


HEWLETT-PACKARD JOURNAL

JUNE 1988



HEWLETT-PACKARD JOURNAL

June 1988 Volume 39 • Number 3

Articles

6 **Statistical Issues in Setting Product Specifications**, by *Sherry L. Read and Timothy R.C. Read*

7 **Robust Estimators**

9 **Propagation of Error with Multiple Sources of Variability**

11 **Boxplots**

12 **Circuit Design Using Statistical Data Analysis**, by *Karen Kafadar and Lynn M. Plouse*

18 **Statistical Calibration of a Vector Demodulator**, by *Karen Kafadar*

24 **Appendix: Solving Nonlinear Least Squares Problems**

26 **An Availability and Reliability Tool for Computer Systems**, by *Wulf D. Rehder*

28 **The Language of Dependability**

30 **Project Management Using Software Reliability Growth Models**, by *Gregory A. Kruger*

36 **A Reliable, Autoloading, Streaming Half-Inch Tape Drive**, by *John W. Dong, Kraig A. Proehl, Ronald L. Abramson, Leslie G. Christie, Jr., and Douglas R. Domel*

43 **Streaming Tape Drive Control Electronics**, by *Bradfred W. Culp, Douglas R. Domel, Wayne T. Gregory, Jeffery J. Kato, Gerod C. Melton, Kraig A. Proehl, David W. Ruska, Virgil K. Russon, and Peter Way*

58 **A One-Gigasample-per-Second Digitizing Oscilloscope**, by *Joe K. Millard*

Editor, Richard P. Dolan • Associate Editor, Charles L. Leath • Assistant Editor, Hans A. Toepfer • Art Director, Photographer, Arvid A. Danielson
Support Supervisor, Susan E. Wright • Administrative Services, Typography, Anne S. LoPresti • European Production Supervisor, Michael Zandwijken

59 **A One-Gigasample-per-Second Analog-to-Digital Converter**, by *John J. Corcoran, Ken Poulton, and Knud L. Knudsen*

60 Repetitive versus Single-Shot Bandwidth

64 Digitizer Hybrid

67 **Front-End Signal Conditioning for a High-Speed Digitizing Oscilloscope**, by *Thomas K. Bohley, Mark E. Mathews, Lewis R. Dove, Joe K. Millard, David W. Bigelow, and Donald D. Skarke*

70 **Digital Filtering in a High-Speed Digitizing Oscilloscope**, by *B. Allen Montijo*

72 Dithering in the HP 54111D

75 Digital Filters

Departments

4 In this Issue

5 What's Ahead

55 Authors

The **Hewlett-Packard Journal** is published bimonthly by the Hewlett-Packard Company to recognize technical contributions made by Hewlett-Packard (HP) personnel. While the information found in this publication is believed to be accurate, the Hewlett-Packard Company makes no warranties, express or implied, as to the accuracy or reliability of such information. The Hewlett-Packard Company disclaims all warranties of merchantability and fitness for a particular purpose and all obligations and liabilities for damages, including but not limited to indirect, special, or consequential damages, attorney's and expert's fees, and court costs, arising out of or in connection with this publication.


Subscriptions: The Hewlett-Packard Journal is distributed free of charge to HP research, design, and manufacturing engineering personnel, as well as to qualified non-HP individuals, libraries, and educational institutions. Please address subscription or change of address requests on printed letterhead (or include a business card) to the HP address on the back cover that is closest to you. When submitting a change of address, please include your zip or postal code and a copy of your old label.

Submissions: Although articles in the Hewlett-Packard Journal are primarily authored by HP employees, articles from non-HP authors dealing with HP-related research or solutions to technical problems made possible by using HP equipment are also considered for publication. Please contact the Editor before submitting such articles. Also, the Hewlett-Packard Journal encourages technical discussions of the topics presenting in recent articles and may publish letters expected to be of interest to readers. Letters should be brief, and are subject to editing by HP.

Copyright © 1988 Hewlett-Packard Company. All rights reserved. Permission to copy without fee all or part of this publication is hereby granted provided that 1) the copies are not made, used, displayed, or distributed for commercial advantage; 2) the Hewlett-Packard Company copyright notice and the title of the publication and date appear on the copies; and 3) a notice stating that the copying is by permission of the Hewlett-Packard Company appears on the copies. Otherwise, no portion of this publication may be produced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage retrieval system without written permission of the Hewlett-Packard Company.

Please address inquiries, submissions, and requests to: Editor, Hewlett-Packard Journal, 3200 Hillview Avenue, Palo Alto, CA 94304, U.S.A.

In this Issue



The voice on the telephone one day last year was HP statistician Tim Read's. Pointing out that not all technical contributions are made by engineers, he made us an offer we couldn't refuse—to invite a group of HP statisticians to submit papers on applications of statistics. The results are five excellent papers that cover an impressive range of topics: designing electronic hardware, calibrating instruments, setting specifications, estimating reliability, and managing software projects. In "Statistical Issues in Setting Product Specifications," page 6, we learn how the science of statistics helps ensure that specifications are useful for describing the product's performance, relevant for competitive comparisons, measurable and verifiable, and valid over the required environmental ranges and time intervals. This paper takes a tutorial approach to its subject. "Circuit Design Using Statistical Data Analysis," page 12, tells how measurements on prototype amplifiers have been analyzed statistically to guide the design of adjustment-free amplifier boards. "Statistical Calibration of a Vector Demodulator," page 18, describes the design of the statistical demodulator calibration algorithm that is programmed into the HP 8981A Vector Modulation Analyzer. The algorithm can be used to calibrate not only the built-in demodulator but also external demodulators. "An Availability and Reliability Tool for Computer Systems," page 26, describes a software program that computes system availability based on reliability data for the system components. The program helps R&D engineers make trade-offs in designing for reliability. "Project Management Using Software Reliability Growth Models," page 30, relates one HP Division's experience with a model-based method for predicting the duration of software system testing. The model also helps management decide when to release a product.

Half-inch reel-to-reel tape recorders are widely used for archival storage of computer data, backing up data stored on disc memory, and interchanging data between systems. The HP 7980A Tape Drive, besides being faster and more reliable than its predecessors, is interesting because it's an autoloading drive. Designed to mount in a rack with the reels and tape path in a horizontal plane rather than the more conventional vertical plane, it takes up much less rack space, but its reels and tape path aren't readily accessible for the user to mount reels and thread tape. However, no problem. All the user need do is open the door and slip in a reel of tape. The HP 7980A automatically threads the tape onto the takeup reel. How does it do that? Its designers tell how in the article on page 36. The drive's state-of-the-art control electronics (page 43) provide features such as variable-velocity rewind, which chops 25% off the 120 seconds it used to take to rewind a 2400-foot reel of tape.

In the February 1988 issue, we presented the design story of a family of digitizing oscilloscopes based on waveform recorder technology. The primary objective in these instruments' design was maximum fidelity within a given bandwidth. Featured this month is a digitizing oscilloscope designed for maximum speed, which means maximum bandwidth for capturing single-shot phenomena. The key element in the HP 54111D Oscilloscope is the advanced one-gigasample-per-second analog-to-digital converter (ADC) shown on the cover. Thanks to this ADC's speed, the HP 54111D can capture and display signals containing frequencies up to 250 megahertz with a single look. This wide single-shot bandwidth is needed for digital system diagnostics, high-energy physics, transient analysis, and other applications. After a brief introduction (page 58) by Joe Millard, who was project manager for the HP 54111D's development, three authors from HP Laboratories describe the fast-ADC technology (page 59). Four digitizers are interleaved in this design, which uses gallium arsenide, silicon bipolar, and silicon NMOS integrated circuits and thick-film hybrid circuit technology. The wideband hybrid preamplifier and attenuator circuits at the instrument's input are described in the article on page 67. Filtering, bandwidth, noise, and waveform reconstruction issues are discussed in the paper on page 70, which shows how the six-bit ADC can provide seven or eight bits of information for sufficiently low-frequency input signals.

-R.P. Dolan

What's Ahead

In the August issue, eight articles will discuss the design of the HP PaintJet Color Graphics Printer, with particular emphasis on the design and manufacturing of the print cartridge. Aluminum gallium arsenide (AlGaAs) red light-emitting diode lamps are the subject of another article. HP-RL, the Hewlett-Packard Representation Language, was used for expert systems research in HP Laboratories until recently. We'll have a paper summarizing its contributions and another paper on MicroScope, an expert system for program analysis, which was originally written using HP-RL.

Statistical Issues in Setting Product Specifications

A primer on the use of statistics in specification setting.

by Sherry L. Read and Timothy R.C. Read

PRODUCT SPECIFICATIONS are frequently the first chance that a potential customer has to evaluate a new product. The purpose of such specifications is to characterize the product's performance and to warrant a minimum performance level to the customer. Consequently, it is extremely important that the specifications accurately reflect the quality and performance of the product. Conservative specifications that downplay the true performance are likely to benefit competitors selling similar products who provide a more realistic assessment of the specifications. On the other hand, specifications that are too ambitious for the product are likely to cause high manufacturing cost and possible customer dissatisfaction.

The choice of product specifications is driven by customer needs, product design, competition, and manufacturing capabilities. Final setting of specifications should account for all these aspects, as well as for statistical evidence of product performance from a production pilot run. In particular, the process of setting specifications needs to reflect the accuracy and repeatability of a single production unit and the variation between different units.

Examples of typical product performance questions are:

- With what level of certainty will the harmonic distortion of this signal generator lie within a certain limit?
- With what confidence will the noise power of a single 100-MHz crystal oscillator vary by less than 10 percent over a two-week period?
- What is the probability that a product meeting test criteria at room temperature will meet or exceed specifications over all warranted environmental conditions?

In this paper we illustrate how statistical methods of data analysis can help product designers, manufacturing and marketing engineers provide efficient and accurate answers to such questions.

Objectives

What are the objectives for the process of setting specifications?

Specifications should be:

- Useful in describing product performance to customers and field offices
- Relevant for competitive comparisons
- Measurable and easily verifiable in production, at service centers, and by customers
- Valid over the appropriate environmental conditions and time intervals.

To meet these objectives and provide maximum information to the customer, product specifications are frequently presented in a two-tier structure.* The top tier describes

warranted specifications and the second tier describes supplemental characteristics. The supplemental characteristics provide information useful in product applications by giving typical, but nonwarranted, performance parameters for the laboratory environment. In this paper, we shall concentrate on the setting of warranted specifications. However, the same general methods can be applied to supplemental characteristics.

Statistics, the science of drawing conclusions and making decisions in the face of uncertainty, aids in achieving many of the objectives for specifications. Based on production pilot units, statistical methods are used to estimate the expected performance of future production units. Environmental effects can be modeled so that not all units have to be tested over all environments. Specific graphical tools provide a clear picture of performance to facilitate the decision process by marketing and manufacturing when setting the final specifications. Statistics is the key to understanding and efficiently quantifying all sources of variation.

Statistical Model

The underlying model for the specifications of many Hewlett-Packard products is illustrated in Fig. 1. In this diagram, production margin is the difference between the average product performance and the test line limit (TLL), which is the pass/fail limit used by the production line at final test under standard environmental conditions. Delta environmental (DE) represents the possible change in performance of the product over the environmental extremes specified on the data sheet (e.g., temperature and humidity), and drift represents the change in performance over the calibration period or other specified time interval. Measurement uncertainty (MU) accounts for possible measurement errors in the equipment used to characterize the product. The customer guardband represents any additional

*This structure was originally proposed to the National Conference of Standards Laboratories (U.S.A.) by John Minck in 1977.




Fig. 1. Representation of the components that make up a warranted product specification. TLL is the test line limit used by the production line at final test.

guardband considered necessary to ensure that in a worst-case scenario, any customer's unit will perform better than the published specification (this additional guardband is sometimes set to zero). The right end of Fig. 1 (SPEC) represents the warranted customer specification as described by the product data sheet.

When finalizing product specifications, it is necessary to look at Fig. 1 from two directions: from the SPEC inward and from the observed pilot data outward. The former direction involves setting internal criteria (the TLLs) to guarantee that the product performance meets specifications over all warranted conditions. The latter direction involves setting specifications that will allow a high yield in manufacturing and simultaneously minimize the chance that good units will be rejected and maximize our confidence that the units passed will perform as specified. We shall focus on the statistical aspects of moving from the production pilot units outward to the preliminary specifications. However, it is important to realize that the final specifications are set by iterating between both approaches to ensure that the final product has a high yield and satisfies the market needs.

Using our experience of specific product introductions, we will describe the relevant statistical methods needed to complete the estimation of each of the quantities in Fig. 1 based on production pilot data. The same approach is also used to adjust specifications on more mature products.

Estimating the Test Line Limit

Based on the assumption that the production pilot units are representative of the product characteristics that can be expected in production, the goal is to use these pilot units to calculate the TLL so that the vast majority of production units will be better than this limit. In particular, we need to ensure that the products will be readily manufacturable over the long term at a low cost.

If we use the convenient assumption that the production pilots come from a Gaussian (normal) distribution with a population mean of μ and standard deviation of σ , then the interval $(\mu - 2\sigma, \mu + 2\sigma)$ contains 95% of the distribution and $(\mu - 3\sigma, \mu + 3\sigma)$ contains more than 99% of the distribution. This normality assumption allows us to model the expected production pilot distribution and the corresponding percentiles using only the two parameters μ and σ .

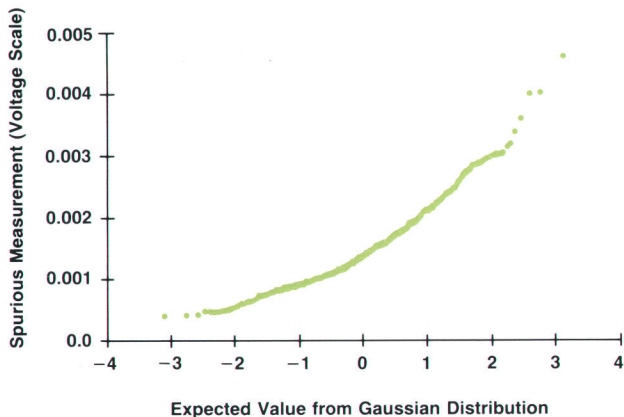


Fig. 2. Quantile-quantile plot of spurious distortion data on a voltage scale.

Robust Estimators

The sample mean and sample standard deviation are the best estimators of the location and scale of a Gaussian or normal distribution. However, a few large errors may occur in the data, resulting in a distribution with heavier tails than the Gaussian. In this case, estimators that place less emphasis on extreme values may provide more appropriate estimates for the mean and standard deviation.

One such set of estimators consists of the trimmed mean and the Winsorized standard deviation. To calculate these estimates, trim off g observations from either end of the set of n ordered observations $x_1 < x_2 < \dots < x_n$, where g is the closest integer to $(0.1)n$. Then the 10% trimmed mean is given by:

$$\bar{x}_t = (x_{g+1} + x_{g+2} + \dots + x_{n-g}) / (n-2g),$$

and the Winsorized standard deviation is calculated as:

$$s_t = \{[(g+1)(x_{g+1} - \bar{x}_t)^2 + (x_{g+2} - \bar{x}_t)^2 + \dots + (x_{n-g-1} - \bar{x}_t)^2 + (g+1)(x_{n-g} - \bar{x}_t)^2] / (n-2g-1)\}^{1/2}.$$

The Winsorized standard deviation formula treats the large (or small) observations that were trimmed as if they were equal to the largest (or smallest) observations in the trimmed sample.

When the trimmed mean and Winsorized standard deviation are used, the effective sample size is the size of the trimmed sample. Hence, when finding the K factors for use in $\bar{x}_t + Ks_t$, enter the tables with $n = n - 2g$.

Other robust estimators, such as the biweight, are described in References 1, 2, and 3 listed on page 11.

While there are good reasons to assume normality, it is important to check the validity of this assumption before proceeding to estimate the percentiles of the production pilot distribution. Histograms are frequently used to verify the symmetry of a distribution. However, the investigation of normality is much clearer when the data is plotted on a Gaussian quantile-quantile plot as illustrated in Figs. 2 and 3.

Here the data values are plotted against the values that would be expected from a Gaussian distribution. If the data




Fig. 3. Quantile-quantile plot of spurious distortion data on a dBc scale.

is Gaussian, the values should lie approximately on a straight line. Such plots as these help to check assumptions, to determine the appropriate scale in which to set specifications, and to detect measurements or units that appear not to conform to the bulk of the data. To illustrate the appropriate choice of scale, Figs. 2 and 3 represent data on spurious distortion for an HP Stanford Park Division product on both the linear (voltage) and the log (dBc) scale, respectively. Fig. 3 shows that the measurements on the log scale are more nearly Gaussian than those on the linear scale in Fig. 2. Consequently, the calculations for determining the TLL of this product should be done in dBc.

A different situation occurs in Fig. 4, which represents a quantile-quantile plot of intermodulation distortion from some pilot-run models of the same product. While the bulk of the data appears to be nearly Gaussian, one instrument (circled) appears to be different, exhibiting extremely good performance. This outlying instrument should be accounted for in the estimation of the mean and standard deviation used to calculate the TLL. The following section outlines the methodology for dealing with such points.

Estimating μ and σ

Consider now replacing the mean μ and the standard deviation σ with their standard estimates

$$\bar{x} = (x_1 + x_2 + \dots + x_n)/n$$

and

$$s = \{[(x_1 - \bar{x})^2 + (x_2 - \bar{x})^2 + \dots + (x_n - \bar{x})^2]/(n - 1)\}^{1/2},$$

respectively, where x_i represents the measurement value from pilot unit i ($i = 1, 2, \dots, n$) and n is the total sample size. We are faced with two potential problems. First, consider the situation where one or two extreme points are found in the sample, such as in Fig. 4 for the intermodulation data. If these points occur only at the opposite extreme from where the TLL is to be set, then under certain circumstances we may wish to use estimates of μ and σ that are less sensitive to (i.e., more robust against) individual extreme departures from a Gaussian distribution. The standard estimates \bar{x} and s given above are notoriously non-robust. An example of how one extreme point can throw off \bar{x} and s is illustrated in Fig. 5 below. Here both sets of

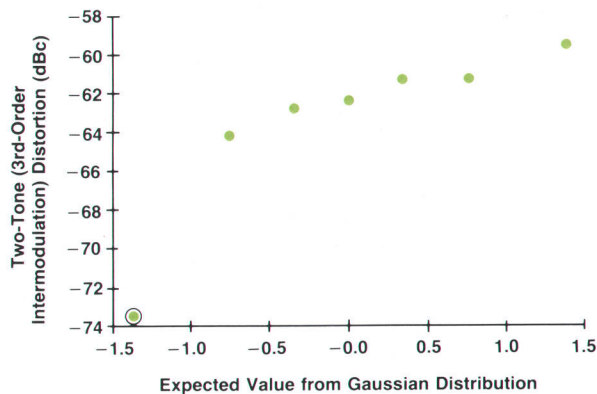


Fig. 4. Quantile-quantile plot of two-tone intermodulation distortion data, showing an extreme point.




Fig. 5. Both sets of observations have the same values for \bar{x} and s , but the values for (b) are inflated because of the extreme point at 15.

observations have the same values for \bar{x} and s . However, clearly, \bar{x} and s are inflated for the set of observations in Fig. 5(b) because of the one extreme point at 15.

There are a number of procedures for robust estimation.^{1,2,3} For example, the standard estimates \bar{x} and s for the intermodulation data in Fig. 4 are -63.53 dBc and 4.59 dBc, respectively, whereas the more robust 10% trimmed mean and Winsorized standard deviation (see "Robust Estimators," page 7) for this data are -62.36 dBc and 1.64 dBc, respectively. When such methods are used, the engineer must be certain that the extreme points are well understood, and could not equally well have occurred on the other extreme of the distribution. Unless the mechanisms causing such extreme observations are understood, robust estimators should be used with extreme caution in this context.

The second potential problem in replacing μ and σ by their estimates is that probability statements such as "the interval $(\mu - 2\sigma, \mu + 2\sigma)$ contains 95% of the distribution" are no longer true in general. The amount of the underlying distribution covered by the interval $(\bar{x} - 2s, \bar{x} + 2s)$ or $(\bar{x} - 3s, \bar{x} + 3s)$ depends on how accurately \bar{x} and s estimate μ and σ , which in turn depends on the sample size n . For large samples, the statement "the interval $(\bar{x} - 2s, \bar{x} + 2s)$ contains 95% of the distribution" will be approximately true (as a result of the Central Limit Theorem). But how large a sample size is needed to achieve reasonable accuracy?

This question is answered by reference to suitable tolerance limit tables.^{4,5} Given the sample size, these tables provide a level of confidence that the interval does indeed cover the requested percentage of the distribution. For example, based on a sample of 10 units, if we want an interval that covers 95% of the distribution with 95% confidence, we would have to use $(\bar{x} - 3.4s, \bar{x} + 3.4s)$ rather than $(\bar{x} - 2s, \bar{x} + 2s)$. This tolerance interval procedure accounts for the random variations in the estimates of both μ and σ , and provides an interval that will include at least a specified percentage of the population with a high degree of confidence.

In conclusion, our procedure for setting the test line limits involves three steps. First, verify the distributional assumption of a sample from a Gaussian distribution. Second, estimate the mean and standard deviation of the per-

Propagation of Error with Multiple Sources of Variability

For a single source of variability, the standard deviation σ can be estimated by:

$$s = \{[(x_1 - \bar{x})^2 + (x_2 - \bar{x})^2 + \dots + (x_n - \bar{x})^2] / (n - 1)\}^{1/2}$$

where the x_i ($i = 1, 2, \dots, n$) represent measurements on n units.

However, it frequently happens that the test characteristic is measured as a known function of other independent measured variables. For example, suppose that y_1 , y_2 , and y_3 are the heights of three types of discs. If one stacks the discs on top of each other, then the height of the stack is $h = y_1 + y_2 + y_3$. Given tolerance limits on the individual discs, what are the tolerance limits on the overall stack height? Define s_1 , s_2 , and s_3 to be the standard deviations of y_1 , y_2 , and y_3 , based on n_1 , n_2 , and n_3 data points, respectively. Then the law of propagation of errors states that the standard deviation of h is given by

$$s_h = \sqrt{s_1^2 + s_2^2 + s_3^2}$$

Note that this formula requires only that the measured variables y_1 , y_2 , and y_3 are independent of one another. Further distributional assumptions are not necessary.

Tolerance intervals can now be constructed for h according to the methods discussed in this paper, using an effective sample size of $n = n_1 + n_2 + n_3 - (3 - 1)$.

formance using a standard or robust procedure. Finally, using the K factors for tolerance intervals (based on the production pilot run sample size), set TLLs at $\bar{x} + Ks$ or $\bar{x} - Ks$ as appropriate. These limits will provide the first step in setting customer specifications and provide knowledge, with high confidence, of what can be expected from future manufacturing, assuming a stable manufacturing process.

Estimating Delta Environmental

To guarantee that the product will meet or exceed its warranted performance specifications over all environmental conditions, it is necessary to estimate the expected change in product performance from the standard test conditions to the worst-case environment. In most cases, this can be estimated as the average measurement in the worst-case environment minus the average measurement in the standard (or test) environment, that is,

$$DE = \bar{x}_{env} - \bar{x}_{std}$$

Some products will experience not only a mean shift over an environmental range, but also an increase in the variability of their performance. This increase in variability can be accounted for by the following equation:

$$DE = \bar{x}_{env} - \bar{x}_{std} + K(s_{env} - s_{std})$$

Fig. 6 illustrates both situations by depicting the distribution of a population of units in the standard operating

environment and in the worst-case environment. If $s_{env} = s_{std}$, as in Fig. 6a, then DE will be equal to $\bar{x}_{env} - \bar{x}_{std}$. If s_{env} is greater than s_{std} (Fig. 6b), then DE will be equal to the difference in some percentile. By accounting for DE, customers can be assured that the units tested to the TLL in the standard environment will meet or exceed specifications over the warranted conditions.

Estimating Measurement Uncertainty

Measurement uncertainty accounts for possible changes in measurements caused by the inherent variation of the test or measurement equipment. The most common method of estimating this variation is to combine the specifications provided on the test equipment data sheets for each piece of equipment integrated into the test system. For example, if there are four pieces of equipment making up the measurement system, with individual specifications $\pm r_1$, $\pm r_2$, $\pm r_3$, $\pm r_4$, then typically we can estimate the respective standard deviations s_1 , s_2 , s_3 , and s_4 as one half or one third of these values (if they are given as 95% or 99% limits, respectively). Using the root-sum-square (RSS) principle (see "Propagation of Error with Multiple Sources of Variability," this page), the measurement uncertainty for the system is calculated as

$$MU = \sqrt{(s_1^2 + s_2^2 + s_3^2 + s_4^2)}$$

Setting Preliminary Specifications

To account for all the sources of variation, the preliminary specifications relative to manufacturing capabilities would be defined as

$$SPEC = \bar{x} + \sqrt{(Ks)^2 + (3MU)^2} + DE,$$

or

$$SPEC = \bar{x} - \sqrt{(Ks)^2 + (3MU)^2} + DE,$$




Fig. 6. Estimation of delta environmental (a) when the mean is a function of the environment but the standard deviation is constant, and (b) when both the mean and the standard deviation are functions of the environment.

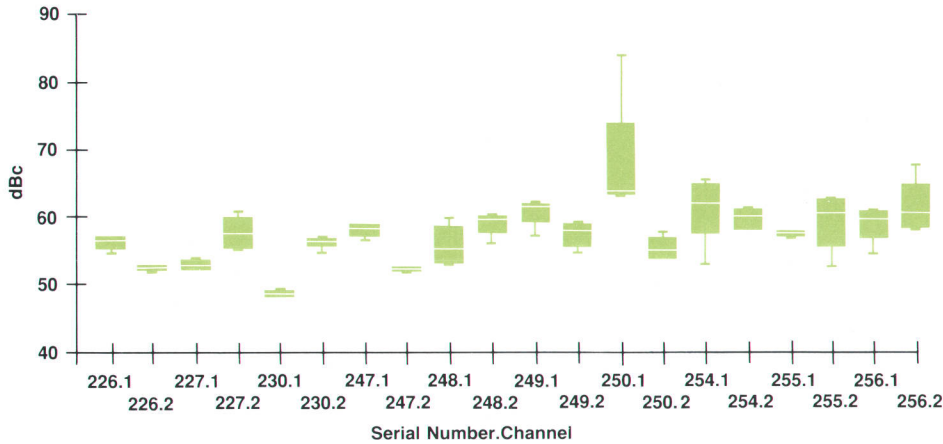


Fig. 7. Boxplots of common-mode rejection data for different instruments and channel numbers over all environmental conditions.

as appropriate. Again we use the RSS principle to combine the unit variation (s^2) and the measurement uncertainty (MU^2). Delta environmental is treated as a bias. This is the preliminary SPEC based on evidence of manufacturing capabilities, but these manufacturing capabilities do not capture all the considerations in choosing the final specifications. Much of the final decision must be based on customer, competition, and business considerations.

Graphical Tools

To facilitate discussion and understanding of the preliminary specifications, graphical displays of the pilot data with TLLs and preliminary specifications are useful. In particular, boxplots (see "Boxplots," page 11) are important tools for displaying the magnitude and dispersion of a number of sets of data. For the HP 5183A Waveform Recorder, the dispersion of measurements for each pilot run instrument over the various environmental conditions is shown in Fig. 7, and the aggregate pilot run data for each of the tested environments is shown in Fig. 8. Fig. 7 allowed the group responsible for setting specifications to identify the particular channels that performed exceptionally well (e.g., #250.1) or showed little variation over temperature and line voltage (e.g., #247.2). Fig. 8 indicates that there is little difference between environments. These graphical tools help identify cases that are different from the rest, a necessary step in improving the consistency of products and the manufacturing process.

Boxplots can be used to evaluate specifications over a range of measurements. For example, Fig. 9 is a plot of the

pilot run results of the same measurement on the various HP 5183A Waveform Recorder ranges. This kind of graphical display helps determine the ranges where the specifications should change. If there are any ranges that are particularly variable, they can be pinpointed for more intensive testing, providing a basis upon which to optimize the efficiency of the internal testing process.

Once preliminary specifications are established, many of the specifications will be as expected and will require little if any discussion, while other specifications may require some specific action. A systematic approach to calculating the preliminary specifications is important. This includes validation of assumptions, a clear understanding of the sources of variation, and graphical displays of all the pertinent information. This approach will do much to ensure that the decisions about final specifications result in a valid representation of a product's performance upon which customers can depend.




Fig. 8. Boxplots of common-mode rejection data for different environments for all instruments.

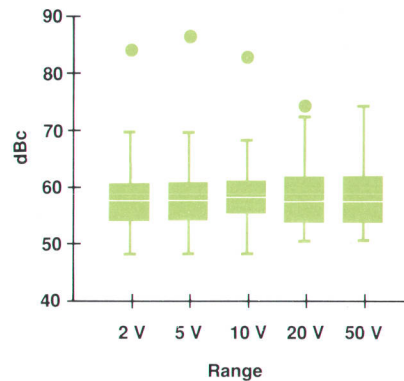


Fig. 9. Boxplots of common-mode rejection data for different ranges for all instruments and environments.

Boxplots

Boxplots are a graphical tool displaying the salient features of the distribution of a set of data—namely, the location, the spread, any asymmetry, and outlying data points.¹ Because of their compactness, boxplots are especially useful for comparing several sets of data on one graph.

Fig. 1 shows a histogram and a boxplot of a single set of data. In this boxplot, the maximum value is not joined to the box by the "whisker" (i.e., the line extending from the box), and is identified separately as an unusual value. If we define the difference between the 75th percentile and the 25th percentile as the interquartile range (IQR), then unusual values are those values that are more than 1.5 IQR away from the appropriate percentile. In the case of an unusual value, the whisker will end at the largest (smallest) data value that is within 1.5 IQR of the 75th (25th) percentile.

Reference

1. P.F. Velleman and D.C. Hoaglin, *Applications, Basics, and Computing of Exploratory Data Analysis*, Duxbury Press, 1981, chapter 3.


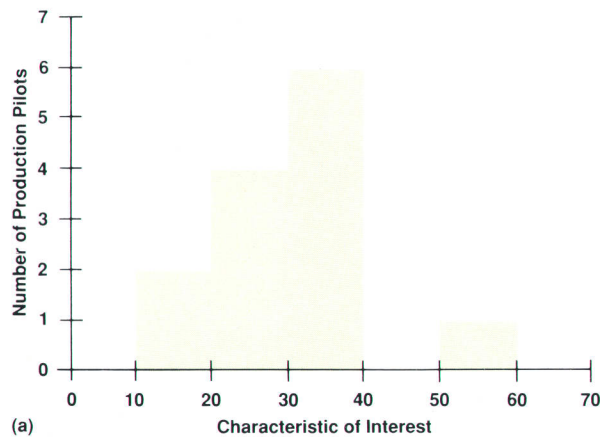


Fig. 1. (a) Histogram and (b) boxplot for the same data.

References

1. F. Mosteller and J.W. Tukey, *Data Analysis and Regression: A Second Course in Statistics*, Addison-Wesley, 1977, chap. 10.
2. D.C. Hoaglin, F. Mosteller, and J.W. Tukey, *Understanding Robust Statistics and Exploratory Data Analysis*, John Wiley & Sons, 1983.
3. F.R. Hampel, E.M. Ronchetti, P.J. Rousseeuw, and W.A. Stahel, *Robust Statistics*, John Wiley & Sons, 1986.
4. M.G. Natrella, *Experimental Statistics*, U.S. National Bureau of Standards Handbook No. 91, Reprinted by John Wiley & Sons, 1985.
5. I. Miller and J.E. Freund, *Probability and Statistics for Engineers*, Prentice-Hall, 1985.

Circuit Design Using Statistical Data Analysis

A methodology for setting limits of uncertainty on amplifier output power based on measured prototype data facilitates the design of amplifiers that minimize output power variations. The result is a printed circuit board requiring no adjustments.

by Karen Kafadar and Lynn M. Plouse

A COMMON SITUATION facing the circuit designer is the selection of design parameters for certain circuits so that specified targets and variations in output power levels are met. This is the case for several circuit boards that are part of the HP 86792A Agile Upconverter, a fast-switching synthesizer that can upconvert complex modulation with up to 32-MHz bandwidth to anywhere in the 10-to-3000-MHz range. These boards contain several circuits such as amplifiers, pads, frequency dividers, power splitters, switches, and filters. While some circuits, such as power splitters and dividers, involve few design choices, other circuits, such as amplifiers and pads, involve parameters that can be chosen in a variety of ways to meet target output levels. The aim is to select values for those parameters so that variation of the output power is minimized. If the overall variation in the output power is small enough, no adjustments are necessary, thereby reducing cost and labor requirements in production.

Statistical methodology has been applied to this project in several ways. Measurements on the different types of circuits provide information on output levels and expected variation. Such data, collected in the prototype phase, is used to select design parameters for amplifiers so variation of the output power can be minimized. Variations caused by other sources, such as temperature and input power variations, are assessed. Finally, the results from the lab pilot phase are compared with those from the prototype phase for prediction of performance in production.

This paper discusses the application of statistical methodology to seven printed circuit boards used in the prototype design phase. This project illustrates the power of statistics and the value of a statistical analysis of data collected early in the design stage. The success of this approach is verified by the close prediction of lab pilot units from the prototype, and, more important, by the labor savings from designing a printed circuit board requiring no adjustments.

Circuit Layout

The new synthesizer upconverts a modulated input signal. Several fast switches allow fast hopping of the carrier. Since each switch drives a mixer, its output level is critical. If the output power is too high, the mixer's spurious outputs (spurs) increase, degrading the overall spur performance

of the instrument. On the other hand, if the switch output power is too low, the signal-to-noise ratio gets worse. Optimally, the output levels of the switches should be within a 3-dB window, regardless of output frequency, temperature, or components used to assemble the circuit.

Often the target output levels on circuit boards of this type are achieved by adjustments for each pole of the switch. Such a procedure requires at least 16 adjustments for these particular boards, depending upon how they were implemented. At a minimum rate of \$12 per adjustment, this is highly costly in production. A better way uses statistical analysis to determine the loss and variation of the circuits in the switch modules. In particular, some of these circuits (e.g., the amplifiers that drive the circuits) can be designed to minimize output variation caused by variation of the incoming power.

Sources of Variation

Having defined the goals of the circuitry, data on various circuits must be collected and analyzed to help realize these goals. Depending upon availability of components and required test time, different circuits of each type are measured. This data provides estimates of output power levels and variations as a function of input power. Several sources of variation in these measurements can be identified:

- Variations in incoming power levels
- Errors in input power level transmitted by the circuit
- Temperature biases on the output
- Circuit component variation
- Variation in the manufacturing of the circuit board material
- Measurement error.

In this study, the incoming power is actually the output power from an earlier portion of the instrument, which itself has a prespecified variation by design. The transmis-




Fig. 1. A typical power path on one board at a single frequency (470 MHz).

sion of the error in this incoming power depends upon the circuit (see next section).

All circuits are required to perform over a temperature range to pass environmental testing; hence the offsets in power levels because of extreme heat (70°C) and cold (5°C) are estimated (see "Temperature Testing," page 14).

Having assembled all circuits on printed circuit boards, variations in lots of board material or components may affect output power; again, such variation is best estimated by taking measurements on different boards (see "Component Variation," page 15).

Finally, this analysis involves two components of variation that come from measurement error. The first arises directly from the test system and is assessed most typically from customer specifications. The second component involves an error imposed by the analysis, because input power levels are not identical for measurements of circuits taken on different occasions (different temperatures, components, etc.). To make fair comparisons of the output levels, most data had to be linearly interpolated to a common input power scale. This common scale was an interval within the span of the data confined to the integers (e.g., -10 dBm, -9 dBm, ..., 10 dBm). The maximum error in linear interpolation is given by one eighth of the second derivative of the function and estimated by one fourth of the second divided difference. Only rarely did this maximum error exceed 0.02 dB. Generally, this error was a negligible component of the overall uncertainty in the power level. However, in other less precise measuring situations, this error might not be negligible.

The variations from these different factors are estimated and combined into an overall standard deviation by a root sum of squares (RSS) method. This method is almost always justified by the fact that the variance (square of the standard deviation) of a sum of independent random quantities equals the sum of the variances, regardless of how the individual random quantities are distributed. In some cases where two factors were measured simultaneously (e.g., temperature and boards), the interaction between them was negligible and was therefore ignored.

For simplicity, and in accordance with recommendations by the Bureau International des Poids et Mesures,¹ all variations are given in terms of plus or minus one standard deviation (SD or σ), except where noted. When a Gaussian (utopian) model for the data applies, 2.576 multiples of the population SD should include 99% of the population. When only ten units have been measured and the Gaussian model is assumed, 3.018 multiples of the estimated SD includes 95% of the population with 90% confidence.² A rough translation of a customer specification of a test system into a 1σ uncertainty is mostly guesswork; often one half to one third of the customer specification is treated as one standard deviation. Customer specification of the power meters in this study is 0.2 dB; hence $\sigma_{ME} \approx 0.07$ dB. (Throughout this paper, σ will refer to an estimate of the variation from a specific source, which will be identified by the subscript on σ .)

Transmission of Error

In this section, we illustrate the statistical methodology by considering one path from one board corresponding to

one frequency input and output; this path is shown in Fig. 1. Power is transferred along this path through a series of circuits that includes a power splitter, an amplifier, a switch, and some pads to absorb excess power. The input power level from the preceding circuitry for this path is estimated at 3.91 dBm ± 0.94 dB. The power splitter incurs a loss of 3.8 dBm ± 0.16 dB. Hence, using RSS, the power preceding the amplifier is at most 0.11 dBm ± 0.95 dB. Increasing the loss of the amplifier input pad will lower the output power but not change the variation of the transmitted power. The goal is to design the amplifier so that the target output power after the switch is -6.5 dBm, with a 99% uncertainty interval of ± 1.5 dB.

