

Accelerating incoherent dedispersion

B. R. Barsdell,^{*} M. Bailes, D. G. Barnes and C. J. Fluke

Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn 3122, Australia

Accepted 2012 January 21. Received 2012 January 18; in original form 2011 August 15

ABSTRACT

Incoherent dedispersion is a computationally intensive problem that appears frequently in pulsar and transient astronomy. For current and future transient pipelines, dedispersion can dominate the total execution time, meaning its computational speed acts as a constraint on the quality and quantity of science results. It is thus critical that the algorithm be able to take advantage of trends in commodity computing hardware. With this goal in mind, we present an analysis of the ‘direct’, ‘tree’ and ‘sub-band’ dedispersion algorithms with respect to their potential for efficient execution on modern graphics processing units (GPUs). We find all three to be excellent candidates, and proceed to describe implementations in C for CUDA using insight gained from the analysis. Using recent CPU and GPU hardware, the transition to the GPU provides a speed-up of nine times for the direct algorithm when compared to an optimized quad-core CPU code. For realistic recent survey parameters, these speeds are high enough that further optimization is unnecessary to achieve real-time processing. Where further speed-ups are desirable, we find that the tree and sub-band algorithms are able to provide three to seven times better performance at the cost of certain smearing, memory consumption and development time trade-offs. We finish with a discussion of the implications of these results for future transient surveys. Our GPU dedispersion code is publicly available as a C library at <http://dedisp.googlecode.com/>.

Key words: methods: data analysis – pulsars: general.

1 INTRODUCTION

With the advent of modern telescopes and digital signal processing back-ends, the time-resolved radio sky has become a rich source of astrophysical information. Observations of pulsars allow us to probe the nature of neutron stars (Lattimer & Prakash 2004), stellar populations (Bhattacharya & van den Heuvel 1991), the Galactic environment (Gaensler et al. 2008), plasma physics and gravitational waves (Lyne et al. 2004). Of equal significance are transient signals such as those from rotating radio transients (McLaughlin et al. 2006) and potentially rare one-off events such as ‘Lorimer bursts’ (Lorimer et al. 2007; Keane et al. 2011), which may correspond to previously unknown phenomena. These observations all depend on the use of significant computing power to search for signals within long, frequency-resolved time series.

As radiation from sources such as pulsars propagates to Earth, it interacts with free electrons in the interstellar medium. This interaction has the effect of delaying the signal in a frequency-dependent manner – signals at lower frequencies are delayed more than those at higher frequencies. Formally, the observed time delay, Δt , between two frequencies ν_1 and ν_2 as a result of dispersion by the

interstellar medium is given by

$$\Delta t = k_{\text{DM}} \text{DM} (\nu_1^{-2} - \nu_2^{-2}), \quad (1)$$

where $k_{\text{DM}} = \frac{e^2}{2\pi m_e c} = 4.148\,808 \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s}$ is the dispersion constant¹ and the frequencies are in MHz. The parameter DM specifies the dispersion measure along the line of sight in pc cm^{-3} , and is defined as

$$\text{DM} \equiv \int_0^d n_e dl, \quad (2)$$

where n_e is the electron number density (cm^{-3}) and d is the distance to the source (pc).

Once a time-varying source has been detected, its dispersion measure can be obtained from observations of its phase as a function of frequency; this in turn allows the distance to the object to be calculated via equation (2), assuming one has a model for the Galactic electron density n_e . When searching for *new* sources, however, one does not know the distance to the object. In these cases, the dispersion measure must be guessed prior to looking for a signal. To avoid excessive smearing of signals in the time series, and a consequent loss of signal-to-noise ratio, survey pipelines typically repeat the

^{*}E-mail: bbarsdel@astro.swin.edu.au

¹We note that the dispersion constant is commonly approximated in the literature as $1/(2.41 \times 10^{-4}) \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s}$.




Figure 1. An illustration of a dispersion trail (top) and its corresponding dedispersion transform (bottom). The darkest horizontal slice in the dedispersion transform gives the correctly dedispersed time series.

process for many trial dispersion measures. This process is referred to as a *dedispersion transform*. An example of the dedispersion transform is shown in Fig. 1.

Computing the dedispersion transform is a computationally expensive task: a simple approach involves a summation across a band of e.g. $\sim 10^3$ frequency channels for each of $\sim 10^3$ (typically) dispersion measures, for each time sample. Given modern sampling intervals of $O(64 \mu\text{s})$, computing this in real-time is a challenging task, especially if the process must be repeated for multiple beams. The prohibitive cost of real-time dedispersion has traditionally necessitated that pulsar and transient survey projects use offline processing.

In this paper we consider three ways in which computation of the dedispersion transform may be accelerated, enabling real-time processing at low cost. First, in Section 2 we demonstrate how modern many-core computing hardware in the form of graphics processing units (GPUs; see Fluke et al. 2011, for an introduction) can provide an order of magnitude more performance over a multi-core central processing unit (CPU) when dedispersing ‘directly’. The use of GPUs for incoherent dedispersion is not an entirely new idea. Dodson et al. (2010) introduced an implementation of such a system as part of their fast transient survey. Magro et al. (2011) described a similar approach and how it may be used to construct a GPU-based real-time transient detection pipeline for modest fractional bandwidths, demonstrating that their GPU dedisperser could outperform a generic code by two orders of magnitude. In this work we provide a thorough analysis of both the direct incoherent dedispersion algorithm itself and the details of its implementation on GPU hardware.

In Section 3 we then consider the use of the ‘tree’ algorithm, a (theoretically) more efficient means of computing the dedispersion transform. To our knowledge, this technique has not previously been implemented on a GPU. We conclude our analysis of dedispersion algorithms in Section 4 with a discussion of the ‘sub-band’ method, a derivative of the direct method.

In Section 5, we report accuracy and timing benchmarks for the three algorithms and compare them to our theoretical results. Finally, we present a discussion of our results, their implications for

future pulsar and transient surveys and a comparison with previous work in Section 6.

2 DIRECT DEDISPERSION

2.1 Introduction

The direct dedispersion algorithm operates by directly summing frequency channels along a quadratic dispersion trail for each time sample and dispersion measure. In detail, the algorithm computes an array of dedispersed time series D from an input data set A according to the following equation:

$$D_{d,t} = \sum_v^{N_v} A_{v,t+\Delta t(d,v)}, \quad (3)$$

where the subscripts d , t and v represent dispersion measure, time sample and frequency channel, respectively, and N_v is the total number of frequency channels. Note that throughout this paper we use the convention that \sum_i^N means the sum over the range $i = 0$ to $i = N - 1$. The function $\Delta t(d, v)$ is a discretized version of equation (1) and gives the time delay relative to the start of the band in whole time samples for a given dispersion measure and frequency channel:

$$\Delta T(v) \equiv \frac{k_{\text{DM}}}{\Delta\tau} \left(\frac{1}{(v_0 + v\Delta v)^2} - \frac{1}{v_0^2} \right), \quad (4)$$

$$\Delta t(d, v) \equiv \text{round}(\text{DM}(d)\Delta T(v)), \quad (5)$$

where $\Delta\tau$ is the time difference in seconds between two adjacent samples (the *sampling interval*), v_0 is the frequency in MHz at the start of the band, Δv is the frequency difference in MHz between two adjacent channels and the function $\text{round}(x)$ means x rounded to the nearest integer. The function or array $\text{DM}(d)$ is used to specify the dispersion measures to be computed. Note that the commonly used central frequency, v_c , and bandwidth, BW , parameters are related by $\text{BW} \equiv N_v\Delta v$ and $v_c \equiv v_0 + \frac{1}{2}\text{BW}$.

After dedispersion, the dedispersed time series $D_{d,t}$ can be searched for periodic or transient signals.

When dedispersing at large DM , the dispersion of a signal can be such that it is smeared significantly within a single frequency channel. Specifically, this occurs when the gradient of a dispersion curve on the time-frequency grid is less than unity (i.e. beyond the ‘diagonal’). Once this effect becomes significant, it becomes somewhat inefficient to continue to dedisperse at the full native time resolution. One option is to reduce the time resolution by a factor of 2 when the DM exceeds the diagonal by adding adjacent pairs of time samples. This process is then repeated at two times the diagonal, four times, etc. We refer to this technique as ‘time-scrunching’. The use of time-scrunching will reduce the overall computational cost, but can also slightly reduce the signal-to-noise ratio if the intrinsic pulse width is comparable to that of the dispersion smear.

2.2 Algorithm analysis

The direct dedispersion algorithm’s summation over N_v frequency channels for each of N_t time samples and N_{DM} dispersion measures gives it a computational complexity of

$$T_{\text{direct}} = O(N_t N_v N_{\text{DM}}). \quad (6)$$

The algorithm was analysed previously for many-core architectures in Barsdell, Barnes & Fluke (2010). The key findings were as follows:

(i) the algorithm is best parallelized over the ‘embarrassingly parallel’ dispersion-measure (d) and time (t) dimensions, with the sum over frequency channels (ν) being performed sequentially,

(ii) the algorithm has a very high theoretical arithmetic intensity, of the same magnitude as the number of dispersion measures computed [typically $O(100-1000)$], and

(iii) the memory access patterns generally exhibit reasonable locality, but their non-trivial nature may make it difficult to achieve a high arithmetic intensity.

While overall the algorithm appears well-positioned to take advantage of massively parallel hardware, we need to perform a deeper analysis to determine the optimal implementation strategy. The pattern in which memory is accessed is often critical to performance on massively parallel architectures, so this is where we now turn our attention.

While the d dimension involves a potentially non-linear mapping of input indices to output indices, the t dimension maintains a contiguous mapping from input to output. This makes the t dimension suitable for efficient memory access operations via *spatial caching*, where groups of adjacent parallel threads access memory all at once. This behaviour typically allows a majority of the available memory bandwidth to be exploited.

