Supercomputing and Big Data for the SKA: challenges and opportunity for Italy

R. Smareglia
Head of ICT & Science Data Management
A. Possenti
Advisor UTG II HPC, archives and data transfer

The II National Workshop of SKA
science and technology
Bologna - 3-5 December 2018

Challenge: **The full system**

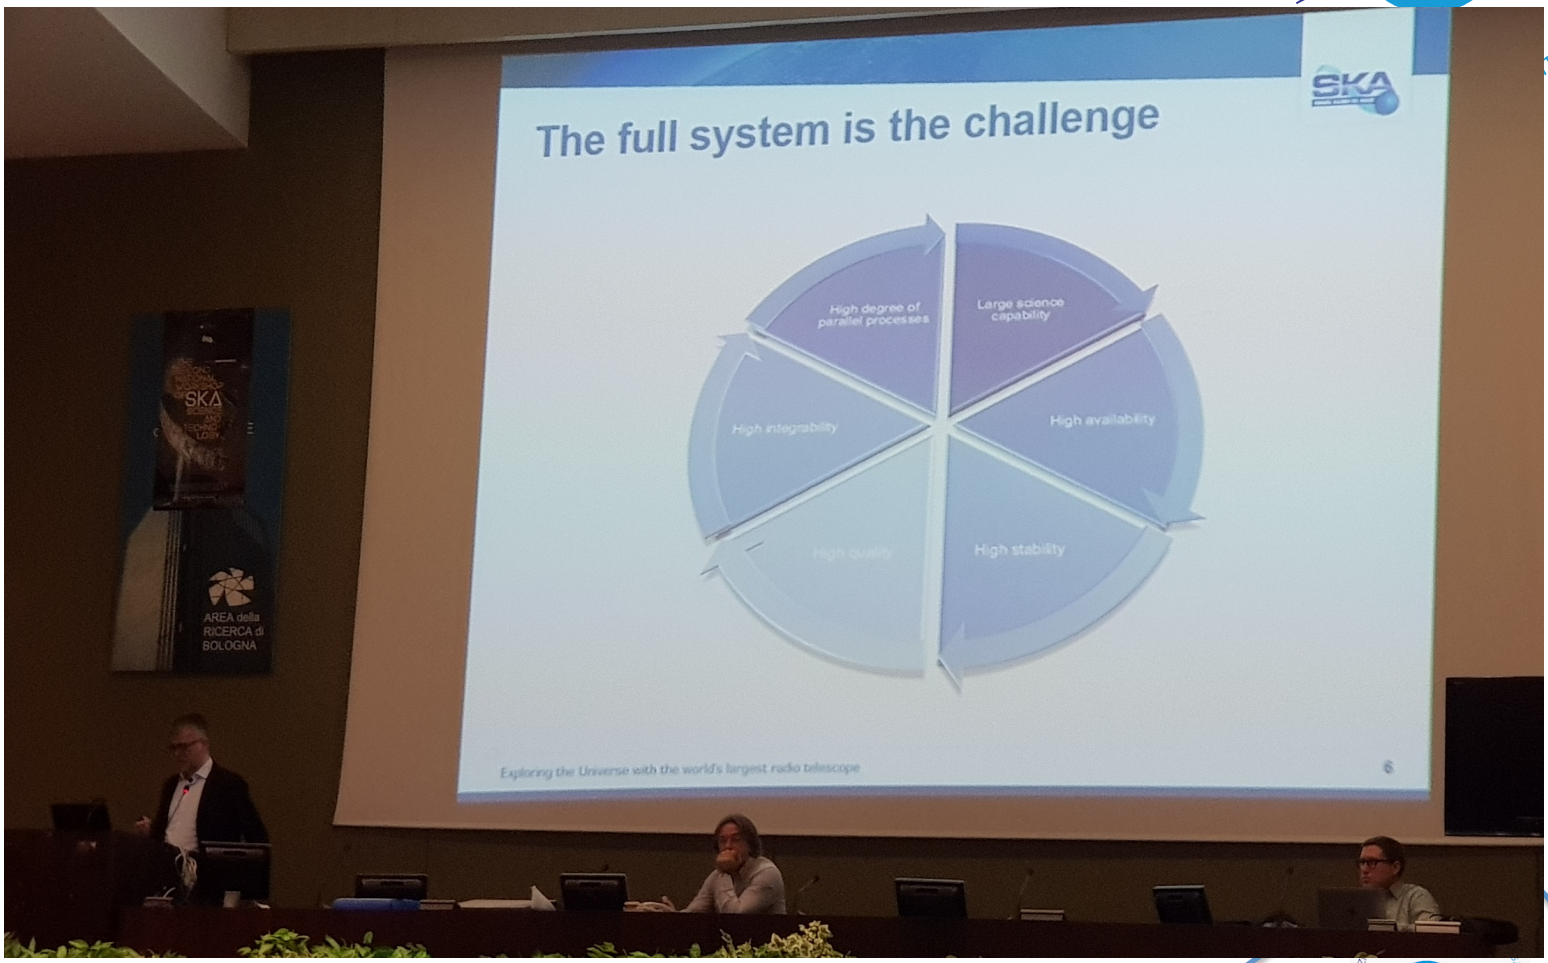

SKA RC Requirement

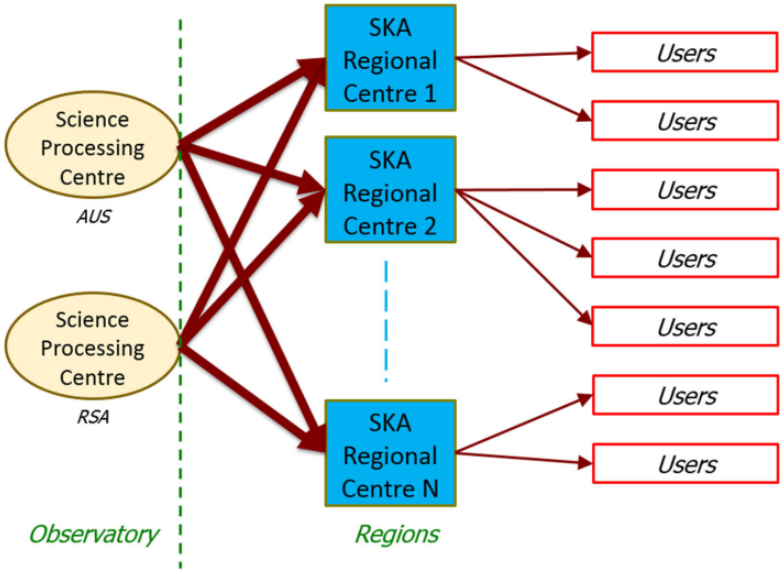

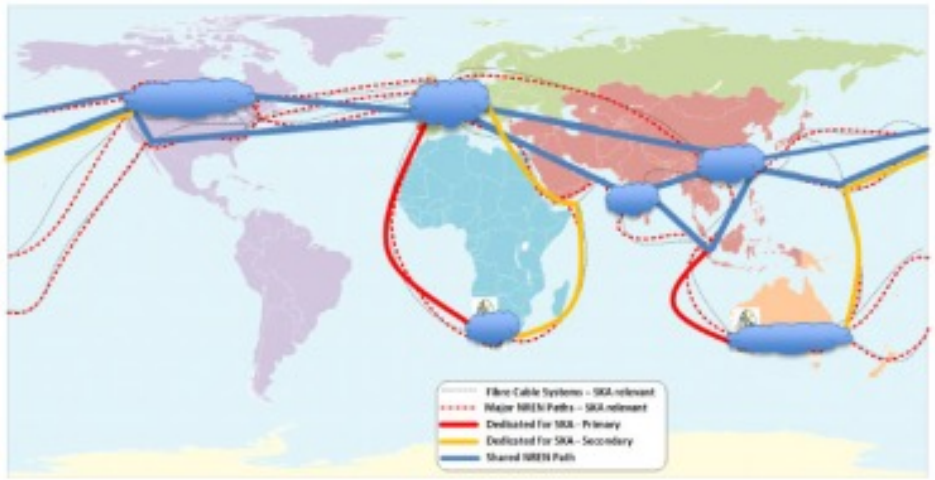

On 2026 ➔

➔ ~ 1 ExaByte of storage

➔ 35-80 PFlops (from 2024)

