New Trends in Workstation I/O

Jim Beetem
Catherine Smith

1. Introduction

The purpose of this paper is to discuss what we in the Data Comm lab feel are the major trends in terminal and workstation I/O, some new tools we have for dealing with newer devices and connection methods, and how programmers and system managers can best connect and utilize these new devices. We are using the term workstation vs. terminal to emphasize the fact that it is not only terminals which are connected over serial communication lines, but also personal computers, printers, plotters, and black boxes.

The paper is organized into four sections. The first section discusses the current connection methods and how we see these methods evolving. The second section mentions features we think will be added, deleted or changed from current terminal drivers. Its emphasis is on features in the point-to-point terminal drivers, since those features are the ones which have evolved more (vs. being planned). The third section discusses trends in support: tools which we are providing for the user, programmer, and system manager. This secion also is concerned with point-to-point connections only. The last section offers some hints and suggestions on how best to program and/or configure non-traditional devices through terminal drivers on the 3000.

2. Current connections & trends

The current connection methods supported by HP for workstations are: point-to-point (via the ADCC and ATP terminal drivers), multipoint (through multipoint terminals and printers or the 2333A cluster controller), X.25 (through PADs or the 2334A cluster controller). There are other methods available from other vendors, some of which work very well. However since these connection methods have not been tested by HP, they are not supported. These other connection methods will not be discussed.

Workstation communication has evolved from teletype communication. Originally this involved infrequent data transmissions, one character at a time at low line speeds. Speeds have gotten faster and workstations much more intelligent which has resulted in new uses and needs for workstation I/O. Applications which do large transfers at line speed such as block mode, downloading, and file transfer are more prevelant. Also, due to the increased amount of data communication in general, the topology of workstation connections has changed. The old choices were either hardwired terminals located close to the system, or terminals connected via modems. A third major change is that the devices being connected over serial lines are different. Instead of terminals, we are finding more personal computers, printers, plotters, and other non-traditional devices. All three of these changes, speed increases, topology changes, and types of devices connected, have put a higher importance on data integrity.

2.1 Point-to-point connections

Many point-to-point terminals connected to a system is the default connection method. Since it is the default connection, almost any application program written for 3000 terminals will run on point-to-point terminals. Connection methods for point-to-point terminals are numerous and usually inexpensive. They can be connected through modems or hardwired. With RS-422 and fiber optics multiplexors they can be hardwired at great distance from the system. Modems for async communication are inexpensive. Another big advantage of

point-to-point workstations is that if something goes wrong (because of the terminal, the application, or the user) it only affects the one terminal. Serial printer support is available for point-to-point connections.

Point-to-point connections have several disadvantages. Since they are point-to-point, there is a separate cable for each device. This requires lots of junction panel space at the system and a lot of cost for the cables and installation. The cost problem is more acute if there are a lot of remote workstations, since there must be two modems per workstation. Also point-to-point terminals have no data integrity since there is no error checking other than parity. This is becoming a bigger problem due to the increase in 8-bit terminal types for foreign language and file transfer applications. The only way that point-to-point devices may be used on other than the primary system is via DSN/DS. Also since point-to-point connections are so flexible, there is a larger variety of applications/device running over point-to-point which makes support difficult since there are more corner cases.

One of the new connection methods for point-to-point workstations is through PBXs. The advantage that this connection method has is that the wiring is already in place for voice communication and so can be used for data communication. It also allows for connection to multiple systems and so enables pooling of systems and system resources (such as modem pools). All applications and devices supported by the point-to-point controllers are supported over PBX connections.

Unfortunately, PBX connections are currently high cost compared to traditional point-to-point connections. PBXs have fewer numbers of switch points than logical connections, so if everyone picked up their phone at once, not everyone would get a dial tone. This means that users who don't disconnect the PBX connection after they have finished communicating with the system, are wasting a valuable resource. PBX connections also have no data integrity other than parity, although the error rate is very low.

2.2 Multipoint terminals

Our current solution for solving some of the deficiencies of point-to-point connections has been multipoint. Since multipoint devices are connected on a single line, cabling costs are reduced. This reduction is especially noticable for remote connection. Only one modem on the system side, and one modem for each drop of devices are needed. Multipoint connections can be either asynchronous or synchronous. Synchronous connections are used for remote connections so that the faster synchronous modems may be used. Asynchronous connections are used for local connections via the factory data link. Not only does the data link provide better imunity to noise, but multipoint also has excellent data integrity via CRC checking.

