Publications of the Astronomical Society of Australia (2026), **43**, e085, 19 pages
doi:10.1017/pasa.2026.10218

CAMBRIDGE UNIVERSITY PRESS

Research Article

The Murchison Widefield Array Phase III upgrade: Sensitivity doubled, number of baselines quadrupled, flexibility enhanced, and EoR observations optimised

Steven J. Tingay[1], Melanie Johnston-Hollitt[2], Randall B. Wayth[1,3,4], Tom A. Booler[1], Jake Jones[1], Yajun Wu[5], Jy Gan[5], Gregory Sleap[1], Andrew McPhail[1], Cary Wintle[1], Andrew Williams[1], Christopher J. Phillips[1], Luke Verduyn[1], David Emrich[1], Phillip Giersch[1], Christopher J. Riseley[6,7], Stefan Duchesne[8], Cathryn M. Trott[1,4,8], Dev Null[9], Bradley W. Meyers[9], Chuneeta D. Nunhokee[1,4], Nichole Barry[10], Lauren Dressler[11], Jade Ducharme[12], Bryna Hazelton[11], Morgan Lee[12], Elias Lilleskov[11], Miguel Morales[11], Jonathan Pober[12], Zhiqiang Shen[5,13], Xiang-ping Wu[14], Xiaoyu Hong[5], Miroslav D. Filipović[15], Steven E. Tremblay[16] and Mia Walker[17]

[1]International Centre for Radio Astronomy Research, Curtin University, Bentley, WA, Australia, [2]Curtin Institute for Data Science, Curtin University, Perth, WA, Australia, [3]SKA Observatory, SKA-Low Science Operations Centre, Kensington, WA, Australia, [4]Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Bentley, Australia, [5]Shanghai Astronomical Observatory, Chinese Academy of Sciences, Shanghai , People's Republic of China, [6]Astronomisches Institut der Ruhr-Universität Bochum (AIRUB), Bochum, Germany, [7]Ruhr Astroparticle and Plasma Physics Center (RAPP Center), Bochum, Germany, [8]Australia Telescope National Facility, CSIRO, Space and Astronomy, Bentley, WA, Australia, [9]Australian SKA Regional Centre (AusSRC), Curtin University, Bentley, WA, Australia, [10]School of Physics, The University of New South Wales, Sydney, NSW, Australia, [11]Department of Physics, University of Washington, WA, USA, [12]Brown University, Providence RI, USA, [13]Key Laboratory of Radio Astronomy and Technology, Chinese Academy of Sciences, Beijing, People's Republic of China, [14]National Astronomical Observatories, Chinese Academy of Sciences, Beijing, China, [15]Western Sydney University, Penrith South DC, NSW, Australia, [16]National Radio Astronomy Observatory, Socorro, NM, USA and [17]Space Science and Technology Centre, Curtin University, Bentley, WA, Australia

Abstract

We describe the latest iteration of upgrades (designated Phase III) to the Murchison Widefield Array (MWA), in the fourth paper in a series that covers the evolution of the telescope from design concept to initial operational facility, and through two major upgrades. As part of the Phase III upgrade of the MWA, we report the completion of work to design, build, and deploy a new fleet of digital receivers that further optimise the MWA for Epoch of Reionisation observations. These receivers complement existing receivers, such that the MWA now supports the full correlation of all 256 antenna tiles currently in the array. This step releases the MWA from the prior constraint of having to correlate only 128 of the 256 tiles at any given time, which means that the maximum instantaneous sensitivity of the MWA is doubled and the maximum number of interferometric baselines is approximately quadrupled. The upgrade is fundamentally enabled by the new MWAX correlator and various other improvements to the MWA sub-systems. In this paper we describe the new digital receivers and the other improvements that result in the Phase III system. A range of operational benefits arise from the upgrade and scientific flexibility is increased. We also comment on the transition from the MWA to the SKA-Low facility near the end of the decade, including a description of some unique science opportunities utilising joint MWA/SKA-Low data during the Science Verification phase of the SKA-Low Array Assembly 2 (AA2) period.

Keywords: Radio Interferometry; Instrumentation; Radio Astronomy

(Received 27 March 2026 UTC; revised 11 May 2026 UTC; accepted 31 May 2026 UTC)

1. Introduction

The Murchison Widefield Array (MWA) is the low frequency Precursor[a] for the Square Kilometre Array (SKA; Dewdney et al. 2009) in Australia, located at the site upon which the low frequency component of the SKA Observatory (SKAO), SKA-Low, is currently being constructed, Inyarrimanha Ilgari Bundara, the CSIRO (Commonwealth Scientific and Industrial Research Organisation) Murchison Radio-astronomy Observatory, in Western Australia.

The original conceptual design for the MWA and some commentary on the early phases of the project are described by Lonsdale et al. (2009) and Tingay (2025), respectively. The as-built MWA, described by Tingay et al. (2013), commenced operations on 20 June 2013. In 2018, the MWA team completed the first major upgrade of the facility, described by Wayth et al. (2018) and post-upgrade referred to as Phase II (meaning the original

Corresponding author: Steven Tingay; Email: s.tingay@curtin.edu.au

Cite this article: Tingay SJ, Johnston-Hollitt M, Wayth RB, Booler TA, Jones J, Wu Y, Gan J, Sleap G, McPhail A, Wintle C, Williams A, Phillips CJ, Verduyn L, Emrich D, Giersch P, Riseley CJ, Duchesne S, Trott CM, Null D, Meyers BW, Nunhokee CD, Barry N, Dressler L, Ducharme J, Hazelton B, Lee M, Lilleskov E, Morales M, Pober J, Shen Z, Wu X, Hong X, Filipović MD, Tremblay SE and Walker M. (2026) The Murchison Widefield Array Phase III upgrade: Sensitivity doubled, number of baselines quadrupled, flexibility enhanced, and EoR observations optimised. *Publications of the Astronomical Society of Australia* **43**, e085, 1–19. https://doi.org/10.1017/pasa.2026.10218

[a]https://www.skao.int/en/explore/precursors-pathfinders.

© The Author(s), 2026. Published by Cambridge University Press on behalf of Astronomical Society of Australia. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

MWA became known as Phase I). As an adjunct to the technical descriptions, papers describing the science aspirations and science outcomes for MWA Phases I and II have been published by Bowman et al. (2013) and Beardsley et al. (2019), respectively. Shortly after the completion of the Phase II upgrade, discussions commenced in late 2019 regarding a Phase III upgrade. Work toward Phase III progressed in stages and was completed in 2025. This paper describes the as-built Phase III MWA system.

Below we provide a brief summary of the progression of the MWA across Phases I and II, to provide context for the Phase III description, but we will not dwell on the details of Phases I and II. The reader is referred to the comprehensive descriptions cited above. However, before that summary, it is worth describing some of the high-level motivations for the MWA, that have guided the technical progression across its lifetime.

1.1. High-level motivations for the MWA's development path

A strong emphasis in the original conceptual design of the MWA was the experiment to detect redshifted emission from the hyperfine transition of neutral hydrogen expected to exist in the early Universe (Furlanetto et al. 2006), the so-called Epoch of Reionisation (EoR). As the realisation of the MWA concept took shape during the late 2000s and early 2010s, this emphasis was retained, but the project goals evolved to support a wider range of science. A somewhat more general-purpose instrument emerged, with the capacity to support future upgrades.

Through Phase II, and now Phase III, the EoR experiment continued to motivate an evolution of the instrument, but the general purpose nature of the MWA has also been retained. The flexible capacity to support upgrades implemented in Phase I has been very effectively exploited to realise balance across a wide range of science goals. This, in turn, has led to the high scientific productivity and impact of the MWA[b] and its ability to support a diverse community. The MWA consortium currently consists of 256 individual members from 39 institutions, in six-member countries: Australia; Canada; China; Japan; Switzerland; and the USA, and seven countries with Associate Members.[c]

In addition to the MWA's technical and scientific goals, as a low frequency radio telescope of significant scale on one of the world's best sites for radio astronomy, the MWA's motivations have also been derived from its status as an SKAO Precursor (an SKAO technology demonstrator located at an SKAO site). Broadly speaking, the SKAO Precursors have helped inform the technical design, prototyping, and testing of the SKAO system, but have also given the global community an opportunity to make early progress in areas of key science for the SKAO, importantly entailing opportunities to grapple with some of the challenges posed by next generation radio telescopes that will culminate in the SKAO. Fulfilling these motivations has helped to grow the SKAO science community and has provided a well-trained new workforce that is now heavily engaged in SKAO construction, commissioning, and operations.

The MWA, being a low frequency Precursor on the SKA-Low site, has been in a unique position to maximally support the path to SKA-Low, hosting multiple generations of SKA-Low station prototypes (Hall et al. 2013; Sutinjo et al. 2015; Wayth et al. 2017; Benthem et al. 2021; Wayth et al. 2022; Macario et al. 2022), that have allowed the design and performance maturation of many SKA-Low sub-systems.

Other SKAO Precursors are HERA (low frequency in South Africa Berkhout et al. 2024), ASKAP (mid frequency in Australia Johnston et al. 2008), and MeerKAT (mid frequency in South Africa, at the SKA-mid site Jonas & MeerKAT Team 2016). At low frequencies (significant capabilities well below approximately 1 GHz), SKAO Pathfinder instruments in the form of LOFAR (van Haarlem et al. 2013), uGMRT (Gupta et al. 2017), NenUFAR (Zarka et al. 2015), CHIME (Bandura et al. 2014), 21CMA (Zheng et al. 2016), Tianlai (Chen et al. 2016), and the LWA (Kassim et al. 2010) have all individually and collectively made similar contributions to the path toward SKA-Low.

1.2. Technical progression of the MWA, Phases I and II

The basic unit of the MWA's aperture is the 'tile', an analog beamformed aperture array consisting of 16 dual-polarised bowtie dipole antennas, arranged on a 4×4 grid, approximately 5 m on a side and attached to a ground plane. The Phase I MWA consisted of 128 tiles, organised around a dense core distribution and a more extended distribution, with a maximum baseline of approximately 3 km. The antenna distribution provided a highly filled (u, v) coverage up to the longest baselines (Tingay et al. 2013), supporting general purpose science, and good surface brightness sensitivity at angular resolutions relevant to EoR science.

Sixteen receiver systems placed in the field each supported eight tiles, digitising and filtering the critically sampled (Nyquist sampled for the relevant bandwidth) radio frequency (RF) signals from the tiles into 24×1.28 MHz 'coarse' channels that could be tuned within the 70–300 MHz accessible frequency range (Prabu et al. 2015). The coarse channels were further processed into 'fine' channels (128×10 kHz channels per coarse channel), before being cross-correlated in a GPU-enabled correlator (Ord et al. 2015). The resulting visibility data were archived and made available to users for processing. Various other observational modes, such as the capture of the fine channel voltage data, were implemented in an early form (Tremblay et al. 2015).

An important decision in the final realisation of Phase I was to build enough infrastucture to eventually support 256 tiles, rather than 128. The project did not possess enough resources to realise a full 256 tile array, but could realise the infrastructure to support a future upgrade. This choice was made, rather than build a (for example) full ∼160 tile array. The reason being that adding additional infrastructure capacity to the array at a later date was deemed likely to be highly expensive and complicated. Whereas the future utilisation of latent infrastructure would be less expensive and could be realised progressively if needed and as resources were secured.

The Phase I MWA operated from 2013 to 2018. In 2018, the Phase II MWA upgrade was completed. In essence, very little changed in the signal path described above for Phase I. The upgrade primarily consisted, as alluded to above, in the doubling of the number of tiles in the array, from 128 to 256. The additional tiles formed a re-configured core distribution (again optimised in pursuit of the EoR signal), but were also distributed to realise longer baselines, with the maximum baseline increased to approximately 5.3 km. The long baselines were enabled by tiles that were powered by independent installations of solar panels plus battery storage.

[b]For the MWA publication library, see https://ui.adsabs.harvard.edu/public-libraries/SB5_iHeZTxCZDCj5kUuR6Q.

[c]https://www.mwatelescope.org/people/.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

Importantly, the receiver count did not change and the correlator did not change, which required a fundamental shift in the MWA's operational model, switching between a 'compact' configuration of 128 tiles (maximum baseline of approximately 750 m) and an 'extended' configuration of 128 tiles (maximum baseline of approximately 5.3 km), requiring the regular physical movement of receivers in the field to achieve the reconfiguration (generally ~twice per year). The Phase II MWA operated between 2018 and 2025.

1.3. MWA Phase III, at a high level

Over the course of 2019–2025, the Phase III MWA upgrade was conceived and progressively deployed. A detailed technical description is the primary point of this paper, but at the high level the goal was to realise a fully correlated 256 tile array. The first major element of Phase III to be deployed was a new and upgraded GPU-accelerated correlator, named MWAX (Morrison et al. 2023), described in Section 5, capable of correlating the signals from 256 tiles (plus reserving capacity for some other signal processing operations). MWAX was deployed in 2022.

The step to a 256 tile correlated array required the development and deployment of a new fleet of receivers, to augment the receivers that had been in place since 2013. Two varieties of new receivers were developed and deployed, such that 18 new receivers, plus the original 16 receivers, are now available to service the 256 tiles (still eight tiles per receiver). Of the 18 new receivers, 16 are of one type and two are of a second type, described in detail in Sections 4.1 and 4.2, respectively. Note that 32 receivers are required for 256 tiles, with two receivers available for sparing from the pool of 34 receivers.

The new receivers utilise modern electronic components and are more capable than the original receivers. A very important feature of the hardware and firmware of the new receivers is that they support an oversampled filterbank architecture as well as the original critically sampled architecture. This means that coarse channel edge artefacts present due to the critically sampled receivers, that presented challenges for the EoR experiment, have been removed in the new receivers. Again, the evolution of the MWA has been guided by the evolving understanding of how to pursue the EoR experiment, as well as general purpose improvements. As such, the new over-sampled receivers have been installed on the core tiles, which produce the short baselines relevant for the EoR experiment. The new receivers can also operate in a critically sampled mode, for compatibility with the original receivers.

