NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

ISBN 0-315-55487-8

Canada

The University of Alberta

Development of a Digital Correlator for Aperture Synthesis

by
David R. Karpa

A Thesis
submitted to the Faculty of Graduate Studies and Research
in partial fulfilment of the requirements for the Degree of
Master of Science

Department of Electrical Engineering

Edmonton, Alberta
Fall of 1989

University of Alberta
Release Form

Name of Author:	David R. Karpa
Title of Thesis:	Development of a Digital Correlator for Aperture Synthesis
Degree:	Master of Science
Year this Degree Granted:	1989

Permission is hereby granted to The University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or reproduced without the author's written permission.

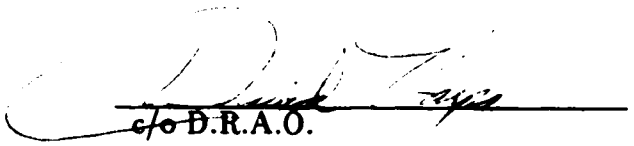

c/o D.R.A.O.
Box 248,
Penticton, B.C.
V2A-6K3

September 11, 1989

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read,
and recommend to the
Faculty of Graduate Studies and Research
for acceptance, a thesis entitled
Development of a Digital Correlator for Aperture Synthesis
submitted by
David R. Karpa
in partial fulfilment of the requirements for the degree of
Master of Science.

D. Routledge
(Supervisor)

T. Landecker
(Supervisor)

J. Vanderlinde

B. J. Kolasa

Date: Oct. 9, 1989

To my best friend, Janiese

Abstract

The Dominion Radio Astrophysical Observatory is expanding and upgrading its synthesis telescope. The existing analog continuum back-end is being replaced by a wider-bandwidth, high-resolution digital correlator.

This thesis describes the development of this digital system. The correlator will support the wider 30 MHz bandwidth, form the cross-polarization products required for polarimetry, and correlate all possible antenna-pairs for 8 dishes. Each of the 1184 correlator channels multiplies two 4-bit numbers at 20 MSPS and can accumulate for several minutes. The 14-level (4-bit) quantization results in a very high correlator efficiency (0.986).

The DSP design uses a parallel-pipelined architecture to maintain high data rates throughout the system (10 billion multiply-accumulate operations per second). The correlation units are modular so the performance of the system scales linearly with size.

The system is designed to require minimal computer control; time-critical and low-level tasks are hidden from the computer. Built-in system test features permit on-line diagnostics for monitoring the performance of the array.

Acknowledgements

The people at the Dominion Radio Astrophysical Observatory have made a lasting impression on me; their kindness and encouragement will never be forgotten. They are to be commended for maintaining a challenging and friendly research environment.

The advice from my supervisors, Tom Landecker and Peter Dewdney at the NRC, as well as David Routledge and Fred Vaneldik at the U of A, kept me off *most* of the dead-end streets. I am just beginning to realise the importance of the administrative skills they demonstrated throughout the project.

The coffee room was where the most important design decisions were discussed: Gary Hovey's engineering expertise was appreciated at every stage of the design; Chris Purton's passion for the science is infectious — a powerful motivator; Everet Sheehan has an amazing tolerance for fickle engineers.

Thank you Elizabeth Landecker for officiating at the arrival ceremonies for Jared and Alanna.

I am grateful for the almost endless financial support of the Herzberg Institute of Astrophysics, The Alberta Heritage Foundation and the Low-Noise Amplifier Group at the University of Alberta. Without them, this project would not have been.

Most of all, I owe a lifelong debt of gratitude to my friend and wife, Janiese, for her steadfast love and support.

Contents

List of Figures

List of Tables

Chapter 1

In the beginning...

1.1 From Babylon to British Columbia

Astronomy is often cited as being one of Man's oldest scientific endeavors. The earliest records are those of Babylonians who meticulously charted the movements of the heavenly bodies. While one might question whether the ancients were sorcerers or scientists, their efforts have formed the rich heritage of the science. The quest to see and understand the heavens is one that is certainly thousands of years old. Through the years, this quest has been aided by technological advances — indeed, astronomy has often been the motivation for such advances. Recent history has been characterized by a rapid growth in technology and the sciences so that today scientists focus on specific aspects of astronomy such as astrometry, astrophysics and interstellar chemistry. Naturally, the partnership between astronomy and technology continues to be a fruitful one as new techniques allow us to see and understand the old universe in new ways.

One such technique is *radio astronomy*. Processes and objects unobserved prior to 1932, when Karl Jansky [Jan33] detected radio emmisions from the galactic plane, now form a vital part of our view of the universe. By 1950, radio astronomy was steadily gaining momentum as more and more scientists realised its significance. At first, better meant bigger; the trend was towards

building larger and larger antennas but by 1964 this was giving way to a technique called aperture synthesis.

Aperture synthesis uses a group of small antennas to mimic the performance of an impossibly large antenna. By operating antennas in concert, we are now able to observe features that would be unseen with a single antenna unless it were miles in extent. The synthesis technique allows astronomers to view extraterrestrial radio phenomena with resolution comparable to (and in some cases surpassing) the performance of optical instruments.

Today, our appetite for understanding the universe is as healthy as it ever was. Often communities of specialists, both scientific and technical, combine their skills and intelligence for greater effectiveness. The Dominion Radio Astrophysical Observatory (DRAO) in British Columbia is evidence of this combined effort. Operating under the National Research Council's Herzberg Institute of Astrophysics, its mandate is to provide the scientific community with high-calibre facilities for radio astronomy. Besides operating and maintaining several radio telescopes, the scientists and engineers at the DRAO are actively involved in developing the future tools for radio astronomy. One aspect of this lies in applying new technology to the design of new instruments and the refurbishment of existing ones. The observatory interacts with the scientific community to the extent that it often enlists the cooperation of universities when challenging problems arise.

This thesis describes the design of a signal processing system for the DRAO aperture synthesis telescope. The system design involved selecting a strategy based on current technology as well proving its feasibility by building and testing representative system components. The system described herein will improve on an existing signal processor which is unable to meet the growing demands of the instrument.

The following chapters will address how the system demands were met and why certain techniques were chosen. To provide a foundation for such discussions, the remainder of this chapter will focus on three areas: first, a review of the basic principles of aperture synthesis relevant to this project and to the synthesis telescope at the DRAO[1]; second, a brief description of

[1]Readers wanting a thorough treatment of the theory and practice of interferome-

DRAO's synthesis telescope with emphasis on the aspects relating to this project; third, an introduction to the performance requirements of the new signal processing system.

1.2 A Review of Interferometry

This section will review some basic properties of interferometry, including instrumental considerations such as system bandwidth and tracking, and will close with a description of how the interferometer is used for polarimetry and aperture synthesis.

Interferometry is a technique where the signals from two antennas are used to generate a multiple fan-beam response—an interference pattern. Typically, the signals are multiplied by each other and integrated for long periods of time. Since each signal has units of volts, their product can be interpreted as power. The response of the interferometer pair will be referred to as the *cross-power*, generally a complex number. Figure 1.1 shows an interferometer pair.

1.2.1 The Monochromatic Interferometer

The interferometer behavior is readily understood by considering the scenario depicted in Figure 1.2. Suppose, for now, antennas 1 and 2 are identical and isotropic.

Consider the effect of a monochromatic radio source of wavelength λ located at infinity in the direction $\hat{k}$ (a unit vector). Since both antennas are equally sensitive in all directions, the response pattern is azimuthally symmetric about the axis defined by the baseline vector $\vec{B}_\lambda$. The subscript λ will indicate that $\vec{B}_\lambda$ is measured in units of the source's wavelength.

First consider the case where the source lies on the meridian plane (the

try and synthesis should look to "Interferometry and Synthesis in Radio Astronomy" by Thompson, Moran and Swenson [TMS86]

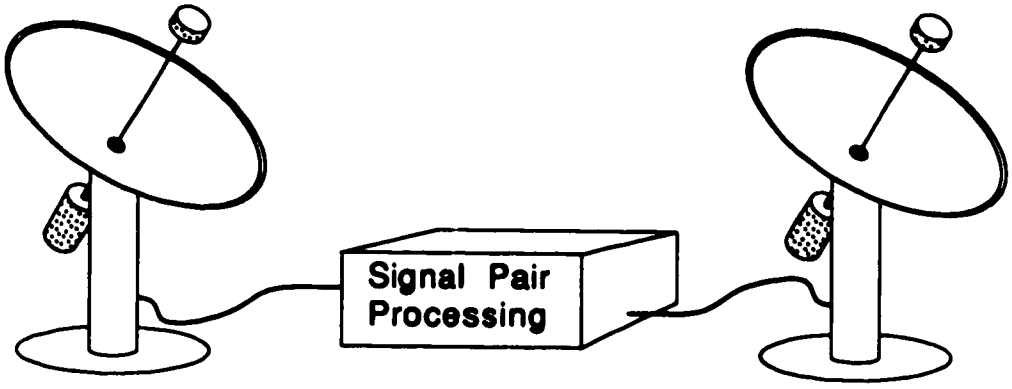

Figure 1.1: **An Interferometer Pair:** The product of two antenna signals is the cross-power. The pair's response is an interference pattern caused by the relative phase delay of signals arriving at the two antennas. The cross-power result can be negative or positive.

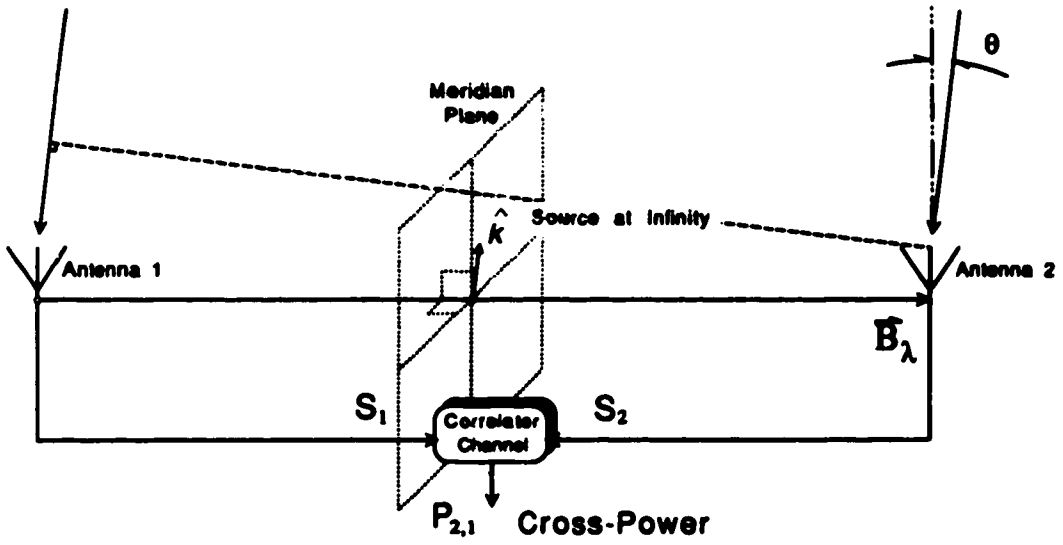

Figure 1.2: **The Isotropic Interferometer Pair:** The isotropic antennas respond uniformly to signals from any direction. The hypothetical transmitter is located at infinity in the direction $\hat{k}$.

meridian plane is the perpendicular bisector of the baseline vector $\vec{B}_\lambda$). In this case signals S_1 and S_2 are identical; the cross-power is equal to the power that would be measured at either antenna independently.

Meridian Case:

$$\begin{aligned} & & S_1 &= S_2 \\ \text{cross-power} & & P_{2,1}|_{meridian} &= S_1 S_2^\dagger \\ & & &= (S_1)^2 = (S_2)^2 \\ & & &= P_m \end{aligned}$$

† denotes conjugation

Now consider the effect of moving the source to a point off the meridian plane in the direction $\hat{k}$. Signal S_2 is now shifted in phase with respect to signal S_1 . The phase shift is produced by the path difference of $\vec{B}_\lambda \cdot \hat{k}$ wavelengths.

$$S_2 = S_1 e^{j2\pi \vec{B}_\lambda \cdot \hat{k}} \tag{1.1}$$

$$P_{2,1} = \mathrm{Re}(P_m e^{-j2\pi \vec{B}_\lambda \cdot \hat{k}}) \tag{1.2}$$

If we define the angle θ as the departure angle from the meridian plane the expression for the interferometer response becomes:

$$P_{2,1} = P_m \cos(2\pi B_\lambda \sin\theta) \tag{1.3}$$

The expression is simplified by defining a new parameter $\zeta = \sin\theta$.

$$P_{2,1} = P_m \cos(2\pi B_\lambda \zeta) \tag{1.4}$$

As shown by equation 1.4 the monochromatic interferometer response pattern is a sinusoidal function of the variable ζ which has values between -1 and $+1$. The antenna separation, B_λ, determines the spacing between the peaks ($+P_m$) and troughs ($-P_m$) of the interferometer response pattern.

It is worthwhile to note here that if a phase lag of $\frac{\pi}{2}$ was added to S_2 in equation 1.1 the response would be a sine function instead of the cosine in equation 1.4. The phase shifted version will be referred to as the quadrature signal (denoted by a prime $'$).

$$S_2' = S_2 e^{-j\frac{\pi}{2}} \tag{1.5}$$

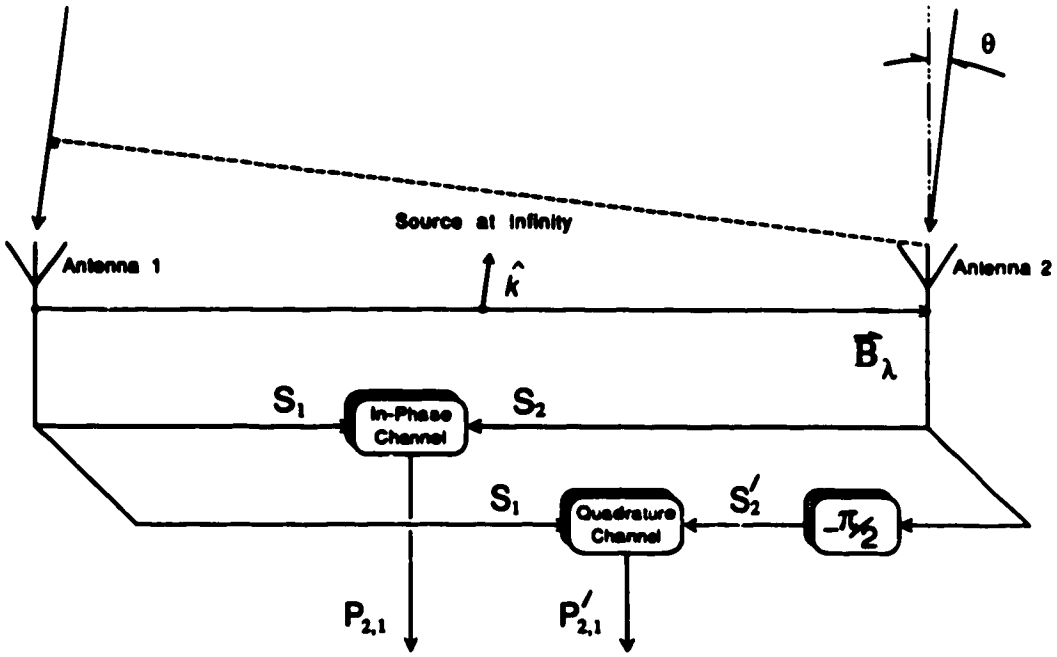

Figure 1.3: The Quadrature Channel Interferometer Pair: The response pattern of the quadrature correlator channel is a *sine* function of the variable ζ. Its peaks coincide with the nulls of the in-phase channel's pattern and vice versa. The measurements made with the two channels are independent, further, they are orthogonal. The power radiated from any point in the sky will be detected by one of these two channels.

This produces the quadrature-channel interferometer response:

$$P'_{2,1} = \mathrm{Re}\left(P_m e^{-j2\pi \vec{B}_\lambda \cdot \hat{k} + j\frac{\pi}{2}}\right) \tag{1.6}$$

$$P'_{2,1} = P_m \sin(2\pi B_\lambda \zeta) \tag{1.7}$$

Figure 1.3 shows a functional representation of the interferometer with the quadrature channel added. Because the peaks and troughs of the in-phase response $P_{2,1}$ coincide with the nulls of the quadrature response, $P'_{2,1}$, (and vice versa), they are independent response patterns.

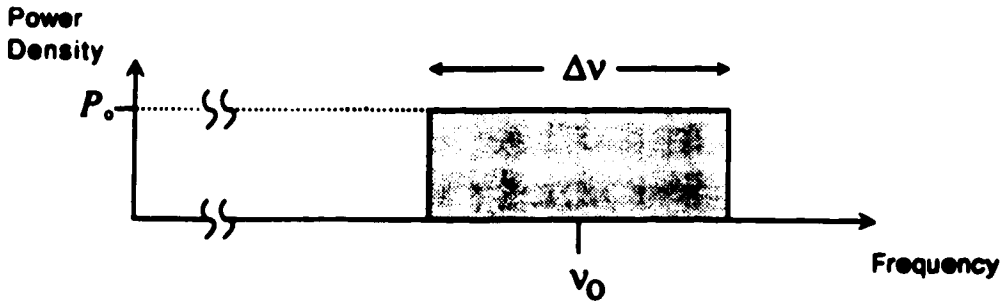

Figure 1.4: **A Rectangular Frequency Response:** The system frequency response pattern has a bandwidth of $\Delta\nu$ at center frequency ν_o. The abscissa represents spectral power density ... $P_m = P_o\Delta\nu$.

1.2.2 The Bandwidth Pattern

We will now examine the effect of the interferometer bandwidth. The simple monochromatic model discussed above can be expanded to include a range of frequencies. If we cast equations 1.4 and 1.7 in exponential form, the response can be represented as a complex value.

$$\hat{P}_{2,1} = P_{2,1} + jP'_{2,1} = P_m e^{j2\pi B_\lambda \zeta} \tag{1.8}$$

A useful result is obtained by examining the interferometer response to a rectangular spectral power distribution such as the one shown in Figure 1.4. Suppose that the power of the imaginary source at infinity is spread uniformly over the bandwidth. To derive the interferometer response in the broadband case we should rewrite equation 1.8 in terms of frequency ν.

$$d\hat{P}_{2,1} = P_o e^{j2\pi \frac{B}{c}\nu\zeta} d\nu \tag{1.9}$$

The broad-band interferometer response pattern is the integral sum of the responses of the elemental interferometers $d\hat{P}$ each with power $P_o d\nu$.

$$\hat{P}(\Delta\nu) = P_o \int_{\nu_o - \frac{\Delta\nu}{2}}^{\nu_o + \frac{\Delta\nu}{2}} e^{j2\pi \frac{B}{c}\nu\zeta} d\nu \tag{1.10}$$

After integration:

$$\hat{P}(\Delta\nu) = \overbrace{(P_o\Delta\nu)}^{1}\underbrace{e^{j2\pi\frac{B}{c}\nu_o\zeta}}_{2}\overbrace{\left(\frac{\sin\left(\pi\Delta\nu\frac{B}{c}\zeta\right)}{\pi\Delta\nu\frac{B}{c}\zeta}\right)}^{3} \tag{1.11}$$

The resulting expression has three terms: term 1 is the total incident power from the source; term 2 is the monochromatic interferometer response at frequency ν_o; term 3 is referred to as the bandwidth pattern (also as the delay beam). The bandwidth pattern attenuates the two-dimensional sinusoidal response (term 2) of the monochromatic interferometer as ζ departs from the meridian plane ($\zeta = 0$). Its effect is most noticeable when the bandwidth and baselines are both large.

1.2.3 The non-isotropic interferometer: tracking the source

By replacing the fictitious isotropic antennas with realistic antennas, further characteristics of interferometers can be developed. Actual antennas have non-uniform directional properties. The directional sensitivity of an antenna is called its primary pattern and the direction of greatest sensitivity is referred to as its main beam. During typical observation sessions, the antennas track the region of interest keeping the main beams centered on a *point* on the celestial sphere; this will be the image center. The response pattern for the non-isotropic interferometer is obtained simply by adding a fourth term (expressing the directional properties of the antennas) to equation 1.11.

$$\hat{\mathbf{P}} = (P_o\Delta\nu)\,e^{j2\pi\frac{B}{c}\nu_o\zeta}\left(\frac{\sin\left(\pi\Delta\nu\frac{B}{c}\zeta\right)}{\pi\Delta\nu\frac{B}{c}\zeta}\right)\cdot\underbrace{G(\zeta,\eta)}_{4} \tag{1.12}$$

In equation 1.12, η is a parameter, like ζ, that results from projecting the curvilinear coordinates θ and ϕ onto a flattened image plane (see Thompson et al. [TMS86, page 79]); ϕ is the departure angle between $\hat{k}$ and the plane containing the baseline and the image center (the baseline plane). The image

plane is tangent to the celestial sphere at the *point* mentioned above (the image, baseline and meridian planes all intersect orthogonally at the image center *point*).

Pointing the Delay Beam

The antenna beam characteristics are defined by $G(\zeta, \eta)$. As the region of interest moves across the sky from east to west, the antenna beams are made to *track* this motion; motorized antenna mounts keep each antenna directed at the field center. Like the physical antenna beams, the bandwidth pattern (delay beam) also has directional properties. A re-examination of the function arguments in equation 1.12 is useful at this point:

$$\frac{B}{c}\zeta = \frac{B}{c}\sin\theta = \tau_g \tag{1.13}$$

Term 3 of equation 1.12 is simplified by substituting equation 1.13, the geometric delay:

$$\frac{\sin\left(\pi\Delta\nu\frac{B}{c}\zeta\right)}{\pi\Delta\nu\frac{B}{c}\zeta} = \frac{\sin\left(\pi\Delta\nu\tau_g\right)}{\pi\Delta\nu\tau_g} \tag{1.14}$$

τ_g is the geometric delay between signals arriving at the two antennas. (The bandwidth pattern has a maximum when the relative delay is zero: $\tau_g = 0$). For this reason, the bandwidth pattern is often called the delay beam. To attain practical sensitivity, the delay beam must also *track* the region of interest. This can be done by artificially delaying the signal from one antenna before the cross-power measurement is made. Figure 1.5 shows the interferometer with a delay compensation system for tracking with the delay-beam. As the earth rotates, the antenna beams track the source (the antennas are in constant motion); the delay-beam is also made to track the source by continually adjusting the instrumental delay τ_i to cancel the geometric delay τ_g. The interferometer response can be written in terms of these delay values:

$$\hat{\mathbf{P}} = (P_o\Delta\nu)\underbrace{e^{j2\pi\nu_o(\tau_g-\tau_i)}}_{1}\overbrace{\operatorname{sinc}\left(\Delta\nu\left(\tau_g-\tau_i\right)\right)}^{2}\cdot G \tag{1.15}$$

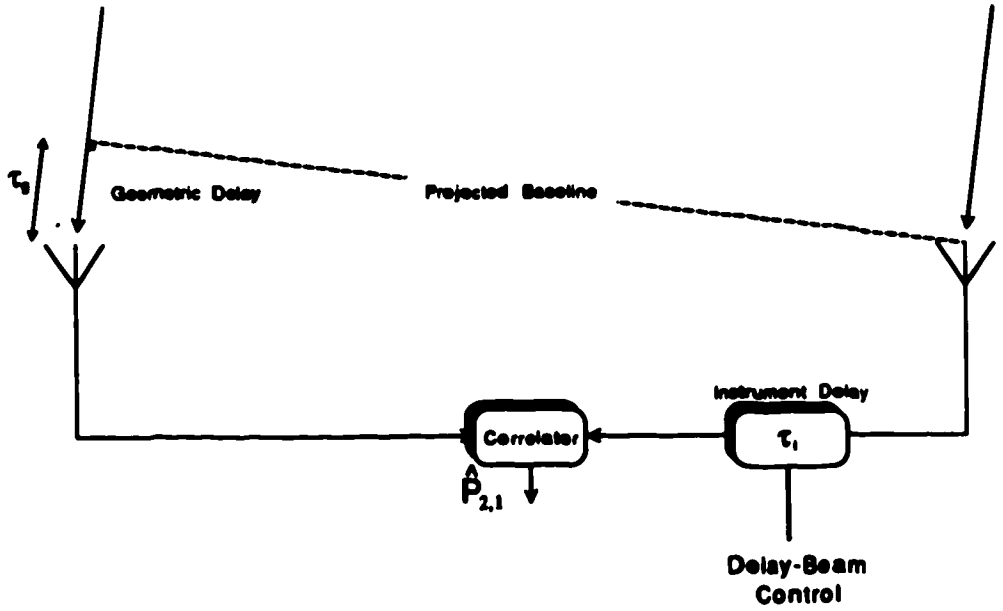

Figure 1.5: Pointing the Delay-Beam: The delay-beam must track the field center. The compensating delay, τ_i, is constantly adjusted to cancel the geometric delay, τ_g, thus keeping the delay-beam aligned with the primary beams. This compensation *projects* the antenna baseline onto a plane facing the source.

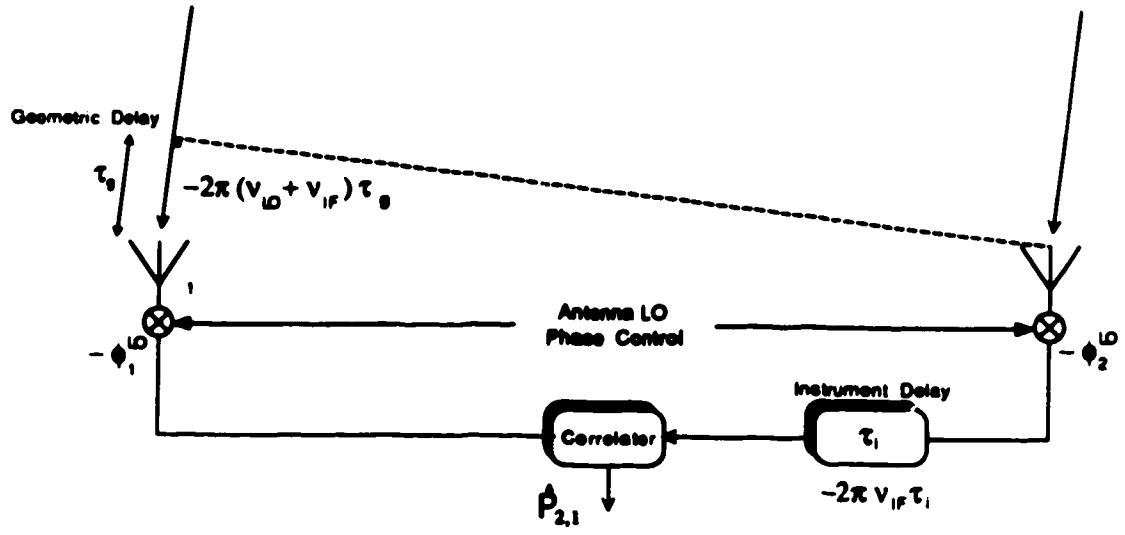

Figure 1.6: Phase Path Equalization: Phase path compensation is applied at the RF-to-IF mix stage at antenna 2. The local oscillator phase difference, $\phi_1 - \phi_2$, is constantly adjusted to cancel the relative phase delay between S_1 and S_2.

Recall that $\tau_g = \tau_i$ is only true for the center of the delay-beam and that term 2 attenuates the response for sources where $\tau_g \neq \tau_i$ (off the meridian plane). Term 1, the isotropic interferometer response, represents the phase difference between the signals from sources in the field of view.

Pointing the Interference Pattern

It is desirable to keep the interference pattern (term 1 of equation 1.15) fixed with respect to the image center to allow long integrations of the cross-power. This is accomplished by equalizing the phase paths from the source to the correlator for each antenna. Achieving this for the tracking interferometer is more difficult than equation 1.15 would imply.

A further step must be taken towards defining the characteristics of an actual interferometer. Because of cable losses and dispersion, the compensating instrumental delay shown in Figure 1.5 is impractical at the high radio

frequencies. Signal delay is often done after the radio frequency (RF) band has been mixed down to a more manageable intermediate frequency (IF). Because the phase path at RF ($2\pi\nu_{RF}\tau_g$) is different from that at IF ($2\pi\nu_{IF}\tau_i$), some phase path compensation is needed. It is possible, by controlling the phase of the local oscillators, to apply the correction during RF-to-IF mixing. The phase path that each signal experiences is shown in Figure 1.6. To track the source with the interference pattern as it moves across the sky the following condition must be maintained: the phase-lags, $\Delta\Phi$, in both signal path must be equal.

$$\Delta\Phi_1 = \Delta\Phi_2 \tag{1.16}$$

Signal 1 phase is a result of the geometric delay at RF, while Signal 2 phase is a result of the instrument's compensating delay at IF (both signal phases are also dependent on their local oscillator phases, ϕ^{LO}):

$$-2\pi\nu_{RF}\tau_g - \phi_1^{LO} = -2\pi\nu_{IF}\tau_i - \phi_2^{LO} \tag{1.17}$$

Since $\nu_{RF} = \nu_{IF} + \nu_{LO}$:

$$-2\pi(\nu_{IF} + \nu_{LO})\tau_g = -2\pi\nu_{IF}\tau_i + (\phi_1^{LO} - \phi_2^{LO}) \tag{1.18}$$

The relative phase between the two local oscillator signals must be adjusted to be:

$$\phi_1^{LO} - \phi_2^{LO} = \underbrace{2\pi\nu_{LO}\tau_g}_{1} + \underbrace{2\pi\nu_{IF}(\tau_g - \tau_i)}_{2} \tag{1.19}$$

Equation 1.19 shows the corrective phase required: term 1 is a result of performing the delay compensation at a frequency other than the original RF frequency; term 2 is zero when $\tau_g = \tau_i$ (this requires the delay system to be very precise). For delay systems where the instrumental delay, τ_i, is not continuously variable[2], the relative phase, $\phi_1^{LO} - \phi_2^{LO}$, can be used to substitute phase for delay; such a correction is exact for only one frequency, usually chosen to be the center frequency of the correlation bandwidth, ν_{IF}^c.

$$\phi_1^{LO} - \phi_2^{LO} = 2\pi\nu_{LO}\tau_g + 2\pi\nu_{IF}^c(\tau_g - \tau_i) \tag{1.20}$$

[2]The compensating delay at the DRAO is a selection of cables of various lengths, each of which can be electronically inserted or removed from the signal path. The computer selects the cable combination which most nearly matches τ_g. Residual phase correction is done for the center of the IF band ($\nu_{IF}^c = 30.0$MHz).

Mapping the Sky

This short discussion will review the interferometer characteristics discussed above in terms of their impact on the imaging process.

Many interferometer measurements are made before there is enough information to make a map of the sky brightness. The mapped region is defined by the primary beam width of the antennas (assumed identical). Within the main beam, the response (for an interferometer pair) is modulated by the striped phase-interference pattern. This striped response pattern allows the complex interferometer to measure a spatial Fourier component of the two-dimensional image brightness. The result, obtained by integrating the cross-power, corresponds directly with the amplitude and phase of the image's Fourier component. The orientation and spacing of the striped response is dependent on the direction and magnitude respectively of the antenna baseline vector with respect to the image coordinates. Two mechanisms are used to alter the antenna baseline vectors: antennas are physically placed to produce the desired baselines (the DRAO has two movable antennas), and the motion of the Earth changes the orientation and projected magnitude of the antenna baseline with respect to the image. When enough spatial frequencies have been measured (at the DRAO this takes about 35 days of observing), the image can be reconstructed using a two-dimensional Fourier transformation.

