

## Minutes of Meeting LOFAR Software

<b>Date:</b>	2009-01-28
<b>Next meeting:</b>	2009-02-04 9:15-10:15
	Multimedia room
<b>Present:</b>	
Andre Gunst	Yes
Ronald Nijboer	Yes
Ruud Overeem	Yes
John Romein	Yes
Michael Wise	Yes

cc: Arnold Meijster, Rob van Nieuwpoort, Arthur Coolen, Jurjen Sluman, Pieter Donker, Chris Broekema, Martin Gels, Joris v. Zwieten, Marcel Loose, Adriaan Renting, Ger van Diepen, Max Avruch, Michiel v. Haarlem, Jan Reitsma, Ger de Bruyn, Arno Schoenmaker, Hanno Holties, Corina Vogt, Jan Noordam, Joe Masters, Lars Bühren, Dion Kant, Johan Hamaker, Maaijke Mevius, Sven Duscha, Jan-David Mol.

### **Remarks previous minutes**

- OLAP section: Instead of one rack, half a rack of BG/P is installed.

### **Announcements**

- LOFAR data processing school is 10, 11, 12 February. They make use of the offline cluster.
- A new developer in Amsterdam arrived (Ken Anderson). He will be working on the pipeline frameworks (replacement of Casey Law).
- A new postdoc is working in Amsterdam (Evert Rol). He will be working on commissioning the imager.
- John Conway sends a PhD student to Astron to work on commissioning of the imager as well (Fabian Batejet).
- In Nijmegen Arthur Crostenji is hired for the cosmic ray application.
- **WAN cabling is currently prepared and installed in February. From 16 February on the current cabling is removed which connects the stations to the BG/P. In that week we probably cannot observe and use the stations. The downtime of the BG/P and offline cluster will be minimized.**
- All LCUs minus two are accepted. The two failing LCUs are sent back.
- The CEP tender process is finally continued again.
- A complete cabinet is integrated in Emmen. The field conditions will not allow flattening the field (partly because of the water and partly because of the frost).
- The 5<sup>th</sup> of February a meeting will be held with primarily the German institutes for the WAN connection from Effelsberg to Groningen.

### **Action item overview**

ID	Date submitted	Description	Owner	Planned date	Status
78	20090121	Take care that as many observations as possible will be performed before half February. Michael will draft a list of observations.	Michael	20090122	Open
79	20090121	Setup meeting to discuss the next steps	Andre	20090128	Open

		concerning the pulsar pipeline.			
80	20090121	Make station – BG/P connection list.	Andre	20090128	Open

Last: 79

## ***Progress***

### **Stations (André):**

Achieved since last meeting:

- Arie is able to integrate and communicate with 2 \* 2 GByte memory modules on the TBB. Now two times more memory can be used than originally foreseen. The next step is to double this again. Also the use of unregistered memory will be investigated (cheaper, but less robust).
- The combs in the cross correlation matrix are solved in BP firmware version 5.5.
- The reason for a non-hermitical cross correlation matrix has most probably to do with rounding effects. The 54 bit words on the RSP board are translated to 32 bit words. Eric will modify this rounding.
- 180 degree phase shifts in the cross correlation matrix at 200 MHz is due to the fact that not all waveform generators start at the same second.
- The length information for the cables between the stations in the superterp are calculated by Arie. Now in total 1 km of cable will be ordered and laid down into the superterp at the same time as the electricity.

Problems / current activities:

- The HBA calibration waits for the first HBA field. This has to be postponed to the next step.
- Dips: Michiel Brentjens can currently predict the dips and the cause appears to be in the BeamServer software. The problem has not been solved yet.

Next actions:

- Continue with LOFAR20

### **OLAP (John):**

Achieved since last meeting:

- The correlator was tested on half a rack of BG/P. It turned out that the asynchronous transpose (that overlaps communication and computation) works much more efficiently than the synchronous transpose, especially on larger partitions.
- The correlator does not support arbitrary numbers of subbands. Whether or not a particular number of subbands is supported depends on the partition size. For large partitions, this becomes too restrictive. Supporting arbitrary numbers of

- subbands (up to 62 per RSP board) is not a fundamental problem, but rather something that requires careful implementation.
- The tied array beamformer is implemented on the BG/P, albeit not optimized for performance. The pulsar people are eager to implement more of their pipeline to the BG/P. We have to think carefully which processing steps can be moved to the BG/P without running out of resources. Currently there is a close collaboration between the users/commissioners and the developers involved. This speeds up the turn around cycle from development to commissioning significantly. Hence, we will give more priority to implementing parts of the offline pipeline of the pulsar application to the online section.

Problems / current activities:

- Martin is busy with CMake and MAC/SAS integration. Neither of them will be finished if we close this Step this month.
- John is busy with bug hunting on the BG/P. One of the issues was solved by using another linux real-time scheduler. Another issue remains: the linux kernels on the I/O nodes occasionally crash when under heavy load. This might be caused by a bug in the interrupt handler of the Ethernet driver or the collective network driver, but we are not sure, since debugging this is extremely difficult.
- Chris is mostly busy with the CEP tender, and occasionally working on the aforementioned kernel bug.

