REPORT OF THE WORKSHOP ON THE AUSTRALIA TELESCOPE CORRELATOR SYSTEM

AT/10.4/003

Held at Parkes - 21,22,23 November 1983

APPENDIX A FFT METHODS TO INCREASE THE NUMBER OF CORRELATOR CHANNELS

CHAPTER 1

INTRODUCTION

This report is the result of a workshop held in Parkes 21st-23rd November 1983. Those attending were;

J.D. O'Sullivan (Chairman), J.G. Ables, J.W. Brooks, D.R. Brown, D.N. Cooper, J.R. Forster, R.H. Frater, A. Henderson, A.J. Hunt, C.E. Jacka, D. le Comte, A.G. Little, R.N. Manchester, M.W. Willing, W.E. Wilson.

The correlator system had been covered in some depth at the Computer Workshop prior to the Correlator Workshop so that a system was on paper. This system was however in urgent need of detailed appraisal on the one hand and alternatives to certain critical areas on the other hand.

We accordingly began the Correlator Workshop with the following general goals in mind;

1. To determine a reasonable specification of the correlation system including fall back options and requirements for day one or earlier operation.
2. To identify any problem areas requiring immediate attention.
3. To provide conceptual designs for various correlator sub-systems.

Due to the magnitude of the task, it was not considered relevant to cover all aspects of the correlator system. Some parts of the system have been quite adequately examined during the Computer Workshop (for the time being). Attention was concentrated on what were seen as major problem areas. In particular, the two areas which received the most attention were the overall correlator structure and the elimination of large scale selection or switching matrices, and the correlator IC, the XCELL.

CHAPTER 2

SCIENTIFIC REQUIREMENTS

The Australia Telescope is conceived as an instrument especially suited to high quality observations over a wide field of view, and also as a spectral line instrument. Furthermore, the ability to perform various types of observation simultaneously is an important aspect of its design. Towards this end, dual-band feeds, independant local oscillators, four IF channels, wide bandwidths and both compact and long baseline arrays are part of the design.

In order to make full use of these features, the backend must be of adequate size and versatility to cope with a wide range of input combinations. A general list of astronomical requirements for the AT correlator is the following (cf. AT/10.4/001, AT/05.4/001):

1. 1 or 2-bit operation
2. 160 MHz (1-bit) or 80 MHz (2-bit) maximum bandwidth
3. Enough modules to:
   - (i) avoid bandwidth smearing for continuum observations at full bandwidth,
   - (ii) provide good velocity coverage and resolution for spectral line observations,
   - (iii) allow simultaneous observation of all polarisation/frequency IF combinations,
   - (iv) allow simultaneous observation of both compact and long baseline arrays,
   - (v) allow simultaneous observation in both line and continuum modes.
4. Enough flexibility to:
   - (i) readout at 1 ms rate,
   - (ii) integrate to 100 s,
   - (iii) process bandwidths ranging from 160 MHz to 0.625 MHz, with a corresponding increase in the number of derived spectral channels (1 to 64),
   - (iv) make maximal use of the available correlation channels for various types of observation.
5. Provide autocorrelation spectra when desired

6. Allow for future expansion.

These requirements have been extensively covered in earlier documents (AT/10.4/001, AT/05.4/001). The fundamental change in the basic correlator structure proposed at this workshop has significantly reduced the complexity of the hardware. It is the aim of this section to examine the impact of these and other changes on the achievement of the astronomical objectives.

2.1 WIDE BANDWIDTH CONTINUUM REQUIREMENTS

For continuum observations, a kind of chromatic aberration limits the undistorted field of view available for a given bandwidth. A fairly conservative specification (~50% smearing near the first null of the primary beam) sets a limit on the percentage bandwidth desirable for wide field continuum imaging at 0.0018. It has been shown (AT/05.4/001) that L-band is the worst case for the AT, and that about 64 spectral channels are required to split a 160 MHz bandwidth into the requisite 2.5 MHz channels.

For the 6 telescopes (15 baselines) of the compact array alone, we require 15x4x64 = 3840 correlator channels to accurately image all 4 Stoke's parameters with an input bandwidth of 160 MHz.

Continuum bandwidths are limited by the restricted data rates available on the radio link for the long baseline array. At present a bandwidth of 10 MHz per IF channel for 1-bit correlation (20 Mbit/s) should be allowed for. The field of view is limited by the lack of planar sampling of the u,v plane (~11 arc sec at L band) rather than bandwidth smearing).

A new suggestion is that the continuum requirements of the LBA could be adequately met by a special purpose continuum correlator section (based of course on identical hardware but rewired). The equivalent of 6 baselines x 4 polarisation channels x 8 frequency channels at 10 MHz would be the minimal requirement.

2.2 SPECTRAL LINE OBSERVATION REQUIREMENTS

The number of channels required for wide band continuum is adequate for most spectral line purposes provided recirculation can be used. Recirculation rates from 1 to 64 are desirable.

In this chapter the term recirculation is used to represent the ratio of operating bandwidth to maximum bandwidth of the correlator and the consequent increase in number of measured channels. In chapter 5, the term utilisation factor is

introduced for the combined effects of the actual recirculation and reconfiguration methods employed by the hardware.

Two additional features should be available for spectral line observing;

1. A sensitive (and thus relatively wideband - 40 MHz) continuum measurement would be essential when observing narrow band molecular lines.
2. A sensitive measure of total spectral flux will be essential when operating with large recirculation factors. The measurement of large numbers of channels with the aid of recirculation will require data selection to allow further processing in the on-line and off-line computer systems.

The first point suggests provision of a separate dedicated set of continuum correlator channels. The second point was discussed at some length with no completely satisfactory solution emerging. Autocorrelation spectra such as measured with a single dish spectrometer would prove difficult to coax to sufficient sensitivity. An alternative might be crosscorrelation spectra obtained from short baselines which would not suffer from severe calibration problems. The clearest solution would appear to be the use of the Parkes telescope in such situations. No specific plans were made to include autocorrelation channels in the correlator specification.

2.2.1 Velocity Coverage And Resolution

The correlator specifications as set out in AT/10.4/001 required 128 frequency channels for 2 polarisation channels per baseline. The total number of channels is 3840 which becomes 245,760 channels when recirculated by a factor 64. A minimum requirement would perhaps be a factor four smaller to give 960 channels.

For comparison, the VLA correlator system boasts 11,772 channels with a maximum recirculation factor of 32. The maximum bandwidth of the VLA correlator system is however 40 MHz, as opposed to the proposed 160 MHz of the AT, and 351 baselines are required in place of 15.

The origin of these numbers can be gauged from the following table showing the total velocity coverage and resolution required for various classes of astronomical objects.

Object	Velocity (km/s)	Resolution (km/s)	Channels
Masers	100	0.1	1024
Galactic Centre	500	1.0	512
Molecular Clouds	100	0.4	256
Dark Clouds	30	0.1	256
Galaxies	1000	4.0	256
Recombination Lines	200	1.5	128

The recirculation technique provides large numbers of channels at low bandwidths. The critical areas are thus high velocity coverage and/or high centre frequency observations.

Velocity range for various bandwidths and observing frequencies.

Bandwidth (MHz)		160	80	40	20	10	5	2.5	1.25	0.63	0.31	0.16
Recirc. Factor		1	2	4	8	16	32	64	64	64	64	64
Channels/ Module		16	32	64	128	256	512	1024	1024	1024	1024	1024
Band	Freq. (GHz)	Velocity (km/s)										
P	.408					7.3K	3.7K	1.8K	919	460	230	115
L	1.5			8K	4K	2K	1K	500	250	125	63	32
S	2.3		10.4K	5.2K	2.6K	1.3K	562	326	163	82	41	20
C	5	9.6K	4.8K	2.4K	1.2K	600	300	150	75	38	19	9
X	10	4.8K	2.4K	1.2K	600	300	150	75	38	19	9	4
K	22	2.2K	1.1K	545	273	136	68	34	17	9	4	2
Q	44	1.1K	545	273	136	68	34	17	9	4	2	1
W	90	533	267	133	67	33	17	8	4	2	1	0.5
F	115	417	209	104	52	26	13	7	3	2	1	0.5

This table shows the velocity coverage of a module. In order to obtain the velocity resolution of a module, divide the table entry by the number of channels per module given at the top of the column. A module is further explained in Chapter 3,

but for the present it suffices to realise that a module is the smallest unit of the correlator selectable for line purposes. Each baseline will consist of 8 modules which may be connected in cascade to provide up to 8x16=128 channels at maximum bandwidth or 8192 channels at maximum recirculation factor.

A modul is velocity range limited beyond 44 GHz if 1000 km/s is the requirement. A single module is resolution limited when a maximum velocity range of 100 km/s is required with 0.1 km/s resolution at 10 GHz and beyond, ie. a module cannot provide 1000 channels with 100 Km/s range beyond X band. Cascading all eight modules for a baseline would ease the limitation to 44 GHz and beyond.

2.2.2 IF Pair Combinations Per Baseline

Each telescope will offer four IF return signals A,B,C,D. While it is tempting to imagine that full switching of each of these signals to the modules comprising a single baseline is required, carefull analysis of the sensible possibilities reveals otherwise.

A possible receiver system might consist of a high and low band feed horn with dual polarisation, a low frequency channel again with dual polarisation and possible narrow bandwidth filtered channels. It must be remembered that neither different centre frequency nor different bandwidth IF signals need ever be involved in a product.

Given that full selection possibilities exist at the receiver to connect A,B,C or D to any receiver IF signal, then the following modes can be identified;

Mode	Frequencies (or Bandwidths)	Combinations	Products	Use **
1	1	AA	1*	5
2	1,2	AA BB	2*	4
3	2,3	AA BB CC	3	7
4	2,3,4	AA BB CC DD	4	3
5	1	AA AB BA BB	4*	1
6	2	AA AB BA BB CC	5	8
7	2,3	AA AB BA BB CC DD	6	6
8	2	AA AB BA BB CC CD DC DD	8	2

Notes: * modes used with LBA,
** priority of useage, 1 corresponds to highest useage.

CHAPTER 3

CORRELATOR STRUCTURE

The correlator structure proposed in AT/10.5/002 included a set of 240 modules, each module being capable of producing 16 frequency channels for one baseline in their basic full bandwidth mode. These modules were connected to the various interferometers of the compact and long baseline arrays by a switching matrix. In its most general form, this switching matrix was required to connect each input of each module to any of the 38 possible I.F. input streams. This problem was investigated in AT/24.1/005 and 006. All indications were that the provision of such a general purpose switching matrix presented a major problem. This workshop proposes a radically different structure, which has developed from a requirement to simplify the pre-correlator switching.

The structure is composed of four distinct correlators, these being :

1. The Compact Array Line (or Wideband Continuum) Correlator - CALC.
2. The Long Baseline Array Line Correlator - LBALC.
3. The Compact Array (Narrowband) Continuum Correlator - CACC.
4. The Long Baseline Array Continuum Correlator - LBACC.

The correlator is now clearly divided into two sections, a telescope based section and an interferometer based section. The telescope based section includes delay tracking, recirculation memories and digital filtering for example, while the actual correlator section is interferometer or baseline based.

Apart from the necessary infra-structure such as control sections, the correlator system should be highly modular. A single interferometer should hopefully only require the sections appropriate to two telescopes and a single baseline.

The entire system is shown in figure 3.1. Shown on the diagram but not discussed during the workshop is a demodulator section. This section is intended to demodulate any phase switching introduced at the telescopes for better rejection of

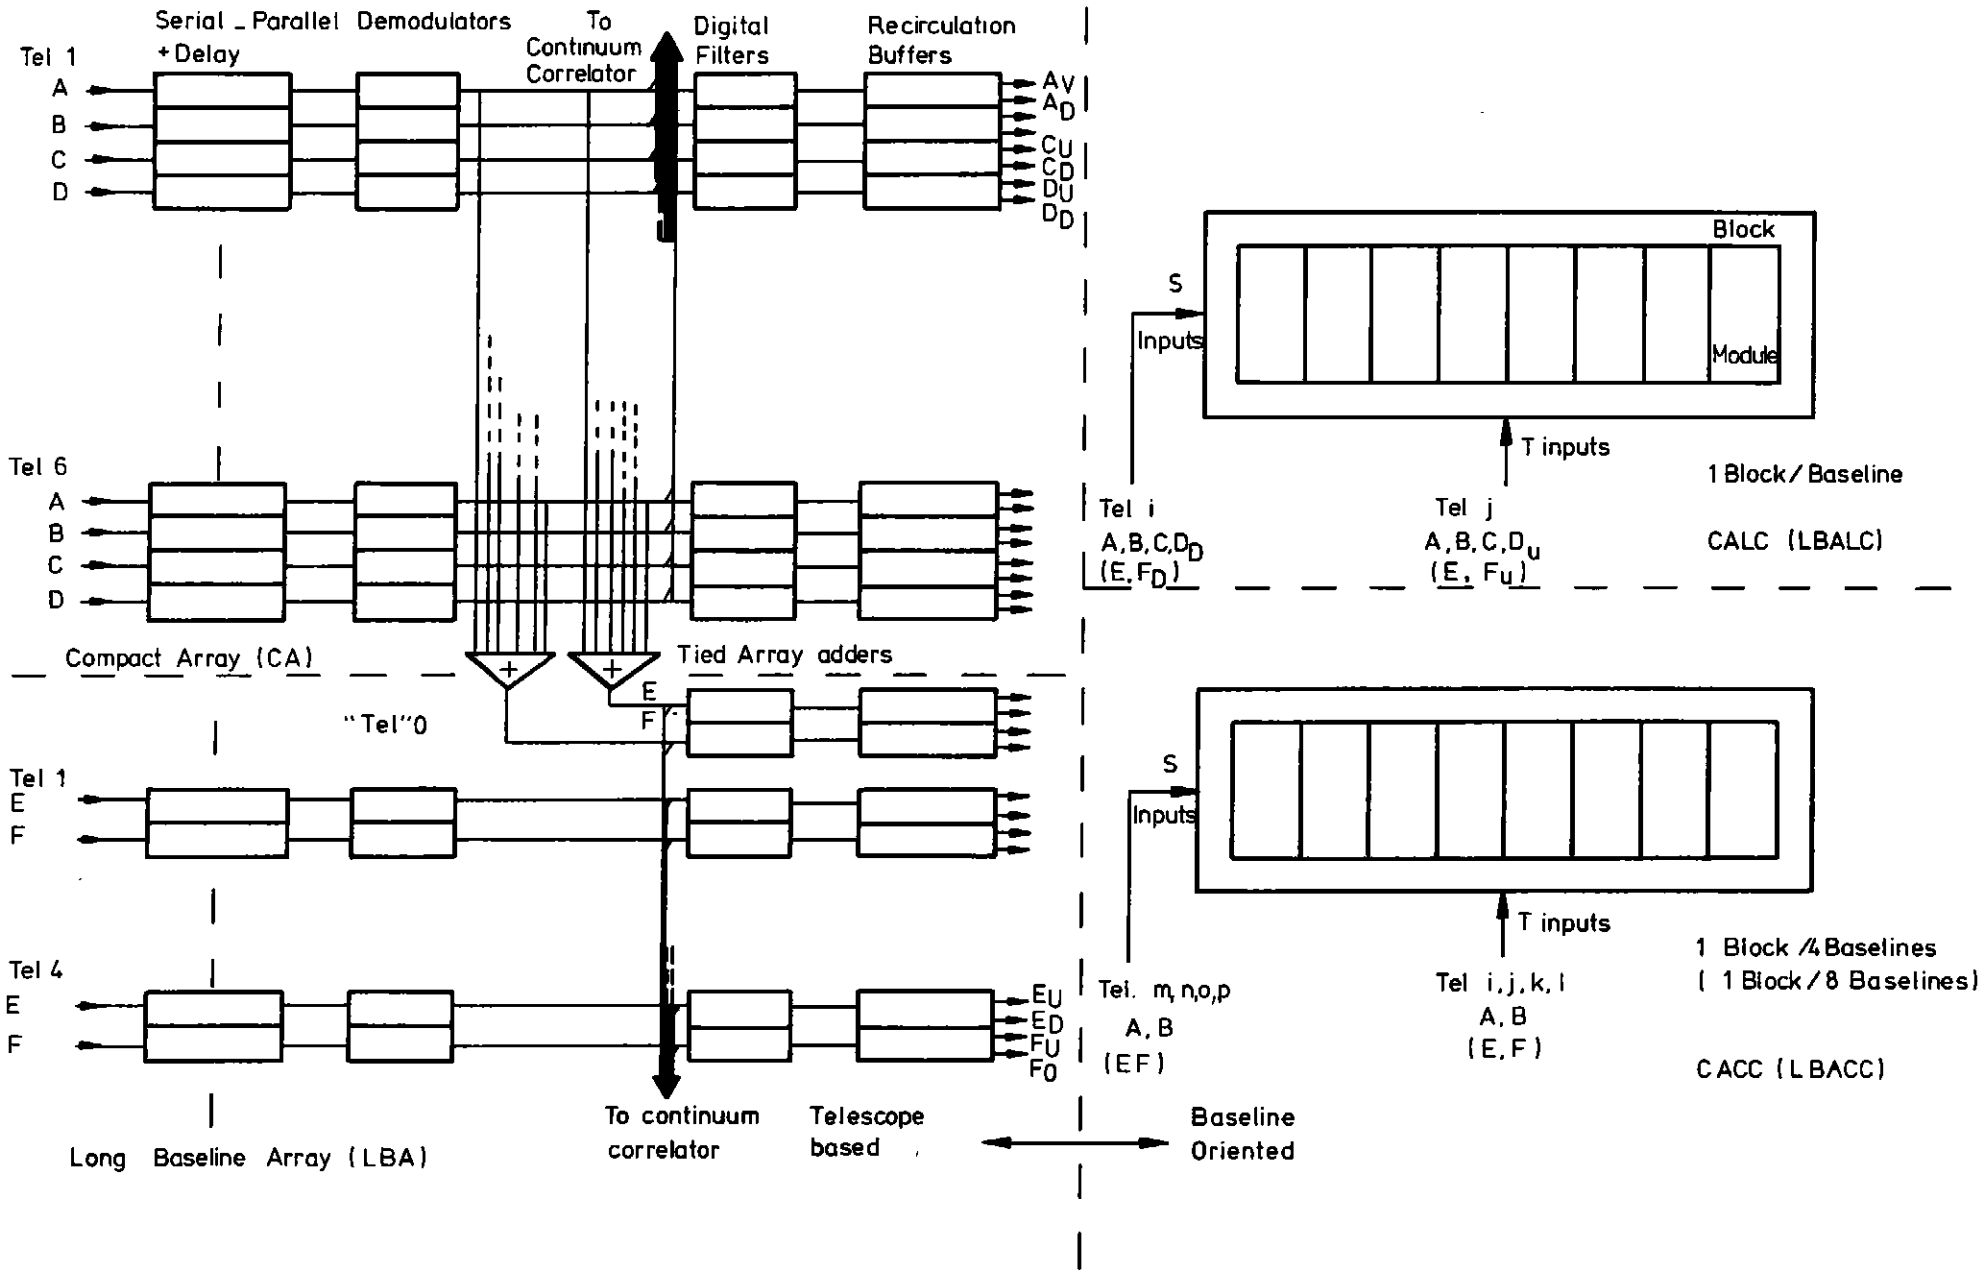

Figure 3.1 : Block diagram of the proposed correlator system

spurious signals. Some thought should also be given to correlator blanking to avoid correlation during instants of switching.