Knowledge of the transmission of error variation through a circuit provides the greatest leverage in reducing overall variation in the output power of the switch. Consider a typical response curve from a switch in Fig. 2. It is clear that the output variation transmitted around P_0 is larger than that around P_1 . To a first approximation, this output variation depends upon the slope of the curve at the input. Taylor's theorem provides the basis for this "law of propagation of error":³

$$\begin{aligned} \text{output} &= f(P) \approx f(P_0) + (P - P_0)f'(P_0) \\ &= \text{constant} + (P - \text{constant}) \times \text{constant}. \end{aligned}$$

Since

$$\begin{aligned} \text{SD}(\text{constant} + X) &= \text{SD}(X) \\ \text{SD}(\text{constant} \times X) &= \text{constant} \times \text{SD}(X) \end{aligned}$$

we have that

$$\text{SD}(\text{output}) \approx |f'(P_0)| \times \text{SD}(\text{input}). \quad (1)$$

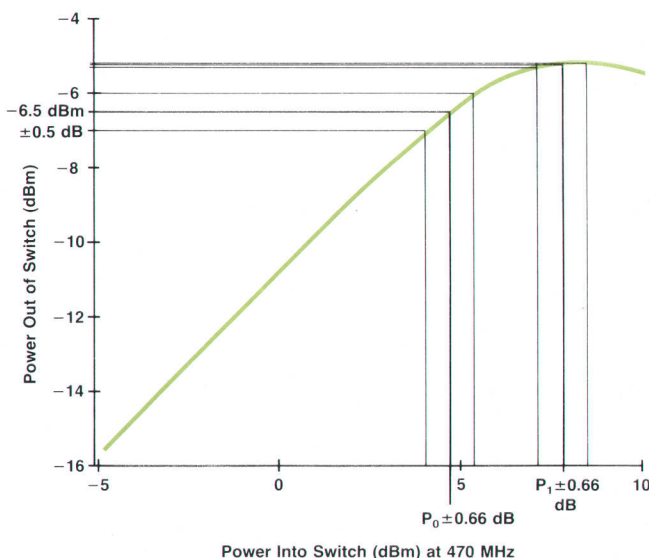


Fig. 2. Typical switch response curve at 470 MHz, showing input power at $P_0 \pm 0.66$ dB to achieve target output of -6.5 ± 0.5 dBm.

Thus, output variation is related to the input variation through the slope of the function relating the output to the input near the point of input power. One may question the accuracy of this variation when we use only the leading term in the Taylor series expansion. Examples indicate that such approximations are in fact quite satisfactory.⁴ Further uses of this law of propagation of error will be demonstrated later in connection with combining several sources of variation into a single uncertainty statement.

To meet other specifications, the switch must operate where its response is linear. Using equation 1, and given the maximum allowable standard deviation of the output and the slope of the curve of the designed switch, one can calculate the maximum allowable variation in input power to the switch. To illustrate, the switch in Fig. 2 has a target output of -6.5 dBm and an allowable variation of ± 1.5 dB around this target. The allowable overall (RSS) standard deviation of the output power is therefore 0.6 dB, since 2.5 multiples of this standard deviation equals the desired ± 1.5 -dB window with 99% confidence. If the output power variation from components is 0.3 dB, then the allowable output variation from transmission of error by the switch is 0.5 dB, since 0.3 dB and 0.5 dB yield an RSS result of approximately 0.6 dB. The slope of the switch function at the input level (denoted P_0 in Fig. 2) corresponding to an output of -6.5 dBm is 0.76 . Thus, the maximum allowable variation in input power to the switch is $0.5/0.76 = 0.66$ dB around a target input power of $P_0 = 4.69$ dBm.

Fig. 3 shows a family of amplifier curves designed with three values of collector current. Notice that the effect of the collector current is to vary the point at which the slope of the curve starts to decrease, that is, the point of amplifier compression. The price for this smaller slope (and hence smaller output variation) is less output power. Recall that sufficient power out of the amplifier is necessary to drive the switch and obtain the target -6.5 dBm out of the switch. (If the amplifier output exceeds 4.69 dBm, the loss of the pad before the switch can be increased to absorb the excess power.) In addition, we seek those curves whose variation at the output of 4.69 dBm or above does not exceed 0.66 dB. Fig. 3 shows that the slope, and hence the output vari-

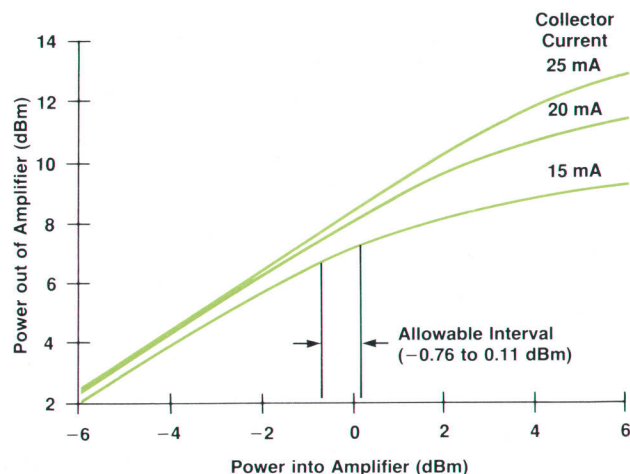


Fig. 3. Amplifier response curves for three values of collector current at 470 MHz.

ation, is smaller at higher input power levels (amplifier compression), but recall that, on average, only 0.11 dBm of power is available at the input to the amplifier.

With a maximum limit of 0.11 dBm of input power, we calculate the slopes of the curves corresponding to the three amplifiers in Fig. 3. These slopes are estimated by successive first differences and are shown in Fig. 4. Notice that the slopes do not drop below the maximum allowable slope of 0.66 until the input power reaches at least -0.76 dBm, 1.7 dBm, or 3.7 dBm, respectively. Only the 15 -mA curve satisfies both the incoming power constraint and the maximum slope constraint.

The allowable interval (-0.76 to 0.11 dBm) is indicated in Fig. 3. Since the slope is smallest at 0.11 dBm, one might reduce the output power of 7.31 dBm after the amplifier with a 2.62 -dB pad. For other engineering reasons, a 1 -dB pad is placed before the amplifier. This increases the slope only slightly to 0.70 . The output of the amplifier at $P_{in} = -0.89$ dBm is 6.62 dBm, so a 2 -dB loss on the pad preceding the switch ensures the target -6.5 dBm at the switch output.

This same analysis is applied to all circuits in the prototype stage of the design. An additional consideration is the effect of temperature on the circuits. This effect is discussed in the next section.

Temperature Testing

Circuits must perform under a variety of environmental conditions. One such set of conditions is extreme temperatures. Measuring at least two circuits of each type at 5°C (cold), 25°C (room), and 70°C (hot) allows an estimate of the effect of temperature on the circuits. This effect on some circuits is minimal; for example, for the power splitter in Fig. 1, the change in power is only ± 0.05 dB, depending on whether the circuit is hot ($-$) or cold ($+$). Likewise, amplifiers are largely unaffected by temperature. Fig. 5 shows a set of six curves: two amplifiers, each measured at room, hot, and cold temperatures. The curves are nearly identical. Heat decreases the output power by an average of only 0.28 dB, and cold increases the power by an average of only 0.33 dB. (The temperature biases are nearly zero when the amplifiers are in compression.) Regardless of




Fig. 4. Slopes of the amplifier curves of Fig. 3.

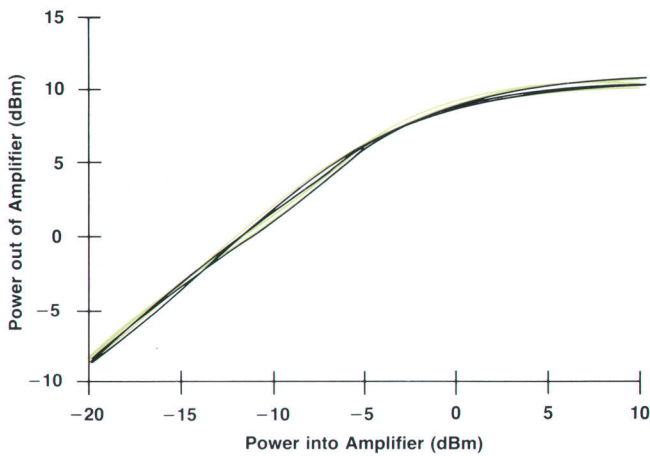


Fig. 5. Six amplifier response curves measured at 5°C, 25°C, and 70°C for two amplifiers.

output frequency, an overall 1σ uncertainty of 0.27 dB accounts for variation in both temperature and amplifier components.

The temperature effect is not negligible for switches. Fig. 6 shows the difference in power between room and hot temperatures and between room and cold temperatures for two switch assemblies. Notice that this effect is not constant for all switches. In the region of interest ($P_{in} = 0$ to 5 dBm), the power at 25°C is roughly 0.8 dB higher than at 70°C for one switch and 0.96 dB higher for another switch. Similarly, one switch is 0.4 dB lower at 25°C than at 5°C, while a second switch is 0.3 dB lower at 25°C. The temperature biases for the two switch assemblies are averaged: $\Delta_{Hot} = -0.88$ dB and $\Delta_{Cold} = 0.35$ dB. The variation in these averages between different switch assemblies constitutes part of the component variation, which is discussed in the following section.




Fig. 6. The upper curves show room-temperature (25°C) output power relative to hot-temperature (70°C) output power as a function of input power for two switches, a and b. The lower curves show room-temperature output power relative to cold-temperature (5°C) output power for the same two switches.

Component Variation

Component variation can be assessed in two ways: manufacturer specification or actual measurement data. The first approach requires a translation of the manufacturer specification into multiples of σ . Usually the best guess is to treat the specification as a 3σ limit. In this study, this approach was limited mostly to characterizing the variation of purchased low-pass and bandpass filters. More often, at least two circuits of each type (power splitters, printed circuit board filters, amplifiers, dividers, switches, XOR gates) were measured to provide more realistic determinations of the uncertainties in the components. Such uncertainty necessarily includes measurement error.

We have seen the effects of component variation both in the amplifiers and in assessing the temperature biases on the switches. The component variation, like the temperature effect, is small for the amplifiers; 1σ estimate of the combined variation is 0.27 dB. The component variation in the switch circuits is higher, as Fig. 7 illustrates. Here, only a portion of the switch curve is plotted to exaggerate the differences in the curves resulting from measuring two different circuits on two different boards. The lower curves, representing measurements of two circuits on board A, are clearly closer to each other than they are to either of the upper curves, representing measurements of two circuits on board B. The board variation is substantial; here, the difference is approximately 0.7 dB. This variation may be caused by various other components on the boards, such as the low-pass filter that follows the switch. On a given board, measured differences between circuits are smaller.

A difference of two measurements can be translated into a 1σ estimate by dividing the difference by $\sqrt{2}$. (In general, for Gaussian data, an estimate of σ from the difference between the largest and smallest of n observations is obtained by dividing this difference by \sqrt{n} , when $n < 10$. See page 317 of reference 4.) This additional variation is included in the overall uncertainty statement, as we discuss in the following section.




Fig. 7. Response curves for two switches on board A (lower curves) and two switches on board B (upper curves).

Combining Uncertainties

The root sum of squares (RSS) rule and the Taylor series rule are special cases of the general law of the propagation of error, which can be used in a general way to combine estimates of variation into a single overall uncertainty statement. The general law of propagation of error says that, for a function of k independent variables

$$y = f(x_1, x_2, \dots, x_k)$$

the variance (square of the standard deviation) of y is approximated by

$$\text{Var}(y) = \sum_{i=1}^k (\partial f / \partial x_i)^2 \text{Var}(x_i)$$

where the partial derivatives are evaluated at the sample means of the k variables (see page 363 of reference 5). Notice that when $k = 1$, taking the square root of both sides leads to equation 1 above. In this paper, the response, output power, involves variations in input power, components, board version, and measurement error. Temperature effects are treated as additive biases. The simplest model

for output power is an additive one:

$$\text{output power} = A(\text{input power}) + \text{component effects} + \text{board effects} + \text{ME}$$

where $A(\text{input power})$ is an amplifier function like those shown in Figs. 3 and 5 and ME stands for measurement error. According to the law of propagation of error,

$$\text{Var}(\text{output}) = A'^2(P_{\text{in}}) \times \text{Var}(\text{input}) + \text{Var}(\text{components}) + \text{Var}(\text{boards}) + \text{Var}(\text{ME})$$

so it follows that

$$\text{SD}(P_{\text{out}}) = [A'^2(P_{\text{in}})\sigma_{\text{in}}^2 + \sigma_{\text{comp}}^2 + \sigma_{\text{board}}^2 + \sigma_{\text{ME}}^2]^{1/2},$$

which is the familiar root sum of squares formula. Notice that nowhere have we assumed a specific probability distribution for any of these variables, only that the variables are independent.

We are now ready to combine the previously identified sources of variation into a single uncertainty statement, applying this law after every circuit. We treat the power levels at the three different temperatures separately, since 5°C and 70°C are extreme conditions, and in a given operating environment, the temperature is not likely to vary over such a broad range. Also, the measurement error σ_{ME} has been included in the measurement of component and board variation.

Consider our previous illustration at room temperature, for which $P_{\text{in}} = 3.91 \text{ dBm} \pm 0.94 \text{ dB}$. The power splitter and the pad incur power losses of $-3.8 \text{ dB} \pm 0.16 \text{ dB}$ and $1 \text{ dB} \pm 0.00 \text{ dB}$ respectively. Thus P_{in} to the amplifier is $-0.89 \text{ dBm} \pm 0.95 \text{ dB}$. The amplifier increases this power to 6.62 dBm , where the slope is 0.7, so the output variation is reduced to $\pm 0.70 \text{ dB}$. Amplifier component and temperature variation of 0.27 dB raises this uncertainty to $\pm 0.75 \text{ dB}$. A 2-dB loss results in $P_{\text{in}} = 4.62 \text{ dBm} \pm 0.75 \text{ dB}$ to the switch, which itself has a slope of 0.76 at P_{in} . Thus the output power is $-6.48 \text{ dBm} \pm 0.57 \text{ dB}$. Combining this with the board variation ($\pm 0.40 \text{ dB}$) by the root sum of squares method, we have a final uncertainty of $\pm 0.70 \text{ dB}$. A similar path of analysis follows from the power levels

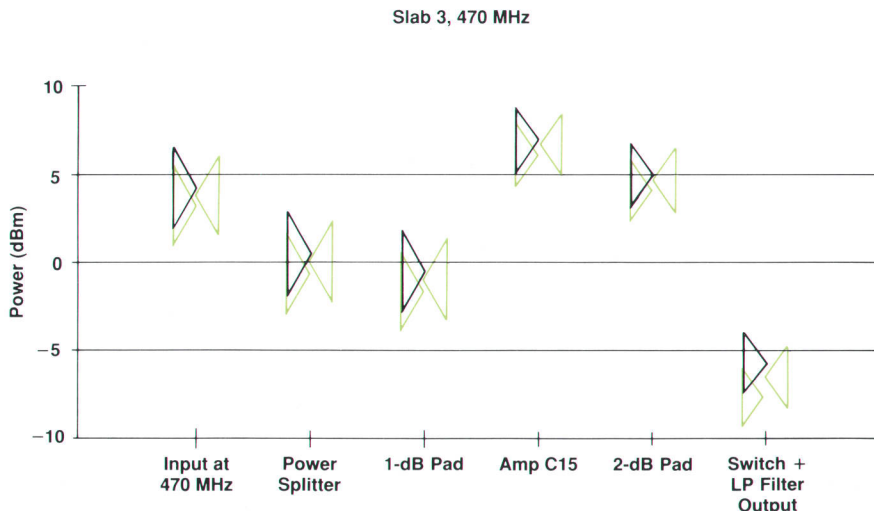


Fig. 8. Triangles indicate the expected values and the 2.5σ limits of the power output from each stage for three temperatures: room (right, color), hot (left, color), and cold (left, black).

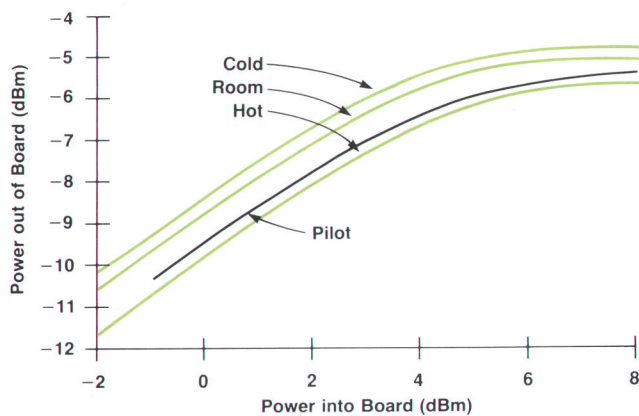


Fig. 9. Room, hot, and cold response curves for an average of four prototype synthesizer boards, and a room-temperature response curve for a lab pilot board.

at hot and cold temperatures.

A graphical display of all these numbers is shown in Fig. 8. The expected power level at each stage is indicated by the center of the triangle on the vertical axis corresponding to each circuit. There is one triangle for each temperature: room on the right, hot and cold on the left. The length of the base of the triangle extends to 2.5 multiples of the standard deviation of the power output at each stage. Thus, considerable information is conveyed in this display: changes in power levels across circuits, variations in power levels at different stages, and differences in power levels at extreme temperatures. The reduction in variation from the beginning to the end of the path can be seen, and the target $P_{out} = -6.5$ dBm is attained within the desired limits at room temperature.

Comparing Prototype Analysis with Lab Pilot Results

The path of power shown in Fig. 1 can be formulated into a single functional relationship between input power to the power splitter and output power from the switch. This functional relationship is computed using linear interpolation in the data collected on amplifiers and switches

and is shown in Fig. 9. A lab pilot board was measured for its power transfer between the initial and final points of the circuitry: it is shown in Fig. 9 for room temperature only. The agreement in the general shape of the curves between observed lab pilot measurements and expected values based on the prototype is excellent. Any absolute difference between the curves can be explained by impedance mismatches; pads can be adjusted accordingly. Once these pad values are ascertained for the lab pilot, target power levels are achieved within the designed limits, without the need for further adjustments of pads in production.

More Complicated Analyses

Not all power paths in the instrument are as straightforward as the one illustrated in this paper. Fig. 10 shows a display similar to Fig. 8 of the uncertainties in a power transfer involving incoming power at one of two different frequencies. While the approach to amplifier design and estimation of component variation is similar, the estimation of variation resulting from input power is complicated by the presence of incoming power at another frequency. Thus, measurements must be taken by varying power levels at both frequencies. Other factors, such as board type and circuit components, can also be varied, and a statistically designed plan to take measurements on several different factors simultaneously is encouraged.⁵

References

1. Report of the Bureau International des Poids et Mésures on the statement of uncertainties, October 21-23, 1980. Partly reprinted in *Metrologia*, Vol. 17, pp. 73-74 (P. Giacimo).
2. H.H. Ku, "Statistical Concepts in Metrology," 1967, reprinted in H.H. Ku, ed., *Precision Measurement and Calibration: Statistical Concepts and Procedures*, U.S. National Bureau of Standards Special Publication 300, 1969.
3. H.H. Ku, "Notes on the Use of Propagation of Error Formulas," 1966, reprinted in H.H. Ku, ed., *Precision Measurement and Calibration: Statistical Concepts and Procedures*, U.S. National Bureau of Standards Special Publication 300, 1969.
4. J.W. Tukey, "Statistical and Quantitative Methodology," 1961, reprinted in L.V. Jones, ed., *The Collected Works of John Tukey*, Volume III, Wadsworth, pp. 143-185.
5. G.E.P. Box, W.G. Hunter, and J.S. Hunter, *Statistics for Experimenters*, Wiley, 1978.

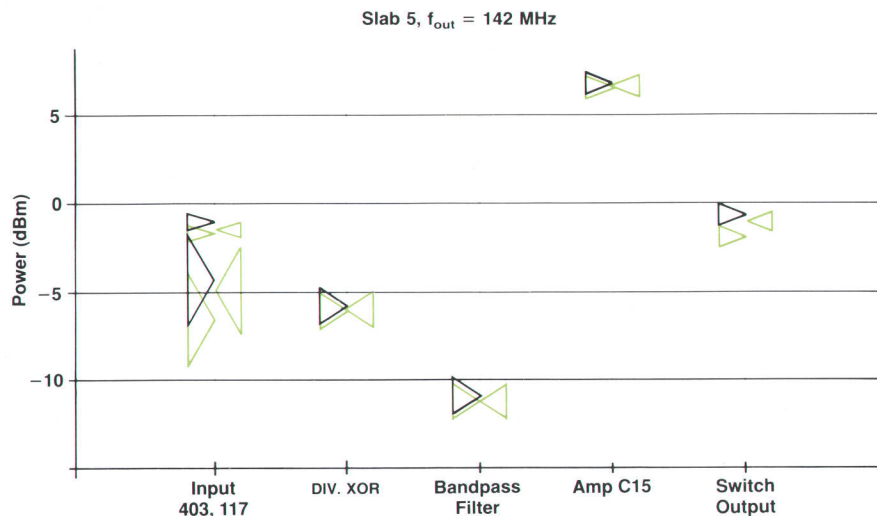


Fig. 10. Uncertainties in a power transfer involving input power at either of two frequencies.

Statistical Calibration of a Vector Demodulator

Circuit performance characteristics must be determined and accounted for in instrument calibration. This paper describes an algorithm for statistical demodulator calibration to guarantee high precision in the demodulated signal.

by Karen Kafadar

VECTOR MODULATION INSTRUMENTS such as the HP 8780A Vector Signal Generator¹ and the HP 8980A Vector Analyzer² provide a test system for measurement applications in such fields as digital microwave radio, communications, and radar. Some of the technologically new developments in this series of products are the modulation capabilities (complex, wideband signals), the high-resolution CRT display, and the 350-MHz bandwidth.

The precision measurement capabilities offered in these products raise the issue of the internal calibration of the instruments. All instruments need to be calibrated to account for differences in components and test conditions. The accuracy and precision of these internal self-adjustments are crucial, since they have direct effects on the characterization of the device under test. The hardware measurements made for purposes of internal calibration must be combined in a logical and intelligent fashion to yield the correct calibration factors.

Statistical methods play an important role in the calibration routine. This paper describes a statistical approach to the calibration of the HP 8981A Vector Modulation Analyzer. The next section briefly introduces the use of statistical methods in such situations. The measurement process for a vector demodulator is then described, using the HP 8981A as an illustration. Two sections provide the technical details of the calibration algorithm both with and without an accurate vector signal generator (e.g., the HP 8780A) as a reference. Finally, examples of real data illustrate the computations in the statistical algorithm and formulas for the uncertainties in the calibration factors. It will be shown that the algorithm can be used to calibrate not only the demodulator in the HP 8981A but also external demodulators.

Statistics in Calibration Routines

Calibration consists of comparing a set of measurements from an uncharacterized instrument (e.g., an HP 8981A) with a defined reference standard (e.g., an HP 8780A) according to a measurement algorithm. Thus the calibration model relates the observed measurement readings to the reference standard. A model is never perfect, and the difference between the model and reality can be characterized according to two types of errors that may be present:

Part of this paper appeared in *MSN & Communications Technology*, March 1988. Reprinted by permission of the publisher. Copyright EW Communications, Inc.

- Systematic errors resulting from incomplete specification of the calibration model
- Random errors, or small, unpredictable fluctuations that affect every measurement but are themselves unmeasurable.

Statistical methods can be used to address both types of errors. First, the calibration should be designed to allow identification of possible systematic departures between the model for measurement and the observed data. Second, the influence of random errors can be assessed when estimating the parameters in the model that relate to the measurement process and their uncertainties. In this paper, we shall assume that the random errors are independent and identically distributed according to a symmetric distribution; that is, that the individual errors cannot be predicted in either size or direction, and that the chances of an erroneous measurement being too large or too small are roughly equal. (Diagnostic tools for checking the validity of these assumptions and the consequences of their violation are discussed in connection with the examples later in this paper.)

Rarely is a single estimate of a target quantity sufficient. For example, reporting a sample mean without its standard error provides no information on the reliability of the data that went into that sample mean. The same is true for estimates of the parameters describing the measurement process. These estimates are the calibration factors, and certain limits of fluctuation may be desirable (e.g., gain adjustment accurate within 1%). Therefore, an important part of the statistics in a calibration algorithm is the derivation of associated measures of uncertainty for the calibration factors.

Vector Analyzer Measurement Process

In the HP 8981A, the phase-modulated signal can be expressed mathematically in terms of its frequency components as

$$\begin{aligned} s(t) &= \sum_k A_k(t) \cos(\omega_k t + \phi_k) \\ &= \sum_k I_k(t) \cos(\omega_k t) + Q_k(t) \sin(\omega_k t) \end{aligned} \quad (1)$$

Equation 1 expresses the signal in terms of its demodulated components $I_k(t)$ (in-phase) and $Q_k(t)$ (quadrature).

Components that vary over time can be observed by examining $I_k(t)$ and $Q_k(t)$ over time at a particular frequency ω_k . Ideally, at a fixed time, a single-frequency signal with constant amplitude is demodulated at

$$I = \cos(a)$$

$$Q = \cos(a - (90^\circ - \phi))$$

where a is the angle of the vector signal and ϕ , the error in the quadrature phase of the demodulation, is zero.

Since the HP 8981A has the ability to display signals in either time or vector form, amplitude and phase information must be recovered accurately. This requires special software functions to estimate and subsequently account for any amplitude or phase offsets. Proper calibration of the HP 8981A thus allows simple measurements of amplitude and phase modulated signals.

Fig. 1 is a block diagram for a simplified demodulator, such as the HP 8981A, which decomposes the modulation on a signal into its in-phase and quadrature components. Accuracy and precision in this decomposition are essential to assess degradation of the signal's amplitude, phase, or frequency, or to assess cross-talk between the two channels. Since I and Q signals are used often in measurement systems, we will discuss the calibration scheme for the Q-vs-I vector representation.

Six sources for adjustments in a vector demodulator can be identified, based on the block diagram of the instrument. These six sources correspond to stages in the transformation of the ideal I/Q signal into the measured, or actual, I/Q signal. The calibration process uses statistics to estimate the sizes and directions of the adjustments attributable to these sources and then corrects the output so that the demodulated I/Q signal agrees with the ideal I/Q signal.

Let (I,Q) denote the pair of points representing the I and Q channels at a particular point in time at a given frequency. The HP 8981A hardware transforms the coordinates (I,Q) into measured coordinates (x,y) through the following series of operations:

$$(I,Q) \rightarrow \text{Rotation} \rightarrow \text{Axis Sheer} \rightarrow \text{Compression} \rightarrow \text{dc Offsets} \rightarrow (x,y)$$

Rotation can be viewed as a phase change in the ideal signal, which can be introduced by any of the components before the demodulation. The axis sheer creates two additional adjustments: quadrature adjustment, resulting from inaccuracy in the 90° split, and gain imbalance, resulting from different power levels in amplifiers or unequal power losses in other components after the split. The I/Q measurements are rescaled by amplifier compression, and dc offsets may be added to the measurements at any stage.

Notationally, let us represent these adjustments as:

- γ = gain imbalance (I-channel/Q-channel)
- ϕ = angular difference between ideal quadrature (90°) and actual quadrature
- θ = angle of rotation (sometimes called lock error)
- ρ = scaling resulting from compression
- I_0 = dc offset in I-channel
- Q_0 = dc offset in Q-channel

The four steps in transforming (I,Q) to (x,y) can be expressed mathematically as follows:

- Rotation through angle θ :

$$\mathbf{R}: \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$$

This transforms (I,Q) into $(I\cos\theta - Q\sin\theta, I\sin\theta + Q\cos\theta)$.

- Sheer transformation (quadrature adjustment ϕ , gain imbalance γ):

$$\mathbf{S}: \begin{bmatrix} \gamma & 0 \\ \sin\phi & \cos\phi \end{bmatrix}$$

This causes two additional adjustments: the ratio in the two channels (I/Q) is now γ , not 1, and the tangent of the angle between the two axes (I,Q) is now $(90^\circ - \phi)$, not 90° .

- The measurements are uniformly compressed according to a factor ρ .
- The measurements are offset by a fixed amount (I_0 in the I-channel, Q_0 in the Q-channel).

Putting these four effects together, the resultant point (x,y) can be expressed as a function of the ideal input (I,Q) as:

$$\begin{bmatrix} x \\ y \end{bmatrix} = \rho \begin{bmatrix} \gamma & 0 \\ \sin\phi & \cos\phi \end{bmatrix} \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} I \\ Q \end{bmatrix} + \begin{bmatrix} I_0 \\ Q_0 \end{bmatrix} + \begin{bmatrix} \epsilon_x \\ \epsilon_y \end{bmatrix} \quad (2)$$

The unobserved variables ϵ_x, ϵ_y are assumed to be uncorrelated random errors in the measurement of the x and y coordinates. That is,

$$x = I_0 + \gamma(\rho\cos\theta)I - \gamma(\rho\sin\theta)Q + \epsilon_x$$

$$= \alpha_0 + \alpha_1 I + \alpha_2 Q + \epsilon_x$$

$$y = Q_0 + (\rho\cos\theta\sin\phi + \rho\sin\theta\cos\phi)I$$

$$+ (-\rho\sin\theta\sin\phi + \rho\cos\theta\cos\phi)Q + \epsilon_y$$

$$= \beta_0 + \beta_1 I + \beta_2 Q + \epsilon_y$$

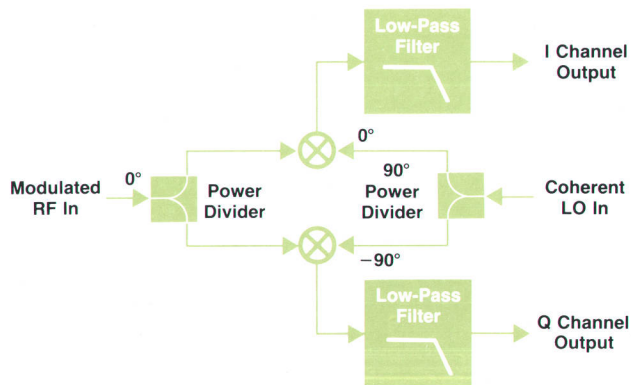


Fig. 1. Simplified vector demodulator.

where:

$$\begin{aligned}\alpha_0 &= I_0 \\ \alpha_1 &= \gamma\rho\cos\theta \\ \alpha_2 &= -\gamma\rho\sin\theta \\ \beta_0 &= Q_0 \\ \beta_1 &= \rho\cos\theta\sin\phi + \rho\sin\theta\cos\phi \\ &= (\alpha_1\sin\phi - \alpha_2\cos\phi)/\gamma \\ \beta_2 &= -\rho\sin\theta\sin\phi + \rho\cos\theta\cos\phi \\ &= (\alpha_2\sin\phi + \alpha_1\cos\phi)/\gamma\end{aligned}$$

The errors of interest can be recovered via the equations:

$$\begin{aligned}I_0 &= \alpha_0 && \text{(I-channel offset)} \\ Q_0 &= \beta_0 && \text{(Q-channel offset)} \\ \rho &= (\beta_1^2 + \beta_2^2)^{1/2} && \text{(Compression)} \\ \theta &= \tan^{-1}(-\alpha_2/\alpha_1) && \text{(Rotation)} \\ \gamma &= [(\alpha_1^2 + \alpha_2^2)/(\beta_1^2 + \beta_2^2)]^{1/2} && \text{(I/Q gain imbalance)} \\ \phi &= \tan^{-1}[(\alpha_1\beta_1 + \alpha_2\beta_2)/(\alpha_1\beta_2 - \alpha_2\beta_1)] && \text{(Quadrature error)}\end{aligned}\quad (3)$$

Notice that, apart from rotation, the ideal I and Q can be recovered via the equations

$$\begin{aligned}I &= (x - I_0)/\gamma\rho \\ Q &= [\gamma(y - Q_0) - \sin\phi(x - I_0)]/\gamma\rho\cos\phi.\end{aligned}$$

Estimation with Accurate Reference

At this point, we need to distinguish between the case where ideal (I,Q) values are provided and the situation where only magnitude information about the (I,Q) points is known. Ideal (I,Q) measurements can be obtained using a precise signal generator such as the HP 8780A, so that signals having several accurately defined phase modes can be guaranteed. Since this situation lends itself to the easier algorithm, it will be discussed first.

It is evident that an input signal of constant frequency and magnitude and variable phase, namely

$$s(t) = A\cos(\omega t + \phi_t), t = 0, 1, \dots$$

results in an I/Q display for which

$$I^2 + Q^2 = (\text{Magnitude})^2,$$

which is a circle. It can be verified that the transformed points (x,y) lie on an ellipse. The parameters of this ellipse are related to the adjustments in the instrument as outlined above.

Since we have six adjustments (parameters), we need a minimum of six phase states to obtain an exact fit for the parameters in equation 2.

The 8PSK mode of the HP 8780A allows rapid generation of an I/Q signal at eight evenly distributed phase states, allowing two extra phase states for model confirmation. The coordinates (I,Q) of the points in any one of these states will be considered ideal, transformed (by the six adjustments) into measured coordinates (x,y). For simplicity, we label the eight ideal states sequentially at those

points having phases $k(45^\circ)$, $k = 0, 1, \dots, 7$ (Fig. 2). Thus,

State	I	Q
0	1	0
1	f	f
2	0	1
3	-f	f
4	-1	0
5	-f	-f
6	0	-1
7	f	-f

where $f = 1/\sqrt{2} = 0.7071068$. (Numerical association with each state is arbitrary.)

Suppose $n/8$ points are measured in each state, so n measurements in all are taken. (This requirement is not essential, but the equations are greatly simplified by encouraging orthogonality into the problem.) Denote the coordinates of these measurements by $(x_1, y_1), \dots, (x_n, y_n)$. Associate the point with its ideal (I,Q) pair as described above: $(I_1, Q_1), \dots, (I_n, Q_n)$. Then we need to fit $\alpha = (\alpha_0, \alpha_1, \alpha_2)'$ and $\beta = (\beta_0, \beta_1, \beta_2)'$ in the two relationships:

$$\begin{aligned}x_i &= \alpha_0 + \alpha_1 I_i + \alpha_2 Q_i + \epsilon_{x_i} \\ y_i &= \beta_0 + \beta_1 I_i + \beta_2 Q_i + \epsilon_{y_i}\end{aligned}$$

The estimates of $(\alpha_0, \alpha_1, \alpha_2)$, namely $(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\alpha}_2)$, can be obtained by a least squares algorithm. If we let M be the matrix of n rows and 3 columns corresponding to the n observations and three "carriers" (constant term, I, Q), that is,

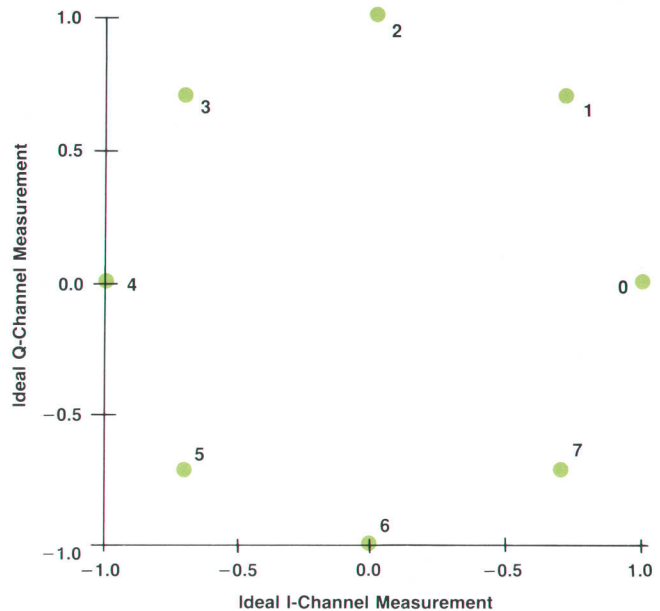


Fig. 2. In-phase (I) and quadrature (Q) modulation states for 8PSK vector modulation.

$$\mathbf{M} = \begin{bmatrix} 1 & I_1 & Q_1 \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ 1 & I_n & Q_n \end{bmatrix}$$

then it is well known³ that

$$\hat{\alpha} = (\mathbf{M}'\mathbf{M})^{-1}\mathbf{M}'\mathbf{x} \quad (4)$$

$$\hat{\beta} = (\mathbf{M}'\mathbf{M})^{-1}\mathbf{M}'\mathbf{y}$$

are the least squares estimates of α and β , where $\mathbf{x} = (x_1, x_2, \dots, x_n)'$ and $\mathbf{y} = (y_1, y_2, \dots, y_n)'$ are column vectors and the prime sign (') indicates transpose. Since there are $n/8$ measurements in each state, and the states are evenly spaced around the circle, the matrix \mathbf{M} is orthogonal, and so

$$\mathbf{M}'\mathbf{M} = \begin{bmatrix} n & 0 & 0 \\ 0 & n/2 & 0 \\ 0 & 0 & n/2 \end{bmatrix}.$$

Furthermore,

$$\mathbf{M}'\mathbf{x} = \left[\sum_{i=1}^n x_i, \sum_{i=1}^n x_i I_i, \sum_{i=1}^n x_i Q_i \right]'$$

and thus

$$\hat{\alpha} = (\mathbf{M}'\mathbf{M})^{-1}\mathbf{M}'\mathbf{x} = [\text{ave}(x), 2\text{ave}(x \cdot I), 2\text{ave}(x \cdot Q)]'$$

$$\hat{\beta} = (\mathbf{M}'\mathbf{M})^{-1}\mathbf{M}'\mathbf{y} = [\text{ave}(y), 2\text{ave}(y \cdot I), 2\text{ave}(y \cdot Q)]'.$$

The estimates of the target parameters ($I_0, Q_0, \rho, \theta, \gamma, \phi$) can then be unscrambled using equations in (3) above.

