SKA Phase 1 Computational Costs 2010-08-27 Workshop Notes

	Speaker	Description	Comments and Actions to be Taken?
1.	Duncan Hall	Presentation: "SKA Phase 1: Costs of Computation"	
2.	Gerry Harp	Perhaps should consider requirement for more rapid dump rates from correlator to match potential requirements to track Geostationary Earth Orbiting (GEO) and Low Earth Orbiting (LEO) satellites and "space junk."	Unclear at this stage, as there are no expressed requirements for GEO or LEO tracking for SKA Phase 1 or subsequent build outs. **Required Action**: Please provide some quantification of the costs versus benefits of such a proposal.
3.	Jan Noordam	Asserted that "we can deal with arbitrary amounts of smearing" and still meet a Dynamic Range (DR) target of 65 dB. This would be advantageous in that the data rate input to the Calibration and Imaging High Performance Computer (CIHPC) could be reduced.	While smearing may be removable, such image processing would require more – probably much more – computation. **Required Action**: Please quantify the level of additional computation required and compare that with reductions in computation resulting from lower input data rates.
4.	Tony Foley	Are 12 bit floating point numbers appropriate to cater for the "headroom" required for Radio Frequency Interference (RFI)?	Unclear at this stage as RFI characterisation for the sites has not yet been completed. **Required Action**: To be considered as part of Receptors and Signal Processing domains?
5.	Jasper Horrell	There have to be computationally cheaper ways of calibrating and imaging.	Correct. A motivator for running this workshop.
6.	Tony Willis	Forms of analogue computation such as sky de-rotators and Focal Plane Arrays (FPAs) should be considered to reduce the computational requirements. [Later in discussion]: what about use of optical computing for e.g. Fourier transformation?	**Required Action**: Please provide some quantification of the costs versus benefits of such proposals.
7.	Ger van Diepen	The size of the m x m subgrid is dependent on required baseline length and field of view.	**Required Action**: Please provide some quantification of the relationships of baseline and field of view to how large m x m must be.

	Speaker	Description	Comments and Actions to be Taken?
8.	Maxim Voronkov	The Discrete Fourier Transform (DFT) may provide some benefit in place of gridding and Fast Fourier Transform (FFT).	**Required Action**: Please provide some quantification of the costs versus benefits of such proposals. A discussion held at the whiteboard – in Dutch – subsequent to the workshop indicated that the computational costs of a DFT approach are not significantly lower than gridding / de-gridding and FFT, refer to the first landscape photo below.
9.	Maxim Voronkov	Rather than talk about "major cycles or loops" we should use the term "iterations over observed data."	Good point. To be adopted in future discussions of computational costs.
10.	Maxim Voronkov (?)	Sanjay Bhatnagar has some thoughts on how to reduce the costs of computation, especially in respect to how many iterations are required over observed data.	**Required Action**: Duncan Hall to contact Sanjay to elicit Sanjay's thoughts.
11.	Ronald Nijboer	"Required iterations over observed data" depend on the fidelity of the Global Sky Model (GSM) and Local Sky Model (LSM) being used for calibration and imaging.	**Required Action**: Please provide some quantification of the relationships of sky model fidelity to how many iterations are required over the observed data.
12.	Jan Noordam	Subtraction of interfering sources in the u-v domain allows "sloppy imaging"	**Required Action**: Please provide some quantification of the costs versus benefits of such proposals.
13.	Maxim Voronkov	Observations to detect transients require wide fields of view, so are computationally expensive.	Good point. **Required Action**: to be borne in mind when science prioritisation occurs.
14.	Maxim Voronkov	The relationship between receptor pointing error and best achievable DR is instrument specific.	Understood, however some idea of the order of magnitude of the relationship for extant instruments would be useful. **Required Action**: Duncan Hall to contact Sanjay Bhatnagar to confirm findings from EVLA.
15.	Not recorded	We need to build up learning experiences from LOFAR and other precursors for perhaps several years before specifying aperture plane arrays for Phase 1.	Good point. To be adopted in future discussions of computational costs.

