

NATIONAL RADIO ASTRONOMY OBSERVATORY
Charlottesville, Virginia 22903

VLBA Correlator Hardware Simulator

by Ray Escoffier, Joe Greenberg, and Chuck Broadward

April 1, 1987

I) Introduction

This memo will describe the "hardware simulator" that is being used in Charlottesville to study the VLBA FX correlator design concept. This simulation project was undertaken to parallel and complement the totally software simulator written by John Benson and described by John in VLBA correlator memo 74 (since special hardware was built, the term "hardware simulator" is used here to describe this simulator and to distinguish it from John's "software simulator" in spite of the fact that all of the simulation done with it is done in software). The concept behind this two-simulator approach was that the more versatile software simulator could be used to study many aspects of the FX concept but was limited, by execution time, from doing integrations deeper than about 100 Msamples, while the hardware simulator was not versatile at all but could be used to do much deeper integrations (to the 1 Gsample range).

The hardware simulator was designed using a Texas Instruments digital signal processing microprocessor programmed to duplicate the intended operation of the proposed gate array butterfly chip implementation of the FX correlator. This microprocessor approach has a number of advantages (and, of course, many disadvantages) over a high level language computer simulation, the three most obvious advantages being;

- 1) A primitive machine programming language that allows both exact bit by bit duplication of the gate array operation and run time optimization.
- 2) The use of the DPS microprocessor instruction set, which is optimized to do precisely the types of operations needed for the FX simulation.
- 3) A dedicated system that can run for weeks if need be, to perform deep integrations.

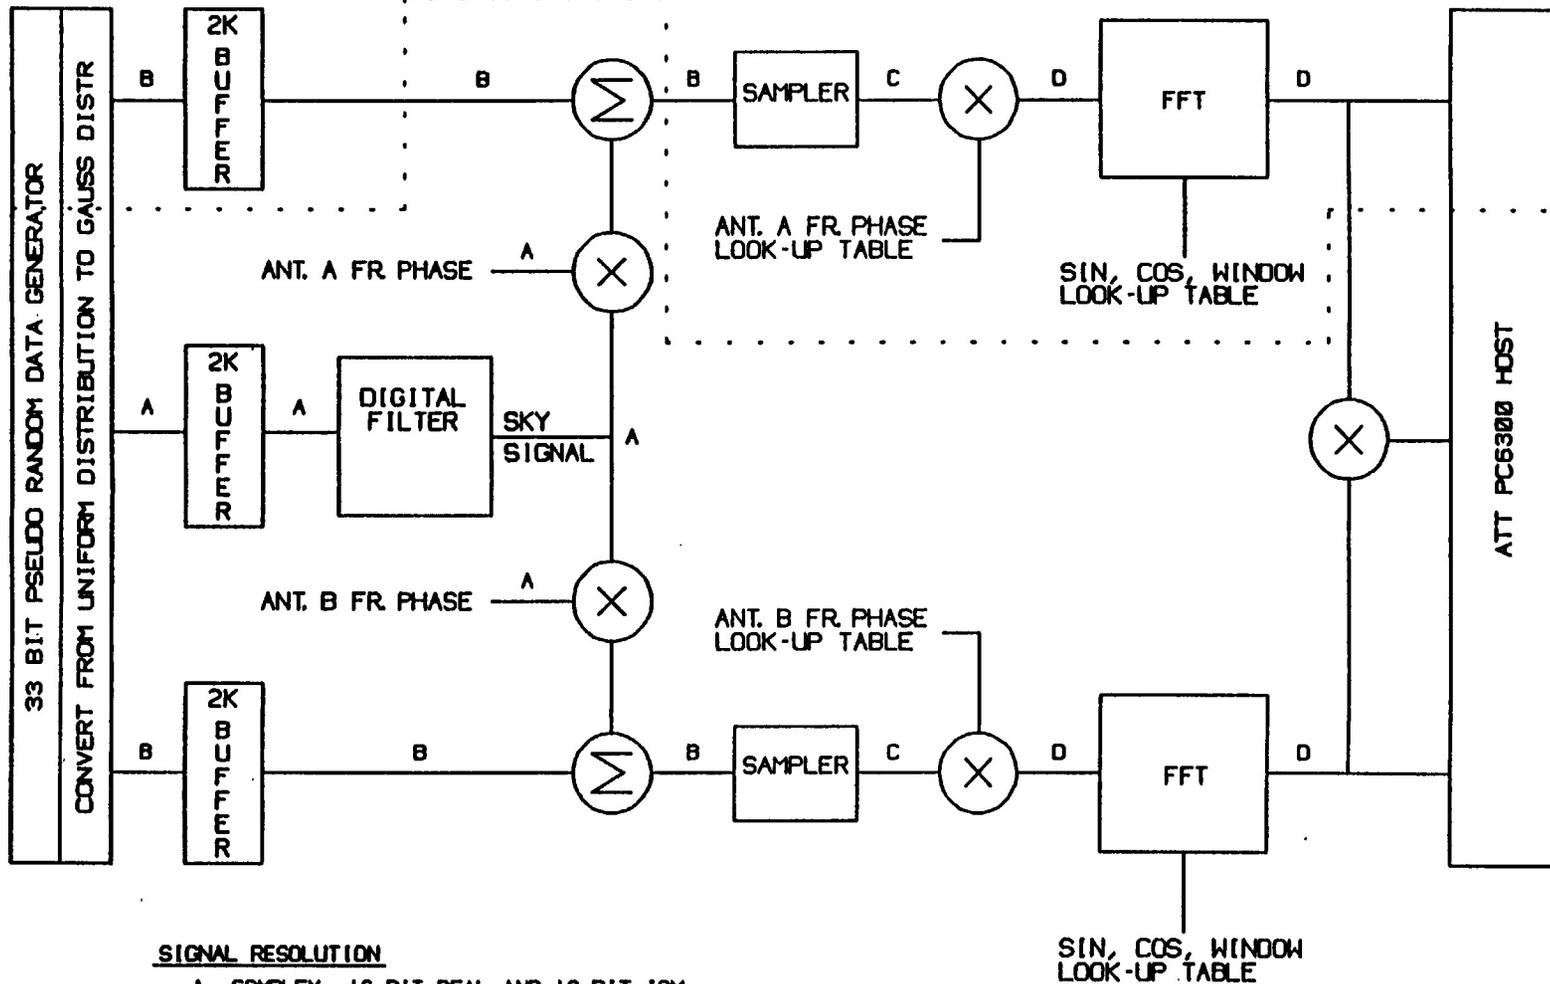
II) Hardware Description

The simulation hardware consists of a small Texas Instruments TMS 32020 microprocessor board (designed by Joe Greenberg) with just enough ROM and RAM memory to perform its function and an 8-bit parallel port into a host AT&T PC 6300 personal computer. The TMS 32020 was programmed (also by Joe Greenberg) to do the functions described schematically in Figure 1.

The first block of Figure 1 shows the generation of pseudo-random data. A 33-bit feedback shift register pseudo-random data generator was programmed to be the source of blocks of random data that would then be processed as inputs for the FFT simulation program. The data generator is used to produce 16-bit random

TMS32020 FX SIMULATOR BLOCK DIAGRAM

STATION ONLY
SIMULATOR



SIGNAL RESOLUTION

- A COMPLEX, 16 BIT REAL AND 16 BIT IGM.
- B 16 BIT REAL
- C 2 BIT REAL
- D COMPLEX FLOATING POINT

SIN, COS, WINDOW
LOOK-UP TABLE

FIGURE 1

numbers. As each 16-bit word is read from the data generator, the shift register is shifted 16 bits so that the next 16-bit word will be totally independent. The data generator used shift register bits 13 and 33 for feedback and will go through 2 to the 33 minus 1 shifts before starting to repeat.

The 16-bit data generator output words are uniformly distributed random variables and the data generator output is used to drive a look-up table conversion program that converts these uniformly distributed words (with only 10 of the 16 bits being used in the conversion process) to 16-bit normally distributed words. The program converts two uniformly distributed random variables, x and y, to one Gaussian distributed random variable, g, using the conversion equation (suggested by Jon Romney);

$$g = \sin(2 \pi x) * \text{sqrt}(-2 \log y)$$

The sin and sqrt(-2 log) functions are implemented by using 1024-entry look-up tables. This Gaussian conversion routine has the advantage of being both simple and fast while the multiplication of two look-up values provides for a range of one million possible output values instead of the 1024 obtainable from reading such small look up tables singly (of course, only 16 bits of the product will be kept and the effective output range is 65536).

In operation, three sets of 2048-point normally distributed data blocks are created. Two of these blocks have real components only and are to represent two receiver noises, while the third complex block is to represent sky signal. Originally, digital filters were to have been used to process the data blocks to simulate realistic data; however, the 16-bit digital filters, easily programmable by the TMS 32020, were not sufficient to produce acceptably flat band passes. Hence, when a flat pass band is required, the filter block is bypassed and Gaussian distributed blocks from the data generator-Gaussianizer are used to drive the FFT inputs directly (and therefore, as will be noticed in the results presented later, no filter side skirts are to be seen in the flat spectra). When sky signals that are not flat are desired, complex tap weights of a digital filter that will produce the desired sky signal spectrum are loaded into the sky signal FIR filter.

The sky signal may be independently fringe rotated by programmable fringe rates in the two digital mixers shown in Figure 1 to simulate realistic receive signals. A programmable correlation coefficient is used to produce antenna signals with the desired fractions of sky signal and antenna noise. The RMS levels of the resulting "antenna outputs" are measured by the host computer and a set of quantization thresholds are down loaded to the TMS microprocessor. Optimum thresholds of 0.9816 of the RMS antenna output voltage as per Fred Schwab's VLBA correlator memo 75 for four-level sampling are used. This memo also gives the optimum sample magnitude-bit weight as being 3.336. Since fringe de-rotation is done by a look-up table complex multiplication, any sample weight is as easy to implement as any other and hence the 3.336 factor was indeed used.

Once the antenna signals are quantized, the signal processing proceeds in a bit by bit simulation of what will happen in the FX correlator hardware. Correlator memos 71 and 72 described the proposed FX correlator FFT architecture and the blocks seen in Figure 1 after the sampler are implementations of this proposed architecture. The signal resolution levels at all points down the signal processing chain are programmable so the requirements and trade offs in signal quantization can be studied. A main reason for undertaking this simulation project was to be

able to establish whether the proposed architecture was sufficient for the FX correlator and to be able to do this studying of the process trade offs.

The simulator has a "station only" mode in which station FFT's are performed and the resulting station spectrum magnitude integrated. No cross multiplication is done in this mode. In station only mode the fringe de-rotator still runs at the rate set.

The host computer was programmed (by Chuck Broadwell) to support the TMS microprocessor in the simulation and to read the station and cross spectrum results and do long term integration. A variety of interactive displays were programmed to aid in the performing and evaluation of the test results.

III) Parameters Used in the Simulation Experiments Below

This memo will present the results of three integrations using the simulator described above. First, however, a brief description of the signal processing details, such as signal quantization levels, is in order.

Station fringe de-rotation phases were calculated to 8-bits of precision and a look-up table fringe de-rotator was used. Signals between FFT butterfly stages were in the 5,5,4 sign-magnitude complex floating point numbering system proposed for the FX correlator and described in VLBA correlator memo 71. The FFT butterfly additions were performed in 16-bit one's complement fixed point arithmetic. Conversion from fixed point to floating point was done at the FFT butterfly output with truncation (not rounding) occurring in the process.

The twiddle factors in the FFT trig table were expressed as 6,6,0 numbers. Preliminary results indicated that a constant 6,6,0 level of twiddle factor resolution was insufficient for the FFT requirements so in reality thirty-two 6,6,0 trig tables that average to an 11,11,0 quantized trig table were used in all of the experiments described below. After every 100 FFT's, a new trig table was down loaded to the TMS simulator from the host continuing through all of the 32 tables in the course of 3200 FFT's. This process of multiple low resolution trig tables that average to higher precision numbers is practical in the actual hardware because trig table generation need only be done once per system and a single large trig table memory is of no consequence in the final cost of the system.

The cross multiplication was done in 5,5,4 precision and long term integration was done by the host in 64-bit floating point.

No window function was used in any of the results presented below.

IV) Preliminary Results

In the figures that follow, plots that show integrated power spectra are given in two versions, the top plot in each figure covers exactly one decade of dynamic range (the two station line spectra are exceptions) and the lower plot shows the same spectrum expanded for better illustration of fine structure.

The first result is a station only integration designed to show the spectral performance of the FFT. The first two months of experimentation with the simulator quickly demonstrated that the station spectrum was the most easily polluted by systematic errors in the FFT calculation. Since most errors we have seen so far

occurring in the FFT process did not produce correlated effects, the cross spectra seem to be much less vulnerable to processing errors or offsets.

Figures 2 and 3 give the result of an 829 Msample single station run. Figure 2 gives the integrated spectral magnitude which is seen to be flat to an RMS error in the 0.18 % range. Figure 3 shows how the variance of the spectral points from the spectrum average varied with integration depth. At least to the depth of this integration, no residual structure (birdies or other features) has emerged from the spectrum and the variance of the points seems to follow a straight line on the log-log plot until deep into the integration. After the 829 Msample integration, the variance plot is 0.027 percent higher (0.184 % actual versus 0.157 % theoretical) than the one over sqrt (number of FFT's) theoretical variance curve.

The second test simulation is represented by the result seen in Figures 4 through 10. This run was a 303 Msamp. two station run with a flat sky spectrum and a correlation coefficient of about 0.2 and station fringe rates of 11,963 HZ (assuming a 32 MHz sample rate) for station 1 and 15,869 HZ for station 2. Figures 4 and 5 give the two station spectra and Figure 6 is the cross power spectrum. Figures 7, 8, and 9 give the variance of the points in the integrated spectra plotted against integration time. In all three of these figures, the variance was calculated by using points 300 through 700 only (a remnant from when there were filter side skirts). Figure 10 shows how the station and cross spectrum values converged with time.

Two curious things can be seen in this set of figures. First, there seems to be some difference in residual spectrum structure between station 1 and station 2 seen in the lower views of Figures 4 and 5. The station 2 spectrum seems to have more high frequency hash in it. We have no explanation for this as yet since both use identical software subroutines with only the fringe rates being different between the two stations. The second point to be noticed is seen in Figure 7 where the variance curve seems to be departing from a straight line after only about 40 Msamp. Since the single station run seemed to say that the FFT program was good to deeper integrations, this result would seem to indicate that there is some structure in the sampled data's spectrum. However, the source of the sky signal is exactly the same as that in the single station experiment (except for the fringe rotator which we removed in trying to find the source of the structure to no effect). We have no answer as yet to this discrepancy.

The results of the last test simulation are seen in Figures 11 through 17. This run was a 322 Msample two station run with a spectral line in the pass band. A correlation coefficient of 0.5 was specified. The line is seen to be fairly broad and this results from the implementation of a 256 tap FIR digital filter in the TMS 32020. The 256 tap filter was decided upon as a reasonable compromise between performance and execution time.

Again Figures 11 and 12 show some kind of systematic fine feature difference between the two stations' spectra when the only difference should be the fringe rates (which were again 11,963 HZ for station 1 and 15,869 HZ for station 2).

