

INTRODUCTION TO THE PROGRAMMER'S GUIDE TO VALDOCS

Developing the integrated software environment of Valdocs (Valuable Documents) has been a process of evolution. From the initial concepts of the HASCI (Human Application Standard Computer Interface) keyboard, to the presentation of those interface concepts on the Epson QX-10, we've been guided by the belief that a comfortable, straightforward design was essential in order to make the computer a tool of man, rather than the other way around.

This manual documents Valdocs release version 1.18 and is intended as a *preliminary* document, illuminating the mysteries of Valdocs and the integrated software within it. This is not a manual for first-time users, nor does it attempt to teach the art of programming to more sophisticated users. The following chapters provide specific information necessary for the experienced programmer to understand WHAT is being done WHERE, and HOW to use the information for the application programs being integrated into the Valdocs environment.

The manual assumes the reader has sufficient experience with TPM-II or CP/M 2.2 to realize the similar nature of the two operating systems and appreciate the differences. TPM, the operating system upon which Valdocs is built, is completely CP/M compatible, and any program running with CP/M 2.2 (or 1.4) will run on the QX-10 through Valdocs. Indeed, programmers should feel free to take full advantage of the ability to shuttle between the two; the extended capabilities of TPM-II will enhance a program's response within the Valdocs environment.

The first part of the manual discusses the conceptual basis of HASCI and Valdocs, and introduces the resident modules, as they are loaded from the disk.

A cursory description of the TPM-II operating system follows--how it is loaded, how memory is organized, how the Zapple Monitor program functions as both I/O handler and debugging tool for the programmer--as well as a description of all of the function calls in both TPM-II and CP/M modes. (This is not, however, the TPM-II manual for learning about TPM-II or operating systems. A TPM-II manual is available either through your Epson dealer or Rising Star Industries.)

An entire document, Sysinit Functions, is devoted to a discussion of the Call 40 Multibank Functions, and the specific capabilities and command parameters that allow use of "back door" subroutines, spooler routines, and interrupt routines.

An explanation of the Video Driver Program, accessed by Function Call 39, follows. This section details the D-opcode system (Drawing Operation Codes), describes the specific Drawing Operation Codes, and tells how to handle screens and manipulate video operations on the QX-10. A brief discussion of character fonts and how they are created is included.

Next, we describe the Valdocs Indexer Program and the Editor, providing the file handling specifications for a Valdocs file, along with the incantations and imbedded sequences used by the Editor to work its magic.

Applications programs, such as the Menu, Scheduler, Calculator, Draw, Print, and Mail, are described in chapters that follow.

Finally, we have provided several important charts, such as the RS232 Pin Assignments, CMOS Clock RAM, and Keyboard Assignments Chart, in the Appendixes.

We reiterate, this manual is in its early stages. Evolution does not stop, and neither have we. We mean to do our best to continue to provide as much information as possible, as soon as possible. Our aim is to improve the Valdocs Programmer's Manual to a point where we can do for manual users, what Valdocs has done for computer users.

Welcome to our machine.

Chapter I

HASCI--THE HUMAN APPLICATIONS STANDARD COMPUTER INTERFACE

By Chris Rutkowski
President - Rising Star Industries

Development and Theory

Many people see the personal computer as merely a cheaper, smaller, and slower version of its larger data-processing relatives. However, it's becoming apparent that the personal computer is an entirely different type of machine, shaped by a technological evolution that should result in computers that work for people, rather than the other way around.

The proposed Human Applications Standard Computer Interface (HASCI) was designed as an important step in that evolutionary process. It is the result of approximately six years of effort, proceeding from the most general considerations to a very specific result. First we'll examine the theory and principles behind the HASCI interface--we'll learn why the interface is needed and what it is generally intended to do. Next, we'll look at the actual implementation and design specifications of the interface.

Theoretical Background

I entered the microcomputer marketplace in 1975, during the infancy of our industry. Then, as now, those of us on the "inside" saw visions of microcomputers gracing every desk in the world, when the industry grew up.

Then as now, the consensus of opinion within the industry was that the microcomputer would be the bright star of the future. We knew it was so, but we couldn't prove it; therefore, financial backing was hard to come by. It is easy to forget that, in 1975, the microcomputer was not yet the darling of the venture-capital set; Wall Street had taken a bath on computer companies just a year or two earlier during a recession, and our claims to have found a magic formula for success fell upon jaundiced ears. The one precept on which everyone seemed to agree was that *no one could predict such a fast-changing market more than a year or so in advance.*

In the intervening years, I've heard that phrase a hundred times or more; I suspect you have too. It's one of those pieces of common wisdom that sounds good in a speech and makes for good press: the media repeat it, the bureaucrats who read the media repeat it, and the media

repeat it again. This sort of publicity is discouraging. Nevertheless, with blinders firmly in place, enterprising companies continue the struggle to design their way into a murky future.

