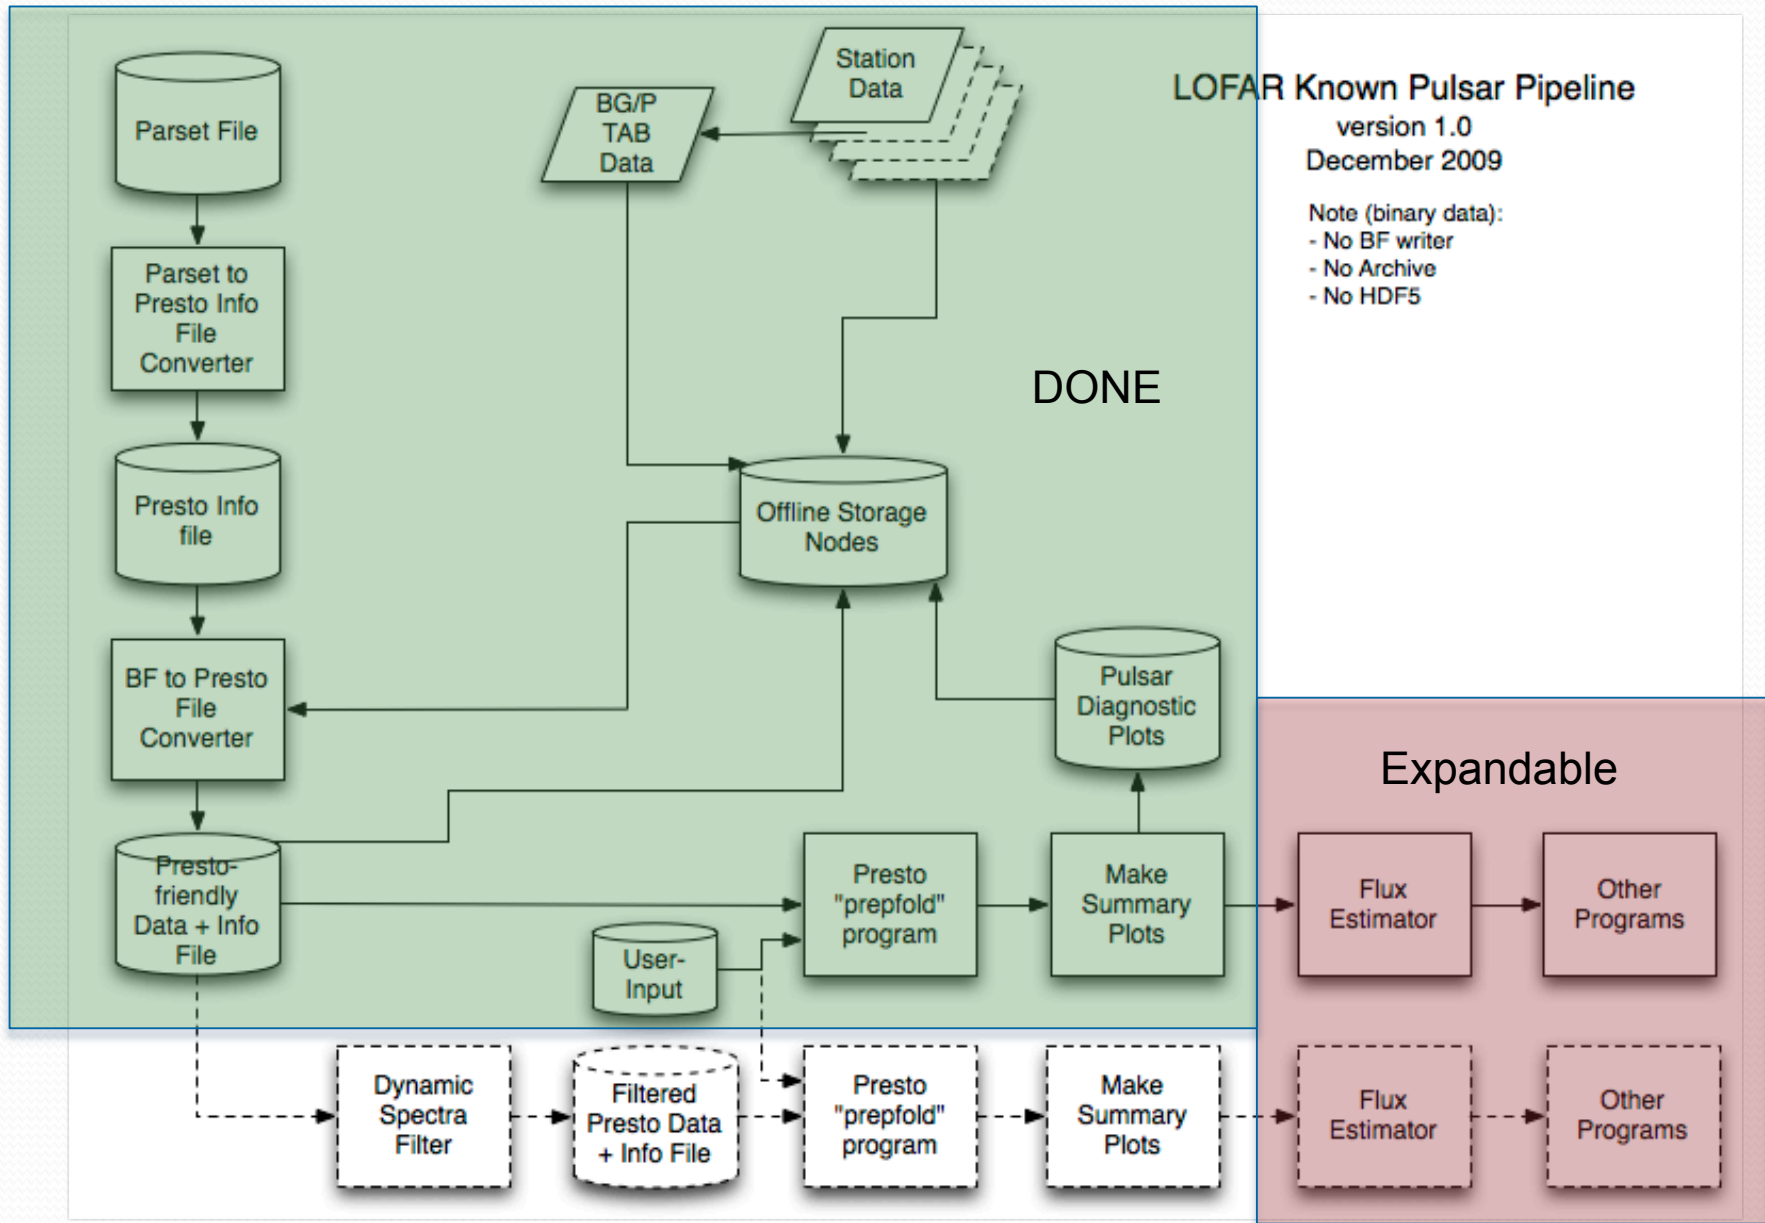


Build/Install/Run the USG pulsar software & pipeline

- Presentation on Pulsar Tools/Pipeline “How To” Build/Install/Run given during the Transients Busy Week (Apr 7, 2010)
LOFAR wiki, Meetings..., Transients - Ongoing Work
<http://www.lofar.org/operations/doku.php?id=science:ksp:transients>
- Pulsar Off-The-Shelf Tools:
 - Single precision FFTW v3.1.2
 - TEMPO v1.1 (<http://www.atnf.csiro.au/research/pulsar/tempo/>)
 - Presto (<http://www.cv.nrao.edu/~sransom/presto/>)
 - SIGPROC v4.3 (<http://sigproc.sourceforge.net/>)
 - Python bindings to PGPLOT (ppgplot) v1.1 (<http://www.astro.rug.nl/~breddels/python/ppgplot/>)
 - [PSRCHIVE v13.0 (<http://sourceforge.net/projects/psrchive/>) - Ubuntu Linux only]
- Pulsar Pipeline & associated scripts:
 - LOFAR BF data processing scripts, programs, BF offline **shell script “pipeline”**
 - Pulsar Pipeline within the Pipeline Framework is progressing (communication with SAS/MAC, multi-core and distributed)
- Pulsar tools and pipeline integrated into USG repository (Daily Build)




