

ATA memo #39
The ATA Imager
By
Wilbert Lynn Urry
University of California
Radio Astronomy Dept.
Jan 28 2002

Abstract

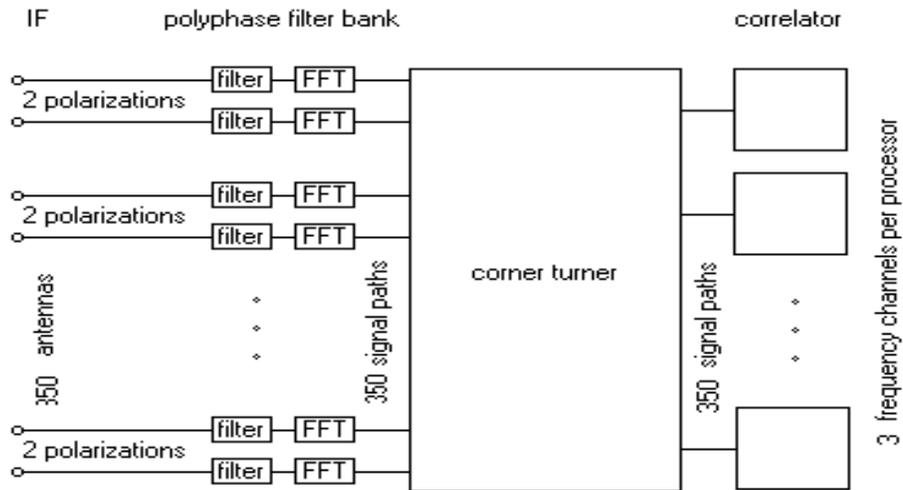
The Allen Telescope Array of 350 6-meter antennas will provide four independently tuned IF processor output ports from each antenna. Emerging from each output port will be two polarizations of 100 MHz bandwidth extracted from the .5 - 11.2 GHz RF band of the telescope front end. The IF processor will provide each output port with quantization to 8 bits and independent tracking of delay and phase for as many as 4 beams. Each beam may support a different "back-end". Different "back-ends" will be produced so that different experiments may take place simultaneously. An imaging "back-end" is described. The imager will rely on an FX correlator architecture. The signal from each antenna will be separated into as many as 1024 frequency channels and then all correlations for each channel will be measured. The channel separation of the FFT will be improved by using a modern polyphase technique. The polarizations of the 100 MHz baseband will pass to a pair of FFT based polyphase filterbanks, which will split the signal into 1024 channels. From the polyphase filter banks, the signals will proceed to a bank of correlators. Each correlator computes all 61425 baselines for each of the frequency channels. A serious problem exists in getting the signals from the polyphase filterbanks to the correlators. Using the innovation of a memoryless corner turner solves this problem. The corner turner converts the 350 streams of 1024 frequency channels into 350 streams of antennas of 3 frequency channels each. Each correlator sees a sequential stream of all of the antennas and can easily compute all of the baselines as the data enters.

We have considered several approaches to the imaging problem. We considered the XF (cross correlation before frequency analysis) approach. XF offers the advantage of coarse quantization with its simplified arithmetic, but in a modern situation where high levels of interference will be encountered, coarse quantization can be overwhelmed. XF has a serious disadvantage compared to FX in the presence of interference. The modern technique of using a polyphase filter bank in place of the FFT used in older designs, has the advantage of extending the width of the weighting function applied to the signal. In an XF design the weighting function is limited to the width of the sum of the lags. In order to gain windowing characteristics that are anyway near acceptable a factor of two more lags in an XF machine is required to achieve the same resolution in a FX machine using a polyphase filter bank. A proper weighting function is absolutely necessary in order to insure good channel separation. XF also offers the advantage of very regular simple design making the design process much quicker. The availability of pretested, predesigned cores for both custom chips and field programmable gate arrays eliminates this advantage. Finally, XF offers the advantage over FX (frequency analysis before cross correlation) of simpler interconnection between the antennas and the correlators¹. We offer the innovation of a corner turner that solves the interconnection problem.