Big Challenges …

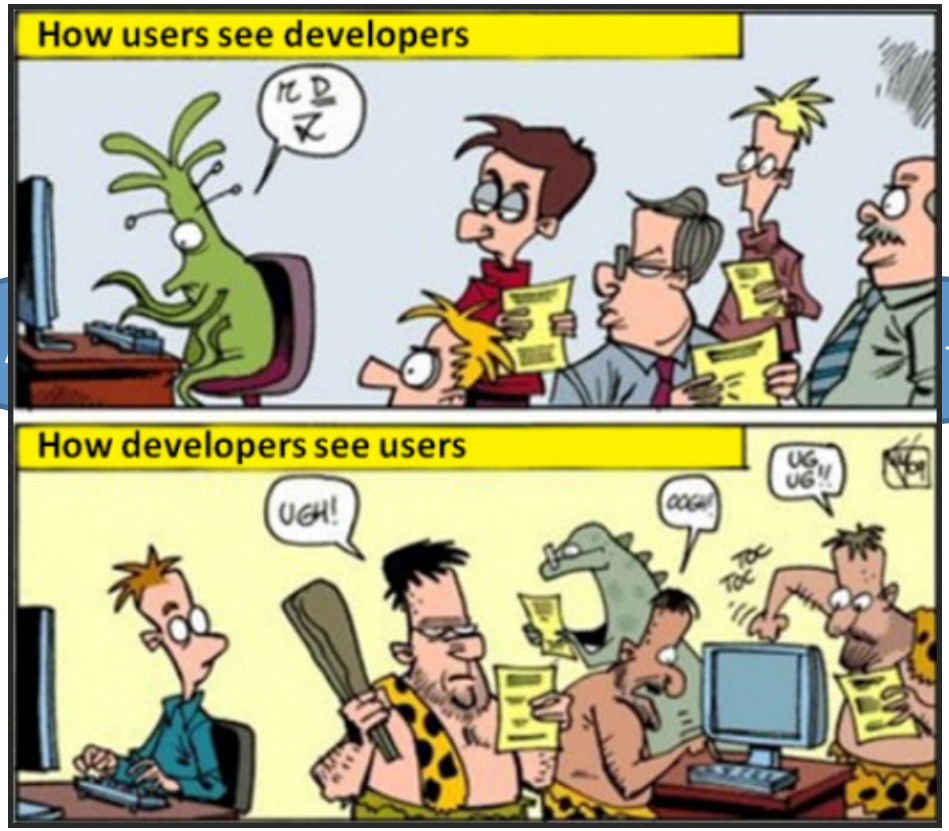

ting

SKA data challenger #1

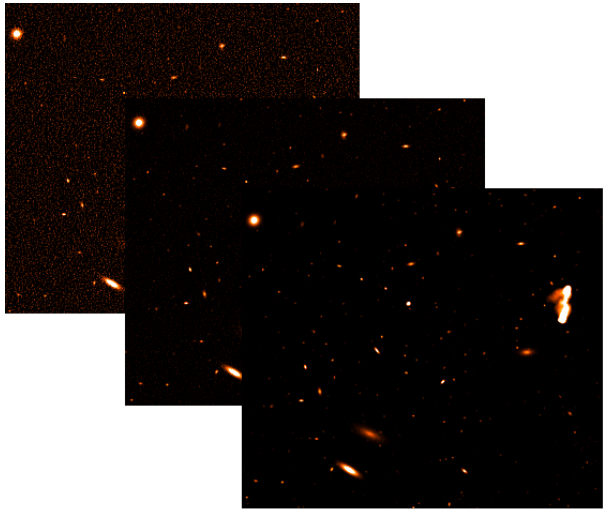

The SKA Science Data Challenge #1 (SDC1) release consists of 9 files, with the format of FITS images. Each file is a simulated SKA continuum image in total intensity of the same field at 3 frequencies (560 MHz, representative of SKA Mid Band 1, 1.4 GHz,

This research used the facilities of the Italian Center for Astronomical Archive (IA2) operated by INAF

Challenge Description

The challenge set for the community is to undertake:

- source finding (RA, Dec) to locate the centroids and/or core positions,
- source property characterization (integrated flux density, possible core fraction, major and minor axis size, major axis position angle)
- source population identification (one of SFG, AGN-steep, AGN-flat)

SKA data challenger #1
➔ Seed of Italy SKA RC

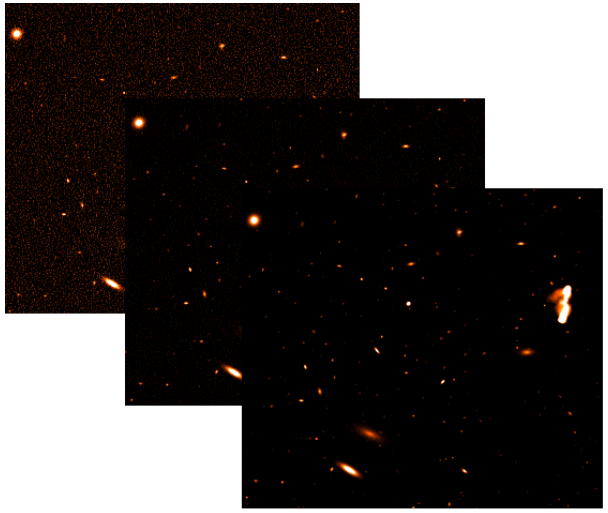

Computing Infrastructure:

- CHIPP
- HPC @ CINECA
- Commercial Cloud

This research used the facilities of the Italian Center for Astronomical Archive (IA2) operated by INAF

User Support:

