THE LOW FREQUENCY ARray

- A TB/h data existing machine -

Gianfranco Brunetti

INAF IRA ISTITUTO DI RADIOASTRONOMIA

Outline

- LOFAR
- Science
- Data flow and computational challenges
- The Italian Roadmap

THE LOw FREQUENCY ARray

Giant digital aperture array radio telescope opening up a new window in the electromagnetic spectum at low radio frequencies
- The largest (area & dataflow) pathfinder toward the SKA -

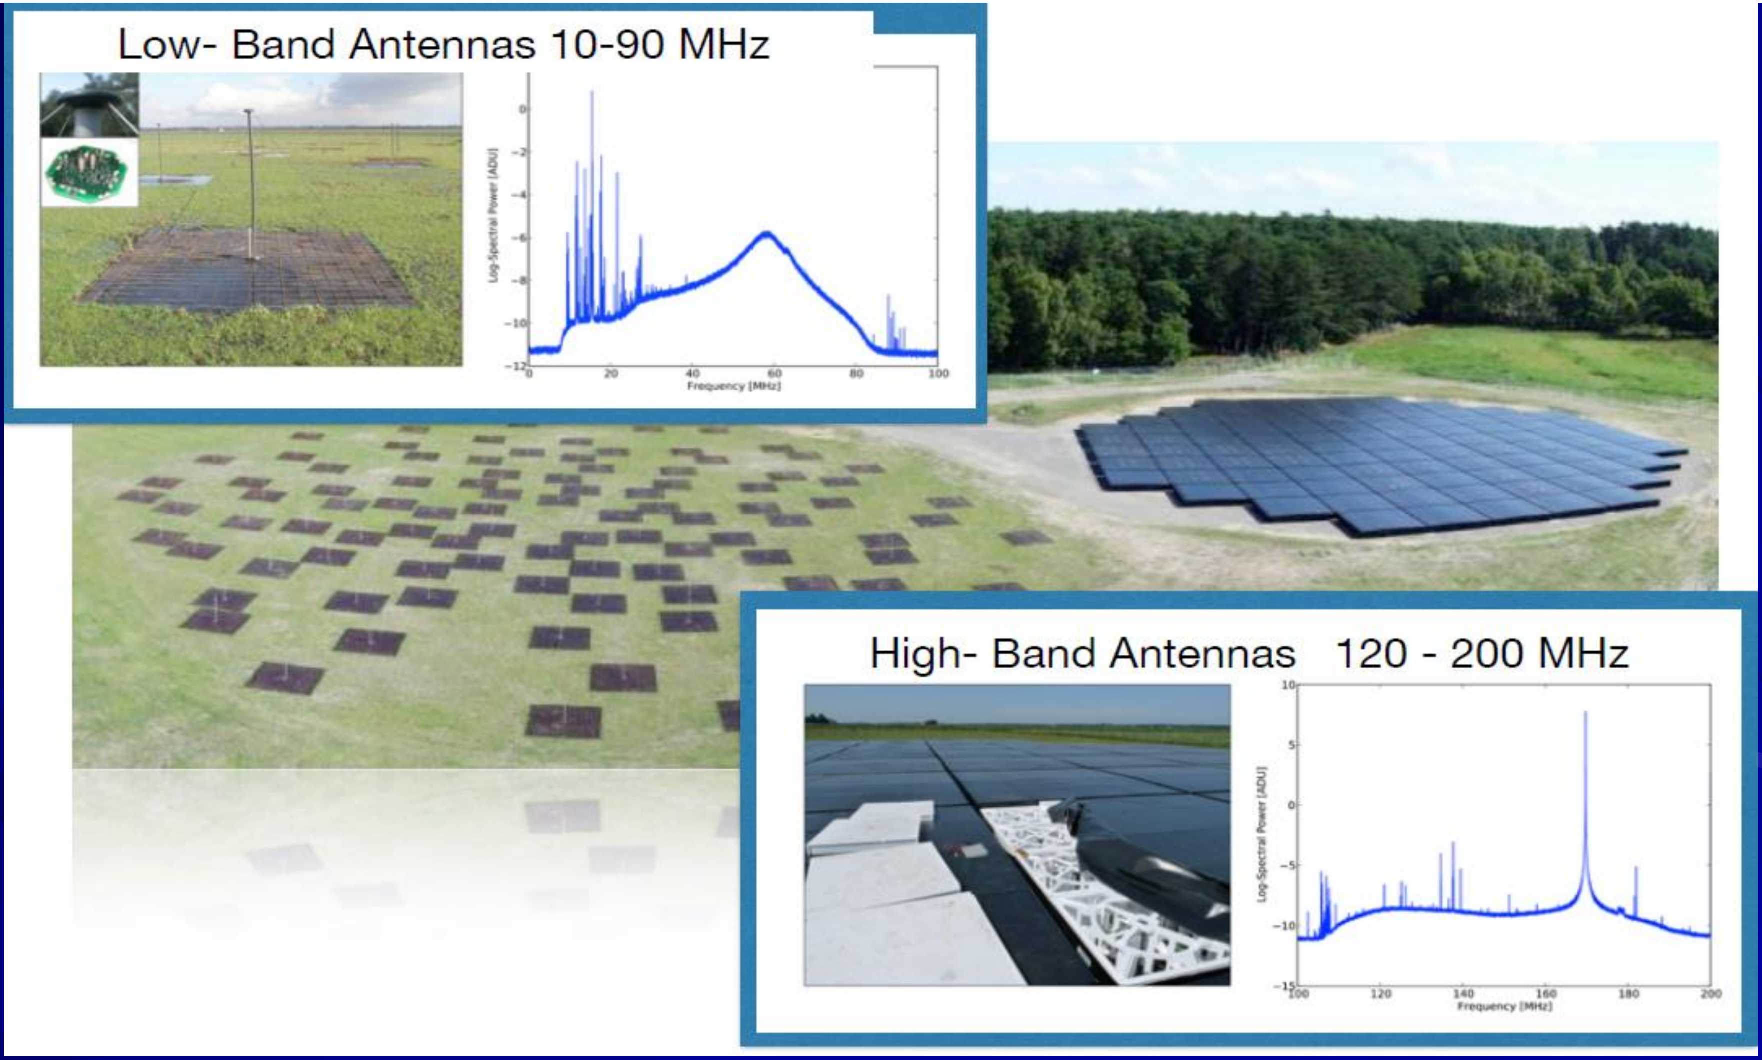

LOFAR Stations






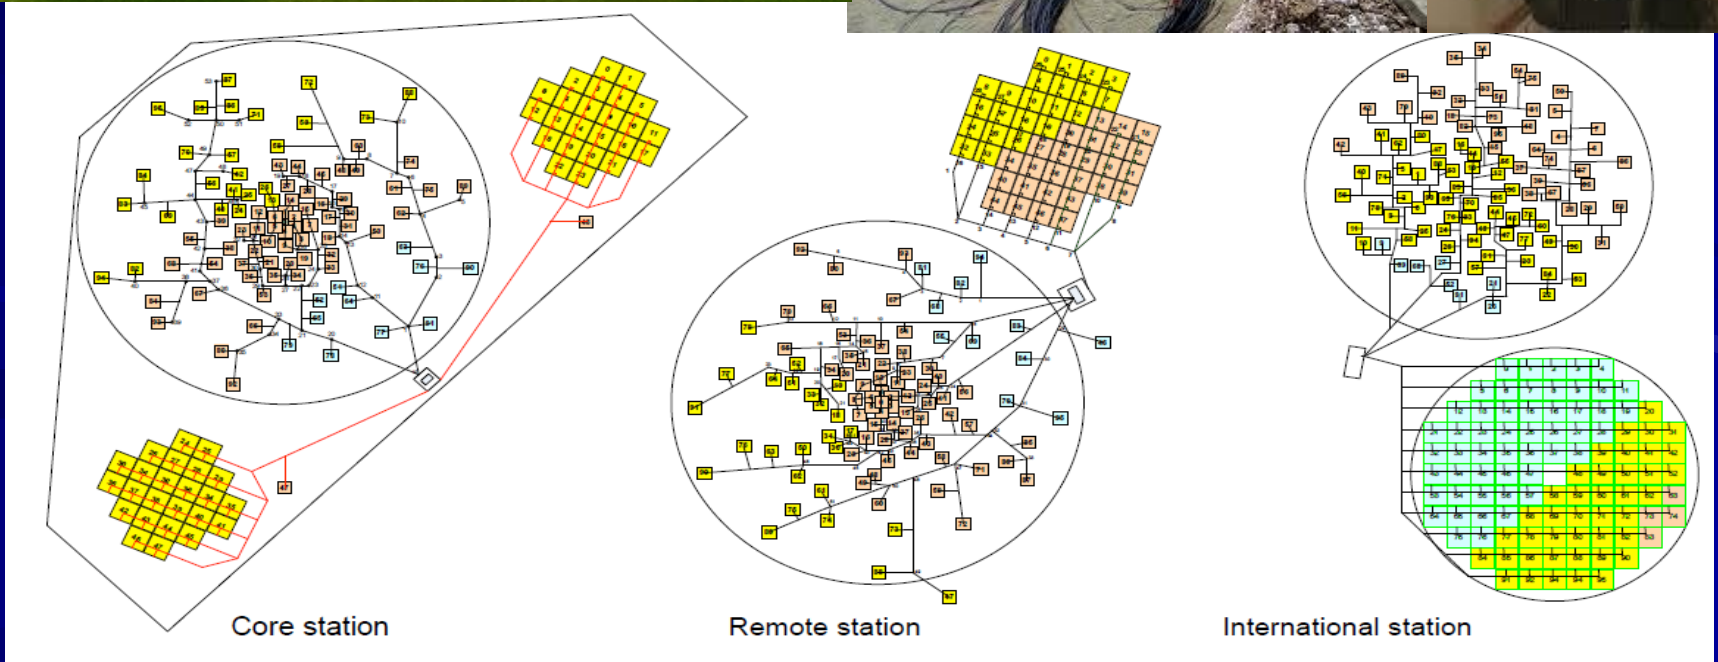

3 scales :

(1)Core ''Superterp'' short baselines (deg scales, arcmin res)
(2) Remote NL Stations 10-100 km baselines (5-10 arcsec res)
(3) International Stations 100-1000 km baselines (sub arcsec)

3 scales :

(1)Core ''Superterp'' short baselines (deg scales, arcmin res)

(2) Remote NL Stations 10-100 km baselines (5-10 arcsec res)

(3) International Stations 100-1000 km baselines (sub arcsec)

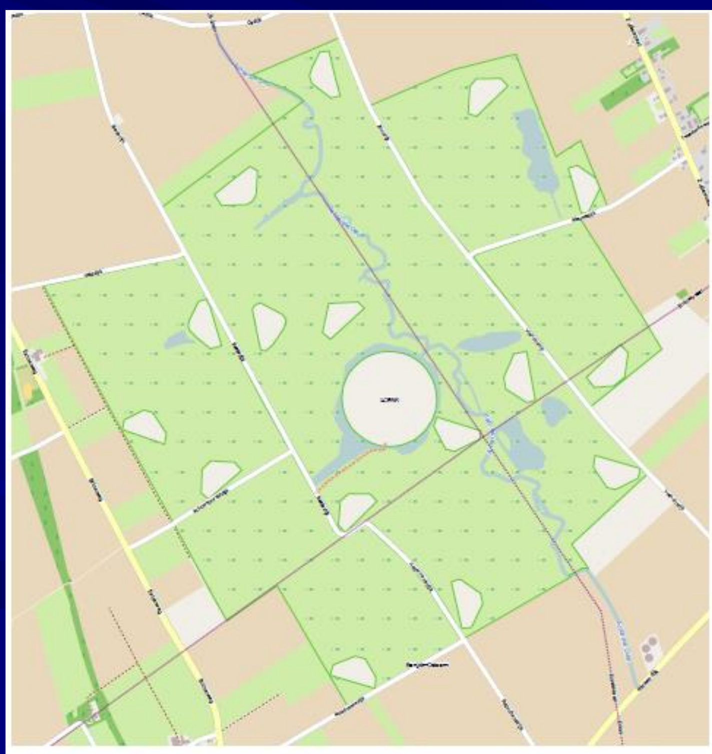

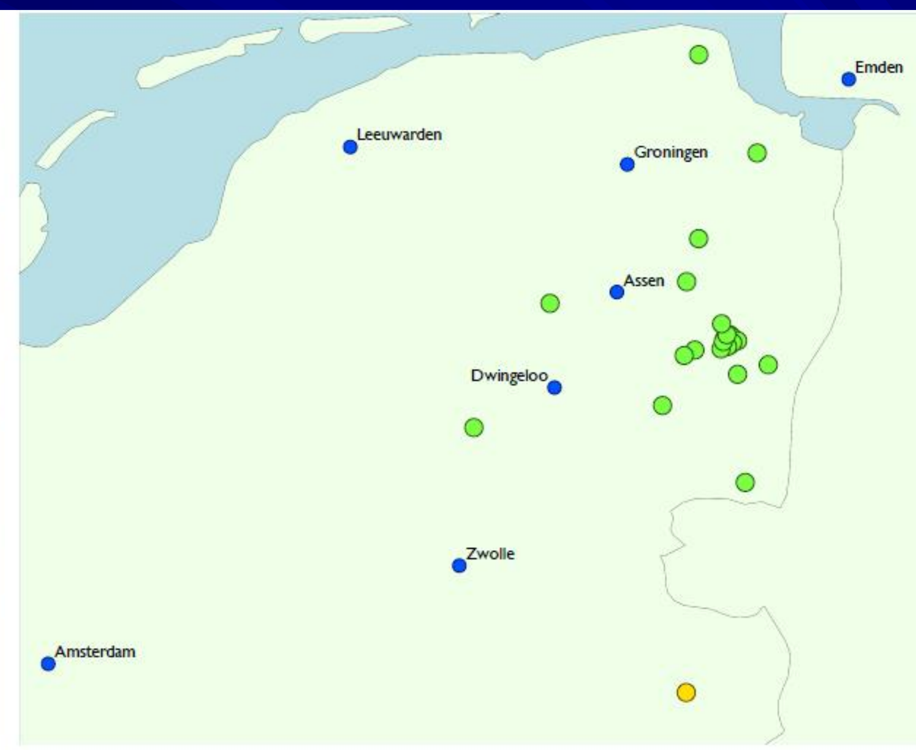

International LOFAR Telescope (ILT)

- Most distant stations 1500 km apart
- 1 new station funded in Latvia
- Ongoing negotiations with INAF/IT

The LOw Frequency ARray

