

Response
of the SKA Pulsar Science Working Group
to the SKA₁ Baseline Design

Version 2.6 – 20 February 2014
Nearly Final Draft – Appendix & References to be completed.

Edited by PSWG Chairs M. Kramer & B. Stappers

Executive Summary

We summarise here the major conclusions from the SKA Science Assessment Workshop (SAW) of the Pulsar Science Working Group (PSWG) that took place at the SPO headquarters at Jodrell Bank on 17th and 18th July 2013. It also includes subsequent analysis which resulted from questions raised during the SAW. It constitutes an assessment of the baseline design and subsequently the Level 1 requirements up to revision 1.

This report encapsulates the joint response of the PSWG and is based on input by a broad representation of the pulsar community that goes beyond the membership of the PSWG. It is obtained from face-to-face discussions at the *IPTA2013* meeting and the Science Assessment Workshop at Jodrell Bank, as well as the extensive discussions via email. Note that contributors and their contact details are listed in the Appendix.

The major conclusions detailed in the follow pages are:

- The PSWG confirmed the current science case for the SKA. While the field has developed enormously in the last 10 years since the SKA Science Book was written, the original vision is still strong, and the science drivers have even been reinforced by the discoveries and progress of the recent past. Consequently, the telescope specifications required are basically the same, but, not necessarily the same as is in the baseline design as yet.
- We find that the distribution of the dishes specified in the baseline design for SKA₁-MID provides a reasonable compromise between our required sensitivity for the pulsar search and the distribution of dishes required for imaging. We recognised the request to modify the dish configuration to allow for improved imaging performance and accept the loss of up to 10% of sensitivity in the core region. We note however that any further losses of sensitivity will result in either: reduced survey speed or a requirement for more beams. Moreover, and most importantly, compensating for sensitivity loss of about 10% with increased observing time results in a doubling of the processing requirements to find the *same* relativistic binary pulsars.
- The PSWG conclude that SKA₁-LOW is an excellent instrument for undertaking pulsar surveys away from the Galactic plane. Our simulation work shows that with a modest number of beams, around 500, which could be traded for beams used with SKA₁-MID, an efficient survey could be undertaken with SKA₁-LOW. We show that this survey will find more pulsars off the Galactic plane than can be found with SKA₁-MID due to the exceptional collecting area available. Our simulations show that a combined SKA₁-MID and SKA₁-LOW survey targeting regions on and

off the Galactic plane respectively finds approximately 20% more pulsars than can be found with SKA₁-MID alone.

- The PSWG also conclude that enabling pulsar timing with SKA₁-LOW would also greatly facilitate the high precision timing and normal pulsar timing projects. The low-frequency capabilities of SKA₁-LOW will be hugely beneficial to the determination of variations in the interstellar medium / dispersion measure which can be the limiting factor to timing precision. The extreme wide-field of view combined with the sub-arraying properties of SKA₁-LOW will enable larger numbers of pulsars to be timed simultaneously addressing the major issue of having sufficient observing time to characterise the large number of pulsars that will be found with SKA₁ .
- For a subset of pulsars, the highest timing precision will be required, implying the usage of the full array. Note that the necessity to time-resolve the timing signatures of compact, fast binaries on timescales as short as 10 to 30 seconds, a loss of sensitivity in pulsar timing will translate directly into a loss science possible with the telescope.
- The PSWG conclude that it should be possible to survey for pulsars in any of the bands of SKA₁-MID and not just the two bands that were mentioned in the first revision of the level one requirements. This ability will enable surveys with SKA₁-MID , both wide area and targeted, to be undertaken at the optimal frequency as determined by the degree of dispersion and scattering along the line of sight. In this context, we emphasize the need for Band 5 for searches in the inner Galaxy and the Galactic Centre
- The PSWG is highly supportive of the expansion of the number of high precision timing beams available in SKA₁-MID to sixteen. As detailed above we would like to have the same number in SKA₁-LOW .
- The PSWG strongly supports sub arraying with both SKA₁-LOW and SKA₁-MID .
- The PSWG strongly supports a VLBI mode being enabled for SKA₁-MID . This would allow for astrometric programs to be undertaken using the SKA₁-MID in combination with other telescopes around the world this will enable precision astrometry of pulsars to be undertaken. Amongst other things precision differential astrometry of pulsars allows astrometric parameters to be accurately measured and subsequently fixed in the pulsar timing model, facilitating improved strong-field tests of gravity.
- The PSWG has also assessed the possibility to discovery pulsar via imaging observation. While this appears as an attractive additional method, our studies conclude that imaging searches will not be able to replace the non-imaging processing prospects and requirements.

1 Introduction

The SKA will allow us to achieve transformational scientific results in our understanding of gravitational physics by discovering and observing pulsars using its high sensitivity, wide field of view, and frequency coverage. With it, we will also explore the variable radio sky in general and a full census of radio pulsars in particular. This will lead to the discovery of previously unknown types of sources and to probe a wide range explosive and dynamic events. Among the pulsars will be fast, spin-stable millisecond pulsars whose period distribution reflects the equation of state of nuclear matter and some of which will serve

as detectors of nano-Hz gravitational waves. Relativistic binary pulsars, particularly those with orbital periods of a few hours or less, will allow strong field tests of General Relativity. The first step toward these unique achievements will be enabled by Phase I of the SKA (SKA₁). Apart from being a transformational telescope in its own right, the first science phase of this unique telescope will set the scene for the experiments to be conducted with the full array. With the newly discovered pulsars and unprecedented sensitivity and timing precision for the already known pulsars, studies of key pulsars will achieve some of the the science goals identified for the 'Pulsar and Gravity' SKA Key Science Project (KSP). Significant advances in our physical understanding of the dynamical gravitational universe will go along with a much enhanced understanding of the Milky Way, its structure and its constituents. This knowledge will be necessary and will allow us to establish the needed techniques and source samples to prepare for full SKA operations and the full galactic census of pulsars.

The requirements for achieving the science possible with the SKA are the sensitive search of the dynamic radio sky and the follow-up timing observations that will eventually extract the science itself. The latter can be significantly aided, if VLBI observations can be made to establish precise astrometric parameters. In order to maximize the yield to allow us to find the rarest objects and to provide the best possible timing precision, certain telescope specifications are demanded. The requirements for VLBI are yet different again. In the following we compare these requirements with the currently established SKA₁ baseline design.

We note that for searching, SKA₁ is not only a simple stepping stone towards Phase 2. Due to limitations in processing power, it is unlikely that the full area of the completed SKA can be utilized for a blind, large-scale survey for some time to come. Hence, SKA₁ with a highly concentrated core represents a significant fraction of collecting area usable for surveys with the full SKA, and significant achievements can be made in pulsar searching in the early science phase. In contrast, in order to achieve the highest precision in pulsar timing, the gain in using the full SKA is enormous when compared to timing observations enabled with Phase I. It is this gain in timing precision for selected KSP objects which will ultimately require the sensitivity of the full SKA.

We point out again that in general loss in sensitivity (i.e. collecting area) cannot be simply compensated by longer integration times. In the best case, a reduction in collecting area that can be phased up coherently may require a significant increase in computing power (both for beam-forming and processing). These costs are typically prohibitive. In the worst case, a loss in sensitivity means a degradation in science possible. The situations are slightly different for timing and searching, and we discuss both appropriately.

2 Pulsar Searching

The pulsar signal is periodic with known periods covering nearly four orders of magnitude (i.e. 1ms to 10s). The pulse duty cycle ranges, typically, from about one degree in rotational phase (though much shorter pulses are observed) to the full pulse period. Hence, in practice, the discovery of pulsars requires a search of the radio sky sampled with high time and frequency resolution in dispersion measure, pulse period and pulse duty cycle. For binary pulsars, a search for acceleration is required additionally. The standard search technique is consequently computational expensive and it involves the Fourier-transform of correspondingly prepared time-series where the search in duty cycle is using a technique known as *harmonic summing*. Apart from these data processing requirements, forming beams for the entire FoV is computationally challenging in the signal processing chain, in particular if the array configuration is not sufficiently compact. In order to ease the computational requirements, we explore different

observing strategies for the search for pulsars, depending on sky (i.e. Galactic) location, frequency (and hence receiver technology), and observing modes. This includes standard pulsars searches and the possibility that pulsar candidates could also be identified in continuum images obtained with the SKA interferometer. The proven standard search techniques are studied in detail by using use-cases that will be presented in the Appendix.

2.1 Impact of loss in (core) sensitivity

The most interesting systems to be found will be highly accelerated. The SKA promises an advance over existing telescopes in the particular ability to allow for short integration times due to its large sensitivity. That means that only a small fraction of the orbit is sampled, so that the pulse frequency may not be much Doppler shifted. In contrast, reduced sensitivities imply longer integration times, which mean more Doppler-smearing that needs to be compensated (if possible at all) by computational means. In the best case, one can assume constant acceleration, and hence a search (in constant) acceleration. But for the most interesting systems, the assumption of constant acceleration is not applicable, so that a more computing-expensive acceleration search is required. In many cases, the required computational effort is prohibitive, meaning that it is impossible to detect corresponding systems. Even in the simplest acceleration search (i.e. constant acceleration as assumed above), the needed computing power scales with the cube of the observing time. In other words, a 10%-reduction in sensitivity implies a required increase in computing efforts by a factor of two in order to find the same systems. This relationship is presented in Fig. 1 as a best-case scenario. In practice, the penalty is much higher, for two reasons. Firstly, the previous assumption of constant acceleration will be wrong for the most compact systems when the integration time becomes a significant fraction of the orbital period (depending also on the unknown eccentricity of the system). Secondly, even if the antennae moved out to larger baseline can be included coherently in effective collecting area, it requires a increased amount of beam forming (see Fig. 2).

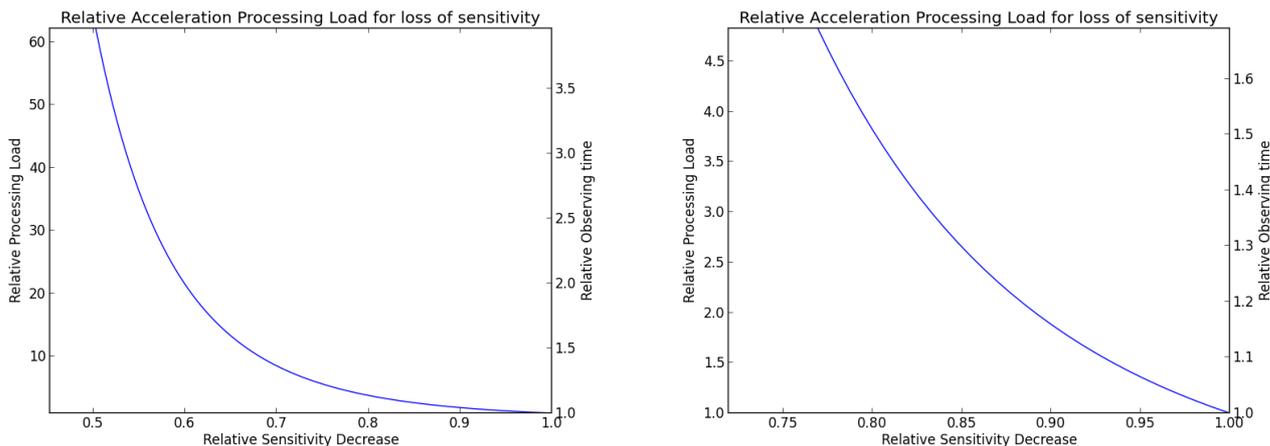


Figure 1: Relative acceleration processing and the implied increase for the simplest possible acceleration processing for loss of sensitivity. The right panel shows a zoomed-in version of the left panel.