We have also considered an FF (frequency analysis followed by Fourier beam forming or direct imaging) approach². Direct imaging is a very promising idea for radio astronomy in that beams are formed on the sky rather than an image which consists of the power in each beam³. The beams may be used for SETI searches or pulsar experiments as well as image creation. The problem with direct imaging is the fact that the computational cost of Fourier beam forming with a sparse antenna array becomes quite high. A sparse antenna array is required for high resolution imaging.

We propose using an FX architecture. Although the concept has been known for 20 years⁴ the largest such correlator previously built is the VLBA's, which handles 20 antennas at 128 MHz total bandwidth and 8192 total spectral channels. Our correlator with 61425 baselines as opposed to

210 will have a larger capacity by more than two orders of magnitude.



System

The system consists of three basic elements. The polyphase filter bank can use as many as 4 independently tuned complex 100 MHz bandwidth streams of 8 bit real and 8 bit imaginary samples. It can separate a 100 MHz stream into as many as 1024 complex channels of 4 bits real and 4 bits imaginary. Each antenna provides two polarizations. There are two filter banks per antenna. The next element is the corner turner. The corner turner converts the data from frequency channel order to antenna order. The third element of the system is the correlator. There are 350 signal paths from the corner turner to the correlator. Each correlator computes all 61425 baselines for each of the frequency channels. In order to keep the data rate from the antennas approximately the same as the data rate entering the correlators, each correlator unit will accommodate 3 frequency channels.

Filter Bank

A problem associated with the FX approach is the poor performance of the fast Fourier transform as a filter bank⁵. This problem of channel separation is also true of lag correlators as well. We propose to precede the FFT with a polyphase structure forming a polyphase filter bank⁶. The polyphase filter bank enables us to provide superior channel separation with very low sidelobes and very low aliasing⁷.

The Fourier transform, by itself, makes a terrible filter bank. If no windowing is applied (a square window) then there is severe leakage between the channels. Strong interference in one channel can mask signal appearing in other channels. Many authors have studied the problem of using a more favorable input window to improve its performance⁸, however performance improvement usually requires overlapping the input sample vectors⁹ to avoid signal loss. The limited window width also implies a loss in resolution. The result is at least a factor of two loss in throughput speed as well as resolution. A relatively new approach of placing a polyphase structure in front of an FFT in order to extend the input window can result in a filter bank with very good properties¹⁰.

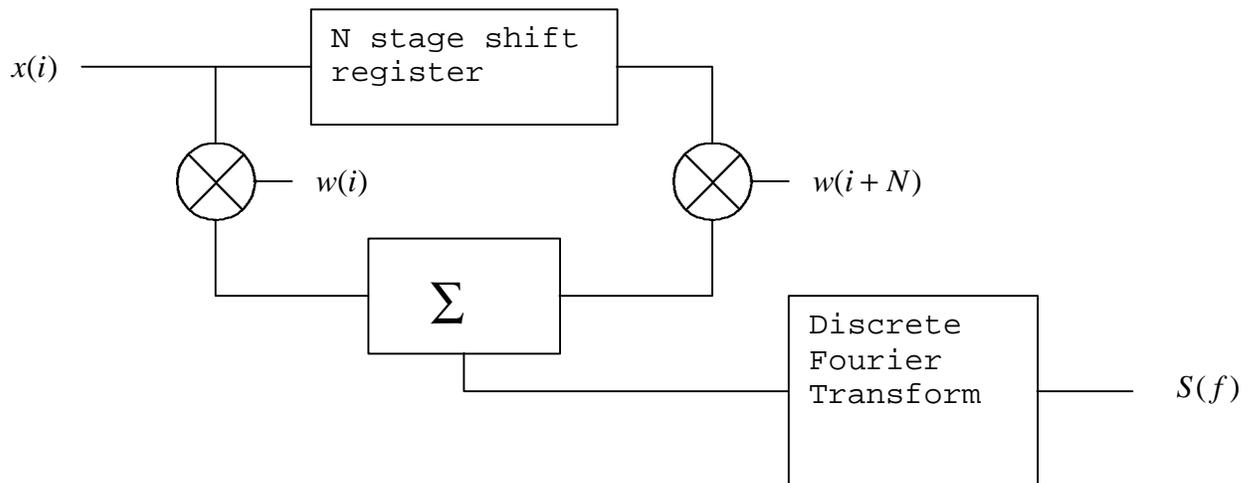
The discrete Fourier transform is usually written as