	Speaker	Description	Comments and Actions to be Taken?
16.	Jan Noordam	u-v coverage has to be very complete to push down side lobe responses; the Point Spread Function (PSF) level of side lobes must be <0.01% which is derived from the square root of the maximum noise to achieve the required DR of order 60 dB (10^6 to 1).	**Required Action**: Duncan Hall to contact Rob Millenaar to discuss Rob's work in this area.
17.	Maxim Voronkov	To achieve high DR – i.e. greater than 60 dB – we will have to consider relativistic effects, as well as tectonic movement within 8 hour periods over longer baselines.	Maxim delivered a presentation which will be made available on the CALIM 2010 Web site. Further information is available at Maxim's personal Web site at: http://www.narrabri.atnf.csiro.au/people/vor010/ **Required Action**: Duncan Hall to include such considerations in estimating costs of computation for SKA Phase 1 and subsequent build outs.
18.	Johan Hamaker	The Jones matrix solutions are not unique: there is unavoidable ambiguity in their solution set. Software solvers will simply pick one of the possible solutions – and there may be other more reasonable solutions. The radio astronomy calibration and imaging community needs to be aware of this characteristic of solutions to the Jones matrices.	Good point. **Required Action**: to be communicated as in this note.
19.	Tony Willis	The radio astronomy community needs to adopt the "correct" definition of Stokes parameters.	Subsequent discussion identified that some radio astronomy papers had complied with URSI and IEEE definitions; but that such compliance was not yet universal. **Required Action**: Not a factor impacting on calibration and imaging to achieve high DR; but nevertheless should be communicated in this note.

	Speaker	Description	Comments and Actions to be Taken?
20.	Maxim Voronkov	"Factor of 2" discussion with respect to polarisation measurements and analysis.	**Required Action**: Not a factor impacting on calibration and imaging to achieve high DR.
21.	Tobia Carozzi	Follow on from his CALIM 2010 presentation; Duncan Hal asked what early results might exist to quantify the error between "exact" Maxwell's equations formulation and current practices.	No results yet available.
22.	Maxim Voronkov	How many science cases really require 65 dB DR?	Good point. **Required Action**: to be borne in mind when science prioritisation occurs.
23.	Steve Torchinsky	Didn't RFI requirements drive the 65 dB DR specification?	**Required Action**: Duncan Hall to investigate.
24.	Stefan Wijnholds	We should not concentrate on DR as the target measure; rather we should focus on what the noise in the final image should be contained to as a target. Stefan offered to circulate a paper on this topic.	Good point. **Required Action**: to be communicated as in this note. Stefan's paper is available at: http://arxiv.org/pdf/1003.2307
25.	Johan Hamaker	Side lobes well outside the primary beam need to be modelled to take into account artefacts from bright sources. We need to know how well this works for e.g. LOFAR and other aperture arrays.	Good point. **Required Action**: to be communicated as in this note. Also see 15 above.
26.	Duncan Hall	Could we work backwards from how much data we are willing to store or archive to arrive at what effort should be put into computation in a pipeline mode of operation?	Chris Broekema commented that only as small subset of e.g. Westerbork Synthesis Radio Telescope (WRST) archived data is ever accessed. Stefan Wijnholds commented that as a survey instrument, SKA Phase 1 should aim to store roughly 2 years of data. Gerry Harp noted that some form of data compression would be desirable, however others noted that noise-like data cannot be efficiently compressed.

