

Tianlai Data Analysis Center

21cm Cosmology

Neutral hydrogen (HI) in the universe is producing copious numbers of radio photons via the hyperfine spin flip transition which produces a narrow line at a wavelength of 21cm. The 21cm line is unique in cosmology in that it is the dominant astronomical line emission over the broad range of frequencies corresponding to cosmological redshifts. So to a good approximation the frequency of a feature can be converted to a Doppler redshift or blueshift without having to first identify the atomic transition. Making a map of the redshift and angle distribution of this line would give us a map of the spatial distribution of HI in the universe. HI is just as good a tracer of the large scale structure (LSS) of the universe as optically bright galaxies which are used in more traditional redshift surveys. Any of these LSS maps can be used to study dark energy e.g. by tracking the angular and redshift scale of baryon acoustic oscillations (BAOs). An advantage of the 21cm technique is that it is very easy to determine very accurate redshifts which is the most difficult part for optical redshift surveys.

Another reason to pursue 21cm LSS mapmaking is its future potential. HI 21cm emission and absorption occurs even before galaxies form, i.e. during the “dark ages”. In principle this technique can be extended to study the LSS in the majority of the cosmological volume which we can only see during their dark ages.

Three main reasons why 21cm is not currently a prominent redshift survey technique is: 1) it wasn't appreciated that making “intensity maps” with telescopes that cannot resolve individual distant galaxies would be useful, 2) foreground emission in these bands is large and the possibility of removing them was not fully explored, 3) it is only the availability of inexpensive fast digital electronics that prevents this from being a “big” (\$100M+) project.

Currently 21cm redshift surveys are at the stage of validating that intensity mapping actually works with pilot surveys. There has been some success with the single dish non-interferometric Green Bank Telescope but it is clear that one can do much better with special purpose radio interferometers. Note that even the pilot projects will survey vast volumes of the universe, larger than current optical redshift surveys, even if the angular resolution and noise of the pilot maps are not as good. There is no reason to expect that 21cm maps could not eventually match or surpass optical maps in quality on BAO scales if they are developed.

The Tianlai Project

Tianlai (translated from Chinese as “heavenly sound”) is a project to make one of the first large scale cosmological map of 21cm emission. The main instrument for the initial pilot stage is a large three cylinder transit radio interferometer (no moving parts) currently undergoing the final stages of construction in western China. First light observations, with a rudimentary part of the telescope working, was made in March 2015. Construction will be completed by the summer of 2015.

The project will map the 21 cm emission over the redshift range of $z \in [0.775, 1.029]$ and over 75% of the sky (the map quality will vary over the survey area). This covers more than 80 Gpc^3 which is larger than the volume surveyed by the Dark Energy Survey although with poorer map resolution and larger map noise. Improvements upon the Tianlai pilot program will improve both resolution and noise. These very large volumes are still less than 1 % of the observable universe.

Overview of Proposal

It is proposed that much of data reduction and analysis for the Tianlai 21cm redshift survey be done at Fermilab. The basic plan is that data tapes are generated for all of the data, either at the telescope site or copied from disk drives in some intermediate location, e.g. the National Astronomical Observatories of China in Beijing. The tapes are shipped to Fermilab where they are “archived”. The order of magnitude of data is 1.5 petabytes per year starting around the Fall of 2015. The data size is linear in the observing time and this estimate assumes no downtime and is thus an upper limit. Note that unlike optical astronomy, radio astronomy observations can be done both day and night. This time stream data will be reduced in an operation which is linear in the data size and will decrease the data by nearly two order’s of magnitude. It is expected that there will be a learning curve and this initial data reduction will be done a few times, more frequently in the beginning when we are still “learning” and the accumulated data size is small. The reduced data will then be processed into science level maps which can then be analyzed by the collaboration and will eventually be made public. The maps are only a few terabytes in size. The computer resources required is data storage and data processing. The later can be done efficiently using Fermilab’s computer farms.

Need for Tianlai Data Analysis Center

In order to analyze the Tianlai data requires computer resources in terms of storage and CPU cycles more than any of the collaborators has at hand. On the other hand these resource requirements are rather small when compared to a number of other experiments, such as particle collider experiments. It makes sense to explore the possibility of using Fermilab’s computer resources to do the Tianlai analysis. The project requires a large amount of storage for storage of the accumulating *time-ordered data* (~1 petabytes/year) which needs to be *reduced* to a *calibrated data cube* (~10 terabyte). This would further be reduced to science data product of

- 1) 3D intensity/polarization map on sky
- 2) catalog of radio transients.

The intensity/polarization map would then be decomposed in

- 1) 3D HI maps with errors, flags, etc.
- a) HI power spectra
- 2) intensity/polarization map of radio foregrounds.