$$S(f) = \frac{1}{N} \sum_{i=0}^{N-1} w(i)x(i)e^{-j2\pi f i / N} \quad (1)$$

The $w(i)$ term is a windowing function that determines the pass-band shape of each of the bins. Since the exponential term is periodic in N , a more general discrete Fourier transform may be written

$$S(f) = \frac{1}{N} \sum_{i=0}^{N-1} \left(\sum_{a=0}^b w(i+aN)x(i+aN) \right) e^{-j2\pi f i / N} \quad (2)$$

Any number of input samples may be used in the construction of $S(f)$ and the window may be unlimited in extent. An implementation of this scheme for $b=1$ is shown below where the shift rate of the silicon is the Niquest rate of the signal.



In order to window the samples of $x(i)$ as they come in, they must be multiplied by the sampled window function $w(i)$. The window function samples are stored in a look-up table and multiply the input samples as they go by. The result is entered into a pipeline Fourier transform in sequence. This arithmetic is needed if any windowing is required. In order to expand the size of the input window, the input sequence is passed to an N stage shift register which delays the input sequence by N, the period of the Fourier transform. Another look-up table is used containing the samples of the extended window. A second multiplier multiplies the delayed input and an adder adds the result to the result of the first multiplier. The sum then enters the Fourier transform in sequence. Any degree of expansion may be added by adding as many shift registers and multipliers as are needed. A factor of two increase in window width requires only two multipliers.

Extending the window is not an option with an XF machine. If the correlation measurement is done first, then you are stuck with the number of lags the machine produces. Any weighting must be applied to the lags that have been measured. The result is loss in resolution and higher side lobes in the spectrum. In order to equal the performance of a polyphase filterbank, a lag correlator must have as many lags as the window width of the polyphase filter bank. FFT cores are available for field programmable gate arrays, as are polyphase filters making the design time required for such a system very short. Preliminary design results indicate that a filter bank using 8-bit arithmetic can achieve side lobes that are 60 db down.

Corner Turner

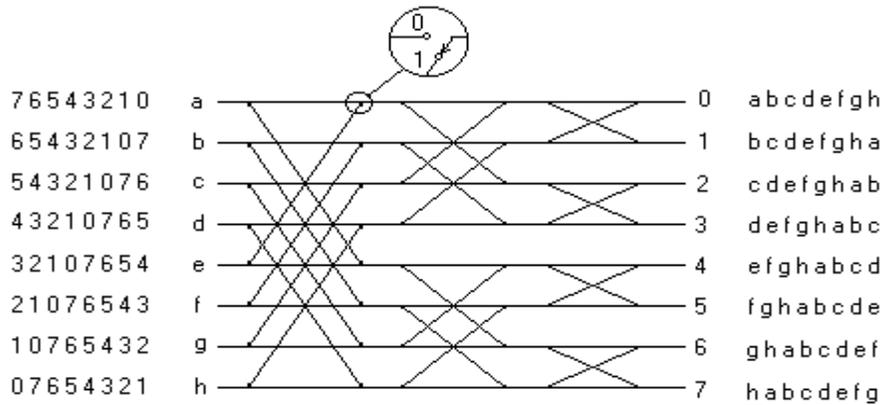
The design of an imaging system with 350 antennas will serve as a proving ground for ideas that are necessary for the future development of even larger arrays. There is, for instance, the problem of connecting 350 antennas to 61425 correlators. While the speed and density of logic continues to accelerate making large antenna arrays feasible in terms of arithmetic capabilities, the task of interconnecting the electronics remains a problem. Indeed the cost of interconnections may dominate the cost of the machine. In

order to cross-correlate all of the antennas, each of the 350 antennas must go to 349 other places where the correlations are done. 122150 independent wide band signal paths (neglecting auto-correlations) are required for an XF machine. The situation becomes worse when a FX machine is considered since each of the separate frequency channels must be independently correlated. Critical attention must be paid to timing, crosstalk, ground-loops and reflections in each one of these lines. The arrival time of each sample from widely separated sources must be carefully controlled. The assembly of even a modest 10-antenna system is a challenging balancing act.