The International LOFAR telescope (ILT) consists of an interferometric array of dipole antenna stations distributed throughout 8 EU Countries: NL, Germany, France, Poland,UK, Sweeden, Ireland, Latvia (50+ Meuro construction +Running costs)

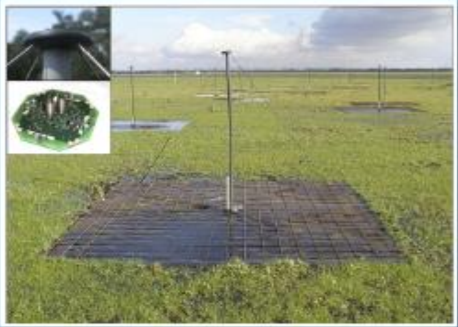

- Operating frequency is 10-250 MHz
- Low band antenna (LBA; 4800 dipole pairs, 96 LBA per station, Area ~ 75200 m^2; 10-90 MHz)

- High Band Antenna (HBA; 47616 dipole pairs, 48/96 tiles per station in NL/EU, Area ~ 57000 m^2; 110-250 MHz)
- Several observing modes (imaging, BF, BF+IM, TBB)

- 96 MHz bandwidth (can be split to perform simultaneous beamforming in different directions)

THE LOW FREQUENCY ARray

The International LOFAR telescope (ILT) consists of an interferometric array of dipole antenna stations distributed throughout 8 EU Countries: NL, Germany, France, Poland,UK, Sweeden, Ireland, Latvia
(50+ Meuro construction +Running costs)

- Operating frequency is 10-250 MHz
- Low band antenna (LBA; 4800 dipole pairs, 96 LBA per station, Area ~ 75200 m^2; 10-90 MHz)

- High Band Antenna (HBA; 47616 dipole pairs, 48/96 tiles per station in NL/EU, Area ~ 57000 m^2; 110-250 MHz)
- Several observing modes (imaging, BF, BF+IM, TBB)

- 96 MHz bandwidth (can be split to perform simultaneous beamforming in different directions)

