SKA1 SCHEDULING AND ARCHIVE CONSTRAINTS

Document number SKA-TEL-SKO-0000951
Document Type .... REP
Revision 01
Author .... SKAO Science Team
Date 2019-03-04
Document Classification .... UNRESTRICTED
Status .... Released

Name	Designation	Affiliation	Signature	
Authored by:				
SKAO Science Team p.p. Robert Braun	Science Director	SKAO		
			Date:	2019-03-05
Owned by:				
Robert Braun	Science Director	SKAO		
			Date:	2019-03-05
Approved by:				
Robert Braun	Science Director	SKAO		
			Date:	2019-03-05
Released by:				
Joseph McMullin	Programme Director	SKAO	J P McMullin	
			Date:	2019-03-06

DOCUMENT HISTORY

Revision	Date Of Issue	Engineering Change Number	Comments
A	2013-08-16	-	First draft release for internal review
01	2019-03-05		First Release

DOCUMENT SOFTWARE

	Package	Version	Filename
Word processor	MS Word	Word 2007	SKA-TEL-SKO-0000951-01_SKA1 Schedule Archive Constraints
Block diagrams			
Other			

ORGANISATION DETAILS

Name	SKA Organisation
Registered Address	Jodrell Bank Observatory Lower Withington Macclesfield Cheshire SK11 9DL United Kingdom Registered in England & Wales Company Number: 07881918
Fax.	+44 (0)161 306 9600
Website	www.skatelescope.org

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AIP....................Advanced Instrumentation Programme
AIPS ....Astronomical Image Processing System
ALMA................Atacama Large Millimetre Array
ATM..................Atmospheric Transmission of Microwaves
BD ....SKA1 Baseline Design
CASA................Common Astronomy Software Applications
CDR..................Critical Design Review
DD ....Direction Dependent (calibration)
DI......................Direction Independent (calibration)
DM....................Dispersion Measure
ECMWF............European Centre for Medium-range Weather Forecasts
EMI ....Electromagnetic Interference
EoR ....Epoch of Reionisation
FRB ....Fast Radio Burst
FWHM ....Full Width Half Maximum
GMRT...............Giant Metre-wave Radio Telescope
GPU..................Graphical Processing Unit
HPC..................High Performance Computing
HPSO ....High Priority Science Objective
ISW...................Integrated Sachs Wolfe effect
JVLA.................Jansky Very Large Array (or simply VLA)
LOFAR ....LOw Frequency ARray
NIP ....Non-Image Processing
PDR..................Preliminary Design Review
PFlops ....Peta Floating Point Operations per Second
PSF ....Point Spread Function
PWV ....Precipitable Water Vapour
RFI....................Radio Frequency Interference
RM....................Rotation Measure
RMS ....Root Mean Square
SED ....Spectral Energy Distribution
SEFD................System Equivalent Flux Density
SKA ....Square Kilometre Array
SKAO ....SKA Organisation
VLA...................Very Large Array
VLBI..................Very Long Baseline Interferometry

1 Introduction

1.1 Purpose of the document

The purpose of this document is to provide estimates of the archive needs associated with a range of use cases for the SKA1 telescopes.

1.2 Scope of the document

In this document, we consider the mix of experiments that can be accommodated with a specific high performance computing deployment scenario. We also consider possible constraints on the fraction of high spectral resolution data products imposed by the network connectivity of the two telescopes and hence the annual archive capacity.

2 References

2.1 Applicable documents

The following documents are applicable to the extent stated herein. In the event of conflict between the contents of the applicable documents and this document, **the applicable documents** shall take precedence.

[AD1] Applicable Document 1
[AD2] Applicable Document 2

2.2 Reference documents

The following documents are referenced in this document. In the event of conflict between the contents of the referenced documents and this document, **this document** shall take precedence.