The remaining memory access issue is the potential use of *temporal caching* to increase the arithmetic intensity of the algorithm. Dedispersion at similar DMs involves accessing similar regions of input data. By pre-loading a block of data into a shared cache, many DMs could be computed before needing to return to main memory for more data. This would increase the arithmetic intensity by a factor proportional to the size of the shared cache, potentially providing a significant performance increase, assuming the algorithm was otherwise limited by available memory bandwidth. The problem with the direct dedispersion algorithm, however, is its non-linear memory access pattern in the d dimension. This behaviour makes a caching scheme difficult to devise, as one must account for threads at different DMs needing to access data at delayed times. Whether temporal caching can truly be used effectively for the direct dedispersion algorithm will depend on details of the implementation.

2.3 Implementation notes

When discussing GPU implementations throughout this paper, we use the terms ‘Fermi’ and ‘pre-Fermi’ GPUs to mean GPUs of the NVIDIA Fermi architecture and those of older architectures, respectively. We consider both architectures in order to study the recent evolution of GPU hardware and gain insight into the future direction of the technology.

We implemented the direct dedispersion algorithm using the C for CUDA platform.² As suggested by the analysis in Section 2.2, the algorithm was parallelized over the dispersion-measure and time dimensions, with each thread summing all N_ν channels sequentially. During the analysis it was also noted that the algorithm’s memory access pattern exhibits good spatial locality in the time dimension, with contiguous output indices mapping to contiguous input indices. We therefore chose time as the fastest-changing (i.e. x) thread dimension, such that reads from global memory would always be from contiguous regions with a unit stride, maximizing throughput. The DM dimension was consequently mapped to the second (i.e. y) thread dimension.

While the memory access pattern is always contiguous, it is not always *aligned*. This is a result of the delays, $\Delta t(d, \nu)$, introduced in the time dimension. At all non-zero DMs, the majority of memory accesses will begin at arbitrary offsets with respect to the internal alignment boundaries of the memory hardware. The consequence of this is that GPUs that do not have built-in caching support may need to split the memory requests into many smaller ones, significantly impacting throughput to the processors. In order to avoid this situation, we made use of the GPU’s *texture memory*, which does support automatic caching. On pre-Fermi GPU hardware, the use of texture memory resulted in a speed-up of around five times compared to using plain device memory, highlighting the importance of understanding the details of an algorithm’s memory access patterns when using these architectures. With the advent of Fermi-class GPUs, however, the situation has improved significantly. These devices contain an L1 cache that provides many of the advantages of using texture memory without having to explicitly refer to a special memory area. Using texture memory on Fermi-class GPUs was slightly *slower* than using plain device memory (with L1 cache enabled), as suggested in the Compute Unified Device Architecture (CUDA) programming guide.³

Input data with fewer bits per sample than the machine word size (currently assumed to be 32 bits) were handled using bit-shifting and masking operations on the GPU. It was found that a convenient format for working with the input data was to transpose the input from time-major order to frequency-major order by whole words, leaving consecutive frequency channels within each word. For example, for the case of two samples per word, the data order would be: $(A_{v_1,t_1} A_{v_2,t_1}), (A_{v_1,t_2} A_{v_2,t_2}), \dots, (A_{v_3,t_1} A_{v_4,t_1}), (A_{v_3,t_2} A_{v_4,t_2}), \dots$, where brackets denote data within a machine word. This format means that time delays are always applied in units of whole words, avoiding the need to deal with intra-word delays.

The thread decomposition was written to allow the shape of the block (i.e. number of DMs or time samples per block) to be tuned. We found that for a block size of 256 threads, optimal performance on a Fermi GPU was achieved when this was divided into eight time samples \times 32 DMs. We interpreted this result as a cache-related effect, where the block shape determines the spread of memory locations accessed by a group of neighbouring threads spread across time-DM space, and the optimum occurs when this spread is minimized. On pre-Fermi GPUs, changing the block shape was found to have very little impact on performance.

To minimize redundant computations, the functions $DM(d)$ and $\Delta T(\nu)$ were pre-computed and stored in look-up tables for the given dispersion measures and frequency channels, respectively. Delays were then computed simply by retrieving values from the two tables and evaluating equation (5), requiring only a single multiplication and a rounding operation. On pre-Fermi GPUs, the table corresponding to $\Delta T(\nu)$ was explicitly stored in the GPU’s *constant memory* space, which provides extremely efficient access when all threads read the same value (this is always the case for our implementation, where frequency channels are traversed sequentially.). On Fermi-generation cards, this explicit use of the constant memory space is unnecessary – constant memory caching is used automatically when the compiler determines it to be possible.

To amortize overheads within the GPU kernel such as index calculations, loop counters and time-delay computations, we allowed each thread to store and sum multiple time samples. Processing

² <http://developer.nvidia.com/object/gpucomputing.html>

³ The current version of the CUDA programming guide is available for download at http://www.nvidia.com/object/cuda_develop.html

four samples per thread was found to significantly reduce the total arithmetic cost without affecting memory throughput. Increasing this number required more registers per thread (a finite resource), and led to diminishing returns; we found four to be the optimal solution for our implementation.

Our implementation was written to support a channel ‘kill mask’, which specifies which frequency channels should be included in the computation and which should be skipped (e.g. to avoid radio frequency interference present within them). While our initial approach was to apply this mask as a conditional statement [e.g. `if(kill_mask[channel]) {sum += data}`], it was found that applying the mask arithmetically (e.g. `sum += data * kill_mask[channel]`) resulted in better performance. This is not particularly surprising given the GPU hardware’s focus on arithmetic throughput rather than branching operations.

Finally, we investigated the possibility of using temporal caching, as discussed in the analysis in Section 2.2. Unlike most CPUs, GPUs provide a manually managed cache (known as *shared memory* on NVIDIA GPUs). This provides additional power and flexibility at the cost of programming effort. We used shared memory to stage a rectangular section of input data (i.e. of time-frequency space) in each thread block. Careful attention was given to the amount of data cached, with additional time samples being loaded to allow for differences in delay across a block. The cost of development was significant, and it remained unclear whether the caching mechanism could be made robust against a variety of input parameters. Further, we found that the overall performance of the code was not significantly altered by the addition of the temporal caching mechanism. We concluded that the additional overheads involved in handling the non-linear memory access patterns (i.e. the mapping of blocks of threads in time-DM space to memory in time-frequency space) negated the performance benefit of staging data in the shared cache. We note, however, that caching may prove beneficial when considering only low DMs (e.g. below the diagonal), where delays vary slowly and memory access patterns remain relatively compact.

In theory it is possible that, via careful re-packing of the input data, one could exploit the bit-level parallelism available in modern computing hardware in addition to the thread-level parallelism. For example, for 2-bit data, packing each 2-bit value into 8 bits would allow four values to be summed in parallel with a single 32-bit addition instruction. In this case, $\frac{2^8-1}{2^2-1} = 85$ additions could be performed before one risked integer overflow. To dedisperse say 1024 channels, one could first sum blocks of 85 channels and then finish the summation by re-packing the partial sums into a larger data type. This would achieve efficient use of the available processing hardware, at the cost of additional implementation complexity and overheads for re-packing and data management. We did not use this technique in our GPU dedispersion codes, although our reference CPU code does exploit this extra parallelism by packing four 2-bit samples into a 64-bit word before dedispersion.

3 TREE DEDISPERSION

3.1 Introduction

The tree dedispersion algorithm, first described by Taylor (1974), attempts to reduce the complexity of the dedispersion computation from $O(N_r N_v N_{DM})$ to $O(N_r N_v \log N_v)$. This significant speed-up is obtained by first regularizing the problem and then exploiting the regularity to allow repeated calculations to be shared between different DMs. While theoretical speed-ups of $O(100)$ are possible, in practice a number of additional overheads arise when working

with real data. These overheads, as well as its increased complexity, have meant that the tree algorithm is rarely used in modern search pipelines. In this work we investigate the tree algorithm in order to assess its usefulness in the age of many-core processors.

In its most basic form, the tree dedispersion algorithm is used to compute the following:

$$D'_{d',t} = \sum_v^{N_v} A_{v,t+\Delta t'(d',v)}, \quad (7)$$

$$\Delta t'(d',v) = \text{round}\left(d' \frac{v}{N_v - 1}\right), \quad (8)$$

for d' in the range $0 \leq d' < N_v$. The regularization is such that the delay function $\Delta t'(d',v)$ is now a linear function of v that ranges from 0 to exactly d' across the band. The DMs computed by the tree algorithm are therefore:

$$\text{DM}(d') = \frac{d'}{\Delta T(N_v - 1)}, \quad (9)$$

where the function $\Delta T(v)$ is that given by equation (4).

