

Data Transport for the Square Kilometre Array

Dr. Brian SMITH

*Optelian Access Networks, 1 Brewer Hunt Way, Ottawa, K2K 2B5, Canada
Tel: 1-613-287-2000, Fax: 1-613-287-2001, Email: brian.smith@optelian.com*

Abstract

The future Square Kilometre Array (SKA) radio telescope will present some of the most significant technical challenges faced by the global radio astronomy community in its attempt to construct the largest terrestrial observing instrument conceived to date. The science goals for the array call for high resolution over a wide frequency range (70 MHz to 10 GHz) with antennas concentrated in a core region, and more sparsely distributed antennas in outer regions. The high-frequency dish array will be installed in the Karoo region of South Africa. A low-frequency aperture array, operating at up to 1 GHz and distributed over a few hundred kilometres, will be installed in Western Australia. Signals from each of the antennas in both the dish and aperture arrays will be fed back to a central correlator for processing.

The engineering challenges associated with this project are significant, including antenna design for short wavelengths, high resolution, wideband analogue-to-digital conversion and advanced signal processing to generate results from the recovered data. This paper concentrates on the high-frequency dish array to be constructed in South Africa, and the challenges associated with data transport between antennas and the central correlator. It discusses the expected requirements for bandwidth, and the potential architectures and transmission technologies that could be used to meet these requirements.

Keywords

Square Kilometre Array, SKA, Optical Transport, Network Applications, High Capacity Networks, DWDM, 100 Gb/s

1. Introduction

Radio telescopes have been gathering information about the formation and structure of the universe for many decades, taking advantage of the transparency of our atmosphere to radio waves, and the fact that celestial sources can be observed both day and night. Early measurements were made on single antennas, until it was realized that, to probe further into the sky, instruments with enhanced resolution would be required.

2. Resolution of radio antennas

One of the major considerations in the design of radio telescopes is that of resolution. This is one of the fundamental drivers in the development of the Square Kilometre Array, so it is helpful to take a step back and look at a comparison of the resolution of typical observing instruments.

Observing Instrument	Resolution
Human eye (20/20 vision)	1 arc minute
Hubble Telescope	0.06 arc seconds
72 m Parabolic Radio Antenna (500 MHz)	28 arc minutes

We can convert the angular resolution numbers we see above into a more understandable frame of reference. For the human eye with perfect 20/20 vision, a resolution of 1 arc minute is equal to being able to discern a 25-cent coin at a distance of 66 m. For the Hubble telescope, a resolution of 0.06 arc seconds is equal to being able to resolve a 25-cent coin at a distance of 70 km.

Unfortunately, a single radio antenna has very poor resolution and the example shows that a parabolic antenna that is 72 m in diameter (which is nonetheless a large structure) at 500 MHz would only be able to recognize the coin if it was at a distance of less than 3 m.

If we were to build a radio antenna with the same angular resolution as the human eye at 500 MHz, it would have to be 2 km in diameter.

If we were to construct a radio antenna with the same angular resolution as the Hubble telescope, it would require a diameter of 2000 km.

It is obviously not feasible to build a single radio antenna with these dimensions, so the only practical way to achieve this is to use smaller, less expensive antennas and connect them as an interferometer array.[Spencer, 1989]The distance between two array elements is called the ‘baseline’.

The collecting area of two small antennas connected in a two-element array is still small compared to a single antenna of equal size, and this has an impact on the sensitivity. However, if many small antennas are connected together with many different baselines, the gaps in collecting area can effectively be filled in and the sensitivity of the instrument improves, as shown in Figure 1, where 18 antennas can have 153 different baselines.




Figure 1: Antenna baselines connected to increase the collecting area of an array

3. The Square Kilometre Array

The Square Kilometre Array will be the largest, most sensitive terrestrial radio telescope in the world with an aperture of up to one million square meters. The scientific goals of the project call for very high resolution over a wide frequency range: from 70 M Hz up to 10 GHz split between a low-frequency aperture array to be constructed in Western Australia, and a high-frequency antenna array to be built in the Karoo region of South Africa.

The parabolic antenna array will be constructed in two phases, with completion targeted for 2020. By the time the second phase is completed, there will be a total of 3000 15-m antennas distributed over a circular region with a radius of 3000 km. 600 antennas will be located in a core region with a radius

of 1 km. An additional 900 antennas will be positioned between 1 and 5 km from the centre, with a further 900 antennas in an intermediate zone between 5 and 180 km. An additional 600 antennas will be located within an outer zone between 180 and 3000 km from the centre. The layout of the antennas in the array will ensure the most even distribution of measurement sensitivity across the full scale of the aperture.

With a maximum baseline of 3000 km and a high-frequency range from 300 MHz to 10 GHz, the antenna array will have a resolution at high frequencies that is an order of magnitude better than the Hubble telescope. In fact, for the maximum baseline with the highest frequency, the resolution of the array is equal to identifying a 25-cent coin over a distance of 2000 km. At the lowest frequency of 300 MHz, the array will have a resolution on par with the Hubble telescope.

