Minutes of Meeting LOFAR Software

Date:	2009-04-22
Next meeting:	2009-04-29 9:15-10:15
	Multimedia room
Present:	
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, 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, Dion Kant, Johan Hamaker, Sven Duscha, Jan-David Mol, Teun Grit, Alwin de Jong.

Remarks previous minutes

-

Announcements

- CS302 is up and running. The complete LBA field is installed and 24 HBA fields are connected. The rest of the tiles will be connected this week. The BG/P connection works as well.
- One week ago the station in Garching was built.
- This week a station in Tautenburg is being built.
- Last Monday and Tuesday the offline and storage cluster have moved to the other side of the BG/P. This means that the racks for the new offline and storage machines can be cleaned.
- Developer from the SOLAR KSP is here this week (Frank Breitling). He will first work on the stand alone mode for the SOLAR application.

Action item overview

ID	Date submitted	Description	Owner	Planned date	Status
82					

Last: 82

Progress

Stations (André):

Achieved since last meeting:

- The time stamp bug in the TBB is solved. However, the firmware is not released yet.

- The work to get the ring splicer fully functional is more then anticipated for MAC/SAS. This will be postponed after Ruud's holiday. The work around is to work with the split HBA fields with 16 MHz bandwidth each.
- The 48 MHz requires a small change in the LCU software as well. Now maximal 54 beamlets per board can be generated, while we need 62.

Problems / current activities:

- The HBA calibration waits for the first HBA field.
- The LBA calibration verification waits for the first LBA field.

Next actions:

- Continue with LOFAR20

OLAP (John):

Achieved since last meeting:

- Support any number of subbands to make the software more flexible is done by Rob. However this caused a bug in the correlator software. Rob is busy solving this.
- Chris is working on the storage performance issues. There was a discussion with Ger and John about this. With the current performance numbers we will not be able to output all the correlated data. Hence an improvement is necessary. To improve the performance instead of MS files, raw binary files should be generated. Ger will provide a CASA interface to this raw format, in such a way that the flagger needs not be changed. The metadata will be in the same format. Only the visibility and the weights will be written in a raw format. The raw data format could also speed up the offline processing tasks. However, then a raw data format data writer is necessary as well in CASA.
- Jan-David and John are working on the beamformer assembly implementation to optimize the tied array beamformer. The bandpass correction was done after the correlation. However this is moved back right after the polyphase filter. Then, it can be done for the full subband instead of at channel level. The estimated optimization is 30-40 times.
- Paper about fast communication between IO nodes and compute nodes for a conference is accepted.
- Correlated paper about performances for supercomputing is submitted.
- Yesterday Arno scheduled an observation using a couple of stations and the OLAP part. The feedback from correlator to MAC/SAS seems to work.
- Martin transformed the BG/L getstats script to a BG/P getstats script.

Problems / current activities:

- There were some issues on cross compilation (building the software for different target hardware than on the system you are actually building on). The tools currently check automatically if boot libraries are available or not and decide to

use it. Sometimes this is not desirable and causes problems. Hence a way of manually overriding this mode is necessary.

Next actions:

- Continue with LOFAR20

Offline pipeline (Ronald):

Achieved since last meeting:

- Pandey works on the region based flagging (=table based flagging). If you know that certain stations are not behaving, than these will be flagged.
- Joris made a Python implemention of the clock phase correction implementation.
- The Australian imager will not be in time for MS^3. Currently we are thinking about using the CASA imager as a backup. Ronald discussed this with Ger. Features like faceting, polarization, channel imaging are already in. However, the CASA imager is not developed yet for distributed processing. The master worker shell is missing. Hence deconvulation can be done only at subband level. Also other features like beam correction must be added.
- Ronald had a discussion about data formats. Ger was not worried about the performance of HDF5 files. The worry of Ger is the required changes of all the programs currently using the MS. This wiIl take a significant amount of time. We have to make an inventory about its pros and cons.

Problems / current activities:

- Global bandpass determination needs a new observation. This can be done only when the self generated RFI at CS010 is gone.
- Joris is busy with solution based flagging.
- John Swinbank is integrating the standard imaging pipeline. He is integrating that with i-Python.
- Evert Rol will work on the regression test suite for the imager. He is currently setting up a test framework and uses BBS as prediction. This will ultimately result in an end to end test loop.
- Ronald is busy with a document about beam modeling. Done so far. Now it waits for real data.

Next actions:

- Continue with LOFAR20

SAS + MAC + SHM (Ruud):

Achieved since last meeting:

- Pieter worked on the weather station, which is currently working.
- Pieter is still working on updating the TBB driver.

- Arthur upgraded SAS/MAC in order to support the new way to specify the antennas in a station (you do not have to specify individual antennas but configurations like inner and outer ring).
- Enhancements of the control framework of MAC have been done which simplifies coding.
- This week Ruud works on analyzing the log streams.
- The coordinate files of Corina can be read.
- There is now a beamserver which takes ITRF antenna coordinates. A 6 hour observation is done. When it is correct, it needs to be optimized. It waits for checking of Michiel Brentjens.

Problems / current activities:

- Both the BF data writer and TBB data writer are not ready for integration into MAC/SAS in this Step.
-

Next actions:

- Continue with LOFAR20

User Software (Michael):

Achieved since last meeting:

- The GSM database implementation is done. It is also hooked to the second part of the pipeline. In fact the pipeline is broken in two: the second gets images which goes through the source detection, while the first one is the flagging, BBS and the imager.
- Testing the accuracy of the source detection module is not done yet.
- Meeting was held about dataformats. Ken is cleaning up documentation and making sample files.
- Lars made upgrades and fixes to the TBB datawriter.
- The imager of Lars works.
- Joe has a first version of the Visit plugin. With this plugin you can visualize the data in 3D.

Problems / current activities:

-

Next actions:

- Continue with LOFAR20

Software integration

Achieved since last meeting:

- Marcel needs still at least a month to finish the integration with 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

ID	Date submitted	Decision
02	20061220	Every Step will start with a Kick-off meeting, in which the complete software team participates.
~~03~~	~~20061220~~	~~The project team starts immediately with the preparations of the next CDR in order to preserve progress of the CS1 realization~~
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 µStations/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	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.

Last: 44

Holidays

- Ruud is off in May for 5 weeks
- André is off in June for 4 weeks
- Ronald: 26 August – 14 September
- John: 17 May – 22 May, …, and 12 July – 18 July
- Michael: 1 June – 5 June

Table round

- André: There were some “rumors” why we do not use a more recent version of Ubuntu for the offline and storage cluster. The reason for this is that till now Ubuntu 8.04 was the most recent version. We are already preparing machines with this OS and do not go to a newer OS.