

THE DEEP SPACE NETWORK STABILITY ANALYZER

Julian C. Breidenthal, Charles A. Greenhall,
Robert L. Hamell, Paul F. Kuhnle

Telecommunications Systems Division
Jet Propulsion Laboratory
California Institute of Technology

14 December 1994

ABSTRACT

A stability analyzer is described, for testing of NASA Deep Space Network installation during flight radio science experiments. The stability analyzer provides realtime measurements of signal properties of general experimental interest: power, phase, and amplitude spectra; Allan deviation; and time series of amplitude, phase shift, and differential phase shift. Input ports are provided for up to four 100 MHz frequency standards and eight baseband analog (> 100 kHz bandwidth) signals. Test results indicate the following upper bounds to noise floors when operating on 100 MHz signals: -145 dBc/Hz for phase noise spectrum further than 200 Hz from carrier, 2.5×10^{-15} ($\tau = 1$ second) and 1.5×10^{-17} ($\tau = 1000$ seconds) for Allan deviation, and 1×10^{-4} degrees for 1-second averages of phase deviation. Four copies of the stability analyzer have been produced, plus one transportable unit for use at non-NASA observatories.

INTRODUCTION

The Deep Space Network (DSN) is called upon to attain high levels of frequency stability for scientific purposes. For instance, the upcoming Cassini mission to Saturn will use the DSN to attempt detection of gravitational radiation, and to observe properties of Saturn's rings, atmosphere, and satellites [1].

These and related investigations [2] measure small perturbations on a radio signal passing between the earth and a distant spacecraft. The Cassini applications are fairly typical, requiring frequency stability of a few parts in 10^{15} (Allan deviation for sampling time $t = 100$ to $10,000$ sec) and single-sided phase noise around -60 dBc/Hz (1 to 10 kHz offset from an 8.4 GHz carrier).

It is challenging to achieve such stabilities in the operational environment faced by the DSN. That environment includes months-long periods of duty; distributed, outdoor, and moving equipment; and competition for observing time. We have found that stability failures can remain hidden in the bulk of DSN activities, only to surface when the scientific experiment is undertaken. This is troublesome because most experiments cannot be repeated.

Therefore the DSN has, in the past, tested its systems using instrumentation suitable for use by specially trained personnel. This approach was expensive, however, and the time to analyze data has often allowed additional diagnostic evidence to disappear, necessitating repeat tests.

We developed a stability analyzer to enable operations personnel to rapidly measure stability in various ways, in order to lower costs and reduce response time. The particular measurements made are: power, phase, and amplitude spectra; Allan deviation; and time series of amplitude, phase shift, and differential phase shift. Our analyzer provides inputs for up to four 100 MHz frequency standards and eight baseband analog (> 100 kHz bandwidth) signals, with the possibility of expanding to accept digital inputs over a local area network. Four copies of the stability analyzer have been produced, plus one transportable unit for use at non-NASA observatories.

INSTRUMENT OVERVIEW

The DSN stability analyzer has two major components: 1) the RF and Analog Assembly, and 2) the Controller Assembly, as depicted in Figure 2.

The RF and Analog Assembly provides the conditioning and conversion of the input analog signals into a signal the controller can analyze. The equipment is installed in two parts: the 100 MHz Interface Assembly and an RF Cabinet Assembly.

The 100 MHz Interface assembly resides as close as possible to the DSN primary frequency standards (usually H-masers). Intentionally, this location is isolated from routine personnel access, as well as from as many environmental influences as possible. The assembly receives four 100 MHz inputs, which can be compared in pairs. The comparison (described further below) results in a 100 kHz signal that is sent over a fiber-optic interface to the RF cabinet. The RF assembly resides in a convenient location for access by test personnel. It provides reference frequency synthesis and distribution, switching among the possible input sources, signal conditioning in the form of amplification, and optionally downconversion with detection of zero crossings.

The Controller Assembly resides next to the RF assembly, and provides an operator interface for selection of the test type and hardware configuration, and for presentation of results. The Controller also controls details of switches and instrumentation, acquires data by means of analog-to-digital converters and a time interval counter, and analyzes the data acquired. Originally, the RF assembly was housed in one rack and the controller equipment was housed in a second rack. These cabinets have since been bolted together to form a double cabinet, and components of each have been swapped to improve ergonomics for the operator. See Figure 1 for a photograph of the double cabinet.

