

A Systematic Mapping Study of Italian Research on Workflows

Marco Aldinucci (UniTO)
Elena Baralis (PoliTO)
Valeria Cardellini (Tor Vergata)
Iacopo Colonnelli (UniTO)
Marco Danelutto (UniPI)
Sergio Decherchi (IIT)

Giuseppe Di Modica (UniBO)
Luca Ferrucci (UniPI)
Marco Gribaudo (PoliMI)
Francesco Iannone (ENEA)
Marco Lapegna (Federico II)
Doriana Medić (UniTO)
Giuseppa Muscianisi (CINECA)

Francesca Righetti (UniPI)
Eva Sciacca (INAF)
Nicola Tonellotto (UniPI)
Mauro Tortonesi (UniFE)
Paolo Trunfio (UniCal)
Tullio Vardanega (UniPA)

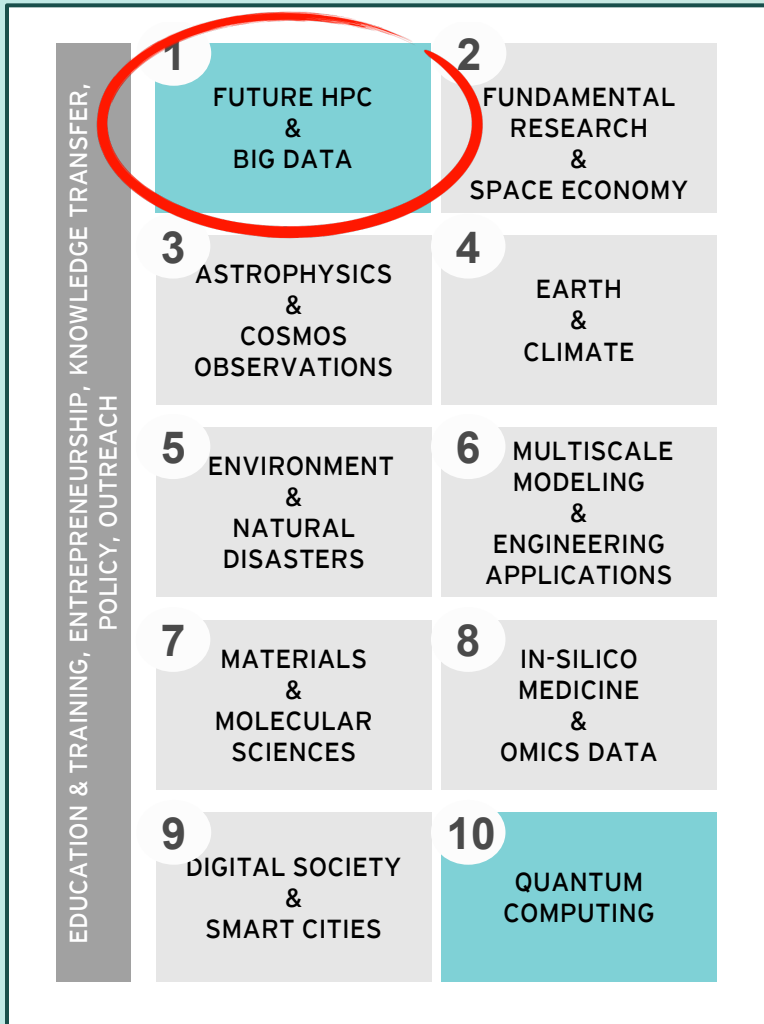
Research questions

- This work aims to answer **three research questions** regarding current and future directions for research on workflow systems:
 - Which are the **main research directions** for Workflow Management Systems (WMSs) in the Computing Continuum?
 - Which research directions are **widespread** in the scientific community?
 - Which research directions **address a critical need** for modern scientific applications?