Standard Errors

We need some measure of uncertainty in our estimates of the parameters. For example, after calculating an estimate of γ , say $\hat{\gamma}$, how far off might that estimate be from the true value? One such measure is the *standard error* (SE) of the estimate. An approximate 95% confidence interval for the estimate is given by

$$\text{estimate} \pm 2(\text{SE}).$$

To calculate the standard errors for the six parameters, we first need a measure of goodness of fit in our model. This can be calculated most simply as a "root mean square" of the differences between the observed measurement x_i or y_i and its fitted value predicted by the model. Denote these differences, or residuals, by:

$$e_{x_i} = x_i - \hat{\alpha}_0 - \hat{\alpha}_1 I_i - \hat{\alpha}_2 Q_i$$

$$e_{y_i} = y_i - \hat{\beta}_0 - \hat{\beta}_1 I_i - \hat{\beta}_2 Q_i$$

and calculate

$$s_x = \sqrt{\sum_{i=1}^n e_{x_i}^2 / (n - 3)}$$

$$s_y = \sqrt{\sum_{i=1}^n e_{y_i}^2 / (n - 3)}.$$

The divisor is $(n - 3)$ instead of n because three parameters have been estimated in each channel. If the departures from the model have a Gaussian distribution, we would expect most of the e_{x_i} to fall within $2s_x$ or $3s_x$, and likewise for e_{y_i} . This criterion may serve as a simple check for unusual measurements (but see "A Caution," page 24).

Using the root mean squares of the model departures (residuals), we can calculate the standard errors of the parameters in the model. Again assuming these residuals are Gaussian distributed, the standard errors of $(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\alpha}_2)$ are given by

$$s_x \sqrt{\text{diagonal of } (\mathbf{M}'\mathbf{M})^{-1}}$$

Thus,

$$\text{SE}(\hat{\alpha}_0) = \text{SE}(\hat{I}_0) = s_x / \sqrt{n}$$

$$\text{SE}(\hat{\beta}_0) = \text{SE}(\hat{Q}_0) = s_y / \sqrt{n}$$

The standard errors for the other target parameters are a little more complicated, because they are functions of $\alpha_1, \alpha_2, \beta_1,$ and β_2 . We rely on propagation of error formulas⁴ for approximations to the standard errors of functions of variables. The approximations are simplified if

$$2s_y^2 / (n\rho^2) \ll 1,$$

a condition that will often be satisfied if the fit is good (s_y is small) and the number of points is moderately large. Then:

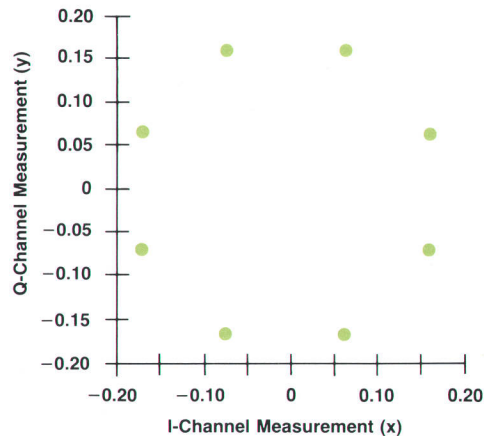


Fig. 3. Data for the first example (see text). Six measurements were taken in each of eight states.

$$SE(\hat{\rho}) \approx s_y \sqrt{2/n}$$

$$SE(\hat{\gamma}) \approx s_x \sqrt{2/n}$$

$$SE(\overline{\tan\theta}) \approx s_x \hat{\rho} \hat{\gamma} \sqrt{2/n/\alpha_1^2}$$

$$SE(\overline{\tan\phi}) \approx [(s_x^2 + s_y^2)(1 + 1/\hat{\gamma}^2)/n]^{1/2}/(\hat{\rho}\cos^2\hat{\phi})$$

where $\overline{\tan\theta}$ and $\overline{\tan\phi}$ are the estimates of $\tan\theta$ and $\tan\phi$. So approximate 95% confidence intervals for the estimates of the errors are:

$$\begin{aligned} \hat{\gamma} &\pm 2SE(\hat{\gamma}) \\ \hat{I}_0 &\pm 2SE(\hat{I}_0) \\ \hat{Q}_0 &\pm 2SE(\hat{Q}_0) \\ \hat{\rho} &\pm 2SE(\hat{\rho}) \\ \tan^{-1}[\overline{\tan\theta} \pm 2SE(\overline{\tan\theta})] \\ \tan^{-1}[\overline{\tan\phi} \pm 2SE(\overline{\tan\phi})]. \end{aligned}$$

This calibration algorithm thus proceeds as follows:

- Measure (x_i, y_i) for $i = 1, \dots, n$
- Associate state (I_i, Q_i) for $i = 1, \dots, n$
- Calculate: $\hat{\alpha} = (\hat{\alpha}_0, \hat{\alpha}_1, \hat{\alpha}_2)'$
 $\hat{\beta} = (\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2)'$
- Calculate target parameters: $\gamma, \phi, \rho, \theta, I_0, Q_0$
- Calculate: goodness of fit (s_x, s_y) and residuals (e_x, e_y)
- Identify unusual residuals
- Calculate standard errors.

Inaccurately Known Phase

We now consider the situation where phase cannot be associated precisely with each measurement. Generally, amplitude information is more reliable, that is,

$$I^2 + Q^2 = (I, Q)'(I, Q) = 1.$$

Two possible signals that satisfy this requirement are:

- A signal with continuously varying phase (e.g., two tones, offset in frequency)
- A signal with eight discrete phases (e.g., an 8PSK modulated signal and its coherent carrier).

Inverting equation 2,

$$\begin{bmatrix} I \\ Q \end{bmatrix} = \rho^{-1} \mathbf{R}^{-1} \mathbf{S}^{-1} \begin{bmatrix} x - I_0 - \epsilon_x \\ y - Q_0 - \epsilon_y \end{bmatrix}$$

and thus

$$(I, Q)'(I, Q) = \tag{5}$$

$$\rho^{-2} (x - I_0 - \epsilon_x, y - Q_0 - \epsilon_y) (\mathbf{S}^{-1})' (\mathbf{S}^{-1}) (x - I_0 - \epsilon_x, y - Q_0 - \epsilon_y),$$

because $(\mathbf{R}'\mathbf{R}) = (\mathbf{R}^{-1})'(\mathbf{R}^{-1})$ is the identity matrix (\mathbf{R} is orthogonal). This can be simplified by noting that

$$\begin{aligned} (\mathbf{S}^{-1})'(\mathbf{S}^{-1}) &= (\gamma \cos\phi)^{-2} \begin{bmatrix} \cos\phi & -\sin\phi \\ 0 & \gamma \end{bmatrix} \begin{bmatrix} \cos\phi & 0 \\ -\sin\phi & \gamma \end{bmatrix} \\ &= (\gamma \cos\phi)^{-2} \begin{bmatrix} 1 & -\gamma \sin\phi \\ -\gamma \sin\phi & \gamma^2 \end{bmatrix} \end{aligned}$$

Let us assume that the errors in measuring x and y have zero means, are uncorrelated with one another, and have common standard deviation σ . Then, on average,

$$\begin{aligned} \text{ave}(\epsilon_x) &= \text{ave}(\epsilon_y) = 0 \\ \text{ave}(\epsilon_x^2) &= \text{ave}(\epsilon_y^2) = \sigma^2 \end{aligned}$$

and so equation 5 becomes:

$$\begin{aligned} (x - I_0)^2 - 2\gamma(\sin\phi)(x - I_0)(y - Q_0) + \gamma^2(y - Q_0)^2 + \\ (\gamma^2 + 1)\sigma^2 - \gamma^2\rho^2\cos^2\phi = 0. \end{aligned} \tag{6}$$

If measurement error is negligible compared to the magnitude of the adjustments, the fourth term in equation 6 can be ignored.

Equation 6 is nonlinear in the parameters, of which there are now only five (clearly the rotation, or lock error, cannot be estimated without relative phase information). Algorithms for nonlinear least squares often rely on gradient methods or on linearization of the problem via a Taylor series expansion. In addition, several methods for approximate confidence intervals for the parameters have been proposed.⁵ These two issues are described in the Appendix on page 24.


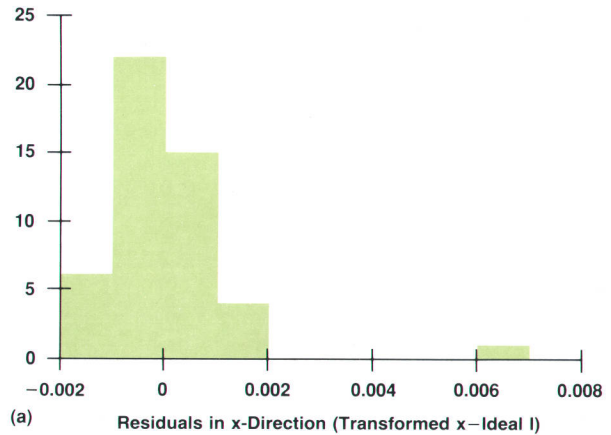


Fig. 4. Histograms of the x and y residuals (e_x, e_y) for the first example. The histograms show the differences between the transformed measurements in each channel and their ideal values as shown graphically in Fig. 2.

Examples

This first example serves as a "control." Six (x,y) measurements were taken in each of eight states (Fig. 3). All calibration factors are essentially at their nominal values.

$$\begin{aligned}(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\alpha}_2) &= (0.000054, 0.16564, -0.068486) \\(\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2) &= (-0.002694, 0.068058, 0.163904) \\ \hat{I}_0 &= 0.00005 \text{ volts (I-channel offset)} \\ \hat{Q}_0 &= -0.00269 \text{ volts (Q-channel offset)} \\ \hat{\rho} &= 0.17747 \text{ volts (Compression)} \\ \hat{\theta} &= \tan^{-1}(0.39206) = 21.41^\circ \text{ (Lock error)} \\ \hat{\gamma} &= 1.0100 \text{ (I/Q gain imbalance)} \\ \hat{\phi} &= \tan^{-1}(0.00152) = 0.0871^\circ \text{ (Quadrature error),}\end{aligned}$$

Root mean squares: $s_x = 0.0012162$; $s_y = 0.0008686$ volts.

Approximate 95% confidence intervals (estimate $\pm 2SE$):

$$\begin{aligned}I_0 &: (-0.00044, 0.00055) \text{ volts} \\ Q_0 &: (-0.00295, -0.00243) \text{ volts} \\ \gamma &: (1.0053, 1.0147) \\ \rho &: (0.17711, 0.17783) \text{ volts} \\ \tan\phi &: (-0.00144, 0.00448) \rightarrow \phi: (-0.08^\circ, 0.26^\circ) \\ \tan\theta &: (0.39153, 0.39253) \rightarrow \theta: (21.38^\circ, 21.43^\circ).\end{aligned}$$

Fig. 4 shows the histograms of the x and y residuals. Clearly there is one aberrant measurement. This measurement occurred in state 4 ($I = -1$, $Q = 0$). A comparison of these estimates with those obtained with a more robust procedure (downweighting this point) verified that the point is not so extreme that it affected the quality of the overall fit.

Using the algorithm for unknown phase (see above),

$$\begin{aligned}I_0 &: 0.00019 (-0.00044, 0.00055) \text{ volts} \\ Q_0 &: -0.00243 (-0.00295, -0.00243) \text{ volts} \\ \gamma &: 1.0100 (1.0053, 1.0147) \\ \rho &: 0.17748 (0.17711, 0.17783) \text{ volts} \\ \tan(\phi) &: 0.00230 (-0.00318, 0.00622) \rightarrow \phi: 0.13^\circ \\ & \quad (-0.18^\circ, 0.36^\circ).\end{aligned}$$

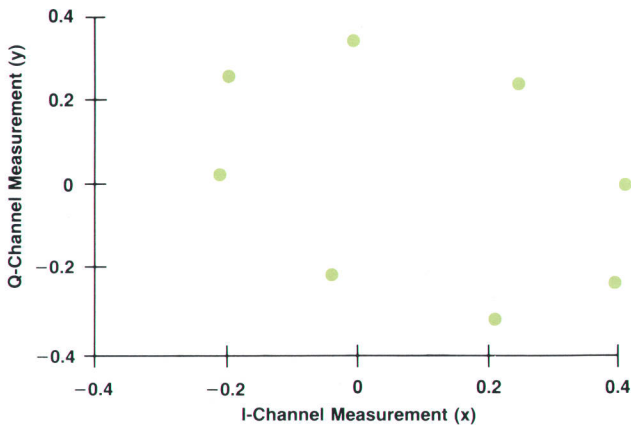


Fig. 5. Data for the second example. Eight measurements were taken in each of eight states.

Notice that these solutions fall well within the confidence intervals of the linear (known-phase) method.

The second example involved a large quadrature adjustment. Eight (x,y) measurements were taken in each of eight states (Fig. 5).

$$\begin{aligned}(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\alpha}_2) &= (0.089621, 0.143384, -0.292603) \\(\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2) &= (0.011144, 0.227279, 0.244292) \\ \hat{I}_0 &= 0.089621 \text{ Volts (I-channel offset)} \\ \hat{Q}_0 &= 0.011144 \text{ Volts (Q-channel offset)} \\ \hat{\rho} &= 0.33367 \text{ Volts (Compression)} \\ \hat{\theta} &= \tan^{-1}(1.11516) = 48.12^\circ \text{ (Lock error)} \\ \hat{\gamma} &= 0.97656 \text{ (I/Q gain ratio)} \\ \hat{\phi} &= \tan^{-1}(-0.36852) = -20.09^\circ \text{ (Quadrature error),}\end{aligned}$$

Root mean squares: $s_x = 0.0019728$; $s_y = 0.0015534$ volts.

Approximate 95% confidence intervals (estimate $\pm 2SE$):

$$\begin{aligned}I_0 &: (0.08913, 0.09011) \text{ volts} \\ Q_0 &: (0.01076, 0.01153) \text{ volts} \\ \gamma &: (0.97586, 0.97726) \\ \rho &: (0.33312, 0.33422) \text{ volts} \\ \tan\phi &: (-0.36263, -0.36891) \rightarrow : (-20.25^\circ, -19.94^\circ) \\ \tan\theta &: (1.10963, 1.12069) \rightarrow : (47.97^\circ, 48.26^\circ)\end{aligned}$$


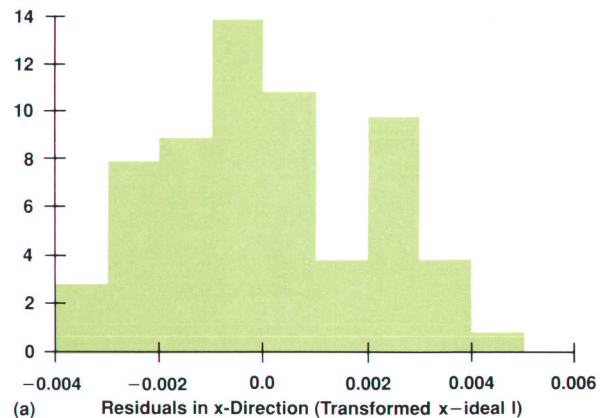


Fig. 6. Histograms of the x and y residuals (e_x , e_y) for the second example.

Appendix

Solving Nonlinear Least Squares Problems

The accompanying paper discusses the solution to estimating the parameters in a model that is linear in those parameters, e.g.,

$$x_i = \alpha_0 + \alpha_1 I_i + \alpha_2 Q_i + \epsilon_i \quad i = 1, \dots, n.$$

The estimates are given explicitly in equation 4 of the paper.

When the model is nonlinear in the parameters, such as in equation 6 of the paper, one approach is to approximate the function by a first-order Taylor series expansion and solve as for the linear problem. For a function, say f , of the data values (x_i, y_i) of only one parameter, say θ ,

$$0 = f(x_i, y_i; \theta) + \epsilon_i \\ = f(x_i, y_i; \theta) + (\theta - \theta_0) df(x_i, y_i; \theta) / d\theta|_{\theta_0} + \epsilon_i$$

from which it follows that

$$\theta = \theta_0 - f(x_i, y_i; \theta_0) / f'(x_i, y_i; \theta_0). \quad (1)$$

Let θ_0 be a prior guess for the parameter θ . Since there are n values of $f(x_i, y_i; \theta_0)$ and $f'(x_i, y_i; \theta_0)$, we can calculate a new θ by replacing the quantities in equation 1 by their averages.

Equation 6 of the paper actually has five parameters: $\theta = (\gamma, \rho, I_0, Q_0, \phi)$. The multivariate extension of Taylor's theorem becomes

$$f(x_i, y_i; \theta) = f(x_i, y_i; \theta_0) + [df(x_i, y_i; \theta_k) / d\theta_k](\theta - \theta_0),$$

where $[df(x_i, y_i; \theta_k) / d\theta_k]$ is a matrix of partial derivatives of the function f with respect to each of the five parameters $\{\theta_k\}$. Let \mathbf{f}_0 denote the column vector of n values of the function $f(x_i, y_i; \theta_0)$, $i = 1, \dots, n$ and let \mathbf{W} be the $n \times 5$ matrix of values of the function $df(x_i, y_i; \theta_k) / d\theta_k$, $i = 1, \dots, n$, $k = 1, \dots, 5$, all evaluated at the current guess θ_0 . Then, ideally, $f(x_i, y_i; \theta) = 0$, so

$$0 \approx \mathbf{f}_0 + \mathbf{W}(\theta - \theta_0) \\ \text{or} \\ -\mathbf{f}_0 \approx \mathbf{W}(\theta - \theta_0). \quad (2)$$

This problem is in the same format as that for the linear least squares problem, where \mathbf{W} was the $n \times 3$ matrix \mathbf{M} of rows consisting of 1s, I_i , and Q_i , and θ was the vector of parameters denoted by α .

Like equation 4 of the accompanying paper, the solution to equation 2 is given by

$$(\theta - \theta_0) = -(\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T \mathbf{f}_0$$

and thus

$$\theta_{\text{new}} = \theta_0 - (\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T \mathbf{f}_0. \quad (3)$$

Equation 3 is iterated until satisfactory convergence is obtained. (For the examples in the accompanying paper, three to four iterations assured convergence to the fifth significant figure.) This method of finding the least squares solution to a nonlinear problem is known as a Gauss-Seidel iteration and has nearly quadratic convergence when the data is known to fit the model well. The algorithm and other methods are described in reference 6 listed on page 25.

The coverage of approximate confidence intervals for parameters solved by nonlinear least squares algorithms such as this one has been investigated (see reference 5 listed on page 25). The simplest and easiest approximation to the standard error of θ_k is given by

$$\hat{\sigma}[\text{kth diagonal element of } (\mathbf{W}^T \mathbf{W})^{-1}]^{1/2}$$

where $\hat{\sigma}$ is the standard deviation of the residuals.

However, an interval of ± 2 of these "standard errors" around the estimated parameters may have coverage as low as 75% instead of the nominal 95%. Hence, for conservative purposes, we recommend using the multiplier 3 for an approximate 95% confidence interval. Further details can be found in references 5 and 6 listed on page 25.

Fig. 6 shows the histograms of the x and y residuals. Using the algorithm for unknown phase (see above),

$$I_0: 0.08871 \text{ volts} \\ Q_0: 0.01187 \text{ volts} \\ \gamma: 0.97723 \\ \rho: 0.33978 \text{ volts} \\ \tan(\phi): -0.36460 \rightarrow -20.03^\circ (-0.18^\circ, 0.36^\circ).$$

Notice that a large quadrature adjustment (-19.25°) and a slight adjustment in the gain imbalance are indicated. The other errors are considered negligible.

A Caution

If the model departures (residuals) are not Gaussian distributed for any reason (e.g., outlying points as in the first example, or misidentified states), or if there are unequal numbers of measurements in each state in the known-phase algorithm, the estimates of the target parameters and their standard errors may be seriously biased. Least squares fitting procedures have a tendency to make all residuals about the same magnitude, even if, in fact, the residuals in all measurements were tiny, except for one. One indication of this case is if the median of the residuals is not zero, even though their average is. But this criterion would be ineffective for detecting two residuals in opposite directions. To reduce the effect of such aberrant measurements, the algorithm is made more robust by assigning a weight to each measurement based on the magnitude of its departure from the model. This feature is essential to avoid serious bias in the calibration factors.^{6,7}

Summary

A vector demodulator can be calibrated using a statistically-based algorithm to determine the size and direction of the adjustments necessary to guarantee accuracy of the output. The uncertainties in these adjustments can be determined based on some assumptions about the measurement process:

- The ideal (I,Q) measurements of the signal in the two channels have been transformed to (x,y) measurements according to the following series of operations (listed in order of occurrence): rotation (lock error), sheer (quadrature, gain imbalance), compression, dc offsets.
- Errors in the measurements in the two channels (x,y) are uncorrelated.
- When the relative phases among the points are known, data is collected in phase-stepping mode for eight states, and the state (0, ..., 7) can be associated with each measurement. It is convenient if the number of measurements is the same in each of the eight states.
- When phase information is known only imprecisely, measurements are made that ideally lie on a circle in the I-Q plane. Input signals must have constant amplitude and varying phase (e.g., continuously, as in a two-tone setup, or discretely, as for 8PSK).

Based on these assumptions, a straightforward linear least squares algorithm can be used to determine the sizes and directions of the adjustments when relative phases of the incoming signal (such as 8PSK from the HP 8780A) are known. If this is not the case, then a linearization of the

nonlinear least squares problem can be solved, with some loss in precision of the adjustments and increased computational time. Since this approach can be used more generally to calibrate external demodulators as well, it is the algorithm currently programmed in the HP 8981A. In both cases, approximate 95% confidence intervals for the adjustments can be derived.

Examples with data collected on the HP 8981A using the HP 8780A as a reference signal illustrate the implementation of the algorithms. The examples illustrate the importance of incorporating an efficient and robust calibration algorithm using statistical principles to control the hardware. Proper modeling of the system combined with the appropriate methodology results in accurate calibration factors as well as an assessment of the uncertainty in these factors.

Acknowledgments

The author would like to thank Kaaren Marquez and Eric McHenry for implementing this algorithm in the breadboard phase, for providing data used in the example, and for their comments on earlier drafts of this paper.

References

1. *Hewlett-Packard Journal*, Vol. 38, no. 11, December 1987, pp. 25-52.
2. *ibid*, pp. 6-24.
3. G.P. Box, W.G. Hunter, and J.S. Hunter, *Statistics for Experimenters*, Wiley, 1978, Appendix 14B.
4. H.H. Ku, "Notes on the use of propagation of error formulas," *NBS Journal of Research*, Vol. 70L, no. 4, 1966, pp. 263-273.
5. J.R. Donaldson and R.B. Schnabel, "Computational experience with confidence regions and confidence intervals for nonlinear least squares," *Technometrics*, Vol. 20, no. 1, 1987, pp. 67-82.
6. W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, 1986 (especially Chapter 14).
7. R. Becker and J.M. Chambers, *S: An Interactive Environment for Data Analysis and Graphics*, Wadsworth, 1984.

An Availability and Reliability Calculation Tool for Computer Systems

This software program helps R&D engineers make trade-offs in designing for reliability. It also provides marketing representatives with a tool to evaluate the reliability and availability of customer-defined systems.

by Wulf D. Rehder

AS OFFICE NETWORKS, computer clusters, multi-processors, and other ensembles of hardware devices are integrated into large systems, the traditional problems of reliability, availability, and serviceability/supportability (the RAS features) gain a new and increased importance. The advent of complex systems that provide a variety of solutions rather than just computations has shifted the computer dependability issue from component reliability requirements to the customer-oriented concept of system availability.

A new algorithm for modeling the RAS features of large systems has been implemented in a prototype software tool for internal use within Hewlett-Packard. Before illustrating the benefits of the RAS tool, it is important to clarify the distinction between device reliability and system availability and serviceability.

Fig. 1 shows how a system might degrade through several stages during which useful work can still be done. For instance, a printer or workstation may be down temporarily, or may not be functioning optimally. In a situation like this the system is to some degree imperfect, but still usable. In most cases, a system with such operational faults is still available for proper, albeit reduced, service.

Designing for this possibility of controlled degradation and high availability in the presence of faults requires the inclusion of means to handle fault diagnosis, system reconfiguration, and operational recovery. Preferably, this is all done intrinsically by a local control or support processor or remotely from a central support facility.

Not every fault results in a failure, even though the system as a whole may contain a transient or permanent error. There exist many examples of such fault-tolerant system behavior. A well-known example of a transient error is an alpha-particle strike of a memory cell causing a one-bit error. If the system is capable of one-bit error correction, then system operation is not affected at all. As another example, deconfiguring certain defective parts, from memory and processor boards to peripherals, does result in a degraded system, but not necessarily in system failure. Similarly, a defect (bug) in a software program need not corrupt the system operation.

On the other hand, if the effect of an error escalates sufficiently to prevent the continuation of acceptable service for the user, then the system state falls below the threshold of availability, and a permanent failure occurs that requires external interference and repair.

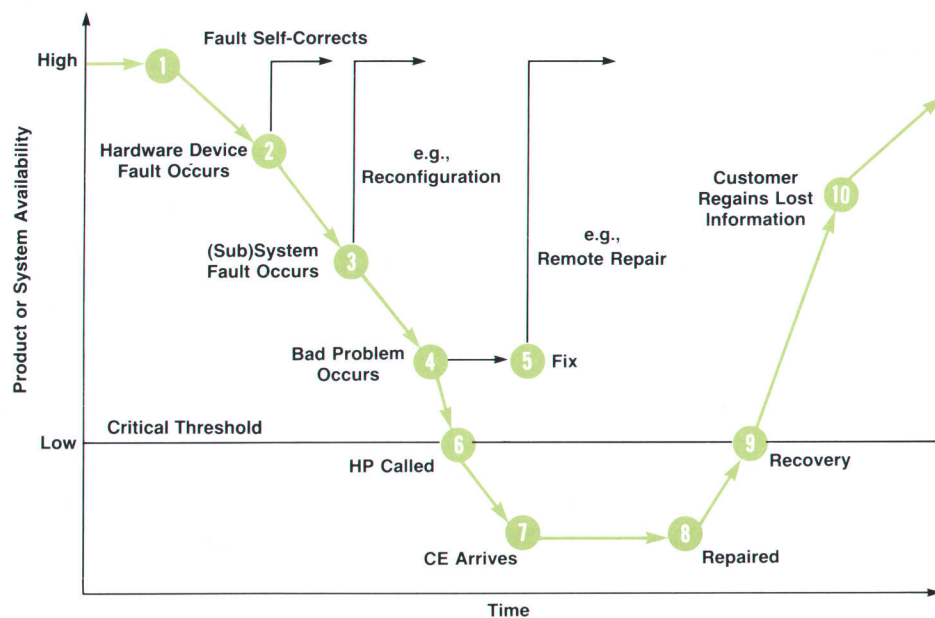


Fig. 1. States of a degrading system.

RAS Calculations and Sensitivity Analysis

For a specific system configuration, the RAS tool computes the availability A, mean time between failures MTBF, annualized failure rate AFR, mean time to repair MTTR, and expected total up and down time per year, given the failure and repair rates of the individual devices. Furthermore, the unit costs are added and the ratio AFR/K\$VAS is computed, where K\$VAS = thousands of dollars of value added shipments.

To make data entry easy and the output format clear and comprehensive, segmented windows on the screen prompt the user for input values and the softkeys are activated to browse back through the data and edit it if needed. Results of the calculation can be inspected on the screen or can be printed in the form shown in Fig. 2. Here an abbreviation such as 2(1) means that out of n = 2 devices present only k = 1 is needed for the system to be operational.

Besides providing the kind of snapshot of the major figures of system merit shown in Fig. 2, the RAS tool's flexible capabilities can be employed for the purpose of a "What if..." sensitivity analysis. It will answer questions about the impact of additional redundancy not only on system availability, but also on its complexity, its cost, and the ratio AFR/K\$VAS. Likewise, some simple trade-off decisions about where to employ redundancy and how to increase service can be made. For instance, will a 2-out-of-3 (triple modular) redundant power system design more effectively increase RAS than, say, selective dual porting?

There are many complex technical issues and questions of resources, market requirements, and competitive pressures that determine the outcome of many trade-offs, probably even more so than the purely numerical suggestions the RAS tool is able to give. But for an initial assessment of the relative impact of different configurations on the RAS features, the tool is valuable, fast, and easy to use.

Redundancy Modeling

The basic model assumption requires that the system configuration be broken up into a series of device classes made up of devices or units or components, which may be field-replaceable units (FRUs) or more complex devices like peripherals. It is in these device classes or modules that the RAS tool allows any kind of k-out-of-n:G redundancy. This means that at least k out of the n units present have to be working (to be G = good) for the specific device class to be considered up, or operational. For instance, a 3-out-of-3:G printer constellation means all 3 printers are necessary, 1-out-of-2:G processors is an assembly of two parallel (masking) processors, and a 92-out-of-118:G setup of workstations means that the device class of 118 workstations is considered available and acceptable to the user even if as many as 26 of them have failed or are off-line.

Another useful property that we believe is new in the RAS tool is the possibility of including device classes composed of several different but functionally equivalent units. The device class of 112 workstations may, for instance, be

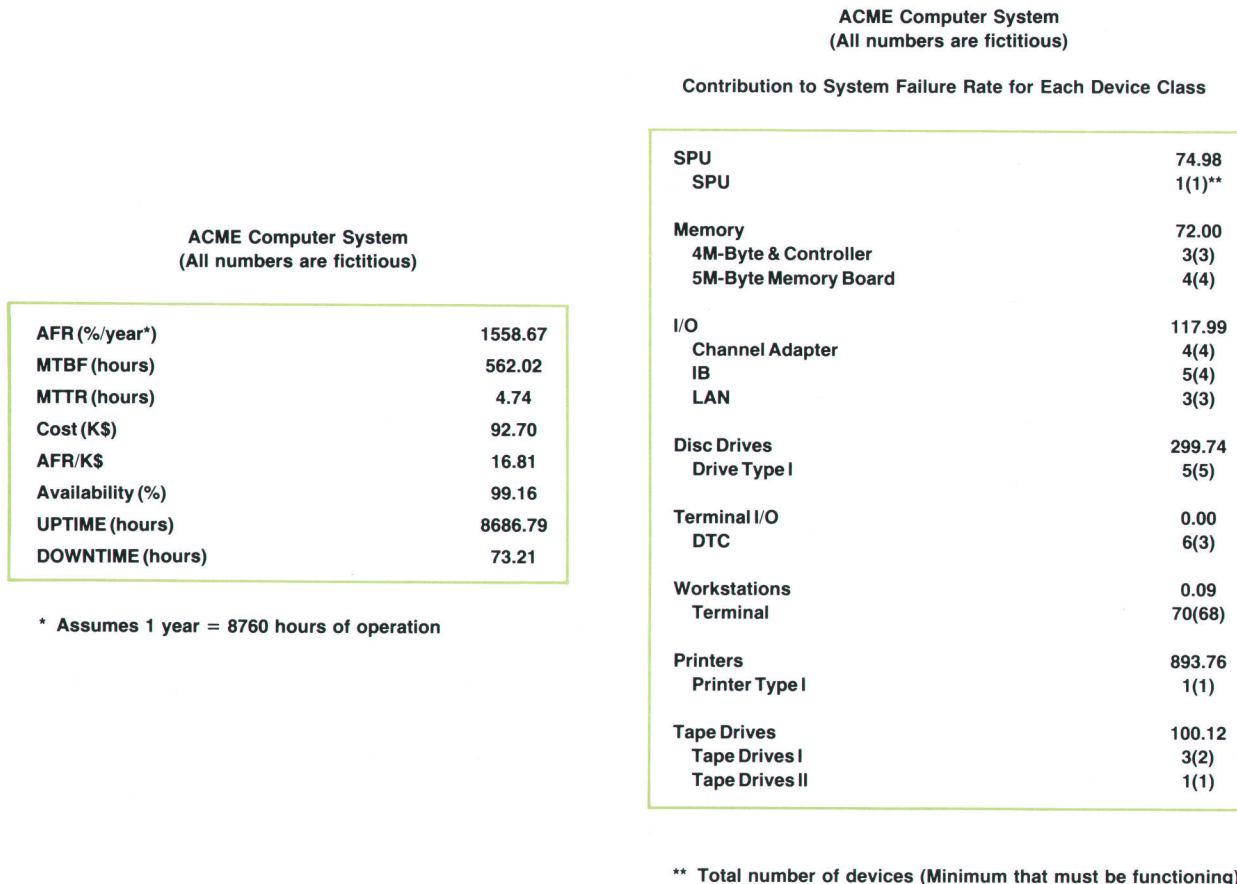


Fig. 2. Results of RAS tool calculations.

The Language of Dependability

The reliability of a system as a function of time, $R(t)$, is the probability that the system will survive throughout the entire interval $[0, t]$. Reliability as a dependability measure is most useful for systems in which repair cannot take place (space missions), or in which the whole system's operation is so critical that even short-time repair or temporary degradations are impossible, or when repair is too expensive.

The availability of a system as a function of time, $A(t)$, is the probability that the system is operational at time t . If the steady-state availability exists, that is, if the system availability stabilizes over time at a constant value A , then A expresses the average time during which the system is capable of performing useful work. Such useful work is done in a machine or system state that is acceptable to and defined by the user or customer. Availability is useful as a figure of merit primarily in systems in which service can be delayed or denied for short periods while the system functions in a degraded mode.

Fault-tolerant computing means the continued correct execution of a specific algorithm in the presence of operational faults. For this to work, the usually detrimental effects of defects must be overcome by employing hardware and/or software techniques. Typical hardware methods include redundancy, but usually software and hardware have to work together for maximal effectiveness. As a consequence, fault-tolerant computing becomes a topic of design and architecture rather than being an issue of hardware components.

Trade-Offs

While high reliability is a no-compromise business, high availability and fault-tolerant computing require delicate trade-offs. A trade-off needs to be made whenever a response variable V (for instance, customer satisfaction) depends on at least two conflicting variables, one of them pulling V up, the other pulling V down (for instance, the benefits versus the costs of high availability). Trade-offs are unavoidable design decisions for complex systems, especially at the high end. Ideally, trade-offs would optimize every important technical and market-driven response variable under the constraints and boundary conditions of resources, expected return on investment, technology, and competitive pressure.

composed of 80 HP 150 Computers and 32 HP Vectras, or any other combination. While the 150 and the Vectra have different failure rates, they are, as workstations, functionally equivalent (even if, for example, one is faster than the other).

Markov Chain Algorithm

McGrady¹ describes an algorithm to calculate the availability of a k -out-of- n : G subsystem with equal or different components. Instead of reproducing this procedure in its generality here, we give a simple example of a parallel subsystem, that is, a 1-out-of-2: G device class. Say that this is the j th subsystem (of a total of m subsystems in the system) and has component availabilities A_{ji} , $i = 1, 2$, which are computed from the respective failure rates λ_{ji} and repair rates ϕ_{ji} :

$$A_{ji} = \phi_{ji} / (\phi_{ji} + \lambda_{ji}).$$

The repair rate ϕ_{ji} is related to the corresponding mean repair time $MTTR_{ji}$ by

$$\phi_{ji} = 8760 / MTTR_{ji},$$

if a full duty cycle of 8760 operating hours per year is assumed. The subsystem availability is

$$A_j = 1 - (1 - A_{j1})(1 - A_{j2}).$$

McGrady's routine to compute such device class availabilities is based on the idea that A_j can be written as

$$A_j = A_{j1}A_{j2}[1 + (1 - A_{j1})/A_{j1} + (1 - A_{j2})/A_{j2}]. \quad (1)$$

Generalizing to more complex k -out-of- n : G subsystems leads to similar expressions. For instance, the corresponding expression for a 2-out-of-3: G device class (sometimes encountered in power systems) is:

$$A_j = A_{j1}A_{j2}A_{j3}[1 + (1 - A_{j1})/A_{j1} + (1 - A_{j2})/A_{j2} + (1 - A_{j3})/A_{j3}]. \quad (2)$$

See reference 1 for more examples.

Finally, the product of all the subsystem availabilities gives the total system availability A_s :

$$A_s = \prod_{j=1}^m A_j$$

The question is how to compute the system failure rate and system MTBF from the knowledge of the subsystem availabilities A_j and the system availability A_s . As it turns out, knowing the subsystem availabilities and their mathematical forms (e.g., equations 1 and 2 above) is sufficient. The method comes from the theory of Markov chains.

Markov chains are capable of modeling the different pos-




Fig. 3. RAS tool block diagram. The tool models degradations and redundancies at the CPU and peripheral levels, and computes reliability measures for general customer-defined configurations.

sible states (e.g., of degradation) of a system. In particular, Markov chain tools allow the calculation of steady-state (matured) failure rates λ_j for all subsystems. The steady-state failure rate λ_j for the j th subsystem can be expressed as a weighted sum involving the device availabilities A_{ji} and the repair and failure rates of the devices in the subsystem. The details are given in Ross.² For example, the failure rate of a parallel subsystem turns out to be:

$$\lambda_j = \lambda_{j1}A_{j1}(1 - A_{j2}) + \lambda_{j2}A_{j2}(1 - A_{j1}),$$

where A_{j1} and A_{j2} can be expressed in terms of the ϕ_{ji} and λ_{ji} of the devices in the subsystem.

The failure rate λ_s of a system that is a series configuration of subsystems is the sum of the subsystem values λ_j .

Finally, the system MTBF, $MTBF_s$, is given by the reciprocal of λ_s :

$$MTBF_s = 1/\lambda_s,$$

which is very familiar to engineers.

The average repair time $MTTR_s$ for the system is equal to

$$MTTR_s = MTBF_s(1 - A_s)/A_s.$$

Fig. 3 is a block diagram showing the major building blocks of the RAS tool.

Lotus Compatibility

Currently the RAS tool is installed in an HP 3000 Series 58 Computer. Future versions are being prepared to run on faster machines. A simplified algorithm for small to medium systems has been implemented in a Lotus® 1-2-3® spreadsheet format to make its use more interactive.

