MASSACHUSETTS INSTITUTE OF TECHNOLOGY

HAYSTACK OBSERVATORY

WESTFORD, MASSACHUSETTS 01886

29 March 1990

Area Code 508
692-4764

MEMO #11

Dr. James M. Moran
Center for Astrophysics
60 Garden Street
Cambridge, MA 02138

Dear Jim,

I enclose ten copies of our final report on the Correlator for the Submillimeter Array.

We are confident that we could build the correlator now using the "Bos" chip to satisfy the requirements in a very cost effective manner.

We are pleased to join with your design team and we are available to make presentations and answer any questions. We look forward to the next phase of detailed design and prototyping.

We have tried to answer all the questions in Colin's letter of February 16, 1990 in this revised report.

Yours sincerely,

Alan

Alan E.E. Rogers

Distribution: (Letter and Report)
Center for Astrophysics
R. Blundell
P. Ho
C. Masson
P. Myers
M. Reid
I. Shapiro
E. Silverberg

Haystack Observatory
J. Connolly
J. Levine
J. Salah
A. Whitney

Design Study

for the

HARVARD-SMITHSONIAN CENTER FOR ASTROPHYSICS

of the

I.F. PROCESSOR AND CORRELATOR FOR THE SUBMILLIMETER WAVELENGTH ARRAY

James I. Levine
Alan E.E. Rogers

Haystack Observatory

29 March 1990

INDEX

SECTION	A.	I.F. PROCESSING AND CORRELATOR	1
	A.1	Introduction	1
	A.2	I.F. Distribution	1
	A.2.1	Baseband Converters and Converter Synthesizers	3
SECTION	B.	REAL TIME DIGITAL PROCESSING	6
	B.1	Introduction	6
	B.1.1	Processing Options	6
	B.1.2	Design Goals	6
	B.2	Processing Requirements	7
	B.3	Processor Design	7
	B.3.1	Baseline Signals	7
	B.3.2	NFRA Correlator Circuit	9
	B.3.3	Multi-processing Techniques	9
	B.4	Physical Organization	11
	B.4.1	Baseline Correlator Chassis	11
	B.4.2	Correlator Module	13
	B.5	Other Processor Requirements	16
	B.5.1	Fourier Transformations	16
	B.5.2	Data/Spectra Links	16
	B.5.3	"Fractional Bit" Correction, Dicke Switching and Interferometer Sideband Separation	16
	B.6	Cost Analysis - Parts	17
	B.7	Cost Analysis - Labor	19
	B.7.1	Hardware	19
	B.7.2	Software	19
	B.7.3	Construction	20
	B.7.4	Summary	20
	B.8	Other Processor Options	21
	B.8.1	The NFRA Correlator Circuit	21
	B.8.2	250 MHz Correlator - CalTech	21
	B.8.3	32 MHz FFT Engine - NRAO	22
	B.8.4	2 GHz FFT Engine - NTT	22
	B.8.5	Custom Design	22
	B.8.5.1	Impact of Semi-custom Chip Design	23
	B.8.5.2	Processing Speed	23
	B.8.5.3	Processor Size, Power	23
	B.8.5.4	Design Effort	23
	B.8.5.5	Schedule	24
	B.8.5.6	Costs	24
	B.8.5.7	Summary	25
	B.8.6	Separate Continuum Correlator	26

SECTION	C.	SUMMARY OF CHOSEN DESIGN	27
	C.1	Summary of XF Hybrid Correlator	27
	C.2	Signal Switching	28
	C.2.1	Examples of Observing Modes	28
	C.3	Cost Estimate and the Cost Equation	29
	C.4	Equivalent Processing Power/Comparison with Other Processors	30
	C.5	Conclusions	31
SECTION	D	APPENDIX	32
	D.1	Parts List for Baseband Converter and L.O. Synthesizers	32
	D.2	Filter Shapes and Linking of Subspectra	33
	D.3	Minimum Number of Lags Needed for Continuum	33
	D.4	Cost Saving Options	33
	D.5	Maximum Sample Rate Selected for NFRA Chip	34

A. I.F. PROCESSING AND CORRELATOR

A.1 Introduction

We have studied various designs for the I.F. processor and correlator for the Harvard-Smithsonian submillimeter array. This first phase is a general design study to provide a cost estimate. The study includes an evaluation of various approaches and their technology. A cost estimate is provided for the approach we best understand and believe is very cost effective.

The basic philosophical principles used in the design were:

1) Preserve a high degree of modularity
2) Minimize power consumption
3) Minimize cost
4) Optimize reliability and ease of maintenance
5) Make flexible use of spectral channels
6) Ensure straightforward modular expansion paths in bandwidth, number of stations and number of spectral points

The chosen approach is a XF correlator preceded by multiple baseband converters to segment the wideband I.F.s. Since many spectral channels are needed for spectral line, we consider it more efficient to also use them for continuum observations, rather than have a separate continuum processor (see Section B.8.6).

Other architectures such as the FX and the AOS are possible and have been examined but are more complex and hence will probably require more engineering development than the hybrid approach. We have studied a range of parameters and technologies possible for the above architecture. For example, we have studied the use of ECL gate arrays (as in the CalTech mm correlator), CMOS gate arrays (like that used in the "Bos" chip), and Japanese GHz A/Ds and FFT engines.

Figure A.1.1 shows a block diagram of the I.F. processor and correlator. The modules are:

I.F. distributor
Baseband converter
Baseband converter synthesizer
Sampler
Digital selector switch
Correlator

A.2 I.F. Distribution

The receiver I.F. signals would be transmitted (probably via fiber optics) to the central processor where the four processor I.F. inputs (A,B,C,D) for each station select any four receiving bands. Since the digital processing is designed to provide all four polarization products the chosen I.F. bands would normally be two bands each with both polarizations. Examples of several observing modes are given in Section C.2.1.

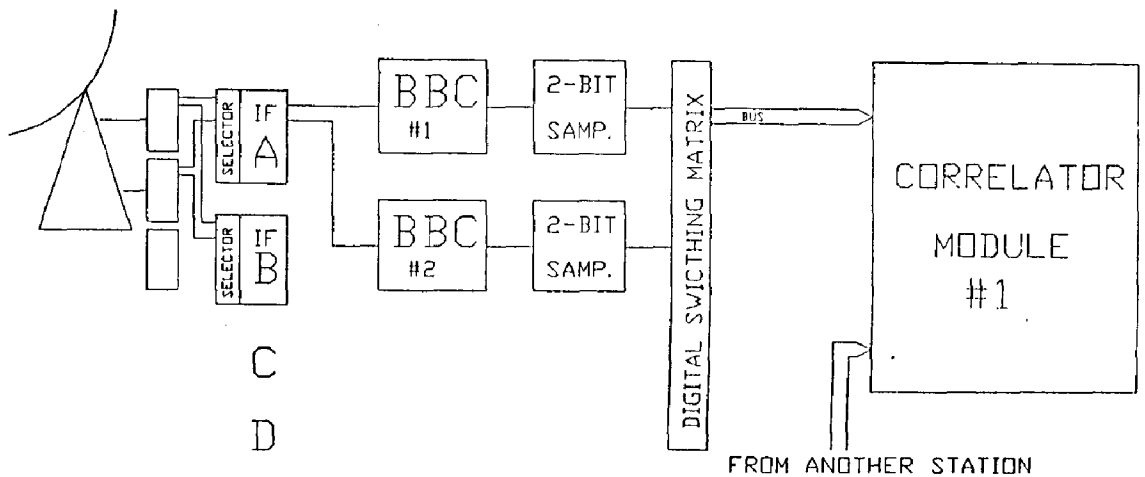

FIG. A.1.1 I.F. PROCESSOR AND CORRELATOR

A.2.1 Baseband Converters/Converter Synthesizers

Each of the four 1 GHz bandwidth I.F. channels would be divided into thirty-two 32-MHz channels with simple baseband converters. Since all the baseband conversion hardware would be placed in a central room, one synthesizer per frequency channel can supply the L.O. oscillator for all stations and polarizations so that separating the L.O. from the image rejection mixers will reduce the number of synthesizers required and make the system insensitive to phase drifts in the synthesizers. Since only 32 synthesizers are required a commercial unit might be considered to save engineering costs - alternately I estimate that the synthesizer shown in Figure A.2.1.1 could be built for about $500 (excluding checkout labor - see Section D.1). While only 32 synthesizers are required, having more would increase the flexibility of the instrument.

If separate synthesizers are used for each I.F. band and they are shared among antennas, there are no constraints on the placement of the "sub" bands. This requires a total of 128 synthesizers in which case it would be necessary to develop an inexpensive module.

Since a large number of baseband converters are needed, considerable effort should be spent in developing an inexpensive single board module. Figure A.2.1.2 shows a circuit diagram for a converter. The estimated parts cost with all commercial components and with many expensive components built in stripline as part of the board is given in Section D.1. With some effort it is not unreasonable to expect a cost (including checkout labor - if test fixtures are built) of around $300 per converter. The availability of very inexpensive ($5 in small quantities), ultra fast amplifiers makes the baseband quadrature network especially simple. Each converter would easily fit on a single board of no more than four inches by six inches. The total volume needed for all 768 converters would be about ten cubic feet. Power consumption would be under 2 KW.

The choice of bandpass filter shape and the linking of "subspectra" from each converter is discussed in Appendix D.2. In general, we propose to follow the approach used in the Kitt Peak hybrid spectrometer, except we propose to use only one sideband from each converter to increase flexibility in the choice of filter overlap.

