

THE FEBRUARY 1981
VOL. 60, NO. 2



**BELL SYSTEM
TECHNICAL JOURNAL**

ISSN0005-8580

R. J. Canniff	A Digital Concentrator for the SLC™ -96 System	121
J. F. Reiser	Compiling Three-Address Code for C Programs	159
A. S. Acampora	Rain Margin Improvement Using Resource Sharing in 12-GHz Satellite Downlinks	167
L. J. Greenstein and B. A. Czekaj	Modeling Multipath Fading Responses Using Multi-Tone Probing Signals and Polynomial Approximation	193
B. Hoadley	The Quality Measurement Plan (QMP)	215
R. D. Gitlin and S. B. Weinstein	Fractionally Spaced Equalization: An Improved Digital Transversal Equalizer	275
	Contributors to This Issue	297
	Papers by Bell Laboratories Authors	299
	Contents, March 1981 Issue	301

THE BELL SYSTEM TECHNICAL JOURNAL

ADVISORY BOARD

D. E. PROCKNOW, *President*, *Western Electric Company*
I. M. ROSS, *President*, *Bell Telephone Laboratories, Incorporated*
W. M. ELLINGHAUS, *President*, *American Telephone and Telegraph Company*

EDITORIAL COMMITTEE

A. A. PENZIAS, *Chairman*
A. G. CHYNOWETH W. B. SMITH
R. P. CLAGETT G. SPIRO
T. H. CROWLEY J. W. TIMKO
I. DORROS I. WELBER
R. A. KELLEY M. P. WILSON

EDITORIAL STAFF

G. E. SCHINDLER, JR., *Editor*
PIERCE WHEELER, *Associate Editor*
JEAN G. CHEE, *Assistant Editor*
H. M. PURVIANCE, *Art Editor*
B. G. GRUBER, *Circulation*

THE BELL SYSTEM TECHNICAL JOURNAL is published monthly, except for the May-June and July-August combined issues, by the American Telephone and Telegraph Company, C. L. Brown, Chairman and Chief Executive Officer; W. M. Ellinghaus, President; V. A. Dwyer, Vice President and Treasurer; F. A. Hutson, Jr., Secretary. Editorial inquiries should be addressed to the Editor, The Bell System Technical Journal, Bell Laboratories, 600 Mountain Ave., Murray Hill, N.J. 07974. Checks for subscriptions should be made payable to The Bell System Technical Journal and should be addressed to Bell Laboratories, Circulation Group, Whippany Road, Whippany, N.J. 07981. Subscriptions \$20.00 per year; single copies \$2.00 each. Foreign postage \$1.00 per year; 15 cents per copy. Printed in U.S.A. Second-class postage paid at New Providence, New Jersey 07974 and additional mailing offices.

© 1981 American Telephone and Telegraph Company

Single copies of material from this issue of The Bell System Technical Journal may be reproduced for personal, noncommercial use. Permission to make multiple copies must be obtained from the editor.

Comments on the technical content of any article or brief are welcome. These and other editorial inquiries should be addressed to the Editor, The Bell System Technical Journal, Bell Laboratories, 600 Mountain Avenue, Murray Hill, N.J., 07974. Comments and inquiries, whether or not published, shall not be regarded as confidential or otherwise restricted in use and will become the property of the American Telephone and Telegraph Company. Comments selected for publication may be edited for brevity, subject to author approval.

THE BELL SYSTEM TECHNICAL JOURNAL

DEVOTED TO THE SCIENTIFIC AND ENGINEERING
ASPECTS OF ELECTRICAL COMMUNICATION

Volume 60

February 1981

Number 2

Copyright © 1981 American Telephone and Telegraph Company. Printed in U.S.A.

A Digital Concentrator for the SLCTM-96 System

By R. J. CANNIFF

(Manuscript received April 22, 1980)

The SLCTM-96 subscriber loop carrier system is a digital subscriber carrier system serving up to 96 single-party customers. The system can be configured with an optional plug-in which digitally concentrates two standard, 1.544 megabit, serial pulse-code modulation (PCM) bit streams into a single stream, thereby concentrating 48 customers onto a single T1 digital line. The concentrator employs a custom NMOS LSI chip providing a full-access, time-slot interchange function. It has microcomputer controllers at the two ends of the system to control time-slot assignments. The development of the concentrator involved challenges in chip design, software design, and performance testing.

I. INTRODUCTION

As the state of the art in electronics has advanced, so have the inroads of electronics into the loop plant as a supplement to or replacement for cable.¹ The application of pair-gain systems to loops has been of particular importance in recent years. The SLCTM-96 subscriber loop carrier systems is the latest in a line of digital loop carrier systems to provide improved features and reduced cost.² The SLC-96 system serves up to 96 single-party subscribers between a Central Office Terminal (COT) and Remote Terminal (RT) using standard pulse-code modulation (PCM) coding over facilities such as T1. System features include a variety of channel units, channel and drop testing provisions, and spare digital line switching. System versatility is further enhanced by three modes of operation (each mode uses one additional T1 line for protection):

Mode I: Carrier only; 96 full-time channels on four working T1 lines.

Mode II: Carrier concentrator; 96 channels on two working T1 lines.

Mode III: Special services only; 48 special service channels on two working T1 lines.

This article discusses the Time Assignment Unit (TAU) which is the concentrator employed in Mode II operation. In a fully equipped Mode II system there are actually two identical, but independent, concentrators, each concentrating 48 lines onto 24 time-slots of a T1 line, and thereby eliminating two main T1 lines (see Fig. 1). In many cases the saved T1 lines provide more advantage than is apparent. The benefits arise in applications limited by apparatus case size for holding repeaters, by small cables where the purpose of pair-gain systems is to avoid adding new cable, and, similarly, in situations that use several systems in parallel along the same route where the number of saved pairs can be very significant. For long-distance systems, of course, the saved lines result in important cost savings.

Section II of this article discusses the traffic-handling ability of the TAU and the traffic administration provisions. This is followed by a functional description of the concentrator operation in Section III. Section IV discusses first the hardware features provided by the custom Time-Slot Interchange (TSI) chip and then enumerates the features provided by software. Selected implementation details for both hardware and software are expanded upon in Section V. Performance testing, that is, the equipment and resources that were used to adequately stress the concentrator to confirm its design, is discussed in Section VI.

II. TRAFFIC CONSIDERATIONS

The concentrator performs a full-access, digital time-slot interchange function. The very conservative two-to-one concentration ratio that was chosen provides for the inclusion of a limited number of dedicated (unconcentrated) special-service circuits as well as a large number of multiparty circuits. It is important to note that a single channel unit for special service occupies the same physical space as a standard unit for dual voice frequency and thus reduces the number of lines to be concentrated by two. Also, the special-service unit requires only a single T1 time slot. Thus, regardless of how many special-service units are inserted, the concentration ratio remains fixed at two to one, but the group size is reduced. The blocking probability for concentrated lines increases as more special-service units are added, because of the reduced group size. By design, the number of special-service channel units is limited to eight and must be physically inserted in the last four channel unit positions of each shelf.

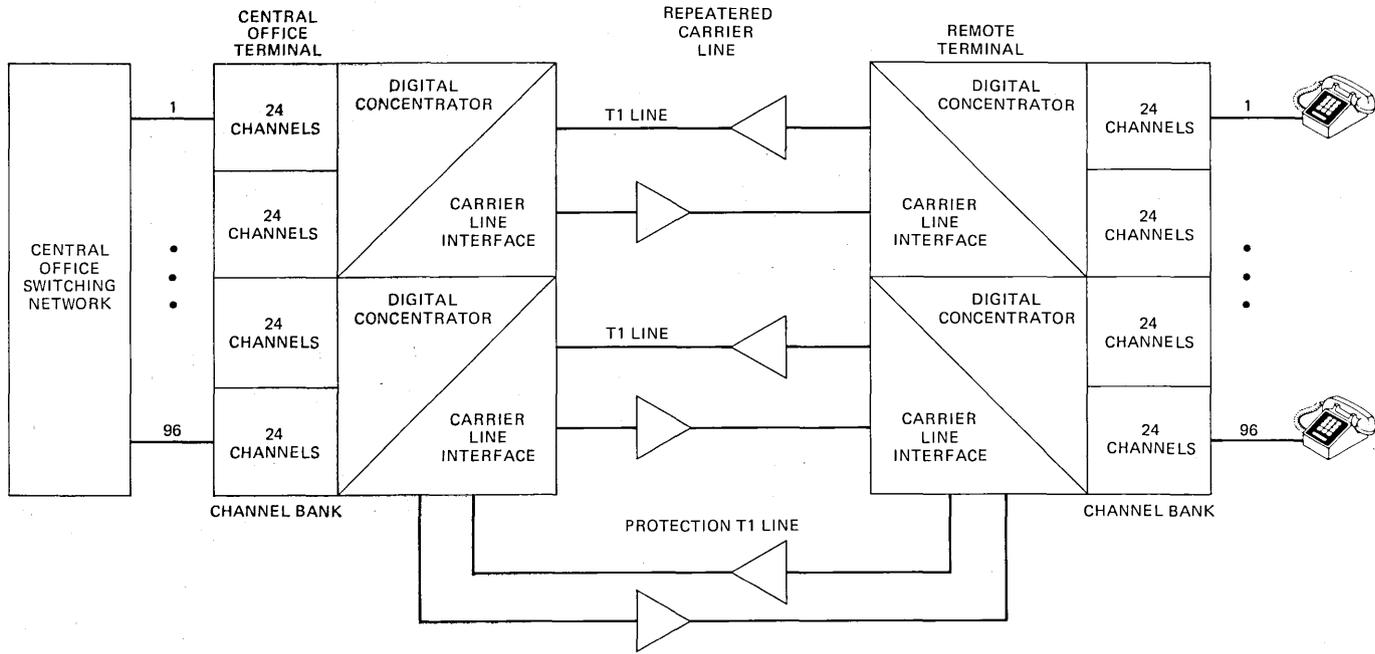


Fig. 1—The SLCTM-96 system-carrier concentrator configuration—Mode 2.

For heavy loading or mandatory monitoring purposes, six traffic administration aids are provided:

(i) A two-digit numeric display is provided on the faceplate of the COT TAU which displays, upon demand, the peak traffic in hundred call seconds per hour (CCS), per concentrated T1 time slot since the internal traffic register was last cleared. The register is cleared by means of a pin jack on the faceplate.

(ii) The same display will also indicate the number of blocked calls, up to a maximum of 15, since the internal blocked calls register was last cleared. This register is cleared by means of the pin jack.

(iii) A traffic lamp is provided on the faceplate, in conjunction with a minor alarm, that indicates there has been two or more blocked calls for two out of three weeks running. This alarm must be manually cleared through the pin jack.

(iv) A relay closure is provided to remote to an electronic switching system (ESS) office that all T1 time slots are in use. The office can then divert terminating calls and provide its own reorder tone while keeping individual line blockage statistics. (In a non-ESS Central Office this diversion will not occur and the TAU will provide a digital reorder tone.)

(v) A second relay is provided that outpulses peak weekly traffic in CCS, once per week, to a remote traffic-monitoring register at the rate of one pulse per second.

(vi) A third relay is provided that outpulses the number of blocked calls as they occur to a remote blocked calls monitoring register at the rate of 1 pulse per second. There is no saturation limit here as there was in the faceplate display.

III. TAU FUNCTIONAL DESCRIPTION

The following definitions are given to clarify all discussion that follows: A line refers to a subscriber loop at the RT or a wire pair appearance at the COT. A channel unit is the physical plug-in serving one or two lines and providing the per-line circuit functions. A channel is the electrical path from a channel unit to the Transmit Receive Unit (TRU) serving it, and the time slot reserved for the line into and out of the TRU on the 1.544 megabit serial PCM busses. Channels enter or leave the concentrator as time slots from or to the TRUs. Time slots on the T1 line interfacing to the concentrator are referred to as trunks, because they provide a limited number of shared paths between the two terminals of the system, in analogy with traditional trunking facilities between central offices.

3.1 System block diagram

Figure 2 shows a simplified system block diagram of the TAUs. The microcomputer controllers are realized using the Bell Laboratories

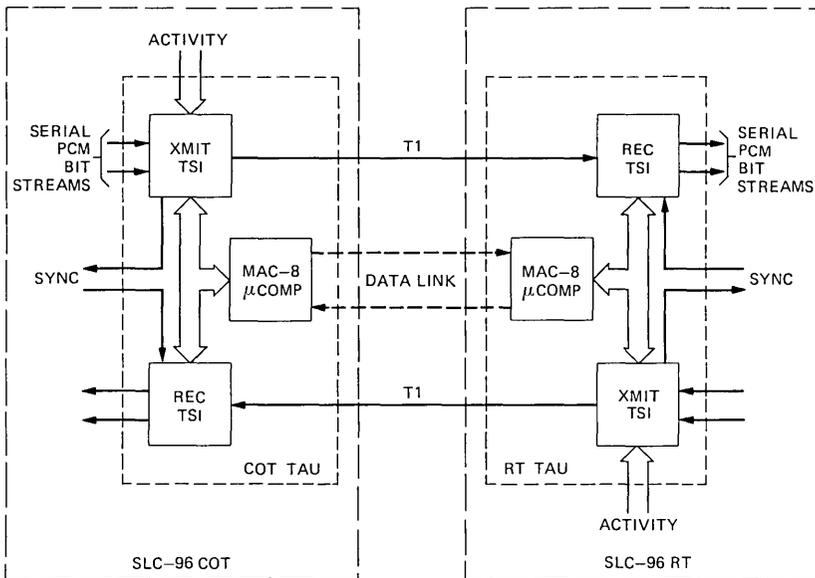


Fig. 2—TAU system block diagram.

designed MAC-8 microprocessor.³ The MAC-8 microcomputers maintain control over the Time-Slot Interchange (TSI) chips and talk to each other over a data-link channel derived from the subframe bits on the T1 line. The transmit TSI selectively combines two 1.544 megabit, serial PCM bit streams, one from each of the TRUs it serves. (A TRU performs the A/D, D/A, and framing tasks for 24 channels.) The resulting T1 signal is sent to the receive TSI at the other end of the system where it is expanded into two streams, for sending to corresponding TRUs. The TSIs provide full access, meaning that any one of the incoming (outgoing) 48 lines or time slots has access to any one of the 24 outgoing (incoming) trunks, where trunks are time slots on the T1 line. The COT TAU handles trunk assignments and, in general, controls the concentrator. The RT TAU acts more as a slave. When the MAC-8 controller assigns a line to a specific trunk, that line will keep the trunk for the duration of the call and in no way inhibits any of the other lines from accessing any of the other trunks.

A line-service request, called "activity," is picked up by the transmit TSI by accessing the A and B bit signaling busses on the system backplane. A and B bits are the standard nomenclature for per-channel signaling that indicates off-hook, ringing, etc. The activity is stored in memory in the TSI from which it can be retrieved by the MAC-8 through the TSI microcomputer address and data ports. Activity at the RT is passed over the data link to the COT where the line/trunk assignments are determined.

The transmit and receive TSIs must be synchronized to the TRUS they serve and appropriate signals are provided for this purpose. No synchronization is assumed between transmit and receive TSIs. Further, no synchronization is assumed between the MAC-8, the TSIs, and the data link. The TSIs are accessed by means of a handshake procedure. The data-link frame signals are polled to determine when message processing is needed.

If all 24 trunks are busy, provision is needed for feeding a fast-busy (overflow) tone to the COT channel units. This is done digitally through the receive TSI, resulting in significant cost savings in the channel units. The receive TSI allows the assignment of up to 24 lines to "busy trunks" whereby the selected lines receive a fast-busy tone in PCM form as read from a code-word table stored in the MAC-8 program memory. Signaling information (A and B bits) is also stored in these codes and thereby allows the channel unit to recognize that it is getting the fast-busy tone and accordingly trip ringing without charge, prior to applying the tone on the blocked customer's line.

3.2 Circuit block diagram

Figure 3 shows a simplified schematic for the COT TAU. The RT TAU is nearly identical with the basic exceptions that Output Port 2 is removed and there is no Read Only Memory 2 (ROM2). The Random Access Memory (RAM), Read Only Memory (ROM), Input Port, and Output Ports connect to the MAC-8 bus as in any normal microcomputer. The custom TSI chips were also designed to connect directly on the bus and appear to the MAC-8 as programmable peripheral chips. The MAC-8 talks to the TSI chips by means of a handshake procedure

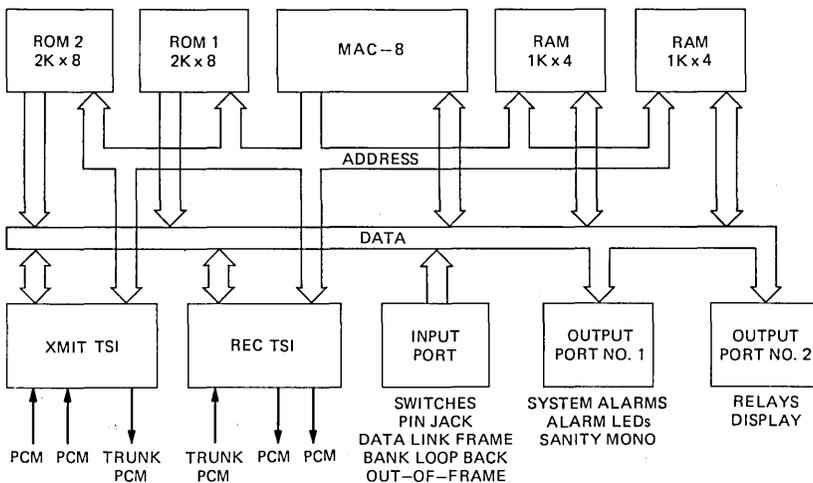


Fig. 3—COT TAU schematic.

using the MAC-8 Data Ready line. The clock frequency for the TSI and MAC-8 is 1.544 MHz. Three to nine wait states result, as required and generated by the TSI on a read or write. The TSI chip generates and responds to the necessary signals for system synchronization.

Output Port 1 provides four alarm light emitting diodes (LEDs) on the faceplate and three system alarms. For the COT TAU, the LEDs indicate alarms for COT, RT, traffic, and special-service channels. The RT TAU has only two LEDs indicating RT alarm and unit alarm for special-service channels. The system alarms are minor, major, and major with channel unit disable. Output Port 1 also provides an output for strobing the sanity monostable, a timeout device that checks on proper sequencing of the program. Output Port 2 provides the dual-digit numeric display for indicating traffic and blocked calls on the faceplate and also services the three relays used for remoting concentrator status. The Input Port provides for the display selection switches and the internal register-clearing pin jack, all mounted on the faceplate. It further provides inputs for the data-link frame signals which are polled to determine data-link message requests. Also, the Input Port allows accessing a receiver out-of-frame signal and a bank loop-back signal.

The COT TAU has 4K bytes of program memory and the RT TAU has 2K bytes. Each are realized using 2K byte ROM chips. At both COT and RT, 1K bytes of RAM are provided, though the COT uses less than one-third of the available memory and the RT uses only about one-fourth. The TAU plug-ins are printed circuit cards measuring approximately 4 inches by 10 inches as pictured in Fig. 4. It was required that the TAUs be sized to physically replace the Line Interface Units for the T1 lines that are not needed in the concentrated mode. Power supplies for the TAU are 5V and 12V, with typical dissipation being about 3 watts. Special requirements had to be met for the RT TAU so that it could work over the temperature range of -40 to $+85$ degrees Celsius.

IV. HARDWARE AND SOFTWARE FEATURES

4.1 Features of the TSI

The TSI chip was designed to be universal, in the sense that it is package lead and microcomputer programmable for use at COT or RT, for either the transmit or receive function. In making a universal chip with all the features mentioned below, it is possible, and very desirable, to reuse pieces of the hardware inside the chip. Thus, for example, pieces of RAM and other hardware used for busy trunk assignments in the receive TSI mode are alternatively used in the transmit TSI mode for the activity and "TNEN" collection (discussed later). Similarly, many input and output package leads serve dual purposes. In addition to the basic features, several very important additional features are

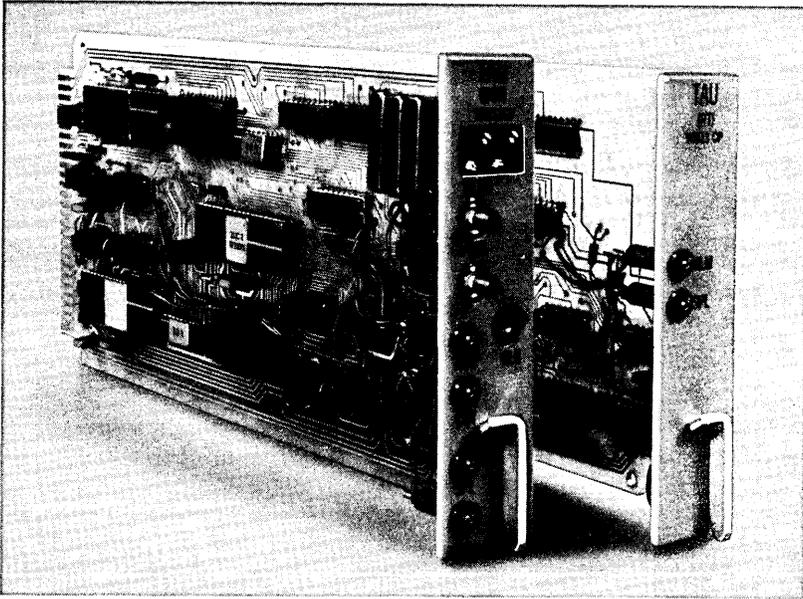


Fig. 4—TAU circuit packs.

provided by the TSI chip. Most of these features cost very little hardware, that is, they increase chip complexity very little, while providing useful features.

Figure 5 is a block diagram of the TSI chip. The heart of the chip is a RAM used for PCM data, trunk assignment information, busy trunk-assignment information, activity data, "TNEN" data, per-line test bits, 1 byte of fast-busy tone buffer, and 2 bytes of ROM. Incoming and outgoing serial PCM data is handled in bytes inside the chip by feeding the data through serial-to-parallel registers (Registers 1 and 2) and parallel-to-serial registers (Registers 2 and 3). The use of Register 2 depends on whether the TSI is performing in the transmit or receive mode. The frame bit is normally stored in the FR FF (frame flip-flop). The main-control logic provides the signals for controlling all the chip's registers, multiplexers, and RAM. Multiplexers are included for selecting address and data for the RAM and for selecting output onto the microcomputer data bus.

4.1.1 Interchanging time slots

The time-slot interchange is executed by the method of writing data into a RAM sequentially and reading it selectively for the transmit concentrator function and vice versa for the receive function. The trunk assignments (24 bytes), which are addresses of the desired PCM line data, are stored in the RAM, yet are used to address the RAM

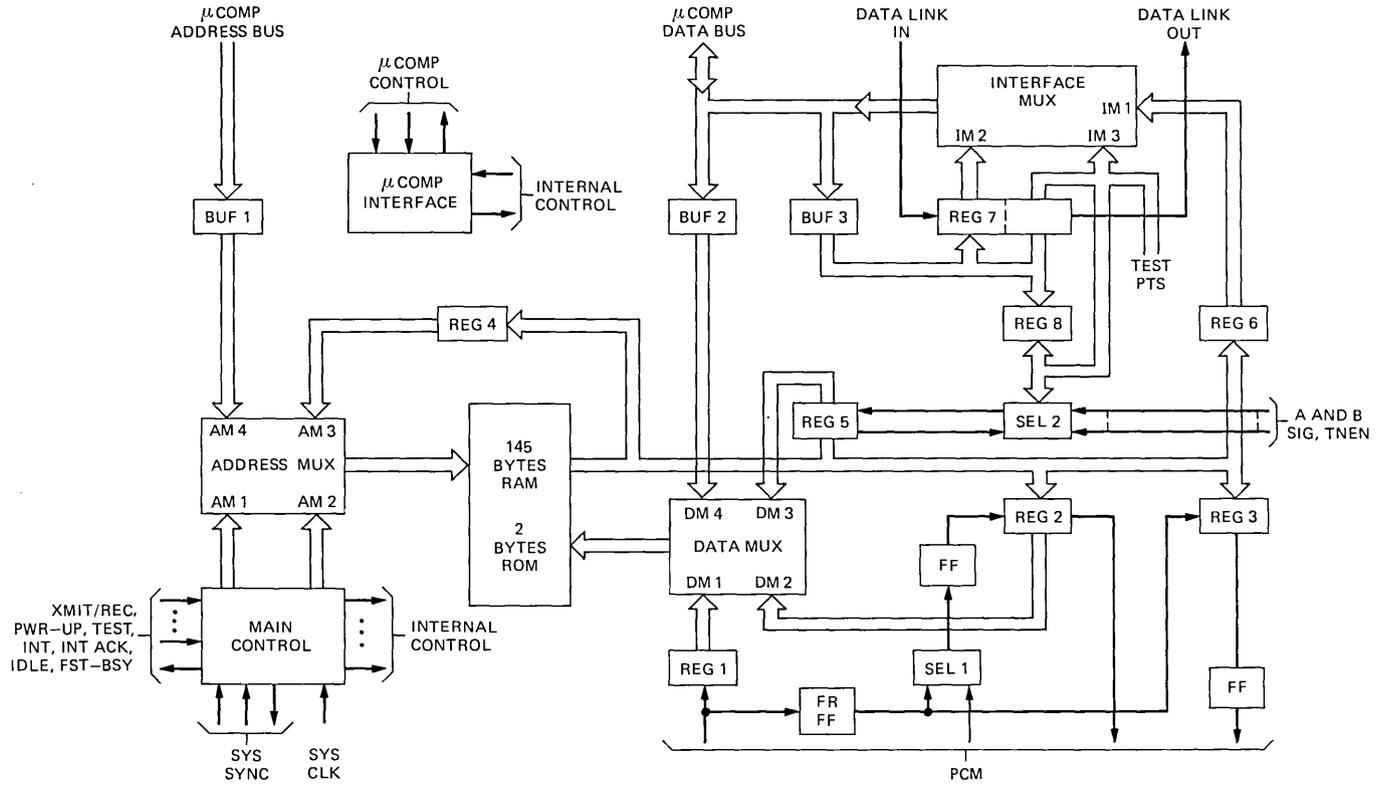


Fig. 5—Time-slot interchange chip.

selectively by a feedback register (Register 4) to the address bus. Each of the trunk (busy trunk) assignments are fed back in turn to address the desired line memory. In the transmit mode, if a trunk is unassigned, the ROM locations are accessed to send idle code on the T1 line. In the receive mode, if a line is unassigned, idle code from the PCM memory is read out for the line or idle is forced by means of a cleared "enable bit" (discussed later).

Since the concentrator is positioned between the TRUS and the T1 line, A and B bit-signaling information is already contained in the PCM bit streams when they are received by the TSI. It is necessary, therefore, that the TSI insure integrity of signaling frames. Signaling frames are the 6th and 12th frames of a 12-frame sequence and contain the per-channel A and B bit signaling information in the least significant bit of the PCM code words. Integrity is maintained by having two RAM sections of PCM data (48 bytes each). While writing (reading) in one section sequentially, the other section is read (written) selectively; the roles are reversed every frame. The frame bit is correspondingly delayed to match up with the outgoing data. The result is that signaling-frame integrity is maintained while the data experiences a fixed delay of approximately two frames (250 microseconds) end-to-end regardless of the line or trunk.

4.1.2 Microcomputer access

The microcomputer has access to the TSI memory through the address and data ports of the TSI chip (see Fig. 5). This allows all the memory locations, including the PCM data memory, to be written and read. The main control block of the TSI contains a frame counter for controlling all chip sequencing. When the counter is in a state not needed for a specific internal function, that clock cycle can be used to respond to a microcomputer read or write request. The result is a variable number of wait states, as mentioned earlier, because the microprocessor can request a TSI access at an arbitrary time. The microcomputer talks to the TSI by means of Read/Write, Select, and Busy leads that connect to the microcomputer interface circuitry. On a read operation the data is held in Register 6 so that the processor can use as much time as necessary to recognize the data. Some of the address space of the TSI is reserved for addressing the data-link register (Register 7) and activity-mode control register (Register 8).

4.1.3 System synchronization

The TSI must be synchronized to the TRUS to know where every bit in the incoming and outgoing bit streams is located. For the transmit function, the TSI puts out a superframe synchronization signal which the TRUS can accept and lock to. For the receive function the TSI

accepts an out-of-frame and synchronization signal from the TRU. The receive function is complicated, however, by the reframing process. Since the TSI exists on the T1 line side of the TRU, the reframing executed by the TRU is carried out after the bit stream has passed through the TSI. This implies the need for special modes of operation while the system is out-of-frame to assure reframe.

4.1.4 Digitally generated, fast-busy tone

A feature of the TSI in the receive mode is the ability to feed a digitally generated, fast-busy tone to a terminating connection when all trunks are busy. In order to receive the fast-busy tone, a line must be assigned to a "busy trunk." The assignment mechanism is nearly identical to assigning a line to a normal trunk. The assignment consists of an address of the line location into which the fast-busy tone bytes are to be placed. These assignments are fed back in turn through Register 4 to select the desired lines. Up to 24 lines can receive the tone byte simultaneously. The source of the current PCM tone byte is a holding register (Register 5) which is indirectly updated every frame by the microcomputer, through the fast-busy tone buffer byte in RAM. Two signaling codes ($A=B=0$ or $A=B=1$) can be sent out with the tone byte by making the least significant bit always 0 or 1. Only one of these codes is used in the TAU for signaling the channel units. The fast-busy tone is simulated by a sequence of 48 PCM coded bytes stored in the program ROM which emulates the dual-tone frequency needed.

4.1.5 A and B signaling bit collection

One very important side feature of the TSI in the transmit mode is that it gathers A and B bit information to supply the microcomputer with line activity information. The simplest and fastest way to collect this data is to tap into the A and B bit busses on the system backplane. The A and B bit data are available there every frame and, because the TSI and the TRU are synchronized, the precise time for each line's A and B bit data is known.

Since the concentrator is interested only in activity and not in the precise A and B bit signaling states, the A and B information is condensed. No activity (idle) is signaled by the channel units as $A = B = 0$ or $A = B = 1$, depending on whether the location is *COT* or *RT*. Thus, the A and B bit collection hardware just looks for a deviation from the idle pattern. It is further desirable to have an elementary filtering effect so that if there is any activity within a certain time period during which the chip is told to collect activity, that activity is caught and held, with the result that the microcomputer is not required to make numerous, closely time-spaced searches for activity.

As a result, the TSI chip offers three modes for activity collection:

sift for zeroes, sift for ones, or no collection. Since the A and B bit information is combined, all 48 lines of activity information are stored in 48 bits or six bytes of the RAM. The collection mode control is set by the microcomputer in Register 8 (see Fig. 5). Selector 2 and Register 5 are used to collect each byte of data before being transferred to the RAM for storage.

4.1.6 TNEN bit collection

Another important side feature of the TSI chip is that it gathers the so-called "TNEN" bits. The per-channel TNEN bit tells the TRU to encode PCM data for the channel or to send digital data off a backplane bus. Since two channels are associated with a single physical channel unit, the two corresponding TNEN bits can be used as an indication of the class of service desired by a channel unit. Thus, four different types of channel unit can be identified based on the permutations of the two TNEN bits, and the concentrator can then take the required action. For example, standard dual voice-frequency units get concentrated service while single voice special and single data special units get a permanent trunk.

The TSI chip in the transmit mode collects TNEN bits in almost the same way it collects activity (using Registers 5 and 8 and Selector 2), because the TNEN bits are available on the backplane bus in the same format as the A and B bits. The only difference is that there is only one TNEN bit per line and no sifting is performed, that is, they are just collected. The TSI cannot collect activity and TNEN simultaneously since they use common hardware. The TNEN bits are stored in different memory bytes in the RAM, however; thus, by means of a fourth collection mode (in addition to the three activity modes), the microcomputer can get a snapshot of the TNEN bits between normal activity collection. This can be done in just a little more than one frame of time so that no significant activity is lost.

4.1.7 Per-channel enable bits

An additional feature of the TSI, and one that has proved to be very powerful for real-time testing, is the provision of per-channel enable bits. The cost of this memory is low because it employs the upper two bits of the trunk assignment bytes which are not needed for addressing the 48 possible lines. Each of the 48 enable bits (two per trunk assignment location) operates independently on its associated line and inhibits PCM writes into the line memory in the transmit TSI or forces idle PCM to the line in the receive TSI mode, regardless of the contents of the PCM line memory. This feature allows the microcomputer to read and write the line memory to check for faults without worrying about the data being overwritten or about extraneous data being sent

to an idle channel. Since the enable bit feature does not inhibit the line from being assigned to a trunk, it allows, with the aid of the microcomputers at both COT and RT, PCM test codes to circulate around the entire connection loop prior to customer cut through. This checks approximately 90 percent of the hardware involved in a connection and, as implemented, adds only about 15 ms to the connect delay time. This feature employs pairs of enable bits as they are read out to Register 4 in conjunction with the trunk assignment for that time slot. While in Register 4, the enable bits can cancel writing data to the RAM from Register 1 (Transmit Mode) or force Registers 2 or 3 to be loaded with idle code (Receive Mode) for the particular lines that correspond to the time slot in question.

4.1.8 Data-link register

Another auxiliary feature built into the TSI to relieve real-time constraints on the processor is an 11-bit data-link shift register (Register 7). The data link, as seen by the concentrator, consists of 11-bit packets of data, every 9 ms, in a serial format. The data-link register is loaded or unloaded in parallel by means of the microcomputer address and data ports, after which it is shifted asynchronously by the data-link clock. The microcomputer polls the data-link frame signals separately to determine when to read or write the shift register.

4.1.9 Initialization (power-up) sequence

A final feature of the TSI is a power-up and initialization sequence. By means of an external RC network, a latch internal to the TSI is set upon power up. Then after the clock starts, a 12-frame sequence must be passed through before the chip comes out of initialization. This dual method of RC timeout and clock timeout assures a robust initialization sequence that assures all memory is initialized and all trunks and busy trunks are deassigned. The initialization sequence is also very useful for manufacturing testing, as is an additional lead that allows breaking up the main-control counter sequence.

4.2 Program features

The software for the TAU MAC-8s was developed on a UNIXTM time-sharing-system (see Ref. 4). A Bell Laboratories microprocessor development tool for the MAC-8, called PLAID, was used for debugging and testing the code.³

The COT and RT programs are written in MAC-8 assembly language and are designed to fit into the available 4K bytes and 2K bytes of ROM, respectively. Assembly language was used, not only because of limited program capacity, but also because of stringent real-time constraints, which exist in part because of the decision to keep hard-

Table I—Code breakdown for the COT TAU functions

Function	Percent of Code
Call processing and message handling	38%
Self-diagnostics	37%
Data base consistency checks	(9%)
Line/trunk fault handling	(10%)
RAM, ROM, MAC-8 tests	(10%)
Initialization and alarm filtering	(8%)
Traffic and blocked calls recording	17%
Channel unit identification	5%
Other	3%

ware at a minimum where software can do the job. For this particular application, these decisions are still justified after the fact. It is also true that understandable code can be written easily in MAC-8 assembly because of its C-like syntax, thereby negating some normal aversion to assembly-level programming.³ A characteristic of the code is that it is very heavy in register instructions, as would be expected for byte and time-efficiency reasons. (The MAC-8 employs 16 general-purpose registers residing in RAM.) Careful attention was given to register usage so that data required over long segments of a routine, or even between subroutines, could remain as register variables.

4.2.1 Software statistics

Table I shows a usage breakdown of the COT TAU code. The fundamental job of call processing and message handling represents only 38 percent of the code. The importance of self-diagnostics is obvious and reflects the concern and effort that was expended in this area.

Table II shows a breakdown of the code in terms of routines, instructions, and bytes. A large percentage of the code runs in response to interrupts generated by data-link message requests. For example, all trunk assignments and deassignments initiated by the COT TAU are triggered by the need to form a new data-link message for transmittal to the RT. A total of 3894 bytes (95 percent) of the available 4096 bytes are used. Table III lists the bytes of RAM required by the COT TAU MAC-8 program. Only 286 bytes (28 percent) of the available 1024 bytes are used.

Table II—Statistics for COT TAU MAC-8 assembly-language program

	Interrupt Processing	Background Processing	Total
Routines	28	6	34
Instructions	923	418	1341
Code (bytes)	2644	1199	3843
Bytes/instruction	2.86	2.87	2.87
Percent of code	69%	31%	100%
Data table (bytes)	49	2	51

Table III—Statistics for COT TAU R/W memory use

Program variables	64 bytes
Line/trunk data base	90 bytes
Stack allowance	100 bytes
MAC-8 registers	32 bytes
Total	286 bytes

4.2.2 Program flow

The COT TAU program can be envisioned as a Main (background) routine that runs when no other processing is needed and interrupt routines that respond to real-time call-processing requests.

The Main routine performs most of the sanity and memory-checking tests. Another job is to manipulate the “TNEN” bits into masks used to force trunk assignments or no assignments in response to the types of channel units that are inserted in the SLC™-96 bank. Several alarm filters are maintained; the Main routine examines these filters and outputs the correct system and faceplate alarms. The traffic and blocked-call calculations are also performed by the Main routine. The results are passed to an interrupt routine, which times the output for displaying on the COT TAU faceplate and for outputting via the relays. The Main routine also performs some housekeeping chores.

An interrupt from the receive TSI occurs every frame (125 μ s). Because of the digitally generated, fast-busy tone feature, whereby PCM code words are read from the processor ROM to the receive TSI, this high-speed interrupt is needed. This interrupt is then counted down to time other functions for the interrupt routines (see Fig. 6). Every 2 ms the Data-Link Polling and PCM Test routine (Poll routine) is entered. This routine polls the data-link frame signals to determine when a data-link message must be transmitted or received. As a result, the Transmit and Receive Message routines are executed every 9 ms. The PCM test portion of the Poll routine refers to a function performed at the time of assigning a line to a trunk. Prior to customer cut-through, test PCM codes are circulated COT to RT to COT. The transfer of the test codes from receive TSI to transmit TSI is done by the processor by sampling at the 2 ms rate. Finally, the two-digit numeric display on the COT TAU faceplate is multiplexed at 6 ms intervals by the Display Mux routine.

4.2.3 Data-link message protocol

As mentioned previously, the concentrator data link consists of 11-bit framed messages. These 11-bit messages are grouped together in a protocol providing error protection by means of redundancy. All messages except “Idle” are 33-bit messages made up of three sequential packets of 11 bits. For digital central-office compatibility, normal

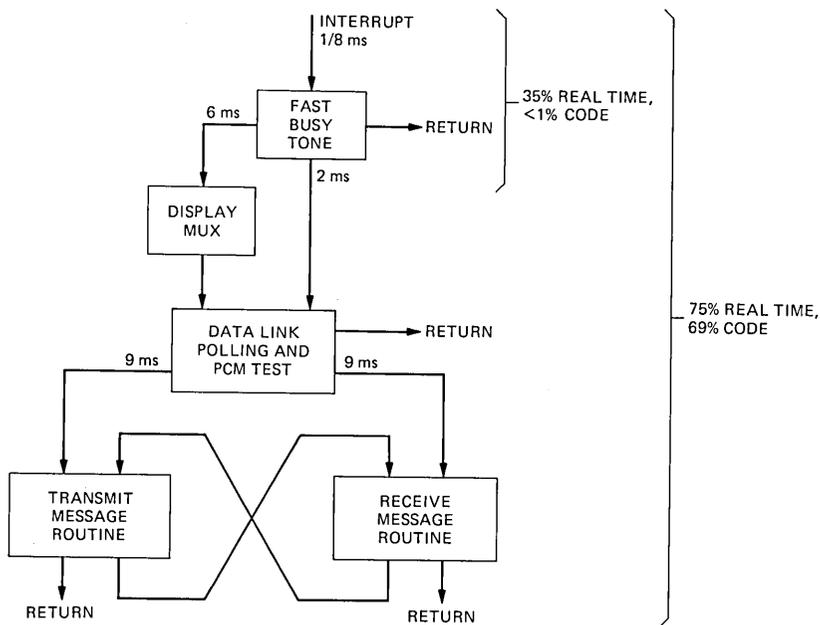


Fig. 6—COT TAU interrupt routines.

messages communicate change-in-state information. Provision is made for updating of assignments and activity as required.

The COT transmits a trunk/line assignment (deassignment) by three identical, sequential 11-bit submessages. The RT looks for a two out of three match to respond, thus providing error protection. The RT similarly sends activity as three identical submessages. Update information is needed to periodically assure that the COT and RT TAUS data bases are in agreement. Updates are sent as a header plus the message and its complement. This biases these messages toward getting through correctly, or not getting through at all, a desirable condition for update messages. The COT sends assignment updates whenever it desires and at the request of the RT by an "Assignment Update Request" message. The RT sends activity updates only at the request of the COT by means of an "Activity Update Request" message. A "Looping Test" message is a periodic message initiated by the cot to test the continuity of the data link. A "No Alarm" message is sent by the RT periodically as a fail-safe way of sending an alarm message to the COT. Care had to be taken in selecting the code words for the messages to assure that message boundaries could be determined, and also to provide error protection across message boundaries since all messages butt end-to-end.

4.2.4 Activity filtering

Activity filtering is provided in both the COT and RT, in addition to the rudimentary filtering provided by the TSI chip. Two bit-up/down counters are employed as filters using unequal attack-decay, with variable thresholds providing hysteresis. These filters provide noise immunity and delay for bridging over dial pulsing, switch-hook flashes, and silent intervals of ringing so that trunks are not deassigned and reassigned during these intervals. These filters negate the need for analog filters on the channel units to perform this function.

4.2.5 Processing trunk requests

The COT controls trunk (busy trunk) assignments and deassignments based on the COT and RT activity. As mentioned earlier, a connection test is done when a line is assigned to a trunk, whereby PCM test codes are circulated COT to RT to COT prior to cut-through of the line. If no trunks are available, the line is assigned to a busy trunk at the COT only and the line then receives the digitally generated fast-busy tone. In the case of a blocked call caused by activity from the RT, that call is transferred to a trunk when it becomes available. Thus, in this case, an RT customer would experience delayed dial tone.

4.2.6 Consistency checks

To assure that all pieces of the data base relating to line and trunk status are correctly correlated, consistency routines have necessarily been implemented. If a conflict arises (for example, a line is assigned to two trunks), corrective action is taken. Such conflicts only arise because of glitches or memory faults, but must be guarded against. In fact, detecting and reacting to "soft" and "hard" errors was one of the most challenging aspects of the software work.

Another related aspect was the requirement that the processor be able to recover from an arbitrary R/W memory state, because it is a stand-alone computer. This required careful consideration and thorough testing to determine that, for instance, the processor would not "hang" if a flag bit accidentally flipped. The ability to detect and react sanely to a genuine fault is also related to these problems. Verifying that software works correctly under the above-mentioned conditions is difficult. The TAU programs were tested by observing the reaction to random data in the processor RAM, and also by forcing bit faults by means of special hardware. These tests were excellent in pointing up several software bugs.

4.2.7 Responding to circuit faults

The TAU software is designed to be very tolerant of R/W memory faults because approximately 65 percent of the R/W memory that is

used (including TSI) is dedicated to per-line functions. Being able to isolate a fault to a single line allows the system to continue operation with a minor alarm condition. The effect is equivalent to a reduced "system crash" failure rate for the TAU plug-in. It is estimated that the equivalent TAU lifetime will be increased by approximately 40 percent because of the fault-responding software. That is, this software would often allow the unit to be replaced before it caused a system crash. This software represents only about a 10 percent overhead in code (cf. Table I).

4.2.8 Identifying channel-unit types

Another software function is processing the "TNEN" information. By associating two TNEN bits, the type of channel unit plugged into a particular physical slot can be determined. The transmit TSI picks up this information and stores it in its own memory from which the processor can obtain it. The result of the processing is essentially two masks. One mask, when combined with the A and B signaling bit activity, generates permanent activity. The other mask forces no activity. Thus a special-service unit plugged into the correct physical slot can be given permanent service or, if plugged into an illegal slot, can be denied service and the condition alarmed. The channel-unit information is also used to condition traffic calculations, since traffic (in CCS) applies only to concentrated trunks.

4.2.9 Measuring traffic and blocked calls

As mentioned previously, a function of the COT software is to calculate and store information related to traffic and blocked calls. A basic software consideration is that some of this information is held for very long periods. The faceplate traffic alarm is based on two or more blocked calls for two out of three weeks running. The traffic and blocked calls displayed on the faceplate are stored indefinitely. Thus, it was necessary to provide storage protection for these pieces of information. For simplicity, the approach used was to triple-store the data and recover them by a two out of three match. This includes not only data, but also the long-term software timers.

4.2.10 Alarm filters

The TAU has the ability to output both minor and major system alarms and to light alarm LEDs on its faceplate. To control these alarm outputs, the software maintains several alarm filters. These filters are up/down counters with a natural decay (down count) built in. To maintain an alarm condition, the appropriate filter must be incremented periodically or fail to be incremented, depending on its use. These alarm filters are maintained for various purposes. For example, one filter checks that the interrupt routines are periodically serviced.

4.2.11 Diagnostic and initialization routines

Finally, the software also includes necessary diagnostic routines. A ROM checksum and a processor maze test are performed continuously. A RAM test is also done continuously, but without slowing down call processing or missing activity. This is done by testing one byte of RAM at a time from an interrupt routine. A few bytes of RAM that are directly involved in the highest-speed interrupt are tested by an indirect method. As alluded to earlier, methods are also used to detect that the processor periodically passes through various portions of the interrupt routines. A sanity monostable is also employed and is strobed on each cycle of the main routine. Power-up initialization routines are provided based on duplicated bytes that are set in the transmit TSI after a power-up or a power supply glitch. Initialization of the data base is not done based solely on a sanity monostable timeout or processor reset since this could be caused by a glitch, which would not be a reason to take down all line/trunk connections.

V. SELECTED IMPLEMENTATION DETAILS

5.1 The time-slot interchange chip

5.1.1 Organization of the TSI RAM

The NMOS TSI chip uses a custom-designed RAM in conjunction with polycells (standard catalog gate functions) for the register and control logic. Figure 7 shows a functional block diagram of the RAM. The operation of the RAM is slaved to a Master Clock (MC) input signal, as is nearly all on-chip circuitry. Address, input data, and output data are all latched internal to the RAM. Separate data input and data output busses exist. The memory is broken into three sections of 49 bytes each. The upper two address bits are decoded to select one of the three sections. Since four possible combinations of the two address leads exist, the fourth combination, not needed for addressing a memory section, is used for addressing the data-link register (Register 7, Fig. 5), activity/TNEN control register (Register 8), and two test bits, thus making these isolated circuits appear as part of the RAM memory space.

The bottom six address leads are decoded to select one of the 49 byte locations within a memory section. To avoid unnecessary transistors and a resulting slowdown of memory operation, a full decoding of the six bits was not done. Valid addresses for the six least-significant bits are decimal 0 to 48 and 63. Addresses 0 to 47 access the 48 bytes used for PCM storage in Sections 1 and 2 of the memory and access the 48 bytes used for trunk and busy-trunk assignments in Section 3. Address 48 or 63 selects the remaining byte. This byte is an all-zero ROM byte in Sections 1 and 2 and a read/write byte in Section 3. The ROM bytes are used to send all-ones (by an inversion) on the T1 line

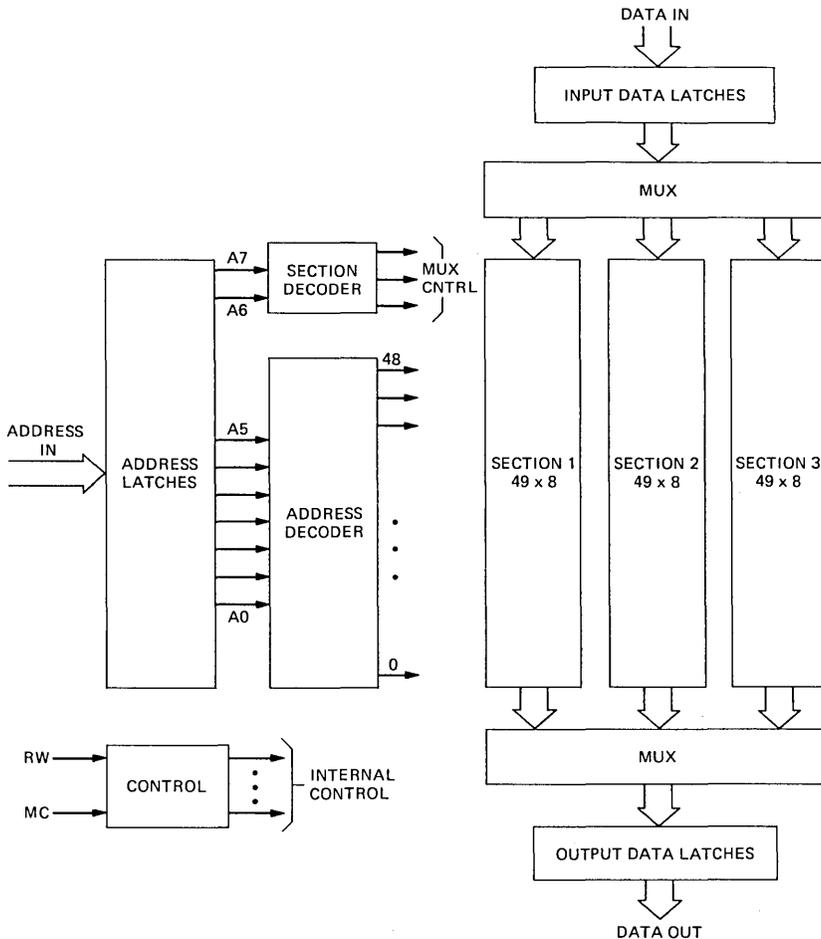


Fig. 7—Block diagram of TSI RAM.

for idle trunks. Unassigned trunks are written to “line” number 48 or 63 which, when fed back, will address one of the ROM bytes. The 49th byte in Section 3 is used as a buffer store for the fast-busy tone bytes written to the TSI by the microcomputer. Since the tone byte must be stable during an entire frame, a buffer is needed to allow the microcomputer the flexibility of writing the next tone byte any time during a frame.

The memory presents a timing interface for the remaining circuitry on the TSI chip (see Fig. 8). For a read or write operation, the address, read/write signal, and input data must be stable at the end of the first half cycle of Master Clock (MC). The memory latches the address, read/write, and input data on the rising edge of MC. For a write

operation, the write will be completed by the end of the MC cycle (Fig. 8). On a read operation, a peculiarity of the memory is that the data appears in the next MC cycle after the cycle that commanded the read. The data becomes valid some time after the start of the cycle and remains valid for the rest of the cycle. Each cycle of MC can be used to perform either a read or write. The TSI design always mandates one or the other. In memory cycles when no useful function is needed, a memory read is performed but the data is not accepted by any of the registers attached to the memory output data bus.

5.1.2 TSI memory layout

For programming purposes, the TSI can be treated simply as a block of memory. Some memory bytes are used for different purposes depending on whether the TSI is functioning as a transmitter or receiver. The 24-trunk assignments are stored in the bottom six bits of the even-addressed bytes in Section 3 of the RAM. The six bits are set to a number between 0 and 47 to indicate which of the 48 lines is assigned to that trunk. If no line is assigned, either 48 or 63 can be written into the six bits. The upper two bits in each of the trunk-assignment locations are called enable bits and operate independent of the trunk assignment. Taken together, there are 48 enable bits (two in each trunk location). Each bit corresponds to a particular line and, if cleared, inhibits PCM read-in to the PCM data locations if the TSI is a transmitter, or forces idle PCM data out of the TSI, regardless of the PCM data memory contents, if the TSI is a receiver.

For receive TSI operation at the COT, the odd-address locations in Section 3 contain the busy trunk assignments. Again, as with the trunks, each location is written to a number between 0 and 47 to assign a line to a busy trunk, or written to 48 or 63 if the busy trunk is not assigned. The upper two bits of the busy trunks have no special use. If

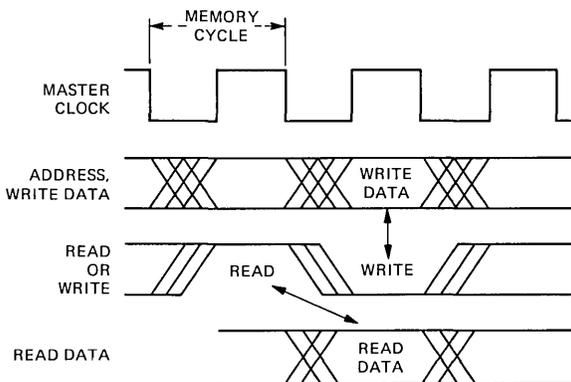


Fig. 8—TSI memory operation.

the TSI is operating as a transmitter, then there are no busy trunks, but some of the same locations are used for holding activity and TNEN bits. Each of the activity or TNEN bytes holds eight data bits, one for each of the eight lines. Thus there are six bytes of activity and six bytes of TNEN to hold data for the 48 lines.

5.1.3 Internal control

The control logic on the TSI chip regulates the flow of data between registers and memory within the TSI chip, and the selection of times appropriate for microcomputer accesses. The main element of this control is a 12-frame counter. By assigning states of this counter to various functions, control of the chip is achieved.

5.1.3.1 Trunk assignments. To understand the assignment of states to specific functions, first consider the steps that must be performed to accomplish the basic time-slot interchange function for a TSI operating in the transmit mode. Serial unconcentrated data is coming in from master and slave TRUS and is being converted to eight bit parallel data in Registers 1 and 2 (see Fig. 5). Every eight cycles of MC (1.544 MHz Master Clock), two PCM words must be written to the memory, one from Register 1 and one from Register 2. Similarly, every eight cycles, one PCM word must be read from the memory and written to Register 3 for transmittal on a T1 trunk. (As noted earlier, the trunk data is read from one section of the PCM memory while the unconcentrated data is being written into the other section.) To read out the trunk data, however, the line assignment for that trunk must first be read out to Register 4, to provide the address for the line data to be read out. Thus, these basic operations consume four out of every eight MC cycles. For the receive TSI operation the trunk data comes into Register 1 and exists unconcentrated through Registers 2 and 3, similarly requiring four cycles.

5.1.3.2 Busy trunk assignments. In the receive TSI operation at the COT where busy trunks are needed, two additional cycles out of every eight are used, one for reading the busy trunk assignment to Register 4 and one for writing the current fast-busy tone byte from holding Register 5 to the PCM line memory of the addressed line. The fast-busy byte is then read out to Register 2 or 3 via the normal sequential readout. The current fast-busy tone byte is read from the fast-busy buffer location in memory to holding Register 5 during the frame bit state (193rd count of the control counter). Thus the receive TSI at the COT uses a total of six out of every eight cycles plus the one extra state occurring each frame.

