

realfast: Real-time, Commensal Fast Transient Surveys with the Very Large Array

C. J. Law¹, G. C. Bower², S. Burke-Spolaor^{3,4,5}, B. J. Butler³, P. Demorest³, A. Halle¹, S. Khudikyan⁶, T. J. W. Lazio⁶,

M. Pokorny³, J. Robnett³, and M. P. Rupen⁷

¹Department of Astronomy and Radio Astronomy Lab, University of California, Berkeley, CA 94720, USA

Academia Sinica Institute of Astronomy and Astrophysics, 645 N. A'ohoku Place, Hilo, HI 96720, USA

³ National Radio Astronomy Observatory, Socorro, NM 87801, USA
⁴ Department of Physics and Astronomy, West Virginia University, Morgantown, WV 26506, USA

⁵ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

⁶ Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

⁷ National Research Council of Canada, Herzberg Astronomy and Astrophysics, Dominion Radio Astrophysical Observatory, P.O. Box 248, Penticton, BC V2A 6J9,

Canada

Received 2017 October 2; revised 2017 December 15; accepted 2018 February 5; published 2018 May 11

Abstract

Radio interferometers have the ability to precisely localize and better characterize the properties of sources. This ability is having a powerful impact on the study of fast radio transients, where a few milliseconds of data is enough to pinpoint a source at cosmological distances. However, recording interferometric data at millisecond cadence produces a terabyte-per-hour data stream that strains networks, computing systems, and archives. This challenge mirrors that of other domains of science, where the science scope is limited by the computational architecture as much as the physical processes at play. Here, we present a solution to this problem in the context of radio transients: *realfast*, a commensal, fast transient search system at the Jansky Very Large Array. *realfast* uses a novel architecture to distribute fast-sampled interferometric data to a 32-node, 64-GPU cluster for real-time imaging and transient detection. By detecting transients in situ, we can trigger the recording of data for those rare, brief instants when the event occurs and reduce the recorded data volume by a factor of 1000. This makes it possible to commensally search a data stream that would otherwise be impossible to record. This system will search for millisecond transients in more than 1000 hr of data per year, potentially localizing several Fast Radio Bursts, pulsars, and other sources of impulsive radio emission. We describe the science scope for *realfast*, the system design, expected outcomes, and ways in which real-time analysis can help in other fields of astrophysics.

Key words: instrumentation: interferometers – methods: data analysis – radio continuum: general – surveys

Supporting material: data behind figure

1. Introduction

The study of fast radio transients is important for a number of fundamental problems in astrophysics, including the discovery of new classes of object like the fast radio burst (FRB, Lorimer et al. 2007), detecting the hidden baryonic matter in the intergalactic medium (McQuinn 2014), and testing general relativity via the formation and evolution of compact objects (Kramer et al. 2006). Some of these systems are very rare, so large surveys are conducted for a chance to find a single source. However, even a single discovery can produce a great scientific payoff by offering a unique view of the physics at play.

Recent technological advances, exemplified by the development of the Jansky Very Large Array (VLA), have enabled access to fast transients through millisecond-timescale radio interferometric imaging. While large, single-dish telescopes have pioneered the field, interferometers will transform it through their ability to precisely localize, efficiently survey, and better characterize source properties. For example, the 100 m Green Bank Telescope and VLA have comparable sensitivities, but the VLA surveys 16 times more sky while localizing sources 10–300 times more precisely. In this way, images from an interferometer function as a high-resolution, multi-beam receiver. However, this powerful capability can only be exploited if data rates larger than 1 TB hr⁻¹ can be searched for hundreds to thousands of hours. A new paradigm

is needed to solve the "needle in the haystack" problem with such massive data streams.

realfast will solve this problem by integrating a compute cluster with the VLA for in situ analysis in support of real-time decision-making. *realfast* is a 32-node, 64-GPU compute cluster that will search images generated on timescales from 1 ms to 1 minute. In situ analysis allows triggered recording of a parallel data stream for later analysis and the rapid announcement of candidates for follow-up observing. By integrating with a high-speed, duplicate data stream of the VLA, this system will turn each observation into a fast transient survey, ultimately encompassing thousands of hours per year.

In Section 2, we describe the science potential of fast radio transients and the challenge in finding them. Section 3 presents the *realfast* transient search algorithm and system design. In Section 4, we discuss the development status and expected performance.