ANALOG ELECTRONICS DESIGN

100 MHz Interface Assembly. This assembly selects the pair of 100 MHz signals to be analyzed and converts the selected signals into a form that can be transported to the low frequency equipment. Figure 3 shows a block diagram.

Output from the 100 MHz assembly is sent on fiber optics to the low frequency equipment to prevent ground loop currents that, could induce spurious signals or noise into signals being measured, or could contaminate the frequency standards outputs. The 100 MHz Interface has four 100 MHz input ports. Two input ports are connected to H-Maser outputs, and one other port is normally used for comparing the stations coherent reference generator 100 MHz output against the H-masers. The 100 MHz signals are selected for measurement using RF relays followed by high reverse isolation amplifiers cascaded with output matrix switches. The combined isolation of both sets of switches and 60

dB reverse isolation of the amplifiers provides more than 150 dB crosstalk isolation between signals.

Switch control commands are sent over fiber optics to the 100 MHz Interface using commercial modems and digital I/O boards to address switch decoders that operate the switches.

The selected pair of inputs are frequency multiplied by 99 and 100 respectively with phaselocked cavity multipliers. The multiplier outputs at 9.9 GHz and 10.0 GHz are mixed to generate 100 MHz. The result is frequency translated to 100 kHz in an offset frequency generator, and sent on a fiber-optic link to the low frequency assembly.

The frequency conversion process yields a single 100 kHz carrier with a phase spectrum containing the relative stability of the 100 MHz inputs, with a 40 dB margin above what would be obtained from direct mixing of one input with the other input, offset by 100 kHz. Amplitude information is lost. Frequency translation to 100 kHz is necessary for the A/D converter, and allows the signal to be transported to the low frequency assembly over low-cost multimode fiber optics.

Low Frequency Interface Assembly. This assembly contains switches that select among baseband receiver signals and the 100 kHz signal from the 100 MHz assembly. The selected signals are routed to measurement ports of the computer system. Figure 4 shows a block diagram.

Baseband signals are selected by matrix switches and sent to programmable attenuators that set levels into the interface amplifiers. Another matrix switch outputs the selected signals to the desired output ports. The frequency translated 100 MHz maser-pair signal is input to the low frequency assembly on multimode fiber. The fiber-optic receiver output is ± 30 kHz bandpass filtered to eliminate aliasing of spectral components, then fed to the output matrix switches. One output of the matrix switch feeds a zero crossing detector for 1-second phase measurements. The zero crossing detector generates a 1 PPS output that is fed over fiber optics to the frequency counter. The other outputs of the matrix switch are sent on coax cable to A/D converters in the VME Assembly for other measurements of signal, phase, and amplitude. A digital I/O assembly receives Rs232 switch commands from the computer to address the switch decoders that actuate the matrix switches and set attenuation values.

Zero Crossing Detector. The stability analyzer employs two methods for phase detection: one method using software processing of A-D samples, and another using a time interval counter [5]. For the second method, we use a new design of zero crossing detector that has reduced time jitter compared to previous designs' [4]. In operation, the zero crossing detector heterodynes the signal to 1 Hz, then processes the 1 Hz output to produce 1 Hz rate, 30 microsecond wide pulses that are sent over fiber optics to the time interval counter.

Time Interval Counter. A HP 5334B Counter is modified to accept inputs from the rear panel, and to accept the fiber-optic signal from the zero crossing detector and a 10 PPS signal from the reference distribution assembly.

Reference Frequency Distribution. This assembly distributes a high-stability 10 MHz station reference to the frequency synthesizers and the time interval counter, and also generates a 10 pulse per second signal used by the time interval counter for phase detection.

Environmental concerns. The stability analyzer has been designed to minimize influence of the environment on measurements. The most environmentally sensitive equipment is placed in the frequency standards room where ambient temperature stability is better than $\pm 1^\circ\text{C}$. All signals between

the 100 MHz assembly and the stability analyzer racks are connected through fiber optics to eliminate groundloops that could induce powerline spurious into measured output. The analog electronics of the 100 MHz and low frequency assemblies are temperature stabilized with a thermoelectric control system that reduces room temperature variations by a factor of 20x. Magnetic shields around the electronic attenuate magnetic fields by more than 20 dB, thereby minimizing pickup of AC powerline harmonics.

