

## Minutes of Meeting LOFAR Software

<b>Date:</b>	2010-06-30
<b>Next meeting:</b>	2010-07-07 9:30-10:30
	Multimedia room
<b>Present:</b>	
Andre Gunst	Yes
Ronald Nijboer	No
Ruud Overeem	Yes
John Romein	Yes
Michael Wise	No
Harm Munk	Yes, later
Hanno Holties	Yes

cc: Arnold Meijster, Rob van Nieuwpoort, Arthur Coolen, Jurjen Sluman, Pieter Donker, Chris Broekema, Joris v. Zwieten, Marcel Loose, Adriaan Renting, Ger van Diepen, Michiel v. Haarlem, Jan Reitsma, Ger de Bruyn, Arno Schoenmaker, Hanno Holties, Corina Vogt, Jan Noordam, Joe Masters, Lars Bähren, Johan Hamaker, Sven Duscha, Jan-David Mol, Teun Grit, Alwin de Jong, Frank Breitling, Anastasia Alexov, Jason Hessels, Joeri van Leeuwen, John McKean, George Heald.

### **Remarks previous minutes**

- Half a line should be removed in the previous minutes: “On behalf of Alwin t”
- 

### **Announcements**

- Chilbolton started with the installation of HBA tiles.
- MAC/SAS has a fully operational test environment again.
- French visitors are here today to talk about the LSS (LOFAR Super Station).
- Next week TBB busy week.
- Nancay: station validation and network will happen next week.
- Sara archive 10 G Amsterdam – Groningen link will be used for Nancay, Chilbolton and Onsala connection.
- Stopday: 5-6 July because of BG/P upgrade and storage node firmware upgrades.

### **Action item overview**

ID	Date submitted	Description	Owner	Planned date	Status
107	20100303	Issue tracker decision. Bugzilla (developers), Craft (system and hardware issues), science support wants to make use of the LOFAR observation tracker. For the user software the favor is for Bugzilla as well. Michael will talk to Antonis to see if this is really what we want. Michael talked to Antonis and Michael got permission to setup Redmine. John S. will give a demo to Antonis, Hanno, Harm and Michael.	Hanno/Michael	20100315	Open
113	20100414	Define end to end quantitative tests for the imager pipeline.	Ronald	20100430	Open

119	20100512	What is the procedure if rack 0 fails and we like to use rack 1? This has to be discussed with CIT. Currently LOFAR can use two racks. Harm Paas can run a script to remove scientific users from a rack. This will be automated, such that an operator can run this as well.	Andre	20100519	Open
122	20100519	A meeting will be organized to finalize beam formed data naming convention. Attendees: J. Hessels, A. Alexov, L. Bähren, J.-M. Griessmeier, A. Renting, H. Holties, R. Nijboer, R. Overeem, A. Gunst, M. Wise, Harm Munk. This is organized at 7 July 2010.	Michael	20100526	Open
123	20100623	A meeting will be held to discuss the PIL library and its compatibility with the parset files.	Michael	20100701	Open
124	20100623	Organize status meeting of all pipelines in September.	Andre	20100701	Open
125	20100623	Next meeting: software architecture offline system and MSSS plan. Discussed lightly since both Michael and Ronald are not in today. Idea is that a small group of people discuss the issues and holes in the current design. Names mentioned are: Adriaan, Alwin, Marcel, Ger, Ruud, John Swinbank, Ken Anderson. Decision has yet to be made who will lead this process. Action remains open.	Andre	20100701	Open
126	20100630	Distributed software inventory list.	Harm	20100707	Open

Last: 126

## **Progress**

### **System Integration**

Achieved since last meeting:

- The beamservers operate under MAC/SAS. Yesterday an observation was done with all four HBA settings. The beamservers can be tested as well in the test environment now.
- The fixes of MAC/SAS done to make it more robust will be tested on the test environment as well.
- Alwin needs a modification on the stations for using the dataslots (=beamlets) in the scheduler. Discussion was held about the urgency: With the new scheduler including the data slot allocation the user gets more freedom to assign subbands to beamlets. Now you specify which data slots are used for all stations at once. And this specification is valid for all stations (cannot be different). If you run multiple observations at the same time with multiple subsets of stations than you do not use all dataslots efficiently. In the new allocation scheme the scheduler determines per antenna field the assignment of subbands to dataslots. Since currently seldom two observations are running at the same time and the efficient

use of resources is not such an issue right now we decide to postpone this after Ruud finishes the beamserver and calserver. So, this waits for another 2 months. The work involved is two - five days. Alwin is able to work around this feature such that he can continue working, testing and improving the scheduler.

- Currently an observation has to run on the same raid number. Hence the operators need to find enough disk space on all storage nodes for this raid number per observations. Since now more stations are online the storage nodes run out of disk space and finding space at the same raid number is a serious limitation. This is all caused by the previous name mask. The scheduler will be made such that you can use different partitions for the different nodes.
- Jan-David has reworked the logger. After this the work on the second transpose will be picked up again.
- Chris is busy defining the phase 2 offline + storage hardware.
- Yesterday (metadata busy week) it was possible to define an observation in MOM and execute the observation via SAS and the scheduler. MOM is connected to SAS and the scheduler reads it from SAS and sent the result back to SAS.
- Pieter has the firmware of the temperature control modified: if temperature in the subrack cabinets becomes critical then 48V is switched off. If 0.2 degrees before the critical temperature the 48V is switched off then the third rack does not need to switch off. Since there the LCU is present you can still communicate with the station.
- Jan-Pieter is busy installing a mechanical construction to the cabinet such that the hot air outflow is not entering the inflow easily.
- Arie Doorduyn is busy implementing firmware for the 4 bit mode. He expects to be done in about two months. By that time Ruud or Pieter should have worked on modifications in the interface to the RSPs as well to set the right registers. Moreover changes in the beamserver will be necessary to support these modes.
- Pieter is working on the beamserver to read in static values for all RCU modes.

Problems / current activities:

- 

Next actions:

- Solutions for the high station temperature in the summer are under investigation.
- A temperature sensor will be installed in the concentrator node as well.
- 4 bit mode implementation ongoing.
- 

### **Imaging Pipeline (Ronald):**

Achieved since last meeting:

- Adriaan is busy trying to start the pipeline via SAS.
-

Problems / current activities:

- Control BBS should be revisited by Marcel. How to deal with failing processing nodes and the use of the global solver. Needs to be prioritized. Becomes an issue when we process with lots of nodes.
- 

Next actions:

- Focus on the minimal required tasks for MSSS.

### **Pulsar Pipeline (Michael):**

Achieved since last meeting:

- 
- 

Problems / current activities:

- Need to identify a new developer for the BF data writer as soon as possible.
- 

Next actions:

- Implement second transpose operation.
- Update BF datawriter.
- 

### **VHECR Pipeline (Michael):**

Achieved since last meeting:

- Cleaning up the CR tools code tree is in progress
- CR tools work now under MAC OS
- Imaging algorithms works in Python
- Source finding algorithms works in Python as well.
- A TBB busy week is scheduled from 5-9 July 2010.

Problems / current activities:

- 

Next actions:

-

## ***Infrastructure (Harm)***

Achieved since last meeting:

- Software inventory was made. Now a list is available with all required software for a release 1.0. If the LTA part is excluded than at the End of October we are ready for 1.0. Including the LTA it takes to the end of May 2011. Suggestion is to distribute this list also to us in order to determine if we have to shift priorities in order to get the 1.0 out as soon as possible.
- During the ILO Harm announced that the system will be three days available for system tests (Monday, Tuesday, Wednesday) and one day to make it stable again for taking Friday + weekend observations. Busy weeks and requests can cause that the observatory deviates from this plan. Arising questions are: Is the balance right? Why not assigning part of the stations for tests during a number of days of the week. And how long is it going to be like this. This will take to at least the end of October and probably the end of the year. The reason to do all this is to reduce the “versnippering” of developers. This scheme only involves the online part. Hence the offline cluster is available through the whole week.
- 

