


HEWLETT-PACKARD JOURNAL



3582A SPECTRUM ANALYZER
HEWLETT-PACKARD



RESET
TRACE 1 STORE RECALL
TRACE 2 STORE RECALL
ASTIG
INTENSITY
CUT OFF
FOCUS
GRAT ILLUM
NOISE SOURCE OUTPUT
LINE
OFF
ON
STATUS LISTEN
TALK
REMOTE
LOCAL
K-Y RECORDER
PLOT
1-16
V6
(RESET)

DISPLAY
AMPLITUDE
SCALE
POSITION
PHASE
TIME
AVERAGE
OFF
RMS
TIME
PEAK
NUMBER
16
37
RESTART
NOISE SOURCE OUTPUT

MARKER
DEL
REF
TRACE SET ENDS
POSITION
TRIGGER
CHANNEL A ONLY
DATA LOADING
LEVEL
SLOPE
REPETITIVE
ADM
COUPLING
A
B
100V MAX
100V MIN

FREQUENCY SPAN
50dBZ
INPUT
CHANNEL A SENSITIVITY
OVERLOAD
INPUT MODE
AUTO
B COUPLING

A High-Resolution, Low-Frequency Spectrum Analyzer

This dual-channel instrument uses digital computation with a microprocessor to make frequency-domain measurements in the 0-to-25.5-kHz range with bandwidths as narrow as 20 mHz, and do it hundreds of times faster than conventional swept-frequency analyzers. At the same time, a number of other important capabilities are obtained.

by Nixon A. Pendergrass and John S. Farnbach

THE RECENT ADVANCES in large-scale integrated-circuit technology have brought substantial improvements to the performance of measuring instruments. In particular, microprocessor control of spectrum analyzers, recently described in these pages,^{1,2} brought significant advances in capability to this class of instrument.

A microprocessor is used in a different way to achieve high performance in a new, moderately-priced dual-channel spectrum analyzer (Fig. 1) for use in the audio and subaudio frequency range (0.02 Hz to 25.5 kHz). The digital technology used in this instrument, Model 3582A, provides greatly increased measurement speed and several measurements not available with analog techniques, such as true-rms averaging and a coherence function.

The measuring bandwidth of the Model 3582A can be as narrow as 0.02 Hz in a 5-Hz frequency span anywhere in the 0.02-Hz-to-25.5-kHz range, and on the 0-to-1-Hz frequency span the bandwidth can be as narrow as 0.004 Hz. The instrument's dynamic range is more than 70 dB with full-scale ranges of +30 dBV down to -50 dBV (+30 V rms to 3 mV rms).


The 3582A processes incoming signals digitally and stores the results for repetitive, flicker-free display on a fully annotated, high-resolution CRT (Fig. 2). The method used to process the incoming signals, however, also derives phase information, a capability not normally associated with spectrum analyzers. Thus, besides obtaining spectrum displays, this analyzer, with its two input channels and phase measurement capability, can be used to measure transfer functions (Fig. 3). This is a measurement already familiar to electrical engineers and it is of growing importance to mechanical engineers, as in determining how a structure reacts to a forcing function. The transfer function phase spectrum is particularly useful for pinpointing resonant frequencies.

To enhance these capabilities, the analyzer has a

built-in noise source that generates all frequencies of interest simultaneously for use as a measurement stimulus. This provides the user with a complete spectrum and network analysis instrument in a single compact package.

Transient Capture and Averaging

Because it uses digital storage, the 3582A can cap-



Cover: Analysis of low-frequency sounds is but one of many uses for the Model 3582A Spectrum Analyzer described in this issue. The spectrum analyzer's coherence function is being used here with a microphone and an accelerometer to find out which parts of a machine contribute to the overall sound level and what their contributions are.

In this Issue:

A High-Resolution, Low-Frequency Spectrum Analyzer, by Nixon A. Pendergrass and John S. Farnbach **page 2**

Window Functions for Spectrum Analysis, by Roger G. Cox, **page 10.**

Designing Programmable Digital Filters for LSI Implementation, by Lynn A. Schmidt **page 15**

Desktop Plotter/Printer Does Both Vector Graphic Plotting and Fast Text Printing, by Majid Azmoon, Jaime H. Bohorquez, and Rick A. Warp **page 24**

Plotter/Printer Interface Languages: HP-GL and ASCII, by Michael P. Trego, **page 26.**




Fig. 1. Model 3582A Spectrum Analyzer works with signals that have frequency components in the range between 0.02 Hz and 25.5 kHz. Digital processing with custom integrated circuits provides high performance in a compact instrument at moderate cost.

ture single-shot transient waveforms for subsequent display in a way similar to digital storage oscilloscopes (Fig. 4). Also, with an appropriate trigger signal, up to 256 repetitions of a waveform can be averaged, improving the signal-to-noise ratio of the waveform's time record by as much as 24 dB. Analysis can then be performed on the averaged waveform to

derive amplitude and phase spectra.

Alternatively, the 3582A can take the power average (RMS AVERAGE) of up to 256 spectra to reduce uncertainties when characterizing signals that have random components. An exponential form of averaging can also be applied to reduce the contributions of older spectra as new spectra are added so changing spectra can be observed while being smoothed by




Fig. 2. The fully annotated display of Model 3582A has four lines of alphanumeric information giving the instrument's measurement configuration. The first two lines from the top give vertical axis information; the bottom two give frequency information. A movable marker (bright dot) reads the frequency and amplitude of the indicated point with three-digit resolution, displaying either absolute or relative values.




Fig. 3. Model 3582A obtains transfer functions, both magnitude and phase, by relating the signal at the output of a device to the signal at the input. The transfer function can be measured over frequency spans as narrow as 5 Hz anywhere within the 25-kHz range of the analyzer.




Fig. 4. Transient waveforms can be captured for display as an amplitude-vs-time record (left). The stored waveform can then be analyzed, such as finding the amplitude spectrum (right).

averaging.

An additional “averaging” mode, referred to as PEAK HOLD, retains the highest value encountered at each frequency as several successive spectra are processed. This allows the analyzer to be used with swept-frequency sources and to perform other tasks such as measuring the amount of frequency drift in a signal.

Coherence

Model 3582A also has a COHERENCE function. Used with RMS AVERAGING, this relates the power spectrum of a signal at the output of a device to the power spectrum of the input signal, giving an indication of how much of the output power is a result of the input (Fig. 5). At frequencies where the coherence function is 1.0 (top line of the CRT graticule), the output is caused entirely by the input. At frequencies where the coherence function is less than 1.0, noise or distortion within the device is contributing to the output. This function is particularly useful for investigating causal relationships in multiple input systems because it can give a measure of how much each input contributes to the output while all inputs are active simultaneously. It also provides insight into the accuracy of transfer function measurements since it gives an indication of how much the measurement is disturbed by noise,

harmonics, and other unwanted signals generated within the system being evaluated.^{3,4}

Display Features

The 3582A gathers input data in one block of read-write memory (RAM) and uses other blocks of RAM for processing and displaying the data. This allows the measurement function to be changed and the result displayed without the need to gather new data. This is important in an analyzer that has narrow resolution bandwidths, since gathering one time record for processing can take up to 250 seconds.

With the exception of time functions, which are displayed singly, any two traces may be displayed at the same time. This enables amplitude and phase to be displayed simultaneously. In addition, any two traces may be stored for later recall. Thus, a newly-acquired transfer function amplitude curve may be compared against a stored curve.

Considerable flexibility exists for display of amplitude information. The display may be linear in volts or gain, or it may be logarithmic in dBV or dB at 10 dB/div or 2 dB/div. The AMPLITUDE REFERENCE LEVEL switch increases the display gain in 10-dB steps for convenient positioning of the displayed information. The reference level and the scale factor are displayed numerically on the CRT for ready reference.




Fig. 5. The coherence function (right) gives a measure of the fraction of power at each frequency in the signal at the output of a device (center) that is caused by the signal at the input (left). Among other uses, this can indicate the validity of a transfer function measurement by showing how much output power comes from sources other than the driving signal.

Phase information is displayed on a $\pm 200^\circ$ vertical scale, giving hysteresis at the $\pm 180^\circ$ boundaries that prevents the phase trace from continually jumping between the top and bottom of the display if the phase measurement jitters across $\pm 180^\circ$. To further prevent ambiguous displays, whenever the signal amplitude falls more than 65 dB below full scale where noise would obscure and confuse phase values, the phase measurement is suppressed and zero phase is displayed.

A marker in the form of a brightened dot may be placed on any displayed trace except time-function and recalled traces. The frequency and amplitude or phase at that point in the spectrum are displayed numerically on the CRT (see Fig. 2). Pressing the MARKER SET REF key stores the marker frequency and amplitude or phase as a reference. Then at any later time, pressing the MARKER REL key causes the present marker value relative to the stored reference to be displayed. If a point on the display is desired as either the START or CENTER frequency of a measured frequency span, placing the marker on that point then causes that frequency to be stored as the new START or CENTER frequency when the SET FREQ key is pressed. Pressing the $\div \sqrt{\text{BW}}$ key gives a reading of spectral density at the MARKER position automatically normalized to a 1-Hz bandwidth for noise density measurements.

HP-IB Compatibility

Model 3582A has an HP-IB* port, enabling remote control of the analyzer in automatic test systems. The HP-IB port also allows reading from and loading into memory, halting and continuing signal processing at specific points, and reading from and loading into the CRT alphanumeric display. It is possible, for example, to off-load the spectrum signature of a rotating machine onto magnetic tape through the HP-IB and at a later time reload it into the analyzer for comparison with the present signature.

When controlled through the HP-IB port by an external controller, such as the Model 9825A Desktop

*HP-IB: The HP Interface Bus, Hewlett-Packard's implementation of IEEE Standard 488-1975 and ANSI Standard MC1.1.

Correction

The last line on page 8 of our August issue was missing. Here it is, in case you'd like to copy it and paste it in:

mixer positive transition. The phase-locked loop has adjusted

The complete section reads: "The condition before triggering is now reestablished, that is, the zeroth count is matched in time to the mixer positive transition. The phase-locked loop has adjusted itself to lock to the new beat frequency."

Also, on page 4, the last line before the title "Hysteresis Arming" should read "In +TI ONLY the 5370A will not measure any time interval less than ten nanoseconds." The word "not" was inadvertently omitted.

And, on page 16, "emitter-follower-logic (EFL)" should have read "emitter-function-logic (EFL)." There's a big difference.

Computer, the Model 3582A gains greatly enhanced signal analysis capability. With user-written software, the knowledgeable user can apply the spectrum analyzer/desktop computer combination to the computation and display of such quantities as auto- and cross-correlation functions, impulse responses, and cross-power spectra.

Displayed traces are also converted to analog form and supplied to an X-Y recorder output.

Internal Details

A simplified block diagram of the Model 3582A Spectrum Analyzer is shown in Fig. 6. Each input signal initially goes through an attenuator and an amplifier that has programmable gain. This is followed by a low-pass filter. The signal then goes to an analog-to-digital converter (ADC), after which all processing is done digitally.

The low-pass filter in each channel is required to prevent aliasing in the analog-to-digital conversion, i.e., it removes high frequencies that would appear as low frequencies when sampled. The ADC sampling rate is 102.4 kHz, so high frequencies appear to be folded about 51.2 kHz. Unless filtered out, input signals above 76.8 kHz would therefore appear on the display as below 25.6 kHz.

The low-pass filtering is done by an active seventh-order, elliptic filter that has a cut-off frequency of 25.6 kHz and a 100 dB/octave roll-off. The stopband above 70 kHz is over 80 dB down while the passband ripple is less than 0.1 dB.

The ADCs are 12-bit successive-approximation converters. To reduce the effects of quantization and associated nonlinearities for low-level signals, an out-of-band 27-kHz sine wave is injected as a dither signal into the signal path before each ADC at a level 45 dB below full scale. Since the dither signal is not synchronized with the sampling, the noise it produces is broadband. A small signal, that otherwise might fall between two quantization levels and not be detected, alters the statistics of the quantization noise so it is resolved by subsequent processing for display. Although the quantization noise is large relative to a -75-dB signal, it is broadband white noise so the narrow passband of the 3582A reduces it to well below -80 dB with respect to full scale, enabling the instrument to consistently detect signals at the -75-dB level.

Digital Heterodyning and Filtering

The outputs of the ADCs are applied to multipliers that digitally "heterodyne" the center of the frequency band of interest down to 0 Hz. Frequencies outside the band of interest can then be removed by low-pass filters. This is the same technique used in the Model 5420A Digital Signal Analyzer to obtain band selectable analysis.⁵

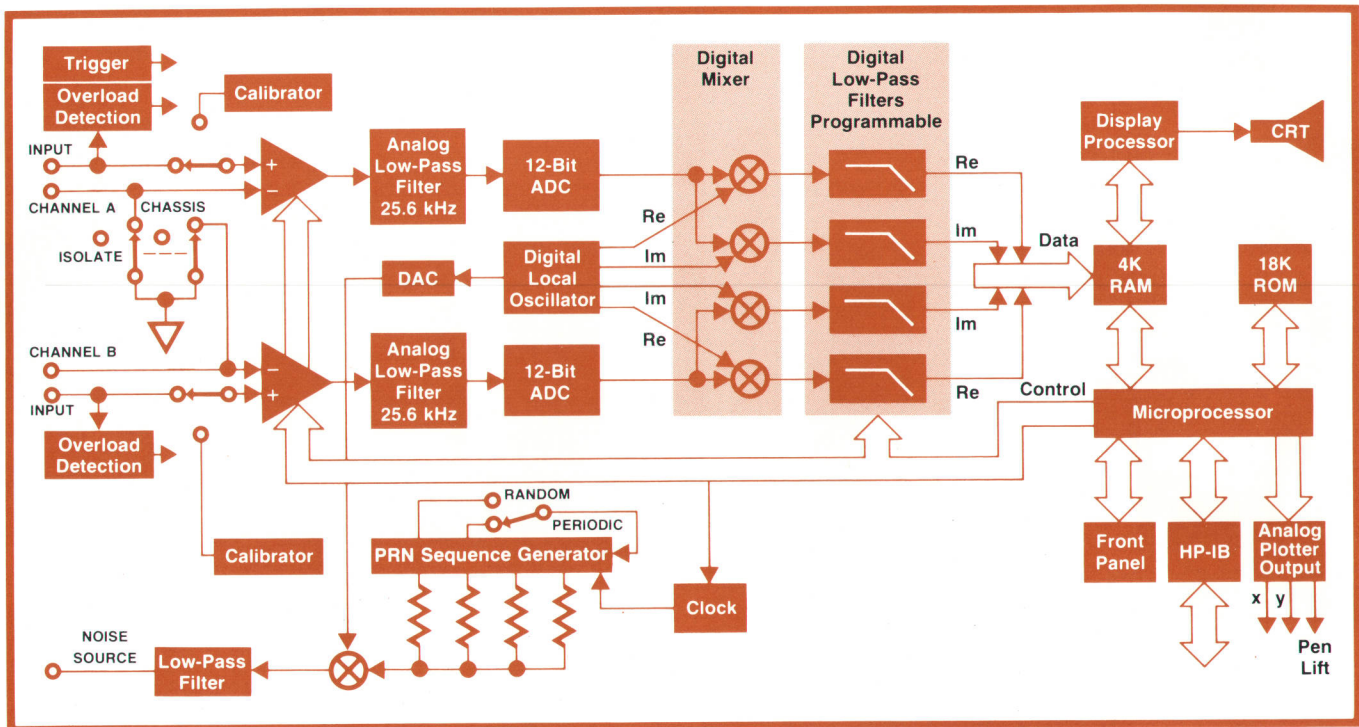


Fig. 6. Simplified block diagram of the Model 3582A Spectrum Analyzer. The entire signal-processing chain from the inputs through the ADCs is isolated electrically from chassis ground and the remainder of the instrument to help eliminate ground loops.

