VLBA Correlator Memo No. 88

(870929)

THE VLBA CORRELATOR PLAYBACK INTERFACE

by Ray Escoffier and Gene Runion
September 28, 1987

I) INTRODUCTION

This memo will describe the playback interface system of the VLBA correlator. This system of the correlator includes what used to be called the DPC which was to be the track to channel conversion logic of the Data Playback System and what used to be called the delay system of the correlator. These two previously separate subsystems of the VLBA processor have been combined into a single function and brought into the racks of the VLBA correlator.

NRAO was greatly assisted in the design of the track recovery logic described in this memo by Hans Hinteregger of Haystack Observatory and we extend our appreciation to Hans for his useful suggestions.

II) THE RECORD SYSTEM FORMATTER

A description of the playback interface system must begin with a brief overview of the record system formatter designed by Haystack Observatory. The formatter at the record station will process the sampler output bit streams in any of several complex ways to produce track signals suitable for recording. This action is taken in order to make efficient use of tape and in order to make the data recoverable upon playback. The playback interface logic must take the playback track signals produced by the playback drive and reconstruct the original sampler streams, thus requiring that it undo what the formatter does to the sampler stream.

Figure 1 gives a basic block diagram of the data flow through the formatter. The specific operations performed on the sampler output data streams are;

1) cross bar switch to assign any sampler output to any recorder track input.

2) convert the sampler bit streams into track bit streams by either multiplexing 1, 2, or 4 sampler bit streams into one track signal or, for higher rate sampling, by breaking a single sampler bit stream into 2 or 4 track signals.

3) perform a barrel roll function to provide a dynamic assignment of track signals to recorder heads.

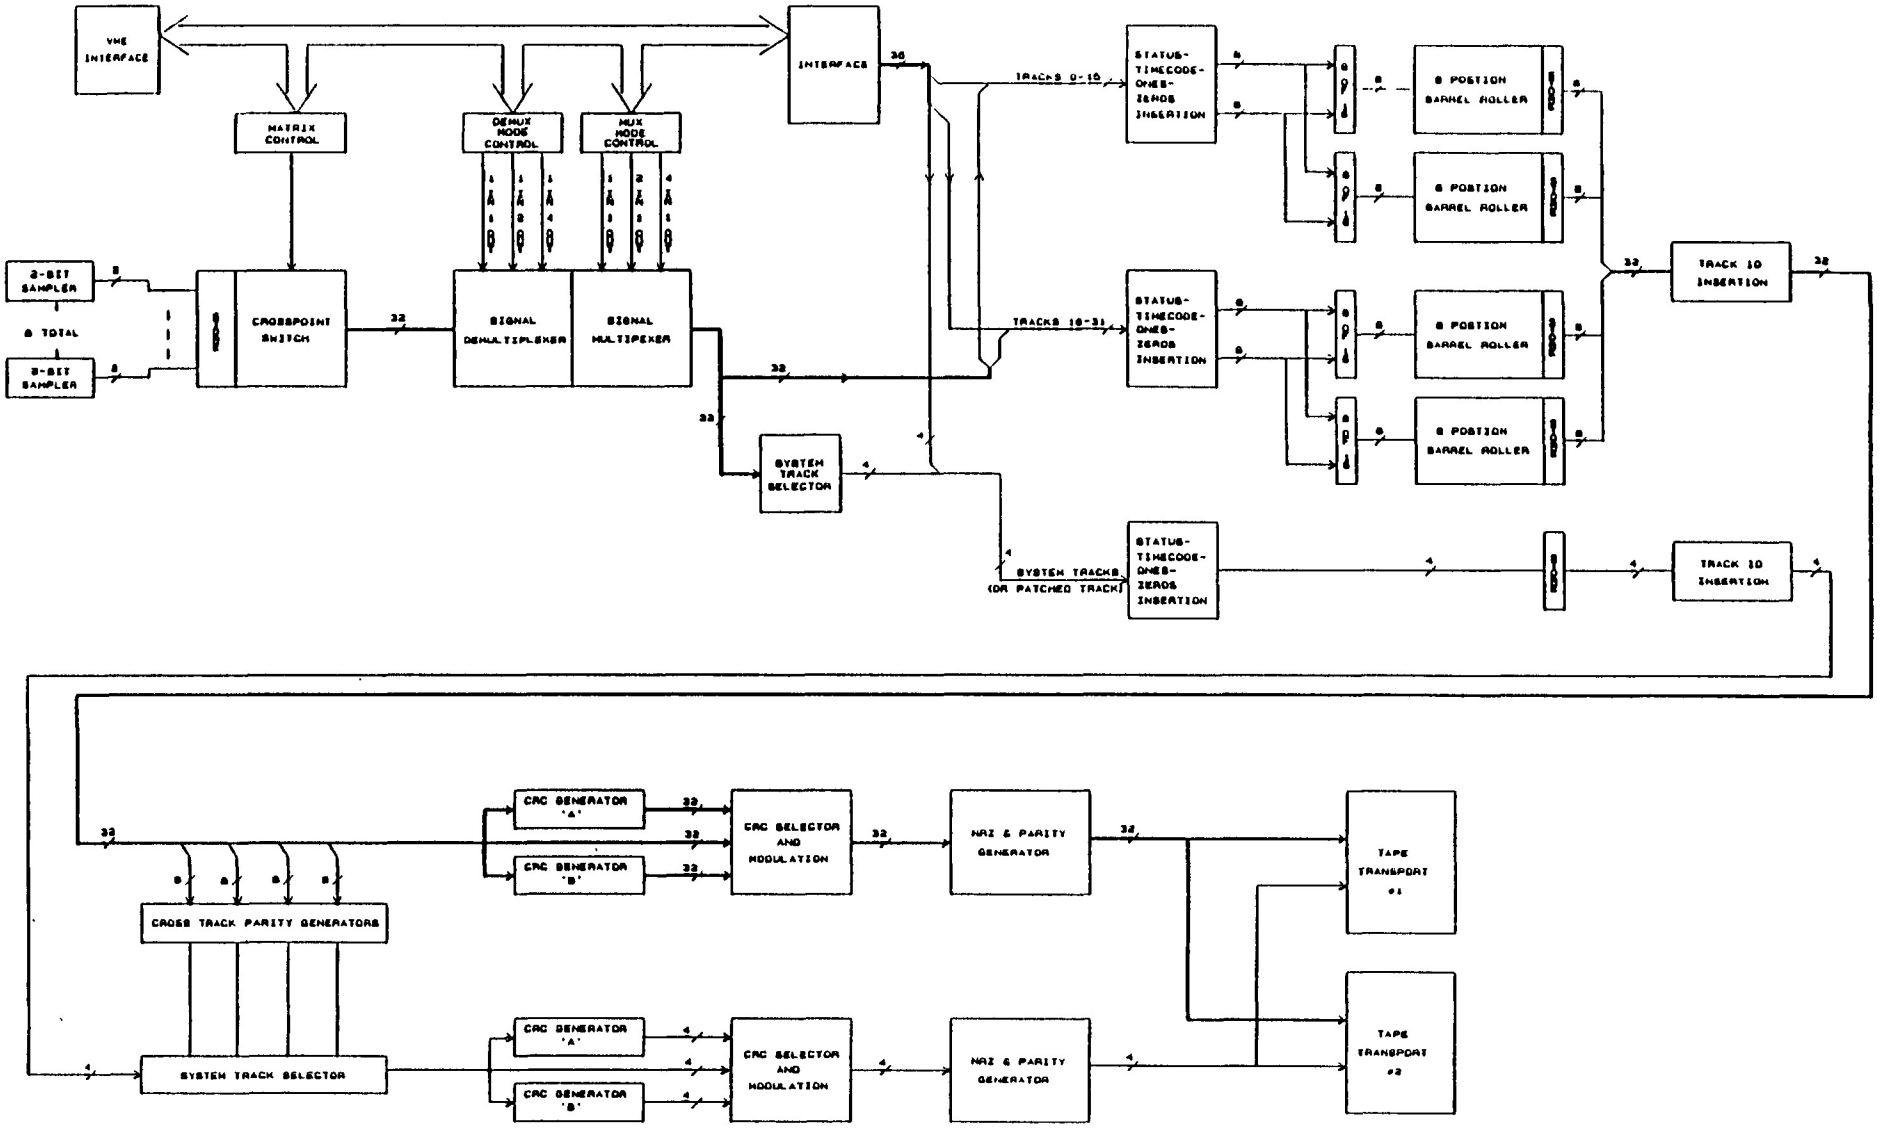
VME INTERFACE
MATRIX CONTROL
DEMUX MODE CONTROL
MUX MODE CONTROL
INTERFACE
36
TRACKS 0-15
TRACKS 16-31
STATUS-TIMECODE-ONES-ZEROS INSERTION
8 POSTION BARREL ROLLER
TRACK ID INSERTION
8-BIT SAMPLER
8 TOTAL
8-BIT SAMPLER
CROSSPOINT SWITCH
32
SIGNAL DEMULTIPLEXER
SIGNAL MULTIPEXER
SYSTEM TRACK SELECTOR
SYSTEM TRACKS (OR PATCHED TRACK)
CRC GENERATOR 'A'
CRC GENERATOR 'B'
CRC SELECTOR AND MODULATION
NRZ & PARITY GENERATOR
TAPE TRANSPORT #1
TAPE TRANSPORT #2
CROSS TRACK PARITY GENERATORS
SYSTEM TRACK SELECTOR

Figure 1

4) generate four system tracks made up either of duplicates of other selected tracks or of parity bits generated across several other tracks (across track parity).