- ALMA RC – Italy
- …

Big Challenges …

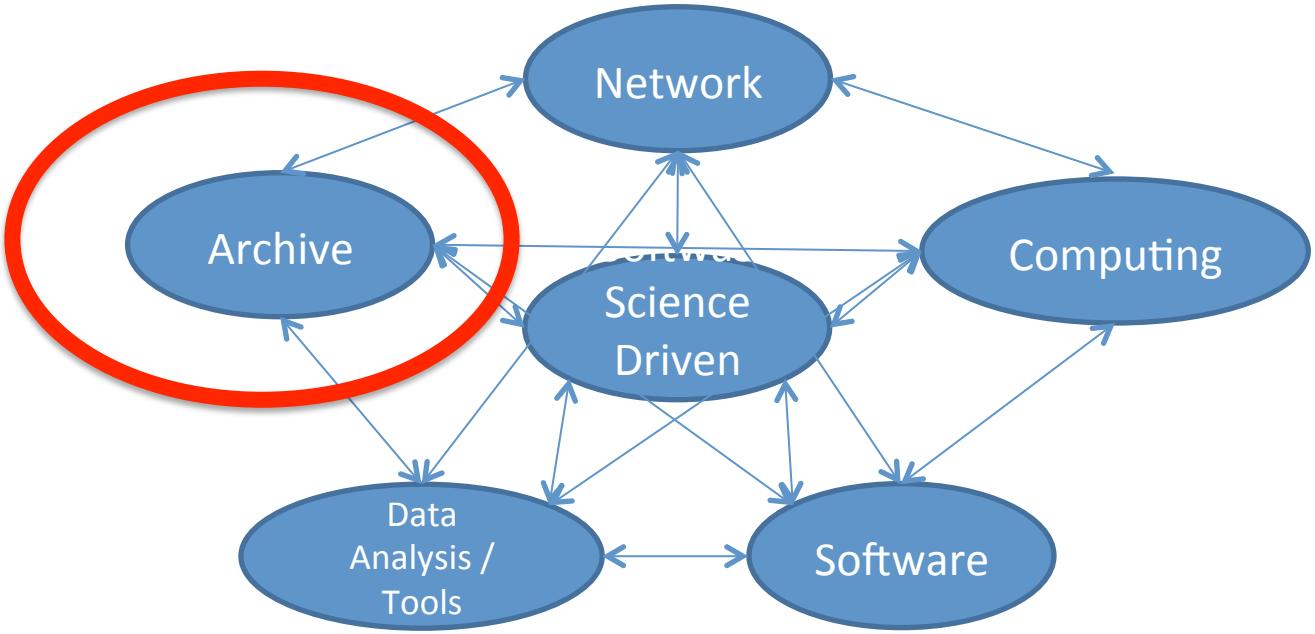

The “pyramid” of scientific data

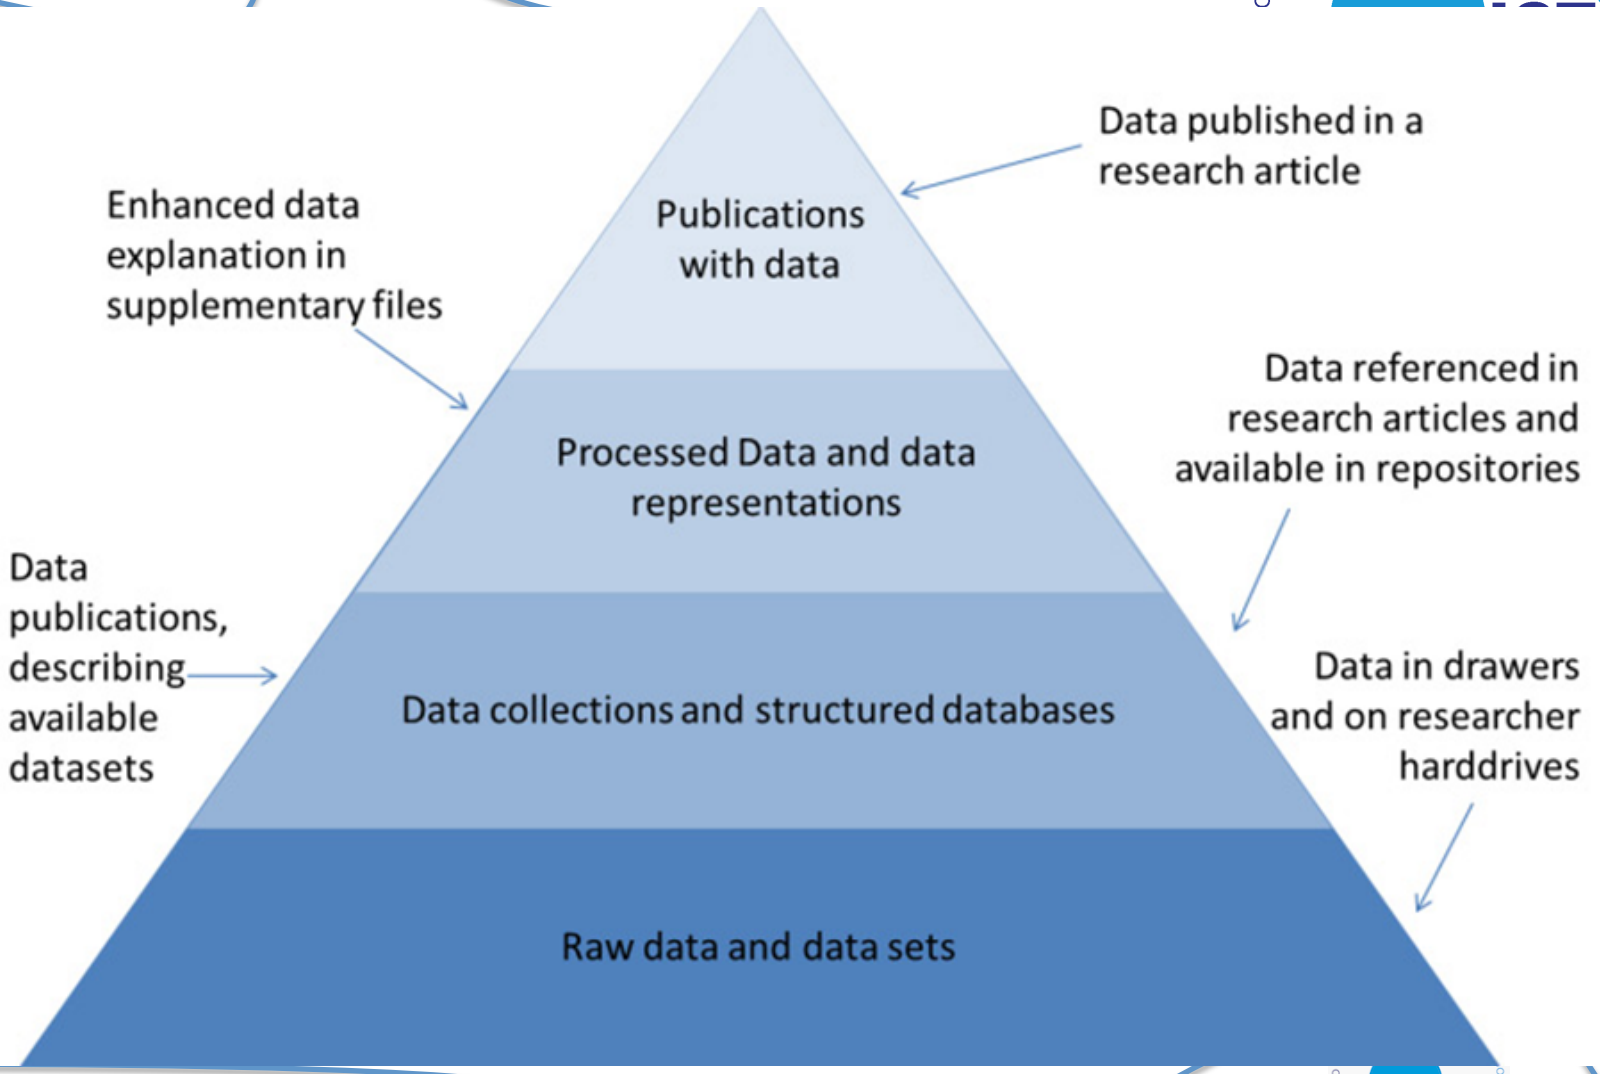

Source: Reilly, S. et al. (2011), Report on Integration of Data and Publications, ODE, Opportunities for Data Exchange

Data Archive: Why it's important

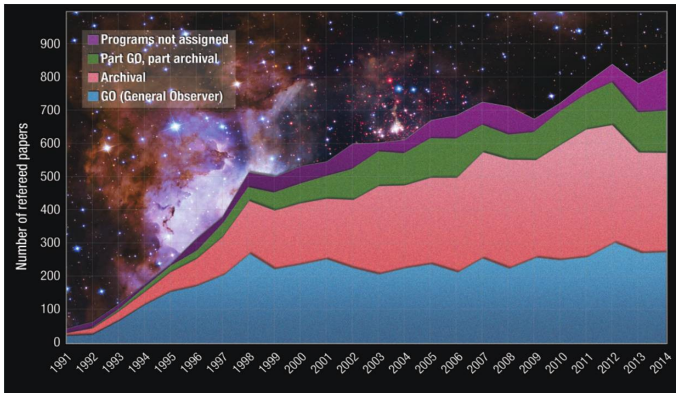

HST Newsletter: "*At the present time, approximately **half of the refereed publications** based on Hubble observations are derived purely **from archival data**, and, every year, this number is slightly higher than the number of publications based on new observations. .... the Hubble Archive has become a goldmine for the astronomical community....*"

Papers only from archives > 35%

Archive are:

- Data Management
- Data Curation
- Data Preservation

Archive are not:

- Data Sharing

Big Challenges …

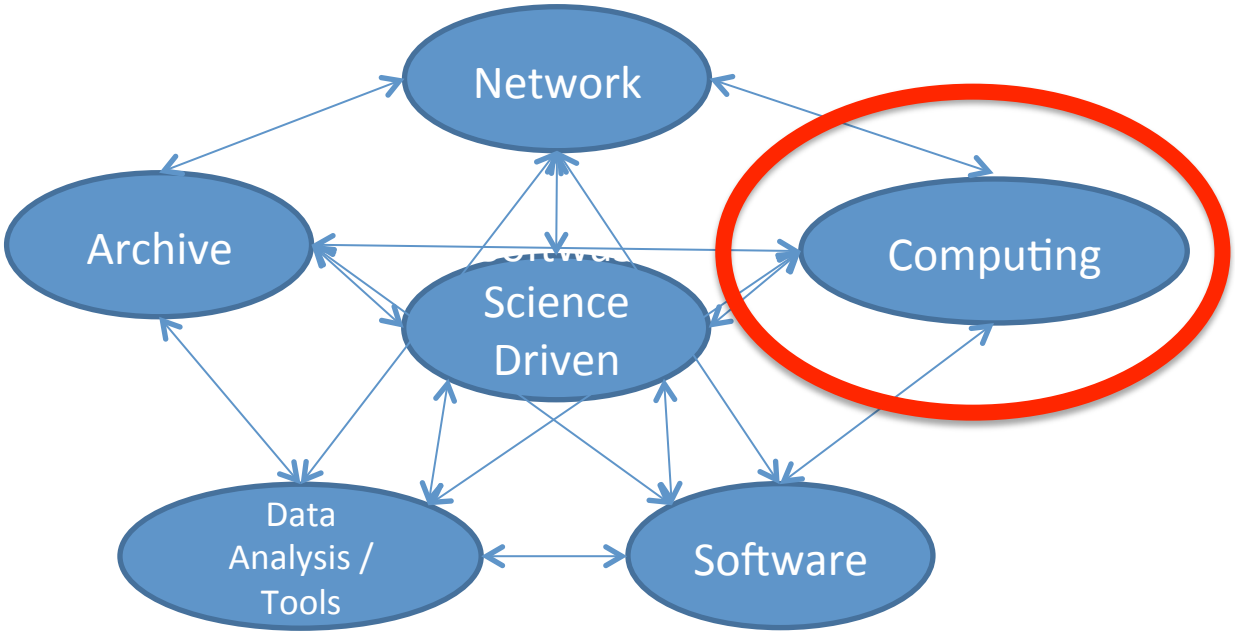

The “pyramid” of computing

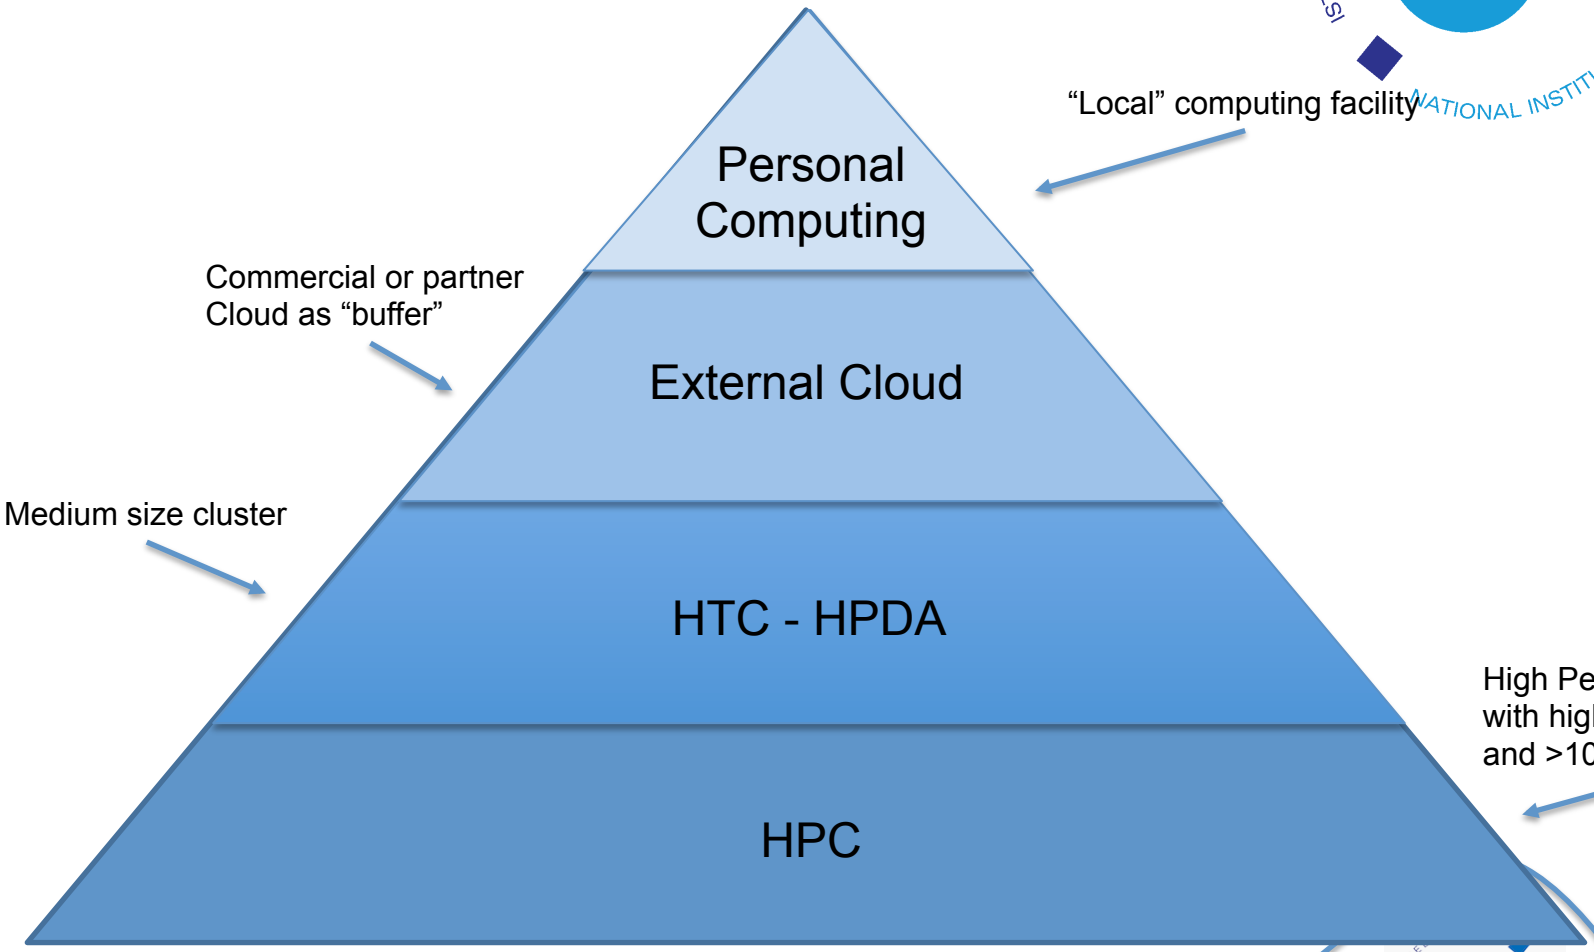

HPC @ INAF

What's happened in the future in HPC

INAF have a MOU with CINECA for using computing infrastructure (50 Million core/hour)

Cineca "sustainable" roadmap toward exascale

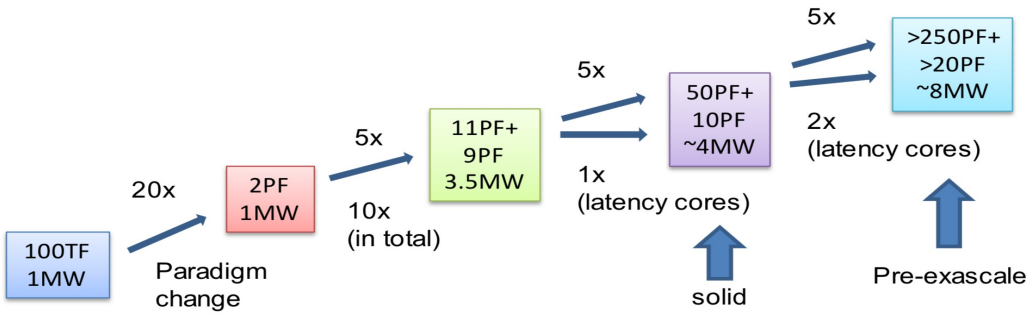

2009	2012/2013	2016/2017	2019/2020	2021/2022
IBM SP6 Power6	Fermi IBM BGQ PowerA2	Marconi Xeon + KNL	Marconi + PPI4HPC + ICEI (PPI-HB)	EuroHPC

Main evolution, from "classical" CPU to accelerators (GPU, TPU, ..)

Old HPC programming will not more usable with the new architecture

INAF - HTC
"Computing Facility"

INAF have some dedicated e-infrastructure for a projects:

- VST
- Planck
- LOFAR → **~ 100 TFlops**
- GAIA
- …

But have also **a generic e-infra**: **CHIPP**

Core	800 [1]	200	192 [2]
Mem/Core	6.4 Gby	4 Gby	2.6 GBy
Network	Infiniband 56 Gbps	Infiniband 10 Gbps	FastEthernet
Fast Storage	240 Tby	70 Tby	24 Tby [3]
Long Storage	> 16 Tby	> 16 Tby	None
CHIPP (%) Core- ours	> 40%	70%	90%

GCP: POC

9 Proof of concepts (PoCs) proposed by our scientists

1. **HTC** computing and software containerisation for **DIAMONDS** (M. Landoni)
2. **HTC** computing for **DIAMONDS** with Kubernetes (M. Landoni)
3. **GPU** computing for **Adaptive Optics** (M. Landoni in collaboration with OA Arcetri)
4. **HPC** computing for **GADGET** (G. Taffoni)
5. **HPC** computing for **Exoclimates** (G. Taffoni)
6. **HPC** computing for GAIA GSR Solver (U. Becciani, A. Vecchiato)
7. Computing for **ALMA** (M. Massardi, A. Giannetti, S. Burkutean)
8. Workflow execution for **GIANO@TNG** pipeline (A. Bignamini)
9. **Euclid** LE3 software in the Google Cloud Platform (D. Tavagnacco)