[RD1] SKA-TEL-SDP-0000040, "Parametric Models of SDP Compute Requirements"
[RD2] Braun, R., 2013, A&A 551, A91, "Understanding synthesis imaging dynamic range"
[RD3] SKA-TEL-SKO-0000641, "SKA1 Error Budgets"
[RD4] SKA-TEL-SKO-0000818, "Anticipated SKA1 Science Performance"
[RD5] Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, AJ, 115, 1693
[RD6] SKA-TEL-SDP-0000038, "SDP System Sizing"
[RD7] SKA-TEL-SKO-0000007, "SKA1 Level 0 Science Requirements"
[RD8] Bonaldi, A, et al. 2018, MNRAS, "The Tiered Radio Extragalactic Continuum Simulation", https://arxiv.org/abs/1805.05222
[RD9] SKA-TEL-SKO-0000941, "Anticipated SKA1 HPC Requirements"

3 SKA1 Archival Needs

3.1 Computational Needs of SKA1 Use Cases

In a previous document [RD9] we have quantified the computational cost associated with data product generation for a wide range of use cases that might be considered for observation with the SKA1 telescopes. We reproduce one of the key outcomes of that study, the instantaneous HPC load as function of use case type, in Figure 1 below.

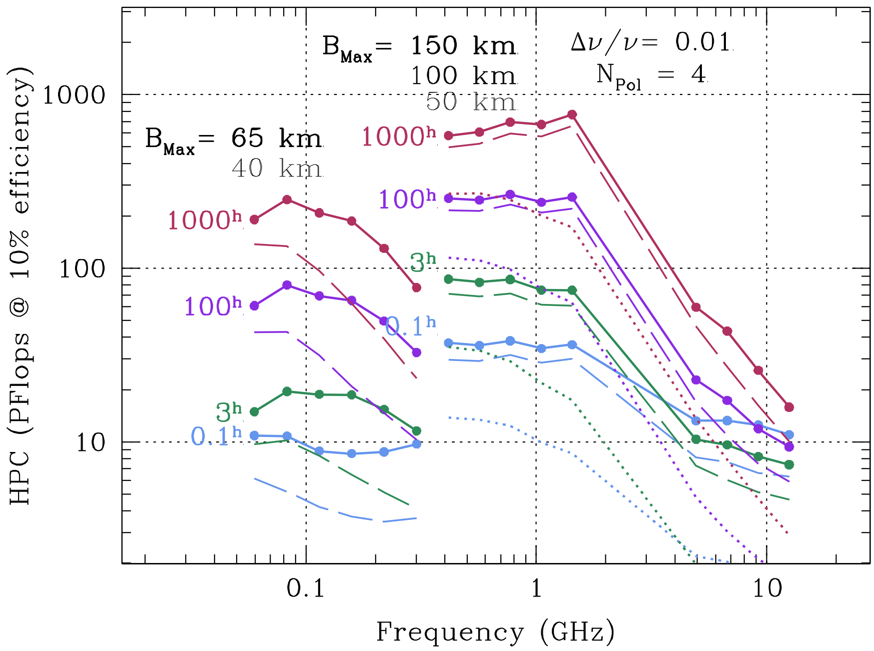

Figure 1. Computational needs for a range of SKA1 generic continuum experiment use cases. The filled circles and solid lines pertain to the largest values of B_{Max}. The four different coloured versions pertain to the four different values of the final depth per pointing direction, T_{Point}, of an experiment as labelled. The dashed and dotted curves of the same colour illustrate the consequence of reductions to B_{Max} to the values indicated.

Some of the most important conclusions of that study were that over the complete range of use cases (a) the computational cost is dominated by the calibration process, rather than by data product preparation and (b) that the main parameters that influence the computational cost are: the final depth of observation, T_{Point}, the maximum baseline, B_{Max}, and the observing frequency, ν_C. Perhaps surprisingly, there is no significant HPC cost difference in the provision of large multi-channel data cubes for spectral line use cases, compared to continuum data products that are more sparsely sampled in frequency. While the HPC cost is essentially equal for these two use case categories, the data product volume is of course not. The generic continuum use case that was considered specified an output data product at intervals of $(\Delta\nu/\nu) = 10^{-2}$ and polarisation number $N_{Pol} = 4$, while the generic spectral use case specified $(\Delta\nu/\nu) = 10^{-4}$ and $N_{Pol} = 2$ for its data products, amounting to a fifty-fold difference in data volume in each frequency sub-band that is observed. For the purpose of determining computational loads we will assume that $N_{Pol} = 4$ during calibration for all use cases. We also acknowledge that much higher spectral sampling $(\Delta\nu/\nu) < 10^{-5}$ will be required in some cases. Our generic "continuum" and "spectral" definitions are only intended to be indicative.