5) insert various types of information into each track data stream on a non-replacement basis.

a) along track parity
b) frame sync pattern
c) time code
d) CRC codes
e) miscellaneous status and information bits

6) provide modulation

The two most important insertions of item 5, above, added by the formatter into a track data stream are the time information and the sync pattern. The formatter blocks the sampler outputs into frames of suitable length and tags each frame with the wall clock time at the record station when bit number one of the frame was taken. This time is essential in synchronizing the data streams from two playback drives from different stations when processing experiments and in merging several playback tracks back into a single high bit rate data stream. The sync pattern, on the other hand, allows the track recovery logic to recover the precise boundaries of the recorded frames.

All of these operations, with the exception of the cross bar switch function, require the playback interface logic to provide some kind of inverse operation.

III) THE DATA PLAYBACK SYSTEM-CORRELATOR SYSTEM DATA INTERFACE

The data input to the correlator will come from 24 playback drives. These 24 units include 20 drives that can support processing up to 20 stations of an experiment and 4 "hot spare" playback units that can be used either to substitute for defective drives or to support efficient tape changing (see section IV for more detail on the spare drive switching concept). Each playback drive will have 36 data and 36 associated clock outputs, one set for each of the 36 magnetic heads on the PBD (playback drive) head stack. Anywhere from one to all 36 of these head outputs may be active at a given time.

Each of the 36 data signals has a number of items, either the result of the action of the formatter at record time or of the PBD at playback time, that must be provided for in order to recover the original sampler data when processing an experiment;

1) each track has a wow and flutter component added to its basic 9 Mbs (approximately) bit rate, thus the need for 36 separate clock signals.

2) each track has an independent head skew of up to a few hundred bits relative to the other tracks being recovered on a given PBD that must be removed before several of these tracks can be multiplexed together to reconstruct a channel.

3) each track has bits inserted into its bit stream on a non-replacement basis. As above, these bits include;

a) along track parity
b) frame sync pattern
c) time code
d) CRC codes
e) miscellaneous status and information bits

After all of the inserted bits above and the mechanically produced variations in the track data have been removed, the result will be continuous track data streams running on an 8-Mbs system clock. These tracks must be processed further by multiplexing several tracks together to reconstruct a high bit rate sampler output or by splitting a track into several data streams to reconstruct the output of several low bit rate samplers.

IV) THE PLAYBACK INTERFACE SYSTEM

A block diagram of the playback interface is seen in figure 2. The major activities depicted in this figure include; individual track processing, across track parity processing, lost track recovery by use of the system tracks, inverse barrel roll, track to channel conversion, and the "channel mix" and "channel buffer" functions to support the FFT system.

Individual track processing is required since each track data signal has a unique clock. A block diagram of the track processing provided in the VLBA correlator for a single playback track is seen in figure 3. The function of this logic is to accomplish all of the operations (wow and flutter removal, head deskew, frame overhead removal, frame buffering, parity error count) required in converting the unstable track signal into a synchronized and continuous data signal. This conversion is shown being accomplished in a single stage. Most of the required operations cited above can be accomplished in a RAM memory with sufficient control over the RAM read and write functions. In addition, if this RAM is large enough, it can also provide the delay function that was previously assigned to a separate delay system. A brief description of the functions

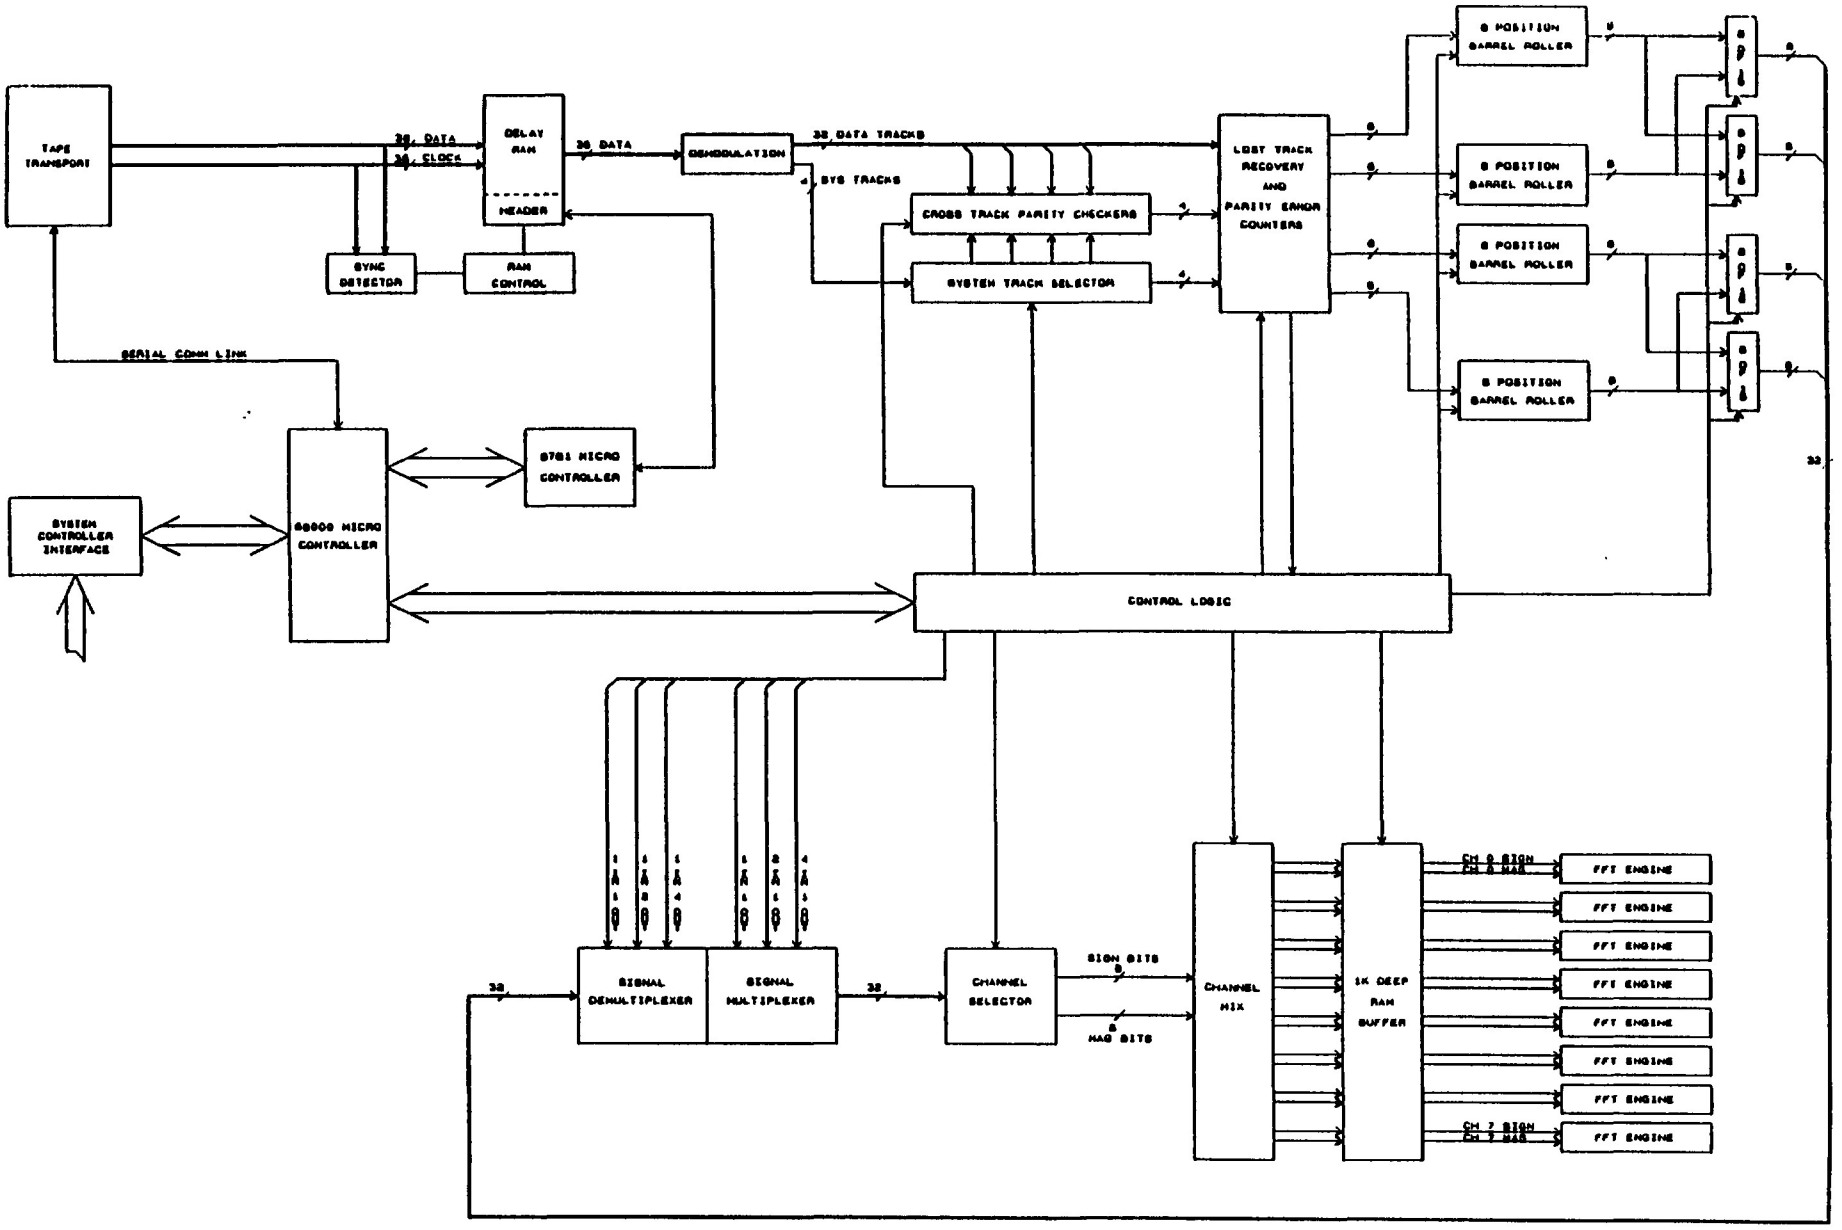
TAPE TRANSPORT
DELAY RAM
HEADER
SYNC DETECTOR
RAM CONTROL
DEMODULATION
CROSS TRACK PARITY CHECKERS
SYSTEM TRACK SELECTOR
LOST TRACK RECOVERY AND PARITY ERROR COUNTERS
8 POSITION BARREL ROLLER
8 POSITION BARREL ROLLER
8 POSITION BARREL ROLLER
8 POSITION BARREL ROLLER
SERIAL COMM LINK
8751 MICRO CONTROLLER
SYSTEM CONTROLLER INTERFACE
68000 MICRO CONTROLLER
CONTROL LOGIC
SIGNAL DEMULTIPLEXER
SIGNAL MULTIPLEXER
CHANNEL SELECTOR
SIGN BITS
MAG BITS
CHANNEL MIX
1K DEEP RAM BUFFER
FFT ENGINE
FFT ENGINE
FFT ENGINE
FFT ENGINE
FFT ENGINE
FFT ENGINE
FFT ENGINE
FFT ENGINE

