

JUNE 1980

HEWLETT-PACKARD JOURNAL



Contents:

3 Electronic Distance Measurement for Industrial and Scientific Applications, by David E. Smith *This new electronic distance meter is designed for harsh environments, moving targets, and computer control.*

11 Industrial Distance Meter Applications, by David E. Smith and Troy L. Brown *Proportional and intelligent position control systems, dynamic monitoring, and static monitoring are discussed.*

20 Mass Storage Management—A Unified Approach, by William A. Hanna *The evolving philosophy of mass storage management software for HP 9800 Series Computers.*

In this Issue:



Modern factories and laboratories are full of systems, processes, and tests that run automatically under the guidance of computers or less elaborate controllers. Many of these activities involve moving objects, such as tools or cranes, that have to be positioned precisely and rapidly. To do its job of positioning such an object, a controller needs position feedback that tells it where the object is at any given time. Among the many devices that are used to keep controllers informed are limit switches, photocells, lead screws, optical encoders, and others. The range of choices is wide, but it narrows considerably when the going gets tough. What if the object is red hot? What if it's moving rapidly? What if it's metres or kilometres away? And what if you still need reliable, continuous position feedback that's accurate within a few millimetres?

For many years, Hewlett-Packard's Civil Engineering Division has been making electronic distance meters that help surveyors measure distances precisely under adverse conditions. These instruments measure distance by bouncing a beam of light off a reflector attached to a target. In the new Model 3850A Industrial Distance Meter, this technology has been extended to include the ability to make measurements rapidly and send them to a computer (page 3). Thus the 3850A can act as a computer's "eyes," helping it track targets that are moving, far away, or too hot to touch. With the 3850A you can track airplanes or monitor dams for ominous movements. The article on page 11 discusses its principal areas of application.

Our cover photo this month shows a 3850A application that's now under development in a steel mill (not the one shown, but one like it). Red-hot I beams come out of a rolling mill until they hit a movable stop, or gauge, that has been set to the proper distance from a stationary hot cutoff saw. The saw then cuts the I beam to length. The 3850A will provide the position feedback in a control system that positions the movable stop.

When you do something over and over again, you sometimes get pretty good at it. If you're designing computer programs to help a computer user file, retrieve, and update information, you'd like to make it as easy as possible for the user by providing all the capabilities the user is likely to need, while eliminating obstacles and irritations, such as differences between storage devices. In the article on page 20, Bill Hanna discusses a philosophy that has evolved over many years of designing mass storage management software for HP 9800 Series Computers.

-R. P. Dolan

Editor, Richard P. Dolan • Associate Editor, Kenneth A. Shaw • Art Director, Photographer, Arvid A. Danielson
Illustrator, Nancy S. Vanderbloom • Administrative Services, Typography, Anne S. LoPresti • European Production Manager, Dick Leeksmas

Electronic Distance Measurement for Industrial and Scientific Applications

This new distance meter features a transducer mode in which it measures the distance to a moving target nine times per second for output to a computer or other controller.

by David E. Smith

UNTIL NOW, ELECTRONIC DISTANCE measurement (EDM) has mainly benefited the land surveyor. EDM instruments with the ability to measure several kilometres with millimetre resolution have vastly improved the surveyor's efficiency, productivity, and measurement reliability.¹ These benefits are now made available to the industrial and scientific user by a new Hewlett-Packard EDM instrument, Model 3850A Industrial Distance Meter (Fig. 1).

The 3850A measures distance to a target optoelectronically, using an infrared light beam. It measures the one-dimensional positions not only of stationary targets but also of moving targets with velocities up to 1600 metres per second. The instrument supplies elapsed time in addition to position data at nine measurements per second. When it is interfaced with a computer, noncontact measurements of the target's position can be made, and the target's velocity and acceleration can be computed. The 3850A's wide dynamic range makes measurement of the position and velocity of an airplane 8,000 metres away as practical as controlling the position of an industrial crane to a resolution of 0.001 metre.

Electronic Distance Measurement Technique

The 3850A uses phase measurement of a modulated infrared beam to measure distance to a retroreflector target. Unlike interferometric distance measurement, which detects changes in a target's position,² the phase measurement technique provides absolute distance measurements, independent of optical path interruptions.

The infrared light beam is modulated at one of three frequencies, 15 MHz, 375 kHz or 3.75 kHz. These frequencies have corresponding modulation wavelengths of 20 metres, 800 metres, and 80 kilometres, respectively. Since the light beam is transmitted to the retroreflector and back, the effective modulation wavelengths, λ_e , are half the actual wavelengths. As the retroreflector moves from 0 to λ_e metres from the 3850A, the phase of the received signal with respect to the transmitted signal varies linearly from 0 to 2π radians. For multiples of λ_e , phase measurements roll over, varying from 0 to 2π radians as the retroreflector moves from $n\lambda_e$ to $(n+1)\lambda_e$, where $n = 1, 2, \dots$. To obtain an absolute distance measurement the 3850A measures the phase of each modulation frequency and merges the three readings into one.

The 3850A measures phase with a resolution of 0.6 mil-




Fig 1. Model 3850A Industrial Distance Meter measures distances to stationary or moving targets using rugged, field-proven technology. Its input/output capabilities include extended remote programmability and data output at nine measurements per second.




Fig. 2. Simplified block diagram of Model 3850A Industrial Distance Meter. The phase difference between the external and internal paths represents the distance to the target. The microprocessor performs computation, control, and input/output functions.

liradians (0.04 degree). This high resolution dictates maximum suppression of any drift in the instrument's electronics. To suppress drift, the 3850A employs two optical paths, generated by splitting the infrared light beam into an internal fixed path and a variable-distance external path. Any phase shift caused by the electronics is common to both paths. The 3850A measures the phase of the internal path and subtracts this reading from each external path reading. Thus the common-mode phase drift is subtracted out.

What's Inside

All measurement signals in the 3850A start at the transmitter (see block diagram, Fig. 2). A temperature-compensated crystal oscillator (TCXO) establishes the instrument's reference transmit frequency of 14,987,090 Hz. This frequency determines the overall distance scale factor, and its stability determines the 3850A's accuracy at long distances. The TCXO's output is stable to within ± 1 ppm over the 3850A's operating temperature range of 0°C to +55°C. Maximum drift is $< \pm 1$ ppm per year, guaranteeing long-term accurate operation.

The output of the TCXO is divided down to produce the other two transmit frequencies, 375 kHz and 3.75 kHz. Under processor control, these signals are multiplexed to the lasing diode driver. The lasing diode assembly (Fig. 3) is a proprietary device developed for HP distance measuring

instruments. The diode produces approximately 1.4 mW of infrared optical power at a wavelength of 840 nm. This output is stabilized by an optical feedback loop built into the diode assembly.

The output of the lasing diode is split into external and internal optical paths using an annular ellipsoidal reflector. The center of the diode's diverging beam passes through a hole in the reflector to become the external path. The rest of the diode's output is focused by the reflector's ellipsoidal surface onto the end of a fiber optic bundle to become the internal path. An optical chopper rotating at nine revolutions per second continuously alternates between paths for processing by the receiver.

During the external path time (when light to the fiber optic bundle is blocked by the chopper) the diode's beam is reflected by a 45° mirror and passes through the transmitter lens. This lens collimates the transmit beam, reducing its divergence to within 0.9 milliradians. The maximum optical power of this transmit beam is 60 μ W, well within safe limits for the human eye.

The transmit beam is reflected by a retroreflector attached to the target. The returning beam passes through the receiver lens and is reflected by an annular mirror and the back of the 45° mirror onto the receiver diode. Depending on the distance to the target and atmospheric conditions, the amplitude of the optical signal at the diode can vary from 60 nW to 20 pW.




Fig. 3. The optical source in the 3850A Industrial Distance Meter is a proprietary HP solid-state lasing diode. An onboard optical feedback loop stabilizes the output power.

During the internal path time (when light to the transmitter optics is blocked by the chopper) light travels through the fiber optic bundle. The bundle's output passes through a variable-density rotary attenuator that can be rotated by a motor to equalize the internal and external path optical signal amplitudes. The equalized output from the attenuator is then focused by the optics onto the receiver diode by reflecting off the back of the 45° mirror.

To recap, at the receiver diode we have two optical signals that alternate at 55.5-ms intervals. One signal represents the length of the external, variable path while the other signal represents the length of the fixed internal path.

The purpose of the photoavalanche receiver diode is threefold. First, the diode converts the optical signal into an electrical current. Second, by means of photoavalanche action the diode amplifies this current by a factor of 35. Third, an electrical current generated by local oscillators is applied to the diode to mix the 15-MHz and 375-kHz signals down to the receiver's 3.75-kHz intermediate frequency. Since phase information is preserved in the mixing process, the 3850A's phase detector can operate at a low frequency to produce high-resolution results.