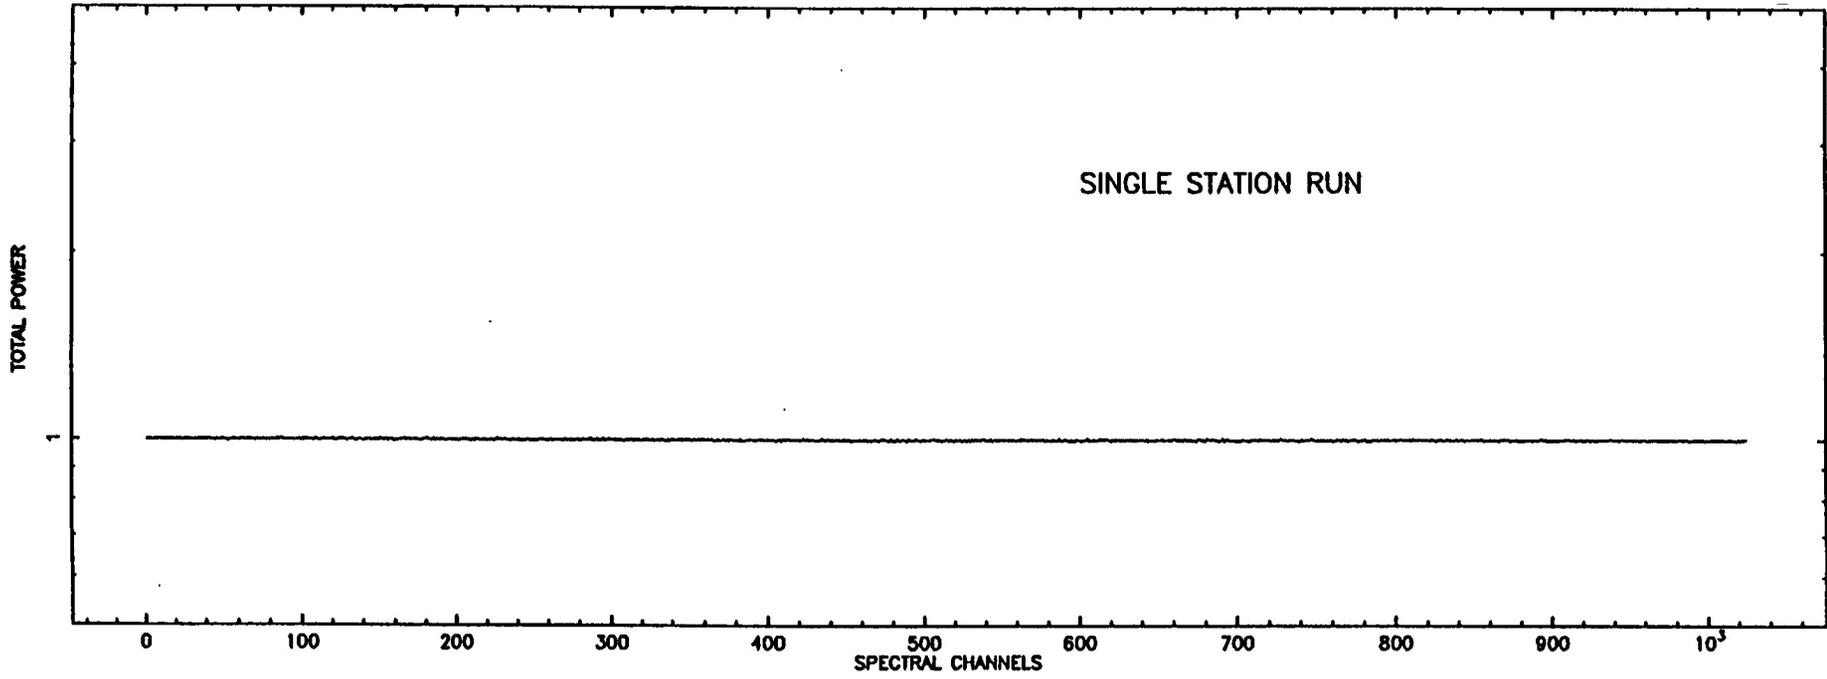
The most disturbing thing seen in Figures 11 and 12, however, are the features seen at spectral channels $1024 - 96$, 512 ± 96 , 256 ± 96 and 768 ± 96 (where the line is centered on spectral point 96). These features, along with the expected third harmonic are in the .05 to .1 percent range.

Otherwise, the cross spectrum shows three or four features down in the .02 percent range.

The variance curves are about as expected, the station variance curve "sees" the birdies fairly early in the integration while the two cross variances seem to hold a straight for the length of the experiment.

Figure 17 shows how the station and cross spectrum values converged with time for the line experiment.

829 meg samples 404,800 FFTs 3/13 to 3/18 + 3/30 to 3/31



829 meg samples 404,800 FFTs 3/13 to 3/18 + 3/30 to 3/31

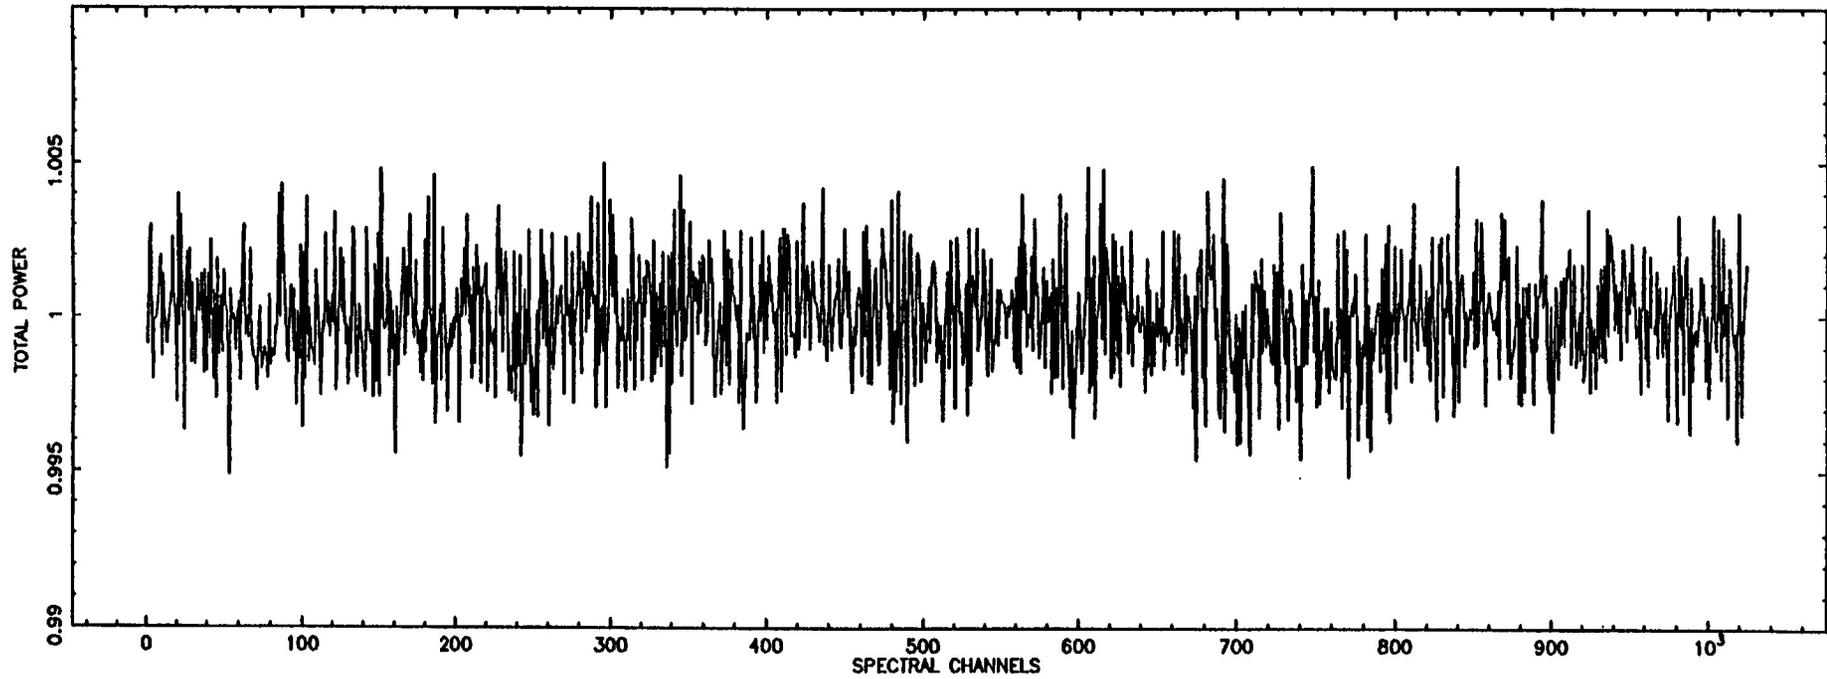


FIGURE 2

829 meg samples 404,800 FFTs 3/13 to 3/16 + 3/30 to 3/31

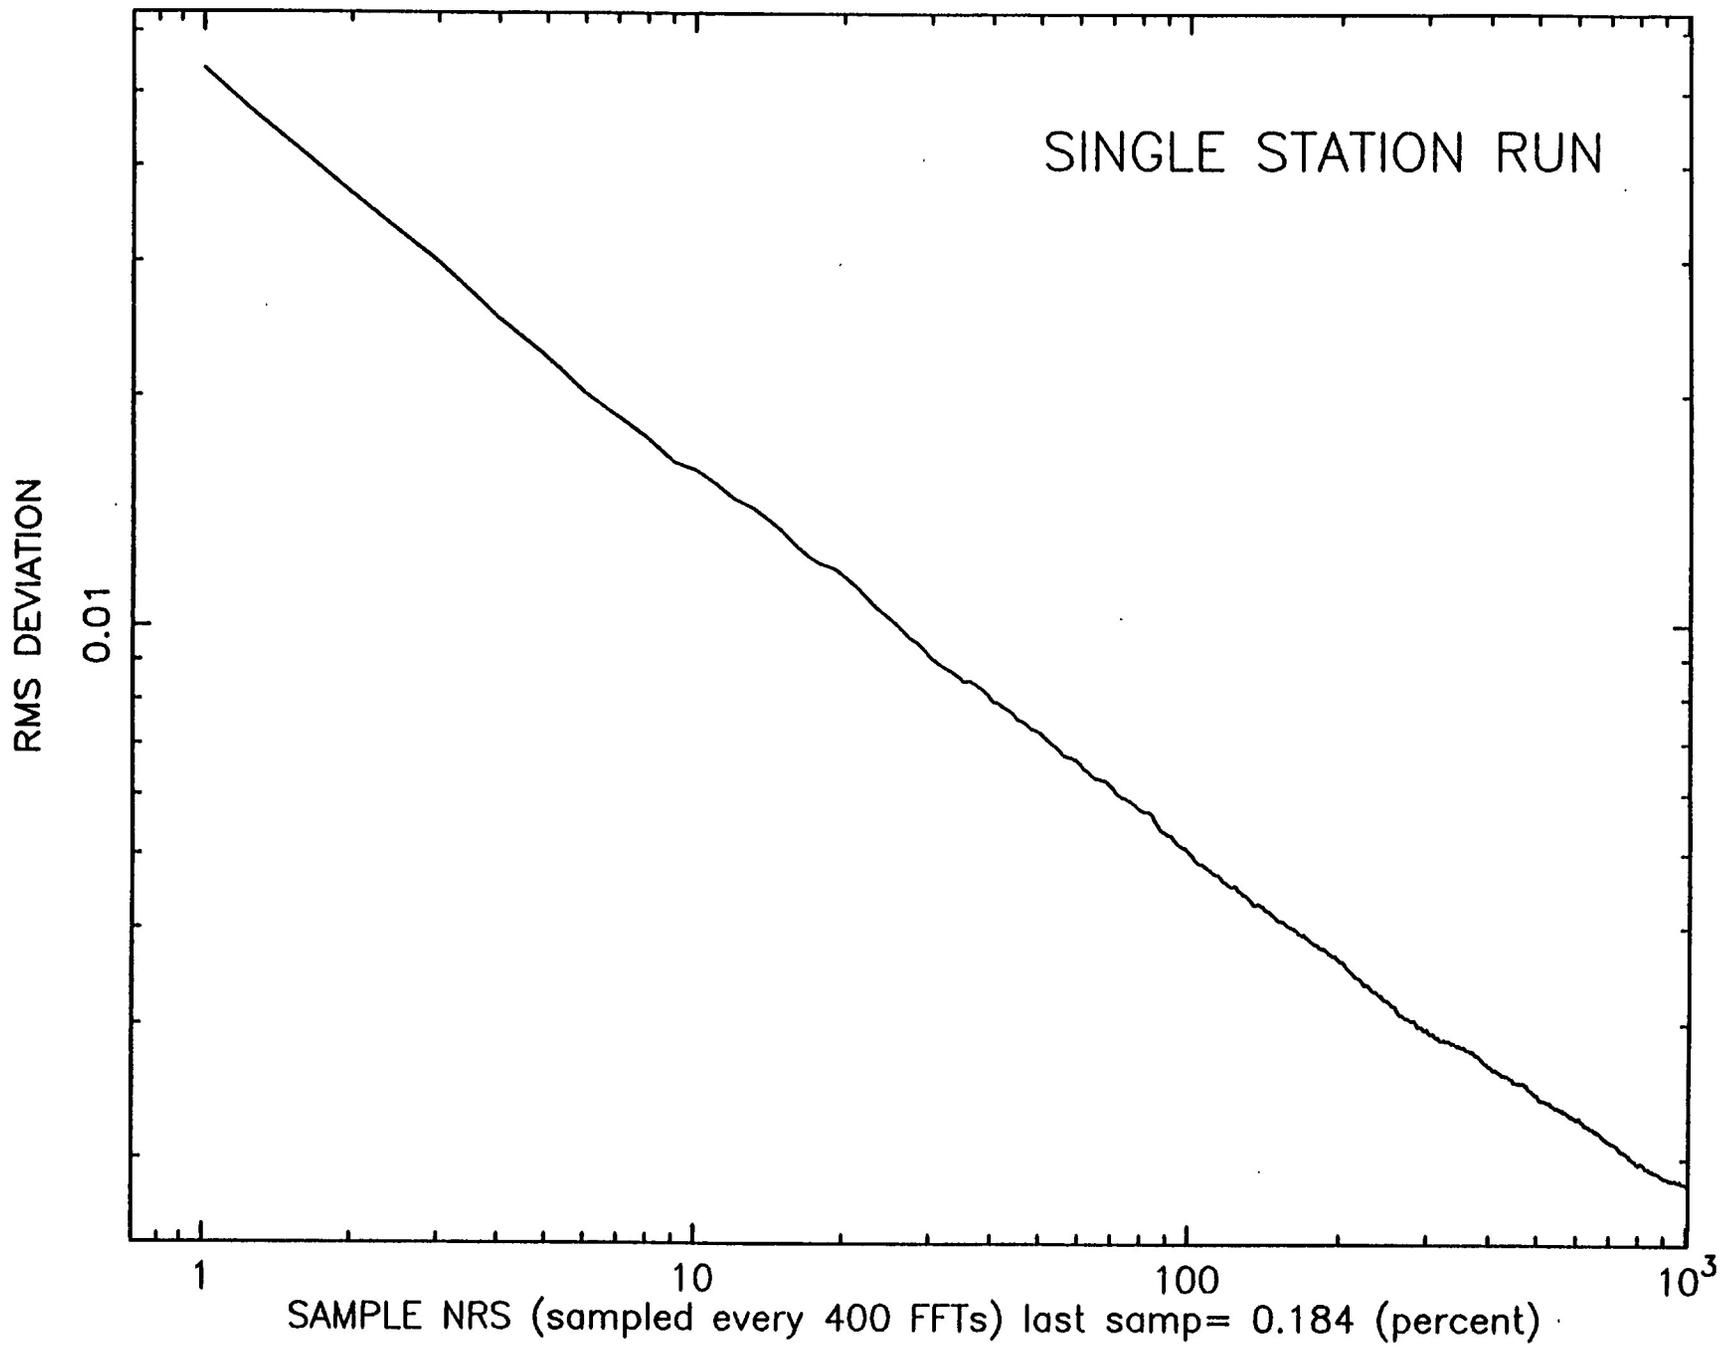
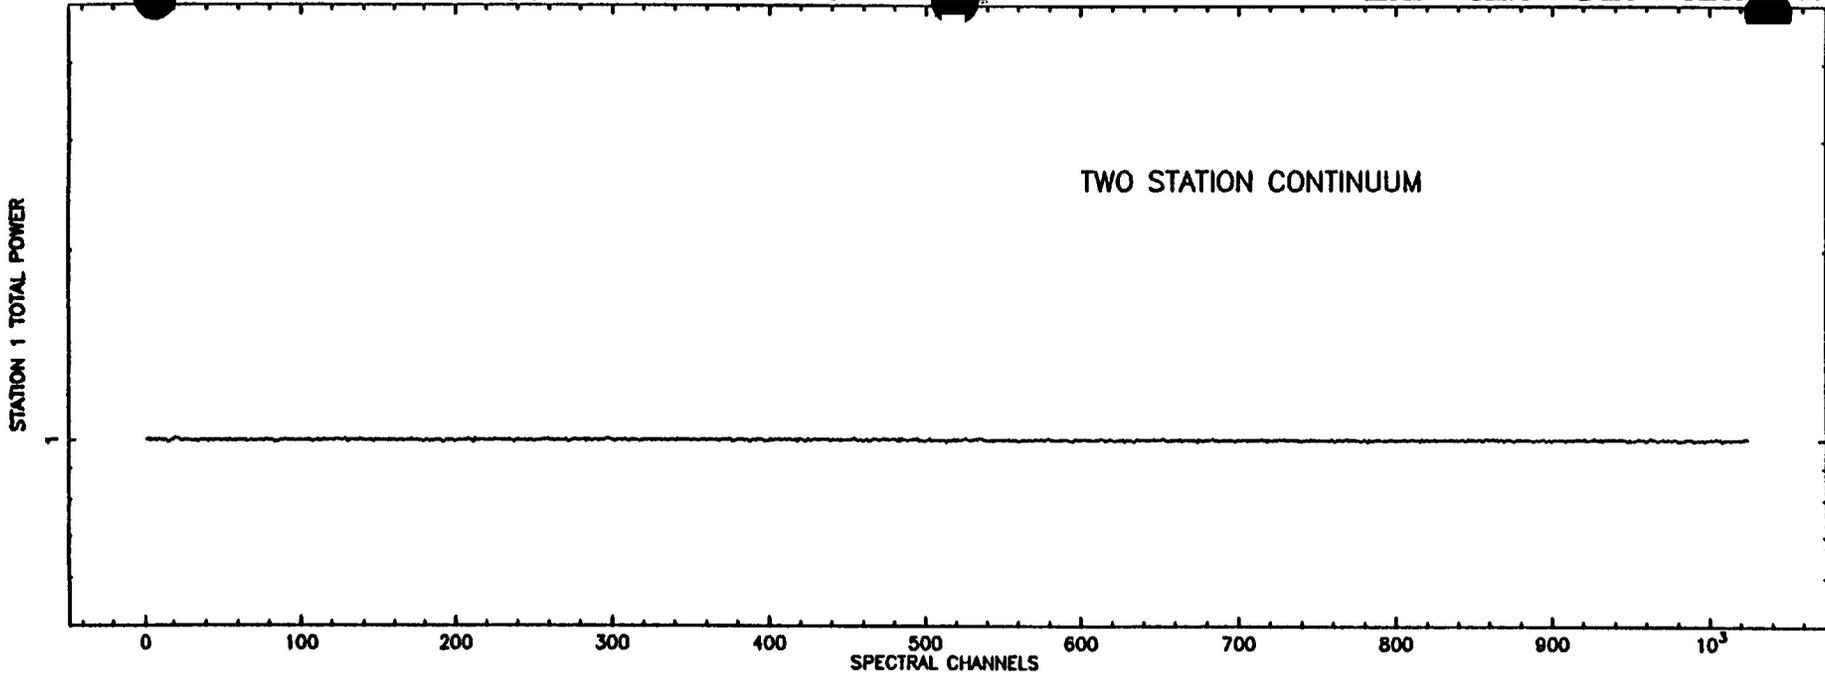