Google Cloud

Big Challenges …

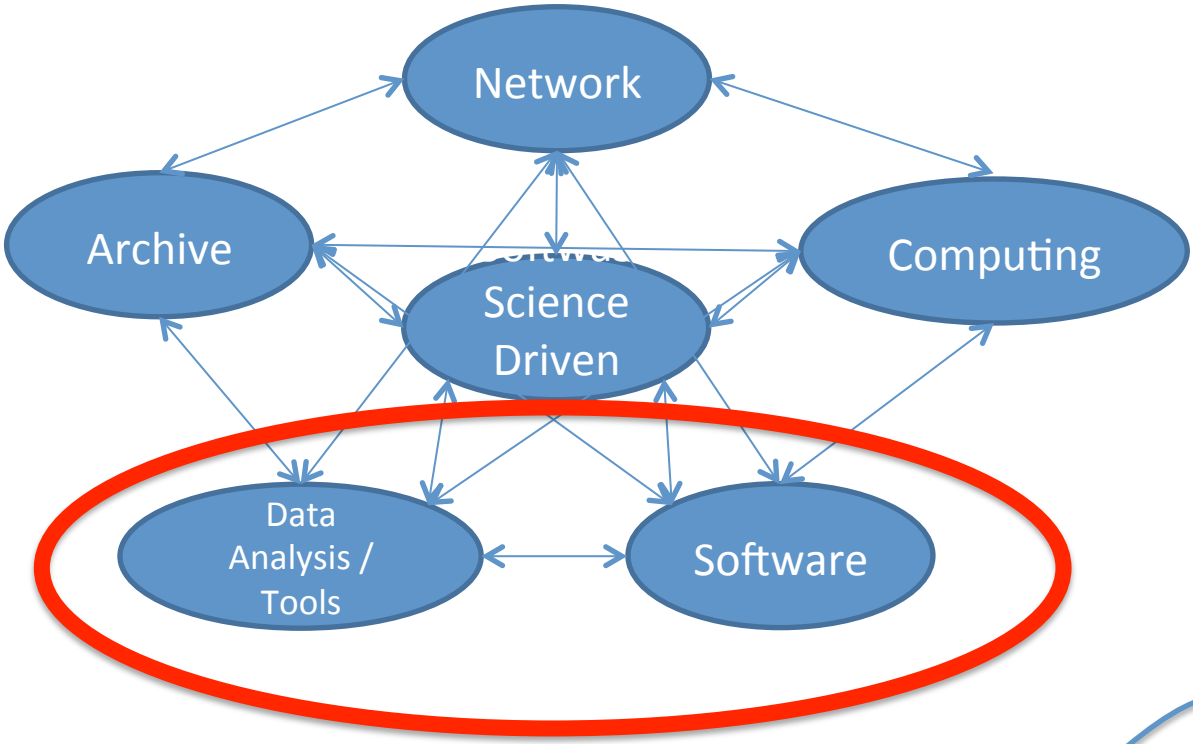

The “pyramid” of software

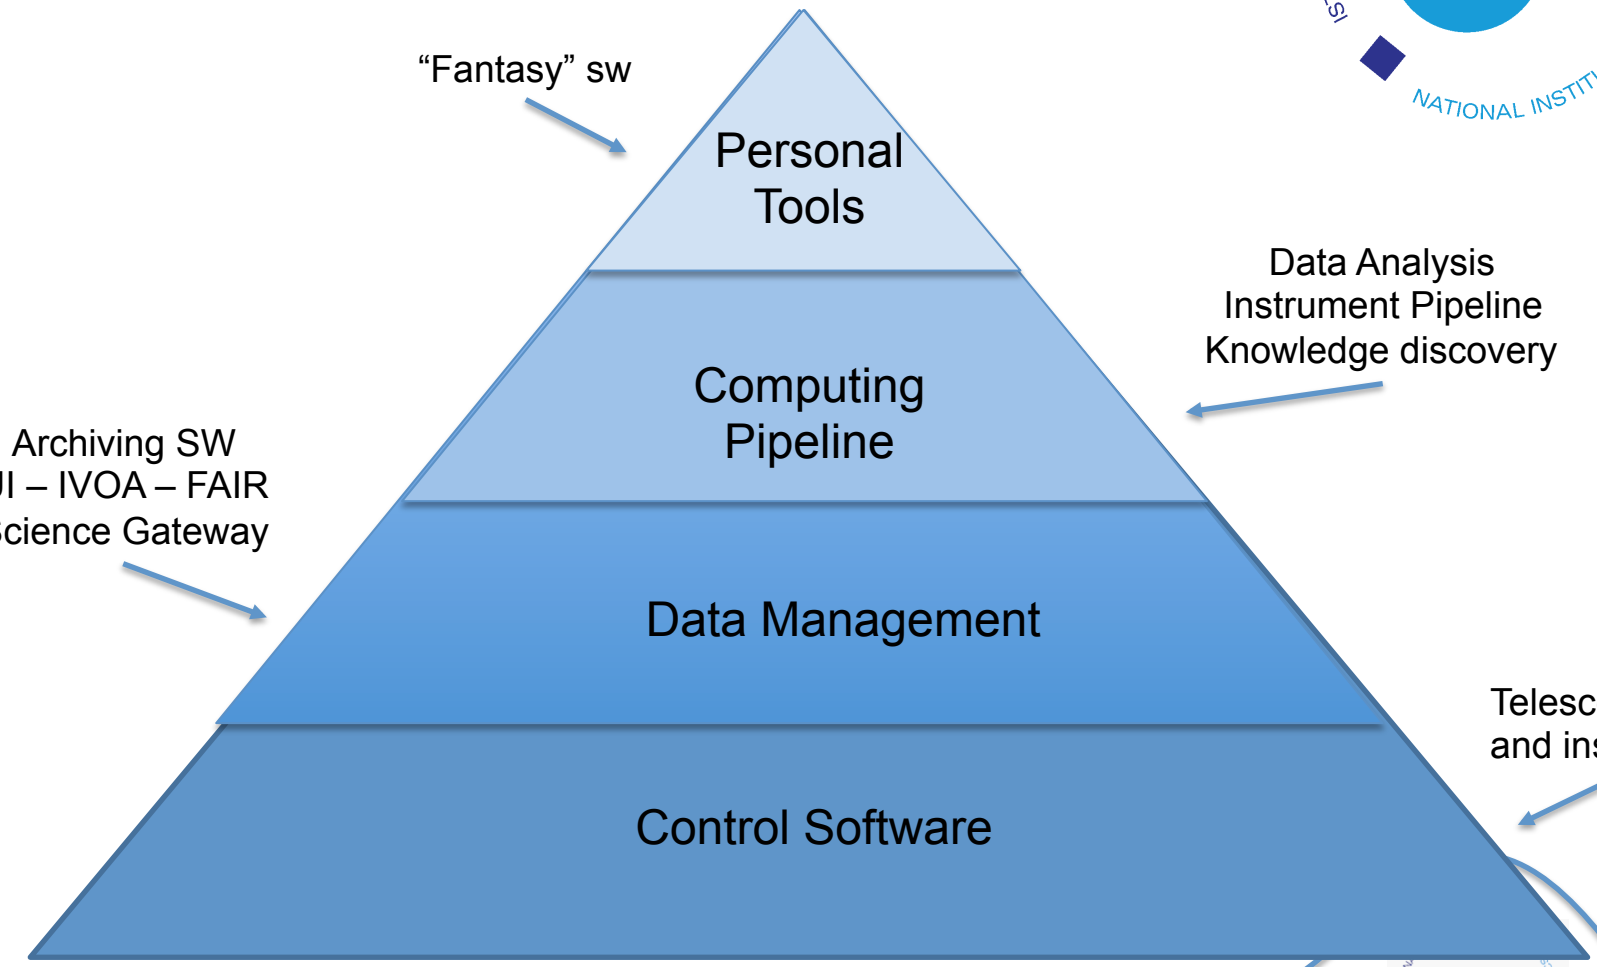

Intellectual Propriety

- **<u>INAF develop more than 100 Package Software,</u>**
  some released to the consortium, some Open, some… to a friends
- Open Source code should be release under Open Licenze:
  - **<u>GNU GPL v3</u>**