The 3.75 kHz current at the receiver diode's output is converted by a transconductance amplifier to a proportional voltage. This is applied to an AGC stage with a dynamic range of 80 dB. The gain of the AGC stage is adjusted to compensate for external path amplitude variance. The output of the AGC stage passes through a four-pole active bandpass filter with a Q of 12.5. The resulting 3.75 kHz sine wave signal is then processed by the limiter.

The amplitude of the external path signal may vary greatly. The heat shimmer that a person observes on a warm day distorts not only visible light but also infrared. These perturbations of the atmospheric index of refraction tend to disturb the transmit beam, sometimes deflecting it off the target at long distances. This deflection results in up to 100% amplitude modulation of the received signal. To avoid processing data that is outside the limiter's 23-dB dynamic range, amplitude sensors monitor the receiver's

A Versatile Computer Interface for Electronic Distance Meters

The 3850A becomes the fourth Hewlett-Packard electronic distance meter (EDM) with data output capability. All four of these instruments use a common 56-bit serial data output format. This format is well suited for transferring surveying data from an EDM to a field data collection device. However, the serial format is neither optimum nor universal enough for direct interface with a computer.

The 38001A HP-IB Distance Meter Interface (Fig. 1) has been developed to meet the need for an optimum computer interface. The 38001A translates the bit-serial data to a byte-serial format for communication via the HP-IB.* It also features software-definable data formatting and provides latched outputs for 3850A programming.

Flexible Data Format

Using the 38001A's format load state, the controller stores in the 38001A the desired format of data transfer from the EDM. When a reading is made available by the EDM, the 38001A first translates the 56 bits into 14 binary-coded decimal digits. These digits are then transferred byte-serially to the controller according to the predefined format. This format can be structured to transfer only some of the digits, or it may group certain digits together to represent the values of individual variables in the controller. Thus if the 14 digits represent five digits of elapsed time, four digits of distance data, and five status digits, the 38001A's data format can be set to transfer these groups into the controller as the values of variables T, D, and A, for instance. The software-definable format enables the user to decipher data efficiently from the 23 operational modes and functions of the four EDMs.

Although the 38001A has been developed to be compatible with all four HP electronic distance meters, it is especially optimized for use with the 3850A.

The 38001A contains eight open-collector outputs used in programming 3850A modes and functions. Using an HP-IB data byte, the controller specifies the states of these latched outputs.

Another feature is the ability of the controller to directly interrogate the status of both the 38001A and the 3850A. When the controller requests the status of the 38001A, the resulting status byte contains information about the interface along with two 3850A status bits. These bits indicate whether the 3850A is operative, the status of the remote/local switch, and the mode selection: DI or DT.

-David Smith

*HP-IB: The HP Interface Bus, compatible with IEEE Standard 488-1978.



Fig. 1. Model 38001A Distance Meter HP-IB Interface makes four HP electronic distance meters compatible with the HP Interface Bus (IEEE 488-1978).

output. These sensors tell the processor when the signal is in-bounds. The sensors also monitor the receiver output amplitude on the internal path. This information is used by the balance circuit to equalize the internal path amplitude by means of a motor-driven variable-density attenuator.

The limiter acts as a precision zero-crossing detector to "square up" the receiver's output. This square wave is sent on to the phase detector.

The phase detector has several functions. It compares the phase of a 3.75 kHz electrical reference from the transmitter with the phase of the internal and external signals. The average of 50 phase comparisons on each path is transferred to the processor, which takes the difference between successive internal and external comparisons. This difference eliminates common-mode phase drifts and yields the true distance to the target. The phase detector also generates time data for use with moving targets. A time mark signal is issued during the 50 phase comparisons on the external path. The leading edge of this mark indicates precisely when the instrument is measuring a stationary or moving target's position. Also, the elapsed time between successive external path comparisons is accumulated and transferred to the processor.

The processor, like the phase detector, is a multifunctional block. It accepts programming inputs from either the front panel or the remote I/O connector. All key signal timing in the instrument is processor controlled. Moreover, the processor processes distance information differently depending upon which operational mode is selected.

In the distance instrument (DI) mode the processor sequentially averages measurements in all three resolution functions (i.e., using all three modulation frequencies). It takes the three averages, merges them into one number representing the absolute distance, and corrects for instrument offset and atmospheric scale factor. The resulting number goes either to the display or to the I/O connector.

In the distance transducer (DT) mode the processor outputs raw distance, elapsed time, and status information

directly to the I/O connector. It allows the external computer to select either raw distance or signal strength using any of the three resolution functions. The processor also generates status digits containing information from the amplitude sensors at the receiver's output.

Extended Remote Programmability

Although digital output is not unique to the 3850A,* the new instrument is the first HP distance meter to be remotely programmable. Not only are the 3850A's major front-panel controls programmable, but the powerful distance transducer mode has been added for measuring moving targets. In this mode the 3850A outputs 14 digits of information about the target to a computer at a rate of nine measurements per second. Four of these digits represent the target's position. Five digits represent the elapsed time since the previous measurement and may be used by the computer to calculate target velocity and acceleration. The last five digits represent the mode and function programmed and the status of the optical path. The optical path status digit can warn the computer immediately of an obstruction in the optical path, a necessity if the 3850A is used as the position transducer in a closed-loop position control system. The 3850A may be interfaced with computers using the 38001A HP-IB Distance Meter Interface (see page 5).

The basic measurement system for static and dynamic monitoring applications consists of the 3850A, the 38001A, a 9825S Computer and a 62012E Power Supply. Fig. 4 shows the first three of these. Additional HP-IB peripherals such as the 59501A Isolated Digital-to-Analog Converter and the 59306A Relay Actuator are available to expand the basic system for closed-loop position control applications.

Operational Modes

The 3850A may be programmed by the computer to operate in one of three different modes: the distance instrument (DI) mode, the distance transducer (DT) mode, or the self-test mode. The DI mode is used for absolute position measurements on stationary targets. Functions in this mode may be selected either locally by means of the 3850A's front-panel controls or remotely via the I/O connector. Thus a measurement may be initiated either from the 3850A's front-panel MEASURE button or remotely from the computer. Once initiated, the measurement proceeds under complete control of the instrument's microprocessor. After a minimum of nine seconds, the microprocessor outputs a distance measurement accurate to within $\pm(5 \text{ mm} + 1 \text{ mm/km})$ mean square error to either the 3850A's display or the computer. This measurement is a composite of at least 24 raw distance measurements using all three resolution functions.

Another function remotely programmable in DI mode is the environmental correction/return signal strength function. This function allows the user to insert a correction factor for the variation of the speed of light with atmospheric temperature and pressure. This function also allows the operator to observe the strength of the signal being received from the target.



Fig. 4. This basic dynamic measurement system can be used for measurements on moving or static targets. For industrial position control applications a digital-to-analog converter or relay actuator may be added.

*Three HP surveying instruments, Model 3808A Medium Range Distance Meter, Model 3810B Medium Range Total Station, and Model 3820A Electronic Total Station, feature digital output.




Fig. 5. Amplitude and phase characteristics of the 3850A receiver's four-pole IF bandpass filter. The filter delay results in an error proportional to target velocity. This error is compensated internally so the 3850A's distance measurement accuracy is independent of target velocity.

The DI mode is well suited for applications such as surveying where measurement speed is not crucial, targets are stationary, and measurement simplicity is paramount.

Measurement of moving targets requires a fast, flexible operating mode. The distance transducer mode gives the 3850A the capabilities needed to measure the dynamics of a moving target. As mentioned previously, this mode features a high-speed data rate of nine measurements per second. Each measurement consists of four digits of raw distance, five digits of elapsed time between successive measurements, and five digits of instrument status information. The computer may select distance information using any of the three resolution functions: high (1 mm), medium (40 mm), or low (4 m). In the highest-resolution function the positional accuracy of $\pm(5 \text{ mm} + 1 \text{ mm/km})$ mean square error for target velocities up to 40 metres per second is the same as stationary-target accuracies.

The third operational mode is the self-test mode. The 3850A's integrity may be verified using this mode. When self-test is initiated, the processor executes an internal test to confirm that the distance measuring system is operational. If a failure occurs, one of 32 failure codes is returned to the computer. These codes allow the user to pinpoint instrument failures.

Measuring Moving Targets

While measurement of a stationary target requires only a static position measurement, measurement of a moving target demands a distance-plus-time measurement. The 3850A specifies this time-distance coordinate pair in two ways.

The first dynamic method of operation is the time sync method. In this case the 3850A's elapsed timer is synchronized to a master system clock by the system controller. The controller resets the elapsed timer at the start of a measurement sequence. The 3850A subsequently outputs the elapsed time between the reset point and the first distance measurement, the time between the first measurement and the second, and so on. Thus, in spite of the fact that the 3850A measures distance asynchronously, the sys-

tem controller can determine when each measurement occurs in real time.