5.1.3.3 Activity/TNEN. In the transmit TSI mode, states other than the four basic states are used for transferring activity and TNEN information from/to Register 5. These states occur once every 32

states rather than once every 8 states, because it takes that long to collect a byte of activity or T_{NEN} data.

5.1.3.4 Microcomputer access. In all modes of operation, no more than six out of the eight states are used for internal operations; the other two states are reserved for potential use by the microcomputer to perform a read or write of the memory. Waiting for an available state gives rise to the wait states encountered by the microcomputer on a TSI access. The two internally unused states out of every eight are evenly spaced to minimize the delay seen by the microcomputer.

5.1.4 System synchronization

5.1.4.1 Transmit TSI mode. In the transmit mode the TSI chip establishes its own reference based on its internal counter state. It then forms a super-frame system synchronization signal for sending to the TRUS (a unique pattern over 12 frames). The signal consists of two pulses, one in frame 2 and one in frame 12. Each pulse must be only 162 ns wide to synchronize the 6.176 MHz count-down circuits in the TRUS (6.176 MHz is the fundamental crystal-controlled clock in the SLC-96 system from which the T₁ rate is derived). The TRUS and the TAU receive the same clock and then the TAU must carefully time the pulse with respect to the clock edges. Because of the short delays involved here, the pulse had to be timed with TTL external to the custom NMOS chip. The TSI chip puts out a pulse that spans two 6.176 MHz periods which is then gated externally to produce the desired pulse. The effect of the pulses that are sent to the TRUS is to synchronize the counters in the TRUS.

As the slave digroup PCM enters the TSI, it is delayed by one clock period by means of an FF (see Fig. 5) so that it arrives one MC cycle later than the master PCM. The outgoing trunk PCM is also delayed one cycle by means of an FF, which exists for matching outgoing PCM streams in the receive mode. The frame bit from the master digroup is picked off and saved in the FR FF for re-insertion into the outgoing trunk PCM bit stream (which is why this stream is called the master).

5.1.4.2 Receive TSI mode. Synchronization in the receive mode is very different from synchronization in the transmit mode, because the timing is controlled by the reframing circuit contained in the TRU. This circuit locks on the frame bit in the received T₁ bit stream and thus can provide the timing reference and reframe signals needed by the receive TSI in the TAU.

The timing reference is provided in the form of a 4 KHz (two frame) clock signal from the master digroup TRU. The edge of this clock waveform is used to trigger the TSI control counter to a predetermined state chosen to synchronize the TSI to the incoming bit stream from the T₁ line interface unit. When calculating the position of this bit

stream relative to the synchronization signal provided by the TRU, the delay in passing through the TSI must be taken into account.

Reframing is a very important consideration for receive TSI operation. Since the reframing process occurs downstream from the TSI, the TSI must assume a special mode of operation during the reframe process. This mode must pass the concentrated trunk bit stream directly through the TSI unchanged to assure that the frame bit is received by the TRUS (the timeslot interchange function can lose the frame bit if the TSI chip is out of sync). Once the TRUS reframes, the receive TSI can resume normal operation. To assure that the TRU is not thrown out-of-frame when the receive TSI returns to its normal sequencing after a reframe, it is necessary that the relative position of the frame bit not change. That is, during the time that the TSI is passing the received concentrated bit stream directly through, it must insert the same number of cycles of delay as it will when it is operating normally. Also note that the scrambled time slots sent to the channel units during reframe are of no consequence because the channel units are not enabled to receive the information.

A further consideration during the reframe process is that the 4 KHz synchronization signal slips as the TRU searches for the frame bit. Thus this signal cannot be allowed to preset the TSI control counter during this time or confusion can result. The TRU provides an out-of-frame signal that can be used to gate the synchronization clock and also force the TSI chip into its special reframe mode of operation. Once the out-of-frame signal indicates the TRU is reframed, the synchronization clock is allowed to once again force the counter state in the receive TSI.

During the reframe mode the TSI performs as follows. All incoming trunk data, instead of being selectively written to the assigned line memory locations, is sequentially written to all the slave PCM memory locations in the same section of PCM memory as the selective writes would have been done. The sequential readout of each line's data in the alternate section of PCM memory continues as normal, except that the slave data is also forced out to the master digroup. In this way the concentrated bit stream is passed directly through the TSI with the same delay experienced by the bit streams in the normal mode of operation. Note that if the receive TSI is out of synchronization (as it is assumed to be if a reframe is needed), then the frame bit is not being put into the frame flip-flop (FR FF). Instead, an arbitrary bit out of the frame sequence is being inserted in the FR FF, depending on the relative state of the TSI control counter. The actual frame bit, then, is being stored somewhere in the PCM memory. The bit stream passes unchanged through the TSI, however, with the bit stream being put back together correctly as it exists from the TSI.

After the reframe process terminates, the control counter is resynchronized as mentioned above. Since the frame bit was arbitrarily located in memory, the resynchronization will, in general, cause the loss of one frame bit. The lost frame bit will be in error with 50 percent probability assuming a random bit is inserted in its place. By selecting the rising edge of the synchronization clock, the frame bit that is lost is a signaling frame bit. Thus there is no concern, during the resynchronization, about approaching the master frame bit error threshold that would throw the TRU back out-of-frame. A signaling frame bit error will, at most, delay the reframing of the signaling bit extraction circuit in the TRU.

One other step must be taken in the TSI to complete the reframing process. Since the slave PCM memory locations were used to hold data while the TSI was passing the concentrated bit stream directly through, it is desirable to clean out this memory, that is, write these memory locations to idle PCM code. This is done automatically by the hardware during the two frames that follow the control counter resynchronization (two frames are needed to initialize both sections of PCM memory).

The TSI adds approximately four frames of time to the reframe process for PCM data. This is an increase of 0.5 ms to an average reframe time of 25 ms. The TSI may also add up to one super frame of delay to the receipt of A and B signaling bits by the channel units because of the erroneous signaling frame bit. This is equivalent, worst case, to the loss of one additional A and B signaling bit.

5.1.5 Chip development

The development of the Time-Slot Interchange (TSI) chip was justified on the basis of cost, power, and space. Further the chip is a universal design that can function as a transmitter or receiver, at COT or RT, so that only a single custom design is required. The breadboard for the TSI was built using 96 off-the-shelf integrated circuits. The chip is realized with NMOS polycells and a custom NMOS RAM using 5 micron rules on a chip 258×367 mils, packaged in a 40-pin DIP. Typical power dissipation is 750 mW. The design cycle from the first paper design to first chips took about two years. The chip contains 432 polycells, 145 bytes of static RAM, and two bytes of ROM. There are about 10,000 transistors, 70 percent of which are used in the RAM. Extensive logic and timing simulations were required to verify the design.

5.1.6 Internal timing considerations

The TSI clock frequency of 1.544 MHz mandated careful consideration of timing delays in the design of subcircuits for the TSI. "Regular power" NMOS polycell gates can give delays of approximately 50 ns. "High-power" and "super-power" gates must be used for shorter delay

times with corresponding increases in power dissipation. Certain critical paths in the TSI did require the use of high-power and super-power gates.

This method of meeting timing constraints, however, was only supplementary to the decision to build a highly synchronous design. By clocking nearly all subcircuits directly with MC (Master Clock), a very clean timing plan was developed. That is, all propagation delays, setup times, and hold times could easily be calculated with respect to an MC edge. Memory operations were also specified with respect to MC (previously discussed). Such a design avoids accumulating long strings of delays that can cause problems at the expense of, in some cases, additional gates. For example, all the registers have their FFs directly clocked by MC. Gating around and between FFs determine the function the register performs such as hold, shift, or load. With this type of design, where all action takes place on a clock edge, the remaining portion of a cycle can be used for propagation delays of the signals that determine what function will be performed when the next clock edge arrives. This design technique is very valuable for high-speed designs (high speed relative to the technology limitations).

Another timing consideration that should be mentioned is the expected clock duty cycle variations in the receive TSI mode. MC in the receive mode is the recovered T1 line clock. As such, duty cycle variations can be expected. Final assumptions were for a 60/40 duty cycle worst case in either direction. This implies a need for the TSI chip to work, equivalently, at a higher frequency (viz. 40/40). Timing must apply for worst-case device, power supply, and temperature variations, as well, leading to the test-clock rate at room temperature of 2.5 MHz.

5.2 COT software details

5.2.1 Real-time constraints

A basic problem in structuring the software was to determine a method for handling the real-time constraints. All the interrupt routines run indirectly off the fast-busy tone interrupt which occurs once every eighth of a millisecond (recall Fig. 6). The message routines can take milliseconds to execute, however. Thus it was necessary to enable the interrupt upon leaving the Fast-Busy Tone routine to enter the other interrupt routines. The result is that several levels of interrupt can exist on the stack. The MAC-8 microprocessor automatically saves the condition register and return address on the stack when an interrupt occurs, and thus conveniently allows nesting of interrupts. Since the Main routine may have been nested in subroutines when it was interrupted and the interrupt routines, themselves, may have been nested in subroutines at the time of another interrupt, the stack must be large enough to hold all the return addresses. A few data values are

also sometimes stored on the stack. Fifty bytes was determined to be adequate by considering worst-case nesting levels and including margin for unexpected levels caused by an interrupt glitch or an accidental bit flip (a subroutine call stores two bytes on the stack, an interrupt three bytes).

5.2.2 Time-share processing

To get a better feeling for how the processing gets time-shared, refer to Fig. 9. This time diagram is an example of the routines that might be processed in response to a sequence of interrupts. The Fast-Busy routine is executed after every interrupt, followed by a return to the routine which had been executing. This time line starts out by assuming the processor is in the Main routine when an interrupt occurs. The interrupt causes the Fast-Busy routine to be executed. After 16 executions of the Fast-Busy routine, a 2 ms timeout occurs causing the Poll routine to be entered (recall Fig. 6). The Poll routine can be interrupted several times, before it finishes, by the Fast-Busy routine. After the Poll routine finishes, it is assumed that the Transmit Message routine needs to be executed. This routine runs for milliseconds and so will be interrupted not only by the Fast-Busy routine but also by the Poll routine. After the Transmit Message routine terminates, it is assumed necessary to process the Receive Message routine and so control switches to it. Finally, the Receive Message routine finishes and control returns to the Main routine.

The Fast-Busy, Poll, and Display Mux routines are fast enough to finish before they would be called upon to re-execute. However, they would not cause serious problems even if they were re-entered because of some glitch. The message routines are much more complex, however, and can cause some serious consequences (confusing trunk assignments, for example) if they are re-entered. Normally these routines would be finished before being called upon again (as necessary to meet the demands of the data links for messages every 9 ms). However, to provide a degree of protection from the havoc that could occur if they were re-entered, a flag is maintained that indicates the message routine needs processing or is in the process of being executed. This flag is set by the Poll routine when it determines that the message routine needs processing, and is cleared at the end of the message routine after all processing is completed. These flags also provide the means by which control can be transferred directly from the Transmit Message routine to the Receive Message routine or vice versa (recall Fig. 6).

5.2.3 Data-link polling

The data link provided the τ AU by the SLC-96 Data Link Unit (DLU) consists of 11 bit packets of serial data every 9 ms. The τ AU is provided with a frame signal that remains high for 2.75 ms and low for

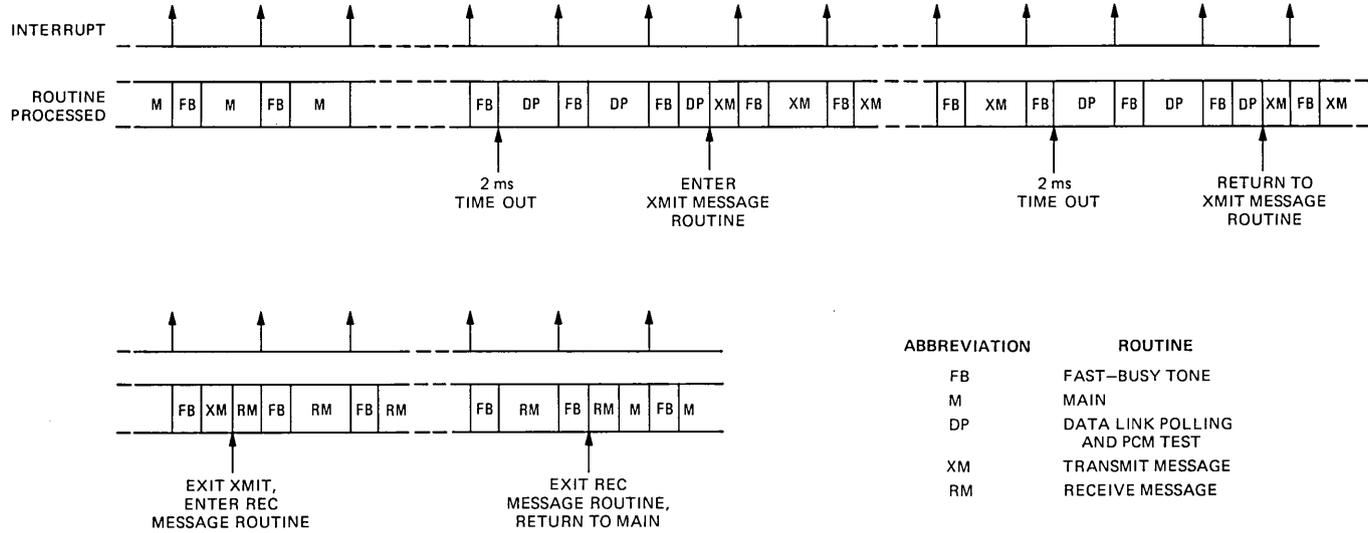


Fig. 9—An example of a routine processing sequence.

6.25 ms. During the high portion of the frame, serial data is transmitted or received at a 4 KHz rate (1 bit every 0.25 millisecond), by combining a 4 KHz clock provided by the DLU with the frame signal. As Fig. 10 indicates, the time while the frame signal is low is available for loading or unloading, in parallel, the shift register provided in the TSI chip. Figure 10 also indicates that the frame signal is polled every 2 ms. The polling is asynchronous and so only one possible phasing of the polling with respect to the frame signal is shown. The polling scheme works by noting 1 to 0 transitions of the frame signal and using this event as a trigger for processing the message routine and loading (unloading) the shift register. Both the Transmit Message routine and the Receive Message routine write or read, respectively, the shift register immediately upon entry; thus, if the routine is delayed or takes a while to process, it is assured that the data-link shift register is loaded or unloaded prior to the next rise in the frame signal.

The actual data-link polling is complicated by the fact that there are two data-link frame signals, one for transmit and one for receive. The data links are asynchronous with respect to each other (or at least out of phase with any random phase), and so all relative phasings must be considered. The message routines are designed to run to completion before transferring control to the other message routine, if needed. Therefore, a message routine can be delayed and real-time constraints must be considered. Since a request to process both message routines may occur at the same time, a priority of processing had to be established.

Figure 11 shows two possible phasings of the data-link frame signals and the processing sequences that could correspondingly result (the Receive Message routine is given priority). Note that message routine

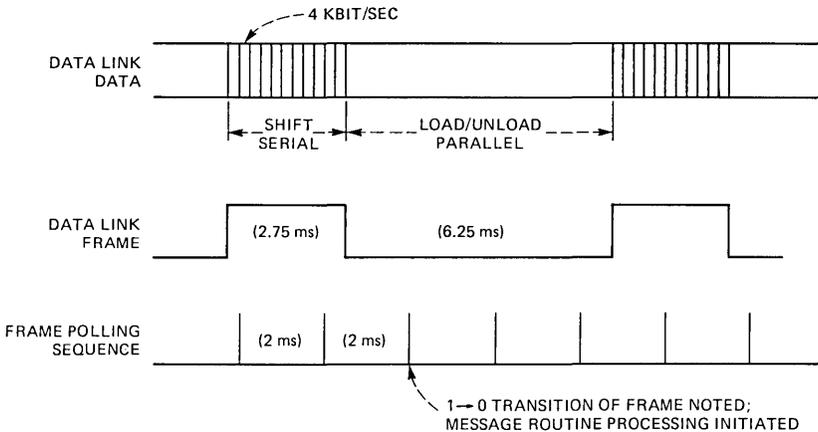
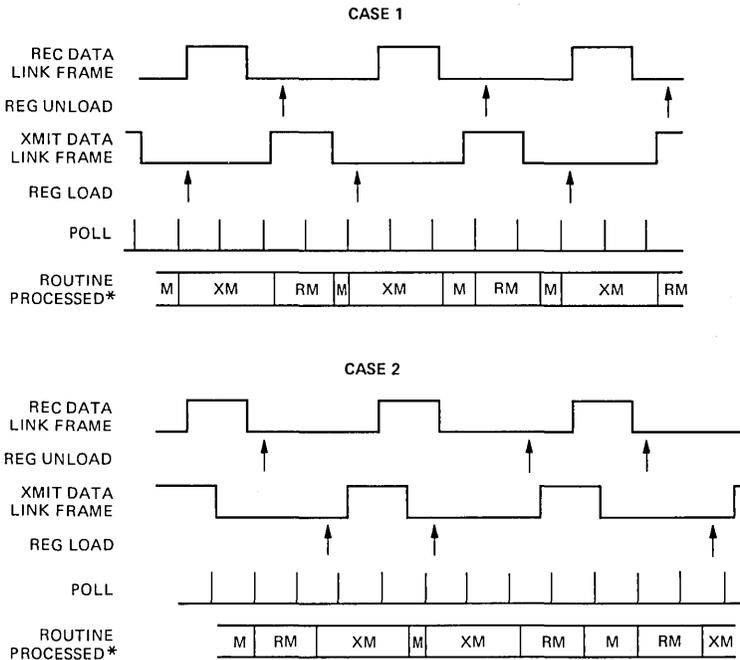


Fig. 10—Data-link message processing.



***NOTES:**

1. M = MAIN ROUTINE
 RM = RECEIVE MESSAGE ROUTINE (3 ms RUN TIME)
 XM = TRANSMIT MESSAGE ROUTINE (4.5 ms RUN TIME)
2. RM HAS PRIORITY OVER XM
3. LOW LEVEL INTERRUPTS IGNORED IN REPRESENTATION

Fig. 11—Data-link message sequencing.

processing need not alternate; instead, the Transmit Message routine may be processed twice in a row followed by the Receive Message routine being processed twice in a row as the result of 2-ms polling of the 9-ms frame interval.

If both frame signals are aligned and the polling was such that the transition from 1 to 0 is not noted until almost 2 ms after it occurs, then the Receive Message routine is constrained to be less than 4.25 ms. Enough time must be allowed to enter the Transmit Message routine and put out the new submessage, prior to the rise of the transmit data-link frame signal.

Since only a single flag is used with each message routine to indicate that the routine needs processing or is in processing, and since requests for processing can occur 8 ms apart, the sum of the execution time of the Transmit and Receive Message routines is constrained to less than 8 ms. If, under some unusual circumstance, the processing takes more

than 8 ms, the result is that the next occurrence of the Transmit or Receive Message routine is not processed and a data-link submessage is therefore lost. But because of the redundancy and error protection built into the data-link messages, a message will probably not be lost.

5.2.4 Time-critical routines

The most time-critical routine is the Fast-Busy routine. Great care was taken to make this routine as fast as possible; two MAC-8 registers are used, one for an auto-increment pointer to the fast-busy word data table (register b4) and one as a temporary holding register for the tone byte (register a5), since a memory-to-memory transfer instruction does not exist. (In MAC-8 assembly code, “b” registers are 16-bit general-purpose registers and “a” registers are 8-bit general-purpose registers.) For the purpose of saving bytes and improving speed throughout the entire TAU program, two eight-bit registers (a3 and a13) are used for flags. One of these flags is used for the silent interval of the fast-busy tone (the tone is pulsed on for a 0.25 s, off for a 0.25 s).

Table IV gives the cycle count on an immediate return path of the Fast-Busy routine. The variable number of cycles for the write instruction is due to the variable number of wait states for a TSI access. Assuming that 67 cycles of each frame (out of the available 193) are used for the Fast-Busy routine implies that 35 percent of the real time is used in this routine.

A 2-ms software timer is derived from the data table pointer (register b4) in the fastest possible way by a test on bit 4. For this scheme to work, the Poll routine, when it is entered, immediately increments the pointer by 16 so that only a single 2-ms timeout is indicated by the bit test on the pointer. Thus the fast-busy code table of 48 bytes is actually stored in ROM as three tables of 16 bytes, separated in address space by 16 locations. Worst case, the data table pointer must be set to its new value in less than one frame of time from the interrupt, to assure it has the correct value when the next interrupt occurs. The values for the fast-busy code words were derived by calculating linear samples of the tone and then using a translation table to companded PCM codes. The least significant bit was forced to always be one as needed for A

Table IV—Listing 1: Fast-Busy routine code fragment

Cycles	Instruction	Comment
10	(interrupt preamble)	
7	if (bit(7, a3)) goto ia1;	/* jump if silent interval */
15	a5 = *b4++;	/* fetch busy word */
16-22	BSYWD = a5;	/* write word to rec TSI */
7	if (bit(4, a4)) goto ia2;	/* check polling timer */
9	ireturn();	
64-70	(Total cycles)	

and B bit signaling to the channel units. The code inverse actually gets sent to the channel units because of a data inversion in passing through the TSI.

The Poll routine is less critical than the Fast-Busy routine by a factor of 16, but it is still important to consider worst-case paths. The routine is held to a minimum, again, by careful use of register instructions and arranging the conditional branches for minimum worst-case path. The execution time of the Transmit and Receive Message routines must also be considered, because they need to be fast enough to service the data-link requests upon demand, to avoid further complication of the algorithms. An example of what can be done to minimize worst-case timing paths is the execution of the consistency checks and updates in the Transmit Message routine on the passes through this routine that do not require the calculation of a new transmit data-link message. That is, since a message is composed of three 11-bit submessages which are calculated and saved until needed, no new message calculation is required while the first two of these submessages are being sent, and so there is real time available for the consistency and update routines to use.

The Main routine has no significant real-time constraints; however, it is still necessary to know the approximate cycle time of the routine to allow an adequate timeout for the sanity monostable. The cycle time of the Main routine is, of course, strongly affected by the running time of the interrupt routines. (It is necessary to strobe the sanity monostable from the Main routine, rather than an interrupt routine, since the latter might continue to strobe the monostable while allowing a return to an unknown loop at an arbitrary location, rather than to the Main routine.)

5.2.5 Line/trunk data base

The processor RAM is used for several purposes as was given in Table III. The line/trunk data base stores copies of the 24 trunk and 24 busy-trunk assignments. These copies are used during assignment searches and acted upon by the consistency routines. The transmit and receive TSI assignments are updated from this copy. These copies are maintained because they can be accessed more quickly than the information in the TSIs.

The line/trunk data base also holds information concerning the 48 lines the concentrator serves. This information is split into six data groups, each group holding information for eight lines. The information is stored in bitwise correlation with the eight bits of line activity as collected in a single byte by the transmit TSI. For each group, seven bytes are stored. Bytes 0 and 1 are the T_{NEN} masks for forcing activity or no activity. Bytes 2 and 3 are the activity filters, stored as least-

significant bits and most-significant bits of two bit up/down counters. Byte 4 stores the status of activity from the RT as received over the data link. The trunk and busy-trunk status in bytes 5 and 6 indicate whether a line is assigned to a trunk or busy trunk, respectively. These bits are maintained to provide a fast method of determining whether an assignment (deassignment) is needed. Without them, one would have to scan all the trunk assignments for every line that had activity to determine which line needed assigning. This could not be done within the time constraints of generating a data-link message "upon demand." All the bytes of data in each group are arranged so that all eight lines can be processed simultaneously by byte operations, thus performing activity filtering, TNEN masking, and service request determination very quickly.

5.2.6 Handling faulty lines and trunks

A significant feature of the software is the ability to allow concentrator operation in the presence of partial faults. Before assignment to either a trunk or busy trunk, extensive local memory checks are made on the trunk (busy trunk) assignment, enable bit, PCM data, and status bit memory locations. A failure causes an alarm to be raised and possibly a line/trunk fault to be stored. If local tests pass, then the PCM looping test is set up for a trunk assignment. A failure of this test can also cause a fault to be stored.

The routines that are provided allow for the detection of memory or connect (disconnect) failures at the time of line/trunk assignment (deassignment). If a particular line/trunk combination fails, that combination is put in a fault store and periodically retried. One such fault gives rise to a minor alarm, two such faults shut down the system. The question is always raised, why not try to assign the line to a different trunk if the first one fails? This is not as easy as it sounds. The basic problem is one of fault isolation. The line/trunk *combination* is fundamental in finding the fault; splitting them up could easily cause a loss in the ability to refind and retest the fault, to maintain an alarm. The result could easily be intermittent alarms or a faulty trunk that wanders from line to line, possibly causing random customer complaints. In any case, it would require a lot more software with diminishing returns. It is also true that, for a given line/trunk combination fault, hardware considerations give a higher probability to the line being at fault than the trunk. In short, it seems very acceptable, and is fairly straightforward in software, to keep the line/trunk combination as a means for maintaining an alarm while allowing all other customers normal service.

A fault consists of a trunk and line pair. The stored trunk number is the lower byte of the address to the trunk or busy-trunk assignments

in the processor RAM. A line number is stored as 0 to 47. If the fault was trapped during a deassignment, the "line" number may be 48 or 63. Both 48 and 63 are used because TSI will allow either one as a assignment number. If memory fails and will not allow deassignment to 63, then 48 is tried. An empty fault store is designated by all ones for the line number and all zeroes for the trunk number.

If a fault is to be retried, the trunk and line numbers are checked for validity and then a jump to the appropriate trunk (busy trunk) assignment (deassignment) algorithm is made. Also, to assure that the fault is not lost because of changes in the faulted line's activity, permanent activity is maintained by setting the activity filter for that line to a count of 3. A fault will be retried approximately every 1.8 s and thereby continue to increment an alarm filter if the fault persists.

5.2.7 PCM looping test

For a trunk assignment, the PCM looping test is always executed. First, the test is initiated in the Transmit Message routine. This process consists of clearing the enable bits in transmit and receive TSI for the line under test, so that the test codes will not be overwritten or sent to the channel units. Also, the first PCM test code (alternating 1's and 0's) is written into the transmit TSI PCM memory locations and a PCM test timer is set for an 80-ms timeout. Since the line is assigned to a trunk at the COT, the test codes are received at the RT as soon as the RT receives the trunk assignment message over the data link. The enable bits at the RT are automatically cleared upon receipt of the assignment message.

The operation of the test at the COT is then picked up by the PCM test portion of the Poll routine. This routine will sample the receive TSI every 2 ms looking for the PCM test code that should be returned by the RT. When the code is received, it is complimented and sent to the RT. When the complimented code is received at the COT, the COT sends a test termination code (all 1's) to the RT for 10 ms and then sets the COT TSI enable bits, thereby cutting through the customer at the COT. The RT will correspondingly set the enable bits at the RT when it receives the test termination code.

If a timeout occurs, the COT will deassign the trunk and store the line/trunk combination for later retry. If the RT fails to see the test termination code after a timeout from receipt of the assignment message, it will simply deassign the trunk (only the COT records the trunk and line that gave rise to a fault). If a trunk (busy trunk) deassignment is requested, the disconnect is performed in a straightforward manner. Some memory tests are performed and may result in the storing of a fault. This fault may involve only a trunk if the problem is in writing the unassigned "line number."

VI. PERFORMANCE TESTING

Performance testing of the TAU was a very big part of the project. When all the hardware and software effort expended in testing is considered, the testing job was nearly as big as the basic job of designing the TAU circuits, custom chip, and software. Because of the complexity and compactness of the design, thorough and sophisticated testing was essential.

6.1 *Early-project testing*

The TAU project began with the design of the TSI. TSI breadboards were built and debugged using a scope and logic analyzer. Because the TSIs use a handshake arrangement in talking to the microcomputer bus, manual switches can be used to simulate the bus signals, and thus a working microcomputer was not needed at this stage. Later, after the TSI design was well on its way to being realized as a chip, the microcomputer designs were finalized and built. The microcomputers were debugged using simple programs and a logic analyzer.

6.2 *Mid-project testing*

Once serious programming began, a versatile test set was needed in addition to the MAC-8 development system (PLAID). A MAC-8-based test set was designed that allows the display of all 24 trunk assignments on numeric displays by examining the data link messages that flow from COT TAU to RT TAU. It also employs A and B bit and T_{NEN} bit generator cards that store information for all 48 lines that the concentrator is working with and allows the operator to manually set these bits to simulate channel units. It also displays the received A and B bits for the selected channels.

This test set was invaluable for tracing bugs in the TAU software as it developed. Some of the intermittent and transitory phenomena was especially visible on the displays. It was also very nice for observing the results of simulated memory faults, which was done with another piece of test hardware, consisting mainly of EPROMs with the selected bits to be faulted marked in the EPROM. Most of the software for the TAU was written and debugged using this MAC-8-test set in combination with the PLAID.

Another capability that was developed later and was very useful in debugging the code was a data-link monitor. This allows the concentrator data-link messages in both directions to be displayed on a CRT in a correlated fashion. The program allows the storage of messages occurring over approximately 2.3 s and allows triggering on a particular message pattern with "don't care" conditions. The stored messages can then be viewed by scrolling forward and backward. The data-link monitor was especially useful for checking out the software that

determines message priority. It was also useful in noting response messages, such as a trunk assignment message leaving the COT TAU in response to an activity message received from the RT TAU. Another use was in viewing widely-spaced periodic messages such as the data-link "Looping Message."

6.3 Late-project testing

A development that paralleled the TAU and which proved very useful to us for system tests and final TAU software tests was the "Traffic Generating System" (TGS). This development was initiated to simulate realistic traffic on digital lines for the testing of a digital switch. It was decided to develop the hardware and software necessary to use the SLC-96 carrier system for simulating traffic on a T1 line. This hardware and software also serves as a debugging tool for the TAU, while the TAU also provided a shakedown test for the traffic-generating system.

"Signaling interface boards" that perform similar to the original MAC-8 test set are used to simulate channel activity in TGS, that is, A, B, and TNEN bits for all channels are stored in a RAM and read out in the proper sequence. Similarly, received A and B bits are stored. One signaling interface board serves 24 channels. The RAM is writable and readable through an I/O port to a DEC-LSI-11. This hardware is flexible, portable, and used simply by plugging the simulator card into a channel unit position in the SLC-96 system bank. The LSI-11 is tied through a satellite processor link to a host UNIX system (see Fig. 12).

Software for TGS allows writing "scripts" in C programming language that describe what a channel does and when it does it. Library routines are available that simulate, for example, dial pulsing and ringing. Many scripts can be run simultaneously using shared library routine code. The system thus allows simulating realistic traffic conditions.

By using TGS and the data-link monitor, the final TAU boards were exercised very thoroughly. One program that was written measured the connect delay time from an A or B signaling bit change. This program was expanded to make thousands of random calls, measure the delays, and store the results in a UNIX file for later graphing of delay distributions. Most of the delay and distribution of delay is attributable to the delay and asynchronism of the data link. Originating calls are delayed more than terminating calls because of the delay in transmitting RT activity to the COT. Average connect delay from the COT is 70 ms, with 100 ms from the RT.

Delay measurements were repeated with a random-error generator used to insert errors on the T1 lines. At an error rate of 2×10^{-4} , which is worse than a functioning SLC-96 system will see, the only noticeable change in the connect delays was an increase of a few milliseconds in the delay. This checked out the error protection built into the data-link messages.

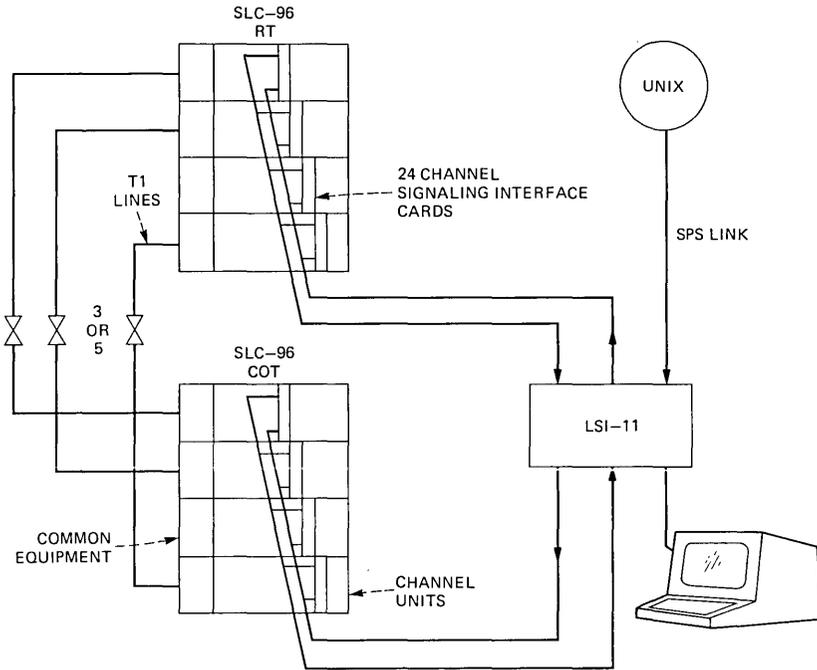


Fig. 12—Traffic generating system for TAU testing.

Another program that was written simulated two simultaneous calls and measured the connect delay of the delayed call. These measurements showed an increase in the connect delay of approximately 50 ms for the second call. Programs were also written to manipulate the T_{NEN} bits, with and without A and B activity, and thereby simulated each type of channel unit in each physical position. Other programs checked blocked-call functioning, the generation of fast-busy tone, and normal ringing and dialing. The traffic generating system was undoubtedly very important in establishing confidence in the final hardware/software design.

VII. CONCLUSION AND ACKNOWLEDGMENTS

The SLC-96 carrier system TAU demonstrates that modern electronics economically provide improved features in the loop plant. The ability to integrate the Time-Slot Interchange function onto a single chip made this development possible. By digital concentration, the TAU reduced the number of T1 lines needed by the SLC-96 system from five to three. Because provision is made for special-service circuits to be given unconcentrated trunks, a separate system is not needed to provide a few special interfaces. Traffic measurement and extensive

maintenance features were also successfully integrated into the firm-ware control of the TAU.

The author wishes to especially acknowledge Lary Range for performing extensive simulations on the TSI chip and John Beck for his software-design assistance. The layout and design of the TSI chip were directed by Gil Mowery. Sam Arnold, Brian Redman, and Doug Corey were responsible for the software design of TGS. Credit is also due to many other individuals who were associated with all phases of this project.

REFERENCES

1. "The Loop Plant," B.S.T.J., 57, No. 4 (April 1978).
2. S. Brolin et al., "Inside the New Digital Subscriber Loop System," Bell Laboratories Record, 58, No. 4 (April 1980).
3. H. D. Rovegno, "A Support Environment for MAC-8 Systems," B.S.T.J., 57, No. 6, Part 2 (July-August 1978), pp. 2251-63.
4. "UNIX Time-Sharing System," B.S.T.J., 57, No. 6, Part 2 (July-August 1978).

Compiling Three-Address Code for C Programs

By J. F. REISER

(Manuscript received February 27, 1980)

This paper describes a post processor that improves the assembly-language code generated by the portable C compiler. The novel ability to change a sequence of two-address instructions into an equivalent three-address instruction distinguishes this particular code improver from other "peephole" improvers. The combined compiler-improver generates good three-address code for the Digital Equipment Corporation VAX-11[®] computer without requiring extensive changes in the compiler itself, which was designed to accommodate machine architectures with at most two addresses per instruction. For typical programs the improver reduces the number of bytes in the instruction stream by 10 to 23 percent. This paper emphasizes the technique used to transform two-address code to three-address code.

I. INTRODUCTION

The portable C compiler¹ is an effective tool for quickly constructing a C compiler² for a general purpose digital computer. With reasonable effort the resulting compiler generates correct code, and the quality of the translation into assembly language is acceptable. However, users frequently demand better code if they anticipate prolonged or extensive use of programs written for a particular application. A post processor that reads the assembly language generated by the compiler and writes better assembly language having the equivalent effect can satisfy much of the demand. (Here "better" code requires fewer bytes for instructions or less time to execute, or both.) This paper describes a program that improves code generated for the Digital Equipment Corporation VAX-11[®] computer, paying particular attention to the technique used to transform two-address codes into three-address codes.

One reason why a code improver can be effective is that the portable C compiler often generates code in the easiest possible correct manner, even if such a code is suboptimal over a wide range of machines. The

compiler expects that a post processor will clean up after it. For example, the compiler translates the C program fragment

```
while (...) {  
    ...  
    if (b > 0) break;  
    ...  
}
```

as if it were written

```
while (...) {  
    ...  
    if (b <= 0) goto L100;  
    goto L101;  
L100:  
    ...  
}  
L101:
```

which contains a conditional jump around an unconditional jump. It should not be difficult to compile the original fragment as if it were

```
while (...) {  
    ...  
    if (b > 0) goto L101;  
    ...  
}  
L101:
```

but the compiler does not do this, so one of the standard tasks for a code improver is to replace “skips over jumps” with jumps on the negated conditions.

Another reason that a code improver can produce better code is that the compiler’s model of code generation may ignore or not take full advantage of architectural features found on a specific machine. The portable C compiler understands one-address instructions and two-address instructions, but does not understand three-address instructions or instructions which use an address as an immediate operand. Similarly, the compiler thrives on certain addressing modes (register, pointer, displacement from a register base) and has difficulty fully exploiting others (auto increment, double indexing).

A code improver can also be effective because C-language statements or compilation on a statement-by-statement basis may be too low level. The concept “turn off bit 15” may have a direct hardware implementation, but must be expressed in C language as a Boolean AND operation. The portable C compiler attempts no analysis of interstatement information flow, nor does it always take advantage of

hardware idioms. A code improver can often perform some flow analysis and recognize more hardware idioms.

The idea of a code improver is not new. "Peephole optimizers" are well known.^{3,4} One C compiler for the PDP-11 computer has had a code improver for many years.* The section of Ref. 5 on the FINAL compilation pass describes a code improver used internally by a BLISS-11 compiler.

The code improver described here makes the portable C compiler usable as the workhorse compiler in a serious production environment. Measurements indicate that for typical programs the improver reduces the number of bytes in the instruction stream by 10 to 23 percent; the novel technique reported here accounts for as much as one-third of the reduction. The time required to execute the code is also reduced by 4 to 8 percent. The improver produces good three-address code from the two-address code generated by the compiler.

II. IMPROVING CODE FOR THE VAX-11

An existing improver of code compiled for the PDP-11 served as a model and outline for the VAX-11 code improver. The improver reads a file of assembly language and divides the file into segments corresponding to C procedures. For each procedure it constructs a doubly-linked list of the instructions and label definitions, with additional links for references to labels. The improver then combs the list, repeatedly trying to apply any one of several incremental transformations. The transformations satisfy a principle of optimality: Any local improvement is guaranteed to be a global improvement at least as large, and conversely, if the program as a whole can be made smaller or faster, then there is a collection of local changes which will account for the improvement. When no further transformation can be made, the improver prints the list and moves on to the next procedure. Many of the transformations depend little on the particular machine. Straightforward adaptation of the old program yielded code to transitively close jumps to jumps, delete instructions that immediately follow unconditional jumps, delete jumps to the immediately following instruction, remove unreferenced or redundant labels, merge common tail sequences, move basic blocks to the point of sole use, and interchange physical order of the consequent and alternative to a test. Simple modifications also produced a program to rotate loops to place a single conditional jump at the bottom, handle skips over jumps, eliminate redundant setting of the condition code, move common antecedents of jumps into the merged tail, eliminate constant tests or tests which are subsumed by a preceding test, exploit add-compare-

* PDP is a registered trademark of Digital Equipment Corporation.

Table I—Translations of $a = b + c$;

PDP-11	VAX-11	Improved VAX-11
mov b, r0		
add c, r0	addl3 b, c, r0	
mov r0, a	movl r0, c	addl3 b, c, a

branch (“DO-loop”) instructions, and remember values already in registers.

III. THREE ADDRESSES FROM TWO

Fully utilizing the three-address instructions available on the VAX-11 presented a new challenge. Table I illustrates a common opportunity to use a three-address instruction. In this example the variables a , b , c are assumed to reside in memory (either global or local) and not in registers. The first column gives a translation for the PDP-11 that cannot be improved in either time or space. (If some of the variables reside in registers, then improvements are possible.) Both the production and the portable C compiler for the PDP-11 produce this translation without the aid of a code improver. The second column contains the code generated by the portable C compiler for the VAX-11. The compiler saves one instruction by doing the work of the first two PDP-11 instructions in one three-address VAX-11 instruction. However, it will not generate the code in the right-most column, where a single instruction suffices for the whole statement. Internally the portable C compiler uses a binary tree to represent each parsed statement. The height of a binary tree with three external nodes (each explicit variable is represented by an external node) must be at least two. Furthermore, the pattern-matching algorithms used by the compiler are restricted to subtrees of height one. (The pattern match has since been generalized to match subtrees of arbitrary height.) Thus the compiler generates two separate instructions for this case. It does have the flexibility to use an instruction with three addresses, but the destination operand of a three-address instruction must always be one of the compiler’s temporary locations, usually a register. The challenge to the code improver is to recognize situations like this one and change the code appropriately.

Table II illustrates a complication. Here the addition and assignment are embedded as an expression whose value is passed as an actual argument in a procedure call. Although the same **addl3** and **movl** instructions appear together, the value in r0 is needed later and r0 cannot be elided. In standard terminology, the value in register r0 is *live*, or alternatively register r0 is *busy*. The improver can elide register usage only when the value in the register is known to be *dead*, or the register is *free*.

For an arbitrary program, determining which registers are free at a given point requires a fair amount of work. The register usage and flow of control through any part of the program can effect whether or not a register is busy in any other part of the program. Code generated by the portable C compiler has a property that makes busy/free analysis much simpler. All registers are free any time the compiler generates a backward branch instruction. The portable C compiler generates code on line, completely translating the current expression or statement before proceeding to the following expression or statement. The use of a temporary expression always occurs physically after its generation. Thus the entire busy/free analysis can be done in a single backward scan over the generated code. The backward scan marks a register busy each time the register is read or used as a source operand. Some instruction occurring closer to the front of the file must have put a live value into the register, or else the register would contain garbage. Analogously, the backward scan marks a register free each time the register is written or used as a destination operand. Since the write destroys whatever used to be in the register, no one could have wanted that dead value.

The backward scan must take precautions to record each use of a temporary register, including the implicit uses. The return instruction `ret` implicitly reads `r0`, the register in which C code returns function values. Thus `r0` is busy just before each `ret`. The overall code-generation strategy of the compiler assumes that each procedure call instruction `calls` writes all the temporary registers. Thus all the temporary registers are free just before a procedure call.

The busy/free information can also be used to eliminate dead code. An instruction that writes only into free registers does no useful work, except possibly for the side effects it causes. If the address computations contain no side effects, then only the condition code could matter. The condition code is set by each nonbranch instruction, so the condition code itself is free unless the instruction which logically follows is a conditional branch.

The backward scan must also be careful with code generated from conditional expressions. There can be no busy registers at the time of a backward jump, as noted earlier. Since the compiler performs no

Table II—Translations of $f(a = b + c)$;

PDP-11	VAX-11	"Improved" (but wrong) VAX-11
<code>mov b, r0</code>		
<code>add c, r0</code>	<code>addl3 b, c, r0</code>	
<code>mov r0, a</code>	<code>movl r0, a</code>	<code>addl3 b, c, a</code>
<code>mov r0, (sp)</code>	<code>pushl r0</code>	<code>pushl r0</code>
<code>jsr pc, f</code>	<code>calls \$1, f</code>	<code>calls \$1, f</code>

Table III—Translation of $x = a ? b : c;$

testl a
jeql L100
movl b, r0
jbr L101
L100: movl c, r0
L101: movl r0, x

interstatement data-flow analysis (and in particular does not recognize common subexpressions), there can be no busy registers at the time of a forward jump generated from an entire C statement. Since labels exist only because jump instructions branch to them, these two facts might suggest that a register cannot be busy at any label, either. A register can, however, be busy at a forward jump (and thus at a label) with one of the values of a conditional expression. Table III illustrates one such situation.

Even though the instruction **movl c,r0** writes r0, the register is busy at the **jbr** because (if *a* is true) it contains the value of *b* to be stored into *x*. Thus the busy/free status of each register must be associated with each label as the label is passed during the backward scan, and retrieved from the corresponding label at each jump. This can be done efficiently by keeping a bit vector associated with each label, initializing all the bits to “free,” and recording busy registers as labels are passed. Because backward jumps have no busy registers and the backward scan encounters the destination label of a forward jump before seeing the jump itself, the bits will always be correct.

In general the code improvements other than insertion of three-address instructions and elimination of dead code by consulting the busy/free information destroy the property that no temporary register is busy at a backward jump. This implies that using a single backward sweep over the code for the entire procedure to determine busy/free is valid only once, at the beginning before other improvements are tried. Fortunately, once is enough.

IV. OTHER USES OF THE BACKWARD SCAN

The backward prescan is also a good time to recognize hardware idioms. The VAX-11 has a number of instructions to set, clear, and test single bits, and to extract contiguous bit fields of arbitrary size. Appropriate uses of these instructions are often concealed in C with various Boolean or shift-and-mask operators or sequences of operators. Computing with the addressing modes by using instructions in which an address is used as an immediate operand often saves time and space. Powerful addressing modes often depend heavily on register usage, and the backward pass is already computing this information. Since the backward scan is performed only once, time will not be

wasted searching for hardware idioms more than once, as part of the general iterative improvement strategy. Table IV gives some example improvements.

V. DEVICE DRIVERS

On the VAX-11, the control and data registers for input/output devices lie in the memory address space. Programs manipulate the registers in much the same way as they manipulate memory, and the assembly-language code for a device driver cannot be identified solely by its form. However, certain instructions and addressing modes do not work properly when addressed to device registers. Generally these are exactly the instructions and addressing modes that the code improver wants to introduce. For example, neither of the first two improvements in Table IV is legal on a device register. Thus the code improver must be told when it is improving the code for a device driver, so it can avoid those improvements that cause problems. Reading or writing a device register typically has side effects that are different from reading or writing a memory location, and other hardware considerations such as bus widths, circuit board area, or number of words of microcode are often important. Yet from a software viewpoint such special cases are irritating and error prone, and it would be desirable to get rid of the complication.

VI. CONCLUSIONS

A single backward scan enables the code improver to determine register usage and introduce three-address instructions where appropriate. The backward scan takes advantage of the fact that all registers are free at each backward jump, a property that would otherwise be considered a weakness in the compiler. The single backward scan also recognizes hardware idioms at a lower cost than previous algorithms.

Table IV—Improvements using VAX-11 hardware idioms

C Code	Raw Translation	Improved Translation
int a; a = 0x8000;	bisl2 \$0x8000, a	jbss \$15, a, L100 L100:
int a, b; b = (a >> 12) & 0xF;	ashl \$-12, a, r0 bicl2 \$-16, r0 movl r0, b	extzv \$12, \$4, a, b
int *p, *q; q = &p [f(x)];	pushl x calls \$1, f ashl \$2, r0, r0 addl2 p, r0 movl r0, q	pushl x calls \$1, f movl *p [r0], q

REFERENCES

1. S. C. Johnson, "A Portable Compiler: Theory and Practice," Conference Record Fifth Ann. ACM Conf. Principles of Programming Languages, Tucson, AZ (January 1978), pp. 97-104.
2. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Englewood Cliffs, New Jersey: Prentice-Hall, 1978.
3. W. M. McKeeman, "Peephole Optimization," *Commun. ACM*, 8 (July 1965), pp. 443-4.
4. A. V. Aho and J. D. Ullman, *Principles of Compiler Design*, Reading, Massachusetts: Addison-Wesley, 1977.
5. W. Wulf, R. K. Johnson, C. B. Weinstock, S. O. Hobbs, and C. M. Geschke, *The Design of an Optimizing Compiler*, New York: Elsevier, 1975.

Rain Margin Improvement Using Resource Sharing in 12-GHz Satellite Downlinks

By A. S. ACAMPORA

(Manuscript received August 6, 1980)

In this paper we consider the effectiveness of sharing a small pool of reserved time slots of a Time-Division-Multiple-Access (TDMA) frame among a large number of ground stations to overcome rain fading. With this approach, a system dynamically assigns time slots from the reserved pool to ground stations experiencing fade depths above the built-in margin. Powerful error correcting codes can be introduced to occupy the extra time slots, providing 10 dB or more of extra fade margin. Because a large number of ground stations are competing for the limited reserved pool, blockage can occur if the number of simultaneous fades exceeds the maximum number that can be accommodated. Some factors that influence the effectiveness of resource sharing are the mutual fade statistics at the various sites, the traffic distribution within the network, the number of earth stations, the size of the reserved pool, and the rain outage objective. Since the mutual fade and traffic statistics are unavailable, we develop models that can be used to find a conservative bound on the required size of the reserved pool. The rain model accounts for diurnal, seasonal, and geographical correlation among attenuation events. Results for a maximum resource-sharing gain of 10 dB show that reserving six percent of the time slots ensures a realized fade gain in excess of 9 dB for a down-link outage objective of 0.005 percent if there are more than 50 ground stations in the network, each with two percent or less of the traffic.

I. INTRODUCTION

In an earlier paper,¹ a shared-resource concept was described for increasing the rain fade margin of a digital satellite system by as much as 10 dB above the design fade margin. With this approach, unused time slots of the Time-Division-Multiple-Access (TDMA) frame are made available to ground stations experiencing rain fading above the

built-in fade margin, and are relinquished when the fade event has ended. Error-correcting coding is introduced to occupy the extra time slots, thereby reducing the carrier-to-noise ratio (CNR) required to maintain the threshold bit-error-rate (BER). Low rain outage is therefore achieved without radiating excessive down-link power. Not only does this conserve satellite power, but, also, interference into the systems of other users of the geosynchronous orbit is minimized. In such an application, the operating speed of the decoder is much lower than the transponder data rate by virtue of the low TDMA duty cycle associated with each ground station.

Because of the infrequency of simultaneous deep fading at multiple sites, a small pool of reserved time slots can often protect a large number of ground stations. The degree of protection so provided is the subject of this current work; we shall evaluate the reduction in rain margin required to achieve a given outage objective when all ground stations in the network are competing for a limited number of shared resources. We restrict our attention to the power-limited down-link since up-link fading can usually be overcome by means of up-link power control.¹ A convolutional code yielding a maximum power saving of 10 dB is assumed throughout. Results are directly applicable to either a wide-area coverage system or a single-scanning beam system,² but the modeling and analytical approach can be extended to study TDMA systems that are fixed-beam satellite-switched, multiple scanning beam, or hybrid-fixed scanning beam.³⁻⁶

Figure 1 shows a typical sequence for interconnecting the various spot-beam footprints. Each interconnection contains one or more time slots during which ground stations within the connected regions communicate on a sequential basis. Although all TDMA time slots can be made available to accommodate normal network traffic demand, the approach taken here specifically reserves a certain number of time slots, shown at the end of the frame, for use exclusively during the rain fade events. By so doing, we have reduced the traffic-handling capability of the system by a small percentage, while guaranteeing that

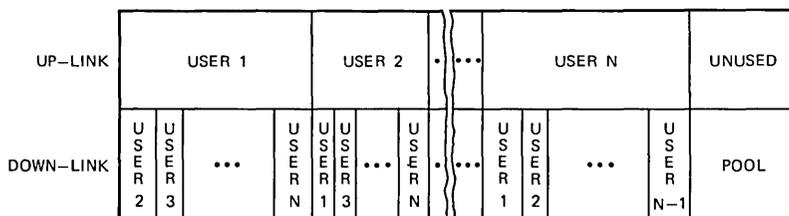


Fig. 1—Typical switching frame showing the interconnections between the system ground stations. An unused pool of reserved time slots is also shown.

some extra time slots will be available as needed to maintain reliability of circuits *already in use* when fade events occur.

The number of reserved slots necessary to provide a specified degree of protection at all stations is dependent upon several factors. Clearly, the utility of resource sharing is dependent upon the joint fade statistics at multiple geographically remote sites; a small number of reserved slots cannot provide protection if the probability of simultaneous excess attenuation at many sites, conditioned upon the occurrence of excess attenuation at any one site, is high. Unfortunately, experimental joint-fade statistics at multiple remote sites are unavailable, and we must resort to modeling to obtain quantitative results.

The utility of resource sharing also depends on the number of ground stations in the network and the traffic distribution among those ground stations. If the number of ground stations is small, then the fraction of TDMA time slots which must be reserved to protect even one site must be large, resulting in an inefficient solution to the rain fade problem. Also, if the traffic distribution is highly nonuniform such that a few ground stations carry a disproportionately large volume of traffic, then again it becomes impractical to reserve enough time slots to protect this small number of large users. In such an event, it might be desirable to protect the large users by some other technique, such as larger antennas or site diversity, and employ resource sharing for the exclusive protection of the much larger number of small users.

Similarly, the effectiveness of resource sharing depends on the geographical distribution of ground stations relative to the profiles of high rain-attenuation regions, and upon the volume of traffic carried by ground stations in high rain-attenuation regions. An additional factor is the relationship between the busy hour, when all time slots not reserved for resource sharing might be expected to be in heavy demand, and the time-of-day occurrence of significant rain-attenuation events.

We shall present both a multiple-site rain-attenuation model and a population-dependent traffic model, upon which is based the subsequent predicted performance of the resource sharing concept. All assumptions implicit in this modeling are addressed in detail in Section II. Section III contains the mathematical analysis of outage based on the rain and traffic models. Section IV contains numerical results of this analysis; the effects of the various factors are presented parametrically. A typical result shows that for a 12/14-GHz network of 100 identical ground stations, a shared resource reservation equal to six percent of the total transponder time slots will provide an outage of ½-hour per year with 9 dB less rain margin than otherwise needed; diurnal, seasonal, and geographical dependencies of joint rain-fade statistics are accounted for in this prediction.

II. JOINT RAIN ATTENUATION AND TRAFFIC MODELS

2.1 Joint rain attenuation model

Figure 2 shows statistics for single-site rain attenuation. It plots the fraction of time that rain fading exceeds the level of the abscissa averaged over one contiguous 12 month interval. This fraction can be interpreted as the probability that any given rain-attenuation level is exceeded. From such single-site curves, we develop a model to be used for predicting joint outages at multiple geographically remote sites.

