Bifrost: A Python/C++ Framework for High-Throughput Stream Processing in Astronomy

Miles D. Cranmer^{1,2,9}, Benjamin R. Barsdell^{1,3}, Danny C. Price^{1,4}, Jayce Dowell⁵,
Hugh Garsden¹, Veronica Dike⁵, Tarraneh Eftekhari¹, Alexander M. Hegedus⁶, Joseph Malins⁵,
Kenneth S. Obenberger⁷, Frank Schinzel^{5,8}, Kevin Stovall^{5,8}, Gregory B. Taylor⁵
and Lincoln J. Greenhill¹

¹Harvard-Smithsonian Center for Astrophysics, 60 Garden Street
Cambridge, MA 02138, USA

²Department of Physics & McGill Space Institute, McGill University, Montreal
QC H3W 2C4, Canada

³NVIDIA Corporation, 2701 San Tomas Expressway, Santa Clara, CA 95050, USA

⁴University of California, Berkeley, Campbell Hall 339, Berkeley CA 94720, USA

⁵Department of Physics and Astronomy, University of New Mexico
Albuquerque, NM 87131, USA

⁶Department of Atmospheric, Oceanic and Space Sciences, University of Michigan
Ann Arbor, MI 48109, USA

⁷Air Force Research Laboratory Space Vehicles Directorate, 3550 Aberdeen Ave. SE
Kirtland AFB, NM 87117, USA

⁸National Radio Astronomy Observatory, Socorro, NM 87801, USA
⁹miles.cranmer@cfa.harvard.edu

Received 2017 August 1; Accepted 2017 September 12; Published 2017 October 12

Radio astronomy observatories with high throughput back end instruments require real-time data processing. While computing hardware continues to advance rapidly, development of real-time processing pipelines remains difficult and time-consuming, which can limit scientific productivity. Motivated by this, we have developed Bifrost: an open-source software framework for rapid pipeline development.^(a) Bifrost combines a high-level Python interface with highly efficient reconfigurable data transport and a library of computing blocks for CPU and GPU processing. The framework is generalizable, but initially it emphasizes the needs of high-throughput radio astronomy pipelines, such as the ability to process data buffers as if they were continuous streams, the capacity to partition processing into distinct data sequences (e.g. separate observations), and the ability to extract specific intervals from buffered data. Computing blocks in the library are designed for applications such as interferometry, pulsar dedispersion and timing, and transient search pipelines. We describe the design and implementation of the Bifrost framework and demonstrate its use as the backbone in the correlation and beamforming back end of the Long Wavelength Array (LWA) station in the Sevilleta National Wildlife Refuge, NM.

Keywords: Instrumentation, miscellaneous-methods, data analysis-methods, numerical-methods, observational.

⁹Corresponding author.

^a<https://github.com/ledatelescope/bifrost>.

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC-BY) License. Further distribution of this work is permitted, provided the original work is properly cited.

1. Introduction.....	3
1.1. Stream processing	3
1.2. Existing packages.....	3
2. Overview of Bifrost.....	4
2.1. Front end.....	4
2.2. Back end	5
2.3. Blocks and pipelines.....	5
2.3.1. Sequences	5
2.3.2. Spans.....	5
2.3.3. Frames	5
3. Details of Bifrost.....	6
3.1. Block and pipeline APIs	6
3.1.1. Pipeline generation and static visualization	6
3.1.2. Views.....	6
3.1.3. Block scopes.....	6
3.1.4. Guarantees	6
3.1.5. Block definitions	6
3.2. Metadata headers	8
3.3. Block execution and concurrency	9
3.3.1. Ctypes and the Python global interpreter lock (GIL)	9
3.3.2. CPU parallelization	9
3.3.3. GPU parallelization	9
3.4. Memory management and ring buffers.....	9
3.4.1. Sequences and spans	10
3.4.2. Ringlets	10
3.4.3. Random access.....	11
4. Bifrost libraries and tools	11
4.1. Some included blocks.....	11
4.2. Some included views	11
4.3. The Bifrost ndarray implementation.....	11
4.4. Bifrost <code>map</code>	11
4.5. Real-time logging	12
4.6. Tools	12
5. Comparison and Evaluation	12
5.1. Performance	12
5.2. Developing with Bifrost	14
5.2.1. Usability comparisons with a handwritten GPU pipeline.....	15
6. LWA Sevilleta Implementation	15
6.1. Packet capture implementation.....	16
6.2. TBN pipeline.....	16
6.3. DRX pipeline	17
7. Future Work.....	19
8. Conclusion	19

1. Introduction

Modern radio observatories are increasingly reliant on real-time processing to overcome data storage and analysis bottlenecks. A trend toward wider bandwidth digital systems, higher bitwidths of digitized signals, and more numerous antennas has driven improvements in both the capabilities and data throughput requirements of radio telescopes (Jones *et al.*, 2012; Hickish *et al.*, 2016; Price, 2016).

Consequently, astronomy survey data volumes are growing rapidly and are increasingly burdensome to analyze (Berriman & Groom, 2011). As an example, the Square Kilometer Array (SKA), scheduled to begin construction in 2018, will produce on the order of 10 exabytes per day in buffer memory (Quinn *et al.*, 2015) — about $5\times$ the 2016 global internet traffic (Cisco, 2016). The inability to store such volumes of data necessitates that they be processed and analyzed in real-time, putting significant demands on computing hardware and software.

As a result of computing demands, many-core accelerators — specifically Graphics Processing Units (GPUs) — have become commonplace in both high-performance computing^(b) and radio telescope installations (see Barsdell *et al.*, 2010), such as the Green Bank Telescope (GBT, Ransom *et al.*, 2009); the Canadian Hydrogen Intensity Mapping Experiment (CHIME, Bandura *et al.*, 2014; Recnik *et al.*, 2015); the Low-Frequency Array (LOFAR, Serylak *et al.*, 2013); the HIPSR system at the CSIRO Parkes Observatory (Price *et al.*, 2016); and the Large-Aperture Experiment to Detect the Dark Ages (LEDA, Kocz *et al.*, 2015).

However, developing software for GPUs often requires expertise in massively-parallel programming, knowledge of the memory hierarchy, and significant development effort. The noted lack of software specialists working in astronomy (discussed in Lupton *et al.* (2001)) complicates the development process more so. In a time when the community is seeing impressive growth in the availability of data, there is a lack of astronomers with an appropriate level of software development experience. In this environment, the creation of a high-level, easy-to-use development platforms is essential to ensure that the community can develop

efficient processing pipelines to manage astronomical data.

In this article, we introduce an answer to this problem: Bifrost, a high-throughput, GPU-enabled stream-processing framework, created to expedite the development of high-performance data processing pipelines in astronomy. This paper will present the core concepts of Bifrost such as the blocks and pipelines abstraction in Sec. 2, the technical details of Bifrost's implementation in Sec. 3, and descriptions of included Bifrost features and example pipelines in Sec. 4. Pipelines made with Bifrost are compared performance-wise, and also regarding usability, in Sec. 5. Finally, the implementation of Bifrost as the back end for the newest Long Wavelength Array (LWA) station, LWA-SV (see Taylor *et al.*, 2012, for the description of the first LWA station, LWA1), is outlined in Sec. 6.

1.1. Stream processing

The Bifrost framework is specifically designed to process streaming data in real-time; that is, to operate upon data with a potentially unlimited extent, such as the time stream of digitized voltages from a telescope, which has no clearly defined start time or end point. In Bifrost, streams of data are processed by creating pipelines: a sequence of modular processing elements, whose outputs feed into the inputs of the subsequent elements. When an algorithm is pipelined, and each processing element defines some finite minimum amount of data to work on, the pipeline can be easily set up for stream processing, by *streaming* the data through each element.

A real-time processing pipeline moves data through a series of different filters, measurements, and transforms, which may reduce raw data and flag for interesting occurrences such as bursts or other transients. Traditionally, such processing pipelines are built from scratch, by manually passing data between different systems. This development process is inefficient, making pipelines time-consuming to implement (demonstrated in Sec. 5.2.1). For this reason, processing frameworks have been adapted in some cases to ease the development cycle for real-time processing pipelines by abstracting the data flow away from the user.