LOFAR KEY SCIENCE PROJECTS

- a revolution in radio astronomy -

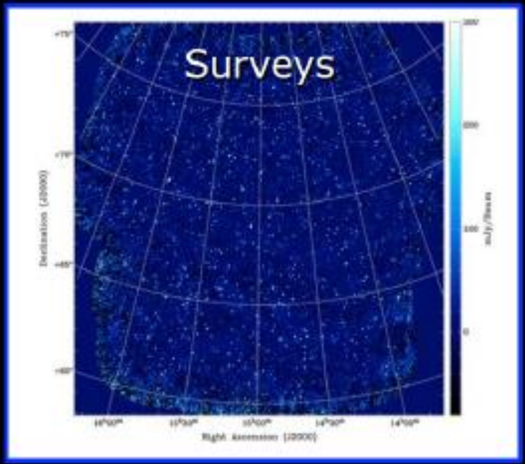

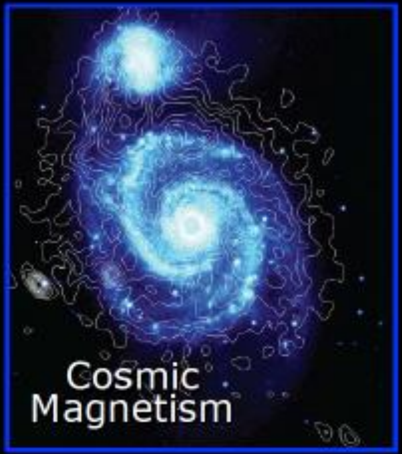

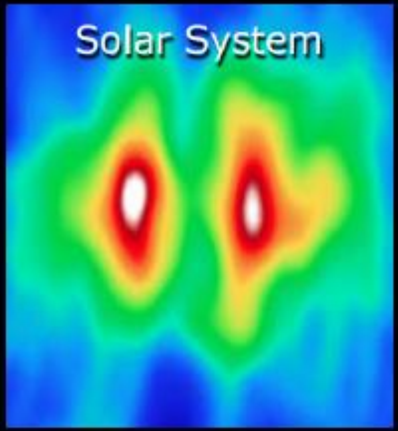

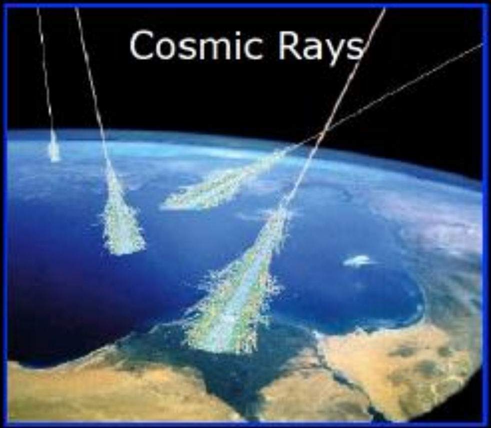

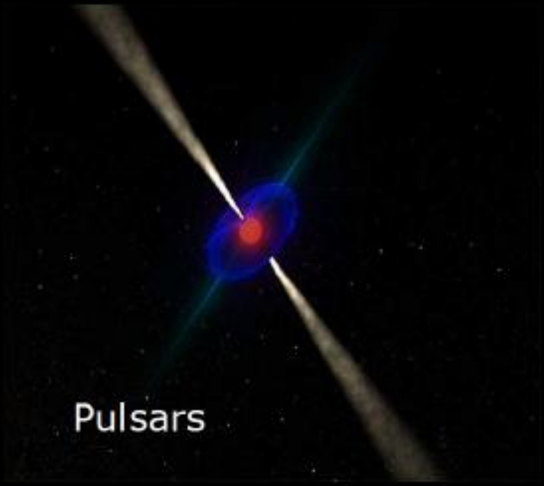

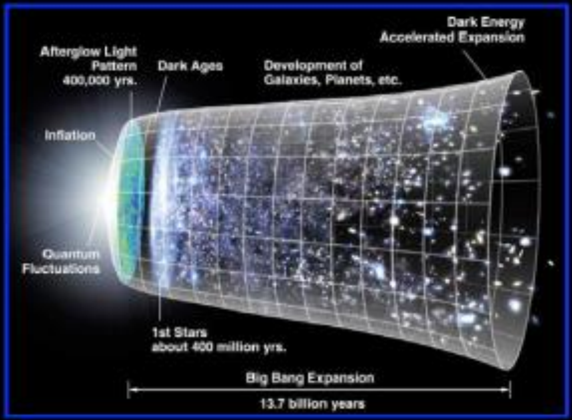

Epoch of Reionization

BIG Challenges with data calibration and analysis

Our enemy: the ionosphere

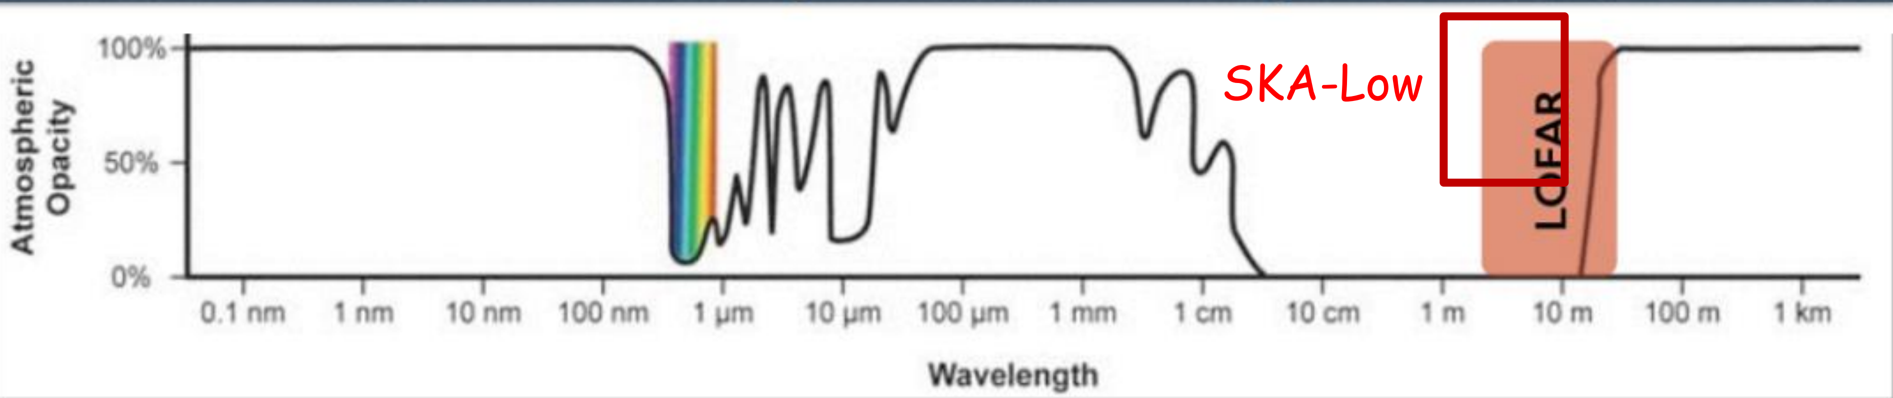

