

---

**Supplementary information**

---

**Discovery of a radio-emitting neutron star  
with an ultra-long spin period of 76 s**

---

In the format provided by the  
authors and unedited

## **Supplementary Information**

## MeerKAT description and configuration for transient searches

The Meer(more) Karoo Array Telescope (MeerKAT) [53–55] operated by the South African Radio Astronomy Observatory (SARAO) comprises 64, 13.5-m antennas distributed over 8-km in the Karoo region in South Africa. 40 of these dishes are concentrated in the inner  $\sim 1$ -km core. The MeerTRAP project is a commensal programme to search for pulsars and fast transients whilst piggy-backing on the large survey programmes at MeerKAT. The ThunderKAT programme is the image-plane transients search programme at MeerKAT and operates in synergy with MeerTRAP for blind transient searches. In the observations presented in this work, MeerKAT operated at a centre frequency of 1284 MHz with a usable bandwidth of  $\sim 770$  MHz in the L-band observations and at a centre frequency of 798 MHz with a usable bandwidth of  $\sim 435$  MHz in the UHF-band observations. MeerKAT observes simultaneously in incoherent and coherent modes using the MeerTRAP backend. In the coherent mode, the voltages from the inner 40 dishes of the  $\sim 1$ -km core of the array are coherently combined to form up to 780 beams on sky with an aggregate field-of-view of  $\sim 0.4$  deg<sup>-2</sup>. In the incoherent mode the intensities of all available (up to the maximum 64) MeerKAT dishes are added to create a less sensitive but much wider field-of-view of  $\sim 1.3$  deg<sup>2</sup>.

The MeerTRAP backend is the association of two systems: the Filterbank and Beamforming User Supplied Equipment (FBFUSE), a many-beam beamformer that was designed and developed at the Max-Planck-Institut für Radioastronomie in Bonn [56, 57], and the Transient User Supplied Equipment (TUSE), a real-time transient detection instrument developed by the MeerTRAP team at the University of Manchester [58]. The input data stream to FBFUSE consists of the complex-valued channels from every dish produced by the MeerKAT F-engine, which are transmitted via the Central BeamFormer (CBF) network; FBFUSE applies the geometric and phase delays (obtained by observing a bright calibrator) before combining the data streams from the dishes into one incoherent beam and up to 780 total intensity tied-array beams. The beams can be placed at any desired locations within the primary beam of the array, but are by default tessellated into a circular tiling centered on the boresight position, and spaced so that the response patterns of neighbouring beams intersect at the 25% peak power point. The beams are then sent over the network to TUSE for processing.

## Transient search pipeline

TUSE consists of 67 Lenovo servers with one head node and 66 compute nodes. Each compute node contains two Intel Xeon CPUs with 16 logical cores each, two Nvidia GeForce 1080 Ti Graphical processing units (GPUs) and 256 GB of DDR4 Random Access Memory (RAM). Each of the nodes is connected to a breakout switch via 10 GbE network interface cards (NIC) that are used to

ingest data from FBFUSE. The data from FBFUSE are received on the NICs as SPEAD2 (<https://casper.ssl.berkeley.edu/wiki/SPEAD>) heaps: small, self-describing units that contain a header and 8KB of data in time-frequency order, attached to a specific frequency sub-band and beam. FBFUSE streams these heaps over a range of multicast network addresses, where each address is associated with a fixed subset of 12 coherent beams in a typical configuration. Each TUSE node then subscribes to a distinct multi-cast address, which effectively distributes the processing load evenly across the TUSE cluster. The data ingest code on a TUSE node collects heaps and writes their contents to POSIX shared memory ring buffers (one per beam), so that the data for each beam are reconstructed in their natural time-frequency order (i.e. with the frequency axis contiguous in memory). Once a data segment has been fully reconstructed, it is ingested by the search pipeline. More details on TUSE will be presented in an upcoming paper.

We utilize the highly-optimized Graphics Processing Unit (GPU)-based ASTROACCELERATE software [59, 60] to search for dispersed signals. The real-time search is performed by incoherently de-dispersing in the DM range 0–5118.4 pc cm<sup>-3</sup> at L-band and 0–2664.9 pc cm<sup>-3</sup> in the UHF-band, divided into multiple sub-ranges with progressively increasing DM steps and time averaging factors. We also search up to maximum boxcar widths of 0.67 s and 0.77 s in the L-band and UHF-band respectively. The extracted candidate files contain raw filterbank data of the dispersed pulse and additional padding of 0.5 s at the start and at the end of the file.

For the targeted observations, where the position of PSR J0901–4046 is known to within an arcsecond, we were able to run in a mode where only 2 nodes with one beam per node were processing the data in real time. Only extracted candidates were saved for further examination. In addition to the real-time processing, two more nodes were used to record the continuous data to disk at a full data rate for further offline processing.

## APSUSE data recording

The Accelerated Pulsar Search User Supply Equipment (APSUSE) instrument consists of 68 high-performance Huawei FusionServer 2288H v5 servers with two head nodes for control. Nodes are split into two classes depending on their primary use. There are eight ingest nodes, designed for the capture and collation of Ethernet data and its writing to disk. Each ingest node has two 3-GHz Intel Xeon Gold 6136 CPUs, a 40-GbE Ethernet adapter and an 56 Gb/s FDR Infiniband adapter. The remaining 60 nodes are used for data analysis, with each having two 2.1-GHz Intel Xeon Silver 4116 CPU and two Nvidia GTX 1080 Ti GPUs. Additionally these nodes have FDR Infiniband adapters. All nodes in the cluster are connected via an Infiniband fabric with each node hosting 8 × 8 TB hard drives, providing a 3.2-PB distributed file-system with 40 GB/s of write performance. The APSUSE ingest nodes were used during PSR J0901–4046 observations to record beams formed by the FBFUSE

instrument. At L-band, data were recorded with 4096 frequency channels over a 856-MHz band centered at 1.284 GHz with a time resolution of 76.56  $\mu$ s. At UHF the data were recorded with 4096 frequency channels over a 544-MHz band centered at 816 MHz with a time resolution of 120.47  $\mu$ s.