Traditional analog heterodyning techniques for shifting a band of frequencies into a fixed bandpass filter have image-frequency problems that become increasingly severe as the band of frequencies approaches 0 Hz where the images become very close to the desired frequency band. In band-selectable analysis, the samples of the input waveform are multiplied digitally by samples of a complex waveform, $\cos 2\pi f_c t - j \sin 2\pi f_c t$, instead of using a real multiplication by $\cos 2\pi f_c t$. The effect of this complex multiplication is to slide the whole frequency spectrum to the left along the frequency axis so the selected center frequency (f_c) is at 0 Hz.⁶ Frequencies that otherwise would become close-in image frequencies thus maintain their relative positions with respect to the desired frequency band and are readily removed by low-pass filtering.

The result of the multiplication is two data streams, one representing the real components of the frequency-shifted spectrum and the other representing the imaginary components (for zero-start frequency spans that require no frequency shift, the LO inputs to the digital mixers are set to +1 in the 3582A so the data streams pass unchanged through the multipliers). The two data streams are low-pass filtered in digital filters to obtain the desired frequency span and then stored in RAM where they are held for subsequent processing. The frequency-shifting and filtering thus enable a narrow band of

frequencies anywhere in the analyzer's frequency range to be isolated for high-resolution examination.

Each filter and associated multiplier (digital mixer) are integrated on a single IC chip (see following article). Putting these complex circuits into integrated circuits reduced the parts count significantly and was a major factor in achieving compact size and moderate cost in this high-performance instrument.

Digital Oscillator

The samples of the waveform, $\cos 2\pi f_c t - j \sin 2\pi f_c t$, used as the "local oscillator" signal, are produced by a digital generator that uses cosine values from a table stored in read-only memory (ROM). Linear interpolation between stored values allows the table length to be confined to 1024 values.

A block diagram of the generator is shown in Fig. 7. A binary representation of the selected center frequency is stored in a latch. This number determines the incremental phase angle between samples of the output cosine wave, and is added repeatedly to the total accumulating in the second latch to select the addresses of 16-bit samples of a full cosine cycle stored in the cosine ROM. This ROM holds 1024 samples.

For a given phase, both $\cos \theta$ and $-\sin \theta$ are computed. Since for low frequencies the phase must be specified to a greater precision than $1/1024$ of a cosine cycle, the computation of $\cos \theta$ is done by using the

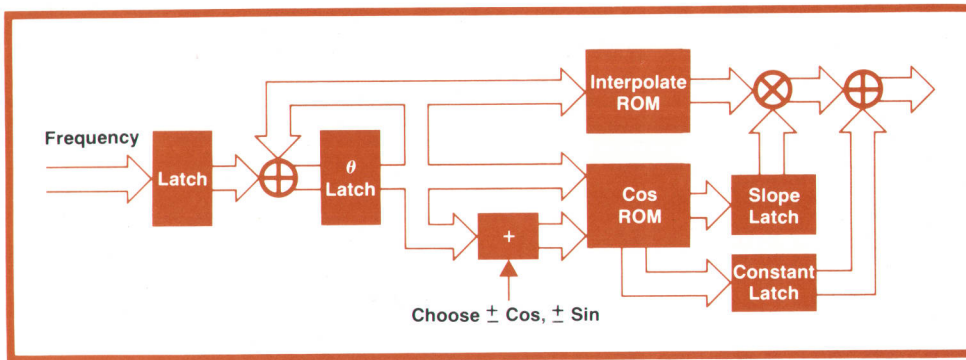


Fig. 7. Block diagram of the "digital oscillator." Instead of waveforms, this generator supplies digital words corresponding to amplitude samples of sine and cosine waveforms.

first two terms of a Taylor series expansion:

$$\cos\theta = \cos(\phi + \epsilon) = \cos\phi + \epsilon \frac{d}{d\phi} \cos\phi = \cos\phi - \epsilon \sin\phi,$$

where ϕ is a phase value in the cosine ROM, and ϵ is the error so $\phi + \epsilon = \theta$. The computation is in fractions of a cycle so ϵ must be converted to radians. These conversions are stored in the interpolate ROM.

The values of $\cos\phi$, ϵ , and $-\sin\phi$ are latched at the outputs of the ROMs and the remainder of the computation is done by a hardware multiplication and addition to produce $\cos\theta$. A similar computation is used to obtain the $-\sin\theta$ output.

Processing the Data

The filtered waveform samples stored in RAM are processed by a fast Fourier transform (FFT) algorithm (see box, next page). The result of the FFT processing is a series of 512 values representing the real and imaginary parts of the input waveform spectrum at 256 points in the selected analysis band (128 points for each channel in dual-channel measurements). These may then be presented on the CRT display as a plot of amplitude vs. frequency. The phase relationships of the frequency components with respect to the waveform reference are also derived from the real and imaginary data and may be presented as a graph of phase vs. frequency.

In effect, the 3582A functions as though the input waveform were applied to a bank of 256 narrow-band filters in parallel and the output of each filter were represented on the display by a discrete dot at the corresponding point on the display frequency scale. A line generator connects adjacent dots to obtain a continuous trace. The measurement can be made with 256-line resolution within narrow frequency spans located anywhere in the basic 0-25.5 kHz range of the instrument when using the digital heterodyne technique to obtain band-selectable analysis ("zoom"). This greatly enhances the analyzer's resolution. For example, the basic spectral line spacing with the 5-Hz frequency span is 0.02 Hz, and this span can be moved anywhere within the 0-25.5-kHz range of the instrument. This gives the equivalent of a resolution of 1,280,000 lines over the full 0-25.5-kHz

range.

Processing Speed

FFT processing, display preparation, and averaging take somewhere between 350 and 600 ms to transform a time record into a displayed spectrum, depending on the instrument's operating and display mode. Thus, the spectrum can be presented almost immediately upon conclusion of data gathering. Furthermore, for frequency spans up to 500 Hz, the instrument has "real-time" operation, that is to say,

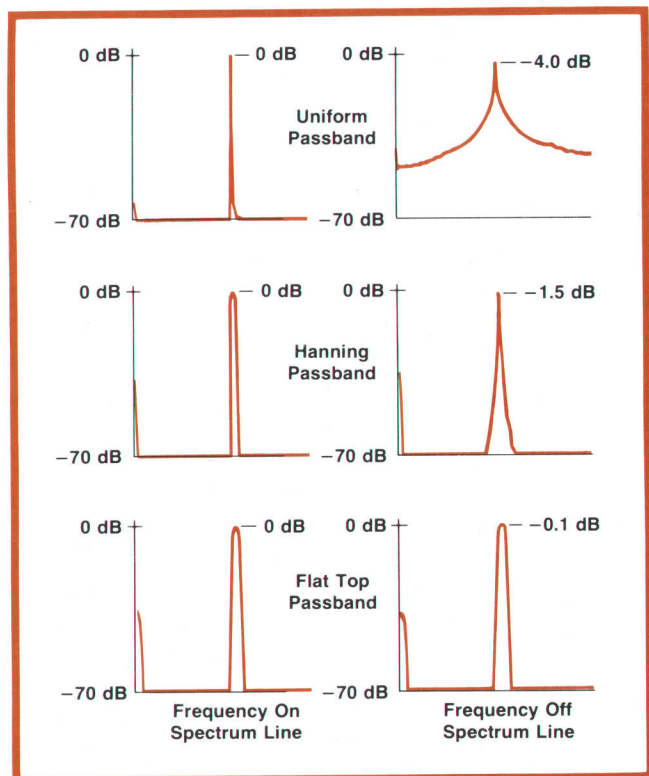


Fig. 8. The effect of the passband shape on spectrum displays of a single tone. Displays in the column at left are the result of tuning the input signal frequency to one of the FFT spectral frequencies. Displays at right result from tuning the input frequency to a point half way between two FFT spectral frequencies where the window length is not an exact multiple of the signal period. In all cases, the input signal amplitude is 0 dBV.

The Fast Fourier Transform and the Model 3582A

Digital processing of spectrum analyzer signals offers several advantages over analog approaches. One of the most significant is that filtering operations can be performed with no drift and with high noise immunity. Complicated arithmetic operations that would be expensive or nearly impossible as analog tasks are readily implemented by digital techniques.

The Model 3582A Spectrum Analyzer uses digital processing with the fast Fourier transform (FFT) to derive a complete 256-line spectrum from just one input record, an especially important capability in an analyzer with very narrow bandwidths. For example, using the 0-1-Hz frequency span, Model 3582A takes 250 seconds to obtain a single time record from which it derives a spectrum with 0.004-Hz resolution. A swept-frequency spectrum analyzer would take over 100 times as long to obtain a spectrum with that resolution.

As explained in the main text, Model 3582A takes samples of the waveform to be analyzed and processes them digitally to obtain a sequence of words that correspond to samples of a low-pass or bandpass filtered version of the input waveform. This sequence is then processed by an FFT algorithm to obtain another sequence of words that corresponds to the real and imaginary parts of the frequency spectrum of the filtered input waveform.

During the design of this instrument, considerable effort was devoted to deriving an FFT algorithm that could be performed by a microprocessor in a reasonable amount of time using a minimum of memory.

Computing the FFT

For a time record consisting of N amplitude samples $x(n)$, where $n=0, 1, 2, \dots, N-1$, the fast Fourier transform calculates the frequency spectrum $X(k)$, where $k=0, 1, 2, \dots, N-1$, at N frequencies:

$$X(k) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi nk/N}$$

To achieve the necessary processing speed, the FFT software in Model 3582A implements this calculation with integer arithmetic.* The software was designed to achieve the highest possible processing speed without unreasonable requirements on the program memory space, and to minimize noise introduced by truncation and rounding without introducing arithmetic overflow problems.

The FFT calculation is not affected by the selected frequency span. Since the digital filters produce data samples at a rate that is properly scaled for the selected frequency span, the FFT software needs to calculate spectra based only on the time and frequency indices n and k . The display processing software applies the proper scaling to the frequency axis of the display after the FFT processing.

*With integer arithmetic, sometimes referred to as fixed-point arithmetic, values are represented by signed integer binary numbers instead of mantissas and exponents (scientific notation). The latter is often called floating-point arithmetic.

In the single-channel baseband mode, the time record consists of 1024 real samples that are formed into an array representing a 512-point complex record by a process known as *scrambling*.¹ The complex FFT algorithm is then applied to this array followed by an *unscrambling* routine that derives the non-negative frequency components of the original time record. This use of scrambling and unscrambling techniques avoids the wasted processing time associated with applying the FFT directly to real data, which would result in calculating both the positive and negative frequency components of a symmetric spectrum.

Further processing speed efficiencies were achieved in the FFT software by implementing the 512-point FFT as a base-8 *tiddle factor* algorithm.² With this organization, most of the multiplications internal to the base-8 transform are by factors of ± 1 and $\pm j$. These do not actually take up processing time so do not slow overall processing speed.

The base-8 FFT organization, however, required significantly more program space than base 2, because considerations of speed required explicit programming of every internal step of the base-8 transform. However, this cost was minimized by using a Sande-Tukey algorithm for the base-8 transform, rather than a Cooley-Tukey. This structure provides an entry into the base-8 subroutine at a point where it performs a dual base-4 transform. With this, a 256-point FFT could be programmed using a base $4 \times 8 \times 8$ tiddle factor algorithm that used the same base-8 subroutine. The 256-point FFT algorithm is necessary for the dual-channel "zoom" mode where two independent, 256-point, complex records are processed to give two 256-point complex spectra.

Also, because of the places in the Sande-Tukey algorithm where multiplications occur, truncation and rounding introduce less noise.

In its final form, the FFT program package consists of about 1100 program words, the bulk of which is devoted to the base-8 subroutine. Much smaller portions are devoted to the control routines for the 512- and 256-point transforms and an even smaller part is devoted to miscellaneous subroutines.

In the zoom mode, either single- or dual-channel, the digital heterodyning supplies the data samples to the FFT processing already in complex form, so the scrambling and unscrambling processes can be bypassed. The FFT algorithm then produces a 512-point complex spectrum when the instrument is in the single-channel mode.

In the dual-channel baseband mode, the time record consists of two 512-point records containing real samples. These are scrambled into a single 512-point complex record that is processed by the 512-point complex FFT algorithm, then unscrambled to give the non-negative frequency components of the two input record spectra.

References

1. E. O. Brigham, "The Fast Fourier Transform," Prentice-Hall 1974, pp 163-169.
2. *ibid.*, pp 191-193.

there is no gap between the end of one time record and the beginning of the next. This is because the instrument can acquire a new time record while simultaneously processing the previous record stored in memory.

Windows

One consequence of making an FFT measurement for a finite length of time is that the length of the time span or "window" during which the measurement is made affects the resulting spectrum. The mathema-

tics of FFT processing deals with signals as though they were periodic with a period equal to the window length. If the signal is indeed periodic, and if the length of the window is an exact multiple of the signal period, then the computed FFT spectrum is similar to the corresponding analog spectrum. When there is no such relationship between the signal and the window, a discontinuity exists that results in broadening the base of response peaks (see Fig. 8). For this reason, FFT analyzers offer a means of attenuating the beginning and end of the window so the signal is brought smoothly to the zero amplitude level at both ends. The modified signal waveform then has no discontinuity between the end of the acquired time record and the beginning of a hypothetical repetition of the same record. The results of FFT processing are thereby made to agree more closely with the true spectrum.

The shape of the window affects the shape of the resolution bandwidth (see box, next page), so the 3582A provides a user with a choice of three window shapes. In keeping with the spectrum analyzer orientation of the instrument, these are labeled PASSBAND SHAPE on the front panel. The three passbands offered are UNIFORM, HANNING, and FLAT TOP.

The UNIFORM passband is basic in that it allows the input signal time record to be treated with a constant gain factor during FFT processing. Since it does not alter the time record, it is used primarily for capturing transient signals. Many transients start at the zero level and decay to the zero level so no window shaping is necessary and the onset of a transient, which may contribute significantly to the spectral properties of the transient, can be left undisturbed (see Fig. 4). The UNIFORM passband provides the narrowest frequency resolution, so it is used whenever the periodicity assumption of the FFT can be satisfied without the end-point discontinuity problem. As will

be explained below, the UNIFORM passband is also useful when the built-in noise source is used as a stimulus for the measurement.