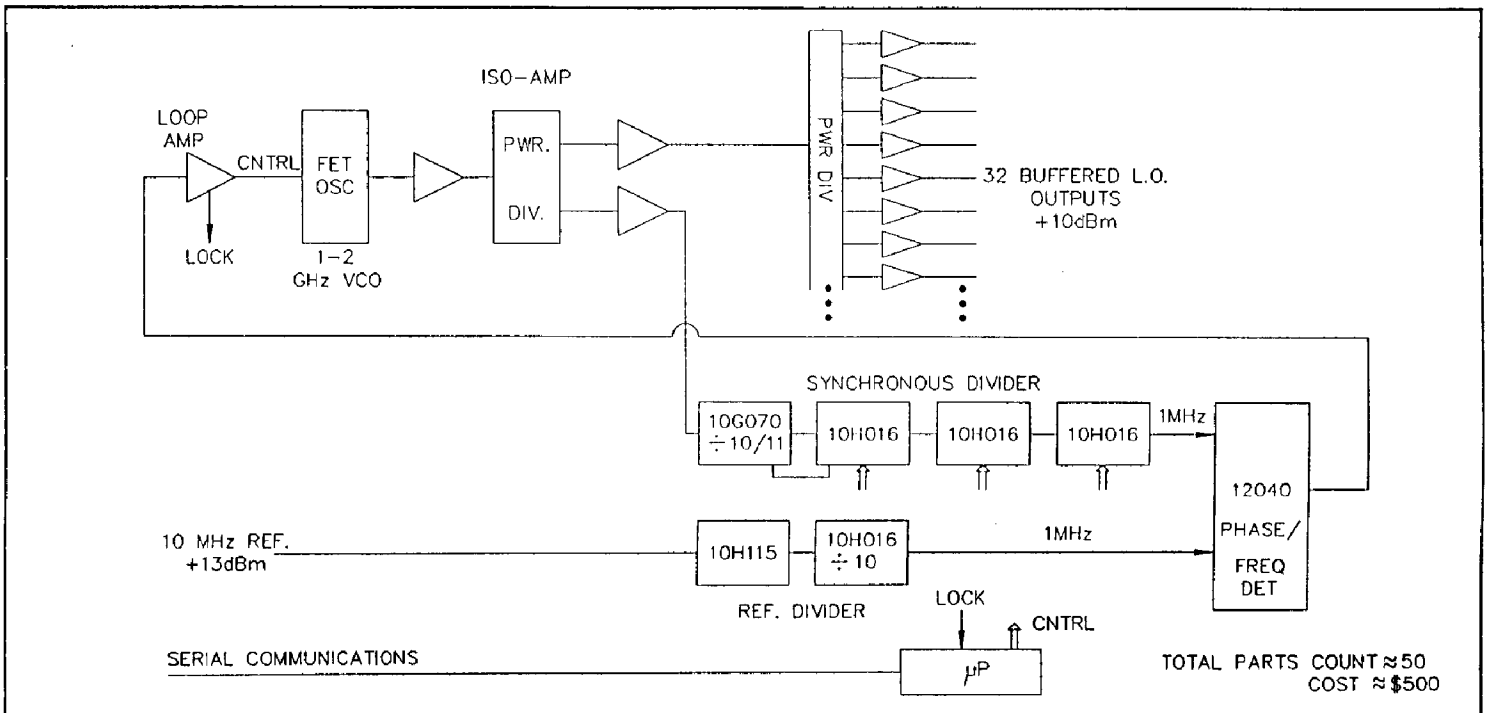

NOTES:

SPECS: FREQ. 1–2 GHz
STEPSIZE 1 MHz
PHASE NOISE <1 DEG
NO. OUTPUTS 32

SHOP NOTES: UNLESS OTHERWISE SPECIFIED

1. DIMENSIONS ARE IN INCHES
2. TOLERANCE ON DIMENSIONS
   FRACTIONAL ± 1/64
   DECIMAL .XX ± .01
   DECIMAL .XXX ± .005
   ANGULAR ± 0°30'
3. SURFACE ROUGHNESS PER MIL-STD-10
4. REMOVE BURRS AND BREAK SHARP EDGES 1/64 MAX.
5. SCREW THREADS PER MIL-STD-8
6. ALL DIMENSIONS TO APPLY BEFORE PLATING OR CONVERSION COATING.

USED ON

NEXT ASSEMBLY

WEIGHT

SCALE NONE

CLASSIFICATION

		DATE
DRAWN FOR	A.E.ROGERS	1/90
DRAWN BY	C.KOSTKA	1/90
CHECKED BY		

APPROVALS: PROJECT; ENGINEER; MATL. & PROCESS; STRUCTURES; THERMAL; MECH. ANALYSIS

NORTHEAST RADIO OBSERVATORY CORPORATION
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS

1-2 GHz
BASEBAND CONVERTER
LO SYNTHESIZER

FILE NAME	DWG. SIZE	DWG. NO.	REV.
BCLOSYNT	A		

Fi[illegible] A.2.1.1

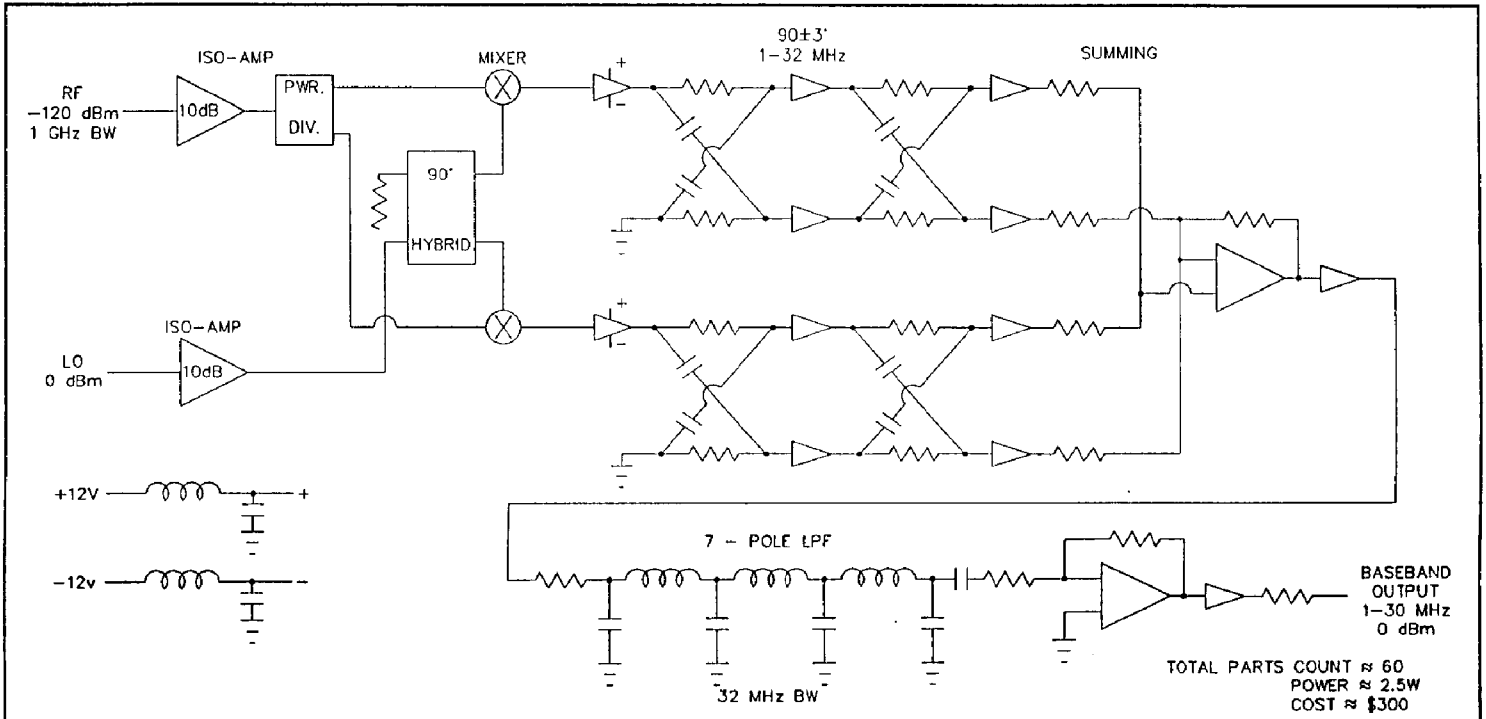

NOTES:

▷ = HA-5002 BUFFER AMP

▷ = 5539 ULTRA H.F. OP. AMP

SPECS: RF 1-2 GHz
LO 1-2 GHz
I.F. RANGE 1-30 MHz (3dB)
IMAGE REJECTION 26 dB
CONVERSION GAIN 35 dB

SHOP NOTES: UNLESS OTHERWISE SPECIFIED

1. DIMENSIONS ARE IN INCHES
2. TOLERANCE ON DIMENSIONS
   FRACTIONAL ± 1/64
   DECIMAL .XX ± .01
   DECIMAL .XXX ± .005
   ANGULAR ± 0°30'
3. SURFACE ROUGHNESS PER MIL-STD-10
4. REMOVE BURRS AND BREAK SHARP EDGES 1/64 MAX.
5. SCREW THREADS PER MIL-STD-8
6. ALL DIMENSIONS TO APPLY BEFORE PLATING OR CONVERSION COATING.

