PROCEEDINGS OF SCIENCE

The Art and Design of the CRAFT (Commensal Real-time Fast Transients) Survey

---

Jean-Pierre Macquart[*]
ICRAR/Curtin
E-mail: `Jean-Pierre.Macquart@icrar.org`

Peter Hall
ICRAR/Curtin
E-mail: `Peter.Hall@icrar.org`

Nathan Clarke
ICRAR/Curtin
E-mail: `Nathan.Clarke@icrar.org`

We describe the objectives and strategy of a commensal survey for fast (<5s) transient radio sources on ASKAP. Short-timescale transients are associated with the most energetic and brightest single events in the Universe. Our objective is to cover the enormous volume of transients parameter space made available by ASKAP, with an unprecedented combination of sensitivity and field of view. Fast timescale transients open new vistas on the physics of high brightness temperature emission, extreme states of matter and the physics of strong gravitational fields. In addition, the detection of extragalactic objects affords us an entirely new and extremely sensitive probe on the huge reservoir of baryons present in the IGM. We outline here our approach to the considerable challenge involved in detecting fast transients, particularly the development of hardware fast enough to dedisperse and search the ASKAP data stream at real-time rates at high time resolution. CRAFT aims to test of many of the key technologies and survey modes proposed for high time resolution widefield science with the SKA.

ISKAF2010 Science Meeting - ISKAF2010
June 10-14, 2010
Assen, the Netherlands

---

[*] Speaker.

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike Licence.

1. Introduction

The CRAFT survey is one of the eight survey science projects intended to be executed on the Australian SKA Pathfinder once it commences scientific operations in 2013. The principal aim of the survey is to detect, localise and characterise impulsive astronomical radio emission at centimetre wavelengths. ASKAP is particularly well suited to a survey of this type for several reasons: its 30 square degree field of view makes it sensitive to rare and possibly exotic events, its $36 \times 12\,$m dishes make it sensitive to relatively faint events, and its operation as an interferometric array, with baselines up to 6 km, allow us to both localise events to better than $10''$ and to verify any putative signal's extraterrestrial origin by ensuring that it is detected at multiple well-separated antennas. CRAFT is one of two projects using ASKAP to search for transient sources; whereas the VAST (Variable And Slow Transients) survey aims to detect signals with durations between 5 seconds and several weeks, the CRAFT survey is optimised for the detection of phenomena on much shorter timescales.

There are two main technical challenges particular to this survey. Firstly, since the ASKAP correlator integration timescale is 5 seconds, much larger than the $< 1\,$ms events we may wish to detect, the survey pipeline must intercept the high time resolution data stream before it reaches the correlator and process the data on a separate dedicated system. Secondly, the events we seek to detect are of sufficiently short duration that their signals are corrupted by interstellar and, potentially, intergalactic dispersion. Since the amount of dispersion (i.e. the dispersion measure, DM) of any transient is, a priori, unknown, this requires dedicated real-time dedispersion of the data stream for a range of possible dispersion measures.

The formidable technical hurdles associated with such an undertaking are offset by the large prospective rewards available. This is attested to by the historical success of high time resolution radioastronomy in attracting Nobel prizes: one for the detection of pulsars [13], and another for their subsequent use to test general relativity [14,24]. In general terms, short-timescale transients are associated with the most energetic and brightest single events in the Universe. They provide Nature's ultimate laboratory; their emission is usually generated by matter under extreme conditions whose properties probe physical regimes that far transcend the range achievable in terrestrial experiments. Even the mere existence of such impulsive emission in some instances can transform our understanding of the behaviour of matter and spacetime under the most extreme conditions.

Our purpose in the present paper is to outline the objectives, both scientific and technical, for a survey for short timescale transients and to describe our current thinking to address the challenges the endeavour entails. It updates some of the strategy outlined in earlier publications (e.g. [18]). In §2 we explore more deeply the scientific motivation for the detection of fast transients. We describe the exotic physics one may encounter in impulsive, cataclysmic events, and what knowledge we stand to acquire about the intervening interstellar and intergalactic media through which the signals of these compact objects must propagate. In §3 we describe our current efforts to hone both our hardware and algorithms to meet the challenge of running a pipeline to detect short-timescale transients on ASKAP, while in §4 we describe how these systems will likely be situated within ASKAP's signal processing chain. We present some final remarks on the overall outlook of the project in §5.