The Challenge Accepted

In 1976, during the first Atlantic City Computer Show (thank you, John Dilks, for your vision), the "can't predict" motto rang loudly in my ears. It was plain at the time that if our industry was to put a microcomputer in every home, the two essentials were money (lots of it) and manufacturing capability. It was also plain that prediction precedes production; no company executives in their right minds would put up the megabucks necessary to develop the microcomputer without knowing where that development would lead. Tooling for extreme mass production costs millions; for that kind of investment a one-year prediction lead was far from adequate.

Thus it was not until 1981 that IBM entered the personal computer market. It is to that company's credit that its machine avoids most if not all of the inanities perpetrated by IBM's peers. Witness the pitiful efforts of most minicomputer companies to introduce personal computers over the last few years; most if not all of these machines were obsolete before the first carton was shipped. The only prediction those companies could make was that their profitability would plummet within a few years if they couldn't penetrate the microcomputer field. And in fact, this has come to pass.

On the subject of market predictability, many heated discussions took place comparing various hardware and software components, but I realized that further arguments on the advantages of one processor over another, one operating system over another, or one language over another were wasted words unless you knew how those items related to the evolutionary path of the industry--the yardstick for measuring potential worth. And I took it upon myself to research the question of prediction.

Research Methodology

I chose a most unscholarly methodology, but one well suited to the task. Rather than dig through stacks of ponderous marketing tomes in dusty libraries and research what had already been done, I reasoned that any worthwhile work was probably buried so deep as to be invisible. After all, if viable principles of predictability (in terms of the computer market) were available, why weren't they in use? I therefore decided to conduct a broad survey of earlier technological industries, narrowed down to those that had reached the mass consumer markets.

I scanned the marketing history of twentieth-century Western civilization, seeking instances where highly technical products were converted into mass-market commodities over a relatively short period of time. If you think this through yourself, you'll find several examples, including radio, television, electric lightbulbs, and of course the automobile.

I soon perceived a pattern in the emergence of these products that either had gone unnoticed before or had been erroneously classified as unimportant. To illustrate this, let's consider how one such product evolved.

Case Study: The Automobile

I ask you to turn your mental clock back to the year 1905 and consider the state of the automobile market at that time.

First, the automobile was nowhere near mass production yet. Most manufacturers were backyard experimenters (I suggest that the phrase "garage shop" must have originated somewhere around here). They were technology freaks working on the hottest gadget then conceivable.

Peruse some of the popular literature of the time; items about the coming wave of horseless carriages abounded. There were literally hundreds of fledgling manufacturers--every bicycle and carriage shop fancied itself to be the next Pullman Company (the coach manufacturer that became very successful making railroad cars). And what cars they made! Although most had four wheels, their similarity to the automobile of today stops there. Some of these contraptions were steered with tillers like a boat, while others had reins like a wagon. A few had three wheels. They had handbrakes on the right and footbrakes on the left; fixed throttles and throttles on the dash. Few if any were closed in with a roof. And not one was truly practical for the average person. (Does this sound familiar?)

It's easy to look back at these early machines and say, "How quaint." It's easy to overlook the fact that every single engineer and user had his or her own idea of perfection. Ideas abounded, and while each no doubt had some validity, no one could agree on what was valid and what wasn't. In modern terminology we would say that the engineers were coming up with possible *design elements* that were combined almost at random into *architectures* (a collection of design elements).

Now turn your mental clock forward to 1925, and consider again the state of the automobile. Things had definitely changed. The auto was in mass production. Hundreds of thousands per year were being added to a blossoming economy.

And more important, we find that every car on the road had a steering wheel and a throttle, brake, and clutch on the floor. It had windshields and headlights. We find that, with the exception of a relatively few details, you would be able to climb into the typical automobile of 1925 and drive it away.

Architectural Stabilization

By 1925, the architecture of the automobile had become standardized. That architecture has not altered significantly in the ensuing 57 years. Today, the products that you see parked on the streets and recognize as automobiles are architecturally identical to each other. No architectural difference exists between a Subaru and a Rolls-Royce.

If you check other technical marketplaces (for example, that of television), you will see that this same phenomenon has occurred. First, independent engineers developed a wide variety of design elements. Then their ideas were assimilated and adapted until, now, the architecture has ceased to change. I call this phenomenon *architectural stabilization*.

In the period following architectural stabilization, the design effort and creativity that were previously engaged in the random creation of architectures is now geared toward the refinement of the design elements that comprise the stabilized architecture.

This point is crucial: a stabilized architecture ends the game of random invention and redirects the tremendous creative energy of engineers to a better-focused goal--the improvement of the design elements. The improvements realized may be quite substantial. For example, consider suspension systems. Before 1925 no automobile had a suspension system truly worthy of the name. In the ensuing years such comfort items evolved beyond all prediction.