FIGURE 3

303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87



303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87

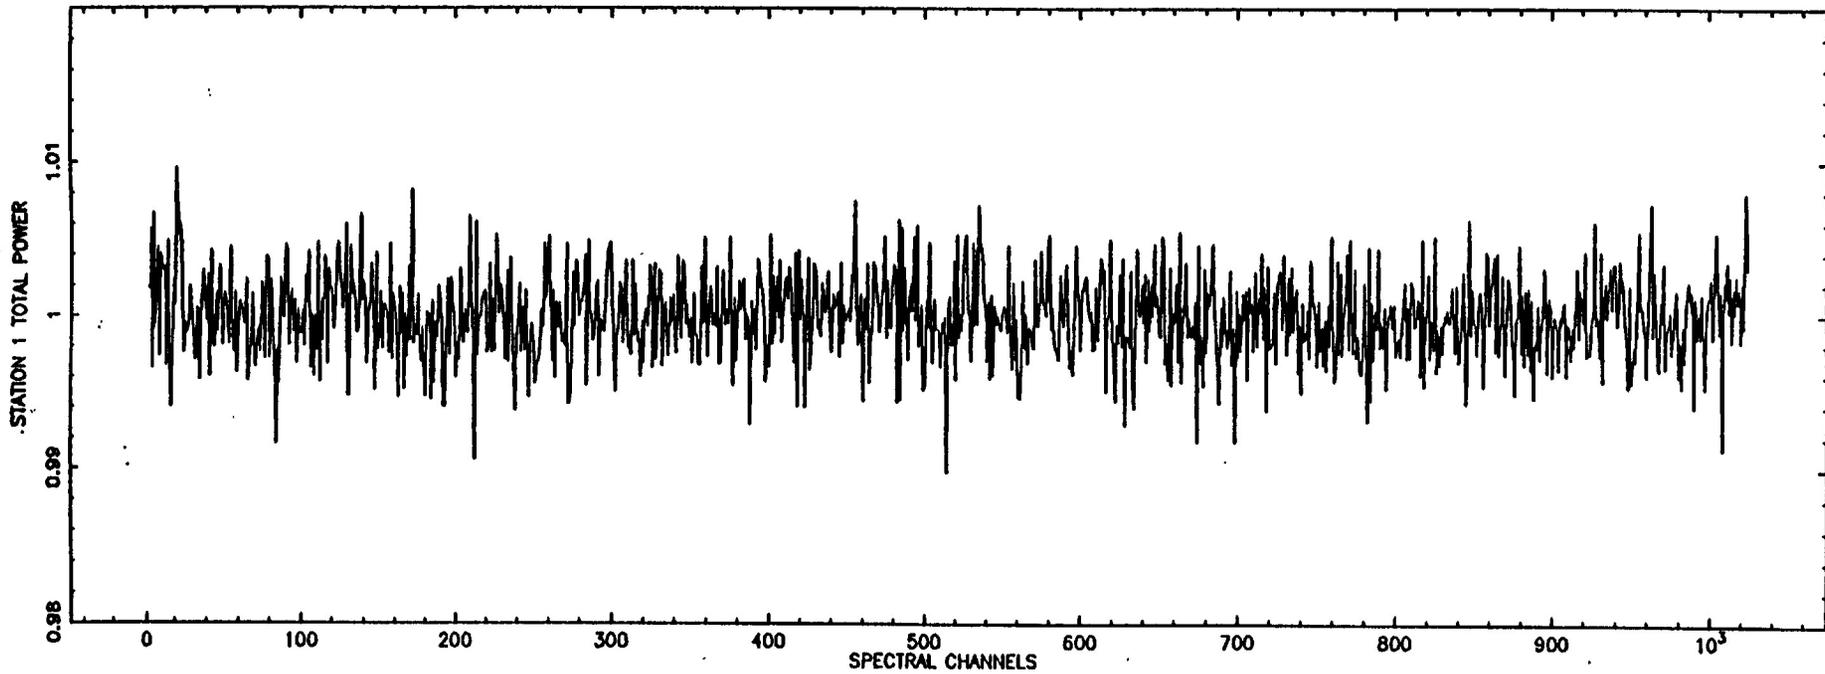
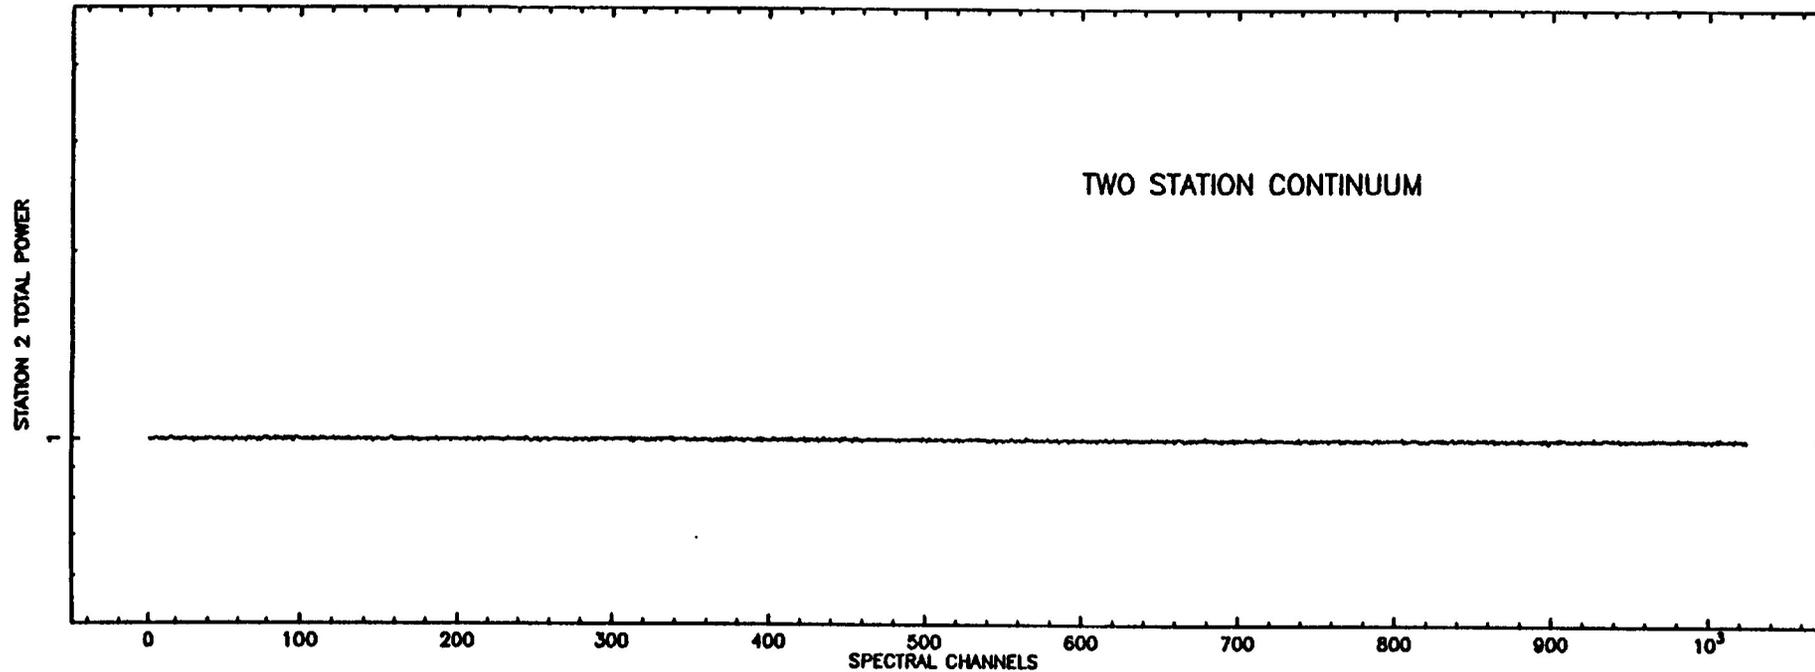


FIGURE 4

303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87



303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87

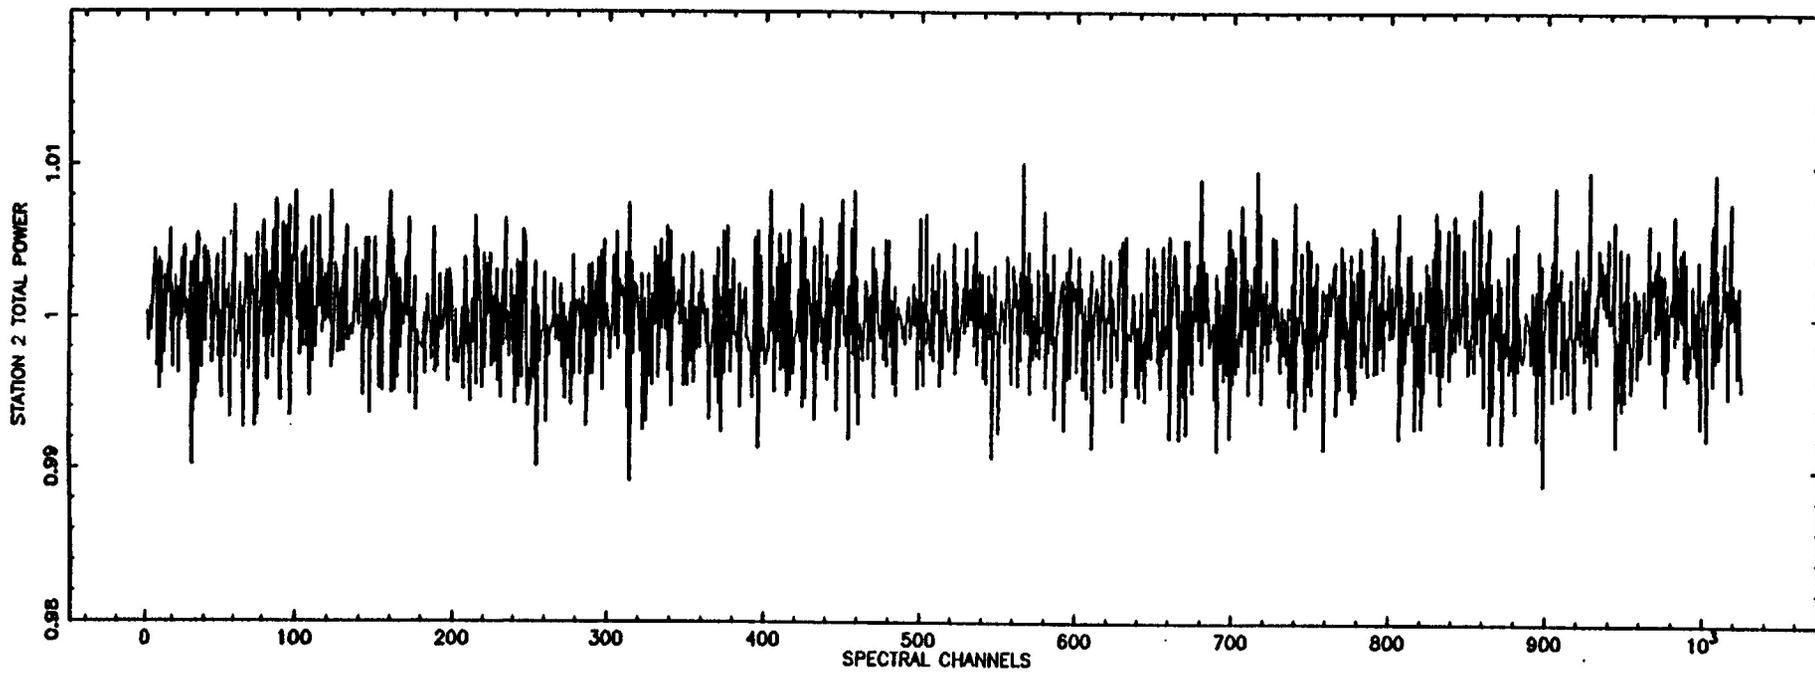
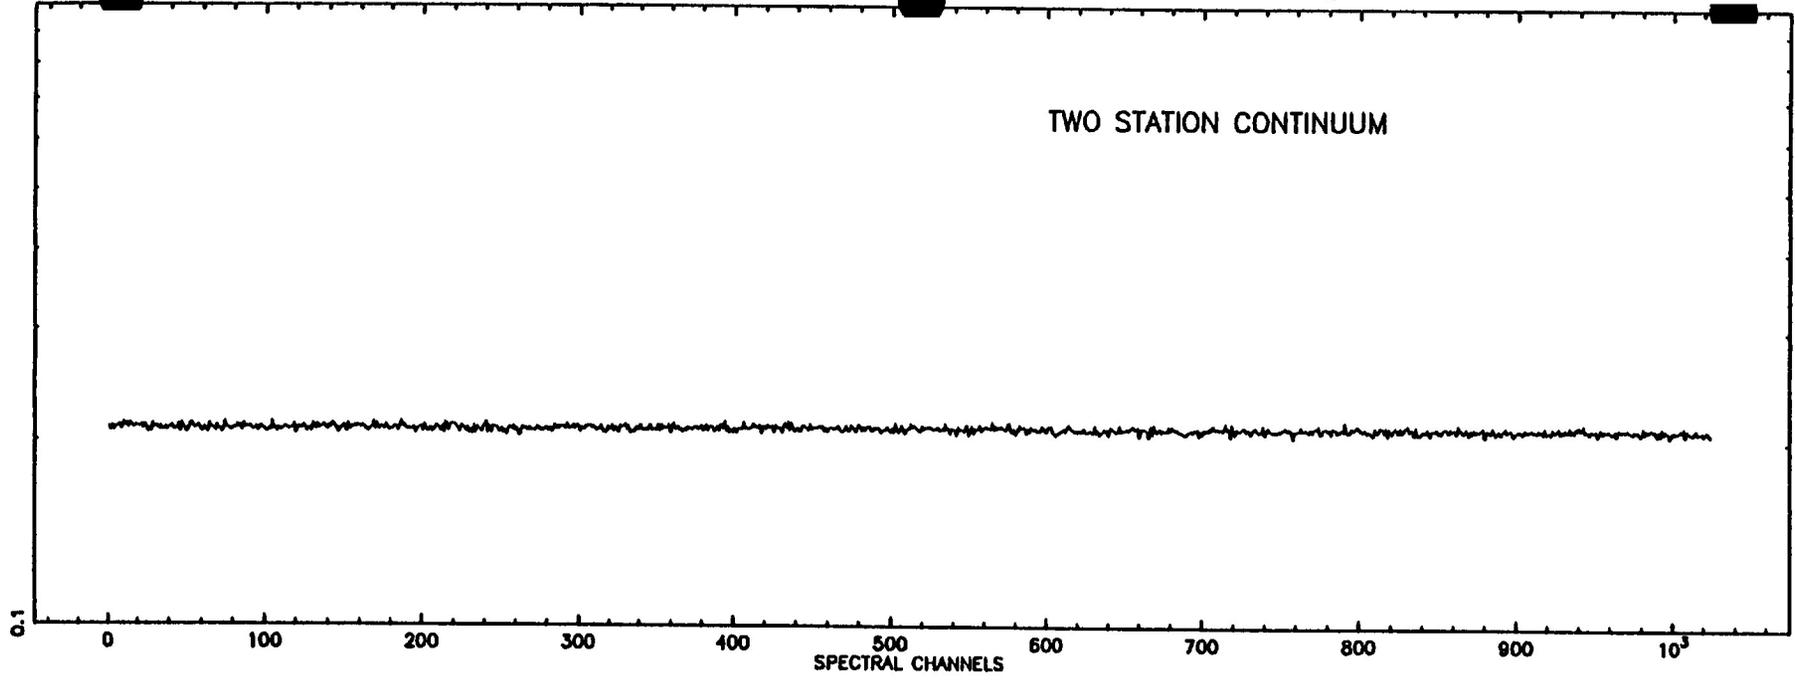


