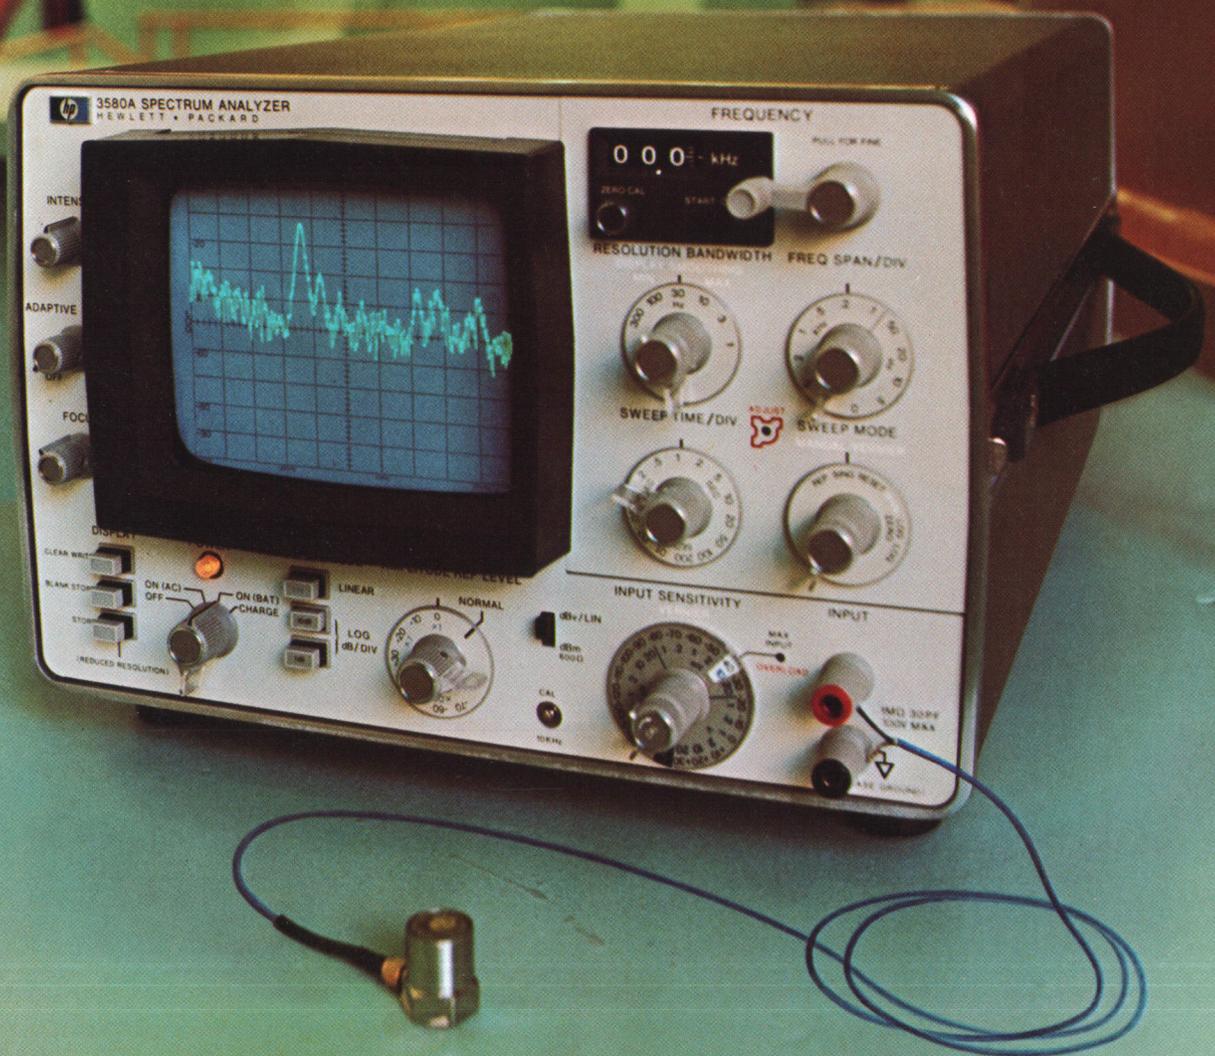


SEPTEMBER 1973

HEWLETT-PACKARD JOURNAL



A Low-Frequency Spectrum Analyzer That Makes Slow Sweeps Practical

Tuning through a 5 Hz-to-50 kHz range with a one-hertz bandwidth must be done slowly, if accurate results are wanted. This new spectrum analyzer speeds up the process while bringing a number of other conveniences to this class of instrument.

By William L. Hale and Gerald E. Weibel

SPECTRUM ANALYZERS AND WAVE ANALYZERS perform the same function—separate and measure the individual frequencies that make up a complex signal. In the low-frequency range, they are useful for analyzing the characteristics of sounds, of mechanical vibrations, of communications signals, and of electrical waveforms in general.

For many applications spectrum analyzers are the most popular because they sweep repetitively through the frequency range of interest, displaying all the signal components at one time as amplitude versus frequency on the cathode-ray tube.

Low frequencies, however, present special problems to these instruments. To allow for the response time of the very narrowband filters needed for good separation of spectral components, the frequency sweeps have to be so slow that the cathode-ray tube does not retain enough of the slow-moving trace to obtain a meaningful display.

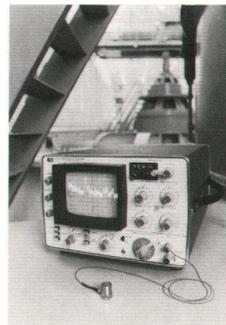
One way around this difficulty has been to use storage CRT's. But now, present-day digital technology makes possible another way of coping with painfully slow sweeps, a way that also brings other conveniences to low-frequency spectrum analysis. Digital storage of a spectrum makes it possible to combine the convenience of non-storage CRT display with the selectivity of a precision low-frequency analyzer. The new HP Model 3580A Spectrum Analyzer (Fig. 1) sweep tunes at a rate appropriate for the selective bandwidth employed, stores the resulting spectrum in a digital memory, and then reads out the stored information repetitively for flickerless display on a conventional CRT.

Fundamentally a Precision Instrument

Basically, the Model 3580A is a high-quality analyzer. Minimum bandwidth of its selective circuits is 1 Hz, and this is useful over the entire 5 Hz-to-

50 kHz range. It has a display range of more than 80 dB, 20 to 40 dB greater than others. Coupled with the high selectivity, this makes it possible to detect low-level spurious signals and powerline related sidebands close to a large spectral line (Fig. 2).

The noise level is less than 30 nV with the 1-Hz filters, allowing the most sensitive range to be 100 nanovolts full scale. The new instrument gives this performance in a compact package that is easily carried and that can be operated on internal batteries.



Cover: *Analysis of mechanical vibrations using an accelerometer as a vibration pickup is one of the several uses suggested for the new Model 3580A Spectrum Analyzer. Our thanks to the personnel of the Bureau of Reclamation, Region 2, U.S. Department of the Interior, for letting us photograph the*

3580A on one of the huge 84 megawatt water-powered generators in the majestic powerplant at Shasta Dam, near Redding, California.

In this issue:

A Low-Frequency Spectrum Analyzer That Makes Slow Sweeps Practical, by William L. Hale and Gerald E. Weibel **page 2**

A High-Performance Beam Tube for Cesium Beam Frequency Standards, by Ronald C. Hyatt, Louis F. Mueller and Terry N. Osterdock **page 14**

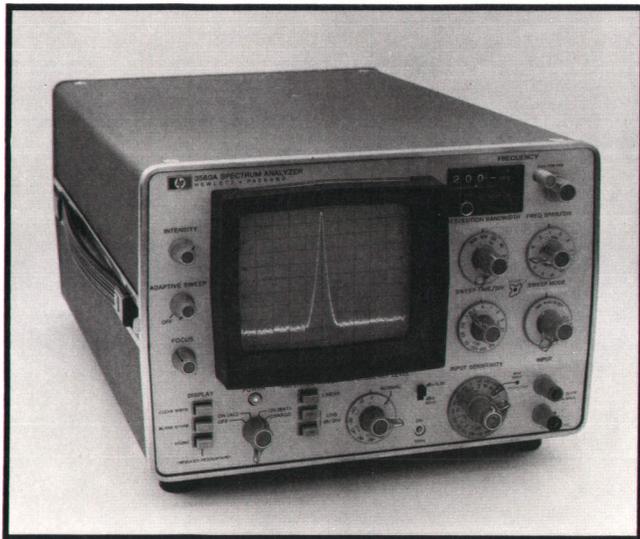


Fig. 1. Model 3580A Spectrum Analyzer spans a 5 Hz-50 kHz frequency range with frequency sweeps optimized for narrow-band selectivity. Even at very slow sweep rates, a flickerless display is obtained through use of digital storage.

Speeding a Slow Sweep

One of the conveniences that the new Model 3580A brings to low-frequency spectrum analysis is adaptive sweep. To the operator, this works like the conventional baseline clipper in that spectral components lying below a selected amplitude level are blanked out (Fig. 3). In the 3580A, wherever the spectrum is below the selected threshold, the sweep is speeded up by a factor of 20 or so. It slows to the selected sweep-tuning rate only when sweeping

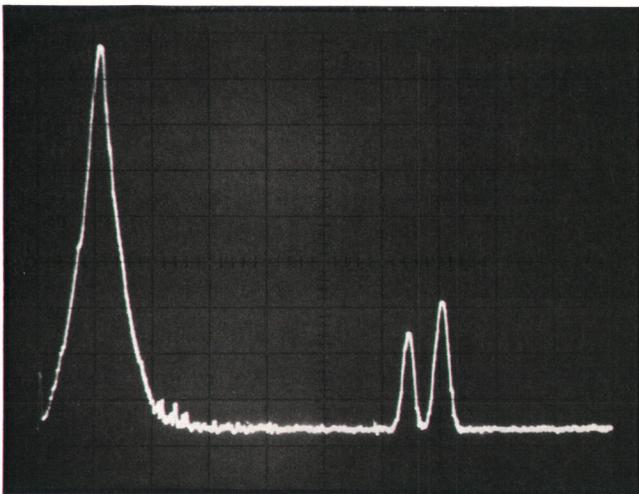


Fig. 2. Wide dynamic range and 1-Hz bandwidth enable the Model 3580A Spectrum Analyzer to resolve closely-spaced, low-level spectral components. Here, the Analyzer separates 54-Hz sidebands, caused by a cooling fan in the equipment rack, from 60-Hz powerline-related sidebands in this 10 kHz oscillator signal. Vertical calibration is 10 dB/div; frequency span is 10 Hz/div.

through spectral components that exceed the threshold by 6 dB. Since many spectra include only a few responses of interest, the new adaptive sweep system reduces the time needed for high-resolution scans of low-frequency spectra. A factor-of-ten reduction in overall sweep time is not uncommon.

Another convenience, made possible by the digital storage is the ability to retain a spectrum for later simultaneous display with a new spectrum for comparison purposes (Fig. 4). The stored spectrum can be displayed continuously, or it can be blanked out and then recalled to the display whenever desired.

How Well It Works

The capabilities of the new Model 3580A Spectrum Analyzer are perhaps most easily described with reference to the front-panel, Fig. 5. The two-speed FREQUENCY tuning crank selects either the beginning or center of the range of frequencies to be examined by sweep tuning, as determined by the switch below the frequency indicator. Tuning resolution is 1 Hz and tuning accuracy is nominally ± 100 Hz. Frequency stability, however, is remarkably good, within ± 10 Hz/hour.

The frequency span covered during a sweep can be as narrow as 50 Hz or as wide as 50 kHz. A zero-frequency span is included for observing amplitude fluctuations at a single frequency for a period of time.

The sweep time is selectable from 0.1 to 2000 seconds per sweep. The selective circuits can have one of six bandwidths from 1 to 300 Hz. If the bandwidth selected is too narrow to respond fully as the frequency is swept, the ADJUST warning light comes on, indicating that either the bandwidth or the sweep time should be increased, or the frequency span should be reduced. If the controls are set so as to cause the instrument to sweep either above 50 kHz or below 0 to "negative" frequencies, the y-axis circuits are blanked so only the baseline is displayed in the out-of-range portion of the spectrum.

The instrument sweeps repetitively or it can be set to sweep once then stop, retaining the spectrum on display. The operator can also tune manually.

For the convenience of those who wish to use the instrument as a frequency-response test set, a single-range logarithmic sweep (from 20 Hz to 43 kHz) is provided for making Bode plots. A tracking oscillator output (0 to 1V into 600 Ω) is available on the rear panel as the measurement stimulus.

Input Flexibility

The input impedance is 1 megohm shunted by 30pF, making it possible to use a standard oscillo-

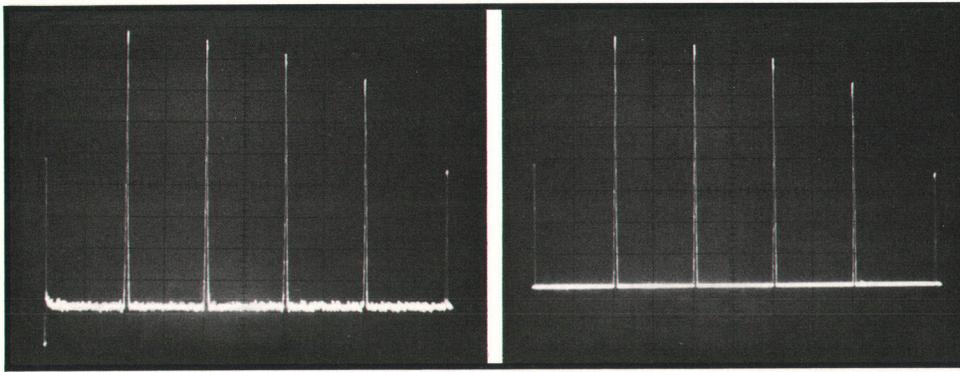


Fig. 3. Adaptive sweep speeds up the frequency scan when responses are below a selected level. It took 100 seconds to scan the spectrum in the left-hand photo without adaptive sweep. With the adaptive sweep threshold set above the noise level (right), the same spectrum was scanned in 12 seconds.

scope probe if higher input impedance is needed and/or if large-amplitude signals are encountered. The maximum input signal level allowed for non-distorting operation without a probe is 30V (no damage occurs with signals up to 100V rms).

As has been common practice, two amplitude controls are provided. The INPUT SENSITIVITY control allows the operator to adjust the signal level to the maximum permitted by the linear range of the input amplifiers, for best signal-to-noise ratio without distortion (the intermodulation resulting from nonlinear operation would add false responses to a spectrum). The AMPLITUDE REF LEVEL control expands the display vertically for examination of fine detail. A mechanical coupling between the AMPLITUDE REF LEVEL control and the indicia of the INPUT SENSITIVITY control enables the overall CRT calibration factor to be read directly from the INPUT SENSITIVITY dial.

