A Digital Spectral Correlator for GMRT

A Thesis

Submitted for the Degree of

Master of Science (Engg.)

in the Faculty of Engineering

By

V. M. Tatke

Centre for Electronics Design and Technology

INDIAN INSTITUTE OF SCIENCE

BANGALORE 560 012 (INDIA)

JANUARY 1998

To my Parents

Who ran to help me when I fell,
And would some pretty story tell,
Or kiss the place to make it well?
My Mother.

–Ann Taylor, 'My Mother'

Honour thy father and thy mother:
that thy days may be long upon the land which
the Lord thy God giveth thee.

–Old Testament, Exodus, XX, 12
(5th. Commandment)

Abstract

A Digital Spectral Correlator for GMRT

The Giant Metrewave Radio Telescope (GMRT), the world's largest radio telescope in the metrewave range, is an array of 30 antennas spread over approximately 25 Sq. Km. and situated about 80 Km. from Pune, India. It is being built as a major national facility for research in Astrophysics.

The GMRT antennas are equipped with highly sensitive superheterodyned radio receivers operating in six frequency bands between 50MHz and 1420 MHz, with a maximum bandwidth of 32MHz. The GMRT receiver electronics outputs four analog signals - two sidebands for each of the two polarizations - of 16MHz bandwidth for each of the 30 antennas. The 120 analog signal streams thus obtained are then input to the *Correlator* - the digital backend for real-time processing of these signals.

The GMRT correlator is required to compute the cross-power spectrum of signals from all 435 ($\frac{30X29}{2}$) antenna pairs as well as their associated auto-power spectra. This required a computing power of about 400 Giga Flops, which necessitated the development of a special purpose hardware, specific to this application. The author was a core member of the team which designed and implemented this complex signal processing system. This thesis describes the basic specifications of the GMRT correlator and the design employed to meet them.

Each of the 120 analog signals is first sampled by the correlator and compensated for path length differences. A real-time Fourier transform is then performed on each of these to compute their spectra. The voltage spectra thus obtained are input to a multiplier array which computes the auto and cross-power spectra for all antenna pairs.

The correlator thus comprises of four major blocks, viz.,

1. The ADC Subsystem : Samples the 120 analog inputs at 32MS/s.

2. The Delay Subsystem : Any value of delay upto 128 μsecs. can be compensated for, in steps of 32ns. These delay values are dynamically adjustable to compensate for effective path length differences produced due to rotation of the earth.

3. The FFT Subsystem : Performs upto 512-point Fourier transform in real-time to compute the spectra of each of the 120 inputs. It thus provides a maximum of 256 spectral channels for each input.

4. The MAC Subsystem : Multiplies the FFT outputs to compute the self and cross-power spectra for all the antenna pairs. It provides a total of 238080 complex accumulator channels.

The discussion presented in the thesis includes the circuit-design and packaging aspects of the various subsystems and the procedures employed to test, validate and integrate them.

Acknowledgements

Designing and building the GMRT Correlator has been, on the whole, an enriching and satisfying experience. I joined the GMRT Correlator team in October 1991 after graduation. Since then, till today, I have had the good fortune to have savoured the bonhomie of the GMRT fraternity in general, and the Correlator team in particular.

First and foremost, I wish to express my heartfelt gratitude for the encouragement and support which Prof. C.R. Subrahmanya has provided me in these past six years. As the head of the correlator group, he maintained an informal, friendly and close relationship with all members of the correlator group, including myself. I also have had the opportunity to exchange notes on topics reaching beyond those imposed by the correlator. I have always found in him a friend, philosopher and guide. I am extremely thankful to him for helping me a great deal with this thesis.

I would like to acknowledge and thank the support provided by Prof. H.S. Jamadagni all through the M.Sc. (Engg.) course. I wish to thank him for his several invaluable suggestions, which he rendered at various times during the course, helping me to see my work done at TIFR in a much broader perspective.

I regard Mr. Abhijit Datta as my mentor as far as the design of the correlator is concerned. I owe him the opportunity I had, to work on all the subsystems of the correlator. We, the members of the correlator group, owe him a well equipped laboratory and a cohesive team which laid the foundation for the present day correlator.

I would like to thank Dr. Yashwant Gupta, for his support and understanding, which has helped me to concentrate on the completion of this course. I am also thankful to him for going through the draft of this thesis and providing his valuable suggestions.

Among my colleauges, I have shared a very close relationship with Mr. Rakesh Malik. And, as a consequence, later I was fortunate to have shared a similar relationship with Mrs. Usha Malik. I owe Rakesh all the understanding of the GMRT and the correlator system, that I have at present. I also owe him a great deal as far as the MAC subsystem is concerned. I am thankful to him to have very cheerfully agreed to work on the MAC subsystem during my coursework at the IISc. And, to both, Usha and Rakesh, I owe a very cosy and homely atmosphere, in which I could completely be at ease.

Mr. Umesh Puranik and Mrs. Alka Bhatnagar, my other colleaugues in the correlator group made the endeavour a pleasure.

I wish to thank Mrs. Madhura Anturkar for having provided encouragement in my moments of sorrow. I am also thankful to her to have gone through my thesis and to have provided her valuable suggestions.

I am grateful to Mr. B. Rajendran and Mr. Ganesh Jangam, who with their skill have crafted our ideas into working systems. Their unflinching support and relentless toil has made it possible for us to see the light at the end of the tunnel.

I wish to thank Mr. Mahesh Burse, who after joining the Correlator group, very diligently took upon himself the task of debugging and integrating the final MAC subsystem, thereby leaving me free to concentrate completely on the completion of this thesis.

I would like to thank Mrs. Sunita Barve for all the help she extended in trying to collect the reference material needed for this thesis.

Among my other colleagues, I wish to thank Mr. Sanjay Bhatnagar, Mr. Divya Oberoi, Mr. Ajit Kulkarni and Mr. Nissim Kanekar to have made the stay in TIFR interesting. I would especially like to thank Mr. Divya Oberoi for having helped me with some of the photographs of the PCBs in this thesis.

In the end, my thoughts go out to my Parents and my Sisters, who have been my source of inspiration and have constantly encouraged me, all through. To them I owe a debt, too enormous to be expressed in words.

Contents

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

Many of the modern buildings in Italy are historically known to have been built out of the pillaged structures of older days. here we may observe a column or a lintel serving the same purpose a second time....I will pursue this rough simile just one step further, which is as much as it will bear. Suppose we were building a house with second-hand materials carted from a dealer's yard, we should often find considerable portions of the same old house to be still grouped together....So in the process of transmission by inheritance, elements derived from the same ancestor are apt to appear in large groups.

– *Sir Francis Galton.*

The question of how the Universe began has fuelled the irrepressible human curiosity since time immemorial. The night sky has always been a challenge for him to comprehend; for Man has always found himself limited by his modest powers of perception, which seem meagre to probe the vast expanse of the Universe which beckons him. Telescopes gave Man, a way to reach out to the distant stars. Today, Radio telescopes extend his reach to realms which he cannot perceive with his eyes. The endeavour to build more, and yet more powerful telescopes is not new. The GMRT is the latest of these. This chapter provides the relevant aspects of Astronomy and radio telescopes, which will form the backdrop for the discussion in the later chapters

1.1 Correlation and Correlators

Ever since Sir Francis Galton (1822-1911), in his endeavour to "quantify" the abstract nature of family traits, introduced the concept of "correlation" to the repertoire of logical techniques at the disposal of mankind, its use has been widespread. Pearson describes Galton's pioneering concept of a correlated system of variates as, the representation "by a single numerical quantity of the degree of relationship, or of partial causality, between the different variables of our ever changing universe".

Correlation owes its near ubiquitous application to the quantity it "measures" - similarity. In all walks of life and reasoning, whenever and wherever, one needs to espouse the cause-effect relationship between events of various kinds, one, often sub-consciously, subscribes to the concept of correlation. One often *observes* certain events occurring either simultaneously or in succession, and *correlates* them to establish some common property between them. In such an event, one says that the two events are *correlated*.

One then naturally expects that science, that branch of knowledge involving systematized observation and experiment, would find this concept indispensable. Of course, Biology interprets correlation as "a mutual relationship or balance existing between different functions, structures, or characteristics in a plant or an animal"; Geology as "the determination of the characteristic or chronological equivalence or relationships between geologic phenomena or events, usually in separated areas" and Statistics as "the strength of the linear dependence between two random variables".

The concept of correlation finds widespread use in various disciplines of engineering. Signal processing applications in various fields use correlation as the basis of their algorithms and techniques. Indeed, correlation finds use in Aviation as CORRELATION DETECTION - a method in which moving aircraft or spacecraft are detected by comparing a signal, point to point, with an internally generated reference; in Military Science as CORRELATION DETECTION and RANGING (CODAR) - a system of submarine detection using sonobuoys dropped from naval aircraft. Variation in time and phase of acoustic signals transmitted by radio from the sonobuoys to the aircraft indicate the location of the submerged submarines; in Material Science as CORRELATION HOLE - an area in a polymer melt where there is a reduced concentration of similar units; in Telecommunications as CORRELATION DISTANCE - an important measure of correlation in which the line of regression that is fitted to the points on a scatter diagram is done in such a way that the sum of the squared deviation from the line is at a minimum. Not only these, but disciplines as diverse as Economy, too, find use for correlation.

A large number of instruments utilise correlation as the crux of their operation; CORRELATION TRACKING SYSTEM - a trajectory measuring system in which various signals obtained from a single source are correlated to determine the phase difference between the signals; CORRELATION ULTRA- SONIC FLOWMETER - a device that measures the rate of fluid flow by recording the amount of time it takes for discontinuities in the fluid to travel between two corresponding transducers that emit and detect high frequency sound.

Today, research in Optoelectronics has produced Optical Correlators that can differentiate between similar and dissimilar faces (Fig.1.1, *Source: [1]*). Global Positioning System (GPS) - originally a concept for the use of the military, revolutionised navigation, and has now found use in innumerable number of civilian applications; places where accurate measurements of either position or time are required, utilise GPS. And, Correlators are at the heart of GPS devices.

The Correlator, as pertaining to electronics, is a logic device that detects weak signals by an operation approximating the computation of the correlation function, in which a reference signal representing the expected signal is constructed, and compared with the receiver input, to distinguish between the noise and the signal components. It is to this variety of the correlator that this thesis pertains.

In order to appreciate the context in which the Correlator system, as described in this thesis is used, a brief introduction to Radio Astronomy is given below.

1.2 Radio Telescopes and Receivers

In this section, an overview of radio astronomy and radio telescopes will be given from the point of view of setting the framework for a later discussion of correlator which is a backend of a radio interferometer. The discussion of radio astronomy is rather simplistic and is just sufficient to give a feel for the quantities which an astronomer wants to measure from a radio telescope. The material presented in this section has been collected from a variety of books and review papers ([2] - [12]).

Radio Astronomy is the study of the Universe through the radio window of the electromagnetic spectrum (Fig. 1.2, *Source: [4]Fig. 1.1*). It is just one of the many branches of astronomy - X-ray, Infrared, Ultraviolet, Optical, Gamma Ray - to name a few; each specialising in utilising a part of the electromagnetic spectrum to perceive the Universe, stars, the space between stars (Inter Stellar Medium), Stellar Systems, their distributions, kinematics and dynamics.

Optical Image Processing

processing holds promise for high performance, compact, and low power computing applications, particularly where the size and power dissipation requirements become critical. Such processors will find place in the airborne and space environment. Photonic Systems Group works with the OPCOM Consortium on some of those issues.

Joint transform correlator with semiconductor multiple quantum well device

Schematic of the JTC

Simple examples of the input image and processed output from the correlator are shown below.

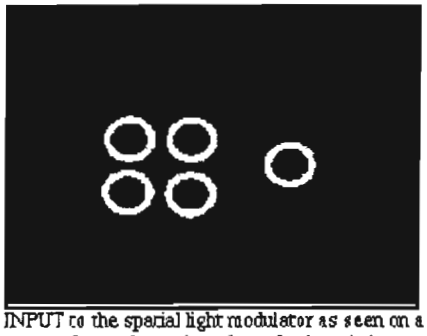

INPUT to the spacial light modulator as seen on a computer monitor. Here five similar circles are input to the correlator. The CCD output shows the region where the four circles on the left in the input correlate with the circle on the right. Cross correlations amongst the four circles show strong peaks in the output although these are close to the DC region in the correlation plane which is to the left of the image shown below.
(CLICK ON PICTURE, need mpeg viewer to see this 10 sec. movie, 54 kbytes).
To download mpeg viewers.

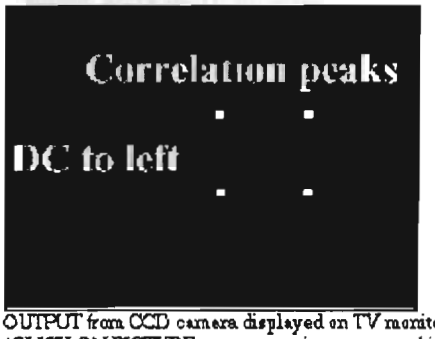

OUTPUT from CCD camera displayed on TV monitor.
(CLICK ON PICTURE, use mpeg viewer to see this 10 sec. movie, 80 kbytes)

Figure 1.1: The Joint Transform Correlator (1 of 5)

INPUT to SLM viewed on a computer screen. This input shows two identical finger prints, the strong correlation peak which results can be seen in the CCD output below. As the movie proceeds the finger print on the right is rotated by 5 degrees in 1 degree increments and then only part of the finger print on the right is input. The change in the correlation peak as the print is rotated and truncated is seen below.
(CLICK ON PICTURE, use mpeg viewer to see this 10 sec. movie, 65 kbytes)

JTC input

OUTPUT from CCD camera.
(CLICK ON PICTURE, use mpeg viewer to see this 10 sec. movie, 6 kbytes)

Face recognition with the correlator

Results of prototype face recognition system based on optical correlator. Image below shows the correlation output from an 8 bit CCD camera when the two similar images of NRC researcher Simon Boothroyd shown at right were input to the system.

Output correlation peak with similar faces as input

Image below shows there is no correlation output when different faces are presented to the system, here on the left Jon Gerrard, Secretary of State for Science, Research and Development, and on the right Brian Garside, president of Optoelectronics Inc.

Output correlation with dissimilar faces as input

!!!! New face recognition movies !!!!

MQW structure and characteristics

Multiple quantum well structure is 2.5 microns thick mounted on a sapphire substrate.
Spectral characterization of the semiconductor MQW element
The multiple quantum well photorefractive device gif - 200 kbytes or jpg - 65 kbytes

More information Dr. Simon Boothroyd

Figure 1.1: The Joint Transform Correlator (2 of 5)

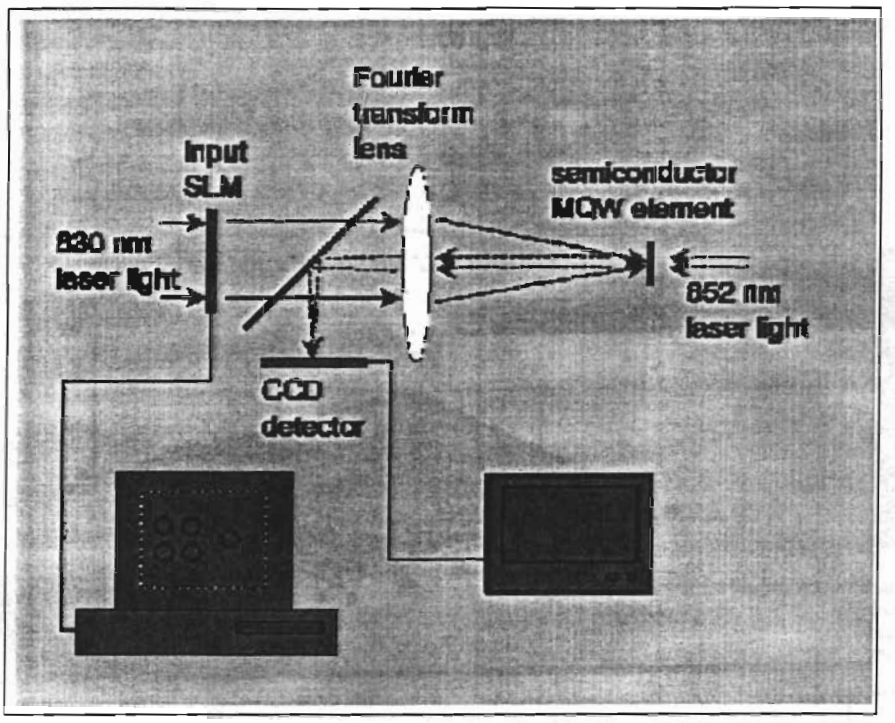

Figure 1.1: Schematic of the Joint Transform Correlator (3 of 5)

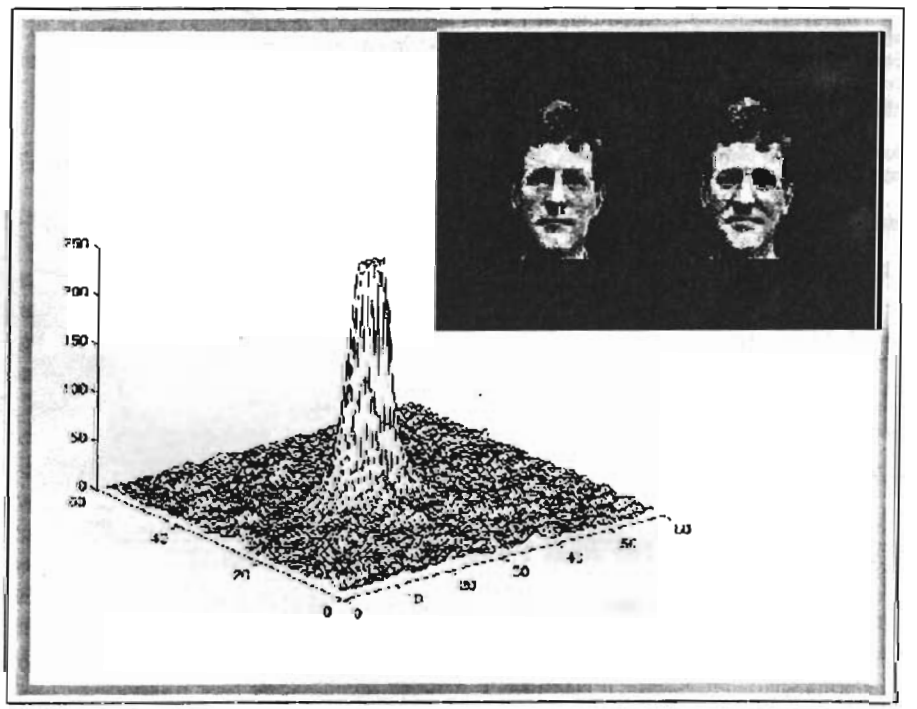

Figure 1.1: Output of the JTC with similar faces as input (4 of 5)

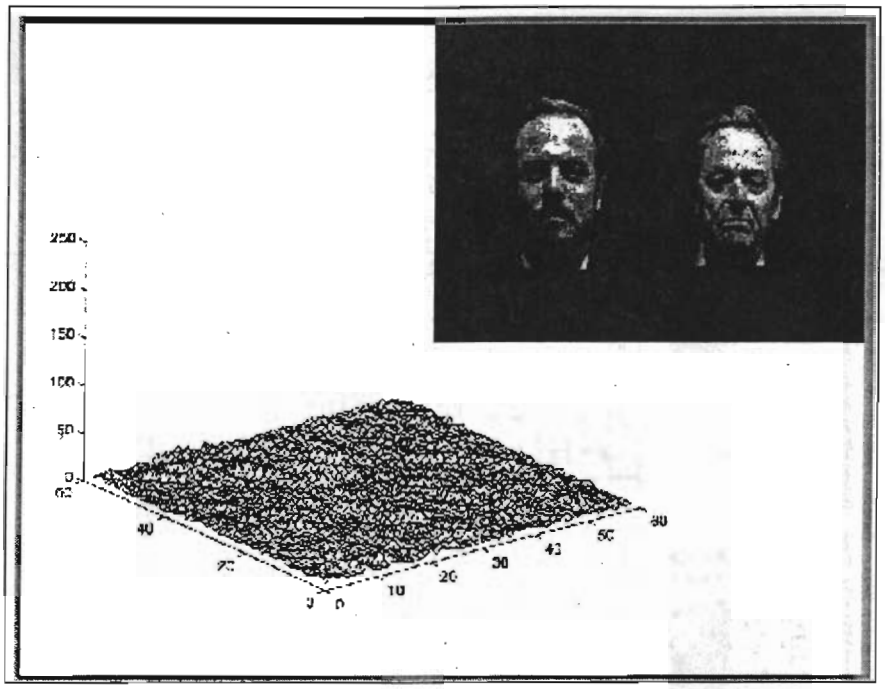

Figure 1.1: Output of the JTC with dissimilar faces as input (5 of 5)

Fig. 1.3 shows the various radio frequencies/bands allocated exclusively for radio astronomy by the World Radiocommunication Conference-1997 (WRC-97).

In the forthcoming sections, a brief introduction to modern Radio telescopes and receivers is presented.

1.2.1 The Radio Telescope

All telescopes - radio, optical, X-ray - couple the radiation from the Universe to the astronomer's measuring devices. The Universe is presented through measurements of intensity and state of polarisation as a function of frequency, angular position and time. The instantaneous values of intensity vary randomly due to the intrinsically stochastic radiation processes, and also because the real world is constantly agitated by quantum fluctuations. Further, the fluctuations arising from extraneous signals added by nature and by man, only add to randomness of the incoming radiation. The measurement process consists of estimating average values for these quantities, and the behaviour of the Universe is then inferred from these estimates. The term "Radio telescope" is used to include the device that measures the received power, the input circuitry, radio receiver and the power measuring device.

A modern Radio Telescope (Fig. 1.4, *Adapted from [4] Fig. 1.6*) consists of one or more antennas which with their feeds collect the radio power from the distant celestial source. After pre-amplification at the feed point, the signal power is conveyed by cables to a main receiver building. Here, it is amplified further, detected and integrated and the output is either displayed or recorded or both. In case of multiple antennas, the power from each may be combined with the others, before processing and subsequent record or display.

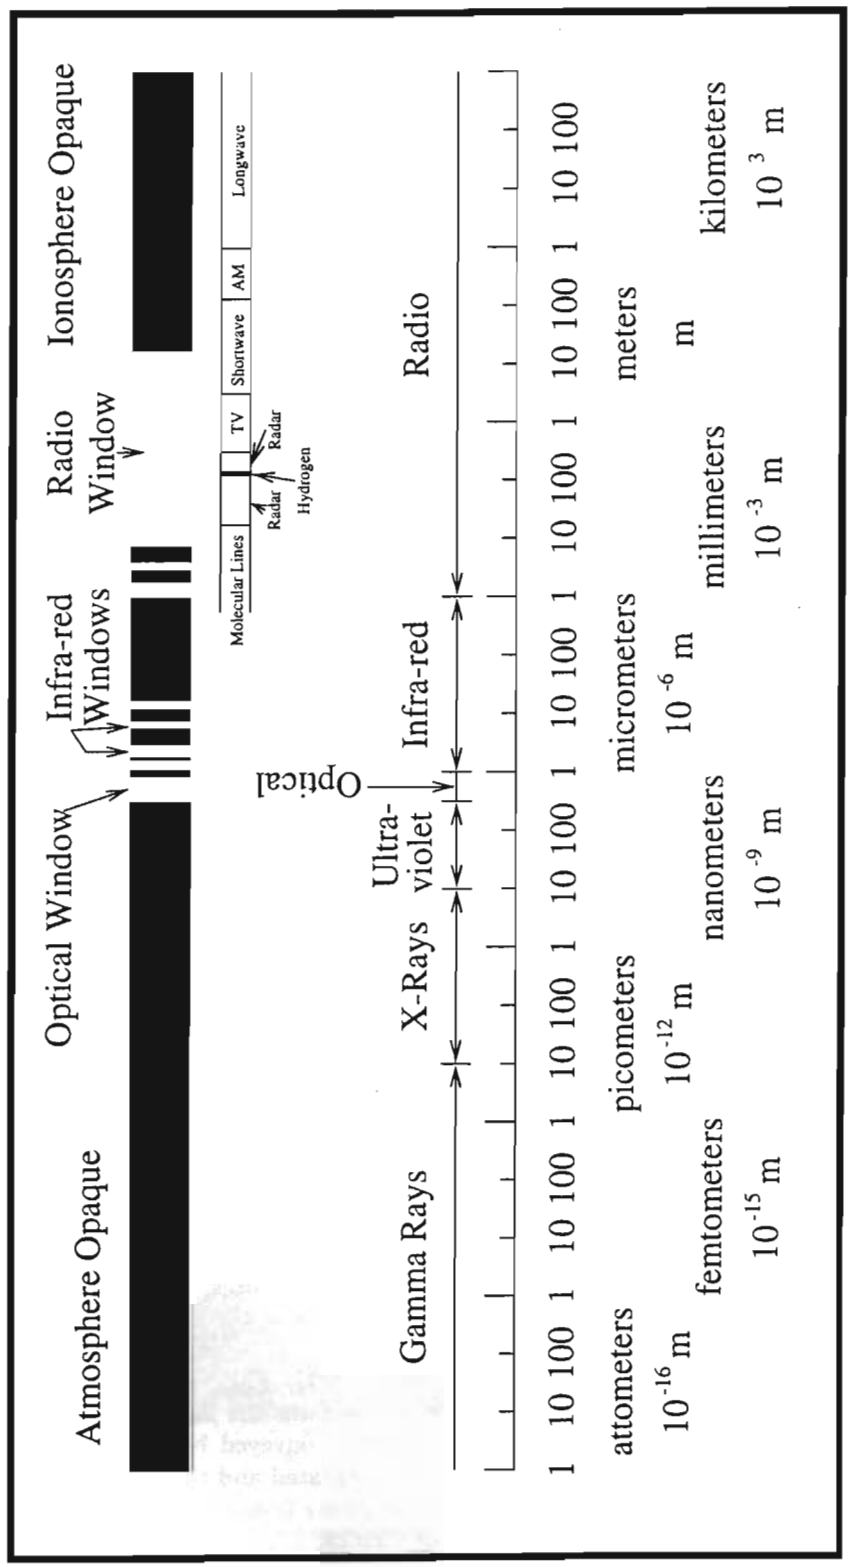

Figure 1.2: The Electromagnetic Spectrum

13,360 - 13,410 kHz, 25,550 - 25,670 kHz, 37.5 - 38.25 MHz, 73 - 74.6 MHz in Regions 1 and 3, 79.75 - 80.25 MHz in Region 3, 150.05 - 153 MHz in Region 1, 322 - 328.6 MHz, 406.1 - 410 MHz, 608 - 614 MHz in Regions 1 and 3, 1330 - 1400 MHz, 1610.6 - 1613.8 MHz, 1660 - 1670 MHz, 1718.8 - 1722.2 MHz, 2655 - 2690 MHz, 3260 - 3267 MHz, 3332 - 3339 MHz, 3345.8 - 3352.5 MHz,	4825 - 4835 MHz, 4950 - 4990 MHz, 4990 - 5000 MHz, 6650 - 6675.2 MHz, 10.6 - 10.68 GHz, 14.47 - 14.5 GHz, 22.01 - 22.21 GHz, 22.21 - 22.5 GHz, 22.81 - 22.86 GHz, 23.07 - 23.12 GHz, 31.2 - 31.3 GHz, 31.5 - 31.8 GHz in Regions 1 and 3, 36.43 - 36.5 GHz, 42.5 - 43.5 GHz, 42.77 - 42.87 GHz, 43.07 - 43.17 GHz, 43.37 - 43.47 GHz, 48.94 - 49.04 GHz, 72.77 - 72.91 GHz, 93.07 - 93.27 GHz,	97.88 - 98.08 GHz, 140.69 - 140.98 GHz, 144.68 - 144.98 GHz, 145.45 - 145.75 GHz, 146.82 - 147.12 GHz, 150 - 151 GHz, 174.42 - 177.4 GHz, 177 - 177.4 GHz, 178.2 - 178.6 GHz, 181 - 181.46 GHz, 186.2 - 186.6 GHz, 250 -251 GHz, 257.5 - 258 GHz, 261 - 265 GHz, 262.24 - 262.76 GHz, 265 - 275 GHz, 265.64 - 266.16 Ghz, 267.34 - 267.86 GHz, 271.74 - 272.26 GHz
Region 1 = Europe and Africa; Region 2 = Western Hemisphere; Region 3 = Asia and Australia.		

Figure 1.3: The Frequencies Allocated for Radio Astronomy

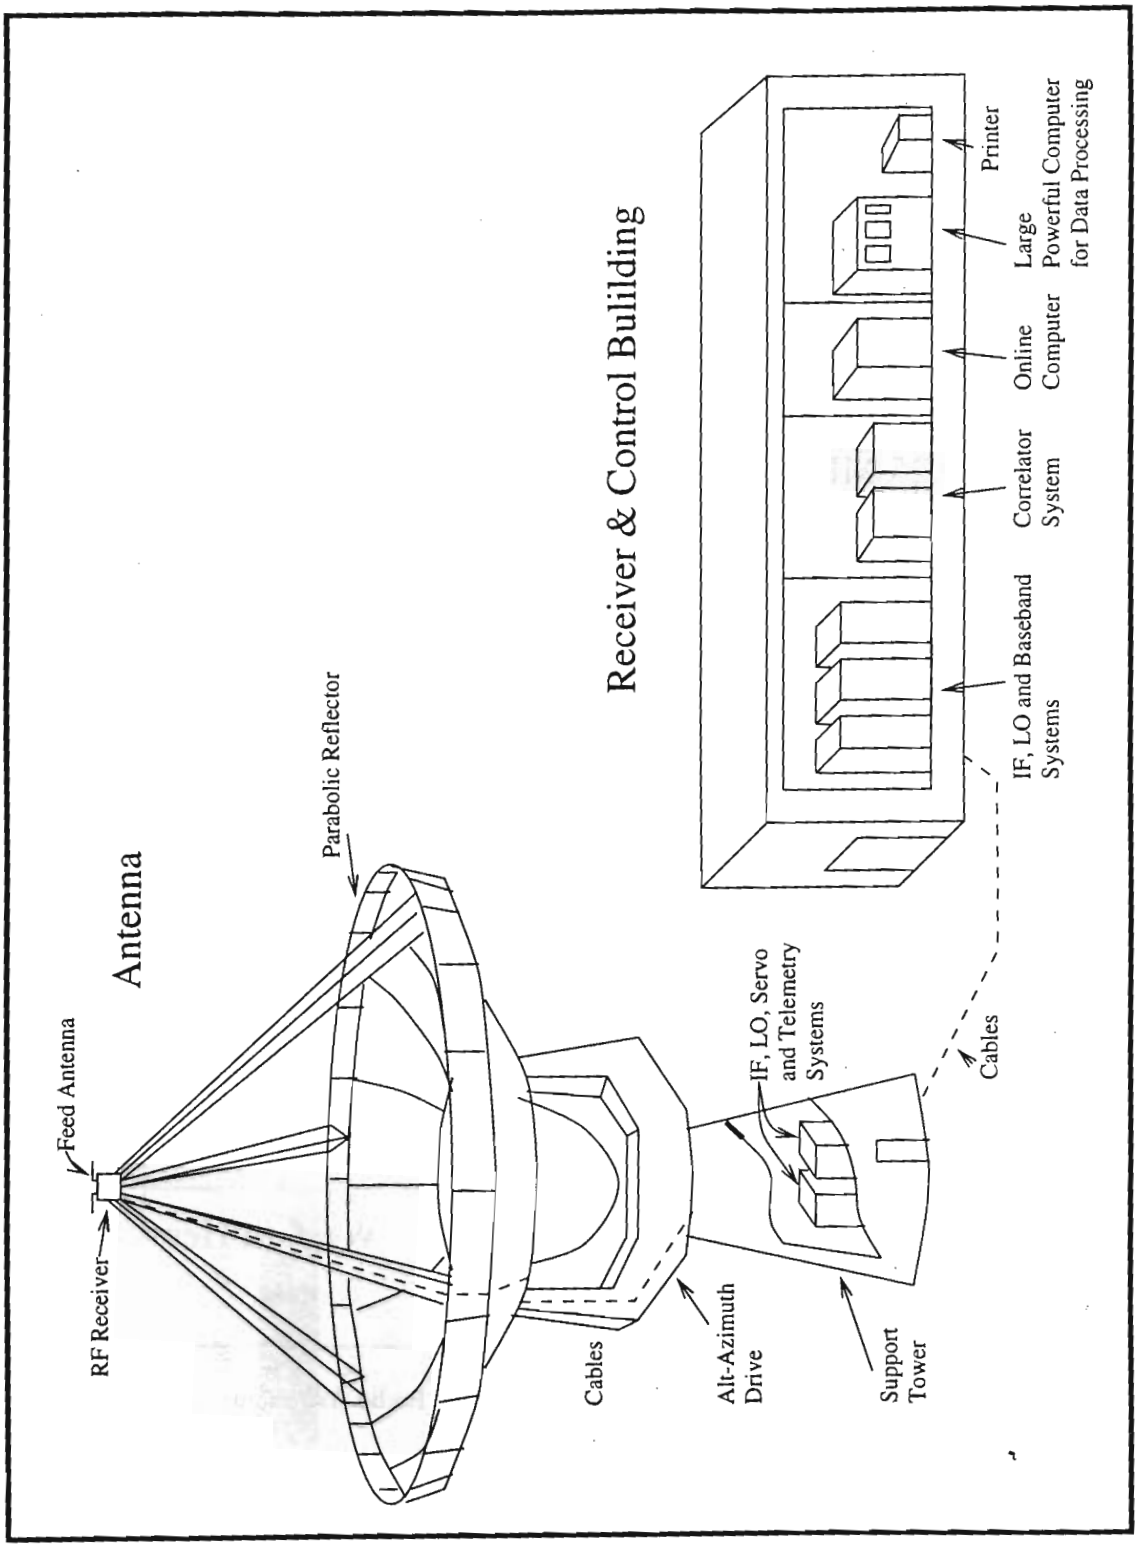

Figure 1.4: A Modern Radio Telescope

1.2.2 How does the Radio Telescope *SEE*?

The observed response of a radio telescope antenna to a sky brightness distribution - the quantity to be measured - is proportional to the convolution of the antenna power pattern and the brightness distribution. Thus,

$$S(\phi_0) = \int_{-\infty}^{\infty} B(\phi)P(\phi_0 - \phi)\, d\phi \tag{1.1}$$

where, $S(\phi_0)$ = observed flux-density distribution
$B(\phi)$ = true source brightness distribution
$P(\phi)$ = normalised antenna power pattern
ϕ_0 = angular displacement

The resolving power of the telescope is limited by diffraction; according to the Rayleigh criterion the angular resolution of a radio telescope is given by

$$\delta \sim \lambda/D \tag{1.2}$$

Hence, for a given wavelength, increasing the diameter D of the radio telescope results in better resolution. As a comparison, if we take the human eye to be sensitive to say 5000Å and its diameter to be approximately 1cm, the resulting angular resolution obtained is

$$\delta = \lambda/D = 5000\text{X}10^{-8}\text{cm}/1\text{cm} = 5\text{X}10^{-5} rads \approx 10 arcsec \tag{1.3}$$

To get the same resolving power at a wavelength of say, 1m, the diameter of an antenna needs to be

$$D = \lambda/\delta = 1\text{m}/(5\text{X}10^{-5}) = 20\text{Km}. \tag{1.4}$$

The largest single dish telescope, as of today, is the Arecibo dish, 305m in diameter. Thus, in order to fulifill the need to obtain higher angular resolutions at practical costs, one employs the technique of interferometry using an array of antenna elements.

In general, it can be shown that the angular resolution of an interferometer is inversely proportional to the separation between the antennas. Hence, the farther the antennas are, narrower is the beam. An interferometer measures a quantity known as visibility $V_0(S_\lambda)$, related to the sky brightness distribution $B(\phi_0)$ by

$$B(\phi_0) = S_0 \int_{-\infty}^{+\infty} V_0(s_\lambda) e^{-j2\pi s_\lambda(\phi_0 - \Delta\phi_0)}\, ds_\lambda \tag{1.5}$$

where, $B(\phi)$ = true source brightness distribution
ϕ_0 = angular displacement
s_λ = $s/\lambda$is the spatial frequency related to the interferometer element spacing s
S_0 = flux density of the source

In practice, an interferometer measures the cross-correlation of the electric fields arriving at a pair of antennas, after suitably compensating for path length differences; these cross-correlations are directly related to the visibility.

A solution of eqn. (1.5) requires observations at suitable intervals out to sufficiently large spacings, a high source signal-to-noise ratio, and no other (confusing) sources of significant flux density in the individual element response pattern. Thus, there are practical limits to the detail with which the source distribution can be determined. The spacing interval need be no smaller

than $1/\alpha$, where α is the full source extent. Increase in detail is obtained by having several antenna elements at various spacings observing the source simultaneously. The output from these elements is then combined to obtain an aperture much larger than that of the individual elements - a technique known as Aperture Synthesis. Another form of aperture synthesis known as the Earth Rotation Aperture Synthesis, makes use of the earth rotation along with the multiplicity of antenna elements to obtain sufficient coverage of spatial frequencies (which define what is known as the u-v plane) in order to obtain a sufficiently detailed image (the brightness distribution) of the source (Fig. 1.5, *Source: [7] p. 210*).

Fig. 1.6 (*Source: [7] p. 211*) is a superposition of the optical and radio images of Centaurus A. The radio image is the central bright region of Fig. 1.5 magnified sixty times. This shows the vast difference between the perceived structure of sources in the optical and radio domains. The optical domains often reveal only a fraction of the corresponding image in the radio domain.

1.2.3 Kinds of Radiation

The voltages induced by cosmic source radiation are generally referred to as signals, although they do not contain information is the usual engineering sense. Such signals are generated by natural processes and have the form of Gaussian random noise. For most radio sources the characteristics of the signals are stationary with time, at least on the scale of minutes or hours typical of the duration of an observation. Gaussian waveforms of this type are assumed to be identical in character to the noise voltages generated in resistors and amplifiers. Such waveforms exhibit ergodicity.

Most of the power is in the form of *continuum radiation*, (Fig. 1.7, *Source: [6] Fig. 1.1*) the power spectrum of which shows slow variations with frequency and may be regarded as constant over the receiving bandwidths of most instruments.

In contrast to continuum radiation *spectral line radiation* (Fig. 1.8, *Source: [5] Fig. 1.2*) is generated at specific frequencies by atomic and molecular processes. A fundamentally important line is that of neutral atomic hydrogen at 1420.405 MHz, which results from the transition between two energy levels of the atom, the separation of which is related to the spin vector of the electron in the magnetic field of the nucleus.

A third kind of radiation is emitted from *pulsars* - an acronym for *PULSA*ting *R*adio *S*tar (Fig. 1.9, *Courtesy: Radio Observations of Mr. Divya Oberoi*). As its name suggests, the radiation from pulsars is not continuous but appears in pulses with a certain duty cycle. The on-time, that is the time when radiation is observed from the pulsar, forms a basis for a broad classification of pulsars - they are loosely classified as millisecond v/s long-period (also called slow) pulsars.

Apart from the intensity of the observed radiation, the other important aspect is its state of polarisation. Polarimetry - the study of the polarisation of the incoming radiation, forms an important part of Radio Astronomy since it helps to reveal the distribution of the magnetic fields within the source. The measure of polarisation that is almost universally used in astronomy is the set of four parameters introduced by Sir George Stokes in 1852. Known as *Stokes parameters*, they are related to the amplitudes of the components of the electric field E_x and E_y, resolved in two perpendicular directions normal to the direction of propagation from a radiating point. Thus if E_x and E_y are represented by $\mathcal{E}_x(t)\cos[2\pi\nu t + \delta_x(t)]$ and $\mathcal{E}_y(t)\cos[2\pi\nu t + \delta_y(t)]$, respectively, the Stokes parameters are defined as,

$$I = \langle \mathcal{E}_x^2(t)\rangle + \langle \mathcal{E}_y^2(t)\rangle \quad (1.6)$$

$$Q = \langle \mathcal{E}_x^2(t)\rangle - \langle \mathcal{E}_y^2(t)\rangle \quad (1.7)$$

$$U = 2\langle \mathcal{E}_x(t)\mathcal{E}_y(t)\cos[\delta_x(t) - \delta_y(t)]\rangle \quad (1.8)$$

$$V = 2\langle \mathcal{E}_x(t)\mathcal{E}_y(t)\sin[\delta_x(t) - \delta_y(t)]\rangle, \quad (1.9)$$

where, the angular brackets denote the expectation or time average. It can be shown that if the antennas provide simultaneous outputs for opposite senses of rotation (R and L, for Right and Left

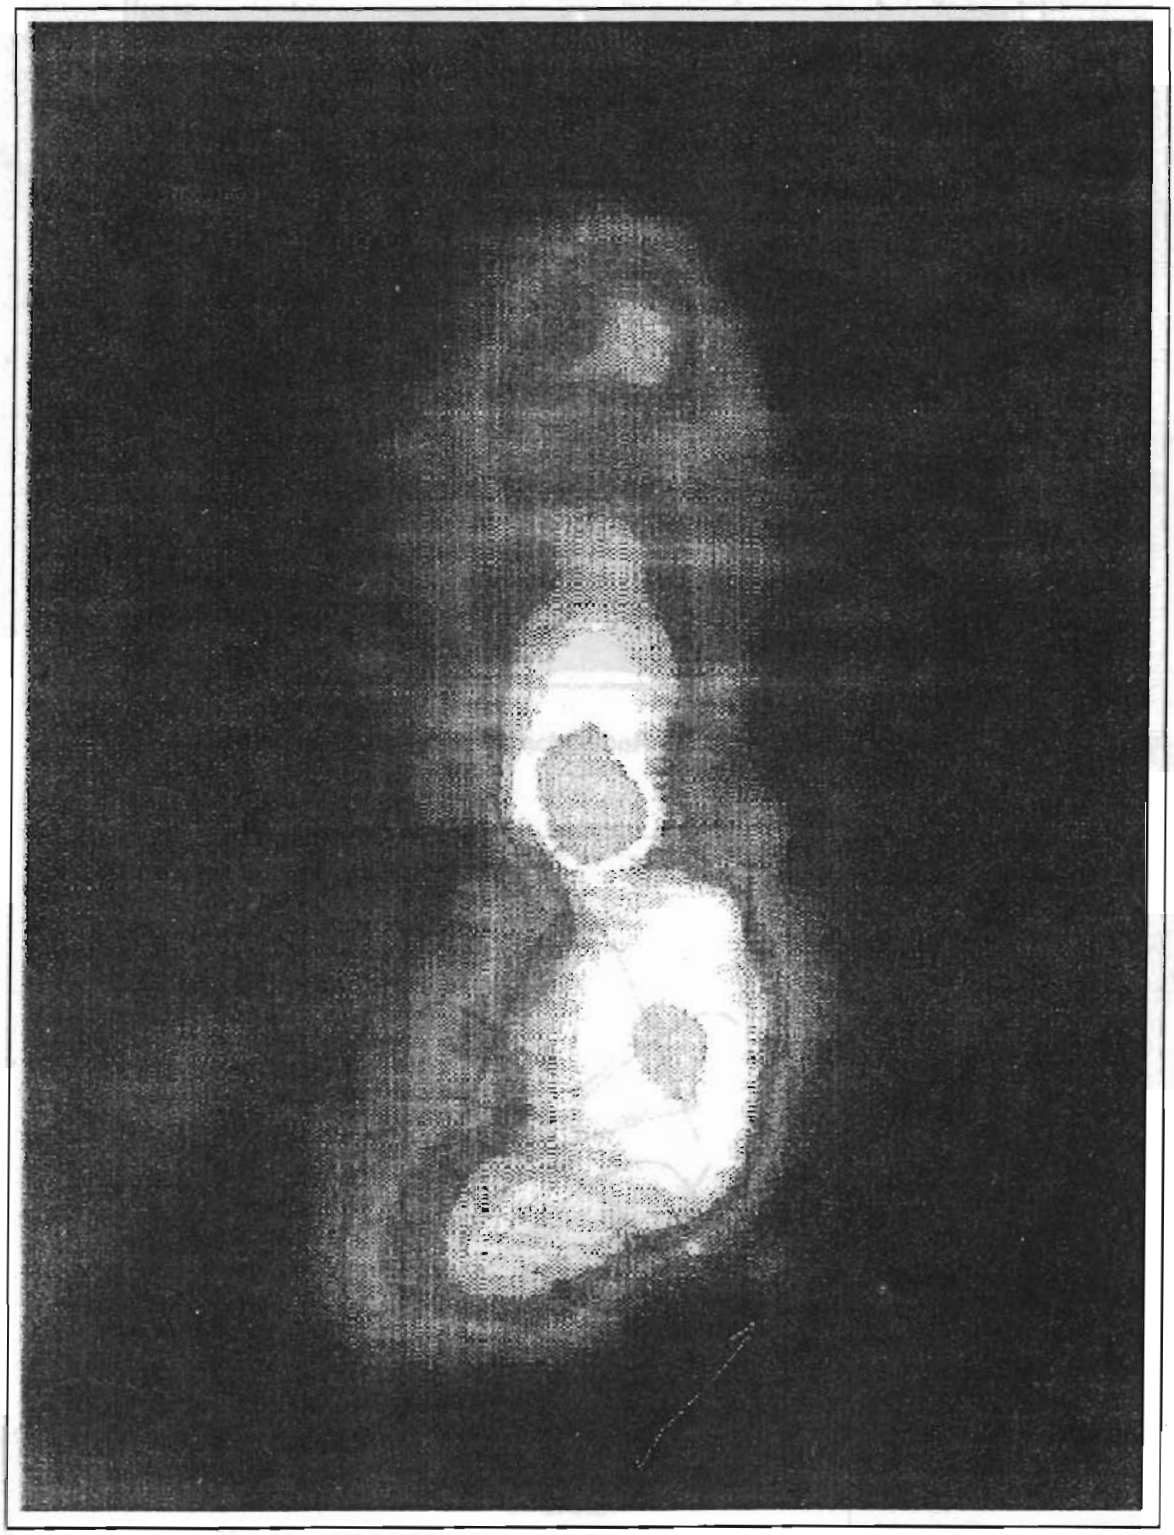

Figure 1.5: A Radio Map : Radio Galaxy Centaurus A

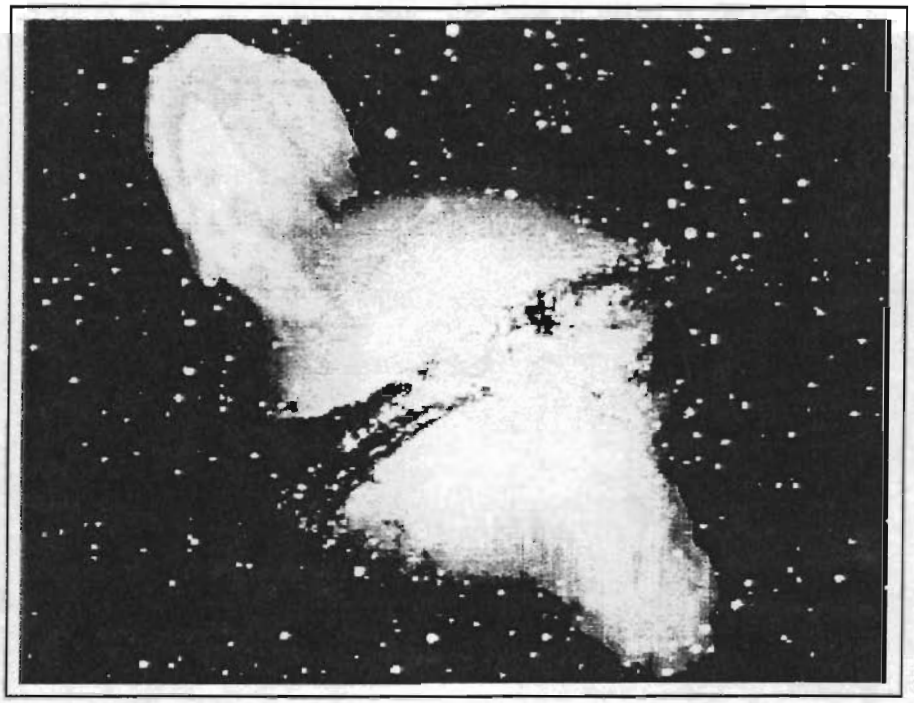

Figure 1.6: An Optical and Radio Map : Radio Galaxy Centaurus A

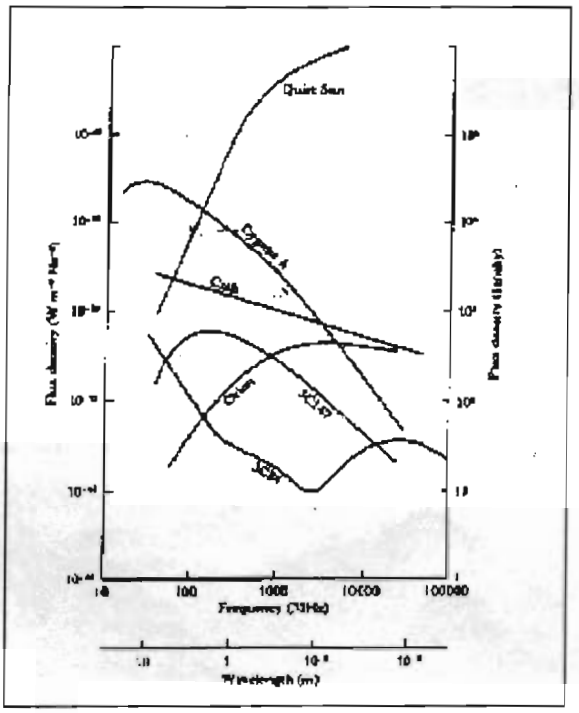

Figure 1.7: Continuum Emission from various Radio Sources

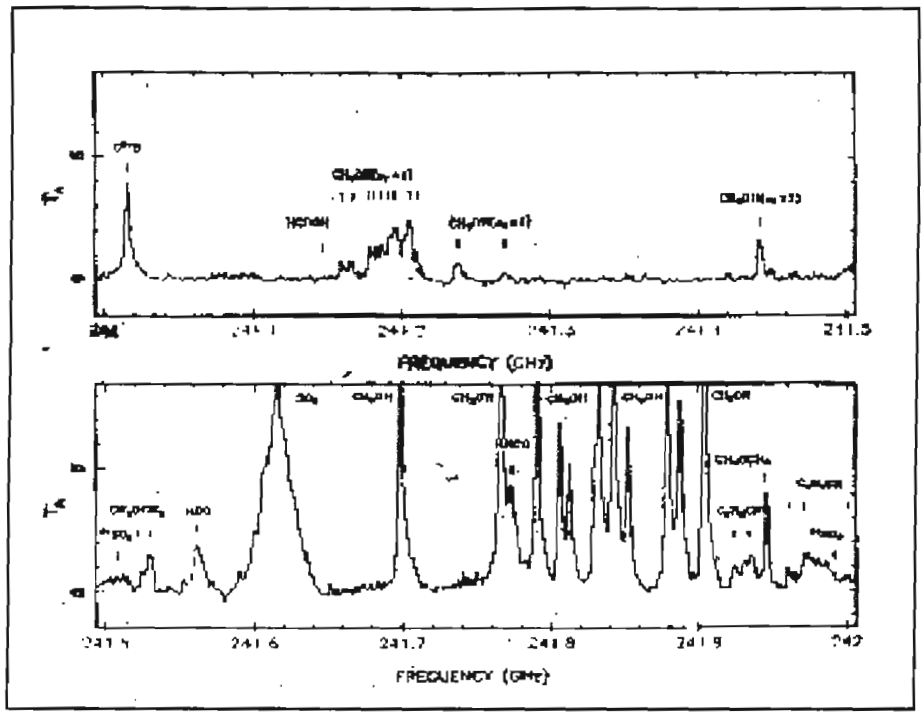

Figure 1.8: Spectral lines from the Orion Nebula

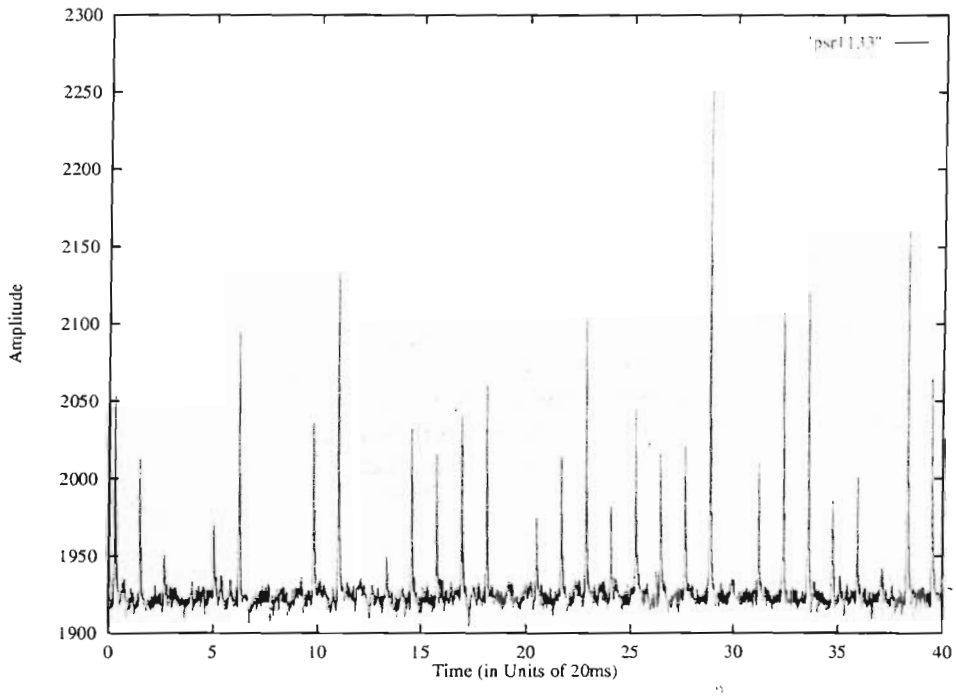

Figure 1.9: Emission from a Pulsar

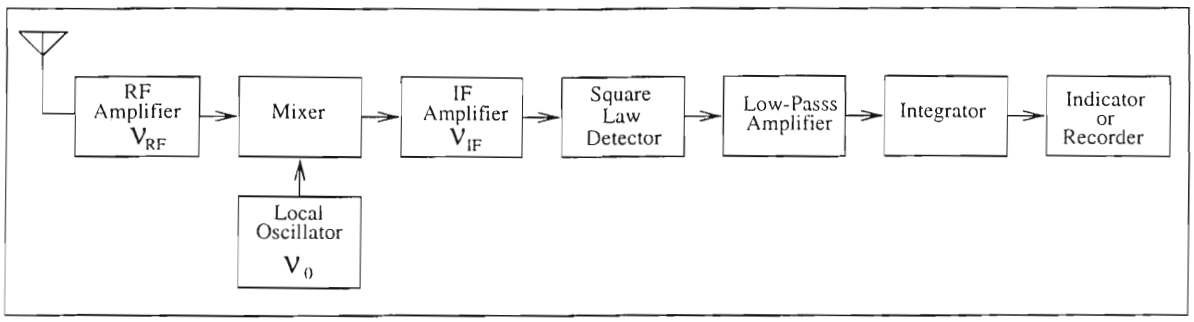

Figure 1.10: A Superheterodyne Radio-Telescope Receiver

Circular polarisation respectively) and four correlators are used for each antenna pair, the outputs can be related to the Stokes parameters by the matrix

$$\begin{bmatrix} \langle RR^* \rangle & \langle RL^* \rangle \\ \langle LR^* \rangle & \langle LL^* \rangle \end{bmatrix} = \begin{bmatrix} I+V & Q+iU \\ Q-iU & I-V \end{bmatrix} \tag{1.10}$$

1.2.4 Radio-Telescope Receivers

The function of a radio-telescope receiver is to detect and measure the radio emission of celestial sources. In most cases, the emission consists of incoherent radiation whose statistical properties do not differ from the noise originating in the receiver or from the the background radiation coupled to the receiver by the antenna. The power level of the signal in radio-telescope receivers is usually quite small, of the order of 10^{-15} to 10^{-20} watt. The power received from the background may be much higher than this, so that both high sensitivity and high stability of the receiver are necessary requirements. In certain cases, the ability to detect the signal spectrum as a function of time, too, is required.