During any given interferometer measurement (lasting several seconds), the striped interference pattern must be held stationary with respect to the image coordinates. This requires continuous precise control of the phase and delay at each antenna.

Polarimetry

Except for spectral-line emissions, the signals received from space are broadband noise. In some cases the signals are partially linearly polarized. To measure (i.e. make a map of) the polarized components of the sky brightness requires dual-polarization antennas and a more elaborate correlator system. The polarization parameters that are typically mapped are the Stoke's

parameters:

I is the total incident power.

U and Q are linearly polarized components of power.

V represents the circularly polarized component.

Some astronomical objects exhibit small degrees of linear polarization. The best measurements of linear polarization are made with circularly polarized receivers (feeds); the antennas at the DRAO receive left and right hand circularly polarized radio waves.

To accomplish the polarization measurement (polarimetry), the interferometer pairs must each form the cross-polarization correlation products (left–right and right–left) in addition to the regular correlation products (right–right and left–left).

Phase Switching

Signals received from space are typically 50 decibels weaker than the noise generated in the front-end antenna electronics; this is the signal level received from the weakest source detectable by the telescope. The correlation process is a very sensitive measurement technique used to detect minute amounts of signal similarity. In the perfect system, the only similarity between two antenna signals would be due to the weak extraterrestrial radiation. In practise, similarity also results from d.c. signal bias, crosstalk between lines or from system sources of EMI (switching power supplies or high-speed logic). A technique for reducing the impact of these system imperfections on the correlation results is called *phase-switching*. Alternate correlation integrations are made with one antenna signal inverted. The inversion is usually done at the antenna focus box by shifting the LO phase by 180°. This inverts the interferometer response; ideally, the effect of the system imperfections will be the same (non-inverted) for successive integrations. Because the sense of the corruption has changed with respect to the desired measurement, the

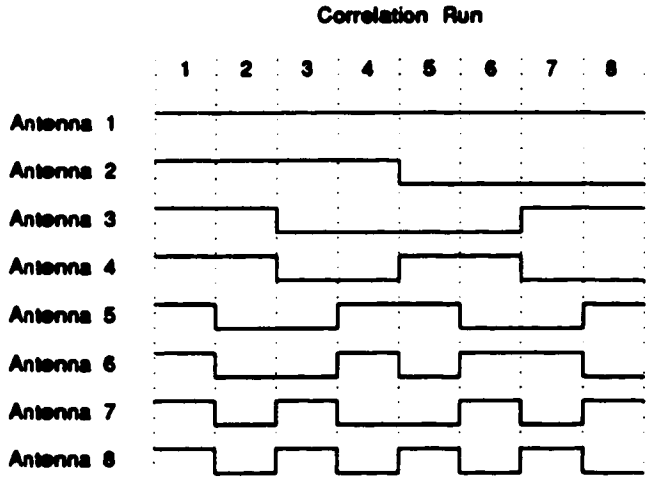

Figure 1.7: Phase Switching for 8 Antennas: The antennas are phase-switched on successive correlation-runs according to the Walsh fuctions.

unwanted component can be eliminated. The technique is effective to the extent that the sources of corruption are stable (stationary).

Phase-switching is currently used at the DRAO[3] where the two fixed dishes are 'correlated' with the two movable ones (track mounted). On alternate correlation-runs (integrations), the track-mounted antennas are phase-switched (relative to the fixed ones).

The phase-switching becomes more complex in an array forming all possible antenna-pair combinations. Walsh functions are bimodal (+ or −) orthogonal patterns which provide an efficient algorithm for phase-switching any number of antennas. Figure 1.7 shows the switching pattern for 8 antennas.

The phase-switch process involves modulating (inverting) the signal as soon as possible in the signal path (at the focus) and then demodulating at a point where the signal is no longer susceptible to contamination. With an

[3] "...to cover a multitude of sins." — T.L. Landecker

analog correlator system, this point is after correlation (in software); with a digital system, the quantized signals can be demodulated (by hardware) prior to correlation. The latter approach is straightforward and it allows multiple integrations to accumulate in the channel hardware (since the sign of the measurement has already been corrected). This would allow much higher (more effective) phase-switch rates to be attained because channel read-out is not required between runs.

1.3 The DRAO Synthesis Telescope

At the Dominion Radio Astrophysical Observatory (DRAO), the main instrument is an earth-rotation synthesis telescope composed of four 9 meter *dish* antennas on a 600 meter east-west baseline. Two of the antennas are track mounted (movable) allowing any antenna spacing to be formed (see Figure 1.8).

The telescope provides simultaneous mapping of two regions of the radio spectrum: the image (field of view) is 2° across at 1.42 GHz and 7° at 408 MHz. The 1.42 GHz region is mapped by two systems: the first, the spectrometer system, produces 128 maps of the HI (neutral hydrogen) spectral line intensity (each map corresponds to a Doppler frequency shift produced by the different relative velocities of the radiating hydrogen); the second system maps the broad-band (continuum) radiation in two 7.5 MHz bandwidths which straddle the spectral line bandwidth.

The 408 MHz system (4 MHz bandwidth) is a continuum mapping system (see Veidt, et al. [Vei84] and Lo, et al. [Lo84] for details about this system). Both fields of view are quite wide by astronomical standards. This unique property of the DRAO telescope makes it well suited to the study of dispersed structures in the interstellar medium. To help the telescope meet the future requirements of the scientific community, the observatory is taking significant steps to improve the instrument's performance.

The most recent step was the addition of dual polarization processing (previously only left-hand circularly polarized signals were detected); this increases the sensitivity of the telescope by a factor of $\sqrt{2}$.

Figure 1.8: The DRAO Synthesis Telescope: Two of the four dishes can be located anywhere on the 300 meter track. A fifth dish (unpainted) is shown under construction.

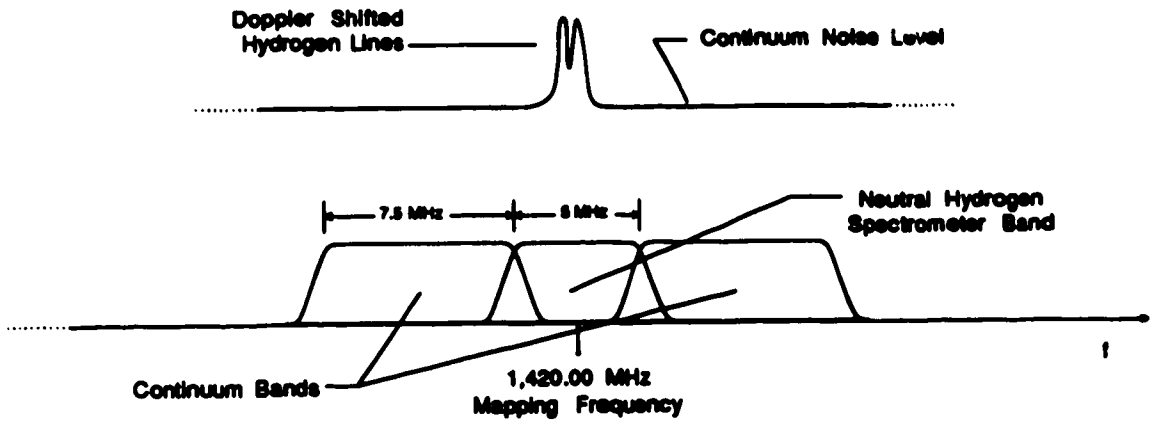

Figure 1.9: Observed Frequencies at 1.42 GHz: The 5 MHz band produces 128 maps corresponding to Doppler shifts of the neutral hydrogen spectral line. The continuum bands measure the *non-spectral* radio brightness. The two 7.5 MHz bands represent the old analog continuum system's frequency response.

A further step will be to use the dual polarization system to perform polarization measurement or polarimetry. This requires more 'back-end' processing; in addition to the visibilities measured by correlating right-right and left-left (present system), the cross-polarization products, right-left and left-right, must be formed for each interferometer pair (8 correlator channels per interferometer pair — 2 channels per complex correlator).

To increase the sensitivity of the telescope's continuum system, the map bandwidth is being increased from the present 15.0 MHz to 30.0 MHz. While this requires some improvement of the telescope's 'front end' (better filters) the 'back-end' processing (correlators) is where most of the improvements are needed. Doubling the system bandwidth will improve the sensitivity of the continuum system by $\sqrt{2}$.

The most visible improvement to the telescope will be three (possibly four) additional antennas. These will be added one at a time, the last scheduled to become active in 1992. This will reduce the observing time from 35 to 12 days. A further advantage of more antennas results from of a technique for *self-calibrating* the telescope. The technique, known as *closure-phase*, be-

comes more effective as the number of antennas in the array increases (see Thompson et al. [TMS86, page 356]).

The closure-parameters, derived from the observations, provide a set of constraints which will guide the subsequent image generation and enhancement. This technique is effective in reducing system phase errors. At the DRAO, the dynamic range of instrument is presently limited by phase errors in the observations. More antennas, together with self calibration techniques such as 'closure-phase', will allow better images (with wider dynamic range) to be produced in less time.

In order for closure-phase to be effective, the telescope must meet an important criterion: phase errors in the observations must reside solely in the antennas (not 'downstream' in the signal distribution and processing systems). This means that the signal distribution paths and correlator channels should be very closely matched. With an an analog system, such performance is not maintainable for even a modest system with 16 correlator channels. Component aging and temperature drift are factors which perpetually change the conditions in the analog system. This is the most desirable aspect of a digital correlation system: after the signals have been digitized (antenna-based) all the processing, whether signal distribution or correlation, is identical and error-free. There are no correlator-based errors in a digital system.

1.4 The New 1420 MHz Continuum Correlator

What are the system-level considerations behind the decision to replace the old continuum correlator system? This project is a direct result of the decision to 'replace' rather than 'reproduce' the existing strategy ... yet another analog system is superseded by its digital counterpart. The items discussed here will reflect the previous section's introduction to the DRAO synthesis telescope and its performance objectives.

1.4.1 The Expanding Array

The main reason for re-evaluating the continuum correlator design is the demand for many more correlators than presently exist. The addition of 4 antennas (and forming all possible antenna-pairs) will increase the number of interferometers from 4 to 21. Adding correlator channels for the cross-polarization products required by polarimetry effectively doubles the requirement to 42. Virtually the entire correlator system would be new, even if the old strategy were copied.

A related reason to evaluate the old system comes as a result of the rapid advances being made in electronics technology. The cost constraints associated with various designs (both digital and analog) have changed dramatically since 1975 when the analog system was designed. The *digital trend* is driven by the still-increasing performance-to-cost ratio of digital devices; a digital system is much cheaper now than it was then.

Analog components have also seen remarkable improvements, implying that a similar strategy to the existing one could provide better performance for an equivalent cost. The system described herein (the new continuum correlator) uses a substantial number of analog components in the IF processing prior to digitization (see Chapter 4). At the time of design (1986) the system strategy represented a balance between analog and digital complexity for the larger array. For example, quadrature channel signals are more easily generated in the IF analog hardware than in digital circuitry (Hilbert transform filters).

The DRAO needs a significantly larger continuum correlator. The new system must be easily manufactured, expandable, and reliable. As well as present system's requirements, the new correlator should provide lity to permit — even encourage — further improvements to 's performance.

1.4.2 Improving the Performance Parameters

The smallest signal detectable by an interferometer pair is proportional to (see Thompson et al [TMS86, page 160]):

$$\Delta s \propto \frac{kT}{\eta_c A_{eff} \sqrt{\Delta \nu \tau}} \tag{1.21}$$

k	is the Boltzmann constant
T	is the system noise temperature
η_c	is the correlator efficiency
A_{eff}	is the effective antenna area
$\Delta \nu$	is the correlated bandwidth
τ	is the correlation time

The sensitivity of the continuum system can be improved by widening each of the 7.5 MHz bands to 15.0 MHz. The disadvantage of the wider correlation bandwidth is the narrow delay beam this produces. For the larger antenna spacings, the sources at the edges of the field are severely attenuated (the improved sensitivity applies only to sources at the field center). This results in a 'smearing' of sources at the edge of the image. The way around this dilemma is to break the bandwidth into smaller sub-bands, each with its own correlation system. Because the normalized antenna baseline, B_{λ_j}, is different for each sub-band (λ_j is the center wavelength of each band), sub-maps must be scaled to compensate. The scaled sub-maps can be combined to obtain a map whose sensitivity has been increased over the entire field of view. Figures 1.11 and 1.10 show three correlation bandshapes and their effect on the delay beam.

The bandwidth processing approach chosen for the new correlator is shown in Figure 1.10 (c). The two 15.0 MHz bandwidths are each filtered into two 7.5 MHz sub-bands. Each of the four sub-bands is digitized and correlated independently.

Another good reason for correlating bands 7.5 MHz wide was the state of technology in 1985. Reasonably priced (less than $100) analog-to-digital

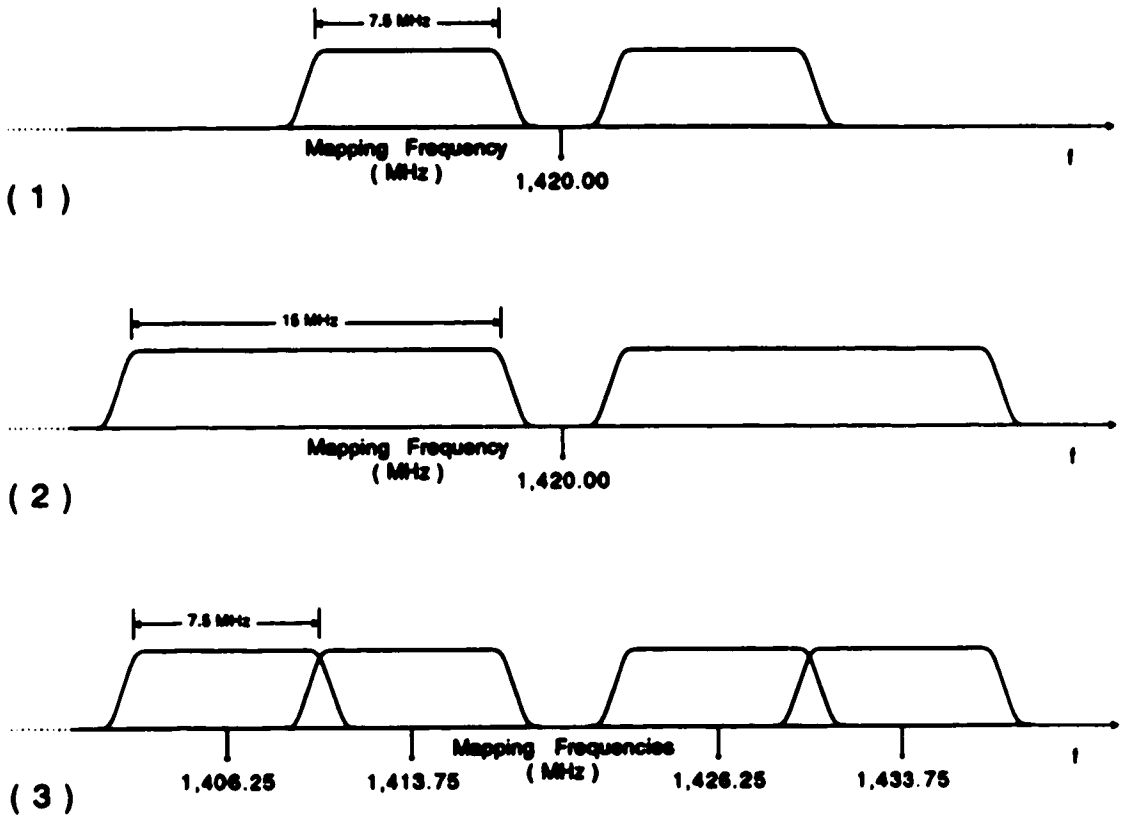

Figure 1.10: Correlation Frequency Options: The frequencies processed by the correlator determine the performance of the telescope: 1 shows the present strategy; 2 is similar but with double the bandwidth; 3 shows the wider bandwidth correlated in four sub-bands.

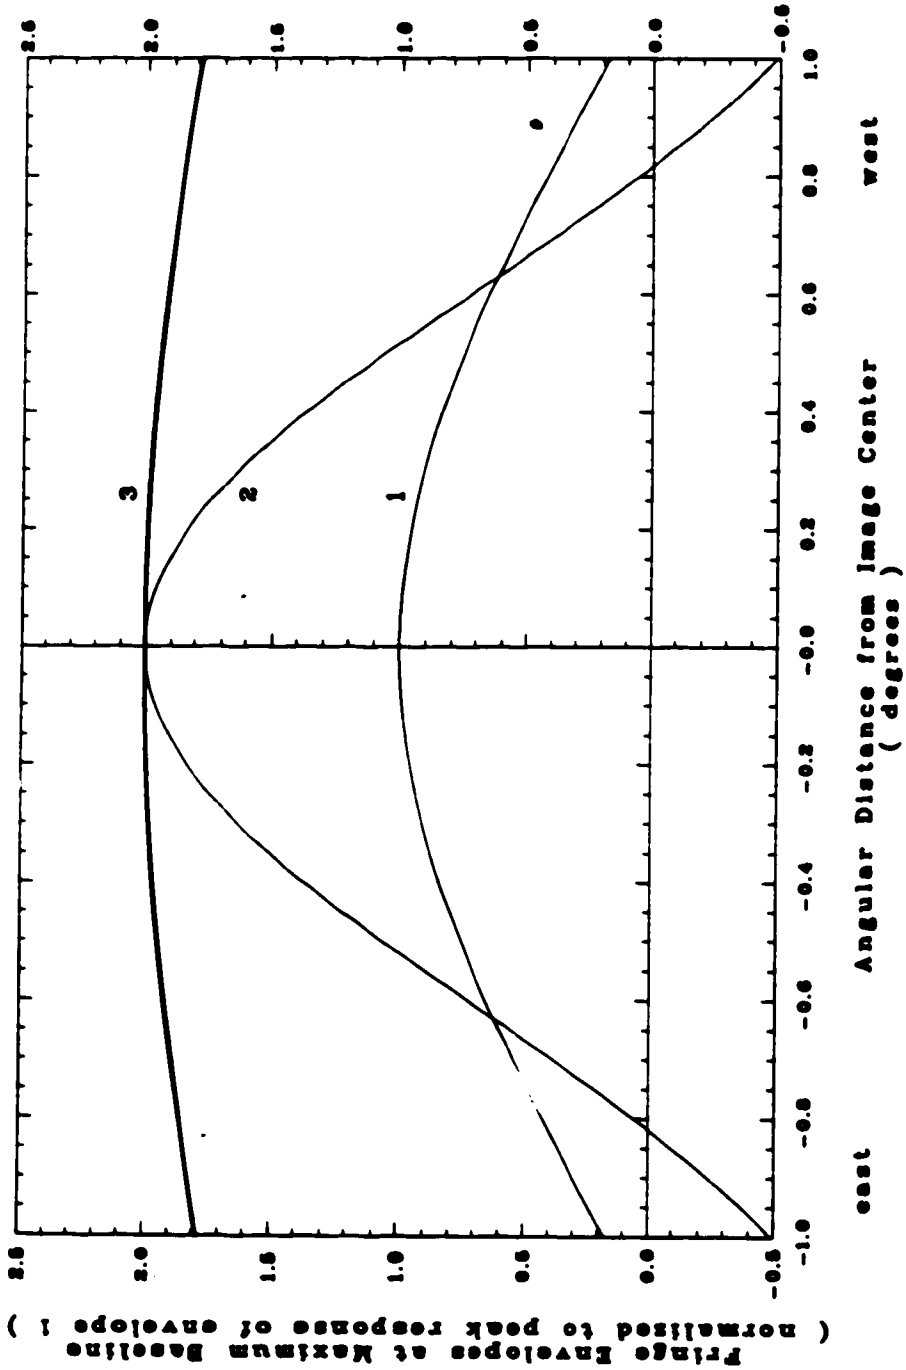

Figure 1.11: **Worst-Case Delay Beams:** The worst case occurs at the maximum baseline (600 metres). The three cases correspond to the correlation bandwidths shown in Figure 1.10.

converters capable of the 7.5 MHz bandwidths were being introduced. The required sample and processing speeds (nominally 20 MSPS) were just slow enough to provide a broad selection of standard digital signal processing components.

The other parameter in equation 1.21 which lies within the scope of this project is the correlator efficiency, η_c (maximum 1). For a digital correlator system, the value of η_c reflects the signal-to-noise penalty incurred by digitizing the antenna signals. For example, a polarity coincidence correlator (1 bit sampling) has an efficiency of $\frac{2}{\pi}$ (shown by Weinreb [Wei63]). Using techniques presented by Klingler and Bowers [KB74] the correlator efficiency for various quantization schemes can be compared [4]; two degradation curves are shown in Figure 1.12.

An ideal analog correlator has an efficiency of 1. The more resolution (bits) an A/D converter has, the more its efficiency resembles that of the ideal analog unit. To maximize efficiency, a 4-bit sampling scheme was chosen for this system. For reasons described later, the signal quantization employs 14 levels (out of 16 possible) to attain a correlator efficiency of $\eta_c = 0.985$. Figure 1.12 shows a comparison of the performance attainable with 2 types of quantization schemes: 2-bit (3 levels) and 4-bit (14 levels) quantization. The penalty paid for more bits is in the complexity of the digital hardware. No substantial improvement in performance is possible beyond 4-bit quantization.

In summary, the four bit/four band system was selected as a baseline for the design because of the following reasons:

- Very high correlation efficiency - virtually indistinguishable from an ideal analog correlator.
- Mature device technology resulting in a broad selection of economical (and available) components for the design.

[4] The correlation efficiency, η_c, is the reciprocal of the degradation factor 'D' described by Klingler and Bowers.

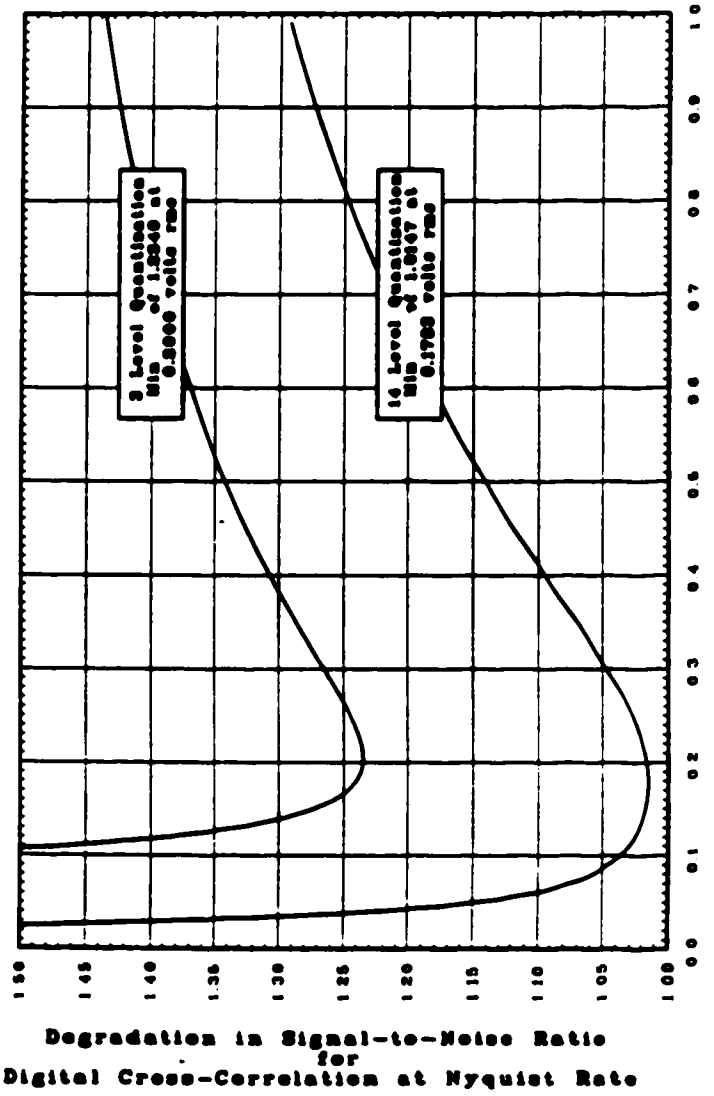

Figure 1.12: Correlator Efficiency: η_c, is the reciprocal of the degradation. The abscissa is the input signal standard deviation, σ: in terms of the A/D output weighting range [−13, +13], an analog input of 0.5 volts corresponds to +16. Optimum correlator performance is attained when the standard deviation of the digital samples is near 5.7.

1.4.3 The Design Baseline

This section will present the system design agreed upon at the outset of the project.

As discussed above, the continuum frequencies would be separated into four sub-bands, each 7.5 MHz wide (see Figure 4.1). This separation applies to both polarizations (left and right) of all eight antennas. Hence each antenna will produce eight independent sub-bands (4 right and 4 left). The bandwidth separations would be done at IF frequencies (12.5 - 47.5 MHz) using bandpass filters; for example a 12.5–20.0 MHz filter isolates the lowest band. Each 7.5 MHz band is then mixed down to baseband (actually, 0.5 - 8.0 MHz) prior to A/D sampling. The quadrature signal ($\frac{\pi}{2}$ phase shifted) is formed during this last mixing stage by splitting each band signal and using LO signals $\frac{\pi}{2}$ apart. This results in 16 separate signals for each antenna's continuum. Sixteen A/D converters per antenna (128 total) are required to sample the signals.

Further system design takes advantage of the symmetry of the design. The processing requirements for each of the four bandwidths are identical.

Designing and testing a correlator system for one band will accomplish the design requirement for the entire system. The correlator system for each band must meet the following requirements:

- All possible interferometer pairs must be formed within the band. This facilitates self-calibration.

  8 antennas form 28 pairs.

- An interferometer pair must measure all polarization visibilities (LL, RR, LR and RL) so that the telescope can measure polarization.

  4 visibilities require 8 correlator channels per pair.

- Correlator channels must multiply and accumulate 4-bit x 4-bit values at a nominal rate of 20 MSPS for intervals as long as 90 seconds. Channel results must be easily accessible to the computer controlling the system.

- For proper phase-switch demodulation, the system must start and stop operations in synchronism with the other observing systems (HI) sharing the antennas.
- The total 4-band system will have at least 896 correlator channels (auto correlation channels included, 1,152) providing a processing rate of 1.8 billion correlation operations per second.

The initial scope of this project included:

1. Development of the analog IF signal chain.
2. Development of an A/D converter unit.
3. Development of the digital correlator system.
4. Development of the software required to control the system and integrate it with the existing DRAO computer.

1.4.4 Project Progress Overview

The volume of work in the project baseline dictated a reduction in scope: since the DRAO had expertise in areas 1 and 4, the project proceeded on fronts 2 and 3.

The IF processing system, area 1, was constructed and tested at the unit-level by DRAO staff and is now ready for system-level testing.

Area 2, the A/D converter design, proceeded as far as preliminary testing of Version 3 prototypes. Further development and testing is being done at the DRAO. Chapter 4 describes the development and testing done on the A/D converter unit. The system is in the latter stages of hardware construction and has passed the required unit level tests.

Area 3, the development of the digital system, proceeded to the successful testing of a subsystem composed of a microprocessor controller and the three fundamental system components: the control/interface card (CIC); the

mother board and chassis; and correlator cards. Using these units, a complete continuum correlator system can be built.

Software development, area 4, will probably proceed in parallel with the development of the A/D converters. Considerable work was done in this area [5] to facilitate the development tests in area 2 and 3 (the 68000 microprocessor controlled all of the test configurations). A microprocessor-to-CIC interface (hardware and software) was designed for use during the testing in area 3; this design will probably meet the needs of the full system.

System integration testing will probably be nearing completion in 1991.

[5]Gary Hovey provided much valuable advice and assistance during these activities.

Chapter 2

The Correlator

This chapter presents the development of the digital correlator[1] beginning with the channel requirements, followed by the particular design choices made. The focus will move from microscopic to macroscopic concluding with the correlation system topology.

The correlator is the 'heart' of an interferometer; it performs the multiplication of signals to obtain the cross-power discussed in Chapter 1. In practice, the cross-power has a very large standard deviation relative to the mean. To achieve a meaningful measurement, the cross power must be averaged over several seconds. The accumulation implied by the average is done, at least in part, by the correlator.

In fact, signals from each antenna are dominated by terrestrial noise. Usually the correlated signal components (power) are several orders of magnitude smaller than signal components produced by amplifier noise and antenna spill-over (sources outside the main beam). The system noises produced by different antennas arise from independent, random processes and so contribute, on average, nothing to the integral of the cross-power. Man-made noise entering through the antenna side-lobes, on the other hand, can

[1]The term *correlator* will be used loosely to refer to a correlator system or to a single channel. One can infer which of the meanings (or both, as in this case) is meant from the context.

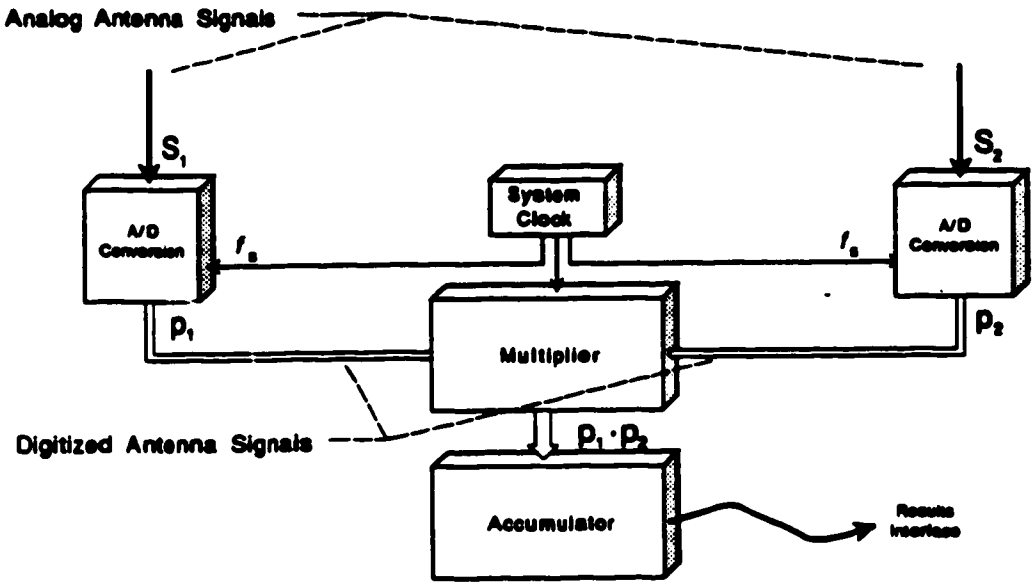

Figure 2.1: Digital Correlation: Antenna signals are digitized, multiplied and accumulated at the rate f_s.

produce a non-zero component in the cross-power.

Often the unwanted signal sources (usually man-made) are fixed with respect to the antennas, behaving like a source at the pole of the celestial sphere.

2.1 The Channel Design

Channel Requirements. Some of the features of a digital correlator are shown in Figure 2.1. Prior to correlation, the antenna signals S_1 and S_2 are sampled (converted to sequences of binary numbers); each sample represents the antenna signal amplitude at an instant in time. For each processing cycle, the digital correlator multiplies two values to attain the instantaneous cross-power. After many cross-power values have been saved in an accumulating buffer, the buffer contents (cross-energy) must be transferred to the computer. Some amount of channel control is required: first, the ability to start and stop correlation; second, the ability to reset the channel (old results must be flushed out). In this system, each correlator channel must measure the cross-power in a 7.5 MHz bandwidth. To make full use of the signal's power requires the sample rate to be at least 15 MSPS. Adding a 30% design margin results in a 20 MSPS processing rate; oversampling compensates somewhat for quantization noise and for short-lived events. This speed requirement is well suited to the capabilities of standard digital hardware. This is also true for data acquisition components; inexpensive A/D converters allowed a 4-bit sampling scheme to be used on all signals.

