Advice on INAF Future Engagements in Computational Astrophysics

G. Brunetti (INAF)
G. Bodo (INAF)
K. Dolag (LMU Munich and MPA Garching, Germany)
K. Heitmann (Argonne National Laboratory, USA)
L. Mayer (CTAC, Inst. for Computational Science, Univ. of Zurich, Switzerland)
A. Possenti (INAF)

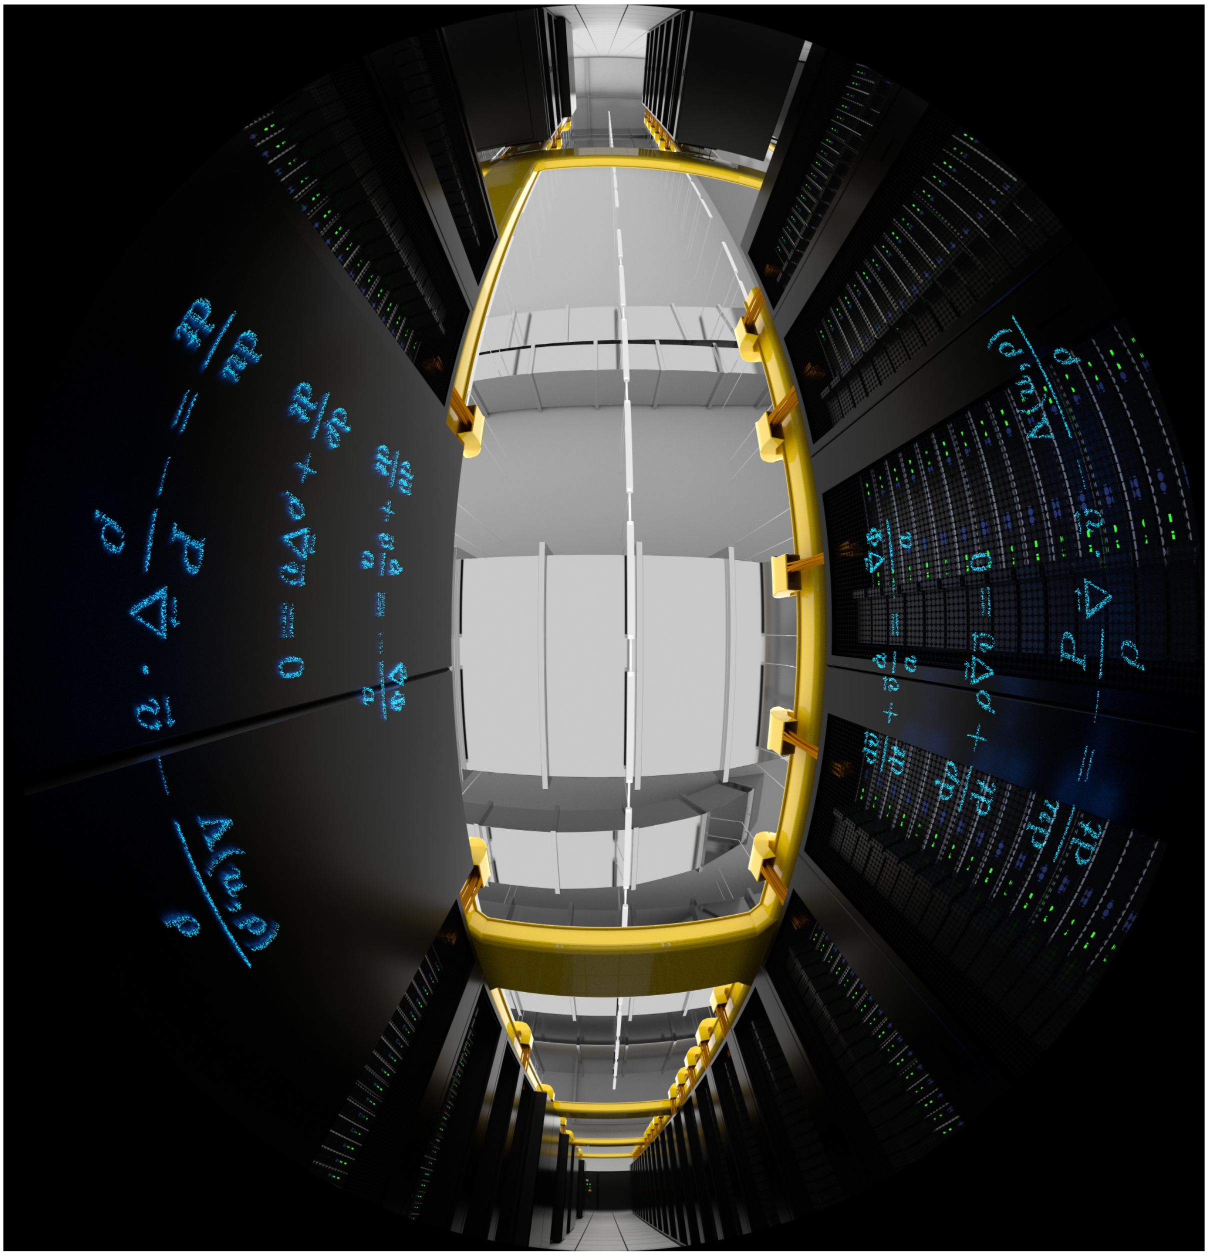

1. Executive Summary

In INAF there is a significant granularity in the distribution of infrastructures for computing and data analysis. This granularity must be preserved for the support of medium-small projects, because it helps the growth of local skills and because in some cases it facilitates fundraising (e.g. EU grants obtained by local researchers, regional funds, etc). However, in the medium-long term, INAF will be deeply involved in large international observational programs that will require the analysis and storage of large amounts of data, and therefore should aim to create and support a community capable of dealing with this new scenario, as well as being competitive in the next generation of numerical simulations (up to simulations on exascale resources).

This requires **coordination at national level and a large investment in a centralized infrastructure** that can grow over time (computing power and storage) and that is **primarily designed to support the large projects** (and programs) in which INAF is involved. These projects include the roadmap towards SKA in the field of radio astronomy, the exploitation of large datasets from ground and space-based surveys in the optical band (and their combination with surveys in other bands also having an Italian leadership), and the use of HPC for theoretical developments, both in the field of cosmology and of astrophysical plasma physics (plasma astrophysics). More generally, a large infrastructure would allow INAF to be more competitive in grasping opportunities in future projects. Specifically, in the medium term (**2022-27**), we suggest an infrastructure of Tier 2-1 scale (i.e. of order 20000-30000 cores, roughly 1 Pflops) with a hybrid architecture (CPU / GPU, medium / fat nodes, 30+ PB of storage), with the possibility of satellite (Tier 3) specialised clusters when necessary. In fact, this infrastructure could be renewed/adjusted (every 3-4 years) on the basis of both scientific considerations and specific needs of INAF. In the long term (**2028+**) part of the infrastructure should evolve into a Tier 1 size infrastructure to match the requirements of a fully deployed SKA RC in which INAF can aim to play a significant role (i.e. of order 3+ Pflops and 70 PB of storage per year).

We also strongly suggest accompanying the infrastructural investment with a **significant investment that can generate a coordinated critical mass of astrophysical researchers** (i.e. collecting 10-15 dedicated FTE) with strong skills in the field of computer science, capable of carrying out frontier research in that field, developing innovative software architectures and porting numerical astrophysical codes on complex (hybrid) infrastructures. This "*computer science division*" is necessary to support large projects (but also medium-small programs on a local scale) and will optimize the pay-back from the investment on the infrastructure. The staff of the software division can be allocated with a level of geographical granularity significantly higher than the HPC and storage infrastructure. This will help in consolidating the current INAF nodes with greater competence in the field of IT, as well as promoting the growth of new nodes.

This large investment (infrastructure, personnel and R&D) and its interplay with large projects and programs with critical IT-needs must be **coordinated and managed by a structure with a strong scientific vision** in order to guarantee both the needs of the large projects for which it was conceived and a clear scientific driver for possibly setting new priorities and a roadmap. We suggest a management model that is anchored to an **executive board**, that *provides* scientific-technological priorities, as well as the coordination and management of infrastructure and activities related to critical IT, a **manager** of the infrastructure(s), and a **coordinator** of the *computer science division*. **Major projects and programs** must have a role in the decision-making flow, possibly being also represented in managerial bodies.

A certain coordination of critical IT activities with local IT activities and services is desirable and should be **included in a global vision**.

2. Introduction

The need for computing power, together with related side facilities, is constantly increasing in modern astrophysics. The need to process a large amount of data of diverse types, within a limited time-slot, is becoming everyday life for astrophysicists as it has been earlier for other physicists and scientists.