Radio-telescope receivers are most commonly of the *superheterodyne* type and is depicted in Fig. 1.10(*Source: [4] Fig. 7.2*). The signal power, having a centre frequency ν_{RF}, is coupled to the receiver by an antenna and is first amplified in a radio-frequency (RF) amplifier. The next stage is a mixer, where the weak signal is mixed with a strong local-oscillator signal at a frequency ν_0 producing an output signal on an intermediate frequency (IF), the IF signal power being directly proportional to the RF signal power. The IF signal is then amplified. The largest part of the gain in a superheterodyne receiver is obtained in this IF amplifier, which also usually determines the pre-detector bandwidth of the receiver. The IF amplifier is followed by a detector, which is normally a square-law device. This means that the output d-c voltage of the detector is directly proportional to the output noise power of the pre-detection section of the receiver. Final stages may consist of a low-pass amplifier or integrator and a data-recording system.

As described earlier, stability of receivers is an important factor in radio-telescope receivers. The total-power receiver can be affected due to gain variations and fluctuations arising along the chain. A *switched receiver* known as *DICKE RECEIVER* is more immune to such variations and is shown in Fig. 1.11, (*Source: [4] Fig. 7.13*).

In a radio interferometer receiver, each antenna of the array may be equipped with its own pre-amplifier or the entire pre-detection section of the receiver (Fig. 1.12, *Source: [4] Fig. 7.18*). The receiver characteristics are assumed to be identical and the antennas pointing in the same direction. A total power receiver in this case consists of suitable phasing of individual elements, adding voltages from each IF output, and passing through a square law detector. Such an array is also called a "phased array". This simple interferometer receiver has the same drawbacks as the total-power receiver. Though, synchronous switching can be used in both receivers to stabilise

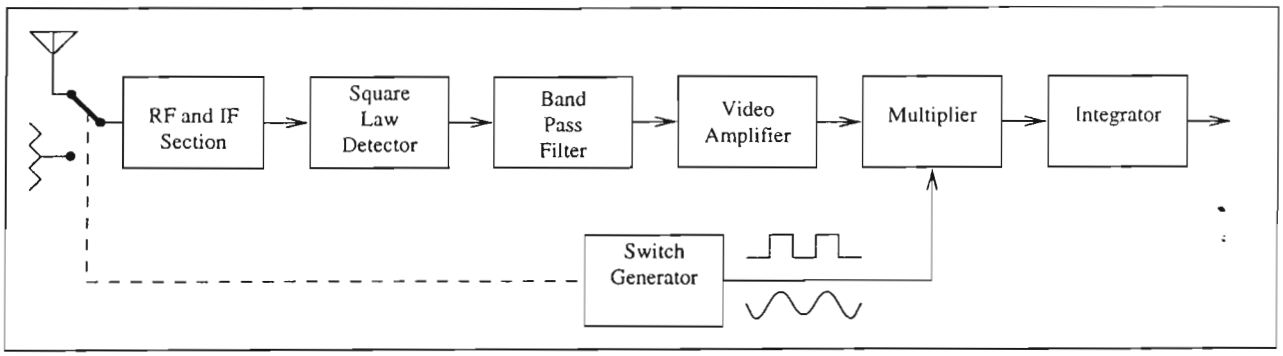

Figure 1.11: A Dicke Receiver

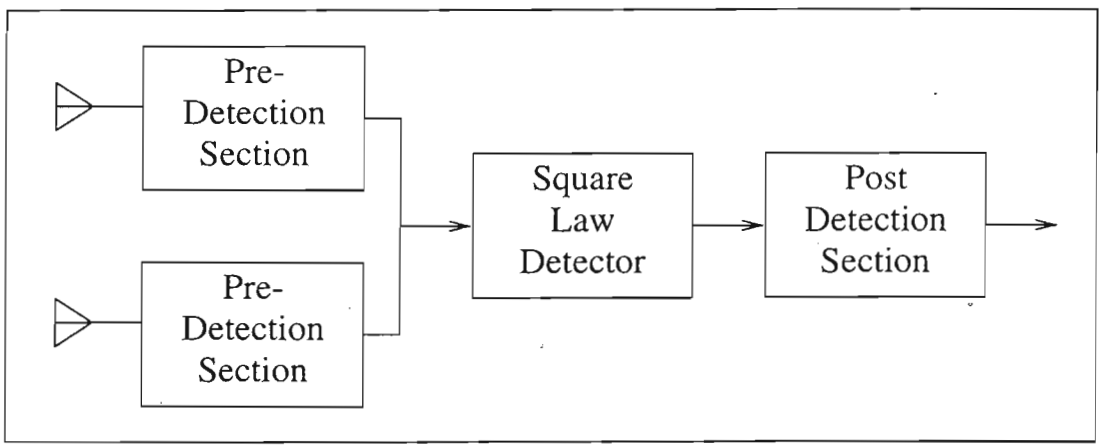

Figure 1.12: A simple Interferometer Receiver

them, a *CORRELATION* method is usually preferred.

Correlation Receivers

In a correlation receiver the IF outputs from the two separate receivers, as in an interferometer, are multiplied instead of added and detected. The correlation receiver is depicted in Fig. 1.13(*Adapted from [9] Fig. 2.1*). If the voltages from antennas 1 and 2 are $V_1(t)$ and $V_2(t)$ then,

$$V_1(t) = V_{c1}(t) + V_{u1}(t), \quad (1.11)$$

$$V_2(t) = V_{c2}(t) + V_{u2}(t), \quad (1.12)$$

$$V_{c2}(t) = V_{c1}(t + \tau_g) \quad (1.13)$$

where $V_{c1}(t)$, $V_{c2}(t)$ denote the correlated part in both antennas due to the source; while $V_{u1}(t)$, $V_{u2}(t)$ denote the uncorrelated components in the two antennas. τ_g is the *geometrical* path delay between the two antennas suffered by the signal as it reaches antenna 2.

The output of the correlator shall then correspond to the time average of the multiplication of the two antenna voltages, thus

$$V_o(t) = \langle V_1(t)V_2(t)\rangle = \langle (V_{c1}(t) + V_{u1}(t))(V_{c2}(t) + V_{u2}(t))\rangle = \langle V_{c1}(t)V_{c2}(t)\rangle \quad (1.14)$$

since, all other products involve uncorrelated voltages whose contribution over time will tend to zero. Here it is assumed that all voltages are zero mean. Thus, the correlation receiver is sensitive only to signals received by both antennas simultaneously. Due to this, voltage gain instabilities will not affect the sensitivity of the correlation receiver; it will only change its calibration.

If we represent $V_1(t)$ and $V_2(t)$ as $v_1 \cos 2\pi\nu(t)$ and $v_2 \cos 2\pi\nu(t - \tau_g)$. The output of the correlator then is,

$$r(\tau_g) = v_1 v_2 \cos 2\pi\nu\tau_g \quad (1.15)$$

τ_g varies slowly as the earth rotates, and results in the oscillation of the output. These oscillations are known as *fringes*. The principal astronomical information from the observation of a cosmic source by the interferometer is contained in the amplitude and phase of these fringes. The amplitude and phase is thus, of the modulation resulting from the interference between the outputs of the two antennas, and not those of the radiofrequency signal. Here it is assumed that the frequency of these fringes is lower than the cutoff provided by the integrator while at the same time filtering out the component at 2ν. $v_1 v_2$ which represents the fringe amplitude is proportional to the received power.

Kinds of Correlators

Generalisation of the above for the case of wide bandwidth signals leads to the development of spectroscopic correlators. Such correlators are required in the observation of spectral line sources in order to measure the line profile within the observed band. And, even in the case of continuum sources, it is often desirable to have not just the cross power over the whole bandwidth, but the cross power spectrum; e.g. in order to properly compensate for bandpass variations or to detect manmade interference dominated by one or more narrow band spikes corresponding to CW. For these reasons, correlators of most synthesis telescopes are required to measure the cross power spectrum on each baseline [1].

There are two major ways of implementing a correlator that measures the cross power in several sub-bands of the received bandwidth, according to whether the frequency analysis is done before or after the cross correlation. Based on the relative sequence of frequency-analysis (F) and cross-correlation (X), the correlators are classified as *FX* (Fig. 1.14, *Adapted from [8], Fig. 4.4*) or *XF* (Fig. 1.15, *Adapted from [8], Fig. 4.5*). The block labelled as "MAC" in the FX Correlator is a "complex correlator" shown in Fig. 1.16(*Adapted from [8], Fig. 4.3*).

[1] A pair of antennas

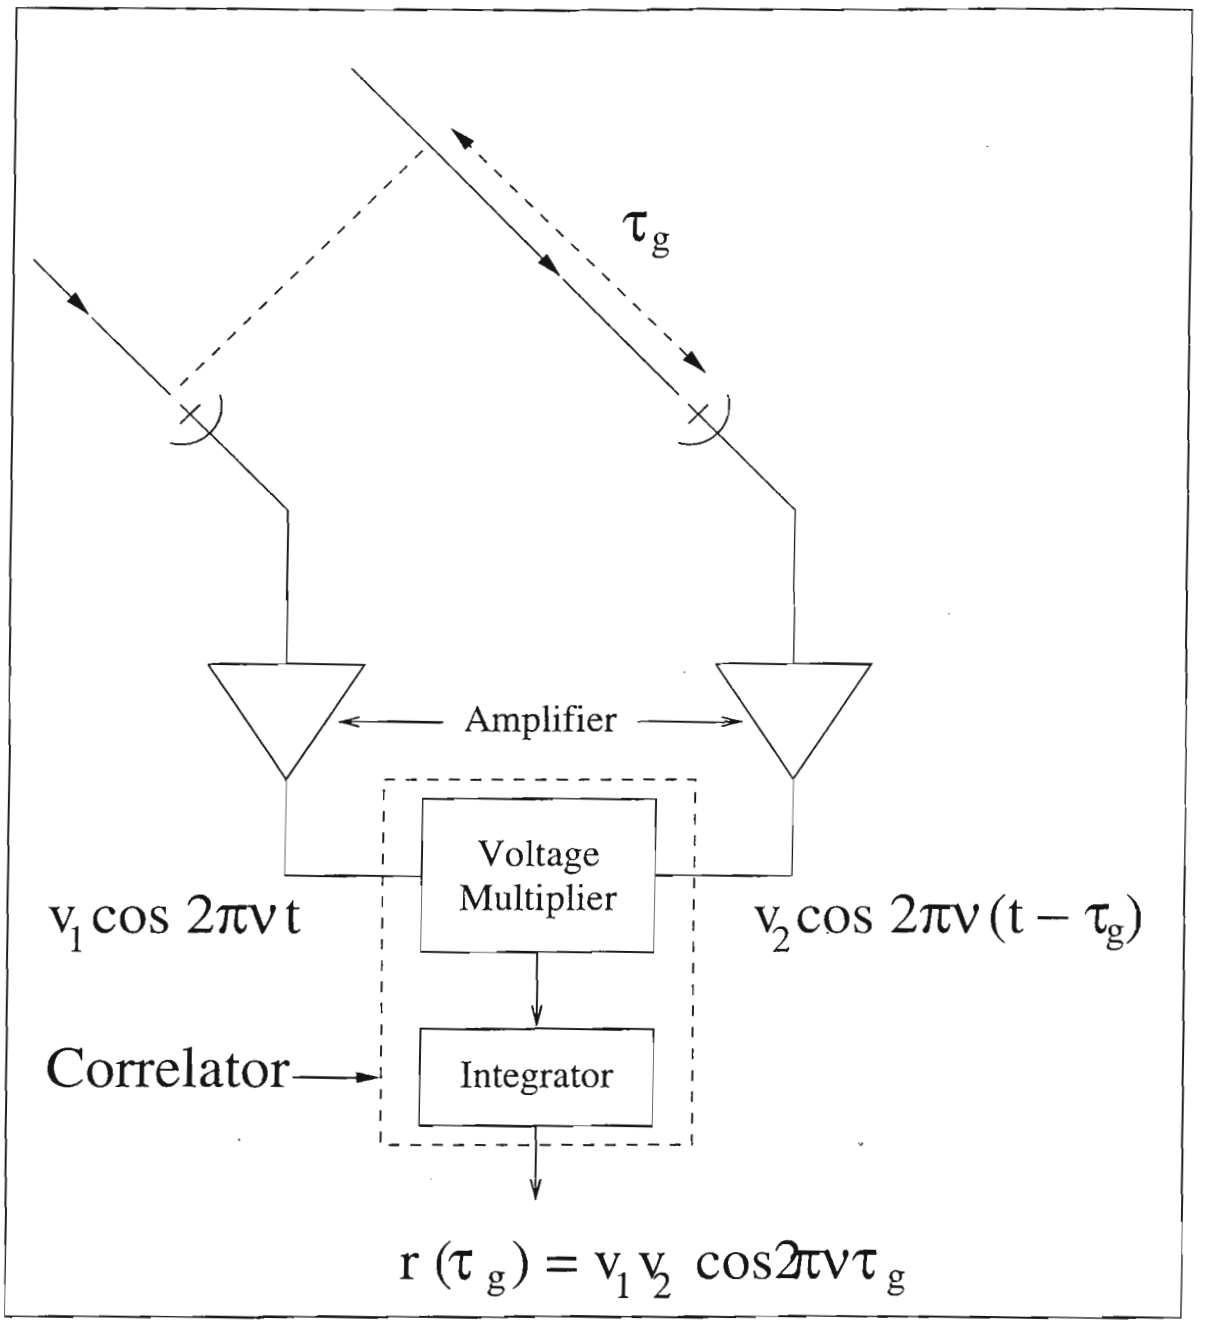

Figure 1.13: A Correlation Receiver

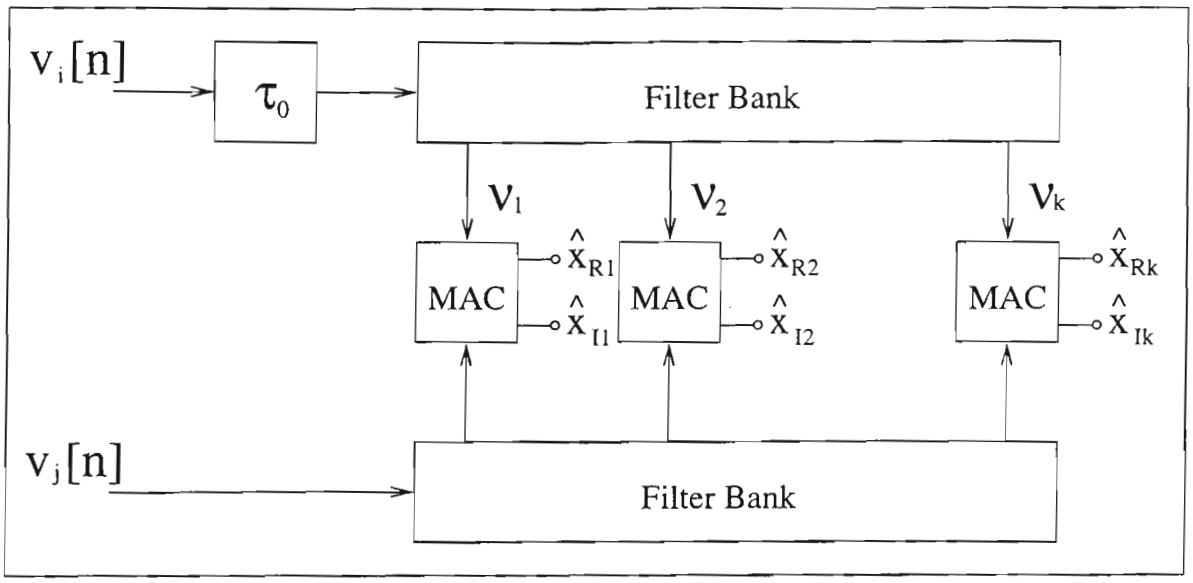

Figure 1.14: A FX Correlator

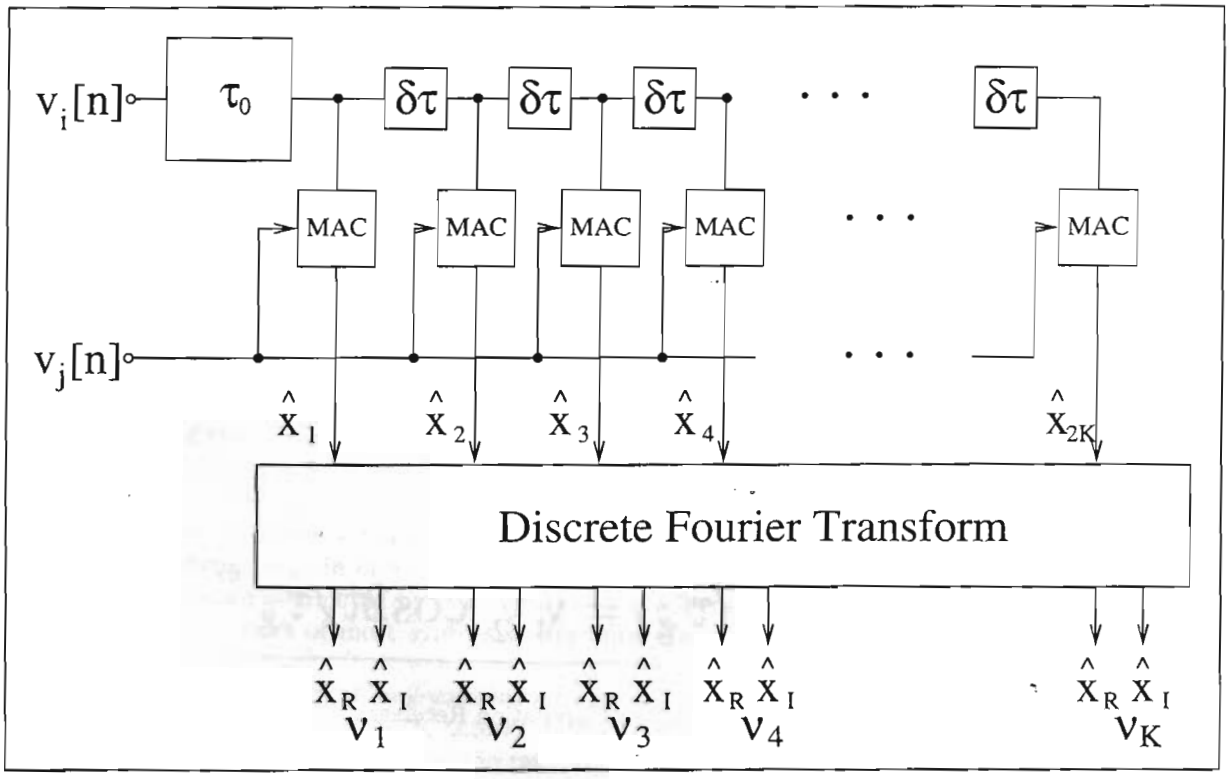

Figure 1.15: A XF Correlator

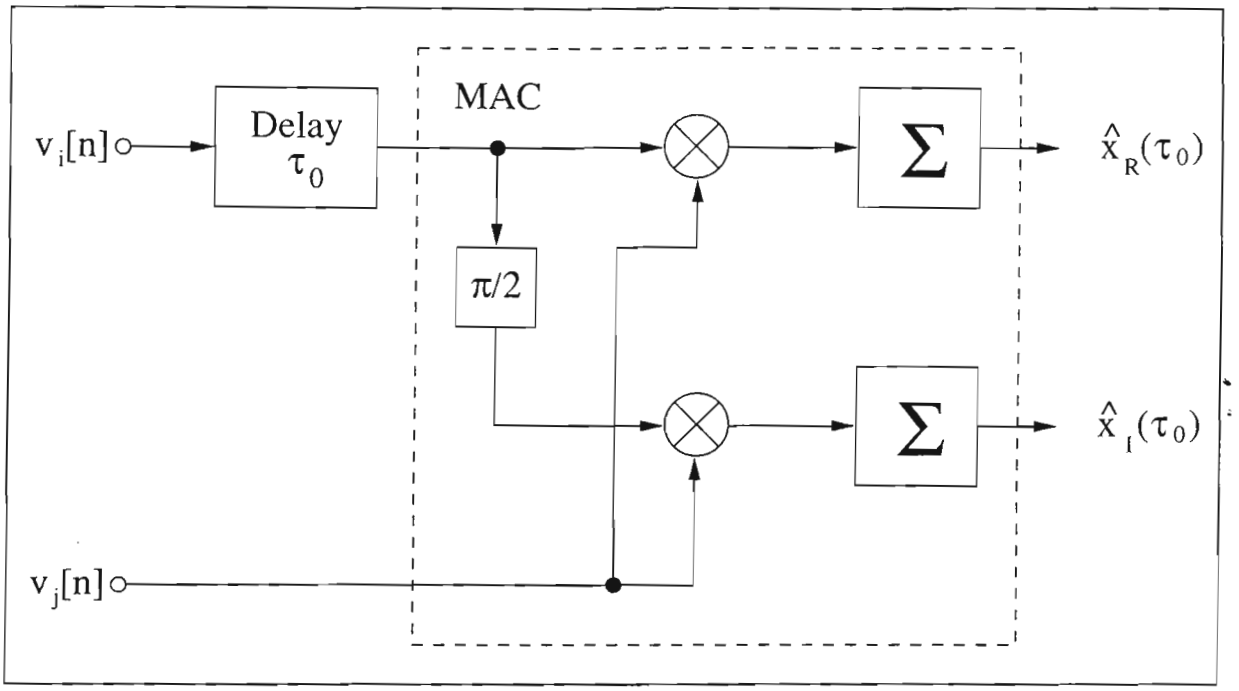

Figure 1.16: A "Complex Correlator"

Major portions of modern correlators are implemented digitally since, (1) digital operations are precisely defined and are repeatable; (2) digital circuits can be exactly replicated at low cost when many identical elements are needed; (3) for the long baselines and wide bandwidths, the delay lines must have a large ratio of length to resolution, and digital delay lines can do this with the necessary accuracy and stability.

1.2.5 Modern Synthesis Arrays

To provide a perspective for the later discussion, a survey of major contemporary modern aperture synthesis arrays is beneficial. In addition to the Giant Metrewave Radio Telescope (GMRT) which is described in the next section, there are at present four such arrays in operation around the world. These are

The Very Large Array (VLA)

Situated on the plains of San Augustin, New Mexico, USA, the VLA [11], the first array of its kind. Its antennas are arranged in a 'Y' shaped formation and are movable on rail tracks, so that the spacings between them can be changed. Each arm of the 'Y' is approximately 21km long. The antennas are equipped with cryogenically cooled receiving systems and are interconnected by low-loss, large-diameter waveguide.
No. of Antennas : 27
Diameter, Type and Mount : 25m, solid dish, altitude-azimuth
Feed Structure : Cassegrain
Correlator : XF

The Westerbork Synthesis Radio Telescope (WSRT)

Situated in Netherlands, the WSRT [11], is an East-West array 1.6km long. The antennas are interconnected by coaxial cables. The specialty of this telescope is its extremely stable receiver system.
No. of Antennas : 12 (10 fixed and 2 movable on rail-tracks)
Diameter, Type and Mount : 25m, wire mesh, equatorial
Feed Structure :Prime Focus supported by quadripods
Correlator : XF

The Australia Telescope (AT)

The AT [10] is a collection of eight antennas located at three different sites in New South Wales. Six identical antennas form the "Compact array" and are located near Narrabri, the seventh is located near Coonabarabran (called the "Mopra" dish) and the eighth is located at Parkes. This array has the distinction of being the sole current source of observations of southern skies at radio wavelengths. The compact array uses optical fibres to interconnect the six antennas.
No. of Antennas : 8. The compact array is made of two sections of rail tracks - Five antennas on a 3-km East-west track and the sixth on a 75m track further 3km to the west.
Diameter, Type and Mount : Compact array and Mopra - 22m, solid dish, altitude-azimuth. Parkes - 64m, Solid dish, altitude-azimuth.
Feed Structure : For all, the Compact array, Mopra and Parkes - Cassegrain
Correlator : XF

The Very Long Baseline Array (VLBA)

The VLBA [12] has its antennas spread all over the USA. This array has the distinction to be the first large array to employ a FX correlator.
No. of Antennas : 10
Diameter, Type and Mount : 25m, solid reflector, altitude-azimuth
Feed Structure : Cassegrain
Correlator : FX

1.3 The Giant Metrewave Radio Telescope (GMRT)

The most powerful radio-astronomical facilities in the world today operate mainly in the centimetre and decimetre part of the radio spectrum. Although there are many exciting and challenging astrophysical problems that are best studied at metre wavelengths, the more widespread prevalence of man-made radio interference at metre wavelengths in the Western world has possibly been responsible for the present neglect of metre-wave radio astronomy. Taking advantage of the much less radio interference in India, a new synthesis telescope, called the Giant Metrewave Radio Telescope (GMRT) [13] is being built as a major national facility for research in radio astronomy . The GMRT has been designed as a major new instrument that would fill the existing gap in radio-astronomy facilities at metre wavelengths. The major driving force behind the proposal of the telescope was detecting the highly redshifted '21-cm' line of neutral hydrogen from protoclusters or protogalaxies to probe the early epochs of the Universe before galaxy formation. The GMRT is also expected to detect and study a large number of millisecond pulsars in an attempt to detect the primordial background of gravitational radiation.

However, since it is the most powerful metrewavelength radio telescope in the world, it has also been built as a versatile instrument which can throw new light on a variety of astrophysical

problems by observing celestial radio sources ranging from planets, stars, nearby galaxies to very distant galaxies and the interplanetary, interstellar and intergalactic media.

The GMRT is an aperture synthesis array, consisting of 30 fully steerable parabolic dishes of 45m diameter each. Six antennas are distributed along each of the three arms of a rough 'Y' and the remaining 12 antennas are more or less randomly placed in a compact cluster near the centre of the 'Y' (Fig. 1.17). The GMRT operates in six frequency ranges centred at 50, 153, 233, 327, 610 and 1420 MHz. The receiver system is of superheterodyne type with a flexible local oscillator system. The backend is the digital correlator which processes the four signals of 16 MHz bandwidth from each antenna, corresponding to the two sidebands and two polarisations of the 30 antennas. Further details on GMRT will be given in Chapter 2.

This chapter prepared a basis for the forthcoming discussion on the design of the backend of the GMRT. Towards this end, a limited background of Astronomy and Radio Telescopes was provided in this chapter. The use of correlators in radio telescope receivers was discussed and a survey of contemporary synthesis arrays was given. The GMRT was introduced as a new major facility for research in metrewave astronomy.

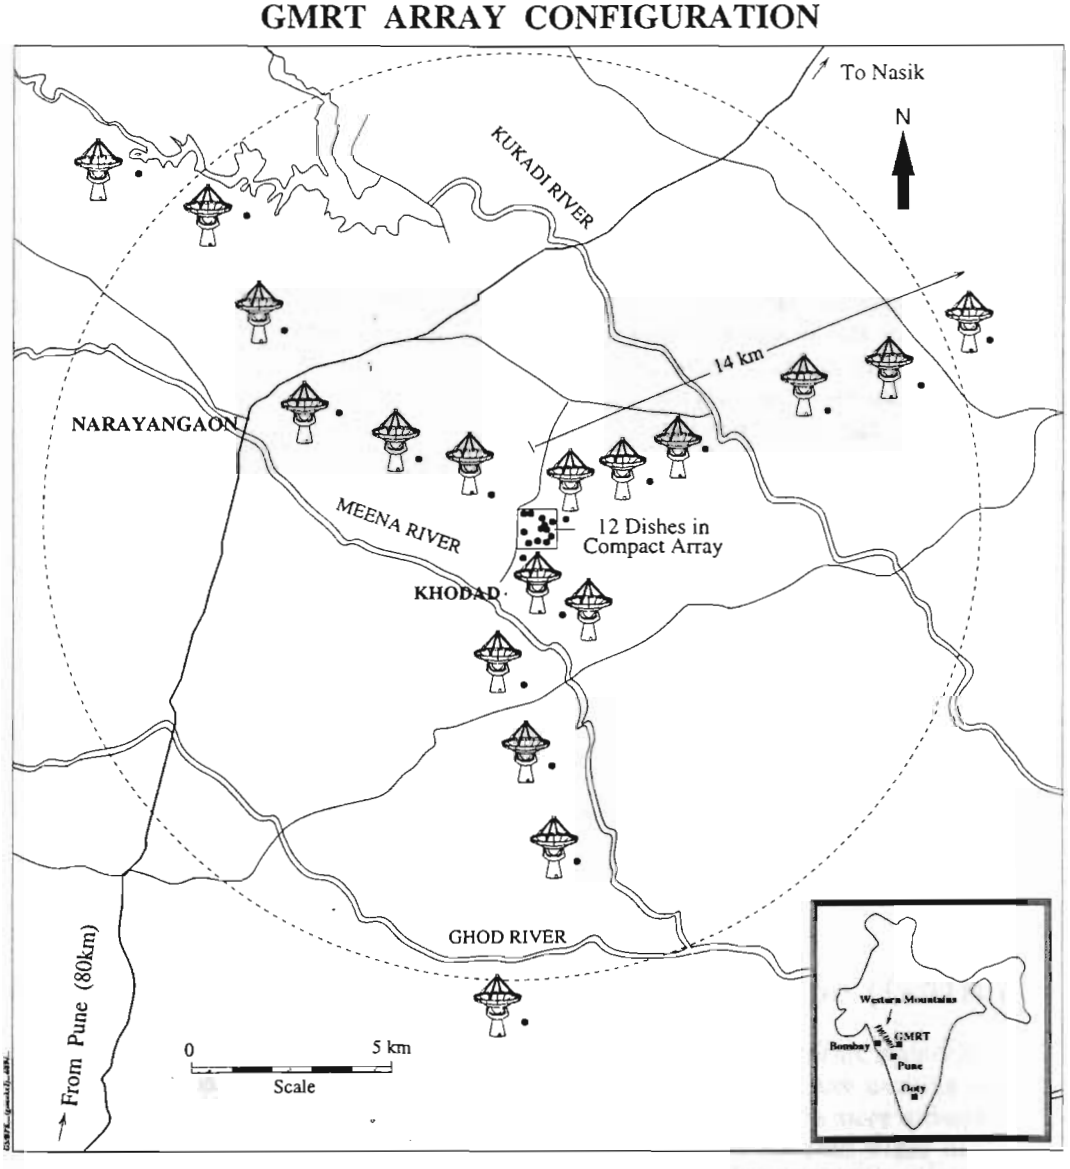

Configuration of the 30 dishes of GMRT. Twelve dishes are located in a compact central array in a region of about 1 km x 1 km and the remaining dishes are being distributed along the 3 arms of an approximate 'Y' configuration. By combining signals from all possible pairs of dishes, this arrangement will allow GMRT to image celestial radio sources with an angular resolution equivalent to that obtainable with a single giant dish, 25 km in diameter.

Figure 1.17: The GMRT Array

Chapter 2

DEFINING THE PROBLEM

To know that we know what we know, and that we do not know what we do not know, that is true knowledge.

-Thoreau, in 'Walden' (quoting Confucius)

The GMRT will be a major facility to explore the metrewavelength realm of the Universe. This chapter will endeavour to bring forth the criteria which went into deciding the type and specifications of its backend - the correlator. A comparison between two design choices, available at the time when the correlator was conceived, is included.

2.1 The Context of the Correlator

The following discussion shall elaborate the implication of the following aspects on the specification of the correlator.

- The Astronomical Objectives to be fulfilled by GMRT,
- The GMRT Receiver System and,
- The Configuration of the GMRT Array.

In order to present such a perspective, the material presented here is based on [13] and the various internal reports of the National Centre for Radio Astrophysics (NCRA). The internal reports of GMRT are not yet complete and the currently available reports can be found in the NCRA library.

2.1.1 Requirements due to Astronomical Objectives

Although the two major objectives of GMRT were protocluster search and the search for new short period pulsars, its versatility and high sensitivity endow it with the potential of making outstanding contributions in many areas of astrophysics, through the continuum and spectral line studies of celestial sources.

It is only in the very recent past (*the VLBA*) that large arrays have used spectroscopic correlators. The major contributing factor for the advent of the spectroscopic correlators has been the phenomenal progress in VLSI technology, whereby it has become feasible to construct Application

Specific Integrated Circuits (ASICs) tailored to the specific needs of the telescopes at costs much lower than ever before.

Thus, to have a correlator which could cater to both kinds of observations with equal ease, a spectroscopic correlator was considered as a prime requirement for GMRT. Secondly, the bandwidth to be processed was originally specified as 16MHz, which was finally increased by a factor of 2 to 32MHz to provide greater sensitivity for the pulsar search and other studies. Also, the protocluster search requirement dictated a frequency resolution of at least 1MHz at 16MHz bandwidth. This meant that the correlator provide at least 16 spectral channels. However, for most spectral line observations and for pulsar studies, a larger number of spectral channels are required. For the GMRT correlator, a maximum of 256 channels was specified.

Additionally, in order to provide for studying the polarisation properties of the observed source, the correlator was required to compute the Stokes parameters from the incoming data. The Correlator was also required to provide a range of integration times to support various kinds of observations.

2.1.2 Requirements due to the Receiver System

The GMRT receiver system too, like almost all systems of this kind, employs a superheterodyned scheme to provide observing capability at six different frequency bands with a maximum bandwidth of 32 MHz. Fig. 2.1 presents a simplified block diagram of the GMRT receiver system. Further details can be found in [13].

The antennas have dual polarised feeds at all the six frequencies. The feeds are mounted on the four faces of a computer controlled rotating turret at the prime focus of the dish. While three of the faces have feeds operating at 153, 327 and 1420 MHz, the fourth one has a dual frequency feed at 233/610 MHz. The 50 MHz feed is mounted on the rim and is always available for use.

The *Front End* consists of the Radio-Frequency (RF), Local-Oscillator (LO) and Intermediate-Frequency (IF) electronics. The linearly polarised signals are first converted to right-hand and left-hand circularly polarised signals in low-noise quadrature hybrids and then amplified in low-noise RF amplifiers. In order to help calibrating the receiver chain, provision exists for injecting a broadband noise of known temperature in the signal, by the RF system. Before transmitting the signal to the antenna base, the signal is phase switched using one of a set of 60 (two polarisations of the 30 antennas) mutually orthogonal Walsh functions, to minimise any coupling between antenna electronics. Low loss cables bring the signals to the antenna base where the RF signals are converted to a 32 MHz wide IF centred at 70 MHz using LO signals tied to a central frequency standard. The LO signals are phase-synchronised using a round-trip phase measurement. A monitor and control telemetry system is used to set the various parameters in the RF, LO and IF systems. The telemetry system also controls the servo system required to point the antenna in the desired direction. The LO and IF signals along with the telemetry signals are transmitted, from and to, a central control room called the Central Electronics Building (CEB). A low loss optical fibre link utilising analog transmission forms the link between the CEB and all antennas. The two 32MHz wide IF signals, one for each polarisation, are further downconverted to two baseband signals per polarisation each of 16MHz bandwidth. A choice of RF, IF and Baseband filters enables any desired bandwidth from 0.0625MHz to 16MHz to be selected in binary steps. The four 16 MHz signals per antenna at baseband are then transmitted to the correlator.

The correlator, therefore, processes 120 signals, each of 16MHz bandwidth from the 30 antennas of the GMRT. These are sampled at 32Ms/s to comply with the Nyquist sampling requirement. Before correlating signals from different stations, Walsh demodulation is required to recover the original signal. The correlator is also required to provide a continuous estimate of the power content of each of the 120 channels, in order to facilitate a continuous calibration of the system gain, using the injected noise at the front end.

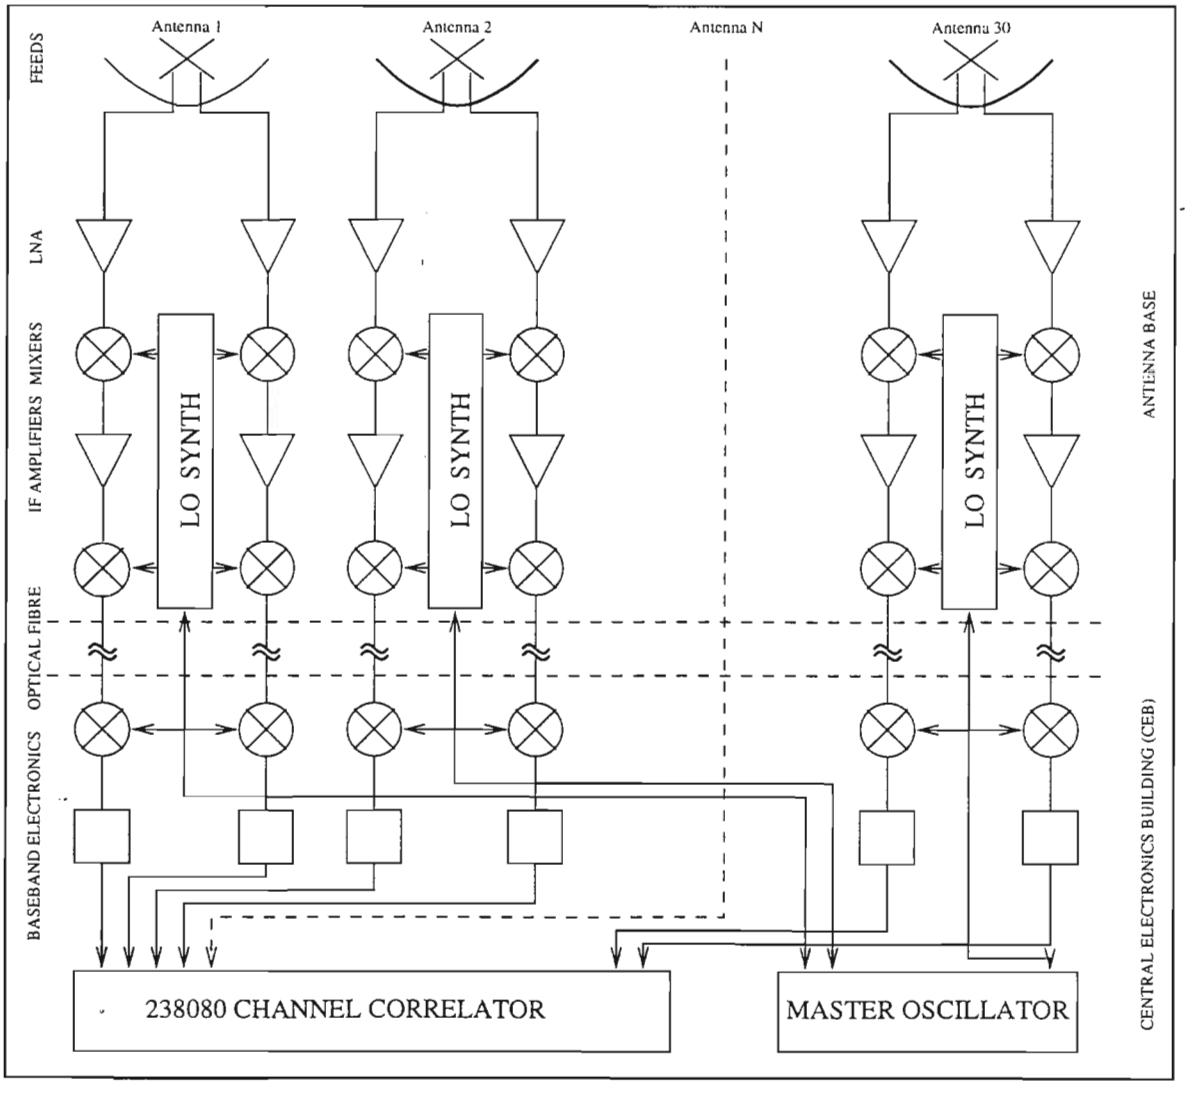

Figure 2.1: Block Diagram of the GMRT Receiver System

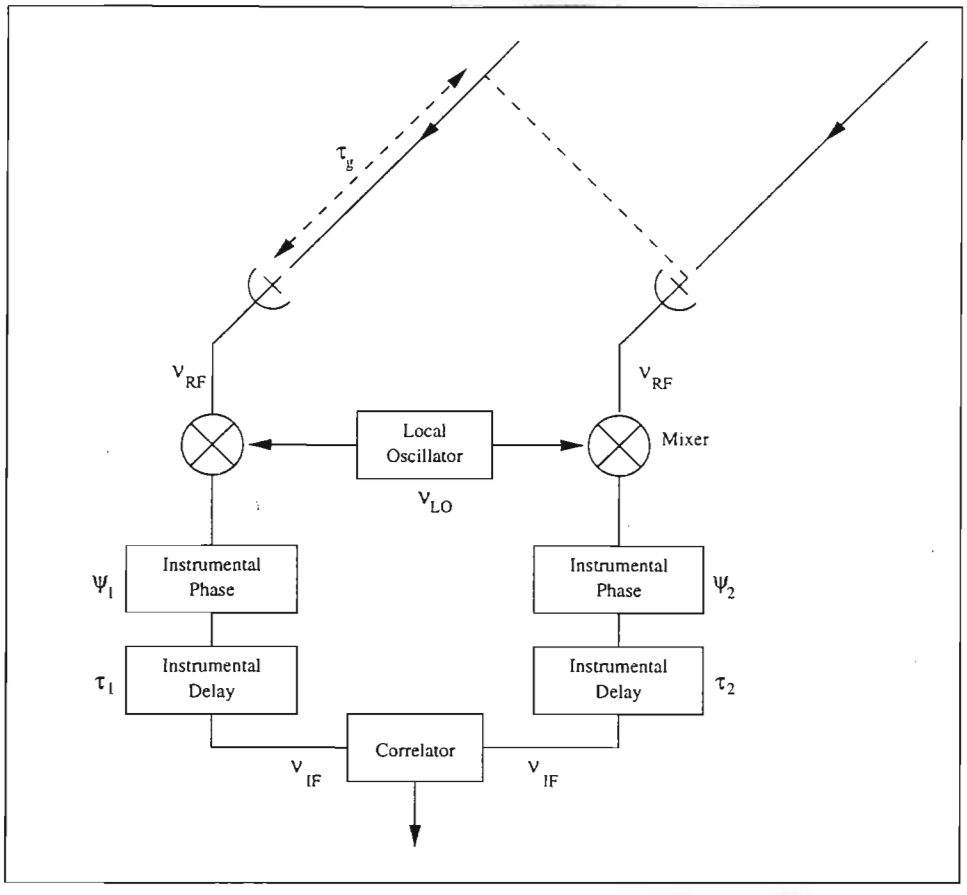

Figure 2.2: An Interferometer System

2.1.3 Requirements due to the Array Configuration

The first and foremost requirement stemming from the array configuration was that of supporting 30 stations. The correlator is thus required to compute spectral correlations arising from the 465 baselines. This number also includes the auto-correlation spectra from the 30 stations.

In order to simplify discussion of array requirements, a brief description of interferometry is given below.

The block diagram of an interferometer system including the sky delay, is shown in Fig. 2.2. Frequency conversion of the incoming signals at RF ν_{RF} with a LO at frequency ν_{LO} is also included. The IF ν_{IF} is related to ν_{RF} and ν_{LO} by

$$\nu_{RF} = \nu_{LO} \pm \nu_{IF} \tag{2.1}$$

Here ν_{LO} is a single valued frequency but ν_{RF} and ν_{IF} refer to bands of width $\Delta\nu$. Considering only one of the resulting sidebands, say the Upper Side-band (USB), the response can be obtained by considering the phase changes ϕ_1 and ϕ_2 imposed on signals received by antennas 1 and 2 before reaching the correlator.

$$\begin{aligned} \phi_1 &= 2\pi(\nu_{LO} + \nu_{IF})\tau_g + 2\pi\nu_{IF}\tau_1 + \psi_1, \\ \phi_2 &= 2\pi\nu_{IF}\tau_2 + \psi_2, \end{aligned} \tag{2.2}$$

where ν_{RF} is replaced by $(\nu_{LO} + \nu_{IF})$ and $(\tau_2 - \tau_1)$ is the instrumental delay that compensates for

τ_g. The response of the interferometer r_u will thus be:

$$\begin{aligned} r_u &= A_0|V| \int_{\nu_{IF_0}-\Delta\nu/2}^{\nu_{IF_0}+\Delta\nu/2} \cos(\phi_1 - \phi_2 - \phi_V) d\nu_{IF}, \\ &= A_0 \Delta\nu |V| \frac{\sin \pi\Delta\nu\Delta\tau}{\pi\Delta\nu\Delta\tau} \cos[2\pi(\nu_{LO}\tau_g + \nu_{IF}\Delta\tau) - \phi_V - \phi_{LO}] \end{aligned} \tag{2.3}$$

where, A_0 = antenna response at the beam centre
V = complex visibility of the source

Here, $\Delta\tau = \tau_g + \tau_1 - \tau_2$, is the tracking error of the compensating delay. It may be noted that τ_g depends on the direction of the celestial source being tracked, which varies continuously due to the earth's rotation. The output fringe oscillations, which result from the time variation of τ_g, depend on the LO frequency ν_{LO} rather than the observing frequency ν_{RF}. These fringe oscillations can be reduced, generally to zero, by varying ψ_1, ψ_2 such that $(2\pi\nu_{LO}\tau_g + \psi_1 - \psi_2)$ remains constant. This procedure of reducing the fringe frequency to zero is known as *fringe stopping.*

In GMRT, both, the fringe stopping and delay compensation are performed digitally within the correlator system. From eqn. 2.3, it is seen that the delay error must be small, $|\Delta\nu\Delta\tau| << 1$ in order to obtain good sensitivity. For the GMRT, a delay tracking accuracy of about 1ns was considered adequate. The maximum fringe frequency is $\sim$ 10Hz; and the maximum range of delays to be compensated is $\sim$ 100 μs.

2.2 The Type of Spectral Correlator

Summarising the requirements outlined above, the correlator should of spectroscopic type with upto 256 spectral channels and should process the 120 signals from the 30 antennas of GMRT to produce data pertaining to the 465 baselines. It should be able to compute the Stokes parameters, should account for delays of upto 100μs and fringe rates $\sim$ 10Hz. The delay compensation accuracy should be around 1ns. The correlator should provide a range of integration times to support various observations.

As was discussed in Chapter 1, spectroscopic correlators can be of either FX or XF types depending on the whether the cross-correlation succeeds or precedes the frequency-analysis. These two types are discussed next, with regard to their conceptual, realisation and economical aspects.

2.2.1 The XF and FX Correlator

The comparison between these two correlator algorithms can be made on several accounts. The information available in the literature ([8]) and [14]) has been viewed from the following three points-of-view.

- The transfer function
- The hardware requirements
- The overall receiver system

The Transfer Function

The XF Correlator The basic XF correlator was shown in Fig. 1.15. Further details of a digital implementation are shown in Fig. 2.3 (*Adapted from [8], Fig. 4.9*) which gives a block diagram for a 2-bit, 3-level quantisation in the digitiser. The main features highlighted here, are

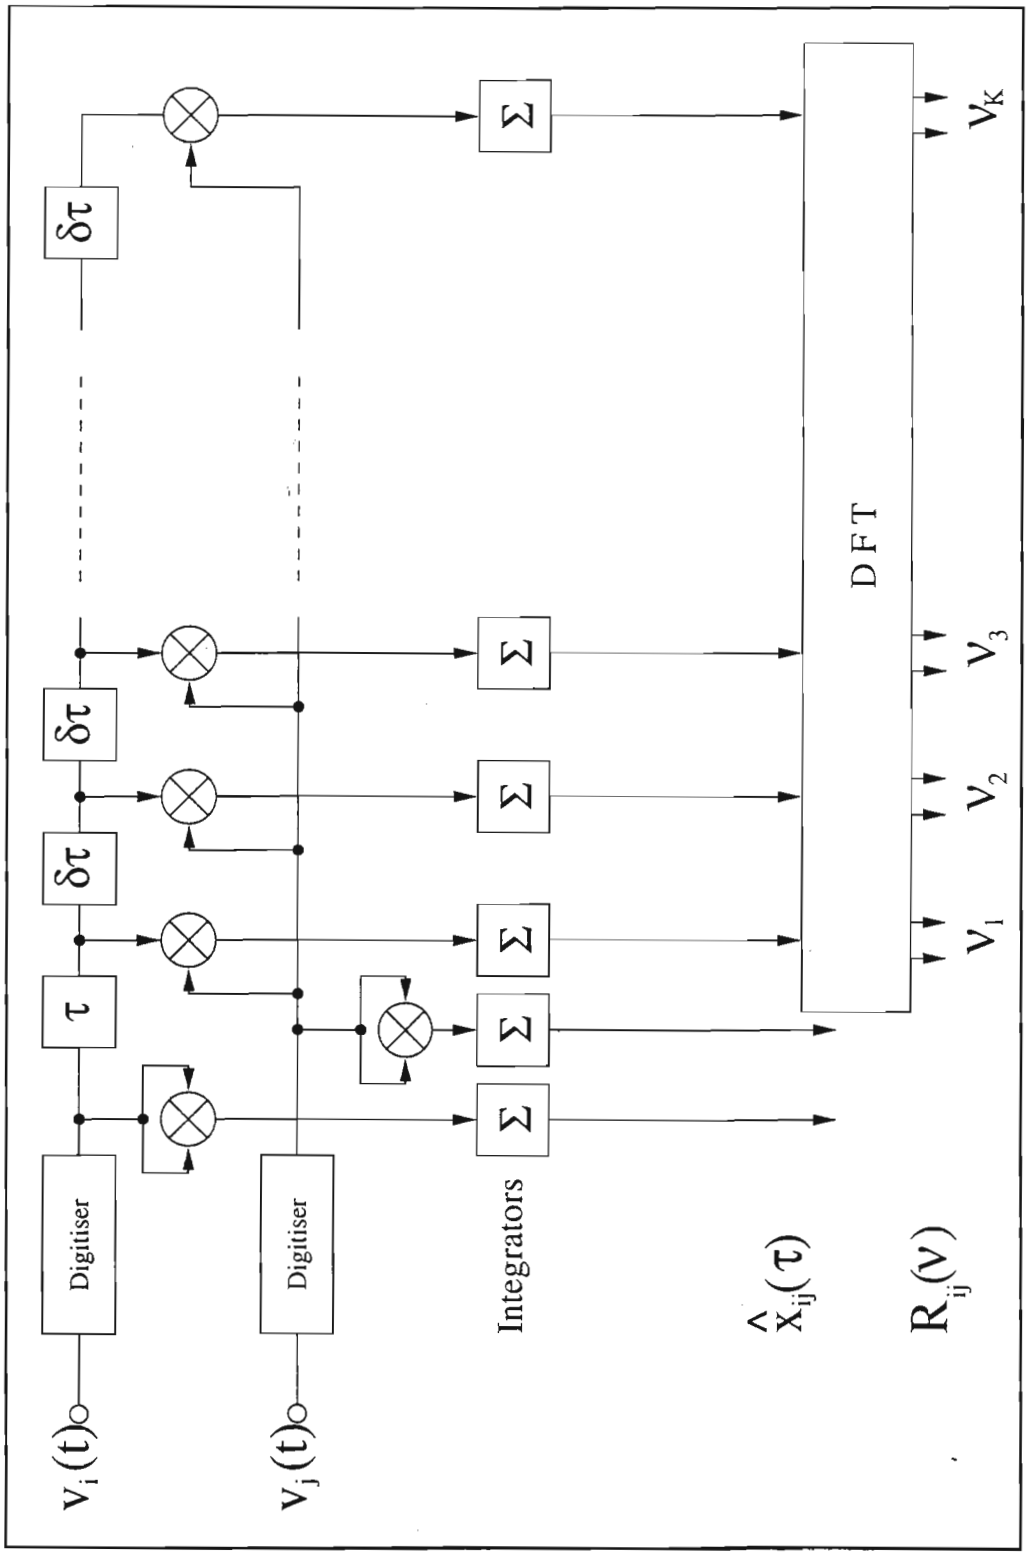

Figure 2.3: A XF Correlator

1. The small delays $\delta\tau$ that determine the spacing of the correlation function measurements are implemented in a shift register, which means that they must be multiples of the sampling interval;
2. By including a self-multiplier for each digitised signal, the power in each signal can be determined.