2. Scientific motivation and some physical impediments to the detection of impulsive emission

In this section we motivate the detection of fast transients using a parameter space argument, and then describe a number of particular transients that the survey would reasonably expect to detect. We provide an update on one particularly exciting sort of recently discovered fast transient and our prospects for furthering our understanding of them with CRAFT. We also briefly discuss the role of propagation effects in terms of the limits they may place on the properties of observable transient emission, and in terms of the information they may yield on the intergalactic medium.

2.1 Parameter Space

There is strong evidence that major discoveries in the field of astronomy, and radioastronomy in particular, follow advances and innovations in instrumentation [12,9]. This is often the direct result of new technology enabling access to hitherto unexplored regions of observational parameter space. In this sense, the science of fast transients represents one of the great frontiers of modern astronomy. Although pulsar astronomy has ostensibly probed part of this domain for several decades, it is only recently that the computational power, I/O bandwidth, and storage capacity has existed to enable searches for one-off transient events over a reasonable range of this parameter space, namely at high time resolution, over a large range in dispersion measures, and over a relatively large field of view.

This is perhaps no more exemplied than in the discovery of a potentially extragalactic burst in Parkes multibeam data reported by Lorimer et al. (2007) [17], and now known as the Lorimer burst (LB). Initial crude estimates indicated that these events occur at a rate of $\sim 225\,\text{day}^{-1}$ over the sky. This prompts questions as to why such seemingly common events were not discovered much earlier in the history of radioastronomy. One answer is that the discovery was likely only after two key technological innovations: (1) the 21 cm multibeam reciever, which increased the FoV of the Parkes telescope by a factor of 13 and, more importantly, (2) the advent of the computational capacity to search the large volume of data generated by this reciever for events at high time resolution. Backend computational power is now perceived as the key enabling technology in this field [1].

The detection of the Lorimer burst also highlights a number of key deficiencies of surveys similar to those conducted at Parkes. 1. Verification: the LB was detected at a single telescope, and simultaneous observation of the same signal at a different location could have confirmed if the signal was indeed extraterrestrial. 2. Field of view: the FoV of most radiotelescopes is small, rendering the detection of even bright but rare events improbable. By comparison, many X-ray and γ-ray telescopes are sensitive to events over a large fraction of the entire sky. 3. Event localization: surveys on single-dish telescopes are too poor at localising events to a small enough region to associate them with known objects.

The CRAFT survey addresses each of these three weaknesses. It represents an increase by more than an order of magnitude in field of view: Figure 1 compares the size of the unexplored parameter space covered by ASKAP relative to other facilities. This means that detections of events as “rare” as the Lorimer burst, which require many hundreds of hours observing time on Parkes, would be detected once every few days. The quoted event rate implies that CRAFT should observe one event every 6.1 days per 30 square degree field of view. As an interferometer with