2. The Promise and Challenge of Fast Radio Transients

2.1. Fast Radio Transient Science

Radio transients are commonly considered "fast" if they last less than about 1 s. This timescale crudely separates slowly evolving synchrotron transients (Readhead 1994; Pietka et al. 2015) from fast, coherent emission processes. The best known class of fast radio transient is the pulsar, a rotating neutron star that is detected through brief intense pulses of light (Hewish et al. 1968; Hulse & Taylor 1975). On fast timescales, radio light propagation effects, such as dispersion⁸ can be measured. Propagation effects have been used to infer the distances and sizes of transients (Bower et al. 2014) and model Galactic structure (Cordes & Lazio 2002).

The newest class of fast radio transient is the FRB, a highly dispersed, millisecond transient. Roughly two dozen FRBs have now been detected by radio telescopes around the world.⁹ The radio signal is similar to that of a pulsar (impulsive, dispersed, scattered), so similar algorithms have been used to detect and study both classes of object. We recently used a prototype version of *realfast* at the VLA to precisely localize an FRB and identify its host galaxy at a redshift of 0.193 (Chatterjee et al. 2017; Tendulkar et al. 2017). This has shown that FRBs function like pulsars, but probe the intergalactic and circumgalactic media (Macquart & Koay 2013; Akahori et al. 2016; McQuinn 2014). With many FRB localizations at high redshift, they could even be used to constrain the dark energy equation of state (Zhou et al. 2014).

A similar approach can be applied in the study of Galactic radio transients. The first demonstration of a fast radio imaging search algorithm was made in a study of a rotating radio transients (RRATs), a potentially new class of pulsar (McLaughlin et al. 2006; Weltevrede et al. 2006; Lyne et al. 2009). We used fast imaging at the VLA to make the first interferometric localization of a RRAT that excluded its association with the optical/IR counterparts expected under the magnetar model (Law et al. 2012). Neutron stars in the Galactic center, globular clusters, or nearby galaxies (Macquart et al. 2010; Rubio-Herrera et al. 2013) may be detectable as millisecond transients. By localizing them to arcsecond precision, we can associate them with multiwavelength counterparts to infer a distance, find companions, or measure the properties of their local environment.

realfast can also be applied to the search for periodic radio emission. Single-dish telescopes currently do both pulsar discovery by searching for periodic emission and characterization by modeling pulse arrival times. Timing analysis provides precise astrometry, measures binary motion, and determines intrinsic spin-down rate of the pulsar. However, regular observations over at least one year are required to separate astrometric effects from spin-down and long-period binary motion (Ransom 2001; Deller et al. 2016). Modern pulsar surveys are discovering pulsars so rapidly that many cannot be completely followed-up with timing observations, potentially missing exotic systems with great potential to test general relativity (e.g., neutron star-black hole binary; Rafikov & Lai 2006; Faucher-Giguère & Loeb 2011). An arcsecond localization that comes via interferometric detection instantly limits the range of spin-down and binary properties by constraining the astrometric timing parameters. An arcsecond localization corresponds to timing precisions on the order of milliseconds ($\delta t \approx 1 \text{ au} \times \delta \theta$), which makes the VLA ideal for sorting ordinary pulsars from more exotic systems.



Figure 1. Schematic representation of interferometric and single-dish approaches to TD. The single-dish approach produces a single data stream (or a small number of data streams, for multi-beam receivers) with dimensions of time vs. frequency that are searched for impulsive signals. An interferometer can produce images for each integration that is searched for a source. The image shows our first demonstration of this concept, which blindly detected and localized RRAT J0628+0909 (Law et al. 2012). Since then, we have used this concept to blindly detect and localize FRB 121102 (Chatterjee et al. 2017; Law et al. 2017).

2.2. The Interferometric Search Problem

It is clear that arcsecond localization can dramatically expand the science scope of fast radio transients. Large, single-dish telescopes localize with a precision of several arcminutes, which is too poor to uniquely identify optical counterparts (Eftekhari & Berger 2017). Interferometers further improve on single-dish data by providing measurements of flux and spectra unaffected by uncertainties about source position within the telescope beam, as well as robust rejection of bad data.

Using interferometers to study fast transients presents a significant data analysis challenge. Millisecond sampling of interferometric data produces data at rates higher than 1 TB hr⁻¹, roughly 100 times larger than is typical (Law et al. 2015). The transient search requires searching each integration with hundreds of matched filters to capture impulses over a range of dispersions and pulse widths (see Figure 1). The traditional collect-and-analyze observing model is not tenable for this science case.