The FLAT TOP passband, developed especially for the 3582A, uses a window with a gain-vs-time characteristic that results in a filter response that is practically flat across adjacent spectral lines. Thus, signal components that lie between FFT spectral lines suffer less than 0.1-dB amplitude degradation. This passband is useful for analyzing the spectra of discrete tones where amplitude accuracy is important. There is considerable overlap of adjacent passband responses, however, so it is more difficult to resolve fine frequency detail with this passband.

The well-known HANNING passband, widely used in FFT analyzers, results in a passband width that is midway between the other two (adjacent responses overlap at -1.5 dB). The passband is narrower than the FLAT TOP passband while the stop-band lobes are at a lower level than the UNIFORM passband. Thus it is most useful where frequency resolution is important but high accuracy in amplitude measurement is not. It is often used for measurements of noise-like spectra that have broad distributions of energy.

The window or time record length is automatically selected with selection of the frequency span so the time record will encompass one complete period of the lowest frequency of interest. Passband shaping is implemented digitally by the microprocessor. In all cases, time records are stored without window shaping, and windows can be changed without requiring the gathering of new data.

Broad Spectrum Source

The built-in signal source is a pseudorandom noise generator. The signal originates in a shift register using feedback taps to produce a two-level (binary) waveform that changes from one level to the other in




Fig. 9. Transfer function measurement (center) of the admittance of a 9.995-kHz crystal (left) using the PERIODIC noise source for excitation and the 0-25-kHz frequency span. Because the noise source's spectral lines are matched to the analyzer's, the measurement result is smooth even though averaging is not used. Much greater detail shows up (right) when the measurement is made over a narrower span (50 Hz centered at 10 kHz) with band-selectable analysis, where both the source and analysis spectral lines are more closely spaced.

Window Functions for Spectrum Analysis

by Roger G. Cox

Window functions are frequently used in conjunction with fast Fourier transform (FFT) analysis to minimize two of the limitations imposed by the FFT: a limited time record provided as input to the FFT; and a limited number of output frequency samples generated by the FFT.

The first limitation gives rise to what is known as the uniform or rectangular window. The uniform window always occurs by default since amplitude samples are taken only during a time interval T . The continuous input signal thus appears to be multiplied by the gating function:

$$G(t) = 1 \text{ for } -T/2 < t < T/2 \text{ and } 0 \text{ for } t < -T/2 \text{ and } T/2 < t.$$

The frequency response $G(s)$ is the Fourier transform of $G(t)$:

$$G(s) = \frac{\sin(\pi s)}{\pi s}, \text{ where } s = fT.$$

Any continuous input signal spectrum will be convolved with $G(s)$ such that when sinusoidal signals are input to the analyzer, instead of seeing a single impulse on the frequency display for each sinusoid, one observes a response with the shape of $G(s)$ at each impulse location. This is why the window function in a FFT analyzer can be regarded as the equivalent of the IF passband in a conventional swept-frequency spectrum analyzer.

Window Shapes

The uniform window (Fig. 1) can be thought of as a window consisting of a dc component (=1) multiplied by the gating function $G(t)$. All window functions must include the gating properties of $G(t)$. If any other window shape is desired, it can be constructed by multiplying the uniform window function $G(t)$ by an arbitrary continuous time function consisting of the sum of several cosine functions. For example, the widely-used Hanning window (Fig. 2) consists of a dc term and a single cosine term:

$$H(T) = G(t)(A_0 + 2A_1 \cos 2\pi t), \text{ where } A_0 = 1 \text{ and } A_1 = 0.5.$$

The frequency response of the Hanning window is the result of the convolution of $G(s)$ with three impulses: one at dc and a pair representing $2A_1 \cos 2\pi t$:


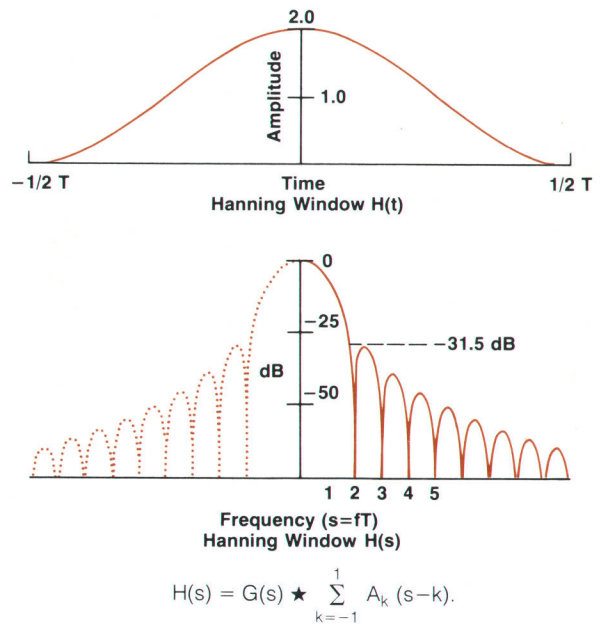


Fig. 1.



To get other frequency response characteristics, more cosine terms can be used. A flat-top passband shape would be desirable because of the second FFT limitation: that of a limited number of frequency sample points being used to evaluate a continuous frequency function. This flat shape is different from




Fig. 3.

the rounded IF response of most conventional spectrum analyzers, where a frequency component being measured is always tuned to the top and center of the IF passband. In a digital analyzer using the FFT, the frequency spectrum is sampled only at discrete intervals. This means that if accurate amplitude measurements are to be made on signals regardless of frequency, the window function's response must be essentially flat over one channel spacing of the FFT ($-1/2 < s < 1/2$). The flat-top window in Model 3582A is flat within 1% over one channel spacing (Fig. 3).

In addition to the flat passband, accurate measurements over a wide dynamic range require that the response roll off rapidly to the maximum dynamic range of the analyzer. The flat-top window in Model 3582A has a rejection of better than 90 dB in the stop band, which occurs only five channel spacings from the center of the window.

To obtain these desirable characteristics, four cosine terms are needed:

$$F(t) = G(t)(A_0 + 2 \sum_{k=1}^4 A_k \cos 2\pi kt)$$

$$F(s) = G(s) \star \sum_{k=-4}^4 A_k \delta(s-k)$$

coincidence with clock pulses, but not necessarily on every clock pulse. Changes from one level to the other occur in pseudorandom fashion, giving the waveform a frequency spectrum that approximates white noise over the analysis band. A low-pass filter with a cutoff frequency much lower than the clock rate limits the frequency band to the flattest part of the spectrum and shapes the output waveform to an approximation of Gaussian noise.


Such a waveform is called pseudorandom because, although locally random in character, the pattern repeats exactly. The spacing between spectral lines is inversely proportional to the length of the pattern (and proportional to clock rate). The 3582A provides the user with a choice of two pattern lengths. With the NOISE SOURCE switch set to PERIODIC, the length of the pattern exactly matches the measurement window length. The spectral lines of the noise source then exactly match the spectral lines of the FFT measurement. The UNIFORM passband can then be used to achieve measurement accuracy without danger of spectrum smearing. Measurements can be made across the entire frequency span without frequency sweeping and without the averaging that is required with true random noise excitation.

If the system has nonlinearities, errors can be introduced when using the PERIODIC noise source because harmonics of the waveform's spectral components fall exactly on other spectral lines. These errors cannot be reduced by averaging. The RANDOM noise source is useful in this situation. Although the RANDOM source is actually periodic and hence pseudorandom, the period of its waveform is over 250,000 times as long as the PERIODIC waveform, so its

Acknowledgments

Ron Potter of Hewlett-Packard's Santa Clara Division provided much valuable help in the derivation of the windows.

Roger G. Cox



Roger Cox joined Hewlett-Packard's Loveland Division in 1970, going to work on synthesizer circuits (one patent) and on a calculator HP-IB interface card. One of the original members of the 3582A design team, he is responsible for the instrument's basic architecture but left the project in 1976 to work on logic state analyzers at the Colorado Springs Division. A native of St. Louis, Missouri, Roger earned a BSEE degree at the University of Wisconsin in 1970. Leisure-time activities include backpacking, ski touring and mountain climbing, and church volunteer work (counselling, visitations, etc.).

include backpacking, ski touring and mountain climbing, and church volunteer work (counselling, visitations, etc.).

spectral lines are far more densely packed and harmonics do not necessarily fall on the spectral lines. Measurement errors due to nonlinearities can therefore be smoothed out by averaging.

Because of the close spacing of the RANDOM source's spectral lines, an additional use for it is in situations where very high-Q resonances may be involved. With the less densely spaced excitation energy lines of the PERIODIC source, a high-Q resonance may fall between lines and the resonance may not be fully excited.

With both types of noise, the noise bandwidth is matched to the selected frequency span. For measurements in the "zoom" mode, the noise source is mixed with the digital oscillator output, converted to analog form, so in all cases the noise spectrum covers the analysis band. Fig. 9 illustrates how this characteristic affects measurement resolution.

Calibrator

Another pseudorandom source within the instrument supplies a calibrating signal to the signal channels when the input SENSITIVITY controls are set to the CAL position. This source generates a pseudorandom sequence 256 clock pulses long that repeats at a 1-kHz rate. This results in a frequency comb with 1-kHz spacing between frequency components and a $\sin x/x$ amplitude distribution. Since the first zero does not occur until 256 kHz, the comb is flat within ± 0.2 dB up to 25 kHz. The calibrating signal is injected into the signal channels immediately after the input attenuators so that all of the amplifiers and filter passbands may be checked.

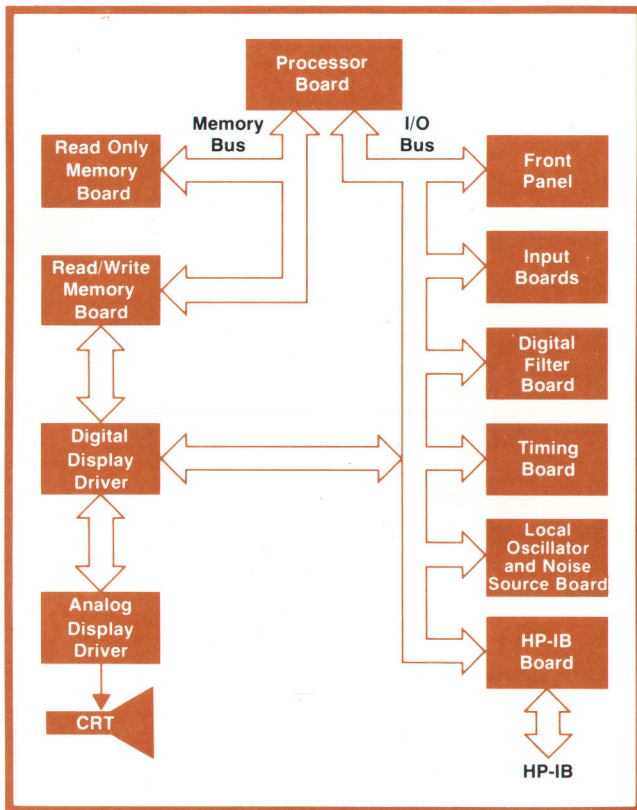


Fig. 10. Block diagram of Model 3582A's control functions.

Instrument Control

A block diagram of the 3582A control functions is shown in Fig. 10. Overall instrument control is handled by the microprocessor on the processor board, which also does the FFT calculations and display conditioning. This is the same 16-bit microprocessor that was developed for the Model 9825A Desktop Computer.⁷

General control of the instrument is accomplished over the processor I/O bus, a 16-bit, parallel, bidirectional bus with four additional control lines. The processor can interrogate the status of the various boards over this bus and program them according to the operating mode of the instrument.

Under local (front-panel) control, the processor interrupts its signal and display processing tasks about 10 times per second to read the status of the front-panel switches. When changes are detected, signal processing may continue with changed parameters, or may be discontinued and restarted depending on the nature of the change. The 10-Hz polling rate is fast enough to enable the instrument to appear to respond instantaneously to switch changes.

Under remote control via the HP-IB, the only change in the control process is that the processor no longer interrogates the front panel. Instead, the HP-IB software routines interrogate the HP-IB board and set internal switch status registers to correspond to the

programmed switch positions.

To be able to offer a high degree of HP-IB control, almost none of the front-panel switches are wired directly to the functions they control. For example, the rotary input sensitivity switches are not connected to the attenuator networks on the input boards. Rather, they are connected to a microprocessor I/O bus port, and the attenuators are switched by means of reed relays via a different port. In this way, the processor can program the input attenuators to agree with either the front-panel switches or the HP-IB selected sensitivity.

The processor board is connected to the ROM and RAM boards by the memory bus, a bidirectional, 16-bit, parallel bus that is separate from the microprocessor I/O bus. The ROM and RAM memories are distinguished only by different high-order address bits.

The bulk of the ROM memory is stored in 16 mask-programmed ROM's, each storing 2K 8-bit bytes. To allow for last minute program improvements that inevitably arise during development of any new program of this size, room was left in the ROM board for four "patch" ROM positions that can accept different types of fuse-link programmable ROMs, depending on jumpers on the board. The 16 mask-programmed ROMs cover the address space from 010000₈ up while the addresses below 010000₈ are assigned to the patch ROM positions.


The software is organized so that small program blocks are linked by addresses stored in the patch ROM. Most of the rest of the patch ROM address space is empty. Hence, to change a program block, it was only necessary to change the address link pointing to that block and write the new block in the blank space in patch ROM. This made it possible to begin instrument production while software revisions were being made without the delays that otherwise would have resulted from waiting for the generation of new masks. Thus, the first production instruments were shipped with the correct software but without the need to equip the entire memory with expensive programmable ROMs.

The RAM memory has 4 K 16-bit words stored in 16 dynamic RAM chips, each organized as 4K 1-bit words. Memory refresh is not performed by the RAM board itself, but rather by the interaction between the RAM board and the digital display driver.

The RAM board is built as a two-port memory, with one port serving the processor over the memory bus, and the other port serving the digital display driver. To present a display, the processor formats both graphic and alphanumeric information into a special area of RAM memory. The digital display driver then reads these words, separates graphic from alphanumeric information, and draws this display on

the CRT using the analog display driver. Since the rate at which the digital display driver reads words from the RAM board is sufficient to keep the memory refreshed, no additional refresh circuitry is needed on the RAM board.

Acknowledgments

Cullen Darnell originally conceived of the Model 3582A Spectrum Analyzer, and project manager Larry Whatley contributed to the overall design with many helpful ideas. Besides those mentioned elsewhere in this issue, others involved in the design were Tomoyuki Akiyama, Mike LaMothe, Jerry Metz, Charlie Potter, and Jim Saar. Industrial design was by Jim Berry. Gary Heimbigner and Chuck Kingsford Smith contributed marketing and applications research. 