Correlator operations are synchronized to the processing clock. This is done using edge triggered D-type flip-flops (DFF's) and counters. Partitioning an algorithm into sequential tasks is called *pipelining*. Channel operations are pipelined to attain a higher processing rate. The four pipeline segments in a channel are:

1. The multiplier: U1 and U2
2. The least significant nibble (4 bits) accumulation: U3 and U4

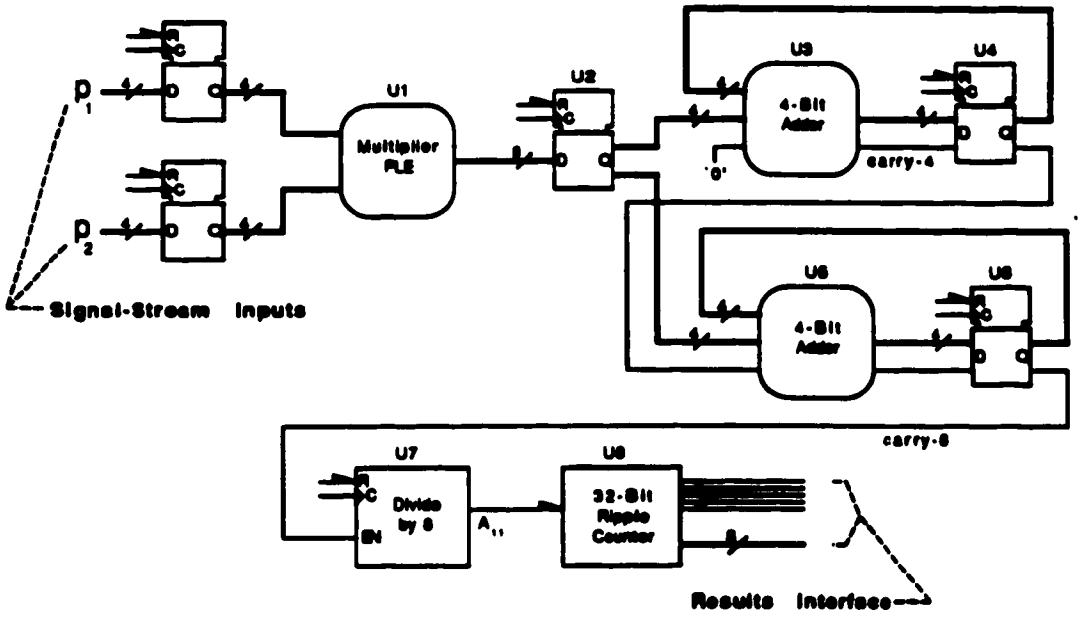

Figure 2.2: A Correlator Channel: Asynchronous channel components are interspersed amongst the synchronous D registers. In general, limiting the complexity of an asynchronous section reduces its propagation delay. The channel has four such sections to allow a faster processing rate.

3. The most significant nibble accumulation: U5 and U6
4. The carry-out counter: U7 and U8

The operation of each pipeline segment is discussed in turn.

2.1.1 The Channel Multiplier

Recent advances in memory technology have made it feasible to use memory devices in high speed logic systems. This channel design uses a 256 x 8-bit PROM as the multiplier. The two 4-bit signals compose an 8-bit address value; stored at the location specified by the address is the cross-power result. Memory devices embedded in logic circuits are referred to as programmable logic elements (PLE's). Compared to the cheapest design alternative (the use of MSI multiplier chips such as the Cray multiplier chips 557/558) the PLE is faster, (25 ns vs. 60 ns) cheaper ($11 vs. $50) and more efficient (1 device vs. 7 devices). A further advantage, which will be discussed at some length, is that the PLE can be tailored to fit the application better than the generic logic devices.

The Multiplication and Offset Algorithm

The algorithm stored in the PLE has channel control features as well as computation features: first we will concentrate on the computation algorithm.

Each of the two data streams entering the correlator channel is 4-bits wide allowing 16 unique values to appear at each of the multiplier inputs. Only 14 of the 16 are used for data values. The other two are reserved for control features (discussed a little later). The multiplier interprets the 14 data values as odd numbers from -13 to +13. The corresponding range of possible products is from -169 to +169. If the multiplier produced a result with a signed number format (eg. 2's complement), the subsequent channel logic would have to add *or subtract* values from the accumulation buffer. Introducing a numerical offset to the multiplier result simplifies the accumulating buffer (only additive operations are required). Consider first

the smallest possible offset which will ensure positive values. Biasing the multiplier output by 169 means the range of products is from 0 to 338. To uniquely represent all the numbers in this range requires 9-bits. Since the product of two odd numbers is also odd we see that using an odd valued offset will always produce an even valued result; the least significant bit is always a '0'. It is not necessary to carry this bit (referred to hereafter as the *phantom bit*) through any computations. The actual offset used is 333 giving a result range from 164 to 502. The multiplier PLE output represents the upper 8 bits of an unsigned binary value within the result range. The 8-bit result, $P_{1,2}$, is passed from the multiplier PLE to the channel accumulator at the rising clock edge. The PLE algorithm is:

$$P = \frac{(n_1 n_2 + 333)}{2} \tag{2.1}$$

where

$$n_1, n_2 \in [-13, -11, -9 \ldots + 7, +9, +11, +13] \tag{2.2}$$

2.1.2 The Channel Accumulator

The input to the channel accumulator is the byte-wide data path flowing from the multiplier PLE. Consider first the least significant 4 bits (nibble) in the path. These are seen by the A-side inputs of a 4-bit adder (U3); the adder's B-side sees the result from the previous cycle's addition. As this 4-bit accumulator overflows it produces 'carry-outs' which are fed to the 'carry-in' input of the adder for the most significant nibble. Notice in Figure 2.2 that the carry-out signals from both adders pass through a DFF; they are seen by the following pipeline segment during the next process cycle. This permits a shorter process cycle since the carry propagation time does not have to incur two adder delays. The rest of the accumulation is done simply by counting the overflow carries from the most significant adder (U5). The first three bits of carry counting are done in U7, a gated synchronous counter. The presence of a carry allows (gates) the counter to increment at the clock rising edge. This device is simply a rate divider to slow down the carry count rate for the last device in the channel. The last device, U8, is a 32-bit ripple counter whose value will eventually be read by the computer.

Each time the third counting bit of U7 makes a high-to-low transition, the ripple counter in U8 increments. Only the final 32 bits of the channel accumulator (the contents of U8) can be read by computer, so the rate dividing counter was kept as small (3 bits) as possible. The ripple-count device 'U8' was selected mainly for its ease of interface.

The channel's accumulation capacity can now be examined. The parallel input portion of the accumulator (adders, etc.) is 9 bits wide (including the phantom bit); the serial accumulation is comprised of a 3-bit rate divider together with the 32-bit ripple counter. In total, the accumulator has a capacity of 44 bits or 2^{44}. Assuming a continuous maximum result from the multiplier (say 512 or 2^9), the channel can record up to $2^{(44-9)} = 2^{35}$ (more than a billion) process cycles without overflowing. At a processing rate of 20 MSPS, this allows 30 minutes of processing which is more than adequate.

2.1.3 Channel Control

Figure 2.3 shows the control network associated with each channel. This section will describe channel control associated with typical operations including:

- starting and stopping correlation
- reading results
- preparing the channel for the next correlation interval

Starting and Stopping Correlation Recall from the discussion of the channel multiplier PLE, that of the 16 possible 4-bit input values, two were not used for data. One of these values, referred to as 'STOP' (1111), is used to stop the processing in the channel. The PLE is programmed to produce zero if 'STOP' appears at either 4-bit input. The zero value is very different from the range of values (164—502) that occur during normal operations. The appearance of a zero at the 4-bit adders stops further accumulation. One processing cycle after the appearance of the zero at the output of U2,

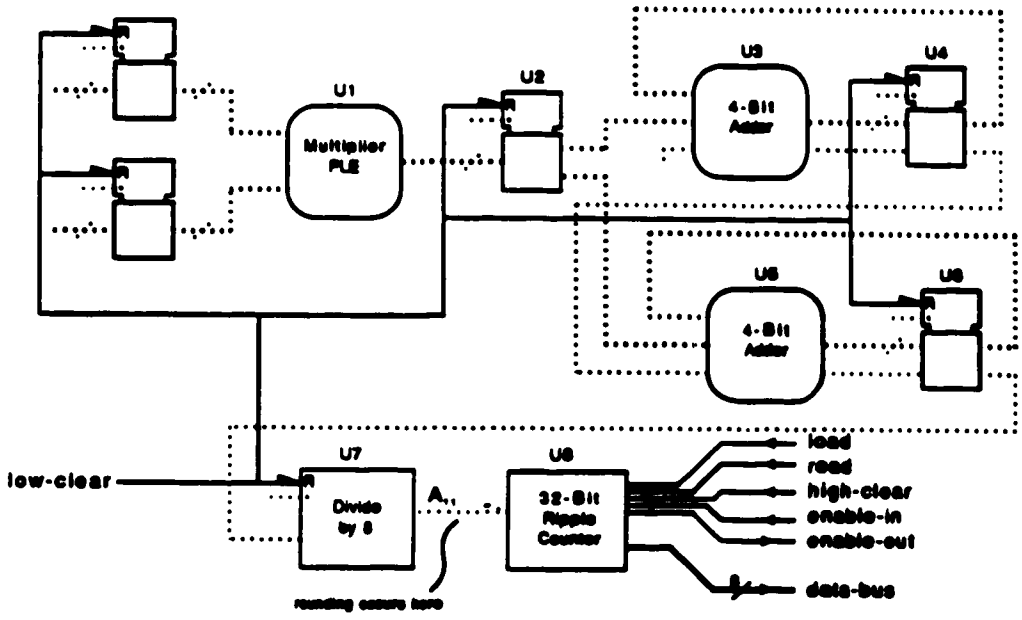

Figure 2.3: Channel Control Signals:

the final 'carry-out' (if any) appears at U7. The next process cycle sends the final count pulse (if any) to U8, the ripple-chip. From 'STOP' arriving at U1 to the last count pulse at U8 takes 3 clock cycles. It is important to recognize that the channel is still operating (process clocks are still running and all the channel devices are still active), it is merely processing zeroes. Hence, channel operation starts and stops depending on the values arriving at the channel multiplier. There are some practical advantages to multiplexing control conditions onto the data distribution network feeding the channel. Because these advantages are system-level considerations, they are best treated later in the chapter.

Reading Results: At the end of a correlation run the processing stops as described above. The contents of the ripple counter need to be read by a computer. Within this device is a 4-byte storage register for saving the 32 bit result. Saving is done by activating the 'load' signal on the device. The result bytes are now ready to be placed on the data bus. Internal and external tri-state drivers are activated by two signals: first an active 'enable-in' condition must be seen; second an active 'read' condition places the value of the first (least significant) byte on the data bus. Successive bytes are displayed as the 'read' line toggles. When 'read' goes inactive for the fourth time, the 'enable-out' signal turns on and further activity on the 'read' line will have no effect. This interface strategy allows an unlimited number of counter devices to share the same bus and control lines simply by daisy-chaining the 'enable' lines from one device to the next.

Clearing the Channel Prior to the start of a new correlation run, the channel results from the previous run must be removed. This is done by exercising the two reset lines 'low-reset' and 'high-reset'. The 'low reset' line clears the 9-bit parallel accumulator as well as the 3-bit rate divider. The 'high-reset' line clears the 32-bit ripple counter. (Notice that the read-out register contents are not affected; the channel can begin another correlation run before the results have been read.) The reason for a separate 'low-reset' line is for rounding the accumulated result. Recall that the lower 12 bits (A_0—A_{11}) of the accumulator are not readable. If they are more than half full, A_{11} is a '1'; in this case we would like to add a count to the upper 32-bit

result so the visible value represents the *rounded* channel result (as opposed to the *truncated* result). By sending the 'low-clear' signal, these bits are set to '0's which generates the count event when needed (high-to-low transition on signal A_{11}). Results returned to the computer are free of bias produced by truncation errors.

2.1.4 Correlator Channel Summary

The correlator channel is capable of operating at a 20 MHz processing clock rate. For each clock cycle the channel will correlate two 4-bit inputs. The channel employs a PLE (look-up table) multiplier and forms all the possible data products. Channel operation (On/Off) is controlled through the data inputs using reserved codes. The accumulator has 44 bits of capacity, with the upper 32 accessible as the correlation result. The 32-bit result can be rounded prior to readout. Because each channel has embedded storage registers, the channel can be cleared and begin the next correlation run before results have been read. This avoids reduction of the observing time by the need for the computer to gather results.

Channel Performance

At this point it is appropriate to give some attention to the characteristics of the channel output. From the statistical nature of the two input signals, we can determine the precision of the channel output. The correlator operates optimally when the standard deviation, σ, of the input signal is 5.71 (on the -13 to +13 scale) [2] . The standard deviation of the product of two such signals increases by $\sqrt{2}$ to +8.07. Averaging products over N independent samples reduces the standard deviation by $\sqrt{N}$. In a typical correlation run (5.6 seconds) roughly 90 million ($5.6 \times 15 \times 10^6$) independent[3] samples will

[2]This value of σ, discussed further in Chapter 4, makes optimum use of the 14 quantiser levels.

[3]For purposes of these calculations, the Nyquist rate, 15 MSPS, is used.

have been taken. The expected channel output is:

$$\langle C \rangle = \frac{(O_A + \langle p_{1,2} \rangle) \times \tau \times f_S}{d_L} \tag{2.3}$$

$\langle C \rangle$ is the expected channel output value.

O_A is the accumulator offset (333).

$\langle p_{1,2} \rangle$ is the expected signal product (+8.07 if $\rho = +1$ or −8.07 if $\rho = -1$).

d_L is a divider which accounts for the 12 invisible channel bits.

τ is correlation-run duration (5.6 seconds).

f_S is sampling rate for each antenna signal (15 MSPS).

The extremes at the correlator output are determined by the perfect *correlation* ($\rho = +1$) and *anti-correlation* ($\rho = -1$) cases:

$$\langle C \rangle_{min} = 6.67 \times 10^6 \tag{2.4}$$

$$\langle C \rangle_{max} = 7.00 \times 10^6 \tag{2.5}$$

The standard deviation at the channel output is related to these values:

$$\sigma_C = \frac{\langle p_{1,2} \rangle \sqrt{\tau \times f_S}}{d_L} = 18.06 \tag{2.6}$$

In general, the uncertainty in the result will be confined to a magnitude within the lower 5 bits of the channel output. The correlation output range will be confined to 19 bits, 14 of which are potentially significant depending on the correlation coefficient ρ. Based on equations 2.3 and 2.6 the detection threshhold is $\rho = 0.011\%$ in 5.6 seconds with the result resolution exceeding that by 4 bits. The digital result is well represented with the 32 output bits.

Contrast the digital correlation results above with the old continuum correlators where a channel's analog output is sampled with a 10-bit A/D converter. Channel sensitivity is limited by the output resolution; some

improvement is had by increasing the channel gain. This is based on the reasonable assumption that the signal correlation will generally be small ($\rho < 5\%$). Nevertheless, correlator effects are seen when an image field has a large dynamic range. The coarse 10-bit result range adds quantization noise to the observations.

The new digital correlator processing strategy will improve both the sensitivity and the dynamic range of the instrument.

2.2 The Correlator Card

With the correlator channel design established, 'larger' aspects of the system can now be addressed. Channels must be *packaged* (mounting, cooling) and *serviced* (electrical connections) in efficient groups. The number of channels required to form an interferometer pair (8) provides a natural design partition for the system; the eight channels, grouped on a printed circuit board, share timing, interface and signal distribution circuitry. This grouping simplifies the traffic between correlation blocks in the system. Interconnecting these boards (correlator cards) forms the processing network required by the antenna array.

2.2.1 The Polarized Interferometer

Figure 2.4 shows the in-phase and quadrature channel correlations associated with the complex interferometer. To discuss details of the correlator board operation, the following notation is needed:

p is an in-phase (unaltered) signal.

q is a quadrature-phase signal ($\frac{\pi}{2}$ phase-shift).

$< \ldots >$ Poisson brackets denote time average (expectation value).

j, k subscripts denote antenna number.

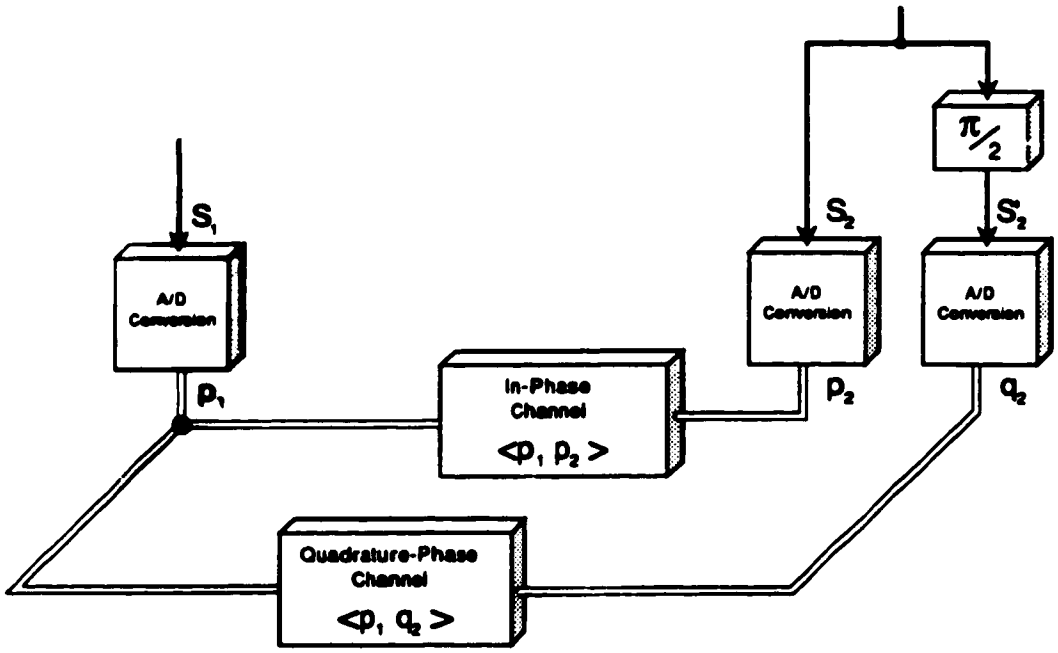

Figure 2.4: **The Complex-Value Correlator:** The channels measure the cross-power vector (a complex number) between antenna signals.

R,L superscripts denote right-hand and left-hand circular polarization.

In this notation, the complex correlator output has the form:

$$\hat{P}^{RL}_{j,k} = \underbrace{\left\langle p^R_j p^L_k \right\rangle}_{1} + \underbrace{j \left\langle p^R_j q^L_k \right\rangle}_{2} \tag{2.7}$$

Figure 2.4 shows this notation. Each expectation value is derived from a correlator channel output, hence, Poisson brackets are synonymous with correlation. The measurement of $\hat{P}_{j,k}$ requires two correlator channels: term 1 is the in-phase channel; term 2 is the quadrature channel.

The measurement of the polarization parameters of the sky brightness requires the following cross-polarization quantities to be measured:

$$\hat{P}^{R,R}_{j,k} = \left\langle p^R_j p^R_k \right\rangle + j \left\langle p^R_j q^R_k \right\rangle$$

$$\begin{aligned}
\hat{P}_{j,k}^{R,L} &= \left\langle p_j^R p_k^L \right\rangle + j \left\langle p_j^R q_k^L \right\rangle \\
\hat{P}_{j,k}^{L,R} &= \left\langle p_j^L p_k^R \right\rangle + j \left\langle p_j^L q_k^R \right\rangle \\
\hat{P}_{j,k}^{L,L} &= \left\langle p_j^L p_k^L \right\rangle + j \left\langle p_j^L q_k^L \right\rangle
\end{aligned} \tag{2.8}$$

For unpolarized or linearly polarized radiation, $\hat{P}_{j,k}^{R,R}$ and $\hat{P}_{j,k}^{L,L}$ are independent measurements of the same quantity; a $\sqrt{2}$ improvement in sensitivity results. Notice that for the eight correlations, six signal streams are required:

From Antenna k: p_k^R, p_k^L, q_k^R, q_k^L

From Antenna j: p_j^R, p_j^L

The four correlation vectors can be used to derive the four Stokes parameters I, Q, U and V: I represents the total incident power, Q and U the power in the linearly polarized component and V the power in the circularly polarized component of radiation.

The relationships between correlations and Stokes parameters (see Thompson et al. [TMS86, page 105]) are, in their simplest form:

$$\begin{aligned}
I &= \frac{1}{2}\left(\hat{P}_{j,k}^{R,R} + P_{j,k}^{L,L}\right) \\
Q &= -j\frac{1}{2}\left(\hat{P}_{j,k}^{R,L} + P_{j,k}^{L,R}\right) \\
U &= \frac{1}{2}\left(\hat{P}_{j,k}^{R,L} - P_{j,k}^{L,R}\right) \\
V &= \frac{1}{2}\left(\hat{P}_{j,k}^{R,R} - P_{j,k}^{L,L}\right)
\end{aligned} \tag{2.9}$$

Stokes parameters (having dimensions of flux density) propagate through space in the same manner as the electromagnetic field, therefore the response pattern of the interferometer applies to these quantities as well. In other words, providing the eight correlator channels shown above will allow mapping (images) of each Stokes parameter.

2.2.2 The Physical Design of the Correlator Card

The eight digital correlator channels are packaged together on the circuit board shown in Figure 2.5. The six 4-bit data streams are labelled '*A*' through '*F*'. The eight channels form the following correlation products:

Channel 1	$\longrightarrow$	$\langle AC \rangle$	$\longrightarrow$	[1]
Channel 2	$\longrightarrow$	$\langle AD \rangle$	$\longrightarrow$	[2]
Channel 3	$\longrightarrow$	$\langle AE \rangle$	$\longrightarrow$	[3]
Channel 4	$\longrightarrow$	$\langle AF \rangle$	$\longrightarrow$	[4]
Channel 5	$\longrightarrow$	$\langle BC \rangle$	$\longrightarrow$	[5]
Channel 6	$\longrightarrow$	$\langle BD \rangle$	$\longrightarrow$	[6]
Channel 7	$\longrightarrow$	$\langle BE \rangle$	$\longrightarrow$	[7]
Channel 8	$\longrightarrow$	$\langle BF \rangle$	$\longrightarrow$	[8]

If streams A and B are used for p_j^R and [illegible] streams C, D, E and F are used for p_k^R, q_k^R, p_k^L and q_k^L respectively, the complex interferometer response vectors from Equation 2.8 (using the loose notation shown above) become:

$$\begin{aligned} \hat{P}_{j,k}^{R,R} &= [1] + j\,[2] \\ \hat{P}_{j,k}^{R,L} &= [3] + j\,[4] \\ \hat{P}_{j,k}^{L,R} &= [5] + j\,[6] \\ \hat{P}_{j,k}^{L,L} &= [7] + j\,[8] \end{aligned} \tag{2.10}$$

The correlator card provides the necessary input, output and control signals for each channel.

Power Distribution: The obvious, though often badly treated, connections on a circuit board are those for power and ground. The digital correlator card design minimizes power system noise and ground-loop problems by using a 4-layer circuit board (one layer is dedicated to power, another to ground), as well as filter capacitors at each device.

Clock Distribution: The signal which is most important to the correct operation of a fast digital system is the system clock. Care was taken to

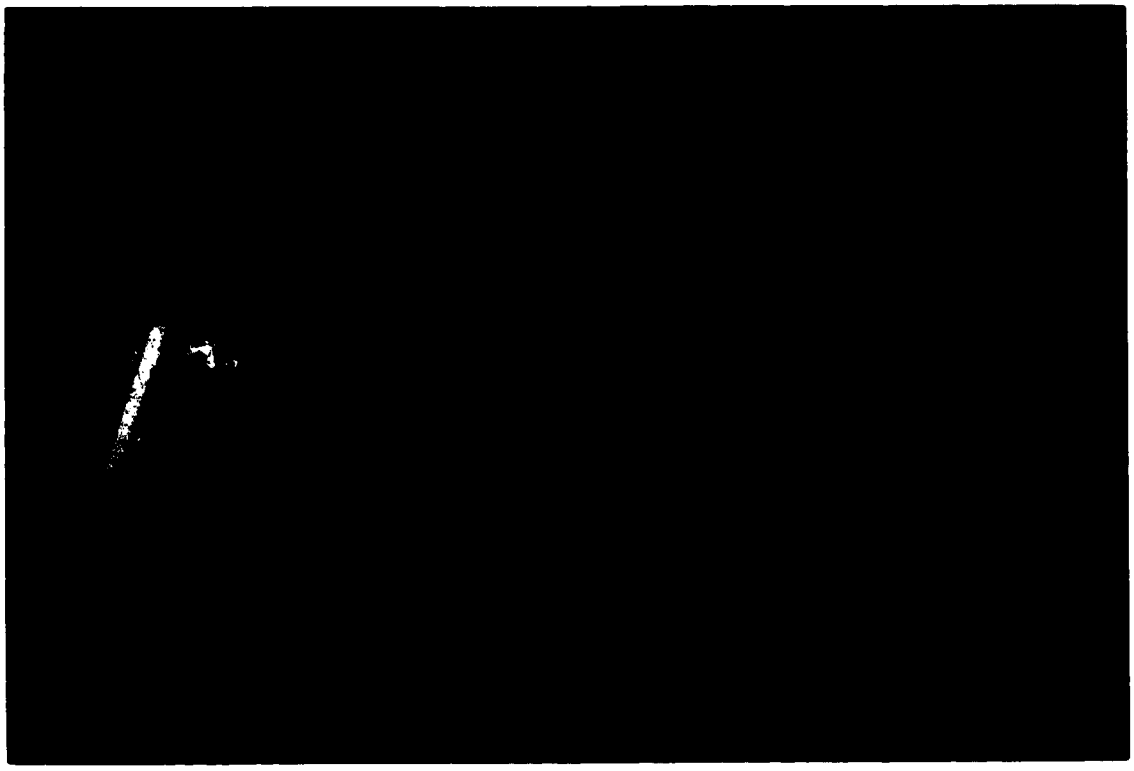

Figure 2.5: The Digital Correlator Card: provides all the signal correlations needed by an interferometer for polarization measurement. The 4-layer circuit boards operate at 20 MHz clock rates.

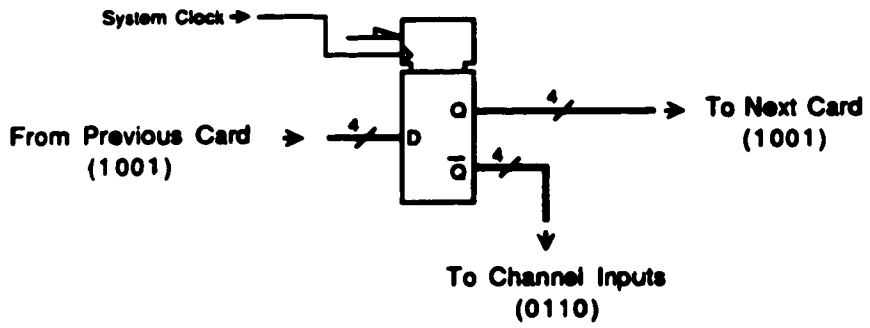

Figure 2.6: Correlator Input Design: At the rising edge of the clock, the value at the input of the D flip-flop is latched and transferred to the inverting (seen by the local correlator channels) and non-inverting (sent to the next correlator card) outputs.

ensure that proper clock distribution was achieved:

- Whenever possible, signal traces stay on one layer.
- Branches on high-speed signal traces were not allowed.
- Signal fan-out never exceeds 8.
- Clocks are driven from a single distribution device to minimize timing skew.

Similar precautions are taken for the other signals on the correlator card, with the exception of the reset control signals which are relatively slow (the signal branching rule was relaxed).

Antenna Signal Distribution: Distributing the digital signal data to the hundreds of channels is an important prerequisite for correlation. An effective strategy is to pass data values in bucket-brigade fashion from one correlator card to the next. To clock values from card to card, each data-stream requires 4 inputs and 4 outputs. Each stream input is connected to a dual-output D-type flip-flop as shown in Figure 2.6. These devices, which synchronize

the data transfer between cards, also distribute the values to the channels on the board. Notice that each data bit gets inverted before it appears at the channel; data-passing between cards uses negative logic (high = 0, low = 1). Inversion to positive logic occurs as the bit values are presented to the channels. A practical reason for this is to avoid power supply surges during changing data conditions; since the DFF is always driving one line high and the other low the device current consumption is relatively constant. Another reason was to simplify the PCB layout by making efficient use of fewer devices.

Figure 2.7 demonstrates the data distribution design by showing several values being passed from card to card:

Because correlation On/Off control is embedded in the data passing network, correlator cards start and stop in domino fashion. Notice that every correlator card on a given stream will see the same data sequence; only the arrival times differ.

Results Interface: After a correlation run, the eight channel results are read by the controlling computer. The data interface at the board level is a direct extension of the single channel design. All the control signals to the channels are connected in parallel with the exception of the 'enable' signals which are connected in series. The read-out protocol is identical to that described in the previous section. There are two support devices associated with board control and interface. One buffers control signals coming to the board. The other is a 'smart' byte-wide tri-state device controlling the correlator card's access to the external data-bus (this is a PAL[4] device). Its tri-state drivers are active only if 'enable-in' is true and 'enable-out' is false. This condition occurs while data is being read from the channels on the board. When all channels have been read, 'enable-out' becomes true and the tri-state drivers relinquish the bus.

