

Pulsars - Enabling the Gravitation KSP of the SKA

M. Kramer^{1,2} and R. Smits^{2*}

¹ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

² University of Manchester, Jodrell Bank Centre for Astrophysics, Turing Building, Oxford Road, Manchester M13 9PL, UK**

Abstract. In a SKADS funded study we have investigated the impact of the SKA design on the prospects to achieve the scientific goals outlined in the Key Science Project on testing gravity using pulsars and black holes. We find that the computing requirements are severe and differ for individual configurations. As today, pulsar science will continue to be limited by the available computing power. However, we present methods to achieve the science goals but point out that every reduction in sensitivity (i.e. A/T) will severely risk the success of this KSP. Unlike other observing modes, we cannot trade collecting area for field-of-view or integration time. This important feature is potentially the most important factor to be considered for pulsar studies when designing the SKA.

1. Introduction

The Square-Kilometre Array (SKA) will be a flexible, multi-purpose observatory that will be able to serve the whole astronomy community as *the* premier imaging and surveying instrument. For this reason, the design needs to ensure that the SKA will be a versatile instrument rather than a fine-tuned experiment. Nevertheless five Key Science Projects (KSPs) have been selected by the international community which must drive essential parts of the design in order to ensure that the science goals can be achieved. It is also important that the design includes natural upgrade paths, so that eventually only the available computing power is constraining the capabilities of the SKA. This is a situation that is well familiar to pulsar astronomers today, who encounter continuously the boundaries of data processing capabilities. This situation will therefore not change, but a number of design decisions should be made if the Key Science related to pulsars and gravitational physics are to be successful. In this contribution, we summarize these design requirements. The presented features are the result of a SKADS funded study in collaboration with international partners, including Ben Stappers, Duncan Lorimer, Jim Cordes and Andrew Faulkner. Related work was done by Aris Karastergiou and Tobia Carozzi. Further details can be found in Smits et al. (2009) and Smits et al. (in prep).

2. Was Einstein right?

In order to motivate the technical requirements, we will briefly summarize the science case which is based on the KSP chapter by Kramer et al. (2004) and the general pulsar chapter by Cordes et al. (2004) of the SKA Science Book edited by Carilli & Rawlings.

* Present address: ASTRON, P.O. Box 2, Schattenberg 17990 AA Dwingeloo, 433 TA Zwiggelte, The Netherlands

** This work was supported by the European Commission Framework Program 6, Project SKADS, Square Kilometre Array Design Studies, contract no 011938.

2.1. Astrometry & Binary Pulsar Tests

Astrometry and in particular pulsar timing of millisecond pulsars should allow us to measure proper motions with a precision of about 100 nas/yr and parallaxes of the order of 10–20 μas , i.e. distances of up to 50–100 kpc. Astrometric measurements are not only useful to potentially trace the pulsars' movement in the gravitational field of the Galaxy or globular clusters, but parallax measurements are also needed to correct for relative acceleration between the pulsar and the solar system barycentre which affects the observed period derivatives (for spin and orbit) of relativistic binaries (e.g. Lorimer & Kramer 2005). This is essential if radiative aspects of a theory of gravity are to be tested. The prospects for tests when using the SKA to study the Double Pulsar (Lyne et al. 2004) are discussed in detail by Kramer & Wex (2009). Essentially, with the SKA we expect an improvement in timing precision by a factor of ~ 100 , leading to precision tests of general relativity that even exceed the precise tests done in the solar system's weak field regime today. We expect to measure the moment-of-inertia of pulsar A in the Double Pulsar system and to detect the previously unmeasured relativistic deformation of the orbit, providing an additional test for general relativity and other theories of gravity (Kramer & Wex 2009).

2.2. Alternative theories & Black hole properties

What makes a binary pulsar with a black-hole companion so interesting is that it has the potential of providing a superb new probe of relativistic gravity. As pointed out by Damour & Esposito-Farese (1998), the discriminating power of this probe might supersede all its present and foreseeable competitors. The reason lies in the fact that such a system would be very sensitive to strong gravitational self-field effects, making it for instance an excellent probe for tensor-scalar theories. Moreover, Wex & Kopeikin (1999) showed that the measurement of classical and relativistic spin-orbit coupling in a pulsar-black hole binary, in principle, allows us to determine the spin and the quadrupole moment of the black hole. This would test the “cos-

mic censorship conjecture” and the “no-hair theorem”. While Wex & Kopeikin (1999) pointed out that with current telescopes such an experiment would be almost impossible to perform (with the possible exception of pulsars about the Galactic Centre black hole), Kramer et al. (2004) demonstrated that the SKA sensitivity should be sufficient. Indeed, this experiment benefits from the SKA sensitivity in multiple ways. On one hand, it provides the required timing precision but it also allows to perform a Galactic Census of pulsars which should eventually deliver the sample of pulsars with a black hole companion.

