

## Full length article

## Fast sampling from Wiener posteriors for image data with dataflow engines

N. Jeffrey<sup>a,b,\*</sup>, A.F. Heavens<sup>b</sup>, P.D. Fortio<sup>b</sup><sup>a</sup> Department of Physics and Astronomy, University College London, Gower Place, London WC1E 6BT, UK<sup>b</sup> Imperial Centre for Inference and Cosmology (ICIC), Imperial College London, London SW7 2AZ, UK

## HIGHLIGHTS

- Quantifying uncertainty of noisy and incomplete image data.
- Messenger field sampling of the Wiener posterior is suited to Dataflow Engines (DFE).
- Speed up > 11 with 8 DFEs (in 1U MPC-X box) compared to 1U server presenting 32 CPUs.
- This can improve dark matter maps and determination of cosmological parameters.

## ARTICLE INFO

## Article history:

Received 11 June 2018

Accepted 3 October 2018

Available online 24 October 2018

## Keywords:

Dataflow engines

Reconfigurable hardware

Data analysis

Bayesian statistics

MCMC

Wiener filter

## ABSTRACT

We use Dataflow Engines (DFE) to construct an efficient Wiener filter of noisy and incomplete image data, and to quickly draw probabilistic samples of the compatible true underlying images from the Wiener posterior. Dataflow computing is a powerful approach using reconfigurable hardware, which can be deeply pipelined and is intrinsically parallel. The unique Wiener-filtered image is the minimum-variance linear estimate of the true image (if the signal and noise covariances are known) and the most probable true image (if the signal and noise are Gaussian distributed). However, many images are compatible with the data with different probabilities, given by the analytic posterior probability distribution referred to as the Wiener posterior. The DFE code also draws large numbers of samples of true images from this posterior, which allows for further statistical analysis. Naive computation of the Wiener-filtered image is impractical for large datasets, as it scales as  $n^3$ , where  $n$  is the number of pixels. We use a messenger field algorithm, which is well suited to a DFE implementation, to draw samples from the Wiener posterior, that is, with the correct probability we draw samples of noiseless images that are compatible with the observed noisy image. The Wiener-filtered image can be obtained by a trivial modification of the algorithm. We demonstrate a lower bound on the speed-up, from drawing  $10^5$  samples of a  $128^2$  image, of  $11.3 \pm 0.8$  with 8 DFEs in a 1U MPC-X box when compared with a 1U server presenting 32 CPU threads. We also discuss a potential application in astronomy, to provide better dark matter maps and improved determination of the parameters of the Universe.

© 2018 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Dataflow computing has recently aided the significant acceleration of many computationally-intensive and data-intensive problems. This paper discusses the use of Dataflow Engines (DFEs) for sampling realisations of noise-free images from the Wiener posterior distribution given noisy and incomplete data, with particular applicability to astronomy and cosmology.

The Wiener filter (Wiener, 1949) is a useful statistical tool in many image analyses, as it is a minimum variance linear filter,

and moreover the filtered data are also the *maximum a posteriori* (MAP) values if the data have Gaussian signal and noise. To be more specific, if the covariance matrices of the noise and signal are known, then the Wiener filtered image has the smallest variance of any linear-filtered image. Mathematically it is straightforward to write down the expression for the Wiener-filtered image, and the covariance of compatible images, but evaluation is problematic as it involves the inversion of large matrices that are in general non-diagonal. As image datasets become larger, naive Wiener methods become unfeasible (requiring approximations such as re-binning to larger pixels or assuming white noise).

By using messenger field algorithms (described in Section 2.2) the Wiener image and posterior can be computed feasibly, with no need to simplify the existing algorithms. Furthermore, the

\* Corresponding author at: Department of Physics and Astronomy, University College London, Gower Place, London WC1E 6BT, UK.  
E-mail address: [niall.jeffrey.15@ucl.ac.uk](mailto:niall.jeffrey.15@ucl.ac.uk) (N. Jeffrey).

repeated operations inherent in drawing samples from the Wiener posterior lend themselves to efficient computation on DFEs, and we demonstrate that by a comparison with an implementation on multiple CPUs.

### 1.1. Data model

Although the typical applications of Wiener filters involve 2D image data, the formalism is general. In any case, we arrange the 2D pixel data as a list, and thus describe it by a data vector  $\mathbf{d}$ , and the true image is similarly described by a vector  $\mathbf{s}$ .

Our linear data model assumes that data  $\mathbf{d}$  and true signal  $\mathbf{s}$  are related by

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{n}. \quad (1)$$

where  $\mathbf{n}$  is random noise, and there is a known linear operator matrix  $\mathbf{A}$ , which in the simplest case is just the identity matrix.