Multipoint also has its disadvantages. Since it does use CRC, there is more overhead than for point-to-point terminals. Also multipoint terminals are inherently block mode, the system does not receive any characters until the user hits enter. Because the command interpreter does not know how to talk to block mode terminals, the multipoint software translates the data to look like it came from a point-to-point terminal. This extra software also adds extra overhead for communicating with multipoint terminals. When a block mode interface like V/3000 is used, no conversion from character mode to block mode is needed, so multipoint runs more effieciently and overhead is reduced.

Since multipoint terminals are inherently block mode, some applications must be changed to run (or to run well) on multipoint terminals. Since multipoint terminals share a single cable, the response time and throughput for one terminal depends on how much of the line bandwidth other devices are using. So what one device does affects more than just itself, it affects all other devices on the line.

Another multipoint connection is via the HP 2333A cluster controller. This allows the user to connect point-to-point terminals to the controller which is connected via multipoint to the system. CRC checking is used to talk to the controller, so data integrity is achieved. Since point-to-point terminals are connected to the cluster controller, any application which can run as terminal type 10 can run on a terminal connected to the 2333A. The 2333A cluster controller allows a wider range of workstations to be connected to the multipoint line.

Workstations connected via the 2333A have some of the same disadvantages as multipoint devices. There is some software overhead on the system to communicate with the 2333A as well as the workstation and all workstations connected to the 2333A are sharing the multipoint line. Also although any software used to run terminal type 10 terminals can be run through the 2333A, no other terminal types are supported and some applications (such as keystroke interaction applications) will not run well.

2.3 Terminals connected over X.25

There are two ways which terminals can be connected over an X.25 packet switched network. The first is via a public PAD (packet assembler disassembler). To use a PAD, the user calls up the local PAD (a local phone call) and then gives the PAD commands to "dial" a

system. This is a very cheap method for long distance communication (if the system end has other uses for the X.25 connection) since the X.25 price is not tied to distance. Instead of being tied to distance, prices for X.25 connections have some fixed cost for connection to the network (small for connections to PADs) and then a very small cost per packet. This connection also is beneficial if the terminal user uses many systems on the network.

On the minus side, PAD terminals have many of the same disadvantages as multipoint terminals. Communication must work its way through the network and so response time is often slow. The response time also depends on the application. An application which runs poorly over multipoint is likely to be worse over an X.25 network. Also all features are not supported over a PAD terminal. Currently the only block mode which can be used is via V/3000. Since terminals are not permanently connected to a system, there is no way for the system to FOPEN a PAD terminal. Since printers can only be FOPENed (they can't logon), printers are not supported through PAD connections.

The second way in which terminals can be connected to an X.25 network is via the HP 2334A cluster controller. This connections method has many of the same advantages and disadvantages as a terminal connected to a public PAD.

Workstations connected to the HP 2334A cluster controller provide a cheap long distance solution if there are many workstations connected at one location. There is an additional expense since the 2334A is connected to the X.25 network as a system, not a PAD. This adds a larger monthly expense in this country, but not much more in Europe. Printers and terminals may be FOPENed since the 2334A cluster controller is connected as a system (a permanent connection).

Terminals connected through the 2334A have the same response time limitations as do terminals connected through a public PAD since data must work its way through the network and the system software. As with PAD terminals, V/3000 is the only block mode supported for terminals connected. The cost of 2334A terminals over PAD terminals only becomes favorable to 2334A terminals if there are a lot of terminals at one location.

3. New Features

As mentioned earlier, there are several trends apparent in the terminal I/O area which must be reflected in the design of future HP systems. First, to handle the demands of personal computers and intelligent workstations, new systems will need I/O facilities that handle data in large blocks with "perfect" data integrity. Second, the traditional asynchronous terminal I/O port will be used to connect an increasing variety of devices, each requiring a different set of characteristics. Many of these devices will not be compatible with HP terminals. Third, as the price of high qulity, high speed printers decreases, HP 3000 customers wish to move these devices closer to the user's work area. Terminal I/O ports (either point-to-point or multipoint) provide the only cost effective solution to this need. Fourth, some of the facilities that are now widely used; many of these facilities will not be implemented on future I/O

In the following section the discussion is aimed at point-to-point terminal connections. Other terminal connections either have solved the problems or the problem never existed for that connection method.