CONTROLLER AND SOFTWARE DESIGN

The Controller Assembly consists of a Sun Microsystem Spare 2 general purpose computing system, with an attached VME computer chassis. The Spare 2 performs the user interface function, hardware control, and the display and logging of test results. The computer includes an Integrix SBus expansion unit, a 1.2 GByte hard disk, a 5.0 Gbyte Exabyte tape drive, a CD ROM reader, along with the usual monitor, keyboard, mouse, and laser printer. Serial ports are used for communication with the analog hardware and a time code translator, and an IEEE 488 bus is used to communicate with two frequency synthesizers and a time interval counter. (One synthesizer supplies the local oscillator for the last downconversion to 1 Hz as shown in Figure 4, while another supplies the sample clock for A-D conversion.) The VME chassis contains a Skybolt 811.6-V vector processor and an Analogic DVX 2503 16-bit, 400 kHz A-p converter.

The Skybolt computer is delivered with its own Unix-based operating system which allows the execution of one user program. We have written the one user program to provide custom real-time multitasking and digital signal processing. The program is designed to accommodate one test at a time, in the form of an execution script including the digital signal processing, along with some small Skybolt system tasks. The code is written in C and Fortran.

The software on the Sun runs with the Unix operating system using a Motif-style window manager environment. Custom screens allow operators to use the stability analyzer with only occasional reference to an instruction manual. Unique test script files are compiled at run-time to control test tasks, which are started in the Sun and executed in the Skybolt. The scripts are written in a custom language, similar to Structured Query Language (SQL), including higher-level operations such as Define, DoWhile, If, etc. The Sun code is written in C, some of which is computer generated by programming tools and utilities, mainly Builder's Xcessory, Lex, and Yacc.

The signal processing software supports tests for Allan deviation of phase and differential phase, time series of phase and amplitude, and spectra of signal, phase, and amplitude. Each of 17 distinct tests can be selected by the operator with a single mouse click on the display. The test configuration parameters (input source, sample rate, averaging time, etc.) are automatically loaded from editable configuration files, and can also be modified at the display by the operator.

The sample clock for A-D conversion comes from a Hewlett Packard 3325A synthesizer, referred to 10 MHz from the Reference Frequency Distribution Assembly. Although the A-D converter can handle 400 kHz, the current limit is 230 kHz because of workarounds for an unresolved hardware incompatibility. Nevertheless, this rate is adequate to handle two of the widest baseband signals (bandwidth 45 kHz) from the Deep Space Network Radio Science open-loop receiver. The frequency span of spectra can vary from 50% of the sample rate down to an arbitrarily small band about the carrier.

SIGNAL PROCESSING ALGORITHMS

Vectorized processing. The signal processing routines run on a single-board computer, the 40 MHz Skybolt, containing an Intel 1860, a floating-point vector processor with its own high-speed data cache. To achieve the best computational throughput on this processor, we avoided recursive operations, such as phase-locked loops and recursive digital filters, in favor of sequential, nonrecursive operations on large arrays, such as element-by-element vector arithmetic, inner products, finite-impulse-response (FIR) digital filters, and the fast Fourier transform (FFT), all of which are supported by Sky Computer's vector library and compiler. Throughput of 25-30 million floating-point operations per second were achieved.

Sampling the video signal. We discuss here only the processing of the signal through the A-D converter; the processing of 1-Hz zero-crossing signals through the counter has previously been documented [5]. The analog "video" signal is specified to be a sinewave with weak sidebands in a known frequency band about the carrier. (The total sideband power should not exceed about -30 dBc.) First, this signal has to be sampled at a such a rate that the sidebands of the digitized signal faithfully reproduce the sidebands of the analog signal. For example, the output of the 100 MHz Interface Assembly is a 100 kHz signal with sidebands between 85 kHz and 115 kHz. If this is sampled at 80 kHz, the sampled signal, which lives in a 40 kHz band, has a carrier at 20 kHz and sidebands between 5 and 35 kHz. The 16-bit A-D necessarily adds its own noise and distortion; fortunately, by adjusting the sample rate one can reduce their effects on measurement results by whitening the noise and moving the aliased harmonic distortion images away from the frequency band of interest.

Overview of Signal Processing. To allow the user to check the overall quality of the signal, we supply a test called "full band spectrum". This test simply computes a spectrum of the sampled signal in the maximum frequency span available, namely, half the sample rate f_s . Also provided are snapshot plots of A-D samples vs time.