The Wiener filter  $\mathbf{W}$  (Wiener 1949, Zaroubi et al. 1995) is given by

$$\mathbf{W} = \mathbf{S}\mathbf{A}^\dagger(\mathbf{A}\mathbf{S}\mathbf{A}^\dagger + \mathbf{N})^{-1}, \quad (2)$$

and the Wiener filtered solution, which is the minimum-variance linearly-obtained solution for the true image, is

$$\mathbf{s}_W = \mathbf{W}\mathbf{d}. \quad (3)$$

In these equations,  $\mathbf{S} = \langle \mathbf{s}\mathbf{s}^\dagger \rangle$  and  $\mathbf{N} = \langle \mathbf{n}\mathbf{n}^\dagger \rangle$  are the signal and noise covariance matrices respectively, which are assumed to be known, and we have assumed that  $\langle \mathbf{s} \rangle = \langle \mathbf{n} \rangle = 0$  for simplicity (this can easily be relaxed). The angle brackets indicate the expectation value, equal to the average over infinitely many realisations of the signal for ergodic fields. If, as we will assume, the pixel noise is uncorrelated, then  $\mathbf{N}$  is diagonal in pixel space. In addition to pixel noise, missing data in a given pixel can be incorporated into the Wiener filter by setting the pixel noise variance to infinity.

As mentioned in Section 2, the Wiener filter reconstruction,  $\mathbf{s}_W$ , is the linear minimum variance filter for a given  $\mathbf{S}$  and  $\mathbf{N}$  regardless of the statistical properties of either the signal or the noise. Note that the Wiener filtered image variance is biased low; e.g. high intensity pixels are suppressed. For Gaussian signal and noise, the Wiener filter additionally becomes the MAP estimate. In addition to computing the MAP estimate, for statistical purposes it is often useful to draw samples of maps, that are compatible with the data, with the appropriate probability. These can be used for subsequent statistical analysis of the true image, such as determining the uncertainty in a given pixel. This is discussed further in Section 2.1.

Calculation of the Wiener filter is challenging due to the inversion of covariance matrices, which may not be diagonal, and can become prohibitively time consuming for large images, especially when one notes that for an  $N \times N$  image, the matrices are  $N^2 \times N^2$  in size.

In some applications the signal is statistically homogeneous, leading to a diagonal signal covariance in the Fourier/harmonic domain, which leads to a route to a solution that does not involve the inversion of large non-diagonal matrices (Elsner and Wandelt, 2013). This is not trivial, since although independent noise has a diagonal covariance matrix in pixel space, it is not diagonal in harmonic space if the dataset has varying noise variance and is thus heteroscedastic. This situation automatically arises if there are missing data, but not only in this case. Therefore, in general there is no natural basis in which both the signal and noise covariance matrices are sparse. It is possible to take advantage of the bases in which the covariance matrices are sparse by using algorithms that employ so-called ‘‘messenger fields’’ (Elsner and Wandelt, 2013) to convey information between harmonic and pixel space.

The messenger field class of algorithms is highly suited to a Dataflow implementation. Using reconfigurable hardware accelerators rather than CPUs helps to deal iteratively with large volumes of data. DFEs have recently been successfully applied to a wide range of scientific problems, including geoscience (Gan et al., 2017), fluid-dynamics (Düben et al., 2015), artificial neural networks (Liang et al., 2018), quantum chemistry (Cooper et al., 2017), and genomics (Arram et al., 2015).

In Section 2, we describe the Wiener filter in a Bayesian framework, and show how messenger fields are used to draw samples from the Wiener posterior probability distribution. In Section 3, we describe Dataflow computing and present our implementation of the Wiener sampler. We present the results in Section 4. In Section 5, we describe our motivation for this work as an application to upcoming large cosmology surveys.

## 2. Theoretical background

### 2.1. Wiener posterior

For the linear model of Eq. (1), the Wiener filter, with  $\mathbf{W}$  given by Eq. (2), is a linear operator which minimises the variance

$$V = \langle (\mathbf{W}\mathbf{d} - \mathbf{s})^\dagger (\mathbf{W}\mathbf{d} - \mathbf{s}) \rangle. \quad (4)$$

From a different starting point, for the Wiener posterior, we begin by assuming a Gaussian likelihood for the pixel noise<sup>1</sup> (Jasche and Lavaux, 2015):

$$Pr(\mathbf{d}|\mathbf{s}, \mathbf{N}) = \frac{1}{\sqrt{(\det 2\pi\mathbf{N})}} \exp\left[-\frac{1}{2}(\mathbf{d} - \mathbf{A}\mathbf{s})^\dagger \mathbf{N}^{-1}(\mathbf{d} - \mathbf{A}\mathbf{s})\right]. \quad (5)$$

Assuming that the prior on the signal is that of a Gaussian random field,

$$Pr(\mathbf{s}|\mathbf{S}) = \frac{1}{\sqrt{(\det 2\pi\mathbf{S})}} \exp\left[-\frac{1}{2}\mathbf{s}^\dagger \mathbf{S}^{-1}\mathbf{s}\right], \quad (6)$$

then using Bayes’ theorem and the fact that  $Pr(\mathbf{d}|\mathbf{s}, \mathbf{s}, \mathbf{N}) = Pr(\mathbf{d}|\mathbf{s}, \mathbf{N})$ , the full Wiener posterior can be found:

$$\begin{aligned} Pr(\mathbf{s}|\mathbf{S}, \mathbf{N}, \mathbf{d}) &= \frac{Pr(\mathbf{d}|\mathbf{s}, \mathbf{N})Pr(\mathbf{s}|\mathbf{S}, \mathbf{N})}{Pr(\mathbf{d}|\mathbf{N})} \\ &= \frac{1}{\sqrt{(\det 2\pi\mathbf{S})}} \frac{1}{\sqrt{(\det 2\pi\mathbf{N})}} \\ &\times \exp\left[-\frac{1}{2}\mathbf{s}^\dagger \mathbf{S}^{-1}\mathbf{s} - \frac{1}{2}(\mathbf{d} - \mathbf{A}\mathbf{s})^\dagger \mathbf{N}^{-1}(\mathbf{d} - \mathbf{A}\mathbf{s})\right] \\ &\propto \exp\left[-\frac{1}{2}(\mathbf{s} - \mathbf{W}\mathbf{d})^\dagger (\mathbf{S}^{-1} + \mathbf{A}^\dagger \mathbf{N}^{-1} \mathbf{A})(\mathbf{s} - \mathbf{W}\mathbf{d})\right]. \end{aligned} \quad (7)$$