Thus we see that while architectural stabilization may seem to limit certain aspects of design, it can and should precipitate a design revolution far more exciting than pure laissez-faire engineering.

The Mechanics of Stabilization

A detailed analysis of the marketing factors that affect each step of a product's life span is beyond the scope of this manual. However, the results of my investigations revealed the following sequence of events leading to architectural stabilization:

*First, engineers (or technical specialists) conceive of a new product class and build it for its own sake.

- *Engineers then use the product.
- *If the product promises to fundamentally revise the quality of life for its users, the number of participating engineers will swell. (They sense the market potential and have visions of wealth and fame.)
- *Eventually, this growing enthusiasm gains popular notice, and certain non-engineers purchase the product. These non-engineers find the architectures designed by engineers to be difficult to use; they recommend improvements *but are willing to undergo difficulty in using the product*. They are "enthusiasts."
- *Increasing demand increases production, which lowers the product's price.
- *People who are not willing to undergo substantial difficulty in using the product purchase it. These users are disappointed by the currently available products. They are consumers--they want the benefits without the difficulties.
- *More communication about the product occurs in the popular media.
- *If the product does not fill a truly fundamental need, its popularity subsides, leaving a core group of enthusiasts that will then grow at a slower rate. The product will show a gradual evolution of architecture across time.
- *If the need for the product is truly fundamental, demand continues to grow, but actual market growth may slacken.
- *This growth of demand (potential market) motivates engineers and enthusiasts to redesign the product to make it easy to operate. In other words, swelling demand precipitates the creation of a human interface that makes the device easy to use.
- *An easy-to-use version of the product finds a ready and willing market.
- *The first manufacturer to implement ease of use soon gains a market edge.
- *Other manufacturers either follow suit or perish.

This sequence, or one closely analogous to it, occurs in the evolution of all product markets. For the microcomputer market, certain factors have become clear. First, the microcomputer market has not yet achieved architectural stabilization. Second, the microcomputer appears to have all the elements necessary to cause architectural stabilization

to occur; that is, its impact on users is of sufficient importance to force stabilization to occur. Third, the microcomputer market has currently reached that step of increased popular demand that should precipitate the development of an easy-to-use version of the product.

It's no accident that human-factors engineering has risen to such prominence over the last year. It is a natural and necessary step in the evolution of the product classification from a *technical specialist's* market to an *enthusiast's* market and finally to a *consumer's* market.

Thus the development of a human interface coupled with mass-production technology should be the key to opening the consumer market for the computer.

Let me digress for a moment to observe that architectural stabilization occurs at many levels of observation, not only with products such as those discussed here but also with subproducts--raw materials and their elemental forms. All undergo microcosmic architectural stabilization. Likewise, stabilization tends to occur in structures far larger than products: nations, families, and businesses. All exhibit variations of this same phenomenon. It thus appears that architectural stabilization is a fundamental mechanism of systems evolution: the imposition of a mutually accommodative interface between two counter efforts, thoughts, forces, or intentions.

In Search of a Human Interface

The many clues that led to the development of our proposed human interface (HASCI) came primarily from fields far removed from the normal realm of computer science. The difficulty was this: before an interface could be designed, the actual relationship of man and computer had to be defined. I had concluded early on that the entire question of artificial intelligence could be ignored in the design of an interface, which was fortunate since no workable definition of intelligence exists. Rather, an interface involves questions of capability: What can people *do*, and what are they *good* at? This approach proved very profitable.

Even if you were offered a million dollars to manually multiply two times two a million times, you would have a very difficult time completing the task; most humans would be psychologically incapable of completing the job. Yet virtually any computer can do it easily and with remarkable speed. Conversely, such problems as "recognize a certain person's voice," solved by almost any infant (especially if the voice belongs to the child's mother) still represent a major challenge to even the finest computers and programmers.

An analysis of these problems suggests that people are much better than computers at recognizing patterns, while computers are much better than people at manipulating symbols.

Following this logic, the ideal relationship of computer and user should involve the computer as a symbolic manipulator and the user as a pattern recognizer.

This explains the overwhelming popularity of word processors and spreadsheet calculators. One manipulates words and letters, the primary symbols of man. The other manipulates numbers, man's second most important symbol set.

It follows that a complete computer for the typical user should provide the facilities for manipulating all the primary symbols of man (words and letters, numbers, general symbols or drawings, and the temporal relationships between these symbols--time).

We usually manipulate these symbols on pieces of paper, which if saved for later reference may be generically called documents. We require a means of storing, retrieving, and indexing these documents and of communicating their contents to some other person.

These considerations gave birth to a hardware-software synthesis. Rather than take the accepted path of generalization--designing the computer interface to accommodate any imaginable task--we conceived of an interface that would be specifically designed for symbolic manipulation tasks as described herein. The HASCI keyboard was the result.