Fast transient science is growing just as interferometers are coming to dominate radio astronomy (Booth et al. 2009; DeBoer et al. 2009; Bandura et al. 2014). An approach taken by many newly designed interferometers is beamforming, in which the interferometric signal is reduced to a number of data streams corresponding to different sky positions within the interferometer's field of view. Each data stream has dimensions of time versus frequency, and thus can be processed independently as if it were the output of a singledish instrument (Staveley-Smith et al. 1996). The resulting data rate, and computational cost, of processing the entire field of view thus scales with the total number of pixels in the image, $N_{\rm pix} \propto (l/D)^2$, where l is the length of the longest interferometer baseline and D is the diameter of each dish. An alternate approach, more traditionally used for interferometric imaging, is to compute "visibilities," the correlation of electromagnetic waves between pairs of antennas. The visibility data stream encodes the same physical information as beamformed data, but is produced at a rate that scales only as the number of baselines (antenna pairs), $N_{bl} = N_{ant}(N_{ant} - 1)/2$, independent of N_{pix} . For sparse interferometers like the VLA, tens of thousands to tens of millions of beams are required to cover the full field of view, so the visibility data rate can be several orders of magnitude smaller than the corresponding

⁸ The cold plasma dispersion law is characterized by a delay of $\tau = 4.15 \text{ ms} (\text{DM}/\nu^2)$, where DM refers to the dispersion measure in units of pc cm⁻³ and ν is the observing frequency in units of GHz.

⁹ See the FRBCat at http://www.astronomy.swin.edu.au/pulsar/frbcat (Petroff et al. 2016).

THE ASTROPHYSICAL JOURNAL SUPPLEMENT SERIES, 236:8 (8pp), 2018 May



Figure 2. Transient search pipeline for *realfast*, known as *rfpipe*. The purple boxes slanted to the right denote sources of data and those slanted to the left denote output data products. The green boxes denote actions and the orange diamond shows a decision point. The dashed boxes show iteration loops over trial values of dispersion and sampling time. Other search algorithms, such as for periodic sources, can be seen as a modification of the innermost iteration loop. Candidate visualizations are generated for candidates above a second, higher threshold.

beamformed data rate. By performing processing operations such as dedispersion in the visibility domain, there is a proportional savings in computational cost. However, closely packed interferometers such as UTMOST (Caleb et al. 2016) or CHIME (Ng et al. 2017) may prefer a beamforming approach.

realfast addresses the challenges of fast interferometric data analysis in three ways:

- 1. In situ analysis. *realfast* integrates the transient search with the data recording system for real-time data triage. This will reduce the recorded data rate by a factor of 10^3 and allow rapid response to new transient detections (TD).
- 2. Commensal observing. In situ analysis also makes it possible to distribute a high-speed copy of all VLA observations to the transient search system.
- 3. Visibility-domain-processing and FFT-imaging. The computational cost of this approach is more efficient than beamforming for interferometers with high spatial resolution like the VLA.

Furthermore, *realfast* software and data will be open, making this novel capability at the world's most sensitive radio interferometer available to astronomers at large.

3. System Description

3.1. Algorithm

Since radio interferometers fully sample electromagnetic radiation, a wide variety of techniques exist to search for transients (Bannister & Cornwell 2011; Law & Bower 2012; Bhat et al. 2013; Obrocka et al. 2015). The principal challenges are driven by data rate, since data must be generated and searched on millisecond timescales for every pair of antennas in the array. Our approach has focused on interferometric imaging because it maintains the ideal sensitivity, is computationally efficient, and can easily be adapted to existing telescopes (Law et al. 2011, 2012).

Figure 2 shows the data flow through the search pipeline, as implemented in the Python package $rfpipe^{10}$ (Law 2017a). The pipeline breaks the problem of studying transient sources into two pieces: detect then analyze (see also Law et al. 2015). Using hierarchical decision-making, the pipeline dramatically reduces the data rate by saving information for a subset of the observation or only saving summary statistics.

Unlike traditional interferometric imaging (McMullin et al. 2007; Noordam & Smirnov 2012), millisecond-long integrations have modest sensitivity ($\sim 10 \text{ mJy}$ in $\sim 10 \text{ ms}$). At this sensitivity, there are roughly two persistent sources detectable per square degree (Becker et al. 1995) and computationally demanding imaging processes (Clark 1980; Schwab 1984) are not needed. The algorithm is further simplified by subtracting the mean visibility in time on timescales less the VLA fringe rate (~ 1 s). This removes constant sources, eliminates the need

¹⁰ *Rfpipe* is open source software and is built on other open source packages, including pwkit (Williams et al. 2017) and astropy (Astropy Collaboration et al. 2013). All *realfast* software can be found at https://github.com/realfastyla.



