/!\ Disclaimer: Activities shared between WP1 and WP2

Given similar goals, milestones, and deliverables: we agreed for this initial phase of the project to share the activities among the two WPs.
About ~20 participants (minutes)
We propose regular monthly calls .

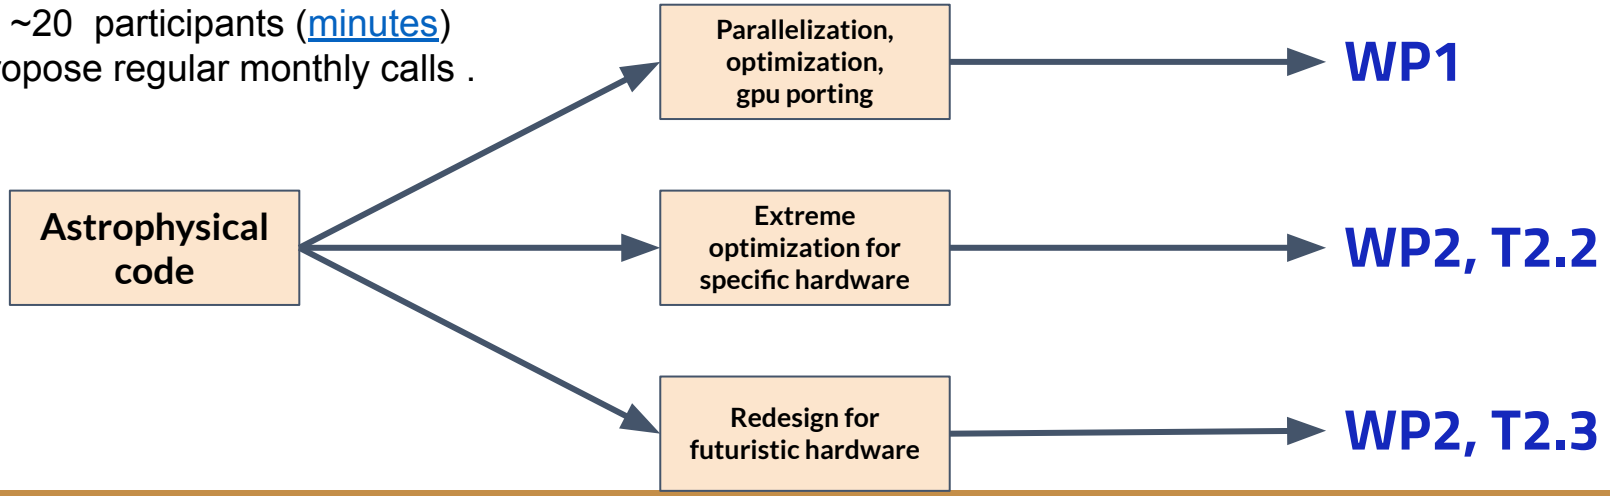

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

ICSC
Centro Nazionale di Ricerca in HPC Big Data and Quantum Computing

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

WP1 - HPC code enabling and optimization

Giuseppe Puglisi & Stefano Della Torre

Spoke 3 General Meeting

12-14 Giugno 2023

Dipartimento di Fisica e Astronomia - Università di Catania

WP1 - HPC codes enabling and optimization

Scope: selects a number of codes that require intensive computational resources to face the next generation of scientific challenges and performs their redesign, reimplementation and optimisation in order to effectively exploit state-of-the-art HPC solutions.

T1.1: **Selection**, Analysis and testing of codes

T1.2: Software **development**, refactoring and optimization

T1.3: Integration, Verification and Validation

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC Big Data and Quantum Computing

