To train the algorithm of full resolutions medical images the GPU computational power required is high, therefore, we would like to acquire at least **2 NVIDIA A100 GPU with 80 GB memory.**

WP 3: Implementation of the automatically segmented ROIs as input for the TPS.

Once we obtain the segmented ROIs from the previous step we will perform TP with the automatically segmented images. The TPS we will use is NeuboronPlan, the system creates the geometry needed for the dosimetric calculation via Monte Carlo code (such as MCNP). For the Monte Carlo code we are going to use the in-house CPU servers we have access to.

	Q1	Q2 (Begins now)	Q3	Q4	Q5	Q6
WP1	Pipeline CT for GBM Pipeline CT fo H&N	Pipeline MR for GBM	Pipeline MR for H&N			
WP2		Training nnUnet for CT images bot GBM and H&N			Training nnUNet for MR images both GBM and H&N	
WP3					Testing for CT and MRi images	Testing real cases

AI_MIGHT

Tasks and computing resources

AI_MIGHT asked for server with 2 NVIDIA A100 80 GB GPUs

The server hosts a total number of 4 GPUs so it has been considered as a shared facility to allocate the GPUs required by other projects, for example NEXT AIM.

As shown also by the Next AIM project the amount of VRAM needed to train a nnUnet on full resolution medical images is a known problem thus prompting our requests for high VRAM new generation GPU's

We did a benchmark on a facility in the University of Pavia using 1 V100 32 GB GPU on a simplified 3D UNET.

We used batch size 1 and just 1 channel for each image. Moreover, we used a simple dice function.

With all these constrains we could study a 256*256*256*1 image only with a shallow network of depth 3.

If we want to use a deep network with depth 7 we are limited on using 256*256*128*1.

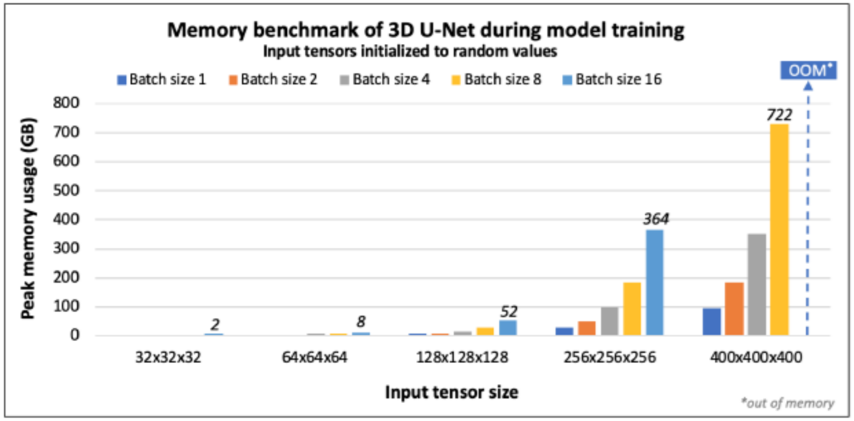

*Figure 3. Benchmarking the memory usage of 3D U-Net model-training over various input tensors sizes on an Intel Xeon Scalable Processor-based server with **1.5 TB system***

https://dl.dell.com/manuals/common/dellemc_overcoming_memory_bottleneck_ai_healthcare.pdf

CT images used for these studies are commonly bigger and MRI images are more complex with more than one channel.

Therefore to analyze full resolution images and not be limited we ask for 2 NVIDIA A100 80 GB GPU's.

next AIM

next AIM: research topics, challenges and implementation

WP1 **Challenge I: no-so-big data**

Strategies for efficient learning with limited data samples.

Evaluation of robustness and reliability of trained models.

WP2 **Challenge II: explainable AI (XAI)**

Make AI results understandable to humans.

Which image/data features were relevant to make a decision?

WP3 **Applications to real-world data samples**

Practical medical data analysis use cases on available samples (public data, private collections, integration of both), where the analysis approaches a, b or both are implemented, and the challenges I, II or both are encountered.

Implementation, test and validation in collaboration with colleagues working in Clinical context

WP4 **Computing resources and SW repository organization**

(ReCaS, IBiSCo, INFN-Cloud + risorse HW locali)

WP5 **Exploitation of research results and communication**

(connection with AIFM, conferences, publications)

==> https://agenda.infn.it/event/30287/

next AIM

next AIM: requested computing resources

We ask for:
- fast access storage (M2), **50 TB** (cost 10kE)
- **2 x NVIDIA A100** (2 x 80 GB) (H100 if available)

- **13 sections involved** (BA, BO, CA, CT, FE, FI, GE, LNS, MI, NA, PI, PV, PD);
- every section needs **computing resources for training algorithms**;
- **12 tasks:** radiomics, machine and deep learning;
- every task is **shared by more than one section**;
- medical images are **high resolution** images (or 3D images) that cannot always be downsampled since details are important.

Why?

We **need shared computing** resources to work together on the same task.

Computing is a pivotal part of the nextAIM experiment: HP hardware allows us to test the **state-of-the-art architectures and algorithms** that may be fundamental for achieving the experiment goals.

Managing such a computing resources will increase and improve our freedom in developing new approaches and positively impact on our expertise in the computing field.

The analysis of 3D images requires **GPUs with high VRAM**.

Some medical exams (e.g. mammography) consist of **many images that should be processed all together**.

We need to decrease as much as possible training time. The training of a 3D V-NET may require many days.

next AIM

next AIM: example of analysis and tasks

If we suppose to process Computed Tomography (CT) scans for the quantification of COVID-19 infection in the lungs::

A typical Chest CT has size: 512x512x100 (the slice number is highly variable from about 50 to 250 slices)

Such 3D images need to be processed in their 3D form since information on the 3 axis is correlated.

Considering the segmentation task, the state-of-the-art architecture is the U-Net.