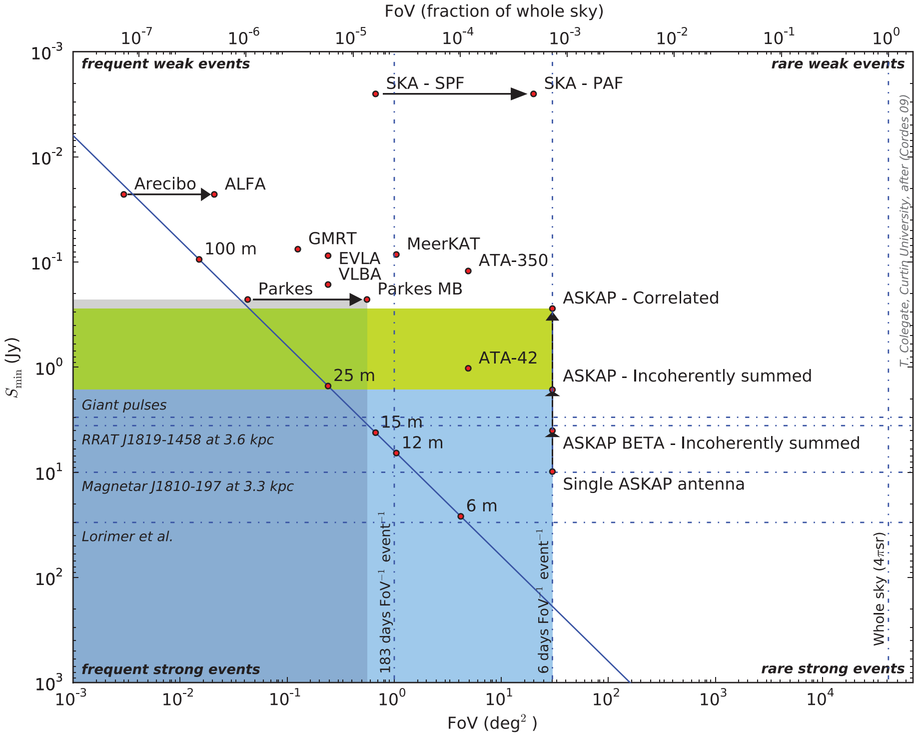

Figure 1: Transients parameter space as probed by ASKAP relative to other present and planned radio astronomy facilities at 1.4 GHz. S_{min} is the $5\,\sigma$ detectable flux density for a bandwidth of 300 MHz and integration time of 1 ms. The solid diagonal line refers to single-pixel, single-reflector telescopes having a system temperature of 25 K and aperture efficiency of 60%. The SKA-SPF line indicates the sensitivity for 2000 single-pixel feed 15-m antennas, extended to phased array feeds (SKA-PAF) with a 20 deg^2 FoV. Adapted from [5].

baselines up to 6 km, it provides an important signal verification capability on the basis that any extraterrestrial signal should be observed by independent well-spaced antennas. It also provides a means of localising the event to a region small enough to usually be able to uniquely associate it with a known object.

2.2 Known Knowns, Known Unknowns, and Unknown Unknowns

The sorts of objects would we expect to detect in a survey for fast transients fall broadly into three categories[1]:

Known Knowns

Giant pulses Some pulsars emit pulses whose strengths exceed the mean pulse intensity by several orders of magnitude. Some giant pulses from the Crab pulsar possess flux densities $\sim 10^3$ Jy at 5 GHz and durations of < 2 ns [10], and numerous other pulsars are now known to exhibit the same phenomenon [4,16,23]. Some other pulsars exhibit a related phenomenon

---

[1] With apologies to Donald Rumsfeld.

of "giant micro-pulses", where an individual component of the pulse sometimes exceeds its mean flux density by factors of ~ 100 (e.g. [16]). CRAFT's contribution lies in its ability to serendipitously detect giant pulses from hitherto unknown sources.

Magnetars Magnetars are neutron stars whose emission is believed to be powered by decay of the object's superstrong magnetic field. Radio emission is episodic and bright. The magnetar J1810-197 has been detected with peak single-pulse amplitudes $\sim$10 Jy and pulse widths ~ 0.15 s. Given its distance of 3.3 kpc [2], single pulses from magnetars like J1810-197 are detectable to $D_{\rm max} = D(S_{\rm pk}/S_{\rm min,SP})^{1/2} \approx 70\,$kpc for ASKAP.

RRATs Pulses from Rotating RAdio Transients are highly intermittent, with the average time between pulses ranging from 3 min to 3 h depending on the object, and fundamental pulse periods between 0.7 and 7 s. The intermittent nature of RRAT emission renders them difficult to detect, and it is likely that the RRAT population actually rivals the number of normal pulsars. ASKAP's large FoV makes it the ideal instrument for addressing the fundamental issue of the total number of these sources.

Known Unknowns

Lorimer bursts The LB was detected as a $\sim 30\,$Jy flare of duration 5 ms. The arrival time of the pulse as a function of frequency suggests that the pulse was dispersed by an extraterrestrial plasma, and the dispersion measure of 375 cm^{-3} pc suggests that the pulse is extragalactic, and travelled through the intergalactic medium (IGM) from a redshift $z > 0.2$. Subsequent reanalysis of Parkes multibeam data has revealed a number of objects with characteristics similar – although not entirely identical – to the LB. However, the $\gg 100\,$Jy flux density of each of the bursts, their quasi-annual cycle, and changes in the frequency-time sweep during the course of some bursts cast doubt on the interpretation of the distance in terms of interstellar interstellar/intergalactic dispersion. These results are presented in Burke-Spolaor et al. (submitted), who suggest that these bursts may instead represent the detection of a new form of terrestrial atmospheric emission.

Annihilating Black Holes An annihilating black hole may produce radio bursts [22]. Previous searches for these [21] have been unsuccessful, but these were severely limited by the technology of the day.