At first glance, one might assume that if two sites are widely separated, then rain events at the two sites occur independently. However, this cannot be the case because in the 12-GHz satellite band, rain attenuation in excess of 5 dB is typically associated with thunderstorm activity which produces intense rainfall. Periods of thunderstorm activity are typically restricted to a four-month interval lasting from June through September, and to an interval of six hours each day lasting from 1 PM to 7 PM. Thus, if we are told that the rain attenuation at one of the sites is, say, 10 dB, then the probability of simultaneous deep fading at the second site must be higher than its yearly average because, at that moment, we are likely to be in the interval when thunderstorm activity normally occurs. Thus, knowledge

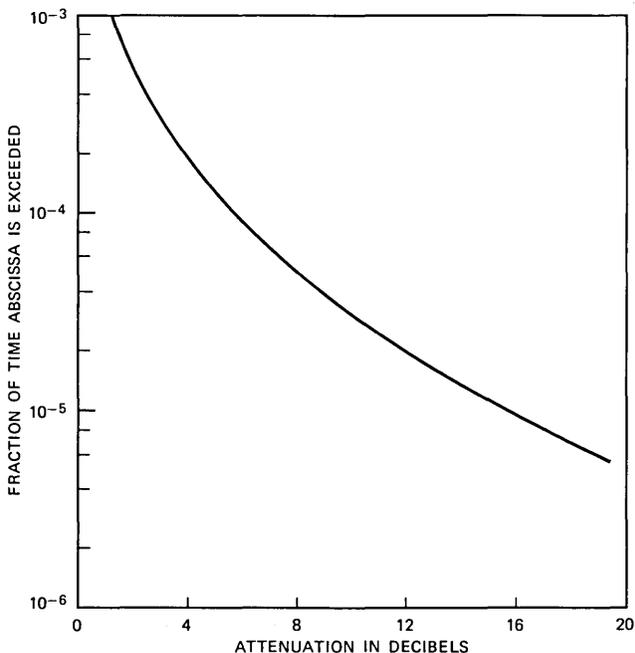


Fig. 2—Typical single-site attenuation curve showing the fraction of time in one year that the attenuation exceeds the abscissa.

of rain attenuation at one site affects the probability of simultaneous attenuation at the second site, and the two events are not independent.

Suppose, however, that we assume that all of the probability of Fig. 1 is attributed to the period of thunderstorm activity. Then, if we restrict our attention to this period only (it is only during this period when resource sharing is needed to combat fading), the probability that attenuation exceeds a given level is about 12 times the yearly averaged value of Fig. 2 (4 out of 12 months per year, 6 out of 24 hours per day). Let $p_1(A)$ be the yearly averaged probability that attenuation exceeds level A at site No. 1. Thus, at a given instant of time t ,

$$p_1(A | t \text{ within thunderstorm period}) = \alpha p_1(A), \quad (1)$$

$$p_1(A | t \text{ outside thunderstorm period}) = 0. \quad (2)$$

In (1), for the illustrative example given above, $\alpha = 12$. Choosing a larger value of α enhances the probability of multiple simultaneous fades since the yearly-averaged probability is then attributed to a narrower interval. In what follows, we conservatively assume the thunderstorm periods to be 3 out of 12 months and 4 out of 24 hours, yielding $\alpha = 24$. The factor α applies at all sites in the satellite network.

Now, given that t is within the thunderstorm-activity period, the events of attenuation exceeding level A at site 1 and level B at site 2 may be assumed to be independent if the two sites are widely separated. For example, knowledge that attenuation exceeds level A in New York during the thunderstorm-activity period provides no information concerning the event that attenuation exceeds level B in Denver, since different independent storms are involved. Thus, for t within the period of thunderstorm activity,

$$p_{1,2}(A, B | t) = p_1(A | t)p_2(B | t) \quad (3)$$

$$= \alpha^2 p_1(A)p_2(B), \quad (4)$$

where $p_{1,2}(A, B | t)$ is the probability that attenuation at site 1 exceeds level A and attenuation at site 2 exceeds B simultaneously at time t . Equation (4) is readily generalized for the case of an arbitrary number of widely separated sites.

For two closely spaced sites, another degree of attenuation event correlation is assumed beyond the seasonal and diurnal correlations just addressed. Here, attenuation at the two sites may be produced by the same storm. We assume that, for closely spaced sites, all fades in excess of some thunderstorm characteristic level A_0 always occur within an H -hour interval of each other, where H is much larger than the typical several minute duration of deep fades. Suppose that a fade of level $A > A_0$ occurs at some site. Then, at a neighboring site, the probability that a fade of level $B > A_0$ is simultaneously occurring is given by $C(B)$, a fade level dependent constant over the H -hour

window. At this second site, the yearly averaged probability of a fade of level $B > A_0$ is attributed exclusively to the H -hour intervals surrounding the events $A > A_0$ at the original site. Then,

$$p_2(B) = \kappa \frac{H}{T} C(B), \quad (5)$$

where T is the number of hours in a year and κ is the average number of events per year that the attenuation level exceeds A_0 . Now,

$$p_{1,2}(A, B | t) = p_2(B | A, t)p_1(A | t). \quad (6)$$

For $A, B > A_0$,

$$p_2(B | A, t) = C(B) = \frac{T}{\kappa H} p_2(B). \quad (7)$$

Thus, defining $\beta = T/\kappa H\alpha$,

$$p_{1,2}(A, B | t) = \beta\alpha^2 p_1(A)p_2(B). \quad (8a)$$

For L closely spaced sites, (8) generalizes to

$$p_{1,2,\dots,L}(A_1, A_2, \dots, A_L) = \beta^{L-1}\alpha^L p_1(A_1) \dots p_L(A_L). \quad (8b)$$

In the following, we will allow the "geographical correlation factor" β to vary between 1 and 6. The extreme value $\beta = 1$ implies that $\kappa H\alpha = T$, and that, within the thunderstorm period, fades occur independently. Recognizing that α is equal to the number of hours in a year divided by the annual number of hours in the thunderstorm-activity period, we see that for $\beta = 1$, the average number of fades per year multiplied by the uncertainty window H equals the number of hours in the thunderstorm-activity period. Hence, knowledge that a fade of level $A > A_0$ is occurring at some site does not restrict the interval within the thunderstorm-activity period when a fade of level $B > A_0$ may occur at a neighboring site. Similarly, the extreme value $\beta = 6$ implies, for example, that knowledge of deep fading at a given site pinpoints the two out of four thunderstorm hours per day and the average of one out of three days during the thunderstorm period when intense rainfall might occur in the region surrounding that site. Conditioned upon an attenuation event at some site, the probability of an attenuation event at a second site within the surrounding region is then six times higher than would be expected from diurnal and seasonal-correlation considerations alone.

The conservatism of the values $\alpha = 24$ and $\beta = 6$ can be demonstrated by applying eq. (8a) to experimental attenuation data obtained with site diversity such as appears in Ref. 7. Averaging (8a) over one full year, we obtain

$$p_{1,2}(A, B) = \frac{1}{\alpha} p_{1,2}(A, B/t) = \beta\alpha p_1(A)p_2(B). \quad (9)$$

For attenuation greater than 5 dB, the site diversity measurements are somewhat more optimistic than predicted by eq. (9). This observation, coupled with ground station separation much wider than used for the site diversity measurements and the physical interpretations given to α and β above, confirms the conservatism of the approach.

To apply the above rain-attenuation event model, we need to know the yearly-averaged attenuation statistics at the location of each site in the network. The following simplification is invoked: We divide the continental United States into three regions such that the yearly averaged attenuation statistics that apply at a representative site in one region are typical for all sites in that region (see Fig. 3). Los Angeles is chosen as representative of the western region (Region 1), where rain attenuation occurs infrequently. New York is chosen as representative of the northeast central region (Region 2), throughout which rain attenuation is moderate. Finally, Atlanta is chosen as representative of the southeast central region (Region 3), where rain attenuation occurs frequently. When applying the model, we make the pessimistic assumption that the “geographical correlation factor,” β , applies at all sites throughout an entire region; for sites located in different regions this factor is neglected. Thus, for this model, “closeness” of two sites is defined by whether both sites are within the same region. The errors incurred because of sites located near each other but on opposite sides of a regional boundary are more than offset by the large “correlation distances” assumed.

2.2 The traffic model

The traffic model used in this analysis is based on a rank ordering of the 100 most populous continental United States cities, as shown in



Fig. 3—Three regional maps of the United States. Regional boundaries are selected such that an attenuation curve for a representative ground station in each region applies throughout that region.

Table I.⁴ We assume that the traffic between any two of these cities is inversely proportional to the product of their indices. Traffic between two cities closer than 500 miles is excluded from the network. Under these assumptions, Region 1 offers 24 percent of the total satellite traffic, while Regions 2 and 3 offer 63 percent and 13 percent, respectively.⁴

The number of ground stations serving a given region is assumed to be proportional to the traffic offered by that region. For example, if the network contains 100 ground stations, then 24 are located in Region 1, 63 are in Region 2, and 13 are in Region 3. All ground stations are assumed to carry identical traffic cross sections. Thus, for the purposes of analysis, the traffic differences among regions or among cities within each region are accounted for by assigning more or fewer identical capacity ground stations, as the case may be, to accommodate the traffic. The modeling and analysis can be extended to networks containing a mixture of large and small traffic ground stations if the large users are protected by site diversity and if resource sharing is applied to protect only the larger number of small users. This extension has not been carried out, however, and the numerical results to be presented are valid only for the former case.

Two additional assumptions are also made. First, all ground stations in a given region are assumed to be identical, i.e., have the same built-in rain fade margin; the built-in margins of ground stations in different regions need not be the same, however. Second, we conservatively assume that the traffic-busy period, when all time slots not reserved for resource sharing are in full-time use, coincides with the thunderstorm-activity period. Thus, the results are intentionally made to be pessimistic since it is precisely when resource sharing is needed that extra time slots not contained in the reserve pool are unavailable; during traffic off-peak hours, when additional slots are available, resource sharing is not needed because rain fading does not occur.

III. RAIN OUTAGE ANALYSIS

3.1 General approach

Consider the three-regional map shown in Fig. 3. Let there be N_1 ground stations or sites located in Region 1, each of which has a built-in fade margin of A_1 dB. Similarly, let there be N_2 and N_3 sites located in Regions 2 and 3, respectively, with fade margins of A_2 and A_3 dB. The extra fade margin provided by resource sharing if extra time slots are available is M dB, that is, M dB is the maximum extra margin provided by the coding approach employed. All sites carry the same volume of traffic, and enough time slots are reserved to accommodate K simultaneous fades. We wish to find the yearly-averaged probability that a given site is operational when all sites in the network are competing, as needed, for the reserved resource-sharing slots. With no

Table I—Rank-ordering of the 100 largest U.S. cities by population

Rank	City	Re- gion	Rela- tive Traf- fic	Rank	City	Re- gion	Rela- tive Traf- fic	Rank	City	Re- gion	Rela- tive Traf- fic	Rank	City	Re- gion	Rela- tive Traf- fic
1	New York	1	60	26	San Bernadino	3	4	51	Ft. Lauderdale	2	2	76	Canton	1	1
2	Los Angeles	3	43	27	Indianapolis	1	4	52	Greensboro	2	2	77	Davenport	1	1
3	Chicago	1	25	28	San Jose	3	3	53	Salt Lake City	3	2	78	El Paso	3	1
4	Philadelphia	1	17	29	New Orleans	2	4	54	Allentown	1	1	79	New Haven	1	1
5	Detroit	1	18	30	Tampa-St. Pete	2	3	55	Nashville	1	2	80	Tucson	3	1
6	San Francisco	3	15	31	Portland	3	3	56	Omaha	1	2	81	W. Palm Beach	2	1
7	Washington	1	10	32	Phoenix	3	3	57	Grand Rapids	1	2	82	Worcester	1	1
8	Boston	1	10	33	Columbus	1	3	58	Youngstown	1	2	83	Wilkes-Barre	1	1
9	Pittsburgh	1	10	34	Providence	1	2	59	Springfield	1	1	84	Peoria	1	1
10	St. Louis	1	10	35	Rochester	1	2	60	Jacksonville	2	2	85	Utica	1	1
11	Baltimore	1	8	36	San Antonio	2	3	61	Richmond	1	2	86	York	1	1
12	Cleveland	1	8	37	Dayton	1	3	62	Wilmington	1	1	87	Bakersfield	3	1
13	Houston	2	8	38	Louisville	1	3	63	Flint	1	2	88	Little Rock	1	1
14	Minneapolis	1	7	39	Sacramento	3	2	64	Tulsa	1	2	89	Columbia	2	1
15	Dallas	2	7	40	Memphis	1	3	65	Orlando	2	2	90	Lancaster	1	1
16	Seattle	3	6	41	Fort Worth	2	3	66	Fresno	3	1	91	Beaumont	2	1
17	Anaheim	3	5	42	Birmingham	2	3	67	Tacoma	3	2	92	Albuquerque	3	1
18	Milwaukee	1	5	43	Albany	1	2	68	Harrisburg	1	1	93	Chattanooga	1	1
19	Atlanta	2	5	44	Toledo	1	2	69	Charlotte	2	1	94	Trenton	1	1
20	Cincinnati	1	5	45	Norfolk	1	2	70	Knoxville	1	1	95	Charleston	2	1
21	San Diego	3	4	46	Akron	1	2	71	Wichita	1	1	96	Binghamton	1	1
22	Buffalo	1	4	47	Hartford	1	2	72	Bridgeport	1	1	97	Greenville	2	1
23	Miami	2	5	48	Oklahoma City	1	2	73	Lansing	1	1	98	Reading	1	1
24	Kansas City	1	4	49	Syracuse	1	2	74	Mobile	2	1	99	Austin	2	1
25	Denver	3	4	50	Gary	1	2	75	Ventura	3	1	100	Shreveport	2	1

loss in generality, we find this probability for a site located in Region 1; a simple permutation of indices enables this result to be applied in either of the two remaining zones.

At any site within Region 1, three disjoint events may occur at any point in time: (1) the fade depth F may be less than A_1 , (2) F may be between A_1 and $A_1 + M$, and (3) F may exceed $A_1 + M$. Call these events E_1 , E_2 , and E_3 . Then

$$P(\text{operational} | t) = \sum_{l=1}^3 P(\text{operational} | E_l, t)P(E_l | t). \quad (10)$$

Clearly,

$$P(\text{operational} | E_1, t) = 1, \quad (11)$$

$$P(\text{operational} | E_3, t) = 0, \quad (12)$$

$$\begin{aligned} P(E_1 | t) \\ = \begin{cases} 1 - \alpha p(A_1), & t \text{ within thunderstorm period,} \\ 1, & \text{otherwise,} \end{cases} \end{aligned} \quad (13)$$

$$\begin{aligned} P(E_2 | t) \\ = \begin{cases} \alpha p(A_1) - \alpha p(A_1 + M), & t \text{ within thunderstorm period,} \\ 0, & \text{otherwise.} \end{cases} \end{aligned} \quad (14)$$

Thus, for t within the thunderstorm period,

$$\begin{aligned} P(\text{operational} | t) = 1 - \alpha p(A_1) \\ + \alpha P(\text{operational} | E_2, t)[p(A_1) - p(A_1 + M)], \end{aligned} \quad (15)$$

and for t outside the thunderstorm period,

$$P(\text{operational} | t) = 1. \quad (16)$$

From (15) and (16), we obtain the result that, averaged over an entire year,

$$\begin{aligned} P(\text{not operational}) = p(A_1) \\ - P(\text{operational} | E_2, \hat{t})[p(A_1) - p(A_1 + M)], \end{aligned} \quad (17)$$

where $\hat{t} = \{t \text{ within thunderstorm-activity period}\}$.

3.2 Derivation of $P(\text{operational} | E_2, t)$

Now, for t within the thunderstorm period, and conditioned on the event E_2 , a particular site will be operational if the number of other sites which need to use the reserve pool is less than $K - 1$. Also, if the number of other sites which need to use the pool is equal to $j \geq K$, then the probability that a particular site is one of the K sites permitted to use the pool is equal to $K/(j + 1)$.

Conditioned on the event E_2 at a particular site, the probability that i_1 additional sites in Region 1 need to use the reserved pool, where $0 \leq i_1 \leq N_1 - 1$, is given by

$$P(i_1 \text{ in Region 1} | E_2) = {}_{N_1-1}C_{i_1}[(\alpha\beta\rho_1)^{i_1}(1 - \alpha\beta\rho_1)^{N_1-1-i_1}], \quad (18)$$

where

$$\rho_1 = P(A_1) - P(A_1 + M) \quad (19)$$

and

$${}_nC_m = \frac{n!}{m!(n-m)!}. \quad (20)$$

Conditioned on the event E_2 at a particular site, the probability that i_2 sites in Region 2 need to use the reserved pool, where $0 \leq i_2 \leq N_2$, is given by the unconditional probability during the thunderstorm-activity period since the events are assumed to be independent. For $i_2 \geq 1$, we find this probability by first defining the three events:

- L_1 : {a particular set of $i_2 - 1$ sites in Region 2 need to use the reserved pool},
- L_2 : {an additional site in Region 2, not included in the particular set of L_1 , needs to use the reserved pool},
- L_3 : {none of the remaining $N_2 - i_2$ sites of Region 2 need to use the reserved pool}.

Then,

$$P(L_1 \cap L_2 \cap L_3) = P(L_1 \cap L_3 | L_2)P(L_2) \quad (21)$$

$$= P(L_1 | L_2)P(L_3 | L_2)P(L_2) \quad (22)$$

$$= (\alpha\beta\rho_2)^{i_2-1}(1 - \alpha\beta\rho_2)^{N_2-i_2}(\alpha\rho_2), \quad (23)$$

where

$$\rho_2 = P(A_2) - P(A_2 + M). \quad (24)$$

Thus, for $i_2 \geq 1$,

$$P(i_2 \text{ in Region 2} | E_2) = {}_{N_2}C_{i_2}(\alpha\rho_1)(\alpha\beta\rho_2)^{i_2-1}(1 - \alpha\beta\rho_2)^{N_2-i_2} \\ = \frac{1}{\beta} \left[{}_{N_2}C_{i_2}(\alpha\beta\rho_2)^{i_2}(1 - \alpha\beta\rho_2)^{N_2-i_2} \right]. \quad (25)$$

The probability that none of the sites in Region 2 need to use the reserved pool is given by

$$P(0 \text{ in Region 2} | E_2) = 1 - \sum_{i_2=1}^{N_2} P(i_2 \text{ in Region 2} | E_2). \quad (26)$$

Substituting (25), and invoking the binomial sum formula,

$$P(0 \text{ in Region 2} | E_2) = 1 - \left(\frac{1 - (1 - \alpha\beta\rho_2)^{N_2}}{\beta} \right). \quad (27)$$

Similarly, for Region 3 and for $i_3 \geq 1$,

$$P(i_3 \text{ in Region 3} | E_2) = {}_{N_3}C_{i_3}(\alpha\rho_3)(\alpha\beta\rho_3)^{i_3-1}(1 - \alpha\beta\rho_3)^{N_3-i_3} \quad (28)$$

and

$$P(0 \text{ in Region 3} | E_3) = 1 - \left(\frac{1 - (1 - \alpha\beta\rho_3)^{N_3}}{\beta} \right), \quad (29)$$

where

$$\rho_3 = P(A_3) - P(A_3 + M). \quad (30)$$

Clearly, conditioned on the thunderstorm period, the event $\{i_2 \text{ sites in Region 2 need to use the reserved pool}\}$ and the event $\{i_3 \text{ sites in Region 3 need to use the reserved pool}\}$ are independent, $0 \leq i_2 \leq N_2$, $0 \leq i_3 \leq N_3$.

Returning to (15), for t within the thunderstorm period, the probability that a particular site is operational, given attenuation between A_1 and $A_1 + M$ dB, can be expressed as the union of the events $\{V_{i_1, i_2, i_3, s}\}$:

$$\{V_{i_1, i_2, i_3, s}\} = \{\mathcal{S}_{i_1} \cap \mathcal{S}_{i_2} \cap \mathcal{S}_{i_3} \cap S\}, \quad (31)$$

$$i_1 = 0, \dots, N_1 - 1, \quad i_2 = 0, \dots, N_2, \quad i_3 = 0, \dots, N_3,$$

where

$$\{\mathcal{S}_{i_1}\} = \{i_1 \text{ additional sites in Region 1 need to use the reserved pool}\}, \quad (32)$$

$$\{\mathcal{S}_{i_2}\} = \{i_2 \text{ sites in Region 2 need to use the reserved pool}\}, \quad (33)$$

$$\{\mathcal{S}_{i_3}\} = \{i_3 \text{ sites in Region 3 need to use the reserved pool}\}, \quad (34)$$

$$\{S\} = \{\text{the particular site has been assigned time slots from the reserved pool}\}. \quad (35)$$

Clearly, for $i'_i \neq i''_i$ or $i'_2 \neq i''_2$ or $i'_3 \neq i''_3$, the events $\{V_{i'_1, i'_2, i'_3, s}\}$ and $\{V_{i''_1, i''_2, i''_3, s}\}$ are disjoint, and the probability of the union of events (31) is equal to the sum of the probabilities of all events in that union.

Let us define the function

$$f(m) = \begin{cases} 1, & m \leq K - 1, \\ \frac{K}{m}, & m \geq K. \end{cases} \quad (36)$$

Then, we conclude that for $t = \hat{t}$,

$$\begin{aligned}
& P(\text{operational} | E_2, \hat{t}) \\
&= \sum_{i_1=0}^{N_1-1} \sum_{i_2=1}^{N_2} \sum_{i_3=1}^{N_3} \{f(i_1 + i_2 + i_3)P(i_1 \text{ in Region 1} | E_2) \\
&\quad \cdot P(i_2 \text{ in Region 2} | E_2)P(i_3 \text{ in Region 3} | E_2)\} \\
&\quad + \left[\frac{(1 - \alpha\beta\rho_2)^{N_2} + \beta - 1}{\beta} \right] \sum_{i_1=0}^{N_1-1} \sum_{i_3=1}^{N_3} \{f(i_1 + i_3) \\
&\quad \cdot P(i_1 \text{ in Region 1} | E_2)P(i_3 \text{ in Region 3} | E_2)\} \\
&\quad + \left[\frac{(1 - \alpha\beta\rho_3)^{N_3} + \beta - 1}{\beta} \right] \sum_{i_1=0}^{N_1-1} \sum_{i_2=1}^{N_2} \{f(i_1 + i_2) \\
&\quad \cdot P(i_1 \text{ in Region 1} | E_2)P(i_2 \text{ in Region 2} | E_2)\} \\
&\quad + \left[\frac{(1 - \alpha\beta\rho_2)^{N_2} + \beta - 1}{\beta} \right] \left[\frac{(1 - \alpha\beta\rho_3)^{N_3} + \beta - 1}{\beta} \right] \\
&\quad \cdot \sum_{i_1=0}^{N_1-1} \{f(i_1)P(i_1 \text{ in Region 1} | E_2)\}, \tag{37}
\end{aligned}$$

where $P(i_1 \text{ in Region 1} | E_2)$, $P(i_2 \text{ in Region 2} | E_2)$, and $P(i_3 \text{ in Region 3} | E_3)$ are given, respectively, by eqs. (18), (25), and (28).

Substituting (37) into (17) yields the desired result for the yearly-averaged probability that a particular site is unavailable versus the built-in fade margins A_1 , A_2 , and A_3 .

IV. RESULTS

We now apply the model of Section II and the analysis of Section III to investigate the utility of resource sharing for a 12-GHz satellite. Initially, we shall assume a satellite location of 100 degrees West longitude. A total of 100 ground stations is assumed, and the thunderstorm-activity factor $\alpha = 24$. Results are obtained for various values of K and β . The yearly-averaged attenuation data used for Regions 1, 2, and 3 are based upon S. Lin's attenuation model^{8,9} for converting long-term rain-rate data, obtained from the United States Weather Bureau, into attenuation predictions. Figure 4 shows derived plots for Los Angeles (Region 1), New York (Region 2), and Atlanta (Region 3). Sensitivity of the results to regionally dependent built-in fade margins, number of ground stations, thunderstorm-activity factor, and satellite orbital position are investigated later on. A coding gain of 10 dB is assumed.

4.1 Baseline results

Figure 5 shows predicted values of yearly fractional outage with resource sharing for a site in Region 2 versus the built-in rain margin,

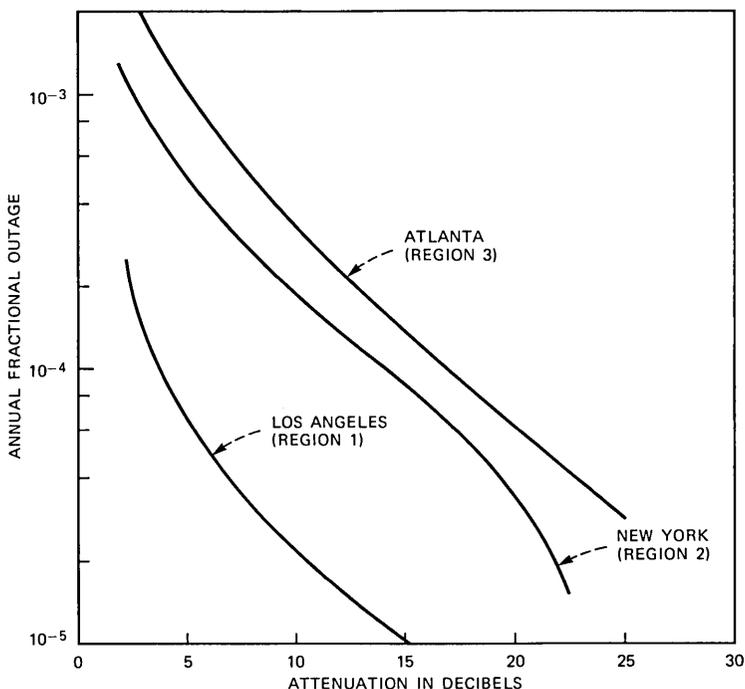


Fig. 4—Yearly average 12-GHz attenuation curves for Los Angeles (Region 1), New York (Region 2), and Atlanta (Region 3). The satellite is at 100°W longitude.

assumed to be common at all ground stations (i.e., all ground stations have the same size antenna). The network contains 100 ground stations, and the “geographical correlation factor,” β , is assumed equal to unity, that is, within a given region, the probability of attenuation during the thunderstorm period conditioned upon an attenuation event at any site is equal to the unconditional probability. The number of simultaneous fades which can be accommodated, K , varies between 1 and 3. Similar curves for $\beta = 2, 4$, and 6 appear in Figs. 6, 7, and 8, respectively.

Figures 5 through 8 show that resource sharing is of increasing utility as the outage objective becomes more stringent. For example, for $K = 2, \beta = 2$, the fade margin gain is 8.8 dB for a down-link outage of 0.01 percent, and increases to 9.8 dB for an outage of 0.005 percent. We see also that the ability to accommodate only a single fade ($K = 1$) restricts the fade margin resource-sharing gain at a down-link outage objective of 0.005 percent to 9.6 dB under the best of conditions ($\beta = 1$). For less favorable conditions ($\beta = 6$), the gain shrinks to 7.6 dB. Thus, the ability to accommodate only a single fade may severely influence the utility of resource sharing. However, if two simultaneous fades can be accommodated, then even for the unfavorable condition

$\beta = 4$, the fade margin gain is 9.4 dB; whereas for $\beta = 1$, the gain becomes 10 dB. If three simultaneous fades can be accommodated, then, even for the extremely unfavorable case $\beta = 6$, a fade margin gain of at least 9.5 dB can be achieved. For this last case, the outage objective of 0.005 percent can be achieved with a built-in margin of 8.8 dB, to compare against 18.4 dB required in the absence of resource sharing.

Figure 9 shows results obtained in Region 3. Again, a 12-GHz satellite located at 100 degrees West longitude and 100 identical ground stations, all with a common fade margin, are assumed. Curves in this illustration apply for $\beta = 1, 2, 4,$ and 6 and for $K = 2, 3,$ and 4 , and show that at an outage objective of 0.005 percent, the fade margin gain is 10 dB for all cases considered. This outage objective can be achieved with a built-in margin of 11.3 dB, to be compared against 21.3 dB required in the absence of resource sharing.

We note that this last result applies only when a built-in fade margin

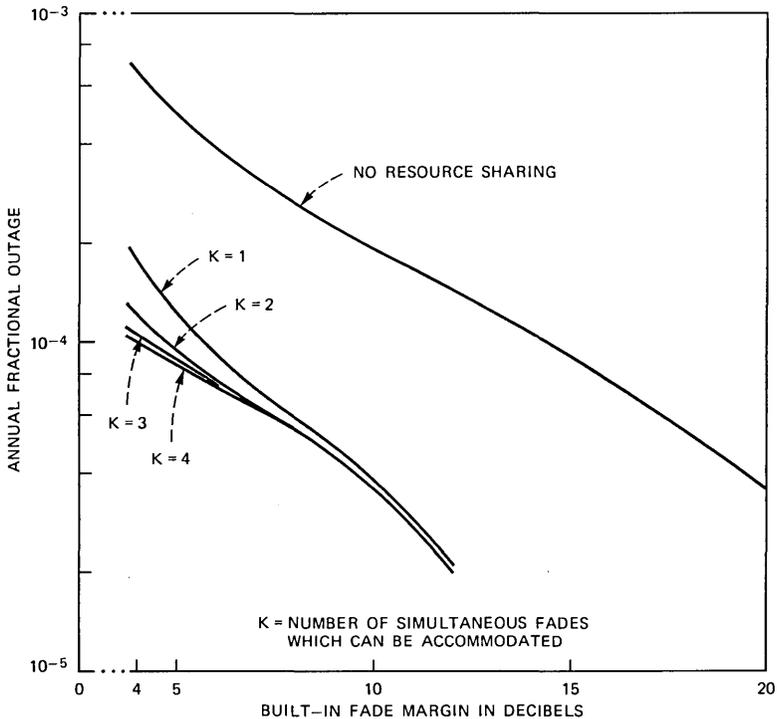


Fig. 5—Rain outage experienced with resource sharing at a site in Region 2 (northeast) versus the built-in fade margin. A maximum possible power saving of 10 dB is assumed. Time-of-day and seasonal thunderstorm activity factor $\alpha = 24$. The geographical factor $\beta = 1$. (Attenuation events at different sites within the thunderstorm activity period are independent.) The satellite is at 100°W longitude, and the network contains 100 identical ground stations.

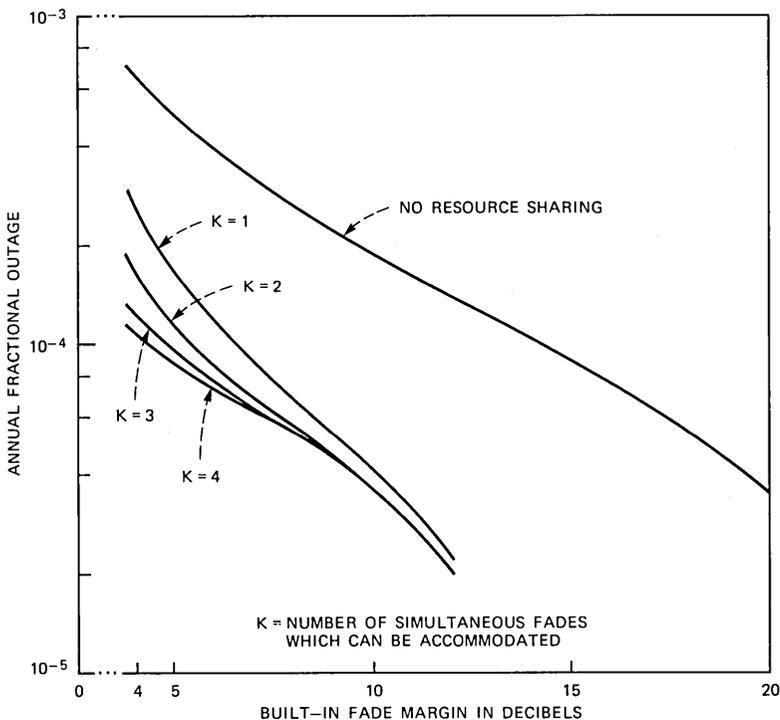


Fig. 6—Same as Fig. 5, except the geographical factor $\beta = 2$. (Given an attenuation event at one site within the thunderstorm activity period, it is twice as likely that a simultaneous attenuation event occurs at any other site within the same region.)

of 11.3 dB is provided at all sites in Regions 1, 2, and 3. From Figs. 5 to 8, we see that in Region 2, a built-in fade margin of 11.3 dB provides an outage lower than the required 0.005 percent. (No results are given for Region 1 because it was found that the objective of 0.005 percent was always readily achieved by virtue of the low level of attenuation prevalent in that region). Thus, we are motivated to consider a case wherein sites with different size antennas are deployed in different regions. The goal here is to optimize the system (i.e., provide the smallest antenna possible in each region) such that a down-link outage objective of precisely 0.005 percent is achieved everywhere.

4.2 Results for nonuniform rain margins

We consider a system wherein the antenna gain at each site in Region 3 exceeds the antenna gain for Region 2 by 2 dB, and the antenna gain at each site in Region 1 is less than the antenna gain in Region 2 by 2 dB. The results are shown in Fig. 10. Again, a 12-GHz satellite at 100 degrees West longitude and a network of 100 ground stations are assumed. Outage is plotted as a function of the built-in

margin in Region 2. Parameters are $K = 2$ or 3 , $\beta = 1$ or 4 . We see that if $K \geq 3$, the outage objective of 0.005 percent can be achieved in Region 3 if the Region 2 built-in margin is 9.3 dB, even for $\beta = 4$. At this level, the outage actually achieved in Region 2 is 0.0044 percent for $\beta = 4$, $K = 3$; an outage of 0.005 percent could have been achieved with a built-in margin of 8.8 dB (within 0.05 dB of that needed to achieve 0.005 percent in Region 3). Thus, this system is close to optimum, except that again, the outage achieved in Region 1 is far lower than the objective. This indicates that the antennas used in Region 1 are still far too large if outage is the only consideration.

For the case of a 500-MHz satellite transponder radiating a power level of 30 watts and a satellite aperture of 15 feet diameter, it has been estimated that a 5 meter diameter earth-station antenna would be needed in Region 2 to provide a rain margin of 16 dB.² Extrapolating to the parameters of Fig. 10 by holding the satellite power, bandwidth and aperture fixed, we find that an antenna diameter of 2.5 meters would be needed in Region 2 to provide a built-in margin of 9.3 dB; the Region 1 antenna diameter assumed in Fig. 10 is 2 dB smaller or 2 meters. Thus, it may be impractical to deploy a smaller antenna in

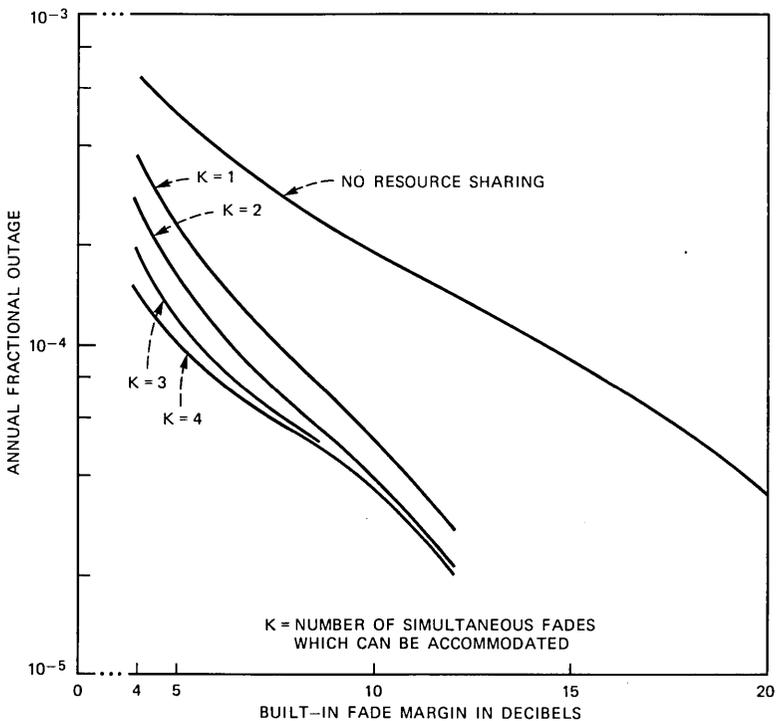


Fig. 7—Same as Fig. 5, except the geographical factor $\beta = 4$.

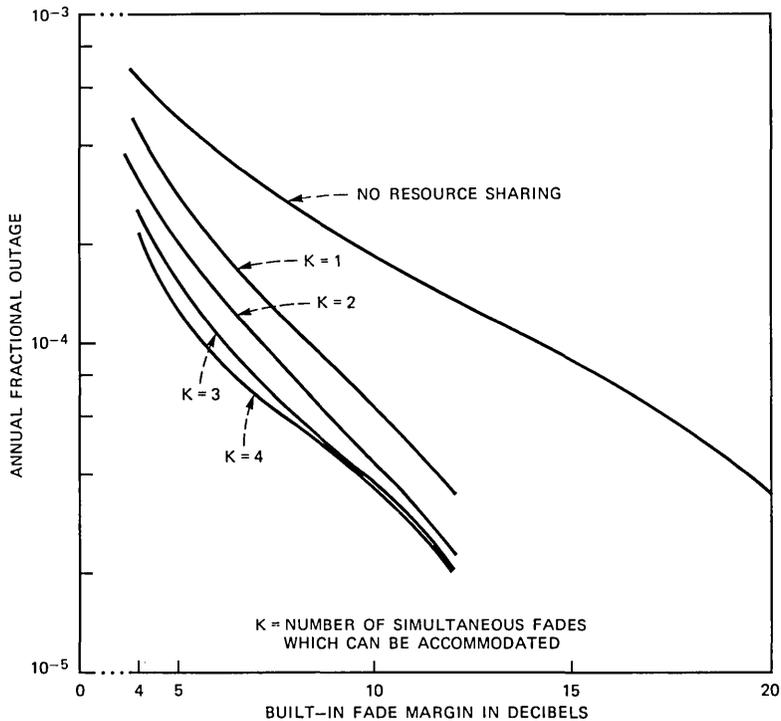


Fig. 8—Same as Fig. 5, except the geographical factor $\beta = 6$.

Region 1 without seriously interfering with other satellites in the geosynchronous orbit. Rather, it might be more advantageous to maintain a larger-than-needed antenna in Region 1 to enable a reduction in the required uplink transmitter power. In summary, the illustration portrayed in Fig. 10 is nearly optimum in a practical sense, and the outage objective of 0.005 percent is achieved everywhere by providing antennas which yield a built-in margin of 7.3 dB in Region 1, 9.3 dB in Region 2, and 11.3 dB in Region 3. The capability to accommodate $K = 3$ simultaneous fades is required for a network with 100 ground stations.

4.3 Sensitivity to number of ground stations

In Fig. 11, we consider a network containing 200 ground stations and plot the achievable outage in Region 2 versus the built-in margin, assumed to be the same at all 200 ground stations. Results for geographical factors β of 1, 2, and 4 and $K = 2$ to 5 appear. Because the number of ground stations has increased by a factor of two above previous cases considered, the number of simultaneous fades which must be accommodated is generally higher to achieve the same resource-sharing advantage. However, this factor is generally smaller

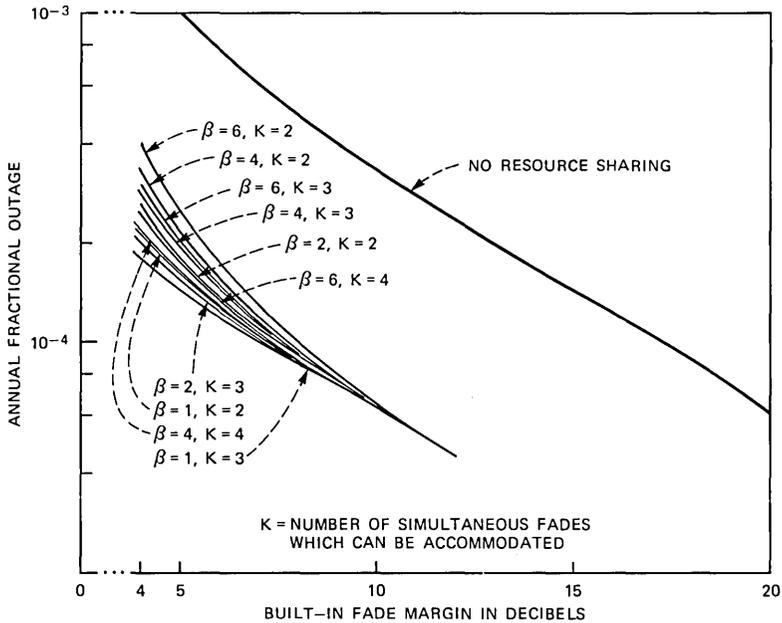


Fig. 9—Rain outage experienced with resource sharing at a site in Region 3 (southeast) versus built-in fade margin. The geographical factor $\beta = 1, 2, 4,$ and 6 . Other conditions are the same as for Fig. 5.

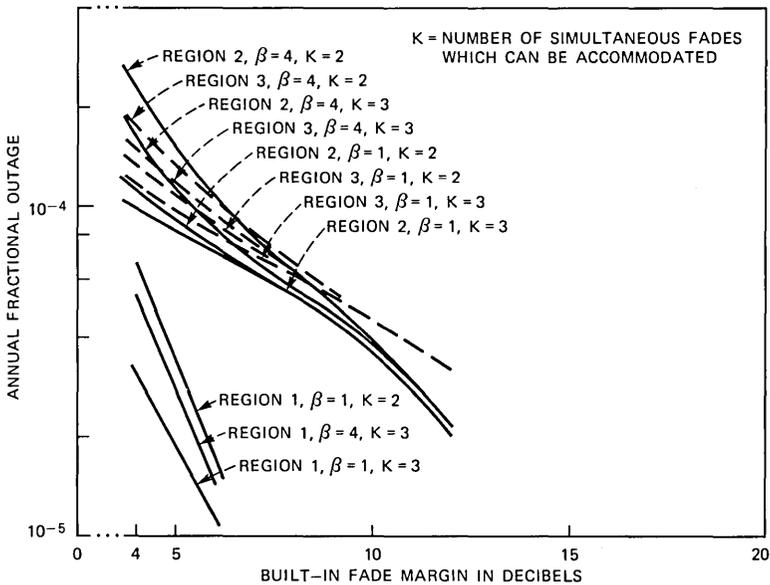


Fig. 10—Rain outage experienced with resource sharing for sites in Regions 1 (west), 2 (northeast), and 3 (southeast) versus built-in fade margin at sites in Region 2. Sites in Region 1 have 2 dB less built-in fade margin than sites in Region 2. Sites in Region 3 have 2 dB more fade margin than sites in Region 2. The geographical factor $\beta = 1$ and 4 . Other conditions are the same as for Fig. 5.

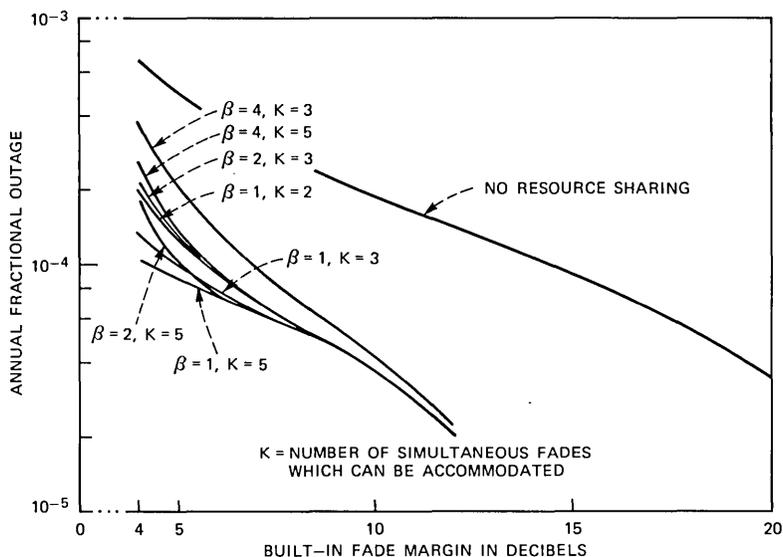


Fig. 11—Rain outage experienced with resource sharing at a site in Region 2 (north-east) versus built-in fade margin. The network contains 200 identical ground stations, and $\beta = 1, 2,$ and 4 . Other conditions are the same as for Fig. 5.

than two, implying that the resource-sharing overhead, expressed as the required number of reserved slots divided by the total number of accesses or interconnections, decreases as the number of users is increased.

Similarly, as shown in Fig. 12, the number of simultaneous fades which must be accommodated to achieve a given level of resource-sharing advantage decreases for a network containing 50, rather than 100, ground stations. Again, results for Region 2 are shown, and a 12-GHz satellite located at 100 degrees West longitude along with a common fade margin at all sites are assumed. The resource-sharing overhead is generally higher for the 50 station network than for the 100 ground station network.

Figure 13 is a composite plot for Region 2 showing the resource-sharing fade margin gain at an outage of 0.005 percent versus the required overhead for $\beta = 1, 2,$ and 4 and for a total number of 50, 100, and 200 ground stations. The four-for-one time slot expansion of Ref. 1, which allows use of a rate $\frac{1}{3}$ convolutional code to provide a coding gain of 10 dB along with extended synchronization preamble, is assumed. Letting the number of ground stations be represented by N , the overhead η , expressed as a percentage, is given by

$$\eta = \frac{3K}{N + 3K} \times 100. \quad (38)$$

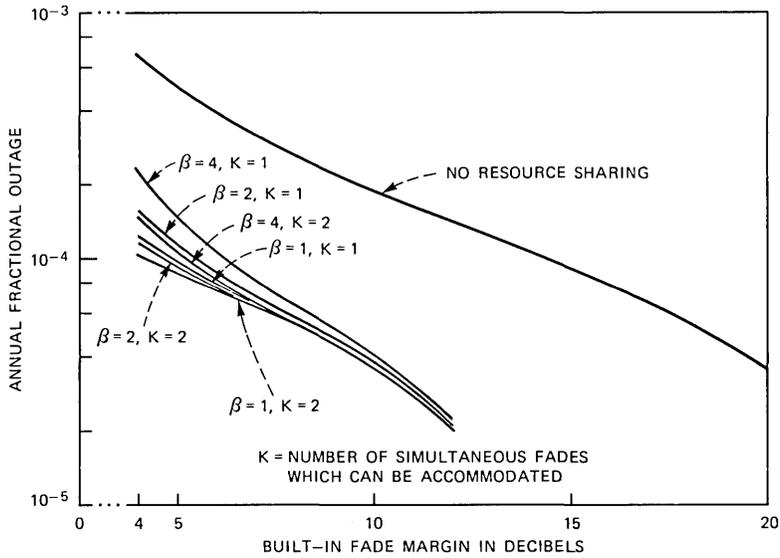


Fig. 12—Same as Fig. 11, except that the network contains 50 identical ground stations.

Again, a 12-GHz satellite located at 100 degrees West longitude is assumed. A similar composite plot can be derived for Region 3. We see that for $\beta \leq 4$, an overhead of 4.3 percent will ensure a fade margin gain greater than 9 dB if the network contains 200 ground stations. An

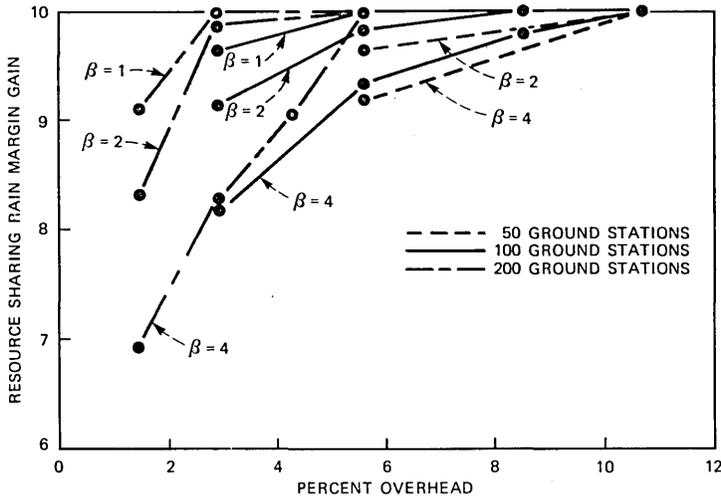


Fig. 13—Composite plot for a site in Region 2 (northeast) showing the resource-sharing fade margin gain at an outage of 0.005 percent versus the TDMA overhead for networks of 50, 100, and 200 ground stations. A maximum possible power saving of 10 dB is assumed. Time-of-day and seasonal thunderstorm activity factor $\alpha = 24$ and the geographical factor $\beta = 1, 2, \text{ and } 4$. The satellite is at 100°W longitude.

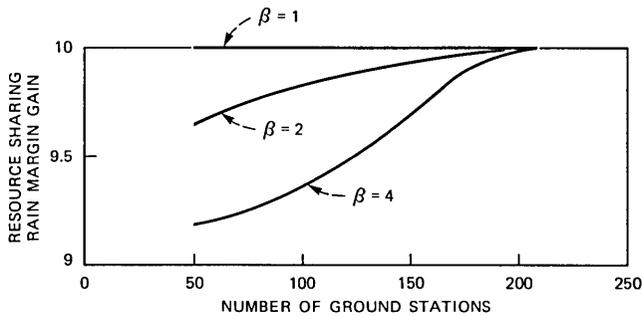


Fig. 14—Composite plot for a site in Region 2 (northeast) showing the resource sharing fade margin gain at an outage of 0.005 percent and a TDMA overhead of 5.66 percent versus the number of ground stations in the network. A maximum possible power saving of 10 dB is assumed. Time-of-day and seasonal thunderstorm activity factor $\alpha = 24$ and the geographical factor $\beta = 1, 2, \text{ and } 4$.

overhead of 5.66 percent is needed to achieve similar results for a network of 100 ground stations. Of course, the required overhead in all cases is reduced if smaller values of β apply.

Figure 14 plots a family of curves for a site in Region 2, showing the resource-sharing fade margin gain, for an outage of 0.005 percent, as a function of the number of ground stations in the network. The overhead is kept constant at 5.66 percent ($K = 1$ for 50 ground stations, $K = 2$ for 100 ground stations, and $K = 4$ for 200 ground stations), and the family parameter is β . Again, we see that resource sharing becomes more effective as the number of ground stations in the network increases.

4.4 Sensitivity to α

In Fig. 15, we investigate the effect of the thunderstorm-activity period factor, α , on the utility of resource sharing. For these curves, we assume that $\alpha = 12$, rather than the value of 24 used for all previous results, implying that the period of heavy attenuation is concentrated over an interval twice as large in time. Thus, at any site, the probability that attenuation exceeds some value A , conditioned on the thunderstorm period, is half its former value. Again, a 12-GHz satellite located at 100 degrees West longitude is assumed, and results apply at a ground station in Region 2. A total of 100 ground stations is assumed, and all ground stations have a common built-in fade margin. We see that, at an outage of 0.005 percent, the ability to accommodate two simultaneous fades provides a fade margin advantage in excess of 9.8 dB for β as high as four. Thus, under the less pessimistic and more realistic assumption that $\alpha = 12$ (4 out of 12 months/year, 6 out of 24 hours/day) rather than 24, we see that the resource-sharing advantage is very close to the maximum possible for an overhead of about six

percent. We note, however, that the gain over $\alpha = 24$ is only 0.4 dB (see Fig. 7).

4.5 Sensitivity to satellite longitude

Finally, in Fig. 16, we plot the yearly average attenuation statistics (based upon Lin's model) for sites in Los Angeles (Region 1), New York (Region 2), and Atlanta (Region 3) for a 12-GHz satellite located at 130 degrees West longitude rather than 100 degrees West as assumed before. Because of the greater slant range to the satellite, the attenuations experienced at sites in Regions 2 and 3 are greater; again, Region 1 experiences very little attenuation. We use these statistics to derive the curves of Fig. 17 which show the fractional outage versus the built-in margin for a site in Region 2 and for various values of K and β . A total of 100 ground stations of common fade margin is assumed, and $\alpha = 24$. Figure 17 shows that for $\beta \leq 4$, capability for accommodating only one fade is sufficient to provide 10 dB of extra protection via resource sharing at an outage objective of 0.005 percent. Because the yearly averaged attenuation statistics are nearly the same in Regions 2 and 3 for this satellite longitude, it is safe to assume that Fig. 17 is reasonably representative of a site in Region 3 as well. Also, because of this similarity between Regions 2 and 3, it appears that resource sharing is more effective for a satellite location at 130 degrees

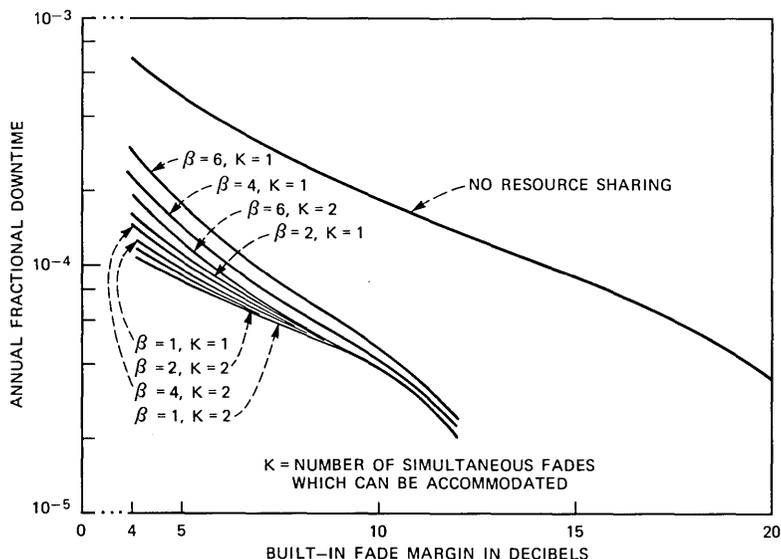


Fig. 15—Rain outage experienced with resource sharing at a site in Region 2 (north-east) versus the built-in fade margin. Time-of-day and seasonal thunderstorm activity factor $\alpha = 12$, and the geographical factor $\beta = 1, 2, 4, \text{ and } 6$. Other conditions are the same as for Fig. 5.