If the computing power is not available, searches for pulsars can only be conducted with reduced

sensitivity. The impact on the so-lost population is difficult to quantify as the result depends not only on the (unknown) luminosity of the sources, but also on the orbital parameters. Population synthesis studies are not overall reliable, but we can be certain that there are more compact orbits than the Double Pulsar ($P_b = 147\text{min}$) or the shortest binaries known ($P_b = 92\text{min}$). It is notably that the Double Pulsar was not detected in the regular pointings of the Parkes Multi-beam Pulsar Survey (PMPS, $t_{\text{int}} = 35\text{min}$) because of Doppler-smearing. Only in the 10-times shorter integrations of the “PH”-survey, it was finally discovered. It is conceivable that most of the strong (little-accelerated) sources are discovered with the currently available sensitivity and computer resources. This would imply that the remaining sources are either relatively weak or very highly accelerated. In other words a loss in SKA sensitivity (the largest that we can ever expect to achieve) would suggest a certain (final) blindness to the most exciting systems.

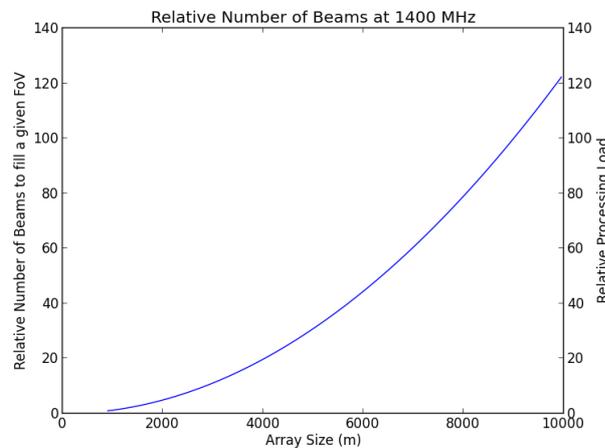


Figure 2: Increase of number of beams and the increase in processing requirements for an increase in array baselines.

Because of the difficulties in processing search data, it has been suggested that a search for candidates as point sources in imaging data may be a potential alternative. We consider this option in the next section.

2.2 Prospects of finding pulsars in imaging data – Karastergiou

As pointed out, traditional pulsar surveys, responsible for the discovery of the majority of all radio pulsars, involve intense processing of high time and frequency resolution time-series data of individual pointings on the sky. The cost of this type of processing is high, and only increased by a further requirement to perform it in real-time.

SKA₁-MID will be designed to perform an all-sky “traditional” survey for pulsars. An SKA₁-MID array will be used to conduct a 5000 sq. deg. survey to $0.35\mu\text{Jy}$ noise rms and sub-1” resolution, and a 30000 sq. deg. survey down to $1\mu\text{Jy}$ noise rms with 1-2” resolution (see Section ??.)

The question is, can the pulsar survey benefit from deep, all-sky continuum surveys to increase the return of new pulsars and keep down the cost of the survey? We will use the SKADS continuum sky simulations [?] to inform a survey strategy.

Radio pulsars have two characteristics that could, in principle, be used to identify them in continuum images. They are intrinsically point sources, although they are broadened by scattering in the anisotropic interstellar medium (ISM). In addition, they are intrinsically steep spectrum sources, with a steep average spectral index, typically quoted to be around -1.6 ± 0.4 [?]. Bates et al. (2013) suggest that there are biases in earlier estimates of the spectral index, and suggest a mean spectral index of -1.4 and a standard deviation of 1. If this is confirmed, the suggested broadness of the distribution makes the sources' spectral index significantly less powerful as a discriminator between pulsars and flatter spectrum sources.

The angular broadening due to scattering by the Galactic ISM is generally smaller than the synthesized beam size of the proposed arrays at frequencies around a 1 GHz. There are counter examples, which consist of sources scattered by their local environment. Scattering is larger on the Galactic plane and at greater distances.

Ignoring the spectral index, which is a valid starting point given the broad distribution of pulsars and the limitations of the SKADS database, the SKADS simulations suggest the following population of sources, detectable at $10\text{-}\sigma$ in a $1\mu\text{Jy}$ -rms survey, per square degree at 1 GHz:

number of AGN per square degree:	5315.2
number of Starbursts per square degree:	10742.9
Total number of sources per square degree:	16058.1

The number of AGNs can be broken down to:

number of RQQ per square degree:	3690.1
number of FRIs per square degree:	1618.7
number of FRIs per square degree''	6.3

There are solid arguments to suggest that most radio quiet quasars will be completely unresolved, as the core and lobes of FRIs can be well associated with each other. That leaves a hard limit of ~ 3690 AGN per square degree. It is not straightforward to calculate how many of the host galaxies to these AGN will also be unresolved, but for $\sim 1''$ resolution, and based on the numbers of unresolved starburst galaxies returned by the SKADS simulation, it is likely that ~ 500 AGNs will have unresolved host galaxies. These numbers depend critically on the configuration of the array performing the survey.

This alternative survey for pulsars requires a deep and high resolution continuum survey, which constitutes a key science goal of SKA. In a document by SKA Science Director R. Braun, addressing science possibilities with the SKA₁ survey telescope, he proposes using the two discriminators mentioned above to pick out steep spectrum sources from the SKA₁ survey continuum survey, and use them for a targeted pulsar search. According to this document: *"If a combination of source angular size, $\theta < 1$ arcsec, and spectral index, $\alpha < 1$, constraints is employed, then its likely that more than 90% of the pulsar population can be retained, while rejecting more than 99% of other source populations."* We find that this statement is unlikely to be true given the numbers above.

It is important to note in this context that the traditional pulsar survey requires a large compact core in order to increase sensitivity within large, tied-array beams, which can be used to search the whole sky within the given time constraints of a pulsar survey. The efficiency of an alternative pulsar survey increases with increasing spatial resolution, so the two types of surveys have different optimal array configurations.

Table 1: Relative number of pulsars found with SKA₁-LOW at different central frequencies assuming a bandwidth of 100 MHz and correcting for the available collecting area as given by the equation $A_{\text{eff}} = 4.8 \times 10^5 \text{ m}^2 \times (f/110 \text{ MHz})^{-2}$.

$f(\text{MHz})$	100	150	200	250	300	350	400	450	500
$N(vL)$	1.0	1.15	1.20	1.17	1.09	1.00	0.91	0.81	0.72
$N(\text{PSRPOP.py})$	1.0	1.38	1.60	1.68	1.72	1.68	1.60	1.50	1.4

It is clear from the SKADS simulation, that the continuum survey cannot be used to replace the traditional pulsar survey. The continuum surveys are not deep enough, and the all sky version has a spatial resolution ($\sim 1 - 2''$) that will likely result in a large number of point-like sources, thus not necessarily reducing the targets on the sky.

Despite the limitations however, if the number of point source targets can be reduced to 500 or less, a very useful shallow survey (compared to the 0.1 mJy kpc² at 10 kpc for the traditional survey) for pulsars could be carried out using the continuum survey data. Off the Galactic plane, where the traditional pulsar survey has lowest efficiency, continuum surveys could play an important role for a first pass of an all sky pulsar survey. With this option, a continuum survey can be an interesting supplement to certain parts of the standard SKA pulsar survey. It does not, however, replace them.

2.3 Searching for pulsars with SKA₁-LOW - Stappers, Cooper, Bates, Thandi, Dodkins, van Leeuwen

In the present baseline design there is no capacity for carrying out pulsar searching with SKA₁-LOW and as we shall show below that is an unfortunate oversight which has important consequences for the overall use of SKA₁ and SKA₁-MID as well. The large collecting area and field-of-view of SKA₁-LOW combine to make it ideal for searching for pulsars, especially away from the Galactic plane as discussed in more detail below.

Using the baseline design as a guide we consider the collecting area of around 500 dishes which extends out to a radius from the core of about 600 m. The collecting area for the log-periodic dipoles is given by $A_{\text{eff}} = 4.8 \times 10^5 \text{ m}^2 \times (f/110 \text{ MHz})^{-2}$ and we used the formula given in the baseline design to calculate the approximate system temperature which is then modulated depending on where in the Galaxy we are pointing.

The first aspect of determining the usefulness of SKA₁-LOW is which frequency is optimal for pulsar search. This will be a tradeoff between the spectral index of pulsars, the temperature of the sky and the collecting area available. Using two different survey simulation codes (incl. *PSRPOP.py*) we find that the optimal frequency lies in the range 200–300 MHz with an assumed bandwidth of 100 MHz (Table 1). Given the uncertainty in the scattering relation and spectral indices we chose to do the remaining simulations at a central frequency of 300 MHz. Comparison of the two different simulation tools show that we have quite good convergence on the total number of pulsars expected to be found at this frequency and so from now on we consider only the simulation results from running *PSRPOP.py*.

To consider the usefulness of surveying with SKA₁-LOW we first compare the number and type of pulsars it will find with those found by SKA₁-MID. Using the parameters given in Table ?? we find that all-sky surveys with SKA₁-MID find about 12,000 normal pulsars and about 1000 MSPs while those with SKA₁-LOW find about 9000 normal pulsars and about 560 MSPs. On the surface of it this

suggests that SKA₁-LOW is not competitive with SKA₁-MID but if we look at the distributions of the number of pulsars as a function of dispersion measure, we see that at some DMs SKA₁-LOW finds more pulsars. The reason for this becomes clear when we consider a side on view of the distribution of pulsars in the Galaxy. This clearly shows that the lower frequencies observed with SKA₁-LOW mean that the dispersion smearing and the scattering in the interstellar medium along lines of sight to and through the Galactic plane reduce the number of pulsars that can be detected there, while the superior collecting area of SKA₁-LOW means that it is able to find more pulsars nearby.

Table 2: A selection of the most relevant survey parameters for both telescopes used in the simulations.. These parameters have been chosen from the baseline design. Note that the bandwidth used is smaller than used in the considerations presented in the frequency comparison table. The survey files also contain: survey degradation factor, sampling time (ms), channel bandwidth (MHz), number of polarizations, full-width half maximum (arcmin), minimum RA (deg), maximum RA (deg) minimum DEC (deg), maximum DEC (deg), minimum Galactic longitude (deg), maximum Galactic longitude (deg). Note in particular the limit on the elevation for SKA₁-LOW which comes directly from the sky visible to the log-periodic antennas.