Gravitational Wave Events It is an open question as to whether gravitational wave events, such as those associated with the coallescence of two black holes, emits EM radiation immediately prior to or after an event. The detection of the electromagnetic signatures of events that generate gravitational waves may be necessary to localize the sources of the emission sufficiently well to uniquely identify them. Gravitational wave events may generate associated electromagnetic emission when matter is present in the vicinity of the system. For example, the in-spiral of a binary neutron star system may produce radio wavelength electromagnetic pulses due to the interaction of the magnetospheres of the neutron stars (e.g. [11]).

SETI and Extraterrestrial emitters ET signals could appear transient, even if intrinsically steady [6]. The hardware configurations and processing techniques necessary for the detection of fast transients also satisfy most of the requirements for the detection of a SETI signal.

Unknown Unknowns
By its very definition, it is impossible to give instances of objects in this category. We can, however, point out that objects such as pulsars and the LB formerly inhabited this category, as they were entirely unanticipated prior to their discovery. Given the new physics associated with the discovery of these objects, they are the most scientifically lucrative.

2.3 The Boon and Bane of Scattering

Objects which vary on sub-second timescales are necessarily compact. This makes their emission subject to a variety of scattering and propagation effects. There are two aspects to this.

1. Induced scattering as a limiting process
Given the sensitivity of our instrumentation, any object that is detectable at Galactic or extragalactic distances is necessarily very bright, and the implied extreme brightness temperature of the emission makes the radiation subject to a number of induced scattering effects which can impede escape of the radiation from the source. These effects, notably induced Compton scattering and induced Raman scattering, potentially limit the brightness of the observed radiation if there is dense matter enshrouding the emission region [19].

2. Dispersion and Temporal Smearing as Probes of the IGM
The implusive nature of the emission means that the relative delay of emission across the observing band, caused by dispersion of the radiation in the interstellar/intergalactic medium, can be large compared to the overall impulse duration. This not only renders the effects of dispersion highly noticeable, it also means that any detection of a short duration transient automatically furnishes a direct estimate of the electron column between the source and observer, and this acts as a surrogate for distance.

The detection of bright transients at extragalactic distances would allow us to probe the ionized Intergalactic Medium (IGM) in much the same way that pulsars have done for the ISM. It is believed that the majority of baryons in the low-redshift, $z \lesssim 1$ universe reside in the tenuous, $\rho \sim 10^{-7}\,\mathrm{cm}^{-3}$ IGM and that at least 45% have not been detected [3,7,20]. The missing baryons are extremely difficult to detect as they are supposed to reside at temperatures $10^5\,\mathrm{K} < T < 10^7\,\mathrm{K}$ in a highly ionized gas. The detection of a number of extragalactic transients offers an excellent opportunity to detect *all* the baryons along each line of sight.

3. Trailblazer surveys and initial algorithm and hardware development

A major challenge in the engineering of instrumentation for observing cosmic transients is the wide-open nature of the observing parameter space. It is not yet clear which volume of the parameter space deserves most attention in the design of the ASKAP system and, to contribute to the maturation of scientific priorities (which in turn lead to the instrumentation specifications), the CRAFT collaboration has begun several 'trailblazer' projects. These are designed to yield increased scientific insight while giving engineers the opportunity to evaluate various approaches to instrument design. The results of these projects will flow to the specification of an ASKAP instrumentation suite while making, in the best estimate, ground-breaking contributions to radio astronomy science and engineering. In this short paper we mention (below) only a few of the projects currently underway, choosing the VLBA experiment as a representative, ambitious science exemplar and the various technology platforms as examples of approaches which may be incorporated in ASKAP. A major aim of CRAFT is to develop techniques and technologies which scale to SKA Phase 1 and beyond; considerable effort is therefore being expended in looking for algorithms and architectures likely to be useful in the era of SKA. With so much to learn about the variable Universe, our first efforts concentrate on incoherent array processing techniques but our thinking is informed by the eventual promise of coherent processing across at least sub-arrays of the SKA.

3.1 The VLBA survey

A novel pilot survey for fast transients is being implemented at the Very Long Baseline Array (VLBA). Since VLBA data are now correlated using the DiFX software correlator [8], a software pipe has been trivially implemented to output the auto-correlation products of all 10 VLBA antennas to a separate transients box during production correlation. This is used to perform a commensal search for fast transient radio sources using 1 ms-averaged auto-correlation data. Whenever a candidate transient event is detected, the raw voltage data is saved for a detailed further analysis. Following on from the success of early trials, a proposal to take data continuously has been submitted by CRAFT members Tingay, Wayth, Brisken and Deller.