Figure 3. System diagram for VLA systems, including *realfast*. The VLA correlator includes WIDAR and the CBE, which together produce the raw data for both primary and commensal observers. The TD system runs the *realfast* search pipeline and is managed by a process that runs on the first node of the cluster. Output products and databases are managed by the portal, which also allows team team members to access and classify new data prior to archiving. Public access to raw data is managed by the NRAO archive, but smaller data products will be served by the *realfast* portal.

for source catalogs, and makes it possible to detect candidate transients by thresholding each image.

The bulk of the work in the search pipeline is in forming a matched filter for the transient signal in time and frequency. A Stokes I image is formed for each DM and time width and thresholded to search for point sources.¹¹ Since constant sources have been subtracted, all remaining point sources must be changing on timescales less than the fringe rate of a few seconds. Point sources with S/Ns greater than a threshold (typically 6σ) are saved to the candidate database and those with S/Ns greater than a higher threshold (typically 7σ) have candidate visualizations made for follow-up inspection.

3.2. Hardware and Software

A key feature of *realfast* is the integration of dedicated computing with the VLA system. The VLA generates data with a large, dedicated computing system called a correlator (see Figure 3). The first stage of correlation is done in Field Programmable Gate Arrays with a system called WIDAR, while the second stage is done in a commodity compute cluster known as the correlator backend (CBE; Perley et al. 2011). The fast transient search will be supported by the TD system, a 32-node, 64-GPU compute cluster on the same infiniband network.

Figure 3 shows how the data stream is duplicated for commensal processing. The standard data path typically gets averaged to a value defined by the primary user (\sim 1 s). This time averaging is done in two stages, first by WIDAR and then by the CBE. *realfast* will configure the correlator to perform less time averaging in WIDAR, which will preserve

millisecond-resolution visibilities in the CBE. Since this configuration produces more visibility data in the CBE and more are distributed over the network, both computational and network limits are important. A *realfast* goal is to generate 1 ms visibilities for the entire array to search data at a rate of 1.4 GB s^{-1} (see Table 1 for example observing configurations).

With fast-sampled visibility data available in the CBE, it is possible to export it over the infiniband network at rates up to 40 Gb s^{-1} . The distribution of visibilities is done by a new protocol for distributing and consuming VLA data via infiniband called *vys* and *vysmaw* (Pokorny & Law 2017). *Vys* is a lightweight messaging protocol that broadcasts the availability of data in CBE memory, while *vysmaw* can optionally read via remote memory access protocols of infiniband networks.

realfast will use *vys* to read data in short time segments. The segments will be the size of the inverse fringe rate (of the order of seconds), which is the largest size for which mean visibility subtraction is accurate. Each node will receive a segment in round robin fashion and process in slower than real time. Real-time processing is maintained over all nodes, but with a latency of (32 nodes $\times \sim$ seconds node⁻¹) $\approx \sim$ minutes. This latency can be reduced by adding processing breadth (more nodes or GPUs), since each segment is independent. Segments will partially overlap in time to keep sensitivity to dispersive delays of all scales at all times. To keep the overlap to less than 10% of the total data volume, the minimum data segment size should be at least 10 times the largest dispersive sweep of roughly 2 s (e.g., for a DM of 2000 pc cm⁻³, bandwidth of 256 MHz at frequency of 1.4 GHz).

The VLA observing system uses TCP to broadcast XML documents that define the configuration of an observation (e.g., antenna positions, correlator configuration, start time, etc.). A single node of the TD cluster will run an asynchronous process

 $[\]frac{11}{realfast}$ will also search for periodic signals by reading a time segment, transforming it to frequencies with a Fourier transform, and imaging each frequency to search for periodic signals.

Freq. Band	Freq. Center/Width (GHz/MHz)	Channels (ms)	Time Resolution	Data rate (MB s ⁻¹)	Sensitivity (10 σ ; Jy ms)	On-sky Detection Rate (hr burst $^{-1}$)
L ^a	1.4/256	256	5	290	0.5	500-1070
L	1.4/256	256	1	1400	1	380-630
S ^a	3.0/1024	256	5	290	0.2	190-2050
S (VLASS)	3.0/1500	1024	50	120	0.5	460-5340
C ^a	6.0/2048	256	5	290	0.13	230-10100
Х	10.0/4096	256	5	290	0.09	320-38450

 Table 1

 Example Observing Configurations

Notes. All configurations assume 27 antennas (351 baselines) and 2 polarizations.

^a This mode was commissioned in dedicated observations.

written in Python to capture these documents, parse them, and determine if they should be searched for transients. If so, then the pipeline state is defined from a set of preferences and observation metadata, as shown in Figure 2. The state is uniquely defined for a set of metadata and preferences, such that metadata and preferences can be associated with any given candidate to reliably reproduce it offline.