With all 256 tiles fully populated with receivers, the MWA can now be flexibly and remotely reconfigured into a range of arrays with different numbers and distributions of tiles, without any physical intervention. This represents a fundamental alleviation of the operational burden of physical reconfigurations. Thus, both scientific and operational benefits of significance have been realised with MWA Phase III.

1.4. The future of the MWA, at a high level

As the SKA-Low construction phase draws to a close near the end of this decade, and the early cycles of SKA-Low science operations are contemplated on the same timescale, the natural end of the MWA's mission is mid-2030. Currently, this is the anticipated end of MWA operations. The Phase III upgrade represents the culmination of the MWA's technical capabilities; the operational objectives until mid-2030 are to maintain this capability and continue to offer the community incremental improvements related to the increased flexibility the Phase III upgrade has enabled.

As briefly explored in Section 9, a scientifically interesting overlap period between MWA Phase III and SKA-Low Array Assembly 2 exists in the 2027–2030 timeframe.

1.5. Structure of this paper

Below we step through the various technical developments that have contributed to the Phase III upgrade, with design information and as-built descriptions. In Section 2, we describe the characteristics of the full 256 tile array, with a focus on the different array options that have been initially defined and offered for observations by users. This gives the interested science audience an upfront feel for the initial possible (u, v) coverages, and therefore imaging performances during Phase III. In Section 3 we provide an overall system architecture overview to orientate the interested technical audience. In Section 4, we describe in detail the two new receiver systems. In Section 5, we describe the MWAX correlator, with an emphasis on improvements and additions to the correlator made since the detailed MWAX description paper was published (Morrison et al. 2023). In Section 6, we describe a modification that supports the overall system, the conversion signal chain connections for some tiles from coaxial cables to fibre. In Section 7 we briefly discuss operational benefits, and science implications and opportunities for the Phase III MWA in Section 8. We wrap up with discussion, including some comments on the transition from MWA Phase III to SKA-Low in Section 9. A full science opportunity paper will appear in due course, to complement this technical description.

2. Tile distribution, initial array options, and example (u, v) coverages

The distribution of the 256 tiles in the MWA has not changed since the Phase II description in Wayth et al. (2018). However, the Phase III upgrade fully populates these 256 tiles with receivers, such that no physical reconfiguration of the array is required in order to select a sub-set of tiles for observations. Such a selection is now possible in a software-defined manner, without physical intervention.

The ability to select sub-sets of tiles can have a range of benefits. The sub-sets can be tailored to support particular science cases and aggregate data rates, for example to select no more than is required for a particular experiment, thereby minimising the impact on the MWA archive and the data handling and processing costs for users.

For initial Phase III operations, five pre-defined array configurations have been made available to users. As Phase III operations progress and mature, it is likely that further configurations will be made available, and ultimately fully user-defined configurations. The initial five configurations are listed below and briefly described with example snapshot (u, v) coverages. In all cases, the (u, v) coverages are presented for a single time point, for a zenith pointing, for a single frequency (at 150 MHz), and with all conjugate (u, v) points shown, in order to simplify the presentation for the reader. The reader needs to keep in mind that realistic (u, v) coverages will generally take advantage of bandwidth synthesis and Earth rotation synthesis to achieve even denser and better filled coverages.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

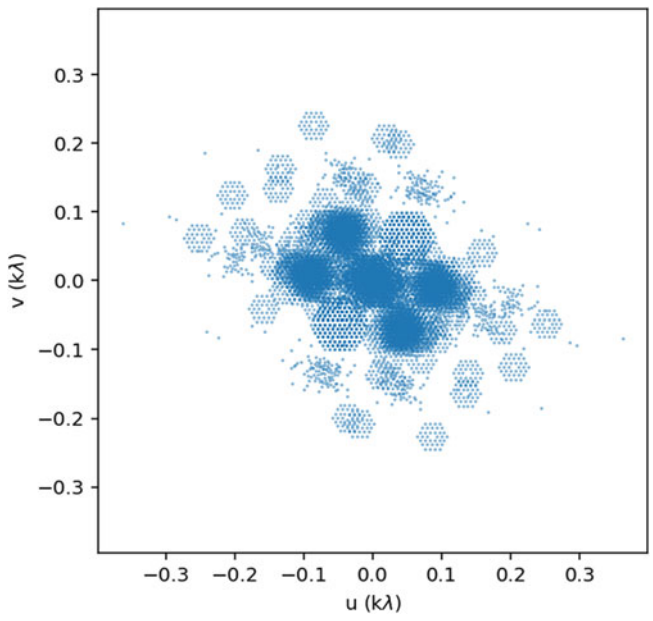


Figure 1. The 'Phase II Compact' array configuration (u, v) coverage.

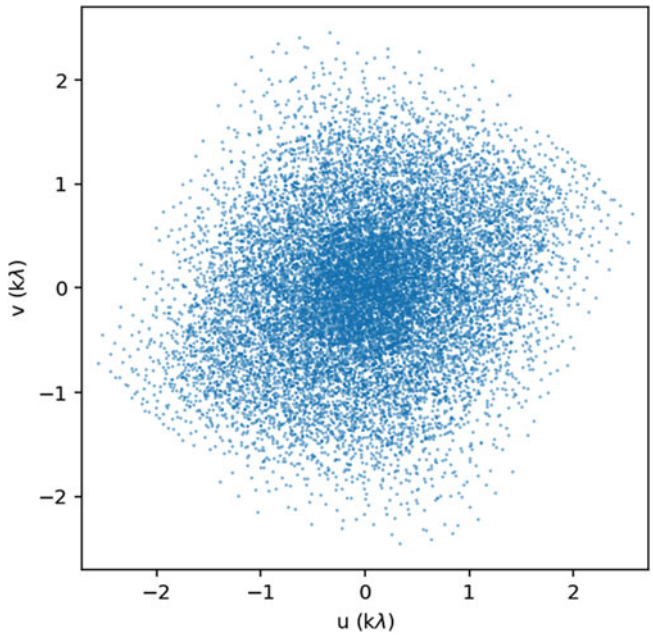


Figure 2. The 'Phase II extended' array configuration (u, v) coverage.

Point spread functions and other parameters (such as data rates, sensitivities etc) corresponding to these (u, v) coverages are not provided, as these products will depend on a range of user choices (for example in weighting schemes for the visibilities) and the intention of this section is to provide a high level and illustrative insight rather than exhaustive performance metrics.

2.1 Phase II compact

This configuration realises a 128 tile compact array utilising only (new) oversampled receivers, configured to maximise surface brightness sensitivity and take full advantage of the greatly improved spectral response of the oversampled receivers. This configuration has been specifically designed to support the EoR experiment, as well as any other observations requiring a compact array and/or high surface brightness imaging (for example, pulsar observations). The characteristic baseline distribution caused by the inclusion of the 'Hex' tiles established in the Phase II upgrade is apparent in the (u, v) coverage (Figure 1), evidenced by the repeated appearance of the hexagonal groupings of (u, v) spacings due to baselines from an individual tile to the sets of 'Hex' tiles.

2.2 Phase II extended

This configuration realises a 128 tile extended array utilising critically sampled receivers and configured for angular resolution and (u, v) coverage, while managing data volume. This configuration produces a (u, v) coverage with maximum angular resolution and uniform filling, suited to applications where angular resolution is prioritised over sensitivity and data volume is a consideration (Figure 2).

2.3. Phase I plus solar

This configuration realises a 184 tile extended array utilising critically sampled receivers and configured for maximum angular resolution and (u, v) coverage (Figure 3), while moderately managing data volume. The configuration adds the Phase II tiles established using individual solar power plants to the 128 tile Phase I configuration to suit applications that require maximum angular resolution and require higher sensitivity, while data volume management is less of a consideration.

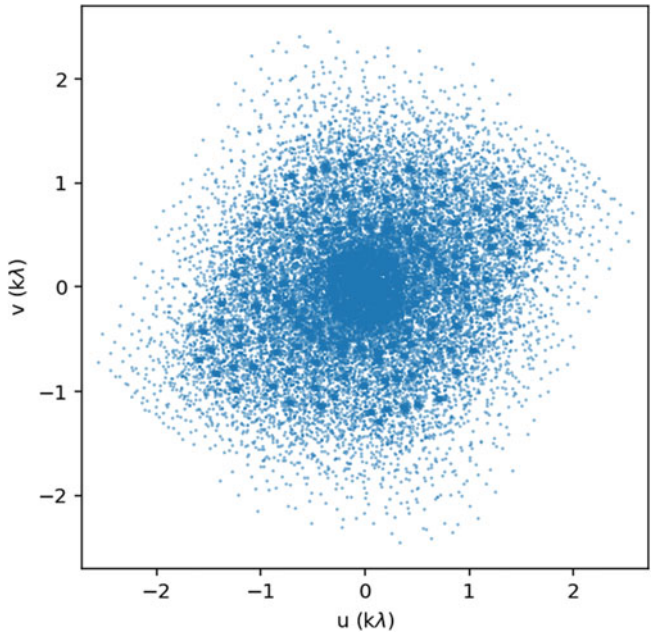


Figure 3. The 'Phase I plus Solar' array configuration (u, v) coverage.

2.4. Phase I plus Hexes

This configuration realises a 200 tile array utilising critically sampled receivers and configured for sensitivity at moderate angular resolutions, with minimal management of data volume.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

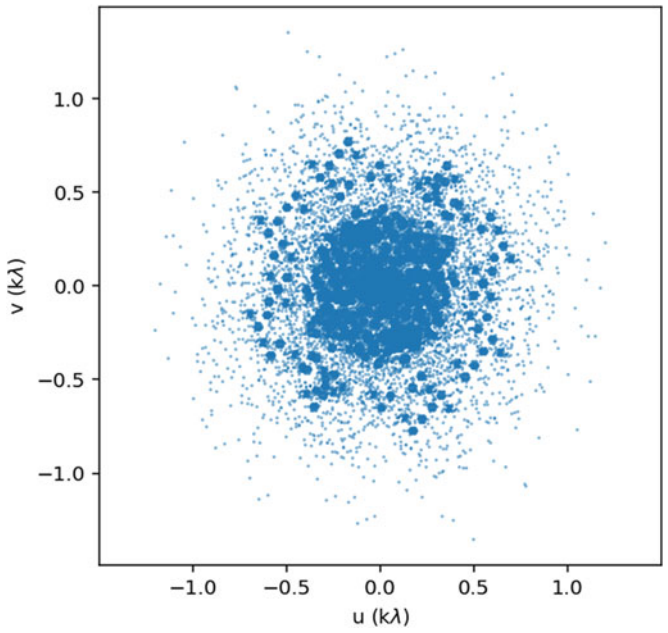


Figure 4. The 'Phase I plus Hexes' array configuration (u, v) coverage.

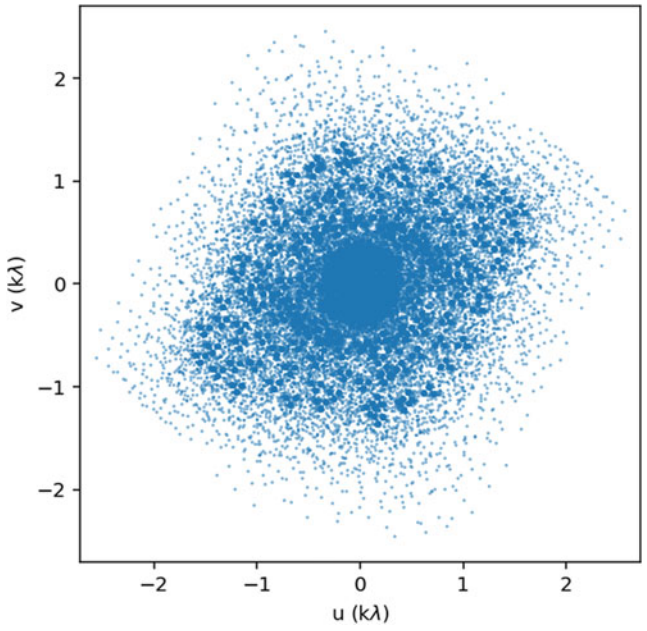


Figure 5. The 'Full Array' array configuration (u, v) coverage.

Analogous to the 'Phase I plus Solar', this configuration adds the core 'Hex' tiles established in Phase II to the 128 tile Phase I configuration to suit applications that require only moderate angular resolution but high sensitivity, while data volume management is less of a consideration. The (u, v) coverage is shown in Figure 4

2.5. Full array

This configuration realises the full 256 tile array (all available tiles) utilising critically sampled receivers, for maximum angular resolution, maximum sensitivity, and maximum (u, v) coverage (Figure 5), at the cost of maximum data volume.

3. System architecture overview

Further details on the architecture components described at a high level in this Section are described in Sections 4, 5, and 6.

The Phase III array is marked by two significant evolutions, and a number of enabling changes to supporting subsystems and infrastructure. The first significant development is the replacement of the original, Phase I, correlator (Ord et al. 2015) with the MWAX correlator (Morrison et al. 2023). Implementation of the MWAX correlator (following a standard Fourier Transform before cross-multiply architecture: FX) lays the foundation for the MWA to correlate more than 128 (dual polarisation antenna tile) inputs. The second is the addition of 16 new digital receivers, developed in collaboration between Shanghai Astronomical Observatory (SHAO) and Curtin University, referred to as SHAO receivers. The addition of the SHAO receivers means that, for the first time, all 256 antenna tiles can be connected to the operational array at the same time.

The Phase III array also includes two additional digital receivers based on National Instruments (NI) FlexRio hardware. Originally developed as evaluation systems, the NI-based receivers have been incorporated into the array to release two Phase I receivers (Prabu et al. 2015) from service, to constitute a spares pool. This was done with a view to prolonging the operational life of the Phase I receiver fleet.

The Phase III array also sees the conversion of additional MWA antenna tiles from electrical signal outputs (on coaxial cable) to optical signal outputs (on fibre optic cable). This change is necessary to overcome practical impediments, including losses in coaxial cables, that otherwise constrain the disposition of receivers, and utilisation of the power and signal infrastructure within the array. Conversion is achieved by introduction of two new sub-systems, the tile interface unit (TIU), which transitions the antenna tile output to RF on fibre (RFoF) and facilitates tile monitoring and control signals, and the pad power supply unit (PPSU), which provides power to a group of converted tiles.

