

CORRELATOR AND CALIBRATION FOR THE CANADIAN HYDROGEN INTENSITY MAPPING EXPERIMENT

Juan Mena-Parra

Doctor of Philosophy

Department of Physics

McGill University
Montreal, Quebec
July 2018

A thesis submitted to McGill University in partial fulfillment of the
requirements of the degree of Doctor of Philosophy

© Juan Mena-Parra, 2018

Abstract

This thesis describes the design, construction, and characterization of the Canadian Hydrogen Intensity Mapping Experiment (CHIME), a transit radio interferometer that will measure the emission of 21 cm radiation from neutral hydrogen in order to study the expansion history of the universe and probe the nature of dark energy. We discuss the scientific motivation of CHIME, followed by a description of the telescope design and characteristics. The CHIME pathfinder, which is a hardware, calibration, and data analysis proof-of-concept for CHIME is also presented. The CHIME correlator processes 2048 inputs across 400 MHz of bandwidth and it is the largest radio correlator that has been built. We have developed a general purpose hardware, firmware, and software framework -the ‘ICE’ system- that has been specialized to implement the data acquisition, F-engine, and the networking engine of this correlator. We describe the design of the ICE system, its use for the CHIME correlator, and demonstrate that the correlator complies with all the requirements for the experiment.

A major challenge of CHIME comes from the calibration needed in order to detect the 21 cm signal in the presence of astrophysical foregrounds that are many orders of magnitude brighter. We have used the pathfinder as a test-bed to develop a number of receiver-gain and digital calibration techniques in order to meet the stringent calibration requirements for full CHIME. We discuss the calibration requirements and present the different calibration strategies, the instrumentation, the current performance on the pathfinder, and the development towards the implementation for full CHIME.

Abrégé

Cette thèse décrit la conception, construction et caractérisation de l'expérience CHIME (Canadian Hydrogen Intensity Mapping Experiment), un radiotélescope interférométrique qui mesurera l'émission des radiations à 21 cm provenant des atomes d'hydrogène neutres afin d'étudier l'histoire de l'expansion de l'univers et la nature de l'énergie sombre. Nous discutons premièrement de la motivation scientifique du télescope CHIME, sa conception et caractérisation. Nous présentons ensuite le 'CHIME pathfinder' qui est un prototype qui a servi de preuve de concept pour le développement du hardware et des techniques d'étalonnage et d'analyse des données du CHIME. Nous décrivons également le corrélateur du CHIME, qui est le plus grand corrélateur radio qui a été construit à ce jour en permettant le calcul en temps réel des tous les produits entre les 2048 entrées du télescope sur une bande passante de 400 MHz. Nous ensuite décrivons le système 'ICE', qui est le matériel et le logiciel spécialisé conçu pour effectuer l'acquisition de données, la séparation des fréquences, la réorganisation des données et leur transmission au corrélateur. Nous démontrons ensuite que le corrélateur satisfait tous les requis de l'expérience CHIME.

L'un des défis les plus importants de CHIME vient de l'étalonnage précis qui est requis pour détecter les signaux de 21 cm en présence des sources astrophysiques qui sont plusieurs ordres de grandeur plus lumineuses. Nous avons utilisé le pathfinder comme une plateforme d'essai pour développer plusieurs techniques d'étalonnage des récepteurs et d'étalonnage numérique du corrélateur Afin de satisfaire les prérequis de CHIME. Nous discutons les prérequis d'étalonnage de CHIME et présentons les différentes techniques d'étalonnage, l'instrumentation, les performances actuelles avec pathfinder, et le développement vers la mise en œuvre du système d'étalonnage de CHIME.

Dedication

Para Jorge Mena-Parra. Hasta luego hermano.

Acknowledgements

I would like to thank my supervisor Matt Dobbs for his guidance and advice throughout my graduate studies. I am grateful for all the support I have received from the Canadian Hydrogen Intensity Mapping Experiment (CHIME) collaboration; in particular Kevin Bandura, Seth Siegel, and Richard Shaw. I would also like to thank Adam Gilbert and Jean-François Cliche for their patience and guidance during this project. I am grateful to all the members of the McGill Cosmology Instrumentation Lab for making my work environment enjoyable, as well as productive. Finally, I wish to thank my family for their support and encouragement throughout my studies.

Contribution of Authors

This thesis was written by me. It is organized into seven chapters including separate chapters for the introduction and the conclusions. Some of the chapters have been adapted from published papers. Since CHIME is a large collaboration, some of the subjects covered in this thesis in order to maintain its logical flow include contributions from other members of CHIME. Therefore, in this section I will clarify my contributions to each chapter, as well as places where work was done by others.

Chapter 1 is based on the introduction from my Master's thesis [1]. The work in this Master's thesis, including the introduction and all the derivations and calculations, was performed entirely by me under the supervision of Matt Dobbs.

In Chapter 2 I present important concepts of radio interferometry that I use throughout this thesis. At the end of this chapter I introduce the m -mode formalism technique for the analysis of transit radio telescopes like CHIME. This technique is very important to understand how CHIME maps the radio sky and to determine the telescope calibration requirements that are presented in Chapter 5. The development of the m -mode formalism was led by Richard Shaw and is not part of my work.

In Chapter 3 I present details of the CHIME instrument. My contributions to the design, construction, and characterization of CHIME and the pathfinder are substantial. I have played a major role in the design and deployment of the correlator (Chapter 4) and also in the calibration of the telescope (Chapters 5-6). Furthermore, I have worked in many other projects throughout my Ph.D. that are not covered thoroughly (or not even mentioned) in this thesis, including the first measurements of the full CHIME beams, the development of the CHIME data analysis pipeline and the analog receiving system. In Figure 3.6 I include two plots that show the improvement in the performance of the CHIME analog receiving system with respect to the pathfinder prototype. These plots were made by Rick Smegal and, although I made contributions to this project, Rick Smegal was the main person in charge of the design of the analog receiving system for CHIME. Some of the information in this chapter is based on the work published in [2], which I co-authored. The design and construction of the CHIME feeds and the X-engine of the CHIME correlator were not part of my work, although I contributed to their installation. These projects were led by Meiling Deng and Keith Vanderlinde, respectively.

Chapter 4 presents the CHIME correlator. I was part of the team that developed the ICE system that is used to implement the digitization, F-engine, and corner-turn network of the correlator. My main contributions to the ICE system are the collaborative design of the digitizer daughter board that specializes the system for CHIME and other radio

interferometry applications, as well as the development of IceCore, the Python-based control software that provides the tools necessary to automatically detect, configure, operate, and monitor the system hardware in an application-specific manner. In addition, I have been the main person responsible for the detailed characterization and quality control of all the modules of the ICE framework. Finally, I performed the full end-to-end characterization, validation, and commissioning of the CHIME and pathfinder correlators.

ICE is a complex system and its design required the expertise of several members of the McGill Cosmology lab. The implementation of CHIME FPGA firmware was done entirely by Jean-François Cliche, the ARM software was developed by Adam Gilbert, and the circuit board design was done by Adam Gilbert, Matt Dobbs, Kevin Bandura, and Jean-François Cliche. These parts of the system were not my work, although I have acquired a comprehensive knowledge of them in order to expose all the functionalities of the ICE hardware and firmware to the end user through IceCore. The description of the ICE system in Chapter 4 is based on the work published in [3, 4], which I co-authored. I was the person in charge of the preparation of [3]. My main contributions for this paper were the description of the ICE software, the application of the ICE system for CHIME, and its use in other interferometry applications like the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX) and very-long-baseline interferometry (VLBI). The preparation of the companion paper [4] was led by Kevin Bandura.

Chapter 5 presents some of the techniques developed to calibrate the CHIME telescope for which I have played a leading role. I developed the hardware, analysis and software for the Broadband Injection Signal (BIS) calibration system in collaboration with Kevin Bandura and Seth Siegel. The analytic derivation of the BIS performance in Appendix E was done by Richard Shaw. My measurements and analysis with the BIS system deployed on the pathfinder contributed to the further refinement of this model and, in particular, to the correct determination of the residual eigenvalue statistics and the BIS signal-to-noise ratio in terms of measurable telescope parameters. The work published in [5] presents the different calibration strategies for CHIME and I wrote the section that describes an early version of the BIS system. The implementation of the digital calibration techniques on the CHIME software was done with the aid of Jean-François Cliche and Kevin Bandura. In this chapter I also present the temperature-based characterization of the analog receiving chain, which is relevant for my work on BIS. Although my contributions to the thermal characterization of the receivers has been limited, I have performed measurements of the stability of the analog receivers on the pathfinder based on both BIS data and bright point source transits. Figure 5.1 shows measurements receiver stability measurements

based on early data from full CHIME. These measurements were performed by Seth Siegel.

Finally, in Chapter 6 I investigate the biasing effect of the quantization in the magnitude and phase of the visibilities measured by digital correlators. The new knowledge derived in this chapter was submitted for publication in [6]. I carried out all the calculations, derivations, and simulations for this research. The paper was written entirely by me with the advice and consultation of Kevin Bandura, Matt Dobbs, Richard Shaw, and Seth Siegel.

Contents

Abstract	iii
Abrégé	iv
Dedication	v
Acknowledgements	vi
Contribution of Authors	viii
List of Tables	xiii
List of Figures	xxv
1 Introduction	1
1.1 Dark Energy	1
1.1.1 Cosmological Model	2
1.2 Baryon Acoustic oscillations	4
1.2.1 Cosmic Sound	4
1.2.2 Standard Ruler	5
1.3 21 cm Radiation	8
1.3.1 21 cm optical depth and brightness temperature	10
1.3.2 Power spectrum	13
1.4 Outline	15
2 21 cm cosmology with transit interferometers	16
2.1 Stationary random processes	16
2.2 The narrowband condition	17
2.3 The two-element interferometer	19
2.3.1 The response of the receiving system	22

2.3.2	Interferometric fringes	24
2.3.3	The response of the interferometer: the complex visibility	27
2.3.4	Sensitivity of a radio interferometer: the radiometer equation	29
2.4	Mapping the 21 cm universe with CHIME	32
2.4.1	The m -mode formalism	32
2.4.2	Power spectrum from m -modes	35
2.5	Conclusions	36
3	The CHIME instrument	37
3.1	Science requirements	37
3.1.1	Frequency coverage	37
3.1.2	Angular resolution	38
3.1.3	Frequency resolution	39
3.1.4	Sensitivity and integration time	39
3.2	The CHIME telescope	40
3.2.1	The CHIME pathfinder	40
3.2.2	Signal path	44
3.3	Analog chain	45
3.3.1	Feeds	45
3.3.2	Low Noise Amplifier	46
3.3.3	Filter amplifier	47
3.4	The CHIME correlator	48
3.4.1	F-engine	48
3.4.2	X-engine	51
3.5	Conclusions	52
4	The ICE system and the CHIME correlator	54
4.1	The ICE system	54
4.1.1	ICE Hardware	56
4.1.2	ICE Firmware and Software	60
4.2	The digitizer, F-engine, and networking engine of the CHIME correlator	61
4.2.1	Digitization	61
4.2.2	Frequency channelization (F-engine)	62
4.2.3	Corner-turn network	63
4.3	Installation of the ICE system for the pathfinder correlator	67
4.4	Validation of the ICE system on the pathfinder	67
4.4.1	Manufacturability	71

4.4.2	Uptime and reliability	72
4.4.3	Data transport integrity	73
4.4.4	Analog performance and data acquisition integrity	76
4.4.5	Timing	81
4.4.6	Power consumption	88
4.5	Installation of the ICE system for the CHIME correlator	88
4.6	Conclusions	90
5	Calibration	93
5.1	Calibration requirements	93
5.2	Receiver gain calibration	94
5.2.1	Thermal characterization	94
5.2.2	Broadband Injection Signal calibration	95
5.3	Digital calibration	104
5.3.1	8-bit ADC calibration	105
5.3.2	Minimization of quantization errors in the computation of the FFT	106
5.3.3	Post-channelization digital calibration	108
5.4	Conclusions	110
6	Quantization bias for digital correlators	112
6.1	Post-channelization quantization in the CHIME correlator	112
6.2	Real-valued quantizer	114
6.3	Complex-valued quantizer	119
6.4	Real-valued correlator	120
6.5	Complex-valued correlator	125
6.6	Implications for radio interferometry	127
6.7	Conclusions	133
7	Conclusions	135
A	Power spectrum of $\delta_{21}(\mathbf{r})$	138
B	The source coherence function	140
C	Flat-sky approximation	144
D	BIS gain estimation	148

E	Details of expected BIS performance	150
E.1	Eigenvalue statistics	150
E.2	Metrics for BIS performance	153
F	$\langle ve \rangle$ for a real quantizer	155
G	Sign of $\langle ve \rangle$	157
	Bibliography	167

List of Tables

3.1	Design parameters for CHIME and the CHIME pathfinder	41
3.2	Design parameters for the CHIME and pathfinder correlator.	52
4.1	Bandwidth-baseline products for large contemporary radio interferometers.	64
4.2	Dynamic performance of the digitizer board in the CHIME band (400-800 MHz) and comparison to the ADC specs (given at 620 MHz only).	77

List of Figures

1.1	Early detection of the BAO peak in the correlation function of galaxies from the Sloan Digital Sky Survey (SDSS) of luminous red galaxies (LRG). Figure reproduced from [16].	6
1.2	Ratio of the matter power spectrum with BAO to the power spectrum without BAO, showing how the peak in the matter correlation function induces oscillations in the matter power spectrum. The forecast uncertainties correspond to the statistical limit of a two-year survey with a CHIME-like instrument at $z = 1.61$ (using redshift bins of size $\Delta z \approx 0.11$). Figure made by Michael Sitwell and adapted from [17].	6
1.3	Since the BAO scale can be determined theoretically, we can determine both the Hubble parameter and the angular diameter distance as function of redshift.	8
1.4	The hyperfine splitting of the ground state energy level of hydrogen. The hyperfine interaction breaks the degeneracy of the $1s_{1/2}$ states.	9
2.1	A basic two-element interferometer consists of two antennas separated by a distance b . The vector connecting the reference points of the two antennas is called the baseline vector \mathbf{b} . The antennas are observing a (far field) point source in the direction of the unit vector $\hat{\mathbf{n}}$. The wavefront of the radio source reaches antenna 2 first and then it reaches antenna 1 after a time τ_g called the geometric delay. The antennas have primary beams $A_1(\hat{\mathbf{n}})$ and $A_2(\hat{\mathbf{n}})$ respectively. Each antenna has a receiver that conditions (amplifies and filters) the antenna signal before sending it to the correlator which is the device that computes the correlation (multiplication and averaging) between the received voltages.	21

2.2	Response of the receiving system of a two-element interferometer to the voltage generated by the incident electric field from a point source. $v_{A1}(t)$ and $v_{A2}(t)$ are the antenna output voltages. The receivers have impulse response $h_1(t)$ and $h_2(t)$ and outputs $v_1(t)$ and $v_2(t)$ respectively. These are the inputs to the correlator which computes the time averaged product of $v_1(t)$ and $v_2(t)$	23
2.3	Normalized response of the two-element interferometer to a band limited signal with bandwidth $\Delta\nu$ and center frequency ν_c , with $\nu_c = 10.2\Delta\nu$. The geometric delay is τ_g . While the rapid oscillation (blue) is determined by ν_c , the envelope of the fringe pattern (red) is determined by $\Delta\nu$	26
3.1	The Dominion Radio Astrophysical Observatory (DRAO) is a facility for science and technology research related to radio astronomy. It currently operates five instruments: CHIME, the CHIME pathfinder, a 26-m fully steerable dish (the John A. Galt telescope), a 7-element east-west interferometric radio telescope (the synthesis telescope) and a solar flux monitor. It also features development laboratories for radio-frequency instrumentation. The observatory is located near Penticton, British Columbia, in a region regulated by the federal government to ensure a radio-quiet environment.	41
3.2	Top: the CHIME telescope, a hybrid cylindrical transit interferometer of four 20×100 m cylindrical dishes instrumented with a total of 1024 dual-polarization feeds. CHIME is currently being commissioned and it observed the first fringes from a point source transit late 2017 (photograph taken by Nolan Denman). Bottom: the CHIME pathfinder, a hardware, calibration, and data analysis proof of concept for CHIME. The pathfinder's first light was late 2013 and its commissioning finished early 2015.	42

3.3	Schematic diagram of the CHIME signal path. The signals from each of the dual-polarization feeds are independently amplified and filtered to 400-800 MHz. The F-engine digitizes and channelizes the signals from the 2048 analog receivers. It also re-arranges the data before sending it to the X-engine that computes the N^2 correlation matrix. These data are sent to the CHIME cosmology pipeline for storage and further processing. The X-engine also performs separate beamforming and upchannelization operations on the high-cadence data for two specialized backends for real-time FRB search and pulsar monitoring.	44
3.4	Schematic diagram of the CHIME analog receiver. At the focus of the cylinder, the sky signal received by the feed is amplified by an LNA with ~ 44 -38 dB gain in the 400-800 MHz band and ~ 30 K noise. After that, the signal is sent through 50 m of low-attenuation (-4 to -6 dB) coaxial cable to the receiver hut, where it is amplified again by $\sim 30 - 36$ dB and filtered to 400-800 MHz. The signal is about -23 dBm at the input of the ADC. This level causes a negligible penalty due to quantization errors, while allowing headroom for receiver gain fluctuations and external RFI.	46
3.5	Left: Photograph of the CHIME cloverleaf feed. Right: Photograph of an array of cloverleaf feeds installed at the focal line of the pathfinder. . . .	47
3.6	Top: Photographs of the CHIME LNA (left) and FLA (right) Bottom: Comparison between the insertion gain of CHIME (blue) and pathfinder (cyan) analog receiving chains (from the input of the LNA to the output of the FLA). The new analog receiver for CHIME has a flatter response over the CHIME band and higher rejection out of band. These measurements and figures were made by Rick Smegal.	49
3.7	The CHIME correlator. See Section 3.4 for a detailed description.	50
4.1	Photograph showing two ICE crates mounted in a standard rack. Each crate is populated with 16 motherboards. The motherboards have daughter boards designed for the CHIME application. The two crates process 512 analog inputs in total.	55

4.2	Top: Photograph of an ICE motherboard. The FPGA is located below the black heat sink visible near the center of the photo. Application specific daughter boards attach to the two slots labeled FMC A and FMC B. The board can be attached to an ICE backplane through the two black connectors seen at the bottom right. Two QSFP+ connectors at the back of the board provide eight 10 GbE data offload links. Bottom: Simplified block diagram of the ICE motherboard.	57
4.3	Left: Photograph showing the ICE backplane from the motherboard mating side. Right: Photograph showing the back of the crate-mounted ICE backplane and one motherboard attached. Eight dual-QSFP+ connectors at the back of the backplane provide data links between crates. Connectors for clock, timestamp, and synchronization signals are also at the back of the backplane.	59
4.4	Photograph of two custom digitizer daughter boards (red) installed on a motherboard (blue). Each daughter board includes eight inputs that can digitize at up to 1.25 GSPS with 8 bits. For CHIME, these inputs sample at 800 MSPS in the second Nyquist zone.	61
4.5	Schematic diagram of the channelizer module. It receives a data stream from the data acquisition module and passes it to the function generator sub-module which can select its source stream from the ADC or from an integrated test pattern generator. The data stream is then passed to a custom PFB/FFT that outputs a frame of 1024 complex frequency samples in a 18+18 bit format. The following scaler sub-module applies a 16+16 bit complex gain to each frequency channel and scales the result to 4+4 bit complex values. A separate prober sub-module can be configured to capture and buffer a subset of the raw ADC or scaler output data and send it back over the control channel for independent data monitoring. .	63
4.6	Schematic diagram of the corner turn. See Section 4.2.3 for a detailed description.	66
4.7	Installation of the first cassette on the west cylinder of the pathfinder in September 2013. Each cassette hosts 4 dual-polarization feeds. The second cassette was installed next to the first one.	68

4.8	Top: Photograph of the two independent 8-input correlators installed on the pathfinder in September 2013. The correlators use a Xilinx Kintex-7 FPGA KC705 evaluation board that used one of our custom FMC digitizer daughter boards. Both the F-engine and the X-engine are implemented on the FPGA. Bottom: Photograph of the first set of autocorrelation spectra obtained with the 8-input correlators. During the installation we found a defective analog receiver (green line), which was fixed afterwards.	69
4.9	Photograph of the ICE-based F-engine for the 16-input correlator installed on the pathfinder in March 2014. The X-engine was implemented on a separate GPU node.	70
4.10	Left: Side view of the ICE crate installed on the pathfinder for the 256-input correlator. Only 32 analog receivers are connected. Right: Bottom view of the ICE crate after connecting the 256 analog receivers.	70
4.11	Example of an eye diagram of the backplane 10 Gbit/s links between all 16 motherboards in a crate. The diagonal has no data because the link between a board and itself is done internally. The diagram is performed by varying the sampling point both in time (x-axis of each mini-plot) and signal amplitude (y-axis of each mini-plot) by discrete amounts and then measuring the BER relative to the optimal sampling point. A good link has a large, well opened center region (blue region of each mini-plot, this is the ‘eye’) where the BER is low. One defective link at transmit slot 12, receive slot 8 was revealed by this particular test. The poor performance of this link was caused by a motherboard connector assembly issue that was identified in this manner and fixed.	75
4.12	Power spectrum for a digitizer-board input when feeding a -6 dBm (1 dBFS) 581 MHz tone. The largest peak in the figure is the input tone. Smaller peaks with red asterisks denote spectral lines that are harmonics of this tone folded into the 400-800 MHz band. Other small peaks are mainly due to RFI, differential nonlinearity of the ADC (deviations of the ADC transfer function from its ideal behavior) and sampling-clock spurs from the on-board crystal used for this test. The cyan line denotes the white noise floor.	78

4.13	Left: Measurements of the power of the IMD3 products (red dots) and fundamental tones (blue dots) as function of the power of the input tones. The IMD3 products increase 3 dB per every dB increase in the power of the fundamental tones. The intersection of the third-order line produced by the IMD3 products with the line produced by the linear term determines the IIP3, which we measured to be ~ 13 dBm. Right: ADC spectra for two input tones 7 dBFS. The fundamental frequencies are $\nu_1 = 632.5$ MHz and $\nu_2 = 652$ MHz. The IMD3 products appear at $2\nu_1 - \nu_2 = 613$ MHz and $2\nu_2 - \nu_1 = 671.5$ MHz. The other tones are mostly harmonics that fold into the sampled band (for this particular test the sampling frequency was set to 850 MSPS).	79
4.14	Measurement of the digitizer-board crosstalk. A 1 dBFS tone is applied to the input 6 of the motherboard the power is measured for the remaining inputs of the motherboard as the tone is swept in frequency. On the remaining inputs we also feed a low amplitude tone (~ 1 bit RMS at 100 MHz, that folds into the 400-800 MHz band at 700 MHz) at a fixed frequency to exercise a few LSBs and smooth the ADC spectrum.	79
4.15	Change in the relative phase as function of frequency for each of the pathfinder inputs after a correlator restart as measured by the BIS calibration system. Each color represents the change in phase, in degrees, for a different input. Vertical stripes represent frequencies contaminated with RFI for which the measurement is not reliable. Points at $\pm 360^\circ$ correspond to either inputs flagged as ‘bad’ or frequency channels corresponding to GPU nodes which are down. For the vast majority of the inputs the relative phase is preserved up to the measurement error (which is approximately 4° for feeds well illuminated by the BIS system). For this particular test there are four inputs that have a jump in their relative phase that appears as a linear phase offset in frequency consistent with a one-sample delay with respect to the rest.	83
4.16	Setup for jitter measurement on the pathfinder ICE crate. A broadband noise source consisting of a terminated LNA connected to a FLA is installed in the pathfinder RF room. The signal is split into 16 copies and send to the same ADC input on the 16 different FPGA motherboards of the crate.	85

4.17	PSD of the delay fluctuations for each motherboard in the pathfinder ICE crate. The thick white line is the PSD of the temperature measured by a thermometer installed near the FPGA crate. The x-axis covers timescales between 1 hour and 20 seconds. The dashed vertical lines denote timescales of 15, 6, and 3 minutes from left to right. Note that the features in the temperature PSD are reproduced in the delay PSD, in particular the two peaks at 5 and 8 minutes corresponding to the RF-room air-conditioner cycle. This analysis was a team effort and this figure was made by Seth Siegel.	86
4.18	RMS of the delay fluctuations as a function of position in crate. Obtained by integrating the delay PSD. Each color corresponds to a different lower bound of the PSD integration (the upper bound is 20 s which is the maximum temporal frequency probed). Note that RMS of the delay fluctuations on time scales below 3 minutes is below ~ 0.6 ps. On 1 minute time scales the RMS of the delay fluctuations is even lower, well within the jitter requirements. This analysis was a team effort and this figure was made by Seth Siegel.	87
4.19	Simplified schematic of the full CHIME site layout. The ICE-based F-engine is installed inside two separate RF huts located between each pair of cylinders to minimize the receiver coaxial-cable length. The GPU-based X-engine is housed in two RF huts adjacent to the east-most cylinder. The FRB backend and the Power Distribution Center (PDC) are housed in a separate RF huts next to the GPU huts.	89
4.20	Photographs of the ICE hardware withing the FPGA huts at different stages of the installation. Top left: First crate installed in the east hut (front view, September 2016). Top right: Four ICE crates mounted in two racks within the east hut (front view, February 2017). Each crate pair processes one cylinder. Bottom left: West hut after installing the ICE crates and routing the data fibers, but before the networking was ready (rear view, May 2017). In addition to the two main racks, each hut has an auxiliary (right-most) rack that contains a GPS and clock distribution unit, a network switch that aggregates the data from the four crates (not shown), a control computer (not shown), and a fully operational spare crate. Bottom right: East hut after installing all the ICE hardware (rear view, May 2017).	91

4.21	First CasA fringes from a single east-west CHIME baseline and a single frequency channel. Top: real component (dotted blue), imaginary component (dotted red), and magnitude of the visibility (green). Bottom: phase of the visibility.	92
5.1	Median absolute deviation (MAD) of fractional amplitude fluctuations of the analog receiver gains in full CHIME from daily CygA transits after subtracting the common-mode variation. The x-axis denotes the feed number (its location in the array). The black dashed lines denote the separation between different cylinders (0-511 for west-most cylinder, 1536-2047 for east-most cylinder). The red dots denote the median value of the fractional amplitude fluctuations over all good frequencies. Shaded regions span the 68-th and 95-th percentile of good frequencies. After removing common mode, the most stable feeds show variations below the $\sim 0.3\%$ level and meet the calibration requirements. The large fluctuations for feeds in the approximate range 280-380 was due to a cassette with defective receivers on the west-most cylinder. These receivers were identified and fixed afterwards. This figure was made by Seth Siegel.	96
5.2	Schematic diagram of the BIS setup for the CHIME pathfinder. Inside the RF room a switched noise source controlled by the ICE crate generates the injection signal that is passed through a signal conditioning and distribution module and split into three copies. One copy of this signal is sent directly to the correlator as reference. The other two copies are sent to broadcasting antennas installed at the base of each cylinder.	97
5.3	Illumination profile (in dB) of the pathfinder feeds using a single wide-beam helical antenna per cylinder. The illumination as function of feed position (x-axis) along the cylinder and frequency (y-axis) is plotted for each cylinder and polarization separately. The horizontal stripes correspond to frequencies flagged as bad (either RFI or bad GPU nodes) while the vertical stripes correspond to feeds flagged as bad (e.g. faulty receiver).104	

5.4	Behavior of an 8-bit quantizer for a real-valued Gaussian input. Top: variance of the quantized output, $\hat{\sigma}^2$, normalized by the unquantized variance σ^2 , as function of the unquantized standard deviation σ . Middle: variance of the quantization error, σ_e^2 , normalized by the unquantized variance. Bottom: magnitude of the correlation coefficient between the input and quantization error, $ \rho_{ve} $. The lower bound of $ \rho_{ve} $ is limited by numerical precision. For each plot, the blue line corresponds to the quantization model developed in Chapter 6 while the red dashed line corresponds to the uncorrelated quantization noise model that is traditionally used in radio astronomy and assumes $\rho_{ve} = 0$	107
5.5	Schematic diagram of the post-channelization digital calibration. The digitized from the ADC is passed to the PFB/FFT sub-module as frames of 2048 8-bit samples. The FFT outputs a frame of 1024 complex values, one per frequency channel, in 18+18 bit format. For each frequency channel, the FFT output is multiplied by a 16 + 16 bit complex gain. The resulting 35 + 35 bit complex number is then shifted to the left a configurable amount of times. The amount of shifts to the left is the same for all the frequency channels. Finally, the 35 + 35 bit complex number is scaled to a 4 + 4 bit complex value by selecting the 4 MSBs in both the real and imaginary component and applying convergent rounding.	109
6.1	Behavior of a quantizer with $N = 15$ levels (left column) and $N = 16$ levels (right column) and a real-valued Gaussian input. From top to bottom row, the plots show the variance of the quantized output, $\hat{\sigma}^2$, the quantization error, σ_e^2 , and the correlation between the input and quantization error, $\langle ve \rangle$, as function of the unquantized standard deviation σ . All the values are normalized with respect to σ^2 . For each plot, the red line corresponds to equations 6.6 - 6.8, the thick blue line shows the results from simulations, and the green dashed line corresponds to the uncorrelated quantization noise model that assumes $\langle ve \rangle = 0$. Note that equations 6.6 - 6.8 predict accurately the results from simulations. When the input σ uses optimally the quantizer's dynamic range the quantization error is very weakly correlated with the input. In this case the uncorrelated quantization noise model provides a very good approximation, introducing only a small bias error.	117

- 6.2 Behavior of a complex-valued quantizer with $N = 15$ levels (left column) and $N = 16$ levels (right column) and a circularly-symmetric Gaussian input. From top to bottom row, the plots show the normalized variance of the quantized output ($\hat{\sigma}^2/\sigma^2$), the quantization error (σ_e^2/σ^2), and the magnitude and phase (in degrees) of the normalized correlation between the input and quantization error ($\langle ve^* \rangle/\sigma^2$). For each plot, the red line is our prediction and the blue line is the result from simulations. There is again excellent agreement between these. Note that $\langle ve^* \rangle$ is always real (in the simulation the imaginary part is consistent with zero at the $\sim 0.15\%$ level), and it is negative (180° phase) for N odd, while it becomes positive (0° phase) in the low σ regime for N even. 121
- 6.3 Results from simulations of $r = \hat{r}_{12}/r_{12}$ as function of σ_1 and σ_2 for different values of ρ for a real correlator with $N = 15$ levels (top row) and $N = 16$ levels (bottom row). The axes for each plot are the unquantized input signal levels and the green solid lines correspond to the highest level of the quantizer above which clipping occurs. The bias in \hat{r}_{12} for moderate values of ρ ($|\rho| \lesssim 0.85$) is below $\sim 0.1\%$ approximately within the inner white square enclosed by the region $\sigma_1 \times \sigma_2 \approx [2^{-0.6}, 2^{0.9}] \times [2^{-0.6}, 2^{0.9}]$. For $|\rho| \gtrsim 0.85$ the bias can increase up to $\sim 4\%$ 124
- 6.4 Correlation between the quantization errors of the two inputs, $\langle e_1 e_2 \rangle$ (normalized by r_{12}), as function of σ_1 , σ_2 and ρ , from simulations. Note that e_1 and e_2 are very weakly correlated as long as at least one of the two inputs is optimally quantized. 125
- 6.5 $|\hat{r}_{12}/r_{12}|$ from simulations as function of σ_1 and σ_2 for different values of $|\rho|$. For moderate values of ρ the bias in the magnitude of $|\hat{r}_{12}|$ is below $\sim 0.1\%$ within the inner square enclosed by the region $\sigma_1 \times \sigma_2 \approx [2^{-0.1}, 2^{1.4}] \times [2^{-0.1}, 2^{1.4}]$, corresponding to region where both inputs are optimally quantized. For $|\rho| \gtrsim 0.85$ this bias can increase up to $\sim 4\%$. . . 128

6.6 $\angle(\hat{r}_{12}/r_{12})$ (in degrees) from simulations as function of σ_1 and σ_2 for different values of $|\rho|$. The bias in the phase of \hat{r}_{12} is negligible when at least one of the inputs is optimally quantized. This is a less stringent requirement than that for the magnitude, which requires both inputs to be optimally quantized. The bias is below $\sim 0.1^\circ$ within the cross-shaped region where either σ_1 or σ_2 are in the approximate interval $[2^{-0.1}, 2^{1.4}]$. When $|\rho|$ is high (last two columns) the bias within this region can rise up to $\sim 1^\circ$. When $|\rho|$ is small (first two columns) the phase bias is below $\sim 0.1^\circ$ for all values of σ_1 and σ_2 , although there are still random fluctuations in the simulation at the \sim sub-degree level for very low values of σ (see text). 128

6.7 Bias due to quantization of the measured visibility as function of the source position θ and the signal-to-system-noise ratio is SNR . The visibility baseline is $b_\lambda = 0.4$ which corresponds to the CHIME normalized feed spacing at 400 MHz. When $\sigma_{sys} = 2^1\Delta$ (left panels), well within the optimal quantization interval, the quantization bias for weak sources ($SNR \lesssim 0.1$) is negligible. This is the regime for CHIME $\sim 99\%$ of the time. When $\sigma_{sys} = 2^2\Delta$ (right panels), which is the optimal input level according to the uncorrelated quantization noise model, the amount of bias increases significantly. 130

6.8 Complex quantization parameter $\eta_q = \hat{R}/R$ as function of the source position θ for $SNR = 0.1$ (top row, this SNR is the approximate upper limit of weak-source regime) and $SNR = 6$ (bottom row, this is the typical SNR of a strong source like the sun). For each plot, the blue labels and dots correspond to the magnitude of η_q and the red labels and dots correspond to its phase in degrees. Note that for $\sigma_{sys} = 2^1\Delta$ (left column), which is well within the optimal quantization interval for $N = 15$ levels, and in the weak-source regime ($SNR \lesssim 0.1$, top left plot), η_q is very close to being real-valued and deviates from unity by less than $\sim 3 \times 10^{-6}$ so the loss of beamforming efficiency due to quantization is negligible. If we set $\sigma_{sys} = 2^2\Delta$ (right column), the beamforming sensitivity reduces significantly even in the weak-source regime. This confirms that for this application the uncorrelated quantization model leads to important deviations from the expected performance of the interferometric array. . . . 132

B.1	When we study the response of an interferometer to an extended source we will assume that the source is spatially incoherent, that is, that radio wavers coming from different regions of the source are uncorrelated.	140
C.1	Geometry to understand the flat-sky approximation. The first and second axes of the coordinate system point in the east and north directions respectively. The third axis completes the right handed coordinate system. In this system the normalized baseline vector b_λ has coordinates (u, v, w) and the brightness temperature distribution of the source is $T(l, m)$. The reference direction is \hat{n}_0 .	145

Chapter 1

Introduction

Dark energy is a property attributed to the large-scale universe in order to explain its accelerated expansion. Understanding the nature of dark energy is one of the biggest challenges of cosmology today, with profound implications for fundamental physics and our understanding of the universe. The Canadian Hydrogen Intensity Mapping Experiment (CHIME¹) is a novel radio telescope designed to study the expansion history of the universe and probe the nature of dark energy. CHIME will map the three-dimensional distribution of neutral hydrogen gas in the universe by directly detecting its redshifted 21 cm radiation. By measuring the scale of the Baryon Acoustic Oscillations across the redshift range $z \approx 0.8$ to $z \approx 2.5$, CHIME will study the epoch when dark energy generated the transition from decelerated to accelerated expansion of the universe. In this chapter we discuss the science that motivated the construction of the CHIME telescope and give context for the subsequent chapters of this thesis.

1.1 Dark Energy

Observations since 1998 have established that the expansion of the universe is accelerating [7, 8]. Results from the Planck mission² provide a precise inventory of the energy constituents in the universe: assuming the standard model of cosmology, the total mass-energy in the universe today contains approximately 5% ordinary matter, 26% dark matter, and 69% the mysterious dark energy which, unlike the other constituents, counters the attractive force of gravity [9, 10].

One possibility is that dark energy is constant in time and uniform in space, taking the form of Einstein’s cosmological constant Λ so its equation of state (the ratio of its

¹<https://chime-experiment.ca/>

²http://www.esa.int/Our_Activities/Space_Science/Planck

pressure to its energy density) is $w = P_\Lambda/\rho_\Lambda = -1$ at all times. Another possibility is that dark energy is a dynamical fluid. In this case its equation of state would not likely be constant, but it would be a function of the scale factor of the universe, $a = (1+z)^{-1}$, where z is redshift. Different theories have different predictions for the evolution of the equation of state.

Whether dynamical or a constant, the equation of state of dark energy can be parametrized with a first order approximation

$$w(a) = \frac{P_\Lambda(a)}{\rho_\Lambda(a)} = w_0 + (1 - a)w_a \quad (1.1)$$

where w_0 is the present value of w , and w_a parametrizes the evolution of $w(a)$. This parametrization is useful if dark energy is important at late times but insignificant at early times [11]. It also provides a framework to compare theory to experiments.

1.1.1 Cosmological Model

For a detailed review of the cosmological model see [12, 13]. On large scales the universe is extremely isotropic and homogeneous. The space-time metric that describes a homogeneous, isotropic expanding universe is the Friedmann-Robertson-Walker (FRW) metric given by (using units with $c = 1$)

$$ds^2 = -dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2(d\theta^2 + \sin^2\theta d\phi^2) \right] \quad (1.2)$$

with (r, θ, ϕ) as the comoving spatial coordinates, a is the scale factor (normalized to 1 today) and k is a constant representing the curvature of the space ($k > 0$ for a closed universe, $k < 0$ for an open universe, and $k = 0$ for a flat universe). Applying the FRW metric to Einstein's field equations of general relativity gives the Friedmann equations

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P) + \frac{\Lambda}{3} \quad (1.3)$$

$$\left(\frac{\dot{a}}{a}\right)^2 = H^2(t) = \frac{8\pi G\rho}{3} - \frac{k}{a^2} + \frac{\Lambda}{3} \quad (1.4)$$

where P and ρ are the mean pressure and density of the contents of the universe and $H(t)$ is the Hubble parameter. Λ is the cosmological constant, which can be interpreted

as an additional component of the universe in the form of vacuum energy (dark energy) $\rho_\Lambda = \Lambda/(8\pi G)$ that is constant. From the last two equations it follows that

$$\dot{\rho} = -3H(\rho + P). \quad (1.5)$$

Assuming that all the components of the energy density are perfect fluids that obey the equation of state

$$P_i = w_i \rho_i \quad (1.6)$$

where w_i is constant, it follows that

$$\rho_i \propto a^{-3(1+w_i)} \quad (1.7)$$

with $w = 0, 1/3, -1$ for radiation, matter and the cosmological constant respectively. We now introduce the critical density, defined as the density for which the spatial geometry is flat,

$$\rho_c = \frac{3H^2}{8\pi G}. \quad (1.8)$$

We can define the density parameters of the different components of the universe today as

$$\Omega_r = \frac{\rho_{r0}}{\rho_{c0}}, \quad \Omega_m = \frac{\rho_{m0}}{\rho_{c0}}, \quad \Omega_\Lambda = \frac{\rho_\Lambda}{\rho_{c0}} = \frac{\Lambda}{3H_0^2} \quad (1.9)$$

where $\rho_{c0} = 3H_0^2/(8\pi G)$ is the critical density today, H_0 is the value of the Hubble parameter today and Ω_r , Ω_m , Ω_Λ are, respectively, the densities of radiation, matter and cosmological constant today relative to the critical density. Analogously, we define the parameter

$$\Omega_k = -\frac{k}{H_0^2} \quad (1.10)$$

so equation 1.4 can be written as

$$\frac{H^2}{H_0^2} = \Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_k a^{-2} + \Omega_\Lambda. \quad (1.11)$$

Ω_k is the ‘spatial curvature density’ today. $\Omega_k < 0$ for a closed universe, $\Omega_k > 0$ for an open universe and $\Omega_k = 0$ for a flat (Euclidean) universe. The spatial geometry of the universe has been measured to be nearly flat. If the equation of state for dark energy has $w \neq -1$, or if w changes with the scale factor, equation 1.11 can be generalized with [11]

$$\Omega_\Lambda \rightarrow \Omega_\Lambda \exp \left\{ 3 \int_a^1 \frac{da'}{a'} [1 + w(a')] \right\}. \quad (1.12)$$

1.2 Baryon Acoustic oscillations

Refer to [14] for a detailed review of Baryon Acoustic Oscillations (BAO) and their relevance for 21 cm cosmology. A brief summary is provided here. The early universe consisted of a hot plasma in which photons and baryons were tightly coupled via Thomson scattering. Quantum mechanical fluctuations that were present at the beginning of the universe and that expanded to scales larger than the Hubble volume during inflation caused the development of overdense regions consisting of baryons, dark matter, and photons. These fluctuations generate sound waves in the hot plasma (acoustic oscillations) of the early universe due to the competing forces of radiation pressure and gravity in the overdensity regions: when an overdensity gravitationally attracts baryons inwards, the photon-baryon plasma is compressed, thus heated, increasing the outward radiation pressure from the photons. After about 380000 years of expansion, the plasma cooled to below ~ 3000 K and the electrons and protons in the plasma could combine to form neutral hydrogen atoms. This recombination happened at a redshift $z \sim 1100$. Recombination removes the pressure from the baryons and effectively ends the wave propagation. These baryon and dark matter perturbations are the seeds of the formation of large scale structure in the universe at late times.

1.2.1 Cosmic Sound

The acoustic oscillations in the primordial plasma are also imprinted in the late-time baryon power spectrum. A simple way to understand this is to consider a single spherical

perturbation common to dark matter, baryons, and photons. The tightly coupled photon-baryon plasma will propagate outwards as an acoustic wave with a speed

$$c_s = \frac{c}{\sqrt{3(1+R)}}, \quad R = 3\frac{\rho_b}{\rho_r} \propto \frac{\Omega_b}{1+z} \quad (1.13)$$

where ρ_b is the baryon density [15]. The dark matter perturbation is left at a small radius from the origin of the overdensity since it only interacts gravitationally. At recombination, electrons and protons in the plasma combine to form neutral hydrogen atoms, photons interact to a much lesser degree with baryons and the pressure on the baryons is removed. This stops the baryon wave propagation while the photons propagate freely, forming what we now observe as the Cosmic Microwave Background (CMB), and leaving a shell of baryonic matter at a fixed characteristic radius from the initial perturbation. This baryon density excess will appear as a bump in the correlation function of the baryon density field at the radius of the spherical shell (Figure 1.1). The BAO feature will also appear as oscillations in the baryon power spectrum (Figure 1.2), since the correlation function and power spectrum form a Fourier pair (see Section 1.3 below). The gravitational interaction between baryons and dark matter implies that dark matter also preferentially clumps on the BAO scale.

The comoving distance a sound wave could have travelled in the photon-baryon fluid before recombination is [18]

$$r_s = \int_{z_{rec}}^{\infty} \frac{c_s dz}{H(z)} = \frac{1}{\sqrt{\Omega_m H_0^2}} \frac{2c}{\sqrt{3z_{eq}R_{eq}}} \ln \left[\frac{\sqrt{1+R_{rec}} + \sqrt{R_{rec}+R_{eq}}}{1 + \sqrt{R_{rec}}} \right] \quad (1.14)$$

where ‘*rec*’ and ‘*eq*’ refer to the epochs of recombination and matter-radiation equality respectively and $1 + z_{eq} = \Omega_m/\Omega_r$. The value of z_{rec} is very insensitive to cosmology since it is mainly determined by atomic physics (e.g. the theory of the ionization state of a gas in thermal equilibrium, [18]). The additional parameters in r_s are accurately determined by measurements in the CMB power spectrum, fixing the comoving sound horizon at $r_s \approx 148$ Mpc [10].

1.2.2 Standard Ruler

Since the scale of the comoving sound horizon is known, it can be used as a standard ruler to learn about dark energy and the expansion history of the universe. We can observe




Figure 1.1: Early detection of the BAO peak in the correlation function of galaxies from the Sloan Digital Sky Survey (SDSS) of luminous red galaxies (LRG). Figure reproduced from [16].




Figure 1.2: Ratio of the matter power spectrum with BAO to the power spectrum without BAO, showing how the peak in the matter correlation function induces oscillations in the matter power spectrum. The forecast uncertainties correspond to the statistical limit of a two-year survey with a CHIME-like instrument at $z = 1.61$ (using redshift bins of size $\Delta z \approx 0.11$). Figure made by Michael Sitwell and adapted from [17].

the preferred clustering scale set by the Baryon Acoustic Oscillations (BAO) at different redshifts to constrain the Hubble parameter and the angular diameter distance.

For a review of distance measures in cosmology refer to [13]. The fundamental distance measure, from which all other distances in the FRW metric are determined is the radial comoving distance

$$\chi(a) = \int_{t(a)}^{t_0} \frac{cdt'}{a(t')} = \int_a^1 \frac{cda'}{a'^2 H(a')} \quad \rightarrow \quad \chi(z) = \int_0^z \frac{cdz'}{H(z')} \quad (1.15)$$

where we used $dt = da/\dot{a}$, $H = \dot{a}/a$ and $a = 1/(1+z)$. Now we introduce the angular diameter distance³

$$d_A(z) = \frac{1}{1+z} \begin{cases} \frac{c}{H_0 \sqrt{|\Omega_k|}} \sinh \left[\frac{H_0}{c} \sqrt{\Omega_k} \chi(z) \right] & \text{if } \Omega_k > 0 \\ \chi(z) & \text{if } \Omega_k = 0 \\ \frac{c}{H_0 \sqrt{|\Omega_k|}} \sin \left[\frac{H_0}{c} \sqrt{-\Omega_k} \chi(z) \right] & \text{if } \Omega_k < 0. \end{cases} \quad (1.16)$$

In general, we need to know both Ω_k and $\chi(z)$ to determine cosmological distances. However, in the particular case of a flat universe cosmological distances are determined by $\chi(z)$ only and $d_A(z) = \chi(z)/(1+z)$. As shown in Figure 1.3, in this case the characteristic BAO scale, $r_{s\parallel}(z)$, along the line of sight provides a measurement of the Hubble parameter through

$$H(z) = \frac{c\Delta z}{r_{s\parallel}(z)} \quad (1.17)$$

while the tangential mode $r_{s\perp}$ provides a measurement of the angular diameter distance

$$d_A(z) = \frac{r_{s\perp}(z)}{\Delta\theta(1+z)}. \quad (1.18)$$

The BAO is a statistical standard ruler in the sense that the preferred scale in the

³The angular diameter distance is defined in the classic way, $d_A = s/\Delta\theta$, where s is the object's physical size and $\Delta\theta$ its angular size as viewed from earth. However, in an expanding universe, d_A depends on the assumed cosmology. For example, in a flat universe, the comoving size of the object is s/a and the comoving distance out to the object is $\chi(a)$, so $\Delta\theta = (s/a)/\chi(a)$. From this result it follows that $d_A^{flat} = a\chi(a) = \chi(z)/(1+z)$.




Figure 1.3: Since the BAO scale can be determined theoretically, we can determine both the Hubble parameter and the angular diameter distance as function of redshift.

matter distribution cannot be directly observed. The matter density distribution is a linear superposition of many acoustic waves like the one described in Section 1.2.1, making it very difficult to visually detect the characteristic scale. However, the BAO signature is still detectable statistically through the two-point correlation function of the matter distribution, $\xi(r)$, which quantifies the excess clustering in the spatial matter distribution as function of relative distance. The BAO will appear as a peak in the correlation function (Figure 1.1) and as oscillations in its power spectrum (Figure 1.2).