A CAL signal for verifying calibration of both axes of the display is provided (see Fig. 5). The

vertical calibration of the display can be either linear or logarithmic. The linear mode allows amplitude measurements in volts. The log mode enables the entire 80-dB dynamic range to be displayed at one time (a front-panel switch allows calibration either in units of dBV or, with an external 600 Ω termination at the input, in dBm units). In the log mode, the 1-dB/div button gives 10 \times expansion. The AMPLITUDE REF LEVEL control then serves as an offset control so any 10 dB segment of the normal display can be brought on screen at 1 dB/cm.

The instrument normally operates on ac line power but, to eliminate ground loops or operate where line power is not conveniently available, it can also operate for more than 5 hours on an optional battery pack (field-installable). The batteries are rechargeable while in the instrument (16 hours for full recharge) but, to keep the power supply at a reasonable size and weight, the instrument was not designed to allow instrument operation at the same time the power supply is recharging the batteries.

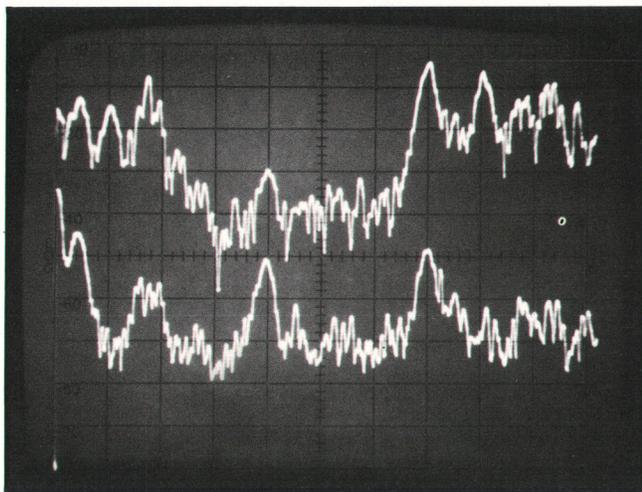


Fig. 4. A spectrum can be stored and then presented simultaneously with a new spectrum for comparison. This photo shows the 0-1 kHz vibration spectrum at two different points on a pump motor (one was offset 20dB to prevent overlap). Comparisons such as this would be helpful in identifying the location of resonances.

Filter Shape

The selective filter is the heart of any wave or spectrum analyzer, as the characteristics of the filter determine the ability of the instrument to separate individual spectral components. The 3-dB bandwidth tells only part of the story, however. How steeply the response curve falls indicates the ability of the analyzer to reject a large component when measuring a small one closely adjacent. The shape factor—the ratio of the -60-dB bandwidth to the -3-dB bandwidth—is a good indicator of this. The Model 3580A has a shape factor of 10 on all bandwidths except the 300 Hz width where it is 8.

But this is still not the whole story. The shape of the response curve is important to a sweeping analyzer, as this determines how long it takes for the filter to respond and stabilize to a change in signal level. The filter response curve preferred for spectrum analyzers is the Gaussian curve as this provides the fastest response without ringing.

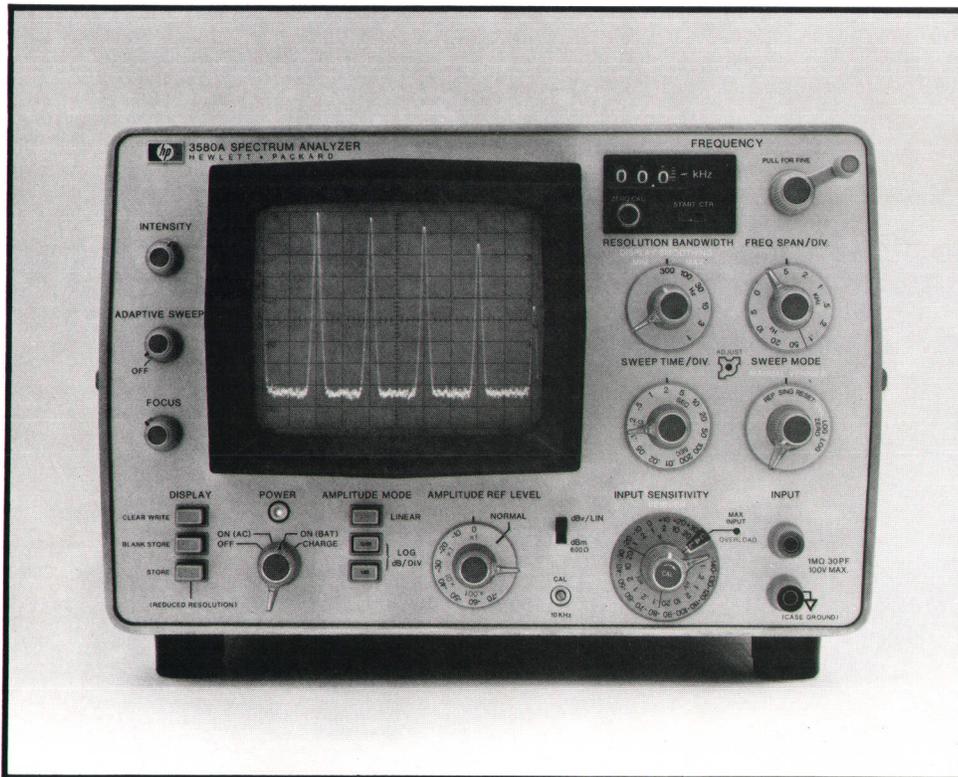


Fig. 5. Model 3580A Spectrum Analyzer displaying the spectrum of its own calibration signal, a crystal-controlled, 10-kHz train of 15 μ s pulses. The even spacing of the spectral lines illustrate the instrument's display linearity. Adjusting the screw-driver CAL control to place the fundamental at the top CRT graticule line calibrates the top-line reference level within 1.5% accuracy.

The filter characteristics desired for the Model 3580A were most readily approximated by cascading five crystal resonators, all tuned to the same center frequency. So that the passband shape will not vary with temperature, the crystals are matched both for temperature coefficient and for the temperature where the curve of frequency vs temperature has zero slope.

FET amplifiers isolate the filter sections from each other. Bandwidth is easily switched by change of the damping resistors that affect circuit Q, with appropriate adjustment of amplifier gain. Overall filter response is shown in Fig. 6.

An important by-product of the filter design is that alignment is straightforward and very fast.

Tuning Range

The crystals available for the selective circuits dictated the choice of 100 kHz for the IF channel, if 1-Hz bandwidth were to be realized. To keep the instrument's costs as low as possible commensurate with the performance expected, a single-conversion scheme is used. This dictated an upper frequency range of 50 kHz to avoid problems with image frequencies.

The selective filters are also a major factor in determining the low-frequency end of the tuning range. The filters determine how close the local oscillator can be tuned to the IF frequency without feeding through the IF. To begin with, good mixer

balance suppresses LO feedthrough (zero beat response is more than 30 dB below the display reference level). With the 1-Hz filter switched in, there is an overall 80 dB suppression of the local oscillator signal when tuned only 5 Hz above the 100-kHz IF. The instrument can be tuned lower, of course, but IF feedthrough must then be accounted for. At 4 Hz, for example, the LO feedthrough is 70 dB below the reference level.

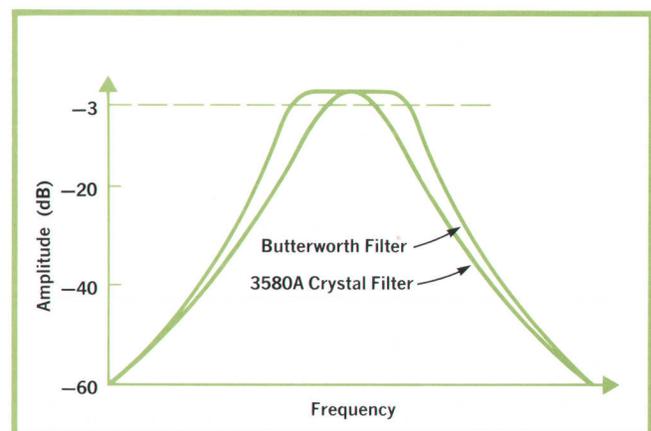


Fig. 6. Overall frequency response of Model 3580A's 1-Hz selective circuits compared to a 3-Hz Butterworth filter that has a shape factor of 4. The two have the same resolution at the -60 dB level but the 1-Hz filter, besides having better resolution at higher levels, also has faster response.

High Stability

To make the use of very narrowband filters practical, an uncommonly "quiet" local oscillator is required. Oscillator instabilities result in an apparent widening of the response curve.

The stability problem is compounded when the oscillator must be tuned electrically. Amplifier noise on the tuning control voltage can contribute as much instability as the oscillator itself.

Then there is the problem of making frequency a linear function of the control voltage, this to make the three-place digital tuning dial read accurately. The design that evolved uses an oscillator with a varactor-tuned LC tank circuit. By operating the oscillator at ten times the LO frequency, or 1-1.5 MHz, LC tank-circuit parameters were found, using a varactor diode of the hyper-abrupt junction type, that gives a voltage/frequency characteristic that is linear within a few percent. A tuning voltage change of only seven volts tunes this oscillator over the desired range.

Noise on the tuning voltage must be less than a few nV/\sqrt{Hz} at noise frequencies below 100 Hz, however. Otherwise, FM noise sidebands would not be held below the -90 dB level. Because of the noise problem, the use of amplifiers in the control loop was avoided, a step that precluded the use of diode shaping networks to linearize the V-to-F characteristics.

Instead, a precision discriminator, as shown in Fig. 7, is used to linearize the V-to-F characteristics.

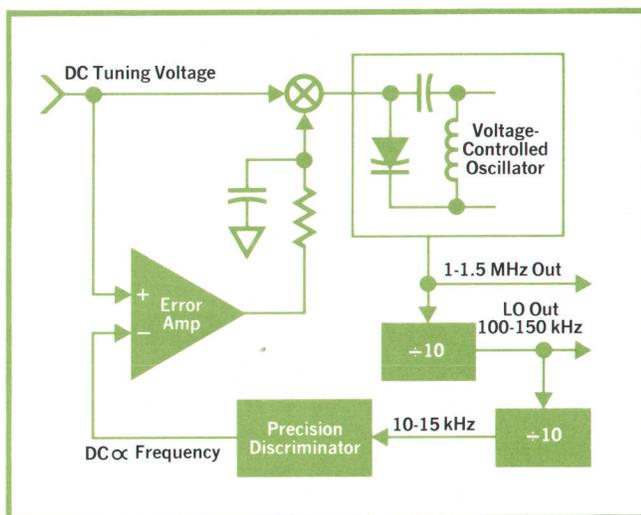


Fig. 7. A precision discriminator makes the oscillator frequency a linear function of the control voltage. The control voltage is a ramp that starts at a dc level determined by the FREQUENCY tuning control and by the START/CENTER switch, and that rises at a rate determined by the SWEEP TIME control to an amplitude determined by the FREQUENCY SPAN control (see Fig. 8).

as well as to improve the accuracy of frequency setting. As a result, tuning linearity of the oscillator is better than the 0.1% linearity of the 10-turn tuning potentiometer (overall linearity at the display is within 1%, see Fig. 5). The discriminator is described in the box on the page opposite.

The oscillator output is divided by 10 to get the needed 100–150 kHz LO signal. Short-term drift of the divided-down signal is less than 1 Hz and noise is 70 dB below the LO signal in a 1-Hz bandwidth 10 Hz from the LO frequency.

The undivided oscillator output is also supplied through a buffer amplifier to a rear-panel connector. This output makes it possible to measure the frequency with an electronic counter in one-tenth the gate time needed to get the same resolution when measuring the LO frequency directly.

A 0–50 kHz tracking oscillator signal is derived by heterodyning the 100–150 kHz LO signal with a 100-kHz crystal oscillator. The crystal is identical to those used in the selective filters and is supplied as part of the matched crystal set. The tracking oscillator output is thus close to the center of the instrument's tuning at all times.

The 100-kHz crystal oscillator also is used to derive the 10-kHz CAL signal output.

The Signal Path

A block diagram of the analog portion of the Model 3580A Spectrum Analyzer is shown in Fig. 8. Control of the amplitude of the input signal is distributed through several stages, as shown in Fig. 9. This allows the use of an input attenuator with 20-dB steps while allowing gain to be switched in 10-dB steps. The 20-dB attenuator has lower output impedance than a 10-dB attenuator and hence lower noise. Thus, the input impedance can be made 1 M Ω rather than the 100 k Ω commonly used.

The input amplifier uses a premium, low-noise FET. The circuit configuration and operating points chosen keep harmonic distortion products more than 95 dB below the signal level. A peak detector warns the user if the signal level exceeds the linear range of the amplifiers.

The mixer is a commercially-available monolithic, balanced modulator-demodulator that works over a 90-dB range.

The final IF amplifier stage is switched. A linear amplifier is used for measurement of signal amplitude on a linear scale, and a logarithmic amplifier is used for the wide-dynamic-range display on a dB scale. The logarithmic amplifier is a hybrid-IC type used in other HP signal-analyzing instruments.^{1,2}

The detector is an averaging type. This is followed by a smoothing filter that is switched along

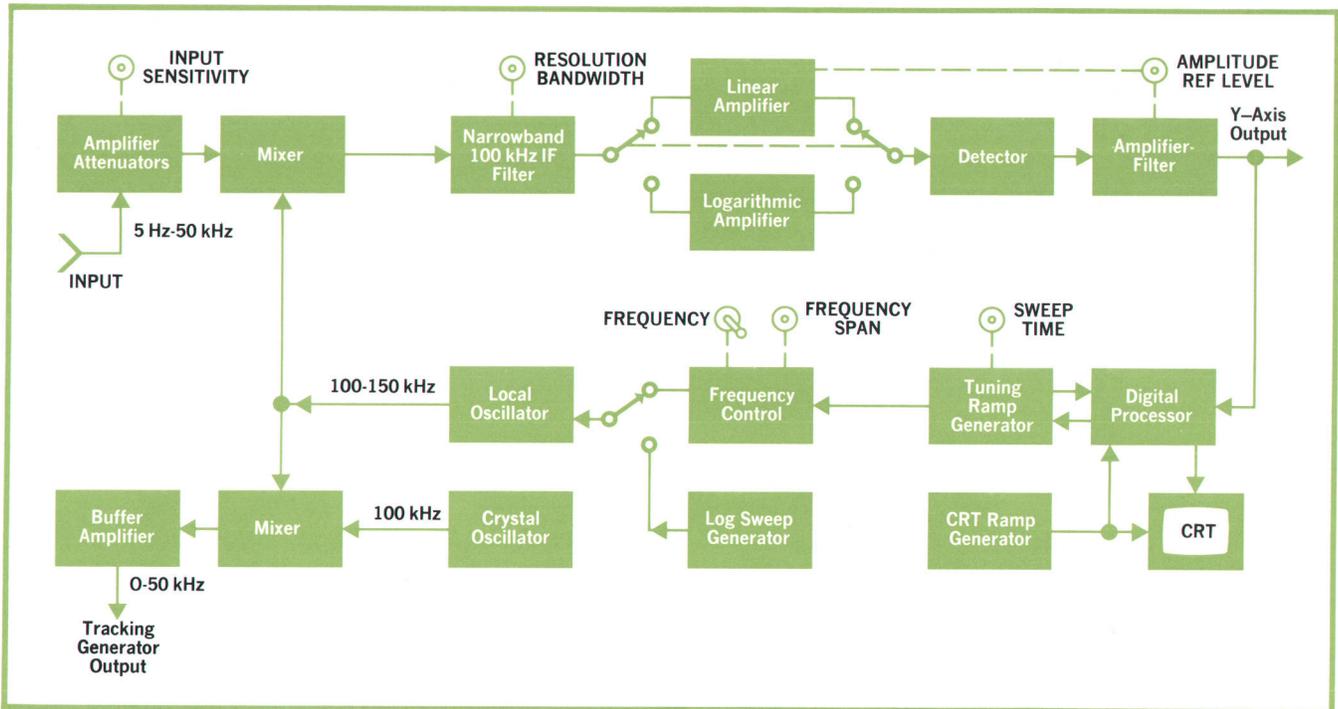


Fig. 8. Block diagram of Model 3580A Spectrum Analyzer. The basic scheme is similar to traditional wave/spectrum analyzers: heterodyne the input signal upwards with a tunable local oscillator signal, apply the resulting difference frequencies to a narrowband IF strip, and provide means for measuring the one that passes.

with the bandwidths. When measuring noisy signals, the degree of smoothing can be increased with a front-panel control in a tradeoff with response time.