The training of the U-Net on about 200 CTs, downsampled to 200x150x100, on a V100 (16 GB VRAM)requires about **5 days with batch size equal to 1.**

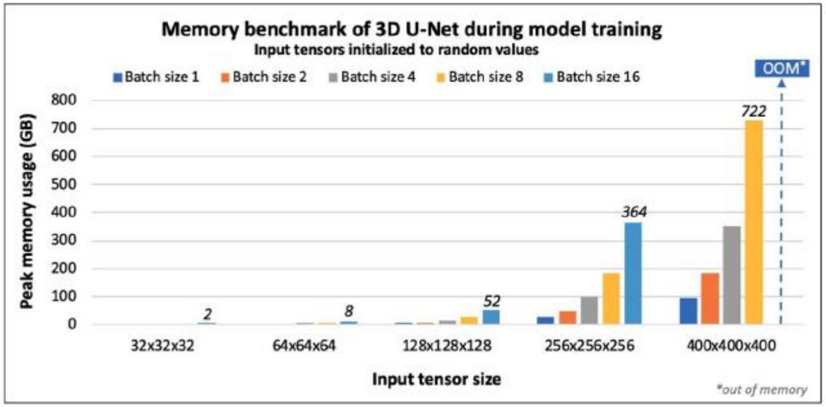

*Figure 3. Benchmarking the memory usage of 3D U-Net model-training over various input tensors sizes on an Intel Xeon Scalable Processor-based server with **1.5 TB system***

The high amount of VRAM required for the training with 3D images is a **well-known problem** in literature.
Image taken from:
https://dl.dell.com/manuals/common/dellemc_overcoming_memory_bottleneck_ai_healthcare.pdf

next AIM

next AIM: example of analysis and tasks

If we suppose to process Computed Tomography (CT) scans for the quantification of COVID-19 infection in the lungs::

A typical Ch
has size:
512x512x1(
(the slice
is highly v
from abo
250 slices

Such 3D

The training of the
about 200
sampled to
0, on a
3
ires about
h batch
to 1.

Many tasks of next-AIM experiment require the **analysis of 3D** images and imply the use of **deep learning**:

- **T3.2** Super-Resolution in DBT;
- **T3.6** Connectivity in functional MRI;
- **T3.7** Radiomics and Deep Learning analysis of CT and patients' data in COVID-19;
- **T3.8** Radiomics and ML-segmentation on Facio-Scapulo-Humeral dystrophy (FSHD) and lung tumor;
- **T3.11** Machine Learning techniques for cardiological applications.

Moreover, some state-of-the-art architecture, such as **DenseNet** requires even more VRAM than U-Net!

ne
ges is
em in

uals/co
oming_
ck_ai_healt
hcare.pdf

*Figure 3. Benchmarking the memory usage of 3D U-Net model-training over various input tensors sizes on an Intel Xeon Scalable Processor-based server with **1.5 TB system***

FRIDA

- Three main research lines need significant computing power:
  - Optimisation of treatment plans taking into account the FLASH effect
  - The study of the beam propagation and collimation in laser plasma acceleration of Very High Energy electrons for therapeutic applications
  - The simulation and modeling of the different stages characteristics of the FLASH effect by means of MC software capable of covering the time scale ranging from the physical interaction of the incoming radiation up to the chemical diffusion of radicals and Oxygen depletion or enrichment that characterise the biology of the flash effect, as well as the DNA repair mechanisms.

Treatment planning (no GPU):
~1000k core/24h for 1 dose calculation in a complete treatment: different energies and shooting points have to be investigated, for each patient…

+ GPUs/normal CPUs with 20GB ram for the treatment optimisation

Prostate treatment 8Gy, IMRT - like, VHEE, 7 fields

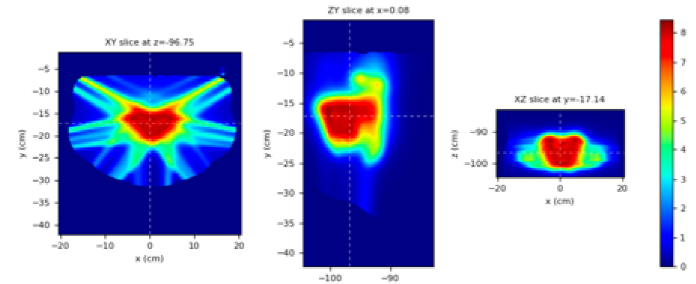

The output requires ~ few TB for each treatment.. [all dose maps with CT granularity have to be stored.]

FRIDA

- Three main research lines need significant computing power:
    - Optimisation of treatment plans taking into account the FLASH effect
    - The study of the beam propagation and collimation in laser plasma acceleration of Very High Energy electrons for therapeutic applications
    - The simulation and modeling of the different stages characteristics of the FLASH effect by means of MC software capable of covering the time scale ranging from the physical interaction of the incoming radiation up to the chemical diffusion of radicals and Oxygen depletion or enrichment that characterise the biology of the flash effect, as well as the DNA repair mechanisms.

Laser plasma beam acceleration studies: 100core/60gg with 2 GB RAM/core memory requests.