The cross power spectrum can be written as an infinite sum due to the sampling theorem, as

$$\begin{aligned} r(\nu) &= \int_{-\infty}^{\infty} x(\tau)e^{-2\pi i\nu(\tau-\tau_0)}d\tau \\ &= \sum_{k=-\infty}^{\infty} x(\tau_0+k\delta\tau)e^{-2\pi i\nu k\delta\tau}\delta\tau, \end{aligned} \tag{2.4}$$

where the second equation holds only within the bandwidth $0 \leq \nu \leq \Delta\nu < 1/(2\delta\tau)$. If the sum is truncated beyond $|k| = K$, the result may be written

$$\begin{aligned} \hat{r}(\nu) &= \sum_{k=-\infty}^{\infty} \sqcap(k/2K)x(\tau_0+k\delta\tau)e^{-2\pi i\nu k\delta\tau}\delta\tau \\ &= \int_{-\infty}^{\infty} \sqcap(\tau/2K\delta\tau)\text{Ш}(\tau/\delta\tau)x(\tau)e^{-2\pi i\nu(\tau-\tau_0)}d\tau \\ &= r(\nu) * \int_{-\infty}^{\infty} \sqcap(\tau/2K\delta\tau)\text{Ш}(\tau/\delta\tau)e^{-2\pi i\nu(\tau-\tau_0)}d\tau, \end{aligned} \tag{2.5}$$

where $\sqcap(.)$ is the unit rectangle function and Ш(.) is the unit sampling function. The last integral may therefore be regarded as the bandpass function of a single channel; for large K, it is approximately $Ksinc(K\nu\delta\tau)$.

Eqn. 2.5 can be computed by the DFT, but only for certain values of ν. Normally, a length-K FFT is used, giving K complex results at $\nu = 0, 1/(K\delta\tau), \ldots, (K-1)/(K\delta\tau)$. Values at intermediate frequencies can be computed by means of a longer FFT with zeros inserted for $x(\tau_0 + k\delta\tau)$ when $|k| \geq K$.

The integrators normally accumulate data for a large number of samples before being read and reset to begin a new measurement. For this reason, the cost of the computing operations done after the integrators, including the FFTs, is generally negligible compared to earlier operations.

The FX Correlator The architecture of a digital FX correlator is shown in Fig. 2.4 (*Source: [8], Fig. 4.10*). Earlier, in Chapter 1, Fig.1.14 was a simpler illustration of a FX correlator. Here the filter banks have been replaced by DFTs that operate on successive segments of the digitised signals. In its simplest form, the correlator computes a length-K transform on non-overlapping segments of $2K$ samples each. The DFT of a length-K segment of the (delayed) samples from the antenna i is given by

$$V_i(\nu) = \sum_{k=-\infty}^{\infty} \sqcap(k/2K)v_i(k\delta t - \tau_0)\delta t, \tag{2.6}$$

for certain frequencies ν; and $V_j(\nu)$ is the similar result for the samples from antenna j (without the τ_0 delay). The estimated cross power spectrum is the average of the products of these transforms at each frequency:

$$\hat{r}(\nu) = \langle V_i(\nu)V_j^*(\nu)\rangle$$

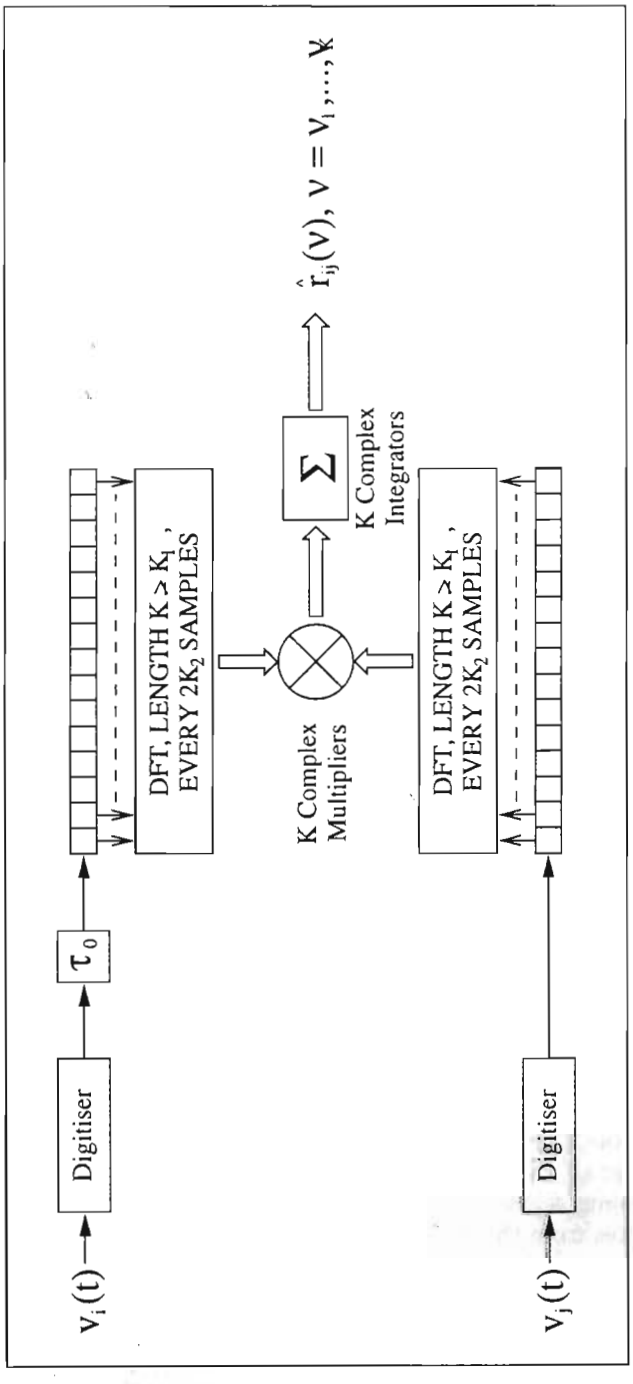

Figure 2.4: A FX Correlator

$$= \sum_{k=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \sqcap(k/2K) \sqcap (k'/2K) \langle v_i(\tau_0 + k\delta t) v_j(k'\delta t) \rangle e^{-2\pi i \nu (k-k')\delta t} \delta t$$
$$= \sum_{k=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \sqcap(k/2K) \sqcap (k'/2K) x_{ij}(\tau_0 + (k-k')\delta t) e^{-2\pi i \nu (k-k')\delta t} \delta t. \tag{2.7}$$

The double sum may be rewritten by letting $l = k - k'$ as

$$\sum_{l=-\infty}^{\infty} W_l R_{ij}(\tau_0 + l\delta t) e^{-2\pi i \nu l \delta t} \delta t, \tag{2.8}$$

where,

$$W_l = \begin{cases} |2K - l|, & \text{if } |l| < 2K \\ 0, & \text{otherwise.} \end{cases} \tag{2.9}$$

Using the unit sampling function Ш(.) to put this into integral form gives

$$\hat{r}(\nu) = r(\nu) * \int_{-\infty}^{\infty} 2K\Lambda(\tau/4K\delta t) Ш(\tau/\delta t) e^{-2\pi i \nu(\tau - \tau_0)} d\tau, \tag{2.10}$$

where $\Lambda(.)$ is the unit triangle function. As before, the integral is the bandpass function of each channel; for large K, it is approximately $2K sinc^2(K\nu\delta t)$.

This shows that the FX correlator response is equivalent to triangularly- weighted measurements of the cross correlation function over a range of $4K\delta t$ in delay. By comparison, the XF correlator response is a uniformly-weighted transform of direct measurements of the cross correlation over a range of $2K\delta t$. In both cases, the weighting can be modified by multiplying the data by an additional weighting function before the DFT, and this can be used to trade resolution for an improved shape of channel bandpass.

Hardware Requirements

In order to describe the cost of a spectral correlator, the following equation can be developed. If each analog filter costs C_a, independent of bandwidth, a $K-$ channel digital filter costs $C_d(K)$ per unit bandwidth (proportional to bandwidth), and each cross-correlator costs C_x per unit bandwidth. Then a complete correlator for N, antennas, total bandwidth $\Delta\nu$, and$M$ frequency channels will cost

$$C_{total} = N[JC_a + C_d(K)\Delta\nu/(JL)] + \frac{1}{2}N(N-1)C_x\Delta\nu/(JK), \tag{2.11}$$

where J is the number of analog filter channels, K is the number of digital filter channels, and L is the number of post-correlation DFT channels. The cost can be minimised over $J, K,$ and$L$ subject to the constraint that the total number of channels is $M = JKL$. Thus, the last term will dominate for sufficiently many antennas N, and it is then advantageous to do as much processing before correlation as possible. The division between analog and digital processing depends on the relative costs of those technologies. A given type of digital circuit has a maximum speed of operation, but any bandwidth can be processed by having enough circuits in parallel; this is why the cost is proportional to bandwidth. When the total bandwidth gets large, it becomes advantageous to use some analog filters.

From these considerations alone, it seems that any large, spectroscopic correlator should be built with pre-correlation frequency analysis - a point where the FX scores over the XF.

Another way of looking at it is by the amount of computation involved in the FX and XF algorithms. In its simplest form, the FX correlator computes a length -K transform on non-overlapping segments of $2K$ samples each. Then the time available to compute one transform is $2K\Delta t$, the number of complex multiplies and adds required per transform is $K \log_2 K$, and one such transform is needed for each of N antennas in an array. Thus, the total computation rate for DFTs is$N\Delta\nu \log_2 K$ complex operations per second (COPS), assuming that $\Delta t = 1/(2\Delta\nu)$. A set of K complex cross multipliers and accumulators is needed for each of $N(N-1)/2$ baselines, but each must perform a multiply and add only every $2K$ samples, leading to a computation rate of $N(N-1)\Delta\nu/4$ COPS. So, the total COPS required are,

$$\begin{aligned} c_{FX} &= \Delta\nu[N \log_2 K + N(N-1)/4] \\ &= N\Delta\nu \left[\frac{N-1}{4} + log_2 K\right] \quad , \end{aligned} \tag{2.12}$$

The equivalent computational requirement for a XF correlator is

$$\begin{aligned} c_{XF} &= \Delta\nu K N(N-1)/2 \\ &= N\Delta\nu . \frac{(N-1)K}{2} \end{aligned} \tag{2.13}$$

if each real multiply-add is counted as equivalent to 1/4 of a COPS [1]

Thus, in the FX architecture, for large systems c_{FX} increases only slowly with $K$and$N$. For GMRT ($N = 30, K = 256$), the ratio $c_{XF}/c_{FX} = 243.41$

The Overall Receiver System

Advantages of the XF Correlator Described below are some of the advantages of the XF correlator which may be treated as a cost penalty for the FX correlator.

Invalid Data In an XF correlator, known invalid samples, e.g., due to interference, can be simply omitted from the correlation, while a corresponding operation is not feasible for the FX which requires to take a Fourier transform of the data in real time. Thus, deletion of a bad data sample requires a real time interpolation to be performed for FX, which is a disadvantage for this scheme.

Sensitivity Since the FFT algorithm used in an FX inherently operates on blocks of data, the samples from different blocks never get correlated (including, for example, the last sample in a block and the first sample in the next block which are adjacent samples). Such a segmentation implementation implies that a given cross correlation is derived from fewer pairs of samples in an FX compared to an XF. This results in an effective loss of sensitivity for FX, which can be recovered only by overlapping adjacent blocks of data. There is thus a net increase in computing for FX to achieve the same sensitivity as XF.

Quantisation Relatively small number of bits sampling are often advantageous for high speed correlators as providing the overall efficiency of the system for a given cost. For instance, even a two-bit three-level quantisation gives about 80% efficiency. In such cases, the precise estimate of the quantisation efficiency can be easily estimated for an XF scheme (and hence corrected before the correlation coefficients are used). Corresponding estimates are less general and often very difficult to obtain for an FX correlator.

[1] The arithmetic operations needed in the XF correlator are much simpler than in the FX correlator. This is so since in the FX more bits of precision are required due to the frequency analysis performed before multiplication. In comparison, most operations in a XF correlator are on 1- or 2-bit samples.

Precision The butterfly processors of the FFT hardware must accommodate a wide dynamic range of spectral forms, from continuum to extremely strong and sharp emission lines. This requires the arithmetic operations of FFT to be much more accurate and employ a relatively large number of bits per word compared to an XF operation.

Advantages of the FX Correlator The following specific advantages of FX scheme would imply additional cost penalties on an equivalent XF implementation.

Fringe Rotation Implementing station-based fringe rotation in a XF correlator requires the transmission and correlation of multi-bit rotated samples to avoid loss of sensitivity and spurious correlation at harmonics of the correct station phase rates, and this has proven impractical. The FX scheme, in contrast, must already produce multi-bit samples at the output of the F stage, and it is a minor additional burden to start the transform with fringe-rotated, complex values of sufficient precision to preclude the two problems mentioned.

Closure Errors The FX correlator is less vulnerable to baseline-dependent systematic effects, since almost all of the signal processing is performed on a station basis.

Fractional-Sample Error Fractional sample error arises due the fact that the delay tracking can compensate for delays only in multiples of the sampling period. An error as large as half a sample period can occur as a consequence. Efficient correction for this effect can easily be implemented as an extra step following the last FFT butterfly in the FX correlator. The XF approach, however, is faced with the much larger task of applying the correction on a baseline basis after transforming; in a practical large system only a statistical correction is feasible, which yields the proper amplitude spectrum but sacrifices sensitivity.

Applying the above discussion to the case of the GMRT array, it was found that the overall cost benefits of FX correlator far exceed those of XF, mainly due to the size of the array. Nevertheless, both options were studied, taking into account the various VLSI chips existing then. Finally, the GMRT correlator adopted a FX architecture and used the same basic computing element which was being designed for the VLBA. A specific mode to be used for the GMRT was also incorporated in its design [2]

2.3 The GMRT Correlator - As a Black Box

The discussion in the previous sections described the environment for the Correlator. This leads to its overall specifications. Seen as a *black box* (Fig. 2.5) the GMRT Correlator performs the following

- Processes signals from 30 stations;
- Each station generates 4 streams of 16MHz bandwidth, pertaining to the Upper and Lower sidebands of the right and left circular polarisations. Hence handles in all 120 streams;
- Processes a maximum of 32MHz bandwidth;
- For calculating the four Stokes parameters, either the bandwidth is restricted to 16 MHz or FX is time-multiplexed between the two pairs of Stokes parameters;
- Provides a maximum of 256 Spectral Channels;
- Compensates a maximum of 128μs. delay with an effective step of 1ns;

[2]The 'Indian' MAC mode. Will be described in Chapter 3

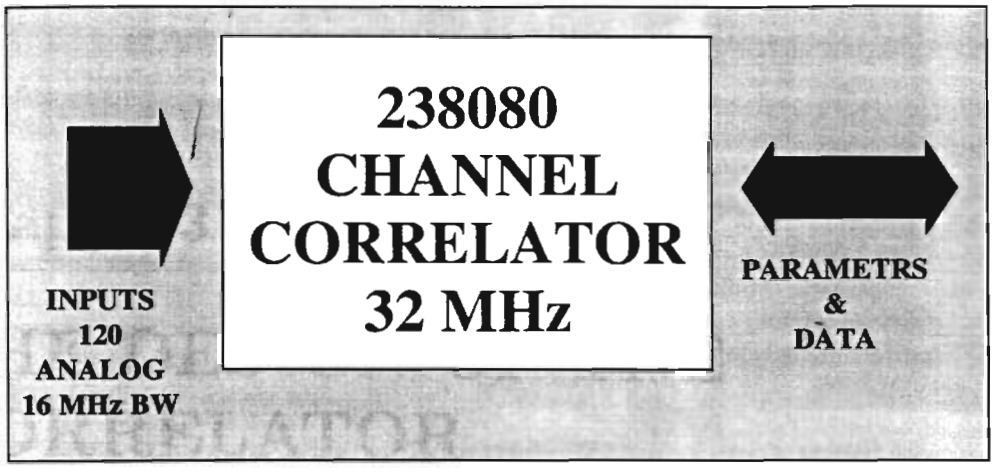

Figure 2.5: The GMRT Correlator : As a "Black Box"

- Provides a nominal integration time of 64ms with 128ms as maximum; integration times as small as 4ms are possible if the number of spectral channels is less than 256;

This chapter discussed the factors which were taken into consideration in formulating the specifications for the GMRT correlator. A comparison of XF and FX algorithms was then made. It was noted that a FX correlator is a judicious choice for the GMRT case. The overall specifications for the GMRT correlator were given. This paves the way for the next logical step - that of realising the specifications. This forms the theme for the following chapter.

Chapter 3

THE DESIGN OF THE CORRELATOR

A plausible impossibility is always preferable to an unconvincing possibility

–Aristotle

This chapter gives the design details of the the GMRT correlator. Detailed specifications of the subsystems are included following which the functional, circuit design and packaging aspects of the subsystems are discussed. The schematics of the various cards given in the appendix complement the discussion provided here.

3.1 A Bird's Eye-View of the Correlator

The design of the correlator was a team effort and the author was one of the core members of the team. Although, the design concepts and the major blocks have been outlined in a series of internal reports, the final system details are not yet published. The following sections present an overview of the GMRT correlator design from the author's own perspective. In addition to the internal reports and the author's own experience, discussions with a large number of colleagues have contributed to forming such a perspective.

Looking top-down on the correlator, one can segregate the system on the basis of functionality. The correlator consists of two major parts as shown in Fig. 3.1.

- The Signal Processing Block: Performs various operations on the signals received
- The Hierarchical Control Block : Conveys user defined parameters and provides certain data necessary for processing of the signals

Further sub-classification of the two major blocks is shown in Fig. 3.2.
The Signal-processing path comprises of four main subsystems.

- The Analog-to-Digital-Conversion (ADC) Subsystem : Digitises the incoming analog signals
- The Delay Subsystem : Compensates for the geometrical and the instrumental delay in the appropriate digitised data streams

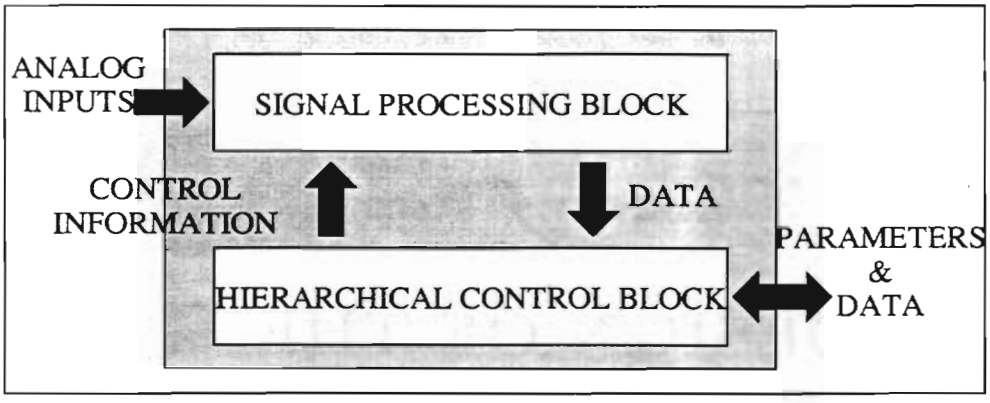

Figure 3.1: The GMRT Correlator: The two main sub-blocks

- The Fast Fourier Transform (FFT) Subsystem : The F part of FX. Performs frequency analysis on the sampled data to provide the spectrum.
- The Multiply and Accumulate (MAC) Subsystem : The X part of FX. Computes the auto and cross correlations between every spectra provided by the FFT Subsystem.

As is evident, these subsystems are identified by the operations they perform on the input signal. Modularity in implementation thus ensures harmony between concept and reality.

The hierarchical control is achieved by distributed elements employing embedded processors at subsystem level, all networked to form the following pyramid.

- Embedded Control Cards: The Delay, FFT and MAC Control Cards form the embedded layer. They are the last in the control chain - the *base* of the control pyramid - and interact directly with the data processing block. They are situated in the subsystem subracks and provide various information to the Delay, FFT and MAC subsystems respectively.
- The *Rack* Control : Handles rack level tasks such as monitoring of temperature, voltage and current. Will also be responsible for remote shutdown of the power-supplies in case of an emergency. Forms a liaison between the Master Control Card and the embedded layer.
- The Master Control Card (MCC) : Performs a variety of tasks. It is the communication link between the Correlator Control Computer and the rest of the correlator. Generates and distributes clocks, synchronisation, and initialisation signals to all control cards.
- The Correlator Control Computer (CCC) : The host of the Correlator control system - the part of the correlator visible to the outside world. This collects information about array configuration, source being tracked, etc., and dynamically evaluates the relevant parameters for various subsystems and communicates these to the MCC.

3.1.1 The Flow of Data

Fig. 3.3 is a block diagram of the correlator showing the path of the signals from the baseband system, through the signal processing chain of the correlator, till they are recorded on tapes for further analysis and processing. The diagram also shows the correlator control system alongside.

The analog signals are input to the ADC subsystem. The signals have a maximum bandwidth of 16 MHz. They are sampled at 32 MHz to comply with the Nyquist sampling requirement. This sampling is constant and does not change even if a signal of less than 16MHz bandwidth is input to the ADC.

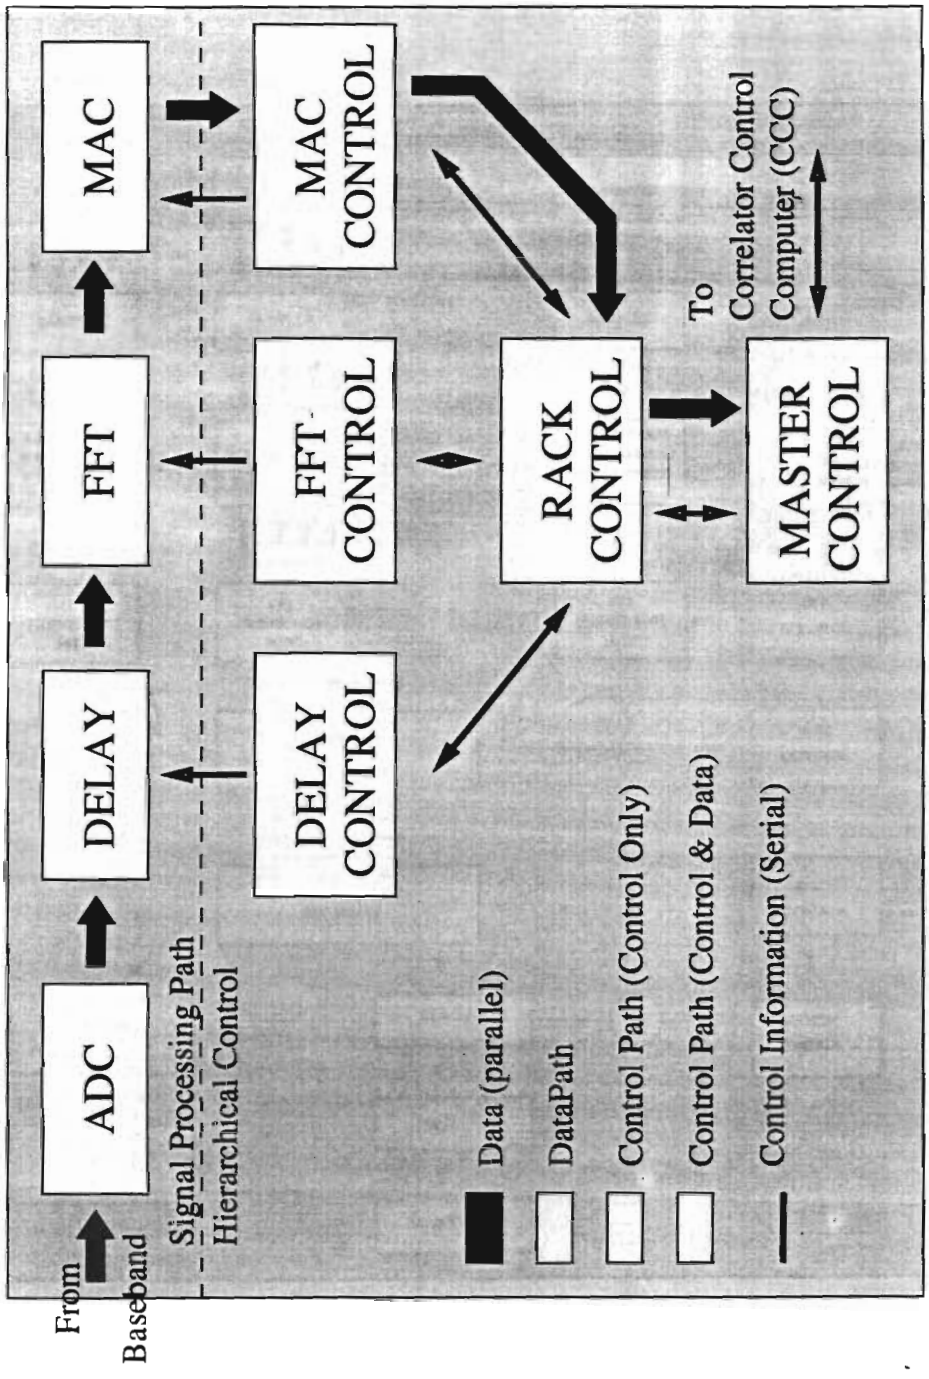

Figure 3.2: The GMRT Correlator: Signal processing and Control Blocks

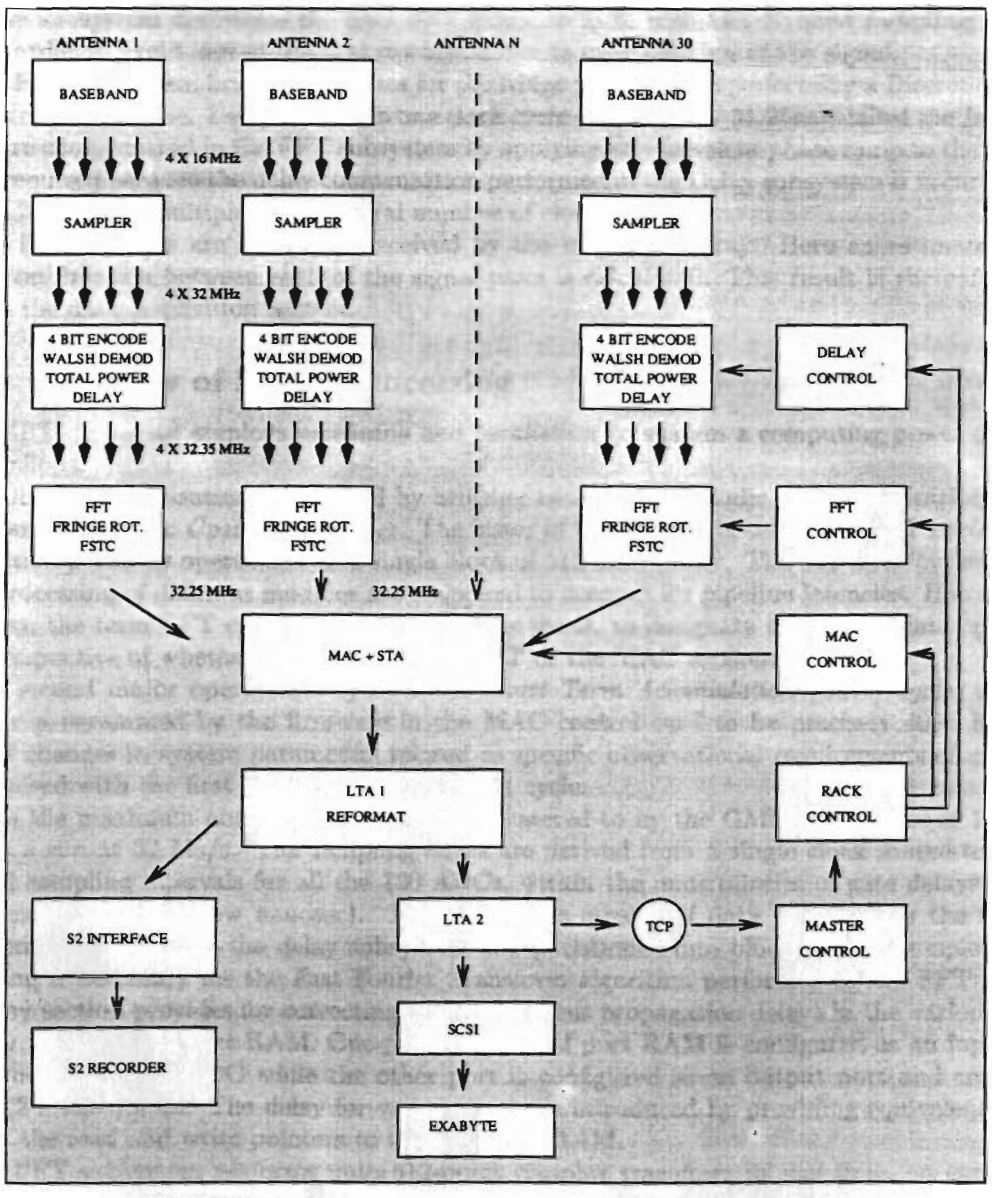

Figure 3.3: The GMRT Correlator: Flow of Data

The digitised signal is then transmitted to the Delay subsystem. Here the signal is first Walsh demodulated and then encoded into a sign-magnitude representation corresponding to the unsigned value which is obtained from the ADC. The signal is hereafter compensated for the delay it has incurred, with respect to a convenient reference antenna or any point in space. The delay subsystem also calculates an estimate of the power content of the signal before it is transmitted to the FFT subsystem for further processing. In case the analog signal was of a bandwidth lower than 16MHz, the delay subsystem decimates the data by a factor so as to maintain Nyquist sampling. This is done in order to avoid loss of spectral resolution due to oversampling of the signal.

The FFT subsystem first compensates for the fringe phase before performing a Discrete Fourier Transform (DFT) on it. Delays less than one clock cycle (1/32MHz = 31.25ns), called the *fractional delays* are compensated in the FFT subsystem by applying an equivalent phase ramp to the output. This is required because the delay compensation performed in the Delay subsystem is accurate only upto 31.25ns or its multiples (an integral number of clock cycles).

The FFT outputs are thereafter received by the multiplier array. Here an estimate of the correlation function between each of the signal pairs is calculated. This result is thereafter read out into the data acquisition system.

3.1.2 Overview of Signal Processing

The GMRT correlator employs *pipelining* and *parallelism* to sustain a computing power of about 400 Gigaflops.

Pipeline synchronisation is achieved by utilising two sets of periodic resynchronisation pulses, which define the basic *Operational cycles*. The faster of these cycles is called the *FFT cycle*, where each section performs operations on a single block of 512 data values. The 512 clock cycles for the signal processing of different modules are staggered to account for pipeline latencies. However, for simplicity, the term FFT cycle will be used in this thesis, to designate the appropriate operating cycle irrespective of whether it is in the delay, FFT or the MAC section.

The second major operational cycle is the *Short Term Accumulation (STA) cycle*, which is generally programmed by the firmware in the MAC control card to be precisely 4096 FFT cycles. All changes in system parameters related to specific observational requirements of users are synchronised with the first clock edge of each STA cycle.

Since the maximum analog signal bandwidth catered to by the GMRT correlator is 16 MHz, the ADCs run at 32 Ms/s. The sampling edges are derived from a single clock source to ensure identical sampling intervals for all the 120 ADCs, within the uncertainties of gate delays of Fast TTL logic (typically a few nanosec). The continuous stream of data provided by the sampler subsystem is acquired by the delay subsystem and partitioned into blocks of 512 samples. Such a blocking is necessary for the Fast Fourier Transform algorithm performed in the FFT section. The delay section provides for correcting for the different propagation delays in the various signal paths by using a dual port RAM. One port of the dual port RAM is configured as an input port and connected to the ADC while the other port is configured as an output port and connected to the FFT subsystem. The delay for various paths is introduced by providing equivalent offsets between the read and write pointers to the dual port RAM.

The FFT subsystem performs upto 512-point complex transform in real time, on each block of 512 data values received from the delay subsystem. Thus the subsystem has 120 FFT engines, each performing a 512 point transform in 512 clock cycles (about 16 μs). The voltage spectra thus obtained are received by the MAC subsystem which multiplies the corresponding spectral channels for every possible pair of antennas, for each polarisation, in each clock cycle. The on-chip accumulators are used by the MAC subsystem to accumulate the products so formed in each FFT cycle. The internal accumulators are organised as two banks of 256 complex words, which alternate between internal and external access for read-out in each STA cycle. The data read out from the MAC system is sent for further processing to the networked file server by the data acquisition system.

In order to sustain the real time operations described above, it is necessary to provide for certain critical operational overheads. The maximum overhead is in the FFT subsystem, which requires four clock cycles for performing its internal operations in each FFT cycle. In order to allow for this, the FFT cycle is stretched to 516 clock cycles in the entire system, where 512 clock cycles are used for signal processing operations while the remaining 4 cycles - called the dead time - are used for the relevant overheads. In order to avoid loss of data during the dead time, the FX system is clocked at 32.25 MHz while the ADCs are clocked at 32 MHz, so that there are precisely 512 samples in the 516 clock cycles of FX. As a consequence, the input port of the delay RAM is clocked at 32 MHz, while the output port is clocked at 32.25 MHz.

It may be noted that the FFT subsystem can only provide a maximum of 256 spectral channels, which is the main limiting factor for the spectral resolution. However, an option is provided where higher spectral resolution can be achieved by processing a smaller bandwidth; e.g., for a bandwidth of 1 MHz, the 256 spectral channels can be separated by 4 kHz. This requires that the input data be sampled at the Nyquist rate appropriate to the bandwidth, e.g., at 2 Ms/s for a 1 MHz bandwidth. In order to avoid the complexity of maintaining synchronisation at many different clock frequencies, the system is always run at 32 MHz (i.e., sampler at 32 MHz and FX at 32.25 MHz). The delay system achieves any required sampling of data by filling the RAM at the appropriate rate. But this introduces a difference in the throughput rates of the input and output section of the delay RAM. This is compensated by introducing an overlap between adjacent blocks of 512 data transmitted to the FFT system.

Parallelism is used most evidently in the MAC subsystem where all the multipliers correlate the data in parallel. Even before the MAC, identical ADC- Delay-FFT Chains process data from the antennas in parallel. The difference being that before the MAC subsystem, each of these chains are totally independent of each other whereas in the MAC the multipliers, spread over some 64 different PCBs operate as one subsystem.

3.1.3 Packaging Considerations

All the printed circuit boards (PCBs) for the GMRT Correlator are based on the DIN-41612 Euro rack standards. The racks used for the system are 24U high and a total of 10 racks comprise the entire system. (Fig. 3.50). Of these 10, six racks house the FFT and MAC subsystems (these racks are therefore called the FX racks), two racks are required for the Delay subsystem and the ADC subsystem requires two racks. The Fig. 3.50 gives the rack layout of these 10 racks. The layout shows only 9 racks (excluding the GAC rack) since the ADC power-supply rack (described later) is omitted.

Each subsystem is packaged in either a 3U or a 6U subrack. PCB's of different subsystems use different subracks. The PCBs in a subrack are interconnected through a backplane which provides power supply and the control signal bus to all the constituent PCBs. The Delay and FX racks are powered by a +5V,200A/ -5V,50A (approx. 1KW) Switched-Mode-Power-Supply (SMPS). The ADC system is powered by +10V,15A/-10V,15A linear power supply. Copper bus-bars, mounted vertically to span the entire height of the rack, connect the power supply to the various subracks. This connection is made on the horizontal bus-bars, provided on each subrack backpanel, to establish multi-point connections to its power and ground planes.

Fan-trays are used in all the racks to provide forced cooling to the subracks. All fan-trays have six 90 CFM (cubic-foot per minute) fans mounted on a 1U rack mountable tray (Fig. 3.4). The positions of the fan-trays is shown in the rack layout of the subsystems, later in the discussion. The top cover of every rack is also provided with four fans - three 90 CFM fans and one 250 CFM fan (Fig. 3.5) - to maintain the flow of air through the rack.

The ADC, FFT and MAC subsystem PCBs have a 280mm card depth, while the Delay subsystem PCBs have a 260mm card depth. Two sets of design parameters were followed for the design of all the PCBs constituting the correlator. The PCBs with lower component density followed a 12 mil-track-width-12 mil-spacing-30/36 mil-via-diameter design rule where as the higher density

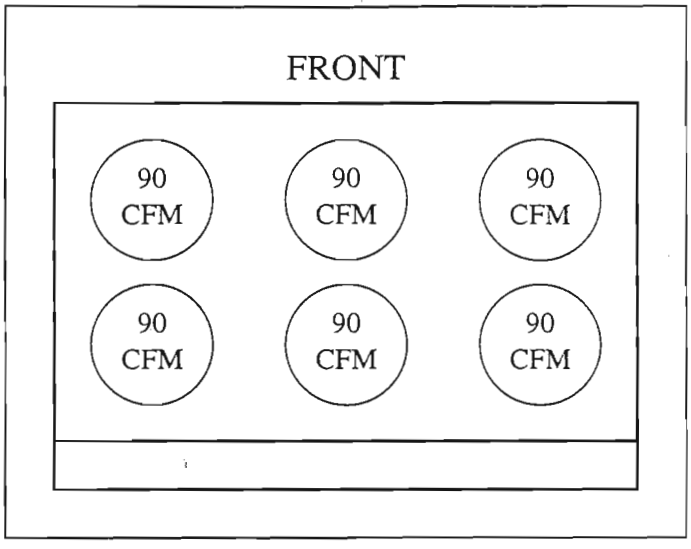

Figure 3.4: Layout of fans in the fan-tray

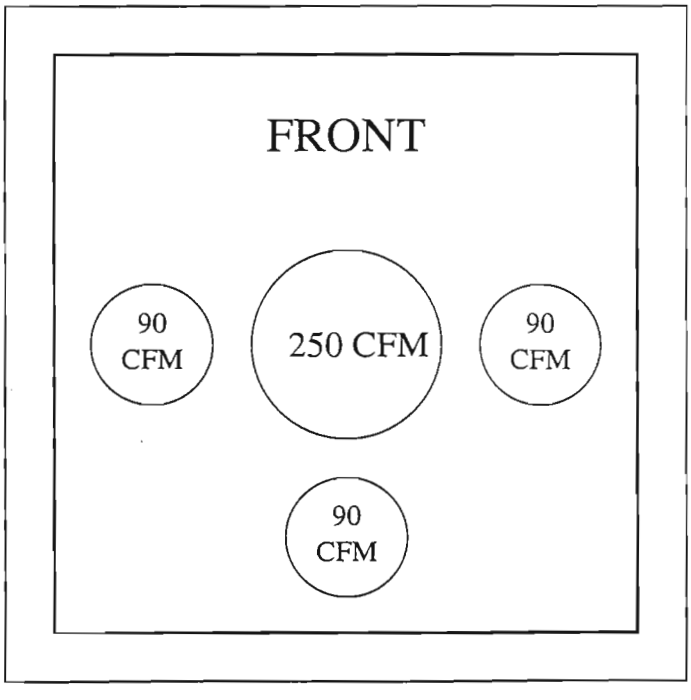

Figure 3.5: Layout of fans in the rack top

Subsystem	*PCB Type*	*Connector Type*	*Input/Output*
ADC	ADC Card	PCB mount Rt.-angle SMA	I
		64-pin Rt.-angle Euro Plug	O
	ADC Backpanel	64-pin St. Euro Header	
DELAY	DPC Card	64-pin Rt.-angle Euro Header	I
		96-pin Rt.-angle Euro Plug	O
	Delay Card	64-pin Rt.-angle Euro Plug	I
		26-pin Rt.-angle FRC	O
	Delay Control	96-pin Rt.-angle Euro Plug	I/O
	Delay Backpanel	64-pin St. Euro Header	
		96-pin St. Euro Plug	
FFT	Main Board	26-pin Rt.-angle FRC	I
		128-pin Rt.-angle Euro Header	I/O
		100-pin Surface Mount Plug	I/O
	Pipeline Board	100-pin Surface Mount Header	I/O
	FFT Control	128-pin Rt.-angle Euro Header	I/O
	FFT Backpanel	128-pin St. Plug	
MAC	MAC Card	160-pin Rt.-angle Euro Header	I/O
	MAC Control	96-pin Rt.-angle Euro Plug	I/O
	MAC Backpanel	160-pin St. Euro Plug	
		96-pin St. Euro Header	
		50-pin AMPMODU Surface Mount Header	
		50-pin AMPMODU Cable Connector	
Dist.	Dist. Card	50-pin St. FRC	O
		50-pin Rt.-angle FRC	O
		96-pin Rt.-angle Euro Plug	I/O
	Dist. Backpanel	50-pin St. FRC	I
		96-pin St. Euro Header	
		50-pin AMPMODU Surface Mount Header	O
		50-pin AMPMODU Cable Connector	

Table 3.1: Connectors used in the GMRT Correlator

boards used a 8 mil- track-width-8 mil-spacing-16 mil-via-diameter design rule.

The analog signals are transmitted from the baseband to the ADC subsystem through a RG-223 shielded co-axial cable. The digitally sampled data are transmitted using $RoundTwist - N - Flat^{TM}$ [1] jacketed shielded cables. These cables carry ECL differential signals to ensure high speed balanced data transmission between the ADC, Delay, FFT and MAC subsystems.

A variety of connectors ranging from the PCB mountable SMA to the 160-pin 5 row EURO to the high density 50mil X 50 mil grid AMPMODU connectors are used in the PCBs ($\sim$ 320 in all) which comprise the GMRT correlator. These are intended to provide a compact solution to the complex packaging problem posed by the GMRT Correlator. Table 3.1 lists all the connectors

[1]This is the trade name for these cables manufactured by Amphenol Corp. under their $SPECTRA - STRIP^R$ range of cable products. The cable is so named since it combines into a single construction, the handling and shielding features of round cable and the mass termination capabilities of flat cable. The cable is composed of twisted pairs with alternating flat sections. Every 20" of cable length has a $2\frac{1}{2}$" flat section to facilitate crimping of IDC connectors. The cable has a foil/braid shield around it and this entire assembly is covered by a round PVC jacket

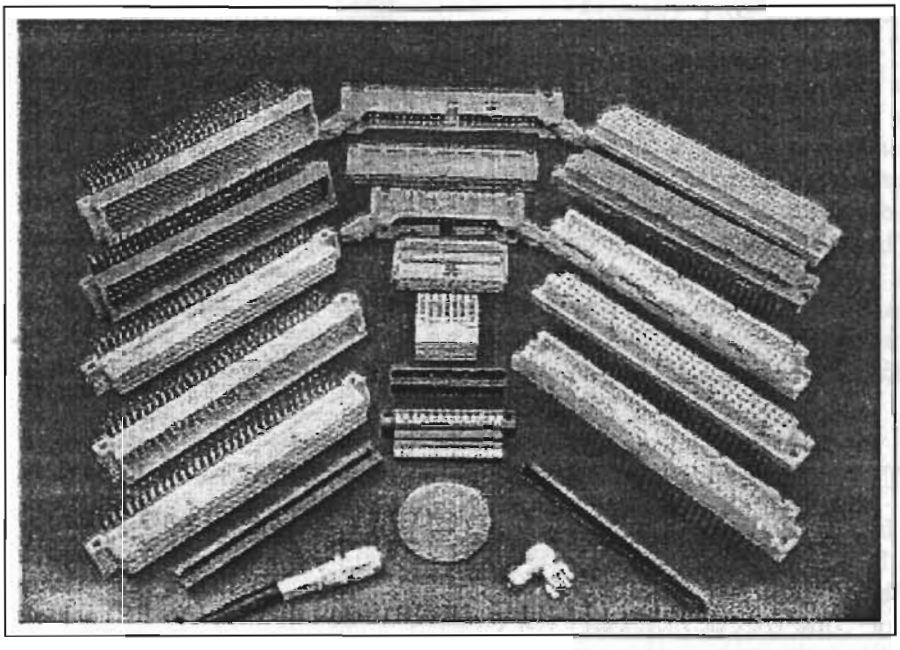

Figure 3.6: Connectors used in the GMRT Correlator

used in the various cards in the correlator system. Fig. 3.6 shows all these connectors.

3.2 The Specifications of the Subsystems

3.2.1 The ADC Subsystem

Sampling Frequency : 32 Ms/s.
Number of Bits : 6 bits.
Thresholds : Uniform, symmetric about mean.
Level-step : 0.3352*rms for an input signal of 0dBm power and 4-bit quantisation.
Output Format : 6 bit ECL.

3.2.2 The Delay Subsystem

Delay Step : 32ns*(BW/16MHz).
Maximum Delay : 128 μs.
Input sampling can be set to any desired value between 32 Ms/s and 125 Ks/s in binary steps, depending on the bandwidth of the signal.
Input : 6 bit ECL differential.
Output : 4 Bit ECL differential.
Buffering Options : Data are sent to FFT engines at a constant rate of 32.25 MHz irrespective of the input sampling rate. The data is sent to the FFT card in blocks of 512 words followed by a dead time of 4 clock cycles. Overlapping of adjacent blocks will occur when input data are sampled at 16 MHz or less.

3.2.3 The FFT Subsystem

Input data Rate : 32.25 MS/s (real numbers).
Input data window : width = FFT length.
Input format : (7,7,2). [2]
Window format : (5,0,4).
Fringe Resolution : 0.7 deg step for fringe angle.
Spectral Channels : 16, 128 or 256 channels.
Input data : 4-bit ECL differential.
Output data : 12 bit Complex, every alternate cycle.
Output Format :(4,4,4).
FFT Clocking Rate : 32.25 MHz.

3.2.4 The MAC Subsystem

Minimum Integration Time : 4 ms (256 FFT Cycles).
Maximum Integration time : 128 ms (8192 FFT Cycles).
Spectral Channels : 128 per baseline.
Modes : Non-polar (RR or LL), Full Polar (RR-RL or LR-LL) and 'Indian' Polar (RR & LL).
IF bandwidth Supported :

- 16 MHz : Full Polarisation (RR-RL or LR-LL mode.
- 32 MHz : Indian Polar (RR-LL) mode only.

Output data: 36-bit complex - one word output per card in each FFT cycle.
Output Format : (15,15,6), The real and imaginary mantissa are in 1's complement, the exponent in 2's complement.

3.3 The ADC Subsystem

It is the front-end of the correlator system. It performs the task of digitising the analog signals. There are no user definable parameters for this subsystem. Hence, there is no control system for the ADC subsystem.

3.3.1 The ADC Subsystem: Circuit Design

The ADC sub-system consists of the ADC cards (Fig. 3.7) and the ADC Backpanel.

1. **ADC Cards**

   This is the active card of the sub-system. Its internal modules are summarised below :

   The Amplifier Section Each ADC card accepts two inputs from the baseband system at 0 dBm power level. These zero-mean signals are AC coupled to the amplifier circuit realised using AD9617 op-amps. Configured as a non-inverting amplifier, it has a gain of 1.5. This value was chosen such that a 6σ signal would occupy the ± 1V (2V peak-to-peak) voltage range of the ADC. The AD9617 provides a low offset gain and is compensated internally. This reduces the circuitry needed to ensure a stable operation.

[2] This is a shorthand notation for the representation of complex floating point numbers used in the FX asic. The three numbers correspond, respectively, to the number of bits in real mantissa, imaginary mantissa and the common exponent. Refer section on FX Asic for further details

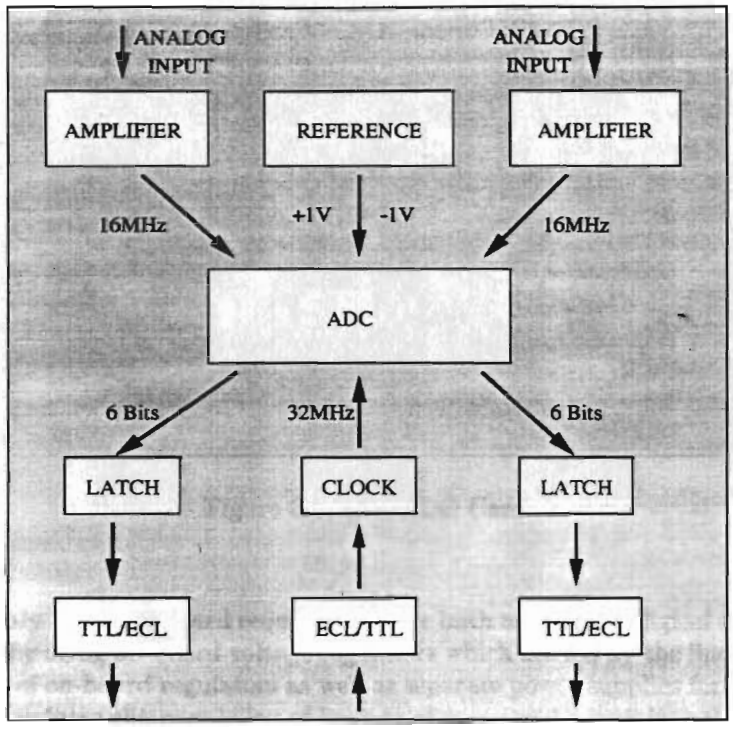

Figure 3.7: The ADC Card : Block Diagram

Reference Generation Circuit The external reference generation circuitry was designed around the Analog Devices two terminal reference AD589. The AD589 serves as a low-cost precision reference to a circuit which uses the AD708 dual op-amp with the matched pair transistors LM3904 and LM3906. The transistors drive the reference ladder of the ADC. The loop thus established ensures stable reference voltage in spite of the changes in the ladder impedance from one ADC to another.

The ADC Chip The Analog Devices AD9058, a dual 8-bit 50MHz Flash ADC forms the core of the ADC card.

Although the AD9058 provides a highly stable and precise internal reference to drive its reference ladder, it is configured to be used with an external reference. This was done since the internal reference requires the input to be uni-polar. A bipolar external reference was used instead to reduce the possibility of introducing DC offsets in the signal.

Protection Circuitry A protection circuit was needed to restrict the maximum input voltage to the ADC to the rated specified of the AD9058. A simple circuit was devised using PIN diodes to restrict voltages to within ± 1.2 V.

Digital Section This section contains latches (74F574) for the ADC output, TTL-to-ECL (MC10124) converters for data and ECL-to-TTL (MC10125) converter for the clock.

The clock is converted to TTL levels and given to the ADC chip. An inverted version of this clock is given to the latches which receive the digitised data from the ADC. The inversion of the clock provides sufficient time for the ADC output to stabilise at the latch input. The latch outputs are given to the TTL-to-ECL converters which then transmit the data to the Delay subsystem.

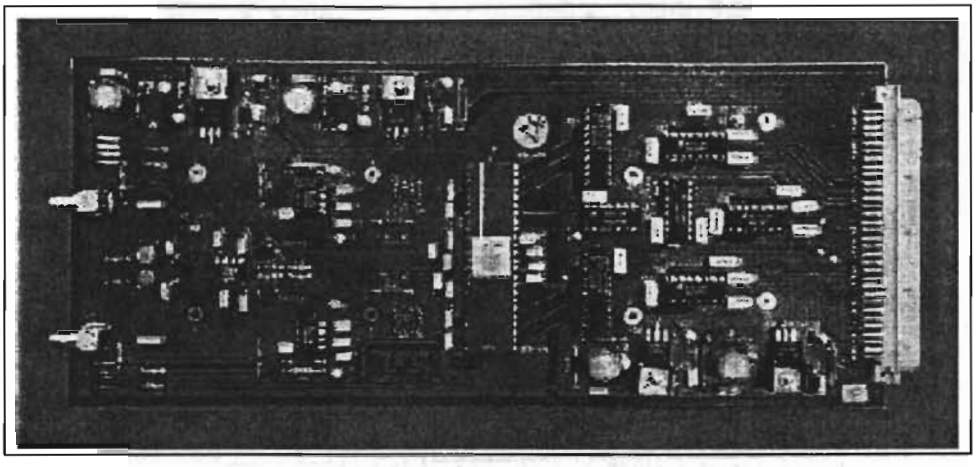

Figure 3.8: The ADC Card

Power Supply The ADC card requires ±5V for both analog and digital sections. These are derived by using on-board voltage regulators which are fed by the linear power supplies. The use of on-board regulators as well as separate power supplies for analog and digital sections reduces the possibility of leakage of correlated power into the analog signal.

2. **The Backpanel**

The ADC cards plug-in to a passive backpanel, to constitute a sub-rack of the ADC sub-system. It distributes power and clock to the ADC cards. The digitised data from the ADC cards is transmitted to the Delay sub-system by connectors plugged directly into the backpanel.