The second dynamic method of operation is the triggered sync method. In this case the 3850A outputs a time mark signal whose leading edge coincides with the point in time when the distance measurement is perceived to have occurred. Each distance reading that the 3850A outputs to the computer is composed of the average of 50 phase comparisons at 3.75 kHz on the external path minus the average of 50 phase comparisons on the internal path. Thus the 3850A actually determines the average target position over a 13-ms aperture time, i.e., $50 \times (1/3750) \text{ s}$. One might conclude that the leading edge of the time mark should occur at the center of this aperture time. However, because of an interesting combination of the Doppler effect and receiver filter delay, the time mark's position is predictably skewed in time.

Dynamic Target Effects

Let us now determine what effect a moving target has on the 3850A. In its highest-resolution function, the 3850A transmits to the target a signal whose effective modulation wavelength is 10 metres. This corresponds to an effective modulation frequency of $\approx 30 \text{ MHz}$. Now assume that the target is moving towards the 3850A at a constant velocity of 40 metres/second. Because of Doppler shift, the frequency perceived by the receiver is approximately:

$$f_{\text{rec}} \approx \left(\frac{c}{c - u_s} \right) f_{\text{trans}},$$

where c is the speed of light and u_s is the target's velocity. Thus

$$f_{\text{rec}} \approx 30 \text{ MHz} + 4 \text{ Hz}$$

and after mixing down,

$$f_{\text{IF}} \approx 3.75 \text{ kHz} - 4 \text{ Hz}$$

Although this Doppler shift of 4 Hz does not appear sig-

nificant, a considerable error is observed when the signal is passed through the receiver's four-pole bandpass filter. At $3.75 \text{ kHz} \pm 4 \text{ Hz}$ this filter exhibits a linear rate of change of phase lag with frequency of $1.47 \text{ milliradian/(rad/s)}$, as shown in Fig. 5. Since 0.63 milliradian of phase shift represents 1 millimetre of distance, the 4-Hz Doppler shift results in a position error of:

$$X_{\text{Doppler}} = 2\pi f_{\text{Doppler}} \cdot \left(\frac{1.47 \text{ mrad}}{\text{rad/s}} \right) \cdot \left(\frac{1 \text{ mm}}{0.63 \text{ mrad}} \right) \\ = 59 \text{ millimetres.}$$

This result means that for a target approaching the instrument at $40 \text{ metres per second}$, the 3850A will perceive the target to be 59 millimetres farther away than it actually is. Similarly, for a target moving away from the 3850A at 40 metres/second , the instrument observes the target to be 59 mm closer than it is. Must the user account for this error, or can the 3850A automatically correct for it?

The Doppler frequency shift produces a position error equal to approximately $(1.47u_s) \text{ mm}$ if we assume that the time of the average measurement is the center of the aperture time. To compensate for this, the 3850A time mark signal has been moved 1.47 ms ahead of the center of the aper-

Dynamic Distance Meter Performance Testing

Whenever a measurement device is developed that surpasses the performance of existing instruments, one has the problem of how to verify its performance. To test a superior instrument with an inferior one is of little value.

How does one verify positional accuracy of $\pm 5 \text{ mm}$ when a target is moving 40 metres/second ? We selected the following approach:

1. Determine theoretically the performance of the 3850A.
2. Develop two independent test instruments that can be calibrated by existing test equipment. These two instruments must use different techniques to establish data independence.
3. Correlate the data gathered from the test instruments with the theory.

The theoretical analysis is complicated, so an HP 3000 Computer and simulation routines were used. Even then, simplifying assumptions were made. Thus, the importance of actual test data increased.

The first test instrument that was developed was the dynamic phase electronic tool (DPET). The DPET electronically generates a signal with a varying phase to simulate a moving target. The phase profile can be changed by external computer control (HP 9825S) to simulate various static positions, velocities, and accelerations.

This signal is injected into the 3850A receiver, bypassing the optics. The 3850A's measured distance is then transferred to the computer. The DPET also transfers the phase information of the generated signal to the computer where it is converted to distance information. Thus, the distance generated by the DPET and the distance measured by the 3850A can be compared.

Because the DPET electronically simulates moving targets and because the electronics are relatively easy to control, the DPET can test the 3850A's reaction to severe conditions such as acceleration of thousands of g 's.

The DPET is calibrated with the HP 5370A Universal Time Interval Counter.¹

The second test instrument is the dynamic target simulator (DTS). Originally, consideration was given to building this test instrument using folded optics, cyclic motion, and the scanning of a continuous target at increasing distances. All of these approaches require that the data obtained be interpreted in terms of what would be expected from a target moving in a straight line. It was finally decided to build a test instrument that more closely simulates a user's application. This not only minimizes the possibility of overlooking some hidden discrepancy but also provides the means to demonstrate various applications for the 3850A.

The DTS (Fig. 1) has a target sled mounted on a 12-metre I beam. A one-horsepower motor propels the target back and forth along the beam via a cable under the direction of control electronics. Attached to the sled is a length of film with 1-millimetre-wide black and clear encoder bars. As the sled travels along the beam, the film passes through a series of optical source/sensor assemblies that are distributed along the length of the beam. The positions of these optical

source sensors are calibrated using the HP 5526A Laser Interferometer. When the bars interrupt the light path of the optical source/sensors, electrical pulses result. These are decoded to determine distance. Hence, the DTS contains a 12-metre linear encoder with 1-mm accuracy.

The DTS stores the distance of the moving target when a time mark signal is received from the 3850A. This distance is then transferred to the computer for comparison with the distance measured by the 3850A.

3850A testing was intensive and took months to complete. Hundreds of thousands of data points were taken. The final results verify that the 3850A can indeed measure distance in the presence of velocity to the accuracy specified.

Reference

1. D.C. Chu, M.S. Allen, and A.S. Foster, "Universal Counter Resolves Picoseconds in Time Interval Measurement", Hewlett-Packard Journal, August 1978.

-Dean Buck




Fig. 1. The Dynamic Target Simulator measures 3850A dynamic position accuracy by moving a retroreflector at velocities up to $15 \text{ metres per second}$.




Fig. 6. Typical 3850A position accuracy for moving targets. The colored curve is the average of 1998 data points for one 3850A as measured by a special test instrument, the Dynamic Target Simulator.

ture time. For the target approaching the instrument at u_s metres/second, its position at the adjusted time mark signal is $(u_s \text{ metres/second}) \cdot (1.47 \text{ ms})$ or $(1.47u_s)$ mm farther away than the position at the center of the aperture time. Thus, the reading that the 3850A produces now coincides precisely with the adjusted time mark signal. Similarly, the same situation exists for receding targets and varying velocities. Thus the 3850A is able to measure the position of a moving target to an accuracy of $\pm(5 \text{ mm} + 1 \text{ mm/km})$ mean square error with no user velocity correction factor.

Dynamic Accuracy Testing

To prove the dynamic positional accuracy of the 3850A, two special test instruments have been developed (see page 8). One of these, the dynamic target simulator (DTS), has been used to observe hundreds of thousands of data points for position accuracy. An example is shown in Fig. 6. Each data point represents the error between the position of the DTS target at the leading edge of the time mark signal and the 3850A's position measurement. This error is plotted against the target's velocity and a first-order curve is fitted to the data. The extremely small value of the proportionality coefficient is an indication of the 3850A's velocity independence.

The DTS may also be used to demonstrate the accuracy of the 3850A under different velocity profiles. Fig. 7a demonstrates the versatility of the 3850A in determining the position, velocity, acceleration, and the resulting position error versus time of a target under constant acceleration. In spite of the nonconstant velocity profile, the position error remains quite small. Fig. 7b shows a similar test run with the DTS programmed to generate a constant velocity.

Although the DTS is limited by its physical properties to a maximum velocity of 15 metres/second, 3850A accuracy is specified for up to 40 metres/second. The physical limitations of the DTS made necessary an electronic simulation of

higher-speed targets. The dynamic phase electronic tool (DPET) was developed to simulate target velocities up to 180 kilometres/second. Accuracy at these higher velocities has been verified using this device.

The Speed Limit

The 3850A uses three internal resolution functions to measure distance. The distance measurements roll over when the rollover interval for a particular resolution function is exceeded. The rollover intervals are 10 m for the high-resolution function and 400 m for the medium-resolution function. Thus a target at 15.241 m or 25.241 m appears to be at 5.241 m in the high-resolution function. For a moving target the external computer must keep track of when these rollover intervals are exceeded. To accomplish this, the computer must be guaranteed that the target has not traveled more than one-half the rollover interval between successive measurements. This Nyquist speed is nominally 40 metres/second for the high-resolution function, 1600 metres/second for the medium-resolution function, and 160 kilometres/second for the low-resolution function.