	SKA1-Mid	SKA1-Low
Centre Frequency (MHz)	1400	300
Antenna Gain K/Jy	6.7	23
Bandwidth (MHz)	300	50
System Temperature (K)	20	46
Integration Time (s)	600	600
Minimum elevation (deg)	10	30
Min abs(Galactic latitude) (deg)	0	0
Max abs(Galactic latitude) (deg)	90	90

2.3.1 Composite Survey

The results of these simulations indicates that a composite survey where the complementary regions are covered by the two telescopes provides the best combination to maximise the number of pulsars that can be found with SKA₁. To give an idea of the best combination of surveys with the two telescopes we performed further simulations as shown in Figure 4 where the area surveyed with a given telescope was modified until a maximum was reached in the number of pulsars discovered. As can be seen the best combination for *normal* pulsars occurs when the Galactic plane, out to about $\pm 5^\circ$ is surveyed with SKA₁-MID (plus the Northern bits of the sky that SKA₁-LOW cannot see) and SKA₁-LOW surveys the rest of the sky that it can see. For the MSPs is a little less clear where the cut-off latitude would be and this is likely due to the greater influence of the DM and scattering on the MSP detectability. It is also important to note also that as expected the total number of pulsars, and MSPs, discovered is more than either survey would be able to find on its own.

An optimal survey strategy would therefore seem to be to search with SKA₁-MID up to Galactic latitudes of about $\pm 10^\circ$ and the region of the sky in the North where SKA₁-LOW cannot reach. SKA₁-LOW would then be used to survey the rest of the sky down to a Galactic latitude of $\pm 5^\circ$. The reason for the overlap in this region is two-fold: the first is to allow of effective cross calibration of the two surveys, to ensure that the relative sensitivities are understood as required for effective modelling of the

population, the second is because as we can see from the MSP simulations that it is apparent that in this region of the sky the two telescopes are finding different sources, SKA₁-LOW faint nearby objects and SKA₁-MID further away objects, and thus this will maximise the return on MSPs, the key target of the survey.

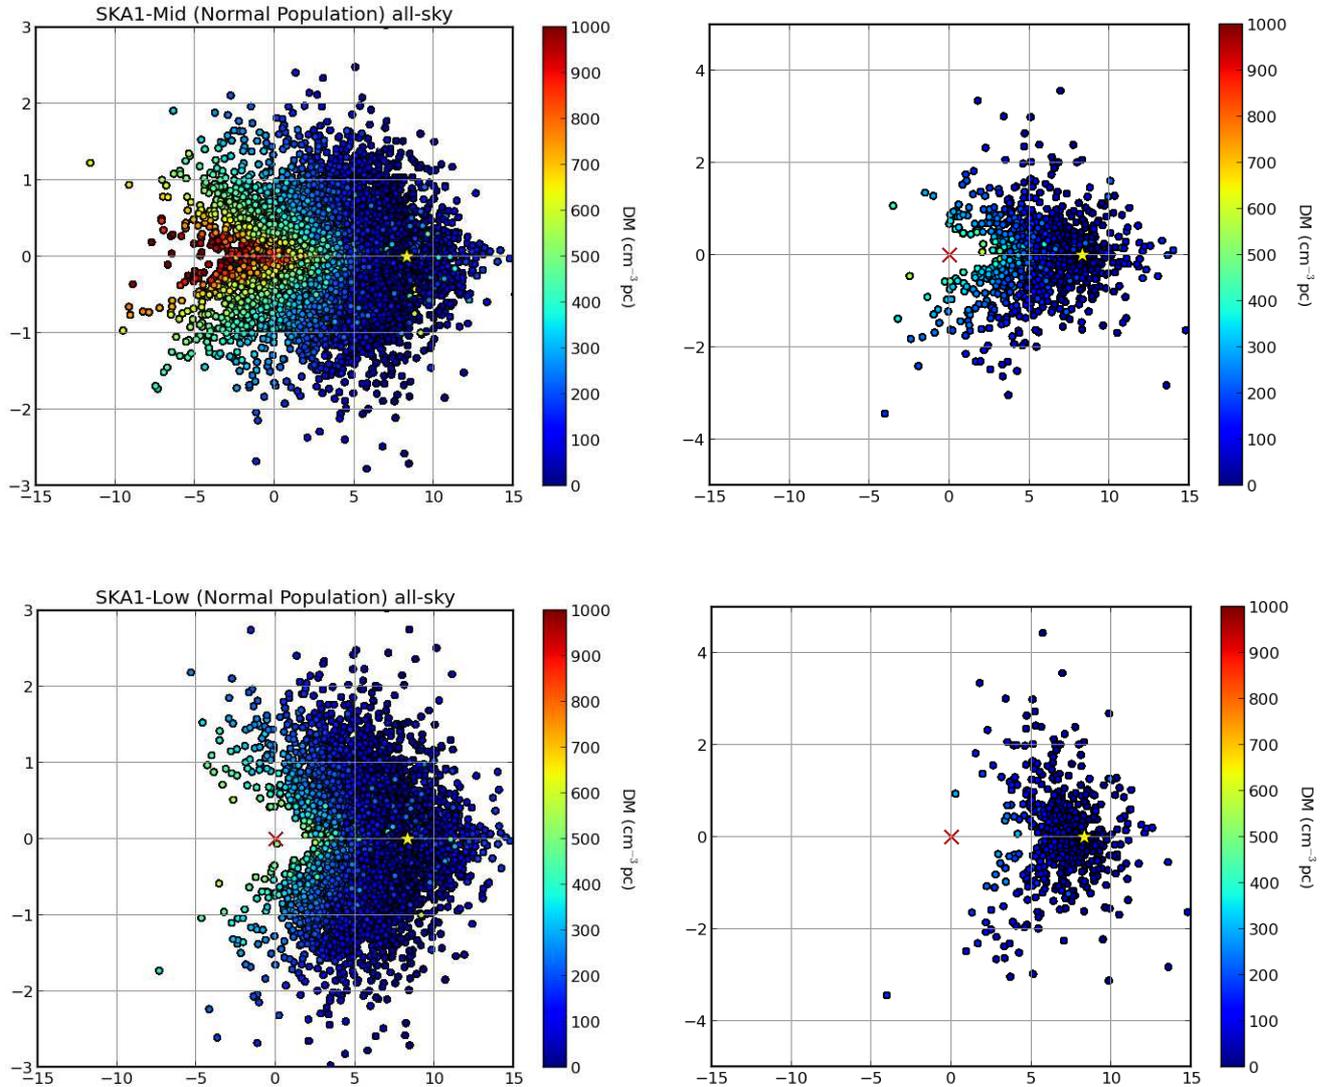


Figure 3: A comparison of the discovered population of pulsars with SKA₁-MID (top left: normal pulsars; top right: millisecond pulsars) and SKA₁-LOW (bottom left: normal pulsars; bottom right: millisecond pulsars) . Simulation parameters for these surveys are given in Table 2

2.3.2 Achieving the Composite Survey

To carry out a large area survey with SKA₁-LOW would require that a *beamformer* and a pulsar search backend be built in Australia. In the present scenario we have considered that we would use the collecting area of SKA₁-LOW out to a radius of 700 m to include a total of 500 stations. Although this is a larger radius than used for SKA₁-MID the significantly lower observing frequency means that the

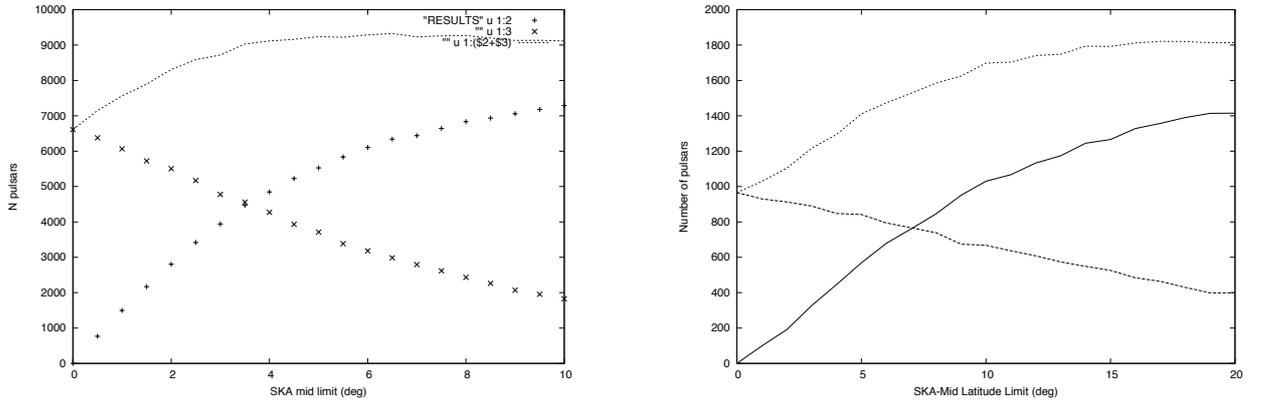


Figure 4: *Left*: Comparing the total normal pulsar yield with SKA₁-LOW (crosses) and SKA₁-MID (plusses) searching different regions of the Galaxy. The SKA mid limit corresponds to the region of the sky, between certain the Galactic latitude limits given, that is surveyed with SKA₁-MID - so when that is zero, then the whole sky is surveyed with SKA₁-LOW . *Right*: The same but now for MSPs. The influence of larger DMs on the detectability of MSPs means that the trade off latitude is a bit higher.

beam size is about 7 times larger for SKA₁-LOW and so to achieve the same survey speed less beams would be required.

We consider here a scenario where we preserve the total number of beams defined in the baseline design for the SKA₁-MID beamformer of about 2048 beams and we split them across both telescopes with SKA₁-LOW having 500 beams and SKA₁-MID having 1500 beams. As the pulsar survey processing *cost* scales approximately linearly with the number of beams this transferring of the beams from SKA₁-MID to SKA₁-LOW is almost cost neutral. A more detailed investigation is needed to know whether there is any additional cost for the beam former to effectively split it into two separate beamformers. We note that this combination of beams and search areas and splitting across two telescopes means that for the nominal total survey time of 2 years that a modest increases in survey integration times would be possible improving sensitivity and further increasing the number of pulsars that will be found. We also note that the in this scenario the decrease in the area to be surveyed with SKA₁-MID would allow for a smaller tied-array beam size. This would then allow a *modest* increase to the radius out to which dishes are included in the tied-array beamforming to achieve an effective collecting area as currently specified in the baseline design and thus allowing for some dishes to be moved from the inner core to the outer core. A pulsar survey with SKA₁-LOW could be carried out during the large fraction of the time when the EOR experiment is unable to observe and thus will significantly improve the efficiency of the telescope.