2.3. Gravitational wave astronomy

While pulsars already provide the (so far, only) indirect evidence for the existence of gravitational waves (GW), they can also be used to detect and study them directly. Timing residuals are affected by a passing GW wave as each pulsar and the Earth can be considered as free masses, whose positions respond to changes in the space-time metric (e.g., Sazhin 1978, Detweiler 1979). For our purposes here, it is important to note that a perturbation of the space-time metric by a GW would lead to timing residuals of the order of $\sigma \sim h_c(f) \times T^{1.5}$ where $h_c(f)$ describes the characteristic amplitude of the GW per unit logarithmic interval of frequency, and T is the total observing time. Importantly, the sensitivity of GW detection using pulsars scales directly with the achieved timing precision.

Despite the apparent simplicity of the experiment, the timing precision required for the detection of GWs is very much at the limit of what is technically possible today. Only a stochastic background (rather than a single GW source) is likely to be detected today, as it would produce a correlated quadrupole signature among a network or “array” of timed millisecond pulsars on the sky. This facilitates the recognition of a signal originating, for instance, from the orbital motion of binary super-massive black holes in the early Universe. A variety of sources is expected to emit GWs in the nHz-frequency range detectable by a “pulsar timing array” (PTA). Current PTA experiments will have a good chance of detecting GWs, but only with the SKA can a detection be guaranteed and the properties of GWs (e.g. spectrum of a stochastic background, polarisation) be studied. Due to the increased timing precision achievable with the SKA and a large number of suitable sources discovered, a PTA with the SKA would achieve a sensitivity that even exceeds that of LISA (Kramer et al. 2004).

Different types of signals can be detected and studied with the SKA, i.e. stochastic, burst, and periodic signals. A stochastic gravitational wave background should arise from a variety of sources. Cosmological sources include inflation, string cosmology, cosmic strings and phase transitions (Kramer et al. 2004). The expected correlation signal among the PTA pulsars can be measured with very high precision, so that the polarisation properties of GW can be probed and compared with the predictions of GR (Lee et al. 2008). The SKA will also detect single sources of GW emission, such as the periodic GW signal of a binary super-massive black hole in the centre of a nearby Galaxy. Here, the SKA and LISA will be complementary: On one hand, both instruments can see sim-

ilar super-massive binary black hole systems but at different evolutionary stages, i.e. LISA will pick up SKA sources close to their merger. On the other hand, the SKA can observe GW emission of systems that are typically more massive than the LISA sources. We should also be able to discover compact pulsar binaries that emit in radio simultaneously with continuous GW emission detectable with LISA. Moreover, the black holes studied with the SKA (stellar and the super-massive black hole in the Galactic Centre) will be complementary and partly overlapping with the population of black holes to be studied with Advanced LIGO and LISA.

Besides these guaranteed GW signals, the SKA should also be able to measure burst emission from, for instance, explosive events. In the GW window, pulsar glitches may also produce observable events, where the SKA would help to pinpoint the exact time of such an event. Fundamental properties of gravitation can also be tested: if GWs are governed by a massive field propagation law, GWs with frequency less than $m_g c^2/h$ cannot propagate. Hence detecting GWs with a 5 to 15 year period (ie. the frequency range covered by PTAs) reduces the upper bound on gravitation mass m_g by at least one order of magnitude (Lee et al., in prep.).

3. One KSP, three components & three challenges

In order to achieve the science goals outlined above, it is necessary to conduct several parts which are complementary in terms of technical requirements but depend on each other in terms of science.

3.1. The experiments

The first essential component is a survey for pulsars that will provide a Galactic census. Depending on the low-end of the luminosity function and the collecting area that is deployable, we should be able to detect a large fraction of every pulsar that directs its beam towards Earth during its rotation. Current estimates suggest that this should be about 20 000 to 30 000 pulsars. In reality, the available collecting area does not only depend on the “raw” square meter of receiving elements on the ground, but also on their spatial configuration and the computing power needed to combine the signals of those. We will discuss this point in more detail later on. This survey for pulsars will be conducted at a frequency between 1-3 GHz, depending on Galactic latitude and frequency overlap of different antenna types. The outcome of this survey will be the sample of accurate millisecond pulsar clocks that we will time as a Pulsar Timing Array to study the nHz-gravitational wave sky. Furthermore, we should uncover the rare pulsar-black hole (BH) systems where a pulsar orbits a stellar BH in the Galactic plane or an intermediate BH in a globular cluster which can be searched deeper.