[4]programmable array logic

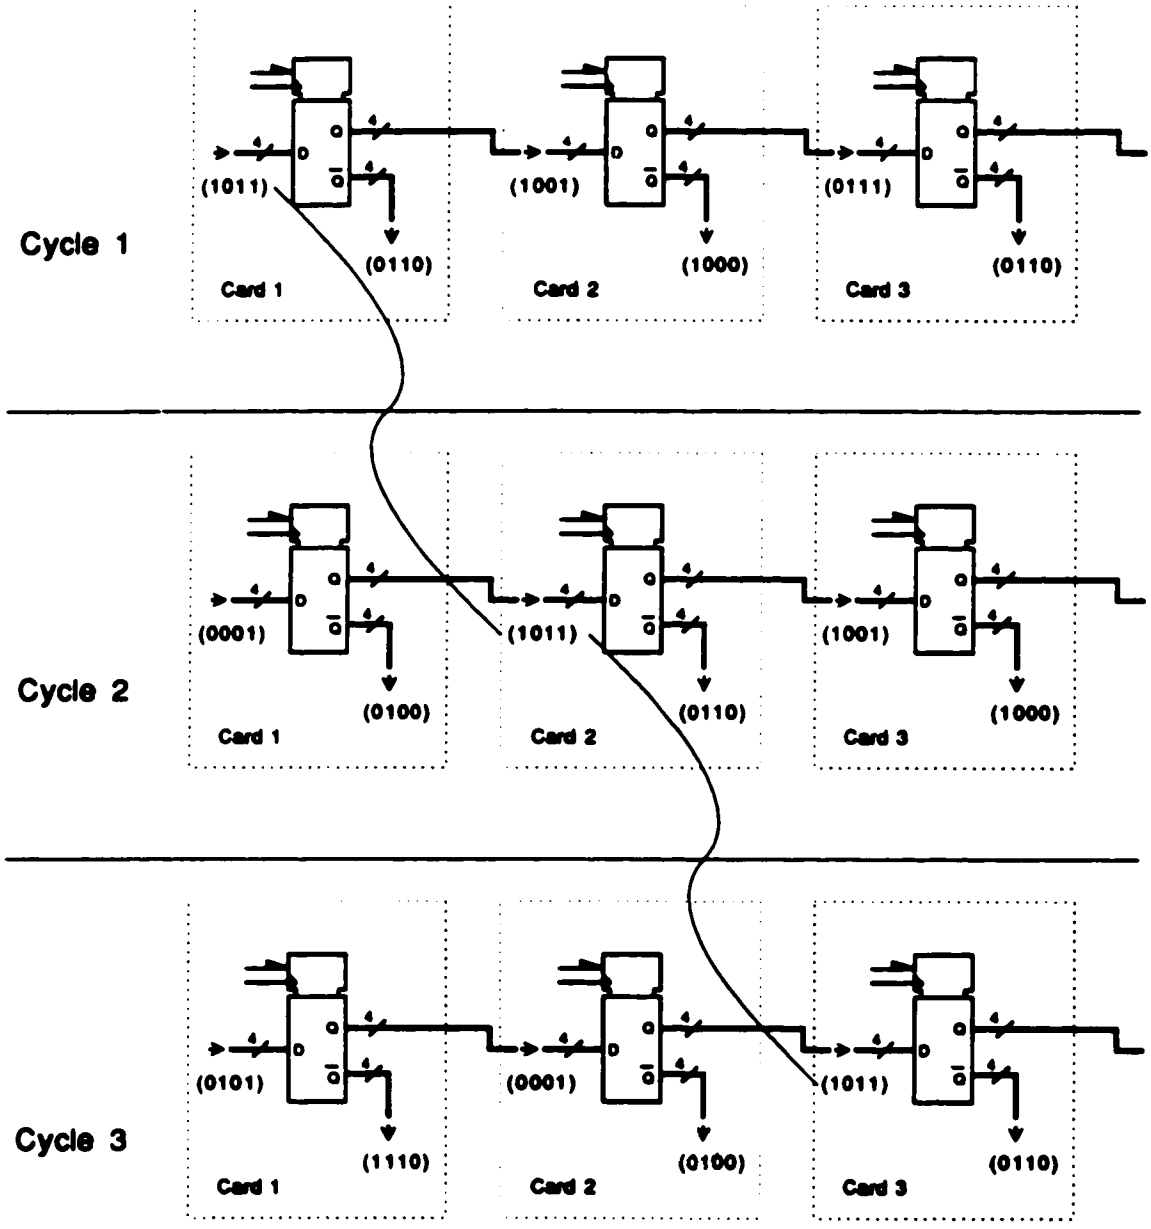

Figure 2.7: **Data Passing Design:** Data-stream values are passed from left to right on successive clock cycles. The bit-inverted values are sent to the correlator channels.

2.2.3 Correlator Card Summary

Each correlator card provides the eight channels required for polarization measurement by an interferometer pair. The six necessary 4-bit data streams are latched on/off the board in synchronism with the processing clock. The four layer printed circuit board provides a reliable power and signal distribution network. All external connections are made at a 96-pin DIN connector; the pin allocation is the following:

timing:	2 inputs	
data streams:	24 inputs	24 outputs
interface bus:		8 outputs
interface control:	5 inputs	1 output
power and ground:	30 inputs	

While the nominal system clock rate is 20 MHz, during tests the correlator board functioned correctly at speeds up to 30 MHz; worst-case limits on component specifications indicate that 25 MHz is the limit of safe operation. The component limiting correlator speed is the multiplier PLE.

A block diagram of the correlator card will be useful during the next section. Figure 2.8 shows all the interconnections needed at the system level.

2.3 Correlation System Topology

We are now able to describe how the correlator cards fit together to meet the telescope's continuum signal processing requirements. This section will also introduce the autocorrelation features built into the correlator card. Autocorrelation can improve the quality of the cross-correlator results.

2.3.1 The Correlation Matrix

Using the correlator card building blocks shown in Figure 2.8 the necessary correlation system can be assembled. As mentioned in the introductory chap-

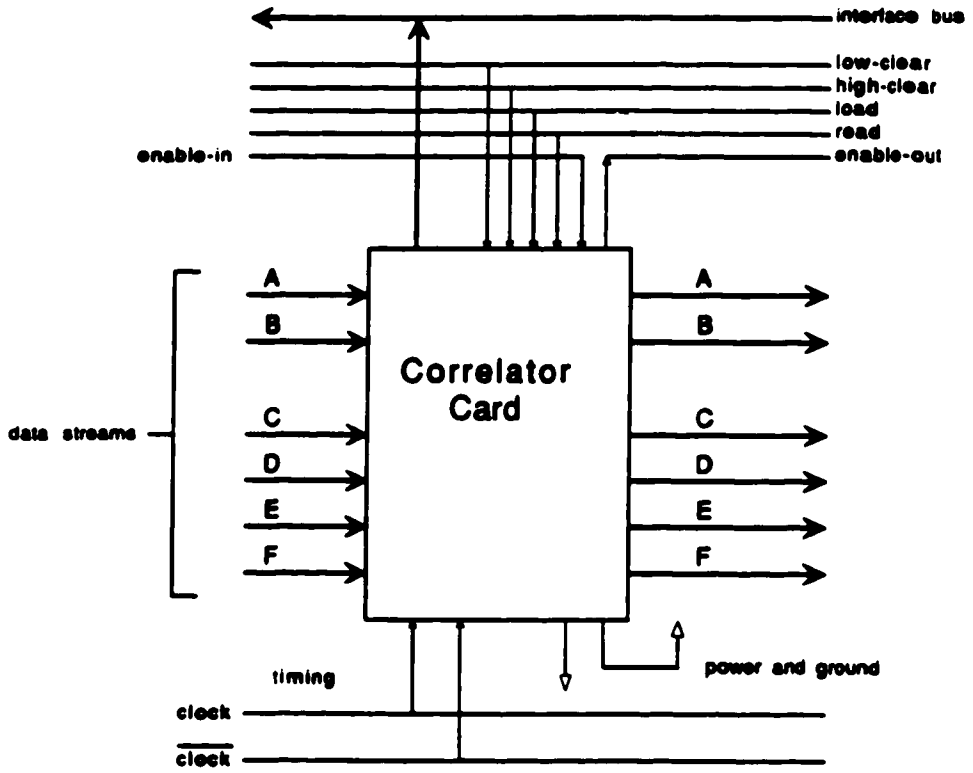

Figure 2.8: **Correlator Card I/O Block Diagram:**

ter, all possible interferometer pairs are to be formed. In other words, data streams from each antenna must be correlated with streams from every other antenna in the array. Figure 2.9 depicts the interconnection pattern used to achieve this.

Observe that for antenna $k \in [5,6,7\ldots]$ all four of the data streams (C_k, D_k, E_k and F_k) are passed horizontally. These are correlated with signal streams A_j and B_j from antenna j ($j < k$). Recall that data values in each stream are passed from card to card (single-source/single-destination). The most important aspect of this data distribution design is the absence of signal fan-out (fan-out = 1). As a result, the data distribution network can be expanded indefinitely.

The system-level (chassis) interface network is a direct extension of the protocol used at the card level. Each row of the matrix shares the results interface: 6 control signals; a byte-wide data bus. The design of the timing network allows adjustment for achieving synchronization between the chassis-rows.

2.3.2 Autocorrelation

The correlator card interconnection pattern suggests the addition of C_{jj} correlation cards. Figure 2.10 shows these blocks in place. Since such units would perform correlation processing on signals from only one antenna, they are referred to as autocorrelation cards. The autocorrelation features will be described next.

Removal of Gain and Offset Errors

An autocorrelation card is used to analyze the four antenna signal streams. These measurements monitor characteristics of the system upstream from the autocorrelators. To some degree, these measurements can be used to correct minor imperfections in the system.

During conversion from analog to digital, a signal will experience some

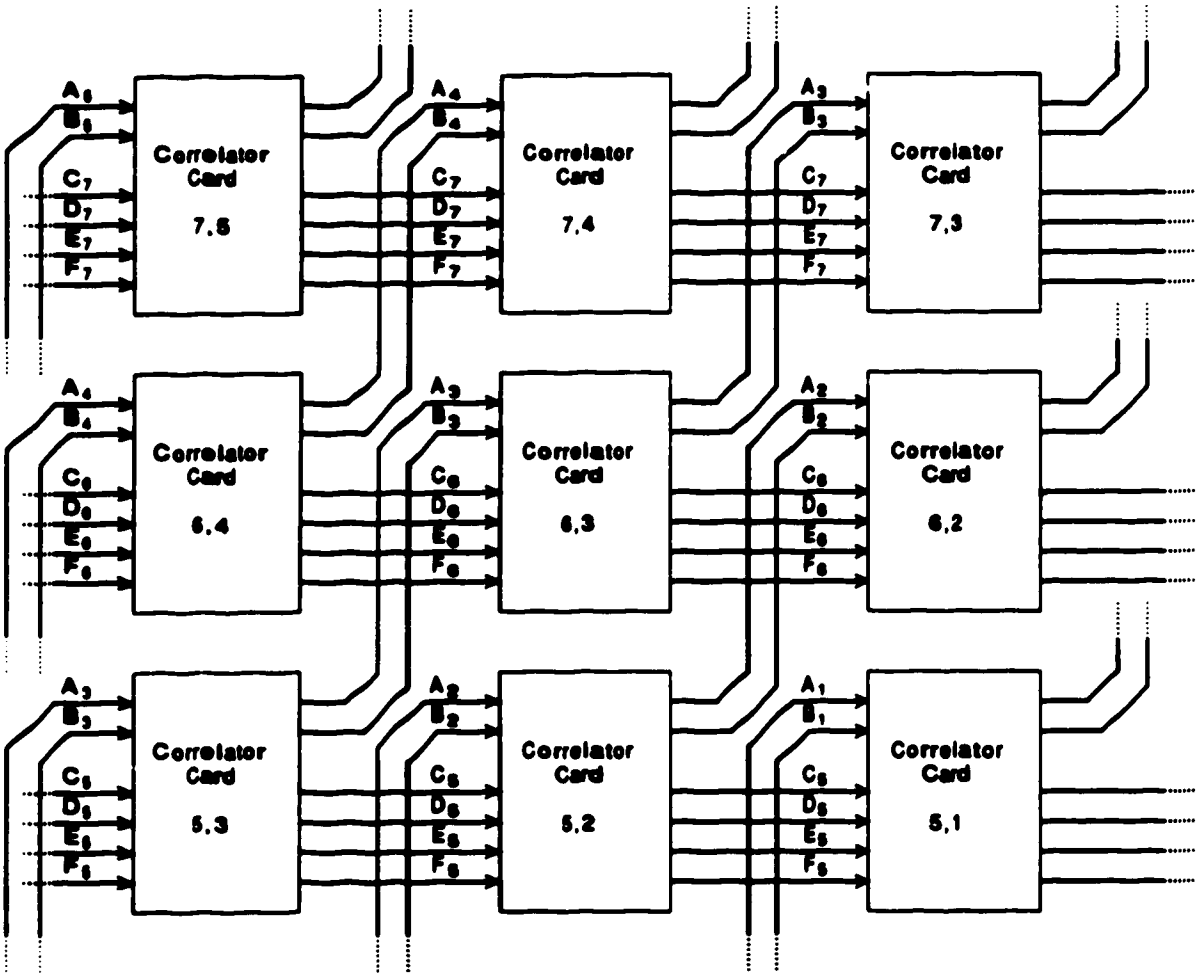

Figure 2.9: **Correlator Interconnections:** Diagram (a) shows the routing of the 4-bit signal streams within the correlator matrix. Streams A_j and B_j flow upwards diagonally; Streams C_k, D_k, E_k and F_k flow horizontally from left to right.

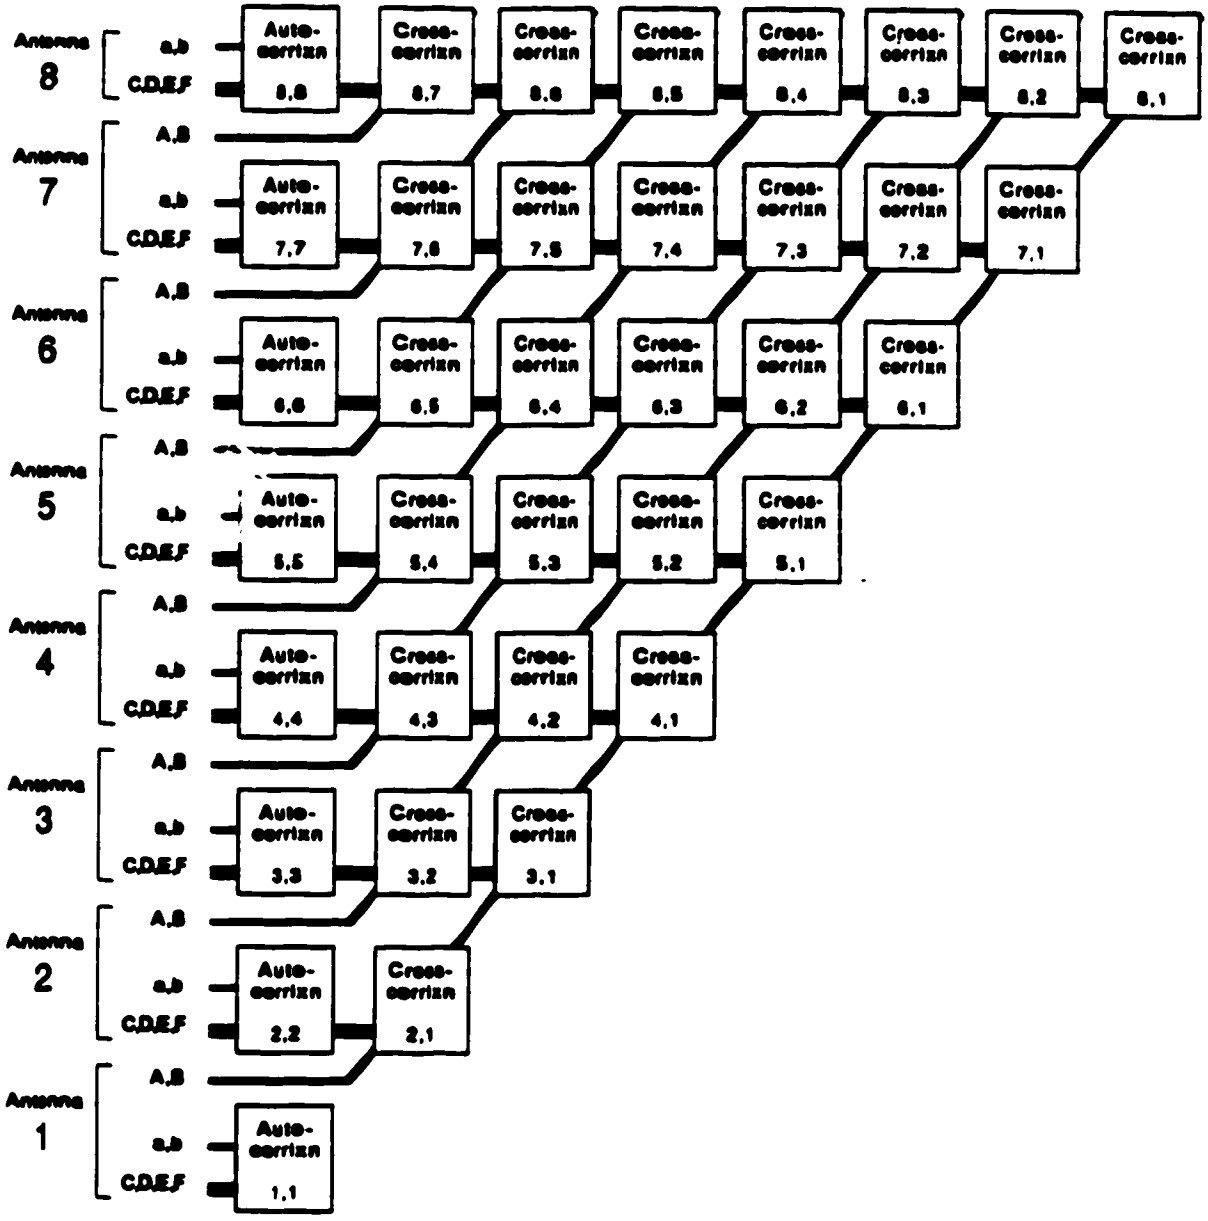

Figure 2.10: Autocorrelation Points in the Matrix: The autocorrelation cards, C_{jj}, require additional streams a_j and b_j used to control the type of autocorrelation processing done for each antenna.

corruption. While quantization noise is inherent in the A/D transformation, additional corruption is a result of imperfect quantization thresholds (see Chapter 4).

One manifestation of this is that the digital output will be *offset* — for example, the probability of observing a positive number will be different from that for a negative number. This will alter the stream expectation value $< p_j^R >$ (the mean).

A second effect is that produced by any changes in the *average* spacing between the quantization thresholds. This is equivalent to changing the gain in the analog path. The effect observable at the quantizer output is some change in the stream power $< (p_j^R)^2 >$ (variance).

We can measure the mean and variance of each of the signal streams using the autocorrelation card. If, during a correlation run, the a_j input is a constant (say +13), channels 1 through 4 will compute the following:

$$\text{Channel } 1_j: \quad C_{jj}^1 = (2^{-12}) \sum^N \left(13 \cdot p_j^L + 333\right) \tag{2.11}$$

$$\text{Channel } 2_j: \quad C_{jj}^2 = (2^{-12}) \sum^N \left(13 \cdot q_j^L + 333\right) \tag{2.12}$$

$$\text{Channel } 3_j: \quad C_{jj}^3 = (2^{-12}) \sum^N \left(13 \cdot p_j^R + 333\right) \tag{2.13}$$

$$\text{Channel } 4_j: \quad C_{jj}^4 = (2^{-12}) \sum^N \left(13 \cdot q_j^R + 333\right) \tag{2.14}$$

The factor 2^{-12} is a result of the 12 hidden low-order channel bits. The expectation value of any signal stream is easily obtained. For example:

$$13 \cdot \left\langle p_j^L \right\rangle = \frac{C_{jj}^1 \cdot 2^{12}}{N} - 333 \tag{2.15}$$

These quantities are offset-errors which skew the correlator results; for example, the result from channel $\left\langle p_j^L \cdot q_k^R \right\rangle$ will be skewed by the amount $\left\langle p_j^L \right\rangle \cdot \left\langle q_k^R \right\rangle$ This skew should be removed from the results (more on this later).

To measure the standard deviation of the signal streams we make use of the control value 'AUTO' (0000)[5] . When a channel sees 'AUTO' at one

[5]This is the second of the two data stream codes reserved for commands: STOP is 1111; AUTO is 0000.

of its inputs, the multiplier PLE result is the square (plus 333) of the data value appearing at the other input. During a correlation run, holding input b_j constant as 'AUTO' produces the following results at channels 5 through 8:

$$\text{Channel } 5_j: \quad C_{jj}^5 = (2^{-12}) \sum^N \left(\left(p_j^L \right)^2 + 333 \right) \tag{2.16}$$

$$\text{Channel } 6_j: \quad C_{jj}^6 = (2^{-12}) \sum^N \left(\left(q_j^L \right)^2 + 333 \right) \tag{2.17}$$

$$\text{Channel } 7_j: \quad C_{jj}^7 = (2^{-12}) \sum^N \left(\left(p_j^R \right)^2 + 333 \right) \tag{2.18}$$

$$\text{Channel } 8_j: \quad C_{jj}^8 = (2^{-12}) \sum^N \left(\left(q_j^R \right)^2 + 333 \right) \tag{2.19}$$

The standard deviation for any signal, say p_j^L, is computed as follows:

$$\left\langle \left(p_j^L \right)^2 \right\rangle = \frac{C_{jj}^5 \cdot 2^{12}}{N} - 333 \tag{2.20}$$

$$\sigma \left(p_j^L \right) = \sqrt{\left\langle \left(p_j^L \right)^2 \right\rangle - \left\langle p_j^L \right\rangle^2} \tag{2.21}$$

Assuming the antenna signals are Gaussian, the optimal[6] value for any signal's standard deviation, such as $\sigma \left(p_j^L \right)$, is near 5.71. Any deviations from this nominal value can be treated as gain errors, in this case $G \left(p_j^L \right)$.

$$G \left(p_j^L \right) = \frac{\sigma \left(p_j^L \right)}{5.71} \tag{2.22}$$

Having measured the offset and gain errors[7] for each signal stream, the cross-power vector values from equation 2.8, for example $P_{j,k}^{L,R}$, can be corrected

[6]This value was discussed in Chapter 1. If the digitally computed standard deviations are close to 5.71, the correlation results will have the highest signal-to-noise ratio.

[7]The offset errors and the gain errors are the low-order effects of (i) imperfect quantisation and (ii) imperfections in the prequantization analog signal processing; while there are also high-order error effects, it is impractical to correct for them (so they are tolerated).

as follows (square brackets denote the correlator channel outputs):

$$\hat{P}_{j,k}^{L,R} = \frac{\frac{2^{12}}{N} \cdot \left[2^{-12} \cdot \sum^{N} \left(p_j^L \cdot p_k^R + 333\right)\right] - 333 - \left\langle p_j^L \right\rangle \cdot \left\langle p_k^R \right\rangle}{G\left(p_j^L\right) \cdot G\left(p_k^R\right)} + j \cdot \frac{\frac{2^{12}}{N} \cdot \left[2^{-12} \cdot \sum^{N} \left(p_j^L \cdot q_k^R + 333\right)\right] - 333 - \left\langle p_j^L \right\rangle \cdot \left\langle q_k^R \right\rangle}{G\left(p_j^L\right) \cdot G\left(q_k^R\right)} \tag{2.23}$$

Here $\hat{P}_{j,k}^{L,R}$ has been corrected for gain and offset errors. Its value is a complex number whose modulus is less than $(8.07)\sqrt{2}$. We can define a complex correlation coefficient, ρ :

$$\hat{\rho}_{j,k}^{R,L} = \frac{\hat{P}_{j,k}^{L,R}}{(8.07)\sqrt{2}} \tag{2.24}$$

The modulus, $\|\hat{\rho}\|$ indicates the degree of similarity between a pair of signals: $\|\hat{\rho}\| = 1$ for exact similarity; $\|\hat{\rho}\| = 0$ for no similarity. The phase angle $\mathrm{Arg}(\hat{\rho})$ indicates the relative phase angle between the similar components of the two signals. In terms of the field of view being mapped, the $\hat{\rho}$ values are proportional to the magnitude and phase of the two-dimensional Fourier components of the sky brightness.

Removal of Quadraturity[8] Errors

In the previous discussion the a_j and b_j inputs of the autocorrelation card were constants (+13 (0111) and 'AUTO' (0000) respectively). We will now examine the information gained by feeding the a_j, b_j inputs with the C_j and E_j signal streams respectively. In this case, the expectation values associated with the eight autocorrelation channels are:

$$\text{Channel } 1_j: \quad C_{jj}^1 \rightsquigarrow \left\langle \left(p_j^L\right)^2 \right\rangle$$

[8] 'Quadraturity' is a term invented somewhere to refer to the accuracy of the $\frac{\pi}{2}$ phase-shift in the quadrature channel.

$$\begin{aligned}
&\text{Channel } 2_j\text{:} \quad C_{jj}^2 \rightsquigarrow \left\langle p_j^L \cdot q_j^L \right\rangle \\
&\text{Channel } 3_j\text{:} \quad C_{jj}^3 \rightsquigarrow \left\langle p_j^L \cdot p_j^R \right\rangle \\
&\text{Channel } 4_j\text{:} \quad C_{jj}^4 \rightsquigarrow \left\langle p_j^L \cdot q_j^R \right\rangle \\
&\text{Channel } 5_j\text{:} \quad C_{jj}^5 \rightsquigarrow \left\langle p_j^R \cdot p_j^L \right\rangle \\
&\text{Channel } 6_j\text{:} \quad C_{jj}^6 \rightsquigarrow \left\langle p_j^R \cdot q_j^L \right\rangle \\
&\text{Channel } 7_j\text{:} \quad C_{jj}^7 \rightsquigarrow \left\langle \left(p_j^R\right)^2 \right\rangle \\
&\text{Channel } 8_j\text{:} \quad C_{jj}^8 \rightsquigarrow \left\langle p_j^R \cdot q_j^R \right\rangle
\end{aligned} \tag{2.25}$$

Channels 3, 4, 5 and 6 correspond to cross-polarization autocorrelation parameters; their interpretation and possible uses will not be discussed here (one can extrapolate from the following discussion—these channels represent the cross-power between nominally orthogonal functions). Channels 1 and 7 provide signal power information equivalent to that obtained using the AUTO control. The channels of interest are 2 and 8. They each provide information about the similarity (or cross-power) between the in-phase, p_j, and the quadrature-phase, q_j, signal streams. In principle, these signals are identical but for the $\pi/2$ phase shift. Therefore, the correlation coefficients, $\left\langle p_j^R \cdot q_j^R \right\rangle$ and $\left\langle p_j^R \cdot q_j^R \right\rangle$, should be zero.

A geometric interpretation is seen by representing the signals in phasor, or function-space, notation (see Figure 2.11). In function-space, the cross-power is a sort of vector dot product. Letting $\vec{p}_j^R$ (magnitude $= 8.07$) lie on the positive real axis implies that, for a perfect $\pi/2$ phase shift, $\vec{q}_j^R$ (same magnitude) lies on the positive imaginary axis. To determine the actual phase-shift value, the gain and offset errors discussed above must be removed from the autocorrelation results:

$$\left\|\vec{p}_j^R\right\| \cdot \left\|\vec{q}_j^R\right\| \cos\left(\theta_j^R\right) = \frac{\frac{2^{12}}{N} \cdot \left[2^{-12} \cdot \sum^N \left(p_j^R \cdot q_j^R + 333\right)\right] - 333 - \left\langle p_j^R \right\rangle \cdot \left\langle q_j^R \right\rangle}{G\left(p_j^R\right) \cdot G\left(q_j^R\right)} \tag{2.26}$$

Again, square brackets denote the channel result; Poisson brackets, $< \cdot >$, denote the offset errors; G values denote gain inaccuracies.

When θ_j^R is not $\pi/2$ the phasors $\vec{p}_j^R$ and $\vec{q}_j^R$ form a non-orthogonal set

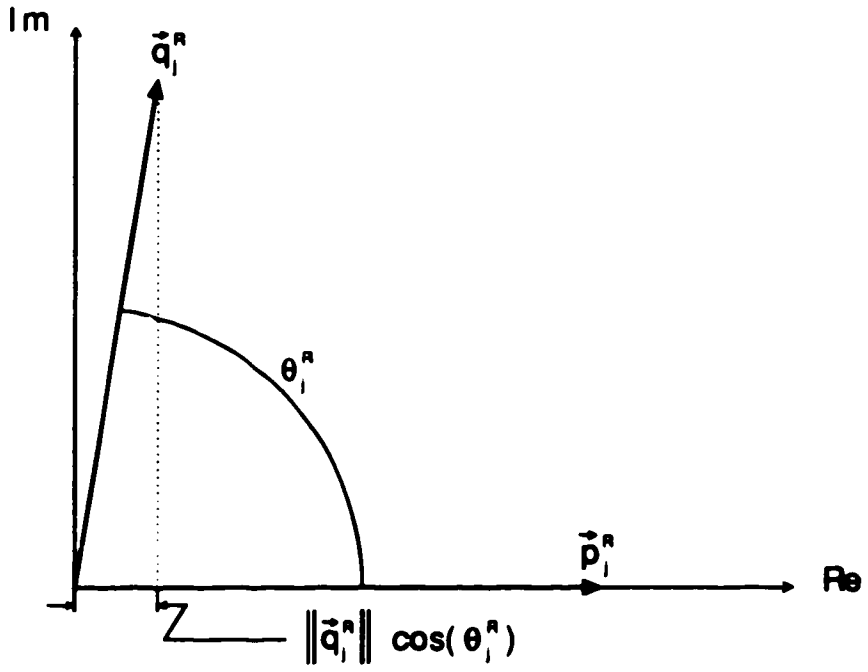

Figure 2.11: **Phasor Signal Representation:** Ideally, $\vec{p}_j^R$ and $\vec{q}_j^R$ are orthogonal functions, in which case, $\theta_j^R = \pi/2$.

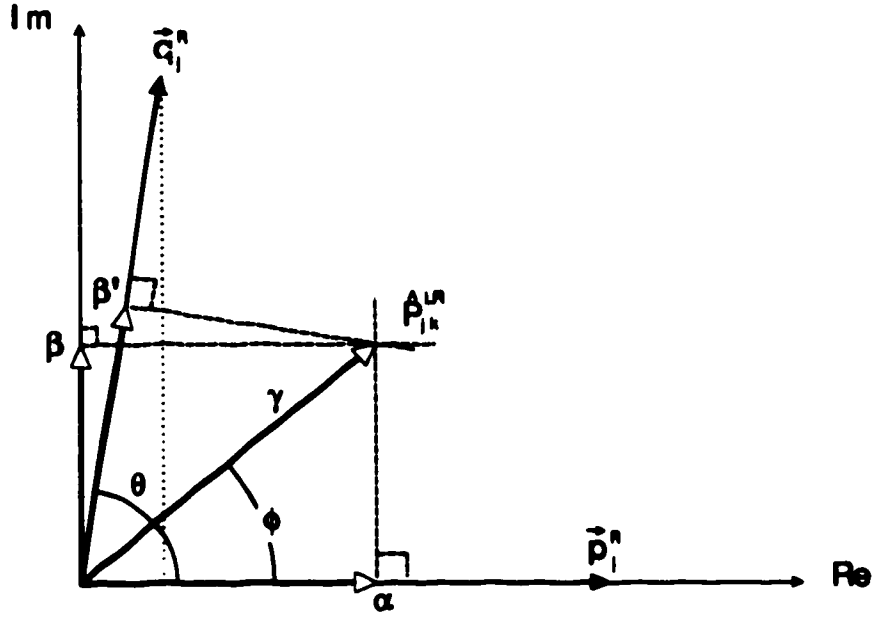

Figure 2.12: **Correlation as Function-Space Projection:** The cross-correlation vector, $\hat{P}_{j,k}$, is the projection of antenna vector $\vec{p}_k$ onto the $[\vec{p}_j, \vec{q}_j]$-plane.

spanning 2 dimensions of the function space. Recall that the correlation quantities like $\hat{P}_{j,k}^{LR}$ are also vectors on the complex plane. The real and imaginary components of $\hat{P}_{j,k}^{LR}$ correspond to the projection of some other function-vector, p_k^L, onto the spanning set $\left[\vec{p}_j^R, \vec{q}_j^R\right]$.