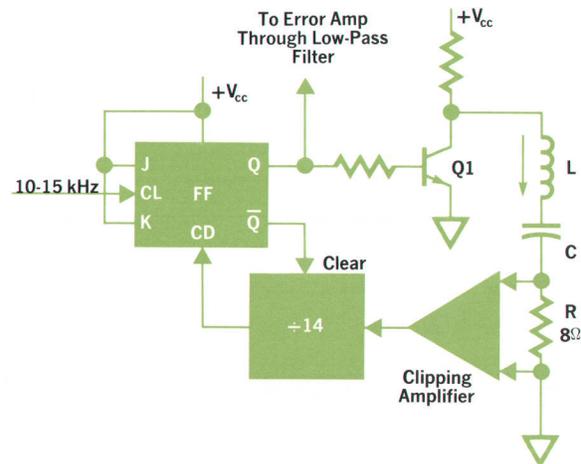
The output signal processor includes switchable gain stages and offset control. From there, the sig-

nal goes to the digital processing circuits. It also goes to the y-axis output for external recorders. This arrangement assures that the scale factors of the CRT display and recorder output always correspond, a convenience in setting up recorder plots. It should be noted, though, that the recorder outputs

Precision Discriminator

The discriminator in the Model 3580A's oscillator control loop is the type that obtains a dc control voltage by averaging the pulsed output of a monostable multivibrator. The multivibrator triggers at each zero crossing of the oscillator waveform.

Generating pulses with the degree of width stability needed to make the dc output a linear function of the multivibrator repetition rate is no job for the conventional one-shot. As shown in the drawing, the pulses are generated at the Q output of a J-K flip-flop. Between pulses, transistor Q1 is cut off and capacitor C charges toward the supply voltage V_{cc} . When an input trigger occurs, the transistor is turned on, providing a low-impedance path for the LC circuit, which then rings at about 280 kHz with the energy stored in C. The lightly-damped sinusoid developed across resistor R is squared up in the clipping amplifier to define zero crossings more accurately, and the resulting pulses are counted by a $\div 14$ counter, the output of which resets the J-K flip-flop. The width of the J-K pulse, which is about 50 μs wide, thus depends on the LC resonant frequency, which is relatively stable.



are tapped off ahead of the digital processor and thus do not supply signals stored in the digital memory.

Digital Processing

Contemporary practice uses shift registers to store waveforms digitally. Instead, the Model 3580A uses a random-access memory (RAM). This is for two reasons. For one, the RAM system can be implemented with CMOS circuits, resulting in a significant reduction in power dissipation. For another, adaptive sweep requires a brief reversal of the frequency sweep, something that would be more difficult to do with shift registers.

A block diagram of the digital processing system is shown in Fig. 10. Note that the system uses two voltage ramps—one that runs repetitively at a 50-Hz rate for the CRT display, and one that runs at a rate set by the front-panel frequency sweep controls for tuning the Analyzer (normalized to a 0-to-4.75 V span).

The RF oscillator in the CRT power supply serves as the clock. This operates somewhere between 55 and 70 kHz, a rate high enough to avoid interference with the measuring circuits.

The RAM is a 1024×8 array. Each of the 1024 addresses corresponds to a position on the CRT x-axis. Horizontal resolution this fine is desirable because of the abrupt changes in y-axis amplitude that occur in a spectrum analyzer display.

The 8 bits at each address represent the amplitude of the y-axis signal. This gives a resolution of 256 levels for the y-axis. The maximum amplitude error that can result from quantization into 256 levels is less than 0.4%, well within the accuracy

requirements of the instrument.

During each period of the Cl1 clock, there is a write phase and a read phase for the RAM. The address multiplexer selects an address corresponding to the frequency sweep position for the write phase and an address corresponding to the CRT sweep position for the read phase. This way, the CRT sweep operates independently of the frequency sweep and may thus operate at a constant rate.*

During the read phase, the data stored at that particular address is latched out to a digital-to-analog converter. The resulting analog voltage is passed through a line generator to the vertical deflection circuits of the CRT. The line generator³ traces a line from sample to sample so the displayed trace is a continuous line, rather than a series of dots (Fig. 11).

Since the CRT is swept at a constant rate, the read address can be obtained from a 10-bit counter that increments one count for each clock pulse. A reset pulse from the counter to the CRT ramp generator synchronizes the two at the start of each CRT sweep. This occurs during the last two clock periods so the display actually has 1022 x-axis segments.

The address for the write phase is obtained from the frequency tuning ramp by way of an analog-to-digital converter. Because of the relative slowness of the frequency sweep, several write phases occur at each x-axis address. For example, on the longest sweep (2000 seconds), the A-to-D converter dwells on each address for about 2 seconds during which time 100,000 write phases occur. Since the

* A side benefit of the constant display rate is that changing the front-panel sweep tuning controls has no effect on the CRT intensity or focus. Once the CRT controls are set for a satisfactory trace, there is no need to bother with them further.

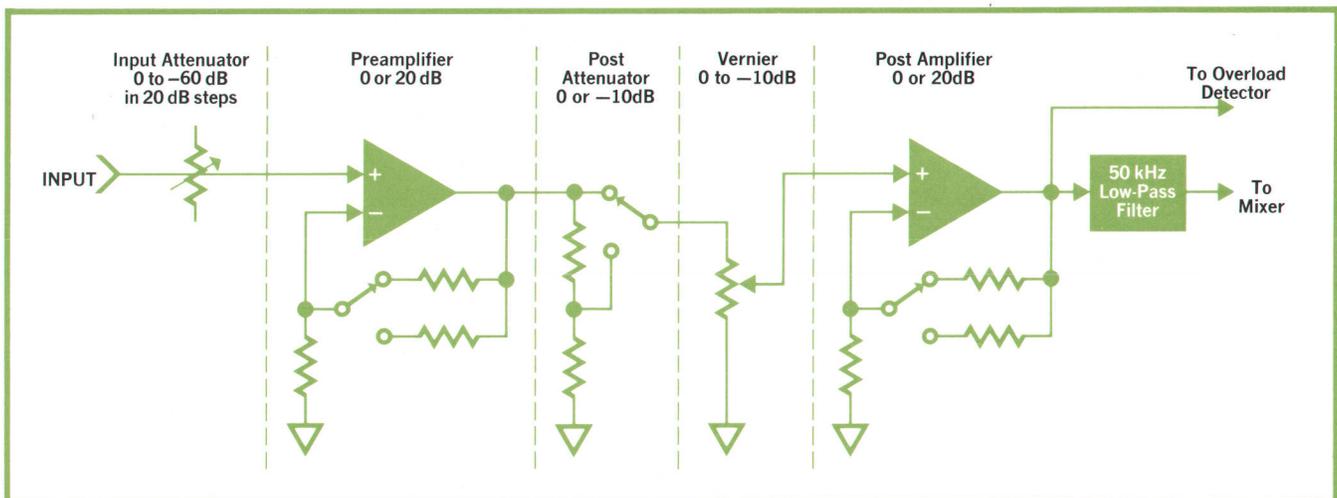


Fig. 9. Input circuits distribute the attenuation and gain functions to allow a wide range of signal amplitudes to be accommodated without introducing distortion or degrading the signal-to-noise ratio.

Permanent Waveform Storage

When the front-panel STORE button is pressed, the digital system looks at every pair of adjacent addresses, compares the numbers stored there, and stores the larger of the two in the even-numbered address. This is to assure that the peak value of every response is maintained in the stored trace. The odd-numbered addresses are then made available for new data.

The additional functions added to the basic digital system of Fig. 10 for doing this are shown in the diagram at right. This system is under control of a four-state logic machine.

When the STORE button is depressed, the four-state machine waits until a carry-out from the CRT address counter indicates completion of the CRT sweep presently in progress, then it goes to its second state. In this state, the address multiplexer is forced to stay switched to the address counter so read and write phases both occur at the same address.

An instant before each latch signal (C12) occurs, the input to the latch contains the data stored in the present memory address and the latch output has the data that was in the previous address. These two numbers are compared, by adding the present address to the complement of the previous address, and if the presently addressed value is the larger of the two, the write control is instructed to allow the output of the latch to be written into memory. When the write phase occurs, the data at the present address will have been transferred to the latch output and it will thus be written back into memory.

If the data at the present address were not the larger, the write control logic would inhibit the writing of data, and the memory contents would remain unchanged.

During the write phase, the least significant bit of the address is forced to zero, if it is not already zero. If the least significant bit is "one," and the data in that address is larger than that in the previous address, the present data will then be written into the previous address. If, on the other hand, the present address were even-numbered, and its data were the larger, the present data would simply be written over itself and no change would occur.

When all the addresses have been scanned in this mode, the four-state machine goes to its third state. Now the system reverts to normal operation with some exceptions. The least significant bit of the frequency sweep A-to-D converter is forced to one, so new information is loaded into odd-numbered addresses only. The CRT sweep address

counter is switched so that in one sweep only the odd-numbered addresses are read out and on the next sweep only the even-numbered addresses are read out. This is done by making the most significant bit of the address counter the least significant bit of the address. At the same time, the ramp speed is doubled.

When the STORE BLANK button is depressed, the least significant bit of the address counter is forced to "one," so only the contents of the odd-numbered addresses are displayed.

When the operator returns the instrument to the non-store mode, the four-state machine initially goes to its fourth state at the time the address counter indicates the end of a sweep. On the next sweep, the contents of the odd-numbered addresses are written into adjacent even-numbered addresses by using the multiplexers as before. At the end of this sweep the four-state machine goes back to its first state, returning the system to normal operation. Now, however, only the most recent spectrum is in memory with each pair of addresses having the same contents. If this step had not been taken, the line generator would have totally confused the display by drawing lines between the two spectra stored at alternate addresses.

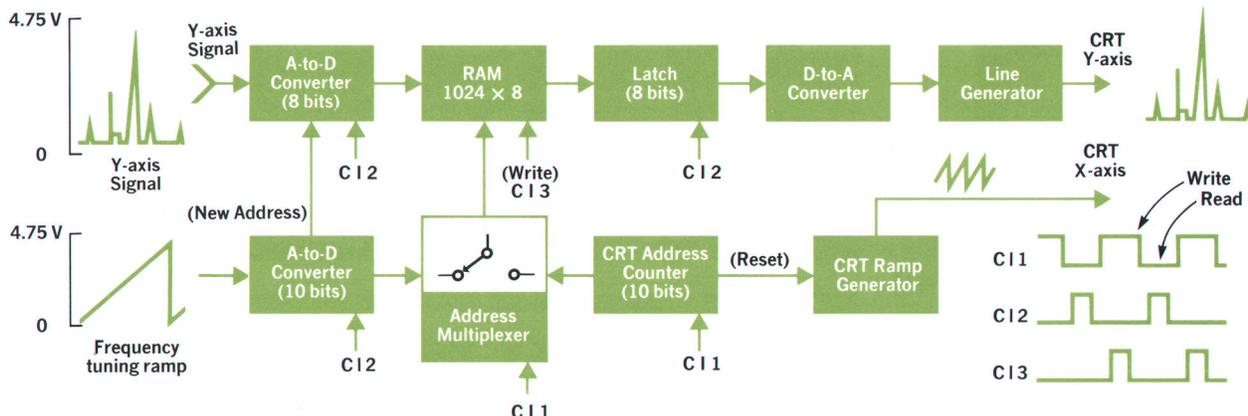
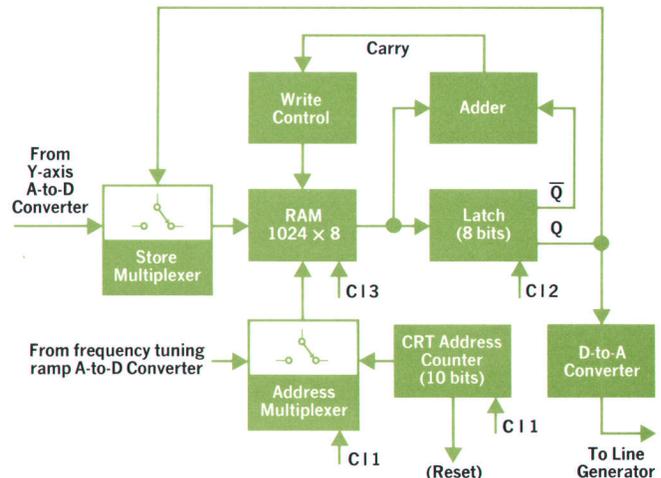


Fig. 10. Digital processing in the Model 3580A

Peak-Detecting Analog-to-Digital Converter

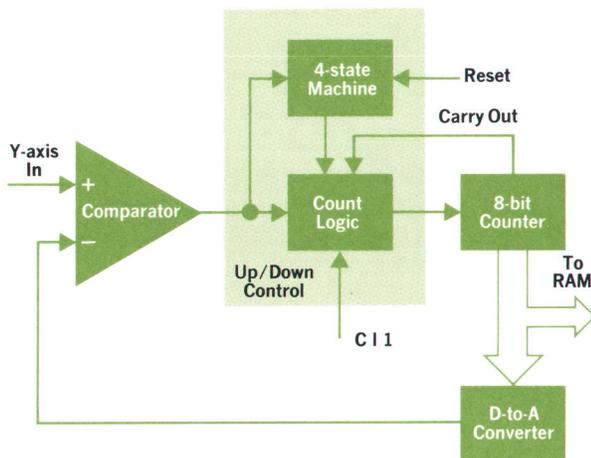
The peak value of the y-axis signal during any segment of the x-axis is selected by the A-to-D converter shown in the block diagram below.

