

HEWLETT-PACKARD JOURNAL



An Automatic Selective Level Measuring Set for Multichannel Communications Systems

Combining microprocessor control with meticulous receiver and synthesizer design results in a powerful new tool for the installation, maintenance, and surveillance of frequency-division multiplexed communications systems.

by J. Reid Urquhart

MAKING RELIABLE MEASUREMENTS on frequency-division multiplexed (FDM) baseband signals presents a two-part problem. First, the frequency at which the measurement is to be made must be determined. Then second, the measurement must be made with high accuracy in the presence of numerous interfering signals.

At first sight the choice of a measurement frequency may seem to be a trivial problem, but finding the center frequency of a particular channel in a communications system that may have as many as 3600 channels is no easy task and could, in fact, occupy more time than making the actual measurement itself. A maintenance engineer may wish, for example, to measure channel 3 in group 2 of supergroup 13 in mastergroup 5 without having any idea of where in the frequency spectrum that channel may lie. The only way to find out has been through the use of large charts that describe the FDM plan, such as the one shown in part in Fig. 1. Usually these only show each channel's carrier frequency, so the engineer must derive the center frequency of the sideband occupying the channel in order to proceed with the measurement. Besides the time taken to hunt for a particular channel's frequency, such an arrangement leads to the possibility of frequent errors.

The problem of translating the FDM description of a channel into a tuning frequency has now been eliminated by a new selective level measuring set (SLMS), the Hewlett-Packard Model 3745A/B (Fig. 2). This instrument has all the details of FDM plans stored in read-only memories. The stored plans are accessed by special keys, so the user merely enters the specifics of a certain channel (channel number, group number, etc.) by way of the keys, and the instrument then automatically tunes itself so the 3.1-kHz filter covers the active region of that channel exactly. It then makes a true-rms measurement of the power in that channel

and displays the results, including the frequency to which it is tuned, on numeric indicators. Furthermore, the instrument is completely autoranging over a 140-dB range, so it displays results directly without the need for setting attenuators manually.



Cover: *Picking one particular channel out of a babble of thousands of frequency-multiplexed communications channels is greatly simplified by a new selective level measuring set that talks the language of the telephone engineer. Micro-*

processor control is the key to this new ability, as described in the series of articles beginning on this page. Our thanks to Farinon Electric for the loan of their facilities for this photo.

In this Issue:

- An Automatic Selective Level Measuring Set for Multichannel Communications Systems, by J. Reid Urquhart* **page 2**
- Designing Precision into a Selective Level Measuring Set, by Hugh P. Walker* **page 8**
- Designing a Quiet Frequency Synthesizer for a Selective Level Measuring Set, by John H. Coster* **page 13**
- Making the Most of Microprocessor Control, by David G. Dack* **page 16**
- Real-Time Multi-User BASIC, by James T. Schultz* **page 19**

TEST TONE AND GROUP REGULATING PILOT FREQUENCIES
CCITT AND NBS-1-1000 INTERNATIONAL TELEPHONE PLAN
(TEST TONE FREQUENCIES DERIVED FROM NBS-1-1000, SEP. 1974, BY 100-KHZ STEPS IN 100-KHZ STEPS AT 100-KHZ)

CH	BAND	GROUP	TEST TONE FREQUENCIES (kHz)											
			100	101	102	103	104	105	106	107	108	109	110	111
Pilot	100-108	1	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
			100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000
100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000	100.000			

Fig. 1. A typical frequency plan for a frequency-division multiplexed communications system. Tables such as this give each channel's virtual (suppressed) carrier frequency from which the channel center or appropriate pilot frequency must be derived for measurements. These tables may list as many as 3600 frequencies.

Measurement Versatility

Although designed particularly to meet the requirements of FDM equipment manufacturers and users, the new SLMS is suited to making measurements wherever signal levels must be determined with high accuracy. Basically, it is a highly accurate, programmable wave analyzer capable of making measurements within a frequency range of 1 kHz to 25 MHz over a level range of -125 to +15 dBm. Besides automatically tuning to selected FDM channel frequencies, it can be tuned to any frequency within the 1-kHz-to-25-MHz range by entering that frequency by way of the numeric keyboard.

High accuracy is achieved by the use of filters that have good passband flatness coupled with very steep roll-off and high out-of-band rejection, and by an autocalibration routine that automatically calibrates the instrument to an internal 1-MHz, -30 dBm standard prior to each sequence of measurements. The Model 3745A/B can measure a pilot tone at -50 dBm anywhere within a 50-kHz-to-20-MHz frequency range with an absolute accuracy of ±0.35 dB, including all errors—attenuator, flatness, and absolute. The absolute error at 0 dBm and 1 MHz is only ±0.05 dB.

World-Wide Application

Two versions of the instrument are available. The "A" version contains the CCITT FDM plans while the



Fig. 2. Model 3745A Selective Level Measuring Set designed for use on CCITT systems is basically a programmable, autoranging, 1-kHz-to-25-MHz, precision wave analyzer with filter bandwidths optimized for telephone applications. Model 3745B, shown on the front cover, is similar but is designed for use on North American systems. In measurements on multiplexed communications systems these instruments automate the selection of tuning frequencies, reducing to minutes procedures that would take hours with manually-tuned instruments.

"B" version has the plans specified by the Bell Telephone System. The stored plans also include the specifics for the various pilot frequencies that control gain in communications systems and that also provide vital information about the workings of a system. A typical CCITT 2700-channel multiplex system may have 225 group, 45 supergroup, and 3 hypergroup pilots, giving a total of 273 pilots, any of which may be selected by pushbutton sequences. The instrument recognizes that a pilot is to be measured if no channel number is entered, tunes itself to the appropriate frequency, and measures the pilot level with a filter bandwidth of 22 Hz.

If the GROUP POWER pushbutton is pressed following the group selection sequence, the instrument tunes itself to the center frequency of the selected 12-channel group and measures the total power in that group with a filter bandwidth of 48 kHz.

The appropriate filter for each of the measurements is automatically selected when the FILTER switch is in the AUTO position, or the filters may be selected manually using the same switch.

Spectrum Analysis

The new SLMS can also function as a spectrum analyzer. It is then useful for finding interfering signals

Operating a Selective Level Measuring Set through a Keyboard

The new Model 3745A/B Selective Level Measuring Set can be tuned to a particular channel by entering information in the same way that a telephone maintenance engineer thinks of a telephone channel, saving him the tedium of deriving the actual frequency. For example, to measure the signal level in channel 1 of group 3 of supergroup 13 in mastergroup 5 of supermastergroup 3, the operator would first press the following key sequence: CH 1 G 3 SG 13 MG 5 SMG 3. Then, when MEAS/HALT is pressed, the instrument automatically tunes itself so the channel filter spans the active region of that channel, and it displays the frequency and signal level.

If it is now desired to select a different channel, it is only necessary to key in the changes. In the above example, to measure channel 3 of group 2 in the same supergroup, the operator would press MEAS/HALT to stop the current measurement, key in CH 3 G 2, and press MEAS/HALT once more to restart the measurement.

If the key sequence does not include a channel number, then when MEAS/HALT is pressed the SLMS tunes to the frequency of the pilot in the specified group. A front-panel annunciator will light up to show that a pilot is being measured. Similarly, if a group number has not been entered, the SLMS tunes to the pilot of the specified supergroup.

When the MEAS/HALT key is preceded by the GROUP POWER key, the instrument tunes to the center frequency of the selected group and measures the total group power.

If the SCAN key is pressed before MEAS/HALT, the instrument first measures the channel specified, then steps to the next higher channel and measures it. It continues stepping to each higher channel until channel 12 is reached, at which time it steps to the next higher group, measuring channel 1 of that group and continuing thus to the next higher channel until all the channels in the FDM plan have been measured. It can scan all the channels this way once and then stop, or it can continue scanning repetitively depending on the setting of the SWEEP switch.

It is not necessary to scan the complete baseband signal every time. If, for instance, the initial scan indicates a problem in supergroup 13, the scan can be limited to that supergroup by first pressing MEAS/HALT to stop the current measurement (all key functions are inhibited when the instrument is making a measurement) and then rekeying SG 13 followed by the decimal point. The decimal point "fixes" all higher-ordered groupings so the scan is confined to the 60 channels of supergroup 13 when MEAS/HALT is again pressed. If the problem should then turn out to be in group 4, then pressing G 4. will confine the scan to the 12

channels of group 4. Finally, by pressing CH 1 SPECT, choosing a step size of, say, 20 Hz, and pressing MEAS/HALT a spectrum analysis of the individual channel is carried out. Thus, a problem can be "zoomed in" on for examination in close detail.

A scan can be conducted manually by using the +↑ and -↓ keys to single-step the scan either up or down. Manual operation is useful for positioning the marker to identify points of interest on the CRT display.

The SLMS is not restricted to measuring at frequencies specified by the FDM plan. Any measurement frequency can be entered by first pressing the FREQ button, after stopping the current measurement with MEAS/HALT, then the frequency in kHz followed by MEAS/HALT again. For example, a measurement is made at 14.38 kHz by the key sequence: MEAS/HALT FREQ 14 . 38 MEAS/HALT.

Spectrum measurements are made by keying in the start and stop frequencies and the step size. To examine the spectrum between 1 and 2 MHz in 3-kHz steps, for example, the key sequence is START FREQ 1 0 0 0 STOP FREQ 2 0 0 0 STEP SIZE 3 SPECT MEAS/HALT.

Start and stop frequencies may also be entered with the FDM keys. For example, if the key sequence G 5 SPECT is entered, the lowest frequency of group 5 will be entered into the start register and the highest frequency into the stop register (assuming that the supergroup and higher groupings were already entered). When an appropriate step size is entered, a spectrum analysis can be performed to find, say, interfering tones lying in frequency regions normally not occupied in the FDM plan.

The dB/dBm key gives a choice of absolute or relative measurements. For relative measurements, a reference is entered by way of the numeric keyboard and the REF LEVEL key.

Measurement limits are specified by first entering the expected measurement level using the REF LEVEL key. Then the upper limit in dB above the reference and the lower limit in dB below the reference are entered using the UPPER LIMIT and LOWER LIMIT keys and the numeric keyboard.

If during scanned measurements it is desired to compare measurements to the previous scan, the TRANS key is pressed immediately before the SCAN or SPECTRUM keys. This facility can be used for spectrum measurements and when scanning pilots and groups, but not when scanning channels.

In all of these operations, the numeric indicators display the numbers as they are entered. Annunciator lights show parameters such as the filter bandwidth in use and the status of the measurement. If non-valid key or switch settings are entered, the numeric indicators display "Error" when MEAS/HALT is pressed and code numbers are displayed to indicate the type of error.

lying in frequency regions normally not used in FDM systems. In this case the user merely keys in the start frequency, the stop frequency, and the size of the frequency step, which may be as small as 10 Hz. Filter selection can be automatic according to the size of the frequency step, or manual.

Alternatively, the instrument may be operated in a scanning mode, in which case it automatically steps through and measures the level of every channel or pilot signal within a range selected by the group/supergroup/mastergroup/supermastergroup keys. This type of measurement is similar to a spectrum analysis but the instrument examines only the points of interest, skipping over the unused portions of the spectrum. For example, it takes only 100 seconds or so to measure all 225 group pilots of a CCITT 2700-channel system whereas a conventional spectrum analyzer, set to 10-Hz bandwidth to give adjacent-channel rejection comparable to the SLMS's pilot filter, would take about one month to scan the same 12-MHz range.

Another advantage of the SCAN mode is that since a discrete measurement is made on each channel, it is easy to identify a channel of interest whereas it is exceedingly difficult to do so with a conventional spectrum analyzer.

Limit Tests

Measurements in the SCAN mode can be used in several different ways. First, it is possible to set high and low limits on the expected level of measurement. If a measurement lies outside the specified range, an annunciator flashes, the scan stops, and the instrument displays the FDM description of the out-of-limit channel and its signal level. Once details have been noted, the scan may be continued. This mode is useful for detecting spurious tones, high- or low-level pilots, overloaded channels and groups, and other out-of-limit conditions.

If the optional accessory printer is used with the SLMS, the instrument can be set to stop the scan just long enough to print the description of an out-of-limit condition, and then continue the scan. This mode is particularly useful for unattended operation as it provides a record that is useful for doing corrective maintenance at a more convenient time (Fig. 3). If desired, all measurements may be recorded and there will be an indication on the printout of those that are out of limits.