References

1. S.N. Holdaway and M.D. Humpherys, "The Next Generation RF Spectrum Analyzer," Hewlett-Packard Journal, June 1978.
2. R. Temple, "A Precision, Digitally-Controlled Spectrum Analyzer for the 20-Hz-to-40-MHz Frequency Range," Hewlett-Packard Journal, June 1978.
3. N. Pendergrass, "Coherence Measurements," Electronics, to be published 14 September 1978.
4. "Measuring the Coherence Function with the HP 3582A Spectrum Analyzer," Hewlett-Packard Application Note number 245-2.
5. R.H. Grote and H.W. McKinney, "Advanced Digital Signal Analyzer Probes Low-Frequency Signals with Ease and Precision," Hewlett-Packard Journal, October 1977.
6. H.W. McKinney, "Band-Selectable Fourier Analysis," Hewlett-Packard Journal, April 1975.
7. W.D. Eads and D.S. Maitland, "High-Performance NMOS LSI Processor," Hewlett-Packard Journal, June 1976.



Nixon A. Pendergrass

Nick Pendergrass earned a BSEE degree at the University of Missouri at Rolla in 1967, and an MSEE degree in 1969 at Purdue University (Indiana) while serving as a graduate teaching assistant. He then joined the U.S. Air Force, graduating from Officer Training School in 1969. While with the Air Force, he was involved in systems analysis of communications systems. Nick returned to the campus in 1972 as a research assistant at the University of California at Berkeley, where he earned a PhD degree in 1975 (his thesis involved digital filters). He joined HP that same year. Married, and with one child, Nick is interested in photography and outdoor sports, including cycling, backpacking, and skiing.

in 1975 (his thesis involved digital filters). He joined HP that same year. Married, and with one child, Nick is interested in photography and outdoor sports, including cycling, backpacking, and skiing.



John S. Farnbach

John Farnbach obtained a BSEE degree at Princeton University in 1965. He designed special-purpose instrumentation while earning an MSEE degree at the University of New Mexico. Returning to Princeton in 1969, he obtained a PhD degree (1972) and then taught electrical engineering (in Spanish) for a year and a half at the Universidad Iberoamericana in Mexico. He then joined the Seismological Lab at the University of Uppsala, Sweden, where he did research on seismic signal processing. He came to HP in 1975, going to work on FFT software. Married, and with two small children, John spends off hours renovating the older home in which they live, in picnicking and skiing, and playing volleyball in regional competitions.

processing. He came to HP in 1975, going to work on FFT software. Married, and with two small children, John spends off hours renovating the older home in which they live, in picnicking and skiing, and playing volleyball in regional competitions.

Hewlett-Packard FFT Analyzers

The Model 3582A Spectrum Analyzer described in the preceding is the newest member of Hewlett-Packard's family of FFT-based analyzers. The others are the Model 5420A Digital Signal Analyzer and the Model 5451C Fourier Analyzer.

The Model 3582A Spectrum Analyzer provides a low-cost means of obtaining the advantages of FFT processing for signal analysis. It can measure spectrum amplitude and phase, transfer functions, and coherence with frequency resolution as fine as 0.02 Hz over the entire 25-kHz frequency range. It is optimized for simplified operation, small package size and weight, and low cost while providing a powerful set of measurement capabilities. It includes many convenience features, such as a choice of three windows and a tracking noise source that reduces the averaging time usually required with random-noise sources.

The Model 5420A Digital Signal Analyzer, described in the

October 1977 issue of the Hewlett-Packard Journal, has additional measurement capabilities. It was designed for those who need auto- and crosscorrelation functions, power and cross spectra, and statistical analyses of signals, in addition to frequency-domain measurements. Pre/post-trigger delay is provided for acoustic, transient, and transfer function measurements, and the analyzer may be operated from an external clock to adapt it to rotating machinery studies.

Model 5420A also provides very flexible display formatting that can give Bode, Nyquist, and Nichols plots as well as linear plots of results. Digital storage for measurement results is provided by its built-in magnetic tape unit. Post-processing capabilities include such operations as integrating an acceleration spectrum to a velocity spectrum, or calculating coherent output power and open-loop gain. It has a built-in true random noise source for stimulus-response measurements. The

5420A's measurement frequency range extends from 25 kHz down to 0.008 Hz and its frequency resolution can be as fine as 40 μ Hz. Its dynamic range is 75 dB.

The Model 5451C Fourier Analyzer, a further elaboration of the Model 5451A described in the June 1972 issue of the Hewlett-Packard Journal, offers a very flexible and powerful alternative that is capable of a full range of analysis functions, including the inverse Fourier transform. It is a fully calibrated, multipurpose, computer-based system and can be obtained with a full range of computer peripherals such as disc memories. This combined with user-written programs allows

each system to be tailored to specific applications. It can also be obtained with up to four input channels and options for such tasks as vibration control, modal analysis, and rotating-machinery signature analysis. Its frequency range extends from dc to 50 kHz and optionally to 100 kHz, and its dynamic range is 75 dB.

These three analyzers cover a wide range of applications in studies concerning mechanical vibrations, structural dynamics, acoustics, underwater sound, communications, and other investigations involving audio and subaudio frequencies.

SPECIFICATIONS

HP Model 3582A Spectrum Analyzer

Frequency

FREQUENCY MODES:
 0-25 kHz SPAN: Measurement is performed over fixed frequency range of 0 Hz to 25 kHz independent of FREQUENCY SPAN control.
 0-START: Measurement is performed over frequency range defined by FREQUENCY SPAN control with fixed start frequency of 0 Hz.
 SET CENTER: Measurement is performed over a frequency range with width determined by FREQUENCY SPAN control and center frequency set by ADJUST control with 1-Hz resolution.
 SET START: Measurement is performed over a frequency range with width determined by FREQUENCY SPAN control and start frequency set by ADJUST control with 1-Hz resolution.

FREQUENCY RANGE: 0.02 Hz to 25.5 kHz.

FREQUENCY SPANS:
 0 START MODE: 1 Hz full scale to 25 kHz full scale in 1-2.5-5-10 sequence.
 SET START OR SET CENTER MODE: 5-Hz span to 25-kHz span in 1-2.5-5-10 sequence.

FREQUENCY ACCURACY: $\pm 0.003\%$ of display center frequency.

FREQUENCY RESOLUTION: Marker resolution is equal to calculated point spacing for selected frequency span and number of channels (see Table).

FILTER PASSBAND SHAPE:

	Flat Top	Hanning	Uniform
3-dB Bandwidth (single-channel):	$(1.4 \pm 0.1)\%$ of span	$(0.58 \pm 0.05)\%$ of span	$(0.35 \pm 0.02)\%$ of span
Shape Factor: [60-dB bandwidth / 3-dB bandwidth]	$2.6 \pm .1$	$9.1 \pm .2$	716 ± 20

FLAT TOP PASSBAND SHAPE provides optimum amplitude accuracy. UNIFORM PASSBAND SHAPE is optimized for use with transients and with the PERIODIC NOISE SOURCE. HANNING PASSBAND SHAPE provides an amplitude-frequency resolution compromise for general noise measurements.

SINGLE-CHANNEL ANALYSIS PARAMETERS:

Frequency Span	Time Record Length (Δt)	Calculated Point Spacing (Δf)	Equivalent Noise Bandwidth		
			Flat Top	Hanning	Uniform
1 Hz	250 s	.004 Hz	14.5 mHz	6.00 mHz	4.00 mHz
2.5 Hz	100 s	.01 Hz	36.3 mHz	15.0 mHz	10.0 mHz
5 Hz	50 s	.02 Hz	72.6 mHz	30.0 mHz	20.0 mHz
10 Hz	25 s	.04 Hz	145 mHz	60.0 mHz	40.0 mHz
25 Hz	10 s	.1 Hz	363 mHz	150 mHz	100 mHz
50 Hz	5 s	.2 Hz	726 mHz	300 mHz	200 mHz
100 Hz	2.5 s	.4 Hz	1.45 Hz	600 mHz	400 mHz
250 Hz	1 s	1 Hz	3.63 Hz	1.5 Hz	1.00 Hz
500 Hz	500 ms	2 Hz	7.26 Hz	3.00 Hz	2.00 Hz
1 kHz	250 ms	4 Hz	14.5 Hz	6.00 Hz	4.00 Hz
2.5 kHz	100 ms	10 Hz	36.3 Hz	15.0 Hz	10.0 Hz
5 kHz	50 ms	20 Hz	72.6 Hz	30.0 Hz	20.0 Hz
10 kHz	25 ms	40 Hz	145 Hz	60.0 Hz	40.0 Hz
25 kHz	10 ms	100 Hz	363 Hz	150 Hz	100 Hz

Corresponding dual-channel parameters are found by doubling calculated point spacing and equivalent noise bandwidth and taking one half the time record length.

Amplitude

MEASUREMENT MODES: 256-point amplitude spectra are measured in single-channel mode. Two 128-point amplitude spectra are measured in dual-channel mode.

DISPLAY MODES:


LOG: 10 dB/major division or 2 dB/major division
 LINEAR: Constant voltage/major division

AMPLITUDE MEASUREMENT RANGE:

LOG: Calibrated attenuator range is +30 dBV to -50 dBV single-tone rms maximum input level in 10-dB ± 0.2 dB steps. Continuous vernier provides >10 dB of uncalibrated attenuation between steps.
 LINEAR: Calibrated attenuator range is +30 V rms to 3 mV single-tone rms maximum input in 1-3-10 sequence. Vernier provides continuous coverage between steps. AMPLITUDE REFERENCE LEVEL provides 8 additional ranges down to 8 μ V full scale.

DYNAMIC RANGE:

DISTORTION PRODUCTS: >70 dB below maximum input level.
 SPURIOUS RESPONSES: >70 dB below maximum input level.



DC RESPONSE: Adjustable to >40 dB below maximum input level with front-panel dc balance adjustment.

AMPLITUDE ACCURACY (Overall accuracy is the sum of accuracy at passband center and filter accuracy):
 ACCURACY AT PASSBAND CENTER: ± 0.5 dB.
 FLAT TOP FILTER: $+0, -0.1$ dB.
 HANNING FILTER: $+0, -1.5$ dB.
 UNIFORM FILTER: $+0, -4.0$ dB.

RESOLUTION:

LOG: 0.1 dB with marker.
 LINEAR: 3 digits with marker.
LINEARITY: ± 0.2 dB $\pm 0.02\%$ of full scale.
AMPLITUDE CALIBRATOR: Internal calibration signal is a line spectrum with nominal 1-kHz frequency spacing and fundamental level of 22 ± 0.2 dBV on log scales and 20 ± 0.5 volts on linear scale.

OVERLOAD LIMITS:

LOG: Overload occurs at 100% of maximum input level which is equal to full scale when AMPLITUDE REFERENCE LEVEL is set to NORMAL. When overload occurs, spurious products may be displayed.

LINEAR: Overload occurs at 100% of maximum input level which, depending on input attenuator setting, is at 6/8 or 5/8 of full scale when AMPLITUDE REFERENCE LEVEL is set to NORMAL. When overload occurs, spurious products may be displayed.

AVERAGING MODES:

RMS: For each calculated frequency point, displayed amplitude is $\sqrt{(1/N) \sum A_i^2(f)}$ and phase is $(1/N) \sum \phi_i(f)$.
PEAK: For each calculated frequency point, displayed amplitude is MAX $A_i(f)$ and phase is corresponding value of retained amplitude point.
TIME: For each time record point the amplitude is $(1/N) \sum A_i(t)$.
 The averaged time record is transformed to give corresponding amplitude and phase.

NUMBER OF AVERAGES: 4 to 256 in binary sequence, and exponential. Exponential averaging in rms mode gives running average with new spectral data weighted 1/4 and previous result by 3/4. Exponential in peak mode gives continuous peak hold operation.

Phase

MEASUREMENT MODES: 256-point phase spectra are measured in single-channel mode. Two 128-point phase spectra are measured in dual-channel mode.

DISPLAY RANGE: From $+200$ degrees to -200 degrees.

ACCURACY: ± 10 degrees.

RESOLUTION:

DISPLAY: 50 degrees/major division.
 MARKER: 1 degree.

Transfer Function

MEASUREMENT MODES: Dual-channel 128-point transfer functions are measured.

DISPLAY MODES:

LOG AMPLITUDE: 10 dB/major division and 2 dB/major division.
 LINEAR AMPLITUDE: Constant floating-point fraction/major division.
 PHASE: Constant 50 degrees/major division.

MEASUREMENT RANGE:

LOG AMPLITUDE: Calibrated ranges of $+160$ dB full scale to -80 dB full scale in 10-dB steps. Uncalibrated verniers provide continuous coverage between steps.
 LINEAR AMPLITUDE: Calibrated ranges of 4.0×10^6 full scale to 4.0×10^{-8} full scale in factor-of-10 steps. Uncalibrated verniers provide continuous coverage between steps.

PHASE DISPLAY RANGE: $+200$ degrees to -200 degrees.

ACCURACY:

AMPLITUDE: ± 0.8 dB.
 PHASE: ± 5 degrees.

RESOLUTION:

LOG AMPLITUDE: 0.1 dB with marker.
 LINEAR AMPLITUDE: 3-digit scientific notation with marker.
 PHASE: 1 degree with marker.

Coherence Function

MEASUREMENT MODE: Dual-channel 128-point coherence functions are measured with rms averaging only.

RANGE: Bottom display line is 0.0 and top line is 1.0.

RESOLUTION:

DISPLAY: 0.125/major division.
 MARKER: 0.01

General

INPUT IMPEDANCE: $10^6 \Omega \pm 5\%$ shunted by ~ 60 pF from input high to low for less than 75% relative humidity.

DC ISOLATION: Input low may be connected to chassis ground or floated up to 30 volts to reduce effects of ground loops on the measurement.

INPUT COUPLING: ac or dc. Low frequency 3-dB roll off of ac coupling is <1 Hz.

COMMON MODE REJECTION:

50 Hz: >60 dB.
 60 Hz: >58 dB.

INPUT CHANNEL CROSSTALK: <-140 dB between channels with 1 k Ω source impedance driving one channel with other channel terminated in 1 k Ω .

TRIGGER MODES:

FREE RUN: New measurement is initiated by completion of previous measurement.
EXTERNAL: Rear-panel switch allows new measurements to be initiated by external TTL pulses.
INPUT SIGNAL: New measurement is initiated when input signal meets specified trigger condition.

TRIGGER CONDITIONS:

SIGNAL CONDITIONS: Triggering can be selected to occur on positive or negative going transitions through trigger level. Trigger level is adjusted between time record overload limits by continuous vernier.

SINGLE/MULTIPLE TRIGGERS: Single-shot triggering is specified by taking instrument out of REPEAT mode. ARM control sensitizes instrument to take another measurement in the single shot mode.

OUTPUT SIGNALS:

X-Y RECORDER:

VERTICAL: 0 to 5.25V $\pm 5\%$.

HORIZONTAL: 0 to 5.25V $\pm 5\%$.

IMPEDANCE: 1 k Ω

PEN LIFT: Contact closure during sweep.

NOISE SOURCE:

PERIODIC: Pseudorandom noise signal with spectral line spacings that match calculated point spacing for selected frequency span. Frequency response is ± 1 dB with UNIFORM passband shape and -10 dBV level.