Systematic mapping study

- We conducted a systematic mapping study of tools targeting **large-scale scientific workflows** and their execution in the **Computing Continuum**
- This report analyzes **25 tools from 9 Italian research institutions** collected among ICSC Spoke 1 partners in the context of FL3
- To answer the last question, the report also collects requirements from **10 scientific applications** developed and maintained by **11 ICSC scientific partners**

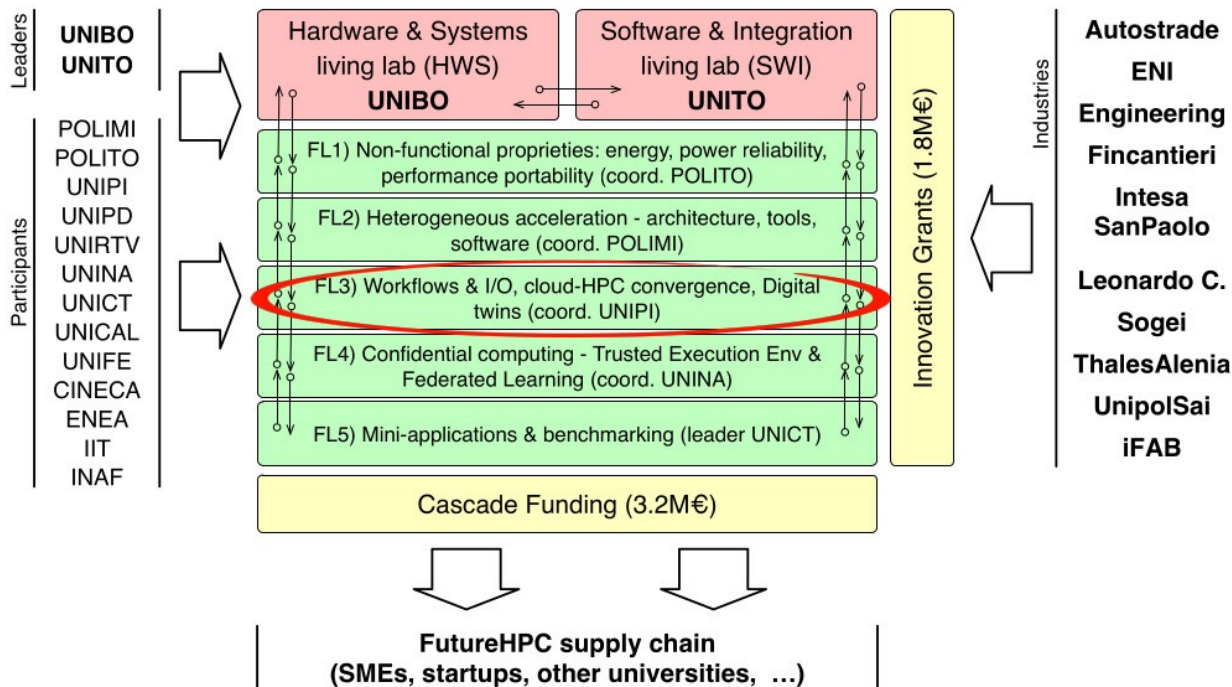
National Center on HPC, BigData and Quantum Computing



- **320M€ funding** (3 years - Sept 1st, 2022) by the NextGenerationEU programme
- **25 research institutes** and **24 large industries** involved
- Organized according to a **Hub&Spoke model**, with ten thematic spokes and one infrastructure spoke