1.3 21 cm Radiation

In a neutral hydrogen atom (HI) the predominant interaction between the electron and the proton is the Coulomb interaction $V(r) \propto e^2/r$, where e is the magnitude of electron charge and r is the electron-proton distance. For a Hamiltonian with this potential energy, the energy levels are given by $E_n = -13.6 \text{ eV}/n^2$, where n is a positive integer. In addition, there is a fine splitting (the *fine structure*) of these energy levels due electron spin and relativistic corrections to this Hamiltonian. There is yet an even weaker interaction that involves the intrinsic spins of the electron and the proton: the electron interacts with the magnetic field generated by the proton's magnetic moment, generating an interaction energy proportional to the intrinsic spin of the two particles.




Figure 1.4: The hyperfine splitting of the ground state energy level of hydrogen. The hyperfine interaction breaks the degeneracy of the $1s_{1/2}$ states.

When two spin $1/2$ particles are combined in absence of orbital angular momentum (which is the case of the ground state of hydrogen), the net state has either spin 1 or spin 0. In fact, there are three spin 1 states, known as *triplet* states, and a single spin 0 state, known as the *singlet* state. From the analysis of the perturbation Hamiltonian, it follows that this intrinsic spin coupling breaks the degeneracy of the $1s_{1/2}$ states in hydrogen, raising the energy of the triplet, and decreasing the energy of the singlet (Figure 1.4). This energy splitting is known as *hyperfine structure*. Using time independent perturbation theory it can be found that the energy difference between the singlet and the triplet states is

$$\Delta E = \frac{8}{3} g_p \frac{m_e}{m_p} \alpha^2 E_0 = 5.88 \cdot 10^{-6} \text{ eV} \quad (1.19)$$

where $E_0 = 13.6 \text{ eV}$ is the (absolute value of the) ground-state energy, g_p is the proton gyromagnetic ratio and α is the fine-structure constant. This energy corresponds to a wavelength $\lambda_0 \approx 21.1 \text{ cm}$ and a frequency $\nu_0 \approx 1420.4 \text{ MHz}$. For a detailed discussion on the quantum mechanics behind the hyperfine splitting of the ground state of hydrogen refer to [19, 20].

The probability of a spontaneous transition from a triplet state to the singlet is extremely small (the spontaneous emission coefficient is $A_{10} = 2.85 \cdot 10^{-15} \text{ s}^{-1}$). However, as the total number of atoms of neutral hydrogen in the Intergalactic Medium (IGM) is very large, this emission line is easily observed by radio telescopes. Also, the lifetime can be considerably shortened by other mechanisms like collisions with other hydrogen atoms and interaction with the CMB.

1.3.1 21 cm optical depth and brightness temperature

In this section we show how 21 cm radiation can be used as a probe for cosmology. The intensity of photons emitted through the 21 cm transition is determined by the spin temperature, T_s , defined through the equation

$$\frac{n_1}{n_0} = \frac{g_1}{g_0} e^{-T_*/T_s} \quad (1.20)$$

where n_1 and n_0 are the number densities of electrons in the triplet and singlet states of the hyperfine level respectively, $g_1 = 3$ and $g_0 = 1$ are the respective statistical weights, and $T_* = h\nu_0/k_B \approx 0.068$ K is the equivalent temperature of the hyperfine transition energy. The spin temperature, whose name reflects the fact that the hyperfine splitting of the ground state of hydrogen arises from the coupling between the electron and proton spins, is just a shorthand for the ratio between the occupation number of the two hyperfine levels. This ratio establishes the intensity of the radiation emerging from a cloud of neutral hydrogen. Since all astrophysical applications have $T_s \gg T_*$, approximately three of four atoms are found in the excited state [21].

Even though, at the redshifts relevant to CHIME, the 21 cm radiation is seen only in emission, it is instructive to go through the full derivation of the 21 cm differential brightness temperature including absorption. If we consider the 21 cm transition as a two level system, then the absorption coefficient of a cloud of hydrogen can be found from the Einstein coefficients (see [22] for details)

$$\alpha(\nu) = \frac{c^2 A_{10}}{8\pi\nu^2} \frac{g_1}{g_0} n_0 \left(1 - \frac{n_1 g_0}{n_0 g_1} \right) \phi(\nu) \quad (1.21)$$

where $\phi(\nu)$ is the line profile; it is sharply peaked at ν_0 and describes the relative effectiveness of frequencies in the neighbourhood of ν_0 for causing transitions. The line profile is defined so that $\int d\nu \phi(\nu) = 1$. From equations 1.20 and 1.21 and using $T_s \gg T_*$, $\alpha(\nu)$ becomes

$$\alpha(\nu) = \frac{c^2 A_{10}}{8\pi\nu^2} \frac{g_1}{g_0} n_0 (1 - e^{-T_*/T_s}) \phi(\nu) \approx \frac{3\hbar c^2 A_{10} n_{HI} \nu_0}{16k_B T_s \nu^2} \phi(\nu). \quad (1.22)$$

In the last expression, n_{HI} is the neutral hydrogen density and there is an extra factor of 1/4 in equation 1.22 to account for the fraction of HI atoms in the hyperfine

singlet state ($n_0 \approx n_{HI}/4$). In order to estimate the optical depth, $\tau(\nu)$, as a function of observed frequency, due to absorption by the IGM along the line of sight to a given redshift, we must integrate $\alpha(\nu)$ along the radiation path length, i.e. from the radiation source at z_{em} to $z = 0$.

If $\alpha(\nu)$ is the absorption coefficient at frequency ν , then the increment of optical depth for photons of this frequency which are redshifted to frequency ν_{obs} when they reach the earth is $d\tau(\nu_{obs}) = \alpha(\nu)dl$, where $dl = cdt$ is the element of proper distance at redshift z and dt is the time interval that radiation takes to travel the path length dl . Then

$$\tau(\nu_{obs}) = \int_0^l dl\alpha(\nu) = \int_0^{z_{em}} \frac{dz c\alpha(\nu)}{H(z)(1+z)} \rightarrow \tau(z) = \int_{\nu_{obs}}^{\nu_{obs}(1+z_{em})} \frac{c\alpha(\nu)}{\nu H[z(\nu)]} d\nu \quad (1.23)$$

where we used the fact that $\nu = \nu_{obs}(1+z)$ so $dz = d\nu/\nu_{obs}$. A simple, yet conceptually important model for $\phi(\nu)$ is to ignore the broadening of the line profile and assume that it can be approximated by a delta function, $\phi(\nu) = \delta(\nu - \nu_0)$, so⁴

$$\begin{aligned} \tau(z) &= \frac{c\alpha(\nu_0)}{\nu_0 H(z)} = \frac{3\hbar c^3 A_{10} n_{HI}(z)}{16k_B \nu_0^2 T_s H(z)} \\ &\approx 6 \cdot 10^{-3} (1 + \delta_b) x_{HI} \left[\frac{T_{CMB}(z)}{T_s} \right] \left(\frac{\Omega_b h^2}{0.022} \right) \left[\frac{\Omega_m (1+z)^3 + \Omega_\Lambda}{5.5} \right]^{-1/2} \\ &\quad \left(\frac{1+z}{2.5} \right)^2 \left(\frac{h}{0.68} \right)^{-1} \end{aligned} \quad (1.24)$$

where $\nu_0 = \nu_{obs}(1+z)$ and we have used $T_{CMB}(z) = 2.73(1+z)$ K and $H(z) \approx H_0 [\Omega_m(1+z)^3 + \Omega_\Lambda]^{1/2}$. The neutral hydrogen density was written as [23, 24] $n_{HI} = x_{HI} \bar{n}_b(1 + \delta_b)$, where $\bar{n}_b = \bar{n}_{b0}(1+z)^3$ is the mean number density of cosmic baryons, with a local baryon overdensity δ_b , and x_{HI} is the fraction of neutral hydrogen. For a more exact calculation of $\tau(z)$ that considers both Hubble expansion and peculiar velocities refer to [21, 25].

Now, in the Rayleigh-Jeans limit, the equation of transfer along the line of sight for a background radiation field with brightness temperature T_{bkg} and a cloud with uniform excitation temperature T_s is (see [22])

$$T'_b(\nu) = T_{bkg}(\nu)e^{-\tau} + T_s(1 - e^{-\tau}). \quad (1.25)$$

⁴Used cosmological parameter values from [10].

Assuming that the radiation background includes only the CMB, then $T_{bkg} = T_{CMB}$ and the differential brightness temperature emergent from the cloud is

$$\Delta T'_b(\nu) = T'_b - T_{CMB} = (T_s - T_{CMB})(1 - e^{-\tau}) \approx (T_s - T_{CMB})\tau \quad (1.26)$$

so the observed differential brightness temperature is

$$\begin{aligned} \Delta T_b(\nu) &\approx \frac{[T_s - T_{CMB}(z)]}{1+z} \tau(z) \\ &\approx 16 (1 + \delta_b) x_{HI} \left[1 - \frac{T_{CMB}(z)}{T_s} \right] \left(\frac{\Omega_b h^2}{0.022} \right) \left[\frac{\Omega_m (1+z)^3 + \Omega_\Lambda}{5.5} \right]^{-1/2} \\ &\quad \left(\frac{1+z}{2.5} \right)^2 \left(\frac{h}{0.68} \right)^{-1} \text{ mK} \end{aligned} \quad (1.27)$$

where in the second step we have used equation 1.24. Note that the 21 cm signal ΔT_b traces the density fluctuations δ_b , making its measurement an excellent probe of cosmology.

Three processes determine the spin temperature [21, 26]: (1) absorption of CMB photons (as well as stimulated emissions); (2) collisions with other hydrogen atoms, electrons and protons; and (3) scattering of Lyman α photons through excitation and de-excitation (Wouthuysen Field Effect). The rate of these processes is fast compared to the de-excitation time of the line, so that to a very good approximation the spin temperature is given by the equilibrium balance of these effects [27]. In this case T_s is a weighted average of the CMB temperature, T_{CMB} , the gas kinetic temperature, T_K , and the effective color temperature of the Ly α radiation field, T_L

$$T_s = \frac{T_{CMB} + y_c T_K + y_L T_L}{1 + y_c + y_L} \quad (1.28)$$

where y_c , y_L are the efficiencies or coupling coefficients for gas collisions and Ly α scattering (we note that $T_L = T_K$ as long as the medium is optically thick to Ly α photons). For a discussion and detailed calculation of the efficiencies see [21, 25, 26, 27].

The evolution of T_s depends on how different mechanisms influence the efficiencies. When any of the efficiencies is very large T_s takes the corresponding temperature value. As explained in [25], after reionization T_s is coupled to T_K , $T_s \gg T_{CMB}$ so ΔT_b is independent of T_s and the 21 cm radiation is seen in emission. In this case the differential

brightness temperature simplifies to

$$\Delta T_b \approx \overline{\Delta T_b}(1 + \delta_b)(1 + \delta_x) \quad (1.29)$$

where $\overline{\Delta T_b}$ is the average differential brightness temperature that depends on redshift,

$$\begin{aligned} \overline{\Delta T_b} \approx & 16 \bar{x}_{HI}(z) \left(\frac{\Omega_b h^2}{0.022} \right) \left[\frac{\Omega_m(1+z)^3 + \Omega_\Lambda}{5.5} \right]^{-1/2} \\ & \left(\frac{1+z}{2.5} \right)^2 \left(\frac{h}{0.68} \right)^{-1} \text{ mK} \end{aligned} \quad (1.30)$$

and we have written $x_{HI} = \bar{x}_{HI}(1 + \delta_x)$, with \bar{x}_{HI} as the average of the neutral fraction of hydrogen and δ_x as the neutral fraction fluctuations.

1.3.2 Power spectrum

To first order in the fluctuations, we can write

$$\Delta T_b(\mathbf{r}) \approx \overline{\Delta T_b}[1 + \delta_b(\mathbf{r}) + \delta_x(\mathbf{r})] = \overline{\Delta T_b}[1 + \delta_{21}(\mathbf{r})] \quad (1.31)$$

where $r = |\mathbf{r}|$ is the radial distance corresponding to the observed frequency (redshift). In general, the variations in the brightness temperature depend on additional parameters, including variations in the Ly α coupling coefficient, kinetic temperature and line-of-sight peculiar velocity gradient. However, if we ignore the velocity term then equation 1.31 is a good approximation for sufficiently low redshifts ($z \lesssim 10$). See [21] for details.

$\overline{\Delta T_b}$ is the zeroth order approximation of the 21 cm signal, averaged over large angular scales at fixed radial distance. In this section we are interested in the 21 cm fluctuations, which are the ones that allow for the measurement of the BAO. The deviation from the mean 21 cm brightness temperature is given by

$$\delta_{21}(\mathbf{r}) = \frac{\Delta T_b - \overline{\Delta T_b}}{\overline{\Delta T_b}} \quad (1.32)$$

which is a zero mean random field. Since each observed frequency corresponds to a different radial shell of the universe, it is convenient to express $\delta_{21}(\mathbf{r})$ in terms of its

spherical harmonic expansion

$$\delta_{21}(\mathbf{r}) = \delta_{21}(\hat{\mathbf{n}}, \nu) = \sum_{l,m} a_{lm}(\nu) Y_{lm}(\hat{\mathbf{n}}) \quad a_{lm}(\nu) = \int_{4\pi} d\Omega \delta_{21}(\hat{\mathbf{n}}, \nu) Y_{lm}^*(\hat{\mathbf{n}}) \quad (1.33)$$

Note that, except for $a_{00}(\nu)$, the spherical harmonic expansion of $\Delta T_b(\mathbf{r})$ is the same as that of $\delta_{21}(\mathbf{r})$ with the additional prefactor $\overline{\Delta T_b}(\nu)$. If we now write $\delta_{21}(\mathbf{r})$ in equation 1.33 in terms of its Fourier Transform

$$\delta_{21}(\mathbf{r}) = \int_{-\infty}^{\infty} \frac{d^3 \mathbf{k}}{(2\pi)^3} \tilde{\delta}_{21}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{r}} \quad \tilde{\delta}_{21}(\mathbf{k}) = \int_{-\infty}^{\infty} d^3 \mathbf{r} \delta_{21}(\mathbf{r}) e^{-i\mathbf{k} \cdot \mathbf{r}}. \quad (1.34)$$

and use the Rayleigh expansion for the $e^{i\mathbf{k} \cdot \mathbf{r}}$ term [28], we obtain

$$a_{lm}(\nu) = 4\pi i^l \int_{-\infty}^{\infty} \frac{d^3 \mathbf{k}}{(2\pi)^3} \tilde{\delta}_{21}(\mathbf{k}) j_l(kr) Y_{lm}^*(\hat{\mathbf{k}}) \quad (1.35)$$

where j_l is the spherical Bessel function of order l .⁵

As shown in Appendix A, for the homogeneous and isotropic random field $\delta_{21}(\mathbf{r})$ its two point correlation function depends on $r = |\mathbf{r}_1 - \mathbf{r}_2|$ only, that is, $\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle = \xi_{21}(r)$. Similarly, its power spectrum $P_{21}(k)$ (the Fourier Transform of $\xi_{21}(r)$) is uniquely defined through the equation

$$\langle \tilde{\delta}_{21}(\mathbf{k}) \tilde{\delta}_{21}^*(\mathbf{k}_1) \rangle = (2\pi)^3 P_{21}(k) \delta(\mathbf{k} - \mathbf{k}_1) \quad (1.36)$$

where $k = |\mathbf{k}|$.

In an analogous way as we defined the power spectrum of $\delta_{21}(\mathbf{r})$ in terms of the correlation of the $\tilde{\delta}_{21}(\mathbf{k})$, we can define its *angular power spectrum*, C_l , in terms of the correlation of the $a_{lm}(\nu)$. Using equations 1.35-1.36 and the orthogonality of the spherical harmonics we obtain

$$\langle a_{lm}(\nu) a_{l_1 m_1}^*(\nu_1) \rangle = \delta_{ll_1} \delta_{mm_1} C_l(\nu, \nu_1) \quad (1.37)$$

where we have defined (with r as the radial distance corresponding to the observed

⁵For details see <http://mathworld.wolfram.com/SphericalBesselFunctionoftheFirstKind.html>.

frequency ν)

$$C_l(\nu, \nu_1) = 4\pi \int_0^\infty \frac{k^2 dk}{2\pi^2} P_{21}(k) j_l(kr) j_l(kr_1). \quad (1.38)$$

For a more detailed calculation that takes into account the finite bandwidth of the experiment refer to [21, 23, 24].

In general, equation 1.38 forms the basis for the analysis that makes it possible to relate the 21 cm fluctuations to cosmology. As explained in [29, 30], after reionization the neutral fraction fluctuations δ_x are proportional to the matter density fluctuations up to a weak (and simple) scale dependency. Since the 21 cm fluctuations allow the direct observation of the matter density field δ_b , the power spectrum can be utilized to probe the geometry of the universe, derive cosmological constraints, and test the equation of state of dark energy by measuring the BAO. In particular, redshifts in the range $z \approx 1 - 3$ are of great interest since they cover the regime in which dark energy begins to dominate the energy budget of the universe [27].

CHIME will be a novel radio telescope that will study the nature of dark energy by mapping the three-dimensional distribution of neutral hydrogen gas in the universe detected from redshifted 21 cm radiation. By measuring the BAO scale across the redshift range $z \approx 0.8 - 2.5$ in both the angular and line-of-sight directions, CHIME will study the epoch when dark energy generated the transition from decelerated to accelerated expansion of the universe.

1.4 Outline

This thesis is structured as follows: Chapter 2 gives an introduction to radio interferometry and 21 cm cosmology with transit interferometers. Chapter 3 describes the CHIME telescope design and characteristics. Chapter 4 describes the CHIME correlator. Chapter 5 discusses the CHIME calibration requirements and presents the calibration techniques that we have developed for CHIME. In Chapter 6 we investigate the biasing effect of quantization in the magnitude and phase of the measured visibilities and give details of the post-channelization digital calibration. Chapter 7 presents the conclusions.

Chapter 2

21 cm cosmology with transit interferometers

CHIME is a transit interferometer, meaning that it has a fixed pointing and observes the sky as the earth rotates. In this chapter we present a detailed analysis of interferometry and introduce several important concepts that will be used throughout this thesis. We also explain how transit interferometers like CHIME map the sky and the 21 cm universe.

2.1 Stationary random processes

The signals from radio astronomical sources are mostly generated by chaotic processes that occur in the source regions: the synchrotron radiation in active galactic nuclei (AGN) caused by the relativistic acceleration of electrons, the thermal radiation in planetary nebulae generated by the thermal motion of charged particles, or 21 cm radiation from clouds of neutral hydrogen due to the transition between the hyperfine energy levels (this radiation is naturally narrow in frequency, but becomes broadband due to, for example, the large scale motion of the gas clouds and the expansion of the universe). In all these cases, the electromagnetic fields incident on an interferometric array as well as the currents and voltages induced in the antennas and receiving systems will look like random noise, so they are properly described as random (stochastic) processes. In almost all cases this noise is Gaussian, i.e. it has a normal probability density function (PDF) and the width of this distribution depends on the brightness of the source. The noise produced in the receiving systems are also well modeled as Gaussian random processes.

Furthermore, when we observe the received voltages from radio astronomical sources, say on an oscilloscope, we usually see that these voltages look like time series that fluctuate around zero, and their time average is also zero, so the random processes are *time*

invariant. In a radio interferometer, the correlator is a device that basically computes time averaged products of pairs of these received voltages. When we check the outputs of the correlator, we usually see that they are approximately constant in time, but that they vary when we introduce relative delays between the received voltages.

All these observations lead us to assume that the signals received from radio astronomical sources can be modeled as random processes that are *stationary* (in the wide sense), meaning that their statistical properties (mean, variance, PDF) do not change with time, and that the cross-correlations are functions of the time difference only, not absolute time. Also from observations (for example, we see that the cross-correlation estimates improve with averaging and that they tend to zero as we increase the relative time delay between received voltages¹), we know that these signals satisfy *ergodicity*, meaning that the time averages (which is what we measure in the correlator) converge to the respective ensemble averages.

For a detailed introduction to the theory of stationary random processes and ergodicity, which we will use intensively in this work, refer for example to [31]. Note that the physical processes that generate the signals we observe in radio astronomy do not satisfy strictly the conditions for stationarity and ergodicity. For example, radio sources can vary significantly in the scale of several years or months (or even shorter time scales). The statistical properties of the measured signals also vary in the scale of hours or minutes due to changes in the response of the receiving systems or due to the rotation of the earth on transit telescopes like CHIME. However, in practice, the conditions for stationarity and ergodicity are generally fulfilled on short observation intervals (at least several seconds, which is the typical integration scale for telescopes like CHIME). This is enough to estimate the statistical properties of the measured signals. Thus, we will assume that the processes are at least quasi-stationary, meaning that the statistical properties may vary over long time scales, while over short time scales these properties are constant.

2.2 The narrowband condition

The cosmic signals that CHIME will observe (e.g. the 21 cm signal and synchrotron emission from our galaxy and other galaxies) are typically broadband in nature. As explained in detail in Chapters 3 and 4, the analog receiving system determines the total bandwidth of interest (400-800 MHz for CHIME) while the digital correlator digi-

¹Loosely speaking, a random process with a correlation function that asymptotically decays to zero satisfies mean ergodicity. Technically, an absolutely integrable autocovariance is a sufficient condition for wide-sense stationary random process to be ergodic in the mean. Correlation ergodicity has additional requirements and we will assume throughout this thesis that those requirements are also satisfied.

tizes the broadband signal at this full bandwidth and then divides it into narrow frequency channels using digital filters. The CHIME polyphase filter bank (PFB) divides the 400 MHz input bandwidth into 1024 contiguous frequency channels which are $\Delta\nu = 400 \text{ MHz}/1024 = 390.625 \text{ kHz}$ wide and that are re-sampled at the appropriate lower rate before performing the correlations between all the antenna pairs. Since both the analog receiver and the digital backend are linear systems (so frequencies are not mixed) then, from the point of view of analysis, the output of each digital filter can be considered as the output of a separate and independent receiving system of bandwidth $\Delta\nu$. Furthermore, since we already know that each frequency corresponds to a different radial slice of the universe when we map with the 21 cm signal, then the output of each digital filter can be effectively considered as the output of a separate experiment that observes a different time in the history of the universe.

It is often convenient for modeling and analysis to assume that the sky signals can be described as a sum of narrowband or quasi-monochromatic signals (or equivalently, that the output of each filter bank can be considered as quasi-monochromatic), so we must determine the conditions under which that assumption is valid. If a real-valued signal $v(t)$ is band-limited to a frequency range $\Delta\nu$ centered at a frequency ν_c with $\Delta\nu \ll \nu_c$ (as is the case for CHIME), then it can be shown (see [32] for details) that $v(t)$ can be written as

$$v(t) = \Re\{x(t)e^{i2\pi\nu_c t}\} \quad (2.1)$$

where $x(t)$ is basically the equivalent lowpass signal of $v(t)$ and is called the *complex envelope*. Sometimes it is easier to manipulate the complex envelope $x(t)$ than $v(t)$. For example, when the bandwidth $\Delta\nu$ is sufficiently narrow then time delays of $v(t)$ across the telescope array can be approximated by phase shifts of its complex envelope. This approximation is used for beamforming, fringe-stopping, and image reconstruction techniques. To determine this narrowband approximation, suppose that $v(t)$ is delayed in time by an amount τ . Then

$$v(t - \tau) = \Re\{x(t - \tau)e^{i2\pi\nu_c t}e^{-i2\pi\nu_c \tau}\}. \quad (2.2)$$

Let $\tilde{x}(\nu)$ be the Fourier transform of $x(t)$. Then

$$x(t) = \int_{-\Delta\nu/2}^{\Delta\nu/2} \tilde{x}(\nu)e^{i2\pi\nu t}d\nu \quad (2.3)$$

while

$$x(t - \tau) = \int_{-\Delta\nu/2}^{\Delta\nu/2} \tilde{x}(\nu) e^{i2\pi\nu t} e^{-i2\pi\nu\tau} d\nu. \quad (2.4)$$

If the condition

$$2\pi\Delta\nu\tau \ll 1 \quad (2.5)$$

holds, then $e^{-i2\pi\nu\tau} \approx 1$ for ν in the interval $[-\Delta\nu/2, \Delta\nu/2]$ and the integrals in equations 2.3 and 2.4 are approximately equal. Thus $x(t - \tau) \approx x(t)$ and equation 2.2 can be written as

$$v(t - \tau) \approx \Re\{x_\tau(t) e^{i2\pi\nu_c t}\} \quad (2.6)$$

where

$$x_\tau(t) = x(t) e^{-i2\pi\nu_c \tau} \quad (2.7)$$

which means that, when the narrowband condition in equation 2.5 holds, a delay of $v(t)$ can be expressed as a phase shift of the complex envelope $x(t)$. This condition is independent of ν_c .

As an example, the narrowband condition for CHIME means that the maximum geometric delay allowed is $\tau_{max} \ll 1/2\pi\Delta\nu \sim 410$ ns. The baselines (antenna separations) for the CHIME telescope range from ~ 0.3 m to ~ 100 m, and therefore the maximum possible geometric delay is $\tau_{max} = b_{max}/c \approx 333$ ns. This means that, although it is safe to use the narrowband approximation for most baselines, the width of the CHIME frequency channels is not narrow enough to satisfy this condition. As shown in Section 2.3.2, a non-zero bandwidth introduces a signal decorrelation and thus, a loss of signal in the visibility data which for the longest baselines and maximum delays is about 3%. In practice, this effect is smaller since the maximum delay corresponds to signals arriving from the horizon which is not within the field of view of the telescope. Also, this decorrelation effect can be corrected by implementing artificial delays, for example within the digital correlator, that compensate for the geometric delay.

2.3 The two-element interferometer

The basic radio interferometer is the two-element interferometer shown in figure 2.1. Larger and more complex interferometric arrays can be treated as collections of indepen-

dent two-element interferometers.

A two-element interferometer consists of two antennas separated by a distance b . The vector connecting the reference points of the two antennas is called the baseline vector \mathbf{b} . The antennas are observing a (far field) point source in the direction of the unit vector $\hat{\mathbf{n}}$. From figure 2.1, the wavefront of the electromagnetic wave from the source reaches antenna 2 first and then it reaches antenna 1 after a time

$$\tau_g = \frac{\mathbf{b} \cdot \hat{\mathbf{n}}}{c} = \frac{b \sin \theta}{c} \quad (2.8)$$

where c is the velocity of light and θ is the angle between $\hat{\mathbf{n}}$ and a straight line perpendicular to \mathbf{b} in the plane formed by $\hat{\mathbf{n}}$ and \mathbf{b} . τ_g is called the *geometric time delay*.

The antennas have primary beams $A_1(\hat{\mathbf{n}})$ and $A_2(\hat{\mathbf{n}})$ respectively. Each antenna has a receiver that conditions (amplifies and filters) the antenna signal before sending it to the correlator which is the device that computes the correlation (multiplication and averaging, see equations 2.9 and 2.10 below) between the received voltages. This correlation provides the interferometric visibility which is the fundamental quantity in radio interferometry.

For two inputs $v_1(t)$ and $v_2(t)$, the output of the correlator after averaging during a period T is

$$\hat{r}_{12} = \frac{1}{T} \int_{-T/2}^{T/2} v_1(t) v_2^*(t) dt. \quad (2.9)$$

In a digital correlator, the inputs are first digitized (sampled and quantized) before multiplication and averaging. In this case, the output of the correlator after integrating N_s samples is

$$\hat{r}_{12} = \frac{1}{N_s} \sum_{n=1}^{N_s} v_1[n] v_2^*[n]. \quad (2.10)$$

In both cases, and in view of ergodicity (time averages being equal to ensemble averages), the estimate \hat{r}_{12} becomes a more accurate measurement of the true correlation

$$r_{12} = \langle v_1(t) v_2^*(t) \rangle \quad (2.11)$$

as the averaging increases, i.e.

$$\lim_{N_s \text{ (or } T) \rightarrow \infty} \hat{r}_{12} = r_{12}. \quad (2.12)$$

We will often assume that the averaging is sufficiently long that we can safely ap-




Figure 2.1: A basic two-element interferometer consists of two antennas separated by a distance b . The vector connecting the reference points of the two antennas is called the baseline vector \mathbf{b} . The antennas are observing a (far field) point source in the direction of the unit vector $\hat{\mathbf{n}}$. The wavefront of the radio source reaches antenna 2 first and then it reaches antenna 1 after a time τ_g called the geometric delay. The antennas have primary beams $A_1(\hat{\mathbf{n}})$ and $A_2(\hat{\mathbf{n}})$ respectively. Each antenna has a receiver that conditions (amplifies and filters) the antenna signal before sending it to the correlator which is the device that computes the correlation (multiplication and averaging) between the received voltages.

proximate the correlator output by r_{12} , except in cases where we want to investigate the effect of finite averaging (e.g. Section 2.3.4).

In some cases we will also assume that the receivers have narrow bandpass filters that pass only frequency components very close to ν . As explained in Section 2.2, in modern digital correlators, the broadband inputs are digitized at full bandwidth and then divided into narrow frequency channels using digital filters, and each filter can be considered as an independent receiver. The outputs of these filters are complex-valued in general.

2.3.1 The response of the receiving system

Let $v_E(t)$ be the voltage generated at the output of an antenna by the incident electric field of single polarization of the radio wave from a point source. Both the electric field and $v_E(t)$ are stationary random processes. Then, from Figure 2.2, the antenna voltages of the two-element interferometer (assuming identical antennas) are $v_{A1}(t) = v_E(t - \tau_g)$ and $v_{A2}(t) = v_E(t)$, where τ_g is the geometric time delay (equation 2.8)².

Let $h_1(t)$ and $h_2(t)$ be the impulse response of the receivers 1 and 2 respectively. Since the receivers are linear (and time-invariant) systems, then the respective outputs for the inputs $v_{A1}(t)$ and $v_{A2}(t)$, which are jointly stationary random processes, are

$$\begin{aligned} v_1(t) &= v_{A1}(t) * h_1(t) = v_E(t - \tau_g) * h_1(t) \text{ and} \\ v_2(t) &= v_{A2}(t) * h_2(t) = v_E(t) * h_2(t). \end{aligned} \quad (2.13)$$

Note that the operator $*$ in the last equation denotes convolution (we use the same symbol as a superscript to denote complex conjugate of a variable, but the meaning will be clear from the context).

The cross-correlation function of the two receiver outputs $v_1(t)$ and $v_2(t)$ is (see [31] for a detailed discussion of stationary random processes in linear systems)

$$r_{v_1 v_2}(\tau) = \langle v_1(t) v_2^*(t - \tau) \rangle = r_{A1 A2}(\tau) * h_1(\tau) * h_2^*(-\tau) \quad (2.14)$$

where $r_{A1 A2}(\tau) = \langle v_{A1}(t) v_{A2}^*(t - \tau) \rangle$ is the cross-correlation function of the antenna outputs $v_{A1}(t)$ and $v_{A2}(t)$.

The frequency domain equivalent of equation 2.14 is

$$\tilde{r}_{v_1 v_2}(\nu) = \tilde{r}_{A1 A2}(\nu) \tilde{h}_1(\nu) \tilde{h}_2^*(\nu) \quad (2.15)$$

²In this section we ignore the effects of receiver noise.




Figure 2.2: Response of the receiving system of a two-element interferometer to the voltage generated by the incident electric field from a point source. $v_{A1}(t)$ and $v_{A2}(t)$ are the antenna output voltages. The receivers have impulse response $h_1(t)$ and $h_2(t)$ and outputs $v_1(t)$ and $v_2(t)$ respectively. These are the inputs to the correlator which computes the time averaged product of $v_1(t)$ and $v_2(t)$.

where $\tilde{r}_{A1A2}(\nu)$ and $\tilde{r}_{v1v2}(\nu)$ are the Fourier transforms of $r_{v1v2}(\tau)$ and $r_{A1A2}(\tau)$ respectively, while $\tilde{h}_1(\nu)$ and $\tilde{h}_2(\nu)$ are the Fourier transforms of $h_1(t)$ and $h_2(t)$ (i.e. the frequency response of the receivers).

Assuming that the averaging is sufficiently long, the output of the correlator is

$$r_{12} = \langle v_1(t)v_2^*(t) \rangle = r_{v1v2}(0) = \int_{-\infty}^{\infty} \tilde{r}_{A1A2}(\nu) \tilde{h}_1(\nu) \tilde{h}_2^*(\nu) d\nu. \quad (2.16)$$

In the last step of equation 2.16 we used $r_{v1v2}(\tau) = \int_{-\infty}^{\infty} \tilde{r}_{v1v2}(\nu) e^{i2\pi\nu\tau} d\nu$.

Now we have to find the relation between r_{12} , which is what the correlator measures, and the power spectrum $\tilde{r}_{EE}(\nu)$ of the voltage $v_E(t)$ generated by the common electric field, which is what we want to measure, and that by the Wiener-Khinchin theorem³ is just the Fourier transform of $r_{EE}(\tau)$, that is

$$r_{EE}(\tau) = \langle v_E(t)v_E^*(t-\tau) \rangle = \int_{-\infty}^{\infty} \tilde{r}_{EE}(\nu) e^{i2\pi\nu\tau} d\nu. \quad (2.17)$$

Note that since

$$r_{A1A2}(\tau) = \langle v_{A1}(t)v_{A2}^*(t-\tau) \rangle = \langle v_E(t-\tau_g)v_E^*(t-\tau) \rangle = r_{EE}(\tau - \tau_g) \quad (2.18)$$

then, by taking Fourier transforms we have

$$\tilde{r}_{A1A2}(\nu) = \tilde{r}_{EE}(\nu) e^{-i2\pi\nu\tau_g}. \quad (2.19)$$

Finally, using equation 2.19 in equation 2.16 we obtain

$$r_{12} = \int_{-\infty}^{\infty} \tilde{r}_{EE}(\nu) \tilde{h}_1(\nu) \tilde{h}_2^*(\nu) e^{-i2\pi\nu\tau_g} d\nu. \quad (2.20)$$

In the particular case of receivers with identical frequency response $\tilde{h}(\nu)$ then equation 2.21 simplifies to

$$r_{12} = \int_{-\infty}^{\infty} \tilde{r}_{EE}(\nu) |\tilde{h}(\nu)|^2 e^{-i2\pi\nu\tau_g} d\nu. \quad (2.21)$$

2.3.2 Interferometric fringes

The receiving system can be modeled as an amplification stage that increases the power of the antenna output signal (which is typically very weak) followed by a filtering stage that

³For details see https://en.wikipedia.org/wiki/Wiener-Khinchin_theorem.

selects the bandwidth of interest. The simplest (but conceptually important) example is to assume that the frequency response of the receiver, $|\tilde{h}(\nu)|^2$, is just that of a rectangular bandpass filter with center frequency ν_c , bandwidth $\Delta\nu$, and unit in-band (power) gain, that is

$$|\tilde{h}(\nu)|^2 = \begin{cases} 1 & \text{if } \nu_c - \frac{\Delta\nu}{2} \leq |\nu| \leq \nu_c + \frac{\Delta\nu}{2} \\ 0 & \text{otherwise.} \end{cases} \quad (2.22)$$

For a broadband radio source, and if the bandwidth $\Delta\nu$ is sufficiently narrow, then it is usually safe to assume that the source power spectrum $\tilde{r}_{EE}(\nu)$ is approximately flat, or white, in the filter passband. For simplicity, let $\tilde{r}_{EE}(\nu) = R_c$ (only its value inside the receiver passband is important). Then using equation 2.22 in equation 2.21 we obtain⁴

$$\begin{aligned} r_{12}(\tau_g) &= 2R_c \Re \left\{ \int_{\nu_c - \frac{\Delta\nu}{2}}^{\nu_c + \frac{\Delta\nu}{2}} e^{-i2\pi\nu\tau_g} d\nu \right\} \\ &= 2R_c \Delta\nu \frac{\sin(\pi\Delta\nu\tau_g)}{\pi\Delta\nu\tau_g} \cos(2\pi\nu_c\tau_g). \end{aligned} \quad (2.23)$$

An example of the geometric delay dependent response of the interferometer to a band limited signal is shown in Figure 2.3. The bandwidth is $\Delta\nu$, the center frequency is ν_c and we have chosen $\nu_c = 10.2\Delta\nu$. The solid blue line shows the normalized response $r_{12}(\tau_g)/r_{12}(0)$. This is the *fringe pattern* that is observed for example, when a point source transits through the field of view of an east-west interferometer, with τ_g tracking the change in geometric delay as the source moves across the sky.

From equation 2.23, the fringes are due to the term $\cos(2\pi\nu_c\tau_g)$, so the fringe spacing is just

$$\Delta\tau_{g,F} = \frac{1}{\nu_c}. \quad (2.24)$$

To find the angular spacing $\Delta\theta_F$ corresponding to the fringe spacing $\Delta\tau_{g,F}$ note that, since the primary beam of CHIME telescope is very narrow in the east-west direction ($\theta_{EW} \sim 2^\circ$), then τ_g (and θ) is small, so equation 2.8 becomes

$$\tau_g \approx \frac{b\theta}{c}. \quad (2.25)$$

⁴Note that we take the real part only and multiply by 2 to normalize for total power since both $\tilde{r}_{EE}(\nu)$ and $|\tilde{h}(\nu)|^2$ are even functions of ν

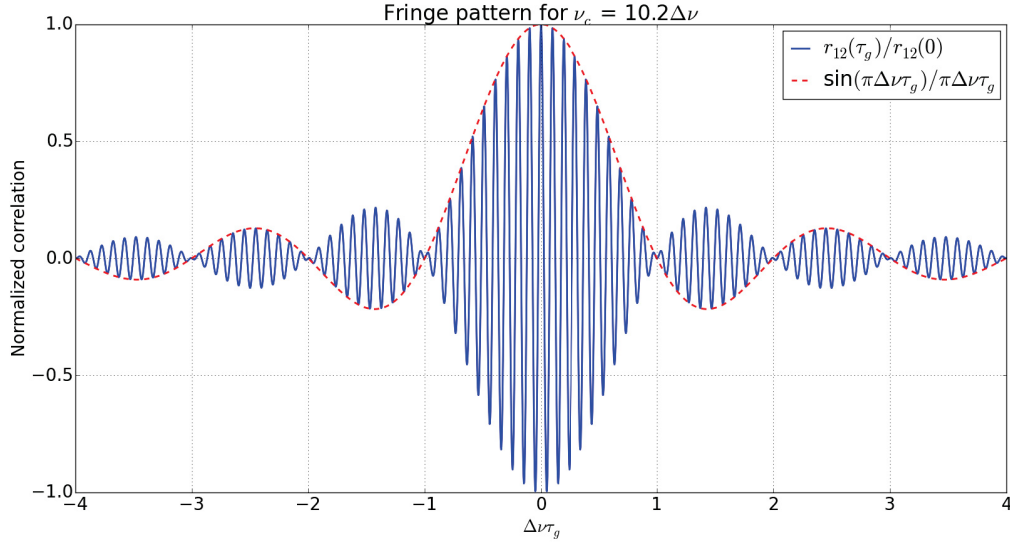


Figure 2.3: Normalized response of the two-element interferometer to a band limited signal with bandwidth $\Delta\nu$ and center frequency ν_c , with $\nu_c = 10.2\Delta\nu$. The geometric delay is τ_g . While the rapid oscillation (blue) is determined by ν_c , the envelope of the fringe pattern (red) is determined by $\Delta\nu$.

Thus, the angular spacing corresponding to the fringe spacing is

$$\Delta\theta_F \approx 2 \left(\frac{c}{b} \right) \left(\frac{\Delta\tau_{g,F}}{2} \right) = \frac{\lambda_c}{b}. \quad (2.26)$$

As explained in Appendix C, the resolution of an interferometer is also given by equation 2.26. For CHIME the fringe spacing at 600 MHz ($\lambda_c = 0.5$ m) is $\sim 1.4^\circ$ for the shortest east-west baselines (20 m) and $\sim 0.5^\circ$ for the longest east-west baselines (60 m).

The fringe pattern is enclosed by an envelope (dashed red line in figure 2.3) which is a sinc function for a rectangular bandpass, but that can change for other band shapes. However, in all cases the envelope is determined by the bandwidth $\Delta\nu$, so the envelope is usually called the *bandwidth pattern*. The width of the main lobe of the bandwidth pattern is called the *coherence interval*, and is given by $\Delta\tau_{g,B} = 2/\Delta\nu$. The angular extent $\Delta\theta_B$ corresponding to the coherence interval is

$$\Delta\theta_B \approx \frac{2c}{b\Delta\nu} \quad (2.27)$$

If the telescopes were isotropic then the coherence interval would determine field of view of the interferometer, basically because outside this interval the signal level is very small. Although $\Delta\theta_B > 180^\circ$ for CHIME, in practice, the east-west field of view of CHIME is much narrower and is determined by the antenna beam (see Chapter 3).

The bandwidth pattern demonstrates the effect of a non-zero bandwidth on the measured signal (see Section 2.2). When the narrowband condition holds from equation 2.23, we have $r_{12}(\tau_g) \rightarrow 2R_c \cos(2\pi\nu_c\tau_g)$, which is the expression we would have obtained if we had assumed monochromatic signals. A non-zero bandwidth introduces a decorrelation term $\text{sinc}(\pi\Delta\nu\tau_g)$ corresponding to a loss of signal. For CHIME, this decorrelation term is of the order of one part in 10^7 for the shortest baselines, and about 3% for the longest baselines and maximum delays.

2.3.3 The response of the interferometer: the complex visibility

In this section we will find the response of the interferometer to an arbitrary extended source. Referring to figure 2.1, let $e(\hat{\mathbf{n}}, t)$ be the electric field intensity received by antenna 2 from a source region of solid angle $d\Omega$, in the direction of the unit vector $\hat{\mathbf{n}}$, and at time t (see Appendix B). Let $v_{A2}(t)$ be the output voltage of antenna 2. The Fourier transforms of $e(\hat{\mathbf{n}}, t)$ and $v_{A2}(t)$ are given by⁵

$$\tilde{e}(\hat{\mathbf{n}}, \nu) = \int_{-\infty}^{\infty} e(\hat{\mathbf{n}}, t) e^{-i2\pi\nu t} dt, \quad \tilde{v}_{A2}(\nu) = \int_{-\infty}^{\infty} v_{A2}(t) e^{-i2\pi\nu t} dt. \quad (2.28)$$

At any given frequency ν , the antenna voltage is a weighted combination of the electric field coming from each direction of the sky, with the weighting function being the antenna beam $A(\hat{\mathbf{n}})$, so we can write⁶

$$\tilde{v}_{A2}(\nu) = \int_{4\pi} \tilde{e}(\hat{\mathbf{n}}, \nu) A_2(\hat{\mathbf{n}}) d\Omega. \quad (2.29)$$

Note that we have extended the range of the integration to the whole sky since $\tilde{e}(\hat{\mathbf{n}}, \nu) = 0$ beyond the edge of the source and also because $A_2(\hat{\mathbf{n}}) = 0$ far from the antenna beam.

The electric field received by antenna 1 from the same region is $e(\hat{\mathbf{n}}, t - \tau_g(\hat{\mathbf{n}}))$, where $\tau_g(\hat{\mathbf{n}}) = \mathbf{b} \cdot \hat{\mathbf{n}}/c$ is the geometric delay. In frequency domain this field is $\tilde{e}(\hat{\mathbf{n}}, \nu) e^{-i2\pi\nu\tau_g(\hat{\mathbf{n}})}$,

⁵Note that $\tilde{e}(\hat{\mathbf{n}}, \nu)$ and $\tilde{v}_{A2}(\nu)$ are Fourier transforms of stationary random processes and, strictly speaking, the integrals in equation 2.28 do not necessarily converge, and ensuring convergence may cause problems when requiring stationarity of the random processes. Since our goal is to find a relation between the stationary random process $e(\hat{\mathbf{n}}, t)$ and the response of the interferometer, we will assume that these integrals converge without getting heavily involved in the mathematical strictness. For a discussion of this topic see for example [33].

⁶Note that the antenna beam is also frequency dependent, although we do not write that dependence explicitly in the equations not only to keep the notation simple, but also because this seems to be the common notation in radio astronomy textbooks.

so

$$\tilde{v}_{A1}(\nu) = \int_{4\pi} \tilde{e}(\hat{\mathbf{n}}, \nu) A_1(\hat{\mathbf{n}}) e^{-i2\pi\nu\tau_g(\hat{\mathbf{n}})} d\Omega. \quad (2.30)$$

In Section 2.3.1 we introduced $\tilde{r}_{A1A2}(\nu)$, which is the Fourier transform of the cross-correlation function of the antenna outputs $r_{A1A2}(\tau) = \langle v_{A1}(t) v_{A2}^*(t - \tau) \rangle$. Now we want to express $\tilde{r}_{A1A2}(\nu)$ in terms of the intensity (or brightness) distribution of the source, I_ν .

First, note that $v_{A1}(t)$ and $v_{A2}(t)$ are jointly stationary random processes which means that their cross-correlation $r_{A1A2}(\tau)$ is a function of the time difference only. This also means that the cross-correlation function of their Fourier transforms $\tilde{v}_{A1}(\nu)$ and $\tilde{v}_{A2}(\nu)$ are uniquely related to $\tilde{r}_{A1A2}(\nu)$ by

$$\langle \tilde{v}_{A1}(\nu) \tilde{v}_{A2}^*(\nu') \rangle = \tilde{r}_{A1A2}(\nu) \delta(\nu - \nu'). \quad (2.31)$$

Equation 2.31 above can be easily obtained with the same procedure used in equation A.2 from Appendix A.

On the other hand, from equations 2.29 and 2.30 we have

$$\langle \tilde{v}_{A1}(\nu) \tilde{v}_{A2}^*(\nu') \rangle = \int_{4\pi} \int_{4\pi} \langle \tilde{e}(\hat{\mathbf{n}}, \nu) \tilde{e}^*(\hat{\mathbf{n}}', \nu') \rangle A_1(\hat{\mathbf{n}}) A_2^*(\hat{\mathbf{n}}') e^{-i2\pi\nu\tau_g(\hat{\mathbf{n}})} d\Omega d\Omega'. \quad (2.32)$$

As explained in Appendix B, most radio sources are *incoherent* which implies that signals from different parts of the source can be treated independently. In this case, $\langle \tilde{e}(\hat{\mathbf{n}}, \nu) \tilde{e}^*(\hat{\mathbf{n}}', \nu') \rangle$ can be written as (equation B.14)

$$\langle \tilde{e}(\hat{\mathbf{n}}, \nu) \tilde{e}^*(\hat{\mathbf{n}}', \nu') \rangle = I_\nu(\hat{\mathbf{n}}) \delta(\nu - \nu') \delta(\hat{\mathbf{n}} - \hat{\mathbf{n}}'). \quad (2.33)$$

Using equation 2.33 above in equation 2.32 we obtain

$$\langle \tilde{v}_{A1}(\nu) \tilde{v}_{A2}^*(\nu') \rangle = \left[\int_{4\pi} I_\nu(\hat{\mathbf{n}}) A_1(\hat{\mathbf{n}}) A_2^*(\hat{\mathbf{n}}) e^{-i2\pi\nu\tau_g(\hat{\mathbf{n}})} d\Omega \right] \delta(\nu - \nu'). \quad (2.34)$$

Comparing equations 2.31 and 2.34 we obtain

$$\tilde{r}_{A1A2}(\nu) = \int_{4\pi} I_\nu(\hat{\mathbf{n}}) A_1(\hat{\mathbf{n}}) A_2^*(\hat{\mathbf{n}}) e^{-i2\pi\nu\tau_g(\hat{\mathbf{n}})} d\Omega. \quad (2.35)$$

Equation 2.35 together with equation 2.16 give us the general response of the inter-

ferometer to any radio source of intensity distribution $I_\nu(\hat{\mathbf{n}})$.

Assuming receivers with very narrow bandpass filters centered at ν then, up to an overall constant (which is equal to $G = h_1(\nu)h_2^*(\nu)$) we can write the output of the correlator r_{12} as

$$r_{12} = \langle v_1(t)v_2^*(t) \rangle = \int_{4\pi} I_\nu(\hat{\mathbf{n}})A_1(\hat{\mathbf{n}})A_2^*(\hat{\mathbf{n}})e^{-i2\pi\mathbf{b}_\lambda \cdot \hat{\mathbf{n}}}d\Omega \quad (2.36)$$

where $\mathbf{b}_\lambda = \mathbf{b}/\lambda$ is the baseline vector normalized by the observed wavelength $\lambda = c/\nu$.