The digital output is derived in a reversible counter. The y-axis input is compared to the number in the counter (converted to an analog value) to indicate whether the counter should count up on the next clock pulse to match the y-axis value, or count down.

At the start of a new x-axis segment, a reset pulse from the frequency sweep A-to-D converter resets the four-state logic machine, shown in more detail at right. Each of the decision boxes asks "Is the output of the y-axis comparator high?". In the case of a rapidly rising signal, where the counter would be trying to catch up to the y-axis, the decision would be "yes" on each clock pulse and on the third clock pulse, the machine would be in state D, where it locks up until reset.

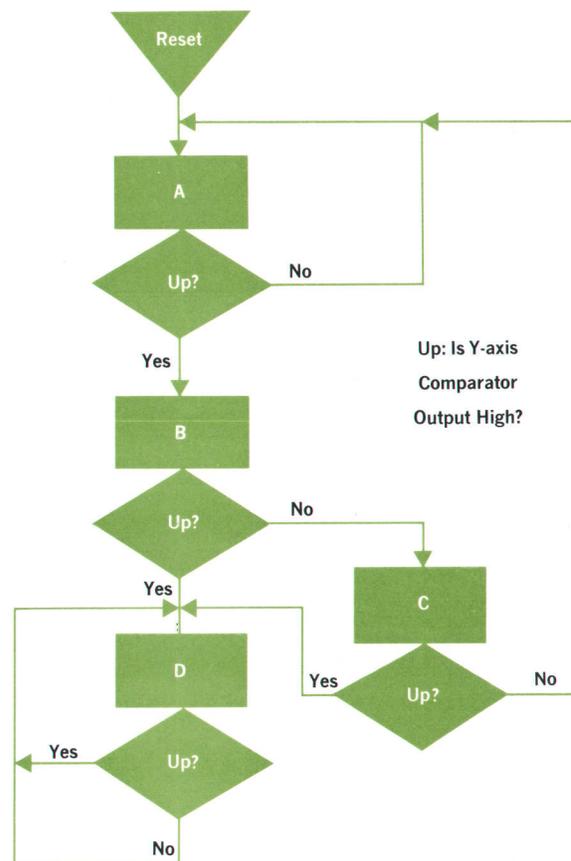
If the y-axis were changing slowly, more clock pulses would elapse before the machine arrives at state D, but once there it would lock up and remain there until reset.

Whenever the machine is in state D, the counter is inhibited from counting down. Thus, during the remainder of the x-axis segment, the counter can count up whenever the y-axis input exceeds the stored value, but it does not count down when the y-axis input is less. The value retained thus represents the peak value of the y-axis signal during that x-axis segment.



On the other hand, if in three clock pulses the y-axis signal moves down more than the analog value of one count, the machine never reaches state D. The value then retained for each x-axis segment is the value at the end of the segment. In this case, the A-to-D converter is not presenting the peak value, but is giving a truer depiction of a rapidly falling signal.

The carry-out signal from the counter is used as a limit stop. When the contents are all 1's, the carry-out inhibits the clock pulses if an up-count is indicated, and when the counter contents are all 0's, it inhibits the clock pulses if a down-count is indicated.



y-axis signal can change during this time, it is important that the value retained in storage as the sweep moves onto the next address is truly representative of the y-axis signal. Since the magnitude of a spectral component is the quantity of interest, the value retained is the peak value. The y-axis A-to-D converter selects the peak value, as described in the box above.

Peak selection also retains the peak value of noise. The switchable smoothing filter averages the

noise before the A-to-D conversion, so noise measurements can be made.

Storing Waveforms

When the front-panel STORE button is pressed, the spectrum on display is retained while a new spectrum is written and displayed. In this mode, one half the memory (512 words) is used for permanent storage and the other half remains receptive to the input.

Pressing BLANK STORE suppresses the display of the stored spectrum without erasing it. Pressing the button again to release it returns the stored waveform to the display. Pressing CLEAR WRITE loads all zeros into the active part of the memory without disturbing a stored spectrum, and returns the frequency sweep to the start frequency.

A description of how the stored-waveform system operates will be found in the box on page 9. The storage system is organized so even-numbered RAM addresses are used for the stored waveform and odd-numbered addresses are used for incoming waveforms. The CRT ramp then runs at twice its normal rate (10 ms per sweep), displaying the even-numbered addresses every other sweep and the odd-numbered addresses on alternate sweeps. This arrangement for alternating the displays allows the line generator to be used when two spectra are displayed.

Horizontal resolution is cut in half in the STORE mode. Since the reduced resolution is adequate for the majority of applications for this instrument, it was felt that the cost and the added power drain of an extra memory for retaining the 1000-segment resolution could not be justified.

Adaptive Sweep

As mentioned previously, adaptive sweep accelerates the tuning rate whenever the instrument's response is below a selected amplitude level. This requires more than a straightforward change of sweep speed, however. Because the sweep speed can be $20\times$ faster during the fast sweep, the sweep could already have tuned past the center of a spectral line before the narrowband filters could respond to a change in signal level.

The adaptive sweep system therefore steps back when a response is encountered, pauses a moment to let transients decay, then starts the slow sweep.

The amount of step back is determined by the IF bandwidth in use and the sweep width.

The slow sweep continues for a time that would normally carry it well beyond the point where the response again drops below the threshold level. Unless the slow sweep were extended, transient sidebands generated by the local oscillator while changing sweep speeds could institute a new step back if they combined with the decaying response to produce a signal higher than the threshold.

The shape of the adaptive sweep control signal is shown in Fig. 12.

The adjustable threshold level can be set anywhere between 0 and 60% of full screen. An analog dc voltage corresponding to the threshold level clamps the baseline of the CRT vertical deflection signal so the display baseline will appear at the selected threshold level.

Power Requirements

The instrument dissipates less than 15 watts, permitting battery operation (total dissipation when operating on ac line power is 35 watts). The CRT filament takes only one watt, and the 5-volt digital circuits were "stacked" for 10-volt operation (12-volt batteries are used) so current for the digital portion could be cut in half. As a result, forced-air cooling is not required. Neither are vent holes needed, thereby increasing the instrument's immunity to dusty environments.

Balanced Input

For those who wish to make measurements on balanced lines, an optional transformer-coupled input is available. CRT calibration is then switchable to be in dBm units with respect to 900Ω or 600Ω , or in volts. The transformer restricts the input to a frequency range of 300 Hz to 20 kHz and frequency

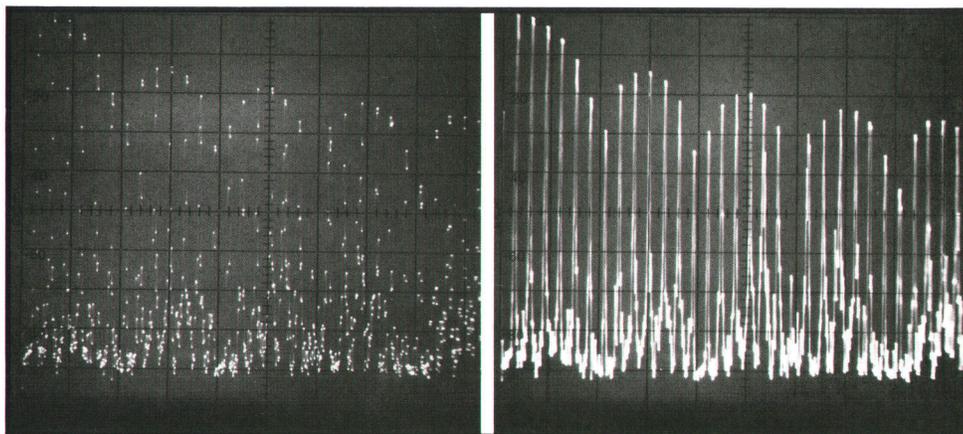


Fig. 11. Without connecting lines between dots, a digitized spectrum is difficult to interpret (left). The line generator clarifies the display (right).

Adapting a Sweep

The adaptive sweep system is controlled by the eight-state logic machine diagrammed at right.

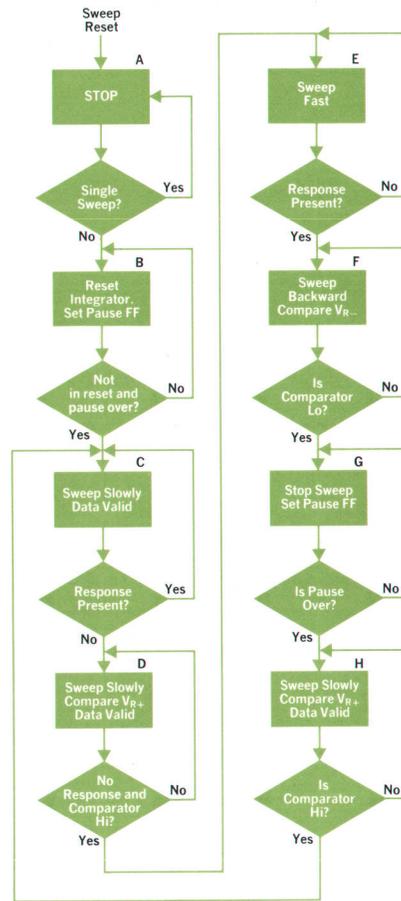
A comparator, shown at far right, determines the backward and forward steps. The ramp input to the comparator is grounded except when comparisons are being made. While the input is grounded, capacitor C charges so at the instant the switch opens, the ramp input starts from ground to track the frequency sweep ramp.

The other input to the comparator is a reference derived from the settings of the BANDWIDTH and SWEEP WIDTH controls. This determines how far the frequency sweep steps back when a response is encountered.

The sequence of events is most easily followed by starting with state E of the logic diagram. Here, the adaptive sweep controller is causing the frequency sweep to move at the fast rate, and it continues to do so as long as the y-axis input does not exceed the threshold (no response present).

When a response is detected, the controller moves on to state F, causing the instrument to sweep backwards while comparing the ramp voltage to the reference (V_{R-}). The backward sweep continues until the ramp falls back an amount equivalent to V_{R-} , which then causes the comparator output to go low. When this occurs, the controller goes to state G, stopping the sweep.

At the conclusion of the pause, the controller starts the slow sweep (state H) and it permits y-axis data to be processed. At the same time, the reference voltage is switched to a positive value (V_{R+}), equal in magnitude to $0.75 V_{R-}$. The slow sweep continues until the ramp input to the comparator rises above its starting level (ground) to a level equal to V_{R+} at which time the controller goes back to state C. If no response is present, the controller steps through state D, returning to the fast sweep in state E.



(Continued)

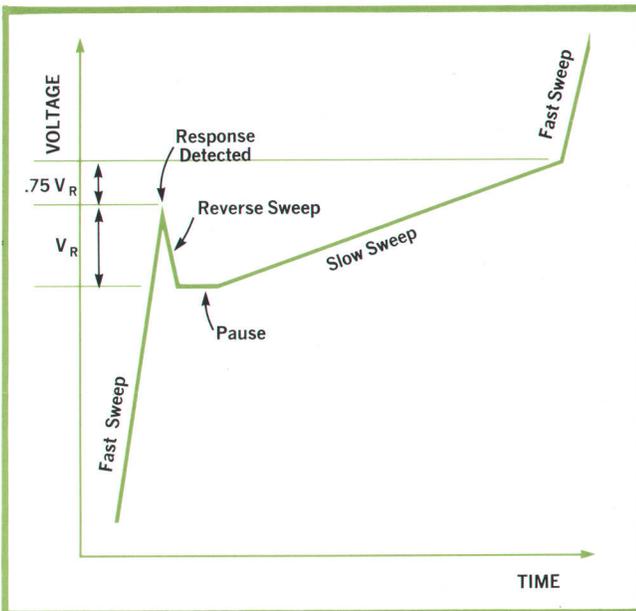


Fig. 12. Sweep tuning ramp steps back when a response is encountered, insuring that the instrument will tune slowly across the full range of the response. The amount of step back is determined by V_R , derived from the settings of the BANDWIDTH AND SWEEP TIME controls.

response over this range is ± 0.5 dB. A similar transformer is provided at the tracking oscillator output so the instrument can also be used to stimulate balanced lines.

Optimum Sweep Rate

A front-panel indicator guides the user in selecting suitable sweep parameters. If the indicator turns on, the instrument is sweeping too fast for the bandwidth in use.

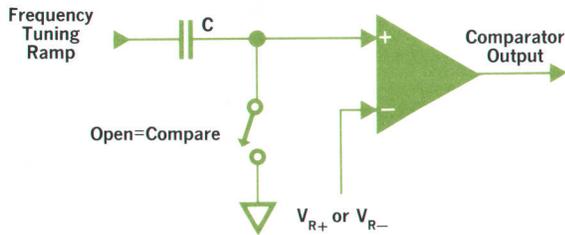
It can be shown that for less than 2% amplitude error, the sweep rate should be no greater than half the square of the bandwidth. Sweep rate depends on the settings of sweep time and frequency span. Thus,

$$\text{Sweep rate} = \frac{\text{Hz/div}}{\text{Time/div}} \leq \frac{1}{2} \text{BW}^2$$

or,

$$\log \left(\frac{1}{2} \text{BW}^2 \right) + \log \left(\frac{\text{Time}}{\text{div}} \right) - \log \left(\frac{\text{Hz}}{\text{div}} \right) > 0$$

This equation is implemented by using weighting resistors at each switch position to supply voltages to a summing amplifier. If the inequality is not met, the amplifier turns on the LED indicator.



Had a response been present at the time the controller moved back to state C, the controller would have stayed there, maintaining the slow sweep until a response is no longer present.

Note that at the start of a sweep (states A and B) the controller pauses, then starts the sweep at the slow rate just in case a response might exist in that part of the spectrum. If no response is present, it quickly moves through states C and D to the fast sweep.

Acknowledgments

Product design was by Jim Saar, with assistance from industrial designer Barry Mathis. Electrical design was by Larry Whatley (project leader), Bill Spaulding, and the authors.