3.3.2 The ADC Subsystem : Packaging

The ADC card (Fig. 3.8) is a 6 layer PCB, 110mm by 260mm in size, 1.6mm thick with 35μ copper plating on all layers. The layer assignments are shown in Fig. 3.9 which depicts schematically, the construction of the ADC card. The PCB is designed with a 12mil track-width and 12mil track-to-track spacing design rule. The via size used is 36mil. Two right-angled PCB mount SMA connectors couple the analog inputs to the PCB. A single 64-pin EURO connector is used for the backpanel. The analog and digital sections of the PCB are isolated from each other to ensure separate power-ground planes for the analog and digital sections. The two sections are coupled only at the ADC power-supply input pins.

The Backpanel (Fig. 3.10) is a 3.2mm, 6 layer PCB. Its construction is shown in Fig. 3.11. Each ADC card is an independent unit and there are no buses running on the backpanel which interconnect the ADC cards. Two *Berg* pins are provided with every ADC card for the ECL clock connections. The connections are wire-wrapped on the pins from a clock distribution card which is mounted adjacent to the backpanel in the rack.

A total of 60 cards are required for the entire system. Each backpanel supports 12 cards, with a inter-card spacing of 7T (= 1.4 inch). Thus, a total of 5 sub-racks make up for the entire requirement. A total of 2 24U racks are used for the ADC system. One rack houses all the 5 subracks with fan-trays for forced cooling while the other rack houses the power supplies required by the 5 subracks. A schematic front view of the subrack is shown in Fig. 3.12. The arrangement of the two racks is shown in Fig. 3.13.

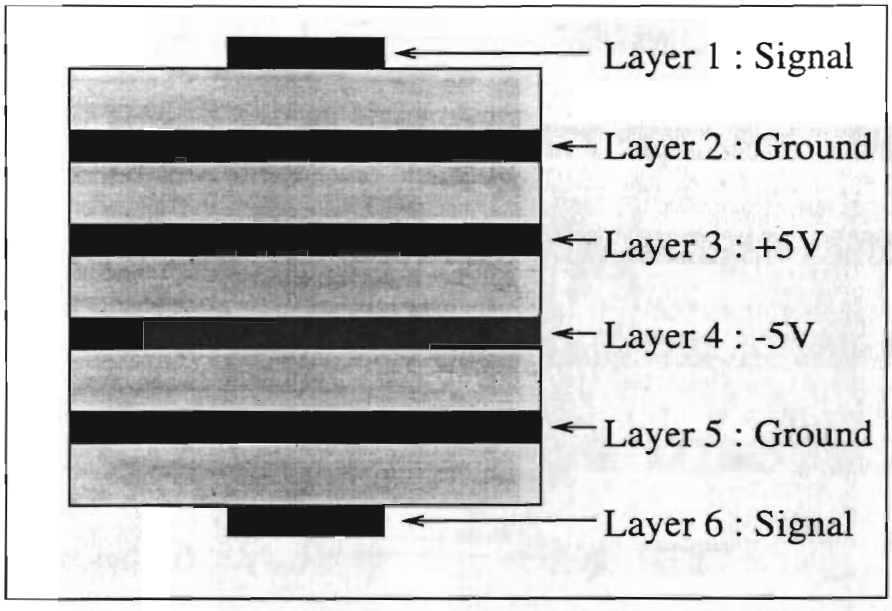

Figure 3.9: The ADC Card Layer Assignments

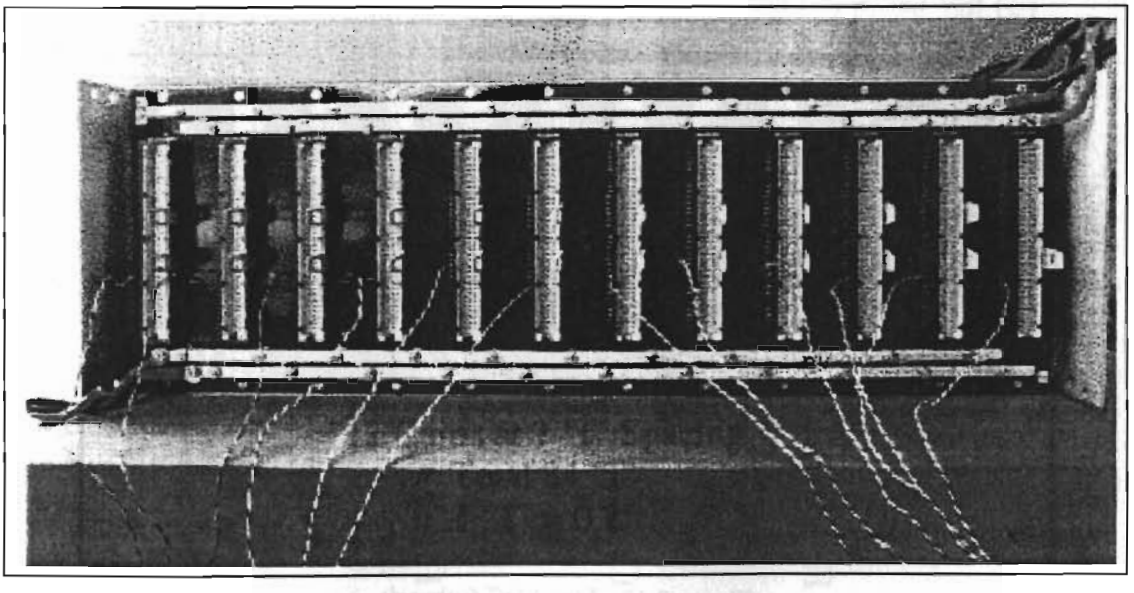

Figure 3.10: The ADC Backpanel

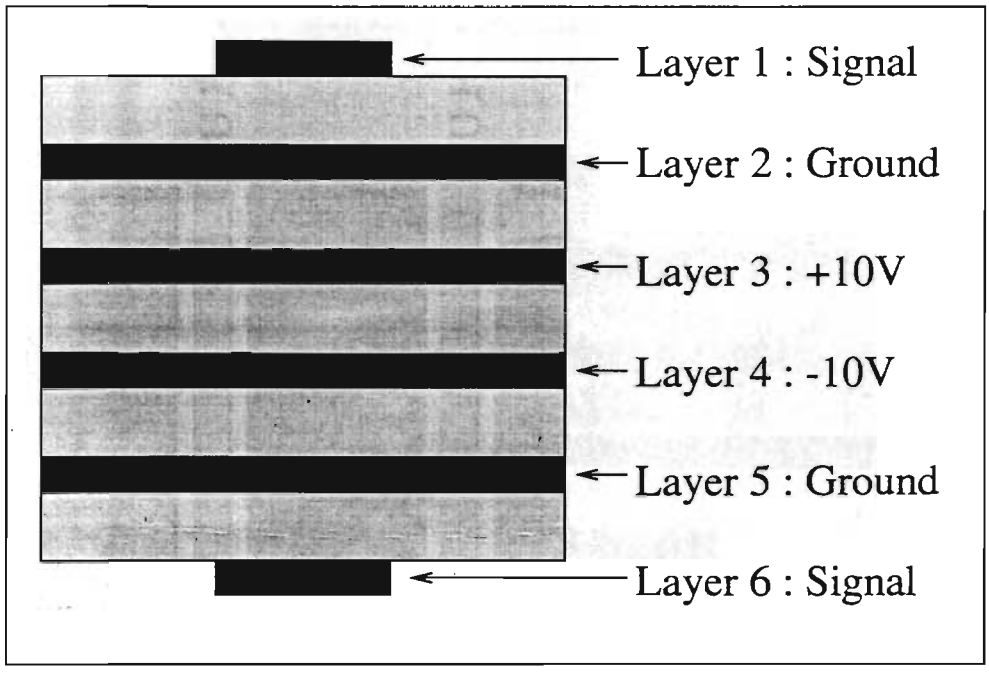

Figure 3.11: The ADC Backpanel : Layer Assignments

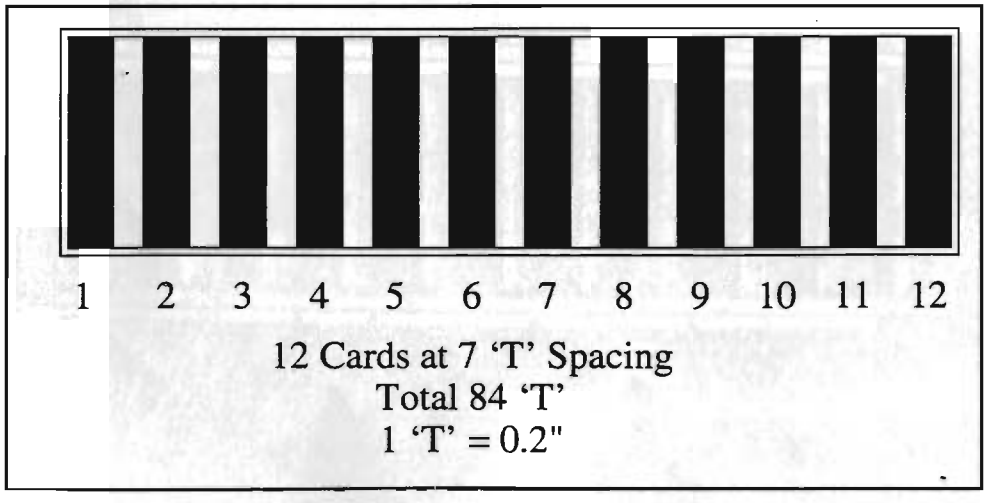

Figure 3.12: The ADC Subrack

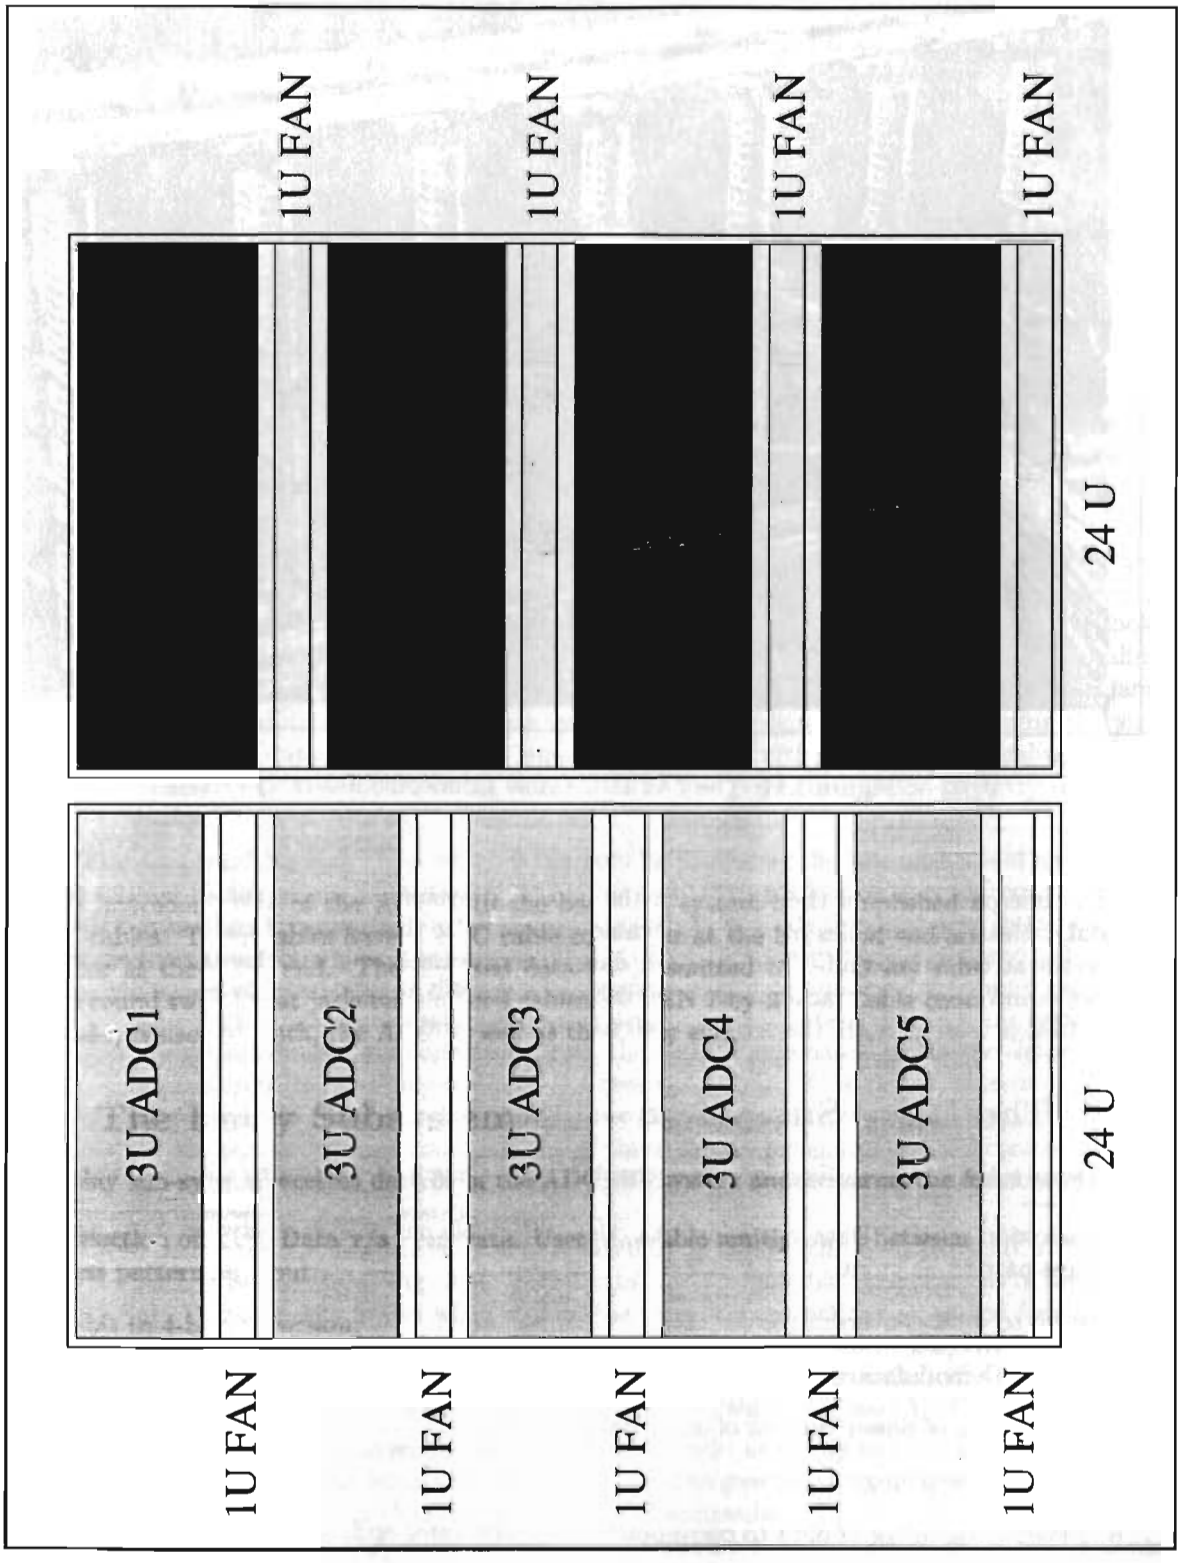

Figure 3.13: The ADC Racks

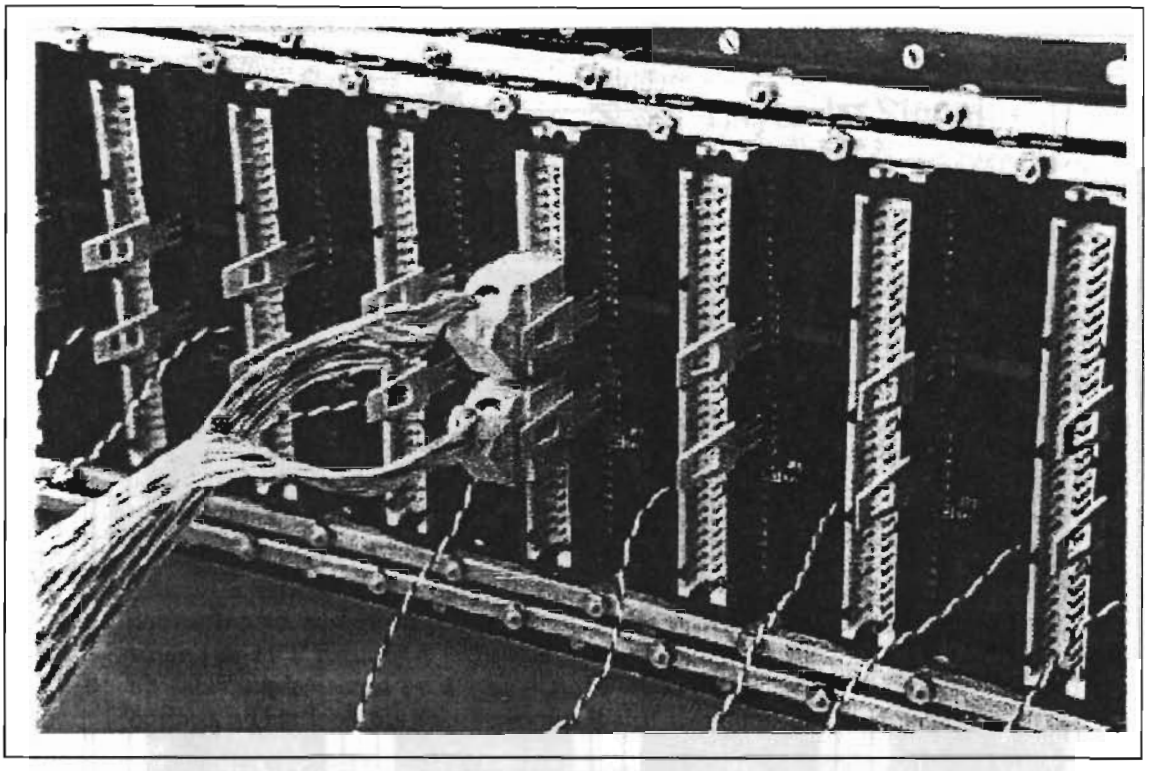

Figure 3.14: The 7X2 IDC Cable Connector Assembly

The interconnection of the ADC with the baseband system is accomplished by 120 RG-223 coaxial cables. These cables have a TNC cable connector at the baseband end and a SMA cable connector at the ADC end. The digitised data is transmitted to the Delay subsystem via 120 14-core round twist-n-flat jacketed shielded cables. A OEN 7-by-2 IDC Cable connector assembly (Fig. 3.14) is used at both, the ADC as well as the Delay end.

3.4 The Delay Subsystem

The Delay sub-system receives data from the ADC sub-system and performs the following tasks:

1. Selection of ADC Data v/s Test Data. User selectable multiplexing between ADC data and test pattern as input.
2. 6-bit to 4-bit conversion;
3. Walsh Demodulation;
4. Estimation of power content of input data;
5. Channel Switching;
6. Effective sampling of data to correspond to Nyquist rates and,
7. Compensation of delay

3.4.1 The Delay Subsystem: Circuit Design

The Delay Sub-system comprises of the DPC, Delay (Fig. 3.15) and Backpanel cards:

1. **The Pre-processor Card (DPC)**

   The DPC comprises of the following sections:

   Input Section This section is consists of ECL-to-TTL (MC10125) converters followed by a multiplexer- cum-latches (74F399). The ECL input from the ADC is converted to TTL levels and then input to the multiplexer. The multiplexer, which is controlled by the Delay control card, selects and latches either the ADC input or a test pattern according to the user's choice. The test pattern is generated by the Delay control using a 2048-word look-up table, which can be defined by the user during initialisation.

   This test input is used for correlator self-tests. This tests the entire chain of the correlator except the ADC. It is thus an extremely important feature and can be used as and when required to ascertain the serviceability of various subsystems. During normal observations, one typically moves the antennas to a calibration source once or twice an hour, which provides time gaps usable for self-tests.

   The Encoder The encoding refers to the 6-bit to 4-bit conversion including re-scaling and conversion to a sign-magnitude form. The ADC generates a 6-bit unsigned output while the FFT requires a 4-bit signed input. This conversion is performed by a look-up table implemented by a synchronous ROM (CY7C245A) with two possible scaling options. This look-up also performs Walsh demodulation and squaring of the data. Walsh modulation being a phase switching can be demodulated by flipping the sign bit of the data as and when a signal synchronised with the Walsh switching pattern is asserted. The square of the data is fed to the power estimation circuitry described below.

   The Channel Switch This switch is required for re-routing the four channels of an antenna to FFT cards in the two sidebands, according to the polarisation or dual frequency mode desired by the user. For instance, in the polarisation mode the same data (obtained from the signal of any one of the two sidebands of the receiver system) is required to be input to the two side-bands of the correlator system. One side-band then computes two of the four Stokes parameters while the other side-band computes the remaining two. When not operating in the polarisation mode, the two FX side-bands of the correlator process signals from the two side-bands of the receiver system. This switch, therefore, receives all four inputs (both side-bands of the two polarisations) from an antenna and produces at the output various combinations of these inputs depending on the control pattern given by the Delay control card. Fig. 3.16 gives the input and output combinations for various values of the control value.

   This switch too, like the encoder is realised using two synchronous ROMs (CY7C245A). As can be seen from Fig. 3.16, only a subset of the possible permutations of the inputs are realised. The states where one of the inputs is reproduced at all the four outputs - namely, the outputs for control values 0, 1, 2 and 3 by which BL, BR, AL and AR respectively, are reproduced at the output - are useful in diagnostics to feed digital copies of an analog input or a test pattern to multiple Delay, FFT and MAC subsystems.

   Power Estimation Conversion to 4-bits in the encoder uses only half the width of the ROM (CY7C245A). The remaining 4-bits were used to generate the square of the same number. This square is then input to an adder-accumulator (adder-74F83, latch- 74F175, accumulator-74F579) combination to implement a realisation of Σx_i^2 where x_i represents a data point in the voltage time series. The estimate of the power content of the signal thus generated is used to calibrate the gain of the receiver chain. It also has a potential usage for interference detection in the future.

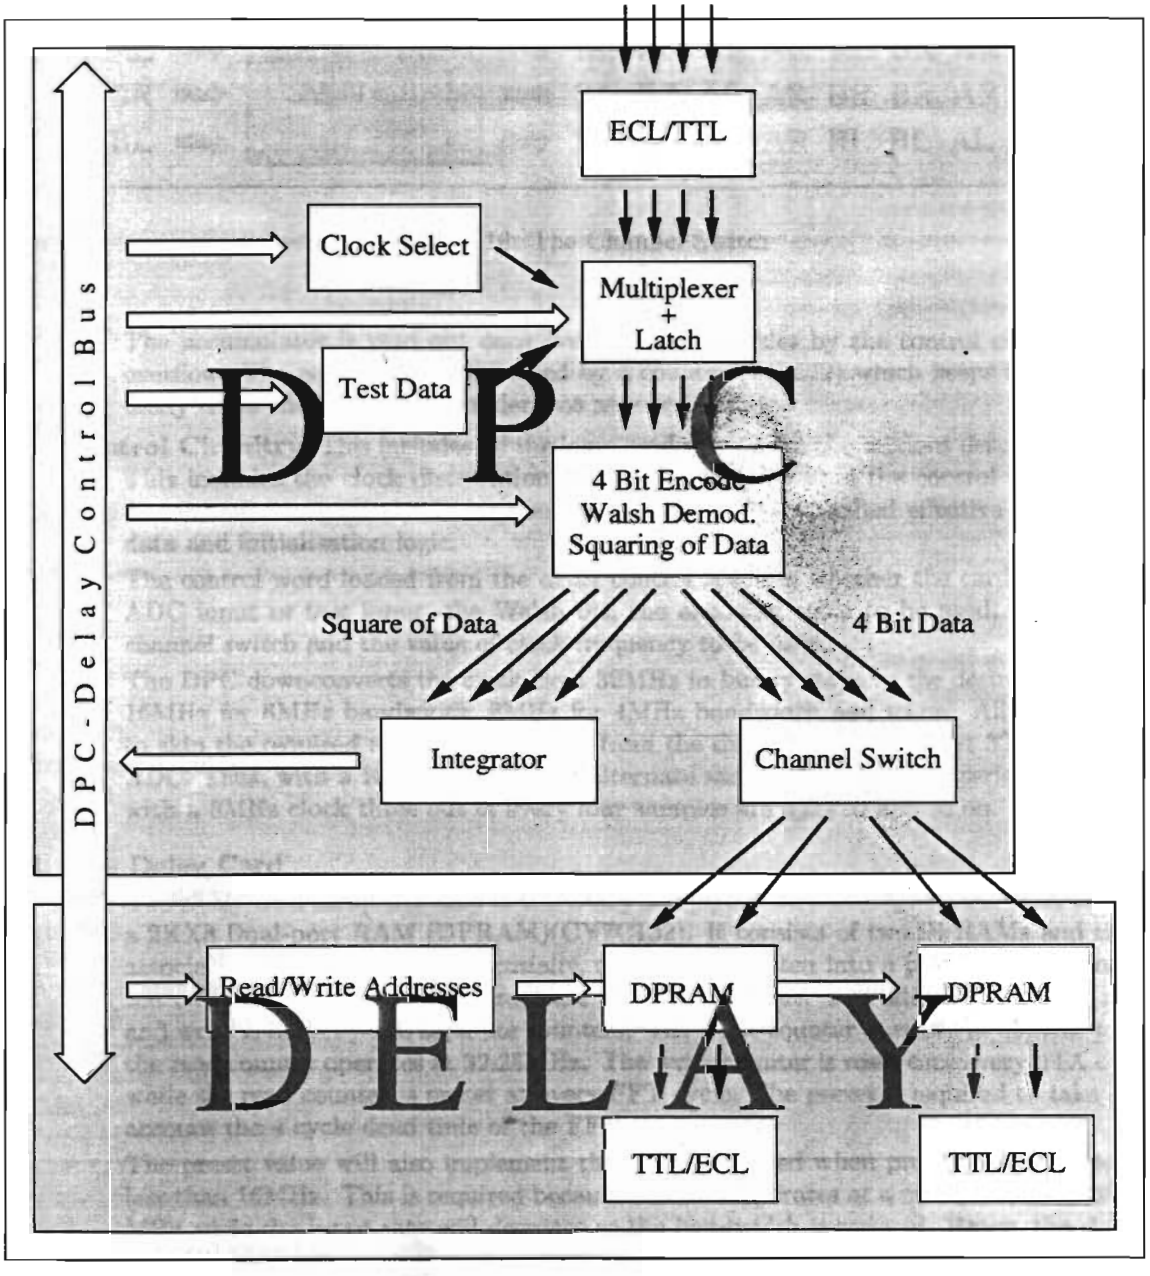

Figure 3.15: The DPC and Delay Cards : Block Diagram

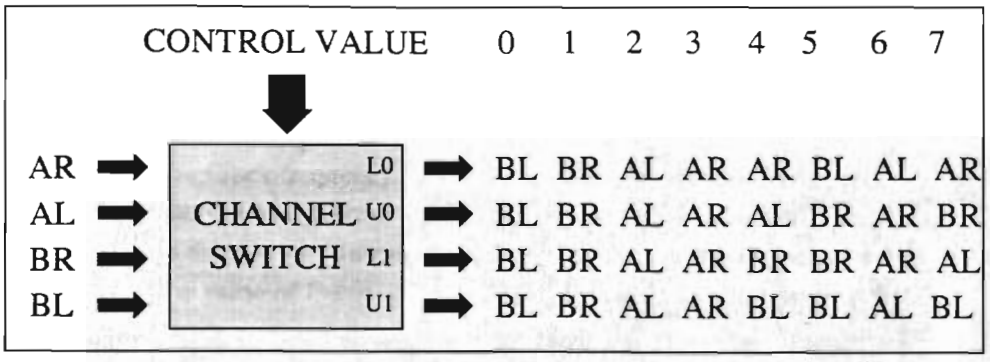

Figure 3.16: The Channel Switch

The accumulator is read out once every few FFT cycles by the control card to avoid overflow. The accumulator is realised by a counter (74F579) which keeps track of how many times the carry of the adder was asserted.

Control Circuitry This includes all the logic used apart from the sections described above. This includes the clock distribution logic, latches which hold the control word loaded from the Delay control, clock generation logic for user specified effective sampling of data and initialisation logic.

The control word loaded from the delay control specifies whether the card is receiving ADC input or test input, the Walsh bit, the encoding table to be used, the value of channel switch and the value of clock frequency to be used.

The DPC downconverts the clock from 32MHz in binary steps to the desired value e.g. 16MHz for 8MHz bandwidth, 8MHz for 4MHz bandwidth and so on. All this does is to skip the required number of samples from the time series sampled at 32MHz by the ADC. Thus, with a 16MHz clock every alternate sample from the time series is ignored, with a 8MHz clock three out of every four samples are ignored and so on.

2. The Delay Card

The Delay card comprises only of the Delay compensation section, which is built around a 2KX8 Dual-port RAM (DPRAM)(CY7C132). It consists of two DPRAMs and there associated control circuits. Essentially, the data is written into a particular location of the DPRAM and after the appropriate delay, is read out. The addressing for the read and write is done by two separate counters. The write counter operates at 32MHz while the read counter operates at 32.25MHz. The write counter is reset once every STA cycle while the read counter is preset at every FFT cycle. The preset is required to take into account the 4 cycle dead time of the FFT sub-system.

The preset value will also implement the overlap needed when processing bandwidths less than 16MHz. This is required because the FFT operates at a constant rate of 32.25 MHz while the input rate will decrease as the bandwidth is reduced. Hence, the slower occurrence of new data points is compensated for, by re-using the old data points.

3. The Delay Control Card

The Delay subsystem has one embedded controller per subrack in the form of the Delay Control card. The control card is designed around the ADSP2105 fixed-point DSP and uses the 20Mbps *INMOS* transputer links for communicating with the rack control. It handles tasks such as communication with the Master Control (through the Rack Control). It thereby ascertains various parameters to be configured in the Delay subsystem. These parameters can be classified to be of two types:

(a) Parameters fixed for an Observation session
   i. Choice of input - ADC or test data
   ii. The test pattern
   iii. The encoding table
   iv. The Clock frequency
   v. The value of channel switch

(b) Parameters which change dynamically during the session
   i. The value of Delay

The parameters which are fixed for a session are loaded at the time of initialisation of the system. The parameters which change, like the delay values, are communicated once every STA cycle to the control card. The delay values are then loaded into the appropriate delay card and are activated precisely at the start of a new STA cycle. The precision in timing is required to avoid discontinuities in the phase of the signal which will occur if delays are changed within a STA cycle.

The control card is also expected to read out the power estimation data fast enough in order to avoid an overflow occurring in the accumulators within the DPC card. This data is accumulated further in the control card, if need be, before sending it up the control chain to the CCC via the MCC for necessary interpretations.

The control card also receives a user specified test pattern, for performing diagnostic tests on the correlator system, which it inputs to the DPC every clock cycle. This pattern is the same for all DPCs in the subrack.

4. The Backpanel

The backpanel for the Delay subsystem accommodates the DPC and Delay cards alongwith the Delay control card. It provides power and system clocks to these cards and interconnects them via the control buses running throughout its length. The backpanel also connects the DPC and Delay cards to integrate them as a set to process outputs from an antenna. The backpanel isolates two such sets from each other except for the control buses which provide various information to these sets.

3.4.2 The Delay Sub-system: Packaging

The DPC and Delay (Fig. 3.17and Fig. 3.18) are both 4-layer PCBs, 100mm by 260mm in size, 1.6mm thick with 35μ copper plating on all layers. The layer assignments are shown in Fig. 3.19 which depicts schematically, the construction of the DPC and Delay cards. Both the cards have the same construction. Both PCBs are designed with 12mil track width and 12mil track-to-track spacing design rule. The via size used is 30mil.

The DPC has one 64-pin EURO in the front for the inputs from the ADCs. Four ADC inputs can be coupled to one DPC. A 96-pin EURO forms the DPC's backpanel connector. This connector is used solely to connect to the Delay cards. The backpanel provides the necessary connection between the two. Each DPC card generates outputs for two delay cards. This is so because the DPC handles four inputs where as the Delay card handles only two. These two inputs correspond to inputs of one FX sideband.

The Delay card takes its input from the DPC via the backpanel through a 64-pin EURO connector. The output of the Delay card is transmitted to the FFT subsystem via a 26-pin FRC connector, located at the front.

The set of three cards, comprising of one DPC and two Delay cards form one set of the delay subsystem. This set processes four signals from an antenna. The inputs and outputs for these sets are all from the front of the subracks.

The Backpanel (Fig. 3.20) is a 3.2mm, 6 layer PCB. Its construction is shown in Fig. 3.21.

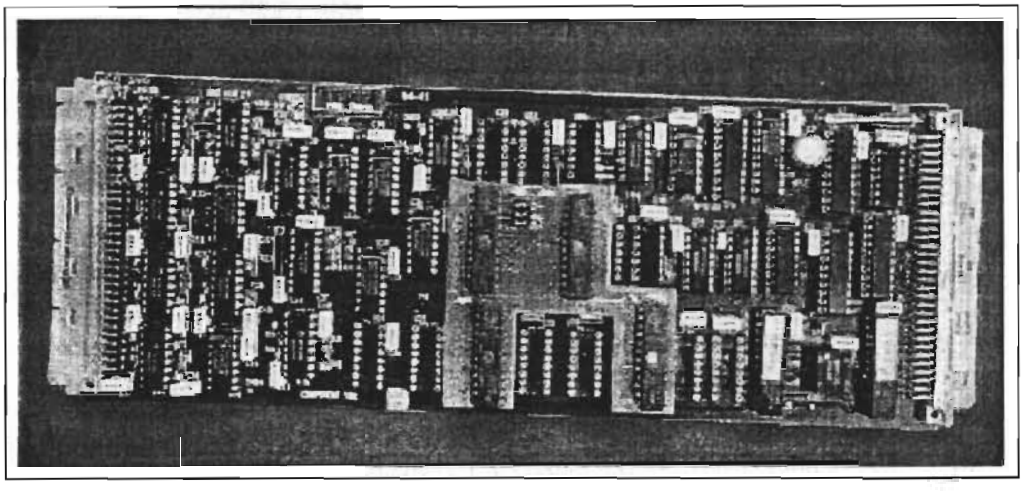

Figure 3.17: The DPC Card

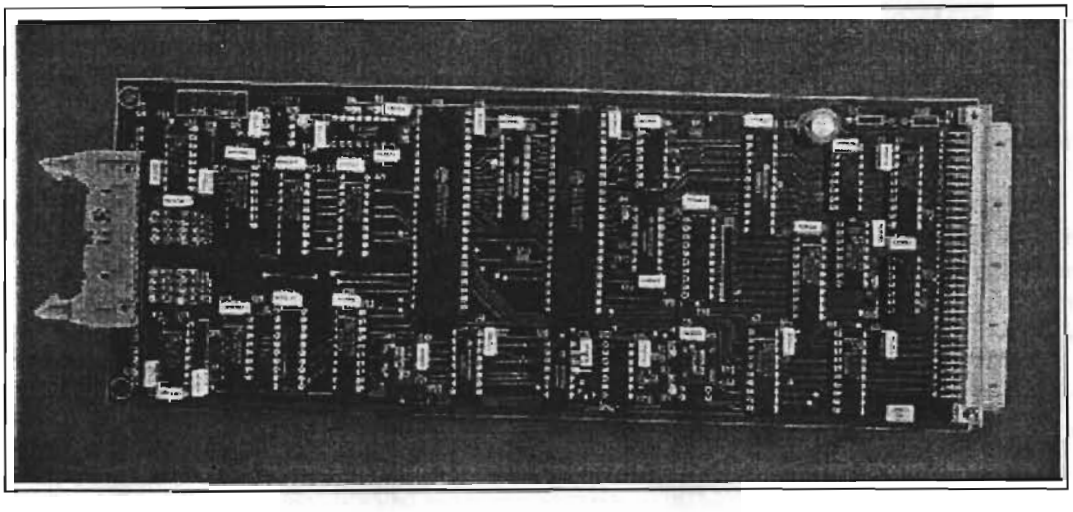

Figure 3.18: The Delay Card

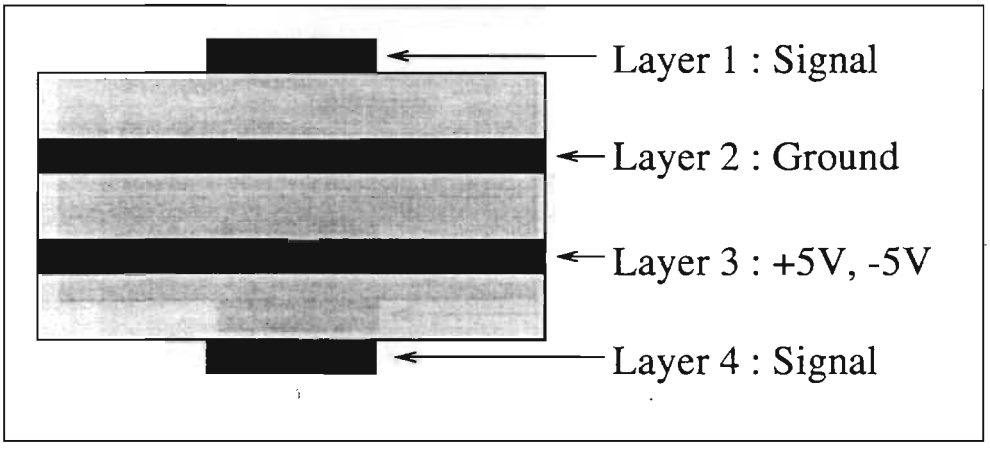

Figure 3.19: The DPC & Delay Card Layer Assignments

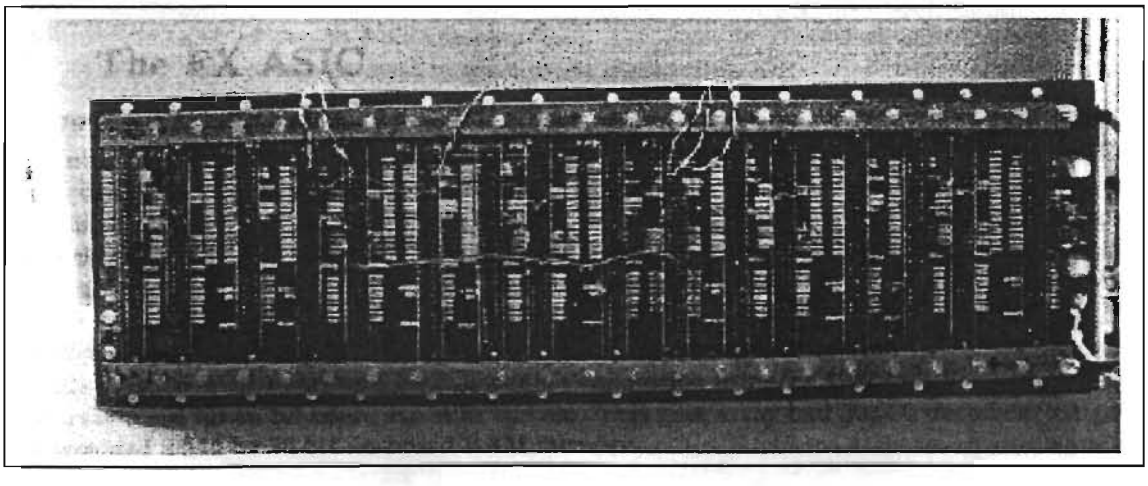

Figure 3.20: The Delay Backpanel

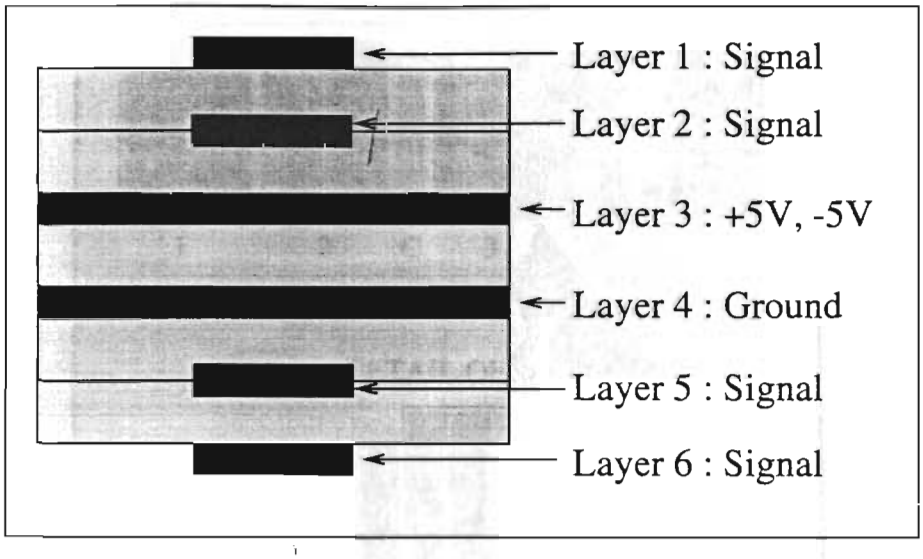

Figure 3.21: The Delay Backpanel : Layer Assignments

Two *Berg* pins are provided for every DPC and Delay card for the ECL clock connections.

A total of 30 DPC-Delay sets are required for the entire system. Each backpanel supports 5 sets in addition to a Delay control card, with an inter-card spacing of 5T (= 1.0 inch) . Thus, a total of 6 sub-racks make up for the entire requirement. A total of two 24U racks are used for the Delay subsystem. Each rack houses 3 subracks plus the power-supply and fan-trays. A schematic front view of the subrack is shown in Fig. 3.22. The arrangement of the rack is shown in Fig. 3.23.

The interconnection of the Delay with the ADC subsystem was described with the ADC subsystem. The Delay cards are connected to the FFT subsystem by 26-core round twist-n-flat jacketed shielded cables. Both ends of this cable use a 26-pin FRC connector.

3.5 The FX ASIC

The FFT and MAC subsystems have an ASIC (referred to as the FX chip here) as their basic computing block. In view of its special architecture, this basic computing element is first described below, before moving on to the subsystems that use it.

This ASIC was designed by National Radio Astronomy Observatory (NRAO), USA for their use in the VLBA Correlator. Later, the GMRT too decided to use the same ASIC. Consequently, a special mode for use by GMRT was also included in the design of the ASIC. The combined requirement of the VLBA and GMRT projects reduced the cost per chip to both its users. The GMRT uses 1550 of these chips and the cost was approximately $82.00 each.

This ASIC was manufactured by LSI Logic Corporation using its 10K gate family 1.5 micron gate array fused on an on-board RAM in the die. This gate array had 30,000 uncommitted gates and provided a 512 by 36-bit on-board RAM memory.

The details of the ASIC and the block diagram are based on the design data provided by NRAO ([15], [16]). The author would particularly like to acknowledge the invaluable discussions with Ray Escoffier and Chuck Broadwell of NRAO which proved pivotal in the understanding of the ASIC.

The ASIC performs, mainly three functions, all at a clock rate of 32MHz

- Radix 4 FFT Butterfly,
- Radix 2 FFT Butterfly and,

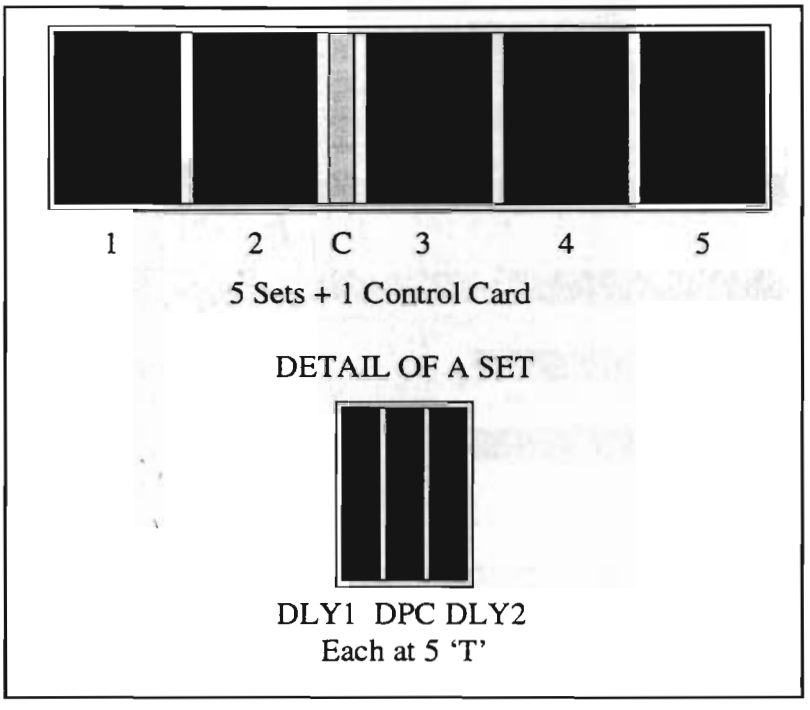

Figure 3.22: The Delay Subrack

- Floating point complex Multiply-accumulations.

The complex floating point signal representation used, has a non-standard form in which the real and imaginary components of a number have a common exponent bit field. Such complex numbers are expressed in a short hand fashion, such as 7,7,4 indicating 7-bit real and 7-bit imaginary mantissa fields and a common 4-bit exponent field. The common exponent has an implied negative sign. The four basic precisions used in the various applications are:

- Data Points being Fourier Transformed: (7,7,4)
- Sin,Cos twiddle factors used in the FFT Butterfly: (5,5,0)
- Time-Domain Window Function : (5,0,4)
- Points being multiply-accumulated: (4,4,4)
- Accumulator Precision: (15,15,6)

The multifunction ASIC can be used for the following applications:

3.5.1 Modes

Radix 4 FFT Butterfly

The ASIC performs a standard radix 4 decimation-in-time FFT butterfly on four digitally sampled complex data points clocked into the chip in four consecutive clock cycles. Two input ports exist into the ASIC, one for the points to be transformed and the other for FFT twiddle factors. A point and its corresponding twiddle factor enter the chip simultaneously on the same clock edge. The ASIC is fully pipelined so as to input and output one complex data point every 32MHz clock

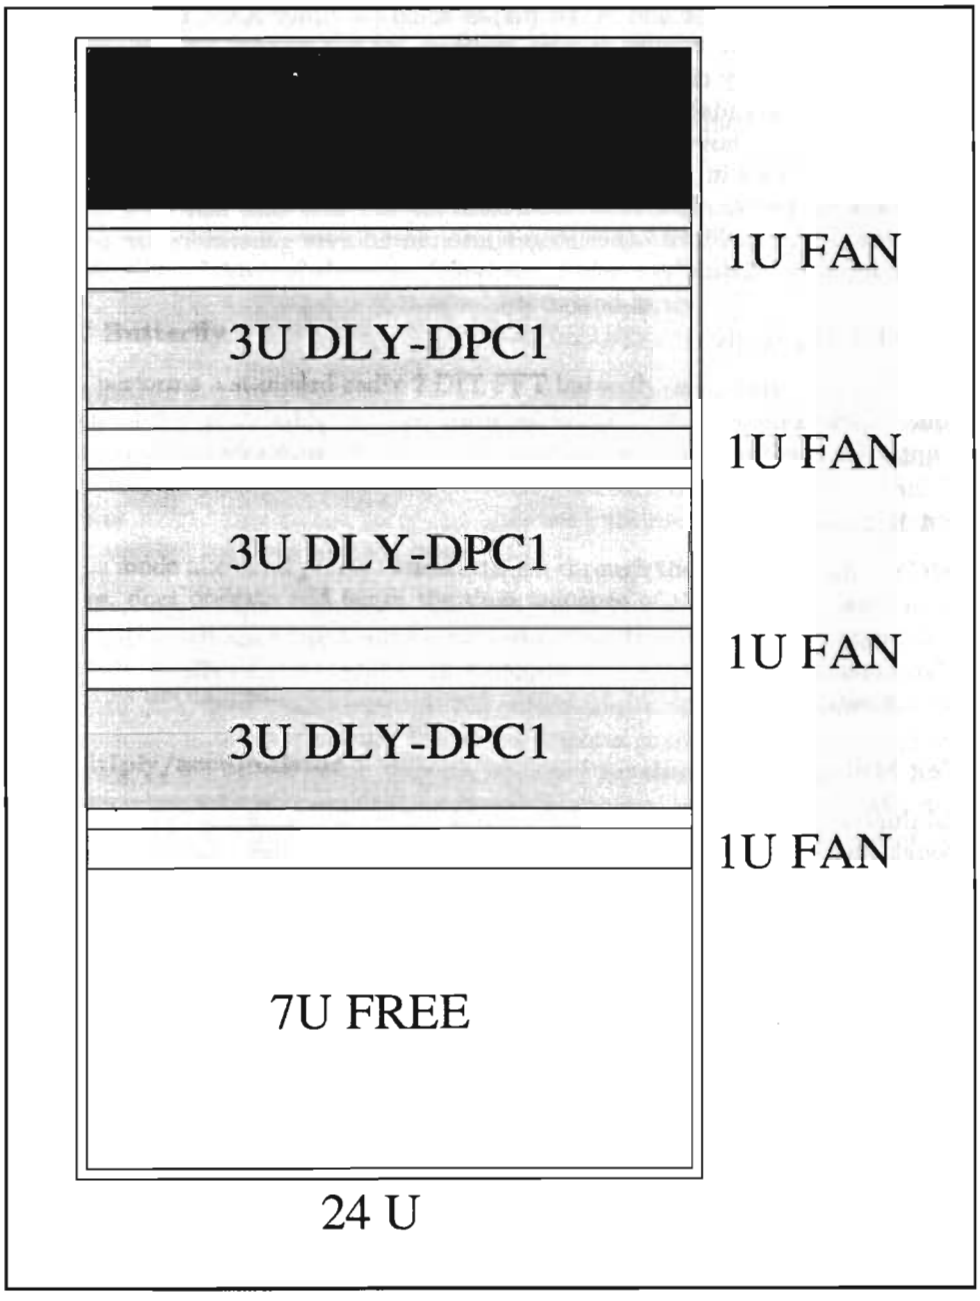

Figure 3.23: The Delay Rack

cycle in addition to performing one complex multiply and add in each clock cycle. The points being input and output from the chip are represented as sign-magnitude floating point numbers.

The on-chip RAM is configured in the FFT mode as two independently addressable 512 deep by 18-bit RAMs providing double buffering of the data points being transformed. The points clocked into the chip are stored into one of the RAMs while the other RAM is read out to the butterfly computing circuitry. The twiddle factors, however, do not go through a storage and are directly clocked into the butterfly circuitry.

The ASICs are cascadable, and hence larger FFTs can be computed by cascading individual radix 4 stages. This is however limited by the internal RAM size to 512 point FFT. The ASICs when cascaded, operate in "stages" and an appropriate addressing sequence is used to order the incoming data properly. The RAM addresses for the first and last FFT butterfly stages are generated externally while for other stages internal address generators are used. The cascaded ASICs thus form a "pipeline".

Radix 2 FFT Butterfly

The ASIC also performs a standard radix 2 DIT FFT butterfly on two digitally sampled complex data points clocked into the chip in two consecutive clock cycles. Most of the details for the radix 4 FFT apply for this function as well.

Radix 4 Bypass

The ASIC in this mode allows all points to flow straight through the chip unaltered. The addressing scheme, however, does operate and hence the time sequence of the output will be different from that at the input. The finite arithmetic precision of the ASIC too affects this time series. Thus, except for a rearrangement of the time sequence and finite precision effects, the ASIC outputs the original time series unchanged.

Complex Multiply/accumulator

In the Multiply-Accumulate (MAC) mode, the ASIC inputs two floating point complex numbers, one through the butterfly data point input port and one through the twiddle factor input port, performs a complex multiplication between the two and adds the complex result into a complex floating point accumulation obtained from the RAM. The accumulation result is stored in the RAM across the entire available width. The RAM thus operates as two banks of 256 deep by 36 bits wide. Three operating modes are available in this application, known as the Non-polar, Polar and the 'Indian' mode.

The Non-Polar Mode In this mode, points to be multiplied enter the chip in pairs, one pair every *second* 32MHz clock. On the average, the RAM is read, a multiplication is done, the result is added to the accumulator and the sum is stored back in RAM in the *two* available clock cycles.

This mode is so termed because, in practice two consecutive clock cycles carry points corresponding to the the right and left circular polarisations. And since, points are read only once every two cycles, information pertaining to only one polarisation gets processed. This mode thus, cannot be used to compute information required for polarisation observations.

If the time series being input to one of the ports of the chip is expressed as $R_1, L_1, R_2, L_2, \ldots$ and the input to the other port as $R'_1, L'_1, R'_2, L'_2, \ldots$ then the output can be either $R_1^* R'_1, R_2^* R'_2, \ldots$ *OR* $L_1^* L'_1, L_2^* L'_2, \ldots$. Here R and L are points from the two polarisations and the subscripts $1, 2, \ldots$ refer to the frequency channels. Hence, this mode is also termed as the *RR Mode.*

The Polar Mode In this mode, point pairs are multiply/accumulated one pair *every* clock cycle. In this mode, two consecutive point pairs entering the chip will be added into the same

accumulation result between reading the accumulator partial sum and writing that sum back into the RAM. Thus, on average, the RAM is read, two multiplications made, both results added via the accumulator to the same accumulation sum, and the new accumulation stored back in RAM each two clock cycles.

Here again if we denote the inputs to the two ports by $R_1, L_1, R_2, L_2, \ldots$ and $R'_1, L'_1, R'_2, L'_2, \ldots$ then the output can be either $R_1^* R'_1 + R_2^* R'_2, R_1^* L'_1 + R_2^* L'_2, \ldots$ *OR* $L_1^* R'_1 + L_2^* R'_2, L_1^* L'_1 + L_2^* L'_2 \ldots$. Hence, this mode is also termed as the *RR-RL Mode.*

This mode is used to calculate the *RR, RL, LR and LL* products to derive the four stokes parameters.