Fundamental Principles

The described theoretical explorations led to the evolution of a number of principles that form the rationale of the HASCI standard. A detailed examination of these principles follows.

The Computer Is a Tool

The computer as symbol processor and the user as pattern recognizer complement each other well. In this arrangement, the weaknesses of each can be ignored; together their strengths form a synergetic whole far more powerful than either, and such a blending of strengths is the functional property of any tool.

A hammer uses the advantage of a steel working face (hardness and mass) combined with the advantages of the human arm (motion and leverage) joined by an interface (the handle)

to perform some task dictated by intellect. Similarly, the computer uses the advantages of electronics (rapid manipulation of symbols) combined with the capabilities of the human mind (pattern recognition), joined together by an interface (keyboard and screen) in order to perform tasks dictated by intellect.

In an ideal situation the relationship of user and tool approaches one of *transparency*. The user is able to apply intellect directly to the task; the tool itself seems to disappear. This transparency is characteristic of all expert applications of tools--everything from hacksaws to racing cars.

Thus, a study of tools as a class can provide us with a set of rules that are applicable to a computer interface:

- *The interface is a means of controlling the tool.
- *The interface must accommodate the needs of both the application and the user.
- *The interface itself must present the information necessary for its use.
- *Mastery of the interface may require practice.
- *With mastery, the interface must become transparent to the user.

Clearly Label the Controls

Televisions are easy to operate. They have a limited number of controls. A stereo may have many more controls--complex models have dozens. But in each case, the controls either produce an immediately observable effect or are very clearly labeled as to their function. In each case, a relatively casual comparison of the controls on the device with the results produced and the *understanding* provided by intellect makes operation almost self-evident. Such is true of all mass-consumed products. However, on the average computer there are numerous functions that are in no way self-evident.

Perform this test: walk up to a computer you're not familiar with and pretend it's the first computer you've ever looked at. Then guess how to save or load a file of information. Get it? No way! You've got to study the manual and learn the code. You're required to learn and memorize the information. A little memory requirement is a positive thing: it makes the skill more valuable. But when you must rely on memory, the interface is effectively in you head rather than on the machine. (Imagine the potential hazards if a power saw were designed this way.)

We therefore see the necessity of providing controls for the major functions of the computer and of clearly labeling these controls. Ideally, activating the controls should

generate an instant feedback to the user: not just an audible "click" to prove the button was pushed, but also a significant change in status (such as a new message on the video display) indicating that something is happening.

Transportable Knowledge

The concept of *transportable operator knowledge* refers to the fact that users of consumer products expect and demand that the skills they acquire in learning to operate one machine be applicable to any machine of the same class.

For example, consider the typewriter. There are minor differences in the placement of certain controls, but a user who has learned on one typewriter can pretty well sit down at any typewriter in the world and type away. This is not because the task is overly simple: a typist must learn to manipulate a hundred or more keys, switches, and levers to operate the machine in differing circumstances. However, the typewriter as an architecture is fully stabilized to perform its appointed task. All typewriters have carriage returns, a means of setting tabs, a margin release, etc. These are sufficiently clear that an inspection of any machine rapidly reveals how to perform these functions.

Now consider the computer. Nearly every software writer and hardware designer has a unique way of telling a computer to save and load a file. Even though virtually every operator needs to perform these functions with great frequency, every time you change machines or programs you have to learn how to save and load all over again. (This is not to say that any one of these ways is wrong; rather, than on a consumer computer the basics should be done in one workable, learnable way.)

It is ironic that the data-processing and computer science industries have given so much attention to transportability of software. The benefits of this transportability appear to accrue primarily to programmers, and while it's understandable that people should create tools that they themselves need, transportable software eases only the programmer's burden. Transportable operator knowledge serves *all* users.

In a similar vein, it becomes clear that arguing the benefits of 16-bit versus 8-bit machines is analogous to arguing the merits of 8-cylinder versus 4-cylinder engines. Your choice should be based on how much payload you expect to haul, *not* whether you get a steering wheel with the vehicle. Performance from the consumer's standpoint is the ease with which desired tasks are accomplished: fast and difficult is still difficult.

When we approach the matter in this light, we realize that consumers will expect computers, both complex and simple, to have interfaces that are virtually identical. For all intents and purposes, anything that can be run on a 68000 microprocessor should be able to run on an 8080; the difference should be in how fast and how much, not *how*.

In terms of operating systems, while Unix may have certain advantages over CP/M (or vice versa), this is of no interest to the average user. Operating systems are tools for programmers. The symbol manipulator should function as an intelligent interpreter between the user and the operating system, and that interpreter should function almost identically on any operating system. (Most applications programs are considered as running *under* an operating system. The interface, however, should be considered as running *over* the operating system. It actually mediates between the operating system and the user just as would a programmer. In this case the interface is the expert who makes the difficult seem easy.)