USED ON

NEXT ASSEMBLY

WEIGHT

SCALE NONE

CLASSIFICATION

DRAWN FOR	A.E.ROGERS	DATE 1/90
DRAWN BY	C.KOSTKA	1/90
CHECKED BY		

APPROVALS: PROJECT, ENGINEER, MATL. & PROCESS, STRUCTURES, THERMAL, MECH. ANALYSIS

NORTHEAST RADIO OBSERVATORY CORPORATION
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS

BASEBAND CONVERTER
IMAGE REJECTION (SSB) MIXER

FILE NAME	DWG. SIZE	DWG. NO.	REV.
BCIRSSBM	A		

Fig. A.2.1.2

B. REAL TIME DIGITAL PROCESSING

B.1 Introduction

Several methods are available to implement the real time data processing tasks for the sub-mm array. All of the options under consideration (so far) hinge on custom or semi-custom integrated circuits just now becoming available.

Since the data processing system is scheduled to begin toward the end of the array project, we propose that a final decision on the design be postponed for as long as possible. This will allow us to take maximum advantage of the circuit experience now being generated as well as any new developments that may, and probably will, surface.

A reasonable design and cost analysis is still required for planning purposes, however.

We have therefore selected a design which could be started tomorrow. All parts and techniques are available now. We fully expect that a complete re-assessment will be made before the final design is committed.

B.1.1 Processing Options

In this design estimate we limit ourselves to the digital processing techniques now available.

Recently, several integrated circuits have been developed which support the data processing requirements of the sub-mm array. Two correlator chips have just entered their production phase and an FFT chip set is in prototype development. A 2 GHz A/D and FFT chip set is reportedly under development in Japan. The A/D converter is complete but the FFT device(s) is still a few years away.

Finally, field programmable gate arrays continue to expand, raising the possibility of a device tailored specifically to SAO's requirements.

Each of these options is discussed more fully in Section B.8.

B.1.2 Design Goals

Our object here is to provide a real design, using real parts with known costs to allow upper and lower level estimates to be made. The components and prices reflect our experience in the construction of the Haystack Wideband Spectrometer (1989-1990).

These costs will change in the next two years, but if the integrated circuit industry continues as it has been going, we can expect to see a reduction in price or an improvement in performance (or both) when the final design phase starts.

One last ground rule - in the following discussions, "bandwidth" refers to the Nyquist bandwidth. "Usable" bandwidth and "effective" bandwidth can be defined by the reader - see Appendix D.2 for further discussion.

B.2 Processing Requirements

Each station in the array will generate a 4 GHz wide band of information coded into 8 G samples/sec of digital data.

The six stations will form 15 possible baselines with the following characteristics:

IF Channels per Station	4 - 2 RCP, 2 LCP
BW/IF Channel	1.024 GHz
Total Bandwidth	4 GHz
BBC's per IF Channel	32 - 32 MHz BW each
BBC's per Station	128 - 32 MHz BW each
Sample Rate	64 M samples/sec
Quantization	2 bits/sample (4 levels)

The wide band signals have been split into 128 bands to match the selected processor capabilities. Higher speed processing modules would reduce the number of baseband converters required but increase the cost of the distribution and processor electronics.

This match of converters to sample rate has been chosen to optimize the expected costs of analog and digital processing elements. Again - this is open to re-evaluation in the final decision phase.

B.3 Processor Design

The processor design selected will handle data on a "baseline" basis. That is, signals from station "X" will be processed against those of station "Y", channel by channel and frequency band by frequency band. Delay RAMs are provided to compensate for the interferometer delay. 100 lags are needed for each Km of baseline delay.

We will now discuss the distribution of the station signals, the circuits employed and the technique, used to process high speed samples using lower speed hardware.

B.3.1 Baseline Signals

The processor is organized on a "baseline" basis (as opposed to a station basis). Each baseline processor will receive 128 signals from each of station "X" and "Y". A representative slice of the processor is shown in Figure B.3.1.

Each polarization of each frequency band is processed in 4-32 channel cross-correlators. Data samples are accepted at 64 M samples/sec. Four correlation functions are derived from each channel pair, yielding 256 points per octal element to the post processing electronics.

One may think of these as pages of a book - each page represents a 32 MHz band in each of two polarization (Video A, B, and Video C, D). There are 128 upstream baseband converters driving the page inputs - therefore there must be 32 pages to accommodate the signals generated.

What we do with the inputs are discussed in the following sections.

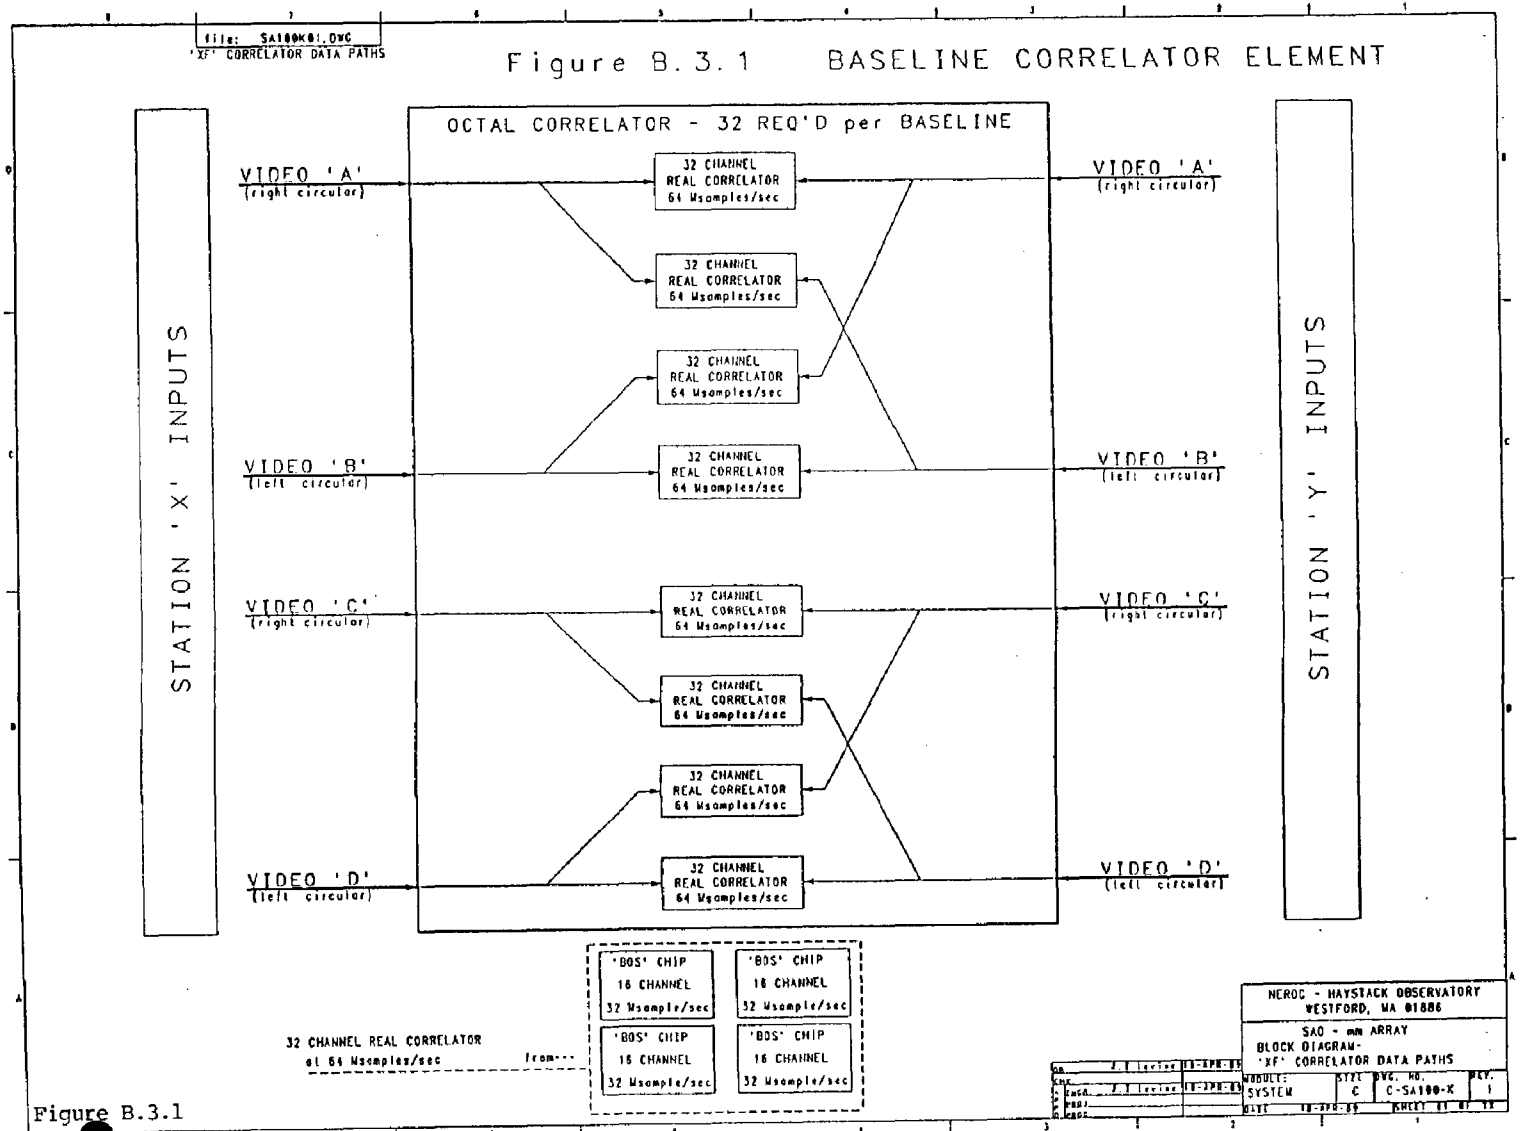
file: SA100K01.DWG
'XF' CORRELATOR DATA PATHS
Figure B.3.1 BASELINE CORRELATOR ELEMENT
OCTAL CORRELATOR - 32 REQ'D per BASELINE
STATION 'X' INPUTS
STATION 'Y' INPUTS
VIDEO 'A' (right circular)
VIDEO 'B' (left circular)
VIDEO 'C' (right circular)
VIDEO 'D' (left circular)
32 CHANNEL REAL CORRELATOR 64 Msamples/sec
'BOS' CHIP 16 CHANNEL 32 Msample/sec
32 CHANNEL REAL CORRELATOR at 64 Msamples/sec
from---
NEROC - HAYSTACK OBSERVATORY
WESTFORD, MA 01886
SAO - mm ARRAY
BLOCK DIAGRAM-
'XF' CORRELATOR DATA PATHS
MODULE: SYSTEM
SIZE C
DWG. NO. C-SA100-K
REV. 1