The tree algorithm is able to evaluate equation (7) for d' in the range $0 \leq d' < N_v$ in just $\log_2 N_v$ steps. It achieves this feat by using a *divide and conquer* approach in the same way as the well-known fast Fourier transform (FFT) algorithm. The tree algorithm is visualized in Fig. 2. We define the computation at each step i as follows:

$$A_{v,t}^0 \equiv A_{v,t} \quad (10)$$

$$A_{2v,t}^{i+1} = A_{\Phi(i,2v),t}^i + A_{\Phi(i,2v+1),t+\Theta(i,2v)}^i \quad (11)$$

$$A_{2^{v+1},t}^{i+1} = A_{\Phi(i,2^v),t}^i + A_{\Phi(i,2^{v+1}),t+\Theta(i,2^{v+1})}^i \quad (12)$$

$$D'_{d',t} = D_{v,t} = A_{v,t}^{\log_2 N_v}. \quad (13)$$

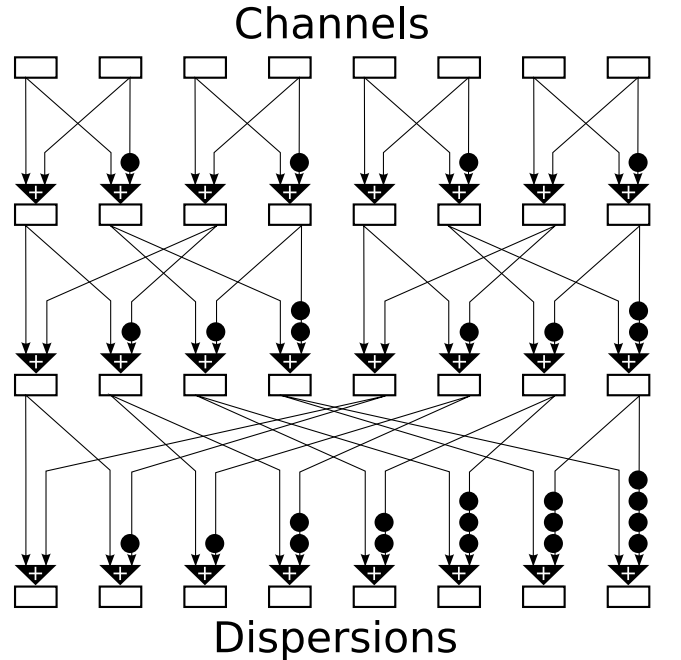


Figure 2. Visualization of the tree dedispersion algorithm. Rectangles represent frequency channels, each containing a time series going ‘into the page’. Arrows indicate the flow of data, triangles represent addition operations and circles indicate unit time delays into the page.

The integer function $\Theta(i, \nu)$ gives the time delay for a given iteration and frequency channel and can be defined as

$$\Theta(i, \nu) \equiv [(\nu \bmod 2^{i+1}) + 1]/2, \quad (14)$$

where mod is the modulus operator, and division is taken to be truncated integer division. The integer function $\Phi(i, \nu)$, which we refer to as the ‘shuffle’ function, re-orders the indices ν according to a pattern defined as follows:

$$\Phi(r, \nu) \equiv (\nu \bmod 2) \times r + \frac{\nu}{2} + \frac{\nu}{2r} \times r, \quad (15)$$

where the parameter $r \equiv 2^i$ is known as the *radix*.

While the tree dedispersion algorithm succeeds in dramatically reducing the computational cost of dedispersion, it has a number of constraints not present in the direct algorithm:

- (i) the computed dispersion trails are linear in frequency, not quadratic as found in nature (see equation 1),
- (ii) the computed dispersion measures are constrained to those given by equation (9), and
- (iii) the number of input frequency channels N_ν (and thus also the number of DMs) must be a power of two.

Constraint (iii) is generally not a significant concern, as it is common for the number of frequency channels to be a power of two, and blank channels can be added when this is not the case. Constraints (i) and (ii) are more problematic, as they prevent the computation of accurate and efficiently distributed dispersion trails. Fortunately there are ways of working around these limitations.

One method is to approximate the dispersion trail with piecewise linear segments by dividing the input data into sub-bands (Manchester et al. 1996). Another approach is to quadratically space the input frequencies by padding with blank channels as a pre-processing step such that the second order term in the dispersion trail is removed (Manchester et al. 2001). These techniques are described in the next two sections.

3.1.1 The piecewise linear tree method

Approximation of the quadratic dispersion curve using piecewise linear segments involves two stages of computation. If the input data are divided into N_s sub-bands of length

$$N'_\nu = \frac{N_\nu}{N_s}, \quad (16)$$

with the n th sub-band starting at frequency channel

$$\nu_n = nN'_\nu, \quad (17)$$

then from equation (7) we see that the tree dedispersion algorithm applied to each sub-band results in the following:

$$S_{n,d',t} = \sum_{\nu'}^{N'_\nu} A_{\nu_n + \nu', t + \Delta t'(d', \nu_n + \nu')}, \quad (18)$$

which we refer to as *stage 1* of the piecewise linear tree method.

In each sub-band, we approximate the quadratic dispersion trail with a linear one. We compute the linear DM in the n th sub-band that approximates the true DM indexed by d as follows:

$$d'_n(d) = \Delta t(d, \nu_{n+1}) - \Delta t(d, \nu_n) \quad (19)$$

$$= \text{round}(\text{DM}(d) [\Delta T(\nu_{n+1}) - \Delta T(\nu_n)]). \quad (20)$$

Applying the constraint $d'_n < N'_\nu$ and noting that the greatest dispersion delay occurs at the end of the band, we obtain a limit on

the DM that the basic piecewise linear tree algorithm can compute. This limit is commonly referred to as the ‘diagonal’ DM, as it corresponds to a dispersion trail in the time-frequency grid with a gradient of unity:⁴

$$\text{DM}(\text{piecewise})_{\text{diag}} = \frac{N'_\nu - \frac{1}{2}}{\Delta T(N_\nu) - \Delta T(N_\nu - N'_\nu)}. \quad (21)$$

A technique for computing larger DMs with the tree algorithm is discussed in Section 3.1.3.

The dedispersed sub-bands can now be combined to approximate the result of equation (3):

$$D_{d,t} \approx \sum_n^{N_s} S_{n,d'_n(d), t + \Delta t'_n(d)}, \quad (22)$$

$$\Delta t'_n(d) = \text{round} \left(\text{DM}(d) \sum_m^n \Delta T(\nu_{m+1}) - \Delta T(\nu_m) \right). \quad (23)$$

This forms *stage 2* of the piecewise linear tree computation.

The use of the tree algorithm with sub-banding introduces an additional source of smearing into the dedispersed time series as a result of approximating the quadratic dispersion curve with a piecewise linear one. We derive an analytic upper limit for this smearing in Appendix A.

3.1.2 The frequency-padded tree method

An alternative approach to adapting the tree algorithm to quadratic dispersion curves is to linearize the input data via a change of frequency coordinates. Formally, the aim is to ‘stretch’ $\Delta T(\nu)$ (equation 4) to a linear function $\Delta T'(\nu') \propto \nu'$. Expanding to the first order around $\nu = 0$, we have

$$\Delta T'(\nu') = \Delta T(0) + \Delta T(\nu') \left. \frac{d}{d\nu} [\Delta T(\nu)] \right|_0. \quad (24)$$

The change of variables $\nu \rightarrow \nu'$ is then found by equating $\Delta T(\nu)$ with its linear approximation, $\Delta T'(\nu')$, and solving for $\nu'(\nu)$, which gives

$$\nu' = \text{round} \left(\frac{1}{2} \frac{\nu_0}{\Delta \nu} \left[1 - \left(1 + \frac{\Delta \nu}{\nu_0} \nu \right)^{-2} \right] \right). \quad (25)$$

Evaluating at $\nu = N_\nu$ gives the total number of frequency channels in the linearized coordinates, which determines the additional computational overhead introduced by the procedure. Note, however, that this number must be rounded up to a power of two before the tree dedispersion algorithm can be applied. For observations with typical effective bandwidths and channel counts that are already a power of two, the frequency padding technique is unlikely to require increasing the total number of channels by more than a factor of 2.

In practice, the linearization procedure is applied by padding the frequency dimension with blank channels such that the real channels are spaced according to equation (25). Once the dispersion trails have been linearized, the tree algorithm can be applied directly.

The ‘diagonal’ DM when using the frequency padding method corresponds to

$$\text{DM}(\text{padded})_{\text{diag}} = \frac{1}{\Delta T(1)}. \quad (26)$$

⁴ Note that the ‘ $\frac{1}{2}$ ’ in equation (21) arises from the round-to-nearest operation in equation (20).

3.1.3 Computing larger DMs

The basic tree dedispersion algorithm computes exactly the DMs specified by equation (9). In practice, however, it is often necessary to search a much larger range of dispersion measures. Fortunately, there are techniques by which this can be achieved without having to resort to using the direct method. The tree algorithm can be made to compute higher DMs by first transforming the input data and then repeating the dedispersion computation. Formally, the following sequence of operations can be used to compute an arbitrary range of DMs.

- (i) Apply the tree algorithm to obtain DMs from zero to DM_{diag} .
- (ii) Impose a time delay across the band.
- (iii) Apply the tree algorithm to obtain DMs from DM_{diag} to $2DM_{\text{diag}}$.
- (iv) Increment the imposed time delay.
- (v) Repeat from step (ii) to obtain DMs up to $2^N DM_{\text{diag}}$.

The imposed time delay is initially a simple diagonal across the band (i.e. $\Delta t = \nu$), and is implemented by incrementing a memory stride value rather than actually shifting memory. While this method enables dedispersion up to arbitrarily large DMs, it does not alter the spacing of DM trials, which remains fixed as per equation (9).

The ‘time-scrunching’ technique, discussed in Section 2.1 for the direct algorithm, can also be applied to the tree algorithm. The procedure is as follows.

- (i) Apply the tree algorithm to obtain DMs from zero to DM_{diag} .
- (ii) Impose a time delay across the band.
- (iii) Apply the tree algorithm to obtain DMs from DM_{diag} to $2DM_{\text{diag}}$.
- (iv) Compress (‘scrunch’) time by a factor of 2 by summing adjacent samples.
- (v) Impose a time delay across the band.
- (vi) Apply the tree algorithm to obtain DMs from $2DM_{\text{diag}}$ to $4DM_{\text{diag}}$.
- (vii) Repeat from step (iv) to obtain DMs up to $2^N DM_{\text{diag}}$.

As with the direct algorithm, the use of time-scrunching provides a performance benefit at the cost of a minor reduction in the signal-to-noise ratio for pulses of intrinsic width near the dispersion measure smearing time.