## Full-time resolution data capture pipeline

We also used the transient buffer mode of the MeerTRAP pipeline to capture complex channelized data in order to investigate the sub-pulse structure and polarization properties of the PSR J0901–4046 single pulses. For each observation, the MJD of the first detection in the real-time pipeline was used as a reference in order to predict the time of arrival of future pulses using the best known pulse period. Then, a list of predicted MJDs were sent as a trigger to the transient buffer capture code ([https://github.com/ewanbarr/psrdada\\_cpp](https://github.com/ewanbarr/psrdada_cpp)) to each beamformer node. On receiving the trigger, the code uses the DM of the source to extract the complex channels around the time of the pulse accounting for the dispersion delay in each frequency channel before saving the data to disk. Each beamformer node saved data for one sub-band and all antennas used in the observation. For each pulse, around 1.5 seconds of complex channel data were stored. For every pulse, the data for each sub-band were combined and the proper complex gain correction applied per antenna and per frequency channel. The complex channels from each antenna were then added coherently, in phase, to form a beam at the phase centre of the telescope. These coherently beamformed complex channels were then used to study the polarization properties. Due to storage constraints only a small fraction of the total number of pulses from PSR J0901–4046 were captured.

## Astrometry

We compare the positions of radio sources in our combined images to those of the Rapid ASKAP Continuum Survey (RACS) [30]. We achieve this by running the PYBDSF [61] source finder on a joint MeerKAT image (made prior to model subtraction from all of the 2 second data), and the RACS image of the corresponding region. We forego the use of the RACS catalogue and run PYDSF directly on the RACS image in order to minimize any differences that may be introduced by the use of differing source finders. Furthermore, we select only sources that are compact in both catalogues (represented by a single point or Gaussian component) in order to minimize any resolution biases. The offsets between 254 MeerKAT sources and their matched RACS counterparts are shown in Supplementary Data Figure 1. The mean offsets in RA and Dec respectively are 0.90 ( $\pm 0.07$ )" and  $-0.55$  ( $\pm 0.05$ )", where the quoted uncertainties are the standard error in the mean. The standard deviations of the offsets in RA and Dec are 1.04" and 0.84" respectively.

RACS' positional accuracy is compared to that of the International Celestial Reference Frame (ICRF) version 3 [62]. Using RACS as an intermediate

step we find negligible ( $-0.052''$ ) systematic offset in RA, and a  $-0.95''$  systematic offset in Dec (approximately 1/6th the size of the restoring beam) between the MeerKAT positions and the ICRF v3.

## Location in the $P - \dot{P}$ parameter space.

Pulsars dissipate their rotational kinetic energy through electromagnetic radiation and a wind of relativistic particles leading to a gradual increase (called spin-down) of the rotation period as they age. A combination of the spin-period and spin-down or period derivative gives an estimate of a pulsar's age and magnetic field strength and places it in the  $P - \dot{P}$  parameter space which is used to compare relative populations of pulsars and to potentially trace their evolution (see Figure 1). Over the lifetime of the pulsar, it is thought to cross over a death line into a region called the pulsar graveyard [63] where it is no longer expected to emit in the radio. The rotational parameters of PSR J0901–4046 place it in the upper right region of the  $P - \dot{P}$  parameter space, below two of the three death-lines in Figure 1. These death-lines are dependent on conditions needed for pair production in the pulsar magnetosphere. The conventional emission model for pulsars assumes the existence of a vacuum gap above the polar cap of a neutron star. In order to sustain pair production, the potential difference across this gap must be sufficiently large and this is no longer possible beyond this death line. As a result, pair production and consequently, radio emission ceases.

Equation 9 from CR93 [22], illustrated by the solid line in Figure 1, represents a model for very curved or twisted magnetic field lines with curvature radius of the magnetic field line,  $r_c \sim R = 10^6$  cm comparable to the radius  $R$  of the neutron star. It also suggests a polar cap size much smaller than that of the pure dipole field case. PSR J0901–4046 lies well beyond this death-line implying that this model is not viable.

The model for curvature radiation from the vacuum gap (represented by Equation 4 in [52] and illustrated by the dot-dashed line in Figure 1), utilizes a multipole magnetic field configuration and relativistic frame-dragging effects to achieve the necessary potential difference for continued pair production. PSR J0901–4046 is also located beyond this death-line implying that an alternative model is required to explain the radio emission.

The space-charge-limited flow (SCLF) model death line for curvature radiation (calculated using Equation 9 in [52] and illustrated by the dashed line in Figure 1) assumes a multipole magnetic field and lies just below the location of PSR J0901–4046 indicating that this model is potentially feasible. Additionally, it has been proposed [64] that different equations of state of a neutron star possibly affect the death-line, thereby leading to heavier radio pulsars surviving beyond the standard death line.

For the simplest assumption of a dipolar magnetic field configuration, the minimum dipole magnetic field strength  $B$  at the surface of a canonical pulsar is,

$$\left( \frac{B}{\text{Gauss}} \right) > 3.2 \times 10^{19} \left( \frac{P\dot{P}}{\text{s}} \right)^{1/2}, \quad (2)$$

which corresponds to  $1.4 \times 10^{14}$  G for PSR J0901–4046. This value is similar to the bulk of the magnetar population in the  $P - \dot{P}$  parameter space (calculated using this same standard formula). Assuming that the source was born with an initial period much shorter than its current period, a simplified estimation of the approximate, or characteristic age,  $\tau$  can be made as,

$$\tau = \frac{P}{2\dot{P}}, \quad (3)$$

which is equivalent to an age of 5.3 Myr for PSR J0901–4046. A magnetar with a spin period of  $\sim 10$  s and a braking index  $n = 3$ , would evolve to the spin period of PSR J0901–4046 in  $\sim 5$  Myr. We note that even if  $n$  were as high as 5, it still gives an age of  $\sim 2$  Myr.