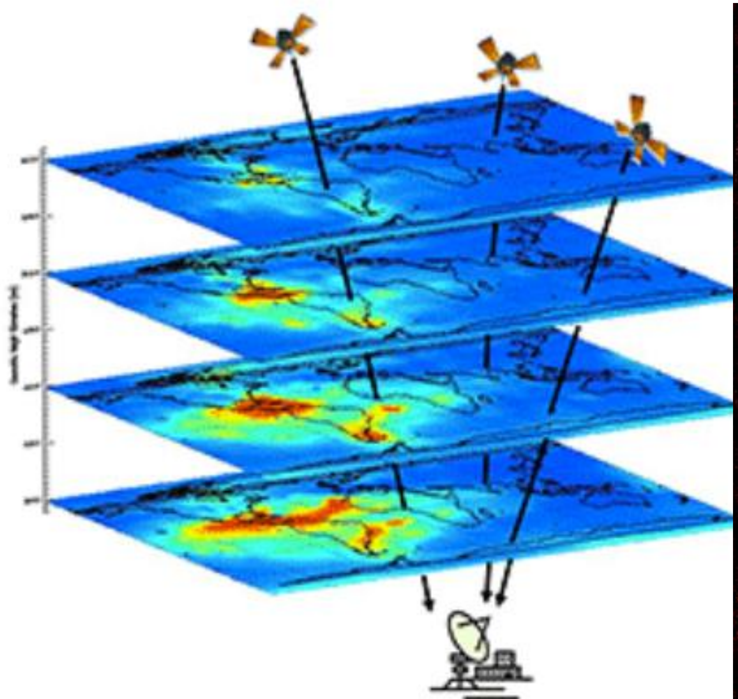

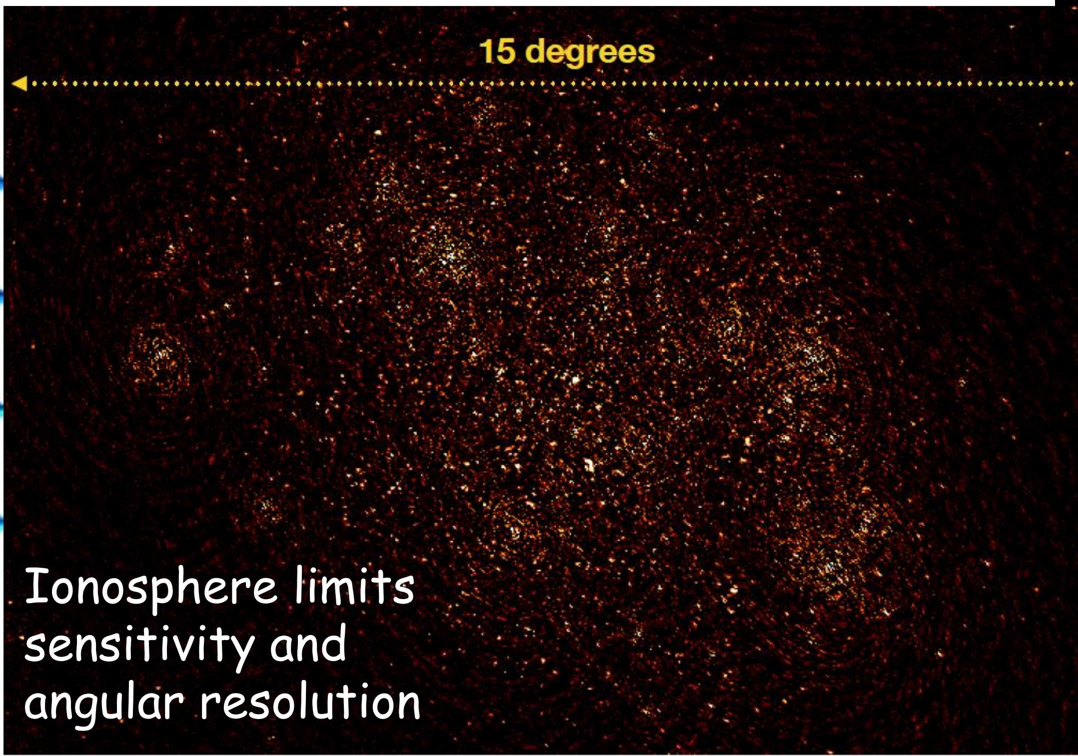

Ionosphere limits sensitivity and angular resolution

BIG Challenges with data calibration and analysis

Our enemy: the ionosphere

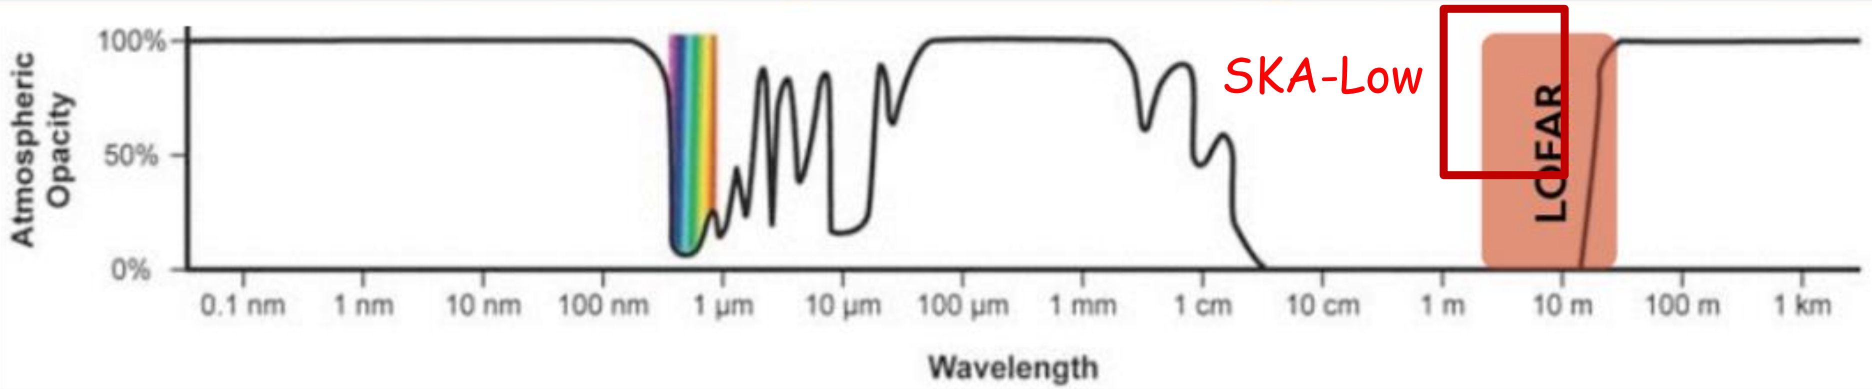

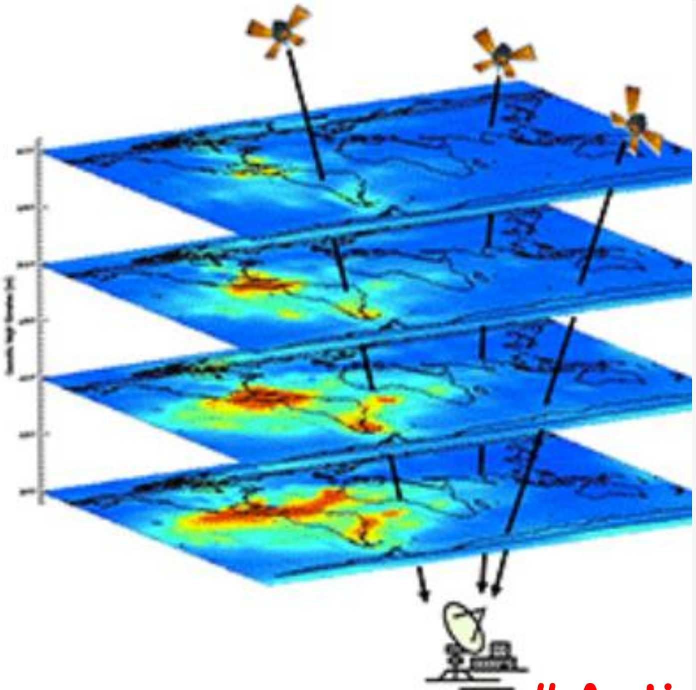

Facet calibration

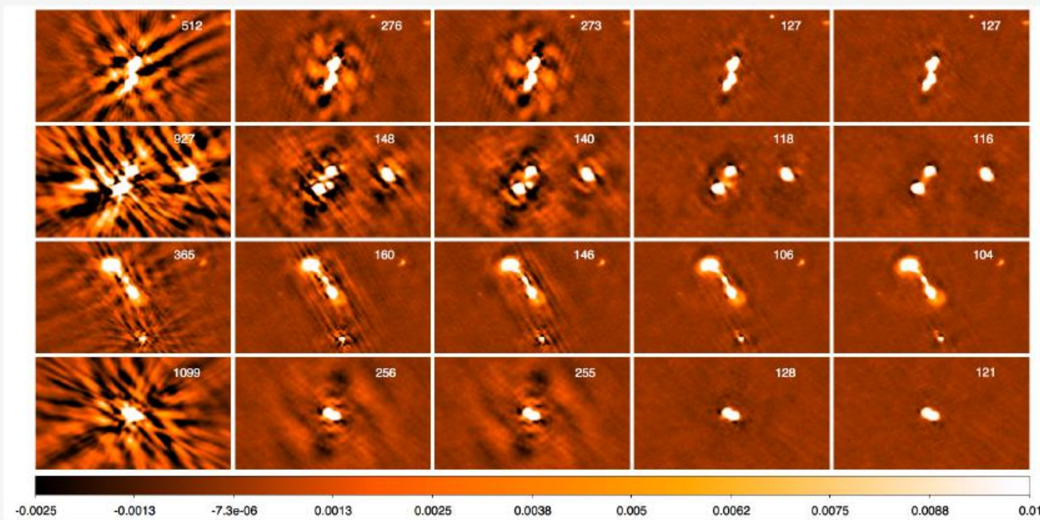

Demonstrating direction dependent calibration (van Weeren R. J., et al., 2016, ApJS, 223, 2)