Figure B.3.1

B.3.2 NFRA Correlator Circuit

Albert Bos and his associates at the Netherlands Foundation for Radio Astronomy (NFRA) designed one of the first gate arrays tailored to coarse quantization correlation. Targeted to support the Anglo-Dutch telescope in Hawaii, this was one of the first gate arrays to support the type of data processing required by the radio astronomy community. The availability of semi-custom gate arrays has opened a new world to the 'casual user'. It is no longer required that one order hundreds of thousands of circuits to justify the setup charges.

Figure B.3.2 displays the block diagram and specifications of the NFRA correlator circuit. (Also referred to as the "Bos" chip in reference to its designer.)

The circuit is manufactured by LSI Systems in California and uses an 8000 gate CMOS array. Initial production prices have dropped from $70 to $40 per chip in the first year. Several facilities are engaged in designing correlator modules using the Bos chip.

The Dwingeloo telescope has designed a 40 M sample/sec 1024 channel correlator module for use at the Anglo-Dutch facility in Hawaii. A smaller copy is being assembled for the Haystack Observatory. Investigators at the University of California, Berkeley, and the Onsala telescope are both designing alternate modules tailored to their requirements. Commercial applications are being explored by Interferometrics, Inc. in the Washington D.C. area.

Many facilities are actually engaged in exploring the uses of these new devices. We expect that the problems that will arise, will be solved, in the next few years. One such problem is the trade-off of chips versus processing speed. A well understood technique is used in the SAO array design.

B.3.3 Multi-processing Techniques

The careful reader has probably noticed that we intend to process 64 M sample/sec data using 50 M sample/sec chips. In this section we attempt to explain the multi-processor technique used.

Simply stated, one can cross-correlate signals at N samples per second by dividing each signal into two channels, each running at N/2 samples/sec. The four signal streams (two streams of odd and even samples) are cross-correlated to yield Odd*Odd, Even*Even, Even*Odd and Odd*Even. Results of the four correlation functions are summed, pairwise, and yield the cross-correlation function expected at N samples per second.

Because of the pairwise summation, only half of the expected points (a 32-point cross-correlation for 4 16-point Bos chips) are generated. A processing factor of 2 is adequate for the sub-mm array project. Factors of 4, 8, and 16 are being designed into the Haystack Spectrometer.

Here too, we expect the design experience of the next two years to help in the final SAO decision process.

file: SA100K03.DWG
NFRA CORRELATOR CHIP

Figure B.3.2 NFRA CORRELATOR CHIP

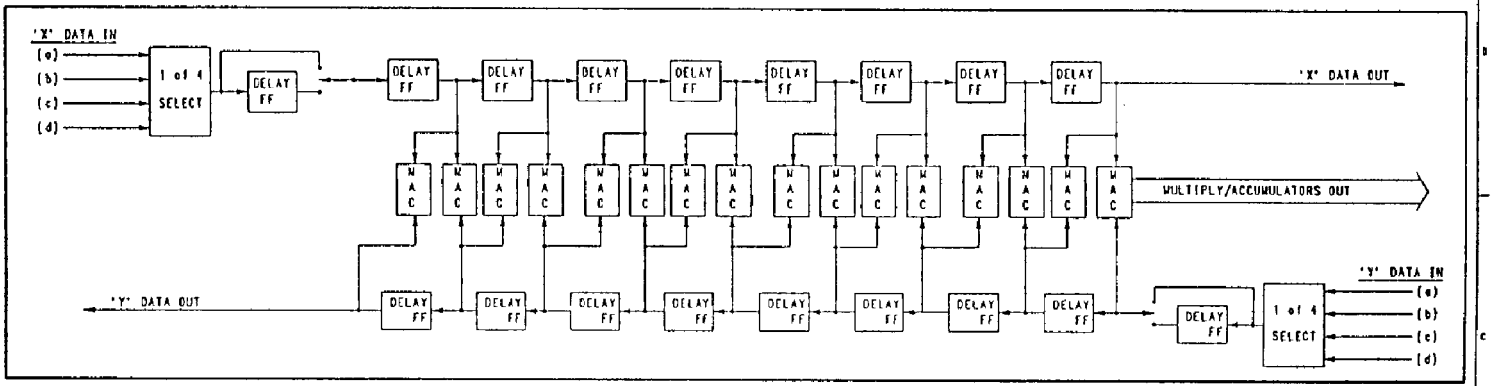

Minimum clock rate	:	55 MHz - worst case commercial
Number of lags	:	16
Multiplication scheme	:	reduced 2*2 bit (4 level)
Degredation factor	:	1.146
Accumulation stages	:	20 or 22
Prescaler stages	:	4 or 6
Data storage	:	16 bit on chip latches - each stage
Power dissipation	:	21 mW per MHz
Cost (1000 pieces) 1989	:	$40

NEROC - HAYSTACK OBSERVATORY
WESTFORD, MA 01886

SAO - mm ARRAY
BLOCK DIAGRAM- NFRA CORRELATOR CHIP

MODEL:	SIZE	DWG. NO.	REV.
SYSTEM	C	C-SA100-K	1

Figure B.3.2

B.4 Physical Organization

In this section we will attempt to give the reader a "feel" for the size of the array real time processor. The bottom line is envisioned as a system approximately twice the size of the Haystack Mark IIIA correlator.

This initial design will require three equipment bays, each containing five VME card cages plus one "contingency" bay for signal switching and communications. The three processor bays would contain the fifteen baseline processor chassis.

B.4.1 Baseline Correlator Chassis

Each of the fifteen baseline correlator chassis will process the signals between station "X" and "Y" as shown in Figure B.4.1. If auto-correlation is simultaneously desired from all six stations another six correlator chassis are needed. Alternately auto-correlations can be provided by assigning some modules to an auto-correlation mode. A baseline chassis will contain sixteen correlator modules plus a control unit and a single board computer (SBC). The modules will conform to the VME specification and will be 9V*280 mm (or 400 mm) size. Chassis size will be 19X17X15 inches (approximately).

Five baseline chassis plus power supplies and cooling equipment will be housed in a standard eight foot equipment bay. Each correlator module will process two frequency pages as described in Section B.3.1.