3.2 HPC-constrained Use Case Distributions

Some examples of generic use-case distributions, providing the fraction of observing time as a function of the use case parameters, with differing total HPC deployments were provided in Figures 7, 8 and 9 of RD9. In those examples, only “uniform” distributions, in which equal fractional observing time was allocated to all relevant use cases, were considered. With this approach, the total HPC deployment scope is the outcome of simply averaging over the full range of use cases with equal weight (ie. fractional observing time).

Ultimately, the problem must be addressed in the opposite sense. A fixed HPC deployment of some scope will be available and it will be desirable to schedule a mix of experiments that is consistent with data product preparation within this fixed scope. Use cases with instantaneous HPC needs that exceed the HPC capacity can still be supported if they are balanced against a suitable number of less demanding use cases within the visibility buffer window. When undertaken in an operational setting, the scheduling will be strongly influenced by the proposal priority that is recommended by a time allocation committee. In the event that the available HPC deployment is highly constrained relative to the typical HPC needs, then the computational cost may well influence the degree to which any use case can be scheduled in practise.

We will consider finding solutions for the fraction of observing time, w_i, that could be allocated to an imaging use case with instantaneous computing cost, F_i, given some assumed time allocation for non-imaging use cases, w_{NIP}, and their estimated computational cost, F_{NIP} = 1 PFlop (as discussed in RD9), subject to the constrained total HPC deployment, F_0,

$$F_0 = \sum_i w_i F_i + w_{NIP} F_{NIP}. \quad (1)$$

The other constraint that should be met is a filled telescope schedule,

$$1 = \sum_i w_i + w_{NIP}. \quad (2)$$

We will search for solutions of the form,

$$w_i = a_0/n + (1 - a_0) (F_0 + w_{NIP} F_{NIP}) F_i^{-\alpha}/n, \quad (3)$$

for a set of “n” use cases, where the fractional time, w_i, is the sum of a constant term and a power law term with index, $-\alpha$, that weights down the most computationally expensive use cases relative to those with lower HPC cost. For a discrete list of specific uses cases, each with associated cost, F_i, a simple parameter search over the two variables, a_0 and α, can be used to obtain the best fitting simultaneous solution to equations 1 and 2.

3.3 Deployment Baseline HPC Constraints

The initial HPC deployment proposed within the SKA1 Deployment Baseline is 50 PFlops, 24th SKA Board Meeting – July 2017
which is a substantial reduction from the 260 PFlops that was specified within the Design Baseline. This proposal took account of the relatively short refresh cycle for HPC that was foreseen within the Operations budget (five years for a complete refresh of 260 PFlops) as well as the likelihood of some ramp up in both observational and data product pipeline efficiency during the first years of operation.

ν_{min} (GHz)	ν_c (GHz)	ν_{max} (GHz)	B_{Max} (km)	T_{Point} (h)	F_i (PFlops)	w_i	V_{Si} (PByte/yr)	η_{Si}
0.05	0.09	0.13	40	0.1	15.5	0.18	152.0	0.82
0.13	0.23	0.35	40	0.1	10.8	0.24	202.7	0.82
0.05	0.09	0.13	40	3	28.2	0.12	3.4	0.86
0.13	0.23	0.35	40	3	15.7	0.18	5.0	0.84
0.05	0.09	0.13	40	100	117.1	0.07	1.5	0.92
0.13	0.23	0.35	40	100	45.9	0.10	2.0	0.90

Table 1. Use-case parameters (ν_{Min}, ν_{Max}, B_{Max}, T_{Point}) considered for SKA1-Low and their associated HPC cost (at 10% computational efficiency), F_i. The schedule fractions, w_i, derived for a 25 PFlop deployment are also listed together with the associated spectral line data product rate, V_{Si}, and maximum allowed spectral line fraction, η_{Si}, for a 300 PByte/year archive capacity.

Comparison with Section 3.6 of RD9, in which use case distributions were shown for an 85 PFlop deployment, demonstrates that a 50 PFlop scenario will be quite highly HPC constrained. We will consider a similar use case mix to that analysed there. As previously, we will assume a scheduling fraction for non-imaging use cases, w_{NIP} = 0.1. This is completely arbitrary and does not in any way reflect a perceived priority. It is done simply to permit elucidation of trends that are being influenced by the mix of *imaging* use cases that have varying computational and data transport needs. Such trends would be diluted if a large w_{NIP} were assumed.