Referring to Figure 2.12 the value measured at the correlator is represented by:

$$\hat{P}_{j,k}^{LR} = \alpha + \jmath\beta' \tag{2.27}$$

The prime, ', is used to designate the erroneous value read from the quadrature channel. The result for an orthogonal spanning set should have been:

$$\perp \hat{P}_{j,k}^{LR} = \alpha + \jmath\beta = \gamma \cdot (\cos\phi + \jmath \sin\phi) \tag{2.28}$$

The erroneous quadrature channel result is:

$$\beta' = \gamma \cos(\phi - \theta)$$

$$= \gamma \cos\theta \cos\phi + \gamma \sin\theta \sin\phi$$

$$\alpha = \gamma \cos\phi \tag{2.29}$$

The corrected (orthogonal) quadrature value, β, can be calculated from β', α and θ:

$$\begin{aligned} \beta &= \gamma \sin\phi \\ \beta &= \frac{\beta' - \alpha\cos\theta}{\sin\theta} \end{aligned} \tag{2.30}$$

The correct cross-power vector can now be calculated:

$$\perp \hat{P}^{LR}_{j,k} = \alpha + j\frac{\beta' - \alpha\cos\theta}{\sin\theta} \tag{2.31}$$

It should be emphasized that this does not correct for any absolute phase errors between antennas. This correction merely ensures that the in-phase and quadrature channel values are orthogonal (they contain completely independent information).

When the offset, gain and quadraturity corrections are applied, the correlator results will contain fewer low-level instrument errors. These are not *data reduction* operations (they are reversible); further, they do not violate the prerequisites for self-calibration—that all phase and gain errors be *antenna based* not *correlator based*.

The penalty paid for the above improvements in data quality is the amount of CPU processing time required to accomplish the tasks. While the calculations required by equations 2.24, 2.26 and 2.31 are straightforward, performing them for 1184 channels every processing interval (5.6 sec.) is a demanding task (in the order of 10^5 floating point operations per second). Operational tests will reveal the magnitude of the improvement these corrective measures have on the data and what the real CPU requirements are to handle such a load.

Auto-Correlation Summary

It should be emphasized that the autocorrelation card is identical to the other correlator cards; its function is specified solely by its location in the system topology. By controlling the values at the a/b inputs, special channel operations[9] can be invoked. This permits measurement of data-stream characteristics like the *mean*,*standard deviation* and *quadraturity*. These can be used to enhance the quality of the results (ortho-normal correlation coefficients are attainable).

Finally, autocorrelation results provide a sensitive assessment of the analog system performance. As such they can be used for non-disruptive diagnostics while the telescope is in operation.

2.4 Correlation System Summary

The design of the correlator started with individual channel performance requirements:

- At the inputs the channel must have the necessary processing capacity (20 MBytes/second).
- At the outputs, the data must accurately represent the correlation results (32 bits of properly rounded data) and provide an efficient system interface.

In the next level of design, eight channels were grouped together to form the correlator card, a unit which satisfies the processing required for all polarization products in an interferometer pair.

The final level of system design establishes an interconnection topology using the correlation cards, to meet the requirements for the expected eight

[9]These streams make use of the command 'AUTO' to measure the standard deviation of all four data paths from each antenna. Signal means are obtained by correlating against a constant such as +13.

antenna array.

Due to the modular nature of the data distribution within the topology, the system performance is not degraded by size. This means that the board-level test results are indicative of the expected system performance. Reliable operation was verified at clock speeds of 25 MHz .

The autocorrelation card provides gain, offset and quadrature phase information which can be used to remove instrumental anomalies from the results. This removal implies a considerable amount of post-correlation processing, most of it requiring floating-point operations and transcendental functions. Whether or not the improvements in correlation results is worth the cost of the development manpower has not yet been determined (hardware costs usually pale by comparison). Assessing the magnitude of potential improvements requires extensive system-level testing.

To provide a smooth transition to the next chapter, we should take an inventory of the signals entering and leaving the correlation matrix. Each antenna row of the matrix requires the connections shown in Table 2.1. Unlisted but not forgotten is the power/ground network. The correlator cards for each row of the matrix plug into a mother-board (the chassis backplane). Like the other circuit boards in the system, the mother-board has two internal copper layers dedicated to power and ground.

This signal inventory represents the boundary between the digital (hardware) signal processing system and the control (hardware and software) system. Recall that the 8 signal streams (A, B, a, b, C, D, E and F) carry control directives as well as data. The following chapter will describe the control system design.

Data Bus	8 Output	for reading correlation results
Control Signals	5 Input	for read-out control and clearing the channels
Read-Out Status	1 Output	to indicate completion of row read-out
Clock Signal	1 Input	to synchronize signal processing
Streams A and B	8 Input	two 4-bit cross-correlation streams passed diagonally
Streams a and b	8 Input	two 4-bit auto-correlation streams
Streams C,D,E,F	16 Input	four 4-bit correlation streams passed horizontally

Table 2.1: **Correlator Matrix Signal Inventory** These are the signals required to operate each row of the correlator matrix.

Chapter 3

Controlling the Correlator System

This chapter will focus on aspects concerned with the control of correlation.

The correlation hardware described in the previous chapter has a set of signal input and output requirements. The first portion of this chapter describes how these are met and what controls are relevant to each signal set. The material here will focus on control hardware. A printed circuit board called the control/interface card (shown in Figure 3.3) was designed to control the correlation matrix. The hardware discussion will close with a review of the system packaging design.

The higher levels of system control are in the realm of software. What will be presented in the latter portions of the chapter is a discussion of the operation scenarios which were anticipated during the design. The design, implementation and testing of the control software for this system still requires several man-years of effort. The task was handed over to DRAO's software engineering group when it became clear that there was insufficient time to include it in this project; however, there were many discussions concerning the software design. [1] The design of this hardware was carried out in

[1] Preliminary software engineering was done in cooperation with Gary Hovey. The design followed an approach suggested by Gomaa [Gom84] for designing real-time multi-

light of the anticipated microprocessor/software requirements and limitations (mentioned throughout the remainder of the chapter).

3.1 Developing the Control/Interface Card

The control/interface cards (CIC's), shown in Figure 3.1, operate between the data-acquisition units (A/D converters) and the correlation matrix; there is one CIC per antenna (matrix row). The computer directs the activities of the CIC's individually or in groups.

At the end of the previous chapter, (Table 2.1), a list of input and output requirements for the correlation matrix was given. Each row has eight 4-bit data stream inputs, a byte-wide output bus, and 5 control lines. The control interface card fills these requirements on the correlator side; on the data-acquisition side are connections for the antenna's 4 A/D units. In the background is the bidirectional link to the continuum system computer.

Recall that the four A/D input streams are the in-phase and quadrature signals for right and left circular polarization. These arrive at the CIC at a 20 MSPS rate in sync with the system clock. As the digital signal values are clocked through the CIC, the 16-level A/D sample streams are 'clipped' to 14 levels to get the *primary* streams (C, D, E, F); the *cross-correlation* streams are copied from the primary streams (A = C or D, B = E or F); the *auto-correlation* streams (a, b) are either copies of A and B or they are computer controlled constants (such as 'AUTO' and +13).

The computer in control of the digital continum system communicates with each CIC via the *computer interface* — a bidirectional byte-wide path for control and commands as well as correlation results and system status. Correlation results are retrieved from the correlator cards by the CICs, to be read later by the computer.

Figure 3.2 shows the internal modules of the CIC. The cards' two mandatory functions are:

---

tasking embedded software.

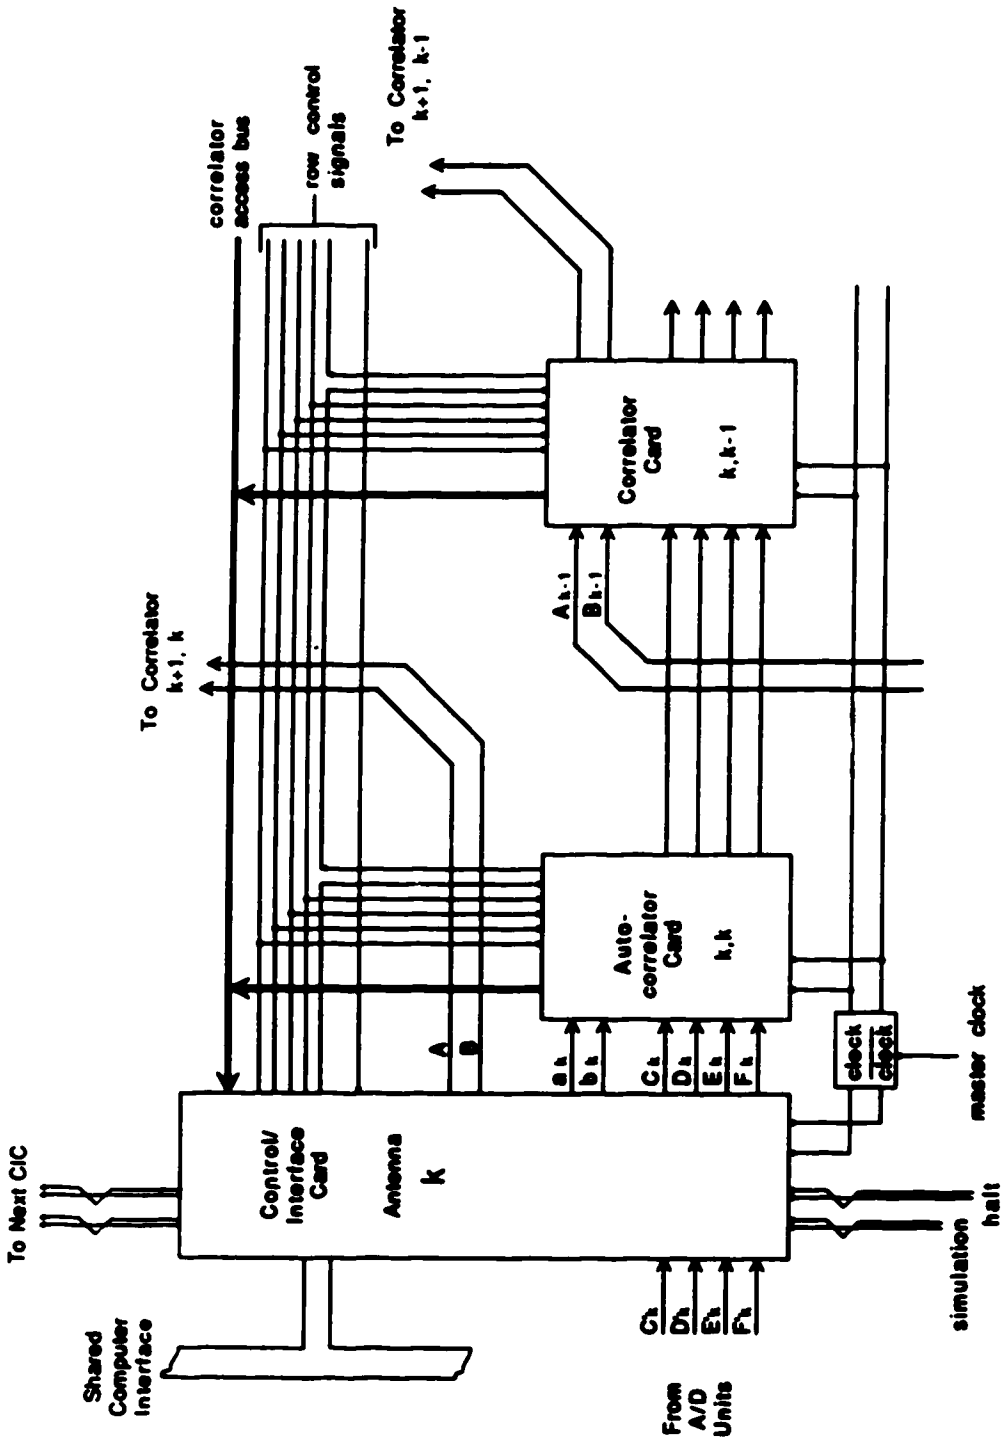

Figure 3.1: Location of the Control/Interface Card: Under computer control, the CIC's route the signal streams from the antenna A/D's into the correlation matrix.

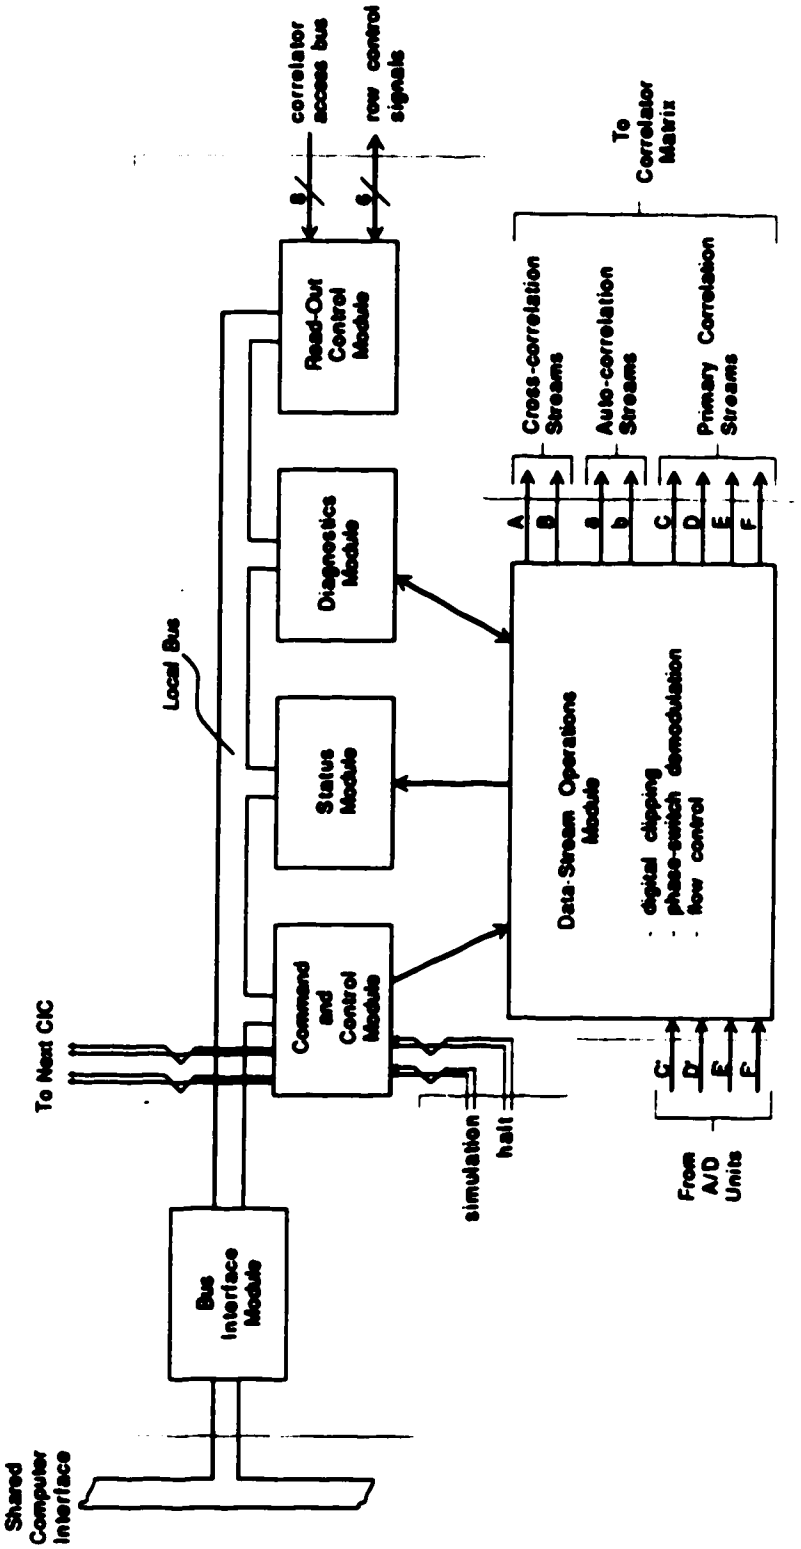

Figure 3.2: Internal Structure of the Control/Interface Card: Most of the modules are on the *local bus*.

1. To *control* the data streams flowing from the A/D converters into the correlation matrix (On, Off, selection, etc.)

2. To provide a data path for results— an *interface* from the correlator matrix to the computer.

The initial design of this card augmented the above items with requirements for status and diagnostic information. The functions were partitioned into related groups. For example, all the data-stream operations were grouped together. These groups became the specifications for the subsystem modules shown in Figure 3.2. Since most of the modules communicate with the computer, the 'bus interface' module was designed first to establish a communication protocol. The rough design was done using standard TTL components for implementing the subsystem framework. PAL's and PLD's were used to complete the design (*glue-logic*); they handle the control signals within and between modules. Of course, the design was an iterative activity; designs were refined until all the modules were compatible.

Most PAL design software allows rudimentary simulation of the device operation; this permitted some verification of each module's feasibility. Once feasibility was established, the design of the printed circuit board began. In the case of some of the simpler modules, no PAL specifications (other than a list of inputs and outputs) were necessary to proceed with the PCB layout. In such cases the PAL's were allocated with some margin for expansion by leaving some pins uncommitted. This provided contingency for unforeseen operational complexity. In fact, the hardware (PCB) design was released before the functional details were known. The final PAL designs and simulations were done during the manufacturing of the PCB. Advantages to using PALs in the design were:

- Savings in PCB real estate and power consumption as a result of the reduced number of devices required. A conservative estimate is a factor of 5:1 reduction in the number of devices required to implement the control functions (glue).

- High level design specification. Using CUPL[2] as a hardware design

[2]CUPL was a trademark of Personal CAD Systems Incorporated.

language, the operation of the circuit is specified as Boolean equations or in a state machine format. Not only is this a faster logic design technique, it is more reliable because the specification is stated directly in terms of the desired operation characteristics; no Karnaugh maps or schematic diagrams are required, hence the opportunities for human error are greatly reduced. The PAL source file is also a "readable" description of the logic design and so provides better documentation.

- PALs permit a great deal of flexibility in the design, even after PCB completion.

The last item was proven during board-level integration testing. Testing of the interaction between modules on the CIC revealed some unexpected conditions and requirements (the PAL design software would only simulate one device at a time). To correct for these, some of the PAL's had to be redesigned, one as many as five times (Rev. F). Fortunately, no PAL-related changes to the PCB (jumpers, cut tracks or new devices) were required (spare device pads were placed on the PCB in case extra logic was needed).

The Control/Interface Card (shown in Figure 3.3) is a six layer circuit board with about 70 chips on it (1,700 solder points). The board level testing revealed three errors (1 schematic error, 2 layout errors) in the metallization artwork.

Correcting these on the 5 prototypes was a delicate task due to the high component density and the hidden layers involved. In subsequent board level testing, the CIC exceeded the 20 MHz requirement (worst case design) and ran error-free at 25 MHz. It is expected that the CIC will perform equally well during system-level testing.

The following sections describe the modules shown in Figure 3.2. Each begins by describing the relevant functional requirements, followed by discussion of the implementation (design and operation).

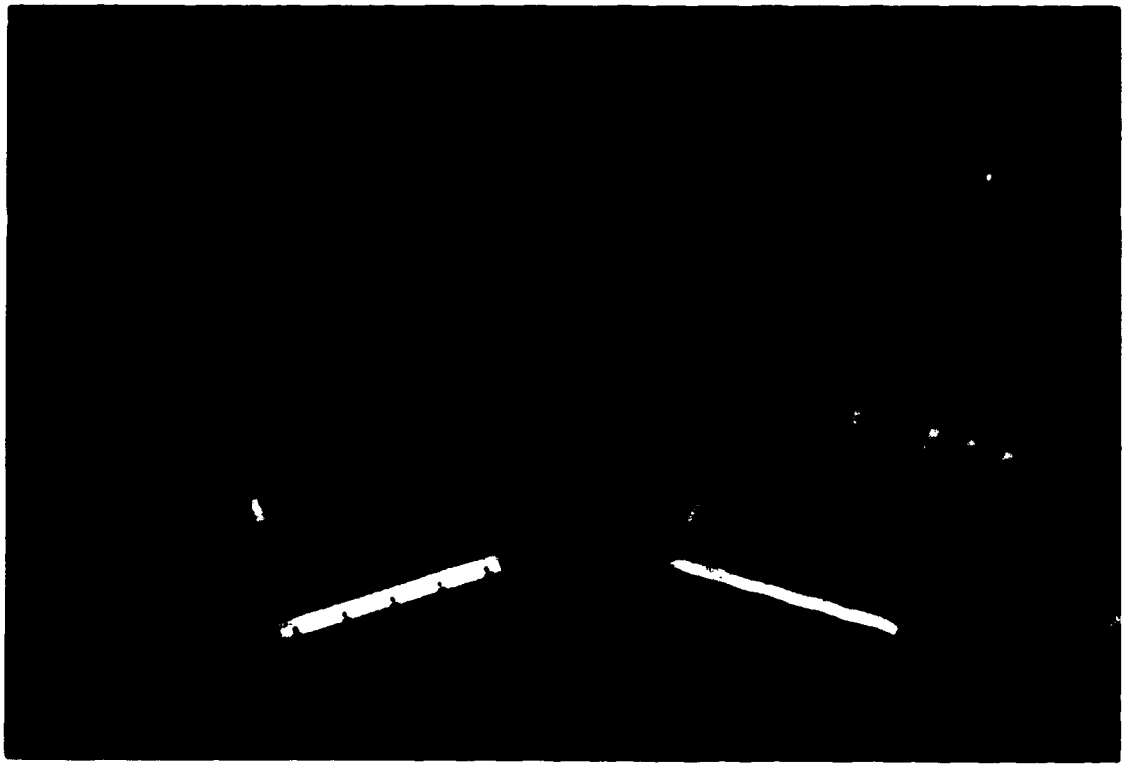

Figure 3.3: **The Control/Interface Card:** This card is a 6-layer PCB. The board design makes extensive use of PALs and PLDs.

3.2 Signal Stream Control

The signal streams flowing from the antenna's A/D converters require some alteration before they can proceed to the correlation matrix. Following is a list of tasks performed by the signal stream control subsystem:

- The 16-level A/D values are 'clipped' to 14 levels to accomadate the two embedded command values, 'STOP' (1111) and 'AUTO' (0000).
- The polarity (inverting or noninverting) of the data streams is controlled to permit phase-switch demodulation prior to correlation.
- On/Off operation of the primary streams (C, D, E and F) is controlled by a synchronous *gate* operation on the 16 signal lines:

  $$\begin{aligned} \text{On} &\leftarrow \text{14 level signed data value} \\ \text{Off} &\leftarrow \text{STOP} = 1111 \end{aligned}$$

  On/Off control occurs simultaneous with other CICs feeding the matrix.
- Cross correlation streams A and B are split off from the primary streams; A comes from either C or D; B comes from either E or F.
- Streams A and B are delayed by one stage more than the primary streams so that the data flow in the matrix is coherent (A and B bypass the delay stage that would have been incurred at the auto-correlation card site; see Fig. 3.1).
- The A and B cross-correlation outputs are disabled during direct coupling between a CIC unit and a correlator card; this allowed integration testing without a completed motherboard. It is possible that this configuration will be useful for servicing faulty cards.
- Autocorrelation streams a and b are selectable from either the A and B streams (for quadraturity measurement) or from a computer-controlled register containing fixed values (for gain and offset measurement). In the former case, the values must be coherent with the primary streams.

- Data-stream encoding is *negative logic* at the CIC outputs. This complies with the data-passing strategy described in Chapter 2.

The process path for the signal streams is shown in Figure 3.4. Two clock stages after the A/D samples arrive at the CIC, they are fed through the *Walsh PROM*. This is a 5 input, 8 output PLD which implements three of the requirements listed above. Of the five address inputs, four are presented with the data value; the fifth is used by the polarity signal which controls demodulation. Together the 5 bits compose the look-up table address. The two 4-bit values (1 byte) that appear at the PROM output are the transformed representations of the input value (14-level coding, demodulated, in negative logic format).

One 4-bit value from each PROM continues on through as the primary data path (C,D,E,F); the other branches off to become the cross-correlation data stream (A,B). The selection of streams A and B is done immediately after the PROM. The final device in the cross-correlation pipeline is a register with tri-state outputs; these outputs are enabled only when the CIC is running in the chassis (discussed at the end of this chapter). If the CIC is connected directly to a correlation card, for bench top testing, these outputs are disabled. The control signal (active low) sees a high signal when connected to a correlator card and a low signal when installed in the chassis.

Autocorrelation stream selection (a,b) is done using more 4-bit selectors between stream A and *a'* a local register value (called the *autocorrelation* register in Figure 3.4).

The final requirement concerns t' On/Off control of the primary data streams. This is done downstream from the Walsh PROMs using registers with reset controls. The signal value is inverted logic at this stage so that an active reset produces a hex:F, i.e. (1111 = 'STOP') at the primary stream output. As discussed in the previous chapter, this value halts correlation.

The requirement to synchronize the On/Off activities with the other CIC units in the matrix is met by the 'Operation Sequencer' block in Figure 3.4. This block of control hardware receives an external signal (the 'HALT' signal) which coordinates the operation of the correlator matrix. This signal is sampled on the rising clock edge; the clock edge also transfers the received

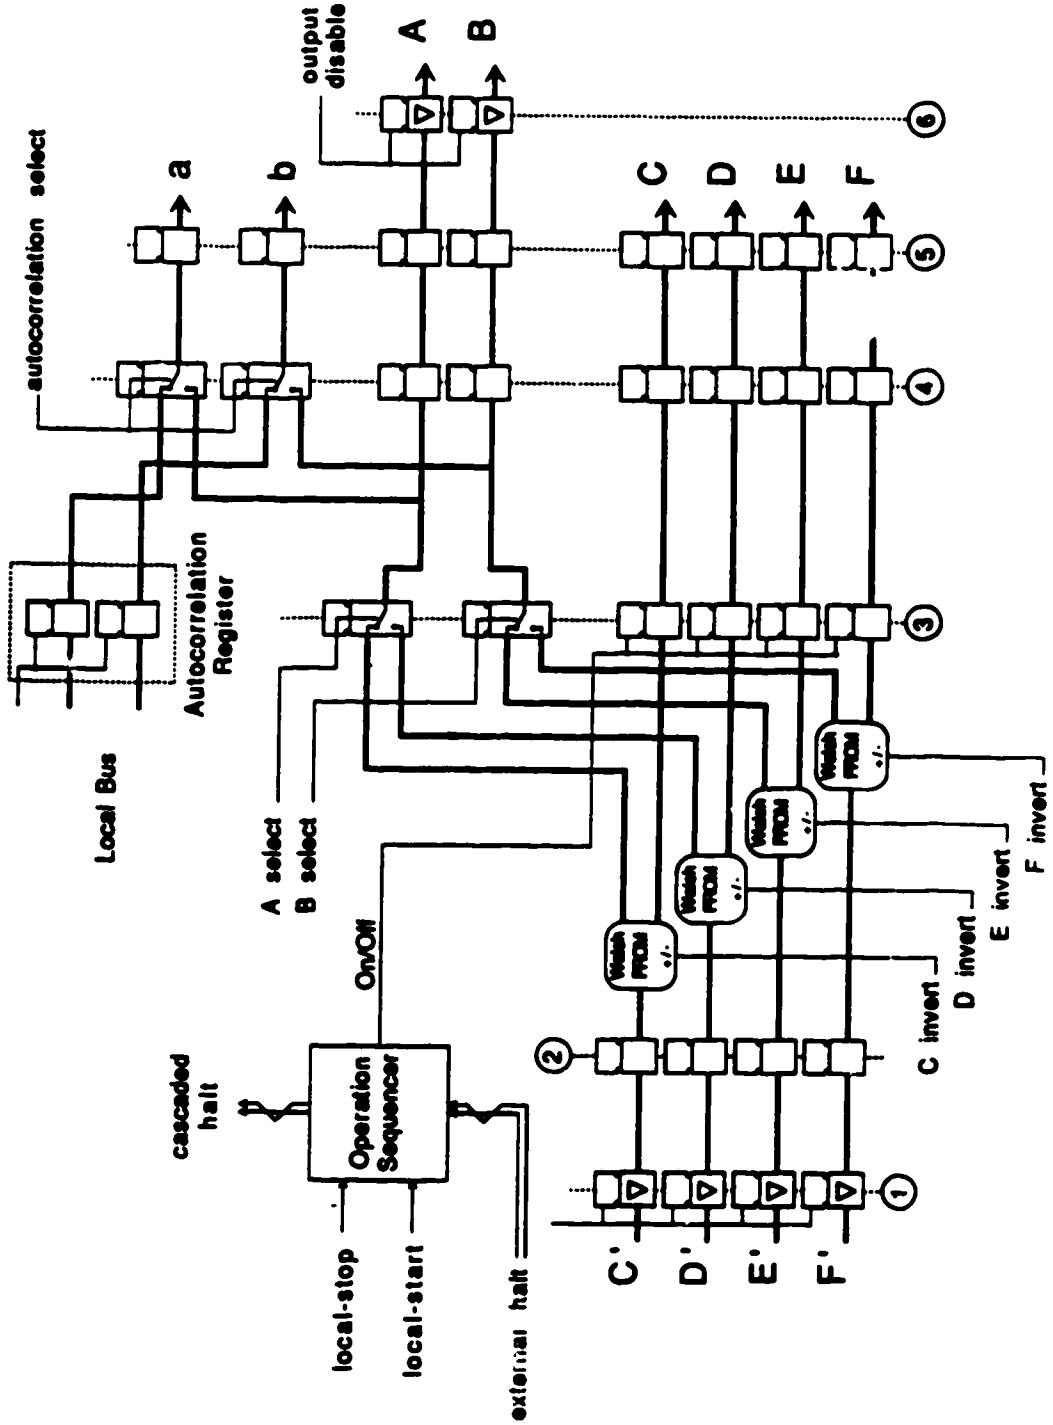

Figure 3.4: Signal Stream Control: Digital signal values arrive from the A/D units in 16-level format. Recoding to a 14-level format is done at the Walsh PROM (look-up table transformation) as is the sign inversion/noninversion required to demodulate the phase-switched signals. The remaining stages are for autocorrelation and cross-correlation stream selection.

signal condition to the output. The signal is passed from one CIC to another on successive clock cycles.

When a CIC detects a change on the 'HALT' line, it begins a count-down sequence and implements the 'HALT' directive N clock cycles later. The CIC next in line will wait $N - 1$ cycles and so on. The result is that all the CIC units will start and stop correlation simultaneously. The number of cycles in the count-down sequence is determined by the CIC's 4-bit address (switch selectable). For example, a CIC card with address h will have a count-down delay of $N = 16 - h$ cycles. The first CIC in the 'HALT' signal chain must have the lowest address and subsequent CICs must have successively larger address values. The original 'HALT' signal[3] is provided to the digital continuum system by the telescope control computer.

For development and testing, the 'Operation Sequencer' block has the option to override the external 'HALT' signal. In this mode, called the 'local mode', the local 'HALT' signal is controlled by the digital continuum system's computer. This permits system-level independence for development and off-line testing. Notice that only one CIC needs to be in 'local mode' for the continuum computer to take control of the entire correlation matrix. As an example, correlator rows 6, 7 and 8 can be taken off-line for testing by placing CIC 6 in 'local mode'. This will not interfere with normal observations using antennas 1 through 5.

Data Stream Control Summary The primary data streams (C, D, E and F) are conditioned and controlled as values are clocked through the CIC on their way to the correlation matrix. Phase-switch demodulation is controlled independently for each of the primary streams (4 control lines required). The cross-correlation and autocorrelation streams are chosen using 3 selection control lines. All the controls are single-bit conditions corresponding to bit locations in each CIC's control registers (discussed later).

The On/Off control of the correlation matrix data streams is simultaneous (same clock cycle) across all the matrix inputs. The 'Operation Sequencer blocks on each CIC accomplish this by postponing the response to the 'HALT'

[3]For historical reasons, the telescope system's 'HALT' signal is called 'SST-Reset'.