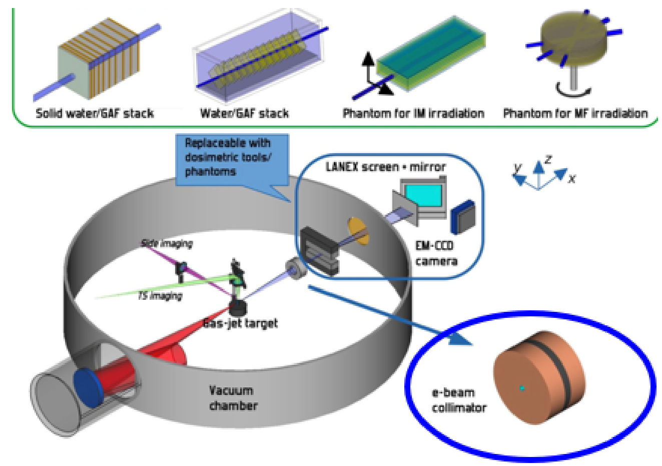

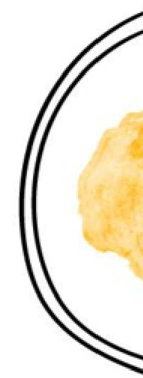

optimization of the collimation system for beamlet shaping;

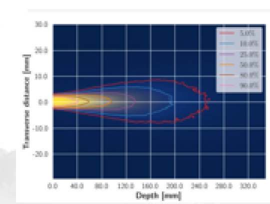

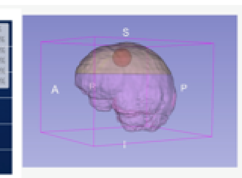

FRIDA

- Three main research lines need significant computing power:
  - Optimisation of treatment plans taking into account the FLASH effect
  - The study of the beam propagation and collimation in laser plasma acceleration of Very High Energy electrons for therapeutic applications
  - The simulation and modeling of the different stages characteristics of the FLASH effect by means of MC software capable of covering the time scale ranging from the physical interaction of the incoming radiation (10^{-15}s) up to the chemical diffusion of radicals and Oxygen depletion or enrichment (10^{-7}s) that characterise the biology of the flash effect, as well as the DNA repair mechanisms [time scale here is hours!].

The involved simulation software are: Microdosimetric Kinetic Model (MKM) developed in INFN-RM3 to generalize the DNA damage repair kinetics (DDRK) and the TRAX-CHEM monte carlo

Kinetic models and homo/heterogeneous stages simulations (FLASH effect modelling): CPU 400k core hour, 6GB/core

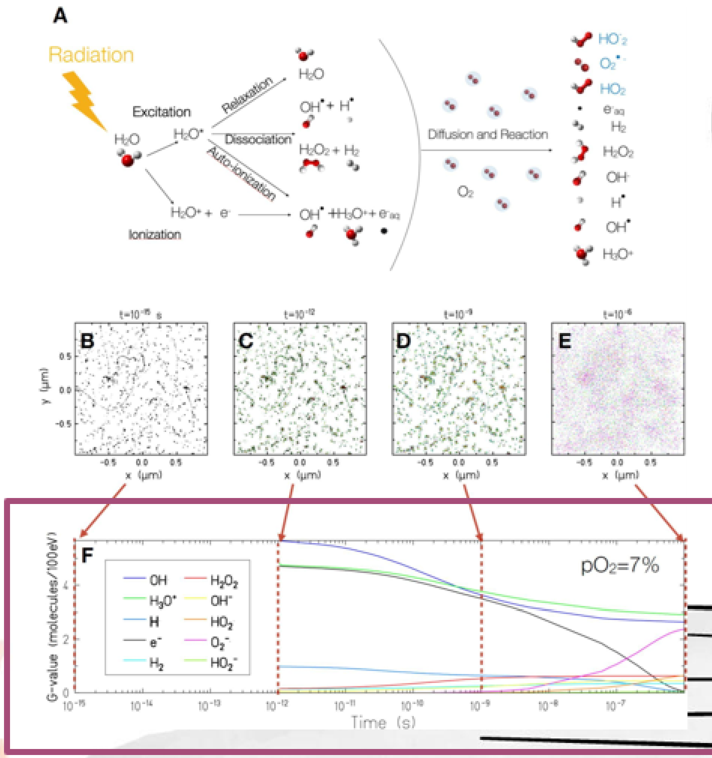

FUSION (FUsion StudIes of prOton boron Neutronless reaction in laser-generated plasma)

R. Testoni, A. Froio, R. Bonifetto
(Dipartimento Energia, Politecnico di Torino)
Responsabile locale: R. Testoni
Totale sede di Torino: 1 FTE

FUSION - Laser-induced p-11B fusion reaction

p(11B, α)8Be

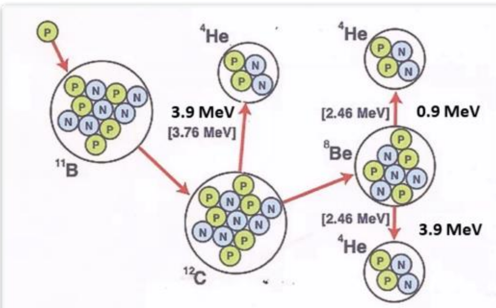

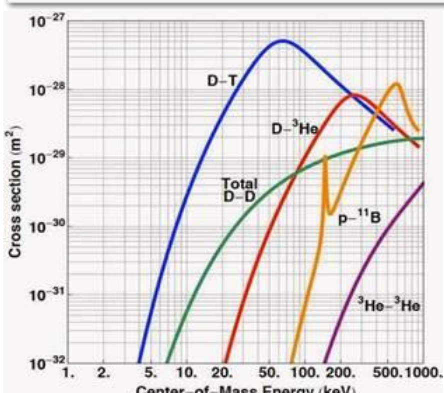

M.Oliphant, L.Rutherford, Proc. R. Soc. London A 141 259 (1933)

- **Neutronless** fusion reactions
- It is considered as a **potential candidate** in inertial fusion scheme
- Reagents **more abundant in nature** with respect to other fusion reactions of interest, and easier to handle
- Interest for the realisation of intense **α sources** for applications

p-^{11}B reaction can be **triggered using high-power laser systems** and could be in future suitable as fuel for Inertial Fusion Energy