Figure 2

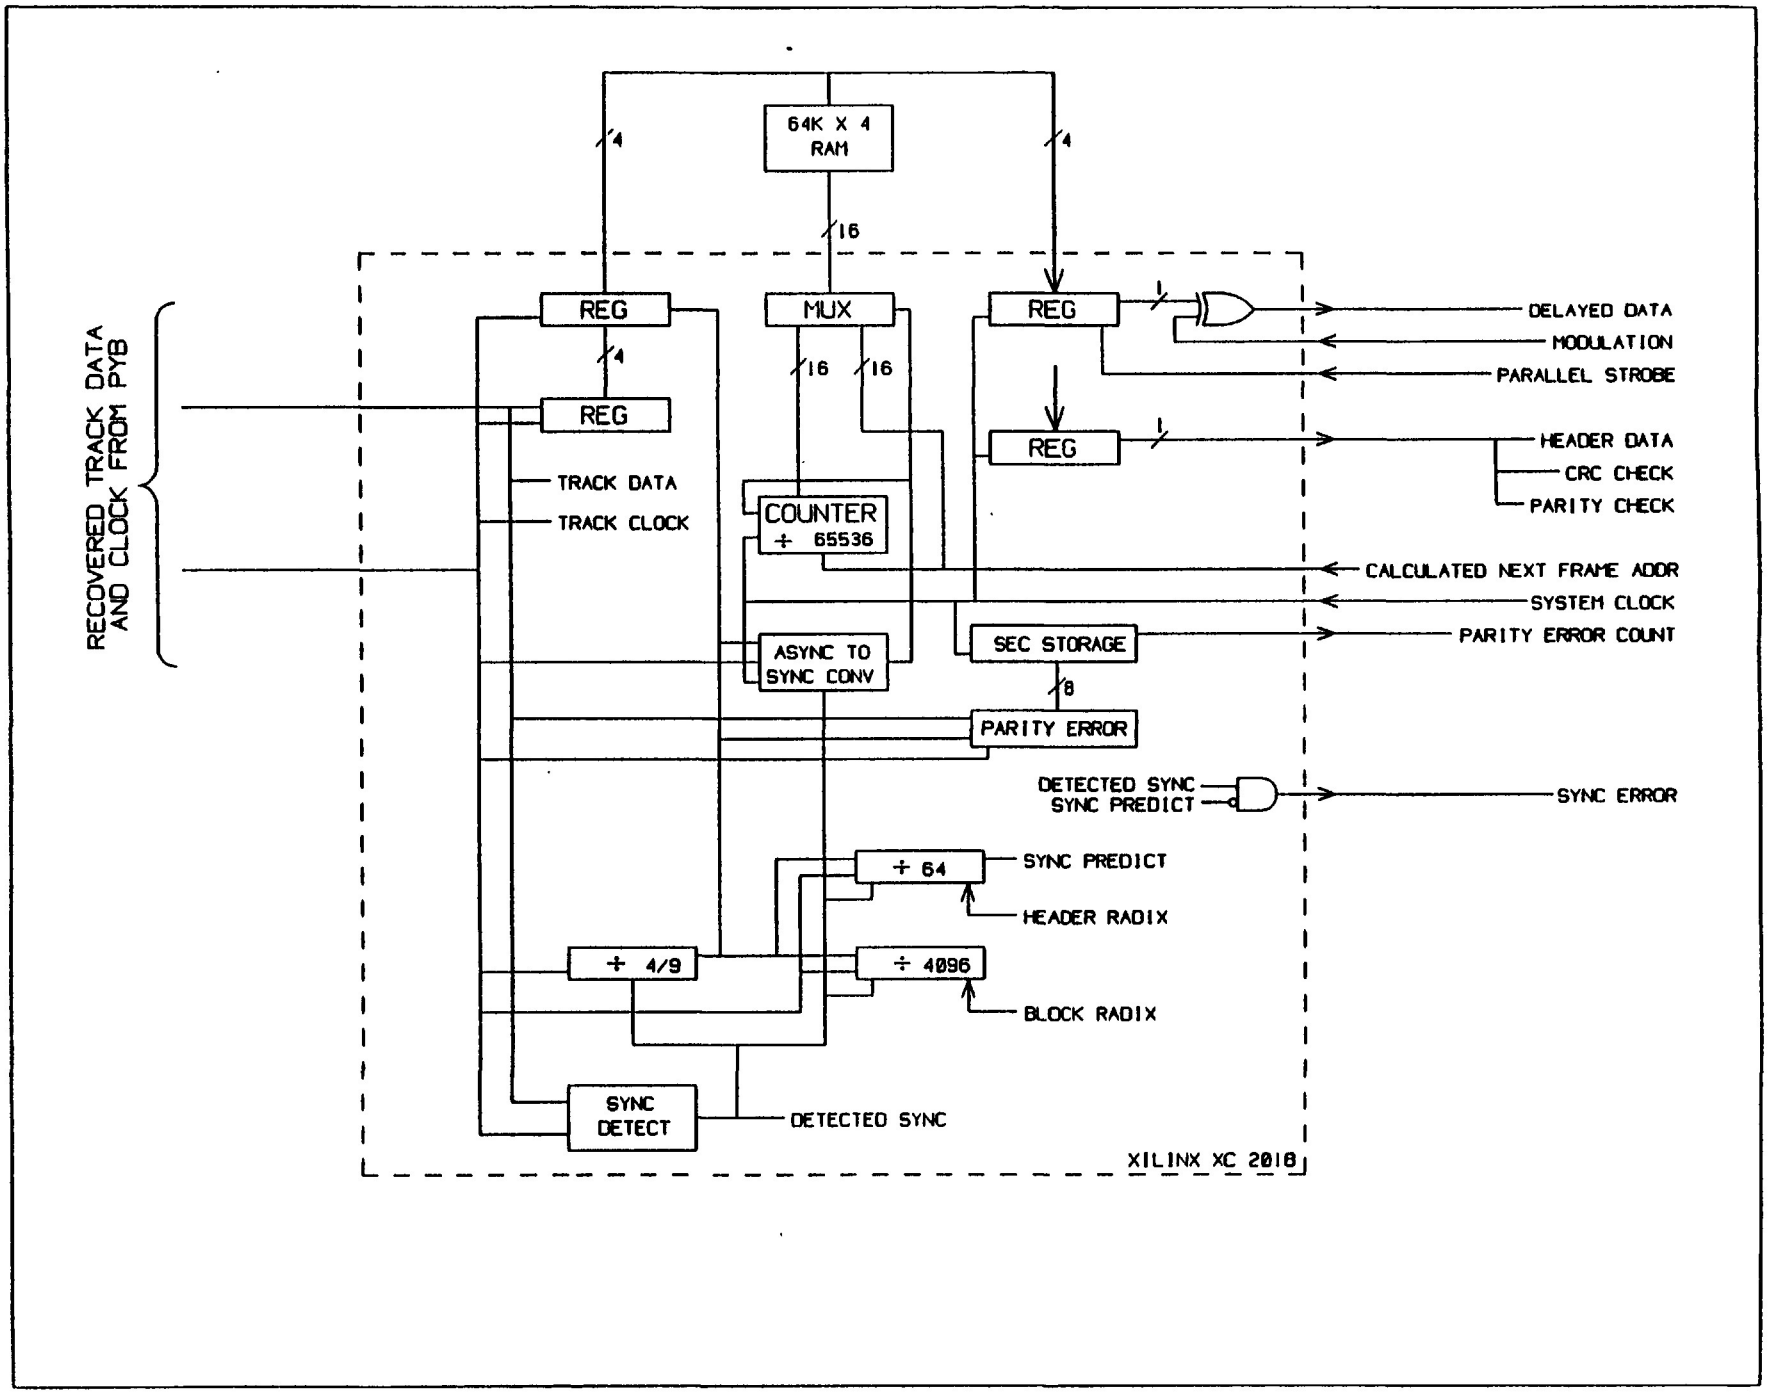
64K X 4
RAM
4
16
4
REG
MUX
REG
4
16
16
REG
REG
1
1
RECOVERED TRACK DATA
AND CLOCK FROM PYB
TRACK DATA
TRACK CLOCK
COUNTER
÷ 65536
ASYNC TO
SYNC CONV
SEC STORAGE
8
PARITY ERROR
÷ 64
÷ 4/9
÷ 4096
SYNC
DETECT
DETECTED SYNC
DELAYED DATA
MODULATION
PARALLEL STROBE
HEADER DATA
CRC CHECK
PARITY CHECK
CALCULATED NEXT FRAME ADDR
SYSTEM CLOCK
PARITY ERROR COUNT
DETECTED SYNC
SYNC PREDICT
SYNC ERROR
SYNC PREDICT
HEADER RADIX
BLOCK RADIX
XILINX XC 2018

Figure 3

depicted in the block diagram of figure 3 will be given here. A more detailed description of the track recovery logic is presented in appendix I.