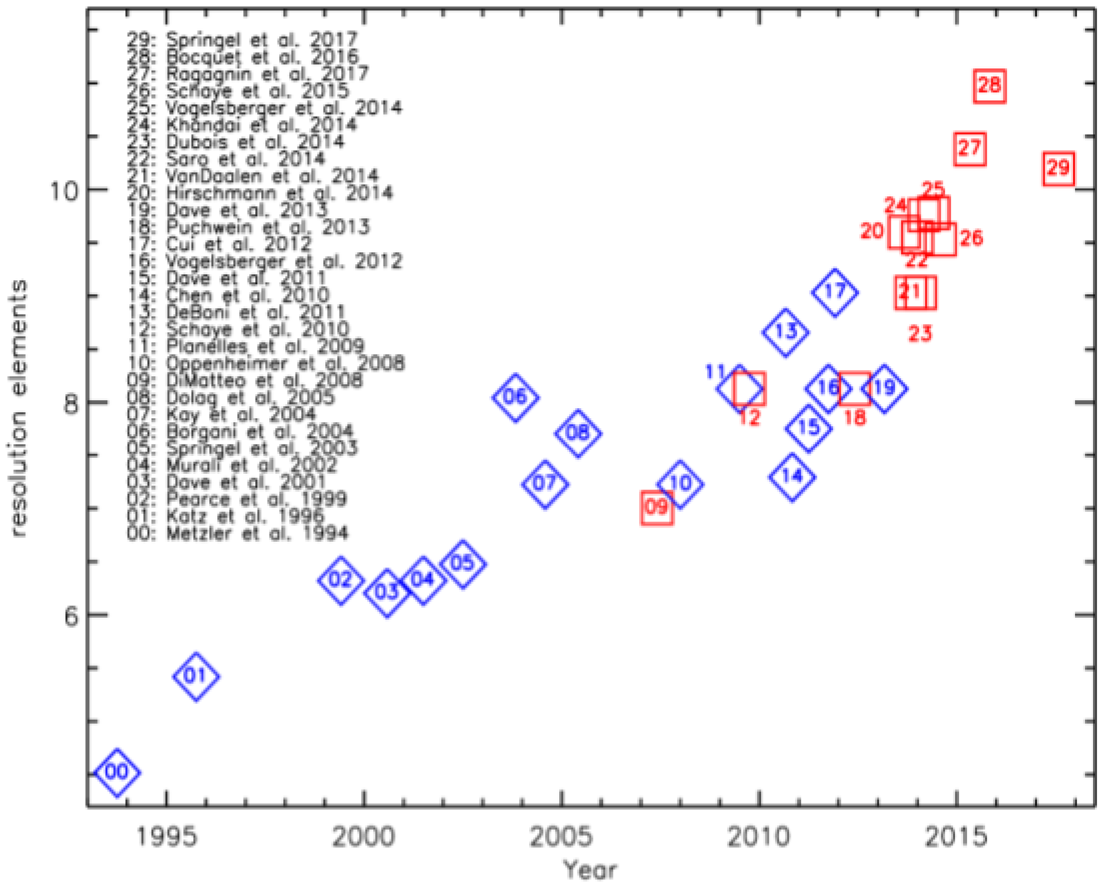

FIGURE 1 : evolution of resolution element in cosmological simulations

Radio telescopes, like ALMA, LOFAR, MeerKAT, and in perspective SKA, Large Cherenkov Telescopes, like ASTRI Mini-Array and in perspective CTA, large optical-NIR facilities, like E-ELT and its suit of instruments, or synoptic survey telescopes and fast transit-search facilities are changing the order of magnitude of the computing and storage facilities required to maximize the scientific output. A clear case is provided by the rapid evolution of the data flow in the field of radio astronomy in the last decade. For example, the pan-European LOFAR radio telescope, of which INAF is a part, generates an unprecedented (for astronomy) data flow of about 200 Gb/s, and the datasets to calibrate and analyze are typically in the range 1-10 TB, i.e. about 100 times larger than the data of previous radio telescopes (e.g. VLA, GMRT).

Another clear case is provided by modern theoretical simulations that indeed may cost several millions of core-hours on large supercomputers, require hundreds of terabytes of memory, and may use petabytes of disk storage. These computational needs are rapidly increasing with time (es., FIGURE 1). Future numerical simulations will have to use innovative codes, which explicitly include the treatment of complex physical and microphysical processes, capable of fully exploiting the high-performance computing power and architecture available in the near future, even with “exa-scale” class infrastructures. These simulations will be crucial in all areas of cosmology and astrophysics to provide theoretical predictions and/or interpret observations with future facilities (e.g. SKA and precursors, Athena, CTA, etc.).

For these reasons INAF is planning a large investment on **critical-IT** in the next years acquiring a **centralized high performance computing facility**, eventually linked to **other distributed smaller scale computing centers**, and possibly external supercomputing centers and commercial providers.
This large investment on critical-IT and the amount of personnel involved in large international projects in INAF require an effective **coordination** and a **managerial model based on a science vision and on a roadmap**.

With this report we have been asked to provide advice on this plan, in particular on the following specific points :

- The general perspectives of the intervention
- The optimization of the overall architecture
- Specific suggestions on the configuration of the infrastructure(s)
- Organization and management of the infrastructure(s)

As a first step, we have analyzed the large international projects in which INAF is involved in the medium to long term, which present obvious challenges in the field of critical IT, and the theoretical activity with numerical simulations by the community. This led to the definition of basic requirements for computational infrastructures operated by INAF.

The second step was to understand how to insert the coordination of these infrastructures within the current management structure of INAF. This led us to outline an overall critical IT investment framework and management model. Although we are not directly called to define an “overall model” of IT coordination and management in INAF, we considered it important to also discuss the relationship between critical IT and activities related to small computational infrastructures managed locally and medium-small IT programs.

Our study is based on the analysis of major INAF projects among those specifically mentioned in the letter of appointment by the DS. Therefore we have not conducted a complete survey of the activities in INAF, so that other large projects and activities in the medium-long term could materialize and require the use of infrastructures for critical IT. For this reason we aimed at defining an investment and management model of critical IT in INAF that is as modular as possible, so that it can be expanded to support further needs over time.

The report is structured as follows: in Sect 3 we summarize the status of INAF computing, in Sect. 4 we discuss the needs of major projects and programs in INAF. In Sect. 5 we explain our general view of the investment and its main goals. In Sect.6 and 7 we discuss the details of the proposed infrastructure(s) and computer science division. In Sect. 8 and 9 we discuss the suggested management model and the roadmap of the investment. Finally our main conclusions are reported in Sect. 10. Details of critical IT needs from SKA roadmap, optical surveys and HPC are reported in Appendix 1,2,3. In Appendix 4 we also report benchmark performance for CPU and GPU that have been adopted in our study to convert cores and GPU into Tflops.

3. Status of INAF computing

In INAF there is a significant granularity in the distribution of infrastructures for scientific computing and data analysis. Currently INAF counts on a distributed configuration of computing clusters (up to 500-1000 cores). Most of these clusters are currently data centers around observational facilities (eg SRT at OACa, radio telescopes at IRA and Cherenkov telescope ASTRI, CTA at OACt, OARm, OAS, VST at OACn), or they support the activity of HPC groups (eg OAPa, OACt, OATs) and laboratories (eg Arcetri for optical turbulence, etc).

IT SERVICES:
an attempt to coordinate local IT resources in INAF at the national level is provided by the activity of the **ICT**, an office within the DS organization **which offers IT services** (e.g., software procurement and management, archives (IA2 based at OATs), IT networks, systems engineering) and **coordinates these services at National level** in INAF also managing agreements and MoU with external IT groups and services (e.g., OpenPOWER, TANGO).
In recent years, ICT has promoted valuable initiatives also offering general **computational/HPC services** (1000+ cores-size) to the entire community (e.g. CHIPP based at OATs and OACt, and in the near future PLEIADI). On larger scales, the use of big computers for theoretical HPC in INAF currently takes place through on-demand access to larger infrastructures (CINECA, also through a service established by the INAF-CINECA MoU, and global commercial suppliers) without having the possibility of an intermediate step on medium-large INAF platforms, e.g. Tier 2.

IT FOR INTERNATIONAL PROJECTS:
A different case is that of the medium-size IT activities that are carried out at INAF within large international projects. The computational needs in the context of these programs derive from well-defined scientific and technological visions. The infrastructures, supported by single (or small networks of) INAF Institutes, are specifically designed to address well-identified scientific objectives. These small infrastructures (1000 core size) constitute aggregates around which astrophysics and technological research is carried out in the several fields within INAF (eg ALMA RC at the IRA, LOFAR-It at the IRA, OATs, OACt, etc.) and have an immediate impact on community growth and productivity. On larger scales, due to the lack of adequate computing facilities in INAF, the support of large programs from space in which INAF is deeply involved (es Euclid ground segment, DPCT Gaia, etc) is delegated to the use of infrastructures that are managed by external bodies (ASI, ALTEC).