FIGURE 5

303.9 meg samples 148,400 FFTs CC=0.2 3-20-87 to 3-23-87

TWO STATION CONTINUUM

CROSS POWER



303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87

CROSS POWER

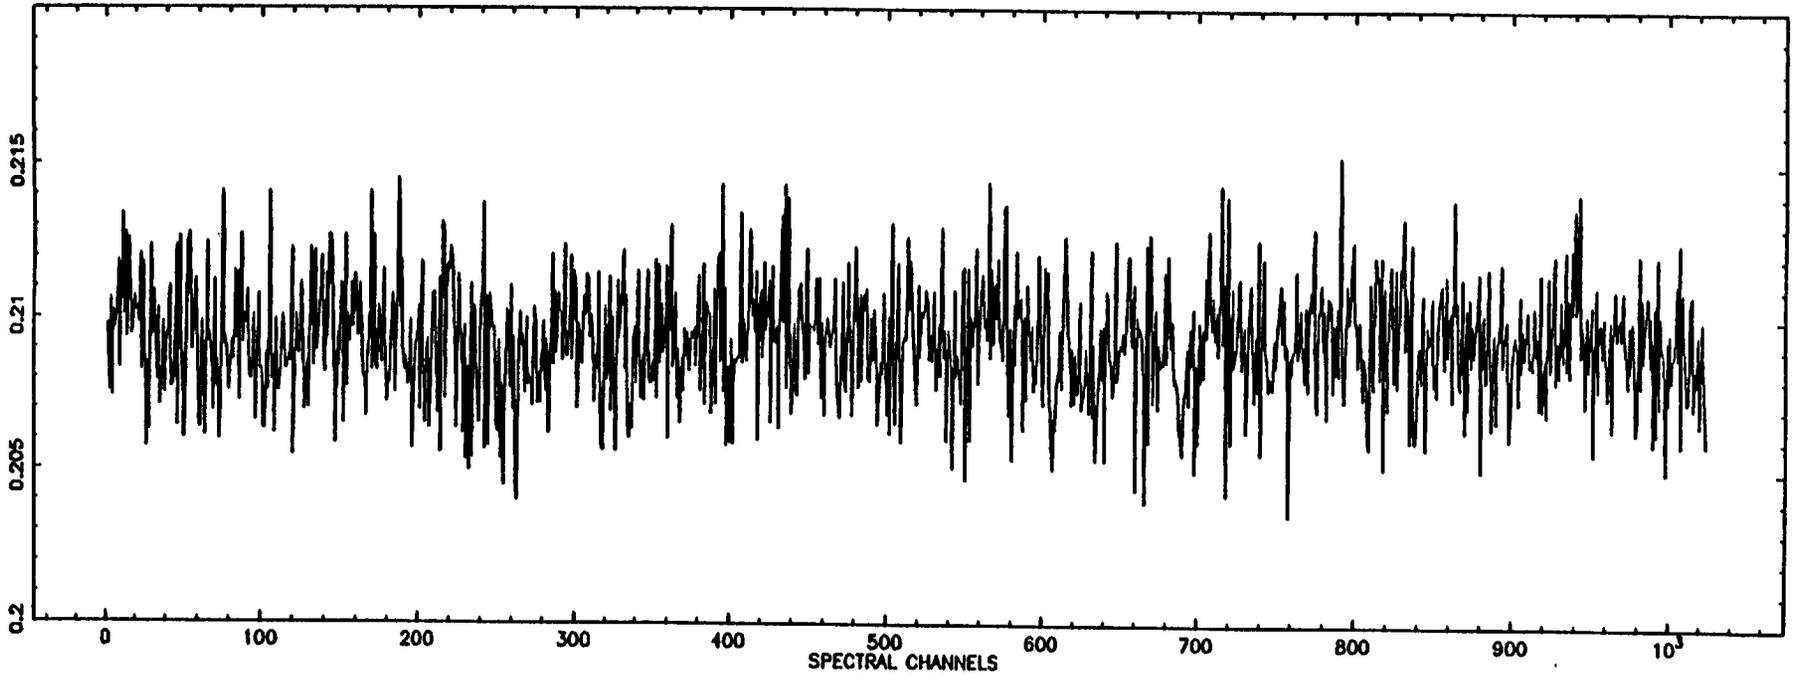


FIGURE 6

303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87

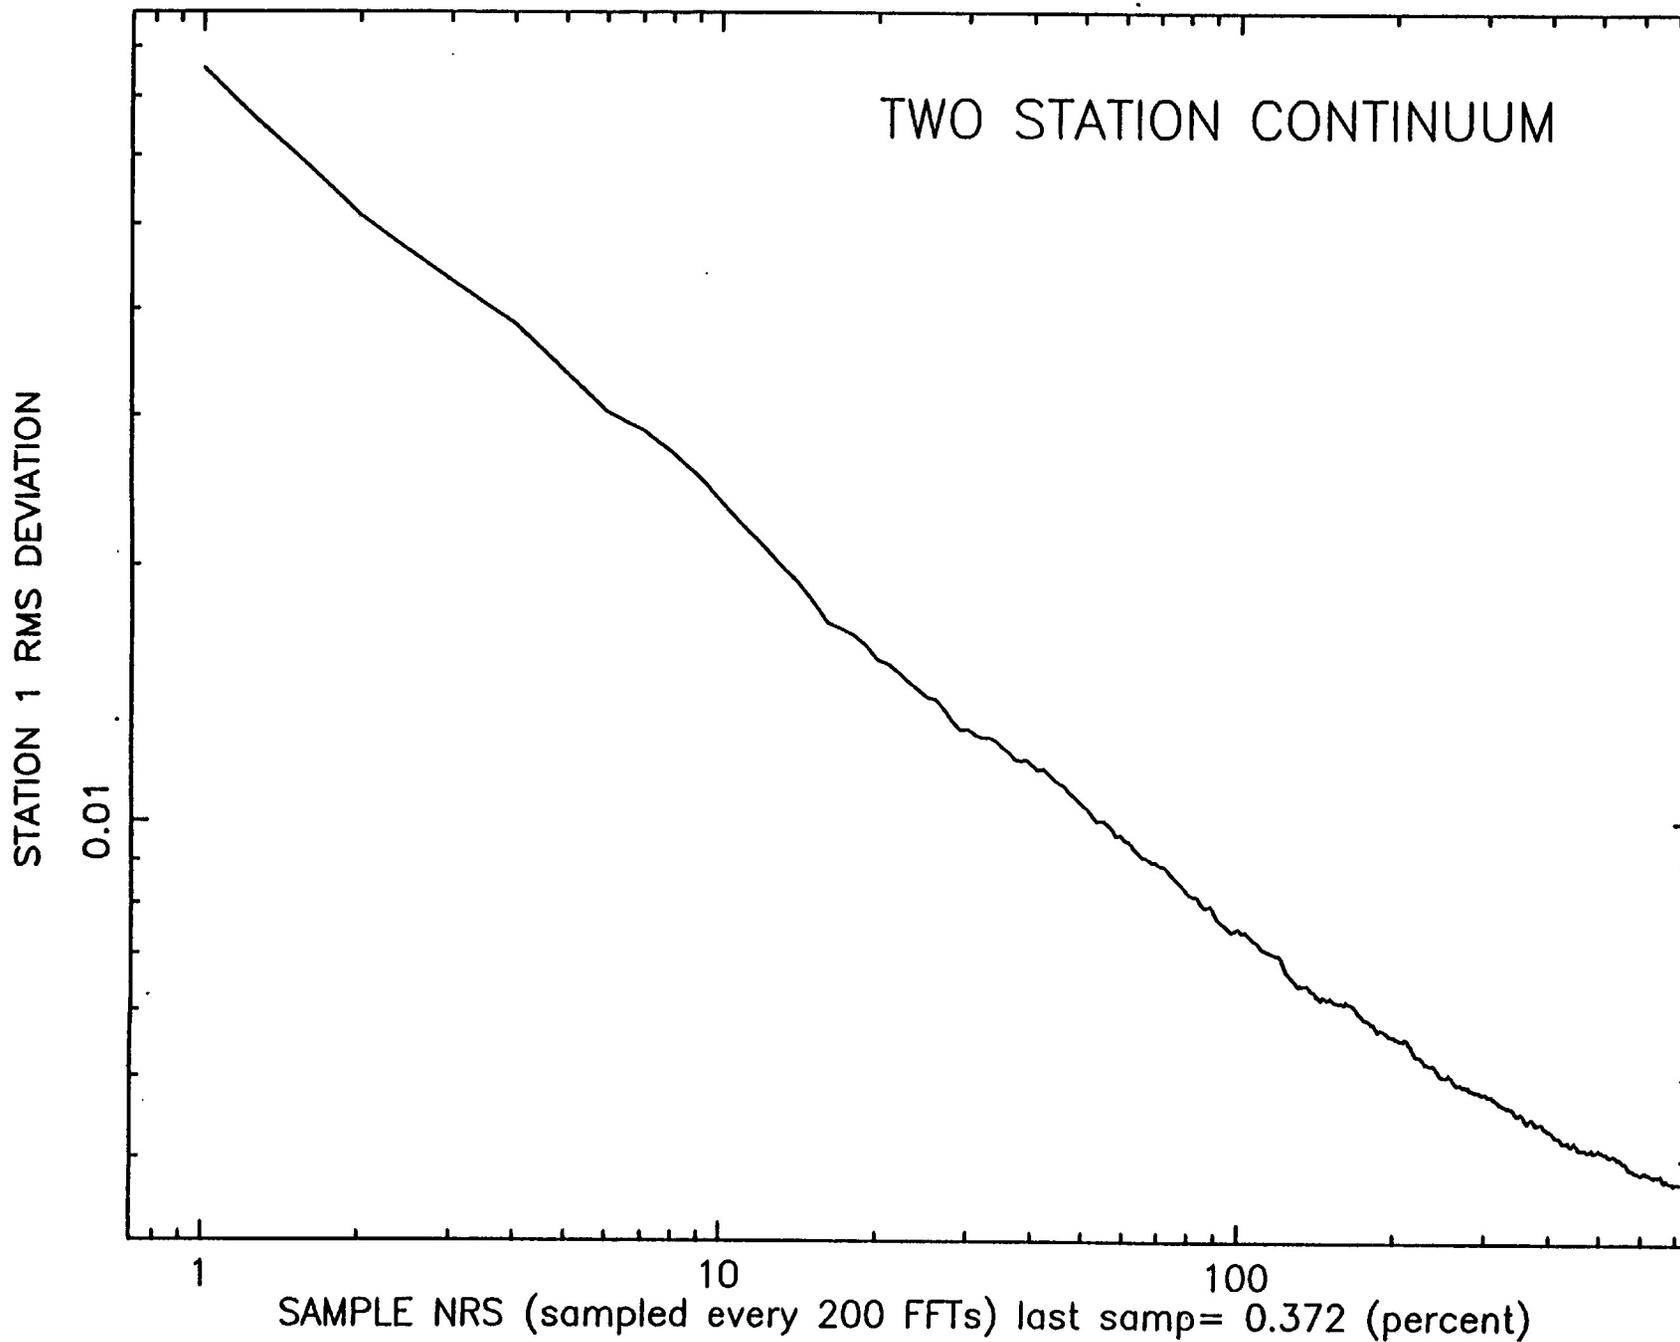


FIGURE 7