The extension to require four essentially independant correlator subsystems is a reflection of the reduced flexibility inherent in the baseline oriented approach where channels cannot be swapped from one baseline to another. Unused channels for a given baseline cannot be utilised elsewhere.

The four separate correlator subsystems are in turn composed of identical modules. All units are composed of "blocks" and "modules", which are defined thus :

MODULE :

A module is an array of 16 XCELL chips, arranged in a 4*4 matrix with 32 S (side) inputs and 32 T (top) inputs. If the S (T) inputs are a set of samples of an input signal A (B) at a sampling period of t, then the module produces a set of 32 samples of the cross-correlation function of A and B, with sampling period in lag space of t. Hence, in its normal mode of operaton it provides the data for 16 frequency channels. There are 1024 separate correlators per module.

BLOCK :

A block is a set of 8 modules. It is normally connected to a single interferometer and is capable of producing all the useful products obtainable from this baseline. Each module within a block can either be connected directly to the block inputs, or to the outputs of another module to provide more channels for a particular product. Each block has an input switching matrix and internal switching to provide this function. A block also contains an output adder, which adds the 8K correlator outputs into an accumulator memory during each correlator cycle.

A detailed description of the four correlators follows.

3.1 COMPACT ARRAY LINE (OR WIDEBAND CONTINUUM) CORRELATOR.

This is a baseline oriented correlator containing one block for each of the 15 baselines. At maximum bandwidth and with 8 products per baseline, it provides 16 channels per product. With fewer products, there are proportionally more channels per product. It will be used for CA line measurements or continuum

observations above 40MHz bandwidth.

3.2 LBA LINE CORRELATOR.

The same as 1. above, with one block for each of the 6 baselines.

3.3 COMPACT ARRAY CONTINUUM CORRELATOR.

This unit is used for CA continuum measurements at 40 MHz bandwidth. (i.e. reference continuum for simultaneous line measurements). In this case, the structure is baseline oriented at the module level, rather than at the block level. Two modules are provided for each baseline. To obtain more than two products per baseline, two signals are time multiplexed onto the inputs of the modules (i.e. alternating samples across the S and T inputs) to provide up to 4 products per module with 8 frequency channels per product per module. Although the block no longer has a functional significance, it is envisaged that the modules in this correlator will be constructed in four "blocks" of 8 modules which will be physically similar to (if not the same as) the line correlator blocks. It is yet to be determined whether the block input switching matrix can be designed to satisfy the requirements of both the line and continuum correlators.

It should be noted that the reconfiguration scheme proposed in AT/10.4/002 is not possible in this correlator.

3.4 LBA CONTINUUM CORRELATOR.

This unit is used for all continuum observations with the LBA. It is similar in structure to the CA continuum correlator, but, because there are no more than 4 products per baseline, it only contains one module per baseline. The entire unit is contained within one block. Once again reconfiguration is not applicable.

A single module used in this way is capable of providing 8 frequency channels at 40 MHz bandwidth. In fact, the radio link will limit bandwidths to 10 MHz in the first instance. Recirculation would then allow 32 frequency channels to be obtained for the smaller bandwidth.

One point requiring further examination is the number of frequency channels. A limited number of lag channels used to produce the frequency resolution will give rise to an imperfect frequency beam or response function. Eight frequency points, for example, would give rise to poor estimation of the sine

channel continuum correlation and effects such as "ghosting" may be of great concern (Bos - IAU/URSI Symposium on Indirect Imaging).

3.5 EXPANSION

Future expansion of the AT correlator may include the following steps.

1. Expansion of the CA Line correlator to 2 blocks per baseline. This involves the addition of an extra 15 blocks.

2. Expansion of the LBA Line correlator to cater for more baselines at the rate of one block per baseline.
   NOTE: The provision of 15 extra blocks would allow either 1. or 2. (to 15 baselines) above. Alternatively, once the LBA was extended to 15 baselines (6 antennas), the blocks in the LBA Line correlator could be used to expand the CA Line correlator. It is envisaged that this would only be done by manually switching the blocks in question.

3. Expansion of the LBA Continuum correlator to provide for more baselines at the rate of 1 block per 8 baselines.

The overall unit count for the AT correlators then becomes:

DEC. 1988	Blocks	Modules	XCELLS
CALC	15	120	1920
LBALC	6	48	768
CACC	4	32	512
LBACC	1	8	128
TOTAL	26	208	3328
EXPANSION			
CALC	15	120	1920
LBALC	15	120	1920
LBACC	2	16	256
	32	256	4096
GRAND TOTAL	58 Blocks	464 Modules	7424 XCELLS

CHAPTER 4

BASELINE ORIENTED CORRELATOR COMPONENTS

This chapter describes the baseline oriented sections of the correlator. Major attention is devoted to the line correlator system and the switching necessary for reconfiguration.

Reconfiguration was seen at the workshop to be essential in view of the total recirculation memory cost if only recirculation was performed (Chapter 5). However the implications are not insignificant for the baseline oriented components as considerable extra delays are required in many signal paths.

The point of perhaps greatest importance to the success of the entire correlator system is covered here also. The XCELL correlator chip is considered at great length with regard to requirements, yield and time scales.

4.1 THE XCELL

The basic correlation element of the proposed Australia Telescope correlator is the 8x8 XCELL chip. This comprises 64 correlator-accumulator structures arranges in an 8x8 matrix. A diagram of the XCELL is given in figure 4.1. A fall-back option of a 4x4 chip is also under consideration, but use of this would involve both a fourfold increase in the number of chips required and a doubling of the minimum size of the recirculation memory needed to achieve the required utilization factor of 64. Because of this, only the 8x8 implementation will be discussed in this note. Obviously the fall-back option would require a considerable decrease in the overall capability of the correlator.

The 8x8 XCELL chip, when operated at a 10MHz basic clock rate, can be considered as capable of measuring 8 lags at 80 MBaud. The 80-MBaud data rate comprises Nyquist sampled data of either 40 MHz bandwidth with 1-bit digitization, or 20 MHz bandwidth with 2-bit digitization. The two bits are input to the XCELL chip in successive 80 MHz clock periods thereby

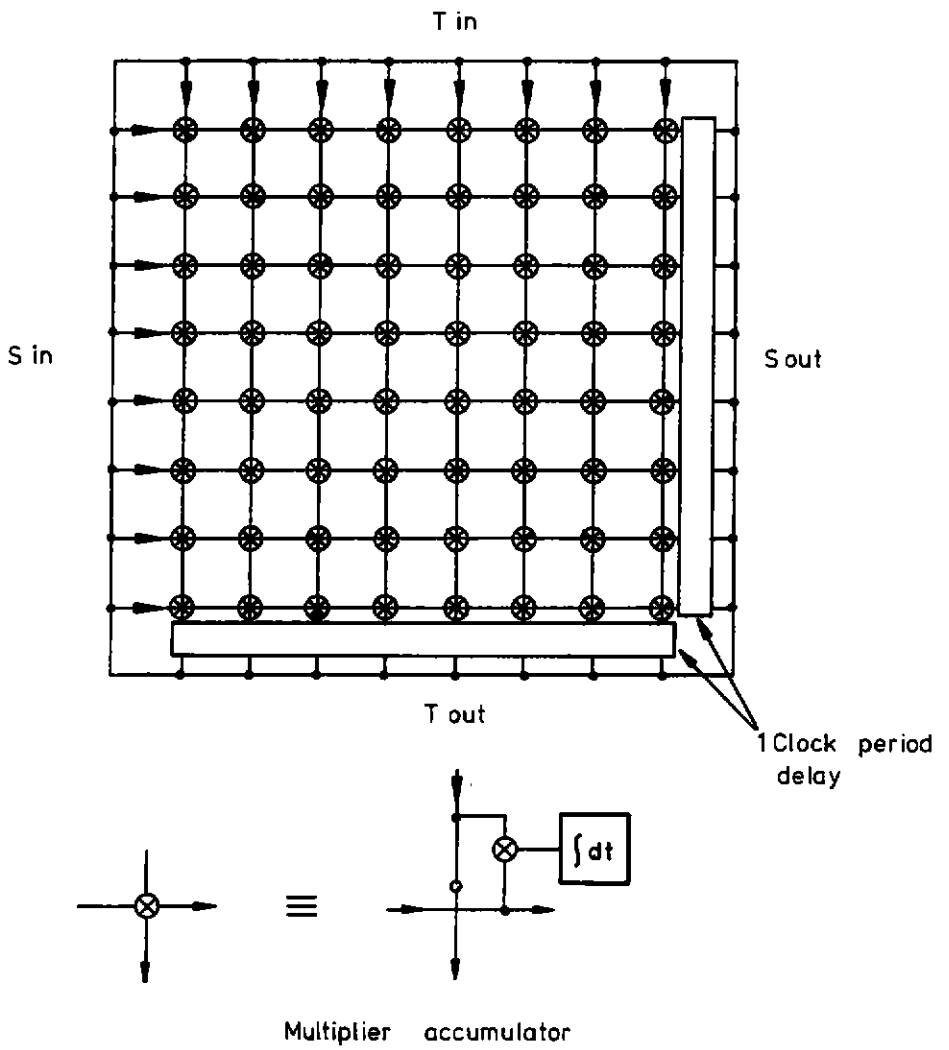

Figure 4.1 : An 8 x 8 XCELL correlator chip, a conceptual organization

requiring a halving of maximum bandwidth. Various methods of performing the 2-bit correlation are at present being examined with a view to fabrication of prototypes in the December AUSMPC multiproject VLSI chip run.

The XCELL chip as presently implemented includes sufficient on-chip accumulation to integrate for a single correlator dump period of 100 ms. The dump period is primarily set by recirculation memory requirements. The possibility of providing much smaller accumulation periods with on-chip accumulator and employing off-chip RAM accumulator to reduce chip area is presently being considered.

Should the XCELL prove incapable of operation at 10 MHz, the impact on the correlator system would be similar to that of failing to produce an 8x8 XCELL. The product of number of channels and speed is the relevant measure of processing power.

4.2 THE MODULE

A MODULE capable of measuring 32 lags at 320 MBaud can be constructed by connecting together 16 XCELLs in a 4x4 array. The method of interconnection is shown in figure 4.2. The module accepts 32 S (side) inputs and 32 T (top) inputs and provides 16*64=1024 accumulated products. In maximum bandwidth mode (see for example AT/10.4/002) each of the 32 inputs for the S and T inputs is provided by a one of 32 parallel bits of a serial to parallel conversion (actually performed at the correlator system input) which lowers the data rates by a factor 32. Figure 4.2 also shows the lags measured in this mode.

It is also possible to reconfigure such a module so that it measures increased numbers of lags (by factors of 2) at correspondingly reduced data rates. Thus a single module can be considered capable of measuring 32 lags at 320 MBaud, 64 lags at 160 MBaud, 128 lags at 80 MBaud, or 256 lags at 40 MBaud for reconfiguration factors (K) of 1, 2, 4, and 8 respectively. It should be remembered that the bandwidth for Nyquist sampled data is either half the Baud rate for 1-bit digitization, or one quarter of the Baud rate for 2-bit digitization.

The details of this reconfiguration are considered in section 4.4.

Each Module, comprising 16 8x8-XCELLs, contains a total of 1024 correlators and can measure a maximum of 256 lags with a reconfiguration factor (K) of 8 but without recirculation. This can be increased to 2048 through recirculation by a further factor of 8; and a cascade of eight modules, fully reconfigured and recirculated can measure the maximum required number of lags (16384) giving a complex spectrum of 8192 points.

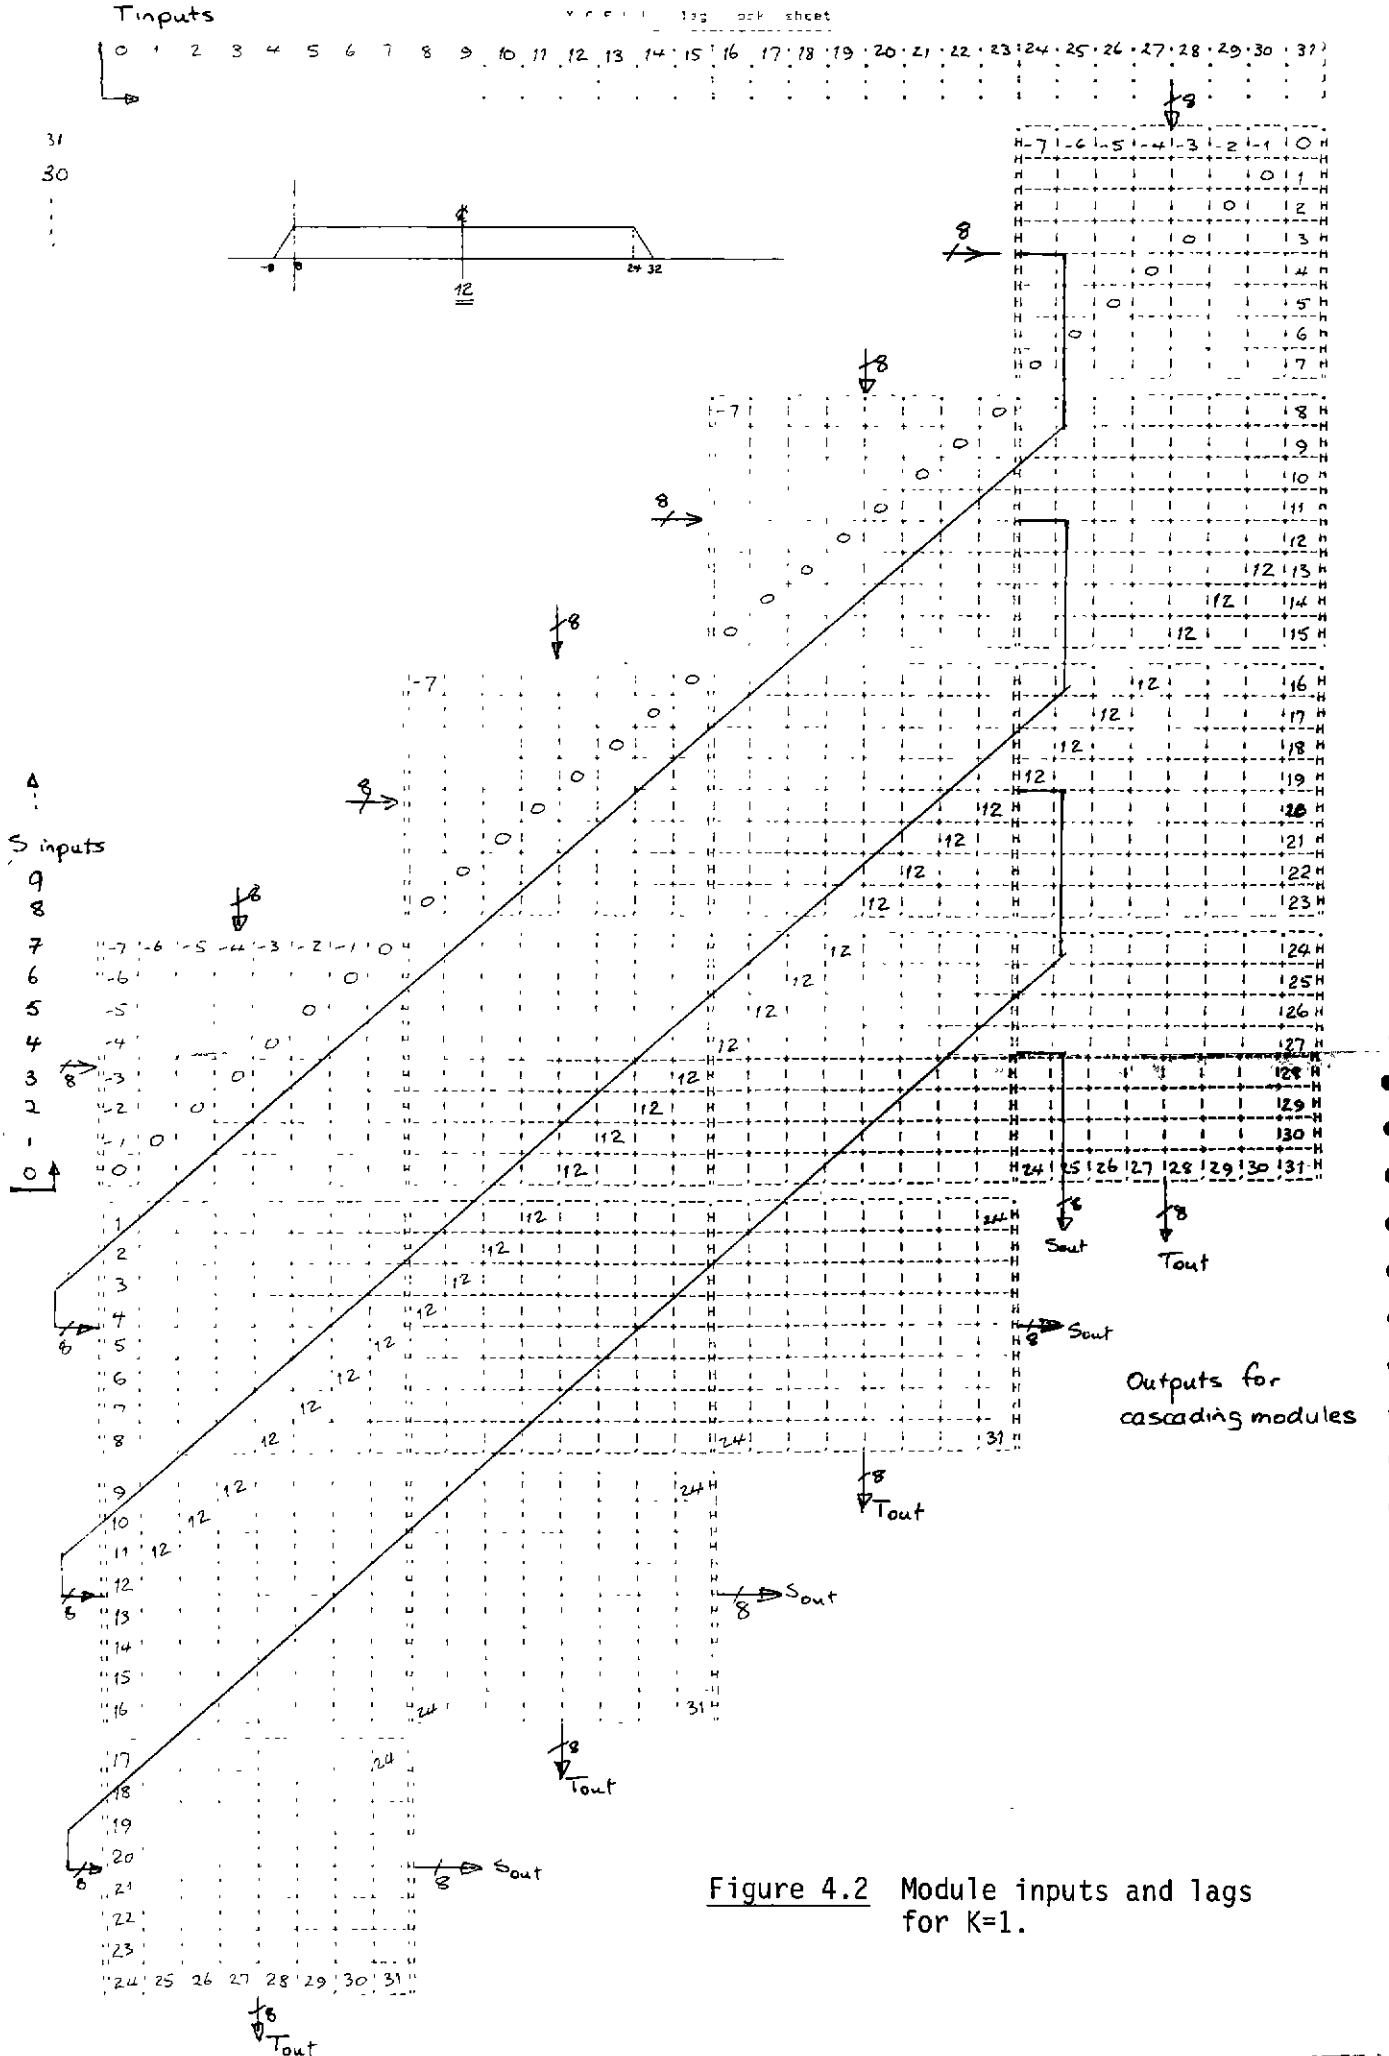

Figure 4.2 Module inputs and lags for K=1.