Both the data cube and maps would then be available to the collaboration at Fermilab and copies shipped to the main collaborating institutions in China, France, and the US.

For reasons described below we expect to perform a data reduction of the time series data several, maybe 10’s of times, during the project, more frequently in the early stages of the project as we learn. At the early stages (say the 1st month of data) the accumulated time series data size is small, less than 100TB. We would propose to store the time series data on tape, needing ~2PB over the 1st two years. The reduction techniques are of order N (proportional to the data size and may well be limited by data I/O), embarrassingly parallel, and could be done efficiently on Fermilab’s computer farms. The data reduction from time series to data cube would be planned in advance and coordinated with CD personnel if that is desired. The much smaller data cube and maps would be available to the collaboration and we would like collaborators on the project to be able to submit jobs at Fermilab to analyze this data.

Note that while significant computational effort is needed to reduce and analyze the data, in terms of computational power a much greater number of FLOPS are performed at the telescope site by digital correlators. The unaveraged correlations are generated at > 1 PB/hour. It is only by averaging over 1 second intervals that this is reduced to the recorded time stream data.

Number of Correlations

From cross correlating N_{feed} dual polarization feeds the number of visibilities is

$$\text{Tianlai Data Center} = \underbrace{\frac{\text{real \&}}{2} \times \frac{\text{imaginary}}{2} \times \frac{\text{polarization}}{3} \times \frac{\text{feed}}{N_{\text{feed}}}}_{\text{After}} \times \frac{\text{pairs}}{(N_{\text{feed}}-1)} = 3 N_{\text{feed}} (N_{\text{feed}} - 1)$$

From auto-correlating N_{feed} dual polarization feeds the number of visibilities is

$$N_{\text{corr}}^a = \left(\underbrace{\frac{\text{real \&}}{\text{imaginary}} \times \frac{\text{only one}}{\text{cross pair}}}_{2} + \underbrace{\frac{\text{only}}{\text{real}} \times \frac{\text{two}}{\text{polarization}}}_{1} \right) \times N_{\text{feed}} = 4 N_{\text{feed}}$$

Adding the number of auto- and cross-correlations we find the total number of correlations is

$$N_{\text{corr}}^{\text{feed}} = N_{\text{corr}}^a + N_{\text{corr}}^x = (3 N_{\text{feed}} + 1) N_{\text{feed}}$$

This is before multiplexing the correlations into different channels.

Suppose we multiplex into N_{ch} channels in which case the total number of correlations we might keep track of is

$$N_{\text{corr}}^{\text{freq}} = (3 N_{\text{feed}} + 1) N_{\text{feed}} N_{\text{ch}}.$$

Maximal Data Rate

We average the multiplexed correlations and then store the averaged correlation functions at a frequency f_{sample} . For a transit telescopes the time between samples should be short enough that the Earth has not rotated significantly in that interval. Let D_{EW} be the maximum east-west dimensions of the interferometric array. The beam patterns on the sky have well sampled azimuthal wavenumbers $m \leq m_{\text{max}} = \frac{D_{\text{EW}}}{\lambda_{\text{min}}}$. There will be some spillover to larger m but with larger noise. The minimum R.A. angular scale well-probed (half a wavelength) is thus $\pi \frac{\lambda_{\text{min}}}{D_{\text{EW}}}$ so one would generally want $f_{\text{sample}} \gg \frac{1}{\text{day}_{\text{sidereal}}} \frac{D_{\text{EW}}}{\lambda_{\text{min}}}$.

One might want to sample even more frequently in order to facilitate downstream (*i.e.* not real-time) elimination or deweighting of RFI (radio frequency interference). The amount of data that is contaminated by or lost to short bursts of RFI is smaller the smaller the sample time. So if there were no costs we would make f_{sample} large. There is also an upper limit to f_{sample} which is given by the spectral resolution: $f_{\text{sample}} < \frac{1}{2} \delta v$.

If we store the time averaged correlations as $\#_{\text{bytes}}$ then the data rate is

$$R_{\text{data}} = (3 N_{\text{feed}} + 1) N_{\text{feed}} N_{\text{ch}} f_{\text{sample}} \#_{\text{bytes}}.$$

Initial Tianlai Cylinder Telescope

For the initial Tianlai cylinder telescope $N_{\text{feed}} = 96$ and $N_{\text{ch}} = 1024$ so

$$N_{\text{corr}}^{\text{feed}} = 27744$$

$$N_{\text{corr}}^{\text{freq}} = 28409856.$$

For initial Tianlai

$$D_{\text{EW}} \approx 30 \text{ m}$$

$$\lambda_{\text{min}} > 21 \text{ cm}$$

$$\delta v = 100 \text{ kHz}$$

so

$$m_{\text{max}} > 142$$

and we require

$$1.6 \text{ mHz} \ll f_{\text{sample}} < 50 \text{ kHz}.$$

For the purpose of sky sampling $f_{\text{sample}} \approx 0.1 \text{ Hz}$ would not be unreasonable but for improved RFI rejection we have chosen $f_{\text{sample}} = 1 \text{ Hz}$. Using 1 byte to store each of the real component of the averaged

correlations will keep the round-off error well below the other sources of noise.