Design Out Technical Choices

In the early days of the S-100 bus, I put together a kit for a serial interface board (the 3P + S). It was quite marvelous and went together easily, that is, until I got to the "jumper operations." There were dozens of options. You could configure the system just about any way you might imagine: number of data bits, parity, stop bits, and so on. All fine except for one small problem--I was a novice computer user and had no possible way of knowing which of these options served my purposes. After a few days of messing about and getting nowhere I asked a computer expert for help. He had the board configured for my system in a matter of minutes.

This highlights a typical problem. Because a computer can be configured in many ways, experts often want to build in every conceivable option because "you never know what the user may want to do with the system." However, we have already accepted the concept that the consumer computer is a tool for manipulating symbols. So we do have an idea of what the user will want to do.

Even a so-called user-friendly system may have an incredible array of choices. I recently bought what was billed as a user-friendly electronic mail system. It offered me options of stop bits and parity and data rate--just like the old 3P + S. It also presented a vast array of choices of how to send the data: compacted format, binary code, straight ASCII (American Standard Code for Information Interchange), and more. The designer of this code apparently confused "user-friendly" with "all possible options accessible." (The term

"user friendly" must surely rate as the inanity of the decade. When was the last time you thought of a tool as "friendly"? "Usable" and "useful" are the appropriate operative terms.)

Burdening the user with decisions concerning technical choices in no way addresses the task to which the tool will be applied, i.e., the manipulation of symbols. The system should automatically test the lines and choose settings appropriate for the circumstances. The user is then free to concentrate on the act of manipulating symbols rather than on the hardware. (This is how transparency is achieved.)

Thus a rule of thumb evolved: technical choices irrelevant to the symbol-manipulation task at hand should be eliminated from the user interface.

Predictability

In order to ease the chore of learning the HASCI system, we have attempted to keep the system as straightforward and predictable as possible. We try to allow different operations to be performed in a similar fashion whenever possible or appropriate. This does not require that there be only one way of doing each function, however.

For example, you can move the video cursor by pressing cursor keys on the HASCI keyboard. These arrow keys, when pressed in combination with the Shift key, or in combination with arguments such as WORD, move the cursor by different units. Even the novice experiences little difficulty with this scheme. Learning is accomplished by inspection and some experimentation.

However, experienced users may find this method cumbersome; moving their fingers from the main keyboard to type on a different group of keys slows them down. For the more-than-casual user, Control-letter functions (where you press a control key and a letter key simultaneously instead of a separate cursor key) are much quicker. Therefore, the HASCI processor also recognizes control key combinations for these same functions.

In this fashion both the novice or occasional user as well as the professional are well accommodated.

Simplicity

In designing a user interface it's important to keep simple things simple. More complex functions may be handled in a more complex manner because these will typically be used by more experienced users.

It's easy for experienced users to forget just how overwhelming a microcomputer can be. We attempt to judge the value of any product solely by the number of features offered for a given price. But what of the neophyte? Novices can assimilate only so much in one gulp, and that gulp is apt to be a small one.

A year and a half ago I tested the concept of a seven-function word processor, analogous to a four-function calculator. My premise was that seven functions are absolutely necessary for a useful screen editor: text entry, moving the cursor, insert character, delete character, save file, load file, and print file. With these functions, you can handle almost any word-processing task. More advanced functions can expand these capabilities and increase ease of use.

I tested the validity of this screen editor on a number of nontechnical users and found that they could be taught these basic functions in a few minutes of verbal instruction. And with only these functions, the system was truly useful. In fact, some of the users never asked if there were more functions. Even such a bare-bones editor proved to be a very useful tool, about as far ahead of a typewriter as the typewriter is ahead of clay tablets and sharp sticks.

I am not recommending that a screen editor be limited to these functions. On the contrary, I believe that constantly increasing the power of the system to manipulate symbols is mandatory and very desirable. However, the basics must not be obscured by the complexities of more advanced functions.

The HASCI standard calls for a selection of the most desirable functions to be placed directly on the keyboard with dedicated function keys. Many users will never venture beyond this--they will never feel the need to do so. More complex functions can be accessed via the use of Control-letter functions for access to specialized menus.

Defang the Computer

Over the years I've seen dozens of ways to get bitten by a computer. For example, one popular computer uses 8-inch drives for increased storage. There's a catch, however: the disks absolutely must be removed from the machine before it is turned off; failure to do so results in absolute and complete loss of all data on every disk in the system. Now it's easy to say, "Always remember to take out the disks," but in fact even experienced users occasionally fail to remember. They get so wrapped up in the job they're doing (as they should) that they forget that the hardware itself needs this critical piece of attention.

Another computer hazard shows up in the use of editors. Have you ever deleted something and then wished you hadn't? Silly question. I know of no more awful feeling than to have just erroneously deleted a document that I put a week's work into. The system should be smart enough to alleviate or entirely eliminate these dangers.