The main job of the DSP is to extract the phase and amplitude modulations from the digitized video signal within a user-selected frequency B of the carrier. Two processes for this are supplied, called medium band and narrow band. Medium band is used for B from $f_s/4$ down to $f_s/256$. Narrow band is used for smaller values of B , with essentially no lower bound except that implied by the user's patience. These processes are described below. First, however, we describe a vectorized algorithm for sinewave analysis that underlies much of the processing.

The Pony Computation. At the heart of the DSP is a simple vectorized algorithm for estimating the frequency, phase, and amplitude of one batch of a sampled sinewave $c_n = A \cos(\omega n + \theta)$. The computation is divided into two parts: Pony 1, which estimates frequency ω , and Pony 2, which estimates A and θ . The Pony 1 computation uses the observation that the noiseless sinewave c_n satisfies the difference equation $c_{n-1} + c_{n+1} = (2 \cos \omega)c_n$. Accordingly, we estimate $2 \cos \omega$ as the regression coefficient of the vector $(x_{n-1} + x_{n+1})$ on the vector

(x_n) , where n runs from 1 to $N-2$. This computation requires only two inner products, of form $\sum x_n^2$ and $\sum x_n x_{n+1}$, plus some scalar arithmetic. For use in Pony 2 and elsewhere, we also generate a complex vector of powers u^n , where $u = \exp(-i\omega)$, by means of a vectorized "powers" algorithm that takes advantage of the Skybolt architecture.

Pony 2 uses W to estimate A and θ by solving the two-parameter least-squares problem $x_n = a \cos \omega n - b \sin \omega n$ for the unknowns a and b . The only vector computation needed is $\sum x_n u^n$. Then A and θ are obtained from $A \exp(i\theta) = a + ib$.

Medium Band Processing. This mode of processing operates by a sequence of mixing and filtering to extract the complex-valued analytic signal, containing only the power from the positive-frequency side of the original waveform [7], from which the amplitude and phase modulations can be extracted by a rectangular-to-polar operation. The Pony 1 calculation estimates the carrier frequency f_c , and a mixing signal $\exp(-i2\pi f_c t)$ is generated by the powers algorithm. After the right-hand part of the carrier is mixed to zero frequency, a FIR lowpass decimation filter is applied to eliminate the other part of the carrier and to select the desired frequency span $(-B, B)$. The result is the desired analytic signal within B of the carrier, shifted to zero frequency.

The analytic signal is the basis of all further processing. If a signal spectrum is wanted, then a two-sided spectrum is generated after removing the DC component (the shifted carrier). If amplitude or phase are wanted, then a rectangular-to-polar operation is applied and the phase sequence unwrapped from $(-\pi, \pi)$. (Wan, Austin, and Vilar [8] give a more efficient unwrapping method.)

Narrow Band Processing. In this mode of computation, the A-D data for the whole run are processed in contiguous batches of size N , each of which is analyzed by both parts of the Pony computation to produce a sample of batch-averaged frequency, amplitude, and phase. The bandwidth of the extracted amplitude and phase samples is $f_g/(2N)$. Because of the efficiency of the Pony computation, the DSP can keep up with the stream of A-D samples at the highest rate of the A-D converter, 400 kHz, although, as mentioned above, the analyzer is currently limited to a total sample rate of 230 kHz.

For computational efficiency, N has to be at least 200. To save storage, we do not allow batches greater than a designated maximum batch size (now 8192). Because we also wish to allow arbitrarily small analysis bandwidths, the batch averages can themselves be averaged together in groups of arbitrary size r to produce samples with bandwidth $f_g/(2Nr)$. In choosing this crude lowpass decimation method, we accepted some aliasing problems to gain simplicity, consistency, and efficiency.

The phase information computed by Pony 1 and 2 is local to each batch, and is known modulo 2π only. We have devised an algorithm to process these local data into a sequence of global phase residuals; it is essentially the same as the algorithm used for processing the 1 Hz zero crossing counter readings [5]. For the algorithm to succeed, the frequency must be changing slowly enough from batch to batch so that the current batch phase can be predicted from earlier ones within π . The algorithm issues an alarm if any prediction error exceeds $\pi/2$ in absolute value.

The low-rate sequence of amplitude and phase residuals extracted by the narrow band process can be subjected to a variety of post-processing functions, including time-series display, spectral estimation, and Allan deviation. For a two-channel test, the phase residuals of the two channels are combined by subtraction (with some adjustments) to give differential phase, which can be post-processed in the same way as single-channel phase residuals..