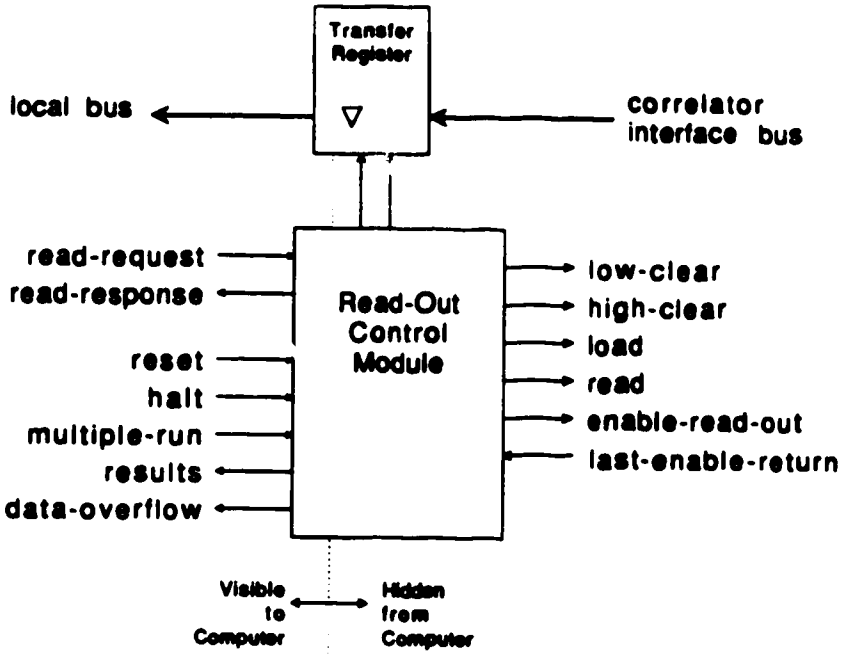

Figure 3.5: **Read-Out Control:** This module assists the computer by (i) prefetching result bytes, (ii) hiding the device-specific timing requirements and (iii) simplifying the interface (2 lines instead of 5).

signal.

3.3 Result Read-Out Control

This section will present the logic circuitry controlling the retrieval of correlation results. Chapter 2 discussed the operation of the five control signals associated with the 8-bit data-bus (see Figure 2.8). The module controlling the readout operates between this correlator matrix bus (shared by the correlator cards) and the data-bus internal to the CIC (the local bus).

The module, shown in Figure 3.5, aids the computer, which has the task of reading correlation results from as many as 1184 channels through 32 CIC's. The following description of operations will illustrate all of the aspects of its design.

When the continuum control computer gains control of the digital continuum system (say, following start-up), the correlation channels need to be prepared for processing. The processor can invoke (on any CIC, at any time) the 'CIC-reset' condition. As well as initializing other subsystems on the CIC, this condition has the following effects on the 'Read-Out Control' module:

- 'low-clear' is activated
- 'high-clear' is activated
- 'load' is activated

These signals are held active until channel contents have been erased. The 'clear' signals erase any residual bits present in the channel accumulators (bits set by previous correlations or by power-up conditions). The 'load' signal clears the channel read-out storage registers. Having done this, the 'Read-Out Control' module leaves activating the three controls and the channels are ready to begin correlation operations.

The very first correlation run begins in response to 'HALT' going inactive. The correlators begin processing data and accumulating results as the valid data values propagate into the matrix. The correlation run will come to an end millions of cycles later when 'HALT' has been activated; operations cease as the matrix is filled with 'STOP' (1111) values. The activation of 'HALT' causes the 'Read-Out Control' module to begin the sequence of actions required after correlation (see Chapter 2 on read-out control):

Post-Correlation Sequence

1. After waiting a few cycles for processing to end, 'low-clear' is activated; this rounds the channel results and erases the low-order accumulator

results.

2. Since step 1 may have generated a count pulse in the ripple counter, a pause for ripple propagation is required (about 1 μsecond).

3. 'load' is pulsed to store the rounded 32-bit channel results in the channel-resident read-out buffers.

4. 'high-clear' is pulsed, flushing results out of the 32-bit counters, preparing the channels for the next correlation run.

The post correlation sequence above takes about 10 μseconds, after which, the entire correlation matrix is ready for another correlation run. Action from the controlling processor has not been required yet.

The 'Read-Out Control' module begins the transfer of results after step 3 of the above sequence. Transfer begins with the module activating the 'enable-read-out' signal (recall that this signal is daisy-chained through all the channels in a matrix row). In adherence to the timing limitations of the 32-bit accumulator devices, the 'scan' signal is activated to place a result byte on the bus. After this byte is latched into a holding register on the CIC, the 'scan' signal is deactivated. When the continuum computer eventually reads the holding register, the module will prefetch the next result byte from the channels. The sequence repeats until the 'last-enable-return' signal goes active indicating all the channels have been read. If, for some reason, the computer has not read all the results by the time the next correlation run ends, the error flag 'data-overrun' is raised. The post-correlation sequence proceeds as usual causing any unread data to be overwritten.

The 'Read-Out Control' module supports one operations alternative. If several correlation runs are to be accumulated in the channel hardware, the 'multiple run' control line causes the following changes to the operation described above:

- Only step 3 of the post-correlation sequence is executed. This permits intermediate results to be read, if desired.

- The 'data-overrun' condition is not possible and so this flag is not used.

After the last correlation run in the series has started, the 'multiple run' flags on each CIC should be lowered to restore the normal post-correlation channel preparations.

As mentioned in Chapter 2, the channels have ample capacity for such multiple-run operations.

Read Out Control Summary: This module provides assistance to the controlling processor in the following areas:

- The routine operations of preparing the correlation matrix for results transfer and subsequent correlation are hidden from the computer. It does not need to *know* about control-line details and hardware timing specifications. These operations are done in parallel; all the CIC units execute these operations as soon as the correlation run ends. This greatly simplifies the software design and relaxes the computer's real-time response requirements.
- The results from the channels are transfered to holding registers (on the CIC's) which have a much faster response time than the channels. This allows quick access to the data for the controlling processor.

The net result of the hardware design is a marked reduction in read-out software complexity, as well as an increased amount of control time available to the control computer for system-level operations. The CIC units manage all the low-level controls associated with the channels (information hiding); the computer simply makes repeated 'reads' from a register location on the CIC to obtain the correlation data.

Post-correlation activities (system down-time) last about 10 μseconds after which the system is ready to begin the next correlation run. The fast turn-around time is due to the result buffering at each channel; results do not have to be read before subsequent processing begins. At present, the interval between correlation runs is determined by the settling time of the RF local oscillator phase-lock system; following a phase-switch (180° phase reversal) the system requires several milliseconds before the antenna LO phases are accurate to within 1°.

This new continuum system provides ample margin for improving the phase-switch efficiency; this would become an important parameter in a faster switching (better noise immunity) system. In such a system, the controlling computer might not be able to retrieve and process the results for each correlation run. For this reason the channel control was designed to accomodate multiple correlation runs; results can accumulate in channel *hardware* instead of computer *software*.

3.4 Hardware Testing and Diagnostic Features

This section concerns the design of hardware features extraneous to the correlation requirements of the system. The reasons for including these features can be grouped in two categories:

Development: The ability to perform tests at the subsystem level allows design faults to be detected earlier in the project.

Diagnostics: Embedded testing features allow the system integrity to be monitored so that faults can be detected, localized and repaired quickly.

In the absence of human and hardware fallibility such features would be unnecessary; the design would work the first time and it would never fail.

This section will focus on three areas where testing features were included in the design of the control interface card (CIC). The first is the most ambitious, the diagnostic RAM system; second is the correlation timer; the last features discussed are those permitting bench top testing.[4]

3.4.1 Diagnostic Data RAM

[4]One aspect of diagnostics, the monitoring of CIC status bits, is best left until the next section on software control.

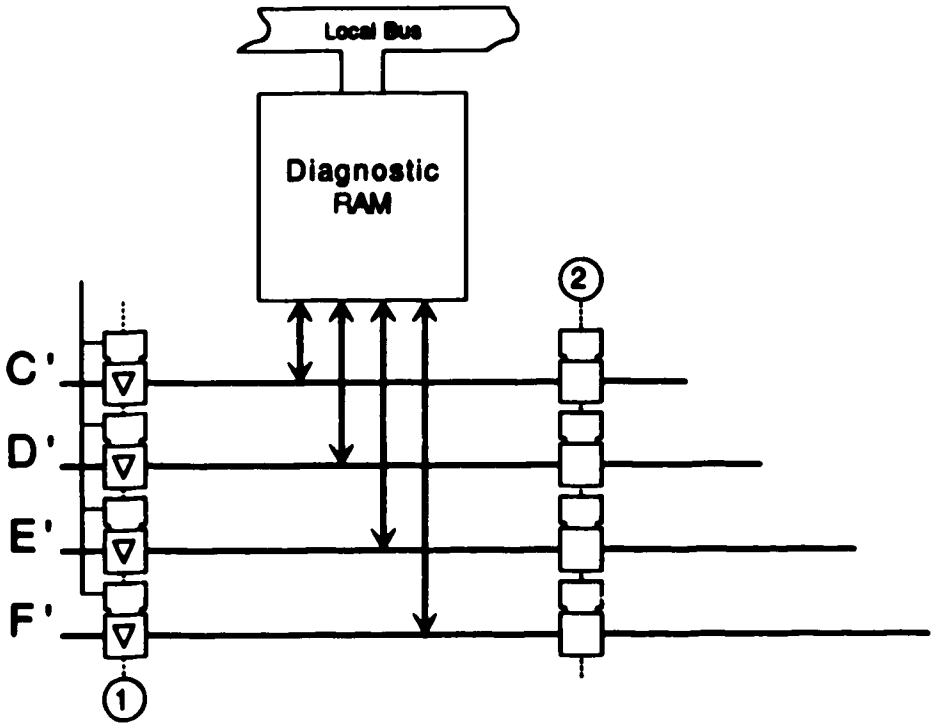

Figure 3.6: **Diagnostic Data-Stream RAM:** Data-stream values can flow *to* or *from* the RAM devices allowing both *upstream* and *downstream* diagnostic operations to be performed. Note that the 'access point' is upstream from the 'Walsh PROM' (see Figure 3.4).

A typical test of a system module, whether hardware or software, consists of feeding the component under test with a predetermined set of inputs while observing the output. The CIC units are located at a strategic point in the digital continuum system — at the boundary between data acquisition and data processing subsystems. Having *access* to the data streams here would allow the testing of both subsystems.

Downstream: Testing the Correlator

To perform input/output tests on the digital correlators requires a source of digital data to mimic the antenna data streams. To provide complete flexibility as to the test sequence used, fast memory chips were selected as the data source. By using memory devices, any sequence of values can be loaded by the slower computer and used for the high-speed test. During correlator testing these chips each send a data sequence of 4,096 4-bit values (one chip per data-stream). The memory is read in sync with the system clock so that the correlators are tested at full speed. The point of memory access to the data stream system is shown in Figure 3.6.

The RAM on each CIC is loaded with its test sequence by the controlling processor. To run a simulation test the first CIC (the Leader) in the matrix must be placed in 'local mode' (discussed in the previous section). Prior to the processor issuing a 'start' command, the 'simulate' flag on the Leader CIC must be set. The flag condition is passed along to the other antenna CIC units (Followers). Its purpose is to inform the Follower CICs that the correlation will only run for 4,096 cycles (the size of the test sequence memory). Each CIC has a 12-bit address counter which begins the run at hex:000 and ends at hex:FFF (4,095). During simulation, the On/Off control of each CIC's data streams is locked to its own address counter (instead of the cascaded 'HALT' signal).

The test sequences can be used to stress different parts of the correlation system. Some examples of system stress tests are:

1. Sequences producing current surge conditions in the devices used for signal distribution. The bin:0111 to bin:1000 transition is an example.

2. Test values selected to produce surge conditions at the multiplier or adder outputs (within the channels).

3. Test sequences which generate large multiplier outputs to produce the highest possible accumulation rates. This will test the performance of the final stages of accumulation.

Since any errors that might occur are likely to be single bit errors, they would probably not show up in a single 4,096-value simulation run. Recall that the 12 least significant channel accumulation bits are not visible to the controlling processor. Repeating the same test sequence 2^{12} times has the effect of left-shifting the channel results by 12 bits. This will give greater error visibility. Elapsed test time for $2^{12} \times 4096$ samples is less than 1 second. **Note** that this test requires the use of the 'multiple-run' option (even on the last simulation) to prevent premature results rounding.

Upstream: Testing the Data Acquisition

The diagnostic RAM can also be used to log the output of the A/D conversion units. A given data stream can be examined by recording 4,096 consecutive samples in the memory. These samples can be processed using the following techniques:

Fourier Analysis of the contents will reveal the spectral power distribution. This will detect bandshape anomalies produced by faulty RF or IF components (see Figure 3.7).

As well, the magnitude of local-oscillator leakage into the signal path can be monitored. Such leakage would show up as a spike on an otherwise 'flat' spectal power pattern. A similar pattern would be expected in the case of man-made radio noise. Even weak signals from satellites, aircraft or automobile ignition can overpower the fainter signals from space. To obtain suitable sensitivity for the detection of low-level man-made noise would require numerous sample, transform and integrate cycles (computer intensive). If this new continuum system (with wider bandwidth) is more susceptible to man-made noise, implementing such

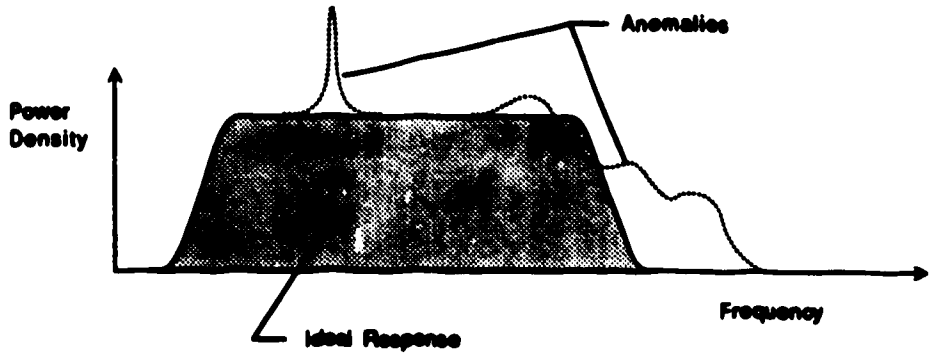

Figure 3.7: **Signal Path Anomolies:** Fourier analysis will detect deviation from the desired spectral power distribution.

an 'interference monitor' (an investment in computer hardware and software development) might be worthwhile. While unable to correct for the contamination, the monitor could 'flag' suspect observations, which would help reduce the noise in the final image.

Histogram Analysis examines the frequency of occurence of each 4-bit data value; this provides information on the A/D converter performance. The data value distribution should reflect the probability density function (PDF) of the analog input signal (nominally Gaussian). During normal operation, the data value histogram should resemble the Gaussian PDF of the antenna signals. During system testing, the dynamic characteristics of the A/D converters can be measured as follows.

An accurately known analog input signal is required. This will allow the quantization thresholds to be deduced from the data value histogram and the input PDF. A sinewave is a particularly useful waveform because the quantization thresholds are easily calculated from the data value histogram.

The normalized histogram values h_j represent the probability of detecting the values, j, at the quantizer output; the quantizer threshold locations are easily obtained as follows:

$$T_j = \sin\left(\pi \sum_{i \le j} h_i - \frac{\pi}{2}\right) \tag{3.1}$$

T_j is the threshold value normalized to the amplitude of the input sine wave; it is the analog location of the boundary between digital value j and value $j+1$. For the 4-bit quantization used in this system, j ranges from -8 to $+7$ (2's complement format). The frequency dependence of the quantization thresholds can be observed by performing this test at several points within the correlation bandwidth.

Summary of the Test RAM Capabilities

The data stream test RAM can be used in two ways: signal sequences can be injected into the correlation matrix (downstream diagnostics); the RAM can record a burst of signal values coming from the data acquisition system (upstream diagnostics).

Figure 3.8 illustrates the testing modes for the signal streams. The integrity of the data acquisition subsystem can be monitored: Fourier analysis focuses on analog performance; data value probability analysis focuses on quantizer performance.

The integrity of the correlator (and CIC components downstream from the signal stream RAM) can be easily tested. Errors will either be isolated, as with the failure of a channel's accumulator, or they will be shared by several channels; 'shared' errors can be traced back through the system to a common fault (perhaps a bad connection or failing component). The 'downstream' configuration was used extensively during the board-level integration testing of the control interface card as well as the correlator cards. The performance benchmarks mentioned in Chapter 2 were achieved using the RAM test feature.

3.4.2 Correlation Interval Timer

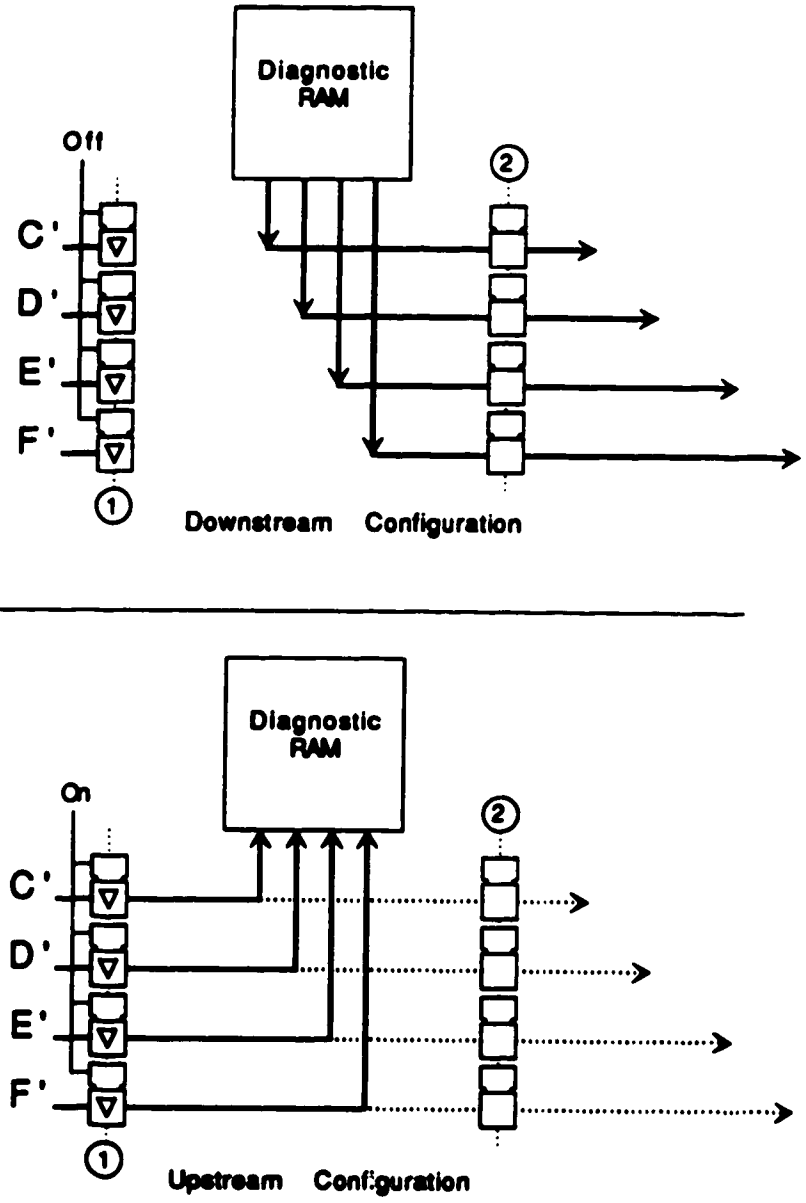

Figure 3.8: **Diagnostic Configurations:** The dashed flow arrow shown in "Upstream Configuration" indicates that the diagnostics will not interfere with normal operations.

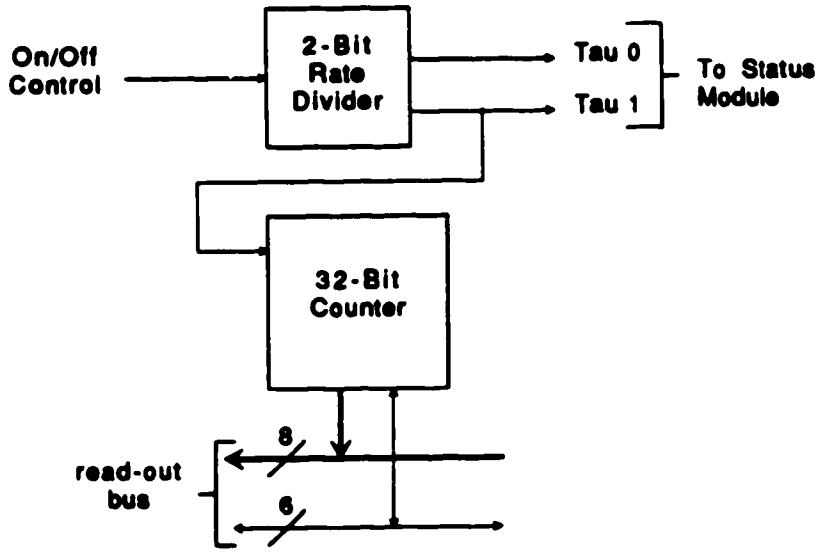

Figure 3.9: **Correlation Interval Timer:** The 'timer' counts the process cycles in a correlation-run; the result is used for system integrity checking. The 'timer' shares the read-out bus with the correlators.

Another diagnostic feature included in the design of the control interface card is the correlation-run interval timer. 'Timer' is a misnomer since the circuitry is actually a counter (see Figure 3.9); the 'timer' counts the number of processing cycles (clock pulses) occurring in a correlation run (signal streams 'On'). The principal component used here is the same 32-bit ripple counter used in the correlator channels. The 'timer' counter shares the read-out control signals and data-bus used to retrieve the channel results; the first four bytes of the read-out sequence are 'timer' bytes. Being an older NMOS device, this counter is relatively slow. The 20 MSPS system clock rate must be scaled down by 4 before the counter can operate correctly. The two bits from the divide-by-4 scaler are sent to the CIC's status registers (see the next section).

The measurement of the correlation interval can be used to verify proper operation of the On/Off control circuitry on each CIC; all the CIC timer values should be identical. Since the correlation run duration is nominally fixed, the values should fall within some small range. Similarly, correlator values should fall within some small range. Also, knowing the exact number of process cycles N permits a precise removal of the accumulator bias ($N \times 333$) from the channel results. Lack of precision at this stage adds noise or bias to the image data.

A final consideration is the normalization of channel results in the case where the correlation-run times vary in length (by greater than say 0.5%). The normalization algorithms in Chapter 2 show how N is used to correct for such variations.

3.4.3 Bench Top Testing

The last aspect of built-in-test capability involves the packaging design of the CIC rather than the unit's function.

At an early point in the design of the CIC, the type and size of its connectors had to be selected. The final system configuration established the packaging requirements baseline (card-cage chassis, backplane interconnect strategy, cabling, etc.). Enhancements to the baseline were made to aid

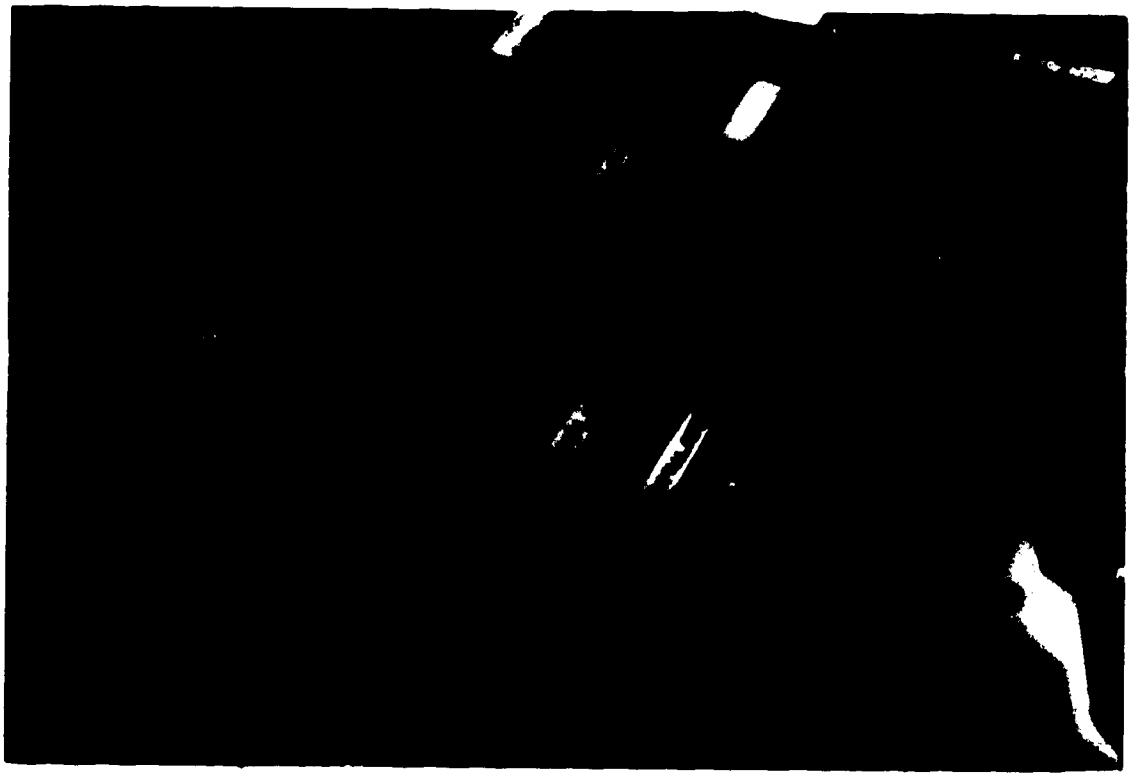

Figure 3.10: Development Features: The indicator lights on the end of the CIC provide visual status information during development testing and operations. The auxiliary power inputs (used in the 'bench-top' configuration) are also shown. Test points for oscilloscope probes are located across the board.

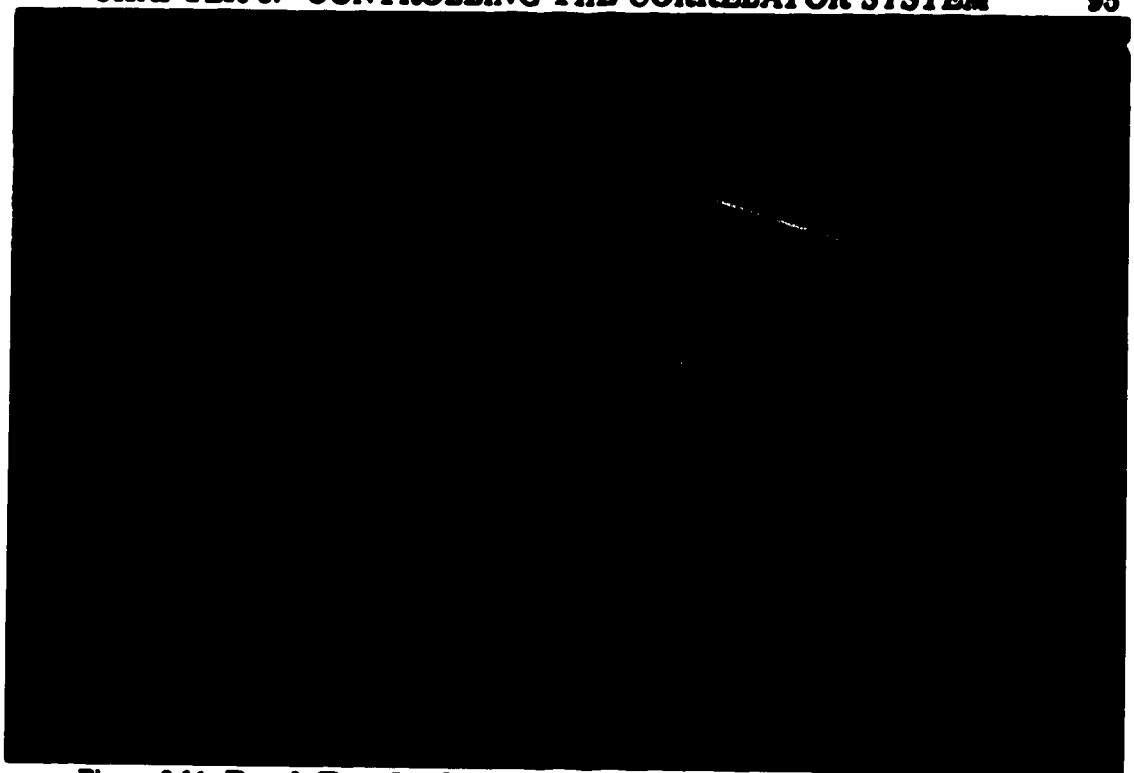

Figure 3.11: Bench-Test Configuration: allows the CIC to be mated directly to a correlator card.

the development and testing of the CIC. It was more efficient to design-in some test features than to simply satisfy the baseline, and then have to design/build test fixtures later. Figure 3.10 shows some of the development and test features built into the board. The main CIC connector (96-pin DIN) was selected and positioned to allow it to be mated directly with a correlator card. Figure 3.11 shows the selection of connectors on the CIC and the correlator boards. Figure 2.5 shows these cards mated.

In the baseline configuration, cards obtain power and ground from the chassis backplane's DIN connectors. In the bench-top configuration, alternate contact points for power and ground are needed on the CIC. Two spade-lug contacts on one end of the board are used to power the CIC which, in turn, powers the correlator card through the mated DIN connectors.

This configuration required the connector pin assignments to be mirror images (every CIC input must mate with its corresponding correlator card

output and vice-versa). The only mismatch occurs with the cross-correlation signal streams (A and B) on both the CIC and correlator card; both cards use these pins as outputs. Conflict was easily avoided by using tri-state drivers on the CIC's A and B output signals. The CIC detects the configuration (backplane or correlator card) and activates the outputs only for the backplane case.

Having this bench-top test capability meant the completion of the chassis and backplane PCB was not a prerequisite to the functional testing and board-level integration. The bench-top configuration can be used in the future during further developments or board repairs.

3.5 Computer Control of the DSP System

This section will present the controlling computer's 'view' of the new digital correlator system. Access to the system is provided by the parallel bus shown in Figure 3.12; up to 16 CIC units can share this link.

Using the link, the computer can perform a variety of operations: control and/or test conditions can be set up; commands (like Start, Stop or Reset) can be issued; system status can be monitored; correlation results can be read. These transactions make use of the 8-bit address bus and the bidirectional byte-wide data bus.

3.5.1 The Computer Access Bus

The computer uses the port shown in Figure 3.12 by issuing an 8-bit address; this selects a particular CIC and the specific transaction (discussed later) desired.

The computer and the selected CIC signal each other using two pairs of handshake lines: one pair is dedicated to computer-to-CIC data transfers (*write* transactions); the other pair is for CIC-to-computer transfers (*read* transactions).