These Phase III changes extend the diversity of system architectures that resulted from the Phase II upgrade (Wayth et al. 2018). Phase III's heterogeneous system architecture therefore includes five signal path configurations (C1–C5), summarised in Figure 6. Elements of the architecture that have been introduced as part of the Phase III upgrade are indicated in green in Figure 6.

Forty-eight tiles have a configuration C1 signal path. This was the signal path configuration for all Phase I tiles. The output of these tiles (from the analogue beamformer) is carried over a long coaxial cable and input to six Phase I receivers (Tingay et al. 2013).

The 56 tiles indicated in C2 are all long baseline tiles added to the array in Phase II. These tiles are accompanied by small solar photovoltaic (PV) array and battery systems to meet their power requirements. The C2 signal path includes a BFIF (beamformer interface) unit. The coaxial cable output of the analogue beamformer is input to a BFIF positioned immediately adjacent. The signal is transitioned to RFoF by electrical-to-optical converters within the BFIF. The BFIF also manages the local PV and battery power system (Wayth et al. 2018). All 56 tiles connect to seven Phase I receivers.

In Phase III, 8 tiles have been converted to configuration C3 by the introduction of a TIU. The coaxial cable output of the analogue beamformer is input to a TIU positioned immediately adjacent. Like the BFIF, the TIU transitions the signal to RFoF via electrical-to-optical converters. This sub-set of antenna tiles connects to one Phase I receiver.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

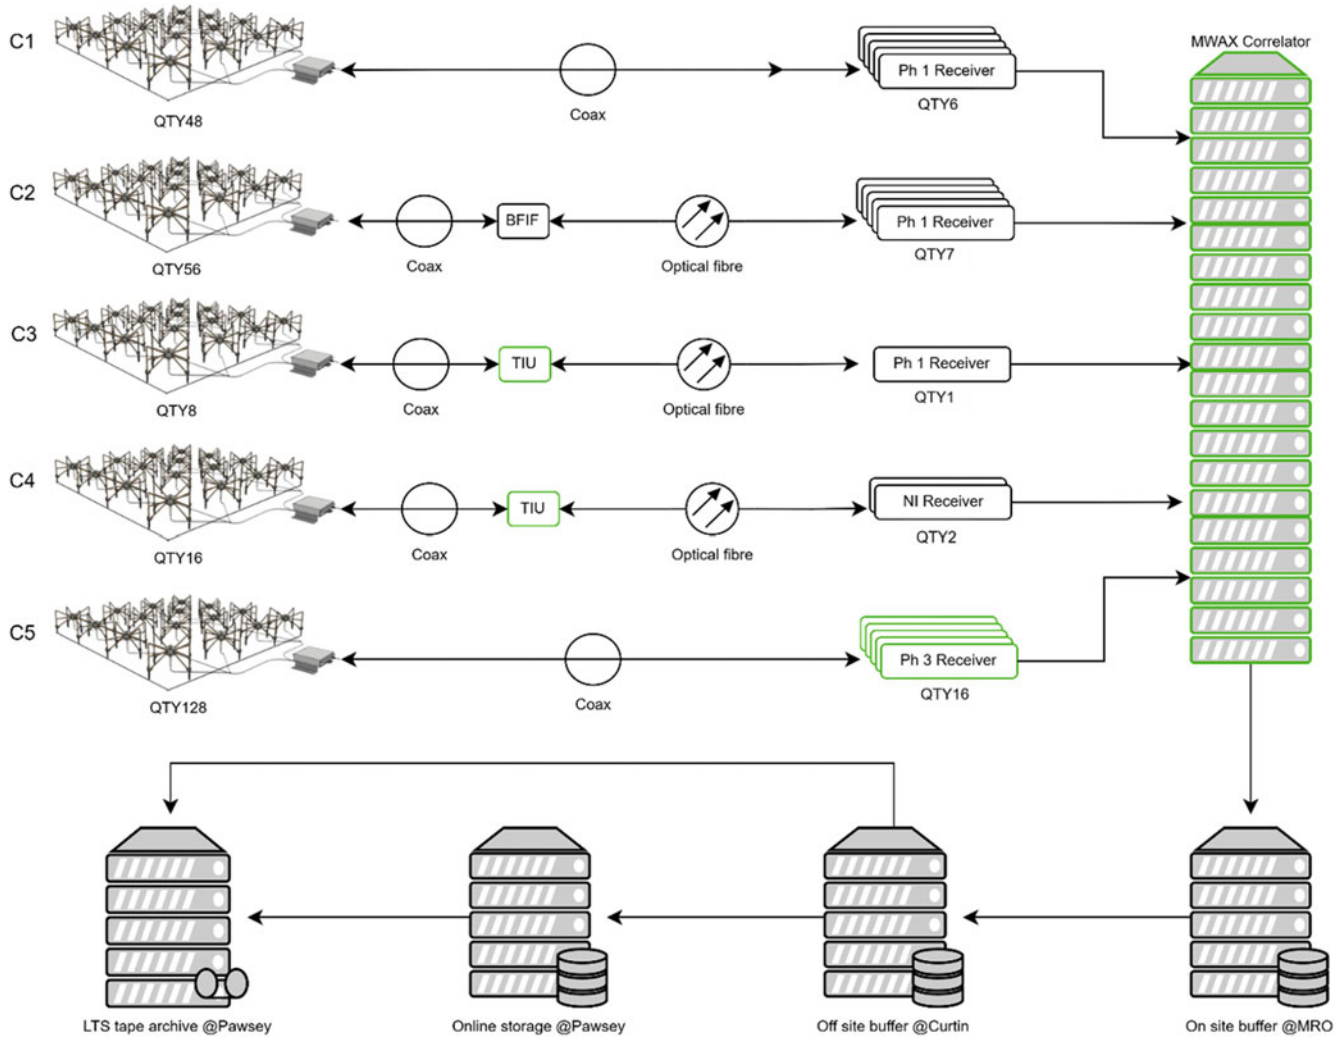


Figure 6. The five different signal path implementations (C1–C5) that exist in the Phase III array. Phase III additions are shown in green.

The 16 tiles that feature signal path configuration C4 also feature the TIU. The path mirrors that of C3 until the receiver. C4 uses two NI-based receivers.

Finally, 128 tiles have a configuration C5 signal path. The output of these tiles is carried over a long coaxial cable and input to 16 SHAO receivers.

The 256 antenna tiles included in the Phase III MWA are thus supported by 14 original Phase I receivers, two NI-based receivers, and 16 SHAO receivers. Configurations C1–C5 feature heterogenous analogue signal paths that utilise different combinations of transmission hardware and cable formats. This diversity of signal path architecture reflects the confluence of technical, budgetary, resource, and schedule considerations that have shaped the MWA project over 13-years of operations, as the technologies have also evolved.

4. New receivers

The Phase III receiver development comprised a large component of the overall Phase III upgrade and has resulted in a significant improvement in the MWA's capabilities. In addition, eight field cabinet enclosures that house these receivers with accompanying hardware comprise the majority of the remainder of the development.

The development of new Phase III receiver hardware was motivated by the need to address the limitations and shortfalls of the Phase I systems, in addition to creating a small spare pool for the event of a receiver failure. This activity presented the opportunity to deliver several key improvements to the MWA's receiver systems in order to meet the objectives of Phase III.

The most significant feature of the receiver upgrade is the introduction of an Oversampled Polyphase Filter Bank (OPFB) mode, which was a core requirement for new Phase III receiver designs. Since the outset of the MWA, Phase I and II data have suffered from aliasing near the coarse channel boundaries, effectively corrupting approximately 10% of the channel bandwidth. This not only reduces the amount the usable bandwidth, but also impacts on EoR science in particular due to the regular spacing of the resulting spectral artifacts (Kolopanis et al. 2023). Ultimately, this issue is caused by out-of-band signals aliasing back into the coarse channel passband, an effect inherent to critically sampled PFBs, as their passband width and sampling rate are equal, leaving no space for the filter to transition between passband and stopband. The OPFB resolves this by increasing the sampling-rate

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

by a factor of 32/25, but leaving the channel spacing and passband width unchanged at 1.28 MHz. This provides redundant regions of padding on each channel edge that can later be discarded without any loss in data, but allows the filter to transition between bands, thus eliminating contamination of the passband by aliased signals. The filter coefficients that are implemented in all Phase III receivers were designed to provide a stop-band attenuation of at least 60 dB and a pass-band ripple of no greater than ± 1 dB. The remaining passband ripple is corrected within the MWAX correlator to produce a flat spectral response (Morrison et al. 2023).

The Phase III receivers also feature major improvements to the dynamic range capability of both the front-end Analog to Digital Converters (ADCs) and the data stream that is ultimately transmitted to the correlator. The dynamic range describes the range of input power levels a system can handle without clipping or distorting the data. A sufficiently high dynamic range is important as it enables the system to be sensitive to weak signals while simultaneously being immune or unaffected by higher power signals such as solar activity or RF interference (RFI). The dynamic range of the Phase I receivers is low relative to modern hardware standards, with only 8-bit ADC front ends and 10-bit complex (5+5i) data streams, increasing the chance of distorted data and high levels of quantisation noise. The requirement for the Phase III receivers specifies a minimum 12-bit ADC and must transmit the channelised data in the form of 16-bit complex numbers ($8 + 8i$). Once ingested by the MWAX correlator, data streams from all receiver types are converted to floating point and are represented by 64-bit complex floats thereafter.

The Phase I receiver's use of proprietary data transport technologies such as 'Xilinx RocketIO' resulted in many constraints for the correlator ingest architecture, such as having few off-the-shelf solutions and requiring an additional media conversion stage prior to the correlator. The Phase III receivers are required to use standard Ethernet networking to transmit data to the correlator. Both new receiver designs (SHAO and NI) use 10 Gigabit Ethernet interfaces, simplifying the supporting infrastructure and allowing the use of off-the-shelf networking equipment. The Phase III receivers also directly transmit the data in the standard User Datagram Protocol (UDP) packet structure (Morrison et al. 2023) required for the MWAX correlator to ingest.

Finally, to observe in full 256 tile mode, the new receivers must operate in conjunction with the legacy Phase I receivers, and therefore they must be capable of operating in a compatibility mode that implements the same sampling rate and PFB of the Phase I receivers. Additionally, the new receivers must be capable of switching between oversampling mode and compatibility mode automatically, quickly, and reliably.

We have two flavours of new receivers that meet the same specifications, as the result of two hardware development paths. The NI receivers are off-the-shelf, and expensive. The SHAO receivers had a longer development time, but are much more cost effective than the NI receivers. The different solutions lend themselves to different project resourcing conditions that may prevail at any given time.

4.1 Shanghai Astronomical Observatory (SHAO) receivers

The SHAO receivers were designed and developed in collaboration with the Shanghai Astronomical Observatory (SHAO) which contributed 16 units in total towards the MWA Phase III upgrade.

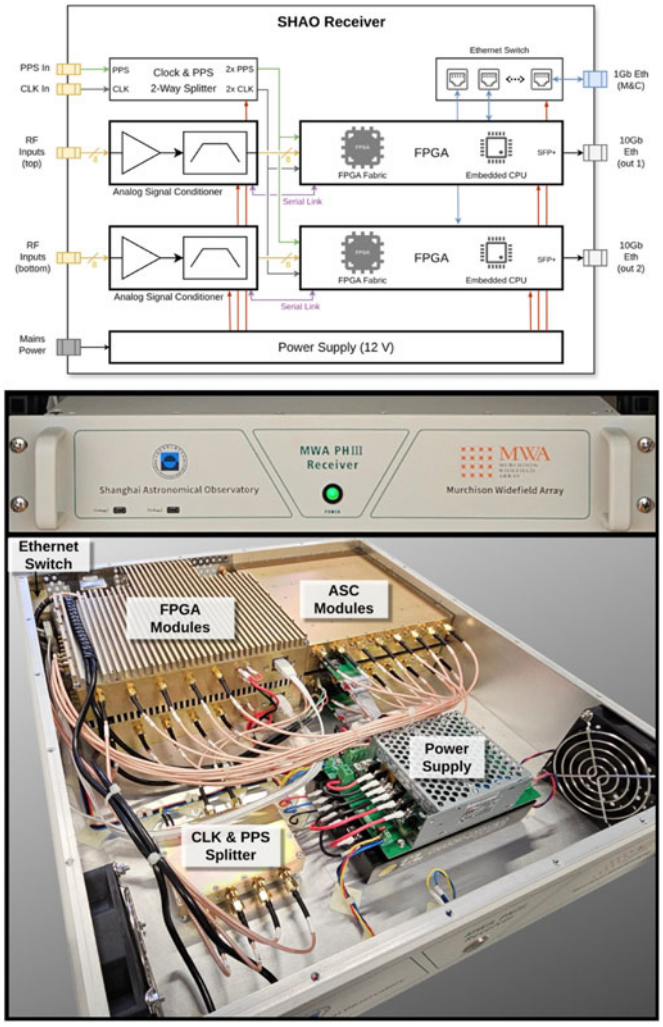


Figure 7. SHAO receiver schematic (top panel) and as-built hardware (bottom panel).

They were designed to be a standalone unit capable of digitising 16 RF inputs or the equivalent of 8 dual-polarisation tiles. Each unit has a 1 Gb Ethernet (RJ45) interface for Monitoring and Control (M&C), is synchronised by a 10 MHz clock and 1 Pulse Per Second (PPS) reference signals, and transmits the channelised data over two 10 Gb Ethernet Enhanced Small Form-factor Pluggable (SFP+) interfaces. Physically, each unit consists of an aluminium enclosure that is rack mountable in a standard 19-inch rack with a height of two Rack Units (RU).

Internally, a single SHAO receiver consists of several sub-modules and is largely split into two parallel sub-receivers with 8 inputs each (see Figure 7). They consist of two Analog Signal Conditioners (ASCs), two Field-Programmable Gate Array (FPGA) Modules, a two-way clock (CLK) and PPS distributor, a three port Ethernet switch, and an integrated 12 V power supply. Finally, each unit is fitted with two fans that force airflow through the enclosure.