The current proposals for the layout of the antennas would follow a circular Gaussian distribution for the 1500 antennas within the core and inner regions, while the intermediate and outer regions would follow five spiral arms, with clusters of antennas placed at uniform logarithmic intervals along each path, as shown in Figure 2 [Millinaar, 2010]. All of the antenna data will be transported back to a correlator near the centre of the core region.




Figure 2: Proposed layout of high frequency antennas

4. Data Transport

The intention is that all of the data from the antennas will be transmitted in real time back to the central correlator for processing. The implication is that significant investment will be required in optical fiber infrastructure to enable this. The region proposed in South Africa was chosen specifically to minimize the amount of terrestrial radio frequency interference, thus the area is not densely populated. For this reason, it is highly unlikely that any pre-installed dark fiber resources can be employed.

The layout proposed (Figure 2) for the intermediate and outer regions is amenable to optical transport of the antenna data via fiber installed along the spiral arms of the infrastructure. The core and inner region connectivity will present a challenge for fiber routing due to the distribution of antennas, but the distances are small. Work is underway to try and optimize the location of fiber installations. (Garci, 2011) To understand how such an optical network might be architected, we need to calculate the expected bandwidth requirements between the antennas and the central correlator. The maximum bandwidth supported by each of the antennas to meet the science goals is 10 GHz per polarisation. The analogue signal will need to be digitized using a sampling rate at least 2 times the

maximum frequency. For many implementations of radio antenna technology, an analogue-to-digital convertor with a resolution of 2 bits per sample is sufficient, but in cases such as these where more tolerance to ground- and space-based RFI is required, a minimum of 4 bits per sample is being proposed (Dewdney, 2009). The minimum data rate for each antenna per polarisation can be calculated as:

$$\text{Rate} = 10\text{GHz} \times 2 \times 4 \times 1.25 = 100\text{Gb/s}$$

Bandwidth
Nyquist sampling rate
Resolution of A/D convertor
line coding

Each antenna must transmit data from two polarisations, so the total data rate for each antenna is 200 Gb/s. Given that the entire structure will have 3000 antennas by the time Phase 2 is complete, suggests a total capacity requirement of 600 Tb/s.

Using recent forecast data, global internet traffic is predicted to reach 100 exabytes (10¹⁸) per month by 2016. Assuming a CAGR (Compound Annual Growth Rate) of 25%, it is estimated that global internet traffic will be 750 Tb/s by 2020.

The implication is that, by the time Phase 2 of the Square Kilometre Array is completed and operational, it will be carrying the equivalent of 80% of the global internet traffic over the South African based antenna array alone.

5. Transport Technology

Given the scale of the data transmission requirements, there is no doubt that DWDM will play a central role in the configuration. Assuming that a single DWDM wavelength is used for each polarisation per antenna, each wavelength will be required to transport a single serial 100-Gb/s signal back to the correlator.

How many wavelengths will be required? If we look at a single spiral arm as an example, and focus on the intermediate and outer regions from 5 to 3000 km from the core, a total of 1150 antennas will be distributed along the five spiral arms of the array. Within the intermediate region, there will be 50 stations with 11 antennas being connected to each station. In the outer region, there will be 25 stations with 24 antennas connected to each station. The connection from each antenna to a station is assumed to be less than a few kilometres as shown in Figure 3.




Figure 3: Potential configuration of stations and antennas in outer regions

If we assume that the stations are equally distributed between the spiral arms, then each arm will have 5 stations of 24 antennas in the outer region, and 10 stations of 11 antennas in the intermediate region.

With 2 polarisations per antenna, the total number of wavelengths required along a single spiral arm of the array will be 460. This is significantly more than a single fiber pair can carry with today's DWDM technology.

The best spectral efficiency available today is based on the ITU-T DWDM grid with 50-GHz wavelength spacing. In the C-band, this allows up to 80 wavelengths to co-propagate on a single fiber pair, so a single spiral arm will require a minimum of six fiber pairs to be trenched and installed along the path between stations.

6. Advanced Modulation for 100 Gb/s

Today's carrier networks are typically based on DWDM technology with either 100-GHz or 50-GHz wavelength spacing. The standard data rate carried in most global networks is 10 Gb/s. For these rates, simple on/off keying (OOK) modulation is acceptable to enable transmission through many cascaded optical filter elements, without significant optical penalties. OOK encodes 1 bit per symbol, so the symbol rate is equal to the bit rate i.e. 10 Gbaud.