Since in the Rayleigh-Jeans limit we can express the intensity I_ν in terms of the brightness temperature T as $I_\nu = k_b T/\lambda^2$, then we can always re-normalize r_{12} (all the constants can be made part of the overall gain G mentioned before) so it has temperature units, which is convenient, for example, when we compare the observed signal to the noise temperature of the receiving system, or to the 21 cm brightness temperature. We can write

$$r_{12} = \langle v_1(t)v_2^*(t) \rangle = \int_{4\pi} T(\hat{\mathbf{n}})|A(\hat{\mathbf{n}})|^2 e^{-i2\pi\mathbf{b}_\lambda \cdot \hat{\mathbf{n}}}d\Omega \quad (2.37)$$

where we assumed identical antennas so $A_1(\hat{\mathbf{n}})A_2^*(\hat{\mathbf{n}}) = |A(\hat{\mathbf{n}})|^2$.

We call equation 2.37 the *complex visibility equation*. The visibility is arguably the fundamental quantity in radio interferometry since it shows how the interferometer observes the radio sky.

Note that, with this temperature-like definition of the visibility, and normalizing the beam such that $\int_{4\pi} |A(\hat{\mathbf{n}})|^2 = 1$, then we ensure that for a sky with uniform brightness temperature T then the autocorrelation of each antenna is just $r_{11} = r_{22} = T$.

2.3.4 Sensitivity of a radio interferometer: the radiometer equation

As discussed in Section 2.3, the correlation between two complex-valued signals v_1 and v_2 is $r_{12} = \langle v_1(t)v_2^*(t) \rangle$ (equation 6.12), while the output of a digital correlator after integrating N_s samples is $\hat{r}_{12} = N_s^{-1} \sum_{n=1}^{N_s} v_1[n]v_2^*[n]$ (equation 2.10).

In the case of CHIME, the signals v_1 and v_2 are the outputs of the respective narrow-band digital filters centered at ν_c . The output of these filters have the frequency content of the sky signals concentrated around ν_c , which are represented by their complex envelope (see Section 2.2). Furthermore, if the inputs of these filters are stationary jointly Gaussian random processes with zero mean (which is the typical case for radio astronomy applications, see Section 2.1), then it can be shown that the output vector $\mathbf{v} = (v_1, v_2)^T$

(the superscript T denotes matrix transpose) is a stationary complex Gaussian process that is *circularly symmetric* [34, 35]. This means that \mathbf{v} has zero mean and zero relation matrix, i.e.

$$\langle \mathbf{v} \rangle = \begin{pmatrix} \langle v_1 \rangle \\ \langle v_2 \rangle \end{pmatrix} = \mathbf{0}, \quad \langle \mathbf{v} \mathbf{v}^T \rangle = \begin{pmatrix} \langle v_1 v_1 \rangle & \langle v_1 v_2 \rangle \\ \langle v_2 v_1 \rangle & \langle v_2 v_2 \rangle \end{pmatrix} = \mathbf{0}. \quad (2.38)$$

Thus, the complex Gaussian process \mathbf{v} is completely specified by its covariance matrix $\langle \mathbf{v} \mathbf{v}^H \rangle$ (the superscript H denotes matrix conjugate transpose).

Since the vectors $\mathbf{v}[n] = (v_1[n], v_2[n])^T$ are samples from the stationary random process \mathbf{v} , then these vectors come from the same distribution. From equation 2.10, this means that $\langle \hat{r}_{12} \rangle = r_{12}$, so the measured correlation \hat{r}_{12} is an unbiased estimator of r_{12} . We also know from ergodicity that $\hat{r}_{12} \rightarrow r_{12}$ as the averaging increases (equation 2.12).

Since the number of averaged samples N_s is finite, the estimates \hat{r}_{12} will fluctuate around the mean r_{12} . The strength of this fluctuation is characterized by the variance of the correlator output

$$\text{var}(\hat{r}_{12}) = \langle (\hat{r}_{12} - \langle \hat{r}_{12} \rangle)(\hat{r}_{12} - \langle \hat{r}_{12} \rangle)^* \rangle = \langle \hat{r}_{12} \hat{r}_{12}^* \rangle - \langle \hat{r}_{12} \rangle \langle \hat{r}_{12} \rangle^*. \quad (2.39)$$

The first term of equation 2.39 is

$$\langle \hat{r}_{12} \hat{r}_{12}^* \rangle = \frac{1}{N_s^2} \sum_{m,n} \langle v_1[n] v_2[n]^* v_1[m]^* v_2[m] \rangle. \quad (2.40)$$

We can use Isserlis theorem [36] to expand the fourth-order moment above

$$\begin{aligned} \langle \hat{r}_{12} \hat{r}_{12}^* \rangle &= \frac{1}{N_s^2} \sum_{m,n} [\langle v_1[n] v_2[n]^* \rangle \langle v_1[m]^* v_2[m] \rangle + \langle v_1[n] v_1[m]^* \rangle \langle v_2[n]^* v_2[m] \rangle + \\ &\quad \langle v_1[n] v_2[m] \rangle \langle v_2[n]^* v_1[m]^* \rangle]. \end{aligned} \quad (2.41)$$

Since $\mathbf{v}[n]$ has zero relation matrix then the last term in equation 2.41 is zero. As for the second term, it is non-zero only when $m = n$ since, for Nyquist sampling of band

limited signals (as is the case of CHIME), the samples are uncorrelated [33], so

$$\begin{aligned}
\langle \hat{r}_{12} \hat{r}_{12}^* \rangle &= \frac{1}{N_s^2} \sum_{m,n} \langle v_1[n] v_2[n]^* \rangle \langle v_1[m]^* v_2[m] \rangle + \frac{1}{N_s^2} \sum_n \langle v_1[n] v_1[n]^* \rangle \langle v_2[n]^* v_2[n] \rangle \\
&= \left(\frac{1}{N_s} \sum_n \langle v_1[n] v_2[n]^* \rangle \right) \left(\frac{1}{N_s} \sum_m \langle v_1[m]^* v_2[m] \rangle \right) + \frac{1}{N_s} \langle \hat{r}_{11} \rangle \langle \hat{r}_{22} \rangle \\
&= \langle \hat{r}_{12} \rangle \langle \hat{r}_{12} \rangle^* + \frac{1}{N_s} \langle \hat{r}_{11} \rangle \langle \hat{r}_{22} \rangle.
\end{aligned} \tag{2.42}$$

From equations 2.39 and 2.42 we obtain

$$var(\hat{r}_{12}) = \frac{1}{N_s} \langle \hat{r}_{11} \rangle \langle \hat{r}_{22} \rangle. \tag{2.43}$$

The result above is called the *radiometer equation*. It is useful, for example, to determine the point source sensitivity of an interferometric array. To write it in a more familiar version, assume that the telescopes (antennas and receivers) are identical so $\hat{r}_{11} = \hat{r}_{22}$. Also, the usual case in radio astronomy is that the sky signals are weak and the correlator inputs are dominated by the system noise. This means that the autocorrelations, which are just measures of the input power, are equal to the *system temperature*⁷, T_{sys} , so we can write

$$\sigma_{T,12} = \frac{T_{sys}}{\sqrt{N_s}} = \frac{T_{sys}}{\sqrt{\Delta\nu\tau}} \tag{2.44}$$

where $\sigma_{T,12}$ is the size of the fluctuations (standard deviation) of \hat{r}_{12} , $\Delta\nu$ is the (total, including negative frequencies) receiver bandwidth, and τ is the integration time. In order to detect the signal from a radio point source, its temperature T_{src} must raise the system temperature (several times) above the output fluctuations $\sigma_{T,12}$, so the signal-to-noise (SNR) of the observation becomes

$$SNR = \frac{T_{src}}{\sigma_{T,12}} = \frac{T_{src}}{T_{sys}} \sqrt{\Delta\nu\tau} \tag{2.45}$$

In general, for an interferometric array consisting of N identical telescopes (so there are $N(N-1)/2$ identical two-element interferometers), the size of the fluctuations for any visibility pair is $\sigma_{T,ij} = T_{sys}/\sqrt{\Delta\nu\tau}$. Since for a point source, and up to a phase, all

⁷The system temperature is a measure of the power level when a radio telescope is pointed at ‘blank’ region of the sky (with no sources). It includes the contributions from the sky background, the noise in antenna (e.g. ground spill-over), and the noise in the receiving system. The system temperature is the mean power level of a radio telescope.

the visibilities measure the same value (the temperature distribution of a point source is a delta function so it has equal power in all spatial frequencies), then we have $N(N-1)/2$ independent measurements of this source so the SNR becomes

$$SNR = \frac{T_{src}}{T_{sys}} \sqrt{\frac{N(N-1)\Delta\nu\tau}{2}}. \quad (2.46)$$

2.4 Mapping the 21 cm universe with CHIME

In contrast to traditional radio interferometry applications that often involve telescopes with a narrow field of view that track and observe small regions of the sky, mapping the universe with the 21 cm line requires rapidly observing large volumes with wide-field radio telescopes. This is because it involves mapping a faint and diffuse cosmic signal that covers all of the sky and needs to be separated from foreground contamination (dominated by synchrotron emission from both our own galaxy and high redshift radio galaxies) that is many orders of magnitude brighter.

CHIME is a transit interferometer, meaning that it has a fixed pointing and observes the sky as the earth rotates (as it transits through the telescope's field of view, a strategy also known as drift scan). Furthermore, since CHIME has a wide field of view in order to map large portions of the sky at once, the conventional flat-sky approximation described in Appendix C is not valid, and a different approach that incorporates the curvature of the sky is required.

In this section we will briefly describe the m -mode formalism, a technique for the analysis of transit radio telescopes in the spherical harmonic domain that is well-suited for wide-field radio interferometers like CHIME. It allows a computationally efficient representation of the data that facilitates map making and foreground subtraction. The m -mode formalism is described in detail in [37].

2.4.1 The m -mode formalism

From Section 2.3.3, the visibility r_{ij} is the instantaneous correlation between antennas i and j ,

$$r_{ij} = \langle v_i v_j^* \rangle = \int_{4\pi} T(\hat{\mathbf{n}}) A_i(\hat{\mathbf{n}}) A_j^*(\hat{\mathbf{n}}) e^{-i2\pi \mathbf{b}_{\lambda,ij} \cdot \hat{\mathbf{n}}} d\Omega \quad (2.47)$$

where $\mathbf{b}_{\lambda,ij} = \mathbf{b}_{ij}/\lambda$ is the baseline vector for antennas i and j normalized by the observed wavelength (\mathbf{b}_{ij} starts on antenna i and ends on antenna j).

Under the narrow-beam and flat-sky approximation, equation 2.47 becomes a Fourier transform mapping between the sky and the uv -plane (see Appendix C). Instead, we realize that as the earth rotates, both the primary beams and the baselines of a transit interferometer rotate relative to the sky. This means that the visibilities are periodic with period equal to one sidereal day. In order to exploit this property, we first write the visibility at time t in terms of the *beam transfer function* B_{ij}

$$r_{ij}(t) = \int_{4\pi} d\Omega' B_{ij}(\theta', \phi', t) T(\theta', \phi') \quad (2.48)$$

where⁸

$$B_{ij}(\theta', \phi', t) = A_i(\theta', \phi', t) A_j^*(\theta', \phi', t) e^{-i2\pi \mathbf{b}_{\lambda, ij} \cdot \hat{\mathbf{n}}}. \quad (2.49)$$

Note that the rotation of the earth creates a linear correspondence between time t and the azimuthal angle ϕ . Also that the ϕ dependence simply rotates B_{ij} about the earth's polar axis, so we can write

$$r_{ij}(\phi) = \int_{4\pi} d\Omega' B_{ij}(\theta', \phi' - \phi) T(\theta', \phi'). \quad (2.50)$$

Since $r_{ij}(\phi)$ is periodic with period 2π , it can be decomposed into its Fourier series, given by

$$\tilde{r}_k^{ij} = \int_{2\pi} \frac{d\phi}{2\pi} r_{ij}(\phi) e^{-ik\phi}. \quad (2.51)$$

The next step is to write B_{ij} and T in terms of their respective spherical harmonic expansions (note that these a_{lm} coefficients will not be the same as in Section 1.3 since T contains not only the 21 cm signal but also the foreground sources)

$$T(\theta', \phi') = \sum_{l,m} a_{lm} Y_{lm}(\theta', \phi') \quad B_{ij}(\theta', \phi') = \sum_{lm} B_{lm}^{ij} Y_{lm}^*(\theta', \phi'). \quad (2.52)$$

Note that we chose to expand B_{ij} in terms of the conjugate spherical harmonics since this simplifies the calculations that follow.

From the definition of the spherical harmonics, $Y_{l_1 m_1}(\theta', \phi') \propto e^{im_1 \phi'} P_{l_1 m_1}(\cos \theta')$, it is clear that

⁸Note that the argument of the exponential in equation 2.49 is also a function of θ' , ϕ' , and t .

$$B_{ij}(\theta', \phi' - \phi) = \sum_{l_1, m_1} B_{l_1 m_1}^{ij} e^{im_1 \phi} Y_{l_1 m_1}^*(\theta', \phi'). \quad (2.53)$$

Using equation 2.53 for B_{ij} and equation 2.52 for T in equation 2.51 we have

$$\begin{aligned} \tilde{r}_k^{ij} &= \int_{2\pi} \frac{d\phi}{2\pi} \left\{ \int_{4\pi} d\Omega' \left[\sum_{l, m, l_1, m_1} a_{lm} B_{l_1 m_1}^{ij} e^{im_1 \phi} Y_{lm}(\theta', \phi') Y_{l_1 m_1}^*(\theta', \phi') \right] \right\} e^{-ik\phi} \\ &= \int_{2\pi} \frac{d\phi}{2\pi} \left\{ \sum_{l, m, l_1, m_1} a_{lm} B_{l_1 m_1}^{ij} e^{im_1 \phi} \left[\underbrace{\int_{4\pi} d\Omega' Y_{lm}(\theta', \phi') Y_{l_1 m_1}^*(\theta', \phi')}_{\delta_{ll_1} \delta_{mm_1}} \right] \right\} e^{-ik\phi} \\ &= \int_{2\pi} \frac{d\phi}{2\pi} \left[\sum_{l, m} a_{lm} B_{lm}^{ij} e^{im\phi} \right] e^{-ik\phi} \\ &= \sum_{l, m} a_{lm} B_{lm}^{ij} \left[\underbrace{\int_{2\pi} \frac{d\phi}{2\pi} e^{-i(k-m)\phi}}_{\delta_{km}} \right] \\ \tilde{r}_m^{ij} &= \sum_l a_{lm} B_{lm}^{ij}. \end{aligned} \quad (2.54)$$

In practice, the visibilities are also corrupted by instrumental noise, so an additional noise term $n_{ij}(\phi)$ must be added to equation 2.50. Assuming that the noise is stationary such that its statistics are independent of ϕ and repeating the procedure above leads to

$$\tilde{r}_m^{ij} = \sum_l a_{lm} B_{lm}^{ij} + n_m^{ij} \quad (2.55)$$

or, using a more compact notation that also includes the additional degrees of freedom (the frequencies ν and feed pairs ij), we can write (see [37] for details)

$$\tilde{\mathbf{r}} = \mathbf{B}\mathbf{a} + \mathbf{n} \quad \forall m. \quad (2.56)$$

Equation 2.56 describes the way an interferometric array (determined by the beam transfer matrices \mathbf{B} and the noise vectors \mathbf{n}) maps the sky (contained in vectors \mathbf{a}) into the observed data (the visibility vectors $\tilde{\mathbf{r}}$). Note that different m modes do not mix

(the stationarity assumption means that the noise m modes are uncorrelated) so we can perform operations for each m independently. Also that the finite size of the instrument ensures that the sums are finite.

2.4.2 Power spectrum from m -modes

With the m -mode formalism the measurement process has been reduced to a linear map with finite degrees of freedom, so operations like map-making, foreground subtraction, and power spectrum estimation become conceptually simple. In particular, the angular power spectrum of the sky signal can be recovered from equation 2.56 directly.

The covariance of the visibilities, $C_{(ij\nu m);(i'j'\nu' m')} = \langle \tilde{r}_{ij\nu}^m \cdot \tilde{r}_{i'j'\nu'}^{m'*} \rangle$, is a prohibitively large matrix (dimension of the order of 10^8 for an experiment like CHIME) since it is the covariance between all possible degrees of freedom: baselines, frequencies and m modes. However, under the assumption that the sky signal is a homogeneous and isotropic random field (see Section 1.3.2), then both the amount of computation and data storage are significantly reduced since this covariance is block diagonal and the statistics can be calculated in an m by m basis. In this case there is a simple linear relationship between the covariance of the data and that of the sky signal that can be written in matrix notation as

$$\mathbf{C} = \mathbf{B}\mathbf{C}_{sky}\mathbf{B}^H + \mathbf{N} \quad \forall m \quad (2.57)$$

where \mathbf{C}_{sky} and \mathbf{N} are the covariance matrices of the sky signal and the instrumental noise, respectively.

As explained in Chapter 1, the covariance of the 21 cm signal encodes the cosmological information that we are interested in (the BAO signal). In practice, \mathbf{C}_{sky} contains both the 21 cm signal (\mathbf{C}_{21}), which we are ultimately trying to recover, and foreground sources (\mathbf{C}_f) like diffuse synchrotron emission from the galaxy and emission from extragalactic point sources. To efficiently extract \mathbf{C}_{21} we also need to understand the two-point statistics of the foregrounds. However, the m -mode formalism allows a powerful foreground removal technique in the form of the Karhunen-Loève transformation to be used, which allows us to identify a basis in which the astrophysical foregrounds and 21 cm signal are maximally separated. For details on foreground removal with the Karhunen-Loève transform refer to [37].

2.5 Conclusions

In this chapter we presented in detail important concepts of radio astronomy and radio interferometry that we use intensively in this thesis. We discussed the use of random processes to model the signals received from radio astronomical sources and introduced the concepts of stationarity and ergodicity. We also discussed the representation of band-limited signals and derived the narrowband approximation that allows us to represent time delays of the sky signals as phase shifts of their equivalent baseband representation. This approximation is commonly used in interferometric techniques including beamforming, fringe-stopping, and image reconstruction. We also determined the limitations of this approximation and its application for CHIME.

We presented a detailed analysis of the two-element interferometer which forms the basis for larger and more complex interferometric arrays. We studied the response of the two-element interferometer and derived the complex visibility equation which shows how the interferometer observes the radio sky. We also presented the radiometer equation that measures the sensitivity of the interferometer.

Finally, we introduced the m -mode formalism technique for the analysis of transit radio telescopes in the spherical harmonic domain that is well-suited for wide-field radio interferometers like CHIME.

Chapter 3

The CHIME instrument

The CHIME telescope consists of an array of four cylindrical reflectors that are instrumented with a total of 1024 dual-polarization feeds. CHIME operates as a transit interferometer that surveys the northern half of the sky every day with an instantaneous field of view of 90 degrees by 1.3-2.6 degrees and a synthesized beam resolution of 12-24 arcminutes. The signals from each feed are amplified using receivers with noise performance below 50 K, filtered to 400-800 MHz, and sampled at 800 MSPS¹. A hybrid FX digital correlator processes the signals from all the feeds and computes the complex-valued correlations which provide the interferometric visibilities. In this chapter we describe the design and characteristics of CHIME and the CHIME pathfinder, which is an early prototype of CHIME consisting of two cylindrical reflectors with a total of 128 dual-polarization feeds.

3.1 Science requirements

The choice of the CHIME design parameters is driven by its science goals in terms of redshift coverage to constrain dark energy parameters, angular and line-of-sight resolution to measure the BAO scale, and sensitivity to detect the 21 cm signal. These science goals determine the frequency coverage, angular resolution, frequency resolution, and integration time of the experiment.

3.1.1 Frequency coverage

The ability of an experiment to measure the properties of dark energy can be quantified by the Dark Energy Task Force (DETF, [11]) Figure of Merit (FOM), defined as the

¹MSPS is mega samples per second.

reciprocal of the error ellipse enclosing the 95% confidence limit of the $w_0 - w_a$ plane, where w_0 and w_a are the dark energy equation of state parameters defined in equation 1.1. A larger FOM indicates a greater accuracy in constraining these parameters.

As demonstrated in [17], the sensitivity of a BAO survey, and thus the improvement of the FOM, depends on measuring the expansion history of the universe before and after dark energy starts to dominate this expansion. This inflection point corresponds to a redshift slightly above $z = 1$, or an observed frequency of 710 MHz. The improvement factor can be increased by reaching lower redshifts, corresponding to operating at higher frequencies. The upper edge of the operating frequency range of CHIME is practically limited by the radio-frequency (RF) interference (RFI) from the cell phone band beginning above 800 MHz (even in the radio quiet zone where CHIME operates this RFI is still present). Note that probing redshifts below ~ 0.8 (790 MHz) provides little improvement to the FOM since these redshifts have already been probed by past and current surveys. As for the CHIME lowest frequency, note that the improvement of the FOM is limited above ~ 2.5 (405 MHz). This is because if dark energy approximates a cosmological constant then it is expected to be sub-dominant to matter density for $1 \lesssim z \lesssim 3$ and its effect is negligible at higher redshifts. However, if dark energy is more important ($w_a \neq 0$) at high redshifts than in the standard model of cosmology (see Chapter 1), then $z \gtrsim 1$ measurements become more useful for dark energy constraints.

3.1.2 Angular resolution

The BAO scale is $r_s \approx 148$ Mpc and appears as oscillations in the baryon power spectrum (see Section 1.2.1). On small scales, this BAO structure is attenuated by nonlinearities such as those induced from nonlinear gravitational clustering. These nonlinearities become significant for spatial scales smaller than ~ 44 Mpc [14], making the reconstruction of the position of the BAO peak less accurate (they shift and broaden the peak). Thus, this is the smallest spatial scale we shall consider, and a Nyquist sampled map up to this scale would require pixels of size $r_{pix} \sim 22$ Mpc.

The angular resolution of an interferometric array with maximum baseline b_{max} is $\Delta\theta = \lambda/b_{max}$ (See Appendix C and equation C.9). The longest baseline for the CHIME telescope is $b_{max} \sim 100$ m, so its angular resolution is $\sim 0.2^\circ$ at redshift 0.8 (800 MHz) and $\sim 0.4^\circ$ at redshift 2.5 (400 MHz). From equation 1.18 this angular resolution $\Delta\theta$ corresponds to a comoving distance $r_\perp = \Delta\theta d_A(1+z)$ perpendicular to the line of sight, so ~ 11 Mpc at redshift 0.8 and ~ 45 Mpc at redshift 2.5 for CHIME. This comoving distance resolution is appropriate to measure the BAO structure since at the lowest

redshifts we can safely oversample the BAO scale and at the highest redshifts we can reach well below the smallest linear scales².

The CHIME pathfinder has $b_{max} \sim 30$ m, corresponding to an angular resolution $\sim 0.7^\circ - 1.4^\circ$. The pathfinder resolution should be appropriate to Nyquist sample the BAO scale up to $z \sim 1.4$ (600 MHz) perpendicular to the line of sight.

3.1.3 Frequency resolution

From equation 1.17, the spatial scale r_{pix} in the line-of-sight direction corresponds to a redshift change $\Delta z = Hr_{pix}/c$. Using $\nu = \nu_0/(1+z)$, where $\nu_0 \approx 1420$ MHz is the rest-frame frequency for the hyperfine transition of hydrogen (Section 1.3), we find that the spatial scale $r_{pix} \sim 22$ Mpc corresponds to a frequency separation

$$\Delta\nu_s = \frac{H(z)r_{pix}}{c} \frac{\nu_0}{(1+z)^2}. \quad (3.1)$$

This frequency separation is ~ 3.5 MHz at redshift 0.8 (800 MHz) and ~ 2.2 MHz at redshift 2.5 (400 MHz). The CHIME correlator splits the 400-800 MHz bandwidth uniformly into 1024 frequency channels ~ 0.4 MHz wide. This frequency resolution is well above the requirements to measure the BAO structure along the line of sight, but necessary for RFI removal as well as for maintaining coherence and avoiding loss of signal for the longest baselines (see Section 2.3.2).

3.1.4 Sensitivity and integration time

Using observations of damped Ly α systems, it is found that the cosmic density of neutral hydrogen gas is $\Omega_{HI} \sim 10^{-3}$ in the redshift range $0.8 \lesssim z \lesssim 2.5$ with a weak dependence on redshift [38]. This sets the average neutral fraction of hydrogen at $\bar{x}_{HI} \sim \Omega_{HI}/\Omega_b \sim 0.02$. Using equation 1.30 we find that the average sky brightness temperature due to 21 cm radiation is $\overline{\Delta T_b} \sim 320$ μ K at $z = 1.5$.

Measuring the 21 cm power spectrum requires detecting variations of the 21 cm sky brightness temperature $\Delta T_b = \overline{\Delta T_b} \delta_{21}$ (equation 1.32), which on ~ 22 Mpc scales are expected to be of order 100 μ K [39].

If we require a telescope sensitivity $\Delta T \sim 100$ μ K per sky pixel then, from the radiometer equation 2.44, the observation time per pixel is $t_{pix} = T_{sys}^2/(\Delta\nu \cdot \Delta T^2) \sim$

²Note that measuring the BAO scale at high redshifts is very challenging, not only because its size decreases with redshift, but also because the angular resolution of the telescope degrades at low frequencies. This means the size (and cost) of CHIME is dominated by how high a redshift we expect to reach.

35 hours, where we used $\Delta\nu \sim 2$ MHz and $T_{\text{sys}} \sim 50$ K, which are typical values for CHIME. For a single telescope with the same angular resolution as CHIME ($\sim 0.2^\circ$ at 800 MHz) surveying the northern half of the sky this would require ~ 2000 years of observation time. On the other hand, CHIME can form 1024 (dual polarization) independent beams by taking linear combinations of the $\sim N^2/2$ visibilities (equivalent to beamforming) with roughly the same noise (up to a factor of $\sqrt{2}$), reducing the integration time to approximately 2 years.

3.2 The CHIME telescope

CHIME is a hybrid cylindrical transit interferometer located at the Dominion Radio Astrophysical Observatory (DRAO³) near Penticton, Canada (Figure 3.1). It is an array of four adjacent cylindrical paraboloidal reflectors 20 m wide by 100 m long oriented north-south for a total collecting area of 8000 m². The cylinders are fixed with no moving parts, forming a transit instrument that surveys the northern half of the sky every day with an instantaneous field of view of $\sim 90^\circ$ by $\sim 1^\circ - 2^\circ$ and a synthesized beam resolution of $\sim 0.2^\circ - 0.4^\circ$.

The cylinder structures are steel, and the reflecting surface is a 19-mm-spacing mesh made of galvanized-steel and bolted to the structure. The mesh spacing achieves a good trade-off between reflecting effectively the sky radio signals and allowing snow to fall through. The measured surface RMS deformation is ~ 9 mm.

The focal line of each cylinder is instrumented with 256 dual polarization feeds placed every ~ 30 cm, giving a total of 2048 sky signals. This spacing is enough to Nyquist sample the north-south field-of-view over almost all the CHIME band.

A summary of the CHIME design parameters is in Table 3.1. A photograph of the CHIME telescope at DRAO is shown in Figure 3.2. The construction of the CHIME telescope is complete and its first light was late 2017. The telescope is now operational and collecting commissioning and early science data.

3.2.1 The CHIME pathfinder

The CHIME pathfinder is presented in detail in [2]. It is a proof-of-concept instrument for CHIME consisting of two cylinders 20 m wide by 37 m long oriented north-south for a total collecting area of ~ 1500 m². Each cylinder is instrumented with 64 dual-polarization feeds for a total of 256 sky signals. The pathfinder maps the northern sky

³<https://www.nrc-cnrc.gc.ca/eng/solutions/facilities/drao.html>.



Figure 3.1: The Dominion Radio Astrophysical Observatory (DRAO) is a facility for science and technology research related to radio astronomy. It currently operates five instruments: CHIME, the CHIME pathfinder, a 26-m fully steerable dish (the John A. Galt telescope), a 7-element east-west interferometric radio telescope (the synthesis telescope) and a solar flux monitor. It also features development laboratories for radio-frequency instrumentation. The observatory is located near Penticton, British Columbia, in a region regulated by the federal government to ensure a radio-quiet environment.

Table 3.1: Design parameters for CHIME and the CHIME pathfinder

Parameter	Pathfinder	Full CHIME
Structure	Two 20×37 m cylinders	Four 20×100 m cylinders
Focal ratio f/D	0.25	0.25
Number of feeds per cylinder	64	256
Feed spacing	30 cm	30 cm
Frequency range	400 MHz - 800 MHz	400 MHz - 800 MHz
Redshift range	2.5 - 0.8	2.5 - 0.8
Frequency resolution	390 kHz, 1024 bins	390 kHz, 1024 bins
E-W Field of View	$2.6^\circ - 1.3^\circ$	$2.6^\circ - 1.3^\circ$
N-S Field of View	90°	90°
Synthesized beam size	$1.4^\circ - 0.7^\circ$	$0.4^\circ - 0.2^\circ$
Number of beams	128 dual polarization	1024 dual polarization
Receiver noise temperature	$\lesssim 50$ K	$\lesssim 50$ K



Figure 3.2: Top: the CHIME telescope, a hybrid cylindrical transit interferometer of four 20×100 m cylindrical dishes instrumented with a total of 1024 dual-polarization feeds. CHIME is currently being commissioned and it observed the first fringes from a point source transit late 2017 (photograph taken by Nolan Denman). Bottom: the CHIME pathfinder, a hardware, calibration, and data analysis proof of concept for CHIME. The pathfinder's first light was late 2013 and its commissioning finished early 2015.

every day with the same field of view of CHIME and a synthesized beam resolution of $\sim 0.7^\circ - 1.4^\circ$ in the 400-800 MHz band.

A summary of the pathfinder design parameters is in Table 3.1 and a photograph of the telescope at DRAO is shown in Figure 3.2. The pathfinder’s first light was late 2013 and its commissioning finished early 2015, overlapping with the start of the construction of the full CHIME structure.

The pathfinder has served as a test-bed to

- Develop instrumentation and test the performance of different CHIME subsystems (e.g. feeds, analog receiver, correlator).
- Demonstrate instrument calibration techniques (e.g. beam calibration, receiver noise calibration, digital calibration).
- Test RFI and foreground subtraction techniques.
- Develop the CHIME data analysis pipeline.

The valuable experience and lessons learned during the commissioning of the pathfinder have been used to improve the final CHIME design. However, the pathfinder has been designed to operate as an independent telescope with the sensitivity to measure the BAO at low redshifts.

Its drift-scan strategy, together with its large field of view, also makes the pathfinder (and full CHIME) an excellent instrument to pursue additional science goals, including pulsar monitoring and detection of radio transients and Fast Radio Bursts (FRBs). In 2016 we performed a 52.85 day incoherent-beam FRB search with the pathfinder, corresponding to one of the largest FRB survey exposures to date, with $\sim 2.4 \times 10^5 \text{ deg}^2 \text{ hours}$. An *incoherent beam* was formed by adding the square of the voltages from all the pathfinder feeds. This beam is less sensitive than a beam formed by the coherent combination of the voltages from all the feeds (a *tied-array beam*), but it has a much larger size, favoring the rapid detection of ultra-bright FRBs if these events were common. This allowed constraining the slope of the FRB distribution, $\alpha = -\partial \log N / \partial \log S$, where S is the flux density and N is the number of events. By not detecting any FRBs, it was possible to rule out $\alpha < 0.9$ with 95% confidence, constraining the number of events brighter than $\sim 220\sqrt{(\tau/\text{ms})} \text{ Jy ms}$ for τ in the range 1.3-100 ms to fewer than $\sim 13 \text{ sky}^{-1} \text{ day}^{-1}$. The details of this wide-field FRB survey and the incoherent beamformer are presented in [40].

3.2.2 Signal path

A schematic diagram of the CHIME signal path is shown in Figure 3.3 and further explained in the sections that follow. The signals from the broad-band dual-polarization feeds are amplified (each polarization separately) using room temperature receivers with a noise performance below 50 K that amplify and filter the signals to 400-800 MHz. The correlator is an FX design (temporal Fourier transform before spatial cross-multiplication of data), where the F-engine digitizes (samples at 800 MSPS and with 8 bits) and channelizes (i.e., divide the 400 MHz input bandwidth into thousands of frequency channels) the signals from the 2048 analog receivers⁴. The F-engine also performs the corner-turn network that rearranges the channelized data before sending it to the X-engine that performs the cross-multiplications and averaging to compute the N^2 correlation matrix. These data are sent to the CHIME cosmology pipeline for storage and further processing.

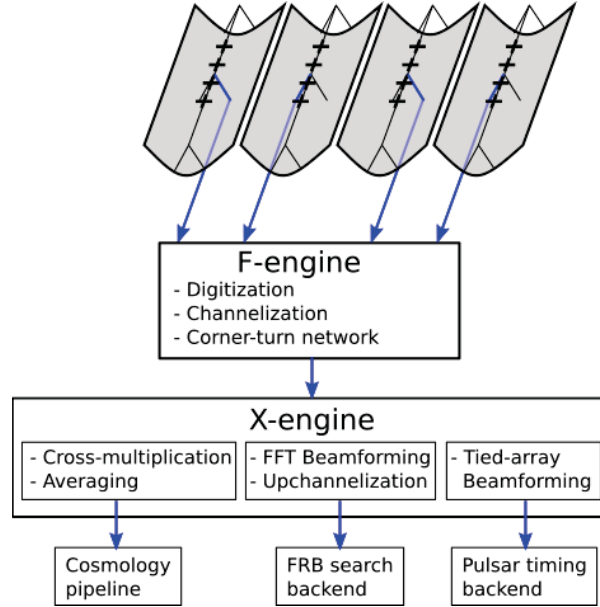


Figure 3.3: Schematic diagram of the CHIME signal path. The signals from each of the dual-polarization feeds are independently amplified and filtered to 400-800 MHz. The F-engine digitizes and channelizes the signals from the 2048 analog receivers. It also re-arranges the data before sending it to the X-engine that computes the N^2 correlation matrix. These data are sent to the CHIME cosmology pipeline for storage and further processing. The X-engine also performs separate beamforming and upchannelization operations on the high-cadence data for two specialized backends for real-time FRB search and pulsar monitoring.

⁴Technically speaking, the F-engine module of the CHIME correlator only performs the frequency channelization. However, in this chapter we use the term F-engine loosely to refer to the system that implements the digitization, channelization, and corner-turn network of the CHIME correlator. The details of this system are presented in Chapter 4.

The X-engine has been upgraded to independently perform additional real-time data processing operations for pulsar timing and FRB science. It forms 10 dual-polarization tied-array beams that allow simultaneous tracking of up to 10 pulsars at different locations. The high-cadence data from these beams are fed into a specialized backend that monitors the dispersion measure, timing, and scintillation of pulsars. A detailed description of the pulsar backend is found in [41]. Separately, the X-engine also implements a Fast Fourier Transform (FFT) beamformer and an additional frequency up-channelization. The high-cadence data from the resulting 1024 dual-polarization and high-frequency-resolution formed beams are fed into a specialized backend that triggers on high-dispersion radio transients to search for FRBs. The CHIME FRB backend is presented in [42], and a detailed description of the FFT beamforming and up-channelization modules of the X-engine is found in [43].

3.3 Analog chain

A schematic diagram of the CHIME analog receiver is shown in Figure 3.4. At the focus of the cylinder, the sky signal received by each feed is amplified by a low noise amplifier (LNA) with a noise temperature below 30 K across the band. Each polarization is amplified separately. The signals are then transmitted 50 m on low-attenuation coaxial cables to nearby receiver huts with RF-shielded rooms, where they are amplified, band-pass filtered to 400-800 MHz, and directly sampled at 800 MSPS with 8 bits. The overall receiver temperature is designed to be below 50 K. This includes ground effects (e.g. ground-spill over the dish edge and through the mesh), but does not include sky contribution. The nominal signal level at the input of the analog-to-digital converters (ADCs) is about -23 dBm, corresponding to a voltage standard deviation of ~ 3.5 bits RMS. As demonstrated in Section 5.3, this level causes a negligible penalty due to quantization errors, while allowing headroom for receiver gain fluctuations and external RFI (see Section 5.3).

3.3.1 Feeds

CHIME uses *cloverleaf antennas* [44] which are compact and broadband dual-polarization feeds made of conventional low loss printed circuit boards. Photographs of the cloverleaf antenna are shown in Figure 3.5. The broad bandwidth is achieved by having petals with smooth curved outer edges. The geometry of the petals was optimized to achieve a return loss better than 10 dB across the CHIME band. For each polarization, differential




Figure 3.4: Schematic diagram of the CHIME analog receiver. At the focus of the cylinder, the sky signal received by the feed is amplified by an LNA with $\sim 44\text{--}38$ dB gain in the 400–800 MHz band and ~ 30 K noise. After that, the signal is sent through 50 m of low-attenuation (-4 to -6 dB) coaxial cable to the receiver hut, where it is amplified again by $\sim 30\text{--}36$ dB and filtered to 400–800 MHz. The signal is about -23 dBm at the input of the ADC. This level causes a negligible penalty due to quantization errors, while allowing headroom for receiver gain fluctuations and external RFI.

signals from pairs of adjacent petals are combined through tuned baluns (the ‘stem’) to form one single-ended output. Thus, each polarization involves currents from the four petals. For details on the feed design and optimization refer to [44].

3.3.2 Low Noise Amplifier

The LNA is attached directly to the feed output. It achieves a noise temperature below 30 K across the CHIME band and its gain decreases smoothly from ~ 44 dB at 400 MHz to ~ 38 dB at 800 MHz. The input matching, output matching, and feedback of the LNA were designed to achieve a low noise temperature. A photograph of the LNA is shown in Figure 3.6.

The CHIME LNA design is an optimized version of the prototype described in [2] that is used for the pathfinder. The main differences with respect to the pathfinder prototype are:

- The CHIME LNA design improves the noise performance by ~ 5 K at the top half of the CHIME band.
- The in-band gain response of the CHIME LNA is very similar to that of the pathfinder. However, the CHIME LNA includes an impedance matching network to produce lower reflection coefficients at the input and output of the LNA and to reduce the out of band gain response, especially at high frequencies.




Figure 3.5: Left: Photograph of the CHIME cloverleaf feed. Right: Photograph of an array of cloverleaf feeds installed at the focal line of the pathfinder.

- The pathfinder has power supplies installed on the focal line to provide power for the LNAs. The CHIME LNA includes an output bias-tee that allows it to receive power from the filter amplifier over the coaxial cable, removing the necessity for power supplies or power cabling on the focal line and allowing LNA power control from the RF room.

3.3.3 Filter amplifier

The output of each LNA is transmitted 50 m on LMR-400 coaxial cables to the RF room where it is connected to a highly linear filter amplifier (FLA). The FLA consists of a custom-made 400-800 MHz anti-aliasing band-pass filter, three amplification stages, and a gain shaping network that flattens the overall gain response of the analog receiver by compensating for the negative slope of the LNA and coaxial cable gain response. A photograph of the FLA is shown in Figure 3.6.

The CHIME FLA design is an optimized version of the prototype described in [2] that is used for the pathfinder. The main differences with respect to the pathfinder prototype are:

- Since the pathfinder FLA has a flat gain response, the overall response of the pathfinder analog receiving system has lower gain at high frequencies due to the LNA and coax contributions, resulting in a non-negligible contribution of the digitization noise at these frequencies. The CHIME design includes a gain shaping

network to flatten the overall gain response of the analog receiver. A comparison between the gain response of the CHIME and pathfinder analog receivers is shown in Figure 3.6.

- The CHIME FLA has a lower noise temperature (~ 300 K) compared to the pathfinder version, especially near the band edges.
- During the pathfinder operation it was found that the FLAs were only conditionally stable at high frequencies, and some FLAs would occasionally oscillate at 4.5 GHz when coupled to the ADCs. This problem has been fixed in the new design, and the CHIME FLAs are now unconditionally stable.
- The pathfinder receiving system operates with an external attenuator at the input of the FLA to adjust the signal level at the input of the ADC and improve the input impedance matching. This attenuation is now integrated in the new CHIME FLA design.
- The input network of the FLA now includes a bias tee that provides power to the LNA through the coaxial cable.

Overall, the CHIME analog receiver has a better dynamic performance compared to the pathfinder analog receiver, achieving lower noise temperature, flatter response over the CHIME band, higher rejection out of band, and unconditional stability. The construction of the pathfinder has proven to be very valuable in allowing the iterative optimization that makes the CHIME analog receiving system close to optimal.

3.4 The CHIME correlator

A schematic diagram of the CHIME correlator is shown in Figure 3.7. It has the structure of an FX correlator, where the data acquisition, Fourier transform channelization and corner-turn networking are performed in field-programmable gate arrays (FPGAs), which are interfaced to a set of graphics processing units (GPUs) that computes the interferometric visibilities.

3.4.1 F-engine

The F-engine is presented in detail in Chapter 4; we describe it briefly here. The F-engine is built around ‘ICE’ motherboards (also known as *iceboards*), which are custom

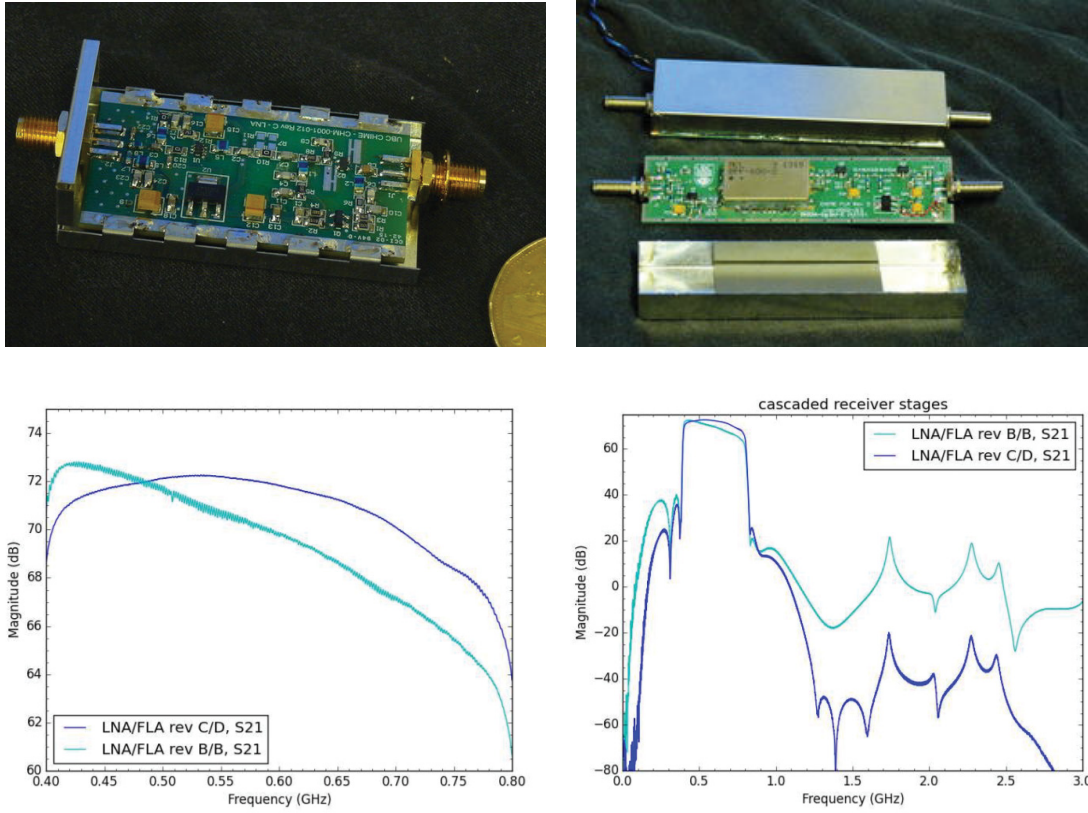


Figure 3.6: Top: Photographs of the CHIME LNA (left) and FLA (right) Bottom: Comparison between the insertion gain of CHIME (blue) and pathfinder (cyan) analog receiving chains (from the input of the LNA to the output of the FLA). The new analog receiver for CHIME has a flatter response over the CHIME band and higher rejection out of band. These measurements and figures were made by Rick Smegal.




Figure 3.7: The CHIME correlator. See Section 3.4 for a detailed description.

FPGA-based processing circuit boards that have been specialized for the CHIME application through custom software, firmware, and digitizer daughter boards. Each ICE motherboard digitizes 16 analog inputs at 800 MSPS with 8 bits, so the 400-800 MHz sky signals are directly sampled using the second Nyquist zone. Each input is then fed to a custom PFB that splits the 400 MHz bandwidth into 1024 frequency bins 390 kHz wide. The output of the F-engine for each frequency bin is a 4-bit real + 4-bit imaginary complex value per frame.

At this point each ICE motherboard has the data from 16 inputs for all the frequency channels. However, computing the correlation matrix requires, in a single place, the data from all the inputs and a subset of frequency channels. This operation is performed in a four-stage corner-turn network. The first stage occurs within each ICE motherboard, where the data are split into 16 subsets, each containing 1/16 of the frequency channels for the 16 inputs. In Figure 3.7 this stage is represented by a red block inside each iceboard module.

Groups of 16 ICE motherboards are packaged in crates (*icecrates*), and all the boards within a crate are interconnected through a high-bandwidth backplane. The second corner-turn stage is a data exchange between the boards in the crate, after which each board has all the data from 256 inputs for 64 frequency channels. In Figure 3.7 this stage is represented by a green block inside each icecrate module.

The third stage is a data exchange between pairs of crates, after which the data from 512 inputs is split into 256 subsets distributed through the ICE motherboards, and each subset contains 4 unique frequency channels. Each crate pair contains all the data for one quarter of the CHIME array, or one cylinder. In Figure 3.7 this stage is represented by a magenta block inside each crate pair module.

3.4.2 X-engine

The cross-multiplication and averaging for all the CHIME inputs takes place in a dedicated GPU-based computing cluster consisting of 256 independent and identical processing nodes. A computer acting as a control system is housed in the same cluster, which serves the software and operating system used by the nodes. This same system aggregates and buffers the data prior to archiving on a data server.

The fourth stage of the corner-turn network is completed within the GPU nodes. Each node receives one frequency subset from each crate pair and recombines the data to compute the correlation matrix for the 2048 inputs and 4 frequency bins. In Figure 3.7 this stage is represented by the blue lines that connect the F-engine and the X-engine.

Table 3.2: Design parameters for the CHIME and pathfinder correlator.

Parameter		Pathfinder	Full CHIME
Number of analog inputs		256	2048
Analog sampling		800 MSPS, 8 bits	800 MSPS, 8 bits
Channelizer type		2048 sample PFB/FFT	2048 sample PFB/FFT
Frequency channels		1024 bins, ~ 390 kHz/bin	1024 bins, ~ 390 kHz/bin
F-engine data path	Input	8 bits	8 bits
	Internal	18+18 bits complex	18+18 bits complex
	Output	4+4 bits complex	4+4 bits complex
Power consumption	F-engine	1.3 kW	10.4 kW
	X-engine	10 kW	260 kW
Internal data rates*	Raw F-engine input	1.6 Tbit/s	13.1 Tbit/s
	F-engine output	0.8 Tbit/s	6.6 Tbit/s
	X-engine operations	13 TcMAC/s	839 TcMAC/s
Output rate (10 s integration)*,*		0.2 Gbit/s	13.8 Gbit/s
Baselines		32896	2098176

*Data rates are for visibility data only and do not include flags and overhead.

*Each component of the output correlation matrix is 4+4 bytes complex.