 # Anticipate the challenges with the SKA LOW #

Science Highlights : from CRs to galaxy clusters

LETTER

doi:10.1038/nature16976

A large light-mass component of cosmic rays at 10^{17}–$10^{17.5}$ electronvolts from radio observations

S. Buitink[1,2], A. Corstanje[2], H. Falcke[2,3,4,5], J. R. Hörandel[2,4], T. Huege[6], A. Nelles[2,7], J. P. Rachen[2], L. Rossetto[2], P. Schellart[2], O. Scholten[8,9], S. ter Veen[3], S. Thoudam[2], T. N. G. Trinh[8], J. Anderson[10], A. Asgekar[3,11], I. M. Avruch[12,13], M. E. Bell[14], M. J. Bentum[3,15], G. Bernardi[16,17], P. Best[18], A. Bonafede[19], F. Breitling[20], J. W. Broderick[21], W. N. Brouw[3,13], M. Brüggen[19], H. R. Butcher[22], D. Carbone[23], B. Ciardi[24], J. E. Conway[25], F. de Gasperin[19], E. de Geus[3,26], A. Deller[3], R.-J. Dettmar[27], G. van Diepen[3], S. Duscha[3], J. Eislöffel[28], D. Engels[29], J. E. Enriquez[3], R. A. Fallows[3], R. Fender[30], C. Ferrari[31], W. Frieswijk[3], M. A. Garrett[3,32], J. M. Grießmeier[33,34], A. W. Gunst[3], M. P. van Haarlem[3], T. E. Hassall[21], G. Heald[3,13], J. W. T. Hessels[3,23], M. Hoeft[28], A. Horneffer[5], M. Iacobelli[3], H. Intema[32,35], E. Juette[27], A. Karastergiou[30], V. I. Kondratiev[3,36], M. Kramer[5,37], M. Kuniyoshi[38], G. Kuper[3], J. van Leeuwen[3,23], G. M. Loose[3], P. Maat[3], G. Mann[20], S. Markoff[23], R. McFadden[3], D. McKay-Bukowski[39,40], J. P. McKean[3,13], M. Mevius[3,13], D. D. Mulcahy[21], H. Munk[3], M. J. Norden[3], E. Orru[3], H. Paas[41], M. Pandey-Pommier[42], V. N. Pandey[3], M. Pietka[30], R. Pizzo[3], A. G. Polatidis[3], W. Reich[5], H. J. A. Röttgering[32], A. M. M. Scaife[21], D. J. Schwarz[43], M. Serylak[30], J. Sluman[3], O. Smirnov[17,44], B. W. Stappers[37], M. Steinmetz[20], A. Stewart[30], J. Swinbank[23,45], M. Tagger[33], Y. Tang[3], C. Tasse[44,46], M. C. Toribio[3,32], R. Vermeulen[3], C. Vocks[20], C. Vogt[3], R. J. van Weeren[16], R. A. M. J. Wijers[23], S. J. Wijnholds[3], M. W. Wise[3,23], O. Wucknitz[5], S. Yatawatta[3], P. Zarka[47] & J. A. Zensus[5]

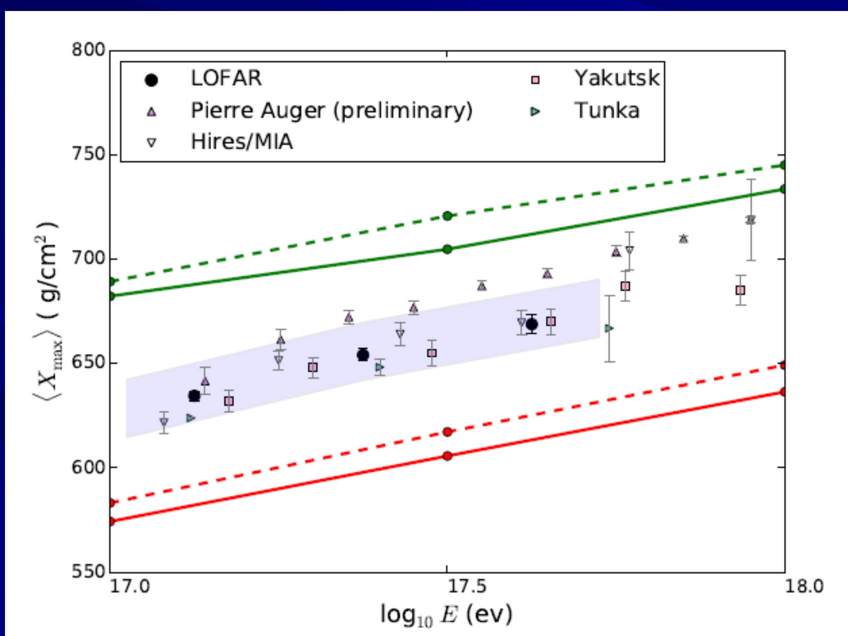

SCIENCE ADVANCES | RESEARCH ARTICLE

PHYSICAL SCIENCES

Gentle reenergization of electrons in merging galaxy clusters

Francesco de Gasperin,[1,2]* Huib T. Intema,[1] Timothy W. Shimwell,[1] Gianfranco Brunetti,[3] Marcus Brüggen,[2] Torsten A. Enßlin,[4] Reinout J. van Weeren,[1,5] Annalisa Bonafede,[2,3] Huub J. A. Röttgering[1]

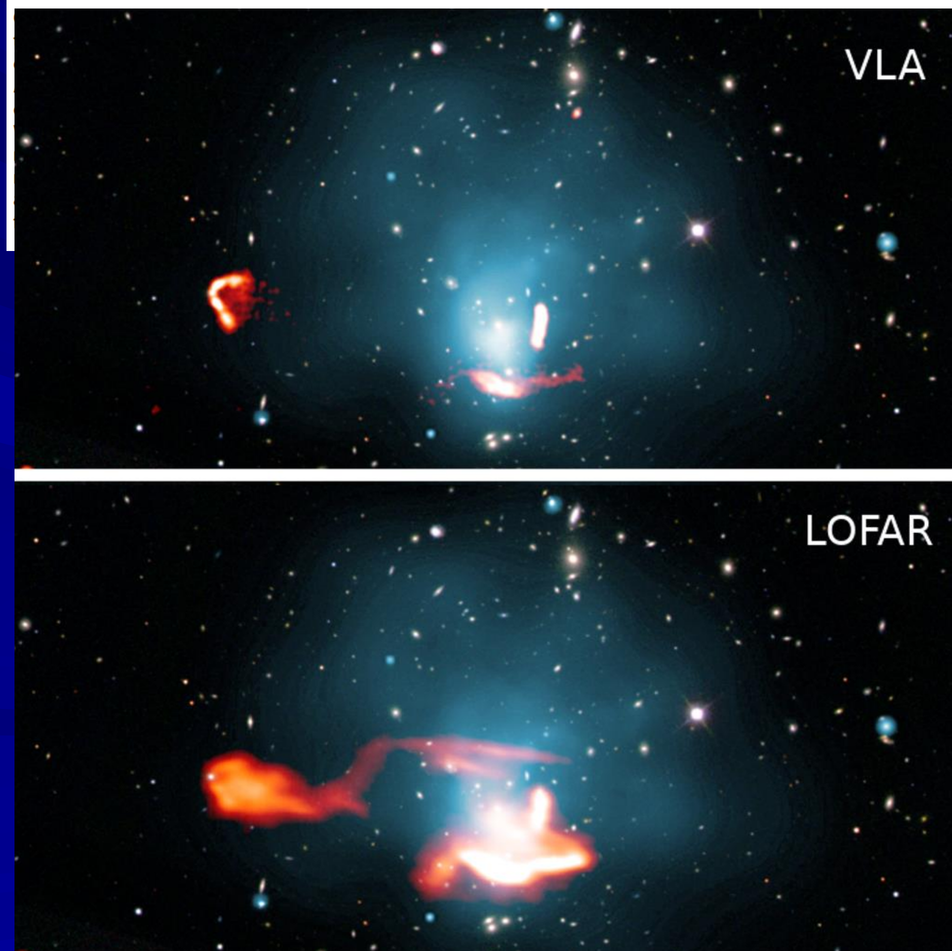

Technical Highlights : extreme resolution !