Pitcher-catcher scheme

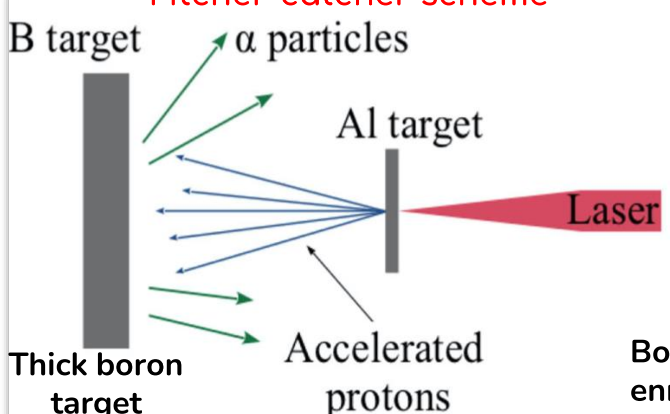

F. Consoli et al., (2020) Front. Phys. 8:561492
D. Margarone et al, Front. Phys. 8 (2020) 343

In-target scheme

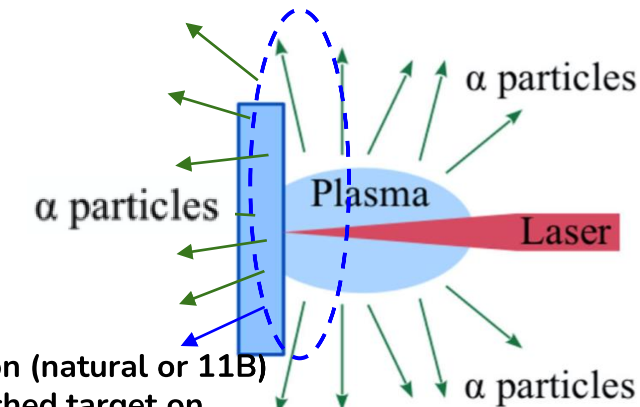

Boron (natural or 11B) enriched target on silicon substrate

A.Picciotto, et al. Physical Review X 4, 031030 (2014)
L.Giuffrida, GAP Cirrone, et al. Phys. Rev. E 101, 013204 (2020)
D.Margarone, et al. Appl. Sci. (2022), 12, 1444

FUSION - modelling and simulation

- The different FUSION experimental activities will benefit of the support from the modelling aspects: the test matrix definition and the interpretation of the results, as well as, where applicable, also the target design.
- Three different modelling techniques will be adopted, namely hydrodynamic (HyD), Monte Carlo (MC) and Particle-In-Cell (PIC) models. They will be properly coupled when required by the physics to be included in the simulation.
- Depending on the target to be designed and tested in combination with a defined laser pulse, the objective and methodology will be different.
- **Politecnico di Torino activities will concentrate on the PIC modeling: computational resources are needed**

The original project 2020-22

The INFRASTRUCTURE

Strong synergy with **INFN Cloud** to deploy and operate a highly customizable **GPU cluster** by employing modern *virtualization* and *container* technologies.

TRAINING and EVENTS

Base and Advanced **Hackathons** organized to ease the access to technologies and concepts useful for ML

e.g. Jupyter, Pandas, Keras, DNN architectures, AI for HEP, GW and Med. Phys.

KNOWLEDGE BASE

Collection of ML **use cases relevant for INFN research**, with open data samples and executable examples

https://confluence.infn.it/display/MLINFN/ML-INFN+Knowledge+Base

The proposed extension 2023-25

Keep the infrastructure updated and ready to host **upcoming AI technologies**

Enhance the *Quality of Service* of the GPU cluster exploring techniques and technologies to **reduce idle time of the GPU resources**.

Consolidate **online training initiatives** and **revive live events** (assuming looser Covid restrictions)

Modernize the KB with PyTorch, JAX and new architectures (GNN, Transformers); organize and document **open datasets** usable for *Data Science*.

Risorse di calcolo 2023

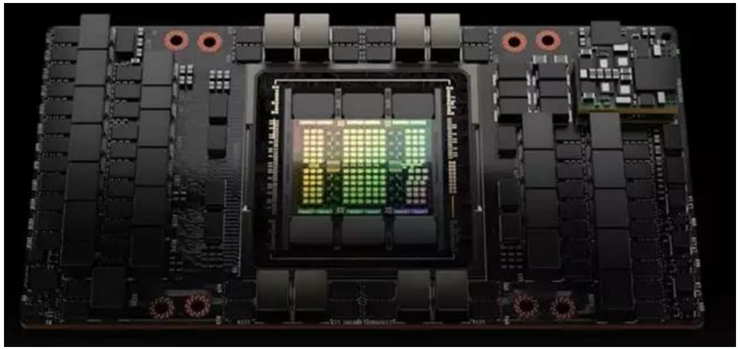

Research focus is on modern computing solutions, including hardware.

To keep the pace with the quickly evolving technological landscape, a frequent modernization of the equipment is of critical importance.

Preliminary interactions with *Huawei* and *nVidia* to establish future collaborations

nVidia Hopper is the upcoming GPU architecture for machine learning workloads in data centers.

ML_INFN will **setup a server** capable of hosting **Hopper GPUs** to provision them through **INFN Cloud** as soon as they land the European market.

More than a half of the GPUs currently hosted by ML_INFN were acquired on external funds.

QUANTEP

QUANtum Technologies Experimental Platform

- **INFN** project on Quantum Technologies (CALL)
- Started in 2021

INFN Sections and Laboratories involved: **LNL**, **MI**, **PG** (Camerino), **PI**, **PV** (Modena e Reggio Emilia), **RM2**, **SA**, **TO**

- Interest and support from: LNGS (LUNA-MV), LABEC (DEFEL), NEST, TYNDALL, Institut Ruđer Bošković (RBI), Micro Photon Devices (MPD), University of Leipzig, Chalmers University of Technology, Physikalisch-Technische Bundesanstalt (PTB).
- 15-17 FTE/year, up to ~ 1 MEuro budget
- Creation of a common **Silicon Photonics** platform for development and characterization of
  - quantum computing circuits;
  - single photon sources;
  - single photon detectors;
  - polarization control circuits.