303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87

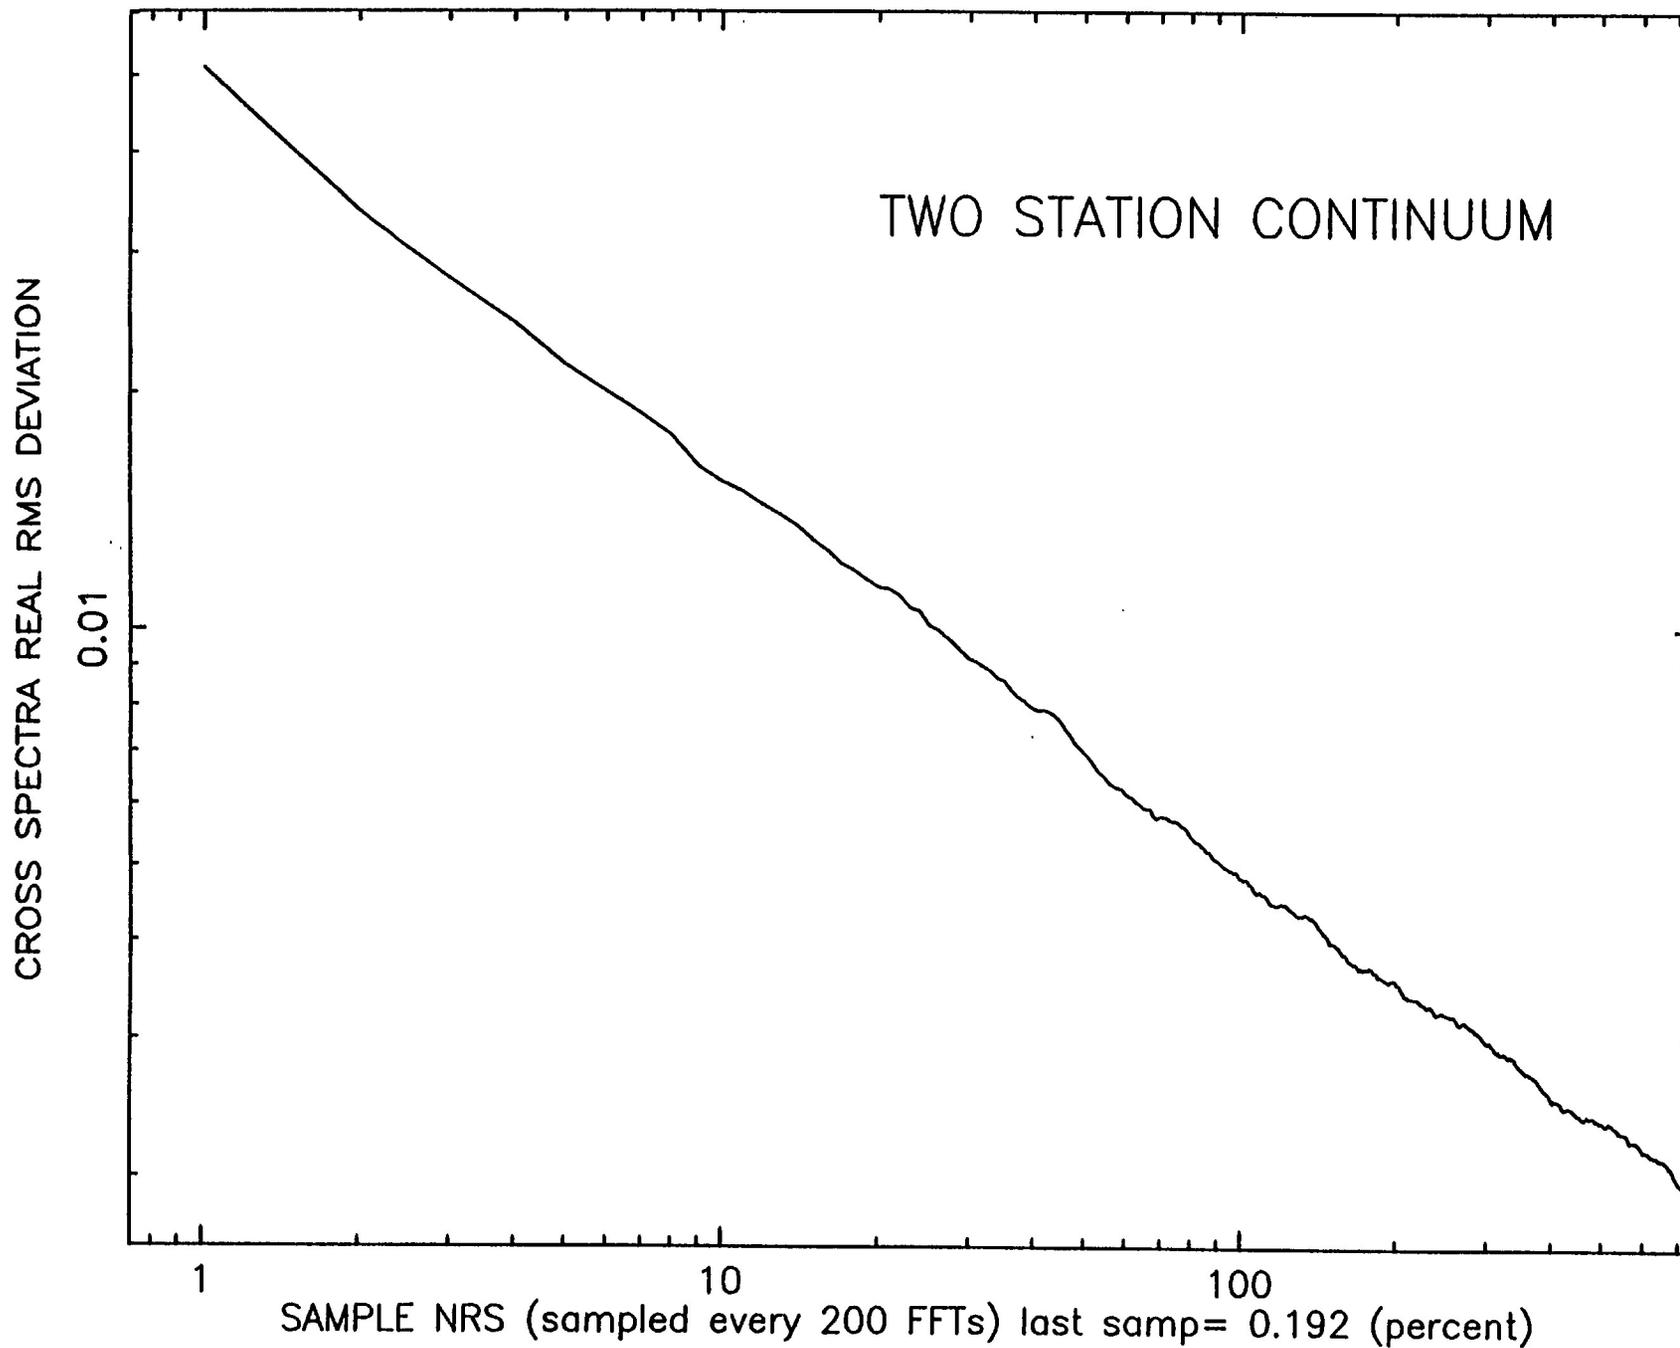


FIGURE 8

303.9 meg samples 148,400 FFTs CC= 0.2 3-20-87 to 3-23-87

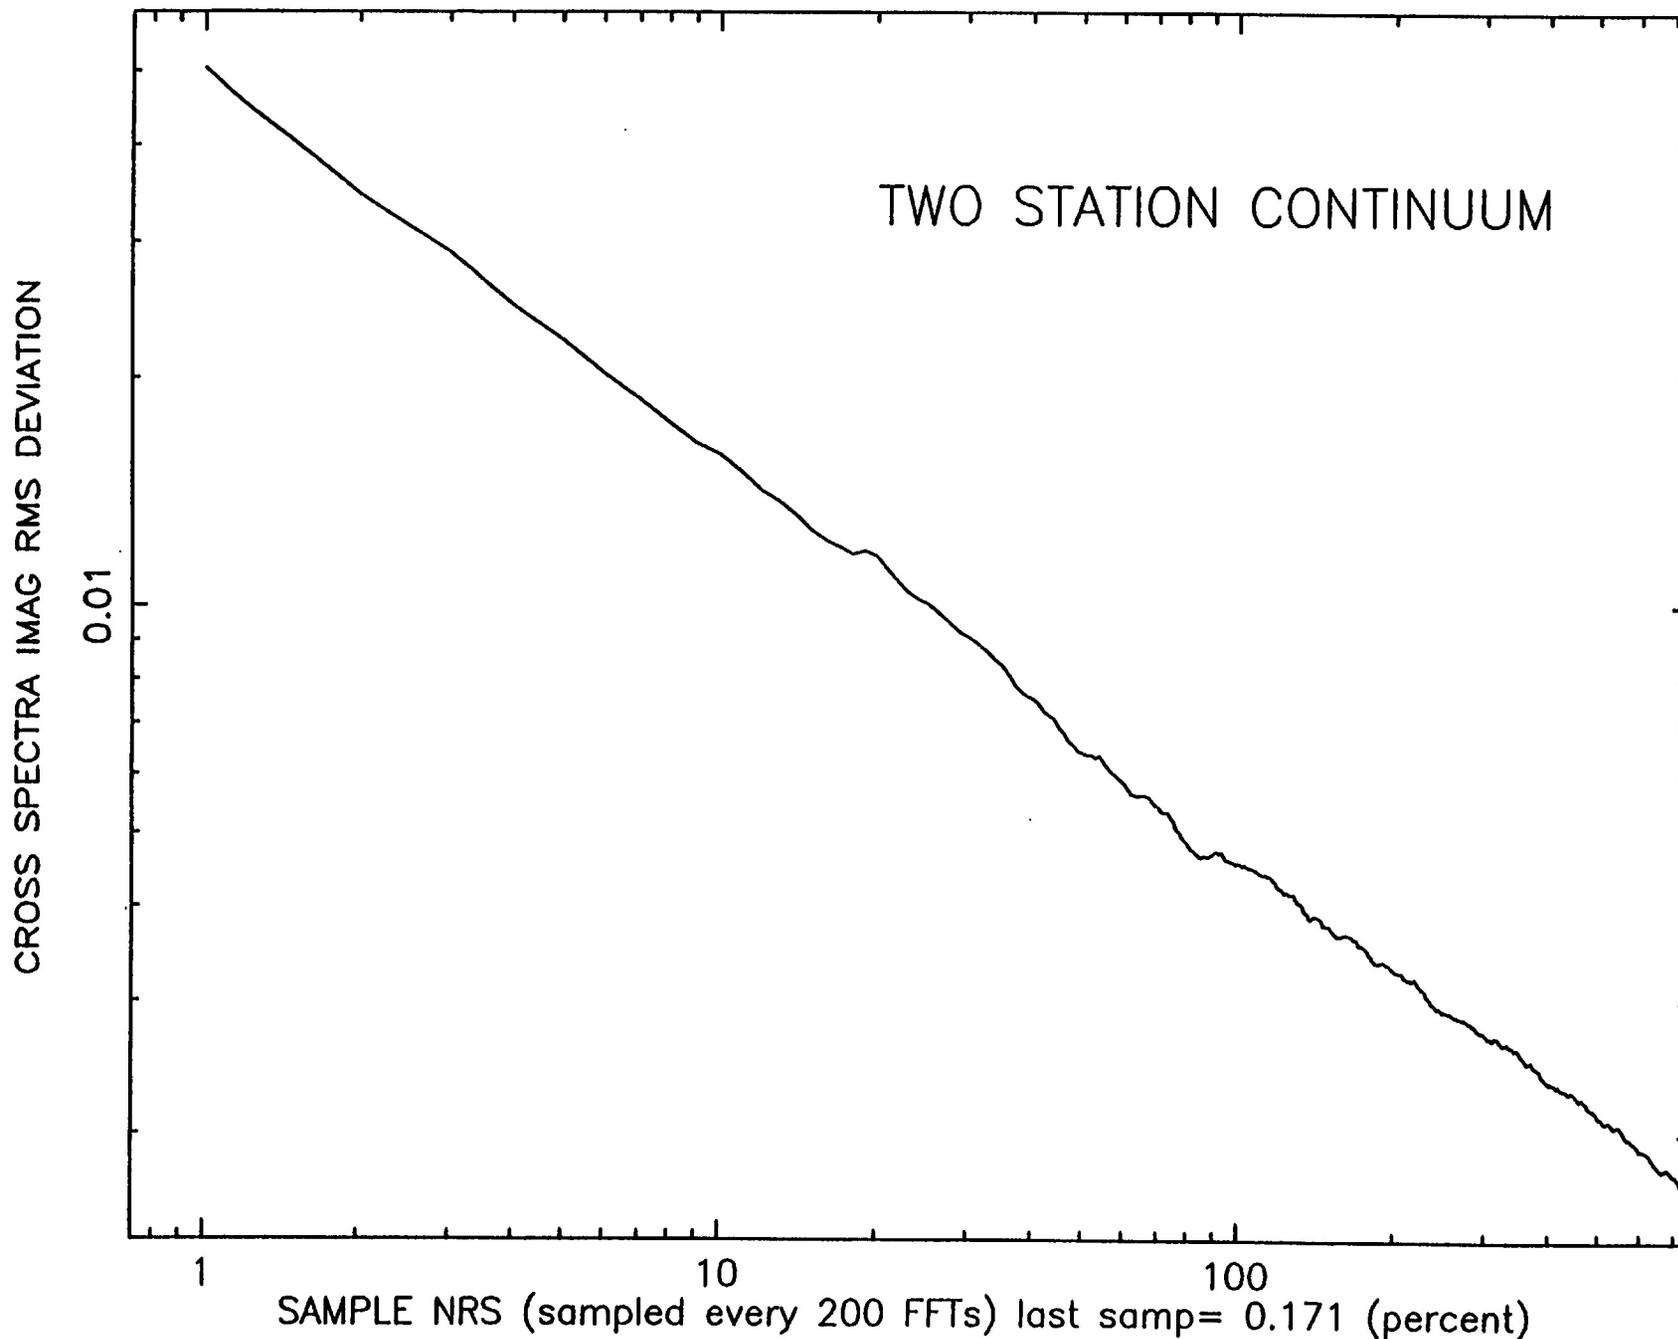


FIGURE 9

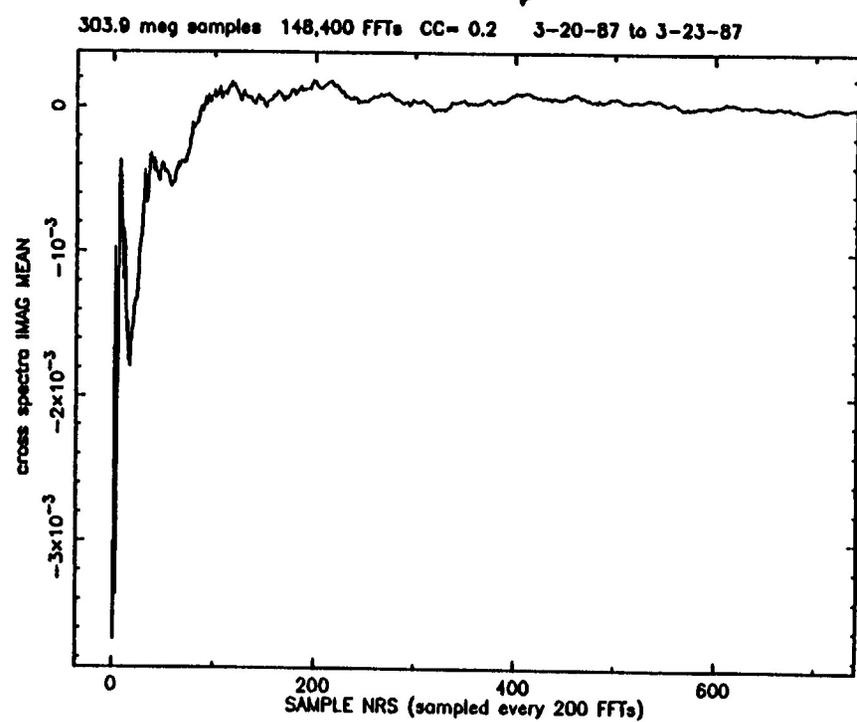
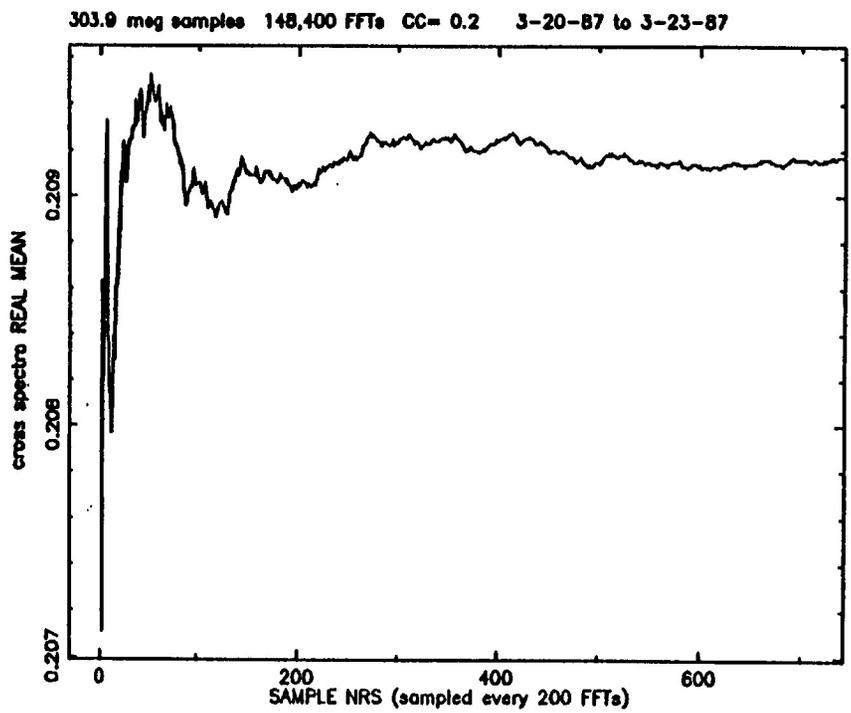
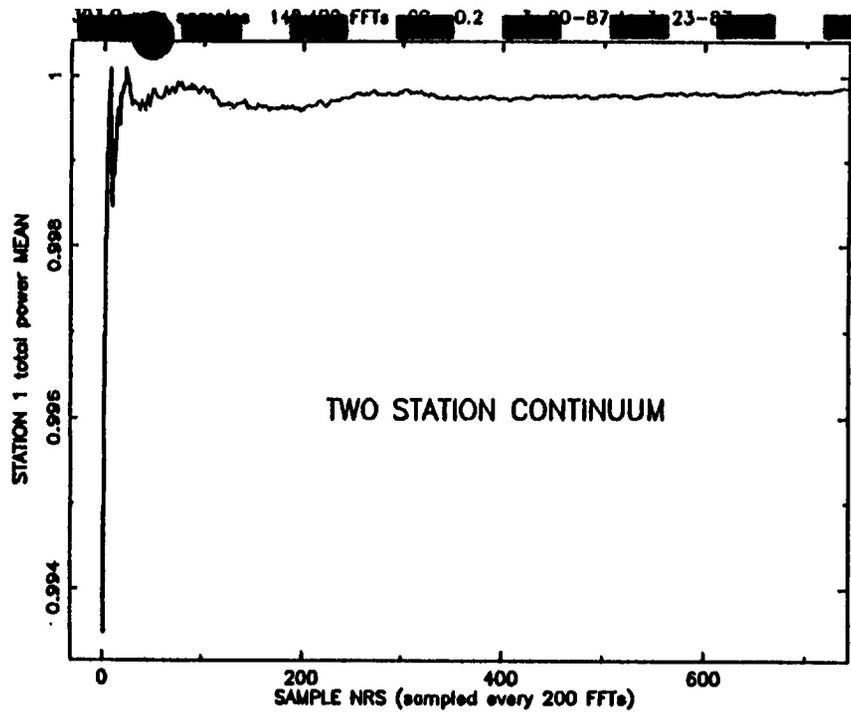