One answer to this problem is to deliberately place a slower menu structure in the way of any potentially destructive action. This often takes the form of a query, such as: "Your action will cause [a certain consequence] to occur. Please confirm this before I continue."

Another solution would allow you to change any decision even after the computer has acted on it. This is expressed as an Undo function key. Literally, this key allows you to undo or reverse your decision. For example, pressing the Undo key within a menu would take you to the prior menu. Pressing Undo within an editor after you had made a deletion would bring back the deletion. However, in order to fully defang the system, you should not allow the operator to undo *everything*. For example, suppose you just typed in three pages of text and pressed the Undo key: would you want the system to Undo your three pages of text? Hardly.

The HASCI concept requires that designers allow people to be people, not machines. Even the best of us occasionally forgets the right sequence or fails to do some required part of a protocol. It is the responsibility of the systems designers to defend the right of users to be human beings.

One shortcoming of many computer systems involves the use of modes. I don't see modes as inherently bad; certainly a human being does only one function at a time--you can't do order entry and write a letter at the same time. However, the problem in most system designs is that it is very difficult to change between functions.

Suppose you are merrily typing away and you need to calculate a few numbers for the document. Should you have to save the file, load the calculator, perform the computation, print the results, and reload the editor, all just to enter the result of your calculation? That's the trouble with modes. They make it difficult to change between functions and trap the user in the complexities of system integration. Common symbol-manipulation tasks and document-manipulation tasks should be accessible with push-button ease. HASCI allows you to change functions at will by pushing the appropriate control. Furthermore, when appropriate, if a prior function is recalled, you should find that function configured as you left it.

In an ideal implementation of HASCI, you should be able to turn the machine off, then power it back up and find it just as you left it, even if it was running a program at the time.

What You See . . .

The phrase "What you see is what you get" summarizes a concept of text display on word processors whereby formatting commands no longer appear as obscure codes imbedded in the on-screen text. Instead, the commands appropriately modify the displayed text so that you can see your specified formats on the screen before you print out hard copy. For example, if you indicate that a line is to be centered, it will appear centered in the displayed text. In addition, if you specify a change in type style, the altered text will appear in a graphic approximation of that style, enabling you to visually distinguish it from the surrounding text.

When we got the first sample of the Epson MX-80 dot matrix printer way back when, it already had a terrific selection of type styles available: emphasized, double-emphasized, compressed, etc. This opened up a whole new era of correspondence-quality printing, where the perfection of a fully formed character is gladly traded off for vastly increased versatility coupled with adequate legibility. The MX-80 was, of course, only the start. The newest printers now offer as many as 60 or 70 different type styles, and they also offer programmable character fonts. We may certainly expect to see the matrix densities of these machines increase very substantially over the next year or two, widening still further their performance gap over the fully formed character printer.

But then, as now, the problem was that the editors and personal computers available were designed to display on their screens only one or at best two or three different type styles--far fewer than even the first MX-80 was capable of printing.

This meant that although the printers had the capability, the computers were far behind in making this capability available in anything resembling an easy-to-use fashion. Most of us have had to settle for inserting control codes using one language-like protocol or another. This is clearly unacceptable because it violates the "easy to learn" maxim.

Here is a case where very useful symbolic manipulation features are very difficult to access. The answer is to design the system with this capability in mind, make these functions easy to access, and at least where desktop units are concerned, place these changes *right on the screen*. This establishes a feedback loop which makes the system easy to operate.

"What you see is what you get" is more than a maxim. It is a crucial consideration in the effort to make the symbol manipulators--computers--easy to use.

Consumer Quality

All the above principles and guidelines add up to make the computer a consumable product. With the computer, as with any good stereo, television, or automobile, we expect to be able to gain access to substantial capabilities with little if any specialized knowledge. Manuals are for reference; you shouldn't need an advanced degree just to open the box. You should be able to set up the computer, hook up the cables in the obvious places, turn it on, and have it work right the first time and every time. Using computers to advantage should be a game that everyone can win.

Beyond Theory

Now that the theory and principles behind the HASCI system have been explained, some obvious questions arise: "How can this idea actually be implemented on a personal computer? What specific keys do we need? What should they do? And what should be displayed on the monitor screen?"

The Menu

Menus present an exceptionally easy way of introducing the newcomer to the operation of a system. They tend to fail, however, on two points: first, some designers create unwieldy menus by trying to throw in everything but the kitchen sink; second, they provide no alternative for experience users who eventually learn the menus cold and find it irritating to have to wait for each menu to appear.

In the HASCI scheme, the problem of cumbersome menus is eliminated by treating the entire computer system as a series of interconnected choices in an inverted tree of decisions. Each branch of the tree represents a possible function that the computer can perform for you. Also, in virtually all cases, the number of choices in a menu is kept below eight. This number of choices has proven to be a perceptual limit for understandability.