4. INAF Major Projects with critical IT challenges

We have analyzed the needs in terms of computational resources and storage of large international projects and programs which present obvious challenges in the field of critical IT. The most relevant projects in which INAF will be involved in the medium-long term are listed in the letter of appointment from the DS. We have analyzed the most demanding programs in that list and derived the basic requirements for the computational infrastructures that are needed in INAF and their organizational model.

There are essentially three groups of major lines facing critical IT needs:

1. **The Roadmap to the SKA:** SKA precursors and pathfinders are making incredible discoveries and charting the roadmap to the SKA. Along this roadmap, the communities involved must meet fundamental technical challenges which include strong computational (computing, storage and software) needs due to the unprecedented size and complexity of the datasets and the data rate that is generated by these facilities. INAF is currently part of **LOFAR and MeerKAT+**, whereas the benchmark of this roadmap in the long term is the **SKA Regional Center**, **the most demanding program in terms of computing power and storage that is foreseen in INAF** (Appendix 1 for details).

2. **Observational cosmology & time-domain:** An important case for our study is INAF's leadership role in ESA's **Euclid space mission,** which is scheduled to launch in 2022. Given the strong investment in terms of

personnel and resources, a full scientific exploitation of this mission in the medium term is mandatory, and is foreseen a significant request of computing power and storage due to the quantity and size of the datasets and catalogs to analyze. Similar IT challenges might be faced by the community connected to the **Rubin-LSST**, as INAF has the potential to acquire significant data-rights from this telescope in the short-medium term. Furthermore through LSST the time-domain astronomy will also lend itself to data challenges and critical IT, starting the large-volume era of transient astronomy. In order to acquire details on the specific computational needs of the Euclid and LSST communities in INAF, we have organized talks from the coordinators of these projects and critically evaluated their needs (Appendix 2).

3. **HPC Theory :** An important point of our study is to understand the computational needs of the broad community in INAF that is involved in theoretical simulations. Theory-driven projects will always make use of Tier 0-1 systems, for which the increase in performances will be matched by an increase in resolution and/or in the complexity of physical processes treated in simulations. Current models of access to supercomputer resources in INAF are based on the INAF-CINECA MoU, ISCRA and PRACE. The roadmap for the development of a world-class European Exascale supercomputer and its intermediate steps on short-mid term (i.e, LEONARDO) offer an extraordinary opportunity for the INAF community to develop and use large new generation simulations in the field of cosmology and astrophysics. However, this requires a large computational investment also in INAF, an intermediate level infrastructure to prepare the community for pre-exascale/exascale supercomputers and for the analysis of future large numerical simulations (Appendix 3 for details).

4. **Cherenkov Arrays:** A mention here is for the research line with Cherenkov arrays. ASTRI Mini-Array in INAF is paving the way to the CTA in the mid-term. The analysis of CTA data and the extensive simulations to support data calibration are expected to be compute intensive. The Italian community is planning to face this challenge through a collaboration between INAF and INFN, a collaboration already established in the field.
To our understanding **the current plan to support CTA foresees the use of the Tier 2 facility that is operated by the INFN in Frascati Labs (Rome)** (http://w3.lnf.infn.it/ricerca/computing/tier2-frascati/).

5. General Vision

INAF has in place a coordination of national IT **services** which also includes access to small clusters. On the other hand, in the past the management of INAF has not had the need to develop a clear vision on critical computing and on how to coordinate the computational needs of medium and large projects in the medium-long term. **A clear vision on critical computing becomes now timely for the growth of the community in the field of computer science and for a strong (effective) role of INAF in the international projects and science programs with large computational needs that we have analysed in the previous section**.

Here we report the main points of our vision :

MAIN GOAL OF THE INVESTMENT :
The investment in large computational infrastructure(s) for critical IT should support the astrophysical community primarily in tackling **the critical computational challenges of major projects and research lines**. Such investment should also provide **a central element for the growth of the community active in the field of computational astrophysics** and for an effective interplay of this community with the rest of INAF.

A CENTRAL INFRASTRUCTURE :
It is clear that the computational needs of each of these projects (Sect.4) require medium-large infrastructures (5000+ core size) that cannot be easily managed locally. Furthermore, by examining the computational needs of the communities involved in the different projects, we realized that the diversity of needs within these communities are as large as those between different communities.
We therefore believe that a ***centralized solution*** (large and with a certain level of heterogeneity) in the next few years **would allow INAF to grasp all the different critical needs more effectively than a solution based on the combination of medium-sized local infrastructures.**
A key point in this case would be the **connectivity to even larger systems** operating in Italy (e.g. CINECA, Tecnopolo and LEONARDO). More generally, however, this also implies the need to establish agreements and strong collaborations with external organizations, leaders in critical IT at national level (CINECA, some universities, etc.). Through these agreements INAF should obtain **critical IT services** (e.g. reserved computing time on large infrastructures, storage, etc.), and aim **especially at establishing scientific collaborations and partnerships** for research and training in the field of IT, and for **system operations** (systems engineering) and **acquisition** of infrastructures.

LOCAL INFRASTRUCTURES & COORDINATION :
At the same time it is **necessary to define a correct level of granularity of the medium-small investments that can be sustained locally at the INAF nodes.** A certain level of granularity of the IT infrastructure in INAF must be

supported. Granularity is in fact vital for the growth of the community. In some cases, local infrastructures can be cost effective solutions due to favorable fundraising and administrative management conditions. **INAF may establish a fruitful network combining a large central infrastructure (Tier 2-Tier 1 size, and its connection with even larger infrastructures) and some specialized medium-(Tier 3 size) systems** that individual communities, or projects, may find convenient to manage locally in INAF; such a model has already been adopted by other communities in Italy, for example in the INFN. Also in this case the connectivity between the various systems is a key point.
The presence of a clear scientific-technological vision is crucial to coordinate the different infrastructures in INAF (Tier 2+3 and, on a medium-long term Tier 1+2+3) and the communities that are involved in the critical IT research lines within major projects. **This coordination should necessarily be present for those local infrastructures connected to major projects**; an even more aggressive vision could possibly envisage the coordination of all the computing and archiving infrastructures that have a use on a national scale (for example DC and RC of national, e.g. SRT, or international, e.g. ALMA, facilities, and INAF archives).

A NEW RESEARCH BRANCH:

INAF participates with personnel in the ASI Space Science Data Center (SSDC), which deals with scientific operations, data processing and storage, supporting various scientific space missions (eg. Swift, AGILE, Fermi, NuSTAR, Herschel etc.).
However -at the moment- **INAF does not have a research branch on computational astrophysics and a division for software support**. There are researchers in INAF with great competence in the field of computational astrophysics and software architecture, but they do not constitute a coordinated network, being mainly associated with specific projects (e.g. EUCLID, GAIA, ALMA, LOFAR, SKA, etc) or local research groups (HPC activities).
We believe that the lack of a coordinated network of computer scientists in INAF may constitute an important limitation to the growth of the community in several fields. The formation of a computational astrophysics division in INAF is even more **necessary in light of the large investment that is planned in critical IT infrastructures and the coordination between infrastructures and large projects with critical IT needs**.

IMPLEMENTATION:

Our vision of critical computing in INAF foresees three main players: a **computer science division**, the **infrastructure(s)** and the **large projects/programs with critical IT needs.** The connection between these three players is key for the future success. Their effective coordination requires an **organizational model capable of defining a scientific and**

technical roadmap, and capable of setting priorities and establishing an effective decision-making flow.

6. Suggested investment for a large computational infrastructure in INAF

In our model the architecture of the critical IT infrastructure is **driven by the needs of large projects and programs** with critical computational needs. In particular, after looking at the major projects, we agreed that the investments on the **short-medium** term should initially be designed primarily to **prepare for the SKA RC** (and SKA roadmap, Sect.4.1), to **facilitate conditions to plan exa-scale simulations from INAF** and to **match the needs of Euclid**.

Furthermore, from our analysis we concluded that there are two types of needs:

1. **type 1**: critical IT support for research programs in INAF aiming at the scientific exploitations of big data from international facilities or advanced numerical simulations. Euclid, LSST and theory-HPC, belong to **type 1.**
2. **type 2**: critical IT support for offering shared computational resources within large international projects. The SKA RC, and potentially part of the activities with SKA precursors and pathfinders (e.g., LOFAR), belong to **type 2**.

The two types of needs require a different strategy for the management and different use of the infrastructures (e.g, conditions deriving from in-kind contributions, different access policy and queues, identification of priorities, protocols, model for the policy and roadmap for the upgrade of the infrastructure). **As a consequence the model should be able to meet both requirements at the same time, or it should foresee two lines of investments**.