As an example, Fig. 4 shows how the device class failure rates change when the configuration from the example in Fig. 2 is altered as follows:

Lotus and 1-2-3 are registered trademarks of Lotus Development Corporation.

FAILURE RATES FOR DIFFERENT CONFIGURATIONS




Fig. 4. "What if..." comparison done with the Lotus version of the RAS tool.

- 2 parallel SPUs instead of 1 SPU
- 2 out of 3 4M-byte memory boards instead of 3 boards
- 3 out of 4 tape drives I instead of 2 out of 3 tape drives I

The Lotus version of the RAS tool can handle the reliability impact of such configuration changes as well as changes in component failure rates, repair times, or prices in a few minutes, giving the designer a quick evaluation of design trade-offs and providing the sales representative in the field with a hands-on tool to discuss RAS options with a customer.

A more ambitious enhancement to include stress factors (e.g., temperature) is in its investigation phase.

Acknowledgments

Many people in HP have contributed to the development of the RAS tool. The project was initiated by Brian Unter, quality manager for the Information Technology Group, and started by Doug Howell and Rick Scherer. Their first version was reviewed by Kathy Chen, who suggested the enhancements that eventually led to the Markov chain approach. Linda McCarthy did most of the software development. Finally, the continuous interest in the customer issue of improved availability by Peter Rosenblatt, general manager of the High-Performance Systems Operation, provided an additional stimulus for us.

References

1. P.W. McGrady, "The availability of a k-out-of-n:G network," *IEEE Transactions on Reliability*, Vol. R-34, no. 5, 1985, pp. 451-452.
2. S.M. Ross, *Introduction to Probability Models*, Academic Press, 1985.

Project Management Using Software Reliability Growth Models

At HP's Lake Stevens Instrument Division, the Goel-Okumoto software reliability growth model has provided a means for more accurately predicting the duration of system testing and determining when to release a new product.

by Gregory A. Kruger

SIGNIFICANT IMPROVEMENT in software reliability calls for innovative methods for developing software, determining its readiness for release, and predicting field performance. This paper focuses on three supporting strategies for improving software quality. First, there is a need for a metric or a set of metrics to help make the decision of when to release the product for customer shipments. Second, accurately estimating the duration of system testing, while not directly contributing to reliability, makes for a smoother introduction of the product to the marketplace.

Finally, achieving significant improvement is easier given the ability to predict field failure rates, or perhaps more realistically, to compare successive software products upon release. Although this third strategy will be discussed in this paper, the emphasis will be on choosing the right software reliability metric and confidently managing the testing effort with the aid of software reliability growth models. For a more thorough discussion of estimating software field failure rates, see Drake and Wolting.¹

Defects per Hour as the Quality Metric

Two principal reliability measures are:

- Defect density expressed as defects per thousand lines of noncomment source statements (KNCSS)
 - Number of new defects found per hour of system testing.
- The first measure will show progress on successive software projects, while the second measure is more helpful for real-time project management.

The user is unlikely to perceive a software product in terms of lines of code, since this is an invisible dimension of the product. The customer perceives quality in terms of how often a problem is encountered. This is analogous to a hardware product. We speak of the mean time between failures (MTBF) for repairable instruments, and in fact, failure rate and its reciprocal, MTBF, are appropriate measures of software quality. It is likely that, however different the sizes in lines of code, if two software products have the same failure rate during operation, they will be perceived as having the same level of reliability by the customer.

Although ascertaining the customer-observed failure rate is difficult at best, measuring the failure rate during system testing is straightforward. For each successive time period, a week for example, record both the engineering hours invested in testing and the number of new defects found.

Dividing the number of defects by the hours spent testing gives an estimate of the "instantaneous" new-defect-finding rate. As defects are found and removed from the software, this rate declines. The decreasing new-defect-finding rate and the corresponding increasing mean time between finding new defects provide the criteria by which progress and the ultimate conclusion of system testing can be judged.

Reliability Growth Modeling

There is considerable statistical literature on modeling the reliability growth process of finding and fixing defects in a software product. Available models can be classified into four broad categories:²

- Input-domain-based models
- Fault-seeding models
- Time-between-failures models
- Fault-count models.

Input-domain-based models can be characterized as a basic sampling theory approach. One defines the input domain of a software system to be the set of all possible input data the software will encounter during operational use. The reliability can then be estimated by taking a representative sample from the input domain and looking at the resultant failure rate when the sample data is input to the system for execution. Although the input domain sampling procedure will not be strictly followed here, the con-




Fig. 1. Typical plot of $m(t)$ and $l(t)$ for the Goel-Okumoto model.

cepts embodied in this technique turn out to be helpful in designing the testing process and understanding the limitations of making estimates of field performance based on defect-finding rates observed during system testing.

The approach in fault-seeding models is to seed a known number of defects into a software product. The software is then tested using a process that presumably has equal probability of finding a seeded or an indigenous defect. The numbers of indigenous and seeded faults found are used to estimate the system reliability. This concept has several implementation problems and draws considerable skepticism from project managers and designers alike.

The preferred data in time-between-failures models is the observed running times between successive failures. Typically it is assumed that the time between two consecutive failures follows a distribution whose parameters are dependent upon the number of defects remaining in the software during that time interval. These distribution parameters can be estimated from the observed time-between-failures data. However, the actual running time between successive failures may be a difficult measure to extract from the software testing process.

In fault-count models, the variable of interest is the number of defects observed per specified time interval. The basic idea is that the number of failures observed per time interval can be modeled according to a Poisson process. The Poisson distribution is widely used to model the number of occurrences of some event in a time or space interval. As previously noted, the defect-finding rate statistic is relatively easy to capture. For this reason, the focus of this paper will be on the category of fault-count models.

John Musa's execution-time model, discussed by Drake and Wolting,¹ is an example of a fault-count model. Although their derivation differs from Musa's, Goel and Okumoto³ propose essentially the same model. Because of its simplicity and intuitiveness and this author's preference for the derivation of the model from the underlying assumption of a Poisson process, the Goel-Okumoto model will be used here.

The Goel-Okumoto Model

The model assumes that between software changes, the number of defects observed during each hour of test will follow the Poisson distribution with a constant average, λ . Note that this defect rate includes the occurrence of both new and repeat defects. If the observed defects are fixed immediately, the Poisson parameter λ will be decreasing every hour of testing (or at least decreasing with every code correction). In this case, we have a nonhomogeneous Poisson process and can expect the defect-finding rate to decrease exponentially until software changes cease. Conversely, as the defect rate declines, the cumulative number of defects found will increase and asymptotically approach a constant. Since there is a lag in the correction of defects, the ideal state of instantaneously fixing defects when found can be approximated by counting only the new defects found each hour.

The functional form of the model is as follows. The cumulative number of defects found by time t is given by

$$m(t) = \alpha(1 - e^{-\beta t})$$

and the instantaneous new-defect-finding rate at time t is given by

$$l(t) = m'(t) = \alpha\beta e^{-\beta t}.$$

In this model, α is the expected total number of defects in the software, β is the initial defect-finding rate divided by the expected total number of defects in the software, t is the cumulative hours of system testing, and $m'(t)$ denotes the derivative of $m(t)$.

Figure 1 shows the typical shape of $l(t)$ and $m(t)$.

It should be noted that the parameter α is actually an estimate of the total number of defects the system testing process is capable of detecting. The input domain concept is useful for understanding this point. If there is some stratum of the input domain that the system testing process is not covering, the parameter α is actually an estimate of the total number of defects minus those generated by data from that specific stratum.

There are three key assumptions underlying this model:

- All new coding is completed before the start of system testing.
- Defects are removed with certainty and without introducing new defects.
- Testing methods and effort are homogeneous, that is, the last hour of testing is as intense as the first.

The first assumption ensures that the defect-finding rate follows a monotonic pattern. While the second assumption is difficult to guarantee, instituting strict security measures will minimize the rate of defect introduction. The phasing of distinct test activities, the introduction of new testing tools, the addition of new team members, and software testing burnout make meeting the third assumption difficult, as well. However, good planning of the testing process and averaging the defect-finding rate over a calendar week will tend to smooth out differences in testing effort.

Fitting the Model on a Project

Consider the application of the Goel-Okumoto model to a 90-KNCSS firmware product. This product had completed system testing eighteen months earlier, but records

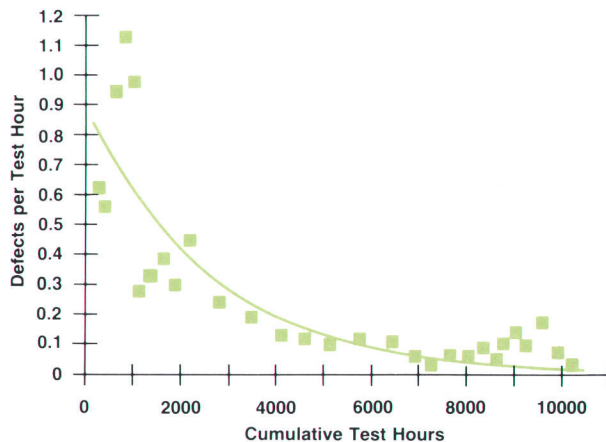


Fig. 2. Weekly defect-finding rates for project A firmware and an $l(t)$ curve fitted to the data.

were available from which a data set could be constructed to test the feasibility of the model. For each of 43 weeks of system testing, the cumulative test hours, the instantaneous new-defect-finding rate, and the cumulative number of defects found were calculated. However, because the number of hours invested in system test during the first several weeks of testing was highly variable, the first 19 weeks of data was condensed into five data points, each representing approximately 200 hours of testing. This left 29 sets of data values, each representing at least 200 consecutive hours of testing. The nonlinear least squares fit of the $l(t)$ model to this data appears in Fig. 2. The obvious departure of the actual defect-finding rate data from the model at about 9000 cumulative hours appears to be explained by the addition of several experienced designers to the testing team. This is a clear violation of the assumption that the testing effort remains homogeneous. Using the parameter estimates obtained from fitting $l(t)$, the $m(t)$ model and cumulative data are plotted in Fig. 3. Parameter estimates and standard errors from fitting $l(t)$ are:

$$\begin{array}{ll} \text{Estimates: } \hat{\alpha} = 2307 & \text{Standard Errors: } \sigma_{\alpha} = 362 \\ \hat{\beta} = 0.000384 & \sigma_{\beta} = 0.000102 \end{array}$$

There are some statistical advantages to fitting $l(t)$ to the defect-finding rate data and using the resulting parameter estimates to plot $m(t)$. By the nature of cumulative data, the residuals from fitting $m(t)$ are serially correlated. However, the residuals from fitting $l(t)$ to the instantaneous defect-finding rates are well-behaved, more closely conforming to the standard assumption of normally distributed errors with mean zero. On the other hand, $l(t)$ can only be fit if there are enough hours of testing and defects found to generate reasonable estimates of the instantaneous defect-finding rate. If a project will be undergoing a small number of test hours so that blocks of time (such as a week) cannot be aggregated to give an estimate of the defect-finding rate, then it is difficult to obtain parameter estimates by the method of least squares. In this case, maximum likelihood estimation should be used to fit the $m(t)$ function to the cumulative data. This should not be interpreted to mean that maximum likelihood estimation is an inferior technique to be employed when the least squares method is not feasible. On the contrary, it can be demonstrated that maximum likelihood estimates are among the best. The intent here is to show what is possible with simple least squares procedures.

Using the Model in Real Time

The project discussed above, project A, demonstrates that real software testing data does follow the nonhomogeneous Poisson process model. Another project, project B, an application software product consisting of 100 KNCSS, demonstrates the potential for using the defect-finding rate metric and the Goel-Okumoto model in managing the system testing process.

Setting the Release Criteria. First, a defect-finding rate objective must be incorporated into the existing set of software release requirements. The total release checklist would then include such things as all functionality complete, all performance specifications met, no known serious defects,

and the defect-finding rate criterion. The question remaining is, "What would be an appropriate value for this defect rate?" Project A concluded system testing with a rate of 0.03 defects per hour. As it turned out, the firmware proved to have extremely high reliability in the field. However, it is not clear that a 0.03 defect-finding rate is appropriate for all projects—particularly on software rather than firmware projects. In addition, arguments can be made that the relative severity of each defect found should be used to express the release criteria in terms of a weighted defect-finding rate. Ultimately, based upon work at other HP divisions and the belief that the defect weighting would be quite severe, a release requirement of 0.04 weighted defects per test hour was established. A team of four project engineers reviewed all defects found and scored each on a 1-9 scale with 1 being a cosmetic flaw and 9 being most serious. These scores were then normalized by dividing by nine. The team recognized that the software reliability modeling process would be dependent upon their ability to weight the defects consistently throughout the duration of the project.

Applying the Model. Once the release requirement is set, the Goel-Okumoto model becomes useful as a real-time management tool. Having fit the model to the most recent data, an estimate for the total test hours required to release is obtained by calculating where the curve crosses the release requirement. The management team can then review weekly plots of the defect-finding rate and estimates of the total test hours to release. These plots enable the team to predict the project's conclusion more accurately and to judge progress along the way. In addition, the engineering staff conducting system testing may find the data to be morale building. Without some measure of progress, system testing is often perceived to be a never-ending process.

Project B ultimately reached the release goal after 5760 hours of testing over 22 calendar weeks. Figs. 4 and 5 show the final plots of the $l(t)$ and $m(t)$ functions fit to Project B system testing data using nonlinear least squares. The parameter estimates and standard errors from fitting $l(t)$ are as follows:

$$\begin{array}{ll} \text{Estimates: } \hat{\alpha} = 1239 & \text{Standard Errors: } \sigma_{\alpha} = 116 \\ \hat{\beta} = 0.000484 & \sigma_{\beta} = 0.000069 \end{array}$$




Fig. 3. Cumulative defects each week for project A firmware and fitted curve.

One can simplify the analysis by first taking the natural log of the defect-finding rate data. A quick check of the $l(t)$ function verifies that taking the log transformation makes the model linear. Fig. 6 shows the regression line fit to the logged data. The benefit of operating with a linear model is that standard techniques can be used to set confidence bounds on the estimates of total system test hours needed to reach the release requirement.⁴ As can be seen from Fig. 7, the estimates of total test hours proved to be both consistent and accurate. The project concluded system testing within 300 hours of the estimate made eleven weeks previously. The accuracy of these estimates far exceeded our expectations.

Estimating Field Failures

Two different but related estimates can be made concerning the number of defects the customer will observe in the first year after release. The first is the expected number of times the user will experience a defect in the system including the recurrence of a previously observed defect. The second is the expected number of unique defects encountered.

New And Repeat Defect Rate. Including repeat defects in estimates of a software system's failure rate can be considered as taking the perspective of the customer. The customer perceives quality in terms of how often a defect is encountered. The user may not be able to tell if two or more defects are actually caused by the same fault in the code. Estimating the combined new and repeat defect rate may be possible using the $l(t)$ model fit during system testing. Recall that the defect-finding rate declines because defects are removed from the software as they are found. Once QA ends, defects are no longer being fixed. Therefore, defects should now be observed according to a Poisson process with a fixed average number of defects per hour of use. The best estimate available for this is the defect-finding rate observed at the conclusion of system testing—0.04 weighted defects per hour for Project B. The question is, how good an estimate is this for what the customer will experience?

The answer is dependent upon the difference between a customer's and an HP engineer's defect-finding effi-




Fig. 4. Weighted weekly defect-finding rates for project B software.




Fig. 5. Cumulative defects each week for project B software.

ciency. It is likely that the test engineer, rather than the customer, is the more efficient at finding defects. If an acceleration factor, K , could be identified expressing the ratio of customer hours to one hour of system testing, the customer's defect-finding rate could be estimated as $1/K$ times the defect rate at the conclusion of system testing.* The simplifying assumption here is that the customer's defect-finding efficiency is constant. Evidence to date suggests that a value of $K=50$ might be a reasonable estimate of the acceleration factor. The estimated combined new and repeat defect-finding rate would then be 0.0008 weighted defects per hour.

New-Defect-Finding Rate. Considering the number of unique defects in the software is taking the perspective of the company; we want to know if code revisions are likely to be necessary. When considering repeat failures, we stopped on the exponential failure rate curve and extended a fixed horizontal line at the defect-finding rate observed at the end of system testing. To estimate unique failures, we continue down the $l(t)$ curve for the first year after release to shipment. Conversely, the cumulative number of defects found should continue up the $m(t)$ curve.

It is a trivial matter to estimate the number of unique defects the system testing team would find in an additional

*See reference 1 for another perspective on acceleration factors.




Fig. 6. Linear model for project B after log transformation.

t hours of testing. This is simply:

$$m(T + t) - (\text{total found so far})$$

where T is the total number of test hours accumulated to date. In fact, the total number of defects remaining in the software can be estimated by

$$\alpha - (\text{total found so far}).$$

For the Project B software, $1239 - 1072 = 167$ weighted defects.

An estimate of how many of these remaining defects will be found in the field can be accomplished by defining a new $m(t)$ function for the customer. The α parameter can be set equal to the estimated number of defects remaining. The β parameter can be obtained if we are once again willing to use an acceleration factor. For the Project B software, the cumulative number of unique defects found after t hours of customer use is given by:

$$m(t) = 167(1 - e^{-0.000484t/50}).$$

For example, in 2000 hours of customer use we could expect 3.2 unique weighted defects to be found.

Field Failure Rates in Perspective. Estimates of customer defect rates and defect totals should be taken with a healthy degree of skepticism. It is a risky proposition to propose acceleration rates and attempt to estimate customer observed defect rates. Experience to date clearly shows that our customers are not finding as many defects as estimated by the model. However, these two projects have not provided sufficient insight for us to estimate confidently the relationship between model projections and customer experience. After more products are released according to this process, comparisons will be drawn between the defect rate observed in-house and the subsequent rate in the field. An empirical estimate of the acceleration factor may be obtained in this way. To date, the results from these two projects are as follows.

Recall that on Project A the analysis was conducted on




Fig. 7. Weekly updated estimates of test hours needed to release project B software.

unweighted data—all defects were counted equally regardless of severity. From the model, the number of unweighted defects estimated to be remaining in Project A is $2307 - 2167 = 140$. After 18 months in the field, there have been three customer reported defects. All of these have been resolved to the satisfaction of our customers. Three defects found out of an estimated 140 remaining suggests a ratio of estimated to actual defects of 47 to 1.

As has been stated, the estimated number of weighted defects remaining on Project B is 167. After 12 months in the field, there have been four defects identified, all of which have been resolved in a subsequent release. After applying the same weighting criteria that was used during system testing, these four defects represent 2.4 weighted defects. Since this product has not been in the field as long as Project A, we could multiply by 1.5 to give 3.6 weighted defects after 18 months. The ratio of estimated to actual defects would then be 46 to 1—amazingly close to the results from Project A but still conjecture at this point.

Conclusions

The defect-finding rate metric and the Goel-Okumoto model proved, on one project, to be a real contribution to managing the system testing phase of the software development process. Judging from the fit to historical data, the model would have been helpful on an earlier project as well. Through weekly plots of the data and the model, the management team received feedback on testing progress and information to aid resource planning. Ultimately, there was a clearer decision as to whether to release the product for customer shipments.

In cases with limited test hours, the straightforward application of the Goel-Okumoto model would be to fit the $m(t)$ model to the unweighted cumulative defect data using the method of maximum likelihood. If there are enough test hours to allow reasonable estimates of the instantaneous defect-finding rates, fitting $l(t)$ using least squares techniques is possible, although not superior to the maximum likelihood procedure. Based upon the results from these two projects, it appears that the modeling process works for both unweighted and weighted defect data. Extrapolations from the model to field conditions have proved to overestimate the number of defects customers will encounter. Further study is required to estimate confidently the relationship between model estimates and customer observed failure rates. However, we are continuing to tighten the failure rate release criterion as new products enter final system test.

An alternative approach to making the release-to-production decision would be to plot the mean time between finding new defects, $1/l(t)$. This metric could then be used to make an economic decision on when continued testing would be too costly.

Acknowledgments

I would like to thank Doug Howell, Karen Kafadar, and Tim Read for their suggestions for making the statistical portions of this paper clear. I am also indebted to Peter Piet for the long telephone conversations in which we worked through the technical details of estimating field failure rates.

References

1. H.D. Drake and D.E. Wolting, "Reliability Theory Applied to Software Testing," *Hewlett-Packard Journal*, Vol. 38, no. 4, April 1987.
2. A.L. Goel, "Software Reliability Models: Assumptions, Limitations, and Applicability," *IEEE Transactions on Software Engineering*, Vol. SE-11, No. 12, December 1985.
3. A.L. Goel and K. Okumoto, "A Time Dependent Error Detection Rate Model for Software Reliability and Other Performance Measures," *IEEE Transactions on Reliability*, Vol. R-28, pp. 206-211, 1979.
4. G.W. Snedecor and W.G. Cochran, *Statistical Methods*, Sixth Edition, Iowa State University Press, 1976.

A Reliable, Autoloading, Streaming Half-Inch Tape Drive

Designed for rack mounting, this compact tape drive cleverly channels air flow to load different-sized half-inch tape reels automatically. It also features higher performance and reliability than its predecessor.

by John W. Dong, Kraig A. Proehl, Ronald L. Abramson, Leslie G. Christie, Jr., and Douglas R. Domel

HP's NEWEST HALF-INCH TAPE DRIVE, the HP 7980A (Fig. 1), is an autoloading, reel-to-reel, horizontally mounted, streaming drive that reads and writes in two standard nine-track formats: 6250 GCR and 1600 PE.^{1,2} This maintains compatibility with previous drives and tapes, an important feature because half-inch tape is still a significant standard in the computer industry for backing up, archiving, and interchanging computer data.³

The HP 7980A provides computer system users with a reliable, low-cost, backup device for large amounts of on-line disc memory. It does this with higher performance and reliability and lower cost than its predecessor, the HP 7978B, which it replaces. It reads and writes at 125 inches per second, roughly 60 percent faster than the HP 7978B. The HP 7980A can rewind a 2400-foot tape in less than 90 seconds, reducing overall data transfer times significantly.

The HP 7980A is 40 percent more reliable, a result of the increased use of VLSI components to reduce parts counts even further than was achieved on the HP 7978B. The monthly maintenance cost of the HP 7980A is half that of the HP 7978B. Horizontally mounted in a standard-width rack cabinet, the HP 7980A is 8.75 inches high, a third the size of the HP 7978B. This saves valuable floor space and allows better use of rack cabinets.

The horizontal mounting means that the user normally cannot access the tape path to load the tape manually. Hence, an autoloading feature was designed into the HP 7980A. The operator simply places the tape reel in the door opening and closes the door. The autoloading sequence starts automatically once the door is closed, leaving the operator free to do other tasks while waiting for the tape to load in a nominal time of half a minute. There is no need for an EZ-LOAD cartridge around the tape reel, which




Fig. 1. The HP 7980A Tape Drive is a half-inch, reel-to-reel, 125-ips streaming tape drive designed for systems with disc backup requirements of greater than 400 megabytes. Like its predecessor, the HP 7978B, this high-performance drive operates with both 6250 GCR and 1600 PE standard formats. The HP 7980A can automatically load and thread any size reel ranging in diameter from six to 10.5 inches. This drive can be horizontally mounted in a 19-inch-wide rack enclosure for better floor-space utilization, and it supports IBM/ANSI-compatible formats for software distribution and data interchange between HP and non-HP systems.

is required for autoloading on the HP 7976A Tape Drive. The HP 7976A autoloads only 10.5-inch tape reels; the HP 7980A autoloads all standard half-inch tape reels from six to 10.5 inches in diameter. The earlier HP 7978A/B Tape Drive must be loaded manually and cannot autoload tapes.

In addition, the HP 7980A supports the use of 3600-foot half-inch tape, under certain guidelines, which the HP 7978B doesn't. This increases data capacity by 50 percent over standard-length tape reels.

The HP 7980A was developed and released in 40 percent less time than the previous tape drive. This was mainly a result of two factors. The first was keeping together an experienced core group of engineers from the HP 7978A/B development team to architect and design the HP 7980A drive. The second was concentrating on a core system development of the HP 7980A, that is, having a well-defined product and not adding additional features and configurations along the way. Such added features tend to prolong development cycles. This does not mean that these additional features are not eventually added, but they are worked on according to need after the core or base system is released.

Tape Path

The HP 7980A is a totally integrated tape drive, simultaneously incorporating a small form factor, an autoload feature, design-for-assembly concepts, low cost, and high reliability. The major design objective established to accomplish these goals was design simplification.

The HP 7980A has a very simple tape path as shown in Fig. 2. There are only two rolling elements: the speed sensor and the buffer arm roller. There is only one additional fixed tape guide. The oxide side of the tape contacts only the tape cleaner, the magnetic tape head, and the tape displacement unit. The tape displacement unit, located between the tape head and the tape cleaner, contacts the oxide side of the tape only during repositioning and while the tape is stopped. The tape displacement unit pushes the tape off the very smooth surfaces of the head and the tape cleaner to prevent the tape's sticking to these smooth surfaces during high temperature and humidity conditions.

The buffer arm assembly (buffer arm, spring, and roller) helps take up slack in the tape during servo starts and stops. It also establishes the tension on the tape. The buffer arm roller and fixed guide, along with the speed sensor, guide the tape in a precise manner over the head. The speed sensor measures the velocity of the tape and feeds it back to the servo system.

The half-inch tape reel is centered, seated, and locked by the supply hub. The autoload blower forces air through the door. The louvers in the door direct the air onto the tape reel to lift the end of the tape. The tape end is then carried by the air flow around the buffer arm roller and the fixed guide. It then goes over the tape cleaner and the magnetic tape head. The tape is finally sucked onto the take-up reel after passing around the speed sensor.

The drive motors are located directly beneath and are attached to the supply and take-up hubs. The blower is located between the two drive motors.

Integrated Autoload and Tape Path Design

The HP 7980A tape path is very simple. This greatly reduces costs and assists the design of the autoload mechanism. It simplifies the task of blowing the tape end off the supply reel and threading it through the tape path and onto the take-up reel.

The casting is an integral part of the tape path and autoload mechanism. The walls of the casting help determine the way the tape is blown around by the air flow created by the blower. The surface roughness of the casting is very important. The surface of the casting around the supply reel area must provide enough friction to permit the supply hub to autocenter the smaller tape reels. Conversely, the casting surface near the buffer arm assembly must be smooth enough so that the tape doesn't catch on any surface features while autoloading. The buffer arm shape is critical to autoloading and servo success. The design of the air deflector on the speed sensor is crucial to how well the tape end moves around the speed sensor roller and attaches to the take-up reel.

The autoloading success rate is greatly affected by the design of the door ramp and the door louvers. These are designed so the HP 7980A can autoloading all standard half-inch reels with varying amounts of tape on each reel.

Holes in the top cover eliminate the need for gaskets or tight tolerances on the door-to-bezel fit and the top cover-to-casting fit. They do this by eliminating air flow reversal, which can occur when autoloading air flows out the door-to-bezel and top-cover-to-casting cracks, rather than down the tape path. The ribs on the top cover fit inside the casting to reduce air leakage and allow loosening of the top-cover-to-casting fit.




Fig. 2. HP 7980A tape path components.

The speed sensor is placed in the tape path for maximum tape wrap on the roller. This prevents the tape from slipping on the roller during tape acceleration and deceleration.

The buffer arm and its spring are designed and placed so that tension variations caused by the tape coming off the supply reel at different angles relative to the arm are minimized. An additional roller could have been used to feed tape at a constant angle to the buffer arm, but was not put in for several reasons: cost, simplicity, and ease of autoloading the tape.

Air System Design. The autoloader air system was designed into the deck casting and the tape path from the beginning. Reduction of air ductwork was a major consideration. The less the air has to pass through winding passages, the lower the air pressure and volume requirements. For this reason, the tape path needs to be simple. Reducing the air pressure and volume reduces the autoloader air blower's size, noise, and cost.

The blower is placed between the two drive motors since that provides the shortest air passage. The air is sucked into the take-up reel, into the blower, and then blown out of the blower into the door area to repeat the path (see Fig. 3).

The door is used as an air passage for several reasons. No space is taken up by air ducting to blow air into the tape path, thus allowing the maximum amount of width for the card cage. In addition, using the door as an air duct allows the air to blow across the entire front of the tape reel for maximum design flexibility. The door louvers were then designed to blow air to autoloader all standard reel sizes.

The air velocity required to move the tape down the tape path was calculated from momentum and Bernoulli equations. The pressure drop was calculated in a similar manner. The blower was then selected based on these calculations and the tape path, with the blower, was measured for pressure drops and velocities. The measurements were fairly close to the theoretically calculated pressure drops and air volumes.

Autoloader Algorithm. The HP 7980A autoloader algorithms are designed to load all sizes of half-inch tape reels in a minimum amount of time, while making sure that the user's tape is treated with as much care as possible. The autoloader process is a combined effort of firmware, mechanical design, and electronic sensors. Four sensors are used by the autoloader process to monitor its progress and to indicate when an error condition is present. The first of these sensors is the door sensor, which consists of two microswitches connected in series that detect whether the front door or top cover is open. When this sensor detects a door closure, the autoloader process begins. Opening the top cover at any point halts the autoloader. Under the supply hub is the reel encoder sensor. This optical sensor detects the three reel encoder flags attached to the supply hub. These flags interrupt the sensor beam when a reel is properly seated on the supply hub. The tape-in-path sensor is an optical sensor which is positioned across the tape path preceding the head. The tape leader will interrupt this sensor beam when it enters the tape path. The final sensor is the optically encoded speed sensor. This sensor provides two quadrature pulse trains which are decoded to give position and velocity information.

Once a door closure is sensed, the autoloading operation begins. The first task is to detect whether a reel is already threaded through the tape path, a condition that is most likely to be present after a power failure. This condition is detected by the tape-in-path sensor and by turning both the supply and take-up motors in opposing directions at a low rate. If a tape is already threaded, activity will be seen on the speed sensor. In this case, the tape does not need to be automatically threaded and the servo loops can be closed.

If a tape is not already threaded, the load fan (blower) is turned on and the supply hub is slowly rotated in a counterclockwise direction to center and seat the reel. At this point, the reel encoder sensor is checked continuously for evidence that the reel is seating itself on the hub properly. When the reel is seated correctly, all three reel encoder flags will interrupt the sensor beam once per revolution. Time interval measurements are made to ensure that all three flags (not just one or two) are present, and that they have the correct relationship. If all three flags are not detected, the supply hub is shaken back and forth quickly in an attempt to get the reel to locate itself properly on the hub. A reel that will not properly seat after shaking a second time will be rejected and a MISLOAD will be reported.

The reel encoder sensor and flags are also used to regulate the speed of the supply hub during these open-loop operations. The dc motors must turn at relatively low speeds when centering a reel and feeding the tape. These speeds require relatively low voltages at the motors. Because of voltage offsets, motor constants, and temperature changes, the rotational speed could not be adequately controlled by a single voltage command. Hence, the voltage command is constantly adjusted based on the time measured between each reel encoder flag pulse. This provides rather gross, but sufficient control over the rotational speed of the motor




Fig. 3. HP 7980A tape path showing air flow for automatically loading and threading a tape.

in its open-loop mode.

After the reel is found to be seated properly on the hub, it is ready to be locked into place. First, however, the tape-in-path sensor is polled to determine if the tape is being seen continuously in the tape path. If this is the case, the reel has been loaded upside down. When the reel is spinning counterclockwise, the tape should only be momentarily flopping into the tape path. If the reel is found to be upside-down, the supply hub is turned in the clockwise direction until the tape clears the path. The door is then opened and the INVERT message is reported.

If the reel is not inverted, the hub locking routine can proceed. The supply hub motion is stopped and the hub lock solenoid is engaged. The hub is then rotated clockwise by applying a voltage ramp. As the hub rotates, three locking feet come up out of the hub and grab the tape reel, holding it securely to the hub. During this locking operation, the reel encoder flags and sensor serve as a check on whether the lock is a success. If the reel is properly locked on the hub, one of the encoder flags will interrupt the sensor beam as the voltage ramp is applied. After a successful lock, the hub lock solenoid is disengaged and the supply hub is free to spin again. This locking cycle is described in more detail later.

The supply hub is now rotated again in the counterclockwise direction in preparation for the inertia check. The inertia check is used to determine the size of the reel being loaded. This information is used to set up some autoloading and servo parameters that can be optimized according to the reel size. The inertia check is performed by applying a step voltage to the supply motor at the instant a reel encoder flag breaks the sensor beam. The time it takes for the next reel encoder flag to come around and break the sensor beam is proportional to the inertia of the reel.

As the hub continues spinning counterclockwise, some additional error conditions are checked. The check for an upside-down reel is repeated. A check for the tape leader being stuck to the reel is done. In the same manner that tape continuously in the tape path signals an inverted reel, a tape that never breaks the tape-in-path sensor beam indicates that the leader is stuck to the reel. This can often occur because of static electricity or because the leader is jammed under the reel flange. If the tape leader is indeed

stuck, the supply reel is spun counterclockwise at high speed in an attempt to free the tape end. Failure to free the tape leader at this point will abort the load process.

Next, the tape is ready to be threaded. The supply hub continues to spin in the counterclockwise direction, and the tape-in-path sensor is monitored for the tape leader. Once the tape leader is sensed in the tape path, the supply reel continues to spin for another half second (to pull the end of the tape back to the beginning of the tape path), at which point it reverses direction and starts feeding the tape down the tape path assisted by the air flow. The tape-in-path sensor is checked continuously to make sure the tape stays in the path during the feeding operation. If the sensor fails to detect the tape in the path, an error condition is flagged and the load is aborted.

As the threading proceeds, the speed sensor is monitored for activity. If the tape has been correctly fed down the tape path and has caught onto the take-up reel, the speed sensor begins to spin. The speed sensor is then used to calculate the amount of tape that is wrapping around the take-up reel. If activity is not seen at the speed sensor within a certain period, or the required number of wraps are not completed within another certain period, the tape will be pulled back out of the tape path by the supply reel and the autoloading will be retried. The HP 7980A will attempt to load the tape in this manner five times before reporting failure to do so.

After it has been determined that the tape is properly threaded, the load fan can be turned off and the servo loops closed. The take-up and supply motors are driven slowly in opposition to tension the tape. The tension arm position is monitored by an analog-to-digital converter (ADC) and, once it has reached the center (0V) position, the tension loop is closed. The tension integrator is turned on and the loop is given about a half second to stabilize. Next, the microprocessor closes the velocity loop and digitally controls its operation. The tension shutdown circuitry is enabled to prevent the drive from damaging itself or the tape during some sort of catastrophic failure. Any failure to establish tension during this process will cause the load operation to be aborted and a MISLOAD reported.

The autoloading process takes 30 seconds if all goes well. Then the drive can be put on-line and all normal reading




Fig. 4. Cross section of supply hub. The foot is rotated upward and outward from the supply hub by holding the hub fixed and rotating the sleeve underneath. This causes the dumbbell-shaped tie rod to move from a slanted position to a position more upright, thus pushing the foot up and out. Reel encoder flags pushed down by the presence of a reel are sensed by the optical sensor to verify proper reel loading.

and writing operations can be performed.

Hub Lock Mechanism

The first major objective in the autoloader process is to determine the presence of, seat, and secure reels of tape. This process is complicated by the availability of four different sizes of half-inch-tape reels ranging in diameter from six inches to 10.5 inches, the latter being the most common size. The tape reel is secured onto a supply hub which suspends the reel at the proper height, holds it secure, and rotates the reel.

When the door of the tape drive is closed, the autoloader algorithm starts immediately. The supply hub rotates in a clockwise direction and an off-center protrusion atop the hub contacts the inner race of the reel of tape. The rotational motion drags the reel in an inward spiral pattern until the reel eventually drops onto the hub's mounting plane. Three small plastic springs (the reel encoder flags mentioned earlier) mounted around the mounting plane are pushed down by the weight of the reel and trip the reel encoder optical sensor. The use of three springs assures planar contact of the reel on the hub, which helps prevent tape edge damage by keeping the reel as parallel as possible to the tape path. A fourth spring may be deflected by the write enable ring on the reel of tape to relay that status as well.

At this point, the reel is secured to the hub by means of three feet that rotate up and out of the hub to clamp the inner race of the reel. The feet are driven by a four-bar linkage that translates the rotational motion of the hub into a vertical rotation of the feet. At the end of their travel, the feet are locked into position in much the same manner as a toggle switch with an overcenter design.

The mechanism that drives the feet up and into place has four basic parts (Fig. 4). The first is the foot itself, which pivots about an axle swaged into the hub base. The foot is driven by a barbell-shaped tie rod which translates the horizontal rotation of the hub into the vertical rotation of the foot. The barbell, in turn, is driven by a sleeve-shaped part that fits around the base of the hub and is held in place by a spring and retaining ring. In normal operation, the sleeve rotates with the hub. During the reel locking stage, a solenoid engages an arm that stops the motion of the sleeve, and while the hub is rotated, causes the sleeve to rotate relative to the hub. The foot is pushed up and into place, depressing the wave spring when the foot contacts the tape reel. As the hub rotation continues, the barbell pops over center and starts to retreat. The relative motion at this point is ended by stops in the hub that limit the motion of the sleeve. The lock cycle is double-checked by the reel encoder optical sensors. Sensor placement is such that one of the plastic springs always interrupts the sensor at the moment the sleeve hits its stops. A neoprene pad embedded in the surface of the foot ensures that the reel does not slip relative to the hub. The hub is unlocked by simply reversing the process. In the case of a power failure, the entire lock or unlock cycle can be performed by manually turning the reel and engaging the arm.

Many obstacles had to be overcome to ensure a successful design. The first of these was materials selection. The material for the plastic springs must have a high endurance limit, low creep at moderately high temperatures, and a

consistent and high spring constant. The material chosen was ULTEM 1000 by General Electric. Material for the sliding parts on the hub was tailored to minimize wear and friction, and to resist deformation under load. The second obstacle was cost. We made extensive use of design for manufacturability to decrease part count and minimize assembly time. The result is a hub with no screws that is completely assembled from one side.