- Or must be released with a DOI - Digital Object Identifiers:
  (persistent, globally unique, resolvable)
- **<u>publications,</u> data, software**
- Resolvable by prepending https://doi.org/
- Citable (unambiguously)
- As a URL, avoids link rot
- Machine-readable when cited
- Repositories offer (linked) metadata (for humans and machines)

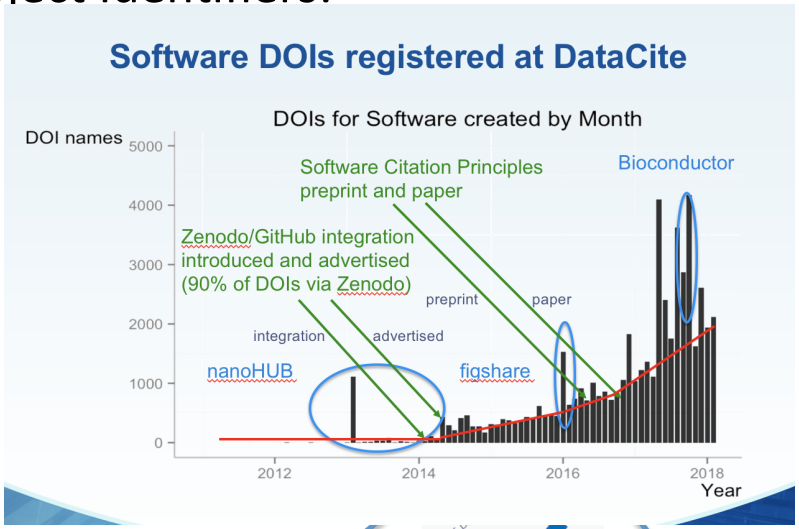

Yesterday, we have heard some of the software-related challenges for having SKA & precursors/path-finders properly and efficiently running

high italian involvement in those challenges

Some examples …

See Mauro Dolci's talk

Valentina Alberti , Matteo Canzari, Matteo Di Carlo ++ Franco Tinarelli, Cristina Knapic, …

- **Local Monitor and Control** (called **TM Services**): contains important TM services that support both OSO (Observatory Science Operations) and TMC (Telescope Manager Control):
- **TM Maintenance**
- **GUI**: The SKA UIs has been tackled starting from the user's perspective..
- **Authentication, Authorisation and Auditing**: activity supporting the whole of the SKA covering the Authentication and Authorisation of all users accessing the resources of the SKA Observatory and Telescopes

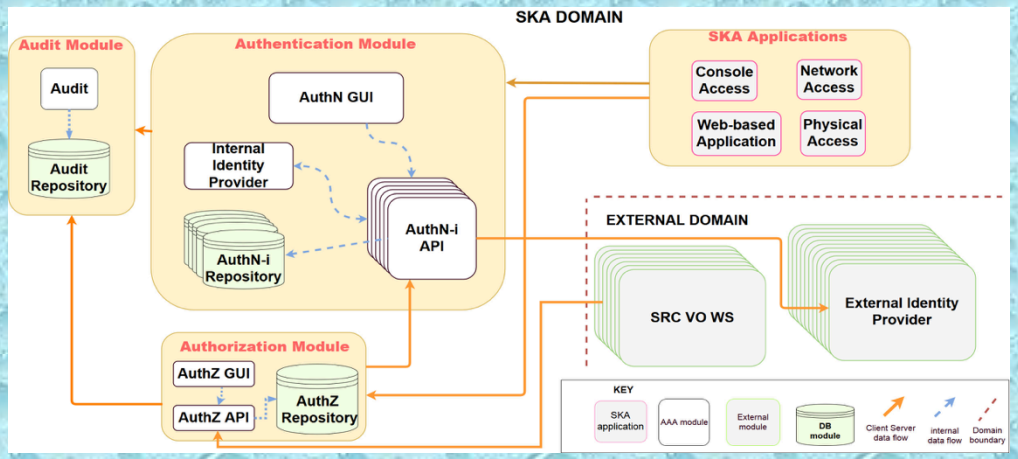

Architecture of **the Autentication** system (AuthN), **Authorization** (authZ) and Auditing, in connection with external and internal services at SKA (***Criistina Knapic*** et al.)

Storage and Databases for Observation Data Archive (ODA)

Ingestion system for the Italian Radio data (but a possible template for a precursor of the SKA radio data portal) in the NADIR RADIO archive and Portals for Medicina, Noto end SRT.

(*Criistina Knapic* et al.)

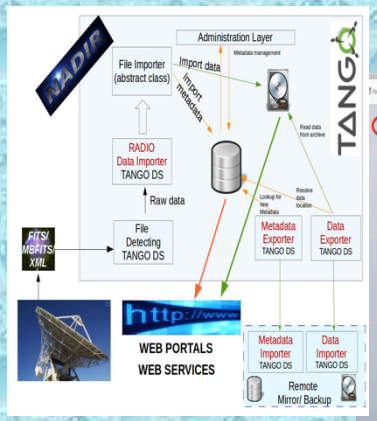

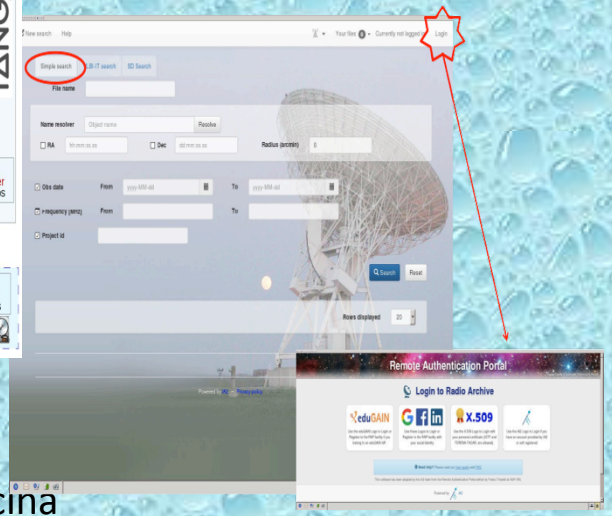

Interaction with Precursors

- Interview with operator from Meerkat, Alma, Lofar and Medicina
- Architecture evaluation for MeerKAT and GMRT

Ongoing work – Bridging

- Working with Safe approach in two different team
- Continuity of the work made during the preconstruction

See Mauro Dolci's talk

Valentina Alberti , Matteo Canzari, Matteo Di Carlo + ...

CSP – PST – PSS

See Comoretto's talk

.... etc ...

See next Taffoni & Becchiani's talk

...now a short summary of some of the software-related challenges for

exploiting SKA & precursors/path-finders science data

particular attention to the (large) Italian involvement in the mentioned challenges

Some software-related challenges to face for fully exploiting science data

- Huge data volumes
- Real-time or quasi-real-time data reduction
- Unprecedented number of sources per pointing to extract and characterize
- Data visualization
- Calibration (especially for Low freq array)
- Polarization Calibration
- RFI excision in presence of very large number of frequency channels
- Scheduling optimization
- Data archiving
- .....

A sample of challenges with Italian involvement

Instrument	Some key software activities developed within INAF
LOFAR	Data working group Pipelines Computing Code optimization & profiling Porting on exascale machines
ASKAP	Caesar: source extraction & parametrization Algorithms to destripe single-dish images Source extraction from combined IR + Radio
MeerKAT	HI/continuum data analysis for inteferometric data Pulsar pipelines and schedule optimization
uGMRT	Optimization of existing pipelines
eMERLIN	Optimization of existing pipelines – combination with JVLA data
JVLA	Optimization of existing pipelines
Other projects	VisiVo: big astronomical data 3D visualization Distributed data and computer center for SKA

Giuliano Taffoni

SKA & LOFAR teams & activities (led @ OATS)

- **LOFAR – Data working group**
  - Design and implementation of LOFAR e-Infrastructure that involves OATs, OACt, IRA and UniTO
  - Definition and coordination of LOFAR pipelines tests
  - Definition of policies for accessing computing resource
  - Definition of usage policies to optimize computing resources
- People involved: ***Gianmarco Maggio, Sara Bertocco + A. Bonafede, A. Botteon***
- **LOFAR pipelines**
  - Porting, testing and tuning of Pre-FACTOR and FACTOR (collaboration with OACt, IRA)
  - Testing and analysis of the "new" direction-dependent self-calibration pipeline
  - Testing storage performances and CPU/RAM performance
  - GOAL: estimate computing and storage requirements on single server and clusters
  - GOAL: knowledge acquisition to support community and on long term contribute to SKA ESDC
- People involved: ***Sara Bertocco, Gianmarco Maggio, Claudio Vuerli, David Goz + A. Bonafede, A. Botteon***