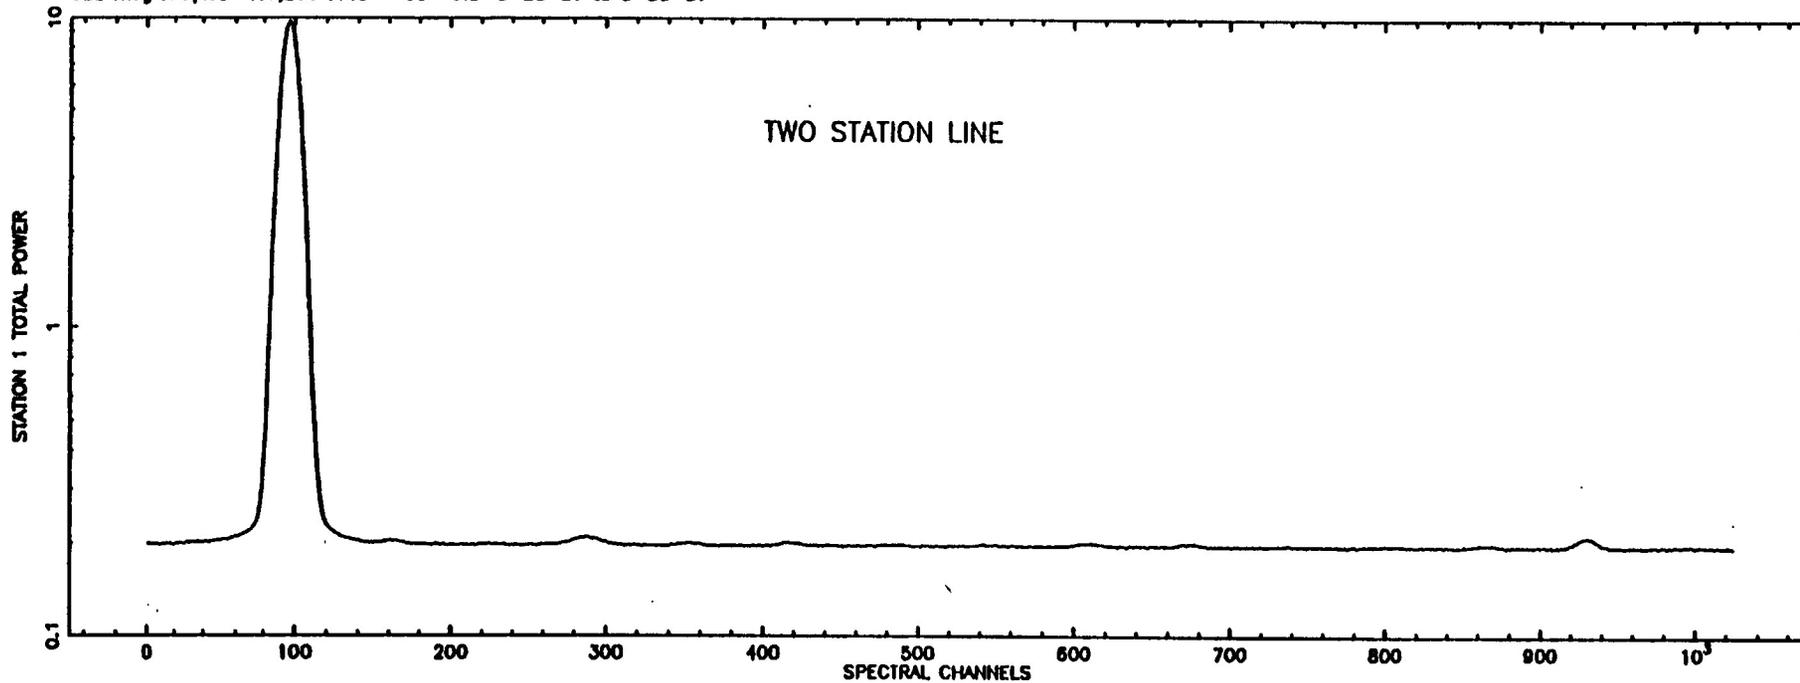
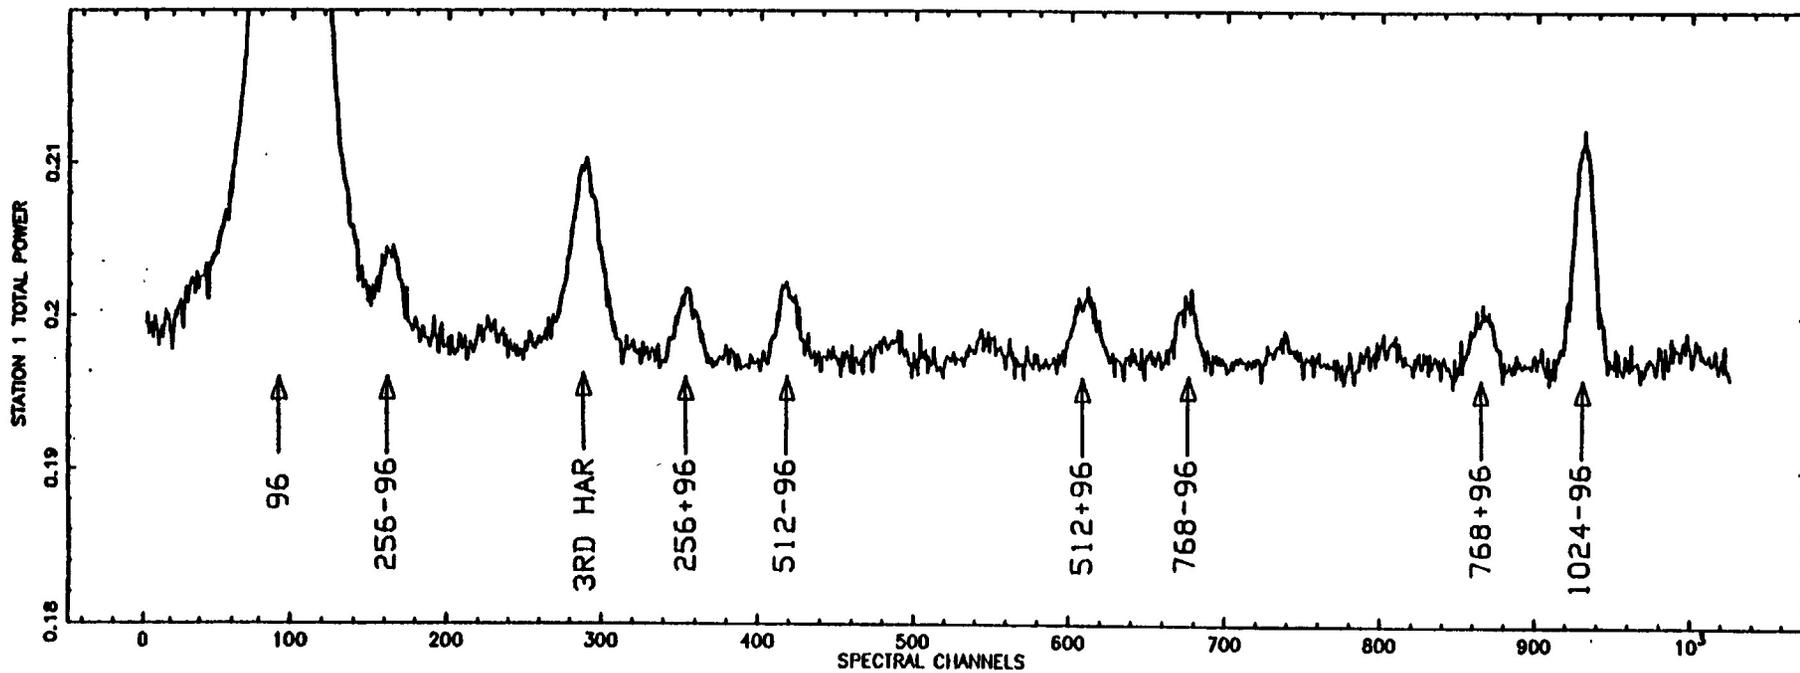


FIGURE 10

322 mag samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87



322 mag samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87



FIGURE

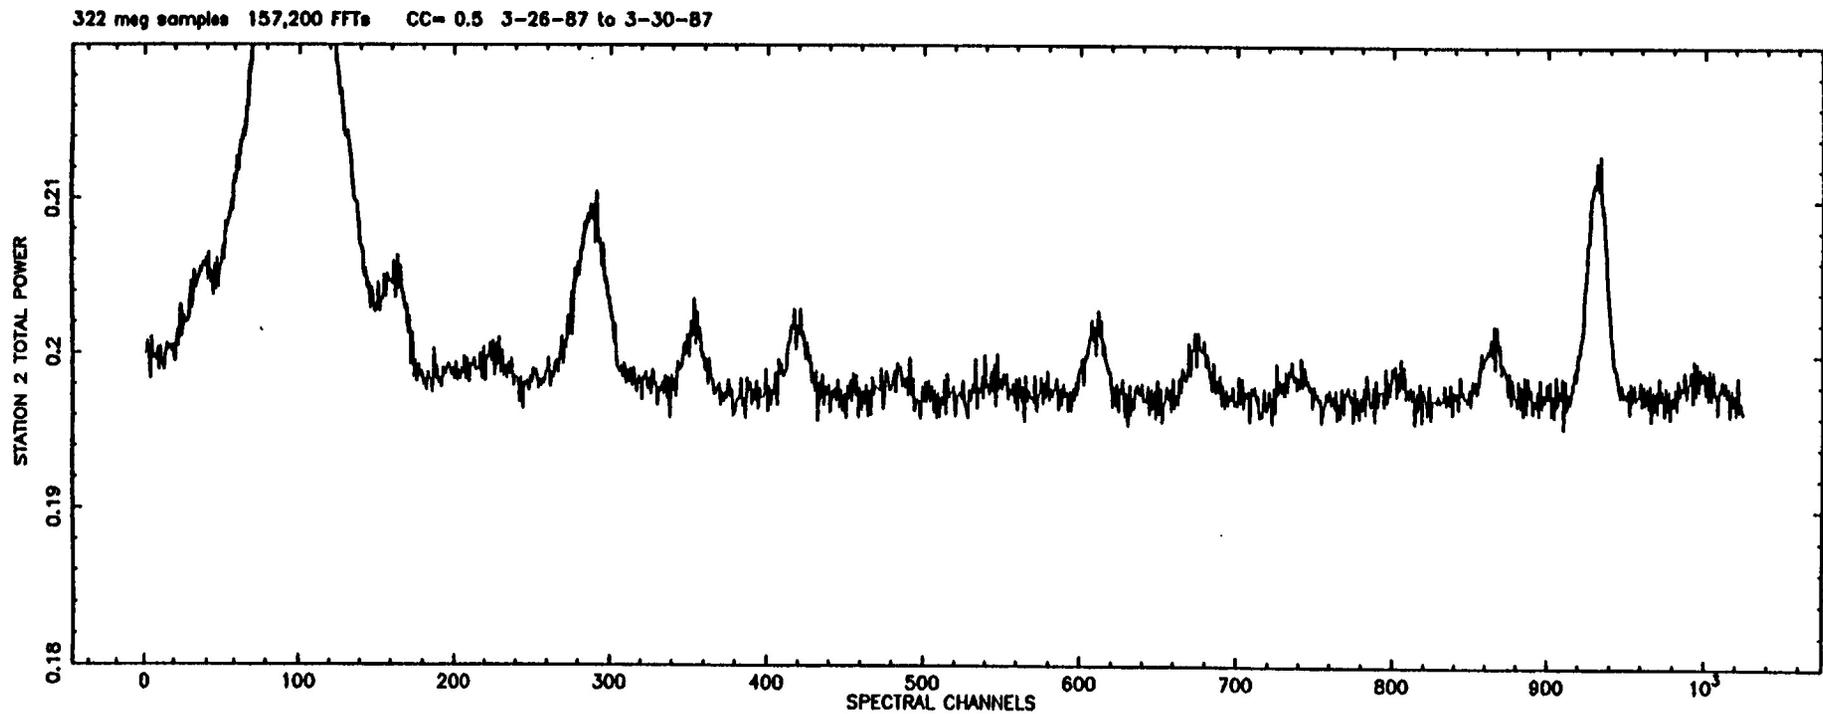
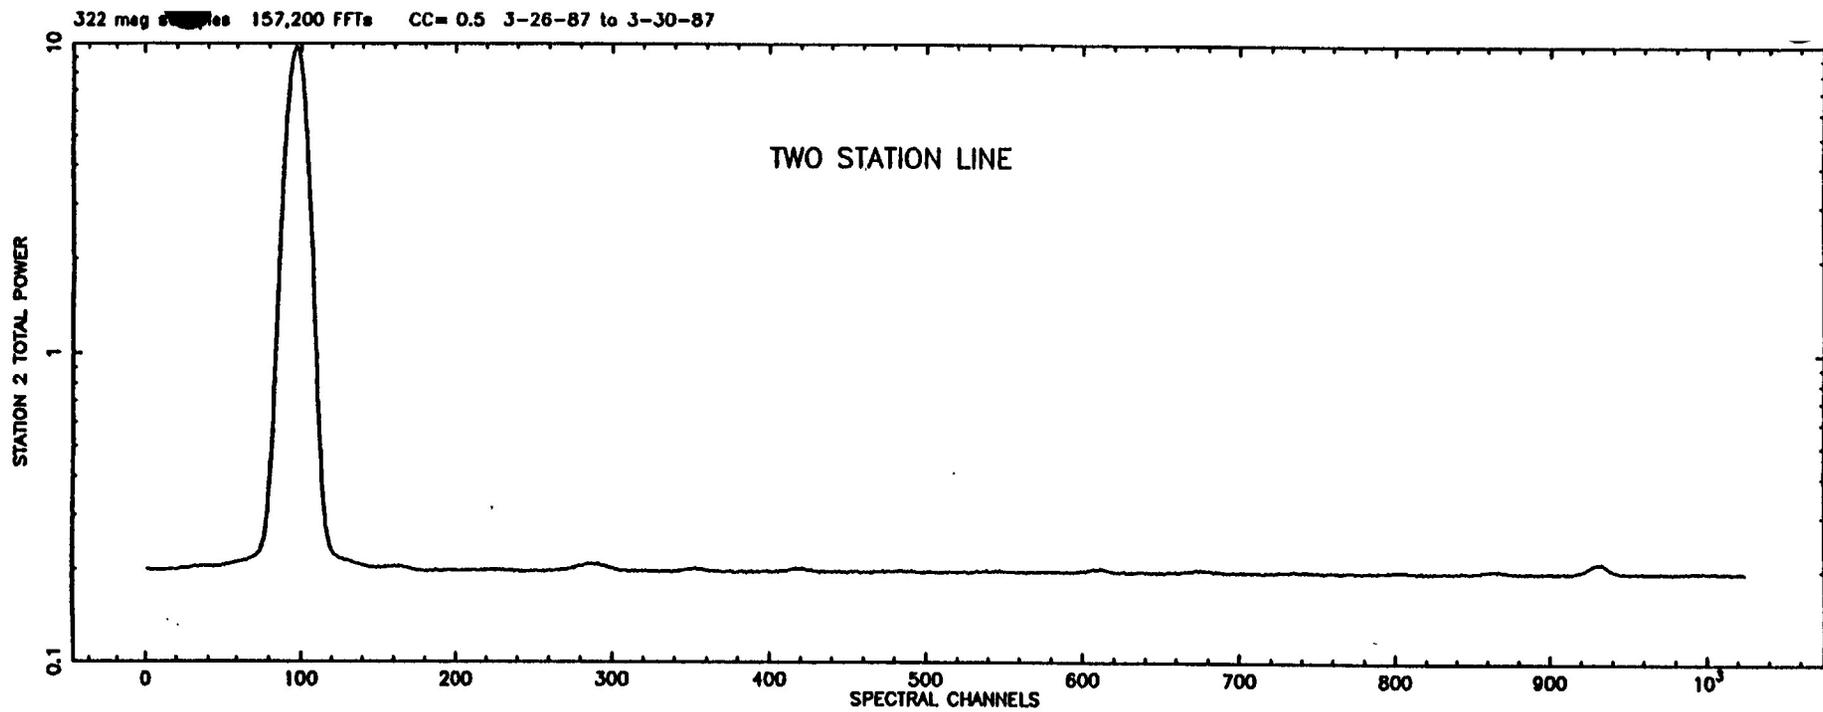
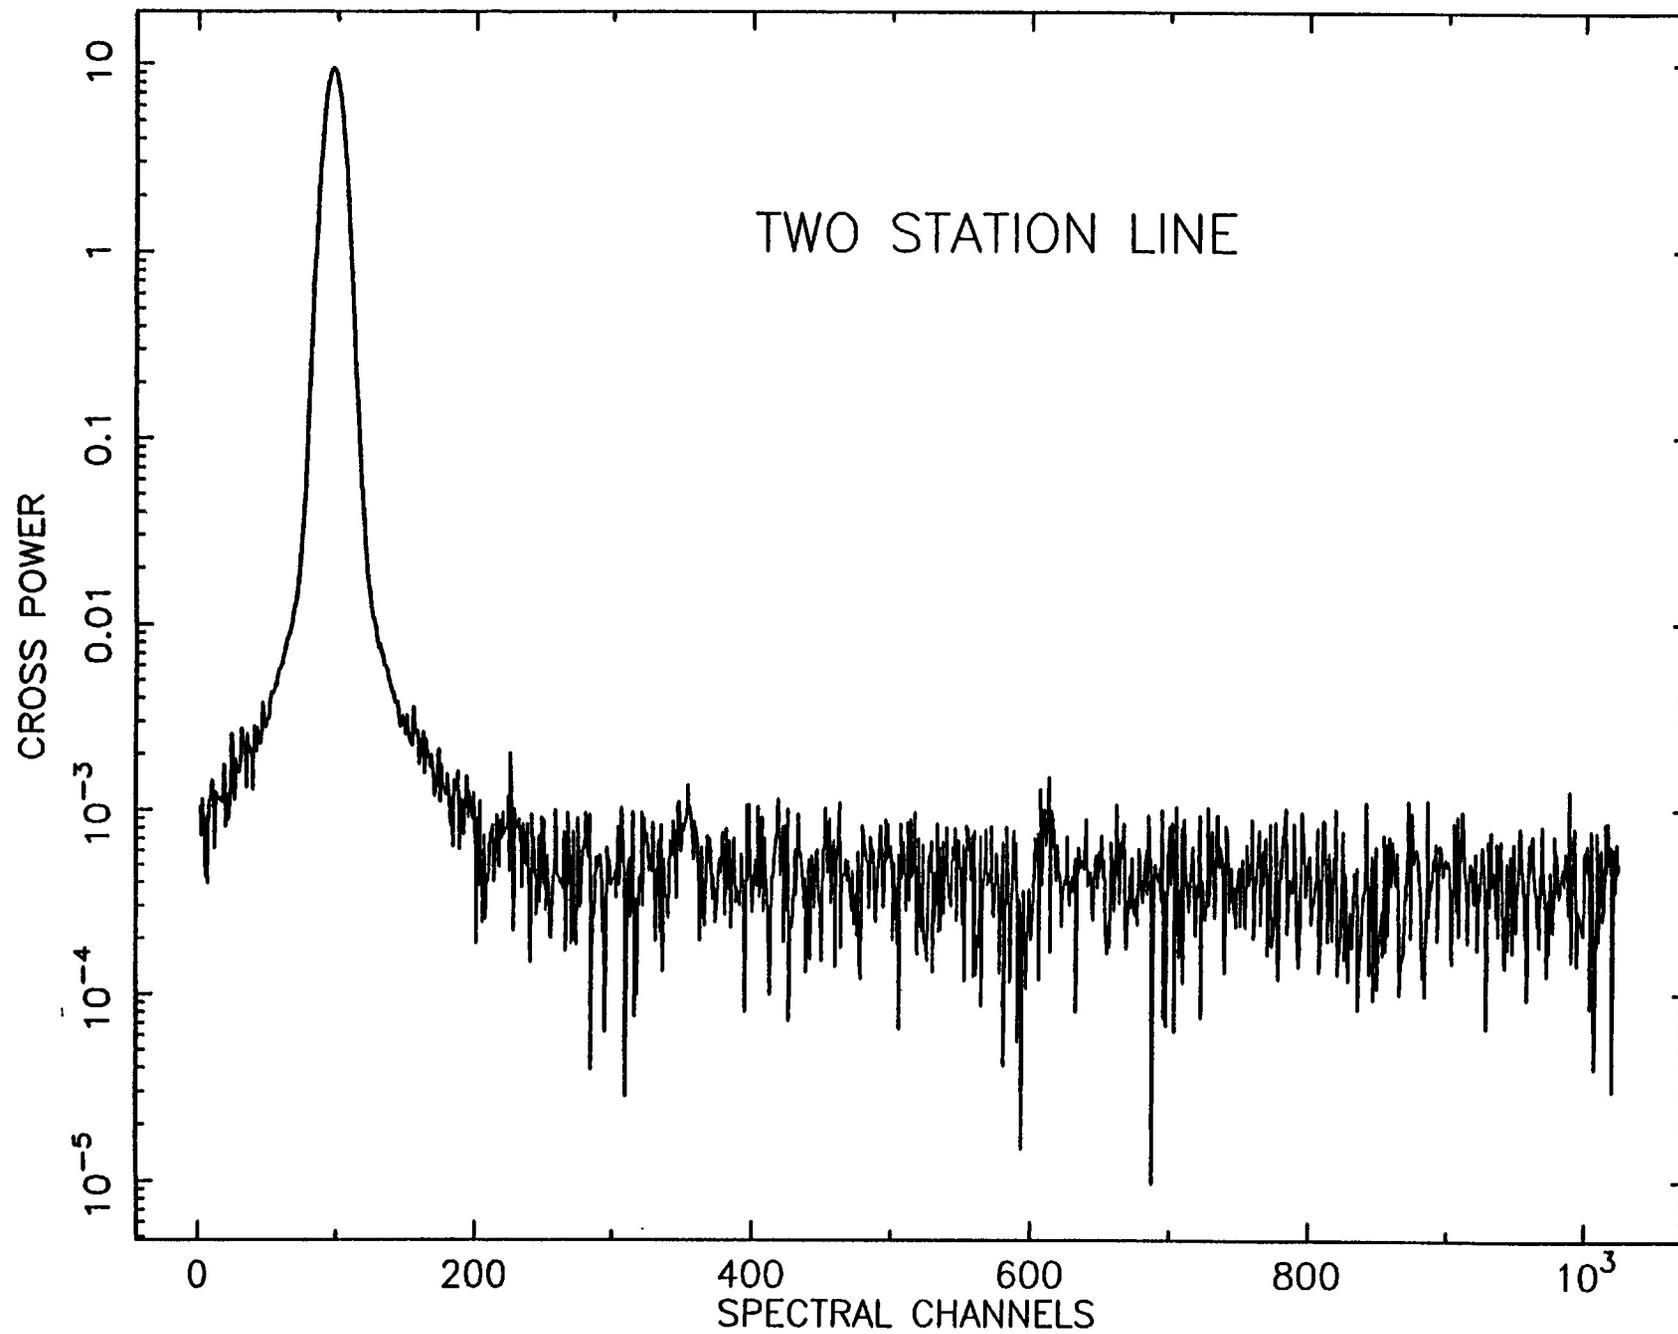