4.3 THE BLOCK

Various configurations called for by the astronomical objectives of the A.T. require up to eight simultaneous cross-spectra to be measured for a single baseline with a resolution of up to 1024 complex frequency points. The post-XCELL processing of the correlated data also fits in well with groups of eight modules. Eight modules, maximally reconfigured and recirculated, require 16K x 32-bits (64KBytes) of buffer memory, which is a reasonable size of memory to either build or purchase on a single card. Each module generates 1024 x 16-bit words of data to be processed each read-out cycle, thus with a read-out cycle period of 20mS and groups of 8 modules, each word of data can claim up to 2.5 microseconds. This is quite a reasonable time to allow for adding each data word to the contents of its corresponding memory location. All of these considerations lead to the conclusion that 16-XCELL modules should be grouped in 8-module BLOCKs. (It should also be noted that BLOCK is an acronym for "BaseLine Oriented Correlator Kluge".)

Each BLOCK comprises 8 modules, two 16K x 32-bit output data buffer memories, an output data adder with associated steering memory and control circuitry, and an input data path switching network to permit connection of up to 4 inputs from each of two telescopes to selected module inputs. It is also worthwhile to note that the same BLOCK can be used, with a modified non-switched set of input data path connections for the special continuum correlators which, by dint of the internal matrix organisation of the XCELL chips, generate all four Stokes parameters for a single baseline/frequency to a resolution of 8 complex frequency points in a single module, or to 16 complex frequencies in just two.

4.3.1 Signal Switching Within A BLOCK

The various modes sufficient to measure all sensible IF combinations for a baseline were described in section 2.2.2. A BLOCK has eight S inputs and eight T inputs which must be provided with signals from four undelayed inputs Au,Bu,Cu,Du and four delayed inputs Ad,Bd,Cd,Dd. The delayed inputs are required to center the correlation function measurement at zero lag or to shift the correlation function for either reconfiguration or recirculation. The undelayed and delayed outputs are to be generated by the recirculation memory and possibly also by an additional unit to provide for reconfiguartion requirements where the recirculation memory is not connected.

The possibility of directly cascading MODULES to increase the number of channels assigned to a product must also be provided.

While it might have been tempting to provide a pair of full 4x8 selection matrices (the undelayed to the S, delayed to T only), this ignores the fact that not all possible combinations are relevant and the order is also irrelevant. It must be remembered that all data paths are 32 bit wide, 10 MHz so a general purpose switch will require considerable hardware.

Figure 4.3 shows a possible simplified switching arrangement. This is tailored specifically to the requirements of eight modes. Some flexibility in choosing numbers of channels per product has been forsaken. The possibilities are summarised in the table;

Mode	Length of Product in MODULES							
	AA	AB	BA	BB	CC	CD	DC	DD
1	8							
2	4			4				
3	2			4	2			
4	2			2	2			2
5	2	2	2	2				
	1	3	3	1				
6	1	1	2	2	2			
7	1	1	1	1	2			2
8	1	1	1	1	1	1	1	1

Clearly there are some limitations here. In particular, modes 2,4,5 and 6 will ideally require more flexibility in assignment of numbers of modules. It is not clear to what extent this would require an extensive reworking of the basic switching as shown.

4.4 LAG DISTRIBUTION AND RECONFIGURATION

4.4.1 Lag Spectrum

The spectrum of lags generated by any particular configuration of XCELLs is not quite as simple as the statements in the second paragraph of section 4.2 might indicate. Those statements correctly indicate the incremental contribution of a single XCELL or module when concatenated with another, but only in the case of K=8 do they indicate the exact size of the measured lag spectrum. In all other cases, each end of the lag spectrum contains lag terms which are only partially measured. These partially measured lag terms correspond to the extreme upper and lower edges of the XCELL array.

The array of XCELLS in a module is chosen as a parallelogram in order to minimise the number of partially measured lags. The parallelogram structure is composed of a

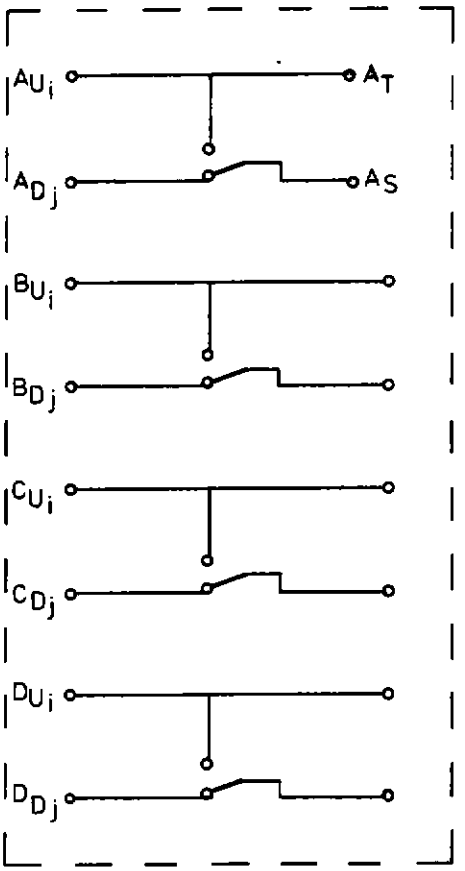

Possible additional switching to measure autocorrelations of telescope i.

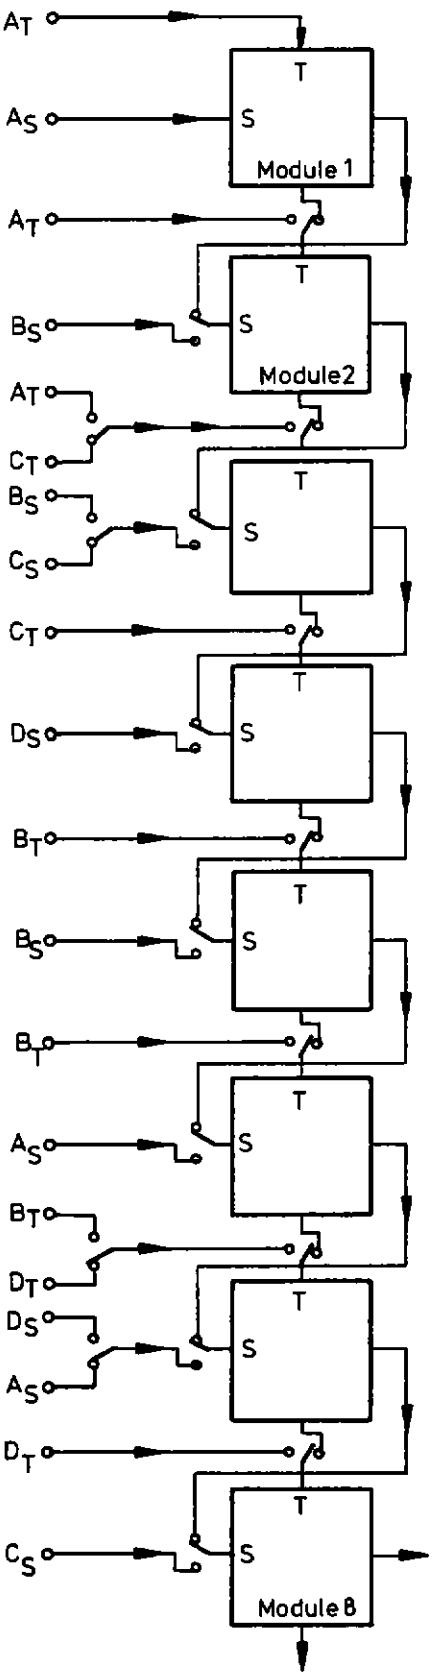

Figure 4.3 : Possible reduced switching of modules within a block to obtain all required products

series of squares, each square being a single XCELL. Each XCELL, operating with K=1, only fully measures the lag term along its diagonal. If this is the lag=0 term, it can be seen that it also measures seven eighths of each of the lag=-1 and lag=+1 terms, six eighths each of the lag=-2 and lag=+2 terms and so forth, measuring only one eighth each of the lag=-7 and lag=+7 terms. This can be stated more briefly as follows: the single XCELL fully measures the lag=0 term and tapers to zero at lag=-8 and lag=+8.

When connected together as an array the internal tapers match so that the missing components in one XCELL are supplied by the adjacent one. Thus the only partial measurements are at the most negative and most positive lags. A 16-XCELL module with K=1, fully measures the lag spectrum from lag=0 to lag=24 tapering to zero at lag=-8 and lag=32. The mid point of this lag spectrum is lag=12. Therefore, in order to measure equal lags in both positive and negative directions, a delay of 12 sample periods must be applied to the relatively undelayed -- or 'S' -- inputs of the module.

Other alternatives would be a square array in which only the lags along the diagonal would be completely measured, but the maximum to minimum measured delay range would be greater. In comparison with the parallelogram array the square array could offer better resolution with the exchange of some sensitivity. It may be noted that no carefull comparison has been performed. Also no analysis has been performed to determine the extent of any aliasing which might occur due to the regular pattern of the discarded samples.

4.4.2 Reconfiguration

The basic idea behind reconfiguration was described in AT/10.4/002. The 8*8 array of multiplier accumulators in the XCELL may be reconfigured as four 4*4 sub-arrays, sixteen 2*2 sub-arrays or sixty four 1*1 sub-arrays with certain restrictions. The structure is restricted by the connectivity of the XCELLs in that each S input is multiplied by all T inputs and vice versa. This defines that the sub-arrays cannot be independant. For example, the data input to the four 4*4 sub-arrays must consist of two groups (or words) of four samples on the S inputs and another two groups of four samples on the T inputs. The delay between words at each of the groups of inputs must be so arranged that a contiguous set of lag terms is measured. This constraint requires that different delays be imposed on the sample words in the S and T inputs. Figures 4.4,4.5 and 4.6 show one set of solutions with the correspondingly generated lags.

The original proposal called for the delays required in each set of inputs to be produced by suitably tapping the signal path as it propagates through the XCELL array. Further

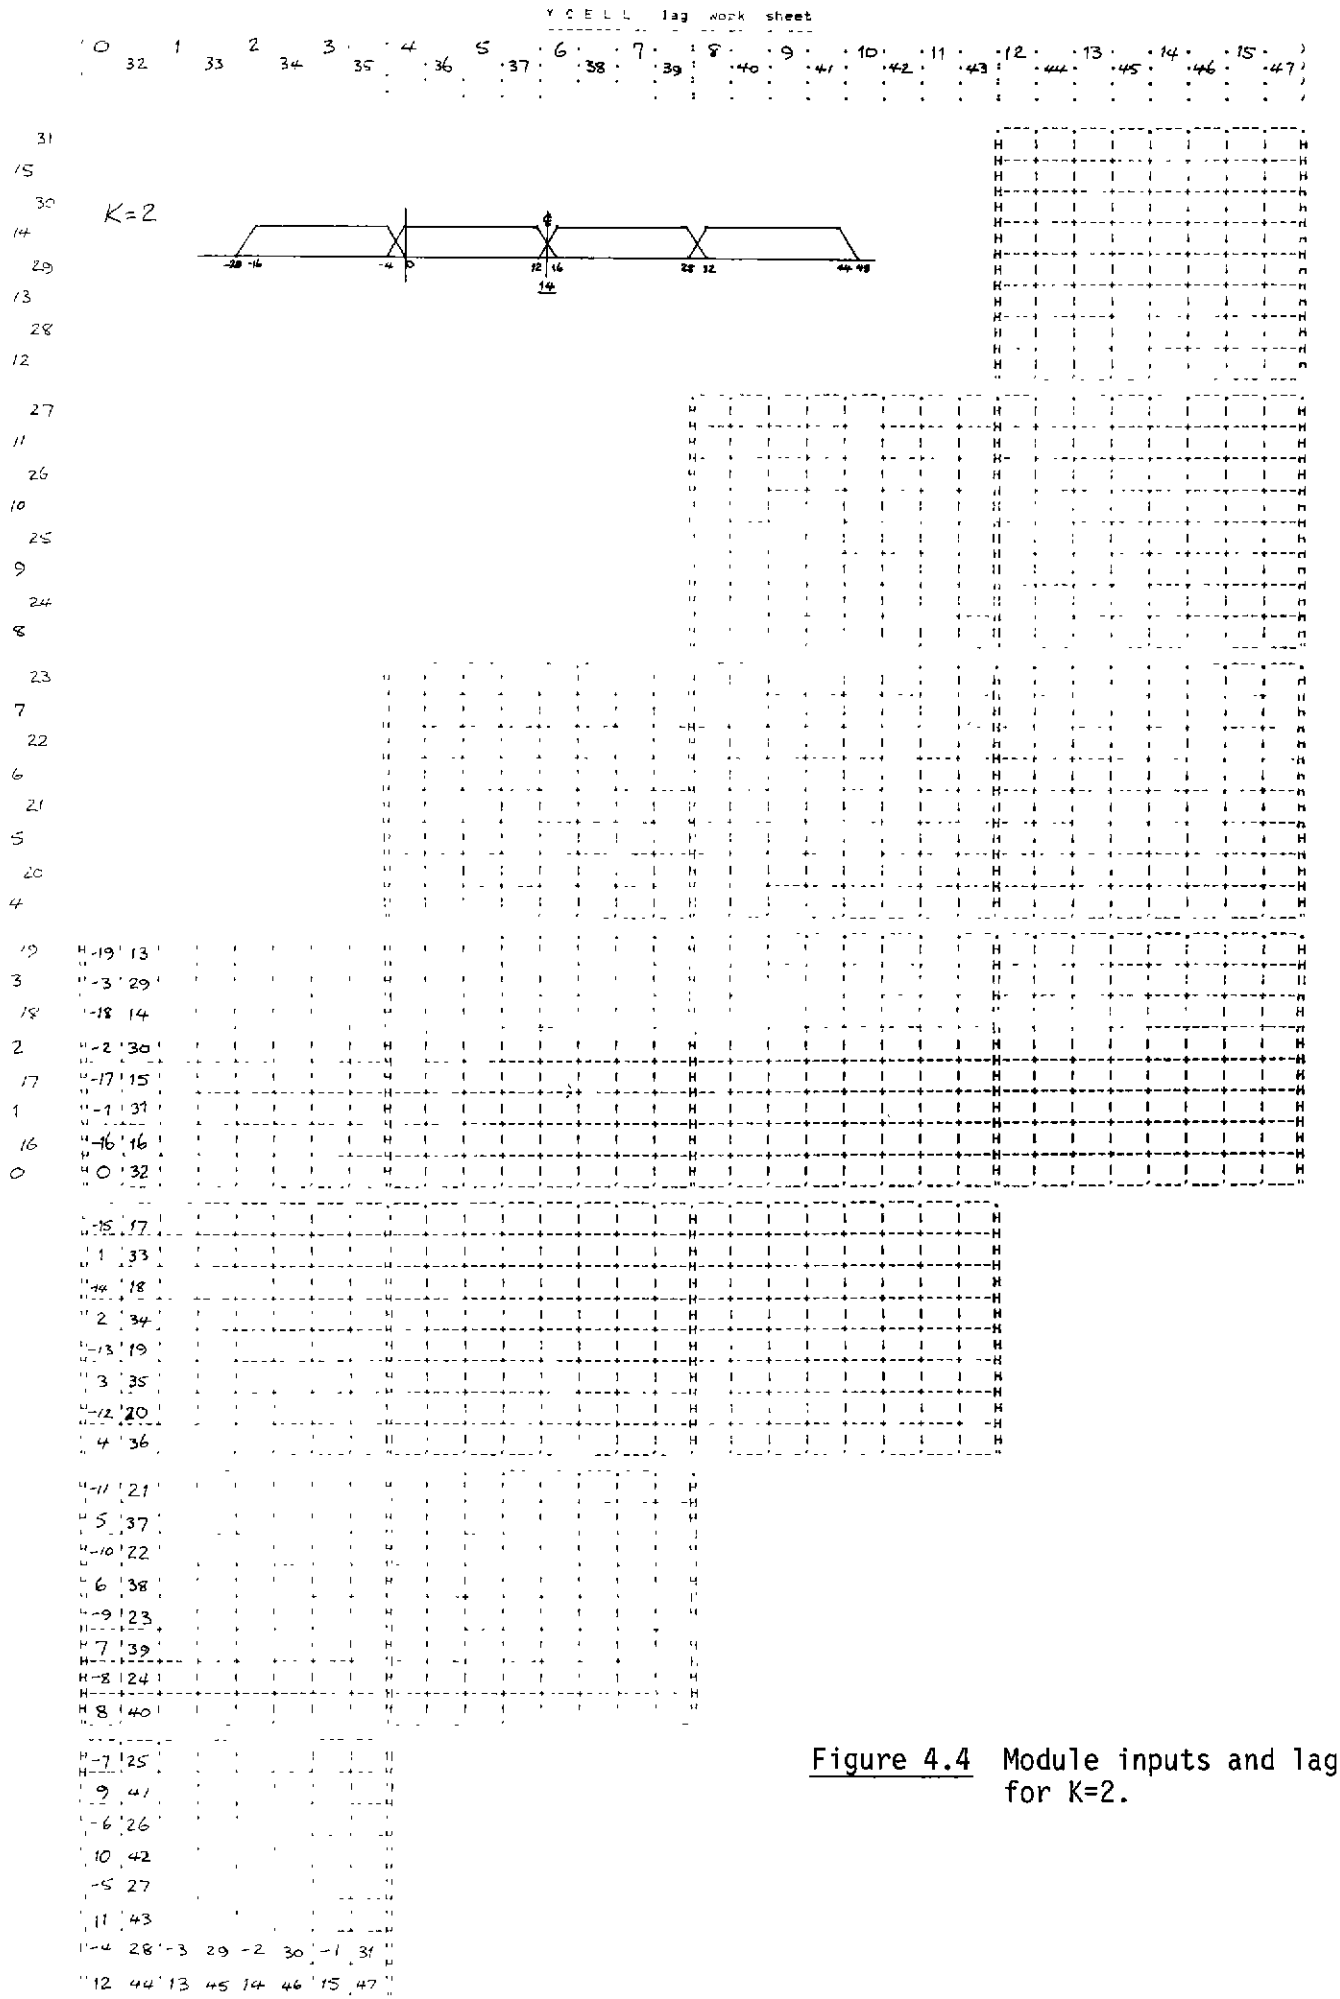

Figure 4.4 Module inputs and lags for K=2.