QUANTEP – Universal Quantum Gates

1 qubit:

Some 1 qubit elementary gates

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \qquad R_\phi = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{pmatrix} \qquad H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

Pauli-X (NOT) gate | Pauli-Z gate | Phase shift gate | Hadamard gate

2 qubits: $a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle \qquad |a|^2 + |b|^2 + |c|^2 + |d|^2 = 1$

The prototype (universal) 2 qubits gate is the Controlled NOT (CNOT) gate

control bit

$$\mathrm{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

target bit

- the control bit is left unchanged
- the output target bit is the XOR of the input control and target bits
- but of course it does much more: it works on the wave function

$$a|00\rangle + b|01\rangle + c|10\rangle + d|11\rangle \longrightarrow a|00\rangle + b|01\rangle + c|11\rangle + d|10\rangle$$

QUANTEP – CNOT Gate

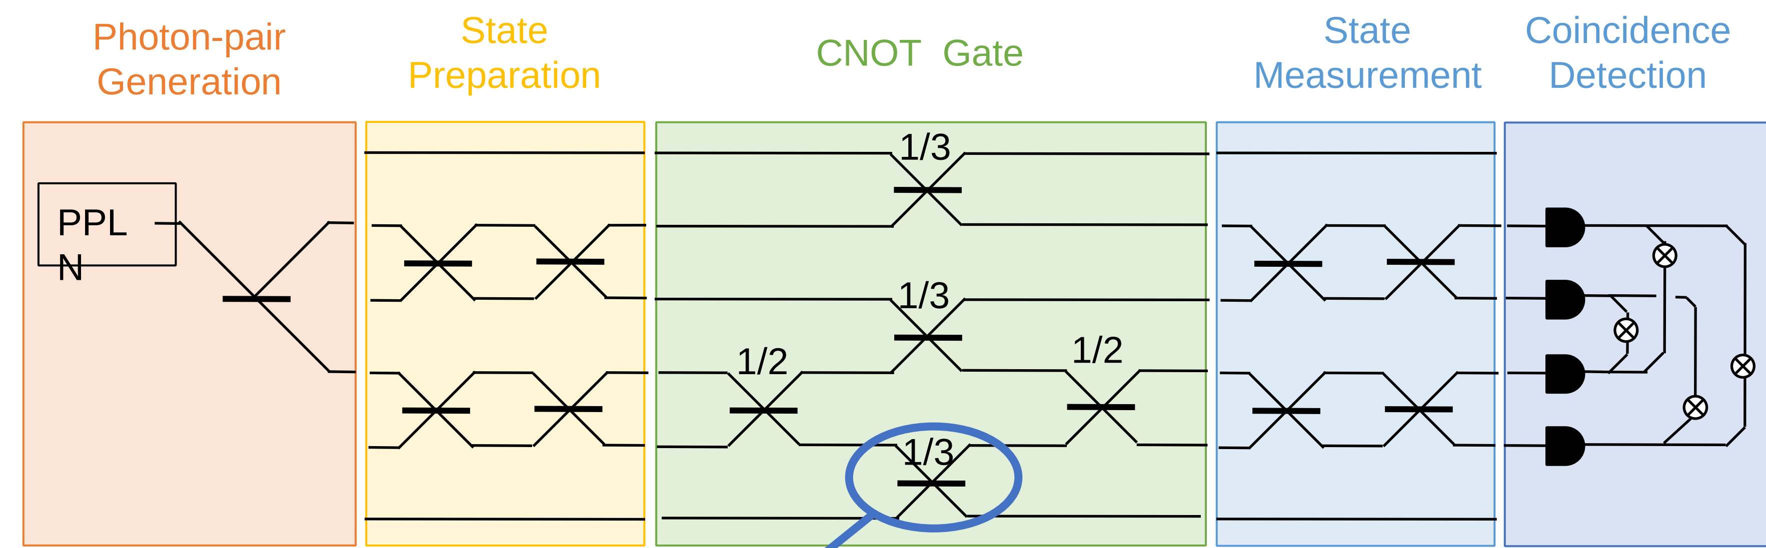

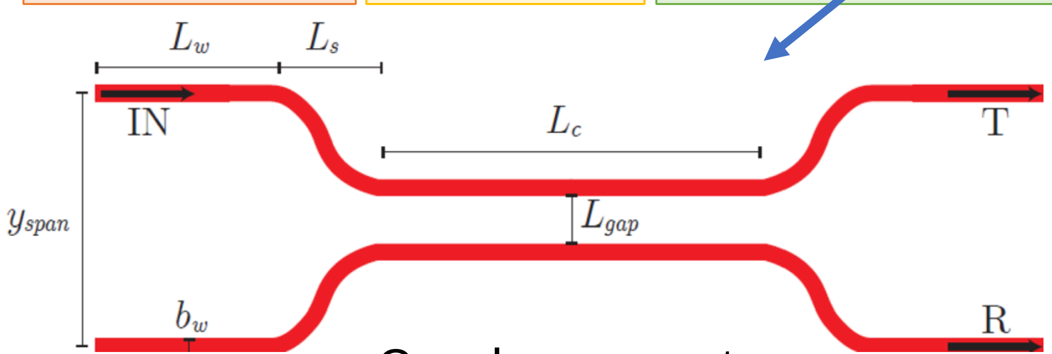

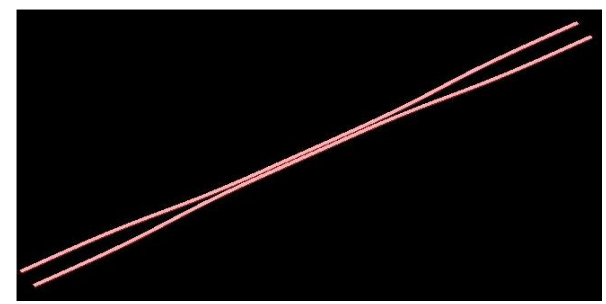

Coupler parameters:
b_w = 0.45 um
b_h = 0.22 um
y_{span} = 2.5 um

QUANTEP – Directional Coupler

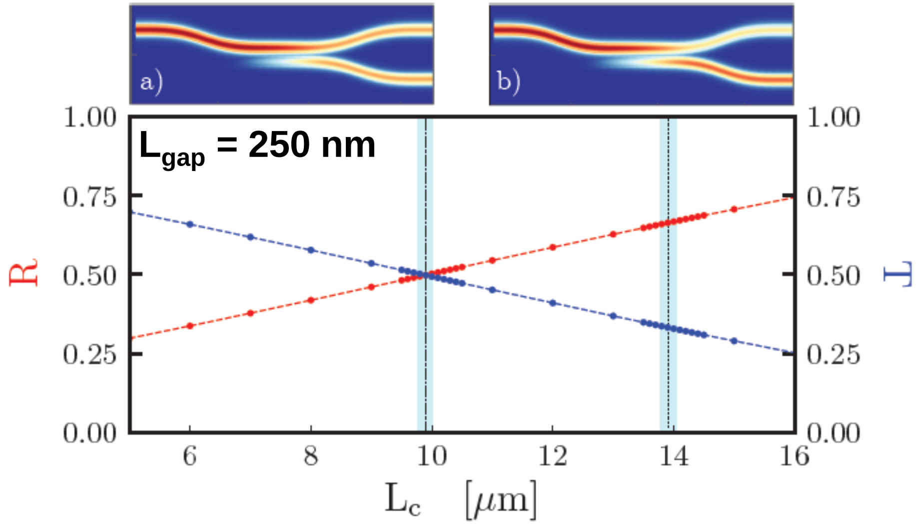

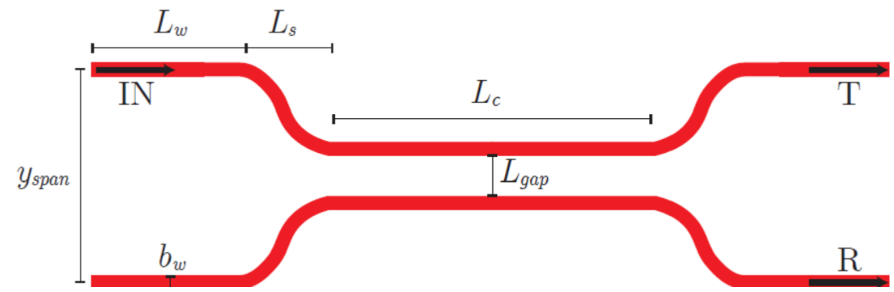

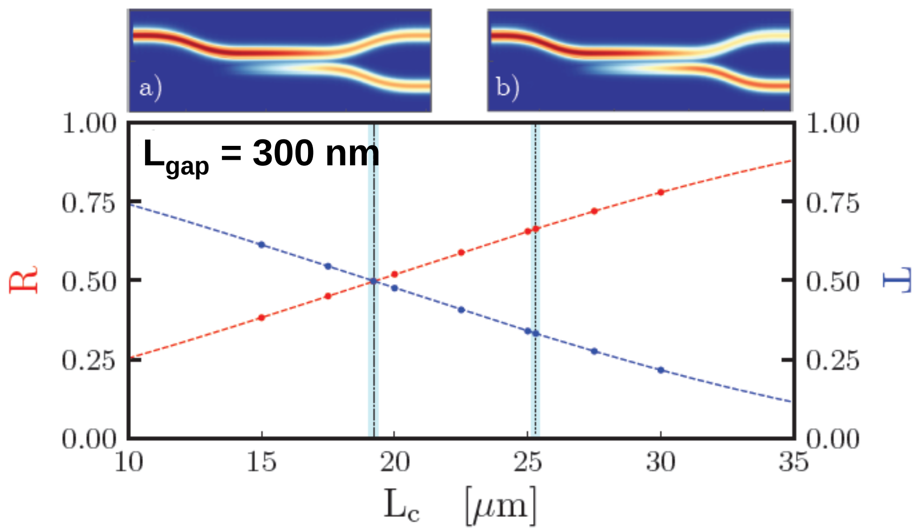

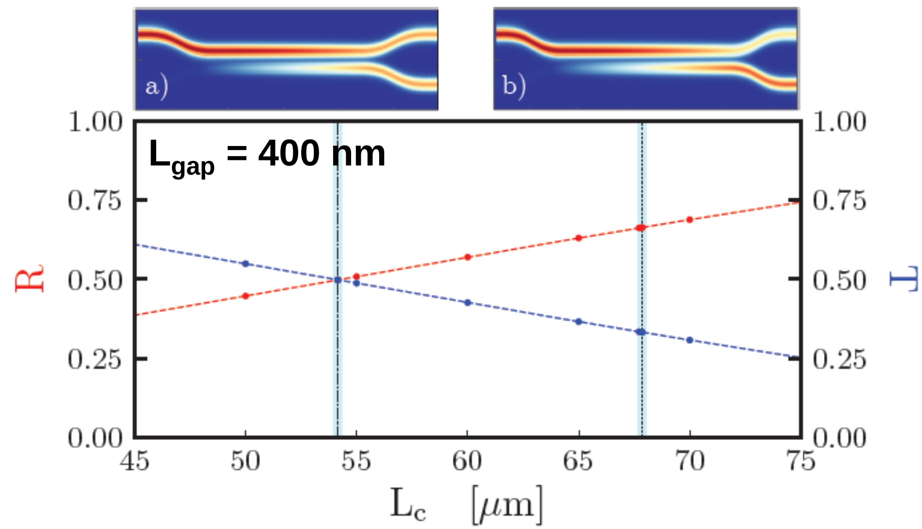

Ansys/Lumerical FDTD: risoluzione delle equazioni di Maxwell con simulazione agli elementi finiti

QUANTEP - Spot Size Converter

The propagation of the mode propagating outside the Spot Size Converter has been studied performing a Lumerical FDTD simulation.