For SKA1-Low, we include two sets of frequency sub-bands (1+2+3) and (4+5+6) which span (50 – 132 MHz) and (132 – 350 MHz) respectively, keep B_{Max} fixed at 40 km and allow for three total depth classes, where T_{Point} = (0.1, 3, 100^h) are ultimately accumulated for each pointing direction (which will determine the quality of calibration that must be achieved). Individual tracking observations are limited to a T_{Obs} = 4^h duration so as to keep zenith angles suitably low. A total of six of such generic use cases span the parameter space, as shown in Table 1.

For SKA1-Mid we will need to consider each of the frequency sub-bands, with $(\Delta\nu/\nu) = 0.32$, individually. As shown in Table 3 of RD9, there are five of such sub-bands spanning 0.35 – 1.54 GHz and another four spanning 4.18 – 14.5 GHz. We will consider both B_{Max} of 50 km and 120 km and allow for the three total depth classes, where $T_{Point} = (0.1, 3, 100^h)$. In this case the maximum tracking length of an observation is taken to be $T_{Obs} = 6^h$. (Adjustments to the maximum assumed tracking length would make only minor changes to the $T_{Point} = 100^h$ case.) Here, there are a total of 54 generic use cases that span the parameter space as shown in Tables 2 and 3.

ν_{min} (GHz)	ν_c (GHz)	ν_{max} (GHz)	B_{Max} (km)	T_{Point} (h)	F_i (PFlops)	w_i	V_i (PByte/yr)	η_{Si}
0.35	0.42	0.48	50	0.1	13.8	0.010	27.3	0.15
0.48	0.57	0.65	50	0.1	13.4	0.010	27.3	0.15
0.65	0.77	0.89	50	0.1	12.3	0.010	27.3	0.15
0.89	1.05	1.21	50	0.1	9.9	0.010	27.4	0.15
1.21	1.43	1.65	50	0.1	8.6	0.010	27.5	0.15
4.18	4.94	5.70	50	0.1	2.2	0.021	59.1	0.15
5.70	6.74	7.78	50	0.1	2.0	0.025	69.7	0.15
7.78	9.20	10.60	50	0.1	1.6	0.047	129.7	0.15
10.60	12.54	14.46	50	0.1	1.4	0.068	189.2	0.15
0.35	0.42	0.48	50	3	35.2	0.010	0.9	0.73
0.48	0.57	0.65	50	3	33.5	0.010	0.9	0.73
0.65	0.77	0.89	50	3	29.1	0.010	0.9	0.73
0.89	1.05	1.21	50	3	21.9	0.010	0.9	0.73
1.21	1.43	1.65	50	3	17.3	0.010	0.9	0.73
4.18	4.94	5.70	50	3	2.0	0.027	2.5	0.45
5.70	6.74	7.78	50	3	1.6	0.046	4.3	0.29
7.78	9.20	10.60	50	3	1.4	0.071	6.6	0.21
10.60	12.54	14.46	50	3	1.3	0.092	8.6	0.17
0.35	0.42	0.48	50	100	114.8	0.010	0.5	0.85
0.48	0.57	0.65	50	100	111.0	0.010	0.5	0.85
0.65	0.77	0.89	50	100	97.9	0.010	0.5	0.85
0.89	1.05	1.21	50	100	76.3	0.010	0.5	0.85
1.21	1.43	1.65	50	100	63.3	0.010	0.5	0.85
4.18	4.94	5.70	50	100	4.8	0.010	0.5	0.84
5.70	6.74	7.78	50	100	3.1	0.013	0.6	0.81
7.78	9.20	10.60	50	100	2.1	0.023	1.1	0.69
10.60	12.54	14.46	50	100	1.7	0.042	1.9	0.53

Table 2. Use-case parameters (ν_{Min}, ν_{Max}, B_{Max}, T_{Point}) considered for SKA1-Mid (part one) and their associated HPC cost (at 10% computational efficiency), F_i. The schedule fractions, w_i, derived for a 25 PFlop deployment are also listed together with the associated spectral line data product rate, V_{Si}, and maximum allowed spectral line fraction, η_{Si}, for a 300 PByte/year archive capacity.