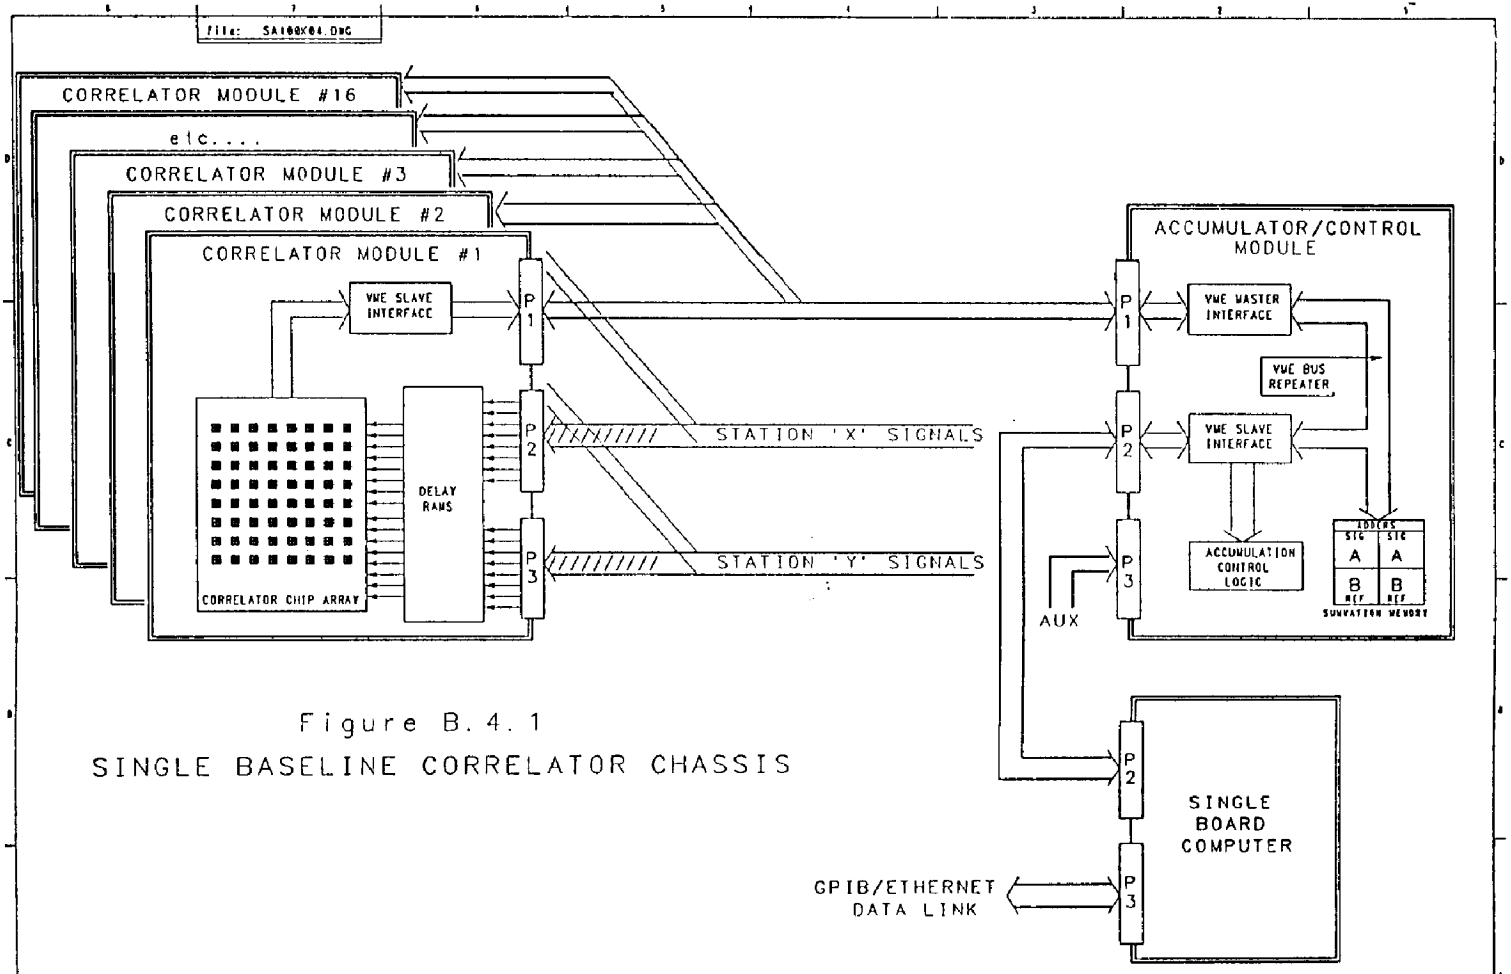

Figure B.4.1

SINGLE BASELINE CORRELATOR CHASSIS

Figure B.4.1

B.4.2 Correlator Module

One possible correlator module layout is shown in Figure B.4.2.1. Each module will contain two octal correlators as shown in Figure B.3.1. Sixteen such modules will support a single baseline processor.

Drive and support circuitry is delivered over one or two VME backplanes to a central control module and single board computer. The control module will be responsible for the application of specific details of the system operation. The commercial SBC will be responsible for command, control and final output using a standard data transfer protocol.

A complete baseline processor would appear as shown in Figure B.4.2.2.

file: SA100A01.DWG

Figure B.4.2.1 MODULE LAYOUT - 9U * 280 mm

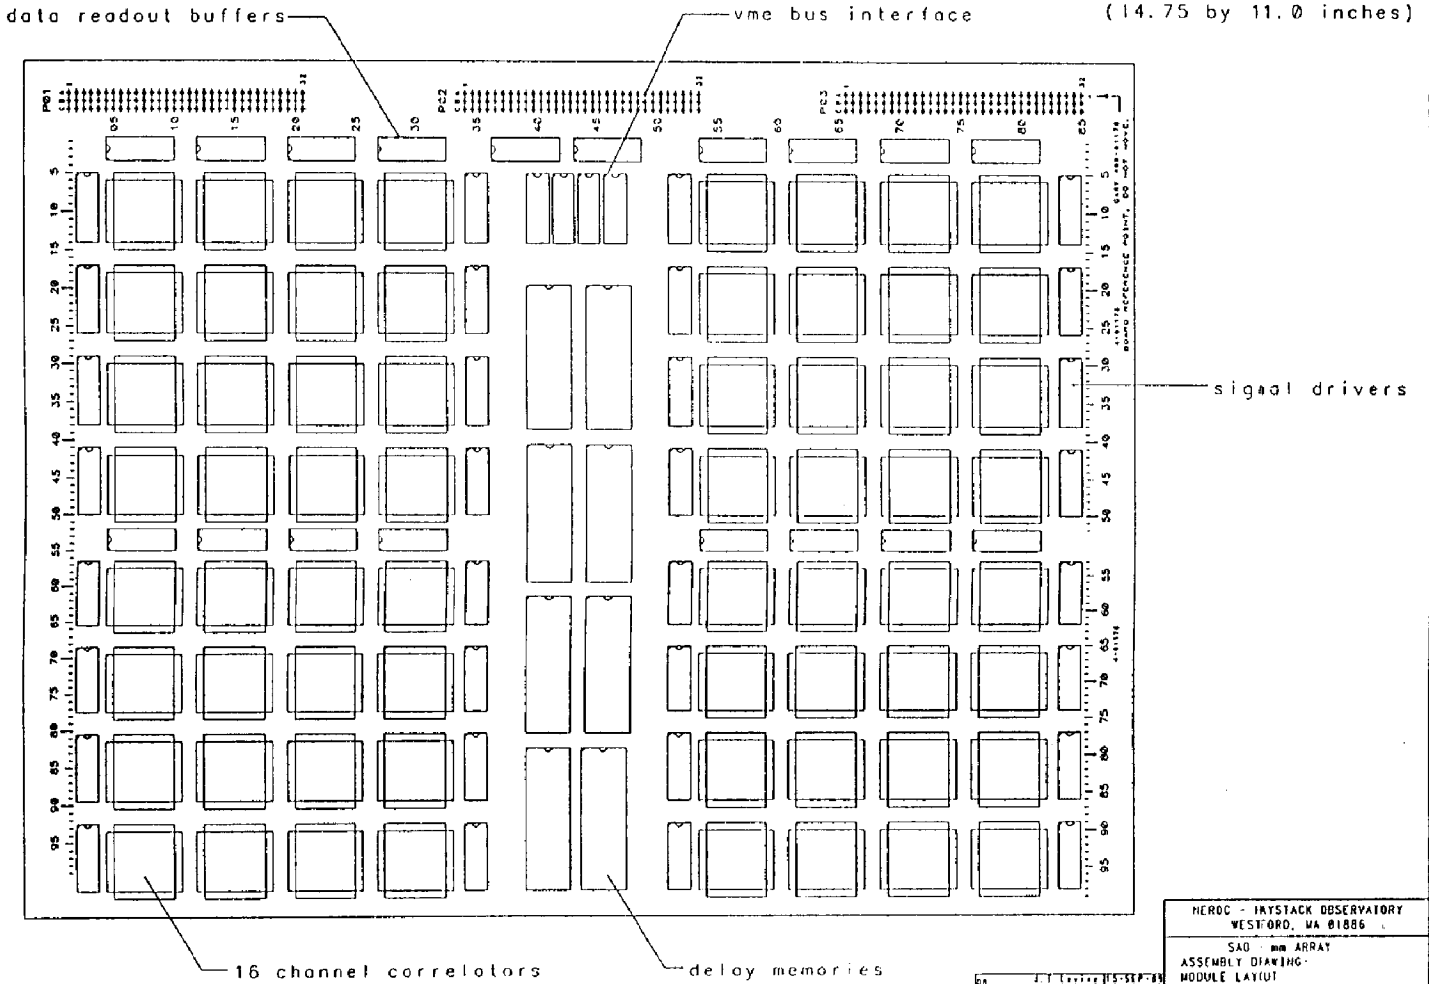

Figure B.4.2.1

D-SA100-A-002

Figure B.4.2.2

CHASSIS ASSEMBLY - SINGLE BASELINE CORRELATOR

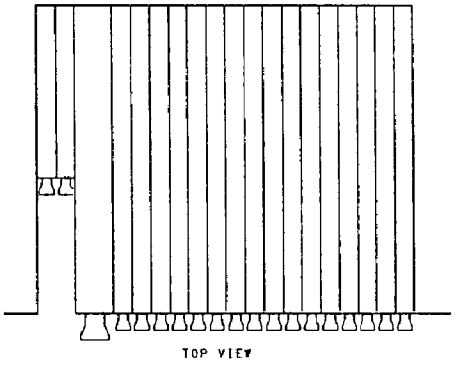

BACKPLANES

A-	20 POSITION VME 'P1'
B-	16 POSITION VIDEO SIGNAL
C-	16 POSITION VIDEO SIGNAL
D-	4 POSITION VME 'P1'
E-	4 POSITION VME 'P2'

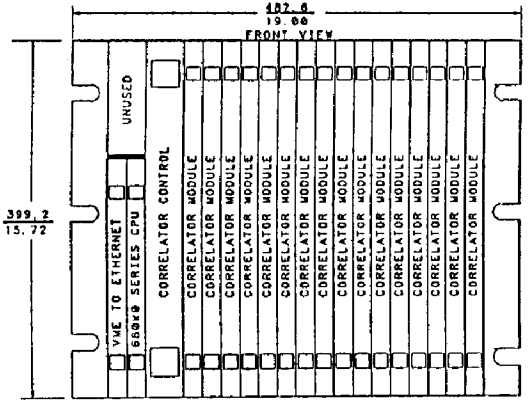

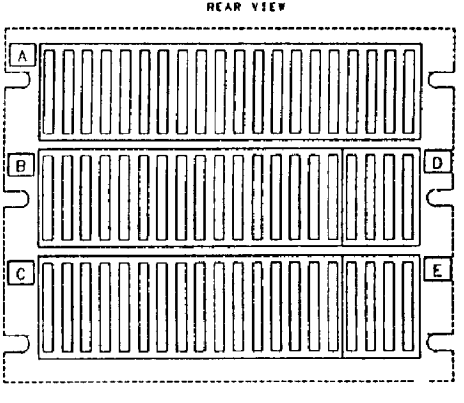

D-SA100-A-002

NRAO - GREEN BANK OBSERVATORY
WESTFORD, MA 01886

SAS 90m ARRAY

CHASSIS LAYOUT -
SINGLE BASELINE CORRELATOR

Figure B.4.2.2

B.5. Other Processor Requirements

There are a number of other processing tasks that have not been mentioned. Compared to the main processor elements, their impact is small but they should be listed.