3.1 A Reliable Link to Work Stations

Personal computers and intelligent workstations often move entire files of data to or from the HP 3000, manipulate the contents and move it back. Their primary concern is that a large amount of data be moved quickly and accurately between the two machines. This is a new requirement in the HP 3000 Terminal I/O environment.

The current terminal I/O interface on the HP 3000 was designed for use with simple "TTY-compatible" terminals. It is intended to transfer one character at a time at human typing speeds inbound and to print data on the display at human reading or scanning speeds. HP has enhanced this basic design to handle VPLUS and HPWORD block mode terminals with reasonable reliability and efficiency. However, the new devices demand much higher thruputs and require more efficient operation to avoid degrading system performance.

The primary source of data integrity in the existing design is inherent in the physical facilities. Within their speed/distance limitations, the RS-232-C and RS-422 standards offer a very high quality link to local area users. Where telephone lines must be used, only parity checking is availale. In the future, HP is committed to using the 8th bit for data not parity, so even this facility will not be available.

The solution to these requirements is an efficient link protocol. It should allow movement of masses of data at close to the capacity of the data comm line connecting the machine. The protocol should use an error detecting mechanism like Cyclical Redundancy Checking (CRC) and automatically trigger retransmission of blocks containing a error. LINK/100, a product that transfers data between the HP 3000 and the HP 100 series of personal computers is a step toward an efficient and reliable workstation link.

3.2 Broader "Connectivity"

On current HP 3000 systems, each terminal I/O port must be configured as a specific "terminal type". The "terminal type" describes the set of facilities that are operational for that port. Any terminal connected to that port must be compatible with those facilities. For example, any terminal connected to a port with the ENQ/ACK handshake enabled must be able to perform that protocol.

HP provides point-to-point terminal types that allow our software packages to operate HP terminals properly. However, customers developing their own software or wishing to connect non-HP devices to the HP 3000 often have problems finding a terminal type that exactly fills their needs.

The solution to this problem is a new HP product, the Work Station Configurator (WSC). This product allows customers to specify their own terminal types as files that reside on the disc. These user-defined terminal types may be configured as the default terminal type in SYSDUMP, specified as Log-on terminal type in the :HELLO command, or changed programmatically at run time. Users cannot add new capabilities to the existing terminal I/O interface, but they can selectively enable or disable almost all of the available facilities.

The new terminal types that can be created by this product (and by the associated enhancements to MPE) allow users with unusual needs to tailor the actions of the terminal I/O interface to their specific device. For example, none of the HP-supplied terminal types allow V/3000 to by used without the ENQ/ACK flow control handshake. A user-defined terminal type can eliminate this restriction, allowing more efficient operation of statistical multiplexers. For printers, the FOPEN-time Initialization string that is written to the device may re-defined, allowing users to initialize the device to any configuration of margins, character size, paper size, pitch, etc that the device supports. Since the default terminal type in the I/O configuration can be temporarily changed by the ;ENV= clause in the :FILE command, subsequent jobs on a spooled device are not affected by the unusual requirements of a single job.

As new capabilities are added to the terminal I/O interface, control of the capability will generally be available to users thru the WSC facility.

This product is intended for knowledgeable users. The WSC allows users to specify any set of the available features. Many of the sets would not be satisfactory for general use. There is no method that the WSC product can use to detect possible problems. For example, a user-defined terminal type with all flow control mechanisms disabled is likely to suffer data overruns at high line speeds. A terminal type that sets the ENQ/ACK handshake block size to 10 characters will operate inefficiently, especially over statistical multiplexers. The detection and prevention of such problems with user-defined terminal types is the user's responsibility.

3.3 Improved Printer Support

As the price of high quality, high speed printers continues to decline, more and more HP 3000 customers want to move these devices into the user's work area. The only economical method of connecting these devices to the HP 3000 is via one of the serial I/O facilities.

Between the point-to-point I/O port and the device, only the physical quality of the data comm line is available to insure data integrity, so that connection of printers must be limited to the distance supported by the Level 1 standard, either RS-232-C or RS-422 (ATP only). Special care must be taken in electrically noisy environments to prevent data errors.

The only exceptions to this restriction are devices that operate as terminal types 19 and 21. These are 7 bit devices that always operate with parity checking enabled, allowing them to detect but not correct data transmission errors. Multipoint printers obtain data integrity through the CRC checking inherent in the multipoint protocol.