3.2 Algorithm analysis

The tree dedispersion algorithm’s computational complexity of $O(N_t N_\nu \log N_\nu)$ breaks down into $\log_2 N_\nu$ sequential steps, with each step involving the computation of $O(N_t N_\nu)$ independent new values, as can be seen in equations (10) to (13). Following the analysis methodology of Barsdell et al. (2010), the algorithm therefore has a *depth complexity* of $O(\log N_\nu)$, meaning it contains this many sequentially dependent operations. Interestingly, this result matches that of the direct algorithm, although the tree algorithm requires significantly less total work. From a theoretical perspective, this implies that the tree algorithm contains less inherent parallelism than the direct algorithm. In practice, however, the number of processors will be small relative to the size of the problem ($N_t N_\nu$), and this reduced inherent parallelism is unlikely to be a concern for performance except when processing very small data sets.

Branching (i.e. conditional statements) within an algorithm can have a significant effect on performance when targeting GPU-like hardware (Barsdell et al. 2010). Fortunately, the tree algorithm is inherently branch-free, with all operations involving only memory

accesses and arithmetic operations. This issue is therefore of no concern in this instance.

The arithmetic intensity of the tree algorithm is determined from the ratio of arithmetic operations to memory operations. To process $N_t N_\nu$ samples, the algorithm involves $N_t N_\nu \log_2 N_\nu$ ‘delay and add’ operations, and produces $N_t N_\nu$ samples of output. In contrast to the direct algorithm, where the theoretical arithmetic intensity was proportional to the number of DMs computed, the tree algorithm requires only $O(\log N_\nu)$ operations per sample. This suggests that the tree algorithm may be unable to exploit GPU-like hardware as efficiently as the direct algorithm. However, the exact arithmetic intensity will depend on constant factors and additional arithmetic overheads, and will only become apparent once the algorithm has been implemented. We defer the discussion of these results to Section 3.3.

Achieving peak arithmetic intensity requires reading input data from ‘slow memory’ into ‘fast memory’ (e.g. from disc into main memory, from main memory into cache, from host memory into GPU memory etc.) only once, before performing all computations within fast memory and writing the results, again just once, back to slow memory. In the tree dedispersion algorithm, this means performing all $\log_2 N_\nu$ steps entirely within fast memory. The feasibility of this will depend on implementation details, the discussion of which we defer to Section 3.3. However, it will be useful to assume that some sub-set of the total computation will fit within this model. We will therefore continue the analysis of the tree algorithm under the assumption that we are computing only a (power-of-two) subset, or *block*, of B_ν channels.

The memory access patterns within the tree algorithm resemble those of the direct algorithm (see Section 2.2). Time samples are always accessed contiguously, with an offset that is essentially arbitrary. In the frequency dimension, memory is accessed according to the shuffle function (equation 15) depicted in Fig. 2, where at any given step of the algorithm the frequency channels ‘interact’ in pairs, the interaction involving their addition with different time delays.

With respect to the goal of achieving peak arithmetic intensity, the key issue for the memory access patterns within the tree algorithm is the extent to which they remain ‘compact’. This is important because it determines the ability to operate on isolated blocks of data independently, which is critical to prolonging the time between successive trips to slow memory. In the frequency dimension, the computation of some local (power-of-two) sub-set of channels B_ν involves accessing only other channels within the same subset. In this sense we can say that the memory access patterns are ‘locally compact’ in channels. In the time dimension, however, we note that the algorithm applies compounding delays (equivalent to offsets in whole time samples). This means that the memory access patterns ‘leak’ forward, with any Local Group of time samples always requiring access to the next group. The amount by which the necessary delays ‘leak’ in time for each channel is given by the integrated delay in that channel after B_ν steps (see Fig. 2). The total integrated delay across B_ν channels is $B_\nu(B_\nu - 1)/2$, which is the number of additional values that must be read into fast memory by the block in order to compute all $\log_2 B_\nu$ steps without needing to return to global memory and apply a global synchronization.

3.3 Implementation notes

As with the direct algorithm, we implemented the tree algorithm on a GPU in C for CUDA. For our first attempt, we took a simple approach where each of the $\log_2 N_\nu$ steps in the computation was performed

by a separate call to a GPU function (or *kernel*). This approach is not ideal, as it is preferable to perform more computation on the device before returning to the host (as per the discussion of arithmetic intensity in Section 3.2), but was necessary in order to guarantee global synchronization across threads on the GPU between steps. This is a result of the lack of global synchronization mechanisms on current GPUs.

Between steps, the integer delay and shuffle functions (equations 14 and 15) were evaluated on the host and stored in look-up tables. These were then copied to constant memory on the device prior to executing the kernel function to compute the step. The use of constant memory ensured retrieval of these values would not be a bottle-neck to performance during the computation of each step of the tree algorithm.

The problem was divided between threads on the GPU by allocating one thread for every time sample and every *pair* of frequency channels. This meant that each thread would compute the delayed sums between two ‘interacting’ channels according to the pattern depicted in Fig. 2 for the current step.

The tree algorithm’s iterative updating behaviour requires that computations at each step be performed ‘out-of-place’, i.e. output must be written to a memory space separate from that of the input to avoid modifying input values before they have been used. We achieved this effect by using a double-buffering scheme, where input and output arrays are swapped after each step.

While the algorithms differ significantly in their details, one point of consistency between the direct and tree methods is the need to apply time delays to the input data. Therefore, just as with our implementation of the direct algorithm, the tree algorithm requires accessing memory locations that are not aligned with internal memory boundaries. As such, we took the same approach as before and mapped the input data to the GPU’s texture memory before launching the device kernel. As noted in Section 2.3, this procedure is unnecessary on Fermi-generation GPUs, as their built-in caches provide the same behaviour automatically.

After successfully implementing the tree algorithm on a GPU using a simple one-step-per-GPU-call approach, we investigated the possibility of computing *multiple* steps of the algorithm on the GPU before returning to the CPU for synchronization. This is possible because current GPUs, while lacking support for global thread synchronization, do support synchronization across local thread groups (or *blocks*). These thread blocks typically contain $O(100)$ threads, and provide mechanisms for synchronization and data-sharing, both of which are required for a more efficient tree dedispersion implementation.

As discussed in Section 3.2, application of the tree algorithm to a block of B_v channels \times B_t time samples requires caching additional values from the next block in time. We used blocks of $B_v \times B_t = 16 \times 16$ threads, each loading both their corresponding data value and required additional values into shared cache. Once all values have been stored, computation of the $\log_2 B_v = 4$ steps proceeds entirely within the shared cache. Using larger thread blocks would allow more steps to be completed within the cache; however, the choice is constrained by the available volume of shared memory (typically around 48 kB). Once the block computation is complete, subsequent steps must be computed using the one-step-per-GPU-call approach described earlier, due to the requirement of global synchronizations.

While theory suggests that an implementation of the tree algorithm exploiting shared memory to perform multiple steps in cache would provide a performance benefit over a simpler implementation, in practice we were unable to achieve a net gain using this

approach. The limitations on block size imposed by the volume of shared memory, the need to load additional data into cache and the logarithmic scaling of steps relative to data size significantly reduce the potential speed-up, and overheads from increased code-complexity quickly erode what remains. For this reason we reverted to the straight-forward implementation of the tree algorithm as our final code for testing and benchmarking.

In addition to the base tree algorithm, we also implemented the sub-band method so as to allow the computation of arbitrary dispersion measures. This was achieved by dividing the computation into two stages. In stage 1, the first $\log_2 N'_v$ steps of the tree algorithm are applied to the input data, which produces the desired N_v/N'_v tree-dedispersed sub-bands. Stage 2 then involves applying an algorithm to combine the dedispersed time series in different sub-bands into approximated quadratic dispersion curves according to equation (22). Stage 2 was implemented on the GPU in much the same way as the direct algorithm, with input data mapped to texture memory (on pre-Fermi GPUs) and delays stored in look-up tables in constant device memory.

The frequency padding approach described in Section 3.1.2 was implemented by constructing an array large enough to hold the stretched frequency coordinates, initializing its elements to zero, and then copying (or *scattering*) the input data into this array according to equation (25). The results of this procedure were then fed to the basic tree dedispersion code to produce the final set of dedispersed time series.

Because the tree algorithm involves sequentially updating the entire data set, the data must remain in their final format for the duration of the computation. This means that low bit-rate data, e.g. 2-bit, must be unpacked (in a pre-processing step) into a format that will not overflow during accumulation. This is in contrast to the direct algorithm, where each sum is independent, and can be stored locally to each thread.

4 SUB-BAND DEDISPERSION

4.1 Introduction

Sub-band dedispersion is the name given to another technique used to compute the dedispersion transform. Like the tree algorithm described in Section 3, the sub-band algorithm attempts to reduce the cost of the computation relative to the direct method; however, rather than exploiting a *regularization* of the dedispersion algorithm, the sub-band method takes a simple *approximation* approach.

In its simplest form, the algorithm involves two processing steps. In the first, the set of trial DMs is approximated by a reduced set of $N_{DMnom} = N_{DM}/N'_{DM}$ ‘nominal’ DMs, each separated by N'_{DM} trial dispersion measures. The *direct* dedispersion algorithm is applied to sub-bands of N'_v channels to compute a dedispersed time series for each nominal DM and sub-band. In the second step, the DM trials near each nominal value are computed by applying the direct algorithm to the ‘miniature filterbanks’ formed by the time series for the sub-bands at each nominal DM. These data have a reduced frequency resolution of $N_{SB} = N_v/N'_v$ channels across the band. The two steps thus operate at reduced dispersion measure and frequency resolution, respectively, resulting in an overall reduction in the computational cost.