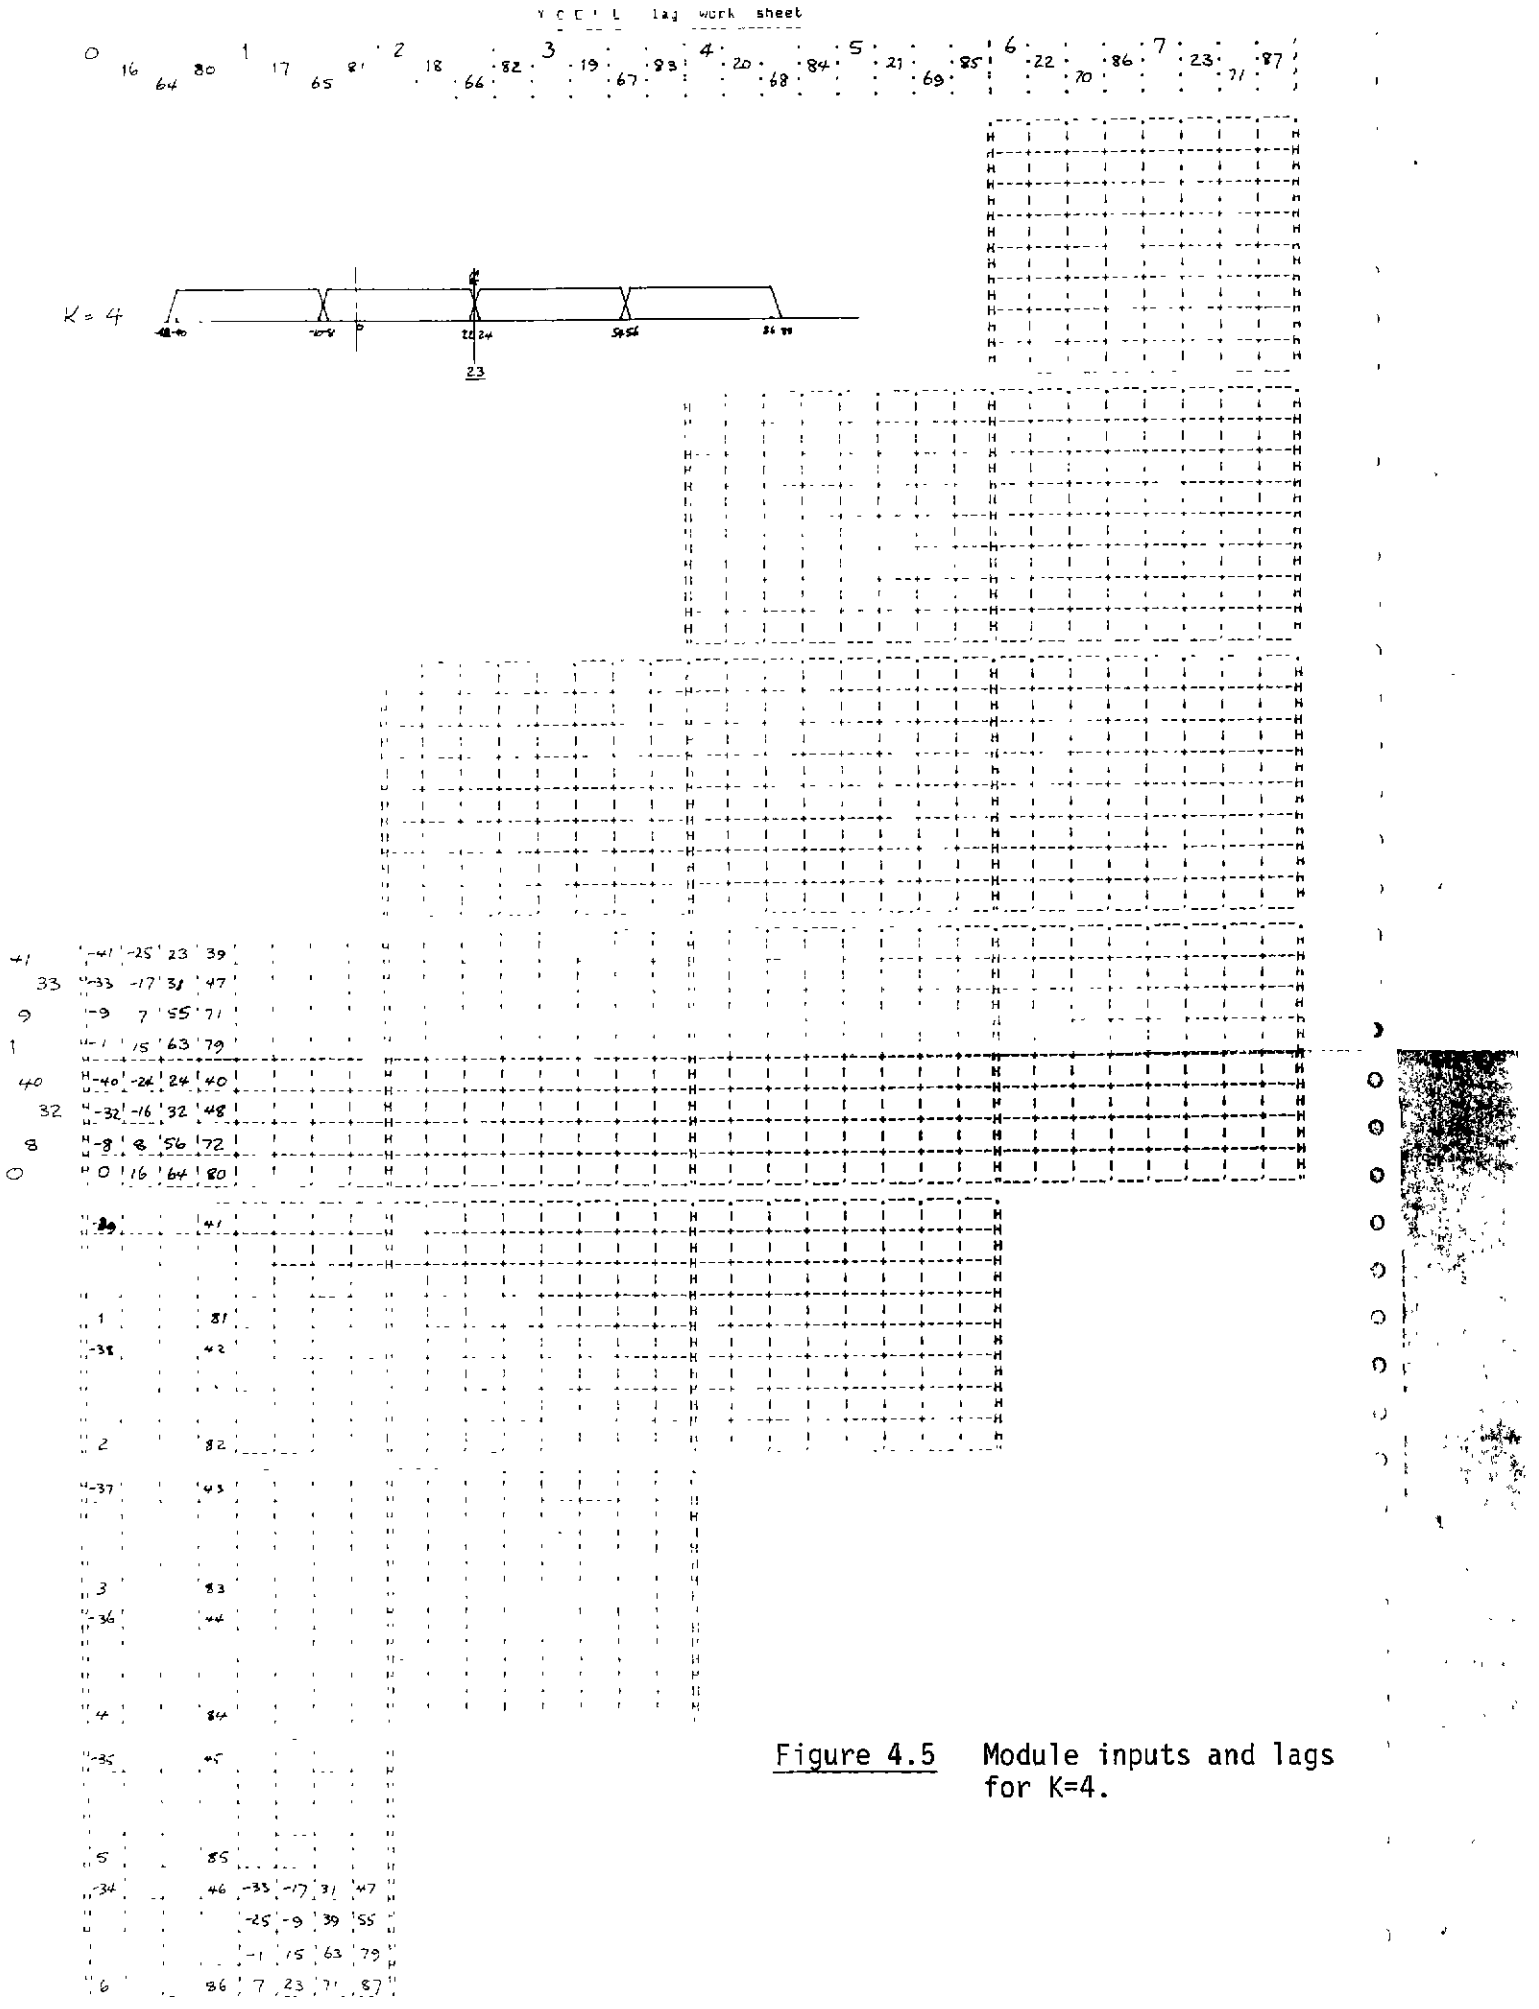

Figure 4.5 Module inputs and lags for K=4.

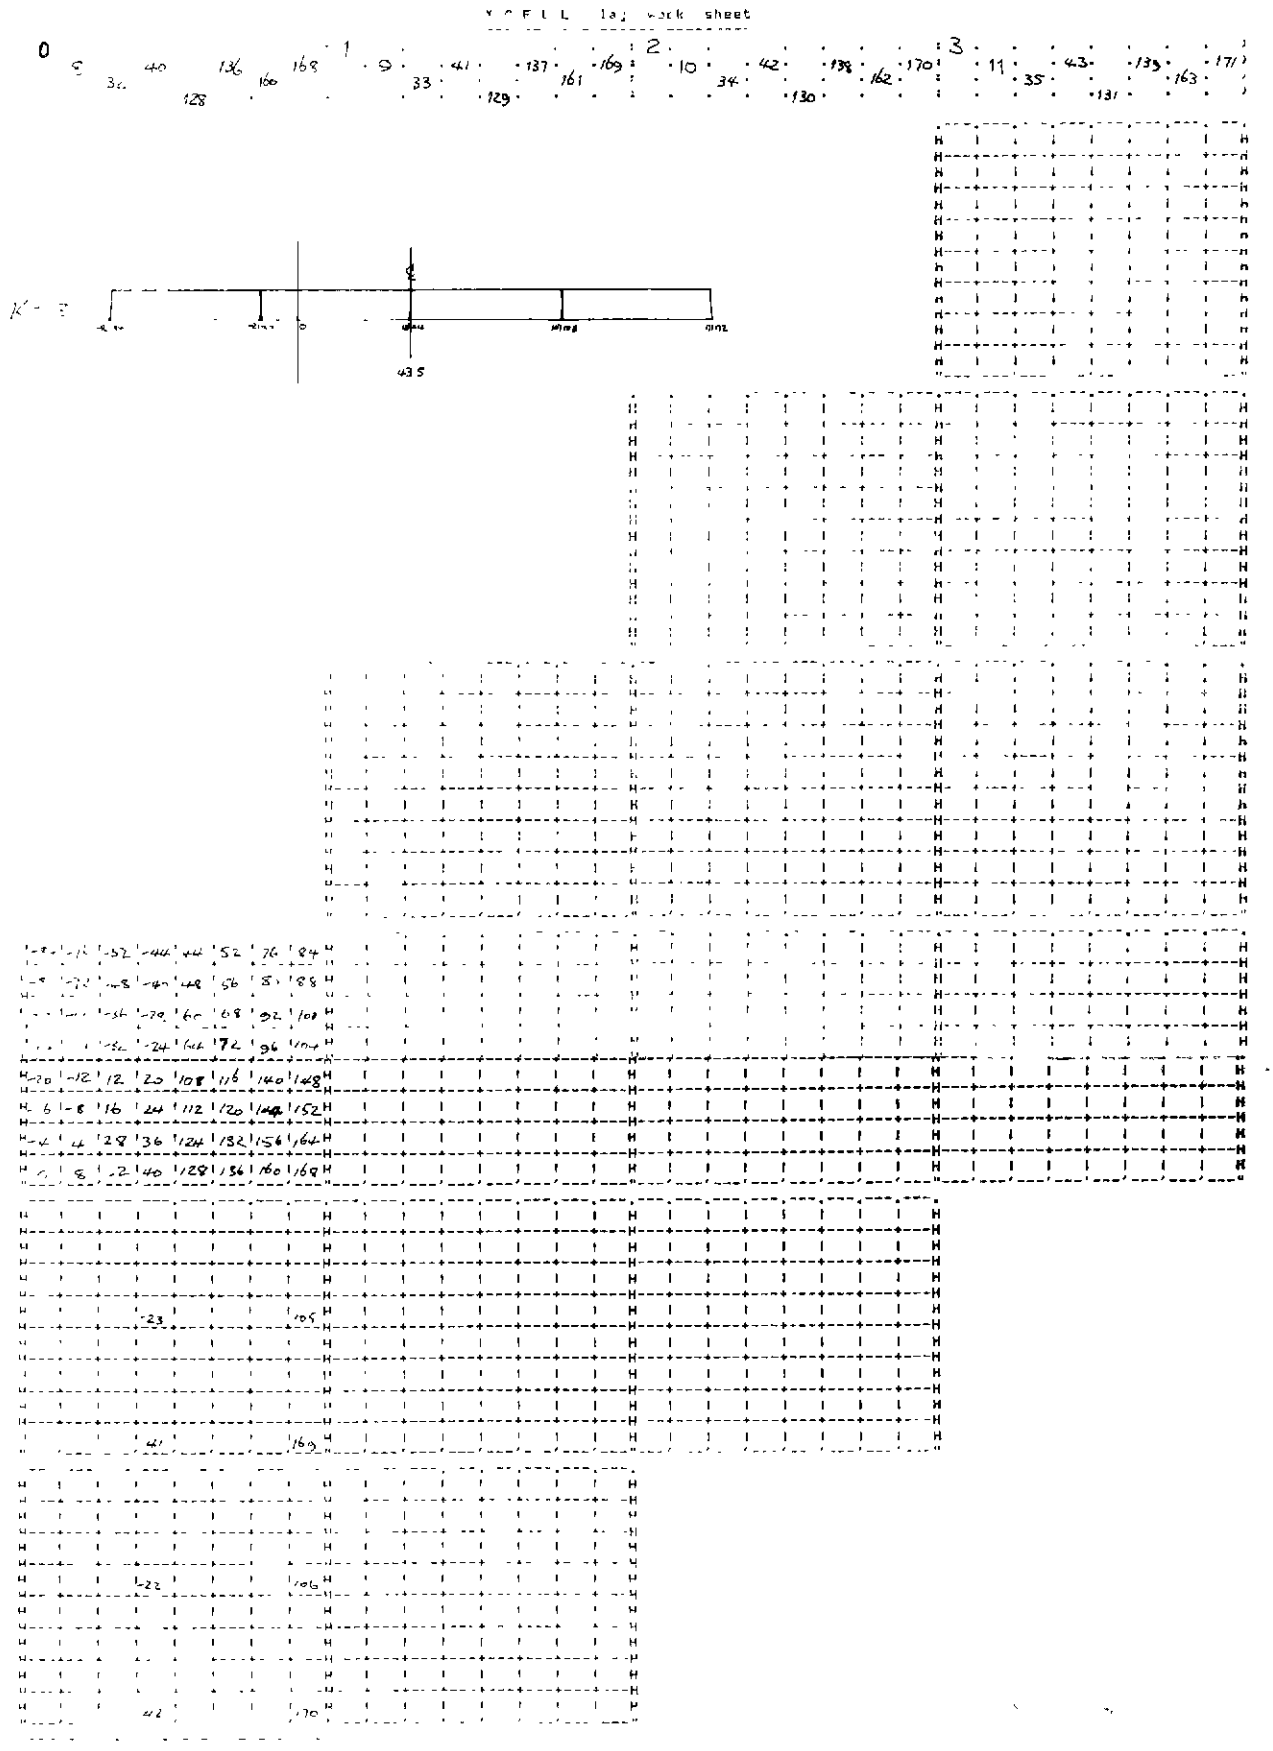

Figure 4.6 Module inputs and lags for K=8.

consideration reveals that this approach runs into trouble for reconfiguration factors greater than K=2.