After looking at the needs of the major projects, we concluded that "commercial cloud" (for both type 1 and 2) is not an option, as it is too expensive in light of the demand in terms of performance and time span of use (e.g., https://www.google.com/url?q=https://www.oversight.gov/sites/default/files/oig-reports/IG-20-011.pdf&sa=D&source=editors&ust=1627307406920000&usg=AOvVaw0-2MlgcuBpzJU_BF__e5s3); of course another obvious drawback is that the cloud would not constitute an opportunity for the development of the technological and scientific know-how in INAF.

One or more infrastructures at INAF are preferred. In this case, if convenient, one can have a hybrid approach, where the interface is designed like the cloud (Virtual Clusters), but the access is to high-performance compute nodes; this is for instance a solution adopted by the CSCS in Switzerland for LHC (and soon for SKA and CTA).

The other point is to understand the level of **granularity of the infrastructure**, both for computing and storage. This of course depends on the needs and architectures needed for supporting the different programs. We identify the following needs and priority for type 1 and 2 in the **short and medium term** :

1. **Type 1:** In the short-mid term, we believe that Euclid and theoretical HPC represent the main drivers for designing this branch of computing infrastructure(s) within INAF. The computing needs from the Rubin-LSST community have been revisited by the commission assuming that the current pipelines can be updated and properly parallelised (at least in part). These requirements appear grossly similar to the needs of Euclid; also the hybrid architecture proposed for Euclid and LSST data analysis do not differ significantly. A source of uncertainty stems from the fact that the precise level of involvement of INAF groups in Rubin-LSST and the specific LSST programs (data-rights) that will be led by INAF are still unclear. Yet we believe that a hybrid **Tier 2 sized** cluster designed for Euclid will also be able to support the LSST communities in INAF as soon as LSST programs will materialize. A **Tier 2 system** is strongly suggested also to support theoretical simulations as a subsequent step with respect to the INAF CHIPP project. The infrastructure will allow the setup preparation and tuning of very large simulations running on Tier 0-1 systems and the postprocessing of the results from these large simulations. Testing and optimising the porting of large simulations to hybrid architectures is also a key motivation for having a hybrid Tier 2 infrastructure in INAF. In general, a Tier 2 system in INAF will also facilitate carrying out research programs in the field of computer science.
2. **Type 2:** the exploitation of **LOFAR data** is currently the largest scientific activity in INAF on the way to the SKA; in the near future a similar activity is expected from the involvement in **MeerKAT+**. In the short term the data analysis within the international LOFAR collaboration will be carried out in Italy by combining the current 4 sites (3 INAF : IRA, OATs, OACt, and the C3S Torino) and part of the PLEIADI (via reserved/priority access), combining a total of about 2000 cores (about 40 medium-fat nodes) and 0.5 PB storage. However the upgrade toward LOFAR 2.0 and LOFAR-VLBI will further boost computational needs by a significant factor in the **short-medium term**. This activity, combined with the analysis of data from other SKA precursors and pathfinders (MeerKAT, ASKAP, and MeerKAT+ in the near future), and with the **computing-intensive activity on the**

SKA-RC proto-network (2023+), deserves a **dedicated Tier 2-3 system with fat nodes and large dedicated storage (10+ PB in 2024+);** storage capacity may also be conveniently offered as part of the network of the international archiving systems of LOFAR and SKA-RC proto-network.

As already mentioned, we realized that the diversity of the needs within the communities are as large as across the communities. Therefore we think that **one combined solution in a centralized location would much better allow INAF to capture all the different needs** than individual, local solutions. A centralized infrastructure would also be more cost efficient (hardware as well as system engineers and cooling systems) than distributing machines in different places.

If the investment is concentrated into a single infrastructure, the architecture should match the different requirements. Both the projects for observational cosmology and the HPC programs have a diversity of needs that can be matched by **hybrid architectures including light, medium and fat nodes and a mix of CPUs and GPUs**. Medium and fat nodes are also required to analyse large datasets from the international SKA precursors and pathfinders (es medium-fat nodes with 8+ GB/core), whereas first experiments using GPU for radio data calibration and analysis are very promising; extensive use of GPU is expected for the SKA RC in the longer term. Fat-mid memory nodes with powerful GPUs would be also useful for ML/AI experiments (Euclid, LSST, radio); for AI/ML purposes only, cheaper GPUs would be sufficient since double-precision performance in that case is not that important anymore.
Storage is critical in the case of observational projects. In fact, both the projects for observational cosmology and the SKA pathfinders (LOFAR) envisage the need of 10+ PB storage (spinning disk and tapes) in the **short-medium term**; in the longer term, it will take an order of magnitude leap to match the storage that is needed for the SKA RC.

Given what is above, in the **short-medium term (2022-27)** we foresee that - if properly designed - **a large Tier 2 system with hybrid architecture and hybrid configuration of nodes (fat, medium) may support both theoretical HPC and large projects (SKA roadmap, Euclid, etc).**
Specifically, we suggest a large infrastructure that may **gradually reach Tier 2-1 size in 2026-27** (20000-30000 cores, with hybrid architecture and about 1+ PFlops in total) and **about 30 PB storage.** Based on our experience, as a reference final configuration (2026-27) we suggest the following percentage (ranges) of distribution of node-types :

- light CPU nodes : 10-30 % (256 GB RAM)

- medium CPU nodes : 30-50 % (512 GB RAM)
- fat CPU nodes : 5-10 % (2 TB RAM)
- medium GPU nodes : 20-30 % (2-4 GPU/node, 512 GB RAM)
- super-fat CPU nodes : 1-2 % (5 TB RAM)

A dynamical partition of the facility (nodes and storage) will be necessary to support **types 1 and 2 needs** at the same time. Furthermore the infrastructure and its priorities and partitions is expected to evolve with time (every 3-4 yrs) **following the evolution of the needs of the critical projects**.

Finally it should be noted that INAF has no internal expertise to manage and operate a computational infrastructure that would in fact be **20+ times greater** than the largest facilities that are currently operated by INAF. A **significant investment in personnel to support this central infrastructure would not be strategic for INAF**. For this reason, we suggest identifying a **convenient site** in which large computational facilities are already operated by specialised groups of systems engineers. A location like **Tecnopolo** (or similar) would be a natural choice, thanks to the environment and to the possibility to have system support from other partners (e.g. CINECA) that routinely manage large (Tier 0-1) systems.

A SKA RC may naturally evolve from this "seed" with the final goal to have in the **long (2028+) term a Tier 1** infrastructure which will satisfy the needs at regime of the Italian pole of the SKA RC network (i.e. a steady state computational capability of 3+ Pflops and 70+ PB of storage “per year”), while having a Tier 2 infrastructure (for continuing to support the other aforementioned needs) as a sub portion of that. We note that this plan makes the selection of a suitable location for the infrastructure even more important. For example, if the infrastructure were located at the Tecnopolo, the possibility could be explored of acquiring a part of the LEONARDO supercomputer (or its update in the long-term, 2028+) to cover the needs of the SKA RC.

7. Computer Science Division

To achieve the scientific objectives and optimize the scientific return from the investment in the computational infrastructures and large international projects, **a strong investment in specialized personnel in the field of computational astrophysics and software architectures is necessary.**
We suggest organizing this staff in a distributed division with a certain granularity, primarily around groups in INAF that are already active in the field. The purposes of this division are mainly:

- **research (R&D)**: coordinate a team of scientists with critical mass that carry out advanced research in the field of computational astrophysics in collaboration with other groups in INAF; this is essential to build a community with a critical mass and adequate technical skills to face the medium to long term challenges in the field of critical (Big Data) data analysis and next generation of numerical simulations.

- **infrastructure**: coordinate numerical and technological experiments aimed at the optimization and design of the architecture of the central infrastructure for critical IT in INAF (possibly also of its satellites, if present). Optimize the science exploitation of the INAF computational infrastructure(s) also through an involvement and support of the activities within the major projects and programs supported by the infrastructure(s).
- **training & fertilization**: establish strong connections with external bodies (e.g. Universities and supercomputing centers) and carry out training actions (schools, support activities) within INAF to optimize the growth of the community. It would be desirable to set up a support service for INAF researchers for the use of very large (Tier 0) facilities abroad.

This division could include researchers and technologists who already work in INAF. However, based on the experience of the "scientific laboratories" operated at international HPC centers and the estimates in the context of SKA RC (scaled to adapt to the size of the investment in INAF), we estimate that the activity of this division should be supported by approximately 10-15 dedicated FTEs (corresponding to approximately 30 technologists and researchers). **Consequently, setting up a computer science division also requires a significant investment to acquire new staff**. We are aware of the difficulty of finding scientists specialized in computer science due to the current recruitment procedures in INAF and the competition with the industry.
Based on our experience in large international programs and with software divisions (e.g. in Max Planck, ICS Zurich, Argonne USA), we suggest ***hiring mainly personnel with an astrophysical-physical background and with strong skills in the field of numerics and software***. A limited number of highly specialized technologists (e.g. software engineers) may possibly be acquired by the division; in this latter case, agreements with external companies are also an option.
We also suggest the use of ***specifically tuned procedures/profiles*** aimed at recruiting personnel for the computer science division ***and related dedicated career paths.*** A possible model could be to encourage institutes/observatories, those that are strongly involved in large programs and projects that require critical IT or with interests in the field of computer science, to select both technologists and/or researchers with specific profiles where the evaluation of their CV will have to balance the two sides of the needed expertise. Alternatively, a number of positions (TD, TI) could be made available directly to the DS and the management of critical computing in INAF.