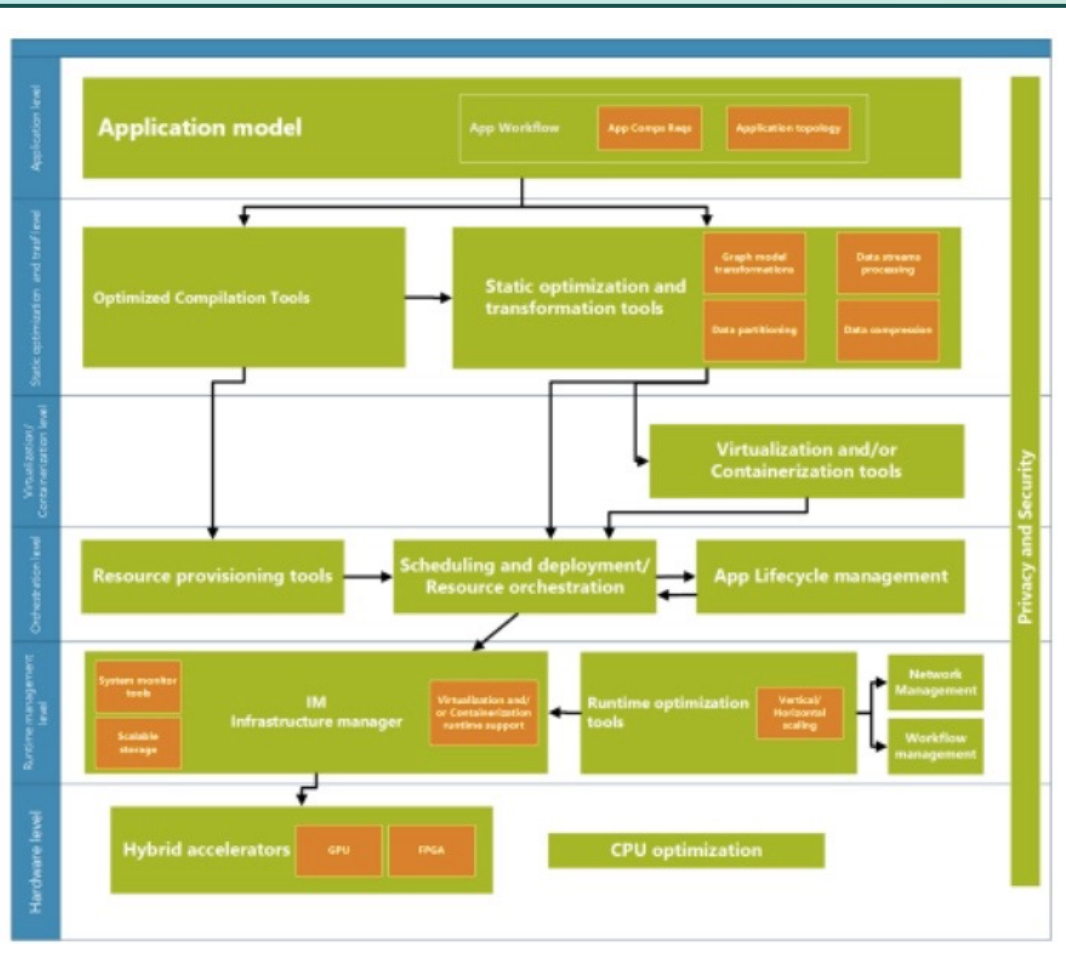
Spoke 1 – FutureHPC & BigData

Big picture of Spoke 1 - FutureHPC & BigData
(financial envelope 21,5M€)



- The objective of Spoke 1 is the **creation of new labs** as an integral part of a National federated centre on a global level with skills aimed at **hardware and software co-planning**
- Spoke 1 is organized into **five scientific flagships (FLs)** and two living labs

Spoke 1 FL3 – Workflows & I/O, cloud-HPC convergence, digital twins



- FL3 is an Italian umbrella community contributing, developing, and maintaining an **integrated set of SW tools**
- FL3 aim at implementing a **unifying software stack for HPC-oriented workflows**, integrating tools and methodologies for cloud-HPC and high-performance storage and IO