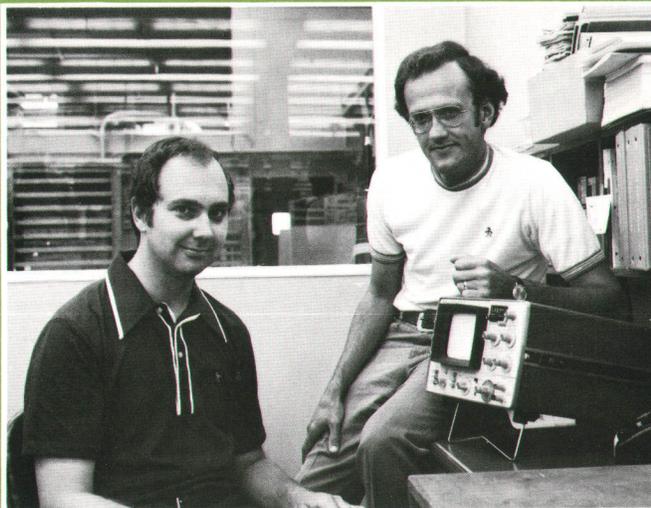
References

1. Hugo Vifian, Frank K. David, and Wayne L. Frederick,

'A "Voltmeter" for the Microwave Engineer,' Hewlett-Packard Journal, November 1972.

2. Gerald E. Nelson, Paul L. Thomas, and Robert L. Atchley, 'Faster Gain-Phase Measurements with New Automatic 50 Hz-to-13 MHz Network Analyzers,' Hewlett-Packard Journal, October 1972.

3. Jerry Weibel and Larry Whatley, 'Digital Storage Improves and Simplifies Analysis of Low-Frequency Signals,' Electronics, June 21, 1973.



Gerald E. Weibel (Right)

Jerry Weibel earned a BSEE degree at the University of Wisconsin in 1967 and then joined HP, going to work on the 746 High Voltage Amplifier for the 745A AC Calibrator. He has since earned his MSEE degree at Colorado State University in the HP Honors program. Jerry worked on the power supplies and the digital controller for the Model 3330A Automatic Synthesizer before tackling the digital system in the 3580A Spectrum Analyzer. Married, but with no offspring so far, he likes golfing and hiking.

William L. Hale (Left)

Bill Hale joined Hewlett-Packard in 1965 upon getting his MSEE degree from Ohio State University. Initially he worked on some digital voltmeter investigative projects then contributed to the 745A AC Calibrator and 746A High Voltage Amplifier before going on to the 3580A Spectrum Analyzer. In his spare time he and his wife like to explore the western hemisphere in his Beech Bonanza.

SPECIFICATIONS HP Model 3580A Spectrum Analyzer

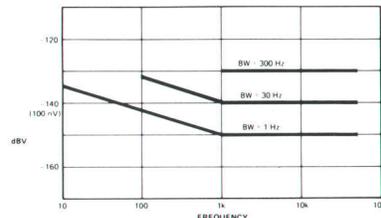
Frequency Characteristics

FREQUENCY RANGE: 5 Hz to 50 kHz.
FREQUENCY DIAL ACCURACY: ± 100 Hz, 20–30°C; ± 300 Hz, 0–55°C.
FREQUENCY STABILITY: typically ± 10 Hz/hr after 1-hr warm-up; ± 5 Hz/°C.
BANDWIDTHS: 1 Hz (20–30°C), 3, 10, 30, 100, 300 Hz.
SHAPE FACTOR: –60 dB/–3 dB ratio: 10 on all bandwidths except 300 Hz, where it is 8.
SWEEP FREQUENCY SPAN: 50 Hz to 50 kHz in 1, 2.5 sequence.
START/CRT switch places frequency selected by tuning dial at left edge or at center of CRT display.
LINEARITY: $\pm 1\%$.
ZERO SCAN: Allows continuous observation of signal amplitude, within selected bandwidth centered at frequency tuning.
LOG SWEEP: one range, 20 Hz to 43 kHz, swept in 5 seconds (SWEEP TIME and FREQ SPAN controls are disabled).
SWEEP TIMES: 0.1 to 2000 seconds.
SWEEP MODE: Repetitive, single, manual, and log. At end of sweep in single mode, frequency rests at upper end of sweep. RESET button or contact closure through rear-panel connector returns sweep to start frequency and starts new sweep. In manual mode, full turn of concentric knob duplicates selected electronic sweep.
ADAPTIVE SWEEP: When response is below threshold level, sweep speed is 20 to 25 times faster than selected sweep. Responses greater than 6 dB above threshold are swept at selected rate. Threshold is adjustable over lower 60% of display.

Amplitude Characteristics

AMPLITUDE RANGE (Input Sensitivity):
LINEAR MODE: 18 full-scale ranges from 0.1 μ V full scale to 20 V full scale. Top CRT graticule line is calibrated to voltage selected; bottom line is 0 volts. Graticule has 10 divisions.
LOG MODE: Top CRT graticule line is calibrated to –140 dB to +30 dB, in eighteen 10-dB steps. Graticule calibration is 10 dB/div. Any 10-dB segment can be expanded to fill screen, using AMPLITUDE REFERENCE as offset control.
LOG MODE CALIBRATION: Switch selected; dBV or, with external 600 Ω input termination, dBm.
LOG MODE DYNAMIC RANGE: >80 dB.

AMPLITUDE ACCURACY:		Log	Linear
FREQUENCY RESPONSE:			
20 Hz–20 kHz		± 3 dB	$\pm 3\%$
5 Hz–50 kHz		± 5 dB	$\pm 5\%$
SWITCHING BETWEEN BANDWIDTHS (25°C):			
3 Hz–300 Hz		± 5 dB	$\pm 5\%$
1 Hz–300 Hz		± 1 dB	$\pm 10\%$
AMPLITUDE DISPLAY:		± 2 dB	$\pm 2\%$
INPUT ATTENUATOR:		± 3 dB	$\pm 3\%$
AMPLITUDE REFERENCE LEVEL: (IF ATTENUATOR)			
Most Sensitive Range:		± 1 dB	$\pm 10\%$
All Other Ranges:		± 1 dB	$\pm 3\%$
MAXIMUM INPUT LEVEL: 100 V rms or ± 100 Vdc.			
INPUT IMPEDANCE: 1 M Ω $\pm 100\%$ shunted by 30 pF.			
NOISE LEVEL:			



NOISE SIDEBANDS: In 1-Hz bandwidth 10 Hz away from CW signal, >70 dB below signal.
SPURIOUS RESPONSES: >80 dB below input reference level.
LINE-RELATED SPURIOUS: >80 dB below input reference level or –140 dBV (0.1 μ V).

IF FEEDTHROUGH: For input levels >10 V, –60 dB; <10 V, –70 dB.

ZERO-BEAT RESPONSE: >30 dB below reference level.

Outputs

X-Y RECORDER OUTPUTS:

Y-AXIS: 0 to +5 V, $\pm 2.5\%$
 X-AXIS: 0 to +5 V, $\pm 2.5\%$
 IMPEDANCE: 1 k Ω

PEN LIFT: Contact closure to ground during sweep.

TRACKING GENERATOR OUTPUT:

AMPLITUDE: 0 to 2 V rms with rear-panel control.
 IMPEDANCE: 600 Ω
 ACCURACY: Within ± 2.5 Hz of bandwidth center.
 FREQUENCY RESPONSE: $\pm 3\%$, 5 Hz to 50 kHz
 TOTAL HARMONIC AND SPURIOUS CONTENTS: 40 dB below 1-volt output level.

L.O. OUTPUT: 100 mV; tunes from 1 to 1.5 MHz as instrument is tuned from 0 to 50 kHz.

General

POWER: 100, 120, 220, 240 V, +5%, –10%; 48–440 Hz; 35 W max.
DIMENSIONS: 11 $\frac{1}{4}$ in W x 8 in H x 18 in D (28.6 x 20.3 x 45.8 cm).
WEIGHT: 27 lb. (12 $\frac{1}{4}$ kg), 35 lb. (15 $\frac{3}{4}$ kg) with optional battery and front cover.

OPTIONS:

001: Battery and splash-proof front-panel cover. Battery operates instrument 5 hours, recharges in 16 hours. Automatic turn-off prevents deep discharge. Useful life is over 100 cycles.
 002: Balanced and floating input and tracking generator output.
INPUT CALIBRATION: dBm 900 Ω , dBm 600 Ω , Volts.
INPUT CMR: >70 dB at 60 Hz
OUTPUT IMPEDANCE: 600 Ω
FREQUENCY RESPONSE: ± 0.5 dB, 300 Hz–20 kHz (Input and output have same response).

PRICES IN U.S.A.: 3580A, \$3800, Option 001, \$255, Option 002, \$60.
MANUFACTURING DIVISION: LOVELAND DIVISION
 815 Fourteenth Street, S.W.
 Loveland, Colorado 80537

A High-Performance Beam Tube for Cesium Beam Frequency Standards

The benefits are significant improvements in accuracy, short-term stability, settability, and sensitivity to external dc magnetic fields.

By Ronald C. Hyatt, Louis F. Mueller and Terry N. Osterdock

THE CESIUM BEAM frequency standard is the most accurate commercially produced instrument in the world. Its output frequency conforms to international standards within parts in 10^{12} . It is a primary frequency standard, which means that any cesium beam standard can be aligned independently and will produce the accepted standard frequency within parts in 10^{12} . Most nations' frequency and time standards are based on commercial or specially built cesium beam standards. Navigation systems (Loran C, Omega, Apollo), communications systems, calibration laboratories, and a variety of scientific endeavors also depend upon the precise time and frequency information provided by these instruments.



Fig. 1. HP 5061A Cesium Beam Frequency Standard is now available with standard or new high-performance cesium beam tubes. The high-performance tube's specifications guarantee 30% better accuracy, ten times better short-term stability, seven times better settability, and a factor of ten reduction in sensitivity to external dc magnetic fields.

Hewlett-Packard produced its first cesium standard, Model 5060A, in 1964. The current version, Model 5061A (Fig. 1), has an accuracy specification of $\pm 1 \times 10^{-11}$, twice as good as that of the 5060A, but it still uses the same basic cesium beam tube design.

Now an entirely new high-performance beam tube has been developed (Fig. 2). Designated 5061A Option 004, the new tube's specifications guarantee 30% better accuracy, ten times better short-term stability, seven times better settability, and a factor of ten reduction in sensitivity to external dc magnetic fields.

The 30% increase in accuracy, from $\pm 1 \times 10^{-11}$ to $\pm 7 \times 10^{-12}$, is the result of increased microwave

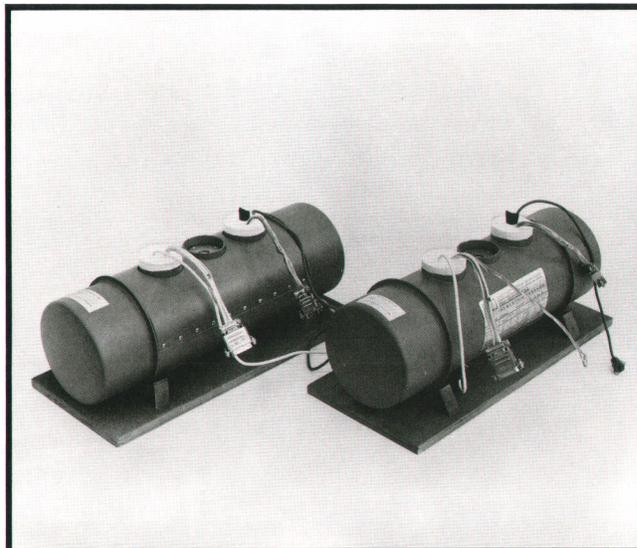


Fig. 2. High-performance beam tube (right) is mechanically and electrically compatible with the standard tube (left). Retrofit kits will be available for upgrading existing 5060A and 5061A standards.

cavity length, a reduction in cavity phase shift by improved cavity design and manufacturing techniques, and an improvement in C field homogeneity. The improved accuracy can be significant, for example, where a remote site needs to be synchronized with a standard to microsecond accuracy. The frequency of "flying clock" visits to the remote site can be considerably reduced if the site has one of the new beam tubes.

The factor of ten improvement in short-term stability for averaging times greater than 0.1 second makes it possible to compare two standards using the new tubes to a one-sigma accuracy of $\pm 1 \times 10^{-12}$ in 100 seconds. With standard beam tubes, it would take two hours to get the same accuracy. The better stability of the new tube results largely from increased cesium beam flux, which improves the signal-to-noise ratio of the output.

When a cesium beam standard needs to be adjusted to exactly the same frequency as a reference standard, the new beam tube makes this possible to within $\pm 1 \times 10^{-13}$, a considerable improvement in settability over the $\pm 7 \times 10^{-13}$ of the standard tube. A finer C field control and a built-in degaussing coil contribute to the improvement.

Better magnetic shielding in the new tube reduces frequency changes caused by a 2 gauss external magnetic field to less than 2 parts in 10^{13} . This is significant in flying clock applications or in laboratories where equipment near the standard is being moved. Another feature, important in portable applications, is the dual beam design of the new tube, which reduces sensitivity to acceleration and vibration.

The new tube is somewhat heavier than previous beam tubes, and considerable attention was given to designing components to meet the shock and vibration requirements often specified for portable equipment.

The improved performance of the new beam tube has been achieved without changing external dimensions, power supplies, or connectors. Retrofit kits will soon be available for present users of 5060A or 5061A Cesium Beam Standards who want to replace their standard tubes with high-performance tubes.

Cesium Beam Frequency Standard Operation

To understand the design of the new high-performance beam tube, it's helpful to know how the tube is used in a cesium beam frequency standard.

Fig. 3 is the basic block diagram of a cesium beam standard. The output of the standard is derived from a 5 MHz voltage controlled crystal oscillator (VCXO) whose frequency is locked to a resonance of the cesium atom corresponding to a change in atomic state.

The output of the VCXO is put into two channels. In one channel the 5 MHz signal is amplified and then synthesized to produce 12.6317715959 MHz. The second channel phase modulates the 5 MHz signal at 137 Hz and then multiplies by 18 to produce 90 MHz. Both signals are sent to the harmonic generator, where the 90 MHz is multiplied by 102 to produce 9180 MHz. This is mixed with the 12.6317715959 MHz signal to produce the resonant frequency of the cesium atom (in a magnetic field of approximately 0.06 gauss), 9.1926317715959 GHz.