For 100 Gb/s serial transmission, On/Off keying is not an option for the following reasons:

- the spectral width of the 100G OOK signal is 10 times larger than that of an equivalent 10G OOK signal, making it too wide to pass through DWDM network elements without excessive optical penalties
- the modulator, driver and receiver electronics required are difficult to fabricate, not readily available, and would be excessively expensive
- for OOK, the tolerance to noise, chromatic dispersion, polarisation mode dispersion and non-linear impairments would limit the transmission reach to a few kilometres before regeneration was required, rendering such a solution uneconomical

Dual Polarisation, Quadrature Phase Shift Keying (DP-QPSK) is the modulation format of choice for 100-Gb/s transmission because it codes 4 bits per symbol using in-phase and quadrature-phase components of two polarisation states (see Figure 4). DP-QPSK reduces the spectral width by a factor of 4, reducing a 127-Gb/s data rate to a symbol rate of 31 Gbaud. This enables transmission through many cascaded DWDM network elements at 50- or 100-GHz spacing with minimal penalty. DP-QPSK technology with coherent detection is capable of transmitting up to 1200 km without regeneration.




Figure 4: Comparison of modulation formats for 100 Gb/s DWDM transmission

7. Potential Network Architecture

The technology is available to allow transmission of serial 100 Gb/s over the distances required for the Square Kilometre Array using a combination of DP-QPSK modulation with coherent detection and standard 50-GHz DWDM infrastructure. From the discussion above, it is clear that the minimum requirement would be six fiber pairs on each of the five spiral arms of the array, with each fiber pair carrying up to 80 wavelengths. With this architecture, the traffic from each antenna at each station can be distributed evenly across the 6 fiber pairs. Figure 5 shows the high level concept of a potential architecture.




Figure 5: Potential architecture for the (a) 5 stations in the outer region and (b) 10 stations in the intermediate region of the spiral arm


Figure 5(a) considers the required layout of the fiber pairs servicing the outer region from 180 to 3000 km. It shows the five stations with an assumed separation of 560km. Intermediate locations would be required between each of the stations to deploy optical amplifiers with optimal spacing. The architecture shown assumes that each fiber pair at each station uses a Reconfigurable Optical Add Drop Multiplexer (ROADM), and that each ROADM in the outer region allows for the addition of 8 wavelengths. This covers the 48 wavelengths required at each of the five stations in the outer region. A ROADM is a necessary element of the design. Each fiber pair will carry up to 80 wavelengths of 100-Gb/s traffic over a 3000 km distance. To achieve this optical reach with minimal regeneration, accurate automatic power balancing at each add/drop location is required; a function that can only be provided by ROADM elements.

Figure 5(b) shows the layout for the same fiber pairs but within the intermediate region of the spiral arm between 5 and 180 km. It shows the 10 antenna stations with an assumed separation of 20 km. The station locations would be a combination of ROADM and amplifier elements as needed to meet the optical performance requirements. Again, the ROADM would be required at each of the station locations to provide optimal spectral flatness across the 80 wavelengths, and ensure maximum un-regenerated reach for each wavelength.


8. Antenna to Station Interconnect

Figure 3 showed the potential configuration of stations along each spiral arm of the array and how the individual antenna elements may be connected to each station. The distance between each of the antennas and the station is likely to be a maximum of 2 km. In this situation, it would be overkill to deploy DWDM and more cost effective to deploy two fiber pairs (one pair per 100 Gb/s) between the stations and each antenna. Assuming that a 100-GbE interface is available on the digitization and packetization equipment at the antenna site, a simple, commercially available LR4 CFP (100G small form factor pluggable transponder) device can be used to deliver the traffic to the station. An LR4 CFP converts a 100-GbE flow into 4 lanes of 28-Gb/s traffic that are then transported over a single

fiber using an internal CWDM (Coarse Wave Division Multiplexer) configuration as shown in Figure 6(a). Figure 6(b) shows the connection over the dual fiber spurs from the antenna to the station and on to the DWDM ROADM interface for long-haul transmission to the correlator.



(a) LR4 CFP Architecture



(b) 100-GbE direct-on fiber connection between antenna and station using CFP

Figure 6: (a) Pluggable CFP technology for 100-GbE interfaces and (b) Connection between antenna site and station using two fiber pairs

9. Conclusion

The Square Kilometre Array is shaping up to be a project of unprecedented scale that will drive the development of new technology, but also push the capabilities of currently available technologies to the limit. The signal and data transport architecture alone will present significant challenges associated with optimal routing and bandwidth capabilities. At first glance, given the science goals, it would seem that real-time optical interconnect within the array is feasible and indeed, can be achieved using technology that is currently commercially available.

References

- Dewdney P.E et al (2009) ‘The Square Kilometre Array’ *Proc. IEEE*. Vol. 97, No. 8 pp. 1482-1496
- Gauci A. et al. (2011) ‘Optimal SKA Dish Configuration using Genetic Algorithms’, *Mon. Not. R. Astron. Soc.* Vol. 000, pp. 1-17.
- Millenaar R.P., Bolton R. (2010) ‘SKA Configurations Design: Status January 2010’ *Doc. # WP3-050.020.000-R-001*. Rev 1.2
- Spencer R.E (1989) ‘Theory of Interferometry and Aperture Synthesis’ In: Felli, M., Spencer R.E, *Very Long Baseline Interferometry*. Springer.