8. Suggested management and protocols

In this report, we recommended a **centralized IT infrastructure** and a ***computer science division***. The effective coupling between these two lines of investment is clearly the important ingredient for future success.
As discussed in previous sections, in order to optimize the return on the investment, **major projects and programs with critical IT needs**, which **in fact motivate and drive the investment plans**, must be supported by the two investment lines and **should play a crucial role in the decision flow**.
The right balance in a managerial model between these **three fundamental elements** is the critical point for the success.

Furthermore, to support the activity of large projects and to coordinate the research lines of the personnel of the computer science division and their coupling with projects and community, a **decision process is necessary which is based on a solid scientific vision**. A scientific and technological roadmap should also guide the design of the architecture of the centralized infrastructure, its development over time and the connection with large external supercomputing centers.

All this suggests a **coherent management of both the infrastructural part and the computer science division**, possibly with a unified coordination to avoid bottlenecks in the decision-making and operational flow.

To our understanding, this complex organization of support and R&D activities in critical IT **does not easily fit into the current managerial structure of INAF**. For example, this activity goes far beyond those that are the INAF ICT activities in its current configuration; in fact ICT activity provides and coordinates **IT services at different levels**, yet the **ICT office does not deal with critical IT or R&D in the critical IT field, nor with establishing a scientific roadmap in this field**.

As a consequence an "*ad hoc*" organization model for critical IT is needed. In this report we will not discuss in detail how the suggested model could fit into the current possible INAF management structures (e.g. DS, UTG, Laboratories, etc). Rather, we provide advice on some important ingredients that may be used by INAF to build an effective operating model.
It should also be stressed that we have been asked to provide advice on the critical IT and not on the overall organization model of IT activities in INAF. However, in order to provide a more general view, we also attempt to explore how our model may be combined with the rest of IT initiatives in INAF, namely

small infrastructures, DC and RC of observational facilities and HPC activity from small groups.

In FIGURE 2 we report the main management layers and connections in our model :

1. Given the current INAF managerial structure, it appears natural to have the management of the investment on critical IT (as a whole) as a duty of the **DS structure**, allowing effective connectivity with other INAF management units;

2. We suggest the presence of a managerial layer (level 1), preferably an **executive board (President/Chair and voting members)**, that provides the scientific-technological roadmap and decides on the priorities. The board should supervise the management and coordination of the central (and satellite) infrastructure (via a *Manager* of that) and of the computer science division (via a *Coordinator* of that). In this context, the board, jointly with the DS, will also decide on the roadmap for the hardware acquisition, as well as evaluate the requirements and decide on the plan for the enrolment of the personnel of the computing division and of its geographical distribution; for this latter reason it may be useful to include also a delegate of INAF Directors in the board.
Also, this board (most probably via a representative, delegated by the DS/INAF President) should keep the relations with the international or national partners and institutions.

3. **Conflicts between board decisions and the needs of major projects established by INAF should be minimized**. For this reason, major projects and ongoing programs should also play a role in the decision-making flow. This situation may be established for example through a **delegate** in the executive board; alternatively, an effective process of interaction between projects and the management at level 1 must be guaranteed.

4. It is desirable to have an **advisory board** (scientific and technological) that may provide advice directly to the level 1 of the management. This council can be composed of representatives from the major projects involved in critical IT, delegates from INAF CSN, and independent scientists and technologists; following the practices adopted by other Institutes abroad, a small number of independent members of the advisory board should be invited from the *science labs* operated at international HPC centers.

5. The level 2 of the management is composed of a **Manager** of the central computing infrastructure and of a **Coordinator** of the IT Division. They should have the responsibility of operations and coordination of the personnel. Both should join the executive board as (non-voting) members. Satellite facilities

(Tier 3 sized) that might be established by major projects and programs, should also have a **referent person**; in this case, coordination between the central infrastructure and its satellites will be a duty of the level 1 of the management.

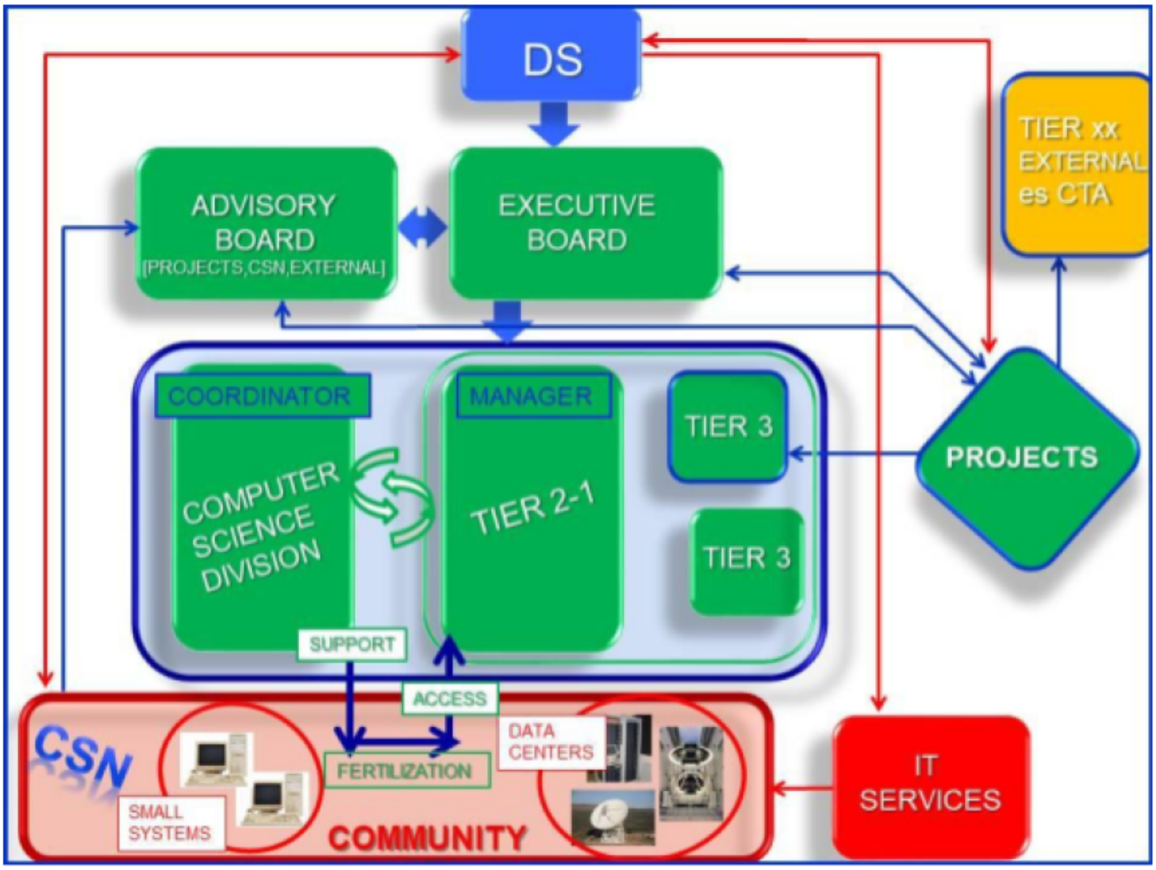

FIGURE 2 : block diagram of the proposed management model. Green boxes indicate IT-critical bodies.

6. External computing infrastructures that may be established by large programs through collaborations between INAF and external partners (e.g., the case of the infrastructure for the CTA) are not included in our model. A certain level of coordination between these external facilities in which INAF is involved and the system of INAF critical infrastructure(s) may be convenient in a number of cases.

7. IT services and small infrastructure(s) that are managed locally (i.e. Institutes or Observatories) should not necessarily be included in our model. Small infrastructure(s) may include DC of INAF facilities (e.g., SRT, TNG etc) and RC of international facilities (e.g., ALMA) that do not have critical IT needs. In the present model a certain level of coordination of these small distributed facilities with critical IT infrastructures in INAF may be easily guaranteed by the DS office and its units. In the event that the projects supported by these facilities evolve and require significant IT needs, INAF

may decide to include them in the critical IT model; in this case an update of the infrastructure(s) will match the additional needs.

8. IT services that are not directly related to critical-IT (licences, commercial software, systems, administrative IT) and their coordination at national level, should not be included in the model for critical IT and may remain a task of the current ICT office.