ν_{min} (GHz)	ν_c (GHz)	ν_{max} (GHz)	B_{Max} (km)	T_{Point} (h)	F_i (PFlops)	w_i	V_i (PByte/yr)	η_{Si}
0.35	0.42	0.48	120	0.1	33.0	0.010	157.1	0.15
0.48	0.57	0.65	120	0.1	32.0	0.010	157.1	0.15
0.65	0.77	0.89	120	0.1	34.6	0.010	157.1	0.15
0.89	1.05	1.21	120	0.1	31.4	0.010	157.1	0.15
1.21	1.43	1.65	120	0.1	33.3	0.010	157.1	0.15
4.18	4.94	5.70	120	0.1	10.4	0.010	157.6	0.15
5.70	6.74	7.78	120	0.1	10.0	0.010	157.7	0.15
7.78	9.20	10.60	120	0.1	8.8	0.010	158.1	0.15
10.60	12.54	14.46	120	0.1	8.3	0.010	158.4	0.15
0.35	0.42	0.48	120	3	77.4	0.010	5.2	0.25
0.48	0.57	0.65	120	3	74.4	0.010	5.2	0.25
0.65	0.77	0.89	120	3	77.3	0.010	5.2	0.25
0.89	1.05	1.21	120	3	67.1	0.010	5.2	0.25
1.21	1.43	1.65	120	3	66.3	0.010	5.2	0.25
4.18	4.94	5.70	120	3	8.8	0.010	5.3	0.24
5.70	6.74	7.78	120	3	8.1	0.010	5.3	0.24
7.78	9.20	10.60	120	3	6.4	0.010	5.4	0.24
10.60	12.54	14.46	120	3	5.8	0.010	5.4	0.24
0.35	0.42	0.48	120	100	229.2	0.010	2.6	0.44
0.48	0.57	0.65	120	100	226.4	0.010	2.6	0.44
0.65	0.77	0.89	120	100	244.4	0.010	2.6	0.44
0.89	1.05	1.21	120	100	221.2	0.010	2.6	0.44
1.21	1.43	1.65	120	100	234.1	0.010	2.6	0.44
4.18	4.94	5.70	120	100	20.2	0.010	2.6	0.44
5.70	6.74	7.78	120	100	13.3	0.010	2.6	0.44
7.78	9.20	10.60	120	100	9.2	0.010	2.6	0.43
10.60	12.54	14.46	120	100	7.3	0.010	2.7	0.43

Table 3. Use-case parameters (ν_{Min}, ν_{Max}, B_{Max}, T_{Point}) considered for SKA1-Mid (part two) and their associated HPC cost (at 10% computational efficiency), F_i. The schedule fractions, w_i, derived for a 25 PFlop deployment are also listed together with the associated spectral line data product rate, V_{Si}, and maximum allowed spectral line fraction, η_{Si}, for a 300 PByte/year archive capacity.

Solving for the values of a_0 and α that satisfy equations 1 and 2 for these use case mixes, each constrained by F_0 = 25 PFlops, yields the scheduling fractions shown in Figure 2, and tabulated in Tables 1, 2 and 3. Smaller symbols are used for the smaller values of B_{Max} and three different colours (as labelled) are used for the three depth classes. The solid curve is the solution for equation 3. The use case distributions are also shown as function of observing frequency in Figure 3. The frequency coverage of each case is indicated with the horizontal error bars.

