

Data processing on the LOFAR EoR GPU cluster

Lofar Status Meeting
Panos Labropoulos
panos@astron.nl



- **The EoR KSP will observe 5 windows with:**
 - ~ 6 beams in two bands (total $64\text{MHz}/10\text{KHz}=6400$ channels)
 - Full stokes
 - 1200 baselines
 - 10 sec integration – 100 nights of 4 hours
- It will accumulate **1.5 PB data**
- Desired dynamic range is $10^6:1$
- 3
- After calibration: **Maximum Likelihood inversion ($O(N^3)$)**

Cluster Layout

Lofareor01 inlog machine
Lofareor05 supporting machine
(alleen system administrators kunnen op deze machine inloggen)

192.168.235.1 t/m 40


192.168.235.41 t/m 80

Impi 192.168.235.101 t/m 140

Impi 192.168.235.141 t/m 180

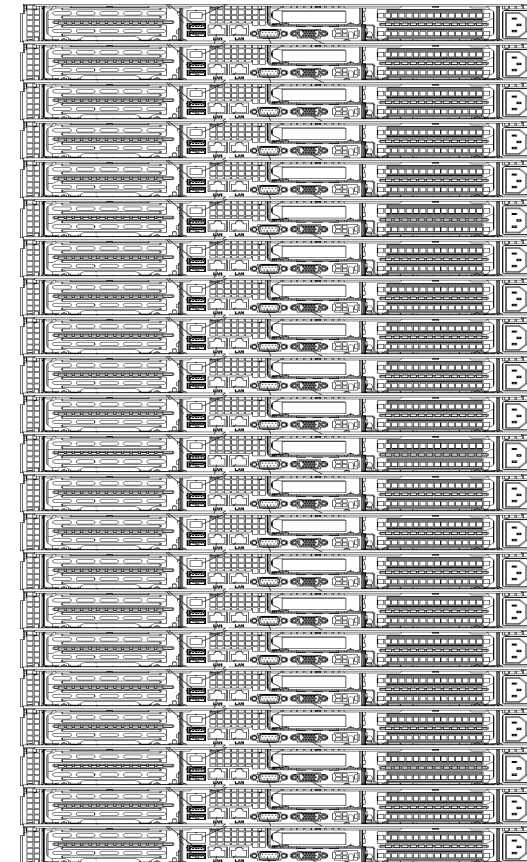


node001
node005
node009
node013
node017
node021
node025
node029
node033
node037
node041
node045
node049
node053
node057
node061
node065
node069
node073
node077




lofareor01
lofareor05

node002
node006
node010
node014
node018
node022
node026
node030
node034
node038
node042
node046
node050
node054
node058
node062
node066
node070
node074
node078



Cluster Layout


Netwerk switches worden verbonden met 4 x 10 gb glas
verbindingen. Totaal 40 gb



Hardware

- 4x Intel Xeon E5520 @ 2.27 GHz
- 12 GB RAM
- 2x 2TB SATA HDDs
- 2x 1GBps NICs
- 2x NVIDIA TESLA C1060 GPUs (4G DDR3 RAM)
- Centos LINUX 5.4
- LOFAR software installed
- Experimental versions (ie BBS with shapelets, GPU accelertion, EoR Imager)

Location




Zernike borg




Land Leven

LOFAR Switch Equipment (D row)





— Already installed
— To be installed





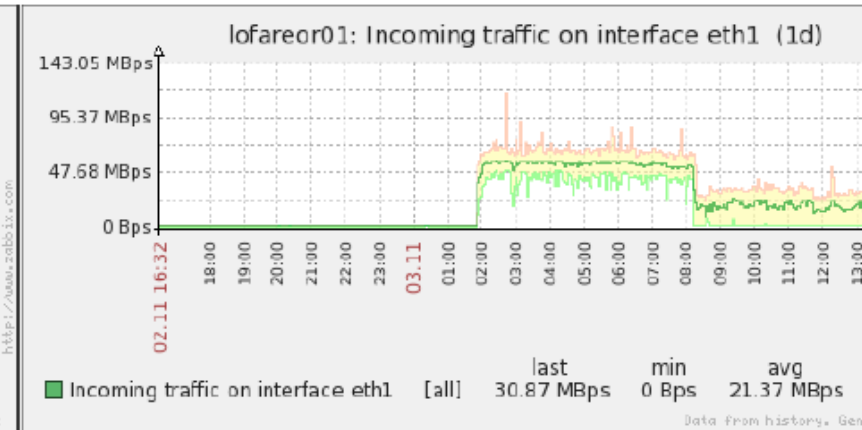

-1GBs to the world

-4x 10GBps connections to the TARGET cluster

Current data rates



Transfer rates




Successfully processed all the sub-bands at full resolution

1. AOFLagger for RFI flagging
2. BBS for calibration using 25 Sources
3. Tried different imagers (mwimager:casa or cimager mode and the CASA imager)
4. Can run a basic pipeline job in one day.
5. Update the LSM using the Duchamp (after having corrected the images using the Hamaker beam)