## Average Pulse Profiles and Pulse Modulation

In Supplementary Data Figure 2 we show the average pulse profiles formed from the observations taken on 2021-04-02 when we observed with both the UHF and L-band separated by less than 2 hours to be able to compare the behaviour quasi-simultaneously between the two frequencies. The profiles are similar at the two frequencies with a smooth leading edge rising up to the peak of the profile and then a second component apparent on the trailing edge. The UHF profile is less smooth at the peak and has a more distinct trailing component which may be to do with the increased modulation (see below) in the UHF single pulses. Although the L-band data is from a single epoch, it is representative of the average pulse seen on other days.

We measured the full-width at half maximum  $W_{50}$  of both the pulses by fitting five and seven von Mises functions (the von Mises distribution resembles a Gaussian distribution, but is cyclic) to the L- and UHF-band profiles respectively, using the PSRSALSA package [65]. This analytic, noise-free description of the profile allows a robust width measurement of the profile shape. The uncertainties on the measurements are estimated by bootstrapping, i.e. by repeatedly adding Gaussian noise with the standard deviation measured from the off-pulse region of the profile, over multiple iterations. The quoted uncertainty is the standard deviation of the iterations. We find that the  $W_{50}$  widths of the two pulses are identical to within the errors and also to within 0.3% of the pulse width indicating that between the two bands there is no evidence for radius-to-frequency mapping. It is interesting to note that the  $W_{50}$  point lies just about the point where the trailing feature is present indicating that too is apparently not evolving dramatically with frequency.

As can be seen in Figure 1 there are a number of different types of pulse shapes that are seen from PSR J0901–4046. To try and quantify the variability we calculated the modulation index and standard deviation as a function of

pulse phase from a series of pulses (see [66] for the definition of these terms). We again used the UHF and L-band data from 2021-04-02 and the data set was the same as that used for the pulsar timing. We formed a fully-frequency averaged stack of single pulses for each of the 30 minute sessions, two each at the two frequencies, using the ephemeris given in Table 1. RFI mitigation was done before frequency averaging and was the same as described for the timing analysis above except that in some cases the pulses were corrupted too much by the RFI and these pulses were removed. We then combined the two 30-minute stacks to form a single stack for each frequency, resulting in about 45 pulses in each case. The single pulses were recorded with 65536 phase bins across a pulse period to give a time resolution of approximately 1.15 ms. The subsequent analysis was all undertaken using the PSRSALSA package. The off-pulse baseline was removed from each single pulse, and in some cases there was a slope in the baseline across the pulse and this was removed for doing a linear fit to the off-pulse region. The data were then gated to the section around the pulse profile as seen in Supplementary Data Figure 2.

There are stark differences between the modulation properties at the two frequencies, which were observed less than 2 hours apart. We note that a similar analysis that used all the L-band pulses over the 6 epochs in 2021 gives a result very similar to the one obtained from this single day L-band data set. We only have this single epoch of UHF data, but given the stability of the L-band properties it is likely that this is representative of the UHF data and even if it is not, no similar outlier is seen in all the previous L-band observations, which still makes this difference very interesting. We note that the S/N at UHF is significantly higher than the L-band data and this is supported by the greater flux measured in the imaging data at UHF (weighted mean of  $169.3 \pm 14$  mJy beam<sup>-1</sup>, compared to  $89.3 \pm 2.7$  mJy beam<sup>-1</sup> at L-band). The modulation and standard deviation at L-band are reminiscent of many other pulsars, with the former peaking at the pulse edges and the latter somewhat tracing the intensity and neither being particularly high. This correlates well with the very low pulse jitter discussed in the timing section and confirms that despite the fact that the individual pulses are radically different, they form a stable average after a surprisingly small number of pulses. In contrast the UHF data show very strong modulation with parts of the leading edge exceeding a modulation index of one and a clear dip where the transition to the trailing component happens in the average pulse profile. We also see that the trailing component is much more highly modulated than it is at L-band. The standard deviation is also very high and traces the average profile indicating that the overall pulse-to-pulse flux is changing significantly. The imaging data also shows a greater standard deviation in flux between pulses at UHF (15 mJy) than at L-band (5 mJy) confirming the phase resolved analysis.

Observations of the 23-second pulsar [19] at 350 MHz showed significant modulation in the pulse profile, that had not been seen at other frequencies, which was particularly evident in the leading component. However, these data



were not contemporaneous and the profile was different than that seen at higher and lower frequencies [67].

## Polarimetric Calibration

The complex channelized data for every pulse from PSR J0901–4046 were processed using a custom-made pipeline to generate beamformed, full-Stokes dynamic spectra (see “Voltage data capture pipeline” for more details). Then, we calibrated the polarization using the methodology given below. We assume that any leakage between the two polarization hands affects only Stokes V (defined as  $V=LL-RR$  using the PSR IEEE convention [78]). We also assume that the delay between the two polarization hands affects only Stokes Q and U. The calibration we apply ignores second-order effects. We perform a brute-force search for the rotation measure that maximizes the linear polarization fraction by using the PSRCHIVE tool `rmfit` [42]. With `rmfit` we select a range of rotation measures to search, in a number of trial steps. The delay between the polarization hands approximately manifests as an offset from the true rotation measure of the source, assuming the delay is frequency-independent. While this method provides a correct rotation measure, it still does not calibrate the polarization position angle (PPA) to its absolute value. The complex gain solutions that are computed by the Science Data Processor pipeline, also account for the phase offset and delays between the two dipoles of the receiver and the phase offsets and delays with respect to the sky [68]. Since the data for each single pulse were only 1.5 seconds long, we ignored the phase offsets because of the rotation of the feed with respect to the sky. In order to confirm whether we have absolute PPA calibration, we also captured complex voltage data for a very well-known, bright pulsar PSR J0437–4715 whose polarization is very well measured at 1.4 GHz. We processed the data for PSR J0437–4715 with same pipeline and corrected the full-Stokes data for rotation measure with the best known RM value taken from [69]. Then we compared the resulting polarization of PSR J0437–4715 with a fully calibrated profile taken from the publicly available EPN pulsar database (<http://www.epta.eu.org/epndb/>) [70]. Extended Data Figure 6 shows examples of fully calibrated single pulses of PSR J0901–4046. It is clear from the 1284 GHz pulses that the PPA follows a S-shaped sweep akin to PPAs observed for canonical radio pulsars. This is more obvious in the phased resolved PPA histogram shown in Supplementary Data Figure 3.