Bearing in mind the need to provide for the use of modules in cascade (for the measurement of more lag channels), two different methods have been considered for providing these delays. In one method the modules are cascaded without intermediate delays in exactly the same way as for K=1. This requires all the delay to be inserted at the top of the concatenation; and for the amount of delay to change with different numbers of concatenated modules as the sub-arrays grow in length. With this method the concatenated sub-arrays measure contiguous lag terms and the delays arrange the spacing of the sub-array lag centres so that an overall contiguous set of delays is measured. Clearly, with this method a separate, configurable delay module would need to be provided for each module or group of modules which might be cascaded.

In the second method, delays are provided at the input of each module in such a way that the module as a whole measures a contiguous set of lag terms independent of the number of modules which are to be concatenated. An additional set of delays is then provided at the module output to permit a similarly configured module to be concatenated for measurement of the adjacent set of lags. With this method any number of modules can be concatenated without changing the size of any delay elements. The required delays, which have a definite size for each value of the reconfiguration factor K, can be built into the modules, indeed into the XCELL chips themselves, and all modules constructed identically.

In order to calculate the amount of delay which has to be added in each path, we might start by considering a single array of XCELLS, or a single sub-array, with or without cascading, as providing a range of lags described by,

Fully measured lags: 0 to M inclusive.
Partially measured lag range: 1-N to -1, M+1 to M+N-1

where M and N are functions of the reconfiguration factor K, and the number of modules cascaded or concatenated C. Then we have,

$$M = 24.C/K + N.(C-1)$$

$$N = 8/K$$

(Note, if the second method described above is used for providing the delays, the value of C is always 1.)

The range of lags may be shifted in the positive sense by adding delay to the T inputs or in the negative sense by adding delay to the S inputs. In order to provide a continuous range of lag channels with no duplication or incompletely measured lags, the delays in the S and T inputs of the sub-arrays must be arranged to give a range of lags separated by units of M+N.

Thus we have lag steps Lstep, given by,

$$Lstep = 32.C/K$$

(Note that, because the modules are rhomboid, the value of Lstep when C=1 is exactly the same as the word size for any value of K.)

A number of solutions for the delays in the S and T inputs are possible. For a given reconfiguration factor K, there are K.K separate sub-arrays. In order for these to measure a continuous set of lags the delay differences at the S and T inputs must be a contiguous range of values chosen from a set whose elements differ in value by Lstep. These values when divided by Lstep can be represented, modulo K.K, by the complete set of numbers which can be expressed in 2.Lgg(K) bits, (i.e. 6 bits for K=8, 4 bits for K=4 and 2 bits for K=2). This arrangement of delay differences can be achieved by inserting ahead of the S inputs a set of K delays equal to Lstep multiplied by the K possible combinations of the powers of 2 represented by one half of the bits, and ahead of the T inputs the K delays derived from the other half. Selection of which bits is assigned to each set gives some freedom to produce solutions with the central delay closest to zero but this may not be the only factor influencing the assignment.

One valuable fact worth noting in achieving a suitable assignment is the following. The bits required to be assigned for smaller values of K can be seen to be a subset of those required for larger values. A saving in hardware to be constructed might well be achieved if the assignment of a particular bit to S or T was to be the same for all values of K for which that bit is required. This still allows considerable freedom in the choice of assignments.

If the second method for achieving the required delay, described above, is used, the lag range of the additional modules can be made contiguous if a delay equal to Lstep.K.K is inserted in either the S or T data path; of course, in this case this is just K.K words, since Lstep is equal to the word size.

One possible solution is given by the delays shown in the table below. This arrangement is achieved by retaining assignments from smaller values of K when dealing with higher values and assigning the higher valued bit of the two new candidates in each case to the T inputs. With this arrangement the lags measured by a single module are as shown in figures 4.2, 4.4, 4.5, and 4.6. The central lags are always positive, so if method 2 is used the output delay is placed in the T path; this ensures that the centre lag of the concatenation is also always positive and thus the compensating delays, required to centre the lag spectrum, need only be placed in the S path.

K	Input Delays multiples of Lstep (samples) S	T	Output Delay Method 2 (T path)	Lstep
1	0 (0C)	0 (0C)	1 (32C)	32C
2	0 (0C) 1 (16C)	0 (0C) 2 (32C)	4 (64C)	16C
4	0 (0C) 1 (8C) 4 (32C) 4+1 (40C)	0 (0C) 2 (16C) 8 (64C) 8+2 (80C)	16 (128C)	8C
8	0 (0C) 1 (4C) 4 (16C) 4+1 (20C) 16 (64C) 16+1 (68C) 16+4 (80C) 16+4+1 (84C)	0 (0C) 2 (16C) 8 (64C) 8+2 (80C) 32 (128C) 32+2 (136C) 32+8 (160C) 32+8+2 (168C)	64 (256C)	4C

4.4.3 Measured Lag Range

From the above it can be seen that minimum and maximum lag can be simply specified in terms of the maximum delay in the S input Smax and the maximum delay in the T input Tmax, and from this the central lag can be calculated. We have then,

Fully measured lags:
Minimum lag = -Smax
Maximum lag = Tmax + M
Partially measured lags:
Minimum lag = -Smax - (N-1)
Maximum lag = Tmax + M + (N-1)

Central Lag = (Tmax - Smax + M)/2
= (Tmax - Smax + 32.C/K)/2 - 4/K

It is then only necessary add an equivalent amount of delay in one of the input paths in order to centre the measured lag range. If the central lag is positive the compensating delay needs to be added in the S input path, if negative in the T path. From the hardware point of view it can thus be seen that the central lag should always have the same sign so that compensating delays need only ever be placed in one input path.

The table below gives the limits of the lag spectra for the four configurations of a single 16-XCELL module using the delay difference bit assignment of the example in section 4.4.2.

K	no measurement at lag=	fully measured at lag=	fully measured until lag=	no measurement at lag=	centre point d(K)
1	-8	0	+24	+32	+12
2	-20	-16	+44	+48	+14
4	-42	-40	+86	+88	+23
8	-85	-84	+171	+172	+43(.5)

Natural lag spectra for various values of the reconfiguration factor (K) for a single 16-XCELL module

Note that, while the delays needed at the module inputs for reconfiguration are whole multiples of the size of the input words, the compensating delays corresponding to the central lag represent fractions of a module input word. Note also, the central lags are expressed in tha above in terms of lags or samples and are thus independent of the number of bits used in digitising the data. Therefore any additional delay used to centre the lag spectrum at lag=0 must take this into account and delay by samples not bits.

The optimum place to provide these extra delays has been discussed but has not yet been absolutely decided. If method 1 is used to provide the reconfiguration delays then a delay unit is required at the input of each concatanation of modules. In the case where each of the 8 modules in a block is used to measure a separate correlation function, 8 separate delay units would be required. In the case where modules are concatenated the size of the delay unit used would need to be adaptable to cope with the length of the concatenation.

If method 2 is used, since the reconfiguration delay required for each module is independant of the number of modules concatenated, the required delays can be built into the modules themselves. The current suggestion, since the size of the delays is a multiple of the input word size, is to actually fabricate the necessary delay elements into the XCELLs themselves together with suitable control lines to select input or output delay and the size required for the selected reconfiguration factor K. This may appear to be a waste of delay elements in that half of the XCELLs which would require neither input nor output delay, but the integration cost for the modules would be reduced by not having to add the extra delay and switching circuitry using SSI/MSI chips on the module

boards. On the other hand, though, we must be sure that adding these delays on the XCELL chips does not make the chip size so large that yield is adversely affected.

In respect of the central lag compensating delay, the logical place to to put this seems to be just before the distribution network which is placed after the recirculation memory or its bypass path on a per-telescope basis. Because all baselines sharing a common set of input data require the same configuration of correlator modules, it is easy to see that they also require the same central lag compensating delay. If the exampled module configuration above is used this delay is then required in the S data path and should be placed just before the S distribution network of each telescope/I.F. data path.

4.5 XCELL COSTS AND PROBLEMS

4.5.1 XCELL Yield

Although the subject of VLSI wafer yield may have major impact on the economics of the correlator, it is an area of great uncertainty. This is due to the fact that wafer yield is the dominant factor in fabrication industry profit and is thus a secretive issue.

Initial calculations assumed a 2-bit, 8x8 Xcell implemented with a 5 micron process, resulting in a 7mm x 7mm chip. Most fabrication lines guarantee a defect density between 1 and 10 per square cm, 10 being used for worst case considerations. A commonly used but naive yield model is:

Yield = exp(-Chip Area * Defect Density)

The result of the initial calculation was:

Yield = exp(-0.7 * 0.7 * 10)
= 0.7 %

Assuming a total of around 30 blocks, the cost of Xcell chips amounted to

XCELL Cost = 30 blocks * 8 modules per block * 16 XCELLS per module
* $500 per wafer / (120 chips per wafer * 0.007 yield)
> $2,000,000

This (admittedly worst case) cost of Xcell chips alone is obviously impractical, and imposes the necessity of examining 3 areas:

4.5.2 The Defect Model

There are many possibilities, depending on the degree of accuracy required, and the types of defects being simulated. Some examples:

1. Yield = exp(-AD) => 0.7% for 8x8, 5 micron cell. Assumes a "delta function" distribution of defect densities

2. Yield = exp(-(AD)**0.5) => 10%. The Seeds yield model used by the VLSI group at Adelaide

3. Yield = ((1-exp(-AD))/AD)**2 => 4%. Assumes a triangular distribution of defect densities.

where A = Chip area in square cm.
D = Defect density per square cm.

The last model is probably the best on physical grounds and seems a good compromise between the two other common alternatives at low yields.

4.5.3 Use Of Redundancy

This involves building extra correlators into the Xcell, and hoping that the added chip area is favourably countered by the ability to bypass one or more defects by choice of connections.

Using model (a), and assuming a perfect Poisson distribution of defects, 1 in 8 and 2 in 8 redundancies add 3% and 8% respectively to the initial yield of 0.7%. This is likely to be grossly optimistic however, as it assumes that defects are uniformly distributed over a wafer. If this does not happen to be the case, e.g. if defects are highly clumped, redundancy is less likely to be applicable.

Other assumptions such as:
1. Redundant areas are functionally independent,
2. Only a small area of the chip cannot be made redundant,
3. A suitable architecture can be developed,
may not be realistic.

It is thus seen to be unwise to rely on redundancy as the only yield booster.

4.5.4 Reducing Chip Area

Since it seems unwise to rely on redundancy, the only course of action is to reduce the chip area. Two directions are considered:

1. Reduced chip function i.e. 4x4 instead of 8x8 Xcell.

2. Smaller geometry i.e. 5 -> 3 micron process

4.6 COST OF INTEGRATED XCELLS

The total costing for the various options is summarised here:

XCELL size	8x8	4x4
Total Cost = Number of Blocks * Block Cost + XCell Mask Cost		
Block Cost = Modules per Block * Module Cost + Block Integration Cost		
Module Cost = XCells per Module * XCell Cost + Module Integration Cost		
XCell Cost = Package Cost + Bonding Cost + Wafer Cost / (Chips per Wafer * Yield)		
Number of Blocks	30	30
Modules per Block	8	8
Block Integration Cost	$500	$1000
XCell Mask Cost	$80,000	$80,000
XCells per Module	16	64
Module Integration Cost	$800	$1600
Package Cost	$20	$20
Bonding Cost	$10	$10
Wafer Cost	$500	$500
Chip Area, sq. mm (5 micron process)	50	12
Chip Area (3 micron process)	18	
Chips per Wafer (5 micron process)	120	600
Chips per Wafer (3 micron process)	380	

The cost of integrated Xcells per block will be calculated for three cases, assuming defect model (c).

	Area	Yield	XCELL Cost	Module Cost	Block Cost	Total Cost
8x8 5 micron	50	0.04	$104+30	$2144+800	$23500+500	$720K+80K
8x8 3 micron	18	0.22	$6+30	$576+800	$11000+500	$345K+80K
4x4 5 micron	12	0.34	$3+30	$2112+1600	$29700+1000	$921K+80K

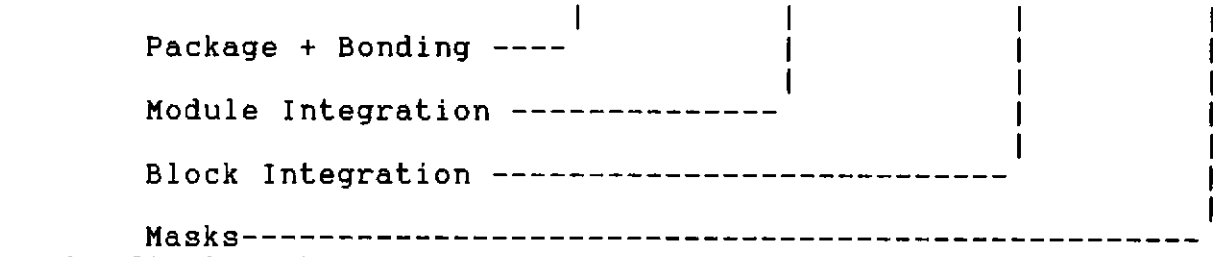

The final costs are:
8x8 5 micron => $800K
8x8 3 micron => $425K
4x4 5 micron => $1000K

The conclusion is clearly that an advanced process such as the 3 micron process is essential, otherwise a significant reduction in correlator capacity would have to be accepted.

4.6.1 Time Scales

The XCELL chip lies on a critical path for the correlator system. Given that considerable uncertainty still exists as to questions of yield and required functionality of the chip, it was felt necessary to identify a last decision date for the XCELL.

The first baseline of the correlator is required in April 1988 with the full correlator available by the end of 1988. Counting back, we obtain the following relevant dates;

1. December 83 - Test of 2-bit correlator structures on AUSMPC (multiproject chip).
2. April, August 84 - Further test structures on AUSMPC runs. In particular, tests of 3 micron structures will be required.
3. End 84 - Final decision date for XCELL including selection of fabricator. Alternatives such as 8x8, 4x4, gate arrays etc. must have been considered by this time.
4. 1st half 85 - Production by chip fabricator.

CHAPTER 5

TELESCOPE BASED CORRELATOR COMPONENTS

This chapter concentrates on a number of sections of the correlator which are telescope based. A number of sections of the correlator are treated here in detail for the first time.

Of particular importance here is the revised costing of the recirculation memory which indicates the necessity for reconfiguration of the correlators in place of recirculation for the highest bandwidths. Apart from the additional correlator complexity, this has implications for future extension of the correlator with Fourier Transform techniques (Appendix A).

One section, shown on figure 3.1, but not discussed during the workshop is the demodulator block shown immediately after the delay tracking. This block is intended to provide demodulation for any phase switch modulation applied at the telescopes. A further aspect is the possible necessity of correlator blanking during switching of noise sources etc.

The digital filters proposed elsewhere are considered and are seen to be a significant extra cost for the correlator.

To allow many modes of operation (see Astronomical Requirements) the following system of IF signals is proposed. Any increase in the the bandwidth of the radio links will require alterations to the system proposed.

From each of the compact array antennas the following bit streams will be selectable on each IF channel:

1. 160 MHz, 1-bit, 320 MBaud

2. 80 MHz, 2-bit, 320 MBaud

3. A time multiplexed signal consisting of;
    1. 40 MHz, 2-bit, 160 MBaud (Continuum)
    2. 10 MHz, 4-bit, 80 MBaud (Tied Array)
    3. 20 MHz, 2-bit, 80 MBaud (Narrow band line)

The last signal mentioned is used only for 20 MHz narrow band line. For lower bandwidth line measurements, the 10 MHz, 4-bit signal for the tied array may be further filtered as discussed in section 5.3.

5.1 DELAY TRACKING - SERIAL TO PARALLEL CONVERSION

In order to process data at 320 Megabits per second it is necessary, not only for the delay tracking but also for the correlator stages downstream to convert the serial input to W parallel data streams. These streams would each have a data rate of Fmax/W bits per second (Fmax=320 Megabits per second).

Unfortunately the delay tracking has to be implemented with a resolution of 1 input bit time, not W bits. Thus the delay control effectively has two sections. The first section will be refered to as "fine delay", whilst the other will be refered to as "coarse delay". Obviously the former deals with delays of less than W bits, whilst the latter deals with delays in multiples of W bits.

The fine delay tracking problems for both the LBA as well as the Compact array are the same. W does not need to be the same during the coarse and fine delay sections, and indeed need not strictly be the same throughout the correlator. For a 320 Mbits/s input, the correlator system calls for W=32, while sections incorporating memory may well operate more slowly and require W=64,128 or 256.

It should also be noted that the delay system will be called upon to handle signals of varying numbers of bits and/or sampling rates. A related problem is the question of separation of the three multiplexed low bandwidth signal possibilities when used and the question of synchronisation of multibit signals in general.

With regard to the latter, it would appear to involve less problems for the delay system if advantage were taken of the interval between basic 5 second integration periods to perform synchronisation with the aid of a known bit sequence.

Finally, it must be realised that the multiple inputs A,B,C,D will each require a separate delay unit while any multiplexed signals such as in point 3 above, may in fact require three separate delay systems.

5.1.1 Fine Delay

There is a spectrum of solutions to this problem. At one end of the spectrum the delay control could be done at the Fmax data rate and serial to parallel conversion take place after. At the other end the serial to parallel conversion could take place as soon as possible with the minimum amount of high speed ECL and the fine delay control be performed with a very large barrel shifter of width W or some related solution.

The first extreme suffers from practical difficulties. Whilst it is possible to buy ECL MSI packages which can clock at speeds higher than 320 MHz, when you interconnect these packages and add up clock-to-output, preset and hold times as well as non-clocked SSI type gate delays the 3.125 nanoseconds of inter-clock delay required by Fmax is easily exceeded. The problem can be circumvented by using self clocking systems or by delay compensation techniques but they are awkward.

The problem with the latter extreme is best described by assuming that W=32. In this instance one requires a 32 bit barrel shifter to extract a 32 bit parallel data stream from a 64 bit parallel stream. To implement this with 16:1 multiplexors would require 64 such packages plus 8 quad 2:1 multiplexors. A possible alternative is given in the discussion of the delay tracking method.