Giuliano Taffoni

SKA & LOFAR teams & activities II (led @ OATS)

- **LOFAR computing**:
  - Implementation of LOFAR cluster partition (queue, authentication, authorization, accounting)
  - Optimization of storage facility to maximize pipeline performance
  - Test of Parallel Filesystem and Object storage
  - GOAL:
    - provide an optimized platform for data reduction and analysis for LOFAR
    - knowledge acquisition to support community and on long term contribute to SKA ESDC
- People involved: ***Sara Bertocco, Gianmarco Maggio***
- **Code Optimization and profiling**
  - Algorithm optimization for Intel and Arm CPU and Clusters
  - Code profiling based on “papi” libraries and proprietary tools
  - Energy profile of codes and algorithms
- People involved: ***Luca Tornatore, David Gox***

Giuliano Taffoni

Exa-Scale projects

- Porting of LOFAR pipeline on exascalable machines
  - Porting on Arm platforms
  - Testing Exascale system-software for parallelizing pipelines (OmpSs or GPI)
  - Testing new accelerators (GPU or FPGA)
- People involved: ***David Goz, Luca Tornatore***
- Preliminary study of distributed data and computing center for SKA
  - Work coordinated by IA2 to study a possible architecture of a distributed environment based on standards for authentication and authorization, access to computing resources and data.
  - Collaboration with IRA (***Tinarelli, Bedosti***) and IA2 (***Knapic***)
- People involved: ***Sara Bertocco.***

ASKAP, LOFAR, SKA team @ OACT

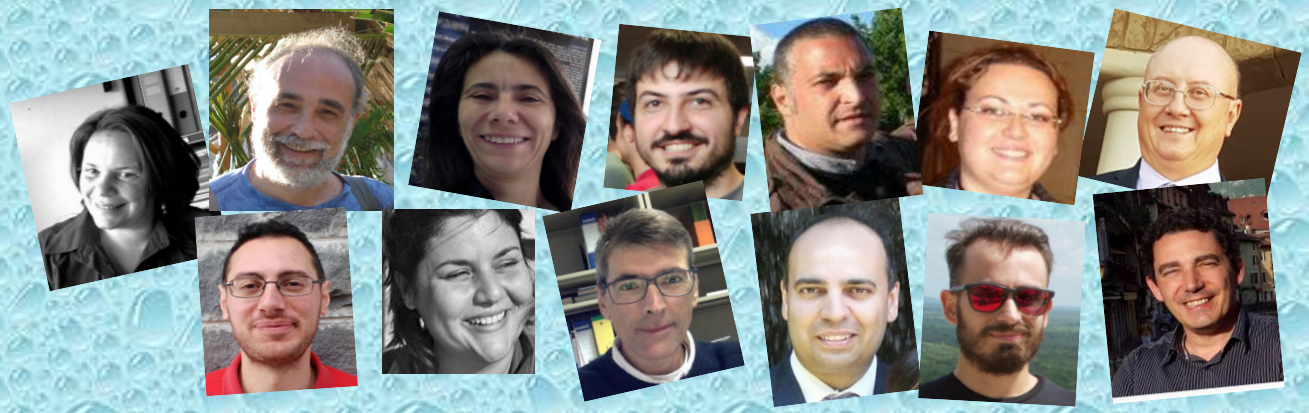

Grazia Umana
Corrado Trigilio
Paolo Leto
Carla Buemi
Milena Bufano
Francesco Cavallaro
Adriano Ingallinera
Francesco Schillirò
Ugo Becciani
Alessandro Costa
Eva Sciacca
Fabio Vitello
Simone Riggi

- **Radio & ICT groups working together since 2012: Successful initial experience for Competence Center**
  - 13 people involved in different projects (see below): 10 TI + 1 TD + 2 AdR
  - Software expertises: data reduction software, HPC/HTC, Visual analytic & Virtual Reality, ...

- **OACT in SKA & precursors (ASKAP, LOFAR) science groups**
  - SKA "*Our Galaxy*" KSP (co-lead: ***G. Umana***)
  - SKA AENEAS WP 3.3 (lead: ***U. Becciani***) & WP 5.3 (lead. ***A. Costa)***
  - ASKAP EMU Galactic Plane KSP (co-lead: ***G. Umana***)
  - ASKAP EMU Dev. Projects "GP Imaging & Diffuse Sources" DP4 (leads: ***S. Riggi, F. Cavallaro),*** DP6 & Parkes GP survey
  - ASKAP SCORPIO Early Science Project (lead: ***C. Trigilio***)
  - LOFAR IT data working group (local ref.: ***A. Costa***)
  - LOFAR Transient & Exoplanets KSP

Software for data reduction, post-processing and visualization (led @ OACT)

- **CAESAR**: A software tool for automated compact & extended source extraction and parametrization from large interferometric maps
  - Ref: S. Riggi, MNRAS 460, 1486 (2016), https://github.com/SKA-INAF/caesar.git
  - Status: tested with ASKAP/ATCA SCORPIO data, testing with simulations ongoing
- **VisIVO: Via Lactea Visual Analytics**. A software tool for big astronomical data 3D visual inspection and analysis for sources characterization.
  - Ref: F. Vitello et al 2018 PASP 130 084503 (2018), https://www.ict.inaf.it/gitlab/fabio.vitello/ViaLacteaVisualAnalytics.git
  - Status: tested with Via Lactea project data. EOSC Pilot: Science demonstrator. Application foresen for ASKAP Radio + IR map
- **Algorithms to destripe single-dish images and** combine to interferometric data
  - Destriping and feather combination status: tested with GBT data.
  - Ref: Ingallinera et al. (2014b)
  - Advanced combination status: testing with simulations ongoing, testing with SCORPIO Parkes +ASKAP data expected in early 2019.
- **Algorithms for extraction of source object** (e.g. bubbles) parameters from IR & radio combined maps for classification purposes
  - Status: tested with SCORPIO ATCA/ASKAP maps, Herschel & WISE IR maps
  - Ref: F. Bufano, MNRAS 473, 3 (2018), A. Ingallinera, in prep.

Evaluation and tuning of precursor pipelines (led @ OACT)

- **ASKAP & LOFAR** pipeline porting, testing and tuning on different infrastructures
  - Prototypes being tested for different Infrastructures: INAF OACT/OATS clusters, CINECA, Google/ OpenStack clouds
  - LOFAR pipelines
    - Calibration pipeline being tested with real data: pre-factor, init-sub, factor
    - Joint collaboration with IRA BO + OATS
    - See **Becciani/Taffoni** presentations
  - ASKAP pipelines
    - imaging continuum pipeline being tested with SCORPIO Early Science data
    - fine tuning of calibration & imaging parameters required in the Galactic Plane
    - See **F. Cavallaro** presentation

- **Multiple goals**
  - estimate pipeline computing requirements for LOFAR IT computing infrastructure design
  - validate the design of SKA ESDC computing model from precursor realistic use cases
  - being prepared to the upcoming ASKAP EMU & LOFAR surveys (short-term) and SKA KSP surveys (long-term)

MEERKATHI

https://github.com/ska-sa/meerkathi

MeerKAT team: led @ OAC

A pipeline for interferometric data reduction
Based on best available software connected together using Stimela
(MeqTrees, AOflagger, WSClean, SoFiA, CASA, …)

Initial motivation
MeerKAT Fornax Survey (Serra et al.)
Cutting-edge data processing techniques (Smirnov et al.)

Used on HI/continuum data from
MeerKAT, VLA, GMRT, WSRT, APERTIF, ASKAP

Hugo, Józsa, Makhathini, Ramaila, Smirnov, Thorat (SARAO, Rhodes); Kleiner, Maccagni, Molnar, Ramatsoku, Serra (INAF - Cagliari); De Blok (ASTRON, Kapteyn); Kamphuis (Bochum)

SKA AFRICA
SQUARE KILOMETRE ARRAY

RHODES UNIVERSITY
Where leaders learn

INAF
ISTITUTO NAZIONALE DI ASTROFISICA
NATIONAL INSTITUTE FOR ASTROPHYSICS

OAC
Osservatorio Astronomico di Cagliari

1.4 GHz image of Fornax A made with MeerKATHI by F. Maccagni
single pointing, MeerKAT-40 SKARAB-4k, 10 MHz band

MEERTIME

https://www.meertime.org

MeerKAT team: led @ OAC

TRAnsients and PUlsars with MeerKAT

https://www.trapum.org

A real-time pipeline for pulsar data (timing and search) reduction.