Limitations

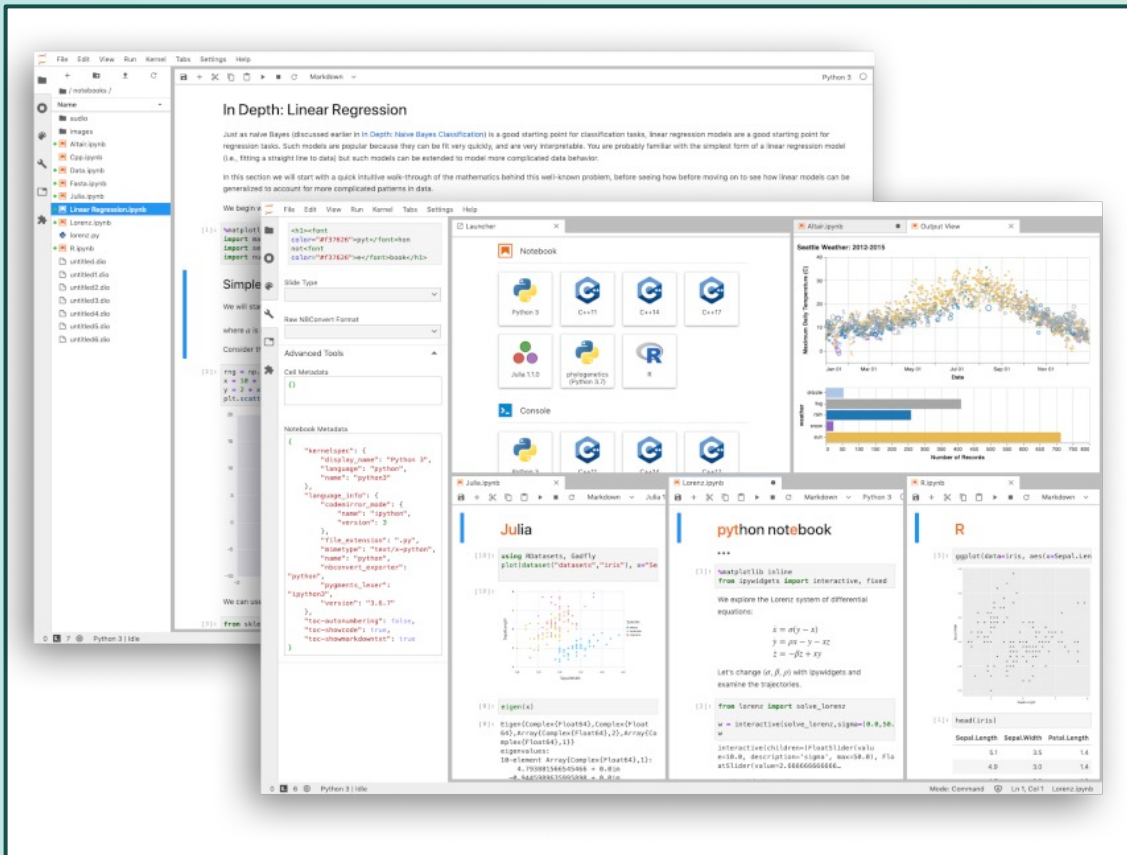
- The study aims to discuss research directions in scientific workflows using the ICSC ecosystem as a **statistical sample** of international research on workflows
- Since it only considers the Italian ICSC ecosystem, it cannot be considered a comprehensive survey of the state-of-the-art workflows at the national or international level

Which are the main research directions for WMSs in the Computing Continuum?

Five classes of tools have been identified: interactive computing, orchestration, energy efficiency, performance portability, and Big Data management

Interactive computing	Orchestration	Energy efficiency	Performance portability	Big Data management
BookedSlurm	TORCH [72]	PESOS [16]	FastFlow [5]	ParSoDA [10]
ICS [19]	INDIGO [23]	Lapegna et al. [41]	Nethuns [12]	MALAGA
Jupyter Workflow [20]	Liqo [37]	De Lucia et al. [46]	INSANE [64]	aMLLibrary [33]
	StreamFlow [21]		CAPIO [47]	WindFlow [49]
	SPF [53]		BLEST-ML [14]	CHD [18]
	BDMaaS+[17]		MLIR [43]	Mingotti et al. [51]
	MoveQUIC [59]			