The instrument can also compare the results of a scan to those of the previous scan, generating an alarm only if a measurement has either gone out of limits or come within limits since the previous scan. This mode is useful, for example, when scanning pilots on a partially loaded system where all supergroups may not be in use, a situation typical of a new

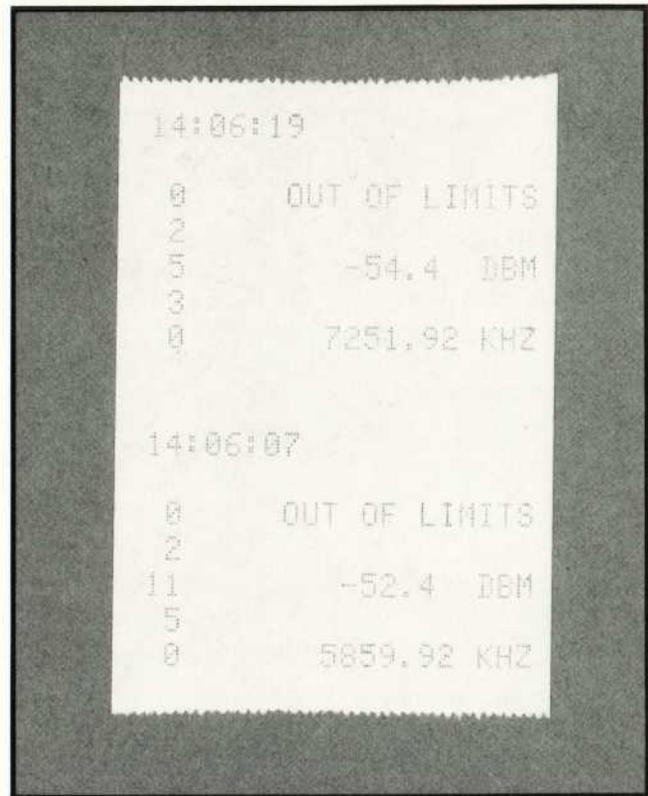


Fig. 3. Typical printout from the optional digital printer lists the time of day and the FDM description of the measured channel (vertical column of numbers).

system being brought into service. A missing pilot generates an alarm on the initial scan but the alarm is suppressed on subsequent scans (unless the missing pilot appears within the limits). This avoids the printout of redundant information.

An X-Y driver option is also available. This provides outputs for driving either an X-Y recorder or a CRT display to derive plots of signal level versus frequency. Each measurement is represented by a dot in the plot. The instrument can store up to 255 measurement points and read them out repetitively at a fast enough rate for flicker-free display on a non-storage CRT (Fig. 4).

Points continue to be stored as scans repeat, the earliest points being replaced by new data when over 255 measurements have been made. It is thus possible to make repetitive scans of group reference pilots within a mastergroup, for example, and display several measurements of each pilot simultaneously. Any short-term level variations in any of the pilots, or of any other signals, would immediately be apparent. A flashing, circular marker indicates the point currently being measured, and its identity is displayed by the front-panel indicators. The marker can be positioned with single-step keys to identify points of interest on the display.

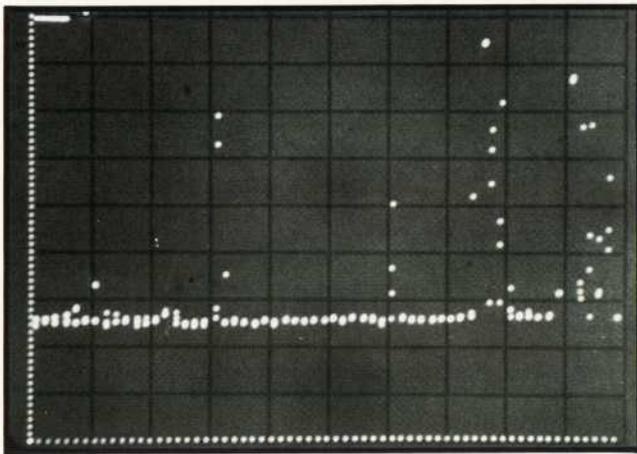


Fig. 4. This display on the optional CRT monitor shows signal levels in the 60 channels of a supergroup. Most channels are not in use and lie at the idle noise level (the short horizontal line at left is the reference level at -20 dBm). Channels with several dots in a vertical line are speech channels while the three bright dots at top are data channels. The vertical scale is 10 dB/div.

All of the measurement modes described for SCAN operation may also be used when the system operates in the SPECTRUM mode.

Stimulus and Response

The new SLMS can also work with a suitable signal source to measure frequency response, gain/loss, harmonic distortion, and crosstalk. All of these measurements can be made automatically when the SLMS and the generator are working together in a calculator-based system. The new SLMS is equipped to work with the HP Interface Bus (Hewlett-Packard's implementation of IEEE Standard 488-1975). This enables its use with a calculator in automatic test systems, allowing complex measurement sequences to be made without operator intervention.

Frequency response can be measured with just the generator and the SLMS. Two modes of operation are possible. In one, the SLMS acting as a controller tunes a compatible generator by way of the HP Interface Bus. In the other, the two operate "open loop", with the SLMS slaving itself to the generator output by way of the signal path.

In open-loop tracking, the SLMS operating in the SPECTRUM mode is set to step through the selected range of frequencies and the generator is set to step through the same range with the same step size but at a slower rate than the SLMS (typically at 1 s per step). The SLMS will not display a measurement result until it makes two successive measurements that lie within selected limits. Once this has been done, it steps to the next frequency and again waits until it is able to complete two more successful measurements. In this way, it follows the signal source without any interconnection other than the signal path.

The open-loop mode can also be used for offset tracking, where the generator and SLMS are not tuned to the same frequency because of a frequency translation in the device under test.

When controlling a generator through the interface bus, the SLMS can also accommodate a situation where the tested device may require an appreciable amount of time to settle down after a change in driving frequency. When the SWEEP switch is changed to the STAB TEST position, the SLMS will not step to a new frequency until it makes two successive measurements that not only lie within the selected limits but that are also within 0.2 dB of each other (within 0.05 dB with the AVERAGING switch set to LONG). This then assures that the item under test has settled at the new frequency before the measurement is completed.

Calculating Non-Standard Pilots

When multiplexed group pilots are being measured, a front-panel switch allows the operator to select the pilot frequency appropriate for the system under test and the SLMS will then calculate and select the correct measurement frequencies. For example, if group 3 of supergroup 13 in mastergroup 5 is selected, the instrument tunes itself to 14,271.91 kHz, the frequency of the reference pilot of that group in the Bell System L4 FDM plan. If, on the other hand, the group number had not been entered, the SLMS would have tuned to 14,175.92 kHz, the supergroup reference pilot in the same system.

Two baseband group pilot frequencies are commonly used in FDM systems: 84.08 kHz and 104.08 kHz. Either of these can be selected by the front-panel PILOT select switch for use by the SLMS when calculating pilot frequencies. There are many other pilot frequencies in use, however, so the SLMS provides for them with a NON-STD position of the PILOT select switch. With this, any other frequency can be entered into a pilot storage register through the numeric keyboard and pressing the NON STD PILOT pushbutton. This frequency is then used for calculating group pilots.

This facility also permits other useful measurements to be made. For example, by entering 0 kHz into the pilot register, the instrument can be set to scan suppressed-carrier channel frequencies, thus measuring carrier leaks. Also, by choice of an appropriate frequency for the pilot register, the SLMS will scan the intersupergroup slots.

Audio Outputs

A demodulated version of the signal in the channel being measured is applied to an internal loudspeaker and to a front-panel connector. The loudspeaker provides a quick method of telling whether the channel is carrying speech, data, tones, or noise. The front-

panel output can be used for other measurements such as jitter, group delay, and transients. The demodulator automatically reinserts a carrier either above or below the signal spectrum as appropriate so the demodulated signal is always "erect" regardless of whether the channel is normal or inverted. This has the advantage of permitting true weighted measurements, either C-message or CCITT psophometric, instead of the usual approximate "equivalent bandwidth" measurements (both of the above-mentioned filters are available as options).

Phase Jitter

An optional feature provides for the measurement of phase jitter on a suitable test tone injected into a voice channel. The reading is displayed on a meter with full-scale ranges of 3° and 30° peak-to-peak. The jitter of carriers and pilots may also be measured.

Instrument Organization

A simplified block diagram of the Model 3745A is shown in Fig. 5. It has three major sections: a precision multiple-conversion superheterodyne receiver, a programmable frequency synthesizer that supplies highly accurate and stable local oscillator signals for the receiver, and a microprocessor-based controller that also includes the LED display. The

microprocessor is itself controlled either from the front-panel keyboard or through the HP Interface Bus. These blocks are described in the following articles.

Acknowledgments

In addition to the various team members mentioned in the articles that follow, special mention must be made of the two product designers, the hardest working members of the team: Harry Elder, who was in charge of the design, and Arthur Thornton. Besides packing a lot of electronics into the instrument while maintaining ease of assembly and access, they coped with at least three impossible requests per day throughout the life of the project. The power supply was designed by Colin Grossart, who also designed a precision attenuator for production test of the SLMS.

Mention must also be made of the project's resident production engineer, Ian Betsworth, who spent two years working with the development team before returning to production with the instrument.

Finally, a large debt of gratitude is owed the many people throughout the telecommunications industry who contributed to the final design with their criticisms and suggestions of early prototypes. Without their help and encouragement, the Model 3745A would not have developed as it did. 

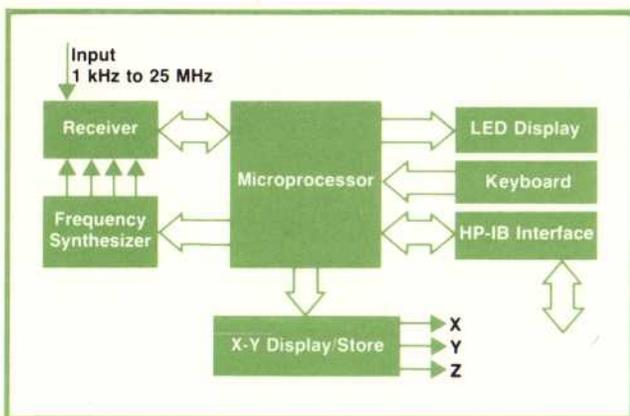
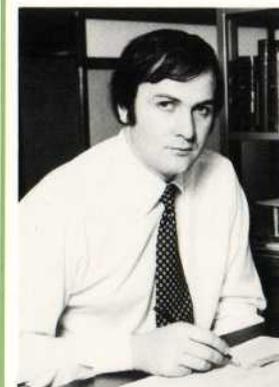


Fig. 5. Organization of the Model 3745A/B Selective Level Measuring Set.



J. Reid Urquhart

Graduating with a BSc degree in Electrical Engineering from Glasgow University in 1965, Reid Urquhart worked two years on satellite communications before joining HP Ltd. He was project leader on the 3710A Microwave Link Analyzer before becoming Level-Measurements Group Leader. The 3745A/B is the first product of that group. He recently transferred to production to serve for a year as line manager in charge of the 3745A/B. Reid and his wife have one small daughter.

ABRIDGED SPECIFICATIONS HP Model 3745A/B Selective Level Measuring Set

FREQUENCY RANGE: 1 kHz to 25 MHz
FREQUENCY ACCURACY:
INITIAL SETTING ACCURACY: within $\pm 1 \times 10^{-7} \pm 1$ Hz
AGING RATE: $< 1.5 \times 10^{-7} \pm 1$ Hz/year

Filters and Measurement Ranges

PILOT FILTER (22 Hz):
MEASUREMENT RANGE: +15 to -125 dBm
RIPPLE OVER 22 Hz BANDWIDTH: < 0.1 dB pk-pk
3-dB BANDWIDTH: 38 Hz $\pm 10\%$
REJECTION AT ± 110 Hz: > 60 dB
REJECTION AT ± 1 kHz: > 80 dB

CHANNEL FILTER (3.1 kHz)
MEASUREMENT RANGE: +15 to -115 dBm
RIPPLE OVER 2.6 kHz BANDWIDTH: < 0.5 dB pk-pk
3-dB BANDWIDTH: 3.1 kHz $\pm 10\%$
VIRTUAL CARRIER REJECTION (± 1.85 kHz): > 55 dB

ADJACENT CHANNEL REJECTION (± 4 kHz): > 67 dB

GROUP FILTER (48 kHz):
MEASUREMENT RANGE: +15 to -75 dBm
3-dB BANDWIDTH: 48 kHz $\pm 15\%$
40-dB BANDWIDTH: 140 kHz $\pm 15\%$
ADJACENT GROUP REJECTION (± 48 kHz): > 25 dB

BROADBAND:
MEASUREMENT RANGE: +15 to -35 dBm
10 kHz - 25 MHz: ± 1 dB Flatness

RETURN LOSS: on 75 Ω unbalanced input > 30 dB (50 kHz to 25 MHz)

RESOLUTION: 0.01 dB or 0.1 dB (selected by switch)

ABSOLUTE ACCURACY: at 0 dBm (at 1 MHz ± 1 Hz): ± 0.05 dB

FLATNESS (referred to 1 MHz and 0 dBm with input signals within 1 Hz of tuning frequency):
50 kHz to 20 MHz: ± 0.15 dB
10 kHz to 25 MHz: ± 0.25 dB
1 kHz to 25 MHz: ± 1 dB

ADDITIONAL ERROR (for measurements in the range +15 to -80 dBm, with respect to accuracy and flatness at 0 dBm)

FOR EACH 10 dB STEP BELOW 0 dBm: ± 0.03 dB
FOR EACH 10 dB STEP ABOVE 0 dBm: ± 0.05 dB
FOR EACH 1 dB STEP ABOVE OR BELOW 0 dBm: ± 0.01 dB
(maximum cumulative error for up to ten 1-dB steps: ± 0.03 dB)

INTERMODULATION AND SPURIOUS PRODUCTS:

INTERMODULATION REJECTION: > 70 dB
SPURIOUS PRODUCTS: either -80 dB with respect to input signal or -115 dBm, whichever is greater.