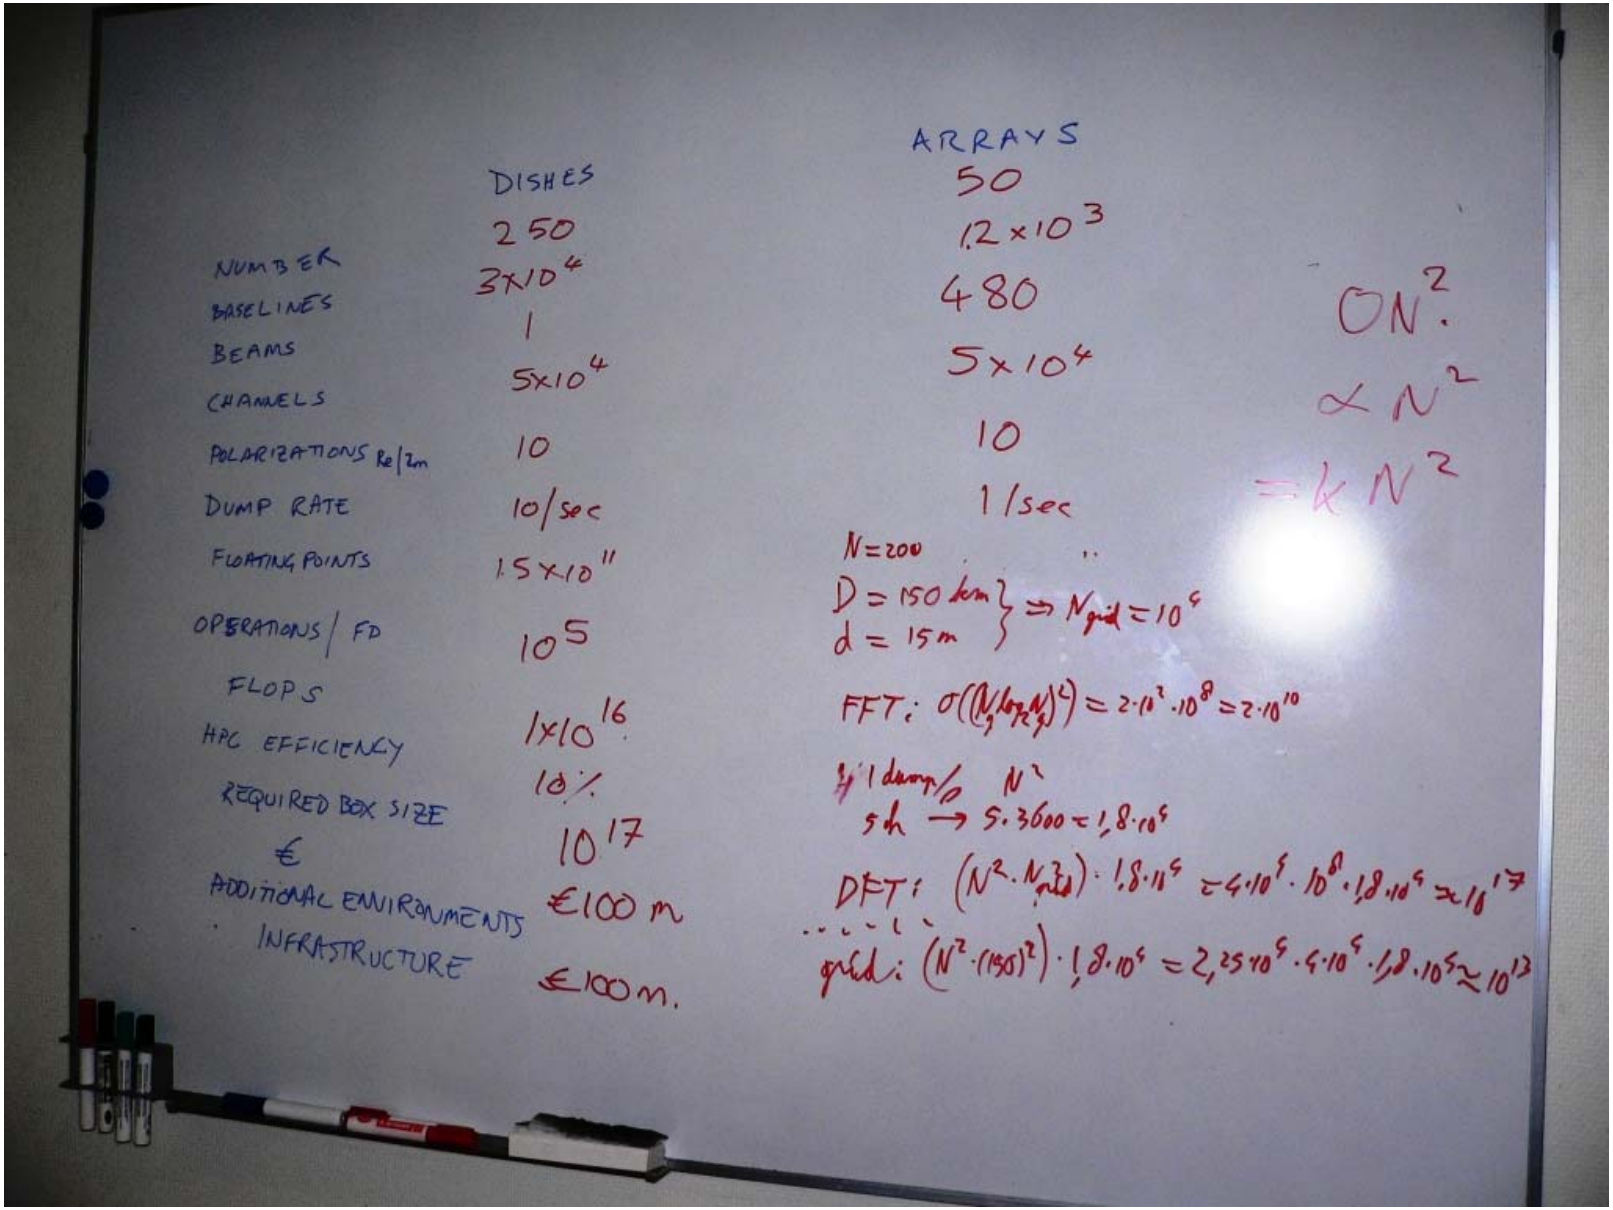
DISHES
ARRAYS
NUMBER
250
50
BASELINES
3x10^4
1.2x10^3
BEAMS
1
480
CHANELS
5x10^4
5x10^4
POLARIZATIONS Re/Im
10
10
DUMP RATE
10/sec
1/sec
FLOATING POINTS
1.5x10^11
OPERATIONS / FD
10^5
FLOPS
1x10^16
HPC EFFICIENCY
10%
REQUIRED BOX SIZE
10^17
€
ADDITIONAL ENVIRONMENTS
€100 m
INFRASTRUCTURE
€100 m.
ON^2
∝ N^2
= k N^2
N=200
D = 150 km
d = 15 m
⇒ N_grid = 10^6
FFT: O((N_g log N_g)^2) = 2·10^2 · 10^8 = 2·10^10
1 dump/s
5h → 5·3600 = 1,8·10^4
DFT: (N^2 · N_grid^2) · 1,8·10^4 = 4·10^5 · 10^8 · 1,8·10^4 ≈ 10^17
grid: (N^2 · (150)^2) · 1,8·10^4 = 2,25·10^5 · 4·10^5 · 1,8·10^4 ≈ 10^13