Here we see that the *maximum a posteriori* (MAP) solution is indeed that of the Wiener reconstruction,  $\mathbf{s} = \mathbf{W}\mathbf{d}$ .

If we can handle the large matrices, realisations of the true underlying signal image  $\mathbf{s}$  can be drawn from the posterior distribution  $Pr(\mathbf{s}|\mathbf{S}, \mathbf{d})$ . The expected mean of these samples is the Wiener-filtered image. Drawing samples from the Wiener posterior clearly also suffers from the need to invert large matrices with no natural sparse basis.

Progress can be made for signal images with statistical properties that are independent of pixel position  $\mathbf{x}$  (i.e. statistically homogeneous signals), for in this case, the Fourier transform of the image  $\mathbf{s}_{\mathbf{x}}$ ,

$$\tilde{s}_{\mathbf{k}} = \sum_{\mathbf{x}} s_{\mathbf{x}} e^{-i\mathbf{k}\cdot\mathbf{x}} \quad (8)$$

<sup>1</sup> We can also argue that if only the covariance and the mean are known, the Gaussian distribution is most appropriate to assume, as it is the maximum entropy distribution.

has a diagonal covariance matrix,

$$\langle \tilde{s}_{\mathbf{k}} \tilde{s}_{\mathbf{k}'}^* \rangle = P(k) \delta_{\mathbf{k}\mathbf{k}'} \quad (9)$$

and  $\delta_{\mathbf{k}\mathbf{k}'}$  is a Kronecker delta for the discrete 2D wavenumbers  $\mathbf{k}$  and  $\mathbf{k}'$ , and  $P(k)$  is the power spectrum, which depends only on the magnitude  $k \equiv |\mathbf{k}|$ . The covariance matrix  $\mathbf{S}$  for the signal is diagonal, with entries given by the appropriate  $P(k)$ .

## 2.2. Messenger fields

The messenger field approach splits the problem into two, performing some operations in harmonic space and some in pixel space, transferring the information using an extra field,  $\mathbf{t}$ , called the messenger field, whose covariance matrix is diagonal in both spaces. The method takes advantage of the diagonal signal covariance matrix in harmonic space and the diagonal noise covariance matrix in pixel space, such that no matrices need to be inverted in a basis in which they are dense.

This field is defined to have zero mean and a covariance matrix proportional to the identity matrix,  $\langle \mathbf{t}\mathbf{t}^\dagger \rangle \propto \mathbf{I}$ , which will always be diagonal in both harmonic and pixel spaces. The Markov Chain Monte Carlo (MCMC) algorithm used in [Jasche and Lavaux \(2015\)](#) is a method which uses the messenger field to draw samples from the Wiener posterior, without inversion of non-diagonal covariance matrices, requiring instead repeated Fourier transforms and inverse Fourier transforms. The algorithm is presented in Algorithm 1. In the limit of large numbers of iterations, this unconditionally converges to drawing samples from the desired distribution.

A sufficient number of samples from the Wiener posterior probability distribution can characterise the statistical properties of the underlying signal given some data.

**Algorithm 1:** Messenger Field Wiener Sampler: an iterative method to draw sample signal images from a Wiener posterior distribution using messenger fields ([Jasche and Lavaux, 2015](#)).

```

1: procedure SAMPLER
2:   for  $t_i$  in  $\mathbf{t}$ :
3:      $t_i = \mu_i^t + \sqrt{(\sigma_i^t)^2} G(0, 1)$ 
4:    $\hat{\mathbf{t}} = \mathcal{F}_{2D}(\mathbf{t})$ 
5:   for  $\hat{s}_k$  in  $\hat{\mathbf{S}}$ :
6:      $\hat{s}_k = \mu_k^s + \sqrt{(\sigma_k^s)^2} G(0, 1)$ 
7:    $\mathbf{s} = \mathcal{F}_{2D}^{-1}(\hat{\mathbf{S}})$ 
8:   Return  $\mathbf{s}$ 
9:   GOTO line 2
10: end procedure

```

Definitions:

- $\mu_i^t = \frac{T_i}{T_i A_i^2 + \bar{N}_i} A_i d_i + \frac{\bar{N}_i}{T_i A_i^2 + \bar{N}_i} s_i$  if  $A_i^2 > 0$
- $(\sigma_i^t)^2 = \frac{T_i \bar{N}_i}{T_i A_i^2 + \bar{N}_i}$  if  $A_i^2 > 0$
- $\mu_k^s = \frac{\hat{s}_k}{\hat{s}_k + T_k} \hat{t}_k$
- $(\sigma_k^s)^2 = \frac{\hat{s}_k \hat{t}_k}{\hat{s}_k + T_k}$
- $\mathbf{T} = \min((\mathbf{A}^{-1})^\dagger \mathbf{N} (\mathbf{A}^{-1})) \mathbf{I}$
- $\bar{\mathbf{N}} = \mathbf{N} - \mathbf{A}^\dagger \mathbf{T} \mathbf{A}$
- $G(0, 1)$  is a zero-mean Gaussian random variate with unit variance.
- $\mathcal{F}_{2D}$  is the 2D Fourier transform and  $\mathcal{F}_{2D}^{-1}$  its inverse.