1.2. Existing packages

The astronomy community has many software packages which encompass different applications

^bAs of November 2016, 40 of the top 50 supercomputers on the Green 500 energy efficiency ranking used accelerators.

and use cases: the Astrophysics Source Code Library^(c) (Allen et al., 2013) now boasts more than 1500 codes, and has seen the cumulative number of citations to its entries double every year since 2012 (see Allen et al., 2017, p. 3) to over 1000 in 2016. The astronomy community has a vast supply of codes, made for specific tasks, but the process of integrating these algorithms together into an efficient real-time pipeline remains a difficult problem, and there is a lack of a standard stream-processing framework for astronomical applications. Numerical computing libraries such as `numpy`^(d) (see Walt et al., 2011) provide high productivity, but lack the performance required for use in cutting-edge applications.

A few astronomy pipeline libraries have been developed, including PELICAN (Mort et al., 2015), PSRDADA^(e) and HASHPIPE.^(f) PELICAN provides a highly abstracted Application Program Interface (API) but is designed for static, quasi-real-time processing purposes. PSRDADA and HASHPIPE are designed for high-throughput stream processing, but suffer from rigidity and low-level C APIs.

There are also frameworks not specifically designed for astronomy, such as GStreamer^(g) which was used in the Laser Interferometer Gravitational-Wave Observatory (LIGO) merger detection pipeline (Messick et al., 2017). However, the use of frameworks specialized for other domains requires laborious porting of astronomy algorithms, file formats, and data types (e.g. support for complex numbers, low-bit integers, etc.).

Bifrost addresses these issues by combining a high-level pipeline interface in Python with a high-performance GPU-enabled back end that contains data structures and algorithms relevant to astronomy data processing. This combination results in a framework that is both highly productive and ready for cutting edge high-throughput applications. Bifrost also includes detailed documentation and example pipelines that are easily adapted to new applications. The use of Bifrost in a real-world application (the digital back end of LWA-SV, discussed in Sec. 6) also led to many refinements and enhancements of the framework.

2. Overview of Bifrost

Bifrost, currently versioned at v0.8.0, combines a high-level pipeline interface with a collection of high-performance data structures and algorithms relevant to astronomy data processing. The framework — written in C++ and Python, with GPU support enabled through NVIDIA’s CUDA platform — is divided into *front end* and *back end* components, which are described in the next two sections.

2.1. Front end

Bifrost’s front end is written as a set of Python modules which provide a high-level interface to the library. The design uses multiple layers of abstraction (all in Python):

- (1) Block/pipeline abstractions,
- (2) Pythonic^(h) wrappers of back end functionality, and
- (3) Direct wrappers of the back end C library.

At the highest level, the block/pipeline abstraction provides a concise and intuitive means of developing processing pipelines. Algorithms are described as blocks whose inputs and outputs can be connected to form pipelines. Use of rich metadata enables blocks to interoperate seamlessly. The block/pipeline abstraction is described in more detail in Sec. 2.3.

Beneath the block/pipeline abstraction sits a set of Pythonic wrappers around back end functionality. These provide a robust and convenient interface to the key components of Bifrost’s back end, such as its GPU-enabled ring buffer (a data structure emulating a wrap-around in memory), N -dimensional-array handling, and data processing algorithms. In addition to supporting the block/pipeline abstraction, this interface is used to implement a large suite of unit tests. An abstract graphic of this block/pipeline/ring buffer approach is shown in Fig. 1.

Finally, Bifrost’s back end C library is made directly accessible to Python via a `ctypes`-based wrapper orchestrated through the third-party `ctypesgen` package.⁽ⁱ⁾ This provides comprehensive

^c<http://ascl.net/>.

^d<https://www.numpy.org>.

^e<http://psrdada.sourceforge.net/>.

^f<https://github.com/david-macmahon/hashpipe>.

^g<https://gstreamer.freedesktop.org/>.

^hWe use *pythonic* here to mean code that takes full advantage of the features of the Python language rather than being one-to-one mappings of the C API.

ⁱ<https://github.com/davidjamesca/ctypesgen>.

Fig. 1. A cartoon illustrating some of the basic concepts of Bifrost: ring buffers (“rings”), blocks, and memory spaces. Each block: “Read”, “FFT” (a Fast Fourier Transform), and “Write”, lives in an independent CPU thread and applies a user-defined operation to a stream of data. Data are communicated from one block to the next through ring buffers — contiguous allocations of memory that emulate wrap-arounds from end to beginning. Ring buffers can be written to and read from different threads simultaneously, which enables blocks to operate asynchronously.

type-safe access to the C library, with an interface generated directly from the library headers.

2.2. Back end

Bifrost’s back end is written in C++ and exposes a C API. It is implemented as a collection of independent modules sharing a common core of type definitions and utility functions. The key modules (described in more detail in Sec. 4) are:

- Ring buffer — data structure for communication between threads in a pipeline.
- FFT — Fourier transform over any axes of an N -dimensional array.
- Linear algebra — matrix multiplication and related routines.
- FDMT — Fast Dispersion Measure Transform (the efficient dedispersion algorithm introduced in Zackay & Ofek, 2014).
- Quantize/unpack — conversions to/from 1–32 bit data types.
- Transpose-efficient data reordering.

- Map-application of arbitrary functions to N -dimensional arrays.
- User Datagram Protocol (UDP) capture/transmit-receive/send data across a network.

The back end API uses a single structure, `BFarray`, to represent N -dimensional arrays, providing a concise and consistent interface to many routines. The structure describes the type and shape of the data, along with other information such as the memory space in which the data are stored (e.g. system memory, CUDA device memory).

2.3. Blocks and pipelines

Bifrost’s high-level interface allows applications to be described as sets of *blocks* connected to form pipelines. Conceptually, blocks are black boxes that process streams of data, with inputs and outputs that can be connected. Blocks with no inputs or outputs are referred to as *source* and *sink* blocks, respectively, while those with both are referred to as *transform* blocks.

The flow of data through a pipeline is logically partitioned using the concepts of *sequences* and *spans*.

2.3.1. Sequences

Sequences define the boundaries between logically-independent intervals of data flow; e.g. when processing a set of files or observations, each file or observation is typically treated as a separate sequence. Sequences each have their own *header*, a collection of metadata describing the contents of the sequence.

2.3.2. Spans

Spans define the intervals of data within a sequence as they are processed by a block. Blocks “acquire” (read) or “reserve” (write) spans of data, process them, and then continue to the next span until the sequence is complete. Importantly, blocks may request spans of any extent and may progress through the sequence at any rate, with complete independence from other blocks in the pipeline.

2.3.3. Frames

Within a span, Bifrost defines a *frame* as a single element of the data stream, analogous to the frames of a video stream. Frames are treated as N -dimensional arrays, defined by a data type and shape. Other metadata can also be included, such as axis

labels, scales, and units, which are used to provide seamless interoperability between blocks.

3. Details of Bifrost

This section provides a look at how Bifrost implements the features and concepts discussed in Sec. 2. Section 3.1 outlines the API for Bifrost usage, and gives an overview of how to create a block within the Bifrost API. Section 3.2 describes how metadata is used in the framework, and Sec. 3.3 discusses how parallelism is achieved on both the CPU and GPU. Finally, Sec. 3.4 details how different memory spaces are used and how the Bifrost ring buffer is implemented.

3.1. Block and pipeline APIs

3.1.1. Pipeline generation and static visualization

In the Bifrost Python API, as exemplified in Fig. 2, the `Pipeline` class is used to set up a processing pipeline from a list of blocks and their connections. Blocks are initialized before the pipeline is executed. This prior initialization allows the user to store settings in the blocks before they are executed by the pipeline. The pipeline object then reads through all of the connections and initializes ring buffers between each connection with the appropriate memory space and allocation.