This experiment possesses a number of virtues from both the scientific and engineering perspective. On the scientific side, the use of 10 telescopes simultaneously, spanning thousands of kilometres, gives a highly robust filter against signals of a non-astronomical origin. Ready recourse to the original voltage data streams makes it possible to perform detailed follow-up once a candidate transient signal is detected. This includes full coherent de-dispersion of the high-resolution time series, as well as cross-correlation of the signals to localise the transient to better than 10 mas accuracy.

On the engineering front, the experiment serves a platform to trial strategies for implementation on higher-throughput instruments such as ASKAP and the SKA. Some of this architecture is detailed below.

3.1.1 Software-based architectures and the VLBA experiment

Although processing capacity limitations (per Watt and per dollar) of software-based systems render them nonviable for investment in high throughput applications such as ASKAP and the SKA, the expediency and versatility of software makes it ideal for developing and trialing new transient detection algorithms. The leading trailblazer architecture therefore consists of a dedicated server-class computer executing a software-based realtime de-dispersion and detection pipeline.

CRAFT collaborators from ICRAR, JPL and NRAO have tailored the software-based transient detection system for testing and demonstration on the VLBA telescope in Socorro, New Mexico, with aspirations of performing a full transients survey later this year. The transient detection system runs in commensal mode alongside ordinary VLBA correlation and receives auto-correlated power spectrum data from the VLBA's DiFX correlator. The system processes the data through several pipeline stages that perform packet re-ordering, missing data filling, calibration signal removal, RFI excision, incoherent de-dispersion and transient detection. Finally, the system schedules dump requests back to the telescope to capture any transient events deemed worthy of closer off-line examination.

3.2 GPU-based architectures

Graphics Processing Units (GPUs) are specialised microprocessors primarily designed for offloading and accelerating the graphics rendering tasks otherwise performed by general purpose microprocessors (i.e. CPUs). Their programmable, high-performance signal processing capabilities make them adaptable to a range of radio astronomy applications. A team at ICRAR are developing a GPU-based transient detection system which constitutes the second of our trailblazer investigations.

Although the system is still under development, GPU kernels implementing the de-dispersion and transient detection algorithms are complete and have been verified using test data. The transient detection algorithm currently employs a peak finding routine, but more elaborate algorithms, such as Duchamp emission search routines, are also being examined. System performance is dictated by I/O bandwidth limitations and the team is looking to improve performance by optimizing the interfaces between key components of the system. As it is now, the system can process up to 18 beams from the 1.4 to 1.7 GHz band with each beam consisting of 1 ms integrated samples of auto-correlated data.

3.3 FPGA-based architectures

Field Programmable Gate Arrays (FPGAs) are integrated circuits consisting of highly configurable logic elements, memories, digital signal processing cells and other generic logic components, all inter-connected via a signalling fabric that is also highly configurable. The ability to reconfigure FPGAs makes them especially useful for developing and testing new designs and for benchmarking alternative algorithms and architectures. FPGA configurations are specified using hardware description languages that are also used in the design of Application-Specific Integrated Circuits (ASICs) and typical development paths often lead from FPGA to ASIC in which productions of many thousands of chips are more economically viable. With the SKA ultimately in mind, FPGA therefore offers a very promising technology on which to base our trailblazer transient detection architectures.

To expedite development, collaborators from JPL and ICRAR are planning to implement preliminary FPGA-based transient detection architectures on a mature, commercial FPGA development platform, taking advantage of vendor support for some of the more universal but challenging design tasks, such as interfacing to external memory. The platform consists of a Linux host system with six PCI-Express slots available to accommodate an array of FPGAs and memory. Each PCI-Express card serves as a high-speed backplane interconnecting up to six daughter-card modules, with each module contributing one Xilinx Virtex-6 FPGA and 512 MBytes of associated memory. Initial demonstrations will likely involve a single platform capable of monitoring in the order of tens of signal beams simultaneously (sufficient for a single ASKAP dish for example), but an array of these platforms would be needed for larger demonstrations and more formal transients searches on BETA and Full ASKAP telescopes.