B.5.1 Fourier Transforms

The system described yields cross-correlation functions. At some point these need to be transformed into the frequency domain for further processing. Compared to the real time correlation task, we feel that the transformation will be a minor task and its location best left for a later decision.

For a worst case scenario, a commercial FFT chip will generate a 1024 point transform in 500 msec and thus allow each correlator module to operate on a four millisecond unload cycle. If that high a rate is not required, FFT chips may be shared or the new data transferred and the transform performed by the array processor.

B.5.2 Data/Spectra Link to Array Processor

The data transfer to the array processor will have to be specified at some point. Preliminary studies indicate that the post-processing tasks greatly exceed the real time data transfer requirements. This is considered to be a decision best left until later.

B.5.3 "Fractional bit" Correlation, Dicke Switching and Interferometer Sideband Separation

The main reasons for relatively short accumulation time is the need for various rapid post correlation operations as follows:

1] "Fractional Bit" Correction
Accounts for the quantization of the delay tracking in units of 31 ns requires that a correction in the cross-spectral domain be made on a time scale short compared with the bit shift time which is 15 seconds for a 20 Km baseline.

2] Dicke Switching
It may be desirable to phase switch or load switch to reduce systematic errors in which case the correlation functions need to be separately accumulated for each part of the switching cycle.

3] Sideband Separation
If the first local oscillators are fixed (and do not track the fringes) the upper and lower sidebands of the front end will have fringe rate of opposite sign (or at least at different rates - depending on L.O. scheme) and can be separated by separately accumulating (in computer memory) the cross-spectra which have been rotated with opposite signs.

B.6 Cost Analysis - Parts

The cost estimates used for this section reflect the 1989 prices paid during the construction of the Haystack Wideband Spectrometer. While integrated circuit parts tend to decrease in cost, mechanical and power expenses are unpredictable. In general, however, we usually expect either a decrease in cost or an increase in performance over a two year period. Sometimes we get both.

Estimated parts cost for the digital processor totals $970K as detailed in Table B.6.1.

File: SA100D04.DWG

Table B.6.1 DIGITAL PROCESSOR COMPONENTS

CURRENT COSTS (1989 $)

NFRA CORRELATOR CHIPS	64*16*15 @ $40	$ 614,400
RAM DELAY GATE ARRAYS	8*16*15 @ $40	$ 76,800
VME INTERFACE GATE ARRAYS	2*16*15 @ $20	$ 9,600
CORRELATOR CIRCUIT PANELS	16*15 @ $750	$ 180,000
CORRELATOR CONTROL PANELS	15 @ $1500	$ 22,500
SINGLE BOARD COMPUTERS	15 @ $1000	$ 15,000
ETHERNET LINKS	15 @ $500	$ 7,500
VME CHASSIS	15 @ $300	$ 4,500
VME BACKPLANES	3*15 @ $300	$ 13,500
SIGNAL BACKPLANES	2*15 @ $200	$ 6,000
RACKS	3 @ $1000	$ 3,000
POWER SUPPLIES	15 @ $750	$ 11,250
COOLING EQUIPMENT	6 @ $200	$ 1,200
RACK CABLES and CONNECTORS	3 @ $1500	$ 4,500
digital processor.....		$ 969,750

NERDC - HAYSTACK OBSERVATORY
WESTFORD, MA 01886

SAO - mm ARRAY

- CURRENT COSTS/ 'XF CORRELATOR

MODULE:	SIZE	DWG. NO.	REV.
SYSTEM	C	C-SA100-D	1

DATE 19-SEP-89 SHEET 01 OF 01

Table B.6.1

B.7 Cost Analysis - Labor

A detailed manpower estimate is difficult to make at this stage of the project. Decisions yet to be made will strongly affect the amount and type of labor required. What follows is a very rough estimate based on our experience constructing the Mark IIIA VLBI processor.

B.7.1 Hardware Design

If we assume that existing integrated circuits will be used as the basic blocks for the system, approximately seven man-years of labor will be required for the hardware design effort. Resources would be allocated as follows:

2 Digital Design Engineers - 18 Months Each
- Signal interconnections
- Correlator modules
- Backplanes and chassis
- Test and checkout facilities

1 Analog Design Engineer - 12 Months
- Baseband Converters
- L.O. Synthesizers
- A/D conversion
- Signal transmission system
  (assume this task is done by Recvr. group)

2 Engineering Aide/Senior Technicians 18 Months Each
- Prototype support service
- Test and diagnostic construction
- Parts and documentation support

B.7.2 Software Design

Until the processor design is finalized it is impossible to decide which tasks will require firmware support (as opposed to hardware). We can state, however, that at least two software engineers will be required during the entire design phase to insure that the system is designed for testability.

The large number of circuits employed makes checkout of a system this size very expensive unless the 'hooks' are installed at the very start. Because of the complexity, one finds that debugging and fault location cannot be done by a $5 per hour technician with a voltmeter. You end up using $30 per hour engineers and costly diagnostic tools.

It is strongly recommended, therefore, that software/firmware funds be allocated at the very beginning to guarantee a testable, robust system.

The software design costs which can be specified at this time total about five man years, distributed as follows:

2	Software Engineers - System design - Test, checkout, diagnostics	-	18 Months Each
2	Scientist/Programmers - System setup - Communication links - Math algorithms	-	12 Months Each

B.7.3 Construction, Assembly, Test

The parts costs detailed in Section B.6 include items such as the insertion of components and the assembly of modules. The software/firmware budgets should cover the bulk of the testing and assembly expenses. We estimate an in-house cost of three man years of senior technician time will be required to coordinate the construction tasks required. The construction phase will require an engineer (18 months) for overall project supervision.

B.7.4 Summary

As stated, these labor estimates are very variable, depending on the final system design.

	# Man-Months
Project Engineer	18
Digital Design Engineer	36
Analog Design Engineer	12
Software Engineer	36
Scientist/Programmer	24
Engineering Aide/ Senior Technician	72

The total labor budget is estimated at about sixteen man years, or a total labor cost (with overhead) of about equal to the materials and services.

B.8 Other Processor Options

As mentioned in the introduction, several candidates exist for use as the basic data processing element. A brief description of each follows:

B.8.1 The NFRA Correlator Circuit

This gate array was developed by the Netherlands Foundation for Radio Astronomy for use in the Anglo-dutch project in Hawaii. Several thousand chips have been produced and the unit cost has dropped from $70 to $40 in the first year. The gate arrays are produced by LSI Systems in California.

Circuit specifications have been presented already (Section B.3.2) and will not be repeated here. A 1024 lag correlator module has been designed at NFRA to support the correlator chips.

Thirty modules have been fabricated and are currently being assembled and tested. The 30 printed circuits will be used at Dwingeloo, the Anglo-Dutch Telescope and the Haystack observatory in on-line spectrometers. Several are being supplied to the University of Massachusetts and MIT Lincoln Laboratory for evaluation studies. A significant amount of design and operational experience will be in place by the time a final selection for the sub-millimeter array is made.

B.8.2 250 MHz Correlator Circuit (CalTech)

A very high speed gate array has been developed at the California Institute of Technology. Using ECL technology, the chips are manufactured by Motorola and currently cost $50 each.

This circuit operates at a 250 MHz sample rate and contains eight correlation channels. Four level (2-bit) samples are processed. A 4-bit pre-scaler buffers each channel's output. The higher order accumulator bits are contained in a second custom chip, also developed at CalTech.

The higher sample rate more than offsets the reduced number of channels and the lack of internal accumulators as compared to, say, the NFRA chip. Operation of a complete module at a much higher clock rate raises problems of its own, however, The higher sample rates reduce the number of sub-channels required per station and may save money in the filter-video converter - A/D converter chain.

Again, we expect that the construction and operational experience that CalTech will gain in the next two years will prove invaluable to us in the final design decision.

B.8.3 32 MHz FFT Engine - NRAO

A fast fourier transform chip set has been developed at the National Radio Astronomy Observatory for use in the VLBA processor project. The basic unit performs a radix-4 FFT butterfly at a 32 MHz sample rate. Several circuits are cascaded to provide a nominal 1024 point transform covering a 16 MHz (Nyquist) bandwidth. The chips are manufactured by LSI Systems in California. Final production prices are not know yet.

Prototype circuits have been received and an FFT 'engine' module designed at NRAO. Modules and chips are now in the checkout/evaluation process.

This chip set supports the 'FX' mode of processing where the Fourier transform is performed on a station basis and only multiplication is required on a baseline basis. Questions regarding the optimum subchannel bandwidth to use depends on the ease of implementing a multi-engine module to process wider bandwidths. Again, NRAO's experience of the next few years will serve as a guide in the final decision process.