After the pipeline class performs this, the user can trigger a static visualization function: `Pipeline.dot_graph()`, which generates a graph of the pipeline using Graphviz. An example of this is shown in Fig. 3.

Stating `Pipeline.run()` executes the pipeline: the pipeline will create a thread for each block, which run until their blocks finish processing, until the user performs a keyboard interruption, or until a custom signal defined by `Pipeline.shutdown_on_signals()` is activated.

3.1.2. Views

In addition to blocks, Bifrost also defines *views*, which allow for simple metadata transformations that require no actual modification of the data stream. Examples include the `astype` view, which allows data to be explicitly reinterpreted with a different data type (analogous to `reinterpret_cast` in C++ or `numpy.ndarray.view` in Python), and the

`split_axes` and `merge_axes` views, which allow modifying the shape of data frames.

3.1.3. Block scopes

Bifrost's high-level pipeline/block API uses a hierarchical design that allows blocks to be nested inside *block scopes*. These constructs encapsulate parameters that are inherited by the blocks within; inherited parameters include the CPU core and GPU to which the blocks are bound. Block scopes also expose a *fuse* flag that configures blocks to share input, output, and temporary buffers; this can significantly reduce the amount of memory required to execute a pipeline, at the cost of the blocks being unable to execute in parallel.

An example of `block_scope` usage is given in Fig. 4.

3.1.4. Guarantees

By default, blocks are *guaranteed*, which means that the block that feeds them will (if necessary) wait for them to finish reading before writing new data. This results in blocks applying back-pressure to the pipeline if their execution is too slow. In some cases, it is desirable to prevent this by making blocks *unguaranteed*; the downside is that their input data will then be ungraciously overwritten if they are too slow, and they will end up *skipping* data. This data-guarantee functionality can be useful for diagnostic or monitoring blocks, or for decoupling different branches of a pipeline.

3.1.5. Block definitions

A block implementation must define two functions: `on_sequence` and `on_data`. The `on_sequence` function is called at the beginning of each sequence and is where the block interprets the input header(s) and generates its output header(s). The `on_data` function is called whenever a new set of input and output spans have been acquired or reserved and is where the block performs its processing operation. While more complex data processing patterns are possible, this simple abstraction is sufficient for most blocks, and greatly simplifies the process of implementing new blocks.

A block can also, optionally, redefine the function `on_skip`, which defines the behavior of the block when data is skipped in a ring if it is set to `guarantee=False` (see Sec. 3.1.4). By default, the output data is set to zero.

```

from copy import deepcopy

import bifrost as bf
import bifrost.blocks as blocks
import bifrost.views as views
import bifrost.pipeline as bfp
from bifrost.ndarray import copy_array

class CopyBlock(bfp.TransformBlock):#
    """Copy the input ring to output ring"""
    def __init__(self, iring, space=None, *args, **kwargs):
        super(CopyBlock, self).__init__(iring, *args, **kwargs)
        if space is None:
            space = self.iring.space
        self.orings = [self.create_ring(space=space)]
    def on_sequence(self, iseq):
        ohdr = deepcopy(iseq.header)
        return ohdr
    def on_data(self, ispan, ospan):
        copy_array(ospan.data, ispan.data)#

raw_data = blocks.read_wav(['audio_file.wav'], gulp_nframe=4096)

gpu_raw_data = CopyBlock(raw_data, space='cuda')#
chunked_data = views.split_axis(gpu_raw_data, 'time', 256, label='fine_time')
fft_output = blocks.fft(chunked_data, axes='fine_time', axis_labels='freq')
squared = blocks.detect(fft_output, mode='scalar')
transposed = blocks.transpose(squared, ['time', 'pol', 'freq'])#

host_transposed = blocks.copy(transposed, space='cuda_host')
quantized = blocks.quantize(host_transposed, 'i8')
blocks.write_sigproc(quantized)

pipeline = bf.get_default_pipeline()#
pipeline.shutdown_on_signals()
pipeline.run()#

```

Bifrost block definition

GPU processing part of pipeline

Start the pipeline

Fig. 2. An example of Bifrost Python syntax which creates a new block to copy memory from the GPU↔CPU, and then uses to implement a pipeline which channelizes an audio file on the GPU, writing it to a filterbank file.

Fig. 3. A graph generated by Bifrost for an example pipeline created in the Python API, the same one given in Fig. 2. Green color represents CUDA memory allocation, and blue color represents system allocation. This visualization uses Graphviz (Ellson *et al.*, 2001), which is wrapped by pyGraphviz.*

Note: *<https://github.com/pyGraphviz/pygraphviz>.

```
import bifrost as bf
...
with bfp.block_scope(core=2, gpu=0, fuse=True):
    data = bf.blocks.transpose(data, ['pol', 'freq', 'time'])
    data = bf.blocks.fdmf(data, max_dm=1000.)
    data = bf.blocks.transpose(data, ['time', 'pol', 'dispersion'])
...
```

Fig. 4. Code snippet which shows the use of a `block_scope` in a dedispersion pipeline. The three blocks inherit the `core` and `gpu` parameters such that they will all be executed on the same CPU core and GPU. Because the `fuse` flag is enabled, they will also share input and output buffers, which reduces memory usage but forces them to execute sequentially.

3.2. Metadata headers

Bifrost makes extensive use of metadata stored in sequence headers. These metadata give semantic meaning to the contents of a ring, and enable blocks to communicate seamlessly with little or no user intervention. While sequence headers may be defined any way the user wishes, so long as they are consistent enough to write blocks which can interpret those headers, Bifrost has a “standard” for headers. Abiding by this (informal) specification allows the user to make use of Bifrost’s pre-defined blocks with ease. A list and descriptions of the built-in Bifrost blocks are given in Sec. 4.

Bifrost’s standard header is as follows, which is inspired by several of the attribute names used for numpy arrays.

- ‘`name`’ — Name the sequence. This could be the filename that was used for data input, for instance.
- ‘`_tensor`’ — Dictionary of the following keys:
 - ‘`dtype`’ — Datatype of the data, in Bifrost format, e.g. ‘`i8`’ for 8-bit signed integers, or ‘`cf32`’ for 32-bit complex floating point numbers.
 - ‘`shape`’ — The shape of each frame of data. E.g. for each time step of a time stream, this would indicate a dimension of the frequency channels, and perhaps another dimension for

polarization. The axis which is “streamed”, such as time, should have a value of `-1`.

- ‘`labels`’ — List of labels for each of the axes in shape, in the same order. Some blocks may assume a specific name for a label; others will work on whatever axis you specify.
- ‘`units`’ — List of physical units each axis of a stream. For a time stream, the axis with `-1` size (the streamed axis), could have units of ‘`s`’ for seconds. Units are made possible using the `pint` library.^(j)
- ‘`scales`’ — List of 2-tuples which specify the offset and spacing between values along each axis (specified in the same order). The first element of a tuple is the starting offset of the axis, such as the starting Unix time of a time stream, while the second element is the spacing between elements along the axis, in units of the corresponding ‘`units`’ entry.

Any other dictionary values can be added after this, to specify observational parameters or anything else that is relevant to blocks downstream in the pipeline.

Figure 5 shows an example header for the `.wav`-reading Block.

```
{
  'frame_rate': 44100,
  'gulp_nframe': 4096,
  '_tensor': {
    'units': ['s', None],
    'dtype': 'i16',
    'shape': [-1, 2],
    'labels': ['time', 'pol'],
    'scales': [[0, 2.2675e-05], None]
  },
  'name': 'my_music.wav',
  'time_tag': 0
}
```

Fig. 5. Example header for a Bifrost sequence, which is in this case, output from `bifrost.blocks.read_wav`. Since a ring buffer in Bifrost does not have a particular data type or shape or size, the header for each sequence on that ring buffer must provide all information necessary to interpret the incoming stream of bytes. The header is read before the data is processed, which lets the block store various settings and options targetted by the information given in the header. Headers should describe at the very least, the data type and shape.

^j<https://github.com/hgrecco/pint>.