## Rotating Vector Model fits

The Rotating Vector Model (RVM; [71]) provides a geometric interpretation of the variation of the PPA with rotational phase. To increase the sensitivity 11 consecutive pulses were summed together. The resulting polarimetric profile is shown in Supplementary Data Figure 4. A rapid swing in PPA is shown near the peak of the profile, with the steepest part of the swing coinciding with the transition between the two main components of the profile. There is a distinct

dip in the degree of linear polarization. Although such a dip can in principle be the result of mixture of orthogonal polarization modes, the same orthogonal mode appears to dominate in this pulse longitude region for all the analyzed pulses (see the PPA distribution in Supplementary Data Figure 3).

The PPA swing as observed in individual pulses is variable, especially where the degree of linear polarization is low. Nevertheless, the steep swing is a stable feature. The PPA swing of the 11 summed pulses was fitted with the RVM following the methodology described in [72]. The best fit (red line in the bottom panel of Supplementary Data Figure 4) is relatively poor (the reduced  $\chi^2$  is 16), with significant deviations from what can be modelled with the RVM. The inability of the RVM to reproduce the observed PPA swing accurately means that the model parameters are poorly constrained. The magnetic inclination angle (with respect to the rotation axis) is unconstrained. What can be constrained is that the line of sight passes very close to the magnetic axis, with an impact parameter  $\beta \lesssim 0.2^\circ$ .

A small impact parameter is expected for a slowly rotating pulsar, as it will have an extremely large light cylinder, and a correspondingly small magnetic pole defined by open magnetic field lines on which particle acceleration and the production of radio emission is expected. The PPA swing therefore suggests that the observed radio emission originates close to the magnetic pole of the neutron star.

## Sub-pulse width evolution

On top of the well-defined quasi-periodic behaviour and substructure shape, multiple pulses show a complex evolution in time. An example of such a pulse can be seen in the 'quasi-periodic' panel of Figure 2. Here the pulse starts with relatively fast oscillations and the oscillations appear to slow down closer to the pulse peak. Interestingly, most quasi-periodic pulses show such 'driven' oscillation suggesting some sort of a 'charge and discharge' mechanism. Such behaviour may be too complex to be easily identifiable within a single ACF, which in this case is more sensitive to the shorter period. In order to investigate the evolution of the components and their quasi-periodicity we perform a continuous wavelet transform (CWT) on pulses detected with the APSUSE instrument as those data possess a finer time resolution (76  $\mu$ s). This approach allows us to analyze our signal in both the time and frequency domains. The CWT of signal  $x(t)$  can be computed using the equation,

$$\text{CWT}(a, b) = \frac{1}{\sqrt{a}} \int_{-\text{inf}}^{+\text{inf}} x(t) \psi^* \left( \frac{t-b}{a} \right) dt, \quad (4)$$

where  $\psi(t)$  is the mother wavelet. The wavelet can then be stretched or compressed by a scale factor  $a$  and shifted by the translational parameter  $b$ . By varying the scale factor, the wavelet transform can be sensitive to features with different sizes. By applying the shift, the transform can localize the signal details in time. In our analysis we use two wavelets to extract different features

from the pulse profiles. The first order derivative of the Gaussian wavelet, later referred to as the G1 wavelet, at a time  $t$  given by

$$\psi_G(t) = C_G t \exp(-t^2), \quad (5)$$

is used to detect sharp changes in the pulse profiles. By using a wavelet best suited to designing sharp transitions in the signal, we can obtain the start and end times of the quasi-periodic emission and derive the widths of individual components. To better understand the evolution of the quasi-period, we detect individual components and localize them in time. We can distinguish separate sub-pulses by using the Ricker wavelet of the form

$$\psi_R(t) = \frac{2}{\sqrt{3\pi^{\frac{1}{4}}}} (1 - t^2) \exp\left(\frac{-t^2}{2}\right) \quad (6)$$

Supplementary Data Figure 5 shows an example wavelet transform analysis performed for one of the pulses with a strong quasi-periodic behaviour. The de-dispersed signal at a time resolution of 76  $\mu$ s is first convolved with a window of size 64 samples. This reduces the amount of unwanted noise and makes the structures of interest more prominent, without negatively impacting the time resolution. The convolved signal is then high-pass filtered to remove the overall pulse envelope. In this case, we use an ideal high-pass filter and fully filter-out frequencies below 5 Hz. The resulting convolved and filtered signal, shown in blue in all the panels of Supplementary Data Figure 5, is passed through the two wavelet transforms described above. In the case of the G1 wavelet the trough-to-peak transition results in a positive value of the wavelet transform at the point of the transition. Conversely, the peak-to-trough transition gives a negative value of the wavelet transform. When using the Ricker wavelet, the transform has its maximum at the centre of every peak. For both transforms we then use a simple peak finding algorithm to find and extract features of interest. The bottom row of Supplementary Data Figure 5 shows the resulting peak and trough width (left panel) and peak-to-peak period (right panel) measurements. For this pulse, the sub-pulse widths and separations are relatively stable and the averaged quasi-period 82.90 ms is consistent with the value of 79.7 ms obtained using the ACF method.