Gantt charts

WP1

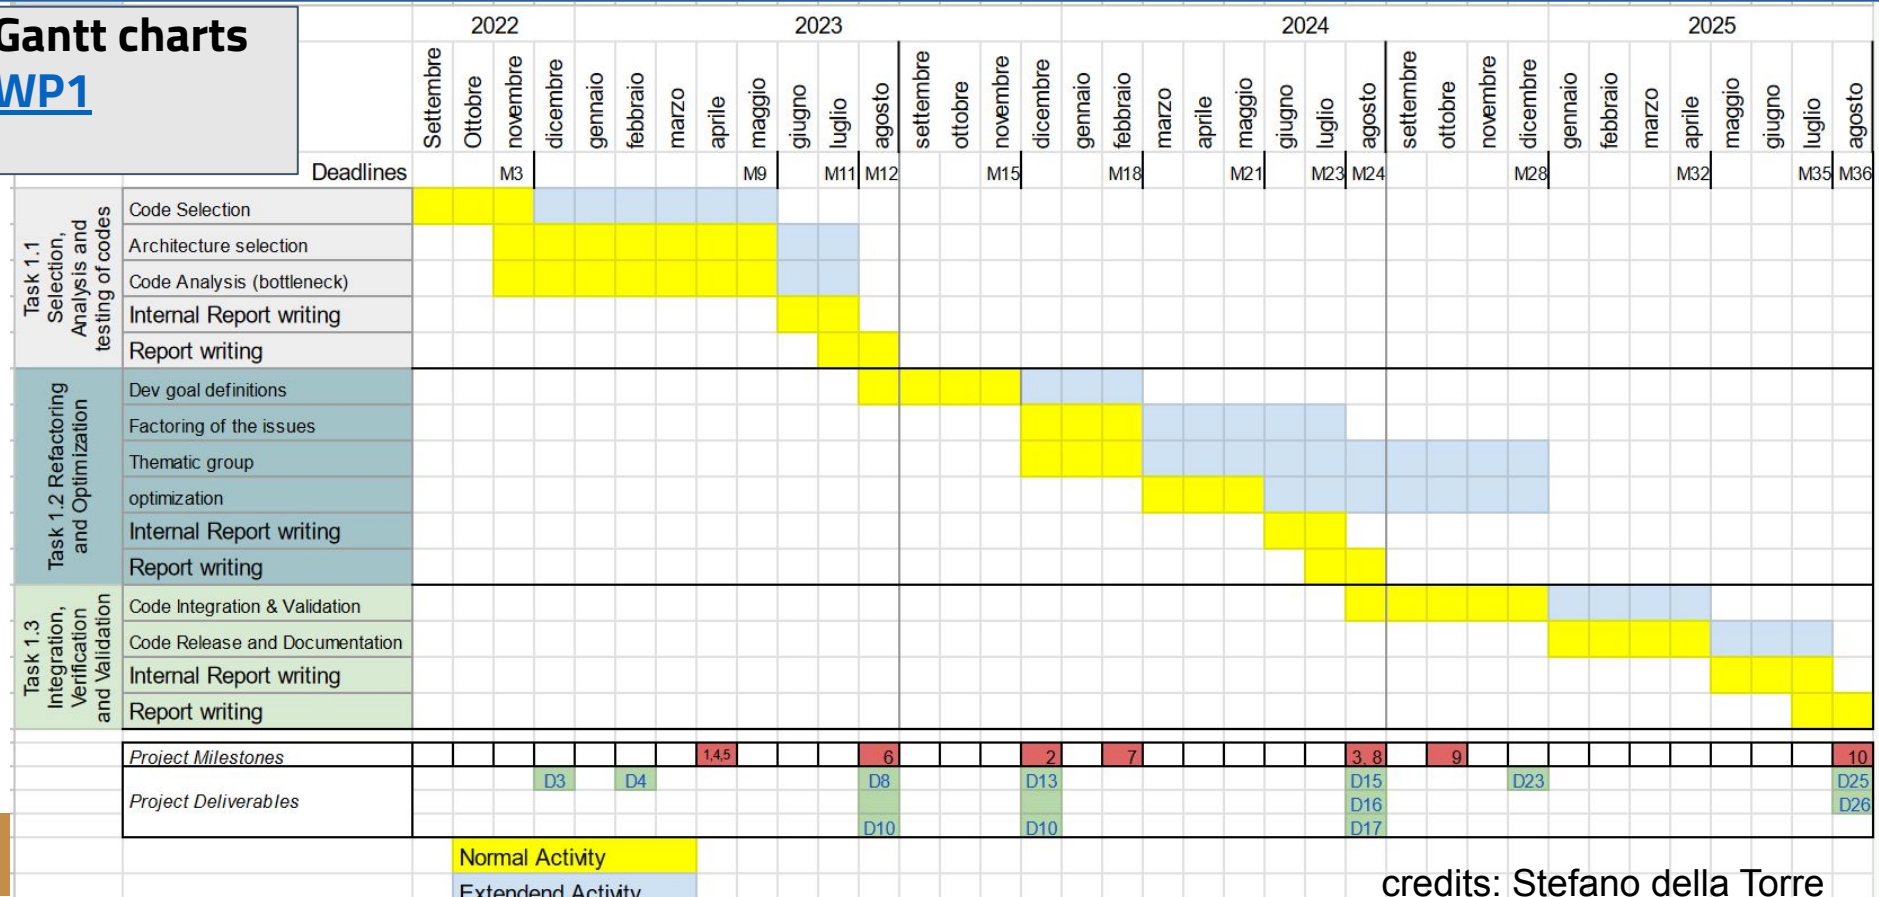

credits: Stefano della Torre

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

ICSC
Centro Nazionale di Ricerca in HPC
Big Data and Quantum Computing

ICSC
Centro Nazionale di Ricerca in HPC,
Big Data and Quantum Computing

WP2 - Design of innovative algorithms, methodologies and codes towards exascale and beyond

Pierluigi Monaco & Mario Spera

Spoke 3 General Meeting

12-14 Giugno 2023

Dipartimento di Fisica e Astronomia - Università di Catania

WP2 - Design of innovative algorithms, methodologies and codes towards exascale and beyond

Scope: This WP identifies innovative algorithms and methodologies upgrading their capability to exploit, and scale on, the exascale and post exascale architectures, reintegrating the resulting improved features in codes, workflows and pipelines. The energy impact will also be specifically considered.

T2.1: **Science cases definition**, algorithms **identification**, parallelism level assessment and profiling

T2.2: **Algorithms co-design** and methodologies to scale-up the capabilities of the algorithms and to find new innovative solutions

T2.3: Design of **new architectural solutions** aimed at the exploitation of post-exascale infrastructure

T2.4: Algorithms and methodologies **integration** into new big-data analysis applications