There are also ingredients that are less clear at this level of the study :

1. In our report we have not discussed the decision-making process within INAF that should determine which projects need to be included in the critical IT model. As already mentioned, the list of ongoing major projects and programs was already present in our engagement letter. Presumably the decision-making process will be based on the interaction between the community (e.g. CSN, CS, etc.) and the INAF management; the participation of the critical IT management in this process would also be desirable.

2. In general major projects and their critical computational needs may have a lifetime that is shorter than the lifetime of the entire investment. Even if the investment must be oriented to support large projects, **projects should have a dynamic role within the organization**, with a **periodic re-evaluation of the reference projects from INAF (3-4 yrs)** (see previous point). Also for this reason it is desirable to have a project delegate within the executive board, rather than a representative for each major project, whereas the composition of the advisory board may be re-evaluated periodically in order to include representatives of the most relevant projects.

3. A certain level of **coordination** between the activities for the **archiving** and **storage** of large infrastructures for critical IT and those devoted to smaller facilities is **desirable**. Given the different scales and management of the systems, in general we do expect that these activities face different challenges. However, we do also expect that the two activities in INAF may share qualified personnel.

A protocol-model for reserved access :

The model of access to the infrastructure(s) is a relevant part of the organization. Several models for reserved access and dynamical partition of the central infrastructure can be foreseen. We suggest a model of reserved access that is based on a **budget of core hours and storage per project that is guaranteed on specific nodes and partitions**. The model will be defined by the level 1 of the management (projects will outline their needs and/or the international agreements for in-kind contribution [type 2]) and operated by the Manager of the central infrastructure. More specifically :

1. **theory-HPC** : a specific budget may be distributed through **calls for proposals** within the community; the use of a large fraction of the infrastructure (exploiting up to Pflop capacity) might be made possible for very large simulations;
2. **major projects** (es. EUCLID + Rubin-LSST) : reserved access to a specific budget of node hours and disk space;
3. **radio** (LOFAR, LOFAR 2.0, MeerKAT+, SKA-RC proto-network) : reserved access to a specific budget of node hours and disk space, part of the budget may be allocated **within a reserved partition** depending on international agreements [type 2];
4. **technical reserved access** : part of the budget (primarily core hours) can be reserved for technical experiments (architecture, partitions, etc).
5. **Open access** : depending on the request, a limited fraction of the budget may be devoted to support observational programs which meet critical IT needs for a limited period of time.

9. Roadmap

In this section we report on our suggestions for the roadmap of the investment.
Clearly it is not possible to suggest a detailed roadmap for a number of reasons :

1. The timing of the roadmap **depends on the possibility of having an ideal site for the central Tier 2-1 infrastructure**. Despite being aware of the existence of agreements to have space dedicated to INAF at the Tecnopolo, the time-scale is currently still uncertain to anchor a roadmap.

2. At the moment the **real investment of resources within Rubin-LSST is not clear**, as a negotiation and evaluation phase is still underway; a **similar situation exists for MeerKAT +**, in which case the role of INAF within the Key Projects is still to be discussed. It is reasonable that these situations will become more clear already by the beginning of 2022.

3. Although the basic timing has been decided, **the technical details of the roadmap towards SKA and SKA RC are not yet fully defined**; a level 2 document on requirements of the SKA RC network and the implementation plan is expected in 2022.
Furthermore, due to the **establishment of the LOFAR ERIC** (expected in 2022-23) and the **LOFAR 2.0 upgrade** (2023-24), INAF might find it convenient to invest in computing and archives for LOFAR/LOFAR2.0. Also in

this case the situation will be more clear by the end of 2021 or beginning of 2022.

Apart from these uncertainties, **it is clear that an investment in critical IT in INAF is urgent**. It is necessary to consolidate the ongoing activities in the field of HPC and data analysis of SKA precursors (potentially exploiting the opportunity provided by ERIC LOFAR and consolidating the scientific weight in LOFAR 2.0 and MeerKAT+). At the same time, it is necessary to be adequately prepared for 2023, when the experimentation activities around the SKA RC networks will begin and when large amounts of data from Euclid and possibly Rubin-LSST begin to be available.

We suggest a roadmap based on 2 phases, a short-mid period (2021-27) and a longer term (2028+).

In the **short-medium term (2021-27),** INAF could identify **two main steps of the roadmap**. A **first step, in the short (2021-24)** term, should be devoted to interventions on the personnel and on programmatic agreements, and on establishing a first seed of a large INAF computational infrastructure. T**he second step (2025-27)** should allow INAF to match the level of investment suggested in Sect.6.

During the first step (2021-24) we foresee a number of **urgent priorities**:

1. It is important to start immediately from the **computer science division**, by organizing a network of researchers with strong skills in the field of critical IT and with strong connections and interest in the major projects underway and in the HPC activity. A gradual acquisition of new staff in these areas should begin immediately, preferably by reinforcing existing groups.

2. The current MoU with CINECA is expiring and will be renewed soon. We believe that INAF should take this opportunity to strengthen connections with CINECA by establishing a framework agreement that combines service and R&D activities in the critical IT field. This is an important step in light of the overall investment that INAF will make in the field and to establish a strong link with CINECA in the context of the Tecnopolo, aiming at sharing spaces and system support for INAF infrastructures.

3. A first investment on a platform that aims at a computing power of about 200-400 Tflops and storage. For reference, this could be a computer with a total of about 200 nodes (8000-10000 cores) with a high percentage of GPU / node (25+%), preferably based on medium-fat nodes (512+ GB), with few super-fat nodes (5TB) for experimental activity. A large number of GPUs is

central for porting activities to large hybrid infrastructures (HPC), furthermore, GPUs are expected to be very important for SKA RC and are common needs to the activities foreseen with SKA precursors, Euclid and LSST.
Disks and tapes are long-lasting investments having a much longer lifetime than processors. Therefore a significant investment in the first step (eg 5-10 PB / year) could be correct and also strategic. For example, the management of storage for any in-kind contributions in the context of large international projects is much simpler than that of computing cores; eventually the storage-infrastructure could then be used for other projects (e.g. SKA RC) in the longer term.

A **warning** regarding this first step clearly derives from the possibility of combining the hardware investment in a **single site**; also for this reason it is crucial to immediately establish agreements with CINECA (i.e., point 2).
In the event that it is not possible to find an ideal site, the hardware could be split (for a limited time and as short as possible) into several poles that are locally operated (at system level), and centrally managed. Initially this would constitute an important drawback especially for HPC activities which could not thus exploit the full computing power of the infrastructure. For this reason we suggest **not splitting the hardware into more than two poles**, initially providing a slightly larger infrastructure (cores) for HPC and observational cosmology, and a slightly smaller one for radio and SKA RC activities, whose management might be more complicated due to the presence of in-kind contribution activities [type 2].
First step would allow experimentation activity around the infrastructure by astrophysicists with strong computer skills (importance of point 1). This will be important also to define the architecture of the Tier 2-1 (Pflop +) infrastructure that should be implemented **during the second step in the medium term (2025-27)**.

In the longer term (2028+) the activity of SKA RC will become predominant, deserving a **step 3 of the investment** to match the extreme computational and storage needs. We note that during the first two steps INAF might have already set up an infrastructure for storage of 30-70 PB which may thus provide a significant seed toward the storage of a SKA RC. As already mentioned during this report, a multi Pflops computer reserved for SKA RC activities could be acquired by INAF as part of very large facilities in the exa-scale era, for example in connection with the renewal of the LEONARDO, presumably expected in 2028+.

10. Conclusions

At the end of April 2021 the Scientific Director of INAF appointed this committee to obtain advice on the scientific, technical and organizational aspects of a large investment in the field of critical IT by INAF.
In about 3 months of work we have acquired a vision of the current state of INAF and of the involvement in large observational projects and theoretical programs with critical IT needs. We also examined the specific needs of larger projects through presentations from the community involved in these projects.

In the light of our information **we have built a scientific, technical and organizational model for the investment on critical IT in INAF**. The model is based on three main ingredients: **large projects** for which investment in critical IT is designed, a **large central Tier2-1 infrastructure** (Tier 3 satellite infrastructures, if necessary) that can grow in the medium-long term based on an assessment of needs and a scientific-technical roadmap, and a **computer science division** that brings together scientific expertise in the field of critical computing. We have suggested an integration model between the computer science division and computational infrastructures aimed at optimizing the overall investment and fertilizing the scientific community in INAF.

Specifically we have also suggested a **management mode**l for the overall investment and the coordination between infrastructures and critical IT activities with local activities and small facilities. This model is based on an **executive committee** (President/chairman and voting members), an **infrastructure manager and a coordinator of the computer science division**. It is also desirable to have an **advisory board** which may include expertise from large projects, the scientific community and the community abroad.