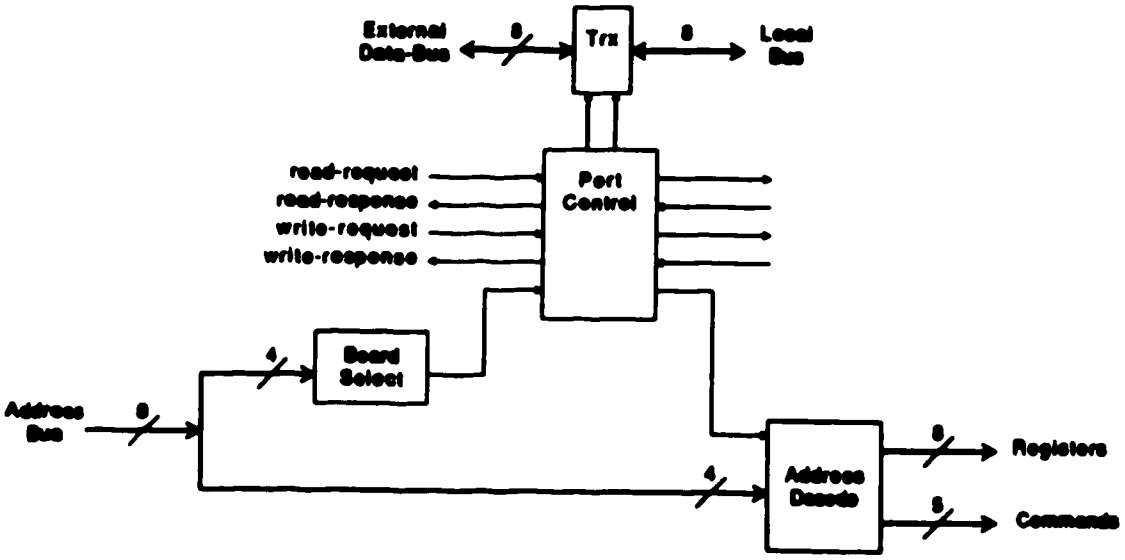

Figure 3.12: **The Computer Access Bus:** For computer control, the CIC's support a shared parallel access bus. Half of the 8 address bits are used to select a CIC (16 possible); the other four bits specify (to the selected CIC) the transaction required. The four hand-shake lines use a double-interlock transfer protocol.

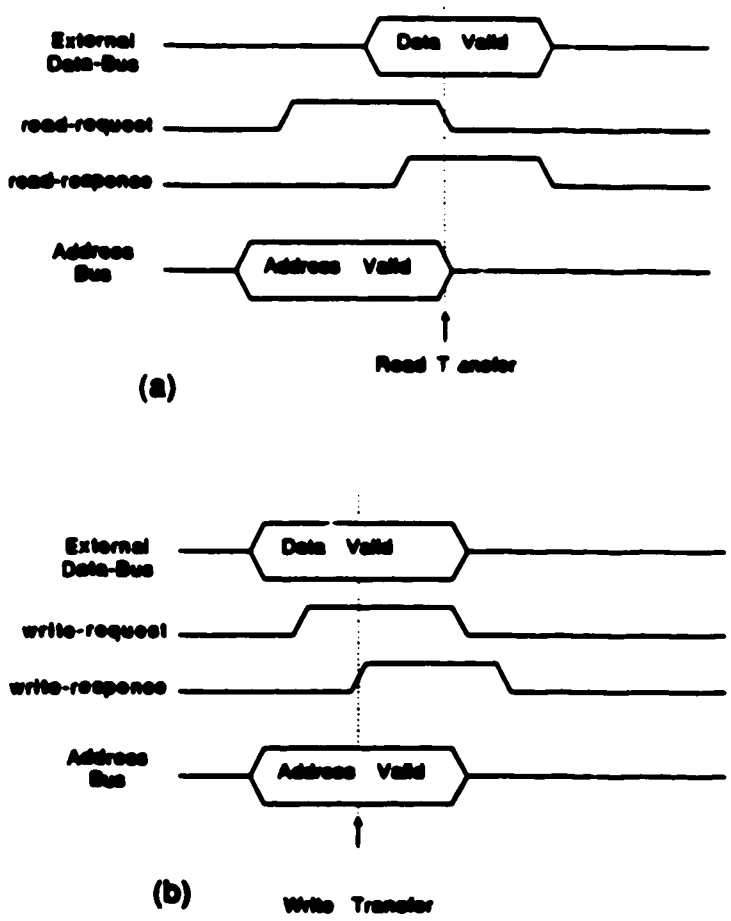

Figure 3.13: Bus Transaction Protocol: Bus activity for (a) read and (b) write transactions is shown. Arrows indicate the instant the receiving party 'captures' the data.

Each handshake signal pair is composed of a 'request' line, controlled by the computer, and a 'response' line, controlled by the CIC. An active 'request' signal indicates a valid address is on the address-bus and in the case of a 'write-request' valid data is on the data-bus as well. When the CIC activates the 'response', it indicates that the appropriate action has been taken: 'write-response' says "I've got the data."; 'read-response' says "The data you want is on the bus.". The computer finishes the transaction by deactivating the 'request' and the CIC indicates it is disengaged when the 'response' is deactivated. This transfer protocol (double-interlock) was chosen because it is asynchronous and robust; it does not depend on the speed of the communicating devices (no further action is taken until response is obtained). Figure 3.13 shows the bus activity during read and write transactions.

3.5.2 The CIC's Internal Address Map

Commands: This discussion concerns the least significant four bits of the 8-bit bus address; they are used to specify locations *within* the selected CIC. Each CIC unit has a memory mapped architecture where some address locations receive data bytes (write only) some transmit bytes (read only), some do both (read/write) and some do neither (address mapped commands). This last group is one where a 'write' transaction to a particular address (hex:8–F) will activate a control line on the selected CIC. For example, to issue a 'CIC-reset' command, the computer writes to CIC address "8". The various commands and their corresponding address locations are shown in Table 3.1.

Commands cause the selected CIC to execute a sequence of operations. The CIC response is to the active-going command transition so that the command duration is not crucial. Commands are only active for the duration of the 'write-request' strobe.

Registers: The CIC registers capable of transfering data bytes are shown in Table 3.2.

The only bidirectional locations are the Test-RAM registers. The computer reads or writes to Test-RAM locations determined by the address

4-bit address	invokes the command
8	CIC-RESET
9	...not used
A	...not used
B	START
C	STOP
D	...not used
E	RESET-RAM-POINTER
F	INCREMENT-RAM-POINTER

Table 3.1: CIC Command Map

4-bit Address	Register Name	Transactions Allowed
0	Status 1	read only
1	Status 2	read only
2	Control 1	write only†
3	Control 2	write only†
4	Autocorrelation	write only†
5	CD-RAM-access	read/write
6	EF-RAM-access	read/write
7	Timer/Cl nnel Results	read only

Table 3.2: CIC Register Map: The three 'write-only' locations, denoted '†', are double-buffered; they are stored normally by the CIC during a transaction but they are not brought into effect until correlation has stopped. This permits the computer to preconfigure the digital correlator system.

bit position	flag name	On action	Off action
0	Init Control 1	initialized	not initialized
1	Autocorrelation Select	quadraturity	gain/offset
2	Stream B Select	choose E	choose F
3	Stream A Select	choose C	choose D
4	Invert C	C inverted	C noninverted
5	Invert D	D inverted	D noninverted
6	Invert E	E inverted	E noninverted
7	Invert F	F inverted	F noninverted

Table 3.3: Control 1 Flags

pointer (a 12-bit counter). The address pointer is shared by the C/D RAM and the E/F RAM. The commands 'reset-RAM-pointer' (hex:E) and 'increment-RAM-pointer' (hex:F) are used by the computer to select the desired location in the 4,096 byte blocks. Each byte transfered to or from the RAM is composed of two concatenated 4-bit signal stream values; for example, values associated with streams E and F are written/read together as one byte.

The reading of timer and channel results was discussed earlier in the chapter. This register location behaves much like an input port; the computer repeats read transactions at this point until all the results have transfered.

3.5.3 The Control Registers

The control registers contain bit flags whose values determine the operating configuration of the CIC. Tables 3.3 and 3.4 show the flag-bit locations in 'Control 1' and 'Control 2'.

The CIC allows the computer to send new control parameters at any time, but it postpones control-flag configuration changes until the correlation run has ended. The CIC provides *double buffering* at these registers. When the computer sends new control information during a correlation run, the control bytes are captured and held in storage registers. As soon as the correlation run is over, the new control flags are brought into effect. This permits configuration control to be pre-loaded; the correlators do not have to

bit position	flag name	On action	Off action
0	Init Control 2	initialized	not initialized
1	...		
2	Simulate	RAM source	A/D source
3	Sample	log A/D values	
4	...		
5	Multiple Run	many runs	one run
6	...		
7	Cascade Control	external mode	local mode

Table 3.4: **Control 2 Flags**

be idle while the computer sets up the next correlation run. Double buffering is also provided for the autocorrelation register.

Initialization Flags: During initialization (a 'CIC-reset' command) the values in both the control registers are cleared. The bit-0 position in both control registers is used as a verification flag. Whenever the computer loads this register, it should 'turn on' (write a '1') this flag bit. The status registers reflect this bit's condition and hence provide a crude verification that some control is in effect.

Stream Selection Flags: Bits 1, 2 and 3 of Control 1 operate the signal stream multiplexers for choosing the cross-correlation and autocorrelation stream sources. Bits 4 through 7 dictate the phase demodulation at each of the four primary signal streams (C,D,E,F).

RAM Operation Flags: In 'Control-2', bits 2 and 3 control the operation of the diagnostic RAM. The different conditions are described in more detail in the next section (Status Registers).

Post-Correlation Mode Flag: Bit-5, when turned 'On', permits multiple correlation runs to be accumulated in channel hardware.

bit position	flag name	On means	Off means
0	Control 1	initialized	not initialized
1	Control 2	initialized	not initialized
2	Autocorrelation	initialized	not initialized
3	Results Ready	more results	none left
4	Data Flow	correlating	idle
5	Sampling	A/D → RAM	C,D,E,F not logged
6	Simulating	using RAM data	using A/D data
7	External 'Halt'	Upstream Idle	Upstream Busy

Table 3.5: **Status 1 Flags**

Local Mode Flag: Bit 7, when 'Off', places the CIC in 'local mode', where the 'Halt' and 'Simulation' controls from the *outside* are ignored. Instead, the CIC responds to the on-board 'Start' and 'Stop' commands (see Table 3.1) issued by the computer. Recall that the computer can control operations of any or all the CIC units (Followers) simply by placing one CIC (the Leader) in 'local mode'; 'Start' and 'Stop' commands issued to the Leader are passed on to the followers as 'Halt' = Off and 'Halt' = On respectively.

3.5.4 The Status Registers

Like the control registers, the status registers contain a variety of bit flags. These flags can be read by the computer to determine the conditions in effect on any CIC. Table 3.5 and Table 3.6 show the bit flags contained in the two registers.

Status Register 1

Initialization Status: The first three flags in Status 1 reflect the initialization status of the Control 1, Control 2, and Autocorrelation registers. If any of these registers has not been initialized by the computer, its status flag will be 'Off'.

bit position	flag name	On means	Off means
0	Timer 0	value 1	value 0
1	Timer 1	value 1	value 0
2	Multiple-Run	incomplete	results final
3	Data Overflow	results lost	results safe
4	Off	CIC off-line	normal
5	On	normal	CIC off-line
6	Off	CIC off-line	normal
7	Follow Mode	External Run	Local Run

Table 3.6: **Status 2 Flags:** These flags represent the conditions in effect at the end of the last correlation run.

Results Status: The 'Results Ready' flag 'On' indicates at least one byte of new correlation-run data hasn't been read.

Correlation Status: 'Data Flow', when 'On', indicates that a correlation run is in progress.

Sampling/Simulating Status: 'Sampling' and 'Simulating' together represent the type of testing being done using the diagnostic RAM feature. Four possible cases exist:

1. Sampling 'Off' and Simulating 'Off': No tests are being run or tests are finished.
2. Sampling 'Off' and Simulating 'On': The RAM contents are being sent to the correlators.
3. Sampling 'On' and Simulating 'Off': The RAM chips are being loaded with data values arriving from the A/D converters. This test mode can be in effect during normal correlation run processing; the first 4,096 samples of the correlation run are captured in the RAM.
4. Sampling 'On' and Simulating 'On': In this mode, only the 4,096 signal stream values (from the A/D units) that are captured in RAM

are allowed to proceed into the correlation matrix. Operations stop automatically after 4,096 samples cycles have elapsed.

Telescope Status: The final status bit in status register 1, 'External Halt', can be used to assess the current state of the telescope operations; if the instrument is in the middle of a correlation run this flag will be 'Off'. This information can be used to synchronize the continuum system with the rest of the telescope on start-up. If for some reason the continuum system goes 'off-line' its recovery will be more predictable if it reconnects during the pause between correlation-runs. Start-up will probably take 2 or 3 correlation run intervals to obtain (from another computer) and load the desired configuration parameters into the CIC units.

The contents of the Status 1 register represent the conditions in effect at the instant of the 'read transaction' (ignoring inherent hardware and software delays). This is NOT the case with the Status 2 register.

Status Register 2

The contents of Status 2 represent the conditions in effect at the END of the last correlation run. The information in this register is associated with the data (results) from the correlation run.

Tau 1 and Tau 2: These flags contain the value of the divide-by-4 scalar described in the Interval Timer discussion earlier. This 2-bit counter slows the count rate seen by the correlation-run interval timer from 20 to 5 MSPS. The exact interval value is made accessable to the computer by providing these 2 bits in the status register. These bits are cleared by 'high-clear', as is the 32-bit Interval Timer chip.

Multiple-Run: This flag indicates to the computer that the previous correlation run is one of several runs being accumulated in channel hardware. When this flag is true, the channel results do not need to be read-out.

Data Overflow: This is an error flag that is raised if the last correlation run ended before the results of the previous run had been read by the computer. This error will occur if the computer is too slow or if there is a fault in the cascaded 'enable read-out' loop (see Section 3.3 on page 81).

Test Pattern: The three flag-bits, Status 2:bits 4,5 and 6 values act as a test pattern; these bits are *hard-wired* 'Off', 'On', and 'Off' respectively. This allows the computer to verify the presence (or absence) of all CIC's in the system.

Follow Mode: The flag 'Follow Mode' is 'Off' if the last correlation run was done with the CIC in 'local mode' (as a Leader). During normal operations where control signals are cascaded from one CIC to the next, this flag will be 'On' for all CIC's.

Summary of Software Control Features

The CIC control features allow the computer to dictate the system configuration by setting control registers on each c d. The double-buffer feature on the control registers (as well as the auto correlation register) allows the controlling processor to set up the configuration for the next correlation run whether or not one is still in progress. Activities on each CIC can be directed by issuing commands (such as 'reset', 'start' or 'stop'). The CIC is sensitive to the command signal transitions; this prevents multiple commands caused by a slow transaction.

System reliability is improved by monitoring important conditions at each CIC. While the system is not less susceptible to faults (commercial components have been used throughout), errors will be detected sooner. This is to be contrasted with the old continuum correlator system where no automated monitor points are provided; system integrity is checked by visually assessing the quality of the night's observations. This new system will reduce the time taken to detect system faults.

3.5.5 Packaging the DSP System

Figure 3.14 shows a prototype of a correlation system chassis. The final system configuration with 4 bandwidth correlators requires 16 of these boxes mounted in two 19-inch racks. The card sites within it correspond to the rows of the system matrix. To make full use of available rack space, the triangular correlation matrix is folded back on itself as shown in Figure 3.15.

Each chassis houses two rows of the folded matrix. The horizontal row connections (including clock distribution, primary data-streams, the results bus, as well as power) are implemented on the chassis backplane, a four layer circuit board. Each chassis has two backplanes, one for each antenna. Cooling fans, as well as the power supply, occupy the 3 inch space directly below the card cage.

The four chassis are stacked on top of one another with vertical connections made using twisted-pair ribbon cable. The signal assignment on the cables provides generous shielding to enhance the integrity of the high-speed TTL data paths; each signal is flanked by ground lines (8 signal lines: 12 ground lines).

The backplane in this chassis differs from most in that there are active components on board (see Figure 3.16). Four of the devices are 4-bit differential line receivers for the digital samples accepted from the A/D converters; A/D signal cables connect at the back of the chassis on the outer ends of the backplanes.

Each of the 8 backplanes has an adjustable clock distribution system; processing in adjacent chassis can be synchronized to within 5 nanoseconds.

The backplane printed circuit board was designed for eight correlator cards. Seven of the eight backplanes in a 4-chassis matrix require fewer than eight correlator slots, so these backplane PCBs must be physically cut to fit the chassis requirements. Figure 3.17 shows a backplane in the foreground with one CIC slot and eight correlator card slots; the one in the background has only one correlator card slot (for auto correlation) beside its CIC slot.

When all four mapping bandwidths are considered, the entire digital sig-

Figure 3.14: **The Prototype System Chassis:** Cooling fans and the power supply occupy the lower part of the chassis (facing page).

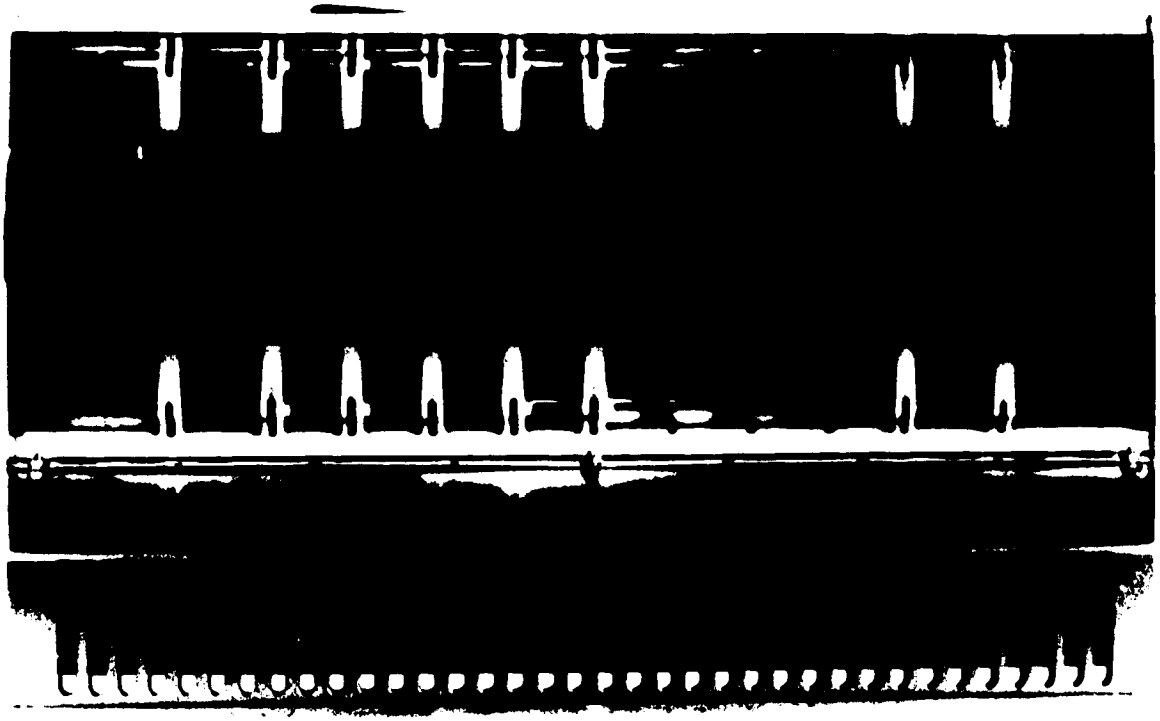

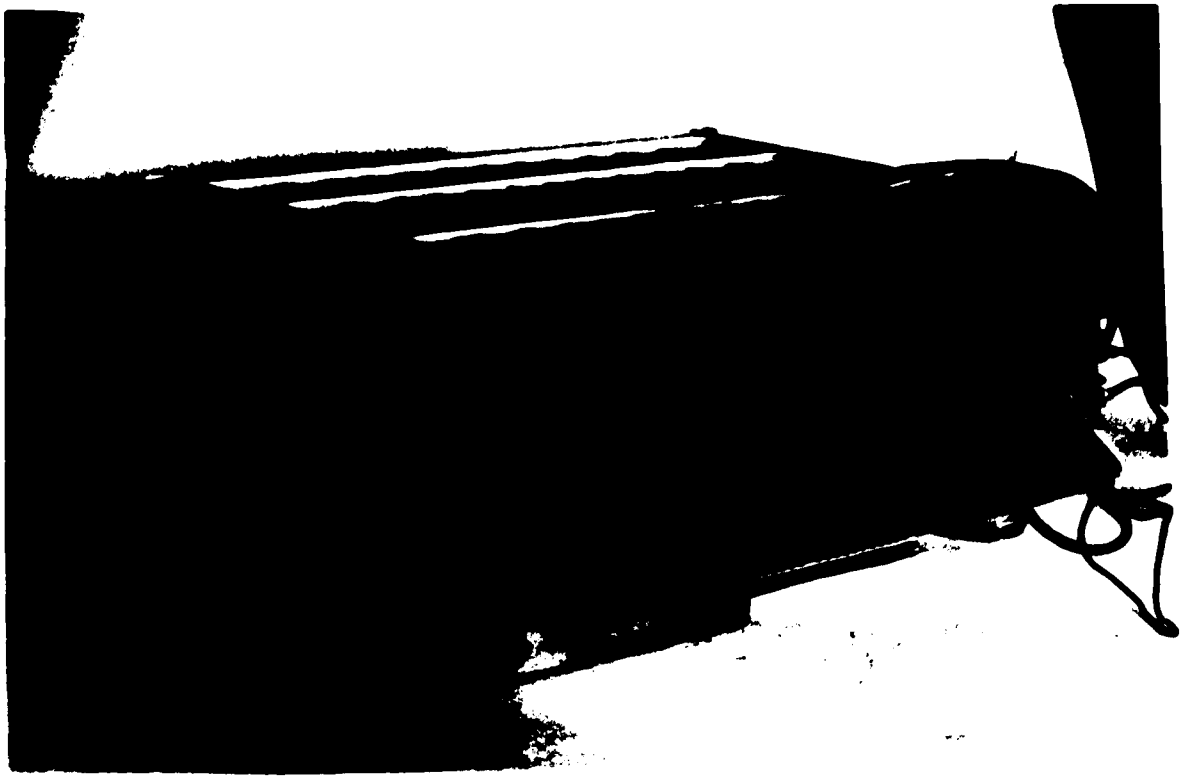

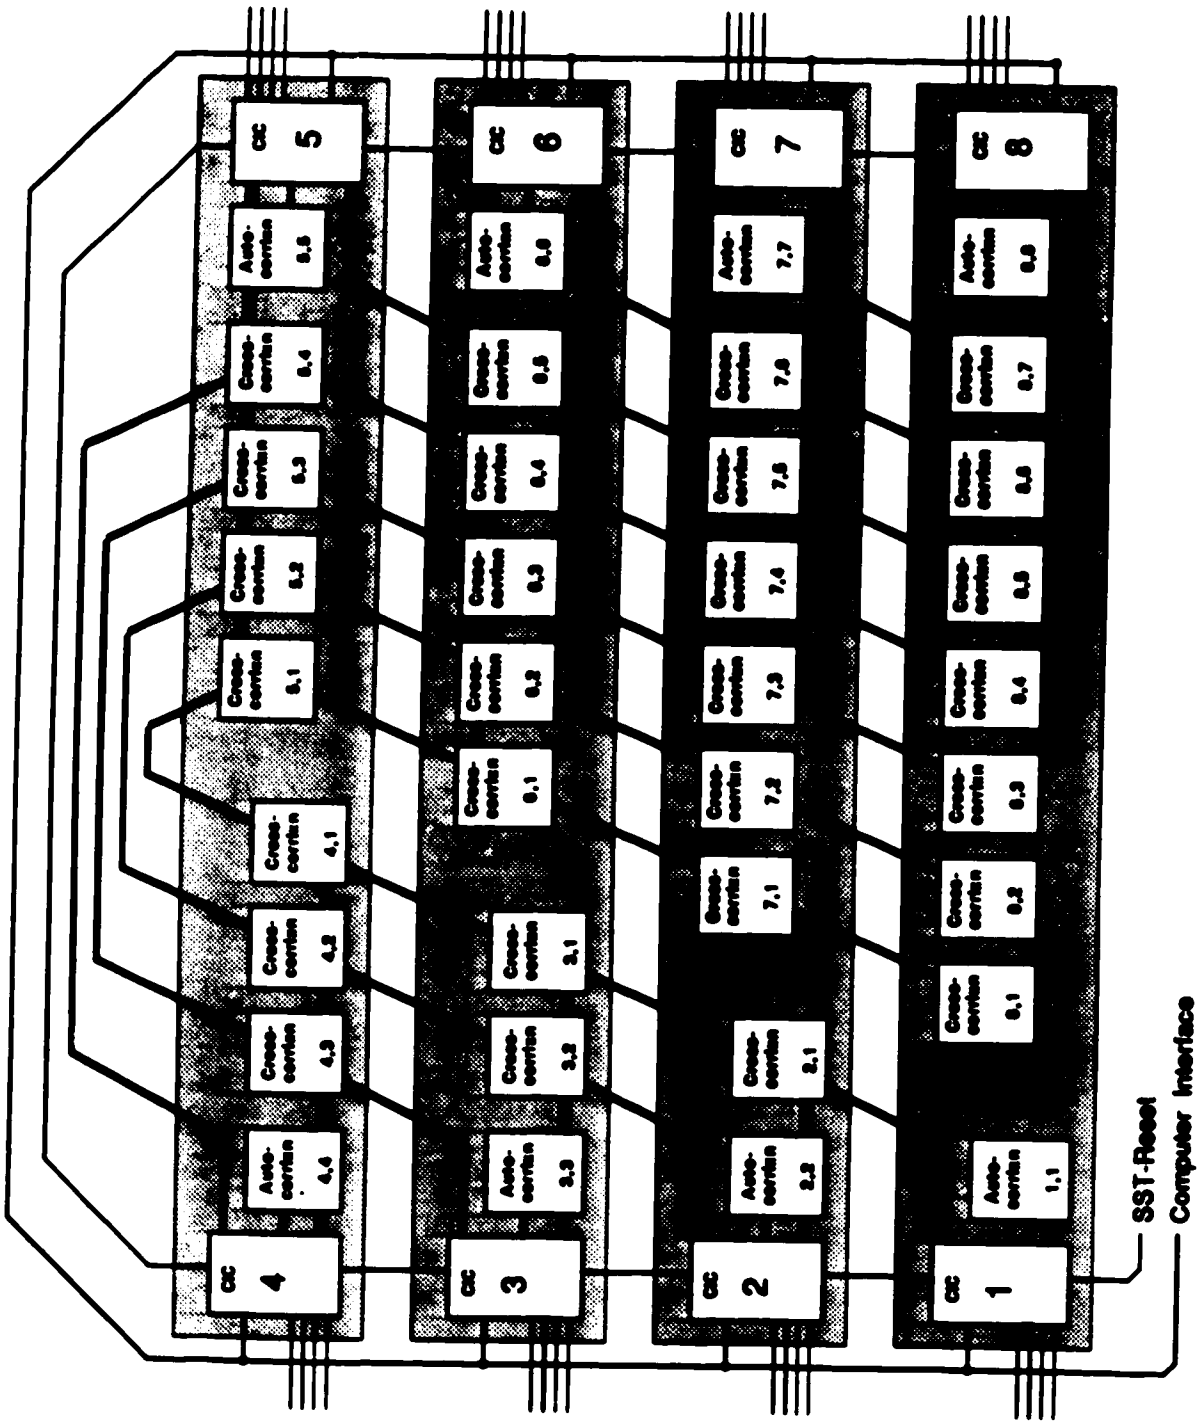

Figure 3.15: The Packaged Correlation Matrix:

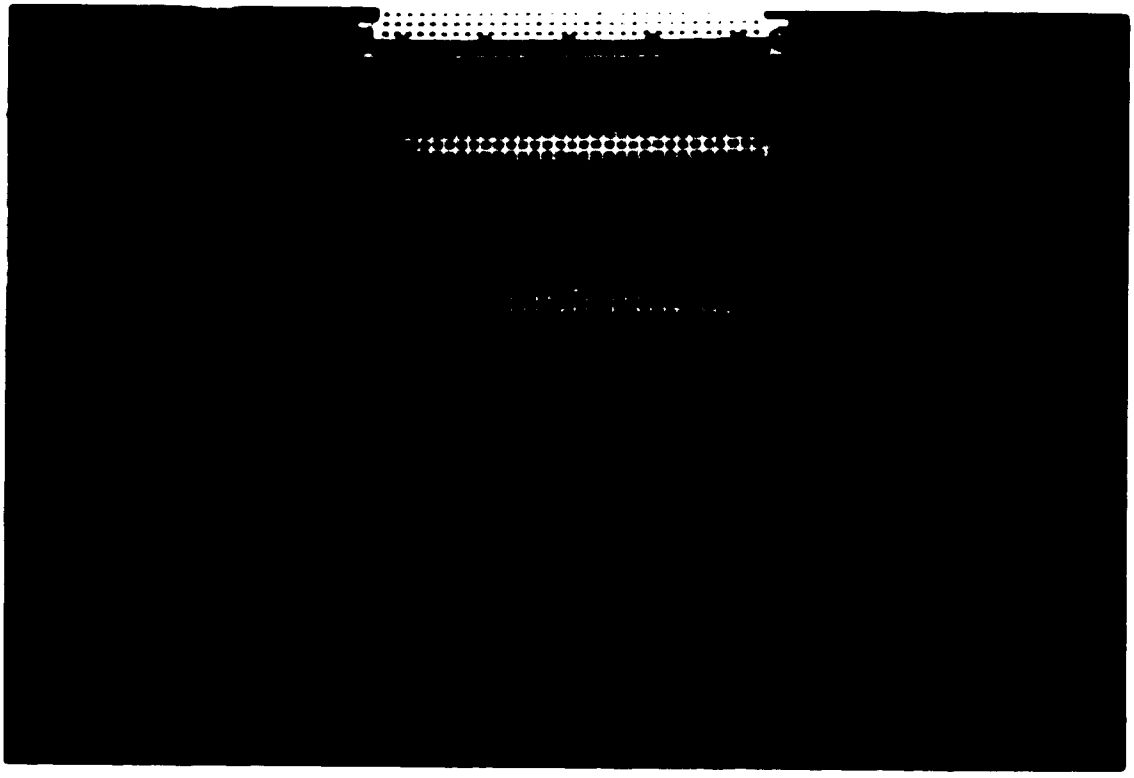

Figure 3.16: The Mother Board: The mounting holes double as power and ground rails. Power indicator LEDs and the ECL-to-TTL receivers for the A/D signals are in the foreground. Components on the right include a comparator, programmable delay unit and buffers for clock synchronisation and distribution.

