SKA Transients Science Assessment Workshop Summary

Version 2, 2014 Feb 21

This is the summary of the Science Assessment Workshop held between representatives of the SKA Transients Science Working Group (TSWG) and the SKA Project Office, in January 2014. It includes a summary of the discussions, the main findings and change requests of the TSWG. It also incorporates, in two appendices, a list of detailed comments on the current SKA1 Baseline Design and Concept of Operations documents.

Overview and headline science

Radio transients are both the sites and signatures of the most extreme phenomena in our Universe: e.g., exploding stars, compact object mergers, black holes and ultra-relativistic flows. They also have the potential to act as probes of the intervening medium on cosmological scales. These transients can be broadly divided into (i) incoherent synchrotron events, which are associated with all explosive kinetic feedback in the Universe, from relativistic jets to supernova explosions, and (ii) coherent bursts, which can achieve much higher surface brightness and hence can be of very short durations. These two classes of transients also broadly match to the observation method: incoherent events have flaring timescales of minutes to years and are typically found in images, coherent bursts can be unresolved on timescales shorter than milliseconds, and are typically found in beamformed time series data. These divisions are not exclusive, however, and very fast imaging (on millisecond timscales) and strong scattering can blur the boundaries between the two. Furthermore, while the above constitute the majority of radio transients, strong variability can also be observed associated with thermal emission (e.g. novae), lensing and scattering, all of which can be picked up by the revised approaches outlined here. In all cases, transient science benefits greatly from monitoring the sky frequently. As such, commensal observing alongside other projects is a powerful strategy.

Rob Fender (incorporating material from Stéphane Corbel and James Miller-Jones), Jean-Pierre Macquart and Zsolt Paragi spoke about this headline science.

Commensal searching and rapid response

Assuming that the collecting area and hence sensitivity of SKA_1 is broadly fixed, the SKA Transients Science Working Group (TSWG) concluded that its highest priority design changes would be to firmly establish commensal searching of all data streams, and the ability to rapidly and flexibly respond to alerts (external or internal), in order to maximise the number of transients discovered, and the likelihood of discovering the rarest events. Our priority changes are listed below; more detailed comments are in Appendix 1.

- **Commensal/real-time detection of radio transients:**
  Put simply, commensal searches of all data streams will increase by an order of magnitude the rate of transient detection, and hence scientific return (in this area). The fastest transients can only be detected in real time, as the medium or long term storage of high time resolution data products is prohibitively expensive for an array (as an aside, real-time detection is also enhanced greatly if a more raw form of the data can be buffered for further offline analysis, e.g. in order to precisely localize the source). All searches can be done in a commensal fashion, regardless of the primary scientific objective (and at no cost to it) and mode of observation, greatly expanding the amount of observing time available. We therefore strongly recommend that real-time image-plane and beamformed transient searches are incorporated into the SKA_1 design.

- **Rapid follow-up of SKA_1 or external transient triggers:** A *responsive* telescope is essential for studies of transients, some of which may only be observable for days or even seconds. As an example, a rapidly conducted, wide search of a gravitational wave trigger could provide the first electromagnetic counterparts to gravitational wave sources.
- **Archival studies:** An easy to access and full-featured archive is essential to the study of longer-timescale transients, both to provide a source of transients which can then be sought at other wavelengths, as well as to facilitate the identification of radio counterparts to sources in other wavebands. In many situations, it is also important to have the lag between observation and archiving be as short as possible, so quick decisions can be made on how to proceed with the further multi-wavelength follow-up of a transient source.

As it stands, none of these three areas are optimally served by the Baseline Design for SKA_1. Fortunately, some limited changes to the Baseline Design can greatly improve the productivity of SKA_1 for transient studies, and we list these at the end of the report. Fully commensal observing should increase the rate of transient detection by an order of magnitude compared with the more conservative 'traditional' approach of dedicated observations. Our best estimate of transient detection rates is that for a fully commensal SKA – i.e. all data streams being searched, all the time – of order one 'interesting' (a very subjective description, roughly corresponding to publishable) transient should be discovered *per day*.