Each component of the output correlation matrix is a 4-byte real + 4-byte imaginary complex value. The nominal integration time for CHIME is 10 s, corresponding to an X-engine output data rate of ~ 150 TB/day. These data are further compressed over redundant baselines by the CHIME cosmology pipeline prior to long-term archiving. After compression, the output data rate is ~ 1 TB/day.

The design and implementation of the X-engine is presented in [45] and details of the GPU kernels that implement efficiently the correlation operation are presented in [46, 47]. The correlator for the pathfinder consists of a single iccrate for the F-engine and a GPU-cluster with 16 nodes for the X-engine (see [2] for details). The design parameters of the correlator for CHIME and the pathfinder are summarized in Table 3.2.

3.5 Conclusions

The CHIME hybrid cylindrical transit interferometer is operational and collecting early science data. Its design parameters (frequency coverage, angular and spectral resolution, and integration time) are driven by its cosmology science goals of detecting the 21 cm signal, measure the BAO scale, and constraining dark energy parameters. CHIME features 1024 broadband dual-polarization feeds, room temperature receivers with noise performance below 50 K, and a powerful hybrid FX correlator where the data acquisition,

Fourier transform channelization and corner-turn networking are performed in FPGAs. These are interfaced to a GPU cluster that computes the interferometric visibilities.

With its drift-scan strategy and large field of view, CHIME is also an excellent platform to pursue additional science goals, including pulsar monitoring and detection of radio transients and FRBs. The CHIME correlator has been upgraded to perform real-time beamforming and up-channelization operations for two separate specialized backends: An FRB backend that triggers on high-dispersion radio transients to search for FRBs, and a pulsar backend that monitors the dispersion measure, timing, and scintillation of pulsars.

The CHIME pathfinder is a proof-of-concept instrument for CHIME with $\sim 10\%$ of its collecting area that has served as a test-bed for developing instrumentation and testing calibration and analysis techniques. Although the pathfinder has proven to be very valuable to optimize the different CHIME subsystems, its design and operating parameters allow its use as an independent telescope with the sensitivity to measure the BAO at low redshifts and probe the time-variable radio sky.

Chapter 4

The ICE system and the CHIME correlator

The CHIME correlator processes 2048 digitizer inputs with 400 MHz of bandwidth and, measured in number of baselines times bandwidth, it is the largest radio correlator that has been built. We have developed the ICE system, an FPGA-based hardware, firmware, and software framework that has been specialized to implement the digitization, F-engine, and the networking engine of the CHIME correlator. In this chapter we describe the design of the ICE system and its use for the CHIME correlator. We also describe the full end-to-end characterization and validation of the pathfinder correlator demonstrating that the system complies with all the requirements for full CHIME. Finally, we give details of the installation of the pathfinder and full CHIME correlators.

4.1 The ICE system

ICE¹ is a hardware, firmware, and software framework developed at the McGill Cosmology Instrumentation Laboratory for next-generation scientific instruments that require large-scale electronics backends for data acquisition, signal processing, and networking.

The design of the ICE system was mainly based on the specifications for two applications: the digitization, F-engine, and corner-turn network for the CHIME radio correlator and the digital frequency multiplexing (DfMux) bolometer readout system for the South Pole Telescope (SPT) and Simons Array ([48]). However, the flexibility of the ICE system has also allowed it to be used in other applications including the correlator for the Real-time Analysis eXperiment (HIRAX², [49]) and very-long-baseline interferometry (VLBI)

¹ICE is the name of the system, and is not an acronym.

²HIRAX is a transit radio interferometer consisting of 1024 six-meter parabolic dishes in a compact

between the CHIME pathfinder and the Algonquin Radio Observatory (ARO) in Ontario. Details on the use of the ICE system for DfMux and other applications are found in [3].




Figure 4.1: Photograph showing two ICE crates mounted in a standard rack. Each crate is populated with 16 motherboards. The motherboards have daughter boards designed for the CHIME application. The two crates process 512 analog inputs in total.

The ICE system consists of FPGA motherboards and their respective daughter boards, crates with custom backplanes for high-density applications, system software, and FPGA firmware. The daughter boards and firmware are mostly application specific.

and highly redundant configuration that is currently under development for deployment in South Africa. HIRAX has the same science goals of CHIME, and will complement it by observing the southern radio sky in the same redshift range.

Figure 4.1 shows part of the ICE system assembled for CHIME, where two fully populated ICE crates are mounted in a standard rack. Each crate is populated with 16 motherboards. Each motherboard has two digitizer daughter boards that specialize the system for the CHIME application. The system in this example can perform the digitization, frequency channelization, and corner-turn operations for 512 inputs (or one CHIME cylinder).

4.1.1 ICE Hardware

FPGA Motherboard

Figure 4.2 shows a photograph and a block diagram of the ICE motherboard which is the basic processing unit of the ICE system and features

- Two connectors to attach application-specific daughter boards.
- An FPGA for real-time signal processing and networking.
- A co-processor running Linux for remote access and high-level interface to the motherboard.
- A control and monitoring network accessible to both the FPGA and the co-processor.

FPGA: The ICE motherboard hosts a Xilinx Kintex-7 420T FPGA that provides the appropriate amount signal processing and networking resources to meet the requirements of the target applications. It provides twenty-eight 10 Gbit/s high-speed serial transceivers that are used to:

- Communicate with other boards within a crate through an ICE backplane (15 transceivers).
- Communicate with other crates through Quad Small Form-factor Pluggable Plus (QSFP+) connectors on the backplane (4 transceivers).
- Offload data through a pair of QSFP+ connectors on the motherboard (8 transceivers).
- Communicate with the FPGA through a 1 GbE or 10 GbE link (1 transceivers).

The CHIME correlator uses these transceivers to perform the corner-turn network operation that re-arranges the digitized and channelized data before sending it to the X-engine for cross-multiplication and averaging. This application is described in Section 4.2.

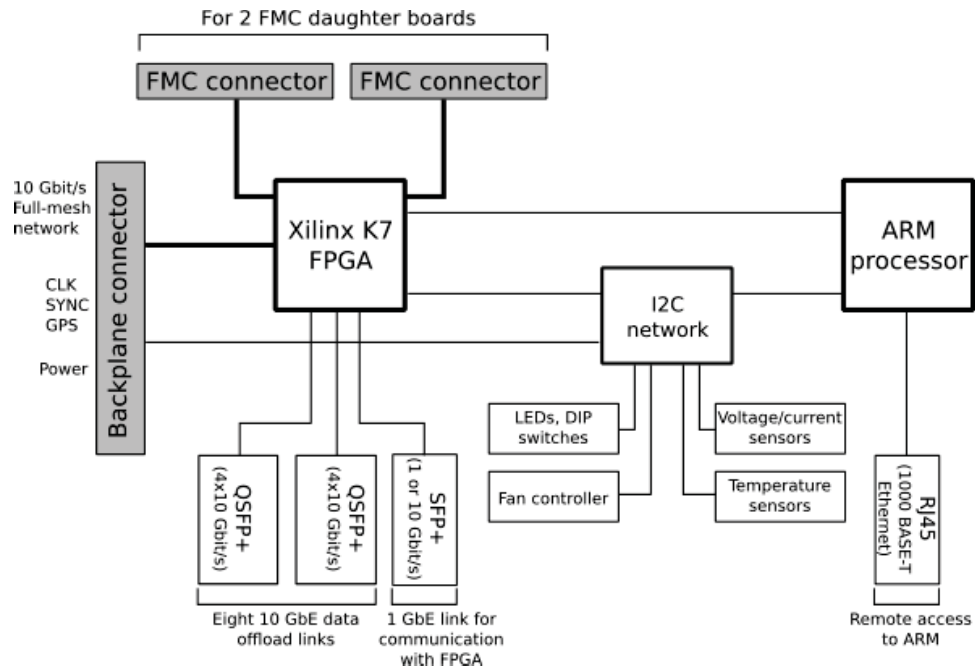



Figure 4.2: Top: Photograph of an ICE motherboard. The FPGA is located below the black heat sink visible near the center of the photo. Application specific daughter boards attach to the two slots labeled FMC A and FMC B. The board can be attached to an ICE backplane through the two black connectors seen at the bottom right. Two QSFP+ connectors at the back of the board provide eight 10 GbE data offload links. Bottom: Simplified block diagram of the ICE motherboard.

ARM Processor: The ICE motherboard includes an ARM³ co-processor (Texas Instruments AM3874) to provide remote access to the motherboard resources. Operations like remote programming of FPGAs, continuous system monitoring, and implementation of high-level functions can be done through the ARM processor. These features not only make the ARM processor particularly useful for managing large arrays of boards, but also allow freeing valuable FPGA resources that can be used for simpler but real-time operations.

FMC daughter boards

By installing different industry-standard FPGA Mezzanine Card⁴ (FMC) daughter boards, the ICE motherboards can be specialized for different applications. Details of the custom daughter board designed for the CHIME application are presented in Section 4.2.1.

Backplane and Crate

For applications like CHIME that require large arrays of FPGA nodes, crates equipped with custom ICE backplanes have been designed to package up to sixteen motherboards. The backplane provides power, clock, timestamp, and synchronization signals to all its motherboards. It also implements a full-mesh passive network that provides a direct 10 Gbit/s link between all the boards within a crate. Finally, the backplane has sixteen QSFP+ connectors (each capable of 4×10 Gbit/s) for communications between motherboards in different crates. Photographs of the backplane are shown in Figure 4.3.

The use of ICE backplanes and crates reduces wiring and volume occupied by the hardware, facilitating the cooling of the boards and simplifying the implementation of the ICE system in high-density applications.

Other hardware features

Other hardware features of the ICE system are:

- **Power system:** To facilitate the implementation for the user, the ICE system operates from a single power rail provided with a voltage in the range 13-20V DC. Switching power supplies within the boards generate all the voltages required to power the FPGA, ARM and peripherals. The voltage and current provided by each power supply can be monitored by either the FPGA or the ARM.

³https://en.wikipedia.org/wiki/ARM_architecture

⁴https://en.wikipedia.org/wiki/FPGA_Mezzanine_Card




Figure 4.3: Left: Photograph showing the ICE backplane from the motherboard mating side. Right: Photograph showing the back of the crate-mounted ICE backplane and one motherboard attached. Eight dual-QSFP+ connectors at the back of the backplane provide data links between crates. Connectors for clock, timestamp, and synchronization signals are also at the back of the backplane.

- **Cooling:** The power consumption of a full crate of ICE motherboards operating at capacity is approximately 1.3 kW. When used in an ICE crate, custom-made fan trays circulate air through the system to keep the FPGAs below their specified maximum temperature (85 C).
- **Timing:** The system operates with a single 10 MHz clock from which all its components (FPGA, ARM co-processor, switching power supplies, peripherals, etc.) derive their timing. The backplane distributes the clocking and synchronization signals to all the boards in the crate through a high-precision sub-picosecond jitter network.
- **Automatic hardware identification system:** The ICE system detects automatically every major hardware component (motherboards, backplanes, daughter boards, QSFP+ cables, etc.) available in the array. When the array is booted, the ARM processors advertise the presence of the hardware on the network.
- **Peripherals:** The ICE motherboards and backplanes are equipped with serial networks that are accessible to both the FPGA and the ARM processor through the I²C protocol⁵, and that allow the user to control peripherals, fans, and power systems. They also allow monitoring voltages, currents and temperatures of different parts of the system.

⁵<https://en.wikipedia.org/wiki/I2C>

4.1.2 ICE Firmware and Software

Three pieces of firmware and software compose the ICE framework: the Python-based control software, the ARM processor software, and the FPGA firmware. Application-specific firmware and software is built upon a core library included with the ICE system that provides the basic structure of these three elements.

Python-based System Control Software

The control software provides the tools necessary to configure, operate, and monitor the hardware across the array in an application-specific manner. It is built upon the *IceCore* Python package which is part of the ICE framework library and provides a basic set of services common to the ICE hardware.

Applications are created by defining Python objects that represent each hardware element and its interface through their attributes and methods. Specific applications are built by extending the core objects (e.g. an `IceBoard` object that represents the ICE motherboard, an `IceCrate` object that represents the backplane, a `Mezzanine` object that represents a generic daughter board) provided by the *IceCore* package. As an example, the basic `IceBoard` object provides methods to program the FPGA, monitor control peripherals, access power systems to monitor voltages and currents, and measure system temperatures. Methods available on the ARM processor are also available and can be run as if they were local methods. The user then extends this `IceBoard` object to add the functions that are available in the specific FPGA firmware.

ARM Software

The ARM software provides network-based access to the ICE system's resources, programs and communicates with the FPGA, and runs high-level non-real-time applications locally.

The ARM runs a custom Linux kernel that contains all the drivers needed to access the hardware in the ICE motherboard. Once Linux has booted, the ARM can obtain an Internet Protocol (IP) address automatically using the Dynamic Host Configuration Protocol (DHCP⁶), announcing its presence and its surrounding hardware on the network, and starting the applications needed to access the system's resources. Access to the ARM functions is done through the *IceCore* package (see the control software sub-section above).

⁶https://en.wikipedia.org/wiki/Dynamic_Host_Configuration_Protocol

FPGA Firmware

The FPGA firmware is the most powerful (and complicated) part of the ICE firmware and software system since it is responsible of the real time signal-processing and associated networking for a given application. Although the FPGA firmware is mostly application specific, the ICE framework provides an FPGA library with core building blocks (e.g. a serial interface with the ARM processor, a basic clock module, and a GPS time encoder/decoder) to help build specific applications. Details of the custom firmware for the CHIME application are presented in Section 4.2.

4.2 The digitizer, F-engine, and networking engine of the CHIME correlator

In this section we describe the specialization of the ICE system for the CHIME high-bandwidth radio correlator, which is an FX design that uses the ICE system to implement the data acquisition, Fourier transform channelization (F-Engine), and the majority of the corner-turn networking.

4.2.1 Digitization

After the 2048 analog signals originating at the 1024 dual-polarization CHIME feeds are amplified and band-pass filtered to 400-800 MHz, they are fed to custom digitizer daughter boards on the ICE motherboards. The daughter boards specialize the motherboards for the radio interferometry application.




Figure 4.4: Photograph of two custom digitizer daughter boards (red) installed on a motherboard (blue). Each daughter board includes eight inputs that can digitize at up to 1.25 GSPS with 8 bits. For CHIME, these inputs sample at 800 MSPS in the second Nyquist zone.

As Figure 4.4 shows, each motherboard hosts two digitizer daughter boards. CHIME requires a total of 128 motherboards to digitize the signals from the 2048 receivers. Each digitizer daughter board has two EV8AQ160 ADCs and provides a total of eight inputs that can digitize at up to 1.25 GSPS with 8 bits⁷. For CHIME, these inputs sample at 800 MSPS in the second Nyquist zone. Each input provides more than 15 dB return-loss performance in the CHIME band. The insertion loss is approximately 2 dB. The ADCs support analog inputs with 500 mV peak-to-peak and provide up to 7.5 Effective Number of Bits (see Section 4.4.4).

The data from each of the 16 digitized inputs of a motherboard is processed by its FPGA with custom firmware. The first signal processing module performs the data acquisition from the ADC. The data from each input is acquired at 800 MSPS through 8 lines, one per bit. The module aligns the data acquisition of each line to compensate for the board and FPGA routing delays. Finally, the data is passed to the *channelizer* module as frames of 2048 8-bit samples.

4.2.2 Frequency channelization (F-engine)

A schematic diagram of the channelizer module for a single digitizer input is shown in Figure 4.5. It receives a data stream from the data acquisition module and passes it to the *function generator* sub-module. The function generator can pass the ADC data down the pipeline or it can insert user-defined packets or other test waveforms. The data stream is then passed to a custom PFB and FFT that uses the CASPER⁸ toolset. The PFB/FFT applies a sinc-Hann window to 4 consecutive data frames, and outputs a frame of 1024 complex values, one per frequency channel, in 18+18 bit format. The following *scaler* sub-module applies a 16+16 bit complex gain to each frequency channel and scales the result to 4+4 bit complex values. This complex gain can be configured by software. The output data from the scaler sub-module includes ADC, FFT and scaling saturation flags which is used to identify RFI. A separate *prober* sub-module within the channelizer can be configured to capture and buffer a subset of the raw ADC or scaler output data and send it back over the control channel for independent data monitoring.

⁷The EV8AQ160 converters have reduced channel modes which allow the sampling rate to be increased to 2.5 GSPS (for two channels per ADC) and 5 GSPS (for a single channel per ADC)

⁸<https://casper.berkeley.edu/>




Figure 4.5: Schematic diagram of the channelizer module. It receives a data stream from the data acquisition module and passes it to the function generator sub-module which can select its source stream from the ADC or from an integrated test pattern generator. The data stream is then passed to a custom PFB/FFT that outputs a frame of 1024 complex frequency samples in a 18+18 bit format. The following scaler sub-module applies a 16+16 bit complex gain to each frequency channel and scales the result to 4+4 bit complex values. A separate prober sub-module can be configured to capture and buffer a subset of the raw ADC or scaler output data and send it back over the control channel for independent data monitoring.

4.2.3 Corner-turn network

After the channelization, each ICE motherboard contains the information of 16 inputs for the 1024 frequency channels. However, each X-engine node requires the data from all the inputs and a subset of the frequency channels in order to compute all the cross-multiplications and averaging required to form the visibilities. Each CHIME GPU node can process four frequency channels. The *corner-turn* operation is the re-arrangement of the data from a configuration where all the information of one input across all the frequency channels is in one location at the F-engine (i.e. an ICE motherboard), to a configuration where all the information from one frequency channel across all inputs is in one location at the X-engine (i.e. a GPU node).

The total data rate, DR , of the corner-turn network is given by

$$DR = N \cdot R_{bits} \cdot BW \quad (4.1)$$

where N is the number of inputs, R_{bits} is the number of bits per frequency channel, and BW is the total frequency bandwidth. Note that DR is independent of the number of frequency channels produced by the F-engine. This is the case for any F-engine that implements Nyquist sampling of the input signals and a critically sampled PFB structure (that is, that each frequency channel is also Nyquist sampled) since, for a Nyquist sampling frequency $\nu_N = 2BW$ and an FFT of length L , there are $L/2$ (positive) frequency channels per input and the output rate for each frequency channel is $2BW/L$ in SPS, so the factor of $L/2$ cancels in DR . See [50, 32] for a detailed introduction to critically sampled PFB structures.

With $N = 2048$ inputs, $BW = 400$ MHz of total bandwidth, and $R_{bits} = 8$ bits per frequency channel (4 bits real + 4 bits imaginary), the CHIME corner-turn network has to re-arrange 6.6 Tbit/s of data before sending to the X-engine that computes the correlation matrix for each frequency channel. As a comparison, this data rate is greater than all of North America’s mobile data traffic in 2017 (5.8 Tbit/s on average [51]).

On the other hand, since there are $n_b = N(N + 1)/2$ baselines (including auto-correlations and the two polarizations), then the X-engine has to compute n_b complex multiply-accumulate (cMAC) operations per frequency channel and per sample from the F-engine, so the computational cost of the N -element correlation is

$$\eta = BW \cdot n_b. \quad (4.2)$$

Table 4.1 shows the bandwidth-baseline product for large contemporary radio interferometers. CHIME is the largest correlator that has been built as measured by η .^{9,10}

Table 4.1: Bandwidth-baseline products for large contemporary radio interferometers.

Instrument	Bandwidth (MHz)	Inputs	η (TcMAC/s)
CHIME	400	2048	839
ALMA*, [52]	8000	128	66
CHIME pathfinder	400	256	13
JVLA*, [53]	8000	54	12
LEDA†, [54]	58	512	8
MWA°, [55]	31	256	1

*Atacama Large Millimeter Array.

*Jansky Very Large Array.

†Large-Aperture Experiment to Detect the Dark Ages.

°Murchison Widefield Array.

Corner-turn architecture

Figure 3.7 shows the overall structure of the corner-turn architecture, while Figure 4.6 shows details of the hardware and FPGA firmware within each stage of the corner turn.

⁹The total computational cost for the CHIME correlator is dominated by the X-engine. The F-engine performs $(L/2)\log_2(L)$ cMAC operations every $L/(2BW)$ seconds (the time to obtain a frame) per input, so the computational cost for the F-engine is $BW \cdot N \cdot \log_2(L)$ operations per second. For CHIME, this value is about 9 TMAC/s, or about 1% of the total computational cost.

¹⁰Note that η alone does not necessarily reflect the total complexity or capability of each correlator. See [33] for a detailed discussion.

The data needs to be re-arranged such that at the end each of the 256 GPU nodes receives $1024/256 = 4$ unique frequency channels from the 2048 inputs.

The corner-turn operation is performed by the ICE system in four stages. The first stage is implemented within each FPGA, where each ICE motherboard digitizes and channelizes 16 inputs, for a raw data rate of 51.2 Gbit/s after the channelization stage. The FPGA creates 16 new data streams, each one having $1024/16 = 64$ frequency channels from each of the 16 inputs.

In the second stage, each motherboard within a crate is assigned 1/16th of the spectrum. Thus, each board keeps one of its stage-1 data streams and sends the remaining 15 streams to the other boards within a crate through the backplane's 10 Gbit/s mesh network. After this data exchange each ICE motherboard within a crate contains the data for a subset of 64 unique frequency channels, but all the 256 inputs of the crate. After adding flags and headers, this results in about 4 Gbit/s of data sent by each backplane link, resulting in about 1 Tbit/s of backplane full-mesh traffic, which is about 40% of the theoretical capacity of the network. After receiving the data from the other boards within a crate the FPGA re-orders the data into two streams, each one containing 32 frequency channels for the 256 crate inputs.

In the third stage of the corner turn, one of the two streams from stage two is sent to a sister motherboard in the adjacent (same rack) crate through the four links provided by one of the sixteen backplane QSFP+ connectors (see right photograph of Figure 4.3). This results in about 7.4 Gbit/s of raw data and flags that are transmitted over each link. The board from the other crate also sends one of its two data streams through the same link. After this inter-crate data exchange, each board within a crate pair contains the data for a subset of 32 unique frequency channels and 512 inputs. This corresponds to the data from one quarter of the CHIME array, or one cylinder. The FPGA then re-orders the data again into eight streams, each containing 4 frequency channels for 512 inputs. Although not used in CHIME, the ICE system also allows the formation of a full-mesh network between motherboards in up to five crates using the backplane QSFP+ links. These configurations are discussed in detail in [4].

In the fourth stage, the motherboard uses the eight links from its two QSFP+ connectors to send each of the eight stage-three streams to a different GPU node. Including flags, this corresponds to about 7.5 Gbit/s of data rate per link. Since the F-engine and the GPU X-engine are located in different buildings, the data of each board is sent over two active 100 m multi-mode optical fiber QSFP+ cables, each terminated into four independent SFP+ connectors. Each SFP+ connector is connected to a different GPU node, and each GPU node receives data from a board in a different crate pair that handles a




Figure 4.6: Schematic diagram of the corner turn. See Section 4.2.3 for a detailed description.

different cylinder, such that at the end it contains the data for a subset of four unique frequencies and 2048 inputs (see Figure 3.7). In the last corner-turn stage the GPU node recombines the data from its four 10 Gbit/s links to compute the correlation matrix for the 2048 inputs and four frequency channels. For details on the final corner-turn step refer to [45, 46, 47].

4.3 Installation of the ICE system for the pathfinder correlator

The construction of the pathfinder cylinder structures started in January 2013 and took approximately 9 months to complete. We installed the first set of 8 dual-polarization feeds and their respective analog receivers in September 2013 (see Figure 4.7). During that trip we also installed the first version of the pathfinder correlator, consisting of two independent 8-input prototype correlators, each implemented on a Xilinx Kintex-7 FPGA KC705 evaluation board that used one of our custom FMC digitizer daughter boards (see Section 4.2.1). For these correlators the number of channels was low enough that both the F-engine and X-engine could be implemented in the FPGA without the need of a separate GPU. These FPGA correlators send the averaged input products directly to a local control computer for further averaging and storage. Photographs of the two 8-input correlators and the first set of measured autocorrelation spectra are shown in Figure 4.8.

In March 2014 we replaced the two 8-input correlators by a single 16-input correlator consisting of a single ICE motherboard that implemented the digitization and F-engine, and a GPU node that implemented the X-engine. A photograph of the ICE system for this correlator is shown in Figure 4.9.

In September 2014 we finished the installation of pathfinder’s 256-input correlator, consisting of a fully populated ICE crate and a cluster of 16 GPU nodes. Photographs of the pathfinder ICE crate after the installation are shown in Figure 4.10. This system initially operated as 16 independent 16-input correlators. By December 2014, the implementation of the second-stage corner-turn within the FPGA firmware was completed. The system operates as a 256-input correlator since then.

4.4 Validation of the ICE system on the pathfinder

After the installation of the pathfinder ICE crate we performed a full end-to-end characterization and validation of the ICE system on the pathfinder correlator to demonstrate



Figure 4.7: Installation of the first cassette on the west cylinder of the pathfinder in September 2013. Each cassette hosts 4 dual-polarization feeds. The second cassette was installed next to the first one.

that the system complies with all the requirements for CHIME and that the hardware was mature enough to proceed with the production of all the motherboards, daughter boards and backplanes for full CHIME and other experiments (e.g. HIRAX). The requirements were grouped in the following broad categories:

- **Manufacturability:** The circuit boards must be demonstrated to be manufacturable with a reasonable yield and cost.
- **Uptime and reliability:** The firmware and hardware that operate CHIME must be reliable so that the system can continually run with minimal monitoring and maintenance once configured. It should run for extended periods without failure. The circuit boards should not degrade with time in their normal operating environment.
- **Data transport integrity:** The system must have a low bit error rate (BER) and packet loss during all the stages of the corner turn. Lost packets must be identified and accounted for by the system.
- **Analog performance and data acquisition integrity:** The analog performance of the digitizer boards must not be a limiting factor for science goals. The data acquisition system must guarantee error-free acquisition of the ADC data.
- **Timing:** The system must provide accurate absolute time. The system clock jitter





Figure 4.8: Top: Photograph of the two independent 8-input correlators installed on the pathfinder in September 2013. The correlators use a Xilinx Kintex-7 FPGA KC705 evaluation board that used one of our custom FMC digitizer daughter boards. Both the F-engine and the X-engine are implemented on the FPGA. Bottom: Photograph of the first set of autocorrelation spectra obtained with the 8-input correlators. During the installation we found a defective analog receiver (green line), which was fixed afterwards.



Figure 4.9: Photograph of the ICE-based F-engine for the 16-input correlator installed on the pathfinder in March 2014. The X-engine was implemented on a separate GPU node.




Figure 4.10: Left: Side view of the ICE crate installed on the pathfinder for the 256-input correlator. Only 32 analog receivers are connected. Right: Bottom view of the ICE crate after connecting the 256 analog receivers.

must meet the requirements for cosmology and additional science (e.g. FRBs, pulsar timing) goals.

- **Power consumption:** The system must have power consumption that is commensurate with the power supplies we plan to use, and be able to operate within component temperature tolerances with the convective cooling from the fans we have selected.

4.4.1 Manufacturability

Requirements

The pathfinder correlator requires a fully populated ICE crate (one backplane, 16 motherboards and 32 digitizer boards) to implement the data acquisition, F-engine, and corner-turn (up to second-stage) operations. For full CHIME, there will be 8 copies of this hardware. The manufacturability requirements for full CHIME were:

1. Successful manufacture and operation of the pathfinder hardware.
2. Demonstration of manufacturability with adequate (>85%) yield.

Performance

The hardware quality control (QC) process consists of:

1. ***Assembly house QC:*** When the boards are manufactured, the assembly house performs basic QC checks as part of the of their internal production process, including visual inspection and X-ray checks of critical regions of the boards. In order to further qualify the boards, we developed a suite of functional tests that the assembly house performs on the boards before they are shipped to McGill. These tests include impedance checks on the power rails, board power up, and ARM memory, boot and networking tests.
2. ***McGill Cosmology lab QC:*** We have developed a Python-based custom QC software that exercises all the functionalities of each board. The software runs on a host computer that connects to the board under test, runs the tests automatically (or directs the operators for manual tests like visual inspections), analyzes the results, records, and archives the results. The tests include board identification and configuration, ARM initialization, FPGA firmware programming, power and temperature measurements, iceboard and backplane transceiver communications, digital delay calculation for the ADC daughter boards, among others.

3. ***Board tracking system:*** During the first test of every board, a data file is automatically generated. This file is used to keep track of all the tests results, as well as repairs and modifications made to the board. This file is automatically pushed in a QC Git¹¹ repository accessible to all board users.

After reaching final revision versions for each board, but before implementing the assembly house QC software, we produced an initial batch of 40 ADC daughter boards and 40 ICE motherboards for the manufacturability test and pathfinder correlator deployment, of which 38 (95% yield) daughter boards and 36 (90% yield) motherboards were operational. Problems with defective boards were mainly due to soldering issues during assembly. These problems were easily identified using the McGill Cosmology lab QC software and fixed by the assembly house. After the implementation of the assembly house QC software we produced a second batch of 90 boards with 100% yield. By the end of the test we had produced 10 backplanes with a 100 % yield.

By the time we finished the validation of the ICE system, the ICE crate installed on the pathfinder (one backplane, 16 motherboards, and 32 daughter boards) had been operating for more than 6 months with no faults, while the rest of the boards has been used extensively in the McGill Cosmology lab and other experiments with no faults.

4.4.2 Uptime and reliability

Requirements

We quantify the uptime and reliability requirements as follows:

1. A fully populated ICE crate should run the 256-input F-engine and corner turn for a test period of at least one week without needing to be restarted or significant data corruption.
2. The entire system should be reliably configured and booted as desired without additional checking and interaction from a user.
3. A fully populated ICE crate should operate in the pathfinder RF room environment for many months without significant board failures or reliability issues that are not understood.

¹¹<https://git-scm.com/>

Performance

During the week of December 20-27 2014 we tested the reliability of the ICE system by running a fully populated ICE crate in 256-input mode inside the pathfinder RF room uninterruptedly and without user interaction. At the end of the testing period the system was still running to specifications.

Additionally, the pathfinder correlator has been routinely run for many days before stopping it without incident, including several 3-week long run periods. The interruptions are typically driven by testing needs, pathfinder hut cooling, and other external factors, but not by ICE system failures or crashes. We operate (start, stop, configure) the ICE system and the correlator remotely and reliably without issues.

The system meets reliability requirements. We have run the pathfinder correlator for three years now without issues that can be attributed to the ICE system hardware and software that are not understood.

4.4.3 Data transport integrity

Requirements

1. Packet loss must result in lower than 1% total data loss. Note that this requirement is much more stringent than 1% packet loss, since losing data from one antenna will result in losing correlations between that antenna and all others.
2. Malformed or lost packets must be identified and accounted for by the overall system.
3. Data frames must be aligned in time to preserve correlation across the system. The digital data transport must not degrade the correlation between inputs.

Performance

The corner-turn system that implements the data transport for CHIME is described in Section 4.2.3. Each corner-turn stage is implemented using the same firmware module that is configurable for each stage. Each stage starts with a frame alignment module that ensures the frames are aligned when passed to the next processing stage. The module has a nominal waiting time after which frames that are not received are considered missing and the remaining data is forwarded with the respective flags.

Each packet contains a header that indicates its geometry (number of channels and inputs), origin (stage, crate, slot and lane), and a GPS timestamp. The packet also has

a trailer with flags indicating whether packets were lost or corrupted during the various stages of the corner turn. Data transfers between motherboards use a custom-written version of the 10 GbE Ethernet protocol that implements a cyclic redundancy check (CRC) to validate the integrity of the packets. The data sent to the GPUs in the last corner-turn stage use standard 10 GbE Ethernet and UDP packets.

The verification of the backplane full-mesh network signal integrity was performed by transmitting a pseudorandom binary sequence (PRBS) between every motherboard in a crate and performing an eye diagram for each link. The results for a particular backplane test are shown in Figure 4.11. The diagram consists on varying the sampling point both in time (x-axis of each mini-plot) and signal amplitude (y-axis of each mini-plot) by discrete amounts and then measuring the BER relative to the optimal sampling point. A good link has a large, well opened center region (blue region of each mini-plot, this is the ‘eye’) where the BER is low. Figure 4.11 shows that, except for the link corresponding to transmit slot 12 and receive slot 8 (indicated with a yellow rectangle), all the links have well opened eyes. The figure also shows how a bad link can be identified. In this case, the high-BER link was caused by a motherboard connector assembly issue that was easily identified in this manner and fixed.

The inter-crate QSFP+ links (the ones used in the third stage of the corner turn) are also tested by measuring their BER with a PRBS sequence. After measuring the BER and adjusting the transmit power of each link to the optimal level we obtain error-free transmission during many-second tests, which implies a BER lower than 10^{-11} .

Finally, the communication between the ICE system and the X-engine was tested with a fully-populated ICE crate pair (each pair can handle data from one quarter of the CHIME array) configured to send data to the four 10G Ethernet ports of one GPU. The channelizer of each input (see Figure 4.5) was configured to send digital test patterns instead of analog samples from the ADC. The performance was measured by counting the number of packets that were rejected due to CRC errors, flagged by the corner-turn firmware, and also by verifying that the test pattern arrived to the right link in the right order. We obtained error-free transmission during many-minute tests, implying that we did not detect errors on the backplane, inter-crate, and GPU corner-turn stages (see Figure 4.6).

These tests demonstrate that the system meets the data transport integrity requirements for CHIME.

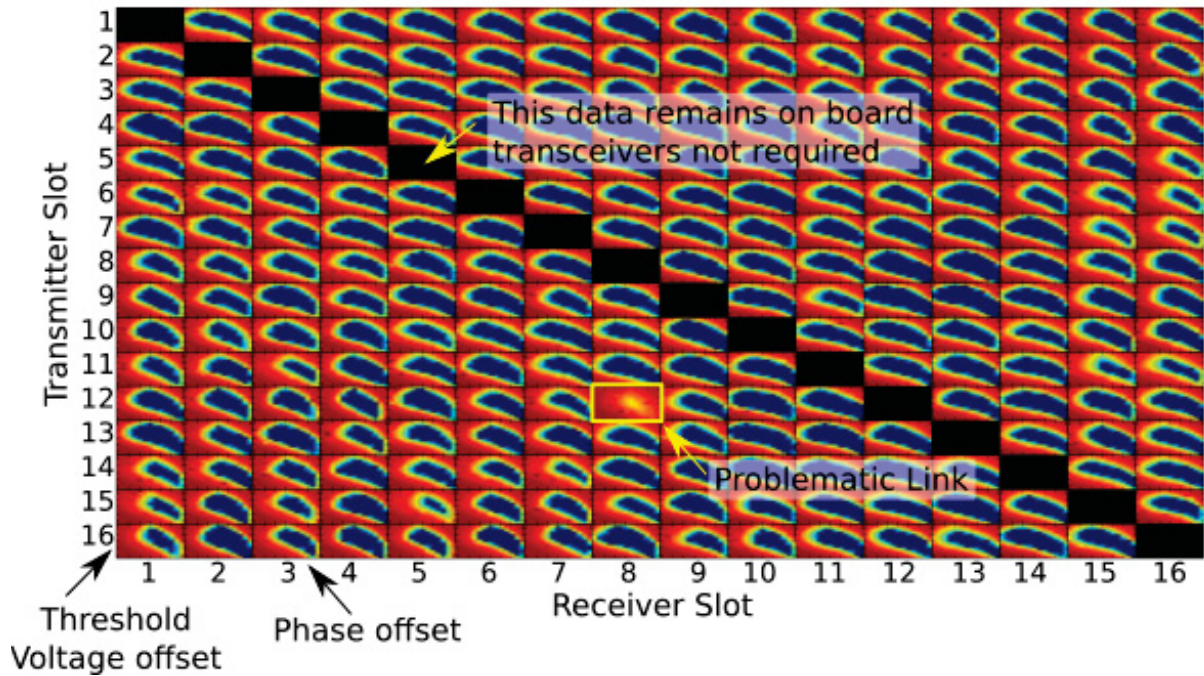


Figure 4.11: Example of an eye diagram of the backplane 10 Gbit/s links between all 16 motherboards in a crate. The diagonal has no data because the link between a board and itself is done internally. The diagram is performed by varying the sampling point both in time (x-axis of each mini-plot) and signal amplitude (y-axis of each mini-plot) by discrete amounts and then measuring the BER relative to the optimal sampling point. A good link has a large, well opened center region (blue region of each mini-plot, this is the ‘eye’) where the BER is low. One defective link at transmit slot 12, receive slot 8 was revealed by this particular test. The poor performance of this link was caused by a motherboard connector assembly issue that was identified in this manner and fixed.

4.4.4 Analog performance and data acquisition integrity

The analog performance of the ICE system for the CHIME correlator is determined by the ADC and its surrounding digitizer board and motherboard electronics. The performance characteristics of the ADC chip are well defined by the manufacturer. Since the chip was selected to meet our cosmology requirements, the requirements for the overall electronic system is to not significantly degrade the ADC performance itself.

Crosstalk is another analog property that can be affected by ADC performance. For CHIME, adjacent feeds can have cross-coupling at the ~ -15 dB level, while feeds on different cylinders have cross-coupling below ~ -70 dB. The crosstalk contribution from the digital backend electronics needs to be small in comparison to other contributors.

Once the analog waveforms are digitized by the ADCs, the digitized values are sent to the FPGA at a very high speed (800 MSPS) through the mezzanine connectors using an 8-bit wide differential bus. The timing of each bit is motherboard and daughter board dependent and must be adjusted carefully to ensure proper acquisition of the digitized values. We require that the FPGA data acquisition system guarantees error-free acquisition of the ADC data.

Requirements

1. ***Digitizer-board dynamic performance:*** The dynamic performance of the digitizer boards must be consistent with the datasheet for the EV8AQ160 8-bit ADC¹², (i.e. is not substantially degraded by the board design) and must not be a limiting factor for the cosmology goals.
2. ***Digitizer-board Crosstalk:*** The crosstalk at the digitizer-board inputs must not be a limiting factor for the crosstalk for the system. Since we are limited by the crosstalk performance internal to the ADC devices, we require minimum degradation of the crosstalk performance at the digitizer-board analog inputs with respect to the ADC specification.
3. ***ADC Data acquisition Integrity:*** The system must ensure reliable data acquisition from ADC to FPGA.

¹²The EV8AQ160 datasheet can be found on <https://www.e2v.com/products/semiconductors/adc/ev8aq160/>

Performance

Digitizer-board dynamic performance: This was determined by measuring standard ADC dynamic performance parameters and comparing to the ADC specification. A summary of the digitizer-board performance is shown in table 4.2, with a comparison to expected ADC performance from the datasheet (which in the CHIME band is given at 620 MHz only).

Table 4.2: Dynamic performance of the digitizer board in the CHIME band (400-800 MHz) and comparison to the ADC specs (given at 620 MHz only).

Parameter	Value	ADC spec
Signal-to-Noise Ratio (SNR)	45.3 - 45.5 dBc	43-46.5 dBc
Total Harmonic Distortion (THD)	47 - 52 dBc	46-56 dBc
Effective Number of Bits (ENOB)	6.9 - 7.2 bits	6.9-7.5 bits
(Input) Third-order Intercept Point (IIP3)	13 dBm	11 dBm

For a tutorial on the specifications for quantifying the ADC dynamic performance see [56]. An example of the method used to measure the Signal-to-Noise Ratio (SNR), Total Harmonic Distortion (THD), and Effective Number of Bits (ENOB) at 581 MHz is shown in Figure 4.12. A 1 dB below full scale (1 dBFS) tone is fed to one of the digitizer-board inputs and the power spectrum is computed. The SNR is the ratio of the input signal power (the largest peak at 581 MHz), S , to the integrated white noise power within the sampled band (cyan line), N . The THD is the ratio of S to the total power in the harmonics of the input signal that fold into the sampled band, T (smaller peaks with red asterisks). For this measurement we use the first 25 harmonics. The ENOB is calculated as

$$\text{ENOB} = \frac{\text{SINAD} - 1.76}{6.02} \quad (4.3)$$

where $\text{SINAD} = S/(N+T)$ is the signal-to-noise and distortion ratio in dBFS. As Table 4.2 shows, all these values are within the ADC specifications.

The (input) third-order intercept point (IIP3) is a figure of merit of the digitizer-board linearity. A detailed explanation of this parameter and its measurement for the digitizer board is presented in [1]. Briefly, we inject two tones at relatively close frequencies into one of the digitizer-board inputs and measure the resulting third-order intermodulation (IMD3) products, which increase 3 dB per every dB increase in the power of the fundamental tones. The intersection of the third-order line produced by the IMD3 products




Figure 4.12: Power spectrum for a digitizer-board input when feeding a -6 dBm (1 dBFS) 581 MHz tone. The largest peak in the figure is the input tone. Smaller peaks with red asterisks denote spectral lines that are harmonics of this tone folded into the 400-800 MHz band. Other small peaks are mainly due to RFI, differential nonlinearity of the ADC (deviations of the ADC transfer function from its ideal behavior) and sampling-clock spurs from the on-board crystal used for this test. The cyan line denotes the white noise floor.

with the line produced by the linear term determines the IIP3, which we measured to be ~ 13 dBm. This measurement is shown in Figure 4.13. The 2 dB degradation in the IIP3 of the digitizer board with respect to the ADC specifications is due to the insertion loss of the board’s single-ended-to-differential conversion that precedes the ADC chip and is consistent with our design expectations.

Crosstalk: For the crosstalk tests, a -6 dBm tone (~ 1 dBFS) is applied to one of the digitizer-board inputs and the power is measured for the remaining inputs of the motherboard as the tone is swept in frequency. On the remaining inputs we also feed a low amplitude tone (~ 1 bit RMS at 100 MHz) at a fixed frequency to exercise a few least significant bits (LSBs) and smooth the ADC spectrum. The crosstalk measurements with respect to the input 6 of the motherboard are shown in Figure 4.14.

For inputs within the same ADC chip (second plot from left to right), the measured crosstalk at 620 MHz (-51 dB, -54 dB, and -70 dB) is broadly consistent with the -55 dB number provided in the datasheet. The circuit board is providing another path for coupling, which degrades the performance by a few dB, as expected. Note that these

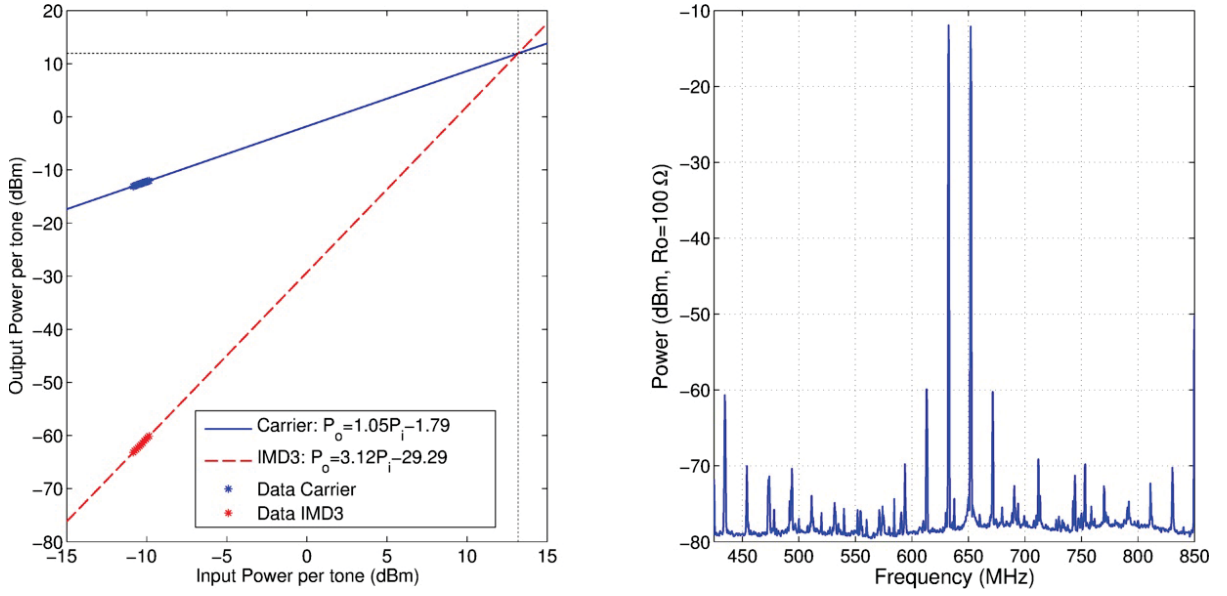


Figure 4.13: Left: Measurements of the power of the IMD3 products (red dots) and fundamental tones (blue dots) as function of the power of the input tones. The IMD3 products increase 3 dB per every dB increase in the power of the fundamental tones. The intersection of the third-order line produced by the IMD3 products with the line produced by the linear term determines the IIP3, which we measured to be ~ 13 dBm. Right: ADC spectra for two input tones 7 dBFS. The fundamental frequencies are $\nu_1 = 632.5$ MHz and $\nu_2 = 652$ MHz. The IMD3 products appear at $2\nu_1 - \nu_2 = 613$ MHz and $2\nu_2 - \nu_1 = 671.5$ MHz. The other tones are mostly harmonics that fold into the sampled band (for this particular test the sampling frequency was set to 850 MSPS).




Figure 4.14: Measurement of the digitizer-board crosstalk. A 1 dBFS tone is applied to the input 6 of the motherboard the power is measured for the remaining inputs of the motherboard as the tone is swept in frequency. On the remaining inputs we also feed a low amplitude tone (~ 1 bit RMS at 100 MHz, that folds into the 400-800 MHz band at 700 MHz) at a fixed frequency to exercise a few LSBs and smooth the ADC spectrum.

measurements provide more detailed information than the datasheet about the crosstalk performance across the CHIME band, which rises at high frequency up to -46 dB for neighboring inputs at 800 MHz. This is within the crosstalk performance requirements since inputs within the same ADC chip are connected to adjacent feeds which have a crosstalk level which is three orders of magnitude higher.

Crosstalk was also measured across the CHIME band for other input-to-input couplings (see Figure 4.14):

- Crosstalk between inputs arising from different ADC chips situated on the same digitizer board (left plot): $\lesssim -61$ dB across the band, with maximum coupling at 475MHz. Across most of the band, the measurement of -70 dB is an upper limit only.
- Crosstalk between inputs arising from different ADC chips on different digitizer boards from the same motherboard (third and fourth plots from left to right): $\lesssim -70$ dB. Across the band, this represents an upper limit.
- Crosstalk between inputs on different motherboards is expected to be lower yet, and measurements indicate the same upper limit of $\lesssim -70$ dB.

Overall, the crosstalk performance is consistent with expectations from the ADC chip and a slight expected degradation (4 dB max) from the circuit boards coupling. It is consistent with our design expectations.

The cross talk levels become lower as the participating analog inputs are separated by circuit boards. Since for full CHIME we use adjacent ADC inputs for feeds with close spatial separation as we have done for the pathfinder, this crosstalk will not be a limiting factor for the system.

ADC Data acquisition integrity: This requirement refers to sampling integrity and digital transport from ADC to FPGA. The FPGA firmware/software provides all the resources to configure the ADC acquisition. This includes

1. configuring the ADC's digital output to consist of (a) samples of the analog input or (b) internal test functions like ramps or pulse trains,
2. measuring the eye diagram of the ADC digital data lines and computing and programming optimum delays for each line to ensure reliable data acquisition,
3. and testing the performance of each ADC bit (e.g. check for stuck or faulty bits).

The data transport and integrity reliability from ADC to FPGA is demonstrated by first measuring the eye diagram and then finding the optimal digital delay of each bit that ensures that all the bits of a given sample are measured correctly and within the same clock edge. This is done by configuring the ADC to generate and send pulses (set all the bits to 1) at regular intervals instead of sending analog sampled data. After the bit delays are set, the ADC is configured to generate full scale triangular ramp functions. The output at the FPGA side (before the FFT) is checked to verify that all the ADC bits are exercised and by the right amount (for a ramp input, the output histogram must be flat, meaning all the LSBs occurred the same number of times) and in the right order (meaning that the delays were set properly).