The sub-band algorithm is implemented in the PRESTO software suite (Ransom 2001) and was recently implemented on a GPU by Magro et al. (2011) (see Section 6.1 for a comparison with their work). Unlike the tree algorithm, the sub-band method is able to compute the dedispersion transform with the same flexibility as

the direct method, making its application to real observational data significantly simpler.

The approximations made by the sub-band algorithm introduce additional smearing into the dedispersed time series. We derive an analytic upper-bound in Appendix B and show that, to the first order, the smearing time t_{SB} is proportional to the product $N'_{\text{DM}}N'_v$ (see equation B2).

4.2 Algorithm analysis

The computational complexity of the sub-band dedispersion algorithm can be computed by summing that of the two steps:

$$T_{\text{SB},1} = N_{\text{SB}} T_{\text{direct}}(N_t, N'_v, N_{\text{DMnom}}) \quad (27)$$

$$T_{\text{SB},2} = N_{\text{DMnom}} T_{\text{direct}}(N_t, N_{\text{SB}}, N'_{\text{DM}}) \quad (28)$$

$$T_{\text{SB}} = T_{\text{SB},1} + T_{\text{SB},2} \quad (29)$$

$$= O \left[N_t N_{\text{DM}} N_v \left(\frac{1}{N'_{\text{DM}}} + \frac{1}{N'_v} \right) \right]. \quad (30)$$

This result can be combined with the knowledge of the smearing introduced by the algorithm to probe the relationship between accuracy and performance. Inserting the smearing constraint $t_{\text{SB}} \propto N'_{\text{DM}}N'_v$ (see Section 4.1) into equation (30), we obtain a second-order expression that is minimized at $N'_{\text{DM}} = N'_v \propto \sqrt{t_{\text{SB}}}$, which amounts to balancing the execution time between the two steps. This result optimizes the time complexity of the algorithm, which then takes the simple form

$$T'_{\text{SB}} = O \left(\frac{N_{\text{DM}} N_v}{\sqrt{t_{\text{SB}}}} \right) \quad (31)$$

and represents a theoretical speed-up over the direct algorithm proportional to the square root of the introduced smearing.

The sub-band algorithm's dependence on the direct algorithm means that it inherits similar algorithmic behaviour. However, as with the tree method, the decrease in computational work afforded by the sub-band approach corresponds to a decrease in the arithmetic intensity of the algorithm. This can be expected to reduce the intrinsic performance of the two sub-band steps relative to the direct algorithm.

One further consideration for the sub-band algorithm is the additional volume of memory required to store the intermediate results produced by the first step. These data consist of time series for each sub-band and nominal DM, giving a space complexity of

$$M_{\text{SB}} = O(N_{\text{SB}} N_{\text{DMnom}}). \quad (32)$$

Assuming the time complexity is optimized as in equation (31), the space complexity becomes

$$M'_{\text{SB}} = O \left(\frac{1}{t_{\text{SB}}} \right), \quad (33)$$

which indicates that the memory consumption increases much faster than the execution time, in direct proportion to the introduced smearing rather than to its square root. This can be expected to place a lower limit on the smearing that can be achieved in practice.

4.3 Implementation notes

A significant advantage of the sub-band algorithm over the tree algorithm is that it involves little more than repeated execution of the

direct algorithm. With sufficient generalization⁵ of our implementation of the direct algorithm, we were able to implement the sub-band method with just two consecutive calls to the direct dedispersion routine and the addition of a temporary data buffer.

In our implementation, the 'intermediate' data (i.e. the outputs of the first step) are stored in the temporary buffer using 32 bits per sample. The second call to the dedispersion routine then reads these values directly before writing the final output using a desired number of bits per sample.

Experimentation showed that optimal performance occurred at a slightly different shape and size of the thread blocks on the GPU compared to the direct algorithm (see Section 2.2). The sub-band kernels operated most efficiently with 128 threads per block divided into 16 time samples and eight DMs. In addition, the optimal choice of the ratio N'_v/N'_{DM} was found to be close to unity, which matches the theoretical result derived in Section 4.2. While these parameters minimized the execution time, the sub-band kernels were still found to perform around 40 per cent less efficiently than the direct kernel. This result is likely due to the reduced arithmetic intensity of the algorithm (see Section 4.2).

5 RESULTS

5.1 Smearing

Our analytic upper-bounds on the increase in smearing due to the use of the piecewise linear tree algorithm (equation A6) and the sub-band algorithm equation (B3) are plotted in the upper panels of Figs 3 and 4, respectively. The reference point (W in equations A6 and B3) was calculated using equations for the smearing during the *direct* dedispersion process⁶ assuming an intrinsic pulse width of 40 μs .

For the piecewise linear tree algorithm, the effective signal smearing at low dispersion measure is dominated by the intrinsic pulse width, the sampling time $\Delta\tau$ and the effect of finite DM sampling. As the DM is increased, however, the effects of finite channel width and the sub-band technique grow, and eventually become dominant. These smearing terms both scale linearly with the dispersion measure, and so the relative contribution of the sub-band method, μ_{SB} , tends to a constant.

The sub-band algorithm exhibits virtually constant smearing as a function of DM due to its dependence on the DM step, which is itself chosen to maintain a fixed fractional smearing. While the general trend mirrors that of the tree algorithm, the sub-band algorithm's smearing is typically around two orders of magnitude worse than its tree counterpart.

5.2 Performance

Our codes as implemented allowed us to directly compute the following.

- (i) Any list of DMs using the direct or sub-band algorithm with no time-scrunching.
- (ii) DMs up to the diagonal (see equation 21) using the piecewise linear tree algorithm.

⁵ The direct dedispersion routine was modified to support 'batching' (simultaneous application to several adjacent data sets) and arbitrary strides through the input and output arrays, trial DMs and channels.

⁶ Levin (private communication).

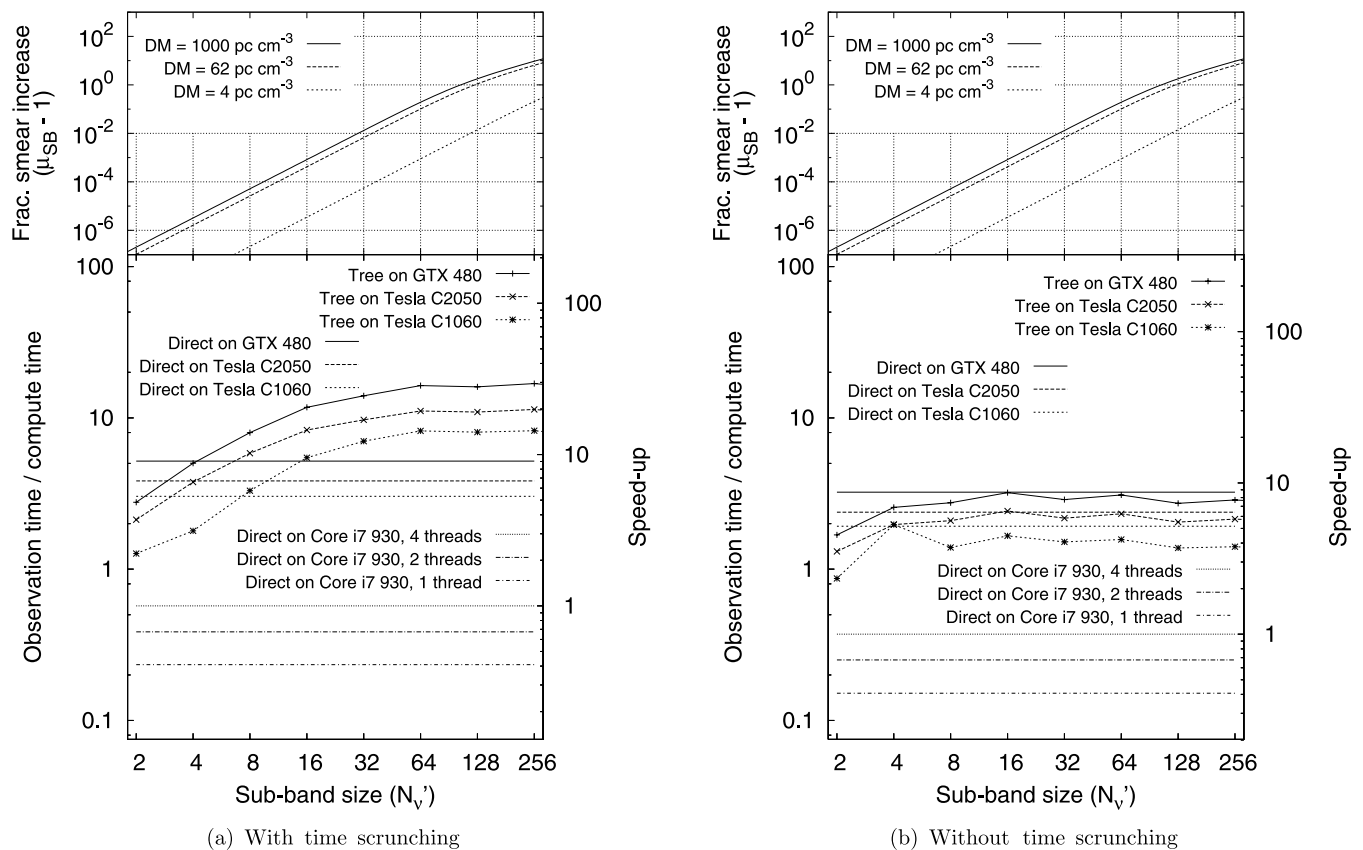


Figure 3. Upper panel: analytic upper-bound on signal degradation of a $40 \mu\text{s}$ pulse due to the piecewise linear tree algorithm as a function of the number of channels per sub-band (see equation A6). Lower panel: performance results for the direct and piecewise linear tree algorithms with (a) and without (b) ‘time-scrunching’ applied. Benchmarks were executed on an Intel Core i7 930 quad-core CPU and NVIDIA Tesla C1060, Tesla C2050 and GeForce GTX 480 GPUs. All results correspond to operations on one minute of input data with the following observing parameters: bits/sample = 2, $\nu_c = 1381.8 \text{ MHz}$, BW = 400 MHz, $N_v = 1024$, $\Delta\tau = 64 \mu\text{s}$. A total of 1196 DM trials were used, spaced non-linearly in the range $0 \leq \text{DM} < 1000 \text{ pc cm}^{-3}$ (see text for details). Error bars are too small to be seen at this scale and are not plotted. Note that performance results are projected from measurements of codes performing sub-sets of the benchmark task (see text for details).