A solution which hits on a compromise between these two is probably the ideal.

5.1.2 Coarse Delay

The LBA has a baseline two orders of magnitude longer than the Compact Array. For this reason the solutions for each are not necessarily the same. I will treat each case separately.

5.1.2.1 The Compact Array

Since the Control room is offset from the centre of the array by 1.5 kilometres (along axis). The maximum differential delay is 30 microseconds. This sets the delay (as bits) at 9600. Putting this in an algebraic form with MDDC denoting the 30 microseconds Maximum Differential Delay for the Compact array and the width W bits, yields:-

$$L = MDDC*Fmax/W = 9600/W$$

where L is the delay FIFO length. A length of 300 locations of 32 bits is required for the compact array.

5.1.2.2 The Long Baseline Array

A similar formula is relevant for the long baseline array. The baseline separations of the long baseline array elements are;

Culgoora	- Siding Springs	= 115 Km
"	- Parkes	= 321 Km
"	- Tidbinbilla	= 565 Km
Parkes	- "	= 274 Km

The maximum baseline is 565 Km but the maximum differential delay is appropriate to a pathlength of roughly 1300 Km, where it must be remembered that all signals are returned to Culgoora. The maximum bandwidth for each of the two IF channels is 10 MHz at 1-bit indicating that W=1 or 2. The total delay is then 4.3 ms maximum which requires 86.7 Kbit delay.

Some consideration could be given here to a modular design such that not all elements are required to provide the maximum delay.

5.1.3 Delay Tracking And Setting

The coarse and fine sections of the delay system must be capable of initialisation to a particular delay setting within a resolution of one sample interval. It must track the delay difference alteration for a period of 5 seconds until the setting is reinitialised prior to the following integration period.

It must be remembered that the clock used to extract the data from the delay tracking unit for each antenna will not be the same as the clock bringing data into the delay tracking unit. This is a consequence of the "Unified Clock Principle" in which the clock and indeed all oscillators sent to each telescope are altered to exactly compensate for the path length differences experienced by the signal.

An idealised delay tracking system is shown in figure 5.1a where a FIFO is used to automatically synchronise the samples arriving from the antenna (plus clock) with the central system clock common to all antennas. This sytem was in fact used for the TEST interferometer at Parkes.

The interesting aspect of the unified clock technique for the present discussion is highlighted by the FIFO delay system possibility. The unified clock method dictates that a one to one correspondence must exist between samples at the antenna and samples arriving at the correlator.

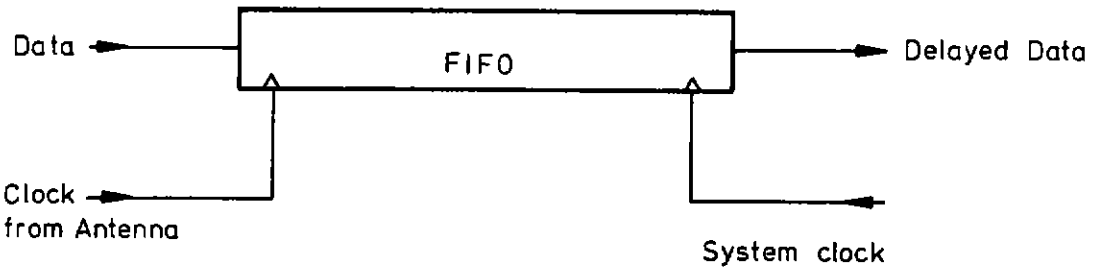

Figure 5.1a : FIFO delay system required by the unified clock technique

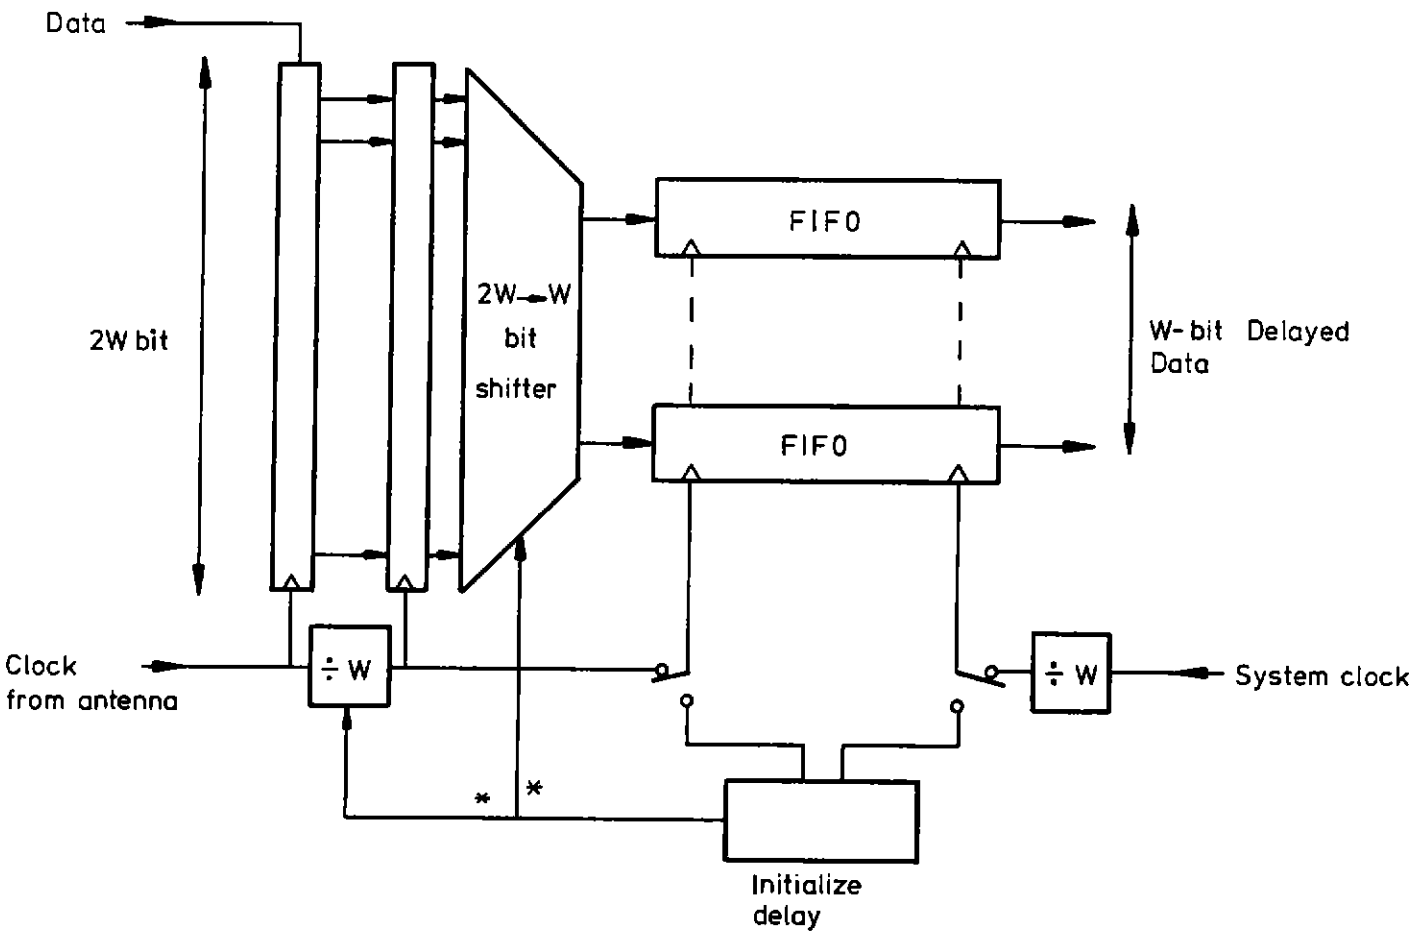

Figure 5.1b : FIFO delay system with W-bit parallel data paths

* Fine delay initialization can be performed either with the barrel shifter or by presetting the ÷W counter

The relative variation of system and antenna clocks during a period T is a quantity determined by the following,

$$L = D.\cos(\emptyset), \text{ maximum path length difference}$$

$$\dot{L} = D.\sin(\emptyset).\dot{\emptyset}$$

$$= 0.44 \text{ m/s} \quad \text{(extreme bound for compact array)}$$
$$= 44 \text{ m/s} \quad \text{(extreme bound for the LBA)}$$

Thus the maximum bounds for the frequency difference between the system and telescope clocks is given by:

$$dF = fo*(\dot{L}/c)$$

where c is the speed of light.

Thus dF = 0.47 Hz for the compact array (fo=320 MHz) and dF = 2.9 Hz for the LBA (fo=20 MHz). Over a 5 second integration time we could expect at most a 2 to 3 sample "creep" on each compact array I.F and at most a 15 sample "creep" for the LBA. It is of interest to note that the FIFO buffers are only called upon to absorb this much smaller relative change in delay occupancy. The delay could thus consist of a small layer of FIFO memory backed up by a bulk storage implemented entirely at system clock rate.

The interaction of the fine and coarse sections of the delay system is more complicated. A possible scheme is indicated in figure 5.1b where two options are shown. The fine delay difference is set either with the aid of the 64 to 32-bit shifter or, more simply, by setting the relative phase of the ÷W counters controlling the serial to parallel conversion.

Since a one to one correspondance between samples at the telescope and at the correlator must be maintained, the fine delay shift at the serial to parallel conversion stage must be maintained until the delay is reinitialised. The fine delay tracking is then automatically maintained as the telescope clock varies.

The delay initialisation then becomes relatively straightforward but is of course, limited to integer multiples of a sample period. The steps in the initialisation are;

1. The fine delay (delay modulo W) is set in the shifter or by presetting the ÷W counter at the input.

2. The coarse delay (delay/W) is set by loading the FIFOs to the required depth.

5.2 THE RECIRCULATION MEMORY

5.2.1 The Recirculation Technique

The recirculation technique will be used in combination with reconfiguration of the correlator modules in order to obtain full utilization of the correlator capacity at lowered input bandwidths. The recirculation technique allows a high speed correlator to perform a number of low speed tasks sequentially. For the AT correlator, this technique will be employed to generate more frequency or lag channels as the input bandwidth is lowered.

Figure 5.2 shows a recirculation system applied to a single correlator. The double input buffers are employed to allow the input data to be played back at the maximum speed allowed by the correlator. The input memory is partitioned into two symmetric blocks. Incoming data is written from one buffer whilst output data is read from the other. The role of the two halves are then reversed when the input buffer has filled and the recirculation cycle is complete.

The slow input rate allows time for the output to the correlator to be repeated a number of times equal to the total recirculation factor. Each complete read of the input buffer corresponds to a correlator integration period or dump time and is performed with a different delay between the delayed (D) and undelayed (U) outputs of the recirculation memory. In this fashion, the total number of lag channels measured can be increased by the same recirculation factor.

At the output of the correlators the integrated results must be stored in distinct locations for each of the lag offsets introduced by the recirculation. This part of the recirculation memory will not be furher considerd in this section

5.2.2 Recirculation Memory Size

It is possible to calculate the size,the length, and the width of the recirculation memory based on a few parameters. These parameters are:-

1. K : The "reconfiguration factor" - the ratio of maximum correlator bandwidth to bandwidth obtained by reconfiguring or reconnecting the inputs and outputs of the correlator modules to extend the range of lags produced by the correlator without requiring multiple accesses of the recirculation memory.

2. R : The "recirculation factor" - the ratio of the output data rate of the recirculation memory to the input data rate.

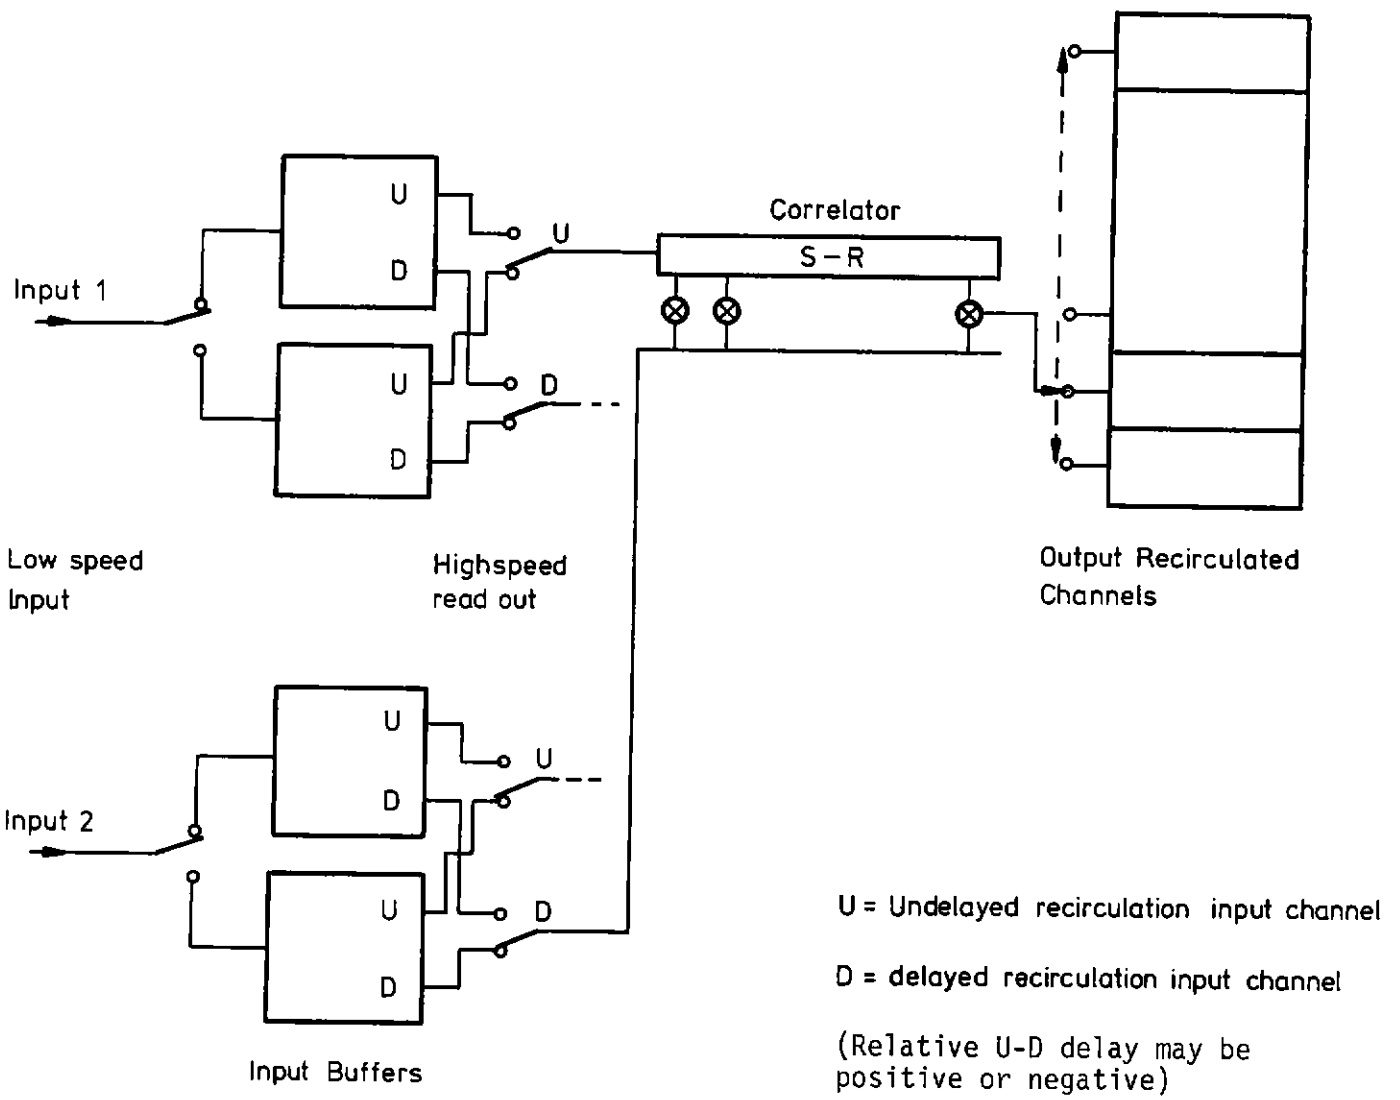

Figure 5.2 : Block diagram of a recirculated correlator system

3. U : The "utilisation factor" - the effective total useage of the correlators for lowered bandwidths, U = K.R.

4. Fmax : The peak data rate of the system. This value is set at 320 megabaud.

5. a : The memory access cycle time.

6. n : The number of data elements required by the correlation system during each correlation cycle. With the current ideas on the XCELL and module design n is two, corresponding to the delayed and undelayed recirculation memory outputs.

7. t : The time between correlator dumps.

The (minimum) size(in bits) of each half of the recirculation memory can be calculated by the formula:-

$$S = Fmax*t/K$$

This is based on the fact that the output data rate of the recirculation memory is Fmax/K and that the whole of the memory needs to be read in the correlator dump time.

We can calculate the width of the recirculation memory based on the output data rate and the memory access time.

$$W = n*(Fmax/K)*a$$

where Fmax/K defines the output bit rate and the width is proportional to the speed of the memory a and how the number of data elements the correlator requires in each correlator cycle (n).

The minimum length of the memory is simply,

$$L = S/W = t/(n*a)$$

The first thing to observe here is that the recirculation factor R has not affected the size, width or length of the memory. Note also that the length is independent of the reconfiguration factor K. It is attractive then to consider the prospect of bypassing the recirculation memory until K reaches its maximum.

Size of recirculation memory for possible recirculation and reconfiguration factors.

U	R	K	S	W	Size(in Kbytes)
1	1	1	6.4M	256	1024*
2	1	2	3.2M	128	512*
4	1	4	1.6M	64	256*
8	1	8	0.8M	32	128
16	2	8	0.8M	32	128
32	4	8	0.8M	32	128
64	8	8	0.8M	32	128

where,

t = 20 ms,
n = 2,
a = 400 ns,
L = 25000.
* is only appropriate if recirculation is used

The size quoted in the table assumes that the memory length has been taken to the next power of 2 above 25000(ie 32*1024). We can calulate the total recirculation memory requirements of the project as being,

Total memory = 2*(No of telescopes)*(No of IFs)*1024/K kilobytes of memory.
= 2*6*4*1024/8 kilobytes (assuming no recirculation with K < 8)
= 6.144 Megabytes for the compact array.
= 2.048 Megabytes for the LBA.

The initial "2" is due to double buffering.

If t is increased to 50 milliseconds we double the memory requirements. The width is independent of t but the minimum length now becomes 62500. This is quite close to 64*1024. It also has the advantage that it conforms to some industry standard configurations.