The 'Indian' Mode This mode was incorporated in the ASIC specifically for the GMRT, hence the name. This mode, in implementation, is similar to the *Polar Mode*, and in effect can be considered a mixture of the *Non-Polar* and *Polar* modes. As its output, it produces $R_1^* R'_1 + R_2^* R'_2, L_1^* L'_1 + L_2^* L'_2, \ldots$. It is thus also termed as the *RR-LL Mode.*

Number Controlled Oscillator (NCO)

The ASIC , in addition to the above, also has a 10-bit slice of a number- controlled-oscillator on board, for external applications. It has two secondary storage registers for loading an initial oscillator phase and an oscillator rate. The secondary storage registers have a serial input and output. The NCO adder carry in and carry out lines are pipelined so as to enable any number of bit slices to be cascaded together to make larger NCO's.

3.5.2 ASIC Block Diagram

The "top-level" block diagram of the FX ASIC is given in the appendix. Further details of the ASIC internals are in a set of 84 schematics [16], which give a hierarchical gate-level description of the constituents of the block diagram. The block diagram consists of the following six general sections, which are shown in the block diagram, enclosed by dotted lines. The two fundamental modes of the device are referred to as FFT mode and Multiplier/Accumulator (MAC) mode. FFT mode includes Radix 4, Radix 2 and Bypass modes. MAC mode includes Non-polar, Polar and Indian Polar modes.

1. Input/Output
2. Complex Multiply
3. (a) Floating Point to Fixed point for FFT mode or,
   (b) Floating Point to Common Exponent for MAC mode
4. (a) FFT Butterfly for FFT mode or,
   (b) Accumulator for MAC mode
5. Fixed point to Floating point conversion
6. Number Controlled Oscillator

Sections 1,2,3 and 4 are involved in both FFT mode and MAC mode. Section 5 is involved only in FFT mode. Section 6 is totally independent of the rest of the device except for the 32MHz clock.

Input/Output

This section provides for the data interface for the device, along with RAM storage. Multiplexers are provided so that the two blocks of RAM may be placed in either the input data stream for FFT mode operations or in the middle of the accumulator logic for the MAC application. The address generators allow the double buffered points to be written in one sequence to one RAM block while previous data is read from the second RAM block using a separate sequence. While operating in certain stages in the FFT mode, external address generators can address the RAMs for both read and write purposes. In MAC mode too, either internal or an external address generator controls the read/write sequences, and the memory is addressed as a 512 X 36 bit RAM.

Complex/Multiply

This section provides a complex multiply of two floating point numbers:

$$(R1 + jI1)(2^{exp1}) * (R2 + jI2)(2^{exp2}) = (RE + jIM)(2^{EX})$$

where,

$$\begin{aligned} RE &= [(R1 * R2) - (I1 * I2)] \\ IM &= [(I1 * R2) + (R1 * I2)] \text{ and} \\ EX &= (exp1 + exp2) \end{aligned}$$

$R1 + jI1$ has (7,7,4) resolution and $R2 + jI2$ has (5,5,0) resolution. The mantissas are in sign/magnitude form. After the multiplications, the sign of the $I1 * I2$ product is inverted (to allow the upper adder to perform a subtraction) and then the RE and IM values are converted from sign/magnitude form to one's complement just prior to the addition function. The addition is one's complement, where both sets of eleven input bits are sign extended to 12 bits, producing a 12-bit output.

The two exponents both represent negative numbers where the sign bit is implied. The range of possible values is from zero to minus 12. After the unsigned addition of the two exponents, the resulting 5-bit number can be thought of as being in sign/magnitude form with a range of zero to minus 24 with the implied sign bit = 1, but for purposes of the next section, it can also be thought as one's complement number that has been negated which has an implied sign bit = 0.

Floating Point to Fixed Point or Common Exponent

This section consists of four shifters. Each has 12 to 15 input and 15 output bits. The possible shifts are between 0 and 15 bits. The direction of shift is towards the LSB, with sign extension. The two inner shifters are involved in FFT and MAC modes. The two outer shifters are involved only in MAC mode.

In FFT mode, the inner shifters need to shift an amount equal to the exponent. The rest of the exponent logic in this section in not active.

In MAC mode, the current result from the complex multiplier needs to be added to the accumulator total, which may be accessed from RAM, or may be in an accumulator loop where two or more accumulations may be done before the result is written back to RAM. In either case, the exponent of the multiplier must be compared to the exponent of the accumulated value to determine which set of shifters (inner or outer) must be shifted.

The adder is used to subtract as follows:

$$(exp\ of\ accum\ result) - (exp\ of\ the\ cmplx\ mult) = result$$

The magnitude of the result determines the number of places to shift. The result of the one's complement subtraction is gated to produce a 4-bit magnitude plus a sign bit, where any difference

of greater than 15 is limited to 15 and the sign bit determines which set of shifters will shift. (A shift of more than 15 places would shift below the LSB, so a shift of 15 is sufficient for any longer shift that might be required.)

The final effect is that the mantissas associated with the smaller (more negative) exponent will be shifted until the exponent equals the larger one, which is selected as the common exponent.

The remaining logic relating to the exponent calculation is used to add one to the exponent if a carry is generated in the next stage, where the actual accumulation occurs.

FFT Butterfly or Accumulator

In FFT mode, this section can calculate either a Radix-4 or a Radix-2 butterfly. The calculations required for these operations are indicated on the block diagram. The input points such as (A+jW) represent the result of a complex multiply of a data point and a twiddle factor.

The sequence of events for calculating the real and imaginary values will be described for a Radix-4 operation.

As the points become available at the output of the shifters, they are clocked into two parallel paths in such a manner as to allow the first adder to produce the required sums and differences and clock the results into the four following parallel paths. This makes the results available for the second adder where they are multiplexed and summed to form the final results.

The butterfly timing from the output of the shifters to the output register following the second adders is shown in Fig. 3.24 and Fig. 3.25.

Each column represents one clock cycle.

Row 1 shows the real part of the products at the output of the shifter. A, C, B and D are the four real numbers required for a single butterfly operation. The number associated with each entry (0, 1, 2, 3 etc.) refers to separate butterfly operations. Every four clock cycles a new butterfly operation is completed. Rows 1a, 2a, etc. represent the corresponding points in the imaginary path.

Row 11 represents the other input to the second adder. Row 12 represents the output of the register that follows the second adder.

For FFT lengths that are not divisible by four, the final FFT butterfly chip in the pipeline must do a Radix-2 butterfly. To accomplish this, the clock enable to the A+C and W+Y registers is changed from 8MHz to 32MHz, and the multiplexers that drive the second adder are controlled so that the adder is "bypassed". Only the first output point of each butterfly operation is calculated.

In MAC mode, the clock enables to registers A, C, W, Y, A+C and W+Y are changed to provide a 32MHz clock to the registers. The first adder does the addition for the accumulation function. There are two fundamental modes of operation for the accumulator, non-polarisation and polarisation.

In non-polarisation mode, an accumulation cycle consists of reading a number from RAM, adding a new single value to that number and writing back to RAM, as indicated in the Fig. 3.26.

In polarisation mode, it is necessary to read two numbers from RAM (RxR and RxL for example), add in two new adjacent points of data, and write the results back to RAM. The Fig. 3.27 represents the timing for a 512 point FFT in polarisation mode, where each FFT produces 256 points to be accumulated.

For each point number x

$$\begin{aligned} TxRR &= (AC\ Px\ RR) + (FFT0\ Px\ RR) \\ RxRR &= TxRR + (FFT1\ Px\ RR) \end{aligned}$$

That is, each result is the sum of the accumulation plus points from two adjacent FFTs.

The data written to RAM is delayed one clock from Dre

Clocks	1	2	3	4	5	6	7	8	9	10
1	A1	C1	B1	D1	A2	C2	B2	D2	A3	C3
2	B0	A1	A1	B1	B1	A2	A2	B2	B2	A3
3	B0	B0	A1	A1	B1	B1	A2	A2	B2	B2
4	D0	D0	C1	C1	D1	D1	C2	C2	D2	D2
5	A0+C0	A0+C0	A0+C0	A1+C1	A1+C1	A1+C1	A1+C1	A2+C2	A2+C2	A2+C2
6	A0-C0	A0-C0	A0-C0	A0-C0	A1-C1	A1-C1	A1-C1	A1-C1	A2-C2	A2-C2
7		B0+D0	B0+D0	B0+D0	B0+D0	B1+D1	B1+D1	B1+D1	B1+D1	B2+D2
8			B0-D0	B0-D0	B0-D0	B0-D0	B1-D1	B1-D1	B1-D1	B1-D1
9	A0+C0	A0+C0	A0-C0	A0-C0	A1+C1	A1+C1	A1-C1	A1-C1	A2+C2	A2+C2
10		A0+C0	A0+C0	A0-C0	A0-C0	A1+C1	A1+C1	A1-C1	A1-C1	A2+C2
11		B0+D0	B0+D0	X0-Z0	X0-Z0	B1+D1	B1+D1	X1-Z1	X1-Z1	B2+D2
12			A0+C0 + B0+D0	A0+C0 - B0+D0	A0-C0 + X0-Z0	A0-C0 - X0-Z0	A1+C1 + B1+D1	A1+C1 - B1+D1	A1-C1 + X1-Z1	A1-C1 - X1-Z1

Figure 3.24: The FFT Radix-4 Computation : Real Channel

Clocks	1	2	3	4	5	6	7	8	9	10
1a	W1	Y1	X1	Z1	W2	Y2	X2	Z2	W3	Y3
2a	X0	W1	W1	X1	X1	W2	W2	X2	X2	W3
3a	X0	X0	W1	W1	X1	X1	W2	W2	X2	X2
4a	Z0	Z0	Y1	Y1	Z1	Z1	Y2	Y2	Z2	Z2
5a	W0+Y0	W0+Y0	W0+Y0	W1+Y1	W1+Y1	W1+Y1	W1+Y1	W2+Y2	W2+Y2	W2+Y2
6a	W0-Y0	W0-Y0	W0-Y0	W0-Y0	W1-Y1	W1-Y1	W1-Y1	W1-Y1	W2-Y2	W2-Y2
7a		X0+Z0	X0+Z0	X0+Z0	X0+Z0	X1+Z1	X1+Z1	X1+Z1	X1+Z1	X2+Z2
8a			X0-Z0	X0-Z0	X0-Z0	X0-Z0	X1-Z1	X1-Z1	X1-Z1	X1-Z1
9a	W0+Y0	W0+Y0	W0-Y0	W0-Y0	W1+Y1	W1+Y1	W1-Y1	W1-Y1	W2+Y2	W2+Y2
10a		W0+Y0	W0+Y0	W0-Y0	W0-Y0	W1+Y1	W1+Y1	W1-Y1	W1-Y1	W2+Y2
11a		X0+Z0	X0+Z0	X0-Z0	X0-Z0	X1+Z1	X1+Z1	X1-Z1	X1-Z1	X2+Z2
12a			W0+Y0 + X0+Z0	W0+Y0 - X0+Z0	W0-Y0 + X0-Z0	W0-Y0 - X0-Z0	W1+Y1 + X1+Z1	W1+Y1 - X1+Z1	W1-Y1 + X1-Z1	W1-Y1 - X1-Z1

Figure 3.25: The FFT Radix-4 Computation : Imaginary Channel

Clock	RAM	DR	REG A	REG C	DRE
1	RD AC P0				
2	WR?	AC P0			
3	RD AC P1		AC P0	NEW P0	
4	WR R0	AC P1			R0 = AC P0 + NEW P0
5	RD AC P2		AC P1	NEW P1	
6	WR R1	AC P2			R1 = AC P1 + NEW P1
7	RD AC P3		AC P2	NEW P2	
8	WR R2	AC P3			R2 = AC P2 + NEW P2
9	RD AC P4		AC P3	NEW P3	
10	WR R3	AC P4			R3 = AC P3 + NEW P3
11	RD AC P5		AC P4	NEW P4	
12	WR R4				R4 = AC P4 + NEW P4

(AC = ACCUMULATOR P0,P1,... = POINT NUMBERS R0,R1,... = RESULTS)

Figure 3.26: The MAC Non-Polar Computation

Clock	RAM	DR	REG A	REG C	DRE
1	RD AC P0 RR				
2	RD AC P0 RL	AC P0 RR			
3	WR?	AC P0 RL	AC P0 RR	F0 P0 RR	
4	WR?		AC P0 RL	F0 P0 RL	T0 RR = AC P0 RR + F0 P0 RR
5	RD AC P1 RR		T0 RR	F1 P0 RR	T0 RL = AC P0 RL + F0 P0 RL
6	RD AC P1 RL	AC P1 RR	T0 RL	F1 P0 RL	R0 RR = T0 RR + F1 P0 RR
7	WR R0 RR	AC P1 RL	AC P1 RR	F0 P1 RR	R0 RL = T0 RL + F1 P0 RL
8	WR R0 RL		AC P1 RL	F0 P1 RL	T1 RR = AC P1 RR + F0 P1 RR
9	RD AC P2 RR		T1 RR	F1 P1 RR	T1 RL = AC P1 RL + F0 P1 RL
10	RD AC P2 RL	AC P2 RR	T1 RL	F1 P1 RL	R1 RR = T1 RR + F1 P1 RR
11	WR R1 RR	AC P2 RL	AC P2 RR	F0 P2 RR	R1 RL = T1 RL + F1 P1 RL
12	WR R1 RL		AC P2 RL	F0 P2 RL	T2 RR = AC P2 RR + F0 P2 RR
13	RD AC P3 RR		T2 RR	F1 P2 RR	T2 RL = AC P2 RL + F0 P2 RL
14	RD AC P3 RL	AC P3 RR	T2 RL	F1 P2 RL	R2 RR = T2 RR + F1 P2 RR
15	WR R2 RR	AC P3 RL	AC P3 RR	F0 P3 RR	R2 RL = T2 RL + F1 P2 RL
16	WR R2 RR		AC P3 RL	F0 P3 RL	T3 RR = AC P3 RR + F0 P3 RR

(AC = ACCUMULATOR T = TEMP R = RESULT F0 = FFT#0 F1 = FFT#1 P0,P1,... = POINT NUMBERS)

Figure 3.27: The MAC Polar Computation

Fixed Point to Floating Point Conversion

For FFT mode, the fixed point results of butterfly operations are converted back to (7,7,4) floating point format in this stage.

The calculation of the exponent provides a capability of adding a programmable value to the exponent to counteract the tendency of the exponent to grow more negative at each stage. (The amount of exponent adjust at each stage of the FFT pipeline will be determined in such a manner as to ensure that overflow does not occur to any significant extent.)

Number Controlled Oscillator

This section is functionally independent of the rest of the device. A 10-bit slice of a wider counter is implemented on each device. Cascading some number of these circuits provides a wider number controlled oscillator that increments at a controlled rate. The application for the NCO is to generate the fringe rotation.

3.5.3 The Control Word

The configuration of the ASIC to perform in any one of the modes described above, is done by loading it with a 4-byte *Program Control Word*. This control word is shifted serially, bit 0 first, into the ASIC. The shift register is a chain of positive-edge triggered D-Flip-flops. The parallel outputs of this shift register are input to a set of 32 gated-D-latches, to provide secondary storage.

The serial shifting of the control word is done using four signals provided solely for this purpose. These are

Control Register Shift Input (CRSI) This is the serial input of the 32-bit shift register.

Control Register Shift Output (CRSO) This is the serial output of the 32-bit shift register.

Control Register ShiFT CLocK (CRSFTCLK) This is the shift register clock. A positive edge of this clock inserts the value asserted at the CRSI into the shift register.

Control Register STroBe (CRSTB) This is the active-high gating signal for the D-latch.

A bit is shifted into the shift register with the bit asserted at the CRSI, the CRSTB asserted and a positive edge on the CRSFTCLK.

The CRSO is used to cascade the shift registers of all ASICs on the same card to form a single shift register.

Functional Description of the Control Word Bits

In the following description of the control word bits, the terms "upper" and "lower" refer to realisation of FFT cards required to perform a 1K or 2K point transform. These are not relevant to the GMRT FFT cards.

Bit 31,30,29 : Unused.

Bit 28, RAMOFF : 1 = De-select RAMs (used for IDD current test).

Bit 27, WEOUT : 1 = Switch an always enabled RAM write pulse to be output from the CRBIT pin.

Bit 26, FFTST0 : 1 = Chip is at stage 0 of the FFT pipeline. 0 = Otherwise.

Bit 25, R41K2K : 1 = Chip is either at stage 0 or stage 4 of FFT pipeline and is involved in a 1K or 2K size FFT. 0 = Otherwise.

Bit 24, SEQ3REV : 1 = Chip is at stage 3 or stage 4 of a lower FFT pipeline. 0 = Otherwise.

Bit 23, EXTIN : 1 = Select the external data (NORM or AUX) as the input to the butterfly (normally only for MAC mode, possible for test purposes). 0 = Select the RAM data as the input to the butterfly (normally for FFT modes).

Bit 22, EXPOFENBL : 1 = Enable the exponent overflow detection signal out of the chip. 0 = Disable the output (default for MAC mode).

Bit 21, R2SUB : 0 = Normal. 1 = Radix 2 "mirror" mode where the A-C result is kept rather than the A+C result. This gives the "negative" part of the spectrum.

Bit 20, R2C0INV : Controls the phase of the R2CNT0 output. This output is used as the mux control signal for the multiplexer at the output of the two pipelines, used to combine outputs of the two pipelines into one stream. Changing the phase of this signal can thus change the order of multiplexing.

Bit 19, NCOFRZ : 1 = Allow the NCO accumulator to be loaded with the data in the normal manner, but "freeze" the NCO operation so that the value will not ramp up. 0 = Allow normal NCO operation.

Bit 18,17,16, WRSTG : Bit 18 forms the MSB and Bit 16 the LSB of the three bit field which defines the address generation sequence to be used for writing to RAM. Values 0 to 4 = Sequence 0 through 4 (set = 4 for MAC mode). 5 = Not used. 6 = External address sequence at stage 0. 7 = Not used.

Bit 15, R4CNT0 : 1 = CNT0 same as CT\0. 0 = CNT0 same as CT0. This bit is low for all MAC modes and for radix 2, 2K in a lower pipeline. It is high for all other FFT modes.

Bit 14,13,12, RDSTG : Bit 14 forms the MSB and Bit 12 the LSB of the three bit field which defines the address generation sequence to be used for reading from RAM. Values 0 to 4 = Sequence 0 through 4 (set = 4 for MAC mode). 5 = Not used. 6 = Not used. 7 = External address sequence at stage 5.

Bit 11,10 ADEXP : Bit 11 forms the MSB and bit 10 the LSB of the two bit field. Value 0 through 3 = the number to "add" to the exponent at the input to the butterfly.

Bit 09, DISCINTWE : 1 = Inhibit the internal RAM write enables. 0 = Enable the internal write enables.

Bit 08, EXTEXT : 1 = Force control of the RAM address multiplexers to come from the chip input pins rather than from the internal control logic. 0 = Allow internal control of the multiplexers.

Bit 07, NM/AUX or UPR/LWR : This is a dual function bit. In certain modes, this bit is set high, to select the normal data input. (If set low, it would select the auxiliary data input.) Set = 1 for MACNP and MACPA0, 0 = for MACPA1.

In some modes, the input toggles back and forth between the norm and aux. For these modes this bit is a don't care, except for radix 4, 1K or 2K modes in stage 0, where this bit is used to specify the upper and lower pipeline. 1 = upper, 0 = lower.

Bit 06, 2K/<2K or BP/NBP or A1/0 : This bit is a triple function bit.

For Radix 2 modes 1 = 2K FFT size. 0 = Less than 2K FFT size.

For Radix 4 modes 1 = Bypass. 0 = Not bypass.

For MAC Polarisation modes 1 = Array 1. 0 = array 0.

Bit 05, R4/R2 or P/NP : This bit is a dual function bit.

For FFT modes 1 = Radix 4. 0 = Radix 2.

For MAC modes 1 = Polarisation. 0 = Non-polarisation.

Bit 04, TOGCEN8 : 0 = FFT. 1 = Normal MAC modes. Set = 0 in MACPA0 for Indian mode.

Bit 03, TEST : 1 = Test mode. 0 = Not test mode.

Bit 02, PGENBL : 1 = Enable the pulsar gate logic. A must for MAC mode. 0 = Disable the pulsar gate logic (chip ignores pulsar gate).

Bit 01, RAMONIN : 1 = Route input data to RAM. 0 = Route MAC accumulator results to RAM.

Bit 00, CARDUSE : This bit is not used on the chip. It is provided for use on the card.

3.5.4 Hardware Signals

The ASIC requires certain control signals to process the data when operating in FFT and MAC modes. These control signals are generated by the respective control cards - FFT control and MAC control for the FFT and MAC subsystems. These signals are distributed to all the cards in a subrack by the backpanel. The control cards in different subracks are synchronised to each other using a synchronising signal generated at a single point and distributed wherever required.

Signals in FFT mode

The control signals required for the FFT operation can be classified into five groups and are described below.

Control Word Signals These as described in the previous section and comprise of CRSI, CRSO, CRSTB and CRSFTCLK. In addition to these, the CRBIT pin on the ASIC is used to check if the ASIC is configured with the correct control word. Since the control word for all ASICs on a card is loaded in as a single pattern, the Bit 27, WEOUT, of the control word is programmed to be a 1 in every alternate ASIC and a 0 in every other. On ASICs which have this bit as a 1, the RAM write signal, which has a similar shape as the clock, is available on the CRBIT pin. On ASICs which have this bit as a 0, the CRBIT is either a LOW or a HIGH. Thus every alternate ASIC in the control word path should exhibit a clock at its CRBIT pin. This bit is the last useful bit (Bits 28,29,30,31 are 0 for all operating modes) in the control word pattern. This in conjunction with Bit 0, CARDUSE - the state of which can be checked on the CRSO pin of the ASIC - provide a reliable mechanism to verify if the control word has been shifted properly.

The control word configures the 5 ASICs of the pipelines to be as stage 0,1,2,3 and 5. Stage 4 is missing since the 1K or 2K FFT modes are not used on the GMRT FFT card. The last ASIC has to be operated in stage 5 and not in stage 4 because the external addresses, which are required in the last stage, can be given only in stages 0 and 5.

The control word also configures the proper read and write stage address sequences for all the ASICs. Stages 0 and 5 get the externally generated sequences while all the rest use internally generated sequences.

The pins for NCO and external addresses are shared, on the ASIC. Hence, the NCO output cannot be obtained from stages 0 and 5, although the counting process of the NCO is still active. Thus, stages 1,2 and 3 are the only chips available for NCO outputs.

Bit 05, R4/R2 or P/NP : This bit is a dual function bit.

For FFT modes 1 = Radix 4. 0 = Radix 2.

For MAC modes 1 = Polarisation. 0 = Non-polarisation.

Bit 04, TOGCEN8 : 0 = FFT. 1 = Normal MAC modes. Set = 0 in MACPA0 for Indian mode.

Bit 03, TEST : 1 = Test mode. 0 = Not test mode.

Bit 02, PGENBL : 1 = Enable the pulsar gate logic. A must for MAC mode. 0 = Disable the pulsar gate logic (chip ignores pulsar gate).

Bit 01, RAMONIN : 1 = Route input data to RAM. 0 = Route MAC accumulator results to RAM.

Bit 00, CARDUSE : This bit is not used on the chip. It is provided for use on the card.

3.5.4 Hardware Signals

The ASIC requires certain control signals to process the data when operating in FFT and MAC modes. These control signals are generated by the respective control cards - FFT control and MAC control for the FFT and MAC subsystems. These signals are distributed to all the cards in a subrack by the backpanel. The control cards in different subracks are synchronised to each other using a synchronising signal generated at a single point and distributed wherever required.

Signals in FFT mode

The control signals required for the FFT operation can be classified into five groups and are described below.

Control Word Signals These as described in the previous section and comprise of CRSI, CRSO, CRSTB and CRSFTCLK. In addition to these, the CRBIT pin on the ASIC is used to check if the ASIC is configured with the correct control word. Since the control word for all ASICs on a card is loaded in as a single pattern, the Bit 27, WEOUT, of the control word is programmed to be a 1 in every alternate ASIC and a 0 in every other. On ASICs which have this bit as a 1, the RAM write signal, which has a similar shape as the clock, is available on the CRBIT pin. On ASICs which have this bit as a 0, the CRBIT is either a LOW or a HIGH. Thus every alternate ASIC in the control word path should exhibit a clock at its CRBIT pin. This bit is the last useful bit (Bits 28,29,30,31 are 0 for all operating modes) in the control word pattern. This in conjunction with Bit 0, CARDUSE - the state of which can be checked on the CRSO pin of the ASIC - provide a reliable mechanism to verify if the control word has been shifted properly.

The control word configures the 5 ASICs of the pipelines to be as stage 0,1,2,3 and 5. Stage 4 is missing since the 1K or 2K FFT modes are not used on the GMRT FFT card. The last ASIC has to be operated in stage 5 and not in stage 4 because the external addresses, which are required in the last stage, can be given only in stages 0 and 5.

The control word also configures the proper read and write stage address sequences for all the ASICs. Stages 0 and 5 get the externally generated sequences while all the rest use internally generated sequences.

The pins for NCO and external addresses are shared, on the ASIC. Hence, the NCO output cannot be obtained from stages 0 and 5, although the counting process of the NCO is still active. Thus, stages 1,2 and 3 are the only chips available for NCO outputs.

The control word also changes the twiddle format in stage 0 from the normal (5,5,0) to (5,0,4) as a provision for a time-domain window to be used. The stage 0 does not require any twiddle factors because their effect is absorbed in the radix-4 logic.

INITs This stands for *INITialisation*. The ASIC has an internal 9-bit counter, the outputs of which are used to generate various control patterns. The INIT acts as a reset to this counter and as a consequence signifies the beginning of a new FFT computation. This is an active-low pulse, one clock cycle wide, occurring every 516 cycles. It is this signal that defines the FFT cycle, the basic operating cycle of the GMRT correlator. Every FFT cycle, a new set of 512 points are processed to produce 256 spectral channels. After every 512 clocks, an additional 4 clocks are required to accomplish the internal bank switching. No data is processed during these 4 cycles. In order to avoid loss of data, the FFTs are run faster by an equivalent amount.

Every ASIC operating in Radix-4 mode requires 11 clock cycles to compute the FFT. This introduces a pipeline delay and hence every subsequent ASIC has its FFT cycle staggered by this amount. The INITs for the 5 stages (0 to 5) are all similar in structure but offset from their preceding stage INIT by 11 clocks.

Twiddle Inputs All stages except stage 0 require twiddle factors for computing the FFT. The stage 0 twiddle input accepts a time-domain window instead, in a (5,0,4) format.The twiddle factors are in the (5,5,0) format.

The twiddle factors or trigonometric tables, often referred to as *trigs* in short, for stages 1,2 and 3 are supplied by the FFT control card. The trigs for stage 5, also incorporate the *Fractional Sample Time Correction (FSTC)* and are stored on the FFT card in RAMs. The stage 5 trigs can be manipulated to provide a radix-2 bypass mode, since this is not provided for, by the logic inside the ASIC. This mode is required when, say, a 256-point transform is to be computed.

External Addresses The stage 0 and stage 5 external addresses are stored on the FFT card in RAMs, which are loaded at the system initialisation stage by the FFT control card. These, unlike the trigs, cannot be distributed globally from the control card because the different FFT cards in the same subrack can, in principle, be computing FFTs of different lengths.

Signals for NCO Control These can be classified into the two following groups

Signals for loading the NCO The NCO requires a 20-bit number to be input. This signifies a 10-bit initial phase and a 10-bit increment. The 20-bit number is shifted serially into a 20-bit register in a manner similar to the control word. The NCO has NCOSI, NCOSO, NCOSRCLK as the NCO Shift Input, Shift Output and Shift Register CLocK respectively. The NCOs of stages 0,1 and 5 are cascaded to form a 30-bit NCO. The NCOs of stages 2 and 3 serve as 9-bit counters, clocking at 32MHz for addressing the several on-board RAMs on the FFT card.

Signals for NCO operation These signals are the NCOLDIN, NCOSTA, NCORW, NCOCE, NCOCLK, NCOCI and NCOCO. The NCOLDIN (NCO LoaD INput) is a load line which strobes the initial oscillator phase and phase increment values from the secondary storage registers to the active registers. The accumulator of the NCO is clock enabled, to allow selectable update rates. An intermediate storage register will store the NCO phase when the NCOSTA (NCO STore Accumulator) is asserted. The accumulator is preset to this stored phase when the NCORW (NCO ReWind) signal is asserted. NCOCE is the NCO Chip Enable signal and NCOCLK is the NCO CLocK which updates the accumulator. NCOCI and NCOCO are the carry-in and carry-out signals which are used to cascade NCOs.

Signals in MAC mode

The signals in MAC mode have two basic periodicities. One set of signals repeat every 516 cycles where as the other set repeats once every *Short Term Accumulation (STA)* cycle. The STA cycle is one higher than the FFT cycle in the operating cycle hierarchy. The highest being the LTA cycle which is one higher than the STA cycle. An STA cycle typically consists of 4096 FFT cycles. This number determines the amount of averaging performed by the MAC system and can be as low as 256 cycles.

The signals can be classified into three groups

Control Word Signals The control word in MAC mode configures the ASICs in one of the three MAC modes - Non-Polar, Polar or Indian Mode. The control word signals are the same as required for the FFT mode.

FFT Cycle Signals These signals have a modulo 516 clock cycle structure and are active within each FFT cycle of any STA cycle.

MACINIT *MAC INITialise.* It is similar in structure to the FFTINIT, active-low and asserted for one clock cycle every 516 cycles. It resets the internal address counter used for reading data out of the RAM. The periodicity can be lesser than 516 cycles in case of integrating spectral channels lesser than 256.

MACWE *MAC Write Enable.* Active-low and similar to MACINIT. This latches the result of integration read out from the accumulator (RAM) into a 36-bit register inside the ASIC.

External Addresses These are of two kinds - one for the incoming data from the FFTs and the other for the results to be read out. The addresses asserted during the active period (non-dead time period) of the FFT cycle are used for reading out the accumulators from the active bank of the RAM. Each address (from 0 to 255) is asserted for two clock periods. The addresses for write back are provided internally.

The address for the read out of the results is asserted for one clock period during the dead-time. This address reads out the corresponding RAM location from the non-active bank of the RAM (that is the bank on which accumulation is over).

BANKSWITCH This signal has periodicities of both types. This is MSB of the external address and its function, as its name conveys, is to select the correct bank of the RAM as appropriate to that instant.

During the 512 active clock periods, integration is being performed on the active bank. Thus the BANKSWITCH selects this bank. During one clock period in the dead-time of 4 clocks, a result is required to be read out from the non-active bank. The BANKSWITCH thus, selects the other bank for this one clock period.

At the end of an STA cycle, the roles for the active and non-active banks are reversed and so is the polarity of the BANKSWITCH signal.

Cardselects The read-out operation of the results of the accumulation are performed in the following sequence - first a particular channel from a particular chip is read out from ALL the cards, next the channel is incremented. On reading out all channels from a particular chip on all cards, the chip number is incremented. All the 256 channels from all the chips on all cards are thus read out. The read-out operation can thus be summarised as a counter: [CARDS, CHANNELS, CHIPS] with cards as the LSB and CHIPS as the MSB. The CARDS increment from 0 to 10, CHANNELS from 0 to 255 and the chips from 0 to 15.

Each Cardselect is thus asserted once *every* FFT cycle to read out a CHANNEL from a CHIP.

AUXOUT The MAC result is 36-bit wide while only 18 output pins are available on the chip. AUXOUT multiplexes the 36-bit result in two halves of 18 bits each out of the chip. The 36-bit result is stored in a 36-bit register on the ASIC using the proper external address and MACWE as the write pulse. This result is in the (15,15,6) format. AUXOUT = 1 brings out the 15-bits of the real mantissa with the 3 LSBs of the 6-bit exponent. AUXOUT = 0 brings out the 15-bits of the imaginary mantissa with the 3 MSBs of the exponent.

The AUXOUT thus is a square wave having 11 periods in each FFT cycle, since it has to toggle once for each of the 11 MAC cards read out during the FFT cycle.

STAREAD A positive edge of this signal is used to latch the 18-bit output from the ASIC into the on-board register. Thus the STAREAD is double the frequency of AUXOUT. For every logic state of the AUXOUT, the STAREAD toggles once.

STA Cycle Signals The signals grouped under this category have a periodicity of one STA cycle. However, though lower in frequencies as compared to the signals active within a FFT cycle, they have to adhere to stringent timing requirements. They have precise instants at which to be asserted and in some cases do not have a latitude of even a single clock cycle.

CLRACC *CLeaR ACCumulator* This signal is asserted for an entire FFT cycle during the *first* FFT cycle of every STA cycle. The CLRACC emulates an accumulator clear operation in the following manner. Instead of explicitly writing a zero into the accumulator, it clears the port of the adder holding the present accumulator value, to zero. The incoming product is thus added to zero before being stored into the accumulator.

BANKSWITCH This signal is used to accomplish the reversal of roles, from active to non-active and vice-versa, of the two RAM banks.

This signal and the one in the previous section are actually one and the same. It is the same hardware line which exhibits the two different events during the STA cycle.

ROW and COLUMN These form the chip select signals for the ASICs. Used for enabling an ASIC for read-out of results after accumulation. The ROW and COLUMN each are fields 2 bits wide, which are decoded on the MAC card to select one of the four ROWs and one of the four COLUMNs.

3.6 The FFT Subsystem

3.6.1 The FFT Subsystem: Circuit Design

The FFT subsystem involves three card types : the FFT card (Fig. 3.28), the FFT control card and the FFT Backpanel. It performs the following tasks :

1. Applies a time-domain window on the data,
2. Compensates for the fringe phase ,
3. Performs a DFT on the input data and,
4. Compensates for the fractional delay difference.

1. The FFT Card

The FFT Card consists of the following sections:

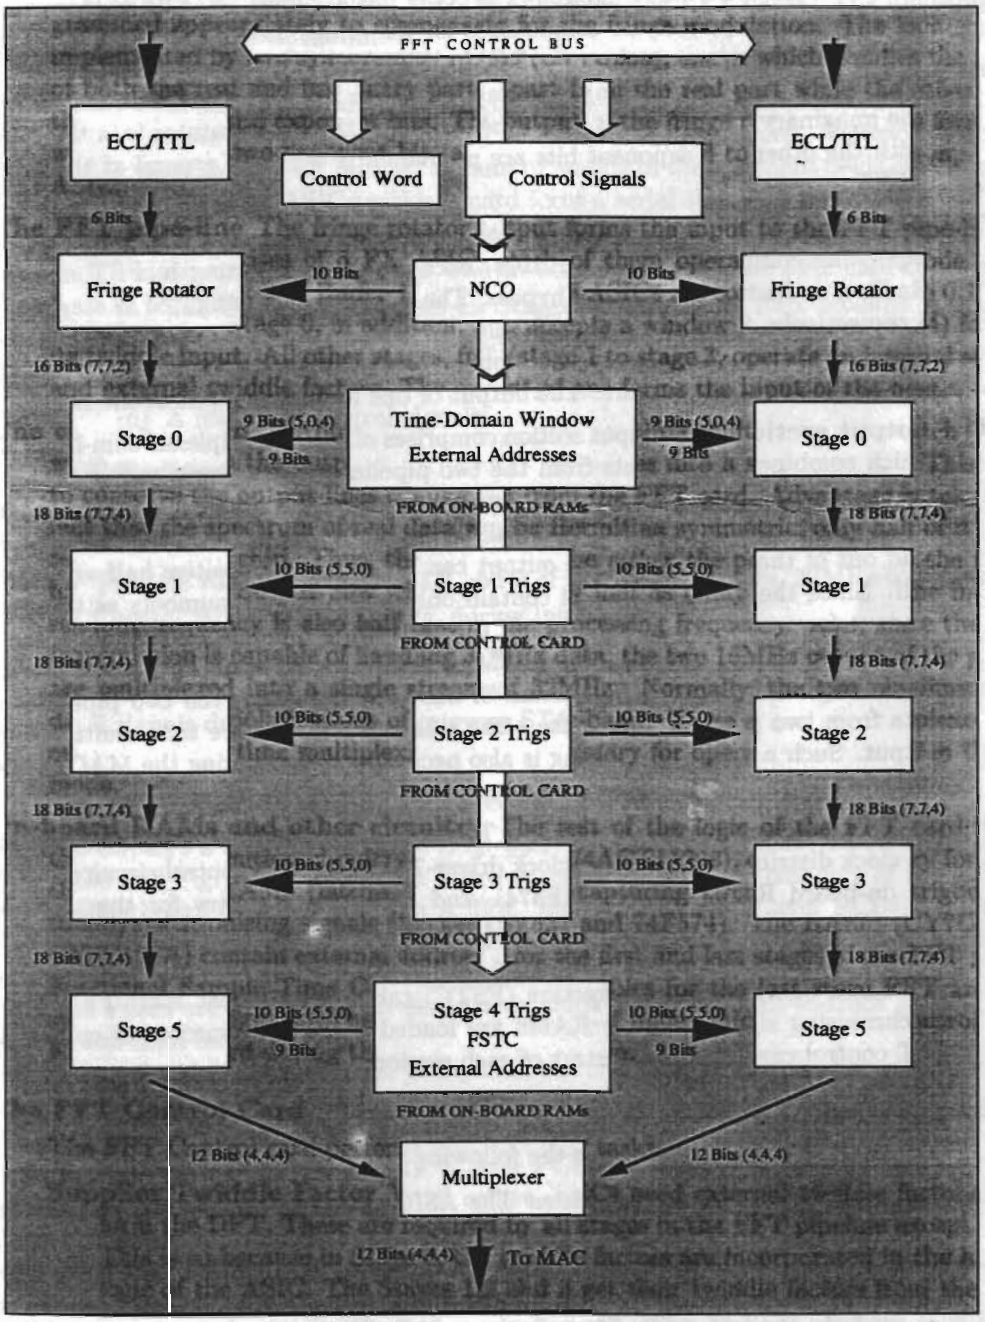

Figure 3.28: The FFT Card : Block Diagram

The input section and the fringe rotator The 4-bit wide ECL data input from the delay sub-system is converted to TTL levels by a MC10125, latched using a 74F574 and fed to the fringe rotator. The fringe rotator is implemented as a lookup table with the data and the NCO as its inputs. The NCO initial phase and rate of increment is programmed appropriately to compensate for the fringe modulation. The lookup table is implemented by two synchronous ROMs (CY7C265), one of which handles the sign bits of both the real and imaginary parts apart from the real part while the other handles the imaginary and exponent bits. The output of the fringe rotator is in the form (7,7,2) with the other two exponent bits are permanently wired to ground at the input of the ASIC.

The FFT pipe-line The fringe rotator output forms the input to the FFT pipe-line. The pipe- line comprises of 5 FX ASICs, each of them operating in FFT mode – either Radix-4 or Radix-2 or Radix-4 bypass. The 5 ASICs are configured as stage 0,1,2,3 and 5 respectively. Stage 0, in addition, also accepts a window function in (5,0,4) format as its twiddle input. All other stages, from stage 1 to stage 3, operate on internal addresses and external twiddle factors. The output of one forms the input of the next.

The output section The output section comprises of three multiplexer-cum-latch (74F399) which combines the outputs from the two pipelines into a single output. This is done to conserve the output lines coming out from the FFT card. Advantage is taken of the fact that the spectrum of real data will be Hermitian symmetric, only half of it need be read out of the chip. Thus, the output can be *either* the positive half *or* the negative half. Since the output now will contain only half as many numbers as the input, the readout frequency is also half that of the processing frequency. And, since the output transmission is capable of handling 32MHz data, the two 16MHz output of the pipelines are multiplexed into a single stream of 32MHz. Normally, the two pipelines process data from two polarisations of the same side-band. These are time-multiplexed in the output. Such a time multiplexing is also necessary for operating the MAC in the polar mode.

On-board RAMs and other circuitry The rest of the logic of the FFT card includes the clock distribution circuitry (clock driver-74ACT11208), control circuitry for loading the on-board RAMs (latches-74F574), and capturing circuitry for the trigonometric tables/synchronising signals (latches-74F821 and 74F574). The RAMS (CY7C199 and CY7C128A) contain external addresses for the first and last stages of the FFT pipeline, Fractional Sample Time Correction (FSTC) tables for the last stage FFT and a few synchronising signals. These RAMs are loaded under microprocessor control by the FFT control card during the start of each session.

2. The FFT Control Card

The FFT Control card performs the following tasks:

Supplies Twiddle Factor Values The ASICs need external twiddle factors to perform the DFT. These are required by all stages in the FFT pipeline except stage 0. This is so because in Stage 0, the twiddle factors are incorporated in the hardware logic of the ASIC. The Stages 1,2 and 3 get their twiddle factors from the control card. In the last stage, Stage 5, the twiddle factors are combined with the FSTC and hence are supplied from the on-board RAM.

Supplies INITs The FFT INITs to all stages, 0 through 5 are generated and distributed by the control card.

Loads Fringe Values The NCO for each pipe-line on every FFT card is programmed with new values of initial rate and increment every STA cycle to accomplish fringe stopping.

Loads the FSTC Values Similar to the fringe values. One number per FFT card, ranging from 0 to 31, pertaining to the amount of FSTC is loaded once every few FFT cycles to compensate for the change in fractional delay.

Initialisation of the FFT Cards At the beginning of every observing session the control card performs the initialisation of all the FFT cards. The following tasks constitute this initialisation

Loading Control Words into the ASICs All ASICs in the FFT cards are loaded with their required control words. These control words, each 32 bits wide, define parameters such as the length of FFTs to be performed, sequence of polarisations in the FFT output, amount of scaling in the exponent fields of the numbers etc. All the ASICs within a card form a serial shift register 320 bits (10 ASICs X 32 bits each) long, into which the FFT control card shifts the control word.

Loading on-board RAMs The on-board RAMs on the FFT card contain

(a) The External Addresses for Stages 0 and 5 of the FFT pipeline

(b) A user definable time domain window

(c) The FSTC tables

(d) A few NCO control signals

3. The FFT Backpanel

Interfaces the FFT cards with the FFT control. Has the following buses through all the cards :

- Twiddle Factors for Stages 1,2 and 3 of the FFTs (30 bits, 32MHz continuous)
- Control Word Signals (4 bits, during initialisation)
- The FFT Inits (5 bits, one clock width pulses every FFT cycle)
- The NCO Signals (NCO control word: 3 signals during initialisation; NCO update: 4 signals during NCO update every STA cycle)

3.6.2 The FFT Sub-system: Packaging

The FFT card employs a two-tier structure (Fig. 3.29) consisting of a main board and a pipeline board (Fig. 3.30 and Fig. 3.31). The main board contains all the logic circuitry of the FFT card while the pipeline board contains only the ASICs. Both boards are 6 layer PCBs, 1.6mm thick with 35 μ copper plating on all layers. The main board is 100mm by 280mm while the pipeline board is 100mm by 246mm. The layer assignments for both boards are same and is shown in Fig. 3.32. Both the PCBs are designed with 8mil track width and 8mil track-to-track spacing design rule. The via size used is 16mils.

The main board which plugs into the backpanel, has one 128-pin EURO (4 rows X 32 pins) as the backpanel connector and one 26-pin FRC for the input from Delay. In addition, it has three 100-pin surface mount connectors for the pipeline board.

The Backpanel (Fig. 3.33) is a 3.2mm, 6 layer PCB. Its construction is the same as that of the FFT card shown in Fig. 3.32. Like all others, two *Berg* pins are provided for ECL clock connections.

A total of 60 FFT cards are required for the entire system. Each backpanel supports 10 FFT cards in addition to two FFT Control cards. The inter-card spacing is 7T (= 1.4 inch). Thus, a total of 6 sub-racks make up for the entire requirement. The 10 FFT cards in each subrack are split into two groups of 5 each. Each group has its own control card. These two groups are independent of each other, except for the power and ground layers of the backpanel. Each FFT subrack is housed in one of the six FX racks. The FX rack is described in greater detail in the discussion on MAC subsystem packaging.

The interconnection of the FFT with the Delay subsystem is already described in the discussion on the Delay subsystem. The FFT cards are connected to the MAC subsystem by 50-core round

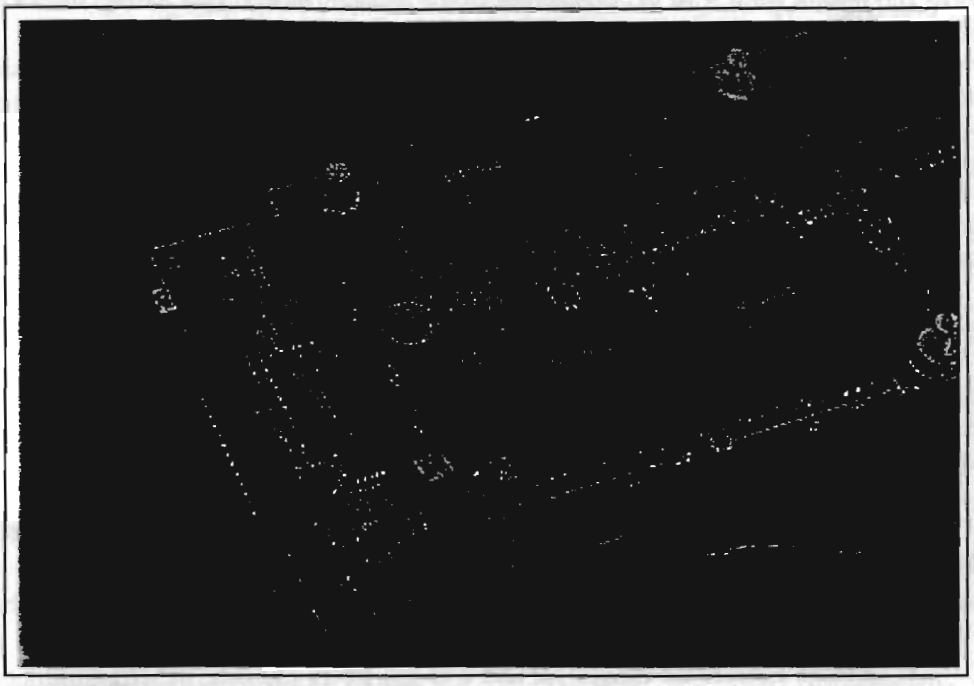

Figure 3.29: The FFT Card

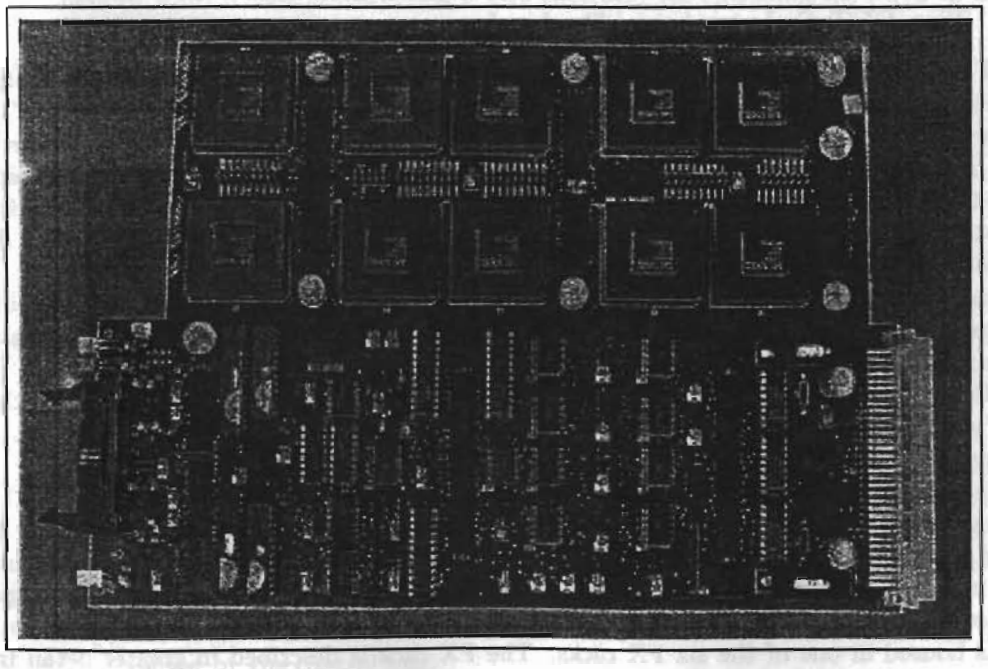

Figure 3.30: FFT Main Board (Lyr.6) and Pipeline Board (Lyr.1)

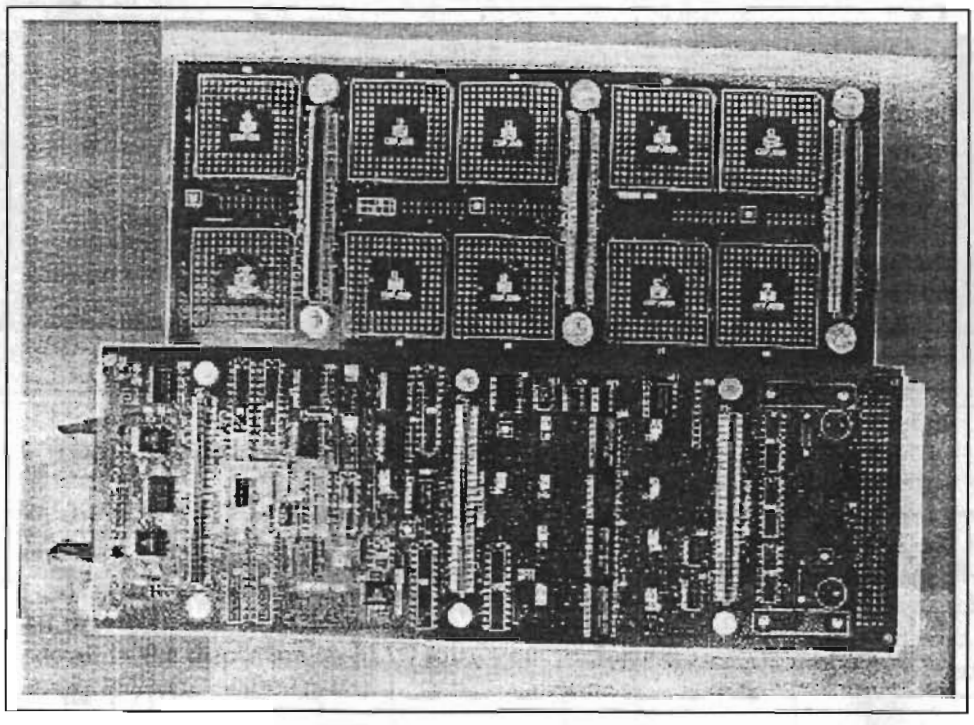

Figure 3.31: FFT Main Board (Lyr.1) and Pipeline Board (Lyr.6)

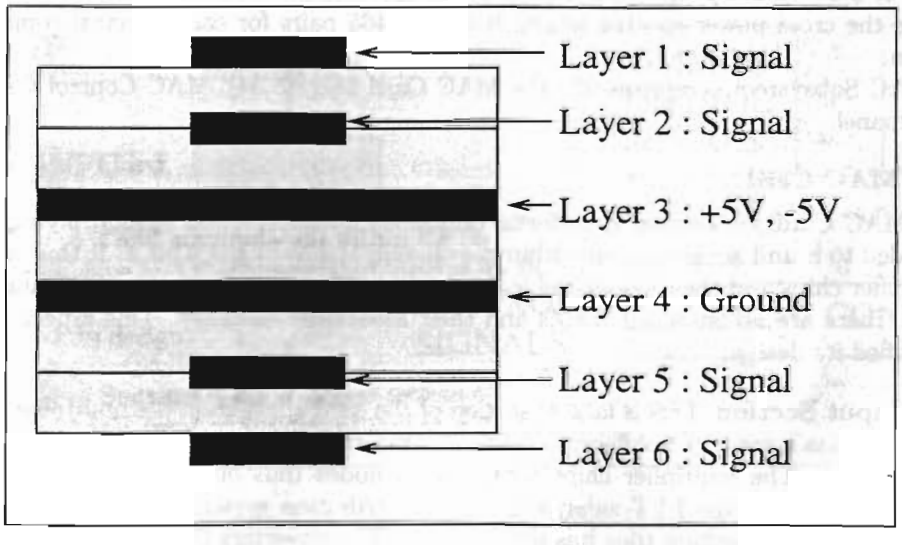

Figure 3.32: FFT Card : Layer Assignments

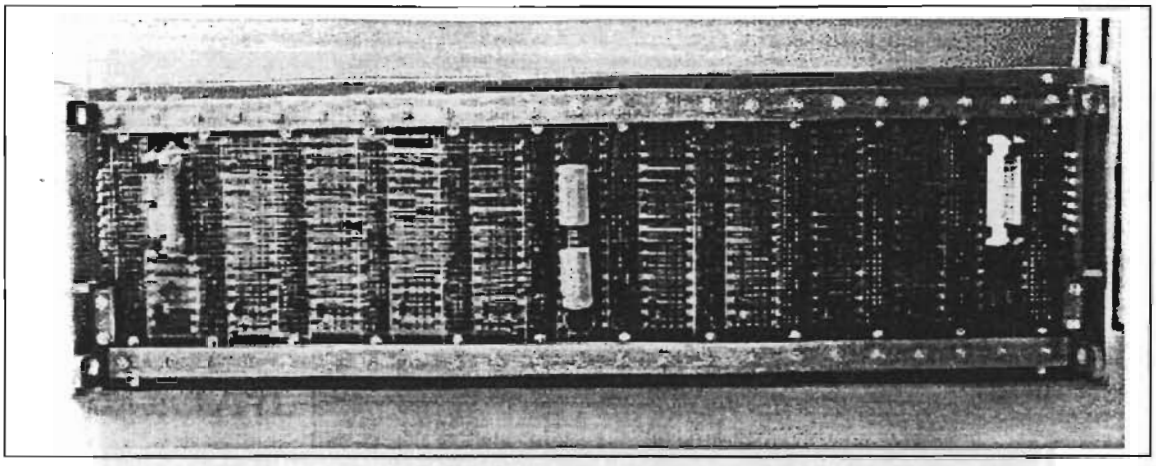

Figure 3.33: FFT Backpanel

twist-n-flat jacketed shielded cables. Two FFT card outputs are combined into one 50-pin FRC connector and is transmitted to a distribution system, as TTL signals interleaved with ground. The system then converts these to ECL and transmits them to the MAC subsystem. This is described in greater detail in the discussion on MAC subsystem packaging.