(iii) DMs up to the diagonal (see equation 26) using the frequency-padded tree algorithm.

A number of timing benchmarks were run to compare the performance of the CPU to the GPU and the direct algorithm to the tree algorithms. Input and dispersion parameters were chosen to reflect a typical scenario as appears in modern pulsar surveys such as the High Time Resolution Universe (HTRU) survey currently underway at the Parkes radio telescope (Keith et al. 2010). The benchmarks involved computing the dedispersion transform of one minute of input data with observing parameters of bits/sample = 2, $\nu_0 = 1581.8 \text{ MHz}$, $\Delta\nu = -0.39062 \text{ MHz}$, $N_v = 1024$, $\Delta\tau = 64 \mu\text{s}$. DM trials were chosen to match those used in the HTRU survey, which were originally derived by applying an analytic constraint on the signal-smearing due to incorrect trial DM.⁷ The chosen set contained 1196 trial DMs in the range $0 \leq \text{DM} < 1000 \text{ pc cm}^{-3}$ with approximately exponential spacing.

For comparison purposes, we benchmarked a reference CPU direct dedispersion code in addition to our GPU codes. The CPU code (named `DEDISPERSE_ALL`) is highly optimized, and uses multiple CPU cores to compute the dedispersion transform (parallelized over the

time dimension) in addition to bit-level parallelism as described in Section 2.3. `DEDISPERSE_ALL` is approximately 60 times more efficient than the generic `DEDISPERSE` routine from `SIGPROC`,⁸ but is only applicable to a limited subset of data formats.

At the time of writing, our dedispersion code-base did not include ‘full-capability’ implementations of all of the discussed algorithms. However, we were able to perform a number of benchmarks that were sufficient to obtain accurate estimates of the performance of complete runs. Timing measurements for our codes were projected to produce a number of derived results representative of the complete benchmark task. The direct/sub-band dedispersion code was able to compute the complete list of desired DMs, but was not able to exploit time-scrunching; results for these algorithms with time scrunching were calculated by assuming that the computation of DMs between two times and four times the diagonal would proceed twice as fast as the computation up to two times the diagonal (as a result of there being half as many time samples), and similarly for four and eight times, etc. up to the maximum desired DM. A simple code to perform the time-scrunching operation (i.e. adding adjacent time samples to reduce the time resolution by a factor of 2) was also benchmarked and factored into the projection. For the tree codes, which were unable to compute DMs beyond the diagonal, timing

⁷ Levin (private communication); see Cordes & McLaughlin (2003) for a similar derivation.

⁸ `sigproc.sourceforge.net`

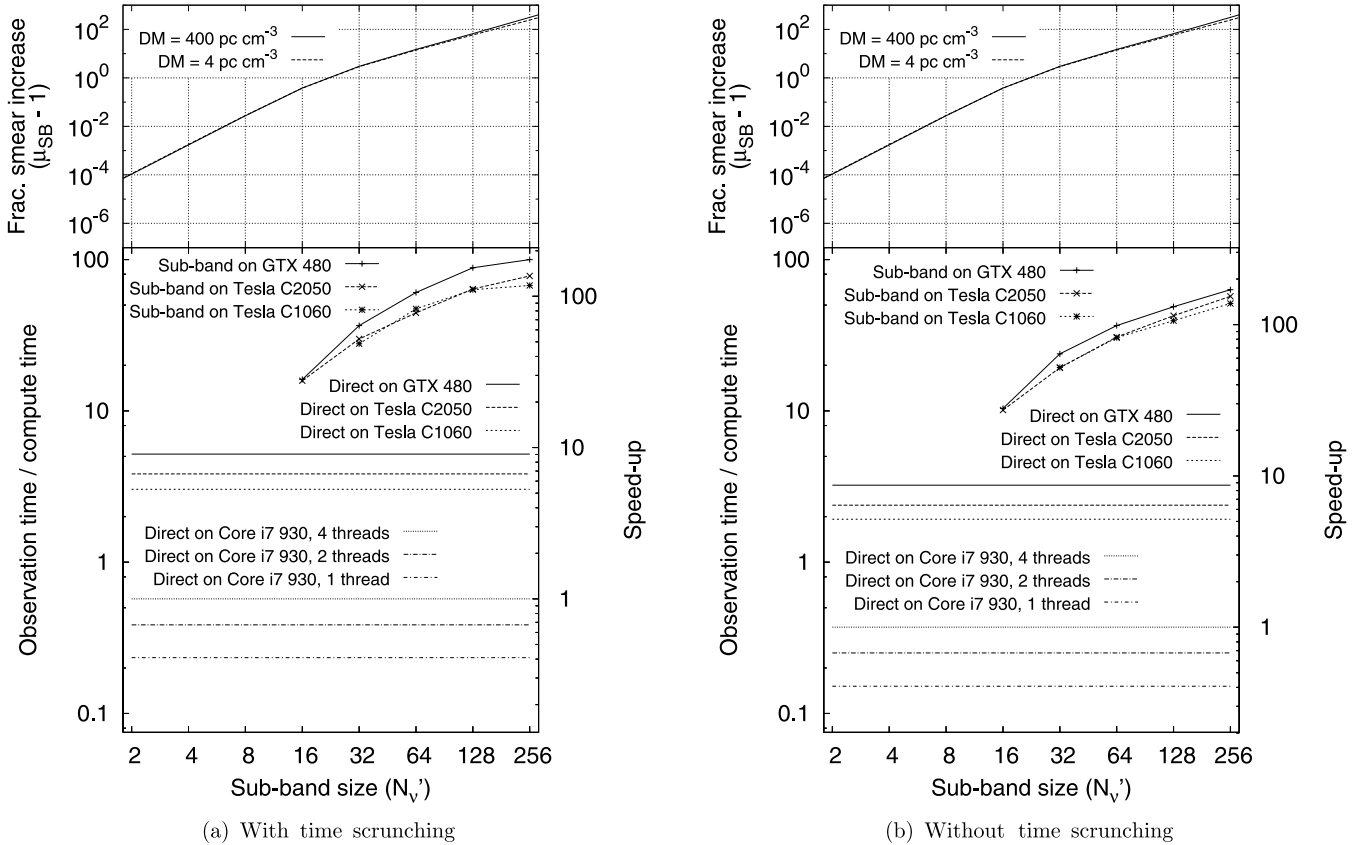


Figure 4. Upper panel: analytic upper-bound on signal degradation of a 40 μs pulse due to the sub-band algorithm as a function of the number of channels per sub-band (see equation B3). Lower panel: performance results for the direct and sub-band methods with (a) and without (b) the use of time-scrunching. See Fig. 3 caption for details. Note that it was not possible to run benchmarks of the sub-band code for $N'_v < 16$ due to memory constraints.

results were projected by scaling as appropriate for the computation of the full set of desired DMs with or without time-scrunching. Individual sections of code were timed separately to allow for different scaling behaviours.

Benchmarks were run on a variety of hardware configurations. CPU benchmarks were run on an Intel i7 930 quad-core CPU (Hyperthreading enabled). GPU benchmarks were run using version 3.2 of the CUDA toolkit on the pre-Fermi generation NVIDIA Tesla C1060 and the Fermi generation NVIDIA Tesla C2050 (error-correcting memory disabled) and GeForce GTX 480 GPUs. Hardware specifications of the GPUs' host machines varied, but were not considered to significantly impact performance measurements other than the copies between host and GPU memory. Benchmarks for these copy operations were averaged across the different machines.

Our derived performance results for the direct and piecewise linear tree codes are plotted in the lower panels of Fig. 3. The performance of the frequency-padded tree code corresponded to almost exactly half that of the piecewise linear tree code at a sub-band size of $N'_v = 1024$; these results were omitted from the plot for clarity.

Performance results for the sub-band dedispersion code are plotted in the lower panels of Fig. 4 along with the results of the direct code for comparison. Due to limits on memory use (see Section 4.2), benchmarks for $N'_v < 16$ were not possible.

Performance was measured by inserting calls to the Unix function `gettimeofday()` before and after relevant sections of code. Calls to the CUDA function `cudaThreadSynchronize()` were inserted where

necessary to ensure that asynchronous GPU functions had completed their execution prior to recording the time.

Several different sections of code were timed independently. These included pre- and post-processing steps (e.g. unpacking, transposing, scaling) and copies between host and GPU memory (in both directions), as well as the dedispersion kernels themselves. Disc I/O and portions of code whose execution time does not scale with the size of the input were not timed (see Section 6 for a discussion of the impact of disc I/O). Timing results represent the total execution time of all timed sections, including memory copies between the host and the device in the case of the GPU codes.

Each benchmark was run 101 times, from which the median execution time was chosen as the final measurement. Recorded uncertainties corresponded to the 5th and 95th percentiles; the error bars are too small to be seen in Figs 3 and 4 and were not plotted.