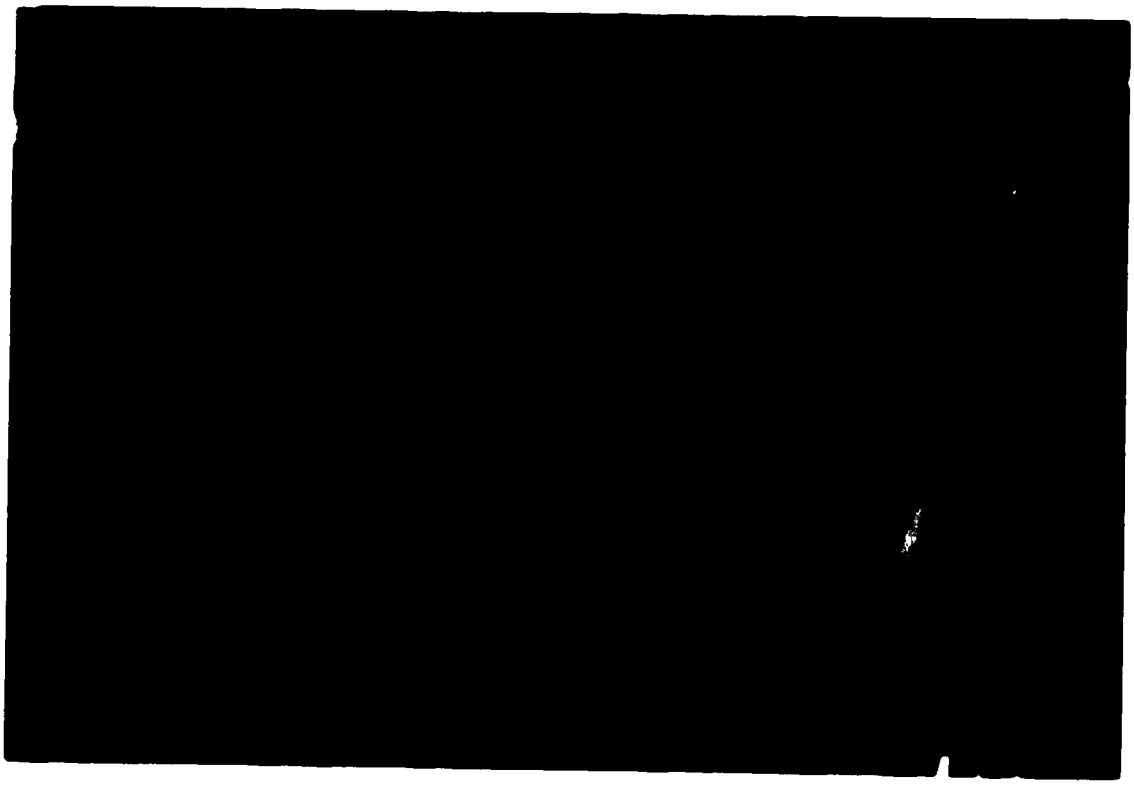

Figure 3.17: **The Baseline Configuration:** In the final system configuration, these would be the backplanes for antennas 8 (foreground) and 1 (background) in the lowest chassis of Figure 3.15. This correlator card is mounted in an autocorrelation site (for Antenna 8). A/D cables connect at the rear.

nal processing system for the continuum will occupy two 19 inch racks, each 8 feet high. Cables feeding into these racks will be:

- Controlling I rocessor 25-Wire Cables (2 or 4).
- Telescope Control Input ('SST Reset'): 1 twisted pair signal for controlling correlation intervals; this is the system-level 'Halt' signal.
- Antenna Signal Stream Inputs: 128 cables; 4-twisted-pairs (4 bits) per cable.
- System Clock Signals: 32 coaxial cables distributing AC-coupled square waves at 20 MSPS.
- AC Power Cables: One per chassis; 120 volts/3 amps.

A comparison of system densities is a crude, but valid, packaging yardstick. The old continuum system houses about 6 correlator channels in a cubic foot; the new system houses about 37 in the same space (this neglects to mention space required for A/D units and all the IF signal processing equipment used to prepare the signals for quantization).

3.6 Summary of the Correlation Control Design

This design (the control hardware) relieves the computer of all time-critical tasks; average throughput, not response time, is the only design driver for the software. Because all of the low-level tasks are handled by the control hardware (CICs), the computer has more time to perform high-level operations such as data limit checking, preprocessing of results or diagnostics.

Including the diagnostic capabilities in the system design was necessary to facilitate the unit-level development and testing. These features will also assist the system-level integration testing. For the system-level testing, the

control features have been designed to allow split-system operation (the preliminary software design takes this into account as well); part of the system can be used for off-line testing while the rest of the system is engaged in observations. This will aid in testing new antennas as they are brought on-line.

The modular design of the control system (and the correlators) will allow the new continuum system to grow as antennas are added.

Chapter 4

Signal Processing and Data Acquisition

This chapter will focus on the antenna signal processing that occurs upstream from the digital continuum correlator system. First, a high-level review of the existing telescope parameters as they pertain to the continuum system will be presented.

Next an overview of the new continuum IF signal processing system will lead to the main discussion regarding the analog-to-digital (A/D) converters. The development of the IF system[1] was being done at the same time as the A/D converter development.

The discussion on design is followed by a summary of the tests leading up to the Version 3 prototypes. Development of the A/D units suffered when the A/D device selected for the project was discontinued by the manufacturer (Advanced Micro Devices AM-6688 4-bit flash converter). Effort was spent *again* selecting and procuring a new device and designing its enclosure.

[1]The DRAO staff had the task of developing the RF and IF hardware required by the new continuum correlator — Ron Casorso was the principal contributor in this area.

4.1 Analog Signal Processing

Telescope Parameters

Recall that continuum signals (pertaining to the system described in this thesis) are received from space in a band 35 MHz wide centred on 1,420 MHz (the neutral hydrogen spectral line). The continuum system does not 'see' the 5 MHz band surrounding the neutral hydrogen line; this band is processed separately by the HI spectrometer system. At each of the telescope's antennas, the received signals are amplified, filtered (for image rejection) and mixed down to an IF frequency of 30.0 MHz before being sent down coaxial cables to a central processing building. Left and right circular polarization signals from each antenna are sent down separate cables. Once inside the building, antenna signals are fed into the cable-delay system (this system steers the 'delay beam' - see Chapter 1) to compensate for geometric delay in wavefront arrival times at each antenna. Each signal is fed through a delay unit consisting of a binary sequence of cable lengths; switches insert (or remove) cables in the signal path. More details about the cable delay system are presented in the paper by Landecker [Lan84].

Following delay compensation, the antenna signals enter the continuum IF signal processing system.

4.1.1 Continuum IF Signal Processing

Figure 4.1 shows the processing strategy being used to separate the signal into four different bandwidths; the four bandpass filters are each 7.5 MHz wide with RF center frequencies at 1,406.25, 1,413.75, 1,426.25 and 1,433.75 MHz. Each of these bands is mixed down to baseband (actually 0.5 - 8.0 MHz); the quadrature signal is produced during the final mix using a $\pi/2$ phase-shifted local oscillator signal.

The output of the IF signal processing is two signals (in-phase and quadrature-phase) for each sub-band in each polarization, resulting in 16 signals per antenna. Each of these must be quantized to 16 levels and sampled at a rate

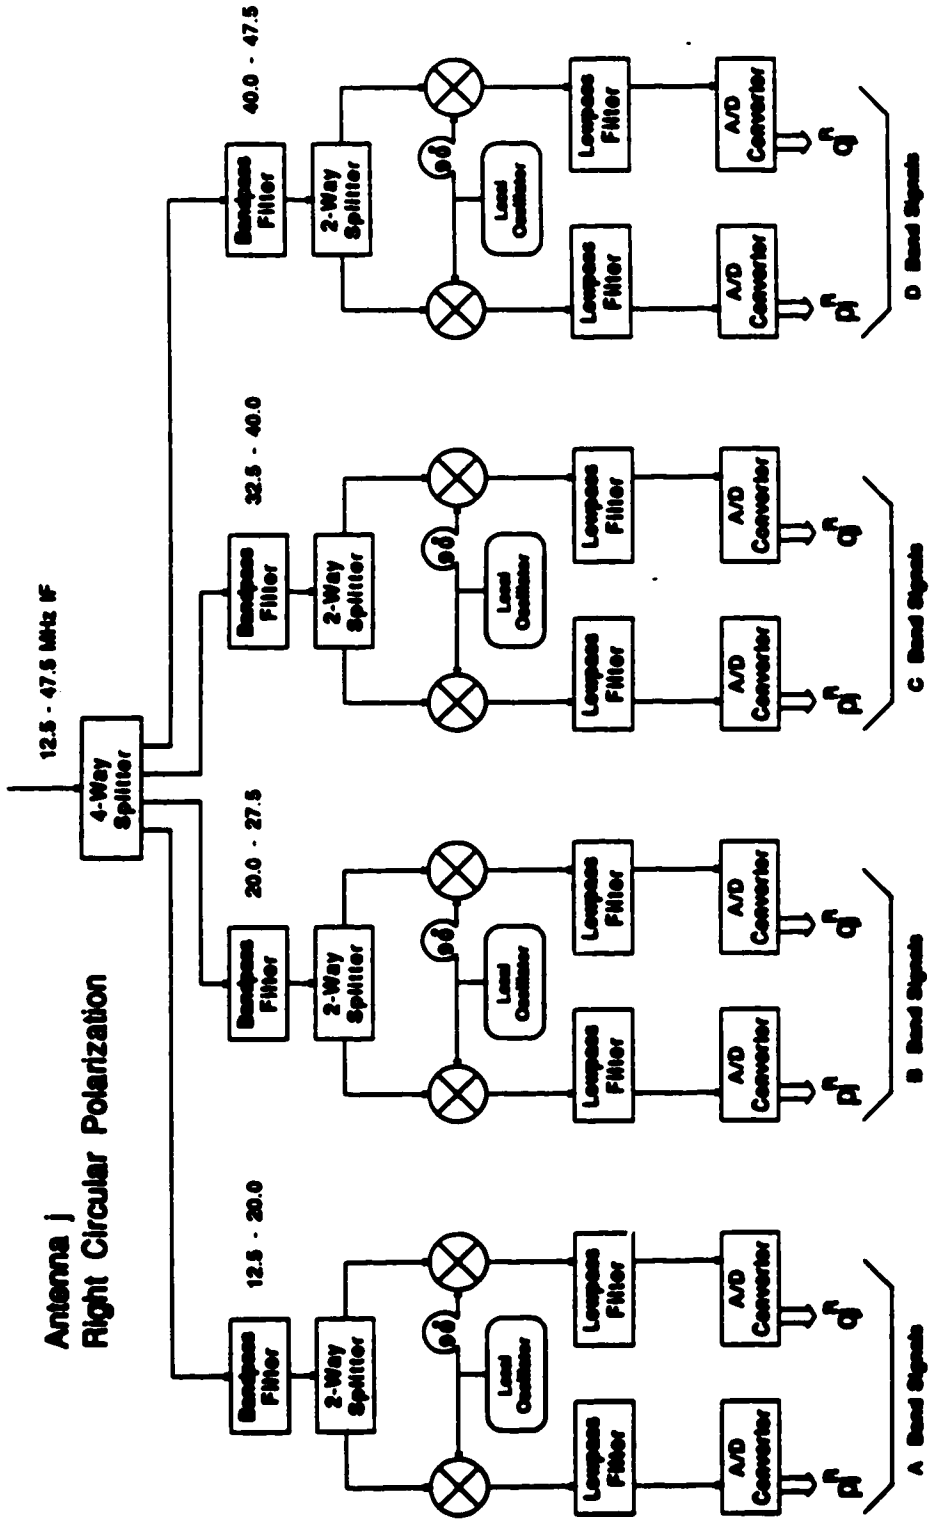

Figure 4.1: The Continuum IF Signal Chain:

exceeding 15 MSPS (the Nyquist rate for the 7.5 MHz sub-band).

4.2 Analog-to-Digital Conversion

The Version 3 prototype A/D units are based on commercial 4-bit flash converters from TRW. These units are functionally quite sim the design difficulty was to anticipate and minimize sources of *crosstalk* osstalk, in a correlation system, refers to any artificially induced commonality between signals. Weak coupling of digital signals onto two analog signal lines will show up as an erroneous correlation result.

Phase-switching eliminates the effect of stationary crosstalk components, i.e. those whose contribution to the correlation result does not change from one correlation run to the next. Examples of stationary crosstalk mechanisms are: d.c. signal bias (digital or analog); IF local oscillator leakage into the signal path (this includes the A/D converter sampling pulses).

The A/D unit design attempts to, first, minimize internal crosstalk and second, shield (from radiated EMI) and isolate (from conducted EMI) the A/D unit from the external, generally non-stationary, crosstalk sources.

Minimizing internal crosstalk requires protecting the sensitive analog circuits from the *noise* generated by the digital components within the A/D enclosure. This is done by encasing the different circuit groups in separate compartments; analog lines providing the signal and reference voltages to the A/D chip enter the digital compartment in close proximity to the device. Internal shielding is provided by the compartment walls and by extensive ground-plane on digital and analog PCBs. Noise conducted on d.c. lines (power, reference) is minimized using independent power supply voltage regulation, filter capacitors, filter feed-through's and ferrite beads. Figure 4.2 shows a block diagram of the circuitry of the A/D units.

The A/D enclosure consists of a central aluminum plate with aluminum boxes mounted on either side (the boxes are cut from extruded tubing and are held to the central plate by their lids). Figure 4.3 shows a photograph of the two Version 3 A/D units. Shielding from external radiative crosstalk sources

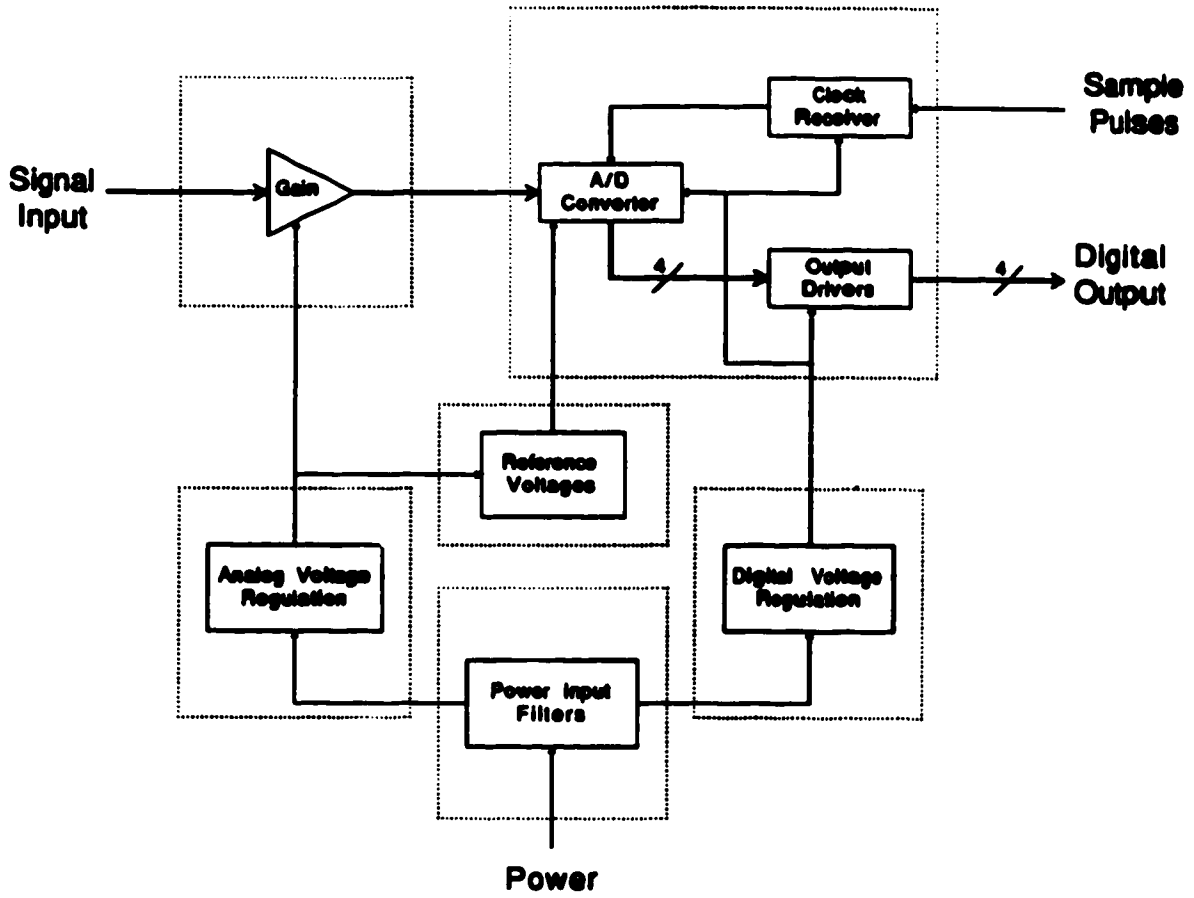

Figure 4.2: The A/D Circuit Groups: The dotted lines indicate compartment boundaries.

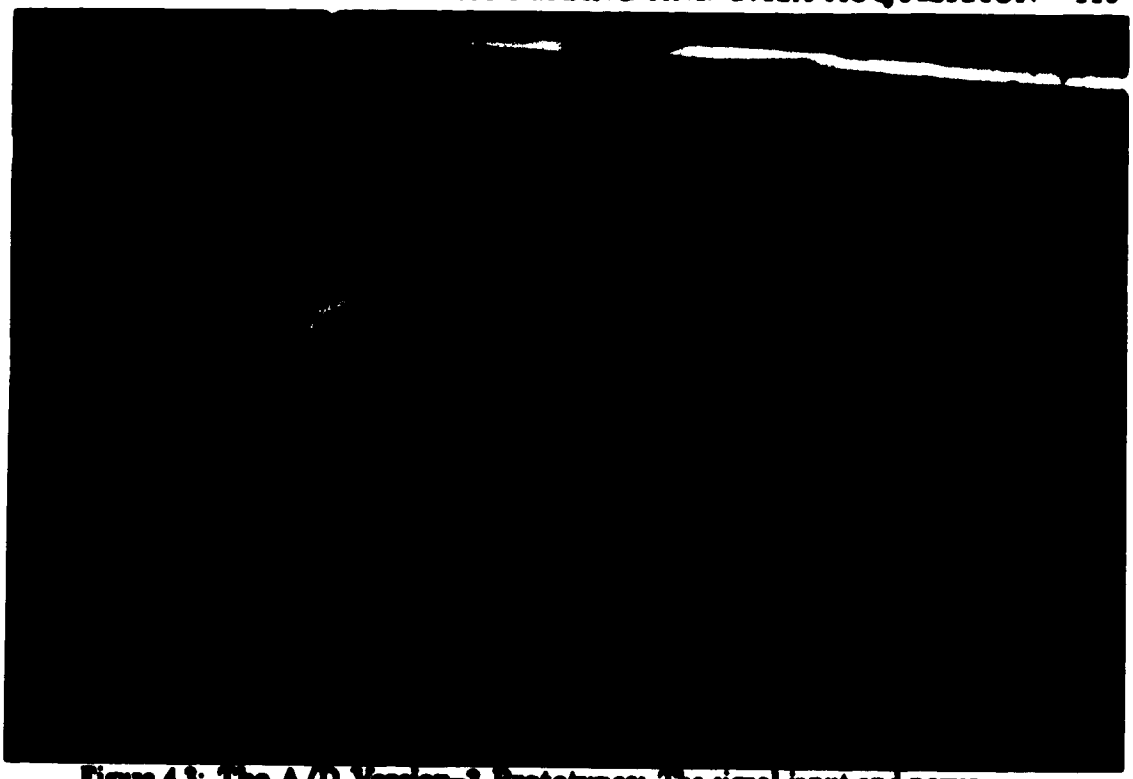

Figure 4.3: The A/D Version–3 Prototypes: The signal input and power connectors are on one end of the unit; digital output and the sample signal input are at the other end.

is accomplished by the enclosure. Filtering and internal voltage regulation on the power distribution lines is aimed at reducing conducted crosstalk entering through the power supply lines.

The design aimed at simplifying considerations involved with operating and maintaining the 128 units. The Version–3 units are compact and stackable: connectors are accessible at the ends of the enclosure. To power a unit, two voltages (± 15 volts) are required (as well as 'return', 0 volts). The clock input is designed to operate with any periodic input signal having a peak-to-peak amplitude greater than 50 mV; square-wave signals provide the best synchronization stability.

4.2.1 Thermal Considerations

The circuits in the A/D unit (operating in closed cells) must rely on conductive cooling mechanisms; conduction of heat to the aluminum chassis is the only means of cooling. While these units operated for extended periods (12 hours at a time) with no apparent damage, attention to thermal design is needed to ensure the components will survive (for more than a year) the conditions inside the compartments.

Another thermal consideration concerns variation in device performance at different temperatures. If internal temperature fluctuations are not damped, they will be a source of nonstationary crosstalk. To avoid this, the thermal time constant of the A/D units should be large. With good conductive heat transfer from components to case, the thermal mass of the aluminum enclosure will stabilize thermal fluctuations.

Thermal potting is probably the simplest way of providing good conductive cooling (compartments are filled with a thermally conductive compound). The disadvantage of thermal potting is that maintainance of the A/D unit will be messy (if thermal grease is used) or impossible (if a solid thermal epoxy is used).

4.2.2 Performance Testing

The development of the A/D prototypes began with very crude (Version-0) lash-up units. As development progressed, each version's performance improved. With the advent of Version-2 (designed around AMD's 6688), a correlation system was required to assess the design's immunity to crosstalk. A prototype digital correlator was built to verify the channel design as well as to test the A/D performance; the prototype digital board also provides data value counters (16 per A/D unit) permitting analysis of the digital probability distribution function.

The crosstalk test configuration is shown in Figure 4.4 Two independent noise sources produce signals with no similarity; each signal is sampled by the A/D's and correlated. Non-zero correlation results reveal the magnitude

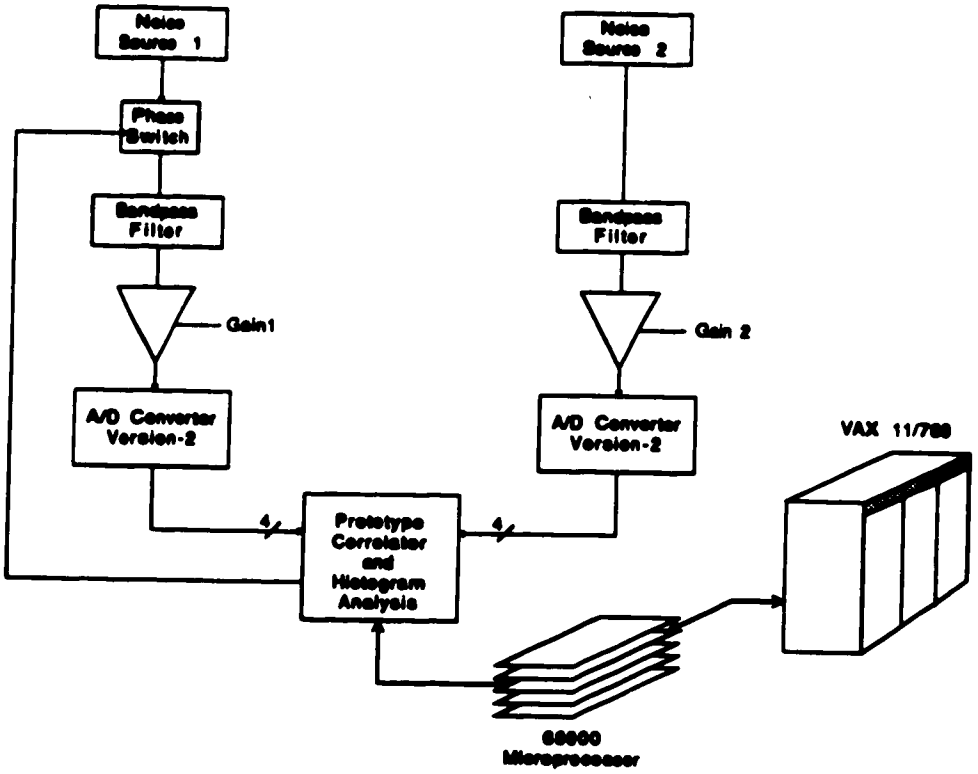

Figure 4.4: The Crosstalk Test Configuration: The testing was automated using the microprocessor. Correlation runs were performed with phase-switching (inversion) on signal 1 using a double balanced mixer as the signal modulator.

of the crosstalk in the test configuration. Even after phase switching a detectable crosstalk component remained; detected crosstalk levels were in the order of 55 db below the signal levels.

Ideally, much more investigation of the crosstalk mechanism(s) should have been done before pursuing further A/D unit design; this investigation would have included a review of the channel design and error analysis of the software computations. Unfortunately, the development schedule was interrupted by AMD's cancellation of V2's A/D chip. It was decided that rather than continue testing, an alternate device should be chosen and that the known V2 design problems should be fixed in Version-3. Several factors pushed this decision:

- Version–2 had several known problems: digital noise on the analog lines; impractical packaging design (not manufacturable); poor thermal characteristics.

- A replacement was needed for the AMD 6688. Following the design of V3, parts would have to be ordered and built; the two units would have to be assembled and tested at a preliminary level. In light of the development still required on the digital system, much of this work could be transferred immediately to DRAO's staff.

- The performance of the prototype digital correlator channel was also suspect. The observed crosstalk may have been caused by truncation of low-order channel bits (this channel design did not 'round' the 32-bit result). As well, the channel design may have been faulty; in the absence of PAL design software, the fuse-map was done by hand (no design verification was possible). At the time there was no means of performing exhaustive algorithmic tests on the digital hardware (this reinforced the need for the diagnostic RAM and the Interval Timer discussed earlier).

The main difference between V2 (one PCB) and V3 is the separation of the circuit groups into shielded compartments. A further benefit of the separation is that, in the event of yet another device change (forced or voluntary), less of the design will require revision. Only circuit groups impacted directly by the change will need rework. It was expected that the improved isolation and filtering would improve the crosstalk immunity of the converter units. The aluminum enclosure of V3 was selected to improve thermal stability.

4.3 A/D Development Summary

Using the alternate chip (TRW's AD–1044), the V3 A/D units had passed preliminary testing when they were handed over to the DRAO for completion. The [illegible] requiring further effort were:

- The thermal design and testing mentioned above.

- The analog input amplification circuit (no work had been done here).
- The redesign of the digital output drivers. The initial V3 design made use of a TTL device for driving the digital outputs. Because ECL line drivers give better performance and internal noise characteristics than TTL, the mother-board receivers were designed for ECL. The A/D output drivers need to be modified to comply with this interface.

Because the digital signal processing system is now functional (and verified by testing) the development testing of the A/D units will be quicker and more thorough; for example, knowing the correlation interval precisely eliminates the possibility of introducing computational bias during offset removal. As well, the upstream diagnostic features will simplify the adjustment of A/D performance parameters (reference voltages and offsets).

Aside from the above considerations, the V3 prototype represents a robust, servicable and economical design.

Chapter 5

Project Summary

This chapter will provide an overview of the project's engineering development phases. There were three design stages required to achieve the system presented in the preceding chapters.

5.1 The Development Stages

Both aspects of this project, the analog to digital signal conversion and the digital signal processing, went through three distinct stages in their development.

5.1.1 Version 1: the Lash-Ups

The Version 1 units were the initial lash-ups built during the initial three months of the project; these were developed as quickly as possible to help identify potential problems.

While the V1 A/D converter operated correctly, there were obvious problems with the design. The noise on the analog input and reference voltages had to be reduced by providing better shielding and isolation from noise

sources (internal and external). It was decided that local power regulation and better input/output signal buffering was required.

The V1 correlator was constructed using standard SSI and MSI parts such as 74F' series TTL and Schottky 2×4-bit multipliers. The channel required 23 components and consumed 7 watts; its limit of reliable operation was only 6 MSPS. The main design problem was in the distribution of the high-speed signal lines, particularly the clock; fan-out and branching would have to be minimized. As well, noise on the power and ground lines indicated the need for filter capacitors and low-inductance power distribution.

5.1.2 Version 2: the Bench Models

Based on the insights gained from Version 1, the Version 2 units were designed. Considerable effort was invested to eliminate the problems seen in V1 and to improve both performance and efficiency. These were to be the final-form prototypes for the proof-of-concept testing.

The V2 A/D units were built on a double-sided PCB. All the power-supply and reference voltages were provided by on-board circuits; this reduced the noise to small, but still visible, levels. The units demonstrated proper performance characteristics at speeds greater than 50 MSPS.

The V2 correlator showed significant improvements over the V1 design: the four 24-pin MSI multiplier chips were replaced with a 20-pin look-up table multiplier PROM; the eight TTL accumulator chips were replaced by a 32-bit LSI counter with built-in storage. Channel performance went up by a factor of four; power consumption and board area went down by 60% and 80% .

The V2 units were suitable for the cross-talk testing shown in Figure 4.4. Even though this test activity was halted prematurely it served to provide valuable ideas, some of which were incorporated in the next generation designs. A significant outcome of this activity was the need for greater control of the test conditions. The facilities required for unit-level and system-level performance verification were not available; they would have to be built. The

best way to do this was to build the diagnostic features into the control system for the correlator. This resulted in features like the Diagnostic RAM and the Processing Interval Timer. Unfortunately, the duration of this project was too short to reap the full benefits of these considerations.

5.1.3 Version 3: the Production Models

The step to Version 3 required attention to many more aspects than were required for the bench model. These units faced the following considerations:

- Packaging Design: partitioning the design into serviecable modules; designing units to meet manufacturing and space constraints; providing maintenance access and easy expansion.
- Service Design: power supplies and cooling.
- Interconnect Design: A/D-to-DSP interface; data distribution network; chassis connector selection and pin allocation; chassis level and system level interface design; system synchronization.
- Control Design: determining system operation characteristics; control distribution; synchronization to telescope control.

Another large activity was the procurement of parts for the system and the contracting of PCB artwork and manufacturing.

The Version 3 implementation transformed the 'bench' concepts into a usable, manufacturable and testable system. Many of the system-level design decisions were based on the results of V1 and V2 testing as well as on techniques used on other systems (HI, 408 MHz, and the 26-meter dish) at the DRAO.

5.2 The Present Status

The digital signal processing system for the new continuum facility is in the final stages of manufacturing (contracted out). Units for 5 antennas are required in the immediate future (20 control/interface cards, 16 chassis units and 50 correlator cards are needed).

The V3 A/D unit is being modified to accept a better and cheaper A/D chip; the Analog Devices AD-9688 is a pin-compatible upgrade of the original device, AMD-6688, selected (then discontinued) for V2.

The IF processing hardware is complete and awaiting system-level integration testing. The continuum system software is in its preliminary design and testing stage.

With the successful test of the digital design (see Chapters 2 and 3), accomplished near the end of this project, the DRAO acquired a system that will (i) efficiently meet the continuum processing requirements for the foreseeable future and, (ii), support the next few years of the telescope's development and testing[1]. After the continuum development activities have ceased, the system will support operations and maintenance with automated diagnostics and performance monitoring.

[1] I want to be there!

Bibliography

[Gom84] H. Gomaa. A software design method for real-time systems. *Communications of the ACM*, 27(9):938–949, September 1984.

[Jan33] Karl G. Jansky. Electrical disturbances apparently of extraterrestrial origin. *Proceedings of the Institute of Radio Engineers*, 21:1387–1398, 1933.

[KB74] Rolf J. Klingler and F. K. Bowers. Quantization noise in correlation spectrometers. *Astron. Astrophys. Suppl.*, 15:373–380, 1974.

[Lan84] T. L. Landecker. A coaxial cable delay system for a synthesis radio telescope. *IEEE Transactions on Instrumentation and Measurements*, IM–33:78–83, 1984.

[Lo84] Wing F. Lo. *A Digital Correlator for an Aperture Synthesis Telescope*. Master's thesis, University of Alberta, Edmonton, Alberta, 1984.

[TMS86] A. Richard Thompson, James M. Moran, and George W. Swenson Jr. *Interferometry and Synthesis in Radio Astronomy*. John Wiley & Sons, Inc., New York, 1986.

[Vei84] Bruce G. Veidt. *A 408 MHz Synthesis Radio Telescope*. Master's thesis, University of Alberta, Edmonton, Alberta, 1984.

[Wei63] Sander Weinreb. *A Digital Spectral Analysis Technique and Its Application to Radio Astronomy, Technical Report No. 412*. PhD thesis, Research Lab. for Electronics, Massachusetts Institute of Technology, Cambridge, Mass., 1963.