HP now supports terminal type 18 for connection of printers that depend only upon the quality of the data comm line for data integrity. The use of terminal type 18 devices across phone lines with modems is NOT supported, since their quality is never high. In many cases there is one error for each one hundred characters transmitted.

3.4 Obsolete Features Not Expected To Be Supported

Many of the features of the existing terminal I/O interface were added to support a specific device; when the device was no longer sold and the need for the feature went away, the feature stayed on. Many of the current features were intended to support simple TTY terminals

and are rarely used today. Implementing these features in the ATP added significantly to its cost; these same features are a major contributor to the data overrun problems experienced by the ADCC at high line speeds. Most of these features will not be supported on the next generation of terminal interface products.

a. Control of Echo from the terminal: typing (ESC) ; and (ESC) : allow the user to enable or disable the input echo facility. Scanning the input stream for the (ESC) character is a major problem. It is very unlikely to be available in the future. Programmatic control will be available; if sufficient demand is present, an MPE command like the :SPEED command could be implemented.

b. Line Feed triggers output of Carriage Return: to facilitate input on certain devices, the HP 3000 will echo two characters, Line Feed and Carriage Return, in response to a single Line Feed character. This feature also requires scanning the input data stream.

c. Substituting Line Feed for Form Feed in Output streams: this feature handled a device that didn't recognize the Form Feed character. That device (the HP 2615) has not be sold since 1974; it will not be supported on future interfaces.

d. Echo the End of Medium character when Backspace is received: another feature unique to the HP 2615 terminal.

e. X-ON/X-OFF Disabled during FREADs: this "feature" of the ATC and ADCC interfaces was not implemented on the ATP and will not be on future interfaces. For the X-ON/X-OFF protocol to operate properly, the X-ON and X-OFF characters must always be special characters that start and stop the flow of data. When it is necessary to pass "binary data", ie. any one of the 256 possible 8 bit characters may exist in the data stream, some other flow control protocol must be used. Defining and implementing such a protocol must be done by the user.

f. Input Data Saved over a Break: at present, a user can type an MPE command, type BREAK, type some other MPE commands (like :FILE), type :RESUME, then type Return and the first MPE command is now executed. It was saved during all of the subsequent transactions. This facility is largely unknown by users and probably will not by implemented on future interfaces.

g. TTY Delays: early terminals were not able to process all characters in "real time" and had no buffering; this required a short delay following these characters. The HP 3000 inserts a short delay after Carriage Return, Line Feed and Form Feed for some terminal types. Support of this feature requires scanning the output data stream, adding to system overhead and slowing the output data rate. This feature will not be supported on future controllers.

h. Slow Line Speeds: current terminal interfaces support most of the following line speeds: 110, 150, 300, 600, 1200, 2400, 4800, 9600 and 19200 bits per second (bps). Future interfaces are unlikely to support speeds of 110 and 150. Support of 300 and 600 bps is probable, but as higher speed modems become cheaper and available throughout the world, support of these speeds becomes less likely.

i. End of File Indicator: typing the string :EOF: from your terminal causes the File System to detect an end of file (EOF) condition on the input file. If you are logged on, the EOF occurs on $STDIN and causes your program and session to terminate immediately. This feature does not appear to be used and adds slightly to the overhead processing of all FREADs. It is likely to be eliminated in future implementations.

4. Trends in Support

There are two new facilities available for point-to-point terminal I/O users that are typical of the support facilities that will be available in future products. The TERMDSM (or TERMinal Diagnostic/Support Monitor) is an on-line support tool for use with the ATP controller (it was named ATPDSM in MPE IV). It will be extended to include the ADCC hardware and the new ADCC driver (HIOTERM2 and HIOASLP2) when MPE V/E is available. A new Point-to-point Workstation I/O Reference Manual will be available with MPE V/E, greatly expanding the information available on users of point-to-point terminals.

4.1 TERMDSM

This support tool is designed for use with the ATP hardware and the HIOXXXX1 software and the ADCC hardware with the HIOXXXX2 software. The facilities available are similar to those included with the ATPDSM (which it replaces).

This On-line tool allows the system manager, system operator and diagnostician to handle many hardware and software problems without halting normal operation of MPE. For example, the new driver software does not call the System Failure routines for most software

errors; instead, the state of the single terminal port affected by the error is dumped to an extra data segment, the port is marked "Broken" in the Device Information Table (DIT), a message is printed on the system console and operation of the port is suspended. As a result, that single port is "broken", but all of the remaining ports can continue normal operation.