Finally, we have suggested a **roadmap of the investment** which foresees in the short term a priority on the personnel of the computer science division and the acquisition of a medium-sized computational facilities that can consolidate ongoing activities in the context of precursors SKA and HPC, and that can also used for the analysis of data from optical surveys (Euclid, LSST) in the medium term.
Starting from 2024-25, we expect an expansion of the infrastructure over time until arriving in 2028-30 to match the requirements of a SKA RC which represents the greatest computational challenge in which INAF is involved.

11. Acknowledgements

Our suggestions were developed in total independence during internal (online) meetings and discussions via email. However, in accordance with the guidelines suggested by the Scientific Director of INAF, we also interacted (via bilateral contacts or talks during our meetings) with several INAF researchers to obtain information on current and future projects and on INAF organization. In this respect we acknowledge very useful discussions with numerous colleagues, including L. A. Antonelli, U. Becciani, F. Bocchino, M. Brescia, A. Fontana, C. Gheller, M. Nanni, R. Smareglia, G. Taffoni, L. Valenziano, A. Zacchei. Members of the committee also attended the “audizioni” of CSNs at the end of May 2021.

APPENDIX 1 : THE ROADMAP TO THE SKA

The scientific roadmap to SKA is based on research activity with precursors and pathfinders. INAF participates in LOFAR (2018+) and more recently in MeerKAT+ (2021+). This participation involves leadership roles within key programs and data-rights, and the involvement of the community in major technical and computational challenges in data analysis and management.
The INAF community is also actively involved in the scientific exploration of other precursors, including MeerKAT and ASKAP.

LOFAR

The pan-European LOFAR telescope is the major precursor of SKA also in terms of data challenge. It generates an unprecedented (for astronomy) data flow of about 200 Gb/s, and the datasets to calibrate and analyze are typically in the range 1-10 TB, i.e. about 100 times larger than the data of previous radio telescopes (e.g. VLA, GMRT).
LOFAR currently represents the largest investment in SKA precursors/pathfinders at INAF, in terms of personnel involved in scientific exploration and data challenge. About 100-200 TB of LOFAR data are analyzed in 2020 and 2021 and a large part of the LoLSS survey data (about 2000 hours of observation at PI INAF) will be analyzed at INAF in 2022-23.
Currently there are 16 medium-fat nodes (384-768 GB / node) dedicated to the analysis of LOFAR data in Italy (distributed over 3 INAF poles and 1 pole at the C3S in Turin) and a storage of about 0.5 PB. A significant update of the infrastructure is expected by 2021 based on priority / reserved access to part of the PLEIADI (reaching a total of about 2000 cores dedicated to LOFAR analysis distributed in about 40 medium-fat nodes). However the upgrade toward LOFAR 2.0 (2024+) and the activity based on the LOFAR-VLBI (expected to enter in production-mode in 2022+) will further boost computational needs by a significant factor in the short-medium term.

The LOFAR community is also active in the field of computer science and pipelines, through the activity of a Data WG that brings together about 15 researchers (TD and TI) with strong computational skills, and specific computational tasks are conducted in collaboration with ASTRON and the Survey Key Program.

An opportunity for the future is also represented by the ongoing activity for the ERIC LOFAR which plans to submit phase 1 of the ERIC in September 2021 and to establish an ERIC by the end of 2022; Italian MUR is fully supporting this activity. Current ERIC agreements (Financial model) open the possibility to form new poles of the international LOFAR archive (currently 60 PB concentrated in NL, Germany and Poland) in the member countries through a scheme based in-kind contribution (150 kEuro per PB / year). It could therefore be convenient (and strategic for the future towards SKA RC) to foresee computational and storage investments for LOFAR within the roadmap for critical computing in INAF.

SKA RC

In the longer term, SKA will produce an enormous amount of scientific data, at an overall rate of about 14 TBy per second. Most of their reduction (in terms of overall volume and in order to make them usable by the scientific community) will take place at the two observation sites. The resulting scientific data will in any case occupy approximately 710 PBy for each year of operation and will require a steady state capability of 22 Pflops of processing power to be analyzed.

The resources needed to maximize the ability to carry out further processing of the aforementioned data are not included in the SKA1 construction and operational budget. In March 2016 the SKAO Board therefore encouraged member states to form a collaborative network of SKA Regional Centers (SRC) to provide those additional functions that will not automatically be available within the SKA1 project. In November 2018 the SKA Board then approved the formation of the “Steering Committee” of the SRCs (called SRCSC), including in it a representative for each member country of the SKA Board. The mission of the SRCSC is: "To guide the definition and creation of a long-term operational partnership between the SKA Observatory and a set of SKA regional centers based on independent resources".

Within this framework, the SRCSC has assembled a series of Working Groups (WGs) with the task of defining and taking the necessary steps over the next two years and beyond. Currently the following 7 groups have been formed: Architecture (WG0); Data Logistic (WG1); Operations (WG2): Federated computing and software services for data analysis (WG3): Science Archive-Virtual Observatory–Implementation of FAIR principles (WG4): Computational systems (WG5); User engagement (WG6)

Timeline

The activities of the SRCSC began in May 2019 and those of the WGs in January 2021. The first fundamental deliverable will be the preparation of a Level 2 Document containing all the indications, emerged from the work of the various WGs regarding the requirements of the SRC network and the implementation plan of the same. This document is expected for the end of

2022. By mid-2023 the entry into operation of a SRC proto-network is also planned, within which to begin to involve the radio astronomical community. The dual objective is on the one hand to test the system, but, on the other hand will be to produce some initial scientific results, starting from data collected at precursor telescopes/pathfinders (first of all LOFAR, Meerkat, Meerkat+, ASKAP, but not only those instruments) or related to particularly complex simulations of SKAO data. In this context, the SRC proto-network will become part of the cyclical "Data Challenge " that SKAO will continue to organize.

Budget ramp up

The cost estimate is based on the preliminary analysis carried out by the SRCSC: when fully operational, the total cost of the worldwide network will be 34.5 M€ per year all inclusive: personnel, maintenance, hardware upgrade. From the growth curve, we obtain a projection of 200 M€ in the reference period 2019-2030. These estimates are affected by various uncertainties, which lead to a range that fluctuates from 150 M€ to 300 M€. Assuming the best available estimate (200 M€), the overall envelope for INAF to maintain a leadership role, and therefore to organize its own SRC hub included in a European network, will certainly be higher than 6% (percentage contribution of Italy to the IGO) and should amount to about 10% of the total cost, that is 20 M€, spread between 2021 and 2030. An (assumed) linear growth from almost zero cost in 2020 up to the cost at regime in 2028 - when the INAF supposed 10% of the yearly total cost of the network will sum to 3.5 M€ - implies starting with an investment of 430 k€ in 2021, 860 k€ in 2022, 1300 k€ in 2023 and then goes up steadily.

Key points for planning

Alongside the fundamental need to continue to train first-rate scientific personnel in radio astronomy, two further ingredients are fundamental to guarantee INAF to play a leadership role in the exploitation of the SKA project: (i) the availability of state-of-the-art and quantitatively adequate IT resources and (ii) the ability to attract and to train a human capital of excellence that knows how to make the best use of the resources referred to in the previous point.

Point (i) As regards the first point, the solution that would maximize the leadership role of the Entity is linked to the creation of an Italian hub for the network of SRCs, possibly to be integrated into a European hub. It should also be integrated/coordinated with the INAF calculation structure. Opportunities emerged in this sense (for example a central hub at the Bologna Technopole and a distributed structure to conduct R&D or for the specialization of the various calculation models) seem promising. Human support for this pole can also be distributed throughout the national territory.

Point (ii) As regards the second aspect, the starting situation in INAF is encouraging from a qualitative point of view. In fact, there is internationally recognized expertise within the organization in key sectors of ICT with a scientific orientation. The maintenance of these excellences is facilitated by the participation of INAF researchers and technologists in various international teams (not only directly connected with SKA), whose activities are focused on the issues in this sector: think for example of the EOSC, IVOA, EuroHPC projects, AENEAS, ARC, etc. The large participation of INAF members in the WGs of the SRCSCs is also a guarantee on the one hand of being able to influence the design of the SRC network and, on the other hand, ensures that the Entity is constantly informed about the most recent R&D developments in areas of direct relevance to SKA.

On the other hand, the White Book prepared by the SRCSC identifies in about 100 the number of FTEs necessary for the maintenance of the worldwide network of the SRCs. Most of these FTEs will have to be provided "in kind" by the member countries of the network. This translates into at least 10 FTEs of INAF personnel who will have to gravitate around the activities of the SRC network in order to maintain a significant role for INAF within the SRC system. This number will probably have to grow, reaching 15 FTEs, if, as desirable, INAF will opt to budget its own Italian hub for the SRC system. Today the number of people in INAF with the necessary technological / IT / software skills is clearly below this numerology and this is true in particular as regards the available FTEs.