The logic block of figure 3 labeled "sync detector" will monitor the incoming track data for the sync word that is in the header bit field. When the sync word is detected, it will cause several counters in the track recovery logic of figure 3 to be initialized. These counters are clocked on the recovered track clock and the state of these counters are subsequently used to predict frame boundaries, the position of future sync patterns in the track data, and to provide input addressing for the RAM buffer. Time is calibrated across the RAM which is large enough to contain several complete frames. Time calibration of the RAM addresses is established on the output side of the buffer by the delay model generator. The model generator develops read addresses for the RAM in order to produce an output data signal that will eventually drive the FFT system of the correlator. When the delay model generator addresses RAM memory location X, it will expect to find a sample taken at time Y to be located in it. The fixed sample period results in the ability to then calibrate the RAM addresses in units of time. The address counter that controls the data input into the RAM is set by calculating the proper RAM address in which to store a given input track data bit on the basis of the time specified by a previous track header time code. The delay function is provided since the delay model generator specifies the bit within the RAM that is required at the correlator input at any given time.

Both the track data and the track header information will be stored in the same RAM. Hence two modes of operation are required when writing into the RAM. During the data portion of a playback frame, only the data bits are written into the RAM, the along track parity bits are stripped off and discarded after parity checking and error counting. During the header portion of a frame, however, all header bits are written into the RAM. A parity check of the header information is to be made after the fact when the header data is read out of the RAM. The track data and track header bits are separated from each other coming out of the RAM by providing two ports out of the RAM buffer as shown in figure 3. The delayed data output is obtained by reading through the RAM addresses where frame data is stored and skipping over the header data addresses. The header data is output from the second port as shown for non-real time header analysis.

An 8751 microprocessor will have access to the header data output from the frame buffer. The logic for each station will have 4 such 8751s with each microprocessor thus being required to supervise the maintenance of up to 9 tracks. An 8751 has the responsibility of reviewing the header information stripped out of the track data and conveying status and relevant information to a 68000 microprocessor which will provide the

principal intelligence for the station. This 68000 micro will have the responsibility to supervise the station synchronization process and to discern the overall station status. The synchronization process will require the 68000 to communicate with the 8751 microprocessors to obtain track status information and with the PBD for speed control. Communication between the 68000 and the PBD will probably by via an RS422 line at 56 kbaud.

The 68000 in each station of the playback interface logic, in order to control the delay function as has been described, will also provide the delay model generation function. The delay model generator will use an 8-term Taylor series expansion algorithm implemented in software for delay model tracking. This model generator is programmed with the Taylor series coefficients by the system controller. There will be 4 total sets of model coefficients in the model generator to support fast model switching or to support multiple phase centers. The model generator is required to receive model parameters, start tracking delay models upon command, and down load fresh delays into the delay lines every 4-msec. The delay model generator is very similar functionally to the fringe model generator which will be described elsewhere.

The only problem associated with performing the delay function in the track domain and not in the channel domain is presented when different phase centers must by observed in a single channel. Each phase center requires a different delay model. Thus there is need for a multipliexer in front of the track recovery logic depicted in figure 3 so that one track may drive two or more such track recovery logic circuits (where each copy of the track data can be provided with a different delay model). Except for this requirement only one delay model would be needed for a station since delay is frequency independent. The requirement for a single track to receive several delay models is easily met, however, and figure 4 shows the logic necessary to support the ability to observe up to 4 phase centers in a single channel. Four tracks drive a multiplexer (made out of a PAL logic chip) and a selection is made from among these four tracks for the sources for four track recovery logic blocks (only one of which is shown in figure 4). On one extreme, where only one phase center is to be observed, these four track signals will flow through the multiplexer stage to drive the four track recovery stages. In the other extreme, only one track will be active and it can be copied into all four track recovery circuits and four phase centers may be observed. Appendix I gives a more detailed description of figure 4.

As stated in section III, the correlator will support up to 20 active DPS stations at any one time. There will be a total of 24 DPS stations at the correlator as illustrated in figure 5. For each 5 active DPS playback units there will be a "hot spare" playback drive which can substitute for one of the 5