By replacing the random variates in Algorithm 1 with zero ( $G(0, 1) \rightarrow 0$ ), the iteration outputs converge to the (unsampled) Wiener filter reconstruction (Eq. (2)). This was first described by [Elsner and Wandelt \(2013\)](#) in the first use of messenger fields. With this small change the code provides Wiener-filtered images, rather than samples from the Wiener posterior.

For each calculation of a Wiener filter or sample drawn from the Wiener posterior,  $\mathcal{O}(n^3)$  operations are required for a length  $n$  data vector using a naive approach. Using the messenger field algorithm, this reduces to  $\mathcal{O}(n \log n)$  for covariance matrices that are diagonal in their respective domains. The naive approach is bottlenecked by the matrix inversion and the messenger field approach is bottlenecked by the harmonic/Fourier transform.

## 3. DFE implementation

### 3.1. DFE system

The standard computing paradigm in the present day still follows the outline of the von Neumann model, often called Control Flow. In a standard setup a Central Processing Unit (CPU) carries out computational operations with data and instructions provided by memory, usually Random Access Memory (RAM). Data and instructions are iteratively passed between memory and CPU.

Dataflow Engines (DFEs) use reconfigurable hardware rather than CPUs to represent a static description of an algorithm with deep hardware pipelines consisting of a series of standard arithmetic and logic operations. DFEs do not need to continually get new instructions from the memory ([Pell et al., 2013](#)). They are therefore intrinsically parallel. The Wiener sampling problem described above has a high volume of data with highly deterministic computation (few “if” statements), so is well suited to DFEs.

Unlike standard CPU-based High Performance Computing (HPC) platforms, DFEs can be reconfigured on occasion to the need of a given algorithm or dataset. For the cost of an initial build time ( $\mathcal{O}(\text{hours})$ ), the speed and efficiency at runtime is improved. These systems allow greater flexibility with memory, data type, and clock frequency.

For example, higher clock frequencies can lead to shorter run time of the compute kernels instantiated on the DFE. This can yield faster execution if the algorithm is compute bound. However, for higher clock frequencies it becomes more difficult to build the reconfiguration bitstream, so the clock frequency can be chosen optimally for a given algorithm.

The CPU code for managing a DFE can be written in C or C++ and runs on a host (a traditional control flow machine). For the DFE, the software is written in Java-like code, which is compiled into the reconfiguration file for the hardware chip. This turns the DFE into a problem-specific hardware accelerator.

Once reconfigured, the DFE accepts data streaming and compute action calls launched by the host CPU code. A single DFE is a PCIe card that can either be available locally on a CPU server or be mounted in a Maxeler MPC-X: a CPU-free 1U server appliance hosting up to 8 DFEs, which is connected to host CPU servers by an infiniband network. Each DFE carries a chip with large amount of reconfigurable logic and on-chip resources (e.g. a Field-Programmable Gate Array, FPGA) with up to 96 GB of on-board DRAM storage. The MAX4 generation cards available to the authors are MAX4 Maia DFEs with Altera Stratix V FPGA and 48 GB of DRAM. Integration of CPU and DFE codes is done by the dedicated compiler as described in [Kos et al. \(2015\)](#).

Dataflow Engines allow user-friendly control over the features of the underlying hardware, so the hardware description can be optimally designed and built for the algorithm at hand. This can lead to large speed-ups at runtime compared to the same algorithm's implementation on a comparative CPU platform.

Time to complete a task is also only one metric of performance among other metrics. Lower clock frequencies mean that DFEs use less power than conventional CPU machines (Gan et al., 2017). Usually FPGAs use an order of magnitude less power than CPUs (Liang et al., 2018). In many applications, it is therefore more cost efficient to use DFEs as it allows more science per Watt.

Another commonly used and increasingly popular alternative to CPU hardware is graphics processing units (GPU), which gain acceleration for vectorised problems using “single instruction, multiple data” (SIMD) architectures and high-clock frequency (Liang et al., 2018). However, they are disadvantaged by their high energy cost. CPUs are more efficient than GPUs, and, as discussed, FPGAs are in turn more efficient than CPUs. GPUs additionally do not benefit from the flexibility that allows reconfigurable DFEs to tailor to a specific algorithm. Their hardware cannot be optimally designed for a given problem.

### 3.2. Implementation

We show the steps taken to generate a typical simulated dataset with the desired properties in Fig. 2. To simulate underlying signals  $\mathbf{s}$ , we generate realisations of square, two-dimensional images, which are in this case real, zero-mean, Gaussian random fields with known power spectra. The real and imaginary parts of  $s_{\mathbf{k}}$  are each drawn randomly from Gaussian distributions with variance  $P(k)/2$ , and reality of the signal is enforced by  $s_{\mathbf{k}}^* = s_{-\mathbf{k}}$ . We simulate square signal maps with  $128^2$  pixels.