Finding the gems of pulsars that promise the science that we are after, is only the first move in the game. No science can be extracted from any of the discovered sources, unless they are “timed”. Timing is the continuous counting of neutron star rotations by means of measuring their pulse arrival times with our telescope on Earth. Any deviation from the expected

POS (SKADS 2009) 008

arrival time reveals potentially unmodeled effects, resulting in the detection and measurement of relativistic effects, the impact of a gravitational wave or other physical phenomena. This assumes, of course, that the pulsar clock itself is stable and rotates smoothly and accurately to a precision that is needed to measure particular effects.

Not all pulsars rotate smoothly. Young and middle-aged pulsars are less level-headed than so-called millisecond pulsars (MSPs) which are older, recycled and much more stable than non-recycled pulsars. Even though they live longer, we expect millisecond pulsars to be outnumbered by their discovered non-recycled young cousins, as their formation require a companion that donates mass for an angular momentum transfer. As we seem to start seeing some variation in the spin-down stability of millisecond pulsars, we require regular timing observations in order to identify good and bad “timers”. Such monitoring is also essential to determine whether a pulsar has a companion or not, what type it is and which orbital parameters the system has. Once the interesting sources have been identified, they will need to be timed with the full array to achieve the best possible precision. The timing therefore consists of two parts: a medium to low-precision timing program to follow-up the 20,000 to 30,000 discovered sources to identify the “goodies”, and a second high-precision timing part of, probably, 100 to 300 sources or so, to extract the science.

Usually, timing a pulsar requires the observation of the source for one full year, so that the pulsar is “viewed” from all parts of the Earth orbit about the Sun. By this, the position of the pulsar can be precisely determined and correlations in the parameter estimation between pulsar spin-down and astrometry can be broken. The timing program may be sped up by obtaining the position earlier, in less than a year, by using the SKA as an interferometer. The VLBI observations are also important for the most precise measurements of orbital decay due to gravitational wave emission or the spin-down behaviour. We therefore want to perform VLBI measurements for all sources in the “high-precision timing programme”.

While the above components describe the survey of the sky for the vast majority of its area, a small piece of empty sky exists that needs to be treated especially: the Galactic Centre (GC). The GC region is not only interesting because a large stellar density suggests the occurrence of BH-MSP systems, but we will also discover pulsars orbiting the super-massive BH in the GC. Tracing the movement and rotation of pulsars in their orbit provides in principle an easier measurement of the BH spin and quadrupole moment as those quantities scale with the square and cube of the masses, respectively. Such a measurement would also extend the mass range of studied BH properties, but the problem is that the interstellar medium in the GC causes severe interstellar scattering due to multi-path propagation. This can only be combated by observing at higher radio frequencies, i.e. frequencies between 10 and 15 GHz should be sufficient. This implies a small beam, and hence a slow survey speed. However, only a small part of the sky needs to be searched and the rewards are huge, making this an integral component of the KSP.

3.2. Three potential problems...

Can we survey the sky? Pulsar surveys are technically challenging in many respects. Firstly, the signals of the individual antennas need to be combined computationally, whereas the basic field-of-view (FoV) of a single (or station) element is pixelized to cover the sky without a gap. The number of pixels needed to cover the FoV can be estimated to be $N_{\text{pix}} \sim (b_{\text{max}}/D)^2 \sim 10^4 - 10^9$, where b_{max} is the maximum baseline between two telescopes and D the diameter of the receiving element (see SKA Memo 97 by Jim Cordes). This alone translates into a requirement of peta-ops per second. However, this beam-formed data should also cover a wide receiving bandwidth that needs to be sampled with fine frequency resolution at fast rates (i.e. about 500 kHz at 100 μ s). Finally, it is likely that the survey data rate is so high (at least initially until technology catches up) that data need to be processed on-the-fly. This problem depends on a number of factors, including the array configuration. It is however already immediately clear that every reduction in collecting area (or A/T) will automatically imply even higher processing cost when we try to find the most relativistic and accelerated system. As processing requirements scale roughly with t^2 where t is the observing time, and since $t \propto (T/A)^2$, any cut in (A/T) by a factor of k increases the needed computing power by a k^4 .

Can we time all pulsars? As underlined above, the follow-up timing of pulsars is of equal importance to the search. In fact, searching and timing should always be considered together, in particular as follow-up timing typically requires much more observing time than the survey itself. This is certainly true for current search projects, implying that it will be nearly impossible to undertake the ambitious task to time 20,000 or more pulsars every two weeks which is the preferred time interval, initially. However, not all pulsars will be discovered at the same time and, most importantly, the array character of the SKA comes to the rescue. Here, aperture arrays are particularly valuable, but also sub-arraying dishes will enormously help.