Further ways of tuning the pulsar yield is given by possible changes to the integration times in both parts of the survey. A given amount of observing time can be split when takes also into account the lesser area that needs to be searched for a given telescope type. If one changes the integration times for SKA₁-LOW to 1800s and SKA₁-MID to 2600s respectively, the relative yield will be different from the results with a fixed integration time of 600s for both telescopes. This optimization can be performed later, when further results from the ongoing surveys like HTRU, LOFAR and GBT-300 MHz are available to make an informed judgement. It is clear, however, that SKA₁-LOW should be equipped

with searching capabilities. *

2.3.3 Searching the Galactic Plane with SKA₁-MID outside Band 2 - Kramer

Flexibility is important for pulsar searching as the requirements depend on specific sources or sky locations to be covered. In addition to the importance of being able to search with SKA₁-LOW as described above, it shall also be possible to undertake a pulsar survey using any of the frequency bands available in SKA₁-MID. Restricting it to a specific choice of receiver band will cause a loss in science capabilities. Therefore, the processing bandwidth for SKA₁-MID should be at least 300 MHz. It should also be possible to survey with less bandwidth if required. In the following we describe the case of searching the Galactic plane with Band 3 (1.65–3.05 GHz), in addition to the “standard” search frequency of 1.3–1.6 GHz which the community has used in the recent Galactic Plane surveys (e.g. PMPS, PALFA, HTRU). Searches for pulsars in the Galactic Centre is likely to require even higher frequencies, such as Band 5.

According to the baseline design, the telescope efficiency and system temperature is identical in Band 2 and 3. The larger beamwidth available (808 MHz vs. 1403 MHz) is potentially extremely useful, if more bandwidth can indeed be processed. That would offset, at least in parts, a larger number of beams that is required while shortening the observing time, which be compensated by the larger bandwidth. Depending on the bandwidth that can be processed, a direct reduction in integration time has, as explained, a huge benefit for acceleration search requirements while maintaining similar sensitivity.

An attempt to summarize the effects is done in Figure 5 where we plot the impact of changing the observing frequency for a search on integration time, total survey time and requirements for acceleration searches. The integration time is scaled with the prerequisite that the same population of pulsars will be found, i.e. it is compensating for a reduction in flux density according to a power law spectrum with an average population spectral index of -1.6 . At the same time, the total survey time increases at higher frequencies due to an decreased beamwidth. The change in integration time also requires larger processing capabilities for an acceleration search. On the other hand, utilizing larger bandwidths would indeed balance these effects significantly.

The plots do not reflect the severe advantage of reduced interstellar scattering at high frequencies. Scattering smears out the pulse and eventually makes it undetectable. With a frequency dependence of typically ν^{-4} , scattering becomes prohibitive at low frequencies, so that the upper part of Band 2 has been chosen as the nominal search frequency for the Galactic plane, where scattering is strongest. For the inner part of the Galactic plane, however, strong scattering reduces the search sensitivity, eventually making it impossible to detect pulsars at others than highest frequencies. For this reason, the most inner part of the Galaxy should be searched at Band 3, where the exact strategy would be refined depending on Galactic longitude and available processing bandwidth. Note that for searches in the Galactic Centre a frequency at Band 5 would be most likely required (see Use Case).

*Note that the figure comparing the two telescopes doing the survey assumes a total available survey time of 2 years for *both* telescopes. Therefore, the integration time also changes as the fraction of the survey done with each telescope changes. Hence, the situation is somewhat more complex than a simple fixed integration time per pointing.

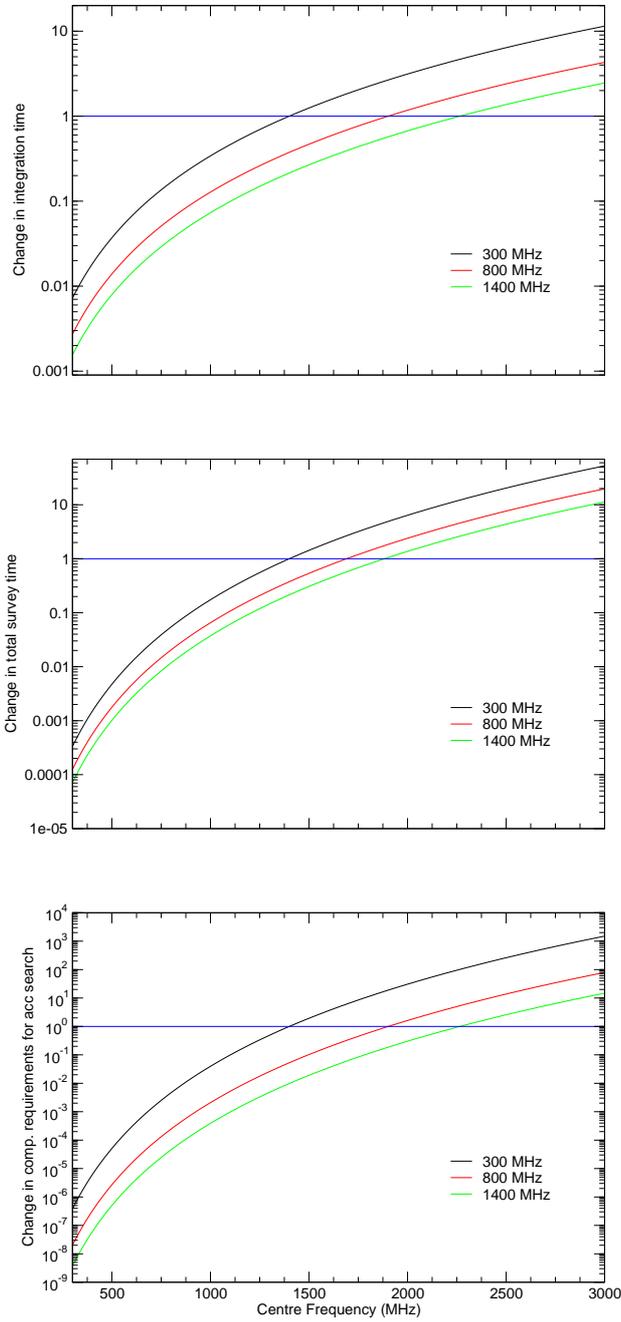


Figure 5: Relative changes in integration time (top), total survey time (middle) and computing requirements in acceleration searches (bottom) for searching at frequencies different from the “standard” search frequency of 1.3-1.6 GHz aiming for the same sensitivity.

2.4 Acceleration searches - Eatough, Ransom

As emphasized repeatedly above, acceleration searches are critical for the detection of the most exciting binary pulsar systems. Without such search methods, the chances of finding systems like double neutron stars and the elusive pulsar black-hole binary (PSR-BH) are greatly reduced. The parameters outlined in the Baseline Design (BD) for a pulsar acceleration search, as part of the non-imaging processing with SKA 1, are *broadly correct* with respect to the number of computational operations required (~ 10 Peta-ops). This is irrespective of time or frequency-domain search methods since both use similar numbers of operations, to within a factor of a few.

A large fraction of the currently known highly relativistic binary systems would be detectable with the BD specified time-domain acceleration search, however, it is worth noting that taking the example of the Double Pulsar system, where the acceleration can reach $\sim 300 \text{ m s}^{-2}$, this pulsar would not be detectable with full sensitivity at all orbital phases. We also note that much larger accelerations, than the $\pm 100 \text{ m s}^{-2}$ specified in the BD, can be expected in PSR-BH systems with compact orbits (e.g. $\sim 2 \text{ hr}$, acceleration $> 1000 \text{ m s}^{-2}$), although the likelihood of finding such compact PSR-BH systems is small due to their relatively short coalescence time. Nevertheless, we consider the acceleration range of the Double Pulsar as the least minimum and strongly suggest to increase the range to even 500 m s^2 for the most important and exciting systems.

In the following we study the impact further, assuming our bare minimum of 300 m s^{-2} and will derive at the following statements

1. Acceleration searching is a critical *real-time* task for pulsar surveys.
2. The number of operations specified in the Baseline Design (~ 10 Peta-ops) is roughly correct.
3. Some acceleration searching is better than none.
4. We will use all available computing power to do as much acceleration searching as possible.
5. It is too soon to decide between the two main forms of coherent acceleration searching (i.e. time-domain or frequency-domain). Both work well but have advantages and disadvantages.
6. Acceleration searching will almost certainly require high-performance computing (HPC) infrastructure (ala CPUs and GPUs or vector accelerators) and not pure hardware (ala FPGAs or ASICs).
7. Given the HPC nature of the requirements, the computing should *not* be designed now, but should be designed and purchased at the last possible moment in order to derive the most benefit from Moore's Law.

2.4.1 Detailed response to the Baseline Design

A consideration, important in determining the amount of computation required, is the acceleration sampling interval for time-domain searches or "step size", δa given by $\delta a = c/fT^2$, where f is the spin frequency, T is the observation length, and c is the speed of light. In the Baseline Design the step size optimally samples for pulsars of spin frequency 500 Hz (spin period 2 ms). Most of the power of pulsar signals is in the higher frequency harmonics, therefore ideally the step size should be optimal for the highest spin frequency harmonic. In highly accelerated systems typically only up to 4 harmonics can be detected. For the BD specified search this corresponds to a frequency of 2000 Hz giving an acceleration

step size smaller by a factor of 4 and a corresponding factor 4 increase in the number of acceleration trials. Both an increase in acceleration range (e.g. to $\pm 1000 \text{ m s}^{-2}$) and finer acceleration steps would result in more computational operations by a factor of ~ 40 . Such an increase would oversample a large volume of the parameter space (for the longer period pulsars) and as such is unlikely to be necessary. We note, though, that the most important thing is to actually perform acceleration searches as any acceleration is better than none. The amount of computations that are actually performed is entirely tuneable to the scale of computational hardware available. *We therefore believe that ~ 10 Peta-ops is a reasonable representation of the level of real-time computation required for a large sky-area pulsar survey.*

While it is still too early to fix the acceleration search algorithm, much wider searches in effective accelerations, at lower computational cost, can be offered by frequency-domain approaches. Assuming a similar number of computational operations as in the BD, an equivalent frequency-domain search could achieve spectral bin drifts, z of $z \sim 300$. The acceleration, a of a pulsed signal is given by $a = zc/fT^2$ (where f is the spin frequency, T is the observation length, and c is the speed of light) so for spin periods larger than ~ 16 ms, this would easily reach the accelerations we might expect in compact PSR-BH systems (e.g. 1000 m s^{-2}). Naturally therefore, the frequency-domain approach is sensitive to highly accelerated longer period pulsars (PSR-BH systems), moderately/highly accelerated mildly recycled pulsars (double neutron star systems), and low acceleration fast spinning pulsars (pulsar white dwarf binaries).

In addition, considerations of the optimal acceleration step size can be neglected in the frequency-domain approach since the space of frequency and frequency derivative are “gridded” uniformly on a bin by bin basis.

In general, the two coherent acceleration search algorithms are still in development for serious HPC use, especially in order to efficiently use GPUs or other vector/massively-parallel accelerators. It is basically too early to make a decision on which algorithm should be preferred, and each has distinct advantages and disadvantages.