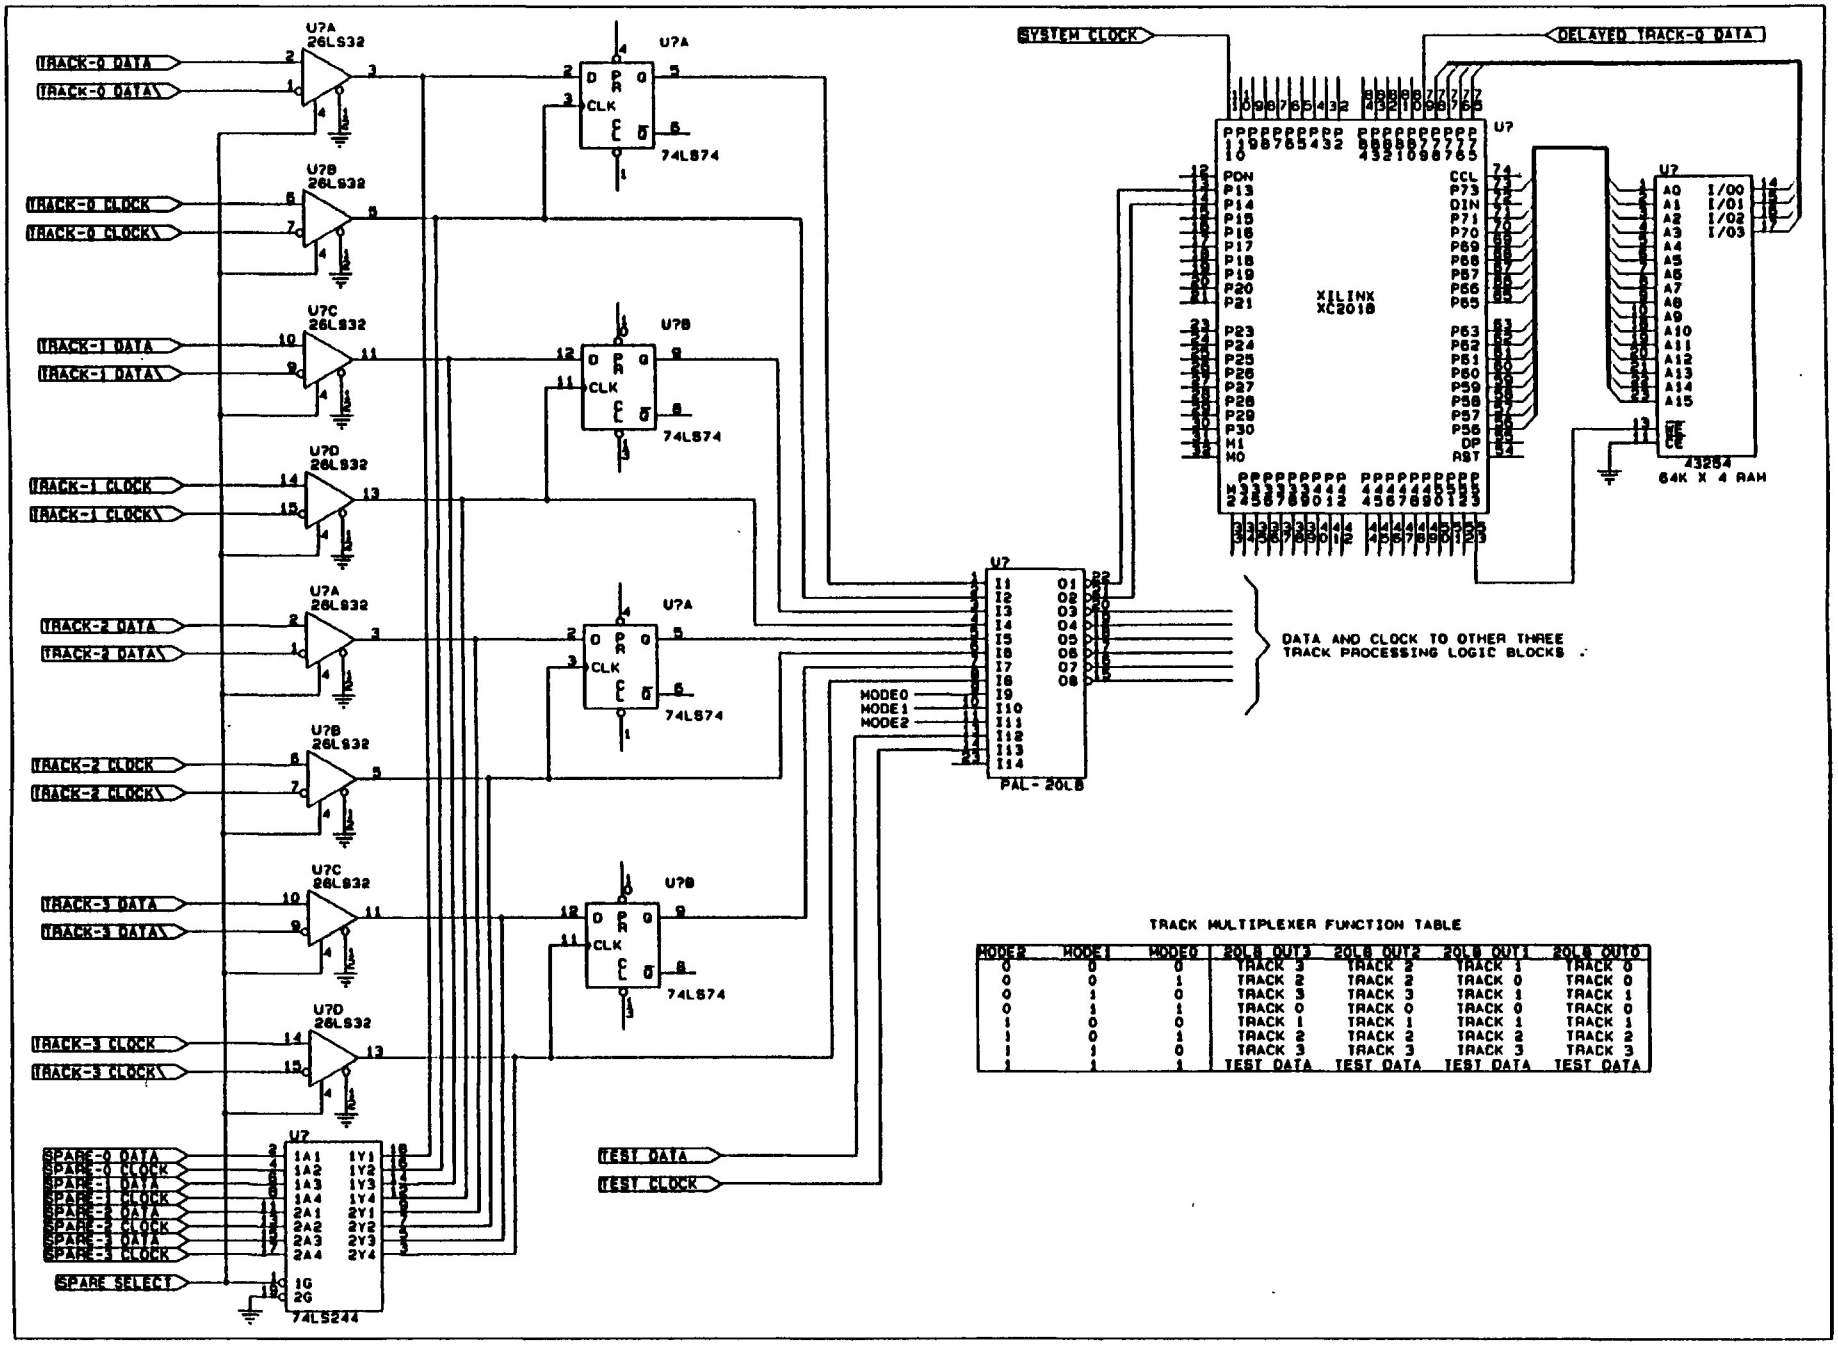

TRACK MULTIPLEXER FUNCTION TABLE

MODE2	MODE1	MODE0	20L8 OUT3	20L8 OUT2	20L8 OUT1	20L8 OUT0
0	0	0	TRACK 3	TRACK 2	TRACK 1	TRACK 0
0	0	1	TRACK 2	TRACK 2	TRACK 0	TRACK 0
0	1	0	TRACK 3	TRACK 3	TRACK 1	TRACK 1
0	1	1	TRACK 0	TRACK 0	TRACK 0	TRACK 0
1	0	0	TRACK 1	TRACK 1	TRACK 1	TRACK 1
1	0	1	TRACK 2	TRACK 2	TRACK 2	TRACK 2
1	1	0	TRACK 3	TRACK 3	TRACK 3	TRACK 3
1	1	1	TEST DATA	TEST DATA	TEST DATA	TEST DATA

Figure 4

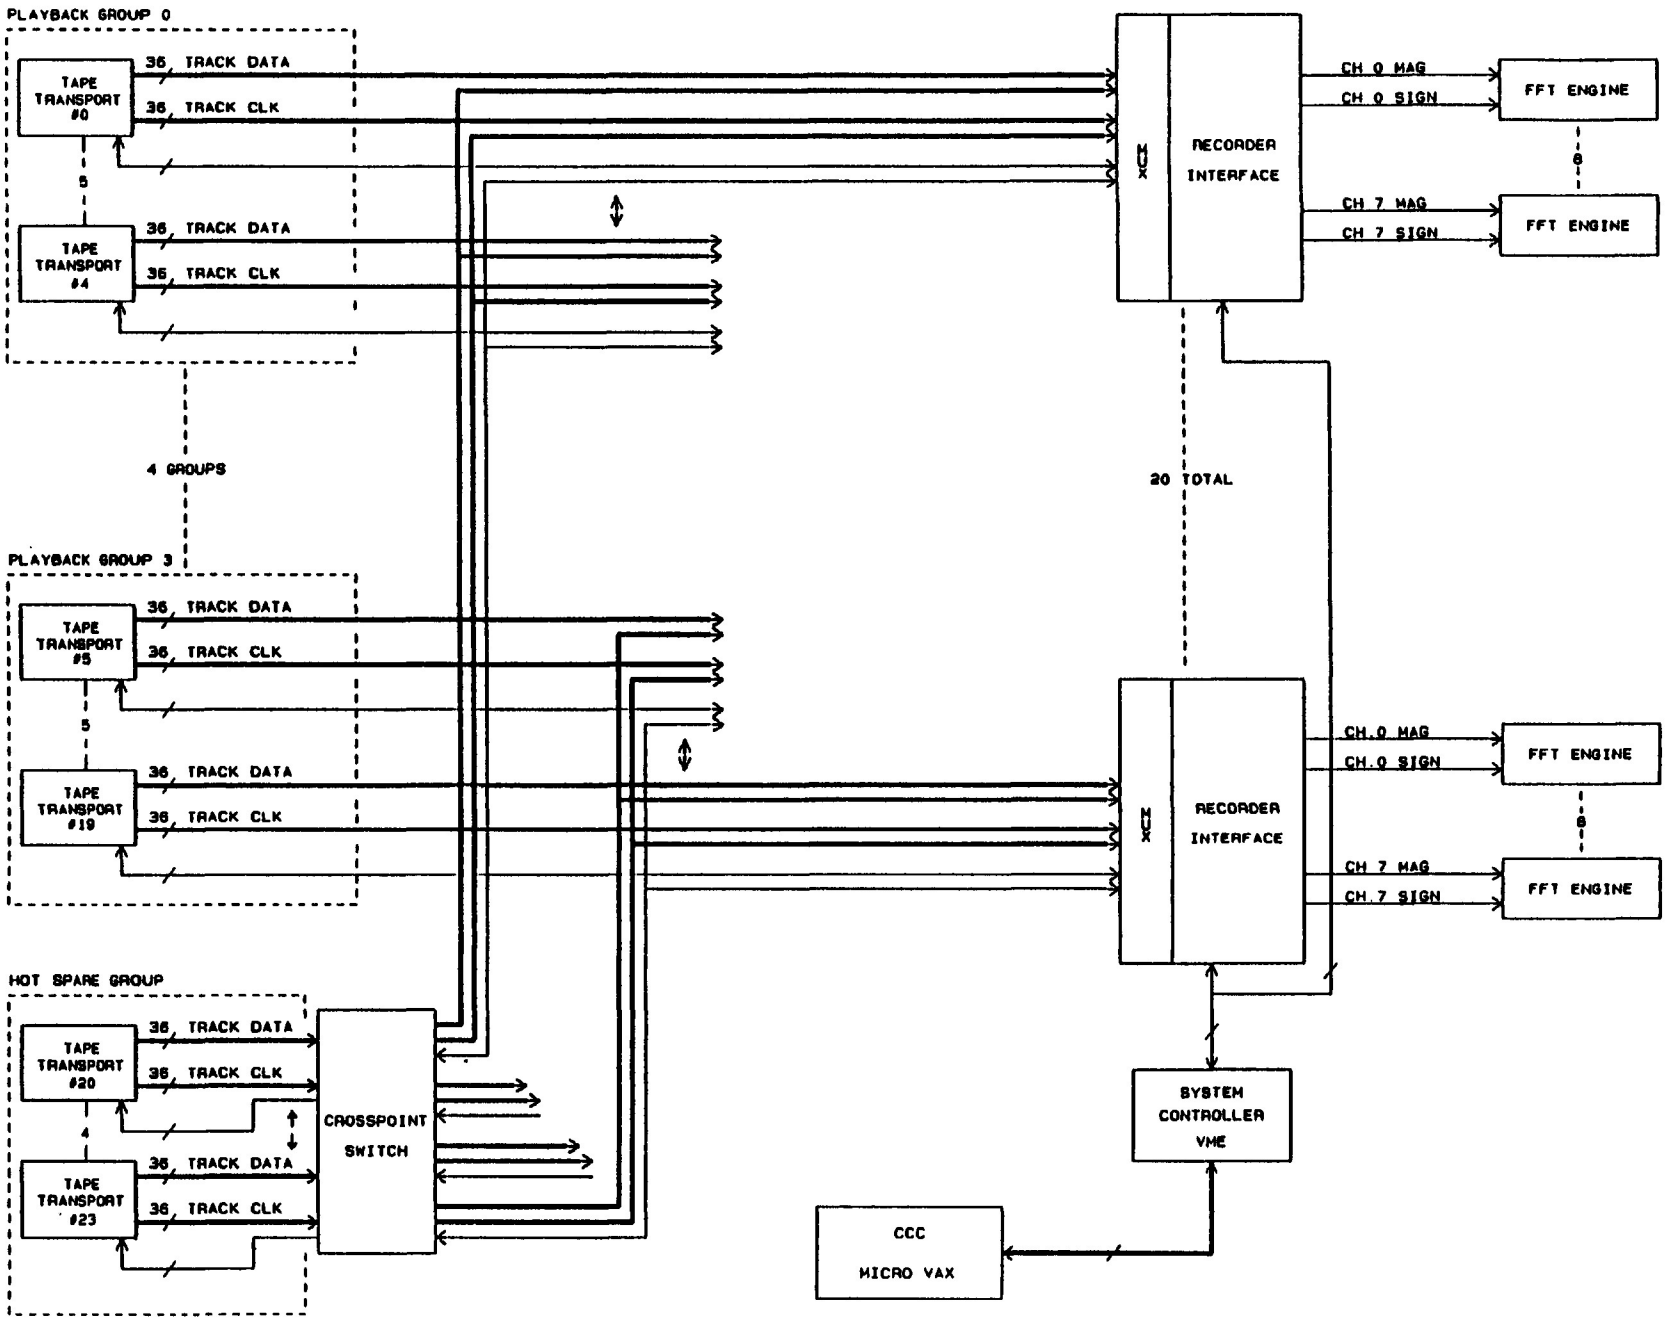
PLAYBACK GROUP 0
TAPE TRANSPORT #0
36 TRACK DATA
36 TRACK CLK
5
TAPE TRANSPORT #4
36 TRACK DATA
36 TRACK CLK
4 GROUPS
PLAYBACK GROUP 3
TAPE TRANSPORT #5
36 TRACK DATA
36 TRACK CLK
5
TAPE TRANSPORT #19
36 TRACK DATA
36 TRACK CLK
HOT SPARE GROUP
TAPE TRANSPORT #20
36 TRACK DATA
36 TRACK CLK
4
TAPE TRANSPORT #23
36 TRACK DATA
36 TRACK CLK
CROSSPOINT SWITCH
MUX
RECORDER INTERFACE
20 TOTAL
CH 0 MAG
CH 0 SIGN
FFT ENGINE
8
CH 7 MAG
CH 7 SIGN
FFT ENGINE
MUX
RECORDER INTERFACE
CH.0 MAG
CH.0 SIGN
FFT ENGINE
8
CH 7 MAG
CH.7 SIGN
FFT ENGINE
SYSTEM CONTROLLER VME
CCC MICRO VAX