POWER: 100, 120, 200, 240 V $\pm 10\%$, 48 Hz to 66 Hz, 200 VA.

DIMENSIONS: 425 mm W \times 268 mm H \times 505 mm D (16.8 \times 10.6 \times 19.9 in.)

WEIGHT: 39 kg (86 lb)

PRICE IN U.S.A.: \$23,625.

MANUFACTURING DIVISION: HEWLETT-PACKARD LIMITED
South Queensferry
West Lothian, Scotland

Designing Precision into a Selective Level Measuring Set

by Hugh P. Walker

LIKE ANY TUNABLE RECEIVER, the purpose of a selective level measuring set (SLMS) is to select a wanted signal with minimum distortion while rejecting unwanted signals to the maximum degree possible. In the case of measurements on frequency-division multiplexed systems, the SLMS must measure voice channels 3.1 kHz wide separated from adjacent channels by only 900 Hz and with carrier leaks and pilot signals only 300 Hz from the band edges. It must reject adjacent traffic by at least 60 dB when checking the background noise in a telephone channel, and reject the pilots and carrier leaks, which are usually more than 10 dB below the traffic level, by at least 50 dB.

The high degree of selectivity is achieved, as in a communications receiver, by converting the frequency of the input signal through a series of intermediate-frequency stages to a fixed frequency where the filtering can be carried out effectively. Apart from

that, the design of an SLMS follows rather different principles from the design of a communications receiver because its purpose is to measure signal levels with high accuracy. Thus, instead of an AGC system to adjust the signal level appropriately, it uses precision attenuators.

Even so, errors in level measurements may arise. These can be caused by several sources, including cumulative errors in attenuators, compression in the amplifier chain, and nonlinearity in the final detector. Additional errors may be caused by various spurious responses when low-level signals are being measured in the presence of strong unwanted signals. Described in the box on page 12, these must be considered when planning the block diagram of the SLMS receiver.

Receiver Plan

A block diagram of the receiver portion of the Model 3745A/B SLMS is shown in Fig. 1. The re-

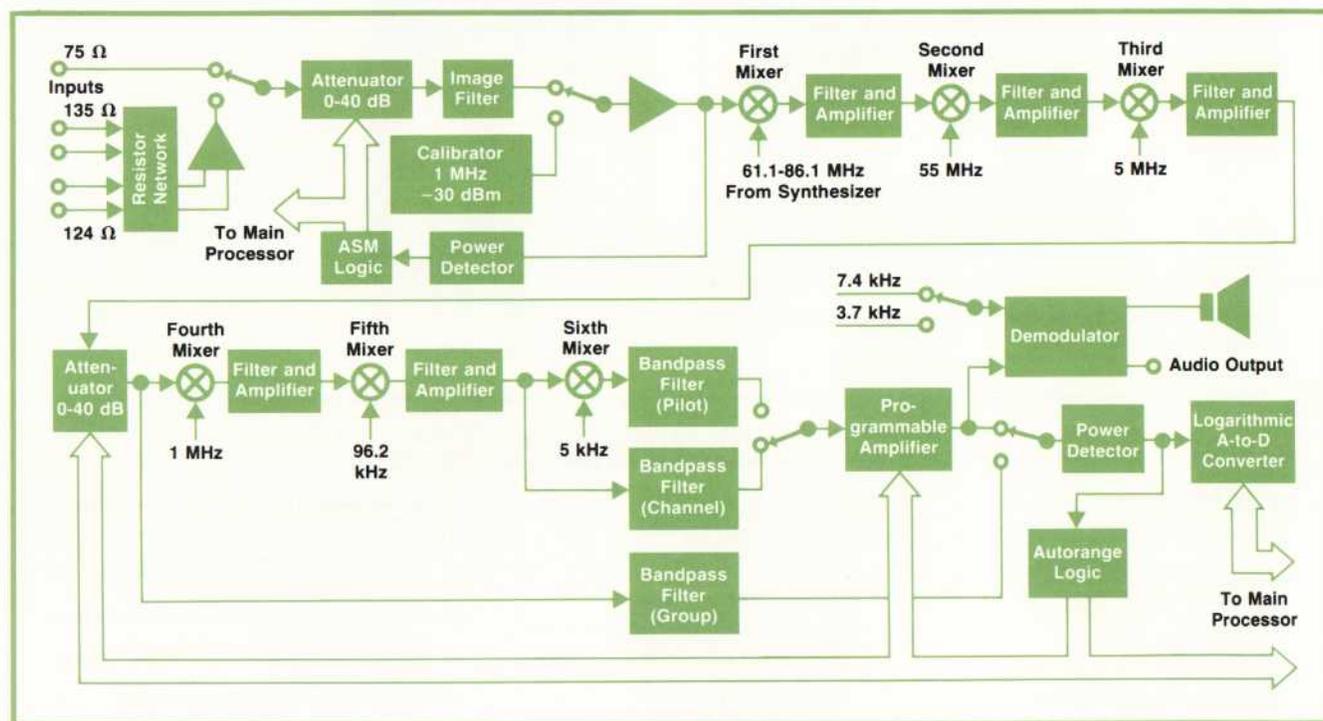


Fig. 1. Block diagram of the receiver portion of the Model 3745A Selective Level Measuring Set.

ceiver has three sections. The first of these is the front end, which includes the autoranging attenuator and the first mixer. The mixer converts input signals between 1 kHz and 25 MHz upwards to a constant 61.10175-MHz intermediate frequency. The second section converts the 61.10175-MHz signal down to 5.55 kHz and 487 Hz for filtering. It includes mixers, amplifiers, the filters, an attenuator and a programmable amplifier for autoranging gain adjustment, and the rms detector. The final section includes a logarithmic A-to-D converter and autoranging control circuits.

Receiver Front End

Two inputs to the receiver are provided. One is an unbalanced 75Ω input and the second is a balanced input, either 124Ω or 135Ω , used mainly with U.S. multiplex equipment. A relay selects one or the other as the active input. The balanced input uses a pair of transistors in a differential-input-to-single-ended-output configuration that rejects longitudinal (common-mode) signals by 40 dB.

A switchable attenuator preceding the first mixer adjusts the input signal level to prevent overload. A sample of the total power reaching the input mixer is taken to a thermopile that provides a proportional dc output. This is checked against high and low thresholds by an algorithmic state machine that in turn controls the relay-switched precision attenuator. The attenuator introduces up to 40 dB of attenuation in 10-dB steps and operates independently of the auto-range circuit, to be described later, so the input circuits are protected against overload at all times. Control lines also feed the main processor to adjust the display of the measured level according to the attenuator setting.

The thermopile output can also be displayed to show the total broadband power in the input signal.

Following the attenuator, a fifth-order elliptic low-pass filter rejects image frequencies in the 120-150 MHz band and also signals at the 61.1-MHz first IF. A fixed bridged-T equalizer compensates for non-flatness in the input cables, attenuator and filters.

The signal next goes to the input amplifier which by means of a "totem-pole" output stage and a large amount of negative feedback achieves low intermodulation, typically -100 dB. Besides buffering the mixer from the input, it provides about 10 dB of gain, thereby improving the receiver noise figure. Adjustable RLC networks in the feedback path compensate for non-flatness at the low and high ends of the frequency band.

The signal is up-converted in the first mixer to a constant 61.10175 MHz, providing a wide separation between the input signal and its image frequency thereby enabling effective rejection of the image fre-

quency. The mixer uses the familiar diode-ring configuration with transmission-line transformers and matched Schottky diodes to maintain a mixer balance of better than 40 dB.

Nonlinearities in such a mixer occur primarily because of phase modulation of the switching function by the input signal. This is minimized by driving the bridge with a 6V p-p square-wave local oscillator signal that has a switching time of only 1 ns. As a result, third-order modulation products are typically only -100 dB.

A profusion of harmonics results from the high-level switching. These could readily cause spurious frequencies with harmonics of the second local oscillator as the receiver is tuned across its range. The spurious frequencies are held below -90 dB by a low-pass filter at the output of the first mixer, by extensive filtering of the power supply lines, and by mounting the input sections in die-cast boxes with bulkhead connectors to reduce ground currents.

IF Section

The 61.1-MHz IF is converted to the final center frequency by a number of frequency conversions, the number of conversions required being determined by the required image rejection and available filter technology. It would be desirable to use as few conversions as possible to minimize the noise, spurious frequencies and non-linearities that each mixer can introduce but this requires a large ratio of adjacent intermediate frequencies, which would then place the image frequencies quite close to the desired frequency. The very high Q of crystal filters would

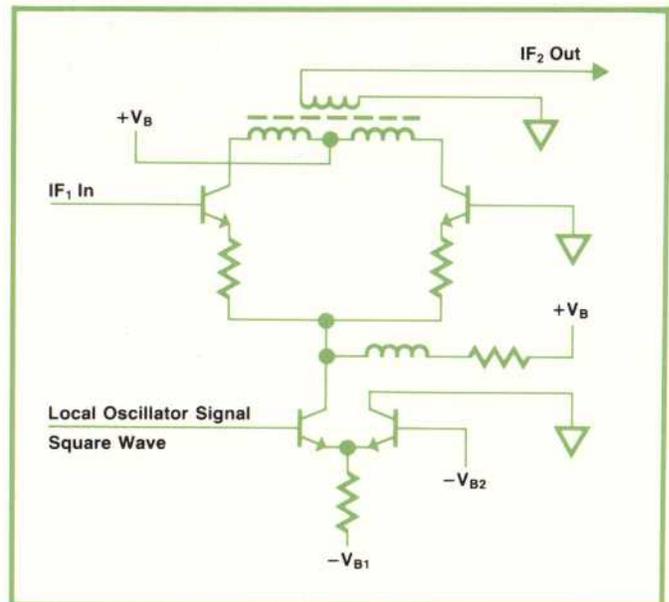


Fig. 2. Transistor mixer uses the common-mode action of a differential amplifier to suppress the local oscillator signal in the output.

enable adequate rejection of close-in image frequencies but it would be very difficult to obtain the flat in-band response required for an SLMS. It was therefore decided to use L-C filters and make adjacent IF ratios 10 or 20 to 1.

The receiver IF's are 61.10175 MHz, 6.10175 MHz, 1.10175 MHz (48-kHz group filter), 101.75 kHz, 5.55 kHz (3.1-kHz channel filter), and 487 Hz (22-Hz pilot filter). The first three conversions require local oscillator frequencies of 55, 5, and 1 MHz, all of which are obtained from the synthesizer reference section. The remaining local oscillator frequencies are 96.2 and 5.063 kHz, both derived from a 6.1568-MHz crystal oscillator.

A diode mixer similar to the first mixer is used for the second mixer. At the lower frequencies encountered in the remaining mixers, transistor switching mixers, of the type shown in Fig. 2, can be used. In this type of circuit, mixing is performed by using the local oscillator signal to switch the current source supplying the differential pair. The local oscillator signal and noise generated during switching transitions are suppressed by common-mode action at the differential output. As a result, for input signals of

700 mV p-p, third-order intermodulation products are approximately -90 dB. This configuration also provides some gain as well as good isolation.

Filtering

Image frequencies and IF feedthrough are rejected by bandpass filters of the type shown in Fig. 3. These have notches below the passband to give about 90 dB suppression of the image frequency of the following mixer. The response of one of the filters is also shown in Fig. 3.

Where possible, a Gaussian filter roll-off characteristic is used so that alignment can be simply a matter of adjusting for maximum response. Simple low-pass filters are also added in some cases to improve the rejection of the higher-frequency harmonics.

Many of the filters use ferrite pot cores. These have good stability and provide high Q but they can also cause intermodulation because of their nonlinear B-H characteristics. This problem was partially alleviated by using larger cores with air gaps. The design goal of less than -90 dB intermodulation products was achieved by holding the signal at a fairly low level. This necessitated keeping thermal noise at an acceptably low level, accomplished by minimizing filter insertion loss and by following each filter with a low-noise amplifier that matches the filter impedance through combinations of shunt and series feedback. The resulting noise figure is about 2 dB.