2.4.2 Hardware considerations - with Bailes, Barr

The acceleration processing stage is the most compute intensive part of the search process, and the complexity of coherent pulsar search algorithms implies that dedicated hardware (such as FPGAs or ASICs) for realtime searches will not be possible. The search processing will almost certainly need to be handled by high-performance computing (HPC) infrastructure, such as normal CPU clusters although likely utilizing GPUs or other vector/massively-parallel accelerators (like Xeon Phi systems).

In order to give a numerical example, we consider a fractional bandwidth from 900-1650 MHz, 2048 channels and $64 \mu\text{s}$ sampling. In this case, the number of trial DMs from 0 to 1000 pc cm^{-3} to avoid more than 25% loss in sensitivity due to dispersion smearing on a 2048 channel filterbank is 2506, if the minimum sampling time is $64 \mu\text{s}$ and the assumed intrinsic width of the narrowest intrinsic feature is $40 \mu\text{s}$ is 2506 trials. We can assume that in most cases for DM exceeding 1000 pc cm^{-3} , scatter broadening mentioned above will dominate the dispersion smearing, so that the extra computational time is negligible.

In order not to lose more than 25% of the SNR, an acceleration search out to $\pm 300 \text{ m s}^{-2}$, we need 200 trials for every $8 \text{ Mpts} \times 64 \mu\text{s} = 512 \text{ s}$ at $\text{DM} = 0$. At current times, on a Kepler K10 one achieves 73 trials/second using both halves of the GPU. The search time is therefore:

$$N_{\text{beams}} \times 200 \times (a_{\text{max}}/[300 \text{ m s}^{-2}]) \times 2506/71/512/N_{\text{GPU}} \times \text{Real time.}$$

For $N_{\text{beams}} = 1$, $a_{\text{max}} = 300 \text{ m s}^{-2}$, $N_{\text{GPU}} = 1$, we get $13\times$ real time per beam / N_{GPU} .

In practice, because of the DM smearing, the real computing cost is much lower as the number of acceleration trials goes down with increasing DM smearing time. Out to DM of 1000 pc cm^{-3} , this averages down to $5.6\times$ real time per beam/ N_{GPU} . Loosely, this scales as $(T_{\text{obs}}/512)^2$. On specific systems, such as the GTX690 (although not available anymore), these times are less by about 40%, resulting in about $3.5\times$ real time. This is for the re-sampling method in the time domain. Note that reducing the maximum loss to just 10%, these numbers scale by a factor of 2.7.

Note that the spacing where one loses less than about 10% of a harmonics power in terms of bins drifted (or "z") is 2. For $\pm 300 \text{ m s}^{-2}$ with 200 trials, that means an acceleration spacing of 3 m s^{-2} . With $a = zc/fT^2$, the frequency where one starts doing worse than what is basically interbinning in the frequency direction (but this time in the fdot direction) is (where $z = 2$):

$$f = 2c/(daT^2)$$

In the above case, we specified T to be 512 s, so that $f = 2 \times 3 \times 10^8 / (3 * 512 * 512) = 762.9 \text{ Hz}$. For a fast spinning millisecond pulsar, we will be assured of getting the majority of the power for the first few harmonics. In general, this should be sufficient, as most highly accelerated pulsars are not detected in more than 7 or 8 harmonics. Nevertheless, the above should be considered as a lower limit only and more acceleration trials are highly recommended.

Concluding, at this stage it would be unwise to design the computer architecture for acceleration searching as this does not depend on the signal transport or imaging/beam-forming stages. The only required input for the acceleration searches are de-dispersed time series of full integrations from the phased array beams (or alternatively from a pixel time series from fast-dumped imaging).

2.5 Quality Assurance / RFI mitigation schemes

Work in progress.

2.6 Storage considerations for candidates - Champion, Chatterjee

The information about candidates that is stored as the output of survey pipeline serves two purposes:

1. To allow an algorithm, neural net or human to decide if the candidate is a likely pulsar.
2. To provide a first TOA for any pulsar discovery.

Current surveys usually create a summary page per candidate that allows the viewer to evaluate the candidate using several plots e.g. phase vs time, phase vs frequency, S/N vs DM and S/N vs acceleration.

This information can be stored as an image file or as the data required to reproduce this information.

As an image this information typically only requires 200kB per candidate. However this is difficult to translate into a machine readable format, cannot be further analysed by removing subbands or sub-integrations and cannot provide a TOA for a pulsar discovery.

Alternatively these plots can be reproduced from a down-sampled data-cube of phase, observation frequency and time. This down-sampling is usually very coarse, assuming that the DM and period have been well estimated by the pipeline. However unlike a conventional survey a survey with the SKA will not have any raw data that can be used to provide higher resolution information if required. However as

the data have already been folded at the candidate period the use of much higher resolution data-cubes would be limited to only improving the quality of a detection and cannot produce a new candidate.

The resolution required for each axis of the data-cube can be looked at independently: The phase resolution should not need to be higher than 256 bins across the profile. This provides enough resolution to resolve 95% of pulsars listed in the pulsar catalogue (W_{50}/P_0). For short period pulsars at high DMs this number can be reduced to account for intra-channel smearing without significant loss of information.

Time information can be used to check that the pulsar is emitting continuously (except in cases of intermittence and scintillation), to reproduce an acceleration vs S/N curve and sub-integrations can be removed to improve a detection. Thus the time resolution is largely dependent on the RFI situation (cadence, duration etc) of the site and how good the RFI mitigation in the pipeline is. Sub-integrations of 10 seconds is fairly common for pulsar observations and allows non-repeating RFI to be removed with moderate loss of data that is not RFI affected. For the 600 second pointings of the survey this gives 60 sub-integrations, this is also enough sub-integrations to allow for an acceleration vs S/N curve to be recreated.

Frequency resolution is used to check that the pulsar is broad-band (except in cases of scintillation), to reproduce a DM vs S/N curve and subbands can be removed to improve a detection. With the time resolution, this is dependent on the RFI situation of the site and upstream RFI mitigation. A number of 32 subbands is a typical value for current surveys and allows a DM curve to be reproduced.

Note that polarisation information is not usually used in pulsar surveys, only the summed intensity is recorded. As an archive file this comes to ~ 1.2 MB per candidate. In addition metadata for each candidate is stored. This is typically less than 1kB (in addition to telescope information e.g. position).

The above gives full flexibility for post-processing. In order to save storage, it can be used a compressed version of candidate files, e.g. one that is loosely based on the structure of a PMPS-''PHCX'' file (Keith et al.). The main differences are:

- 8 bits per sample in the plots, i.e. sufficiently large for greyscale plotting etc.
- Keeping both frequency-scrunched data and time-scrunched data to a certain resolution, not the full cube.

This gives about 28 kbytes. Keeping the full cube, but still only using 8 bits, results in 254 kbytes, which is about a factor of 4 smaller than what is outlined above, which uses 32-bit values. In general, we suggest to keep as much information as possible but would also be able to work with a bit-reduced version if necessary.

In contrast to timing observations, search observations do not have a significant need for calibration. However, for normalization purposes, one may want to save the bandpass shape and monitor its stability. Nevertheless, the information will not add any significant overhead.

The effort needed in order to examine data for RFI mitigation can be substantial, but the data storage / timescale needs depend on the level of detail that can be stored per candidate as discussed above. However, this information only needs to be stored as long as candidates have been identified, which needs to be done quasi-instantaneously. Hence, maintaining a rolling buffer would be mostly sufficient as long as the buffer covered the timescale for data processing (so selected data chunks stored for days, not seconds or minutes). In order to monitor the long-term environment, sample data need to be kept for timescales of days and weeks.

3 Pulsar Timing

To elaborate further, the timing precision of pulsars is governed by a number of factors, one of which is sheer integration time on a given object to obtain a TOA. Over the set of pulsars that will be timed over many years, which will encompass a range of flux densities over the sample, different configurations can be allocated. These may include subarrays for strong pulsars, and the full array for weak ones. Multiple pixel feed arrays will increase the flexibility. With Phase I, the timing throughput will not be as large as with the full SKA and it is certain that the complete set of goals of the Pulsar & Gravity KSP will not be achieved in Phase I. In the case of a pulsar-black hole system discovered in a Phase I survey, the integration time cannot be necessarily traded for collecting area, as the full SKA sensitivity is needed to reach an instantaneous timing precision that is sufficient to map the pulsar orbit around its companion with appropriate resolution in orbital phase. Similarly, full-SKA sensitivity may be needed to discover fast binary pulsars in the first place, until sufficient computing power for full-orbit search techniques is available.

The Phase I of the SKA will also be crucial in constructing an enlarged pulsar timing array (PTA) before it is fully exploited with the completed SKA. For most known pulsars we can expect the timing precision to increase by a factor of 10 due the availability of an Arecibo-sized collecting area. Together with more suitable pulsar being discovered and timed, the sensitivity of the PTA should exceed the current activities by about a factor of 100. Achieving the final PTA sensitivity of the full SKA will require precise polarization and gain calibration techniques and the observations of pulsars that behave a stable pulsar clocks. The establishment of these procedures and the identification of stable pulsars is a crucial task for the Phase I observations.

3.1 Impact of loss in sensitivity - Kramer

In Section 2.1, we have discussed severe impact of the loss of overall sensitivity on pulsar searching and the type of discoverable systems. The corresponding response was triggered by a discussion of whether dishes from the core of SKA₁-MID can be moved out to intermediate or outer baselines. The impact on pulsar timing can be equally severe.

In order to exploit the most important and exciting binary systems for testing theories of gravity, we need to measure the pulse times-of-arrival (TOAs) with sufficient time resolution across the binary orbit. This requires a certain amount of raw sensitivity, as it will not be possible to integrate longer as this smears orbital phase information. For compact and eccentric orbits, in particular near periastron, the important TOAs are spaced very closely in time and we simply need raw sensitivity to resolve the orbit. It is very difficult to quantify the impact of, say, 10% loss in sensitivity, as we cannot predict how strong the rare systems are going to be. However, given that we may have found the strong pulsars already, they are probably weak, even on SKA scales. Nevertheless, sensitivity for timing (and moving out dishes) may not be a problem for pulsar timing, *if* it is possible to phase-up the whole array (i.e. beyond the core) for specific directions on longer baselines.

3.2 Operational considerations – Karastergiou, Huynh

For pulsar timing experiments with SKA₁-mid dishes, observations across multiple frequency bands are essential in monitoring and removing variability in each pulsar's dispersion measure (DM). DM variability is typically smaller than 1 part in 10000 (with exceptions, e.g. Keith et al. 2012), however may result in significant residuals in the timing. The reason for this is that a change in DM results

not only in a different relative delay between the two ends of the observing band, it also results in a different total delay of the signal at the centre frequency compared to a signal at a reference frequency (for example, infinite frequency). It is therefore essential to conduct pulsar timing observations across multiple frequencies, to accurately measure the DM and DM variability and account for it in the timing models.

As a consequence, it is necessary to have the ability to switch the receivers on the SKA₁MID dishes during each epoch of timing observations, to measure the DM with the highest possible precision. This precision is determined by the total extent of the frequency coverage and the time resolution of the observations.