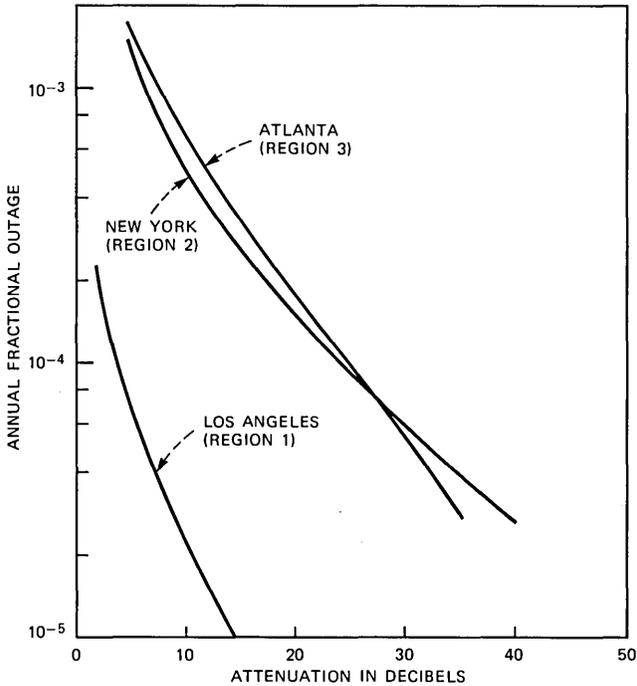


Fig. 16—Yearly average 12-GHz attenuation curves for Los Angeles (Region 1), New York (Region 2), and Atlanta (Region 3). The satellite is at 130°W longitude.

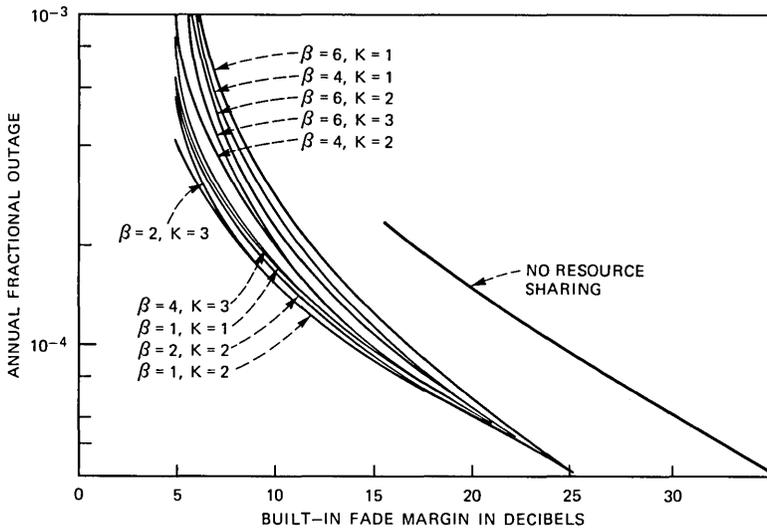


Fig. 17—Rain outage experienced with resource sharing at a site in Region 2 (north-east) versus the built-in fade margin. The satellite is at 130°W longitude, and the time-of-day and seasonal thunderstorm activity factor $\alpha = 24$. Other conditions are the same as for Fig. 5.

West, compared against 100 degrees West, although the required built-in margin is much greater. This is because, when we position the satellite at 100 degrees West, the higher attenuation experienced in Region 3 relative to Region 2 presents more competition for the reserved time slots, which is disadvantageous for a site in Region 2.

V. SUMMARY AND CONCLUSION

In this paper, we have attempted to evaluate the effects of the size of the reserved pool, the number of ground stations, correlation among attenuation events, rain outage objective, and satellite position on the utility of resource sharing to provide 10 dB of extra down-link margin against rain fading. A single wideband satellite transponder such as might be associated with a scanning spot-beam system was assumed, but the joint rain attenuation and traffic modeling and the analytical approach can be extended to study multiple spot beam frequency reuse concepts. In general, it was found that for the cases examined, a pooled resource overhead equal to six percent of the available TDMA time slots is adequate to ensure a fade margin gain in excess of 9 dB for an outage objective of 0.005 percent if there are more than 50 ground stations in the network. The difference between the actual gain realized and the maximum gain of 10 dB possible with the coding approach employed arises from the effects of many ground stations in competition for a small number of shared resources.

The traffic model employed for this study assumes that all ground stations carry similar traffic cross sections. Although the analytical approach can be modified to reflect the effects of users with different traffic cross sections, such an approach might become numerically unwieldy if the number of user classes assumed becomes too large. However, a simple overbound on the required overhead can readily be obtained if for a given number of users, we apply the results obtained by assuming a smaller number of users. For example, suppose the network contains 200 ground stations, and the amount of traffic carried by the ground stations are within a factor of four of each other. Then, the required overhead calculated for a network of 50 ground stations is an overbound of that needed for the 200 ground station network.

The predicted results depend, of course, on the rain model assumed; modeling is necessary because statistics of multiple fade events for a large number of geographically dispersed sites are unavailable. Fortunately, in a TDMA system designed on the resource-sharing concept, it is a simple matter to alter the size of the reserved pool in response to real operating experience.

Three final observations will be offered. First, it appears from these results that resource sharing is most effective in the southeastern portion of the United States because the traffic demand from that

region is smaller than that presented by other regions. This is indeed fortunate because it is precisely in that region, where attenuation is high, that resource sharing can offer the most saving. Second, if the satellite receives a western orbital-slot assignment, resource-sharing again assumes an important role because of the large rain margins (and, therefore, the large satellite radiated power or large ground stations) required to provide a suitably low rain outage. Finally, we note that resource sharing can provide high-system reliability (low-rain outage) while maintaining a sufficiently low-satellite effective radiated power to satisfy interference constraints imposed by the presence of other users of the geosynchronous orbit. It is perhaps in this last regard that resource sharing will prove most valuable, by enabling coexistence of spot-beam systems with existing wide-area coverage systems.

REFERENCES

1. A. S. Acampora, "A Shared Resource TDMA Approach to Increase the Rain Margin of 12/14 GHz-Satellite Systems," *B.S.T.J.*, 58, No. 9 (November 1979), pp. 2097-111.
2. D. O. Reudink and Y. S. Yeh, "A Scanning Spot Beam Satellite System," *B.S.T.J.*, 56, No. 8 (October 1977), pp. 1549-60.
3. A. S. Acampora and B. R. Davis, "Efficient Utilization of Satellite Transponders via Time-Division Multibeam Scanning," *B.S.T.J.*, 57, No. 8 (October 1978), pp. 2901-14.
4. A. S. Acampora, C. Dragone, and D. O. Reudink, "A Satellite System with Limited-Scan Spot Beams," *IEEE Trans. Commun., Com-27*, No. 10 (October 1979), pp. 1406-15.
5. D. O. Reudink and Y. S. Yeh, "The Organization and Synchronization of a Switched Spot-Beam System," *Fourth Int. Conf. Digital Satellite Commun.*, Montreal, October 1978.
6. D. O. Reudink, A. S. Acampora, and Y. S. Yeh, "Methods for Achieving High Capacity Universal Service Satellites," *Nat. Telecom. Conf.*, Birmingham, December 1978.
7. S. H. Lin, H. J. Bergman, and M. V. Pursley, "Rain Attenuation on Earth-Satellite Paths—Summary of 10-Year Experiments and Studies," *B.S.T.J.*, 59, No. 2 (February 1980), pp. 183-228.
8. S. H. Lin, "Nationwide Long-Term Rain-Rate Statistics and Empirical Calculation of 11-GHz Microwave Attenuation," *B.S.T.J.*, 56, No. 9 (November 1977), pp. 1581-1604.
9. S. H. Lin, "Empirical Rain Attenuation Model for Earth-Satellite Paths," *IEEE Trans. Comm., COM-27*, No. 5 (May 1979).

Modeling Multipath Fading Responses Using Multitone Probing Signals and Polynomial Approximation

By L. J. GREENSTEIN and B. A. CZEKAJ

(Manuscript received August 7, 1980)

We show in quite a general way that highly accurate modeling of multipath fading responses is possible using low-order complex polynomials. This applies to all terrestrial radio systems in the channelized common carrier bands below 15 GHz, where channel widths are 40 MHz or less. The context of the study is a new multipath experiment being conducted in New Jersey over a 23-mile path at 11 GHz. The transmitted signal consists of up to nine tones in a 40-MHz bandwidth. These tones are coherently processed, sampled, and digitized in the receiver and recorded, during fading events, for later off-line reductions. Simple routines can be used to determine polynomial coefficients from these recorded data. This paper describes the signal processing and data reduction methods and analyzes them to assess the accuracy of polynomial fitting. The analysis uses a mean-square error measure and assumes a representative form for the underlying response function. Our results predict that the vast majority of multipath fading responses can be accurately approximated over bandwidths of 40 (62) MHz using first- (second-) order complex polynomials.

I. INTRODUCTION

Multipath fading (hereafter abbreviated MPF) on terrestrial microwave paths can be a major cause of outage in digital radio systems.¹⁻⁴ Numerous efforts have been aimed at understanding, analyzing, and correcting this source of disruption, and some have led to new statistical models for MPF responses.⁵⁻⁹

The particular model that inspires the present work approximates the MPF response by a low-order complex polynomial in frequency.⁸ For a particular 26-mile path in Georgia, it was shown that a first-order polynomial suffices to characterize the fading response in a 25-

MHz band centered near 6 GHz. The joint probability distribution for the polynomial coefficients was derived for that path, thus permitting a complete statistical description of the MPF response.

Now another experiment is being instrumented, this time for a 23-mile path in New Jersey operating in the 11-GHz band. The aim of the new experiment is to add to, and in several respects improve upon, the data base used to quantify the earlier polynomial model. The improvements include higher measurement signal to noise ratios (SNRs), higher sampling rates (20 measurements per second rather than 5), coherent processing to obtain phase information (previously absent), and a wider measurement bandwidth (40 MHz rather than 25 MHz).

Given the highly variable nature of multipath fading, such improved measurements for a new path in a different frequency band and locale should add importantly to our knowledge of this phenomenon.

The basic design of the experiment can be simply stated: As many as nine coherently-phased tones within a 40-MHz bandwidth are transmitted from Murray Hill and coherently demodulated in a receiver at Crawford Hill; the demodulated tones are sampled, digitized, and screened by a desktop computer/controller; and the digitized data, if deemed interesting, are recorded on magnetic tape for later off-line processing.

The recorded data will be in a form that facilitates polynomial approximation using simple, efficient computer routines. The data will be quite general in form, however, i.e., amenable to modeling via any mathematical approximation considered promising.

The present study evaluates the accuracy of polynomial approximation, relating it to the experiment parameters and to the methods of signal processing and data reduction. Section II describes the signal processing in the transmitter and receiver, and derives signal and noise relationships used in the subsequent error analyses. Section III describes the methods of polynomial fitting to be considered, and defines the mean-square error measures that will be used to evaluate them.

Section IV analyzes the errors in the polynomial fitting caused by measurement noise, and Section V analyzes the errors caused by finite sampling of the frequency response. In general, the errors increase with the bandwidth over which the fitting is done. In analyzing the errors caused by finite sampling, we assume a general form for the MPF response function that has been applied successfully in other data-fitting studies,⁵ and assume either worst-case or typical values for the function parameters.

The mean-square error calculations permit predictions of the maximum bandwidths for which polynomial fitting is valid. Section VI summarizes the results computed, under rather stringent mean-square error requirements, for polynomial orders of one, two, four, six, and eight.

II. SIGNAL PROCESSING ANALYSIS FOR THE MPF EXPERIMENT

2.1 Propagation path and radio channels

Multipath fading responses are to be measured on a 23-mile path between Murray Hill and Crawford Hill, similar to the one used by Crawford and Jakes in their earlier experiments.¹⁰ The transmitting antenna at Murray Hill is 655 feet above sea level, and the receiving antenna atop Crawford Hill is 425 feet above sea level. An experimental license has been obtained to operate over this path in three 40-MHz channels within the 11-GHz common carrier band. These channels are centered at 11.465, 11.545, and 11.625 GHz. The initial measurements will be in the band centered at 11.545 GHz.

We describe here the signal processing relationships that underlie the experiment design. The details of circuitry, components, and equipment will be reported separately by those who have developed the MPF measurement system.

2.2 Transmitted signal

The transmitted signal is created by the two-stage upconversion of a baseband signal having the form

$$b(t) = d_0 + \sum_{n=1}^{N/2} d_n \cos(n\Delta\omega t + \theta_n), \quad (1)$$

where N is even and the other parameters will be discussed. The upconversion places the signal in an RF channel centered at radian frequency $\omega_c = 2\pi f_c$ ($f_c = 11.545$ GHz). Hence, the transmitted signal is

$$V_T(t) = \sum_{n=-N/2}^{N/2} \sqrt{2p_n} \cos(\omega_c t + n\Delta\omega t + \theta_n), \quad (2)$$

where p_n is the power of the n th transmitted tone and a total of $N + 1$ tones are transmitted. From (1), we see that p_0 is proportional to d_0^2 and that p_n ($n \neq 0$) bears the same proportionality to $d_n^2/4$.

The variation of p_n with n is clearly symmetrical about $n = 0$ because it derives from amplitude settings of the baseband tones. Nonuniform variations of p_n can easily be compensated for via baseband adjustments in the receiver. In Section IV we consider nonuniform variations for which receiver noise effects are minimized.

The frequency spacing between transmitted tones, Δf , may be 5, 10, or 20 MHz. Since the transmission is confined to a channel of 40-MHz width, and occupies a bandwidth $N\Delta f$, we have the constraints $(N + 1) \leq 9$ tones when $\Delta f = 5$ MHz; $(N + 1) \leq 5$ tones when $\Delta f = 10$ MHz; and $(N + 1) \leq 3$ tones when $\Delta f = 20$ MHz. We will consider the four particular combinations $N = 2$, $\Delta f = 20$ MHz; $N = 4$, $\Delta f = 10$ MHz; $N = 6$, $\Delta f = 5$ MHz; and $N = 8$, $\Delta f = 5$ MHz.

Finally, we mention that the $N/2$ baseband tones are derived from

a common 5-MHz reference, and so can be relatively phased in any manner desired. For purposes of analysis, we will assume all θ_n 's to be zero here; since any phasings in the transmitter are easily undone in the receiver, no generality is lost. One criterion for choosing the actual θ_n 's is minimization of the peak factor of the RF signal (2). The baseband phase adjustments that accomplish this have been derived for $N = 2, 4, 6,$ and 8 .¹¹ We will use the resulting minimized peak factors in making noise calculations later.

2.3 Response of the propagation medium

We denote the complex signal gain of the propagation medium by

$$F(\omega) = \frac{\text{Complex gain at radian frequency } \omega_c + \omega}{\text{Gain magnitude } (g_0) \text{ during nonfading}}. \quad (3)$$

The quantity g_0 can be computed from familiar radio path equations. Note that ω is measured from the center of the channel. During nonfading periods, $|F(\omega)| = 1$ throughout the channel bandwidth; during multipath fading, $F(\omega)$ varies with ω in a randomly time-varying manner.

The function $F(\omega)$ contains two phase factors of no interest to us. One is $\exp(j\phi_0)$, where ϕ_0 is the phase shift through the medium at $\omega = 0$; the other is $\exp(-j\omega t_p)$, where t_p is the nominal propagation time along the path. (For a 25-mile path, $t_p \doteq 0.13$ ms.) The investigation of multipath fading can be simplified, with no loss of information, by removing these two factors. Thus the function of interest to us is

$$H(\omega) = F(\omega) \exp[j(-\phi_0 + \omega t_p)]. \quad (4)$$

The aim of our modeling effort is to find suitable functions for approximating $H(\omega)$, and to statistically characterize the parameters of those functions.

We will see in Section 2.5 that the response function actually sampled by the measurement system is

$$G(\omega) = F(\omega) \exp\left[-j\left(\psi + \omega \frac{\theta}{\Delta\omega}\right)\right], \quad (5)$$

where ψ and θ are the (possibly) random or unknown phases of frequency references in the receiver. To obtain samples of the desired function [$H(\omega)$] from samples of the measured function [$G(\omega)$] will therefore require performing the operation

$$H(\omega) = G(\omega) \exp\left[\underbrace{j(\psi - \phi_0)}_{\Delta\Phi} + j\omega \underbrace{\left(t_p + \frac{\theta}{\Delta\omega}\right)}_{\Delta T}\right] \quad (6)$$

at each of the sampling frequencies $[\omega = 0, \pm\Delta\omega, \dots, \pm(N/2)\Delta\omega]$. We will show later how to accomplish this in the data processing.

2.4 Decompositions of $F(\omega)$, $G(\omega)$, and $H(\omega)$

We demonstrate here a useful decomposition for complex response functions such as $F(\omega)$, $G(\omega)$, and $H(\omega)$. We will treat only $H(\omega)$, noting that the same mathematics and notation apply to $F(\omega)$ and $G(\omega)$.

Since ω is measured from an arbitrary microwave frequency ($2\pi f_c$), there is no physical reason to assume complex conjugate symmetry for $H(\omega)$. In its most general form, $H(\omega)$ can be expressed as

$$H(\omega) = H_a(\omega) + jH_b(\omega), \quad (7)$$

where $H_a(\omega)$ and $H_b(\omega)$ are each functions having complex conjugate symmetry. Accordingly, we can write

$$H_a(\omega) = \underbrace{H_{ar}(\omega)}_{\substack{\text{Real,} \\ \text{Even}}} + j\underbrace{H_{ai}(\omega)}_{\substack{\text{Imaginary,} \\ \text{Odd}}}, \quad H_b(\omega) = \underbrace{H_{br}(\omega)}_{\substack{\text{Real,} \\ \text{Even}}} + j\underbrace{H_{bi}(\omega)}_{\substack{\text{Imaginary,} \\ \text{Odd}}}. \quad (8)$$

By transmitting and coherently receiving $N + 1$ tones spaced by $\Delta\omega$, one could in theory obtain measurements of the two even functions at $\omega = 0, \Delta\omega, \dots, (N/2)\Delta\omega$; and of the two odd functions at $\omega = \Delta\omega, \dots, (N/2)\Delta\omega$. [The total number of samples, $2(N + 1)$, is consistent with measuring the amplitudes and phases of the $N + 1$ received tones.] In reality, the receiver obtains these samples for the corresponding set of G functions, which differ from the H functions if ψ and θ are not both 0. Obtaining H samples from G samples is discussed in Section 2.6.

Another departure of the receiver outputs from the desired samples is the presence of measurement noise. We will defer the introduction of noise until Section 2.7.

2.5 Signal processing in the receiver

The receiver input at RF is

$$V_R(t) = g_0 \sum_{n=-N/2}^{N/2} \sqrt{2p_n} |F_n| \cos(\omega_c t + n\Delta\omega t + \phi_n), \quad (9)$$

where $|F_n|$ and ϕ_n are the magnitude and phase of $F(n\Delta\omega)$; g_0 is the normal (nonfading) path gain, (3); and we have used (2) with all θ_n 's assumed to be zero.

The signal goes through a two-stage down conversion which amounts to quadrature demodulation. That is, two baseband outputs are obtained which correspond to mixing $V_R(t)$ with $2 \cos(\omega_c t + \psi)$ and with $-2 \sin(\omega_c t + \psi)$. A nonzero value of ψ signifies that the RF and IF

references in the receiver are not in phase synchronism with those in the transmitter.

Each of the baseband signals consists of a dc component plus sinusoids at $\omega = \Delta\omega, \dots, (N/2)\Delta\omega$. The n th sinusoid in each of these signals goes through quadrature demodulation, via the local references $\cos[n(\Delta\omega t + \theta)]$ and $-\sin[n(\Delta\omega t + \theta)]$, to produce two more dc outputs. These references are all derived from a 5-MHz source in the receiver, and nonzero θ signifies that this source is not in phase synchronism with the one in the transmitter.

Using (5) and ordinary trigonometric identities, the following statements can be proven:

(i) The dc outputs produced by demodulation via

$$\begin{pmatrix} 2 \cos(\omega_c t + \psi) \\ -2 \sin(\omega_c t + \psi) \end{pmatrix}$$

are

$$\begin{pmatrix} I_0 \\ Q_0 \end{pmatrix} \triangleq \begin{pmatrix} \sqrt{2p_0} g_0 G_{ar}(0) \\ \sqrt{2p_0} g_0 G_{br}(0) \end{pmatrix}. \quad (10)$$

(ii) The dc outputs produced by demodulation via

$$\begin{pmatrix} 2 \cos(\omega_c t + \psi) \\ -2 \sin(\omega_c t + \psi) \end{pmatrix}$$

and $\cos[n(\Delta\omega t + \theta)]$ are

$$\begin{pmatrix} (I - I)_n \\ (Q - I)_n \end{pmatrix} \triangleq \begin{pmatrix} \sqrt{2p_n} g_0 G_{ar}(n\Delta\omega) \\ \sqrt{2p_n} g_0 G_{br}(n\Delta\omega) \end{pmatrix}, \quad n = 1, N/2. \quad (11)$$

(iii) The dc outputs produced by demodulation via

$$\begin{pmatrix} 2 \cos(\omega_c t + \psi) \\ -2 \sin(\omega_c t + \psi) \end{pmatrix}$$

and $-\sin[n(\Delta\omega t + \theta)]$ are

$$\begin{pmatrix} (I - Q)_n \\ (Q - Q)_n \end{pmatrix} \triangleq \begin{pmatrix} \sqrt{2p_n} g_0 G_{ai}(n\Delta\omega) \\ \sqrt{2p_n} g_0 G_{bi}(n\Delta\omega) \end{pmatrix}, \quad n = 1, N/2. \quad (12)$$

Thus, the dc outputs of the receiver correspond to evenly-spaced frequency samples of the four G functions $G_{ar}(\omega)$, $G_{ai}(\omega)$, $G_{br}(\omega)$, and $G_{bi}(\omega)$. Each of the $2(N + 1)$ dc components is low-pass filtered, with a noise bandwidth of 100 Hz; time-sampled 20 times per second and quantized with 14-bit precision by a Datel System 256 acquisition system,* and passed to a HP 9845A computer for real-time screening.†

* Datel Corporation.

† Hewlett-Packard.

Any set of samples that seems interesting, or is part of a sequence that seems interesting, is recorded on magnetic tape for subsequent off-line processing.

2.6 Relating the G and H functions

By using (6), the H functions ($H_{ar}(\omega)$, $H_{ai}(\omega)$, etc.) can be easily obtained from the corresponding G functions once $\Delta\Phi$ and ΔT are specified. Defining a new function $G(\omega; \Delta\Phi) \triangleq G(\omega) \exp(j\Delta\Phi)$, we first perform the matrix operation

$$\begin{bmatrix} G_{ar}(\omega; \Delta\Phi) \\ G_{br}(\omega; \Delta\Phi) \end{bmatrix} = \begin{bmatrix} \cos \Delta\Phi & -\sin \Delta\Phi \\ \sin \Delta\Phi & \cos \Delta\Phi \end{bmatrix} \begin{bmatrix} G_{ar}(\omega) \\ G_{br}(\omega) \end{bmatrix}. \quad (13)$$

An identical equation relates $G_{ai}(\omega; \Delta\Phi)$ and $G_{bi}(\omega; \Delta\Phi)$ to $G_{ai}(\omega)$ and $G_{bi}(\omega)$. The H functions are then obtained from these new G functions, for specified ΔT , as follows:

$$\begin{bmatrix} H_{ar}(\omega) \\ H_{ai}(\omega) \end{bmatrix} = \begin{bmatrix} \cos \omega\Delta T & -\sin \omega\Delta T \\ \sin \omega\Delta T & \cos \omega\Delta T \end{bmatrix} \begin{bmatrix} G_{ar}(\omega; \Delta\Phi) \\ G_{ai}(\omega; \Delta\Phi) \end{bmatrix}. \quad (14)$$

An identical equation relates $H_{br}(\omega)$ and $H_{bi}(\omega)$ to $G_{br}(\omega; \Delta\Phi)$ and $G_{bi}(\omega; \Delta\Phi)$.

The operation indicated by (13) leads to the result $H_{br}(0) = 0$, i.e., the phase response of the function to be analyzed is forced to zero at $\omega = 0$. [To see this, combine (6), (12) and (13) for $\omega = 0$.] This is a welcome simplification in the data and entails no loss of useful information. Fortunately, $\sin \Delta\Phi$ and $\cos \Delta\Phi$ are readily obtained from the measured G samples, for

$$\sin \Delta\Phi = -\frac{G_{br}(0)}{\sqrt{G_{ar}^2(0) + G_{br}^2(0)}}, \quad \cos \Delta\Phi = \frac{G_{ar}(0)}{\sqrt{G_{ar}^2(0) + G_{br}^2(0)}}. \quad (15)$$

To see this, use (5) at $\omega = 0$ and recall that $\Delta\Phi = \psi - \phi_0$.]

By way of contrast, the value of ΔT to use in (14) is not so readily specified or determined. Yet, to get the full benefit of polynomial modeling (i.e., accurate fitting using low-order functions), ΔT must be carefully chosen.* We have arrived at a criterion for choosing ΔT based upon the following data reduction procedure: For a given ΔT , (14) is applied and the resulting H samples are fitted by a finite-order complex polynomial in $j\omega$. We consider that value of ΔT to be optimal for which the polynomial fitting is best, in some least-squares sense

* To see why, consider the example $G(\omega) = \exp(-j\omega t_0)$, where $t_0 \approx 1/B$, and B is the bandwidth over which $H(\omega) = G(\omega) \exp(j\omega\Delta T)$ is to be fitted. If $\Delta T = 0$, $H(\omega) = \cos \omega t_0 - j \sin \omega t_0$, and high-order polynomials are needed for accurate fitting over the bandwidth. If $\Delta T = t_0$, however, $H(\omega) = 1 + j\omega$, and low-order polynomials are quite sufficient.

defined later. In Section 5.4, we will identify a data-derived measure that accurately predicts the optimal ΔT .

2.7 Measurement noise

We have shown that the dc receiver outputs are proportional to frequency samples of the function $G(\omega)$, and that the unwanted phase factors that distinguish $G(\omega)$ from $H(\omega)$ can be removed in the data processing. Not so readily removed are the noises associated with the digitized outputs. These consist of both additive Gaussian noise from the input and components of the receiver and quantizing noise from the 14-bit analog-to-digital conversions.

Receiver noise produces an additive random component for each dc output defined by (10) to (12). These $2(N + 1)$ noises are zero-mean and mutually independent. All have the same variance except those associated with (10), i.e., $n = 0$, for which the variance is 3 dB higher. These findings follow from the receiver processing described in Section 2.5.

We shall now assume that each of the dc outputs in (10) to (12) is adjusted by a factor $1/(g_0 \sqrt{2p_n})$, $n = 0, N/2$, before being digitized. Accordingly, the variance of the Gaussian noise associated with a given output sample is

$$\sigma_G^2 = \begin{cases} kTbN_F/p_0g_0^2, & n = 0, \\ kTbN_F/2p_n g_0^2, & n = 1, N/2. \end{cases} \quad (16)$$

Table I defines the quantities in (16) and gives values for each. Using those data and assuming uniform tone powers, we obtain the following

Table I—System parameters used in noise analysis

Parameter Definition and Assumed Value(s)	
g_0^2	Power gain between transmitter output and first receiver amplifier stage; includes clear-air path loss and circuit and waveguide losses: $10 \log g_0^2 = -74$ dB
kT	Thermal noise density at receiver input: $10 \log kT = -174$ dBm/Hz
b	Noise bandwidth of receiver processing for each tone: $b = 100$ Hz
N_F	Noise figure of first receiver amplifier stage: $10 \log N_F = 2$ dB
P_p	Peak transmitter power, constrained to meet out-of-band emission requirements: $10 \log P_p = 14$ dBm
P_F	Transmitter peak factor, minimized via phase adjustments of the baseband tones before upconversion. For uniform tone powers, $10 \log P_F = \begin{cases} 7.8 \text{ dB}, & N = 2, \\ 6.3 \text{ dB}, & N = 4, \\ 6.3 \text{ dB}, & N = 6, \\ 6.6 \text{ dB}, & N = 8 \end{cases}$
\bar{P}	Average transmitter power: $\bar{P} = P_p/P_F$. For $N = 8$, $10 \log \bar{P} = 7.4$ dBm
p_n	Average power in transmitted tone at frequency $\pm n\Delta\omega$ from center. For uniform tone powers and $N = 8$, $10 \log p_n = -1.8$ dBm

result: σ_G^2 in dB, for $n \neq 0$, lies in the range $-79 \text{ dB} \pm 2 \text{ dB}$, where the precise value depends on N ; for $n = 0$, σ_G^2 is 3 dB higher.

Assuming the dc outputs are amplitude-adjusted as indicated, the input to every 14-bit A/D conversion is precisely a sample of a G -function. During normal propagation, the G samples lie within ± 1 [see (3) and (5)]. To provide some room for excess gain, we assume quantizer amplitude limits set at ± 1.60 (4-dB margin). As a result, the quantizing error for each digitized sample can be characterized as an additive noise uniformly distributed on $[-\Delta/2, \Delta/2]$, where $\Delta = 2 \times 1.6 \times 2^{-14} = 1.95 \times 10^{-4}$. The quantizing noise variance is then

$$\sigma_Q^2 = \Delta^2/12 = 3.18 \times 10^{-9} (-85 \text{ dB}), \quad \text{every sample.} \quad (17)$$

Comparing this with σ_G^2 above, we find justification for ignoring quantization effects. Alternately, they can be accounted for using an approximate correction factor given in Section IV.

To simplify matters further, we introduce the notation

$$H_{ar,n} \triangleq H_{ar}(n\Delta\omega), \quad H_{br,n} \triangleq H_{br}(n\Delta\omega), \quad \text{etc.} \quad (18)$$

These are the quantities produced by (13) and (14) for any specified combination of $\Delta\Phi$ and ΔT . We account for the noisiness of the H samples via the notation

$$\hat{H}_{ar,n} \triangleq H_{ar,n} + \zeta_{ar,n}, \quad \hat{H}_{br,n} \triangleq H_{br,n} + \zeta_{br,n}, \quad \text{etc.,} \quad (19)$$

where $\zeta_{ar,n}$, $\zeta_{br,n}$, etc., are the random noise samples. Since the ζ s are produced by phase rotations of the noises associated with the G samples, they are all Gaussian, zero-mean, and mutually independent, just like the original noises. Moreover, their variances are identical to σ_G^2 , as given by (16). We thus have an accurate, simple description for the noisiness of the data to be processed.

III. POLYNOMIAL FITTING AND ERROR MEASURES

3.1 Fitting polynomials to the \hat{H} samples

The implicit assumption of the polynomial fitting approach is that, over some finite bandwidth B centered on $f = 0$, the response function $H(\omega)$ can be accurately approximated by a low-order complex polynomial, i.e.,

$$H(\omega) \cong \tilde{H}(\omega) \triangleq \sum_{k=0}^M (A_k + jB_k)(j\omega)^k, \quad |\omega| \leq \pi B. \quad (20)$$

Using (7) and (8), we can break this representation down as follows:

$$H_a(\omega) \cong \sum_{k=0}^M A_k(j\omega)^k = \underbrace{\sum_{k=0}^M A_k(j\omega)^k}_{\text{Fit to } H_{ar}(\omega)} + \underbrace{\sum_{k=1}^M A_k(j\omega)^k}_{\text{Fit to } jH_{ai}(\omega)}, \quad (21)$$

$$H_b(\omega) \cong \sum_{k=0}^M B_k(j\omega)^k = \underbrace{\sum_{\substack{k=0 \\ \text{Even}}}^M B_k(j\omega)^k}_{\text{Fit to } H_{br}(\omega)} + \underbrace{\sum_{\substack{k=1 \\ \text{Odd}}}^M B_k(j\omega)^k}_{\text{Fit to } jH_{bi}(\omega)}. \quad (22)$$

The A_k 's and B_k 's are slowly varying random coefficients; in any given measurement interval, they collectively characterize the short-term frequency response of the propagation medium.

The fittings indicated above can be done, for every 50-ms measurement interval, by using the $2(N + 1) \hat{H}$ samples obtained in that interval. The way the fitting is done depends on the values of M and N . We now consider three possible cases.

Case 1: $M = N$, with $N = 2, 4, 6$, or 8

The \hat{H} samples obtained using $N + 1$ tones can be fitted precisely using an N th-order complex polynomial. Thus, when $M = N$, fitting consists of matching each summation in (21) and (22) to the appropriate \hat{H} samples at the sample frequencies. The resulting equations for the A_k 's are as follows (identical equations apply to the B_k 's, with \hat{H}_{br} 's and \hat{H}_{bi} 's replacing the \hat{H}_{ar} 's and \hat{H}_{ai} 's):

$$A_k = \begin{cases} \hat{H}_{ar,0}, & k = 0, \\ \frac{1}{(\Delta\omega)^k} \sum_{n=1}^{N/2} D_{k/2,n}^e [\hat{H}_{ar,n} - \hat{H}_{ar,0}], & k = 2, 4, \dots, N, \\ \frac{1}{(\Delta\omega)^k} \sum_{n=1}^{N/2} D_{(k+1)/2,n}^o \hat{H}_{ai,n}, & k = 1, 3, \dots, N-1, \end{cases} \quad (23)$$

where $D_{l,m}^e$ and $D_{l,m}^o$ are the (l, m) th elements of the $N/2 \times N/2$ matrices $[D^e]$ and $[D^o]$, respectively; $[D^e]$ and $[D^o]$ are the inverses of the matrices $[d^e]$ and $[d^o]$, respectively; and the (m, l) th elements of $[d^e]$ and $[d^o]$ are

$$d_{m,l}^e = (-1)^l m^{2l} \quad (24)$$

and

$$d_{m,l}^o = -(-1)^l m^{2l-1}. \quad (25)$$

The matrices $[D^e]$ and $[D^o]$ for $N = 2, 4, 6$, and 8 are given in Table II. Note for future reference that the derived A_k 's and B_k 's are weighted sums of the \hat{H} samples.

Since N can be as high as eight, this method of fitting suggests the possibility of eighth-order polynomial modeling. Earlier studies, however, suggest that this order is unnecessarily high for bandwidths of 20 to 40 MHz.^{8,12} Reductions using $M = N = 8$ may therefore involve excessive demands on data storage and analysis, and unduly complicate

Table II—Matrices for evaluating eq. (23)

N	Matrix [D^*]				Matrix [D^n]			
2	[1]				[1]			
4	$-\frac{4}{3}$	$\frac{1}{12}$			$\frac{4}{3}$	$-\frac{1}{6}$		
	$-\frac{1}{3}$	$\frac{1}{12}$			$\frac{1}{3}$	$-\frac{1}{6}$		
6	$-\frac{3}{2}$	$\frac{3}{20}$	$-\frac{1}{90}$		$\frac{3}{2}$	$-\frac{3}{10}$	$\frac{1}{30}$	
	$-\frac{13}{24}$	$\frac{1}{6}$	$-\frac{1}{72}$		$\frac{13}{24}$	$-\frac{1}{3}$	$\frac{1}{24}$	
	$-\frac{1}{24}$	$\frac{1}{60}$	$-\frac{1}{360}$		$\frac{1}{24}$	$-\frac{1}{30}$	$\frac{1}{120}$	
8	$-\frac{8}{5}$	$\frac{1}{5}$	$-\frac{8}{315}$	$\frac{1}{560}$	$\frac{8}{5}$	$-\frac{2}{5}$	$\frac{8}{105}$	$-\frac{1}{140}$
	$-\frac{61}{90}$	$\frac{169}{720}$	$-\frac{1}{30}$	$\frac{7}{2880}$	$\frac{61}{90}$	$-\frac{169}{360}$	$\frac{1}{10}$	$-\frac{7}{720}$
	$-\frac{29}{360}$	$\frac{13}{360}$	$-\frac{1}{120}$	$\frac{1}{1440}$	$\frac{29}{360}$	$-\frac{13}{180}$	$\frac{1}{40}$	$-\frac{1}{360}$
	$-\frac{1}{360}$	$\frac{1}{720}$	$-\frac{1}{2520}$	$\frac{1}{20160}$	$\frac{1}{360}$	$-\frac{1}{360}$	$\frac{1}{840}$	$-\frac{1}{5040}$

model development and usage. For this reason, we also consider the combinations $M = N = 6$, $\Delta f = 5$ MHz; $M = N = 4$, $\Delta f = 10$ MHz; and $M = N = 2$, $\Delta f = 20$ MHz.

Case 2: $M = 2$, with $N = 8$

Another possibility is to assume a low-order polynomial while using all nine tones, in which event the fittings indicated in (21) and (22) are done by least-squares methods. Compared with the case $M = N = 2$, this approach requires more data storage and analysis. On the other hand, it leads to better fitting accuracy and protects against loss of tones caused by equipment problems. Noise effects can be slightly worse for this case, despite the noise-averaging it affords, because the power per tone must be lower to satisfy the transmitter power constraint.

The least-squares polynomial fitting method is well known,¹³ and so we give only the results:

$$A_0 = \frac{59}{231} \left[\hat{H}_{ar,0} + 2 \sum_{n=1}^4 \left(1 - \frac{5n^2}{59} \right) \hat{H}_{ar,n} \right], \quad (26)$$

$$A_1 = \frac{1}{30\Delta\omega} \sum_{n=1}^4 n \hat{H}_{ai,n}, \quad (27)$$

$$A_2 = \frac{5}{231\Delta\omega^2} \left[\hat{H}_{ar,0} + 2 \sum_{n=1}^4 \left(1 - \frac{3n^2}{20} \right) \hat{H}_{ar,n} \right]. \quad (28)$$

The equations for B_0 , B_1 , and B_2 are identical, with $\hat{H}_{br,n}$ and $\hat{H}_{bi,n}$ replacing $\hat{H}_{ar,n}$ and $\hat{H}_{ai,n}$. Again, the A_k 's and B_k 's are weighted sums of the \hat{H} samples.

Case 3: $M = 1$, with $N = 8$

The case $M = 1$ represents the ultimate in modeling simplicity, namely, a response function which is first order in $j\omega$. Strong experimental evidence for such a polynomial has been reported.^{8,9} Assuming least-squares fitting to nine tones, we find that

$$A_0 = \frac{1}{9} \left[\hat{H}_{ar,0} + 2 \sum_{n=1}^4 \hat{H}_{ar,n} \right]; \tag{29}$$

A_1 is given by (27); and the same equations give B_0 and B_1 , with $\hat{H}_{br,n}$ and $\hat{H}_{bi,n}$ replacing $\hat{H}_{ar,n}$ and $\hat{H}_{ai,n}$.

3.2 Error measures

The MPF response function, as reconstructed from the data-derived A_k 's and B_k 's, is

$$\hat{H}(\omega) = \sum_{k=0}^M (A_k + jB_k)(j\omega)^k. \tag{30}$$

Since the coefficients are weighted sums of noisy H samples, (19), we can write $\hat{H}(\omega)$ in the form

$$\hat{H}(\omega) = \underbrace{\tilde{H}(\omega)}_{\substack{\text{Weighted sum} \\ \text{of true} \\ H \text{ samples}}} + \underbrace{Z(\omega)}_{\substack{\text{Weighted sum} \\ \text{of random} \\ \text{noise samples}}} \tag{31}$$

Whereas $\tilde{H}(\omega)$ is noiseless, it is not identical to the true response function, $H(\omega)$, but is a polynomial approximation.

We now define the error function

$$\begin{aligned} \epsilon(\omega) &\triangleq \hat{H}(\omega) - H(\omega) \\ &= \underbrace{[\tilde{H}(\omega) - H(\omega)]}_{\substack{\text{Approximation} \\ \text{error, } \epsilon(\text{approx})}} + \underbrace{Z(\omega)}_{\substack{\text{Noise error,} \\ \epsilon(\text{noise})}} \end{aligned} \tag{32}$$

If B is some bandwidth about f_c over which $H(\omega)$ is to be characterized, then the mean-square error for a given polynomial fit can be defined as

$$\begin{aligned} \bar{\epsilon}^2 &\triangleq \frac{1}{B} \int_{-B/2}^{B/2} |\epsilon(\omega)|^2 df \\ &= \underbrace{\frac{1}{B} \int_{-B/2}^{B/2} |\tilde{H}(\omega) - H(\omega)|^2 df}_{\epsilon^2(\text{approx})} + \underbrace{\frac{1}{B} \int_{-B/2}^{B/2} |Z(\omega)|^2 df}_{\epsilon^2(\text{noise})}. \end{aligned} \tag{33}$$

A useful normalizing quantity for these mean-square errors is the mean-square gain of the fading channel, i.e.,

$$\overline{H^2} \triangleq \frac{1}{B} \int_{-B/2}^{B/2} |H(\omega)|^2 df. \quad (34)$$

Radio system experience has shown that, for $B \geq 20$ MHz, $\overline{H^2}$ is seldom lower than 10^{-4} (40-dB total power fade).

Let us now interpret B as the bandwidth over which the intended receiver output (undistorted by MPF) has a roughly uniform power spectrum, and outside which the spectral content is small. Accordingly, ϵ^2 accurately represents the mean-square error in predicting the output signal using $\hat{H}(\omega)$ for the MPF response and $\overline{H^2}$ accurately represents the true mean-square output. It is thus reasonable to say that a given polynomial fit is valid if $\overline{\epsilon^2(\text{noise})}/\overline{H^2}$ and $\overline{\epsilon^2(\text{approx})}/\overline{H^2}$ are both 10^{-3} or less. Since $\overline{\epsilon^2(\text{noise})}$ is independent of $H(\omega)$, and since $\overline{H^2}$ is seldom lower than 10^{-4} , our requirement for $\overline{\epsilon^2(\text{noise})}$ is that it be 10^{-7} or less.

IV. MEASUREMENT NOISE ERRORS

4.1 Method of analysis

We compute $\overline{\epsilon^2(\text{noise})}$ by forming $Z(\omega)$ in (31), using (30) and the governing formulas for A_k and B_k . For simplicity, we use the first-order model (Case 3: $M = 1$, $N = 8$) to exemplify the approach.

From (30) and (31),

$$Z(\omega) = \sum_{k=0}^1 (\delta A_k + j\delta B_k)(j\omega)^k, \quad (35)$$

where δA_0 and δA_1 are the noises in A_0 and A_1 , which are computed using (27) and (29); and similar definitions and equations apply for δB_0 and δB_1 . These noises are related to the noise components $\zeta_{ar,0}$, $\zeta_{ar,1}$, etc., in the \hat{H} samples of (27) and (29). Thus,

$$\delta A_0 = \frac{1}{9} \left[\zeta_{ar,0} + 2 \sum_{n=1}^4 \zeta_{ar,n} \right], \quad \delta A_1 = \frac{1}{30\Delta\omega} \sum_{n=1}^4 n\zeta_{ai,n}, \quad (36)$$

and similarly for δB_0 and δB_1 . We now proceed in the obvious way: δA_0 , δA_1 , etc., are combined with (35); $|Z(\omega)|^2$ is formed; and the integration in (33) is performed to obtain $\overline{\epsilon^2(\text{noise})}$. In doing so, we make use of the statistical independence among all the noise samples. We also use (16) for their variances. The result is of the form

$$\overline{\epsilon^2(\text{noise})} = \frac{kTbN_F}{2g_0^2} \sum_{n=0}^{N/2} \frac{S_n(B/\Delta f)}{p_n}, \quad (37)$$

where, for the case under study, $N = 8$ and

$$S_n(B/\Delta f) = \frac{4}{81} \left[1 + \frac{3}{1600} \left(\frac{B}{\Delta f} \right)^2 n^2 \right], \quad n = 0, N/2. \quad (38)$$

For all of the other cases considered, the analytical approach and the form of the result are as shown above; only the specific equations for $S_n(B/\Delta f)$ vary.

The final step in computing $\overline{\epsilon^2(\text{noise})}$ is to specify the values of the transmitted tone powers. We have considered two approaches, namely (i) assuming all p_n 's to be equal, adding up to some specified value (\bar{P}) of average transmitted power; and (ii) choosing the p_n 's so as to minimize $\overline{\epsilon^2(\text{noise})}$, subject to the same average power specification. In the first approach, we apportion power according to the rule

$$p_n = \frac{\bar{P}}{N+1}, \quad \text{all } n. \quad (39)$$

In the second approach, we use the method of Lagrangian multipliers to minimize (37), subject to the constraint

$$p_0 + 2 \sum_{n=1}^{N/2} p_n = \bar{P}. \quad (40)$$

The result in that case is

$$p_n = \frac{\sqrt{(1 + \delta_n^0) S_n(B/\Delta f)} \bar{P}}{\sqrt{2S_0(B/\Delta f)} + 2 \sum_{n=1}^{N/2} \sqrt{S_n(B/\Delta f)}}, \quad \text{all } n, \quad (41)$$

where δ_n^0 is the Kronecker delta function. Note that the optimal variation of p_n with n depends on the bandwidth ratio, $B/\Delta f$.

We have reduced (37) to numerical results using the parameter values in Table I. We now present our findings.

4.2 Results for Case 1 ($M = N = 2, 4, 6, \text{ or } 8$)

Figure 1 shows curves of $\overline{\epsilon^2(\text{noise})}$ vs. B for the various possibilities under Case 1. We can make the following observations:

(i) The noise errors and their rate-of-growth with B depend strongly on the polynomial order M , both being less for lower M . From a noise standpoint then, M should be chosen to be low. The competing factor influencing this choice is the approximation error, which we consider later.

(ii) Optimizing p_n does not improve significantly on the use of uniform tone powers. There is no compelling reason, therefore, to taper the p_n 's. On the other hand, doing so entails no price in complexity and may offer benefits, e.g., tapering of p_n may reduce out-of-band spurious tones caused by transmitter nonlinearities. We will not explore this topic here, but present in Table III some examples of optimal p_n variations.

(iii) Values of $\overline{\epsilon^2(\text{noise})}$ less than 10^{-7} can be attained, using either uniform or optimal p_n 's, for bandwidths up to 34 MHz or more, depending on M . For purposes of modeling, errors of this magnitude can be regarded as negligible, as noted in Section 3.2.

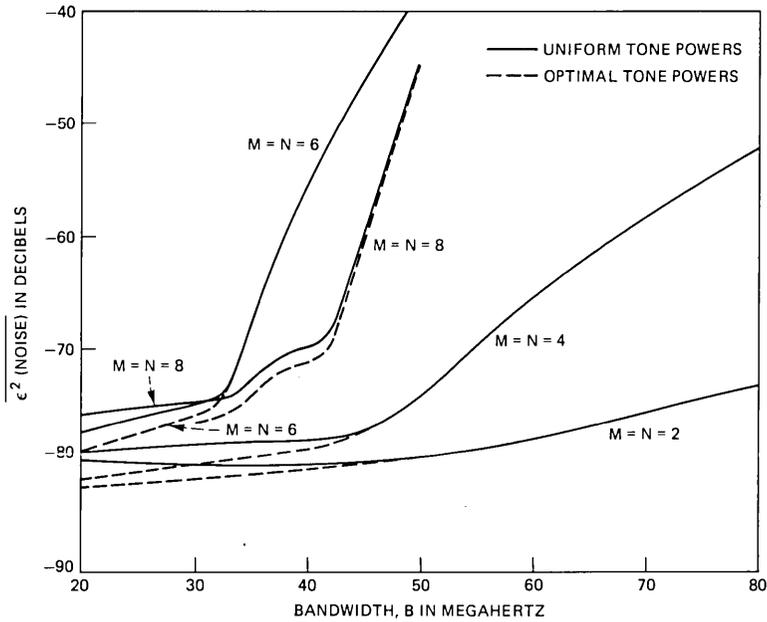


Fig. 1—Mean-square noise errors, $\epsilon^2(\text{noise})$ in dB, for Case 1 ($M = N = 2, 4, 6$ or 8). For each N , the curve for optimal tone powers (dotted) merges with the one for uniform tone powers (solid).

4.3 Results for Cases 2 and 3 ($M = 2, N = 8$ and $M = 1, N = 8$)

Figure 2 shows curves of $\epsilon^2(\text{noise})$ vs. B for Cases 2 and 3. These results show the effects of noise when first- and second-order polynomials are formed via least-squares fitting to nine received tones. For purposes of comparison, Fig. 2 also repeats the results for $M = N = 2$ under Case 1, i.e., second-order matching to three received tones. We note the following:

- (i) As before, fitting $H(\omega)$ with a lower-order polynomial leads to smaller noise-related errors.
- (ii) Comparing the two ways of getting $M = 2$, noise errors are

Table III—Optimal tone powers for Case 1, $B = 40$ MHz (in dB above power of central tone)

Frequency \ N	± 5 MHz	± 10 MHz	± 15 MHz	± 20 MHz
2	—	—	—	-3.02
4	—	-0.10	—	-4.05
6	-0.80	-3.30	-7.60	—
8	-0.36	-1.36	-2.63	-7.95

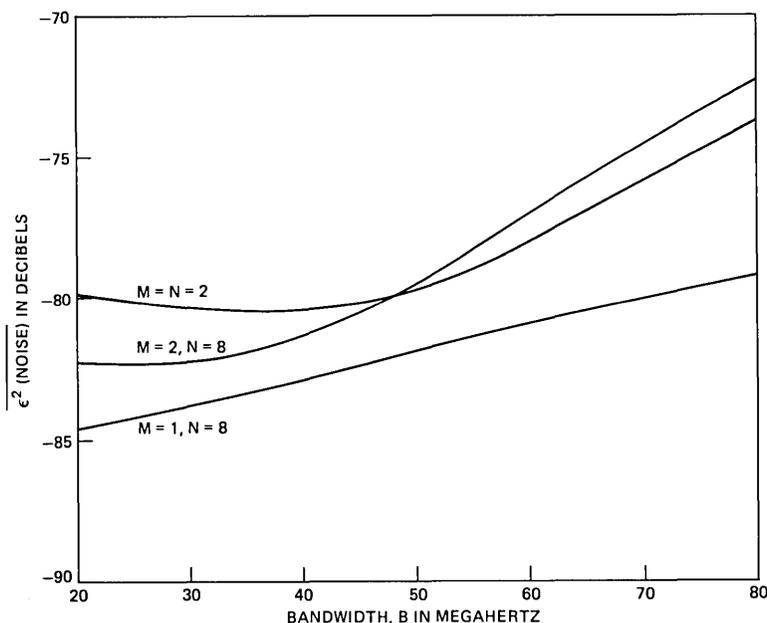


Fig. 2.—Mean-square noise errors, $\epsilon^2(\text{noise})$ in dB, for Cases 2 and 3 ($M = 2, N = 8$ and $M = 1, N = 8$). Results shown are for uniform tone powers; results for optimal tone powers are lower by less than 0.5 dB.

sometimes smaller when $N = 2$ and sometimes smaller when $N = 8$. The explanation lies in two offsetting factors: Using nine tones leads to less power per tone, but affords averaging benefits not available when using three tones. Depending on B , one or the other of these factors dominates.

(iii) For all three low-order polynomial approaches considered, $\overline{\epsilon^2(\text{noise})}$ is less than 10^{-7} for bandwidths up to 90 MHz or more.

4.4 Generalizations and extensions

We emphasize that $\overline{\epsilon^2(\text{noise})}$ in (37) scales readily in the quantities $bN_F/\bar{P}g_0^2$ and $B/\Delta f$. Therefore, the curves of Figs. 1 and 2, however particularized to the present experiment parameters, can easily be scaled to reflect other conditions.

With this in mind, we can account in a simple and fairly accurate way for both quantization noise (previously neglected) and changes in \bar{P} , g_0 , b , and N_F . We assume the quantization and Gaussian noise powers to be additive, invoke (16) and (17), and obtain the following rule: If $bN_F/\bar{P}g_0^2$ is changed by a factor C from the one obtained using Table I, then $\overline{\epsilon^2(\text{noise})}$ should be adjusted by the factor $(C + 0.25)$. For $C = 1$ (no change in $bN_F/\bar{P}g_0^2$), quantizing noise adds roughly 25 percent (1 dB) to the mean-square noise error.

V. APPROXIMATION ERRORS

5.1 Method of analysis

The approximation error defined in (32) results from sampling $H(\omega)$ at a finite number of discretely-spaced frequencies. The mean-square error, as defined in (33), depends on B , Δf , M , N , and—unlike the noise error—on the specific function being approximated. We consider a highly accurate polynomial approximation to be one for which η , defined by

$$\eta \triangleq \overline{\epsilon^2(\text{approx})/H^2}, \quad (42)$$

is less than or equal to 10^{-3} . Our aim here is to calculate η using a function for $H(\omega)$ that is “representative,” in a sense to be described shortly. The method of computation, given $H(\omega)$, follows directly from the relationships in Section III.

The basic form assumed for $H(\omega)$ is the one used by W. D. Rummler in previous studies.^{5,6,9} To fit measured amplitude responses in a 25.3-MHz bandwidth near 6 GHz, Rummler expressed $H(\omega)$ by the function $a[1 - b \exp(j\theta) \exp(-j\omega\Delta\tau)]$, where a , b , θ , and $\Delta\tau$ are function parameters to be chosen. For each of 24,920 response measurements, these four parameters were chosen to give a least-squares functional fit to the data. Error analyses of the results (Fig. 17 of Ref. 5) revealed a high degree of fitting accuracy over nearly all of the data records. For this reason, we use the same function here to represent the actual underlying response of multipath fading.

Rummler also analyzed the effect of fixing the delay parameter $\Delta\tau$ at 6.31 ns, and choosing (a, b, θ) to fit each data record. While not as good in all cases, the “fixed delay” model was found to be highly accurate over at least 98 percent of the data base. With this in mind, we conjecture that the above function, with $\Delta\tau \leq 10$ ns, can be used to represent the vast majority of MPF responses. We will invoke this conjecture later in tabulating bandwidths over which polynomials of specified order can be used to model multipath fading.

Our approach in what follows will be to find η for specified B and $\Delta\tau$, with b and θ chosen to give worst-case polynomial fitting results. This should provide conservative estimates of the limitations of polynomial fitting. We will also assume a phase adjustment to the function used by Rummler, to enforce consistency with the data processing approach described earlier. In particular, we specify the response function to be

$$H(\omega) = a(1 - be^{j\theta}e^{-j\omega\Delta\tau}) \exp[j(\Delta\Phi + \omega\Delta T)], \quad (43)$$

where $\Delta\Phi$ and ΔT were defined in Section II. As noted there, the data reduction process chooses $\Delta\Phi$ in such a way that $H(0)$ is made real, and chooses ΔT in such a way that η is minimized for a given M -th-

order polynomial fit. Accordingly, we specify that $\exp(j\Delta\Phi) = (1 - be^{-j\theta})/\sqrt{1 + b^2 - 2b \cos\theta}$; search over ΔT so as to minimize η for a given fitting polynomial and given (b, θ) ; and search over b and θ so as to maximize that minimized η .^{*} Thus, we obtain the measure

$$\eta_0 \triangleq \text{Max}_{(b, \theta)} \left\{ \text{Min}_{\Delta T}(\eta) \right\}, \quad \text{worst-case } \eta, \text{ given } \Delta\tau \text{ and } B. \quad (44)$$

We have done this for each of the different combinations of M and N described earlier, using $B\Delta\tau$ as a variable and B as a parameter. We now present our findings.