The access cycle time a is very important. The memory cycle time including bus arbitration, address decoding, memory refresh etc is included in the 400 nanoseconds mentioned above.

To avoid refreshing it might be possible to use an suitable addressing scheme. The recirculation memory consists of two buffers, i.e. the input and output buffers. The input data rate is related to the output data rate by R (by definition). The input data rate thus is

= Fmax/(K*R) [output rate = Fmax/K]

When a single location of a 16K or 64K dynamic ram is read, 128(16K) or 256(64K) locations are actually refreshed. If one can arrange the addressing so that adjacent memory locations are physically accessing unique blocks of refreshed memory and that in the time it takes to fill or empty the input or output buffer all locations are refreshed then a refresh controller is not required.

Even with this method, the avoidance of refresh places constraints on the size of the memory. If we assume that the memory must be refreshed every 2 milliseconds then we can write

2 millseconds >= (L/256)*(W/(Fmax/(K*R)))
for the input buffer.
and

2 millseconds >= (L/256)*(W/(Fmax/K))
for the output buffer.

wher L/256 is the number of cells (256 wide) refreshed by a row access in a 64K dynamic memory.

Obviously the most critical of these two is the input data rate equation. We have a formula for L and it should be sensible to bring this into the inequality. If we assume that W is 32 this also reduces the problem. For the worst case R=8.

2 milliseconds >= ((t/(n*a))/256)*(8*8*32/(320*1000000))
N.B. K should be 8 before any recirculation

Assuming n=2 and a =400 nanoseconds this reduces to

t =< 64 milliseconds.

5.2.3 Revised Costing Of Recirculation Memory

The recirculation memory as described comprises a very large volume of high speed random access memory. If it is assumed that commercially available memory boards are used, then the total memory cost will be of the order $35K for the compact and long baseline arrays.

It might be emphasised however that such factors as the memory cycle time and the ability to reconfigure the correlators are critical to the memory size and hence total cost calculation. For clarity, the parameters used are repeated,
1. correlator dump time t = 50 ms,
2. memory cycle time a = 400 ns,
3. reconfiguration to K = 8 before recirculation,
4. 64Kbit memory chip based.

With these parameters we obtain the final recirculation memory size;

Size of recirculation memory for anticipated recirculation and reconfiguration factors.

U	R	K	S	W	Size(in Kbytes)
1	1	1	16.0M	256	2048*
2	1	2	8.0M	128	1024*
4	1	4	4.0M	64	512*
8	1	8	2.0M	32	256
16	2	8	2.0M	32	256
32	4	8	2.0M	32	256
64	8	8	2.0M	32	256

where,
L = 62500.
* is only appropriate if recirculation is used

It may be noted that the choice of memory size and width as shown in the table above has significant impact on the possibility of using transform methods to expand the capacity of the correlator in the future. Ideally, the recirculating buffer design should provide for future expansion.

5.3 COMPENSATING DELAYS AND DATA DISTRIBUTION

Delays to compensate for the shift in central lag which occurs in the correlator modules are logically placed at the output of the per-telescope data paths in that data stream which is destined to reach the S inputs of correlator modules. This point is just after the recirculation memory and just before the distribution buffers which are required to change the data

streams from per-telescope to per-baseline. The reasons for this were briefly discussed at the end of section 4.4.3.

5.4 TIED ARRAY

The objective is to configure the Compact Array so that it appears as a single telescope. This implies that the six telescope outputs must be added in phase with correct delays. Phase and delay corrections for two fequencies at each telescope must be determined by source observation. Provision must therefore be made for applying these corrections. The normal phase and delay tracking provided for array operation will automatically ensure that the Tied Array beam will be in the correct position. It should be noted that the Tied Array beam is a fan beam.

In order to tie the telescopes together we must provide 4-bit signals to avoid loss of sensitivity due to quantization noise. The output from the Tied Array will be 2-bits.

The bandwidth required for the Tied Array is 10 MHz to correlate with the remote antennas whose bandwidth is restricted by the radio link. Thus the two signals to be correlated are digitized at 10 MHz 2-bits.

Because simultaneous operation at two frequencies is required it will be necessary to have two Tied Array outputs; one at each frequency.

The tied array summing unit is required to perform a sum of up to 6 signals each 10 MHz and 4-bit. Antennas must be capable of elimination from the sum by switching controlled centrally. In addition, some means of IF selection (the A,B,C or D inputs) should be provided.

The summation will give a 7-bit result in the extreme case. The result may be converted into a 1-, 2-, or 4-bit result most flexibly with the use of a look up table which may be loaded from the central control. The choice of outbut bits will depend on further use of the summed signal. For example, if the summed signal is required as input to the digital filter stages (assumed present for the moment), then a 4-bit sum will be desirable.

Detailed costing has never been performed for the tied array units. A preliminary costing, plus the fact that only two units are required indicates that the tied array option should present no budgetary problems.

5.5 DIGITAL FILTERING.

It is proposed to further filter the 10 MHz 4-bit data stream in order to produce selected bandwidths of: 5, 2.5, 1.25, 0.625 MHz. It was originally planned to use a digital filter at the inputs to the correlator to perform this filtering. The questions that now have been raised are:

1. Can the filtering be done digitally at this frequency?
2. Are digital filters in fact necessary, and if not, would they really offer any advantages over other filtering means?

The preference for using a digital filter at the central site rather than using an analogue implementation before the digitization at the telescope, was based upon the concern with the stability associated with analogue filters consistent with the desire to preserve the high dynamic range of the resultant map. Digital filtering schemes have the advantage that in principle they can be made to be very stable. However it has now been suggested that the real cause of instability in current aperture synthesis telescopes (VLA and Westerbork) is not due to inadequate performance of the filters, but rather the frequency performance of the whole telescope system, e.g. drift in IF bandpasses. Hence the concept of using analogue filters has not yet been ruled out. However if analogue filters were to be used, they should be implemented at each telescope before digitization, because the advantages of using digital transmission of the IF over the fibre-optic system are considerable. An analogue IF transmission system would involve a multiplexed IF in order to transmit both narrow and wide band data simultaneously.

Document AT/24.1/002 describes some aspects of digital filters, and suggests that it is essential to work from either 3 or 4 bit input data in order to prevent an excess of sensitivity degradation. A 3-bit input quantization results in a sensitivity degradation of 4 percent. Note however that the output has to be either one or two bits. Given that the bit data rate is 320 Mbits/sec, the natural starting point for digital filter is therefore the 10 MHz bandwidth input stream. The Digital filter could of course be placed at the telescope so that even higher bandwidths might be possible.

For analysis of the feasability of a digital filter the following filter parameters will be used:

1. Passband flatness of 1dB
2. Stop Band rejection of 40dB
3. Transition frequency range of order of 0.1 of the filter bandwidth

An empirical formula for the filter tap length is given in Crochiere and Rabiner (Proc. IEEE, Vol. 69, No. 3, pp300-331, 1981) for the design of filters for decimation. From this can be obtained,

$$N = 2(A.D)/C$$

where,

A is a factor dependant on the required flatness and stopband rejection, and is approximately equal to 2 - 2.5
D is the required decimation factor,
C is the transition frequency rage.

The minimum bandwidth of 0.625 MHz requires a maximum reduction ratio of 16. This leads to a filter with the characteristics shown in the table.

Digital Filter Parameters for 10MHz Input Bandwidth.

Decimation Factor	Input Rate	Output Bandwidth	Output Rate	Number of Taps
16	20 MHz	0.625 MHz	1.25 MHz	700
8	20 MHz	1.25 MHz	2.5 MHz	350
4	20 MHz	2.50 MHz	5.0 MHz	175
2	20 MHz	5.00 MHz	10.0 MHz	87

The above figures assume that a low pass digital filter is required. For bandpass implementations the number of filter taps is doubled to obtain the same bandwidth. Practical implementations would also tend to use the nearest power of two taplength. For this reason, some easing of the criteria noted above should be possible to allow a maximum length of 512 taps.

The filter structure described in AT/24.1/002 is applicable to the problem although it was designed for the more extreme case of 320 Msamples/sec input rate. The structure is shown in figure 5.3.

The filter characteristics are determined by loading the digital filter coefficients, or rather the precomputed products of inputs with the coefficients into the memory of the filter. For extreme bandwidth reduction (decimation) factors, the full tap length must be implemented, and the full adder tree is necessary. The entire adder tree works at the minimum output rate noted in the table above.

As successively smaller decimation factors are required, the adder tree becomes split into 2 X 256 input, 4 X 128 input, and so forth. The increased output rate is absorbed by using the split adder trees to function in a time-multiplexed fashion on successive output samples.

Straightforward costing of the system yields $5K - $7K. The reduced bandwidth requirements mean that the adder tree of figure 5.3 is capable of operating at many times the speed required. An alternative to the full parallel implementation is shown in figure 5.4 where a single adder tree is employed to

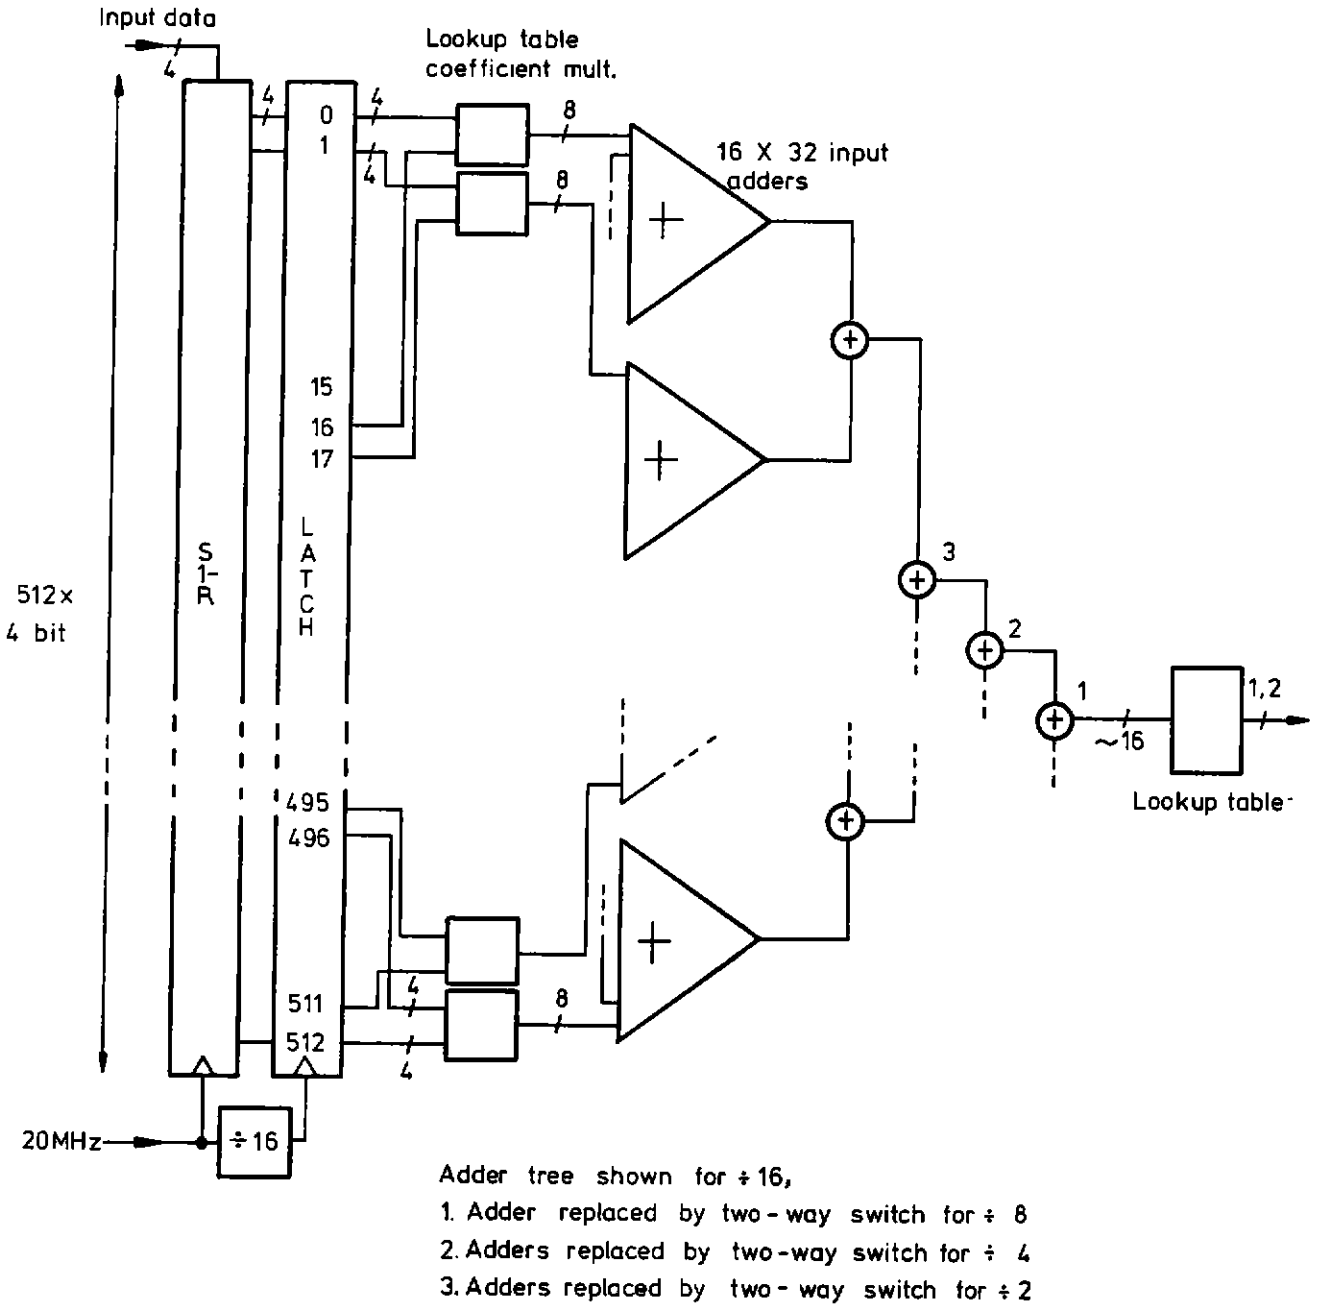

Figure 5.3 : Fully parallel implementation of a decimation digital filter

perform the required sum sequentially. The reduction in component count should reduce the cost to $2K per filter In total however, (4 IFs X 6 telescopes) filters are required for the compact array and (2 IFs X 4 telescopes) filters are required for the LBA system, yielding a total of $64K. This has not been costed within the correlator previously.

For comparison, if the filtering is done at the antenna there will be some extra electronics required for:

1. Further conversion to a suitable frequency for filtering; either baseband or passband.
2. Time multiplexors or 4 extra fibres to transmit both wideband and narrowband simultaneously. However, the extra cost of fibres, transmitters, and receivers would be prohibitive.

We should therefore allow total cost of about $5000 per antenna for the four IFs required. To this we have to add the filter costs which are estimated as ($150 X 9 filters X 4 per antenna) or $5,400. So the total cost of the analogue filter system is of the order of $62,400. It should be noted that this has not been costed within the receiver budget either.

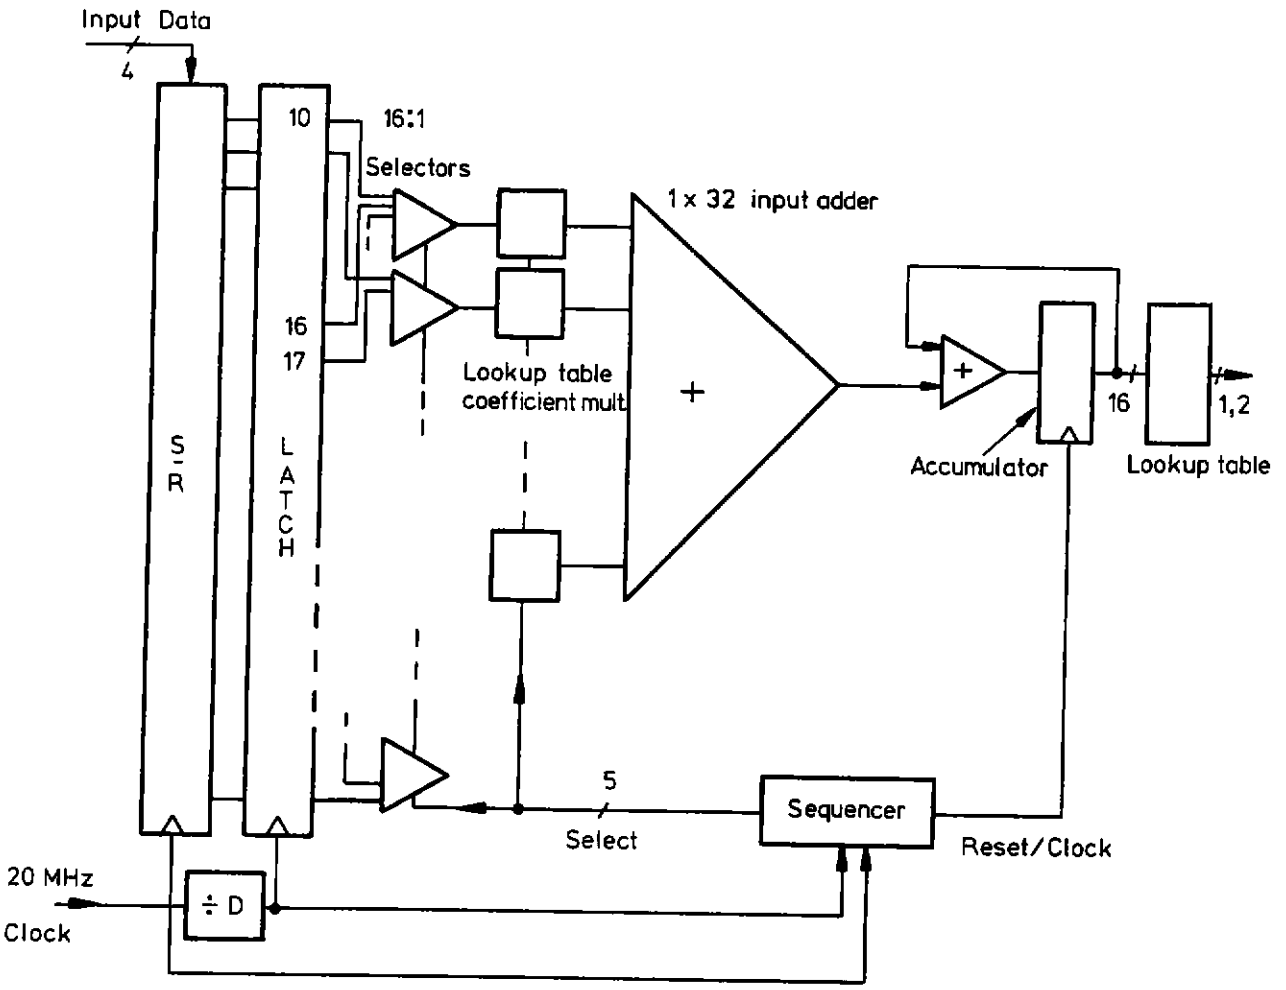

Figure 5.4 : Serial-Parallel implementation of a decimation (factor D) digital filter

APPENDIX A

FFT METHODS TO INCREASE THE NUMBER OF CORRELATOR CHANNELS

The total number of channels for the AT correlator system is ultimately limited by cost considerations, particularly those due to the XCELL chip (Chapter 4). The possibility exists that at some point in the future the effective correlation capacity of the correlator could be increased by an add-on processor unit which preprocesses the data input to the correlator.