Any imaging system must have access to all of the antennas. Our basic interconnection strategy is to break up the high bandwidth signal emerging from the each of the antennas into several parallel lower bandwidth slices. A transformation is applied to the data. A stream of frequency slices for each antenna is converted into a stream of antennas for each frequency slice. The slices are then routed to independent imagers¹¹ that work on all of the antennas for each slice. The device that performs this transformation is a corner turner.

A corner turner is a device that performs a matrix transpose¹². Data is read into the rows of an array and, once it is filled, the same data is read out of the columns. Corner turners are usually implemented using memory. All of the frequency slices are read into the memory in frequency sequence from all of the antennas. As soon as all of the frequency slices have been read in, the memory is reconfigured to read out. Another memory continues to read in the next batch of frequency slices so that no data is lost. The order of the read out is changed so that the data is in antenna sequence. All of the antenna information can now be directed to a single correlator that can calculate all of the baselines for a single frequency slice. A separate correlator can be provided for each slice. Notice that once the samples have been placed in proper order by the corner turner, their arrival time at the correlator is no longer important so that each of the correlators is independent. A serious flaw of other correlator schemes is the fact that all of the correlation elements must be in exact synchronism.

The use of a corner turner applied to a correlator system is unique and the use of a memory to accomplish the turn would be adequate to produce an economic system. We



have developed a corner turner that requires no memory and uses only switches. It promises to provide a system that is less costly and more versatile than one using memory.

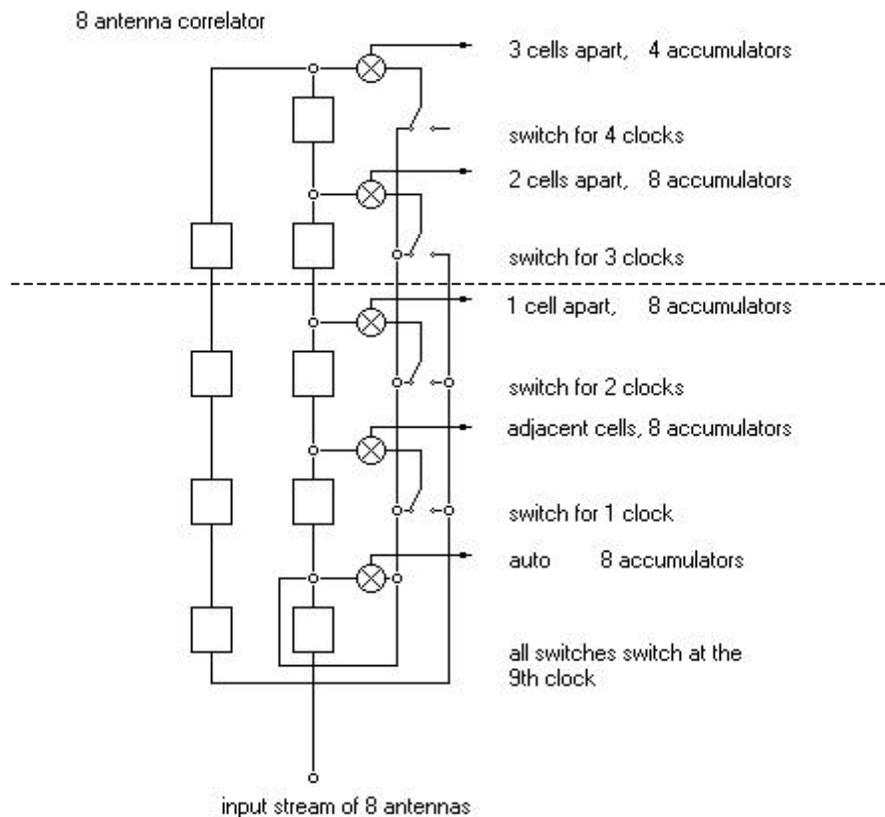
The configuration has the appearance of a modulo 2 fast Fourier transform butterfly arrangement. At the right end of each butterfly, a switch is located rather than the phase rotation and add of the fast Fourier transform. When the switch is in the upper position it will be called the zero position. When it is in the lower position it will be called the one position. The data enters the array from the left in the diagram. The switches all change position after the passage of one cell time. In order for the scheme to work, the data must be skewed in time or order from one antenna to the next. When the data emerges from each port, at the right, the order of the antennas will be the same but they will be skewed in time or order. This data skewing should not, in general, present a problem. It takes eight cell times, $t=0-7$, for the data to complete one turn. The switch positions for each cell time are indicated in the following table. The numbers across the top of the table represent the cell times. The letters to the left represent the antenna row. The pattern of ones and zeros represents the switch state for each of the switch locations.