Problems / current activities:

- 

Next actions:

- 

## ***User Data and Archive (Hanno)***

Achieved since last meeting:

- An imaging pipeline meeting was held with SARA, CIT and LOFAR at the 8<sup>th</sup> of June.
- An Astrowise LOFAR workshop is held 16, 18, 21 June.
- Marcel is buy installing BBS at SARA. After that he will work on installing the imager there or installing software on the Groningen compute cluster (another one then the offline cluster).
- Adriaan is today busy with a data challenge to Juelich. However the usable bandwidth is limited.
- A request for a compute cluster for BiG Grid in Groningen connected to the Target storage systems waits on feedback from the BiG Grid management.
- Use case / blue print is made for data flows, storage and processing.
- One of the worries is the bandwidth in between the Target storage and Target compute cluster. According to IBM this can be scaled up.

- Today there will be an internal meeting to talk about the required 10 G connections to the archives and LOFAR stations and how to acquire those.

- 

Problems / current activities:

- 

Next actions:

- 

## Decisions

ID	Date submitted	Decision
02	20061220	Every Step will start with a Kick-off meeting, in which the complete software team participates.
<del>03</del>	<del>20061220</del>	<del>The project team starts immediately with the preparations of the next CDR in order to preserve progress of the CS1 realization</del>
04	20070116	This meeting will take place every week on Tuesday 11:00. The existing software team meeting with all developers will stop to exist.
05	20070130	Step 1 will be changed to 16 subbands instead of 32 subbands.
06	20070130	Step 2 will contain a multiple node BBS. 6 $\mu$ Stations/Station will be postponed. Instead of this, 32 subbands measurements will be realized.
07	20070206	Step 1 will support 160 MHz observations. The other steps will support 200 MHz as well.
08	20070424	Step 2 will support 16 subbands @ 200MHz and 24 MHz at 160 MHz
09	20070424	During the rest of step two, OLAP will only support observations during the weekend.
10	20070522	The number of subbands per Measurement Set is set to 6 or 8 default.
11	20070522	Scheduler activities will be preferably activated in Q4 2007.
12	20070522	Procure, three Local Control Units to accommodate 12 microstations in CS010 in a quick way.
13	20070529	Integrate version numbers in all software.
14	20070529	Distinguish the software between a production version and an engineering version (partly now already the case).
15	20070605	All developed software under CVS will be transferred to Subversion. The main reason for this is that Subversion supports the integration of version numbers in the executables. In this way you can always retrieve which software is used for a certain build. First the impact of the transfer will be investigated by Marcel.
16	20070619	Marcel Loose will be the librarian of the LOFAR software. The available time for this will be shared with his BBS work.
17	20070710	The known pulsar survey mode will be the next mode to support (not in its full extent but partly on-line and off-line).
18	20070710	The temporarily off-line part of the known pulsar mode pipeline will not be under control of SAS/MAC. This will be put under control of SAS/MAC as soon as that software is available in the on-line part of the system.
19	20070814	Joe Masters makes the routine to read in the TBB data.