Accounting for the larger number of channels than usually before: RFI zapping

Performing real time search

Optimization of the observing schedule to exploit beam-forming and very rapid switching among the targets

Alessandro Ridolfi, Federico Abbate, Marta Burgay, Andrea Possenti, Ewan Barr, Ryan Shannon + Meertime team, Trapum team

SKA AFRICA
SQUARE KILOMETRE ARRAY

INAF
ISTITUTO NAZIONALE DI ASTROFISICA
NATIONAL INSTITUTE FOR ASTROPHYSICS

OAC
Osservatorio Astronomico di Cagliari

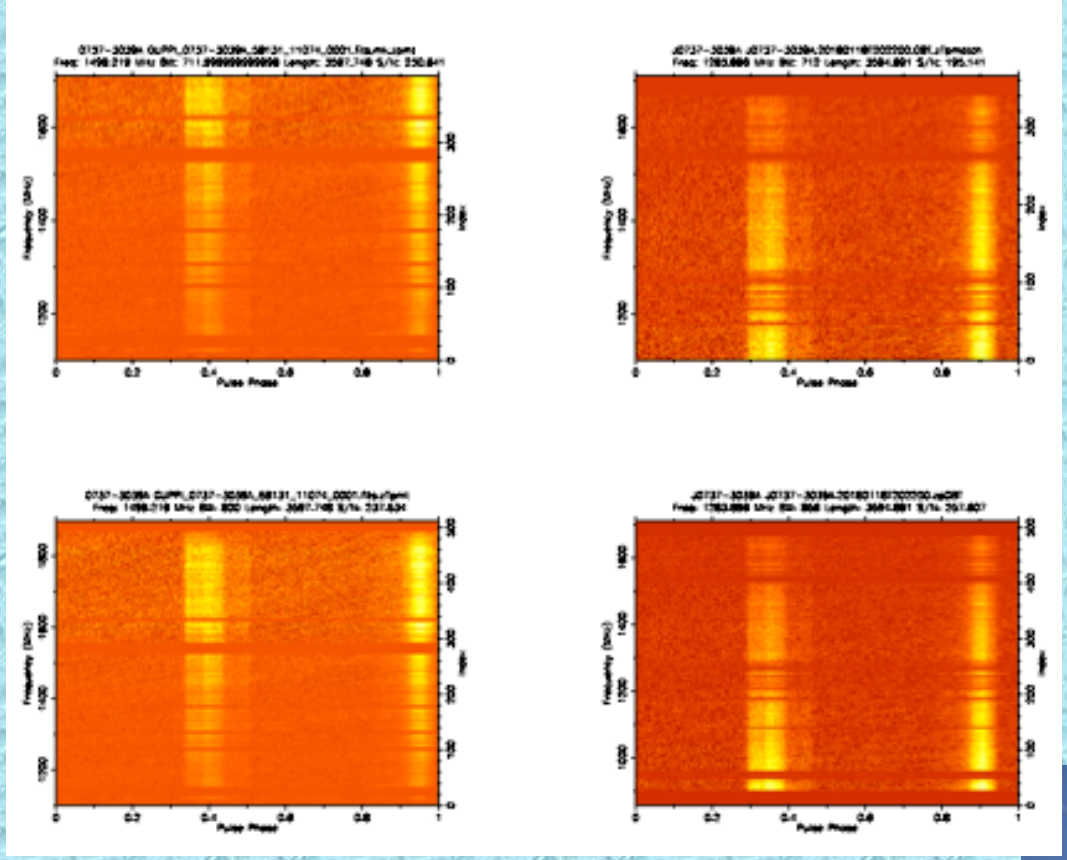

1.4 GHz images of the Double Pulsar collected at MeerKAT-16 and GBT

Federico Abbate

Frequency band	GBT	MeerKAT
Matched band	224	195
Full band	237	257

JVLA, uGRMT, MeerKAT, LOFAR, ASKAP, eMERLIN teams and activities (led @ IRA)

JVLA

Marco Bondi (surveys - COSMOS field)
Monica Orienti (radio galaxies, jelly fish galaxies)

- pipelines developed by NRAO-CALTECH for JVLA observations are used only in the initial stages of calibration (a-priori calibration and flagging) for continuum data;
- polarization data are not included in any pipeline, so that is done from scratch;
- self calibration is standard and is usually carried out using AIPS
- imaging is done in CASA to best account for the broad band
- weeks to months are necessary to get the final result, depending on the size of the dataset, and on the number of pointings. For the 3 GHz COSMOS project (384 hours, 64 pointings) a full FTE for a year has been necessary.

uGMRT

Tiziana Venturi, Daniele Dallacasa & Beatrice Terni de Gregory

(galaxy clusters and radio galaxies in clusters)

- initial calibration is done using "flagcal", the pipeline developed by NCRA (not public)
- self calibration and further editing are best done in AIPS using a standard approach and calibrating sub-bands individually
- imaging is done in CASA to account for the broad band
- weeks to months are necessary to get a final satisfactory image

MeerKAT

Gianni Bernardi, T. Venturi, Daniele Dallacasa & Beatrice Terni de Gregory

(galaxy clusters and radio galaxies in clusters)

- initial calibration done using the pipeline developed by Cape Town & INAF-OACa
- self calibration and imaging are done with a standart approach using CASA
- sub-bands are self calibrated and a final imaging is made
- weeks to months are necessary to get a final satisfactory image

LOFAR

Andrea Botteon (& LOFAR galaxy cluster people at IRA)

- data reduction made using the pipelines developed within LOFAR
- Initial steps with Pre-Factor, which performs the initial direction independent calibration. Pre-factor is in practice two different pipelines: Pre-Facet and Initial-Subtraction. The input of the whole process is the raw dataset. The full Pre-Factor stage requires a few days of computing time, and considerably reduces the size of the dataset (from several TB to few hundred GB).

- The oputput of Pre-Factor is the input of the second step in the data analysis, which is carried out using DDFacet, the pipeline developed by Tasse, Hardcastle and Shimwell. DDFacet produces images of the full primary beam. It takes about a week to provide the final image, but intermediate steps can be inspected. The data output is of the order of 2 TB. It reaches the best rms available these days.

- The above refers to one of the typical targets of LoTSS (120-168 MHz). At present the effort concentrates on the extraction of small fields around the targets of interest (of the order of 20x20 arcmin^2) from the full FoV of the pointing (6x6 deg^2). After the extraction, some more calibration and re-imaging are performed, and the final image is ready for the analysis. This latter file is small (few GB) and the noise in the small field of interest can be as good as 60-70 muJy/b.

eMERLIN + JVLA

- **eMERGE legacy project**
  - 1.4 and 5 GHz RC coverage of GOODS-N region, sub-)uJy rms,sub-arcsec resolution (0.05-0.2 arcesc): 900 hr allocated at eMERLIN
  - reduction packages: AIPS, CASA, wsclean + AOflagger for RFI flagging
  - ***I. Prandoni***: lead of 5 GHz survey
  - other people involved @INAF-IRA: ***M. Bondi, D. Guidetti***
  - data reduction of 5 GHz data @IRA:
    - dedicated multi-core computer (Merlino), data storage: 20 Tb (external funding)

ASKAP

- **EMU legacy project**: RC 1.4 GHz all sky survey (Norris+2011)
  - 10 uJy rms, 10 arcsec resolution
  - data reduction pipeline: ASKAPsoft
  - ***I. Prandoni***: lead of RQ AGN Working Group + other people involved: RQ AGN WG (~40 people)
  - current focus GAMA-23 60 sq. degr. region
    - debugging of ASKAPsoft pipeline + image/source catalogue validation
  - EMU management request to install and run pipeline at other host institutes

    → this needs dedicated HPC (best if coordinated at INAF level)

MEERKAT

- **MIGHTEE key project (Jarvis+2017)**

  -1.4 GHz coverage of 3 well-known fields (for a total of 20 sq. degr), ~1 uJy/b rms, 6" resol

  - dedicated pipeline at IDIA

  - survey starting now, leadership and level of data processing contribution under discussion

LOFAR

- **3-tier survey**
  - dedicated pipeline (now installed at INAF institutes)
  - ***I. Prandoni***:
    - lead of LH tier-2 field (~30 sq. degr, 25 uJy/b rms)
    - lead of source extraction and characterization; source counts for first data release (see dedicated A&A issue, in press)
    - lead of multi-band forced photometry for LH for deep fields data release (expected for second-half of 2019)
    - work done in collaboration with 2 postdocs hired at INAF-IRA: ***M. Brienza; M. Bonato***

Thanks!