3.3. Block execution and concurrency

Bifrost involves multiple types of parallelization to take advantage of all processors available. The first level of parallelization — across different CPU cores — is handled by the threads, and by native Bifrost core-binding control. Each block executes in a separate thread. The second level of parallelization, which is used for GPU processing, works by calling CUDA kernels from within a Bifrost block, accessed via `ctypes`.

3.3.1. `ctypes` and the Python global interpreter lock (GIL)

Most Python implementations (e.g. the default, most commonly used CPython) are interpreters, meaning that the user's Python installation has an executable (the Python interpreter) which reads through the Python code at execution time, and executes machine instructions directly based on that executable's interpretation of each line of code. The GIL of Python is a lock on the interpreter so that it cannot execute two lines of code simultaneously. This means that parallelizing Python code is impossible unless the GIL is released.

This means that Bifrost must evoke external functionality to escape the GIL and execute truly multi-threaded pipelines.

`ctypes` is a native library for Python which allows the interpreter to execute arbitrary C functions from a dynamically loaded shared library. `ctypes` releases the GIL while the C function is executing, allowing it to run with true concurrency. Python's `numpy` library is an epitomic example of `ctypes`: most `numpy` functions release the GIL, so calling these functions allows parallelization within a pipeline.

`ctypes` wrappers for Bifrost's back-end are currently generated using the third-party `ctypesgen`^(k) library, which parses Bifrost's C headers and outputs Python code when the library is built.

3.3.2. CPU parallelization

The first level of parallelization is done across CPU cores. Python, where Bifrost pipelines are generated and controlled from, has a native module called `threading`. This module allows Bifrost to put each block into a separate thread, so that they can call algorithms without requiring these calls to be

sequential in code. This is represented visually in Fig. 6. As discussed in Sec. 3.3.1, the use of `ctypes` allows Bifrost to avoid Python's GIL, so that each block can be bound to a specific CPU core, and process data at the same time. An entire section of a pipeline can be bound to a CPU core using a block scope (see Sec. 3.1.3).

3.3.3. GPU parallelization

Fig. 6. A cartoon illustrating the two types of parallelization intrinsic to Bifrost — CPU, which maps process blocks to different CPU cores, and GPU, which allow CUDA functions to be called on ring buffers.

3.4. Memory management and ring buffers

A ring buffer (also called a circular buffer) is a data structure created to emulate a wrap-around in memory. A diagram of the ring buffer in Bifrost is given in Fig. 7. A ring buffer is a contiguous block of memory that pretends as though it has an arbitrary beginning and no end. When two processes are working on a contiguous block of memory that is not circular, with one process writing and the other reading, the writing process must either stop at the end of the block, allocate more memory, or go back to the start of the block of memory. With a ring buffer, the writing process wraps around to the start of the block of memory and continues to write seamlessly. This wrap-around is emulated with a region of memory called the “ghost region”, which is an extra allocation of memory at the end of a ring buffer that is synchronized with the beginning of the ring buffer. This circular interaction gives the term its name. Ring buffers have been used successfully as a standard data structure for buffer management in other astronomy stream processing software in the past, such as PSRDADA,^(l) PELICAN (Mort *et al.*, 2015), kotekan (Recnik *et al.*, 2015),

^k<https://github.com/davidjamesca/ctypesgen>.

^l<http://psrdada.sourceforge.net/>.

Fig. 7. Bifrost ring buffer implementation, showing ringlet sub-structure. Each ringlet is contiguous in memory and follows the previous ringlet. The ghosted and ghost regions are identical in memory content, and are synchronized with `memcpy`, so that a ring reader or writer can always read contiguous chunks of memory, even if its gulp size is not a divisor of the total ring size (and so partially overlaps the ghost region at the end of the ring buffer). The gray color in the diagram represents memory which has not been written to yet in this example.

and GUPPI (Ransom et al., 2009), which has evolved into the more general HASHPIPE.^(m)

In Bifrost, ring buffers form the connections between blocks in a directed graph. A ring has a single block that writes to it and can be read from multiple blocks at the same time. In addition to this, Bifrost ring buffers are dynamic: they can be resized by the user during processing to allocate more memory for storage.

The *gulp size* of a ring buffer refers to how many bytes of memory a block should read in from the ring buffer at once. The *buffer factor* of a ring buffer relates to the size of the ring buffer in memory as a multiplicative factor of the *gulp size*. By default, the buffer size is 4, which means that 4 times the *gulp_size* is the total memory of a ring buffer. As the *gulp_size* changes with a new incoming sequence, the ring buffer will resize accordingly. Different buffer sizes can also be set by the user.

Block fusion, which is discussed in Sec. 3.1.3, uses a buffer size of 1 for the rings between blocks, allowing things to operate sequentially. This is useful when memory is scarce, especially for GPU-only parts of a pipeline.

Bifrost ring buffers can be bound to one of the following memory spaces:

- `system` — host memory allocated with `aligned_alloc`
- `cuda_host` — pinned host memory allocated with `cudaHostAlloc`
- `cuda` — device memory allocated with `cudaMalloc`
- `cuda_managed` — unified memory allocated with `cudaMallocManaged`

Ring buffers store their memory binding in the “space” attribute: when a Bifrost block reads in a

ring buffer, they can check this attribute to trigger CPU or GPU processing depending on the space of the ring buffer. Bifrost blocks can then call the appropriate C functions on that data using `ctypes`.

Copying data from host memory to GPU memory is as easy as inserting a copy block into a pipeline, specifying the desired output space, such as `g_data = copy(h_data, space='cuda')`. These memory spaces are shown graphically in Fig. 1, and also in Fig. 6.

3.4.1. Sequences and spans

Bifrost ring buffers are implemented as contiguous memory allocations. Users of a ring buffer can request it to be resized in order to satisfy a given minimum total size as well as a maximum guaranteed-contiguous span. Blocks use this mechanism to specify the gulp size in which they will access memory in the ring.

Sequence metadata, including their starting and (eventually) ending offset, name, time tag and header, are stored with the ring. Reading is always done relative to the beginning of a sequence, and is bound by the beginning and end of the sequence (attempts to read outside the bound of a sequence will result in either no data or an error, depending on context).

3.4.2. Ringlets

To support data streams in which the temporal dimension is the fastest-changing dimension in memory, Bifrost rings use a concept dubbed “ringlets”. Ringlets act like a second dimension on a ring, emulating multiple individual rings running parallel to each other. The ability to store data in a time-fastest manner is important for the efficient execution of some algorithms. Moving data from time-slowest to/from time-fastest requires a transpose operation (the transpose block is a common feature of data-processing pipelines).

^m <https://github.com/david-macmahon/hashpipe>.

3.4.3. Random access

A common use-case in real-time data processing pipelines is to perform triggered dumps of buffered data. Bifrost ring buffers support random access to sequences by name or time tag which allows the user to easily inspect and save a specific portion of data that are still in memory.

4. Bifrost Libraries and Tools

4.1. Some included blocks

Bifrost currently has the following built-in blocks:

- **accumulate** — Accumulate frames one at a time.
- **copy** — Copy data, possibly between memory spaces (e.g. CPU → GPU).
- **detect** — Apply square-law detection to I/Q data to form polarization products.
- **fdmt** — Fast Dispersion Measure Transform (see [Zackay & Ofek \(2014\)](#)).
- **fft** — Fast Fourier Transform, using the CUFFT library.⁽ⁿ⁾
- **fftshift** — Shift the frequency center of an array along its axes.
- **print_header** — Print the header of each new incoming sequence to stdout.
- **transpose** — Reorder the axes of N -dimensional data.
- **read_sigproc/write_sigproc** — Read or write data in SIGPROC files.^(o)
- **read_wav/write_wav** — Read or write data in Wave files.
- **read_audio** — Read data from an audio stream using the portaudio library.^(p)
- **read_guppi** — Read data from GUPPI RAW^(q) files.
- **serialize** — Write data to a Bifrost-specific file structure, with headers as JavaScript Object Notation (JSON), and data streams as raw binary.
- **quantize** — Quantize data to a given smaller data type.
- **unpack** — Unpack packed data to a given larger data type.
- **reverse** — Reverse data along a specific axis.