There is therefore the need to start as soon as possible a plan of aggregation to INAF of figures who can immediately support the existing expertise and fill the obvious staff gaps that will emerge in the years to come. The main skills required will be related to: the deep understanding of the operation of the raw data acquisition systems at the 2 SKA observation offices, the management and development of the complex systems that will oversee the analysis of SKA data and their archiving, software development for the various scientific cases, as well as the interaction with national users in the preparation and management of SKA observation programs.

APPENDIX 2 : Euclid and Rubin-LSST

On June 10 we invited representatives of the Euclid and Rubin-LSST projects to discuss the computational needs of these projects. The presentations are available at the following links:

Insert link to Euclid presentation

Insert link to Rubin-LSST presentation

In brief summary (details in the presentations and in the discussions presented in the main body of our report) the most relevant Euclid computing needs for critical IT at INAF refer to the exploitation of Euclid results, i.e. programs based on data that are already reduced. A significant part of this full exploitation (including in part also the use of ML/AI) is planned in conjunction with other surveys (including Rubin-LSST, SKA, SKA precursors or pathfinders, etc). Since the data of these complementary surveys will become available at different times (for example SKA data will become available in 5-7 years after the Euclid data), the resources could be modulated/allocated to the project over time according to a roadmap. Presented specs are quite precise, aiming at about 10000 processing cores including fat nodes (about 1 TB RAM each node at least for 10 nodes) and with an explicit request of GPUs in the platform architecture.

The situation of the Italian involvement in the Rubin-LSST project is currently more uncertain simply because (part of) the process that will define such involvement is still in progress. The most important variable is the amount of data rights that will be acquired by the INAF community and the specific projects that will be led by INAF (or the projects with stronger participation of INAF personnel). Consequently this leads to an uncertainty in quantifying the actual needs for critical IT at INAF that will materialize in the next few years. We have assumed the optimistic scenario in which the outcome of the ongoing process that will define the data-rights and programs of INAF will match the desires of the community.
In this case the computing needs refer to advanced data analysis tools, primarily image processing, to be applied to very large datasets. The entire project is expected to require initially about 150 Tflops (First Data Release) that matches essentially the capabilities of a Tier 2-3 level platform. Up to 10 percent of this computing power may be requested for the analysis of the

survey of the Galactic Bulge, which is among the most important targets of interest of INAF. Computing needs of the entire Rubin-LSST project will grow with time matching the capacity of a Tier 1 system in about 10 years.
In brief summary, the architecture required for Rubin-LSST is in line with Euclid's requirements (see details in the presentations). Although current LSST scientific pipelines are CPU-based, an INAF platform with hybrid configuration (CPU and GPU) in perspective (future porting to GPU) and for data exploitation (including ML / AI) is desirable. Formally the Rubin-LSST community in INAF also put forward the request for a significant number of super-FAT nodes (request based however on the use of current software, for example the current version of the ALLFRAME), however the development of parallel codes would significantly reduce these needs.

APPENDIX 3 : HPC Theory

The Community

A perspective on the Italian community that makes use of High Performance Computing (HPC) for theoretical studies can be obtained from the utilization of the resources that have been made available to the community through the 2017-2020 Memorandum of Understanding (MoU) “New Frontiers in Astrophysics: HPC and new generation Data exploration” between INAF and CINECA. From the projects and the papers that acknowledge the use of those resources, we can estimate a number of about 130 researchers involved and, restricting to the more actively involved researchers, the number decreases to about 70-80. The groups involved are spread over all INAF structures. In the figure below the projects are divided in relation to the astrophysical problems under study (left panel) and to the type of physics and numerical codes that are used (right panel).

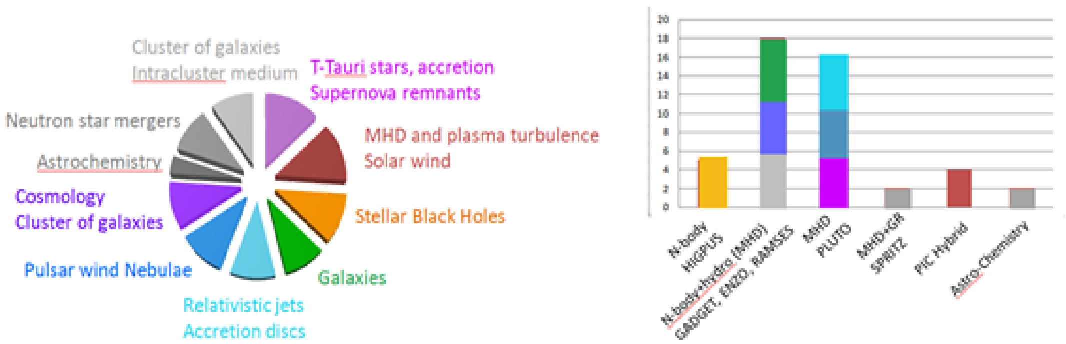

From the panel on the left, we can see that the astrophysical topics cover a very wide range, from cosmology, to high energy astrophysics, stellar astrophysics and heliophysics. On the right we can see that, from the point of view of the numerical codes and the involved physics, there is a prevalence of two classes: codes that couple N-body and (magneto)hydro-dynamics that are used for cosmological problems and purely magnetohydrodynamic codes that are mainly used for high energy astrophysics. One point that has to be stressed is that there is also a very active involvement in the code development.

Access to HPC resources

As already mentioned, the INAF-CINECA MoU guarantees to INAF a certain amount of CINECA resources. This agreement is the last of a series that started in the early 2000 and helped to foster a strong computational

community in INAF. The computing time available under the current MoU (2021) corresponds to 12.5 Mcore-hours on the Tier0 system Marconi 100. The resources are distributed to researchers on a competitive basis, with two calls per year, with three different kinds of projects: type A with a maximum of 1.5 Mh and 1 year duration, type B with a maximum of 0.25 Mh and duration of six months and test with maximum of 25 kh and duration of one month. The average number of projects presented for each call is about 10 with an oversubscription of a factor of about 1.5 in the requested computing time.
A second possibility is provided by the INAF CHIPP project that makes available two small systems located in Trieste (800 cores) and Catania (200 cores). The computational requests have a limit of 200 kh, with two calls per year. Theoretical programs cover about 80% both in terms of approved projects and computational time.
CINECA resources are accessible also through the Italian SuperComputing Resource Allocation (ISCRA) program run by CINECA. This program is open to all Italian scientists, with two calls per year and two classes of projects, Class B with a maximum of 1.5 Mh and a duration of one year and Class C with a maximum of 60 kh and a duration of nine months. On average, each year the astrophysical community gets 13% of the total allocated time with about 10 projects. The total time allocated to astrophysical projects is approximately equal to the time obtained through the MoU.
Finally, large computational projects may access all Tier 0 systems in Europe through the Partnership for Advanced Computing in EU (PRACE). Typical allocations are of the order of tens of Mh, as an example the minimum allocation on Marconi 100 is 35 Mh. In the last ten years there have been 17 astrophysics projects from Italy, that constitute about 10% of the total.

Towards the Exascale

At present, the main computing systems at CINECA are Marconi 100, a Tier 0 system with hybrid CPU/GPU architecture and a peak performance of 32 PFlop/s, and Galileo 100 (available for production from July 2021), a Tier 1 system with a majority of CPU only nodes. The future is leading towards the Exascale. The EU initiative EuroHPC aims at deploying a European world-class Exascale supercomputer. The first step is procuring and deploying by 2021 three pre-exascale supercomputers in the EU. One of the selected sites for those is CINECA. The pre-exascale system LEONARDO at CINECA should be available for production in the first months of 2022 and will have a hybrid CPU/GPU architecture and will deliver a peak performance of 250 Pflop/s.

Perspectives

Theory-driven projects will always make use of Tier 0/ Tier 1 systems, for which the increase in performances will be matched by an increase in resolution and/or in the complexity of physical processes treated in simulations. The current model of access to resources with the INAF-CINECA MoU has been successful, with the advantage of a very short time span between the presentation of proposals and the availability of resources.

A Tier 2 system can be however very useful for tests, setup preparation and tuning since performing these tasks on the bigger machines may be hampered by the long waiting times in queues. The convenience of this approach has been shown by the results of the INAF CHIPP project. In addition it can be used for postprocessing of the simulation results. For this last task an adequate storage capacity is required.

A critical issue is the porting of codes to the new hybrid architectures. At present, the porting has been partial with mixed results in terms of performances. This task requires a strong investment in person power and specific competences.

APPENDIX 4 : conversion schemes between CPU, cores, GPU and Tflops

In the following we report peak performance of modern CPU (with typical 20-28 cores per CPU) (upper) and GPU (bottom). These performances have been used in our report to estimate the performance (Pflops) of the proposed infrastructure(s).