Integrated Tape Path

All the tape path components, except for the reel motors, are mounted onto a precision head plate which is mounted on the deck casting. The magnetic tape head and the tape cleaner are permanently mounted onto the head plate in a manner similar to that used for the earlier HP 7978A Tape Drive.⁴ This greatly reduces the number of tight tolerances required on the deck casting because of tape path requirements. The buffer arm and speed sensor assemblies are removable. This allows for easy replacement of the head plate, speed sensor, and buffer arm assemblies. Each of these assemblies is designed to be modular and interchangeable. There are no service adjustments on the HP 7980A because any adjustments required are done at the factory. No routine maintenance is required, except for regular cleaning of the tape path.

The tape path components perform several tasks in guiding the tape over the magnetic tape head. They make sure the tape is wrapped around the head properly and consistently. They also guide the tape past the head at a precise angle (skew) and at a proper height (tracking) so that the written tape is interchangeable with other tape drives. Normally, there are two precision stationary or fixed guides that perform these functions. The HP 7980A tape path incorporates one precision guide into the buffer arm assembly and the other guide is integrated with the speed sensor roller.

Speed Sensor

The speed sensor integrates the functions of an optical

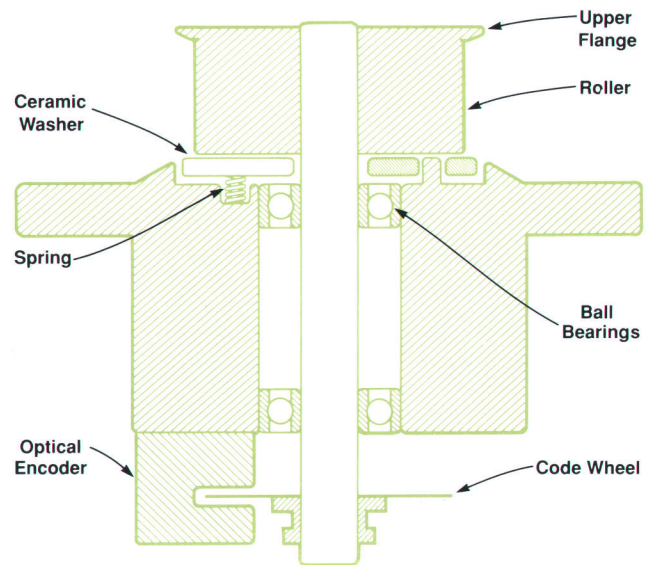


Fig. 5. Cross section of speed sensor.

encoder and a precision tape guide. The particular challenge was to create a rolling element, needed for transmitting tape speed to the encoder, that would guide the tape with the accuracy of a fixed guide.

Fundamentally, a rolling guide is simply a vertical cylinder with flanges on the top and bottom. The upper and lower flanges restrict the vertical limits of the tape's travel. According to the ANSI standard, the width of a half-inch tape must be between 0.496 and 0.500 inch. Clearly, the distance between the roller flanges must be at least 0.500 inch not to damage maximum-width tapes, but this allows a 0.496-inch-wide tape to wander 0.004 inch between flanges. However, in the simplified tape path the speed sensor must also guide the tape for proper skew and tracking across the head. This makes 0.004 inch of vertical wander unacceptable.

A traditional solution of this problem is to put a fixed guide between the roller and the head. A fixed guide is a nonrotating cylinder with a top flange set to a precise height. A spring-loaded washer pushes the bottom edge of the tape so that the top reference edge of the tape stays in contact with the top flange. This eliminates tape wander since the distance between the washer and the top flange can vary with tape width.

The HP 7980A speed sensor combines features of both rolling and fixed guides. Fig. 5 shows a cross section of the speed sensor. The roller, which includes a top flange, spins with the tape and transmits tape speed (through the shaft) to the optical encoder. The roller has no bottom flange, but there is a spring-loaded ceramic washer below the roller. The washer is restrained from rotating, but can move vertically to push the tape against the fixed height of the roller's upper flange. Thus, the roller is fixed vertically, but moves rotationally, while the washer is fixed rotationally and moves vertically.

To keep the bottom of the spinning roller from rubbing on the non-spinning washer, the roller width is slightly smaller than the minimum tape width. Thus, the semicircular arc of tape wrapped around the roller pushes the ceramic washer down and away from the roller. The spring is placed under the center of gravity of this semicircular arc of tape. This balances the forces on the washer so that it remains horizontal and does not damage the bottom edge of the tape.

Buffer Arm Assembly

The design of the buffer arm assembly presented several

special challenges. The buffer arm assembly must provide tape buffering, proper tape tension, a servo position signal, proper tape guidance, and an overtension shutdown signal. The buffer must be lightweight enough to maintain servo bandwidth, but strong enough to hold the tape height within a tenth of a millimeter. Because the HP 7980A is an autoloading drive, the buffer arm must also act as an air dam to load the tape correctly. To achieve these requirements, a thin-wall aluminum die-cast part is used. Its flexibility allows an air baffle, a spring post, a stop, and a slot for overtension shutdown to be incorporated into a single part.

During testing, it was discovered that the tape resonant frequency dropped when the drive repositioned frequently. It was determined that during such reposition cycles more air became entrapped in the tape stack. As a result, the effective tape length was increased, decreasing its spring constant. To overcome this problem, a Coulomb damper is used. A cantilever spring is placed between the buffer arm and fixed guide so that it produces approximately four inch-ounces of frictional torque. This dissipates enough energy from the system to allow servo stability without affecting other parameters.

In a typical tape drive, the tape, after leaving the buffer assembly, enters a fixed guide for proper skew alignment. To minimize space requirements, simplify service, and reduce parts count, the fixed guide is combined with the buffer assembly. The buffer arm pivots about the center of the fixed guide as can be seen in Fig. 6. The fixed guide, made of a stainless-steel ring and two ceramic washers, is bonded on the buffer base to a 0.015-mm tolerance. This allows the buffer assembly to provide all the tape guidance in the front half of the tape path, but still be removed by loosening only three screws.

Another feature is the way the position of the buffer arm is sensed. The sensing assembly had to be small enough to pass through a 40-mm-diameter hole in the head plate, but require no adjustment if the buffer is replaced. To satisfy these requirements, a small ceramic magnet is mounted on the buffer shaft. Changes in the rotating magnetic field are sensed by a linear Hall-effect IC. The plastic magnet holder allows each magnet to be bonded in its correct calibrated position. With this arrangement a signal that is linear within $\pm 8\%$ over the buffer arm range is achieved, requiring no further calibration for interchange.




Fig. 6. Cross section of buffer arm assembly.

Design for Manufacturability

The HP 7980A tape drive is designed to be built easily. The major improvement over the previous HP 7978B drive is its size. The HP 7980A is about a third the size and half the weight of the HP 7978B. This makes it much easier to build and handle.

All of the major components are attached to the main deck casting. The main deck casting is moved down the manufacturing line as major components and assemblies are attached to it. It is then flipped over to add components onto the top and tested. Most of the components and assemblies are attached to the casting in a downward movement and fastened with a common self-tapping screw.

Reduction of parts is another major improvement. The number of mechanical parts was reduced from approximately 610 on the HP 7978B to 340 on the HP 7980A. The number of unique mechanical parts was reduced from approximately 260 to 140. All these factors help substantially reduce the time to build the HP 7980A. It is, therefore, a much more manufacturable machine than its predecessor, the HP 7978B.

Acknowledgments

Thanks to Hoyle Curtis and Don DiTommaso for obtaining the necessary funding for the success of the HP 7980A and for fending off additional features and configurations so the core system could be delivered on time.

There were many difficult mechanical issues during the design of the HP 7980A Tape Drive. Thanks to Tom Bendon for managing the resolutions to these issues. The many inputs from our manufacturing team were invaluable and contributed greatly to the manufacturability of the HP 7980A. Thanks to John Meredith, Gregg Schmidtke, Mel Crane, and Lee Devlin, our manufacturing team. Our tooling engineers, Dave Halbert and Jesse Gerrard, provided much needed input into the tooling and design of the many metal and molded plastic parts in the drive. Robert "Bob" Archer and Dave Lundgren did the early product design and Jim Dow did the industrial design. Dave Jones did the rack design and the HP 7980A rack slides. Many thanks to Dan Dauner for finishing the latter stage of the product design, which is the most difficult phase, and for wrapping it up for manufacturing.

To those not specifically mentioned, thank you for your contribution to the success of the HP 7980A.

References

1. ANSI X3.54-1976 GCR.
2. ANSI X3.39-1973 PE.
3. H.L. Curtis and R.T. Turley, "Development of a High-Performance Half-Inch Tape Drive," *Hewlett-Packard Journal*, Vol. 36, no. 3, March 1985.
4. J.W. Dong, D.J. VanMaren and R.D. Emmerich, "Streaming Tape Drive Hardware Design," *ibid.*

Streaming Tape Drive Control Electronics

by Bradfred W. Culp, Douglas R. Domel, Wayne T. Gregory, Jeffery J. Kato, Gerod C. Melton, Kraig A. Proehl, David W. Ruska, Virgil K. Russon, and Peter Way

IN DEVELOPING THE HP 7980A Tape Drive, the design team leveraged its expertise gained from the development of previous streaming tape drives. Design goals included a low factory cost, a small form factor, an auto-loading capability, and a selection of interfaces. To achieve these goals, advancements were made in many key areas of tape drive design. These areas included the drive architecture, controller software, read/write electronics, tape servo system, and front panel.

Three-Box Control Architecture

The HP 7980A Tape Drive incorporates a control architecture based on the three major functional blocks of a streaming tape drive: interface, buffer, and drive. Each functional block contains a microprocessor-based controller and data path electronics to handle user data. Fig. 1 shows the three-box architecture of the HP 7980A.

Interface. The interface block handles all communications with the host computer. The HP 7980A contains an HP-IB (IEEE 488) interface. Currently, two other interfaces (PERTEC and SCSI) have been developed. Drives with one or the other of these interfaces are available as the HP 88780A. Each interface accepts tape commands from the host, transfers the user data, and returns the appropriate status. To perform a task, the interface first relays the tape command to the buffer controller. User data is then sent to or obtained from the buffer memory. Upon completion of the task, status is obtained from the buffer controller.

Buffer. The buffer block is responsible for streaming performance. The buffer controller oversees the use of the 512K-byte buffer memory. The buffer accepts tape commands from the interface and forwards them to the drive controller. The controller transfers user data between the interface, the buffer memory, and the drive's formatter. Status is obtained from the drive controller and forwarded back to the interface. Streaming performance is maintained through the use of immediate response on write operations and read-ahead on read operations. Immediate response on writes is implemented by the return of status as soon as a

full record of data is accepted into the buffer, but before it is written onto the tape. This allows the host to free its memory and begin transferring the next record. When a write error occurs on the tape, the buffer will perform the necessary write retries independently of the interface and host computer. Read-ahead on reads is implemented by issuing additional read commands to the drive so that streaming is maintained until either the buffer is full, a command other than a read is received from the interface, or a read error occurs. This frees the host and interface from the constraint of performing all command and status processing overhead within the interblock gap time. Overhead processing time can be averaged out over several records. The buffer also performs all necessary read retries independently of the interface and host computer. Incorrect data is not sent to the host.

Drive. The drive block executes all tape commands. The drive controller oversees the servo system, motor drive, tape path, formatter, read/write, and front panel. The drive controller accepts tape commands from the buffer controller and executes them. User read data is read from the tape, corrected and checked by the formatter, and then sent to the buffer memory. User write data is accepted from the buffer memory, encoded by the formatter, and then written to the tape. Read-after-write occurs on all write commands to ensure error-free writes. Status is sent to the buffer controller. In executing a tape command, the drive controller must control tape speed, tension, and position. The drive controller relies upon the servo system electronics for closed-loop control of tape speed and tension. Tape positioning is performed in conjunction with the read electronics and formatter and is based on block boundaries. The drive controller also manages user requests and information through the front panel. User requests requiring buffer and/or interface involvement are passed to these functional blocks.

Each functional block is connected by a common communication link (CCL) to an adjacent functional block. The CCL encompasses common hardware and common soft-

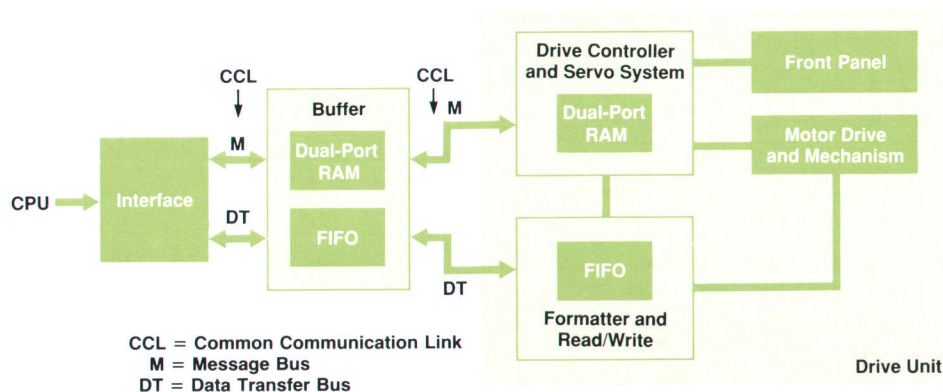


Fig. 1. HP 7980A architecture.

ware definitions. All of the optional interfaces communicate with the buffer using the same CCL. The buffer communicates with the drive using a similar CCL. The CCL software definition will be discussed later.

The CCL hardware uses a message bus for the transfer of commands and status, and a data transfer bus for the transfer of user data to and from the tape. The message bus interfaces the controller of a master functional block (closer to the host) to a dual-port RAM located on the controller of a subordinate functional block (closer to the drive). Fig. 1 shows the location of the dual-port RAM for each CCL. This dual-port RAM is directly and fully addressable by the microprocessors of both controllers. Commands and status are passed between the controllers by writing and reading specific memory locations. The data transfer bus interfaces the data path of a master functional block to FIFO (first-in, first-out) buffers located on the data path of a subordinate functional block. Separate inbound and outbound FIFOs are used to facilitate loopback testing. User data is passed along the data transfer bus in a sequential manner. Odd parity is maintained throughout with parity checking occurring as data is received across a functional boundary. An end-of-data bit is asserted to flag the last byte of each data record.

Interfaces

Important to the tape drive's acceptance by OEM customers (as the HP 88780A) is a selection of interfaces. The drive currently supports HP-IB, SCSI, and PERTEC. In the development of these interfaces, common hardware and software designs were shared. All three interfaces contain a 6809 microprocessor-based controller. About 30 percent of the circuitry is common to all three interfaces. As mentioned above, a common communication link definition is used to connect an interface to the buffer. Each interface completely isolates the rest of the drive from host, interface, and protocol specific requirements. By developing these interfaces in parallel it became clear what functionality belonged in the interface as opposed to the buffer or drive functional blocks.

The development of three sets of interface code in the time allotted for one was a major undertaking. There are 24K bytes of HP-IB code, 40K bytes of SCSI code, and 17K bytes of PERTEC code. Design time was reduced through the use of a common firmware architecture and the leveraging of common code. These interfaces share the same top-level firmware structure. About 6K bytes of code is common among the interfaces. The common code contains dual-port RAM communication routines and command processing of front-panel requests including diagnostics. With the exception of a power-on routine and a table-access module, all of the code is written in high-level languages. The use of high-level languages increased design efficiency and code readability. Overhead processing time became an issue with the PERTEC interface. Instead of rewriting the code in assembly language, overhead time was removed by rearranging the order of tasks during reads and writes. An increase in parallelism between interface and buffer was obtained, saving about two milliseconds per command.

The three-box architecture gives the HP 7980A and HP 88780A designers flexibility. Not only can new interfaces

be easily designed and supported, but also options and enhancements can be added easily. For example, a new buffer design can be integrated without affecting any of the interfaces or the drive. Parallel development, debug, and test of the functional boxes increased design efficiency by minimizing the impact that problems in any one area had on the entire project. Because these boxes are functional entities, they could be turned on and debugged separately. Separate turn-on and debug included functional testing as well as the testing of communications with neighboring blocks. The integration of the first complete HP 7980A produced a working drive within a month.

Common Communication Link Software

The common communication link (CCL) is a single interface definition for two separate interface links. The HP 7980A contains three processors, a host interface controller, a buffer controller, and a device controller, performing specific tasks. The main function of the buffer controller is to intercept read and write commands and buffer the data within its data buffer to provide fast data access. For operations not involving data buffering or command queuing, the buffer controller typically passes the operation (command) directly through. By having common interface links on both the interface and the drive, the complexity of executing unbuffered commands is greatly reduced.

The CCL is designed to have a common command set for the half-inch tape format. Although the HP-IB interface is the primary host interface for the HP 7980A, the CCL command set does not preclude support of other interfaces such as SCSI, PERTEC, or a future high-performance interface. The buffer controller and drive controller implement a superset of commands that represents the operations needed to support most half-inch tape interface protocols.

CCL protocol is implemented by a 1K-byte dual-port RAM. This dual-port RAM is accessible by both the master (closer to the host CPU) and the subordinate (closer to the drive). Access is controlled by message control locations which synchronize the passage of messages across the CCL.

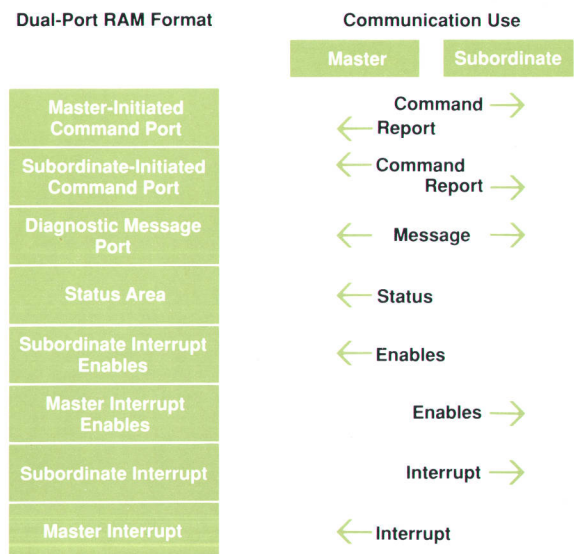


Fig. 2. Dual-port RAM use.

The dual-port RAM is divided into two command ports, a diagnostic message port, a status area, two interrupt controls, and two interrupt locations (see Fig. 2). The command ports support bidirectional command transfers. A master-initiated command port passes commands from the master to the subordinate, while the subordinate-initiated command port passes commands from the subordinate to the master. The master-initiated command port typically executes commands generated by the host, while the subordinate-initiated command port typically executes commands generated from the front panel.

A command port contains one of two message types: a command sent from the initiator or the report returned by the target. The contents of the command port are determined by a message control location.

The diagnostic message area is used during certain diagnostic operations when data messages larger than the command ports (64 bytes) need to be passed. A diagnostic message can be up to 512 bytes long, followed by a two-byte checksum.

The status area is used to maintain information on the drive that changes infrequently. This area contains information such as tape loaded, write protected tape, and the density of the current tape.

The interrupt control locations are used to specify areas of the dual-port RAM that require that an interrupt be sent notifying the other controller of any change in their contents. Interrupts are most often requested on changes to the command port message controls (indicating that a command is ready or a report has been returned) and changes to the status area. The use of programmable interrupts gives more freedom for each controller's implementation. The interrupt locations allow the master and subordinate to interrupt each other independently. The interrupt location is written with a value that indicates the area of dual-port RAM that has been changed.

The CCL commands are divided into six major opcode ranges: write, read, space, general, configuration, and diagnostic commands. Each command range is further subdivided into the individual commands supported within that group. For example, the write command group consists of write data record, write file mark, write gap, and write density ID. The opcode ranges simplify command decoding. Parameters on CCL commands are dependent on the opcode and vary in size from none to 57 bytes. For most opcodes, the number of parameter bytes is fixed. For example, the write record request opcode contains one three-byte parameter specifying the maximum length of the data to be sent, the write tape mark opcode contains no parameters, and the write density ID contains a one-byte parameter specifying the density to be written.

The report type returned on a CCL command depends on the command sent. All tape commands return a normal report containing status flags, tape position flags, a byte count field, a retry count, and an error code. The status flags quickly tell whether the command was successful, successful with error recovery, unsuccessful, or rejected. The tape position flags indicate the current position with regard to beginning of tape (BOT), end of tape (EOT), and end of file (EOF). The byte count contains the length of a write record or read record. The retry count indicates

whether physical retries were needed during the operation, and the error code contains detail on a command failure or command rejection.

Non-tape commands return reports specific to their operation. For example, the diagnostic test commands return a diagnostic report that can contain the test number that failed, the failure code, and up to two field-replaceable units (FRUs) that could have caused the failure.

High-Performance Buffer Software

The HP 7980A is a high-performance streaming tape drive. It is also very flexible in its ability to be adapted to different environments and to use performance features in a variety of ways. High performance and flexibility of use in the HP 7980A are achieved by the features provided within the data buffer.

A streaming tape drive mechanism such as the drive controller requires a constant stream of data if it is to maintain high performance. In a typical host-drive interaction, irregular or slow data rates can interrupt the constant stream of data and cause the tape to stop and perform a costly repositioning cycle. The data buffer addresses the data rate problems by providing an asynchronous link between the interface and drive, allowing each to operate independently at its own data and command rates for limited time intervals. The HP 7980A has a large buffer in which to hold data as data rates are matched, making the asynchronous link possible. The buffer holds 512K bytes of data, enough not only to match variations in data rates, but also to allow the data buffer subsystem to accept data throughout a repositioning cycle at an average rate of 400 kilobytes/second.

The data buffer packs variable-sized records into its buffer and places an entry in a queue indicating the location and size of the corresponding record within the data buffer. The queue is large enough to allow the buffer to be completely filled without reaching queue limitations except on small records.

The buffer maintains two independent processes, one that interacts with the drive and another that interacts with the interface. Any transactions the host computer makes with the HP 7980A Tape Drive are made with the interface side of the data buffer. The host transaction is complete when data has passed between the host and the buffer and necessary reporting and status information has passed. The operations necessary to write or read the data from the tape occur independently of the host interactions.

During write operations data passes into the buffer from the host. As soon as data is in the buffer, the command is reported on with an immediate response report indicating that the operation is complete. The queue contains the size of each record. Data held in the buffer is then written to the tape automatically when a sufficient amount of data has been accumulated or when an appropriate time-out has occurred. The drive is streamed as long as data is available in the buffer.

When a read command is received from the host, the buffer goes into a read-ahead mode during which the drive side of the buffer controller issues continuous reads to the drive controller. Data passes into the buffer and reports are placed into the queue. This time they contain not only

record size, but also information regarding the success of the read attempt. Data is taken from the buffer and reports are taken from the queue to satisfy all read commands from the host. As data is removed by the interface, the drive tries to keep the buffer full, managing tape motion as necessary.

The queuing/buffering algorithm implemented in the HP 7980A allows a great deal of flexibility with the contents of the buffer. When data is read from the buffer, pointers are moved forward in both the data buffer and queue. Until the information is physically overwritten, it is possible to recover the information by moving the pointers back again. This capability is built into the read-ahead functions of the buffer controller. If a backspace command is received while the drive has been performing read-aheads, the buffer first attempts to recover the data by moving pointers back. This electronic backspace, as it is termed, gives a great performance advantage in that it eliminates the need to reverse the tape physically to position back one record.

Increased Buffer Flexibility. The data buffer has other features that provide greater flexibility in the use of its performance capabilities. Two of these features are greater access to buffer functions and the ability to change key performance parameters. Flexibility is also obtained through the use of manual buffer operations and through configurability.

Buffer operations described to this point involve two sides of the buffer. The drive side performs either delayed or anticipatory actions automatically to provide quick response to the host. Data is disposed of according to conventions in streaming conditions.

Manual commands offered by the buffer provide the same capabilities that occur automatically during a read or write operation, with each operation broken down into suboperations. During manual operations, only the specified suboperations are performed, with no automatic actions taken and with data disposition specified, not assumed. For example, a single write command is broken down into write-to-buffer, write-to-tape, and remove-record-from-buffer suboperations. Manual buffer operations are much more cumbersome to use and are not optimized for high-performance streaming, but offer flexibility in certain key areas such as diagnostics and in recovering buffered data.

Diagnostics make heavy use of manual commands. Data loopback operations are performed by writing to the buffer and then reading data back to the interface. The drive provides wellness and error rate tests using manual commands to write and read tapes locally. Diagnostic access to manual commands also allows access to buffer and tape operations from the front panel.

Recovering buffered data also uses manual commands. When a hard error occurs, or if the tape is stopped at EOT with unwritten data, the data in the buffer can be recovered by issuing read-from-buffer commands. Data recovery is not a typical operation, but is possible using manual commands. The flexibility of manual commands also allows current and future interfaces to implement interface specific operations.

The performance of the drive is dependent not only on the large buffer, but also on how the drive is set up to use the buffer. For example, parameters such as the trip point at which the drive starts writing data out of the buffer to the tape, the maximum record size to be written, or the

length of time writes are held can all affect performance.

The HP 7980A is a highly configurable drive. Most variables that can affect performance and many that determine the personality of the drive can be configured and maintained within nonvolatile RAM on the buffer board. Configurations are used by the drive controller, the buffer controller, and the interface controller. A section of configurations is maintained and defined differently for each interface.

The configuration values are distributed to the appropriate subsystems when the drive is powered up. The interface also maintains default ROM configuration values. In the event that battery backup power is removed from the nonvolatile memory, default values are obtained from the interface.

Through configurations, a drive can be set up with a distinct personality and set of performance features for a particular user. The configuration system also maintains a set of locks which can be used to lock the configuration to the current value and prevent it from being changed.

Concurrent State Machines. The data buffer holds a special position within the drive in that it maintains CCL communications with four ports, three of them entirely independent. Each port has its own characteristics and needs.

The downstream interface port is the source of all host commands. Commands received here must be processed quickly since all streaming operations into and out of the buffer occur here. Streaming commands also have the characteristic that when one is received, the next command will probably be the same.

All streaming commands pass to the drive through the downstream device port. Quick response at this port is critical. If a command is not sent to the drive when a report is received, the drive controller may be forced to stop streaming and reposition the tape.

The upstream drive and interface ports are not independent since there is no queuing of upstream commands. Commands received from the drive are either fully or partially processed by the buffer controller. If necessary, they are then passed up to the interface controller. Commands received from the drive are irregular and are typically initiated as a result of human interactions. Hence, processing speed only needs to be fast enough for human response.

Resets can also be received as asynchronous events from either the drive or the interface. Resets are infrequent and, like upstream commands, do not require fast processing response like downstream commands.

The processing needs of the different ports led to an architecture within the buffer controller that allows maximum independence and specialization for the processing of each of the ports. Each major input source is handled by a separate program which controls the source and can be tailored to meet the streaming or background needs of each port. An operating system provides a set of concurrent processes for the programs to run within. Each program runs independently within its own time slice, with time apportioned for critical needs as they arise.

The programs are implemented as state machines and the operating system as a concurrent state machine driver. State machines have the characteristic that they perform an action when a particular event occurs, then proceed to the next state and wait for another event to occur. State

machines can be controlled effectively by an operating system. The operating system performs state transitions and calls to state execution modules. It also allocates time appropriately between the different processes during the waiting periods. Within the time slice given to each process, each of the potential events for the current state can be checked. If an event has occurred, the associated action for that event is also taken. Actions are coded as straight-line operations without loops. At the conclusion of an action, processing returns to the operating system for state transitions and for other processes.

The interface state machine handles all communication between the downstream interface port and the queue and buffer hardware. The drive state machine handles downstream communication between queue and drive port, including all automatic retries of tape operations. The subordinate state machine handles upstream commands in both the interface and drive ports. The reset state machine handles all resets regardless of location.

The operating system has no watchdog timer nor the associated timing and communication race conditions associated with unexpectedly losing the processor to another time slice. Instead, the operating system allows each process to complete its action before passing control to the next process. Time for the heavy needs of the interface state machine and the drive state machine is automatically allocated by allowing the processes to take the time that they need. Processes such as the subordinate state machine and the reset state machine are effectively turned off by having only a single event to check on. They have no com-

plexity until the event occurs.

Each of the processes is independent, but each can have communication with other processes and affect their operation by passing messages or commands. All communication between processes occurs as events in the receiving process caused by actions in the sending process. Communication locations are tightly controlled data structures, the queue being a prime example. The receiving process looks at the communication locations for conditions that trigger events in themselves and that require action to be taken.

The buffer controller functions are divided into individual programs, which are independent, have their own specialization and definition of tasks, and have localized access to hardware, ports, and internal communication locations. This provides the structure needed to tackle the complex tasks of the buffer controller. Each of the pieces has defined functions and boundaries for good process definition and implementation.

Integrated Read/Write System

The design objective for the HP 7980A read/write system was to lower cost while improving performance over previous tape products. This objective was met through the reduction of printed circuit board space, the real-time optimization of read gains, the calibration of each read channel, and the simplification of the write electronics.

The HP 7980A read/write board replaces the equivalent functions of earlier read, write, and write formatter boards. This reduction in printed circuit board space was achieved

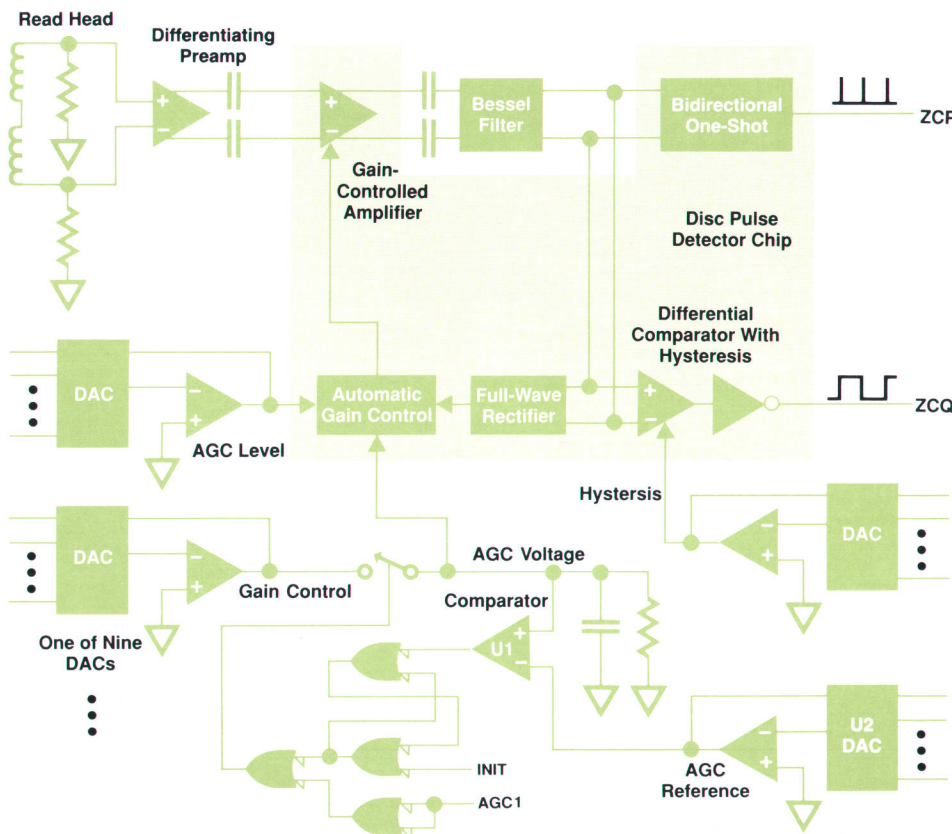


Fig. 3. Analog read channel.

by using surface-mount components, combining the preamplification and differentiation functions into one stage, using an integrated disc pulse detector chip, and developing a semicustom phase-locked loop chip.

The read channel shown in Fig. 3 consists of a differentiating preamplifier, a gain-controlled amplifier, an automatic gain control circuit, a gap clamp circuit, a zero-cross pulse circuit, a zero-cross qualifier circuit, a calibration circuit, and a fourth-order Bessel filter. The input to each channel is the output of a single-track read coil. Two signals, a zero-cross pulse ZCP and a zero-cross qualifier ZCQ, make up the output of each read channel. The zero-cross pulse is generated each time the differentiated input signal passes through zero. The zero-cross qualifier goes to a logical zero when the differentiated input signal becomes more negative than the lower hysteresis level and goes to a logical one when the differentiated input signal becomes more positive than the upper hysteresis level. This information is used by the phase-locked loop to recreate the digital data as it was recorded.

Differentiating Preamplifier. The preamplifier is an NE592 differential video amplifier configured as a differentiator. The 200-kHz (PE format) gain is set for 32 dB with a phase shift of 88 degrees, and the 555-kHz (GCR format) gain is set for 40 dB with a phase shift of 86 degrees.

A typical read channel has a preamplifier stage followed by an AGC amplifier and differentiating amplifier. To reduce board space and parts count, the preamplifier and differentiating amplifier are combined. The layout of the differentiating preamplifier is very critical to achieving high common-mode rejection and retaining a high signal-to-noise ratio. Fig. 4 shows the difference between the constant gain and the differentiating configurations for the NE592. In the differentiating configuration, any leakage capacitance from the junction of capacitor C1 and resistor R1 results in the two emitters seeing different impedances to ground. This unbalances the differential front end of the NE592, resulting in a decrease in common-mode rejection. The board layout techniques shown in Fig. 5 are used to reduce the noise level and retain the common-mode rejection of the NE592. The read and write sections have board-level isolation of 5V, ± 12 V, and ground. Each of the read-channel NE592 front ends has individual paths to the bottom ground plane, individual lines for power, and minimum trace runs near the differentiating preamplifiers. The head signals from the J2 connector to the NE592 inputs alternate between layers 2 and 3 which are located between ground plane layers 1 and 4.

Disc Pulse Detector Chip. The preamplifier is ac coupled to the AGC amplifier, which drives a differential input, fourth-order, passive Bessel filter. The output of this filter drives the AGC, zero-cross pulse, and zero-cross qualifier circuits. The AGC circuit measures the input signal and compares it against the AGC level (an external dc voltage). This voltage difference controls the gain of the AGC amplifier to make the peak-to-peak differential voltage output of the Bessel filter equal to four times that of the AGC level. The differentiating preamplifier turns peaks into zero-cross data so that a comparator can be used to locate the read signal peaks. If the read signal exhibits a tendency to return to the baseline between peaks, the comparator

could respond to noise near the baseline. To avoid this problem, a zero-cross qualifier circuit is used. This circuit is a comparator with externally controlled hysteresis. The differential signal must go above or below the dc hysteresis value before the comparator will switch states. The comparator output is fed to the phase-locked loop as zero-cross qualified data. The output of a bidirectional one-shot is fed to the phase-locked loop as zero-cross data that indicates the start of a data window. The phase-locked loop uses the zero-cross data to generate a pulse that is half a bit window long. This pulse is used to clock the zero-cross qualified data into a latch. The output of this latch represents the data as it was put on the tape.

Automatic Calibration. During track activity (data records), each track independently adjusts its gain via the AGC voltage to an optimum amplitude. The AGC voltage required by each channel while reading data is determined using comparator U1 and DAC U2 as shown in Fig. 3. During gap conditions (no track activity), the AGC voltage is overridden by individual gain control DACs. This allows each channel to be preset to a gain that produces an output signal of optimum amplitude if a nominal input signal is present. Detection of data by the phase-locked loop signals the channel to enter the AGC mode until another gap is detected or the gain increases above a predetermined value. DAC U2 and comparator U1 in Fig. 3 are used to determine when the gain is above the predetermined value. During writes the gain is also held constant so that low-amplitude read signals cause a write retry.

Track-to-track variation in the AGC voltage is dependent on the head and read channel combination and is independent of tape-to-tape variation. Track-to-track variation (gain profile) from a high-output tape to a low-output tape can be closely modeled using a single gain offset value across all tracks. Gain profiles of GCR and PE densities differ, and therefore require separate gain profiles for each density. Automatic calibration is performed at the beginning of each tape when the tape density is being identified or being written. This calibration in GCR involves all nine tracks

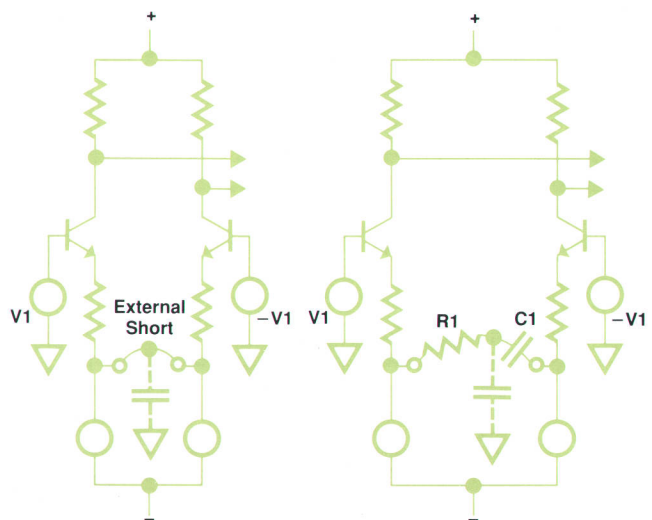


Fig. 4. Constant gain (left) and dV/dt (right) configurations for NE592 preamplifier.

during the ARA burst portion of the GCR density ID, while in PE it involves only the PE density identification track. Reference DAC U2 in Fig. 3 is stepped upwards until comparator U1 indicates the reference voltage is equal to the AGC voltage for this track. This procedure is repeated two more times and the median voltage for each track is saved as one of the raw gain voltage levels for the current tape. The average gain of the raw gain voltage levels is calculated and saved as one of the average gain voltage levels for the current tape. A gain profile calibration is performed using the tape drive diagnostics. An all-ones pattern at the appropriate density is written for 100 feet beginning at BOT. The pattern is read and the gain voltage for each track is determined 256 times. The average gain voltage for each track is the gain profile for the specific density. Each time the average gain voltage level is calculated during a tape ID routine, the gain profile for the specific density is normalized such that its average value across all nine tracks matches the average gain voltage level for the current tape.