Gantt charts

WP2

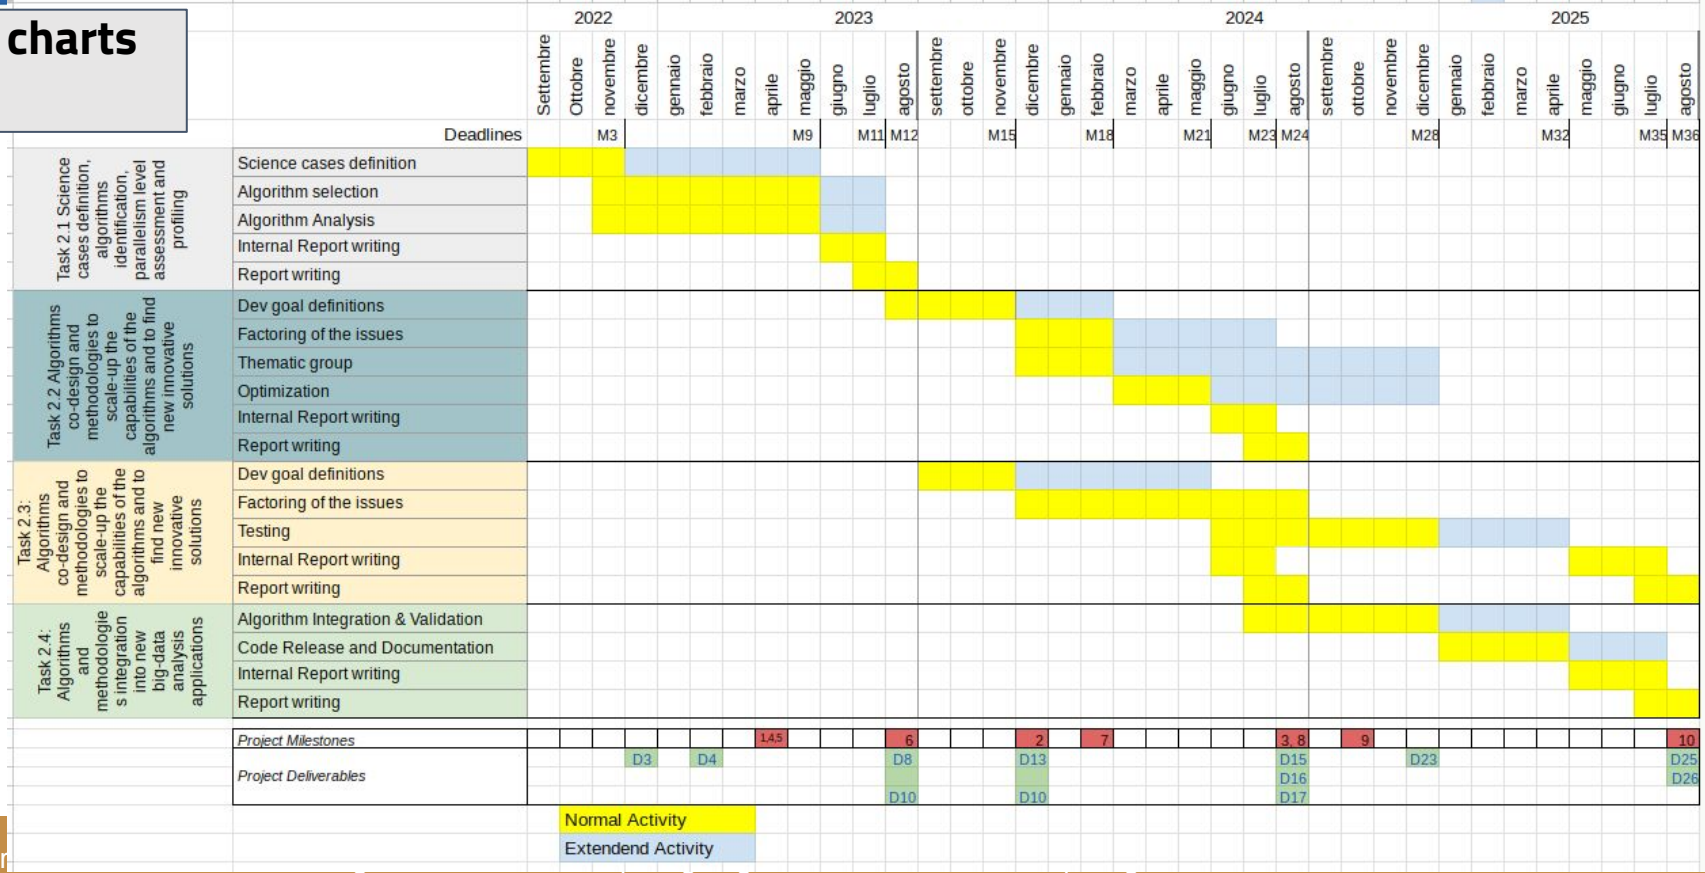

WP1 and WP2 - Deliverables

Internal deliverables

Within M6: Internal Note (tabular form) listing the scientific cases studied with contributor and developer list

Within M11:Internal Note (tabular form) with codes profiling and development goals

Within M23: Internal Note (tabular form) with bottleneck identification for each studied codes, debug procedure and implemented strategies