The measurement bandwidths are provided by filters designed to meet the requirements of FDM communications systems. The group filter, intended for a quick scan of group powers when checking for high levels, is a fourth-order Chebyshev filter with a bandwidth of 48 kHz between -3 -dB points. The channel filter, centered on 5.55 kHz, required an eleventh-order elliptic design to achieve the necessary discrimination. It is flat across 3.1 kHz, rejects

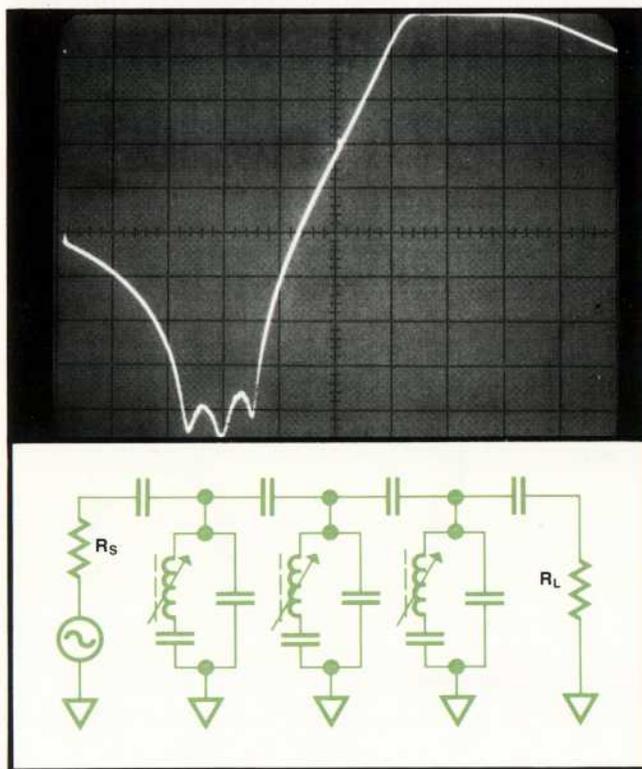


Fig. 3. Configuration of a typical IF bandpass filter used in the new SLMS. The filter has stopband notches below the passband that suppress image frequencies and IF feedthrough. The CRT display shows the frequency response of the 101.75-kHz bandpass filter which has staggered notches to broaden the stopband (vertical scale: 10 dB/div; horizontal scale: 1 kHz/div).

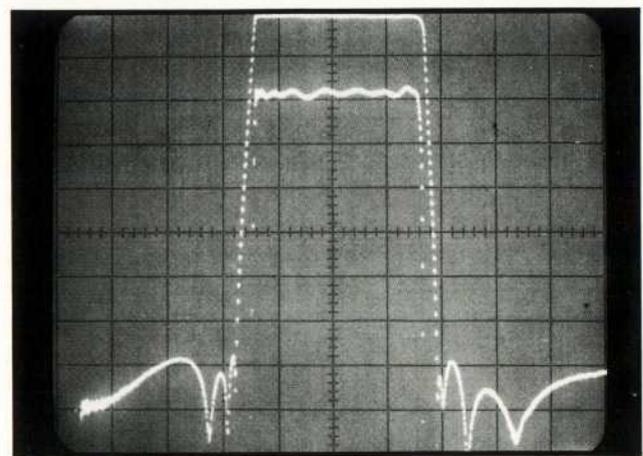


Fig. 4. Double exposure photo shows the response of the channel filter at 10 dB/div and at 1 dB/div. (Horizontal scale: 1 kHz/div).

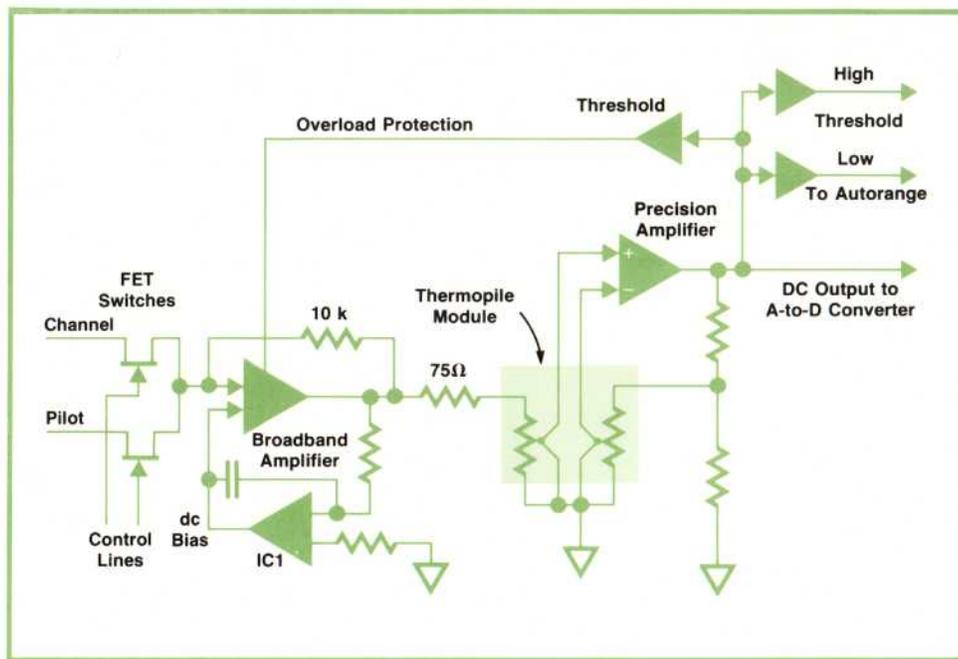


Fig. 5. Power level detector matches a precise fraction of its dc output to the rms value of the input signal in this feedback arrangement. To preserve the excellent linearity of the thermopile (<0.01 dB error over a 10-dB range), a precision IC operational amplifier that minimizes offsets and thermal drift is used.

adjacent channels by more than 67 dB, and suppresses any carrier leak (virtual carrier) by more than 55 dB. It is a lossless, equally terminated ladder filter that uses a two-section, adjustable equalizer to compensate for the drooping response corners caused by finite Q. The response is shown in Fig. 4. Alignment is straightforward using the stopband notches as tuning points.

Single-sideband demodulation of the channel signal is carried out by mixing the filtered channel with a 3.7-kHz (USB) or 7.4-kHz (LSB) carrier and low-pass filtering the result. The carriers are obtained by division of the 6.1568-MHz crystal oscillator frequency.

The narrowband pilot filter is a fourth-order Chebyshev design, flat with less than 0.1 dB p-p ripple across a 22-Hz bandwidth centered on 487 Hz. The 3-dB bandwidth is 38 Hz and adjacent pilot rejection (± 60 Hz) is greater than 40 dB. It is an active filter using integrator-loop resonators that is an exact analog of a passive L-C ladder design enjoying the same advantage of low sensitivity to component tolerances, unlike the conventional realization by stagger-tuned sections. The filter is tuned easily by decoupling each resonator and tuning to 487 Hz.

Level Measurement

Gain adjustment in the IF section is provided in 10-dB steps by a precision 40-dB attenuator in the 1.10175-MHz IF path preceding the 48-kHz group filter. It uses diode-switched π pads. A programmable audio amplifier gives precise 10-dB steps up to 50 dB following final filter selection. Gain switching is under control of the autoranging logic, described later.

Level measurement over the 10-dB range between gain steps is carried out by the circuit shown in Fig. 5.

The filtered signal is applied to the heating element of one of a matched pair of thin-film thermopiles sharing a common substrate. The other thermopile is heated by the divided-down output of a precision dc amplifier. The thermopile outputs are compared at the input to this amplifier so the dc output automatically adjusts to make the two thermopile outputs equal. Thus, the amplifier dc output is proportional to the rms value of the ac signal supplied to the measuring thermopile.

The amplifier that drives the thermopile handles signals from 487 kHz to more than 1 MHz and can reproduce noise peaks 13 dB above the normal 10-dBm maximum output level with low distortion (-80 dB, third order). The output is dc coupled to the thermopile, a zero dc output level being maintained by IC1 in the feedback path.

Measurement Sequence

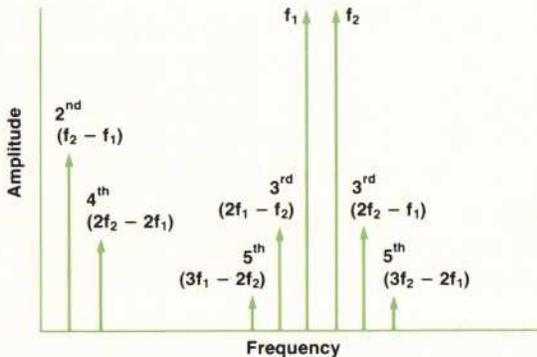
The sequence of events during a measurement is controlled by the autorange circuit. Once the main processor initiates a measurement, it takes no further action until it receives a data transfer signal, at which time it displays the measurement results.

When a measurement is initiated, first the synthesizer is tuned according to the indicated measurement frequency. This takes 16 ms. Tuning is followed by a series of up to nine autoranging steps for setting the IF attenuator and programmable amplifier according to the high and low threshold levels from the thermopile detector. The thresholds are set approximately 13 dB apart to provide enough hysteresis to prevent display instability around the range change points. Each autorange step requires 32 or 64 ms, depending on the filter selected.

Spurious Responses That Can Perturb a Measurement

Errors can occur in measurements of signal level because of the presence of various spurious signals within the measuring instrument. The spurious signals described here must be considered in the design of a selective level measuring set.

Intermodulation and harmonic responses result from non-linearities in the amplifiers and mixers. The non-linearities cause signals lying outside the selected bandwidth to produce low-level signals within the band.



The drawing shows a typical situation involving two signals, f_1 and f_2 , processed by a nonlinear device. The intermodulation products include frequencies that lie close to the signals, the third-order terms ($2f_1 - f_2$, $2f_2 - f_1$) being particularly troublesome. Products higher than fifth order are usually negligible.

Each of the signals at the output of a nonlinear device may be represented by the polynomial:

$$V_0 = AV_{in} + BV_{in}^2 + CV_{in}^3 + \dots$$

For a 10-dB increase in signal level in a given circuit, second-order terms increase by 20 dB and third-order terms by 30 dB. This relationship was used during the development of the new Model 3745A/B SLMS by operating the circuits 10 dB above working level to make the intermodulation products readily measurable on a spectrum analyzer. For a compression in the fundamental of less than 0.01 dB, third-order intermodulation

products should not be greater than -70 dB.

Image and IF responses occur as a result of each frequency conversion. If, for example, a 6-MHz input is mixed with a 5-MHz local oscillator signal to produce a 1-MHz IF, the 4-MHz image frequency will also cause a 1-MHz response. In addition, if 1-MHz were present at the input, it could leak into the IF because of incomplete mixer balance, the usual case.

Both the image and so-called IF responses can be suppressed by suitable filtering, but the greater the ratio between adjacent intermediate frequencies in a multiconversion receiver, the closer the image will be to the wanted signal, and the more difficult it becomes to suppress.

Harmonic responses may also occur. If a 16-MHz signal were present at the input in the above example, it could mix with the 15-MHz third harmonic of the 5-MHz local oscillator signal (strongly present in a switching mixer) to give a 1-MHz IF.

All of these responses can be minimized by bandpass filtering and appropriate choice of intermediate frequencies.

Mixer spurious occur when harmonics of a signal at the input mix with local-oscillator harmonics (particularly odd harmonics in a switching mixer) to give a response at the IF frequency. To return to the previous example, the third harmonic of 8 MHz could mix with the fifth harmonic of 5 MHz to produce 1 MHz.

Harmonics of the input signal may be caused by the preceding circuitry, or by the mixer itself in which case filtering of the input would not help. This is usually a problem only with broadband mixers such as that found at a receiver front end.

Reciprocal mixing is an effect that occurs when a small signal close to a strong, unwanted signal is being selected. The problem occurs when there are phase-noise sidebands on the local-oscillator signal. These are transferred to the unwanted signal during the frequency translation and may then fall within the passband of the receiver even though the unwanted signal itself is rejected. In effect, the local-oscillator phase-noise sidebands have the effect of broadening the receiver passband and may well be the limiting factor in selectivity rather than the shape factor of the filters. Synthesizers, used because of their programmability and better long-term stability, tend to be poorer than a well-designed tuned oscillator and therefore require careful design for this application, as discussed in the article beginning on page 13.

Once an in-range condition is established, further autorange periods are allowed for ensuring that the detector has settled. When a third consecutive period has elapsed in a valid state, the A-to-D conversion is initiated. This conversion combines the functions of averaging, logarithmic conversion, and analog-to-digital conversion. The digital result is transferred to the processor where it is combined with the attenuator settings for display of the measured level.

Acknowledgments

Donald Stewart was responsible for the design of the main IF strip and Boyd Williamson designed most of the signal processing and filtering. Appreciation is also due David Arnold who devised the logarithmic A-to-D converter and the autoranging logic. 



Hugh P. Walker

A graduate of Churchill College Cambridge (BA 1969) and Southampton University (MSc 1970), where he specialized in circuit design, Hugh Walker joined HP Ltd. in 1970. He worked first on the 3710A MLA System and then in 1972 became project leader for the 3745A/B SLMS receiver. Outside work, his main interest is music both as a listener and as a pianist. During the summer he enjoys gardening and travelling by train and bicycle. Winter is the time for reading and woodwork.

Designing a Quiet Frequency Synthesizer for a Selective Level Measuring Set

by John H. Coster

THE HIGH TUNING ACCURACY and resolution required for the Model 3745A Selective Level Measuring Set, along with the requirement for microprocessor control, dictated the use of a frequency synthesizer for the local oscillator.

The Model 3745A was designed so that only the first local oscillator needs to be tunable while the remaining local-oscillator frequencies are fixed. A first local-oscillator frequency range of 61.10275 to 86.10175 MHz was established by the input tuning range of 1 kHz to 25 MHz and the first intermediate frequency of 61.10175 MHz. Because the narrowest filter had a 3-dB bandwidth of 38 Hz, a frequency resolution of 10 Hz was chosen for the synthesizer.