To "repair" the broken port, the System Manager or Operator should first use the DUMP facility in TERMDSM to print the data dumped to the extra data segment and then mail it with a software service request to his local HP Service office. (This report allows the factory to find and fix the software problem.) Then, the broken port can be RESET, aborting the user's session, and the user can log on again to resume processing. In a few cases, the RESET fails. This most often happens because the "critical" bit is set for this process. When this happens, a WARMSTART is required to clear the port.

The DIAGnostic facilities of TERMDSM allows some hardware diagnostics to run on-line, affecting only as many users of the system as absolutely necessary. (No on-line diagnostics exist for the ADCC hardware.) The loopback capabilities allow troubleshooting of intermittent data loss or garbling problems.

The DISPlay command allows knowledgeable users to look at a snapshot of the MPE I/O tables that relate to a terminal port. This information can be very valuable to help determine why a port or process appears to be "hung".

There are several contributions that this tool makes. First, it allows HP to obtain data to fix software problems without serious inconvience to our customers. Second, it allows customers to fix simple software problems themselves, without waiting for an HP Customer Engineer to come to the site. Third, the diagnostic and display facilities allow troubleshooting of terminal I/O problems in either HP hardware or software or in customer applications without bringing down the entire system.

This tool is typical of the approach HP expects to take in future products, where facilities are provided that allow the customer to do more of the system support work, lowering the cost of support by minimizing the need for expensive on-site visits by HP support personnel.

4.2 New Point-to-point Workstation I/O Manual

The new Point-to-Point Workstation I/O Reference Manual (30000-90250) will be available with MPE V/E. It expands greatly the previously available information on the use of the MPE intrinsics to perform terminal I/O. It describes in detail what many of the intrinsics do and gives hints on how to accomplish many tasks within the constraints imposed by MPE.

5. Programming for new workstations

One of the problems mentioned for a lot of the terminal connections was response time for some applications. The problem is that connections with good error correction (currently multipoint and X.25 but there will be more coming) have a fixed overhead (both CPU overhead and bandwidth use) per block of data sent. When an application sends a fixed amount of data in several small writes (less than 20 characters), the performance perceived by the terminal user is much worse than if the application sent it all at once. What is even worse are applications which do single character reads. These types of applications can be made very friendly on a single user system. The user sees immediate response to any input. Unfortunately on either a heavily loaded system or over a network, the user has to wait for every keystroke to be handled by the system or sent through the network.

Another common problem is data overruns on the system side. The flow control for character mode devices was designed when the problem was sending too much data to a terminal. With personal computers becoming more prevelant, there is a need for stopping the device from sending too much data to the system. Input from terminals in character mode was always seen to be dependent on how fast the user could type. Now some terminal emulators or personal computers send large amounts of data at line speed. This can overflow the controller (the point-to-point controllers in the system or the 2333A and 2334A cluster controllers). This can be avoided by using the block mode handshake if the transfer is not a binary transfer.

Another reason for data overruns on some of the point-to-point controllers is due to special characters. The character which presents the most difficulty is ESC. Since ESC ;/ESC : is used to turn on or off echo, if an ESC is received, the controller must check the next character to see if it is a semicolon or a colon.

If the transfer is binary, the block mode handshake characters will be treated as normal data (that is the purpose of binary mode). Binary mode is needed for file transfers, so the only way to not overrun the controller is to make sure that no more 128 characters are sent in one block.

Another problem we had been seeing a lot of is devices which do not fit any of our supplied terminal types. If a user wanted to connect some black box that didn't do ENQ/ACK handshake but did need the DC1 read trigger, there was no way to do it. This last problem has been solved with the new Work Station Configurator user definable terminal types.

Jim Beetem is a project manager in IND's Data Comm lab where he was responsible for the ATP product and the INP software. His undergraduate degree is in Mechanical Engineering from MIT; he has an MBA from Stanford where he has done graduate work in Computer Science. Jim is married and has two small children, who occupy most of his "free" time.

Catherine Smith is also a project manager in IND's Data Comm lab. She joined HP in 1978 after receiving her BS degree in Electrical Engineering and Computer Science. She has worked on software for multipoint terminals, point-to-point terminals, and the 2333A cluster controller. Catherine, her husband, and her teenage son spend as much time as possible skiing.