RANDOM: Random noise signal; noise spectrum is band limited and band translated to match selected measurement.

LEVEL: Vernier control from <10 mV to >500 mV rms into load of $\geq 50 \Omega$.

IMPEDANCE: $<2 \Omega$.

IMPULSE SOURCE: TTL low-to-high pulse with period equal to time record length.

PROGRAMMING: All analyzer front panel controls except CRT controls, NOISE SOURCE LEVEL and TYPE, TRIGGER LEVEL, AMPLITUDE VERNIERS, and GROUND ISOLATION are remotely programmable via interface bus (HP-IB).

DATA INPUT: Time records, amplitude and phase spectra, coherence functions, transfer functions and alphanumeric text can be input to the analyzer via HP-IB.
DATA OUTPUT: Time records, amplitude and phase spectra, coherence functions, transfer functions alphanumeric text, marker values and control settings can be output via HP-IB.

ENVIRONMENTAL:

OPERATING TEMPERATURE: 0°C to $+55^\circ$ C.

HUMIDITY: To 95% relative humidity at 40°C.

OPERATING ALTITUDE: 4600 meters (15,000 feet).

SHOCK: 30 G, 11 ms half sine wave on each of six sides.

VIBRATION: 10 Hz to 55 Hz at 0.010 inch peak-to-peak excursion.

OPERATING POWER: Switch selection of 110V or 230V $+5, -10\%$, 48-66 Hz; less than 150 VA.

SIZE: 188 mm H \times 425 mm W \times 552 mm D (7 1/2" \times 16 1/2" \times 21 3/4" in.).

PRICE IN U.S.A.: \$10,000.

MANUFACTURING DIVISION: LOVELAND INSTRUMENT DIVISION
 P.O. Box 301
 815 Fourteenth Street S.W.
 Loveland, Colorado 80537 U.S.A.

Designing Programmable Digital Filters for LSI Implementation

by Lynn A. Schmidt

THE POWER OF DIGITAL SIGNAL PROCESSING now becoming available through the use of large-scale integrated circuits is giving low-frequency spectrum analyzers unprecedented flexibility and computational power at relatively low cost. For example, the HP Model 3582A, described in the preceding article, derives much of its signal-processing capabilities from four single-chip LSI digital filters. This article describes how it was possible to implement these filters on LSI chips.

The arithmetic, storage, and speed requirements of the digital filtering proposed for the Model 3582A presented formidable technical challenges. The requirement for doing three to six million multiplications per second far exceeded the capabilities of available microprocessors. Additional complications were presented by the need for 18- to 21-bit precision in the computations. Thus, a dedicated hardware signal processor of some kind appeared to be the only viable solution.

This processor could have been built with off-the-shelf components, but the cost in terms of component count, board space, power, and reliability would have been high. Implementing the processor by large-scale integration (LSI), however, promised more than a 80% hardware savings, coupled with lower power and higher reliability from the use of fewer components. The development of a custom NMOS digital filter chip was therefore undertaken.

The underlying problem was how to design the filter processor with enough power to do the job within the confines of a reasonably sized chip.

Digital Filtering in Spectrum Analysis

The term "digital filter" refers to a computational process by which a sequence of numbers, usually samples of an analog signal, is transformed into a second sequence of numbers (see box, page 17). In a spectrum analyzer this process corresponds to either low-pass or bandpass filtering.

Digital filtering in a spectrum analyzer also provides a practical means of implementing band-selectable analysis or "zoom" to analyze the spectrum of a narrow frequency band centered on any frequency within the range of the analyzer. The translation of the frequency band to baseband needed for

band-selectable analysis is possible by analog techniques but is often impractical because of the high cost of ensuring good gain and phase match between the real and imaginary signal channels. Two or more digital filters, on the other hand, can be matched in gain and phase exactly, regardless of temperature, aging, or production variations.

Span Control

Model 3582A uses a particular type of digital filter called a decimation filter, a type that plays a fundamental role in analyzers based on the fast Fourier transform (FFT).

As explained in the box on page 7, the first step in FFT processing is the collection of N equally-spaced-in-time samples of the analog input during a measurement time interval T . The N samples are used in computing the discrete Fourier transform of the input signal. The result of the transform can be likened to the response of a bank of N narrow bandpass filters spaced at center frequencies of $1/T$ Hz. The responses of these filters are plotted on the analyzer's CRT to generate the amplitude-vs-frequency display.

An important characteristic of the FFT frequency analysis is that the frequency span width and resolution is directly proportional to the sampling rate ($f_s = N/T$). In accordance with sampled-data theory, the input signal must be band limited to less than $f_s/2$ to avoid aliasing but because of practical filter limitations, the input signal bandwidth is usually limited to something between $f_s/4$ and $f_s/3$.

Spectrum analyzers need to change their frequency span over a broad range to accommodate a variety of input signals. Model 3582A, for example, has fourteen spans ranging from 1 Hz to 25 kHz in a 1-2.5-5 sequence. Providing fourteen sample rates poses no problem, but providing an analog anti-aliasing filter with fourteen programmable bandwidths does. In the Model 3582A, digital filters following the A-to-D converter limit the bandwidth to the frequency span desired. The input signal to the A-to-D converter is band-limited to 25.6 kHz by a fixed, analog, low-pass filter regardless of the frequency span desired, and samples of the input are taken at a fixed rate of 102.4 kHz. Sample words at the output of the digital filters are selectively discarded to effect a resampling at a




Fig. 1. Specifications for the band-limiting filters to be used in the Model 3582A Spectrum Analyzer. The filter is to have 14 programmable cut-off frequencies (f_c) ranging from 25 kHz down to 1 Hz.

lower rate of four times the desired frequency span. This process of simultaneous digital low-pass filtering and sample rate reduction is called decimation filtering. Only the band to be analyzed is presented to the FFT processor. Since the resample rate tracks the span width, the FFT algorithm remains the same regardless of the selected span.

Choosing a Filter

The basic requirements for the digital filters in the Model 3582A are shown in Fig. 1. The passband corner frequency, since it sets the frequency span of the analysis, needs to be variable from 25 kHz (no filtering) down to 1 Hz in a 25-10-5-2.5 sequence. To provide proper anti-aliasing, the stopband needs to start at three times the cut-off frequency and should be at least -80 dB relative to the passband to assure that aliased signals are suppressed to an undetectable level. The passband ripple, however, can be on the order of ± 0.5 dB since the effects of ripple can be compensated for by weighting the spectral frequency components after the FFT analysis. This compensation is accomplished simply by multiplying the spectrum by the inverse of the ripple function. Phase nonlinearities introduced by the filter can be removed in an identical fashion.

It soon became obvious that it was impractical to meet these requirements with a single filter. The higher reduction ratios (input bandwidth/output bandwidth) that would be needed for the narrow span widths would require coefficient word lengths greater than 30 bits if a recursive infinite-impulse-response (IIR) low-pass filter were used. If a non-recursive finite-impulse-response (FIR) filter were used, the number of coefficients required would be greater than 100,000.

Coefficient words get very long with IIR filters because the poles of the transfer function in the z plane group closer and closer to the unit circle at $z = +1$ as the bandwidth is reduced. The relative positions of these poles must be maintained precisely to preserve both stability and the desired transfer function. With finite arithmetic, the poles can only assume certain quantized positions so, to obtain the desired transfer function as the bandwidth is made narrower, the precision of the arithmetic needs to be increased.

With FIR low-pass filters, the transfer function is determined by zeros in the z plane. The order of the filter, or equivalently the length of its impulse response, grows in size inversely with bandwidth. The order (number of multiplication coefficients) required for the 1-Hz bandwidth is about 10,000 times that required for the 10-kHz bandwidth.

LSI chip area relates directly to the complexity of the processing algorithms for the LSI filters, so a reduction in processing complexity was a major goal for this project. The solution was to implement the filter as a cascade of several filters with interstage decimation of the sample rate, a common practice with decimation filtering. Each filter stage then needs to accomplish only a modest sample rate reduction with a consequent reduction in the complexity of the filter




Fig. 2. Interleaved processing scheme for the digital filter. Two complete filter computations are performed during each input sample period: one for the first filter stage, and one for one of the downstream filters according to this example schedule.

stage. In addition, the downstream filters in the cascade compute at reduced sample rates, further simplifying computational requirements.

Interleaved Multiplications

In the final filter design, a cascade of eight filters in series is used with the output selected from any of the eight. Each filter stage is allowed to compute either a decimation by two or a decimation by five, giving overall decimation ratios of $2^m \times 5^n$ where $m + n \leq 8$. This easily satisfies the requirement for the 14 frequency spans with decimation ratios ranging from 1 to 25000.

An interesting aspect of this filter structure is found upon examination of the total computational requirements. Consider the case where all eight filters are each performing a decimation by 2. The first filter computes at the input sample rate. The second filter, however, needs to compute at only one-half the rate because of the sample rate reduction. In like manner, the third filter computes at one-fourth the rate, the fourth at one-eighth the rate, and so on, forming a geometric series, $1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots$, that converges on 2. This means that the total computational requirement of all eight filter stages can be satisfied by a single processor operating at twice the rate required by the first filter stage. Actually, hardware is needed for only one filter stage. The processor has eight channels of memory to store intermediate states of the filters but only one set of arithmetic hardware. Fig. 2 illustrates how the processing steps are scheduled so the eight filter stages can timeshare the hardware.

IIR or FIR?

Much has been written in the literature about the virtues of both IIR and FIR filters. The primary differences between these two classes of filters as far as decimation is concerned are listed below in Table 1.

Evidently the advantages of FIR filters have resulted in more decimation filters being implemented with FIR approaches than IIR, but in our case the necessity to implement the filters in LSI required some special considerations.

Table 1

FIR	IIR
Only need to compute at output rate	Must compute every input sample
Linear phase	Not linear phase
More storage	Less storage
4-8 multiplications/input	8-12 multiplications/input
N=16 to 38/per stage	5th order structure/per stage
More coefficients	Fewer coefficients
Less noise/bit	More noise/bit

In making the choice between FIR and IIR, let us consider the elements that require chip area in an LSI digital filter. By far the largest consumer of chip area would be the multipliers. For example, a 12×16 -bit multiplier requires nearly 10^6 square micrometres of chip area. A second consumer is memory. Dynamic memory takes less space than static, and serial shift registers take much less space than random access memory.

A third large consumer of chip area is interconnect. We have found that about 40% of the area of most LSI chips is devoted to simply interconnecting the various active devices. Serial arithmetic structures bring an advantage here since only a single interconnect is needed to transfer a 24-bit word. This, however, is at the expense of the total transfer time.

We chose to use serial arithmetic to minimize interconnects, and dynamic shift-register memory. However, the key element by far in our design strategy became the multipliers and associated coefficients.

What is a Digital Filter?

Digital filtering is a computational process or algorithm by which a sampled signal or sequence of numbers, acting as an input, is transformed into a second sequence of numbers called the output. The computational process may correspond to high-pass, low-pass, bandpass, or bandstop filtering, integration, differentiation, or something else. The second sequence can be used for further processing, as in a fast-Fourier-transform analyzer, or it can be converted to an analog signal, producing a filtered version of the original analog signal.

The digital approach offers many advantages over analog approaches:

- Changes resulting from variations in component values—normally associated with filter capacitors and resistors as a result of temperature or aging—are non-existent.
- Periodic calibration is eliminated.
- The performance from unit to unit is stable and repeatable.
- Great flexibility is available since filter response can be al-

tered by changing arithmetic coefficients.

- Arbitrarily high precision can be achieved, limited only by the number of bits involved.
- Very small size, low power, and low cost are possible by large-scale integration.

An Example

Let us consider the following practical situation: the need to reject very strong 60-Hz interference contaminating a 10-Hz signal of interest. The obvious solution is to build a bandstop or notch filter that rejects the 60-Hz component while passing the 10-Hz signal with little or no alteration.

First a familiar analog RC notch filter will be examined. A commonly used RC notch filter network is the "twin-T" shown in Fig. 1 (next page). This familiar network derives its bandstop characteristics from a pair of transmission zeros located in the s plane on the imaginary axis at 60 Hz. It also has two poles on

the negative real axis, but they hardly affect the bandstop characteristics. Applying the waveform shown in color in Fig. 1 to the notch filter results in the output shown in black. As can be seen, the twin-T circuit rejects all of the 60-Hz interference and leaves the desired 10-Hz signal.

Changing the center frequency of the notch requires changes in the values of all three capacitors, all three resistors, or the two shunt elements, R' and C' , together. The position and depth of the notch is very sensitive to parameter changes in the components of the twin-T, so precision components must be used in building the filter if good rejection at exactly 60 Hz is to be maintained.

The Digital Filter

The digital filter solution to this problem is simple yet just as effective as the linear filter and is considerably more flexible. In this case the input signal is sampled by an analog-to-digital converter at a rate of 500 samples per second. That is, the input signal level is measured at intervals of 2 ms and the result forms the input sequence to our digital filter.

Fig. 2 shows the digital counterpart to the twin-T notch circuit. This filter is made using two registers, a digital multiplier, and a 3-input adder. As the difference equation in Fig. 2 shows, the current output sample word $[y(n)]$ is formed by summing the current input sample word $[x(n)]$ with the previous input $[x(n-1)]$ multiplied by the constant -1.46 , and with the second previous input $[x(n-2)]$. When this filter is fed sample words at

a rate of 500 samples/s, 60-Hz components in the analog input signal will be eliminated from the regenerated analog output signal.

Just as the Laplace transform can be used to analyze linear filters, the z transform can be used to analyze digital filters. The filter shown in Fig. 2 derives its bandstop characteristics from a pair of transmission zeros located in the z plane on the unit circle at a point corresponding to 60 Hz. No poles are indicated in this transfer function, but these can be implemented when desired.





Fig. 2 also shows two cycles of the sampled 10-Hz input signal (in color) which is contaminated as before by 60-Hz interference. Passing this input sequence through the digital bandstop filter results in the output sequence shown in black. It is evident that the filter has rejected all of the 60-Hz interference leaving the desired 10-Hz signal. (It is a simple exercise on a calculator or computer to generate this input signal and compute the output sequence $y(n)$ according to the difference equation shown in Fig. 2).

One advantage that the digital filter has in this application is that the position and depth of the bandstop notch will remain constant and will not drift with temperature or age.

The notch center frequency can be changed in two ways. The first way is simply to change the multiplication coefficient. For example, changing the -1.46 coefficient to -1.62 changes the notch center frequency to 50 Hz.

The second way is to change the sample rate. In this particular example, the notch center frequency is 12% of the sample rate. Reducing the sample rate to 417 Hz reduces the notch center to 50 Hz.