The indirect synthesis technique used in many other synthesizers^{1,2} is also used here because of its simpler structure and lower cost. The relatively slow frequency-switching speed of an indirect synthesizer as compared to a direct synthesizer is not important in view of the time required in the SLMS for filter

settling, autoranging, and A-to-D conversion.

Synthesizer Organization

The general plan of the synthesizer is depicted in Fig. 1. There are three divider loops (N1-N3), two summing loops (SL1, SL2), and one reference loop. All the frequencies are referenced to the master oscillator, a 10-MHz, oven-controlled, quartz oscillator module (HP Model H36-10544) of high stability (aging rate $<1.5 \times 10^7$ /year) and high spectral purity (SSB phase noise more than 115 dB/Hz down 10 Hz from carrier).

The 110-MHz reference-loop oscillator is phase-locked to the master oscillator. It uses a fifth overtone crystal and has a small pull range (± 1.5 kHz or 13 ppm), typical of an overtone crystal oscillator. To allow operation over a practical ambient temperature range, the crystal is enclosed in a specially developed miniature on-board oven set at 80°C. The oven uses proportional control and reaches operating tempera-

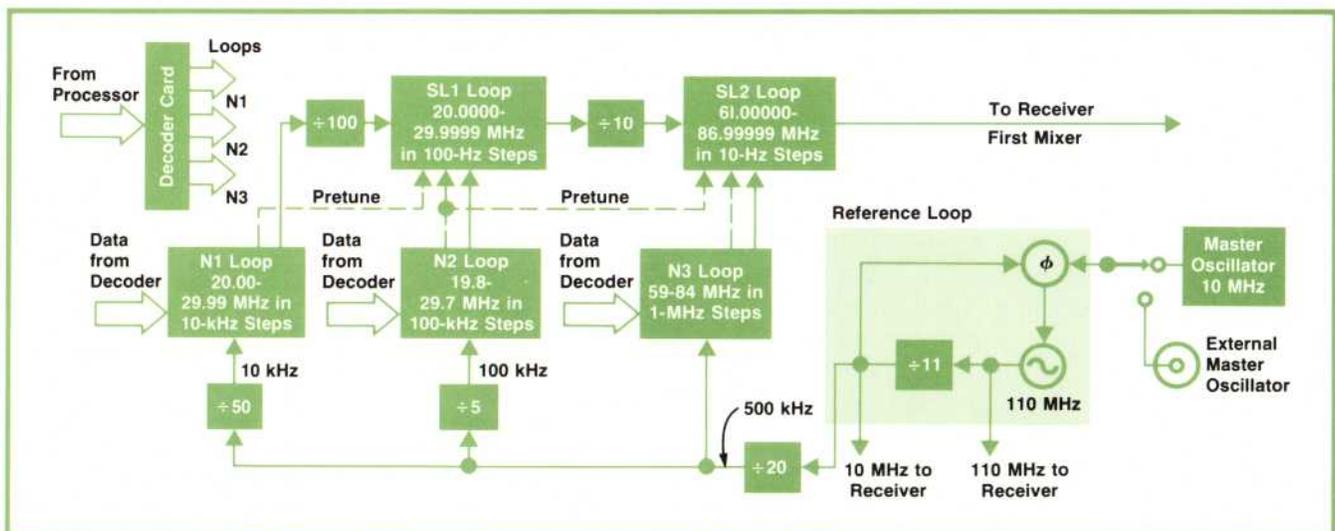


Fig. 1. Block diagram of the synthesizer used for generating local oscillator frequencies in the Model 3745A Selective Level Measuring Set. It is tuned by digit-serial data sent from the instrument's processor and stored in shift registers on the decoder board.

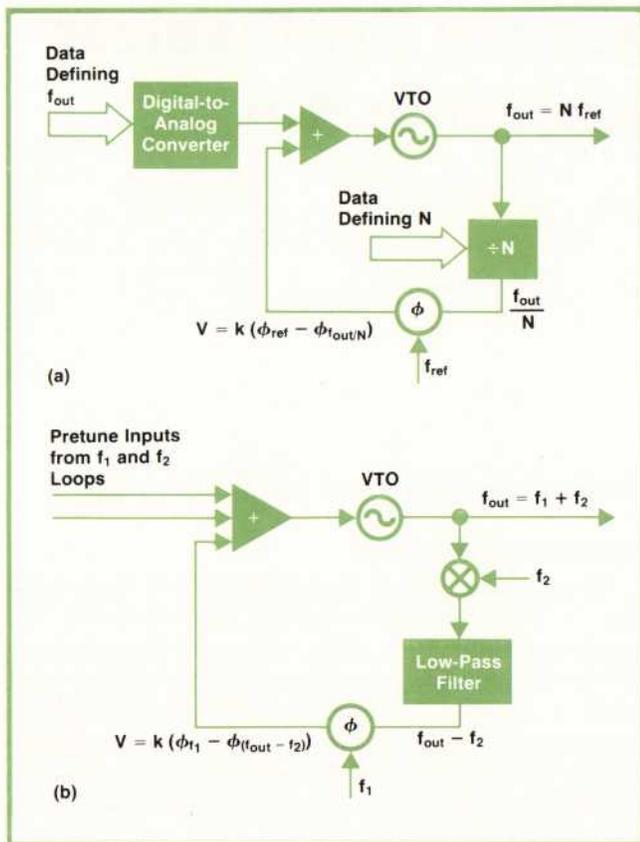


Fig. 2. Divider loops (a) use digital dividers to obtain a sub-multiple of the output frequency. The result of comparing this to the reference frequency controls the voltage-tuned oscillator (VTO), giving an output frequency that is an exact multiple of the reference frequency. Summing loops (b) use a VTO controlled by a phase comparator to derive a frequency that is the sum of a divider loop output and a decade sub-multiple of the previous frequency (see Fig. 1).

ture in about four minutes after switch-on, one minute sooner than the master oscillator.

The 110-MHz reference also supplies the frequencies needed to derive the fixed second, third, and fourth local-oscillator frequencies for the receiver.

Each divider loop consists of a voltage-tuned oscillator that is phase-locked to a reference via a programmable frequency divider, as shown in Fig. 2a. The output frequency is thus a multiple of the reference frequency, the multiplication factor being determined by the data fed to the programmable divider. To speed-up a change of frequency, the oscillator is pretuned by a voltage derived from a digital-to-analog converter controlled by the digital data.

Each summing loop (Fig. 2b) derives a frequency that is in effect the sum of the frequency from its adjacent divider loop and a submultiple of the preceding frequency.

Spurious Signals

Obviously there are several oscillators and many

frequencies involved in deriving the single output frequency of the synthesizer for the SLMS. To achieve accuracy in level measurements, the synthesizer had to meet stringent requirements with regard to spurious outputs and phase noise. A spurious signal that lies within the selected SLMS bandwidth as seen at the first IF would give rise to a receiver output when no input is present. A spurious signal outside this range can give rise to a spurious response by reciprocal mixing if an input is present (see box, page 12). To ensure that the Model 3745A would meet its specifications, spurious signals inside the IF bandwidth had to be more than 102 dB below the synthesizer output, which is at +6 dBm, and those outside the band had to be at least 80 dB down.

For the same reasons, spurious signals accompanying the fixed local-oscillator signals from the reference loop had to be more than 102 dB down.

There are two principal paths by which spurious signals can reach the synthesizer output, assuming no coupling through power supplies or ground loops. The first is by way of the amplifiers in the second summing loop (SL2). The second is via the control voltage to the SL2 oscillator. Any spurious signals on this line will modulate the oscillator, causing unwanted sidebands.

An example is shown by Fig. 3. The synthesizer is tuned to 86 MHz, which causes the output of the N3 loop to be 84 MHz. The 84-MHz signal, typically at +6 dBm, is applied to the L port of the SL2 mixer and it appears at the R port but reduced by the degree of L-to-R mixer balance. It may then reach the synthesizer output, reduced further in level by the reverse isolation of amplifier 1 but amplified by the gain in amplifier 2. The mixer balance is typically 46 dB and the sum of the isolation and gain of the amplifier is typically -75 dB, giving an output of -121 dB referenced to +6 dBm, or an absolute level of -115 dBm. This is well within the requirements of the

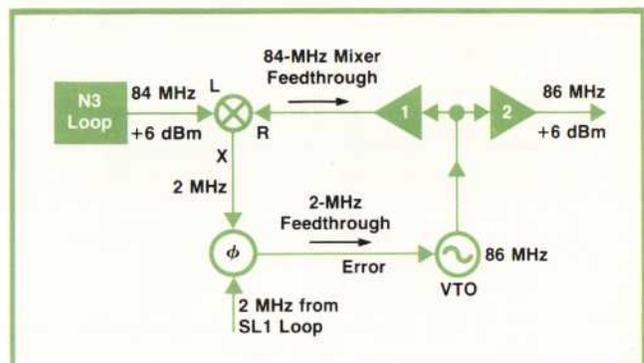


Fig. 3. Spurious outputs can be generated at the SL2 output by direct feedthrough of the N3 frequency by way of the mixer and by modulation of the VTO by feedthrough of the phase-detector switching frequency.

synthesizer.

This same example also illustrates how a spurious output can be generated by way of the oscillator control line (error signal). With the synthesizer tuned to 86 MHz, the phase detector operates at 2 MHz. The phase detector is a digital device that generates a 2-MHz pulse train with the pulse width proportional to the phase difference between the input signals. Thus, there can be a 2-MHz component on the oscillator's dc control line that would frequency modulate the oscillator, giving a first pair of sidebands at 84 and 88 MHz. Obviously, good filtering of the dc control line is a necessity. In the SLMS synthesizer, unwanted signals on this line are rejected 110 dB.

The N3 loop frequencies lie outside the IF bandwidth because they occur in 1-MHz steps so feed-through of these signals causes spurious responses only by reciprocal mixing when an input signal is present. An in-band spurious can be generated in the SL2 mixer, however, causing a spurious response when no input is present. When the synthesizer is tuned to 66.89825 MHz, for example, the N3 output is 64.00000 MHz and the phase detector operates at 2.89825 MHz. The currents circulating in the mixer as a result of these two signals can also mix, producing at the R port a 61.10175-MHz difference frequency, the same as the IF center frequency. If the reverse isolation of amplifier 1 is insufficient, this frequency could appear at the synthesizer output. Measurements indicate that in the new SLMS this spurious is approximately -114 dB referenced to the synthesizer output.

Spurious signals can also be generated by unwanted signals from the SL2 mixer beating with harmonics of the wanted signal in the phase detector. The beat frequency appears on the oscillator control line, and modulates the oscillator. This is prevented by appropriate filtering of the mixer output.

Phase Noise

Phase noise on the synthesizer output would degrade the receiver selectivity by reciprocal mixing, as discussed on page 12. In any of the loops within the synthesizer, phase noise on the input frequencies appears on the loop output, modified by the loop's low-pass transfer function.

In the case of a divide loop, phase noise on the reference input is increased by $20 \log_{10} N$ dB at the output because of the effective multiplication of the reference frequency.³ Five of the synthesizer loops are locked to the 500-kHz reference, so this signal was made as noise-free as possible by use of a crystal oscillator and by use of low-noise Schottky diodes in the phase detector.

The outputs of the N1, N2, and SL1 loops are divided down in frequency, reducing their phase noise.

Thus the N3 and SL2 loops are the main contributors to output phase noise so special care was taken to minimize the phase noise produced by these two loops. In this type of phase-lock loop, the VTO phase noise close in to the carrier is greater than the reference but further away, it drops below the reference. The loop response was therefore designed to cut off at the crossover point so below that point, the noise is limited to that of the reference and above, the noise is that of the VTO. Thus, total noise output is minimized.

The pretune loops are filtered to suppress any noise from the pretune digital-to-analog converters. These filters would normally increase the settling time of the loops, but this is avoided by shorting out the filters during frequency changes.

As a general precaution, the power supplies are heavily decoupled by filters on a separate board to prevent power supply noise generated in the rest of instrument—e.g. the strobed display—from reaching the loop circuits.

As a result of these measures, the synthesizer output phase noise is typically -76 dB referred to the carrier in a 3.1-kHz band centered 4 kHz from the carrier. This is equivalent to about -110 dB/ $\sqrt{\text{Hz}}$, 4 kHz from the carrier.

Acknowledgments

Thanks are due Cullen Darnell, John Gibbs, and Bob Temple, all of HP's Loveland Instrument Division, for assistance given during the initial stages of synthesizer development. 🍷

References

1. J.C. Shanahan, "Uniting Signal Generation and Signal Synthesis," Hewlett-Packard Journal, December 1971.
2. C.A. Kingsford-Smith, "The Incremental Sweep Generator—Point-by-Point Accuracy with Swept-Frequency Convenience," Hewlett-Packard Journal, July 1972.
3. R.B. Shields, "Review of the Specification and Measurement of Short-Term Stability," Microwave Journal, June 1969.



John H. Coster

John Coster graduated from the University College of North Wales in 1967 with First Class Honours and then went on to three years of postgraduate research. He joined HPLtd. in 1971 to work on projects related to microwave link analyzers following which he moved to the 3745A/B synthesizer. Married, but with no children, John's interests include sailing, squash, hi-fi, and woodwork.