Standard Cell Phase-Locked Loop. When we decided to do a follow-on design to the HP 7978A Tape Drive, our experience with its design afforded us a great opportunity to make technological advancements in the areas of price, performance, and reliability. With this in mind, we selected areas of the HP 7978A design where we could get the best return using some sort of LSI integration. One area was in

the clock recovery section of the HP 7978A.

Designing a multidensity phase-locked loop for a nine-track tape drive presents a problem when one is trying to minimize cost and maximize performance. However, with this as a design criterion, it was evident that some sort of integration was needed. What we decided on was a digital/analog hybrid phase-locked loop design. This gave us the flexibility to set important third-order loop parameters using normal analog loop parts, that is, an op amp, filter, and VCO. We then wanted to use LSI technology to integrate all the digital logic associated with the recovery scheme. We chose HP's own CMOS standard cell process whose density allowed us to put three tracks of phase-locked loop logic into one chip, and thus use only three chips per system.

The major elements of the standard cell are shown in Fig. 6. Two programmable digital delay lines are used in the multidensity design, one for qualification of incoming read channel flux data, the other for calculating the phase error to be used by the phase detector. A pattern detector is used to detect the presence of data to control the phase-locked loop as well as the automatic gain control circuitry of the read channel. A two-state phase detector that can switch between frequency-lock and phase-lock modes virtually eliminates the chances of a harmonic lockup. An output encoder is used to synchronize all the channel out-

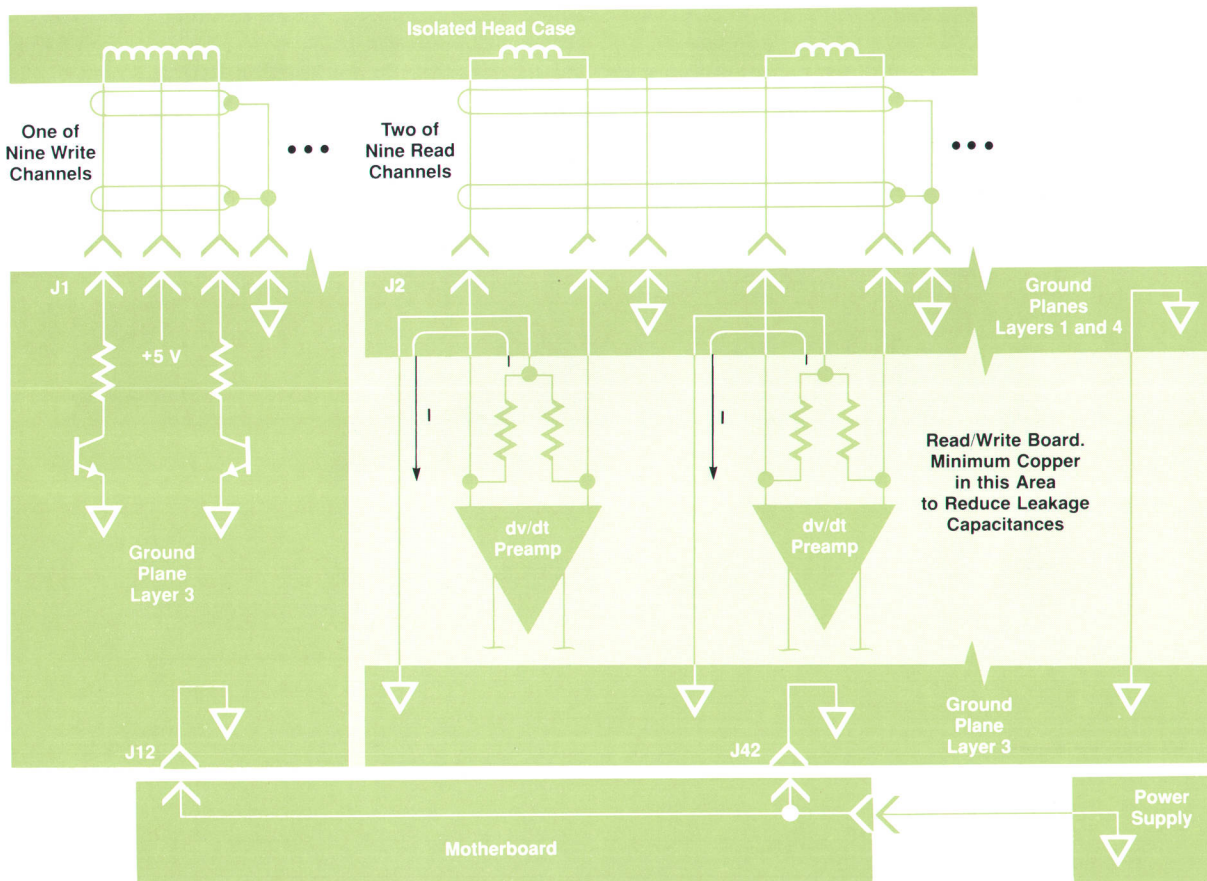


Fig. 5. The layout of the ground planes for the NE592 front ends is critical for reducing noise and maintaining common mode rejection.

puts with a common system clock. The bus interface section contains four control registers for full microprocessor controllability. A center-frequency generator provides a reference when there is an absence of data.

The standard cell is implemented using HP's CMOS process. The implementation and testing of parts took only six engineer-months. Our first silicon version was production quality. Because the chip was flexible and usable the first time, we designed it into the HP 7978B as a cost-reduction action before the HP 7980A was introduced. The cost savings achieved in the HP 7978B was enough to pay for the chip's development cost. The gains achieved in the HP 7980A over the original HP 7978A phase-locked loop include a reduction in printed circuit board area by a factor of 10, an improvement in reliability by a factor of 8, a reduction in factory cost by a factor of 5, and an increase in performance by a factor of 2.

Digitally Controlled Write. Each of the nine data signals from the write formatter is turned into two write currents, each driving one half of the write coil for one track. The positive half of the write current is controlled by the signals GCR, PE, and COMP while the negative half is controlled by the signals GCR*, PE*, and COMP*. These six signals are generated for each channel by a PROM which is driven by common address signals comp, bp/g*, and dack. The other two address lines, data-1 and data-2, are channel dependent. data-2 lags data-1 by one bit window time period, allowing the PROM to know the value of the preceding data bit. For data equal to one, GCR is asserted for the full bit window time. PE is asserted for 100 percent of the bit window time in PE mode but for only 16 percent of the bit window time in GCR mode. If the preceding bit was a zero, COMP is asserted for 80 percent of the bit window time in GCR mode or for 90 percent of the bit window time in PE mode. For data equal to zero the control signals GCR* and PE* are asserted instead of the signals GCR and PE. If the preceding bit was a one, COMP* will be asserted.

Each track has two groups of three control signals. These signals control a group of three open-collector drivers. The open-collector drivers are connected to one end of a write coil through current-limiting resistors. The center tap of each write coil is connected to 5V. The half-coil current is then the sum of the three currents generated by each open-collector driver pulling its resistor to ground. This technique allows the write currents to be generated using 5V TTL logic.

To guarantee the write and erase currents will turn on only when commanded to, a POWER GOOD signal is hard-wired to a protection circuit. The POWER GOOD signal is a signal from the power supply which will not be valid until voltages are at rated values. The protection circuit turns on an enhancement mode n-channel MOSFET which supplies five volts to the center taps of the nine write coils.

The same type of circuit is used in the erase circuit to turn on an npn transistor which provides a ground for the erase circuit. The dc erase current must ramp to the steady-state current in less than 0.2 ms and stay at this value until turned off. The erase current is controlled by two circuits. The two circuits operate in parallel until the steady-state current level is reached. The ramp circuit is then turned off, leaving the steady-state current circuit on until the erase circuit is turned off.

Hybrid Servo System

The HP 7980A servo system loads and unloads the tape automatically, controls the tape velocity and tension, and performs diagnostic functions. The drive controller uses the servo system in conjunction with the read electronics and formatter to perform accurate control of tape position. The servo moves the motors independently during the tape loading process under the control of a 6809 microprocessor. Sensors monitor the progress of the tape and check for error conditions. Tension is established and the velocity and tension servos are enabled. The tape is under closed-




Fig. 6. Block diagram of standard cell for phase-locked loop chip.

loop control thereafter. The tape velocity is primarily controlled by the take-up motor, and the tension is controlled only by the supply motor. A block diagram of the servo system is shown in Fig. 7.

Closed-Loop Operation. Tape velocity is controlled by an integral plus proportional control scheme. This yields a zero steady-state velocity error. The long-term tape velocity error is well under one percent. The controller is implemented as a digital-analog hybrid to take advantage of the best features of each type of control. In the HP 7980A, a design requirement for the servo is that the 6809 microprocessor be able to manage all the drive control functions and also control the velocity with minimal involvement. For this reason, the error signal and integral portion of the loop are performed by the processor, and the proportional feedback and multiplication by coefficients are performed by analog circuits. An optical position encoder is used by the processor to measure linear tape velocity. This is done by counting the number of pulses of the encoder every 5 ms. The measured velocity is compared with a velocity command at each sample period and accumulated over time. The resulting value is sent to a 12-bit DAC. The processor has complete control over the velocity of the tape by changing the value of the velocity command. This scheme performs the integral portion of the control with high accuracy. The processor involvement is minimal because only simple fixed-point additions and subtractions are required. The proportional part of the control is done in analog form. The optical encoder drives a tachometer circuit which gives a bipolar voltage proportional to tape velocity. The tachometer circuit does not have to be very accurate since it is only used to generate a proportional feedback term. The velocity control accuracy is determined

by the integral term computed by the processor. The integral and proportional terms are summed with appropriate scaling factors by an operational amplifier and the resulting signal is then sent to the motor amplifier.

The output voltage of the velocity controller drives the take-up and supply motors. This reduces buffer arm deflections during velocity ramps. If the velocity and tension servos were completely decoupled, the velocity servo would start a ramp with the take-up motor, and the tension arm would then deflect and start a ramp of the supply motor. This would require large arm deflections. In the HP 7980A the velocity control voltage is scaled and sent to both motors. The buffer arm then only needs to deflect a small amount to correct the tape velocity at the supply motor.

The variable-velocity rewind is an important feature of the HP 7980A tape drive because it allows a reduction of rewind time for a 2400-foot reel from 120 seconds for a constant-velocity drive to 90 seconds. The maximum rewind velocity attainable by a tape drive primarily depends on the maximum voltage available to drive the motors. Each motor generates a back emf proportional to its angular velocity. This back emf counteracts the supply voltage and limits the maximum angular velocity the motor can achieve.

The tape radius of this tape drive can vary over a 2:1 ratio from beginning of the tape to its end. At each end one or the other motor can limit the maximum rewind velocity attainable to about 250 ips. In the middle of the rewind, the tape radius for both motors is roughly equal, and the tape can attain velocities of over 400 ips. In the HP 7980A, the processor maximizes the velocity of the drive during rewind by reading the voltages sent to the motors with an ADC, comparing them against the supply




Fig. 7. Block diagram of servo controller.

voltage, and adjusting the servo velocity command several times a second to use the available supply voltage most effectively. Both servo loops operate at all times; only the velocity command is varied during the rewind.

The tension loop measures the deflection of the buffer arm and corrects the velocity of the supply motor until the arm returns to its nominal position. This effectively matches the linear tape velocities at both reels as the tape winds from one reel to the other. If both reels were identical, no buffer arm deflection would be required. The buffer arm applies a relatively constant force of 10 ounces to the tape and thus passively controls the tension. The arm force is constant as the arm travels ± 1 inch. The velocity and tension loops are approximately decoupled since each motor sees a relatively constant load of 10 ounces.

The tension loop is implemented completely in analog form. The buffer arm uses a linear Hall-effect sensor and a magnet to convert position to a voltage. This voltage is processed by an analog compensator and combined with the velocity control voltage to drive the supply motor. The key feature of the compensator is that it has an integrating term. This always returns the arm to its null point after a ramp. It also reduces arm deflections during ramps and eliminates final velocity-dependent arm deflections. If the integrating term were not present, the high rewind velocities of the HP 7980A would not be attainable.

Open-Loop Operation and Tape Sensors. The autoloading feature of the HP 7980A requires individual control of the two motors until the tape is loaded. The velocity and tension control loops are each followed by a gain/load control stage. In the gain mode, this stage can output a voltage proportional to an 8-bit number. The processor can perform the autoloading operations in this mode. When the loops are closed, these stages are switched to the gain mode. In this mode they operate as programmable gain stages. The processor can adjust the velocity and tension loop gains to account for different reel sizes. The gain/load stage is implemented with an 8-bit multiplying DAC. In the gain mode, the voltages generated by the servo loops are fed to the reference inputs of the DAC, and the processor adjusts the loop gain by writing to the DAC. In the load mode, the reference voltage is fixed, and the processor controls the motor drive voltage by writing to the DAC.

Several sensors monitor the progress of the tape during the autoloading process. A three-line encoder on the supply hub has several functions. It determines whether the feet on the supply hub have properly engaged the reel, it determines the presence of the reel, it is used to control the velocity of the supply hub during the tape threading operation, and it is used by the processor to determine the reel inertia. The three-line encoder is also used with the velocity encoder during closed-loop operation to determine the supply reel radius instantaneously.

The beginning-of-tape (BOT) and end-of-tape (EOT) signals are determined on half-inch tapes by reflective markers. The markers are detected by the tape drive using two reflective sensor assemblies. Some tapes are almost as reflective as the markers, and this has been a problem for marker detection circuits. The HP 7980A employs a circuit that responds to changes in reflectivity, rather than absolute levels, to detect the markers. This means that the circuit

is insensitive to component variations and can handle a wide variety of tapes without adjustments. The circuit is shown in Fig. 8. The sensor consists of an LED and phototransistor pair. The phototransistor receives infrared light generated by the LED which bounces off the tape. As a marker passes the sensor, the phototransistor emits a current pulse. This current is converted to a voltage V_1 by amplifier A1. V_3 is an attenuated version of V_1 . V_2 is a filtered version of V_1 . A2 is an amplifier used as a comparator. A sufficiently large pulse at V_1 will cause a pulse at V_0 since V_2 does not rise as rapidly as V_3 . The output does not respond to slow variations in tape reflectivity or extraneous light sources, but it does respond to the rapid change across a reflective marker.

A tape-in-path sensor determines when tape has entered the path. The sensor is a phototransistor and LED pair that shines infrared light across the tape entry point near the supply reel. The sensor is designed without adjustments by using a pulsed technique for detection. The detection circuit is similar to the circuit of Fig. 8. In this case the LED is pulsed by the processor for a short time with high current. If there is tape in the path, no output is detected at V_0 . Otherwise, a fast change is detected at V_0 .

An eight-port ADC is used to perform diagnostics on the analog circuits. It measures the buffer arm position during load and signals the processor to close the servo loops. It measures the motor voltages during rewind to adapt the rewind velocity. It is also used during self-tests to perform a motor-drive loopback, a tachometer loopback, and a 48V supply voltage check.

Tape Positioning Control. Unlike a stop/start drive, streaming tape drives do not have the ability to stop and then restart the tape within the distance of an interblock gap. Each time the tape is stopped, a streaming tape drive must reverse the tape until it is well ahead of the next record to be read. The tape can then be ramped up to speed and the read electronics activated upon reaching the interblock gap where stopping was initiated. While stopping and starting the tape the read head may traverse an unknown number of records. Unfortunately, the 6250 GCR and 1600 PE tape formats do not support a standard mechanism for identifying individual records. Therefore, tape repositioning must be performed based on physical distance along the tape. The HP 7980A Tape Drive uses an optical shaft encoder mounted to the capstan wheel in the tape path. This en-

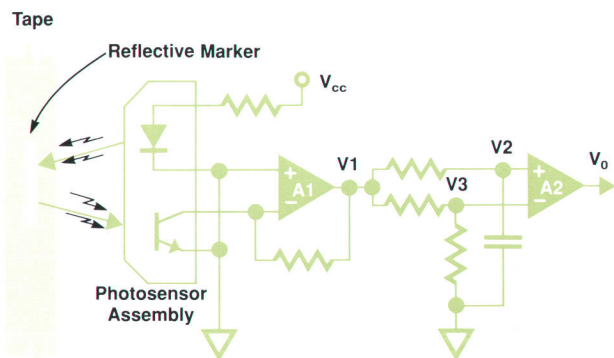


Fig. 8. Beginning and end-of-tape sensor circuit.

coder provides two quadrature pulse trains. A quadrature decoder chip (QDC) converts these pulses into a relative tape positioning count with the use of an up/down counter.

Because of the high density and small gap size of the GCR recording format, it is important that relative tape positions supplied by the QDC be as accurate as possible. When a block-to-gap or gap-to-block boundary is sensed by the read circuitry, it is essential that the position obtained by the QDC at that instant be captured and stored away for later use. Although the QDC is a bus-compatible device, time delays from the occurrence of a block boundary to the point at which the microprocessor senses that event and reads the position can introduce unacceptable errors. To avoid the latencies caused by the additional microprocessor cycles, position capture circuitry was designed around the QDC. This circuitry monitors inputs from the read formatter which signal when a block boundary has been crossed. When the event occurs, external hardware extracts the position information from the QDC and stores it in the appropriate location of a 4×12 -bit register file. The position is captured within $2.4 \mu\text{s}$ of the event. This translates into a maximum position error of only 0.0003 inch.

Upon stopping the tape, the last valid block-to-gap position is used to mark the gap that follows the last block read. The tape is stopped, then restarted in reverse. Monitoring the up/down counter, the drive controller waits until the read head is in the target interblock gap. The read electronics is then activated and the tape position at the next gap-to-block boundary is latched. This position will be the same as the block-to-gap position previously latched plus any positioning error encountered in stopping and starting of the tape. The tape is then brought to a stop and restarted again in the forward direction upon receiving a new read command. Again, using the up/down counter to monitor tape position, the drive controller activates the read electronics in the interblock gap before the next record to be read. By recapturing the block-to-gap position in reverse as a gap-to-block position, the drive controller prevents the accumulation of tape positioning error. Each time a position is recaptured, the amount of positioning error is calculated by subtracting the two tape position counts. If this error becomes excessive it is reported to the host, thus indicating a problem with the drive.

Vacuum-Fluorescent Front Panel

We had two basic objectives in designing the HP 7980A's front-panel display. The first was to provide an HP-quality solution, especially since the display is a focal point for the user of the drive. We analyzed a number of factors, ranging from the angle of the display to its color and clarity. The second objective was to provide a convenient means of displaying tape drive status that would be viewable from across the length of an average computer room. Because of the long-distance viewing requirement, special attention was placed upon character size, color, and contrast, and upon the physical placement of annunciators. After a thorough analysis of the options available to us, we converged on a design employing a vacuum-fluorescent display.

Display Operation. In simple terms, the vacuum-fluorescent display (VFD) is a directly heated triode vacuum tube.

The VFD consists of an electron-emitting cathode filament, a screened phosphor anode, and a grid (located between the anode and cathode) to control electron flow. Therefore, the VFD is structurally similar to old radio or television vacuum tubes, yet it performs a totally different function. Besides the fact that VFDs are optical displays and not amplifiers, VFDs are often constructed with multiple triode structures included within the same evacuated tube. This feature allows the implementation of cost-effective, complex displays. Examples of these inexpensive yet complex VFDs can be found in all areas of modern consumer electronics.

VFD segments are illuminated when electrons emitted by the heated cathode (filament) are accelerated by the positively biased anode and are allowed by the positively biased grid, or mesh, to collide with the screened phosphor. This causes photons to be emitted from the phosphor at a frequency that depends on the type of phosphor used. This feature allows the design of multicolored displays by screening different phosphor patterns on the same display. Although this multicolored feature is attractive, it is not totally free. Standard available phosphors emit photons at luminous efficiencies that can vary by as much as 50 percent. This causes a brightness balance problem between different phosphors on the same display. Brightness is typically balanced using optical filters, different power supplies, or different duty cycles.

The cathode filament, a thin oxide-coated tungsten wire, emits thermal electrons when heated to around 600°C by a filament power supply. The filament power supply is a major part of the power consumed by a VFD.

The VFD display used in the HP 7980A (Fig. 9) uses an ac filament supply providing 2.8Vrms and 0.47W at 715 Hz, a 30Vdc, 0.6W power supply which provides power for the anode and grid drivers, and a 40-pin IC which provides level shifting (0-5V to 0-30V) and digital control for 18 anodes and seven grids.

Why a Vacuum-Fluorescent Display? There are five reasons for choosing a vacuum-fluorescent display—cost, reliability, versatility, market separation, and human factors. The VFD is very cost-effective compared to other options available. One other display technology, the liquid-crystal dis-




Fig. 9. Vacuum-fluorescent panel of HP 7980A.

play, was slightly less expensive, but caused concern regarding low-light viewability and electrostatic discharge sensitivity during assembly.

Reliability is probably the VFD's greatest asset. With a glass-encased body and an evacuated interior, it is impervious to electrostatic discharge damage and can endure harsh environments. Its low power consumption relative to LED displays also allows enhanced reliability.

Since the VFD's light output is multicolored, it can be easily filtered to almost any color desired with standard colored lenses. In the HP 7980A Tape Drive, the blue-green output is filtered to almost white without significantly affecting the amber light output.

Many other tape drives have LED displays and annunciators. By using a more consumer-oriented display (the VFD is standard fare on VCRs), we were able to achieve a more contemporary look.

The contrast of the VFD's bright phosphor with the dark filtered background gives excellent viewability from up to twenty feet while not being too bright at close range.

Acknowledgments

The authors would like to acknowledge the efforts and contributions given to the HP 7980A by the rest of the design team. The team included Dave Ellis, Joel Larner, Elen Hunt, Joe McCarthy, John McCarthy, Alfred Nat-rasevski, John Scherr, Dave Van Maren, and Sally Wang. We would also like to thank our managers Don DiTommaso, Mark Gembarowski, and Mike Tremblay.

Authors

June 1988

6 Setting Specifications

Sherry L. Read



Sherry Read is production planning manager at HP's Santa Clara Division. She joined HP's Spokane Division in 1983 with a bachelor's degree in biometry from Cornell University (1979) and a master's degree in statistics from Colorado State University (1983). After she relocated to Santa Clara, she added an MBA degree from California State University at San Jose (1987). Before joining HP, she worked as a statistician for the Forintek Canada Corporation, a nonprofit forest products research laboratory. She's a member of the American Statistical Association.

Timothy R. C. Read



Tim Read first joined HP as a statistician for manufacturing at HP's Stanford Park Division and, in the ensuing four years, became manager of statistical engineering with responsibility for statistical consulting in R&D, manufacturing, and marketing, as well as quality information systems. Tim was born in Ashbourne, England, but earned his bachelor's degree in mathematics (with honors, 1978) and PhD degree in statistics (1982) at the Flinders University of South Australia. Before joining HP, Tim held visiting assistant professorships of statistics at the University of Wisconsin and Colorado State University. Tim has authored and coauthored 18 papers and a book on various aspects of the theory and application of statistics.

12 Statistical Circuit Design

Karen Kafadar



Karen Kafadar has contributed many professional papers to statistics journals and to the Encyclopedia of Statistical Sciences. As a statistician for the Stanford Park Division, she was closely involved in the development of both the HP 8981A Vector Modulation Analyzer and the HP 86792A Agile Upconverter. Her other product development projects since she joined HP in 1983 include a peak power meter and the HP 8770A Arbitrary Waveform Synthesizer. Karen was born in Evergreen Park, Illinois. After earning her BS and MS degrees in mathematics and statistics from Stanford University in 1975 and her PhD from Princeton, she first accepted a position as assistant professor at Oregon State University, but soon moved to the National Bureau of Standards as a mathematical statistician. Her special interests include exploratory data analysis, robust methods, and spectrum analysis. For recreation, she enjoys swimming and bicycling.

Lynn M. Plouse



In the two years before she joined HP's Stanford Park Division, Lynn Plouse served internships in both marketing and R&D at the Signal Analysis Division, where she worked on microwave mixer modeling, a field that provided the subject of her master's thesis at the University of California at Davis. After earning her MSEE degree in 1983, she joined a small team of Stanford Park circuit designers who shared in developing the HP 86792A Agile Upconverter. A native of Pasco, Washington, Lynn is married, teaches an analog circuits course at Cogswell Polytechnical College, and spends her leisure hours running, water- and snowskiing, camping, or golfing.

18 Statistical Calibration

Karen Kafadar

Author's biography appears elsewhere in this section.

26 Availability Tool

Wulf D. Rehder



Wulf Rehder, who describes his place of origin as "a tiny village in Northern Germany," pursued studies at universities on two continents: Hamburg, Freiburg, Berkeley, and Berlin, where he earned his PhD in 1978. Ancient languages, mathematics, and physics are among his subjects of study, and he has held various teaching positions, most recently as professor of mathematics and computer sciences at California State University at San Jose. He joined HP's System Technology Division as a statistician in 1986, working on computer system availability modeling and, more recently, on subsystem

simulation in R&D. He is a prolific writer and has published some 35 papers on mathematics, statistics, philosophy, and linguistics. He's working on his third book. Wulf's hobbies include the study of the middle ages, especially the 11th century. He also specializes in early 19th-century literature.

30 Software Project Management


Gregory A. Kruger



In the productivity section of HP's Lake Stevens Instrument Division, statistician Greg Kruger's main responsibilities include software metrics, software reliability modeling, and hardware metrics. When he first joined HP at the Loveland Instrument Division almost seven years ago, Greg's assignment included implementing statistical quality control practices in manufacturing and training people to use and understand them. Later he was commissioned to spread Total Quality Control practices throughout the Lake Stevens Instrument Division. Greg was born in Waterloo, Iowa. His BS in mathematics/statistics (1979) and MS in statistics (1981) are both from Iowa State University. He is married, has two children, and offers much of his leisure time to his duties as deacon of his church. Greg is an avid archer, serves on the board of directors of the Washington State Bowhunters, and edits a newsletter on the subject. Vocal music is another of his interests.

36 Autoloading Tape Drive

Douglas R. Domel



Doug Domel was born in Georgetown, Texas and studied mechanical engineering at Texas A&M University (BSME 1979) and Colorado State University (MSME 1981). He joined HP in 1981 and contributed to the mechanical design of the HP 7978A, 7978B, and 7980A Tape Drives. His work on the HP 7980A has resulted in a patent application for the autoloading hub. His professional interests include solid mechanics and computer simulations. He lives in Fort Collins, Colorado, is married, has two children and enjoys hunting, fishing, and camping.

Kraig A. Proehl




Kraig Proehl began his career at HP's Waltham Division while working on an electrical engineering degree. He received a BSEE from Purdue University in 1983 and joined HP full time that same year, working as an R&D and production engineer on the HP 7974 Tape Drive. He contributed to the autoloading and streaming design for the HP 7980A Tape Drive. Kraig was born in Trotwood, Ohio, and currently resides in Loveland, Colorado. He is married and has a daughter. His hobbies include woodworking, sports, and photography.

Ronald L. Abramson




Ron Abramson was born in Spokane, Washington, and after a tour of duty with the U.S. Army in Europe, he studied mechanical engineering at the University of Washington (BSME 1984). He joined HP in 1984 and worked on the HP 7978B Tape Drive and the speed sensor and product design for the HP 7980A Tape Drive. Ron lives in Boise, Idaho, is married, and has one daughter. Outside work he enjoys tennis, cycling, and reading.

John W. Dong



John Dong was born in Phoenix, Arizona, and studied mechanical engineering at the Massachusetts Institute of Technology (BSME 1975 and MSME 1976). With HP since 1976, he has served as a production engineer, as a product designer for the HP 9876 Thermal Printer, and as a mechanical design engineer for the HP 7978A and 7980A Tape Drives. His work on the HP 7980A has resulted in two patent applications for the door assembly and the integrated tape path. John lives in Greeley, Colorado, and enjoys the usual Colorado outdoor activities of hiking, skiing, camping, and sailing. He is engaged to be married this summer.


Leslie G. Christie Jr.



Leslie Christie was born in Fayetteville, Arkansas, and attended the University of Arkansas where he studied mechanical engineering (BSME 1970) and engineering science (MSES 1979). Before joining HP, he taught mechanical engineering courses at Kansas State University and was a design engineer for a company that made fishing tackle. He came to HP in 1984, contributed to the HP 7978A Tape Drive, and worked on the buffer assembly for the HP 7980A Tape Drive. His work on the HP 7978A Tape Drive resulted in a patent for an anti-adhesion apparatus for magnetic-tape drives. Leslie lives in Greeley, Colorado, is married, and has two children. His free time is mainly taken up with family activities.

43 — Tape Drive Electronics

Bradford W. Culp



Brad Culp completed work for his BSEE degree from Virginia Polytechnic and State University in 1978, and with an HP fellowship earned his MSEE from the University of Illinois at Champaign-Urbana in 1984. He joined HP in 1978 and worked on the control-

ler design for a color inkjet printer. He also designed the controller for the HP 7978A and 7980A Tape Drives. Brad was born in Ogden, Utah, and now lives in Greeley, Colorado. His interests include bridge, chess, skiing, hiking, softball, and volleyball.


Douglas R. Domel

Author's biography appears elsewhere in this section.

Kraig A. Proehl

Author's biography appears elsewhere in this section.

Wayne Thomas Gregory



Tom Gregory joined HP's Civil Engineering Division in 1979, after receiving a BSEE degree from Arizona State University, and worked as the production engineer for the HP 3805, 3808, and 3810 distance and angle meters. After transferring to Greeley Division R&D, he worked on the HP 7978A Tape Drive and was responsible for the power and display electronics for the HP 7980A Tape Drive. He was born in Oceanside, California, and now lives with his wife in Fort Collins, Colorado. His outside interests include hiking, skiing, and table tennis.

Gerod C. Melton




Gerod Melton joined HP in 1979 with a background in radar and navigation flight test instrumentation. He worked first as a production engineer at the Loveland Instrument Division, and then became an R&D engineer at the Greeley Division, where he was responsible for the design of the read/write electronics for the HP 7980A Tape Drive. He studied electrical engineering at Texas Tech University (BSEE 1969). Gerod lives in Loveland, Colorado, is married, and has one daughter. His interests include hiking, cross-country skiing, photography, and traveling with his family.

Peter Way



Currently a product manager with HP's Greeley Storage Division, Peter Way contributed to the servo design and testing of the HP 7978A, 9144A, and 7980A Tape Drives. He joined HP in 1979 after receiving an MS degree in aeronautics and astronautics from the Massachusetts Institute of Technology. He also holds a BSE degree in aerospace and mechanical engineering from Princeton University (1977). Peter was born in Caracas, Venezuela, is married, and lives in Fort Collins, Colorado. He is active in local church and civic activities and enjoys Hobie Cat racing, kayaking, and motorcycling.

Jeffery J. Kato




Jeff Kato was born in Havre, Montana and studied electrical engineering at Montana State University (BSEE 1982). He joined HP in 1982 and worked on the read channel electronics for the HP 7978B Tape Drive and the phase-locked loop chip for the HP 7980A Tape Drive. Jeff is married and enjoys skiing, golfing, basketball, and softball. He is also active in the United Way Partner's Organization.

Virgil K. Russon



Born in Lehi, Utah, Virgil Russon studied mechanical engineering at Brigham Young University (BSME 1979 and MSME 1979). He joined HP in 1981 and worked on the firmware for the HP 7978A, 7978B, and 7980A Tape Drives. Virgil lives in Greeley, Colorado with his wife and five children. He is a scoutmaster and active in his church. He enjoys commercial and fine art and is currently landscaping his two-acre lot.

David W. Ruska



Born in Detroit, Michigan, David Ruska received a BSEE from Michigan Technological University in 1982 and joined HP the same year. He developed the firmware for the HP-IB interface of the HP 7978A Tape Drive and developed the software architecture and firmware for the drive and front panel of the HP 7980A Tape Drive. David currently lives in Greeley, Colorado. His interests include raising Arabian horses, Bible study, playing keyboards, and arranging and performing Messianic music.

58 — Digitizing Oscilloscope

Joe K. Millard

Author's biography appears elsewhere in this section.

59 — 1-Gigasample/Second ADC


John J. Corcoran



When John came to HP in 1969, he joined an R&D team in the Santa Clara Division doing laser interferometer work. He moved to HP Laboratories in 1972, where he is now a project manager. He has been involved in integrated-circuit design for analog-to-digital converters and communications circuits, using bipolar, NMOS, CCD, and GaAs technologies. John holds a BSEE degree from the University of Iowa and an MSEE degree from Stanford. He is a member of the IEEE and serves on the program

committee of the International Solid-State Circuits Conference. He lives in Portola Valley, California, and enjoys skiing, tennis, and winemaking.

Ken Poulton



Ken Poulton's primary function at HP Laboratories is to design and test analog-to-digital converter components and IC process test vehicles. He came to HP just after receiving his MSEE and BS in physics from Stanford in 1980. During his school years, he

worked at Ampex on SAW devices. Ken is a member of the IEEE and has published a paper describing the 1-GHz ADC system featured in this issue of the HP Journal in the Journal of Solid-State Circuits and the Proceedings of the ISSCC. Circuits for analog/digital conversion and CAD tools and methodologies are Ken's foremost professional interests. He was born in Oakland, California, and now lives with his wife and one-year-old son in Palo Alto. His favorite leisure activities are skiing and rock climbing.

Knud L. Knudsen




In the 25 years since Knud Knudsen joined the HP Laboratories, he has been working on a variety of product developments, among them the HP 1980 Oscilloscope Measurement System, the HP 970A Digital Multimeter, and the HP 8450 UV-Visible Spec-

trophotometer. More recently, Knud has been responsible for the development of the silicon digitizer chip described in this issue of the HP Journal. The chip is used in the HP 54201, HP 54111D, and HP 54112 Digital Oscilloscopes, and in an optional board for the HP 16500A Logic Analyzer. Knud was born in Skjern, Denmark, and his degree in electronics engineering (1957) is from the Aarhus Elektroteknikum there. Before moving to the United States, he did research work on radio receivers at the Technical University in Copenhagen. Knud is named as an inventor or coinventor on several patents for transistor-noise-measuring circuitry and a particle analyzer. He is a member of the IEEE, and his technical interests focus on signal-to-noise optimization in measurement circuits, optical detection schemes, and solid-state circuits ranging from CCD sensors to high-speed bipolar circuits. Knud is married and has two grown children. His avocations include playing the piano, making wine, tennis, skiing, and soccer.

67 — Signal Conditioning


Lewis R. Dove



Lewis Dove came to the HP Colorado Springs Technology Center in 1982, after receiving his BSEE degree from the University of Arizona. Currently, he attends his alma mater for a master's degree in science and electrical engineering on an HP resi-

dent fellowship. In his first two years at HP, Lewis was an R&D product engineer working on thick-film hybrid components for the HP 54100 and HP 54201 Digitizing Oscilloscopes. In 1985, he became a marketing and applications engineer. Hybrid techniques and electronic packaging continue to be his focal interests. Lewis was born in Tucson, Arizona, where he still lives when he's attending school; when he works at the Technology Center, he lives in Colorado Springs. He is married. A pilot in his off-hours, he also enjoys outdoor activities like hiking, camping, running, and skiing.


Mark E. Mathews



Before joining HP's Colorado Springs Technology Center in 1984, Mark Mathews had been involved in research on infrared materials, NMOS process development, and test engineering. At HP, he handles both production and R&D engineering for a

variety of hybrid products. He earned his BS degree in physics at the University of California at Riverside (1971) and his MSEE degree at California State Polytechnic University at Pomona (1976). A paper he wrote about PbTe/ZrO MIS structures for infrared sensing appeared in the Journal of Applied Physics. Mark was born in Hamilton, Ontario; he is married and has a daughter. In his off-hours, he plays the guitar for the Colorado Springs Jazz Ensemble.


Thomas K. Bohley



Thomas Bohley's BSEE and MSEE degrees are both from the University of Missouri. Since joining HP's Colorado Springs Division in 1966, he has served as engineer on the HP 1300 Series X-Y Displays and as project leader on HP 1701, 1707, and 1980 Oscillo-

scopes and the HP 1332A Display. Tom's primary interest is analog circuit design, and his work resulted in three patents for trigger circuits. He recently accepted a position with a medical laser company. Born in St. Louis, Missouri, he presently lives with his wife in Colorado Springs. As an avid horseman, Tom likes foxhunting and competes in dressage and combined training, but is also interested in skiing, hiking, and scuba diving.

Joe K. Millard




For the past four years, Joe Millard has been project manager for the HP 54111D Digital Oscilloscope at HP's Colorado Springs Division, which he joined in 1972. His earlier development projects focused on IC design for the HP 1740 Series of Oscillo-

scopes. He's now a member of the technical staff at the Colorado IC Division. Born in Athens,

Tennessee, Joe earned his PhD degree in electrical engineering from the University of Tennessee in 1970; his MSEE and BSEE are from the same institution. Before he came to HP, he worked on nuclear instrumentation at the Oak Ridge National Laboratories. His work has provided the subjects for some 11 papers. In off-hours, Joe enjoys hiking and four-wheeling the back roads in the Colorado mountains. He is married and has two sons.


David W. Bigelow



David Bigelow's BSEE degree is from the University of Minnesota (1967). After receiving his PhD degree from Stanford in 1972, he joined HP's Microwave Technology Center in Palo Alto, California, where his responsibilities included thin-film process develop-

ment. He later moved to the Colorado Springs Technology Center, developing integrated-circuit processes and serving as IC process engineering manager and as IC production manager. Presently, he is working as a circuit design engineer. A member of Sigma Xi, Dave was born in Rochester, Minnesota. He's married, has three children, and is active in his church in Colorado Springs. He enjoys hiking to and from work—"except for the rattlesnakes," as he emphatically points out.