Figure 5

to support efficient tape changing or in the event of a DPS failure. The four spare drives feed a cross bar switch as shown in figure 5 so that any of the four may be assigned to a given set of 5 active DPS stations. In summary, the 24 playback drives may be grouped into four sets of 6 units. Five of each group of 6 playback units have dedicated paths through the correlator. The sixth unit can be selected freely from the set of 4 spare playback units via the cross bar switch and can be switched in place of any of the 5 when supplying input drive into the correlator.

Other items shown in figure 2 correspond to inverses of operations that occurred in the formatter. Thus the barrel roll and the track to channel conversion will reverse procedures performed by the formatter at record time.

If a station drive has a defective head in its head stack, data from the track assigned to the bad head may be recovered in one of two ways. If the defect had been known of in advance, the data for that track could have been duplicated in one of the system tracks and recovery can be accomplished by switching in the system track in place of the bad track. Otherwise, if the defect had not been known of in advance, and if across track parity was being used, then the across track parity errors could be used to reconstruct the data of the lost track by exclusive-oring the bad track data with the parity error flag. Both of these options are supported in the playback interface system.

V) PLAYBACK INTERFACE LOGIC TO FFT INTERFACE

The descriptions of section IV explained how the playback interface logic will reconstruct the sampler outputs from the playback data. A small amount of additional data processing is still required to deliver the proper drive into the correlator. There are a number of correlator options that affect the requirements of the playback interface output stages. The varying bit rate of the reconstructed sampler signals will also affect the playback interface logic to FFT input.

The samplers at an antenna can be set to run at 250 or 500 ks/s, 1, 2, 4, 8, 16, or 32 Ms/s. A four to one or two to one playback speedup factor for the lower data rates will reduce the sample rate variations at the correlator input to 1, 2, 4, 8, 16, or 32 Ms/s, however. The total bit rate input to the correlator is set by this sampler data rate and the number of active channels per station to be processed. This aggregate data rate into the FFT system of the correlator must be reconciled to the capacity of the FFT system in order to determine the number of processing options that are allowable in processing a given experiment.

The processing options mentioned above are those described in VLBA correlator memo 71, i.e., transform overlapping, transform zero padding, and transform interleaving. Where excess FFT capicity exists, there is also the possibility to map several phase centers within an antenna beam. Each of the items above requires some support of the FFT system by the playback interface system.

Two blocks of figure 2 are required to support the FFT system data processing. The "channel mix" block of figure 2 allows the data from one channel to drive several FFT engines. This ability is required by processing options such as interleaving and 2048-point transforms. The "channel buffer" block of figure 2 contains a 1024-bit RAM that will be used to store samples temporarily in order to support other processing options such as offset transforms and zero padding and to convert between the varying channel bit rate and the constant 32-Mbs FFT system input data rate.

A) Interleaving

The processing option called interleaving can be used when the IF signal is oversampled. This option allows the observer to have independent Fourier transforms and cross-multiplications performed on odd and even samples in the data. The resulting cross-spectra will be accumulated together yielding an increase in sensitivity. This option requires the separation of the samples into odd and even data samples.

B) Overlapping

The overlapping option in high sample rate experiments is provided for by copying the samples of a given channel into two (or more) RAM stages in the channel mix logic and programming the RAMs to provide the fractional FFT delay. This operation requires more than one FFT engine to be dedicated to the processing of one channel. For low sample rate experiments, overlapping will come as a natural consequence of the change of clock rate that occurs at the 1K RAM.

C) Zero Padding

The zero padding option is used to make the FFT reversible by padding the transform with one half zero samples. Since one half of a transform does not contain antenna samples, transforms must be performed with boundaries in the data of one half the transform length so as not to throw away half of the information. This option is, like the overlapping option, obtained by copying a channels' samples into two delay RAMs and driving two FFT engines from the samples of one channel. In practice, this option is performed exactly like the factor of two

overlap option above with the zero pad being supplied by the window generator.

D) Phase Centers

Multiple phase centers in the telescope beam can be processed simultaneously by again driving two or more FFT engines with the samples from one channel, this time with different models in the delay and fringe model generators. Figure 4 illustrated how this option is provided for in the VLBA correlator (the duplication of a channel into more that one FFT engine must takes place not in the channel mix block but at the input to the track processing logic).

E) 2048 Point FFT Support

One additional requirement of the playback interface system is to support the FFT system in performing 2048-point Fourier transforms. The custom gate array chips that will perform the actual Fourier transforms have two 1024 X 18 RAMs on board to allow for double buffering the points as they proceed down the FFT chains. Double buffering is required because of the need to process points in a bit-reversed time order. To perform 2048 point transforms, however, this RAM size is insufficient. Hence, two FFT chains are required to perform 2048-point transforms with a point shuffle occurring between the two at the last butterfly stage input. The requirement that this operation imposes on the delay system is exactly the same as the interleave option, the first 1024 points of a 2048-point FFT are supplied by one delay RAM into one FFT chain and the second 1024 points are supplied into the second FFT chain by a second delay RAM. However, in terms of transform multiplicity, 2048-point transforms are "free" since the transform occurs in one half the time normally expected.

APPENDIX I DETAILED DESCRIPTION OF TRACK PROCESSING

Refer again to figures 3 and 4. These figures illustrate the hardware required to process the output of a playback transport. The logic shown in the dashed line of figure 3 will be implemented in a single integrated circuit. This IC is the Xilinx XC-2018 which is a programmable logic chip.

Consider first figure 4 which gives a more global view of the track recovery logic. The first function to be seen in figure 4 is the differential receivers and flip-flops that receive and capture the track data from a playback drive. The three-state output of the differential line receivers form a multiplexing stage where spare drives may be substituted for the normal input into the track recovery logic. This substitution is

accomplished by disabling the differential receivers' three-state output and enabling the spare three-state buffer. Next, a PAL multiplexer stage allows tracks within a group of four to be duplicated into multiple track recovery circuits to support multiple phase center observations. Also at this stage, a pseudo random test signal may be switched into the correlator input. This last path will allow for system testing during non-processing periods such as at the end of a tape pass when the drives are busy reversing direction or to verify the correlator system prior to starting the processing of an experiment.

The first function after the multiplexers (input of figure 3) is the serial in, parallel out register with secondary storage. This register converts the input data into 4-bit nibbles for storage in the 64K by 4 RAM. The divide by 4/9 block develops the strobes that transfer data from the input serial register into secondary storage each time four bits are shifted into the chip. During the data portion of a frame, the parity bits will be rejected by the parallel strobe and hence will not be written into the RAM. During the header portion of a frame, however, all of the header bits, parity bits included, are stored. Two additional counters, the divide by 64 counter and the divide by 4096 counter provide for complete frame event timing (64 and 4096 being the maximum counter radices for header and frame lengths it can accommodate). The divide by 64 counter counts off 4-bit nibbles in the header portion of a frame and the 4096 counter counts off 9-bit parity bytes during the data portion of a frame. These counters predict header boundaries and the sync time. The programmable radix counter receives the established frame and header count lengths for the tape format selected for the VLBA via the Xilinx chip personality program bits.

The onboard sync detector logic will scan the incoming bit stream for the sync pattern. Perfect compliance with the sync word is required for detection by this internal sync detector. When the header data is read out of the RAM, a more sophisticated sync detector chip (a TRW TDC 1023) will monitor the data stream with a sync template. This chip has an internal 64-bit register that will be programmed by the local microprocessor before an observation is to be processed. A second internal 64-bit mask register, also programmed before hand, will allow only the selected bits in the 64-bit field to be compared with the template. A programmable threshold will establish the level of template agreement that constitutes sync detection. Since this second sync search is not done in real time, it cannot help in the initial synchronization of a track but it will help in identifying sync slips more quickly or more definitely than the real time sync detection.

On initial synchronization, the on board sync detector must continuously test the data stream for the sync pattern. Upon achieving synchronization, however, the sync detector can be

limited to a sync search within the header only. This limitation will desensitize the system to false sync detections in data and with the sync pattern located at about mid header, this narrowing of the sync search boundaries should be acceptable in recovering from typical sync slips.

