Minutes of Meeting LOFAR Software

Date:	2010-02-24
Next meeting:	2010-03-03 9:30-10:30
	Multimedia room
Present:	
Andre Gunst	Yes
Ronald Nijboer	Yes
Ruud Overeem	Yes
John Romein	No
Michael Wise	Yes
Harm Munk	Yes
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

-
-

Announcements

-

Action item overview

ID	Date submitted	Description	Owner	Planned date	Status
99	20100113	Report on release management. For station software a procedure is established but not written down yet.	Harm	20100127	Open
104	20100224	Verification test suite should be defined to test the storage manager.	Ronald	20100315	Open

Last: 103

Progress

System Integration

Achieved since last meeting:

- Number of validated stations is now 21.
- ITRF beamserver progress: Two new antenna position files have been received from Michiel Brentjens with new coordinates.

- Not much progress in the CalServer. Currently Ruud prepared a new development machine which is equal to the systems in the field.

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.

Imaging Pipeline (Ronald):

Achieved since last meeting:

- Adriaan is working on the MS definition of LOFAR.
- Preparation on the dress rehearsal is ongoing.
- John Swinbank is making progress to interface the pipeline with MAC. Currently he is communicating to an echo server. Communicates not yet to a MAC controller. For the dress rehearsal connection to MAC is essential. Hanno likes to have a connection to MOM as well primarily for the data challenges to ship the data to the archives during the dress rehearsal. The aims of the rehearsal are (1) take the data and process it well, (2) running the pipeline and checking if the headers are filled with the right metadata and (3) to archive the processed data.
- 40% performance gain was reached in DPPP by Ger van Diepen. Multiple steps in a single run can be done now. The parset file connected to DPPP changed quite a bit. Parameters have new names and are simplified.
- Make source db was running faster on the front end nodes than the compute nodes. We have to find out why?
- Multiple read/write performance should be tested for the second phase hardware including predefined access patterns.
- The issue with the UVW coordinates turns out an issue in the storage manager. First there was a sign flip (fixed) and the storage manager uses a different phase centre than BBS. This is probably also the cause for the difference between Casa and diff map. Hence urgent action is required to define a test suite which can be used as reference. Before any release all the tests in the test suite should be passed. This saves a lot of debugging later on.
- Major cycle is done for source py components.
- Bas thinks a first ionospheric calibration will be ready for the dress rehearsal.
- Evert investigated the speed difference between Casa and Cimager. By changing parameters it can run as fast as the Casa imager. However the reason for this is not quite clear yet.
- Hanno Spreeuw wrote a document about the source finding module.
- Bart Scheers has been busy with the sky model database of the LSM.
- Evert wrote a script around the Casa imager to use it in the pipeline.

-

Problems / current activities:

-
-

Next actions:

- Focus on the minimal required tasks for MSSS.

Pulsar Pipeline (Michael):

Achieved since last meeting:

- Second transpose will be delayed because of the absence of Jan-David. Hence online version is on hold and we are going to focus on the offline version (storing data from the station directly to the storage nodes).
- A first version of the second transpose is implemented.
- The format of the input stream was communicated.
- Jan-David was also working on the MAC/SAS connection.
- Lars finished the data classes and Lars and Anastasia are busy testing it.
- All offline components are wrapped.
- Wrapping in pipeline framework is ongoing. Ken Anderson is now working on a single threaded version.

Problems / current activities:

-

Next actions:

- Implement second transpose operation.
- Update BF datawriter.
- Integration of pulsar beamformed observation in MOM via a Parset from SAS.
-

VHECR Pipeline (Michael):

Achieved since last meeting:

- The TBB control is in good shape.
- The LCU trigger code is in test.
- Joerg was here as well last week to talk about the particle detector. The particle detectors produce another stream of data. Its additional metadata. This needs to be defined. The software to produce it is there. The code is already on the repository.

- There is a lot of trigger information which should be attached as metadata.
- The TBB datawriter is more or less ok. The same underlying infrastructure is used for that.
- Martin is working on the offline components and wrapping them up.
-

Problems / current activities:

-

Next actions:

-

Infrastructure (Harm)

Achieved since last meeting:

-
-

Problems / current activities:

-

Next actions:

-
-

User Data and Archive (Hanno)

Achieved since last meeting:

- This week we were succesfull from proposal to data capture during a busy week. There were some issues with the station software and later on with CEP.

Problems / current activities:

- Determine where data is stored and from there we can ship it to Amsterdam.
- CEP datamanagement/uses. For a large observation we need a lot of storage nodes (half of them). Rethink the subcluster assignments. Antonis will come with a solution.
- Resolve open issues with the archive.
- A quick look at Juelich: data transport protocol is not working yet.
-

Next actions:

-

Decisions

ID	Date	Decision

	submitted	
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.
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.

Last: 49

Table round

- Michael: We need an issue tracker soon. Antonis is pushing this, but Hanno will push as well.