Frequency cube

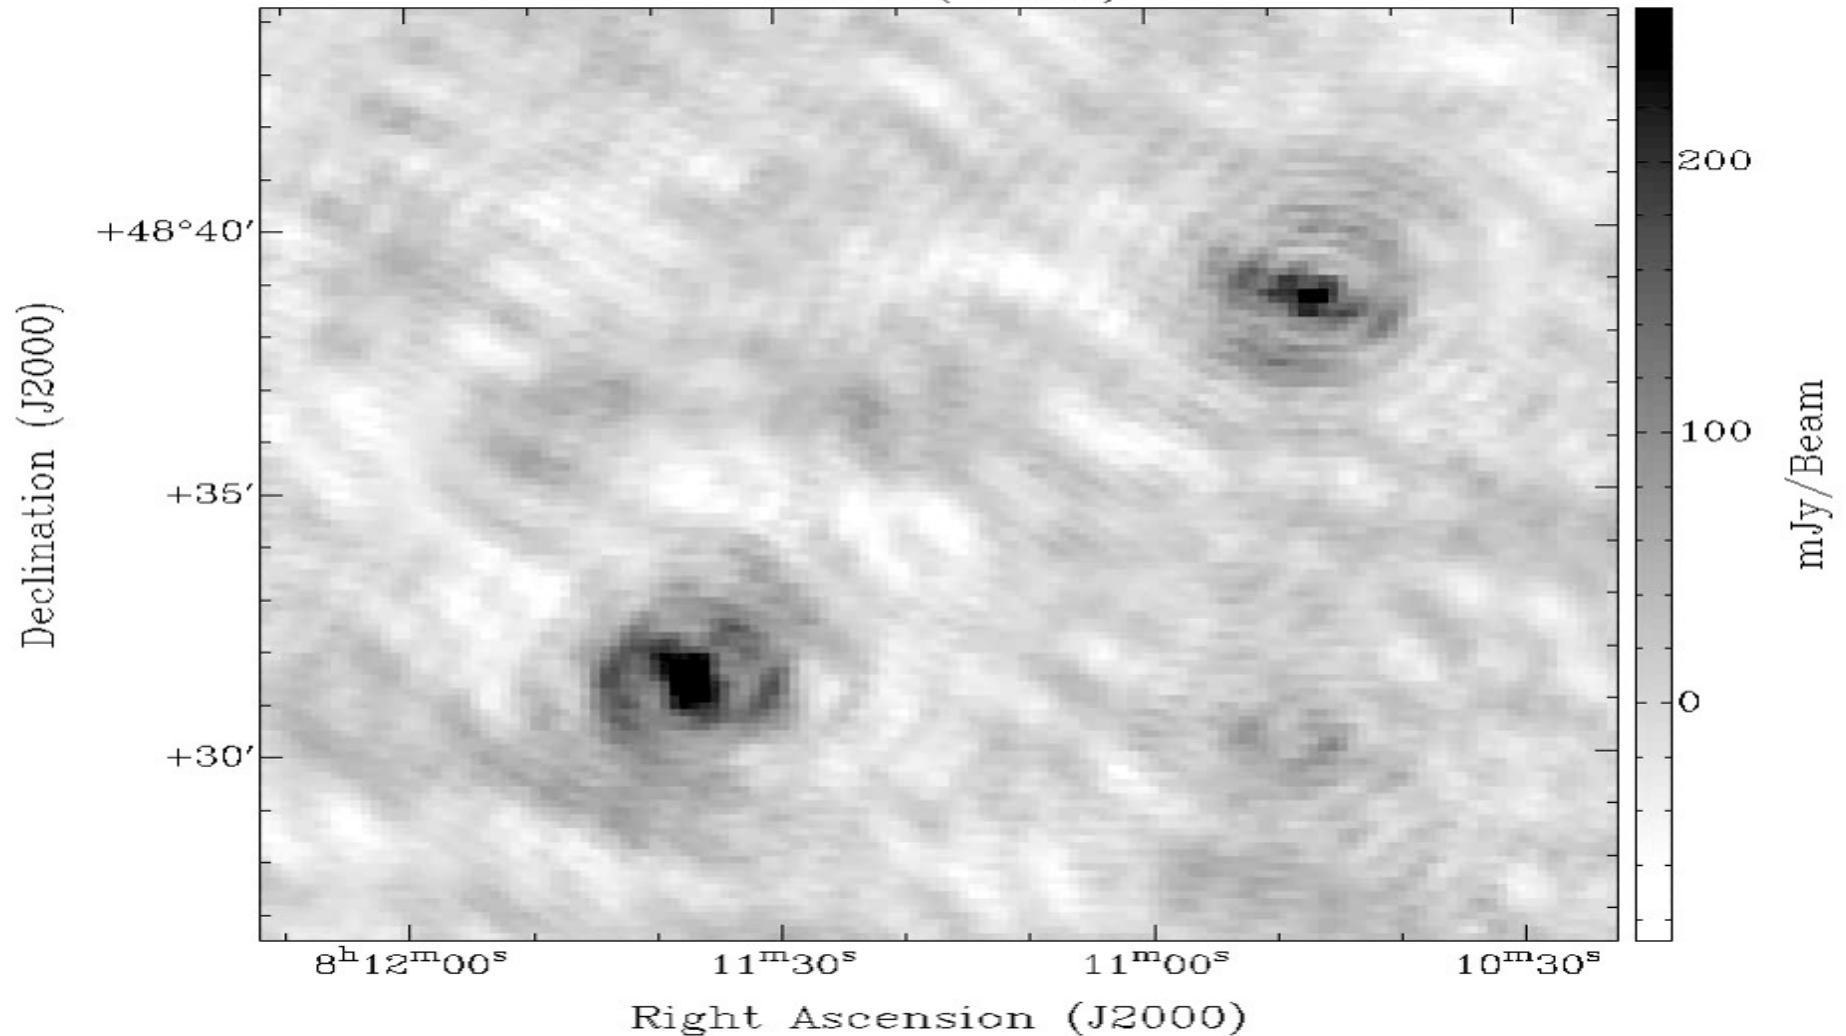
Frequency: 114.954 MHz



Frequency cube: zoom in

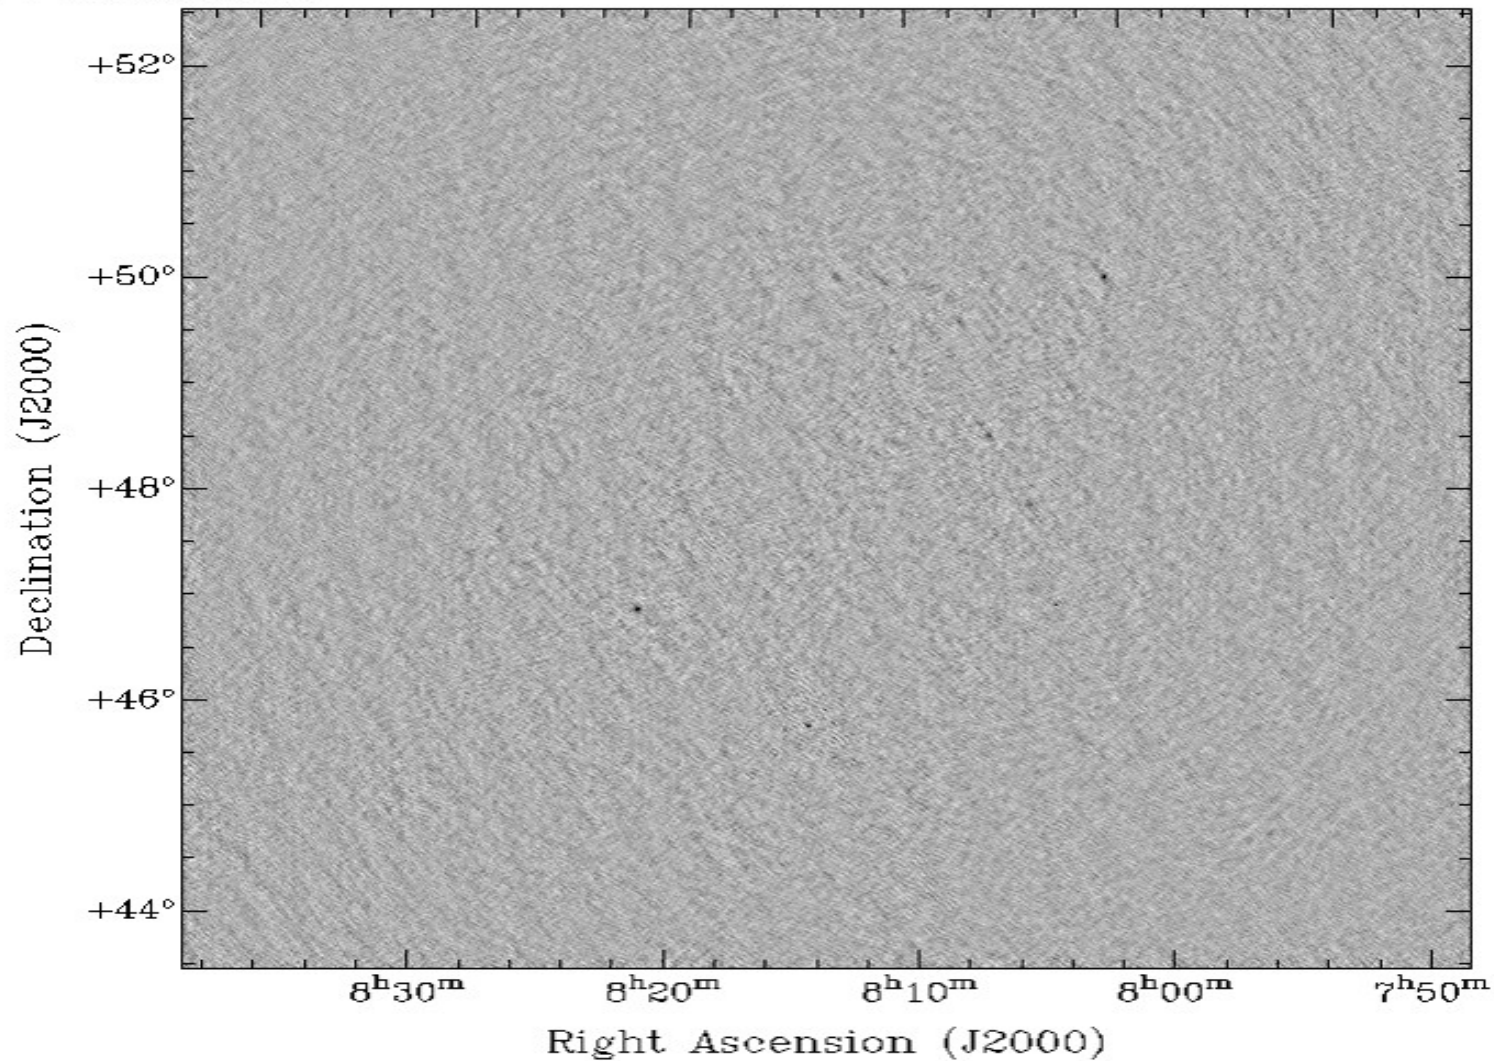
Frequency: 114.954 MHz

3C196 (22006)




Snapshots (12 min.)

STOKES: 1.000000e+00



Snapshot (zoom in)


STOKES: 1.000000e+00



Snapshot (zoom in more)

STOKES: 1.000000e+00

Snapshot



Initial conclusions


Errors look as half circles (phase error)+ they follow the same pattern along the frequency direction -> They must be related to the beams

Factor of 10 away from the theoretical sensitivity limit (6 mJy noise instead of 0.6mJy)

Long time scale error around 140 MHz: corrugations in the image


Difference map (2 sub-bands)

- 10mJy@132MHz noise
- Almost consistent with noise
- Spectral structure around bright sources




Residuals



- $\sim 6\text{mJy}$
- Few point sources can be seen
- Consistent with noise
- Gaussian histogram
- 10 SB average



Noise behaviour



Hint!



- Lots of computing power, but still...
- DATA HUNGRY (need better connection to the CEP cluster)
- Can provide useful feedback and
- Reduce the load on the CEP cluster
- Testbed for the final EoR data processing

Thanks to E. Tiesinga, O. Martinez and S. Yatawatta