We have found working delays with this method for the 2048 bits operating in the pathfinder (256 8-bit ADCs). The delay settings have proven to be quite stable over time and operating temperature. We have not found a meaningful variation in the computed delays over operating temperature or time. The delays only need to be calculated once.

4.4.5 Timing

Although the cosmology goals alone do not pose a very stringent requirement on the accuracy of absolute time, our calibration strategy includes using known pulsars as they transit through the CHIME beam to measure the primary beam of each antenna, which requires sub-millisecond absolute timing. The FRB and pulsar science requires microsecond level absolute timing, whereas measurements of arrival times for pulsar timing arrays are typically at the 100 ns RMS level [57]. The latter motivates achieving absolute timing at the 2 – 20 ns level.

Jitter specifications are driven by two requirements: (a) the ADC requires a jitter performance of ~ 0.5 ps RMS (typically on millisecond time scales) to achieve its dynamic range, and (b) jitter affects the sensitivity of the interferometer¹³. The ADC requirements are more stringent than the interferometry ones.

¹³If the voltage signals from antennas 1 and 2 are represented by $v_1 e^{i\phi_1(t)}$ and $v_2 e^{i\phi_2(t)}$, where the ϕ terms are phase errors due to jitter, then the measured correlation is $r_{12} = \langle v_1 v_2^* e^{i\delta\phi(t)} \rangle = v_1 v_2^* \langle e^{i\delta\phi(t)} \rangle$ where $\phi_1(t) - \phi_2(t) = \delta\phi(t) = 2\pi\nu \cdot \delta\tau(t)$ is the relative phase error due to jitter. If $\delta\tau(t)$ is Gaussian with zero mean then

$$\frac{r_{12}}{v_1 v_2^*} = e^{-(2\pi\nu \cdot \sigma_\tau)^2 / 2}$$

Where σ_τ is the jitter RMS. At CHIME's highest frequency (800 MHz), a 1 ps RMS timing jitter results in a signal loss below $\sim 10^{-5}$.

Requirements

1. ***Sync:*** After synchronizing the ICE system and starting the data acquisition, all inputs must start sampling within a few 1.6 GHz clock edges from one another, and this edge must be measurable through the calibration. We note that, although it is not an absolute requirement, having each channel start on a reproducible relative clock edge after each turn-on would simplify the phase calibration of the instrument.
2. ***Absolute time:*** In order to calibrate on second scale pulsars, the sync absolute time must be accurate to milliseconds. For millisecond pulsar observations with gating, absolute time to microseconds required.
3. ***Jitter:*** Clock jitter of the entire system must be less than ~ 0.5 ps on millisecond time scales to meet the CHIME calibration requirements. This jitter requirement is system wide (backplane, motherboards, daughter boards, and ADCs). Since directly measuring jitter at this level is very challenging and we do not have access to equipment capable of measuring this directly, we use a simplified version of the BIS calibration system described in Chapter 5 to infer the jitter performance.
4. ***Longterm timing:*** Timing accuracy must be maintained across multiple days of observation. Phase shifts must be slow enough that they can be calibrated on the integration cadence.

Performance

Sync: We have measured the quality of the ICE system synchronization by using the Broadband Injection Signal (BIS) calibration system described in Chapter 5. First, the system is calibrated and the relative phases between the 256 inputs are determined. After that, the correlator is restarted (the correlator is stopped and the ICE system is power-cycled and reconfigured) and the system is calibrated again. If all the inputs start the sampling on the same clock edge then the relative phases are preserved up to the measurement error. An input that starts on a different clock edge will have a jump in its relative phase that appears as a linear phase offset in frequency, and the slope of that phase determines the number of clock edges ‘missed’ by that input. Figure 4.15 shows the change in the relative phase as function of frequency for each of the pathfinder inputs after a correlator restart as measured by the BIS system. Each color represents the change in phase for a different input. For this particular restart, there are four inputs that show a jump in their phase consistent with a one-sample delay.

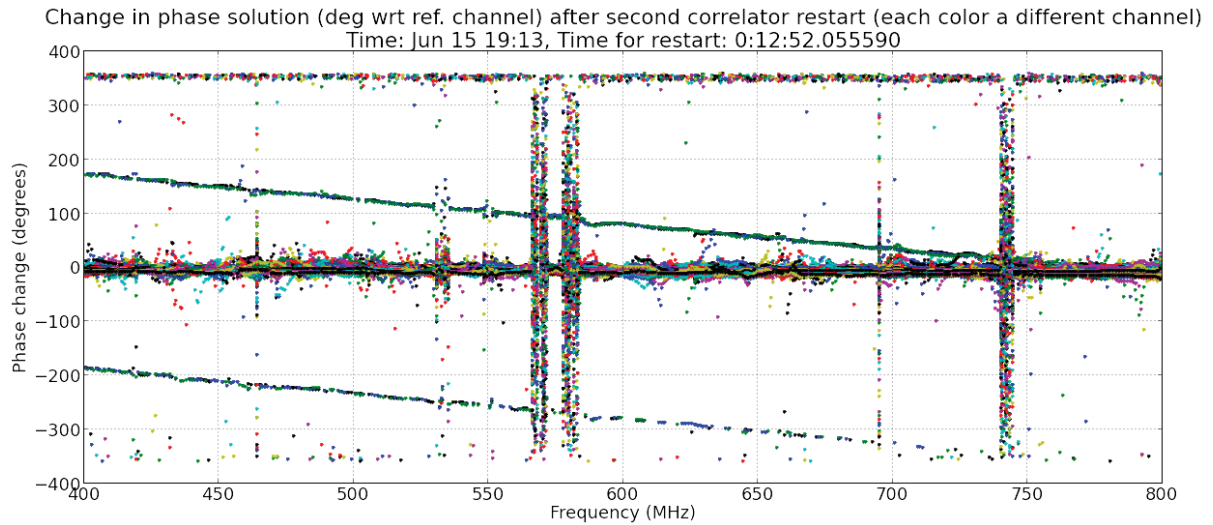


Figure 4.15: Change in the relative phase as function of frequency for each of the pathfinder inputs after a correlator restart as measured by the BIS calibration system. Each color represents the change in phase, in degrees, for a different input. Vertical stripes represent frequencies contaminated with RFI for which the measurement is not reliable. Points at $\pm 360^\circ$ correspond to either inputs flagged as ‘bad’ or frequency channels corresponding to GPU nodes which are down. For the vast majority of the inputs the relative phase is preserved up to the measurement error (which is approximately 4° for feeds well illuminated by the BIS system). For this particular test there are four inputs that have a jump in their relative phase that appears as a linear phase offset in frequency consistent with a one-sample delay with respect to the rest.

After repeatedly restarting the pathfinder, the vast majority of the 256 ADC inputs did not see a jump in their relative phases. In 2 out of 4 repeated restarts, a single ADC chip (4 inputs) out of the 64 in the pathfinder started one sample off relative to the others. The ADC that was off by one sample was different in the two instances. Note that we have not measured any jumps in the synchronization after the system has been started and has been running. This means that the relative phases between inputs are stable.

The phase of the sync edge that is distributed to each digitizer board is controlled by the FPGA. Since there is just one sync signal for each digitizer board which itself has two ADC chips, the phase needs to be optimized as a compromise for the best value for the pair of ADCs on each digitizer board. The ICE system synchronization can be optimized further and firmware/software work is ongoing to improve the quality of the syncing.

Absolute time: Absolute time is provided by a GPS source and distributed to each motherboard across the backplane. Each board within a crate measures absolute time independently, and these measurements agree to 10 ns, which is the current accuracy of the GPS timestamps set by the firmware. We intend to add extra digits to the timestamp that is encoded with the data, since this may be useful for pulsar timing array data.

The absolute timing offset between the GPS timestamp of a sample and the time when the sample was digitized has been measured by sending a pulsed signal from the GPS source to the inputs of a motherboard. The board is synced to start collecting data on an integer GPS second and each pulse of the input signal also coincides with an integer GPS second. The absolute timing offset for a given input is determined by locating the pulse in the raw ADC data and comparing the timestamp of the respective sample to the nearest integer second. The absolute timing offset is in the range 133 – 155 ns for all the inputs, and the timing difference between two inputs is repeatable to within about 1.25 ns. This meets expectations and exceeds requirements. These offsets only need to be calculated once.

Jitter: For a detailed review of jitter and phase noise refer to [56]. The first jitter measurement on a fully populated ICE crate was performed in the pathfinder using a simplified version of the BIS calibration system described in Chapter 5. The setup is shown in Figure 4.16. A BIS source consisting of a terminated LNA connected to a FLA is installed in the pathfinder RF room. The signal is split into 16 copies and sent to the same ADC input on the 16 different FPGA motherboards of the crate. We then collected ~ 11 h of cross-correlation data overnight at 20 s cadence. After that, we used the BIS calibration algorithm to determine the phase of each of the 16 ADC inputs (one

per board) as function of time and frequency.




Figure 4.16: Setup for jitter measurement on the pathfinder ICE crate. A broadband noise source consisting of a terminated LNA connected to a FLA is installed in the pathfinder RF room. The signal is split into 16 copies and send to the same ADC input on the 16 different FPGA motherboards of the crate.

The clock jitter and drift¹⁴ can be modeled as a time varying delay between the signals input to different motherboards. For jitter, the phase fluctuations $\delta\phi$ are related to the time-delay fluctuations $\delta\tau$ as

$$\delta\phi(\nu, t) = 2\pi\nu \cdot \delta\tau(t). \quad (4.4)$$

Since the size of the phase fluctuations scale linearly with frequency, a simple linear regression¹⁵ can be used to construct a delay template $\delta\hat{\tau}(t)$ for each board by using all the frequency information from the phase at time sample t . The jitter-induced phase fluctuations are then corrected by subtracting the phase fluctuations $\delta\hat{\phi}$ estimated from these delay templates from the visibility data. This is done in a frequency-by-frequency basis.

The power spectral density (PSD) of the measured delay fluctuations for each motherboard is shown in Figure 4.17. The thick white line is the PSD of the temperature measured by a thermometer installed near the FPGA crate. The x-axis covers timescales between 1 hour and 20 seconds. The dashed vertical lines denote timescales of 15, 6, and 3 minutes from left to right. Note that the features in the temperature PSD are repro-

¹⁴Here we call jitter the fast phase fluctuations and drift the slow ones.

¹⁵https://en.wikipedia.org/wiki/Simple_linear_regression

duced in the delay PSD, in particular the two peaks at 5 and 8 minutes corresponding to the RF-room air-conditioner cycle.




Figure 4.17: PSD of the delay fluctuations for each motherboard in the pathfinder ICE crate. The thick white line is the PSD of the temperature measured by a thermometer installed near the FPGA crate. The x-axis covers timescales between 1 hour and 20 seconds. The dashed vertical lines denote timescales of 15, 6, and 3 minutes from left to right. Note that the features in the temperature PSD are reproduced in the delay PSD, in particular the two peaks at 5 and 8 minutes corresponding to the RF-room air-conditioner cycle. This analysis was a team effort and this figure was made by Seth Siegel.

The jitter is the RMS of the delay fluctuations, which is obtained by integrating the delay PSD from Figure 4.17. The RMS of the delay fluctuations on different time scales as function of the position in the ICE crate is shown in Figure 4.18. Each color corresponds to a different lower bound of the PSD integration (the upper bound is 20 s which is the maximum temporal frequency probed). Note that the RMS is minimum towards the center of the crate, and increases near the edges¹⁶. The jitter is also highly correlated with the FPGA die temperature of each particular motherboard, which depends not only on the slot-by-slot air flow variations in the crate that are caused by fan edges lining up with boards slots, but also on the thermal coupling between the FPGA die and its heat sink.

Note that RMS of the delay fluctuations on time scales below 3 minutes is below ~ 0.6 ps. On millisecond time scales the RMS of the delay fluctuations is well within

¹⁶For this calculation all the phases were referenced to the mean value over the 16 motherboards. Changing the reference will change the delay RMS profile in Figure 4.18.

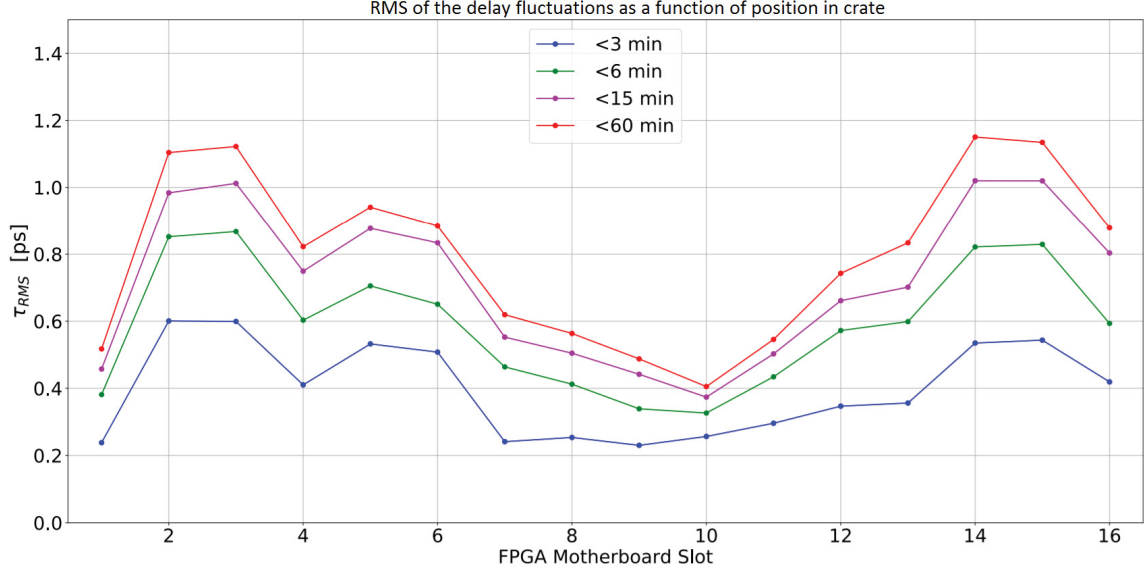


Figure 4.18: RMS of the delay fluctuations as a function of position in crate. Obtained by integrating the delay PSD. Each color corresponds to a different lower bound of the PSD integration (the upper bound is 20 s which is the maximum temporal frequency probed). Note that RMS of the delay fluctuations on time scales below 3 minutes is below ~ 0.6 ps. On 1 minute time scales the RMS of the delay fluctuations is even lower, well within the jitter requirements. This analysis was a team effort and this figure was made by Seth Siegel.

the jitter requirements. However, on minute time scales this sub-picosecond jitter is still an important source of phase fluctuations that must be corrected in order to meet the CHIME calibration requirements (see Chapter 5). We are currently using this setup on the pathfinder to measure and correct the intra-crate timing jitter. Also, an enhanced system based on this design is currently installed on full CHIME to track and correct the intra and inter-crate jitter of the ICE system.

Longterm timing: Long time-scale measurements of input-to-input phase have been made with the pathfinder, using both the BIS calibration system and the jitter measurement setup described above. As explained in the jitter sub-section, the long time-scale phase variations due to the clock drift are consistent temperature changes on the FPGA dies. No large, sudden changes in phase (as might be expected from a digital system if, for example, a clock edge was missed) that can be attributed to the ICE system have been observed. The long term variation is at the sub-degree level. The CHIME calibration system is designed to measure these variations.

4.4.6 Power consumption

Requirements

1. The system is designed to be powered by one Agilent N5764A power supply per crate, capable of delivering 1520 W at 20 V. The power consumption for one crate fully populated ICE crate operating the 256-input channelizer and corner-turn firmware must be commensurate with this power supply, \lesssim 1300 W typical.
2. The FPGA internal temperature of the warmest board must be < 85 C to meet the recommended operating conditions of the FPGA¹⁷.

Performance

The power consumption of the ICE crate operating on the pathfinder with 20 V input voltage draws approximately 1250 W after configuring the ICE system for 256-input mode operation, meeting the power consumption specification.

The FPGA temperatures are continuously monitored. With the pathfinder air-conditioning system fully operational, the temperature of most FPGA dies is within the requirements. During the hottest days of the year, the coolest FPGAs are well below ~ 80 C, while the warmest FPGA can reach ~ 88 C. As explained before, the FPGA temperature depends on both the particular airflow profile of the slot in which the motherboard is located and the effective thermal coupling of the FPGA die to its heat sink. Although the FPGA specification recommends operating conditions below 85 C, we have not yet seen any degradation in the performance of the warmest FPGAs. However, continuous operation under these conditions can reduce their lifetime. With this in mind, the chilling system of the full-CHIME FPGA rooms has been improved to further cool the FPGAs and meet the temperature requirements.

4.5 Installation of the ICE system for the CHIME correlator

A simplified schematic of the full CHIME site layout is shown in Figure 4.19. The ICE-based F-engine is installed inside two separate receiver huts (RF-shielded rooms installed within shipping containers). Each FPGA receiver hut serves two cylinders and is placed

¹⁷The Kintex-7 FPGA datasheet can be found on https://www.xilinx.com/support/documentation/data_sheets/ds182.Kintex-7_Data_Sheet.pdf

between them. This configuration minimizes the receiver coaxial-cable length. The GPU-based X-engine is housed in two RF huts adjacent to the east-most cylinder. The 256 optical fibers that communicate the FPGA and GPU engines (each multifiber contains four fiber strands that connect to four different GPU nodes) are carried within a cable tray that goes underneath the cylinders and above the huts. Also within the cable tray are the coaxial cables that distribute the clock and absolute time signals to the FPGA huts. The FRB backend is housed in a separate RF hut next to the GPU huts, as is the Power Distribution Center (PDC) that provides the power for the full experiment. The cosmology pipeline backend and the pulsar timing backends (see Figure 3.3) are housed in a shielded room in the DRAO building and are not shown in Figure 4.19.




Figure 4.19: Simplified schematic of the full CHIME site layout. The ICE-based F-engine is installed inside two separate RF huts located between each pair of cylinders to minimize the receiver coaxial-cable length. The GPU-based X-engine is housed in two RF huts adjacent to the east-most cylinder. The FRB backend and the Power Distribution Center (PDC) are housed in a separate RF huts next to the GPU huts.

The construction of the full CHIME cylinder structures started in January 2015 and took approximately 9 months to complete. The installation of the FPGA and GPU RF-rooms within the shipping containers was completed by May 2016 and the campaign to install the full CHIME correlator started in September 2016. The installation of the ICE hardware, cooling system and networking in the east receiver hut was completed in May 2017. We installed the FPGA-GPU fiber tray in December 2016, and we routed and

connected all the data fibers and timing cables between April and July 2017. By August 2017 all the ICE hardware for CHIME had been deployed and the first on-site operational tests with the full (east and west) ICE system had been performed. Photographs of the ICE hardware within the FPGA huts at different stages of the installation are shown in Figure 4.20. The installation of the GPU hardware ran in parallel with that of the ICE system. We observed the first fringes from a CasA transit with the full correlator on August 31st 2017. This is shown in Figure 4.21. At present, the correlator is operational and collecting commissioning and early science data.

The installation of the feed cassettes, analog receivers, and the weatherproofing of the focal line of the cylinders started in September 2015 and finished in September 2017. The last cassette was installed during the ‘first light’ ceremony held on September 7th 2017 to inaugurate the operational phase of the CHIME telescope¹⁸.

4.6 Conclusions

In this chapter we presented the ICE system, a hardware, firmware, and software framework that implements large arrays of FPGA-based data acquisition, signal processing and networking for the next generation of radio to sub-millimeter telescopes. We also described in detail the CHIME correlator which uses the ICE system to implement the data acquisition, Fourier transform channelization, and corner-turn networking. The correlator processes a massive amount of data: the F-engine raw input rate is ~ 13.1 Tbit/s, the corner-turn network has to rearrange ~ 6.6 Tbit/s of data, and the X-engine performs ~ 839 TcMAC/s. Measured in number of baselines times bandwidth, the correlation operation for CHIME is larger than for any currently existing telescope array.

Finally, we described the installation of the CHIME correlator, and demonstrated that it complies with all the requirements for CHIME by performing tests of manufacturability, uptime and reliability, data transport integrity, analog performance and data acquisition integrity, timing, and power consumption.

¹⁸To see ‘CHIME in the News’ go to <https://chime-experiment.ca/news>




Figure 4.20: Photographs of the ICE hardware withing the FPGA huts at different stages of the installation. Top left: First crate installed in the east hut (front view, September 2016). Top right: Four ICE crates mounted in two racks within the east hut (front view, February 2017). Each crate pair processes one cylinder. Bottom left: West hut after installing the ICE crates and routing the data fibers, but before the networking was ready (rear view, May 2017). In addition to the two main racks, each hut has an auxiliary (right-most) rack that contains a GPS and clock distribution unit, a network switch that aggregates the data from the four crates (not shown), a control computer (not shown), and a fully operational spare crate. Bottom right: East hut after installing all the ICE hardware (rear view, May 2017).




Figure 4.21: First CasA fringes from a single east-west CHIME baseline and a single frequency channel. Top: real component (dotted blue), imaginary component (dotted red), and magnitude of the visibility (green). Bottom: phase of the visibility.

Chapter 5

Calibration

A major challenge for CHIME comes from the calibration needed in order to detect the 21 cm signal in the presence of bright astrophysical foregrounds that are many orders of magnitude brighter. The pathfinder has been used to develop techniques that allow us to meet the stringent calibration requirements for full CHIME. Among these, we have developed a Broadband Injection Signal (BIS) that allows us to measure and correct the receiver gain fluctuations. We have also developed digital calibration techniques that minimize the negative effects of the data compression that occurs at different stages of the digital correlator. In this chapter, we discuss the calibration requirements for CHIME and present in detail the BIS technique and the digital calibration techniques including the theory, instrumentation, current performance on the pathfinder, and the development towards the implementation for full CHIME. The information on digital calibration presented in this chapter is complemented with the information in Chapter 6.

5.1 Calibration requirements

The m -mode formalism described in Section 2.4 allows a computationally efficient representation of the data from transit telescopes like CHIME that enables a highly effective method to separate the 21 cm signal from the contaminating foregrounds based on their statistics. However, the level of success of this foreground filter depends on the level of detail with which we understand our instrument, including the primary beams and the fluctuations in the receiver gains. An imperfect knowledge of the telescope leads to an imperfect foreground removal, and thus to a biasing in the estimated 21 cm power spectrum.

The effects in the recovered power spectrum of both receiver gain fluctuations and an inaccurate model for the primary beam are investigated in [58] where, based on end-to-

end simulations of a pathfinder-like instrument, it is found that all the primary beams must be known to $\sim 0.1\%$ and the complex receiver gains to $\sim 1\%$ on minute time scales in order to recover an unbiased power spectrum. More recent simulations suggest that for full CHIME we should aim for gain calibration requirements around 0.3% .

We have used the pathfinder as a test-bed to develop calibration strategies that allow us to address these stringent requirements and understand the CHIME instrument. These calibration techniques are classified into three main categories:

- Beam characterization.
- Receiver gain calibration.
- Digital calibration.

Beam characterization techniques include the observation of bright point sources, pulsars, and holographic measurements using the DRAO John A. Galt 26 m telescope. These techniques are presented in detail in [5, 59]. Receiver-gain and digital calibration techniques are presented in the following sections.

5.2 Receiver gain calibration

The stability of the receiving system determines the time scales over which the receiver gain calibration needs to be performed. This scale is typically several minutes, much faster than the time scale for point-source transits (several hours). For the analog system, phase variations are dominated by temperature effects on the coaxial cables, while amplitude variations have contributions from the amplifiers and coaxial cables. For the digital system, variations are mainly due to timing jitter and clock drift. In addition to the correction of the visibility data based on a temperature model developed for the analog receivers that we describe briefly as background information in Section 5.2.1, we have implemented a BIS technique on the pathfinder to calibrate the receiver gains and stabilize the telescope.¹

5.2.1 Thermal characterization

The thermal susceptibility of samples of the different components in the CHIME analog receiver chain (LNA, coaxial cables, FLAs) has been characterized in a temperature-

¹There are other calibration methods currently under investigation for CHIME that take advantage of the high redundancy of the CHIME array to determine both the receiver gains and the true-sky visibilities. These methods are presented in detail in [17].

controlled chamber. Currently, a subset of the receivers installed on CHIME and the pathfinder is equipped with thermal sensors that allow us to monitor temperatures on time scales faster than the gains typically vary.

On the pathfinder, the analog receivers have been found to be stable over time, with typical gain variations below $\sim 5\%$ over multiple days. Similar results have been obtained for full CHIME from recent observations. The dominant source of variance is common to all the receivers and is driven by temperature changes in the analog receiver components, specially in the LNAs and the cables which are in an outdoor environment. Of these two components, the variations of the cables dominate since the LNAs are enclosed in the feed cassettes which provide some crude weather protection. The effect of the FLAs is sub-dominant since these are located inside the temperature controlled RF-room. The thermal susceptibility of the analog receiver chain inferred from the common-mode variations is consistent with the lab measurements.

Figure 5.1 shows the amplitude fluctuations of the full CHIME receiver gains after subtracting the common-mode variation. The data for these measurements consisted of 31 consecutive days between April and May 2018. The gains are obtained by eigenvalue decomposition of the correlation matrix (see Section 5.2.2 and Appendix D) during daily CygA transits. Note that the receivers are remarkably stable. After removing the common mode, most feeds show amplitude fluctuations below the $\sim 0.3\%$ level and meet the calibration requirements. This suggests that a simple common-mode measurement from ambient temperature may be sufficient to meet our requirements. However, phase calibration additionally requires the measurement of the timing jitter using a system like the one described in Section 4.4.5.

5.2.2 Broadband Injection Signal calibration

Since the receiver gains might vary on time scales much faster than the transit of bright point sources, we are developing a BIS calibration technique that allows us to track and correct in real time for the variations in the complex receiver gains in a way that is similar to the calibration with point sources. The technique consists of injecting a broadband signal across the interferometric array. This signal is measured and correlated with all the inputs of the array. The signal is switched on and off on time scales much faster than the characteristic receiver fluctuation time scales to remove the sky contribution. Finally, the receiver gains are recovered from the sky-subtracted correlation matrix for every frequency in the CHIME band.

A diagram of the setup for the BIS calibration on the pathfinder is shown in Figure 5.2.




Figure 5.1: Median absolute deviation (MAD) of fractional amplitude fluctuations of the analog receiver gains in full CHIME from daily CygA transits after subtracting the common-mode variation. The x-axis denotes the feed number (its location in the array). The black dashed lines denote the separation between different cylinders (0-511 for west-most cylinder, 1536-2047 for east-most cylinder). The red dots denote the median value of the fractional amplitude fluctuations over all good frequencies. Shaded regions span the 68-th and 95-th percentile of good frequencies. After removing common mode, the most stable feeds show variations below the $\sim 0.3\%$ level and meet the calibration requirements. The large fluctuations for feeds in the approximate range 280-380 was due to a cassette with defective receivers on the west-most cylinder. These receivers were identified and fixed afterwards. This figure was made by Seth Siegel.

A simple way to generate a bright broadband source is to amplify the noise produced by an amplifier. The BIS source for the pathfinder consists of a terminated and band-pass filtered LNA. This BIS source is installed inside the RF room. The output of this source is connected to an RF switch that is controlled by one of the motherboards of the pathfinder ICE crate. The switched signal is then passed through a signal conditioning and distribution module that includes equalization, matching, signal level adjustment, delay adjustment and signal splitting, such that at the output of this module we have three copies of the injection signal. One copy of the signal is connected directly to one of the inputs of the correlator. This is the reference input that is used to track variations in the injected signal and to subtract the BIS signal from the visibility data after calibration such that it does not contribute significantly to the system noise. Two additional copies are sent to custom-made broadcasting helical antennas installed at the base of each cylinder.




Figure 5.2: Schematic diagram of the BIS setup for the CHIME pathfinder. Inside the RF room a switched noise source controlled by the ICE crate generates the injection signal that is passed through a signal conditioning and distribution module and split into three copies. One copy of this signal is sent directly to the correlator as reference. The other two copies are sent to broadcasting antennas installed at the base of each cylinder.

Gain estimation

From Section 2.3.3, we can write the measured visibility between feeds i and j when the BIS source is ‘off’ as

$$r_{ij}^{off} = g_i g_j^* r_{ij} + n_{ij}^{off} \quad (5.1)$$

where r_{ij} is the ‘true-sky’ visibility given by equation 2.37, g_i and g_j are the receiver gains for feeds i and j , and n_{ij}^{off} is the system noise term². When the BIS source is ‘on’, it is seen by the array as an additional radio source uncorrelated with the sky, so the measured visibility is

$$r_{ij}^{on} = \gamma_i \gamma_j^* + g_i g_j^* r_{ij} + n_{ij}^{on} \quad (5.2)$$

where the term $\gamma_i = \alpha_i g_i \sigma_{BIS}$ includes the BIS source (voltage) level σ_{BIS} , the (time-dependent) receiver gain term g_i , which is the quantity we want to measure, and a term α_i that represents the static transfer function of the injected signal and defines the amount of BIS signal ‘seen’ by each feed (it contains the effects of the distribution system and also the direction-dependent coupling between the BIS transmitter and the CHIME feeds). The α_i terms are determined later in the cosmology pipeline by performing an absolute calibration using sky point sources. The n_{ij}^{on} term is another realization of the system noise (not to be confused with the noise from the BIS source).

By switching the BIS source on and off on time scales much faster than those over which the sky changes, we can remove the effect of the sky completely by subtracting r_{ij}^{off} from r_{ij}^{on} , obtaining

$$R_{ij} = r_{ij}^{on} - r_{ij}^{off} = \gamma_i \gamma_j^* + N_{ij} \quad (5.3)$$

where $N_{ij} = n_{ij}^{on} - n_{ij}^{off}$ is the residual noise term from the subtraction. The noise term is expected to be Gaussian with

$$\langle N_{ij} \rangle = 0, \quad \langle N_{ij} N_{kl}^* \rangle = \mathcal{N}_{(ij),(kl)} \quad (5.4)$$

where \mathcal{N} is the covariance matrix of the full set of noise correlation products.

In the case of telescopes like CHIME where the sky (and BIS) signals are weak and the system temperature, T_{sys} , is dominated by the uncorrelated noise of the analog receiving system (T_{sys} is basically the measured power in the autocorrelations), then the noise covariance matrix \mathcal{N} is diagonal with [60]

$$\mathcal{N}_{(ij),(kl)} = \delta_{ik} \delta_{jl} \sigma_{ii} \sigma_{jj} \quad (5.5)$$

where $\sigma_{ii} = T_{sys,i} / \sqrt{\Delta\nu\tau}$, $\Delta\nu$ is the width of the frequency channel under analysis, and τ as the integration time. Note that this is equivalent to say that all the correlation

²Note that all these terms depend on frequency. We do not write that dependence explicitly to keep the notation simple.

products are uncorrelated and that the variance in the correlation product ij is given by the radiometer equation 2.43

$$\sigma_{ij}^2 = \frac{T_{sys,i}T_{sys,j}}{\Delta\nu\tau} = \sigma_{ii}\sigma_{jj}. \quad (5.6)$$

By making the change of variables

$$R'_{ij} = \sqrt{\frac{1}{T_{sys,i}T_{sys,j}}} R_{ij} \quad \gamma'_i = \sqrt{\frac{1}{T_{sys,i}}} \gamma_i \quad (5.7)$$

we transform the problem in equations 5.3 and 5.4 into a problem with a noise covariance that is just a scalar matrix

$$R'_{ij} = \gamma'_i \gamma'^{*}_j + N'_{ij} \quad \langle N'_{ij} N'^{*}_{kl} \rangle = \delta_{ik} \delta_{jl} \sigma^2 \quad (5.8)$$

where $\sigma = 1/\sqrt{\Delta\nu\tau}$. We call γ'_i the *weighted gain*. In practice, this change of variables is done by normalizing correlation products by the geometric mean of the autocorrelations in the ‘on’ sample³, so $T_{sys,i} = T_{sys,i}^{on}$.

See [34, 61] for a detailed introduction to maximum likelihood estimation (MLE). With this uncorrelated, uniform-variance noise weighting σ^2 , the gain vector $\tilde{\gamma}'$ that maximizes the likelihood function $\mathcal{L}(\mathbf{R}'|\gamma')$ is the vector that minimizes the chi-squared, χ^2 , given by

$$\chi^2 = \frac{1}{\sigma^2} \sum_{ij} |R'_{ij} - \gamma'_i \gamma'^{*}_j|^2. \quad (5.9)$$

As shown in Appendix D, the MLE solution is given by

$$\tilde{\gamma}' = \underset{\gamma'}{\operatorname{argmin}} \sum_{ij} |R'_{ij} - \gamma'_i \gamma'^{*}_j|^2 = \sqrt{\lambda_1} \hat{\mathbf{v}}_1 \quad (5.10)$$

where λ_1 is the largest eigenvalue of the sky-subtracted and normalized correlation matrix \mathbf{R}' and $\hat{\mathbf{v}}_1$ is the corresponding (unit-norm) eigenvector⁴. This MLE solution is unbiased.

Finally note that, although the analysis above (and the one in Appendices D and E) is done in the context of BIS calibration, the results are very general, and can also be applied for other calibration methods including bright point-source and pulsar transits. In these cases the analysis is identical but α_i becomes the phase term due to the geometric time delay and σ_{BIS}^2 becomes the temperature of the sky source. For bright point sources

³Note that we normalize by the autocorrelations before subtraction. In the regime where we are dominated by T_{sys} we can use the autocorrelations from either the ‘on’ or the ‘off’ sample.

⁴The eigendecomposition exists because \mathbf{R}' is hermitian.

the ‘on’ and ‘off’ samples correspond to measurements with the sky source in and out of the CHIME primary beam respectively. For pulsars, the ‘on’ samples are synchronized with the received pulses.

Performance requirements

A detailed analysis of the expected BIS performance from the eigenvalue decomposition described above is presented in Appendix E. We present the main results in this section.

The accuracy in the estimation of the gain solutions with BIS calibration is determined by the eigenvalue spectrum of the sky-subtracted and normalized correlation matrix \mathbf{R}' . As a metric of the BIS calibration performance we define the *BIS signal-to-noise ratio*, SNR_{BIS} , as

$$SNR_{BIS} = \sqrt{N-1} \frac{\langle \lambda_1 \rangle}{\langle \hat{\sigma}_r^2 \rangle^{1/2}} \quad (5.11)$$

where N is the number of feeds (including the BIS reference) and $\hat{\sigma}_r^2$ is the sample variance of the remaining eigenvalues (which in the absence of noise would all be zero). The SNR_{BIS} can be written in terms of known or measurable telescope and BIS parameters as

$$SNR_{BIS} = \sqrt{\Delta\nu\tau} \sum_{i=1}^N f_i \quad (5.12)$$

where

$$f_i = \frac{T_{sys}^{on} - T_{sys}^{off}}{T_{sys}^{on}} = \frac{r_{ii}^{on} - r_{ii}^{off}}{r_{ii}^{on}} \quad (5.13)$$

is the fraction of the total signal ‘seen’ by feed i that comes from the BIS signal, also known as the *illumination* of feed i . Note that the illumination is a measurable quantity that can be determined from the autocorrelations in the ‘on’ and ‘off’ samples.

In general, an accurate determination of the receiver gains requires achieving a high SNR_{BIS} . As shown in Appendix E, the errors on the estimated gains are approximately uncorrelated and the expected fractional errors for the receiver gain solutions are given by

$$\delta g_i = \sqrt{\frac{\sigma}{SNR_{BIS} \cdot f_i}} = \left(\Delta\nu\tau \cdot f_i \sum_{n=1}^N f_n \right)^{-1/2}. \quad (5.14)$$

Equation 5.14 gives the expected fractional error in the receiver gain solutions in terms of telescope parameters like bandwidth and integration time, and BIS calibration parameters like the illumination of the array. These parameters need to be adjusted to obtain $\delta g_i \lesssim 0.3\%$ which is the requirement to achieve the cosmology science goals (see Section 5.1). Note that the gain error of feed i depends not only on how well the feed ‘sees’ the injected signal, but also on the illumination of the full array. Also that $0 \leq f_i \leq 1$ and that including a reference input that measures the BIS source directly means that this reference input has $f = 1$ so we can write

$$\delta g_i = \left[\Delta\nu\tau \cdot f_i \left(1 + \sum_{n=2}^N f_n \right) \right]^{-1/2} \quad (5.15)$$

where we chose input 1 as the BIS source reference.

Current pathfinder parameters have $\tau \sim 20$ s, $\Delta\nu \sim 0.4$ MHz, and $N \sim 200$. If all the CHIME feeds could be illuminated such that the injected signal was $\sim 1\%$ of the total power absorbed by each feed (i.e. $f_i \sim 0.01$ for $i > 1$), then the SNR_{BIS} would be about 39 dB and the expected gain errors would be $\delta g_i \sim 0.2\%$ which is about the minimum level of accuracy we aim to achieve. If the injected signal was $\sim 10\%$ (i.e. $f_i \sim 0.1$ for $i > 1$), then $SNR_{BIS} \approx 48$ dB and the expected gain errors would be $\delta g_i \sim 0.02\%$, well below the minimum calibration requirements. As for full CHIME, typical parameters are $\tau \sim 10$ s and $N \sim 2000$, so a $\sim 5\%$ illumination ($f_i = 0.05$ for $i > 1$) would be enough to achieve $\delta g_i \sim 0.02\%$.

Note that although equation 5.15 suggests that the gain errors can be reduced by increasing the injected signal level, this situation is not desirable since: 1) it would mean that, for each feed, T_{sys} is dominated by the noise source (so equation 5.7 would not hold), 2) there is risk of saturating the analog receivers, and 3) there is risk of introducing a significant bias in the magnitude and phase of the measured visibilities due to large deviations of the correlator input signal levels from the nominal values (see Chapter 6).

Status and current development

Since the first tests in 2014 with 16 feeds on the pathfinder that was presented in [5] the BIS system has been under constant improvement, progressively evolving towards a system that is more robust and stable, synchronized with the data acquisition, and that can be easily extended for full CHIME.

The main improvements in the BIS system with respect to the system presented in

[5] are:

- In the early stages of the BIS system the RF switch was controlled by a separate signal generator and the switching was independent of the data acquisition and integration, which resulted in a considerable fraction of the correlated data needing to be discarded due to on-off transitions in the middle of the integrations. For this same reason we could only switch on time scales larger than the nominal visibility integration time (~ 20 s). For the current system the RF switch is controlled by one of the motherboards in the pathfinder ICE crate. The ICE firmware and software now allow the configuration of one of the board's general purpose outputs as a Pulse Width Modulation (PWM) control signal that can be connected directly to the switch control driver. The PWM signal is synchronized with the data acquisition, and both its period and duty cycle (in number of frames) can be configured through software. Using this capability it is possible to switch at cadences faster than the integration time, improving both the sky signal subtraction and the immunity to RFI. The nominal switching time is ~ 80 ms, and the X-engine accumulates and saves separately the visibility data during the on- and off-periods of the BIS source.
- Each motherboard in the ICE system can be used to control different BIS sources independently. On the pathfinder, we currently use the ICE crate to independently operate two additional and much simpler BIS calibration systems: one that is used to perform the dual-receiver and polarization calibration of the 26 m telescope, and another that is used to measure and correct for the intra-crate timing jitter of the ICE system. The pathfinder intra-crate jitter measurement system was described in Section 4.4.5. An improved and more robust system based on the pathfinder design is currently installed on full CHIME to track and correct the intra and inter-crate jitter of the ICE system, including the hut-to-hut jitter.
- We have developed custom software along with the hardware for BIS configuration and operation, as well as for data analysis and gain estimation. This software has been integrated to the CHIME cosmology pipeline as part of the calibration.

We are currently using BIS calibration successfully to measure and correct the intra and inter-crate timing jitter (see Section 4.4.5), which is the dominant source of phase noise on full CHIME. This is a simplified version of the BIS system presented in this section that does not require broadcasting the calibration signal accross the array. Since the CHIME analog receivers have been found to be much more stable than anticipated by the initial design (Section 5.2.1), with gain fluctuations that can be adequately corrected

by tracking temperature variations, the implementation of the complete version of the BIS system on full CHIME may not be necessary. Given this, the present strategy for gain calibration on full CHIME is that the thermal model will be used to track the common-mode variations of the analog receiving system and the BIS calibration will be used for tracking the inter and intra-crate timing jitter of the ICE system, including the determination of the relative phases between the inputs after each restart of the correlator.

The full BIS system that includes the broadcasting antennas is currently implemented in the pathfinder. If we were to deploy this system for full CHIME, we would first need to handle two known issues before the system can efficiently calibrate the telescope. The first one is related to the broadcasting antennas: we have found that our hand-built helix antennas are susceptible to weather-induced variations (e.g. vibrations due to wind) that cause fluctuations in the beam coupling between the helical antenna and the CHIME feeds at a level that is comparable to the fluctuations that we are trying to measure (which are sub-percent level fluctuations in the receiving system). The broadcasting antennas for CHIME would require a compact and sturdy design that makes them more stable and less susceptible to weather-induced variations⁵.

The second challenge of the full BIS system for full CHIME is related to the uniform illumination of the CHIME feeds since we must ensure that the signal level seen by each feed is enough to meet the calibration requirements. Figure 5.3 shows the typical illumination profile (in dB) of the pathfinder feeds using a single wide-beam helical antenna per cylinder. The illumination as a function of feed position along the cylinder and frequency is plotted for each cylinder and polarization separately. Note that achieving the required illumination for the furthest feeds comes at the expense of a relatively high signal level for the feeds closest to the broadcasting antenna. This high signal level regime for the closest feeds is undesirable for the reasons explained in the performance requirements sub-section above. If we were to deploy a BIS system for full CHIME, a more uniform illumination of the CHIME feeds could be obtained by using several broadcasting antennas along the cylinder reflectors and switching them on and off sequentially, effectively time-multiplexing the calibration along the feed line. The switch control capability that we have added to the ICE firmware and software is designed to allow for this sort of time-multiplexing of the calibration signals.

⁵Note that solving the problem of stabilizing a single BIS broadcasting antenna is much simpler than that of stabilizing hundreds of analog receivers.




Figure 5.3: Illumination profile (in dB) of the pathfinder feeds using a single wide-beam helical antenna per cylinder. The illumination as function of feed position (x-axis) along the cylinder and frequency (y-axis) is plotted for each cylinder and polarization separately. The horizontal stripes correspond to frequencies flagged as bad (either RFI or bad GPU nodes) while the vertical stripes correspond to feeds flagged as bad (e.g. faulty receiver).

5.3 Digital calibration

In CHIME, the analog signals from the dual-polarization feeds are digitized (sampled and quantized) before they are fed into the digital correlator to compute the complex visibilities. The CHIME correlator contains several quantization stages where the signal amplitude is encoded with a finite set of discrete values. This quantization process introduces errors in the measurement of brightness and position on the sky and must be addressed before the calibration of the beams and receiver gains. The three correlator stages where it is critical to address the effects of quantization are:

1. 8-bit digitization of the analog signals.
2. Computation of the FFT.
3. Data reduction to 4-bit real + 4-bit imaginary after the channelization stage.

5.3.1 8-bit ADC calibration

The 8-bit ADC calibration refers to the adjustment of the analog input signal level to minimize the effect of the quantization errors⁶ introduced by the 8-bit quantizer. If the input signal level is too low, it does not use all the available dynamic range of the quantizer and the output is dominated by the granular noise introduced by the spacing between the possible quantizer outputs. On the other hand, if the signal level is so high that it exceeds the maximum range of the quantizer, then the signal is clipped and the output is dominated by this overload distortion. In these two regimes the quantization error is also correlated with the input.

Figure 5.4 shows the behavior of an 8-bit quantizer for a real-valued Gaussian input with zero-mean and variance σ^2 . The top plot shows the fractional increase in the variance of the signal, $\hat{\sigma}^2/\sigma^2$, that results from quantization, as function of the unquantized standard deviation σ in units of the quantization step Δ . The middle plot shows the variance of the quantization error, σ_e^2 , relative to σ^2 , and the bottom plot is the magnitude of the correlation coefficient between the input and quantization error, $|\rho_{ve}|$. For each plot, the blue line corresponds to the quantization model developed in Chapter 6 while the red dashed line corresponds to the commonly used approximation where the quantization noise is assumed to be uncorrelated with the input. The x-axis of the plots is in base 2 logarithmic scale, so the exponents can be interpreted as the input level in bits RMS.

Note that the $\rho_{ve} = 0$ model that is traditionally used in radio astronomy (red dashed line) deviates significantly from the model developed in Chapter 6 (blue line), where it is also shown that this model provides a more accurate representation of the behavior of a real-valued 8-bit quantizer in the very low and very high signal level regimes. Under the $\rho_{ve} = 0$ model, the output variance is $\hat{\sigma}^2 = \sigma^2 + \sigma_e^2$, and maximizing the quantizer's efficiency is equivalent to minimizing σ_e^2/σ^2 . As the middle plot shows, this occurs when $\sigma \approx 2^5\Delta$, or 5 bits RMS. Under the more realistic model that takes into account the input-error correlation, maximizing the quantizer's efficiency requires keeping both σ_e^2/σ^2 and ρ_{ve} at negligible levels. In CHIME, the gain of the analog receiving system has been optimized so the nominal signal level at the input of the ADC is approximately 3.5 bits RMS⁷. As Figure 5.4 shows, at this level the quantization noise is negligible compared to the input (the increase in variance is well below 0.1%) and the quantizer is within the optimal region of operation where the input-error correlation is very weak

⁶The quantization error is the difference between the input value and its quantized value.

⁷The overall gain of the analog receiving system is fixed to this optimal value (~ 70 dB, see Figure 3.6) and, except for external attenuators that can be added to the analog chain, cannot be modified.

($|\rho_{ve}| \ll 10^{-14}$) so the model of additive uncorrelated quantization noise provides a very accurate approximation. This signal level also allows headroom for changes in the signal level due to bright point sources, receiver gain fluctuations, and RFI, while still keeping the quantizer within the optimal region of operation.

5.3.2 Minimization of quantization errors in the computation of the FFT

The FFT is an efficient algorithm for the evaluation of Discrete Fourier Transform (DFT) of a signal that reduces the number of operations in the computation of DFT of length L from $O(L^2)$ to $O(L \cdot \log_2 L)$. This is effectively achieved by decomposing the L -point DFT into successively smaller DFTs (see [32] for a detailed review of the DFT and its efficient computation using FFT algorithms).

When L is a power of two ($L = 2^\gamma$ where γ is a natural number), as is the case for CHIME, then the FFT is computed in $\gamma = \log_2 L$ stages. The use of registers with finite length in the computation of the FFT introduces round-off errors at each stage that result in an overall reduction of the signal-to-quantization-error ratio (SQR) at the FFT output. As shown in [32, 62, 63] this effect can be minimized with an appropriate scaling of the input sequence. Briefly: for an FFT computed with b -bit precision (18+18 bits for CHIME), and in order to guarantee that there are no overflows at any stage of an L -point FFT, it is necessary and sufficient to divide the input sequence by L . In this case the SQR in any value the FFT is

$$SQR \propto \frac{2^{2b}}{L^2}. \quad (5.16)$$

The scaling is responsible for the $1/L^2$ decrease in the SQR , which is a rather severe penalty. However, the scaling can be done in a more clever way to prevent overflow and obtain an SQR that decreases as $1/L$, instead of $1/L^2$. The CHIME FPGA firmware and software allow the configuration of an optional division by 2 (right shift) at each stage of the FFT. Thus, the same $1/L$ scaling can be introduced by dividing by 2 at each of the γ stages. This scaling scheme not only prevents overflow at all the stages of the FFT, but also improves its quantization performance since the SQR becomes

$$SQR \propto \frac{2^{2b}}{L}. \quad (5.17)$$

Intuitively, this stage-by-stage scaling scheme has a better SQR performance because it keeps the signal level as high as possible at each stage compared to the quantization




Figure 5.4: Behavior of an 8-bit quantizer for a real-valued Gaussian input. Top: variance of the quantized output, $\hat{\sigma}^2$, normalized by the unquantized variance σ^2 , as function of the unquantized standard deviation σ . Middle: variance of the quantization error, σ_e^2 , normalized by the unquantized variance. Bottom: magnitude of the correlation coefficient between the input and quantization error, $|\rho_{ve}|$. The lower bound of $|\rho_{ve}|$ is limited by numerical precision. For each plot, the blue line corresponds to the quantization model developed in Chapter 6 while the red dashed line corresponds to the uncorrelated quantization noise model that is traditionally used in radio astronomy and assumes $\rho_{ve} = 0$.

errors.