The ASC modules were specifically designed for the MWA's requirements and have two main functions, filtering and gain adjustment. The first stage provides a bandpass filter that acts as an anti-aliasing filter for the ADCs. The passband of the filter ranges between 50 and 300 MHz with no more than 1 dB of ripple, and the stopband provides at least 60 dB of attenuation. The second stage

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

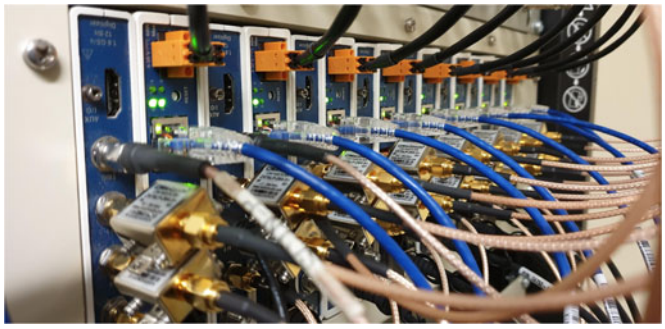


Figure 8. The as-built NI receiver, showing the simple nature of the COTS equipment and the straightforward packaging, in contrast to bespoke systems.



Figure 9. Design drawing of the environmental enclosure of the field cabinet enclosure, showing the major elements described in the text. Left panel shows air baffles. Middle panel shows protected ingress points. Right panel shows the enclosure empty and open.

provides independently configurable gains for each input between −40 dB and +20 dB, which is essential for the equalisation of power between tiles that have different cable lengths.

The FPGA modules consist primarily of the HTG-ZRF8 board from HiTech Global that hosts Xilinx ZYNQ UltraScale+ RFSoC ZU28DR FPGAs. These RF System on Chip (RFSoC) FPGAs have eight built-in ADCs accessible from the FPGA fabric in addition to an embedded ARM Cortex CPU that is used for M&C. They also have a daughter board that provides additional SFP+ ports for data transmission. These boards are housed within a custom-made RF shielded enclosure that provides both heat dissipation and electromagnetic compatibility (EMC) shielding.

The FPGA firmware was custom designed specifically to meet the needs for the MWA Phase III. Each receiver mode, oversampling mode and compatibility mode, represents a different version of the firmware that is programmed onto the FPGA when transitioning between modes. The management of the FPGA firmware and configuration is handled by the on-board CPU which runs a lightweight Linux distribution (PetaLinux) and custom software that interacts with the MWA M&C system.

4.2 National Instruments (NI) receivers

The National Instruments or 'NI' receiver is the second variant of new receiver also capable of meeting the requirements for Phase III (Figure 8). The MWA has a total of two NI receivers, each made up of several sub-components where one receiver comprises all the equipment required to digitise eight dual polarisation tiles outputs. Unlike the SHAO receivers, the NI receiver variant makes significant use of commercial off-the-shelf (COTS) equipment.

The NI receiver is primarily based on the National Instrument 7935R FlexRIO FPGA controller, and the NI-5772 Analog to Digital Converter Module. Each FlexRIO module, with its ADC sub-module, has two RF inputs capable of digitising a single dual polarisation tile output, and has an assortment of IO options. Of these, the system utilises the external ADC clock input (655.36 Msps), a trigger input used for 1PPS synchronisation, a 1 Gb Ethernet port for M&C, and a 10 Gb Ethernet SFP+ port.

A full NI receiver is comprised of eight individual FlexRIO modules that are physically housed in a 4U rack mountable enclosure alongside supporting hardware, most of which is concerned with either splitting or aggregating signals eight ways. Each receiver consists of a 24 V power supply module with eight individually switchable outputs that allows for a properly sequenced power up of the system. It also consists of a timing panel that hosts an eight-way Mini Circuits power splitter to distribute the external clock and a BG7TBL square wave distribution module to split the 1 PPS signal eight ways. Finally, each receiver requires both a 1 Gb and a 10 Gb Ethernet aggregation switch for the M&C and data output interfaces, respectively.

The firmware architecture for the NI receivers is similar to that for the SHAO receivers, in that they also implement a PFB filter, Digital Gain, and UDP Packetisation stages alongside additional telemetry features. The NI receivers are also primarily controlled by an on board CPU that can manage the configuration of the FPGA and communicate with the wider MWA M&C system. The firmware is developed using the National Instruments LABView FPGA software package.

4.3 Field cabinet enclosure

The field cabinet enclosure was designed to host rack mountable equipment within an environmentally managed cabinet and have sufficient EMC shielding such that is can be located in proximity to tiles (see Section 4.3.2). As part of the upgrade, eight field cabinet enclosures were fabricated and deployed. With two SHAO receivers per cabinet, this enabled the additional 16 receivers to be connected to 128 of the core distribution tiles. A field cabinet enclosure consists of two parts, an inner RF shielded cabinet and an outer environmental enclosure to protect the inner cabinet from the elements.

4.3.1 Environmental enclosure

The environmental enclosure is designed to house and protect the RF-shielded cabinet, as well as power, fibre, and coaxial cable connections. It comprises a galvanised square tube frame with white powder coated steel doors and roof structure. One end of the enclosure protects the gland plate and wave guides housed on the inner shielded cabinet that facilitate cable entries. A large, hinged hood on the enclosure protects these sensitive connection points. The other end features a set of large baffles that facilitate airflow to the air-conditioning plant installed at that end of the enclosure. The powder coating used on the environmental enclosure was specifically selected to improve thermal heat flows to the edges of the panels where it is dissipated via natural airflow. This feature contributes to a reduced internal heat load, moderating the demand placed on mechanical cooling. Figure 9 shows the primary features of the enclosure.

The design of the Phase III environmental enclosure is strongly informed by lessons learned from Phase I. The Phase III enclosure stands taller. The position and orientation of the internal equipment racks, and the clearance provided by the large hinged access

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

panels, ensure that a maintainer has comfortable standing and/or seated access to all major sub-systems. The hinged access panels have gas struts fitted to ensure that they can't close rapidly and cause injury to personnel. With the access panels raised, the enclosure also provides a measure of shade when the sun is high in the sky. This saves maintenance personnel having to locate and erect temporary shade as they move around the site.

Finally, the Phase III environmental enclosure is designed to be tamper evident. The major access panels are designed to be lockable, with all hinges recessed and fixings hidden. When the enclosure is properly closed and locked, accessing the sub-systems it houses results in readily identifiable, physical damage.

4.3.2 RF shielded cabinet

Like the environmental enclosure, the design of the Phase III RF shielded cabinet was substantially motivated by lessons learned from the configuration and packaging of the Phase I receiver system. The sub-systems installed within the equivalent Phase I cabinet were racked horizontally, and the cabinet was set low to the ground. This arrangement resulted in poor ergonomics which made installing and maintaining equipment within the cabinet awkward and taxing. The difficult access also increased the likelihood of poor physical connections at system interfaces.

The internal capacity and configuration of the Phase III cabinet are designed to maximise flexibility and utility, rather than tailored to a particular receiver implementation and specification. The internal arrangement of the cabinet is based on two standard 22RU equipment racks. The racks are positioned vertically, side-by-side, in order that equipment can be installed and accessed easily.

The cabinet features large access doors on both long sides. These doors provide excellent working access to both ends of rack mounted equipment. The two equipment racks are positioned toward one end of the cabinet. The external face of this end features a gland plate with connectors to receive coaxial cable inputs, and waveguides for ingress of fibre optic cables. The air-conditioning plant that maintains the operating temperature within the cabinet is positioned at the other end. This end features a hinged door that allows the air-conditioning plant to be installed and removed. The exterior of this door features honeycomb RF gaskets that allow air to flow but prevent RF emission from escaping the cabinet. The air-conditioner will operate in outside ambient air temperatures up to 55 degrees Celsius.

The primary function of the RF shielded cabinet is to ensure that RF noise generated by the receivers and the various electronic sub-systems that support them doesn't escape into the environment. Radiated RF emission must be supressed to a level that is at least 20 dB below MILSTD461F (the adopted US Military Standard that describes how to test equipment for electromagnetic compatibility: Satav & Rama Sarma 2008) between 30 MHz and 3 GHz. This ensures that RF noise doesn't compromise the performance of the MWA, or affect other nearby instrumentation. Protection of the RF environment is mandated in the site license agreement that governs the MWA's activities at the site.

The cabinet is constructed from steel. The exterior and interior surfaces are powder coated, except for some masked areas that require electrical contact such as door seals and bulkhead connectors. Fully welded seams, RF gaskets integrated into all doors, and the use of waveguides for optical fibre cable entry ensure the cabinet provides sufficient levels of RF suppression. The use of (grounded) bulkhead connectors for coaxial cable inputs prevents conducted emissions from escaping on cables.



Figure 10. RF shielded cabinet. Left panel, cabinet open. Middle panel, cabinet closed. Right panel, cabinet installed inside environmental enclosure.



Figure 11. A rendering of the RF shielded cabinet inside the environmental enclosure.



Figure 12. Cabinet control unit.

Figure 10 shows the design of the shielded cabinet and Figure 11 shows a render of the RF shielded cabinet inside the environmental enclosure.

4.4. Supporting subsystems

A range of subsystems are required to support the primary elements of the new receivers, including devices to manage, control, and monitor the receivers, as well as devices to integrate the receivers into the wider MWA system architecture. In this section, we briefly describe these supporting subsystems.

4.4.1 Cabinet control unit

Internal environmental monitoring is facilitated by a custom designed and developed Cabinet Control Unit (CCU) (Figure 12). The CCU is a 2RU rack mounted network device based on a Raspberry Pi (RPi5) microprocessor. The RPi5 is supported by a small printed circuit board (PCB) mounted uninterruptible power supply (UPS) capable of supplying at least ten minutes of power. The small UPS enables the RPi5 to safely shutdown and protect software and firmware in the event of sudden power loss.

Environmental data are collected from an array of sensors and from other instruments within the cabinet. The data are collated and presented to telescope operators via a Grafana based GUI.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

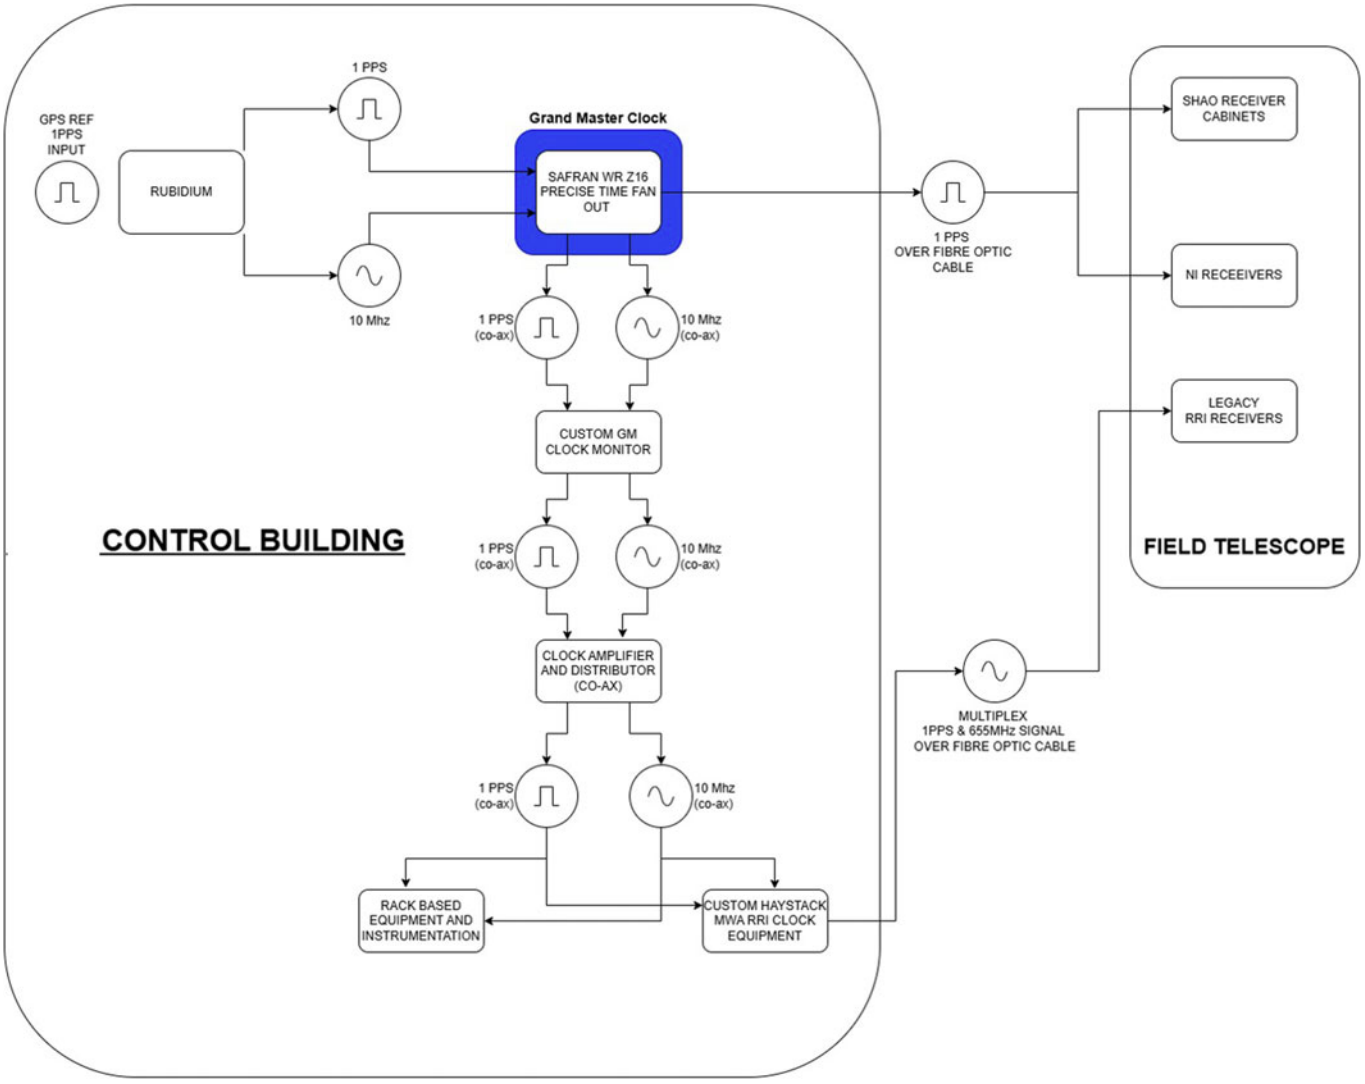


Figure 13. Phase III MWA clock and timing distribution architecture.