Spectral Estimation. Direct FFT-based spectral estimation methods are used [9]. The sequence of operations applied to a data buffer is detrending, tapering, zero-padding to a power-of-2 FFT size, applying a real or complex FFT, squaring the magnitude, equalizing the lowpass decimation filter, and scaling. Some of these elements are discussed below. A sequence of spectral estimates can be averaged to produce a run spectrum with greater statistical stability.

Spectral density of signal, phase, or fractional amplitude deviation is displayed in units of dBc/Hz, i.e., single-sideband power per Hz relative to total (carrier) power, expressed in decibels. Thus, a phase spectrum shows $S_{\varphi}(f) = S_{\varphi}(f)/2$. A signal spectrum shows both sidebands.

Each spectrum produced by the analyzer has an associated resolution bandwidth b , which is just the two-sided noise bandwidth of the spectral window. The power of a narrow spectral line in dBc equals its displayed level in dBc/Hz plus $10 \log_{10} b$. Both b and $10 \log_{10} b$ are reported to the user.

Detrending. Before applying the FFT to a data array, this analyzer preconditions the array by subtracting a linear fit obtained by drawing a straight line between the centroids of the first sixth and the last sixth of the graph of data vs time. This procedure removes both the level and slope divergencies characteristic of certain processes with stationary second increments [10], and allows the average of many array spectra to converge to a stable run spectrum. Walls, Percival, and Ireland [11], who preconditioned their data by subtracting the mean, found that the estimated spectrum for noise with a true f^{-4} spectrum depended on the number of array spectra that were averaged. For full band spectrum, no detrending is needed because most of the energy is in the carrier.

Data Tapering. To avoid problems of energy leakage from high portions of a spectrum into lower portions, each data array is multiplied by a tapering sequence drawn from a family of functions called discrete prolate spheroidal sequences (DPSS). (Actually, we use a set of convenient approximations, the "trig prolates" developed by Greenhall [13].) For full band and medium band spectra, we use a single bell-shaped taper from this family. For narrow band spectra we use a nonadaptive, unweighed version of Thomson's multiple-taper method [12][9]. An array of detrended data is tapered by four orthogonal tapering sequences, giving rise to four distinct "eigenspectra", $S_0(f)$ through $S_3(f)$. These are averaged to produce the spectral estimate $S(f)$ for the array. In a broadband noise region, the $S_k(f)$ are approximately uncorrelated, and hence $S(f)$ has about one-fourth the variance of each $S_k(f)$. For a given frequency resolution, the desired statistical stability is achieved from fewer data arrays.

Allan Deviation. From a stream of narrow band or 1 Hz zero crossing phase residuals, the analyzer produces estimates of Allan deviation with estimated drift removed, using the simple three-point drift estimator recommended by

Weiss and Hackman [14]. The required 'r-overlapped sums for first and second moments of second T-differences of phase are accumulated in real time.

To generate conservative error bars for plus or minus one standard deviation of Allan variance, we assumed a random-walk-frequency model of phase noise. Using a method of Greenhall [15], we carried out a numerical computation of ν , the equivalent degrees of freedom of the drift-removed Allan variance estimator, as a function of M, the number of summands. The sequence of ν vs M was fit with a simple empirical formula. Then, if σ is the estimated Allan deviation, the reported error bar is

$$\sigma(1 \pm (2/\nu)^{1/2})^{1/2}.$$

Because of severe negative bias of the drift-removed estimator for small M, results are reported only for $M \geq 4$.

TEST METHODS

A series of tests of the stability analyzer were conducted at JPL's Frequency Standards Laboratory in order to demonstrate first, that the results of the stability analyzer agree with those of other measurement equipment, and second, that it meets its noise floor requirements. Noise floor results are given in Table 1 below.

Allan deviation runs of length at least 24 hours were carried out on pairs of 100 MHz frequency standards. The results were compared to those from an existing FSL Allan Deviation test set and found to agree within 5%. The noise floor was measured by splitting the single output of an H-maser and applying it to two inputs of the stability analyzer. These tests were carried out in both zero crossing detector mode and the narrow band phase modes.