As with the GPUs, the main performance bottle-neck for the FPGA platform is at the I/Os, in particular the interfaces between the FPGAs and their associated SDRAM memories. While the

FPGAs have ample resources for processing the data, they do not have sufficient internal memory to apply the necessary delays to the data for de-dispersion. For this the FPGAs must utilize the commodious SDRAM resources dedicated to each of them. The SDRAM interfaces are 32 bits wide and (with the -2 and -3 speed-grade parts) are capable of operating at the DDR3-1066 JEDEC standard rate. This amounts to a raw data rate of 4.26 GBytes/s to memory for each FPGA. In practice however the utilization of raw SDRAM bandwidth is often significantly less than 100% due to limitations of the interface (such as minimum burst size) and the team is investigating algorithms and architectures that make sufficient use of the available bandwidth to achieve our performance targets.

4. ASKAP Hardware

Figure 2 shows a simplified version of the ASKAP signal path. Some transients science will undoubtedly come from using the correlator in its highest time resolution mode (~ 5 s visibility dumps). However, we envisage the bulk of the investigations relying on one or more additional processing paths having access to summed array data in either coherent or incoherent summation modes; these modes give maximum time resolutions of about 0.5 ms and 3 ns, respectively. As noted in the previous section, most initial thinking involves incoherent addition across the array of dishes. Of course, with the 30× field of view expansion technology employed in ASKAP, a full implementation of such a mode requires that each of the 30 beams is processed discretely, making even an incoherent processor a formidable undertaking.

In the ASKAP architecture it is convenient to extract incoherent data from each dish at the phased array feed (PAF) beamformer. As part of the normal FX correlation architecture, digital data is coarse-channelized into 1 MHz segments at this stage and, for telescope timing purposes, a buffer of several seconds is inserted into the signal path. It is a relatively simple matter to intercept the samples at the buffer output, square them to form power estimates and accumulate them for a selected integration period - probably around 0.5 ms minimum based on PAF beamformer internal considerations. Thus, channelized, integrated samples can be obtained on a ‘per beam’ basis from each antenna. This information is communicated to the parallel time domain processing chain via 1 GbE ports on the beamformer; information from all dishes is then distributed via a switch to the required signal processing engines. If the real-time time domain system detects an event of sufficient significance, data can be frozen in the buffer, allowing the sythesis imaging machinery (or future coherent time domain processor) to ‘replay’ the event for deeper analysis. Current plans call for a buffer length of only a few seconds but promising results from CRAFT trailblazers will be a strong motivation for seeking funds to fit extra buffer memory at each dish.

While blind searches are most practically tackled by the incoherent processing approach, coherent processing offers a substantial (factor of six) gain in targetted ASKAP observations. In the foreseeable future, these observations will be enabled via tied-array machinery put in place ostensibly for VLBI or pulsar timing. The field of view in each array beam will be small, but up to four pencil beams will be available within the PAF field. For certain monitoring or similar transients observations, possibly invoking off-line or limited-bandwidth processing, this mode may be a powerful one.

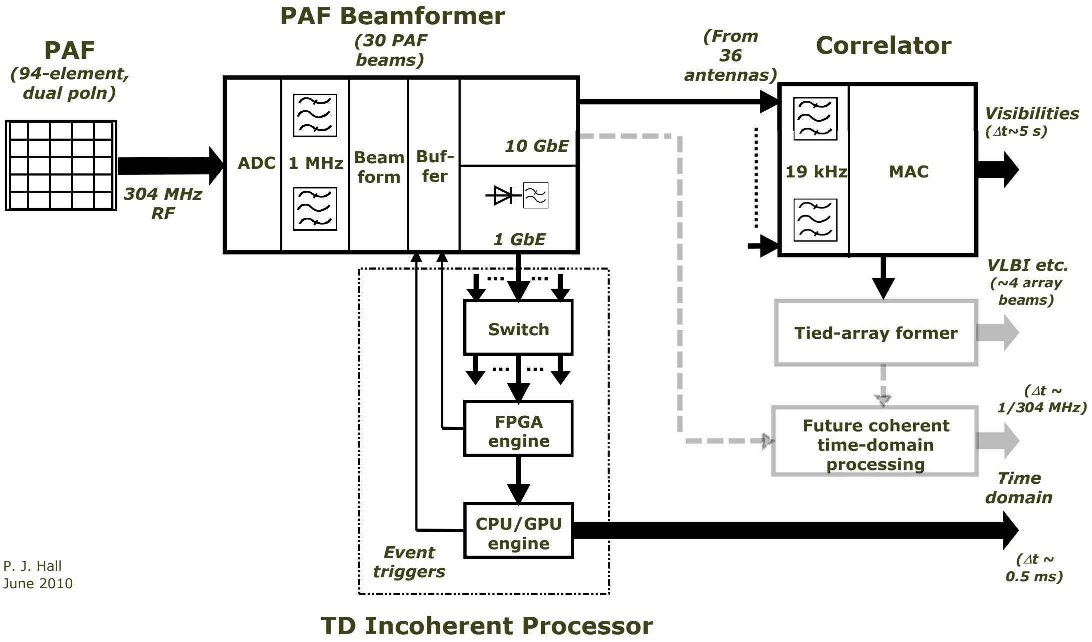

Figure 2: A block diagram showing how CRAFT hardware would fit within the ASKAP hardware.