Interactive computing

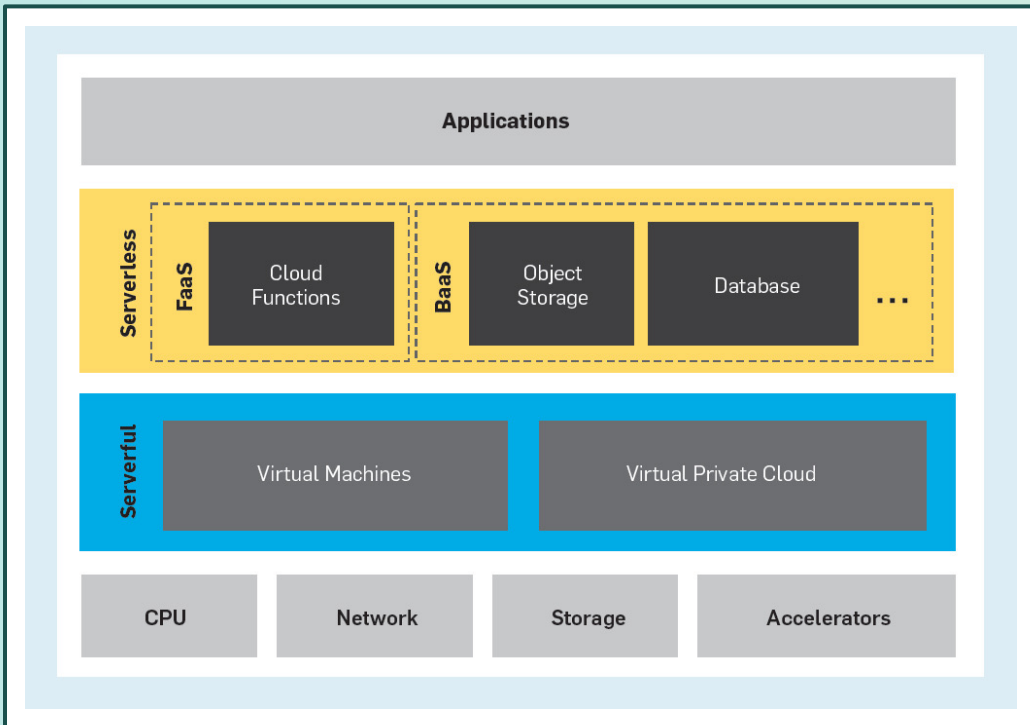


- **Jupyter Notebooks** are a promising technology enabling interactive workflows in HPC infrastructures
- Jupyter supports interactive computing in **several languages** can be offered as a service on Cloud platforms (JupyterHub, Google Colaboratory)

Interactive computing

- Still, enabling Jupyter-based workflows as a service **on HPC facilities** poses three main challenges:
 - Interactive computing requires **on-demand resource provisioning**, while HPC facilities offer batched executions through queue managers
 - The ZeroMQ transport layer requires a **bidirectional TCP connection** between the publicly exposed frontend and the air-gapped worker nodes
 - The standard execution flow of Notebook cells is **purely sequential**, preventing users from modelling applications as workflow graphs

Orchestration



- The **Function-as-a-Service (FaaS)** paradigm recently percolated to Edge and Fog environments, pushing towards a Cloud-Edge Continuum
- It gained traction in scientific workflow management, either as a **primary execution infrastructure** or **combined with HPC facilities** in hybrid settings

Samuel Kounev, Nikolas Herbst, Cristina L. Abad, Alexandru Iosup, Ian T. Foster, Prashant Shenoy, Omer F. Rana, Andrew A. Chien:
Serverless Computing: What It Is, and What It Is Not? Communications of the ACM 66(9): 80-92 (2023)



Orchestration

- FaaS moves all the deployment and life-cycle management aspects **to the provider (WMS) side**. When moving to the Continuum, additional aspects emerge:
 - Orchestration algorithms and tools must be capable of guaranteeing **near real-time responses** for function invocations independent of the underlying deployment infrastructure
 - **Placement decisions** become vital as invocation performance rises from the combination of available computing power and near-data processing
 - **Efficient migration strategies** are crucial whenever data sources expose high dynamicity in their generation rate

Energy efficiency




- Energy consumption is a key indicator in **the whole spectrum** of the Computing Continuum
- On the HPC side, there is growing attention on **measuring and reducing the carbon footprint** of computational research, even if initiatives promoting sustainable HPC have existed for several years (e.g., the Green500 List)

Energy efficiency

- With the entire Computing Continuum available, ensuring energy efficiency is a challenging task:
 - A viable solution is to adopt **energy-aware placement algorithms**, which try to minimize the carbon footprint of workload executions without violating QoS requirements
 - Another possibility is to **move computations on Edge sensors** whenever possible. Besides relying on low-power hardware, this strategy also removes data transfers, saving additional energy
 - Efficiently exploiting this class of devices requires **resource-constrained algorithms and implementations**

Performance portability



- Performance portability is the sum of two opposing forces: **gaining portability**, which requires high-level abstractions, and **maximising performance**, which requires a deep knowledge of the target architecture
- The **heterogeneity** of hardware accelerators and the increasing **modularity** of modern scientific applications made performance portability libraries crucial for any large-scale scientific application that targets production usage