Internal cabinet cooling is provided by a robust industrial standard electronics cabinet cooling system. Significant thermal modelling was undertaken to optimise the overall physical size of the cabinet and the interaction with the dynamics of cooling performance. The optimisation focussed on balancing the internal volume of the cabinet with the capacity of the internal air mass to absorb thermal heat and moderate the cycling of the cooling system in order to minimise power consumption.

4.4.2 Timing module for clock distribution

Each SHAO receiver requires 1PPS and 10 MHz clock signals. The MWA Phase I and Phase II clock architecture reticulated a 1PPS and a bespoke Phase I receiver clock frequency of 163.84 MHz provided by custom timing equipment supplied by the Haystack Observatory. The addition of the SHAO receiver required an upgrade to the MWA clock architecture to facilitate three clock signals to be communicated across the MWA site.

Lessons learned from the previous years of operations identified that effective clock and 1PPS monitoring was a missing component in the management of the MWA. Two important factors constrained the development of the proposed Phase III clock system.

The first was the distance between the clock and timing source and the telescope. The clock and timing source were generated within the CSIRO Control building, located 5 km from the MWA core. This significant separation distance necessitated fibre optic clock and timing distribution. The second constraint was that the clock and timing architecture must not only support the telescope but also support other rack mounted equipment and instrumentation within the control building via coaxial cable only. These two constraints, combined with a preference for a COTS based solution, shaped equipment purchases and the development of the Phase III clock and timing architecture.

The built solution established a Grand Master Clock (GMC) at the apex of the clock and timing signal architecture (Figure 13). The GMC acts as the primary clock with all other clocks being subordinate. The GMC selected was a SAFRAN Z16 time fan out that provided a mix of coaxial cable and fibre optic clock and timing outputs based on the White Rabbit technology.[d]

[d]https://white-rabbit.web.cern.ch/.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press



Figure 14. Field Cabinet Clock Module showing the WR_LEN input 1PPS supply fibre optic port and the banks of 10 MHz and 1PPS outputs.

A custom designed clock monitor unit is installed on the coaxial cable outputs of the GMC prior to distribution to rack equipment and prior to the supply of a clock and timing signal to the custom made Haystack timing instruments.

The fibre optic distribution capability of the GMC allows the transmission of the 1PPS signal to each SHAO enclosure. The 1PPS signal is received and managed within each SHAO enclosure by the Field Cabinet Clock Distribution module. This unit is very similar to the GMC monitor and is depicted in Figure 14. The Field Cabinet Clock Distribution module is a 2RU rack mounted network device also supported by a RPi5 microprocessor with a small UPS. Similar to the CCU, the RPi5 enables telescope operators to monitor the clock and timing pulse via a Grafana based GUI.

The fibre optic 1PPS pulse is received within the Field Cabinet Clock Distribution module by a subordinate SAFRAN solutions WR-LEN module. 1PPS and 10 MHz signals are distributed to two banks of three outputs for each signal. Within the Timing Module, an output signal from each of the 1PPS and 10 MHz outputs is monitored by the RPi5. 1PPS and 10 MHz clock and timing signals are distributed to each SHAO receiver unit using the remaining outputs.

4.4.3 Power distribution and filtering module

The site is located in the Australian Radio Quiet Zone Western Australia.[e] Within this Zone, all installed equipment must be designed and engineered to minimise electromagnetic emissions to protect radio astronomy observations. Consequently, the power supply and internal power distribution must be designed and developed to minimise the possibility of any conducted emissions entering the cabinet on the mains supply, but also prevent any conducted emissions leaving the cabinet on the mains supply. This requirement for filtering of mains power on the supply and return paths also requires internal electromagnetic shielding of all filtering components to prevent the impact of electromagnetic interference on the SHAO receivers.

The power distribution and shielding module is located very close to the mains supply (230 VAC) connection point. The module consists of a single mains supply, an inline mains noise filter, and two filtered IEC mains outputs. The module is a completely enclosed metal box with a RF gasket around the lid access point (Figure 15).

Two managed, networked AC power rails distribute power to all internal cabinet equipment and also enable monitoring of power consumption of each device and remote power cycling of sub-systems.

[e]https://www.industry.gov.au/science-technology-and-innovation/space-and-astronomy/ska-project-australia/australian-radio-quiet-zone-wa.

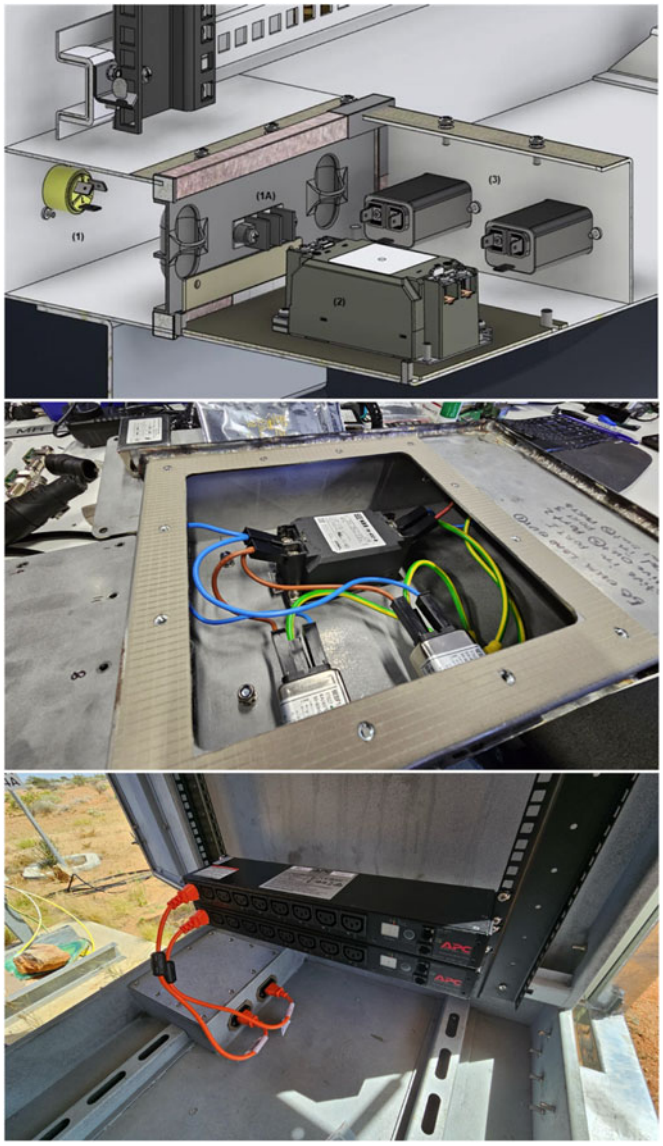


Figure 15. Upper image: Cross sectional view of Power distribution and shielding module. Mains power is supplied through socket at (1) and across to a rail at (1A), inline noise filter at (2), and two IEC outlets at (3). Middle image: Populated power filter and distribution module with RF gasket fitted. Lower image: Enclosed Power filtering and distribution module with managed, networked AC power rails attached.

4.4.4 Beamformer controller and power supply

Sitting at the nexus between the SHAO receiver and the MWA beamformer is the MWA Beamformer Controller (BFC) and Power Supply Module (PSU). Within the Phase I receivers, the MWA beamformer is powered and controlled from sub-systems that formed an integral part of the receiver. However, the SHAO receiver does not provide any form of power supply or control communications to the MWA Beamformer. Consequently, a new device was required that could provide DC power to the beamformer, communicate analogue M&C commands to analogue beamformers connected to SHAO receivers, and receive the analogue RF signals from the beamformer.

Lessons learnt from previous years of telescope operations identified that the legacy DC power, M&C, and RF signal chain were very susceptible to failure during periods of lightning activity.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

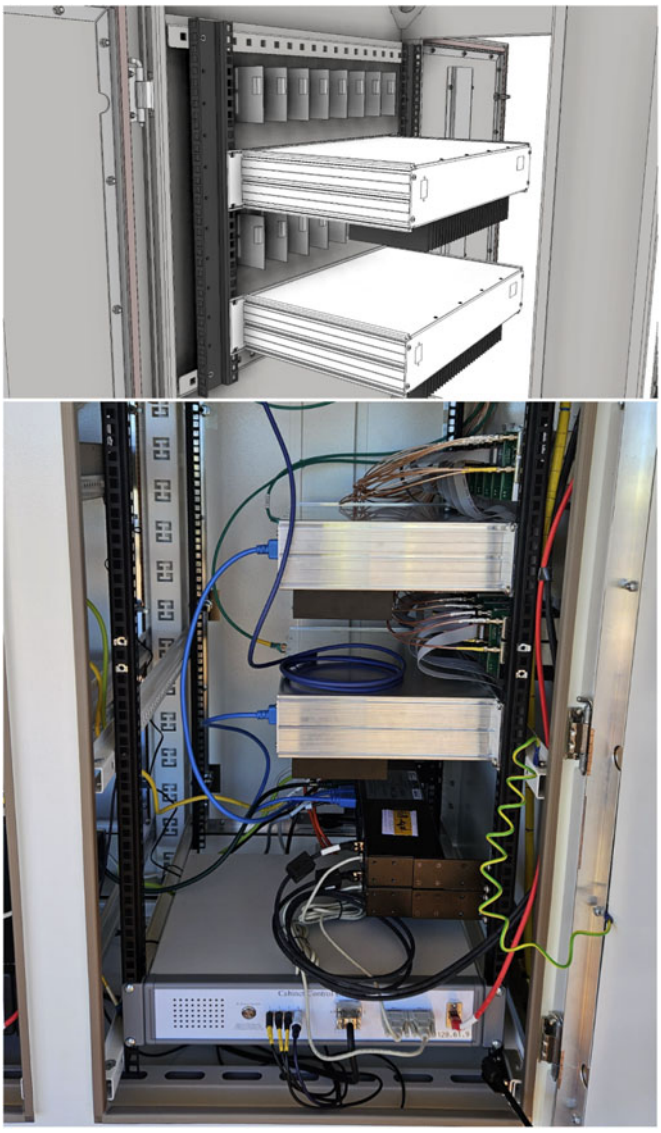


Figure 16. Top image: Design drawing of built rack mount solution. Bottom image: (top portion of image) two BFC PSU (silver boxes) connected with two banks of eight DoC PCBs via ribbon cable and coaxial cables. The bottom of the picture shows the CCU. This image demonstrates the constrained environment for airflow movement.

With the need identified for a new source of DC power, M&C communications, and RF signal management for the SHAO receiver, a design requirement was introduced to harden the 'receiver – beamformer' nexus against lightning. The intent of this design requirement was two fold – improve the resilience of the MWA in adverse weather conditions and control maintenance and repair costs.

The design development phase faced three significant constraints. First, the BFC would have to be installed in the side of the cabinet furthest from the cooling system. This had two implications – the BFC would be on the north facing end of the cabinet (highest solar load) and would have to operate in temperatures and airflow conditions significantly different to equipment installed closer to the cooling system. Second, the built solution had to re-use the existing coaxial cables between beamformers and receiver. The cost to replace this cable could not be supported by the SHAO receiver project budget. Finally, the design effort occurred during the COVID-19 period. The significant disruption to global manufacturing and logistics that occurred during that period made sourcing the microprocessors preferred for the BFC impossible. This forced the adoption of an alternate solution. The design team adapted a BFC implementation developed for the Engineering Development Array (EDA; Wayth et al. 2017), a retired SKA-Low benchmarking array.

The deployed solution is a custom designed rack mounted 2RU network-managed power supply unit supporting eight 'Data over Co-ax' (DoC) PCBs. Each DoC PCB is connected to an MWA beamformer via coaxial cables. The BFC PSU provides M&C communications and DC power to the DoC PCB via a ribbon cable, whilst RF analog signals are communicated directly to the SHAO receiver via coaxial cable with push on/ pull off (DIN 1.0/2.3) connectors.

The BFC solution incorporates several key features and improvements. First, the original EDA Cosel multiple output switch mode PSU was replaced with a new AC/DC switch-mode design based around a high-efficiency PCB mounted converter module and downstream DC–DC conversion. The high temperature tolerance and low thermal dissipation enabled the use of a large passive finned heatsink without requiring any changes to internal airflow within the cabinet. Second, communications between the wider telescope M&C network and the beamformer through the DoC PCB was achieved through the integration of a microcontroller based PCB within the BFC, assisting with reliability. The microcontrollers have no operating system or filesystem that can become corrupted – instead, the firmware is run directly from flash memory inside the microcontroller, This is read-only except for the rare occasions when a new version is pushed to the device. Third, a re-design of the DoC PCB reduced susceptibility to damage by lightning. Since their deployment in July 2025, signal paths featuring the hardened DoC PCB have endured several significant lightning events without sustaining any failures. Finally, the modularity of the BFC system means turnaround time in the event of failure is significantly reduced. Push on ribbon cable connectors, push on coaxial cable connectors and a single stand-alone box for the BFC means that, in the event of a complete PSU failure, a new unit can be installed very quickly.

All components described above can be seen in the as-built form in Figure 16.

5. MWAX

The MWAX correlator (Morrison et al. 2023) was commissioned in early 2022 and operated successfully up until the end of Phase II. During development and testing, simulated MWA Phase III observations were processed by the correlator, ensuring the system would be capable of correlating up to 256 tiles on the MWAX hardware in real-time. As each new receiver was brought online during Phase III deployment, the number of tiles MWAX had to correlate increased by eight. At each increment, the system was tested to confirm that the software, hardware, and network could handle the extra load.