Running the Pulsar “Shell-Script” Pipeline

- Examples with processing times for typical 1-hour observation (note, this does not scale to N-hour observations); multi-core script:
 - 1) Basic run with no additional switches; uses 8 cores → takes 5 min
 - `make_subs_SAS_Ncore.sh -id L2010_o6296 -p B2111+46 -o /net/sub5/lse013/data4/LOFAR_PULSAR_ARCHIVE_lse013/L2010_o6296_red`
 - 2) Basic run using 4 cores; deleting previous results → takes 10 min
 - `make_subs_SAS_Ncore.sh -core 4 -del -id L2010_o6296 -p B2111+46 -o L2010_o6296_red`
 - 3) Run with ALL subbands option and RFI option turned on → takes 17 min
 - `make_subs_SAS_Ncore.sh -id L2010_o6296 -p B2111+46 -all -rfi -o L2010_o6296_red`

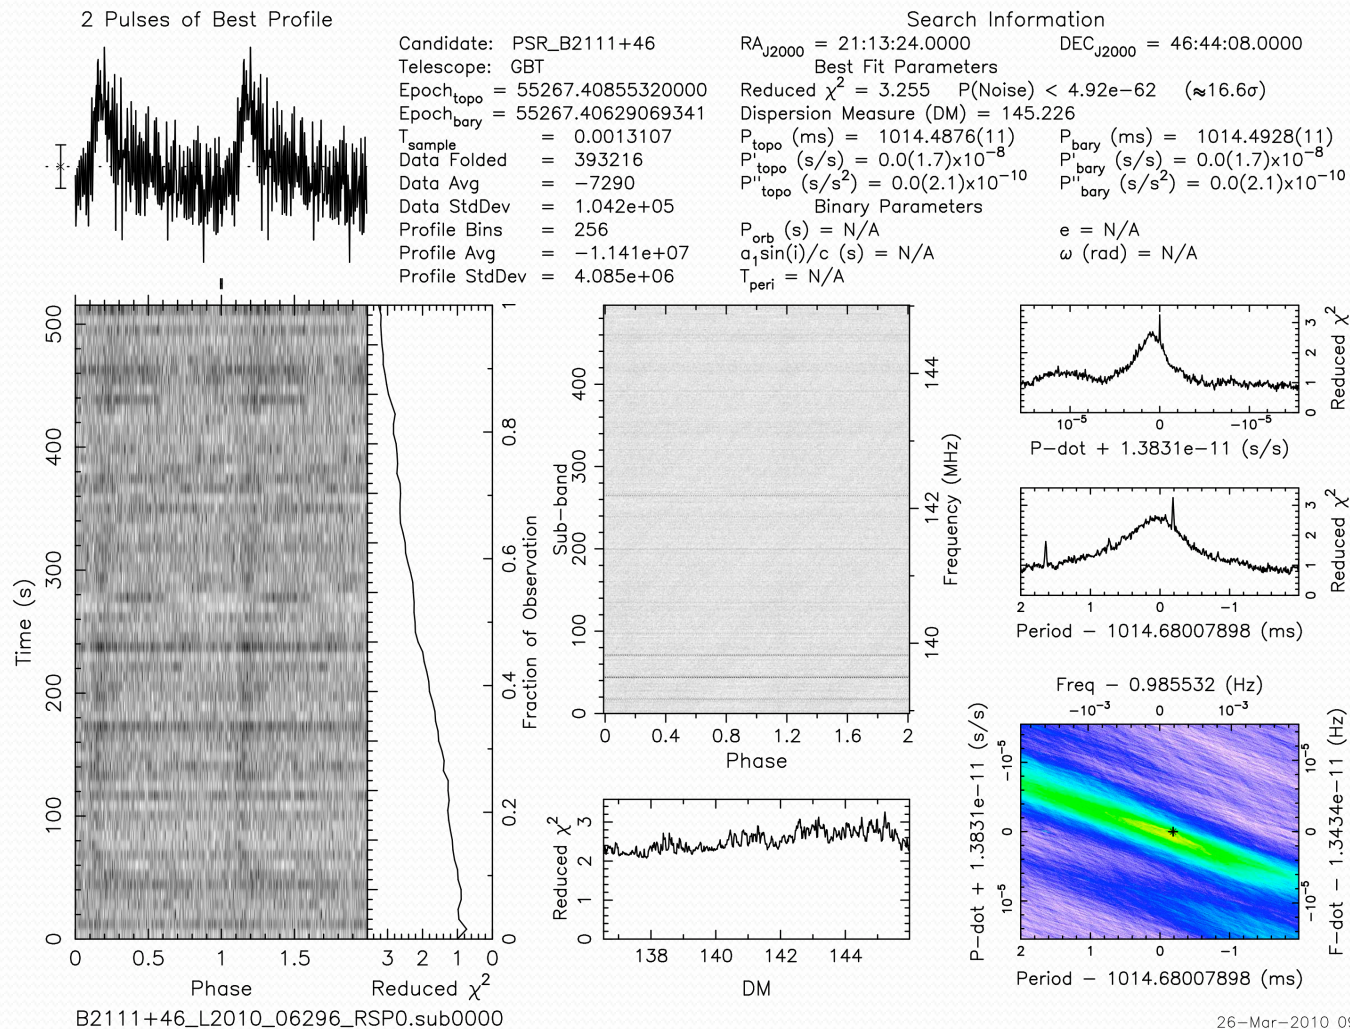
- Pipeline Output Products:

- Pipeline logfile (command summary, pipeline version, processing time)
- Summary plots of N-core splits of pulse intensity
- RFI-summary performed per split to show noisy subbands (channels)
- Presto pulsar profile plots



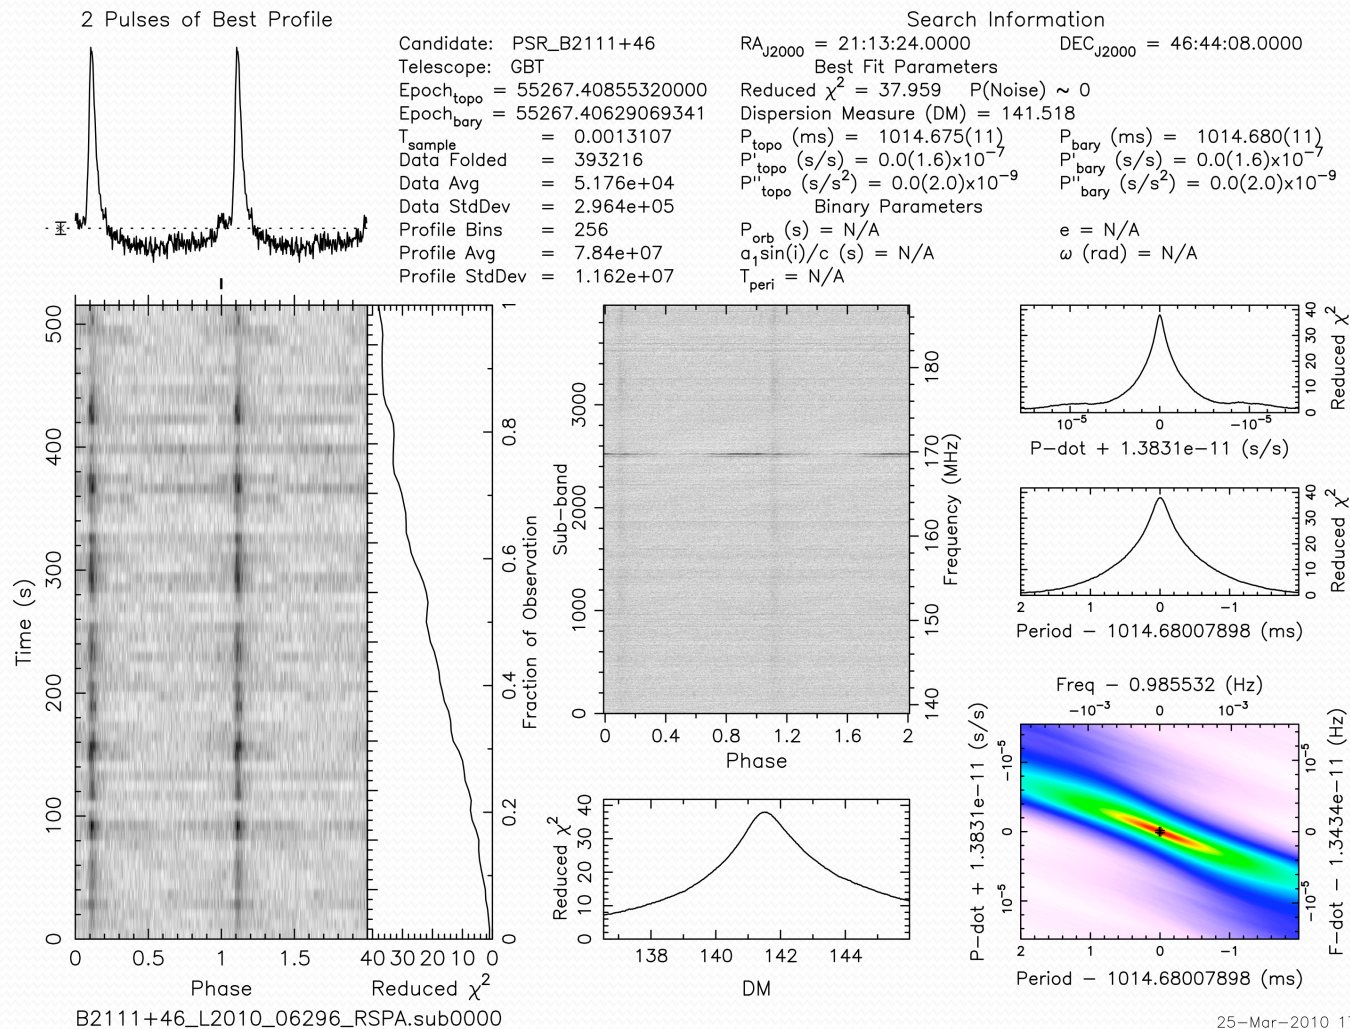
Split Results - Pulsar "Shell-Script" Pipeline