In summary, for certain applications digital filters offer advantages in terms of flexibility, stability, and simplicity of control when compared to their analog counterparts.



$$y(n) = x(n) - 1.46x(n-1) + x(n-2)$$

$$H(z) = \frac{Y(z)}{X(z)} = 1 - 1.46z^{-1} + z^{-2}$$


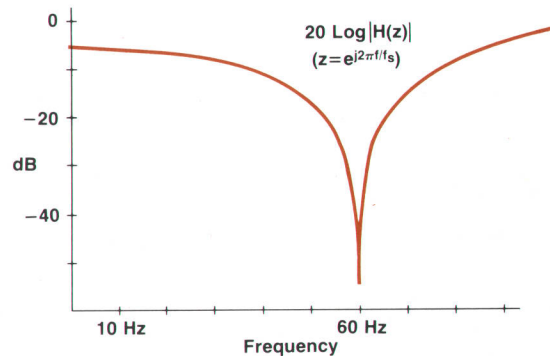


Fig. 2.




Fig. 3. Comparison of discrete values between 0 and 2 available with 7-bit binary code and with canonical-signed-digit code limited to three or fewer non-zero bits.

With reference to Table I, we found that FIR filters require fewer multiplications but these involve more coefficients. IIR filters require less storage but longer coefficients. Linear phase was not of prime importance because the analyzer would have the ability to correct for both gain and phase after the FFT analysis.

To achieve the required multiplication rate with NMOS circuits, several multipliers operating in parallel would have to be included on the chip. Minimizing the number of coefficients that each multiplier had to work with would further simplify the circuitry. Some further space-saving developments with respect to IIR coefficients, to be described next, occurred. Thus an IIR filter design was chosen.

Coefficients

The usual binary coefficients can be represented by the expression:

$$X = \sum_{j=0}^{B-1} x_j 2^j,$$

a weighted sum of B powers of 2 where the weights, x_j , or bits as we know them, can take on the values 0 or 1.

Another method of representing coefficients, called canonical signed digits (CSD), allows the bits to take on the values 0, 1, and $\bar{1}$ (-1). During the 1950's in the early days of computers, this code was discussed in many publications with reference to fast multiplication.¹

It has been shown that for any integer X, there exists a unique representation in CSD code in which no two consecutive bits are non-zero.¹ Furthermore, the CSD representation has the least number of non-zero bits. For example, decimal 31 becomes 011111 in binary but in CSD it is 10000 $\bar{1}$ —only two non-zero bits as compared to five in binary.

It can be shown that CSD representation generally requires 33% fewer non-zero bits than binary. This bit minimization feature of CSD was very important to the realization of the LSI filter. Consider the operation of multiplication. Although there are many multiplication algorithms, the fundamental technique involves shifting and adding. CSD requires fewer additions than straight binary because an add is needed only for the non-zero bit positions in the multiplier word. Although using a CSD multiplier and a binary

multiplicand requires a capability for subtraction, this is just as easy to do as addition. (These techniques result in mathematical operations closely related to the well known Booth's algorithm for multiplication.)

This bit-minimization feature along with further developments to be discussed allowed the implementation of separate multiplier circuits for each of the 12 coefficients in the filter structure.

Modified CSD's

Further reductions in multiplier complexity are possible by limiting the number of non-zero bits in each CSD word to three, a representation we chose to call sparse canonical signed digits. This means, of course, that certain values cannot be represented. However, it turns out that this is not a severe restriction. Fig. 3 indicates all the CSD numbers with three or fewer non-zero bits available between the values of 0 and +2 compared to 7-bit binary (fractional powers of two are present and restricted to no smaller than 2^{-6}). An identical set exists for the range -2 to 0. Most IIR digital filters are designed using second-order sections having coefficient values between +2 and -2, a result of the necessity to keep the poles and zeros on or within the unit circle in the z plane. Thus, very few coefficient values would be missed by using




Fig. 4. Second-order canonic digital filter section. The boxes labelled z^{-1} are storage registers that delay output sample words by one sample period.




Fig. 5. (a) Discrete poles realizable with 8-bit canonical-signed-digit code limited to three non-zero bits in the filter of Fig. 5. (b) Poles realizable with CSD code limited to three non-zero bits when nine shift positions are allowed.

CSD code restricted to three non-zero bits.

Let us now examine the distribution of poles available with sparse CSD in a filter design. Fig. 4 indicates the configuration of a second-order canonic digital filter section (canonic meaning that this structure has the least number of memory elements with respect to other second-order structures; it also involves the least number of multiplications).

The pole and zero locations can be found by substituting all possible sparse CSD values for coefficients a and b in the equations of the z -transform transfer function in Fig. 4. Fig. 5a shows the location of the poles in the first quadrant resulting from this substitution. Since this analysis simply involves finding the root positions of a second-order polynomial, Fig. 5a also shows the location of available zeros. Zeros useful for shaping the stop band of the filter all fall on the unit circle. The missing bands correspond to the missing digits in sparse CSD, as indicated in Fig. 3.

In this analysis the number of non-zero digits was restricted to three and the maximum shift positions to eight, giving the least-significant digit a weight of $\pm 2^{-7}$. The density of pole locations can be increased by allowing more shift positions. Fig. 5b shows the result of allowing an additional shift position but still only three non-zero bits, where the least significant digit has a weight of $\pm 2^{-8}$.

Which Poles and Zeros?

The problem now was to find appropriate filter transfer functions using the available pole and zero positions of Fig. 5b. An iterative procedure was used. Initially, a prototype analog filter was selected from one of the filter design handbooks. Using the bilinear z -transform, the transfer function $H(s)$ of this filter was converted to $H(z)$. Then, with the aid of an HP

2100 disc-based computer system, the pole positions of the z transform were iteratively adjusted to the neighboring sparse CSD positions of Fig. 5b and the resulting frequency responses evaluated. The optimization strategy was directed towards keeping the stopband rejection to greater than 80 dB while minimizing the passband ripple. Equal positive and




Fig. 6. Transfer functions of the decimation filter selected with the help of a computer iterative procedure. The procedure derived the functions using pole-zero positions selected from those plotted in Fig. 5b.

negative excursions of the ripple were desired in both the pass- and stopbands.

This procedure resulted in finding sparse CSD coefficients that contain an average of 1.7 non-zero bits (2.3 for poles, 1.1 for zeros). This compares to the 4 to 7 non-zero bits that normally would be required for conventional binary.

The chosen transfer functions are shown in Fig. 6. The decimation-by-5 filter requires a fifth-order elliptic algorithm having five poles and five zeros and the decimation-by-2 filter uses a fourth-order elliptic algorithm. Although the latter has a less-than-optimum transfer function, it was chosen for reasons to be explained later.

Scaling and Word Lengths

Two other items were considered as an integral part of the design of the filter structure. One concerns overflows at the internal summing nodes. Not only does an overflow cause distortion in the output, but in recursive digital filters an overflow can cause the filter to go into an oscillatory mode that is sustained by repeated overflows. Called overflow oscillation, it

is well documented in the literature.

The usual way of dealing with this problem is to scale the signal levels between filter sections to keep the arithmetic values within bounds. The scale factors can be included as part of the feed-forward coefficients that define the zeros. We chose the so-called safe scaling, a conservative approach wherein the theoretical maximum value attainable at each node is determined and the scaling factors are adjusted to keep these values within bounds. Other strategies are available but they require additional overload detection and limiter circuits to prevent oscillations. Safe scaling also insured that no input signal regardless of shape or crest factor could cause overloads in the internal structure of the filter.

Secondly, analysis of the effects of rounding or truncation played a fundamental role in the design. For example, a 27-bit product results from multiplying a 16-bit data word by a 12-bit coefficient. Ordinarily, the product would be trimmed back to 16 bits to conserve storage and chip area but internal noise is generated in filters at every point where this trimming occurs. These noises are summed at the output

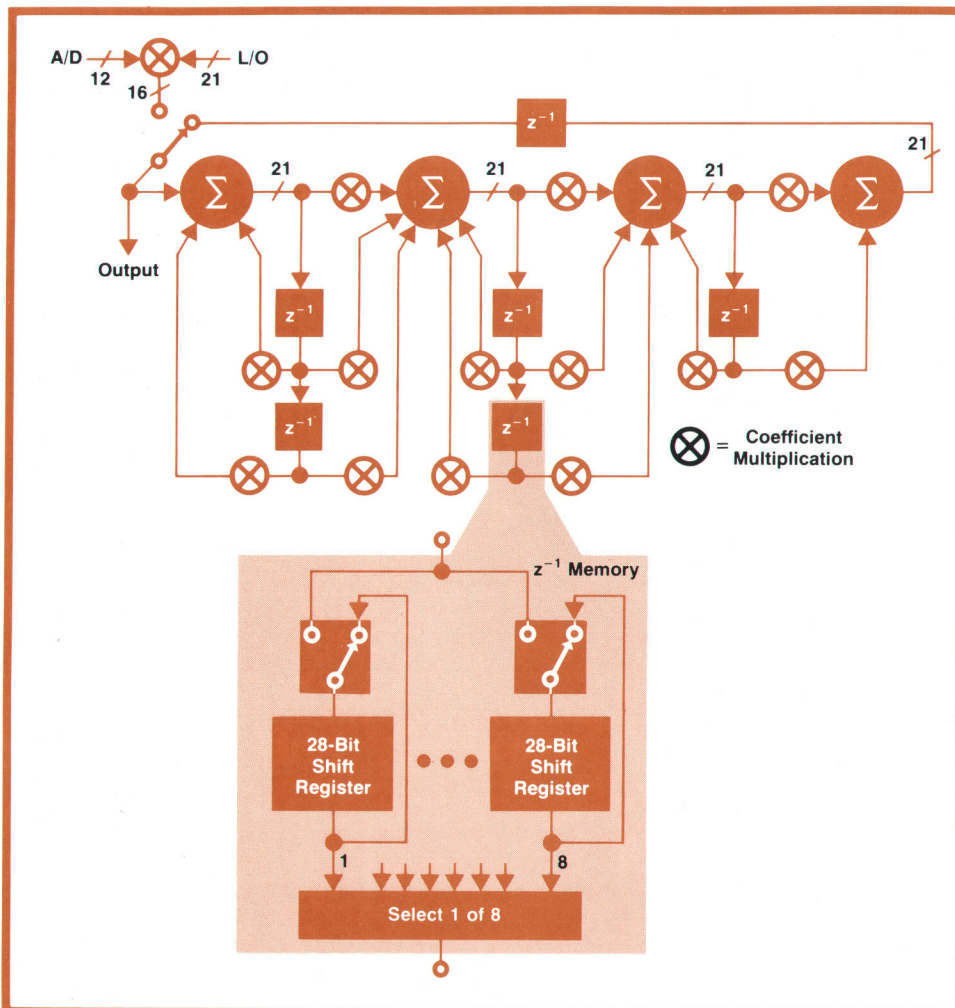


Fig. 7. Block diagram of the selected filter. This one structure is timeshared by the eight implied filter sections. Eight channels of memory in each z^{-1} block store results from the various filter section computations.

to form the total self noise of the filter. Analysis and simulation showed that to keep this self-noise at an insignificant level compared to the signal, an interstage word length of 21 bits would be required.

Optimum Structure

The final structure of the filter is related to both the self noise and the scaling. The decimation filter has two second-order zeros, two second-order poles, one first-order zero, and one first-order pole. We chose to implement the filter as three cascaded sections: two second-order and one first-order.

The relative positions of these sections can be permuted within the filter structure, and the poles and zeros can be paired in many combinations. The relative noise-power figures for the various permutations were derived theoretically (and verified once the filter was completed). The quietest structure could then be determined. Actually, the next-to-quietest structure was chosen instead of the quietest because the simpler coefficients used with the non-optimum decimation-by-2 algorithm, mentioned previously, could be formed as a subset of the decimation-by-5 coefficients. This allowed a single processor to be used for both the decimation by 2 and the decimation by 5. Savings in chip area was the reason for this decision.

The resulting arithmetic structure is shown in Fig. 7. The two second-order sections are followed by the first-order section. The multiplier at the filter input multiplies the 12-bit ADC output with the 16-bit "local oscillator" sample to perform the frequency translation for band-selectable analysis (zoom).

The data word length is 21 bits. Of these, 17 bits are necessary for a single filter's operation, an average of $1\frac{1}{2}$ bits are necessary for the combined insertion loss of the eight cascaded filters, and $1\frac{1}{2}$ bits compensate for the accumulated self noise. One bit is added for timing considerations.

As shown in Fig. 7, each memory block has eight shift-register memories to serve as the z^{-1} (sample) delays for the internal nodes of the filters. The feedback memory collects and saves filter outputs to be used as inputs to the subsequent downstream filter operations in the cascade.

Software Emulation

A bit-for-bit model of the filter structure was programmed in software and exercised extensively on the computer to determine the filter's true performance. Roundoff noise was studied and it agreed with predicted performance. Overloading was tested with worst-case inputs.

Probably the most important outcome of this portion of the design was the discovery of small signal limit-cycle problems. These are non-zero outputs

even though the input is zero, a consequence of the truncation or rounding of the results of arithmetic operations. Since the normal sample-to-sample changes in the filter input are large compared to the rounding error, the effect is simply the addition of noise, as previously discussed. The limit-cycle non-zero outputs can be either a constant value (dead-band), or oscillatory in nature and they are well-documented in the digital-filter literature.

During early investigations of the filter algorithms, analysis indicated that the data word lengths were sufficient to keep the limit-cycle amplitudes at acceptable levels. However, in experiments with the software model, more severe limit cycles appeared. Investigation revealed that several filters in the cascade would limit cycle in a "synchronous" manner such that the limit-cycle output of the last filter in the cascade was well above the noise level.

Rather than increase the data word length to reduce the limit cycle amplitude, we chose to inject a low-




Fig. 8. Shift register memory cell used in the LSI implementation of the digital filter. A bit is stored in parasitic capacitance C_1 when clock ϕ_2 occurs, and it is transferred to capacitance C_2 on clock ϕ_1 .

level dither signal into the rounding input of the summation nodes of the two second-order sections. The dither variance is sufficient to break up the limit cycles but low enough to be invisible compared to the signal. As a further precaution, the dominant spectral energy of the dither signal was placed in the filter's transition band.

Hardware Simulation

The digital filters were then simulated in hardware. This required about 450 TTL IC's per filter, this many being required because the simulated filters were designed to be as close as possible to the proposed LSI design. Many of the details and problems of the distributed control and timing system were debugged and improved with the aid of these filters. Along with the debugging performed with the software model, this proved the integrity of the design before the LSI design was complete.

The simulated filters also enabled lab prototypes of the spectrum analyzer to be built and evaluated long before the filter LSI filter chips became available.

LSI Implementation

The final design was implemented in LSI with the Hewlett-Packard NMOS-II process.² Typical of the kind of circuits used is the one-bit memory cell shown in Fig. 8. This design was chosen because it uses less chip area and loads the clock less than other designs. Twenty-eight of these are cascaded to form each of the eight circulating shift-register memories in each delay element shown in Fig. 7. These registers operate at a 6-MHz clock rate.





Fig. 9. Digital filter chip contains about 16,500 transistors and measures 4.73 x 5.27 mm.

Other circuits on the chip include full adders, counters, R-S flip-flops, D-latches, ROMs, and assorted logic gates. A key element in the design phase was the inclusion of about seventy-five 20 μ m-square test pads at critical points. During the debug phase, digital sequences at these points resulting from a test input were compared to sequences generated by the software simulator. Debugging could therefore be performed in a manner similar to signal tracing. This saved a substantial amount of time during the chip turn-on phase.