Does the sensitivity of the SKA really translate into timing precision? In our experiments using the SKA, we can expect to gain in several ways, enabling science at a complete new level that is not simply a continuation of what has been done before, but what represents a new quality of pulsar science. Firstly, the SKA should provide us with many more sources that can be used in a PTA experiment, allowing us to sample the expected correlation curve between MSP timing residuals finely. At the same time, we expect the timing precision to increase by one to two orders of magnitude. With wide frequency coverage and the ability to instantaneously measure dispersion and rotation measures for a large number of sources regularly, we will be able to remove interstellar weather effects effectively. Finally, with many beams on the sky, we can afford to also time our pulsars more frequently, again increasing the timing performance. Among these many improvements, the one about increasing timing precision deserves a further look. It lies at the centre of most of our performance predictions and basically assumes that timing precision scales with telescope sensitivity. However, a number of effects will need to be controlled to achieve this:

calibration, interstellar weather, pulse jitter and long-term timing noise.

4. Surveying the pulsar sky

It is extremely important to realize that pulsar surveys are very different from other, e.g. continuum surveys: *it is practically impossible to trade sensitivity (read, collecting area) for integration time or FoV!* The reason is that we are looking for time-variable signals which are ideally modulated by Doppler-shifts due to some orbital motion. The faster this orbital motion and its change, the more interesting are the pulsars – and the more difficult it is to find them with long integration times. In the worst case, the periodic signal is smeared out and diminished, so that it becomes undetectable. As indicated earlier, one can try to compensate for the orbital motion with a search in acceleration, but this becomes computationally expensive as the number of trials scales with the square of the number of samples. A cut in sensitivity essentially deselects against the most interesting sources. A detailed analysis on the impact of the discoverable population of binary pulsars will be presented elsewhere (Smits et al., in prep.). A similar severe degradation in science capabilities due to reduced sensitivities also occurs for pulsar timing, as discussed further below. Concentrating firstly on the computational and data transport requirements to pixelize the FoV for searching, those results were presented by Smits et al. (2009).

4.1. SKA configuration

As in SKA Memo 100, we assume the SKA to consist of a sparse aperture array of tiled dipoles in the frequency range of 70 to 500 MHz and above 500 MHz it consists of one of the following three implementations:

- A 3 000 15-m dishes with a single-pixel feed, a sensitivity of 0.6 SKA units, $T_{\text{sys}}=30$ K and 70% efficiency covering the frequency range of 500 MHz to 10 GHz.
- B 2 000 15-m dishes with phased array feeds from 500 MHz to 1.5 GHz, a sensitivity of 0.35 SKA units, a field of view (FoV) of 20 deg^2 , $T_{\text{sys}}=35$ K and 70% efficiency and a single-pixel feed from 1.5 to 10 GHz, with $T_{\text{sys}}=30$ K.
- C A combination of a dense aperture array (AA) with a FoV of 250 deg^2 , a sensitivity of 0.5 SKA units, covering the frequency range of 500 to 800 MHz and 2 400 15-m dishes with a single-pixel feed covering the frequency range of 800 MHz to 10 GHz, a sensitivity of 0.5 SKA units, $T_{\text{sys}}=30$ K and 70% efficiency.

We assume that 20% of the elements will be placed within a 1 km radius and 50% within a 5 km radius. The sparse aperture array below 500 MHz will be ignored, as we will not consider using it for pulsar surveys or timing in this paper. We will refer to the dishes from configurations A and B (which have a lower frequency limit of 500 MHz) as 500-MHz dishes and we will refer to the dishes from configuration C (which have a lower frequency limit of 800 MHz) as 800-MHz dishes.

The FoV of a survey has a maximum that depends on the characteristics of the elements used. In the case of circular

dishes with a one pixel receiver, the FoV can be approximated as $(\lambda/D_{\text{dish}})^2$ steradians, where λ is the wavelength at which the survey is undertaken and D_{dish} is the dish diameter. However, by placing a phased array feed at the focal point of the dish, the FoV can be extended by a factor which we denote ξ . The total FoV then no longer depends on λ . In the case of the Aperture Array (AA), the FoV of the elements is about half the sky. The actual size of the FoV that can be obtained will be limited by the available computational resources. This is because the signals from the elements will be combined coherently, resulting in ‘pencil beams’, the size of which scales with $1/D_{\text{el}}^2$, where D_{el} is the distance between the furthest elements that are used in an observation. To obtain a large synthesised FoV it is therefore necessary to restrict the beam forming to using only the elements in the core of the telescope (so that $D_{\text{el}} = D_{\text{core}}$), which will, however, reduce the sensitivity.