- 1 of 8 splits (31 Subbands of 248) – Presto’s prepfold summary:



“All” Results - Pulsar “Shell-Script” Pipeline

- “All” Subbands – cumulative Presto’s prefold on all 248 subbands (x 16 channels per subband → 3968 files to fold)



Pitfalls: Imaging Pipeline & LOFAR Software

- Daily Build (use Loflm) can be “in-process”/fail, but there is no warning that you should use previous build; tracking down this problem takes time. Suggest a STABLE Build (weekly/blessed/well-used by commissioners/users); suggest automated test-suites with the daily build.
- Pipeline Framework’s user-environment and setup was evolving until recently; also heavily depends on the success of Daily Build. System needs to be fully documented; users need to understand that this is a heavily evolving system.
- Usually end up using other peoples’ configuration and parameter files, which may be out of date, missing information, wrong altogether, belong to specific software version or specific sources, etc. Configuration files / Parameter files / Cluster Description files for Pipelines and tools all need to be grouped into annotated examples and placed into one location for reference.
- Communication of known sw and hw problems is limited and seems to be in-house. (i.e. ‘conjugate’ Imaging data run through DPPP before Apr 2, 2010; i.e. system boost libraries updated, breaking many of the Daily Build tools). Use the Forum to announce issues, general problems, solutions, etc; software group could perhaps give regular updates at the LSM.
- You need an army of experts to help you process your data! Be prepared for 100’s of emails/parameter-settings/opinions/failures/iterations.

Pitfalls: Imaging Pipeline & LOFAR Software

- NFS/cluster-node blips can bite you at any time and send you down the wrong error-chasing path; **sometimes re-running will make things magically work!**
- Processing 248 imaging subbands with hundreds of possible parameter settings is daunting for those starting out. **Start small – process one subband and use the Imaging Pipeline to iterate and then process all the subbands once you determine optimum settings/parameters.**
- Slew of home-grown tools in people's home directories – not visible across sub-clusters and no good to people who don't know about them. (i.e. ~rol/sw/bin is mentioned in the Imaging Cookbook, but user's files are only accessible on sub3 nodes; **expanded to sub5 nodes**). **Suggest all home directories be visible from all subclusters; user tools should be in common repository.**
- Limited cross-sub-cluster access for “Imaging” vs “Pulsar” folks prevents Imaging scientists from helping Pulsar group with our data processing, since they cannot log onto Pulsar subclusters. **Open login access to all subclusters to larger group of users.**
- Large datasets are spread across all sub-cluster nodes; sub-cluster login limitations prevent users from processing their data until the LOFAR support staff can move them by hand. **Automate/script the process of consolidating large observations.**

“Successes”: Imaging/Pulsar Pipeline & LOFAR Software

- Imaging Pipeline can successfully be run on these data, and by someone other than John Swinbank
- Pipeline Framework setup/environment “How To” documentation is coming together (soon)
- Parameter sets for non-bright, not in the phase-center, images are being put together for future reference (prep for MSSS)
- Imaging Tools are being exercised by non-imaging folks, on non-imaging subcluster nodes
- Bugs are being flushed out (Imaging & BF)
- Imaging Cookbook corrections are being made
- Pulsar BF data can be checked in ~6 min after typical observation
- ...
- Pulsar Group is learning about Imaging with some (slow) success!

Thanks for all the assistance!