Each segment of data is scheduled for processing by the *distributed* library,¹² a part of the *dask* parallel computing framework. *Distributed* builds a directed acyclic graph for the pipeline to find the most efficient way to schedule the individual stages of the pipeline, given constraints on memory, processors, and GPU availability. When nodes are added or removed from the cluster, *distributed* automatically redistributes pipeline jobs. Many functions (e.g., dedispersion) are written as kernels with numba¹³ such that they can be run on multi-core CPUs. The most computationally demanding stage of processing is the 2d FFT for imaging, which is implemented on the GPU using *pycuda*.

3.3. Products and Services

Figures 2 and 3 summarize the products and services associated with *realfast*. The data products are:

- 1. Candidate statistics database. Statistical properties will be measured from the image and spectrum of each candidate above a threshold. These properties will summarize the properties of the candidate and in many cases help distinguish good astrophysical events from interference. Since these properties are easy to calculate and save, they are calculated at a relatively low threshold ($\sim 6\sigma$) and kept indefinitely.
- 2. Candidate plot. Initially, a candidate can only be verified by visual inspection of its image and spectrum. These will be calculated for candidates above a somewhat higher threshold ($\sim 7\sigma$), as they are larger and require more time to generate. They will be associated with candidates in the statistics database and kept indefinitely.
- 3. Raw data cutout. Each candidate with a plot will also be associated with raw visibility data for a small time window around the candidate. These are relatively large (several GB per candidate), so they will be deleted if they are clearly false positives. If they are worth keeping, they are archived in the standard NRAO archived and associated with a candidate in the statistics database.

- 4. Data quality database. Since *realfast* intends to throw away roughly 99.9% of the fast visibility data, we must keep a reliable measure of the amount of data searched and its quality. The data quality database will include regular measures of the fraction of configuration of the data, the amount of data flagged due to interference, and the sensitivity of a typical image.
- 5. Mock source database—The search pipeline will randomly insert mock transient sources into data as a test of its efficiency and sensitivity. The mock source is inserted after calibration, so it only tests the algorithm. This also makes it impossible to accidentally identify a mock transient as an astrophysical transient.

Of these products, the most important is the raw visibility data that will be created for each candidate transient. We require a full reanalysis to validate any candidate transient, so raw data, observing metadata, and calibration files must be saved for offline analysis. For a typical commensal data set, we expect a raw data rate of 1.4 GB s^{-1} and a size of several GB per candidate (see Figure 1). An example of this data product is available on dataverse , which shows nine visibility data files and associated calibration products for bursts detected from FRB 121102.

The *realfast* services are summarized in Figure 3 and example views are shown in Figure 4. They include:

- 1. Candidate search and review. We expect to generate candidate visualizations for a few tens of events per day, most of which will be false positives triggered by terrestrial interference and thermal noise. Raw data with good candidate events must be sent to the NRAO archive, while raw data from bad candidates must be deleted. Figure 4 shows that a web application will allow *realfast* team members to search and tag candidate transients to trigger this data management process. This is supported by a database indexed with *elasticsearch*.¹⁴
- 2. Candidate meta-analysis via Jupyterhub. The typical VLA observation lasts about an hour and is composed of many separate minute-scale segments with identical configuration. Figure 4 shows a Jupyter notebook we use to summarize candidate transients from a whole observation to provide context on the overall data quality of observation. We will use Jupyterhub (an extension of IPython; Pérez & Granger 2007) to serve these and other computational notebooks to the collaboration.

¹² See https://github.com/dask/distributed.

¹³ See http://numba.pydata.org.

¹⁴ See https://www.elastic.co.



Generate candidate and noise plots



Figure 4. (Top) a view of the searchable interface of the candidate database. The candidate plot shows the image and spectrum of the candidate, in this case a pulse from FRB 121102. (Bottom) a portion of a Jupyter notebook that summarizes candidate transients during an hour-long observation on MJD 57970. The notebook includes interactive visualizations and can support reanalysis on demand. A static version of this notebook is available as a supplement to this paper as "example_notebook.ipynb." The data used to create this figure are available.

 VOEvent server. Astrophysical transients will be of immediate interest to the community and will be distributed by the FRB VOEvent protocol (E. Petroff et al. 2018, in preparation; see https://github.com/ ebpetroff/FRB_VOEvent).

Of these services, the most important is the candidate review service. This service was inspired by the V-FASTR project that searched for fast radio transients in VLBA data (Wayth et al. 2011; Burke-Spolaor & Trott 2016). Initially, our team will manually inspect each candidate to remove false positives that do not need to be archived or analyzed further, but we ultimately aim to automate the process with classification based on candidate statistics. With experience, we will reduce the archived data rate by a factor of 1000 and speed up the time from discovery to confirmation.