Time series of differential phase were tested using a HP 3326 dual channel synthesizer as the input source. The two outputs of the synthesizer were manually steered in frequency to produce phase drifts of known amplitude. Comparison was made to the results from a HP 8508 phase meter, and also to a strip chart recording the phase difference. This last signal was developed by simple mixing between the two outputs of the synthesizer. These tests were also run with both channels of the synthesizer set at the same frequency for at least 15 hours, to observe the noise floor.

Spectra were tested in a variety of ways. The signal sources were two H-masers for the 100 MHz inputs, one H-maser and an HP 8662 synthesizer, or one or two HP 3325 synthesizers for the baseband analog inputs. In the latter two cases, one synthesizer was modulated either by another synthesizer to simulate spurious signals, or by a HP 3561 noise source to simulate phase noise. The spectrum was then compared to the results from a HP 3589 or 3561 spectrum analyzer. The results agreed within a typical 2 dB peak-to-peak variation between spectral bins. For noise floor tests, a single H-maser signal was divided and applied for comparison at two inputs.

Table 1. Stability Analyzer Noise Floors

Test	Input Source	T	σ
Allan Deviation	100 MHz	1 sec	6e-15
		10 sec	2e-15

		100 sec	2e-16
		1000 sec	3e-17
Phase Spectrum	100 MHz	Freq.	Spectral Density
		1 Hz	-126 dBc/Hz
		10 Hz	-135 dBc/Hz
		>100 Hz	-142 dBc/Hz
Signal Spectrum	baseband	1 Hz	-92 dBc/Hz
		10 Hz	-97 dBc/Hz
		>100 Hz	-98 dBc/Hz
Phase Spectrum		1 Hz	-98 dBc/Hz
		10 Hz	-104 dBc/Hz
		>100 Hz	-105 dBc/Hz
Amplitude Spectrum		1 Hz	-70 dBc/Hz
		10 Hz	-85 dBc/Hz
		>100 Hz	-88 dBc/Hz
Diff. Phase	baseband	Avg. Time	Phase Error
		1 r3ec	<0.001 deg rms
		1000 sec	<0.04 deg rms

ACKNOWLEDGMENTS

We would like to thank Gerard Benenyan, Michael Grimm, Diana Howell, Nancy Key, Barron Latham, Beverly St. Ange, Eric Theis, and John Vitek for their contributions to this project.

The work described in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] A.J. Kliore et al, "Investigation description and science requirement document, Cassini radio science team", pp 6 & 13, private communication, 10 Feb 1993
- [2] S.W. Asmar and N.A. Renzetti, The Deep Space Network as an Instrument for Radio Science Research, Publication 80-93, rev 1, Jet Propulsion Laboratory, Pasadena, CA, 15 Apr 1993
- [3] P.F. Kuhnle, "NASA/JPL Deep Space Network frequency and timing", Proceedings of the 21st Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, 1991
- [4] G.J. Dick, P.F. Kuhnle, and R.L. Sydnor, "Zero crossing detector with submicrosecond jitter and crosstalk", Proceedings of the 21st Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, 1990
- [5] C. Greenhall, "A method for using a time interval counter to measure frequency stability", IEEE Trans UFFC, vol 6, pp 478-480, 1989
- [6] s. Marple, Digital Spectral Analysis with Applications, Prentice-Hall, 1987
- [7] D. Vakman, "Computer measuring of frequency stability and the analytic signal", IEEE Trans Instrum Meas, vol 43, pp 668-671., 1994
- [8] K.-W. Wan, J. Austin, and E. Vilar, "A novel approach to the simultaneous measurement of phase and amplitude of oscillators", Proc 44th Freq

- Control Symp, pp 140-144, 1990
- [9] D. Percival and A. Walden, Spectral Analysis for Physical Applications, Cambridge, 1993
 - [10] P. Lesage and C. Audoin, "Characterization and measurement of time and frequency stability", Radio Science, vol 14, pp 521-539, 1979
 - [11] F. Walls, D. Percival, and W. Ireland, "Biases and variances of several FFT spectral estimators as a function of noise type and number of samples", Proc 43rd Freq Control Symp, pp 336-341, 1989
 - [12] D. Thomson, "Spectrum estimation and harmonic analysis", Proc IEEE, vol 70, pp 1055-1096, 1982
 - [13] C. Greenhall, "Orthogonal sets of data windows constructed from trigonometric polynomials", IEEE Trans ASSP, vol 38, pp 870-872, 1990
 - [14] M. Weiss and C. Hackman, "Confidence on the three-point estimator of frequency drift", Proc 24th PTTI Meeting, pp 451-460, 1992
 - [15] C. Greenhall, "The fundamental structure function of oscillator noise models", Proc 14th PTTI Meeting, pp 281-294, 1982