The problem of menus that make you wait is solved by allowing you to input menu selections as fast as you can make them; thus, the tedium of sitting through long, familiar menus is entirely eliminated.

The Choices

When dealing with HASCI, as with any computer system, your first choice is whether or not you want to use the computer. If you do, you must of course turn the machine on. When power is first applied, the system comes up automatically as a word processor—you needn't access the operating system.

The HASCI keyboard controls are divided into seven main groups of keys. Of these, the following three groups are typical of many contemporary keyboards in their configuration and layout:

- *typing keys
- *editing and cursor-movement keys
- *the numeric keypad

The remaining four groups take the place of programmable-function keys (which have always had such clever names as F1, F2, F3, and so on). These groups give you access to the most essential functions of the system:

- *system controls
- *file controls
- *applications controls
- *typestyle controls

Each of these four groups is clearly labeled on the keyboard itself. In contrast, the first three groups are self-explanatory and are not labeled.

This arrangement of the keyboard provides the first menu level of the system: you choose the *group* whose function title (or self-evident application) most closely matches your needs.

People should be able to guess which group, and which key within each group, performs any given function. The titles of the groups and the individual keys on the HASCI keyboard have been chosen to facilitate this capability. (The keyboard has been tested on a large number of people unfamiliar with computers; virtually everyone was able to correctly guess the intended function of each key the first time.) In addition, after selecting any given key, the effect on the system is immediately obvious. And, if all else fails, the HASCI system has a Help key. Thus the HASCI system is nearly manual independent.

The second menu level involves choosing an individual key from among the seven groups on the keyboard. As mentioned before, the keys of the first three groups are already fairly familiar. Of greater interest are the individual keys of the four function control groups.

The System Controls

The four system controls affect the execution of a system program already in progress:

Stop takes the place of more usual Pause and Break keys. When pressed, it effectively halts system execution and asks if you wish to stop or continue.

Help provides you with specific information relating to the nature of the choices available at any point in the decision tree. Your options are explained in some detail. Additionally, you may access information about any specific function.

Copydisk lets you do just that: copy a disk. (We didn't use a Backup key because inexperienced computer users expect Backup to make the machine go backward.) Although Copydisk is a fundamentally necessary function in any floppy-disk-based system, this key might not be used in other implementations of HASCI.

Undo is an "undecided" key. At any point, virtually any decision can be undone with this key. It protects you from accidental deletions and also allows you to skip rapidly back up a menu tree.

The File Controls

File controls allow you to easily manipulate your files (i.e., the places where your documents are kept):

Store places a document you've created into the mass storage files.

Retrieve is the complement of Store. It allows you to procure a specific document for further symbol manipulation.

Print allows you to print the contents of any document on the system printer. Numerous print-time options are provided.

Index may be the most novel and useful key on the machine. It displays an index of all files in the system. All files are filed by date and time or sequence of creation. This information is automatically assigned by the system. The name of the file or index reference is requested by the computer in response to the Store command; you may specify a reference of up to eight words in length. When Index is pressed, you are offered three choices. You may view the index (1) sequentially by date and time of creation; (2) alphabetically by index reference; or (3) alphabetically

cross-indexed, with every word in each reference cross-referenced to every other word. (This is exactly what most people wish they could do with their manual systems.)

Mail accesses a complete electronic-mail system such as the Valdocs system, which implements the HASCI system on the Epson QX-10.

The Applications Keys

The applications keys cover the entire family of symbol processors. If you recall, a computer is basically a symbol manipulator, and there are four kinds of symbols that we need to manipulate: words and letters, numbers, graphic symbols, and the temporal relationships among these symbols (time).

The manipulation of words is accomplished with the typing keys and is essentially self-evident. Of the four keys in this group, three are dedicated to the remaining symbol types. These keys are labeled Calc (for calculator), Draw (for graphics utilities), and Sched (for schedule).

The nature of each of these programs is flexible: while a four-function calculator may be enough for me, you may require a sophisticated scientific processor, and a spreadsheet calculator may be ideal for someone else. Likewise, some people have simple appointment-scheduling needs, while others require complicated systems such as the Performance Evaluation Review Technique (PERT) or the Critical-Path Method (CPM). The same is true for the Draw utility. Thus, no standard exists for these functions. However, HASCI standardizes the means by which one enters and departs from any symbol processor.

You switch from one application type to another by pressing the appropriate key. On larger systems these keys would have indicator lights to show when they were selected. The selection would also be clearly indicated on the screen.

The fourth key of this group is labeled Menu. As you might guess, it is the garbage can; everything else is found there: languages, utilities, all the stuff that normally clutters up a directory listing. Ideally, any programs resident on the particular operating system that had not been converted to use the functions and protocols of HASCI would appear under Menu. In other words, a HASCI system running over CP/M should be capable of running any standard CP/M software. The same would be true for a system running Unix or any other operating system.