Conclusion

In conclusion, the 3850A Industrial Distance Meter has extended electronic distance measurement capability to include dynamic position measurement. With the 38001A HP-IB Distance Meter Interface, the 3850A can be computer controlled to output high-speed position and elapsed time data. Software packages now under development will enhance the system for use in automated position control, dynamic monitoring, and static monitoring.

Acknowledgments

The excellence of an instrument is determined by the people who design and develop it. Therefore, it is quite




Fig. 7. Measurements of the dynamic characteristics of moving targets using the 3850A.

appropriate to thank the following people for their excellence of talent and effort in the creation of the 3850A and the 38001A: Jerry Bybee for his diligence as mechanical design leader, Dave Rustici for the overall electronic design of the 3850A, Dean Buck for the electronic design of the DTS and

DPET and for 3850A dynamic testing, Rick Frey for the mechanical design of the DTS, Troy Brown for 3850A application software, Rod Harris for the overall mechanical design of the 3850A, Jerry Wasinger for the electronic development of the 38001A, Arnold Joslin for industrial de-



David E. Smith

Dave Smith is project manager for the 3850A Industrial Distance Meter and the 38001A HP-IB Interface. With HP since 1970, he's contributed to the design of the 3805A Distance Meter and the 3810A Total Station, and served as project leader for the 3808A Distance Meter and as project manager for the 3810B Total Station. He was born in Ardmore, Oklahoma, and received his BSEE degree from the University of Oklahoma in 1970. Four years later, he received his MSEE degree from Colorado State University. He's married, lives in Fort Collins, Colorado, and is

interested in woodworking, audio design, auto repair, and pipe organ design.

sign of the 3850A and 38001A, and Randy Waite for the original design concept for the 38001A. Special acknowledgments to Bill Smith and Rick Warren for the original design concepts for the 3850A.

References

1. M.L. Bullock and R.E. Warren, "Electronic Total Station Speeds Survey Operations," Hewlett-Packard Journal, April 1976.
2. R.R. Baldwin, G.B. Gordon, and A.F. Rudé, "Remote Laser Interferometry," Hewlett-Packard Journal, December 1971.

Industrial Distance Meter Applications

by David E. Smith and Troy L. Brown

THE PROGRAMMING FLEXIBILITY and fast measurement rate of the 3850A Industrial Distance Meter make it useful in many applications in the areas of automated position control, dynamic monitoring, and static monitoring. For example, when used as a replacement for limit switches in the positioning of industrial cranes, the 3850A can improve reliability, speed, and positioning accuracy while vastly decreasing labor-intensive installation costs. Accurate, noncontact measurement of the dynamic characteristics of automobiles, ships, and airplanes can be achieved optoelectronically, and vital safety information such as the position of an earthen dam can be monitored and relayed to a central location.

In the sections that follow, these three areas of potential 3850A applications—proportional and intelligent position control, dynamic monitoring, and static monitoring—will

be discussed. Software for the 9825S Computer is being developed for each of these areas.

Proportional Position Control

Fig. 1 shows a block diagram of a 3850A system for proportional position control. The computer generates a numerical error proportional to the difference between the desired position and the actual position of the target as measured by the 3850A. This numerical error is transformed to an error voltage by the digital-to-analog converter (DAC). Amplified by a power driver, this voltage forces a motor to position the movable object (with retroreflector attached) so that the error is minimized. The computer can be programmed to provide virtually any system response—overdamped, critically damped, or underdamped—merely by changing the error proportionality




Fig. 1. Block diagram of a proportional position control system using the 3850A Industrial Distance Meter for position feedback. The process to be positioned might be an industrial crane or trolley.




Fig. 2. Position control at low velocities may be impossible if linear velocity gain (K) is assumed, because the actual system has a deadband caused by static friction.

constant or gain.

How is the optimum gain determined? To compute the optimum gain, we must first determine the transfer function P/V , of the process. Working in the complex frequency (s) domain, we observe the following typical transfer function:

$$\frac{\text{Actual Position}}{\text{Input Voltage}} = \frac{P(s)}{V(s)} = \frac{kab}{s(s+a)(s+b)}$$

where k is the object's velocity per unit voltage input, a is the inverse mass time constant of the motor and load, and b is the inverse inductance time constant resulting from the motor's electrical resistance and inductance. The 9825S software being developed for this application assumes a transfer function of this form and automatically determines and implements the optimum gain.

In reality, the system parameters, k , a , and b are not constant but can vary with time, temperature, loading, and

so on. Since this variation severely complicates the determination of optimum gain, most designers elect to treat them as constants. Worst-case values are generally chosen so that the system remains stable. However, the assumption that the velocity gain factor, k , is constant results in severe performance degradation. This gain factor actually has a deadband caused by static friction (Fig. 2). The deadband can result in the inability of the system to move the object for small errors.

Characterization, Linearization, and Optimization

Determination of the overall optimum system gain constant is a three-step process. First, the 3850A is used to characterize the open-loop step response of the motor and load. Second, the nonlinear velocity gain constant is linearized by the computer. Third, the optimum overall gain is derived for the linearized system.

To characterize the system's parameters, the computer applies various step voltages to the motor. The resulting 3850A position and elapsed time data is stored by the computer. Velocity information is computed, and the resulting discrete data points may be plotted as a function of step voltage input (Fig. 3a). Information concerning the mass time constant of the motor and load can be extracted from the step response data and plotted as a function of velocity (Fig. 3b).

The next step is to linearize the system's velocity gain. To accomplish this, a curve is fitted to the discrete velocity versus voltage relationship (Fig. 4a). In the 9825S software being developed for this application, this curve is a ninth-order Taylor series. From this curve the computer then calculates an inverse relationship (Fig. 4b) that produces a linear relationship (Fig. 4c) when multiplied by the nonlinear velocity gain. For fast access, this linearization function is implemented in the 9825S as a lookup table. A curve is also fitted to the mass time constant relationship to find the time constant at zero velocity.

The last step is to determine the optimum value for the overall gain. So far we have assumed that the error sent to the motor is a function of position only. However, the con-





Fig. 3. A software package for the 9825S Computer is under development that will help the user determine and implement the optimum gain for a position control system. The first step is discrete characterization of the process to be controlled. This is done by plotting a series of measured data points as shown here.




Fig. 4. The second step in designing a position control system is to linearize the motor's characteristics. (a) A curve is fitted to the measured data. (b) An inverse curve is computed. (c) The product of curves (a) and (b) is linear.

trol system's response improves if velocity as well as position feedback information is included in this error. An equally valid approach to position-plus-velocity feedback is to make the DAC's output, V , a function of the present and previous errors:

$$V = (\text{position feedback gain}) \cdot (\text{present error}) + (\text{velocity feedback gain}) \cdot (\text{present error minus previous error})$$

or

$$V = G \cdot (\text{present error}) + H \cdot (\text{previous error})$$

where G and H are constants. Thus $G+H$ is the position feedback gain and $-H$ is the velocity feedback gain.

Since the 3850A samples the position of the object under control nine times per second, the system forms a sampled data control loop. To optimize the loop, we must analyze the system in the z domain rather than the s domain. For a transfer function P/V of the form given above, transforming to the z domain, writing the closed-loop transfer function, and setting its denominator equal to zero results in the pole equation:




Fig. 5. Final step in position control system design is to determine the optimum value of the overall gain to minimize response time. To solve the actual system equations in real time to find this minimum takes too much computation time. Therefore the actual system functions are approximated by these universal functions. Optimum loop gain is a function of the time constants a and b , velocity gain k , and sampling period T .




Fig. 6. Detailed block diagram of the control loop using sampled position data from the 3850A. The linearizing and optimizing are done by computer software, making adjustments easy.

$$z^4 + z^3(GkU - 1 - \tau_a - \tau_b) + z^2(GkV + HkU + \tau_a\tau_b + \tau_a + \tau_b) + z(GkW + HkV - \tau_a\tau_b) + HkW = 0$$

where:

$$\tau_a = e^{-aT}$$

$$\tau_b = e^{-bT}$$

T = sampling period

$$U = (\tau_a + \tau_b - 1) \left(\frac{1}{a} + \frac{1}{b} \right) + T + \frac{\tau_a/b^2 - \tau_b/a^2}{1/a - 1/b}$$

$$V = \left(\frac{1}{a} + \frac{1}{b} \right) (1 - \tau_a - \tau_b - \tau_a\tau_b) - T(\tau_a + \tau_b) + \frac{2(\tau_b/a^2 - \tau_a/b^2)}{1/a - 1/b}$$

$$W = \tau_a \tau_b \left(T + \frac{1}{a} + \frac{1}{b} \right) + \frac{\tau_a/b^2 - \tau_b/a^2}{1/a - 1/b}$$

The roots of this equation determine the response time, R , as a function of G and H . Optimum performance occurs when R is minimized, but the determination of this minimum in three dimensions requires excessive computation time. To save time, discrete values of G and H that minimize R have been calculated for about 10,000 different sets of values for a and b , and approximating functions have been determined that best fit the results. These functions are incorporated in the 9825S software now being developed. They are universal in that they are valid for any system of the type shown in Fig. 1 that has 0, 1, or 2 time constants and a transfer function P/V of the form given above. The surfaces of these universal approximating functions are shown in Fig. 5. The computer stores only the coefficients that describe these surfaces and uses these coef-

ficients together with the constants from the characterization routine to establish the best loop gain. The final form of the control loop is shown in Fig. 6.