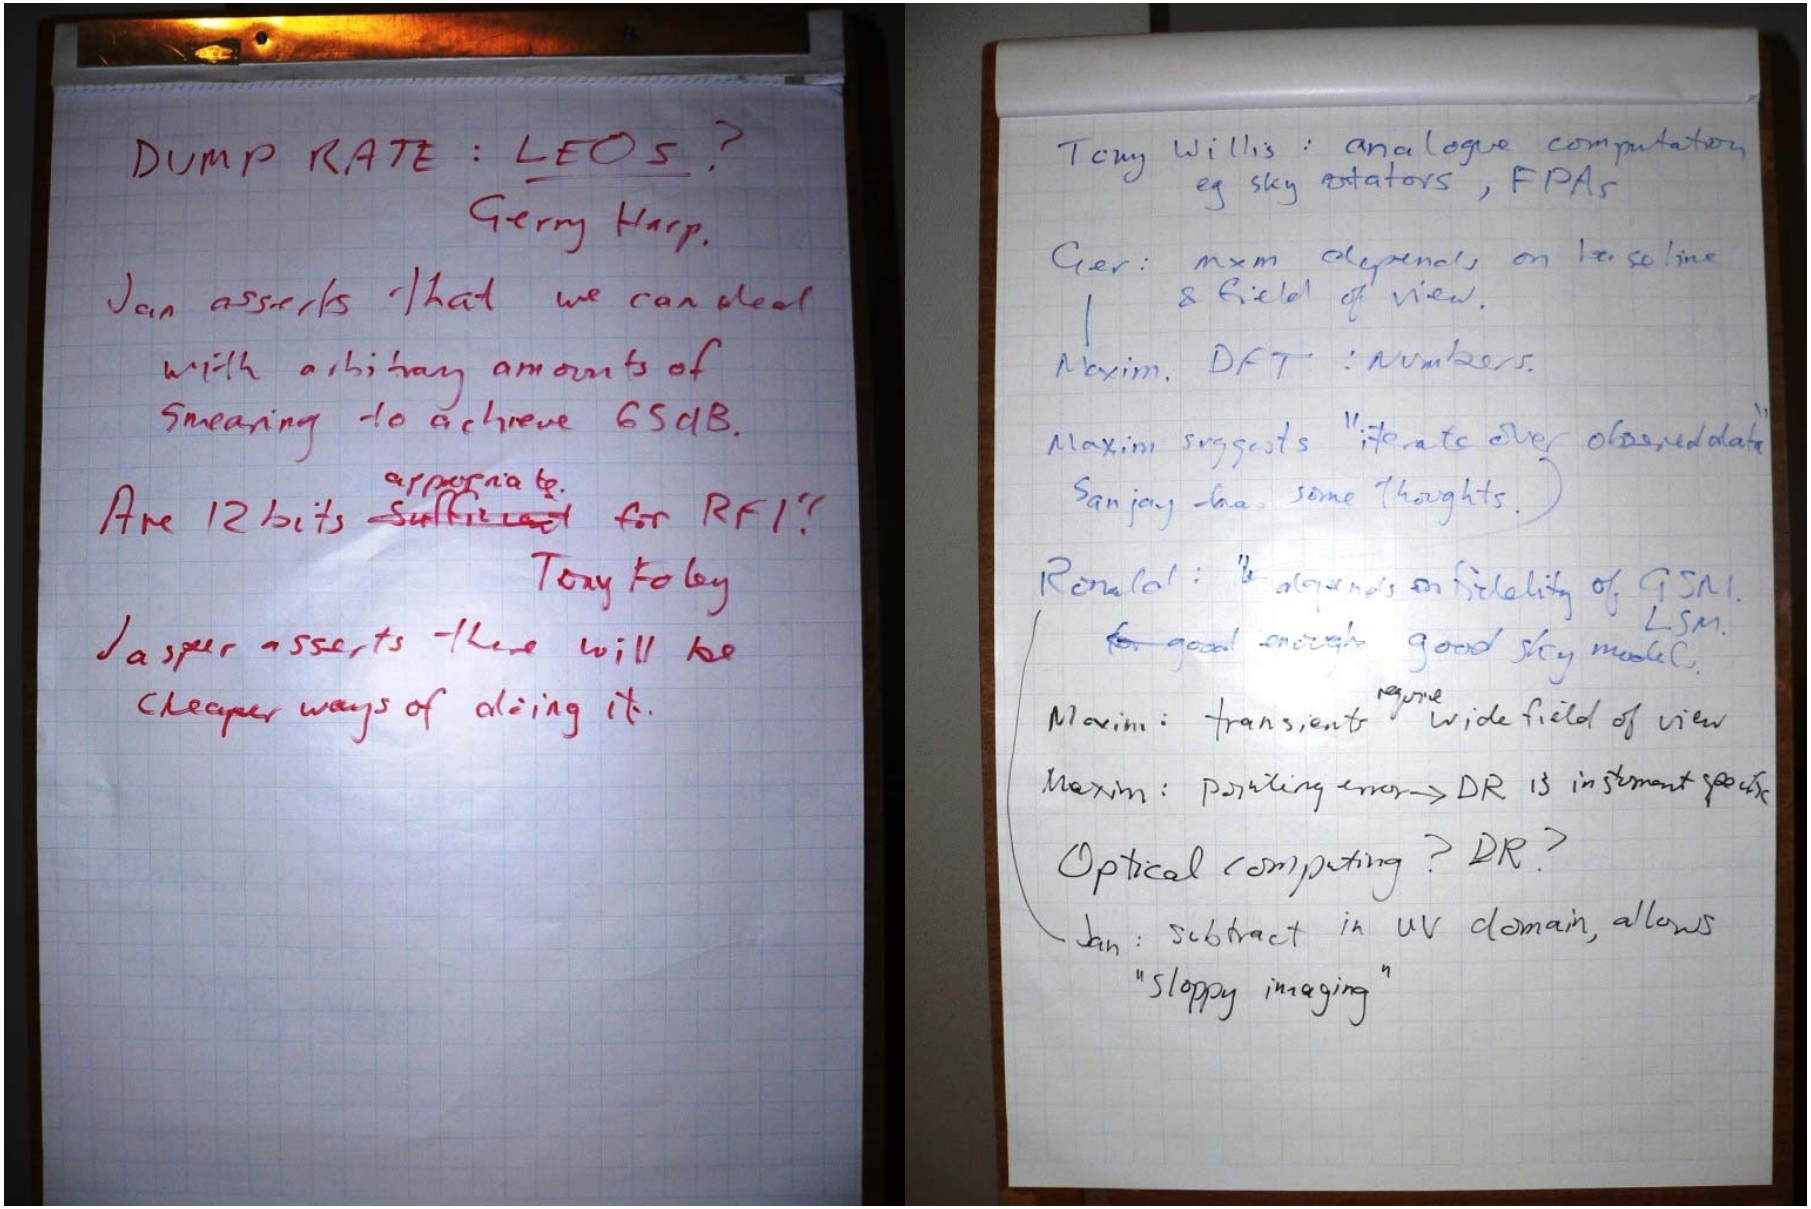
DUMP RATE : LEOS ?
Gerry Harp.
Jan asserts that we can deal with arbitrary amounts of smearing to achieve 65dB.
Are 12 bits appropriate for RFI?
Tony Foley
Jasper asserts there will be cheaper ways of doing it.
Tony Willis : analogue computation eg sky rotators, FPAs
Ger : nxm depends on baseline & field of view.
Maxim. DFT : Numbers.
Maxim suggests "iterate over observed data"
Sanjay has some thoughts.
Ronald : depends on fidelity of GSM. LSM.
good enough good sky models.
Maxim : transients require wide field of view
Maxim : pointing error → DR is instrument specific
Optical computing ? DR ?
Jan : subtract in UV domain, allows "sloppy imaging"