We repeat this process for a sample of the detected pulses. Of those, 39 had at least 1 clearly defined sub-pulse that we were able to measure the width of. Supplementary Data Figure 6a shows the distribution of 145 sub-pulse widths that we have obtained with the wavelet method. The solid vertical line represents the median sub-pulse width of 49.00 ms. The dashed vertical lines represent one median absolute deviation limits. We can see that the majority of sub-pulse widths are concentrated around the median value, with only a small number of detections beyond 100 ms. The largest width detected within our sub-pulse sample was 120.43 ms, which is close to 50% of the FWHM at L-band reported in Table 1.

Supplementary Data Figure 6b shows the relationship between the quasi-period of the sub-pulse oscillations and the width of the components of the

individual pulses. To reduce the complexity, only median values of component widths and their quasi-periods for every pulse are included. The solid black line represents a theoretical behaviour where the width of the sub-pulse is exactly half of the quasi-period. This describes a sub-pulse emission that spends half of the period in the ‘on’ state, followed by half of the period in the ‘off’ state. For our sample of pulses, only a handful can be classified as existing in such a state, with the divergent behaviour away from this half-period line for higher quasi-periods.

## Extant data

We looked at archival synthesis images from various imaging surveys in the Southern Hemisphere. We searched for point source like emission in the closest fields in the TIFR-GMRT Sky Survey (TGSS) observed on 2016 March 15 at 150 MHz [? ], the Sydney University Molonglo Sky Survey (SUMSS) observed between February 2003 and July 2007 at 843 MHz [? ] and the Rapid ASKAP Continuum Survey (RACS) observed on 2019 April 30 at 887.5 MHz [? ]. We did not find any significant emission at the best known radio position of PSR J0901–4046 in all these surveys and report a 3-sigma upper limit of 10.2 mJy, 6.6 mJy and 1.5 mJy for TGSS, SUMSS and RACS respectively. Taking the pulse-averaged flux density at 1.4 GHz as 200  $\mu$ Jy and the measured spectral index of  $-1.7$ , i.e. approximately consistent with those of known rotation-powered pulsars [73], the expected flux densities would be 8.9 mJy, 0.47 mJy and 0.43 mJy for TGSS, SUMSS and RACS, respectively. Additionally, we examined the literature for pulsar time-domain surveys that had covered the PSR J0901–4046 field in the past. The field was indeed observed as part of the Parkes Multibeam Pulsar Survey (PMPS, [31]) in 1999 March (pointing centre offset by about 2.2 arcmin) and there were three further PMPS pointings near it in 1999 July/August, albeit offset by about 12 - 14 arcmin. We obtained the data and searched them for single pulses and periodic emission, both blindly and by folding using the known ephemeris. We repeated the analysis without applying a high-pass filter, as was originally done in the default PMPS data processing, on the dedispersed time series. Neither the single-pulse, nor the periodicity search resulted in a detection. This is not surprising, as the data from the primary pointing are heavily affected by RFI. There was also a high-pass filter consisting of a 2-component RC time constant included in the data acquisition system which meant that longer timescale structure than 0.9 seconds would be removed from the data. However, as the pulse width is less than this, some harmonics would still get through into the spectrum. The sensitivity penalty suffered because of this would therefore depend to a large extent on the pulse width, as well as the period. We estimate the 8-sigma flux density upper limit assuming a 1% duty cycle to be about 0.3 mJy. This too suggests that the non-detections in archival data (both synthesis imaging and time-domain) are not unexpected.

## Energy distributions


We examine the possibility of various types of pulses corresponding to variations in the emission mechanism by computing pulse energy distributions. For this analysis we group the quasi-periodic and partially nulling pulses together, as all quasi-periodic pulses are partially nulling but not vice versa. It is not straightforward to separate these two classes without some sort of a qualitative metric which is beyond the scope of this paper. The average on-pulse energy distribution for each pulse archetype across both frequencies and all epochs is shown in Supplementary Data Figure 8. We use the PSRSALSA software suite to model pulse profile envelopes for the quasi-periodic and partially nulling pulses using a low-pass filter (see Supplementary Data Figure 7) to estimate the energy a pulse would have without the dropouts in power (Panel 2 from the top in Supplementary Data Figure 8). We lose  $\sim 40\%$  of the energy to the observed dropouts. When this loss is accounted for by the profile envelopes, the energy distribution is not very different from the distributions of the ‘normal’ or ‘split-peak’ pulses. These oscillations are likely linked to the radio emission production itself, or a periodic absorption mechanism that can suppress the radio emission. The brightest observed pulses are classified under the split-peak category and are UHF-detections which is consistent with the observed spectral index reported in Table 1.

## Optical Follow-up


High speed photometry of the 17th mag Gaia optical source near the radio position of PSR J0901–4046 was obtained on 7 and 9 January 2021, using the 1-m telescope of the South African Astronomical Observatory and the SHOC camera. On 7 January 2021, the source was observed with no filter and 5 second continuous exposures for 2.66 hr (start time BJD 245 9222.41766). The second observation on 9 January 2021 again used no filter, and exposure times of 10 seconds continuous for 4.13 hr (start time BJD 245 9224.37231). The second observation showed some evidence for long term variations on time scales longer than the observation length, but no photometric modulations were seen at 76 seconds.

Spectra of the same Gaia optical counterpart candidate was undertaken with SALT [74] on 2021 February 2. Two consecutive 1200 s exposures were obtained, beginning at 01:25:53 UTC, with the Robert Stobie Spectrograph [75] which used the PG900 VPH grating, covering the region 3920–6990 Å at a mean resolution of 5.7 Å with a 1''5 slit width. The spectra were reduced using the PYSALT package [76], (<https://astronomers.salt.ac.za/software/pysalt-documentation/>), which corrects bias, gain, amplifier cross-talk and cosmetic defects and finally mosaics the three CCDs comprising the detector. The spectral extraction, wavelength calibration and background subtraction were all undertaken using standard IRAF routines, as was the relative flux calibration. The latter was achieved using an observation of EG21, taken

on 31 January 2021. The SALT spectrum of PSR J0901–4046 was flux calibrated using an observation of the spectrophotometric flux standard, EG21, observed with the same grating setup two nights prior. Due to the inherent design of SALT, in particular the moving and variable entrance pupil, it is impossible to determine absolute fluxes of spectra. However, the relative fluxes and shape of spectra that are calibrated are reliable. In Supplementary Data Figure 9 we show the spectrum of the Gaia optical counterpart candidate, which is typical of a reddened A-type star and therefore unlikely to be associated with PSR J0901–4046.