B.8.4 2 GHz A/D and FFT Engine (NTT)

At a recent STAG meeting, it was reported that Nippon Telephone and Telegraph was in the process of designing an impressive 2 GHz Fourier transform chip set consisting of A/D converter and FFT processor. The specifications attempt an astounding 500,000 points covering a 1 GHz bandwidth. the estimated cost is an equally astounding $750,000 per engine. (These are 'day 1' estimates and are by no means final.)

Even so, it appears that extremely wide bandwidth processing may be on the way. This may reduce the processor's front end costs significantly, but, because we are so early in the design cycle, nothing is firm enough to make decisions on. Lower resolution devices at a reduced price are anticipated.

B.8.5 Custom Design

A final possibility exists to create a new circuit, customized to the sub-mm array's requirement or to modify one of the existing chips to be more cost effective in this application.

This discussion will depend on the design and setup costs involved in the 1992 time period.

B.8.5.1 Impact of Semi-custom Chip Design

Another area requiring significant effort to evaluate concerns a possible semi-custom chip design, targeted specifically to the sub mm array's requirements. Consultations with just one of the manufacturers yielded the following preliminary information. An additional 6 man month effort plus a $20,000 design tool investment will be required before such estimates can be considered firm.

The most obvious design would quadruple the number of channels per chip and allow time-demultiplexed correlations to be performed in a single circuit. This would reduce the system chip count by almost a factor of 4. Each integrated circuit would produce a 32 point auto-correlation function on data sampled at up to 75 MHz, depending on the final design. At lower sample rates, a single chip would yield a maximum of 64 output points. The accumulator length would be expanded to allow up to 10 minute integration cycles and input signal switching could be tailored to the sub-mm-array requirements. Such a design, using the NFRA chip architecture is now possible, according to the manufacturer, LSI Systems.

B.8.5.2 Processing Speed

The new generation of gate arrays have an improved gate speed 35% faster than the previous systems. Much of this will be expended due to the increased interconnections on the improved circuit. Only a moderate increase in processing speed is expected.

B.8.5.3 Processor Size, Power

Power dissipation per gate might be reduced by as much as 30%; even so, the higher gate density per circuit would still increase the chip dissipation by a factor of three. Without additional cooling mechanisms, no significant system size reduction is expected.

B.8.5.4 Design Effort

Improved design tools now allow the use of high density 'pre-canned' macro functions which may significantly reduce the design and verification labor involved. Whether the initial NFRA cell design can be transported intact into the newer gate arrays remains to be seen. It may have to be 're-coded' but the initial design can always serve as a template.

The labor intensive effort of test vector generation may be significantly reduced by the use of the vendor's macro logic functions which include exhaustive test, verification and simulation software support. Silicon compiler design software will assist in fault coverage test generation vectors required but, even so, 2 to 3 man years of design labor will probably be required to implement an expanded version of the NFRA correlator chip. As you might expect, the design support software is not free.

Another approach would be to use the manufacturer's engineering staff for the detailed design. This would still require at least a man year of in-house labor for system definition and liaison. As you might again expect, the contracted design staff is not free. An accurate task description and vendor negotiations would require 6 man-months of in-house effort.

B.8.5.5 Schedule

A 12 month lead time will be required from initial design to prototype delivery. An additional 6 months of familiarization could also be necessary, depending on the experience of the design staff. A pre-design period of 6 months (1 man) would have to be scheduled to generate the design specifications, establish the design tool facility and begin vendor liaison.

We feel that the custom circuit design period can overlap system design by 6 months. Even so, this indicates that the preliminary investigation should begin 1 year early, with the goal of having firm specifications in hand 6 months before the detailed system design begins.

If the decision were made to proceed with a new semi-custom array, the next 6 months would be dedicated to the detailed chip design. The final 6 months of the custom effort will be spent completing the detailed design, generating test and acceptance criteria and building the necessary chip test facilities. Power consumption, pinout and rate specifications would have to be fixed by the start of the system design.

B.8.5.6 Costs

Preliminary Investigations:

Six man-months of engineering will be required to complete the initial effort and install the design tools required. Software costs should be about $20.000 assuming a 'SUN' or equal platform is already available.

NRE - Capital Equipment:

Based on the experience of the previous designs, we can expect non-recurring engineering charges of $100,000. We can also expect to spend another $15,000 on test fixtures and jigs to verify the prototype operation.

Labor:

A total of three man years of effort will be required for the semi-custom design:

Engineer	6 months	preliminary investigation
Engineer	12 months	chip design
Scientist/Programmer	9 months	test/verification vectors
Engineer/Programmer	9 months	prototype test fixtures

Parts:

The expanded correlator chip will cost around $200 each in quantities of 4000 as compared to $160 for a set of 4 existing chips. When the installation costs are factored in, the parts costs become equal.

B.8.5.7 Summary

There doesn't appear to be any obvious cost reduction to be gained from a simple duplication of this existing design into a denser package. In fact, development costs and design labor will result in an additional $150,000, plus 3 man-years of overall system cost.

Only small speed and power advantages would be gained from a denser design. If simple cooling methods are employed, no reduction in overall system size is to be expected.

Other architectures and manufacturers may result in some advantages but will require a 6 man-month effort and is beyond the scope of this report. New approaches will require additional labor costs and a higher risk factor than estimated in this section.

A new chip design will require a minimum lead time of 1 year prior to detailed system design, 6 months of preliminary effort and 6 months of initial circuit design.

B.8.6 Separate Continuum Correlator

A separate wideband continuum correlator could be used. To cover a 4 GHz bandwidth the I.F.s could be split into 8 500-MHz-wide bands. Wider bands would make the delay compensation very difficult and even 500 MHz bandwidth would require fiber optics delays for the larger delays. In order to obtain an approximate estimate for the cost of a continuum system consider a correlator with the following modules:

Module	Quality Required
Delay Line Section	6 x 8 x 16 = 768
Quadrature Multiplier	15 x 8 x 2 = 240

Each delay line section (1 bit out of 16 needed for each station) would provide a delay of 2^X x 50 picoseconds and high isolation switching to bypass the delay. The delay section could use coax cables for bits with delay less than 10 nanoseconds but would require fiber optics delay paths for the larger delays (in order to obtain a flat response over 500 MHz). Coax sections might cost about $200 each while fiber optics sections would cost more like $1,000 each owing to the cost of modulator and demodulator components. Quadrature multipliers would contain a quadrature hybrid and two analog multipliers. The analog multipliers might use sum and difference square law detectors using the identify

$$(A + B)^2 - (A - B)^2 = 4\,AB$$

and might cost about $200 each. The total parts cost of about $300,000 would have to be added and would involve some substantial development of the wideband delay sections. The delay sections would have to be temperature-controlled and trimmed to obtain the exact powers of 2 values needed. With computers, racks, and power supplies, the continuum correlator is likely to cost at least half of the proposed hybrid correlator.

Given that a separate continuum correlator is a substantial system, it would seem wiser to build one correlator for both continuum and spectral line. Another consideration is that it will probably be easier to achieve high accuracy in spectral line astrometry if the continuum calibrations use the same I.F. processing electronics.

C. SUMMARY OF CHOSEN DESIGN

C.1 XF Hybrid Correlator - Summary

A hybrid filterbank XF correlator is proposed for the sub-millimeter array as a system that could be built today with a well defined cost. The correlator would accept wide IF bands and segment them into smaller bands using many identical programmable baseband converters (BBC). The baseband outputs would then be correlated using VLSI gate array correlators. With the following design parameters:

#I.F. channels/station	4
normally	2 RCP and 2 LCP per station
I.F. bandwidth/channel	1 GHz (1-2 GHz range)
#BBCs (for 6 stations)	6 x 128 = 768
#BBC L.O.s	32
Filter bandwidth	32 MHz
#32 lag (64 Ms/sec) correlators	15x4x32x2=3840
Quantization	2 bits/sample

This correlator would provide:

Total # spectral channels	61,440
# spectral channels/baseline/ polarization	1024
all polarization products	RR,LL,RL,LR
Bandwidth in each polarization	2 GHz
Total bandwidth/station	4 GHz
Resolution	2 MHz
	(0.7 Km/sec at 860 GHz)

Switching is provided to allow higher spectral resolution by reducing the number of polarization products, or total bandwidth to a limit of 8 KHz (0.01 Km/sec at 220 GHz) when a single 32 MHz bandwidth in only one polarization.

Fringe rotation would be applied to the local oscillator at each antenna, at least to a precision of a small fraction of the inverse of the correlator accumulator dump rate. Residual fringe rotation and/or phase switching demodulation would be applied in fast firmware or software following correlation. All delay compensation and delay buffering to allow many "unit" correlators to be cascaded for additional spectral resolution would be performed in digital hardware.

C.2 Signal Switching

The first level of switching is provided at I.F. input. Each of the 4 I.F. inputs (labelled A, B, C, D) can independently select any of the available receiver I.F. channels. The next level of switching is provided at the output of the 2-bit samplers (also known as A/D converters). Thus correlator modules can be connected to baseband converters in a manner which uses all converters or fewer converters. The final level of switching is provided in the "Bos" chip itself. This final level of switching allows correlators to be "daisy chained" to increase the number of lags and hence increase the spectral resolution. To simplify the "daisy chain" it is assumed that the number of lags will increase in steps of two.

C.2.1 Examples of Observing Modes

The 3-levels of switching provide a lot of flexibility and a very large number of possible modes. The following constants are assumed:

1] Resolution goes from 2 MHz to 7.8 KHz in steps of 2.
2] Baseband filters are fixed at 32 MHz bandwidth.
3] Baseband converter frequency settings are quantized in 1 MHz steps.
4] The maximum number of spectral channels (all baselines and polarizations) is 61,440.
5] Correlator lags are 1/32 MHz = 31.3 ns apart.