Within M35: final Internal Note (tabular form) with **codes profiling, development goals, bottleneck identification, debug procedure, implemented strategies and final results (publications, libraries, services deployed, performance gain, energy impact)**

external deliverables

M12: yearly status report

M24: yearly status report

M36: final report

Activities shared between WP1 and WP2

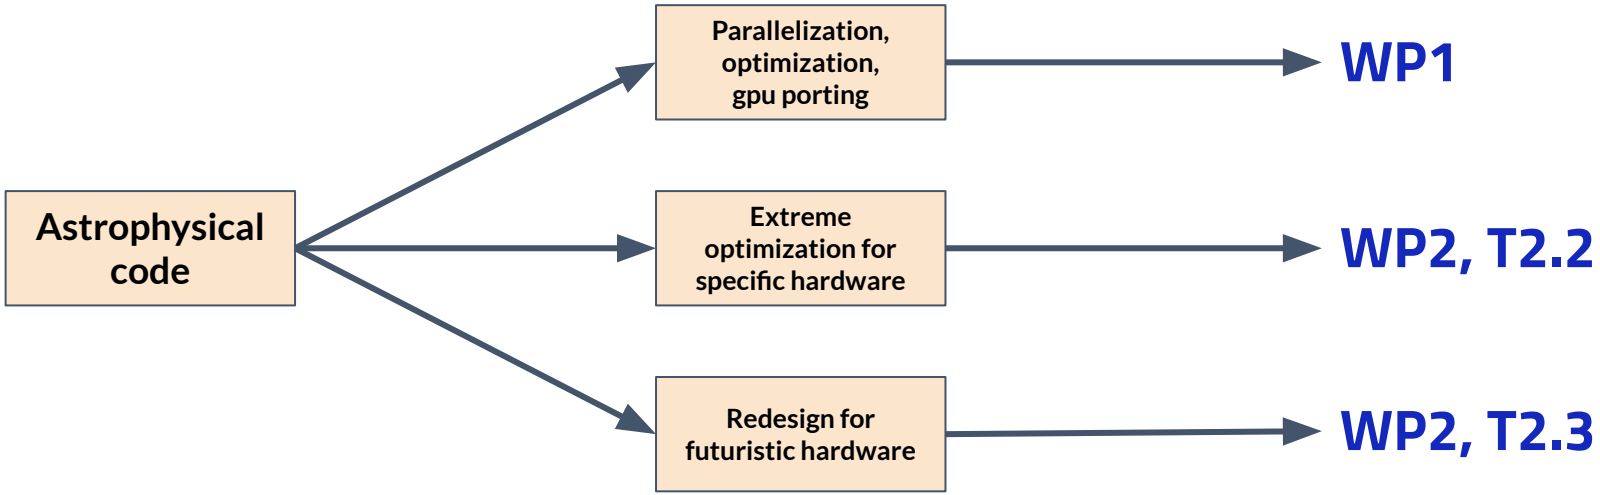

From the slides of Rome kick-off

WP1 and WP2 should be considered as parts of the same work:

- re-engineering a code for better performance is WP1,
- extreme optimization on a part of a code is WP2.

By december 2022 we need a **draft workplan**, that details WG structure, deliverables, milestones.
Deliverables **must be delivered** to have access to the next chunk of funds, so we need to be very careful in specifying what we are going to do. We need to prioritise our plans, e.g.:

- **CERTI**: algorithms that are critical to the success of the project,
- **PROBABILI**: algorithms that will plausibly be delivered,
- **POSSIBILI**: algorithms that we wish to improve, but success depends on extra resources.

Task 2.1: with WP1 we are surveying our participants to decide which codes/algorithms we select for the three priority classes, splitting the work into WP1 and WP2 (also a-posteriori). We can use open calls to

People joining the effort will be divided into two classes:

- **developers**,
- **consultants**.

Once algorithms are selected, we need to provide **use cases to WP4** for the DMP, and receive recommendations for good practices of data management

Towards a (final) list of codes !

List of codes

- 21 codes !
- All of the participants have filled the form
- The progress on activities is updated directly in the form
- **/!\ July 15th: Deadline for all the Spoke3 participants to write the report activities**
- Instructions on how to fill the table are provided in the same file

The idea is to keep the software table updated, adding columns with all informations relevant for the year and final reports

codes profiling, development goals, bottleneck identification, debug procedure, implemented strategies and final results

- Bisogna dire alle persone di essere specifici rispetto alle attivita' pianificate per codice nel contesto HPC e centro nazionale!!

	Partner istituzionale	Codice	Classe	Attivita` pianificate per il codice	WG del codice	Referente scientifico	Personale da assumere su XC	FTE (PM totali)	Altri developers	FTE (PM totali)	Consultants	FTE (PM totali)
13	UniTS	OpenGadget	CERTO	WP1: porting su architetture di tipo GPU; sviluppo di modelli sotto-griglia.	2 - Lagrangian Codes (WP1, WP2)	Stefano Borgani	RTDA + PhD	36			Stefano Borgani	3
14				WP2: ottimizzazione spinta del tree, neighbour search, Particle-Mesh							Matteo Viel	4
15		Pinocchio	CERTO	WP1: porting su GPU delle FFT e del calcolo dei tempi di collasso.	2 - Lagrangian Codes (WP1, WP2)	Pierluigi Monaco	RTDA + PhD	36			Pierluigi Monaco	3
16				WP2: sviluppo del codice Particle-Mesh								
17	UNITO	PLUTO										
18												
19	INFN	Cosmica	POSSIBILE	WP1: porting su GPU dell'algoritmo monte carlo e ottimizzazione	9 - Semi-numerical codes (WP1, WP2)	Stefano Della Torre	PhD	12			Stefano Della Torre	3
20							tecnologo FE	9			Massimiliano Lattanzi	1.15
21		oltzmann (e.g. can	CERTO	Porting su GPU e ottimizzazione	9 - Semi-numerical codes (WP1, WP2)	assimiliano Lattar					Martina Gerbino	2
22												
23							tecnologo FE	6			Massimiliano Lattanzi	1.15
24		istematiche strum	PROBABILE	viluppo e ottimizzazione del codi	9 - Semi-numerical codes (WP1, WP2)	assimiliano Lattar					Martina Gerbino	2
25											Giovanni Signorelli	1
26	INAF	HPC Imaging	CERTO	Parallelizzazione e ottimizzazione di un codice per realizzare l'imaging di grandi moli di dati raccolti da LOFAR e altri radiotelescopi.	8 - Data-reduction & imaging (WP1, WP2)	Gianfranco Brunetti	PhD (Emanuele de Rubeis)	30	Claudio Gheller	2		