5.2 Results for Case 1 ($M = N = 2, 4, 6, \text{ or } 8$)

The worst-case combination of (b, θ) for any $B, \Delta\tau$, and polynomial order is found to be $(1.0, 0.0)$. This combination corresponds to total fading at the band center ($\omega = 0$). Also, the optimal ΔT for $b = 1.0$ is precisely $\Delta\tau/2$.

The variations of η_0 with $B\Delta\tau$, with M and B as parameters, are shown in Fig. 3. As expected, η_0 is a strongly decreasing function of M , a strongly increasing function of $B\Delta\tau$, and a weakly increasing function of B alone for given $B\Delta\tau$.

A useful empirical formula for these results, accurate to within 40 percent for $B \leq 80$ MHz and $B\Delta\tau \leq 1.2$, is

$$\eta_0 = k(M, B) 10^{-M} (B\Delta\tau)^{2M} [1 + \alpha(B\Delta\tau)], \quad (45)$$

where $\alpha = 0.57$ and $k(M, B)$ is a fairly mild function of M and B . Numerical values for $k(M, B)$ are given in Table IV.

5.3 Cases 2 and 3 ($M = 2, N = 8$ and $M = 1, N = 8$)

Curves of η_0 for these cases are shown in Fig. 4, where the results for $M = N = 2$ are repeated for comparison purposes. There is an evident improvement when second-order polynomials are derived using nine tones instead of three. An empirical formula for these results (again, accurate to within 40 percent for $B \leq 80$ MHz and $B\Delta\tau \leq 1.2$) is given by (45), where $\alpha = 0.32$; the power of $(B\Delta\tau)$ is replaced by 4; and numerical values for $k(M, B)$ are given in Table IV.

5.4 Optimizing ΔT using measured data

In minimizing η with respect to ΔT , (44), we used computer search procedures and exploited our assumed knowledge of the underlying $H(\omega)$. In an actual measurement the latter is not possible, as the only information available for optimizing ΔT is the set of measured G samples.

^{*} The amplitude factor a has no impact on fitting accuracy and so is set to unity for purposes of this study.

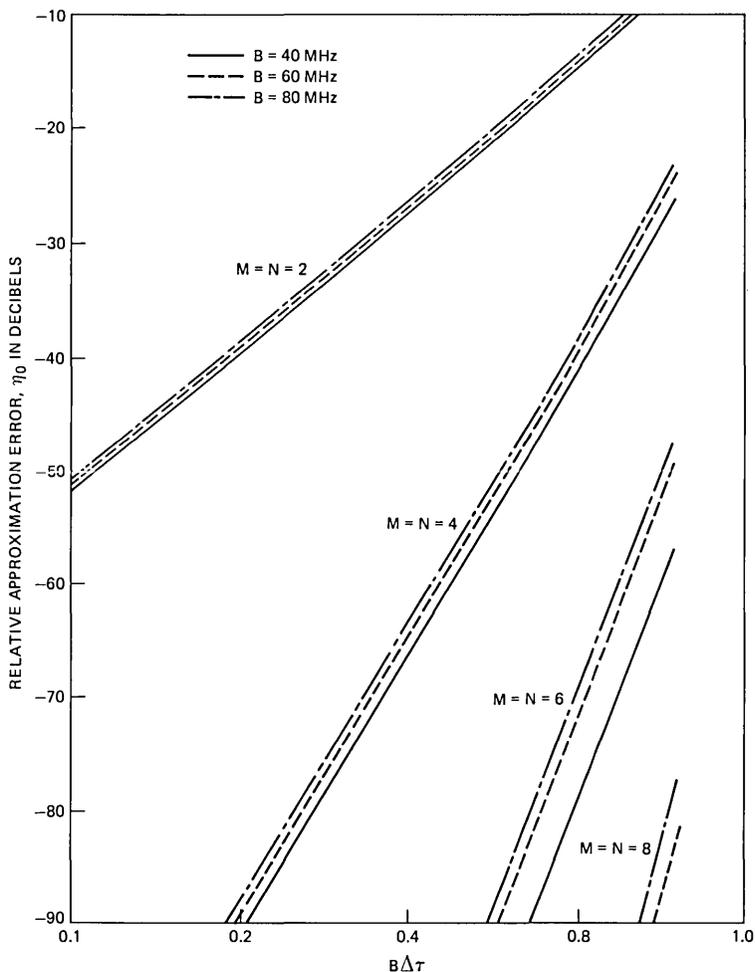


Fig. 3—Relative mean-square approximation errors, η_0 in dB, for Case 1 ($M = N = 2, 4, 6, \text{ or } 8$). The variable is $B\Delta\tau$, where $\Delta\tau$ is the delay spread quantity in Eq. (43). Bandwidth, B , and M are parameters.

By experimenting with different strategies, we have identified a computationally efficient procedure that uses these samples to optimize ΔT . It consists of computing $C_2 \triangleq (A_2^2 + B_2^2)$ as a function of ΔT and choosing that ΔT for which C_2 is a minimum. The procedure is even simpler than it might seem; for, if $A_{n,o}$ is A_n computed for $\Delta T = 0$, then A_2 for any other ΔT is just

$$A_2 = A_{2,o} + A_{1,o}\Delta T + \frac{1}{2}A_{0,o}\Delta T^2, \quad (46)$$

and an identical relationship applies to B_2 . These results can be

Table IV—Values of $k(M, B)$ in eq. (45)

Case	Case 1 ($M = N$)				Cases 2 and 3 ($N = 8$)	
$B \setminus M$	8	6	4	2	2	1
40 MHz	1.31×10^{-4}	0.138	3.32	6.39	1.47	1.51
60 MHz	1.95×10^{-2}	0.734	5.31	6.98	2.56	1.78
80 MHz	5.83×10^{-2}	1.14	6.17	7.20	4.26	2.26

derived using (6), (21), and (22), and the power series expansion of $\exp(j\omega\Delta T)$.

Thus, "optimizing" ΔT amounts to finding the minimum of a fourth-order polynomial in that quantity. This can be done very efficiently

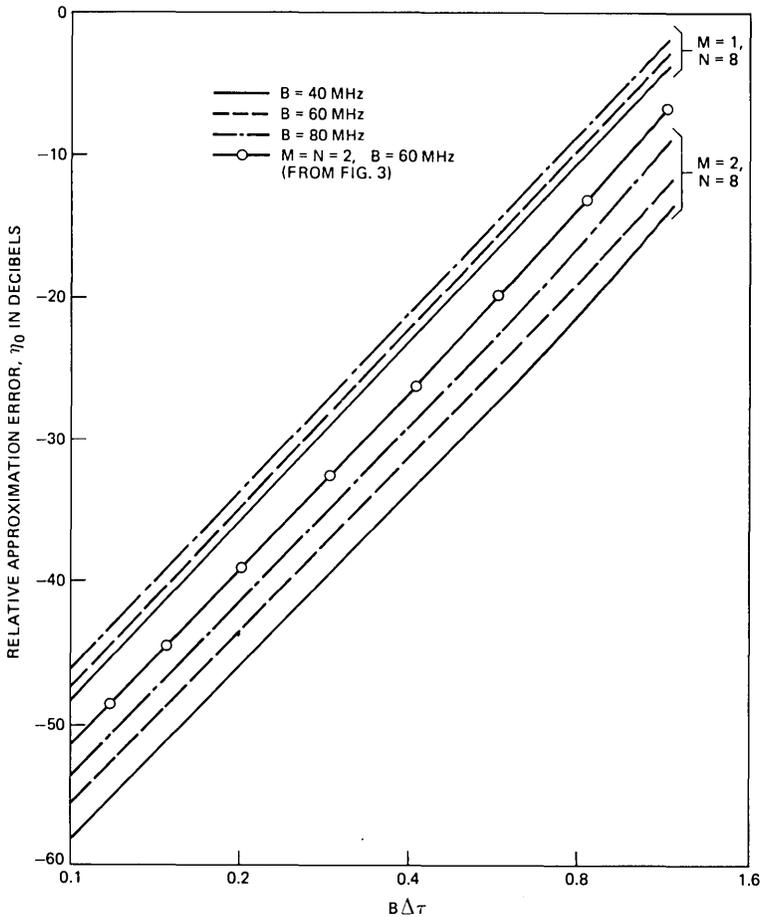


Fig. 4—Relative mean-square approximation error, η_0 in dB, for Cases 2 and 3 ($M = 2, N = 8$ and $M = 1, N = 8$). Same variable and parameters as in Fig. 3.

using computer algorithms. Our investigation of this procedure consisted of minimizing C_2 with respect to ΔT for numerous $(b, \theta, \Delta\tau)$, and using the resulting values of ΔT to compute η . We obtained remarkably similar answers, in most cases, to those obtained with ΔT rigorously optimized [i.e., via computer search using the known function $H(\omega)$]. The worst departures occur when η is either so large as to be of no interest or so small that the error increases do not matter.

Other measures derivable from the data, such as $[(A_2 - A_1^2)^2 + (B_2 - B_1^2)^2]$, may be even more appropriate. Quantities such as this, or C_2 as defined above, are reliable indicators of the "curvature" in $H(\omega)$; choosing Δt to minimize them should therefore maximize the fitting accuracy of low-order polynomials. Our results give empirical support for this principle.

VI. SUMMARY AND CONCLUSION

The major findings of this study are summarized in Table V. The first row gives, for each combination of M and N considered, the bandwidth below which polynomial fitting yields $\epsilon^2(\text{noise}) \leq 10^{-7}$. These bandwidths are derived for uniform tone powers and the parameter values listed in Table I. They cannot be increased much without dramatic (and unlikely) improvements in transmitter power, noise bandwidths, and quantizing precision.

The second row of Table V gives the corresponding bandwidths below which $\eta \leq 10^{-3}$. They are derived for the response function (43), with b and θ having worst-case values and $\Delta\tau = 10$ ns. The third row gives the maximum bandwidth which satisfy *both* mean-square error requirements. Since these requirements are quite stringent, we present in the fourth row the results of relaxing both of them by 6 dB (factor of four).

These tabulations suggest that the vast majority of MPF responses can be accurately approximated, over bandwidths of 40 (62) MHz, by

Table V—Bandwidths, in MHz, over which $H(\omega)$ can be accurately characterized

Case	Case 1 ($M = N$)				Cases 2 and 3 ($N = 8$)	
Requirement \ M	8	6	4	2	2	1
$\epsilon^2(\text{noise}) \leq 10^{-7}$	40	34	55	96	90	180
$\eta \leq 10^{-3}$	140	125	96	35	48	28
Both requirements	40	34	55	35	48	28
Requirements relaxed 6 dB	43	37	62	48	62	40

Notes:

(1) Results assume $\Delta\tau = 10$ ns.

(2) Results in excess of 80 MHz are extrapolated estimates.

complex polynomial functions of just first (second) order. We surmise, then, that the appropriate modeling approach is to form first- and second-order polynomials using nine-tone measurements ($N = 8$) and least-squares data fitting. With these low orders, there is maximal simplicity in developing and using the model, with no substantial sacrifice in accuracy. By using nine tones, moreover, polynomials of higher order (up to eight) can also be deduced from the data, thereby enabling the theoretical predictions given here to be fully tested.

VII. ACKNOWLEDGMENT

We are grateful for the helpful comments of W. D. Rummler.

REFERENCES

1. W. C. Jakes, Jr., "An Approximate Method to Estimate an Upper Bound on the Effect of Multipath Delay Distortion on Digital Transmission," *IEEE Trans. Commun., COM-27*, No. 1 (January 1979).
2. L. J. Greenstein and V. K. Prabhu, "Analysis of Multipath Outage with Applications in 90-Mb/s PSK Systems at 6 and 11 GHz," *IEEE Trans. Commun., COM-27*, No. 1 (January 1979).
3. C. W. Anderson, S. Barber, and R. Patel, "The Effect of Selective Fading on Digital Radio," *Int. Conf. Commun. 1978, Conf. Record, 2* (1978), pp. 33.5.1-6.
4. C. W. Lundgren and W. D. Rummler, "Digital Radio Outage Due to Selective Fading-Observation vs. Prediction from Laboratory Simulation," *B.S.T.J.*, *58*, No. 7 (May-June 1979), pp. 1073-100.
5. W. D. Rummler, "A New Selective Fading Model: Application to Propagation Data," *B.S.T.J.*, *58*, No. 7 (May-June 1979), pp. 1037-71.
6. W. D. Rummler, "Extensions of the Multipath Fading Channel Model," *Int. Conf. Commun. 1979, Conf. Record, 2* (1979), pp. 32.2.1-5.
7. L. J. Greenstein and B. A. Czekaj, "A Statistical Model for Multipath Fading Channel Responses," *Int. Conf. Commun. 1979, Conf. Record, 2* (1979), pp. 32.1.1-5.
8. L. J. Greenstein and B. A. Czekaj, "A Polynomial Model for Multipath Fading Channel Responses," *B.S.T.J.*, *59*, No. 7 (September 1980), pp. 1197-225.
9. W. D. Rummler, "Advances in Multipath Channel Modeling," *Int. Conf. Commun. 1980, Conf. Record, 3* (1980), pp. 52.3.1-5.
10. A. B. Crawford and W. C. Jakes, Jr., "Selective Fading of Microwaves," *B.S.T.J.*, *31*, No. 1 (January 1952), pp. 68-90.
11. L. J. Greenstein and P. J. Fitzgerald, "Phasing Multitone Signals to Minimize Peak Factors," (to be published).
12. L. J. Greenstein, "A Multipath Fading Channel Model for Terrestrial Digital Radio Systems," *IEEE Trans. Commun., COM-26*, No. 8 (August 1978), pp. 1247-50.
13. J. R. Green and D. Margerison, *Statistical Treatment of Experimental Data*, New York: Elsevier, 1977, Chapter 15.

The Quality Measurement Plan (QMP)

By BRUCE HOADLEY

(Manuscript received August 5, 1980)

This paper describes the Quality Measurement Plan (QMP), a recently implemented system for reporting the quality assurance audit results to Bell System management. QMP replaces the T-rate system, which evolved from the pioneering statistical work of Shewhart and Dodge during the 1920's and 1930's at Bell Laboratories. Box and whisker plots are used for graphically displaying confidence intervals for the quality of the current production. The confidence interval is computed from both current and past data and is derived from a new Bayesian approach to the empirical Bayes problem for Poisson observations. Here we discuss the rationale, mathematical derivations, dynamics, operating characteristics, and many comparative examples. We show that QMP reduces statistical errors relative to the earlier T-rate system.

I. INTRODUCTION

1.1 Quality assurance

The responsibility of the Bell Laboratories Quality Assurance Center (QAC) is "to ensure that the communications products designed by Bell Laboratories and bought by Bell System operating companies from Western Electric Company, Incorporated will meet quality standards and will perform as the designers intended."¹ This obviates the need for each operating company to carry out its own acceptance inspection.

To meet this responsibility, the QAC works with its Western Electric (WE) agents, the Quality Assurance Directorate (QAD),² and Purchased Products Inspection (PPI) organizations. However, as stated in Ref. 1, "The primary responsibility for quality lies with the line organizations: Bell Laboratories for the quality of design and Western Electric for the quality of manufacture, installation, and repair." The quality assurance organizations conduct independent activities to assure quality to the operating companies.

1.2 Quality assurance audit

The quality assurance organizations have two major activities. The first is to conduct quality audits where products change hands, either within WE or between WE and the operating companies. Examples are manufacturing, installation, and repair audits. The second concerns a collection of field quality monitoring activities. Examples are the Product Performance Surveys. These are designed sample surveys of reported field troubles.

An audit is a highly structured system of inspections done on a sampling basis. The ingredients of an audit are: (i) sampling method, (ii) scope of inspection, (iii) quality standards, (iv) nonconformance procedures, (v) defect assessment practices, (vi) quality rating method, and (vii) report formats.

The sampling method along with the scope of inspection determines what tests will be performed on what units of product or attributes of product. The statistics and economics of sampling, the engineering requirements, and the field effect of defects play the central roles in determining the sampling and the scope of inspection.

The quality standards are numerical values expressed in defects, defectives or demerits per unit. They are set by the QAC in consultation with the QAD. The standards are target values, reflecting a tradeoff between first cost and maintenance costs.

The nonconformance procedures are rules for detecting and disposing of audited lots that are excessively defective with respect to a particular set of engineering requirements.

The defect assessment practices are a set of transformations that map defects found into defects assessed for quality rating purposes. A terminal strip may have all ten connections off by one position, but, the consequences of these ten defects found are much less than ten independent occurrences of this defect. Therefore, less than ten defects are assessed.

The quality rating method and report formats determine how the results of the audit are presented to Bell System management. For example, a product is reported as "Below Normal," when it fails a statistical test of the hypothesis that the quality standard is being met.

1.3 The quality measurement plan (QMP)

The statistical foundations of the audit ingredients were developed by Shewhart, Dodge, and others, starting in the 1920's and continuing through to the middle 1950's. This work was documented in the literature in Refs. 3 to 6.

In recent years, research has been carried out to evaluate the application of modern statistical theories to the audit ingredients. An important idea is summarized in an article by Efron and Morris⁷ which

explains a paradox discovered by Stein.⁸ When you have samples from similar populations, the individual sample characteristics are not the best estimates of the individual population characteristics. Total error is reduced by shrinking the individual sample characteristics part way towards the grand mean over all samples. Efron and Morris used baseball batting averages to illustrate the point. But the problem of estimating percent defective in quality assurance is the same problem. And you are always concerned with similar populations—for example, the population of design-line telephones produced for each of several months.

This idea was originally explored in Ref. 9. The idea has now evolved into the Quality Measurement Plan (QMP). QMP is the recently implemented system for conducting three of the audit ingredients: defect assessment, quality rating, and quality reporting.

As a quick introduction to QMP, consider Fig. 1. This is a comparison of the QMP reporting format (Fig. 1a) with the old *T*-rate reporting format (Fig. 1b). Each year is divided into eight periods. On the bottom, the *T*-rate is plotted for each period and it measures the difference between the observed and standard defect rates in units of sampling standard deviation (given standard quality). The idea is that if the *T*-rate is, e.g., less than negative two, then the hypothesis of standard quality is rejected. Section II considers the exact rules for exception reporting under the *T*-rate system.

Under QMP, a box and whisker chart is plotted each period. The box chart is a graphical representation of the posterior distribution of current population quality on an index scale. The index value one is the standard on the index scale and the value two means twice as many defects as expected under the quality standard. The posterior probability that the population index is larger than the top whisker is 0.99. The top of the box, the bottom of the box, and the bottom whisker correspond to the probabilities 0.95, 0.05, and 0.01, respectively.

The heavy “dot” is a Bayes estimate of the long run *process average*; the “cross” is the observed value in the current sample; and the “dash” near the middle of the box is the posterior mean of the current population index and is called the *Best Measure* of current quality. The process averages, “dots,” are joined to show trends.

Although the *T*-rate chart and the QMP chart often convey similar messages, there are differences. The QMP chart provides a measure of quality; the *T*-rate chart does not. For example, in 7806 (Period 6 of 1978) both charts imply that the quality is substandard, but the QMP chart also implies that the population index is somewhere between one and two.

QMP and the *T*-rate use the past data in very different ways. QMP

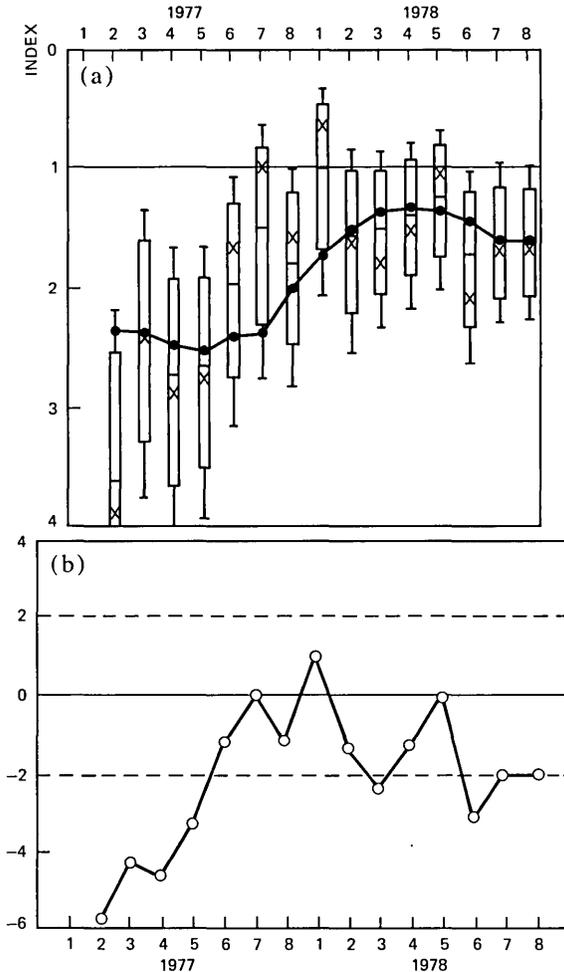


Fig. 1—QMP versus the T -rate. The box and whisker plot in (a) is the QMP replacement of the T -rate. One is standard on the index scale; two is a defect rate of twice standard. The box and whisker are 90 and 98 percent confidence intervals for production during the period; the “crosses” are the indices in the samples; the “dots” are process averages; and the “dashes” in the middle of the boxes are Best Measures of current quality derived from empirical Bayes theory. (b) is a time series of T -rates. Each point measures the difference between observed quality and expected quality on a standard deviation scale. Notice that the sixth period of 1977 and the fourth period of 1978 are the same in the T -rate chart but quite different in the QMP chart.

uses the past sample indices, but makes an inference about current quality. The T -rate system uses runs criteria based on attributes of the T -rate, such as “less than zero,” and can make an inference about past quality. In Fig. 2, 7707, the T -rate signals an exception, because six T -rates in a row are less than zero, indicating that quality has not

been standard for all six periods. But for QMP, the standard is well within the box, indicating normal current quality. The different treatment of past data is also illustrated in Fig. 1. Comparing 7706 with 7804 reveals very similar T -rates, but QMP box charts with different messages.

The T -rate system is based on the assumption that the total number of defects in a rating period has a normal distribution. QMP is based on the Poisson distribution. This difference is important for small audits, as shown in Section VII.

QMP was on trial for two years and was applied to 20,000 sets of audit data. The relatively simple QMP algorithm published in Ref. 9

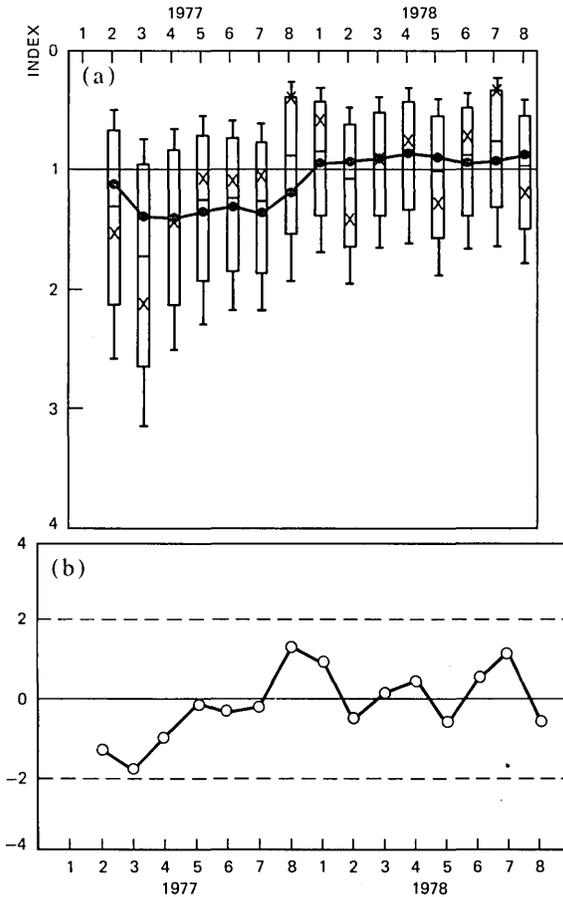


Fig. 2—A weak T -rate exception. The seventh period of 1977 was reported as a quality exception because six T -rates in a row were less than zero. For QMP, it would have been reported as normal. This is because QMP provides a statistical inference about current production only, even though past data is used.

was used originally. This simple algorithm worked for most data sets, but not all (e.g., zero defects in every period). The relatively complex algorithm discussed in Section IV is the result of a lengthy fine-tuning process, designed to make the algorithm work for every case. This is why the full power of Bayes theorem with empirically based prior distributions had to be used.

1.4 Relationship to the empirical Bayes approach

Note that in the QMP box chart, the Best Measure always lies between the estimated process average and the current sample index. The Best Measure is a shrinkage of the sample index towards the estimated process average. In 7706 of Fig. 1a, the shrinkage is away from standard; but, in 7804, it is towards the standard.

The Best Measure is related to the class of estimators described by Efron and Morris.¹⁰ In the cited reference, they provide a foundation for Stein's paradox with an empirical Bayes approach. In Ref. 7, they used baseball data to illustrate Stein's paradox. There is a clear analogy between percent defective in a quality assurance application and a baseball batting average. The data in Ref. 7 was for many players at a given point in time. The QMP algorithm works with the data for one product over time. So a better baseball analogy would be one player over time.

Table I contains batting average data for Thurman Munson from 1970 through 1978. This data was collected and analyzed by S. G. Crawford and is displayed graphically in Fig. 3. The "crosses" are Munson's batting averages reported on the last Sunday of April for each year. The "boxes" are Munson's batting averages at the end of the season. The dashed line is the average of the "crosses."

The early season averages are analogous to the audit data. The averages are the results from small samples of the populations. The populations are the finite populations of "at bats" for each season. In

Table I—Batting average data for Thurman Munson

Year	Reported* on Last Sunday in April			End of Season			QMP Estimate of Batting Average
	At Bats	Hits	Batting Average	At Bats	Hits	Batting Average	
1970	44	7	0.159	453	137	0.302	0.165
1971	37	6	0.162	451	113	0.251	0.168
1972	31	11	0.355	511	143	0.280	0.288
1973	56	20	0.357	519	156	0.301	0.315
1974	72	16	0.222	517	135	0.261	0.233
1975	48	18	0.375	597	190	0.318	0.323
1976	40	13	0.325	616	186	0.302	0.308
1977	42	9	0.214	595	183	0.308	0.273
1978	63	17	0.270	617	183	0.297	0.283

* AP statistics.

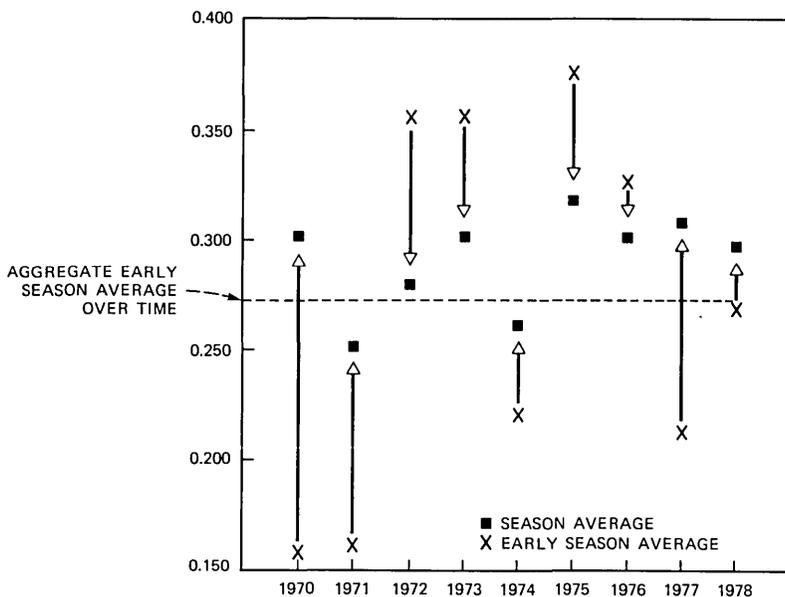


Fig. 3—Batting averages for Thurman Munson. For each year, the movement from the early season average (the sample) to the season average (the population) is always in the direction of the time average of the samples. This suggests strongly that by shrinking the samples towards their time average, one can obtain improved estimates of the populations.

the audit, we are interested in making a statistical inference each period about the current population. So our problem in Fig. 3 is to make a statistical inference each year about the season batting average using only the early season averages observed to date.

As an estimate, one would be tempted to use the maximum likelihood estimate, the early season average. But, in every year, the movement of the batting average from the early season to the season end is in the direction of the aggregate early season average over time. So, paradoxically, the early season averages from other years seem to be relevant to the current season average. It is clear from the data, that a better estimate of the season average is some kind of shrinkage of the early season average towards the aggregate early season average over time. And the amount of shrinkage can depend only on the available data—the early season averages.

What we really have here is a multivariate problem. We observe a nine-dimensional vector of observations whose mean is a vector of population characteristics, one of which we are particularly interested in. Stein⁸ showed (for the normal distribution) that the maximum likelihood estimate of the vector is inadmissible. Why this is true manifests itself in baseball lore. A player that starts the season rela-

tively hot, usually cools off; and a player that starts in a relative slump usually improves. This is due to the nature of sampling error. The hot player is usually partially lucky and the slumping player is usually partially unlucky.

QMP is based on the concepts illustrated by the Munson data. We saw in Fig. 1a that the Best Measure of the population index is between the current observed index and the estimated long-run process average.

The approach used for QMP is actually Bayesian empirical Bayes. The shrinkage factor used is a Bayes estimate of an optimal shrinkage factor. So the Best Measure has the form

$$W \begin{bmatrix} \text{estimated} \\ \text{process} \\ \text{average} \end{bmatrix} + (1 - W) \begin{bmatrix} \text{current} \\ \text{sample} \\ \text{index} \end{bmatrix},$$

where W is a Bayes estimate of

$$\frac{[\text{sampling variance}]}{[\text{sampling variance}] + [\text{process variance}]}$$

The bigger the sampling variance is, relative to the process variance, the more weight is put on the estimated process average.

There are two advantages to the Bayesian empirical Bayes approach over the approach in Ref. 10. One is that the weight, W , is always strictly between zero and one. This is because W is a Bayes estimate of an unknown optimal weight, ω , which has a nondegenerate posterior distribution on the interval $[0, 1]$. The approach taken in Ref. 10 is to use maximum likelihood estimates of ω , which can be one; i.e., total shrinkage to the process average.

The second advantage is that an interval estimate of the current population index can be constructed from its posterior distribution. Most of the literature (e.g., Ref. 10) treats the estimation problem thoroughly, but it provides little guidance for the interval estimation problem.

The QMP algorithm is applied to the Munson data and the QMP estimates of the season averages are given in Table I. The sums of the absolute errors for the maximum likelihood estimates (April averages) and the QMP estimates are 0.603 and 0.331, respectively—a forty-five percent improvement. Notice that the QMP estimates for 1970 and 1971 are close to the April averages. This is because there was no history on Munson. The reduction in total absolute error for the years 1973 through 1978 was sixty-five percent, because of the benefit of history.

1.5 Objectives

This paper is intended to document QMP. It contains the rationale for changing the rating system, a synopsis of QMP features, mathemat-

ical derivations of the rating formulas, the dynamics of QMP, the operating characteristics of QMP, many examples, and the QMP reporting format.

Readers who are interested only in the mathematics of QMP and how it relates to empirical Bayes, may skip Section II. Readers, who are not interested in the mathematical derivation of QMP, may skip Section IV.

II. T-RATE SYSTEM

To understand the rationale behind QMP, one must first understand the *T*-rate system. From this we shall see where things have changed and where things have remained the same.

2.1 Finding defects

The sampling methods along with the scope of inspection provide for a sample of units of count for each set of inspections. A unit of count is either a unit of product or a unit of a product's attribute such as solderless wrapped connections.

The result of conducting a set of inspections is a list of defects found and their descriptions. Frequently, underlying a defect is a variable measurement* that falls outside a range. QMP does not affect the process of finding defects.

2.2 Assessing defects

The defects found sometimes occur in clusters for which the effect of the cluster is nonadditive; i.e., the effect is less than the sum of the effects of the individual defects occurring by themselves. In this case, the number of defects assessed for rating purposes is less than the number found. The defect assessment practices for the *T*-rate system evolved over a 50-year period, so these practices were based on a variety of criteria and engineering judgements. The defect assessment practices under QMP amount to a redesign of the practices using a single principle, which is described in Section 3.1.

2.3 Defect weighting and demerits

The defects assessed are transformed into demerits or defectives or may remain as simple unweighted defects. In an audit based on demerits, each defect assessed is assigned a number of demerits: 100, 50, 10, or 1 for *A*, *B*, *C*, or *D* weight defects, respectively. Guidelines for assigning demerit weights are contained in numerous general and

* For rating transmission characteristics of exchange area cable, some variables measurements are used directly without conversion to defects. We do not treat this case here.

special purpose demerit lists. The principles underlying these demerit lists are described by Dodge in Refs. 5 and 6. In an audit based on defectives, all defects found in a unit of product are analyzed to determine if the unit is considered defective. The assessment is either one or zero defectives. These transformations to demerits, defectives, or defects are not affected by QMP.

2.4 Quality standards

For any set of inspections, the quality engineers in the QAC have established quality standards. To do this, they considered audit scope, shop capability, field performance, economics, complexity, etc. The philosophy of standards is described in Ref. 3. For audits based on defects or defectives, the standards are expressed in defects or defectives per unit. For audits based on demerits, the standards are derived from fundamental defect per unit of count standards for A, B, C, D -type defects. In addition, we use Poisson as the standard distribution of the number of type A defects (for example).

To make this clear, let's consider a simple example. Suppose in a sample of size n , there are X_A, X_B, X_C, X_D -type A, B, C, D defects. The definition of standard quality is that X_A, X_B, \dots are independent and have Poisson distributions with means $n\lambda_A, n\lambda_B, \dots$. The number of demerits in the sample is

$$D = 100X_A + 50X_B + 10X_C + X_D.$$

The mean and variance of D , given standard quality, are

$$\begin{aligned} E(D|S) &= 100(n\lambda_A) + 50(n\lambda_B) + \dots \\ &= n[100\lambda_A + 50\lambda_B + 10\lambda_C + \lambda_D] \\ &= nU_s \\ V(D|S) &= (100)^2(n\lambda_A) + (50)^2(n\lambda_B) + \dots \\ &= n[(100)^2\lambda_A + (50)^2\lambda_B + (10)^2\lambda_C + \lambda_D] \\ &= nC_s. \end{aligned}$$

The notation " $D|S$ " reads " D conditional on S ."

Note that U_s is the demerit per unit standard and C_s is a variance per unit standard. These are the numbers that would be published in the official list of standards called the Master Reference list. The quality standards are not affected by QMP.

2.5 Rating classes and periods

For the purpose of reporting quality results to management, the products are grouped into rating classes. An example is: ESS No. 1

wired equipment, functional test, at Dallas.* The results of all the inspections associated with this rating class are aggregated over a time period called a *rating period*. A rating period is about six weeks long and there are eight rating periods per year. QMP does not affect the rating classes or periods.

2.6 The *T*-rate

The advantage of having quality standards is that observed quality results can be statistically compared to the standards. In the *T*-rate system, this is done with a statistic called the *T*-rate.

For a given rating class, let Q denote the total number of defects, defectives, or demerits that are observed in all the inspections conducted on all the subproducts during a rating period. Because there are quality standards for each set of inspections on each product subclass, it is possible to compute the standard mean and variance of Q , denoted by $E(Q|S)$, $V(Q|S)$. The *T*-rate is

$$T\text{-rate} = \frac{E(Q|S) - Q}{\sqrt{V(Q|S)}}.$$

It measures the difference between the observed result and its standard in units of statistical standard deviation.

For each rating period, the *T*-rate is plotted in the control chart format shown in Fig. 1b. The control limits of ± 2 are reasonable under the assumption that Q has an approximate normal distribution. Then the standard distribution of Q is the "standard normal," and excursions outside the control limits are rare under standard quality. For large audit sample sizes, this approximation follows from the central limit theorem. As we shall see, the approximation is poor for small sample sizes.

2.7 Reports, Below Normals and ALERTs

The fundamental reports to *WE* management are books of *T*-rate control charts for all rating classes. However, every rating period, a summary booklet is prepared. The summary consists of various aggregate quality performance indices and an exception report which lists rating classes that are having quality problems.

There are two kinds of exceptions: Below Normal (BN) and ALERT. These are based on statistical tests of the hypothesis that quality is at standard. The rules for BN and ALERT are based on six consecutive *T*-rates, t_1, \dots, t_6 , where t_6 is the current *T*-rate. The rules use the

* Technically, this is called a scoring class in quality assurance documentation. Here, rating class means scoring class.

following runs criteria:

SCAN(S): $t_1 < 0, \dots, t_6 < 0,$

341 (T): $t_6 < -1$ and at least two of the set $\{t_3, t_4, t_5\}$ are less than $-1.$

Finally, the rules for BN and ALERT are:

Below Normal (BN): One of the following two conditions is satisfied:

(1) $t_6 < -3$

(2) $-3 \leq t_6 < -2$ and at least one of the following three conditions hold:

(i) SCAN

(ii) 341

(iii) At least one of the set $\{t_2, t_3, t_4, t_5\}$ is less than $-2.$

ALERT: SCAN or 341 but not BN.

In Fig. 1b, examples of BN are 7806 and 7803. Examples of ALERT are 7808 and 7804.

Both the fundamental report formats and the rules for BN and ALERT are different under QMP

2.8 Pros and cons of the T-rate

The advantage of the T-rate is its simplicity. It can be calculated manually. Exceptions can be identified by inspection. The fact that the T-rate has been used for so long is a testimonial to its advantages.

However, the T-rate does have problems.* The T-rate does not measure quality. A T-rate of -6 does not mean that quality is twice as bad as when the t-rate is $-3.$ The T-rate is only a measure of statistical evidence with respect to the hypothesis of standard quality. This subtle statistical point is often misunderstood by report readers. Years of explanations have not cleared up the confusion.

Another problem is that the ALERT (SCAN and 341) rules are tests of hypothesis about quality trends, not current quality. Consider Fig. 2. You can assert that quality was probably substandard sometime between 7702 and 7707. You cannot, however, assert that quality is substandard in 7707. The QMP result for 7707 is normal.

In addition, the rules for ALERT and to some extent BN depend on attributes of past T-rates rather than their exact values. For example, five consecutive past T-rates at -1.0 are treated exactly like five consecutive T-rates at $-0.1.$ This was done for statistical robustness. But statistical information is lost. There are no "outliers" in the audit data. Defects assessed were in the product. Many defects assessed

* Although the foundation of the T-rate system was laid by Shewhart⁴ and Dodge,⁵ the details are the results of contributions by many people over 50 years.

mean substandard quality at the time of assessment. It is possible that very unusual circumstances caused the defects. But it is intended that the audit flag such unusual circumstances.

The significance level of the T -rate hypothesis test depends on sample size and can be very large. Suppose that we have a simple test defect audit with a sample size of 32 units and a standard of 0.005 defects per unit. The expected number of defects is $(32)(0.005) = 0.16$. For one defect observed, the T -rate is

$$t = \frac{0.16 - 1}{\sqrt{0.16}} = -2.1.$$

So, every time there is a defect, the T -rate exceeds the control limit.

Now, assuming standard quality, the number of defects has a Poisson distribution with mean = 0.16. The Poisson probability of one or more defects is 0.15. So even when the standard is being met, there is a 15 percent chance of the T -rate dropping below -2.0 . In statistical terms, we have a biased test (i.e., there is no reasonable upper bound on the significance level).

Clearly, it is not reasonable to take action every time the audit finds a defect. So special rules called modification treatments have evolved to handle cases like the one just described. Some of these modification treatments are statistically sound, others are not. This detracts from the desired objectivity of our quality rating.

In a sense, QMP is orthogonal to the T -rate. On the one hand, QMP cannot easily be computed manually. On the other hand, QMP does not have any of the disadvantages described above for the T -rate. The basic message of the QMP box chart (Fig. 1a) is unambiguous and exceptions can be identified by inspection.

III. OVERVIEW OF QMP

As described in the introduction, QMP is the new way of conducting three of the audit ingredients: defect assessment, quality rating, and quality reporting. This section contains an overview of QMP. Mathematical derivations and detailed analyses of QMP are left for later sections.

3.1 Defect assessment practices

Defect assessment practices have two parts. Part one is a description of those situations where fewer defects are assessed than are found. Part two is a formula for the number of defects assessed.

In QMP, the principle for part one is: Normally all defects found in the quality assurance audits are assessed. Occasionally a cluster of two or more defects is found for which the seriousness of the cluster is less than the seriousness implied by individually assessing every defect in

the cluster. Such a cluster shall be called *reducible*. Seriousness is measured from the customer's point of view. The audit attempts to measure seriousness as if the auditor is the customer. So if defects are found and corrected as a result of the audit, *no* adjustment in assessment is necessary. More specifically, a reducible cluster is a collection (on one audited unit) of

- [1] dependent identical defects that the customer will
- [2] almost surely discover in its entirety when a small part of the cluster is discovered and
- [3] will correct or otherwise account for en masse, so that
- [4] total seriousness is better represented by assessing d_a defects (computed by the assessment formula), rather than the number found.

In [1], we use the word dependent in a statistical sense. Defects are dependent if they occur in a short interval of time and are systematically introduced by a common feature of the production process.

Ideally, the assessment associated with a reducible cluster of defects should depend on the situation. Over time, lists of reducible clusters and their assessments could be catalogued and added to the demerit lists. But, for now, there is no list of reducible clusters, so an assessment formula is needed.

For QMP, the assessment formula has the general form

$$d_a = AN + 1,$$

where AN stands for "Allowance Number." In turn, AN has the general form $AN = e + 3\sqrt{e}$, rounded down to an integer or to the closest integer, where e can be interpreted as an expected number of defects. The computation of e and the rounding depend on the audit. For some audits, tradition has prevailed, and for other audits, methods of computing e were developed for QMP.

As an example, consider a single relay for which three contacts are defective (class B defect). The traditional method of computing e for an apparatus audit is

$$e = (12)(0.005) = 0.06,$$

where 12 is the number of contacts in the relay and 0.005 is a traditional generic standard per unit of count for class-B defects. The traditional rounding is down, so $AN = (0.06) + 3\sqrt{0.06} = 0.79$, rounded down to zero. Hence, one class-B defect is assessed.

Another example is a reducible cluster of loose terminations found on a bay of equipment in a transmission installation performance audit. In this case, e is just the quality standard for the bay in defects per unit, and the rounding is to the nearest integer.

As you have gathered by now, defect assessment is an art not a science. The principles and rules described here have empirical validity. In practice, they usually lead to reasonable assessments.

3.2 Equivalent defects and expectancy

A complicating factor in the analyses of audit results is that defects, defectives, and demerits are different. But, are they really? The answer is no; because, for statistical purposes, they can all be transformed into equivalent defects that have approximate Poisson distributions.

Suppose we have a quality measure Q (total defects, defectives, or demerits). Let E_s and V_s denote the standard mean (called expectancy) and variance of Q . So the T -rate is $T = (E_s - Q)\sqrt{V_s}$.

Now define

$$X = \text{equivalent defects} = \frac{Q}{V_s/E_s}$$

and

$$\begin{aligned} e &= \text{equivalent expectancy} = \text{standard mean of } X \\ &= \frac{E_s}{V_s/E_s} = \frac{E_s^2}{V_s}. \end{aligned}$$

If all defects have Poisson distributions and are occurring at θ times the standard rate, then it can be shown that

$$E[X | \theta] = V(X | \theta) = e\theta;$$

hence, X has an approximate Poisson distribution with mean $e\theta$.

As an example, consider the demerits case. The total number of demerits has the general form

$$D = \sum w_i X_i,$$

where the w_i 's are known weights and the X_i 's have Poisson distributions. Assume that the mean of X_i is $e_i\theta$, where e_i is the standard mean of X_i and θ is the population quality expressed on an index scale. So $\theta = 2$ means that all types of defects are occurring at twice the rate expected.

The mean and variance of D are

$$\begin{aligned} E(D) &= \sum w_i E(X_i) \\ &= \sum w_i (e_i \theta) \\ &= \theta E_s \end{aligned}$$

and

$$\begin{aligned}V(D) &= \sum w_i^2 V(X_i) \\ &= \sum w_i^2 (e_i \theta) \\ &= \theta V_s,\end{aligned}$$

where E_s and V_s are the standard mean and variance, respectively, of D .

The mean and variance of equivalent defects, X , are

$$\begin{aligned}E(X) &= \frac{E(D)}{V_s/E_s} \\ &= \frac{\theta E_s}{V_s/E_s} \\ &= \theta e\end{aligned}$$

and

$$\begin{aligned}V(X) &= \frac{V(D)}{[V_s/E_s]^2} \\ &= \frac{\theta V_s E_s^2}{V_s^2} \\ &= \theta e.\end{aligned}$$

The mean and variance of X are equal; so, X has an approximate Poisson distribution with mean $e\theta$. Of course, it is not exact; because, X is not always integer valued. But, this Poisson approximation for equivalent defects is better than the normal approximation implied by the T -rate system. It is the Poisson approximation in QMP that obviates the need for the modification treatments discussed in Section 2.8.

A similar analysis works approximately for the defectives case. So, any aggregate of demerits, defectives, or defects can be transformed into equivalent defects. Just use the standard expectancy and variance as illustrated above for demerits.

3.3 Statistical foundations of QMP

The algorithm used for computing the QMP box chart shown in Fig. 1a was derived from a Bayesian analysis of a particular statistical model. In this Section we describe the model and put it in perspective. This will provide an appreciation for how the box charts can be interpreted and why they are a useful management tool.

3.3.1 QMP model

For rating period t , let x_t = equivalent defects in the audit sample, e_t = equivalent expectancy of the audit sample, θ_t = population index,

as defined in Section 3.2. Based on the discussion in Section 3.2, we assume that the conditional distribution of x_t given θ_t is Poisson with mean $e_t\theta_t$; i.e.,

$$x_t | \theta_t \sim \text{Poisson}(e_t\theta_t).$$

In Fig. 3 we see that the season average varies from year to year. Some of that variation is due to the fact that the season is itself a sample from a conceptual infinite population of at bats. The rest of the variation is due to changes in ability, competition, etc., that are caused by numerous factors that may or may not be identifiable. The important concept is that the time series of season averages is a stochastic process. For QMP we assume that the time series of θ_t 's is an unknown stochastic process.

For reasons that are partly statistical and partly administrative, we have decided to restrict our use of past data to five periods. The main administrative reason is that the T -rate system used the past five periods. So all of the T -rate administrative rules that dealt with missing data and reinitialization of rating classes can be used in QMP. Statistically, QMP works as well for six periods as it does for eight periods (one year).

A consequence of using only six periods of data is that no useful inference can be made about possible complex structure in the stochastic process of θ_t 's. So we assume simply that the θ_t 's are a random sample from an unknown distribution called the *process distribution*. Furthermore, six observations are not enough to make fine inferences about the family of this unknown distribution. So for mathematical simplicity we assume it to be a gamma distribution with unknown mean = θ and variance = γ^2 (Appendix A); i.e.,

$$\theta_t \stackrel{\text{iid}}{\sim} \text{Gamma}\left(\frac{\theta^2}{\gamma^2}, \frac{\gamma^2}{\theta}\right),$$

$$t = 1, \dots, T(\text{current period}).$$

The parameters θ^2/γ^2 and γ^2/θ are the shape and scale parameters of the gamma distribution. We use the names

$$\theta = \text{process average,}$$

$$\gamma^2 = \text{process variance.}$$

This choice of a unimodal distribution reflects our experience that usually many independent factors affect quality; so there is a central limit theorem effect.

We are assuming that the process average is unknown but fixed. In reality, it may be changing. We handle this by using a moving window of six periods of data. But this treats the past data symmetrically. An

alternative would be some kind of exponential smoothing or Kalman filtering. My colleague M. S. Phadke is developing a generalization to QMP based on a random walk model for the process average.

The model so far is an empirical Bayes model.¹⁰ The parameter of interest is the current population index, θ_T , which has a distribution called the process distribution. Bayesians would call it the prior distribution if it were known. But we must use all the data to make an inference about the unknown process distribution. So, the model is called empirical Bayes.

Efron and Morris¹⁰ take a classical approach to the empirical Bayes model. They use classical methods of inference for the unknown process distribution. QMP is based on a Bayesian approach to the empirical Bayes model. Each product has its own process mean and variance. These vary from product to product. By analyzing many products, we can model this variation by a prior distribution for (θ, γ^2) .

Summarizing, our model is

$$x_t | \theta_t \sim \text{Poisson}(e_t \theta_t), \quad t = 1, \dots, T,$$

$$\theta_t \stackrel{\text{iid}}{\sim} \text{Gamma}\left(\frac{\theta^2}{\gamma^2}, \frac{\gamma^2}{\theta}\right), (\theta, \gamma^2) \sim \text{prior distribution } \rho(\theta, \gamma^2).$$

For now, $\rho(\theta, \gamma^2)$ remains general.

This is a full Bayesian model. It specifies the joint distribution of all variables. The quality rating in QMP is based on the posterior distribution of θ_T given $\mathbf{x} = (x_1, \dots, x_T)$.

3.3.2 The model in perspective

Quality rating in QMP is based on posterior probabilities given the audit data. Of course these probabilities depend on the model. But how do we know the model is right?

It is important to understand that we are not doing data analysis with QMP. In data analysis, each set of data is treated uniquely. Probabilities cannot be computed. Objective decisions cannot be made.

A requirement of quality rating is a specific rule that defines quality exceptions and a figure of merit (e.g., a probability) associated with an exception. A statistical model provides both. QMP could have been based on a more elaborate model. Our model represents a compromise between simplicity and believability.

So our exception decisions are at least consistent with one simple model of reality. The probabilities are conditional on that model. Otherwise, they can only be interpreted as figures of merit.

We have imbedded the simple hypothesis of a Poisson distribution with a standard mean into a class of alternatives. The alternatives are

Poisson distributions with nonstandard means. Much more complicated alternatives can be included: e.g., the class of negative binomial distributions, and our probabilities would change a little. But QMP has achieved a kind of empirical validity. The exceptions being identified are accepted by the managers being rated. And for the products declared normal, there is a model (i.e., our model) that affords the standard hypothesis some credence.

3.3.3 Posterior distribution of current quality

We show in Section IV that it is computationally impractical to derive the exact posterior distribution of θ_T . The best we can do is approximate the posterior mean and variance of θ_T .

The posterior mean and variance of θ_T are derived in Section IV. The posterior mean is

$$\begin{aligned}\hat{\theta}_T &= E(\theta_T | \mathbf{x}) \\ &\doteq \hat{\omega}_T \hat{\theta} + (1 - \hat{\omega}_T) I_T,\end{aligned}$$

where

$$\begin{aligned}\hat{\theta} &\doteq E(\theta | \mathbf{x}), \\ \hat{\omega}_T &\doteq E(\omega_T | \mathbf{x}), \\ \omega_T &= \frac{\theta/e_T}{\theta/e_T + \gamma^2}.\end{aligned}$$

The posterior mean, $\hat{\theta}_T$, is a weighted average between the estimated process average, $\hat{\theta}$, and the defect index, I_T , of the current sample. It is the dynamics of the weight, ω_T , that makes the Bayes estimate work so well. For any t , the sampling variance of I_t is

$$\begin{aligned}V(I_t | \theta_t) &= V\left(\frac{x_t}{e_t} \middle| \theta_t\right) \\ &= \frac{1}{e_t^2} V(x_t | \theta_t) \\ &= \frac{1}{e_t^2} (e_t \theta_t) \\ &= \theta_t / e_t.\end{aligned}$$

The expected value of this is

$$E[\theta_t / e_t] = \theta / e_t.$$

So the weight, ω_T , is

$$\frac{[\text{expected sampling variance}]}{[\text{expected sampling variance}] + [\text{process variance}]}$$

If the process is relatively stable, then the process variance is relatively small and the weight is mostly on the process average; but if the process is relatively unstable, then the process variance is relatively large and the weight is mostly on the current sample index. The reverse is true of the sampling variance. If it is relatively large (e.g., small expectancy), then the current data is weak and the weight is mostly on the process average; but if the sampling variance is relatively small (e.g., large expectancy), then the weight is mostly on the current sample index. In other words, ω_T , is a monotonic function of the ratio of expected sampling variance to process variance.

The posterior variance of θ_T is

$$V_T \doteq (1 - \hat{\omega}_T)\hat{\theta}_T/e_T + \hat{\omega}_T^2 V(\theta | \mathbf{x}) + (\hat{\theta} - I_T)^2 V(\omega_T | \mathbf{x}).$$

If the process average and variance were known, then the posterior variance of θ_T would be $(1 - \omega_T)\hat{\theta}_T/e_T$ (Appendix B). So the first term is just an estimate of this. But since the process average and variance are not known, the posterior variance has two additional terms. One contains the posterior variance of the process average and the other contains the posterior variance of the weight.

The first term dominates. A large $\hat{\omega}_T$ (relatively stable process), a small $\hat{\theta}_T$ (good current quality), and a large e_T (large audit) all tend to make the posterior variance of θ_T small (the box chart short).

If $\hat{\omega}_T$ is small, then the second term is negligible. This is because the past data is not used much, so the uncertainty about the process average is irrelevant.

If the current sample index is far from the process average, then the third term can be important. This is because outlying observations add to our uncertainty as to what is happening.

If the process average and variance were known, then the posterior distribution would be gamma (Appendix B). So we approximate the posterior distribution with a gamma fitted by the method of moments. The parameters of the fitted gamma are

$$\begin{aligned} \alpha &= \text{shape parameter} \\ &= \hat{\theta}_T^2 / V_T, \\ \tau &= \text{scale parameter} \\ &= V_T / \hat{\theta}_T, \end{aligned}$$

and the posterior cumulative distribution function is

$$\Pr\{\theta_T \leq y | \mathbf{x}\} = G_\alpha(y/\tau)$$

(Appendix A).