Our intention is to outline the top-level form of the ASKAP time domain instrumentation suite by Q4, 2010. We expect some trailblazer results to be available at that point, and over the following 12 months more experience with the trailblazers will enable refinement of the ASKAP design. Trailblazer equipment will also be deployed on the ASKAP BETA telescope as soon as feasible, in order to learn early lessons about the telescope and MRO site. During 2012-13, ICRAR and other CRAFT partners will produce a detailed, costed design study for the full ASKAP time domain instrumentation suite. While full funding is available for the trailblazers and design study, more resources (at the several million AUD level) will be required for the full implementation.

4.1 The Impulsive RFI environment at Boolardy

In a detection system which incurs a time penalty every time a buffer readout is called it is necessary to limit the number of triggers to a manageable rate. However, we would like to be certain that the events the system triggers on have a high likelihood of being genuine; it is undesirable to limit the number of triggers merely by raising the trigger S/N until the event rate is sufficiently low.

As a precursor to CRAFT, we have undertaken to characterise the short-timescale RFI environment at Boolardy in an attempt to estimate the number of false triggers our system will need to process. Data were analysed from an 8 hour duration VLBI observation of Centaurus A conducted at 1.4 GHz with 64 MHz bandwidth. Data were collected with the first 12 m ASKAP dish at Boolardy, current equipped with a single-pixel feed, and compared to data taken at the Parkes 64 m telescope as part of the same experiment. The Parkes voltages were recorded with 2-bit digitisation, whereas the Boolardy data were recorded with 8-bit sampling. Construction was occurring at Boolardy during the observations, and no particular efforts were made to observe radio quietness.

We examined the statistics of the total powers on timescales of 10 μs. Over identical time ranges, 770 events were detected at Boolardy and 143661 events were detected at Parkes that

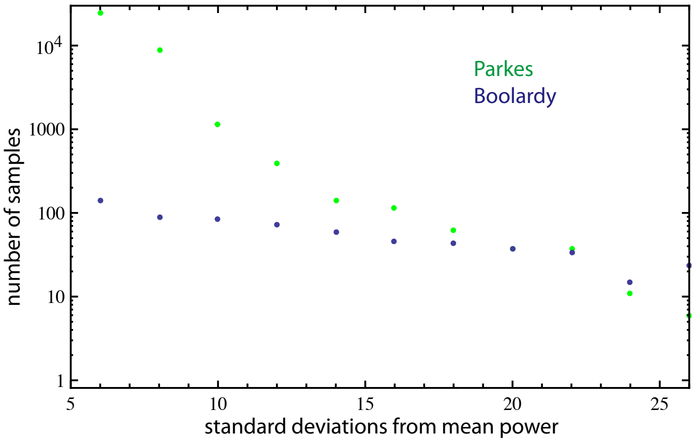

Figure 3: A histogram of the number of events exceeding 6 standard deviations in power from the mean on a 10 μs timescale observed during the same 8 h duration of a VLBI observation of Centaurus A with both the 12 m ASKAP dish at Boolardy and the 64 m dish at Parkes.

deviated from the local mean by more than 6 standard deviations. Thermal noise fluctuations should yield roughly 3 events over the same time range. Although, it is difficult to make a direct comparison between the two sites, it does illustrate the fundamental fact that, for a campaign to detect fast transients with these two telescopes at their respective sites, the false trigger rate at Boolardy is over two orders of magnitude lower than at Parkes. Figure 3 shows a histogram of the number of samples as a function of amplitude. The amplitude distribution for Boolardy is essentially flat, showing that when one obtains 6σ deviations, one also tends to equally observe 20- and 30-σ deviations. Further investigation have revealed that the RFI is in fact highly clustered in time and that it is often unresolved on 10 μs timescales.