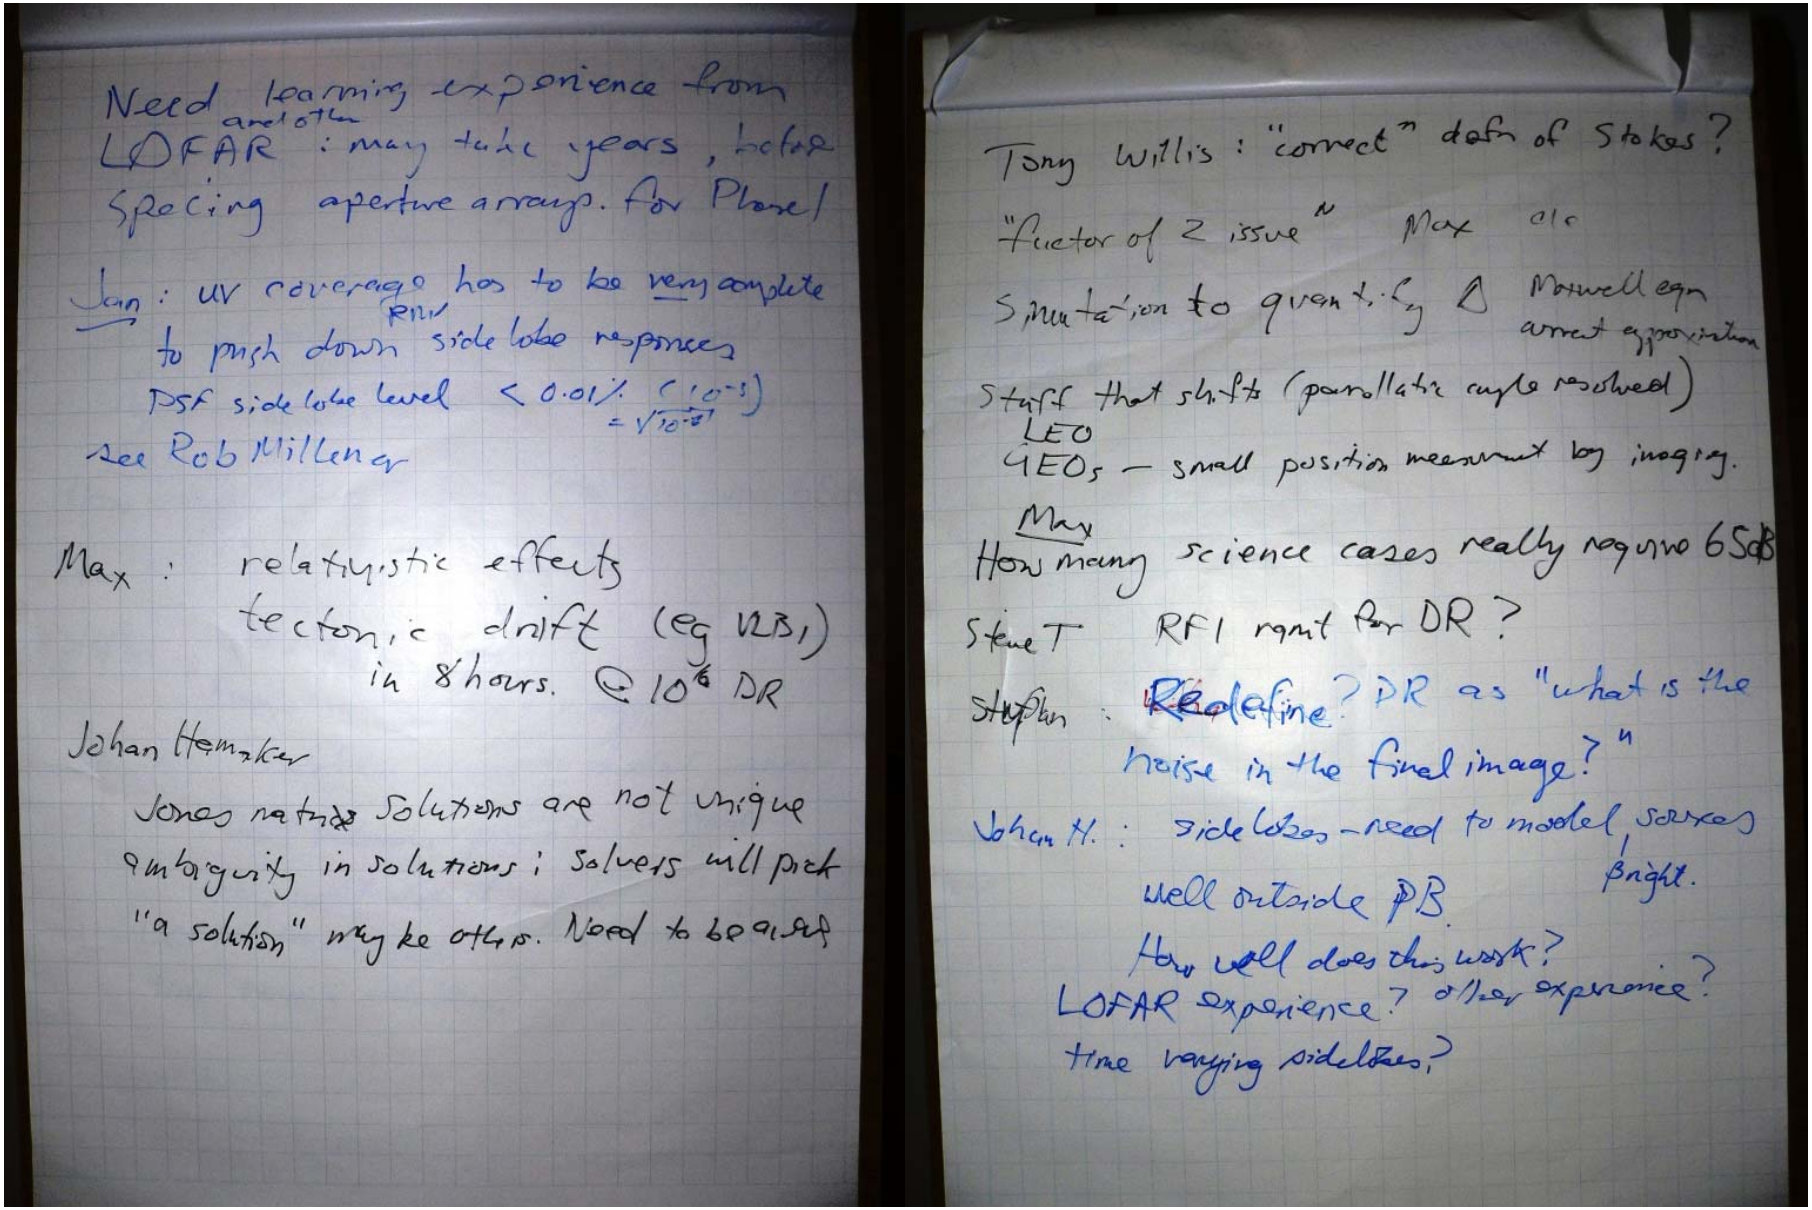

Need learning experience from LOFAR and others : may take years, before Spacing aperture arrays. for Phase1

Jan : uv coverage has to be very complete to push down RMS side lobe responses
PSF side lobe level < 0.01% (10^{-3}) $= \sqrt{10^{-6}}$
see Rob Millenaar

Max : relativistic effects
tectonic drift (eg VLBI)
in 8 hours. @ 10^6 DR

Johan Hamaker
Jones matrix Solutions are not unique
ambiguity in solutions ; solvers will pick "a solution" may be other. Need to be aware

Tony Willis : "correct" defn of Stokes?
"factor of 2 issue" Max etc
Simulation to quantify Δ Maxwell eqn correct approximation
Stuff that shifts (parallactic angle resolved)
LEO
GEOs – small position measurement by imaging.