6 DISCUSSION

The lower panel of Fig. 3(a) shows a number of interesting performance trends. As expected, the slowest computation speeds come from the direct dedispersion code running on the CPU. Here, some scaling is achieved via the use of multiple cores, but the speed-up is limited to around 2.5 times when using all four. This is likely due to saturation of the available memory bandwidth.

Looking at the corresponding results on a GPU, a large performance advantage is clear. The GTX 480 achieves a nine times speed-up over the quad-core CPU, and even the last-generation Tesla C1060 manages a factor of 5. The fact that a single GPU is

Table 1. Summary of host↔GPU memory copy times.

Code	Copy time	Fraction of total time
Direct	0.62 s	<5 per cent
Tree	1.05 s	<30 per cent
Sub-band	0.62 s	10 to 65 per cent

able to compute the dedispersion transform in less than a third of the real observation time makes it an attractive option for real-time detection pipelines.

A further performance boost is seen in the transition to the tree algorithm. Computation speed is projected to exceed that of the direct code for almost all choices of N'_v , peaking at around three times at $N'_v = 64$. Performance is seen to scale approximately linearly for $N'_v < 32$, before peaking and then decreasing very slowly for $N'_v > 64$. This behaviour is explained by the relative contributions of the two stages of the computation. For small N'_v , the second, ‘sub-band combination’, stage dominates the total execution time [scaling as $O(1/N'_v)$]. At large N'_v the execution time of the second stage becomes small relative to the first, and scaling follows that of the basic tree algorithm [i.e. $O(\log N'_v)$].

The results of the sub-band algorithm in Fig. 4(a) also show a significant performance advantage over the direct algorithm. The computable benchmarks start at $N'_v=16$ with around the same performance as the tree code. From there, performance rapidly increases as the size of the sub-bands is increased, eventually tailing off around $N'_v=256$ with a speed-up of approximately 20 times over the direct code. At such high speeds, the time spent in the GPU kernel is less than the time spent transferring the data into and out of the GPU. The significance of this effect for each of the three algorithms is given in Table 1.

The results discussed so far have assumed the use of the time-scrunching technique during the dedispersion computation. If time-scrunching is *not* used, the projected performance results change significantly [see lower panels Figs 3(b) and 4(b)]. Without the use of time-scrunching, the direct dedispersion codes perform around 1.6 times slower, and similar results are seen for the sub-band code. The tree codes, however, are much more severely affected, and perform five times slower when time-scrunching is not employed. This striking result can be explained by the inflexibilities of the tree algorithm discussed in Section 3.1. At large dispersion measure, the direct algorithm allows one to sample DM space very thinly. The tree algorithms, however, do not – they will always compute DM trials at a fixed spacing (see equation 9). This means that the tree algorithms are effectively over-computing the problem, which leads to the erosion of their original theoretical performance advantage. The use of time-scrunching emulates the thin DM-space sampling of the direct code, and allows the tree codes to maintain an advantage.

While the piecewise linear tree code and the sub-band code are seen to provide significant speed-ups over the direct code, their performance leads come at the cost of introducing additional smearing into the dedispersed signal. Our analytic results for the magnitude of the smearing due to the tree code (upper panels of Fig. 3) show that for the chosen observing parameters, the total smear is expected to increase by less than 10 per cent for all $N'_v \leq 64$ at a DM of 1000 pc cm^{-3} . Given that peak performance of the tree code also corresponded to $N'_v = 64$, we conclude that this is the optimal choice of sub-band size for such observations.

The smearing introduced by the sub-band code (upper panels of Fig. 4) is significantly worse, increasing the signal degradation

by three orders of magnitude more than the tree code. Here, the total smear is expected to increase by around 40 per cent at $N'_v=16$, and at $N'_v=32$ the increase in smearing reaches 300 per cent. While these results are upper limits, it is unlikely that sub-band sizes of more than $N'_v=32$ will produce acceptable results in practical scenarios.

In contrast to the piecewise linear code, the frequency-padded tree code showed only a modest speed-up of around 1.5 times over the direct approach due to its doubling of the number of frequency channels. Given that the sub-band algorithm has a minimal impact on signal quality, we conclude that the frequency-padding technique is an inferior option.

It is also important to consider the *development cost* of the algorithms we have discussed. While the tree code has shown both high performance and accuracy, it is also considerably more complex than the other algorithms. The tree algorithm in its base form, as discussed in Section 3.1, is much less intuitive than the direct algorithm (e.g. the memory access patterns in Fig. 2). This fact alone makes implementation more difficult. The situation gets significantly worse when one must adapt the tree algorithm to work in practical scenarios, with quadratic dispersion curves and arbitrary DM trials. Here, the algorithm’s inflexibility makes implementation a daunting task. We note that our own implementations are as yet incomplete. By comparison, implementation of the direct code is relatively straightforward, and the sub-band code requires only minimal changes. Development time must play a role in any decision to use one dedispersion algorithm over another.

The three algorithms we have discussed each show relative strengths and weaknesses. The direct algorithm makes for a relatively straightforward move to the GPU architecture with no concerns regarding accuracy, and offers a speed-up of up to 10 times over an efficient CPU code. However, its performance is convincingly beaten by the tree and sub-band methods. The tree method is able to provide significantly better performance with only a minimal loss of signal quality; however, it comes with a high cost of development that may outweigh its advantages. Finally, the sub-band method combines excellent performance with an easy implementation, but is let down by the substantial smearing it introduces into the dedispersed signal. The optimal choice of algorithm will therefore depend on which factors are most important to a particular project. While there is no clear best choice among the three different algorithms, we emphasize that between the two hardware architectures the GPU clearly outperforms the CPU.

When comparing the use of a GPU to a CPU, it is interesting to note that our final GPU implementation of the direct dedispersion algorithm on a Fermi-class device is, relatively speaking, a simple code. While it was necessary in both the pre-Fermi GPU and multi-core CPU implementations to use non-trivial optimization techniques (e.g. texture memory, bit-packing, etc.), the optimal implementation on current-generation, Fermi, GPU hardware was also the simplest or ‘obvious’ implementation. This demonstrates how far the (now rather misnamed) graphics processing unit has come in its ability to act as a general-purpose processor.

In addition to the performance advantage offered by GPUs today, we expect our implementations of the dedispersion problem to scale well to future architectures with little to no code modification. The introduction of the current generation of GPU hardware brought with it both a significant performance increase and an equally significant reduction in programming complexity. We expect these trends to continue when the next generation of GPUs is released, and see a promising future for these architectures and the applications that make use of them.

While we have only discussed single-GPU implementations of dedispersion, it would in theory be a simple matter to make use of multiple GPUs, e.g. via time-division multiplexing of the input data or allocation of a sub-set of beams to each GPU. As long as the total execution time is dominated by the GPU dedispersion kernel, the effects of multiple GPUs within a machine sharing resources such as CPU cycles and PCI-Express bandwidth are expected to be negligible. However, as shown in Table 1, the tree and sub-band codes are in some circumstances so efficient that host↔device memory copy times become a significant fraction of the total run time. In these situations, the use of multiple GPUs within a single host machine may influence the overall performance due to reduced PCI-Express bandwidth.

Disc I/O is another factor that can contribute to the total execution time of a dedispersion process. Typical server-class machines have disc read/write speeds of only around 100 MB s^{-1} , while our GPU dedispersion codes are capable of producing 8-bit time series at well over twice this rate. If dedispersion is performed in an offline fashion, where time series are read from and written to disc before and after dedispersion, then it is likely that disc performance will become the bottle-neck. The use of multiple GPUs within a machine may exacerbate this effect. However, for real-time processing pipelines where data are kept in memory between operations, the dedispersion kernel can be expected to dominate the execution time. This is particularly important for transient search pipelines, where acceleration searching is not necessary and dedispersion is typically the most time-consuming operation.

The potential impact of limited PCI-Express bandwidth or disc I/O performance highlights the need to remember Amdahl's Law when considering further speed-ups in the dedispersion codes: the achievable speed-up is limited by the largest bottle-neck. The tree and sub-band codes are already on the verge of being dominated by the host↔device memory copies, meaning that further optimization of their kernels will provide diminishing returns. While disc and memory bandwidths will no-doubt continue to increase in the future, we expect the ratio of arithmetic performance to memory performance to get worse rather than better.

The application of GPUs to the problem of dedispersion has produced speed-ups of an order of magnitude. The implications of this result for current and future surveys are significant. Current projects often execute pulsar and transient search pipelines offline due to limited computational resources. This results in event detections being made long after the time of the events themselves, limiting analysis and confirmation power to what can be gleaned from archived data alone. A real-time detection pipeline, made possible by a GPU-powered dedispersion code, could instead trigger systems to record invaluable baseband data during significant events, or alert other observatories to perform follow-up observations over a range of wavelengths. Real-time detection capabilities will also be crucial for next-generation telescopes such as the Square Kilometre Array pathfinder programs ASKAP and MeerKAT. The use of GPUs promises significant reductions in the set-up and running costs of real-time pulsar and transient processing pipelines, and could be the enabling factor in the construction of ever larger systems in the future.

6.1 Comparison with other work

Magro et al. (2011) recently reported on a GPU code that could achieve very high (>100 times) speed-ups over the dedispersion routines in SIGPROC and PRESTO (Ransom 2001) whereas our work only finds improvements of factors of 10–30 over DEDISPERSE_ALL.

Table 2. Timing comparisons for direct GPU dedispersion of the ‘toy observation’ defined in Magro et al. (2011) ($\nu_c=610 \text{ MHz}$, $\text{BW}=20 \text{ MHz}$, $N_v=256$, $\Delta\tau=12.8 \mu\text{s}$, $N_{\text{DM}}=500$, $0 \leq \text{DM} < 60 \text{ pc cm}^{-3}$). All benchmarks were executed on a Tesla C1060 GPU.