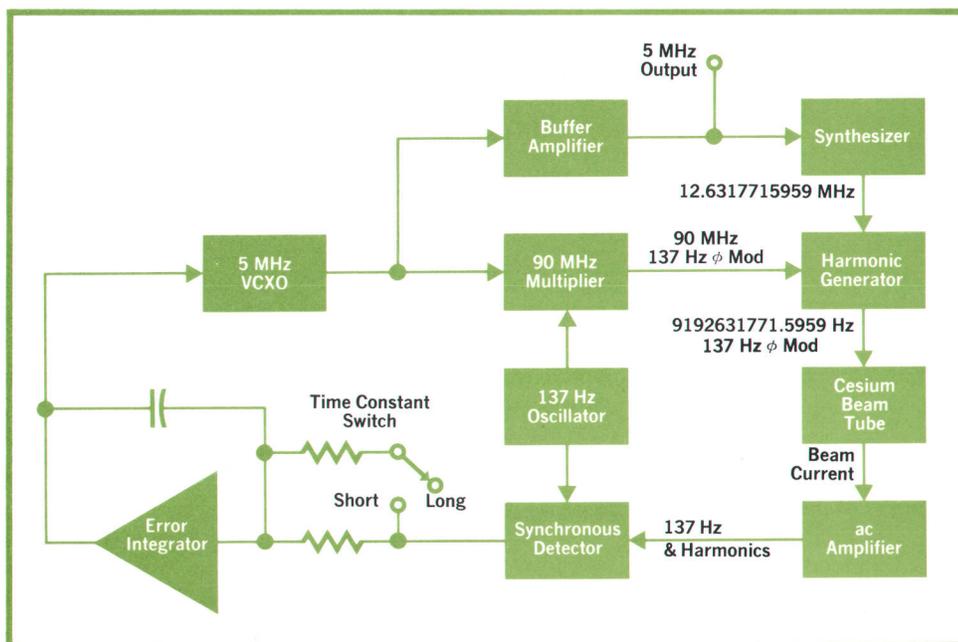


Fig. 3. Output of the cesium beam frequency standard comes from a 5 MHz voltage controlled crystal oscillator (VCXO). The VCXO frequency is locked to an invariant resonance of the cesium atom undergoing a change of atomic state in the beam tube.

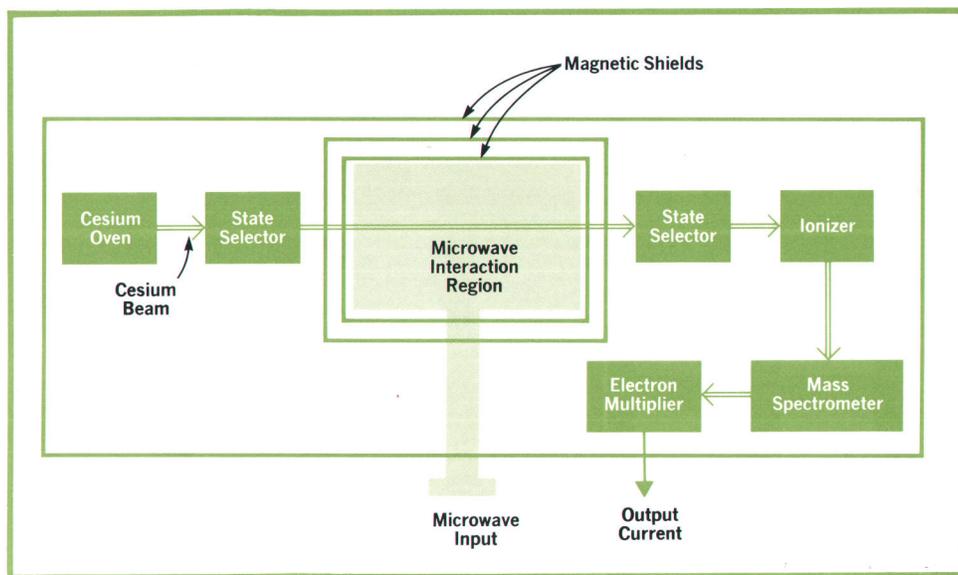


Fig. 4. Major components of a cesium beam tube. The cesium atoms absorb microwave energy and change state in the interaction region. The output, a current proportional to the number of atoms that change state, is sharply dependent on the microwave frequency.

The interaction of an oscillating magnetic field at this frequency with the cesium atoms in the beam tube produces the appropriate atomic state changes and provides an output current consisting of an average dc value and ac components at 137 Hz and harmonics.

If the VCXO frequency is exactly 5 MHz the beam tube input frequency is exactly 9.1926317715959 GHz and the output current is modulated at even harmonics only, primarily the second harmonic, 274 Hz. If the VCXO frequency differs from 5 MHz the output current modulation will have a 137 Hz component whose amplitude indicates the difference between the beam tube input frequency and the cesium resonance frequency, and whose phase indicates the sign of the difference.

The output current is amplified and synchronously detected using as a reference the oscillator that generates the 137 Hz modulation. The synchronous detector and an integrator produce an error signal that is used to adjust the 5 MHz VCXO. Thus, the VCXO is controlled through the servo loop containing the beam tube to produce a very stable 5 MHz output signal.

High-Performance Beam Tube

The new high-performance cesium beam tube is a dual beam passive atomic resonator that has a Q of approximately 2.5×10^7 and an output signal-to-noise ratio of approximately 4500. All the non-electronic components necessary for cesium beam tube operation are contained inside the beam tube vacuum envelope, which is also a structural member and provides the means of mounting the beam tube in the instrument.

The functional components inside the tube (see

Fig. 4) are: the cesium oven, which produces two ribbon-like beams of Cs atoms; the state selector magnets, in which a magnetic field gradient separates the cesium atoms into two energy groups; the microwave interaction region, in which the change of atomic energy state occurs; the hot-wire detector, which ionizes the cesium atoms that reach it by virtue of having changed state; the mass spectrometer, which spatially separates cesium ions from noise-producing ions such as potassium; the electron multiplier, which amplifies the 10^{-12} ampere signal to a level where its amplified shot noise is larger than the noise levels in the subsequent electronic circuits; an internal getter ion pump, used to maintain a high vacuum inside the tube; getters for collecting the expended cesium; and a set of magnetic shields, which surround the interaction region with a minimum of three layers of magnetic shielding.

There are also coils inside the magnetic shields. One is used to degauss the shields, and a set of coils generates the uniform magnetic field (C field) necessary for operation of the beam tube.

The beam-optic geometry used in the tube was computer optimized to provide the best combination of narrow line width and high signal-to-noise ratio. The main reason for using dual beams is to minimize variations in the beam current caused by acceleration or rotation. The two beams also give twice as much signal, resulting in a 41% improvement in the signal-to-noise ratio.

Cesium Oven

The oven is the source of cesium for the beams and must perform several functions. First, it must form and aim the beams, and for this function a

dual multitube collimator is used. Second, it must provide a means of controlling the temperature of the cesium so the desired beam intensity can be achieved.

Third, the cesium, a metal that melts at 29°C , must be contained in such a manner that liquid cesium cannot escape from the oven. To accomplish this, the oven uses a stainless steel "sponge" and an anti-spill maze, followed by a conductance limiter to assure that even if a droplet of liquid cesium were ever to reach the chamber behind the collimator, it would evaporate out through the collimator faster than it could be replenished from the reservoir.

Finally, the oven must completely contain the cesium supply during the tube bake (400°C) so that no cesium can escape the reservoir prior to the completion of the beam-tube baking and degassing procedures. To accomplish this, the cesium is contained in an ampoule that is closed by a thin stainless steel rupture disc. After processing, this disc is opened by discharging a capacitor through it, thereby melting a hole that allows the cesium to escape from the reservoir. The cesium is cleaned and carefully degassed prior to being sealed into the ampoule both

to prevent bursting the rupture disc during bakeout and to keep foreign gas from being liberated along with the cesium during normal operation.

State Selector Magnets

The state selector magnets use the variation of cesium-atom ground-state energy levels with magnetic field, illustrated in Fig. 5. At zero field there are only two energy levels for neutral cesium atoms. However, the application of a small magnetic field, 0.06 gauss for example, causes these two levels to split into a total of 16 hyperfine levels, the higher-energy ($F=4$) group into nine and the lower energy ($F=3$) group into seven. At higher magnetic fields, say 10,000 gauss, there are also two distinct groups, but each has eight sub-levels. The $4, -4$ sub-level is now grouped with the $F=3$ sub-levels.

The state selector magnets produce a strongly inhomogeneous magnetic field with a maximum value of approximately 10,000 gauss. Since the energy is a function of the field and the field is a function of position the atoms experience a force. Those in the $F=3$ sub-levels and the $F=4, m_F=-4$ sub-level are deflected to the region of higher magnetic field, while those in the remaining $F=4$ sub-levels are deflected to the region of lower field. Thus the magnets spatially separate these two groups of atoms.

Fig. 6 shows a representation of the dual beam tube in which the solid lines indicate the path of the atoms that contribute to the signal, and the dashed lines the path of the atoms that do not. Getters are used to capture the cesium atoms in the unwanted paths to prevent their diffusing through the tube to where they might be detected, thus decreasing the signal-to-noise ratio.

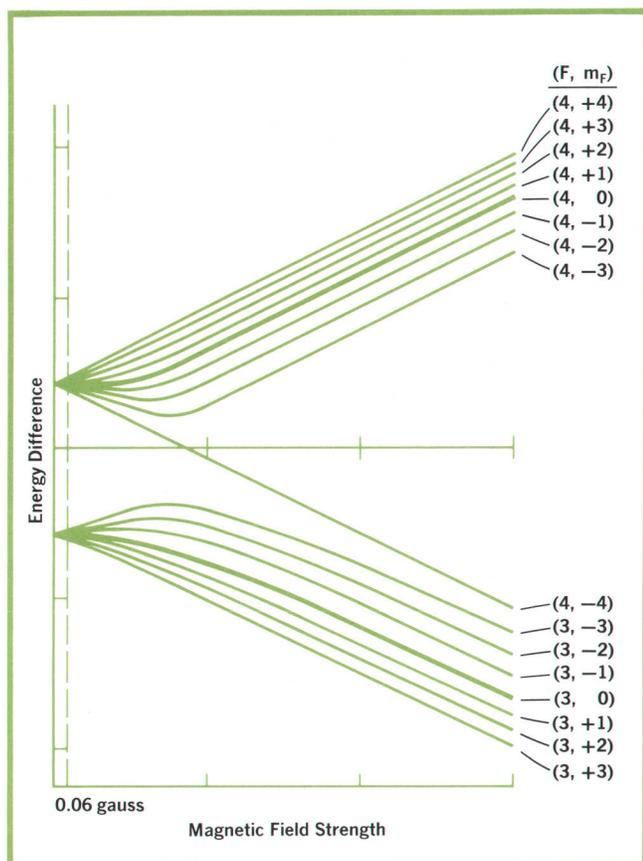


Fig. 5. Energy levels of cesium are functions of magnetic field. For frequency control, the cesium atoms are induced to change from the $3,0$ state to the $4,0$ state.

Resonance Transition for Frequency Control

For the control of frequency, the cesium atom is required to perform a resonance absorption of energy from the microwave exciting signal, corresponding to a transition from the $3,0$ state to the $4,0$ state.* The frequency of this transition is defined for unperturbed atoms at rest as 9,192,631,770 Hz.

In any practical device, however, this frequency is changed slightly by several effects. Two that are unavoidable result from the small applied magnetic field and the relativistic time dilation that occurs because of the velocity of the atoms in the beam. The correction for velocity is the transverse or second-order Doppler correction and is approximately one part in 10^{13} . The dependence of the frequency on magnetic field is given by $f = 9192631770 + 8.7026 \times 10^{-10} f_z^2$, where f and f_z are in hertz, f_z ,

*The difference in energy (ΔE) between the two states is related to the transition frequency f by $\Delta E = hf$, where h is Planck's constant.

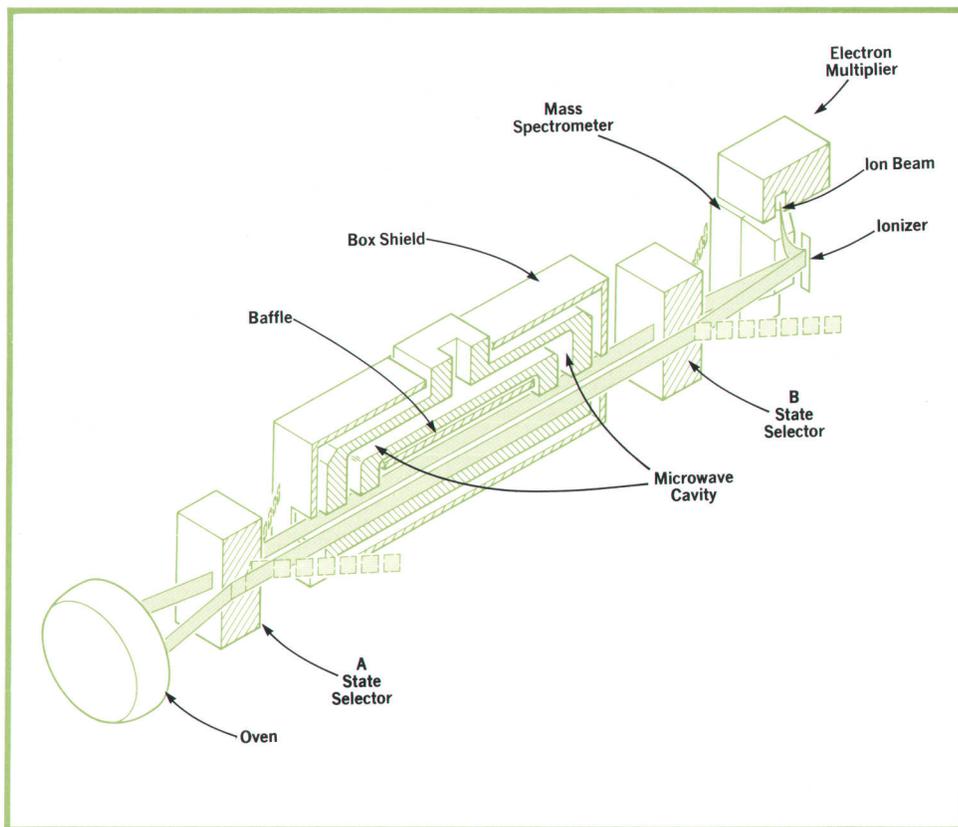


Fig. 6. Cross-sectional representation of the new high-performance cesium beam tube. Dual cesium beams give higher beam flux, which improves the signal-to-noise ratio and short-term stability. A longer interaction region sharpens the resonance curve and improves accuracy. A new box-shaped C field shield and an improved microwave cavity also contribute to accuracy.

the Zeeman frequency, is the average difference in frequency between the 3,0 to 4,0 transition and the adjacent field-dependent transitions. f_z has a nearly linear dependence on magnetic field. In HP standards $f_z = 42,823$ Hz. When all corrections are applied the frequency of the 3,0 to 4,0 transition in HP cesium beam standards is 9,192,631,771.5959 Hz.

For small fields, the frequency of the 3,0 to 4,0 transition has only a quadratic dependence on magnetic field, which means that the slope is zero at zero field. The consequent lower sensitivity to small magnetic fields is the primary reason for choosing the 3,0 to 4,0 transition, which is also known as the "field-independent transition" and the "clock transition."

