

Minutes of Meeting LOFAR Software

Date:	2008-07-02
Next meeting:	2008-07-09 9:15-10:15
	Paviljoen West room
Present:	
Andre Gunst	No
Ronald Nijboer	Yes
Ruud Overeem	No
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 Bähren, Dion Kant, Johan Hamaker, Maaijke Mevius

Remarks previous minutes

- The comment in the previous minutes about how the TAB mode will deal with missing station data was inaccurate. A decision about how the system should deal with data dropouts has not been decided yet and is still being considered.

Announcements

- BG/P is installed in Groningen. Successful link-back tests have been conducted; however not all the switches have been installed.
- A GSM design meeting was held on 2008-07-02.

Action item overview

ID	Date submitted	Description	Owner	Planned date	Status
49	20080116	Simultaneous data storage of TBB and in OLAP to validate inverse poly phase filter bank of Kalpana. [Update] Lars and Adriaan have tried to track down the data from this observation. There appears to be no entry in the observation log. Clearly we need a more robust procedure to request and track observations.	Michael/Andre	20080204	Open
60	20080528	Plan GSM meeting. A date is suggested.	Michael	20080604	Closed
62	20080603	Design/decision stand alone mode of E-LOFAR stations. The idea is to reconfigure the switch by a SNMP script when we are going from an observation using all stations to an observation excluding the EU stations. The script changes the routing in de switch for security reasons. The machines of the international partners can be connected to the	Andre	20080611	Open

		LCU via VNC access. The VNC server will be brought down if we are going to non-stand alone mode.			
63	20080611	Discussion with RedHat about LCU licenses for all LOFAR stations.	Ruud	20080730	Open
64	20080611	We need MAC addresses with an unique vendor ID to prevent problems with other hardware in the future. A vendor ID is requested by system management.	Andre	20080618	Closed
65	20080611	Define validation tests for the imager which can be executed by for example Casey.	Ronald	20080618	Open

Last: 65

Progress

Stations (André):

Achieved since last meeting:

-

Problems / current activities:

- Eric is busy implementing the 16 bit to 4 bit option. Currently the problem is that the FPGA is almost completely full.
- 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.
- Long distance delay tracking observations have been done. James Anderson will look for fringes.

Next actions:

- Step 5
- Finalize steps for LOFAR20

OLAP (John):

Achieved since last meeting:

- John has been working on the communication protocol between the compute and I/O nodes. This layer replaces ZOID. Protocol is basically finished and seems to achieve the desired wire speed. Additional testing is still needed.
- Chris has been busy working on the BG/P network design.
- Martin has been working with Lars to investigate the CMAKE build environment. Initial experiences are very positive. Much of the OLAP software can now be built using CMAKE. Build performance is much better with a speed-up on the order of a factor of 10. Not all of the functionality of the RUB system has been replicated in CMAKE yet however.

- Martin has finished the raw data conversion program which was originally running on the BG/L. That means raw station data can be recorded and exported to users, while the BG is down.
- Rob is holiday for two weeks.

Problems / current activities:

- Rob busy implementing the TAB mode, especially for the superstation. The issue of how to best handle data drop-outs is still being explored.
- Robustness for failing disks is not included yet (Arnold Meijsters).

Next actions:

- Step 5

Offline pipeline (Ronald):

Achieved since last meeting:

- Adriaan has a new version of the DP³ tool ready which fixes the segmentation fault issue. Pandey plans to test the new version in the coming week.
- Pandey has been testing the online bandpass correction. There are still issues with the procedure and he sees remnants of the polyphase filter in the data at the 5% level. Possible problem with the filter coefficients.
- Ger has successfully built a parallel version of the CIMAGER on his Mac. He plans to build a parallel version on the cluster in Groningen in the next few days. Testing is still needed.
- Ger and Alexander have continued to discuss the design for the mosaicing tool based on the AIPS++ code. One issue is where will the information on the primary beam originate?
- Marcel is on holiday.
- Joris starts his holiday for two weeks.