4.2. Some included views

Bifrost views are all built upon a single function, **block_view**, which takes a block and a “header transform,” which is a function acting on a header, and applies that function to the block to modify the outgoing sequence headers. These views are easy to create, but Bifrost has a few built-in ones to speed up development:

- **rename_axis** — Change the label string on a specific axis.
- **split_axis** — Split an axis into chunks to form a new axis.
- **merge_axes** — The inverse of **split_axis**: merge two axes into one.
- **astype** — Change the datatype without changing the underlying data.
- **expand_dims** — Add an extra dimension to the shape.

4.3. The Bifrost ndarray implementation

Much of Bifrost is built around a feature of multi-dimensional array class, which inherits from **numpy.ndarray** for compatibility. Bifrost’s **ndarray** is the main part of the interface between ring buffers and the functions which compute on their data and provides an easy way to interface with both CUDA allocations of memory and **numpy** arrays.

Bifrost’s **ndarray** is cached with the C-struct: **BFarray**, which supports many datatypes and memory spaces, and is used to link with all of the C/C++/CUDA back end of Bifrost.

ndarray also has a simple function, **as_GPUArray**, which produces a **PyCUDA**^(r) GPU allocation object, for easy interfacing with **PyCUDA** functions ([Klöckner et al., 2012](#)).

4.4. Bifrost map

Bifrost’s **map** routine allows arbitrary functions to be applied to sets of ND-arrays using the GPU. The user provides a set of named arrays along with a code string describing how their elements are to be transformed, and the **map** routine generates and executes a corresponding GPU kernel.

The following example demonstrates the use of **map** from Python to multiply an array by 3:

```
bf.map('y = x * 3', {'x': x, 'y': y})
```

ⁿ<https://developer.nvidia.com/cufft>.

^o<http://sigproc.sourceforge.net/>.

^p<http://www.portaudio.com>.

^q<https://github.com/UCBerkeleySETI/breakthrough/blob/master/doc/RAW-File-Format.md>.

^r<https://github.com/inducer/pycuda>.

In this example, x and y are `bf.ndarray` objects whose data reside in GPU memory, and the operation is applied to each element of the arrays. More complex operations can be performed using explicit indexing, such as in this example that computes the outer product of two arrays:

```
bf.map('c(i,j) = a(i) * b(j)',
      axis_names=('i', 'j'),
      data={'a': a, 'b': b, 'c': c})
```

Other supported features include negative indexing, broadcasting, support for scalar parameters, and complex math.

Internally, the `map` routine uses the NVRTC^(s) library to just-in-time compile CUDA code into PTX, which is then passed to the CUDA driver for execution. Compiled kernels are cached in thread-local storage so that subsequent calls are fast. The routine is part of the Bifrost back end, and can be called from C or Python.

Due to its flexibility, `map` is used to implement some blocks in the Bifrost library, including `FFTShiftBlock`, `ReverseBlock`, `DetectBlock`, and `AccumulateBlock`. Using `map` avoids the need to write custom kernels for each of these tasks, which is a time-consuming and error-prone process.

4.5. Real-time logging

Bifrost implements lightweight real-time logging using a `proc`-like interface (`procfs` is a special Linux filesystem that lists system processes in a directory tree, and is used for `ps` and `top`). This logging captures various aspects of each block, e.g. CPU core binding and timing information, and is useful not only for debugging but also profiling and monitoring running pipelines. The log files are written to the Linux shared memory temporary file storage in a hierarchical structure. Similar to `proc`, each running pipeline is a directory named by its process ID. Within this pipeline directory, there are sub-directories, one for each block which implements logging. The log files are found within these block directories and store information in text files are *key:value* pairs. The log files are structured so that each file describes a particular aspect of the block. For example, the `bind` files contain how many and which CPUs or GPUs are being used by a block while the `perf` stores timing information on

how long it takes to load data from a ring buffer. This distribution of logging information makes it possible for the pipeline to update only a subset of files when new information is available and simplifies writing software that aggregates and analyzes the data.

4.6. Tools

Bifrost also comes with several pipeline-monitoring tools, which can be used to expedite development with Bifrost. These are found in the `tools` folder in the Bifrost repo. They make use of Bifrost's logging in Linux shared memory, i.e. `/dev/shm`. The tools included with Bifrost are:

- `like_ps.py` — Shows which blocks are running in a Bifrost pipeline, just like `ps`.
- `like_top.py` — Provides a per-block analysis of the processing load and timing of each of the elements in a Bifrost pipeline in quasi-realtime, just like `top`.
- `like_bmon.py` — Provides information on packet loss, which is useful for UDP packet capture pipelines.
- `pipeline2dot.py` — Creates graph files describing a running pipeline suitable for plotting.
- `getirq.py` — Displays system interrupt core bindings.
- `setirq.py` — Sets system interrupt core bindings.
- `getsiblings.py` — Shows which cores are hyperthread siblings.

5. Comparison and Evaluation

5.1. Performance

In Figs. 8 and 9, the performance of Bifrost is compared over various parameter settings with a serial pipeline implemented with the package `scikit-cuda`^(t) (Givon et al., 2015). Both `scikit-cuda` and Bifrost have Python frontends, wrapping a `cuFFT` back end, meaning that differences in performance should measure the speedup due to pipeline parallelism, or slowness due to Python overhead.

Due to the pipeline parallelism made available by Bifrost, performance increases are seen in Fig. 9, despite both packages using the same back end `cuFFT` kernel.

^s<http://docs.nvidia.com/cuda/nvrtc/index.html>.

^t<https://github.com/lebedov/scikit-cuda>.

Fig. 8. Performance timings of GPU pipelines for Bifrost and scikit-cuda. Each pixel is an individual benchmark. The pipeline for each package reads 32-bit floating point data from a binary file, copies the data to the GPU, then does a particular number of FFTs. The data is read, copied, and transformed in “gulps” of a specific size. Both pipelines are written in Python, have no tuning of settings, and are set to use the same file size, gulp sizes, number of copies, and number of batched FFTs. The times are averaged over 56 trials, and the significant figures reflect standard deviation about the average timing for each test. Explicit combined errors are shown in Fig. 9. The number inside each pixel represents the corresponding measurement for each benchmark.

In Fig. 8, a number of trends can be seen in the lower figure (Bifrost). As seen in each 5×6 pixel sub-grid, from left-to-right, a greater gulp size results in significantly increased performance of Bifrost, up to a certain point where the GPU becomes saturated, and pipeline parallelism loses effectiveness (the far right column of each sub-grid). This is because smaller gulp sizes result in more Python operations per kernel execution, and slow down a pipeline. From bottom-to-top in each sub-grid, more batched FFT’s result in more efficient computation due to extra parallelism. From the bottom row of sub-grids to the top row, a slight performance decrease is seen, simply because the GPU is performing twice as many FFT’s.

In the top plot of Fig. 8, the scikit-cuda pipeline shows more consistent, but much slower values on the majority of benchmarks. The consistency is due to the simpler operation of the library, which acts as

a thin wrapper on CUDA operations. This means there will be fewer potential Python operations per kernel execution, even on smaller gulp sizes. However, the relative slowness is due to the lack of pipeline parallelism.

Then, in Fig. 9, this speedup is shown across the same grid of benchmarks, highlighting the speedup Bifrost achieves when using a modest gulp size, and demonstrating the magnitude of speedup one can get on a GPU with pipeline parallelism.

The benchmarks described were measured on an Amazon Web Services (AWS) p2.xlarge instance, running Ubuntu 16.04, inside a docker container using the `nvidia/cuda`^(u) base instance, running CUDA 8.0. The p2.xlarge instances run a Tesla K80 GPU on an Intel Xeon E5-2686v4 CPU,

^u<https://github.com/NVIDIA/nvidia-docker>.