Making the Most of Microprocessor Control

by David G. Dack

JUST A FEW YEARS AGO, AN INSTRUMENT that performed the way the Model 3745A Selective Level Measuring Set does would most probably have been under the control of a minicomputer. The built-in frequency plans, the automatic sequencing of measurements, and the automatic control of peripherals could have been realized only by a minicomputer controlling a programmable receiver and a separate synthesizer.

Today, all of the measurement functions are concentrated into a single unit that can be taken to the signals to be tested. The computer has shrunk to three printed circuit boards occupying but a small fraction of the instrument volume. Indeed, there is no need at all for the user to be aware that a computer is there. The all-important human interface is accomplished with pushbuttons labeled with words dedicated to the communication testing problem, rather than with a general-purpose alphanumeric keyboard. The whole impression is one of operating a single instrument, rather than a computer system.

System Architecture

During the design phase, however, the systems

aspects of the instrument were emphasized rather than deemphasized. All the usual systems analysis techniques were required to match computation time to data rates, split tasks between hardware and software, and define precisely the interface rules between various modules.

The processor organization is shown in Fig. 1. The final choice of architecture centered around a commercially-available, 8-bit microprocessor. A direct-memory access scheme was used, however, for the annunciators and 7-segment numeric displays. Without direct memory access, the processor would have had to spend all of its time strobing the displays rather than performing its primary function of calculating the receiver tuning frequencies. Even so, it was necessary to "pipeline" the operations of the processor and receiver so that the processor can be working out the next tuning frequency while the receiver is making a measurement. With the "pipeline" always full, no time is wasted by the processor in waiting for other sections to complete their tasks. In other respects, the basic design of the microprocessor system conforms to current practice and need not be de-

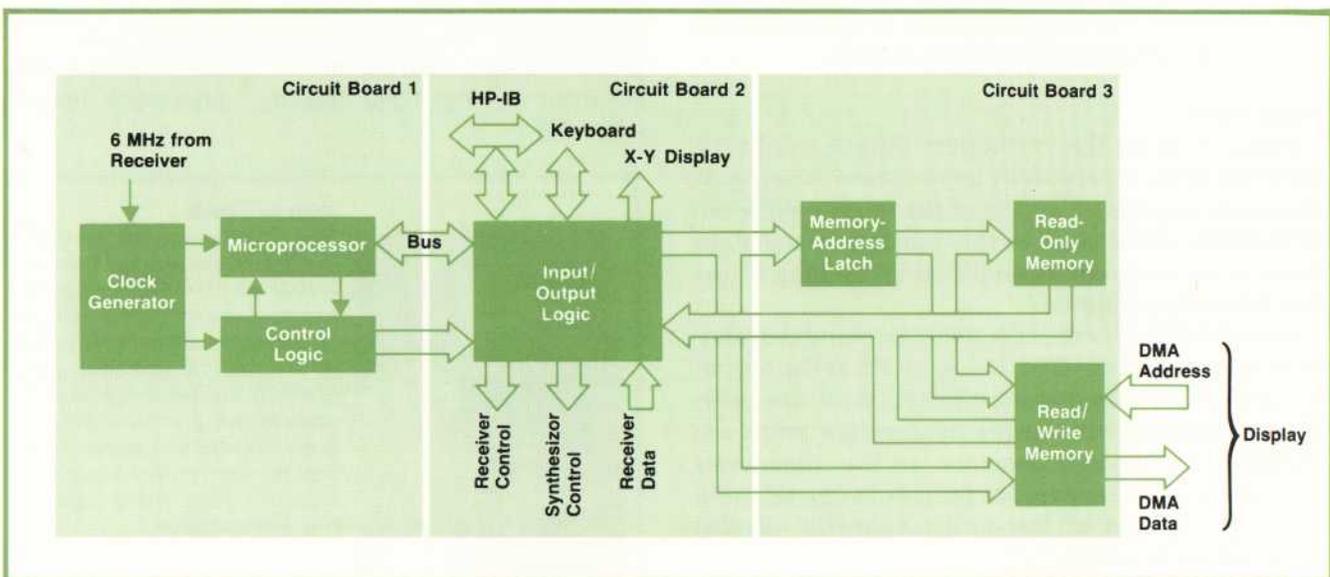


Fig. 1. Organization of the processor in the Model 3745A/B SLMS.

scribed in detail here.

The CPU and supporting chips are on one circuit board. A second board has input/output circuits for distributing control signals to the rest of the instrument. The memory, which consists of five 16K-bit ROM chips and five 1K-bit RAM chips, is confined to a third board. A major advantage of confining all the memory to one board was that it could be replaced by a minicomputer during prototype development. The minicomputer served as a programmable memory that allowed refinement of the software before it had to be committed to mask-programmed ROMs.

The three processor circuit boards are mounted on the back of the swing-up keyboard, forming a unit that can be assembled and tested independently of the rest of the instrument.

LED Display

An obvious approach to activating the twenty-two 7-segment digits and 19 annunciators of the display was to use a strobing technique. The direct memory access takes advantage of the fact that access time of the semiconductor memory is small compared to the time the memory is not being used by the microprocessor.

Each digit or annunciator is allocated a specific location in RAM. A number to be displayed is placed in the appropriate memory location by the microprocessor once only, so no further overhead in the system software is required. The display processor waits until the microprocessor is not using the RAM, then it supplies the data to the display control board for the strobed display.

Interfacing

From the beginning, the HP Interface Bus, Hewlett-Packard's implementation of IEEE Standard 488-1975, was considered an essential part of the instrument. By adapting this standard as a means of controlling peripherals and receiving commands from external controllers, both hardware and software could be made simpler than that required for dedicated interfaces for each peripheral.

The IEEE standard specifies state diagrams that indicate the state transitions and actions to be initiated when interface messages are transferred, but the designer is free to implement these state diagrams in any way available to him. The approach taken with the SLMS was to identify those functions that require responses within a few hundred nanoseconds, which would be too fast for today's microprocessors, and implement them in hardware. The remaining state diagrams were then implemented in software.

Serviceability

There is more to the design task than developing an instrument that can make certain measurements. The designers must recognize that during the lifetime of

an instrument it will need servicing. The design should not only stress reliability to minimize servicing, but it should simplify as much as possible the servicing that is required.

The microprocessor gave the Model 3745A SLMS greater measurement and functional capability but made the instrument more complex technically. However, the microprocessor also makes it possible to include self-test routines that greatly simplify servicing procedures. For example, pressing the key sequence MEAS/HALT TRANS dBm MEAS-HALT causes the instrument to measure its own 1-MHz, -30-dBm, internal standard. For this test to be successful, the three major sections of the receiver—processor, receiver, synthesizer—must be in working order.

Synthesizer operation can be quickly verified independently by unplugging the synthesizer control board and reinserting it upside down. This hardwires the control lines of the synthesizer to give a frequency of 62.10175 MHz, tuning the receiver to 1 MHz. If operation is not correct, each of the synthesizer loops can be checked with a counter to take the fault finding down to the loop level.

The display can also operate independently. When the normal data and clock inputs are disconnected, an internal clock takes over and operates the display. As shown in Fig. 2, a data switch allows static data to be placed on the input lines for checking the single and 7-segment LEDs, and a plug-in card provides a dynamic display for display verification and troubleshooting. In this way, the display can be checked independently and then used with confidence for microprocessor initiated tests of the rest of the system.

Microprocessor Tests

Once the display operation is verified, it may be used with the microprocessor for verifying operation of the microprocessor itself. A bank of switches on the processor board behind the front panel, accessed by swinging up the panel, is used to select simple programs for exercising the processor. Most of the processor functions can be shown to be working correctly by running these programs.

The most important of these is the ROM test. Since the entire operation of the instrument depends on the contents of the read-only memories, it is clearly evident that there must be some way of checking that all 10,000 of the program words are correct. Fortunately, the majority of common faults can be detected at virtually no cost by simply including a single check word in each ROM. This check word, which can be anywhere in the ROM, is selected to make the parity of each column of data in the ROM equal to one. The test program then checks that the parity is indeed one and indicates on the display whether or not each

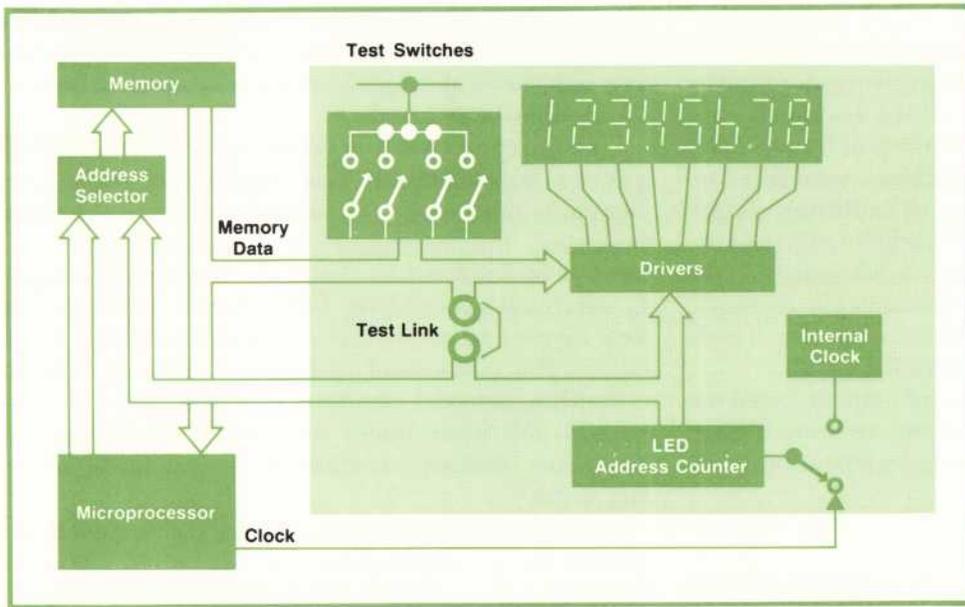


Fig. 2. Direct memory access allows display of measurement information during times that the microprocessor is not using memory. When the display is unplugged from the processor, it can operate from its own internal clock with data placed on the input lines by the test switches or by a link to the address lines.

ROM passes the test.

This simple test detects a surprising number of faults, for example:

- Any single bit error;
- Any word error;
- Any output line stuck at 1 or 0;
- Any address line stuck at 1 or 0.

In addition, there is a high degree of probability of detecting a number of other faults, for example:

- Address lines shorted together;
- Data lines shorted together;
- Multiple random errors.

The ROM test gives a high degree of confidence that the ROM programs are error-free before using them to test other functions, and ultimately to operate the instrument.

A test program for the random-access memory loads an incrementing number into sequential RAM locations and reads the numbers back later, checking for correctness. The program repeats continuously, changing the starting values so eventually every possible combination of data bits is loaded into every RAM location. If a fault is located, the test program loops on that point so the address of the fault can be found with an oscilloscope.

Since the display is also looking at the memory through direct memory access, it displays a sequence of digits when the program is running and halts if a fault is found.

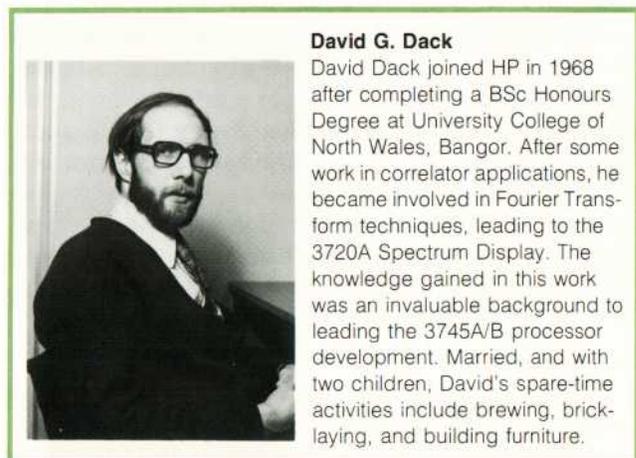
A keyboard test assigns a two-digit display number to each pushbutton; display of the appropriate digits indicates proper operation as each pushbutton is pressed. The program also assigns each switch to one of the digits in the display. As the switch is moved, the digit changes from 0 to 4, for a five-position switch, to indicate proper operation. Since none

of the switches has more than five positions, the number 5 is used to indicate an open circuit or that a switch is stuck halfway between two positions.

Once microprocessor operation has been verified, other programs can be used to test other parts of the system. Another program checks the operation of the optional X-Y display output. The program generates a sequence of consecutive numbers that causes the CRT display to trace diagonal lines. Thus every bit in the D-to-A converters is exercised, and any dropped bits or nonlinearities are immediately apparent.

Acknowledgments

The design team included David Warren, David Arnold, and Kevin Bradford. Robin Myles contributed the software for the interface bus and Ralph Hodgson managed to squeeze the world's FDM plans into a smaller space than was thought possible. 📄



David G. Dack

David Dack joined HP in 1968 after completing a BSc Honours Degree at University College of North Wales, Bangor. After some work in correlator applications, he became involved in Fourier Transform techniques, leading to the 3720A Spectrum Display. The knowledge gained in this work was an invaluable background to leading the 3745A/B processor development. Married, and with two children, David's spare-time activities include brewing, brick-laying, and building furniture.