The MWAX system has been incrementally improved during Phase III development. The following features and enhancements are now deployed in the system.

5.1. Cable and geometric delay correction and fringe stopping

The cable and geometric delay correction and fringe stopping feature, as described in Morrison et al. (2023), are now fully

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

implemented as described in that work. Observations can be scheduled to apply all, some, or none of these corrections. The main benefit of this feature is the ability to use more aggressive averaging of the output visibilities, significantly reducing the data volume, and fringe stopping ensures minimal decorrelation due to the Earth's rotation.

5.2. Oversampled coarse channels

The passband response of the original 16 receivers exhibit steep attenuation at the edges of the coarse channels, resulting in approximately 10% of the channel being unusable (Morrison et al. 2023), due to the critically sampled receiver.

The SHAO and NI receivers both support an oversampling mode. In oversampling mode, the MWAX UDP captures can now handle eight seconds of data from a single 1.6384 MHz coarse channel. Figure 17 provides a comparison of the critically sampled receiver passband and the oversampled receiver passband, demonstrating that the oversampled receivers maintain a band free from aliasing across the full 1.28 MHz coarse channel bandwidth, removing the band edge artifacts experienced in data from the critically sampled receivers.

The top panel of Figure 17 shows the passbands for the critically sampled (blue) and oversampled (yellow) cases in comparison. The unshaded area represents the frequency range of a single coarse band. The response of the critically sampled case shows a band edge drop off within the coarse channel, whereas the oversampled case does not. The middle and lower panels show the passbands and aliased responses for the critically sampled and oversampled cases, respectively. These panels show that the aliased response is of a low level inside the coarse channel for the oversampled case, and high and within the band for the critically sampled case. The high in-band aliasing for the critically sampled case causes artifacts at the coarse band edges that are problematic for the EoR experiment.

5.3. Automated switching from critically sampled to oversampled coarse channels

Switching between critically sampled and oversampling modes entails three main tasks, all triggered when a special previously-scheduled 'correlator mode change' observation occurs. We describe the tasks below, in the case of the SHAO receivers. The tasks are similar for the NI receivers:

1. Setting up the SHAO receiver firmware to deliver oversampled packets. This is done in the SHAO receiver handler process, running on the M&C server. The handler process for each SHAO receiver detects each new observation in the schedule as it starts, and passes the correct observing parameters to the receiver FPGAs via the firmare API. When a correlator mode change observation starts, the handler process uses the SHAO firmware API to boot into the appropriate mode. This firmware swap takes a few tens of seconds.
2. Reconfiguring the telescope to include only SHAO receivers and tiles connected to them (for oversampling), or all receivers (for critical sampling). This is done by the 'MWAX oversampling switcher' process running on the M&C server. When a correlator mode change observation starts, each receiver is flagged as 'active' or 'inactive' in the database as required for the new mode. This reconfiguration takes a few tens of seconds, including new observations dynamically scheduled to set up the new configuration.
3. Stopping, reconfiguring, and restarting the MWAX correlator, using subfiles (temporary packet storage areas in shared memory) of the appropriate size for either 128 oversampled tiles, or 256 critically sampled tiles. This process is by far the most time consuming, and takes 10–15 min to complete, during which time no observations can be made. This process is also managed by the 'MWAX oversampling switcher' when correlator mode change observations occur.

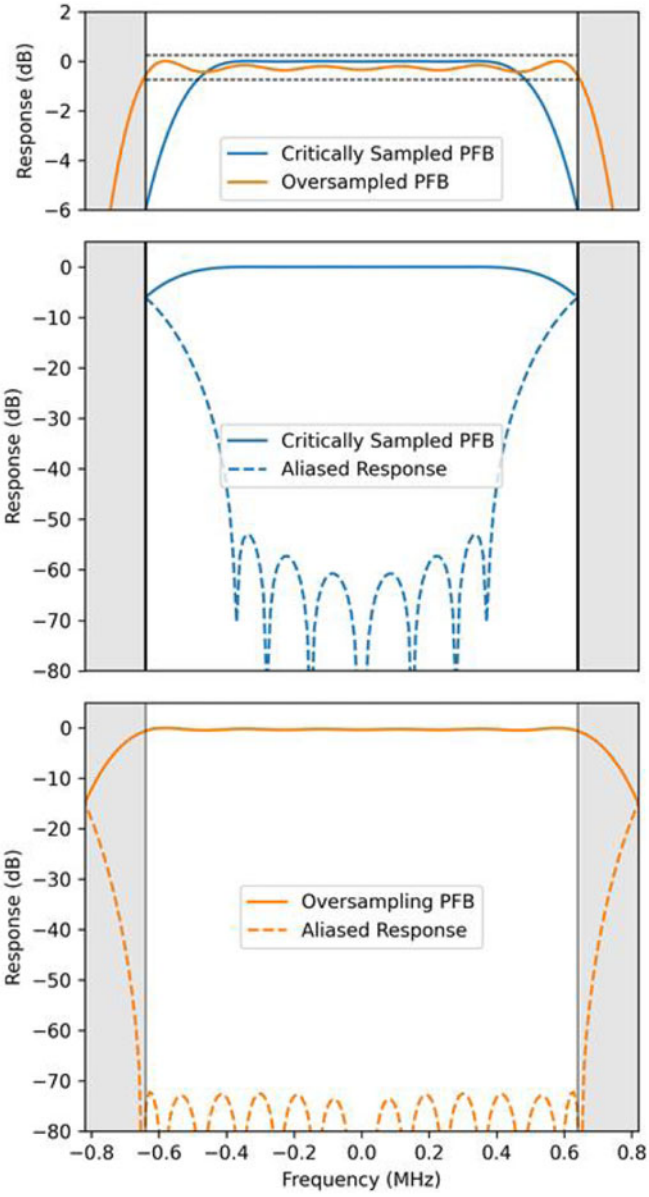


Figure 17. Comparison of the passband response for the critically sampled (Phase I/II) receivers and the Phase III (SHAO and NI) receivers, showing the recovery of the band edges in the oversampled receiver case. Details in the text.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

5.4 Near-real-time calibration

In an extension of the work of Sokolowski et al. (2020), a new, seven-server GPU calibration cluster and pipeline (called 'Calvin') was commissioned on the site in June 2025 to perform both near-real-time calibration of observations as they are produced by the correlator, and on-demand calibration of historical observations via requests from users through the MWA All-Sky Virtual Observatory (MWA ASVO) data portal. Calibration solutions are generated by the Hyperdrive[f] calibration software and are applied when observations are delivered through ASVO via the Birli[g] software package (Jordan et al. 2025). The calibration solutions are a key diagnostic tool for the MWA Operations team, to monitor the array's health, and they will also be used in the currently in-development MWAX real-time coherent beamformer.

via requests from users through [ASVO]. Calibration solutions are generated by the Hyperdrive calibration software and are applied when observations are delivered through ASVO via the Birli software package.

5.5 Improved archiving data integrity

The ten 'mwacache' archive servers situated in one of Curtin University's data centres were replaced in June 2025 with ten new servers and an updated archiving mechanism. The original mwacache servers transferred data to the Pawsey Supercomputing Centre (where the MWA's archive is located) long-term storage (LTS) system using the Boto3 Python library and, although this worked well, the data transfer process did not explicitly re-read the data and verify the checksum at the Pawsey LTS destination to ensure each file was transferred successfully, leading to some very rare data corruption issues. After extensive testing and research, Rclone[h] was chosen as the replacement transfer tool for ingest into Pawsey's LTS. Rclone's copy command is used to transfer the data, then upon detection of a successful exit code, 'rclone check' is called to verify the checksum of the file at the destination. If this check or the copy fails, the file is re-queued and attempted again later.

Each of the eight active mwacache servers transfer three data files each in parallel to the Pawsey LTS at an average rate of 0.5 Gbps per file (1.5 Gbps per server/12 Gbps in aggregate across all 24 streams). A typical 100 GB MWA correlator observation takes just over one minute to transfer to the archive, and once at Pawsey, it is available for download via the MWA ASVO data portal a few minutes later.

5.6 Implications for the MWA data archive

As of writing, the MWA data archive at the Pawsey Supercomputing Centre has 56.9 PB of capacity, which is a mix of Ceph object storage and magnetic tape storage. The archive consumes 49 PB and has 7.9 PB available storage. The data in the MWA archive consist of over 600 000 observations, involving 24 000 h of observations and producing over 37 million files since operations commenced in mid-2013.

The Phase III upgrade increases the maximum number of tiles in an observation from 128 up to 256. For MWA voltage capture system (VCS) observations this increases the data volume per unit time by a factor of two. For correlated observations, this increases the data volume by just under a factor of four as the number of visibilities approximately quadruples with the doubling of the number of tiles.

In anticipation of increased Phase III data volumes and the lack of prospects for increasing the storage allocated to the MWA, some mitigation strategies were put in place to ensure the archive could be managed in a sustainable manner.

First, the MWA Data Retention Policy[i] was adopted. The aim of the policy is to manage the availability of archive storage so that new observations can be made, while deleting the oldest, least useful data. The MWA Data Manager and Principal Scientists nominate a list of observations which they assess to be old and/or of low value and start a four-week consultation period where MWA members can review the deletion list and request them to be retained if they believe they still hold scientific value. The policy has been activated twice since it was adopted by the MWA Board in 2019 and ensures a balance is maintained between new observations and preservation of as much scientifically useful data as possible.

Second, fringe-stopping (see Section 5.1) is an effective method to reduce the volume of archived visibilities, by allowing more aggressive frequency and time averaging of correlated visibilities. Without fringe-stopping, a standard MWA correlator observation at 256 tiles would not be able to have greater than 10 kHz fine channel width or 0.5 s integration time, due to visibility decorrelation caused by the Earth's rotation, and would require 6.46 GB per second of storage. However, when fringe-stopped, a more modest 40 kHz, 2 s correlator mode could be used which would not suffer materially from decoherence and result in only 0.4 GB per second of data generated. Thus, in this example, fringe-stopping produces data volumes sixteen times smaller than without fringe-stopping, which more than mitigates the nominal quadrupling of visibility data due to the increase from 128 to 256 tiles.

Finally, another major mitigation for the growth in the MWA archive is real-time voltage beamforming. As of writing, an MWAX real-time beamformer mode is under development. The beamformer will reduce the need to capture and store tile-based voltages, which are produced at 114 TB per hour, for offline voltage beamforming. Instead, the MWAX system will be able to form several coherent and incoherent beams in real-time, generating dual-polarisation, complex voltage beams in VDIF format[j] and Stokes I Filterbank format files for incoherent beams (Lorimer 2011). The resulting beamformed data products consume much less space than VCS or even correlated visibilities (approximately 1/n per beam compared with VCS observations, where n is the number of tiles). In addition to reducing archive growth, a real-time beamforming capability will enable the MWA to participate in VLBI experiments and contribute to the fast transient community more effectively through rapid response follow-up observations of pulsars and similar objects.

6. RFoF conversion

Power distribution around the MWA site is a fixed constraint on the array layout. The power available at different access points around the site determines where powered electronics can be

[f] https://github.com/MWATelescope/mwa_hyperdrive.
[g] https://github.com/MWATelescope/Birli.
[h] https://rclone.org/.
[i] https://www.mwatelescope.org/wp-content/uploads/2025/12/MWA-DATA-ACCESS-POLICY-v1.3.pdf.
[j] https://vlbi.org/wp-content/uploads/2019/03/VDIF_specification_Release_1.1.1.pdf.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

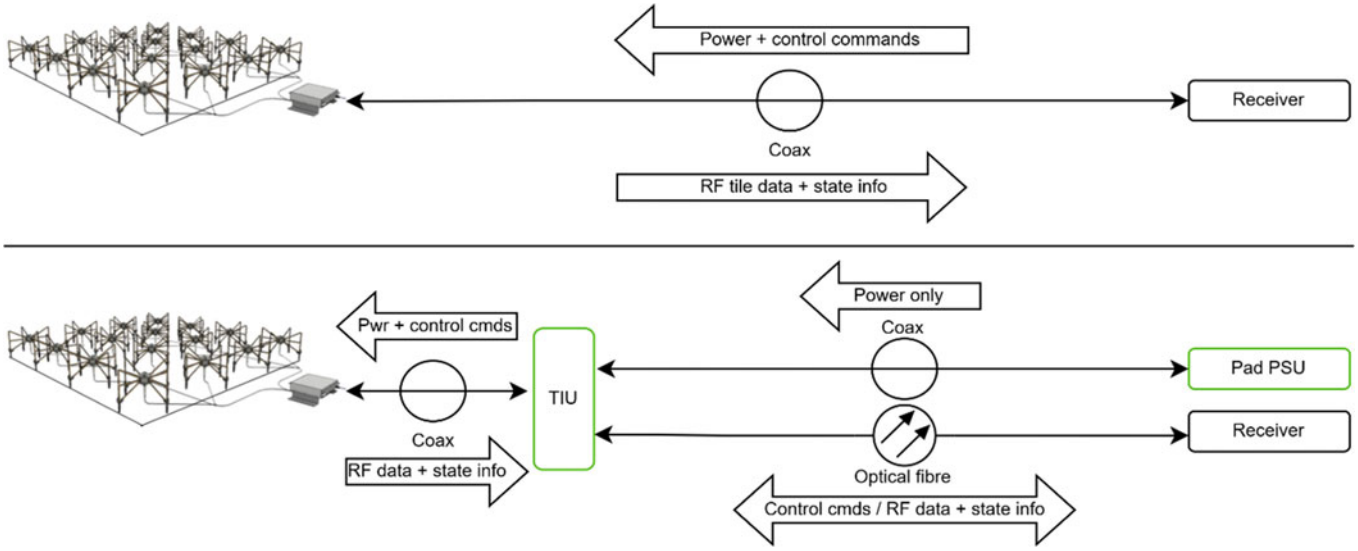


Figure 18. Top panel: Coax tile configuration. Bottom panel: Phase III converted RFoF tile configuration. Phase III additions are shown in green.

positioned. In Phase I, receiver enclosures were positioned at 230 V power access points. Receivers supplied 48V power to all Phase I tiles via coaxial cable. Resistive cable losses determined the maximum viable distance between receiver and tile of approximately 500 m (Tingay et al. 2013). The development of a small, tile-level, photovoltaic and battery system enabled the deployment of 56 long baseline tiles as part of the Phase II upgrade (Wayth et al. 2018). For these tiles, power was generated at the tile location and the tile output signal transmitted on surface-laid optical fibres (RF over Fibre: RFoF).