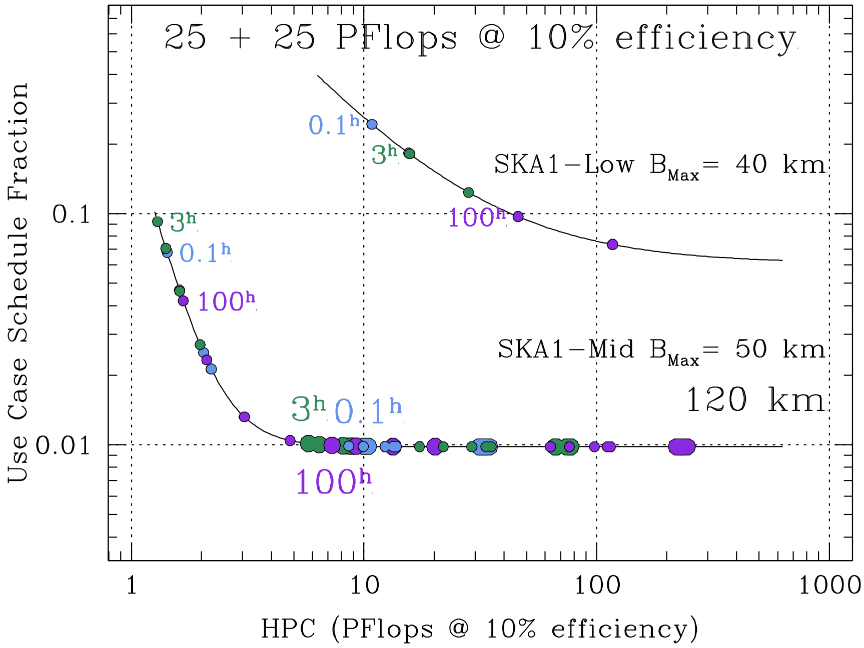

Figure 2. Use case scheduling fraction as function of computational load for the case of a 25 PFlop (at 10% assumed efficiency) deployment for each of SKA1-Low and SKA1-Mid.

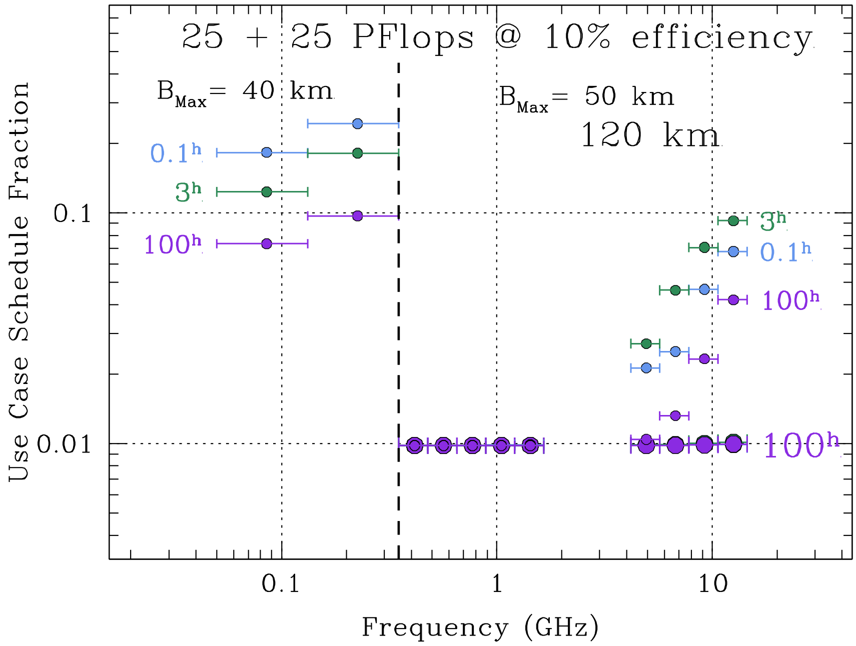

Figure 3. Use case scheduling fraction as function of observing frequency for the case of a 25 PFlop (at 10% assumed efficiency) deployment for each of SKA1-Low and SKA1-Mid.

3.4 Deployment Baseline Archive Constraints

As noted in Section 3.1, while there may be no significant difference in HPC cost for continuum and spectral line use cases, there is a substantial difference in the associated data product volume. It is anticipated that the long-term archive of SKA data products will be hosted within a network of SKA Regional Centres and that the initial data transport limit will be constrained to a single 100 Gbit/s link for each of SKA1-Low and SKA1-Mid. This translates into an annual limit on data product transport and storage of about 300 PByte/year (assuming a 75% average link utilisation) for each telescope.

The data product volume per year, V_i, for each imaging use case can be determined from its specifications,

$$V_i = N_{py}\, N_{Sub}\, N_{Chan}\, N_{Pol}\, b_{pix}\, N_{pix}^2, \tag{4}$$

where $N_{py} = h_{py} / T_{obs}$, is the number of observations that would fill a year, the number of hours in a year, h_{py}, and the length of a single telescope track for that observation, T_{obs}, (which will be less than T_{Point} for the deepest integrations). Further, N_{Sub} is the number of sub-bands, N_{Chan} the number of spectral channels, N_{Pol} the number of polarisations, $b_{pix} = 4$, is the number of bytes per pixel and N_{pix} the number of pixels per spectral channel. The number of image plane pixels is given by [eg. RD1],

$$N_{pix} = 10\, B_{Max}\, Q_{pix}\, Q_{FoV}\, (\nu_{Max}/\nu_{Min}) / (\pi\, D_S), \tag{5}$$

in terms of the maximum baseline, B_{Max}, the synthesised beam oversampling factor, $Q_{pix} = 2.5$, the fractional Field of View, $Q_{FoV} = 1$, the ratio of maximum to minimum frequency per sub-band, $(\nu_{Max}/\nu_{Min}) = 1.37$, and the dish or station diameter, D_S.