Related to the above is the fact that, for some pulsars, the collecting area of a fraction of the array may be sufficient to produce high S/N timing data. The ability to create independent sub-arrays increases the observing efficiency. Each sub-array may then be used to observe a single pulsar, with different receiver between the arrays, or to observe multiple pulsars with a single receiver across all sub-arrays.

3.2.1 Science background from the DRM

Interstellar propagation effects have strong frequency dependences. As You et al. (2007) discuss, a DM variation of order $5 \times 10^5 \text{ pc cm}^3$ produces a time delay of order 100 ns, comparable to the level of timing precision required. Both the solar wind and the Galactic interstellar medium can produce DM variations, in some cases over the relatively short time scales of order 30 minutes. Simultaneous, or nearly simultaneous, multi-frequency observations, in which the center frequencies of the observations differ by of order the observation frequency, can exploit the frequency dependence of the propagation effects, particularly dispersion smearing, to mitigate or remove them. The optimal set of frequencies is likely to be direction dependent, and even the number of frequencies required is a topic of active investigation. Current results suggest that multiple frequencies will be required.

3.2.2 Comments on the Baseline Design

We note the following:

- There is no mention in the baseline design document of time scales between receivers changes. The following is however mentioned: *"Space at the focus for five independent receivers."* (in the context of SKA-I MID). Nevertheless, a discussion on having, on stand-by, multiple cryogenically cooled receivers in the SKA-I mid dishes for DM determination during pulsar timing observations is missing.
- Regarding sub-arrays the baseline design states: *"All of the SKA1 telescopes can operate independently as multiple sub-arrays (i.e., collecting area split and allocated to separate, concurrently observing programmes)."* In particular, for SKA₁-MID, *"The number of sub-arrays can be as large as the number of antennas."* and *"The beamformer must be designed to support beams from any number of sub-arrays."*
- Further it states for Level 1 requirements: *"Timing resolution. The timing resolution for Pulsar timing shall be better than 100ns."* We note that there is no additional requirement in the Level-1 requirements document, compared to the baseline design document.

- Regarding sub-arrays, it is stated *”Beam-former sub-array support. The central beam-former shall support beams from one to sixteen sub-arrays”*

Consequently, we suggest:

- Pulsar timing with SKA₁ has the objective of obtaining sub-100 ns timing residuals for as many millisecond pulsars possible. Given that uncorrected DM variations can result in a systematic effect in the timing residuals, with a magnitude comparable to or worse than the target, it is essential that each timing observation be used to measure an extremely accurate DM. Highest accuracy can be achieved by using multiple bands. It is therefore essential that receiver switches are possible on a timescale shorter than the timescale of DM variations.
- For optimal performance, we suggest the capability to switch between the 5 nominal receivers of SKA₁mid within 10 minutes.
- The exact observing strategy, i.e. time per frequency band, can be well determined and is different from source to source (Keith et al. 2012).
- The sub-arraying capabilities as stated in the Level-1 requirements document, that is a total of 16 sub-arrays, ensure efficient observing of multiple pulsars at multiple frequency
- One should be able to phase-up the total area of SKA₁-MID to produce beams in particular directions.

4 Pulsar Astrometry – Deller, Chatterjee with Bignall & Paragi

The limited baselines (maximum 50-200 km) of SKA₁-mid and SKA₁-survey mean that when operated in isolation, they do not have sufficient angular resolution for precision astrometry. Precision differential astrometry of pulsars allows astrometric parameters (position, proper motion and parallax) to be accurately measured and subsequently fixed in the pulsar timing model, facilitating improved strong-field tests of gravity, one of the key SKA science goals. Moreover, if a (astrometric) VLBI capability is enabled for SKA₁, it will also find many other applications; studying the neutron star equation of state, better modelling the Galactic electron density distribution, calibrating luminosities from gamma-ray pulsars and low-mass x-ray binaries, studying the distance and distribution of OH masers and low-mass star forming regions, frame ties with the optical reference frame, and detailed studies of AGN and core collapse supernovae, to name but a few.

This section first describes a use case of SKA₁ participation in a sensitive VLBI array for differential astrometry, followed by a description of the technical specifications required to meet these capabilities. We note that the technical requirements for astrometric observations are in general more demanding than most other VLBI use cases, and so the requirements detailed below will already enable most VLBI applications. We focus here on observations at 1.6 GHz, but as noted in the text, there are many applications where a higher frequency would be desirable to reduce the effects of ionospheric errors and scattering in the Galactic plane, or to target specific lines such as methanol (6.7 GHz, high mass star forming regions).

4.1 Astrometric VLBI observations utilizing SKA₁-MID

Differential astrometry (also known as relative astrometry) measures the position of a target source in a reference frame defined by one or more nearby calibrator sources. In radio astrometry, these calibrator sources are typically far more distant than the target source, and so in the absence of structural evolution in the calibrator source(s) this frame can be considered stationary.

The precision of the positional measurement is determined by three main considerations:

1. The noise-limited position fit, with an error σ_N ;
2. The interpolation-limited accuracy of the registration of the position within the calibrator frame, with an error σ_R
3. The stationarity of the calibrator frame, which can be perturbed by structure evolution in the calibrator source(s), with an error σ_F . σ_N is linearly proportional to the synthesized beam-size and inversely proportional to the S/N of the target detection. Reducing this error source demands long baselines and high sensitivity; higher observing frequency is useful if feasible.

The dominant contributions to σ_R are the unmodeled ionosphere (at frequencies below 5 GHz) and the unmodeled troposphere (at frequencies above 8 GHz). Positional wander due to refractive scintillation in the interstellar medium is also a small contribution at low frequencies, but is typically smaller than the ionospheric errors. At small target-calibrator separations, the error is roughly linearly proportional to the target-calibrator separation (although this is not always true for refractive positional wander, depending on the pulsar distance and calibrator characteristics). Reducing this error demands (preferably multiple) calibrator sources as close as possible to the target, which demands high sensitivity (there are many more faint calibrator sources than bright calibrator sources). Alternatively, the observing frequency can be increased, since both refractive image wander and ionospheric errors are inversely proportional to the frequency squared. Observing at times of low ionospheric activity is also desirable, but not always possible. In order to average away local effects, a large number of baselines (and thus many stations in the VLBI array) is desirable.

σ_F is the most difficult term to predict a priori, as previous astrometric campaigns have rarely if ever been dominated by this contribution. However, as σ_N and σ_R are reduced when employing the higher sensitivity of SKA₁, it is possible that σ_F will dominate. The best mitigation strategies are expected to be the use of many calibrators, so that any sources with gross structure evolution can be identified and removed from the calibrator frame, and small effects from the other sources will average away in a \sqrt{N} fashion.

Current pulsar astrometric campaigns typically make use of 1-3 calibrator sources within a radius of ~ 20 of the target pulsar. With these setups, astrometric accuracy of ~ 10 microarcseconds for parallax (a 10% distance at 10 kpc) is possible. With the order-of-magnitude increase in sensitivity available in an array utilizing SKA₁, we expect that a greater number of closer calibrator sources will be used (5–10 sources within 5–10 arcminutes). Coupled with the reduction in N due to the raw sensitivity improvement, this will considerably reduce R and F and make parallax accuracies of a few microarcseconds possible (10% distance error at ~ 25 kpc). This exceeds the planned accuracy for the GAIA mission in the optical by a factor of a few, although a far more limited number of sources will be possible.

5 Advantages/Disadvantages of simultaneous subarraying at different frequencies - Hessels

Some of the observing modes, pulsar searching and timing, may benefit from the possibility to split up the full arrays in sub-arrays. In the following, we discuss advantages and disadvantages.

Advantages of sub-arraying are:

1. The large number of SKA pulsar discoveries (of the order of 10,000 in SKA₁) will require follow-up timing observation. A basic timing model requires ~ 15 observations over the course of 1 year. More observations may be needed if the source is in a binary system. Though some sources may be bright enough to follow-up with other radio telescopes, it is likely that most sources will require the SKA. Sub-arraying can help alleviate the observing load for sources that are easily detectable with a fraction of the collecting area. It may also be advantageous in some cases to run pulsar timing experiments in parallel with imaging observations to boost the observing efficiency.
2. Sub-arraying can allow large instantaneous frequency coverage by using different subsets of the array with different receivers. For SKA₁-MID, the array could be split in 3 parts (not necessarily of equal collecting area) to provide 0.35 - 3 GHz instantaneous coverage using the envisioned Bands 1-3. Such wide frequency coverage may be important for precision modelling of propagation effects in the interstellar medium that limit the pulsar timing precision. In other words, the timing precision in some cases may benefit more from a larger instantaneous bandwidth than by having a large single collecting area.
3. While the high-filling-factor cores of SKA₁-Mid/Low are being used for pulsar surveying, the antennas on the longer baselines can be combined either incoherently or coherently (in some situations, and depending on calibration) in order to simultaneously observe known sources. Coherent sums of the longer baselines will produce small tied-array fields-of-view, but for known sources with reasonably well-known positions this will not be an issue.
4. Observing multiple sources simultaneously may have some calibration applications, such as removing the ionospheric effects to the highest-possible precision. Especially in the commissioning stages of the telescope, simultaneous observations may help debug processing issues.
5. Small fractions of the array can be used as a sub-array to test new observing modes and debug existing ones while production science observations are ongoing.

Disadvantages of sub-arraying are:

1. Splitting the collecting area into sub-arrays reduces the raw sensitivity in a linear way. Trying to recover the lost sensitivity through increased bandwidth or integration time is not necessarily fruitful because the sensitivity only scales as the square-root of these quantities. Thus, it is ultimately a question of whether the science goals are better served by raw sensitivity and short observations or lower instantaneous sensitivity but longer dwell time. Both ends of the spectrum are scientifically interesting. For example, if single pulses are needed, then instantaneous S/N is likely most important. Conversely, if enough pulses must be acquired that the cumulative profile is stable to some high level of precision, then longer dwell time may win out.

2. Designing an interferometric telescope that is flexible enough to provide sub-arrays requires more careful system design and a low-level degree of modularity in the system architecture. This may be more costly or difficult to implement. On the other hand, preparing a system capable of sub-arraying enforces a system design that is arguably better thought-out and "future proof" in terms of its adaptability to new observing concepts.

In summary, we conclude that sub-arraying is not useful for all applications, but the system should provide this possibility to allow for time-efficient usage of the telescope, in particular for timing a large number of pulsars, where the total required observing time may be dominated by pulse-phase jitter and resulting profile stabilization time-scales.

6 Overall considerations on data products, rates and volumes - Possenti

Previously, we have already discussed the storage requirements for storing and processing search candidates. We want to elaborate on this here further, in particular with the possibility that some of this information will be needed archived to assess data quality. For searching, we refer to the Pulsar Survey Scientific and Technical Requirements, as reported in the SKA₁-Design Reference Mission document (Table 5-1 and 5-2).

We also consider the amount of Data Products expected from the Pulsar Timing Scientific and Technical Requirements (Table 6-1 and 6-2). That allows us to estimate the amount of storage needed for Archiving the results of these projects and comment on that.