In summary, the 3850A may be used as a position transducer in both open-loop system characterization and closed-loop system control. A software package for the 9825S Computer will be available that enables the user to characterize, linearize, optimize, and control a system easily and quickly.

Intelligent Control System

The 3850A is well adapted for applications not only in proportional control systems but also in nonlinear control. With the 3850A's continuous position update, a system can be developed that not only characterizes system parameters but also constantly adjusts loop gain to compensate for parametric drift. One such system is shown in Fig. 7.

Although this block diagram is similar in appearance to a proportional control loop, a major distinction is the motor control. This motor has at most two speeds plus a friction brake. The computer minimizes the error between the desired position and the 3850A reading by manipulating the direction, speed, and brake control lines.

This system is similar to an open-loop limit switch control system (Fig. 8). Both systems need to know the system's stopping distance to insure repeatable positioning. However, the 3850A-based system exhibits superior performance in two areas: position repeatability and adaptability to system change. Limit switches provide only discrete distance information. The 3850A will continuously provide position error information so that the system can converge to within a millimetre of the desired position. Moreover, this information is available to the computer for updating the stopping-distance parameter. If the system's




Fig. 7. A nonlinear control system uses a motor controller that is different from the controller in a proportional control system. Continuous position feedback from the 3850A allows the computer to learn changing system parameters and adapt for optimum control.




Fig. 8. A position control system based on limit switches provides information only at switch locations. The system must be shut down for adjustments if inertia, friction, or other parameters change.

friction or inertial characteristic changes, the computer senses the change and corrects the stopping distance. This adaptability eliminates the downtime necessary for re-adjustment of a limit switch system.

Learning System Parameters

At system turn-on the 3850A-based control loop must "learn to walk before it runs." The relationship between stopping distance, S , and steady state velocity is unknown and must be determined. This is accomplished by programming the system to move sequentially to various locations. Like a child learning to walk, the system initially moves with repeated attempts to the desired position. Each time a new attempt is made, data from the 3850A is stored. As data points are accumulated, a curve with basis functions u and u^2 (Fig. 9) is fitted to the data (Fig. 10). Using the curve, the system extrapolates stopping distances so that faster response can be achieved. Thus the system learns how to operate optimally at maximum speed. Fig. 11 demonstrates this learning capability at turn-on and shows how the system adapts to a parametric change.

The majority of industrial positional control systems presently employ two-speed motor control rather than proportional control. In the past it has been difficult to obtain the reliable continuous position information necessary to implement proportional control. The 3850A now makes this feasible.

Proportional control has the advantage of being able to




Fig. 10. To learn the relationship between stopping distance and steady-state velocity, the 3850A-based control loop is programmed to move the object to various locations. Data points are accumulated from these attempts to position the object, and a truncated power series is fitted to this data. Later this series is used to compute and update stopping distance.

position faster for a given accuracy than a two-speed controller. To demonstrate this fact, the one-horsepower motor of the Dynamic Target Simulator (see page 8) was configured with an off-the-shelf proportional (regenerative) motor controller. Software was written to simulate either a truly proportional control system or a two-speed system with a 10:1 speed ratio. With positioning error as the independent variable, several test runs were observed to determine the average amount of time needed by the system to position to the specified accuracy. The results shown in Fig. 12 verify the superiority of the proportional system. For a given positional accuracy the proportional system can position the target at least twice as fast as the two-speed controller. This shows that using proportional control the ultimate speed of a process can be achieved without sacrificing positioning accuracy.

Dynamic Monitoring

The 3850A's nine-measurement-per-second update rate



Fig. 9. Two types of braking phenomena are present in non-linear systems. Total stopping distance S is a superposition of the two mechanisms: $S_T = Gu_0 + Hu_0^2$, where G and H are constants.




Fig. 11. Number of attempts necessary for the 3850A-based control system to position an object to one-millimetre repeatability. Moves 0 through 49 show how the number of attempts decreases as the computer learns the system's parameters after turn-on. Between moves 49 and 50 the motor's deadband was increased and its gain was decreased. Moves beyond 50 show how the computer adapts to this change.

and eight-kilometre range make it an ideal instrument for measuring a moving vehicle's dynamic characteristics. The equipment for one such application, measurement of an aircraft's velocity, is shown in Fig. 13. As the aircraft flies towards the instrument, an operator aims the 3850A and holds it on target. For this application the 3850A is programmed for its medium-resolution function (0.04-metre resolution, 400-metre rollover interval). Raw data is stored for post-processing by a standard software pack.

One of the main problems encountered in this application is holding the instrument on target. Since perfect track-

ing of the target is difficult, lapses in data occur. This periodic loss of data has two effects on the 3850A: degradation of signal-to-noise ratio in computing velocity and uncertainty in rollover correction.

The 3850A measures position and elapsed time. The computer can calculate velocity, u , from this data:

$$u_k = \frac{X_k - X_{k-1}}{T_k - T_{k-1}}$$

where X_k and X_{k-1} are present and previous position measurements, and T_k and T_{k-1} are corresponding times. Since this computation is comparable to differentiation, any noise in the measurement will be greatly amplified. To offset this noise gain we can average n measurements to obtain:

$$u_{avg} = \frac{1}{n-1} \sum_{i=2}^n \frac{X_i - X_{i-1}}{T_i - T_{i-1}}$$

We might alternately choose to average the velocity measurements by taking only the first and last readings:

$$u'_{avg} = \frac{X_n - X_1}{T_n - T_1}$$

Not only is this formula simpler, it has less noise gain than the previous computation for interrupted data. However, the second computation is more vulnerable to error caused by rollover uncertainty.

When the rollover interval is exceeded, the distance reading rolls over, as explained in the preceding article. Fig. 14 illustrates this phenomenon. Rollover must be taken into account by the computer for both position and velocity measurements. Since the 3850A samples the target's posi-




Fig. 12. Results of tests using HP's Dynamic Target Simulator show that a proportional control system can position an object with given accuracy twice as fast as a two-speed system.



Fig. 13. Equipment used to measure the velocity of an airplane. The 3850A was an experimental model modified to produce an audible tone when pointed at the retroreflectors. The fixture holding the 3850A was also a special experimental model. (a) Experimental dynamic monitoring system. (b) Mounting of retroreflectors on the aircraft's wing.

tion, there is a Nyquist speed, u_N , above which aliasing occurs. This speed is:

$$u_N = \frac{X_{\text{roll}}}{2T}$$

where X_{roll} is the rollover interval (10 m or 400 m) and T is the sampling interval. In the medium resolution function for successive readings:

$$\begin{aligned} u_{N \text{ ideal}} &= \frac{400}{2(1/9)} \\ &= 1800 \text{ metres/second}^* \end{aligned}$$

*This is specified as 1600 metres per second to allow for deviations from the nominal measurement rate of nine per second.




Fig. 14. Exceeding the 3850A's speed limit causes aliasing in the velocity measurement. For accurate measurements, the target must not travel farther than $\frac{1}{2}X_{\text{roll}}$ between samples.

However, if misaiming causes the loss of m readings:

$$u_N = \frac{u_{N \text{ ideal}}}{m+1}$$

The Nyquist rate has been reduced by a factor of $m+1$ from the ideal. Thus, loss of readings can restrict the useful velocity range of the instrument.

To counteract this problem, we can take advantage of the fact that the first velocity averaging formula has good rollover uncertainty rejection. We can compute an approximate velocity:




Fig. 15. Velocity of an airplane measured by an experimental 3850A system. The curves demonstrate the results of averaging using techniques discussed in the text.



Fig. 16. Equipment used to monitor earth movement includes stable mountings for both the 3850A and the retroreflector.

$$u_{\text{approx}} = \frac{1}{j-1} \left(\sum_{i=2}^j \frac{X_i - X_{i-1}}{T_i - T_{i-1}} \right)$$

where $j < n$. This approximate velocity can be used to compute a correction factor α .

$$\alpha = \frac{X_1 + u_{\text{approx}}(T_n - T_1) - X_n}{X_{\text{roll}}}$$

When α is rounded to the nearest integer, A , it can be used as a rollover correction factor in the second velocity averaging equation. This equation now becomes:

$$u'_{\text{avg}} = \frac{X_n - X_1}{T_n - T_1} + \frac{A \cdot X_{\text{roll}}}{T_n - T_1}$$

These formulas will be incorporated into the dynamic monitoring software pack being developed for the 9825S Computer. The results of an experimental test run are shown in Fig. 15. During this test the distance from the 3850A to the aircraft varied from three to five kilometres.