Although this is a great improvement, it is still a worst case scenario, since division at each stage is only necessary to prevent overflow for inputs with maximally peaked spectra, i.e., for inputs where the power is contained within a few frequency channels (e.g. a constant input, or a sinusoid). For inputs with smoother spectra, as is the nominal case for CHIME, it is possible to reduce the number of divisions by 2 and thus improve the *SQR*. The determination of the optimal scaling scheme for signals with smooth spectra is investigated in [63], where the probability of occurrence of different scaling patterns and their associated *SQR* is estimated by transforming many white sequences, and keeping track of the stages at which divisions by 2 are required to avoid overflow. Based on this result we find that, for an FFT with $\gamma = 11$ stages like the one for CHIME, the optimal scaling pattern is SSSNSNSNSNS, meaning that divisions by 2 take place at stages 1, 2, 3, 5, 7, 9, 11, with no scalings at the remaining stages. This FFT scaling pattern is set during the initialization of the correlator before each data acquisition.

5.3.3 Post-channelization digital calibration

As explained in Chapter 4, the complex-valued output stream from the FFT at each frequency channel is truncated to $4 + 4$ bit complex values before the data is re-arranged in the corner-turn network. This truncation inevitably results in a loss of information which, in the case of typical radio astronomy signals, is small and justified by other advantages like the possibility of using a larger bandwidth with fewer bits. This complex-valued quantization step also introduces a bias in the magnitude and phase of the measured correlations which translates into errors in the measurement of source brightness and position in the sky, affecting both the system calibration and image reconstruction. In Chapter 6 we investigate in detail the biasing effect of the post-channelization quantization and determine the optimal signal level that reduces this effect to negligible levels. By setting the signal level to two quantization steps ($\sigma = 2^1 \Delta$, or 1 bit RMS) before the complex quantization takes place, the correlation between the input and the quantization error is very weak ($|\rho_{ve}| \ll 10^{-3}$), the quantization noise is kept relatively low (the increase in output variance is about 4%), and the resulting bias in the measured correlations due to quantization is negligible (below $\sim 10^{-6}$ for the magnitude and $\sim 10^{-11}$ degrees phase).

The post-channelization calibration module applies a digital gain to each frequency channel in order to adjust the signal level before this 4-bit complex-valued quantization takes place in the scaler module (see Section 4.2.2). An schematic diagram of this cal-

ibration process is shown in Figure 5.5. Since the input to this module are $18 + 18$ bit samples from the PFB/FFT, the calibration not only has to set the signal level of each frequency channel to the optimal value, but it also has to select the set of $4 + 4$ bits that contain the sky information. As an example, for a typical frequency channel the most significant bits (MSBs) of the $18 + 18$ FFT output are zeros, while for a channel corrupted by strong RFI the MSBs may not be zero.




Figure 5.5: Schematic diagram of the post-channelization digital calibration. The digitized from the ADC is passed to the PFB/FFT sub-module as frames of 2048 8-bit samples. The FFT outputs a frame of 1024 complex values, one per frequency channel, in $18+18$ bit format. For each frequency channel, the FFT output is multiplied by a $16 + 16$ bit complex gain. The resulting $35 + 35$ bit complex number is then shifted to the left a configurable amount of times. The amount of shifts to the left is the same for all the frequency channels. Finally, the $35 + 35$ bit complex number is scaled to a $4 + 4$ bit complex value by selecting the 4 MSBs in both the real and imaginary component and applying convergent rounding.

The calibration for each input and frequency channel is performed in two stages. First, the $18 + 18$ bit output of the PFB/FFT sub-module of the channelizer is multiplied in the scaler sub-module by a $16 + 16$ bit complex gain for each frequency channel. Then, the resulting $35 + 35$ bit complex number⁸ can be shifted to the left (multiplied by two) a configurable amount of times before it is finally scaled to a $4 + 4$ bit complex value by selecting the 4 MSBs in both the real and imaginary component and applying convergent rounding. The amount of shifts to the left is the same for all the frequency channels.

The pair formed by the $16 + 16$ bit complex gain and the number of shifts to the left is the digital gain that is applied in the post-channelization calibration. This gains can

⁸Note that the complex multiplication requires 35 bits in the real and imaginary components of the result since the operation involves both multiplication and addition.

be configured through software, and are adjusted so the signal level of the PFB/FFT output is about 1 bit RMS before the 4-bit quantization takes place. This is done in two stages: first, for each input the digital gain is calculated independently for each frequency channel by measuring the RMS of the channelized data. After that, the gains are smoothed in frequency using an iterative algorithm that combines polynomial fitting and Fourier smoothing. RFI frequency channels are treated separately. The details of this algorithm are presented in [64]. The final digital gains are applied to the channelized data. They are also stored and passed to the cosmology pipeline so they can be accounted for in the visibility data.

In principle the digital gains only need to be calculated once for each input and frequency channel. In practice for CHIME, the digital gains are updated when there are changes in the analog receiving system (e.g. an analog input that was down and has been fixed) or in the RFI environment (e.g. a ‘clean’ frequency channel that is suddenly contaminated by RFI or vice versa).

5.4 Conclusions

In this chapter we presented the calibration requirements for CHIME and the different techniques that we have developed to calibrate the telescope. Although highly effective at separating the 21 cm signal from foregrounds, the m -mode formalism described in Chapter 2 requires a detailed knowledge of the instrument, including the complex receiver gains and the primary beams. An imperfect calibration of the telescope leads to an imperfect foreground removal, and thus to a biased estimate of the 21 cm power spectrum. In order to avoid significant power spectrum biases, the complex receiver gains must be known to $\sim 0.3\%$ within each minute.

We have used the pathfinder as a test-bed to develop a number of receiver-gain and digital calibration techniques in order to address these stringent requirements and understand the CHIME instrument.

For receiver gain calibration, the thermal susceptibility of the different components of the analog receiver chain has been characterized. The dominant source of gain variation is common to all the receivers and is driven by temperature changes in the analog receiver components. After subtracting the common-mode drift most feeds show fluctuations below the 0.3% level and meet the calibration requirements.

We developed and implemented a BIS system that injects broadband calibration signals across the array to measure and correct for the receiver gain fluctuations. We presented the theory of BIS calibration, which is very general and can also be applied to

other calibration methods including bright point-source and pulsar transits. We also gave details of the BIS instrumentation and its improvements with respect to the early system presented in [5]. The new BIS system is synchronized with the data acquisition and allows switching at cadences much faster than the integration time, improving the sky background subtraction and the immunity to RFI. Each motherboard in the ICE system can be used to control different BIS sources independently and we have developed a much simpler BIS system that is currently used to track and correct the intra and inter-crate jitter of the ICE system on the pathfinder and full CHIME.

Since the CHIME receivers have been found to be more stable than anticipated, the implementation of the full version of the BIS system on full CHIME may not be necessary, although the BIS system for timing jitter is still required. The present strategy for gain calibration on full CHIME is that a thermal model will be used to track the common mode and calibrate the analog receiving system and the BIS calibration will be used to track the timing jitter of the ICE system.

For digital calibration, we have developed efficient calibration techniques that minimize the effects of quantization errors at the critical correlator stages. For the 8-bit digitization of the analog signals, the gain of the analog receiving system has been optimized so the nominal signal level at the input of the ADC is approximately 3.5 bits RMS, reducing the effects of the quantization noise to negligible levels. For the computation of the FFT, we have implemented a scaling scheme for the FFT stages that maximizes the signal-to-quantization-error ratio at the FFT output. For the post-channelization 4-bit complex quantization we have implemented a calibration algorithm that adjusts the signal level for each frequency channel before this 4-bit complex-valued quantization takes place, effectively minimizing the quantization errors. This calibration is based on the quantization model developed in Chapter 6.

Chapter 6

Quantization bias for digital correlators

In this chapter we investigate the biasing effect of the post-channelization complex quantization in the magnitude and phase of the measured visibilities for circularly symmetric Gaussian inputs, which is the typical case for CHIME and other radio astronomy applications. Using the new knowledge derived in this chapter and submitted for publication in [6], We determine the optimal interval for the input signal level for which the correlation between the input and the quantization error is very weak and the model of additive uncorrelated quantization noise provides a very accurate approximation. We also determine the quantization bias in the visibilities measured by a digital correlator and derive the conditions under which the bias in the magnitude and phase of the measured correlation is negligible with respect to the unquantized values. Finally, we show how the post-channelization quantization reduces the point-source sensitivity of a CHIME-like array and how this effect can be reduced to negligible levels with a proper calibration of the system noise level at the input of the quantizer.

6.1 Post-channelization quantization in the CHIME correlator

The CHIME correlator is an FX design, where the F-engine divides the 400 MHz input bandwidth into thousands of frequency channels. The data from each frequency channel is then quantized to 4 bits (4 bits real + 4 bits imaginary) before being reorganized by the corner-turn network and fed into the X-engine that computes the full N^2 correlation matrix (see Chapter 4 for details on the CHIME correlator).

This post-channelization quantization process introduces a bias in the magnitude and phase of the measured correlations which translates into errors in the measurement of source brightness and position on the sky, affecting both the system calibration and image reconstruction. This bias in the correlations will show up as an amplitude dependent (and direction independent) gain term that must be addressed before beam and receiver gain calibration. As we will show, this effect can be significant for large deviations from optimal signal levels or large correlation coefficients, which means that it is critical to understand the bias. This implies understanding the statistics of the quantization error and its correlation with the quantizer input.

In most cases the quantization error is modeled as additive stationary white noise that has a uniform bounded distribution and is uncorrelated with the input. In general, this model provides a very good approximation when the quantization step size is small, the input signal traverses many quantization steps between successive samples and the effect of clipping (for input values outside the quantizer's dynamic range) is small or negligible. In this case [65] derives formulas for the fractional increase in the variance of a white Gaussian real input signal that results from quantization with many levels (eight or more) and provides tables with the optimal input signal levels that minimize this effect. Although the uncorrelated quantization error model is still very accurate even for significant deviations from the optimal signal level (the range depending on the number of levels), the model breaks when the input signal level is too small or when it is too high and the effect of clipping is important (i.e. when the fraction of samples that lie outside the quantizer limits is significant). More important, it leads to the incorrect conclusion that the magnitude and phase of the quantized correlation remain unbiased. We will show that this is not the case in the signal regimes described above and when the correlation coefficient is large.

The effect of quantization on correlators has been studied in the past for quantization with few levels (e.g. [66] for two levels, [67] for three levels, [68] for four levels). For many levels, [65] studies the loss in efficiency in a correlator resulting from quantization with eight or more levels for real Gaussian inputs assuming that the quantization error is uncorrelated with the unquantized input, while [69] finds the component of the quantization noise that is uncorrelated with the input and calculates the loss in efficiency due to this component. [33] presents a detailed discussion on these methods. Recent work from [70] generalizes the Van Vleck quantization correction for two-level correlators to correlators with multilevel quantization and Gaussian inputs. Since it is not always computationally feasible to implement this correction, in this chapter we investigate in detail the biasing effect of quantization on the magnitude and phase of the measured correlations

and determine the conditions under which this effect is negligible so the correction is not necessary. In order to do that, we calculate the contribution of each quantization level to the correlation between the input and the quantization error in the case of a single quantizer, and in the case of two quantizers with different inputs we calculate the correlation between quantization errors for every pair of quantization levels¹. We then use these results to calculate the effect of the quantization errors on the measured correlation of a real and complex-valued correlator.

The results of this chapter are general and apply to any digital correlator. However, we will focus our analysis and simulations on the 4-bit real + 4-bit imaginary complex-valued quantization at the channelization stage of the CHIME correlator. We are particularly interested in the effects of quantization in the high signal level and high correlation regimes which are relevant for CHIME when the antenna temperature and thus the correlator input signal can increase significantly relative to the optimal level (typically determined at night hours or when observing a relatively quiet region of the sky), for example during bright point source transits (e.g. the sun) and point-source calibration, and during complex receiver gain calibration using the BIS technique (see Chapter 5).

6.2 Real-valued quantizer

We will assume a quantizer with uniformly spaced levels and an odd symmetric transfer function (same number of levels above and below zero). This means that if the number of levels N is odd then the quantizer has a level at zero (mid-tread) and if N is even it has a threshold at zero (mid-riser). We do not consider non-uniform quantization steps for optimization. The CHIME case, which we assume as an example, corresponds to $N = 15$ (levels at -7, -6, ..., 6, 7) for the complex channelization stage. In general, the quantizer levels are (in units of the quantization step Δ)

$$k_i = -\frac{N+1}{2} + i, \text{ for } i = 1, \dots, N \quad (6.1)$$

and the decision thresholds are

$$y_0 = -\infty, \quad y_N = \infty, \quad y_i = k_i + \frac{1}{2} = -\frac{N}{2} + i, \text{ for } i = 1, \dots, N-1. \quad (6.2)$$

¹This method differs from [33] and [70] since it does not use Price's theorem [71], a very useful tool for estimating the expectation of nonlinear functions of jointly Gaussian random variables. The approach used in this chapter applies to generic probability density functions, and can be used for example, to investigate the effect of quantization in the presence of RFI, although that analysis is left for future work.

Let v be the (real-valued) input of the quantizer. For the i -th quantization level, the correlation between the input and the quantization error e is [72]

$$\langle ve \rangle_i = \int_{y_{i-1}}^{y_i} (k_i - v) v f(v) dv \quad (6.3)$$

where v is in quantization step units and has probability density function (PDF) $f(v)$. Each input sample can fall in only one quantization slot so events that take place in the various slots are mutually exclusive. This means that we can write the correlation between the input v and the quantization error $e = \hat{v} - v$ as (\hat{v} is the quantizer output)

$$\langle ve \rangle = \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v) v f(v) dv. \quad (6.4)$$

Similarly, the quantization error variance $\sigma_e^2 = \langle e^2 \rangle$ can be written as

$$\sigma_e^2 = \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v)^2 f(v) dv. \quad (6.5)$$

As equation 6.4 shows, the calculation of $\langle ve \rangle$ depends on the input PDF. If v is an independent and identically distributed (IID) Gaussian process with zero mean, then equation 6.4 can be written in a more concrete form

$$\begin{aligned} \langle ve \rangle &= \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v) v \mathcal{N}(v|\sigma^2) dv = \sigma^2 \left[-1 + \sum_{i=1}^{N-1} \mathcal{N}\left(-\frac{N}{2} + i \middle| \sigma^2\right) \right] \\ &= \begin{cases} \sigma^2 \left[-1 + 2 \sum_{i=0}^{\frac{N-3}{2}} \mathcal{N}\left(\frac{1}{2} + i \middle| \sigma^2\right) \right] & \text{if } N \text{ odd} \\ \sigma^2 \left[-1 + \frac{1}{\sqrt{2\pi\sigma^2}} + 2 \sum_{i=0}^{\frac{N-4}{2}} \mathcal{N}(1 + i \middle| \sigma^2) \right] & \text{if } N \text{ even} \end{cases} \end{aligned} \quad (6.6)$$

where $\mathcal{N}(v|\sigma^2) = (2\pi\sigma^2)^{-1/2} e^{-v^2/(2\sigma^2)}$ is the Gaussian PDF, σ is in units of the quantization step Δ , and it is clear that the summation term is zero for $N = 2$. Appendix A provides a derivation for equation 6.6.

It is also clear from the symmetry of the quantizer and the input PDF that both e and \hat{v} have zero mean. Using the same procedure we find the variance of the quantization

error $\sigma_e^2 = \langle e^2 \rangle$ as (see Appendix A for details)

$$\begin{aligned} \sigma_e^2 &= \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v)^2 \mathcal{N}(v|\sigma^2) dv \\ &= -2\langle ve \rangle - \sigma^2 + \left(\frac{N-1}{2}\right)^2 - \sum_{i=1}^{N-1} \left(-\frac{N}{2} + i\right) \operatorname{erf}\left(\frac{-N/2 + i}{\sqrt{2\sigma^2}}\right). \end{aligned} \quad (6.7)$$

where $\operatorname{erf}(v)$ is the error function. The quantized output variance $\hat{\sigma}^2$ follows from equations 6.6 and 6.7

$$\begin{aligned} \hat{\sigma}^2 &= \langle (v + e)^2 \rangle = \sigma_e^2 + \sigma^2 + 2\langle ve \rangle = \left(\frac{N-1}{2}\right)^2 - \sum_{i=1}^{N-1} \left(-\frac{N}{2} + i\right) \operatorname{erf}\left(\frac{-N/2 + i}{\sqrt{2\sigma^2}}\right) \\ &= \begin{cases} \left(\frac{N-1}{2}\right)^2 - 2 \sum_{i=0}^{\frac{N-3}{2}} \left(\frac{1}{2} + i\right) \operatorname{erf}\left(\frac{1/2 + i}{\sqrt{2\sigma^2}}\right) & \text{if } N \text{ odd} \\ \left(\frac{N-1}{2}\right)^2 - 2 \sum_{i=0}^{\frac{N-4}{2}} (1 + i) \operatorname{erf}\left(\frac{1 + i}{\sqrt{2\sigma^2}}\right) & \text{if } N \text{ even.} \end{cases} \end{aligned} \quad (6.8)$$

Results from simulations where we verify equations 6.6 - 6.8 for the case of a real quantizer with $N = 15$ levels (left column) and $N = 16$ levels (right column) and a real Gaussian input are shown in Figure 6.1. From top to bottom row, the plots show the variance of the quantized output, $\hat{\sigma}^2$, the quantization error, σ_e^2 , and the correlation between the input and quantization error $\langle ve \rangle$ as function of the unquantized standard deviation σ . All the values are normalized with respect to σ^2 . For easier visualization of the results, especially in the low and high signal level regimes, the x-axis is in logarithmic scale (base 2, so the exponents can be interpreted as bits RMS). For each plot, the red line corresponds to equations 6.6 - 6.8 and the blue line (made thicker so it can be distinguished from the red line) corresponds to the results from simulations where, for each value of σ , 10^6 samples of a Gaussian input are quantized with N levels and then the statistics of the input, output and quantization error are calculated. As reference, we also include the green dashed line that shows the expected behavior from the uncorrelated quantization noise model that assumes $\langle ve \rangle = 0$ (see [33] for a detailed discussion). The black solid vertical line corresponds to the highest level of the quantizer (7 for $N = 15$ and 7.5 for $N = 16$) above which clipping occurs.

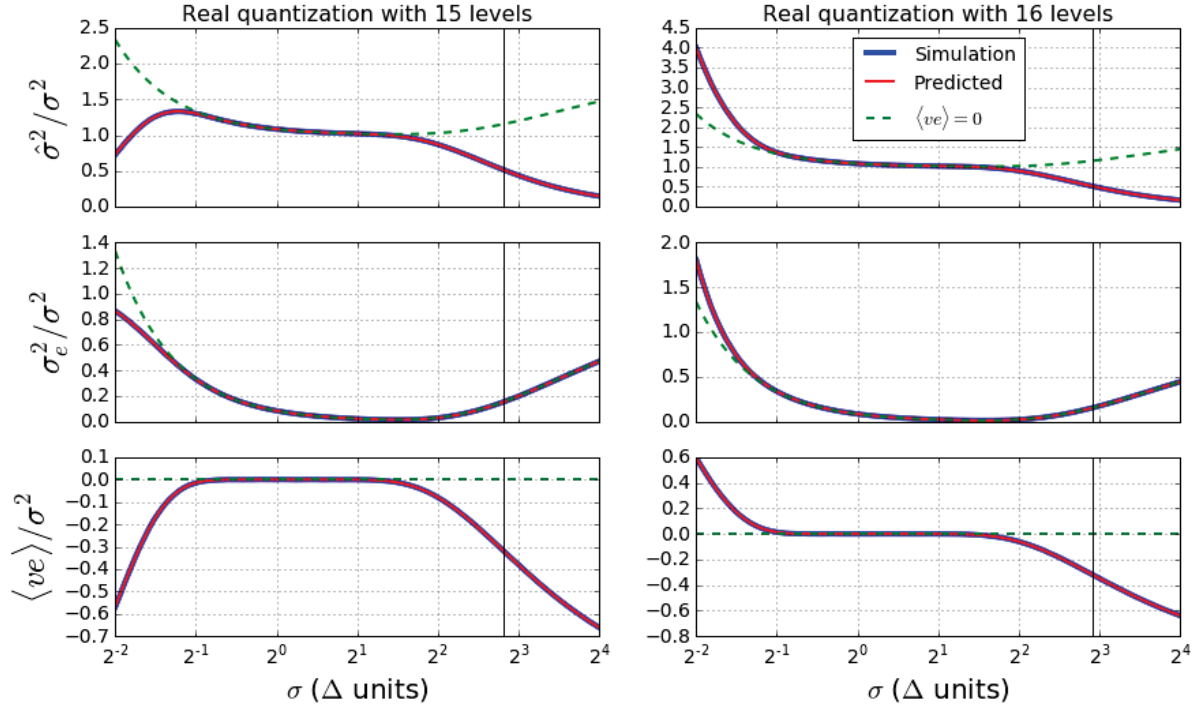


Figure 6.1: Behavior of a quantizer with $N = 15$ levels (left column) and $N = 16$ levels (right column) and a real-valued Gaussian input. From top to bottom row, the plots show the variance of the quantized output, $\hat{\sigma}^2$, the quantization error, σ_e^2 , and the correlation between the input and quantization error, $\langle ve \rangle$, as function of the unquantized standard deviation σ . All the values are normalized with respect to σ^2 . For each plot, the red line corresponds to equations 6.6 - 6.8, the thick blue line shows the results from simulations, and the green dashed line corresponds to the uncorrelated quantization noise model that assumes $\langle ve \rangle = 0$. Note that equations 6.6 - 6.8 predict accurately the results from simulations. When the input σ uses optimally the quantizer's dynamic range the quantization error is very weakly correlated with the input. In this case the uncorrelated quantization noise model provides a very good approximation, introducing only a small bias error.

The first thing to note from Figure 6.1 is that equations 6.6 - 6.8 predict accurately the results from simulations regarding $\hat{\sigma}^2$, σ_e^2 , and $\langle ve \rangle$. Also that the uncorrelated quantization noise model provides an excellent approximation in the interval where $\langle ve \rangle \rightarrow 0$. For $N = 15$, the value of σ that minimizes the magnitude of the input-error correlation coefficient, $\rho_{ve} = \langle ve \rangle / (\sigma \sigma_e)$, is $\sigma \approx 2^{0.14} \Delta$. At this point $|\rho_{ve}| \approx 5.5 \times 10^{-10}$. For $N = 16$ we have $\rho_{ve} = 0$ for $\sigma \approx 2^{0.2} \Delta$. In both cases the minimum of $|\rho_{ve}|$ is broad so there is effectively a σ -interval, which we denote the interval of optimal quantization, for which the correlation between the input and quantization error is very weak and the uncorrelated quantization error model provides a very accurate approximation (the error in the calculated quantization parameters is negligible). The length of this interval depends on N and on the tolerance required by each specific application. For example, if we require that $|\rho_{ve}| \lesssim 10^{-3}$ for $N = 15$, then the interval of optimal quantization is, approximately, $[2^{-0.6}, 2^{0.9}]$. Within this interval the values of $\hat{\sigma}^2$ and σ_e^2 from the uncorrelated quantization error model agree with the values from equations 6.6 - 6.8 at the $\sim 0.07\%$ level. The performance of the $N = 16$ quantizer within this interval is similar².

Also note that, even in the high- σ regime, where the quantization error resulting from clipping dominates and is correlated with the input, the uncorrelated quantization noise model also predicts with high accuracy the contribution of this overload error to σ_e^2 as the middle plot shows. However, it cannot track the quantized standard deviation (top plot) since in this regime $\langle ve \rangle < 0$ which eventually makes $\hat{\sigma}^2 / \sigma^2 < 1$ for large inputs. In the low- σ regime, when $\sigma \lesssim 1/2$, the uncorrelated quantization noise model deviates from equations 6.6 - 6.8 for two reasons: first, it is no longer true that the quantization error is uniformly distributed in the interval $[-\Delta/2, \Delta/2]$, and second, the behavior is now closer to that of a 3-bit (N odd) or 2-bit (N even) quantizer, so the quantization error is again correlated with the input. As N increases, both the interval of optimal quantization and the accuracy of the uncorrelated quantization error model increase.

Finally, note that $\langle ve \rangle$ is negative (it approaches zero asymptotically) for $N = 15$ while it becomes positive in the low signal level regime for $N = 16$. Since the sum $S_o = \sum_{i=0}^{(N-3)/2} \mathcal{N}(1/2 + i|\sigma^2|)$ in equation 6.6 is positive and bounded above by $1/2$ ($S_o < 1/2$, a proof is provided in Appendix G) then $\langle ve \rangle$ is always negative for N odd. Furthermore, $\langle ve \rangle \in (-\sigma^2, 0)$ in this case. On the other hand, for N even, the sum $S_e = \sum_{i=0}^{(N-4)/2} \mathcal{N}(1 + i|\sigma^2|)$ is also positive and bounded above by $1/2$, but the term

²The CHIME correlator also has a quantizer with $N = 255$ levels at the digitization stage (see Section 5.3.1). For this quantizer the interval of optimal quantization is much broader, spanning several bits, and the correlation between the input and quantization error over this interval is even weaker ($|\rho_{ve}| \ll 10^{-14}$). The effects of this correlation are negligible compared to the $N = 15$ complex-valued quantizer at the channelization stage.

$1/(\sqrt{2\pi}\sigma^2)$ becomes arbitrarily large as σ decreases. Thus, $\langle ve \rangle$ is always positive and unbounded for N even in the low- σ regime.

6.3 Complex-valued quantizer

In the CHIME correlator, the (real-valued) analog signal of each input is first digitized and then passed through the F-engine that implements a PFB which splits the 400 MHz-wide input into 1024 frequency bins, each 390 kHz wide. The output of each frequency bin is a complex-valued signal and its real and imaginary parts are separately quantized with 15 levels before the data is re-arranged and sent to the X-engine for cross-multiplication and integration. In this section we extend the results of Section 6.2 to the case of an N -level complex-valued quantizer, where the real and imaginary parts of the input are separately quantized with N levels. In this case we assume that the input $v = v_r + jv_i$ is a complex and circularly-symmetric Gaussian process such that $\langle v_r v_i \rangle = 0$ and $\langle v_r^2 \rangle = \langle v_i^2 \rangle = \langle |v|^2 \rangle / 2$ where $\langle |v|^2 \rangle = \sigma^2$ is the unquantized standard deviation (see Section 2.3.4). As in Section 6.2 we are interested in the standard deviation of the quantization error, $e = e_r + je_i$, and its correlation with the input. In this case we have

$$\begin{aligned} \langle ve^* \rangle &= \langle (v_r + jv_i)(e_r + je_i)^* \rangle \\ &= \langle v_r e_r \rangle + \langle v_i e_i \rangle + j(-\langle v_r e_i \rangle + \langle v_i e_r \rangle). \end{aligned} \quad (6.9)$$

The circular symmetry of v (its real and imaginary part are uncorrelated and have identical statistics) implies that $\langle v_r e_r \rangle = \langle v_i e_i \rangle$. As for $\langle v_r e_i \rangle$, note that, for the m -th imaginary quantization level we have

$$\begin{aligned} \langle v_r e_i \rangle_m &= \int_{y_{m-1}}^{y_m} \int_{-\infty}^{\infty} (k_m - v_i) v_r f(v_r, v_i) dv_r dv_i \\ &= \int_{y_{m-1}}^{y_m} (k_m - v_i) \mathcal{N}\left(v_i \left| \frac{\sigma^2}{2} \right.\right) dv_i \int_{-\infty}^{\infty} v_r \mathcal{N}\left(v_r \left| \frac{\sigma^2}{2} \right.\right) dv_r \\ &= 0. \end{aligned} \quad (6.10)$$

Thus $\langle v_r e_i \rangle = 0$ and, for the same reason, $\langle v_i e_r \rangle = 0$. This means that $\langle ve^* \rangle$ is real and, from equation 6.6

$$\langle ve^* \rangle = 2\langle v_r e_r \rangle = \sigma^2 \left[-1 + \sum_{i=1}^{N-1} \mathcal{N}\left(-\frac{N}{2} + i \left| \frac{\sigma^2}{2} \right.\right) \right]. \quad (6.11)$$

From the circular symmetry of v it also follows that $\sigma_e^2 = 2\langle e_r^2 \rangle$ and $\hat{\sigma}^2 = 2\langle \hat{v}_r^2 \rangle$, so similar expressions for σ_e^2/σ^2 and $\hat{\sigma}^2/\sigma^2$ in the complex case are obtained from equations 6.6 - 6.8 by changing $\sigma^2 \rightarrow \sigma^2/2$.

Results from simulations and comparison to our prediction for the complex-valued quantizer with $N = 15$ levels (left column) and $N = 16$ levels (right column) are shown in Figure 6.2. From top to bottom row, the plots show the normalized variance of the quantized output ($\hat{\sigma}^2/\sigma^2$), the quantization error (σ_e^2/σ^2), and the magnitude and phase (in degrees) of the normalized correlation between the input and quantization error ($\langle ve^* \rangle/\sigma^2$). For each plot, the red line is our prediction and the blue line corresponds to the results from simulations where, for each value of σ , 10^6 samples of a complex and circularly-symmetric Gaussian input are quantized with N levels (real and imaginary parts quantized separately) and then the statistics of the input, output and quantization error are calculated.

There is again excellent agreement between the simulations and the predictions. The correlation between the input and quantization error, $\langle ve^* \rangle$, is always real (in the simulation the imaginary part is consistent with zero at the $\sim 0.15\%$ level). Furthermore, it is always negative (180° phase) for N odd, while it becomes positive (0° phase) in the low σ regime for N even. In this case the optimal quantization interval corresponding to $|\rho_{ve}| \lesssim 10^{-3}$ is approximately $[2^{-0.1}, 2^{1.4}]$ (the interval shifts by $\sqrt{2}$ with respect to the real-valued case).

6.4 Real-valued correlator

The correlation between two real-valued quantized inputs \hat{v}_1 and \hat{v}_2 , is

$$\hat{r}_{12} = \langle \hat{v}_1 \hat{v}_2 \rangle. \quad (6.12)$$

The output of a real-valued digital correlator after integrating N_s samples is

$$\hat{r}_{12, N_s} = \frac{1}{N_s} \sum_{n=1}^{N_s} \hat{v}_1[n] \hat{v}_2[n]. \quad (6.13)$$

Since the quantized sample vector $(\hat{v}_1[n], \hat{v}_2[n])$ comes from the IID joint Gaussian process (v_1, v_2) , then $\langle \hat{r}_{12, N_s} \rangle = \hat{r}_{12}$ so the measured correlation \hat{r}_{12, N_s} is an unbiased estimator of \hat{r}_{12} . Henceforth we will refer to \hat{r}_{12} as the output of the digital correlator.

Note that we already investigated the behavior of $\hat{r}_{11} = \hat{\sigma}_1^2$ and $\hat{r}_{22} = \hat{\sigma}_2^2$ in Section 6.2 (the result in this case is the same because the marginal PDFs of v_1 and v_2 are




Figure 6.2: Behavior of a complex-valued quantizer with $N = 15$ levels (left column) and $N = 16$ levels (right column) and a circularly-symmetric Gaussian input. From top to bottom row, the plots show the normalized variance of the quantized output ($\hat{\sigma}^2/\sigma^2$), the quantization error (σ_e^2/σ^2), and the magnitude and phase (in degrees) of the normalized correlation between the input and quantization error ($\langle ve^* \rangle/\sigma^2$). For each plot, the red line is our prediction and the blue line is the result from simulations. There is again excellent agreement between these. Note that $\langle ve^* \rangle$ is always real (in the simulation the imaginary part is consistent with zero at the $\sim 0.15\%$ level), and it is negative (180° phase) for N odd, while it becomes positive (0° phase) in the low σ regime for N even.

independent of the correlation between inputs). Now we are interested in \hat{r}_{12} and its relation to $r_{12} = \langle v_1 v_2 \rangle$ which is the correlation between the unquantized inputs v_1 and v_2 and what we ultimately want to measure. We can write \hat{r}_{12} as

$$\begin{aligned}\hat{r}_{12} &= \langle (v_1 + e_1)(v_2 + e_2) \rangle \\ &= r_{12} + \langle v_1 e_2 \rangle + \langle e_1 v_2 \rangle + \langle e_1 e_2 \rangle\end{aligned}\tag{6.14}$$

where

$$\langle v_1 e_2 \rangle = \sum_{i=1}^N \int_{y_{i-1}}^{y_i} \int_{-\infty}^{\infty} (k_i - v_2) v_1 f(v_1, v_2) dv_1 dv_2 \tag{6.15}$$

and

$$\langle e_1 e_2 \rangle = \sum_{j=1}^N \sum_{i=1}^N \int_{y_{j-1}}^{y_j} \int_{y_{i-1}}^{y_i} (k_i - v_1)(k_j - v_2) f(v_1, v_2) dv_1 dv_2. \tag{6.16}$$

$\langle e_1 v_2 \rangle$ is defined as in equation 6.15. If the samples from v_1 and v_2 come from a zero-mean joint Gaussian PDF

$$\mathcal{N}(v_1, v_2 | \sigma_1^2, \sigma_2^2, \rho) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} e^{-\frac{1}{2(1-\rho^2)} \left[\frac{v_1^2}{\sigma_1^2} + \frac{v_2^2}{\sigma_2^2} - \frac{2\rho v_1 v_2}{\sigma_1 \sigma_2} \right]} \tag{6.17}$$

where $\rho = \langle v_1 v_2 \rangle / (\sigma_1 \sigma_2)$, then equation 6.15 can be simplified

$$\begin{aligned}\langle v_1 e_2 \rangle &= \sum_{i=1}^N \int_{y_{i-1}}^{y_i} \int_{-\infty}^{\infty} (k_i - v_2) v_1 \mathcal{N}(v_1, v_2 | \sigma_1^2, \sigma_2^2, \rho) dv_1 dv_2 \\ &= \rho \frac{\sigma_1}{\sigma_2} \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v_2) v_2 \mathcal{N}(v_2 | \sigma_2^2) dv_2 \\ &= \rho \frac{\sigma_1}{\sigma_2} \langle v_2 e_2 \rangle \\ &= r_{12} \frac{\langle v_2 e_2 \rangle}{\sigma_2^2}.\end{aligned}\tag{6.18}$$

Similarly $\langle e_1 v_2 \rangle = r_{12} \langle v_1 e_1 \rangle / \sigma_1^2$. Note that with this result we can find both $\langle v_1 e_2 \rangle$ and $\langle e_1 v_2 \rangle$, which are correlations between mixed input-error terms, using equation 6.6

for the correlation between an input and its respective quantization error.

As for $\langle e_1 e_2 \rangle$ in equation 6.16, it can be simplified in the case when ρ is small, since in this regime we have

$$\mathcal{N}(v_1, v_2 | \sigma_1^2, \sigma_2^2, \rho) \big|_{\rho \ll 1} \approx \mathcal{N}(v_1 | \sigma_1^2) \mathcal{N}(v_2 | \sigma_2^2) \left(1 + \frac{\rho v_1 v_2}{\sigma_1 \sigma_2} \right) \quad (6.19)$$

so

$$\begin{aligned} \langle e_1 e_2 \rangle \big|_{\rho \ll 1} &\approx \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v_1) \mathcal{N}(v_1 | \sigma_1^2) dv_1 \sum_{j=1}^N \int_{y_{j-1}}^{y_j} (k_j - v_2) \mathcal{N}(v_2 | \sigma_2^2) dv_2 \\ &\quad + \frac{\rho}{\sigma_1 \sigma_2} \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v_1) v_1 \mathcal{N}(v_1 | \sigma_1^2) dv_1 \sum_{j=1}^N \int_{y_{j-1}}^{y_j} (k_j - v_2) v_2 \mathcal{N}(v_2 | \sigma_2^2) dv_2 \\ &= \langle e_1 \rangle \langle e_2 \rangle + \frac{\rho}{\sigma_1 \sigma_2} \langle v_1 e_1 \rangle \langle v_2 e_2 \rangle \\ &= r_{12} \frac{\langle v_1 e_1 \rangle}{\sigma_1^2} \frac{\langle v_2 e_2 \rangle}{\sigma_2^2}. \end{aligned} \quad (6.20)$$

Equation 6.20 will be useful when we analyze the phase behavior of the complex-valued correlator.

From equations 6.14 and 6.18 we can write

$$\hat{r}_{12} = r_{12} \left(1 + \frac{\langle v_1 e_1 \rangle}{\sigma_1^2} + \frac{\langle v_2 e_2 \rangle}{\sigma_2^2} \right) + \langle e_1 e_2 \rangle \quad (6.21)$$

and, using equation 6.20

$$\begin{aligned} \hat{r}_{12} \big|_{\rho \ll 1} &\approx r_{12} \left(1 + \frac{\langle v_1 e_1 \rangle}{\sigma_1^2} + \frac{\langle v_2 e_2 \rangle}{\sigma_2^2} + \frac{\langle v_1 e_1 \rangle}{\sigma_1^2} \frac{\langle v_2 e_2 \rangle}{\sigma_2^2} \right) \\ &= r_{12} \left(1 + \frac{\langle v_1 e_1 \rangle}{\sigma_1^2} \right) \left(1 + \frac{\langle v_2 e_2 \rangle}{\sigma_2^2} \right). \end{aligned} \quad (6.22)$$

The behavior from simulations of the normalized and quantized input correlation $r = \hat{r}_{12}/r_{12}$ and the contribution of the correlation between the quantization errors of the two inputs, $\langle e_1 e_2 \rangle$ (also normalized by r_{12}) are shown in Figures 6.3 and 6.4 respectively. For each value of σ_1 , σ_2 and ρ , 10^7 sample vectors $(v_1[n], v_2[n])$ from the joint Gaussian distribution in equation 6.17 are quantized with N levels and then both \hat{r}_{12} and $\langle e_1 e_2 \rangle$ are

calculated and normalized by the (measured) unquantized input correlation r_{12} . The axes for each plot are the unquantized input signal levels and the green solid lines correspond to the highest level of the quantizer (7 for $N = 15$ and 7.5 for $N = 16$) above which clipping occurs.

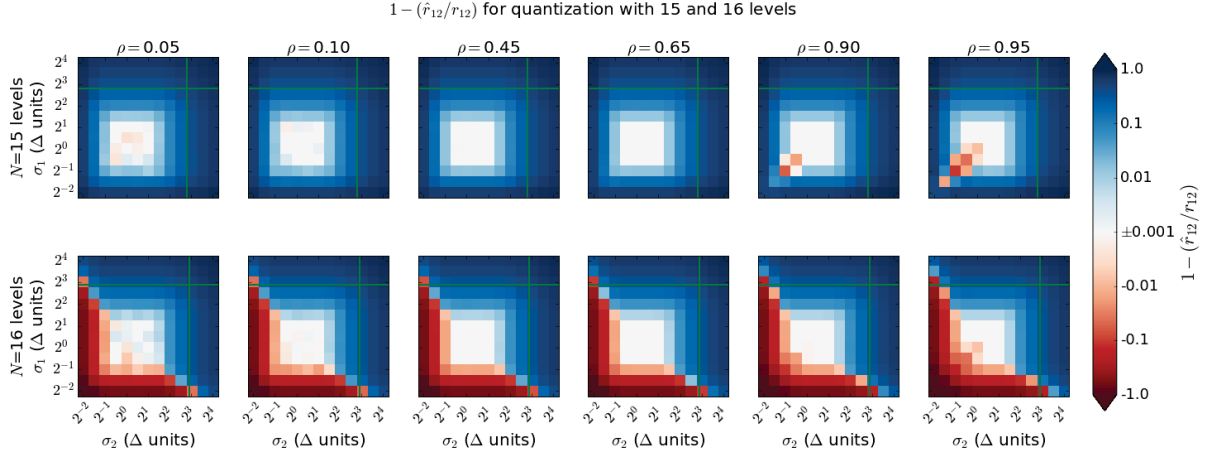


Figure 6.3: Results from simulations of $r = \hat{r}_{12}/r_{12}$ as function of σ_1 and σ_2 for different values of ρ for a real correlator with $N = 15$ levels (top row) and $N = 16$ levels (bottom row). The axes for each plot are the unquantized input signal levels and the green solid lines correspond to the highest level of the quantizer above which clipping occurs. The bias in \hat{r}_{12} for moderate values of ρ ($|\rho| \lesssim 0.85$) is below $\sim 0.1\%$ approximately within the inner white square enclosed by the region $\sigma_1 \times \sigma_2 \approx [2^{-0.6}, 2^{0.9}] \times [2^{-0.6}, 2^{0.9}]$. For $|\rho| \gtrsim 0.85$ the bias can increase up to $\sim 4\%$.

With 10^7 samples, the values in each pixel of Figure 6.3 agree with the values from equation 6.21 with unbiased error fluctuations below $\sim 1\%$. The worst case corresponds to low values of σ and ρ where r_{12} is very small. These results confirm that equation 6.21 accurately reproduces the relation between \hat{r}_{12} and r_{12} for the real-valued correlator.

For moderate values of ρ ($|\rho| \lesssim 0.85$) the bias in \hat{r}_{12} (Figure 6.3) is below $\sim 0.1\%$ (values from equation 6.21) approximately within the inner white square enclosed by the region $\sigma_1 \times \sigma_2 \approx [2^{-0.6}, 2^{0.9}] \times [2^{-0.6}, 2^{0.9}]$, corresponding to the region where both inputs are optimally quantized (see Section 6.2). For $|\rho| \gtrsim 0.85$ the bias within this region can increase up to $\sim 4\%$.

The most important feature from Figure 6.4 is that e_1 and e_2 are weakly correlated as long as at least one of the two inputs is approximately uncorrelated with its respective quantization error (either $\langle v_1 e_1 \rangle$ or $\langle v_2 e_2 \rangle$ is negligible). Another way to say this is that e_1 and e_2 are weakly correlated as long as at least one of the two inputs is optimally quantized, i.e., when the model of additive uncorrelated quantization noise is (approximately) valid. Note that this is what one would intuitively assume using the nominal model of

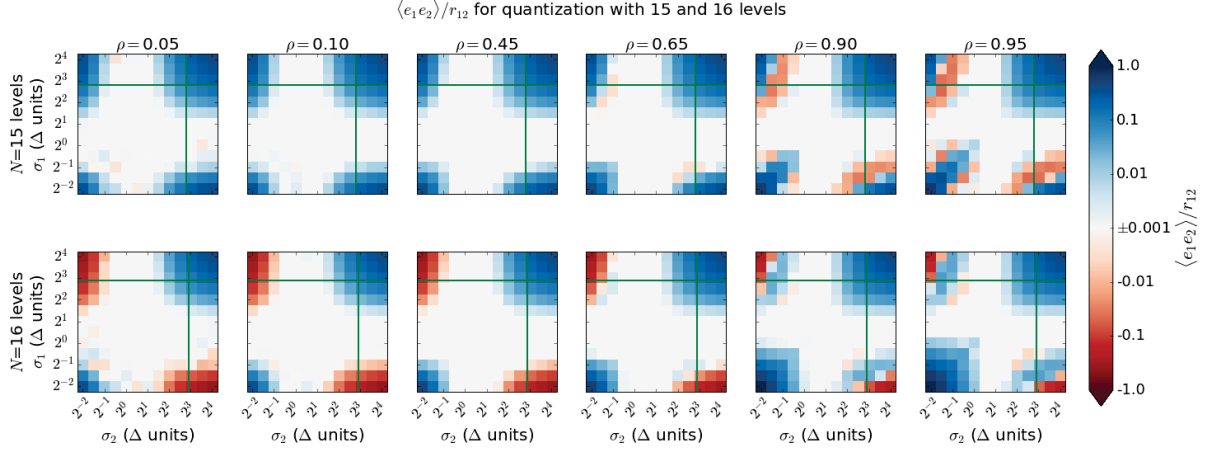


Figure 6.4: Correlation between the quantization errors of the two inputs, $\langle e_1 e_2 \rangle$ (normalized by r_{12}), as function of σ_1 , σ_2 and ρ , from simulations. Note that e_1 and e_2 are very weakly correlated as long as at least one of the two inputs is optimally quantized.

additive uncorrelated quantization error. We will use this result when we analyze the phase of the measured correlation in a complex correlator.

6.5 Complex-valued correlator

Now we extend the results from Section 6.4 to the case when the correlator inputs are complex-valued, such as for the complex channelization stage of the CHIME correlator, where the digitized inputs are channelized using a PFB that splits the 400 MHz-wide input into 1024 narrow frequency bins. The complex-valued output of each frequency bin is quantized with $N = 15$ levels for both the real and imaginary parts. Finally, the quantized signals are sent to the correlator that measures complex-valued correlation between quantized inputs, $\hat{r}_{12} = \langle \hat{v}_1 \hat{v}_2^* \rangle$. We are ultimately interested in $r_{12} = \langle v_1 v_2^* \rangle$ so we need to find a relation between these.

As in Section 6.3, we assume that $\mathbf{v} = (v_1, v_2)$ is a complex and circularly-symmetric Gaussian process. Then

$$\begin{aligned} \hat{r}_{12} &= \langle (\hat{v}_{1r} + j\hat{v}_{1i})(\hat{v}_{2r} - j\hat{v}_{2i}) \rangle \\ &= \langle \hat{v}_{1r}\hat{v}_{2r} \rangle + \langle \hat{v}_{1i}\hat{v}_{2i} \rangle + j(-\langle \hat{v}_{1r}\hat{v}_{2i} \rangle + \langle \hat{v}_{1i}\hat{v}_{2r} \rangle) \end{aligned} \quad (6.23)$$

The circular symmetry of \mathbf{v} implies that $\langle \hat{v}_{1r}\hat{v}_{2r} \rangle = \langle \hat{v}_{1i}\hat{v}_{2i} \rangle$ and $-\langle \hat{v}_{1r}\hat{v}_{2i} \rangle = \langle \hat{v}_{1i}\hat{v}_{2r} \rangle$

so

$$\begin{aligned}\hat{r}_{12} &= 2 [\langle \hat{v}_{1r} \hat{v}_{2r} \rangle + j \langle \hat{v}_{1i} \hat{v}_{2r} \rangle] \\ &= 2(\hat{r}_{1r,2r} + j \hat{r}_{1i,2r}).\end{aligned}\tag{6.24}$$

For $\hat{r}_{1r,2r} = \langle \hat{v}_{1r} \hat{v}_{2r} \rangle$ and $\hat{r}_{1i,2r} = \langle \hat{v}_{1i} \hat{v}_{2r} \rangle$ which are real, we can use equation 6.21 so

$$\begin{aligned}\hat{r}_{12} &= 2 \left\{ \left[r_{1r,2r} \left(1 + \frac{\langle v_{1r} e_{1r} \rangle}{\sigma_{1r}^2} + \frac{\langle v_{2r} e_{2r} \rangle}{\sigma_{2r}^2} \right) + \langle e_{1r} e_{2r} \rangle \right] + \right. \\ &\quad \left. j \left[r_{1i,2r} \left(1 + \frac{\langle v_{1i} e_{1i} \rangle}{\sigma_{1i}^2} + \frac{\langle v_{2r} e_{2r} \rangle}{\sigma_{2r}^2} \right) + \langle e_{1i} e_{2r} \rangle \right] \right\} \\ &= 2(r_{1r,2r} + j r_{1i,2r}) \left(1 + \frac{\langle v_{1r} e_{1r} \rangle}{\sigma_{1r}^2} + \frac{\langle v_{2r} e_{2r} \rangle}{\sigma_{2r}^2} \right) + 2(\langle e_{1r} e_{2r} \rangle + j \langle e_{1i} e_{2r} \rangle) \\ &= r_{12} \left(1 + \frac{\langle v_{1r} e_{1r} \rangle}{\sigma_{1r}^2} + \frac{\langle v_{2r} e_{2r} \rangle}{\sigma_{2r}^2} \right) + 2(\langle e_{1r} e_{2r} \rangle + j \langle e_{1i} e_{2r} \rangle)\end{aligned}\tag{6.25}$$

where in the second step we used the fact that $\langle v_{1r} e_{1r} \rangle / \sigma_{1r}^2 = \langle v_{1i} e_{1i} \rangle / \sigma_{1i}^2$ and in the third step we used $r_{12} = 2(r_{1r,2r} + j r_{1i,2r})$. All these follow from circular symmetry. Note that all the terms in equation 6.25 can be obtained from equations 6.6 and 6.16 using $\sigma_{1r}^2 = \sigma_{1i}^2 = \sigma_1^2/2$ and $\sigma_{2r}^2 = \sigma_{2i}^2 = \sigma_2^2/2$.