Fig. 9. Performance comparison of GPU pipelines for Bifrost and scikit-cuda, using the timings shown in Fig. 8. The bottom plot shows the combined error from each pixel's standard deviation about the average timing for both of the plots in Fig. 8.

with 61 GiB of memory. The benchmarks were performed using Bifrost version v0.8.0-benchmark, with all optimizations turned on. scikit-cuda was downloaded at git commit hash b32a679.

5.2. Developing with Bifrost

Bifrost provides a variety of interfaces for working with the library that help improve the productivity of writing new software. The most obvious example of this is the high-level Python blocks which wrap many of the intricacies of the pipeline mechanics. The library of existing blocks also provides many common operations needed for pipeline operation, both for real-time and off-line processing. These blocks also serve as a practical reference for the development of new blocks.

The low-level Python modules of Bifrost are suitable for advanced usage or pipelines where performance is more critical. These modules provide lightweight wrappers around the C++ library calls. These modules, when combined with direct access

of data through BFarray's, allow for application and algorithm-specific tuning. Finally, at the lowest level, the C++ library provides an interface for developing the more performance-critical applications.

In addition to increasing development productivity, Bifrost also provides a number of monitoring points for running pipelines. The monitoring points are useful for not only debugging and tuning pipelines but also for overall monitoring of production software. The monitoring points are implemented in the high-level Python blocks and provide information on the time required for acquiring data from a ring, the amount of time spent processing the data, and the time required to reserve a span in the output ring. For network related blocks, such as `bifrost.udp.capture`, other monitoring points are exposed that profile the packet statistics and can be used to look for packet loss. Bifrost also provides a collection of UNIX-style tools for polling the monitoring points. For example, `like_top.py` provides a per-block analysis of the processing load and timing

in quasi-realtime. For packet capture monitoring the `like_bmon.py` script provides information on any packet loss. These are outlined in Sec. 4.6. Similar to the high-level blocks, these tools also provide working examples for the development of application-specific monitoring software.

5.2.1. Usability comparisons with a handwritten GPU pipeline

A Bifrost implementation of a pipeline which channelizes data generated by the Green Bank Ultimate Pulsar Processing Instrument (GUPPI raw data)^(v) into filterbank data (from the SIGPROC project^(w)) is compared with unmodified handwritten legacy GPU pipeline code written as part of the Breakthrough Listen SETI (Search for Extraterrestrial Intelligence) project (see Worden *et al.*, 2017; MacMahon *et al.*, 2017) software stack aboard the GBT.

This comparison of syntax is meant to be a numerical way to quantify ease-of-development, and makes use of the following three metrics:

- Source Lines of Code (SLOC), which is the number of lines of code contain an executable statement (i.e. no comments or blanks). This measures the verbosity required to develop the pipeline, which is roughly correlated to the total development effort. This measure excludes the effort required for debugging and benchmarking code.
- Tokens, which is the total number of words, constants and operators appearing in the code. This is an alternative to SLOC in measuring the verbosity required to develop a pipeline, as it counts verbosity in a way that stresses not vertical length but rather the number of independent grammatical symbols in a statement. Because of how a programming habit such as breaking up each long function call across several lines might significantly impact the SLOC score, counting the tokens would only account for complexity in terms of number of words, not in terms of lines of code.
- Cyclomatic Complexity Number (CCN) (see McCabe, 1976), which, in some sense, measures the number of different paths the control flow can move through code (i.e. entering an *if* statement or not entering an *if* statement are two

Table 1. Comparison of a Bifrost version of a pipeline channelizing radio data into the frequency domain, against the original legacy pipeline. The comparison shown is only by code syntax, with three metrics that quantify development effort. All three metrics are inversely correlated with development effort.

	Files	SLOC	CCN	Tokens
Without Bifrost:	4	1329	265	12638
With Bifrost:	3	177	10	396

Notes: SLOC — Source Lines of Code (no blanks or comments).

CCN — Cyclomatic Complexity Number, accumulated over each function.

Tokens — Number of words, operators, or constants.

different paths through code). This provides a way to roughly quantify how difficult it will be to test some piece of functionality. CCN can complement the SLOC and tokens by quantifying the complexity of code rather than the verbosity of code, which can demonstrate how difficult it was to actually debug and test a segment of code. CCN is measured using the tool Lizard.^(x)

The scores on each metric for the pipeline with Bifrost and for the legacy code are given in Table 1. Bifrost outclasses the legacy C/CUDA code in all categories, having $7.5\times$ fewer source lines of code, $26.5\times$ smaller cyclomatic complexity number, and $31.9\times$ fewer number of tokens. If these metrics can be used together to quantify development effort, then it has been shown that Bifrost greatly eases the development process, and makes the final pipeline far less complicated than it had been in the handwritten version.

6. LWA Sevilleta Implementation

The LWA is a proposed low-frequency instrument sited in the southwest United States that will provide arc-second resolution for the frequency range of 10 to 88 MHz by cross-correlating the beamformed output of 53 stations. The first station of the LWA, LWA1 (see Taylor *et al.*, 2012) is co-located with the Karl G. Jansky Very Large Array (VLA) in the state of New Mexico while the second station, LWA-SV (LWA Sevilleta), is located ~ 75 km northeast on the Sevilleta National Wildlife Refuge.

^vSee <https://safe.nrao.edu/wiki/bin/view/CICADA/GUPPiUsersGuide>.

^wSee <http://sigproc.sourceforge.net/>.

^x<https://github.com/terryyin/lizard>.

Both stations consist of 256 dual-polarization dipole antennas spread over a 100 by 110 m area. LWA1 uses a custom Field-programmable gate array-based (FPGA) digital back end while LWA-SV uses a Bifrost-based back end running on commodity hardware.

The LWA-SV digital back end consists of 16 Reconfigurable Open Architecture Computing Hardware version 2 (ROACH2) boards and seven dual-GPU servers. The ROACH2 boards each contain two analog-to-digital converter (ADC) ADC16×256-8 digitizer cards and accept analog signals from 16 dual-polarization antennas. After digitization at 204.8 Ms/s, the signals are Fourier transformed into 4096 channels (25 kHz spectral and 40 μ s time resolution), sub-banded into three spectral tunings, and then sent to the GPU servers. On the servers, the data are further processed into the three main modes of operation for LWA-SV.

The first two modes, Transient Buffer Narrow-band (TBN) and Transient Buffer Full band (TBF), provide data from all 256 dipoles while the final mode, beamformed Data Receiver (DRX), provides time domain beamformed data. These modes are all implemented as Bifrost pipelines. The two all-dipole modes differ in the type of data returned and the duty cycle. TBN provides a 100 kHz bandwidth time series of complex voltages with a 100% duty cycle which is suitable for correlation. In contrast, TBF provides limited duration (< 1s) dumps of the raw “F-engine” (converts to frequency domain) data over a much larger bandwidth (~ 20 MHz) with a lower duty cycle ($\approx 0.5\%$). The TBF data are suitable for triggered acquisitions on external triggers, e.g. lightning or cosmic rays. The beamforming mode generates complex voltage timeseries data that are coherently summed to maximize the gain in a particular direction. This mode provides dual polarization data for two spectral tunings, each of ≈ 10 MHz, that are suitable for cross-correlation with LWA1.

The details of the Bifrost pipelines for all three modes are provided below.

6.1. Packet capture implementation

The ROACH2 boards transmit the frequency domain data to the servers using UDP packets. Each packet consists of a 128-bit header that specifies basic metadata, i.e. the source of the packet, frequency, and time. The data section consists of up to 256 channels from all 16 dual-polarization inputs

to the board represented at 4+4 complex signed integers. Since the size of the packets scale with bandwidth, each ROACH2 board provides a constant rate of 25 k packets/s.

The Bifrost `udp_capture` block is used to capture these packets, unpack the payload, and reorder the packets sent from the various ROACH2 boards. The capture code is optimized to use single instruction, multiple data (SIMD) operations and the Mellanox Messaging Accelerator (VMA) library to improve the efficiency of the capture. These, combined with interrupt binding and ethernet interface tuning, allow each GPU node to capture ≈ 13 Gb/s and 800 k packets/s. After capture and unpacking, the data are placed into a Bifrost ring for additional processing.