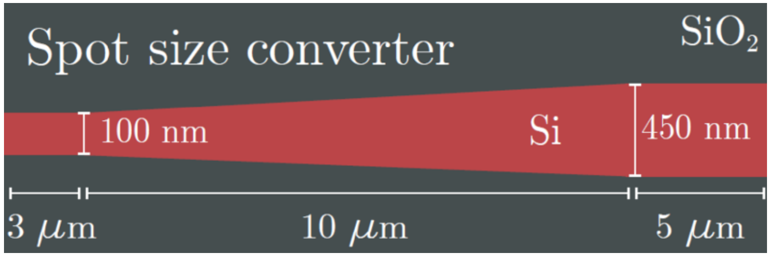

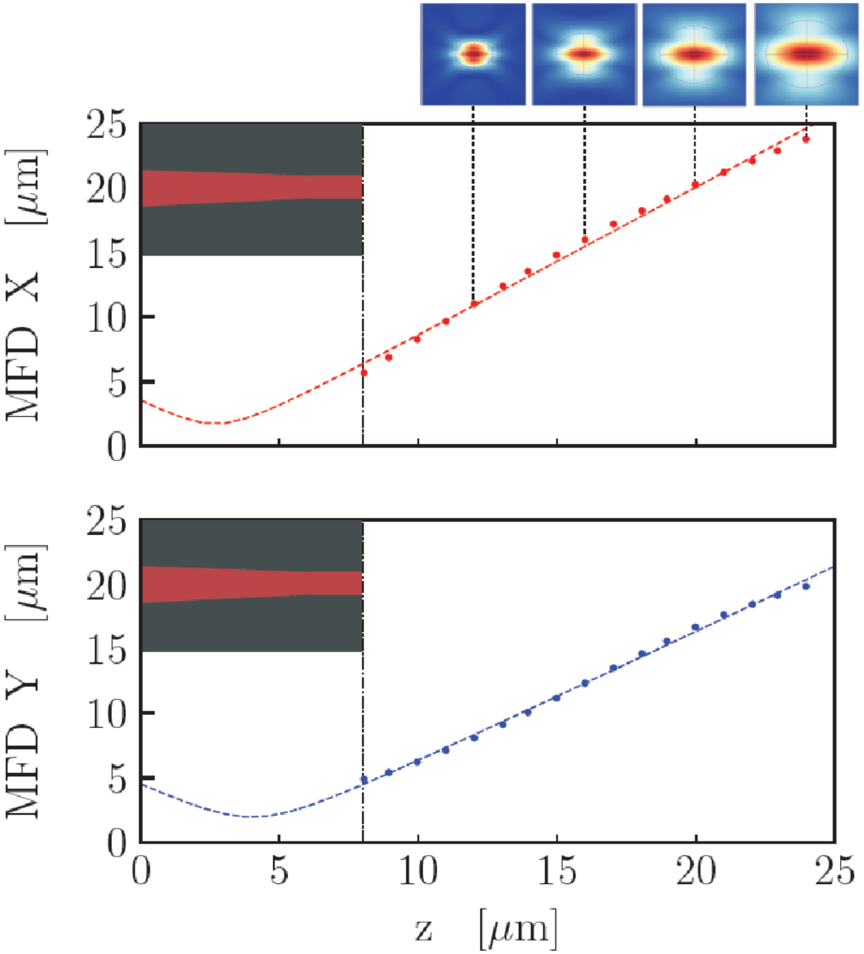

Ansys/Lumerical FDTD: risoluzione delle equazioni di Maxwell con simulazione agli elementi finiti

QUANTEP – Polarization Control

with nanowires, graphene and other 2D materials

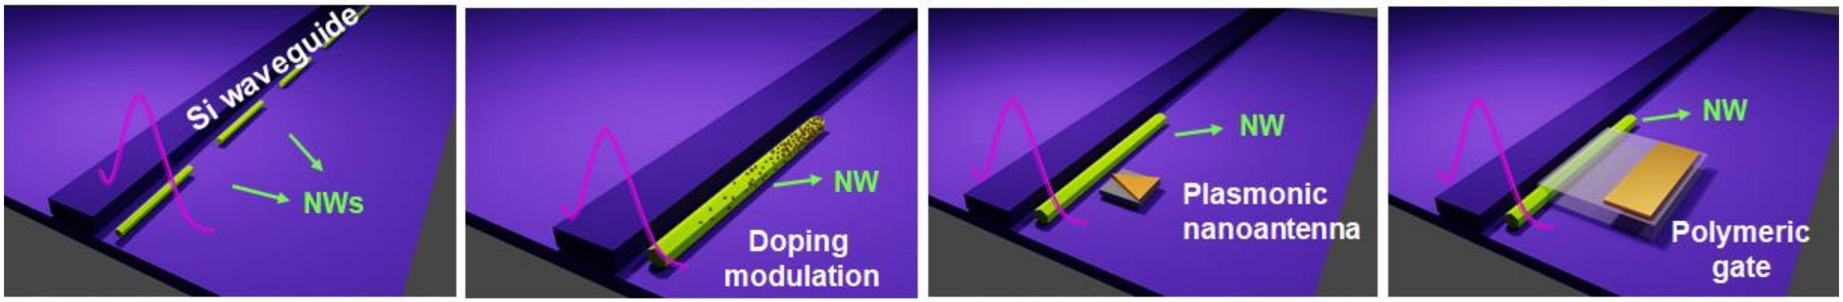

An example: array of nanowires

- Si-wg= 450x220 nm
- NW length= 5um
- NW radius= 60nm
- Material= InP
- Number of NWs= 10
- NW separation (d) = 500nm
- TE -> TM conversion