The image, and therefore the vectors  $\mathbf{k}$  and  $\mathbf{x}$ , are two dimensional, so the transforms employ a 2D DFT. In practice the fast Fourier transform (FFT) algorithm (Cooley and Tukey, 1965) is used to evaluate the coefficients.

The datasets are generated according to the linear model of Eq. (1). For simplicity, we do not apply the linear operator (setting  $\mathbf{A} = \mathbf{I}$ ), though this could be included for a given application. The first panel of Fig. 2 shows an initial power spectrum,  $P(k)$ , from which we generate our real, Gaussian field as the signal map.

The noise is independent between pixels and is drawn from a Gaussian distribution where the noise variance varies across the data. We assume that the noise variance is known. We mask some of the pixels to represent missing data. The Wiener filter and the Wiener posterior treat the missing data as a special case of infinite noise. Infinite noise variance, in the region of the missing data, is set to be  $10^8$ , as an effective infinity.

On both CPU and DFE, we implement the messenger field algorithm (Algorithm 1) to draw samples of signal from the Wiener posterior (Eq. (7)), using 5 different datasets at each iteration. This reflects Alsing et al. (2017), where multiple chains were run in parallel to test convergence. In Fig. 1, the value of the same pixel in 5 independent chains with different initial values can be shown to converge after a sufficient number of iterations. The period during which the chains have not converged is known as burn-in, and using these samples reduces the influence of the initial starting point. Subsequent points are not converged immediately, therefore it is essential to have multiple chains, to check convergence and improve statistics.

On DFEs it is possible to instantiate fixed point and IEEE-like floating point arithmetic units of an arbitrarily chosen bitwidth, with greater flexibility beyond the standard options of single- or double-precision floating-point. Reducing the bitwidth of the number representation results in less reconfigurable logic spent on single arithmetic operations. This allows the user to instantiate more arithmetic units to fit the budget of reconfigurable space available on the chip, which may be used to implement more complex logic, or to replicate the computational pipeline; the latter reduces time for solution due to increased parallelism, but at the cost of reduced precision. In the implementation presented in this




Fig. 1. The absolute value of the same pixel at each iteration in 5 independent MCMC chains of the messenger field algorithm. The initial value of each pixel is different, to show convergence after a sufficient number of iterations.

paper, we use single-precision floating point format on both the DFE and on the CPU, to compare more easily the results.

The CPU code, written in C, uses a Box–Muller transformation to generate pairs of normally-distributed random variates for use in the algorithm. This custom-written implementation was shown to be consistently faster than the std C++ Gaussian random number generator in unit tests. Our implementation is slightly faster as we only ever generate one pair of zero mean and unit variance Gaussian random numbers at each iteration. The DFE uses the Gaussian random number generator from a dedicated dataflow library.<sup>2</sup> This small difference changes the overall time measurement little, as the fraction of time spent generating random numbers is small in this algorithm.

The 2D FFT from the FFTW3 package (Frigo and Johnson, 2005) was used for the CPU code, optimised with Advanced Vector Extensions (AVX2) available on the CPU hardware (see Section 4.2). A dedicated dataflow FFT library was used for the DFE.<sup>2</sup>

## 4. Results

### 4.1. Wiener posterior properties

As described in Section 2.2, samples from the Wiener posterior sampling algorithm without the random variates converge to the Wiener filter solution (Eq. (2)). In the left panel of Fig. 3 the Wiener filter reconstruction from this method is shown for the data generated in Fig. 2. By doing this we tested that the CPU and DFE outputs are identical up to computational precision.


A second test also provided the DFE with a vector of random Gaussian variates, shared with the CPU, where the output samples from the Wiener posterior were shown to be the same within computational precision.

In the centre panel of Fig. 3, the mean of the  $10^5$  samples from the Wiener posterior can be seen. By comparing to the Wiener filtered image in the left panel, one can see that the Wiener filtered solution is indistinguishable from the mean of the samples from the Wiener posterior, as expected. Due to sample variance, the mean of samples from Wiener posterior is not exactly equal to the Wiener filter, though for an infinitely large number of samples it would be.

In Fig. 3, the variance of the same  $10^5$  samples can be seen in the right panel. The variance in the region of missing data is high,

<sup>2</sup> MaxPower ([maxeler.com/mymaxeler](http://maxeler.com/mymaxeler) requires Maxeler account).







**Fig. 2.** This figure shows our data model, and gives an example realisation of a simulated dataset. We begin with a **Signal Power Spectrum**,  $P(k)$ , from which we generate a real, Gaussian random field as a **Signal Map**. We then take a **Noise Variance Map**, whose values vary across the data, from which we generate a **Noise Map** of Gaussian, independent pixel noise. The noise is added to the signal to generate the **Data No Mask**. We mask pixels representing missing data in **Data with Mask**.

as expected, but constrained by the signal covariance. The structure of the variance of the samples matches the structure of the noise variance map (see Fig. 2) as expected.

By drawing sufficient samples from the full posterior probability distribution, the code can characterise it very well, not just providing its mean and covariance.