In agreement with the Figure 27 of the SKA₁ System Baseline Design document, note that in this context the kind of Data Products to be stored will be a combination of Level 3 (time series) and Level 4 and Level 5 data, where Level 4 will mostly consist of partially processed data and Level 5 of lists of observations/results included in a database structure.

6.1 Search Data Products

According to the SKA Design Reference Mission document, the expected Data Products in the context of the pulsar surveys will have to include *the sky position, pulse-average total intensity flux density, period, DM, and, if a binary orbit is found, a preliminary indication of the orbital parameters, as well as with the de-dispersed time series for each candidate*. All these Data-Products are of the Level 4 and/Level 5 (according to the definition reported in Figure 27 of the SKA1 System Baseline Design document), but for the last one (time series) which belongs to the Level 3 category. The last dominates, by far, the budget of the long term storage requirements for Data Products and Archiving.

We here assume the reference survey parameters reported in the Pulsar Survey Scientific and Technical Requirements (Table 5-1 and Table 5-2 of the Design Reference Mission document), i.e. temporal resolution of $50\mu\text{sec}$, integration time of 600 sec, full coverage of the entire visible sky above an elevation limit of ~ 10 deg. The first two requirements implies the collection of a time series of 12×10^6 samples for each of the N_{chans} frequency channels which the bandwidth is split into. In the case of the SKA1-mid pulsar survey (see Table 13 of the Baseline Design), the required sky coverage translates in about 100000 observations per survey, with about 800 beams per observation between 0 and 5 degrees latitude and about 2000 beams per observation above 5 degrees latitude, i.e. the full visible sky will be entirely tessellated by looking at about 2×10^8 positions. In this case, assuming to sample at N_{bits} and to retain N_{cands} per each sky position, the Data Products to be long term Archived will be

$$\text{Storage Requirement (SKA1 – mid Pulsar Survey)} \sim 300 N_{bits} \times N_{chans} \times N_{cands} \text{ TBy}$$

According to the Design Reference Mission document $N_{chans} = 1$. However, retaining the candidate time series in a significant number of frequency sub-bands will be very valuable as a diagnostic tool for investigating it off-line (i.e. via automatized procedures applied to the candidate and to the data later collected for confirming that the candidate is really a pulsar). Ideally, one would like to retain the N_{chans} time series for all the pulsar candidates down to a low signal-to-noise threshold. That would be even more important for the SKA pulsar surveys, since the vast majority of the raw data resulting from them will not be stored. However, given the large demand of required storage, the value of N_{cands} will have to be limited. On the basis of the recent pulsar surveys [to be better constrained with the examination of the output of the data analysis of various ongoing surveys] a reference choice is that of adopting $N_{chans} \times N_{cands} \sim 100$. Also assuming $N_{bits} = 8$ (in order to preserve sensitivity in the beam forming procedure), it turns out the need of at least *300 PB* for storing the Data Products of the SKA1-mid Pulsar Survey. At current prices this sums to a total of about 10 Million euro, but it is expected to be at least a factor 10 less in a 5-yr time scale, representing a negligible fraction of the cost of the project.

Besides the $N_{chans} \times N_{cands}$ candidates (the top-ranked in the candidate reliability list) for which the entire N_{chans} time series will be preserved, the data for additional lower-ranked candidates could be stored after folding the N_{chans} time series in a series of sub-integrations, adopting the constant spin period of each candidate. With this procedure, it will be formed what is usually called a "Pulsar Archive" (in the sense of the most common pulsar software for handling them, namely `psrchive`). The number of bins for producing the folded profiles, as well as the number of sub-integrations and frequency channels, can be optimized on the basis of the spin period and DM of each candidate, accounting for the time resolution of the collected data and the possible occurrence of Doppler smearing due to orbital motion. The rapidly spinning pulsars requires a larger number of sub-integrations than the slow ones, but this is partially compensated by the smaller number of needed bins for the folded profiles. Thus the required amount of storage is similar for the various classes of pulsars and, for a given choice of N_{chans} , it results a factor ~ 100 smaller [this estimate can be improved with the study of the candidate population in ongoing surveys or with a simulation] than that necessary for storing the original time series.

Given what above, at the price of simply doubling the total data storage requirements (up to about 600 PBy), it may be possible to permanently archive partially processed data for ~ 100 times more candidates, thus including almost all the candidates resulting from a given position in the sky. This may be very useful for unveiling ultra short period binaries. In fact, their high acceleration strongly smears the signal when the N_{chans} time series are folded at the best (constant) period resulting from the non accelerated analysis and this implies that the candidate may slip down the ranking list. However, if the pulsar is bright enough, the signature of the orbital motion will neatly appear when inspecting (either visually or automatically) the "Pulsar Archive" associated with the low-ranked candidate. A simple (and computationally light) cross correlation of the Pulsar Archives resulting from a two-pass survey could thus pick up and immediately unveil at least some of the most aimed targets of the pulsar search experiment (i.e. the pulsars in a ultra short orbit).

6.2 Timing Data Products

According to the Design Reference Mission document, the expected Data Products in the context of the pulsar timing will have to include *folded pulse profiles, at full polarization, and pulse times of arrival, with estimates of the uncertainties on those, in the solar system barycenter for a specified ephemeris, as derived from full polarization observations, for a specified list of timed pulsars, which specifies the cadence and integration time for each pulsar.*

Any pulsar to be timed has its own optimal combination of observing parameters as well as parameters for storing the folded pulse profiles. We here follow the general indications of the reference Survey parameters reported in the Pulsar Survey Scientific and Technical Requirements (Table 6-1 and Table 6-2 of the SKA1Design Reference Mission document), as well as the parameters of Table 14 of the System Baseline Design document. Thus, as a reference set of parameters, we adopt the hypothesis of 48 hours a month devoted to pulsar timing, with a typical observation lasting 10 min and looking at 20 targets (some of them within the same beam among the 10 formed beams), with temporal resolution of 100 ns. In term of rate of observed targets R_{targ} , this translates in $R_{targ} = 48 \text{ hr} \times (60 \text{ min}/10 \text{ min}) \times 12 \text{ months} \times 20 \text{ targets} \sim 70000 \text{ targets/yr}$.

In order to fully exploit the required temporal resolution, the folded profiles are expected to have between 10^4 and 10^6 bins for rapidly spinning pulsars. As a reference value we adopt $N_{bin} = 10^5 \times N_{bin,5}$ where $N_{bin,5}$ is the number of bins in units of 10^5 . We assume as a reference that the baseband data will be acquired in N_{chan} sub-bands and hence will then produce $N_{chan} = 10 \times N_{chan,1}$ channels, where $N_{chan,1}$ is the number of channels in units of 10^1 . For rapidly spinning pulsars, it is preferable to store the coherently dedispersed folded profiles in short-duration sub-integrations, both for isolating and clipping sub-integrations affected by RFI and/or for monitoring short term variations in the emission. As a guideline, we can adopt 1-sec long sub-integration, which, for a 10-min total integration, translates in $N_{sub} = 600 \times N_{sub,600}$ sub-integrations for each targeted pulsar, where $N_{sub,600}$ is the number of sub-integrations in units of 600. Also, the folded profiles will have to contain full Stokes information, which implies an additional increase of a factor $N_{stokes} = 4$ in the size of the archived products. Finally, we assume that the folded profiles will be represented with $N_{bit} = 8 \times N_{bit,8}$ where $N_{bit,8}$ is the number of bits in units of 8. With all these parameters set, the amount of storage required for one timing observation of one target would be $B_{targ} = 10^5 \times N_{bin,5} \times 10 \times N_{chan,1} \times 600 \times N_{sub,600} \times N_{stokes} \times N_{bit,8} \text{ Bytes} = 2.4 \text{ GBy}$.

In summary, a reference value for the size of the stored data resulting from one year of timing observations of rapidly spinning pulsars is

$$R_{targ} \times B_{targ} \sim \mathbf{170 \text{ TBy/yr}}$$

It's worth noting that the value above is likely an upper limit for coherently dedispersed timing data, since the data storage requirements for slower rotating pulsars (which may represent a fraction of the total population to be timed) will be significantly less than B_{targ} and the rate R_{targ} has been calculated for 100% efficiency in the observations, i.e. non including the time required for repointing and/or switching between the observing bands. Therefore, even assuming a timing program lasting 10 years, the total requirements in term of storage of the results of coherently dedispersed timing observations will remain at least a factor 100 below the requirements for a 2-year pulsar survey.

A caveat to the estimate above might be that one will succeed to use single-pulse analysis at full timing resolution for performing the regular timing of the most relevant targets. On the basis of the present wisdom, this will likely not be a effective option. However, if adopted, for rapidly spinning pulsars, the product of $N_{sub} \times N_{bin}$ would increase a factor ~ 100 with respect to the reference values

adopted above, affecting, by the same factor, also the total data storage requirements.

6.3 Comments on storage costs

Costs can be estimated based on an extrapolation from the rate of decrease of the prices for hard disks during last 30 years. We are now at the level of about 0.03 Euro/GB=30 Euro/TB for the media-drive. Of course, the "server" for making the disks promptly available (and always online) and the related software would be a cost on top of that, which is more difficult to quantify. Another cost is related to personnel and replacement parts for maintaining the system. On the basis of a study of 2007 (REF), all the additional costs may at most triple the cost of the disk-drives on a yearly base. In summary, for a conservative estimate, one could increase the total cost mentioned above by up to a factor 3 for each of the 2 years of the survey.

We note that tapes (although less handy in the use) are likely still cheaper. In 2007, according to the aforementioned study, there was an overall factor 3 better for tapes than disks. However, we consider tapes (and the required maintenance costs) as an unlikely option for the SKA.

In summary, making the assumptions above, it may require an investment of order 1-3 MEuro/year for establishing and preserving the entire pulsar database. Nevertheless, costs will decrease every year.

Appendix A – Use Cases

6.4 Pulsar astrometry use case

The distance to a pulsar determines the size of the parallax signature, which in turn determines the astrometric precision required, and hence the maximum allowable values for σ_N , σ_R and σ_F . Below, we give a typical example, and note that this can easily be extrapolated to pulsars that are at different distances or are brighter/fainter.

Consider a moderately faint pulsar with a period-averaged flux of 40 μ Jy at 1.4 GHz, at a distance of 15 kpc. If the pulsar has a 10% duty cycle (typical) then the equivalent flux density after accounting for the S/N improvement of pulsar gating will be 120 μ Jy. The parallax signature at 15 kpc is 67 μ as, so a 10% distance requires a parallax accuracy of 6.7 μ as. Such accuracies are not even possible for bright pulsars today, but should be achievable with a SKA₁-enabled VLBI array for even faint pulsars - the vast majority of pulsars known today are brighter than the 40 μ Jy used in this example. [†]

Assuming a typical observing campaign (8 epochs spread over 18 months), a parallax accuracy of 6.7 μ as demands accuracies on each individual epoch of around 20 μ as. For σ_N , assuming a beamsize of 5 mas (which requires 8,000 km baselines, see below), this demands a S/N of 125, and hence a thermal noise of $\sim 1\mu$ Jy. This thermal noise can be reached by a VLBI array including both SKA₁-mid and SKA₁-survey in around 3 hours on-source, assuming 500 MHz of bandwidth, sampled in dual polarisation. Observing at higher frequency would reduce the required S/N, because of the smaller synthesised beam size.