Examples were given during the workshop by Rob Fender (AMI), Richard Armstrong (KAT-7, plans for MeerKAT), and Aris Karastergiou (ARTEMIS/LOFAR international stations) of currently operating near-real-time commensal and/or rapid-response systems. Jason Hessels (incorporating material from Cath Trott) spoke on the need for Transient Buffers. John Swinbank spoke on the need to design flexibility into the system from day one, informed by experience with LOFAR, as well as the (two-way) interaction of the system with the *VOEvent* framework for reporting and responding to transients.

It was noted that as it stands the *Concept of Operations* (CoO) document (revision B, 2013-10-29) either does not support the flexibility outlined above, or is vague and noncommittal enough that future iterations could remove such support. The TSWG strongly urges that a revised CoO document should fully embrace the potential for new and flexible observing modes, *in order to maximise the scientific return of the SKA* (its own stated objective). Appendix 2 lists more detailed comments on the CoO.

Level 0 requirements

These considerations have led, as a direct result of the discussions at the workshop, to the two following Level 0 Science Requirements:

> **All components of SKA_1 will provide a capability for searching of data streams (imaging and/or beam-formed) for astrophysical transients and be able to report unambiguous detections with a latency of less than 100 seconds.**

This will increase by an order of magnitude the rate of transients detection, and the chances of finding the most rare events. Rapid dissemination of the brightest alerts will maximise the global scientific return via multi-wavelength follow up.

All components of SKA_1 will be able to respond to alerts, which may be human- or robotically-generated, on a timescale of <30 seconds. SKA_1-Mid and SKA1-Survey will begin slewing on this timescale; SKA1-Low will begin beam-forming in the new direction on this timescale.

This will maximise the scientific return of transients studies, and greatly enhance SKA_1's ability to work productively in concert with other contemporaneous facilities such as LSST, CTA, ground-based gravitational wave interferometers and space-based missions, all of which have a strong focus on astrophysical transients.

Regarding the first of these L0 requirements, we note that for *internal* triggers – such as the real time detection of a FRB – we would need a much shorter response latency of order a few seconds (or a prohibitively large raw data buffer).

Higher frequencies and higher angular resolution

The importance of high angular resolution and – independently – high frequencies, was discussed extensively at the workshop, and is not satisfied simply by the incorporation of VLBI modes into SKA1, which we strongly support (see below). High angular resolution is essential for spatially resolving nearby transients (see below), as well as providing sufficiently good localisation for identification of counterparts at other wavelengths and/or regions within host galaxies, which can provide strong clues as to the nature of the progenitor. Independently of angular resolution, at higher frequencies synchrotron events peak earlier and more sharply, making them both easier to find and to correlate with trigger events at higher frequencies (e.g. science coordinated with LSST). For this reason we strongly support inclusion of (at least) Band 4 in the SKA_1 design.

Zsolt Paragi and Rob Fender spoke about the need for higher angular resolution and higher frequencies for synchrotron transients.

Support for VLBI modes

The TSWG strongly supports the inclusion of VLBI modes in SKA_1. Much transient science has come from VLBI observations, notably with the high sensitivity EVN and VLBA arrays, probing the rapid evolution of relativistic jets from stellar mass black holes, GRB afterglows and supernovae. VLBI-capable SKA_1 elements could be incoporated into an expanded EVN and/or African VLBI Network.

Zsolt Paragi and Adam Deller spoke about the need for VLBI and the relative ease of its incorporation into the SKA_1 baseline design.

Specific comments on each component of SKA_1

All three components of SKA_1 are likely to discover large numbers of transients.

1. SKA_1 Low

The overriding request for a design change for SKA_1 Low is for a beamformer with multiple beams. It is noted that if SKA_1 Low is adapted to facilitate pulsar surveys, then little further adjustment would be needed to have a telescope capable of also finding fast radio transients.

This discussion was led by Jason Hessels.

2. SKA_1 Mid

For SKA1 Mid our main change request is for simultaneous imaging and beamformed modes, running all the time. This maximises the exploration of transient parameter space (for example, for the same population of extragalactic FRBs, imaging/beamforming are optimal for discovery depending on whether or not there is strong/weak scattering in the IGM). Multiple VLBI beams are highly desirable.