=> Internal deliverable

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1	- Bisogna dire alle persone di essere specifici rispetto alle attivita' pianificate per codice nel contesto HPC e centro nazionale!!															
2		**Partner istituzionale**	**Codice**	**Classe**	**Attivita` pianificate per il codice**	**WG del codice**	**Referente scientifico**	**Personale da assumere su XC**	**FTE (PM totali)**	**Altri developers**	**FTE (PM totali)**	**Consultants**	**FTE (PM totali)**			
13		UniTS	OpenGadget	CERTO	WP1: porting su architetture di tipo GPU; sviluppo di modelli sotto-griglia.	2 - Lagrangian Codes (WP1, WP2)	Stefano Borgani	RTDA + PhD	36			Stefano Borgani	3			
14					WP2: ottimizzazione spinta del tree, neighbour search, Particle-Mesh							Matteo Viel	4			
15			Pinocchio	CERTO	WP1: porting su GPU delle FFT e del calcolo dei tempi di collasso.	2 - Lagrangian Codes (WP1, WP2)	Pierluigi Monaco	RTDA + PhD	36			Pierluigi Monaco	3			
16					WP2: sviluppo del codice Particle-Mesh											
17		UNITO	PLUTO													
18																
19		INFN	Cosmica	POSSIBILE	WP1: porting su GPU dell'algoritmo monte carlo e ottimizzazione	9 - Semi-numerical codes (WP1, WP2)	Stefano Della Torre	PhD	12			Stefano Della Torre	3			
20								tecnologo FE	9			Massimiliano Lattanzi	1.15			
21			oltzmann (e.g. can	CERTO	Porting su GPU e ottimizzazione	9 - Semi-numerical codes (WP1, WP2)	assimiliano Lattar					Martina Gerbino	2			
22																
23								tecnologo FE	6			Massimiliano Lattanzi	1.15			
24			istematiche strum	PROBABILE	viluppo e ottimizzazione del codi	9 - Semi-numerical codes (WP1, WP2)	assimiliano Lattar					Martina Gerbino	2			
25												Giovanni Signorelli	1			
26		INAF	HPC Imaging	CERTO	Parallelizzazione e ottimizzazione di un codice per realizzare l'imaging di grandi moli di dati raccolti da LOFAR e altri radiotelescopi.	8 - Data-reduction & imaging (WP1, WP2)	Gianfranco Brunetti	PhD (Emanuele de Rubeis)	30	Claudio Gheller	2					

Partner	Participating Personnel	Hired Personnel	#codes proposed	#CERTI
UniTS	Stefano Borgani, Pierluigi Monaco	RTDA: Milena Valentini; PhD: Alice Damiano, Luca Cappelli (Spoke 10), Marius Lepinzan (Spoke 1)	2	2
INAF	Matteo Bachetti, Aldo Bonomo, Gianfranco Brunetti, Francesco Calura, Claudio Gheller, Giacomo Mulas, Diego Turrini	(11 TD da qui a fine anno)	6	5
SISSA	Andrea Lapi, Mario Spera, Matteo Viel	AdR: Tommaso Ronconi, Chiara Moretti, Elena Sarpa (da agosto)	3	2
UniTO	Andrea Mignone	AdR: Loic Chantry; PhD: Alessio Suriano	1	1
UniCT	Giuseppe Puglisi, Loreto Di Donato, Alessandro Lanzafame	(1 PhD)	1	1
INFN	Stefano Della Torre, Martina Gerbino, Massimiliano Lattanzi	TD: Paolo Campeti (+1 PhD)	3	1
SNS	Andrea Pallottini, Andrei Mesinger, Julien Bloino	(6/8 PhD, da ottobre)	3	3
RomaTOV	Herve' Bourdin, Pasquale Mazzotta, Marina Migliaccio, Francesco Tombesi, NIcola Vittorio	TD: Marco Faltelli, Simone Ferretti; PhD: Avinash Anand	2	2