Stage	Magro et al. (2011)	This work	Ratio
Corner turn	112 ms	7 ms	$16\times$
De-dispersion	4500 ms	1959 ms	$2.29\times$
GPU→CPU copy	220 ms	144 ms	$1.52\times$
Total	4832 ms	2110 ms	$2.29\times$

There are two key reasons for the apparent discrepancy in speed. First, the SIGPROC routine was never written to optimize performance but rather to produce reliable dedispersed data streams from a very large number of different backends. Inspection of the innermost loop reveals a conditional test that prohibits parallelization, and a two-dimensional array that is computationally expensive. Secondly, SIGPROC only produces one DM per file read, which is very inefficient. We believe that these factors explain the very large speed-ups reported by Magro et al. In our own benchmarks, we have found our CPU comparison code DEDISPERSE_ALL to be ~ 60 times faster than SIGPROC. For comparison, this puts our direct GPU code at ~ 300 times faster than SIGPROC when using the same Tesla C1060 model GPU as Magro et al.

Direct comparison of our GPU results with those of Magro et al. is difficult, as the details of the CPU code, the method of counting floating point operations and the observing parameters used in their performance plots are not clear. However, we have benchmarked our GPU code on the ‘toy observation’ presented in Section 5 of their paper. The execution times are compared in Table 2. Magro et al. did not specify the number of bits per sample used in their benchmark; we chose to use 8 bits/sample, but found no significant difference when using 32 bits/sample. We found our implementation of the direct dedispersion algorithm to be ~ 2.3 times faster than that reported in their work. Possible factors contributing to this difference include our use of texture memory, two-dimensional thread blocks and allocation of multiple samples per thread. The performance results of our implementation of the sub-band dedispersion algorithm generally agree with those of Magro et al., although the impact of the additional smearing is not quantified in their work.

In summary, we agree with Magro et al. that GPUs offer great promise in incoherent dedispersion. The benefit over that of CPUs is, however, closer to the ratio of their memory bandwidths (~ 10 times) than the factor of 100 reported in their paper, which relied on comparison with a non-optimized single-threaded CPU code.

6.2 Code availability

We have packaged our GPU implementation of the direct incoherent dedispersion algorithm into a C library that we make available to the community.⁹ The application programming interface (API) was modelled on that of the FFTW library,¹⁰ which was found to be a convenient fit. The library requires the NVIDIA CUDA Toolkit, but places no requirements on the host application, allowing easy integration into existing C/C++/FORTRAN, etc. codes. While the library

⁹ Our library and its source code are available at <http://dedisp.googlecode.com/>

¹⁰ <http://www.fftw.org>

currently uses the direct dedispersion algorithm, we may consider adding support for a tree or sub-band algorithm in future.

7 CONCLUSIONS

We have analysed the direct, tree and sub-band dedispersion algorithms and found all three to be good matches for massively parallel computing architectures such as GPUs. Implementations of the three algorithms were written for the current and previous generations of GPU hardware, with the more recent devices providing benefits in terms of both performance and ease of development. Timing results showed a nine times speed-up over a multi-core CPU when executing the direct dedispersion algorithm on a GPU. Using the tree algorithm with a piecewise linear approximation technique results in some additional smearing of the input signal, but was projected to provide a further three times speed-up at a very modest level of signal-loss. The sub-band method provides a means of obtaining even greater speed-ups, but imposes significant additional smearing on the dedispersed signal. These results have significant implications for current and future radio pulsar and transient surveys, and promise to dramatically lower the cost barrier to the deployment of real-time detection pipelines.

ACKNOWLEDGMENTS

We would like to thank Lina Levin and Willem Van Straten for very helpful discussions relating to pulsar searching, Mike Keith for valuable information regarding the tree dedispersion algorithm and Paul Coster for help in testing our dedispersion code. We would also like to thank the referee Scott Ransom for his very helpful comments and suggestions for the paper.

REFERENCES

- Barsdell B. R., Barnes D. G., Fluke C. J., 2010, MNRAS, 408, 1936
 Bhattacharya D., van den Heuvel E. P. J., 1991, Phys. Rep., 203, 1
 Cordes J. M., McLaughlin M. A., 2003, ApJ, 596, 1142
 Dodson R., Harris C., Pal S., Wayth R., 2010, in ISKAF2010 Science Meeting, Proc. Sci.
 Fluke C. J., Barnes D. G., Barsdell B. R., Hassan A. H., 2011, Pub. Astron. Soc. Australia, 28, 15
 Gaensler B. M., Madsen G. J., Chatterjee S., Mao S. A., 2008, Pub. Astron. Soc. Australia, 25, 184
 Keane E. F., Kramer M., Lyne A. G., Stappers B. W., McLaughlin M. A., 2011, MNRAS, 415, 3065
 Keith M. J. et al., 2010, MNRAS, 409, 619
 Lattimer J. M., Prakash M., 2004, Sci, 304, 536
 Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, Sci, 318, 777
 Lyne A. G. et al., 2004, Sci, 303, 1153
 McLaughlin M. A. et al., 2006, Nat, 439, 817
 Magro A., Karastergiou A., Salvini S., Mort B., Dulwich F., Zarb Adamis K., 2011, MNRAS, 417, 2642
 Manchester R. et al., 1996, MNRAS, 279, 1235
 Manchester R. et al., 2001, MNRAS, 328, 17
 Ransom S. M., 2001, PhD thesis, Harvard University
 Taylor J. H., 1974, A&AS, 15, 367

APPENDIX A: ERROR ANALYSIS FOR THE TREE DEDISPERSION ALGORITHM

Here we derive an expression for the maximum error introduced by the use of the piecewise linear tree dedispersion algorithm.

The deviation of a function $f(x)$ from a linear approximation between $x = x_0$ and $x = x_1$ is bounded by

$$\varepsilon_f \leq \frac{1}{8}(x_1 - x_0)^2 \max_{x_0 \leq x \leq x_1} \left| \frac{d^2}{dx^2} f(x) \right|, \quad (\text{A1})$$

which shows that the error is proportional to the square of the step size and the second derivative of the function. For the dedispersion problem, the second derivative of the delay function with respect to frequency is given by

$$\frac{\partial^2}{\partial \nu^2} \Delta t(d, \nu) = \text{DM}(d) \frac{d^2}{d\nu^2} \Delta T(\nu) \quad (\text{A2})$$

$$= 6 \text{DM}(d) k_{\text{DM}} \frac{\Delta \nu^2}{\nu_0^4} \left(1 + \frac{\Delta \nu}{\nu_0} \right)^{-4}, \quad (\text{A3})$$

which has greater magnitude at lower frequencies. Evaluating at the lowest frequency in the band, $\nu = N_\nu$, and substituting into equation (A1) along with the sub-band size N'_ν , one finds the error to be bounded by

$$t_{\text{tree}} \equiv \varepsilon_{\Delta t} \leq \frac{3}{4} \text{DM} \frac{k_{\text{DM}}}{\nu_0^2} \left(\frac{N'_\nu}{N_\nu} \right)^2 \frac{\lambda^2}{(1 + \lambda)^4}, \quad (\text{A4})$$

where $\lambda \equiv \frac{\Delta \nu}{\nu_0} N_\nu$ is a proxy for the *fractional bandwidth*, a measure of the width of the antenna band.

If the smearing as a result of using the direct algorithm is quantified as the effective width, W , of an observed pulse, then the piecewise linear tree algorithm is expected to produce a signal with an effective width of

$$W_{\text{tree}} = \sqrt{W^2 + t_{\text{tree}}^2}, \quad (\text{A5})$$

giving a relative smearing of

$$\mu_{\text{tree}} \equiv \frac{W_{\text{tree}}}{W} = \frac{\sqrt{W^2 + t_{\text{tree}}^2}}{W}. \quad (\text{A6})$$

In contrast to the use of a piecewise linear approximation, the use of a change of frequency coordinates (‘frequency padding’) to linearize the dispersion trails results in no additional sources of smear.

APPENDIX B: ERROR ANALYSIS FOR THE SUB-BAND DEDISPERSION ALGORITHM

Here we derive an expression for the maximum error introduced by the use of the sub-band dedispersion algorithm.

The smearing introduced into a dedispersed time series due to an approximation to the dispersion curve is bounded by the maximum temporal deviation of the approximation from the exact curve. The maximum change in delay across a sub-band is $\Delta t(\text{DM}, N_\nu) - \Delta t(\text{DM}, N_\nu - N'_\nu)$; the difference in this value between two nominal DMs then gives the smearing time:

$$t_{\text{SB}} \leq \Delta \text{DM}_{\text{nom}} [\Delta T(N_\nu) - \Delta T(N_\nu - N'_\nu)] \quad (\text{B1})$$

$$= N'_\nu \Delta \text{DM} \frac{k_{\text{DM}}}{\nu_0^2} \left[-2 \frac{N'_\nu}{N_\nu} \frac{\lambda}{(1 + \lambda)^3} + O\left(\frac{N'_\nu}{N_\nu}\right)^2 \right], \quad (\text{B2})$$

where the second form is obtained through Taylor expansion in powers of $\frac{N'_\nu}{N_\nu}$ around zero. Note that this derivation assumes the dispersion curve is approximated by aligning the

‘early’ edge of each sub-band. An alternative approach is to centre the sub-bands on the curve, which reduces the smearing by approximately two times but adds complexity to the implementation.

As with the tree algorithm, we can define the relative smearing of the sub-band algorithm with respect to the direct algorithm as

$$\mu_{\text{SB}} \equiv \frac{W_{\text{SB}}}{W} = \frac{\sqrt{W^2 + t_{\text{SB}}^2}}{W}, \quad (\text{B3})$$

where, as before, W is the effective width of an observed pulse after direct dedispersion.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.