4.2. Survey simulations

The results of survey simulation presented by Smits et al. (2009) and summarized here, do not yet consider the impact on the population of discovered binary pulsars. In general, the optimal frequency for a pulsar survey depends on Galactic latitude. We find that the optimal survey frequency decreases with lower sensitivity. The reason for this is that for lower sensitivity, the average distance to the detected pulsar is smaller. Thus, frequency dependent propagation effects are smaller, allowing lower frequencies to be used. It is surprising that the optimal centre frequencies for normal and millisecond pulsars are similar, since scattering effects on the pulsar profile have a larger impact on the detection of millisecond pulsars than on the detection of normal pulsars. Thus it can be expected that in the Galactic plane, higher observation frequencies would favour the detection of millisecond pulsars. We performed four simulations corresponding to possible pulsar surveys. For the dishes they were performed at the optimal centre frequencies, taking the frequency limits into account. The frequency range of the AA was held fixed at 500 to 800 MHz. The simulations are:

- An all-sky survey with the AA (including the Galactic plane) and a Galactic plane survey with the 800-MHz dishes. This shows the fraction of pulsars that can be found with configuration C.
- An all-sky survey with just the AA. This shows the fraction of pulsars that will be found with configuration C, if we do not include the dishes in the survey.
- An all-sky survey with the 500-MHz dishes. This shows the fraction of pulsars that can be found with configurations A or B.
- A Galactic plane survey with the 500-MHz dishes. This shows the fraction of pulsars that will be found with configuration A or B, without performing a survey outside the Galactic plane.

We assumed a bandwidth of 500 MHz for the dishes and 300 MHz for the AA and an integration time of 1800 s. The simulations were limited to isolated pulsars only. Fig. 1 shows

POS (SKADS 2009) 008

the fraction of normal and millisecond pulsars that were detected in the simulation of 4 different surveys as a function of sensitivity. Only 1.5% of the pulsars could not be detected due to the declination of the SKA of -30° . As described below, we expect that initially only the 1-km core of the SKA can be used for the pulsar survey. This leads to a sensitivity of 0.1 SKA units, which is similar to having a fully steerable Arecibo-class telescope in the southern hemisphere. Performing an all-sky survey with only the 1-km core of the SKA, will detect about 14 000 normal pulsars out of a possible 30 000 and about 6 000 millisecond pulsars out of a possible 30 000, regardless of the SKA implementation. The total observation time of the survey, however, does depend on the implementation. Using single-pixel feed dishes to perform an all-sky survey or a survey of just the Galactic plane will take 600 days and 30 days, respectively. With phased array feed dishes or an AA the total survey time can be much less, depending on the available computation power.




Fig. 1: Fraction of normal (top) and millisecond (bottom) pulsars detected from different pulsar survey simulations as a function of sensitivity of the SKA. For both the 500-MHz and 800-MHz dishes, the surveys were performed at the optimal centre frequency, taking the frequency limits into account. The frequency range of the AA was 500 to 800 MHz.

5. Timing the discoveries

Here we summarize implications of particular design choices on the feasibility of timing all the discovered pulsars. This part of the experiment consists of two parts: the relatively low-precision timing to identify pulsars required for gravitational wave detection and tests of general relativity, and the high-precision part to do exactly that. It is important to note that for the latter, *again we cannot trade sensitivity for observing time!* The reason is that we need to measure a precise pulse time-of-arrival for sufficiently small integration times, since we need to resolve a potential short orbit. If that is not possible due to a lack of sensitivity, the fruits of having discovered exciting short binary pulsars cannot be harvested and the KSP cannot be completed. A quantitative treatment of the problem is in preparation (Smits et al., in prep.).

In the low-precision part of the timing project, pulsars need to be timed ideally once every two weeks for about 6 months to a signal-to-noise ratio of about 9 (which is only a nominal target to characterize the new pulsars). The minimum timing duration is also restricted by the pulsar stabilization time, which is the time it takes to obtain a stable integrated profile. Assuming that they will be timed using dishes, the total duration of timing all newly discovered pulsars just once, depends on how the observations are performed, i.e. at which points on the sky the telescope is pointed. To estimate the total duration of this timing process, we therefore suggest the following scheme that minimizes the average observation time per pointing, by grouping dim pulsars together in one FoV. Let ρ_{fov} be the angle of the FoV.