**Supplementary Data Figure 1:** Offsets between the MeerKAT positions and those of the Rapid ASKAP Continuum Survey [30]. We find negligible systematic offset of  $-0.052''$  in RA, and a  $-0.95''$  systematic offset in Dec between the MeerKAT positions and the ICRF v3.




**Supplementary Data Figure 2:** Pulse modulation properties of PSR J0901–4046 at UHF and L-band. A comparison between the modulation (blue), standard deviation (orange) and average pulse profiles (dark) for the L-band (top panel) and UHF (bottom panel) data taken on 2021-04-02. There are some baseline variations still in the data due to the long period and so the standard deviation is truncated at a value of 0.15 and modulation indices are only plotted when the error is less than 0.1. The profiles are scaled to have a peak amplitude of 1 and off-pulse mean of 0.

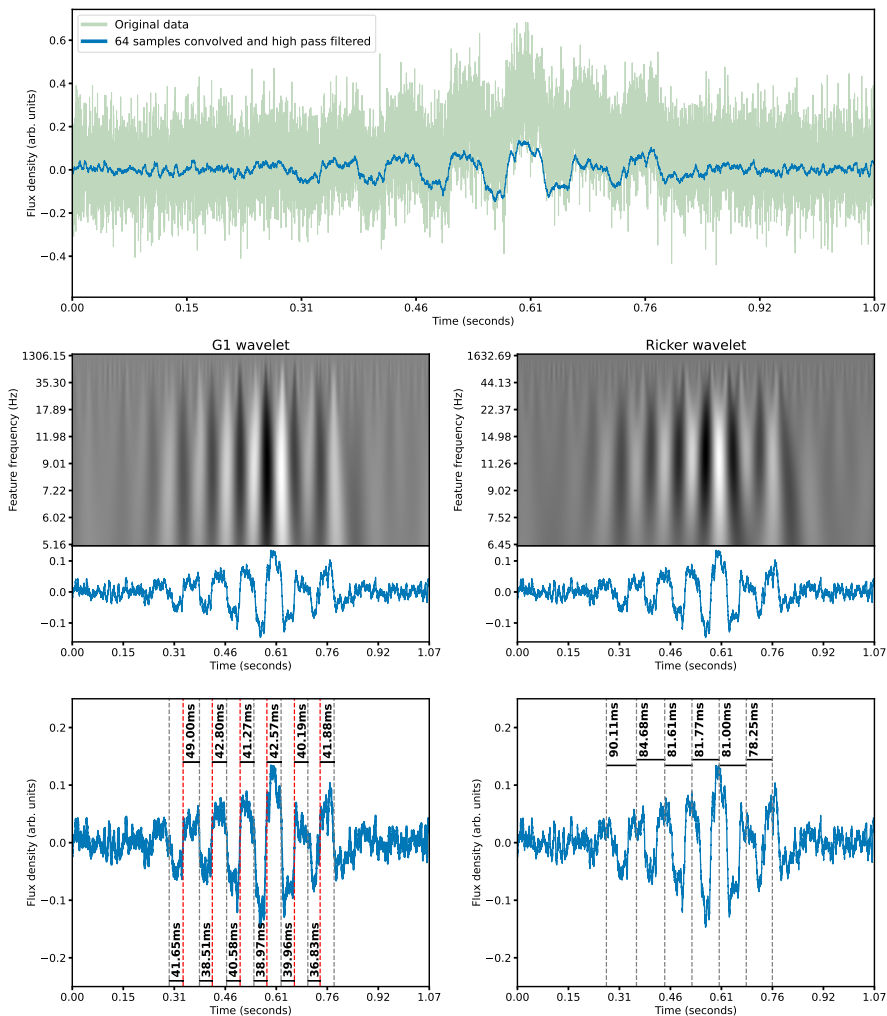




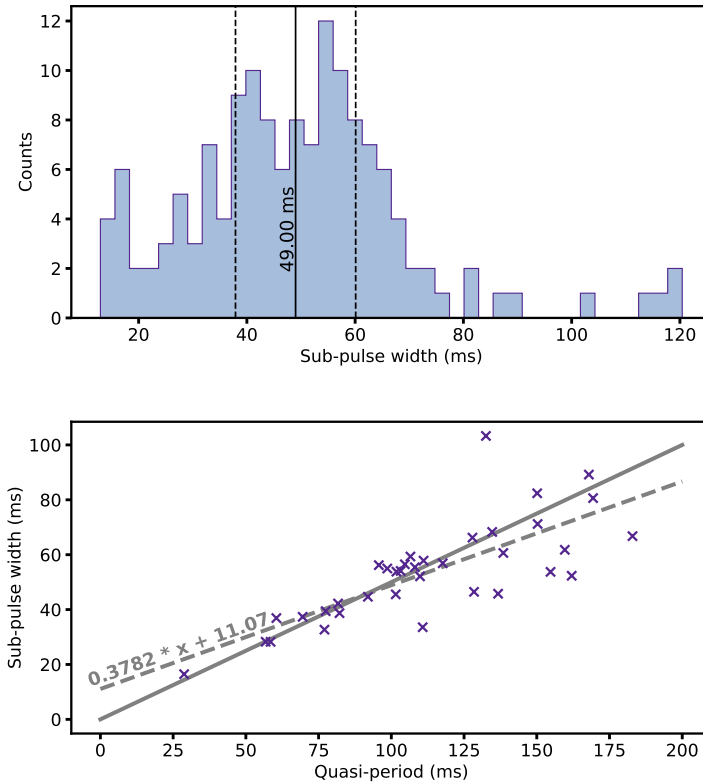
**Supplementary Data Figure 3:** Phase resolved polarization position angle histogram. The top panel shows the phase resolved PPA histogram. The heat map of the PPA clearly traces out a S-shaped sweep reminiscent of the PPA swing seen in canonical radio pulsars. Bottom panel shows the average profile of the 23 pulses of PSR J0901–4046 added together.