- Unique Telescope (even) in the SKA era -

Morabito + 2016

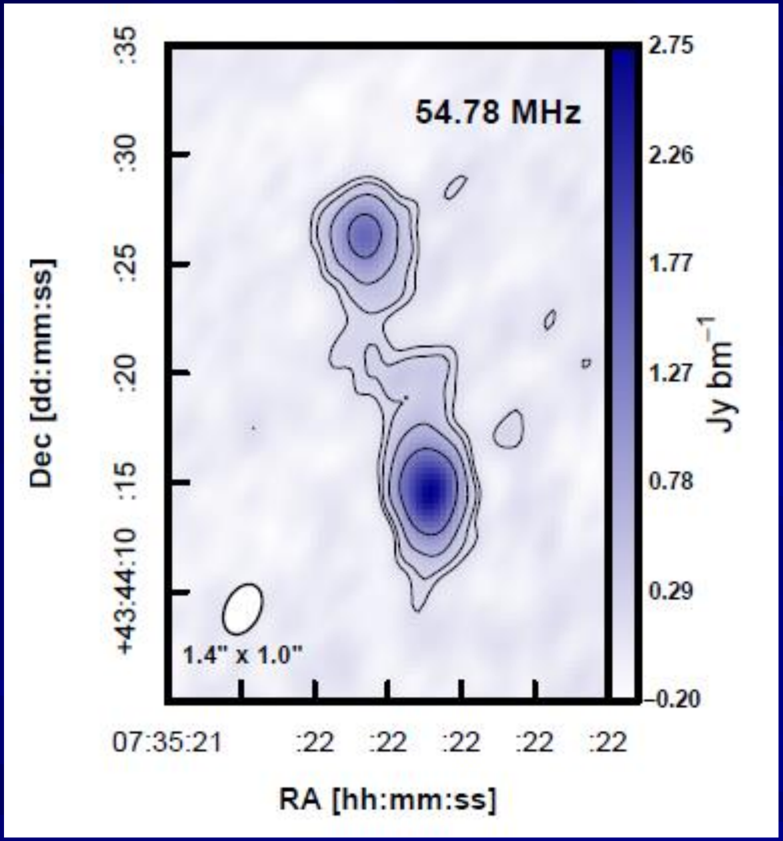

4C 43.15 (z=2.4)
1 arcsec res

Compact sources
SNRs ?

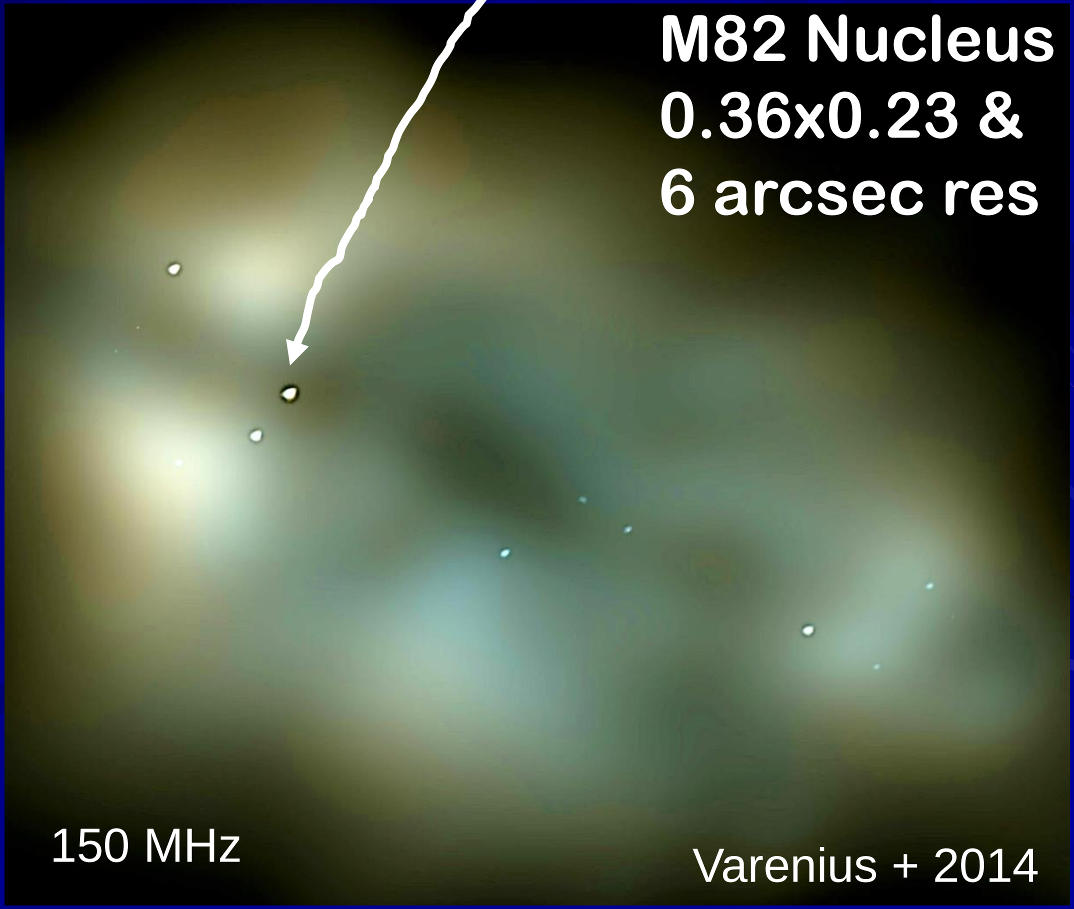

Data-flow & archive challenges

1.7 Tbyte/s
Distributed across stations

28 Gbyte/s

Station signals collected in the station cabinets

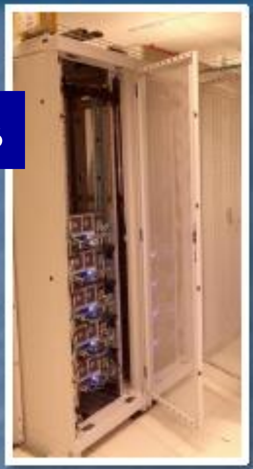

Signal sent to COBALT for correlation

CEP4
(IM and BF pipelines)

Data sent to CEP4 for initial RO processing

Products sent to the long-term archive (3 sites: NL, D, PL)

- Data handling challenge:
  - Data flow from all antennas combined: 1.7 Tbyte/s
  - To COBALT from station after beamforming: 28 Gbyte/s
  - Correlator output to disk: between 2-10 Gbyte/s
  - Data storage challenges: ~ few TB/h
- LOFAR is the first of a number of new astronomical facilities dealing with the transport, processing and storage of these large amounts of data and therefore represents an important technological pathfinder for the SKA

Data-flow cha

1.7 Tbyte/s
Distributed across stations

28 Gbyte/s

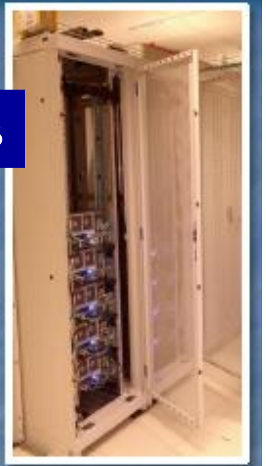

Station signals collected in the station cabinets

Signal sent to COBALT for correlation

On January 1st 2013 ASTRON started the COBALT (COrrelator and Beamforming Application platform for the Lofar Telescope) project to develop a CPU-GPU based system as the central correlator and beamforming platform for the International LOFAR Telescope. The COBALT system is a co-design of both Commercial-Of-The-Shelf hardware components and ASTRON written software. The system consists of 8 production nodes and 1 development / test node, each consisting of 2 CPUs (Intel Dual Xeon E5) and 2 GPUs (NVIDIA K10) housed in a DELL T620 box and connected by an FDR Infiniband Switching network. This gives a balanced system where each CPU connects to one GPU, one Infiniband port and two Ethernet ports and both CPUs within one node are also connected. The cooling of the GPU cards in this set-up turned out to be an issue. Special air ducts were designed and constructed by the CIT Groningen en the ASTRON Mechanical Department, which provided enough cooling for the cards. The full system passed certification by DELL at the end of 2013.