Positions opened on ICSC

Hired

About to Hire

Not yet hiring

SISSA

1. Galapy - AdR
2. HI-emulator - AdR

UniTS

1. OpenGadget - RTDA + PhD
2. Pinocchio - PhD (Spoke 1)

INAF

1. HPC Imaging - PhD
2. Stingray - TD
3. NP_TMCode - TD
4. Mercury-Arxes - TD
5. GUIBRUSH® - TD
6. RAMSES - TD

INFN

1. COSMICA - PhD
2. CMB related codes - TD

SNS

1. RAMSES - PhD
2. 21cmFAST- PhD
3. Gaussian - PhD

UniTO

1. PLUTO - AdR + PhD

UniToV

1. Litebird_sim - PhD
2. Sparse representations-2TD

Thematic groups proposed for Spoke3

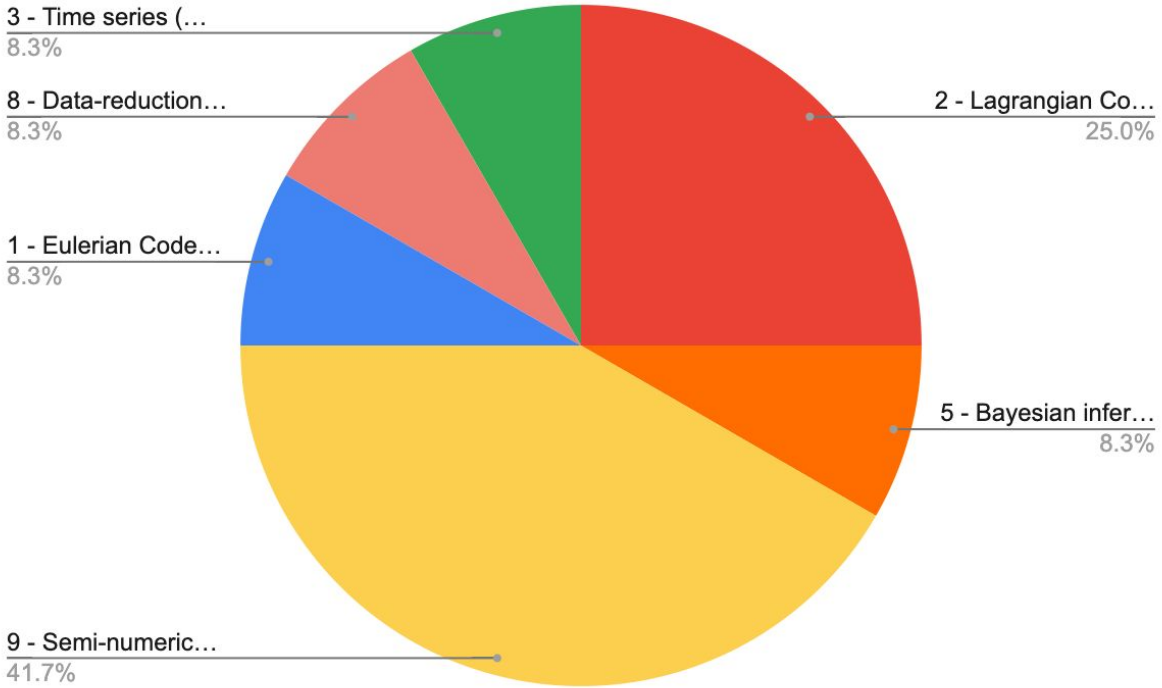

Thematic groups (TG)