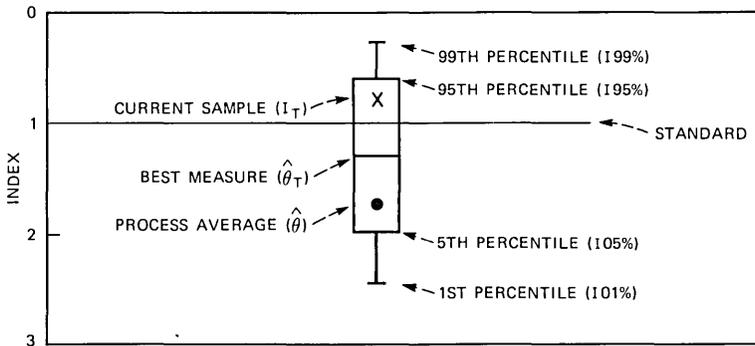


Fig. 4—QMP box and whisker chart. This is a graphical representation of the posterior distribution for current production given the six most recent periods of audit data. The whiskers display the 99th and 1st percentiles and the box displays the 95th and 5th percentiles. The Best Measure is the posterior mean or Bayes estimate. It is a weighted average of the process average (“dot”) and the current sample (“X”). The weight is the ratio of sampling variance to total variance. If all the variance is due to sampling, then the production is stable and the process average is the Best Measure of current quality. If the sampling variance is zero, then the current sample is the Best Measure.

3.4 QMP reports

3.4.1 QMP box chart

The QMP box and whisker chart is shown in Fig. 4. $I99\%$, \dots , $I01\%$ are defined by

$$\begin{aligned}
 1 - G_{\alpha}(I99\%/\tau) &= 0.99, \\
 &\vdots \\
 1 - G_{\alpha}(I01\%/\tau) &= 0.01.
 \end{aligned}$$

So, e.g., *a posteriori*, there is a 99 percent chance that θ_T is larger than $I99\%$.

3.4.2 QMP Below Normal and ALERT definitions

In QMP, a rating class is Below Normal (BN) if

$$I99\% > 1;$$

i.e., the posterior probability that the product is substandard exceeds 0.99. Substandard means $\theta_T > 1$. A rating class is on ALERT if

$$I99\% \leq 1 < I95\%;$$

i.e., the posterior probability that the product is substandard exceeds 0.95 but not 0.99.

These definitions are illustrated graphically in Fig. 5, which is oriented like the location summary in Fig. 6.

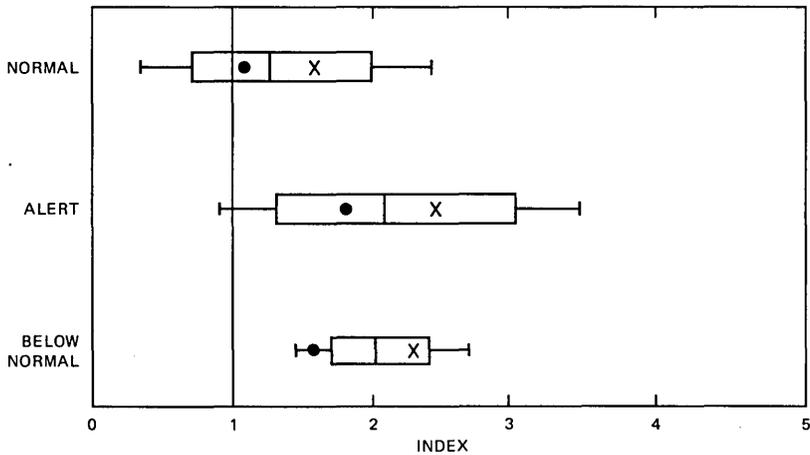


Fig. 5—QMP exceptions. Below Normal means that the probability of substandard quality exceeds 0.99. For ALERT, the probability exceeds 0.95 but not 0.99. Normal is not a quality exception.

3.4.3 QMP report formats

There are two report formats for QMP results. One is a time series of box charts illustrated in Fig. 1a. The estimated process averages are joined. The other is a location summary for the current rating period. This is illustrated in Fig. 6. It orders the rating classes by Best Measure for the current period. Another ordering that will be used is by rating class name.

Western Electric, Bell Laboratories, and American Telephone and Telegraph management will receive all QMP results. Operating company management will receive QMP results on those rating classes that are of direct interest to them. Examples of results provided to the operating companies are the quality of repaired telephone sets and installed switching systems.

3.5 Advantages of QMP

Many of the advantages of QMP relate to the disadvantages of the *T*-rate system (Section 2.8). QMP provides a direct measure of quality. If a rating class is Below Normal, one can tell how bad the quality is. QMP uses past and current data to make an inference about current quality not past quality. If a rating class is on ALERT, then it is over 95 percent probable that there is a quality problem *now*. QMP does not use runs criteria, but uses the actual equivalent defects observed. This provides more statistical efficiency and therefore shorter interval estimates. QMP is robust against statistical "jitter." It does not over-react to a few defects. Consequently, there is no need for special

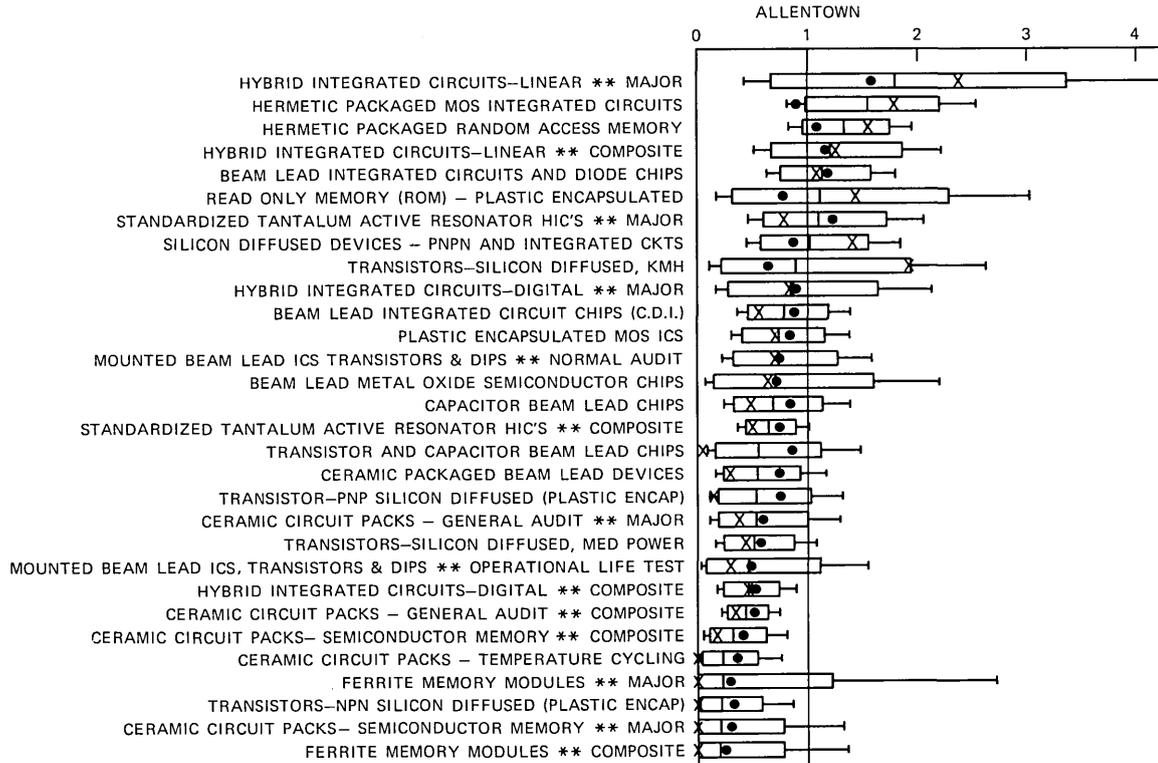


Fig. 6—QMP location summary.

modification treatments. This way we retain statistical objectivity conditional on our model.

Another advantage compared to the T -rate is that QMP provides a lower producer's risk and consequently a more accurate list of exceptions. This is supported by data presented in Section VI.

Finally, QMP will allow us to unify our reporting to Bell System management. In the past, the T -rate statistic did not meet the needs of the operating companies; so, we developed a collection of special reports for the operating companies. Since the QMP report format does meet operating company needs, the relevant subset of all the results is a useful report.

IV. MATHEMATICAL DERIVATION OF QMP

4.1 Exact solution

We are interested in the posterior distribution of θ_T given \mathbf{x} , for the model described in Section 3.3.1. Now $\Pr\{\theta_T \leq y | \mathbf{x}\} = \int_0^\infty \int_0^\infty \Pr\{\theta_T \leq y | \theta, \gamma^2, x_T\} \rho(\theta, \gamma^2 | \mathbf{x}) d\theta d\gamma^2$, where $\rho(\theta, \gamma^2 | \mathbf{x})$ is the posterior distribution of θ, γ^2 given \mathbf{x} .

From Appendix B, we know that the distribution of θ_T given θ, γ^2 , and x_T is gamma; so, $\Pr\{\theta_T \leq y | \theta, \gamma^2, x_T\}$ can be expressed in terms of an incomplete gamma function.

By Bayes theorem,

$$\rho(\theta, \gamma^2 | \mathbf{x}) = \frac{\rho(\theta, \gamma^2) L(\theta, \gamma^2)}{\int_0^\infty \int_0^\infty \rho(\theta, \gamma^2) L(\theta, \gamma^2) d\theta d\gamma^2},$$

where $\rho(\theta, \gamma^2)$ is the prior density of θ, γ^2 and $L(\theta, \gamma^2)$ is the likelihood function. Since x_t given θ_t is Poisson and θ_t given θ, γ^2 is gamma, it follows that x_t given θ, γ^2 is negative binomial; hence,

$$L(\theta, \gamma^2) = \prod_{t=1}^T L_t(\theta, \gamma^2),$$

where

$$L_t(\theta, \gamma^2) = \frac{\Gamma(x_t + \theta^2/\gamma^2)}{x_t! \Gamma(\theta^2/\gamma^2)} [1 - \omega_t(\theta, \gamma^2)]^{x_t} [\omega_t(\theta, \gamma^2)]^{\theta^2/\gamma^2},$$

$$\omega_t(\theta, \gamma^2) = \frac{\theta/e_t}{\theta/e_t + \gamma^2}.$$

So the posterior distribution of θ_T is a complex triple integral that has to be inverted to compute the QMP box chart. The posterior mean and variance of θ_T can be expressed in terms of several double integrals. There are more than 1,000 rating classes that have to be analyzed each

period, so computational efficiency is important. This is why we have developed an efficient heuristic solution to the problem.

4.2 Empirical priors for process parameters

It is clear from Section 4.1 that prior distributions for θ and γ^2 are needed. In the fourth rating period of 1979, we applied an earlier version of the QMP algorithm to over 1,000 rating classes. This provided over 1,000 estimates of θ and γ^2 and empirical distributions of these estimates. The empirical mean and variance of the θ estimates were 0.75 and 0.17, respectively. The empirical mean, variance, and mode of the γ^2 estimates were 0.28, 0.19, and 0.05, respectively.

In the remainder of Section IV, we use 1 as a mean value of θ instead of 0.75. This is because 1 is the desired standard value that minimizes first cost plus maintenance costs. Under QMP, the shops will be able to operate on the average closer to 1, because the producer's risk (see Section 6.2) is smaller than for the T -rate. Also, more defects are assessed under QMP than for the T -rate (see Sections 2.2 and 3.1).

4.3 Posterior mean of current quality

For the model described in Section 3.3.1,

$$\begin{aligned}\hat{\theta}_T &= E(\theta_T | \mathbf{x}) \\ &= E[E(\theta_T | \theta, \gamma^2, \mathbf{x}) | \mathbf{x}].\end{aligned}$$

Conditioning on θ and γ^2 means that the process distribution is known. So by Theorem B.1 in Appendix B,

$$\hat{\theta}_T = E[\omega_T \theta + (1 - \omega_T) I_T | \mathbf{x}]. \quad (1)$$

To calculate this posterior expectation exactly requires a double integral. But a posterior expectation, $E[\cdot | \mathbf{x}]$, can be viewed as an estimate of the operand “ \cdot ”, because it is the Bayes estimate. So all we need are estimates of ω_T and θ .

4.3.1 Moment estimates of process parameters

As argued in Section 4.1, given θ and γ^2 , x_t has a negative binomial distribution. We show in Appendix D, eqs. (56) and (58), that

$$\begin{aligned}E(I_t) &= \theta, \\ E(Y_t) &= \gamma^2,\end{aligned}$$

where

$$Y_t = (I_t - \theta)^2 - I_t/e_t. \quad (2)$$

So we have many independent estimates of θ and γ^2 . A general method of combining independent estimates of parameters is a

weighted average, where the weights are proportional to the reciprocal of the variances of the individual independent estimates. Such estimates of θ and γ^2 are

$$\bar{\theta} = \sum_t p_t I_t, \quad (3)$$

$$\bar{\gamma}^2 = \sum_t q_t Y_t, \quad (4)$$

where

$$\sum_t p_t = \sum_t q_t = 1$$

and

$$p_t \propto 1/V(I_t),$$

$$q_t \propto 1/V(Y_t).$$

Notice that Y_t depends on θ . So in the application, we replace θ by an estimate.

Now $V(I_t)$ and $V(Y_t)$ depend on the unknown parameters θ and γ^2 . The important consideration in setting the weights p_t and q_t is their general behavior as e_t varies. So for simplicity (to avoid iteration), we choose $\theta = 1$ and $\gamma^2 = 1/4$, which are empirically-determined mean values of these process parameters (see Section 4.2).

In Appendix D, we derive formulas for $V(I_t)$ and $V(Y_t)$. Plugging $\theta = 1$ and $\gamma^2 = 1/4$ into eqs. (56) and (59) yields

$$p_t \propto f_t = \frac{1}{V(I_t)} = 1 / \left(\frac{1}{e_t} + \frac{1}{4} \right) = \frac{e_t}{1 + e_t/4}, \quad (5)$$

$$q_t \propto g_t = \frac{1}{V(Y_t)} = 1 / \left[\frac{2.5}{e_t^2} + \frac{1.5}{e_t} + 0.22 \right] \\ = \frac{e_t^2}{2.5 + 1.5e_t + (0.22)e_t^2}. \quad (6)$$

Note that for small e_t , $f_t \propto e_t$; but for large e_t , the f_t 's and therefore the weights, p_t 's, are all about equal. This is because for any large e_t , $I_t \doteq \theta$, and we are trying to estimate the average of the θ_t 's.

4.3.2 Bayes estimate of the process average

In the case $I_t \equiv 0$ for all t , there is a problem with the estimate $\bar{\theta}$. If we plug $\bar{\theta} = 0$ into (1), then $\hat{\theta}_T = 0$. But $\hat{\theta}_T$ is a posterior mean of a positive parameter, so it cannot be zero. The correct method of handling this problem is to start with a proper prior distribution on the process average, θ . But then the mathematics and the computations become complicated.

So we assert that we have prior information that is equivalent to observing some “prior data,” x_0 and e_0 . Then a Bayes type estimate has the form

$$\hat{\theta} = \sum_{t=0}^T p_t I_t, \quad (7)$$

which has the same form as the moment estimate, $\bar{\theta}$, but uses all the data including the “prior data.”

To choose values for x_0 and e_0 , consider $T = 1$. A generic form of a Bayes type estimate of θ is

$$wE(\theta) + (1 - w)I_1,$$

where

$$w = \frac{V(I_1)}{V(I_1) + V(\theta)}.$$

Setting this generic form equal to (7) yields

$$E(\theta) = x_0/e_0,$$

$$V(\theta) = 1/e_0 + 1/4.$$

From Section 4.2, $E(\theta) = 1$; and we conservatively choose $V(\theta)$ to be 1.25 (we do not want our prior observations of θ estimates to preclude large future values of θ). This implies $x_0 = e_0 = 1$.

4.3.3 Bayes estimate of weight

Now define an estimate of γ^2 analogous to (7),

$$\hat{\gamma}_1^2 = \sum_{t=0}^T q_t (I_t - \hat{\theta})^2 - \sum_{t=0}^T q_t (I_t/e_t),$$

as suggested by eqs. (2) and (4).^{*} The first term is the total variance about the process average and the second term is a weighted average of estimated sampling variances. [Recall from Section 3.3.3 that $V(I_t | \theta_t) = \theta_t/e_t$, which can be estimated by I_t/e_t .] We denote the average sampling variance by

$$\sigma^2 = \sum_{t=0}^T q_t (I_t/e_t). \quad (8)$$

The problem with $\hat{\gamma}_1^2$ as an estimate of process variance is that it can be negative. To solve this problem, we use the results in Appendix C. Assume σ^2 is a known constant, and define the unknown weight as

$$\omega = \frac{\sigma^2}{\sigma^2 + \gamma^2}.$$

^{*} Note that we treat the “prior data” as real data.

To apply Appendix C, we must find a statistic, SS , and a degree of freedom, df , so that, approximately,

$$\frac{\omega}{\sigma^2} (SS) \sim \chi_{df}^2.$$

Originally, we just assumed approximate normality of I_t and took $SS = (df + 1) \sum_{t=0}^T q_t (I_t - \hat{\theta})^2$ and $df = T$. But we found unusual sets of data for which the number of defects allowed (before declaration of Below Normal) was a decreasing function of expectancy in short ranges of small expectancy. We dubbed this the "QMP wiggle."

To solve this problem, we approximate the sampling distribution of

$$Z = \sum_{t=0}^T q_t (I_t - \hat{\theta})^2$$

by a scaled chi-square with degrees of freedom deduced by the method of moments.

Let

$$Z_1 = \sum_{t=0}^T q_t (I_t - \theta)^2,$$

and try an SS of the form

$$uZ,$$

where u is an unknown constant. The two moment equations that have to be satisfied are

$$E \left[\frac{\omega}{\sigma^2} (uZ) \right] = df, \tag{9}$$

$$\frac{E^2[(\omega/\sigma^2)(uZ)]}{V[(\omega/\sigma^2)(uZ)]} = \frac{E^2[\chi_{df}^2]}{V[\chi_{df}^2]}.$$

And the second equation is

$$\frac{2E^2[Z]}{V[Z]} = df. \tag{10}$$

Inspired by well-known normal theory, we use the approximations

$$E(Z) = \left(\frac{df}{df + 1} \right) E(Z_1), \tag{11}$$

$$\frac{2E^2(Z)}{V(Z)} = - \frac{2E^2(Z_1)}{V(Z_1)} - 1. \tag{12}$$

Now using eqs. (11) and (56), eq. (9) becomes

$$E \left[\frac{\omega}{\sigma^2} (uZ) \right] = \frac{\omega}{\sigma^2} u \left(\frac{df}{df + 1} \right) \sum q_t (\gamma^2 + \theta/e_t)$$

$$\begin{aligned}
& \doteq \frac{\omega}{\sigma^2} u\left(\frac{df}{df+1}\right) (\gamma^2 + \sigma^2) \\
& = u\left(\frac{df}{df+1}\right) \\
& = df;
\end{aligned}$$

hence,

$$u = df + 1.$$

As for eq. (12), the mean and variance of Z_1 depend on θ and γ^2 . So to avoid iteration, we now select $\theta = 1$ and $\gamma^2 = 0$,* which were empirically determined in Section 4.2. Then by eqs. (12), (56), and (57),

$$\begin{aligned}
\frac{2E^2[Z]}{V[Z]} &= \frac{2[\sum q_i(1/e_i)]^2}{\sum q_i^2(1/e_i^3 + 2/e_i^2)} - 1 \\
&= df.
\end{aligned} \tag{13}$$

So our statistic ss is

$$ss = (df + 1) \sum q_i(I_i - \hat{\theta})^2,$$

where df is given by eq. (13).

Now apply the Corollary to Theorem C.1 in Appendix C to get

$$\omega | ss \sim C - \text{Gamma}\left(a, \frac{\sigma^2}{b}, 1\right),$$

where

$$a = a_0 + \frac{df}{2}, \quad b = b_0 + \frac{ss}{2}. \tag{14}$$

Define

$$\begin{aligned}
S^2 &= b/a \\
&= \frac{2b_0 + (df + 1) \sum_{t=0}^T q_t(I_t - \hat{\theta})^2}{2a_0 + df},
\end{aligned} \tag{15}$$

$$R = S^2/\sigma^2. \tag{16}$$

Now apply Theorem C.2 in Appendix C to get

$$E(\omega | \mathbf{x}) = \frac{1}{RF} = \frac{\sigma^2}{FS^2}, \tag{17}$$

* The choice of $\gamma^2 = 0$ here may seem inconsistent with choice of $\gamma^2 = 1/4$ in the definition of f_i and g_i . There it was necessary to take a positive value (the empirical mean across products) of γ^2 to get the correct behavior for large e_i . Here it was not necessary, so for simplicity, we took the approximate empirical mode across products, $\gamma^2 = 0$.

$$V(\omega | \mathbf{x}) = G. \quad (18)$$

To determine the parameters (a_0 , b_0) of the prior distribution of ω , we first develop an empirical distribution of estimated ω 's across many rating classes, which have a mean and variance of 0.6 and 0.03, respectively. To be conservative, we inflate the variance, shrink the mean, and select the prior mean and variance of ω to be 0.55 and 0.045, respectively. The parameters a_0 and b_0 are then solutions to (see Appendix C)

$$\frac{1}{R_0 F} = 0.55,$$

$$G = 0.045,$$

where F and G are defined in Theorem C.2 in terms of a_0 and $R_0 = b_0/a_0\sigma^2$. A numerical analysis yields

$$a_0 = 4.5,$$

$$R_0 = b_0/a_0\sigma^2 = 1.6,$$

or

$$a_0 = 4.5, \quad (19)$$

$$b_0 = (7.2)\sigma^2. \quad (20)$$

Now we define an estimate, $\hat{\gamma}^2$, of the process variance by

$$E(\omega | \mathbf{x}) = \frac{\sigma^2}{\sigma^2 + \hat{\gamma}^2}.$$

So by eq. (14)

$$\begin{aligned} \hat{\gamma}^2 &= FS^2 - \sigma^2 \\ &= (FR - 1)\sigma^2. \end{aligned} \quad (21)$$

This is our improvement of the moments estimate. The inflation factor F prevents $\hat{\gamma}^2$ from being negative or zero. It can be shown that if R is large, F is approximately one; but if R is small, $FR - 1$ is positive and F is large.

We are now in a position to estimate ω_T by

$$\hat{\omega}_T = \frac{\sigma_T^2}{\sigma_T^2 + \hat{\gamma}^2}, \quad (22)$$

where

$$\sigma_T^2 = \hat{\theta}/e_T, \quad (23)$$

and our approximation to (1) is

$$\hat{\theta}_T = \hat{\omega}_T \hat{\theta} + (1 - \hat{\omega}_T)I_T. \quad (24)$$

4.4 Posterior variance of current quality

For the model described in Section 3.3.1,

$$V_T = V(\theta_T | \mathbf{x}) = E[V(\theta_T | \theta, \gamma^2, \mathbf{x}) | \mathbf{x}] + V[E(\theta_T | \theta, \gamma^2, \mathbf{x}) | \mathbf{x}].$$

Conditioning on θ and γ^2 amounts to the process distribution being known. So by Theorem B.1, in Appendix B,

$$V_T = E[(1 - \omega_T)E(\theta_T | \theta, \gamma^2, \mathbf{x})/e_T | \mathbf{x}] + V[\omega_T(\theta - I_T) | \mathbf{x}].$$

Conditioning on γ^2 in the second term yields

$$\begin{aligned} V_T &= E[(1 - \omega_T)E(\theta_T | \theta, \gamma^2, \mathbf{x})/e_T | \mathbf{x}] \\ &\quad + E[V[\omega_T(\theta - I_T) | \gamma^2, \mathbf{x}] | \mathbf{x}] \\ &\quad + V[E[\omega_T(\theta - I_T) | \gamma^2, \mathbf{x}] | \mathbf{x}], \end{aligned} \quad (25)$$

so the posterior variance has three components.

4.4.1 First component

The first component is approximated by regarding the posterior expectation operator as an estimation operator, and it is

$$(1 - \hat{\omega}_T)\hat{\theta}_T/e_T. \quad (26)$$

4.4.2 Second component

To approximate the second component, we first approximate $V[\omega_T(\theta - I_T) | \gamma^2, \mathbf{x}]$. Since ω_T depends primarily on γ^2 and e_T , we shall consider ω_T a constant. So

$$V[\omega_T(\theta - I_T) | \gamma^2, \mathbf{x}] \doteq \omega_T^2 V(\theta | \gamma^2, \mathbf{x}). \quad (27)$$

We use the approximation

$$V(\theta | \gamma^2, \mathbf{x}) \doteq V(\hat{\theta} | \gamma^2, \theta = \hat{\theta}). \quad (28)$$

Now by eq. (57),

$$\begin{aligned} V(\hat{\theta} | \gamma^2, \theta) &= V(\sum p_i I_i | \gamma^2, \theta) \\ &= \sum p_i^2 V(I_i | \gamma^2, \theta) \\ &= \sum p_i^2 [\gamma^2 + \theta/e_i]. \end{aligned} \quad (29)$$

Plugging eqs. (28) and (29) into eq. (27) yields

$$V[\omega_T(\theta - I_T) | \gamma^2, \mathbf{x}] \doteq \omega_T^2 \sum_{i=0}^T p_i^2 [\gamma^2 + \hat{\theta}/e_i].$$

Again, treating the posterior expectation operator as an estimation operator, we get for the second component of eq. (25)

$$\hat{\omega}_T^2 \sum_{i=0}^T p_i^2 [\hat{\gamma}^2 + \hat{\theta}/e_i]. \quad (30)$$

4.4.3 Third component

For the third component in eq. (25), we first approximate $E[\omega_T(\theta - I_T) | \gamma^2, \mathbf{x}]$ by $\bar{\omega}_T(\hat{\theta} - I_T)$, where

$$\bar{\omega}_T = \frac{\hat{\theta}/e_T}{\hat{\theta}/e_T + \gamma^2}.$$

So the third component in eq. (25) is

$$(\hat{\theta} - I_T)^2 V(\bar{\omega}_T | \mathbf{x}). \quad (31)$$

If we define

$$\begin{aligned} r_T &= \frac{\hat{\theta}/e_T}{\sigma^2}, \\ \omega &= \frac{\sigma^2}{\sigma^2 + \gamma^2}, \end{aligned} \quad (32)$$

then

$$\bar{\omega}_T = \frac{r_T \omega}{(r_T - 1)\omega + 1} = h(\omega).$$

So

$$V(\bar{\omega}_T | \mathbf{x}) \doteq [h'(\hat{\omega})]^2 V(\omega | \mathbf{x}), \quad (33)$$

where

$$\begin{aligned} \hat{\omega} &= \frac{\sigma^2}{\sigma^2 + \hat{\gamma}^2}, \\ h'(\hat{\omega}) &= \frac{r_T}{[(r_T - 1)\hat{\omega} + 1]^2}. \end{aligned} \quad (34)$$

Equations (31), (33), and (15) imply that the third component of eq. (25) is

$$\frac{r_T^2(\hat{\theta} - I_T)^2}{[(r_T - 1)\hat{\omega} + 1]^4} G. \quad (35)$$

Putting eqs. (26), (30), and (35) together implies that the approximate posterior variance of θ_T is

$$\begin{aligned} V_T &= (1 - \hat{\omega}_T)(\hat{\theta}_T/e_T) \\ &+ \hat{\omega}_T^2 \sum_{i=0}^T p_i^2 [\hat{\gamma}^2 + \hat{\theta}/e_i] + \frac{r_T^2(\hat{\theta} - I_T)^2}{[(r_T - 1)\hat{\omega} + 1]^4} G. \end{aligned} \quad (36)$$

4.5 QMP algorithm

Here we summarize the QMP formulas. On the right side of the formulas are the section numbers or equation numbers where the formulas were derived.

The audit data for $t = 1, \dots, T$ is the following:

Q_t = Attribute quality measure in the sample, period t (total defects, defectives, or demerits),

E_{St} = expected value of Q_t given standard quality,

V_{St} = Sampling variance of Q_t given standard quality.

For each period compute the following:

Equivalent defects:

$$x_t = \frac{Q_t}{V_{St}/E_{St}}, \quad (\text{Section 3.2})$$

Equivalent expectancy:

$$e_t = E_{St}^2/V_{St}.$$

For the "prior data" ($t = 0$), let $x_0 = e_0 = 1$.

For $t = 0, \dots, T$, compute the following: (Section 4.3.2)

Sample index:

$$I_t = x_t/e_t,$$

Weighting factors for computing process average and variance:

$$f_t = \frac{e_t}{1 + e_t/4}, \quad (5)$$

$$g_t = \frac{e_t^2}{2.5 + 1.5e_t + (0.22)e_t^2}, \quad (6)$$

Corresponding weights:

$$p_t = f_t/\sum f_t, \quad (5)$$

$$q_t = g_t/\sum g_t. \quad (6)$$

Over all periods $t = 0, \dots, T$ compute the following:

Process average:

$$\hat{\theta} = (\sum p_t I_t) \quad (7)$$

Degrees of freedom:

$$df = \frac{2 [\sum q_t (1/e_t)]^2}{\sum q_t^2 (1/e_t^3 + 2/e_t^2)} - 1, \quad (13)$$

Total observed variance:

$$S^2 = \frac{(14.4)\sigma^2 + (df + 1) \sum q_t (I_t - \hat{\theta})^2}{9 + df}, \quad (15), (19), (20)$$

Estimated average sampling variance:

$$\sigma^2 = \sum q_t(I_t/e_t), \quad (8)$$

Variance ratio:

$$R = S^2/\sigma^2, \quad (16)$$

F and G :

$$a = 4.5 + \frac{df}{2}, \quad (14), (19)$$

$$B = \sum_{i=0}^{\infty} T(i), \quad T(0) = 1,$$

$$T(i) = T(i-1) \left[\frac{aR}{a+i} \right],$$

$$F = \frac{B}{B-1}, \quad (17), (50)$$

$$G = \frac{1}{RF} \left[\left(\frac{a+1}{aR} \right) - (F-1) - \frac{1}{RF} \right], \quad (18), (54)$$

Current sampling variance:

$$\sigma_T^2 = \hat{\theta}/e_T \quad (23)$$

Sampling variance ratio:

$$r_T = \sigma_T^2/\sigma^2 \quad (32)$$

Process variance:

$$\hat{\gamma}^2 = FS^2 - \sigma^2 = (FR - 1)\sigma^2, \quad (21)$$

$$\hat{\omega}_T = \sigma_T^2/(\sigma_T^2 + \hat{\gamma}^2), \quad (22)$$

Weights:

$$\hat{\omega} = \sigma^2/(\sigma^2 + \hat{\gamma}^2) = 1/FR, \quad (34)$$

Best measure of current quality:

$$\hat{\theta}_T = \hat{\omega}_T \hat{\theta} + (1 - \hat{\omega}_T) I_T, \quad (24)$$

Posterior variance of current quality:

$$V_T = (1 - \hat{\omega}_T) (\hat{\theta}_T/e_T) + \hat{\omega}_T^2 \sum p_i^2 \left[\hat{\gamma}^2 + \frac{\hat{\theta}}{e_t} \right] + \frac{r_T^2 (\hat{\theta} - I_T)^2}{[(r_T - 1)\hat{\omega} + 1]^4} G, \quad (36)$$

Box chart percentiles:

$$\alpha = \hat{\theta}_T^2 / V_T, \quad (\text{Section 3.3.3})$$

$$\tau = V_t / \hat{\theta}_T,$$

I99%, *I95%*, *I05%*, *I01%* defined by:

$$1 - G_\alpha(I99\%/\tau) = 0.99, \quad (\text{Section 3.4.1})$$

.

.

.

$$1 - G_\alpha(I01\%/\tau) = 0.01.$$

V. QMP DYNAMICS

The Best Measure and the box chart percentiles are nonlinear functions of all the data, so the dynamic behavior of these results can appear to be complex. But this complex behavior is desirable and can be explained. This section characterizes the fundamental dynamics of QMP by example.

5.1 Dynamics of sudden change

Since QMP is partially based on a long run average, it is natural to be concerned about responsiveness of the box chart to sudden change. If there is a sudden degradation of quality, Quality Assurance would like to detect it. If the producer solves a chronic quality problem, they would like their exceptions to disappear. Figures 7 and 8 illustrate the QMP dynamics of sudden change.

The history data in Fig. 7 is a typical history for a product that is meeting the quality standard. The equivalent expectancy of five is average for a manufacturing audit. The history is plotted on a *T*-rate chart along with six possible values for the current *T*-rate (labeled *A* through *F*). So the current period is anywhere from standard (*T*-rate = 0) to well below standard (Index = 3.24, *T*-rate = -5).

The right side of Fig. 7 shows the six possible current results plotted in QMP box-chart form. The box chart labeled *A* is the result of combining current result *A* with the past five periods. The box chart labeled *F* is the result of combining current result *F* with the same past history.

As you can see, the QMP result becomes ALERT at about *T*-rate = -3 (letter D) and becomes BN at about *T*-rate = -4 (letter E). For the *T*-rate method of rating, you would have a BN at *T*-rate = -3. The good past history has the effect of tempering the result of a *T*-rate = -3.

It is informative to study the relative behavior of the current sample

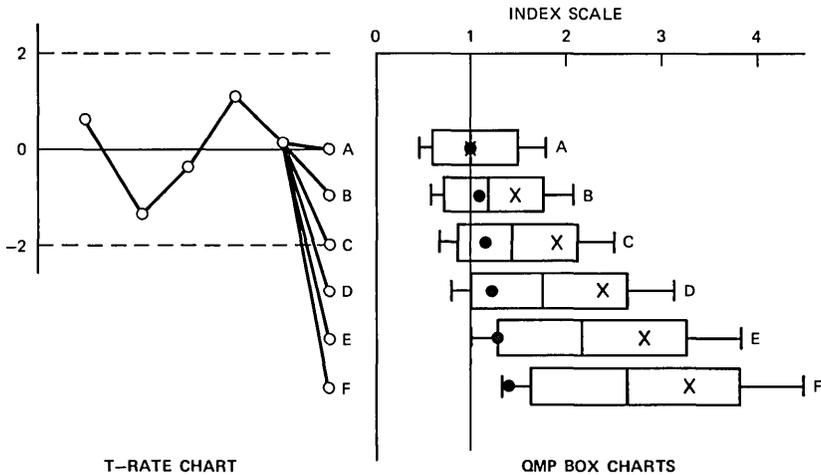


Fig. 7—Dynamics of sudden degradation. The six QMP box charts (labeled A through F) result from the analysis of six time series of data, which all have the same *past* history, but have different current values, as shown in the T-rate chart. A QMP ALERT is triggered at a T-rate of -3 (letter D) and a QMP Below Normal is triggered at a T-rate of -4 (letter E). So a good past history tempers an observed change. Notice that from A to F, the Best Measure swings towards the sample value. This results from increasing evidence of an unstable process (expected number of defects equals 5 for this chart).

index, process average, and Best Measure as the current value goes from A to F. The current index changes a lot (from 1.00 to 3.24) and the process average changes a little (from 1.00 to 1.38), both in a linear way. The Best Measure also changes substantially, but in a nonlinear way. It changes slowly at first and then speeds up. This is because the weight is changing from 0.71 to 0.32. The weight changes, because as the data becomes more and more inconsistent with the past the process becomes more and more unstable, while the current sampling variance changes slowly in proportion to the process average.

Figure 8 is the dual of Fig. 7. It illustrates the dynamics of sudden improvement. For the first five periods plotted, the process average is centered on an index of two. Then an improvement takes place and from the sixth period on, the sample index is at the standard value of one.

For the first five periods plotted, the rate is BN four times and ALERT once. In the sixth period there is a sudden improvement and the sample index goes to standard. Immediately, there is a jump in the Best Measure and the rate is no longer BN. Because of the increase in process variance, the weight changes from 0.69 to 0.61, putting more weight on the current good result. The posterior variance stays about the same [$\hat{\theta}_T$ gets smaller but $(1 - \hat{\omega}_T)$ gets larger].

For the next five periods the sample index stays at standard. During

these periods both the process average and the Best Measure gradually move up towards the standard.

5.2 Bogie charts

A Bogie chart is a graphical device for tracking quality assurance audit data during a rating period. Figures 9 and 10 are examples of Bogie charts. The vertical axis is an index scale and the horizontal axis is an equivalent expectancy scale. During the rating period, as the audit sample size builds up, the sample equivalent expectancy increases. So the horizontal axis can also be viewed as a time axis.

The Bogie curves labeled ALERT and BN are plots of the indices in the current sample for which 195% and 199% (the 95th and 99th percentiles) are exactly one, respectively. So the Bogie curves depend on the past history. The past histories associated with Figs. 9 and 10 have average indices of 0.92 and 4.89, respectively. The variance of the past histories were 0.69 and 5.36, respectively.

To use the Bogie chart, you plot continuously through the period the sample index as a function of the equivalent expectancy in the sample (see Fig. 9). Anytime this plot falls below the ALERT or BN curve, the rate is ALERT or BN at the plotted equivalent expectancy. Then to bail the rate out, the plotted sample index must get above the Bogie curves before the end of the period. For example, in Fig. 9, if the period had ended at an equivalent expectancy of three, then the rate would be ALERT. If it had ended at an equivalent expectancy of five,

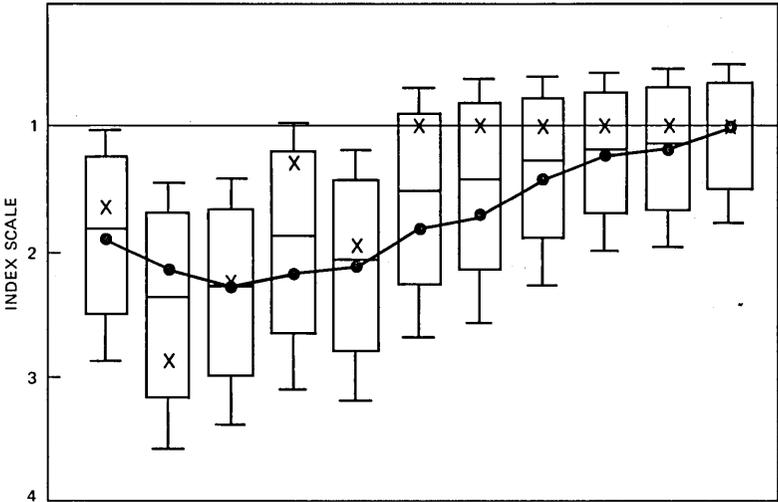


Fig. 8—Dynamics of sudden improvement. As soon as the sample value becomes standard, the product is no longer in the quality exception report (expected number of defects equals 5 for this chart).

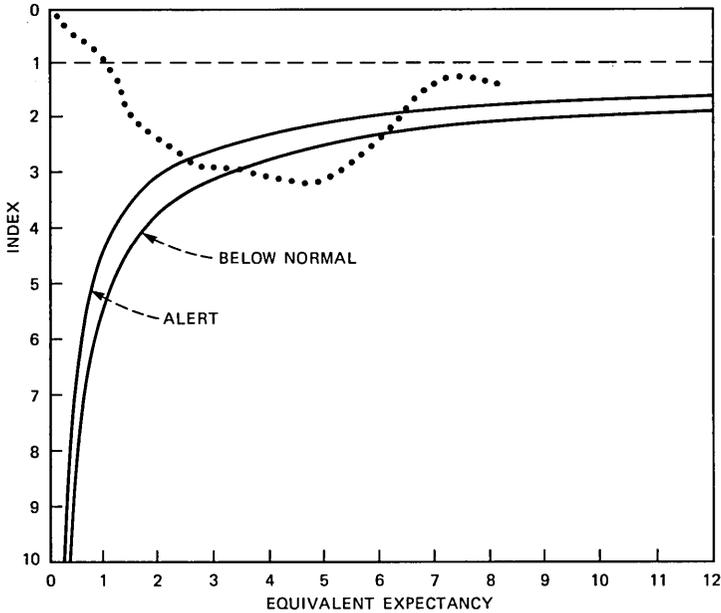


Fig. 9—Index Bogie chart for a good past history. Equivalent expectancy is a measure of how many defects are expected in the sample; so, equivalent expectancy increases with sample size. During a rating period, as the sample size increases, one can track the observed sample index (dotted curve) and compare it to Below Normal and ALERT thresholds.

then the rate would be BN. But the period ended at an equivalent expectancy of eight and there is no exception.

The ALERT Bogie curve in Fig. 10 is interesting. It starts at zero, so you start the period on ALERT. The past history is so bad, that in the

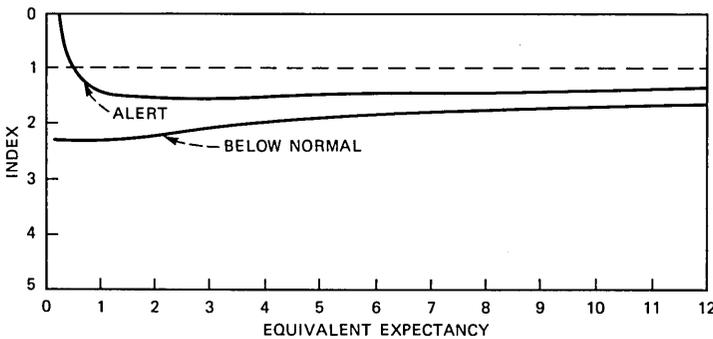


Fig. 10—Index Bogie chart for a substandard past history. The Below Normal and ALERT thresholds are very tight. At the beginning of the period, the product is on ALERT until proven otherwise.

absence of any current data the probability that the current quality will be substandard exceeds 0.95.

5.3 Bogie contour plots

For a fixed past history and current equivalent expectancy, there is a BN Bogie for the current sample index. If the sample index is worse than the BN Bogie, then the product is BN. Figure 11 is a contour plot

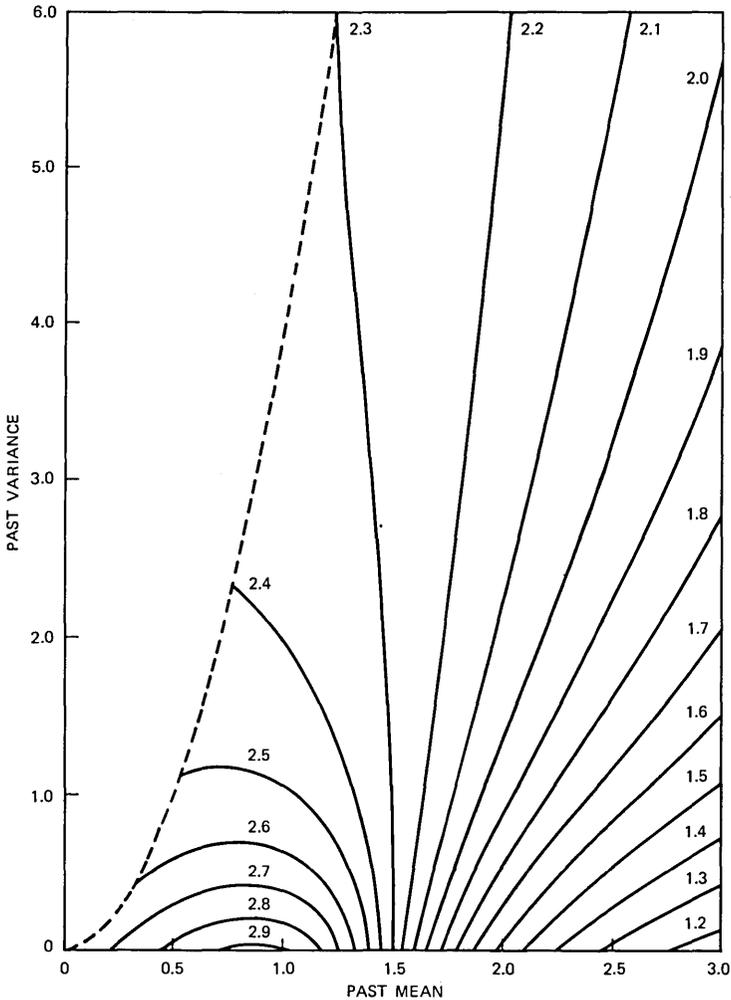


Fig. 11—Below Normal Bogie contour plot. If the past mean is 0.8 and the past variance is 0.7 (on an index scale), then the product is on the contour labeled 2.6. This means that if the current sample index exceeds 2.6, the product will be Below Normal (equivalent expectancy equals 5 for this chart).

of the BN Bogie for an equivalent expectancy of five. The axes are the mean and variance of the five past values of the sample index; i.e.,

$$\bar{I}_p = (1/5) \sum_{t=1}^5 I_t,$$

$$S_p^2 = (1/5) \sum_{t=1}^5 (I_t - \bar{I}_p)^2,$$

where I_t is the sample index in past period t . For given values of \bar{I}_p and S_p^2 , we used a standard pattern of I_t 's to compute the Bogie. The results are insensitive to pattern. The dashed curve is an upper bound for S_p^2 .

To see how the contour plot works, consider an example. Suppose $\bar{I}_p = 0.8$ and $S_p^2 = 0.7$. The point (0.8, 0.7) falls on the contour labeled 2.6. This means that if the current sample index exceeds 2.6, then the product will be BN. The contour labeled 2.6 is the set of all pairs (\bar{I}_p , S_p^2) that yield a BN Bogie of 2.6. The T -rate associated with a BN Bogie of 2.6 is -3.6 , as shown in Table II.

This contour plot summarizes the BN behavior of QMP for an equiv-

Table II—Index to T -rate conversion table

Index (I)	T -Rate (T)*		
	Equivalent Expectancy (e)		
	1	5	10
1.0	0	0	0
1.1	-0.1	-0.2	-0.3
1.2	-0.2	-0.4	-0.6
1.3	-0.3	-0.7	-0.9
1.4	-0.4	-0.9	-1.3
1.5	-0.5	-1.1	-1.6
1.6	-0.6	-1.3	-1.9
1.7	-0.7	-1.6	-2.2
1.8	-0.8	-1.8	-2.5
1.9	-0.9	-2.0	-2.8
2.0	-1.0	-2.2	-3.2
2.1	-1.1	-2.5	-3.5
2.2	-1.2	-2.7	-3.8
2.3	-1.3	-2.9	-4.1
2.4	-1.4	-3.1	-4.4
2.5	-1.5	-3.4	
2.6	-1.6	-3.6	
2.7	-1.7	-3.8	
2.8	-1.8	-4.0	
2.9	-1.9	-4.2	
3.0	-2.0	-4.5	
3.1	-2.1		
3.2	-2.2		
3.3	-2.3		
3.4	-2.4		
3.5	-2.5		

* $T = \sqrt{e}(1 - I)$.

alent expectancy of five. As \bar{I}_p gets larger than one, the BN Bogie gets smaller. If \bar{I}_p exceeds 1.6, then the BN Bogie is smaller than 2.34, which corresponds to a T -rate of -3 . So in T -rate terms, BN triggers earlier than a T -rate of -3 .

For \bar{I}_p less than 1.4, as S_p^2 gets larger, the BN Bogie gets smaller. This is because large S_p^2 implies large process variance which makes an observed deviation more likely to be significant.

For very small S_p^2 , as you move from $\bar{I}_p = 0$ to $\bar{I}_p = 1$, the BN Bogie increases from 2.6 (T -rate = -3.6) to 2.9 (T -rate = -4.2). This is an apparent paradox. The better the process average, the less cushion the producer gets.

This is *not* a paradox, but an important characteristic of QMP. Remember with QMP we are making an inference about current quality, not long-run quality. If we have a stable past with $\bar{I}_p = 0.2$, and we suddenly get a sample index of 2.7, then this is very strong evidence that the process has changed and very probably become worse than standard. If we have a stable past with $\bar{I}_p = 1$, and we suddenly get a sample defect index of 2.7, then the evidence of change is not as strong as with $\bar{I}_p = 0.2$. The weight we put on the past data depends on how consistent the past is with the present.

Notice that the maximum BN Bogie is 2.92 and occurs at $\bar{I}_p = 0.85$ and $S_p^2 = 0$. It would be a *mistake* for the producer to conclude from the contour plot that he should control his process at $\bar{I}_p = 0.85$ and $S_p^2 = 0$. He cannot achieve $S_p^2 = 0$. The sample index has substantial sampling variance that the producer cannot control.

The Bogie contour plots provide the engineer with a manual tool to forecast the number of demerits that will be allowed by the end of a period. So we have published a book of BN and ALERT Bogie contour plots for equivalent expectancies from 0.5 to 25.

5.4 Nonlinearity of QMP

It is tempting to conjecture that if both the process average and the current sample index for one rate are worse than for another, then the Best Measure will also be worse. This is because the Best Measure is a weighted average between the process average and the current sample index. But, since the weight depends on the data nonlinearly, the conjecture is *not true*.

To illustrate this, consider Fig. 12. The six sample indices in Chart B are uniformly worse than the six sample indices in Chart A. But the Best Measure in Chart B is better than for Chart A. The reason is that the weight in B is 0.54 vs 0.12 for Chart A.

VI. OPERATING CHARACTERISTICS

The T -rate and QMP methods of rating are similar in some respects, but there are major differences. In this section, these differences are

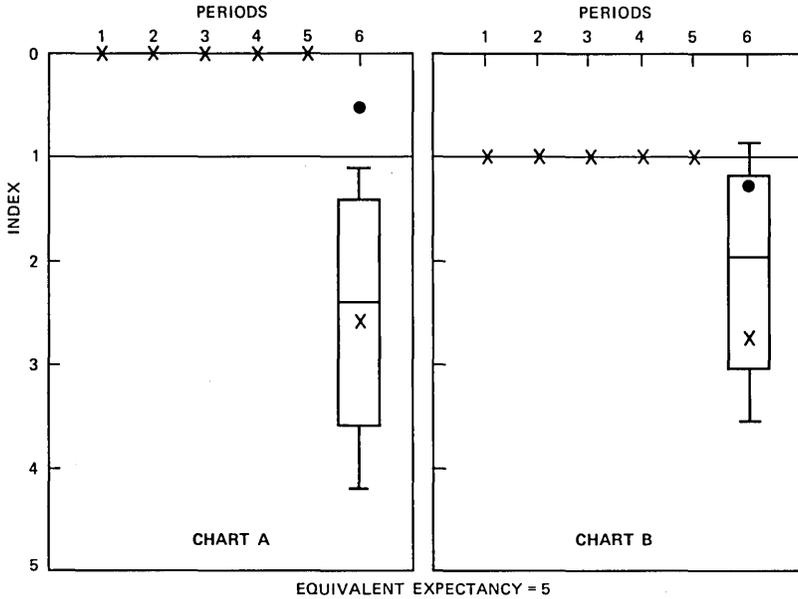


Fig. 12—Nonlinearity of QMP. The sample indices in Chart B are uniformly worse than the sample indices in Chart A; but, the QMP result in Chart B is better than for Chart A. In Chart A, the data provides very strong statistical evidence of an unstable process, so the past data is used very little in estimating current quality. This is not as pronounced in Chart B.

explored using operating characteristics. The differences are a result of different rating formulas and assessment practices.

6.1 Ranges of probability substandard for T-rate exceptions

A QMP analysis of a rating class provides a probability, PS , that the rating class is substandard. For a typical rating period analyzed in detail, we computed PS for all T-rate Below Normals and ALERTS. Table III shows the results.

So we find that for T-rate BNS, the QMP PS is typically high (greater than 0.97); but, there can be an occasional low PS (e.g., 0.75). However, for T-rate ALERT's, the QMP PS is frequently low (e.g., 0.85). This is because the T-rate ALERT is an indicator of long-run quality, not current quality.

Table III—Ranges of probability substandard for T-rate exceptions

Exception	Range of Probability	Outlier Probability
Below Normal	0.97 → 1.00	0.75
ALERT	0.83 → 0.99	0.59

6.2 Producer's risk and exceptions

Any list of rating classes that is put in an exception report has a producer's risk. It is the fraction of rating classes on the list whose population quality meets the standard. For a given period, let θ_i = population index, rating class i , $i = 1, \dots, I$. Label the rating classes so that product 1 through product L are on the exception list.

Having done QMP for each rating class, we have a posterior distribution for each θ_i . Now let

$$u_i = \begin{cases} 1, & \text{if } \theta_i \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

The number of rating classes on the list whose population quality meets the standard is

$$\sum_{i=1}^L u_i,$$

with posterior expected value

$$\sum_{i=1}^L \Pr\{\theta_i \leq 1\}.$$

Hence,

$$[\text{producer's risk}]^* = \frac{\sum_{i=1}^L \Pr\{\theta_i \leq 1\}}{L}.$$

In QMP, there is an exception list for each threshold probability (TP). TP = 0.95 corresponds to the list of all QMP BNS and ALERTS. Figure 13 shows the QMP producer's risk and number of exceptions as a function of TP for the manufacturing audits in a particular period. The smaller TP, the bigger the exception list and the bigger the producer's risk. Also, note that the producer's risk must be less than 1-TP.

The set of all T -rate BNS and ALERTS is another exception list, whose producer's risk is 0.037. This is relatively large because some individual ALERTS have relatively large probabilities (e.g., 0.15) of being standard. The number of T -rate exceptions (BN + ALERT) is shown to be 34.

Of course to implement QMP, a particular TP had to be chosen. The TP that would match the T -rate producer's risk is about 0.885. But that would lead to an unreasonable (70 percent) increase in exceptions, and a producer's risk of 0.037 is considered too high for this type of exception reporting, because of the high cost of false alarms. So we took TP = 0.95, a reasonable balance between producer's risk and size of the exception report.

* This is not the classical definition of producer's risk.

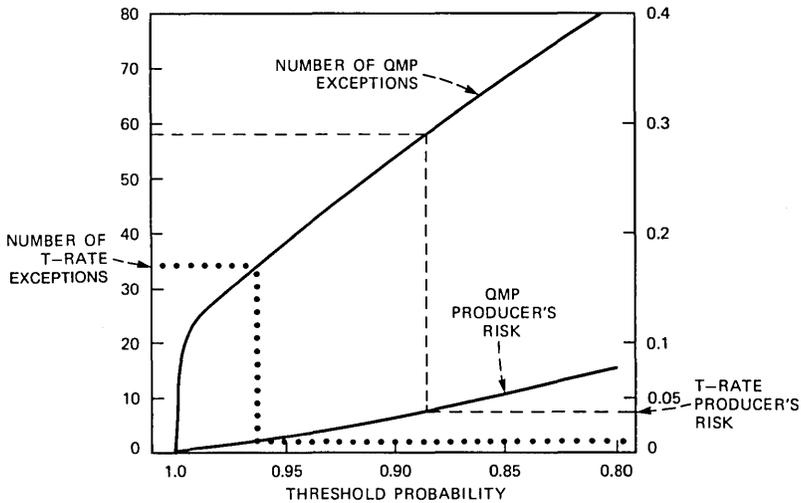


Fig. 13—Operating characteristics of QMP versus the *T*-rate. As the QMP threshold probability for exceptions (currently set at 0.95) is lowered, the number of exceptions and the producer's risk (for a particular rating period) both increase. The number of exceptions and producer's risk for the *T*-rate were 34 and 0.037. For threshold probabilities between 0.96 and 0.89, QMP has more exceptions and lower producer's risk than the *T*-rate.