1. All pulsars are placed on a long list, in order of increasing brightness.
2. All pulsars within an angle of ρ_{fov} from the first pulsar of the list (the dimmest pulsar) are placed on a short list.
3. We now consider the dimmest pulsar on the short list. If possible, we restrict the pointing of the FoV to include this pulsar and we remove this pulsar from both the short list and the long list. If the pointing of the FoV is already too restricted to include this pulsar, it is removed from the short list only.
4. Step 3 is repeated until the short list is empty. The resulting pointing is stored.
5. Step 2 through 4 are repeated until the long list is empty.

This timing optimization method saves about 30 to 50% of observation time over a simple grid approach. To compare timing performances, we used this scheme to estimate the time to obtain a single timing point of 14 000 pulsars from an all-sky survey to a signal-to-noise ratio of 9 for each of the SKA-configurations, assuming usage of the full collecting area and a pulsar stabilization time of about 200 pulses. This leads to observation times of 6 days for configuration A and 1.5 days for both configurations B and C. For configuration C the AA was used to time all the pulsars that were detected by the AA in the survey. These estimates do not include densely spaced ‘resolving’ observations that are required to obtain the first phase coherent timing solution.

To perform strong field tests of General Relativity and gravitational wave detection, as defined in the KSP, it is essential to time a large number of specific pulsars to signal-to-noise ratios of up to 100 on a regular basis. To estimate how much observation time is required for such high-precision timing, we have to distinguish between timing millisecond pulsars (for the Pulsar Timing Array) and timing pulsars in a binary with a neutron star or a black hole companion. We calculated the observation time for timing the 250 millisecond pulsars with the best signal-to-noise ratio's for the three configurations. We assumed that for timing purposes, the nearly full collecting area (out to several hundreds of kilometres from the core) of the SKA can be used. As a conservative estimate (subject to current studies), we timed every pulsar for at least 5 minutes to ensure a stable pulsar profile to minimize the error in the time of arrival. With the single-pixel feed 15-m dishes this takes about 20 hours, the 15-m dishes with phased array feeds take about 15 hours and the AA only takes about 6 hours. However, this assumes that the polarisation purity after calibration of the AA is similar to that of the dishes, which might not be the case. To estimate the maximum required observation time for timing pulsars in a binary with a neutron star or black hole, we will assume a (somewhat optimistic) number of 200 such binaries that are potentially detectable in the Galaxy. We further assume that the characteristics of the pulsars in these binaries are similar to isolated pulsars. This leads to a detection of 90 compact binary pulsars. With the single-pixel feed 15-m dishes it takes about 2 days to time these binary pulsars, the 15-m dishes with phased array feeds take about 1.5 days and the AA only takes about 8 hours (once again, assuming similar polarisation purity as the dishes). However, timing only the brightest 80% of the binary pulsars, would take about half the observation time. It should be noted that these are only initial estimates to give an indication of the maximum required observation time for high-precision timing and to compare the performance of the different configurations. A paper containing a more detailed study on high-precision timing is in preparation (Smits et al. in prep.).

6. Computational requirements

Pulsar science with the SKA will always, as today, be limited by the available computing power. For a pulsar survey, we require the coherent addition of the signals from the elements in the core of the SKA and to form sufficient pencil beams to create the required FoV. These pencil beams will produce a large amount of data which will need to be analyzed. In the following, we summarize the required computation power and expected data rates for a configuration of 3000 15-m dishes with a bandwidth of 500 MHz and 2 polarizations and for the AA with a frequency range from 500 to 800 MHz and 2 polarizations. Details can be found again in Smits et al. (2009).

6.1. Beam forming

We estimate the number of operations per second (ops), N_{osb} , to fill the entire FoV of a dish with pencil beams as^a:

$$N_{\text{osb}} = F_c N_{\text{dish}} N_{\text{pol}} B \xi \left(\frac{D_{\text{core}}}{D_{\text{dish}}} \right)^2, \quad (1)$$

where F_c is the fraction of dishes inside the core, N_{dish} is the total number of dishes, N_{pol} is the number of polarizations, B is the bandwidth, D_{dish} is the diameter of the dishes, D_{core} is the diameter of the core and ξ is the factor by which the phased array feed extends the FoV. To fill the FoV of single-pixel 15-m dishes in a 1-km core requires 4400 pencil beam. The number of operations per second to beam form a 1-km core containing 20% of 2400 single-pixel 15-m dishes then becomes 2.2×10^{15} . The number of operations for beam forming the full FoV of 15-m dishes with a phased array feed are simply a factor ξ higher, corresponding to the increase in FoV. For the SKA implementation B, this corresponds to $\xi = 31$. The number of operations for beam forming can be reduced by beam forming in two stages, if we assume that the dishes in the core are positioned such that they can form sub-arrays that all have the same size. In the first stage, beam forming of the full FoV is performed for each sub-array. In the second stage the final pencil beams are formed by coherently adding the corresponding beams of each sub-array formed in the first stage. For the same parameters as above, the number of operations per second becomes 2.0×10^{14} . For the beam forming of the AA, the calculations are slightly different. Initial beam forming will be performed at the stations themselves, leading to beams equivalent to those of a 60-m dish for each station. Demanding a total FoV of 3 deg² and assuming a total collecting area of 500 000 m², the number of operations becomes:

$$N_{\text{osb}} = F_c N_{\text{dish}} N_{\text{pol}} B \cdot 3 \left(\frac{\pi}{180c} \right)^2 D_{\text{core}}^2 v_{\text{max}}^2, \quad (2)$$

where c is the speed of light and $N_{\text{dish}} \approx 177$. For a 1-km core, this leads to a required computation power of 1.4×10^{14} ops. Fig. 2 shows the number of operations per second for the beam forming as a function of core diameter for the 15-m dishes with single-pixel feed and the AA.

As an alternative to beam forming by coherently adding the signals from the dishes up to a certain core size, it is also possible to perform the beam forming by incoherently adding the signals from sub-arrays. This process is similar to beam forming in two stages as mentioned above, except that in the second stage the beams that were formed in the first stage are added incoherently. This leads to a much larger beam size, which reduces the required computation power for beam forming significantly. It also reduces the total data rate and the required computation power for the data analysis. The drawback is that it reduces the sensitivity of the telescope by a factor of $\sqrt{N_{\text{sa}}}$,

^a Note that for a phased array feed this function is actually frequency dependent, because the total FoV of the dish then becomes constant, yet the FoV of the pencil beams is given by $(\lambda/D_{\text{tel}})^2$. However, for the values of bandwidth and observation frequency used in this paper, Eq. 1 is accurate.




Fig. 2: The number of operations per second required to perform the beam forming for the 15-m dishes and the AA. It is assumed that there are 2400 15-m dishes and the bandwidth is 500 MHz. The FoV of 0.64 deg^2 corresponds to filling the natural FoV of the dishes ($\xi = 1$) at a frequency of 1.4 GHz. For the AA we assume a FoV of 3 deg^2 and a total collecting area of $500\,000 \text{ m}^2$. The frequency range is 0.5 to 0.8 GHz. In all cases the number of polarizations is 2. The thick black line corresponds to the beam forming of the 15-m dishes in one stage. The thin black line corresponds to the beam forming of the 15-m dishes in two stages. The striped/dotted line corresponds to the beam forming of the AA.

where N_{sa} is the number of sub-arrays, which will be several hundred. However, this can be partially compensated as this method allows utilizing all elements that are placed in sub-arrays of the same size, which can possibly be as much as 5 times the collecting area of the core of the SKA.

6.2. Data analysis

There are two factors which impact the total data volume to be analyzed. Firstly, the FoV will be split up in many pencil beams, each of which will need to be searched for pulsar signals. Secondly, the SKA will be able to see the majority of the sky, all of which we want to search for pulsars and pulsar binaries. There are two ways to achieve this. The first option is to analyze the data as it is received, immediately discarding the raw data after analysis. This requires the analysis to take place in real time. The second option is to store all the data from part of the survey and analyze them at any pace that we see fit before continuing with the next part of the survey. Both approaches pose serious technical challenges, which we studied in detail in Smits et al. (2009). The results are shown in Fig. 3, where we present the data rate from a pulsar survey for the 15-m dishes and the AA as a function of core diameter, assuming $t_{\text{samp}}=100 \mu\text{s}$, $\text{DM}_{\text{max}}=1000 \text{ cm}^{-3} \text{ pc}$ for the dishes and $\text{DM}_{\text{max}}=500 \text{ cm}^{-3} \text{ pc}$ for the AA, $N_{\text{pol}}=1$ (sum of 2 polarizations) and 2 bit digitization. The AA was assumed to operate on the full frequency range of 0.5 to 0.8 GHz. The frequency range for the dishes was assumed to be 1 to 1.5 GHz. For a 1-km core, the data rate for the 15-m dishes with single-pixel feed is 4.7×10^{11} bytes per second and for the AA it is 1.6×10^{12} bytes

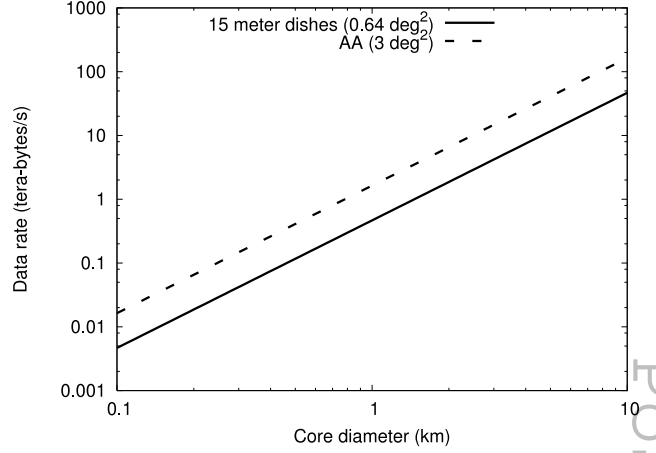


Fig. 3: Data rate from pulsar surveys using the 15-m dishes or the AA, as a function of core diameter.