The transition occurs inside a shielded structure containing a microwave cavity. The cavity has two arms. Cesium atoms that pass the first state selector magnet, or A magnet, enter the shielded structure, interact with the microwave energy in the first arm of the cavity, drift for a time, interact again with the microwave energy in the second arm of the cavity, and then leave the shielded structure. The closer the microwave frequency is to the correct frequency, the more atoms will make the desired transition.

C Field

Cesium-atom transitions are possible between any

pair of sub-levels shown in Fig. 5 for which the second quantum number either does not change or changes by ± 1 ($\Delta m_F = 0$ or ± 1). It would be desirable to operate in a field sufficiently small that all the transitions behaved as a single transition. This would give about seven times as much signal but would require that the field be kept lower than about 10^{-8} gauss to avoid errors greater than 1 part in 10^{12} caused by line asymmetry. Fields this low cannot be obtained easily, so this approach is not used. Instead, a small field of about 0.06 gauss is applied and this separates the $\Delta m_F = 0$ transitions by about 40 kHz in frequency so that any unbalanced residual overlap causes less than 1 part in 10^{13} error. This field is called the C field.

Since the desired transition does have a small quadratic dependence on magnetic field the C field must be stable with time and uniform throughout the region between the arms of the microwave cavity. Sudden changes in field direction must be avoided everywhere since these might induce low-frequency transitions in the atoms that would mask the desired microwave transition. This effect becomes more severe as the C field is reduced.

The microwave transition overlap and the low-frequency transition problems push the choice of C field in the high field direction. The quadratic dependence of the desired transition on C field and the consequent requirements on field stability and

homogeneity push the choice in the low field direction. Approximately 0.06 gauss is a fairly good compromise. The exact value is chosen to simplify frequency synthesis of the microwave signal.

The C field and shield design of the new high-performance beam tube are a significant departure from previous commercial practice. The magnetic structure of the standard tube has a U-shaped cross section. A coil is wound around the bottom of the U to generate the C field. There are rather severe edge effects, both at the ends and at the open edge of the U. This structure is surrounded by a box that has holes to let the beam through and the microwave power in. A cylindrical shield around the center of the tube completes the shielding on the standard beam tube.

In the new beam tube a patented box-shaped-cross-section C field shield is used. The field is generated by a coil wound on the outside of the microwave cavity structure. The cavity provides support and positioning for the wires, which run the length of the cavity just inside the top and bottom of the box. Thus the winding forms a close-fitting solenoid inside the box. The hole where the microwave power is admitted is magnetically isolated by a baffle parallel to the top and bottom of the box. A winding on the baffle excites it so that it shunts no flux and therefore does not upset the magnetic field generated by the main C field winding.

This box-shaped shield with its smaller openings has smaller edge effects and thus more useable space, permitting a longer drift region between the cavity ends.

A second shield surrounds this C field shield, and one more shield completely surrounds the entire internal structure including oven, detector, state selectors, electron multiplier and getter ion pump, providing increased shielding to external fields and

Field Direction	Fractional Frequency Change from Zero Field
Top	4.2×10^{-14}
Bottom	-3.8×10^{-14}
Right Side	-7.5×10^{-14}
Left Side	1×10^{-13}
Front	4×10^{-15}
Rear	2×10^{-15}

Fig. 7. Improved magnetic shield design reduces sensitivity to external fields. Shown here is measured magnetic susceptibility of a typical high-performance tube for all six orientations of a dc magnetic field equal to 2 gauss. Specified maximum fractional frequency change is $\pm 2 \times 10^{-13}$

increased immunity to magnetic materials outside the tube. The multiple shields also reduce external magnetic fields produced by the permanent magnet components in the tube.

The improvements in the magnetic shielding of the beam tube markedly reduce sensitivity to external dc magnetic fields. The magnetic susceptibility is shown in Fig. 7 for all six orientations of a dc magnetic field equal to 2 gauss. Frequency changes are less than $\pm 2 \times 10^{-13}$ for all orientations.

Microwave Cavity

The two-armed microwave cavity is a machined unit made from two identical halves fastened together. Machined cavities are used because this method provides the required symmetry tolerances with better reproducibility and less cost than the previous method of bending and brazing monel waveguide. The new cavity is made to symmetry tolerances of 0.002 inch between the two arms. For reasons of electrical conductivity, high-conductivity oxygen-free copper is used.

The field uniformity desired for the C field requires that the sides of the shield be held parallel within 0.005 inch. The machined cavity matches these tolerances nicely.

Detecting the Transition

In the second state selection magnet, or B magnet, cesium atoms that have made the transition from the 3,0 state to the 4,0 state are separated from the other atoms that were passed by the A magnet but have not made a transition. When the microwave frequency is correct, atoms that were in the 3,0 state are now in the 4,0 state and are directed by the B magnet toward the ionizer. The other atoms are directed away from the ionizer.

The ionizer is a heated tantalum ribbon upon which the cesium atoms are ionized and then evaporated. The cesium ions then pass through a mass spectrometer. The function of the mass spectrometer is to remove common contaminants such as potassium that might cause noise bursts that could overload the amplifiers in the frequency standard and cause loss of lock.

The ions that pass through the mass spectrometer are accelerated into a multistage electron multiplier, where the ion current is converted to an electron current and amplified. The output current of the electron multiplier is carried by a coaxial cable to the signal processing electronics.

Output Current

Fig. 8 is a recording of the output current of a typical high-performance beam tube as the micro-

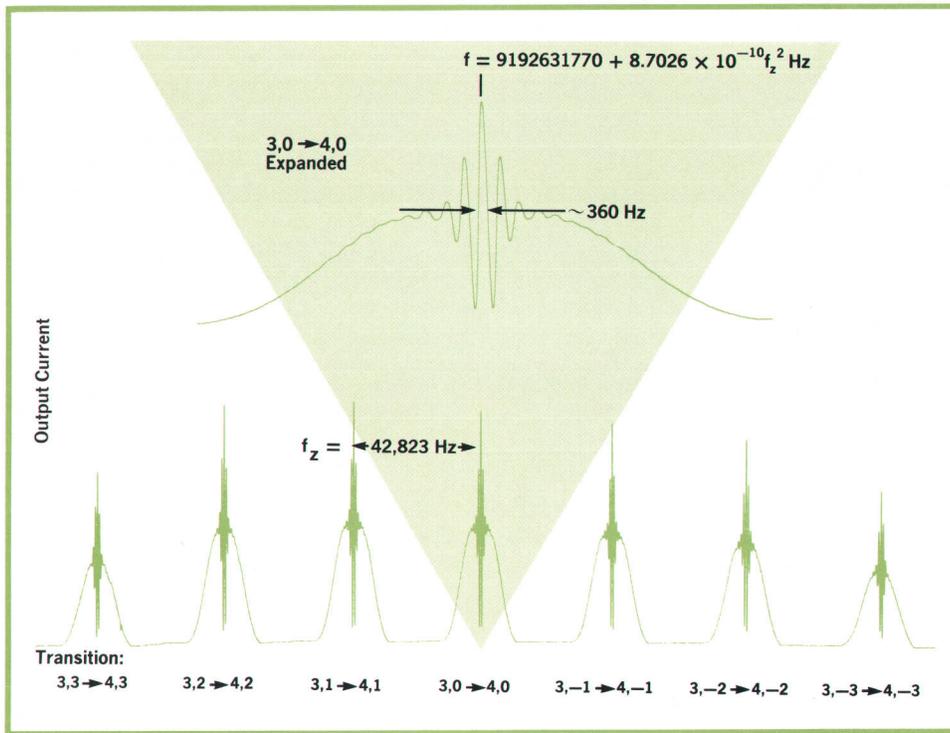


Fig. 8. High-performance beam tube output current versus frequency. Each peak is labeled with the name of the atomic transition that produced it. 360-Hz line width of the 3,0 to 4,0 peak is equivalent to a Q of approximately 2.5×10^7 .

wave frequency is swept. Each of the seven response pedestals is labeled with the transition that produced it. The center line of the center pedestal is the one used for frequency control. The other pedestals are also the result of transitions in which the second quantum number does not change, such as 3,-3 to 4,-3, which is the lowest-frequency microwave transition.

The fine structure (Ramsey pattern) on the top of each pedestal is caused by the difference in phase between the applied microwave signal and the free precession of the cesium atom. The precession, corresponding to the desired transitions being partially induced, is stimulated when the atom passes through the first arm of the cavity and its phase at that time matches the phase of the microwave signal. If the frequency of precession and the applied microwave frequency are equal, the phase difference is zero when the atom reaches the second arm of the cavity and the transition is further induced, leading to a maximum signal. Secondary maxima occur when the frequency difference is such that the phase difference is $\pm 2\pi$, $\pm 4\pi$, and so on. Minima occur when the phase difference is $\pm \pi$, $\pm 3\pi$, and so on. The spread in velocity of the cesium atoms causes the amplitudes of the maxima and minima to decrease rapidly with increasing frequency difference.

Narrow Line Width for Accuracy

The accuracy and stability of the cesium beam standard are related to the width of the center peak

in Fig. 8; the narrower this line width, the better the accuracy and stability. The new tube's line width is narrower, so the center of the resonance is more accurately determined by the frequency lock loop.*

The line width is determined by the length of time the atoms spend in the region between the cavity arms. To maximize this time, slow atoms are selected and the distance between the arms is made as large as possible. The velocity distribution of the selected atoms is below the most probable velocity for the beam, but not so low as to cause unacceptable reductions in signal, leading to poor signal-to-noise ratio.

The most obvious step in increasing the distance between the arms is to reduce the space at the ends of the tube that are not used in the beam path. The design length for the new beam path is 15.25 inches out of an overall beam tube length of 16 inches. This compares with 14 inches for the standard tube. It was possible to design the oven and detector parts to accomplish this, yet retain the same location of parts for cabling and microwave inputs as in the standard tube, thereby retaining mechanical interchangeability.

Short-Term Stability Improvement

One of the important performance parameters

*Besides narrower line width, there are two other contributions to improved accuracy in the new beam tube. One is the improved cavity, in which differential phase shifts between the cavity arms are smaller. The second is the new C field structure, which gives a more homogeneous field than earlier designs.

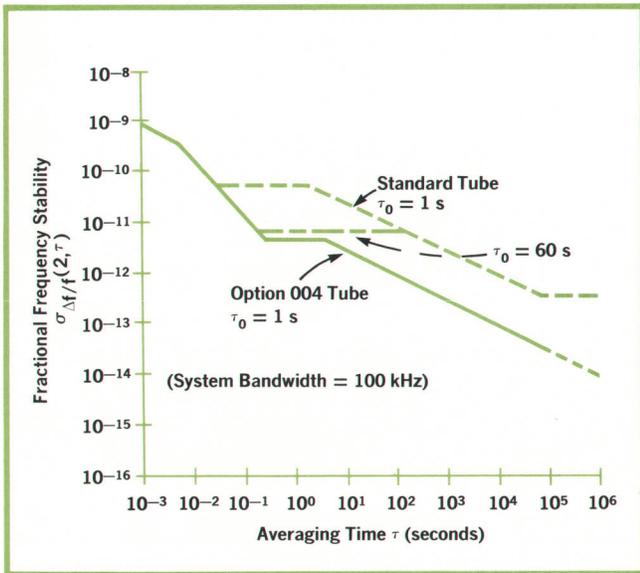


Fig. 9. Frequency stability of high-performance and standard beam tubes. Short-term stability specification of the tube for averaging times of one second or more is ten times better than that of the standard tube and is comparable to that of rubidium-vapor frequency standards.

for a cesium beam standard is the fractional frequency fluctuations of the output as a function of time. The ideal standard would have no variations in output frequency, regardless of averaging time. However, due to finite signal-to-noise ratios in the quartz oscillator, cesium resonator, and electronics, perturbations in the output frequency do occur.

Fig. 9 compares the stability of the high performance tube to that of the standard tube. The frequency stability is dependent on the inverse of

the square root of the averaging time τ (see Appendix), i.e.,

$$\sigma_y(2, \tau) = \frac{K}{\sqrt{\tau}}$$

This equation holds down to a few parts in 10^{14} .

The value of K has been reduced by a factor of ten for the high-performance resonator. Therefore, to achieve a given stability, the required averaging time has been reduced by a factor of 100.

This is a very significant improvement in averaging or measurement time. Because of it, the average frequency difference between two high performance standards can be measured to the same accuracy in 1/100th of the time required for standard tubes, or the accuracy can be improved by a factor of ten in the same measurement time.

In Fig. 10, the accumulated phase difference between two high-performance instruments is compared with the accumulated phase difference between a high-performance tube and a standard instrument. The improvement in frequency stability with the high-performance tube is evident.

Degausser Improves Settability

Another potential source of frequency instability is magnetic domain relaxation that can occur in the magnetic shielding and C field structure. When a change is made in the C field, the shields take a relatively long time to relax, or reach equilibrium. Relaxation can also occur due to changes in external magnetic fields, temperature changes, power interruptions, and shock.

Relaxation changes the value of the static mag-

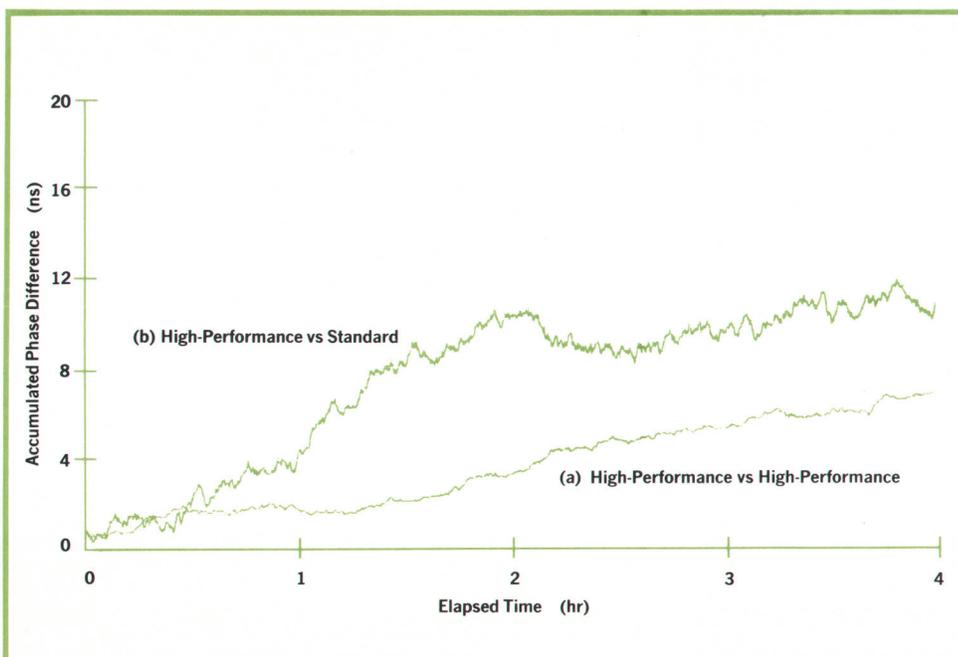


Fig. 10. Accumulated phase difference between two high-performance cesium beam frequency standards (a) and between a high-performance and a standard instrument (b). Quietness of (a) demonstrates the stability improvement of the new beam tube.

netic field established by the stable current flowing through the C field winding. Since the frequency of the resonator depends on the value of the C field, relaxation produces frequency changes in the standard.