Donald D. Skarke




The product design of oscilloscopes and digitizing oscilloscopes, including the HP 180, 1726, 1980, and 54111D, along with CRTs and accessories, has been Don Skarke's main professional interest since he joined HP in 1965. His work on the HP 180 re-

sulted in a patent for a plug-in grounding system. Don was born in Garwood, Texas, and received his BSEE degree from the University of Texas. He served as a sergeant in the U.S. Air Force from 1952 to 1956. Don has three children and three grandchildren and lives in Colorado Springs, Colorado. In his spare time, he enjoys building live steam models, bicycling and building bicycles, and collecting watches.

70 — Digital Filtering

B. Allen Montijo



Allen Montijo came to the HP Colorado Springs Division in 1981, after having worked there the previous summer as a student. He received his BSEE degree from Princeton University in 1981 and his MSEE degree in 1985 from Stanford University. A design engineer

in the R&D laboratory, Allen's professional interests include time-domain digital signal processing and the interaction of subassemblies in system design. He was born in Denver, Colorado, is married, and has a daughter. Allen enjoys photography, skiing, and hiking.

A One-Gigasample-per-Second Digitizing Oscilloscope

This instrument's high sampling rate makes it particularly useful for analyzing high-speed, one-shot occurrences. A blend of state-of-the-art designs was required to achieve this performance.

by Joe K. Millard

OVER THE PAST FOUR YEARS, the demand for digitizing oscilloscopes has grown at a rate of about 30 percent per year. For CRT-based analog oscilloscopes, this rate has dropped to about six percent. The growing preference for digitizing oscilloscopes is driven by several key advantages of their architecture. First, continuous digitization of analog waveforms allows the user to view and analyze pretrigger waveform information. This feature is essential in diagnosing faults in digital systems that lead to a failure from which a trigger is derived. Second, total bus control of the digitizing oscilloscope, including acquisition and transfer of waveform data, meets the growing needs of automatic test and measurement systems. Third, waveform record lengths are not time-limited by the CRT display, but only by the available high-speed memory which can be the equivalent of many CRT screen widths.

Until now, the primary disadvantage of digitizing oscilloscopes has been the lack of large single-shot bandwidths which are essential for digital system diagnostics, high-energy physics studies, transient interference analysis, and many other applications. The minimum bandwidth required for a large number of these applications is in the

range of 200 to 500 MHz.

The obvious problem in producing a large single-shot (frequently called real-time) bandwidth is the difficulty of sampling and analog-to-digital (A-to-D) conversion at rates of about four times the desired bandwidth. Although the Nyquist criterion only requires a sample rate of twice the signal frequency for accurate reconstruction of sine waves, a ratio of about 4:1 (sample rate/bandwidth) is required for accurate reconstruction of arbitrary waveforms. Attempts to raise the signal bandwidth to values closer to the Nyquist frequency result in increasing overshoot and ringing of the reconstructed system step response as the Nyquist limit is approached.

The HP 54111D Oscilloscope (Fig. 1) was developed to address the high-bandwidth single-shot applications indicated above. The one-gigasample-per-second sampling and A-to-D conversion technology used in the product was developed by HP Laboratories (see article on next page). The high-speed NMOS memory required for long waveform records at minimal power was developed at Colorado Springs Division and processed at the Loveland Technology Center. A general-purpose front-end attenuator and preamplifier




Fig. 1. The HP 54111D Digitizing Oscilloscope can sample high-speed transient signals at 1 gigasamples. It features a 250-MHz single-shot bandwidth and a 500-MHz repetitive bandwidth with up to 8 bits of vertical resolution.

were also designed at Colorado Springs Division (see article on page 67) using custom bipolar ICs and thick-film hybrid technologies from the Santa Clara Technology Center and the Colorado Springs Technology Center, respectively. Descriptions of several of these developments will be presented in the following pages. Digital signal processing techniques incorporated for optimizing signal quality will also be described (see article on page 70).

A One-Gigasample-per-Second Analog-to-Digital Converter

by John J. Corcoran, Ken Poulton, and Knud L. Knudsen

EACH CHANNEL of Hewlett-Packard's Model 54111D Oscilloscope contains a one-gigasample-per-second analog-to-digital converter (ADC) which digitizes the signal from the preamplifier and writes the result to digital memory. The characteristics of this ADC to a large extent determine the horizontal and vertical resolution of the oscilloscope. The one-gigasample-per-second sample

rate, for instance, directly determines the achievable bandwidth in single-shot operation, an impressive 250 MHz.

Four main requirements shaped the architecture of the HP 54111D ADC:

- An 8000-sample memory depth was desired to allow the capture of an eight-microsecond time record at the full

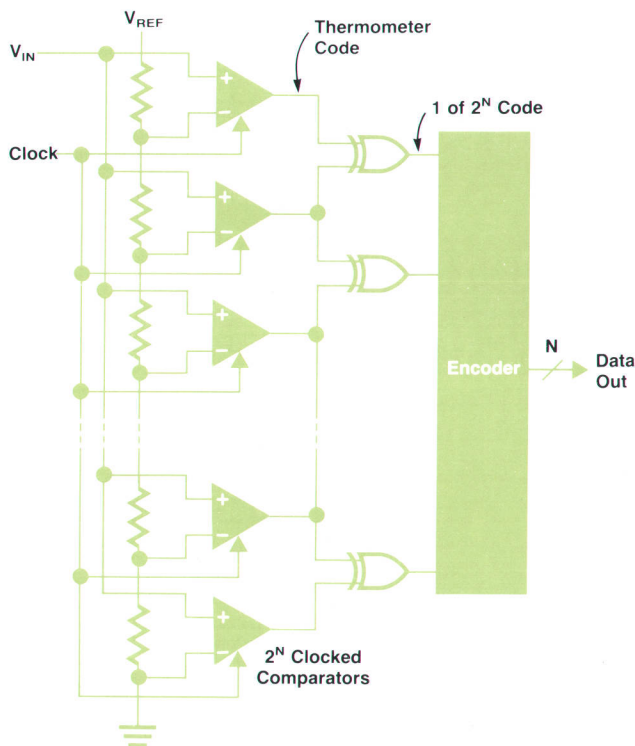


Fig. 1. A traditional flash analog-to-digital converter. This architecture requires 2^N comparators and 2^N exclusive-OR gates.

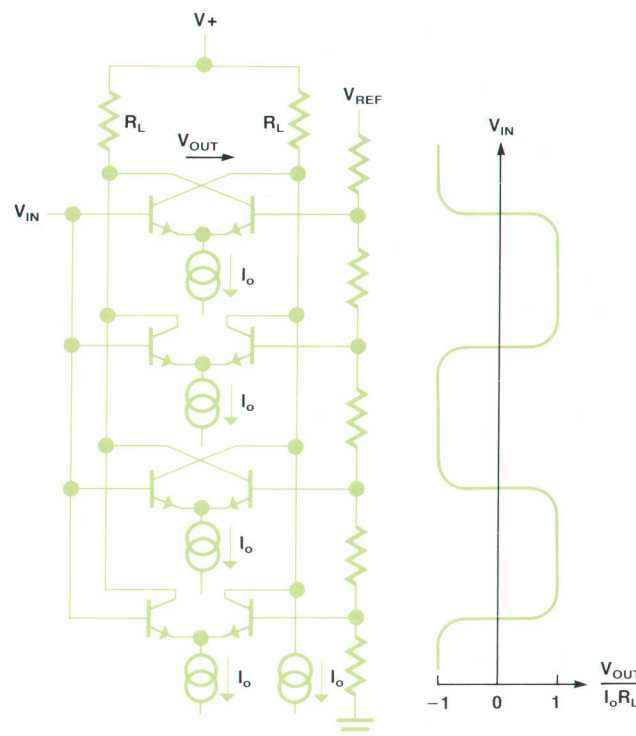


Fig. 2. An analog encoder. Encoding four reference comparisons into one analog voltage enables this ADC to be implemented with fewer devices and operate at lower power.

Repetitive versus Single-Shot Bandwidth

There is often confusion among users of digitizing oscilloscopes about the definitions of repetitive bandwidth and single-shot bandwidth, and how these two quite different specifications relate to the sampling rate of the ADC used in the instrument.

Single-shot bandwidth is important for those applications where the signal occurs infrequently, and it is necessary to digitize the entire event in a single acquisition cycle. These applications include those for which traditional storage oscilloscopes were used in the past, and a host of others made possible by the digital properties of the new instruments.

Repetitive bandwidth applies when the signal occurs often enough that the waveform can be built up from tens or hundreds of successive acquisition cycles to make a composite image on the screen. These applications include those where a traditional (nonstorage) oscilloscope was used in the past. The key difference is that single-shot bandwidth is proportional to the ADC sampling rate, and repetitive bandwidth is quite independent of it.

Single-shot bandwidth is directly related to the ADC sample rate by the Nyquist sampling theorem. This theorem states that to reconstruct a waveform accurately from a limited set of samples of that waveform, the sampling rate must be more than twice the highest frequency component in the waveform. If this condition is violated, frequency components of the input that are above the Nyquist frequency (half the sampling frequency) will be translated by the sampling process to different frequencies below the Nyquist frequency. This frequency shifting (also known as aliasing) causes errors that cannot be corrected in the sampled data. To avoid these errors, the signal bandwidth must be limited to less than half the sampling rate, either by its intrinsic characteristics or by an anti-aliasing filter in the signal path. Then, to reconstruct the entire input signal correctly from the limited set of samples taken, a digital low-pass filter is used to interpolate between the actual sample points. Both the anti-aliasing filter and this reconstruction filter must have significant attenuation at the Nyquist frequency to avoid errors. As a practical matter, to avoid excessive overshoot in the step response of these filters, their composite bandwidth cannot exceed about one-fourth the sample frequency. Thus the filter requirements reduce the achievable single-shot bandwidth from the theoretical one-half of the sampling rate to the more practical one-fourth of the sampling rate.

Repetitive bandwidth, by contrast, is totally independent of the A-to-D sampling rate. In repetitive sampling, an image of the signal is built up by interleaving the samples acquired in multiple occurrences of the input signal. The proper placement (on the time axis) of each set of samples is guaranteed by a special circuit in the instrument that measures the time delay between the internal sampling clock and the trigger, which occurs at the same point on the signal for each acquisition. The composite image thus built up has a very high effective sampling rate, and there is no need for either anti-aliasing or reconstruction filters. The bandwidth in this mode then depends only on analog components like the input attenuator, preamplifier, and sample-and-hold circuit, and not on the A-to-D sampling rate.

John J. Corcoran
Project Manager
Hewlett-Packard Laboratories

single-shot bandwidth.

- The power consumption of the analog-to-digital subsystem was required to be less than 25 watts per channel because of instrument power constraints.
- The resolution of the ADC had to be consistent with accurate waveform reconstruction.
- A 1-GHz analog-to-digital conversion bandwidth was desired so that the overall goal of a 500-MHz instrument bandwidth for repetitive signals could be readily achieved (see box at left for the distinction between single-shot and repetitive bandwidth).

The requirement for an 8000-sample memory system made it difficult to use an indirect A-to-D architecture, that is, one that requires intermediate analog storage. In an indirect A-to-D system, the samples are captured at the full sample rate and stored in an analog memory (for example, a charge-coupled device). The samples are then read out slowly and digitized by a low-speed ADC. This architecture offers the attractive advantages of a lower-speed A-to-D converter and digital memory. However, errors that increase with the storage time of the analog storage device and the power and complexity of the associated clock drivers limit the total practical storage capacity. Hence, the decision was made to build a direct ADC, where the input signal is digitized at the full sample rate and the data is stored directly in high-speed digital memory. The memory depth then can be increased without affecting the analog accuracy by simply adding digital memory chips.

The power consumption constraint immediately made it clear that custom integrated circuits would be required throughout the A-to-D system. A board-level implementation, using commercially available components, would have required too much power and probably could not have met the speed objectives. Next, each IC would need to be fabricated in a state-of-the-art technology with excellent power-delay product, whose performance characteristics are well-matched to the needs of the individual ICs. This led to the use of three different IC technologies:

- Gallium arsenide (GaAs) is used for the front end of the system, where the speed requirements are the greatest.
- Silicon bipolar circuitry is used to digitize the signals, because of its excellent device matching and high transconductance.
- Silicon NMOS is used for the memory, since it offers an excellent combination of speed and device density.

The silicon bipolar digitizer IC, which uses a modified

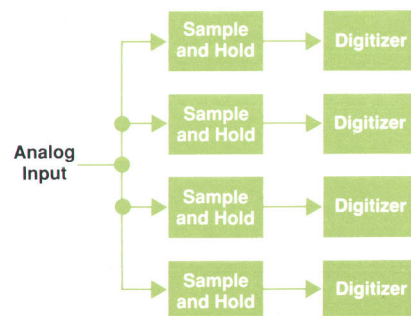


Fig. 3. A simply interleaved A-to-D system. Each digitizer is driven from a separate sample-and-hold circuit.

flash architecture, sets the resolution of the system. In the flash architecture there is one comparator for each input level to be resolved. Thus the transistor count and power are proportional to 2^N , where N is the number of bits of resolution. It was decided early that a six-bit digitizer could both meet the power budget of the system and provide quite acceptable waveform reconstruction, especially when combined with some dithering and digital signal processing features (see the article on page 70). With these additions, the effective resolution of the system is seven or even eight bits for slowly varying signals, where the higher resolution is most useful.

To achieve the 1-GHz A-to-D system bandwidth goal, it became quite clear that a sample-and-hold circuit would be required. A flash converter intrinsically samples the input signal in its clocked comparators, but the usable input bandwidth is typically much less than 1 GHz. For the digitizer designed for this system, the bandwidth for accurate conversion without an external sample-and-hold circuit is about 100 MHz. Thus, it was necessary to add an external sample-and-hold circuit, and investigation showed that only a GaAs FET IC could supply the needed bandwidth and acquisition time.

Digitizer

The silicon bipolar integrated circuit implementing the digitizing function for the system uses a flash architecture, which generally yields the highest A-to-D sample rate possible within a given IC technology. However, even with HP's oxide-isolated 5-GHz IC process (the fastest bipolar process available at the time), a digitizer could not be built that could be clocked at the required one-gigasample/second rate. An even more severe constraint was the requirement to drive and settle the input capacitance of the digitizer with the step-like input signals supplied by the sample-and-hold circuit. Simulations showed that the best achievable settling time, with margins appropriate for a production instrument, was about 4 ns. Thus it was decided to build a digitizer that could be clocked at 250 megasamples/second (a 4-ns period), and interleave four of these to achieve the required net sample rate of 1 GHz.

A block diagram of a traditional flash ADC is shown in Fig. 1. For six-bit resolution, 64 clocked comparators are required. The reference voltage V_{REF} , equal to the full-scale input of the converter, is divided by the resistor network and distributed to the comparators. When the comparators are clocked, their outputs go to a logical one or zero state, depending on the relative value of the comparator reference and input signal at the time the clock occurred, resulting in a "thermometer code" with all zeros above the comparator nearest the input voltage, and all ones below that comparator. The thermometer code contains a single point where an adjacent one and zero occur, and this is detected by one of the exclusive-OR gates. The high exclusive-OR output then drives the encoder to produce the correct binary value, which is the output of the ADC.

Although a flash converter is very fast, it has several disadvantages. The device count and power are both high, because of the fully parallel nature of the architecture. For the same reason, the input and clock capacitances are high, and these lead to practical problems in driving the con-

verter from other circuits. To minimize these problems, we use a circuit that we call an analog encoder^{1,2} to reduce the number of clocked comparators required.

A block diagram of an analog encoder is shown in Fig. 2. It consists of four differential pairs, connected with collectors alternately crossed as shown. The input signal drives one side and the reference voltages drive the other side of the differential pairs. The resulting transfer function of this circuit, also shown in Fig. 2, changes between a high and a low level four times. These changes occur when the input voltage is equal to one of the reference voltages applied to the differential pairs. This circuit essentially implements the functionality of four comparators and four exclusive-OR circuits in the traditional flash converter, except that the output is analog. Therefore, a single clocked comparator is placed after the analog encoder. The net result is that four differential pairs and one clocked comparator replace four clocked comparators and four exclusive-OR gates. This is a net gain of about a factor of four in transistor count and a factor of at least two in power. Furthermore, since the number of clocked comparators is about four times lower, the clock input capacitance is similarly reduced.

The six-bit digitizer is implemented using seventeen of these analog encoders and seventeen clocked comparators. To produce the least-significant bits of the output data (where more than four transitions must be encoded), the outputs of several encoder-comparator groups, with appropriately interleaved thresholds, are logically combined.

The digital encoder produces a Gray-coded output. Gray code is a method of coding 2^N levels in N bits. Its main characteristic is that a transition from any level to a neighboring level causes only one output bit to change. The more common binary coding has points where going from one level to the next can change all output bits simultaneously. If a comparator in a binary digitizer should be undecided because its reference voltage is equal to the input, then the output can produce large errors if it is interpreted differently by the encoding paths for the different output bits affected. But in a Gray-coded digitizer, each comparator affects only one output bit, and hence, if a comparator should be undecided, the output can only be interpreted to be one of two adjacent codes, which can be considered to be equally correct measures of the input voltage.

Although analog encoded architecture reduces the de-




Fig. 4. Effect of uneven sample timing in an interleaved system.

vice count, power, and clock capacitance, it does nothing to reduce the high capacitance at the analog input, which is typical of flash converters. Simulations showed it would be impossible to drive and settle the expected input capacitance through the inductance of the bond wire used to connect the digitizer's input to its package. To eliminate this problem, an on-chip input buffer, optimized for step settling, was added to the digitizer. This buffer reduces the effective input capacitance from 25 pF to 4 pF and allows settling in the required 4 ns.

The resulting digitizer chip² is quite small, measuring only 2.2×3.5 mm. This chip is a complete ADC in its own right; it can be clocked at rates up to 400 MHz, and it can be used without a sample-and-hold circuit on signal bandwidths of up to 100 MHz. It is currently used in the HP 54201 and HP 54112 Digital Oscilloscopes and in an optional board for the HP 16500A Logic Analyzer. Power dissipation of the digitizer is 2.5 watts, so the four chips required for the 1-GHz A-to-D system consume a total of 10 watts.

Interleaving Digitizers

To meet the required sampling rate of 1 GHz, four of the silicon digitizers are interleaved, each operating at 250 MHz. A simply interleaved A-to-D system is shown in Fig. 3, where each digitizer has its own sample-and-hold circuit. The problem with this arrangement is that the four samplers sample the input signal separately, so any errors in their relative timing will be reflected as inaccurate sampling. As shown in Fig. 4, a deviation Δt in the desired timing of sample 3 results in an error ΔV in the voltage displayed

for that sample. For high-slew-rate signals, such as a full-scale, 500-MHz input signal, maintaining six-bit accuracy requires that the relative timing be accurate to 5 ps. With the IC and packaging processes available, timing errors of around 50 ps were expected.

To avoid the 5-ps timing accuracy requirement, a fifth sample-and-hold circuit, operating at 1 GHz, was placed in front of the four sample-and-hold circuits operating at 250 MHz (Fig. 5). This first-rank sample-and-hold circuit samples for 500 ps, acquiring the input signal on its hold capacitor. Then it goes into hold for 500 ps, which creates a flat spot on an otherwise quickly moving waveform. The four second-rank sample-and-hold circuits each sample for 1 ns and hold for 3 ns. They will each finish their acquisition time during one of the 500-ps flat spots. Since they finish their acquisition on a signal that is no longer moving, the timing of their clock signals is no longer critical. Producing a sufficiently stable clock for the single first-rank sample-and-hold circuit is relatively easy.

Sample-and-Hold Circuits

The heart of each sample-and-hold circuit is the diode bridge sampling gate.³ A simplified schematic is shown in Fig. 6. In the sample mode, I_{CONTROL} flows down through the four diodes, D1-D4, forming the bridge. If the voltage V_{IN} is the same as V_{OUT} , the current is split evenly between diodes D1 and D3 and no charging of the 0.3-pF hold capacitor takes place. However, if V_{IN} is higher than V_{OUT} , the voltages across D1 and D4 will be smaller than the voltages across D3 and D2. Hence, the current in D3 will be larger than the current in D4, which causes a net current




Fig. 5. Sample-and-hold chip block diagram. All buffers are unity-gain source followers. The bridges are diode bridge sampling gates.

flowing into the hold capacitor until V_{IN} and V_{OUT} are the same.

The bridge is turned off (put into hold mode) by reversing the current $I_{CONTROL}$. In this case the current flows through diodes D5 and D6, which places a reverse bias on diodes D1-D4, breaking the connection between V_{IN} and V_{OUT} .

A diode bridge sampling gate has several key advantages over simpler circuits such as pass-FET circuits:

- The circuit turns on and off very quickly.
- The diode voltages need only decrease by approximately 200 mV ($8V_T$) to turn off the bridge effectively.
- The circuit is intrinsically balanced, so there is (to first order) no clock feedthrough onto V_{OUT} .
- The hold isolation is quite good. Diodes D1, D5, and D3 effectively form a series-shunt-series isolation circuit in the off state, as do diodes D2, D6, and D4.

A distinction is sometimes made between a sample-and-hold and a track-and-hold approach. If the charging of the hold capacitor is allowed to occur until V_{IN} and V_{OUT} equalize, the circuit is often called a track-and-hold circuit. If the on interval is shorter than the settling time, the circuit is called a sample-and-hold circuit. Using GaAs ICs allowed us to design our circuit as a track-and-hold circuit—the voltage on the hold capacitor settles completely to V_{IN} during the 500-ps on time. The track-and-hold circuit has a higher gain than a sample-and-hold circuit and it is insensitive to variations in the pulse width of the sampling clock. In the rest of this article, we shall use the terms track-and-hold and sample-and-hold interchangeably.

HP's GaAs IC process uses 1- μ m gate lengths to build depletion MESFETs (metal-semiconductor FETs) with a pinchoff voltage V_p of $-2V$ and a unity-current-gain frequency f_T of 14 GHz. Since a MESFET contains a Schottky diode junction, a fast diode ($R_{ON}C_{OFF} = 1$ ps) can be made from a MESFET by connecting its drain and source together as one electrode and using the MESFET's gate as the other. This process also includes integrated resistors and capacitors. A GaAs FET process is particularly well-suited to building sample-and-hold circuits because:

- There are fast Schottky diodes for building diode bridge

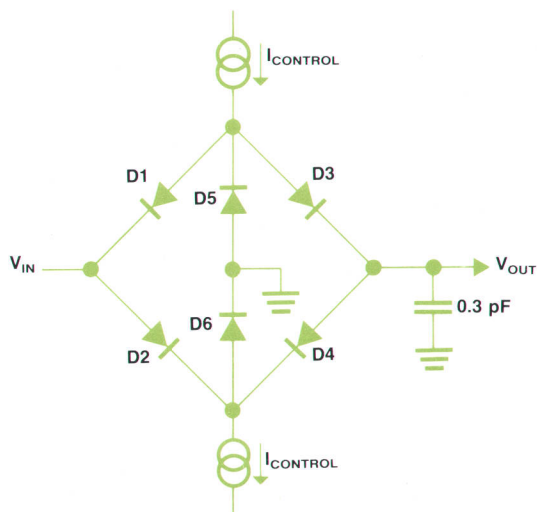


Fig. 6. Simplified schematic of the diode bridge sampling gate. $I_{CONTROL}$ reverses direction to turn the sampling gate off.

circuits.

- The FET speed allows construction of fast switching circuits for control of the bridges and very fast source follower circuits for buffering the inputs and outputs of the bridges.
- The semi-insulating GaAs substrate reduces parasitic capacitances.

This process is described in more detail in reference 4.

The five sample-and-hold circuits shown in Fig. 5 are integrated on one GaAs chip. A companion clock generation chip, also built in GaAs, accepts the input clock (at 1 GHz or below), amplifies it to GaAs levels (3V swings), and provides the clocks for the five sample-and-hold circuits and the four digitizers. Various characteristics of the chips used in this ADC are given in Table I.

Table I

Chip:	Sampler	Control	Digitizer
Technology:	GaAs D-MESFET 1- μ m Gate	GaAs D-MESFET 1- μ m Gate	Si Bipolar 5-GHz f_T
Chip Size:	1.4 \times 1.7 mm	1.1 \times 1.7 mm	2.2 \times 3.5 mm
Device Count:	460	390	920
Number of Pads:	40	35	52
Power:	2.9W	2.5W	2.5W

Increasing Resolution

The major problem with an interleaved ADC is the need to provide good matching between the four digitizers, both in time and in voltage. The need for matching in time is addressed by the use of two ranks of sample-and-hold circuits. The need for matching in voltage is addressed by controlling the voltage references used by the four digitizers with individual digital-to-analog converters that are set by the system processor during a self-calibration cycle.

However, having multiple digitizers can be an advantage. The four 6-bit digitizers' reference voltages can be offset from each other in steps of 1/4 LSB (least-significant bit), that is, in steps of 1/256th of the input voltage range. This causes the four digitizers to be interleaved in voltage, as well as in time. This allows 256 levels to be resolved for slowly changing input signals, so the resolution for low-frequency input signals can be eight bits rather than six bits.

To display a resolution of 256 levels, the data from all four digitizers must be combined. This can be done with a nonlinear smoothing filter, which moves each sample value towards the average of four adjacent samples, but never moves a sample by more than half a six-bit LSB (see "Dithering in the HP 54111D," page 72). Although this mechanism can give true 8-bit resolution on sine waves only if the signal frequency is below 1 MHz, it preserves the instrument signal bandwidth and provides a significant enhancement in resolution on the flat portions of pulse waveforms of any speed.

System Performance

The frequency response of the ADC, sampling at 1 GHz,

(continued on page 65)

Digitizer Hybrid

The three-inch-square digitizer hybrid circuit (see Fig. 1) in the HP 54111D Digitizing Oscilloscope contains the GaAs track-and-hold chip, the GaAs clock and mode control chip, and the four six-bit, Gray code, bipolar analog-to-digital flash converters that operate at 250 million conversions per second. Two main contributions that the thick-film technology makes to the circuit are high-frequency performance and high power dissipation.




Fig. 1. The one-gigasample-per-second digitizer hybrid circuit used in the HP 54111D Digitizing Oscilloscope. The GaAs sample-and-hold and clock generation chips are in the center. The four 250-MHz bipolar digitizer chips are located in each corner.

The high-frequency performance is made possible by mounting the IC chips directly onto the thick-film substrate and then wire-bonding the chip-to-substrate connections. This eliminates the need for an intermediate level of packaging, which would




Fig. 2. Close-up view of the two GaAs IC chips in the center of the digitizer hybrid circuit of Fig. 1.

add delay and parasitics to the circuit. Fig. 2 shows a close-up view of the two GaAs chips mounted in the center of the substrate.

The track-and-hold and clock controller chips are wire-bonded directly together (chip-to-chip). This keeps the lead lengths between them as short as possible to minimize delay, the parasitic inductance of each lead, and the skew between the lines.

The circuit uses 50 and 75-ohm transmission lines for the analog input signals and clock lines, respectively. Thick-film technology provides well-controlled impedances for transmission lines. It also makes possible the fabrication of very small resistors to minimize parasitics.

The second contribution that thick-film technology makes to the circuit is in dissipating power—lots of power. The hybrid circuit dissipates 17 watts. The two GaAs chips, dissipating 6W, are attached directly to a heat spreader that fits through a hole in the center of the substrate (refer to Fig. 2). This heat spreader has a thermal coefficient of expansion that nearly matches that of the 96% alumina substrate and has a thermal conduction value between copper and tungsten. Attaching the chips directly to the heat spreader results in a decrease in junction temperature




Fig. 3. Data acquisition board for the HP 54111D. The digitizer hybrid circuit is located beneath the shield in the center.

of nearly 20°C compared to mounting the chips on the alumina. The chips are mounted to the heat spreader using conductive epoxy. Although the epoxy process has a higher thermal resistance than a solder process, it is less costly and simpler.

The 148 interconnections between the hybrid circuit and the printed circuit board are made using flex circuits. Flex circuits consist of flat beryllium copper conductors etched on a polyimide film. The four flex circuit strips are vapor phase soldered around the perimeter of the substrate, and the other ends of the strips are soldered to the printed circuit board. The flex circuits provide excellent high-frequency performance and allow stress relief between the hybrid circuit and the printed circuit board on which the hybrid is mounted.

A low-profile, large-area heat sink mounts on the back of and extends through the printed circuit board. The heat sink contacts

the bottom of the hybrid circuit at the heat spreader and directly beneath each of the four analog-to-digital converters, each of which dissipates nearly 3W. Five ceramic lids protect the hybrid circuit's wire-bonded IC chips. A shield covers the hybrid circuit and presses down on these lids, ensuring that good contact exists between the heat sink and the back of the hybrid circuit. Fig. 3 shows a picture of the data acquisition board. The digitizer hybrid circuit is underneath the metal shield in the center and the flex circuit leads are just visible around the perimeter of the shield.

Lewis R. Dove

Mark E. Mathews

Development Engineers

Colorado Integrated Circuits Division

(continued from page 63)

is shown in Fig. 7. The 3-dB bandwidth is 1.7 GHz and the distortion is below -40 dBc (1/3 LSB for 6-bit resolution) for input signals up to 500 MHz. (The performance measurements quoted here are for the A-to-D system alone, and do not include the HP 54111D's input attenuator and preamplifier circuits.) Any timing misalignments between the four digitizers would show up as frequency components (sampling products) around 250 MHz; these components are suppressed to below -40 dBc by the use of the two-rank sample-and-hold architecture.

Effective bits of resolution is a measure of analog-to-digital conversion accuracy;⁵ it measures an ADC's performance in digitizing a sine wave, compared against a hypothetical ideal ADC of that resolution. Any noise, non-linearity, or missing codes will cause an ADC's effective bits of resolution to fall below its nominal resolution. Thus, a perfect, noiseless, 6-bit ADC will show 6.0 effective bits of resolution; a poor one might show 3.5 effective bits.

The resolution, in effective bits, achieved by this ADC while sampling at 1 GHz is shown in Fig. 8. Here, 5.2 effective bits are measured all the way up to a 1-GHz input frequency. The limiting factor in this case is the mismatch between the four digitizers, since a single digitizer achieves 5.9 effective bits of resolution at 250 megasamples/second.

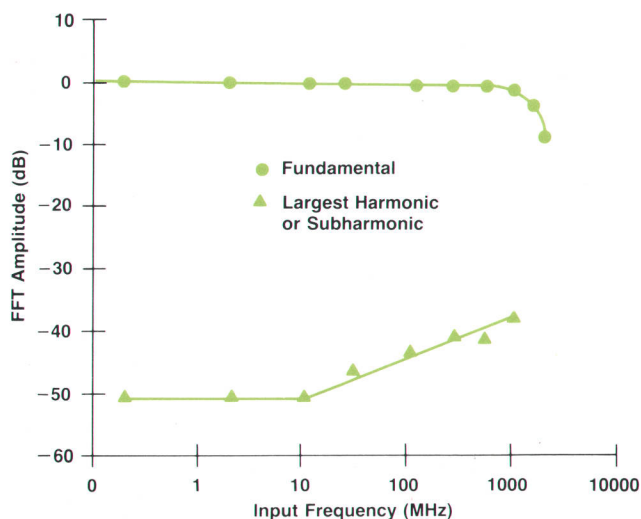


Fig. 7. ADC frequency response and distortion.

At lower input frequencies, the number of effective bits rises to nearly 7 as the higher resolution afforded by the nonlinear smoothing filter comes into effect. At low frequencies, the number of effective bits is limited mainly by the nonlinearities of the digitizers.

Fig. 9 shows the result of a beat frequency test using the voltage interleaving and nonlinear filtering described above. The larger sine wave is a 250-kHz sinusoidal signal sampled at 1 GHz. The smaller sine wave is the result of sampling a 1.000250-GHz sinusoidal signal. Even at a 1-GHz input frequency, the digitized waveform is clean with no visible aberrations.

Packaging

The A-to-D system (that is, the GaAs sample-and-hold chip, the GaAs clock generation chip, and the four bipolar digitizer chips) is packaged on a single thick-film hybrid substrate (see box on page 64). The two GaAs chips are side by side in the center and surrounded by monoblock bypass capacitors and thick-film termination resistors. The four digitizer chips are placed near the four corners of the hybrid, along with their associated capacitors and resistors.

Acknowledgments

We gratefully acknowledge the contributions of Roman Kagarlitsky, who built the test system for the prototype




Fig. 8. ADC effective bits of resolution. Single ADC data represents every fourth point from an interleaved 1-gigasample-per-second record.




Fig. 9. Reconstructed sine waves with resolution enhanced by interleaving the digitizers in voltage.

A-to-D system and wrote the software for it, Paul Clark and Don Hiller, who contributed to the design of the digitizer chip, Mike Kreikemeier, who designed the early versions of the hybrid packaging, Pauline Prather and Dorothy Hol-lars, for amazing bonding work, the GaAs IC group at HP's Microwave Technology Division and the bipolar IC group

at the Santa Clara Technology Center, Joe Millard and his team at Colorado Springs Division, and Tom Hornak, for his ideas, enthusiasm, and guidance.

References

1. T.W. Henry and M.P. Morgenthaler, "Direct Flash Analog-to-Digital Converter and Method," *U.S. Patent 4,386,339*, Hewlett-Packard Company.
2. J. Corcoran and K. Knudsen, "A 400-MHz 6b A-to-D," *ISSCC Digest of Technical Papers*, February 1984, pp. 294-295.
3. J.R. Gray and S.C. Kitsopoulos, "A Precision Sample and Hold Circuit with Subnanosecond Switching," *IEEE Transactions on Circuit Theory*, September 1964, pp. 389-396.
4. R.L. Van Tuyl, et al., "A Manufacturing Process for Analog and Digital GaAs Integrated Circuits," *IEEE Transactions on Micro-wave Theory and Techniques*, Vol. MTT-30, no. 7, July 1982, pp. 935-942.
5. B. Peetz, "Dynamic Testing of Waveform Recorders," *IEEE Transactions on Instrumentation and Measurement*, Vol. IM-32, no. 1, March 1983.
6. J. Corcoran, K. Poulton, and T. Hornak, "A 1-GHz 6b A-to-D System," *ISSCC Digest of Technical Papers*, February 1987, pp. 102-103.
7. K. Poulton, J. Corcoran, and T. Hornak, "A 1-GHz 6-bit A-to-D System," *IEEE Journal of Solid-State Circuits*, Vol. SC-22, no. 6, December 1987.

Front-End Signal Conditioning for a High-Speed Digitizing Oscilloscope

by Thomas K. Bohley, Mark E. Mathews, Lewis R. Dove, Joe K. Millard, David W. Bigelow, and Donald D. Skarke

CRITICAL TO THE ACHIEVEMENT OF high-performance in any digitizing oscilloscope is the design of the front-end signal conditioning stages. The design of these stages in the HP 54111D Oscilloscope was particularly difficult because of the HP 54111D's 500-MHz required signal bandwidth. Key elements in the design are the attenuator, the switch assembly, the hybrid circuit upon which the attenuator is fabricated, and the preamplifier.

Attenuator

The theory of operation of a compensated one-megohm attenuator, shown in Fig. 1, can be found in any text on linear waveshaping. Why then have they not been used in 500-MHz oscilloscopes before? The answer can be found in Fig. 2; not in the divide-by-10 attenuator sections that can be cascaded, but rather in the straight-through divide-by-1 path. It is in that path that the inevitable series parasitic inductance is at a maximum, and it is in that path that the series inductance and the input capacitance of the following preamplifier form a passive resonant circuit that introduces ringing on input pulses and limits the bandwidth.

An effective solution to this problem is to raise the resonant frequency high enough so that the oscilloscope's amplifier no longer responds to that resonant frequency. Generally, attempting to damp the resonant circuit by introducing a series resistance in the inductive path results in creating an RC corner in the passband of the following amplifier.

To raise the resonant frequency we must make both the series inductance and shunt capacitance as small as possible. The shunt capacitance, determined mostly by the input FET, is already small since a high-frequency FET has been chosen. This leaves the series inductance to be minimized and the most effective way of doing that is simply to limit the physical length of the straight-through path.

The attenuator is constructed on a thick-film 0.5×1.0 -inch substrate with the input and output separated by about 0.5 inch. To build an attenuator of this size would not be possible without one more technique—moving the ac coupling capacitor from the attenuator itself, where the capaci-




Fig. 1. Compensated attenuator circuit.



Fig. 2. Input attenuator with $\div 10$ sections and straight-through $\div 1$ path.

tor also contributes series inductance, to the preamplifier where it is placed in a low-frequency path. The divide-by-10 sections require considerable care to ensure that they, because of their own parasitic inductance (Fig. 3), also do not form resonant circuits. In this attenuator, resistors R3 and R4 are sized to damp the parasitic inductance associated with the capacitors and form a resistive divide-by-10 attenuator for signals well beyond the amplifier passband. Resistors R1 and R2 carry only relatively low-speed signals and are less important to the high-frequency performance of the attenuator than capacitors C1 and C2. The divide-by-10 sections have greater bandwidth than the straight-through path and must be intentionally slowed down to keep the response as seen on screen consistent from range to range.

Switch Assembly

The challenges in designing the switch assembly were to provide small size, repeatable characteristics, reliability, long life, and remote switching. The small size required that contact movement be minimal and repeatable. Otherwise, contact stress would be excessive and electrical characteristics would vary. Minimal contact movement also complicated the remote control requirement. The solution to these problems was accomplished by the design shown in Fig. 4. Contact lift is controlled by the shuttle width, contact pin location, and dimensions of the formed contact. These parameters are controlled by the tooling used in their manufacture and thus are repeatable. Shuttle travel is not critical and is easily within the bounds of the solenoids and associated switching apparatus. The contact




Fig. 3. Resistors R3 and R4 added to circuit of Fig. 1 to damp parasitic inductances associated with capacitors C1 and C2.