- Data handling challenge:
  - Data flow from all antennas combined: 1.7 Tbyte/s
  - To COBALT from station after beamforming: 28 Gbyte/s
  - Correlator output to disk: between 2-10 Gbyte/s
  - Data storage challenges: ~ few TB/h
- LOFAR is the first of a number of new astronomical facilities dealing with the transport, processing and storage of these large amounts of data and therefore represents an important technological pathfinder for the SKA

Data-flow & archive chalenges

1.7 Tbyte/s
Distributed across stations

Station signals collected in the station cabinets

28 Gbyte/s

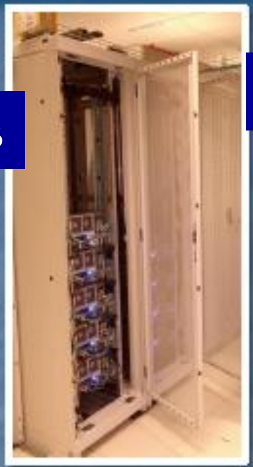

Signal sent to COBALT for correlation

2-10 Gbyte/s

CEP4
(IM and BF pipelines)

Data sent to CEP4 for initial RO processing

MS format
Averaging
Flaggings

Products sent to the long-term archive (3 sites: NL, D, PL)

- Data handling challenge:
  - Data flow from all antennas combined: 1.7 Tbyte/s
  - To COBALT from station after beamforming: 28 Gbyte/s
  - Correlator output to disk: between 2-10 Gbyte/s
  - Data storage challenges: ~ few TB/h
- LOFAR is the first of a number of new astronomical facilities dealing with the transport, processing and storage of these large amounts of data and therefore represents an important technological pathfinder for the SKA

Data-flow & archive ch

1.7 Tbyte/s
Distributed across stations

28 Gbyte/s

Station signals collected in the station cabinets

Signal sent to COBALT for correlation

2-10 Gbyte/s

CEP4 (IM and BF pipelines)

Data sent to C for initial R processing

MS format
Averaging
Flaggings

CEP4

The Lofar phase 4 cluster (CEP4) is adopted to store the raw observation data and process them through the standard data pipelines. Processed data products are made available to the user via the Long-Term Archive, but may also be copied to the CEP3 cluster upon request for further analysis by the user in the original proposal. Due to the intensive nature of the standard data pipelines and the need for these compute resources to be allocated and scheduled by Radio Observatory staff, access to the resources on CEP4 is strictly limited to the Radio Observatory. In the following, a short description of the computing characteristics/performances of the new cluster is given.

The Lofar Phase 4 cluster consists of

- 50 compute nodes (called cpu01..cpu50)
- 4 GPU nodes (gpu01..gpu04)
- 18 storage nodes (data01..data18)
- 2 meta-data nodes (meta01..meta02)
- 2 head nodes (head01..head02)
- 1 management node (mgmt01)

Each node is reachable as XXXX.cep4.control.lofar. Users are only allowed on head01 and head02.

Each compute node consists of:

- CPU: Intel Xeon E5-2680v3 2.5 GHz (12 cores, HyperThreading disabled)
- Memory: 256GB @ 2133 MHz
- Disk: 2x 300GB 10Krpm SAS RAID
- Network: 2x 1GbE, 2x 10GbE, 1x FDR InfiniBand

Each GPU node consists of:

- CPU: Intel Xeon E5-2630v3 2.4 GHz (8 cores, HyperThreading disabled)
- Memory: 320GB @ 1866 MHz
- Disk: 2x 300GB 10Krpm SAS RAID + 2x 6TB 7.2Krpm SAS RAID
- Network: 2x 1GbE, 2x 10GbE, 1x FDR InifiniBand

Each head node consists of:

- CPU: Intel Xeon E5-2603v3 1.6 GHz (6 cores, HyperThreading disabled)
- Memory: 128GB @ 1600 MHz
- Disk: 2x 300GB 10Krpm SAS RAID + 2x 6TB 7.2Krpm SAS RAID
- Network: 2x 1GbE, 2x 10GbE, 1x FDR InifiniBand

The other nodes are not accessible (storage, meta-data, and management nodes).

Storage: The storage and meta-data nodes provide a ~2PB LustreFS global filesystem through the InfiniBand network to all nodes in /data, thus implying that all nodes see the same data.

Processing: CEP4 uses a SLURM batch scheduling system to schedule and run all observation and processing jobs on the cluster.

- Data handling challenge:
  - Data flow from all antennas combined: 1.7 Tbyte/s
  - To COBALT from station after beamforming: 28 Gbyte/s
  - Correlator output to disk: between 2-10 Gbyte/s
  - Data storage challenges: ~ few TB/h
- LOFAR is the first of a number of new astronomical facilities dealing with the transport, processing and storage of these large amounts of data and therefore represents an important technological pathfinder for the SKA

Data-flow & archive chalenges

1.7 Tbyte/s
Distributed across stations

28 Gbyte/s

Station signals collected in the station cabinets

Signal sent to COBALT for correlation

2-10 Gbyte/s

CEP4
(IM and BF pipelines)

Data sent to CEP4 for initial RO processing

MS format
Averaging
Flaggings

storage
2 TB/h

Products sent to the long-term archive (3 sites: NL, D, PL)

- Data handling challenge:
  - Data flow from all antennas combined: 1.7 Tbyte/s
  - To COBALT from station after beamforming: 28 Gbyte/s
  - Correlator output to disk: between 2-10 Gbyte/s
  - Data storage challenges: ~ few TB/h
- LOFAR is the first of a number of new astronomical facilities dealing with the transport, processing and storage of these large amounts of data and therefore represents an important technological pathfinder for the SKA

Data-flow & archive chalenges

1.7 Tbyte/s
Distributed across stations

28 Gbyte/s

2-10 Gbyte/s

CEP4
(IM and BF pipelines)

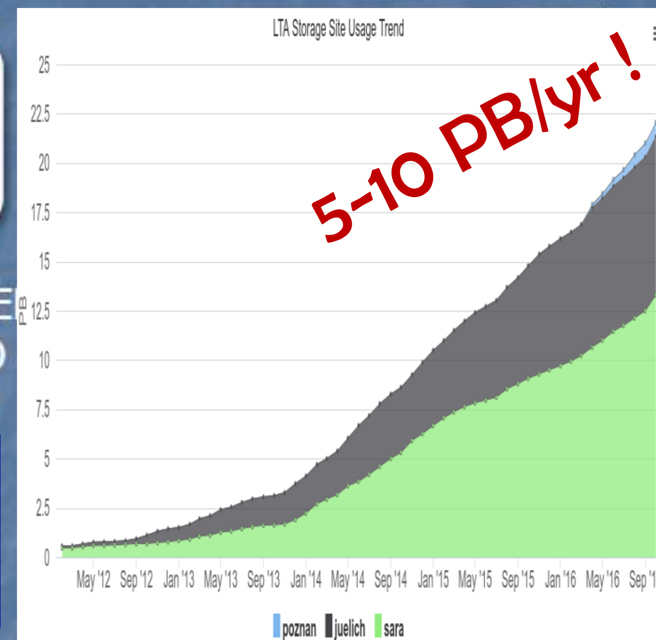

Station signals collected in the station cabinets

Signal sent to COBALT for correlation

Data sent to CEP4 for initial RO processing

MS format
Averaging
Flaggings

- Data handling challenge:
  - Data flow from all antennas combined: 1.7 Tbyte/s
  - To COBALT from station after beamforming: 28 Gbyte/s
  - Correlator output to disk: between 2-10 Gbyte/s
  - Data storage challenges: ~ few TB/h
- LOFAR is the first of a number of new astronomical facilities dealing with the transport, processing and storage of these large amounts of data and therefore represents an important technological pathfinder for the SKA

User point of view … 8 hrs obs [1s, 16/SB]

Station signals collected in the station cabinets

Signal sent to COBALT for correlation

CEP4
(IM and BF pipelines)

Data sent to CE[P4] for initial RO[ ] processing

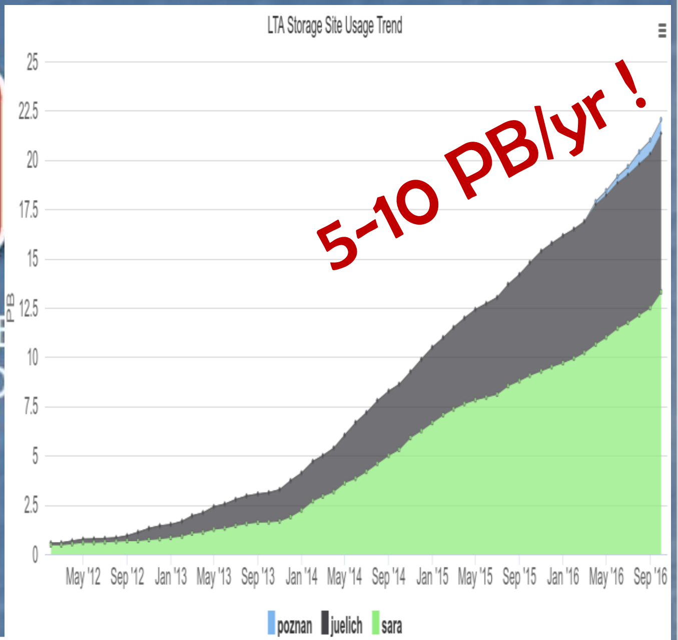

…………………… ILT
COBALT : 30 TB (40 TB Int Stations)
CEP4 : 16 TB (20 TB Int Stations)

…………………… CEP3 OR Local Facilities :
PRE FACTOR (DI calib): from 16TB to 500 GB [200 GB RAM, few days]
INITIAL SUBTRACTION: 500 GB [up to 400 GB RAM, few days]
FACTOR (DD calib) + IMAGING : 700 GB [embarassing parallel, RAM…..]