Real-Time Multi-User BASIC

Easy-to-learn, interactive BASIC is now available as a programming language for HP computer systems operating under the real-time executive operating systems, RTE-II and RTE-III.

by James T. Schultz

BASIC, THE EASY-TO-LEARN, interactive computer programming language, has played a significant role in Hewlett-Packard computer products since 1968. In one form or another, it is available as a programming language on HP 2000-Series Time-shared Systems, HP 9500-Series Automatic Test Systems, HP 8500-Series Network and Spectrum Analyzers, HP 3000 Computer Systems, and the HP 9830A Calculator. However, until recently only a limited, core-based version was available for real-time applications such as sensor data acquisition, laboratory automation, and process control.

Real-time BASIC, a new BASIC system for HP 9600-Series Real-Time Measurement and Control Systems and HP 9700 Distributed Systems, now places at the disposal of the BASIC user all the capabilities of a sophisticated, disc-based, real-time multiprogramming computer system. Real-time BASIC features multi-terminal operation, real-time scheduling of BASIC language tasks, communication between BASIC and non-BASIC programs and data, scheduling of other BASIC or non-BASIC programs from a BASIC program, input and output to a wide variety of peripheral devices, simulation of subroutine calls for checkout of programs, and tracing and breakpoint setting for program debugging.

A primary consideration in the design of real-time BASIC was that it be as quick and easy as possible to develop programs for storing, protecting, and retrieving data. Also taken into consideration was the fact that a great many persons are familiar with HP automatic test system BASIC as well as HP 2000 and HP 3000 time-shared BASIC and their large existing program libraries. For these reasons a high degree of compatibility was maintained between real-time BASIC and existing HP BASIC systems.

Real-time multi-user BASIC operates under HP real-time executive operating systems, RTE-II or RTE-

III.^{1,2} These operating systems provide the BASIC user with access to other powerful languages such as FORTRAN or ALGOL. But the new BASIC is not a real-time language merely because it operates under RTE; what makes it real-time is that intrinsic to it is a scheduler that allows the sequence of execution of BASIC statements to be governed by the time of day or by external events.

9600-Series Systems can have up to four real-time BASIC terminals, each operating independently of the others. Response time of approximately two seconds can be expected from any terminal, depending on what other operations and processes are being executed by the system. Available to each terminal are measurement and control devices and standard peripherals such as line printers and magnetic tapes. Each terminal can lock specific devices so as to isolate the user from interference from other terminals or programs. Each terminal can communicate with other terminals via disc files, provided security codes assigned by the users match. For example, one BASIC program may be taking data and placing it in a disc file while another BASIC program on another terminal may be reading that disc file and formatting a report.

Subroutines and Functions

Frequently a procedure is used several times within a program. Instead of rewriting the procedure each time it is used, BASIC allows the user to write the procedure once and then invoke that segment of code, or subroutine, whenever the procedure is needed. BASIC has two types of subroutines, internal and external. Internal subroutines are written in BASIC and accessed by the GOSUB statement. External subroutines may be written in assembly language, FORTRAN, or ALGOL, and are accessed by the CALL statement.

Any user can attach subroutines and functions to BASIC to be used on his terminal. With this capability the engineer or scientist can write his own special programs in FORTRAN, ALGOL or assembly language, or use one of HP's general-purpose programs provided with BASIC to perform measurements and thereby extend the usefulness of BASIC. The BASIC system may have up to 26 separate and distinct sets of subroutines and functions.

Individual BASIC programmers can modify the repertoire of BASIC-callable subprograms easily without shutting down the system or interfering in any way with other terminals or processes. All of the user-defined subroutines attached to BASIC are generated as disc-resident programs and may operate in a separate partition of memory from BASIC, and therefore do not require BASIC to be swapped to the disc prior to their own execution. Up to 32 separate disc-resident programs may be attached to any terminal, and a total of 64 subroutines and functions may be grouped together in any one disc-resident program (see Fig. 1). Usually the overlays, as the disc-resident programs that contain the subprograms are called, have related routines grouped together. For example, digital input and output subprograms would be grouped with the bit manipulation subprograms, thereby avoiding excessive disc loading time during program loops.

As may be expected on a system that has more than one terminal, the time will come when a user finds he can run BASIC but someone else is using the mea-

surement devices, or possibly he is preparing a program on the central computer of a distributed system for use on a satellite. In either case the user would like to check out and debug his program as fully as possible before running the program with instrumentation. Real-time BASIC has a mode called "simulation" to help the user in these cases. While in simulation mode any subroutine call is executed as a no-operation, the subroutine's name and parameter list are printed out on the terminal, and execution of the program is halted. The user may display or modify any of the variables in the subroutine's parameter list and continue program execution. In this way overall program flow may be checked without instrumentation.

Of primary importance to the engineer is the capability to input and output digital and analog data. Real-time BASIC makes this very simple by providing the user with an extensive library of BASIC-callable subroutines. Included are subroutines to input and output bits, bytes, and words of data, to respond to synchronous and asynchronous interrupts, and to perform conversions and formatting of data. Also provided are routines for communication between BASIC and a wide variety of peripheral devices.

Scheduling Tasks

Real-time BASIC's scheduler makes it possible to execute predetermined tasks at a specified time or upon occurrence of some event. A task may be any group of BASIC statements terminated by a RETURN

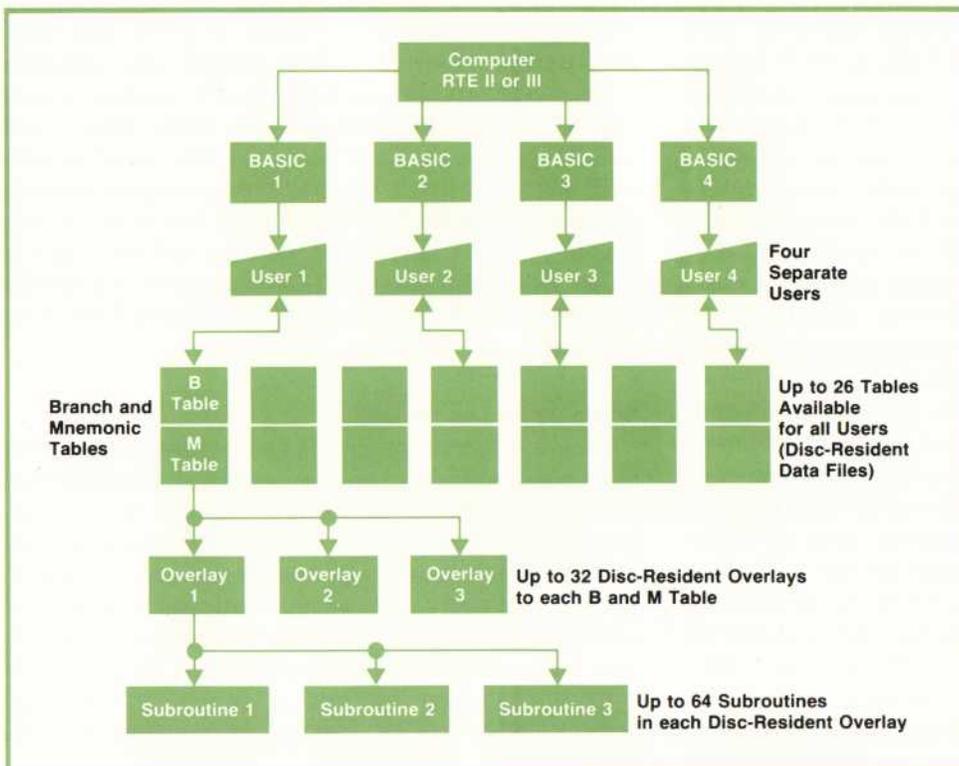


Fig. 1. Real-time BASIC supports up to four independent terminals. Any user can attach subroutines and functions to BASIC to be used on his terminal. The system may have up to 26 sets of subroutines and functions. Each set is identified by a branch table and a mnemonic table and may consist of as many as 32 overlays, each containing up to 64 subroutines and functions.

statement. There are task control subroutines to start a task after a specified delay, to turn on a task at a specific time of day, to execute a task upon receipt of an external interrupt or when a key is pressed at a terminal, to set the priority of execution of a task, and to disable or enable scheduled tasks. The scheduling functions of BASIC provide the means to monitor and log results of processes, respond to alarms, and perform many time-dependent operations. An example of a program to log data on a remote terminal is:

```

10 CALL TRNON (1000, 084500)      Start Task 1000 at 8:45:00 AM
20 P1=11                          Initialize Peripheral Unit Number 11

1000 REM ***** LOGGING TASK *****
1010 CALL START (1000,60)         Restart Task 1000 in 60 seconds
1010 CALL AISQV (5,1,V(1),E)     Read ADC Channels 1 through 5
1020 FOR I=1 to 5
1030 PRINT#P1: "CHANNEL"; I;"=";V(I) Print Voltages On Peripheral Unit #11
1050 NEXT I
1060 RETURN                       Terminate Task 1000

```

This example illustrates task scheduling, instrumentation input and output, and peripheral independence.

Generally BASIC is run interactively from a CRT terminal or teleprinter. Sometimes, however, the user may want to run BASIC under control of a predefined set of BASIC commands stored in a disc file. This might be desirable where a BASIC program developed by a computer professional is being used by manufacturing or clerical personnel. In this case the program developer can prepare a predefined file of BASIC commands using the text editor or file manager. The file contains BASIC commands exactly as if input from the keyboard. The BASIC interpreter is then scheduled either from the terminal or from a computer program to perform the desired function using the name of the command file as a parameter in the scheduling request.

Interpreter Organization

The real-time BASIC interpreter consists of nine modules, a main module and eight segments. Within the main module are routines that perform I/O, formatting, symbol table manipulation, and general-purpose functions required by the segments. The segments perform the following functions:

- Syntax checking and interpretive code construction
- Program and error listing
- Pre-execution checking and table construction
- Program execution
- Program execution of statements that do not require fast execution time such as STOP, END, PAUSE, or ASSIGN
- User command execution
- Program debugging aids

The interpreter is organized in this fashion to provide maximum efficiency in memory use and minimum time spent on disc accesses. Suppose a user requests that a program be executed, using the RUN command. Control is transferred to the command segment, where the RUN command is parsed. Control is then transferred to the pre-execution segments where the necessary tables are constructed and initial verification of the program is performed. Then the execute segment is requested and the program is executed until an END or STOP statement or an error is encountered, whereupon the END or error segment is executed. Total elapsed access time is less than one-half second.

The BASIC interpreter is a table-driven translator. Processing is based on the fact that source statements beginning with letters are commands, and statements starting with numbers are program statements. The first word of a statement or a command uniquely determines which statement or command has been entered; the one exception is the implied LET statement, which is treated as a special case.

Processing begins with a search for commands or statements in the appropriate tables. If no match is found, control is passed to the mnemonic table to determine if the source statement is a subroutine or function call. If there is still no match, control transfers to the syntax routine for the implied LET statement. If a match does occur, the address for the syntax routine is supplied by the statement branch table. Control is passed to that routine to check the syntax and at the same time to build interpretive code for

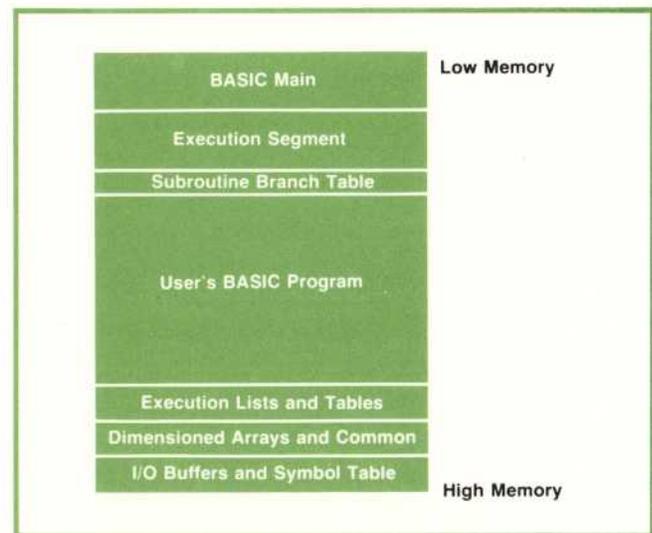


Fig. 2. The real-time BASIC interpreter consists of nine modules: a main module and eight segments. Each user program has its own copy of the main module and the execution segment, arranged in memory as shown here. The interpreter is arranged this way to provide maximum efficiency in memory use and minimum time spent in disc accesses.

the rest of the statement. If the statement calls for a user subroutine, control is passed to the subroutine syntax routine and interpretive code is built using information provided by the subroutine table generator. If the source statement is a command no interpretive code is built and the command is executed immediately.

Source statements are prepared for execution by converting them into operator/operand word pairs for reduced storage space and fast interpretation at execution time. Also, the statements are arranged in sequential order in memory. The program is arranged in memory as shown in Fig. 2.