3.7 The MAC Subsystem

3.7.1 The MAC Subsystem : Circuit Design

The spectra provided by the FFT subsystem are correlated to provide the cross power spectrum by the Multiply and Accumulate Subsystem. This is the crucial subsystem responsible for the realisation of a 30-by-30 array. Inputs from all the 30 stations of GMRT are fed into this array to compute the cross power spectra arising from the 465 pairs for each spectral component and polarisation.

The MAC Subsystem comprises of : the MAC Card (Fig. 3.34), MAC Control Card and the MAC Backpanel.

1. **The MAC Card**

   The MAC Card, conceptually, performs only one operation - that of multiplying the inputs provided to it and accumulation within the on-chip RAMs of the ASICs. It thus has only the multiplier chips and their associated logic as its constituents. Unlike the FFT and the Delay cards, there are no on-board RAMs and their associated circuitry. This aspect has greatly simplified its design.

   The Input Section This is largest section of the card apart from the multiplier chips. The card has a total of 8 different inputs which form the 4 rows and 4 columns of a 4-by-4 matrix. The multiplier chips form the 16 nodes thus obtained. These 8 inputs are received from the FFT subsystem via a distribution system (described later) as ECL signals. This section thus has the ECL-to-TTL converters (MC10125) which feed a set of latches (74F821 and 74F574) to capture the data on the card with the on-board clock. These data latches then feed the multiplier chips.

   The Multiplier Chips The 16 multiplier chips are the FX ASICs operating in one of the three MAC modes viz. Non-polar, Polar or Indian Polar. Of these, the Indian polar

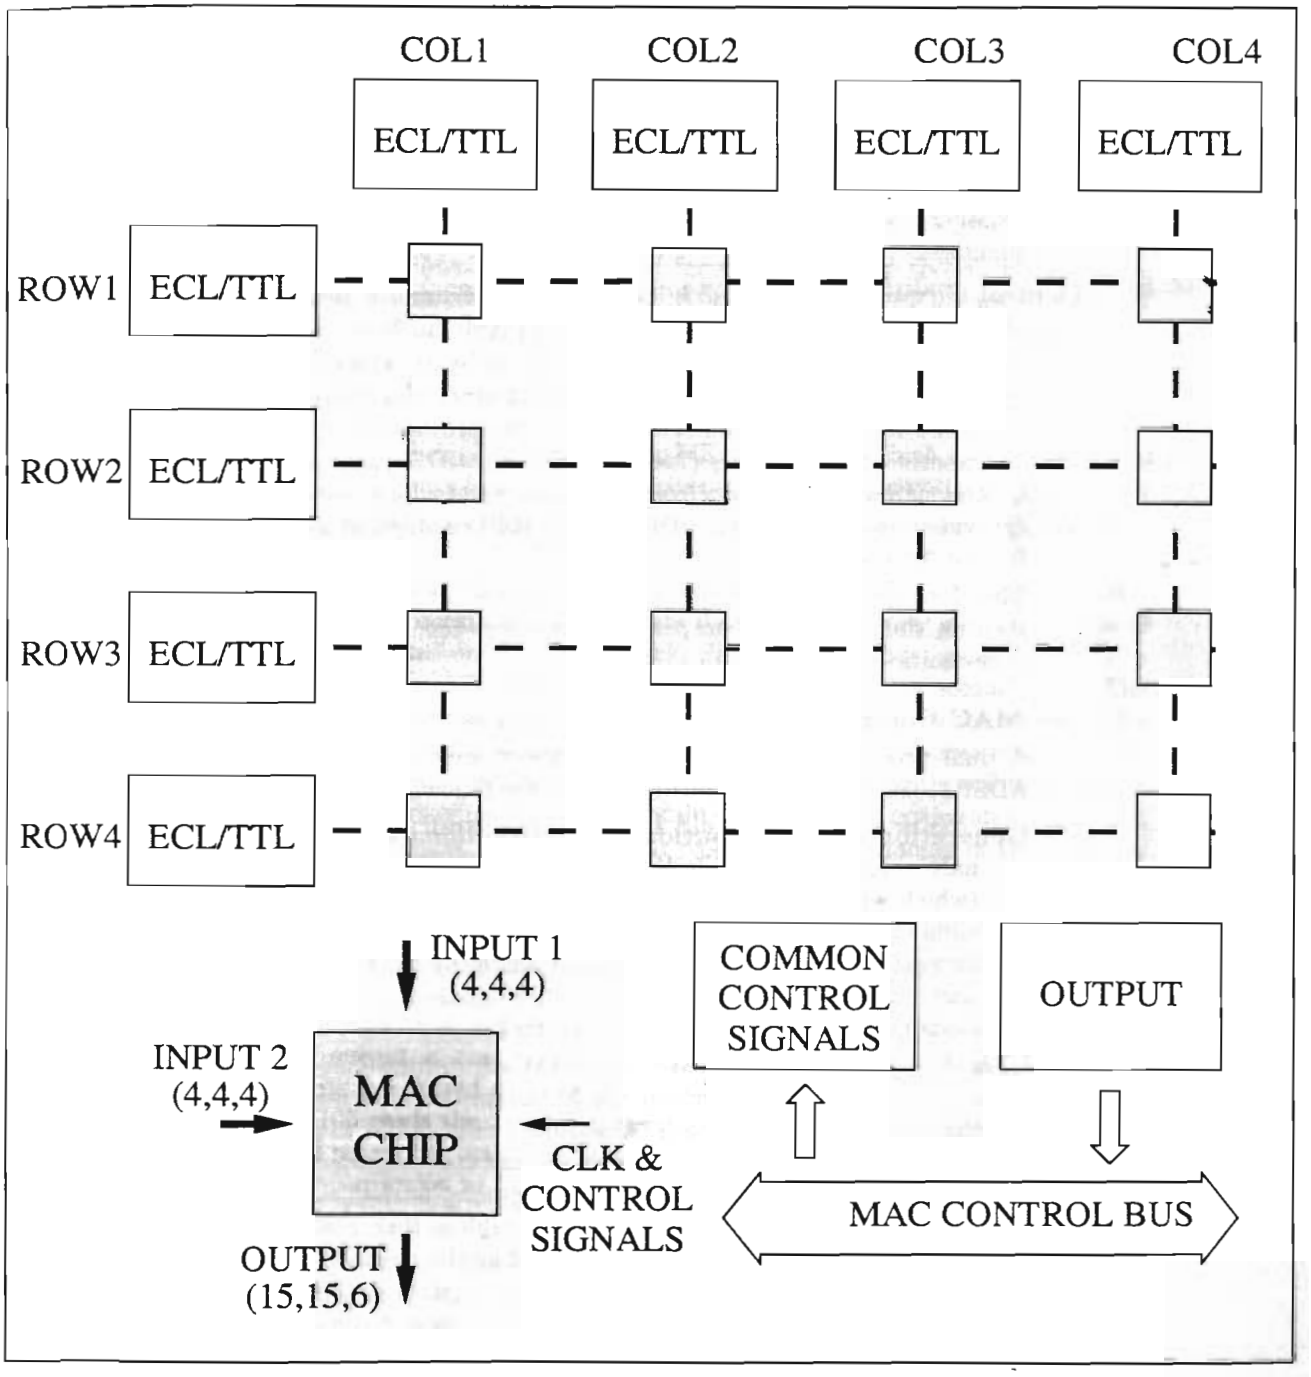

Figure 3.34: The MAC Card Block Diagram

mode is the default operating mode while the Non-polar is the most scarcely used mode. The polar mode will be used whenever calculation of Stokes parameters is required.

The multipliers operate in parallel, accumulating the products till they are read-out. The read-out is time-multiplexed, with all the chips on the same bus. A chip select enables each chip and an address provided to it reads out a result. Since the result is 36-bit wide while the output port is only 18-bits wide, the result is read out in two parts. The real and imaginary parts of the complex number are read in two consecutive words. The common exponent is split into two parts of 3-bits each and is combined with the 15-bits real or imaginary mantissa to form an 18 bit word. The LSB 3-bits of the exponent is combined with the real mantissa and the MSB 3-bits with the imaginary mantissa.

Control Circuitry This includes logic required for capture and distribution of clocks and the various controls signals required by the ASIC in MAC mode.

The clock is received as an ECL signal. It is then converted to TTL level (by a MC10125) and distributed to all ASICs and other logic chips by a pair of clock drivers (74ACT11208). Of the two clock drivers, one provides clocks only to the ASICs where as the other to the latches (74F821, 74F574, 74F175, 74F824 which capture data and control signals). The clocks from these two drivers differ from each other by an amount (provided by the delay line MDLDL-TTL-10F) required to provide sufficient setup time for the ASICs.

The data read from the the ASICs is latched at the card edge (by 74F823) before transmitting it on the backpanel. The cards too, like the chips, are read out in a time-multiplexed fashion.

2. The MAC Control Card

A dual processor card employing the fixed point ADSP2105 and the floating point ADSP21020. The MAC control performs the following tasks

Generation and Distribution of Control Signals The generation of the control signals is done using certain basic information such as number of cards in the subrack (which will be a constant in the field), the amount of STA to be performed, the number of channels to be read out, the kind of MAC mode required etc. This information is provided to the MAC Control by the CCC via the MCC. The MAC control then calculates the various other signals on the basis of this information and loads the RAMs appropriately. This task is performed by the ADSP2105.

LTA The STA output is read from MAC cards into the MAC control card. The MAC outputs are stored on the MAC control in a dual port RAM from which the ADSP2105 reads the STA numbers. The ADSP2105 then unscrambles [3] the outputs and passes the real and imaginary mantissas, as a 21-bit number, to a look up table for conversion to IEEE floating-point format.
The converted output from the look-up table is then read by the ADSP21020. The 21020 performs the specified LTA by adding these results(Σx_i). It can also compute Σx_i^2 to calculate the variance of the incoming data stream. The LTA output is read out into the rack control card which transmits it to the Data Acquisition System (DAS) [4].

The MAC Backpanel

[3]The STA output comes as two numbers of 18-bits each. The first contains 15-bits of real mantissa as 1's complement number as well as the 3 LSBs of the 6-bit exponent. The second number contains 15-bits of imaginary mantissa, also as a 1's complement number plus the 3 MSBs of the exponent. The 6-bit exponent is a 2's complement number. The MAC control puts the 6-bit exponent with each of the mantissas but leaves their formats unchanged

[4]A description of the present DAS is given in Chapter 5.

The MAC subsystem is divided into six subracks - one in each FX rack. Each subsystem accommodates 11 MAC cards and a control card within a subrack. A total of 192 x 11 signals are required to be accepted by this subrack from externally located FFT subsystems, and distributed to the 11 cards after some re-ordering of bits to suit the MAC card pin assignments. In view of the large number of signals, wire wrapping would have been very cumbersome and also difficult to maintain. Hence, a PCB was specially fabricated instead. In view of the large power dissipation (about 550W per subrack) to be accommodated, and in view of the complexity of signal routing, a 16-layer PCB was designed. Keeping in mind the relatively small number of such PCBs required in the system, a cost-effective solution was arrived at by making all the MAC backplanes identical.

The distribution of MAC cards adapted for this purpose is shown in Fig. 3.35. In this Figure, the numbers '1' through 'A' denote the 10 FFT cards. Numbers without a prime (') denote FFTs from first rack. Numbers with a single prime denote FFTs from second rack while two primes denote FFTs from the third rack. The resulting arrangement ensured that no cable needed to cross a rack to connect an FFT output to a MAC card. Also, the cable lengths were kept to a minimum, but of equal length. This set of three racks processes one of the two side-bands. The other side-band is processed by an identical set of three FX racks. Both sets of the three FX racks thus realise a 30-X-30 array of multipliers shown schematically in the Fig. 3.36.

It can be seen that second rack has the maximum number of cross correlators and also that most of its inputs are derived from the neighbouring racks. The backpanel design was thus based on the second rack, since it is the most generic one. The arrangement for first and third racks can be derived from the second rack. The Fig. 3.37 shows the actual 30-by-30 array realisation according to the arrangement of the backpanels. Ideally, all results should have occupied the same triangle. The mixture is the price paid for the optimisation. A few of the results will thus require a sign flip in their imaginary parts to obtain their complex conjugates so that all numbers can be thought of as originating from the same triangle of the 30-by-30 matrix.

3. The Distribution System

A careful look at the Fig. 3.35 reveals that a maximum of 9 inputs have to be sourced by each FFT output. For example, 5' (i.e. the fifth FFT card in Rack2) drives 8 loads - 5 in Rack1 (MAC cards 1, 5, 8 and 10) plus 3 in Rack3 (MAC cards 4,5 and 6). Similarly, 9 (i.e. the ninth FFT card in Rack1) drives 9 loads - 6 in Rack1 (MAC cards 3, 4, 7, 8 and 9) and 3 in Rack2 (MAC cards 3, 4 and 9). As a conservative figure, it was decided to allow sourcing of a maximum of three fan-outs per ECL output. This requires that each FFT output generate 3 copies to drive the 9 inputs required by the MAC array.

This was accomplished using a distribution system which is shown schematically in Fig. 3.38. The FFT outputs from the two adjacent cards are combined into one 50-pin FRC connector. This is done using a PCB which is an add-on to the FFT backpanel. The 12-output bits from both cards are interleaved with ground and transmitted to the distribution system. This system receives these signals via its backpanel and converts them to ECL levels. It uses two independent sets of TTL-to-ECL converters (MC10124) to produce two ECL copies of the same TTL signal. One is transmitted to the Pulsar Machine [5] via a 50-pin FRC connector. The other ECL copy is used to generate the three copies required by the MAC array. These are obtained by passing each of the ECL signals through a 1:3 ECL-ECL (MC10H116) driver. The 144 outputs thus obtained are then transmitted to the MAC array.

[5]This is a specialised hardware for supporting the observation of Pulsars

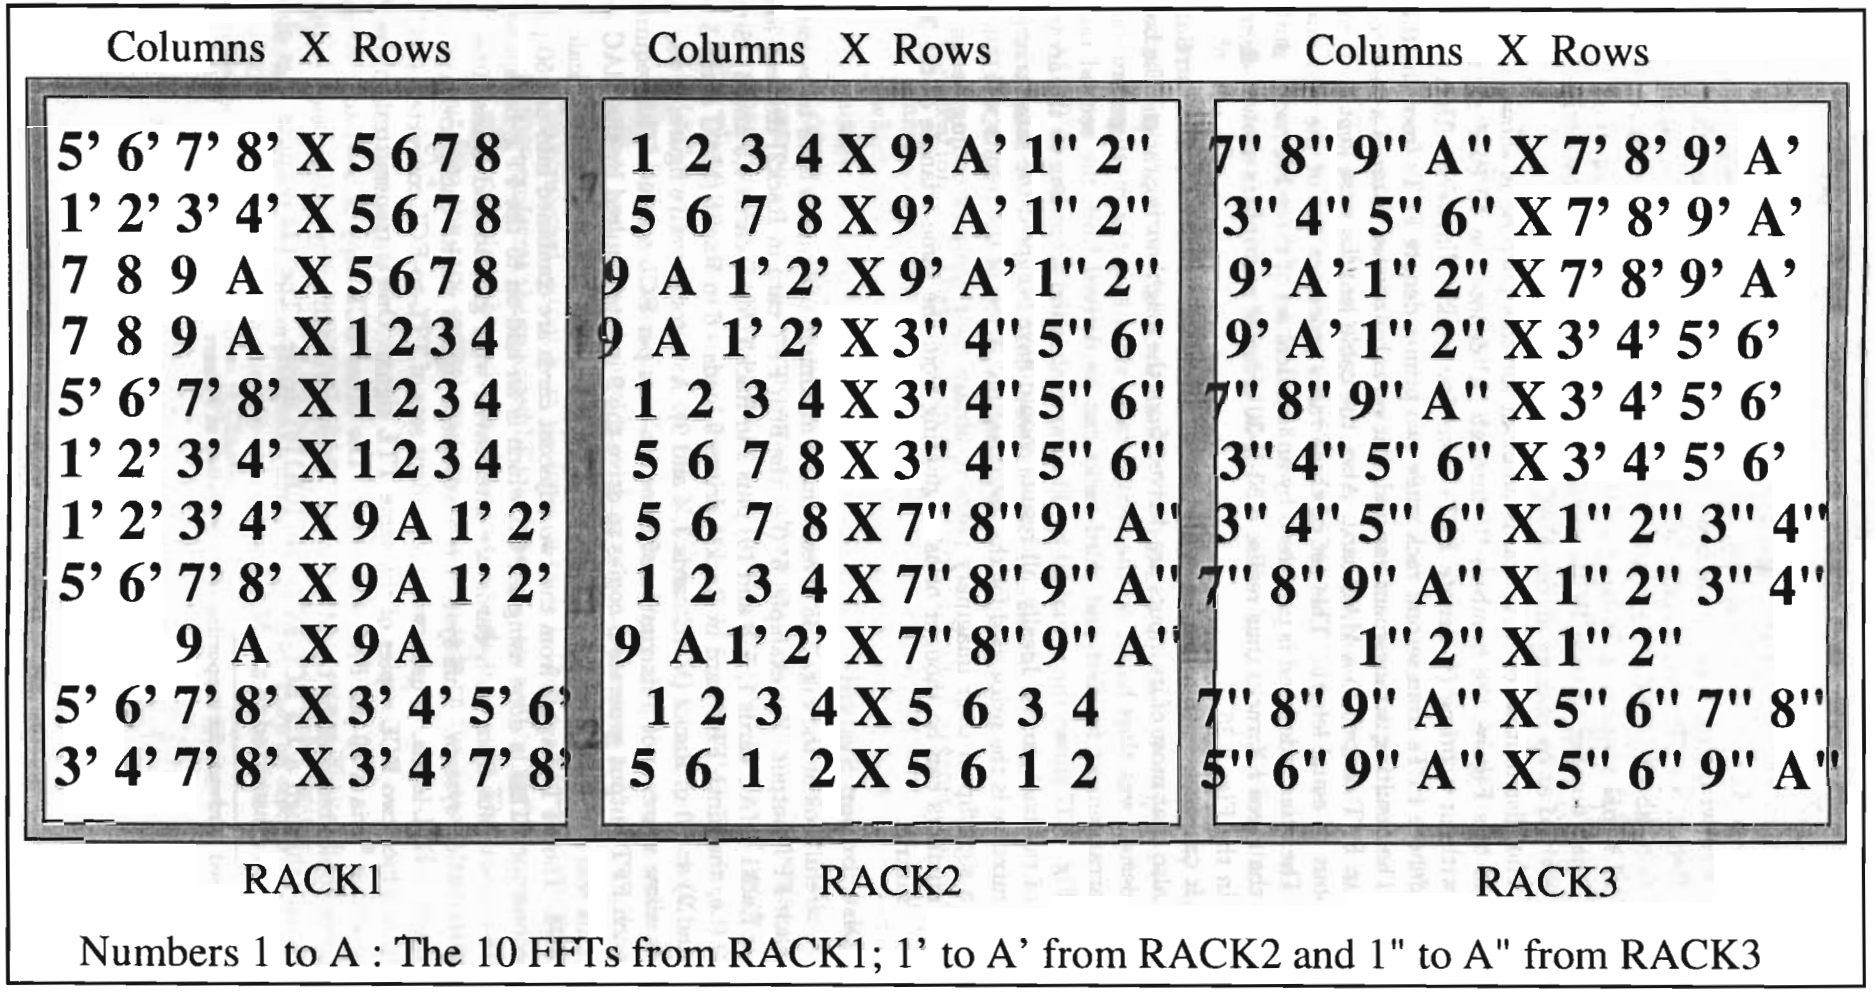

Numbers 1 to A : The 10 FFTs from RACK1; 1' to A' from RACK2 and 1" to A" from RACK3

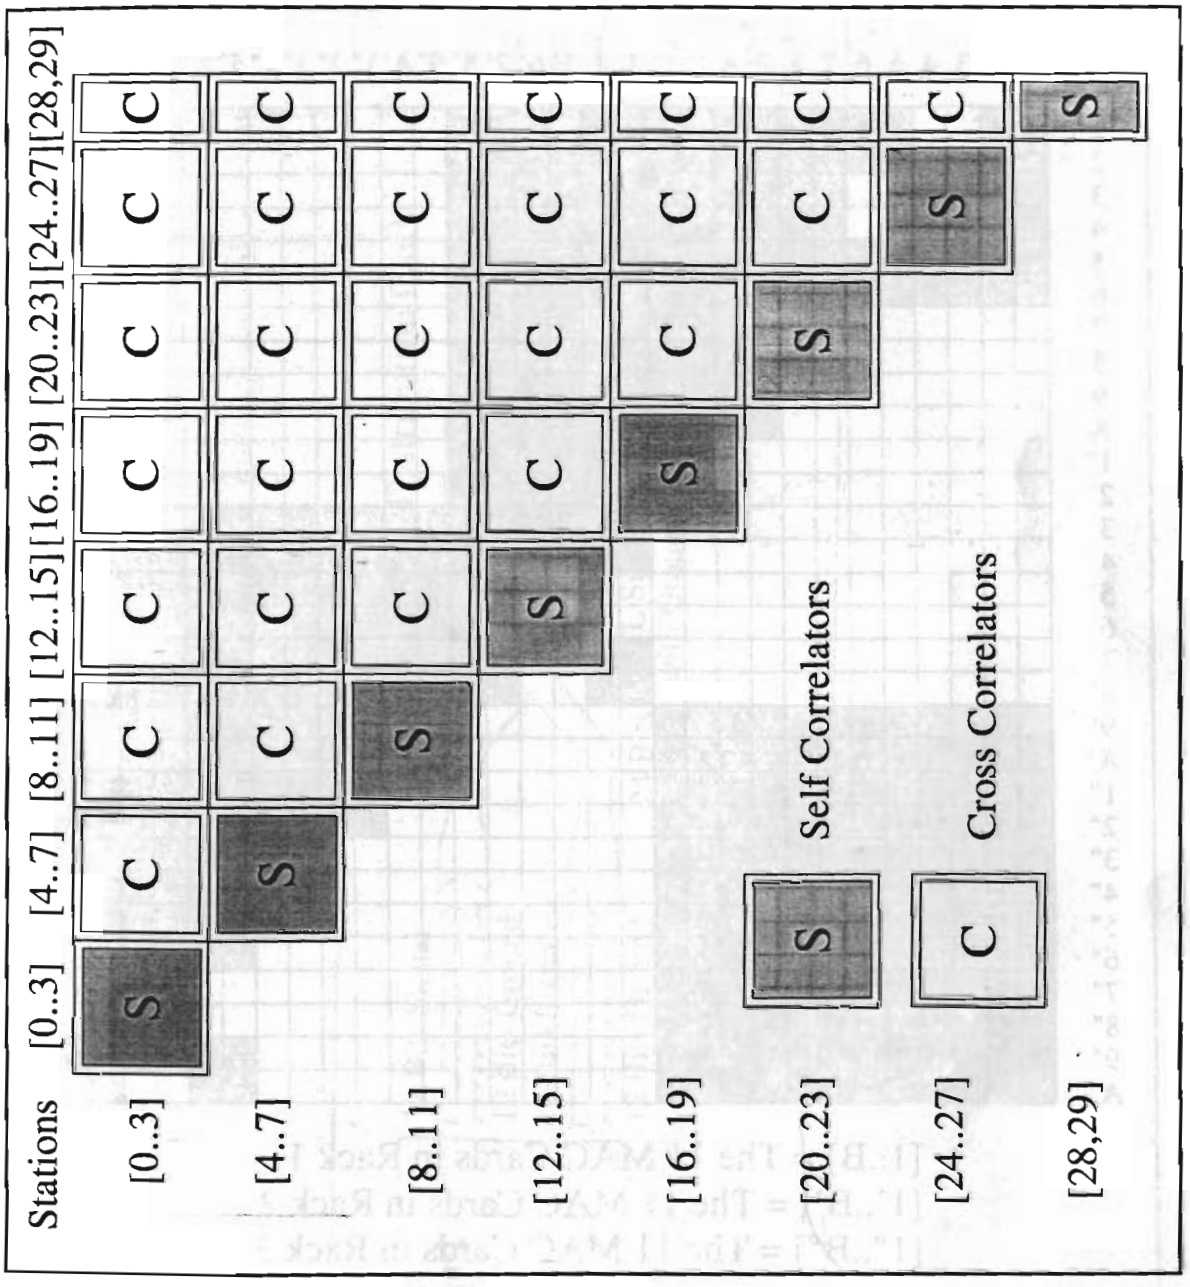

Figure 3.36: The 30-X-30 Array of Multipliers

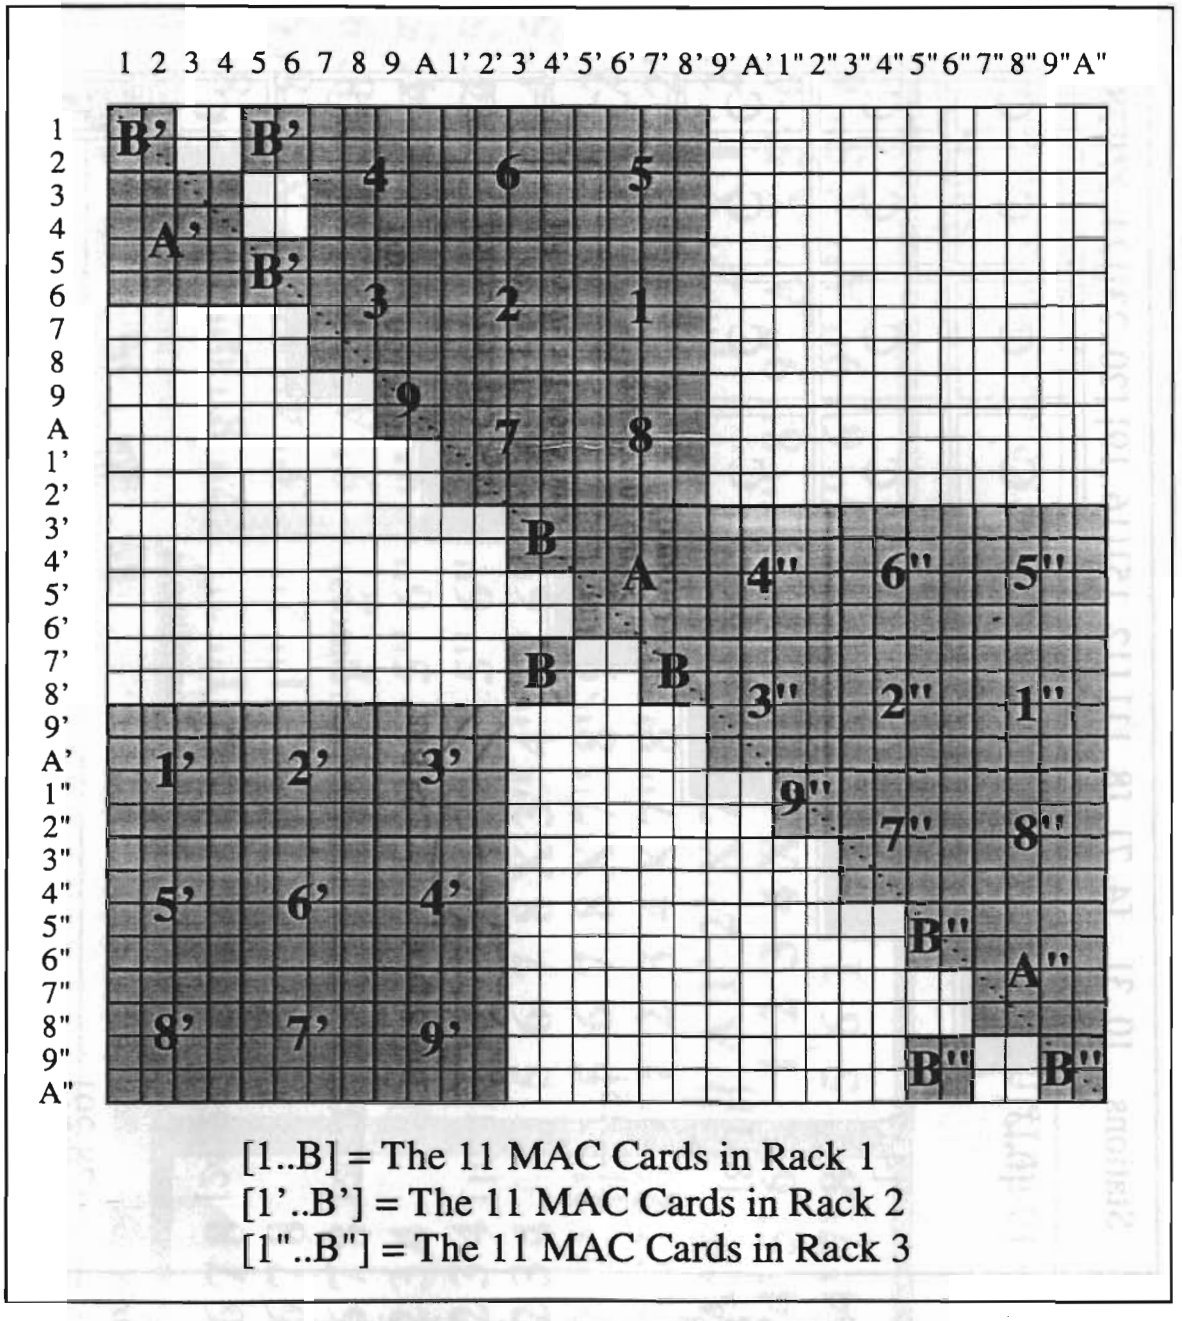

Figure 3.37: The Actual 30-X-30 Array Realisation

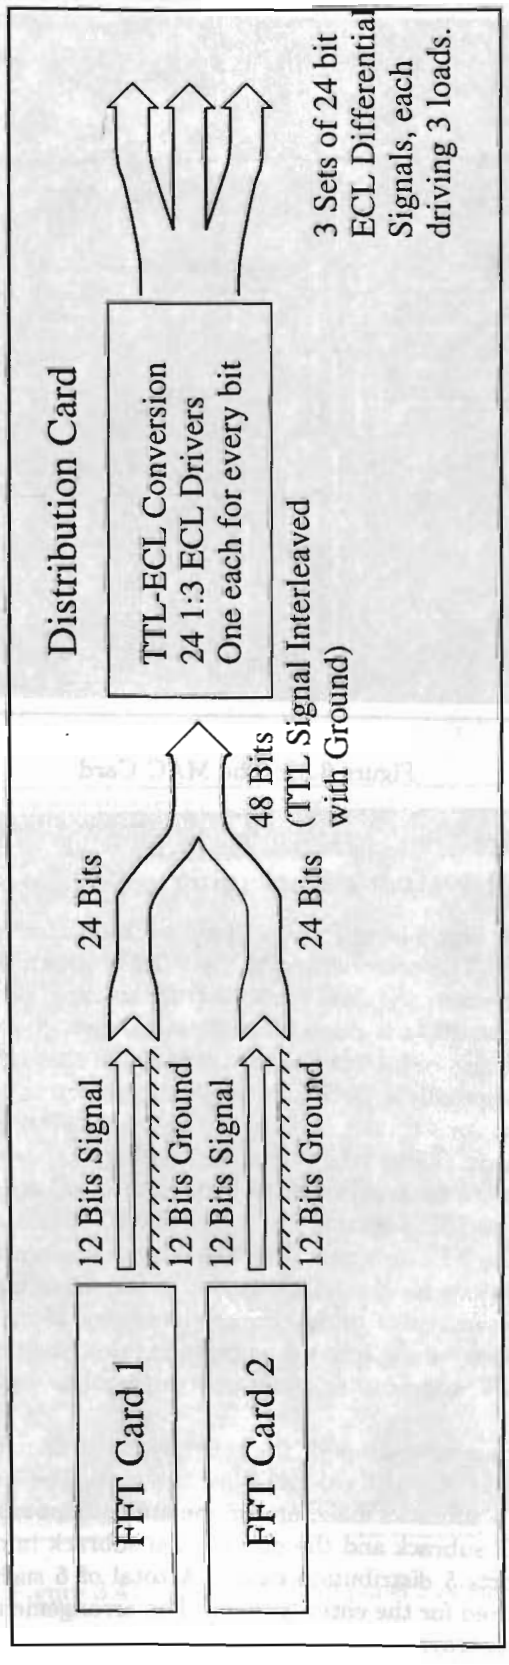

Figure 3.38: The Distribution System

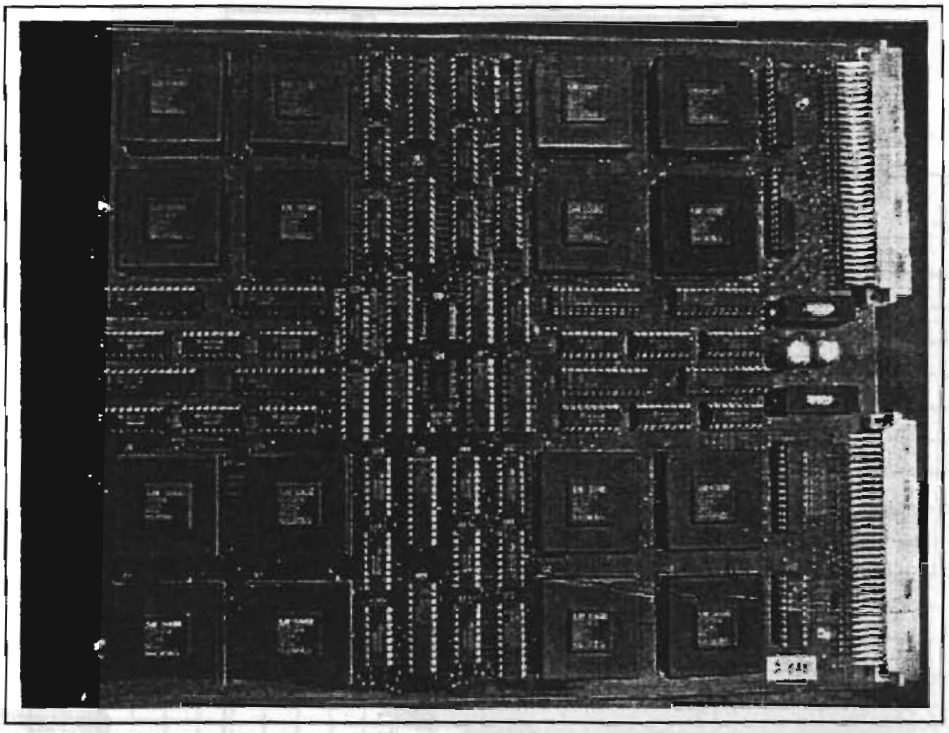

Figure 3.39: The MAC Card

3.7.2 The MAC Subsystem : Packaging

The MAC Card (Fig. 3.39) is a 6-layer PCB, 233mm by 280mm in size, 1.6mm thick with 35 μ copper plating on all layers. The construction of the PCB is shown in Fig. 3.40. The MAC card is designed with 8mil track width and 8mil track-to-track spacing. The via size used is 16mils.

The MAC Card uses two 160-pin Euro Connectors (5 rows X 32 pins) to connect with the backpanel. Both, the input and output to the card is through these 320-pins.

The MAC backpanel is actually a *Backpanel system* as shown in the Fig. 3.41, built around a 3.2mm, 16 layer PCB whose construction is shown in Fig. 3.42. The backpanel system is made up of a set of four different cards. Three sets of smaller auxiliary cards (Fig. 3.43) plug into the 16-layer PCB which is used as the backpanel for the MAC and MAC control cards. Two out of these three are essentially SMT-to-DIP adapters for the 50-mil X 50-mil AMPMODU Surface Mount connectors which receive the FFT outputs. The third is a termination card carrying line-to-line and line-to-ground terminations for the ECL signals. In all, 23 of the SMT-to-DIP adapters are used on the backpanel. Seventeen of these occupy the centre of the MAC backpanel Fig. 3.44 while the other 6 plug-in *horizontally* into the upper and lower ends of the backpanel (Fig. 3.45). The termination card too, like the adapter, plugs-in orthogonal to the plane of the backpanel (Fig. 3.45).

A total of 66 MAC cards are required for the entire system. Each backpanel supports 11 MAC Cards in addition to one MAC Control card and one free slot. The inter-card spacing is of 6T (= 1.2inch). Thus, a total of 6 subracks make up for the entire requirement. Each MAC subrack is housed along with the FFT subrack and the distribution subrack in one of the 6 FX racks. Each distribution subrack supports 5 distribution cards. A total of 6 such subracks housing in all 30 distribution cards are required for the entire system. The arrangement of various subracks as well as the power-supply and fan trays are shown in Fig. 3.46.

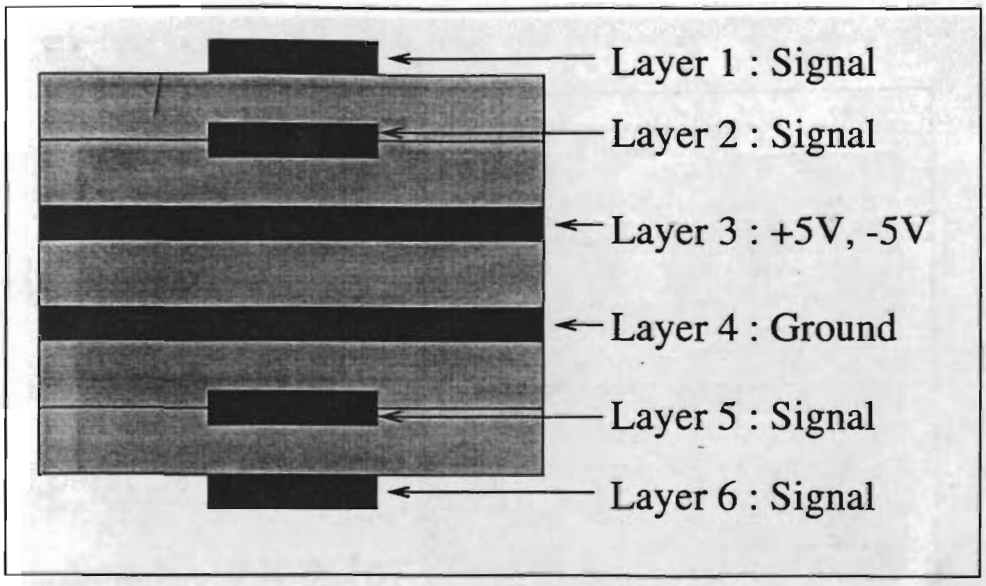

Figure 3.40: The MAC Card : Layer Assignments

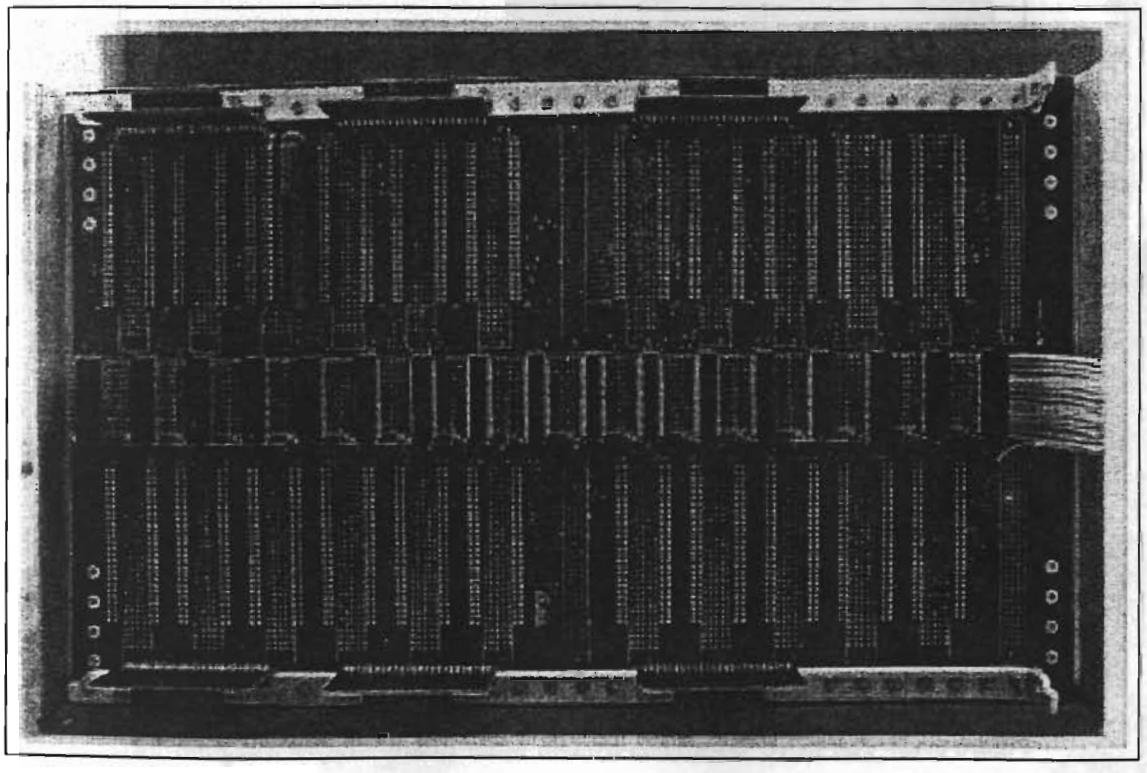

Figure 3.41: The MAC Backpanel System

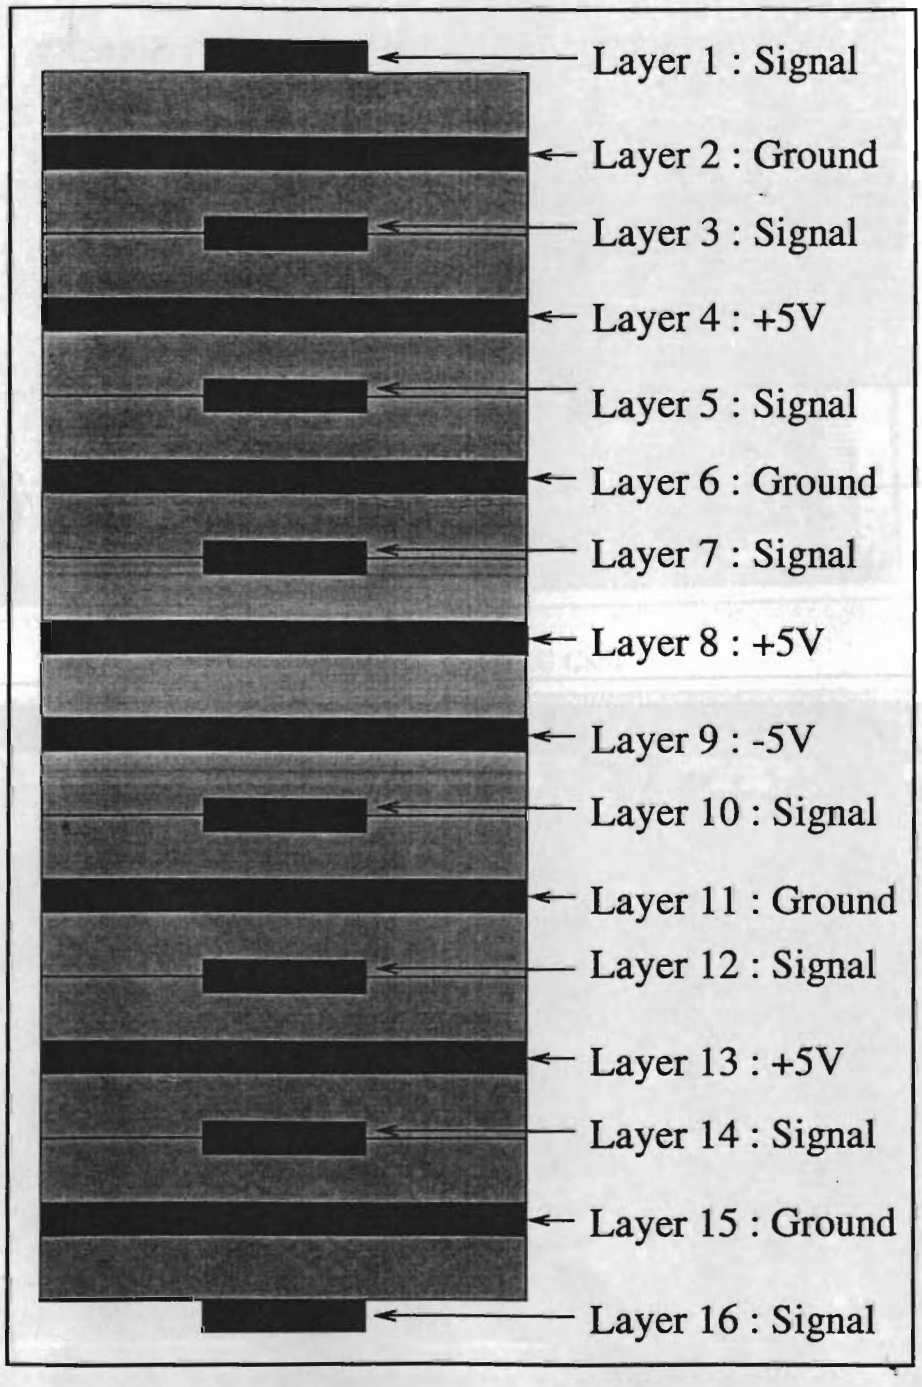

Figure 3.42: The MAC Backpanel System Main Board : Layer Assignments

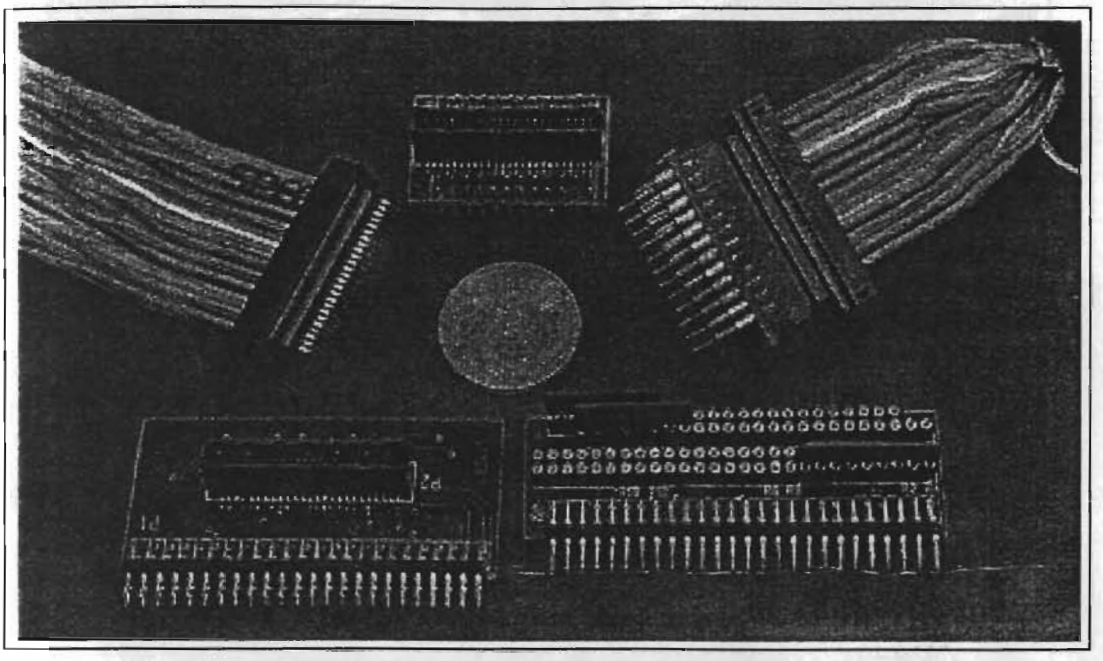

Figure 3.43: The MAC Backpanel System : Auxiliary Cards

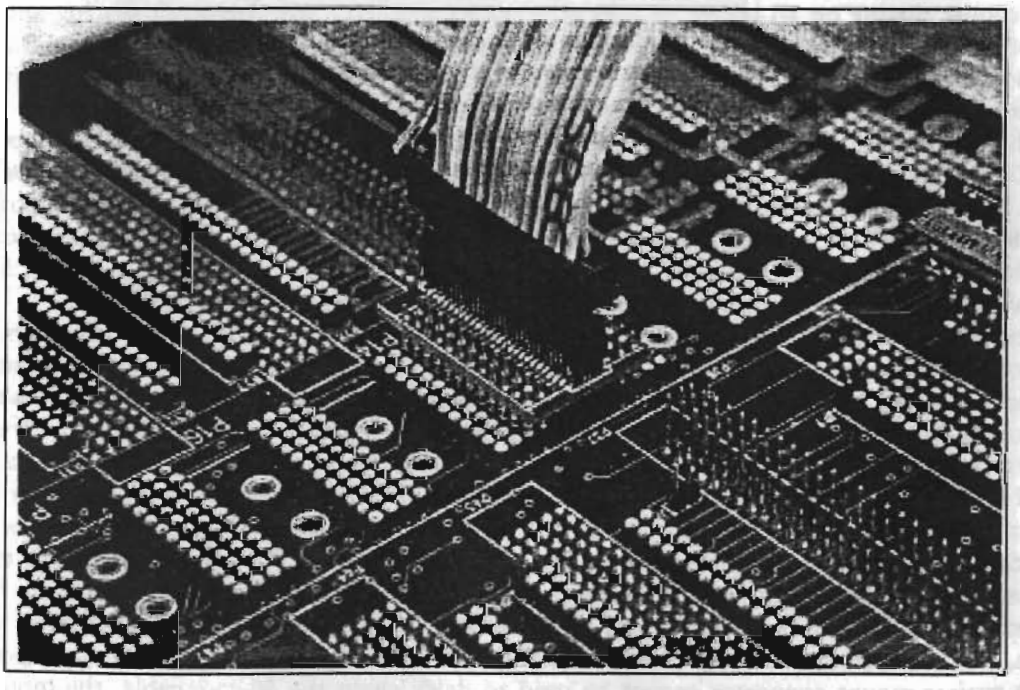

Figure 3.44: The MAC Backpanel Auxiliary Cards (SMT-to-DIP Type I)

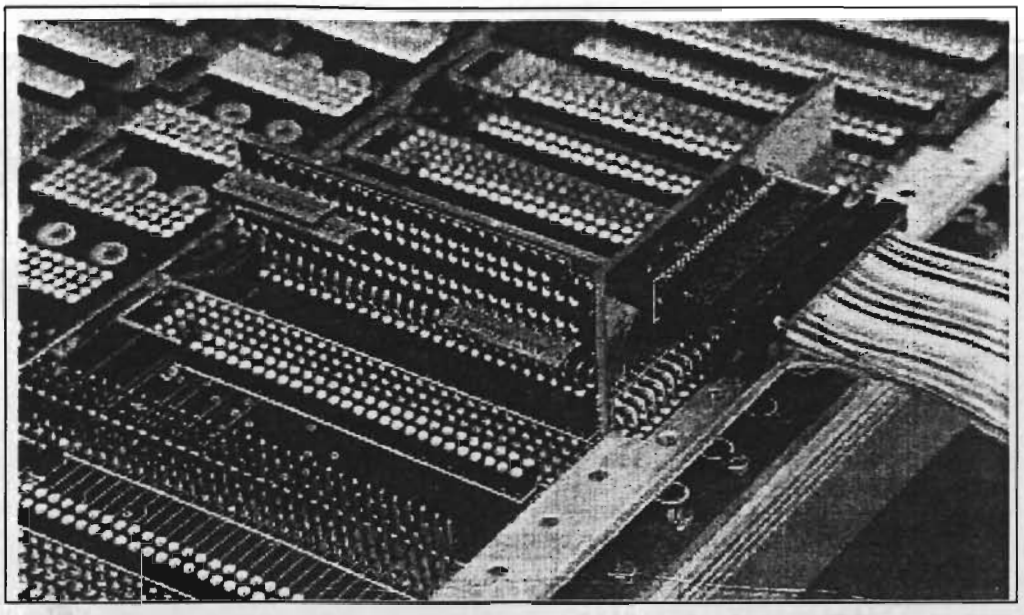

Figure 3.45: The MAC Backpanel Auxiliary Cards (SMT-to-DIP Type II)

The Distribution System

This system consists of the translation card and its backpanel.

The translation card (Fig. 3.47) is a 8-layer PCB, 100mm by 280mm in size, 1.6mm thick with 35 μ copper plating on all layers. It uses 8mil tracks with 8mil track-to-track spacing and a via size of 16mils. It connects with the distribution backpanel through a 96-pin Euro connector. In addition it uses two 50-pin FRC connectors to connect with the backpanel. Another 50-pin FRC connector, which is placed in the front of the PCB, connects with the Pulsar machine.

The distribution backpanel is a 3.2mm, 6-layer PCB (Fig. 3.48). It receives FFT inputs on the 50-pin FRC connector and uses 3 surface-mount connectors to carry the outputs to the MAC array. These surface-mount connectors are mounted on the backpanel using the same adapters as used by the MAC backpanel. Two additional 50-pin FRC connectors bring in signals from the distribution card. The backpanel is, in fact, a set of 5 independent 'distributors' - one each for a pair of FFT cards. These 5 share only the power and ground planes of the backpanel.