4. Latest Developments

4.1. Status

As of late 2017, most major *realfast* software components have been either developed or prototyped. The *vys* protocol has been tested in commensal observations with integration times of 50 ms. This data can be distributed to a dedicated compute node that runs the transient search pipeline running on a prototype *realfast* server with two GPUs (nVidia GeForce GTX 1080) and an early version of a GPU-accelerated search pipeline. This pipeline detects bright pulsars and generates visualizations of candidate transients.

Much of the correlator and search pipeline has been prototyped in non-commensal mode over the last few years. Table 1 shows example *realfast* correlator configurations, including some that have been demonstrated in non-commensal (dedicated) observations for FRBs (Law et al. 2015; Chatterjee et al. 2017). Data was searched with an earlier version of the pipeline *rtpipe* (Law 2017b), which is designed for multi-core CPUs by dedicating each core to a set of integrations in shared memory. Profiling of the old pipeline shows that the dedispersion and imaging stages combined take about 1 ms per 512 \times 512 image. This is dominated by the FFT and scales as expected for larger images, while a smaller image processing time is dominated by dedispersion. Processing time scales sublinearly with the number of cores, suggesting that memory movement reduces efficiency by as much as 50% for 20-core CPUs. The new pipeline will improve on this by using GPUs, a different job scheduling scheme, and a different CPU core parallelization scheme.

Future development will focus on testing faster commensal data streams, integrating with the VLA observing system and archiving, and improving data services in response to test observations. Our goal is to perform the first end-to-end commensal science observing by the middle of 2018 on a 10% scale GPU cluster. The full GPU cluster will be purchased and installed by the third quarter of 2018.

4.2. Performance Goals

As a commensal project, the *realfast* science potential depends on the distribution of VLA observing time. In a typical year, the VLA observes for about 6000 hr at frequencies from 1 to 50 GHz. The antenna configuration is changed between four possible states from "A" (36 km baselines) to "D" (1 km baselines). Roughly 40% of observations are made at frequencies less than 10 GHz, which are most sensitive to fast, coherent emission processes and have a relatively large field of view. About 70% of observations make relatively light use of the correlators (continuum-like mode), which are the easiest for *realfast* to use. Considering both effects, we expect a typical year of commensal observing to search ~1500 hr summed over all antenna configurations.

Table 1 shows example correlator configurations that we intend to search commensally, as well as their expected sensitivity, and FRB detection rate. These rates are calculated based on the formulation of Burke-Spolaor & Bannister (2014) to extrapolate from one observing setup with a given FRB rate to a rate under a new observing setup. As a baseline FRB rate, we use that of Champion et al. (2016) and scale to the *realfast* setups listed in Table 1. We make two alterations in the Burke-Spolaor & Bannister (2014) formulation, which is that (1) we do not take into account Galactic effects (i.e., we only consider rate-scaling effects caused by instrumentation sensitivity and channelization differences) and (2) we use a corrected version of their Equation (3). As shown in Chawla et al. (2017), the scaling of the relative detection rate should instead scale with frequency as:

$$\frac{N_{\rm I}}{N_2} \propto \left(\frac{\nu_1}{\nu_2}\right)^{\alpha\gamma},\tag{1}$$

where α is the spectral index and γ is the power-law slope of the cumulative flux distribution.

The FRB detection rate ranges reflect uncertainty in FRB properties, such as average DM (assumed between 400 and 800 pc cm⁻³), intrinsic width, and scattering timescale of FRBs (between 0 and 3 ms at 1 GHz), a source count index of -1 (Vedantham et al. 2016; Lawrence et al. 2017), and a spectral index range of -1.0-1.0. The two 1.4 GHz modes show the currently commissioned mode with 5 ms integration time and a goal of 1 ms integration time. Using the formulation as described above, we make predictions gridding across a range of combinations of these values and report the minimum and maximum outcome for detection rates.

Modes will be commissioned prior to science observing to ensure that the correlator can safely generate both primary and commensal data. The highest priority modes for *realfast* will be at low frequencies (for sensitivity to coherent transients), use continuum-like spectral configurations (to minimize demand on the correlator), and compact antenna configurations (to minimize demand on the search pipeline). *realfast* will specifically target commensal searches during VLASS, an ongoing, 5500 hr sky survey from 2 to 4 GHz (VLA Sky Survey Science Group 2017, in preparation). Based on past correlator commissioning, we conservatively estimate commensal observations can be made with 50 ms integrations. This would give a 1σ sensitivity of 1 mJy per integration, equivalent to a 10σ fluence limit of 0.05 Jy ms. The VLASS detection rate translates to between 1 and 12 FRBs discovered for the entire 5500 hr survey.