A detection of the sync pattern by the onboard sync detector will cause three counters to be initialized. The divide by 4/9, the divide by 64 and the divide by 4096 counters will all be jam-set to the appropriate count numbers within the header for the sync pattern. In addition, the divide by 65536 counter that supplies the address for track data to be written into the main data buffer will be set to an appropriate address as supplied by the microprocessor.

This last principal is the heart of the operation and deserves detailed explanation. A microprocessor will maintain the count labeled in figure 3 as CALCULATED NEXT FRAME ADDR. This microprocessor has responsibility to keep time for a given set of tracks. At some point within a tape synchronization procedure, the header data written into the RAM will be found to have reliable information in it. The test for reliability includes;

1) a good CRC check
2) no parity errors within the header (taking into account that the sync word may be parity encoded opposite from the sense of other data within the header).
3) a sync pattern in the proper location
4) other data within the header makes sense
5) the time code makes sense

Once this reliability test is passed, the microprocessor can take the recovered time from that header as the PBD time and use the known frame length to calculate the time to be expected for the sync word of the next header. Having this calculated time and knowing the relation between time and the RAM addresses as established by the delay model generator (i.e. that the delay model will expect a sample taken at time Y to be stored in RAM location X), the micro can then calculate the RAM address at which the first bit following the sync word of the next frame should be stored. This calculated RAM address will then be left waiting for the next sync detection to initialize the input RAM address counter (if that sync detection does not come, the micro can use its internal timer to know that the attempt has failed and that it must try again). Eventually the counter will be properly initialized and the system will operate smoothly.

When a sync slip occurs, the RAM counter will be re-set by the next sync word (if proper sync detection occurs) and the microprocessor will know that a problem existed both by non-coincidence of sync predict and sync detect and by a garbled

header and can flag the previous frame as invalid when data from the frame is output from the RAM. By the time the microprocessor knows of the existence of a problem, the re-sync operation should already have taken place. The micro can then judge the success of the re-sync by examining the time code just after it occurred and the time codes of succeeding frame headers.

If a sync pattern is missed in a system already in sync, the frame buffer is long enough that the microprocessor may withhold judgment on the validity of its contents for a few frames if desired. Thus the system may ride through a bad header if, when proper header information is reestablished, no sync slip is detected and the parity counts of the two frames are not excessive. Such strategies are, of course, determined in software and exact implementation need not by considered here.

The microprocessor will have to keep a short file of data validity information in memory that reflects its judgement concerning the validity of the contents of all frames stored in the frame buffer RAM at any given time. An address register will also have to be maintained at the output of the buffer that knows the RAM addresses of the first bit of each frame in the RAM so the system can assert or remove appropriate validity bits on exact frame boundaries and so that the header bits will be skipped over when reading the RAM to reconstruct the track data before frame overhead insertion. An additional requirement of this information is to ensure synchronization of the inverse barrel roll function at frame boundaries.

The frame buffer will work on the stable system clock, allowing station wide generation of control signals. In order to work the RAM on the system clock, the block in figure 3 labeled ASYNC TO SYNC CONV is required to synchronize the asynchronous track data RAM write operation and to jam load the RAM address counter with the CALCULATED NEXT FRAME ADDR count upon sync detection. The RAM write operation is required every time four bits are strobed into the secondary storage of the input serial to parallel converter. When the secondary storage register is full, an asynchronous (to the system clock) write request signal for the RAM will be generated. The RAM control logic will monitor the write request line and service it by generating a RAM write pulse on the next available (system) clock cycle. Since this RAM is time shared between the data input port and two output ports, RAM contention logic is required to allocate access to the memory.

As many functions as possible, such as the header CRC and parity checks, are left until after the reception of the header when time is not of the essence. Thus the header output port of the track RAM in figure 3 can be paralleled with the header outputs of other tracks and sync check, the CRC check, parity check and serial to parallel conversion for the

microprocessor can be done in hardware time shared among 9 tracks.

In all, the track buffer/delay function requires about 3 ICs per track plus other miscellaneous shared logic and a shared microprocessor. Figure 6 gives a possible card layout of a 9U VME card that can process 18 of the 36 tracks of one PBD.

Figure 7 gives a floor plan of the Xilinx XC-2018 chip showing how this chip can be used for this application. The Xilinx chip has 100 programmable logic elements arranged in a ten by ten array in the chip center and 74 programmable I/O logic blocks along the chip edges. The internal logic elements each have 4 logic inputs, a global clock input and two outputs. Two modes of operation in these internal reconfigurable logic blocks are possible. If only one of the two outputs of a logic element is used, the element can be configured to implement any logic equation of four inputs and one output. If both outputs are used, each can implement any logic equation that contains any three of the four inputs. One flip-flop is contained in the element and feedback into the logic element from this flip-flop is possible. Each I/O block can be programmed to be an input with or without storage or a 3-state output.

The Xilinx chip is given a personality by loading an internal RAM with design program words. Since the storage of the personality is in RAM memory it must be re-programmed every time power to the chip is cycled. This fact allows for re-design (to accommodate changes of format by changing the frame counter radix, for example) via firmware changes.

Appendix II TAPE DATA FORMAT

The proposed design will not be able to accommodate all of the versatility built into the VLBA formatter designed by Haystack Observatory. Where the flexible features in the formatter could be easily and inexpensively accommodated, however, they were supported in this design. Some restrictions on the tape format that influence the hardware design, however, are listed below (bytes expressed below are 9 bit bytes):

1) The frame length will be between 2048 and 4096 bytes long.
2) The header field will be contiguous and will be between 16 and 32 bytes in length.
3) The 36-bit MARK III sync word will be used.
4) There will be a 12 or 16 bit CRC field over the header portion of a frame and no CRC field over the data portion of a frame.
5) There will be eight bit parity.

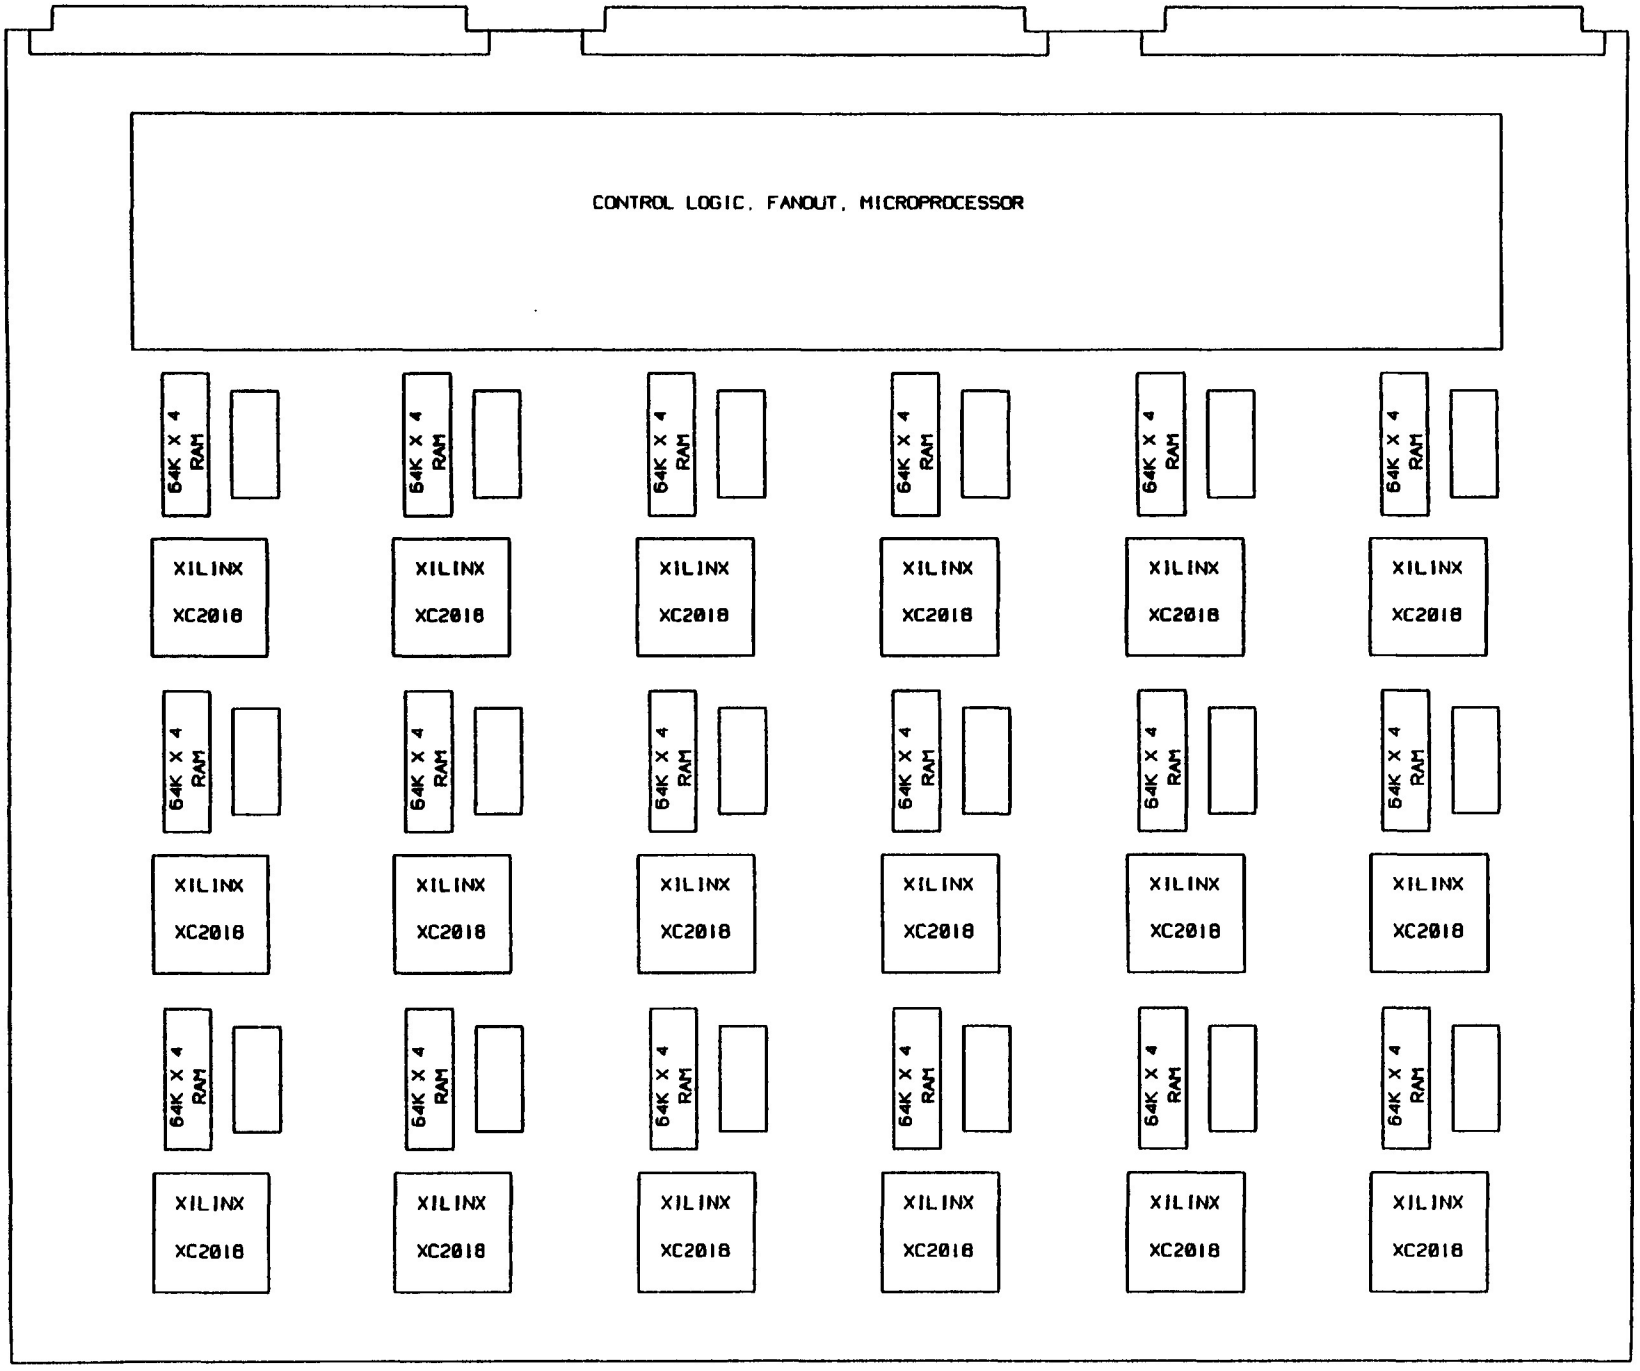
CONTROL LOGIC, FANOUT, MICROPROCESSOR
64K X 4 RAM
XILINX XC2018