Static Monitoring

In recent years the failure of earthen dams has emphasized the need for monitoring earth movement. This type of monitoring has typically involved the periodic surveying of points on the dam's structure. With its good stability, eight-kilometre range, and computer interface, the 3850A can now provide continuous, automatic monitoring for such structures.

Fig. 16 depicts the equipment for a typical dam-monitoring situation. A retroreflector is mounted on the dam's surface with the 3850A mounted far enough away

from the dam so that it is referenced to stable rock strata. In the distance transducer mode, the 3850A can furnish up to 32,400 position measurements per hour to the computer. By monitoring atmospheric temperature and pressure transducers, the computer can automatically correct the distance measurements for environmental changes. By averaging measurements, earth movements less than 1 millimetre can be detected. In the event of unexpected movement, the computer can relay an alarm to a central location via a common-carrier interface.

Long-term stability testing of the 3850A is presently in progress. To date, a drift of 1 mm in one month has been observed for one test instrument.

Troy L. Brown



Troy Brown is developing software for monitoring and control applications of the 3850A Industrial Distance Meter. His work has resulted in a patent application on the optimization approximating functions described in the accompanying article. A member of IEEE and the IEEE Automatic Control Society, he came to HP in 1979, having designed and published on servo electronics for high-resolution Michelson and field-widened interferometer spectrometers. Troy was born in Logan, Utah and attended Utah State University, graduating with BS and MS degrees in electrical engineering in 1978 and 1979. Now living in Loveland, Colorado, he's married and has an 18-month-old daughter. He's fluent in German, teaches Sunday school, and enjoys woodworking, gardening, and motorcycles.

SPECIFICATIONS

HP Model 3850A Industrial Distance Meter

	DISTANCE TRANSDUCER (DT) MODE ¹	DISTANCE INSTRUMENT (DI) MODE
POSITIONAL REPEATABILITY: ² (mse, 6 measurements averaged)	1 mm	N/A
VELOCITY ACCURACY: (mse, computed from successive measurements)	$\pm(0.05 \text{ metres/second})^3$	N/A
READING RATE:	9 measurements/second	Single Measurement: 9 seconds/measurement, minimum Repeat Measurement: 2 seconds/measurement, minimum
RESOLUTION (ROLLOVER INTERVAL):	High: 1 mm (10 m) Medium: 40 mm (400 m) Low: 4 m (40 km)	1 mm (40 km)
REMOTELY PROGRAMMABLE MODES AND FUNCTIONS:	DI vs DT mode Resolution function Return signal strength Reset elapsed timer	Single measurement Repeat measurement Standard deviation Environmental correction/return signal strength Initiate measure Transfer display Self test
DIGITAL OUTPUT: (56 bit serial, grouped as 14 BCD coded digits)		
DATA:	Elapsed time: 5 digits Raw distance: 4 digits	8 digits plus sign
STATUS:	Resolution selected Optical path status	Function selected Solid/flashing display Feet/metres
TIME MARK SIGNAL:	Single line, open collector output; leading edge denotes measurement time.	N/A

NOTES:

1. Unless otherwise stated, all DT mode specifications apply only for the high-resolution function under optimum optical conditions. Moving target measurements are referenced to time mark signal with target velocity less than 40 metres per second and acceleration less than 34 metres per second per second.
2. Positional repeatability is the ability of the instrument to repeat a measured distance under the same conditions within a 24-hour interval, at a constant ($\pm 1^\circ\text{C}$) temperature, distance less than 100 metres. Repeatability and resolution are functions of the number of measurements averaged.
3. High-resolution function

General

POSITIONAL ACCURACY:

(mse, 0° to 55°C , single measurement)
 $\pm(5 \text{ mm} + 1 \text{ mm/km})$

RANGE

MINIMUM: 0.5 metres.
 MAXIMUM (under good conditions):
 8 kilometres to 2 triple prism assemblies,
 2 kilometres to 1 single prism.

STABILITY: (mse, $23^\circ\text{C} \pm 1^\circ\text{C}$, less than 100 metres) $\pm 0.3 \text{ mm/day}$.

MAXIMUM TARGET VELOCITY (DT mode only):

HIGH RESOLUTION FUNCTION: 40 metres/second.
 MEDIUM RESOLUTION FUNCTION: 1600 metres/second.

ELAPSED TIME COUNTER (Transducer mode only):

RESOLUTION 0.1 millisecond.
 ROLLOVER INTERVAL: 10 seconds.
 RESET:

Internally: time reset synchronously with time mark.
 Externally: time reset asynchronously.

LIGHT SOURCE:

TYPE: Solid state GaAs lasing diode (non-visible).
 POWER OUTPUT: Complies with DHEW radiation performance standards, 21CFR, subchapter J and German UVV/VBG and VDE 0836 safety standards.
 BEAM DIVERGENCE: nominally 0.9 milliradians, full angle.
 ENVIRONMENTAL:
 Operating temperature: 0°C to $+55^\circ\text{C}$.
 Storage temperature: -40°C to $+70^\circ\text{C}$.

INSTALLATION:

POWER REQUIRED: 11 to 15 Vdc, 14 watts, ungrounded.
 MECHANICAL INTERFACING: optional 11435B Industrial Yoke.
 AIMING: 18x, erect image telescope.

DIMENSIONS: 150 mm $\times$ 290 mm $\times$ 348 mm (5.9 $\times$ 11.4 $\times$ 13.7 in).

WEIGHTS: 3850A: 6.6 kg (14.5 lb).

11435B: 6.1 kg (13.4 lb).

PRICES IN U.S.A.: 3850A, \$14,500; 11435B, \$595.

HP Model 38001A HP-IB Distance Meter Interface

INSTRUMENT COMPATIBILITY: 38001A interfaces with the following HP distance meters

3808A Medium Range Distance Meter
 3810B Medium Range Total Station
 3820A Electronic Total Station
 3850A Industrial Distance Meter

DISTANCE METER FUNCTIONS SUPPORTED:

ALL INSTRUMENTS: 56-bit serial data output, grouped as 14 binary coded decimal digits.
 3808A and 3810B: Initiation of remote measurement and data output.
 3850A: All remotely programmable modes and functions.

HP-IB FUNCTION SUBSETS SUPPORTED: SH1, AH1, T6, TE0, L4, LE0, SR1, RL0, PP0, DT1.

FORMAT LOAD STATE: Controller may define grouping, order, and amount of data from distance meter.

TRIGGERING:

EXTERNAL: controller must initiate each data transfer between the 38001A and the distance meter.

INTERNAL: data is automatically transferred from distance meter to 38001A.

TIME MARK OUTPUT: Open collector output whose leading edge indicates time of measurement (3850A only).

PROGRAMMABLE OUTPUTS: 8 open collector, latched outputs may be programmed by controller.

ENVIRONMENTAL:

OPERATING TEMPERATURE: 0°C to $+55^\circ\text{C}$.
 STORAGE TEMPERATURE: -40°C to $+70^\circ\text{C}$.

POWER REQUIREMENTS:

VOLTAGE: -10% , $+5\%$: 100V, 120V, 220V, 240V; switch selectable
 FREQUENCY: 48-440 Hz.
 POWER: less than 20 VA.

DIMENSIONS: 213 mm $\times$ 102 mm $\times$ 273 mm (8.375 $\times$ 4 $\times$ 10.75 in).

WEIGHT: 3.4 kg (7.5 lb).

PRICE IN U.S.A.: \$1850.

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

Mass Storage Management—A Unified Approach

After several years of work on mass storage management software at HP's Desktop Computer Division, a generalized philosophy has emerged.

by William A. Hanna

SINCE THE EARLIEST DAYS of computers, designers have had to find new ways of storing information, particularly segments of programs, that did not fit into the main memory. This has led to continuing advances in overlaying and segmentation of software. Data storage, on the other hand, did not receive much attention initially, but when computers were applied to data processing and file management, mass storage devices became a necessity. These devices have evolved in size and speed in the last 30 years from paper cards, paper tapes, and magnetic cards to magnetic tapes and magnetic discs with storage capacities in the hundreds of megabytes and transfer rates in the megabytes/second range.