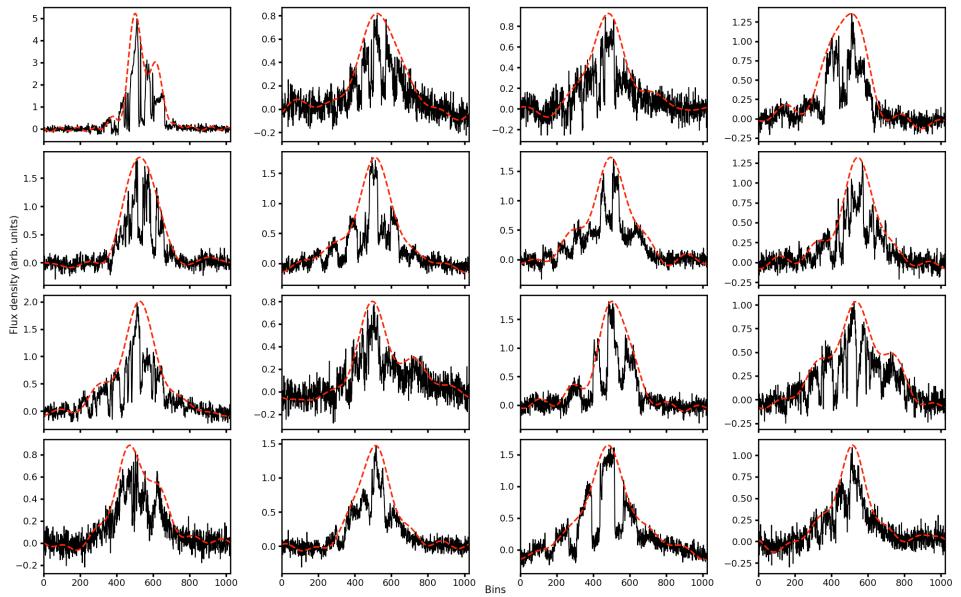
**Supplementary Data Figure 4:** The polarimetric profile of PSR J0901–4046 obtained by summing 11 consecutive individual pulses recorded on the first of February 2021. Top panel: The normalized Stokes  $I$  profile, the degree of linear polarization and Stokes  $V$  are shown as the black, red and green curves respectively. Bottom panel: The position angle (points with error bars) as detected with  $2\sigma$  confidence and the best fit RVM (red line).



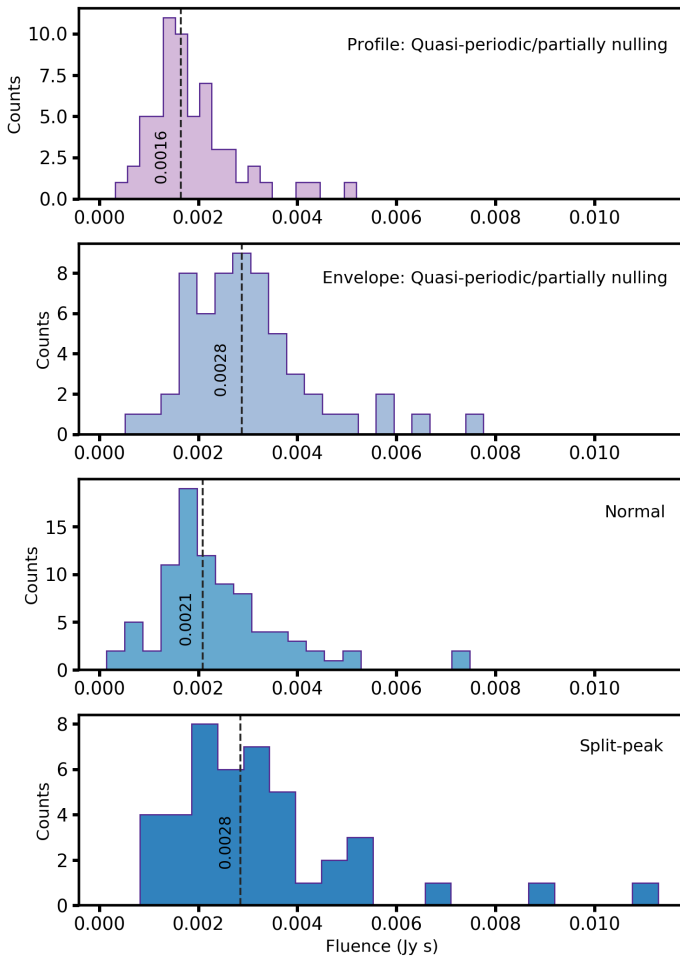
**Supplementary Data Figure 5:** Example wavelet transform analysis for one of the pulses exhibiting a strong quasi-periodic behaviour. Two wavelets were used to estimate the emission properties. The first order derivative of the Gaussian wavelet was used to estimate the positions of peak-to-trough transitions, which provided an estimate of the individual sub-pulse widths and their evolution with time. The Ricker wavelet was used to determine the distance between consecutive sup-pulse peaks, providing independent and consistent estimates of quasi-periods to those obtained with the ACF method.




**Supplementary Data Figure 6:** *Top panel:* Sub-pulse width distributions for widths estimated using the wavelet method. The solid vertical line represents a median sub-pulse width of 49.00 ms. *Bottom panel:* Quasi-period – sub-pulse width relationship. The solid grey line represents a behaviour where the sub-pulses widths are half of their quasi-period, which is a case in e.g. sinusoidal oscillations. The dashed line represents a least squares fit, which shows a deviation from sinusoidal oscillations.