This discussion was led by Ben Stappers.

3. SKA_1 Survey

SKA_1 Survey can maximise its harvest of transients by building in to the design fast imaging modes. While ~10 sec is fine for synchrotron transients, imaging as fast as 1ms would be extremely interesting for fast coherent bursts. A transient buffer of at least 30 seconds is required (these requirements are also relevant to SKA_1-Mid).

This discussion was led by Jean-Pierre Macquart.

We note that in the very likely absence of the ability to tile out the entire SKA1 Survey FoV with coherent beams, we can only maximise the science by having fast imaging and/or simultaneous wide-field incoherent beam.

Use cases

It was established at the workshop that *use cases* would be a good tool to make sure that all sides understood the envisaged use and deployment of the telescope(s) for different science goals. This is particularly important from the point-of-view of transient science, which is arguably driven as much by how the telescope can be used in a responsive and commensal way, as it is by basic parameters like sensitivity and field-of-view. Several such use cases were identified for the TSWG, including:

- FRBs – detailed implementation of searches on all components of SKA1

(this has been circulated to the SPO project scientists)

- ToO followup of LIGO events (both coherent and incoherent)
- 1-minute real-time imaging for synchrotron transients
- VLBI – astrometry of Galactic explosions, SN kicks
- VLBI – resolving ultrarelativistic explosions in extragalactic transients

Appendix 1:
Detailed itemised comments on the SKA1 baseline design:

1. We strongly recommend the capability to do commensal/piggybacking searches of all data streams to look for transients, and to respond rapidly to external (ToO) alerts. This infrastructure will require revised data paths and real-time processing to dynamically identify signals of interest. Specifically, **w**e need to be able to (re-)image our data over the full range of timescales (e.g. from 1 sec to several hours) to filter-match any possible events. Signal diversity (many channels, polarizations, beams) is critical to have confidence in rare/unique transient events. In some cases, this is more important than raw sensitivity, which typically comes at the expense of diversity (e.g., by averaging in frequency, time, polarisation). In parallel to this, flexibility in compute hardware allows for implementation of new algorithms for as-yet unknown classes of transients (in fact, flexibility and modularity along the whole signal path so one can inspect data at a variety of stages between the antenna and final correlator products). For SKA1-low: access to station voltages (ideally, we would like [a subset of] antenna voltages if all the signals are aggregated at the central processing hut). For SKA1-survey: access to PAF beamformer voltages. For SKA1-mid: access to individual dish voltages
2. Simultaneous imaging and beamformed modes (for SKA1 Low and SKA1 Mid) will

greatly increase discovery space and scientific exploitation for fast radio transients.

3. Ability to do many short observations with little to no overhead in between (e.g. for monitoring known sources) → short slew/software reconfiguration timescales. We note that existing slew speed requirements are probably good enough.
4. Snapshot *u*-*v* coverage: the current design has good instantaneous *u*-*v* coverage, which is important for high-time-resolution imaging. Any future baseline design changes should continue to implement this.
5. There needs to be a stronger emphasis on producing a full-featured data archive with easy data access and in situ visualization and processing capabilities. The current Baseline Design is focussed on the hardware and real-time data reduction, but how scientists access the resulting data (and how fast and easily they can access it) also deserve very careful consideration. LOFAR is an example of a large project which failed to plan properly for an effective archive.
6. Sub-arraying is highly desirable. This would be used for monitoring large numbers of sources simultaneously or for responding rapidly to external triggers without requiring the entire SKA collecting area.
7. We strongly support the (already suggested) return of *all* SKA1 Low antenna signals to the central processing facility. This would facilitate the implementation of future all-sky transient searches.
8. SKA1 Low: EOR will be using approximately 20% of the observing time. This strongly argues to build a telescope that is suitable for a much broader range of science, including transients.
9. SKA1 Low: We strongly support the provision of multiple beams. Sixteen beams, even with conserved bandwidth, doubles the survey/transients Figure of Merit.
10. Baseline design: “For example, an external transient trigger to permit the delay buffers to be stored when a transient is suspected to have occurred; a stream of a small fraction of the incoming data streams could be provided in raw form to a facility for experiments to be carried out that cannot be precisely defined in advance.”: this sounds good, but “could be provided” doesn’t sound like a firm design feature or guarantee that it will happen.
11. Baseline design: “Transient triggers and third party equipment are optionally detailed by may not be present in the final design. Third party equipment potentially includes (but may not be limited to) transient processors.”: sounds good, but doesn’t sound like a firm design feature or guarantee that it will happen.
12. We often talk about how transient science requires a responsive and flexible telescope. Given that SKA-Mid and SKA-Survey will grow out of MeerKAT and ASKAP, how much do we need to worry about limitations to the system’s flexibility that are already introduced by inheriting these legacy systems? Do we need to be on top of guiding the MeerKAT and ASKAP design to prevent future problems with the SKA? Are just the dishes being reused, or will the SKA inherit the entire associated telescope management system and data flow architecture?
13. Baseline Design: “The beamformer must be designed to support beams from any number of sub-array.”: Will the correlator also be able to run in a sub-arraying mode?
14. Baseline Design: “Third party equipment potentially includes (but is not limited to) transient processors. It is anticipated that potential access to a limited subset of data will be made available providing the provision of spigots.”: good.
15. Baseline Design: “The buffer at the front end of the correlator may provide