It should be recognized that these curves depend on the particular set of audits being analyzed. For example, the curves depend on the audit sample sizes. It would be possible to lower sample sizes, decrease the threshold probability, and still maintain a comparably sized exception report with a reasonable producer's risk.

Note that consumer's risk is not analyzed in this paper. Consumer's risk is more relevant to acceptance sampling than to an audit. The main purpose of the audit is to provide quality results to management including a compact exception report of high integrity. The Western Electric quality control organizations have primary responsibility for the quality of each individual lot of product.

VII. EXAMPLES OF QMP

Here we explore specific examples that illustrate the similarities and differences between QMP and the *T*-rate. In the examples, both QMP and *T*-rate results are based on the same defect data. For the actual implementation of QMP, the defect assessment rules will be different than they are for the *T*-rate as explained in Section 3.1. The intent of this section is to compare how the two rating methods work on the same data.

The examples are shown in Figs. 1, 2, and 14 through 17. These figures show a comparison between the time series of T -rates and QMP box charts. Table IV contains summaries of the QMP calculations for the particular periods that will be discussed in the following text.

The QMP calculations shown do not use 1976 data. The box chart for the first period of data available is not shown except in Fig. 15. Period 7706 is the first period for which five periods of past data are used in the QMP box charts. So the comparisons made in this section will involve periods 7706 through 7808.

7.1 Agreement with T -rate

Figure 14 illustrates a T -rate borderline* in 7806 preceded by a good history. Since the equivalent expectancy (2.78) is fairly small and the process is fairly stable, the Best Measure (1.81) is heavily weighted (0.65) towards the process average (1.32). The posterior variance (0.36) is fairly large, so I95% is better than standard. However, in the next period, the T -rate plummets to -4.8 and the process average drops to 1.77. Now the rate is clearly BN.

7.2 Disagreement with T -rate

In Fig. 15, 7802, the T -rate is -3.8 (BN) but there is no exception for QMP. One reason is that QMP is based on the assumption that equivalent defects have a Poisson distribution. A T -rate of -3.8 is very significant for a normal distribution, but not as significant for a Poisson distribution with an equivalent expectancy of 0.29. For a normal distribution, the probability, given standard quality, of being below -3.8 is 0.000072. Now the observed number of equivalent defects in 7802 is 2.36. The approximate Poisson probability of exceeding 2.36 equivalent defects given an equivalent expectancy of 0.29 is 0.15—very different from 0.000072.

Another reason is that the QMP result for 7802 is based on one period of data. Rather than using the sample defect index (8.00) as the process average, we use a Bayes estimate [eq. (7)] of 2.77.

Figure 16 is a similar example. In 7708 the T -rate of -2.8 is BN because in 7705 the T -rate was -2.7 . But again, the -2.8 T -rate overstates the significance. The equivalent expectancy is only 0.23. Also, the weight (0.60) on the process average (1.81) adjusts the sample index (6.81) to the more moderate Best Measure (3.83). This, together with the large posterior variance (8.47), implies a comfortable I95% of 0.56.

Figure 1 illustrates how two similar T -rates, both on ALERT, can be either a QMP BN or normal. Compare 7708 with 7804. The sample

* $-3 \leq T\text{-rate} < -2$, but a good history.

Table IV—Summary data for QMP examples

Figure	Period	I_T	$\hat{\theta}$	e_T	σ_T^2	S^2	σ^2	$\hat{\gamma}^2$	$\hat{\omega}_T$	$\hat{\theta}_T$	V_T	195%	199%
14	7806	2.74	1.32	2.78	0.48	0.45	0.42	0.25	0.65	1.81	0.36	0.95	0.71
14	7807	3.52	1.77	3.60	0.49	0.89	0.55	0.50	0.50	2.65	0.56	1.55	1.22
15	7802	8.00	2.77	0.29	9.55	6.80	3.85	3.55	0.73	4.19	6.97	0.47	0.96
16	7708	6.81	1.81	0.23	7.81	9.64	7.44	5.30	0.60	3.83	8.47	0.56	0.21
1	7708	1.57	2.00	7.19	0.29	0.47	0.24	0.26	0.51	1.79	0.15	1.20	1.009
1	7802	1.61	1.51	7.19	0.21	0.44	0.21	0.26	0.45	1.57	0.14	1.01	0.84
1	7804	1.50	1.32	6.71	0.20	0.16	0.19	0.10	0.67	1.38	0.091	0.93	0.78
2	7707	1.04	1.35	4.92	0.28	0.14	0.29	0.16	0.70	1.26	0.11	0.76	0.61
17	7807	3.10	1.10	1.40	0.78	1.36	0.91	0.75	0.51	2.08	1.02	0.74	0.46

indices of 1.57 and 1.50 are very similar, but the process averages of 2.00 and 1.32 are very different and the weights of 0.51 and 0.67 are different. Hence, the Best Measures are very different and the conclusions are very different.

Figure 2 illustrates a “weak” ALERT under the T -rate. The T -rate in 7705 through 7707 are -0.1 , -0.2 , and -0.1 , respectively. Although it

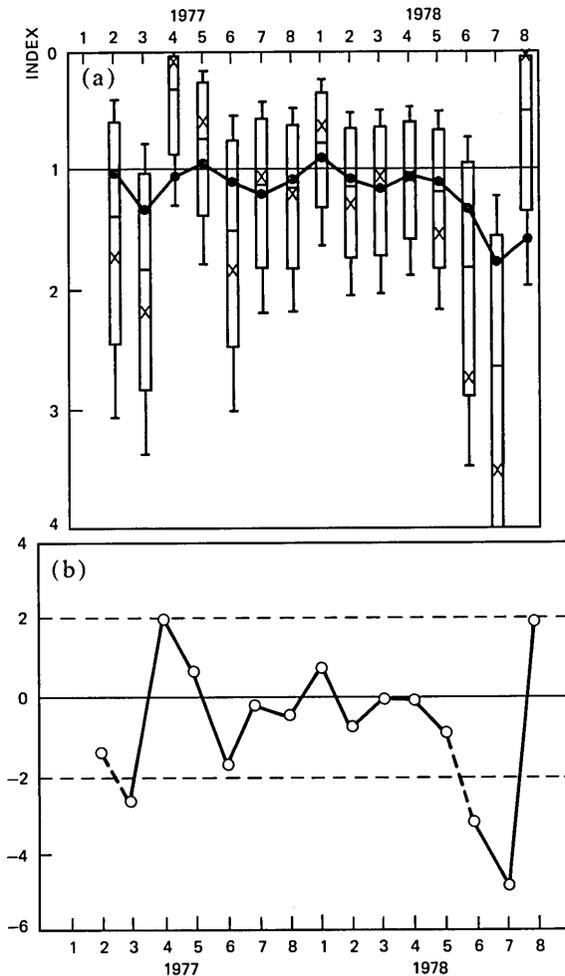


Fig. 14—Example of agreement. Throughout 1978, QMP and the T -rate are in agreement. The drop in the sixth period was called “borderline” under the T -rate, because it was the first excursion below -2 and it was moderate. The QMP box chart conveys the same borderline message. In the seventh period, the product was Below Normal for both systems.

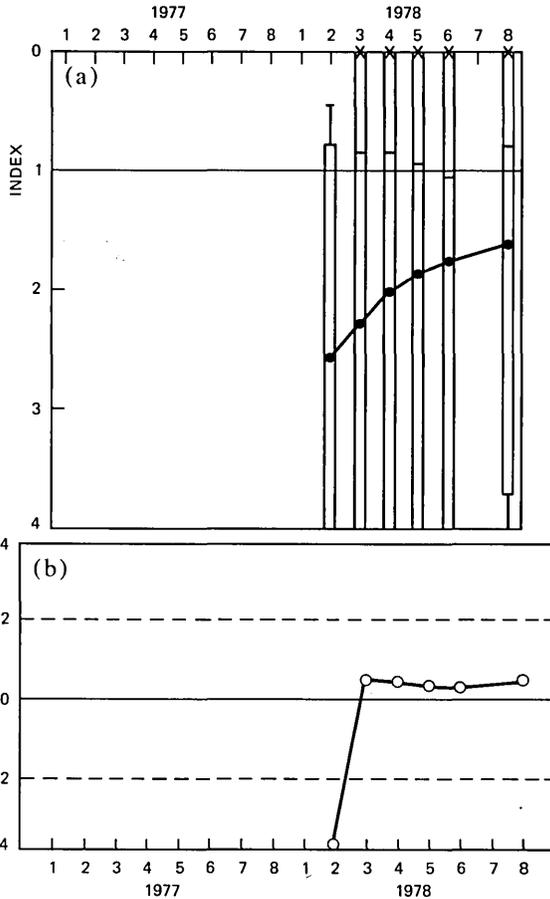


Fig. 15—Poisson versus Gaussian assumption. In the second period of 1978, the expected number of defects in the sample was 0.29 and the observed number of equivalent defects was 2.36. Under the Gaussian assumption, the observed significance level is 0.000072 (i.e., T -rate = -3.8); but, under Poisson, the level is 0.15. This explains why the QMP box chart contains the standard.

is unlikely that the quality standard was being met in *every* period from 7702 through 7707, it is *not* unlikely (probability of 0.23) that the quality standard was being met in 7707.

7.3 Modification treatment

The T -rate system had modification treatments that resulted from the statistical deficiencies of the T -rate (see Section 2.8). There are no modification treatments in QMP. The Poisson model and the stabilizing effect of shifting the sample index towards the process average alleviate the need for modification treatments.

In Fig. 17, the 7807 unmodified T -rate is -2.5 . It is modified to $+0.6$ because of the “isolated” A weight (100 demerits) defect. Under QMP, the process average (1.10) is only slightly substandard, the weight (0.51) is medium, and the equivalent expectancy (1.40) is small. All this implies a safe $I95\%$ (0.74) without modification.

7.4 Venn diagram of BNS and ALERTs

In the Venn diagram of Fig. 18, BNS are shown by circles and ALERTs are shown by rectangles. QMP results are shown by dashed lines and T -rate results are shown by solid lines. Every rating class that is BN or

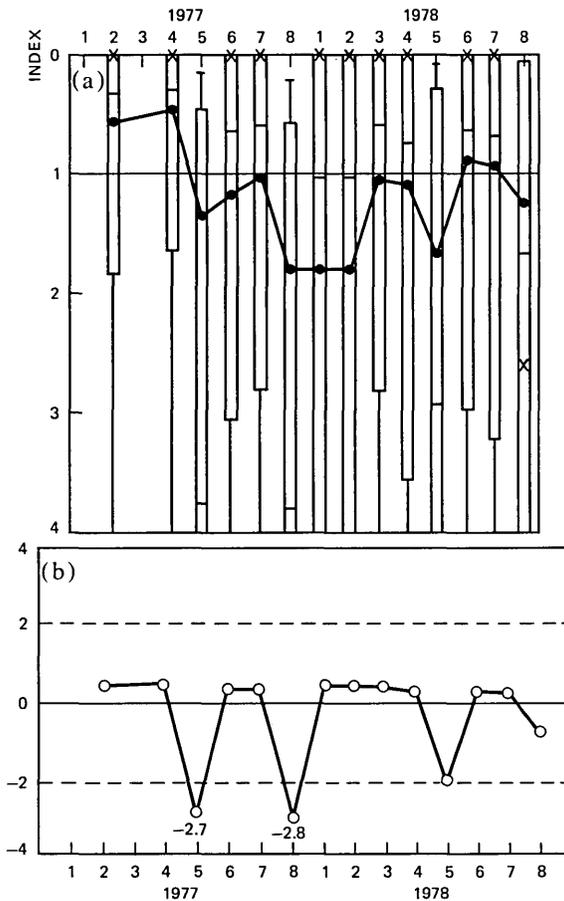


Fig. 16—Statistical jitter in the T -rate. With small samples and zero defects, the T -rate is slightly larger than zero. Every time a defect is found, the T -rate jitters. The message in the QMP chart is that there is too much uncertainty to reach any conclusions.

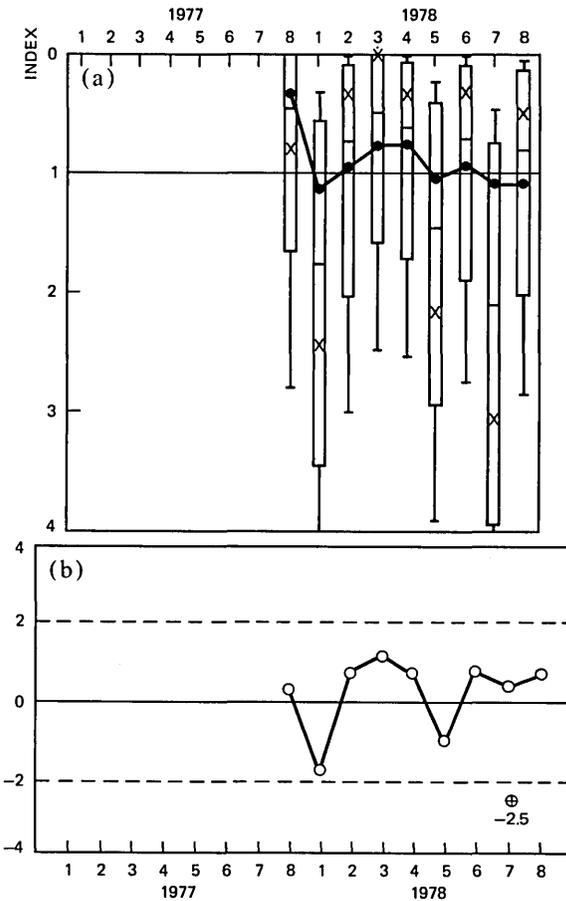


Fig. 17—A case of *T*-rate modification treatment. Because the *T*-rate is biased for small samples, modification treatments were needed to compensate (seventh period, 1978). QMP mathematics obviates the need for modification.

ALERT under QMP or the *T*-rate is represented in the Venn diagram.

Ten rating classes were BN under both methods of rating. Five rating classes were BN under the *T*-rate but ALERT under QMP. There were 16 rating classes that were ALERT under the *T*-rate but normal under QMP. This indicates a major difference. ALERT under the *T*-rate is strong evidence that the quality standards for the current period or some of the past periods have not been met. But it *does not* necessarily imply strong evidence that the quality standard for the current period has not been met. ALERT under QMP implies more than a 95 percent chance that the current quality standards have not been met.

VIII. ACKNOWLEDGMENTS

Many members of the Bell Laboratories Quality Assurance Center and the Western Electric Quality Assurance Directorate have made important contributions to the development of QMP.

B. T. Howard and E. Fuchs originally gained the support of the senior management of Bell Laboratories and Western Electric. R. A. Peters and L. E. Bray followed this with many presentations to Western Electric management. The role of liason among all the organizations involved was handled by J. M. Wier.

As for technical contributions, R. A. Senior did the original data analyses that transformed the idea into a concrete proposal. C. S. Sherrerd developed the prototype QMP reporting system. S. G. Crawford developed the summary reports and numerical analyses algorithms that permit efficient computation of a thousand QMP charts each period. S. W. Roberts, Jr., developed the QMP defect assessment practices. M. S. Phadke made contributions to the mathematical theory of QMP.

Several members of the Quality Information Systems groups have been developing the data bases and software associated with the official implementation of QMP. They are E. W. Hinds, V. A. Partridge, G. D. Rosen, P. A. Douglas, J. H. Carey, P. D. Ting, and S. Kadakia.

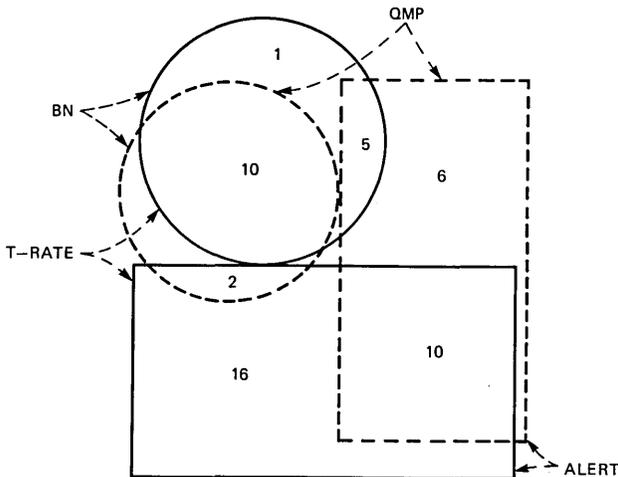


Fig. 18—Venn diagram of exceptions. The Venn diagram accounts for all QMP and T-rate exceptions for a particular period using defects assessed under the T-rate. The lists of ALERTS under the two systems are quite different (only 10 out of 32 in common).

Western Electric Quality Assurance Headquarters people administered and analyzed the QMP trial and prepared the material for numerous presentations. They are C. Popik, N. O. Dickerson, N. Linardakis, T. M. Ferme, D. Snyder, H. M. Cook, E. Hoffman, and S. Chory.

APPENDIX A

The Gamma Distribution

A random variable Y has a standard gamma distribution if

$$\Pr\{Y \leq y\} = G_\alpha(y) = \int_0^y \frac{1}{\Gamma(\alpha)} x^{\alpha-1} e^{-x} dx,$$

$\alpha = \text{shape parameter.} \tag{37}$

A random variable $X = \tau Y$ has a gamma distribution with shape parameter α and scale parameter τ . We write

$$X \sim \text{Gamma}(\alpha, \tau)$$

and

$$\Pr\{X \leq x\} = G_\alpha\left(\frac{x}{\tau}\right).$$

The probability density of X is

$$\frac{1}{\tau^\alpha \Gamma(\alpha)} x^{\alpha-1} e^{-x/\tau}.$$

The mean and variance of X are

$$E(X) = \tau\alpha, \quad V(X) = \tau^2\alpha;$$

hence

$$\alpha = E^2(X)/V(X), \quad \tau = V(X)/E(X).$$

A chi-squared random variable with ν degrees of freedom has a Gamma distribution; namely,

$$\chi_\nu^2 \sim \text{Gamma}\left(\frac{\nu}{2}, 2\right).$$

APPENDIX B

The Poisson-Gamma Bayesian Model

Theorem B.1: Assume

$$x_t | \theta_t \sim \text{Poisson}(e_t \theta_t), \quad e_t \text{ known, } \theta_t \text{ unknown}$$

and

$$\theta_t \sim \text{Gamma}\left(\frac{\theta^2}{\gamma^2}, \frac{\gamma^2}{\theta}\right)$$

(i.e., mean = θ , variance = γ^2).

Then

$$\theta_t | x_t \sim \text{Gamma}\left(\frac{\bar{\theta}_t^2}{V_t}, \frac{V_t}{\bar{\theta}_t}\right),$$

where

$$\begin{aligned}\bar{\theta}_t &= E(\theta_t | x_t) \\ &= \omega_t \theta + (1 - \omega_t) I_t, \\ I_t &= x_t / e_t, \\ \omega_t &= \frac{\theta / e_t}{\theta / e_t + \gamma^2}, \\ V_t &= V(\theta_t | x_t) \\ &= (1 - \omega_t) \bar{\theta}_t / e_t.\end{aligned}$$

Proof: The sampling distribution of x_t is:

$$f(x_t | \theta_t) = \frac{(e_t \theta_t)^{x_t} \exp\{-e_t \theta_t\}}{x_t!}. \quad (38)$$

The process (prior) distribution of θ_t is*:

$$\begin{aligned}p_0(\theta_t) &= \frac{e_0^{x_0}}{\Gamma(x_0)} \theta_t^{x_0-1} e^{-e_0 \theta_t}, \\ \theta &= x_0 / e_0, \quad \gamma^2 = x_0 / e_0^2.\end{aligned} \quad (39)$$

By Bayes theorem, the posterior density of θ_t is proportional to the product of equations (39) and (38), which is in turn proportional to

$$[\theta_t^{x_0-1} e^{-e_0 \theta_t}] [\theta_t^{x_t} e^{-e_t \theta_t}] = \theta_t^{x_0+x_t-1} \exp[-(e_0 + e_t) \theta_t]. \quad (40)$$

We recognize eq. (40) as proportional to a Gamma density. So the posterior distribution is Gamma with shape parameter $x_0 + x_t$ and scale parameter $1/(e_0 + e_t)$. And the posterior mean and variance are

$$\bar{\theta}_t = \frac{x_0 + x_t}{e_0 + e_t}, \quad (41)$$

$$V_t = \frac{\bar{\theta}_t}{e_0 + e_t}. \quad (42)$$

* Here, x_0 and e_0 are *not* the same as the "prior data" introduced in Section 4.3.2.

Now multiply the numerator and denominator in both eqs. (41) and (42) by θ/e_0e_t . Theorem B.1 follows. Q.E.D.

APPENDIX C

Chi-Square, Gamma Bayesian Model

Theorem C.1: Assume there is a statistic, ss , for which

$$\frac{\omega}{\sigma^2} (ss) \left| \omega \sim \chi^2_\nu \quad (\text{chi-square, } \nu \text{ degrees of freedom}) \right.$$

σ^2 known, ω unknown

and

$$\omega \sim \text{Gamma}\left(a_0, \frac{\sigma^2}{b_0}\right), \quad a_0, b_0 \text{ known.}$$

Then

$$\omega | ss \sim \text{Gamma}\left(a, \frac{\sigma^2}{b}\right),$$

where

$$a = a_0 + \frac{\nu}{2},$$

$$b = b_0 + \frac{ss}{2}.$$

Proof: The sampling density of ss given ω is

$$f(ss | \omega) = \frac{1}{(2\sigma^2/\omega)^{\nu/2} \Gamma(\nu/2)} (ss)^{(\nu/2) - 1} \exp\left[-\left(\frac{ss}{2\sigma^2/\omega}\right)\right]. \quad (43)$$

The prior density of ω is

$$\rho_0(\omega) = \frac{1}{(\sigma^2/b_0)^{a_0} \Gamma(a_0)} \omega^{a_0-1} \exp\left[-\left(\frac{\omega}{\sigma^2/b_0}\right)\right]. \quad (44)$$

By Bayes theorem, the posterior density of ω is proportional to the product of eqs. (43) and (44):

$$\left\{ \omega^{\nu/2} \exp\left[-\left(\frac{\omega}{2\sigma^2/ss}\right)\right] \right\} \left\{ \omega^{a_0-1} \exp\left[-\left(\frac{\omega}{\sigma^2/b_0}\right)\right] \right\}$$

$$= \omega^{a-1} \exp\left[-\left(\frac{\omega}{\sigma^2/b}\right)\right].$$

Q.E.D.

Definition: Let $X \sim \text{Gamma}(\alpha, \tau)$. Denote the conditional distribution of X given $X \leq c$ by

$$C - \text{Gamma}(\alpha, \tau, c).$$

Corollary: If instead,

$$\omega \sim C - \text{Gamma}\left(a_0, \frac{\sigma^2}{b_0}, 1\right),$$

then

$$\omega | \text{SS} \sim C - \text{Gamma}\left(a, \frac{\sigma^2}{b}, 1\right).$$

Theorem C.2: If $\omega \sim C - \text{Gamma}(a, \sigma^2/b, 1)$, then

$$E(\omega) = \frac{1}{RF},$$

$$V(\omega) = G,$$

where

$$R = \frac{b/a}{\sigma^2},$$

$$F = \frac{G_a(aR)}{G_{a+1}(aR)} \quad [\text{see (37)}],$$

$$G = \frac{1}{RF} \left[\left(\frac{a+1}{a} \right) \frac{G_{a+2}(aR)}{RG_{a+1}(aR)} - \frac{1}{RF} \right]. \quad (45)$$

Proof: Note

$$\frac{b}{\sigma^2} \omega = aR\omega \sim C - \text{Gamma}\left(a, 1, \frac{1}{aR}\right).$$

So

$$\begin{aligned} E(\omega) &= \frac{1}{aR} E(aR\omega) \\ &= \frac{1}{aRG_a(aR)} \int_0^{aR} y \frac{y^{a-1} e^{-y}}{\Gamma(a)} dy \\ &= \frac{\Gamma(a+1)}{aRG_a(aR)\Gamma(a)} \int_0^{aR} \frac{y^{(a+1)-1} e^{-y}}{\Gamma(a+1)} dy \\ &= \frac{aG_{a+1}(aR)}{aRG_a(aR)} \\ &= 1/RF. \end{aligned} \quad (46)$$

Now

$$\begin{aligned}
 E(\omega^2) &= \frac{1}{(aR)^2} E[(aR\omega)^2] \\
 &= \frac{1}{(aR)^2 G_a(aR)} \int_0^{aR} y^2 \frac{y^{a-1} e^{-y}}{\Gamma(a)} dy \\
 &= \frac{\Gamma(a+2)}{(aR)^2 G_a(aR) \Gamma(a)} \int_0^{aR} \frac{y^{(a+2)-1} e^{-y}}{\Gamma(a+2)} dy \\
 &= \frac{(a+1)a G_{a+2}(aR)}{(aR)^2 G_a(aR)} \\
 &= \frac{(a+1) G_{a+1}(aR) G_{a+2}(aR)}{aR^2 G_a(aR) G_{a+1}(aR)} \\
 &= \left(\frac{a+1}{a} \right) \frac{G_{a+2}(aR)}{R^2 F G_{a+1}(aR)}. \tag{47}
 \end{aligned}$$

This along with eq. (46) implies $V(\omega) = G$.

Computational formula for F

Let

$$g_a(x) = \frac{1}{\Gamma(a)} x^{a-1} e^{-x} dx.$$

From Ref. 11, page 262, 6.5.21,

$$G_{a+1}(x) = G_a(x) - \left(\frac{x}{a} \right) g_a(x). \tag{48}$$

Now define

$$\begin{aligned}
 B_a(x) &= \sum_{i=0}^{\infty} T(i), \quad T(0) = 1, \\
 T(i) &= T(i-1) \left[\frac{x}{a+i} \right] = 1 + \frac{x}{a+1} + \frac{x^2}{(a+1)(a+2)} + \dots
 \end{aligned}$$

By Ref. 12, page 3,

$$G_a(x) = \left(\frac{x}{a} \right) g_a(x) B_a(x). \tag{49}$$

Putting eqs. (48) and (49) together implies

$$F_a(x) \triangleq \frac{G_a(x)}{G_{a+1}(x)}$$

$$\begin{aligned}
 &= \frac{(x/a)g_a(x)B_a(x)}{(x/a)g_a(x)B_a(x) - (x/a)g_a(x)} \\
 &= \frac{B_a(x)}{B_a(x) - 1}.
 \end{aligned}$$

So

$$F = \frac{B_a(aR)}{B_a(aR) - 1}. \quad (50)$$

Computational formula for G

Directly by definition, it follows that

$$B_{a+1}(x) = \left(\frac{a+1}{x}\right)[B_a(x) - 1]. \quad (51)$$

Therefore,

$$\begin{aligned}
 \frac{1}{F_{a+1}(x)} &= 1 - \frac{1}{B_{a+1}(x)} \\
 &= 1 - \frac{1}{[(a+1)/x][B_a(x) - 1]} \\
 &= 1 - \left(\frac{x}{a+1}\right)[F_a(x) - 1].
 \end{aligned} \quad (52)$$

Now plug eq. (52) into the first term in the square bracket of eq. (45) and get

$$\begin{aligned}
 \left(\frac{a+1}{a}\right) \frac{G_{a+2}(aR)}{RG_{a+1}(aR)} &= \left(\frac{a+1}{a}\right) \frac{1}{RF_{a+1}(aR)} \\
 &= \left(\frac{a+1}{a}\right) \frac{1}{R} \left[1 - \left(\frac{aR}{a+1}\right)(F-1)\right] \\
 &= \frac{a+1}{aR} - (F-1).
 \end{aligned} \quad (53)$$

So

$$G = \frac{1}{RF} \left[\left(\frac{a+1}{aR}\right) - (F-1) - \frac{1}{RF} \right]. \quad (54)$$

APPENDIX D

Moments of functions of the sample index

If

$$x_t | \theta_t \sim \text{Poisson}(e_t \theta_t),$$

$$\theta_t | \theta, \gamma^2 \sim \text{Gamma}\left(\frac{\theta^2}{\gamma^2}, \frac{\gamma^2}{\theta}\right),$$

then $x_t | \theta, \gamma^2$ is a negative binomial with density

$$f(x_t | \theta, \gamma^2) = \frac{\Gamma(x_t + \theta^2/\gamma^2)}{x_t! \Gamma(\theta^2/\gamma^2)} \left[\frac{1}{1 + \theta/e_t \gamma^2} \right]^{x_t} \left[\frac{\theta/e_t \gamma^2}{1 + \theta/e_t \gamma^2} \right]^{\theta^2/\gamma^2}.$$

Let

$$\mu_1 = \text{mean of } x_t,$$

$$\mu_\nu = \nu\text{th central moment of } x_t, \quad \nu = 2, 3, \dots$$

Then according to (Ref. 11, page 929),

$$\mu_1 = \alpha P,$$

$$\mu_2 = \alpha P Q,$$

$$\mu_3 = \alpha P Q [Q + P],$$

$$\mu_4 = \alpha P Q + A(\alpha P Q)^2, \tag{55}$$

where

$$\alpha = \theta^2/\gamma^2,*$$

$$P = \gamma^2 e_t/\theta,$$

$$Q = 1 + P,$$

$$A = 3 + 6\gamma^2/\theta^2.$$

Now let

$$\xi_1 = \text{mean of } I_t,$$

$$\xi_\nu = \nu\text{th central moment of } I_t, \quad \nu = 2, 3, \dots$$

It follows from (55) that

$$\xi_1 = \theta,$$

$$\xi_2 = \gamma^2 + \frac{\theta}{e_t},$$

$$\xi_3 = \frac{2\gamma^4}{\theta} + \frac{3\gamma^2}{e_t} + \frac{\theta}{e_t^2},$$

$$\xi_4 = A\gamma^4 + \frac{2A\theta\gamma^2}{e_t} + \frac{A\theta^2 + \gamma^2}{e_t^2} + \frac{\theta}{e_t^3}. \tag{56}$$

* A different α from the one in the main text.

An application of these formulas is

$$\begin{aligned}
 V[(I_t - \theta)^2] &= \xi_t - \xi_2^2 \\
 &= (A - 1)\gamma^4 + \frac{2(A - 1)\theta\gamma^2}{e_t} + \frac{[(A - 1)\theta^2 + \gamma^2]}{e_t^2} + \frac{\theta}{e_t^3}. \quad (57)
 \end{aligned}$$

Now define

$$Y_t = (I_t - \theta)^2 - I_t/e_t.$$

Further applications of (56) are

$$\begin{aligned}
 E(Y_t) &= \xi_2 - \xi_1/e_t \\
 &= \gamma^2 \quad (58)
 \end{aligned}$$

and

$$\begin{aligned}
 V(Y_t) &= E(Y_t^2) - \gamma^4 \\
 &= E\left[(I_t - \theta)^2 - \frac{1}{e_t}(I_t - \theta) - \frac{\theta}{e_t}\right]^2 - \gamma^4 \\
 &= (A - 1)\gamma^4 + \frac{[2(A - 1)\theta\gamma^2 - 4\gamma^4/\theta]}{e_t} \\
 &\quad + \frac{[(A - 1)\theta^2 - 4\gamma^2]}{e_t^2}. \quad (59)
 \end{aligned}$$

REFERENCES

1. E. Fuchs and B. T. Howard, "Quality Assurance Tradition and Change," *Bell Laboratories Record*, 56, No. 9 (October 1978), pp. 226-31.
2. R. Peters and I. O. Karraker, "Why You Can Trust Your Telephone," *Ind. Res.* 17, No. 11 (November 15, 1975), pp. 47-51.
3. W. A. Shewhart, "Nature and Origin of Standards of Quality," *B.S.T.J.*, 37, No. 1 (January 1958), pp. 1-22.
4. W. A. Shewhart, *Economic Control of Quality of Manufactured Product*, New York: D. Van Nostrand, 1931.
5. H. F. Dodge, "A Method of Rating Manufactured Product," *B.S.T.J.*, 7 (April 1928), pp. 350-68.
6. H. F. Dodge and M. N. Torrey, "A Check Inspection and Demerit Rating Plan," *Ind. Qual. Contr.* 13, No. 1 (July 1956), pp. 1-8.
7. B. Efron and C. Morris, "Stein's Paradox in Statistics," *Sci. Am.*, 236, No. 5 (May 1977), pp. 119-27.
8. C. Stein, "Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution," *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, California: University of California Press, 1955, pp. 197-206.
9. Bruce Hoadley, "An Empirical Bayes Approach to Quality Assurance," *33rd Annual Technical Conference Transactions of ASAC*, May 14-16, 1979, pp. 257-63.
10. B. Efron and C. Morris, "Stein's Estimation Rule and Its Competitors—An Empirical Bayes Approach," *J. Am. Statist. Ass.*, 68, No. 341 (March 1973), pp. 117-30.
11. M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*, New York: Dover, 1965.
12. M. B. Wilk, R. Gnanadesikan, and M. J. Huyett, "Probability Plots for the Gamma Distribution," *Technometrics*, 4, No. 1 (February 1962), pp. 1-20.

Fractionally-Spaced Equalization: An Improved Digital Transversal Equalizer

By R. D. GITLIN and S. B. WEINSTEIN

(Manuscript received September 12, 1980)

Here we describe and demonstrate, via analysis and simulation, the performance improvement of voice-grade modems which use a Fractionally-Spaced Equalizer (FSE) instead of a conventional synchronous equalizer. The reason for this superior performance is that the FSE adaptively realizes the optimum linear receiver; consequently it can effectively compensate for more severe delay distortion than the conventional adaptive equalizer, which suffers from aliasing effects. An additional advantage of the FSE is that data transmission can begin with an arbitrary sampling phase, since the equalizer synthesizes the correct delay during adaptation. We show that an FSE combined with a decision feedback section, which can mitigate the effect of severe amplitude distortion, can compensate for a wide range of linear distortion. At 9.6 kbit/s, the FSE provides a 2 to 3 dB gain in output signal-to-noise ratio, relative to the synchronous equalizer, over worst-case private-line channels. This translates to a theoretical improvement of approximately two orders of magnitude in bit error rate.

I. INTRODUCTION

As is well known,^{1,2} high-speed (≥ 4.8 kbit/s) voiceband modems must employ some sort of adaptive equalization to achieve reliable performance in the presence of linear distortion and additive noise. The equalizers are invariably implemented using transversal filters, but the question of how the taps should be spaced has been, and still is, of great theoretical as well as practical interest. Conventionally, the equalizer taps are spaced at the reciprocal of the signaling rate. While it has been known theoretically that this synchronous structure does not, by itself, realize the optimum linear filter, it has up to this time provided adequate performance. The continuing demand for improved

performance at 9.6 kbit/s has renewed interest in adaptive equalizers whose taps are spaced closer than the reciprocal of twice the highest frequency component in the baseband signal.³⁻⁷ As we shall demonstrate, such Fractionally-Spaced Equalizers (FSEs) are able to compensate much more effectively for delay distortion than the conventional synchronous equalizers. Consequently, we will show that the performance of a FSE, with a sufficient number of taps, is almost independent of the channel delay distortion, and thus of the receiver sampling phase. More generally, the FSE is able to adaptively realize, in one device, the optimum linear receiver, which is known to be the cascade of a matched filter and a synchronously-spaced equalizer.⁸

The purpose of this paper is to report the results of an in-depth comparative analytical and simulation study of FSEs and the conventional synchronous equalizer. We also evaluate the performance of an equalizer which results when a decision-feedback section, which is particularly effective in compensating for amplitude distortion, is combined with an FSE. We present simulation results that compare the performance of practical-length synchronous and fractionally-spaced equalizers over a variety of voice-grade private-line channels.

Many years have elapsed between Lucky's invention of the adaptive synchronous equalizer,⁹ Gersho's³ and Brady's^{4,10} early work on FSEs, and our present interest in fractionally-spaced equalization. This is due to both the increased complexity required to implement the FSE, and the relatively satisfactory performance of the conventional synchronous equalizer. Recent investigators have regarded the FSE primarily as a means for mitigating the timing jitter produced by an envelope-derived timing recovery system.^{5,11} Our viewpoint, however, is that this property is just an example of the salient feature of the FSE—the ability to effectively compensate for an extremely wide range of delay distortion, and to deal more effectively with amplitude distortion than the synchronous equalizer.

In Section II we describe why an FSE has the ability to compensate for an arbitrary receiver sampling phase. Performance, as measured by the equalized mean-squared error, of an infinitely-long passband FSE is derived in Section III, and the corresponding results for a finite-length equalizer are described in Section IV. Simulation results, for typical voice-grade channels, are presented in Section V, and these results are used to compare the performance of the synchronous equalizer, the FSE, and the FSE with a decision-feedback section.

II. BASEBAND DESCRIPTION OF FRACTIONALLY-SPACED EQUALIZERS

We begin with a brief discussion of the ability of an FSE to compensate for any receiver timing phase. To do this we need the transfer function of a baseband fractionally-spaced equalizer. Consider the

received signal

$$r(t) = \sum_m a_m f(t - mT) + \nu(t), \quad (1)$$

where $\{a_n\}$ is the discrete multilevel data sequence, $1/T$ is the symbol rate, $f(t)$ is the system pulse response, and $\nu(t)$ is additive noise. As shown in Fig. 1, we denote excess bandwidth of the pulse $f(t)$ by α . The input to a *conventional synchronous* digital equalizer are samples of the filtered received signal at the instants $t = nT + \hat{\tau}$, i.e.,

$$r(nT + \hat{\tau}) = r_n(\hat{\tau}) = \sum_m a_m f(nT - mT + \hat{\tau}) + \nu(nT + \hat{\tau}). \quad (2)$$

The noiseless output of this nonrecursive digital filter, with tap weights $\{c_l\}$, is the sample sequence

$$u(nT + \hat{\tau}) = \sum_m a_m h(nT - mT + \hat{\tau}), \quad (3)$$

where the equalized pulse samples, $h(nT)$, have a (Nyquist-equivalent) Fourier transform

$$\begin{aligned} H_T(\omega) &= \sum_l c_l e^{-j\omega l T} \sum_k F\left(\omega + k \frac{2\pi}{T}\right) \exp\left[j\left(\omega + k \frac{2\pi}{T}\right) \hat{\tau}\right] \\ &= C_T(\omega) F_T(\omega). \end{aligned} \quad (4)$$

Here, $F_T(\omega)$ is the aliased spectrum of $F(\omega)$, $C_T(\omega)$ is the (periodic) transfer function of the equalizer, and, ideally, the equalizer output is the data symbol, i.e., $u(nT + \hat{\tau}) = a_n$.

Recall⁸ that the Nyquist-equivalent or folded (aliased) spectrum is the relevant transform when dealing with sampled-data systems. In particular, since $C_T(\omega) = C_T(\omega + k2\pi/T)$, the synchronously-spaced equalizer can only act to modify $F_T(\omega)$, as opposed to directly modifying $F(\omega)^{j\omega\hat{\tau}}$. In other words, the synchronous equalizer cannot exercise independent control over both sides of the rolloff region about $\omega =$

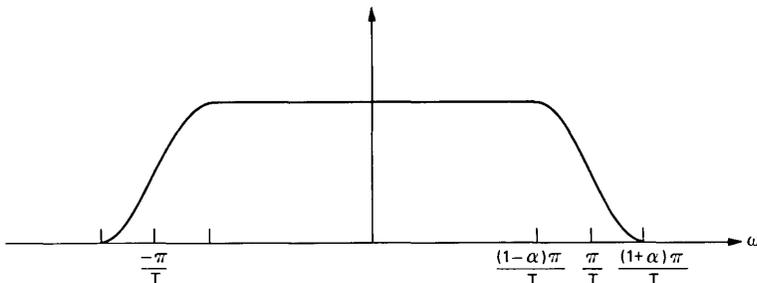


Fig. 1—Fourier transform $F(\omega)$ of baseband pulse $f(t)$ in eq. (1).

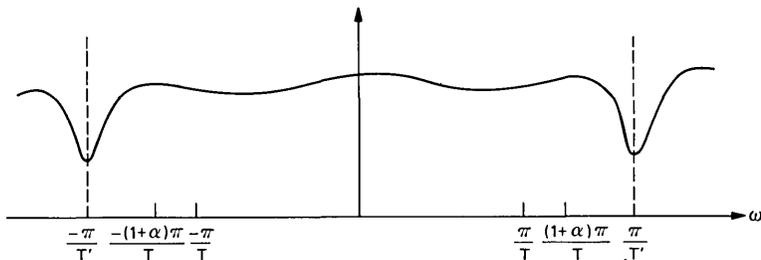


Fig. 2—Representative transfer function of a fractionally-spaced equalizer [tap spacing = $T' T/(1 + \alpha)$].

π/T . If, because of a severe phase characteristic and a poor choice of $\hat{\tau}$, a null is created in the rolloff portion* of the folded spectrum $F_T(\omega)$, then all the conventional equalizer can do to compensate for this null is to synthesize a rather large gain in the affected region; this leads to a severe performance degradation because of the noise enhancement at these frequencies.

Consider, on the other hand, a receiver which uses a *fractionally-spaced* equalizer with taps spaced $T' < T/(1 + \alpha)$ seconds apart. This equalizer has the (periodic) transfer function:

$$C_{T'}(\omega) = \sum_l c_l e^{-j\omega l T'} \quad (5)$$

Note that, if $\pi/T' \geq (1 + \alpha)\pi/T$, then the first repetition interval of the transfer function $C_{T'}(\omega)$ includes the rolloff portion of the spectrum, as shown in Fig. 2. We assume that, for digital implementation purposes, T' is generally an appropriate rational fraction of T . For an FSE receiver the equalizer input is sampled at the rate T' , but the equalizer output is still sampled at the rate T , since data decisions are made at symbol intervals. The equalized spectrum, just *prior* to the output sampler, is periodic (with period $2\pi/T'$) and is given by

$$H_{T'}(\omega) = C_{T'}(\omega) \sum_k F\left(\omega + k \frac{2\pi}{T'}\right) \exp\left[j\left(\omega + k \frac{2\pi}{T'}\right) \hat{\tau}\right], \quad (6)$$

and for systems where $\pi/T' \geq (1 + \alpha)\pi/T$ only the $k = 0$ term survives, i.e.,

$$H_{T'}(\omega) = C_{T'}(\omega) F(\omega) e^{j\omega \hat{\tau}}, \quad |\omega| \leq \frac{\pi}{T'}. \quad (7)$$

The salient aspect of (7) is that $C_{T'}(\omega)$ acts on $F(\omega) e^{j\omega \hat{\tau}}$ *before* aliasing, with respect to the output sampling rate, is performed. Thus $C_{T'}(\omega)$ can compensate for any timing phase—or phase distortion—by syn-

* This is the frequency range $(1 - \alpha)\pi/T \leq \omega \leq (1 + \alpha)\pi/T$.

thesizing a transfer characteristic of the form $e^{j\omega\hat{\tau}}$. Clearly, such compensation is highly desirable since it minimizes noise enhancement and avoids the extreme sensitivity to timing phase associated with the conventional equalizer.* After sampling the equalizer output at the rate $1/T$, the output spectrum is periodic with period $2\pi/T$ and is given by

$$\begin{aligned} H_T(\omega) &= \sum_l H_{T'}\left(\omega + l \frac{2\pi}{T}\right) \\ &= \sum_l C_{T'}\left(\omega + l \frac{2\pi}{T}\right) F\left(\omega + l \frac{2\pi}{T}\right) \exp\left[-j\left(\omega + l \frac{2\pi}{T}\right)\hat{\tau}\right]. \end{aligned} \quad (8)$$

Note that (8) differs from (4) in that it is the sum of equalized aliased components rather than an equalization of an already-formed sum of aliased components.

It is evident that an FSE is capable of much more than compensating for a poor choice of timing phase. With a properly chosen tap spacing ($T' \leq [1/(1 + \alpha)]T$), the FSE has the capabilities of an analog filter. Hence the FSE can be configured as the best linear receiver. In Section III we derive the structure and performance of such a receiver for a passband modem.

III. PERFORMANCE AND STRUCTURE OF THE OPTIMUM FSE

3.1 QAM systems

The receiver minimizing the mean-squared error is known to consist of a matched filter followed by a synchronous sampler.⁸ Our discussion in Section II demonstrated the equivalence of an appropriately sampled fractionally-spaced equalizer with an analog receiver. We begin by writing the transmitted signal $s(t)$, in an in-phase and quadrature, or quadrature amplitude modulated (QAM), data transmission system as the *real* part of the analytical signal

$$\tilde{s}(t) = s(t) + j\check{s}(t) = \sum_n \tilde{d}_n p(t - nT) e^{j\omega_c t}, \quad (9)$$

where \tilde{d}_n denotes the complex† discrete-multilevel data sequence, $a_n + jb_n$, $p(t)$ is the (generally real) baseband transmitter pulse shaping, $1/T$ is the symbol rate, ω_c is the radian carrier frequency, and $\check{s}(t)$ is the Hilbert transform of $s(t)$. In our presentation we will make extensive use of complex notation to denote either passband or in-phase and quadrature signals, as well as system pulse responses. A discussion

* In the synchronous equalizer, a "bad" timing phase is one which produces nulls in the folded spectrum of $F_T(\omega)$ of (4). In the FSE, a "good" timing phase is generated, regardless of the input sampling epoch, such that the FSE does a minimum of amplitude enhancement.

† The overtilde, \sim , is used to denote complex signals and samples.

of this approach is presented in the appendix. As shown in Fig. 3, $s(t)$ is transmitted through the passband (around ω_c) channel $x(t)$, with impulse response

$$\begin{aligned} x(t) &= x_1(t) \cos \omega_c t - x_2(t) \sin \omega_c t \\ &= \operatorname{Re}\{(x_1(t) + jx_2(t))e^{j\omega_c t}\} = \operatorname{Re}\{\tilde{x}_B(t)e^{j\omega_c t}\}, \end{aligned} \quad (10)$$

where the complex baseband-equivalent channel is defined by

$$\tilde{x}_B(t) = x_1(t) + jx_2(t). \quad (11)$$

Thus the received analytic signal has the representation

$$\tilde{r}(t) = r(t) + j\check{r}(t) = \sum_n \tilde{d}_n \tilde{f}_B(t - nT) e^{j(\omega_c t + \theta)} + \tilde{v}(t) e^{j\omega_c t}, \quad (12)$$

where $r(t)$ and $\check{r}(t)$ are the in-phase and quadrature components of the received signal (and are a Hilbert transform pair), $\tilde{f}_B(t)$ is the baseband-equivalent received pulse which is given by the convolution of $\tilde{x}_B(t)$ with $p(t)$, θ is the channel phase shift, and $\tilde{v}(t)$ is the complex noise signal.

At this point we may consider either a passband equalizer, which operates directly on $\tilde{r}(t)$, or a baseband receiver which processes $\tilde{r}(t) \exp[-j(\omega_c t + \theta)]$ —assuming that carrier-phase coherence² has been established. From a mathematical viewpoint both systems are equivalent, and here we find it convenient to filter the demodulated signal,

$$\tilde{q}(t) = \tilde{r}(t) e^{-j(\omega_c t + \theta)} = \sum_n \tilde{d}_n \tilde{f}_B(t - nT), \quad (13)$$

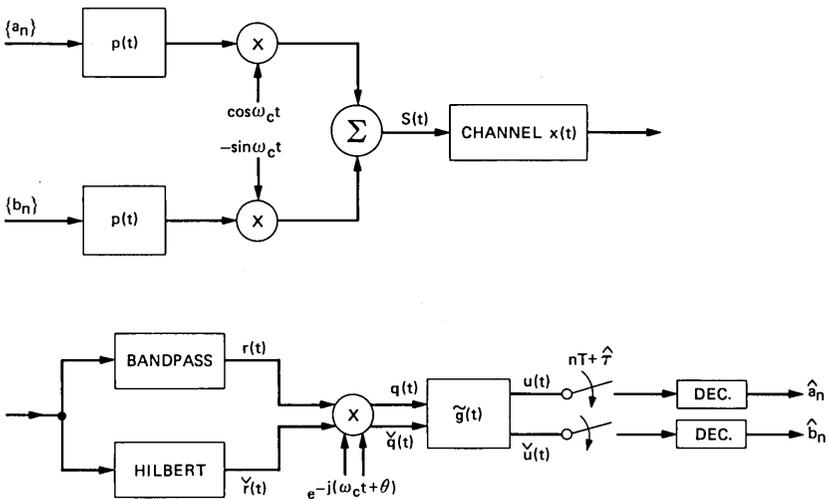


Fig. 3—QAM data transmission system. Variables with overtilde are complex, i.e., they have in-phase and quadrature components.

by the receiving filter

$$\tilde{g}(t) = g_1(t) + jg_2(t). \quad (14)$$

In writing (14) we have used the notation $g_1(t)$ and $g_2(t)$ rather than $g(t)$ and $\check{g}(t)$ to emphasize that the receiving filter does *not* correspond to an analytic pulse. Note, however, that the equalized signal at the filter output is given by the analytic signal

$$\begin{aligned} \tilde{u}(t) &= \tilde{q}(t) \circledast \tilde{g}(t) = (q(t) + j\check{q}(t)) \circledast (g_1(t) + jg_2(t)) \\ &= (q(t) \circledast g_1(t) - \check{q}(t) \circledast g_2(t)) \\ &\quad + j(\check{q}(t) \circledast g_1(t) + q(t) \circledast g_2(t)) \\ &= u(t) + j\check{u}(t), \end{aligned} \quad (15)$$

i.e., $u(t)$ and $\check{u}(t)$ are a Hilbert transform pair. As shown in Fig. 3 the in-phase and quadrature output signals $u(t)$ and $\check{u}(t)$ are synchronously sampled at $t = nT + \hat{\tau}$ and quantized to provide the data decisions \hat{a}_n and \hat{b}_n .

3.2 The mean-squared error

Our attention now turns to finding the linear filter, $\tilde{g}(t)$, which minimizes the output mean-squared error. The output or equalized, mean-squared error (MSE), which is the performance measure commonly used for in-phase and quadrature data transmission systems, is given as

$$\begin{aligned} \mathcal{E} &= E(|\tilde{e}_n|^2) = E\{e_n^2 + \check{e}_n^2\} \\ &= E\{(u(nT) - a_n)^2 + (\check{u}(nT) - b_n)^2\} \\ &= E(|\tilde{u}(nT) - \tilde{d}_n|^2), \end{aligned} \quad (16)$$

where E denotes the ensemble average with respect to the data symbols and the additive noise, \tilde{e}_n is the complex error sample, and e_n and \check{e}_n are the in-phase and quadrature errors, respectively. For convenience, we have absorbed the receiver's sampling phase, $\hat{\tau}$, into the pre-equalizer pulse response by incorporating the transfer function $e^{j\omega\hat{\tau}}$ into the transform of $\tilde{f}_B(t)$. In terms of the equalized pulse, $\tilde{h}(t)$, defined by

$$\tilde{h}(t) = \tilde{f}_B(t) \circledast \tilde{g}(t), \quad (17)$$

the filter output is

$$\tilde{u}(t) = \sum_n \tilde{d}_n \tilde{h}(t - nT) + \tilde{v}(t), \quad (18)$$

where the filtered noise, $\tilde{v}(t)$, is defined by

$$\tilde{v}(t) = \tilde{v}(t) \circledast \tilde{g}(t). \quad (19)$$

With these definitions in mind we can write the MSE as

$$\begin{aligned} \mathcal{E} &= E\{(\bar{u}(nT) - \bar{d}_n)(\bar{u}(nT) - \bar{d}_n)^*\} \\ &= E\{\bar{u}(nT)\bar{u}^*(nT) - \bar{d}_n\bar{u}^*(nT) - \bar{d}_n^*\bar{u}(nT) + \bar{d}_n\bar{d}_n^*\}. \end{aligned} \quad (20)$$

Using the independence of the data symbols and the independence of the noise samples, $\bar{v}(nT)$, the terms in (20) are readily evaluated. The first term is the quadratic form

$$\begin{aligned} E\{\bar{u}(nT)\bar{u}^*(nT)\} &= E[|\bar{d}_n|^2] \sum_m |\bar{h}(mT)|^2 + E[\bar{v}(nT)\bar{v}^*(nT)] \\ &= E[|\bar{d}_n|^2] \int \int [\bar{A}(t, \tau) \\ &\quad + \sigma^2\delta(t - \tau)]\bar{g}^*(t)\bar{g}(\tau) dt d\tau, \end{aligned} \quad (21)$$

where $\sigma^2 = E[|\bar{v}(nT)|^2]/E[|\bar{d}_n|^2]$, and the Hermitian kernel, $\bar{A}(t, \tau)$, is given by

$$\bar{A}(t, \tau) = \sum_n \bar{f}_B(nT - t)\bar{f}_B^*(nT - \tau). \quad (22)$$

Note that (21) can be written compactly as the quadratic form $(\bar{\mathbf{g}}^*, \bar{\mathcal{A}}\bar{\mathbf{g}})$, where $\bar{\mathbf{g}}$ corresponds to $\bar{g}(t)$, $\bar{\mathcal{A}}$ is the Hermitian integral operator with kernel $\bar{A}(t, \tau) + \sigma^2\delta(t - \tau)$, and (\mathbf{x}, \mathbf{y}) denotes the inner product $\int x(t)y(t) dt$. By straightforward evaluations, the second and third terms in (20) are seen to be

$$\begin{aligned} E(\bar{d}_n\bar{u}^*(nT)) &= E[|\bar{d}_n|^2] \int \bar{f}_B^*(-t)\bar{g}^*(t) dt, \\ E(\bar{d}_n^*\bar{u}(nT)) &= E[|\bar{d}_n|^2] \int \bar{f}_B(-t)\bar{g}(t) dt. \end{aligned} \quad (23)$$

Combining (20) to (23) the MSE, normalized by $E[|\bar{d}_n|^2]$, has the compact representation

$$\mathcal{E} = (\bar{\mathbf{g}}^*, \bar{\mathcal{A}}\bar{\mathbf{g}}) - (\bar{\mathbf{f}}_B^*, \bar{\mathbf{g}}^*) - (\bar{\mathbf{f}}_B, \bar{\mathbf{g}}) + 1, \quad (24)$$

which is a quadratic form, where $\bar{\mathcal{A}}$ is recognized as being a Hermitian operator (its kernel is conjugate symmetric).