5. Conclusions

realfast is a system that will commensally search VLA data for fast radio transients. We expect to search thousands of hours of data from the world's most sensitive radio interferometer, potentially localizing several FRBs to arcsecond precision and associating them with host galaxies. This potential has already been demonstrated in the near-real-time detection and precision localization of an FRB by a prototype version of *realfast*. The complete system will open discovery potential by continually searching for fast transients toward VLA targets such as galaxies, YSOs, AGN, TDEs, SNe, and GRBs. *realfast* results and data products will be shared in near real time to bring multiplicative benefits to VLA observations.

This system is an example of how in situ analysis can enable access to data-intensive science. By integrating dedicated computing with the observing system, *realfast* can triage the data flow through a series of decisions that can be made based on available computation and a well-defined model of the transient signal. This is supported by a flexible and extensible design that allows other kinds of commensal analysis, such as periodicity searches for pulsars or flagging to improve the data quality for primary observers. As one of the first large commensal data analysis concepts deployed at an interferometer, *realfast* can drive the design of other interferometers under development around the world.

The growth of astronomical surveys has produced a torrent of data that has driven discovery via machine learning (Bloom et al. 2012) and inference (Lang et al. 2010; Foreman-Mackey et al. 2015). In contrast, *realfast* demonstrates the strategy of asking specific questions of data prior to recording it. This strategy is not new to astronomy, as it was pioneered in data-intensive fields such as astroparticle physics (Gorham et al. 2004; Falcke et al. 2005). However, as instruments for time-domain astrophysics continue to grow in sophistication (LSST Science Collaboration et al. 2009; Harding et al. 2016), triage will continue to be an effective strategy to manage the data deluge.

We thank Bridget Andersen, the NRAO, and especially the VLA staff for their support of *realfast* development.

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. *realfast* is supported by the NSF Advanced Technology and Instrumentation program under award 1611606.

Facility: EVLA.

Software: vysmaw (Pokorny & Law 2017), rfpipe (Law 2017a), pwkit (Williams et al. 2017), astropy (Astropy Collaboration et al. 2013).

ORCID iDs

- C. J. Law () https://orcid.org/0000-0002-4119-9963
- G. C. Bower https://orcid.org/0000-0003-4056-9982

S. Burke-Spolaor https://orcid.org/0000-0003-4052-7838

References

- Akahori, T., Ryu, D., & Gaensler, B. M. 2016, ApJ, 824, 105
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
- Bandura, K., Addison, G. E., Amiri, M., et al. 2014, Proc. SPIE, 9145, 914522 Bannister, K. W., & Cornwell, T. J. 2011, ApJS, 196, 16
- Becker, R. H., White, R. L., & Helfand, D. J. 1995, ApJ, 450, 559
- Bhat, N. D. R., Chengalur, J. N., Cox, P. J., et al. 2013, ApJS, 206, 2
- Bloom, J. S., Richards, J. W., Nugent, P. E., et al. 2012, PASP, 124, 1175
- Booth, R. S., de Blok, W. J. G., Jonas, J. L., & Fanaroff, B. 2009, arXiv:0910.2935
- Bower, G. C., Deller, A., Demorest, P., et al. 2014, ApJL, 780, L2
- Burke-Spolaor, S., & Bannister, K. W. 2014, ApJ, 792, 19