We can use equation 6.25 to draw some important conclusions regarding how quantization affects the magnitude and phase of r_{12} . We can write

$$\hat{r}_{12} = \alpha r_{12} + \beta, \quad \alpha = \left(1 + \frac{\langle v_{1r} e_{1r} \rangle}{\sigma_{1r}^2} + \frac{\langle v_{2r} e_{2r} \rangle}{\sigma_{2r}^2} \right), \quad \beta = 2(\langle e_{1r} e_{2r} \rangle + j \langle e_{1i} e_{2r} \rangle)\tag{6.26}$$

Note that α is real, independent of ρ , and only contributes to the biasing of the magnitude of \hat{r}_{12} . On the other hand, β is complex in general and affects both the magnitude and phase of \hat{r}_{12} .

Quantization will bias the magnitude of \hat{r}_{12} except when $\alpha = 1$ and $\beta = 0$. This occurs approximately when both inputs are optimally quantized since in this case $\langle v_{1r} e_{1r} \rangle \rightarrow 0$, $\langle v_{2r} e_{2r} \rangle \rightarrow 0$ (so $\alpha \rightarrow 1$, see Section 6.2 and Figure 6.1), and also $\langle e_{1r} e_{2r} \rangle \rightarrow 0$, $\langle e_{1i} e_{2r} \rangle \rightarrow 0$ (so $\beta \rightarrow 0$, see Section 6.4 and Figure 6.4).

Quantization will bias the phase of \hat{r}_{12} except in two cases: the first case is when $\beta = 0$, which occurs approximately when at least one of the inputs is optimally quantized (see Section 6.4 and Figure 6.4). Note that this is a less stringent requirement than that for

unbiased magnitude, which requires both inputs to be optimally quantized.

The second case for negligible bias in the phase of \hat{r}_{12} occurs when $\rho \ll 1$ since using equation 6.20 in equation 6.25 we have

$$\begin{aligned}\hat{r}_{12}|_{\rho \ll 1} &\approx 2r_{1r,2r} \left(1 + \frac{\langle v_{1r}e_{1r} \rangle}{\sigma_{1r}^2}\right) \left(1 + \frac{\langle v_{2r}e_{2r} \rangle}{\sigma_{2r}^2}\right) + 2jr_{1i,2r} \left(1 + \frac{\langle v_{1i}e_{1i} \rangle}{\sigma_{1i}^2}\right) \left(1 + \frac{\langle v_{2r}e_{2r} \rangle}{\sigma_{2r}^2}\right) \\ &= r_{12} \left(1 + \frac{\langle v_{1r}e_{1r} \rangle}{\sigma_{1r}^2}\right) \left(1 + \frac{\langle v_{2r}e_{2r} \rangle}{\sigma_{2r}^2}\right).\end{aligned}\tag{6.27}$$

Since the factors that multiply r_{12} are real then $\angle(\hat{r}_{12}) = \angle(r_{12})$.

Figures 6.5 and 6.6 show results from simulations of \hat{r}_{12}/r_{12} (magnitude and phase respectively. The phase is in degrees). The method is the same as in Section 6.2, but this time the 10^7 sample vectors $(v_1[n], v_2[n])$ are drawn from a circularly symmetric Gaussian distribution. We only vary the magnitude of ρ , keeping its phase fixed at 75 degrees.

Note that equations 6.25-6.27 predict accurately the behavior of the magnitude and phase of \hat{r}_{12} . For moderate values of ρ ($|\rho| \lesssim 0.85$) the bias in the magnitude (Figure 6.5) is below $\sim 0.1\%$ roughly within the inner square enclosed by the region $\sigma_1 \times \sigma_2 \approx [2^{-0.1}, 2^{1.4}] \times [2^{-0.1}, 2^{1.4}]$, corresponding to region where both inputs are optimally quantized (see Section 6.3). For $|\rho| \gtrsim 0.85$ the bias within this region can increase up to $\sim 4\%$.

As for the phase (Figure 6.6), the bias is below $\sim 0.1^\circ$ within the cross-shaped region where either σ_1 or σ_2 are optimally quantized. When $|\rho| \gtrsim 0.85$ the bias within this region can rise up to $\sim 1^\circ$. When $\rho \lesssim 0.1$ (first two columns of Figure 6.6) the phase bias is below $\sim 0.1^\circ$ (values from equation 6.25) for all values of σ_1 and σ_2 as predicted by equation 6.27, although there are still random fluctuations in the simulation at the \sim sub-degree level for very low values of σ (bottom and left edges of the plots) for reasons explained in Section 6.4.

6.6 Implications for radio interferometry

The results above have important implications for radio interferometry. Quantization will have a significant biasing effect on the visibility magnitude unless both inputs are optimally quantized, which can be a stringent requirement (both signal levels need to be in the region where the uncorrelated quantization model is valid). However, we have found that the bias in the visibility phase is negligible even in conditions as extreme as

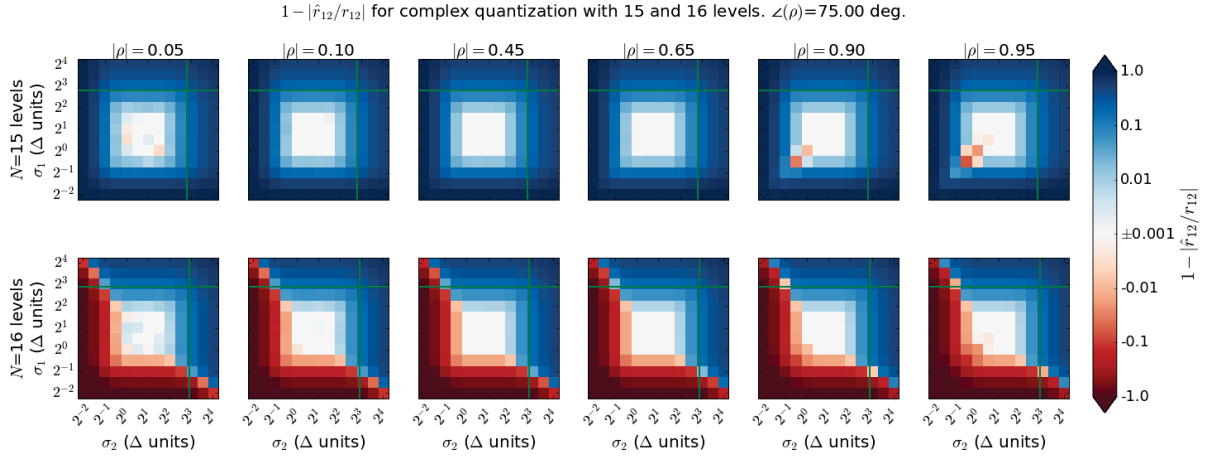


Figure 6.5: $|\hat{r}_{12}/r_{12}|$ from simulations as function of σ_1 and σ_2 for different values of $|\rho|$. For moderate values of ρ the bias in the magnitude of $|\hat{r}_{12}|$ is below $\sim 0.1\%$ within the inner square enclosed by the region $\sigma_1 \times \sigma_2 \approx [2^{-0.1}, 2^{1.4}] \times [2^{-0.1}, 2^{1.4}]$, corresponding to region where both inputs are optimally quantized. For $|\rho| \gtrsim 0.85$ this bias can increase up to $\sim 4\%$.




Figure 6.6: $\angle(\hat{r}_{12}/r_{12})$ (in degrees) from simulations as function of σ_1 and σ_2 for different values of $|\rho|$. The bias in the phase of \hat{r}_{12} is negligible when at least one of the inputs is optimally quantized. This is a less stringent requirement than that for the magnitude, which requires both inputs to be optimally quantized. The bias is below $\sim 0.1^\circ$ within the cross-shaped region where either σ_1 or σ_2 are in the approximate interval $[2^{-0.1}, 2^{1.4}]$. When $|\rho|$ is high (last two columns) the bias within this region can rise up to $\sim 1^\circ$. When $|\rho|$ is small (first two columns) the phase bias is below $\sim 0.1^\circ$ for all values of σ_1 and σ_2 , although there are still random fluctuations in the simulation at the \sim sub-degree level for very low values of σ (see text).

when one of the inputs is suffering from severe clipping, or even when both inputs are severely clipped in the case of weak sources ($|\rho| \ll 1$). The same conditions apply when one or both input levels are very low (note that any of these extreme conditions will affect the signal-to-noise ratio of the measured visibility even if the phase is unbiased, but that analysis is left for future work). An accurate determination of the visibility phase is critical for beamforming, fringe stopping, and image reconstruction techniques.

For the particular case of CHIME, in which the sky signals are weak and the correlator inputs are dominated by the noise of the analog receiving system, the correlation coefficient is typically low ($|\rho| \lesssim 0.1$) even for the brightest radio point sources such as CasA, CygA, and TauA, but excluding the sun. This means that, except for the time when the sun is in the primary beam of the CHIME telescope (~ 14 minutes per day), all the visibility phases will have negligible bias due to quantization.

The quantization bias also has an effect on the beamformed sensitivity of a radio interferometric array. To illustrate this, consider a one-dimensional array consisting of uniformly spaced feeds located at positions $0, 1, \dots, N_f - 1$, in units of the normalized feed spacing $b_\lambda = b/\lambda$, where λ is the observed wavelength. This example corresponds to one of the cylinders of the CHIME telescope, where the feeds are uniformly spaced along the axis of the cylinder. The cylinder axis (and thus the linear array) is oriented North-South (N-S), so the resolution in the N-S direction is provided by the correlations between feeds. We will assume that all the feeds have identical beams that are N-S isotropic and receivers with system noise σ_{sys}^2 , although the generalization is straightforward.

For a point source on the meridian with noise temperature σ^2 such that the signal-to-system-noise ratio is $SNR = \sigma^2/\sigma_{sys}^2$, the unquantized autocorrelations for each feed are identical and equal to

$$r_{ii} = \sigma_{sys}^2(1 + SNR), \quad i = 0, 1, \dots, N_f - 1 \quad (6.28)$$

while the unquantized visibility and correlation coefficient between feeds i and j are

$$r_{ij} = SNR \cdot \sigma_{sys}^2 e^{-j2\pi(i-j)b_\lambda \sin \theta}, \quad \rho_{ij} = \frac{e^{-j2\pi(i-j)b_\lambda \sin \theta}}{1 + \frac{1}{SNR}} \quad (6.29)$$

where θ is the source zenith angle and we have assumed uncorrelated system noise between feeds.

The bias due to quantization of the measured visibility as function of SNR and θ for $i = j + 1$ (consecutive feeds) is shown in Figure 6.7. We use $b_\lambda = 0.4$ which corresponds to the CHIME normalized feed spacing at 400 MHz. These results are obtained directly from equation 6.25.




Figure 6.7: Bias due to quantization of the measured visibility as function of the source position θ and the signal-to-system-noise ratio is SNR . The visibility baseline is $b_\lambda = 0.4$ which corresponds to the CHIME normalized feed spacing at 400 MHz. When $\sigma_{sys} = 2^1 \Delta$ (left panels), well within the optimal quantization interval, the quantization bias for weak sources ($SNR \lesssim 0.1$) is negligible. This is the regime for CHIME $\sim 99\%$ of the time. When $\sigma_{sys} = 2^2 \Delta$ (right panels), which is the optimal input level according to the uncorrelated quantization noise model, the amount of bias increases significantly.

To illustrate the difference between equation 6.25 and the uncorrelated quantization noise model, and the importance of optimizing the input signal level of the quantizer, Figure 6.7 shows the bias for two different values of σ_{sys} : $2^1 \Delta$ (left panels) and $2^2 \Delta$ (right panels). For a system-noise dominated telescope like CHIME, the correlator inputs are calibrated so σ_{sys} corresponds to the optimal input level of the quantizer in order to minimize the effects of quantization. For $N = 15$, the optimal input level according to the uncorrelated quantization noise model is $\sigma_{sys} \approx 2^2 \Delta$, corresponding to the point where σ_e is minimum (see second row of Figure 6.2). On the other hand, equation 6.25 suggests that a better choice for σ_{sys} should be more centered around the optimal quantization interval $[2^{-0.1} \Delta, 2^{1.4} \Delta]$. The CHIME post-channelization calibration module uses $\sigma_{sys} \approx 2^1 \Delta$, which is well within this interval while still keeping σ_e relatively low (see second and third rows of Figure 6.2, if σ_{sys} is too close to the lower end of the interval then the contribution of σ_e is significant).

Note that the SNR sets the overall amount of bias due to quantization since this parameter defines both r_{ii} and $|\rho_{ij}|$ (equations 6.28 and 6.29). For $\sigma_{sys} = 2^1 \Delta$ and $|\rho_{ij}| \lesssim 0.1$ (so $SNR \lesssim 0.1$) the magnitude bias is $\lesssim 10^{-6}$ and the phase bias is $\lesssim 10^{-11}$ degrees, too small to have any significant impact that requires the generalized Van Vleck

correction from [70]. As mentioned before, this is the regime for CHIME $\sim 99\%$ of the time. However, when the sun is in the main beam ($\sim 1\%$ of the time), the SNR can be as high as ~ 6 (green line in Figure 6.7), corresponding to a magnitude bias of $\sim 11\%$ and a phase bias of up to $\sim 0.15^\circ$. Although the CHIME cosmology data pipeline masks out the sun time, this data is still very useful for beam mapping purposes. The quantization bias is significant enough in this case to justify the implementation of either the generalized Van Vleck correction or a digital calibration scheme that dynamically adjusts the input signal level during strong-source transits such that the quantizer always operates in the optimal quantization regime³.

When $\sigma_{sys} = 2^2\Delta$ (right side of Figure 6.7) the amount of bias increases significantly even in the weak-source regime. For $SNR \sim 0.1$ the magnitude bias is $\sim 3\%$ and the phase bias is $\sim 3 \times 10^{-3}$ degrees, while for $SNR \sim 6$ the magnitude bias is $\sim 56\%$ and the phase bias is $\sim 1^\circ$, demonstrating that for this particular application the uncorrelated quantization noise model must be used carefully since it can introduce important effects in the measured visibilities.

The quantization bias also depends on the position of the source and the baseline. These parameters determine $\angle(\rho_{ij})$ which affects the measured visibility \hat{r}_{ij} through the second term of equation 6.25. As Figure 6.7 shows, the position dependence manifests as fringes as a function of $l = \sin \theta$, where the baseline determines the quantization fringe rate.

We can use the $N_f(N_f-1)/2$ visibilities (excluding the autocorrelations) to beamform in the direction of the source. Since for $k = i-j$ fixed there are (N_f-k) identical baselines, then we can write the quantized beamformed output as

$$\hat{R} = \sum_{i>j}^{N_f-1} \hat{r}_{ij} e^{j2\pi(i-j)b_\lambda \sin \theta} = \sum_{k=1}^{N_f-1} (N_f - k) \hat{r}_k e^{j2\pi k b_\lambda \sin \theta} \quad (6.30)$$

while the unquantized beamformed output is

$$R = \sum_{k=1}^{N_f-1} (N_f - k) r_k e^{j2\pi k b_\lambda \sin \theta} = \frac{N_f(N_f - 1)}{2} SNR \sigma_{sys}^2. \quad (6.31)$$

³Note that we are assuming that the sun is a point source to simplify the analysis since we are interested in studying the behavior of quantization for strong sources. Although, strictly speaking, the sun is an extended source for CHIME, for observations with the CHIME pathfinder this is an adequate approximation.

We can define a complex quantization parameter

$$\eta_q = \frac{\hat{R}}{R} \quad (6.32)$$

as a measure of the beamforming efficiency due to quantization. Figure 6.8 shows the magnitude and phase of η_q as function of the source position θ for $SNR = 0.1$ (approximate upper limit of weak-source regime) and $SNR = 6$ (typical strong source like the sun). We used $N_f = 32$ and kept b_λ fixed at 0.4.

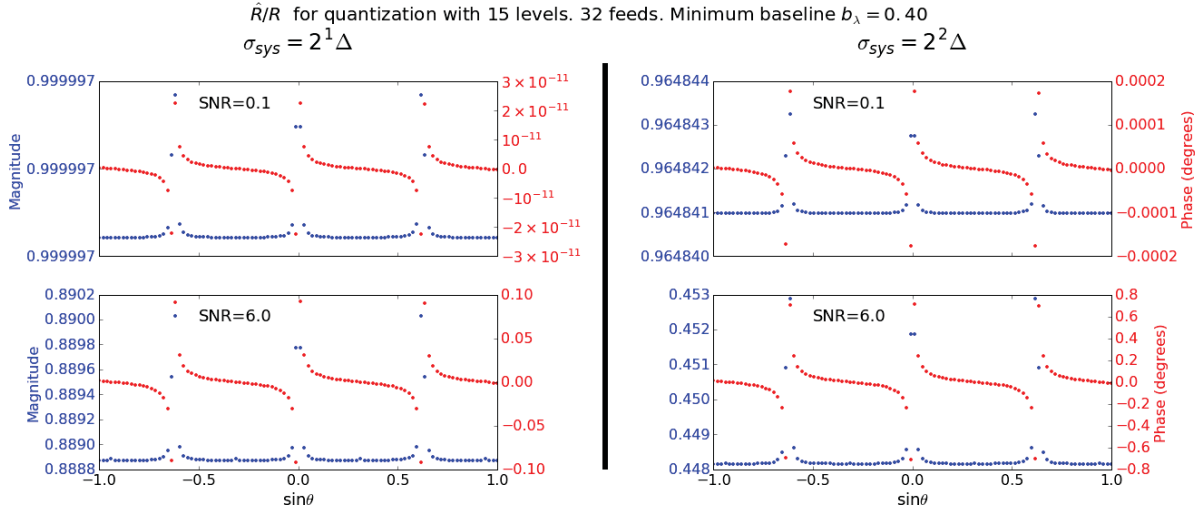


Figure 6.8: Complex quantization parameter $\eta_q = \hat{R}/R$ as function of the source position θ for $SNR = 0.1$ (top row, this SNR is the approximate upper limit of weak-source regime) and $SNR = 6$ (bottom row, this is the typical SNR of a strong source like the sun). For each plot, the blue labels and dots correspond to the magnitude of η_q and the red labels and dots correspond to its phase in degrees. Note that for $\sigma_{sys} = 2^1\Delta$ (left column), which is well within the optimal quantization interval for $N = 15$ levels, and in the weak-source regime ($SNR \lesssim 0.1$, top left plot), η_q is very close to being real-valued and deviates from unity by less than $\sim 3 \times 10^{-6}$ so the loss of beamforming efficiency due to quantization is negligible. If we set $\sigma_{sys} = 2^2\Delta$ (right column), the beamforming sensitivity reduces significantly even in the weak-source regime. This confirms that for this application the uncorrelated quantization model leads to important deviations from the expected performance of the interferometric array.

The most important feature from Figure 6.8 is that for $\sigma_{sys} = 2^1\Delta$ (left column) and in the weak-source regime ($SNR \lesssim 0.1$, top left plot) the loss of beamforming efficiency due to quantization is negligible (η_q is very close to being real-valued and deviates from unity by less than $\sim 3 \times 10^{-6}$). However, for a strong source like the sun the beamforming efficiency decreases below $\sim 89\%$ (bottom left plot). When σ_{sys} is set to $2^2\Delta$ (right column) the beamforming sensitivity reduces to $\sim 96\%$ and $\sim 45\%$ for $SNR = 0.1$

and 6 respectively, confirming that the uncorrelated quantization noise model leads to important deviations from the expected interferometer performance.

6.7 Conclusions

In this chapter we investigated the correlation between the input and the quantization error of a digitizer with uniformly spaced levels and an odd symmetric transfer function, which is the case for CHIME and other radio astronomy applications. We then used these results to explore the biasing effect of quantization in the correlation measured by a complex-valued digital correlator.

We showed that, for a complex-valued quantizer with a circularly symmetric Gaussian input, the correlation between the input and the quantization error is always real. It is always negative when the number of levels N of the quantizer is odd, while for N even this correlation is positive in the low signal level regime. In both cases there is an interval for the signal level σ (which we denote the interval of optimal quantization) for which this input-error correlation is very weak and the uncorrelated quantization error model provides a very accurate approximation. The length of the optimal quantization interval depends on N and on the tolerance required by each specific application.

With these results we determined the quantization bias in the correlations measured by a digital correlator and derived the conditions under which the bias in the magnitude and phase of the measured correlation is negligible with respect to the unquantized values: we demonstrated that the magnitude bias is negligible only if both unquantized inputs are optimally quantized, while the phase bias is negligible when 1) at least one of the inputs is optimally quantized, or when 2) the correlation coefficient ρ between the unquantized inputs is small.

These results are important for radio interferometry where the correlations measured by the digital correlator provide the interferometric visibilities. Although quantization will bias significantly the visibility magnitude unless both inputs are optimally quantized, which can be a stringent requirement, we showed that the bias in the visibility phase is negligible even in extreme conditions such as when one of the inputs is in the high signal-level regime with large amounts of clipping or when it is in the low signal-level regime where the contribution of the quantization error to the quantized output is very high. Even when both inputs are far from the optimal quantization regime (either because of extreme clipping or very low signal level) the phase quantization bias is negligible for weak sources ($|\rho| \ll 1$). This is the typical case for interferometers like CHIME where the analog inputs are dominated by the receiver noise. In this regime all the visibility

phases will be approximately unbiased regardless of the signal levels.

Finally, we demonstrated using a specific example corresponding to a CHIME-like array of antennas that quantization reduces the point-source sensitivity of a radio interferometric array. For a system-noise dominated telescope like CHIME, this effect can be reduced to negligible levels in the weak-source regime with a suitable scaling of the system noise level at the input of the quantizer.

The CHIME post-channelization calibration module described in Section 5.3.3 adjusts the signal level for each frequency channel so the nominal RMS at the input of the 4-bit complex quantizer is approximately two quantization levels ($\sigma = 2^1 \Delta$). This level is well within the optimal quantization interval where the input-error correlation is very weak ($|\rho_{ve}| \ll 10^{-3}$) so the model of additive uncorrelated quantization noise provides a very accurate approximation. At this level the quantization noise is kept relatively low, with a fractional increase in the output variance of about 4% (see Figure 6.2). Furthermore, the bias in the measured correlations due to quantization is negligible: the magnitude bias is below $\sim 10^{-6}$ and the phase bias is below $\sim 10^{-11}$ degrees in the weak source regime. This signal level also allows headroom for changes in the signal level due to bright point sources and receiver gain fluctuations, while still keeping the quantizer within the optimal region of operation.

We are currently developing the extension of this analysis to other probability density functions beyond the nominal Gaussian case. In particular, we are investigating the effect of quantization in the presence of RFI. We are also investigating how quantization affects the signal-to-noise ratio of the measured visibilities when the unquantized inputs are far from the optimal quantization regime, either because of extreme clipping or very low signal level.

Chapter 7

Conclusions

In this thesis we presented important new hardware, analysis and software contributions to the design, construction, and characterization of the Canadian Hydrogen Intensity Mapping Experiment (CHIME). CHIME is a groundbreaking instrument designed to map the large scale structure of neutral hydrogen in the universe by directly detecting its redshifted 21 cm radiation. By measuring the BAO scale across the redshift range $z \approx 0.8 - 2.5$ in both the angular and line-of-sight directions, CHIME will study the epoch where the expansion history of the universe transitions from one dominated by the attractive force of ordinary gravity to one dominated by dark energy.

CHIME consists of four cylindrical reflectors populated with a total of 1024 dual-polarization feeds. It operates as a transit interferometer that maps approximately half of the sky each day with a synthesized beam resolution of 12-24 arcminutes. Its drift-scan strategy, large field of view, and powerful correlator also make CHIME an excellent instrument to study the time-variable sky, including pulsar monitoring and detection of radio transients and FRBs.

The CHIME correlator processes 2048 digitizer inputs across 400 MHz of bandwidth and, measured in number of baselines times bandwidth, it is the largest radio correlator that has been built, performing 8.4×10^{14} complex operations per second. The correlator is based on an FX design. The frequency channelization is performed in FPGAs, which are interfaced to a GPU-based X-engine that computes the correlation matrix. We have developed the ICE system, an FPGA-based general purpose hardware, firmware, and software framework that has been specialized to implement the data acquisition, F-engine, and the networking engine that re-arranges the data before being sent to the X-engine for cross-multiplication and averaging. We presented details of the ICE-system design and its use for the CHIME correlator. We gave details of the deployment of the correlator and demonstrated that it complies with all the requirements for CHIME by

performing tests of manufacturability, uptime and reliability, data transport integrity, analog performance and data acquisition integrity, timing, and power consumption.

The most important challenge for CHIME is separating the 21 cm signal from astrophysical foregrounds that are many orders of magnitude brighter. Our ability to reconstruct the 21 cm power spectrum in the presence of foregrounds requires detailed knowledge of the instrument, including a precise measurement of the complex receiver gains and a detailed calibration of the digital correlator. In order to address these stringent requirements and understand the CHIME instrument, a number of receiver-gain and digital calibration techniques have been developed using the CHIME pathfinder which is a proof-of-concept instrument for full CHIME with 20% of its collecting area and 128 dual-polarization feeds.

For receiver gain calibration we implemented a Broadband Injection Signal (BIS) system that injects broadband calibration signals across the array to measure and correct for the receiver gain fluctuations. We presented the BIS formalism, which can be extended to other calibration methods including bright point-source and pulsar transits. We also presented the BIS instrumentation and its improvements with respect to the early system presented in [5]. These improvements include a system that is synchronized with the data acquisition, a faster switching cadence to improve the sky background and RFI subtraction, and the possibility to control multiple BIS sources independently. We also deployed a simpler version of the BIS system that is currently used to correct the intra and inter-crate jitter of the ICE system on the pathfinder and full CHIME.

Since the CHIME analog receivers have been found to be more stable than initially anticipated, a detailed thermal characterization of the analog receivers may be enough to correct gain fluctuations at the level of the calibration requirements, in which case the implementation of the complete version of the BIS system on full CHIME may not be necessary. However, BIS will still be used to correct the timing jitter which is the dominant source of phase noise on full CHIME.

For digital calibration, we have developed efficient calibration techniques that minimize the effects of quantization errors at three critical correlator stages: the 8-bit digitization of the analog signals, the computation of the FFT, and the post-channelization 4-bit complex quantization. For the digitization stage, the gain of the analog receiving system has been optimized so the nominal signal level at the input of the ADC is approximately 3.5 bits RMS. At this level the quantization noise is negligible (the increase in variance is well below 0.1%) and the 8-bit quantizer is within the optimal region of operation where the input-error correlation is very weak ($|\rho_{ve}| \ll 10^{-14}$). For the FFT computation, the FPGA firmware and software allow the configuration of different scal-

ing schemes for the FFT stages. We have implemented an optimal scaling pattern that maximizes the signal-to-quantization-error of the FFT output for signals with a smooth spectra, as is the nominal case for CHIME. For the post-channelization quantization stage we have implemented an efficient post-channelization calibration algorithm that applies a digital gain to each frequency channel in order to minimize the effects of the 4-bit complex-valued quantization that takes place at the output of the FFT. By setting the signal level of each frequency channel to 1 bit RMS, the correlation between the input and the complex quantization error is very weak ($|\rho_{ve}| \ll 10^{-3}$), the quantization noise is kept relatively low (the increase in output variance is about 4%), and the resulting bias in the measured correlations due to quantization is negligible (below $\sim 10^{-6}$ for the magnitude and $\sim 10^{-11}$ degrees phase).

The optimization and calibration algorithms developed for the digitization and post-channelization quantization stages above are based on the results of the detailed analysis that we have performed on the biasing effect of quantization in the correlations measured by digital correlators. This analysis goes beyond the commonly used approximation where the quantization noise is assumed to be uncorrelated with the input. We demonstrated that, for real and complex-valued quantizers, there is an optimal interval for the input signal level for which the correlation between the input and the quantization error is very weak and the uncorrelated quantization error model provides a very accurate approximation. We also determined the conditions under which the magnitude and phase of the measured correlation have negligible bias with respect to the unquantized values: we demonstrated that the magnitude bias is negligible only if both unquantized inputs are optimally quantized, while the phase bias is negligible when 1) at least one of the inputs is optimally quantized, or when 2) the correlation coefficient between the unquantized inputs is small. Finally, we demonstrated how, for a system-noise dominated telescope like CHIME, this biasing effect can be reduced to negligible levels with a suitable scaling of the signal level at the input of the quantizer. We are currently developing the extension of this analysis to the case where the astronomical signal is contaminated by RFI. We are also investigating how quantization affects the signal-to-noise ratio of the measured visibilities when the unquantized inputs are far from the optimal quantization regime, either because of extreme clipping or very low signal level.

The CHIME telescope is operational and collecting early science data. There is an ongoing effort to characterize the instrument to the level required to detect the 21 cm power spectrum. The detailed analysis and knowledge from this thesis will serve to optimize this characterization and ensure that the experiment completes its science goals of understanding the nature of dark energy and radio transients.

Appendix A

Power spectrum of $\delta_{21}(\mathbf{r})$

Let $\delta_{21}(\mathbf{r})$ be the random field defined in equation 1.32. We say that $\delta_{21}(\mathbf{r})$ is homogeneous if its mean (zero in this case) and covariance (equal to its correlation function since it is a zero-mean random field) $\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle$ are invariant under translations. For the covariance this means

$$\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle = \langle \delta_{21}(\mathbf{r}_1 + \delta \mathbf{r}) \delta_{21}(\mathbf{r}_2 + \delta \mathbf{r}) \rangle. \quad (\text{A.1})$$

If $\delta_{21}(\mathbf{r})$ is homogeneous it can be shown that its two-point correlation function depends on $\mathbf{r}_{12} = \mathbf{r}_1 - \mathbf{r}_2$ only, that is, $\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle = \xi_{21}(\mathbf{r}_{12})$. In this case the correlation of $\tilde{\delta}_{21}(\mathbf{k})$ becomes

$$\begin{aligned} \langle \tilde{\delta}_{21}(\mathbf{k}_1) \tilde{\delta}_{21}^*(\mathbf{k}_2) \rangle &= \int_{-\infty}^{\infty} d^3 \mathbf{r}_1 \int_{-\infty}^{\infty} d^3 \mathbf{r}_2 \langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle e^{-i(\mathbf{k}_1 \cdot \mathbf{r}_1 - \mathbf{k}_2 \cdot \mathbf{r}_2)} \\ &= \int_{-\infty}^{\infty} d^3 \mathbf{r} \int_{-\infty}^{\infty} d^3 \mathbf{r}_2 \xi_{21}(\mathbf{r}) e^{-i(\mathbf{k}_1 \cdot \mathbf{r} + \mathbf{k}_1 \cdot \mathbf{r}_2 - \mathbf{k}_2 \cdot \mathbf{r}_2)} \\ &= \int_{-\infty}^{\infty} d^3 \mathbf{r} \xi_{21}(\mathbf{r}) e^{-i\mathbf{k}_1 \cdot \mathbf{r}} \int_{-\infty}^{\infty} d^3 \mathbf{r}_2 e^{-i(\mathbf{k}_1 - \mathbf{k}_2) \cdot \mathbf{r}_2} \\ &= (2\pi)^3 P_{21}(\mathbf{k}_1) \delta(\mathbf{k}_1 - \mathbf{k}_2) \end{aligned} \quad (\text{A.2})$$

where $P_{21}(\mathbf{k}) = \int_{-\infty}^{\infty} d^3 \mathbf{r} \xi_{21}(\mathbf{r}) e^{-i\mathbf{k} \cdot \mathbf{r}}$ is the Fourier Transform of $\xi_{21}(\mathbf{r})$ (the power spectrum of $\delta_{21}(\mathbf{r})$).

We say that $\delta_{21}(\mathbf{r})$ is isotropic if its mean and covariance $\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle$ are invariant under rotations. For the covariance this means

$$\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle = \langle \delta_{21}(g\mathbf{r}_1) \delta_{21}(g\mathbf{r}_2) \rangle \quad g \in \text{SO}(3). \quad (\text{A.3})$$

where $\text{SO}(3)$ is the group of all rotations in a three-dimensional space.

If $\delta_{21}(\mathbf{r})$ is both homogeneous and isotropic it can be shown that its two point correlation function depends on $r_{12} = |\mathbf{r}_1 - \mathbf{r}_2|$ only, that is, $\langle \delta_{21}(\mathbf{r}_1) \delta_{21}(\mathbf{r}_2) \rangle = \xi_{21}(r_{12})$. In this case $P_{21}(\mathbf{k})$ is a function of $k = |\mathbf{k}|$ only since

$$\begin{aligned} P_{21}(\mathbf{k}) &= \int_{-\infty}^{\infty} d^3\mathbf{r} \xi_{21}(r) e^{-i\mathbf{k} \cdot \mathbf{r}} = \int_0^{\infty} r^2 dr \xi_{21}(r) \int_{4\pi} d\Omega e^{-i\mathbf{k} \cdot \mathbf{r}} \\ &= 4\pi \int_0^{\infty} r^2 dr \xi_{21}(r) j_0(kr) = P_{21}(k) \end{aligned} \quad (\text{A.4})$$

where in the second step we used the Rayleigh expansion for the $e^{-i\mathbf{k} \cdot \mathbf{r}}$ term [28]. In this case, equation A.2 becomes

$$\left\langle \tilde{\delta}_{21}(\mathbf{k}_1) \tilde{\delta}_{21}^*(\mathbf{k}_2) \right\rangle = (2\pi)^3 P_{21}(k_1) \delta(\mathbf{k}_1 - \mathbf{k}_2). \quad (\text{A.5})$$

Appendix B

The source coherence function

When we study the response of an interferometer to an extended source, we need to address the problem of how radio waves from different regions of the source correlate, and how the incident electromagnetic field relates to the intensity (or brightness) distribution of the source, I_ν .




Figure B.1: When we study the response of an interferometer to an extended source we will assume that the source is spatially incoherent, that is, that radio waves coming from different regions of the source are uncorrelated.

Let $e(\hat{n}, t)d\Omega$ be the single polarization component of the electric field intensity received by an antenna from a source region of solid angle $d\Omega$, in the direction of the unit vector \hat{n} , and at time t (see Figure B.1). Then the total incident electric field from the

source is

$$E(t) = \int_{\Omega_{src}} e(\hat{\mathbf{n}}, t) d\Omega. \quad (\text{B.1})$$

Now consider electric field received by the antenna from two different regions $d\Omega$ and $d\Omega'$ of the source, in the directions of $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}'$. Assuming that these fields are jointly stationary random processes, then the cross-correlation function of the two fields is called the *source coherence function* $\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)$, that is

$$\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau) = \langle e(\hat{\mathbf{n}}, t) e^*(\hat{\mathbf{n}}', t - \tau) \rangle. \quad (\text{B.2})$$

The normalized source coherence function is defined as

$$\gamma_N(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau) = \frac{\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)}{\sqrt{\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}, 0) \gamma(\hat{\mathbf{n}}', \hat{\mathbf{n}}', 0)}}. \quad (\text{B.3})$$

It is clear that $0 \leq |\gamma_N(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)| \leq 1$. We say that the source is *completely coherent* when $|\gamma_N(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)| = 1$ for all $\hat{\mathbf{n}}, \hat{\mathbf{n}}'$ and τ . When $|\gamma_N(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)| = 0$ for all $\hat{\mathbf{n}}, \hat{\mathbf{n}}'$ and τ we say the source is *incoherent*. In all other cases we say that the source is *partially coherent*.

For most astronomical sources, the fields from different regions are generated by independent random microscopic processes occurring there, so radio waves from different regions are essentially uncorrelated. This means that most radio sources are incoherent and signals from different parts of the source can be treated independently. We will assume that radio sources are incoherent, so the coherence function can be written as

$$\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau) = \gamma(\hat{\mathbf{n}}, \tau) \delta(\hat{\mathbf{n}} - \hat{\mathbf{n}}'). \quad (\text{B.4})$$

The correlation function of the total incident electric field $E(t)$ is

$$\begin{aligned} r_{EE}(\tau) &= \langle E(t) E^*(t - \tau) \rangle = \int_{\Omega_{src}} \int_{\Omega_{src}} \langle e(\hat{\mathbf{n}}, t) e(\hat{\mathbf{n}}', t - \tau) \rangle d\Omega d\Omega' \\ &= \int_{\Omega_{src}} \gamma(\hat{\mathbf{n}}, \tau) d\Omega \end{aligned} \quad (\text{B.5})$$

where in the last step we used equations B.2 and B.4. Thus, the average energy (power)

of $E(t)$ is

$$\langle |E(t)|^2 \rangle = r_{EE}(0) = \int_{\Omega_{src}} \gamma(\hat{\mathbf{n}}, 0) d\Omega \quad (\text{B.6})$$

or, introducing the Fourier transform of the source coherence function, $\tilde{\gamma}(\hat{\mathbf{n}}, \nu)$

$$\langle |E(t)|^2 \rangle = \int_{\Omega_{src}} \int_{-\infty}^{\infty} \tilde{\gamma}(\hat{\mathbf{n}}, \nu) d\nu d\Omega \quad (\text{B.7})$$

This means that we can write the incident power from a source region of solid angle $d\Omega$, in the direction $\hat{\mathbf{n}}$, and in the frequency range $d\nu$ as (this is the value for single polarization. For a randomly polarized electric field it is half the total power of the field)

$$\langle |\mathbf{E}(\hat{\mathbf{n}})|^2 \rangle = \tilde{\gamma}(\hat{\mathbf{n}}, \nu) d\nu d\Omega \quad (\text{B.8})$$

Since, up to a scale factor that we will ignore for simplicity (we can always rescale at the end), $\langle |\mathbf{E}(\hat{\mathbf{n}})|^2 \rangle$ is also the magnitude of the average Poynting vector $\langle S_\nu(\hat{\mathbf{n}}) \rangle$ (average power per unit area) coming from the small solid angle $d\Omega$, in the direction $\hat{\mathbf{n}}$, and in the frequency range $d\nu$ (see for example [22] for details), then we can write

$$\langle S_\nu(\hat{\mathbf{n}}) \rangle = \tilde{\gamma}(\hat{\mathbf{n}}, \nu) d\nu d\Omega \quad (\text{B.9})$$

This means that the average power dW measured by a detector (e.g. a telescope) of area dA in the direction of $\hat{\mathbf{n}}$, due to $\langle S_\nu(\hat{\mathbf{n}}) \rangle$ is

$$dW = \tilde{\gamma}(\hat{\mathbf{n}}, \nu) dA d\nu d\Omega \quad (\text{B.10})$$

From equation B.10 and since dA , $d\nu$, and $d\Omega$ are arbitrary, we conclude that $\tilde{\gamma}(\hat{\mathbf{n}}, \nu)$ is just the intensity, or brightness, $I_\nu(\hat{\mathbf{n}})$, of the source, that is

$$I_\nu(\hat{\mathbf{n}}) = \tilde{\gamma}(\hat{\mathbf{n}}, \nu). \quad (\text{B.11})$$

Equation B.11 relates Fourier transform of the source coherence function $\tilde{\gamma}(\hat{\mathbf{n}}, \nu)$ to

the intensity distribution of the source, $I_\nu(\hat{\mathbf{n}})$, which is a quantity that we can measure in radio astronomy (e.g. with a radio telescope).

Let $\tilde{e}(\hat{\mathbf{n}}, \nu)$ be the Fourier transform of $e(\hat{\mathbf{n}}, t)$

$$\tilde{e}(\hat{\mathbf{n}}, \nu) = \int_{-\infty}^{\infty} e(\hat{\mathbf{n}}, t) e^{-i2\pi\nu t} dt. \quad (\text{B.12})$$

Note that since $e(\hat{\mathbf{n}}, t)$ is a stationary random process then its power spectrum is uniquely defined through the equation

$$\langle \tilde{e}(\hat{\mathbf{n}}, \nu) \tilde{e}^*(\hat{\mathbf{n}}', \nu') \rangle = \tilde{\gamma}(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \nu) \delta(\nu - \nu') \quad (\text{B.13})$$

where $\tilde{\gamma}(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \nu)$ is the Fourier transform of the source coherence function $\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)$ from equation B.2. Equation B.13 above can be easily obtained with the same procedure used in equation A.2 from Appendix A.

Since for an incoherent source $\gamma(\hat{\mathbf{n}}, \hat{\mathbf{n}}', \tau)$ satisfies equation B.4 then we have

$$\begin{aligned} \langle \tilde{e}(\hat{\mathbf{n}}, \nu) \tilde{e}^*(\hat{\mathbf{n}}', \nu') \rangle &= \tilde{\gamma}(\hat{\mathbf{n}}, \nu) \delta(\nu - \nu') \delta(\hat{\mathbf{n}} - \hat{\mathbf{n}}') \\ &= I_\nu(\hat{\mathbf{n}}) \delta(\nu - \nu') \delta(\hat{\mathbf{n}} - \hat{\mathbf{n}}') \end{aligned} \quad (\text{B.14})$$

where in the last step we used equation B.11.

Appendix C

Flat-sky approximation

Traditional radio interferometry applications typically involve mapping objects that subtend a small angle on the sky. In these cases it is reasonable to assume that the region of the sky under observation is approximately flat, or equivalently, that we can project the celestial sphere onto a plane that is a tangent to the field center. When the conditions for this flat-sky approximation are met, equation 2.37 that relates the source temperature distribution to the measured visibilities, and that is difficult to solve in general, becomes a two-dimensional Fourier transform that, in principle, can be solved to recover $T(\hat{\mathbf{n}})$.

Even though we do not use the flat-sky approximation to map the sky with CHIME, it is still instructive to study this mapping approach, not only due to its historical importance, but also to understand key interferometry concepts like the uv -plane and the angular resolution of an interferometer. We can also have a better understanding of the limitations of this technique, which motivates the spherical harmonic transit telescope mapping approach that we describe in Section 2.4, and that is better suited for wide-field radio interferometers like CHIME.

Let us introduce a coordinate system where the first and second axes of the coordinate system point in the east and north directions respectively. The third axis completes the right handed coordinate system and points towards the source reference direction $\hat{\mathbf{n}}_0$ (see Figure C.1).

In this coordinate system the baseline vector \mathbf{b} has components (u, v, w) or, normalizing by the observed wavelength

$$\mathbf{b} = \left(\frac{u}{\lambda}, \frac{v}{\lambda}, \frac{w}{\lambda} \right) = (u_\lambda, v_\lambda, w_\lambda). \quad (\text{C.1})$$

For a given region of the source in the direction of $\hat{\mathbf{n}}$, the vector $\boldsymbol{\sigma} = \hat{\mathbf{n}} - \hat{\mathbf{n}}_0$ denotes the offset vector from the reference direction to the direction of this region. Since $\boldsymbol{\sigma}$ is




Figure C.1: Geometry to understand the flat-sky approximation. The first and second axes of the coordinate system point in the east and north directions respectively. The third axis completes the right handed coordinate system. In this system the normalized baseline vector b_λ has coordinates (u, v, w) and the brightness temperature distribution of the source is $T(l, m)$. The reference direction is $\hat{\mathbf{n}}_0$.

the vector difference of two unit vectors, its components in the new coordinate system can be written as

$$\boldsymbol{\sigma} = (l, m, \sqrt{1 - l^2 - m^2} - 1). \quad (\text{C.2})$$

The inner product in the argument of the exponential of equation 2.37 can be written as

$$\mathbf{b}_\lambda \cdot \hat{\mathbf{n}} = \mathbf{b}_\lambda \cdot \hat{\mathbf{n}}_0 + \mathbf{b}_\lambda \cdot \boldsymbol{\sigma}. \quad (\text{C.3})$$

The first term in equation C.3 generates a term $e^{-i2\pi\mathbf{b}_\lambda \cdot \hat{\mathbf{n}}_0}$ that can be taken out of the integral in equation 2.37. For an interferometer that tracks the source as it moves due to the rotation of the earth, this term will generate a fringe pattern in the measured visibility. Since we can always stop this fringing (*fringe stopping*) in r_{12} by multiplying r_{12} by $e^{i2\pi\mathbf{b}_\lambda \cdot \hat{\mathbf{n}}_0}$, then we will focus on the integral term that we can write as

$$\mathcal{V}_{12} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} T(l, m) |A(l, m)|^2 e^{-i2\pi[u_\lambda l + v_\lambda m + w_\lambda(\sqrt{1-l^2-m^2}-1)]} \frac{dldm}{\sqrt{1-l^2-m^2}} \quad (\text{C.4})$$

where the solid angle element $d\Omega$ was written in terms of its tangent plane projection as $dldm/\sqrt{1-l^2-m^2}$.

If we are mapping a very narrow region of the sky, such that $l, m \ll 1$ then we can safely approximate $d\Omega \approx dldm$. However, we must be careful when discarding terms of order l^2 and m^2 in the argument of the exponential. We can write

$$2\pi w_\lambda(\sqrt{1-l^2-m^2}-1) \approx \pi w_\lambda(l^2+m^2). \quad (\text{C.5})$$

When neglecting the phase term in equation C.5, it is not enough to require that it is much smaller than the other phase terms in the argument of the exponential in equation C.4. We can safely neglect this term only if it is absolutely small as a phase. Let θ_F be the field of view of the mapping. That means that the angle between $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}_0$ does not exceed $\theta_F/2$. In that case we have $l^2 + m^2 \leq (\theta_F/2)^2$. Also, since $w_\lambda \leq b_\lambda$, then the phase error introduced by the l^2 and m^2 terms is

$$\pi w_\lambda(l^2 + m^2) \leq \pi \frac{b}{\lambda} \left(\frac{\theta_F}{2} \right)^2. \quad (\text{C.6})$$

If we require that the phase errors cannot exceed, say, 0.1 rad, then from equation

C.6 we must have

$$\theta_F \lesssim 0.36 \sqrt{\frac{\lambda}{b}}. \quad (\text{C.7})$$

If the flat-sky condition in equation C.7 is met then equation C.4 becomes

$$\mathcal{V}_{12}(u_\lambda, v_\lambda) \approx \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} T(l, m) |A(l, m)|^2 e^{-i2\pi(u_\lambda l + v_\lambda m)} dl dm. \quad (\text{C.8})$$

It is clear that under the flat-sky approximation \mathcal{V}_{12} is just the two-dimensional Fourier transform of the temperature distribution $T(l, m)$, multiplied by the beam function $|A(l, m)|^2$, at the *spatial frequency* pair (u_λ, v_λ) . Thus, under this approximation the interferometer is just a "spatial filter" that selects the spatial frequency (u_λ, v_λ) from the temperature distribution $T(l, m)$ weighted by the interferometer beam.

