Astronomy and Computing 18 (2017) 40–46

Contents lists available at ScienceDirect

Astronomy and Computing

journal homepage: www.elsevier.com/locate/ascom

Full length article

Enabling pulsar and fast transient searches using coherent dedispersion

C.G. Bassa[a,*], Z. Pleunis[b], J.W.T. Hessels[a,b]

[a] ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands
[b] Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH, Amsterdam, The Netherlands

ARTICLE INFO

Article history:
Received 4 July 2016
Accepted 31 January 2017
Available online 9 February 2017

Keywords:
Methods
Data analysis—pulsars
General

ABSTRACT

We present an implementation of the coherent dedispersion algorithm capable of dedispersing high-time-resolution radio observations to many different dispersion measures (DMs). This approach allows the removal of the dispersive effects of the interstellar medium and enables searches for pulsed emission from pulsars and other millisecond-duration transients at low observing frequencies and/or high DMs where time broadening of the signal due to dispersive smearing would otherwise severely reduce the sensitivity. The implementation, called `cdmt`, for *coherent dispersion measure trials*, exploits the parallel processing capability of general-purpose graphics processing units to accelerate the computations. We describe the coherent dedispersion algorithm and detail how `cdmt` implements the algorithm to efficiently compute many coherent DM trials. We apply the concept of a semi-coherent dedispersion search, where coherently dedispersed trials at coarsely separated DMs are subsequently incoherently dedispersed at finer steps in DM. The software is used in an ongoing LOFAR pilot survey to test the feasibility of performing semi-coherent dedispersion searches for millisecond pulsars at 135 MHz. This pilot survey has led to the discovery of a radio millisecond pulsar—the first at these low frequencies. This is the first time that such a broad and comprehensive search in DM-space has been done using coherent dedispersion, and we argue that future low-frequency pulsar searches using this approach are both scientifically compelling and feasible. Finally, we compare the performance of `cdmt` with other available alternatives.

1. Introduction

Advances in electronics, computing and networking, primarily following Moore's law, have enabled the realization of a new type of radio telescope. Instead of placing receivers at the focus of movable reflecting dishes, these new telescopes employ a large number of stationary dipole antennas to create what is called an *aperture array*. The signals from these antennas are combined digitally in a correlator to create images of the sky, or in a beamformer to form beams on the sky. Three major digital aperture arrays operating at low radio frequencies are presently in operation: the LOw Frequency Array (LOFAR; van Haarlem et al., 2013), the Murchison Wide Field Array (MWA; Lonsdale et al., 2009; Tingay et al., 2013) and the Long Wavelength Array (LWA; Ellingson et al., 2009, 2013), while the low-frequency component of the Square Kilometre Array (SKA1-Low) is being planned (Braun et al., 2015). To maximize sensitivity while keeping the number of individual antennas, and hence cost, low, these aperture arrays operate at long wavelengths, and hence low observing frequencies (below 300 MHz).

One of the key science areas at these low observing frequencies is the study of radio pulsars. Radio emission from these highly magnetized, rotating neutron stars exhibits a very steep spectrum (Maron et al., 2000; Bates et al., 2013), typically peaking or turning over between 100 to 200 MHz (e.g. Malofeev et al., 1994). Surveys at these low frequencies have the prospect of discovering new pulsars that are too faint to be detected in surveys performed at higher observing frequencies and can take advantage of large fields of view (e.g. Coenen et al., 2014). Of particular interest are radio pulsars spinning at millisecond spin periods. These millisecond pulsars provide unparalleled precision for measuring neutron star masses, performing precision tests of General Relativity, understanding binary evolution, and detecting gravitational waves (e.g. Kramer et al., 2006; Demorest et al., 2010; Ransom et al., 2014; Antoniadis et al., 2013; Verbiest et al., 2016).

Radio emission propagating through the ionized interstellar medium suffers from dispersion, introducing a frequency dependent time delay over the requisite large bandwidths of radio astronomical observations. As a result, pulsed signals, such as those of

* Corresponding author.
E-mail address: bassa@astron.nl (C.G. Bassa).

http://dx.doi.org/10.1016/j.ascom.2017.01.004
2213-1337/

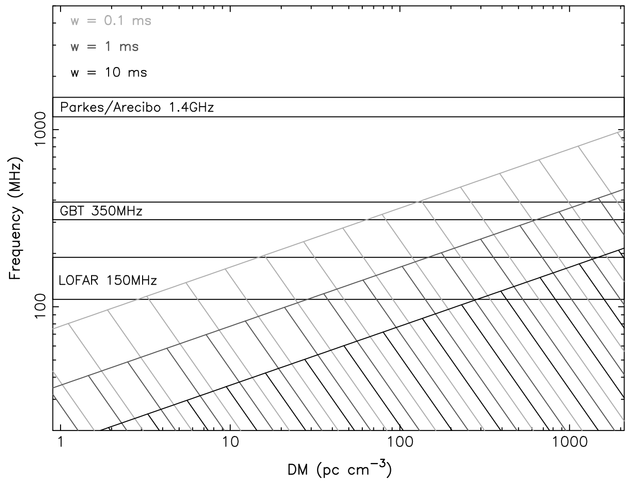

Fig. 1. The effect of residual dispersion smearing in frequency channels as a function of dispersion measure (DM) and frequency. The diagonal lines denote the 3σ detection limit of an undispersed input pulse of 10σ with pulse full-width-half-maxima of $w = 0.1$, 1 and 10 ms. For frequencies below, or DMs above, these limits, denoted by the hashed areas, the pulse is no longer detectable. The channel size is set at $\Delta\nu = 0.02$ MHz which corresponds to a time resolution of $\Delta t = \Delta\nu^{-1} = 50\,\mu$s. This time resolution is what is typically used in current millisecond pulsar searches. Using narrower channels would adversely reduce the time resolution. The frequency bands and central frequencies of representative radio telescopes are shown with the horizontal lines.