- Burke-Spolaor, S., Trott, C. M., Brisken, W. F., et al. 2016, ApJ, 826, 223
- Caleb, M., Flynn, C., Bailes, M., et al. 2016, MNRAS, 458, 718
- Champion, D. J., Petroff, E., Kramer, M., et al. 2016, MNRAS, 460, L30
- Chatterjee, S., Law, C. J., Wharton, R. S., et al. 2017, Natur, 541, 58 Chawla, P., Kaspi, V. M., Josephy, A., et al. 2017, ApJ, 844, 140
- Clark, B. G. 1980, A&A, 89, 377
- Cordes, J. M., & Lazio, T. J. W. 2002, arXiv:astro-ph/0207156
- DeBoer, D. R., Gough, R. G., Bunton, J. D., et al. 2009, IEEEP, 97, 1507
- Deller, A. T., Vigeland, S. J., Kaplan, D. L., et al. 2016, ApJ, 828, 8
- Eftekhari, T., & Berger, E. 2017, ApJ, 849, 162
- Falcke, H., Apel, W. D., Badea, A. F., et al. 2005, Natur, 435, 313
- Faucher-Giguère, C.-A., & Loeb, A. 2011, MNRAS, 415, 3951
- Foreman-Mackey, D., Montet, B. T., Hogg, D. W., et al. 2015, ApJ, 806, 215
- Gorham, P. W., Hebert, C. L., Liewer, K. M., et al. 2004, PhRvL, 93, 041101 Harding, L. K., Hallinan, G., Milburn, J., et al. 2016, MNRAS, 457, 3036
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Natur, 217, 709
- Hulse, R. A., & Taylor, J. H. 1975, ApJL, 195, L51
- Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2006, Sci, 314, 97
- Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, AJ, 139, 1782
- Law, C. J. 2017a, rfpipe, Astrophysics Source Code Library, ascl:1710.002
- Law, C. J. 2017b, rtpipe: Searching for Fast Radio Transients in Interferometric Data, Astrophysics Source Code Library, ascl:1706.002
- Law, C. J., Abruzzo, M. W., Bassa, C. G., et al. 2017, arXiv:1705.07553
- Law, C. J., & Bower, G. C. 2012, ApJ, 749, 143
- Law, C. J., Bower, G. C., Burke-Spolaor, S., et al. 2015, ApJ, 807, 16
- Law, C. J., Bower, G. C., Pokorny, M., Rupen, M. P., & Sowinski, K. 2012, ApJ, 760, 124
- Law, C. J., Jones, G., Backer, D. C., et al. 2011, ApJ, 742, 12
- Lawrence, E., Vander Wiel, S., Law, C., Burke Spolaor, S., & Bower, G. C. 2017, AJ, 154, 117
- Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, Sci, 318, 777
- LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
- Lyne, A. G., McLaughlin, M. A., Keane, E. F., et al. 2009, MNRAS, 400, 1439
- Macquart, J.-P., Kanekar, N., Frail, D. A., & Ransom, S. M. 2010, ApJ, 715, 939
- Macquart, J.-P., & Koay, J. Y. 2013, ApJ, 776, 125
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, Natur, 439, 817
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in ASP Conf. Ser. 376, Astronomical Data Analysis Software and Systems XVI, ed. R. A. Shaw, F. Hill, & D. J. Bell (San Francisco, CA: ASP), 127 McQuinn, M. 2014, ApJL, 780, L33
- Ng, C., Vanderlinde, K., Paradise, A., et al. 2017, arXiv:1702.04728
- Noordam, J. E., & Smirnov, O. M. 2012, MeqTrees: Software package for implementing Measurement Equations, Astrophysics Source Code Library, ascl:1209.010
- Obrocka, M., Stappers, B., & Wilkinson, P. 2015, A&A, 579, A69
- Pérez, F., & Granger, B. E. 2007, CSE, 9, 21
- Perley, R. A., Chandler, C. J., Butler, B. J., & Wrobel, J. M. 2011, ApJL, 739, L1
- Petroff, E., Barr, E. D., Jameson, A., et al. 2016, PASA, 33, e045
- Pietka, M., Fender, R. P., & Keane, E. F. 2015, MNRAS, 446, 3687
- Pokorny, M., & Law, C. J. 2017, vysmaw: Fast visibility stream muncher, Astrophysics Source Code Library, ascl:1710.001
- Rafikov, R. R., & Lai, D. 2006, PhRvD, 73, 063003
- Ransom, S. M. 2001, PhD thesis, Harvard Univ.
- Readhead, A. C. S. 1994, ApJ, 426, 51
- Rubio-Herrera, E., Stappers, B. W., Hessels, J. W. T., & Braun, R. 2013, MNRAS, 428, 2857
- Schwab, F. R. 1984, in Indirect Imaging. Measurement and Processing for Indirect Imaging, ed. J. A. Roberts (Cambridge: Cambridge Univ. Press), 333
- Staveley-Smith, L., Wilson, W. E., Bird, T. S., et al. 1996, PASA, 13, 243
- Tendulkar, S. P., Bassa, C. G., Cordes, J. M., et al. 2017, ApJL, 834, L7
- Vedantham, H. K., Ravi, V., Hallinan, G., & Shannon, R. M. 2016, ApJ, 830, 75
- Wayth, R. B., Brisken, W. F., Deller, A. T., et al. 2011, ApJ, 735, 97
- Weitevrede, P., Stappers, B. W., Rankin, J. M., & Wright, G. A. E. 2006, ApJL, 645, L149
- Williams, P. K. G., Clavel, M., Newton, E., & Ryzhkov, D. 2017, pwkit: Astronomical utilities in Python, Astrophysics Source Code Library, ascl:1704.001
- Zhou, B., Li, X., Wang, T., Fan, Y.-Z., & Wei, D.-M. 2014, PhRvD, 89, 107303

8