It has always been the case in the evolution of computers that hardware design progresses very fast, while somewhat less attention is given to improving the software. Although one should expect more progress in hardware design, it is my opinion that mass storage management software has not yet reached its full potential in taking advantage of state-of-the-art hardware. Thus it seems that it is now the software designer's turn to take a hard look at the improved hardware capabilities of the latest mass storage devices, and to decide how to take advantage of them. For example, some discs are used as surface-mode devices (that is, the software organizes the files sequentially on one surface at a time), when in fact they could be better used as cylinder-mode devices by arranging the files to span more than one surface of the disc pack. Since the head assembly of the disc drive moves as one unit, this change can reduce the access time by a factor equal to the number of surfaces in the pack in some applications.

Another pitfall in software design is to regard the individual mass memory devices as separate entities and to specify drives on an individual basis. This approach wastes valuable development time and address space.

A unified mass storage approach can speed up data processing considerably, and so it is worthwhile to focus on managing mass storage efficiently and effectively. One approach is to consider each disc as a one-dimensional array, or vector, of records, choosing the size of each record as the smallest burst of data that has to be addressed at any given time. This statement is subject to the following restrictions:

1. The size of each record must be an integer multiple of the smallest disc-sector size in the system.
2. To accommodate the expected differences in the transfer rates of different mass storage devices, the computer mainframe should have sufficient space in primary storage to

buffer at least one full record, and it is preferable to have a double buffer area to be able to execute a closed-loop verify.

A Unified Mass Storage Approach

Now let us consider a possible system of mass storage management in some detail. Many of the features to be described have been implemented in the mass memory management software for the HP 9800 Series Computers, particularly for Systems 25, 35, and 45 of that family (Models 9825, 9835, and 9845). However, no single system uses all of these features.

We shall assume that all secondary storage devices in the system have the following properties:

1. Each device has the same record size, R , which is an integer multiple of the hard sector size.
2. Each device is treated as a black box, that is, as a vector of records.
3. On all devices, the first record (record #0) is reserved as a system identification record, or system table.
4. A standard set of data is written into the system table to identify the specific secondary storage medium and the drive that handles it.
5. A directory of names is allocated on the medium. The system table has entries that indicate the location and the size of the directory.
6. A backup directory is allocated with the address and size indicated in the system table.
7. A table of available space, or availability table, and a backup (spare) availability table are allocated, and their addresses and sizes are indicated in the system table.
8. A system area may be allocated and indicated in the system table. The system area is subtracted from the space available to the user. It can be addressed only by privileged instructions as an operating system resource.

The system table is highly dependent on the operating system. A possible general system table would contain the following entries:

1. Device ID—to identify the medium and drive and possibly the mainframe.
2. Directory location—to indicate where the name directory's first record starts.
3. Directory size—the number of records in the directory.
4. Spare directory location.
5. Spare directory size.
6. Location of the available space table.
7. Size of the available space table.
8. Spare available space table location.

9. Spare available space table size.
10. Location of the system area.
11. Size of the system area.
12. Location of the user area.
13. Size of the user area.
14. Date of creation of the volume.
15. Purge date of the volume.

Other information can also be included in the system table for added convenience, such as a table of defective records, information on multiple-disc volumes, authorization, and so on. Fig. 1 shows the flexible disc layout used in the HP 9825 Computer system.

The *directory of names* keeps track of the files created in a particular account. The system guarantees that no duplicate names are present in the account. The directory also contains the file location, size, type, password (security code), and any other information pertinent to the particular system and/or application. An entry in the directory can be as simple as the name, location, and size of the file. On the other hand, entries in more sophisticated directories may include not only the type and password, but also information about how often the file has been accessed, and even some accounting information for charging the user for the space. A tree structure can be introduced to the directory to indicate the accounts, groups, users, and user files of a particular system.

The *available space table*, also called the availability table, provides a convenient way to keep track of any unused space on the medium. This table makes it easier and faster to create new files by looking for available space entries and deleting them as they are used up. Some mass storage management systems do not use an availability

table and waste a considerable amount of time trying to locate unused space on a medium each time a new file is created.

The *system area* on a mass storage medium should be optional. If present, it should be indicated in the system table with one entry for its location and one for its size. The system area can store operating-system-related data and programs. The latter can include the operating system code, system library routines, language compilers, bootstraps for other devices in the system, and so on.

User Instructions for Mass Storage Management

Some instructions for mass storage management are absolutely necessary, while others are mere conveniences for making the user's task easier and may be omitted if code space is scarce. We shall describe them here, classified according to their importance.

Essential Mass Storage Management Instructions. The minimal set contains the following instructions: CREATE, PURGE, CATALOGUE, SAVE, GET (LINK), READ, WRITE, and INITIALIZE.

An experienced computer user is usually familiar with most of these instructions; for others, a brief description follows. It is important to note that the functions of these instructions can be varied to suit the system used.

CREATE is used to create (open) a new file. It should include the file's name and size and, optionally, a device identifier if the user wishes to use a particular mass storage drive for that file. Some systems prefer to use OPEN or BUILD for the name of this instruction. Depending on the system's sophistication, a directory entry for a new file can range from simply the file's name, type, size, and address to an entry containing also the creation date, security code, expiration date, and other information.

PURGE, sometimes called KILL or DELETE, will remove a file entry from the directory of the first device in which the file's name is found. If the file is secured, the user must include the security code in addition to the file's name. In most systems, the "quick and dirty" method used is simply to delete the file's name from the directory and consider the file purged. Some systems require that the file's data area be scratched as well, for security or other reasons.

CATALOGUE, also called LISTFILE or LISTNAME, is an important feature of a secondary storage management system. Without it, the files associated with a particular account cannot be listed conveniently. CATALOGUE should be flexible enough to allow the user to list the file names on any preferred list device (CRT, line printer, etc.) with the option to choose subgroups of files in any account according to given keys. For example, it should be possible to catalogue all files with names that start with a certain string of characters or end in a string of characters, or have a common string of characters anywhere in their names, or fit a certain location mask. More options mean more programming needed to implement the catalogue. Other features of the catalogue might include the ability to list in alphabetical order, in order of creation dates, or in any other order that can be established by examining various fields of the directory entry. Fig. 2 shows samples of directory listings using CAT on the HP 9835/9845 Computer systems.




Fig. 1. Disc areas in the 9825 Computer system.

NAME	PRO	TYPE	REC/FILE	BYTES/REC	ADDRESS
------	-----	------	----------	-----------	---------

INIT		DATA	2	256	5
GARBAG		DATA	1	256	7
XDIAG		DATA	5	256	8
XWRITE		DATA	16	256	13
XREAD		DATA	14	256	29
LUNAR		PROG	46	256	43
LANDER		PRPG	44	256	89
XCORR		DATA	8	100	133
AWARD		PROG	7	256	137
REWRIT		ALL	9	256	148
EXPAND		DATA	8	80	153
HEAT01		DATA	7	256	156
HEAT02		KEYS	1	256	171
DUMP		ALL	25	256	177
GARBAG		DATA	1	256	210
BLKJK		DATA	1	256	425

(a)

NAME	PRO	TYPE	REC/FILE	BYTES/REC	ADDRESS
------	-----	------	----------	-----------	---------

F8			65		
BLKJK		DATA	69	256	34/14
SEN		DATA	35	58	33/14
EXPAND		DATA	5	256	31/15
DATA		DATA	4	256	21/01
COPY		DATA	100	256	21/05
REP		DATA	65	58	33/22
GEORGE		DATA	6	256	33/08
XWRITE		DATA	7	256	34/07

(b)

Fig. 2. (a) An example of a directory listing showing single-dimension array addressing. (b) An example of a directory listing showing physical disc addresses.

SAVE is an instruction to place a program (usually user level) into secondary storage, so that it can be recalled later. Other names for this function are KEEP and STORE. Some systems require the creation of a file before using SAVE to store a program in that file. It is preferable to let the system create the file for the user when SAVE is executed on program files, because the system can decide the exact size of the file needed to save the program in memory. In either case, the directory entry for a SAVE file should contain the appropriate file-type information.

GET (LINK) is the opposite of SAVE. It retrieves a program file from mass storage. Other names can also be used. Sometimes a distinction is made between GET and LINK, in that LINK preserves variables declared before the LINK operation is executed, while GET can destroy variables declared before its execution.

INITIALIZE is a multifunction instruction that is both drive and system dependent. The functions of INITIALIZE are as follows:

1. Divide (format) the medium into cylinders, tracks, and sectors (Fig. 3). A cylinder consists of one or more tracks, depending on the number of heads. Formatting consists of

writing soft addresses in the address portions of the sectors. For some drives, it suffices to check the present address for correctness. Otherwise, the address is written on the address portion of every sector of a track.

2. Run test patterns in and out of a track as fast as the system allows to check for bad spots in the medium and for spots that are sensitive to extreme flux transitions in the medium (marginal spots). All zeros, all ones, alternating zeros and ones, and alternating double zeros and double ones are some of the extreme-bit patterns that are often used.