WP1/2 Codes vs. TGs

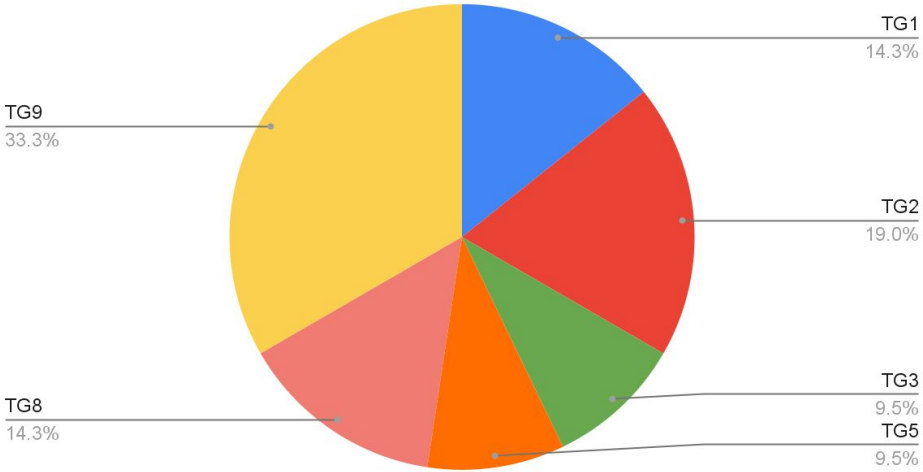

WP1/2 PMs vs. TGs

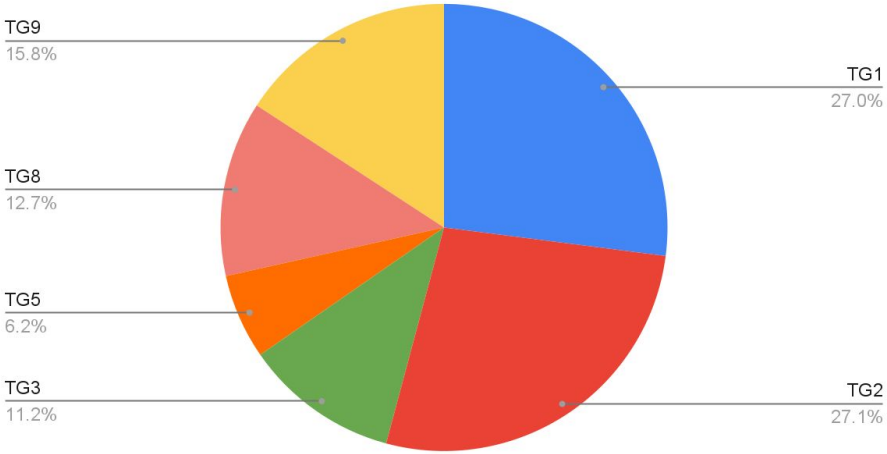

Code tracking

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1	Milestone parts							M3							
2	**Partner istituzionale**	**Codice**	**Classe**	**Attivita` pianificate per il codice**	**WP 1**	**WP2**	**WG del codice**	**Referente scientifico**	**Personale da assumere su XC**	**FTE (PM totali)**	**Altri developers**	**FTE (PM totali)**	**Consultants**	**FTE (PM totali)**	
3	SISSA	ISTEDDAS	CERTO	Implementazione MPI multi nodo, multi-gpu, accoppiamento codice evoluzione stellare, ottimizzazione regolarizzazione algoritmica multithread			2 - Lagrangian Codes (WP1, WP2)	Mario Spera	---	0	Mario Spera	6	Alessandro Bressan	6	
4											Mattia Mencagli	0			
5											Cecilia Sgalletta	0			
6						x					Alessandro Trani	0			
7											Cristiano Ugolini	0			
8											Francesco Gabrielli	0			
9											Ugo N. Di Carlo	0			
10		GALAPY	CERTO	implementazione auto-differenziazione likelihood per metodi bayesiani basati su gradiente, investigazione set-up multi-processore in memoria condivisa per parallelizzazione del modello, inferenza diretta delle posterior tramite Machine Learning			5 - Bayesian inference(WP2 , WP3)	Andrea Lapi	AdR	24	Tommaso Ronconi		Andrea Lapi	9	
11					x						Andrea Lapi		Alessandro Bressan	3	
12											Martina Torsello (Ph.D.)		Mario Spera	3	
13		HI-emulator	PROBABILE	Implementazione MPI del codice, accoppiamento del codice a modelli di foregrounds incluso Principal Component Analysis per foreground removal, implementazione di Gaussian Processes e Machine Learning NN per l'interpolazione del codice su simulazioni idrodinamiche inclusi modelli beyond LCDM, Interfaccia con Bayesian statistics MCMC estimator per estrarre cosmo, astro a parametri di nuisance.	x	x	9 - Semi-numerical codes (WP1, WP2)	Matteo Viel	AdR	24	Matteo Viel	2	Matteo Viel	2	
14	UniTS	OpenGadget	CERTO	WP1: porting su architetture di tipo GPU; sviluppo di modelli sotto-griglia.	x		2 - Lagrangian Codes (WP1, WP2)	Stefano Borgani	Milena Valentini (RTDA) Alice Damiano (PhD)	72	Antonio Ragagnin; Klaus Dolag; Luca Tornatore; Giuseppe Murante	0	Stefano Borgani	3	
15				WP2: ottimizzazione spinta del tree, neighbour search, Particle-Mesh		x					Luca Tornatore	0	Matteo Viel	4	
16		Pinocchio	CERTO	WP1: porting su GPU delle FFT e del calcolo dei tempi di collasso.	x		2 - Lagrangian Codes (WP1, WP2)	Pierluigi Monaco	RTDA + PhD	36			Pierluigi Monaco	3	