**Fig. 3.** The input data, signal power spectrum and noise variance are shown in Fig. 2. *Left panel:* the left panel is the Wiener filter solution, where the random variates are not included in Algorithm 1. *Centre panel:* Mean of  $10^5$  samples from the Wiener posterior distribution evaluated using Algorithm 1. *Right panel:* Variance of the same set of samples.



**Fig. 4.** Time taken to run the sampling algorithm for a given number of iterations, where each iteration returns a sample from each of the 5 chains. Each data point is the mean of 10 runs (with the DFE data overhead removed), and  $1\sigma$  error bars with Akima (1970) interpolation for the error envelope. 32 CPU threads vs. 8 DFEs, run in parallel on MPC-X.

#### 4.2. CPU vs. DFE, speed

We compare the speed of the CPU+DFE implementation of the Wiener sampler (Algorithm 1) to the pure CPU implementation. Both were run on an Intel(R) Xeon(R) E5-2650 v2 @ 2.60 GHz server (2 sockets, 8 dual-thread cores per socket) presenting 32 CPU threads, which is connected to a MPC-X node at the STFC Hartree Centre. A single MPC-X box contains 8 MAX4 (Maia) DFEs. The clock frequency for the DFE implementation was chosen to be 200 MHz.

The sampler was timed for increasing number of iterations on both the CPU and the DFE, up to  $10^5$  iterations. Samples of 5 images are returned at each iteration. The time was measured from the CPU from the start to the end of the algorithm’s execution. At each number of iterations, the algorithm was repeated 10 times and the measured times averaged.

Fig. 4 shows the time to perform the algorithm for a 1U MPC-X with 8 DFEs against a 1U server presenting 32 CPU threads. The time as a function of number of iterations is linear for both CPU and DFE. The DFE has an initial overhead (with an average of 4.0 s) as the data is loaded onto the hardware, which is removed from the DFE time. Errors are obtained from 10 runs of the code. For the low run times, the DFE times have larger error-bars than the CPU, due to larger variance in the DFE data loading time; the relative effect of this decreases with longer running times.

In order to parallelise the problem, we run independent MCMC chains. We measure the time to generate a given number of samples by running the MCMC on 32 memory independent CPU

threads for the CPU-only code. For the CPU+DFE code, we measure the time to generate a given number of samples on 8 DFEs by splitting the work across 8 CPU threads.

In this work we have one CPU thread orchestrating one DFE. In future work we need to support an  $N:1$  ratio of  $N$  CPU threads served by a single DFE. This will help to utilise all the CPU computational capacity as well as all DFEs. Therefore, the speed-up of the CPU+DFE implementation in this work is a lower bound—here there is scope for considerable further acceleration.

From each independent MCMC chain some number of initial samples are unusable due to burn-in and are discarded. As each MCMC chain (run in parallel CPU threads) must discard the same number of initial samples, running 32 chains gives 32 times more unusable samples than a single chains with the same number of iterations. The 32 parallel CPUs will therefore have to discard four times more samples than 8 parallel DFE-accelerated CPU threads due to burn-in. This is also a reason why the time measurement from this MPC-X box parallel test should be interpreted as a lower bound on the potential speed up from DFEs.

We measure the lower bound on the parallel DFE speed-up to be  $11.3 \pm 0.8$ , where we have again used 10 time measurements to estimate the error.


### 5. Potential applications to cosmology

In this section we discuss some potential use cases in cosmology, although the algorithm and implementation are general and could be used in a number of contexts.

#### 5.1. Power spectrum inference

A common problem is to extract information from the power spectrum,  $P(k)$ , of an underlying field,  $\mathbf{s}$ , as defined in Eq. (9), and an extension of the DFE code can allow this. For a zero-mean Gaussian random field, the power spectrum contains all the statistical information that defines the field. The specific aim for power spectrum inference is to calculate the posterior probability distribution of the power spectrum given a set of data.

The standard model of cosmology predicts that the density field of the early universe will be a Gaussian random field, which persists for large cosmological scales in the late universe. Estimating the power spectrum is therefore a standard tool in many cosmological analyses with different data sets, including the early universe through Cosmic Microwave Background (CMB) radiation data (Planck Collaboration et al., 2016). A posterior probability distribution of the power spectrum of the density field in turn leads to posterior probability distributions for the cosmological parameters. These usually include, but are not limited to: the matter density  $\Omega_m$ , the Dark Energy density  $\Omega_\Lambda$ , the Dark Energy equation of state parameter  $w$ , and the Hubble parameter  $H_0$ .



**Fig. 5.** The Bayesian hierarchical forward model as described by Alsing et al. (2017) for signal image,  $\mathbf{s}$ , and power spectrum,  $P(k)$ , inference using the messenger field,  $\mathbf{t}$ . The work described in this paper uses Dataflow Engines can focus on the nodes of this network that do not include the power spectrum: the power spectrum is assumed and kept constant, and samples of the signal image are drawn.

The Bayesian hierarchical inference models described by Jasche and Lavaux (2015), Alsing et al. (2016), and Alsing et al. (2017) infer the posterior distributions of either the power spectrum or cosmological parameters in addition to samples from the field; see Fig. 5. With Gibbs sampling, samples of both the power spectrum and the image are drawn, keeping the other temporarily fixed. For a given power spectrum, large numbers of samples from the posterior probability distribution of the underlying signal  $\mathbf{s}$  can be drawn efficiently using Dataflow Engines, leading to better constraints on the power spectrum and hence on cosmological parameters.