In Phase III 24 additional tiles have been converted from coaxial cable to optical fibre output so that receivers can be located where there is power available to support them. This is achieved by the introduction of a TIU (Tile Interface Unit) and Pad PSU (Power Supply Unit). Figure 18 shows schematic views of a coax configured tile and an RFoF converted tile.

6.1. Tile Interface Unit (TIU)

Every MWA tile includes an analogue beamformer that points the tile beam on the sky by applying selective delays to the RF signal received by each antenna and combines the 32 signals (one signal per polarisation from 16 dipole antennas) into two RF outputs—one for each polarisation Tingay et al. (2013). The TIU (Figure 19) interfaces to the analogue beamformer. It takes the two RF outputs from the beamformer as inputs via two short coaxial cables. The same cables deliver 48 VDC and pointing commands (delay settings) from the TIU to the beamformer (Figure 18, bottom). Electrical-to-optical converters inside the TIU transition the RF signal onto a pair of single mode optical fibres. The optical modules operate at 1310 nm. The system does not utilise any form of optical multiplexing. The TIU receives pointing commands, generated by the telescope monitor and control system, via a receiver, and 48 VDC from a Pad PSU.

6.2. Pad Power Supply Unit (Pad PSU)

Pad PSUs (Figure 20) are collocated with receivers at 230 V access points. They convert 230 V site power to 48 VDC for delivery

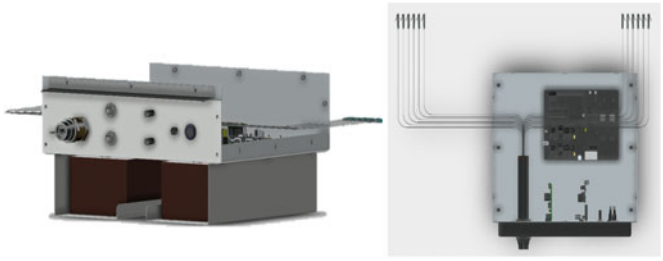


Figure 19. Design drawing of the Tile Interface Unit.

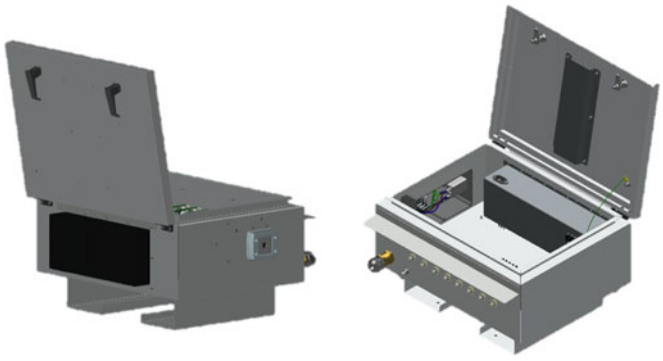


Figure 20. Design drawing of the Pad Power Supply Unit.

to the TIU. Transmission to the TIU is via the coaxial cable that previously connected the antenna tile to a receiver (Figure 18, top).

7. Operational benefits

The Phase III upgrade confers a number of operational benefits, along with the capability and performance improvements that motivated the upgrade. In Phase II the MWA was manually reconfigured to switch between compact and extended arrays to optimise for different observational performance characteristics. This was necessary because the Phase I receiver fleet, which was still in use in Phase II, was only able to support 128 tiles.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

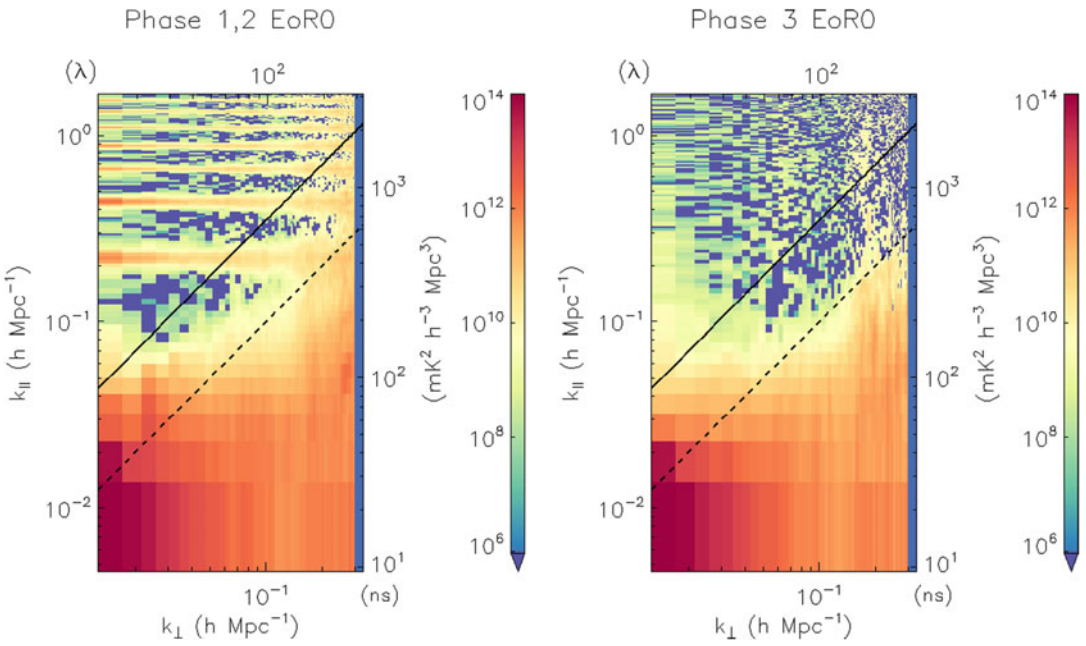


Figure 21. Two dimensional EoR power spectra, from Phase I/II data (left) and preliminary Phase III data (right), showing the improvements to the $k_\parallel$ modes due to removal of aliasing artifacts from coarse channel edges. In both cases, one hour of data is used, from the east-west oriented polarisation visibilities obtained from observations of the EOR0 field (RA = 0 h, Dec = −27 deg.) in the 167–197 MHz frequency range. Further details in the text.

To reconfigure the array, receivers were disconnected and physically relocated. This process was time consuming, costly, and involved significant risk of damage to delicate equipment. In Phase III all 256 tiles are connected to a receiver and can be included in the operational array. The observing array can be reconfigured in software, eliminating the need to manually reconfigure. Not having to manually reconfigure the array also results in a decrease in the number of days per year when the MWA site is attended. Reduced site attendance translates directly to more time on sky, and reduced operational expense.

Over the course of its operational life the MWA has suffered significant lightning damage. The MWA's primary vulnerability to lightning results from the long, surface laid, coaxial cables that carry both power and RF signals between many of the tiles and their receivers. Any lightning in the area can readily induce currents into these cables, which span between 90 and 524 m. The resulting high-voltage surge propagates to the dual purpose (power and RF) PCBs on either end of the cable and can destroy the sensitive RF components. In Phase III, the tiles converted from coaxial to optical fibre (for RF and control signals) become more resilient to lightning damage. Each tile still has a long, surface laid, coaxial cable but it no longer carries the RF signal. The electronics on the end of the power-only link are more robust against the surges generated by lightning events. Funding permitting, more tiles will be converted to optical fibre (for RF and control signals) in the future, making the telescope more resilient against one of its primary failure modes. As noted in Section 6, 80 tiles are currently supported by RFoF, leaving up to 176 available for RFoF upgrade in the future.

8. Science implications

The upgraded Phase III MWA offers a range of fundamental and incremental improvements for science. An example of a fundamental improvement is the new fleet of oversampled receivers that have been attached to the set of core MWA tiles, supporting the short baselines of the array and hence of primary relevance to the EoR experiment. The improvement lies in the removal of the coarse channel edge artifacts that have been present for the EoR experiment during Phases I and II of the MWA, that significantly disrupt efforts to detect the EoR signal. EoR observations performed during Phase III commissioning and early operations have been subject to preliminary processing, and can be shown here to illustrate the nature of this improvement. Figure 21 shows two-dimensional power spectra from the Phase I/II MWA, alongside the same from Phase III observations. In both cases, one hour of data is used, from the east-west oriented polarisation visibilities obtained from observations of the EOR0 field (RA = 0 h, Dec = −27 deg.) in the 167–197 MHz frequency range. The data were identically processed with calibration and peeling using Hyperdrive (Jordan et al. 2025) and the power spectra estimated with CHIPS (Trott et al. 2016), following, for example, Nunhokee et al. (2025). As can be seen in the Figure, the so-called EoR window is filled with a uniform noise-like signal for the Phase III observations, but suffers from periodic artifacts in the frequency dimension (line-of-sight) in the Phase I/II observations due to the coarse channel edge effects caused by the critically sampled receivers used in Phase I and II.

The cleaner EoR window should allow for deeper integrations, and a closer examination of the noise properties in the EoR window, that will hopefully allow upper limits that reach closer to the expectations of theoretical EoR models.

Broadly applicable to most of the science undertaken with the MWA, the ability to utilise all 256 tiles simultaneously enables a factor of two increase in raw sensitivity. Perhaps more important than the potential of this additional available sensitivity is the flexibility allowed by having all 256 tiles simultaneously available. As noted in the early sections of this paper, users now have the

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

ability to switch between different sub-sets of tiles, either to minimise the data volume generated for their science requirements, or to optimise an array for those requirements. The configurations that allow both good angular resolution and surface brightness sensitivity are beneficial, for example for galaxy cluster studies, supernova remnants, and other objects that contain complex structure on different angualar scales.

One of the highest impact areas of science with the MWA has been the emergence of Long Period Transients (Hurley-Walker et al. 2022, 2023) as an exciting (and unexpected) new class of radio transient. Effective searches for transients are best done with the MWA's long baseline configurations, which in Phase II were unavailable for large fractions of the year when the MWA was dedicated to EoR observations. The ability to perform software-defined reconfigurations between short and long baseline MWA configurations now allows the ability to make observations for transients at short notice, when required.

While, early in Phase III operations, the reconfiguration of the MWA will be limited to a small (five) set of pre-defined configurations, and conservatively scheduled in optimised blocks, as the MWA Operations team gains experience in reconfiguration and continues to build incremental improvements to the MWA's software systems, it will be possible over time to shift toward a more dynamic and user-defined exploitation of this flexibility.

The enhanced EoR capability and greatly improved flexibility are both fundamentally enabled by the MWAX correlator, using modern GPU accelerated computing. As well as supporting the correlation of all 256 tiles simultaneously, and the reconfigurable flexibility described above, MWAX also has the capacity to perform other non-correlation signal processing tasks. Primarily important among these will be the ability to generate real-time voltage beams using up to the full 256 tiles. During Phases I and II, the only way to form voltage beams for a range of science applications (e.g. pulsar studies, gamma-ray burst observations, fast radio burst studies etc) has been to record all voltages from all tiles and perform offline beamforming using high performance computing. Due to the very large data volumes involved and the high cost of the computing task, exploitation of voltage beams has been limited and inefficient.

Having the ability to form real-time voltage beams using MWAX will massively reduce the time and effort required to undertake relevant science, and greatly relieve users of data handling and processing burdens. The development of a real-time beamformer using MWAX (based on the approach described in Ord et al. 2019) is in advanced stages at the time of writing and is expected to be a capability released to users in approximately mid-2026.

9. Discussion and conclusions

The evolution of the capabilities of the MWA into Phase III represents a continuation of both episodic step change in capability and gradual evolution/improvement. Each phase of upgrade has been informed by both operational experience and the direction of travel of science requirements and opportunities, as these have inevitably also evolved over time. The result over the course of 13 years since the commencement of operations has been a facility that has remained relevant and responsive to the needs of users, for a very low total cost of ownership relative to the scientific output and impact.

As likely the final major MWA upgrade, Phase III also represents the culmination of the facility and close to the fullest possible realisation of the potential of the base technologies and approaches. As noted in the introduction, we anticipate that the MWA will operate in its Phase III format until mid-2030, albeit likely with minor improvements and adjustments as noted in this paper to fully exploit the Phase III potential. Thus, the MWA will have its peak performance period for the next approximately five years.

As a Precursor for the SKA-Low facility, by 2030 the MWA will have completed its mission (and much of the MWA hardware, including as part of the Phase III upgrade, will be at end of life), and the SKA-Low capabilities will outstrip those of the MWA. This makes 2030 a natural point of transition from the MWA to the SKA-Low, when it comes to Southern Hemipshere low frequency radio astronomy. It is worth examining some dimensions of this transition between now and 2030, in order to illuminate the points at which particular capabilities partially or fully transition, but also to point out some interesting scientific synergies between the Phase III MWA and the emerging SKA-Low. Such a summary may assist low frequency users of the MWA and SKA-Low in their medium term planning.

9.1 Evolution from MWA Phase III to SKA-Low Phase I

Timelines to SKA-Low science have been recently released, with details available via the SKAO website,[k] and linked pages. In summary, SKA-Low will enter into Science Verification phases based on the staged rollout of the SKA-Low instrument. According to the SKAO website, Science Verification is described thus:

Science Verification will be the first opportunity for the community to access SKA data. Six months before Science Verification begins, the Observatory will issue a Call for Science Verification ideas. All Science Verification data will be made publicly available and community participation during this phase will be essential for validating Observatory processes and shaping effective workflows ahead of regular operations. The SKA Regional Centres Network (SRCNet) Project delivery timeline is aiming to be ready for first users at the start of Science Verification, with testing at appropriate scales starting in the first half of 2026.

Following Science Verification, the first opportunity for PI-led proposals and science team ownership of SKA-Low data will commence with a shared-risk Cycle 0 period in 2030, leading eventually to the commencement of Key Science Projects (KSPs) in 2033.