	0	1	2	3	4	5	6	7
a	0 0 0	1 1 1	1 1 0	1 0 1	1 0 0	0 1 1	0 1 0	0 0 1
b	0 0 1	0 0 0	1 1 1	1 1 0	1 0 1	1 0 0	0 1 1	0 1 0
c	0 1 0	0 0 1	0 0 0	1 1 1	1 1 0	1 0 1	1 0 0	0 1 1
d	0 1 1	0 1 0	0 0 1	0 0 0	1 1 1	1 1 0	1 0 1	1 0 0
e	1 0 0	0 1 1	0 1 0	0 0 1	0 0 0	1 1 1	1 1 0	1 0 1
f	1 0 1	1 0 0	0 1 1	0 1 0	0 0 1	0 0 0	1 1 1	1 1 0
g	1 1 0	1 0 1	1 0 0	0 1 1	0 1 0	0 0 1	0 0 0	1 1 1
h	1 1 1	1 1 0	1 0 1	1 0 0	0 1 1	0 1 0	0 0 1	0 0 0

Like the FFT, the above switch technique may be used with any radix. A radix three switch stage would consist of three position switches each position labeled from top to bottom 0, 1 and 2. One can also construct a mixed radix corner turner.

The Correlator

After the data leaves the corner turner it passes to the correlator back end. Each correlator processes one slice of the frequency information from each of the



antennas. The data enters the correlator as a stream of 350 antenna data points. It is necessary to multiply all combinations of antenna pairs together and accumulate the results for each pair. The figure illustrates a method for doing this for 8 antennas. The configuration looks very much like a lag correlator except that the samples are in antenna sequence rather than time sequence. The antennas

may be visualized as being in a circle and all spacings must be measured. Each of the squares represents a stage of shift register that is large enough to hold one frequency slice. (Possibly an 8 bit complex number with 4 bits real and 4 bits imaginary) The first conjugate multiplier is set up to multiply the frequency sample from each antenna as it comes in by itself. The second multiplies samples from adjacent antennas. The third, multiplies antennas that are spaced one antenna apart and so on. Each multiplier sends its result to a drum accumulator that advances to the next number as each product is generated.

At the start of the integration, the shift register and drum memories hold all zeros. As the shift register fills, the multipliers proceed until all eight stages are full. On the ninth clock pulse the multipliers all switch to multiplying the last stage of the shift register except for the first and last multiplier. The shift register loads the next stream of 8 data points while the first stream is finishing. The multiplier switches return to their normal positions as the second stream enters the shift register. The second multiplier switch returns at the 10th pulse, the third at the 11th and so on. At the end of the integration, zeros are loaded into the shift register so that the last stream may finish its contribution to the result. The multiply/accumulates continue until the last data point leaves the shift register. If continuous integration is required, it is only necessary to read out the drum accumulators in the proper order in a zero after read operation.

This architecture minimizes the number of interconnections between chips. If there isn't sufficient space in one FPGA, the circuit can be easily extended by cutting it at the dotted line as indicated on the diagram.

The Use of Field Programmable Gate Arrays

The traditional approach to aperture synthesis array design is to develop a custom VLSI integrated circuit, which is then produced in large quantity. In the five-year plus design cycle for designing such a system, the integrated circuit technology becomes obsolete. At the time the array comes on line, faster, denser, less expensive chips become available. We propose an investment in hardware description software. By using programmable logic the software can be compiled and run in the most recent

technology avoiding time consuming foundry runs. Software becomes a capital cost while hardware becomes an operating cost of our instrument.