6.2. TBN pipeline

The TBN pipeline runs on six of the GPU servers and receives 200 kHz of data from three ROACH2 boards (48 dual-polarization antennas) at a data rate of about 18 MB/s. The pipeline consists of the four stages shown in Fig. 10: packet capture, unpacking, the “T-engine” (which converts frequency-domain channelized data back into a time stream), and a packetizer. The packet capture operation is described above. The unpacking stage converts the 4 + 4 bit complex data into a more standard 8 + 8 bit complex data using `bifrost.quantize`.^(y) This 8 + 8 bit data is then passed to the T-engine where it is filtered down to the 100 kHz needed for TBN.

Internally, the T-engine consists of several steps. First, the data is coerced to a 64-bit complex floating point representation and then reduced in bandwidth to 100 kHz. After the bandwidth reduction, the data are shifted along the channel axis in preparation for the inverse Fourier transform. The inverse transform is carried out on the CPU using the `scipy.fftpack.ifft`^(z) function. Once the data is in the time domain, it is combined with a complex mixer to provide sub-channel tuning capabilities and then quantized to 8+8 bit signed complex integers. The quantization is performed using the `bifrost.quantize` module. The repacked data are then sent to the packetizer where they are

^yFor historical reasons, the LWA-SV pipelines are implemented using the low-level Python wrappers of the C++ library functions.

^z<https://github.com/scipy/scipy>.

Fig. 10. TBN pipeline on the Advanced Digital Processor at LWA-SV generated by the `pipeline2dot.py` utility included with Bifrost. The ovals and diamonds represent network capture and transmit process, respectively. The labeled rectangles denote computing blocks. Each block is labeled with the operation name as well as the CPU core binding being used. Finally, the labeled arrows represent the ring buffers and the associated data type contained within each ring.

formatted as TBN data and sent to a data recorder. The output data rate is ≈ 18 MB/s per server.

6.3. DRX pipeline

Unlike TBN, the DRX pipeline captures data from all 16 ROACH2 boards. Six of the GPU servers runs two DRX pipelines, one for each spectral tuning of the beamformer, that capture 10.8 MHz of bandwidth at a data rate of ~ 800 MB/s. The DRX pipeline shown in Fig. 11 is also more complicated than TBN since it also provides the TBF mode. After the packet capture, the data is copied into a secondary ring using the `bifrost.memory` module to provide the TBF data buffer. The copy is done in a non-guaranteed manner in order to preserve the state of the TBF buffer when a dump is triggered. Since Bifrost supports branching pipelines, the captured data is also ingested by the GPU-based beamformer module which implements a phase-and-sum beamformer based on the `ls1.misc.beamformer` module from the LWA Software Library

Fig. 11. DRX beamforming pipeline on the Advanced Digital Processor at LWA-SV. The labeling is the same as in Fig. 10.

(Dowell *et al.*, 2012). In order to reduce the copy bandwidth to the GPU, the beamformer directly uses the 4+4 complex signed integer data.

After beamforming, the data are still in the frequency domain and need to be transformed into complex voltages. This is done using separate T-engine pipelines, Fig. 12, that run on the remaining

Fig. 12. DRX T-engine pipeline on the Advanced Digital Processor at LWA-SV. The labeling is the same as in Fig. 10.

Fig. 13. Detection of the pulsar PSR B0834+06 at LWA-SV, displayed using PRESTO (Ransom, 2001), the middle two panels show the pulse as a function of frequency and the best-fit dispersion measure for the observation. This detection offers evidence that the LWA-SV pipelines, written with Bifrost, are performing correctly.

seventh GPU server. The transition to the time domain is needed for better delay compensation, i.e. less than the F-engine dump time, when cross-correlating with other LWA stations. The time domain is also beneficial to other areas of research, particularly pulsar studies, for the higher temporal resolution. In order to move the data from the DRX pipelines to the T-engine pipelines, the data are packetized as UDP packets and transmitted using the `bifrost.udp_transmit` block at a rate of about 50 MB/s per server. In the T-engine pipeline, the data are aggregated from all six GPU servers and transformed into the time domain similar to TBN. However, due to the larger bandwidth of DRX, the inverse transform is carried out on the GPU using the `bifrost.fft` module. The time domain data are then quantized to 4+4 bit complex signed integers and transmitted to the data

recorder. The final data rate out of the T-engine is roughly 18 MB/s per spectral tuning.

To validate the Bifrost-based beamformer system, we observed the pulsar PSR B0834+06 with LWA-SV on 2017-03-18. Due to the transient nature of the pulsar, detection requires post-processing using a pulsar dedispersion and folding analysis pipeline; here, we use PRESTO^(aa) (Ransom, 2001). Successful detection of the pulsar requires system stability and for frequency and time metadata to be correctly propagated through the beamformer pipeline. Figure 13 shows a detection of the pulsar PSR B0834+06 made with one spectral tuning of the LWA-SV beamformer. The derived value for the dispersion measure (DM) as observed is $DM = 12.889$, with period $P_{\text{bary}} = 1273.7736(67)$ ms

^{aa}<https://github.com/scottransom/presto>.

($\chi^2_{\text{red}} = 7.862$ for $\text{DOF}_{\text{eff}} = 60.45$). The observed values are in reasonable agreement with those listed in the ATNF pulsar database^(bb) (see Manchester *et al.*, 2016, where $P_{\text{bary}} = 1273.768\,291\,578\,5(7)$ ms (Hobbs *et al.*, 2004), and $\text{DM} = 12.8640(4)$ (Stovall *et al.*, 2015)). Future work involves making this process into a real-time pulsar detection pipeline, incorporating the existing FDMT block in Bifrost.

7. Future Work

Bifrost is under active development and continues to accrue new features and functionality at a rapid pace. Our near-term plans include extending its capabilities to provide comprehensive support for correlation, beamforming, and pulsar and transient search pipelines. In the longer term, we aim to add real-time remote monitoring and control capabilities, accessible via web-based graphical interfaces. Such features would further reduce the cost and effort of developing and deploying new pipelines.

8. Conclusion

Bifrost is an open-source software framework for rapidly composing processing pipelines, built to emphasize the high-throughput needs of radio astronomers. The framework comes with a built-in block and view library which can be used to code many simple pipelines and which also provides functionality for interferometry, pulsar dedispersion and timing, and transient search pipelines. The framework also comes with a just-in-time compilation feature for quick development and deployment of CUDA kernels, and several tools for pipeline debugging and monitoring.

Modern radio telescopes rely on high-throughput processing to meet real-time requirements. Such pipelines are difficult to develop and debug. Bifrost makes the development process easier by abstracting away a flexible ring-buffer implementation, and providing a high-level API to facilitate rapid composition of pipelines through Python. Bifrost comes with full GPU support and can be used to exploit GPU processing without knowledge of CUDA. By exploiting pipeline parallelism, Bifrost gets significantly faster speeds for GPU computation versus regular serial pipelines created with other packages.

Bifrost can be downloaded from <https://github.com/ledatelescope/bifrost>. Code reference documentation and tutorials can be found from the links given in the repository.

Acknowledgments

Bifrost development was supported by NSF Grants OCI/1060067, OIA/1125087, AST/1139974 and AST-1106059.

Miles Cranmer would like to thank Joseph Schull and Anna Yang for their support of his research and for their establishment of a generous award program at McGill University.

Construction of the LWA has been supported by the Office of Naval Research under Contract N00014-07-C-0147 and by AFOSR. Operational support of the LWA-SV station is provided by the Air Force Research Laboratory.