The FFT to MAC connections are established via the distribution system in two ways. The FFT output is connected to the distribution input by 50-core flat cables. The distribution subrack connects with the MAC by 50-core round twist-n-flat jacketed shielded cables. The inputs to the rows of the MAC cards go directly from the distribution subrack to the MAC subrack. Here both ends of the cables have the 50-pin surface-mount cable connector. This connector connects with one of the 6 auxiliary cards which plug in horizontally to the MAC backpanel. On the other hand, the inputs to the columns of the MAC cards go from the distribution subrack to an additional card. This card (Fig. 3.49) accomplishes the 1:3 fan-out for each input. For example, referring to Fig. 3.35 one of the inputs of 5' (out of the three generated by distribution subrack) going to Rack1 connects to MAC cards 1, 5 and 8. All these three MAC cards need to be on the *same* cable. Hence, the requirement to daisy chain these inputs. Since, the cable connector corresponding to the surface-mount connector cannot be used to daisy chain the 50-core cable, the inputs are brought to this card from the distribution subrack. Additional cables connect this card to the MAC backplane.

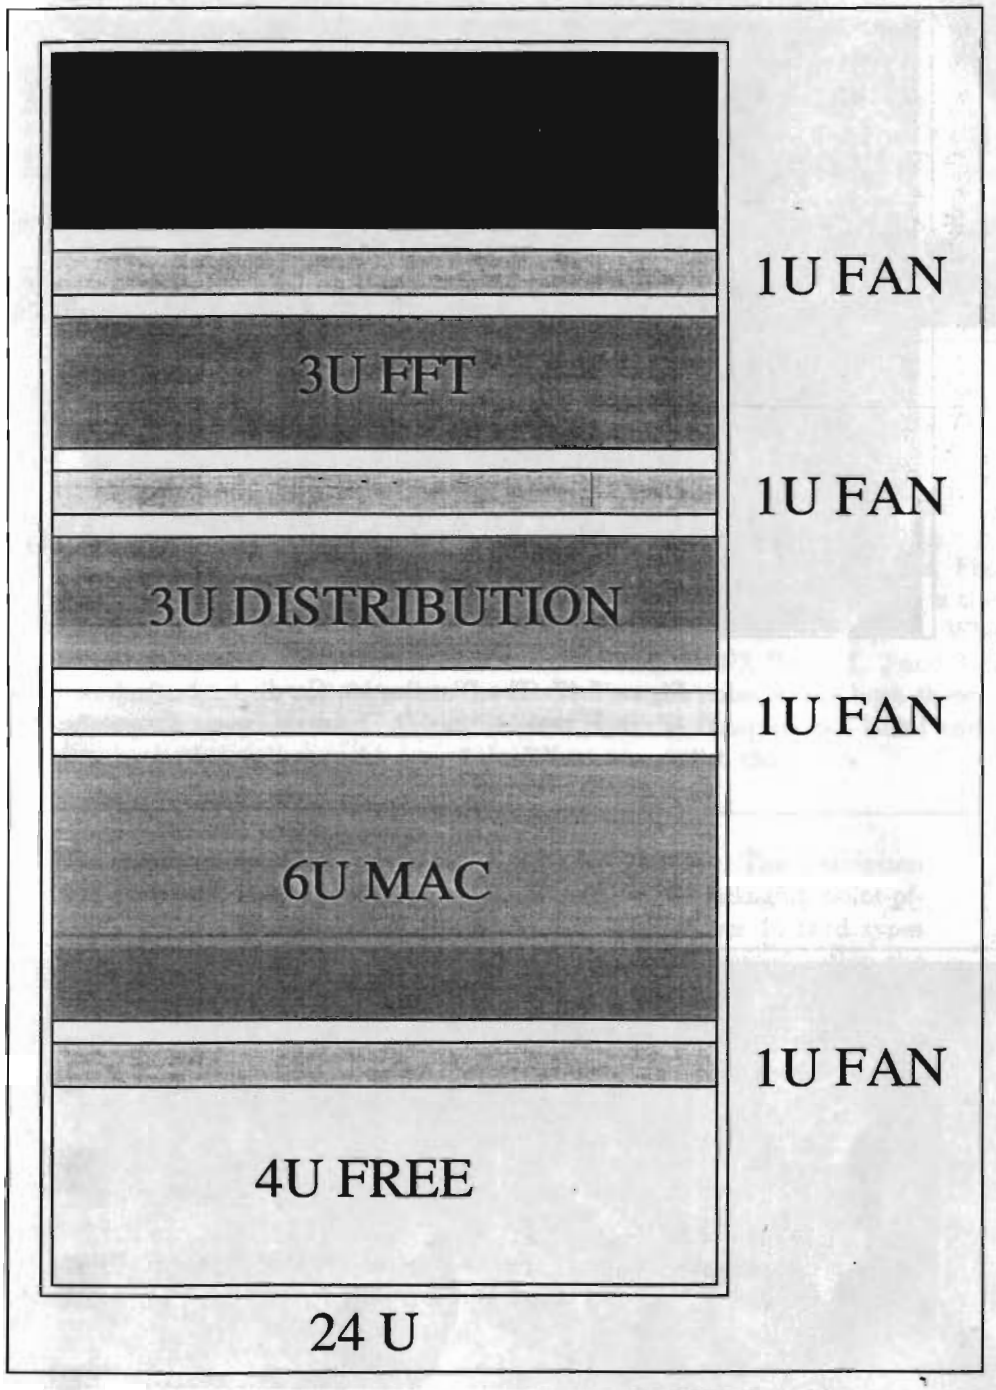

Figure 3.46: The FX Rack

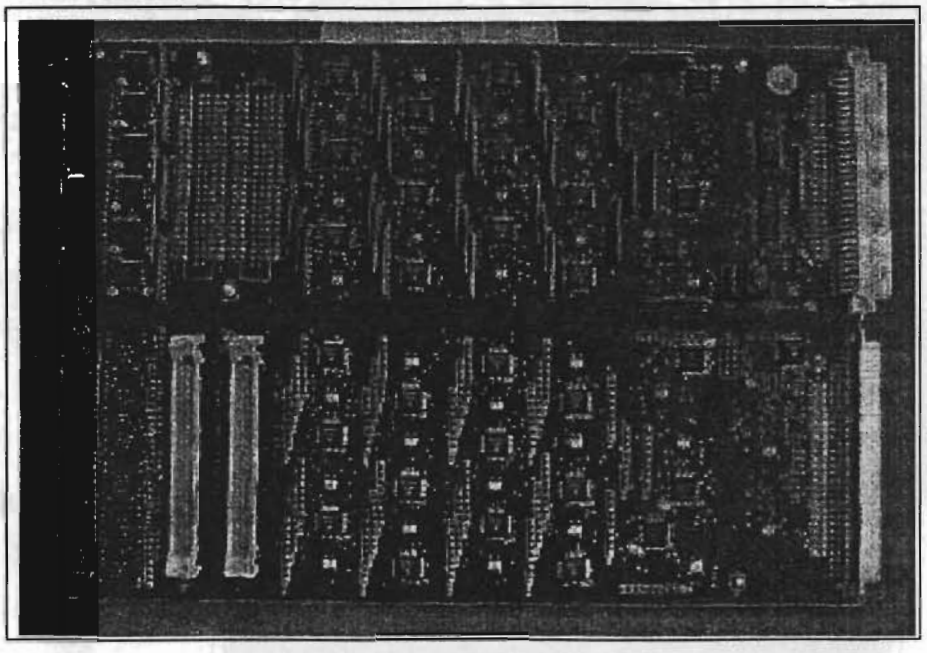

Figure 3.47: The Distribution Card

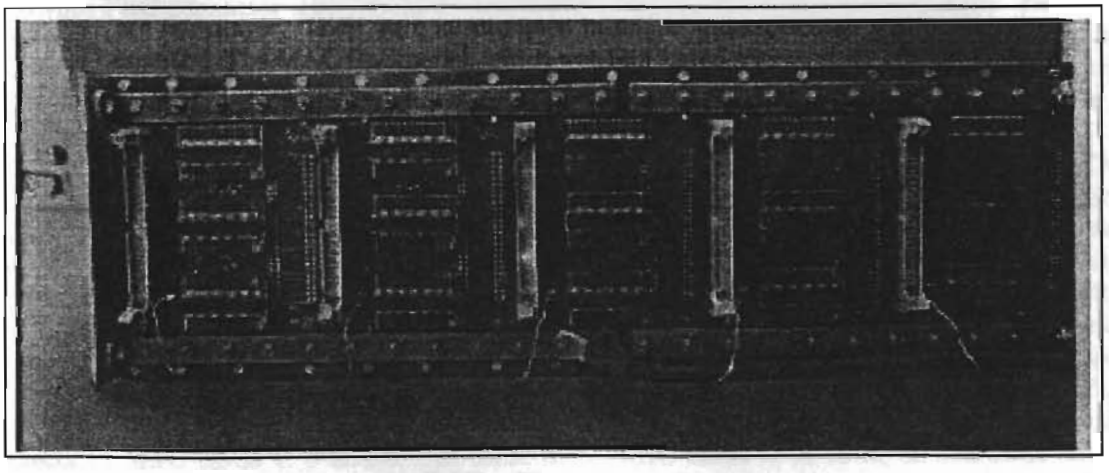

Figure 3.48: The Distribution Backpanel

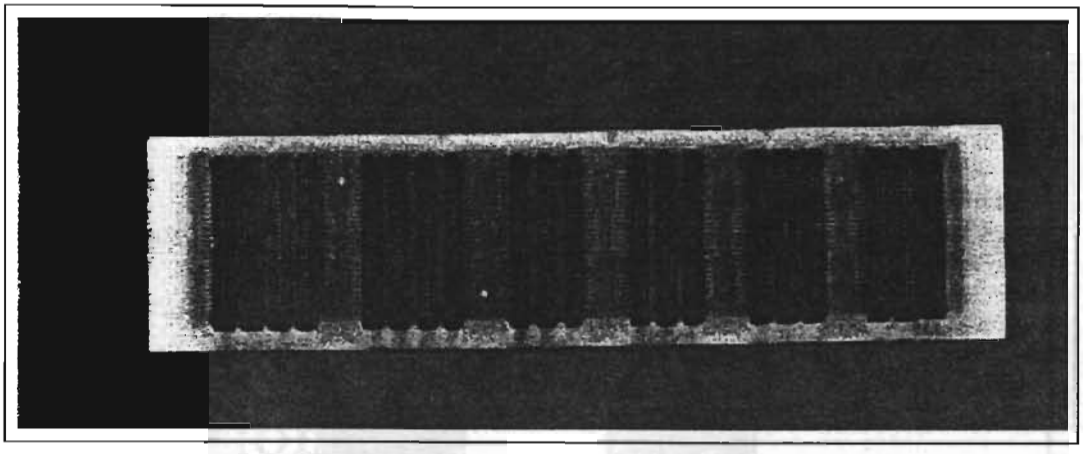

Figure 3.49: The 50-pin FRC Card for Daisy Chain

3.8 The Correlator

As was described in section 3.1.3, the entire correlator system is packaged in 10 racks. Fig. 3.50 shows the layout of these racks with their interconnections. The block labelled GAC is the front end of the pulsar machine. The ADC power-supply rack is excluded from this figure. When the MACs are operating in either Non-polar or Indian Polar mode, the FX Racks 1, 2 and 3 process one sideband while Racks 4, 5, and 6 process the other. For, the polar mode, both these sets of three racks process the same sideband. As can be seen, both the Delay racks - DLY1 and DLY2 are connected to both of the three-racks-sets of the FX to accomplish this.

The design of the GMRT correlator system was described. The description was presented from the circuit design as well as the packaging point-of-view. As can be seen, all of the about 320 cards of over 15 card types are multilayer. They were all designed in-house and fabricated within the country. The MAC Backpanel, which was fabricated overseas, stands a sole exception, and its reason is discussed in Chapter 5. The large number of interconnections required to realise the 30-by-30 array is one of the major contributing factor to the complexity of the system.

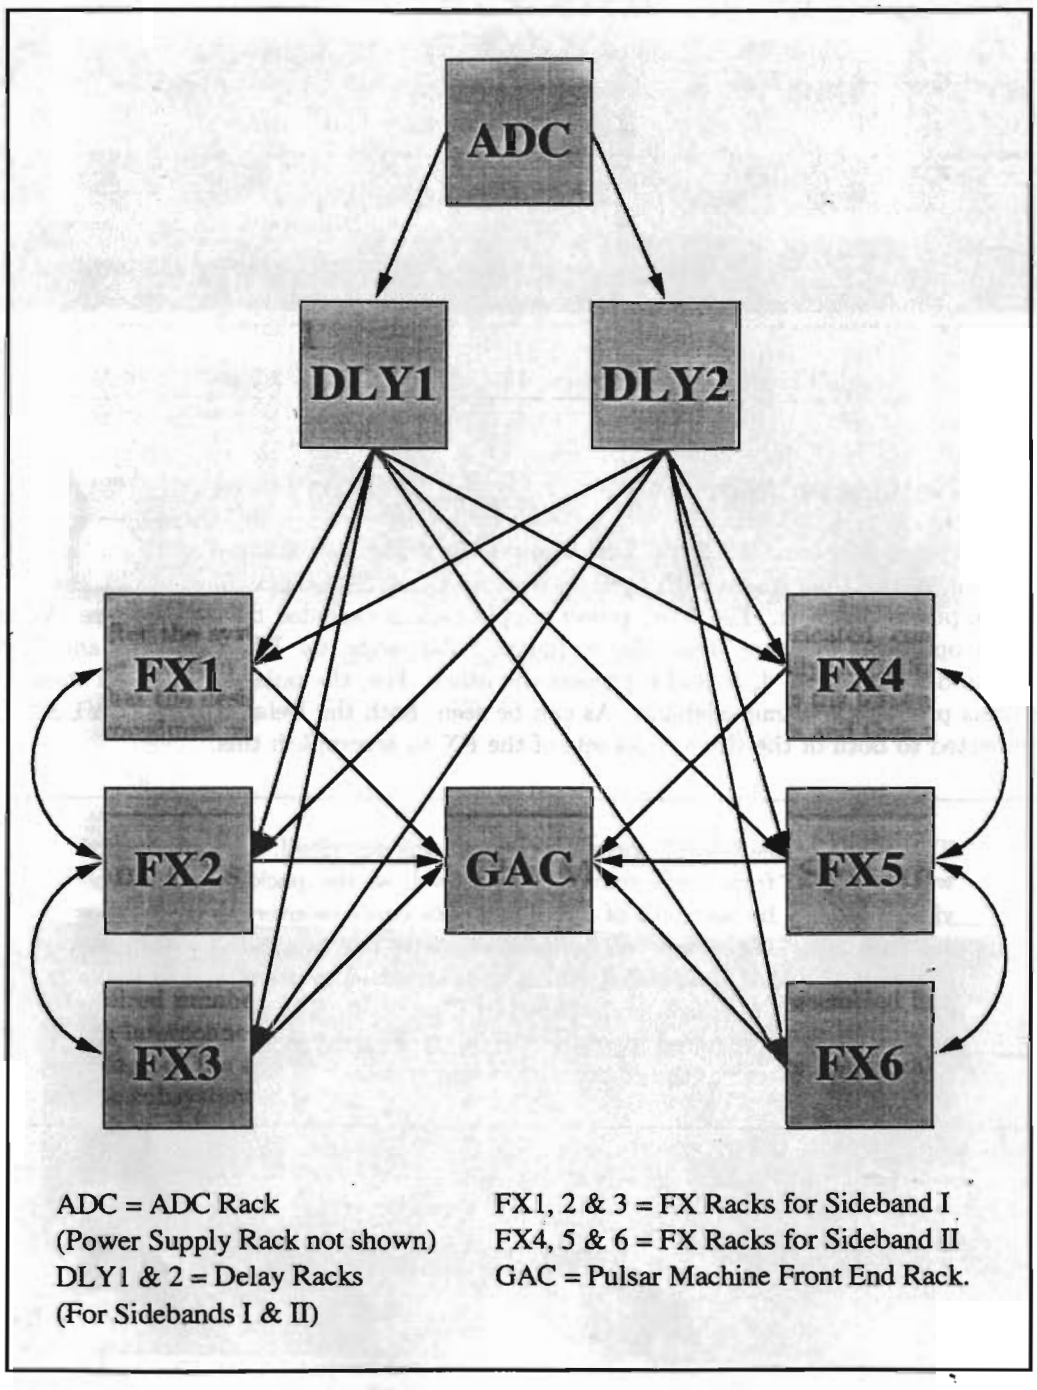

Figure 3.50: The Rack Layout of Correlator System

Chapter 4

Validation, Testing and Integration

'Tis strange but true; for truth is always strange, - Stranger than fiction

–Byron, 'Don Juan'

> After the system and its constituents are designed and fabricated, comes the question of validating and verifying whether they actually do perform what the design intended them to. This chapter shall discuss the tests and procedures used to validate first, the GMRT correlator cards and then the entire system

4.1 From IC's to Racks : The Making of a System

The validation of the system is done following a bottom-up approach. First, each card of all sub-systems is tested out. These cards are then assembled to make a sub-rack (also called a 'bin'). Once the required number and type of sub-racks are tested, they are assembled in racks. These racks are then interconnected together to form the complete system.

At each stage of the assembly, certain tests are performed to ensure functional integrity and stability of the subsystem.

Functional tests Given a certain pattern as the input and the function of the device-under-test, be it a card or a sub-rack or the system as a whole, the expected output can be simulated. The Validation of the device-under-test then implies that the actual output be identical to the expected one.

The stability and repeatability Once the correctness of the output is ascertained, the next requirement is that the output be repeatable and stable over extended periods of time.

Repeatability implies that the same output pattern should be obtained in response to the same input (data and control input). This is ensured by power cycling the system, by reloading the control inputs, by cycling the input pattern, etc.

Stability implies the output be identical over extended period of time (a few hours to days). This ensures that the system is devoid of phenomena such as reflections in the transmission of data, jitter, grounding, etc.

For the functional tests, simulation of the FFT and MAC operations were performed to validate the output of the FFT card and the output of the ASIC in MAC mode. For the ADC card, the expected dynamic range of the output was also obtained in a similar fashion. For the stability tests, generally the method adopted was to form a template of the output by taking a "snapshot" and then compare the output with this template.

The tests themselves have been continuously evolving for the past six years using the experience gained on the test benches and the feedback from the field. These tests are thus a result of the collective experience of the correlator team. The tests mentioned below have been selected by the author to provide an overview and are not necessarily exhaustive.

Th tests can be classified, as outlined above, into three groups, viz.,

- Card Level Tests
- Subrack Level Tests
- System Level Tests

4.2 Card Level Tests

4.2.1 General Tests

Tests independent of the card type can be put under this category. Some of these are checking the connectivity of various tracks/vias (via is a term used for PTHs, *P*lated *T*hrough *H*oles) and checking out power/ground shorts.

A bulk of this category of tests is done at the PCB factory as a part of the *Bare Board Tests (BBT)*. Most of the connectivity problems are weeded out at this stage since they mostly arise due to defects in the PCB manufacturing process. It is to be emphasised here that the BBT are done as a matter of routine, only on the "production" cards (which simply means a large quantity). The prototype cards, which typically have quantities of less than 10 , do not go through the BBT. The BBT itself is of two kinds, one, using a "template" method wherein one of the boards of the lot forms the reference and the rest are required to be identical to it; and two, the "netlist" type wherein the connectivity is generated from the netlist. The second type, as is obvious, is a more fool proof method than the first. There is a caveat to the BBTs : they are done only on the through-hole section of the boards. The surface mount section of the board, if any, evades the BBT.

The post BBT defects in the cards are then mostly a result of the component assembly process. Both techniques, manual and automatic assembly, have been used on the Correlator cards. Manual assembly has been used for cards with lesser quantities and with cards which need to serve as templates for the automatic assembly plants. All control cards come under the former category because of their small quantities. The automatic assembly techniques have been used on the ADC, Delay, FFT and MAC cards. A template for each of these was manually assembled and sent to the assembly plant. The major defect in these class of boards has been the "dry-solder" - a term used to signify an unreliable contact between the device (socket) pin and the PCB pad, arising due to improper soldering.

Of the two kinds, the dry-solder defect has been observed more than the other. However, the combined frequency of occurrence of both these kinds has been rare. And, keeping in mind the limitations in resources including person-power and other logistics, a practical compromise followed, which was by and large of the "go-no-go" kind, that is, populate the assembled card with components and subject it to the test pattern. The boards which clear the test are cleared for subrack-level tests and the ones which fail are then looked into one by one. The defects described above were discovered at this stage.

4.2.2 Tests for ADC Cards

There are mainly two aspects which are tested out in the ADC cards, the first being the conversion *per se* and second, the DC offset of the analog signal. These are the basic tests that are performed to clear a card. There are a few other tests, one of which is to test the coupling between the channels, which have been performed on some cards mainly to ascertain the design limitations. These tests are done only on few cards and are not done as a matter of routine certification of cards for use. Such tests are described later in this sub-section.

The conversion is tested by injecting a zero mean sine wave from a function generator and following it through the entire path - the amplifiers, the ADC, the latches and the TTL-ECL converters. This tests out all the components and also the performance of the card. A test jig is made with a ADC backpanel mounted on a subrack with the power supply and clock connected. The output is acquired from the backpanel connector on a logic analyser or on a specialised test bench. [1] The data acquired from either of these sources can be read on a PC where a FFT analysis shows the performance of the card in terms of the spectral dynamic range obtained.

The signal after passing through the amplifier may have a DC offset, a large value of which can degrade the performance of the FFTs due to spectral leakage. This offset is tuned out in two stages. The first stage is during the ADC card testing where the potentiometers in the power supply section and in the reference generation section are tuned appropriately to make the input signal appear as a zero mean signal to the ADC. This is achieved by tuning the absolute reference voltages at the two ends of the ADC reference ladder. The final voltages obtained at the two ends of the ladder are, in general, asymmetric to compensate for the DC offset of the signal. This stage of tuning is done using an oscilloscope by observing the MSB of the ADC output. For a zero mean signal, the MSB should be a square wave (i.e. a duty cycle of 50%). The second stage of tuning is done after the entire processing chain is integrated. The amplitude of the DC channel of the auto-correlator chip corresponding to the ADC card being tuned, is observed. Further tuning, if required, is done so as to minimise the DC channel amplitude. This is done, since it has been empirically determined to be both, more sensitive as well as easier (as compared to the oscilloscope method).

4.2.3 Tests for Delay/DPC Cards

The card-level test for the Delay-DPC cards is to inject a pattern from the Delay Control card into the DPC card and trace its flow through the logic of the two cards upto the output of the Delay card. The test pattern, viz., a sequence of numbers 0,1,2,4 and 15 (0,1,2,4 and F in hexadecimal), is chosen so as to excite all the four bits of the Delay-DPC system. This tests out the DPC section. The test pattern being a time-series, can test out the Delay card by creating a relative shift between outputs of various channels.

It has been observed that these tests are sufficient to ensure a simple yet complete test for the Delay-DPC system.

4.2.4 Tests for FFT Cards

The card level tests for the FFT card can be classified into the ones that test out the main board and the ones that test out the pipeline board. The tests for the main board check out all the latches and RAMs on it. This is accomplished by a set of test programs running on the FFT control card. The "pipeline-emulator" card can be used to emulate the NCOs in order to address the RAMs on the main board. After the main board is tested out completely, the pipeline board is connected to

[1] This test bench was made as a student project for testing of the FFT cards. It acquires the FFT card outputs on an acquisition card which has 128 KB of RAM. After the RAM is filled completely, it is read by a PC. After completion of the reading the PC puts the card in acquire mode and the cycle goes on as required. The arrangement of the acquisition card is fairly generic and for testing the ADC cards, their output is connected to the acquisition card input.

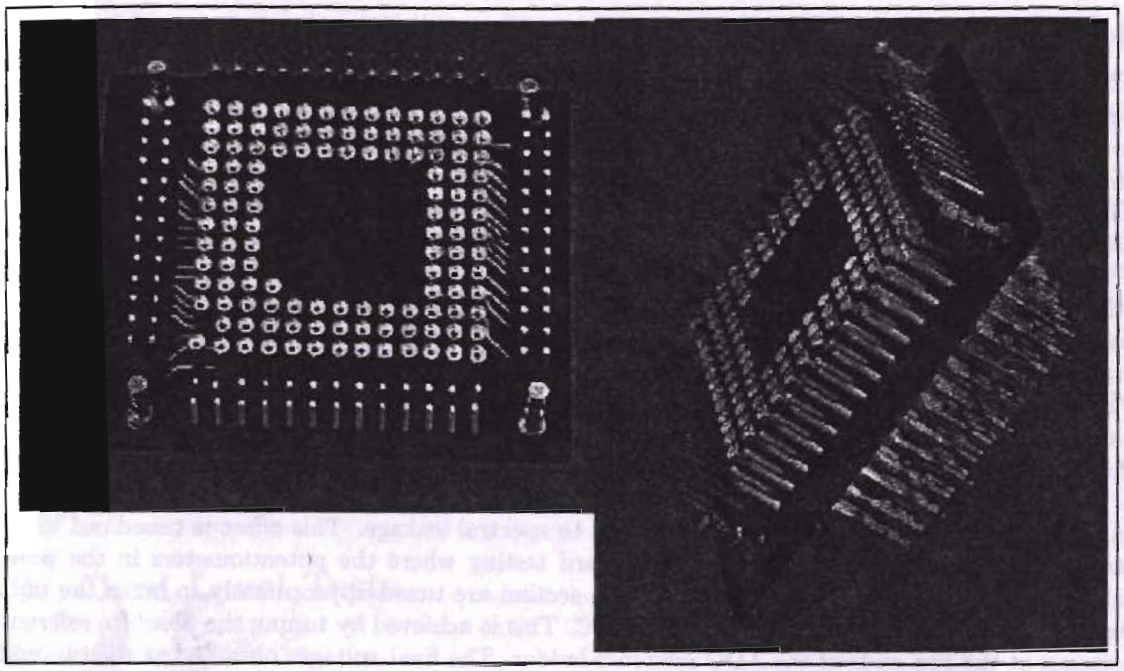

Figure 4.1: Test Grid for the ASIC

it. The signals reaching the ASICs can be tested out on the oscilloscope before putting the ASICs in their sockets. These signals obviously do not include the ones which require the NCO to be operational. In order to test the card with the ASICs operating in their configured states, a special test grid (Fig. 4.1) has been made. The array of pins surrounding the PGA socket is provided for connecting logic analyser pins. To facilitate the testing procedure, a set of test points is available on the pipeline board. The input and output of each stage from both pipelines is brought out on these test points.

Mainly two patterns are used for the test. The first one being a constant input, which just means a DC signal. The output should be a single pulse. This tests out the correctness and stability of the trigonometric tables, the final stage external readout address and the various stage INITs. Incorrectness of the trigonometric tables can stem from either an incorrect value or incorrect timing of the trig table sequence with respect to the addresses and INITs. Instability in the trigonometric tables can arise due to improper latching resulting in a jitter. The same holds for the external addresses and INITs too. The DC test, however, does not test out the timing of the input external addresses. This is so because, as long as the input addresses cycle through all address values, the output will be correct. The addresses, thus, need not be in the correct order. To test out the input address timing, the inverse of the above test is performed. A pulse is fed at location zero of the time series to obtain a constant or DC at the output.

The next test is to input a digitised sine wave to the pipeline. The sine wave values are stored in the fringe rotator ROM (look-up table) and the NCO is programmed as a 32MHz counter. The output of the FFT card should be a spike at the corresponding frequency channel. If 'X' cycles of the sine wave are programmed in one FFT cycle, the output spike will be in the 'X'th. frequency channel. Since, the input is generated totally from within the FFT card, any instability of the output indicates improper latching of the trig tables, INITs or external addresses. This can be resolved by properly phasing the FFT card clock with respect to the FFT Control card outputs.

4.2.5 Tests for MAC Cards

The tests for the MAC cards are designed to test for faulty components and to test out the correctness of the timing relationships between various control signals. The tests consist of two kinds of static (constant) patterns and one pattern of the non-constant type. The two constant patterns are the complementary hex numbers 0x55 and 0xAA. When asserted on a bus, the adjacent lines carry data of complementary logic levels. These help to identify the faulty components (ECL-TTL converters, latches, buffers and ASICs).

The non-constant pattern is a time series with these two patterns alternating. This creates a square wave on all the data lines of the bus. Two adjacent lines are always out of phase by 180°. This pattern tests out the capture of the data and control signals on the card. This is so because the two patterns correspond to the two polarisations and hence they get accumulated separately (in the default RR-LL mode). Thus every alternate number in the output should be the same. If due to some reason, be it faulty capture of data or the control signals, these two cannot be isolated from each other the output will be different.

4.3 Subrack Level Tests

Subrack level tests are performed for all subsystems to ensure that all cards in the subrack can work in tandem. Identical patterns are input to all the cards and the outputs are observed. One of the criteria for the choice of test patterns is to make sure that the transmission of the control signals on the backpanel is proper.

4.3.1 ADC, Delay and FFT cards

For the ADC subsystem, the subrack level testing in not much different from the card level testing. The subrack is populated with all the 12 cards. An analog signal is input to these through a splitter (generally only a 1:2 splitter is has been used in the lab, although on the field a 1:12 splitter also has been used) and the output is observed.

The Delay subsystem consists of 5 DPC-Delay sets per subrack. The subrack level testing consists of driving all the 5 DPCs with the test data from the control card. The test pattern still remains the same but the degradation of signal quality of the control card output due to loading by the cards is the main issue. It also tests out the clock distribution system of the subrack since the clocks on all the cards are assumed to be in phase with each other. This is of concern because in the event of a skew between the delay card outputs, their capture on the FFT cards will dictate the tolerance within which these two subsystems can operate.

The FFT subsystem has 10 FFT cards per subrack. These are split into two 5 card sets. Each set has its own control card. Thus, the first test for the FFT subrack is to observe how well the two control cards synchronise with each other and to verify that their outputs agree exactly with each other since both of them get the same clock. The next part is to tune the FFT card clocks with respect to the FFT control clock. This decides the set-up and hold times for the FFT card latches which capture the control card signals (INITs and trigs). For the twiddle and control signals the fan-out is lesser as compared to the delay (and also the MAC) since each control card drives only 5 cards. The outputs from the two sets should thus be identical both temporally and in value. The skew, if any, between the FFT card outputs is of serious concern since it severely limits the range for tuning the clocks on the MAC subsystem. This is so since each FFT card drives 8/9 MAC inputs on different MAC cards distributed in different racks.

4.3.2 MAC cards

The subrack level testing of MAC cards is somewhat different from the rest, not in principle but in practice. The aim here still remains the same - put all 11 MAC cards in the subrack and observe

Figure 4.2: The FFT Simulator card

the performance in terms of repeatability and stability. The problem here is to provide the required number of inputs reliably. One of the simple tests is to feed identical patterns into all 11 MAC cards, and so all the 176 MAC chips should behave identically.

The FFT subrack is used as the pattern generator in this case. Instead of using the FFT cards, a smaller and much less complex card was designed using ROMs and counters to generate any required pattern. This card is an add-on to the FFT main board (Fig. 4.2). Ten of these add-ons were used as pattern generators. An additional constraint was the non-availability of the FFT output distribution system at the time when the first MAC subrack was tested. This meant that only 9 out of the 11 MAC cards could be excited at any given time, although all 11 MAC cards were present in the subrack. Two passes are thus required to test out the entire MAC subrack. The output is observed on the logic analyser which operates in an acquire and compare mode. In this mode, the analyser compares the recent acquisition pattern with a reference pattern. In the event of any differences between the two, the logic analyser stops the acquisition. In case of no differences, the acquisition continues. Necessary tuning of the system (terminations, clock phase etc.) is done to obtain a condition that the acquisition never stops.

4.4 System Level Tests

The system level tests are done, part by part, as the system is integrated in three stages, viz.,

- FX tests
- Delay and FX tests
- ADC, Delay and FX tests

4.4.1 FX tests

As the first step towards the complete system, the FFT and MAC subsystems are integrated as one FX system. This integration is mainly tuning the synchronisation signals so that the FFT and MAC form a pipeline. The first channel of the FFT output then becomes the first channel of the MAC as well. This synchronisation is achieved by using an external signal to synchronise the FFT and MAC control cards. The clocks of the two subsystems also need to be tuned to account for the delay in transmission of FFT outputs to the MAC inputs.

The FFT subsystem is input with digitised sine wave of various frequencies and the MAC output is observed. The MAC output should exhibit a spike in the appropriate frequency channel corresponding to the FFT input. The sine waves of different frequencies are generated from the same fringe rotator data by changing the NCO rate. Thus by programming a single cycle of a sine wave in the fringe rotator ROM and appropriately programming the NCO rate, the FFT output can be made to peak in any of the 256 frequency channels. The same behavior should be evident in the MAC output. The FX setup is then left for "endurance" testing - identical input is fed to the FFT, cycle after cycle, and the MAC output is observed on a logic analyser or on an acquisition system. The MAC output too, should remain identical, cycle after cycle, over extended periods of time. If all FFT cards produce identical outputs, all MAC chips will behave as auto-correlators which can be verified by looking at the phase of the output. It should be zero. This tests out both, the transmission of data from FFT to MAC and the transmission of data within the MAC subsystem (which means data coming through the distribution system onto the MAC backplane and then being distributed on the MAC backplane).

Once the FX system is integrated, the basic computing block of the correlator system is ready. The Delay and ADC subsystems are then integrated with the FX system.

4.4.2 Delay and FX tests

Here again, like the FX integration, the Delay subsystem is tuned with respect to the FFT subsystem. The Delay subsystem has a provision to inject test data from the Delay control card into the DPC subsystem. This then provides a self-contained and completely digital setup. A pattern generated through the Delay subsystem can be delayed by different amounts to produce the corresponding phase difference in the MAC output. The patterns usually used are the DC and sine wave. The DC input produces a spike at the first channel of the MAC output while the sine wave produces the spike at the channel corresponding to its frequency. The sine wave is shifted by different values of delays to produce different phase differences compared to the sine wave in the "reference" delay card (with a delay of zero).

This tests out the entire digital section of the correlator. It validates the transfer of data between the Delay and FX subsystems. This arrangement provides a way to accomplish correlator diagnostics every once in a while. This can be programmed into the correlator control system which can put the system into the diagnostic mode whenever possible - even during routine observations.

4.4.3 ADC, Delay and FX tests

The inclusion of ADC subsystem is the final part of the system integration. The ADC clocks are tuned to ensure that the ADC data gets captured properly into the Delay subsystem. Once done, analog signals are input into the ADC and the response is observed at the MAC output. Typical signals input into the ADC are analog sine waves of various frequencies and bandlimited random noise. A de-correlation test is performed at this point before connecting the antenna signals to the ADC input. Noise of a certain bandwidth, say $\Delta\nu$(typically 8MHz), is input to the ADC and a sufficiently large delay of the order of $1/\Delta\nu$ is inserted into one of the the ADC paths. The output of the corresponding cross-correlator chip in MAC is observed. The output should be almost de-correlated. Inserting delays less than $1/\Delta\nu$ should give some correlated power in the cross chip. With a delay of zero, the cross chip should be the same as the auto-correlator chip.

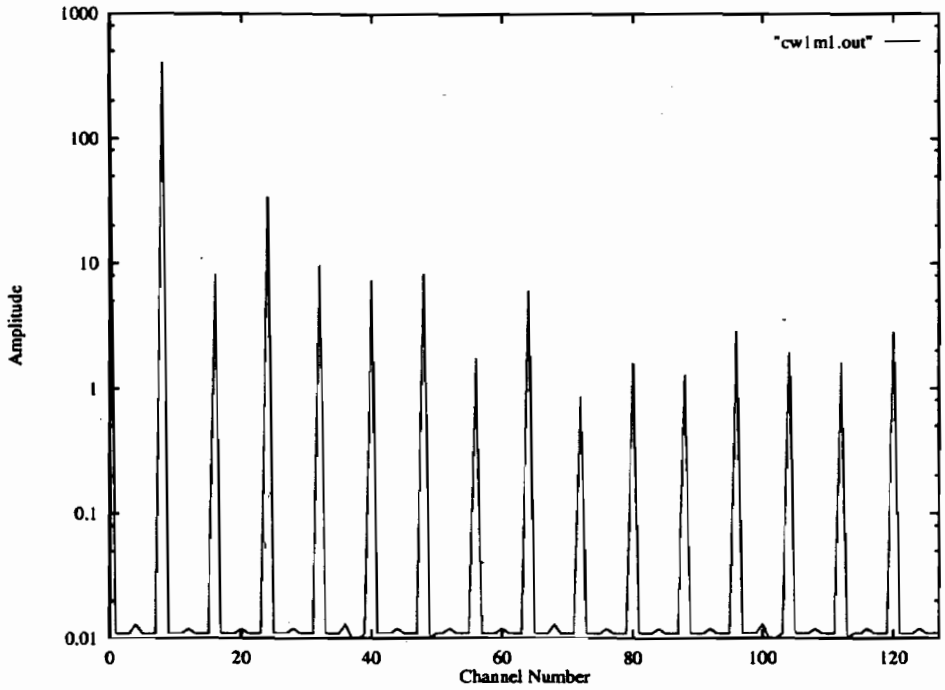

Figure 4.3: A 1 MHz Sine Wave

A few of the results of the system level tests are shown in the following plots. These plots were obtained with an arbitrary waveform generator being used to generate the inputs to the 4 ADC cards (through a splitter). The output was acquired after 64ms STA in the MAC. For the sine waves of 1MHz and 10MHz frequencies as well as for the analog noise of 5MHz and 10MHz bandwidths, all delay cards had identical delays. Introducing a one clock cycle delay between Inputs 1 & 2, 2 & 3 and 3 & 4 (i.e input 3 is delayed 2 clocks with respect to input 1 and input 4 is delayed by 3 clocks), the amplitudes do not show any significant difference but different phase ramps are observed. Increasing the delay to 64 clock cycles, de-correlation effects are observed. The corresponding phases are also shown.

Completion of validation and testing of the system is followed by its commissioning at the GMRT site. As of today, the eight station single sideband correlator installed at the site has been performing satisfactorily; though a lot still needs to be accomplished. The integration of the final 30 station single sideband correlator is in progress in the laboratory at NCRA, Pune. It is expected to be over by February 1998. Looking back over the past years, one finds certain unusual features in the system. The next and concluding chapter of this thesis discusses some of such features.

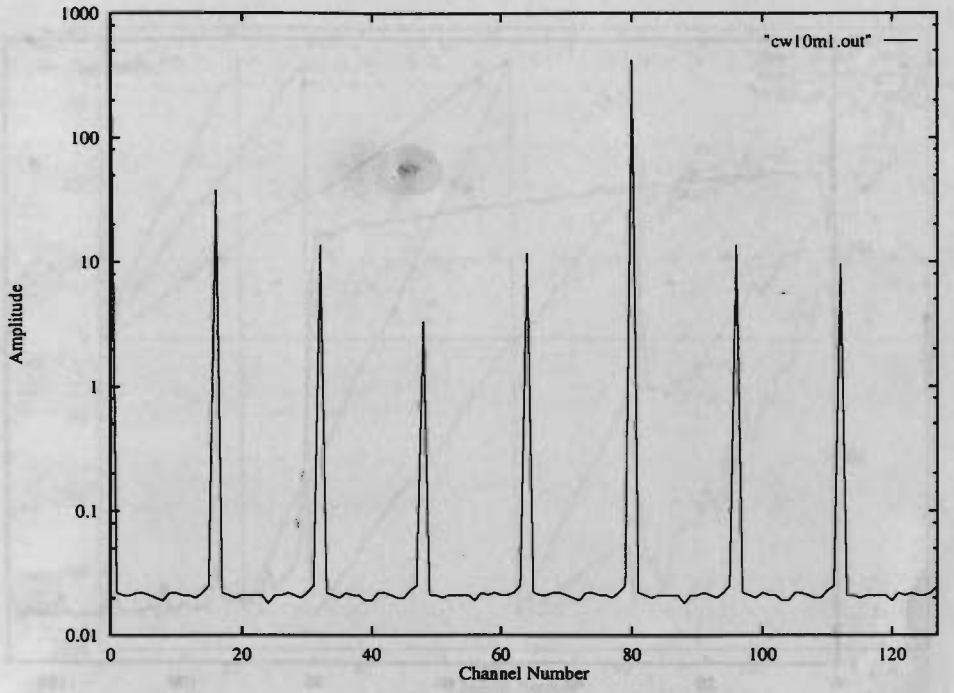

Figure 4.4: A 10 MHz Sine Wave

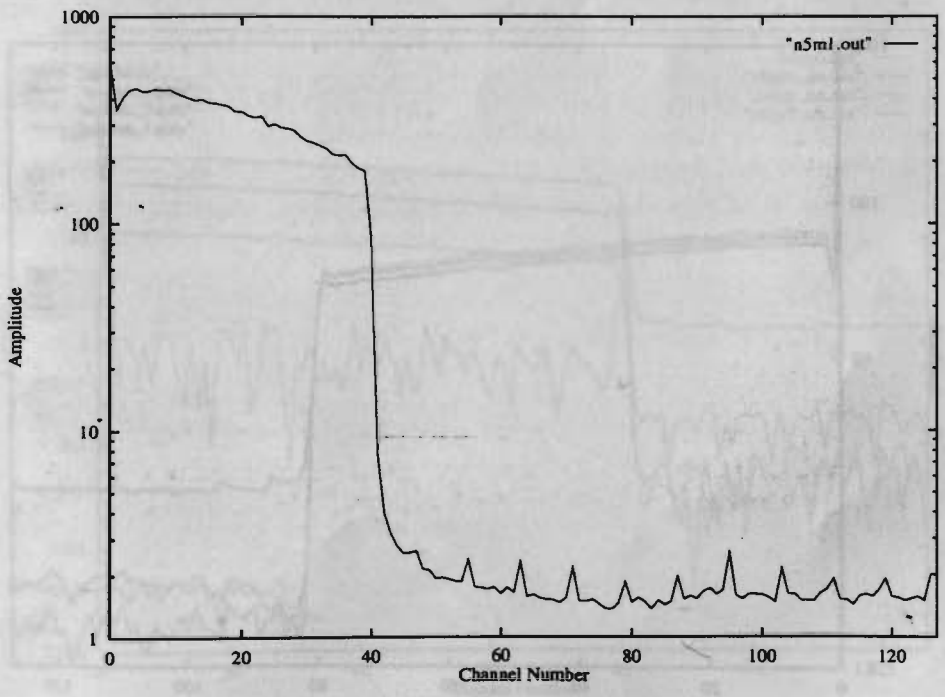

Figure 4.5: A 5MHz Bandwidth Noise

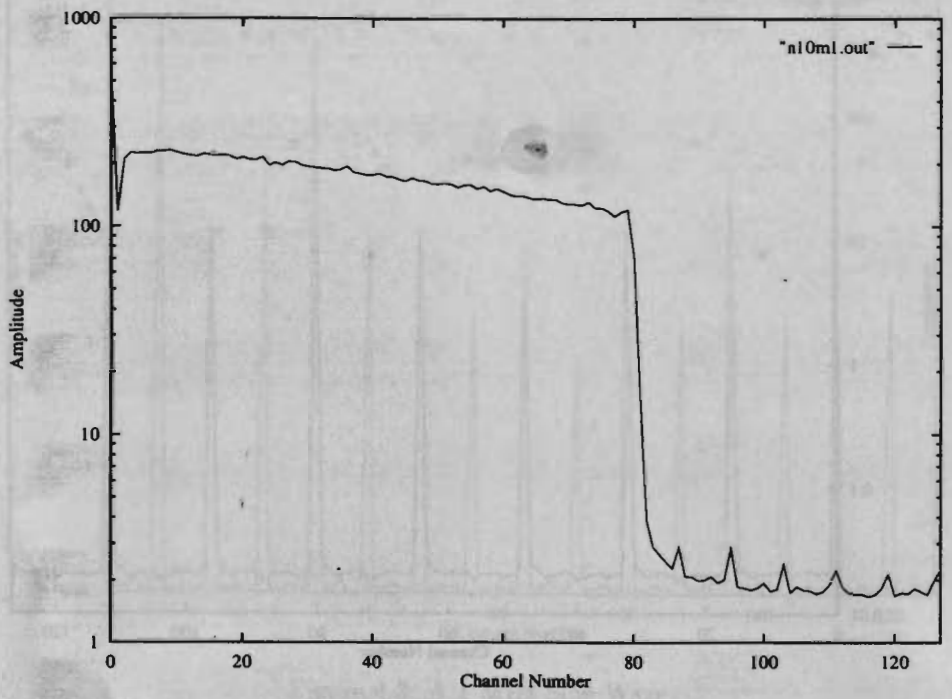

Figure 4.6: A 10MHz Bandwidth Noise

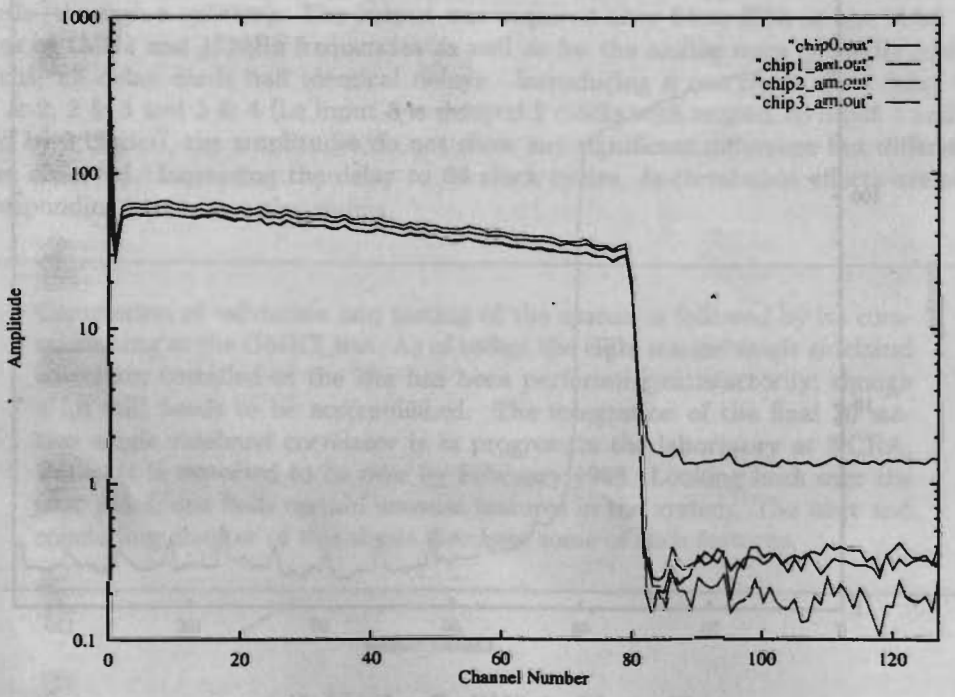

Figure 4.7: Amplitudes of the 4 Inputs with one clock delay

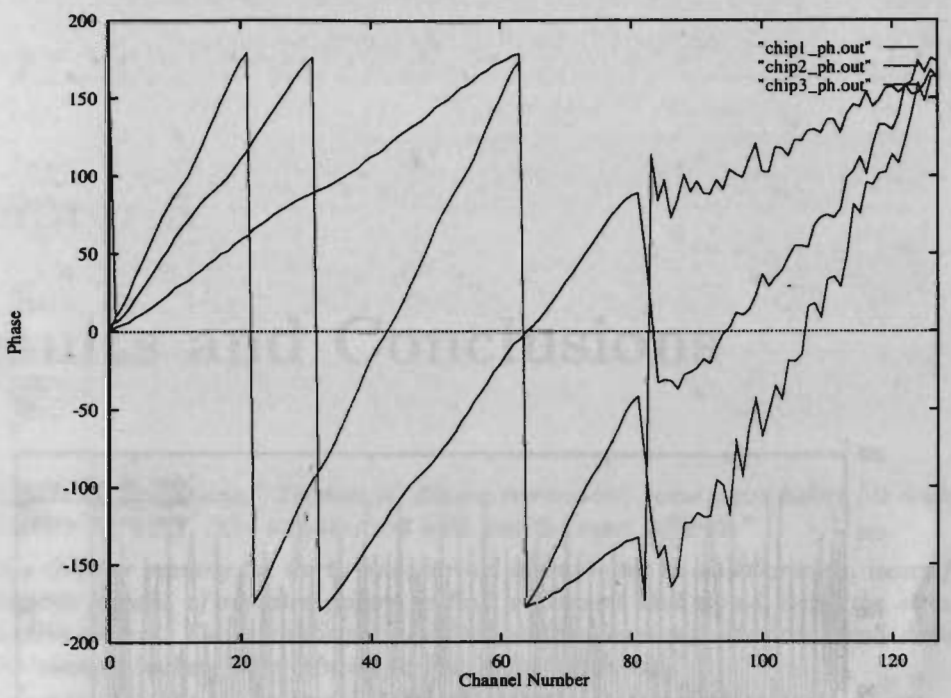

Figure 4.8: Phases of the 4 Inputs with one clock delay

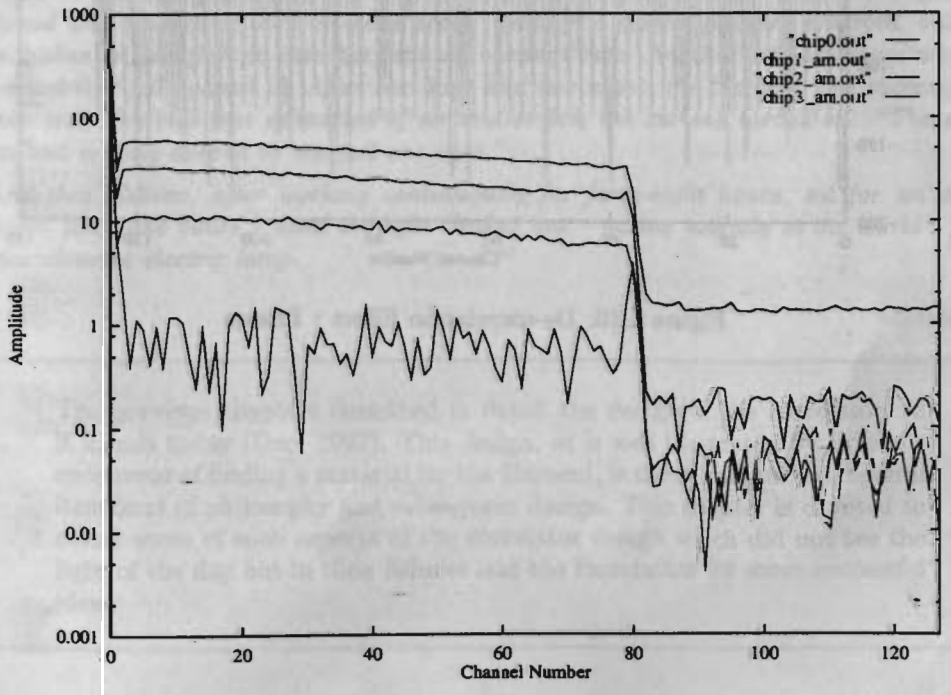

Figure 4.9: De-correlation Effect : Amplitudes

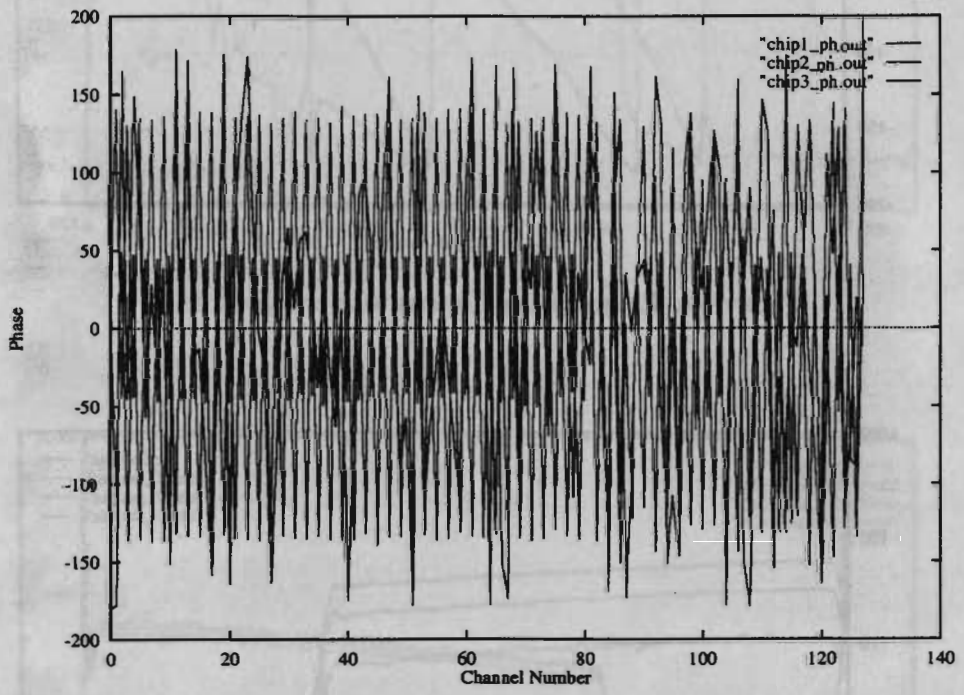

Figure 4.10: De-correlation Effect : Phases

Chapter 5

Results and Conclusions

"Of all my inventions," Thomas A. Edison reminisced, some years before his death on October 18, 1931, "the incandescent light was the most difficult."

One October evening the thirty-two-year-old inventor sat in his laboratory, weary from thirteen months of repeated failure to find a filament that would stand the stress of electric current. The scientific press, at first politely skeptical, were now openly derisive. Discouraged backers were refusing to put up further funds.