Max
How many science cases really require 65dB

Steve T RFI remit for DR?

Stephen : Redefine? DR as "what is the noise in the final image?"

Johan H. : side lobes – need to model sources well outside PB. bright.
How well does this work?
LOFAR experience? other experience?
time varying sidelobes?

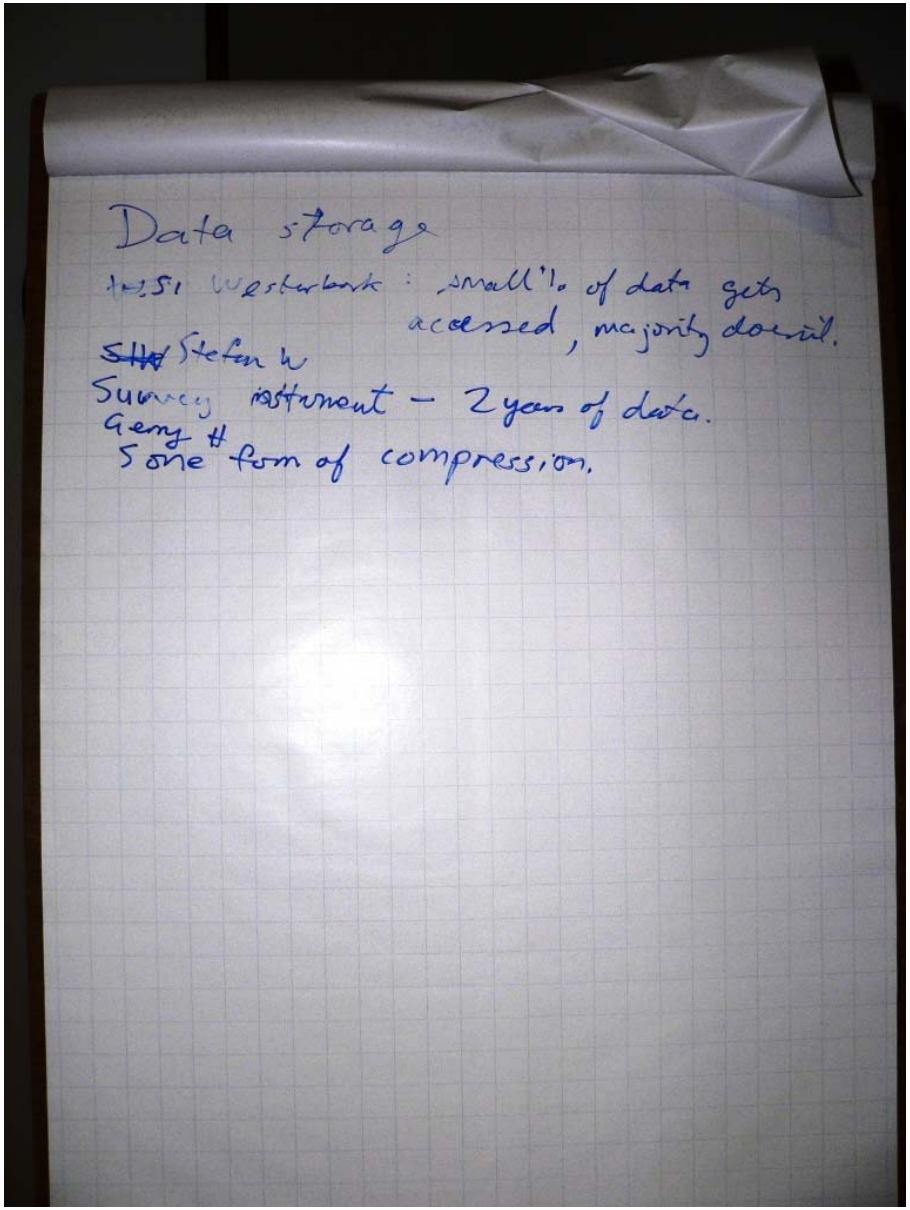
Data storage
1:51 Westerbork : small % of data gets accessed, majority down.
Stefan W
Survey instrument – 2 years of data.
Geny # 5 one form of compression.