Anticipated in 2027 will be the so-called AA2 (Array Assembly 2) stage of SKA-Low, which will comprise of 68 SKA-Low stations distributed over an area supporting a maximum baseline of 65 km. Details are provided in SKAO memo SKAO-TEL-0002299.[l] The so-called AA* stage of SKA-Low will represent the instrument that enters Cycle 0 and proceeds into the KSP era, and is due to be released for Science Verification in 2029. As such, in the remainder of this section, we focus specifically on AA2, as the Array Assembly that will primarily overlap with the remainder of the MWA's operational lifetime. At most, AA* will overlap with the MWA by ∼12 months, thus likely not offering a major opportunity to consider transitioning or complementary capabilities between the two facilities.

[k] https://www.skao.int/en/science-users/timeline-science.
[l] https://zenodo.org/records/16951020.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

The 68 SKA-Low AA2 stations will be grouped into clusters that occupy 16 locations along the extended spiral arms of the SKA-Low configuration, and are thus referred to as 'remote' stations. AA2 will not involve any of the 'core' SKA-Low locations; the core will be populated subsequent to AA2, leading to AA*, which is currently envisaged to comprise of 257 stations.[m] As such, while the sensitivity of AA2 is relatively high, the AA2 (u, v) coverage is effective driven by only 16 geographically distinct locations, with very few short baselines. AA2 is therefore largely an array with high sensitivity to small angular scale structures, and very little sensitivity to large angular scale structures.

MWA antennas meet the SKA-Low sensitivity requirements over a significant fraction of the SKA-Low frequency range (Sokolowski et al. 2022). Thus, to zero order, one SKA-Low station is equivalent in sensitivity to 16 MWA tiles. The AA2 sensitivity (for the same bandwidth) is therefore approximately a factor of four higher than for the MWA. However, all of the MWA's baselines are short, relative to the long baselines realised by AA2. The MWA's filled (u, v) coverage and relatively short baseline distribution means that it has superb low surface brightness sensitivity. This suggests a level of complementarity between the Phase III MWA and SKA-Low AA2.

In order to explore this complementarity, we utilised the 'SKA SciOps Array Configuration and templates,'[n] a python package developed by the SKAO to provide tools to simulate (u, v) coverages, synthesised beams, and other characteristics. We used these tools to generate a zenith pointed snapshot (u, v) coverage at a single frequency of 150 MHz for AA2, and show it in Figure 22, along with the Phase III MWA (u, v) coverage from Figure 5. The long baseline (aside from the grouping of very short baselines due to intra-cluster station baselines), but relatively sparse, nature of the AA2 coverage can be seen, along with the short baseline but filled coverage of the Phase III MWA. Importantly there exist a reasonable number of overlapping (u, v) points between the two coverages. For both facilities, (u, v) coverages improve via the use of bandwidth and Earth-rotation synthesis

This complementarity admits some science opportunities to utilise both arrays together, with the MWA representing the equivalent of 16 SKA-Low stations of sensitivity on short baselines, and AA2 representing 68 stations of sensitivity on predominantly long baselines, with very short AA2 baselines due to intra-cluster station baselines. Although the two arrays are not cross-correlated, a joint deconvolution of AA2 and MWA visibilities, similar to that achieved between Phase I and Phase II MWA datasets (Mantovanini et al. 2025), should be possible, to support a range of imaging applications.

The capabilities envisaged to be released under an AA2 Science Verification program will support such complementary imaging applications, as AA2 users will be able to access calibrated, averaged, and gridded visibilities in both continuum and spectral line modes. Aside from science applications, the ability to jointly deconvolve AA2 and Phase III MWA data should provide a powerful cross-check of performance between the two arrays.

Another area in which Phase III MWA and AA2 will maintain a level of synergy is for Pulsar Searching (PSS) and Timing (PST). As described above, the MWA Phase III upgrade will soon include a real-time voltage beamforming capability implemented using the MWAX correlator. Likewise, AA2 is envisaged to support a single voltage beam. However, as the AA2 configuration is approximately 10 times more distributed than the MWA configuration, the AA2 beam area will be 100 times smaller than for the MWA. The AA2 beams will, however, be four times more sensitive (for the same bandwidth). The voltage beamforming capabilities of the Phase III MWA and AA2 may be complementary in the sense of the MWA being used to cover a larger area of sky more quickly for pulsar searching, but with lower sensitivity than AA2, with AA2 being used for higher sensitivity follow-up observation of any targets of interest. Again, in terms of Science Verification, simultaneous observations of pulsars with both AA2 and Phase III MWA are likely to be able to provide powerful cross-checks of science performance.

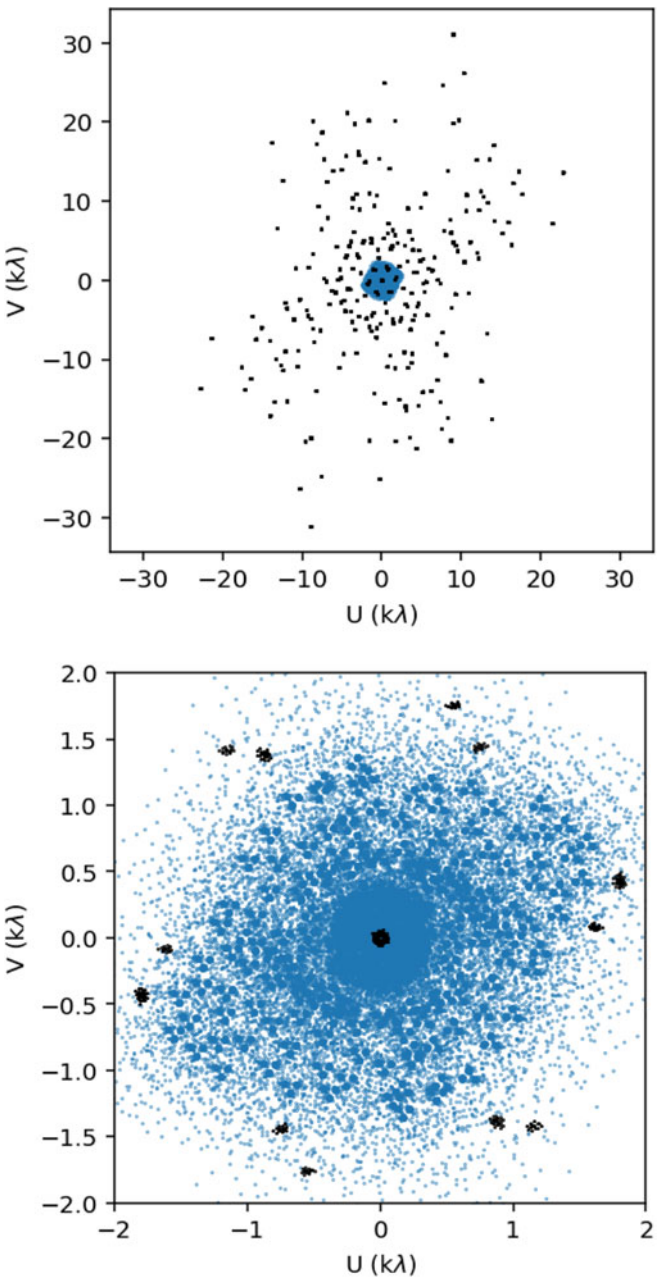


Figure 22. MWA (blue) and SKA-Low AA2 (black) (u, v) coverages for a single frequency at 150 MHz and for a zenith pointed snapshot. Top: full (u, v) coverages. Bottom: zoomed in (u, v) coverages.

[m] https://mailchi.mp/skao/skao-pulse-newsletter-9169539.
[n] https://gitlab.com/ska-telescope/ost/ska-ost-array-config.

https://doi.org/10.1017/pasa.2026.10218 Published online by Cambridge University Press

While some areas of very interesting synergy and complementarity exist between Phase III MWA and SKA-Low AA2, there will be some areas of science that will experience gaps in capability when the MWA ceases operations in 2030. One particular area is in so-called fast imaging, for example for the observation of Long Period Transients, with imaging timescales from seconds to minutes. SKA-Low is not projected to support fast imaging modes in Science Verification until Cycle 2 (2032). Fast imaging should be supported for AA2 due to the availability of visibility data, but the supply of visibilities is not supported subsequently, from AA* (2029). So, without the MWA fast imaging modes, and a gap between AA2 and Cycle 2, some areas of science may experience two to three years of pause.

Overall, however, the brief analysis and summary presented here shows that the Phase III MWA capabilities generally reasonably smoothly transition to those of AA* on a timescale of 2029/30. Moreover, it is clear that the capabilities of SKA-Low AA2 and Phase III MWA have a high level of complementarity in the period 2027–2029. Interestingly, this opens opportunities for combined science during the first SKA-Low Science Verification phase, and also offers opportunities to undertake detailed performance cross-checks between a well-understood and tested system in the MWA to the new system in AA2, which should be extremely helpful to quickly establish and verify SKA-Low capabilities.

Acknowledgements. We thank the anonymous referee for a close reading of the paper, and for offering comments that assisted in providing additional clarity for non-expert readers. This scientific work uses data obtained from Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory. We acknowledge the Wajarri Yamaji as the Traditional Owners and Native Title Holders of the observatory site. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited. We acknowledge the Pawsey Supercomputing Centre which is supported by the Western Australian and Australian Governments. Part of this work was supported by the Australian SKA Regional Centre (AusSRC), Australia's portion of the international SKA Regional Centre Network (SRCNet), funded by the Australian Government through the Department of Industry, Science, and Resources (DISR; grant SKARC000001). AusSRC is an equal collaboration between CSIRO – Australia's national science agency, Curtin University, the Pawsey Supercomputing Research Centre, and the University of Western Australia. We acknowledge funding from the National SKA Program of China, grant number SQ2020SKA0110100, which supported the SHAO receiver development.

References

Bandura, K., et al. 2014, in Ground-based and Airborne Telescopes V, Vol. 9145, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. L. M. Stepp, R. Gilmozzi, & H. J. Hall, 914522. doi: 10.1117/12.2054950. arXiv: 1406.2288 [astro-ph.IM]
Beardsley, A. P., et al. 2019, PASA, 36, e050
Benthem, P., et al. 2021, A&A, 655, A5
Berkhout, L. M., et al. 2024, 136, 045002. doi:10.1088/1538-3873/ad3122. arXiv: 2401.04304 [astro-ph.IM]
Bowman, J. D., et al. PASA, 30, e031
Chen, Z. P., et al. 2016, in Ground-based and Airborne Telescopes VI, Vol. 9906, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. H. J. Hall, R. Gilmozzi, & H. K. Marshall, 99065W. doi:10.1117/12.2232570
Dewdney, P. E., et al. 2009, Proc. IEEE, 97, 1482
Furlanetto, S. R., Peng Oh, S., & Briggs, F. H. 2006, PhR, 433, 181
Gupta, Y., et al. 2017, CurSci, 113, 707
Hall, P. J. 2013, et al. in International Conference on Electromagnetics in Advanced Applications (ICEAA, 1, 1. doi:10.1109/ICEAA.2013.6632250
Hurley-Walker, N., et al. 2023, Natur, 619, 487
Hurley-Walker, N., et al. 2022, Natur, 601, 526
Johnston, S., et al. 2008, ExA, 22, 151
Jonas, J., & MeerKAT Team. 2016, in MeerKAT Science: On the Pathway to the SKA, 1. doi:10.22323/1.277.0001
Jordan, C. H., et al. 2025, in *URSI Asia-Pacific Radio Science Meeting (AP-RASC)*, 1–4. doi:10.46620/URSIAPRASC25/LSCN1310
Kassim, N., et al. 2010, in Advanced Maui Optical and Space Surveillance Technologies Conference, ed. S. Ryan, E59
Kolopanis, M., et al. 2023, MNRAS, 521, 5120
Lonsdale, C. J., et al. 2009, IEEE Proc., 97, 1497
Lorimer, D. R. 2011, Astrophysics Source Code Library, record ascl:1107.016
Macario, G., et al. 2022, JATIS, 8, 011014
Mantovanini, S., et al. 2025, 42, e137. doi:10.1017/pasa.2025.10094
Morrison, I. S., et al. 2023, PASA, 40, e019
Nunhokee, C. D., et al. 2025, 989, 57. doi:10.3847/1538-4357/adda45. arXiv: 2505.09097 [astro-ph.CO]
Ord, S. M., et al. 2019, 36, e030. doi:10.1017/pasa.2019.17. arXiv: 1905.01826 [astro-ph.IM]
Ord, S. M., et al. 2015, PASA, 32, e006
Prabu, T., et al. 2015, ExA, 39, 73
Satav, S. M., & Rama Sarma, V. V. 2008, 10th International Conference on Electromagnetic Interference & Compatibility, 559
Sokolowski, M., et al. 2022, 39, e015e015. doi:10.1017/pasa.2021.63. arXiv: 2204.05873 [astro-ph.IM]
Sokolowski, M., et al. 2020, 37, e021. doi:10.1017/pasa.2020.17. arXiv: 2005.02041 [astro-ph.IM]
Sutinjo, A. T., et al. 2015, IEEE TAP, 63, 5433
Tingay, S. J. 2025, Galaxies, 13, 107
Tingay, S. J., et al. 2013, PASA, 30, e007
Tremblay, S. E., et al. 2015, PASA, 32, e005
Trott, C. M., et al. 2016, 818, 139. doi:10.3847/0004-637X/818/2/139. arXiv: 1601.02073 [astro-ph.IM]
van Haarlem, M. P., et al. 2013, A&A, 556, A2
Wayth, R., et al. JATIS, 8, 011010
Wayth, R., et al. 2017, PASA, 34, e034
Wayth, R. B., et al. 2018, PASA, 35, e033
Zarka, P., et al. 2015, in International Conference on Antenna Theory and Techniques (ICATT, ICATT 2015, 7136773. doi:10.1109/ICATT.2015.7136773
Zheng, Q., et al. 2016, 832, 190. doi: 10.3847/0004-637X/832/2/190. arXiv: 1602.06624 [astro-ph.GA]