Input E-field

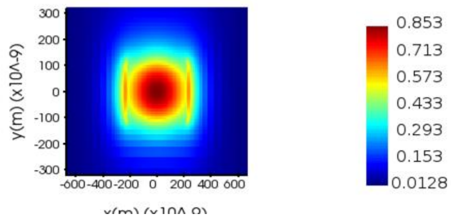

Output E-fields
wl = 1.54 um

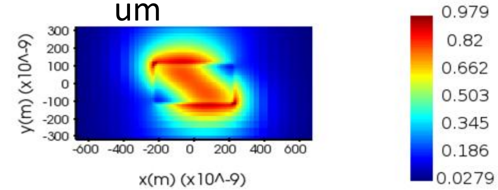

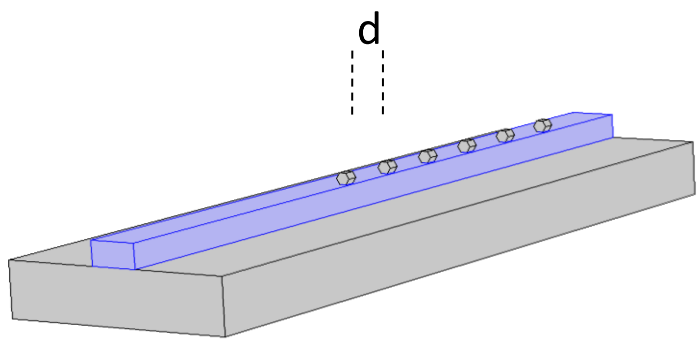

wl = 1.56 um

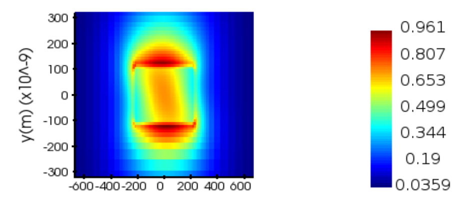

Ansys/Lumerical FDTD: risoluzione delle equazioni di Maxwell con simulazione agli elementi finiti

SETIN2

Proposta Nuovo Esperimento

Sigla: SETIN2
Durata Proposta: 3 Anni
Area di Ricerca: Interdisciplinare
Responsabile Nazione: Davide Cassi (Parma)
Unita' Partecipanti: Napoli, Parma, Roma Tor Vergata e Salerno

The project aims at understanding the fundamental mechanisms involved in the phase transitions of two-dimensional arrays of superconducting islands connected by Josephson junctions. In the resulting networks, due to the interplay between the long-range quantum coherence and topological properties (in the graph theory sense), new quantum orders emerge. The network connectivity becomes much more relevant than the real dimensionality and new physics is expected to arise, as predicted by existing theoretical models.

WP 1: Theoretical study of new topological arrays

Coordinator: Francesco Romeo
Contributors: NAPOLI, PARMA and SALERNO Unit.

The theoretical characterization of selected tree-like arrays will be carried out by studying the physical properties of some basic models (from pure hopping to Bose Hubbard, Ginzburg-Landau approach, etc.), with both analytical and numerical methods, mainly providing predictions on the dependence of measurable quantities.

WP 2: Experimental study of new topological arrays

Coordinator: Roberto Russo (NAPOLI)
Contributors: NAPOLI and ROMA2 Units.

The WP aims to measure JJ arrays based on standard and well established Nb technology with new topological structures defined by the theoretical group as well as to produce and to measure arrays based on materials other than niobium such as aluminum and Niobium nitride as well as Dayem Bridge instead of Josephson tunnel Junctions.

WP 3: Computational study of new topological arrays

Coordinator: Roberto Alfieri (PARMA)
Contributors: PARMA and ROMA2 Units.

The WP mainly concerns complex optimization problems and aims to determine the most effective topology according to which to arrange a given number of junctions, to obtain the greatest improvement of superconducting properties.

SETIN2

WP 3: **Computational** study of new topological arrays

This task includes the development of a prototype [M18] using parallel programming techniques and will be able to exploit modern parallel architectures (multicore and multinode) and libraries (openMP and MPI). The program will calculate key parameters established on the basis of the models provided by the theoretical WP and we will numerically study their dependence on the network topology.

The **execution of the developed tool [M36]** will require the following computational resources:
- HPC cluster of the University of Parma (developed in collaboration with INFN), will be used for the development, testing and processing of configurations of limited size.

- For "medium" configurations (50-100 junctions) the project foresees the exploitation of a dedicated computing node, with a high amount of memory and computing power.

  A suitable computing node to this purposes is the following (which can be found in the CONSIP catalog):
  Server HPE DL560 with 4 Intel 6252N processors and 1TB ram
  Or in a cloud architecture
  2023: 96 core x 12 ore x 120 gg = 140k core hour - RAM: < 10GB x core
  2024 e 2025 : 96 core x 20 ore x 240 gg = 460k core hour - RAM: < 10GB x core
- At the moment it is not clear which is the best approach, in terms of calculation algorithms, for the study of systems of medium-large dimensions. The choice of the best approach to use will be made in 2023, using now traditional approaches (Monte Carlo etc) but they also do not exclude the possibility, necessity of having to develop new algorithms.

Conclusioni

- Primo “esperimento”, stiamo imparando
- Data la tipologia delle sigle in CSN5 (2/3 anni di durata) è fondamentale che le richieste vengano rese disponibili a gennaio, con approvazione delle richieste a settembre.
  - Ambiente software molto variegato
  - Gestione licenze (che succeed con l’uso su Cloud?)
- Questo è vero anche (e soprattutto) in questa prima interazione:
  - Caso AI_MIGHT (richiesta SJ sbloccata nel 2022)
- La possibilità di utilizzare risorse su IBISCO sembra l’unico percorso possible per raggiungere lo scopo il 2023
  - E’ verosimile che queste vengano messe a disposizione per un uso effettivo nel giro di qualche settimana?
  - In caso negativo esistono alternative?