Italian Roadmap to.. LOFAR

Credits S. Tingay

- SKA precursors/pathfinders drive frontier research
- IT community is leader in ideas
- IT community is the only one without a SKA pathfinder/precursor !
- IT is deeply involved in SKA Low technical engineering

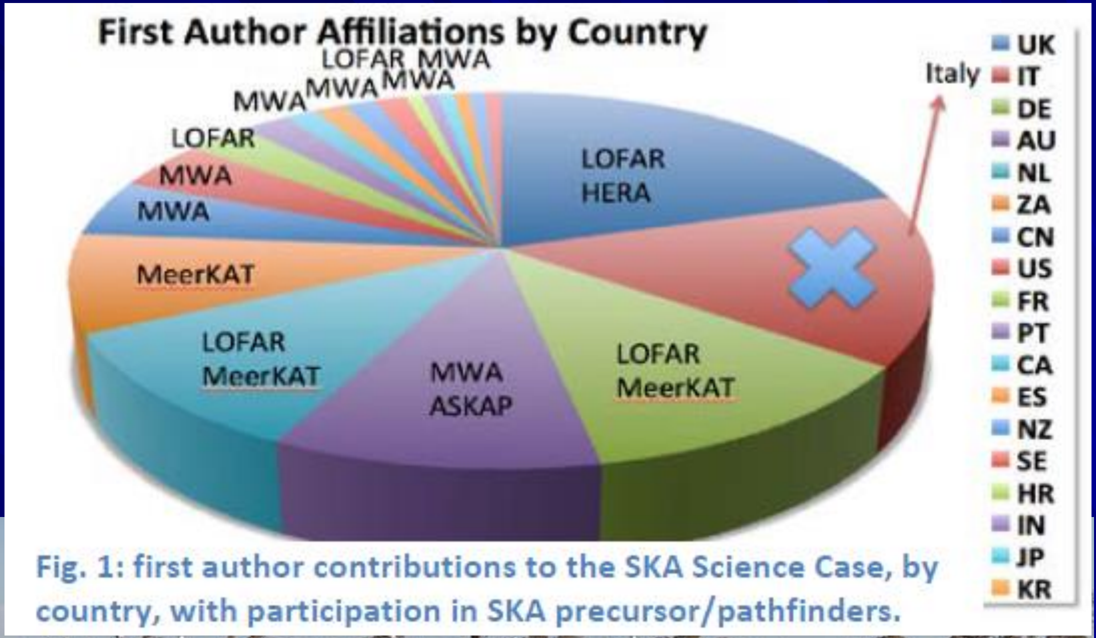

Fig. 1: first author contributions to the SKA Science Case, by country, with participation in SKA precursor/pathfinders.

Urgent action : join a leading SKA precursor/pathfinder to carry out frontier research, meet challenges and bring up a strong community for the SKA era

Italian Roadmap -on behalf of UTG2-

Join LOFAR asap to join Key Science Programs and meet data and computational challenges

Acquire and install 1+ LOFAR stations in IT

Ongoing agreement (ASTRON) to join ILT

(1) IT-LOFAR Consortium : INAF, UniTO, …
(2) Technological involvement in LOFAR 2.0 [2FTE x 3 yrs]
(3) MoU for 1 LOFAR 2.0 Station [installation 2020-2021]
(4) Central running costs 92 kE/year [2018+]

2+ ME

ONGOING ACTION IN UTG2 :

Evaluation of resources , plans for organization (software, computing, storage), technological involvment in LOFAR 2.0, optimization of the scientific impact (training fellowships, post-doc)
[GB + Tech WG: Becciani, Bolli, Bonafede, Monari, Nanni, Perini, Taffoni]

TAKE HOME MESSAGE

- LOFAR IS THE BIGGEST SKA PATHFINDER, OPERATED BY 8 EU COUNTRIES
- WILL ALLOW FRONTIER SCIENCE & DISCOVERIES IN THIS DECADE
- IS FACEING DATA ANALYSIS AND COMPUTATIONAL CHALLENGES IN VIEW OF THE SKA-LOW
- WILL BE BETTER THAN SKA-LOW FOR LONG BASELINES (HIGH RESOLUTION)
- INAF IS CONDUCTING NEGOTIATIONS WITH THE AIM/HOPE TO JOIN LOFAR IN 2018

User point of view … 8 hrs obs [1s, 16/SB]

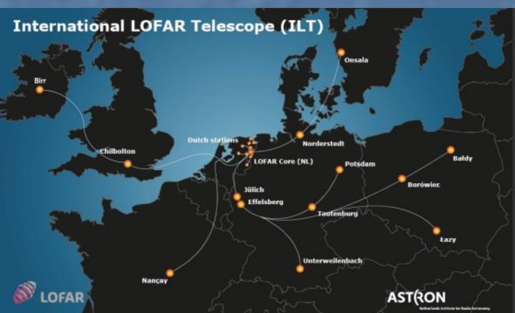

Station signals collected in the station cabinets

Signal sent to COBALT for correlation

CEP4 (IM and BF pipelines)

Data sent to CE
for initial RO
processing

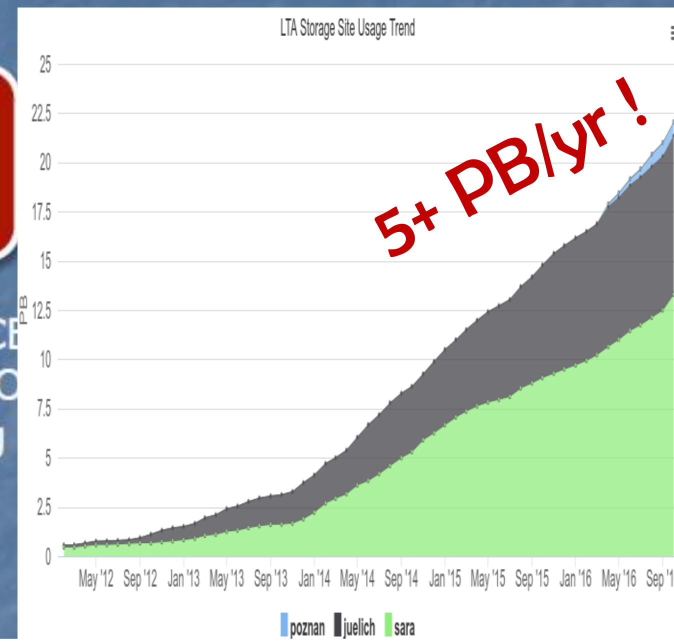

…………………… ILT

COBALT : 30 TB (40 TB Int Stations)

CEP4 : 16 TB (20 TB Int Stations)

…………………… CEP3 OR Local Facilities :

PRE FACTOR (DI calib): from 16TB to 500 GB [200 GB RAM, few days]

INITIAL SUBTRACTION: 500 GB [up to 400 GB RAM, few days]

FACTOR (DD calib) + IMAGING : 700 GB [embarassing parallel, RAM…..]

Phase 1

Phase 2

Possible Organization (work in progress)

UniTO C3S (+upgrade):
Phase 1 Analysis (10-50 TB)

Distributed Centers:
Phase 2 Analysis
4-5 nodes, multi-user
(x2 CPU 16-32 cores
x128+ GB RAM)

Tiger Team:
Update software
Installation
Development