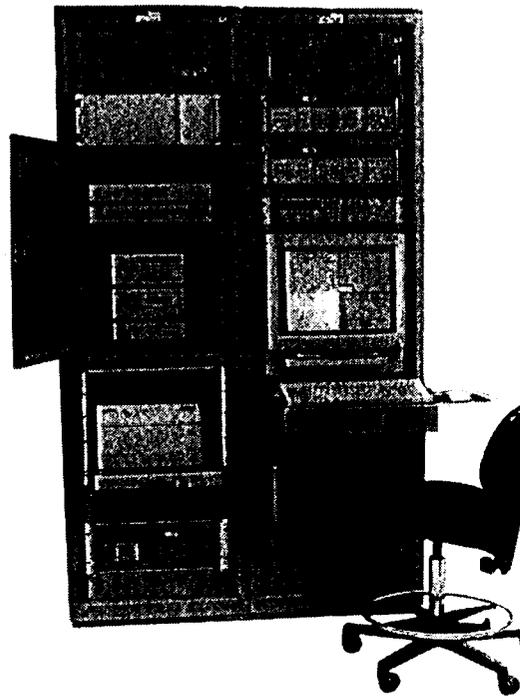


Fig. 1. Stability analyzer rack arrangement.

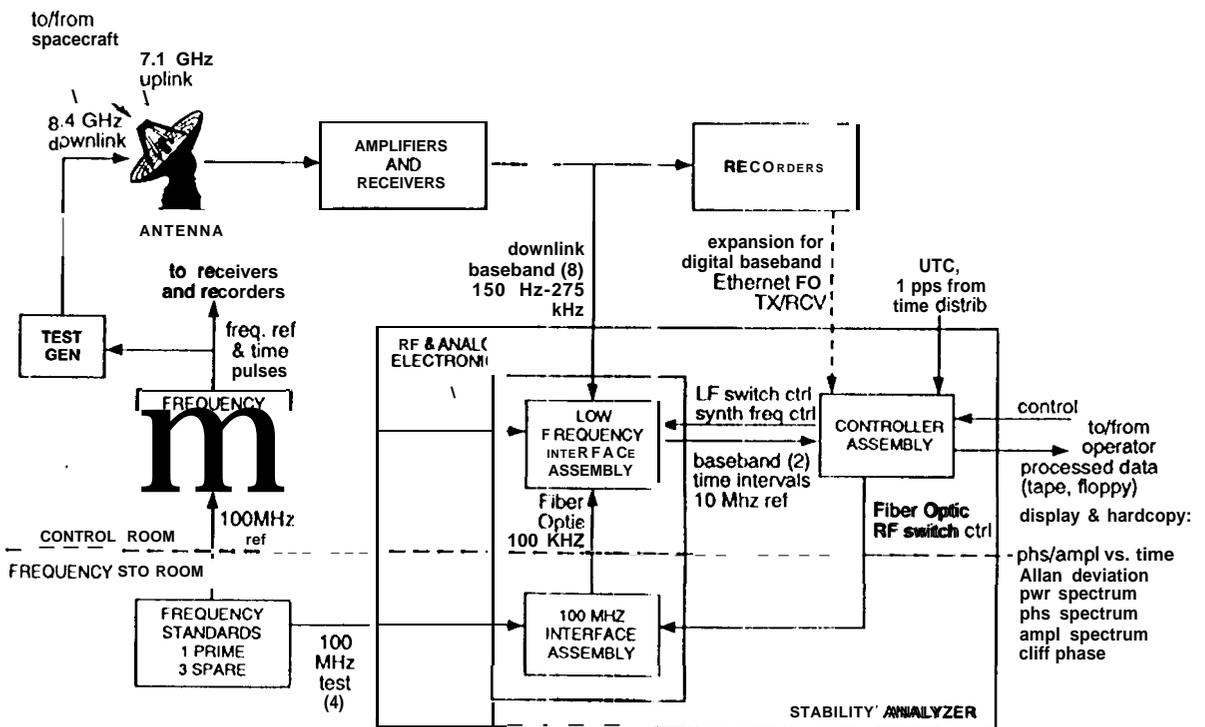


Fig. 2. stability analyzer in a typical DSN installation.