**Supplementary Data Figure 7:** Sample of fitted pulse profile envelopes for quasi-periodic and partially nulling pulses. The average energies of the profiles themselves as well as modelled envelopes are represented by the solid and dashed lines respectively.



**Supplementary Data Figure 8:** Pulse energy distribution of the different archetypes of pulses.  $\sim 40\%$  of the energy is lost to the dropouts/dips seen in the quasi-periodic and partially nulling pulses in panel 1. When this is accounted for by modelling the pulse envelope (panel 2), the resulting distribution is similar to the other distributions in panels 3 and 4. This suggests that the pulses with dropouts/dips are not drastically brighter than the other types.



**Supplementary Data Figure 9:** SALT spectrum of coincident Gaia source, ruled out as the likely optical counterpart candidate. See text for details.

## References

- [53] Jonas, J. & MeerKAT Team *The MeerKAT Radio Telescope*, 1 (2016).
- [54] Camilo, F. *et al.* Revival of the Magnetar PSR J1622-4950: Observations with MeerKAT, Parkes, XMM-Newton, Swift, Chandra, and NuSTAR. *Astrophys. J.* **856**, 180 (2018).
- [55] Mauch, T. *et al.* The 1.28 GHz MeerKAT DEEP2 Image. *Astrophys. J.* **888**, 61 (2020).
- [56] Barr, E. D. Weltevrede, P., Perera, B. B. P., Preston, L. L. & Sanidas, S. (eds) *An S-band Receiver and Backend System for MeerKAT*. (eds Weltevrede, P., Perera, B. B. P., Preston, L. L. & Sanidas, S.) *Pulsar Astrophysics the Next Fifty Years*, Vol. 337, 175–178 (2018).
- [57] Chen, W., Barr, E., Karuppusamy, R., Kramer, M. & Stappers, B. Wide field beamformed observation with MeerKAT. *Journal of Astronomical Instrumentation* **10**, 2150013-178 (2021).
- [58] Caleb, M. *et al.* Simultaneous multi-telescope observations of FRB 121102. *Mon. Not. R. Astron. Soc.* **496**, 4565–4573 (2020).
- [59] Armour, W. *et al.* *A gpu-based survey for millisecond radio transients using artemis*, Vol. 461, 33–36 (ASP, 2011).
- [60] Adámek, K. & Armour, W. Single-pulse detection algorithms for real-time fast radio burst searches using gpus. *The Astrophysical Journal Supplement Series* **247**, 56 (2020).
- [61] Mohan, N. & Rafferty, D. PyBDSF: Python Blob Detection and Source Finder (2015). [1502.007](#).
- [62] Charlot, P. *et al.* The third realization of the International Celestial Reference Frame by very long baseline interferometry. *Astron. Astrophys* **644**, A159 (2020).
- [63] Lorimer, D. R. & Kramer, M. *Handbook of Pulsar Astronomy* (2012).
- [64] Zhou, X., Tong, H., Zhu, C. & Wang, N. Dependence of pulsar death line on the equation of state. *Mon. Not. R. Astron. Soc.* **472**, 2403–2409 (2017).
- [65] Weltevrede, P. Investigation of the bi-drifting subpulses of radio pulsar B1839-04 utilising the open-source data-analysis project PSRSALSA. *Astron. Astrophys* **590**, A109 (2016).



- [66] Weltevrede, P., Edwards, R. T. & Stappers, B. W. The subpulse modulation properties of pulsars at 21 cm. *Astron. Astrophys* **445**, 243–272 (2006).
- [67] Agar, C. H. *et al.* A broadband radio study of PSR J0250+5854: the slowest-spinning radio pulsar known. *Mon. Not. R. Astron. Soc.* **508**, 1102–1114 (2021).
- [68] Serylak, M. *et al.* The thousand-pulsar-array programme on MeerKAT IV: Polarization properties of young, energetic pulsars. *Mon. Not. R. Astron. Soc.* **505**, 4483–4495 (2021).
- [69] Osłowski, S., van Straten, W., Bailes, M., Jameson, A. & Hobbs, G. Timing, polarimetry and physics of the bright, nearby millisecond pulsar PSR J0437-4715 - a single-pulse perspective. *Mon. Not. R. Astron. Soc.* **441**, 3148–3160 (2014).
- [70] Dai, S. *et al.* A study of multifrequency polarization pulse profiles of millisecond pulsars. *Mon. Not. R. Astron. Soc.* **449**, 3223–3262 (2015).
- [71] Radhakrishnan, V. & Cooke, D. J. Magnetic Poles and the Polarization Structure of Pulsar Radiation. *Astrophys. J.* **3**, 225 (1969).
- [72] Rookyard, S. C., Weltevrede, P. & Johnston, S. Constraints on viewing geometries from radio observations of  $\gamma$ -ray-loud pulsars using a novel method. *Mon. Not. R. Astron. Soc.* **446**, 3367–3388 (2015).
- [73] Jankowski, F. *et al.* Spectral properties of 441 radio pulsars. *Mon. Not. R. Astron. Soc.* **473**, 4436–4458 (2018).
- [74] Buckley, D. A. H., Swart, G. P. & Meiring, J. G. *Completion and commissioning of the Southern African Large Telescope*, Vol. 6267 of *Proc. SPIE*, 62670Z (2006).
- [75] Burgh, E. B. *et al.* Iye, M. & Moorwood, A. F. M. (eds) *Prime Focus Imaging Spectrograph for the Southern African Large Telescope: optical design*. (eds Iye, M. & Moorwood, A. F. M.) *Instrument Design and Performance for Optical/Infrared Ground-based Telescopes*, Vol. 4841 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 1463–1471 (2003).
- [76] Crawford, S. M. *et al.* Silva, D. R., Peck, A. B. & Soifer, B. T. (eds) *PySALT: the SALT science pipeline*. (eds Silva, D. R., Peck, A. B. & Soifer, B. T.) *Observatory Operations: Strategies, Processes, and Systems III*, Vol. 7737 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 773725 (2010).