We can thus write the walltime cylinder data rate as

$$R_{\text{data}}^{\text{cyl}} = (3 \times 96 + 1) 96 \times 1024 f_{\text{sample}}^{\text{cyl}} \#_{\text{bytes}} f_{\text{uptime}}^{\text{cyl}}$$

$$= 815 \frac{\text{TB}}{\text{ys}} \frac{f_{\text{sample}}^{\text{cyl}}}{1 \text{Hz}} \frac{\#_{\text{bytes}}}{1 \text{byte}} \frac{f_{\text{uptime}}^{\text{cyl}}}{1}$$

where TB = terabyte = 2^{40} byte, f_{uptime} is the fraction of the time the telescope is operating and we have used the sidereal year $\text{year}_S = 365.256$ days.

Initial Tianlai Dish Telescope

In addition to the cylinder telescope there is an array of 16 dishes each with a single dual polarization feeds. This will be used for various “experiments” but the correlator characteristics will be no different than for the cylinder array. For this interferometric array

$$R_{\text{data}}^{\text{dish}} = (3 \times 16 + 1) 16 \times 1024 f_{\text{sample}} \#_{\text{bytes}} f_{\text{uptime}} = 23 \frac{\text{TB}}{\text{year}_S} \frac{f_{\text{sample}}^{\text{dish}}}{1 \text{Hz}} \frac{\#_{\text{bytes}}}{1 \text{byte}} \frac{f_{\text{uptime}}^{\text{dish}}}{1}$$

This only adds 3% to the cylinder array data rate.

Total Data Rate

Combining the cylinder and dish array assuming a weighted average uptime and sample rate (heavily weighted toward the cylinder) is

$$R_{\text{data}}^{\text{total}} = 838 \frac{\text{TB}}{\text{ys}} \frac{f_{\text{sample}} f_{\text{uptime}}}{1 \text{Hz}}$$

The instantaneous data rate *i.e.* when everything is on

$$R_{\text{data}}^{\text{total}} = 28 \frac{\text{MB}}{\text{sec}} \frac{f_{\text{sample}}}{1 \text{Hz}} = 2.3 \frac{\text{TB}}{\text{day}_S} \frac{f_{\text{sample}}}{1 \text{Hz}}$$

where day_S is the a sidereal day which is 3.93... minutes short of a normal day. The sidereal day is a natural unit of data since in that time the Earth rotates once in an inertial frame and thus once with respect to the distant stars. This gives one full sample of the survey. To get the noise levels down we will need to average over 100s of days. This data rate will increase when the telescope is upgraded which might happen after a year of operation. We are only planning for the 1st year in this proposal.

Data To Tape

The time averaged correlations will initially be stored on disk drives but on site there will only be enough storage for a few days. It has been considered that one would ship the drives to Beijing with data and then back to the site for reuse. Another possibility is to write to tapes on-site. Presumably tapes are more secure transport medium although the read/write speed is slower.

One might store this data on LTO-6 tapes which holds 2.4 TB of data (uncompressed). Presumably we would use WORM (write once read many) tapes. The (uncompressed) write speed is 160 MB/sec which is five times the nominal instantaneous data rate so the tape drives would be fairly heavily used. One would need only one tape drive online however one probably want to keep at least one drive in reserve in case of failure. This is especially true if the tapes are written at the fairly remote telescope site. One would want a system which could detect a tape drive failure promptly.

One can purchase LTO-6 tapes for about US\$40 per tape in the U.S. from various manufacturers and vendors. One could store a day's data on one tape so tape cost would be about US\$15k per year for $f_{\text{uptime}} \approx 1$. A single tape drive currently costs about US\$2k. So long as the MTF is less than two months (~100 tape writes) the data drive cost would be less than the tape cost. These single tape drives require human operators for each tape to load and unload the tape and probably also to start the

tape write process. One would need to write a tape twice a day on average. The time to write a tape is roughly 4 hours. In terms of hardware the tape write system would probably cost \approx US\$20k/year.

Fermilab has its own data storage system including tape robots. I do not know if the tapes generated in China could be used directly or need to be “transcribed” onto a different medium. Whether the shipped tapes could be used as an archive or could be recycled (sent back to China for reuse) is unclear.