### 5.2. Cosmological mass mapping

Due to the local curvature of spacetime by the matter, images of distant galaxies are deformed by the inhomogeneous matter distribution along the line of sight. This is called gravitational lensing. Any matter can contribute to the lensing effect, making it a direct probe of non-visible dark matter (Kaiser and Squires, 1993). Reconstructing this density field facilitates the study of the dark matter physics, its relationship with visible matter, and can provide novel approaches to extract additional cosmological information.

The model describing this process when in the linear regime, known as weak gravitational lensing, is fully described by our linear model in Eq. (1). The data  $\mathbf{d}$  are images where the pixel values are the mean of galaxy shapes within that pixel. The signal  $\mathbf{s}$  is a weighted, projected density field<sup>3</sup> in the foreground of the observed galaxies. The pixel noise, due to the intrinsic random galaxy shapes, is approximately Gaussian. The density field in the late universe on large cosmological scales is also approximately Gaussian.

From data with these properties, the large-scale density field from weak lensing shape measurements can be principally recovered with a Wiener filter. In Jeffrey et al. (2018) the messenger field Wiener filter algorithm is applied to Dark Energy Survey gravitational lensing data to generate a mass map image of the underlying density field. The Wiener filter method has also long been an established tool for reconstructing the underlying density field using only galaxy positions (Lahav et al., 1994), rather than using lensing data. Obtaining a large number of samples of the Wiener distribution, as is described in this work, then gives a posterior probability distribution of the density field in each pixel.

### 5.3. Future data requirements for cosmology

With current cosmic shear data,<sup>4</sup> Alsing et al. (2017) were able to use CPUs to generate samples from the posterior probability distributions of the underlying cosmic shear signal images and the

<sup>3</sup> The weighted, projected density field in mass mapping is called convergence and is denoted by  $\kappa$ .

<sup>4</sup> Cosmic shear is the spin-2 complex field manifested as the coherent distortion of galaxy shapes due to gravitational lensing. It is a function of a linear projection into 2D of the 3D density field.

power spectrum, using the Bayesian hierarchical model shown in Fig. 5. 10 chains were run in parallel to a length of  $10^5$  samples.

Current and future cosmic shear surveys DES (DES Collaboration et al., 2017), LSST (LSST Dark Energy Science Collaboration, 2012), and Euclid (Amendola et al., 2018) expect orders of magnitude of increase in data volume. The European Space Agency project Euclid expects to observe over  $10^9$  galaxies useable for cosmic shear, compared to  $\sim 3 \times 10^6$  with the CFHTLenS data used by Alsing et al. (2017). This leap in data size requires novel computational approaches to previously tractable problems. Here, Dataflow Engines can provide a solution.

## 6. Discussion

We have demonstrated a speed-up of at least  $11.3 \pm 0.8$  for generating  $10^5$  samples of the Wiener posterior of possible images compatible with an observed noisy image of  $128^2$  pixels, using 8 DFEs in a 1U MPC-X box and comparing with a 1U server presenting 32 CPU threads.

Future extensions could be to include the full Bayesian hierarchical model shown in Fig. 5, to further exploit the increased speed afforded to us by the Dataflow approach. This would lead to better constraints on the inferred cosmological parameters through samples of the power spectrum.

For data requirements of future cosmological surveys it would be useful to Wiener filter and draw samples of the Wiener posterior from data which have more than of  $128^2$  pixels per image. The image size in this work is constrained by the size of an FFT problem that fits within the fast FMEM on-chip memory ( $\sim 6$  MB). We expect that future versions of the dataflow FFT library will provide the option to use off-chip memory (48 GB) as an FFT buffer. We could then expect to be able to Fourier transform images of size  $2^{15} \times 2^{15}$ . This would increase the scientific applicability of a single DFE dramatically.

Implementing large scale Bayesian methods for cosmological parameter estimation on Dataflow Engines is a promising solution to the problem of increasingly large datasets from future surveys. This implementation of a Wiener sampler has broad application for inference or de-noising from any images or dataset with similar properties to those described here.

The authors plan to make the code public on [appgallery.maxeler.com](https://appgallery.maxeler.com) soon.

## Acknowledgements

We thank Pavel Burovskiy, Vitali Averbukh, Georgi Gaydadjev and Edward Edmondson for useful discussions and comments. NJ acknowledges support from the UK Science and Technology Research Council (STFC) Grant No. ST/M001334/1. We acknowledge use of Hartree Centre resources in this work. The STFC Hartree Centre is a research collaboratory in association with IBM providing High Performance Computing platforms funded by the UK's investment in e-Infrastructure. The Centre aims to develop and demonstrate next generation software, optimised to take advantage of the move towards exa-scale computing.