Multi-Terminal Operation

Real-Time BASIC provides multi-terminal operation; in other words, more than one user may be using a BASIC interpreter concurrently in the same computer system. This is possible because the RTE operating system can swap the interpreter to the disc, or in the case of RTE-III can maintain several copies of BASIC in separate partitions of memory (Fig. 3). Because the interpreter usually has a human being in the loop, with considerably slower response time than a computer, a great amount of idle time is available. Therefore, whenever the user is thinking, control may be swapped from the currently executing BASIC to one that may reside on the disc or in another partition of memory. Studies have shown that the average user think time is approximately 25 seconds between entries, assuming his activity is a mixture of loading, running, modifying and saving programs.³ The average swap time for the BASIC inter-

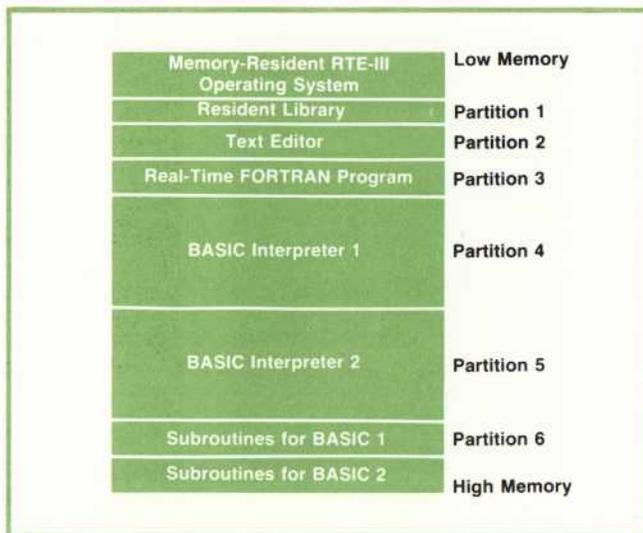


Fig. 3. Real-time BASIC is an option to the HP real-time executive systems, RTE-II and RTE-III. RTE-III can maintain several copies of the BASIC interpreter in separate partitions of memory. This is a simplified RTE-III memory map with two BASIC interpreters simultaneously resident in memory.

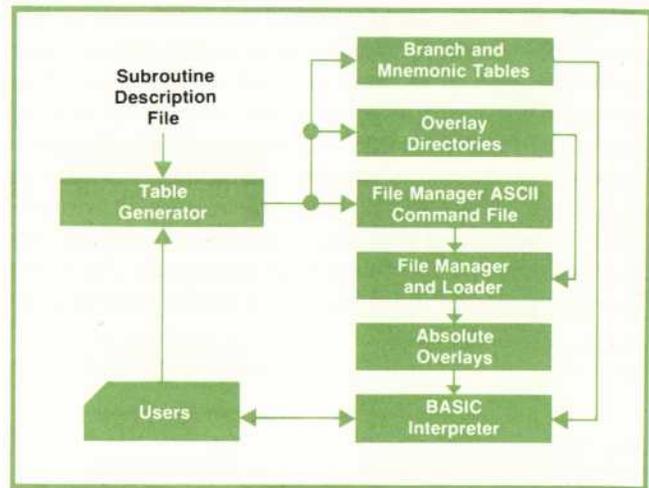


Fig. 4. The process of producing overlays and branch and mnemonic tables is completely independent of the BASIC interpreter. Thus a user can configure his own set of subroutine tables without interfering with other terminals using BASIC.

preter ranges from 100 to 250 milliseconds depending on whether BASIC is running under RTE-II or RTE-III and on what type of disc drive is being used.

The multi-terminal monitor, a module of the RTE operating system, provides each terminal access to the system for program scheduling, peripheral control, and system status monitoring.

Adding Subroutines

Providing a simple method for the BASIC programmer to add special or general-purpose subroutines and functions to BASIC was considered of utmost importance during the design of this product. Definition of a subroutine for use with BASIC requires two distinct phases, a description or table generation phase and a table specification phase. Phase 1 provides a means for the programmer to describe his subroutine: name, parameters used, whether function or subroutine, entry point name if different from name, and the name of the file where it is located. Also specified are the subroutines to be gathered together into each overlay. The subroutine descriptions are placed into a file and then sent to the BASIC table generator, which produces a set of disc files. These are:

- Binary mnemonic table for use during BASIC syntax definition
- Binary branch table for use during execution of the BASIC program
- Relocatable directories for each overlay
- Transfer file which when used with the file manager and loader produces absolute binary overlays from the directories.

The input record to the table generator takes the form:

Name (parameter description), overlay number, partition size (RTE-III only), entry point name, file name, subroutine or function declaration

A specific example might look like this:

```
AISQV(I,J,RVA,R), OV=2, ENT=AISQV, FIL=A2313.
```

The table generation process is shown in Fig. 4. The process of producing overlays and branch and mnemonic tables is completely independent of the interpreter. The table generator, file manager, and loader can all be operated without interfering with any users of BASIC. Therefore, it is possible for a user to configure his set of subroutine tables while other terminals using BASIC are unaware of the configuration process.

The second phase of subroutine definition consists of specifying the newly created tables to the user's specific BASIC interpreter. After BASIC is initiated, the programmer can at any time specify the subroutine tables that are needed for his BASIC program. This is done with the TABLES command. If the user does not use the TABLES command he still can use BASIC but no subroutines will be accessible.

The TABLES command causes the disc files created by the table generator to be loaded dynamically along with BASIC. The mnemonic table is loaded with the syntax segment and the branch table is loaded with the execution segment. Using two tables saves memory space because the mnemonic table is not resident in memory during the execution phase of the BASIC program.

Scheduler Operation

The scheduling of tasks within a BASIC program is performed by the BASIC scheduler. This process requires three data storage areas: the trap table, the

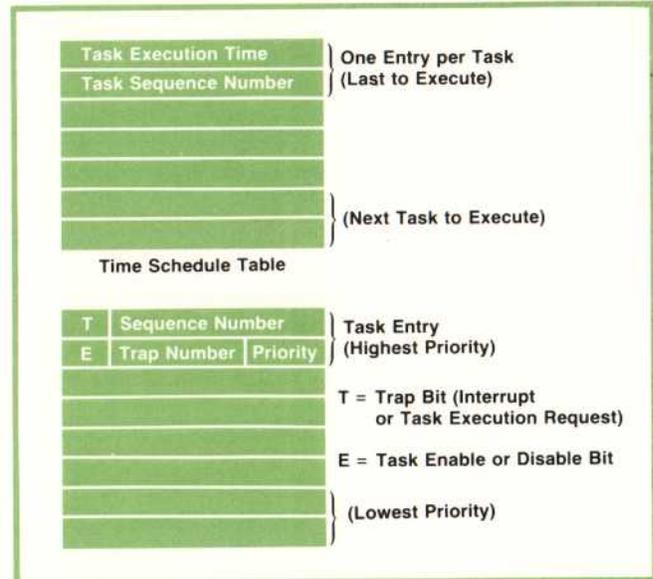


Fig. 5. The BASIC scheduler allows program execution to be controlled by the time of day or by external events. This is what makes the new BASIC a real-time system. Scheduling of tasks requires three tables: the trap table, the time schedule table, and a priority stack. These tables are interrogated before execution of each BASIC program statement.

time schedule table, and the priority stack. The trap table (Fig. 5) has one entry per task, ordered by priority. Each entry contains information associated with the state of a task. The time schedule table (Fig. 5) has one entry per CALL START or CALL TRNON subroutine request, ordered by time of execution. Each entry contains the line number of the task and the time to begin executing that task. The priority stack has an entry for the priority of the currently executing task and one for each suspended task. When a task completes

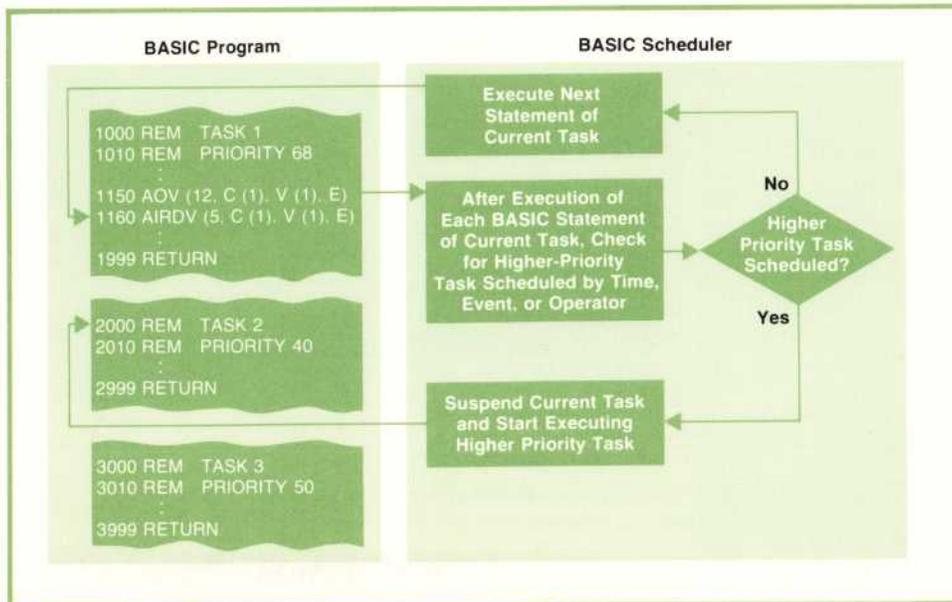


Fig. 6. Operation of the real-time BASIC scheduler as a typical program executes. After execution of each program statement, the scheduler checks for time and event interrupts. It examines task timing and priorities and either continues running the current task or begins a newly scheduled task having higher priority.

execution, the GOSUB return stack in BASIC is checked to get back to the last interrupted task and the priority stack is checked to get the priority of that task.

The trap table, time schedule table, and priority stack are interrogated prior to the execution of each BASIC program statement (see Fig. 6). If after searching the tables a task is found requiring execution then control is passed to that task in the form of an implied GOSUB. The stack discipline is adequate since only a task of higher priority can cause suspension of the currently executing task, and changes of priority do not affect the execution of a task once it has been initiated.

Acknowledgments

The real-time BASIC system would not have been possible without the cooperation of many talented people. The detailed design and development of the BASIC table generator was provided by Adele Gadol, who somehow managed to concurrently develop RTE spooling. Jack Cooley made a tremendous contribution with his design of the real-time BASIC scheduler. Special thanks go to David Means of HP Laboratories who used pre-release versions of real-time BASIC in his integrated circuit test system and made a significant contribution in the development of the subroutine and disc file interfaces. Also special thanks to those responsible for getting BASIC into production: Mary Eicher and Sandy Martensen for writing the manual, Larry Smith and Joe Diesel for preparing training courses, and Van Diehl of product marketing whose encouragement and enthusiasm really made real-time BASIC possible. 

References

1. G.A. Anzinger and A.M. Gadol, "A Real-Time Operating System with Multi-Terminal and Batch/Spool Capabilities," Hewlett-Packard Journal, December 1975.

2. L.W. Averett, "Real-Time Executive System Manages Large Memories," Hewlett-Packard Journal, December 1975.
3. N. Mack and L.E. Shar, "A Supersystem for Basic Time-sharing," Hewlett-Packard Journal, December 1974.

HP Multi-User Real-Time BASIC

ORDERING INFORMATION: The Multi-User Real-Time BASIC system is available as option Y15 for 9600 or 9700 systems equipped with the RTE-II or RTE-III operating system. It is also available as option Y15 for the 92001A RTE-II or 92060A RTE-III Real-Time Executive operating system package, or as the 92101A Multi-User Real-Time BASIC subsystem for use with existing RTE-II/RTE-III systems.

PRICE IN U.S.A.: Real-Time BASIC, \$1000.00.

MANUFACTURING DIVISION: DATA SYSTEMS DIVISION
11000 Wolfe Road
Cupertino, California 95014 U.S.A.

James T. Schultz



Real-Time BASIC project manager Jim Schultz came to HP in 1971 with eight years' experience in the design of control systems, real-time operating systems, and interface hardware and software. His HP projects have included a high-level language for automatic test systems and the BASIC interpreter for the core-based RTE-B operating system. A native of Newport, Rhode Island, Jim attended the University of Tennessee, graduating in 1963 with a BSEE degree.

He enjoys backpacking, flying, woodworking, and bicycling, and has a ten-year-old daughter who shares his activities. He lives in Mountain View, California.

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

HEWLETT-PACKARD JOURNAL
JANUARY 1976 Volume 27 • Number 5

Technical Information from the Laboratories of
Hewlett-Packard Company

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

Editorial Director • Howard L. Roberts
Managing Editor • Richard P. Dolan
Art Director, Photographer • Arvid A. Danielson
Illustrator • Sue M. Perez
Administrative Services, Typography • Anne S. LoPresti
European Production Manager • Michel Foglia

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

209108G EDR862BLACAAA 192
MR C A BLACKBURN
JOHN HOPKINS UNIVERSITY
APPLIED PHYSICS LAB
8621 GEORGIA AVE
SILVER SPRING MD 20910

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.