the ability to freeze and download the data flagged by external transient detection equipment.”: good, but “may” is not very reassuring.

16. In general, the Baseline Design makes some interesting comments about what technical capabilities can be added for transients, but these are mostly non-committal statements. What are we supposed to do here? Certain considerations have not explicitly been left out, but at the same time they are worded in such a way that they could immediately be cut in the (likely?) event of a de-scope.

17. The appropriate Figure of Merit (FoM) is not being used. For TSWG (and indeed for other WGs) the appropriate survey FoM should go as $s^{-3/2}$ not s^{-2} (all other terms the same), where s is the r.m.s. sensitivity. This places a greater emphasis on field of view compared to raw sensitivity.

Appendix 2:
Detailed itemised comments on the SKA1 Concept of Operations document:

1. Section 2.

Somewhat disappointed to see the 'old' Key Science Programs listed here, since we are in the process of reviewing/revising these in 2014.

2. Section 2.6

"T*he primary success metric for the SKA Observatory will be the significance of its role in making fundamental scientific discoveries and facilitating overall scientific progress, expressed as high impact, peer-reviewed scientific papers using SKA data.A8 Additional success metrics such as the total number of users, etc., will also be developed and measured.*"
– all of these metrics will be greatly enhanced by commensal analysis of data and public release of alerts.

3. Section 3.3

Strongly prefer first Access Policy option, namely <u>Open Access</u>.

4. Section 4.3

"*Some time-domain processing will also take place in near-real time. Detailed definition of data processing requirements and limitations for both imaging and time-domain processing will occur during the design phase of the SKA.*"
– encouraging but too vague

5. "*Real-time processing will be carried out only where warranted (e.g. events that require a decision, action or response). Otherwise the processing system will be designed to keep up with the average load.*"– discouraging, but equally vague
6. Section 4.4

"*The SKA Observatory will calibrate SKA data and make science-ready data and ancillary products available to the users.*"
– "Science-ready" is not a meaningful term in isolation. What is "Science-ready" will differ wildly depending on the observing mode and science case, and will need to be defined appropriately. In some cases, such as transient spigots, the data products will assuredly not be delivered science ready by the observatory, and that is ok.

7. Section 4.5

"*The final distributed data products of the Observatory, together with relevant metadata, will be stored in a science archive system that incorporates appropriate data delivery and interoperability interfaces, including Virtual Observatory interfaces*"
– the reference to VO is appropriate. VO also supports VOEvents of course.

8. Section 5.5

"*As the full system capability will not be delivered on day one and technology advances will inevitably lead to incremental system upgrades, the telescopes should be designed to cater for this.*" – good
"..*but the SKA design should assume that substantial parts of the infrastructure will undergo cycles of minor and major maintenance as well as upgrading to follow the scientific and technical state-of-the-art. The length of these cycles will probably differ between components of the SKA telescopes.*" – good(ish)

9. Section 6.3

"*These outreach activities will be coordinated (not controlled) by the Global Headquarters. The GHQ will, however, also have staff engaged in public relations activities. The SKA Observatory, through the GHQ, will have responsibility for control of the SKA brand.*"

The scale of outreach performed by GHQ is completely unclear. How exactly is the observatory / GHQ going to “control the brand?” What does that mean?