## References

- Akima, H., 1970. A new method of interpolation and smooth curve fitting based on local procedures. *J. ACM* 17 (4), 589–602.
- Alsing, J., Heavens, A., Jaffe, A.H., 2017. Cosmological parameters, shear maps and power spectra from CFHTLenS using Bayesian hierarchical inference. *Mon. Not. R. Astron. Soc.* 466, 3272–3292. doi:10.1093/mnras/stw3161, arXiv:1607.00008.
- Alsing, J., Heavens, A., Jaffe, A.H., Kiessling, A., Wandelt, B., Hoffmann, T., 2016. Hierarchical cosmic shear power spectrum inference. *Mon. Not. R. Astron. Soc.* 455, 4452–4466. doi:10.1093/mnras/stv2501, arXiv:1505.07840.
- Amendola, L., Appleby, S., Avgoustidis, A., Bacon, D., Baker, T., Baldi, M., Bartolo, N., Blanchard, A., Bonvin, C., Borgani, S., et al., (Euclid Theory Working Group), 2018. Cosmology and fundamental physics with the Euclid satellite. *Living Rev. Relativ.* 21, 2. doi:10.1007/s41114-017-0010-3, arXiv:1606.00180.
- Arram, J., Pflanzner, M., Kaplan, T., Luk, W., 2015. FPGA acceleration of reference-based compression for genomic data. In: *Field Programmable Technology (FPT), 2015 International Conference on*. IEEE, pp. 9–16.
- Cooley, J.W., Tukey, J.W., 1965. An algorithm for the machine calculation of complex Fourier series. *Math. Comp.* 19 (90), 297–301.
- Cooper, B., Girdlestone, S., Burovskiy, P., Gaydadjiev, G., Averbukh, V., Knowles, P.J., Luk, W., 2017. Quantum chemistry in dataflow: Density-fitting MP2. *J. Chem. Theory Comput.* 13 (11), 5265–5272, PMID: 29019679. doi:10.1021/acs.jctc.7b00649.
- DES Collaboration, Abbott, et al., 2017. Dark energy survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. ArXiv e-prints arXiv:1708.01530.
- Düben, P.D., Russell, F.P., Niu, X., Luk, W., Palmer, T., 2015. On the use of programmable hardware and reduced numerical precision in earth-system modeling. *J. Adv. Model. Earth Syst.* 7 (3), 1393–1408.
- Elsner, F., Wandelt, B.D., 2013. Efficient Wiener filtering without preconditioning. *Astron. Astrophys.* 549, A111. doi:10.1051/0004-6361/201220586, arXiv:1210.4931.
- Frigo, M., Johnson, S.G., 2005. The design and implementation of FFTW3. *Proc. IEEE* 93 (2), 216–231, Special issue on “Program Generation, Optimization, and Platform Adaptation”.
- Gan, L., Fu, H., Mencer, O., Luk, W., Yang, G., 2017. Data flow computing in geoscience applications. *Adv. Comput. Creat. Comput. DataFlow SuperComput.* 125–158. doi:10.1016/bs.adcom.2016.09.005.
- Jasche, J., Lavaux, G., 2015. Matrix-free large-scale Bayesian inference in cosmology. *Mon. Not. R. Astron. Soc.* 447, 1204–1212. doi:10.1093/mnras/stu2479, arXiv:1402.1763.
- Jeffrey, N., Abdalla, F.B., Lahav, O., Lanusse, F., Starck, J.L., Leonard, A., Kirk, D., Chang, C., Baxter, E., Kacprzak, T., Seitz, S., Vikram, V., Whiteway, L., et al., 2018. Improving weak lensing mass map reconstructions using Gaussian and sparsity priors: application to DES SV. *Mon. Not. R. Astron. Soc.* 479, 2871–2888. doi:10.1093/mnras/sty1252, arXiv:1801.08945.
- Kaiser, N., Squires, G., 1993. Mapping the dark matter with weak gravitational lensing. *Astrophys. J.* 404, 441–450. doi:10.1086/172297.
- Kos, A., Ranković, V., Tomažič, S., 2015. Sorting networks on maxeler dataflow supercomputing systems. *Adv. Comput. Dataflow Process.* 139–186. doi:10.1016/bs.adcom.2014.10.001.
- Lahav, O., Fisher, K.B., Hoffman, Y., Scharf, C.A., Zaroubi, S., 1994. Wiener reconstruction of all-sky galaxy surveys in spherical harmonics. *Astrophys. J. Lett.* 423, L93. doi:10.1086/187244, arXiv:astro-ph/9311059.
- Liang, S., Yin, S., Liu, L., Luk, W., Wei, S., 2018. FP-BNN: Binarized neural network on FPGA. *Neurocomputing* 275, 1072–1086.
- LSST Dark Energy Science Collaboration, 2012. Large Synoptic Survey Telescope: Dark Energy Science Collaboration. ArXiv e-prints arXiv:1211.0310.
- Pell, O., Mencer, O., Tsoi, K.H., Luk, W., 2013. Maximum performance computing with dataflow engines. In: *High-Performance Computing using FPGAs*. Springer, pp. 747–774.
- Planck Collaboration, Ade, P.A.R., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A.J., Barreiro, R.B., Bartlett, J.G., et al., 2016. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.* 594, A13. doi:10.1051/0004-6361/201525830, arXiv:1502.01589.
- Wiener, N., 1949. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*, vol. 7. MIT Press, Cambridge, MA.
- Zaroubi, S., Hoffman, Y., Fisher, K.B., Lahav, O., 1995. Wiener reconstruction of the large-scale structure. *Agron. J.* 449, 446. doi:10.1086/176070, arXiv:astro-ph/9410080.