Fig. 4: Operations per second required to perform a real time analysis of pulsar surveys using the 15-m dishes or the AA, as a function of core diameter.

per second. Both values scale linearly with the FoV. We can estimate the number of operations to search one ‘pencil beam’ for accelerated periodic sources as one Fourier Transform of all the samples in the observation, noting that the number of operations for the analysis scales as $N_{\text{samp}}^3 \log(N_{\text{samp}})$. As mentioned, this means that increasing the observation length is computationally very expensive. Fig. 4 shows the number of operations per second required to perform a real time analysis of a pulsar survey with the 15-m dishes and the AA as a function of core diameter, assuming 100 trial accelerations, a sampling time of $100 \mu\text{s}$ and an observation time of 1800 s. For a 1-km core, the required computation power becomes 1.2×10^{16} and 4.0×10^{16} ops for the 15-m dishes and the AA, respectively. These values are lower estimates of the actual number of operations, as there will be significant contributions from dedispersing, harmonic folding and possibly other processes that will contribute to the total number of operations. These additional processes are all on the order of N_{samp} and we expect their total contribution to be similar to one Fourier Transform at most.

POS (SKADS 2009) 008

7. The pulsar clock

The key to the success of the KSP is the increase of timing precision with SKA sensitivity. We assume here that post-detection calibration will be sufficient to produce pulse profiles that do not suffer instrumental effects. SKADS funded studies on this subject have been presented Carozzi & Woan (2009) and Karastergiou & Carozzi (in prep.). Remaining contributions to “timing noise” could arise from the interstellar medium, template matching algorithms, pulse jitter, and rotational instabilities. We should be able to correct for the interstellar medium using techniques presented by Stinebring and collaborators (e.g. Hemberger & Stinebring 2008), while recent work by Liu et al. (in prep.) verifies the validity of our TOA measurements procedures. Results presented by Verbiest et al. (2009) suggest that a sufficient number of MSPs exists that do not show any low-frequency noise that could limit the use of these pulsars for decade-long gravitational-wave detection effort. Even more exciting, Lyne et al. (submitted) argue for a magnetospheric origin of timing noise in young pulsars and put forward a method that may, with SKA sensitivity, produce the perfect pulsar clock.

8. Conclusions

The capabilities to find and exploit pulsars for scientific questions addressed by the KSP will continue to be limited by computing resources. While the computing power expected to be available in the early years of SKA operation should be sufficient to achieve the KSP goals, this offers of course the opportunity to vastly improve on the results over the lifetime of the SKA, in particular if a significant fraction of the collecting area is provided by AAs. However, this is only true if the collecting area is sufficiently large in the first place. Any reduction in A/T will lead directly (and with a power law dependence) to a decrease in science capabilities. Indeed, if the reduction is too large (say, 50%) the success is highly threatened.

References

- Cordes, J.M., et al. 2004, *New Astr. Rev.*, 48, 1413
 Carozzi, T.D., & Woan, G. 2009, *MNRAS*, 395, 1558
 Damour, T., & Esposito-Farese, G. 1998, *Phys. Rev. D.*, 58, 2001
 Detweiler, S. 1979, *ApJ*, 234, 1100
 Hemberger, D.A., & Stinebring, D.R. 2008, *ApJ*, 674, 37
 Kramer, M., et al. 2004, *New Astr. Rev.*, 48, 993
 Kramer, M., & Wex, N., 2009, *Class. Quant. Grav.*, 26, 3001
 Lee, K.J., et al. 2008, *ApJ*, 685, 1304
 Lorimer, D.R., & Kramer, M. 2005, *Handbook of Pulsar Astronomy*, CUP, Cambridge
 Lyne, A.G., et al. 2004, *Science*, 303, 1153
 Sazhin, M.V. 1978, *Sov. Astr.*, 22, 36
 Smits, R., et al. 2009, *A&A*, 493, 1161
 Verbiest, J.P., et al. 2009, *MNRAS*, 400, 951
 Wex, N., & Kopeikin, S. 1999, *ApJ*, 514, 388