Using programmable logic, configuration changes can be made more efficiently. A custom chip must have all possible configurations designed into it. A field programmable gate array can be optimized for one configuration. New configurations consist of completely new programs with new optimizations.

OBSERVING MODES

Bandwidth	Channels	Polarization products	Bits	Polyphase
400 MHz	256	2	1	No
200 MHz	512	2	2	yes
100 MHz	1024	2	4	yes
200 MHz	256	4	1	Yes
100 MHz	512	4	2	Yes
50 MHz	512	4	4	yes

The above sample of observing modes represents just a small sample of the possibilities. At the input to the polyphase filterbank there will be four independently tuned 100 MHz passbands. By changing the programming of the FPGA we can choose to use all or just one of those bands producing 400, 200 or 100 MHz bandwidths. The smaller bandwidth of 50 MHz can be introduced by preceding the polyphase filterbank within the same chip with a finite impulse response half band decimation filter. The logic following the decimation filter continues to operate at 100 MHz on data that is now sampled at a 50 MHz rate so that twice as many operations may be performed on the samples with the same logic. Some of these modes will be stored in memory on the polyphase filterbank card so that the modes may be changed relatively quickly. A more extensive collection of modes will be available by reprogramming the memories from the control computer.

Traditionally, when a system comes online, improvements or modifications are very difficult because the system must come down interrupting observing until the new hardware is bolted into place. Field programmable logic lends itself to ongoing system development. Indeed, the system can come on line before all of the software is

developed. A system upgrade takes no longer than reloading the chips with new code.

Field programmable gate arrays are off the shelf items. They have withstood the tests of many consumers. Any hidden flaws are likely to have been found and eliminated. A custom chip, while less expensive, has not been tested as extensively. There could easily be a problem that doesn't show up until much later causing the system to be unreliable.

¹ Escoffier, R. "The MMA Correlator" NRAO Millimeter Array Memo Series No. 166

² Iwase, S.; Okita, T.; Yamazaki, T.; Daishido, T.; Asuma, K. "A high-speed parallel FFT-processor for a radio astronomy telescope." Electronics and Communications in Japan, Part 3 (Fundamental Electronic Science), vol.72, (no.5), May 1989. p.109-17.

³ Nakajima, J.; Otobe, E.; Nishibori, K.; Kobayashi, H.; Tanaka, N.; Saitoh, T.; Watanabe, N.; Aramaki, Y.; Hoshikawa, T.; Asuma, K.; Daishido, T. "One-dimensional imaging with the Waseda FFT radio telescope." Publications of the Astronomical Society of Japan, vol. 45, (no. 3), 1993. p. 477-85.

⁴ Chikada, Y.; Ishiguro, M.; Hirabayashi, H.; Morimoto, M.; Morita, K.-I.; Kanzawa, T.; Iwashita, H.; Nakazima, K.; Ishikawa, S.-I.; Takahashi, T.; Handa, K.; Kasuga, T.; Okumura, S.; Miyazawa, T.; Nakazuru, T.; Miura, K.; Nagasawa, S. "A 6*320-MHz 1024-channel FFT cross-spectrum analyzer for radio astronomy." Proceedings of the IEEE, vol.75, (no9), Sept.1987. p.1203-10.

⁵ Romney, J.D. "Theory of Correlation in VLBI" ASP Conference Series, Vol 82, {\it Very Long Baseline Interferometry and the VLBA}, J.A. Zensus et al., editors; 1995.

⁶ Urry, W.L. "The FFT as a Filter Bank" ATA memo 10

⁷ Bunton, J. "An Improved FX Correlator" ALMA Memo 342

⁸ Harris, F.J. "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform" Proceedings of the IEEE, vol. 66, No. 1, Jan. 1978

⁹ Welch P.D. "The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms," IEEE Trans. Audio and Electroacoust., vol. AU-15, pp. 70-73, June 1967

¹⁰ Fiore, P.D. "Low-Complexity Implementation of a Polyphase Filter Bank," Digital Signal Processing 8, 126-135 1998

¹¹ Corner turners work well for direct imaging as well. It is my understanding that the new Waseda 64 element array uses one.

¹² Urry, W.L. "A Corner Turner Architecture" ATA Memo #14