The final chip is shown in Fig. 9. Of special interest here are the areas devoted to multipliers. All of the arithmetic required to perform the decimation-by-2 and decimation-by-5 algorithms is located in the central portion of the chip. They perform the equivalent of 13 multiplications every 5 μ s, using the sparse canonical signed digits. By contrast, the input "mixer" (multiplier) at the lower left of the chip is of standard design, it performs one 12 x 16 multiplication every 5 μ s.

Acknowledgments

No IC is designed without the help of many people and the full support of an organization. My thanks to all. Special thanks go to Lou Scheffer who performed most of the chip device design and layout, to Yvette Norman and other members of the LID computer-aided artwork department for their expert help in managing the multitude of rectangles that went into the chip design, and to Dick Toftness and John Stanback of the IC engineering department for their process and design guidance. 

References

1. G.W. Reitweisner, "Binary Arithmetic," Advances in Computers, Vol. 1, Academic Press 1960.
2. J.E. Dewese and T.R. Ligon, "An NMOS Process for High-Performance LSI Circuits," Hewlett-Packard Journal, November 1977.



Lynn A. Schmidt

A native of Colorado, Lynn Schmidt did his undergraduate work at Rensselaer Polytechnic Institute and the Rochester (N.Y.) Institute of Technology, earning a BSEE degree at the latter in 1969. He then worked in sonar signal processing, obtaining an MSEE degree at Syracuse University along the way (1972). Lynn joined HP in 1974 and did some circuit design in voltmeters and signal sources before joining the 3582A project team. His outside interests include skiing and making solid

oak furniture for his home. He and his wife and three children live in Loveland, Colorado.

Desktop Plotter/Printer Does Both Vector Graphic Plotting and Fast Text Printing

This HP-IB desktop hardcopy unit has a bidirectional paper drive for long-axis plots and unattended plotting. It offers user unit scaling, graph rotation, printer capabilities, seven dashed-line fonts, English and European character sets, and user-definable characters.

by Majid Azmoon, Jaime H. Bohorquez, and Rick A. Warp

MOST PEOPLE WHO NEED HARD COPIES of computer-generated data use two separate instruments, a vector plotter for graphics and a printer for reports or program listings. Now there is a cost-effective alternative. A new microprocessor-controlled desktop hardcopy unit, Model 7245A Plotter/Printer (Fig. 1), has both capabilities. Thus it can not only replace two separate instruments, but also imbed text in graphics and combine reports and plots on a single sheet of paper (Fig. 2).

Operating as a plotter, the 7245A has most of the features of Models 9872A and 7221A X-Y Plotters.^{1,2,3} These include absolute and relative plotting, internal drawn-character generation for labeling plots, dashed-line fonts, and point digitizing. Five drawn-

character sets are standard. The plotter/printer also has automatic page advance, long-axis plotting capability, and user unit scaling. Option 001 adds circle and axis generation.

In addition to these plotter features, the 7245A operating in plotter mode can print 7×9 and 14×9 dot-matrix characters in four orthogonal directions for fast plot labeling (Fig. 3). In fact, it can do everything in plotter mode that it can do in printer mode, as described in the next paragraph, except that it cannot do tabbing in plotter mode, and the interface language is different. Operating as a plotter, the 7245A accepts instructions over the HP Interface Bus* in HP's standard graphics language, HP-GL (see box,

*The HP-IB is Hewlett-Packard's implementation of IEEE standard 488-1975 (ANSI standard MC1.1).




Fig. 1. Model 7245A Plotter/Printer combines the features of a fast thermal dot matrix printer with the capabilities of an X-Y plotter that can produce long-axis plots. Its 61-metre roll of paper allows it to operate unattended for long periods.




Fig. 2. The plotter/printer can imbed text in graphics and combine reports and plots on a single sheet of paper. Plots can be labeled with either drawn characters or dot matrix characters.

page 26). In this respect, it is software compatible with the 9872A.

Operating as a printer, the plotter/printer accepts standard ASCII line-printer codes for such functions as form feed, font change, and control of margins and tabbing. It prints 7x9 dot-matrix characters at 38 characters per second. Characters can be underlined while they are being printed, and 88 columns can be printed across the 216-mm-wide paper. A larger 14x9 dot matrix is used to print titles at 19 characters per second in a 44-column format. Option 001 adds a




Fig. 3. For fast plot labeling, the plotter/printer can print dot matrix characters in four orthogonal directions.

132-column font and the ability to reproduce a 720-bit-wide raster-scan picture from the HP-IB (Fig. 4). Standard dot-matrix character sets include a 128-character ASCII set and special characters for six European languages.

The 7245A has a bidirectional paper drive that can produce plots as long as five metres and return to the starting point from any point in any direction with a repeatability of 0.25 mm. The highest plotting speed is 0.25 m/s (10 in/s) on each axis. A moving thin-film thermal printhead serves for both matrix printing and vector plotting. The thermally-sensitive paper comes in rolls 61 m (200 ft) long, adequate for many unattended applications. Self-test features and a 120-character buffer are built-in.

Model 7245A Plotter/Printer is expected to find applications in engineering design, production testing, data acquisition, process monitoring, and business and medical plotting.

Design Considerations

The perceived applications for a product with both printing and plotting capability demanded considerable mechanical capability. In a production test envi-

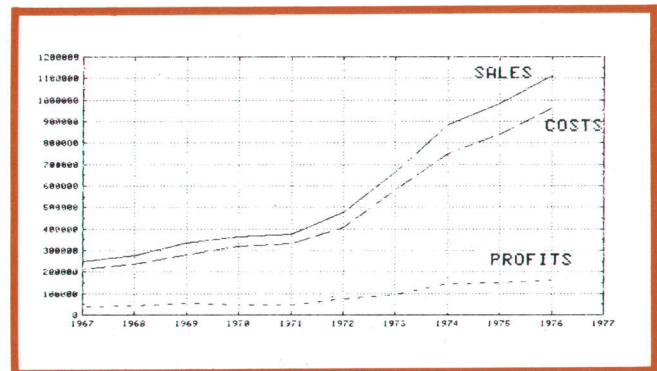


Fig. 4. Option 001 adds the ability to reproduce a raster scan CRT display from the HP Interface Bus.

ronment, for instance, an unattended plotting capability is desirable, so the characteristics of a product can be reported and plotted, then automatically repeated for the next product. This requires the ability to advance pages automatically. In making Gant charts for project management, plots may be many pages long to account for the many time periods involved, and this requires a machine that can make a plot with one very long axis. It also means that the machine must be able to return to a given position on the plot from perhaps three meters away with the same accuracy as if it had come from 20 mm away. These requirements for unattended and long-axis plotting put significant constraints on the development of the paper drive for the 7245A.

In developing the 7245A Plotter/Printer and the 9872A and 7221A X-Y Plotters, many pieces of technology development were shared. For instance, the step motors and the microstep drive used in the 7245A are identical to those used in the 9872A and the 7221A.⁴ The microprocessor systems are the same, and the language is the same as the 9872A's, with some extensions to account for the increased capability of the 7245A.

In addition, some new technologies were required for the dual capabilities of printing and plotting. A thermal printhead was developed to achieve both the plotting quality demanded and a printing speed reasonable for the applications. It was also essential that a bidirectional paper drive be developed to provide the long-axis and unattended plotting capabilities.

While the plotter/printer is a vector device when plotting, it is a raster device when being used as a printer. That is, it prints horizontal rows of dots on the paper to form characters, 12 rows at a time. This capability can be generalized to any dot-matrix information. With Option 001, the 7245A has the ability to accept general raster information from the HP-IB. The first products to make use of the 7245A in this mode are the HP 2647A and 2648A Graphics Terminals.

With the increasing number of instruments that have built-in intelligence, the question arises whether an instrument can drive a graphics device on the HP-IB without any other controller in the system. The 7245A provides for this with a listen-only mode. The instrument need only send the appropriate graphics commands in HP-GL to the 7245A. It does not have to address the 7245A as a listener on the HP-IB. This can significantly decrease the cost of the customer's solution, since no controller is needed to do plotting. The 7245A's scientific input format and internal scaling capability enhance the usefulness of this feature.

Thin-Film Plotter/Printer Printhead

The 7245A Plotter/Printer uses a single thin-film

printhead for dot-matrix printing, vector plotting, and drawn characters. Heat generated when thin-film resistors are energized is applied to areas of heat sensitive paper, causing a reaction that changes the color of the paper at the point of applied heat. By constant application of heat, continuous lines can be produced to form drawn characters or plotted lines. If instead the resistors are pulsed, independent dots can be formed.

The 7245A prints a standard half-shifted 7×9 dot matrix character set. By printing each vertical row twice, a 14×9 matrix font is generated that can be used for printing titles. The expanded font is of the same height, but twice the width of the standard characters, giving the printing a bold appearance. Expanded characters can be printed at 19 characters per second and standard characters can be printed at 38 characters per second. Option 001 adds a 132-column 5×7 dot matrix character set printed at 64 characters per second.

Text can be printed in any of the four orthogonal directions with or without underline. An example of a lower-case s, underlined to show orientation, is shown in Fig. 5. The configuration of the text printing resistors is shown in the center of Fig. 5. This array is the equivalent of a diagonal line of resistors and is used because a true diagonal line was difficult to reproduce photolithographically. An array of this type allows the resistors to be placed without an apparent gap between resistors. The diagonal array allows text printing in two directions, and by approx-




Fig. 5. Dot matrix printhead resistors are arranged in a diamond pattern to allow printing in any of the four orthogonal directions, with or without underline.

appropriate addressing of the print elements, text can be printed in four directions.

The design of the thermal printhead was dictated by several requirements. The primary requirement was for the printhead to be capable of high-quality, fast, high-accuracy plotting. To achieve a high-quality plotted line, the heat element has to have high and constant heat transfer to the surface of the paper. High thermal efficiency guarantees that, during continuous plotting, the printhead assembly temperature does not rise above the thermal paper threshold. These requirements for the plotting element are met by using relatively thick thermal insulating material (glass) to support the thin-film resistors.

A secondary requirement was for the thermal printhead to be capable of fast text printing. To print dot-matrix characters quickly the thermal time constant of the print resistors has to be very short. This requirement for the matrix printing elements calls for a very thin glaze that allows rapid heating and cooling of the printing elements. Early in the project the decision was made to assure that the 7245A was an excellent plotter even at the sacrifice of very fast text printing. The glaze thickness was selected to achieve a high-quality plotted line with acceptable printed text.

As shown in Fig. 6, the thin-film printhead has an alumina (aluminum oxide) substrate 10 mm by 16 mm. The substrate is 1 mm thick and is coated with 70 micrometres of glass as the insulating layer. The glaze is etched to leave raised areas that act as mesas for the thin-film resistors. The 10-micrometre-high mesas improve the contact between the heating elements and the surface of the paper. The resistor film (tantalum aluminum) and the conductor film (alumi-

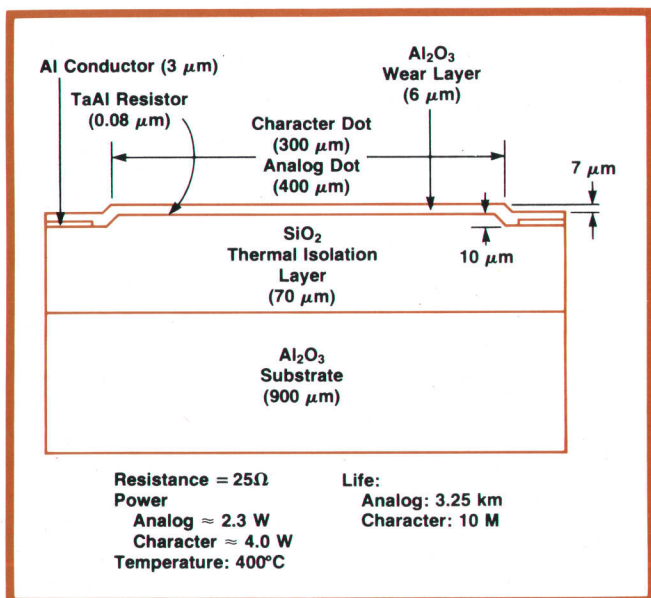


Fig. 6. Cross-section of the thin-film thermal printhead.




Fig. 7. Photograph of the thermal printhead, showing the array of dot matrix printing resistors and the single, larger plotting resistor.

num) are then vacuum-deposited by sputtering. Patterns are etched using conventional photolithographic techniques. A protective layer of aluminum oxide is sputtered to provide both chemical and mechanical protection of the thin-film printhead.

Fig. 7 is a photograph of the printhead showing the array of dot-matrix print resistors and the separate plotting resistor. Each printing resistor is approximately 340 micrometres square and the plotting resistor is 450 micrometres square. Mounted on the printhead is a clear plastic sight with crosshairs. This cursor can be used to position the printhead with pinpoint accuracy.

The thermal paper used by the 7245A employs the two-component dye color reaction method of image formation. The two component compounds are pulverized to a diameter of several micrometres, mixed with a binder, and coated on the paper surface. The completeness of the color-change reaction depends on the temperature reached by the heat sensitive layer. During plotting, the power to the resistor has to be modulated as a function of velocity. To insure rapid temperature rise, the power is increased during initial turn-on or during rapid acceleration.

Because of the compromise glaze thickness, a more complicated method of power modulation had to be used for the printing resistors. Each individual power pulse is modulated to insure rapid temperature rise without excessive surface temperature that might prove damaging to the resistor elements. Modulation of the power envelope minimizes the effects of the gradual heating of the printhead assembly during a continuous power pulse train. Print element power modulation is shown in Fig. 8. Two different modulation time constants are used: τ_0 is the individual power pulse modulation time constant and τ_1 is

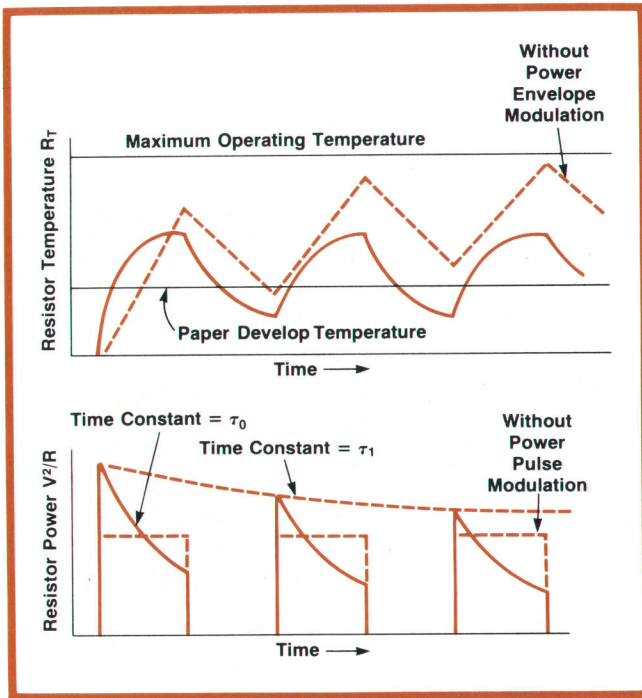


Fig. 8. Power to the printhead resistors is modulated to insure rapid temperature rise without excessive temperatures. The individual power pulses and the envelope of the pulses are both modulated.

the power envelope modulation time constant. This type of power modulation has been shown to improve print quality and printhead reliability.