First science results: flash talks

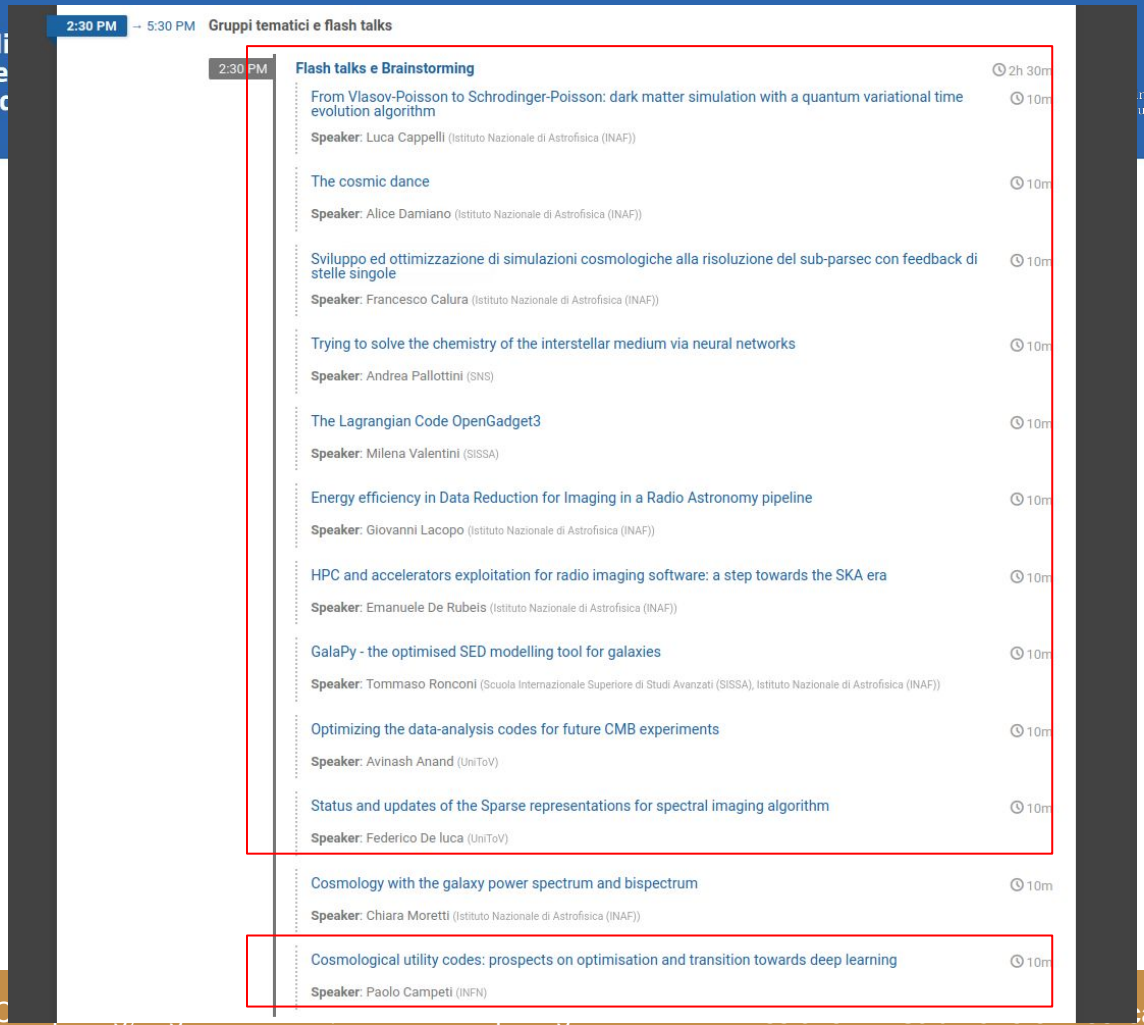

2:30 PM → 5:30 PM **Gruppi tematici e flash talks**

2:30 PM **Flash talks e Brainstorming** 2h 30m

From Vlasov-Poisson to Schrodinger-Poisson: dark matter simulation with a quantum variational time evolution algorithm 10m
Speaker: Luca Cappelli (Istituto Nazionale di Astrofisica (INAF))

The cosmic dance 10m
Speaker: Alice Damiano (Istituto Nazionale di Astrofisica (INAF))

Sviluppo ed ottimizzazione di simulazioni cosmologiche alla risoluzione del sub-parsec con feedback di stelle singole 10m
Speaker: Francesco Calura (Istituto Nazionale di Astrofisica (INAF))

Trying to solve the chemistry of the interstellar medium via neural networks 10m
Speaker: Andrea Pallottini (SNS)

The Lagrangian Code OpenGadget3 10m
Speaker: Milena Valentini (SISSA)

Energy efficiency in Data Reduction for Imaging in a Radio Astronomy pipeline 10m
Speaker: Giovanni Lacopo (Istituto Nazionale di Astrofisica (INAF))

HPC and accelerators exploitation for radio imaging software: a step towards the SKA era 10m
Speaker: Emanuele De Rubeis (Istituto Nazionale di Astrofisica (INAF))

GalaPy - the optimised SED modelling tool for galaxies 10m
Speaker: Tommaso Ronconi (Scuola Internazionale Superiore di Studi Avanzati (SISSA), Istituto Nazionale di Astrofisica (INAF))

Optimizing the data-analysis codes for future CMB experiments 10m
Speaker: Avinash Anand (UniToV)

Status and updates of the Sparse representations for spectral imaging algorithm 10m
Speaker: Federico De luca (UniToV)

Cosmology with the galaxy power spectrum and bispectrum 10m
Speaker: Chiara Moretti (Istituto Nazionale di Astrofisica (INAF))

Cosmological utility codes: prospects on optimisation and transition towards deep learning 10m
Speaker: Paolo Campeti (INFN)

Bandi a cascata

1. Enabling on innovative hardware and software (RISC-V processors and RISC-V-based computing accelerators) for the astrophysical characterization of gravitational-wave transients (WP2)
2. Parallelization and testing of simulation codes for cosmological instrumentation and optimization (WP1, 2)
3. Porting and optimization of codes for cosmological simulations on architectures with accelerators, primarily GPUs, but also FPGA. (WP1,2)
4. High Performance I/O: experimentation and optimization of I/O strategies adopting innovative High Performance storage devices, parallel filesystems or object storage solutions, I/O libraries (in particular HDF5 and ADIOS) (WP1,2)
5. The objective of this Open Call is to accelerate the performance of the XMAP, WINE and TEPID . The project will involve identifying the computationally intensive sections of the codes and optimizing them for execution on GPUs. (WP1,2)
6. Development of (multi)GPU-accelerated simulations of physical models based on SDE Monte Carlo algorithms (WP1,2)
7. Innovative Hybrid Accelerated Architectures (WP1,2)

What will the Hub see of our work (in the deliverables)?

For each code:

- a very brief (10 lines?) project definition,
- a progress metric (percentage?),
- a link to a github (?) repository,
- a link to a redmine (?) page ,
- a (two-page?) report with figures demonstrating results,
- published papers.

Critical points

- **Hardware:** we need to access HPC resources to develop our codes, **ASAP**
- **Hardware:** we need precise specs to emulate the system that we will use, **ASAP**
- **Hardware and community:** it will be hard to keep momentum and build a community if computing resources must be acquired on best-effort basis
- **Science and community:** a meeting more focused on science (including science from industrial partners) would be good, later in October, with sessions based on thematic groups
- **Key science projects**

Spoke Leader INAF
Co-Leader INFN

Soggetti Pubblici Affiliati	
Università Roma Tor Vergata	
Università di Trieste	
Università di Torino	
Università di Catania	
Scuola Normale Superiore - Pisa	
Sissa - Trieste	

Imprese Aderenti allo Spoke	
Intesa SanPaolo	
UnipolSai	
Sogei	
IFAB	
Leonardo	