The annual data volume is given by,

$$D_0 = \textstyle\sum_i w_i\, V_i, \tag{6}$$

where the scheduling fractions, w_i, allowed by the available HPC deployment have been determined in Section 3.3 above. The data volume for each use case category can be further broken down into its two primary components,

$$V_i = \eta_{Ci}\, V_{Ci} + \eta_{Si}\, V_{Si}, \tag{7}$$

where V_{Ci} applies to continuum use cases (for which we take $N_{Chan} = 30$ and $N_{Pol} = 4$ per sub-band) and V_{Si} to spectral line use cases (for which $N_{Chan} = 3000$ and $N_{Pol} = 2$ per sub-band). The relative fraction of spectral, η_{Si}, and continuum, η_{Ci}, use cases, subject to,

$$1 = \eta_{Ci} + \eta_{Si}, \tag{8}$$

will determine the actual data volume. In cases where the total annual data volume would otherwise exceed the anticipated 300 PByte/year limit for each telescope it might be necessary to introduce an upper limit to the fraction of spectral experiments, $\eta_{Si} < 1$. In this case, we will search for solutions of the form,

$$\eta_{Si} = [b_0 + (1 - b_0) \exp(-V_{Si}/V_e)], \qquad (9)$$

where the maximum allowed spectral-line fraction is composed of a minimum value, b_0, and one that scales as an exponential of the spectral data volume, V_{Si}, normalised to V_e. The best fitting values of b_0 and V_e were found, in a discrete parameter search, to satisfy equations 6, 7 and 8. The resulting constraints on the maximum line fraction as function of the spectral line data rate are shown in Figure 4 and tabulated in Tables 1, 2 and 3. The solid curves in the figure are the solutions for b_0 and V_e substituted into equation 9. Observations with SKA1-Low are only very mildly constrained in their allowed spectral fraction, while the observations with SKA1-Mid, particularly those with short T_{Point} and large B_{Max} are strongly constrained.

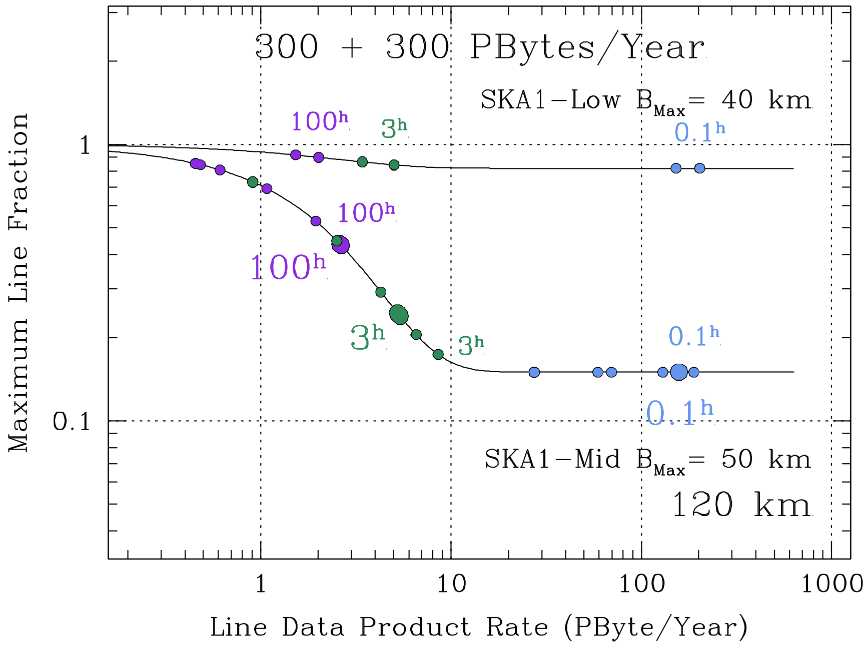

Figure 4. The maximum allowed spectral line fraction as function of data product volume for the case of a 300 PByte/year constraint on archive accumulation for each of SKA1-Low and SKA1-Mid.