3. Clear all data areas of the acceptable sectors to a known state.

4. For any defective sector, the entire track containing that sector is usually declared defective. This is done in one of the following ways:

- The address portions of all sectors on the track are marked as unused; this is how flexible discs are handled. The controller of such a system is designed so that it understands how to skip a track that has address bits indicating that it is unused. Also, the controller's logic must contain information on how many unused tracks it can allow before it declares the whole medium unusable.
- A spare area is allocated on the medium, and a mechanism is established to link a defective track to a spare track, and back again. The linkage mechanism is partly accomplished through the use of two extra bits in the address field. These are called the "d" and "s" bits. By setting the "d" bit in the address field of a sector and writing the address of the replacement sector in the sector address portion, the controller is instructed to seek for a spare every time it encounters the defective sector. For proper linkage, the spare sector should have its "s" bit set in the sector address portion that contains the address of the defective sector. The "d" bit always has precedence over the "s" bit. A defective track linked to itself represents an error condition.

5. Write the system table into sector zero of the medium and INITIALIZE the directory of names and the table of available space.

6. Optionally, INITIALIZE may write on the disc some system-area routines or other coding. One can also protect these areas in two ways:

- By making the space unavailable to the user, the system area can be used exclusively for systems programs.
- By setting a bit called the "p" bit in the address area of the sector, the sector can be write-protected, unless the user has the capability to clear the "p" bit through a privileged system program.

Optional Mass Storage Management Instructions. There is almost no limit to the content of the mass storage management instruction set. A selected group the author has implemented contains: PURGEALL, RESAVE, STORE, RESTORE, STOREALL, LOAD, LOADALL, ASSIGN#, READ#, WRITE#, REPACK, DIRECTPACK, LISTAVAILABLESPACES, SELECT a particular drive, COLD LOAD, QUICK DUMP.

The following is a brief description of these functions, which may enhance the use of mass storage devices. They can be used to:

1. Make the instruction set more convenient to use; e.g., RESAVE does the functions of executing PURGE and SAVE.
2. Save a routine in binary form rather than in source




Fig. 3. (a) Records and tracks on a flexible disc. (b) Interleaving records on a flexible disc to improve throughput. In addition to an alternating numbering system, the location of the beginning record (record 0) of each track is skewed to avoid a revolution when the drive accesses a new track. For example, after record 29, track 0 is accessed, then record 0, track 1 is accessed without an extra revolution.

form, which for a particular machine configuration may work faster; e.g., one can use STORE instead of SAVE.

3. Dump the entire machine environment, so that a task can be interrupted and later brought back to its state at the time of interruption. STOREALL is such an instruction. This makes it possible to run a multitasking system or to initiate a harmless shutdown if power is interrupted or a hardware failure occurs.

4. Save special program-aid data; e.g., STORE BIN, STORE KEYS, etc.

STORE places the binary image of a program into the memory. It is necessary to have the same system configuration and exact memory size to be able to reload that program.

LOAD gives access to a program previously saved by using STORE. If the program is in the wrong configuration or memory size, it may or may not LOAD, but it will always

fail to run as intended, even if it seems to LOAD.

STORE KEYS is used to save a special-function key (softkey) definition. In HP machines, softkeys are either typing aids, short programs, system commands for a system monitor, or terminal-control commands. By storing the keys' definitions, the user can save typing time when the same program is being executed later.

LOAD KEYS is used to load special-function key definitions from mass storage.

STORE BIN is used to store binary programs, which are usually written in assembly language and which can be called in the high-level language used. Binary programs require special handling to find their assigned locations in the memory map of the system. A STORE BIN instruction is a necessity if binary programs are to be used in a system.

LOAD BIN loads a binary program previously saved, and relocates it if necessary to conform with the memory map of the system's primary storage. LOAD BIN also links the binary program to the operating system if necessary, and updates the memory map to reflect the presence and space of the binary program.

STOREALL dumps the entire contents of primary storage and all registers, flags, semaphores, etc. With STOREALL, a complete representation of the state of the entire system is written into secondary storage.

LOADALL reestablishes the state of the machine at the time that STOREALL was executed. The machine continues its operation into its next state as if no STOREALL had occurred. STOREALL/LOADALL has proved to be a useful combination of complementary instructions when power interruptions can occur, or when multitasking is to be implemented.

A mass storage management system should give an error message if the user tries to store data using a name already assigned. The only exception is when the user wants to update a file and discard the old information. The normal procedure is to PURGE and SAVE, which is a two-step process. In HP systems the RESAVE instruction was introduced to take care of such a situation in one step. Files have to be of the same type for RESAVE to be executed. Also, sufficient space should be available; otherwise a message should indicate that RESAVE did not occur and an explanation should be given. A proper RESAVE should go through the following steps:

- If the types do not match, explain the error.
- If the size of the new file is larger than that of the old one, explain the error.
- Purge the old file.
- Save the new file.

PURGEALL empties all the files on a medium, or all the files in all devices in a system for a particular group. The user of an HP 7925 disc pack in an HP 9845 Computer system will appreciate PURGEALL if it becomes necessary to empty all 19,024 files that are available in one disc pack of that system.

Disc Protocol

To minimize the effort needed to develop drivers for mass storage devices, the disc hardware is designed to respond to a minimal set of commands, as follows:

- SEEK
- READ (full sector)

- WRITE (full sector)
- WRITE WITH VERIFY
- LOAD STATUS
- COLD LOAD READ
- INITIALIZE
- CLEAR
- READ SELF TEST.

This disc protocol is implemented in the recent HP AMIGO interface protocol that makes use of the HP Interface Bus (IEEE Standard 488-1978).¹ Its purpose is to make it possible to develop channel programs of a general nature that can be executed for different devices in the same class. For example, only one disc driver may be needed, regardless of the disc type. When a transaction needs to be made specific for the particular hardware (e.g., INITIALIZE), the hardware is recognized by examining the response of the device to the interfacing command IDENTIFY, which is normally performed during the system power-on sequence.

Conclusion

The investment in system software for mass storage management in data processing systems has been increasing over the years. Adding data base management to the functions of mass storage management can make that software as big and as complex as the entire operating system. For such a monumental task, it becomes increasingly important to find ways to economize; the unified mass storage approach does help in that regard. Other advances in hardware design can make the task of writing system software for mass storage easier. For example, unified disc commands can be used for the various models and sizes of discs. In some cases, a standard mass storage management method used in a processor can be improved for the next generation of the same processor family, thus providing some leverage to reduce the engineering effort. For example, the software for the 9885 double-density flexible disc for the HP 9825 Computer requires 1K words of main system software plus 12K words of bootstraps stored on the disc. This same code was adopted for the 9831 Computer, the bootstraps eliminated, and almost all of the same capability brought into the main system in 5K words of code. For 8K words of code, the same instruction set plus a few extra instructions were incorporated in the 9835/

William A. Hanna



Bill Hanna was born in Cairo, Egypt, and attended Cairo University, graduating with BS degrees in electrical engineering and mathematics in 1967 and 1969, respectively. He received his MS degree in electrical engineering from Iowa State University in 1971 and his PhD degree in electrical engineering from Montana State University in 1974. With HP since 1974, he's developed string manipulation code and mass storage management software for 9800 Series Computers. A member of IEEE and author of several papers and reports on parallel processing, he's

taught algebra, engineering mechanics, computer science and electronics at Cairo University, ISU, and MSU. Bill plays tennis and enjoys classical music, reading, travel, and studying history. He teaches Arabic at Colorado State University, participates in orientation programs on the Middle East, serves as president of the Colorado chapter of the American Coptic Association, and is active in the Coptic Church. Married to a medical doctor, he has two sons and lives in Loveland, Colorado.

9845 Computers. Also, 8K more words were added for low-level drivers of the various disc drives (9885, 7905, 7920, and 7925) because of variations in the way these drives are controlled. The low-level drivers could have been cut at least 50% if the command set had been exactly the same for all of these disc drives.

Acknowledgments

I would like to thank Jan Kok for implementing some of the 9825 code, Dave Delano for doing the same on the 9835, Gene Burmeister for his work on the 9845, and the management of the Desktop Computer Division for making it all possible.

References

1. W.G. Matheson, "A Computer Input/Output System Based on the HP Interface Bus," Hewlett-Packard Journal, July 1979.

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

HEWLETT-PACKARD JOURNAL

JUNE 1980 Volume 31 • Number 6

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard Company, 1501 Page Mill Road
Palo Alto, California 94304 U.S.A.
Hewlett-Packard Central Mailing Department
Van Heuven Goedhartlaan 121
1180 AM Amstelveen The Netherlands
Yokogawa-Hewlett-Packard Ltd., Suginami-Ku
Tokyo 168 Japan

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

020002081066&BLAC&CA00
MR C A BLACKBURN
JOHN HOPKINS UNIV
APPLIED PHYSICS LAB
JOHNS HOPKINS RD
LAUREL MD 20810

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