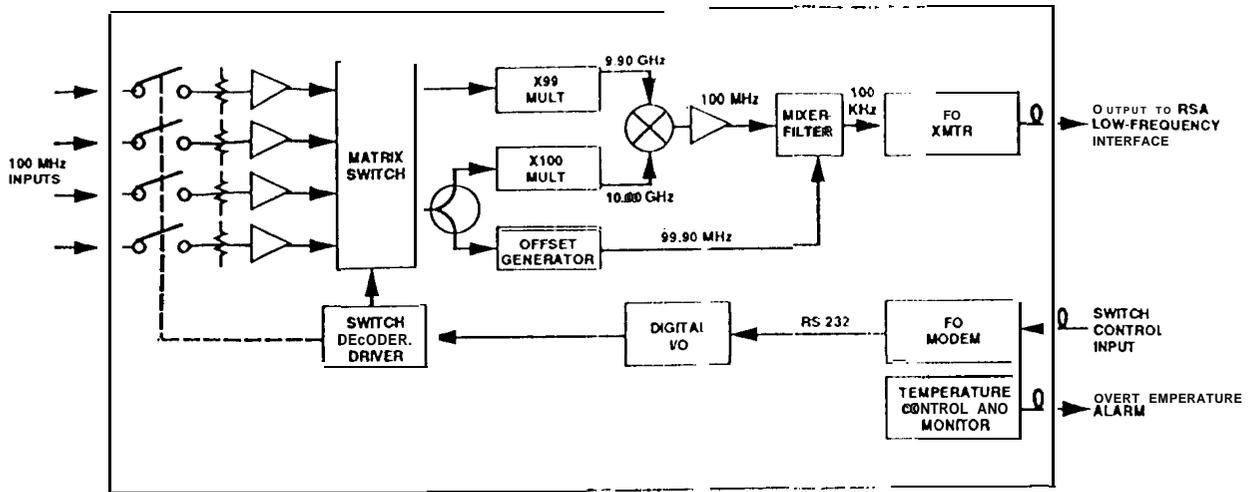


Fig. 3. 100 MHz interface block diagram.

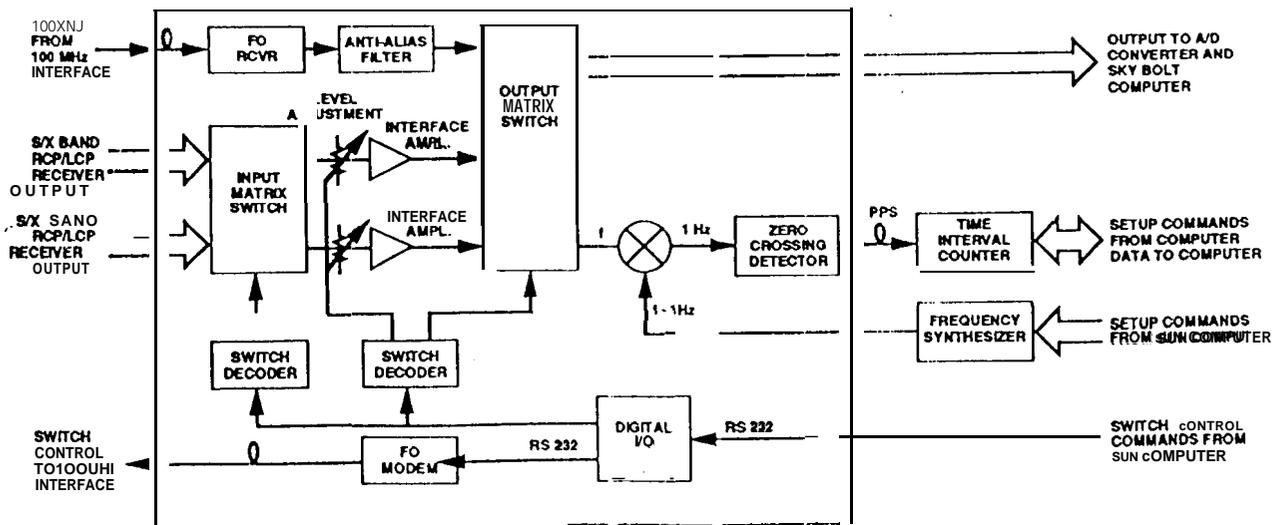


Fig. 4. Low frequency interface block diagram.