Next actions:

- Continue with LOFAR20

## **Offline pipeline (Ronald):**

Achieved since last meeting:

- Pandey finished a document about flagging on baseline length.
- Maaijke gave a presentation about the BBS ionospheric SPAM tests. The first tests have been finished.
- Facet corrected imaging is done by ATNF. Currently we are able to image in facets, but we cannot apply yet a correction per facet. The other ATNF tasks are not finished and lower priority is given on those.
- In the offline group a decision is made concerning the source modeling: Transient source modeling tool under Python will be used for source modeling.
- Define A-team subtraction is written in the strategy document. Also the clock phase correction is described there, which does not work in all cases. For these exceptional cases a strategy has to be defined.
- Adriaan has written a document about flagging on baseline length and station number. This is currently under review. After that this will be implemented.
- BBS global solver is partly tested (solved over multiple subbands). Tests for actual observations must be done.

Problems / current activities:

- Channel imaging takes more work then expected.
- Evert Rol will work on the regression test suite for the imager.
- Ronald is busy with a document about beam modeling.
- John Swinbank reported that Pyreb can be used as Python interface to Casa core now as well.
- Distributing parameter database will be finished at the end of the week.
- Joris is working on solution based flagging. He is implementing the method of Sarod for this.
- Imager regression tests will be defined by Ger and Malte.

Next actions:

- Continue with LOFAR20

### **SAS + MAC + SHM (Ruud):**

Achieved since last meeting:

- Currently the HBA tile temperatures are measured with temperature sensors in two tiles.
- Pieter tried to generate spikes to check if Andreas cosmic ray algorithm was working. However he did not succeed in it. This is not a desired situation. The cosmic ray group should start to commission these themselves. All work regarding the detection algorithms on the TBBs is on a hold until somebody is starting the commissioning of these.
- Switch MAC addresses: Ruud needs a list wherein the station to BG/P connection is written. Andre will take action on this.
- The database for the antenna coordinates is implemented including history. The input is an Azimuth file. This one is converted to ITRF and from that the antenna positions are generated in an appropriate format. Ruud is waiting for an example of the Azimuth file.
- The RSPDriver is currently able to read in cable delays and gains. Dependent on the RCU mode the right delays and gains are applied (as a function of frequency).
- Arthur was busy solving issues concerning the PVSS connection to a station. Normally you call a function dbconnect to connect to a station. The PVSS function does not check if the connection is succesfull. So, the application cannot assume that the station is indeed connected. Hence Arthur added checks for this (PVSS did not support this).

Problems / current activities:

- Metadata flow work is ongoing. Arno is busy with a mechanism to distribute the static meta data from SAS to the stations. The mechanism is written.

Next actions:

- Continue with LOFAR20

### **User Software (Michael):**

Achieved since last meeting:

- Prototype mosaicing script works with some issues still to be solved.
- Solar KSP starts to work on monitoring dynamic spectra.

Problems / current activities:

- Casey will design a set of scientific validation tests for the CIMAGER. Casey is busy to automate the first five tests.
- Lars is continuing to work on the CR near-field imager and is solving the last bugs.
- Sven is working the RM synthesis prototype

Next actions:

- Continue with LOFAR20

### **Software integration**

Achieved since last meeting:

- Marcel, Martin and Lars are busy merging to Cmake.

Problems / current activities:

- Compile a list of anticipated data products and calibration or metadata files associated with each of the pipelines. It is a task on the task list.

Next actions:

- A test program will be initiated to verify the functioning of the LOFAR software in a more structured way. In OLAP it is possible to store the raw station data and feed this into the pipeline later on. This makes it possible to define a standard data set, which can be applied to the pipeline as soon as major software changes have been taken place.
- 

### **Decisions**

<b>ID</b>	<b>Date submitted</b>	<b>Decision</b>
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^3.
41	20081203	We will modify the build environment to cmake from now on.
42	20080129	Transient source modeling tool under Python will be used for source modeling.
43	20080129	Delay deadline of Step 2 to 26 February 2009.

Last: 43

## **Closing Step 1**

Currently not all tasks have been finished as planned. The reasons for this are:

1. Christmas holiday of 2 weeks + 1 week extra for a couple of people
2. Step 2 started 1.5 week later than originally planned
3. People which were exchanged (especially in the USG tasks)
4. Re-ordering of priority for Cmake
5. Additional tasks popped up during Step 2 which have higher priority (for most of the cases because station roll out waited for these tasks to finish).
6. **Some** people do not feel the urgency since: (1) the station roll out is delayed and (2) minor observations are done because they stall due to bug fixing.

We cannot do much about 1. to 5.

So, how to cure reason 6:

- plan as much observations as possible before the network change in Groningen is installed.
- think of the fact that in the station roll out team everything is done to speed up the roll out when the field conditions are ok (so now frost, rain and eggs). One of the ways to speed up the roll out is that the cabinets are currently integrated in house (in Emmen). In this way we can transport full cabinets to the field and are independent of weather conditions. Furthermore for this activity extra personell is hired to integrate those cabinets. And yes, it takes a while before the full machinery is running but if it runs ...
- the full network is currently installed in Groningen and thereafter in the Concentrator node. As soon that works, hooking up stations will be a matter of

- patching the right cables and configuring the hardware. With other words things can go quickly then (provided the normal problems of course).
- the CEP tender is continued again, that means available horse power in a short time, lets be ready for that.

Sorry for this preach and hope that the sense of urgency will be back soon for “**some people**” (because I know that most people have already a sense of urgency). We rather prefer the statement: “the software is ready but the hardware is late” then the other way around for a change.

Due to all those reasons mentioned we decided to postpone the closing of Step 2 to 26 February 2009. The current status if we are closing Step 2 now will be send around. Let's turn the red bars in beautiful green colors.

### ***Table round***

-