We give a few examples of observing modes:

1] Continuum

Total Bandwidth	4 GHz
Resolution	2 MHz
Full Polarization	RR, LL, LR, RL on all 4 I.F.s

Note: First level switching allows the 4 I.F.s to come from any combination of available receiver I.F.s.

2] High Resolution Spectral Line/Single Polarization

Total Bandwidth	32 MHz (1 BBC per station)
Resolution	7.8 KHz

3] Spectral Line/Full Polarization

Total Bandwidth	64 MHz (2 BBCs per station)
Resolution	31.2 KHz
Full Polarization	RR, LL, LR, RL

4] Medium Resolution Spectral line/Full Polarization

Total Bandwidth	512 MHz (16 BBCs per station)
Resolution	250 KHz (0.25 Km/s at 300 GHz)
Full Polarization	RR, LL, LR, RL

5] Mixed Mode - Simultaneous Continuum and High Resolution Spectral Line

a] Continuum

Total Bandwidth	4 GHz
Resolution	2 MHz
Single Polarization	

Note:4 GHz could be four separate wavelength bands of 1 GHz each.

plus

b] Total Bandwidth	32 MHz
Resolution	15.6 KHz
Single Polarization	

Note: This mode example uses an equal amount of correlator hardware for continuum and spectral line.

C.3 Cost estimate and the cost equation

For this approach the hardware is dominated by the BBCs and correlators for such a large system. At present the "Bos" chip is a very cost effective correlator and 4 of these chips are needed for each 32 lag (64 Ms/sec) correlator. The following are estimates for the materials and services if parts were to be procured in 1990:

Digital Processor (From Table 6.1)		970K$
BBC	$300 (768 pieces)	231K$
BBC L.O.s	$500 (32 pieces)	16K$
Samplers		10K$
I.F. Distribution Power supplies, computers, etc.		200K$
		1427K$

Projections suggest that the hardware cost will come down (for example we expect the Bos chip to be available for $30 each in 1991 - if procured in large quantities).

A correlator which avoids special development could be built in two years. In practice it would probably be wise to build a "subset" in 18 months, test it thoroughly for six months, and then complete the system in another six months for a total of two and one half years.

The NRAO cost equation for the cost of the hybrid filterbank-correlator can be approximated by the equation (Weinreb 1983):

$$C = \frac{BMN^2}{J} + 0.6\,JN + 200 \quad (K\$)$$

Where B is the total bandwidth in GHz, M is the total number of spectral channels, J is the number of filters per antenna, and N is the number of antennas. The first term accounts

for the digital correlator, assuming 10^9 multiplications per second per dollar. Note that 2 BMN^2/J is just the bit multiplication rate needed for cross- and auto-correlation. The second term accounts for the analog filters and the third term for fixed costs.

If we substitute B=2, J=128, N=6 and M=1024 we get 1236 K$ or 1812 K$ if we increase the first factor by two to account for full polarization processing.

C.4 Equivalent Processing Power

In order to compare the proposed processor with other computing/signal processing systems it is useful to look at its computing power.

Digital Processor
Each of the 15,360 "Bos" chips in the proposed system has 16 multiplier/accumulators running at a clock rate of 32 MHz for a total of

8 Tera 2-bit Operations Per Second 8 TOPS

I.F. Processor
The I.F. processor provides the equivalent of a digital filter so that if it were implemented in digital hardware it would require

$4 \times 2 \times 10^9 \times 6 \times 8 \times \log_2 32 =$

2 Tera Floating Point Multiplies Per Second 2 TFLOPS

where

4	=	# of I.F.s
2×10^9	=	2 x 1 GHz BW
6	=	# Stations
8	=	Equivalent # Bits Precision
$\log_2 32$	=	# Multiplies in 32-point FFT

Other Machines

JPL WBSA SETI Processor is

5 Giga Operations Per Second	5 GOPS

An Intel 80486 is

30 Mega Instructions Per Second	30 MIPS
Most Powerful Cray-2 Supercomputer	1 GFLOPS
Mark IIIA VLBI Processor	100 GOPS
VLA Correlator	300 GOPS
VLBA Correlator	400 GFLOPS

While it is hard to compare floating point (FLOPS) operations with simple operations (OPS) it is clear that the proposed processor is a very powerful machine with very low cost (≈$250 per MOPS).

C.5 Conclusions

It is fortunate for us that this project comes at this time. It just happens to coincide with the availability of reasonably priced gate array circuits which perform the math functions required by the sub-mm array.

The custom/semi-custom digital field is moving so fast that we cannot decide on the best course of action to take two years hence. Reasonable estimates are still required however. We have designed a system using existing parts with known costs so as to provide upper and lower limits for planning purposes.

By 1992-1993 we expect that the ground rules will change and we all will be much smarter. We also expect that it will not cost more that $1600K in materials and services to build the sub-mm array processor. Labor costs are more uncertain - but should be less than the materials and services.

D. APPENDIX

D.1 Parts List for Baseband Converter and L.O. Synthesizers

Baseband Converter

Part	Description	Manufacturer	Cost	Qty.	Total
HA-5002	Buffer	Harris	5	12	60
NE-5539	Op.Amp	Signetics	6	2	12
TFM-11	Mixer*	Mini-Circuits	40	2	80
PDF-2A-1000	Pwr.Div.*	Merrimack	60	1	60
QFK-2-1.5GK	Hybrid*	Merrimack	50	1	50
MSA-0420	Amp	Avantek	10	2	20
PC Board			40	1	40
Connectors			5	3	15
Resistors			1	18	18
Capacitors			1	15	15
Inductors			5	3	15
TOTAL (All Commercial Components)					$385
TOTAL (With Components Marked with * in Stripline)					($195)

L.O. Synthesizer

Part	Description	Manufacturer	Cost	Qty.	Total
10H016	ECL	Motorola	5	4	20
10H115	ECL	Motorola	1	1	1
10G070	GaAs	GigabitLogic	200	1	200
12040	ECL	Motorola	1	1	1
	μP	Intel	50	1	50
	AMPS	Avantek	5	35	175
	PC Board		40	1	40
	Misc.				20
TOTAL					$507

D.2 Filter Shapes and Linking of Subspectra

The experience of the Kitt Peak hybrid spectrometer will be extensively used in the design of this instrument. NRAO has issued a special memo series on the hybrid spectrometer (known as HYSPEC). By using only one sideband from each converter, we will have the advantage of being free to choose the amount of filter overlap thereby offering some advantage over the Kitt Peak hybrid spectrometer. Eight pole filters should provide an acceptable shape function [30 dB bandwidth = 1.33 x (1 dB bandwidth) - see Weinreb, IEEE Trans. on IM-34, No. 4, Dec 1985]. A combination of total power measurement from the zero lag and comparison of the overlapping spectral points could be used to piece together the subspectra (see A. Emrich - Hyspec Memo #13). The task of linking the subspectra for the single antenna spectra usually requires higher accuracy than linking of the cross spectra.

The imperfect filter shape results in some inefficiency so that the effect bandwidth when operating with some overlap between subspectra would be about 75% of 32 MHz or 24 MHz. However, the continuum sensitivity when operating with overlap is only reduced by about 3%. (See "Sensitivity of VLBI", Rogers, et al, NASA Publication CP-2115, 1980.)

D.3 Minimum Number of Lags Needed for Continuum

We have proposed 32 lag real correlator whereas a complex correlator with perfect (unquantized) delay tracking requires only two (the zero lag in the "sine" and "cosine" channels) as in the VLA. With 32 lags the loss of energy from signal outside the ±16 lag range is under 1%. Reducing the number of lags from 32 to 16 increases the loss to about 2% and another factor of 2 reduction to 8 lags would increase the loss to 3%. However the Bos chip provides 32 lags when used as a quad (see Section B.3.3) and the number of lags can only be reduced to 16 by using a single chip to process 16 MHz bandwidth which would require twice the number of baseband converters. This would result in a cost reduction of about 200 K$. This cost saving option should be considered if the reduction of the widest bandpass, available without splicing subspectra is acceptable. Further reductions in the number of spectral channels cannot be made with the Bos chip without reducing the continuum bandwidth or giving up simultaneous polarimetry.

D.4 Cost Savings Options

While the proposed system is modular and can be easily adjusted to handle more stations and/or more bandwidth, there are some reduction options which would also reduce the capability. For example:

1] Reduced Number of Correlators

The number of correlators could be reduced by a factor of 2 by elimination of the simultaneous polarization cross-products. In this case, polarization studies using all 4 products would be limited to 2 GHz maximum bandwidth.

2] Sharing of Local Oscillators

The cost estimate is given for a correlator which already shares baseband L.O. synthesizers as much as possible providing more synthesizers (as already discussed in A.2.1) is an option which would cost more.

D.5 Maximum Sample Rate Selected for NFRA Chip

The final choice of sample rate (and bandwidth processed) may increase above the 32 Ms/sec used for this estimate. A detailed timing analysis based on the final system specifications is required but we can make some reasonable predictions using the NFRA chip.

While a single circuit is tested to 55 Ms/s, we do not expect to achieve that rate in a full system or even in a single module. Modules already exist however, which do support 40 Ms/s over 64 chips. A 50 Ms/s unit is in its design phase. Unavoidable interconnection delays and clock skews in a large system will most likely reduce these rates also. A reasonable estimate of these limits are:

single chip	55 Msamples/sec
single module (32-64 chips)	50
single chassis (16 modules)	40-45
single rack (4-8 chassis)	35-45