Printhead Gimbal

One of the factors affecting printing quality is contact between the head and the paper. Maintaining good paper-resistor contact over a large area is a problem. Not only must all the resistors touch the paper, but ideally each resistor should see an identical force. The total force from all resistors is called the head force.

A rubber roller can be compressed to provide a flat area of contact, but the force gradient from center to edge is very severe and print consistency is far from adequate. A flat platen, with perfect alignment between head and paper, could provide equal force on each element, but practically speaking, perfect alignment is very difficult, if not impossible, to achieve.

In this machine the head is free to move in two axes, thus making it self-align to a rigid flat platen. This method has its own set of problems, because the rotation or gimbal point is above the surface of the paper. Friction forces produced by the head on the paper introduce moments about the gimbaling point and these moments are counteracted by an unequal force distribution between head and platen (Fig. 9).

The ideal situation would be to have the gimbal

point on the surface of the paper. A four-bar linkage can accomplish this, but involves quite a few moving parts, creating manufacturing and reliability problems. The 7245A does not use a four-bar linkage. Instead, the gimbal point is arranged as close to the surface of the paper as possible and then the moments about this point are minimized by offsetting the gimbal point from the center of the printhead. In this way it was possible to optimize printing in the left-to-right direction. Experimental compromises were then made to obtain adequate printing in the other three directions.

The optimum horizontal position for left-to-right printing is computed as follows (see Fig. 9). Assume a linear force distribution $w(x)$, that is,

$$w(x) = a + bx.$$

Equating vertical forces, the head force F_H is

$$F_H = \int_0^l w(x) dx = \int_0^l (a + bx) dx$$

Balancing the moments about the gimbal point,

$$-F_f d = \int_0^{ql} (ql - x)w(x) dx - \int_{ql}^l (x - ql)w(x) dx$$

$$-F_f d = \int_0^l (ql - x)(ax + b) dx$$

where F_f is the friction force, d is the minimum possible distance from the gimbal point to the paper surface, and $0 < ql < 1$. Solving for a and b , the coefficients




Fig. 9. Printhead gimbal point location is optimized to maintain good paper-resistor contact and equalize the forces seen by the individual print resistors.

of the linear force distribution,

$$a = -\frac{6F_Hq}{l} + \frac{4F_H}{l} + \frac{6F_f d}{l^2}$$

$$b = \frac{12F_f d}{l^3} + \frac{12qF_H}{l^2} - \frac{6F_H}{l^2}$$

For no force gradient, $b = 0$.

$$\frac{12F_f d}{l^3} + \frac{12qF_H}{l^2} - \frac{6F_H}{l^2} = 0$$

Solving for q in this equation gives the optimum horizontal position for left-to-right printing.

$$q = \frac{1}{2} - \frac{F_f d}{F_H l}$$

Bidirectional Paper Drive

The bidirectional paper drive represents a vital contribution to this product. This technology, combined with the thermal printhead, provides vector graphics, long-axis plotting, and unattended plotting, as well as high-quality printing.

For high-quality plotting and printing, a head force of 150 grams is needed to maintain good contact between the printhead and the paper. This head force gives the paper a tendency to buckle. In the 7245A, the buckling problem is overcome by using a vacuum system that tensions the paper uniformly in both directions along the paper length.

One conventional means of moving a sheet of paper is the friction drive, as in a typewriter. A similar friction paper drive was investigated for the 7245A. Essentially, this system consisted of two rubber rollers driving the paper, a writing platen in the middle, and a few spring-loaded idlers to keep the paper on the rollers.

The problems encountered with this system were reliability of the drive in a wide environmental range and accuracy and repeatability in long-axis plotting. Any side-to-side variations in a roller's diameter or any variation in idler forces during the paper movement created a tendency for the paper to move toward one side or the other, resulting in gross inaccuracies in long-axis plotting or complete failure of the drive. To avoid this problem in a friction drive means balancing a number of large vector forces, trying to resolve them into a small force to cause the paper to track precisely. This is extremely difficult if not impossible to obtain in a wide range of environments.

Repeatability, accuracy, and reliability were achieved by using a sprocketed drive system driven by a step motor (Fig. 10). This system has two sets of

sprockets, front and rear. The front set consists of two sprockets and a stationary platen in the middle and is driven directly for better accuracy. The rear set is a sprocketed drum and is driven by a belt from the same step motor. The two sets of sprockets are synchronized to provide identical surface velocities.

The continuous sheet of paper wraps around the bottom of the drum, then around the platen and back to the top of the drum, forming a chamber in which a vacuum is drawn by a centrifugal fan (Fig. 11). Regardless of the direction of paper movement, the vacuum acts evenly on both sides of the platen, tensioning the paper uniformly to overcome the frictional drag from the head force and the stationary platen. The tension in the paper is linearly related to the vacuum pressure, which in this design is 0.2 inch of water.

Changes in humidity and temperature markedly affect the strength and dimensional stability of paper. This paper drive is capable of operating at 40°C and 95% relative humidity as well as at 40°C and 15% relative humidity.

Paper movement on the stationary platen results in enough friction to make the system over-damped. Consequently, this axis of the drive, unlike the other axis, does not need a mechanical damper to reduce the amplitudes of resonances. The mechanical damper used on the other (X) axis is the same as the one used in the 9872A X-Y Plotter.⁴

When the sheet of paper is driven back into the machine, it is stored in a paper chamber instead of being rolled back onto the supply roll. Five metres of paper can be stored in this chamber and driven out and back many times. This mechanism is not so compatible with z-fold paper as with rolled paper, primar-




Fig. 10 Sprocketed drive system provides accuracy repeatability, and reliability. The drive is bidirectional; paper can back up as much as five metres. Repeatability is 0.25 mm from any point in any direction.





Fig. 11. A vacuum in the chamber formed by the paper and the drums tensions the paper uniformly to overcome the frictional drag from the head force and the stationary platen. Acceleration buffer helps the relatively small motor accelerate the large roll of paper.

ily because the static charge created in the paper by paper movements on the stationary platen causes the folded sheets to stick to each other in the paper chamber. When this happens, extra sheets of paper may move through the sprockets. Using a roll of paper eliminates this problem. However, the step motor does not have enough torque to accelerate a full roll of paper. To make use of the available torque without reducing the paper acceleration, an acceleration buffer is placed between the roll and the drum (Fig. 11). This device is basically a shaft that swings in an arc against a spring force. When the paper is pulled by the drum, the shaft is displaced, allowing the roll of paper to accelerate at a lower rate than the rest of the system. The roll then continues to accelerate until the acceleration buffer returns to its stable position and the roll has caught up with the rest of the system.


Acknowledgments

It is difficult to thank the many people to whom we are indebted on a project this large, but here goes. Norm Johnson's leadership of the firmware group and help in product definition kept us on the right path. His group made the hardware sing. They were Mark Allen, Tom Halpenny, and Larry Hennessee. Dave Shelley was a contributor in every discipline on the project. The power supply was done by Dave, Myron Son, and Dave Ellement. It was packaged by Dick Kemplin and cooled by Dave Perach. Steve White implemented our processor, memory, and drive electronics, while Hal Beach conceived the power modulation scheme on the printhead drive. Terry Flower and John Morton did much of the paper drive. John is now keeping the 7245A alive and well in production engineering. Bill Sperry did the development and characterization of the thin-film process and thin-film sputtering systems. Bill Ebert of HP Laboratories

and Beth Nidzicko transferred the process from HPL to San Diego. Terry Siden and Tom Young of our tooling and process group logged many miles and hours obtaining tooled parts for the project. Their expertise was vital. Pat Fobes became virtually a part of the lab project team with his contributions as our product marketing engineer. 

References


1. L.G. Brunetti, "A New Family of Intelligent Multi-Color X-Y Plotters," Hewlett-Packard Journal, September 1977.
2. T.H. Daniels and L.W. Hennessee, "Easy-to-Use Interface Language Controls HP-IB Plotters," Hewlett-Packard Jour-



Majid Azmoon

Majid Azmoon was responsible for the paper drive, head drive and writing system mechanics for the 7245A plotter/printer, and the mechanical characterization and line quality for the 9872A four-color plotter. With HP since 1973, Maj was named inventor on the patent for the 7245A bidirectional paper drive and is currently project leader for development of a new plotter. He graduated from California Polytechnic University in 1969 with a BSME degree and went on to get his MSME degree in

1971 from the University of Southern California. Since graduation in 1971, he has done additional graduate study in mechanical engineering at USC. Born in Tehran, Iran, Maj now lives in Poway, California, with his wife and son (age nine months) in the house that he built himself. His favorite pastimes are swimming, cooking, landscaping, racquetball, woodworking and building houses.



Jaime H. Bohorquez

An HP employee since 1973, Jim Bohorquez was project manager for the development of the thin-film process, facility and thermal printhead for the 7245A plotter/printer. Before joining HP, Jim designed several types of integrated circuits, ranging from low-power linear circuits to a family of sub-nanosecond mainframe logic circuits. Originally from Bogota, Colombia, Jim received his BSEE degree from University of California at Berkeley. He was a United States Navy officer for four years,

has taught courses in electronics and integrated circuit design and is named inventor on two patents. Jim lives in Escondido, California, with his wife and two daughters, ages nine and four. In his leisure time, he enjoys family outings, photography, and restoring his classic Porsche 356C.

nal, September 1977.

3. D.A. Bones and M.L. Patterson, "Remote Terminal Plotter Offers Simple Programming and Efficient Communications," Hewlett-Packard Journal, September 1977.

4. M.L. Patterson, R.D. Haselby, and R.M. Kemplin, "Speed, Precision, and Smoothness Characterize Four-Color Plotter Pen Drive System," Hewlett-Packard Journal, September 1977.



Rick A. Warp

Rick Warp is a Stanford University graduate who received his BSME degree and MSME degree simultaneously in 1971. He was project manager for the 7245A plotter/printer and also did mechanical design on the 7155B portable strip chart recorder. A seven-year HP employee, Rick is now product marketing engineer at San Diego Division and is named inventor on two patents related to the 7245A. Originally from Tacoma, Washington, Rick now makes his home in San Diego, California, with

his wife and their six-month-old daughter. He is "addicted" to tennis and enjoys woodworking, home improvement and "watching his daughter grow."



Michael P. Trego

Mick Trego was born in Gillette, Wyoming, and received his BSEE degree in 1962 and MSEE degree in 1963, both from the University of Wyoming. He was responsible for the design of the HP Interface Bus hardware and software for both the 7245A plotter/printer and the 9872A four-color plotter. He also designed the read head for the 7260A and 7261A optical mark card readers. Before joining HP in 1971, Mick did aerospace work for another firm, primarily in digital design. Mick lives in San Diego,

California, with his wife, daughter, 15, and son, 13, and when he's not water skiing, snow skiing or dune bugging, he's active taking photographs, camping with his family and repairing his buggy.

SPECIFICATIONS
HP Model 7245A Plotter/Printer

PLOTTING/PRINTING AREAS

MAXIMUM: 188 mm x 5 m (7.4 in x 16.4 ft) with full paper return.
188 mm x 61 m (7.4 in x 200 ft) without full paper return.
ENGLISH PAGE: 188 x 279 mm (7.4 x 11 in).
METRIC PAGE: 188 x 298.5 mm (7.4 x 11.75 in).

ADDRESSABLE DYNAMIC RANGE: ±1 x 10⁻³⁰ scaled units.

PLOTTING ACCURACY: ±0.2% of deflection ±0.35 mm (±0.014 in). Includes linearity and repeatability.

REPEATABILITY: 0.25 mm (0.01 in) from any given point.

MOTOR RESOLUTION: 0.016 mm (0.0006 in).

MAXIMUM PLOTTING SPEED

ACCELERATION: 4.48 m/s² (14.7 ft/s²)

VELOCITY

Pen Off: 513 mm/s (20.2 in/s) in each axis
725 mm/s (28.6 in/s) at 45° angle
Pen On: 256 mm/s (10.1 in/s) in each axis
363 mm/s (14.3 in/s) at 45° angle

PRINTHEAD POSITIONING

WITH ARROW BUTTON: 6.1 mm/s (.24 in/s)

WITH ARROW AND FAST BUTTONS: 95.5 mm/s (3.76 in/s).

CHARACTER PRINTING SPEED: 7 x 9 dot matrix character at 38 char/s; 14 x 9 dot matrix character at 19 char/s; with option 001, 5 x 9 dot matrix character at 63 char/s; 2.5 mm (0.1 in) character typically drawn at 3 char/s.

OFF-SCALE PLOTTING: When off-scale data is received by the 7245A, the unit will automatically calculate the mechanical limit intercept of that vector and proceed to that point. As additional off-scale data is received, the 7245A will monitor the location of this data and resume plotting, once on-scale data is received, by again calculating the new mechanical limit intercept and plotting from that limit to the on-scale data point. Plotting accuracy and repeatability specifications are preserved.

PRINTHEAD CONTROL: Local control by front-panel switches or remote control by HP-IB controller program commands.

OPTIONS: 001: CRT Raster Dump and HP-GL Enhancements.

POWER REQUIREMENTS:

SOURCE: 100V -10%, +5%
120V -10%, +5%
220V -10%, +5%
240V -10%, +5%

Internally selectable

FREQUENCY: 48-66 Hz

CONSUMPTION: 100W/2.8A
120V/2.5A
220V/1.3A
240V/1.2A
300 W max

ENVIRONMENTAL RANGE:

TEMPERATURE: 0°C to 55°C.

RELATIVE HUMIDITY: 5% to 95% (below 40°C).

DIMENSIONS: 201 mm H x 442 mm W x 483 mm D (7.9 x 17.4 x 19 in).

WEIGHT: 19.1 kg (42 lb).

CUBE: 0.11 m³ (3.96 ft³).

PRICES IN U.S.A.: Model 7245A, \$4600.00

Option 001, \$250.00.

MANUFACTURING DIVISION: SAN DIEGO DIVISION

16399 West Bernardo Drive
San Diego, California 92127 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 1978 Volume 29 • Number 13

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard Company, 1501 Page Mill Road
Palo Alto, California 94304 U.S.A.
Hewlett-Packard Central Mailing Department
Van Heuven Goedhartlaan 121
Amstelveen-1134 The Netherlands
Yokogawa-Hewlett-Packard Ltd., Suginami-Ku
Tokyo 168 Japan

Editorial Director • Howard L. Roberts

Managing Editor • Richard P. Dolan

Art Director, Photographer • Arvid A. Danielson

Illustrator • Susan E. Wright

Administrative Services, Typography • Anne S. LoPresti

European Production Manager • Dick Leeksa

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.