FIGURE 12

322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87



FIGURE

322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87

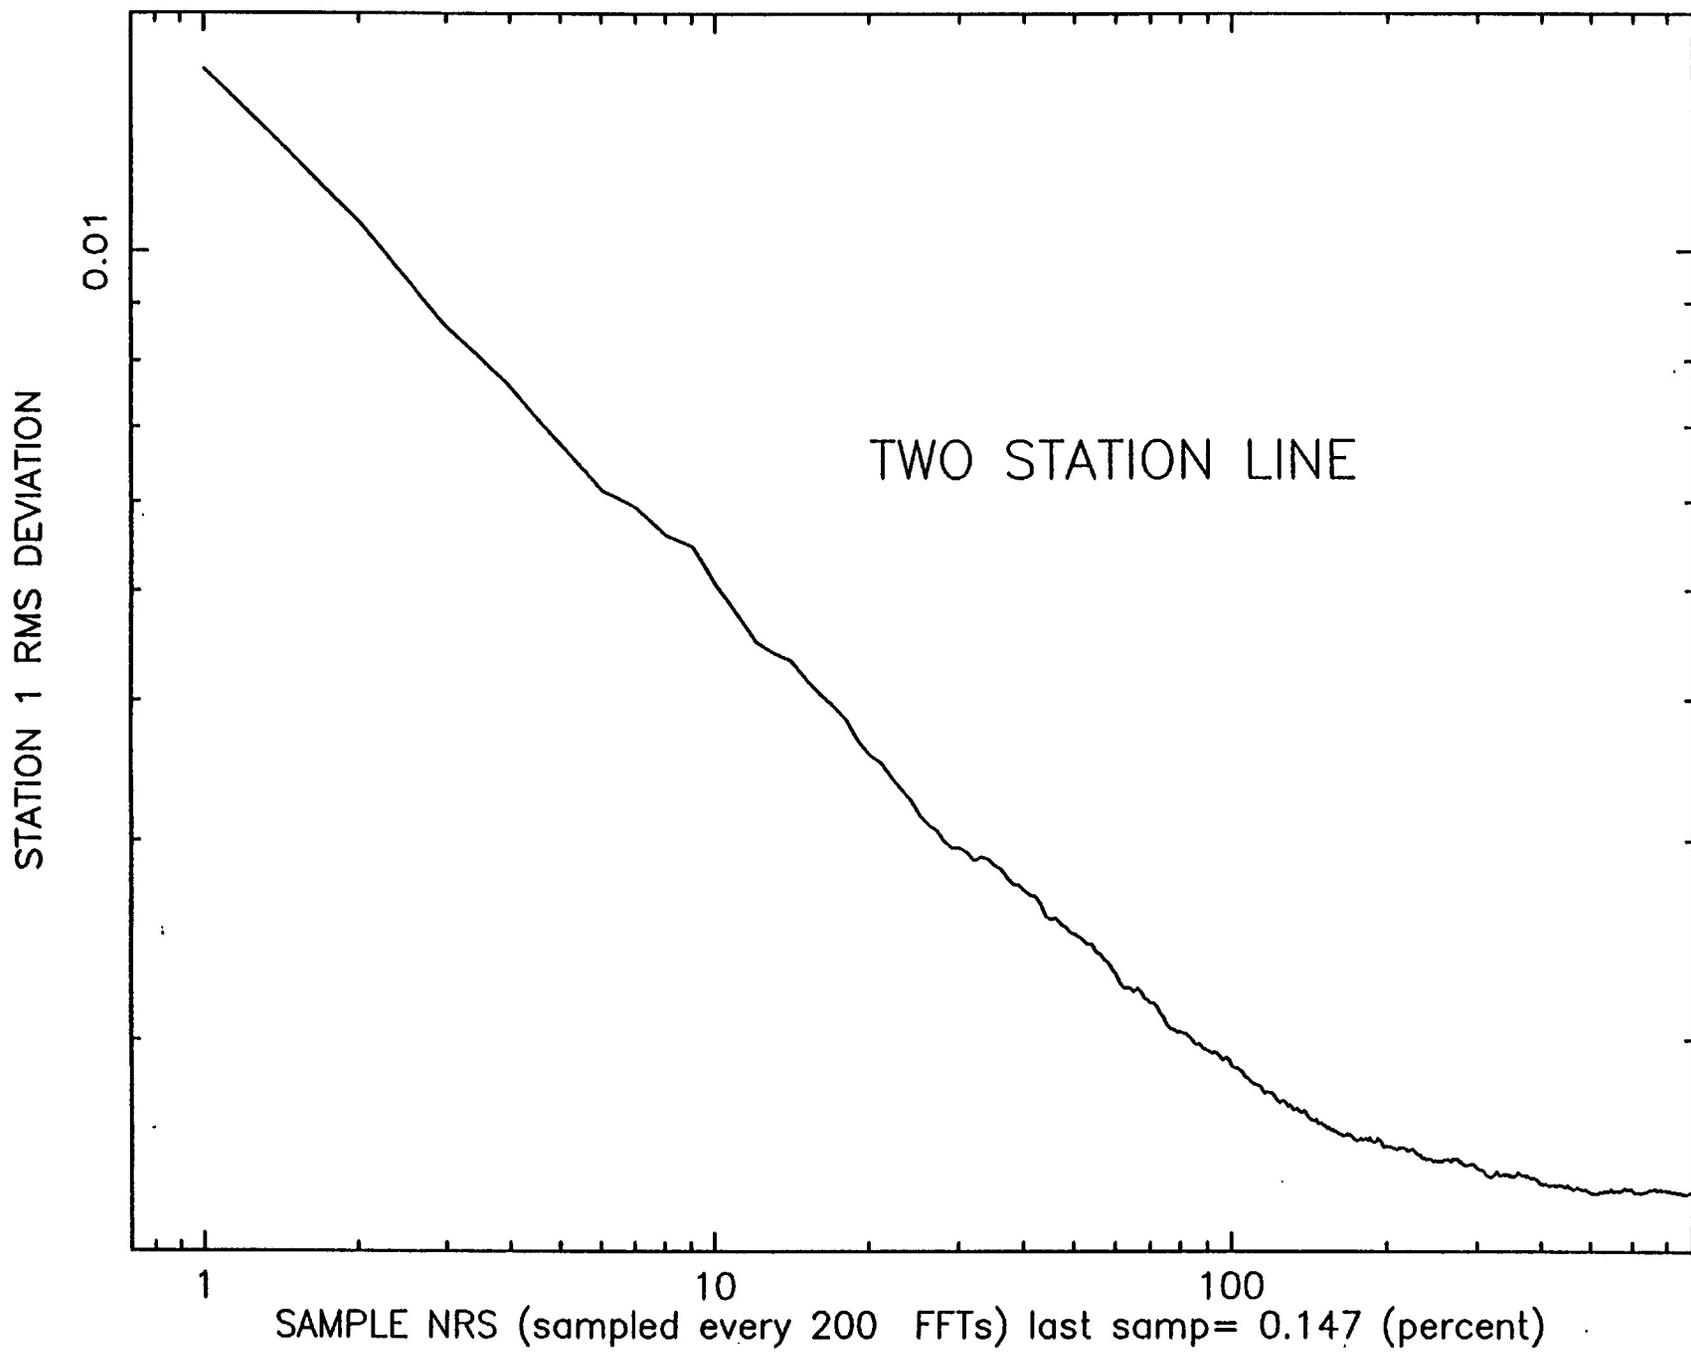


FIGURE 14

322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87

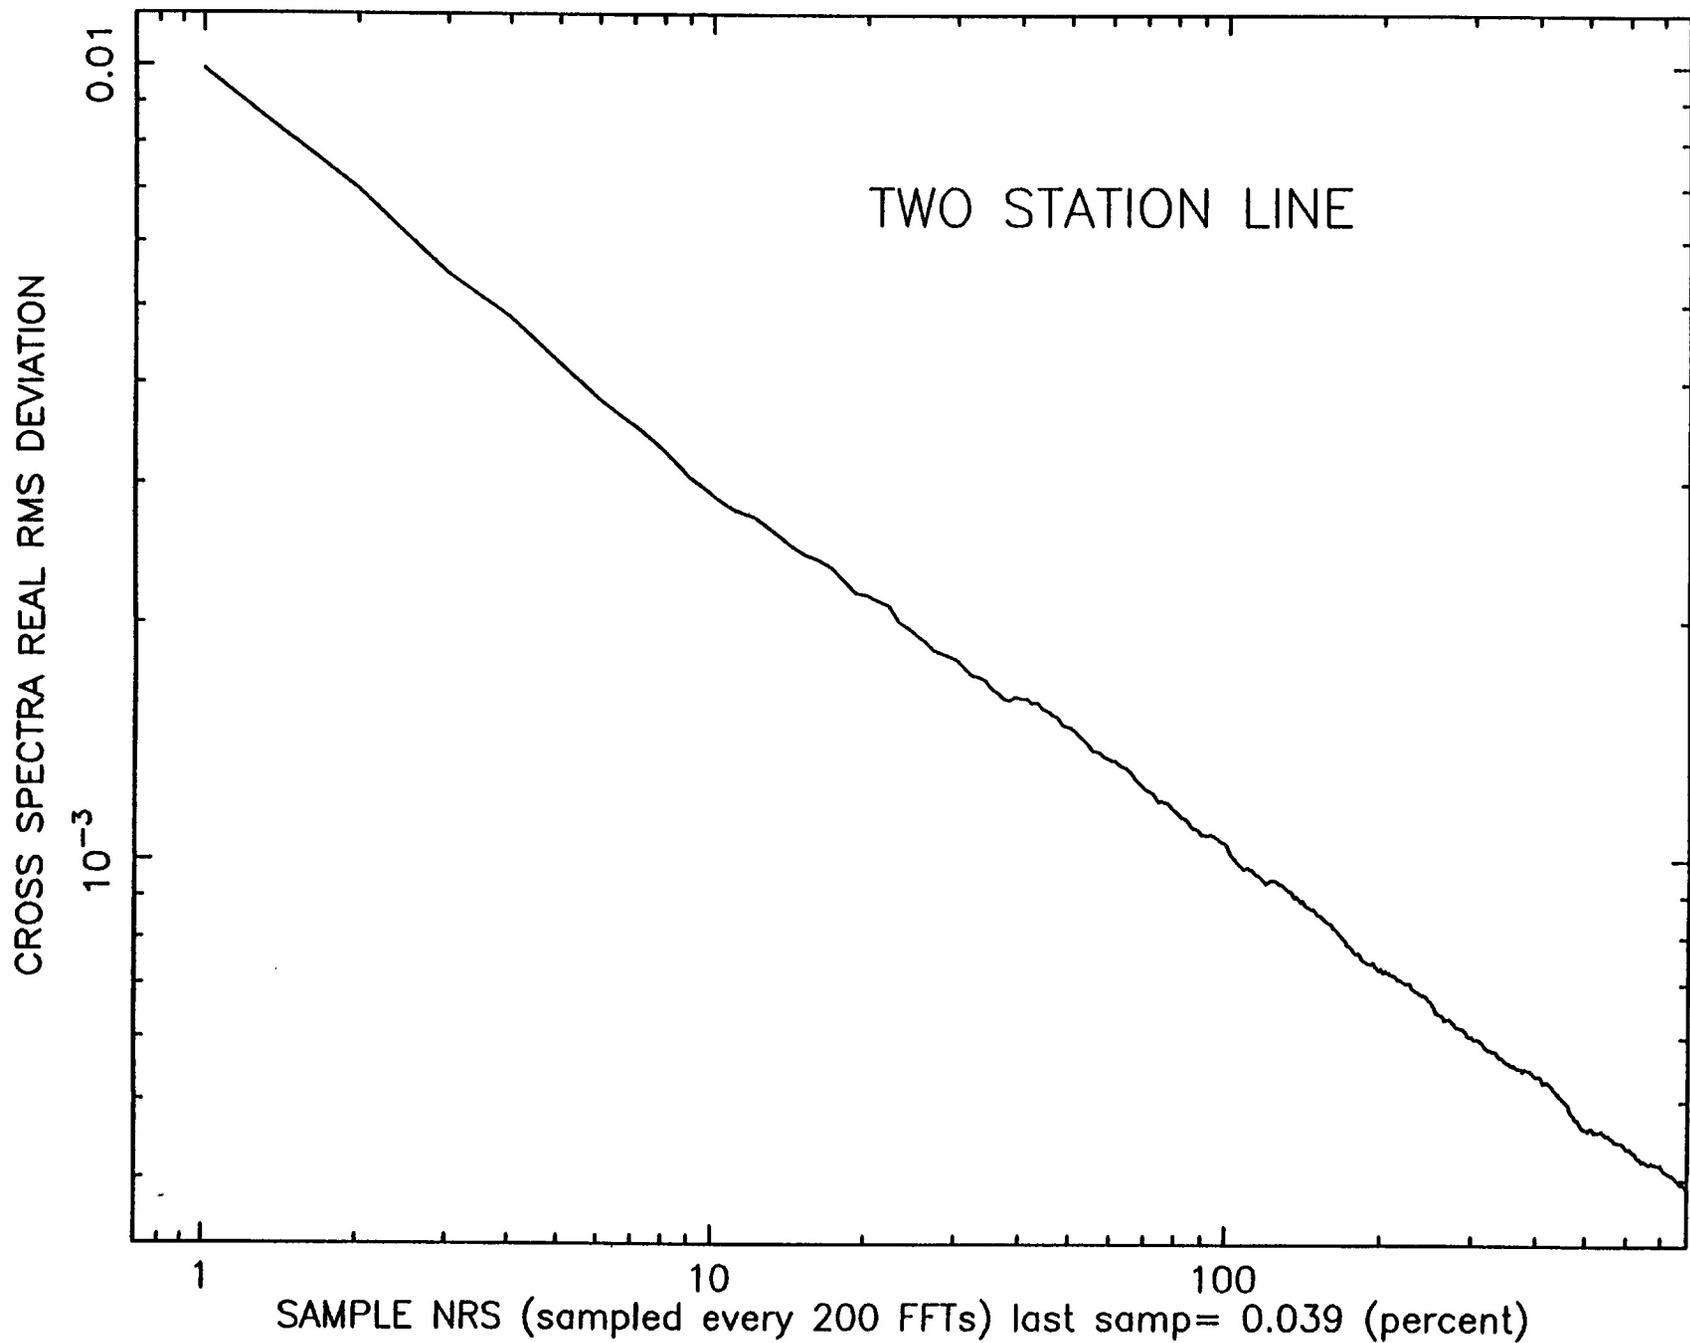


FIGURE 15

322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87