pulsars and fast transients (a generic term used for other sources of millisecond-duration radio pulses; e.g. Lorimer et al., 2007), have a specific dispersion measure (DM) which relates directly to the column density of free electrons between the source and the observer. When searching for new pulsars and fast transients, correcting for this dispersion is typically done through *incoherent dedispersion*, where the dispersive delays are removed by time shifting the time-series of individual, narrow, frequency channels by an amount appropriate to the DM of the source. Though this corrects for dispersion between channels, the dispersion within the finite bandwidth of the individual channels is not corrected for. Nonetheless, the computational efficiency of the technique has been a practical necessity compared to more accurate approaches.

A priori, the DM of a new pulsar or fast transient is unknown, and the data must be dedispersed to a broad range of different DMs. Models for the Galactic electron density (e.g. Cordes and Lazio, 2002) can be used to estimate the maximum DM towards a given direction though –to enable sensitivity to extra-Galactic fast transients –most ongoing surveys search up to a maximum DM of several thousand pc cm^{-3}. Depending on the frequency and time resolution of the input data, several thousand DM trials may need to be computed (Cordes and McLaughlin, 2003).[1] Though dedispersing that many DM trials is a computationally expensive task, recent implementations of incoherent dedispersion algorithms on graphics processing units (GPUs) are fast enough to allow real-time processing (Magro et al., 2011; Barsdell et al., 2012; Zackay and Ofek, 2014; Sclocco et al., 2016).

At low frequencies and/or high DMs, incoherent dedispersion can lead to significant smearing of the pulse (in time) within a channel (see Fig. 1). The effect of dispersion can be completely removed through *coherent dedispersion*. This approach convolves the input signals with the inverse of the transfer function of the interstellar medium. This convolution must be performed before the signal is detected (squared) as the phase information, in addition to the amplitude, is required. Hence, the data rate and computational requirements for coherent dedispersion are typically larger than the filterbanked data used for incoherent dedispersion, as for the latter the two polarizations can be squared and frequency and/or time resolution can be reduced. Because of these higher data rates, coherent dedispersion is presently only used for observing either known pulsars or when searching for pulsars in globular clusters with known DMs (see Prager et al., 2016 and references therein).

Here we present `cdmt`, for *coherent dispersion measure trials*, which implements the coherent dedispersion algorithm to perform coherent dedispersion to many dispersion measure trials in parallel on GPUs. This software allows us to control the residual dispersion smearing within a channel and retain both high time and high frequency resolution when searching for pulsars and fast transients. In a semi-coherent dedispersion search, the input data can be coherently dedispersed to several coarsely separated trial DMs, each of which is then incoherently dedispersed with finer DM steps around the coherent trial DM. Though the total number of incoherent DM trials, and hence processing requirements, will increase, this approach allows us to search for millisecond pulsars at lower observing frequencies than were previously possible – thus probing a new astrophysical parameter space.

The paper is structured as follows. The coherent dedispersion algorithm, combined with channelizing the data, is described in Section 2. Our implementation of the algorithm is outlined in Section 3. In Section 4 we provide an application example of the software and we report on the performance in Section 5. Finally, we discuss prospects for `cdmt` in Section 6.

2. Algorithm description

The coherent dedispersion algorithm as implemented in `cdmt` is that of a convolving synthetic filterbank. This implementation performs coherent dedispersion as a complex multiplication in the frequency domain and combines dedispersion with channelization. For a detailed description of the coherent dedispersion algorithm and different filterbanking versions, we refer to the implementation in the `dspsr` software package, which is detailed in van Straten and Bailes (2011).[2] Here, we closely follow the description outlined in these references to explain our implementation. A schematic representation of the implementation is depicted in Fig. 2.

Coherent dedispersion is a convolution of the raw signal voltages (Nyquist sampled time-series) with the inverse of the transfer function of the interstellar medium (ISM). The convolution is most efficiently performed as a multiplication in the frequency domain through the discrete convolution theorem (see Press et al., 1992 Chapter 13.1). In addition to the convolution, the convolving synthetic filterbank trades time resolution for frequency resolution to create a user-defined number of channels n_c over the input bandwidth. This channelization step is combined with coherent dedispersion by performing a large forward Fourier transform of $N_{\rm bin}$ samples prior to dedispersion, followed by n_c backward Fourier transforms of $N_{\rm bin}/n_c$ samples to provide the channelization.

The transfer function of the ISM, when modeled as a cold tenuous plasma, is defined in the frequency domain (Hankins, 1971; Hankins and Rickett, 1975, see also Lorimer and Kramer, 2012) as:

$$H(\nu + \nu_0) = \exp\left[\frac{2\pi i \nu^2 k_{\rm DM} {\rm DM}}{\nu_0^2(\nu + \nu_0)}\right]. \quad (1)$$

Here ν_0 is the center frequency of a subband or channel of bandwidth $\Delta\nu$ (both in MHz), while ν is the frequency offset within the

[1] See Lyon et al., 2016 and http://www.jb.man.ac.uk/pulsar/Surveys.html for the parameters of ongoing and historic pulsar and fast transient surveys.

[2] See also Willem van Straten's Ph.D. thesis at http://astronomy.swin.edu.au/~wvanstra/papers/thesis.html.