To recover $T(l, m)$ we must perform an inverse Fourier transform. However, that inversion is possible only if we have a continuous function $\mathcal{V}_{12}(u_\lambda, v_\lambda)$ over the uv -plane. For example, the rotation of the earth changes the projection of the interferometer baseline in the uv -plane, which help us to sample the visibilities at different (u_λ, v_λ) points (see [33] and [73] for a detailed discussion). Or we can have a collection of two-element interferometers with different baselines (an interferometric array) to sample the uv -plane at different points simultaneously. Or we can have a combination of these two methods.

In any case, the best we can ever do with our array is to cover the area of the uv -plane contained within a circle of radius b_{\max}/λ , where b_{\max} is the maximum baseline separation of the array. This means that $\mathcal{V}_{12}(u_\lambda, v_\lambda)$ cannot have spatial frequency components higher than b_{\max}/λ , so the inverse Fourier transform cannot restore spatial structures smaller than λ/b_{\max} . It is for this reason that the angular resolution of an interferometric array is given by

$$\Delta\theta_R = \frac{\lambda}{b_{\max}}. \quad (\text{C.9})$$

Note that, in general, satisfying the flat-sky conditions requires interferometers with very narrow field of view. For experiments like CHIME that have large fields of view in order to map large volumes of the sky quickly, the tangent plane approximation is not applicable, so a different approach that takes into account the curvature of the sky must be used. That approach, called the m -mode formalism, is described in Section 2.4.

Finally, note that if the region of the source is very narrow then the flat-sky conditions are also satisfied, even for CHIME. It is for this reason that equation C.9 still applies for the angular resolution of CHIME.

Appendix D

BIS gain estimation

Here we want to solve the problem in equation 5.8

$$R'_{ij} = \gamma'_i \gamma'^{*}_j + N'_{ij} \quad (\text{D.1})$$

where the data matrix \mathbf{R}' is hermitian but not necessarily positive (semi-) definite since it is the result of subtracting (and normalizing) the correlation matrices from the ‘on’ and ‘off’ BIS samples. The noise term N'_{ij} is assumed to be Gaussian with

$$\langle N'_{ij} \rangle = 0, \quad \langle N'_{ij} N'^{*}_{kl} \rangle = \mathcal{N}'_{(ij),(kl)} = \delta_{ik} \delta_{jl} \sigma^2 \quad (\text{D.2})$$

where $\sigma = 1/\sqrt{\Delta\nu\tau}$. The weighted gain vector $\boldsymbol{\gamma}'$ is defined as

$$\gamma'_i = \sqrt{\frac{1}{T_{sys,i}^{on}}} \alpha_i g_i \sigma_{BIS} \quad (\text{D.3})$$

where σ_{BIS} is the RMS of the BIS source, g_i is the time-dependent receiver gain term of feed i , which is the quantity we want to measure, α_i is the static transfer function of the injected signal, and $T_{sys,i}^{on}$ is the system temperature of feed i in the ‘on’ sample of the BIS signal (see Section 5.2.2).

To find the MLE $\boldsymbol{\gamma}'$ solution for equation D.1 we have to minimize the chi-squared given by

$$\begin{aligned} \chi^2 &= \sum_{ijkl} (R'_{ij} - \gamma'_i \gamma'^{*}_j)^* \mathcal{N}_{(ij),(kl)}'^{-1} (R'_{kl} - \gamma'_k \gamma'^{*}_l) \\ &= \frac{1}{\sigma^2} \sum_{ij} |R'_{ij} - \gamma'_i \gamma'^{*}_j|^2. \end{aligned} \quad (\text{D.4})$$

Minimizing χ^2 with respect to γ'_k we have

$$\frac{\partial \chi^2}{\partial \gamma'_k} = -\frac{2}{\sigma^2} \sum_j \gamma_j'^* (R'_{jk} - \gamma_j' \gamma_k'^*) = 0 \quad \rightarrow \quad \sum_j R'_{kj} \gamma_j' = \|\gamma'\|^2 \gamma_k' \quad (\text{D.5})$$

or, in matrix notation

$$\mathbf{R}' \gamma' = \|\gamma'\|^2 \gamma'. \quad (\text{D.6})$$

The solutions to equation D.6 (the critical points of χ^2) are either $\gamma' = \mathbf{0}$ or γ' is an eigenvector of \mathbf{R}' with positive eigenvalue $\|\gamma'\|^2$. This means that if λ is a positive eigenvalue with corresponding unit-norm eigenvector $\hat{\mathbf{v}}$, then $\gamma' = e^{i\theta} \sqrt{\lambda} \hat{\mathbf{v}}$ is a solution to equation D.6 (the arbitrary phase θ is irrelevant since it cancels out in equation D.4). It is trivial to verify that these are local minima of χ^2 .

To find the global minimum it is convenient to write χ^2 in terms of the Frobenius norm¹

$$\begin{aligned} \chi^2 &= \|\mathbf{R}' - \gamma' \gamma'^H\|^2 = \text{Tr} \left[\left(\mathbf{R}' - \gamma' \gamma'^H \right) \left(\mathbf{R}' - \gamma' \gamma'^H \right)^H \right] \\ &= \|\mathbf{R}'\|^2 - 2 \text{Tr} \left(\gamma'^H \mathbf{R}' \gamma' \right) + \|\gamma'\|^4. \end{aligned} \quad (\text{D.7})$$

It is easy to verify that, as long as \mathbf{R}' has at least one positive eigenvalue (which applies in this case since the noise term is only a small perturbation that prevents \mathbf{R}' from being rank one), then the global minimum of χ^2 is given by

$$\chi_{min}^2 = \|\mathbf{R}'\|^2 - \lambda_1^2 \quad (\text{D.8})$$

where λ_1 is the largest (positive) eigenvalue of \mathbf{R}' . Thus, the MLE solution to equation D.1 is

$$\tilde{\gamma}' = \sqrt{\lambda_1} \hat{\mathbf{v}}_1 \quad (\text{D.9})$$

where $\hat{\mathbf{v}}_1$ is the (unit-norm) eigenvector corresponding to λ_1 .

¹The Frobenius norm of a matrix is the square root of the sum of the squares of each element in the matrix. See <http://mathworld.wolfram.com/FrobeniusNorm.html>

Appendix E

Details of expected BIS performance

E.1 Eigenvalue statistics

We can write equation 5.8 in matrix form as

$$\mathbf{R}' = \boldsymbol{\gamma}'\boldsymbol{\gamma}'^H + \mathbf{N}' \quad (\text{E.1})$$

where \mathbf{N}' is the noise perturbation matrix which is a specific realization of the noise with uniform diagonal covariance given by equation D.2.

From Appendix D, we know that this equation can be solved for the weighted gain vector $\boldsymbol{\gamma}'$ (which leads us to the receiver gain vector \mathbf{g}) from the eigendecomposition of the matrix \mathbf{R}' . In this appendix we will investigate the statistics of the gain solution that will allow us to quantify the performance of the BIS method and determine the expected error in the gain solutions in terms of known or measurable telescope and BIS parameters like bandwidth, integration time, and BIS illumination of the CHIME feeds. This is an elegant derivation done by Richard Shaw that I include in this thesis with his permission.

We will assume that \mathbf{N}' is a small perturbation that prevents \mathbf{R}' from being rank one. In this case we can use (degenerate) eigenvalue perturbation theory to estimate the effect of \mathbf{N}' on the accuracy with which we can recover $\boldsymbol{\gamma}'$. See [20, 74] for a detailed introduction to eigenvalue perturbation theory.

In the unperturbed case, i.e., in the absence of the residual noise \mathbf{N}' , the only nonzero eigenvalue of $\boldsymbol{\gamma}'\boldsymbol{\gamma}'^H$ is $\lambda_1^0 = \|\boldsymbol{\gamma}'\|^2$ and $\hat{\mathbf{v}}_1^0 = \hat{\boldsymbol{\gamma}}'$ is the corresponding (normalized) eigenvector. The corrected nonzero eigenvalue λ_1 is, to first order,

$$\lambda_1 = \lambda_1^0 + \hat{\mathbf{v}}_1^{0H} \mathbf{N}' \hat{\mathbf{v}}_1^0 = \|\boldsymbol{\gamma}'\|^2 + \hat{\boldsymbol{\gamma}}'^H \mathbf{N}' \hat{\boldsymbol{\gamma}}'. \quad (\text{E.2})$$

Since the other eigenvalues are degenerate in the unperturbed case, they are found as the solutions to the eigenvalue problem of the projected noise perturbation

$$[\mathbf{P}\mathbf{N}'\mathbf{P}^H] \mathbf{u}_i = \lambda_i \mathbf{u}_i, \quad i > 1 \quad (\text{E.3})$$

where the matrix \mathbf{P} projects onto the null space of the unperturbed matrix or, equivalently, projects onto the subspace orthogonal to $\hat{\boldsymbol{\gamma}}'$.

Now, the corrected eigenvector for λ_1 is, to first order,

$$\begin{aligned} \hat{\mathbf{v}}_1 &= \hat{\mathbf{v}}_1^0 + \sum_{i=2}^N \hat{\mathbf{v}}_i^0 \left(\frac{\hat{\mathbf{v}}_i^{0H} \mathbf{N}' \hat{\mathbf{v}}_1^0}{\lambda_1^0 - \lambda_i^0} \right) = \hat{\boldsymbol{\gamma}}' + \frac{1}{\|\boldsymbol{\gamma}'\|^2} \sum_{i=2}^N \hat{\mathbf{v}}_i^0 (\hat{\mathbf{v}}_i^{0H} \mathbf{N}' \hat{\boldsymbol{\gamma}}') \\ &= \hat{\boldsymbol{\gamma}}' + \frac{1}{\|\boldsymbol{\gamma}'\|^2} \left[\sum_{i=1}^N \hat{\mathbf{v}}_i^0 (\hat{\mathbf{v}}_i^{0H} \mathbf{N}' \hat{\boldsymbol{\gamma}}') - \hat{\boldsymbol{\gamma}}' (\hat{\boldsymbol{\gamma}}'^H \mathbf{N}' \hat{\boldsymbol{\gamma}}') \right] \\ &= \hat{\boldsymbol{\gamma}}' + \frac{1}{\|\boldsymbol{\gamma}'\|^2} (\mathbf{I} - \hat{\boldsymbol{\gamma}}' \hat{\boldsymbol{\gamma}}'^H) \mathbf{N}' \hat{\boldsymbol{\gamma}}' \end{aligned} \quad (\text{E.4})$$

where in the last step we used the fact that $\{\hat{\mathbf{v}}_i^0\}_{i=1,2,\dots,N}$ is an orthonormal basis of \mathbb{C}^N .

Now we can calculate the statistics of $\hat{\mathbf{v}}_1$ and the λ_i 's. From equation E.2 and using equation D.2 we find that the mean of the largest eigenvalue is

$$\langle \lambda_1 \rangle = \|\boldsymbol{\gamma}'\|^2 \quad (\text{E.5})$$

and its variance is

$$\begin{aligned} \langle (\lambda_1 - \langle \lambda_1 \rangle)^2 \rangle &= \langle \hat{\boldsymbol{\gamma}}'^H \mathbf{N}' \hat{\boldsymbol{\gamma}}' \hat{\boldsymbol{\gamma}}' \hat{\boldsymbol{\gamma}}'^H \mathbf{N}' \hat{\boldsymbol{\gamma}}' \rangle = \sum_{ijkl} \hat{\gamma}'_i^* \hat{\gamma}'_j \hat{\gamma}'_k^* \hat{\gamma}'_l \langle N'_{ij} N'_{kl} \rangle \\ &= \sigma^2 \sum_{ijkl} \hat{\gamma}'_i^* \hat{\gamma}'_j \hat{\gamma}'_k^* \hat{\gamma}'_l \delta_{il} \delta_{jk} \\ &= \sigma^2 \end{aligned} \quad (\text{E.6})$$

For the remaining (residual) eigenvalues we calculate the expectation of the eigenvalue sample statistics. Using matrix trace properties we find that the expectation of the sample

mean is

$$\begin{aligned} \left\langle \frac{1}{N-1} \sum_{i=2}^N \lambda_i \right\rangle &= \frac{1}{N-1} \langle \text{Tr} [\mathbf{P} \mathbf{N}' \mathbf{P}^H] \rangle = \frac{1}{N-1} \sum_{ijk} P_{ij} P_{ik} \langle N_{jk} \rangle \\ &= 0 \end{aligned} \quad (\text{E.7})$$

while for the sample variance we have

$$\begin{aligned} \langle \hat{\sigma}_r^2 \rangle &= \left\langle \frac{1}{N-1} \sum_{i=2}^N \lambda_i^2 \right\rangle \\ &= \frac{1}{N-1} \langle \text{Tr} [\mathbf{P} \mathbf{N}' \mathbf{P}^H \mathbf{P} \mathbf{N}' \mathbf{P}^H] \rangle \\ &= \frac{1}{N-1} \langle \text{Tr} [\mathbf{P}^H \mathbf{P} \mathbf{N}' \mathbf{P}^H \mathbf{P} \mathbf{N}'] \rangle. \end{aligned} \quad (\text{E.8})$$

Nothing that $\mathbf{P}^H \mathbf{P} = \mathbf{I} - \hat{\gamma}' \hat{\gamma}'^H$, we can simplify the equation above to obtain

$$\begin{aligned} \langle \hat{\sigma}_r^2 \rangle &= \frac{1}{N-1} \sum_{ijkl} (\delta_{ij} - \hat{\gamma}'_i \hat{\gamma}'_j^*) (\delta_{kl} - \hat{\gamma}'_k \hat{\gamma}'_l^*) \langle N'_{jk} N'_{li} \rangle \\ &= \frac{1}{N-1} \sum_{ijkl} (\delta_{ij} - \hat{\gamma}'_i \hat{\gamma}'_j^*) (\delta_{kl} - \hat{\gamma}'_k \hat{\gamma}'_l^*) \delta_{ij} \delta_{kl} \\ &= (N-1) \sigma^2 \end{aligned} \quad (\text{E.9})$$

Now we want to know the statistics of the gain vector $\tilde{\gamma}'$ inferred from the solution. Plugging equations E.2 and E.4 in equation D.9 we find

$$\begin{aligned} \tilde{\gamma}' &= \lambda_1^{1/2} \hat{\mathbf{v}}_1 \\ &= \left(\|\gamma'\|^2 + \hat{\gamma}'^H \mathbf{N}' \hat{\gamma}' \right)^{1/2} \left[\hat{\gamma}' + \frac{1}{\|\gamma'\|^2} \left(\mathbf{I} - \hat{\gamma}' \hat{\gamma}'^H \right) \mathbf{N}' \hat{\gamma}' \right] \\ &\approx \gamma' + \frac{1}{\|\gamma'\|} \left(\mathbf{I} - \frac{1}{2} \hat{\gamma}' \hat{\gamma}'^H \right) \mathbf{N}' \hat{\gamma}' \end{aligned} \quad (\text{E.10})$$

to first order in \mathbf{N}' . From the last equation it is clear that

$$\langle \tilde{\gamma}' \rangle = \gamma' \quad (\text{E.11})$$

so the estimator $\tilde{\gamma}'$ is unbiased. Its covariance is

$$\left\langle (\tilde{\gamma}' - \langle \tilde{\gamma}' \rangle) (\tilde{\gamma}' - \langle \tilde{\gamma}' \rangle)^H \right\rangle = \frac{1}{\|\gamma'\|^2} \left(\mathbf{I} - \frac{1}{2} \hat{\gamma}' \hat{\gamma}'^H \right) \left\langle \mathbf{N}' \hat{\gamma}' \hat{\gamma}'^H \mathbf{N}' \right\rangle \left(\mathbf{I} - \frac{1}{2} \hat{\gamma}' \hat{\gamma}'^H \right). \quad (\text{E.12})$$

Noting that the quantity

$$\begin{aligned} \left\langle \mathbf{N}' \hat{\gamma}' \hat{\gamma}'^H \mathbf{N}' \right\rangle_{jk} &= \sum_{lm} \hat{\gamma}'_l \hat{\gamma}'_m^* \left\langle N'_{jl} N'_{mk} \right\rangle = \sigma^2 \sum_{lm} \hat{\gamma}'_l \hat{\gamma}'_m^* \delta_{jk} \delta_{lm} \\ &= \sigma^2 \delta_{jk} \end{aligned} \quad (\text{E.13})$$

we have that $\left\langle \mathbf{N}' \hat{\gamma}' \hat{\gamma}'^H \mathbf{N}' \right\rangle = \sigma^2 \mathbf{I}$, so the covariance can be simplified to obtain

$$\left\langle (\tilde{\gamma}' - \langle \tilde{\gamma}' \rangle) (\tilde{\gamma}' - \langle \tilde{\gamma}' \rangle)^H \right\rangle = \frac{\sigma^2}{\|\gamma'\|^2} \left(\mathbf{I} - \frac{3}{4} \hat{\gamma}' \hat{\gamma}'^H \right). \quad (\text{E.14})$$

E.2 Metrics for BIS performance

As a metric of the expected BIS performance we define the *BIS signal-to-noise ratio*, SNR_{BIS} , in terms of the ratio between the largest eigenvalue and the standard deviation of the residual eigenvalues, which is effectively the ratio of the injected signal to the instrumental noise term

$$SNR_{BIS} = \sqrt{N-1} \frac{\langle \lambda_1 \rangle}{\langle \hat{\sigma}_r^2 \rangle^{1/2}} = \frac{\|\gamma'\|^2}{\sigma}. \quad (\text{E.15})$$

From equations D.3 and E.5 we have

$$\langle \lambda_1 \rangle = \|\gamma'\|^2 = \sum_{i=1}^N \frac{|\alpha_i g_i \sigma_{BIS}|^2}{T_{sys,i}^{on}}. \quad (\text{E.16})$$

Note that the term $T_{BIS,i} = |\alpha_i g_i \sigma_{BIS}|^2$ is the temperature of the BIS signal seen feed i , and is effectively $T_{BIS,i} = T_{sys,i}^{on} - T_{sys,i}^{off}$. By defining the *illumination* of feed i , f_i , as the fraction of the total signal seen by feed i that comes from the BIS signal

$$f_i = \frac{T_{BIS,i}}{T_{sys,i}^{on}} = \frac{T_{sys,i}^{on} - T_{sys,i}^{off}}{T_{sys,i}^{on}} \quad (\text{E.17})$$

we can write $\langle \lambda_1 \rangle$ in terms of the illumination of the CHIME feeds as

$$\langle \lambda_1 \rangle = \|\boldsymbol{\gamma}'\|^2 = \sum_{i=1}^N f_i. \quad (\text{E.18})$$

Note that the illumination is a measurable quantity that can be determined from the autocorrelations in the ‘on’ and ‘off’ samples. Similarly, we can write SNR_{BIS} in terms of telescope and BIS parameters as

$$SNR_{BIS} = \sqrt{\Delta\nu\tau} \sum_{i=1}^N f_i. \quad (\text{E.19})$$

For the gain covariance in equation E.14, note that it contains a correlated term proportional to $\hat{\boldsymbol{\gamma}}'\hat{\boldsymbol{\gamma}}'^H$. However, for $N \gg 1$ and assuming a relatively uniform gain vector (so the gain magnitudes are similar for all the receivers), this correlated term is small so we can write

$$\left\langle (\tilde{\boldsymbol{\gamma}}' - \langle \tilde{\boldsymbol{\gamma}}' \rangle) (\tilde{\boldsymbol{\gamma}}' - \langle \tilde{\boldsymbol{\gamma}}' \rangle)^H \right\rangle \approx \frac{\sigma}{SNR_{BIS}} \mathbf{I}. \quad (\text{E.20})$$

This means that the expected error on the solution for the weighted gain γ'_i is

$$\sigma_{\gamma'_i} = \sqrt{\frac{\sigma}{SNR_{BIS}}} = \left(\Delta\nu\tau \sum_{n=1}^N f_n \right)^{-1/2} \quad (\text{E.21})$$

As explained in Section 5.1, we are interested in the fractional receiver gain errors, $\delta g_i = \sigma_{g_i}/|g_i|$, since those determine how well we can recover the 21 cm power spectrum. Using equations D.3, E.17, and E.21 we obtain

$$\begin{aligned} \delta g_i &= \frac{\sigma_{g_i}}{|g_i|} = \frac{\sigma_{\gamma'_i}}{|g_i|} \sqrt{\frac{T_{sys,i}^{on}}{|\alpha_i|^2 \sigma_{BIS}^2}} = \sqrt{\frac{\sigma}{SNR_{BIS} \cdot f_i}} \\ &= \left(\Delta\nu\tau \cdot f_i \sum_{n=1}^N f_n \right)^{-1/2}. \end{aligned} \quad (\text{E.22})$$

For a given bandwidth, integration time, and illumination profile for the CHIME feeds, equation E.22 gives the expected fractional error in the receiver gain solutions. See Section 5.2.2 for specific examples for CHIME and the pathfinder.

Appendix F

$\langle ve \rangle$ for a real quantizer

Here we derive equation 6.6

$$\langle ve \rangle = \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v) v \mathcal{N}(v|\sigma^2) dv. \quad (\text{F.1})$$

Evaluating the integral and re-arranging

$$\begin{aligned} \langle ve \rangle = \sigma^2 \sum_{i=1}^N \left\{ (y_i - k_i) \mathcal{N}(y_i|\sigma^2) - (y_{i-1} - k_i) \mathcal{N}(y_{i-1}|\sigma^2) \right. \\ \left. - \frac{1}{2} \left[\text{erf} \left(\frac{y_i}{\sqrt{2}\sigma} \right) - \text{erf} \left(\frac{y_{i-1}}{\sqrt{2}\sigma} \right) \right] \right\}. \end{aligned} \quad (\text{F.2})$$

The summation of the erf terms in the square brackets gives 2. As for the first two terms of equation F.2, note that $y_i - k_i = 1/2$ and $y_{i-1} - k_i = -1/2$. Simplifying we obtain

$$\langle ve \rangle = \sigma^2 \left[-1 + \sum_{i=1}^{N-1} \mathcal{N} \left(-\frac{N}{2} + i \middle| \sigma^2 \right) \right]. \quad (\text{F.3})$$

Since $\mathcal{N}(v|\sigma^2)$ is an even function we can also write

$$\langle ve \rangle = \begin{cases} \sigma^2 \left[-1 + 2 \sum_{i=0}^{\frac{N-3}{2}} \mathcal{N}\left(\frac{1}{2} + i \middle| \sigma^2\right) \right] & \text{if } N \text{ odd} \\ \sigma^2 \left[-1 + \frac{1}{\sqrt{2\pi\sigma^2}} + 2 \sum_{i=0}^{\frac{N-4}{2}} \mathcal{N}(1 + i \middle| \sigma^2) \right] & \text{if } N \text{ even} \end{cases} \quad (\text{F.4})$$

where it is clear that the summation term is zero for $N = 2$.

To find $\sigma_e^2 = \langle e^2 \rangle$ in equation 6.7 we follow the same procedure

$$\sigma_e^2 = \sum_{i=1}^N \int_{y_{i-1}}^{y_i} (k_i - v)^2 \mathcal{N}(v|\sigma^2) dv = -\langle ve \rangle + \sum_{i=1}^N \int_{y_{i-1}}^{y_i} k_i (k_i - v) \mathcal{N}(v|\sigma^2) dv. \quad (\text{F.5})$$

Evaluating the integral

$$\sigma_e^2 = -\langle ve \rangle + \sum_{i=1}^N \left\{ \frac{k_i^2}{2} \left[\text{erf}\left(\frac{y_i}{\sqrt{2\sigma^2}}\right) - \text{erf}\left(\frac{y_{i-1}}{\sqrt{2\sigma^2}}\right) \right] + k_i \sigma^2 [\mathcal{N}(y_i|\sigma^2) - \mathcal{N}(y_{i-1}|\sigma^2)] \right\}. \quad (\text{F.6})$$

The summation of the second term inside the curly brackets gives $-\sigma^2 \sum_{i=1}^{N-1} \mathcal{N}(y_i|\sigma^2) = -\langle ve \rangle - \sigma^2$. As for the erf terms, after re-arranging we obtain

$$\sigma_e^2 = -2\langle ve \rangle - \sigma^2 + \left(\frac{N-1}{2}\right)^2 - \sum_{i=1}^{N-1} \frac{1}{2} (k_{i+1}^2 - k_i^2) \text{erf}\left(\frac{y_i}{\sqrt{2\sigma^2}}\right). \quad (\text{F.7})$$

Finally, using equation 6.1 we obtain

$$\sigma_e^2 = -2\langle ve \rangle - \sigma^2 + \left(\frac{N-1}{2}\right)^2 - \sum_{i=1}^{N-1} \left(-\frac{N}{2} + i\right) \text{erf}\left(\frac{-N/2 + i}{\sqrt{2\sigma^2}}\right). \quad (\text{F.8})$$

Equation 6.8 in Section 6.2 follows from the two results above.

Appendix G

Sign of $\langle ve \rangle$

To show that $\langle ve \rangle$ in equation 6.6 is always negative for a quantizer with an odd number of levels (N odd), it is enough to show that $S_o = \sum_{i=0}^M \mathcal{N}(1/2 + i | \sigma^2) < 1/2$ for all M positive integer and $\sigma > 0$ real (it is clear that $S_o > 0$). Note that

$$S_o < \sum_{i=0}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(1/2+i)^2/(2\sigma^2)} = \frac{\vartheta_2(q)}{2\sqrt{2\pi\sigma^2}} \quad (\text{G.1})$$

where ϑ_2 is the Jacobi theta function¹ and

$$q = e^{-1/(2\sigma^2)} = e^{-\pi K'/K} \quad (\text{G.2})$$

where $K(k)$ is the complete elliptic integral of the first kind², k is the elliptic modulus, and $K'(k) = K(\sqrt{1-k^2})$. The functions $\vartheta_2(q)$ and $K(k)$ are related through $\vartheta_2^2(q) = 2kK(k)/\pi$. Also, from equation G.2 we have $\sigma^2 = K/(2\pi K')$. Using these results in equation G.1, we find that

$$S_o < \frac{1}{2} \sqrt{\frac{2kK'}{\pi}}. \quad (\text{G.3})$$

The function $f(k) = 2kK'/\pi$ is a strictly increasing function of k and maps the k -interval $(0, 1)$ (corresponding to $\sigma \in (0, \infty)$) to the interval $(0, 1)$. Thus $f(k) < 1$ in this interval and it follows that $S_o < 1/2$.

¹For details see <http://mathworld.wolfram.com/JacobiThetaFunctions.html>

²For details see <http://mathworld.wolfram.com/CompleteEllipticIntegraloftheFirstKind.html>

For the quantizer with an even number of levels (N even), note that the sum $S_e = \sum_{i=0}^M \mathcal{N}(1 + i|\sigma^2)$ is just the right Riemann sum of $\mathcal{N}(v|\sigma^2)$ over the interval $[0, M]$. Since \mathcal{N} is a strictly decreasing function over this interval then it follows that

$$S_e < \int_0^M \frac{1}{\sqrt{2\pi\sigma^2}} e^{-v^2/(2\sigma^2)} dv = \frac{1}{2} \operatorname{erf} \left(\frac{M}{\sqrt{2\sigma^2}} \right) < \frac{1}{2}. \quad (\text{G.4})$$

Thus, the summation (last) term of equation 6.6 for N even is also positive and bounded above by $1/2$. Since the term $1/(\sqrt{2\pi\sigma^2})$ of this equation becomes arbitrarily large as σ decreases, then $\langle ve \rangle$ eventually becomes positive for N even in the low- σ regime.

Bibliography

- [1] J. Mena-Parra, “A radio-frequency-over-fiber link for large-array radio astronomy applications,” Master’s thesis, McGill University, 2013.
- [2] K. Bandura, G. E. Addison, M. Amiri, J. R. Bond, D. Campbell-Wilson, L. Connor, J.-F. Cliche, G. Davis, M. Deng, N. Denman, M. Dobbs, M. Fandino, K. Gibbs, A. Gilbert, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Höfer, P. Klages, T. L. Landecker, K. Masui, J. Mena Parra, L. B. Newburgh, U.-l. Pen, J. B. Peterson, A. Recnik, J. R. Shaw, K. Sigurdson, M. Sitwell, G. Smecher, R. Smegal, K. Vanderlinde, and D. Wiebe, “Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder,” in *Ground-based and Airborne Telescopes V*, vol. 9145 of *Proc. SPIE*, p. 914522, July 2014.
- [3] K. Bandura, A. N. Bender, J. F. Cliche, T. de Haan, M. A. Dobbs, A. J. Gilbert, S. Griffin, G. Hsyu, D. Ittah, J. Mena-Parra, J. Montgomery, T. Pinsonneault-Marotte, S. Siegel, G. Smecher, Q. Y. Tang, K. Vanderlinde, and N. Whitehorn, “ICE: A Scalable, Low-Cost FPGA-Based Telescope Signal Processing and Networking System,” *JAI special issue on Digital Signal Processing in Radio Astronomy*, vol. 5, p. 1641005, Mar. 2016.
- [4] K. Bandura, J. F. Cliche, M. A. Dobbs, A. J. Gilbert, D. Ittah, J. Mena Parra, and G. Smecher, “ICE-Based Custom Full-Mesh Network for the CHIME High Bandwidth Radio Astronomy Correlator,” *Journal of Astronomical Instrumentation*, vol. 5, p. 1641004, Mar. 2016.
- [5] L. B. Newburgh, G. E. Addison, M. Amiri, K. Bandura, J. R. Bond, L. Connor, J.-F. Cliche, G. Davis, M. Deng, N. Denman, M. Dobbs, M. Fandino, H. Fong, K. Gibbs, A. Gilbert, E. Griffin, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Höfer, P. Klages, T. Landecker, K. Masui, J. Mena-Parra, U.-L. Pen, J. Peterson, A. Recnik, J. R. Shaw, K. Sigurdson, M. Sitwell, G. Smecher, R. Smegal, K. Vanderlinde, and D. Wiebe, “Calibrating CHIME: a new radio interferometer to probe dark energy,”

- in *Ground-based and Airborne Telescopes V*, vol. 9145 of *Proc. SPIE*, p. 91454V, July 2014.
- [6] J. Mena-Parra, K. Bandura, M. A. Dobbs, J. R. Shaw, and S. Siegel, “Quantization bias for digital correlators,” *To appear in JAI*, Mar. 2018. arXiv:1803.04296.
- [7] S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and T. S. C. Project, “Measurements of Ω and Λ from 42 High-Redshift Supernovae,” *ApJ*, vol. 517, pp. 565–586, June 1999.
- [8] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry, “Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant,” *AJ*, vol. 116, pp. 1009–1038, Sept. 1998.
- [9] Planck Collaboration, R. Adam, P. A. R. Ade, N. Aghanim, Y. Akrami, M. I. R. Alves, F. Argüeso, M. Arnaud, F. Arroja, M. Ashdown, and et al., “Planck 2015 results. I. Overview of products and scientific results,” *A&A*, vol. 594, p. A1, Sept. 2016.
- [10] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al., “Planck 2015 results. XIII. Cosmological parameters,” *A&A*, vol. 594, p. A13, Sept. 2016.
- [11] A. Albrecht, G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, J. C. Mather, S. Staggs, and N. B. Suntzeff, “Report of the Dark Energy Task Force,” *ArXiv Astrophysics e-prints*, Sept. 2006.
- [12] S. Carroll, *Spacetime and geometry: an introduction to general relativity*. Addison-Wesley, 2003.
- [13] S. Dodelson, *Modern cosmology*. Academic Press, 2003.

- [14] B. Bassett and R. Hlozek, *Baryon acoustic oscillations*, p. 246. 2010.
- [15] D. J. Eisenstein, H.-J. Seo, and M. White, “On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter,” *ApJ*, vol. 664, pp. 660–674, Aug. 2007.
- [16] D. J. Eisenstein, I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol, R. Scranton, H.-J. Seo, M. Tegmark, Z. Zheng, S. F. Anderson, J. Annis, N. Bahcall, J. Brinkmann, S. Burles, F. J. Castander, A. Connolly, I. Csabai, M. Doi, M. Fukugita, J. A. Frieman, K. Glazebrook, J. E. Gunn, J. S. Hendry, G. Hennessy, Z. Ivezić, S. Kent, G. R. Knapp, H. Lin, Y.-S. Loh, R. H. Lupton, B. Margon, T. A. McKay, A. Meiksin, J. A. Munn, A. Pope, M. W. Richmond, D. Schlegel, D. P. Schneider, K. Shimasaku, C. Stoughton, M. A. Strauss, M. SubbaRao, A. S. Szalay, I. Szapudi, D. L. Tucker, B. Yanny, and D. G. York, “Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies,” *ApJ*, vol. 633, pp. 560–574, Nov. 2005.
- [17] M. Sitwell, *Models and probes of the early and dark universe: inflation and 21-cm radiation in cosmology*. PhD thesis, University of British Columbia, 2014.
- [18] D. Eisenstein and M. White, “Theoretical uncertainty in baryon oscillations,” *PhRvD*, vol. 70, p. 103523, Nov. 2004.
- [19] D. Griffiths, *Introduction to quantum mechanics*. Pearson Prentice Hall, second ed., 2004.
- [20] J. Townsend, *A modern approach to quantum mechanics*. Univ Science Books, second ed., 2012.
- [21] S. R. Furlanetto, S. P. Oh, and F. H. Briggs, “Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe,” *PhR*, vol. 433, pp. 181–301, Oct. 2006.
- [22] G. Rybicki and A. Lightman, *Radiative processes in astrophysics*. John Wiley & Sons, 1979.
- [23] M. G. Santos, A. Cooray, and L. Knox, “Multifrequency Analysis of 21 Centimeter Fluctuations from the Era of Reionization,” *ApJ*, vol. 625, pp. 575–587, June 2005.
- [24] M. Zaldarriaga, S. R. Furlanetto, and L. Hernquist, “21 Centimeter Fluctuations from Cosmic Gas at High Redshifts,” *ApJ*, vol. 608, pp. 622–635, June 2004.

- [25] S. Zaroubi, “The Epoch of Reionization,” in *The First Galaxies* (T. Wiklind, B. Mobasher, and V. Bromm, eds.), vol. 396 of *Astrophysics and Space Science Library*, p. 45, 2013.
- [26] G. B. Field, “Excitation of the Hydrogen 21-CM Line,” *Proceedings of the IRE*, vol. 46, pp. 240–250, Jan. 1958.
- [27] J. R. Pritchard and A. Loeb, “21 cm cosmology in the 21st century,” *Reports on Progress in Physics*, vol. 75, p. 086901, Aug. 2012.
- [28] R. Mehrem, “The Plane Wave Expansion, Infinite Integrals and Identities involving Spherical Bessel Functions,” *ArXiv e-prints*, Sept. 2009.
- [29] M. F. Morales and J. S. B. Wyithe, “Reionization and Cosmology with 21-cm Fluctuations,” *ARA&A*, vol. 48, pp. 127–171, Sept. 2010.
- [30] J. S. B. Wyithe and A. Loeb, “The 21-cm power spectrum after reionization,” *MNRAS*, vol. 397, pp. 1926–1934, Aug. 2009.
- [31] A. Papoulis, *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, 3 ed., 1991.
- [32] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing*. Pearson Prentice Hall, 4 ed., 2007.
- [33] A. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*. Springer International Publishing, 3 ed., 2017.
- [34] H. Kobayashi, B. Mark, and W. Turin, *Probability, Random Processes, and Statistical Analysis*. Cambridge University Press, 1 ed., 2012.
- [35] F. D. Neeser and J. L. Massey, “Proper complex random processes with applications to information theory,” *IEEE Transactions on Information Theory*, vol. 39, pp. 1293–1302, Jul 1993.
- [36] L. Isserlis, “On certain probable errors and correlation coefficients of multiple frequency distributions with skew regression,” *Biometrika*, vol. 11, no. 3, pp. 185–190, 1916.
- [37] J. R. Shaw, K. Sigurdson, U.-L. Pen, A. Stebbins, and M. Sitwell, “All-sky Interferometry with Spherical Harmonic Transit Telescopes,” *ApJ*, vol. 781, p. 57, Feb. 2014.

- [38] A. M. Wolfe, E. Gawiser, and J. X. Prochaska, “Damped Ly α Systems,” *ARA&A*, vol. 43, pp. 861–918, Sept. 2005.
- [39] T.-C. Chang, U.-L. Pen, J. B. Peterson, and P. McDonald, “Baryon Acoustic Oscillation Intensity Mapping of Dark Energy,” *Physical Review Letters*, vol. 100, p. 091303, Mar. 2008.
- [40] CHIME Scientific Collaboration, M. Amiri, K. Bandura, P. Berger, J. R. Bond, J.-F. Cliche, L. Connor, M. Deng, N. Denman, M. Dobbs, R. S. Domagalski, M. Fandino, A. J. Gilbert, D. C. Good, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Höfer, G. Hsyu, P. Klages, T. L. Landecker, K. Masui, J. Mena-Parra, L. Newburgh, N. Oppermann, U.-L. Pen, J. B. Peterson, T. Pinsonneault-Marotte, A. Renard, J. R. Shaw, S. R. Siegel, K. Sigurdson, K. M. Smith, E. Storer, I. Tretyakov, K. Vanderlinde, and D. V. Wiebe, “Limits on the ultra-bright fast radio burst population from the chime pathfinder,” *The Astrophysical Journal*, vol. 844, no. 2, p. 161, 2017.
- [41] C. Ng, “Pulsar science with the CHIME telescope,” *ArXiv e-prints*, Nov. 2017. 1711.02104.
- [42] The CHIME/FRB Collaboration, :, M. Amiri, K. Bandura, P. Berger, M. Bhardwaj, M. M. Boyce, P. J. Boyle, C. Brar, M. Burhanpurkar, P. Chawla, J. Chowdhury, J. F. Cliche, M. D. Cranmer, D. Cubranic, M. Deng, N. Denman, M. Dobbs, M. Fandino, E. Fonseca, B. M. Gaensler, U. Giri, A. J. Gilbert, D. C. Good, S. Guliani, M. Halpern, G. Hinshaw, C. Hofer, A. Josephy, V. M. Kaspi, T. L. Landecker, D. Lang, H. Liao, K. W. Masui, J. Mena-Parra, A. Naidu, L. B. Newburgh, C. Ng, C. Patel, U.-L. Pen, T. Pinsonneault-Marotte, Z. Pleunis, M. Rafiei Ravandi, S. M. Ransom, A. Renard, P. Scholz, K. Sigurdson, S. R. Siegel, K. M. Smith, I. H. Stairs, S. P. Tendulkar, K. Vanderlinde, and D. V. Wiebe, “The CHIME Fast Radio Burst Project: System Overview,” *ArXiv e-prints*, Mar. 2018. 1803.11235.
- [43] C. Ng, K. Vanderlinde, A. Paradise, P. Klages, K. Masui, K. Smith, K. Bandura, P. J. Boyle, M. Dobbs, V. Kaspi, A. Renard, J. R. Shaw, I. Stairs, and I. Tretyakov, “CHIME FRB: An application of FFT beamforming for a radio telescope,” *ArXiv e-prints*, Feb. 2017. 1702.04728.
- [44] M. Deng, D. Campbell-Wilson, and for the CHIME Collaboration, “The cloverleaf antenna: A compact wide-bandwidth dual-polarization feed for CHIME,” *ArXiv e-prints*, Aug. 2017. 1708.08521.

- [45] N. Denman, M. Amiri, K. Bandura, J.-F. Cliche, L. Connor, M. Dobbs, M. Fandino, M. Halpern, A. Hincks, G. Hinshaw, C. Höfer, P. Klages, K. Masui, J. Mena Parra, L. Newburgh, A. Recnik, J. R. Shaw, K. Sigurdson, K. Smith, and K. Vanderlinde, “A GPU-based correlator X-engine implemented on the CHIME Pathfinder,” *Application-specific Systems, Architectures and Processors (ASAP), 2015 IEEE 26th International Conference on*, pp. 35–40, 2015.
- [46] A. Recnik, K. Bandura, N. Denman, A. D. Hincks, G. Hinshaw, P. Klages, U.-L. Pen, and K. Vanderlinde, “An efficient real-time data pipeline for the CHIME Pathfinder radio telescope X-engine,” *Application-specific Systems, Architectures and Processors (ASAP), 2015 IEEE 26th International Conference on*, pp. 57–61, 2015.
- [47] P. Klages, K. Bandura, N. Denman, A. Recnik, J. Sievers, and K. Vanderlinde, “GPU kernels for high-speed 4-bit astrophysical data processing,” *Application-specific Systems, Architectures and Processors (ASAP), 2015 IEEE 26th International Conference on*, pp. 164–165, 2015.
- [48] A. N. Bender, J.-F. Cliche, T. de Haan, M. A. Dobbs, A. J. Gilbert, J. Montgomery, N. Rowlands, G. M. Smecher, K. Smith, and A. Wilson, “Digital frequency domain multiplexing readout electronics for the next generation of millimeter telescopes,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, vol. 9153 of *Proc. SPIE*, p. 91531A, July 2014.
- [49] L. B. Newburgh, K. Bandura, M. A. Bucher, T.-C. Chang, H. C. Chiang, J. F. Cliche, R. Davé, M. Dobbs, C. Clarkson, K. M. Ganga, T. Gogo, A. Gumba, N. Gupta, M. Hilton, B. Johnstone, A. Karastergiou, M. Kunz, D. Lokhorst, R. Maartens, S. Macpherson, M. Mdallalose, K. Moodley, L. Ngwenya, J. M. Parra, J. Peterson, O. Recnik, B. Saliwanchik, M. G. Santos, J. L. Sievers, O. Smirnov, P. Stronkhorst, R. Taylor, K. Vanderlinde, G. Van Vuuren, A. Weltman, and A. Witzemann, “HIRAX: a probe of dark energy and radio transients,” in *Ground-based and Airborne Telescopes VI*, vol. 9906 of *Proc. SPIE*, p. 99065X, Aug. 2016.
- [50] R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*. the University of Michigan: Prentice-Hall, 4, illustrated, reprint ed., 1983.
- [51] Cisco white paper, “Cisco visual networking index: Forecast and methodology, 2016-2021,” tech. rep., Cisco, 2017.

- [52] R. P. Escoffier, G. Comoretto, J. C. Webber, A. Baudry, C. M. Broadwell, J. H. Greenberg, R. R. Treacy, P. Cais, B. Quertier, P. Camino, A. Bos, and A. W. Gunst, “The ALMA correlator,” *A&A*, vol. 462, pp. 801–810, Feb. 2007.
- [53] R. A. Perley, C. J. Chandler, B. J. Butler, and J. M. Wrobel, “The Expanded Very Large Array: A New Telescope for New Science,” *ApJL*, vol. 739, p. L1, Sept. 2011.
- [54] J. Kocz, L. J. Greenhill, B. R. Barsdell, D. Price, G. Bernardi, S. Bourke, M. A. Clark, J. Craig, M. Dexter, J. Dowell, T. Eftekhari, S. Ellingson, G. Hallinan, J. Hartman, A. Jameson, D. MacMahon, G. Taylor, F. Schinzel, and D. Werthimer, “Digital Signal Processing Using Stream High Performance Computing: A 512-Input Broadband Correlator for Radio Astronomy,” *Journal of Astronomical Instrumentation*, vol. 4, p. 1550003, Mar. 2015.
- [55] S. M. Ord, B. Crosse, D. Emrich, D. Pallot, R. B. Wayth, M. A. Clark, S. E. Tremblay, W. Arcus, D. Barnes, M. Bell, G. Bernardi, N. D. R. Bhat, J. D. Bowman, F. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. A. Deshpande, L. deSouza, A. Ewell-Wice, L. Feng, R. Goeke, L. J. Greenhill, B. J. Hazelton, D. Herne, J. N. Hewitt, L. Hindson, N. Hurley-Walker, D. Jacobs, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, N. Kudryavtseva, E. Lenc, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, A. Offringa, J. Pathikulangara, B. Pindor, T. Prabu, P. Procopio, R. A. Remillard, J. Riding, A. E. E. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. Udaya Shankar, K. S. Srivani, J. Stevens, R. Subrahmanyam, S. J. Tingay, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe, “The Murchison Widefield Array Correlator,” *PASA*, vol. 32, p. e006, Mar. 2015.
- [56] W. Kester, *Data Conversion Handbook*. Elsevier/Newnes, 2005.
- [57] F. A. Jenet, G. B. Hobbs, K. J. Lee, and R. N. Manchester, “Detecting the Stochastic Gravitational Wave Background Using Pulsar Timing,” *ApJL*, vol. 625, pp. L123–L126, June 2005.
- [58] J. R. Shaw, K. Sigurdson, M. Sitwell, A. Stebbins, and U.-L. Pen, “Coaxing cosmic 21 cm fluctuations from the polarized sky using m -mode analysis,” *PhRvD*, vol. 91, p. 083514, Apr. 2015.
- [59] P. Berger, L. B. Newburgh, M. Amiri, K. Bandura, J.-F. Cliche, L. Connor, M. Deng, N. Denman, M. Dobbs, M. Fandino, A. J. Gilbert, D. Good, M. Halpern, D. Hanna,

- A. D. Hincks, G. Hinshaw, C. Höfer, A. M. Johnson, T. L. Landecker, K. W. Masui, J. Mena Parra, N. Oppermann, U.-L. Pen, J. B. Peterson, A. Recnik, T. Robishaw, J. R. Shaw, S. Siegel, K. Sigurdson, K. Smith, E. Storer, I. Tretyakov, K. Van Gassen, K. Vanderlinde, and D. Wiebe, “Holographic beam mapping of the CHIME pathfinder array,” in *Ground-based and Airborne Telescopes VI*, vol. 9906 of *Proc. SPIE*, p. 99060D, Aug. 2016.
- [60] K. Masui, M. Amiri, L. Connor, M. Deng, M. Fandino, C. Höfer, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, J. M. Parra, L. B. Newburgh, J. R. Shaw, and K. Vanderlinde, “A compression scheme for radio data in high performance computing,” *Astronomy and Computing*, vol. 12, pp. 181–190, Sept. 2015.
- [61] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall, 1 ed., 1993.
- [62] A. Oppenheim, R. Schaffer, and J. Buck, *Discrete-time signal processing*. Prentice Hall, 2 ed., 1999.
- [63] C. J. Weinstein, “Quantization effects in digital filters,” tech. rep., MIT Lexington Lincoln Lab, no. 468, 1969.
- [64] G. Hsyu, “Calibration and testing of the components of the canadian hydrogen intensity mapping experiment,” Master’s thesis, McGill University, 2016.
- [65] A. Thompson, “Quantization efficiency for eight or more sampling levels,” 1998. MMA memo 220, National Radio Astronomy Observatory.
- [66] J. H. V. Vleck and D. Middleton, “The spectrum of clipped noise,” *Proceedings of the IEEE*, vol. 54, pp. 2–19, Jan 1966.
- [67] S. R. Kulkarni and C. Heiles, “How to obtain the true correlation from a 3-level digital correlator,” vol. 85, pp. 1413–1420, 09 1980.
- [68] B. F. C. Cooper, “Correlators with two-bit quantization,” *Australian Journal of Physics*, vol. 23, pp. 521–527, Aug. 1970.
- [69] A. R. Thompson, D. T. Emerson, and F. R. Schwab, “Convenient formulas for quantization efficiency,” *Radio Science*, vol. 42, pp. 1–5, June 2007.
- [70] L. V. Benkevitch, A. E. E. Rogers, C. J. Lonsdale, *et al.*, “Van Vleck correction generalization for complex correlators with multilevel quantization,” *ArXiv e-prints*, Aug. 2016. 1607.02059.

- [71] R. Price, “A useful theorem for nonlinear devices having gaussian inputs,” *IRE Transactions on Information Theory*, vol. 4, pp. 69–72, June 1958.
- [72] M. F. Wagdy, “Effect of various dither forms on quantization errors of ideal a/d converters,” *IEEE Transactions on Instrumentation and Measurement*, vol. 38, pp. 850–855, Aug 1989.
- [73] T. Wilson, K. Rohlfs, and S. Hüttemeister, *Tools of Radio Astronomy*. Springer, 5 ed., 2009.
- [74] J. J. Sakurai, *Modern quantum mechanics*. Pearson, second ed., 2010.