The maximum line fractions are shown as function of the observing frequency in Figure 5. The frequency coverage of each case is indicated with the horizontal error bars.

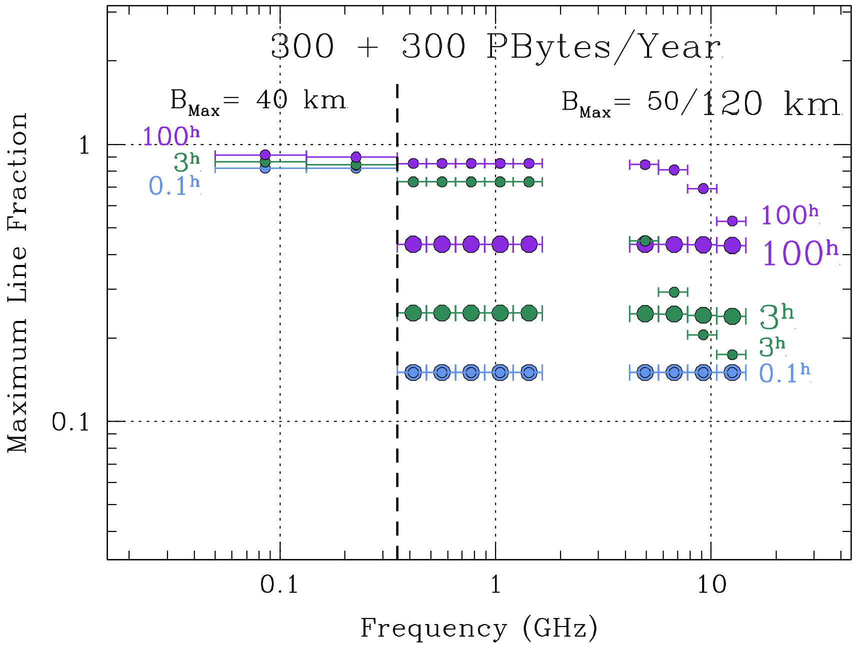

Figure 5. The maximum allowed spectral line fraction as function of observing frequency for the case of a 300 PByte/year constraint on archive accumulation for each of SKA1-Low and SKA1-Mid.

4 Conclusions

4.1 Implications of a Highly Constrained HPC Deployment

It has been demonstrated in Section 3.3 that a moderately wide range of use cases can be supported by even the 50 PFlops HPC system (operating at 10% computational efficiency) that is specified within the Deployment Baseline. Load averaging should permit even computationally demanding experiments to be undertaken, calibrated and imaged for some fraction of the time. Relative to a completely unconstrained scenario, it will typically be necessary to accept some limitations on the simultaneously processed bandwidth. The generic use cases that we have considered have in most cases been restricted to only about one half of the total bandwidth provided by the correlator; only three out of six sub-bands in the case of SKA1-Low and only one out of the (typically) two sub-bands per single pixel feed available on SKA1-Mid. Within this constraint, it is possible to provide use case schedule fractions on both telescopes that are within about a factor of two of the "uniformly" scheduled rate (1/6 for SKA1-Low) and (1/54 for SKA1-Mid) that is simply the reciprocal of the number of relevant use cases. The other short term restriction on use cases that is adopted here is the deferment of the most HPC challenging use case category (the 1000^h integrations) to a later date. Both of these limitations would be eliminated with either or both of a larger HPC deployment or achievement of a significantly higher computational efficiency (as discussed in RD9).

The anticipated data transport limitation imposed by a single 100 Gbit/s link (transmitting about 300 PBytes per year) for each telescope also introduces significant constraints on the use case mix that can be supported, as shown in Section 3.4. Only about 15% of the shortest observations undertaken with SKA1-Mid could provide archived data products at high spectral resolution. Long observations with SKA1-Mid and all observations with SKA1-Low are only mildly constrained in the spectral resolution of archived data products that can be accommodated. Achieving a more natural match to the telescope capabilities would require about a 50% increase in network bandwidth and archive capacity in the case of SKA1-Low (to account for larger processed bandwidths) and a 5 – 10-fold increase for SKA1-Mid (for both larger bandwidth and reduction of the current spectral constraints).