5. Conclusion

The discovery of the fast, bright transient reported by Lorimer et al. (2007) has precipitated the realization that many similar classes of objects may remain to be detected. Little thought has been devoted to the best means of systematically exploring the largest possible volume of this parameter space. We are conducting a number of trailblazer experiments to determine over which regions of parameter space future surveys and hardware detection suites should concentrate their efforts.

The survey for fast transients on ASKAP will represent the culmination of a long chain of efforts in searching for the optimal means of exploring this largely uncharted volume of parameter space. In turn, the CRAFT survey will act as a precursor to the yet more ambitious transients surveys planned on the SKA, providing guidance on optimal hardware, algorithms and experience in exploiting the telescope commensally with other key SKA science.

References

[1] G.C. Bower, *Mining for the Ephemeral*, *Science*, 318, 759 (2007)

[2] F. Camilo et al., *Transient pulsed radio emission from a magnetar*, *Nature*, 442, 892 (2006)

[3] R. Cen & J.P. Ostriker *Where are the baryons?* ,*ApJ*, 514, 1 (1999)

[4] I. Cognard, J.A. Shrauner, J.H. Taylor & S.E. Thorsett, *Giant radio pulses from a millisecond pulsar ApJ*, 457, L81 (1996)

[5] Cordes, J.M., 2009, SKA Memo 97, Revised 04/09, http://www.skatelescope.org/pages/memos

[6] J.M. Cordes, T.J.W. Lazio & C. Sagan, *Scintillation-induced intermittency in SETI*, *ApJ*, 487, 782 (1997)

[7] Davé et al., *Baryons in the warm-hot Intergalactic Medium ApJ*, 552, 473 (2001)

[8] A.T. Deller, S.J. Tingay, M. Bailes & C. West, *DiFX: a software correlator for very long baseline interferometry using multiprocessor computing environments*, *PASP*, 119, 318 (2007)

[9] R.D. Ekers *Big and Small*, *PoS*, sps5, 007 (2009)

[10] T.H. Hankins, J.S. Kern, J.C. Weatherall & J.A. Eilek, *Nanosecond radio bursts from strong plasma turbulence in the Crab pulsar*, *Nature*, 422, 141 (2003)

[11] B.M.S. Hansen & M. Lyutikov, *Radio and X-ray signatures of merging neutron stars*, , MNRAS, 322, 695 (2001)

[12] M. Harwit, *Cosmic Discovery - The Search, Scope & Heritage of Astronomy*, Basic Books Inc., New York 1981.

[13] A. Hewish, et al. *Observation of a Pulsating Radio Source, Nature*, **217**, 709 (1968)

[14] R.A. Hulse & J.H. Taylor, *Discovery of a pulsar in a binary system, ApJ*, 195, L53 (1975)

[15] S. Johnston & R.W. Romani, *Giant pulses from PSR B0540-69 in the Large Magellanic Cloud, ApJ*, 590, L95 (2003)

[16] S. Johnston, W. van Straten, M. Kramer & M. Bailes, *High time resolution observations of the Vela pulsar*, *ApJ*, 549, L101 (2001)

[17] Lorimer, D.R., Bailes, M., McLaughlin, M.A., Narkevic, D.J. & Crawford, F. *A bright millisecond radio burst of extragalactic origin, Science*, **318**, 777 (2007) [arXiv:0709.4301]

[18] J.-P. Macquart, *The Commensal Real-Time ASKAP Fast Transients (CRAFT) survey*, *PASA*, **27**(3), 272 (2010) [arXiv.org/1001.2958]

[19] J.-P. Macquart, *On the detectability of prompt coherent gamma-ray burst radio emission*, *ApJ Lett*, **658**, L1 (2007) [arXiv:astro-ph/0702098]

[20] F. Nicastro, S. Mathur & M. Elvis *Missing baryons and the warm-hot Intergalactic Medium*, *Science*, 319, 55 (2008)

[21] J.D. O'Sullivan, R.D. Ekers & P.A. Shaver, *Limits on cosmic radio bursts with microsecond time scales*, *Nature*, 276, 590 (1978)

[22] M.J. Rees, *A better way of searching for black-hole explosions*, *Nature*, 266, 333 (1977)

[23] R.W. Romani & S. Johnston, S., *Giant pulses from the millisecond pulsar B1821-24*, *ApJ*, 557, L93 (2001)

[24] J.H. Taylor & J.M. Weisberg, J. M. *A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16*, *ApJ*, 253, 908 (1982)