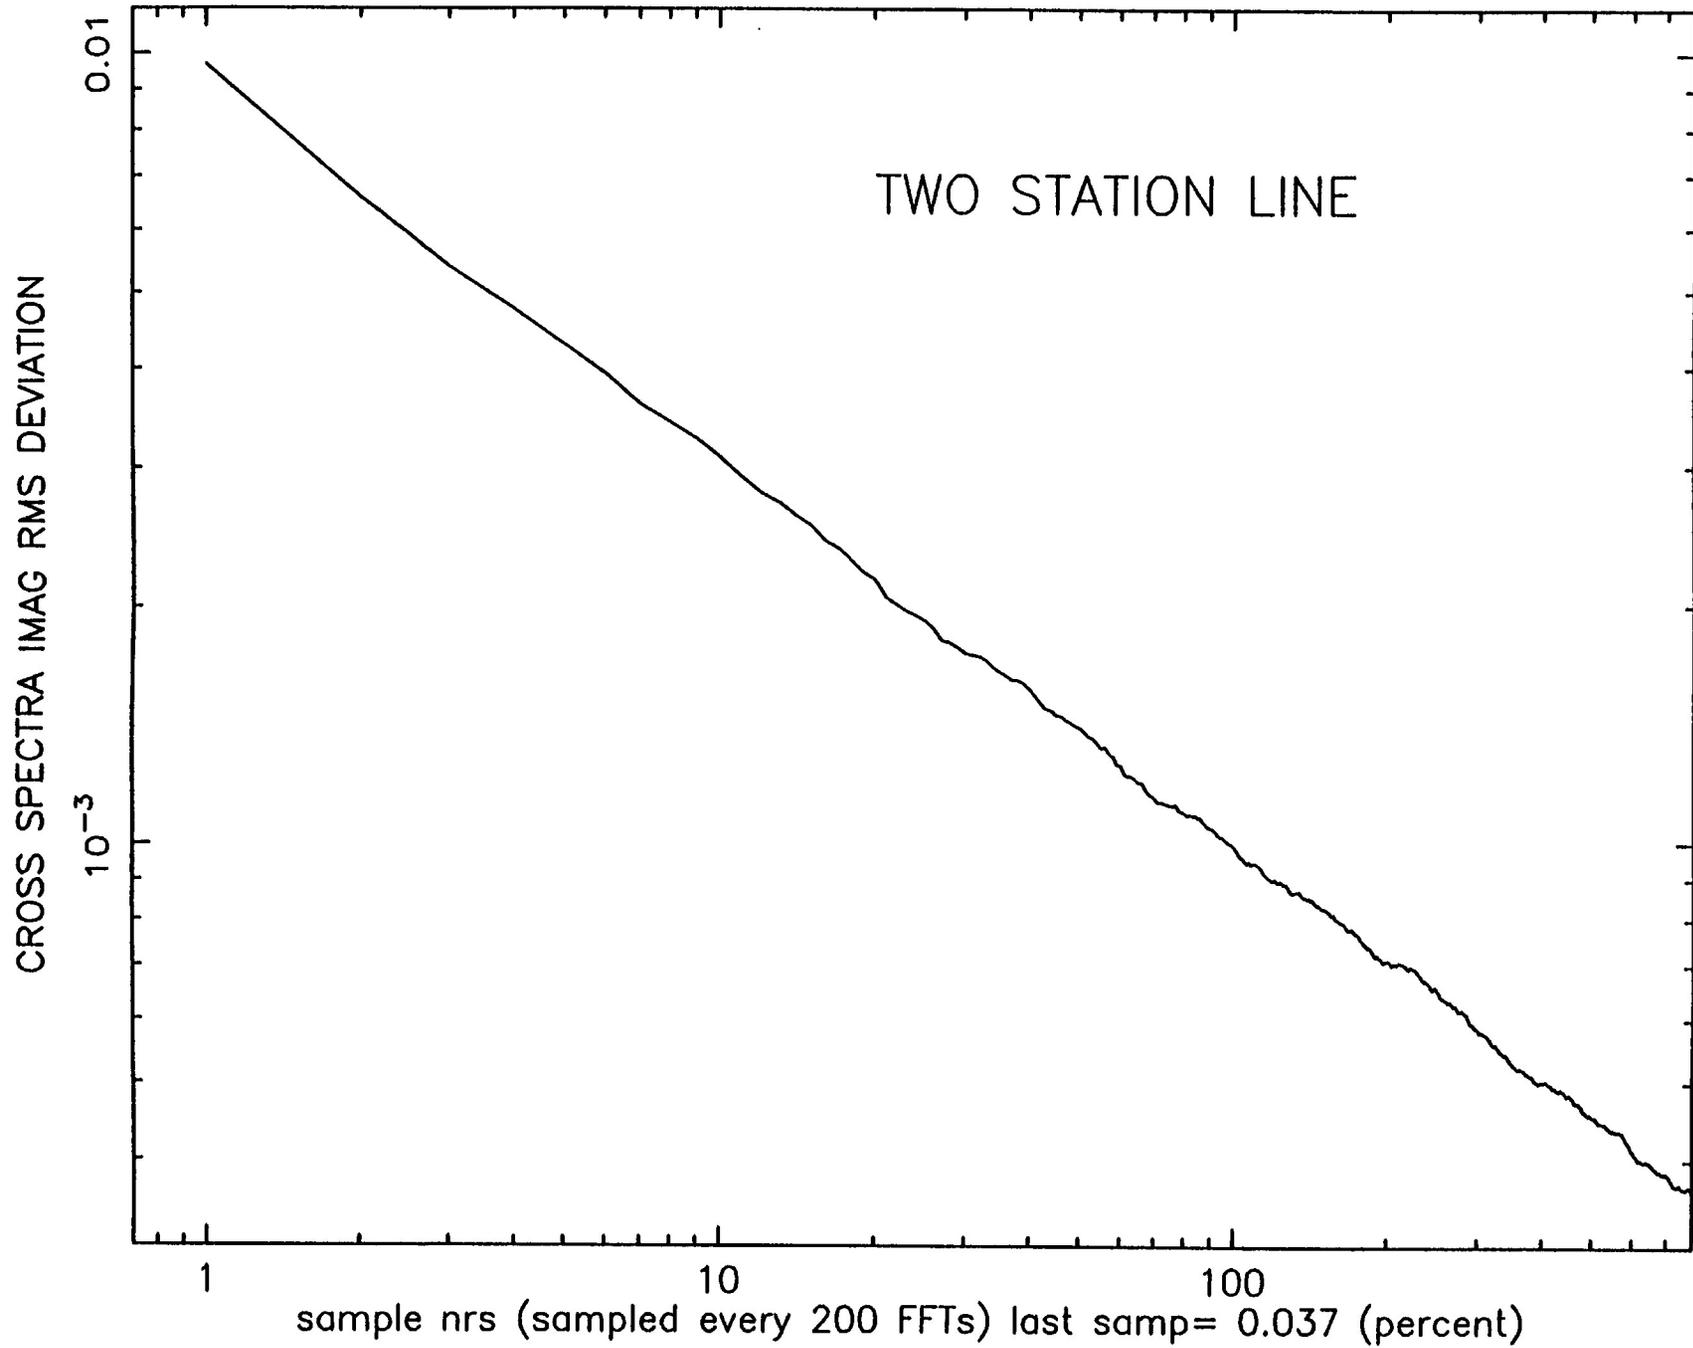
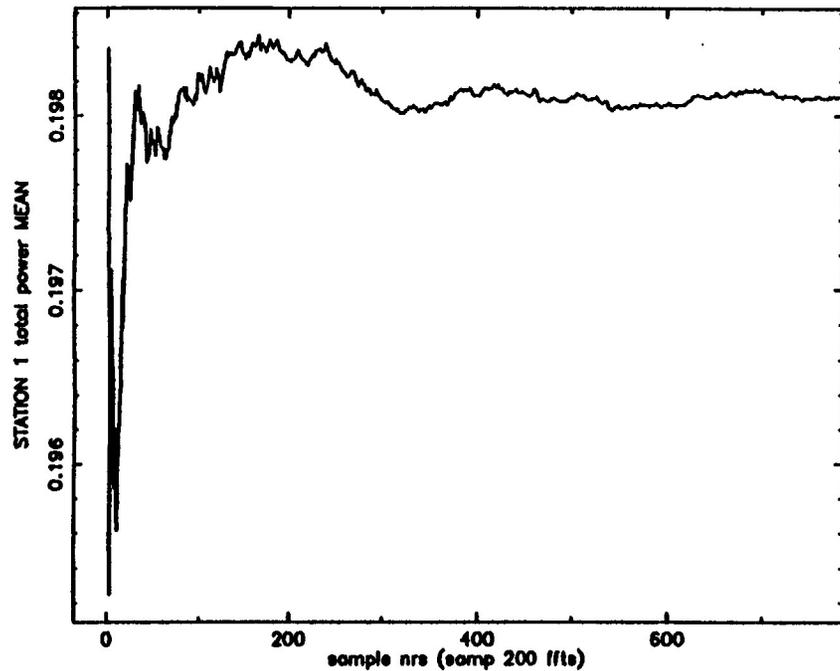
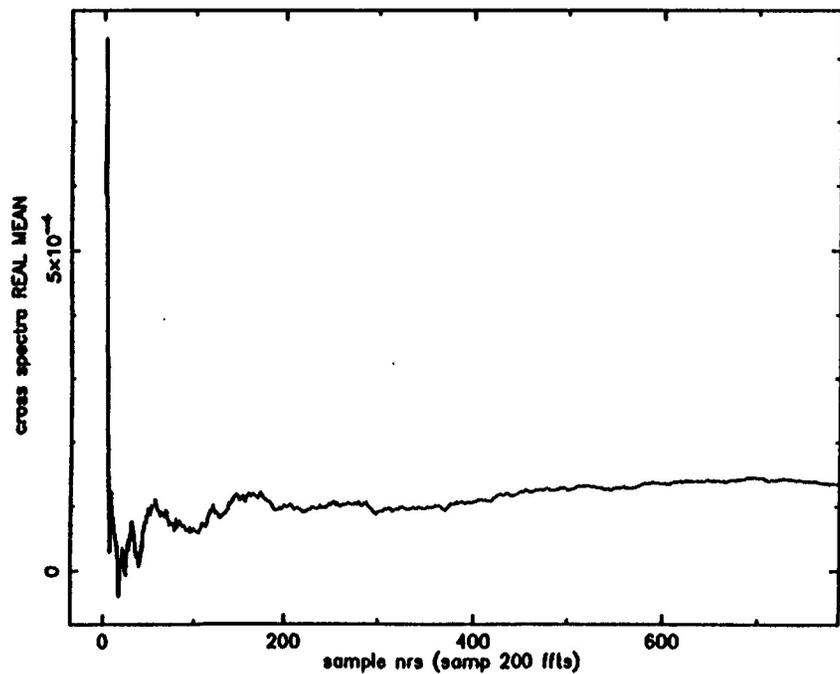


FIGURE 16

322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87



322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87



322 meg samples 157,200 FFTs CC= 0.5 3-26-87 to 3-30-87