Performance portability

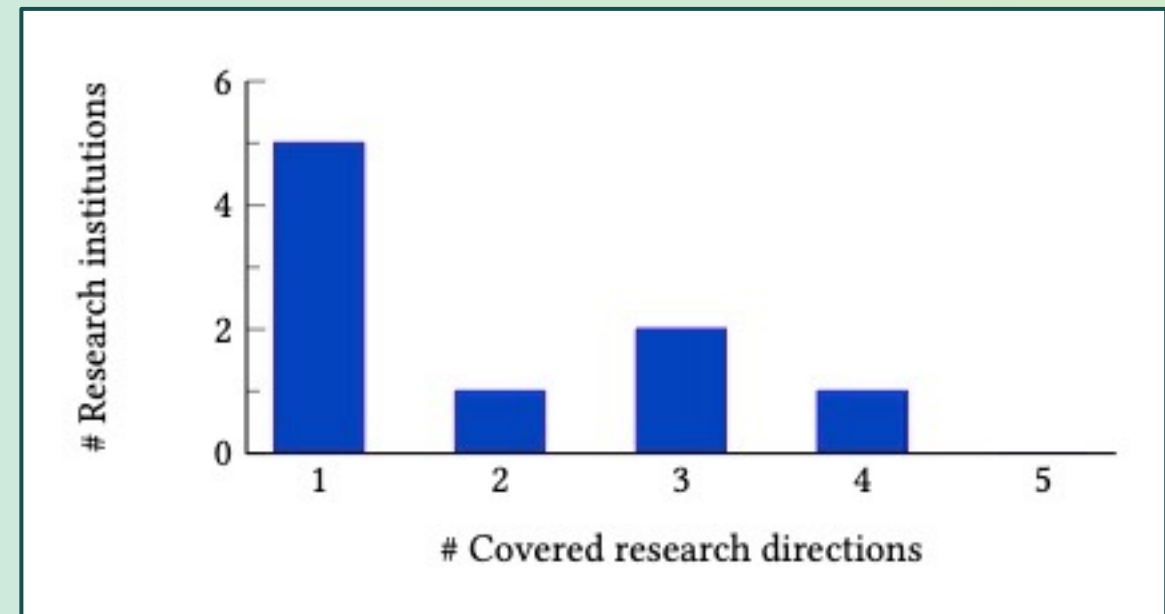
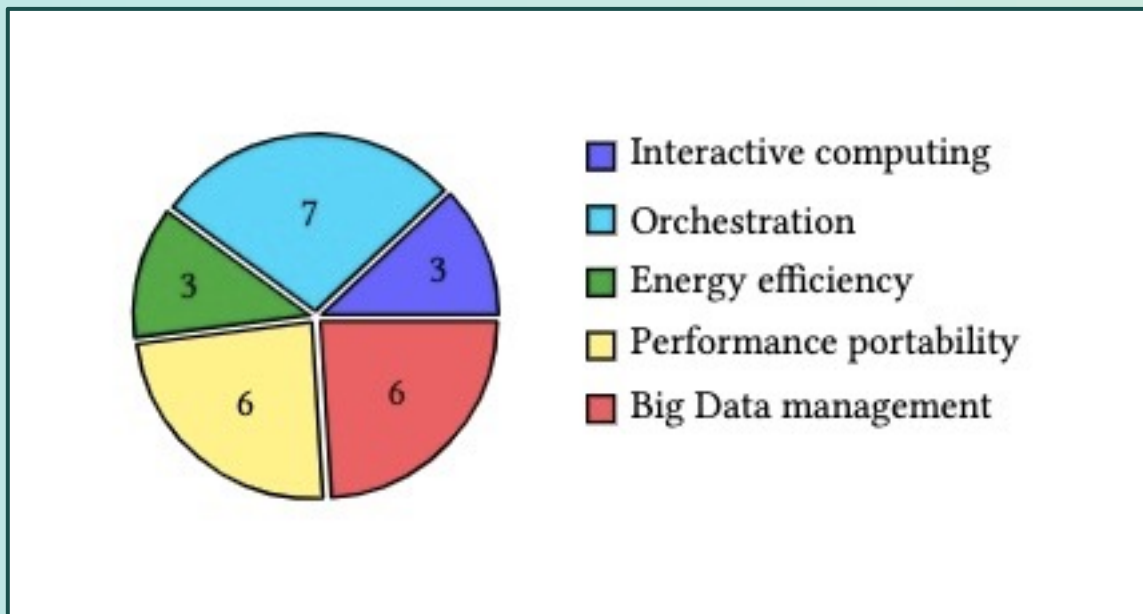
- A performance portability library is composed of **two key elements**:
 - A **programming model** provides developers with abstractions between the application and one or more low-level resources, e.g., network, memory, storage, or data structures.
 - Each abstraction is then translated into an **efficient implementation** optimized for a specific target execution environment, e.g., a high-end network with smart NICs, a high-bandwidth burst buffer, or a distributed
 - Depending on the library, this translation can happen at **compile time** or **runtime**