3.3 The optimum filter

The MSE, given by (24), is minimized by taking the gradient with respect to $\bar{\mathbf{g}}$. The optimum filter is given as the solution of the integral equation,

$$\int [\bar{A}(t, \tau) + \sigma^2\delta(t - \tau)]\bar{g}(\tau) d\tau = \bar{f}_B(-t), \quad (25a)$$

or the equivalent operator equation

$$\tilde{\mathcal{A}}\tilde{\mathbf{g}} = \tilde{\mathbf{f}}_B. \quad (25b)$$

It turns out that the solution

$$\tilde{\mathbf{g}}_{\text{opt}} = \tilde{\mathcal{A}}^{-1}\tilde{\mathbf{f}}_B \quad (26)$$

can be explicitly determined. This is accomplished by writing the left-hand side of (25) as

$$\begin{aligned} & \int \left(\sum_n \tilde{f}_B(nT-t)\tilde{f}_B^*(nT-\tau) + \sigma^2\delta(t-\tau) \right) \tilde{g}(\tau) d\tau \\ &= \sum_n \tilde{f}_B(nT-t) \int \tilde{f}_B^*(nT-\tau)\tilde{g}(\tau) d\tau + \sigma^2\tilde{g}(t) \\ &= \sum_n \tilde{z}_n\tilde{f}_B(nT-t) + \sigma^2\tilde{g}(t), \quad (27) \end{aligned}$$

where $\tilde{z}_n = \int \tilde{f}_B^*(nT-\tau)\tilde{g}(\tau)d\tau$ are the equalized pulse samples. Equating (27) to the right-hand side of (25) gives

$$\sum_n \tilde{z}_n\tilde{f}_B(nT-t) + \sigma^2\tilde{g}(t) = \tilde{f}_B(-t), \quad (28)$$

which indicates that the optimum filter has the representation

$$\tilde{g}_{\text{opt}}(t) = \sum_n \tilde{c}_n\tilde{f}_B(nT-t), \quad (29)$$

where the \tilde{c}_n 's are to be determined. The solutions as represented by (29), is recognized as the cascade³ of a synchronously-sampled filter matched to $\tilde{f}_B(t)$, and a synchronously-spaced tapped delay line with weights $\{\tilde{c}_n\}$. To solve for the $\{\tilde{c}_n\}$ we substitute (29) into (25), giving

$$\begin{aligned} & \int \sum_n \tilde{f}_B(nT-t)\tilde{f}_B^*(nT-\tau) \sum_m \tilde{c}_m\tilde{f}_B(mT-t) d\tau \\ &+ \sigma^2 \sum_n \tilde{c}_n\tilde{f}_B(nT-t) = \tilde{f}_B(-t), \quad (30) \end{aligned}$$

and if we define the channel-correlation function,

$$\tilde{f}_{m-n} = \int \tilde{f}_B^*(nT-\tau)\tilde{f}_B(mT-\tau) d\tau, \quad (31)$$

then we can rewrite (30) as

$$\sum_n \sum_m \tilde{f}_{n-m}\tilde{c}_m\tilde{f}_B(nT-t) + \sigma^2 \sum_n \tilde{c}_n\tilde{f}_B(nT-t) = \tilde{f}_B(-t). \quad (32)$$

Taking Fourier transforms on both sides of (32), with respect to the

continuous variable t , gives

$$\sum_n \sum_m \tilde{f}_{n-m} \tilde{c}_m \tilde{F}_B^*(\omega) e^{j\omega n T} + \sigma^2 \sum_n \tilde{c}_n e^{j\omega n T} \tilde{F}_B^*(\omega) = \tilde{F}_B^*(\omega), \quad (33)$$

where $\tilde{F}_B(\omega)$ is the transform of $\tilde{f}_B(t)$. Dividing through by $\tilde{F}_B^*(\omega)$, over the region where the channel does not vanish, we can rewrite (33) as

$$\tilde{\mathcal{F}}(\omega) \tilde{C}_T(\omega) + \sigma^2 \tilde{C}_T(\omega) = 1, \quad (34)$$

where the following Fourier transforms, with respect to the discrete-time variables, are identified by

$$\begin{aligned} \tilde{C}_T(\omega) &= \sum_n \tilde{c}_n e^{j\omega n T}, \\ \tilde{\mathcal{F}}(\omega) &= \sum_n \tilde{f}_n e^{j\omega n T} = \frac{1}{T} \sum_l \left| \tilde{F}_B \left(\omega + l \frac{2\pi}{T} \right) \right|^2. \end{aligned} \quad (35)$$

The transform $\tilde{\mathcal{F}}(\omega)$ corresponds to the synchronously-sampled matched filter pulse, $\tilde{f}_B(t) \otimes \tilde{f}_B^*(t)$, and $\tilde{C}_T(\omega)$ is the transform of the coefficients of the synchronously-spaced tapped delay line. From (34) we have that

$$\tilde{C}_T(\omega) = \frac{1}{\tilde{\mathcal{F}}(\omega) + \sigma^2} = \frac{1}{(1/T) \sum_l |\tilde{F}_B(\omega + l 2\pi/T)|^2 + \sigma^2}, \quad (36)$$

and thus the optimum linear receiver has the transform

$$\tilde{G}_{\text{opt}}(\omega) = \tilde{C}_T(\omega) \tilde{F}_B^*(\omega) = \frac{\tilde{F}_B^*(\omega)}{(1/T) \sum_l |\tilde{F}_B(\omega + l 2\pi/T)|^2 + \sigma^2}. \quad (37)$$

The final transform of interest is that of the equalized baseband-equivalent pulse, which is

$$\tilde{H}(\omega) = \frac{|\tilde{F}_B(\omega)|^2}{(1/T) \sum_l |\tilde{F}_B(\omega + l 2\pi/T)|^2 + \sigma^2}. \quad (38)$$

Since $\tilde{H}(\omega)$ is real, the real part $h_1(t)$ and imaginary part $h_2(t)$ of its inverse Fourier transform are even and odd functions of time, respectively. Moreover, as $\sigma^2 \rightarrow 0$ it is also clear that $H_{\text{eq}}(\omega) = (1/T) \sum_k \tilde{H}(\omega + k 2\pi/T) = 1$, i.e., not surprisingly, the equalized channel is Nyquist. From (24) to (26) it follows that the minimized MSE is

$$\begin{aligned} \mathcal{E}_{\text{opt}} &= 1 - (\tilde{\mathbf{f}}_B, \tilde{\mathcal{A}}^{-1} \tilde{\mathbf{f}}_B) = 1 - (\tilde{\mathbf{f}}_B, \tilde{\mathbf{g}}) \\ &= 1 - \int_{-\infty}^{\infty} \tilde{F}_B(\omega) \tilde{G}_{\text{opt}}(\omega) \frac{d\omega}{2\pi}, \end{aligned} \quad (39)$$

which can be rewritten as

$$\mathcal{E}_{\text{opt}} = 1 - \int_{-\pi/T}^{\pi/T} \frac{\sum_k |\tilde{F}_B(\omega + k 2\pi/T)|^2}{(1/T) \sum_l |\tilde{F}_B(\omega + l 2\pi/T)|^2 + \sigma^2} \frac{d\omega}{2\pi}. \quad (40)$$

In summary, eqs. (37) to (40) give a complete description of the performance and structure of the optimum linear receiver. The structure is equivalent to an infinitely-long fractionally-spaced equalizer whose taps are spaced close enough to accommodate the bandwidth of the transmitted signal. Finally note that, as expected, the phase characteristics of the channel, including the timing phase, do not enter into the expression for the minimum MSE; thus, the steady-state system performance is independent of these characteristics.

IV. THE FINITE-LENGTH FRACTIONALLY SPACED EQUALIZER

In this section we first describe the steady-state performance of finite-length FSES, and then demonstrate that, even as the noise vanishes, such an equalizer always has a unique tap setting.

4.1 Steady-state performance

Here we consider the mean-squared error of a finite-length fractionally spaced equalizer. The demodulated received signal,* (13), is sampled at the rate $1/T'$, and thus the equalizer input is

$$q(nT') = \sum_m \tilde{d}_m \tilde{f}_B(nT' - mT) + \nu(nT'). \quad (41)$$

We make a slight change in notation by letting \tilde{c}_n denote the complex equalizer taps—thus the transfer function $\tilde{G}(\omega)$ of the previous section is replaced by $\tilde{C}(\omega)$. The equalizer output, which is only needed at the synchronous instants, is given by

$$\tilde{u}(nT) = \sum_{m=-N}^N \tilde{c}_m \tilde{q}(nT - mT'), \quad (42)$$

where the equalizer has $2N + 1$ complex taps. For the finite-length equalizer the MSE is written compactly as

$$\begin{aligned} \mathcal{E} &= E(e_n^2 + \tilde{e}_n^2) = E(|\tilde{u}(nT) - \tilde{d}_n|^2) \\ &= E\{|\tilde{c}' \tilde{\mathbf{q}}_n - \tilde{d}_n|^2\}, \end{aligned} \quad (43)$$

* A passband equalizer, in which the demodulator follows the equalizer, has the same performance against linear distortion as does the baseband equalizer. The passband and baseband equalizer differ in their performance in the presence of phase jitter.

where the tap vector and the delay-line sample vector are given by

$$\begin{aligned}\tilde{\mathbf{c}}' &= (\tilde{c}_{-N}, \dots, \tilde{c}_N), \\ \tilde{\mathbf{q}}'_n &= (\tilde{q}(nT + NT'), \dots, \tilde{q}(nT), \dots, \tilde{q}(nT - NT')), \quad (44)\end{aligned}$$

and the vectors with an asterisk will denote the transposed conjugate vector. Performing the indicated expectation gives

$$\mathcal{E} = \tilde{\mathbf{c}}^* \tilde{\mathbf{A}} \tilde{\mathbf{c}} - (\tilde{\mathbf{c}}^* \mathbf{f}_B + \tilde{\mathbf{f}}_B^* \tilde{\mathbf{c}}) + \sigma_d^2, \quad (45)^\dagger$$

where the $(2N + 1) \times (2N + 1)$ Hermitian channel-correlation matrix, the $(2N + 1) \times 1$ channel vector, and the data power are defined, respectively, by

$$\begin{aligned}\tilde{\mathbf{A}} &= E(\tilde{\mathbf{q}}_n \tilde{\mathbf{q}}_n^*), \\ \tilde{\mathbf{f}}_B &= E(\tilde{\mathbf{d}}_n^* \tilde{\mathbf{q}}_n), \\ \sigma_d^2 &= E(|\tilde{d}_n|^2).\end{aligned} \quad (46)$$

It is interesting to compute the kl th element, \tilde{A}_{kl} , of the channel-correlation matrix; a direct calculation gives

$$\begin{aligned}\tilde{A}_{kl} &= E(\tilde{q}(nT - kT') \tilde{q}^*(nT - lT')) \\ &= \sigma_d^2 \sum_m \tilde{f}_B(mT - kT') \tilde{f}_B^*(mT - lT') + \sigma^2 \delta_{k-l}, \quad (47)\end{aligned}$$

where δ_{k-l} is the Kronecker delta. Note that in contrast to the synchronous equalizer, the channel-correlation matrix is Hermitian but *not* Toeplitz. To explicitly see the non-Toeplitz nature of the $\tilde{\mathbf{A}}$ matrix we can rewrite (47) in the frequency domain as

$$\begin{aligned}& \frac{1}{\sigma_d^2} (\tilde{A}_{k-l} - \sigma^2 \delta_{k-l}) \\ &= \int_{-\pi/T}^{\pi/T} \left\{ \sum_m \tilde{F}_B \left(\omega + m \frac{2\pi}{T} \right) \exp \left[-j \left(\omega + m \frac{2\pi}{T} \right) kT' \right] \right\} \\ & \quad \cdot \left\{ \sum_n \tilde{F}_B^* \left(\omega + n \frac{2\pi}{T} \right) \exp \left[j \left(\omega + n \frac{2\pi}{T} \right) lT' \right] \right\} \frac{d\omega}{2\pi} \\ &= \int_{-\pi/T}^{\pi/T} e^{-j\omega(k-l)T'} \left[\sum_m \tilde{F}_B \left(\omega + m \frac{2\pi}{T} \right) e^{-jm k 2\pi T'/T} \right] \\ & \quad \cdot \left[\sum_n \tilde{F}_B^* \left(\omega + n \frac{2\pi}{T} \right) e^{jn l 2\pi T'/T} \right] \frac{d\omega}{2\pi}, \quad (48)\end{aligned}$$

[†] Recall that for any two complex vectors $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{v}}$, $\tilde{\mathbf{x}}^* \tilde{\mathbf{v}} = (\tilde{\mathbf{v}}^* \tilde{\mathbf{x}})^*$.

and for systems with nonzero excess bandwidth the bracketed terms depend on k and l individually, rather than on $k - l$. Recall that for the synchronous equalizer $T' = T$ and A_{kl} depends only on $k - l$.

For completeness note that the channel vector $\tilde{\mathbf{f}}_B$ has the l th element

$$(\tilde{\mathbf{f}}_B)_l = E(\tilde{d}_n \tilde{q}(nT - lT')) = E(|\tilde{d}_n|^2) \tilde{f}_B(lT'). \quad (49)$$

In terms of the above parameters, it is evident from (45) that the optimum tap setting is[†]

$$\tilde{\mathbf{c}}_{\text{opt}} = \tilde{A}^{-1} \tilde{\mathbf{f}}_B \quad (50)$$

and the minimized MSE is

$$\mathcal{E}_{\text{opt}} = 1 - \tilde{\mathbf{f}}_B^* \tilde{A}^{-1} \tilde{\mathbf{f}}_B. \quad (51)$$

4.2 The adaptive algorithm

As with the conventional passband equalizer,¹⁻² the adaptive control algorithm makes use of the gradient of the sum of the squared in-phase error and the squared quadrature error with respect to the tap weights. Taking these derivatives, and writing the result in complex notation gives the adjustment algorithm

$$\tilde{\mathbf{c}}_{n+1} = \tilde{\mathbf{c}}_n - \alpha \tilde{e}_n \tilde{\mathbf{q}}_n^*, \quad n = 0, 1, 2, \dots, \quad (52)$$

where $\tilde{\mathbf{c}}_n$ is the complex tap vector at the n th iteration, and α is a positive number, called the step size, which affects the algorithm's rate of convergence and the fluctuation about the minimum-attainable steady-state MSE. Note that the algorithm is updated once per symbol interval, but it is conceivable that adjustments could be made more frequently if the mid-symbol output levels can be interpolated reasonably well. Reference 12 gives a detailed analytic and experimental treatment of the convergence rate and some of the dynamic aspects of FSES.

4.3 Does a finite-length fractionally-spaced equalizer have a unique tap setting?

To answer the question posed by the title of this section we return to the baseband data transmission system discussed in Section II. The transmitted spectrum, as shown in Fig. 1, is bandlimited to $(1 + \alpha) \pi / T$ rad/s, where the rolloff factor, α , varies from 0 to 1. From Fig. 2 it should also be evident that when the noise becomes vanishingly small, there is legitimate concern as to what function(s) the equalizer will

[†] A little care must be exercised in differentiating the MSE with respect to $\tilde{\mathbf{c}}$, since $\tilde{\mathbf{c}}^* \tilde{\mathbf{c}}$ is *not* an analytic function of $\tilde{\mathbf{c}}$. The most compact approach is to differentiate the MSE, with respect to the real and imaginary components of $\tilde{\mathbf{c}}$, and to then interpret the gradient as a complex vector.

synthesize in the region $(1 + \alpha) \pi/T < \omega < 2\pi/T$, where there is no signal energy. In terms of the channel correlation matrix given by (47), we note that the matrix A is the sum of two matrices, and, as will be evident from the discussion which follows, the channel-dependent component of A is always positive semidefinite. Since the other component of the channel-correlation matrix, $\sigma^2 I$, is positive definite, then A will also be positive definite, and we can conclude that when there is noise present the optimum tap setting is unique.

We now consider the situation as the noise becomes vanishingly small; clearly from (50), the optimum tap setting is unique if and only if A is nonsingular. A sufficient condition for A to be nonsingular is the nonvanishing of the quadratic form $\mathbf{u}'A\mathbf{u}$, for any *nonzero* test vector \mathbf{u} with components $\{u_i\}$. Let us consider in detail this quadratic form, which we write from (47) as

$$\begin{aligned} \mathbf{u}'A\mathbf{u} &= \sum_{m,n=-N}^N \sum_{n=-N}^N u_m A_{mn} u_n \\ &= \sum_{m,n=-N}^N u_m u_n \sum_{l=-\infty}^{\infty} f(lT - nT') f(lT - mT') \\ &= \sum_{l=-\infty}^{\infty} \left[\sum_{m=-N}^N u_m f(lT - mT') \right]^2 \geq 0. \end{aligned} \quad (53)$$

The above inequality establishes the positive semidefinite nature of the matrix A , and we see from, (53) that $\mathbf{u}'A\mathbf{u}$ can vanish only if *

$$\sum_{m=-N}^N u_m f(lT - mT') = 0, \quad l = 0, \pm 1, \pm 2, \dots \quad (54)$$

If we define the periodic Fourier transform

$$U_{T'}(\omega) = \sum_{m=-N}^N u_m e^{j\omega m T'}, \quad |\omega| \leq \frac{\pi}{T'}, \quad (55)$$

then we can proceed further by noting that

$$\begin{aligned} \sum_{m=-N}^N u_m f(lT - mT') &= \sum_{m=-N}^N u_m \int_{-\infty}^{\infty} F(\omega) e^{j\omega(lT - mT')} \frac{d\omega}{2\pi} \\ &= \int_{-\infty}^{\infty} \left[\sum_{m=-N}^N u_m e^{-j\omega m T'} \right] F(\omega) e^{-j\omega l T} \frac{d\omega}{2\pi} \\ &= \int_{-\infty}^{\infty} U_{T'}(\omega) F(\omega) e^{-j\omega l T} \frac{d\omega}{2\pi} \end{aligned}$$

* The authors gratefully acknowledge discussions with J. E. Mazo which led to this development.

$$\begin{aligned}
&= \sum_k \int_{(2k-1)(\pi/T)}^{(2k+1)(\pi/T)} U_{T'}(\omega) F(\omega) e^{-j\omega l T} \frac{d\omega}{2\pi} \\
&= \int_{-\pi/T}^{\pi/T} \left[\sum_k U_{T'}\left(\omega + \frac{k2\pi}{T}\right) F\left(\omega + \frac{k2\pi}{T}\right) \right] \\
&\quad \cdot e^{-j\omega l T} \frac{d\omega}{2\pi}. \tag{56}
\end{aligned}$$

The right-hand side of (56) is recognized as the sample, at $t = lT$, of a function whose Fourier transform $Z_{\text{eq}}(\omega)$ is contained in the brackets. Now if (56) is to be zero for every value of l , then it must be that the Fourier transform inside the integral vanishes completely, i.e.,

$$Z_{\text{eq}}(\omega) \equiv \sum_k U_{T'}\left(\omega + \frac{k2\pi}{T}\right) F\left(\omega + \frac{k2\pi}{T}\right) = 0, \quad |\omega| \leq \frac{\pi}{T}. \tag{57}$$

For less than 100 percent excess bandwidth, note that only the $k = 0, \pm 1$ terms contribute to the above sum. However, in the nonrolloff region, $|\omega| \leq (1 - \alpha) \pi/T$, only the $k = 0$ term influences the sum. For channels which do not vanish over the entire nonrolloff region, it is clear that for $Z_{\text{eq}}(\omega)$ to vanish it is required that $U_{T'}(\omega)$ vanish at least over the entire nonrolloff region. Since $U_{T'}(\omega)$ is a finite-term Fourier series, it cannot vanish over an interval without vanishing everywhere, which in turn would again make $\mathbf{u} = 0$. Note that if the channel vanished over a portion of the nonrolloff region, then since $U_{T'}(\omega)$ is a finite-term Fourier series, its energy could not be totally concentrated in the region where there was no channel energy. Thus, the solution would still be unique. However, it is worth noting that in the extreme case of 100 percent excess bandwidth, $Z_{\text{eq}}(\omega)$ can vanish. For example, consider a constant $F(\omega)$, with $U_{T'}(\omega) = \cos(\omega T/2)$. It is apparent from (57) that $Z_{\text{eq}}(\omega) \equiv 0$. Thus for a *finite-length* FSE with an excess bandwidth of less than 100 percent, we can conclude that *even as the noise becomes vanishingly small the A matrix is nonsingular and there is a unique optimum tap setting*.

Note that for a finite-length synchronous equalizer where $T' = T$, (57) indicates that since $U_T(\omega + (k2\pi/T)) = U_T(\omega)$, we can conclude that if the folded-channel spectrum does not vanish completely then there is always a unique tap setting. Reference 12 showed that as the equalizer becomes infinitely long there is an infinitude of equalizer tap vectors which achieve the same minimum mean-squared error. Consequently, it is to be expected that a FSE equalizer with a "large" number of tap weights will have many tap vectors which produce essentially the same MSE.

V. PERFORMANCE OF FRACTIONALLY-SPACED EQUALIZERS

To illustrate the advantages of fractionally-spaced equalization over synchronous equalization, a number of computer simulation runs were made for different equalizer configurations and for channel distortions of varying severity. The system tested was the 9.6 kbit/s QAM system, shown in Fig. 3, having a symbol rate of 2400/s and an excess bandwidth of 12 percent, and the transmitted-symbol alphabet, $\{\pm 1, \pm 3\}$. For each run the steady-state mean-squared error was measured after a sufficiently long period of adaptation, and the FSE was of the $T/2$ type.

Amplitude and delay-distortion characteristics are illustrated in Fig. 4 for the three linear channels which were simulated. The "Good" channel is of low distortion, and well within the limits of standard conditioning, e.g., the "Basic" conditioning¹³ illustrated in Fig. 5. The "Bad-Phase" and "Bad-Slope" channels have, respectively, severe phase distortion and severe amplitude distortion, placing these channels just outside the defining boundaries of basic-conditioned channels.

Figures 6, 7, and 8 compare the performance of a 24-tap synchronous (T) equalizer and a 48-tap $T/2$ equalizer on the three test channels. Performance is examined for five timing epochs within a symbol interval, and is measured by the output signal-to-noise ratio, defined as

$$\text{SNR}_{\text{out}} = \log \frac{P_{BB}}{\text{MSE}}, \quad (58)$$

where P_{BB} is the received baseband average signal power (a constant), and MSE is the measured output mean-squared error. The received signal is normalized so that the ratio of the signal power at the output of the receiving filters to the power of the additive noise, at the same point in the system, is 28 dB. Thus if the equalizer could "undo" the channel distortion without enhancing the noise, then the output SNR would be 28 dB. It is apparent that the performance of the fractionally-spaced equalizer is almost independent of the timing epoch, in sharp contrast to that of the synchronous equalizer. This confirms the prediction of the analysis, culminating in expression (40) for the minimum mean-squared error, which is independent of the sampling epoch. It is also significant that the performance of the fractionally-spaced equalizer on the "Bad-Phase" channel is significantly better than that achieved by the synchronous equalizer even for the best sampling phase. The capability of the FSE for phase equalization, before folding the spectrum about the Nyquist frequency, is seen to be an important advantage on channels with severe phase distortion. With the addition of a decision-feedback equalizer (DFE),¹⁴ with feed-

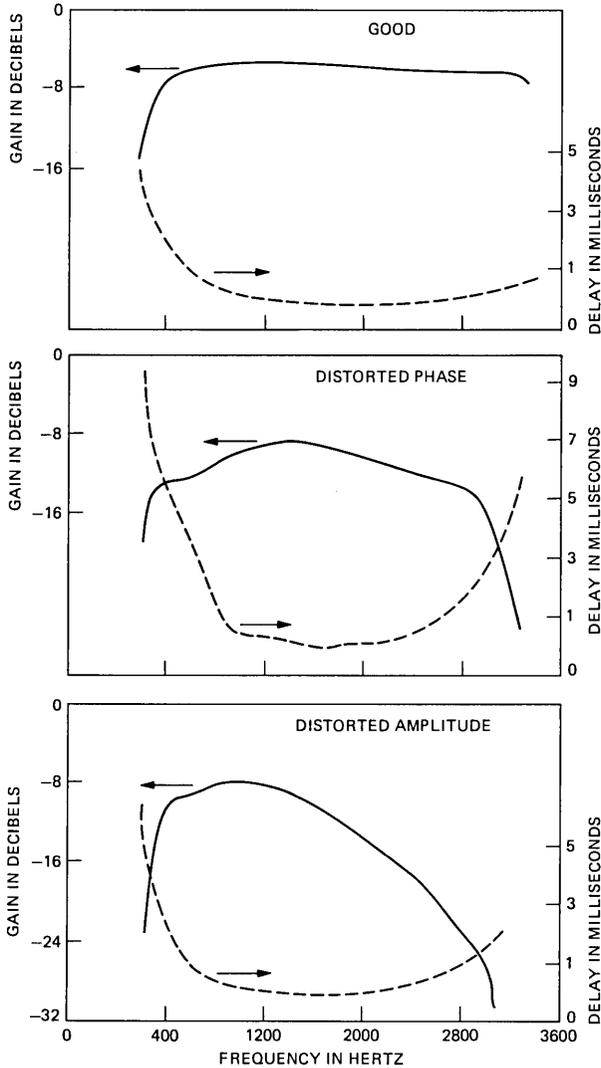


Fig. 4—Characteristics of simulated transmission channels.

back taps $\{f_i\}_{i=1}^N$, shown in Fig. 9, compensation for severe amplitude distortion is also improved, as illustrated in Fig. 8.

The simulation of an FSE with $3T/4$ tap spacing [still less than $T/(1 + \alpha)$, where $\alpha = 0.12$ was the percentage of excess bandwidth] resulted in performance comparable to that of the $T/2$ equalizer. A $3T/4$ equalizer needs only $2/3$ as many taps as a $T/2$ equalizer to span

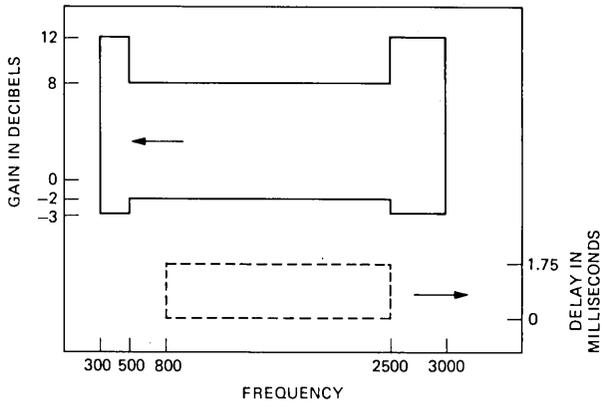


Fig. 5—Defining boundaries of “Basic”-conditioned channels.

a given channel dispersion, which cannot only reduce implementation complexity, but also improve steady-state performance when digital resolution is a consideration.¹²

VI. CONCLUSIONS

We have shown, both analytically and by simulation, that the fractionally-spaced equalizer provides virtual independence from tim-

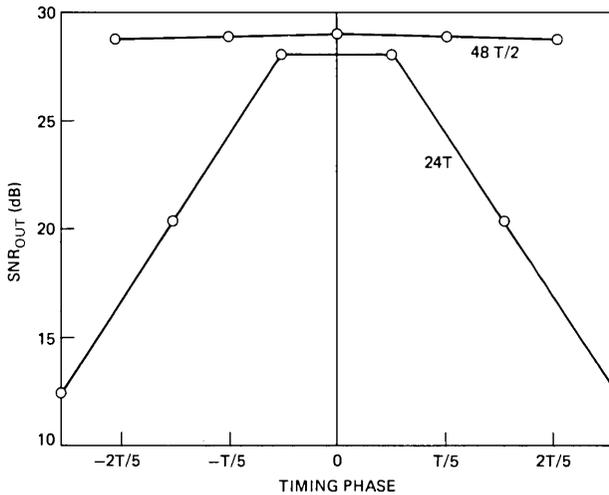


Fig. 6—Performance versus sampling phase of 24-tap synchronous (T) equalizer and 48-tap ($T/2$) equalizer on “Good” channel of Fig. 4. Results from computer simulation of 9600 bps QAM modem.

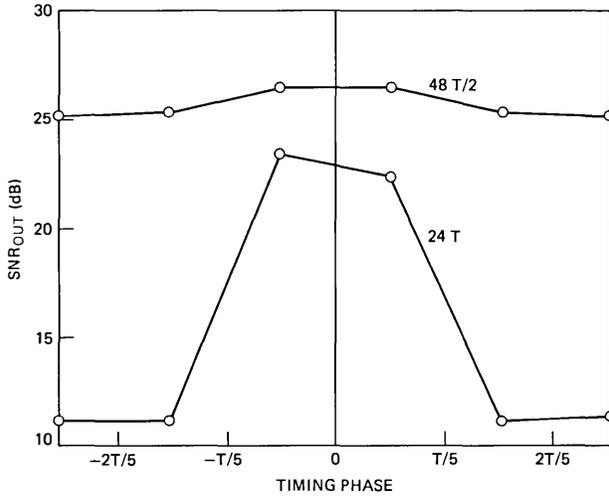


Fig. 7—Performance on “Bad-Phase” channel of Fig. 4.

ing epoch, and significantly improves steady-state performance on severely phase-distorted channels. Implementation of the FSES increased number of taps, with respect to a synchronous equalizer with the same total time span, is well within the capabilities of current digital signal-processing technology. The performance degradation introduced in a digital implementation by using a larger number of taps is more than compensated for by the FSES capability of adaptively

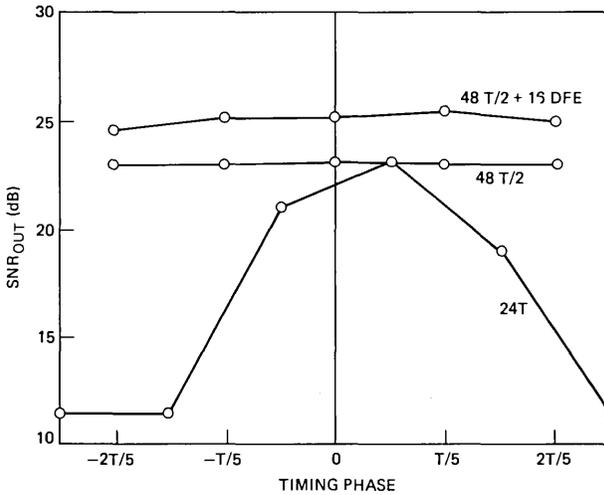


Fig. 8—Performances on “Bad-Slope” channel. The top curve is for a receiver which incorporates both a 48-tap FSE and a 16-tap decision-feedback equalizer (DFE).

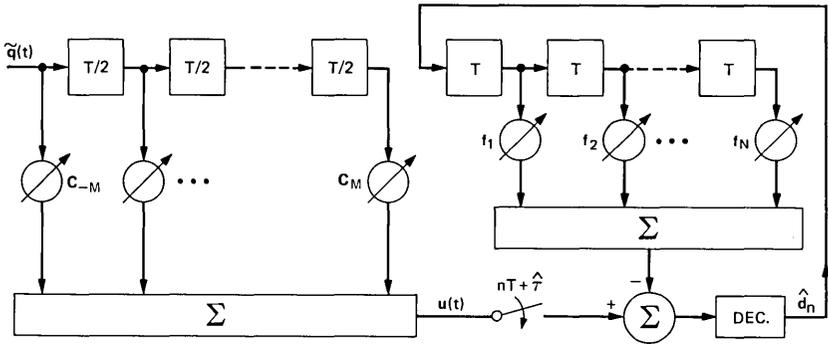


Fig. 9—QAM data receiver combining a fractionally-spaced equalizer and a decision feedback equalizer.

realizing, in one structure, the optimum receiving filter consisting of a matched filter followed by a delay line tapped at symbol intervals.

APPENDIX

Complex Notation and Passband Communication Systems

The purpose of this appendix is to review and organize the compact description provided by complex notation for the discussion of in-phase and quadrature data communications systems. In such a system the transmitted waveform is of the form

$$s(t) = \sum_n a_n p(t - nT) \cos \omega_c t - \sum_n b_n p(t - nT) \sin \omega_c t, \quad (59)$$

where $\{a_n\}$ and $\{b_n\}$ are the in-phase and quadrature data streams, $1/T$ is the symbol rate, $p(\cdot)$ is a bandlimited pulse, and ω_c is the radian carrier frequency. The signal $s(t)$ can be written as the real part of the complex analytic signal

$$\tilde{s}(t) = \sum_n \tilde{d}_n p(t - nT) e^{j\omega_c t}, \quad (60)$$

where

$$\tilde{d}_n = a_n + jb_n.$$

Recall that a signal is analytic if it only has power at positive (or negative) frequencies.¹⁵ The analytic signal with positive frequency content is

$$\tilde{s}(t) = s(t) + j\check{s}(t), \quad (61)$$

where $\check{s}(t)$ is the Hilbert transform of $s(t)$. We now describe the output signal when $s(t)$ is transmitted through a linear passband channel that is band-limited with impulse response

$$x(t) = 2x_1(t) \cos \omega_c t - 2x_2(t) \sin \omega_c t, \quad (62)$$

where $x_1(t)$ and $x_2(t)$ are real. If we denote the Fourier transform of $x(t)$ by

$$X(\omega) = |X(\omega)| e^{j\angle X(\omega)}, \quad (63)$$

then the in-phase pulse response, $x_1(t)$, and the quadrature pulse response, $x_2(t)$, are given by

$$\begin{aligned} x_1(t) &= \int_{-\omega_c}^{\infty} |X(\omega + \omega_c)| \cos[\angle X(\omega + \omega_c) + \omega t] \frac{d\omega}{2\pi}, \\ x_2(t) &= \int_{\omega_c}^{\infty} |X(\omega + \omega_c)| \sin[\angle X(\omega + \omega_c) + \omega t] \frac{d\omega}{2\pi}. \end{aligned} \quad (64)$$

Thus, $x_1(t)$ and $x_2(t)$ are determined by the positive spectral content of the real pulse $x(t)$. In general, the baseband pulses $x_1(t)$ and $x_2(t)$ are unrelated [except through (64)], but if the transfer function $X(\omega)$ has even amplitude symmetry and odd-phase symmetry about the carrier frequency, then $x_2(t) = 0$. Also note that, in general, $x_1(t)$ and $x_2(t)$ are *not* a Hilbert transform pair.

Given the in-pulse and quadrature pulses $x_1(t)$ and $x_2(t)$, we define the analytic pulse

$$\tilde{x}(t) = 2[x_1(t) + jx_2(t)]e^{j\omega_c t} \quad (65)$$

$$= 2\tilde{x}_B(t)e^{j\omega_c t}, \quad (66)$$

where, as noted above, the complex baseband-equivalent pulse, $\tilde{x}_B(t) \equiv x_1(t) + jx_2(t)$, is *not* necessarily analytic. The pulse $\tilde{x}_B(t)$ has a Fourier transform $X(\omega + \omega_c)$, $\omega > -\omega_c$; i.e., the transform is the positive frequency portion of $X(\omega)$ shifted down to the origin. The channel output signal, $s(t) \circledast x(t)$, is of course the $\text{Re}[\tilde{s}(t) \circledast \tilde{x}(t)]$, and we have that

$$\tilde{s}(t) \circledast \tilde{x}(t) = e^{j\omega_c t} \sum_n \tilde{d}_n [\tilde{p}(t - nT) \circledast \tilde{x}_B(t)], \quad (67)$$

where we have allowed the transmitted pulse, $\tilde{p}(t)$, to be "complex." By a complex transmitted pulse we mean that the pulse input is a two-dimensional vector and a cross-coupled operation defines the filter; i.e., if the filter input is the two-tuple vector $(z_1(t), z_2(t))$, which we use to define an equivalent complex signal $\tilde{z}(t) = z_1(t) + jz_2(t)$, then the output vector is $\tilde{u}(t) = \tilde{p}(t) \circledast \tilde{z}(t)$, where the outputs, $u_1(t)$ and $u_2(t)$, are the real and imaginary parts of $\tilde{u}(t)$.

At the receiver, coherent quadrature demodulation by $\cos\omega_c t$ and $\sin\omega_c t$, followed by low-pass filtering, provides the in-phase and quadrature signals; these component signals can also be derived by forming the two-tuple vector composed of the in-phase and quadrature signals (the latter signal is simply the Hilbert transform of the received signal),

and “rotating” the two tuple by $\omega_c t$ radians.¹⁶ This latter operation corresponds to simply multiplying (67) by $\exp(-j\omega_c t)$; i.e., the demodulation signal, $\tilde{r}(t)$, is given by

$$\tilde{r}(t) = \sum_n \tilde{d}_n \tilde{f}_B(t - nT), \quad (68)$$

where the baseband equivalent pulse $\tilde{f}_B(t) \equiv \tilde{p}(t) \otimes \tilde{x}_B(t)$.

Thus, given the channel $x(t)$, we see that from a linear distortion viewpoint, the system is characterized by the (equivalent) baseband pulse $\tilde{f}_B(t) = \tilde{p}(t) \otimes \tilde{x}_B(t)$. If the demodulated signal, $\tilde{r}(t)$, is further filtered, or equalized, by a lattice-type filter,[†] $\tilde{g}(t)$, the overall pulse shape will be $\tilde{p}(t) \otimes \tilde{x}_B(t) \otimes \tilde{g}(t)$.

To summarize, we have illustrated the convenience of complex notation for representing in-phase and quadrature passband signals in terms of the equivalent baseband channel and, explicitly, the carrier frequency. Receiver operations of demodulation and lattice equalization are then easy to visualize and are compactly described as complex multiplications and convolution, respectively.

REFERENCES

1. R. D. Gitlin, E. Y. Ho, and J. E. Mazo, “Passband Equalization of Differentially Phase-Modulated Data Signals,” *B.S.T.J.*, 52, No. 2 (February 1973), pp. 219–38.
2. D. D. Falconer, “Jointly Adaptive Equalization and Carrier Recovery in Two-Dimensional Digital Communication Systems,” *B.S.T.J.*, 55, No. 3 (March 1976), pp. 317–34.
3. A. Gersho, private communication, 1968.
4. D. M. Brady, “An Adaptive Coherent Diversity Receiver for Data Transmission Through Dispersive Media,” Conference Record, ICC 1970 (June 1970), pp. 21–35–21–40.
5. G. Ungerboeck, “Fractional Tap-Spacing Equalizers and Consequences for Clock Recovery for Data Modems,” *IEEE Trans. Commun.*, COM-24, No. 8 (August 1976).
6. J. E. Mazo, private communication.
7. O. Macchi and L. Guidoux, “A New Equalizer and Double Sampling Equalizer,” *Ann. Telecomm.*, 30 (1975), pp. 331–8.
8. R. W. Lucky, J. Salz, and E. J. Weldon, Jr., *Principles of Data Communication*, New York: McGraw-Hill, 1968.
9. R. W. Lucky, “Automatic Equalization for Digital Communications,” *B.S.T.J.*, 44, No. X (April 1965), pp. 547–88.
10. D. M. Brady, Adaptive Signal Processor for Diversity Radio Receivers, U.S. Patent No. 3,633,107, January 4, 1972.
11. S. U. H. Qureshi and G. D. Forney, Jr., “Performance and Properties of a $T/2$ Equalizer,” Conference Record, NTC 1977 (December 1977).
12. R. D. Gitlin and S. B. Weinstein, “On the Required Tap Weight Precision for Digitally-Implemented Adaptive Equalizers,” *B.S.T.J.*, 58, No. 2 (February 1979), pp. 301–21.
13. Bell System Technical Reference, “Data Communications Using Voiceband Private Line Channels,” PUB 41004, October 1973.
14. J. Salz, “Optimum Mean-Square Decision Feedback Equalization,” *B.S.T.J.*, 52, No. 8 (October 1973), pp. 1431–73.
15. H. E. Rowe, *Signals and Noise in Communication Systems*, New York: Van Nostrand, 1965.
16. D. A. Spaulding, “A New Digital Coherent Demodulator,” *IEEE Trans. Commun.*, COM-21, No. 3 (March 1973), pp. 237–8.

† By a lattice filter, $\tilde{g}(t) = g_1(t) + jg_2(t)$, we mean that if the in-phase and quadrature inputs are $r(t)$ and $\tilde{r}(t)$, then the filter outputs are $r(t) \otimes g_1(t) - \tilde{r}(t) \otimes g_2(t)$ and $\tilde{r}(t) \otimes g_1(t) + r(t) \otimes g_2(t)$, respectively.

Contributors to This Issue

Anthony S. Acampora, B.S.E.E., 1968, M.S.E.E., 1970, Ph.D., 1973, Polytechnic Institute of Brooklyn; Bell Laboratories, 1968—. At Bell Laboratories, Mr. Acampora initially worked in the fields of high-power microwave transmitters, radar system studies, and signal processing. Since 1974, he has been studying high-capacity digital satellite systems. His current research interests are modulation and coding theory, time division multiple access methods, and efficient frequency re-use techniques. Member, Eta Kappa Nu, Sigma Xi, IEEE.

Ronald J. Canniff, M.S.E.E., 1972, University of Michigan; Bell Laboratories, 1972—. Prior to his development work on the digital concentrator for the SLCTM-96 system, Mr. Canniff spent nearly five years working on voiceband modems for digital subscriber carrier systems. This included work on delta modulation and active filters for the per-channel codecs of the SLCTM-40 carrier system. He is currently supervising a group concerned with digital signal processing functions for a digital switching machine. Member, Eta Kappa Nu, Tau Beta Pi.

Barbara A. Czekaj, B.S. (computer science), 1980, Monmouth College; Bell Laboratories, 1978—. Ms. Czekaj worked part-time for two years in the Radio Research Laboratory, involved in reductions of multipath fading data and computer analyses of digital radio systems. In May 1980 she became a member of the Business Communications Systems Engineering Center. Member, Association for Computing Machinery.

Richard D. Gitlin, B.E.E., 1964, City College of New York; M.S., 1965, and D. Eng. Sc., 1969, Columbia University; Bell Laboratories, 1969—. Mr. Gitlin is supervisor of the Data Techniques Group in the Data Communications Laboratory. He is a member of the Communication Theory Committee of the IEEE Communications Society, edi-

tor for Communication Theory of the IEEE Transactions on Communications, and is a member of the Editorial Advisory Board of the Proceedings of the IEEE. Senior Member, IEEE; Member, Sigma Xi, Eta Kappa Nu, Tau Beta Pi.

Larry J. Greenstein, B.S.E.E., 1958, M.S.E.E., 1961, and Ph.D. (E.E.), 1967, Illinois Institute of Technology; Bell Laboratories, 1970—. Since joining Bell Laboratories, Mr. Greenstein has worked on digital encoding and transmission. His most recent work is in the areas of satellites, mobile radio, and terrestrial digital radio. He currently heads the Satellite Systems Research Department. Member, Eta Kappa Nu, Sigma Xi, Tau Beta Pi; Senior Member, IEEE; associate editor, IEEE Transactions on Communications.

Bruce Hoadley, B.S. (math), 1961, North Carolina State; Ph.D. (statistics), 1965, University of California; Assistant Professor, Statistics, University of California, Berkeley, 1965–1966; Bell Laboratories, 1966—. Mr. Hoadley was initially a member of the Statistics Department. In 1970, he became a supervisor in the Operations Research Department. Since 1975, he has been supervisor of the Assurance and Reliability Theory and Systems Group. Member, American Statistical Association, American Society for Quality Control, Phi Kappa Phi, Tau Beta Pi, Pi Mu Epsilon, Sigma Xi.

John F. Reiser, B.S. (mathematics), 1973, Michigan State University; Ph.D. (computer science), 1977, Stanford University; Bell Laboratories, 1977—. Mr. Reiser is a member of the Interactive Computer Systems Research Department. His research interests include operating systems, programming systems and languages, and algorithms. Member, Mathematical Association of America, Association for Computing Machinery.

Steven B. Weinstein, B.S.E.E., 1960, Massachusetts Institute of Technology; M.S.E.E., 1962, University of Michigan; Ph.D. (E.E.), 1966, University of California at Berkeley; Philips Research Laboratories, Eindhoven, Netherlands, 1967–1968; Bell Laboratories, 1968–1980. Mr. Weinstein's technical interests include data communications, data communications security, and microcomputer systems. Senior member, IEEE.

Papers by Bell Laboratories Authors

PHYSICAL SCIENCES

Accommodation and Formation of $[11\bar{2}1]$ Twins in Co Single Crystals. S. Vaidya and S. Mahajan, *Acta Met*, 28 (July 1980), pp 1123-31.

Alpha-Particle-Induced Soft Errors and 64K Dynamic RAM Design Interaction. R. J. McPartland, J. T. Nelson, and W. R. Huber, *Proc Int Reliability Phys Symp* (1980), pp 261-7.

Annealing Technique for LEC Grown Twin-Free InP Crystals. W. A. Bonner, *J Electrochem Soc*, 127 (August 1980), pp 1798-1800.

Broken Hexagonal Symmetry in the Incommensurate Charge-Density Wave Structure of $2H-TaSe_2$. R. M. Fleming, D. E. Mocton, D. B. McWhan, and F. J. DiSalvo, *Phys Rev Lett*, 45 (August 18, 1980), pp 576-9.

Brownian Dynamics Study of Polymer Conformational Transitions. E. Helfand, Z. R. Wasserman, and T. A. Weber, *Macromolecules*, 13 (May-June 1980), pp 526-33.

Epitaxial Crystallization from the Melt of Polyethylene and Polypropylene on Mica. A. J. Lovinger, *Am Chem Soc, Polymer Preprints*, 21, No. 2 (1980), pp 253-4.

Glide Ahead of Terminating $[10\bar{1}2]$ Twins in Co. S. Vaidya and S. Mahajan, *Scr Met*, 14, No. 6 (1980), pp 623-6.

Kinetics of Conformational Transition in Chain Molecules. J. Skolnick and E. Helfand, *J Chem Phys*, 72 (May 15, 1980), pp 5489-5500.

Low-Frequency Dynamics of a Cooperative Jahn-Teller Transition: $TbVO_4$. R. T. Harley, K. B. Lyons, and P. A. Fleury, *J Phys C*, 13 (1980), p L447.

Molecular, Electronic, and Crystal Structure of Naphtho[1,8-cd:4,5-c'd']bis[1,2,6]selenadiazine. A. Gieren, V. Lamm, R. C. Haddon, and M. L. Kaplan, *J Am Chem Soc*, 102 (July 16, 1980), pp 5070-3.

A Note on the Effect of Incident Beam Convergence on Quantitative Electron Energy Loss Spectroscopy. D. C. Joy, D. M. Maher, and R. C. Farrow, *Proc Microbeam Analysis 1980*, 1 (August 1980), pp 154-7.

Observation of Direct Excitation of High Orbital Angular Momentum High Rydberg LEDs by Threshold Energy Electron Collision. S. Tarr, J. Schiavone, and R. Freund, *Phys Rev Lett*, 44 (June 23, 1980), pp 1660-3.

Photoelectron Spectrum, Including that of Auger Electrons, of Fano Resonances in Atoms. Y. Yafet, *Phys Rev B*, 21 (June 1, 1980), pp 5023-30.

Picosecond Optoelectronic Detection, Sampling, and Correlation Measurements in Amorphous Semiconductors. D. H. Auston, A. M. Johnson, P. R. Smith, and J. C. Bean, *Appl Phys Lett*, 37 (August 15, 1980), pp 371-3.

The Preparation of Hindered Cuprates from Aldehyde Tosylhydrazones. S. H. Bertz, *Tetrahedron Lett*, 21 (1980), pp 3151-4.

Role of Twinning in Cavitation Erosion of Cobalt Single Crystals. S. Vaidya, S. Mahajan, and C. M. Preece, *Met Trans*, 11A (July 1980), pp 1139-50.

MATHEMATICS

The Complexity of Coloring Circular Arcs. M. R. Garey, D. S. Johnson, G. L. Miller, and C. H. Papadimitriou, *SIAM J Algebraic Discrete Meth*, 1 (June 1980), pp 216-27.

ENGINEERING

- Block Coding of Graphics: A Tutorial Review.** M. Kunt and O. Johnsen, *Proc IEEE*, 68 (July 1980), pp 770-86.
- Coding Method for Vector Representation of Engineering Drawings.** K. Ramachandran, *Proc IEEE*, 68 (July 1980), pp 813-7.
- Electron Channeling Patterns in the SEM.** D. C. Joy and R. C. Farrow, *ISI Electron Opt News*, 1 (July 1980), pp 4-10.
- An I.C. Design Station Needs a High Performance Color Graphics Display.** N. Weste and B. Ackland, 17th Design Automation Conf (June 1980), pp 285-91.
- Planar Electrically-Symmetric n-Way Hybrid Power Dividers/Combiners.** A. A. M. Saleh, *IEEE Trans Microwave Theory Tech* (June 1980), pp 555-63.
- Planar, Multiport, Quadrature-Like Power Dividers/Combiners.** A. A. M. Saleh, *IEEE Int Microwave Symp Digest, MTT-S* (May 1980), pp 483-6.
- Real Time Animation Playback on a Frame Store Display System.** B. Ackland and N. Weste, *Siggraph '80 Conf Proc* (July 1980), pp 182-8.
- Short-Circuit Memory in Electrochromic Displays.** G. Beni, *Appl Phys Lett*, 37 (July 1, 1980), pp 106-8.
- Simple Flying Spot Scanner for Electron Beam Lithography on a Scanning Electron Microscope Without Beam Blanking Capability.** P. Grabbe, *Rev Sci Instrum*, 51 (July 1980), pp 992-4.
- Single-Mode Stabilization by Traps in Semiconductor Lasers.** J. A. Copeland, *IEEE J. Quant Electron*, QE-16 (July 1980), pp 721-7.
- Theorems on Match and Isolation in Multiport Networks.** A. A. M. Saleh, *IEEE Trans Microwave Theory Tech*, MTT-28 (April 1980), pp 428-9.
- Warp Reduction During Wave Soldering.** R. J. Waltner, *Conf Proc Inst Interconnecting Packaging Electron Circuits* (April 1980), pp 1-19.

SOCIAL AND LIFE SCIENCES

- Application of Clustering Techniques to Speaker Independent Word Recognition.** S. E. Levinson, L. R. Rabiner, A. E. Rosenberg, and J. G. Wilpon, *IEEE Trans Acoustics, Speech, and Signal Processing*, ASSP-27 (April 1979), pp 134-40.
- Cochlear Models: Evidence in Support of Mechanical Nonlinearity and a Second Filter (A Review).** J. L. Hall, *Hear Res*, 2 (1980), pp 455-64.
- Cochlear Models: Two-Tone Suppression and the Second Filter.** J. L. Hall, *J Acoust Soc Am*, 67 (May 1980), pp 1722-8.
- A Conversational Mode Airline Information and Reservation System Using Speech Input and Output.** S. E. Levinson and K. L. Shipley, *Proc 1980 ICASSP*, 1 (April 1980), pp 203-8.
- Frequency Selectivity of the Cochlea for Formant Peaks at High Signal Levels.** J. L. Hall, *J Acoust Soc Am*, 68 (August 1980), pp 480-1.
- Loudness of Noise in the Presence of Tones: Measurements and Nonlinear Model Results.** J. L. Hall and M. R. Schroeder, *Psychophysical, Physiological and Behavioral Studies in Hearing*, edited by G. van der Brink and F. A. Bilsen, The Netherlands: Delft University, 1980, pp 329-32.
- Measuring and Modeling User Satisfaction With Telephone Switching and Transmission Performance.** H. S. Cohen, *Int Symp Human Factors in Telecommunications*, 9th Proc (1980), pp 237-42.
- A New System for Continuous Speech Regulation—Preliminary Results.** S. E. Levinson and A. E. Rosenberg, *Proc 1979 Int Conf Acoustics, Speech and Signal Processing*, 1 (April 1979), pp 239-43.
- Some Notes on Reading: Talkers, Material and Reading Rate.** N. Umeda and A. M. S. Quinn, *J Speech Hear Res*, 23-1 (March 1980), pp 56-72.

MANAGEMENT AND ECONOMICS

- Collusive Behavior in Noncooperative Epsilon Equilibria of Oligopolies with Long but Finite Lives.** R. Radner, *J. Economic Theory*, 22 (April 1980), pp 136-54.

Contents, March 1981

- C. S. Myers,
L. R. Rabiner, and
A. E. Rosenberg** On the Use of Dynamic Time Warping for Word Spotting and Connected Word Recognition
- K. Nassau** The Material Dispersion Zero in Infrared Optical Waveguide Materials
- I. W. Sandberg** On Newton-Direction Algorithms and Diffeomorphisms
- N. Gehani** Program Development by Stepwise Refinement and Related Topics
- O. Johnsen** A New Code for Transmission of Ordered Dithered Pictures
- O. Johnsen and
A. N. Netravali** An Extension of the CCITT Facsimile Codes for Dithered Pictures
- L. D. White,
R. W. Coons,
and R. C. Strum** A 200 Hz to 30 MHz Computer-Operated Impedance/Admittance Bridge (COZY)





THE BELL SYSTEM TECHNICAL JOURNAL is abstracted or indexed by *Abstract Journal in Earthquake Engineering*, *Applied Mechanics Review*, *Applied Science & Technology Index*, *Chemical Abstracts*, *Computer Abstracts*, *Current Contents/Engineering, Technology & Applied Sciences*, *Current Index to Statistics*, *Current Papers in Electrical & Electronic Engineering*, *Current Papers on Computers & Control*, *Electronics & Communications Abstracts Journal*, *The Engineering Index*, *International Aerospace Abstracts*, *Journal of Current Laser Abstracts*, *Language and Language Behavior Abstracts*, *Mathematical Reviews*, *Science Abstracts (Series A, Physics Abstracts; Series B, Electrical and Electronic Abstracts; and Series C, Computer & Control Abstracts)*, *Science Citation Index*, *Sociological Abstracts*, *Social Welfare, Social Planning and Social Development*, and *Solid State Abstracts Journal*. Reproductions of the Journal by years are available in microform from University Microfilms, 300 N. Zeeb Road, Ann Arbor, Michigan 48106.



Bell System