The correlator is required to measure the average product of two sampled signals for a range of lags or delay differences. This procedure is functionally equivalent to the performance of repeated convolutions and requires a number of multiplications given by,

$$P = fo.N \text{ per second}$$

where the sampling rate is fo and N lag channels, ie. N/2 frequency channels are required. It is assumed here that the multiplication is performed at the frequency fo. If this is not the case, then more multipliers are required in direct ratio to the multiplier frequency and the total effective number of multiplications remains constant.

It is well known that if the input signal is first split into M bands and sampled at appropriately reduced rates, then the required frequency channels can be determined (with some edge effects). The number of multiplications becomes,

$$P = fo/M \text{ per second} \;.\; N/M \text{ per band} \;.\; M \text{ bands}$$

$$= fo.N/M \text{ per second}$$

We could therefore place a filter bank in front of the correlator and boost the correlator performance by a factor M. Unfortunately, this runs into several problems;

1. The correlator must be split into M blocks in some way.
2. Full correction of the effects that an ideal filter bank would produce is impossible.
3. A digital realisation of an ideal rectangular filter is complex (see Chapter 5) and imperfections will introduce further problems.

The first of these problems can be solved with suitable use of the recirculation and reconfiguration techniques. The second and third represent sufficient reason to drop the filter bank method and to examine other possibilities.

The observation that the correlation is functionally equivalent to a convolution suggests that several methods may be used. Of these, the FFT method is a very good possibility. That and others are described in O'Sullivan J.D., "Efficient Digital Spectrometers - a survey of possibilities", Netherlands Foundation for Radio Asronomy, Note 375.

The FFT method uses an FFT to efficiently compute a convolution. That is, instead of requiring all possible products of samples displaced by a given lag, only products of corresponding frequency points of the FFT are required. The number of correlator operations becomes,

$$P = 4 \text{ mult. per complex product } . M \text{ complex outputs } . N/M \text{ channels per product}$$

$$= 4fo.N/M \text{ per second}$$

where a 2M point transform which produces M distinct complex outputs has been assumed. The complex product requires four real products (or 3 if the transform output is suitably modified). Only M complex outputs are distinct since the real input implies a transform output with positive and negative frequencies as complex conjugates.

The total number of lag channels produced is N and the Fourier transform method is a factor 4 worse than the ideal filter bank method.

Figure A.1 shows a direct implementation of the FFT method. Note that the serial to parallel conversion and FFT are both telescope based. The data rates are all reduced by a factor M where the transform is a 2M point FFT. The obvious disadvantage with this method is that the correlators are required to be split into 4M groups of N/M channels. This method might be termed the correlator speed reduction mode.

The correlators are operated at a considerably reduced speed so the oportunity exists to exchange speed for a small group of channels to produce a corresponding increase in the number of channels. The recirculation technique is ideally suited for this task although the reconfiguration technique is equivalent but limited.

Figure A.2 shows the recirculated or channel increase mode of operation. The inputs are Fourier transformed and stored in the recirculation memory. Each of the M complex frequency points is read out for the entire correlator integration period and the M points are correlated in turn.

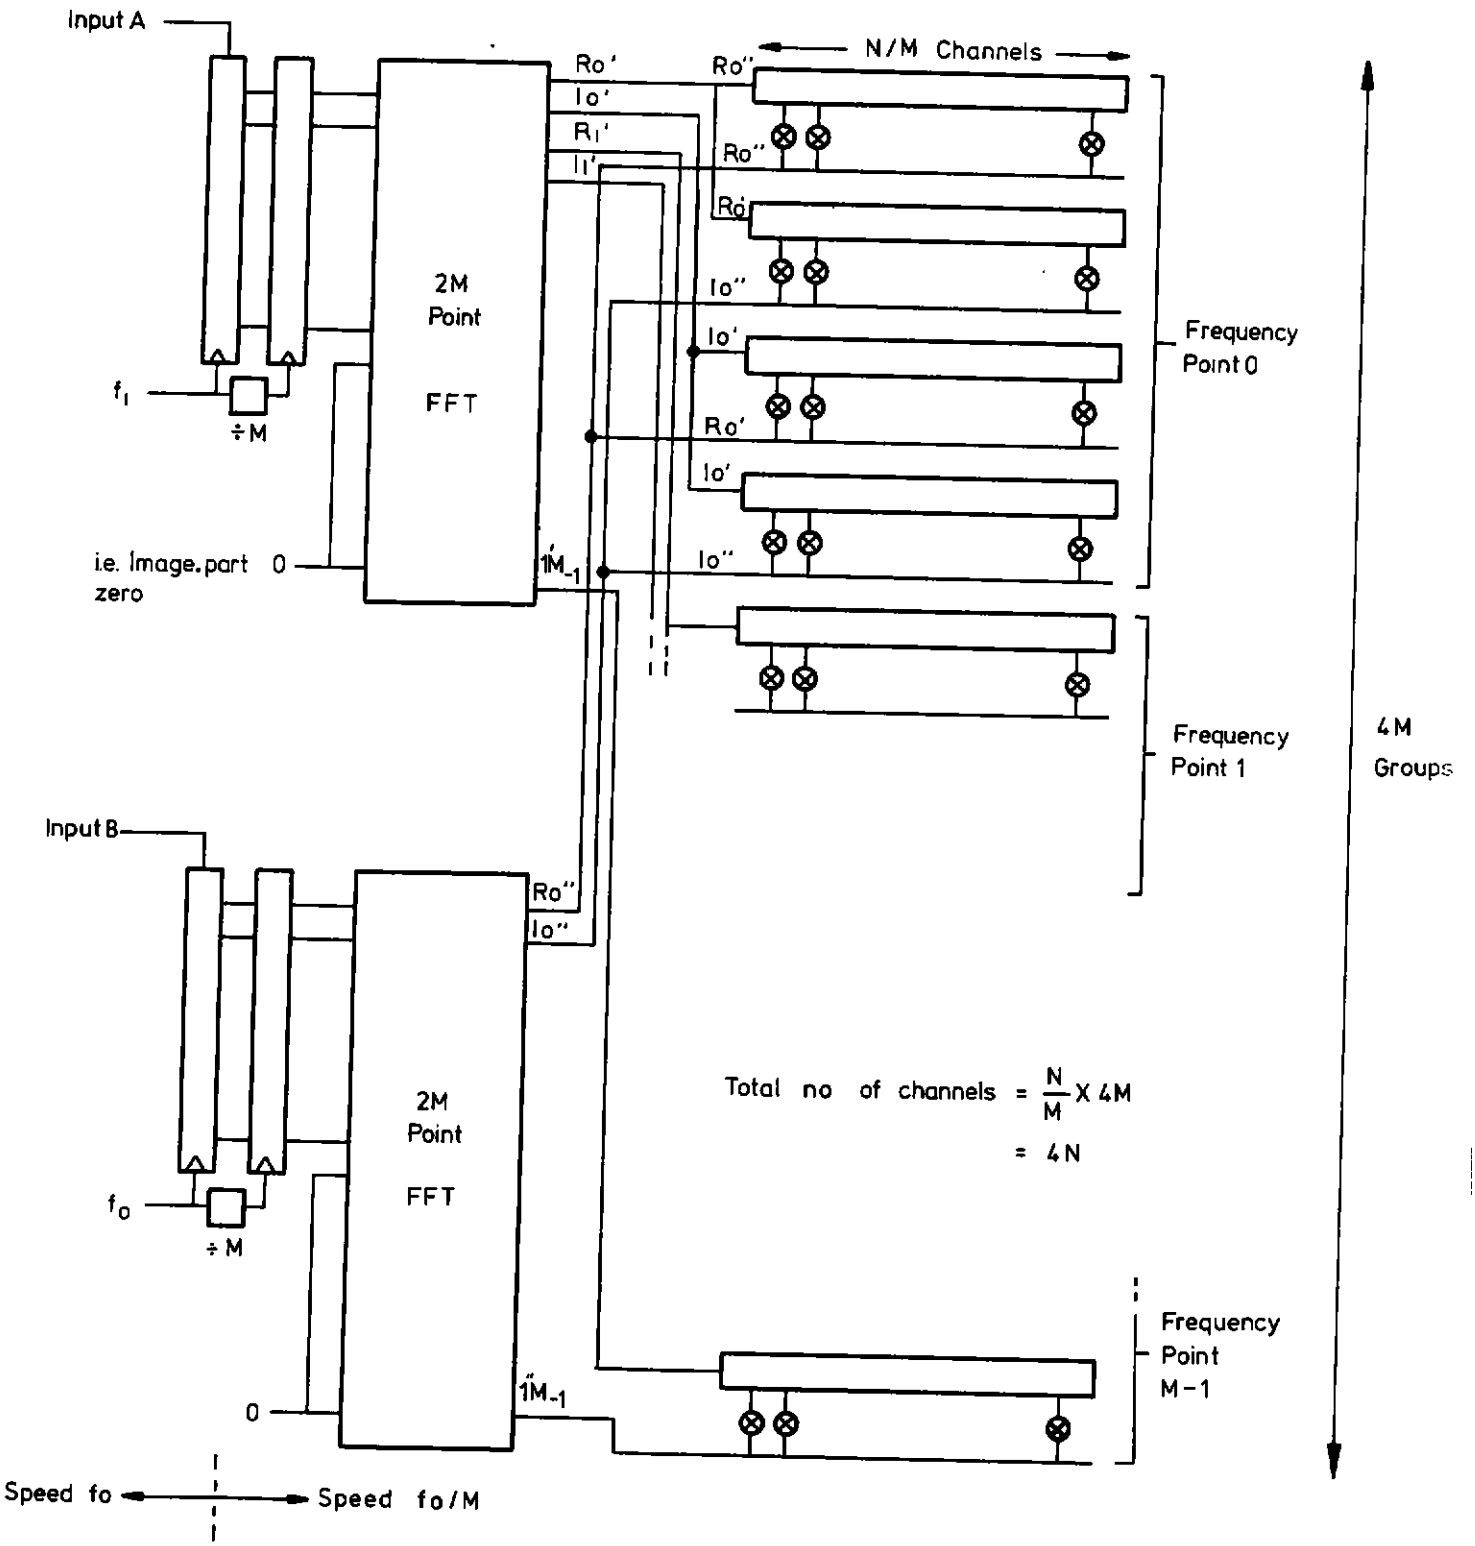

Figure A.1 : The FFT technique applied in the correlator speed reduction mode

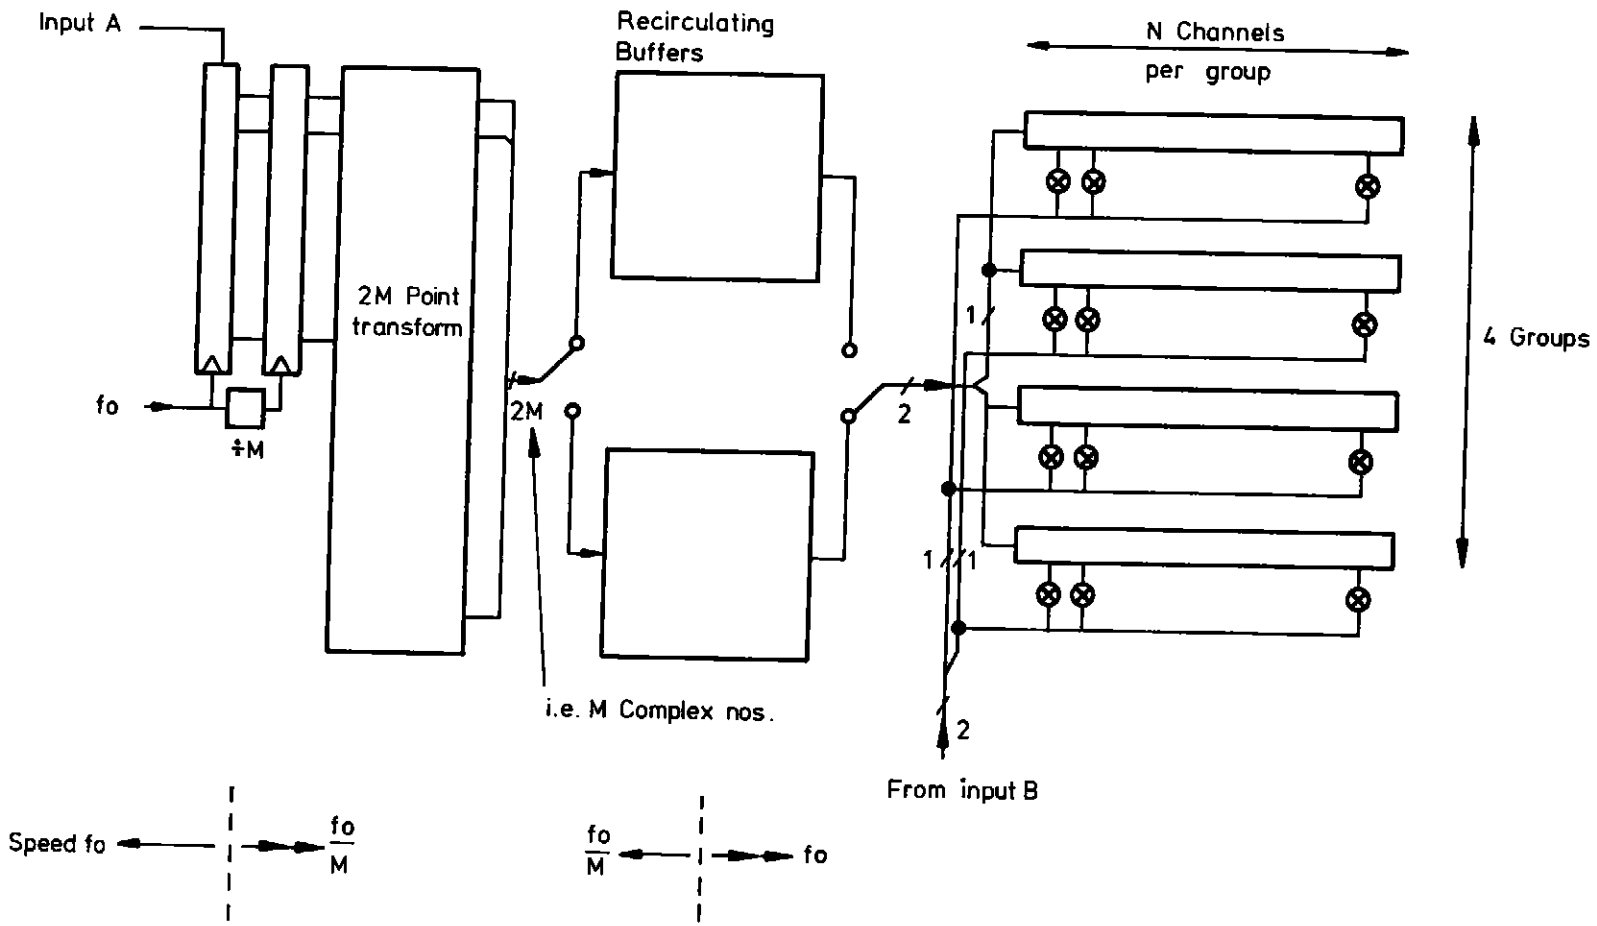

Figure A.2: The FFT technique applied in the channel increase mode

flexible addressing capabilities in the recirculation memory specifications.

The usefulness of transform methods will depend considerably on the expense of performing FFTs on high speed data streams. Were the transform free, then obvious advantages accrue from choosing M as large as possible. Given that the number of input and output bits is limited, potential exists for extensive use of lookup table methods. A combination of lookup table and CORDIC methods (Despain A.M., IEEE Tr. Computers, Vol. C-28, No. 5, May 1979) might for example allow a 2M=48 point FFT at 10 MHz rate. This would employ a small number of VLSI circuits.

AT present the FFT methods appear to give the potential for cheap upgrade of the AT correlator system capacity at some future date.

A number of points are worth noting with respect to the use of recirculation. With an FFT method, 4M products or groups of correlators are required. A recirculation cycle must be of length M.t where t is the correlator dump time as defined in Chapter 5 and M is the number of frequency points which must be separately correlated.

During a recirculation cycle, the FFT of successive groups of M input samples is performed and the complex results stored in the recirculation memory. The size of the recirculation memory is accordingly,

$$S = 2fo.M.t$$

where 2M transform results (M complex points) must be stored per M inputs.

The recirculation memory readout operates at the same frequency as the input. During each correlator dump time, one of the FFT complex frequency points is read out. Only M complex frequency points are required so time exists to perform M separate reads of 1/Mth of the data in the recirculation memory.

The complex product is carried out by splitting the correlator into four groups of N/4 channels. The ability to split the input connections as required for reconfiguration by a factor of K=2 or as required for the continuum correlator are suitable.

A major practical problem caused by the doubling of the input data rate by the transform concerns the size and speed of the recirculation memory. Given the sizes of 5.1.4.3, where 256 kbyte total memory per input channel with S=2.0E6 is suggested, the correlator dump time to just allow filling the recirculation memory is given by,

$$t = 2.0E6 \;/\; (2M \text{ points per transform} \cdot fo/M \text{ transforms per second})$$

$$= 1.0E6/2fo \text{ seconds}$$

It is evident that the correlator dump time must be reduced (by a factor 2) to maintain operation at the rate allowed by the recirculation memory which corresponds to fo=Fmax/8=40 MHz sampling rate.

A further complication is produced by the speed of reading from or writing to the recirculation memory which is double that assumed in Chapter 5 due to the need to store complex numbers. This would require a memory width twice that suggested earlier or a corresponding reducion in bandwidth of maximum operation.

Possibilties involving combinations of recirculation and reconfiguration have not as yet been examined in detail. At present,it would appear prudent to allow for expansion and