BigData management

- The Big Data domain requires **highly-expressive dataflow operators** to support batch, micro-batch, and streaming execution models, and **advanced data structures** to enable distributed in-memory computing
- Moving Big Data analytics and Deep Learning pipelines to the Continuum poses new challenges:
 - Domain experts need tools and algorithms to support **pluggable data processing operators** for diverse data types, from images to graphs to geospatial information
 - At runtime, such operators should coordinate their execution from Edge to Cloud **in a transparent way**
 - They should be able to target **multi-core and distributed architectures** and to exploit **heterogeneous hardware devices** when available

Which research directions are widespread in the scientific community?

- The effort is **quite balanced** among the different research directions, with no single, predominant research effort
- More than half of the involved institutions cover **a single research topic**, and no institutions span the whole set of identified directions



Which research directions address a critical need for modern scientific applications?

- **Ten scientific applications** have been collected from **eleven ICSC partners**
- Application providers were asked to identify, among the collected tools, those that they deemed valuable to improve the current status of their workload, with a specific focus on **workflow execution in a Computing Continuum environment**

Which research directions address a critical need for modern scientific applications?

		3.1	3.2	3.3	3.4	3.5	3.6	3.7	3.8	3.9	3.10
Interactive computing	BookedSlurm										
	ICS		✓							✓	
	Jupyter Workflow		✓					✓			
Orchestration	TORCH										
	INDIGO				✓				✓		
	Liqo				✓				✓		
	StreamFlow		✓	✓							✓
	SPF										
	BDMaaS+ MoveQUIC				✓	✓			✓	✓	
Energy efficiency	PESOS					✓					
	Lapegna et al. De Lucia et al.										
Performance portability	FastFlow	✓									
	Nethuns		✓				✓				
	INSANE										
	CAPIO		✓				✓				
	BLEST-ML										
	MLIR										✓
Big Data management	ParSoDA	✓								✓	
	MALAGA										
	aMLLibrary							✓		✓	
	WindFlow	✓									
	CHD										
	Mingotti et al.							✓			

3.1: compression of petascale collections of textual and source-code files

3.2: Astrophysics data analysis and visualization

3.3: Genomic variant calling pipeline

3.4: Edge-Cloud Continuum federation infrastructure

3.5: Serverledge: QoS-Aware FaaS in the Edge-Cloud Continuum

3.6: Improving I/O phases in computational modelling of Galaxy Formation


3.7: WorldDynamics.jl

3.8: Optimized deployment of Cloud-native applications in the Cloud Continuum

3.9: Anomalous subgroup characterization with DivExplorer

3.10: Compilation flow and deployment strategy targeting HPC RISC-V accelerators

Which research directions address a critical need for modern scientific applications?



- Developing a **solid orchestration infrastructure** targeting Computing Continuum is critical
- Interactive computing, performance portability, and Big Data management have aroused **significant interest**
- 2/3 of the collected tools for energy efficiency target **specific algorithms**

Wrapping up

- This work collected and analyzed **25 tools and 10 applications** from several Italian research institutions in the context of the ICSC Spoke 1 initiative.
- The tools were clustered into **five research directions**. With the exception of **energy efficiency**, these directions overlap with recent literature's beliefs about future directions of scientific workflows
- More than half of the involved institutions cover a **single research topic**. **Collaborative initiatives** are crucial for providing direct links between highly specialized groups and building a holistic research ecosystem
- Application providers highlight a prominent interest in **advanced workflow orchestration** and a still significant interest in all research directions but energy efficiency, partly due to the domain-specific nature of the collected tools

What's next?

The next phases of Spoke 1 will be focused on **implementing** the proposed tool integrations and **testing** them on pipelines from applicative Spokes and industries

Stay tuned!

<https://www.supercomputing-icsc.it/en/icsc-home/>