Problems / current activities:

- Marcel is busy with low level adaptations for the ACC.
- The distributed imager runs only on one node, which was not the intention. Appropriate action is needed to distribute the imager.
- Pandey will commission the Global Solver.
- Joris is busy to implement Johan's beam model.
- Treatment of parameters in BBS and mapping of parameters to grid needs to speed up.
- Ger discussed the facet imaging and how to give that priority for the Australians.
- MS description file lists how the data is distributed over machines. Similar files are necessary for PAR tables and images. The global controller needs this information.

- Ronald is continuing with the strategy document and has identified four different approaches.
- Maaijke will implement SPAM in BBS (needs fitting on non-UV data). Is still busy testing if you have to do peeling or making simultaneous solutions.
- Pandey will test the HBA dipole beam model in BBS

Next actions:

- Step 5
- Step 6: commission the Global Solver

SAS + MAC + SHM (Ruud):

Achieved since last meeting:

- Ruud's suggested MAC/IP address scheme is approved.
- HBA delays can be read back. The send delays and the used delays at the tiles can both be read back.
- Stefan and Pieter have tested all basic functions. Two routines out of six are build and tested already with the basic functions.
- Arthur and Ruud are almost finished by getting all information to the database (except for 3 MAC controllers). All important logging and all hardware can be seen and are validated in the Navigator 2.0. The station information can be for 95 percent read and also the alarm system works. The archiving works as well.
- MACInformationServer is modified for SHM. Now, suspicious components detected by SHM can be coupled to the MAC subsystem.
- Distribution of PVSS data to SAS works.
-

Problems / current activities:

- The information of the MCU should be connected to the database too.
- As it stands now: real significant different observations on BG/P has as a consequence that the RSP Driver must be restarted with another configuration file. This is the case if the station data must be send to different IOs of the BG/P from observation to observation.
- Ruud made a proposal to change the MAC/IP addresses. System management requested a vendor ID. The current proposal is:
as : tr : on : xx : yy : zz, where astron is the vendor ID.

xx: the LOFAR physical location

- 00 = core
- 01 = arm 1
- 02 = arm 2
- 03 = arm 3
- 04 = arm 4
- 05 = arm 5

06 = EU

yy: following number from 1 to 127 maximal

zz: the board ID

00: RSP0

01: RSO1

.

.

07: RSP7

20: TBB0

21: TBB1

.

.

2B: TBB11

- Metadata flow work is ongoing. Arno is busy with a mechanism to distribute the static meta data from SAS to the stations. It would be great if that can be finished as well in Step 5.
- Test programs for the RSPs should be started by the LCU. This is necessary for the assembly company which is building the subrack.
- Extra status registers which are implemented in the FPGAs are not driven yet by the LCU.
- The RCU is currently able to drive the power supply of the HBA independent of the modem (which is desirable from an EMC perspective). The remainder left is that the LCU can also drive this bit.

Next actions:

- Step 5
- Step 6: software suitable for ITRF coordinates, coordinates available centrally in a SAS database.
- Wish list: horizon vector in the stations, beamtracking suitable for the solar system.

User Software (Michael):

Achieved since last meeting:

- Lars has been supporting Martin with testing the CMAKE build environment on the OLAP software. Good progress has been made.
- Lars has also been familiarizing himself with the protocol for making TBB dumps so he can support Andreas and the CR team with data collection when necessary. He has also been working with Adriaan to try and track down the missing Kalpana dataset.

- Joe is back from paternity leave. He has been finishing up the update of the DAL to use version 1.8 of the HDF5 library and continuing to work on the database access layer.
- Sven has produced a first version of the specification for the RM synthesis tool. A few issues need to be settled before moving on to the implementation.
- Lars, Joe, Casey, John, and Sven met to discuss the image cube format. Based on that meeting, a new LOFAR image cube ICD has been drafted. The current draft is up on the USG web page. It still needs some work.
- Casey, John, and Bart demonstrated a prototype for the GSM database implementation. The demo included all the major radio catalogs and utilized a MySQL database. A web interface as well as python access routines were shown.
- Alexander supported the LIONS group with getting MeqTrees updated in Leiden.