Keeping σ_R below 20 μ as per-epoch at 1.6 GHz requires at least one calibrator within ~ 4 arcminutes, but preferably more than one (see Deller et al., 2013, ApJ, 770, 145). From the mJIVE-20 program (Deller & Middelberg, 2013, AJ, submitted) we know that the density of VLBI sources with a minimum peak flux density of 1 mJy is around 12 per square degree, and so on average we can expect

[†]SKA₁ will itself discover many new pulsars and many of these will be fainter than this example, which will limit the radius to which 10% parallax accuracies could be achieved to smaller than the 15 kpc used in this example.

a calibrator source brighter than 1 mJy within a separation of around 3 arcminutes. The baseline sensitivity of SKA₁-mid to a 25-m class VLBI antenna is around 0.1 mJy in 90 seconds, allowing a $10\text{-}\sigma$ calibration. Accordingly, with SKA₁ sensitivity, it should generally be possible to find sufficiently close calibrators to keep R below an acceptable limit. We note that in some cases, observing at a higher frequency will be preferable (e.g., targets in the Galactic plane where scattering is a concern) and in this case more distant calibrators will be acceptable, since the ionospheric error contribution will be reduced.

As already stated, the contribution of σ_F is difficult to estimate. The attainment of 15σ as parallax precision with single calibrators and existing VLBI arrays (where F was not thought to be a significant contribution to the error budget) implies that at least some calibrators are stable at the 10σ level. In this case, 2-3 calibrators should suffice. However, in general, it will be desirable to target at least 5-10 calibrator sources (necessarily spread over an area of around a half a square degree, assuming $> 1\text{-mJy}$ sources are to be used) to allow outlier sources to be identified and removed.

The total observing time for such a campaign would be $8 \times 3 = 24$ hours, assuming the slewing overheads to be negligible (which will be true in most cases; usually there will be a sufficiently bright source within the primary beam).

6.4.1 Prioritised requirements for SKA₁ to support operations in VLBI mode

Obtaining the best sensitivity for SKA₁-mid and SKA₁-survey as VLBI elements requiring phasing up as many elements as possible. Assuming the full core (~ 2 km in each case) is phased up, then the field of view will be limited to around 20–30 arcseconds depending on the frequency. Accordingly, a separate tied array beam will typically be required for the target and each calibrator source, since rarely will two sources be separated by less than 30 arcseconds. Since astrometric observations demand multiple calibrators, this places a requirement for multiple tied array beams for both SKA₁-mid and SKA₁-survey. Ideally, at least 10 independently steerable full-bandwidth tied array beams should be possible, although 4 beams at full bandwidth should be considered an absolute minimum, with scope for trading beams against bandwidth at larger numbers of beams (to a minimum of 10). An antenna mask should be possible for each beam separately, to allow e.g. a larger but slightly less sensitive beam to be formed by including only antennas within the central 500 m, for example. *This is requirement 1.*

The tied array data needs to be packaged in a format which is easily understood by the subsequent VLBI correlation system, including metadata such as timestamps. A format such as VDIF (<http://vlbi.org/vdif/>) would be acceptable. A selectable range of quantisation would be desirable, but 2 bit quantisation is likely to be by far the most commonly used mode. The frequency channelization should be at most moderate (channels should not be narrower than a few MHz), and flexibility in choosing the band edge frequencies and bandwidths is desirable (although compatibility can be maintained in the VLBI correlator if necessary). *This is requirement 2.*

For VLBI operations, the local oscillator and sampler needs to be driven by a sufficiently stable clock, to avoid adding phase noise to the sampled VLBI data stream. A hydrogen maser (or newer technology with equivalent or better accuracy) would be sufficient. The pulsar timing case already demands an accurate local time standard, and so this requirement should not necessitate any additional or improved equipment. *This is requirement 3.*

The VLBI-formatted data packets from each tied array beam then need to be transported to a VLBI correlator facility. Each tied array beam will have a data rate of 4 Gbps, assuming 500 MHz bandwidth, dual polarization and 2 bit sampling, and so if 10 beams are to be used then the total data

rate is 40 Gbps per SKA site (SKA₁-mid core, SKA₁-survey core, and any SKA₁-mid remote stations used). The location of the VLBI correlation facility is not yet known - reasonable options might be the current facilities for the Long Baseline Array [LBA] in Western Australia or the European VLBI Network [EVN] in the Netherlands - which means that this requirement cannot be fully specified. One possibility is that the data transport occurs in real time, in which case a fibre link with a capacity > 40 Gbps would be needed from each SKA₁ VLBI station to the correlation facility. Alternatively, data could be recorded locally to VLBI recorders and the media either physically transported or transferred after the observations at a lower data rate (as is more common for VLBI observations today). On a timescale of ~ 5 years it is to be expected that recording 40 Gbps would be possible on a small number of VLBI recorders, as current technology can already record 4 Gbps on a single recorder. *This is requirement 4*; its provision could be discussed with external partners in the VLBI network. In order to provide short VLBI baselines, it is desirable that a small number of the most distant SKA₁-mid remote stations can also produce tied array output in VLBI format. This is not a hard requirement, but is strongly desirable.

In order to aid calibration of the ionosphere, a GPS receiver should be placed at each SKA site. This is likely already needed for time maintenance for pulsar timing studies in any case. This is not a hard requirement, but is desirable.

Depending on the trade-offs between source flux density, primary field of view, scattering, calibrator availability and ionospheric conditions, a significant number of sources might be better observed at frequencies above the currently planned SKA₁-mid band 3. Astrometry of high mass star forming regions (6.7 GHz methanol line) is one such high-priority case. Thus, the VLBI capability of SKA₁ would be significantly enhanced by the addition of higher frequency bands to SKA₁-mid. This is not a hard requirement from the pulsar perspective; the majority of pulsar targets will be feasible at 1.6 GHz, but it is entirely possible that one or more individual sources of very high priority will not be feasible at this low frequency.

In order to maximize mutual visibility, it is desirable that observations down to low elevations are possible, even if performance is somewhat degraded (higher system temperature due to spillover and shadowing, for instance). As an example, the mutual visibility on the 9000 km long SKA₁-mid to ATCA baseline for a source at a declination of -30 degrees (which is indicative of the SKA₁-mid to east coast Australia baselines), the on-source time would increase by an hour (from 3.7 hours to 4.6 hours) if the SKA₁-mid elevation limit is 5 degrees instead of 15 degrees. This is not a hard requirement, but is desirable.

Finally, it is desirable that regular SKA₁ imaging is possible during VLBI observations (i.e., regular local imaging operations, simultaneous with formation of ~ 10 beams). In many cases, SKA₁-mid and -survey images of the field at arcsecond resolution will add considerable value to the observations. This is not a hard requirement, but is highly recommended for operational efficiency.

6.4.2 Operational considerations

VLBI observations obviously require other antennas to be available and committed to observing. The sparseness of available VLBI antennas with good mutual visibility coverage with SKA₁-mid and SKA₁-survey is a potential problem for (astrometric) VLBI observations. Antennas from the Australian LBA and the East Asian VLBI Network should be available, as well as the EVN for more northerly sources, however this still leaves a relatively small number of baselines, and in particular a lack of intermediate baselines (length 1000-8000 km) to SKA₁-mid. Table 1 shows baseline lengths from SKA₁-mid and SKA₁-survey to a number of existing and potential VLBI stations. As Table 1 shows, the commissioning of one or more antennas as part of the African VLBI Network would make a considerable improvement

to the uv coverage hole for SKA₁-mid.

Table 3: Baseline lengths to SKA₁-mid and SKA₁-survey. Existing telescopes are shown in plain font and 4 potential African VLBI network telescopes (out of ~ 25 possible) are shown in italics. The lack of existing baselines to SKA₁-mid of length 1000–7000 km is apparent.

Station	Baseline length (km) to SKA ₁ -mid	Baseline length (km) to SKA ₁ -survey
Hartebeesthoek (EVN)	800	8000
Noto (EVN)	7100	10700
Medicina (EVN)	7800	11000
Hobart (LBA)	9000	3300
Parkes (LBA)	9700	3000
ATCA (LBA)	9900	3100
Shanghai (EVN)	10600	6000
Zambia	1800	8500
Kenya	3600	8300
Cameroon	4000	10300
Nigeria	4600	10500
Ghana	4900	11000

In addition, it is extremely desirable that a coordinated approach to proposing and scheduling VLBI observations is made available, and this operational mode should be considered well in advance of actual SKA1 availability. A situation in which separate proposals needed to be sent to SKA₁-mid, SKA₁-survey and a multitude of other telescopes would be extremely inefficient, as would a system in which SKA₁-mid and SKA₁-survey needed to be scheduled separately to the rest of the VLBI network. Accordingly, the manner in which SKA₁resources can be operationally made available to VLBI operations should be investigated.

The same process that addresses joint proposal and scheduling concerns could also easily answer the question of where the VLBI data will be correlated. The computational requirements will be very limited, barely larger in scope than current-day VLBI instruments, even when considering the overhead of ~ 10 beams per SKA site. Accordingly, it will be possible to process the VLBI data with existing VLBI software correlator code running on modest computational resources, but where the correlation centre(s) are located (and how the data will be transported there) are questions which would have to be settled.

Finally, accurate terrestrial positions for the SKA₁VLBI reference points will be required. The presence of GPS receivers and local surveying will be useful for this, but geodetic VLBI observations remain the gold standard for obtaining the mm-level precision positions which are necessary for good astrometry. Accordingly, the SKA₁VLBI stations should participate occasionally in geodetic VLBI sessions, and this is another driver for a higher frequency capability (since geodetic VLBI observations are typically carried out between 2 and 10 GHz).

6.5 Searching Galactic plane for pulsars (SAK1-mid) Michael Kramer

to be added.

6.6 Searching off-Galactic plane for pulsars (SKA1-mid/SKA1-low) Mike Keith

to be added.

6.7 Searches with SKA-survey or input to (v) Simon Johnston

to be added.

6.8 Single pulse searching for both (i) and (ii) Maura McLaughlin

to be added.

6.9 Searching globular clusters and/or interesting point sources (e.g. Fermi) Scott Ransom & Paulo Freire

to be added.

6.10 Searching external Galaxies Joe Lazio & Jim Cordes

to be added.

6.11 Follow-up timing of search discoveries David Champion

to be added.

6.12 Timing of interesting sources, e.g. glitches, intermittents etc Patrick Weltevrede

to be added.

6.13 High-precision timing for GW sources Gemma Janssen & Dick Manchester

Appendix B — Contributor Contact Details

to be added