References

- Allen, A., DuPrie, K., Berriman, B. *et al.* [2013] “Astrophysics source code library,” p. 387, <http://adsabs.harvard.edu/abs/2013ASPC..475..387A>, arXiv:1212.1916.
- Allen, A., DuPrie, K., Schmidt, J. & Ryan, P. W. [2017] “2016 annual report — Astrophysics source code library,” http://ascl.net/wordpress/wp-content/uploads/2015/03/ASCL_ann_report.2017_March_final.pdf.
- Bandura, K., Addison, G. E., Amiri, M. *et al.* [2014] “Canadian hydrogen intensity mapping experiment (CHIME) pathfinder,” (eprint: arXiv:1406.2288), p. 914522, doi:10.1117/12.2054950, <http://adsabs.harvard.edu/abs/2014SPIE.9145E..22B>.
- Barsdell, B. R., Barnes, D. G. & Fluke, C. J. [2010] *Mon. Not. R. Astron. Soc.* **408**, 1936, doi: 10.1111/j.1365-2966.2010.17257.x, <http://mnras.oxfordjournals.org/content/408/3/1936>.
- Berriman, G. B. & Groom, S. L. [2011] *Commun. ACM* **54**, 52, doi: 10.1145/2043174.2043190, <http://doi.acm.org/10.1145/2043174.2043190>.
- Cisco [2016] “The Zettabyte Era - Trends and Analysis”, <http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.html>.
- Dowell, J., Wood, D., Stovall, K. *et al.* [2012] *J. Astron. Instrum.* **1**, 1250006, doi: 10.1142/S2251171712500067, <http://adsabs.harvard.edu/abs/2012JAI.....150006D>.
- Ellson, J., Gansner, E., Koutsofios, L., North, S. C. & Woodhull, G. [2001] “Graphviz — Open source graph drawing tools,” *Graph Drawing* (Springer, Berlin, Heidelberg), p. 483, http://link.springer.com/chapter/10.1007/3-540-45848-4_57, doi: 10.1007/3-540-45848-4_57.
- Givon, L., Unterthiner, T., Erichson, N. B., Chiang, D. W., Larson, E., Pfister, L. A., Dieleman, S., Lee, G. R., van der Walt, S., Moldovan, T. M., Bastien, F., jz1536, Shi, X., Schlüter, J., Thomas, B., Capdevila, C. & Rubinsteyn, A.

^{bb}<http://www.atnf.csiro.au/research/pulsar/psrcat/>.

- [2015] “scikit-cuda 0.5.1”, doi:10.5281/zenodo.40565, <https://doi.org/10.5281/zenodo.40565>.
- Hickish, J., Abdurashidova, Z., Ali, Z., Buch, K. D., Chaudhari, S. C., Chen, H., Dexter, M., Domagalski, R. S., Ford, J., Foster, G., George, D., Greenberg, J., Greenhill, L., Isaacson, A., Jiang, H., Jones, G., Kapp, F., Kriel, H., Lacasse, R., Lutomirski, A., MacMahon, D., Manley, J., Martens, A., McCullough, R., Muley, M. V., New, W., Parsons, A., Price, D. C., Primiani, R. A., Ray, J., Siemion, A., Van Tonder, V., Vertatschitsch, L., Wagner, M., Weintraub, J. & Werthimer, D. [2016] *J. Astron. Instrum.* **05**, 1641001, doi: 10.1142/S2251171716410014, <http://www.worldscientific.com/doi/abs/10.1142/S2251171716410014>.
- Hobbs, G., Lyne, A., Kramer, M., Martin, C. & Jordan, C. [2004] *Mon. Not. R. Astron. Soc.* **353**, 1311.
- Jones, D. L. et al. [2012] “Big data challenges for large radio arrays”, *2012 IEEE Aerospace Conf.*, p. 1, doi: 10.1109/AERO.2012.6187090.
- Klöckner, A. et al. [2012] *Parallel Comput.* **38**, 157, doi: 10.1016/j.parco.2011.09.001, <http://dx.doi.org/10.1016/j.parco.2011.09.001>.
- Kocz, J. et al. [2015] *J. Astron. Instrum.* **4**, 1550003, doi: 10.1142/S2251171715500038, <http://adsabs.harvard.edu/abs/2015JAI.....450003K>.
- Lupton, R., Gunn, J. E., Ivezić, Z., Knapp, G. R. & Kent, S. [2001] “The SDSS imaging pipelines,” (eprint: arXiv:astro-ph/0101420), p. 269, <http://adsabs.harvard.edu/abs/2001ASPC..238..269L>.
- MacMahon, D. H. E., Price, D. C., Lebofsky, M. et al. [2017] arXiv:1707.06024 [astro-ph] <http://arxiv.org/abs/1707.06024>, arXiv: 1707.06024.
- Manchester, R. N., Hobbs, G. B., Teoh, A. & Hobbs, M. [2016] *VizieR Online Data Catalog*.
- McCabe, T. J. [1976] *IEEE Trans. Softw. Eng.* **SE-2**, 308, doi: 10.1109/TSE.1976.233837.
- Messick, C., Blackburn, K., Brady, P. et al. [2017] *Phys. Rev. D* **95**, 042001, doi: 10.1103/PhysRevD.95.042001, <http://adsabs.harvard.edu/abs/2017PhRvD..95d2001M>.
- Mort, B., Dulwich, F., Williams, C. & Salvini, S. [2015] *Astrophysics Source Code Library*, ascl:1507.003, <http://adsabs.harvard.edu/abs/2015ascl.soft07003M>.
- Price, D. C. [2016] *ArXiv e-prints* **1607**, arXiv:1607.03579, <http://adsabs.harvard.edu/abs/2016arXiv160703579P>.
- Price, D. C., Staveley-Smith, L., Bailes, M. et al. [2016] *J. Astron. Instrum.* **5**, 1641007, doi: 10.1142/S22511717 16410075, <http://adsabs.harvard.edu/abs/2016JAI.....54 1007P>.
- Quinn, P., Axelrod, T., Bird, I. et al. [2015] *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 147, <http://adsabs.harvard.edu/abs/2015aska.confE.147Q>.
- Ransom, S. M. [2001] “New search techniques for binary pulsars,” *Bull. Am. Astron. Soc.*, **33**, 1484, <http://adsabs.harvard.edu/abs/2001AAS...19911903R>.
- Ransom, S. M., Demorest, P., Ford, J. et al. [2009] “GUPPI: Green bank ultimate pulsar processing instrument,” *Am. Astron. Soc. id.* 605.08, <http://adsabs.harvard.edu/abs/2009AAS...21460508R>.
- Recnik, A., Bandura, K., Denman, N. et al. [2015] *ArXiv e-prints* **1503**, arXiv:1503.06189, <http://adsabs.harvard.edu/abs/2015arXiv150306189R>.
- Serylak, M., Karastergiou, A., Williams, C. et al. [2013] “Observations of transients and pulsars with LOFAR international stations and the ARTEMIS backend,” (eprint: arXiv:1210.4318), p. 492, doi:10.1017/S1743921312024623, <http://adsabs.harvard.edu/abs/2013IAUS..291..492S>.
- Stovall, K., Ray, P., Blythe, J., Dowell, J., Eftekhari, T. et al. [2015] *Astrophys. J.* **808**, 156.
- Taylor, G. B., Ellingson, S. W., Kassim, N. E., Craig, J., Dowell, J. et al. [2012] *J. Astron. Instrum.* **1**, 1250004, doi: 10.1142/S2251171712500043, <http://adsabs.harvard.edu/abs/2012JAI.....150004T>.
- Walt, S. V. D., Colbert, S. C. & Varoquaux, G. [2011] “The numpy array: A structure for efficient numerical computation,” *Comput. Sci. Eng.* **13**, 22, doi: 10.1109/MCSE.2011.37, <http://aip.scitation.org/doi/abs/10.1109/MCSE.2011.37>.
- Worden, S. P., Drew, J., Siemion, A. et al. [2017] *Acta Astronaut.* **139**, 98, doi: <http://dx.doi.org/10.1016/j.actaastro.2017.06.008>, <http://www.sciencedirect.com/science/article/pii/S0094576517303144>.
- Zackay, B. & Ofek, E. O. [2014] *ArXiv e-prints* **1411**, arXiv:1411.5373, <http://adsabs.harvard.edu/abs/2014arXiv1411.5373Z>.