Problems / current activities:

- Members of the Magnetism KSP are exploring using the VisIVO package to visualize RM synthesis cubes.
- Alexander is writing up a design for a radio image cube mosaicing tool based on the AIPS++ routine.
- Lars is busy with the near field imager (skymapper).

Next actions:

- Step 5
- Finalize steps for LOFAR20

Holidays

Ruud: 26 Juni – 15 Juli 2008

John: From ~26 August 3 weeks and structural one day off from ~26 August onwards.

Ronald: 25 August to 12 September

Andre: 25 June – 18 July

Michael: 16-27 June

Software integration

Achieved since last meeting:

- Martin, Marcel and Lars are busy to compile LOFAR code with cmake.

Problems / current activities:

- A request of Michiel Brentjes was to keep a clean trunk. John suggests to tag the production release.
- Marcel is busy with streamlining the build environment

- Compile a list of anticipated data products and calibration or metadata files associated with each of the pipelines.
- LOFAR development software needs to be build in Kubuntu (Michael has volunteered)

Next actions:

- Define the length of Step 4.
- Step 2+: 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.

Scaling up work for 20 stations

After Step 5, one full week will be used to update the documentation of the software. After Step 5, we will aim to LOFAR-20 and hence we start from Step 1 again.

Current list of tasks to do and when:

Stations:

- For the stations the split HBA field boards must be produced and tested in the field. Step 2.
- The firmware for the HBA field split must be written. Step 2.
- Additions in the LCU software for the HBA field split are required. Step 2.
- One central clock should be implemented in the super station. Step 2.
- HBA calibration should be in place. Phase 2.
- Definition of test suite to be run by the LCU. Step 1.
- ARP in the TBB. Step 1.
- RSP boards suitable to send data when message of LCU is. Step 1.
- Conversion to ITRF coordinates. Step 1.

OLAP:

- Should be made suitable to handle multiple independent RSP datastreams without correlating all of them (now OLAP correlates also the different RSP datastreams from the same station to support microstations). Step 5.
- The network surrounding BG/L should be modified in order to split BG/L in four independent partitions. Will be done after the BG/L decision. Step 5.
- Possibly optimizations in the IO nodes to cope with the 32 MHz bandwidth are necessary. Will be done after the BG/L decision. Step 1.
- Making the storage section fault-tolerant and more efficient by using multiple compute cores. Step 5.
- Tied-array beamforming definition will be done in Step 5 and implementation in Step 1 for super stations.
- Multiple TABs: Step 3 ...

- Snapshot imaging support (changing beams in time). Step 2.
- Multiple observations or one observation with multiple beams which change as a function of time. Step 3+
- Porting correlator to BG/P (incl. ZOID communication). Step 1.
- Other observation type support:
- Conversion to ITRF coordinates. Step 1.
- Script on central machine to check station broadcasts (Chris?). Step 1.

Offline:

- Station model. Step 1.
- Beam validation. Step 2.
- Implementation station beam. Step 3.
- Ionospheric approach for 20 km distance between stations
- Multi-beam processing for clock calibration. Step 1.
- Will be detailed in the offline meeting.
- Conversion to ITRF coordinates. Step 1.
-
-

SAS/MAC:

- Updating to Navigator 2.0 to enhanced functionalities in Step 1.
- Making screens for the 20 station configuration. Step 2
- Proof reaction times. Step 1.
- Ring splitter control. Step 2.
- Performance measurements triggers. Step 1.
- Scheduler? Step later.
- Switch MAC addresses flexible. Step 2.
- Reconfigure international switches for stand alone mode (starting SNMP script)
- Hardware monitor and software monitor for central machines. For BG/P it is necessary to check if it is still up.
- MAC uses feedback from SHM. Step 3.
- Conversion to ITRF coordinates. Step 1.
-

USG:

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 CSI 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 μ 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.
--	--	---

Last: 34

Table round

- Ronald: Can people at the VU work on the scheduler? Ruud: They could do research into sorting algorithms. We are not using real-time scheduling for networks.