Figure 6

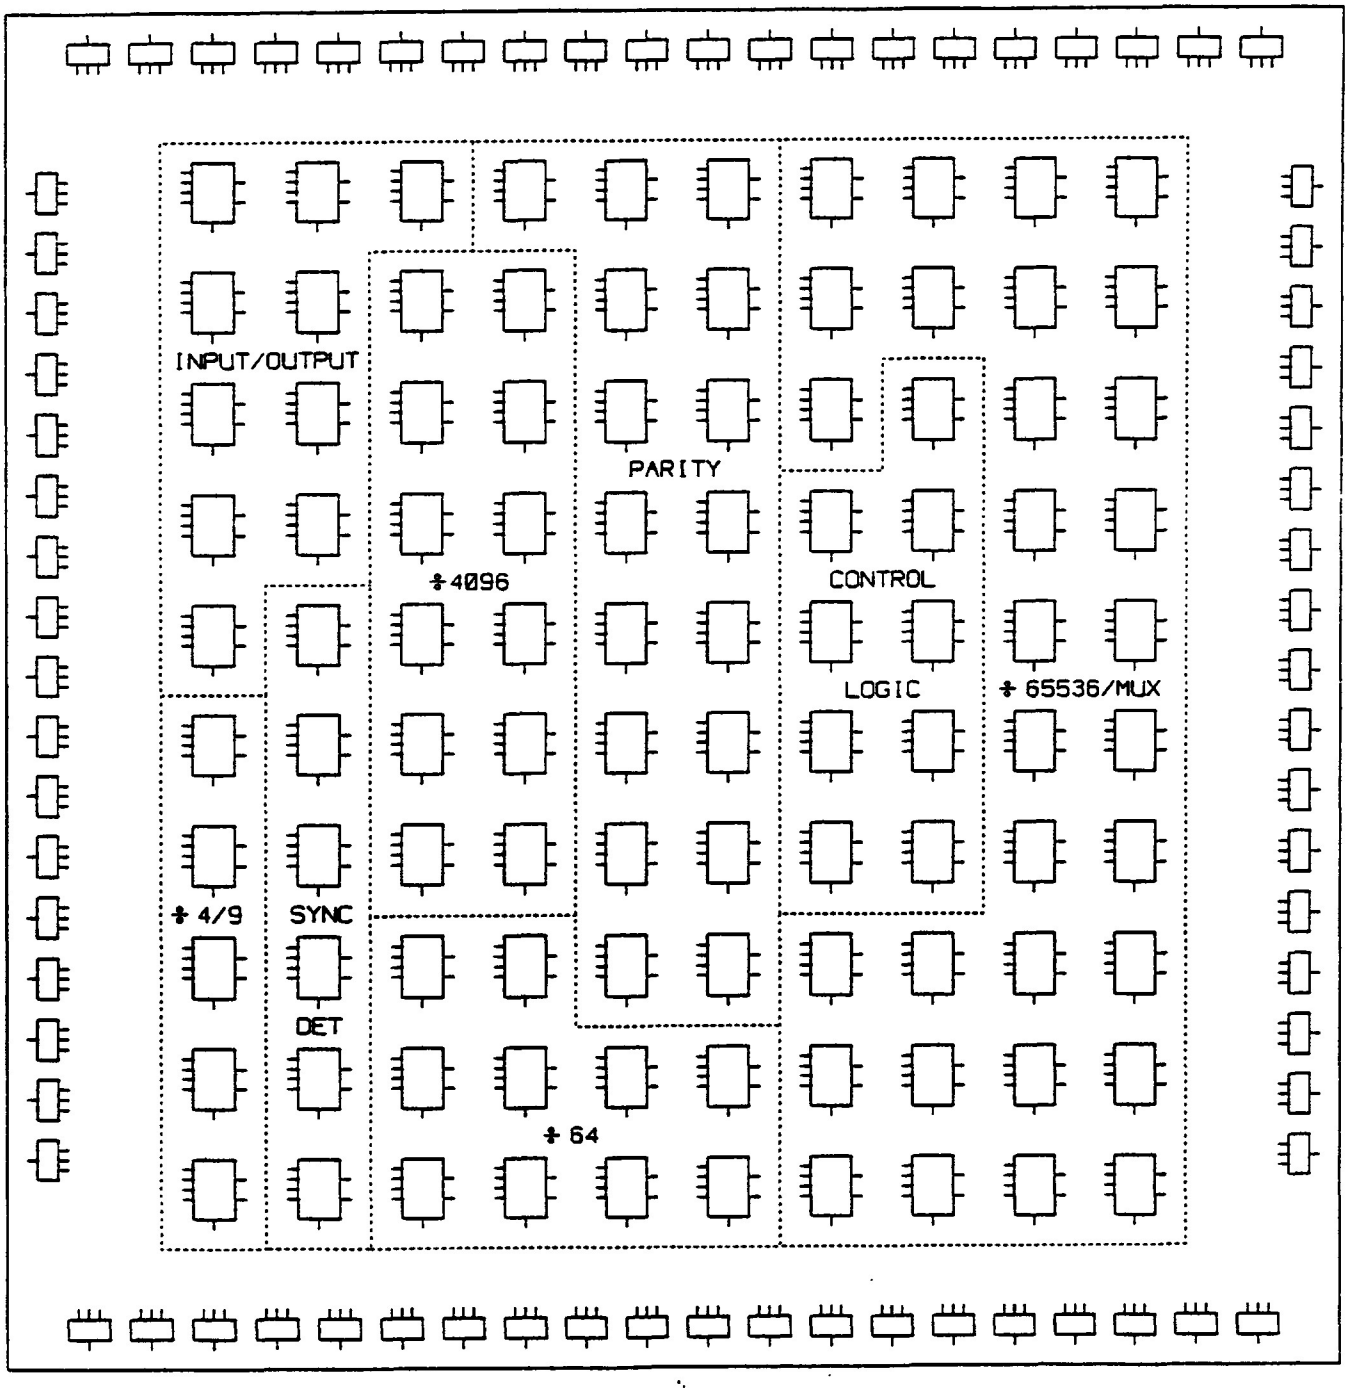

FLOOR PLAN OF XILINX CHIP

Figure 7

Most other information required to define the tape format will primarily influence the software and considerations of the full format specification can be left until later.

It is envisioned that a specific tape format, with fixed numbers for the frame and header lengths, etc, will shortly be defined and that only that format (and possibly the MARK III format if this does not involve an excessive cost) will be supported by the VLBA. However, if an early format turns out to be deficient in some respect, a change of format can be easily accommodated using the design proposed here with no hardware changes by re-programming the Xilinx chip (as long as the limits above are observed).