Idly Edison picked up a bit of lampblack mixed with tar, rolled it into a thread. "Thread," he mused, "...thread...thread...carbonized cotton thread." He had tried every known metal.

It required five hours to carbonize a length of a thread. The first one broke before it could be removed from the mould; likewise a second and a third. An entire spool of thread was consumed; then a second spool. Finally, a perfect filament emerged, only to be broken in an effort to insert it into the vacuum tube. Another was destroyed when a screwdriver fell against it. After two days and two nights the filament was successfully inserted. The bulb was exhausted of air and sealed, the current turned on. "The sight we had so long desired to see met our eyes."

And then Edison, after working continuously for forty-eight hours, sat for an additional forty-five hours – until the light blinked out – gazing intently at the world's first incandescent electric lamp.

–Gospel Banner

The previous chapters described in detail the design of the correlator, as it stands today (Dec. 1997). This design, as is well illustrated by Edison's endeavour of finding a material for the filament, is the culmination of several iterations of philosophy and subsequent design. This chapter is devoted to reveal some of such aspects of the correlator design which did not see the light of the day but in their failures laid the foundation for more successful ideas.

> Perfection should always be the goal. But it always remains a myth, a mirage which one endures and aspires to achieve, but nevertheless falls short of it; for nobody and nothing is perfect. The correlator design team too, aspired to make the correlator blemishless in all respects, an epitome of perfection. Constraints of time and resources often compel Man to adopt pragmatism - almost all correlator cards have at least one "strap" on them - a signature of relinquished dreams. This chapter also tries to put across some odd features of the correlator system which are the remnants of mistakes committed in the quest of perfection!!!

5.1 Experiences : Sweet and Sour

There can be many ways to bring forth a discussion which embraces the six to seven year long development cycle of the GMRT correlator. One could describe it in a chronological fashion dating back to the beginning of the design or one could view it on a sub-system basis and describe the reasons for apparent oddities, if any, existing today in the sub-system. Here the second option is chosen.

5.1.1 The ADCs

The ADC card was the first to be designed. This card has had two prototypes before the final version. This card, being the interface between the analog and digital realms, has always evoked keen and concerned responses from both, the engineers and the astronomers. The design of the ADC card therefore has seen many a tumultuous discussion. A few of the features of the present card are outlined below.

Use of Sub-miniature fuses The power supply sections for analog and digital ± 5V, on the ADC card, use PCB mountable sub-miniature fuses. This was done mainly to conserve "real-estate" on the ADC card. Unfortunately, this plan has backfired due to the poor quality of these fuses, available in the local market. The fuse rating is supposed to be 500mA, but some of these activate at currents as low as 200-250mA. Secondly, the leads of these fuses do not make a reliable contact with its body, as a consequence of which a "live" fuse before mounting on the PCB can show an "open" circuit after mounting. Bending of these leads causes discontinuity in the fuse !!!. The solution : Either import the fuses or remove them and replace them with shorts. We chose the latter.

Use of Isolation transformers Isolation transformers are being used, as of now, between the baseband outputs and the ADC inputs to isolate the grounds of the two systems. This need was felt when it was discovered that some correlated noise was present in the signals. This was traced down to a variable potential difference between the grounds of the baseband and the correlator systems. The use of these transformers has eliminated the correlated power but introduced a new burden of mounting and packaging these transformers in a manner to avoid making them into chronic nightmares for the maintenance personnel.

The size of the ADC sub-rack is non-standard Although the ADC subrack is shown to be 3U in size in chapter 3, it actually is about $3\frac{1}{2}$U. This has come about because of using a card dimension of 110mm with an edge-connector card guide instead of the normal one. Finally, however, the normal card guide was used due to which the subrack size had to be increased making it non-standard. The change in subrack size created a ripple effect; disturbing the position of the fixtures mounting subracks, which in turn shifted the fan-trays. The holes on the rack rails also had to be altered to take into account the increased size. In the end, the

sub-rack and fan-tray combination was made into a 5U unit (instead of the normal 4U) and then mounted. The extra 1U thus required (5 units of 5U each in a 24U rack) was obtained by shifting the first subrack a little above the normal mounting position.

Clock distribution The initial plan for providing clocks to the ADCs was to send them from the DPC card. It did not take long after the first few iterations of system integration to realise that the scheme would not work. The individual delays of the 120 cables along with their associated drivers and receivers would make it difficult to maintain the sampling edges within tight tolerances. The clock distribution would thus be much better by using the clock card instead. Thus the DPCs still have the remnants of the earlier clock distribution scheme while the ADCs have their clocks coming from the clock distribution cards.

The 7X2 IDC connector assembly This connector assembly was chosen to provide a compact connector which could plug into the Euro connector directly. The connector initially identified was for crimping a flat cable. The connector which finally got procured was for individual wire crimping. This particular connector was declared obsolete and subsequently removed from the production list by its manufacturers. This made the procurement of the proper crimping tool a harrowing experience. A single tool finally did make its appearance but it still meant that somebody had to crimp each of the 120 X 14 strands of wire onto the 120 connectors!!!

5.1.2 The Delay/DPCs

The delay subsystem has had two prototypes. The first prototype essentially had both DPC and Delay on the same card. The card thus became extremely dense. Later, in the second prototype the DPC and Delay section were separated into two different cards. Four signals were brought into the DPC while the Delay handled only two. The three card set was thus required to process all the four inputs to the DPC.

One of the most troublesome mistakes made in the delay subsystem was the use of an asynchronous ROM in the DPC to realise the encoder look-up table. The problem was further aggravated by the fact that the ROM was very critically operating when clocked at 32 MHz, and a pin-compatible faster device was not identifiable. The output thus contained spurious pulses at random, confusing the computing chain following it. A piggy-back card using synchronous Cypress CY7C245A was later on added to rectify the error.

Another mistake made in the delay subsystem was not to latch the carry bit of the adder in the power estimation section.

5.1.3 The FFTs

The FFT subsystem has had two prototypes out of which the first never left the PCB factory. This was the first time, forced by time constraints, the prototypes were sent abroad for manufacture. This first prototype of the FFT card was a single 10-Layer card with the ASICs occupying some 80% of the board area. The rest 20% was the add-on cards, mounted orthogonal to the plane of the board, like SIMMs. These contained all the peripheral circuitry which as of today occupies the main board. The board used 8mil tracks with 12mil drills and 8mil track-to-track and track-to-pad spacing. The PCB factory launched these boards, only to discover shorts on the inner layers later in the process. This lot of boards were rejected and a new one launched afresh, only to find the same result. The factory, at this point gave up. This factory was known to have the expertise to deliver high-end low-volume boards with fast delivery times. Their failure, that too twice in quick succession, sent the signal loud and clear to somehow reduce the density of the board and increase tolerances while routing so as to make the board less complex to manufacture. Our endeavour was to try as much as possible, and get the production boards manufactured within India. Something

desperate thus, had to be done to make the boards simpler and routine enough to be manufactured within India.

The two tier structure was the result. The second prototype thus had a main board and a pipeline board. The main board was designed as an 8 layer board and had only the peripheral circuitry while the pipeline board was a 6 layer board having only the ASICs. Three SBus connectors were used to interconnect these two boards. These prototype boards were manufactured by the same overseas factory which had earlier tried to make the 10 layer board. The boards were assembled and tested successfully. These boards were the first ones in the correlator system to use surface mount components. They were then the natural test beds for trying out assembly of SMT components. It was found that the pad sizes used for the SMT components on these boards were critical, and were particularly inconvenient for manual soldering or rework in case of field problems. Slightly longer pads would have made in-house assembly of these cards very much simpler. The cards due to this reason gave lot of contact problems between the pads and the component leads.

The next version, which was the final version of the FFT cards, used SMT connectors for interconnecting the main and the pipeline boards, thereby reducing two layers from the main board. Both boards were now designed as 6 layered PCBs. The SMT pads, too, were lengthened to facilitate easier and more reliable assembly of SMT components. These boards, however, owing to their large quantity, were assembled using automatic assembly techniques. The contact problems in these boards, were rare if not absent.

5.1.4 The MACs

The MAC cards have had only one prototype. The only deficiency in it was the absence of the card ejectors, which happened due to an oversight. The MAC cards use two five-row Euro connectors. A total of 320 pins make it very difficult to pull out the card without the ejectors. The final version was the same as its prototype except for the addition of mounting holes for the card ejectors.

The MAC backplane did not have any prototype. The earlier backplane designed could cater only to three cards. The 11 card backplane thus is the only card in the correlator not to have any prototype. Delivery schedule constraints forced overseas manufacture of this 16 layer PCB. This card thus also has the distinction of being the only card in the final correlator to be manufactured overseas.

5.1.5 The Clock Distribution System

Right from the inception of the correlator, the clock generation and distribution has been an issue which has been procrastinated every time it came up, mainly because the non-availability of the twin clocks (32 and 32.25MHz) did not hamper the development in any manner. Even from the astronomy point-of-view, only a few kinds of observations get seriously affected due to this non-availability. In effect, the non-availability translates to the loss of 4 data points every FFT cycle. This makes it slightly complicated to interpret the phased array output of the pulsar machine for time resolutions better than one FFT cycle (16μs).

Various schemes for the generation of these twin frequencies have been tried out. These include generating these twin frequencies using filters, using discrete PLLs, using clock generator ICs and using a DDS (Direct Digital Synthesiser). Of these the first two were tried out around five years ago where as the latter two are the more recent ones (past two years). The main options available as of now are : either use a clock generator chip (Cypress ICD2051) or use a DDS (Analog Devices AD9955).

The ICD2051 synthesises the frequency depending on the value of a control word loaded into it and the frequency of the reference signal input to it. It uses PLLs to synthesise the output. It is a dual clock generator and has two independent PLLs controlled by two independent control words. The reference input, however, is common. This chip comes as a 20-pin DIP SO-IC package. The earlier idea was to use this chip to generate the required clock (either 32 or 32.25MHz, as the case

may be) on the backplane with the reference input being distributed to all backplanes from a single point. This idea failed because any two clocks thus generated could, on power-on, come up with opposite phases. This, however, may not be a hindrance if both the clocks are generated at only a single point (e.g. using a single ICD2051 chip), and the distribution system provides for a field-adjustable delay at each subrack to properly compensate for pipeline delays. In order to provide for such delays, the original clocks are first converted to sine-waves using a low pass filters followed by adequate buffers, and distributed to each subrack in a Christmas-tree arrangement. At each subrack, the clock distribution card allows for the necessary delay adjustment before converting the sine-wave to square wave for distribution to the cards as ECL signals.

For certain time-critical observations, it is necessary to lock the correlator clock to a stable reference frequency. As of now, opinion seems to be divided on whether it is always desirable to "tie" the time-reference and the frequency reference of the local oscillator system to the same source. In the GMRT, the local oscillator system generates a phase-stable 100 MHz signal which can optionally be locked to an external 5 MHz signal. Depending on whether it is desired to derive the correlator clock by down-converting this 100 MHz reference from the local oscillator system or by upconverting from a 5/10 MHz signal from a Rubidium standard available at the observatory, the DDS-based or ICD2051-based clock generator can be used. Both possibilities of generating the two clocks have been tested and, as of now (December 1997), the DDS is being used.

5.1.6 The LTA and DAS

What was planned...

The output of data from the GMRT correlator has four distinct stages. The first stage will occur during the 4-clock-cycle dead time in each FFT cycle of 516 clock cycles, and the next stage will occur during the active period of duration 512 clock cycles in each FFT cycle. The other stages are asynchronous with respect to the correlator system clock.

The first stage is essentially an internal operation where every ASIC in the multiplier cards will output one STA result, a complex word of 36 bits into separate buffers (2 x 18 bits per ASIC) provided for, within each multiplier chip.

In the second stage, one STA result (36-bit complex) from every ASIC is loaded into the appropriate MAC control card, converted to a 32-bit complex word - (13,13,6) format - and stored in appropriate locations in 32-bit memory banks provided in each MAC control card. There will be 4 memory banks in each MAC control card, but only two will be active during any given FFT cycle. The other two banks in each MAC control card can be freely accessed by any external circuitry.

In the third stage, the STA results are read from the memory banks in the MAC control cards into the long-term-accumulator (LTA). The LTA will provide various options of adding the data. In particular, they can be added across the spectral channels, thus reducing the number of spectral channels over the given band, or/and add contiguous samples of each spectral channel to provide any desired integration. The net result of these options is to effectively reduce the data rate to within about 1 MB/s corresponding to the maximum rate specified for sustained operation of the data acquisition from the entire correlator system. This stage is carried out within the MAC control card which includes a fixed point DSP chip ADSP 2105 for data read out and control and a floating point DSP chip ADSP 21020 for the LTA operation.

In the fourth stage, the results from LTA within each rack will be accessed by the processor-based cards - one for each FX rack on the average. The software in these cards provides for selection, regrouping and further integration depending on the user-specifications for a given observing session. The data are then expected to be pooled together into a buffer card - perhaps identical to the other processor cards, but running a different software. The final data would be sent for direct recording on a SCSI device (Exabyte tape), with simultaneous availability of a subset on the network for monitoring purposes. The final acquisition system was to be networked

with other Unix systems and data logging and monitoring is done through a set of co-operative processes running on a TCP network.

What happened...

While the strategy for first two stages hardly underwent any change during the evolution of the correlator design, the other two stages went through several iterations. Originally, the LTA and data acquisition were planned to be implemented on a combination of a set of cards on a VME bus, with a control processor running a real time operating system VxWorks. However, for the first 4-antenna system (Mark1), a patch-up was done using a chain of transputer cards - a set of about 40 processes were running on a total of 7 transputers. This provided a combination of simulation of many functions intended subsequently for hardware implementation as well as the first generation software which was finally intended to run on a network of Unix computers. This experience, together with the fact the a large number of transputers and memories were available from an obsolete project, made the choice of transputers very attractive for rack control and data distribution. In order to help smooth interfacing with the transputer-based rack control cards, the final design of the delay, FFT and MAC controls based on the DSP chip ADSP 2105 were all provided with a transputer link using IMSC012 as the parallel to serial converter. To help laboratory development, an ISA card was also designed with four INMOS links in the card and device drivers were developed for this card on both the DOS and Linux operating systems. The rack control cards were first prototyped, and final cards designed and fabricated. The MAC control card included a separate section consisting of dual-port memories, look-up-tables and a floating point DSP chip ADSP 21020 for carrying out the LTA. There was a circuitry for assisting accurate time-stamping with the help of counters whose outputs were latched at each STA pulse, and got re-initiated with an externally provided minute pulse from the GPS receiver available in the observatory.

However, at this stage, the engineers with expertise on transputers left the group, and the pressure for fast delivery was mounting on the others. By this time, there was also the sudden increase in the computing power due to the availability of low cost Pentium PCs, and the PCI provided bandwidths for data acquisition eliminating the limitations of ISA/EISA buses. The new team were reluctant to use the unfamiliar transputer-based cards and decided to merge the third and fourth stages into Pentium PCs, based on commercially available PCI cards for fast data acquisition. Part of the reason for taking the LTA outside the correlator hardware was also an unfortunate slippage in the design of MAC control cards due to which the ADSP 21020 did not get the power/ground connections routed during fabrication!

In the system currently being commissioned, the section of MAC control card containing the look-up tables for IEEE conversion, ADSP 21020 circuitry and a time-stamping logic have not been wired. Only the MAC control part of this card was retained, and the data are directly fed to the PCI card through a cable, thus shifting the job of data selection, conversion and LTA to a chain of Pentium PCs running Linux operating system.

5.1.7 System Integration

Till date four correlators, of gradually increasing complexity, have been integrated. These, for ease of description, can be referred to as Mark 0,1,2 and 3 correlators.

The Mark0 correlator was a concept demonstrator, having only a single card each for ADC, Delay, and FFT subsystems. A single MAC chip correlated the spectra from the two pipelines. This system was made to demonstrate the GMRT interferometer. The delay setting was static and was done using DIP switches. FFT size was fixed at 512 points. MAC integration time was 64ms and it was used in the Non-polar mode.

The Mark1 correlator was a 4 station single side-band correlator. Variable delay setting was used for the first time. The MAC subsystem had a choice of two integration times - 4ms and 64ms.

FFT size was fixed at 512 points. MAC was used in the Non-polar mode.

The Mark2 correlator was a 8 station single side-band correlator with the MACs used in the Indian polar mode. This was the first prototype of a completely in-house built correlator with a first generation control system. This correlator formed the test bed for a number of strategies and provided invaluable experience in the areas of system integration, test strategies, control strategies and reliability. The 4 station subset of this correlator performed satisfactorily but the 8 station configuration was beset by reliability problems. The source of these problems were traced to the non-availability of clock tuning mechanism and 'SMT pad-component lead' contact problems in the FFT cards. The experience gained from this correlator was used to refine the testing and validation procedures of cards as well as the entire system. This system also provided a profound insight into the control system configuration. All aspects of the control system viz. hardware, software, protocols were scrutinised and subsequent modifications were made to the second generation control system which is currently in use.

The Mark3 correlator was similar to the Mark2 in being a 8 station correlator. The difference lay in the constituents. The feedback obtained from the Mark2 system was used in this to enhance its reliability. All the subsystems, except the MAC, used the final version cards, which in itself was instrumental in improving the reliability of the overall system. The control system used was the second generation, which used DSP cards as the embedded control cards (the first generation system used micro-controllers and transputers). These control cards were linked to a PC which was networked to the online computer and file server. The system used a clock tuning mechanism to provide phase shifts between clocks of the various subsystems and also within the subsystems. A diagnostic capability, in the form of certain self tests, was enabled in the system. This has proved to be of immense value (this 8 station correlator has been performing satisfactorily for the past entire year) in the field, to ascertain the health of the system, if need be.

5.2 The other Roads to Rome !!

All roads lead to Rome, true, but was ours the right one? At the fag end of the design of the correlator, one can indulge in the luxury of introspection in retrospect. The following discussion endeavours to explore the alternate ways which could have been chosen to realise the correlator. Some out of these did have the good fortune of being contemplated upon, but were not adopted. A few "why's", in such cases, are also given.

5.2.1 Modularity, but with a difference !!

The system, as designed, is modular on a functional basis. The subsystems of a particular kind are placed in one rack or at least a subrack. Modularity could also have been achieved by segregating the system on a station basis. All subsystems processing signals from the same station could have been packaged into one subrack, except for the MAC. The ADC, Delay and FFT for different stations could then have fed the multiplier array. The system could then, possibly have been more scalable. And, maybe interconnections would have been much easier.

One argument against this is that, for instance, if the ADCs do require some shielding, then it is much easier to put all the eggs in one basket and shield *that* basket rather than shielding every egg separately!! Thus, a conscious decision was taken not to mix the ADC cards with any other cards right from the beginning. As for the Delay and FFT subsystems, an antenna-oriented backplane would have required the backplanes to evolve with the different versions of DPC, Delay and FFT cards. Since these cards imposed varying requirements of evolving the design by prototyping individual modules, it was not considered practical to ensure these cards to evolve simultaneously and wait till all these are finalised before the backplane design could be finalised, Further, the multiplier was too complicated and inherently array-oriented to be handled in an antenna-oriented approach. Such considerations deterred us from carefully analysing the antenna-oriented backplane as a possible alternative to functionality-oriented backplane which is currently adapted.

In retrospect, perhaps a compromise could have emerged if we had the experience to begin with. In view of the patches made on the delay system to get them to function in the present system, it is very likely that this subsystem will soon be redesigned keeping in mind the technological advances over the last few years. In such a case, it may be worth considering an integrated delay-FFT backplane as part of the next phase. This may be of some advantage from the point of view of maintenance, where compartmentalisation can often be a hindrance.

5.2.2 Would XF have helped !!

There!!! The genie raises its head again!!! Perhaps, yes, but at definitely more cost. One advantage, from a management point of view, is that in a XF correlator, one could, in principle have made a 30 antenna correlator but for a single channel and then later scaled it to the required number of channels. In the present, FX situation, all channels are *always* available (although they are hardly ever used!; at least in the beginning) but to get all the stations, one has to wait and wait and wait!!! Wait, till the entire gamut of cables, cards and racks are assembled. This lack of station-wise scalability provides ample room and conducive conditions for anxiety, then cynicism, followed by criticism, finally aspersions and ultimately frustration to breed.

5.2.3 So near, yet so far ...

All through the past six years, the correlator team has been hoping for the day to dawn when they can sit back in leisure and enjoy the nostalgia of building the fastest signal processing system in India. As our peers in the other groups of GMRT look on, the thirty station correlator slowly takes shape. The greatest hurdle as of now seems to be to handle the cable I/O of the FX rack. We hope to surmount this in the coming days and realise a long cherished dream.

On the whole, the author feels that building the correlator has been an enriching experience. He would like to thank all members of the team, past and present, to have made it such a pleasurable endeavour. He also acknowledges the help extended by members of the peer groups in making of the correlator.

That completes the story of the GMRT correlator. The strides made by technology in the years since the correlator was conceived, may make this particular realisation look somewhat naive, but given the context, resources and expertise available at the disposal of the design team six years ago, it still does look as a sensible proposition.

Bibliography

[1] http://www.sao.nrc.ca/inms/photonics/english/optpro.html

[2] Bernard F. Burke and Francis Graham-Smith – An introduction to Radio Astronomy. Cambridge University Press.

[3] Kristen Rohlfs – Tools of Radio Astronomy. Springer-Verlag.

[4] John D. Kraus – Radio Astronomy. Cygnus-Quasar Books.

[5] A. Richard Thompson, James M. Moran and George W. Swenson – Interferometry and Synthesis in Radio Astronomy. John Wiley and Sons.

[6] W.N. Christiansen and J.A. Hogbom – Radio Telescopes. Cambridge University Press.

[7] Nigel Henbest and Michael Marten – The New Astronomy. Cambridge University Press.

[8] Larry D'Addario – "Cross Correlators" Chapter 4, Synthesis Imaging in Radio Astronomy. Astronomical Society of the Pacific Conference Series Volume 6.

[9] A. Richard Thompson – "The Interferometer in Practice" Chapter 2, Synthesis Imaging in Radio Astronomy. Astronomical Society of the Pacific Conference Series Volume 6.

[10] Special Issue: The Australia Telescope – Journal of Electrical and Electronics Engineering, Australia.

[11] Instrumentation and Techniques for Radio Astronomy. IEEE Press.

[12] P.J. Napier – "VLBA Design" Chapter 4, Very Long Baseline Interferometry and the VLBA. Astronomical Society of the Pacific Conference Series Volume 82.

[13] G.Swarup, S. Ananthakrishnan, V.K. Kapahi, A.P. Rao, C.R. Subrahmanya and V.K. Kulkarni – The Giant Metrewave Radio Telescope. Current Science, Vol. 60, No. 2, 25 January 1991.

[14] J.D. Romney – Introduction to the Spectral-Domain ("FX") Correlator. VLBA Correlator Memo No. 60.

[15] Ray Escoffier – VLBA Correlator ASIC Chip Specification. VLBA Correlator Memo No. 87.

[16] ASIC Block Diagram, Schematics (84 Sheets) and other information.

Appendix A

Schematics

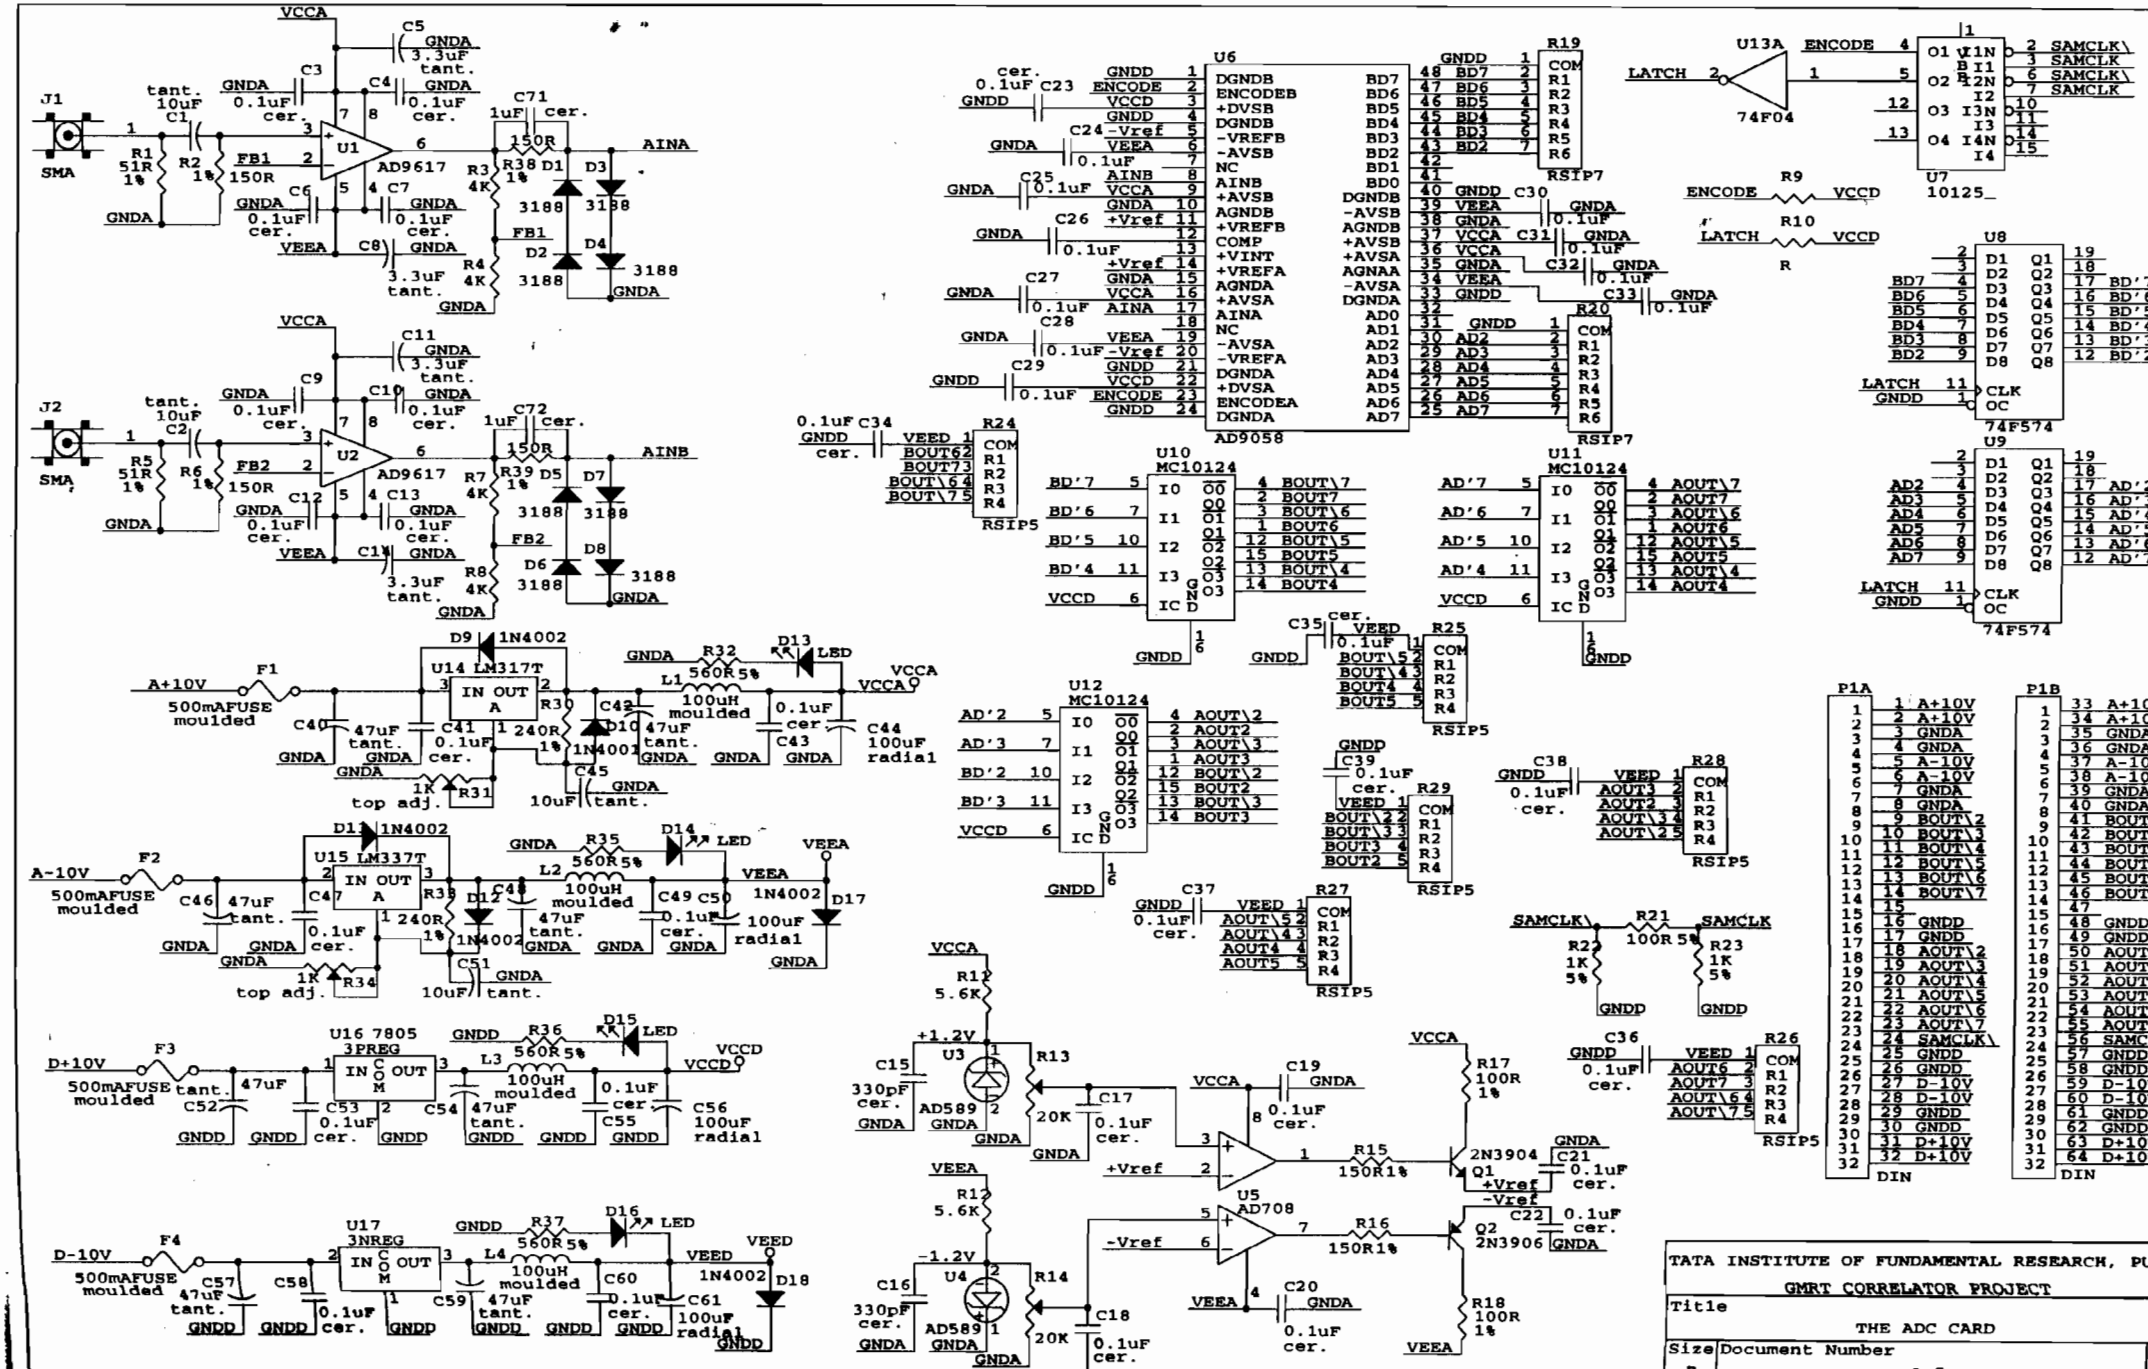
TATA INSTITUTE OF FUNDAMENTAL RESEARCH, PU
GMRT CORRELATOR PROJECT
Title
THE ADC CARD
Size
Document Number

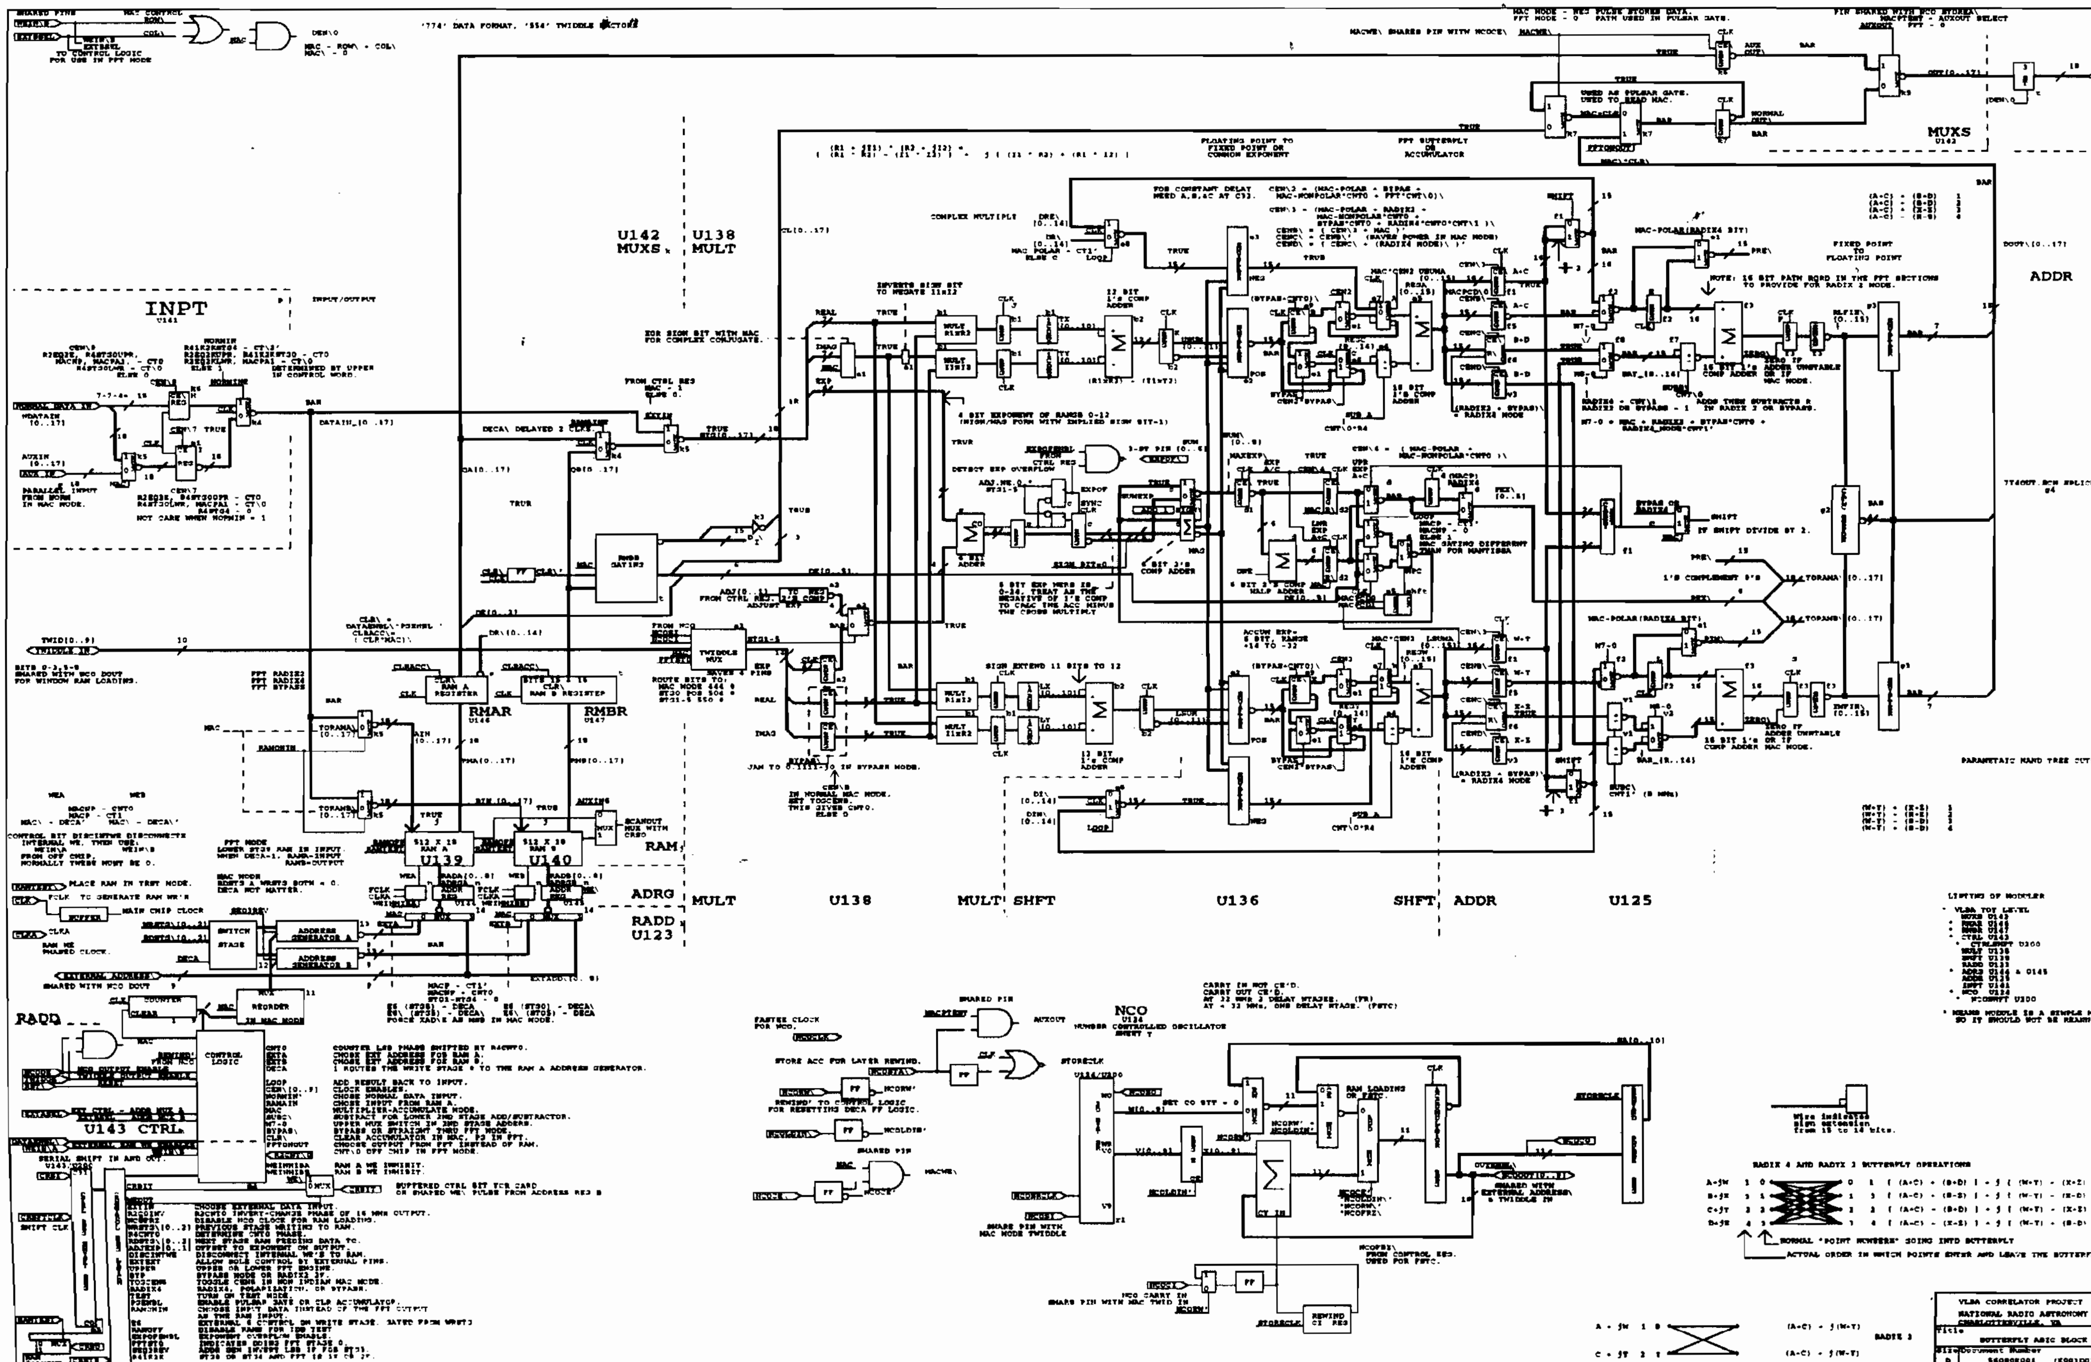

SHARED PINS
TO CONTROL LOGIC FOR USE IN FFT MODE
MAC = ROW\ + COL\
MAC\ = 0
'774' DATA FORMAT, '554' TWIDDLE FACTORS
MAC MODE = NEW PULSE STORES DATA.
FFT MODE = 0 PATH USED IN PULSAR GATE.
MACWE\ SHARES PIN WITH NCOCE\
USED AS PULSAR GATE.
USED TO READ MAC.
FFT BUTTERFLY OR ACCUMULATOR
MUXS
U142
INPT
U141
INPUT/OUTPUT
U142 MUXS
U138 MULT
COMPLEX MULTIPLY
FLOATING POINT TO FIXED POINT OR COMMON EXPONENT
FOR CONSTANT DELAY NEED A,B,&C AT C33.
FIXED POINT TO FLOATING POINT
NOTE: 16 BIT PATH NEED IN THE FFT SECTIONS TO PROVIDE FOR RADIX 2 MODE.
ADDR
INVERTS SIGN BIT TO NEGATE I1xI2
12 BIT 1'S COMP ADDER
FOR SIGN BIT WITH MAC FOR COMPLEX CONJUGATE.
4 BIT EXPONENT OF RANGE 0-12
DETECT EXP OVERFLOW
16 BIT 1'S COMP ADDER
IF SHIFT DIVIDE BY 2.
1'S COMPLEMENT 0'S
MAC-POLAR(RADIX4 BIT)
TWIDDLE MUX
ROUTE BITS TO:
MAC MODE 444
ST30 FOR 504
ST31-5 550
ACCUM EXP=
5 BIT, RANGE
+14 TO -32
SIGN EXTEND 11 BITS TO 12
JAM TO 0.1111-10 IN BYPASS MODE.
TWID[0..9]
TWIDDLE IN
BITS 0-3,5-9
SHARED WITH NCO DOUT
FOR WINDOW RAM LOADING.
RAM A REGISTER
RAM B REGISTER
RMAR
U146
RMBR
U147
512 X 18 RAM A
U139
512 X 18 RAM B
U140
RAM
ADRG
RADD
U123
MULT
U138
MULT
SHFT
U136
SHFT
ADDR
U125
PLACE RAM IN TEST MODE.
MAIN CHIP CLOCK
SWITCH STAGE
ADDRESS GENERATOR A
ADDRESS GENERATOR B
EXTERNAL ADDRESS
SHARED WITH NCO DOUT
RADD
REORDER IN MAC MODE
CONTROL LOGIC
U143 CTRL
SERIAL SHIFT IN AND OUT.
COUNTER LSB PHASE SHIFTED BY R4CNTO.
CHOSE EXT ADDRESS FOR RAM A.
CHOSE EXT ADDRESS FOR RAM B.
ADD RESULT BACK TO INPUT.
CLOCK ENABLES.
CHOSE NORMAL DATA INPUT.
CHOSE INPUT FROM RAM A.
MULTIPLIER-ACCUMULATE MODE.
BYPASS OR STRAIGHT THRU FFT MODE.
CLEAR ACCUMULATOR IN MAC, P3 IN FFT.
CHOOSE OUTPUT FROM FFT INSTEAD OF RAM.
RAM A WE INHIBIT.
RAM B WE INHIBIT.
CHOOSE EXTERNAL DATA INPUT.
NCO
U134
NUMBER CONTROLLED OSCILLATOR
FASTER CLOCK FOR NCO.
STORE ACC FOR LATER REWIND.
SHARED PIN
STORECLK
RAM LOADING OR FSTC.
SHARED WITH EXTERNAL ADDRESS & TWIDDLE IN
REWIND CI REG
NCO CARRY IN
SHARE PIN WITH MAC TWID IN
PARAMETRIC NAND TREE OUT
LISTING OF MODULES
Wire indicates sign extension from 15 to 16 bits.
RADIX 4 AND RADIX 2 BUTTERFLY OPERATIONS
NORMAL "POINT NUMBERS" GOING INTO BUTTERFLY
ACTUAL ORDER IN WHICH POINTS ENTER AND LEAVE THE BUTTERFLY
RADIX 2
VLBA CORRELATOR PROJECT
NATIONAL RADIO ASTRONOMY
CHARLOTTESVILLE, VA
Title
BUTTERFLY ASIC BLOCK
Size
Document Number
D

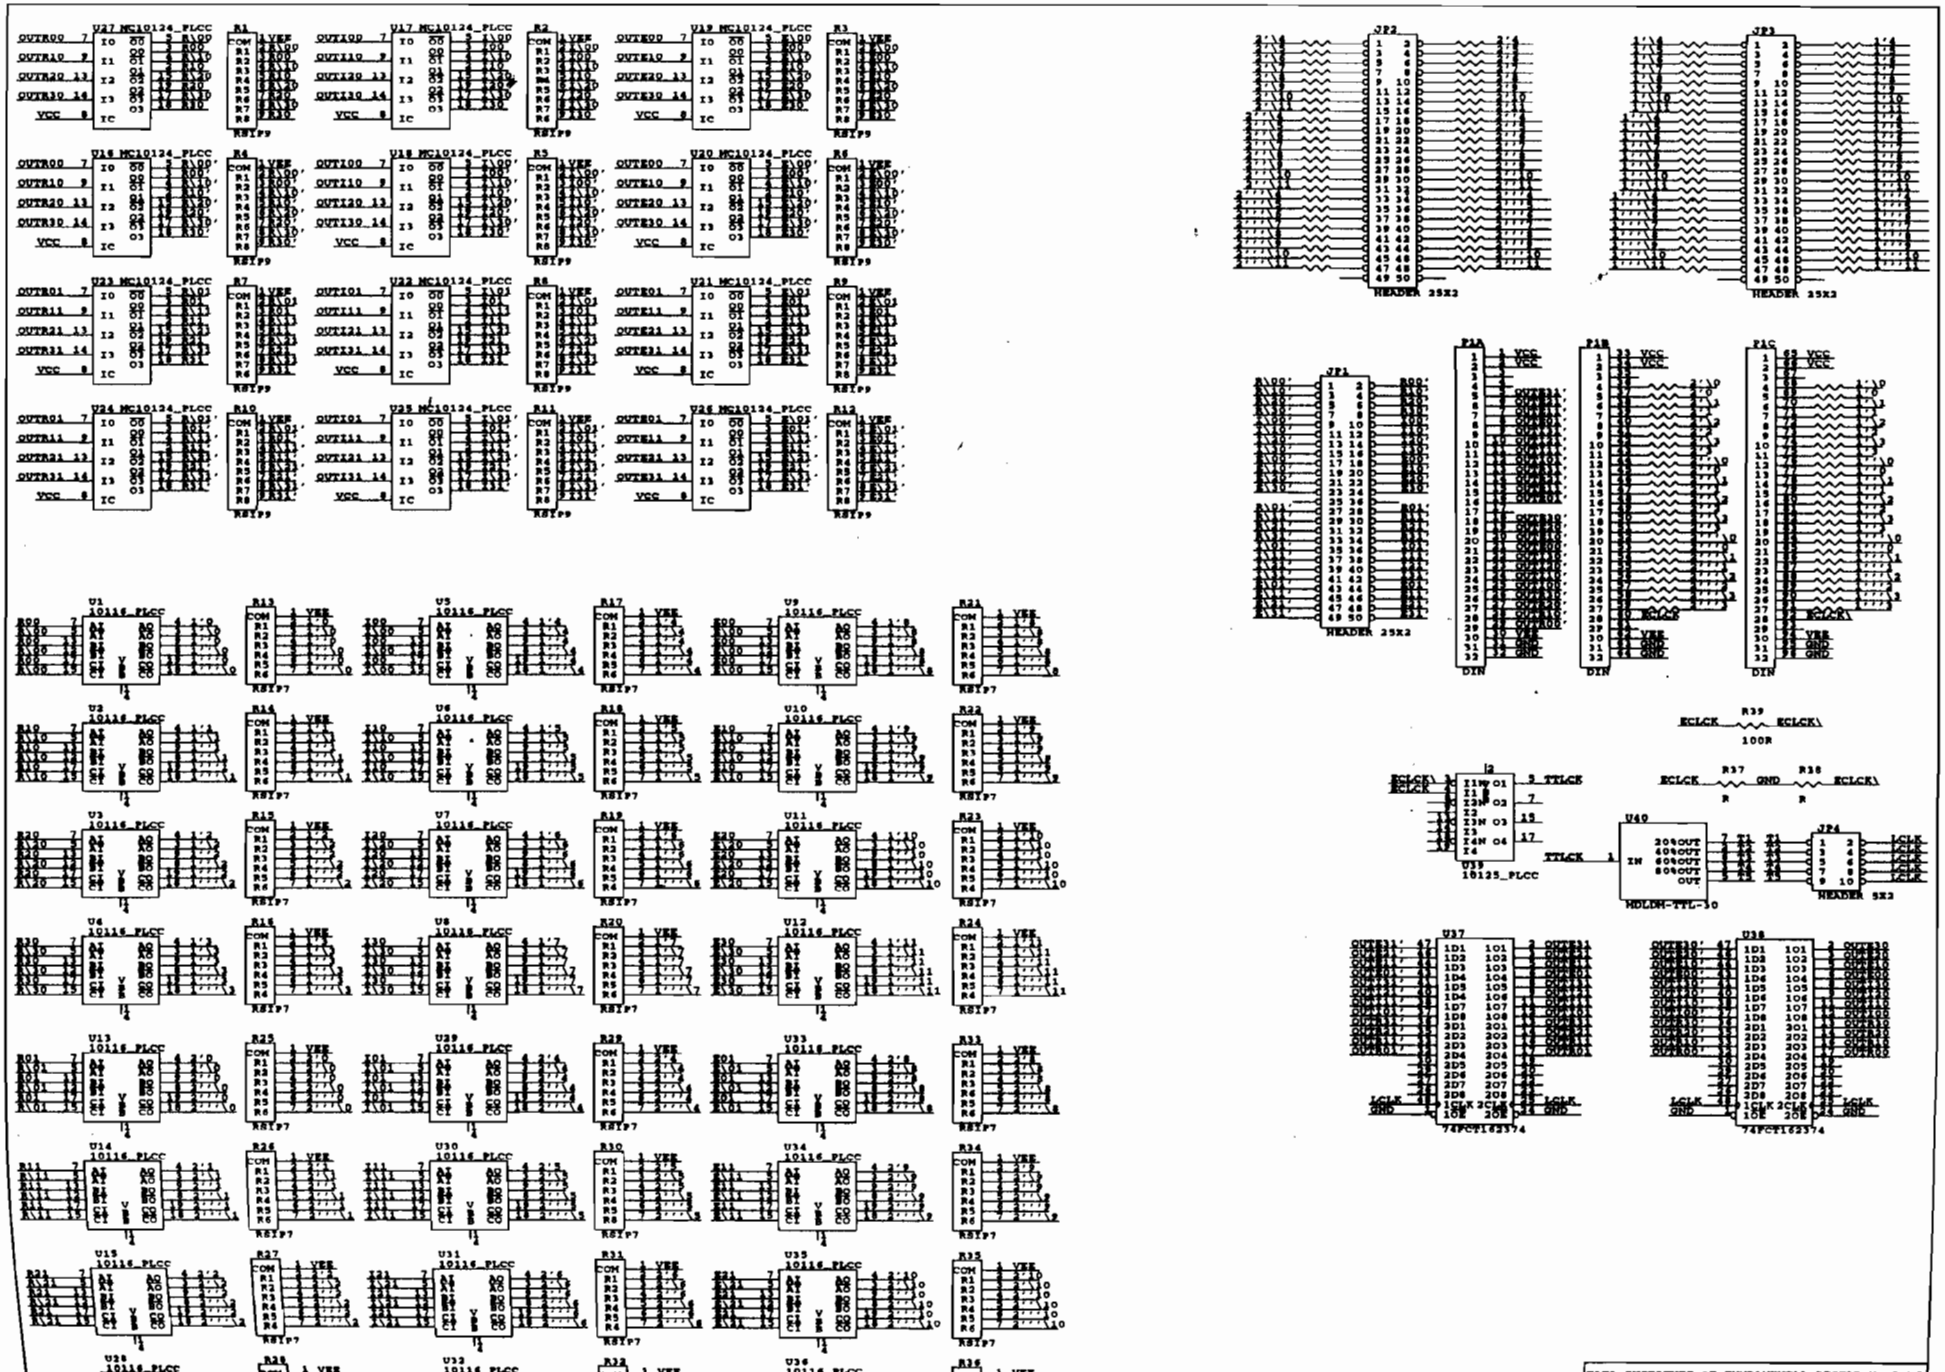

TATA INSTITUTE OF FUNDAMENTAL RESEARCH, PUNE	
GMRT CORRELATOR PROJECT	
Title	TTL-ECL CONVERSION & FAN OUT GENERATION