20	20071002	Fault tolerance of the system (mainly OLAP) is put at the top of the priority list after closing the SAS-MAC and CEP integration.
21	20071123	Kubuntu 7.10 desktop 64 bit OS is chosen for all machines except the BG/L and MAC/SAS machines
22	20071123	Station calibration work is smeared out over Step 4 and Step 5.
23	20071123	Global bandpass shape is moved to Step 5 because of its low priority.
24	20071211	Multiple beams per observation will be implemented instead of multiple observations (this is consistent with the plan).
25	20071211	Step 3 will be closed next Thursday. Any open items will be finished in Step 4.
26	20080130	Multiple beams are defined as multiple directions with the same set of antennas. Hence, only the angle, subbands and beamlets can be modified per beam.
27	20080206	Step 4 and Step 5 for MAC/SAS will be changed. The control of the offline pipeline will be postponed because the offline subsystems are not fixed yet. Currently the definition and design of the metadata flows will be set as goal for Step 4 and the implementation of the metadata flow will be the end goal of Step 5. Hence, after Step 5 (part of) the metadata is included in the Measurement Set.
28	20080213	Currently a single subband and single beam is stored in a Measurement Set. As soon as we are ready for mosaicing this probably should be changed in the future.
29	20080220	For storing the raw station beams the sanitizing operations like input buffer will be included in the online part. For this OLAP has to give operational support or instructions to the observers how to start up manually such observations. Since, this is an between solution this will not be automated via SAS/MAC.
30	20080227	Weekly build environment will be updated and automated.
31	20080227	After Step 5 the software documentation will be updated and obsolete packages will be removed.
32	20080423	Basically two Low Band modes will be supported initially: a LBL and LBH mode. The connection between antennas and RCUs have to be chosen such that those to modes make sense.
33	20080528	The position of all individual dipoles will be made available centrally in the database.
34	20080603	The data format of the positions will be delivered in ETRS coordinates by the roll out team. However, the data format of the positions will be stored in ITRF format in the LOFAR databases. Hence, all software and configuration files dealing with coordinates must be made compatible with the ITRF dataformat. Hans van de Marel is responsible to convert the ETRS coordinates to ITRF coordinates for the LOFAR system.
35	20080903	Kubuntu will be installed on LOFAR18, which will serve as a software development machine.
36	20081022	Station cabinet will be heated (if necessary) to 10 degrees Celsius (for the LCU).
37	20081029	We will transfer the build environment to cmake.
38	20081029	Step 1 will be closed at 11 November.
39	20081112	Bugs found in the field have the highest priority to solve. Bugs which take more than a week to solve will be added to the task list and prioritized in the software meeting. During bug solving tests should be written up, which proves the correct behavior. These tests will result in a procedure to check the functionality when new soft/firm ware is loaded.
40	20081126	The 4 bit mode will be supported after MS <sup>3</sup> .
41	20081203	We will modify the build environment to cmake from now on.
42	20090129	Transient source modeling tool under Python will be used for source modeling.
43	20090129	Delay deadline of Step 2 to 26 February 2009.
44	20090209	Remote Stations including the ring splitter near the core will be renamed to CS stations.
45	20090813	No connection from the Dwingeloo test environment to Groningen is necessary anymore.
46	20090825	Create a Bugzilla environment for the USG software.

47	20090825	Use one subcluster per group, contactpersons and guidelines defined (see section Software integration).
48	20090909	Use the filter range names of MAC/SAS for the ICDs and the archive model.
49	20100116	HBA beam pointing: we decided that one observation is prime and determines the HBA beam. The other observations will be ranked. An additional field for the HBA beam pointing can be set. If this field is not set, then an average of all digital beams will be made within the prime observation.
50	20100303	Changes in definitions which are used in various places in the system will be decided in this meeting.
51	20100303	HBA_ONE and HBA_TWO will be renamed to HBA_ZERO and HBA_ONE for consistency reasons.
52	20100317	Change HBA_BOTH into HBA_DUAL (using two HBA ears independently) and add HBA_ALL to indicate both HBA fields will be added at station level (so treated as one field).
53	20100317	The software should be documented more. However we decide not to set this as priority now and accept this as a risk we take.
54	20100317	The CImager will be the imager used in LOFAR. This is the only one which scales up.
55	20100331	The name HBA_ALL will be replaced by HBA_JOINED.
56	20100407	It was decided earlier to have only one pointing per station beam (fixed in time).
57	20100630	Dataslot allocation scheme at stations will be implemented after the beamsrver and calserver are done.

Last: 57

## ***Holidays***

Ruud: October or November

Ronald: 17 June – 9 July

John: 19-+3.5 week

Michael: 4-20 August

Hanno: 12 June – 6 August

Andre: 2-27 August

## ***Table round***

•