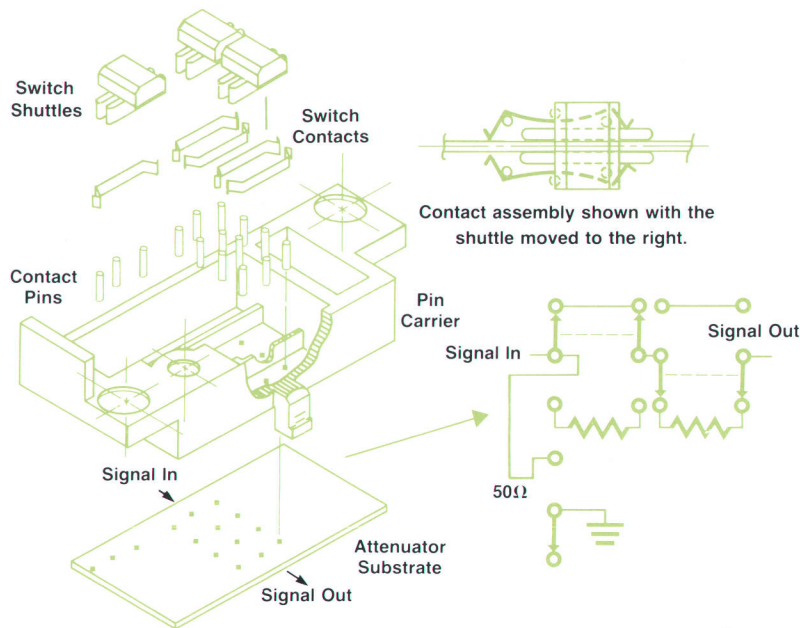


Fig. 4. Attenuator switch assembly for HP 54111D.

pins are held and positioned by the pin carrier and connect directly to the attenuator substrate to keep the electrical path short. The switch contacts and contact pins are made of high-gold-content alloys for reliability and to reduce vulnerability to plating variations and wear.

Attenuator Hybrid Circuit

The attenuator hybrid circuit is shown in Fig. 5. This hybrid circuit has a switchable input impedance of one megohm or 50 ohms. There are three decades of attenuation at one megohm with 1:1, 1:10, and 1:100 divider ratios. Switch contacts are made as described in Fig. 4 using small pins that extend through the hybrid substrate. The electromechanical actuators move a shuttle to open or close the appropriate contact.

The attenuator hybrid circuit's bandwidth exceeds 500 MHz. The high bandwidth is achieved by implementing the circuit elements in thick-film patterns and by miniaturizing the contact assembly, both of which minimize parasitic inductance and capacitance.

The contact pins are attached to the ceramic substrate with conductive epoxy. A print process using a stencil was developed to apply epoxy automatically to the fifteen pins and into the pins' holes. To keep the epoxy from running completely through the holes to the backside of the substrate, positive air pressure is maintained on the backside. The pressure is turned off after the print squeegee passes over the holes so that it does not blow the epoxy back out of the holes.

Most of the thick-film resistors are passively trimmed to a predetermined resistance value. However, one resistor is actively trimmed to set a circuit parameter. The resistor is used in the divider network of a circuit that senses dangerous input overvoltages which could harm the 50-ohm, 1% input resistor. The divider resistor is actively trimmed by applying 10 volts to the input and trimming until 100 mV is obtained at the divider's output. Active trimming is typically much less expensive than the alternative of manually

adjusting a variable component.

Preamp Hybrid Circuit

The preamplifier hybrid circuit shown in Fig. 6 is mounted in the attenuator/preamp assembly. This amplifier boosts the signal from the attenuator to the amplitude required by the GaAs track-and-hold stage. The hybrid circuit has one-megohm input impedance, a bandwidth of 750 MHz, and an overall gain of 36 with a 50-ohm output load. It also provides trigger pickoff and conditioning. The 1.5×3.0 -inch hybrid circuit dissipates 5 watts. This hybrid contains a custom input MOSFET, two custom bipolar ICs, and two pnp and one npn discrete microwave transistor chips.

The custom MOSFET designed at the Loveland Technology Center includes a limited overload protection circuit. The first custom bipolar chip is a preamplifier that has a




Fig. 5. Top (left) and bottom (right) views of attenuator hybrid circuit.

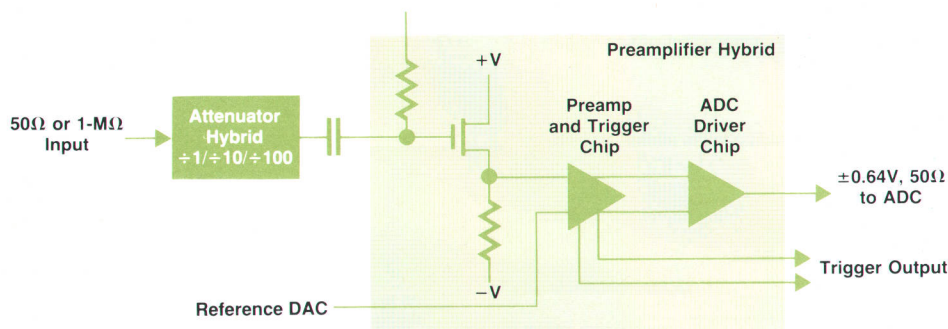


Fig. 6. Preamplifier hybrid circuit diagram.




Fig. 7. Preamplifier hybrid circuit (right) and attenuator (left) mounted in the preamp/attenuator assembly.

programmable 1,2,4 attenuator and a continuous vernier. This allows a user to adjust the input signal amplitude to full screen for maximum resolution. In addition, this chip generates a logic trigger signal used to reduce trigger jitter caused by noise contamination of the analog sync used in most instruments. Finally, an auxiliary input on the preamplifier chip can be switched on to calibrate the entire vertical deflection path from the input of the preamplifier chip through the analog-to-digital converters (ADCs), thus ensuring excellent gain and offset stability.

The second custom bipolar chip is a driver chip which provides the final gain needed to drive the GaAs track-and-hold circuits. Included on the chip is a selectable $\times 5$ amplifier.

As a final part of the signal conditioning the differential signal from the ADC driver chip is level-shifted to near ground and converted to a single-ended signal. The circuit used is a pnp level-shift circuit. The final output voltage

of the preamplifier hybrid is ± 0.64 volt with a 50-ohm load.

The preamplifier hybrid circuit is connected to the attenuator hybrid circuit through a compressed leaf spring that contacts gold pads on each of the hybrid circuits. Fig. 7 shows the placement of the preamplifier next to the attenuator in the assembly (the attenuator is to the left of the preamplifier). The trigger output from the preamplifier IC drives 50-ohm differential transmission lines on the thick-film substrate. These lines lead to two SMC connectors epoxied to the back of the substrate. The connector pins protrude through holes in the substrate so that they can be soldered to the transmission lines on top. The linear signal output is connected to a third backside SMC connector through a 50-ohm microstrip. All of the bias and ground connections to the hybrid circuit are made via two dual in-line connectors. These are soldered through holes in the substrate.

Digital Filtering in a High-Speed Digitizing Oscilloscope

by B. Allen Montijo

IN THE LAST FEW YEARS the digitizing oscilloscope has emerged as a useful tool for waveform acquisition and characterization of repetitive events. As a real-time (single-shot) tool, misconceptions about its abilities and limitations are restricting its use. Confusion over "effective bits" and "interpolation" causes the user to question the validity of data and measurements. To add to the confusion, digital signal processing techniques support concepts such as trading excess bandwidth for extra resolution.

Two events occur when a waveform is digitized; the signal is sampled and it is converted to a digital format. Sampling should be thought of as an analog process. This procedure transforms the signal from the s-plane to the z-plane and can create aliasing. Once aliasing occurs, no operation can recover the original signal without some knowledge about the signal. After the signal is sampled, an analog-to-digital converter (ADC) generates a digital code that represents the analog sample. This process is simply a format conversion, similar to converting a real number to an integer on a computer. The conversion adds quantization noise to the signal which is dependent upon the quality of the ADC and the number of bits the ADC uses to represent the sample. The digitized signal can be thought of as the sampled signal with noise added. This article deals with the signal processing aspects of digitizing and assumes that aliasing is not a concern. With this assumption, the digitized signal contains all of the information necessary to reconstruct the original continuous signal, with some additional noise. Then the digital samples can be interpreted as the original analog signal with quantization noise added. This concept forms the basis for believing the information displayed by a digitizing instrument. So long as you know what the signal processing is doing to your signal and that aliasing is not a problem, you can believe what is on your screen.

Effective Bits and Noise

Effective bits has become a popular term for describing the performance of digitizing instruments. The term compares the noise characteristics of the instrument with those of an ideal ADC. To find an expression for effective bits, begin with the equation for the mean squared error of an ideal ADC. The error can range from $-1/2$ to $+1/2$ of a quantization level (Q) with equal probability. This error is squared, integrated, and divided by the interval:

$$s_{\text{adc}}^2 = (1/Q) \int_{-Q/2}^{Q/2} x^2 dx = Q^2/12 \quad (1)$$

If the ADC covers a voltage range of A and has N bits, then

$$Q = A/2^N \quad (2)$$

so that

$$s_{\text{adc}}^2 = A^2/(12 \times 4^N) \quad (3)$$

The equation can be solved for N to give:

$$N = \log_4 (A^2/s_{\text{adc}}^2) - 1.7925 \quad (4)$$

For a full-scale sine wave input, the rms voltage is

$$B = A/2\sqrt{2}. \quad (5)$$

Then

$$N = \log_4 (B^2/s_{\text{adc}}^2) - 0.2925 \quad (6)$$

To find effective bits for a system, the noise term s_{adc}^2 is replaced with the total system noise $s_{\text{system}}^2 = s_{\text{analog}}^2 + s_{\text{adc}}^2$:

$$N_{\text{eff}} = \log_4 (B^2/s_{\text{system}}^2) - 0.2925 \quad (7)$$

The variable term (B^2/s_{system}^2) is simply a signal-to-noise ratio (SNR). Effective bits is an expression of SNR; a factor of 4 in SNR (a factor of 2 in voltage) is one effective bit.

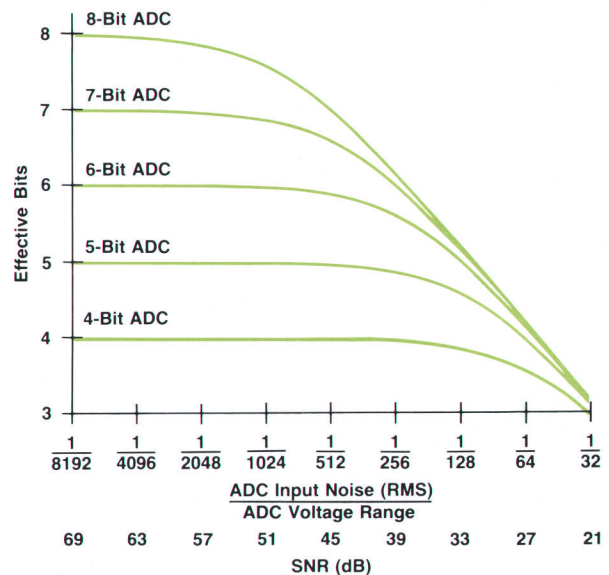


Fig. 1. Effective bits versus input noise for ideal analog-to-digital converters. The dB scale is for a full-scale sine wave input.

The SNR for an analog system can be plugged into equation 7 for comparison with a digitizing system.

For a digitizing system, the SNR is often found using the sine-wave curve-fit test. A spectrally pure sine wave is digitized, then the best-fit sine wave is calculated for the data. The best-fit routine finds the ideal sine wave that minimizes the rms error between the data and the ideal sine wave. This procedure cancels any systematic errors in offset, gain, frequency (sampling and signal), and phase. The residual rms error is substituted into the equation as the system noise voltage s_{system} to obtain effective bits at the signal frequency. The test is repeated at various signal frequencies to obtain a curve of effective bits versus frequency.

The test is not as straightforward as it sounds. Signal amplitude is an important parameter because the slew rate of the signal and the amount of distortion depend upon the amplitude. High slew rates put stress on the analog components and the ADC, and magnify the aperture jitter. Distortion is important because the harmonics generated are considered noise. For most statistical operations, more data is better. In this case, more data places stricter requirements on the phase noise of the sampler and the test signal. Since the ideal sine wave is spectrally pure, any relative phase noise between the sampling clock and the signal decreases the number of effective bits. Finally, the purity of the test signal must be better than the unit under test, preferably by a factor of ten. This often requires the use of filters in the signal path.

Care must also be taken in gathering data. The signal frequency should be selected so that each output code appears several times in the data; it is not a good test to input exactly the Nyquist frequency since only two output codes (ideally) will be produced by the ADC.

The system noise being measured consists of two components: analog noise and quantization noise. Digital processing also adds a noise component, but this will be ignored since it is insignificant if the processing is done with enough bits. Analog component blocks, such as signal conditioning and sample-and-hold circuitry, create the analog noise while the ADC creates the quantization noise. An often neglected source of analog noise is the user's signal. Without special processing the best SNR that can be hoped for is the SNR at the instrument input. Fig. 1 demonstrates how the number of effective bits is determined by both sources of noise. The x axis represents the rms noise at the input of the ADC relative to the ADC's voltage range (also expressed in dB). Peak-to-peak noise is the rms noise multiplied by five. Noise is shown increasing to the right. The




Fig. 2. Linear model of digitizing system showing the sources of noise.

y axis displays the effective bits for ideal ADCs of four to eight bits. On the horizontal portions of the curves, system performance is limited by the ADC, while on the slanted sections, the limitation is analog noise. All the curves overlap when analog noise is dominant. This indicates that extra bits of resolution are of no value if they are only measuring noise.

The overall signal-to-noise ratio, rather than the number of bits, is the important parameter. The effective resolution and measurement accuracy of a system depend on this value. More ADC bits may give a prettier display, but may be meaningless as well as costly.

Improving SNR by Filtering

The seven and eight-bit filters in the HP 54111D Oscilloscope are low-pass filters that allow the user to trade off excess bandwidth for improved noise performance. Although this is often done in analog systems, digitization adds complications.

How do spectrum analyzers achieve a high signal-to-noise ratio? By limiting the bandwidth during a measurement, most of the noise is filtered out, allowing the detection of small signals. For a wideband system, such as an oscilloscope, the major analog noise contributor can be assumed to be white in nature, that is, every frequency band of a given width contains the same noise power. This is reflected in the typical specification for noise, which is expressed as volts per square-root hertz, or in terms of power as watts per hertz. The noise is present at all frequencies within the bandwidth of the system, whether or not a signal actually exists. For a real-time sampled system, the maximum bandwidth is the Nyquist frequency f_n . If the full system bandwidth is not needed for a signal, then the noise can be reduced by limiting the bandwidth, thereby increasing the SNR. If the noise is white, there is a direct correlation between bandwidth reduction and noise reduction. SNR improvement by a factor of 4 gives one effective bit improvement; that is, filtering out 75 percent of the noise improves the performance by one effective bit.

The linear model for the system is shown in Fig. 2. Analog noise is added to the signal by the preamplifier, then quantization noise is added by the ADC. The quantization noise may be a function of the input, but the model is still linear. Superposition can be used to create the model shown in Fig. 3. Each component of the digitized signal can be analyzed easily. Assume that the filter is a perfect low-pass filter at one fourth of the Nyquist frequency. If the input is within the bandwidth of the filter, it passes through the




Fig. 3. Superposition applied to the serial model of Fig. 2 results in this parallel model of an ADC system with noise sources.

Dithering in the HP 54111D

In general, dithering is the addition of a small signal, possibly noise, to a system to improve measurement accuracy. For example, many DVMS average a number of samples to give a more accurate measurement. Random noise helps the averaging process work, but a controlled signal provides better accuracy with fewer samples.

Most ADCs have reference inputs that allow the user to determine the voltage range for conversions. All of the output codes are referenced to this voltage range. If the references are raised by one volt, then the meanings of all of the output codes are also raised by one volt. That is, it takes one more volt on the input to get a particular code. Alternatively, one volt can be subtracted from the signal before the ADC to achieve the same effect. To get a particular code, the input signal must be one volt higher. The ADC acts as a buffer with differential inputs. The input minus the reference determines the output code. Adding a signal to the user's signal can be accomplished by either summing the two signals before the ADC or by subtracting the signal from the ADC references. Therefore, the dither signal can be injected by subtracting it from the ADC references.

The dither signal desired has increments of LSB/4 between samples to achieve an equivalent resolution of eight bits. With LSB/4 steps, the period of the dither signal should be four samples. The HP 54111D uses four ADCs sampling 90 degrees out of phase to achieve a 4x improvement in sample rate. Therefore, if the dither works correctly, each ADC will "see" the same point of the dither signal each time it samples. Since the period of the dither signal is $4/f_s$, all frequencies within the signal are at multiples of $f_s/4$. Each ADC is sampling at $f_s/4$, so the dither signal spectrum appears at multiples of the individual ADC sampling frequency. All of these frequencies will alias to dc. To each of the ADCs, the dither signal appears to be a dc signal. Then the desired result can be obtained by adding the appropriate dc signal to the references of each ADC. This gives a stable, accurate dither signal for optimum performance.

To prevent increasing the rms quantization noise, the dither signal must be subtracted from the ADC output code. The six-bit ADC code is augmented with two bits so that the four dither

values can be represented. With the final eight-bit code, the mean quantization and rms errors remain unchanged.

The order of the dither values is almost as important as the step size. The pattern should give consistent results whether an edge is rising or falling. The frequency content should also be spread out for the greatest attenuation by the filters (only two frequencies, f_n and $f_n/2$, are available with the given period). The dither pattern at the ADC reference that best fits this requirement is: 0, 1/2, 1/4, 3/4. This gives dither bits of 00, 10, 01, 11. Note that the larger part of the signal, the MSB, is changing at f_n while the LSB is changing at $f_n/2$.

Dithering improves all modes of operation in the HP 54111D. The seven-bit and eight-bit modes are described in the accompanying article. In the six-bit (real-time) mode, a nonlinear smoothing filter improves the accuracy of low-slew-rate signals with only a minimal effect on fast edges. The filter averages four samples (two before and one after the current sample). If the average is within LSB/2 (six-bit) of the current sample, then the current output is the average. Otherwise the current output is the current sample changed toward the average by LSB/2. Although the smoothing filter would work with random noise, the quality of the dither signal improves the performance.

In repetitive mode, dithering improves averaging by providing a resolution of LSB/4 with repeated acquisitions. Persistence is the mode with the least improvement. Dithering allows all levels of the screen to be filled in and increases the accuracy of some measurements made by the user. For example, an eye pattern is a multivalued function, so averaging mode does not work correctly. When the user measures a portion of the waveform, the dithering provides LSB/4 resolution.

Dithering could be carried further, providing LSB/8 or LSB/16 resolution with a longer signal period. However, since real ADCs are not perfect, a practical limit exists for any ADC. The ADCs in the HP 54111D are tested to LSB/4 accuracy, or one LSB at eight bits. With this specification, dithering to LSB/4 resolution is an improvement over LSB/2, but LSB/8 would not provide a significant advantage (particularly since the HP 54111D's four-ADC scheme accomplishes LSB/4 dithering so easily).

filter untouched. The analog noise is white, so it is reduced by a factor of 4.

For many signals, quantization noise will appear to be white noise, uncorrelated with the input. For the moment, assume that this is the case. Then the filtering process will reduce the quantization noise by a factor of 4. Since both noise components have been reduced by 4, the SNR has been improved by a factor of 4, or one effective bit. Similarly, the eight-bit filter has a bandwidth of one sixteenth of the Nyquist frequency and improves the SNR by a factor of 16, or two effective bits.

Statistically, this can be explained as follows. The filter is an FIR (finite impulse response) filter. The filter performs a weighted moving average on the input to find the output. Let the filter have N coefficients, where N is odd, and let $M = (N - 1)/2$. The filter coefficients are h_i , where $i = -M, \dots, M$. To achieve unity gain at dc:

$$\sum_{i=-M}^M h_i = 1 \quad (8)$$

The filtering equation for a symmetrical filter is:

$$y_k = \sum_{i=-M}^M h_i x_{k+i} \quad (9)$$

A simple example of this would be the filter defined by the coefficients 0.25, 0.50, 0.25. With input x and output y, the value for y_k is

$$y_k = (1/4)x_{k-1} + (1/2)x_k + (1/4)x_{k+1} \quad (10)$$

If the filter input x is the noise, the noise at the output is:

$$\begin{aligned} s_y^2 &= \text{var}(y_k) = \text{var} \left[\sum_{i=-M}^M h_i x_{k+i} \right] \\ &= \sum_{i=-M}^M h_i^2 s_i^2 + 2 \sum_{j=i+1}^M h_i h_j [\text{cov}(s_i, s_j)] \end{aligned} \quad (11)$$

where var is variance and cov is covariance. The noise characteristic for all samples is identical, so $s_i = s_{in}$. If the noise is white, all noise samples are independent, so $cov(s_i, s_j) = 0$. The equation reduces to:

$$s_y^2 = \sum_{i=-M}^M h_i^2 s_{in}^2 = s_{in}^2 \sum_{i=-M}^M h_i^2 \quad (12)$$

The output noise power is the input noise power times the sum of the filter coefficients squared. To gain an effective bit, this sum must be less than 0.25. This is not difficult to achieve since all of the coefficients have a magnitude less than one and most are positive. The filter given above reduces the noise to $(1/4)^2 + (1/2)^2 + (1/4)^2 = 37.5$ percent of its original value. However, the filter has poor passband characteristics so that the signal is also reduced at all frequencies but dc; the SNR is improved by 8/3 only at dc.

When is the overall system noise white? This occurs when either white analog noise or white quantization noise dominates. With either of these cases, the improvement in SNR is straightforward. Quantization noise is generally

considered to be white whenever the signal changes by at least a few quantization levels between samples. In this case, the error from one sample to the next will take on a random pattern and the error signal will appear to be uncorrelated. In general, the more bits in the ADC, the higher the signal frequency, and the larger the signal amplitude, the more random the quantization noise.

Subharmonics can cause problems creating white noise, depending upon the order of the subharmonic and the filter length. For example, an input at one eighth the sampling frequency f_s will give at most eight different codes (in the absence of analog noise). The quantization error for every eighth sample is correlated. If the filter has a length longer than eight, then the covariance term in the noise equation is nonzero. The worst-case signal for generating white quantization noise is dc with no analog noise. For this signal, the ADC outputs the same code for each sample and the quantization error is exactly the same (this is the degenerative case of a sampling frequency subharmonic). The quantization noise signal is a dc signal—certainly a poor approximation of white noise. For a dc signal, the covariance term in equation 11 exactly cancels the increase in the SNR.

In the HP 54111D, this problem is solved by adding a dithering signal to the analog signal before the ADC (see the box on page 72). The signal is precisely designed to be three fourths of an LSB peak to peak with its entire frequency content in the stop band of the low-pass filters. Sampling takes place so that four successive conversions sample the dither signal at LSB/4 intervals. The six-bit code from the ADC is augmented with two bits to cancel the error from the dither signal. For a dc signal, four successive samples give a resolution of LSB/4 for eight bits of resolution. For example, assume that an ADC has a Q level of one volt and that the code is expressed in volts such that code N represents N volts and will be present whenever the input voltage is $\geq(N-0.5)$ volts and $<(N+0.5)$ volts. If the input voltage is 1.5V, then the ADC outputs code 2 for an error of 0.5V. If dithering is added, then analog voltages of 0, -0.5, -0.25, and -0.75 volt are added to successive samples. The voltages that the ADC sees are 1.5, 1, 1.25, and 0.75 volts so that it outputs codes 2, 1, 1, 1. Note that the average value of the samples is 1.25V. The average value of the dither signal is -0.375V, so the average value of the signal is 0.375 volts higher than indicated. That is, the measured voltage is 1.625 volts for an error of 0.125V.

This example gives the maximum error for both cases. This is a factor of 4 in voltage, 16 in power, or two effective bits! If we augment the output codes with two bits to cancel the dither, then the outputs are 2, 1.5, 1.25, and 1.75 for an average value of 1.625 volts. When the outputs are averaged (filtered), the extra two bits neither help nor hurt the resolution; the extra resolution comes from providing enough extra code thresholds over four samples to match the number of thresholds for an eight-bit ADC.

In the frequency domain, the dithering signal pushes most of the noise at dc (15/16 of it) to the dithering frequencies. Since these frequencies are in the stop band of the filters, the noise is eliminated from the output. Examination of the dithering characteristics in equation 11 shows that

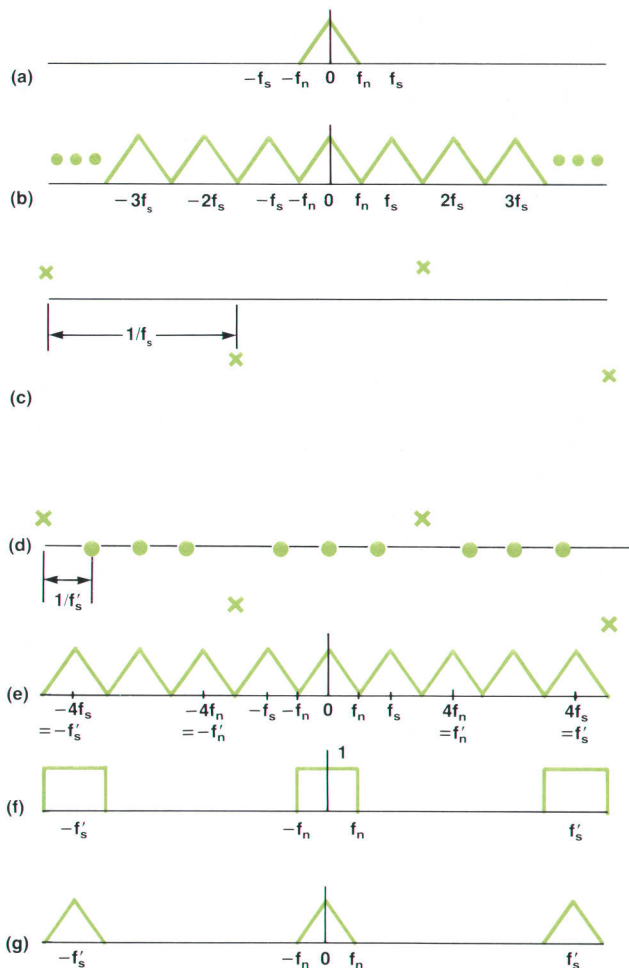


Fig. 4. Interpolation by a factor of 4. (a) Example signal spectrum before sampling. (b) Spectrum after sampling at frequency f_s . (c) Sampled data points. (d) Adding three zeros between points shown in (c). (e) New spectrum. (f) Desired low-pass filter. (g) Final result.

the covariance term is negative so that the overall improvement is better than with white noise.

Under conditions where white noise dominates the input, the dithering signal has no effect; it simply appears as a high-frequency signal that is filtered out by the low-pass filter. Hence, the dithering has taken the worst-case signal and turned it into the best-case signal without hurting the original best case. A signal at any frequency is improved in relation to the quantity of noise that is correlated.

The filters used in the HP 54111D do not come close to ideal low-pass filters (see the box on page 75). The goal was to optimize SNR improvement while maintaining a good step response. The desired SNR improvement determines the minimum stop-band width: $3f_n/4$ for the seven-bit filter and $15f_n/16$ for the eight-bit filter. The stop-band attenuation is not critical; above 20 dB it is easier to improve the SNR by widening the stop band than by increasing the attenuation. (Infinite attenuation is not much improvement when the major source of noise is in the passband.) The maximum amount of overshoot and ringing on the step response is known. These factors determine the width of the transition band and, in turn, the passband. To provide the required SNR improvement over as wide a frequency band as possible, the passband is as flat as possible. The bandwidths of the seven-bit and eight-bit filters are 100 MHz and 25 MHz at full sample rate and they provide one and two extra effective bits to frequencies of 83 MHz and 18 MHz, respectively.

The seven-bit and eight-bit filters in the HP 54111D, along with dithering, provide increases of one and two effective bits at the expense of bandwidth. In Fig. 1, the filters move the operating point at least one or two ADC lines up as a result of reducing quantization noise and one or two tics (6 or 12 dB) to the left as a result of reducing analog noise. Note that any noise in the user's signal is included whenever measurements are made and that the filters reduce this noise too. The HP 54111D has better effective-bits performance in the seven-bit filter mode than it would have with an eight-bit ADC without filters, although a penalty is exacted in lower bandwidth.

Interpolation Process

One of the most questioned functions of digitizing oscilloscopes is interpolation. Users feel uncomfortable with the "fictitious" data that is placed between samples to fill in the waveform. Most interpolators are linear functions that can be described in the same sequence of steps. This includes common filters such as $\sin(\pi x)/\pi x$ (i.e., sinc(x)), Gaussian, and Nth-order polynomials that fit the sample points exactly (LaGrange interpolation). A specific example of interpolation by a factor of four will demonstrate the steps in the frequency domain.

For simplicity, assume that the analog signal has a triangular spectrum as shown in Fig. 4a. After sampling the signal, the spectrum is repeated at intervals of f_s , the sampling frequency, as in Fig. 4b. Mathematically, any combination of the spectrums could account for the sampled waveform. The purpose of interpolation is to increase the sample density as though the signal were sampled at a higher rate. The first step in interpolation is to increase

the sample density. Zero-valued data points are inserted between the actual data points (Fig. 4c) to get the desired density. For the example, three zeros are added to increase the density by a factor of four (Fig. 4d). In the frequency domain, this simply places the new sample rate, f'_s , at $4f_s$ and the new Nyquist frequency f'_n at $2f_s$ (Fig. 4e). There are now four images of the original spectrum below the Nyquist frequency. This means that a combination of spectrums is required to produce the zero-filled signal. The SNR is 1/3, since there are three undesired spectrums to the one desired spectrum. The final step is to filter the result to leave the original spectrum at a higher effective sample rate (Figs. 4f and 4g). This low-pass filter is often a sinc function or a Gaussian function.

The most difficult concept of the process to understand is the spectrum of the zero-filled signal. A simple example may clear up the confusion. The signal

$$\frac{1}{4} + \cos(2\pi f_s t)/2 + \cos(4\pi f_s t)/4$$

can be used to modulate the original signal before it is sampled (Fig. 5a). If the signal is sampled at frequency f'_s in phase with the modulation signal, the result is identical to Fig. 4d. The spectrum of the analog signal before the digitizer is the convolution of the original signal spectrum (Fig. 5b) with the spectrum of the modulation signal (Fig. 5c), giving the "zero-filled" spectrum (Fig. 5d). After sampling, the energy from $2f_s$ to $2.5f_s$ aliases back to the range $2f_s$ to $1.5f_s$, doubling the height of the short triangle. Sampling also causes this new spectrum below $2f_s$ to be copied




Fig. 5. (a) Multiplying input signal by modulation signal. Original spectrum (b) convolved with modulation signal spectrum (c) gives a "zero-filled" spectrum (d). The result after sampling at f'_s is identical to Fig. 4e.

Digital Filters

The digital filters in the HP 54111D are FIR (finite impulse response) filters. The FIR filter has no unsampled analog counterpart since it has no poles—only zeros. One of the advantages of FIR filters is that they can easily be designed with perfectly linear phase, providing a constant delay through the filter regardless of frequency. Constant delay is important for making accurate measurements. Without constant delay, the apparent time between two waveform edges could be a function of the difference in edge speed. For example, in a digital circuit this might be a function of logic family or signal loading.

One characteristic of linear phase is that the step response is antisymmetrical; the preshoot is a mirror image of the overshoot. In an oscilloscope, this appears as noncausal behavior since the filtering routine takes the constant delay into account when placing the filtered points on screen. Although this is unsettling when first encountered, it is an advantage over an analog filter since the user can often separate the matching overshoot/preshoot response from the actual signal overshoot.

Another characteristic of the digital filters is that the bandwidth is related to the sample rate. When the sample rate changes, the bandwidth changes. Some users consider this an advantage and some a disadvantage. Generally speaking, the sample rate is not lowered until the screen shows only one or two periods of a waveform. In this situation, the lower bandwidth does not appreciably alter the signal, but still provides the increase in effective bits. For example, at 100 ns/div, the sample rate drops to 500 megasamples per second. Now the seven-bit filter has a bandwidth of 50 MHz and a rise time of approximately 7 ns. The largest change in response occurs for a perfect step, where one sample is on the bottom of the step and the next is on the top. In this case, the filtered rise time will be 7 ns, or less than a tenth of a division. The rise time for the eight-bit filter will be about 30 ns, or one third of a division.

The FIR filters achieve high throughput by using lookup tables. A table exists for each filter with each coefficient value multiplied by all possible data values. At run time, the microprocessor need only perform lookup and add operations. The sizes of the tables were minimized by taking into account properties of the filters and data. The linear phase characteristics are achieved by making the filter symmetrical about the center coefficient. This allows a reduction in size by a factor of two. The eight-bit data values are shifted to a two's complement format to give 128 nonzero magnitudes instead of 255. Altogether a data reduction of nearly four to one was accomplished. Further reduction is possible, but would lower the throughput of the instrument. Throughput could have been increased by using a math coprocessor or a digital signal processing chip, but this was not done because of cost and space considerations.

Many of the limitations to digital filters parallel limitations to analog filters. The characteristic of most concern is that the steepness of the transition region between the passband edge and the stop-band edge determines the amount of overshoot and ringing in the step response. The filters in the HP 54111D were designed with this in mind, trading off bandwidth for pulse fidelity.

at intervals of $4f_s$ (f'_s), giving a sampled spectrum identical to Fig. 4e.

The requirements for the HP 54111D's interpolation filter are more stringent than for the seven-bit and eight-bit filters. The filter can interpolate up to a factor of 100. For this case, the SNR before low-pass filtering is 1/99, or -20 dB. The filter stopband must attenuate the 99 spectral replicas to achieve an SNR of 42 dB. This minimizes the reduction in effective bits and the measurement uncertainty in both time and voltage. Since there is no anti-aliasing filter in the instrument (to allow a 500-MHz repetitive bandwidth), the stopband must begin near the Nyquist frequency of 500 MHz. To maintain good pulse fidelity, the maximum amount of overshoot and ringing is determined. As with the other filters, these quantities determine the width of the transition band and the passband. The filter is down 2 dB at 250 MHz and 48 dB at 515 MHz with only 6% overshoot.

The sinc(x) and Gaussian filters often used are extreme types of filters. The sinc(x) filter gives the maximum bandwidth (f_n), but has a poor step response. The Gaussian filter adds no overshoot to the signal, but requires a very wide transition band, giving an unnecessarily narrow passband. The HP 54111D's interpolation filter is somewhere between these two extremes. One advantage to having a bandwidth of $f_s/4$ rather than $f_s/2$ is that the transition region provides some anti-aliasing protection. Any signals in the 500-to-750-MHz range will alias to the 250-to-500-MHz region. The attenuation of the transition band will perform an anti-aliasing function for this frequency range. Beyond 750 MHz the attenuation of the analog signal path provides aliasing protection.

A side effect of having a bandwidth lower than the Nyquist frequency is that filtered points do not always overlay the actual sampled points. Many users believe that this is a shortcoming of the interpolation, although it is a necessary phenomenon. Any movement in samples results from reduced bandwidth after interpolation; no one would expect the bandwidth-limited trace on an analog oscilloscope to look like the full-bandwidth trace. To leave the actual samples in place, the filter's impulse response must have the exact zero crossings of the sinc(x) filter and its characteristics will approach those of the sinc(x) filter. It is very limiting on the filter design and fidelity to enforce such a requirement.

Interpolation is not a difficult process, but must be done with care. The result is consistent with the response of the interpolation filter so long as aliasing is not a concern.

Summary

Digitizing a waveform includes two separate effects: sampling and format conversion. Sampling is an analog process that transports the signal from the s-plane to the z-plane, resulting in the possibility of aliasing. The format conversion simply changes the signal to a form that can be handled by a computer, adding noise in the process. If aliasing can be ignored, the digitized data contains all of the information necessary to reconstruct the signal, with a little added noise. For any measurement system, the signal-to-noise ratio is an important parameter. With digitizing instruments, this is more important than the actual number of

bits or the resolution of the display. Since digitized data is merely a different representation of a waveform, the same techniques that have been used to improve the SNR in analog systems can be used with digital systems. The seven-bit and eight-bit filters are digital equivalents of the bandwidth-limit button on traditional analog oscilloscopes.

A function that has often been questioned by digitizing oscilloscope users is interpolation. Waveform fidelity is highly dependent upon the quality of the interpolation filter. There is a direct trade-off between the bandwidth-to-sampling-rate ratio and the quality of interpolation. The HP 54111D does not present the best filter for all types of signals, but it does provide a high-bandwidth, good-quality filter that works very well for the types of signals viewed most often on oscilloscopes—sine waves, pulse trains, and logic signals.

The digital filters in the HP 54111D provide benefits that are not easily available in analog counterparts. First, the filters have no delay at any frequency and timing measurements are not altered because of nonlinear phase characteristics. Second, a stable system response is achieved, not only from unit to unit, but with time and temperature.

Address Correction Requested
Hewlett-Packard Company, 3200 Hillview
Avenue, Palo Alto, California 94304

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

HEWLETT-PACKARD JOURNAL

June 1988 Volume 39 • Number 3

**Technical Information from the Laboratories of
Hewlett-Packard Company**

Hewlett-Packard Company, 3200 Hillview Avenue
Palo Alto, California 94304 U.S.A.

Hewlett-Packard Central Mailing Department
P.O. Box 529, Startbaan 16

1180 AM Amstelveen, The Netherlands

Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan
Hewlett-Packard (Canada) Ltd.

6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

CHANGE OF ADDRESS: To subscribe, change your address, or delete your name from our mailing list, send your request to Hewlett-Packard Journal, 3200 Hillview Avenue, Palo Alto, CA 94304 U.S.A. Include your old address label, if any. Allow 60 days.