For this reason, the improved tube has a built-in degausser coil that provides a means of hastening the relaxation of the magnetic domains to equilibrium. This complete degaussing provides improved repeatability and settability.

Degausser Accessory Model 10638A provides low-frequency alternating current to the degausser coil in the Cs resonator. The degausser should be used after initial turn-on of the instrument, exposure to external magnetic fields, and any change to the C field control. Degaussing takes about 20 minutes.

The degausser has two maximum current levels. The high current level is used for the initial degauss or any time the standard has been turned off. The high-level degauss should always be followed by a low-level degauss. The low current level produces negligible effect on the frequency of the standard and so may be used without interrupting the use of the standard as a clock. This is very important, because it allows very small, precise changes in frequency to be made by C field adjustments without interrupting the standard and without a long settling time.

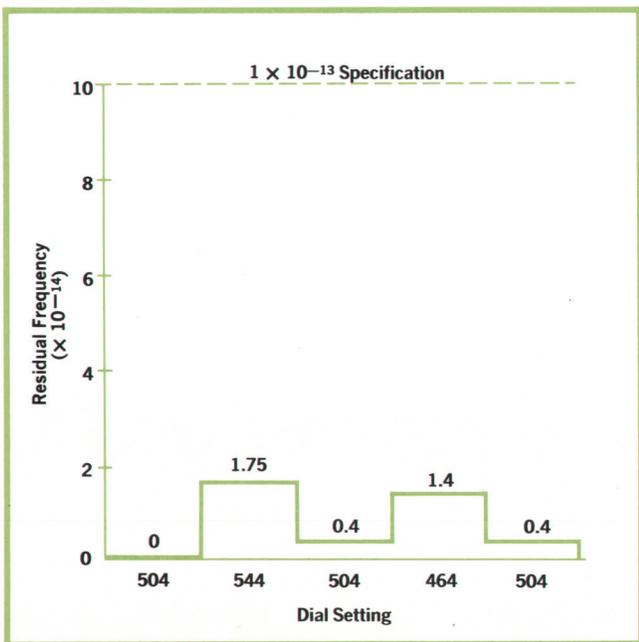


Fig. 11. Typical settability of the new high-performance cesium beam frequency standard. C field dial was changed by 40 divisions (1×10^{-13}) in each direction and returned to the starting point. Residual errors were as shown. Observation time was sufficient for a one-sigma accuracy of 1×10^{-14} .

Current 5061A Cesium Beam Frequency Standards are equipped with a higher-resolution C field control. This control makes it possible to take advantage of the improved settability of the high-performance tube and degausser to make very small changes in output frequency. The resolution of the C field control has been increased to 2.5×10^{-14} from 3×10^{-13} .

Settability and repeatability of the new tube are demonstrated by the actual performance shown in Fig. 11. The C field control was changed by 40 divisions (1×10^{-12}) and the residual error is shown as a function of C field control. In all cases the error is less than 1×10^{-13} .

Acknowledgments

The high-performance cesium beam tube development was directed by Len Cutler, now HP Physical Research Laboratory director, who also made many technical contributions to the concept and design. The initial design work was done at HP's Frequency and Time Division East in Beverly, Massachusetts in 1968 and 1969. In July, 1969, the beam tube development effort was transferred to HP Laboratories in Palo Alto, California. Beam tube manufacturing was transferred to the Santa Clara Division and combined with all other frequency standards activities directed at that time by Lee Bodily.

The late Joseph Holloway's contributions to this product were many. Robert Kern, formerly with Frequency and Time Division East, contributed to the early design efforts. Richard Lacey of HP Laboratories provided many valuable theoretical calculations.

The efforts of several people at Santa Clara Division in achieving a producible product are noteworthy: Dominick de Simone, production manager; Ernie Riberdy, tool engineer; Jack Elmberg, documentation; and Gary Seavey, production engineer.

The machine programming of Jorgen Hagglof, manufacturing division, was invaluable in the manufacture of the resonant cavity.

Joseph H. Holloway (1929-1971)

The scientific career of Joseph Holloway spanned the twenty-year period during which atomic beam resonance technology emerged from the academic laboratory to provide the basis for the extremely accurate and stable frequency and time standards available today. Dr. Holloway used the atomic beam method as a tool in his dissertation research at Massachusetts Institute of Technology and he then became the leading scientist in the practical development of the present highly refined instruments.

Following his graduate studies and the award of the Doctor of Philosophy degree in June 1956, Joe joined National Company in Melrose, Massachusetts where the industrial development of cesium atomic beam standards had been started two years earlier by Richard Daly. Professor Jerrold Zacharias, founder and at that time Director of the Atomic Beam Laboratory at M.I.T., was very much involved as a consultant during these industrial beginnings and his efforts provided the initiating force.

In 1960 Joe joined Varian Associates in Beverly, Massachusetts where his principal work contributed to the refinement of cesium resonators with empha-

sis upon compact, rugged design and high performance. The products of these developments were incorporated into new cesium frequency standards introduced by Hewlett-Packard Company, Ebauches, and Pickard and Burns. In 1967 Joe became a member of Hewlett-Packard after HP acquired the Varian frequency standard activity located in Massachusetts, and in 1969 he became a staff member of Hewlett-Packard Laboratories in Palo Alto.

Holloway's contributions to the industrial development of cesium frequency standards were directed toward all the elements of the resonator and the related interfaces with the electronic system.

Appendix

Short Term Stability

Frequency stability can be characterized in the frequency or the time domain. Spectral density of frequency and phase are the most common parameters in the frequency domain. Sample variance or standard deviation is the most common parameter in the time domain.

Consider a signal whose instantaneous voltage $V(t)$ may be written

$$V(t) = [V_o + \varepsilon(t)] \sin [2\pi\nu_o t + \phi(t)] \quad (1)$$

V_o = nominal amplitude
 ν_o = nominal frequency

If $\varepsilon(t)$ and $\phi(t)$ are small, then the instantaneous fractional frequency deviation can be defined as

$$\Delta f/f = y(t) = \frac{1}{2\pi\nu_o} \frac{d\phi(t)}{dt} \quad (2)$$

The spectral density of frequency fluctuations is then $S_y(f)$ and the spectral density of phase is $S\phi(f)$.

These two quantities are related as follows:

$$S_y(f) = \frac{f^2}{\nu_o^2} S\phi(f) \quad (3)$$

In the time domain, the sample variance is the most useful measure of the time dependence of the frequency:

$$\langle \sigma_y^2(N, T, \tau) \rangle = \left\langle \frac{1}{N-1} \sum_{n=1}^N \left(\bar{y}_n - \frac{1}{N} \sum_{k=1}^N \bar{y}_k \right)^2 \right\rangle \quad (4)$$

where $\bar{y}_k = \frac{\phi(t_k + \tau) - \phi(t_k)}{2\pi\nu_o\tau}$

$\langle \sigma_y^2 \rangle$ denotes the infinite time average.

However, as $N \rightarrow \infty$, σ_y^2 does not converge for $S_y(f)$ containing terms dependent on f^n where $-\infty < n \leq -1$.

To compare data on a common base, it is important to specify N and T . The preferred definition is $N=2$ and $T=\tau$. This definition, known as the Allan Variance, may be written

$$\sigma_y^2(2, \tau) = \left[\frac{(\bar{y}_{k+1} - \bar{y}_k)^2}{2} \right] \quad (5)$$

In a real situation the infinite time average cannot be realized, so a finite number of samples must be used. Then

$$\sigma_y^2(2, \tau) \approx \frac{1}{2M} \sum_{m=1}^M \left(\bar{y}_{k+1} - \bar{y}_k \right)^2 \quad (6)$$

The relation between $S_y(f)$ and σ_y^2 is

$$\sigma_y^2(2, \tau) = 2 \int_0^{\infty} df S_y(f) \left\{ \frac{[\sin^2 \pi f \tau]}{(\pi f \tau)^2} \right\} \left\{ 1 - \frac{\sin^2 2\pi f \tau}{4 \sin^2 \pi f \tau} \right\} \quad (7)$$

The dominant spectral density for the cesium resonator is given by

$$S_y(f) = h_o \quad (8)$$

that is, the frequency noise is white noise. Thus

$$\sigma_y^2(2, \tau) = h_o/2\tau \quad (9)$$

This dependence on averaging time holds for

$$10^0 < \tau < 10^5 \text{ seconds}$$

for the 5061A.

Reference: "Characterization of Frequency Stability," U.S. National Bureau of Standards Technical Note 394, October 1970.



Ronald C. Hyatt (Center)

Ron Hyatt is manager of precision frequency sources development, marketing, and production at HP's Santa Clara Division. He joined HP in 1964 as a development engineer in the precision frequency sources laboratory, became engineering section manager in 1968, and assumed his present responsibilities earlier this year. In 1966 and again in 1967, Ron was a participant in HP flying clock trips, carrying cesium beam standards to timekeeping laboratories around the world. A native of Texas, Ron received his BSEE degree from Texas Technological University in 1962, then came to California and earned his MSEE degree at Stanford University in 1963. He and his wife and two sons now live in Los Altos, California. Tennis and trout fishing head his list of leisure-time activities.

Terry N. Osterdock (Left)

Terry Osterdock is applications engineer for precision frequency sources at HP's Santa Clara Division. By the time this appears in print, he'll have an MSEE degree from the University of Santa Clara to add to the other two degrees he's already received from that school—a BEE in 1964 and an MBA in 1966. Before he joined HP in 1972, Terry had been a missile project controller, a U.S. Army data communications officer, and an engineering instructor. He's a member of IEEE, a registered professional engineer in the state of North Carolina, a member of the American Society of Enologists, and an amateur photographer. Born in Oakland, California, Terry now lives in San Jose with his wife and two children.

Louis F. Mueller (Right)

Lou Mueller joined what is now HP's precision frequency sources group in 1962, and has been heavily involved in hydrogen maser and, more recently, cesium beam tube development projects. Lou is a 1955 graduate of Washington University in St. Louis, Missouri, his birthplace; he holds the AB degree in physics. He's co-authored a number of papers on frequency control and holds a patent on laser machining. Before getting into precision frequency sources, he did cryogenic design and chemical instrumentation research. Hiking, backpacking, and bicycling are among Lou's favorite recreational activities. He and his wife and four children live in Palo Alto, California.

SPECIFICATIONS

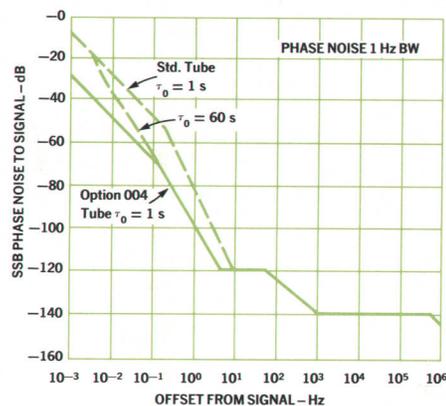
HP 5061A Cesium Beam Standard Option 004 High Performance Cesium Beam Tube

	Standard Beam Tube	Option 004 High-Performance Beam Tube
ACCURACY: maintained over a temperature range of 0 to 50°C and magnetic fields up to 2 gauss or any combination thereof.	$\pm 1 \times 10^{-11}$	$\pm 7 \times 10^{-12}$
REPRODUCIBILITY	$\pm 5 \times 10^{-12}$	$\pm 3 \times 10^{-12}$
SETTABILITY (Frequency)	$\pm 7 \times 10^{-13}$	$\pm 1 \times 10^{-13}$ (with 10638A Degausser)
LONG-TERM STABILITY (for life of cesium beam tube)	$\pm 5 \times 10^{-12}$	$\pm 3 \times 10^{-12}$
DC MAGNETIC FIELD STABILITY: frequency change, any orientation in a 2 gauss field.	$\pm 2 \times 10^{-12}$	$\pm 2 \times 10^{-13}$
TIME CONSTANT: (quartz oscillator control loop)	1 and 60 s	1 s
WARM-UP TIME (at 25°C)	45 min	30 min
BEAM TUBE WARRANTY	3 years	14 months (10,000 hr)

FREQUENCY STABILITY: See Fig. 9, page 21.

τ (s)	$\sigma_{\Delta f/f(2,\tau)}$ for $\tau_0 = 1$ s	
	Standard Tube	Option 004 Tube
10 ⁻³	8.2×10^{-10}	8.2×10^{-10}
10 ⁻²	1.5×10^{-10}	1.5×10^{-10}
10 ⁰	5.6×10^{-11}	5×10^{-12}
10 ¹	2.5×10^{-11}	2.7×10^{-12}
10 ²	8×10^{-12}	8.5×10^{-13}
10 ³	2.5×10^{-12}	2.7×10^{-13}
10 ⁴	8×10^{-13}	8.5×10^{-14}

PHASE NOISE:



Δf	dB for $\tau_0 = 1$ s	
	Standard Tube	Option 004 Tube
10 ⁻³	-8	-18
10 ⁻²	-28	-48
10 ⁰	-82	-96
10 ¹	-120	-120
10 ²	-125	-125
10 ³	-140	-140
10 ⁶	-146	-146

PRICES IN U.S.A.: HP Model 5061A with standard beam tube, \$16,700
Option 004 High-Performance Beam Tube, add \$2160

MANUFACTURING DIVISION: SANTA CLARA DIVISION
5301 Stevens Creek Boulevard
Santa Clara, California 95050 U.S.A.

Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, California 94304

Bulk Rate
U.S. Postage
Paid
Hewlett-Packard
Company

HEWLETT-PACKARD JOURNAL

SEPTEMBER 1973 Volume 2 • Number 1

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard S.A., CH-1217 Meyrin 2
Geneva, Switzerland
Yokogawa-Hewlett-Packard Ltd., Shibuya-Ku
Tokyo 151 Japan

Editorial Director • Howard L. Roberts
Managing Editor • Richard P. Dolan
Contributing Editors • Ross H. Snyder,
Laurence D. Shergalis

Art Director, Photographer • Arvid A. Danielson
Art Assistant • Sue M. Reinheimer
Administrative Services • Anne S. LoPresti
European Production Manager • Kurt Hungerbühler

CHANGE OF ADDRESS: To change your address or delete your name from our mailing list please send us your old address label (it peels off). Send changes to Hewlett-Packard Journal, 1501 Page Mill Road, Palo Alto, California 94304 U.S.A. Allow 60 days.