10. Section 7.3

“*Three precursor telescopes are present or under construction on the SKA sites in the host countries: MeerKAT in South Africa and The Australian SKA Pathfinder (ASKAP) and Murchison Widefield Array (MWA) in Australia. Portions of MeerKAT and ASKAP and their infrastructure will be included in the SKA, maximising their use based on a cost benefit analysis.A23 Disruption to the precursor science programmes will be minimised while expediting the build-out of SKA Phase 1 at both sites.A24 The planned precursor programs are expected to take approximately 5 years at a high level of availability.*”
– bit too vague. How much of e.g. MeerKAT control systems will be morphed in SKA1 Mid control systems?

11. Section 8.1

“*Normal observing will consume most observing time for both Key Science Projects and PI Projects (section 3). Projects observed in this mode will be characterised by the lack of special requirements for equipment, support, or scheduling. PIs and project team members will not normally participate in the observing or receive the data in real time.*”
– concerning. The document really needs to steer away from these defensive statements.

12. Section 8.2

“*Some time-dependent phenomena will require observations made at a fixed time. The GHQ will schedule such observations*.” – encouraging but too vague

13. Section 8.3 and section 8.4

“*For example, where the SKA maintains appropriately documented standard interfaces, custom hardware may be used for data acquisition, or user-provided software may replace the standard data pipeline. Data may be provided in real time*.”
This section does not make clear the distinction between triggered proposals submitted at a normal deadline that require immediate automatic override, triggered proposals submitted at a normal deadline where the response is less time-sensitive (e.g. within 1 day/1 week/1 month, and hence could be approved/denied after the trigger submission) and out-of-deadline (“DDT”) proposals in response to a truly unexpected event.

14. Would a transients commensal real-time search be a “custom experiment”? We hope not. It if is, then we need clarity on how it relates to section 4.5 (data archive). Would a VOEvent or other alert message generated from this experiment be an archivable SKA data product? Would a lightcurve derived from image data and used to inform production of the alert?

15. Section 8.5

“*Commensal observing is defined as observing where two or more projects share the same dataset. Because Observatory resources will still be required to support commensal observing, the normal processes of review will apply to each project. Commensal projects shall not be allowed to exploit data in ways allocated to other projects sharing the same data. An extension of this use case is a project that requests further processing of archival data, within its proprietary period, for other goals than the primary program; if such projects can be facilitated, they will likewise go through the normal processes of review. The question of rights to serendipitously discovered phenomena is subject to future discussions on the Access Policy.*”
This is one of the most important paragraphs. It does not appreciate the scope or openness of what we’re proposing, and will lead to a reduction in the scientific outputs of the SKA. We are committed to “serendipitous discoveries” becoming public alerts.

16. Section 8.6

"*VLBI will initially be a subclass of custom observing.*"

How long is initially? Why does it have to be custom?

17. Section 8.8

"The scientific need for subarrays will be determined as a result of an analysis of use cases.", with a footnote "Position paper needed from SWG.". – subarrays are absolutely essential, for multiple science cases, including responding to transients, and VLBI observations (again relevant to transients). This should be changed. "Position paper" is never defined, but it is really not necessary (and here the ConOps conflicts with other documents, where subarrays are already listed as a requirement, albeit with a lot of TBDs on the exact specs).

Participants:

Transients SWG: R Fender (chair), Ben Stappers, Adam Deller, Jason Hessels, Jean-Pierre Macquart, Zsolt Paragi, Aris Karastergiou, Evan Keane, John Swinbank, Gregg Hallinan, Richard Armstrong

SKA Office: R Braun, T Bourke, J Wagg, J Green, P Dewdney (partly), W Cotton (partly), + 2 others

Workshop held 27-28 January 2014, Jodrell Bank