The Typestyle Keys

You can alter the symbol type style displayed on the screen in alphanumeric by using one of the four typestyle control keys: Italic, Bold, Size, and Style.

The Italic and Bold keys operate immediately. If either key is pressed while typing, all subsequent text entered will assume that type style. Pressing the key again reverts to the previous type style. The Size and Style keys access menus that allow you to select from whatever choices are supported by the terminal and printer used. In regard to Style, a machine must have at least one font; however, two (one serif and one sansserif font) would be desirable.

The Third Menu Level

The third menu level occurs after a function has been selected by pressing its key. In some cases, there is no third level; the functions act immediately. Examples include the cursor keys and the Italic and Undo keys. Other keys may have one or more levels of menu existing beyond the keyboard. These levels are indicated on the display screen.

Screen Standardization

The screen layout should be essentially identical from menu to menu and present all necessary information in an easy-to-understand manner. The HASCI screen is divided into three *windows*, each of which contains a specific type of information.

The *document window* contains the main document, which holds the symbols under inspection or manipulation. When the machine is first powered up, the display resembles that of a word processor--the document window fills the screen.

The document window may contain visual devices to simplify the manipulation of the symbol type in question. In the case of a word processor, this window contains a *ruler line*, which marks column positions, shows current column position, shows tab settings, and so forth. The window also contains a status line showing the name of the document under inspection along with such mundane items as date and time.

When you're browsing through a file, examining a directory, or performing some similar task where you may wish to select from among many choices, the document being examined for these choices (for example, an index) would appear in the document window. If you are to make a choice from such a document, a cursor will appear to indicate that a selection is expected. But there is never more than one cursor at a time on the screen.

The *interaction window* appears only when the machine requires some discrete information or a specific response. It always appears below the document window. On an 80-column by 25-line screen, this window is 80 columns by 8 lines in size.

Two classes of interaction can occur. In the first, the computer may request a string of typed characters. For example, the system may ask, "What is your name?" The question is presented in the interaction window, along with a cursor indicating where your response will be entered. In the second class of interaction the computer requests a selection from a menu. All menus appear in the interaction window. Whenever you have to make a decision, the system prompts that explain the choices appear in the interaction window.

The *prompt window* is a small window at the bottom of the display that contains brief reminders (prompts) or flags of use for any given situation. They are optional with the software designer.

Rules for Menus

Menus must follow certain rules. First, menus should always appear in the same place on the screen. Second, menus should be designed so that you may indicate your choice by one of two standard methods: type the first letter of the first word and press the Return key, or move the cursor until it is over that letter and press Return. A third, optional method, which can be activated by a software switch, would be to type the letter without pressing Return to activate the choice immediately. The first two schemes allow the casual user simple and failsafe means of choosing from a menu, and the third method allows experienced users to reduce the number of keystrokes and access menu choices more rapidly.

Finally, menus should be organized so that the most common choices occur first in position and potentially destructive choices occur last.

The HASCI Keyboard

Like it or not, the keyboard is with us to stay. In designing a keyboard, we chose the format of the typical office typewriter. The key positions are identical and the feel is similar, so anyone familiar with a typewriter should be reasonably comfortable with the HASCI keyboard. However, by adding just a few additional keys, we were able to have the keyboard generate an entire 8-bit superset of the ASCII (American Standard Code for Information Interchange) character set. Thus the HASCI system is upward compatible with ASCII-based systems; a computer using HASCI can run any standard software.

Included in the extended ASCII is a set of standard graphic characters. These allow the creation of accented letters (by actually overtyping one character on top of another) and line drawings for boxes and forms. Also included are some Greek and special-purpose mathematical symbols. You gain access to these characters by pressing a Graphic-Shift key, which converts the normal typing keys to symbol generators. The first set of these symbols should be printed, etched, stamped, or otherwise marked on the front of the key caps in a color similar to that of the key cap. This should not be a high-contrast color; such treatment causes visual distraction and fatigue.

In addition to this primary set, one may simultaneously press Shift and Graphic Shift to access a second set of graphic characters. Most of these are logically related to their unshifted character. For example, all *line* symbols have a *double-line* counterpart. Thus, while the second set is not shown on the keycaps, it is easily learned.

Types of Physical Controls

We have avoided using any controls other than keys and push buttons in the current HASCI standard (although voice recognition may certainly be incorporated when appropriate). Of two primary motivations, the first was familiarity. Contrary to a current myth, keyboards are extremely familiar objects in our society, and a vast number of potential computer users are already familiar with their use; no other practical means of entering textual data into a computer exists today. Second, the HASCI keyboard must be available on portable computers as well as on fixed desktop units. If the interfaces on portable and fixed units are substantially different, the concept of transportable knowledge would be violated.

HASCI allows a number of ways of performing certain actions. No one method is suited to all possible environments. Accordingly, it is quite possible to have a Xerox-style "mouse" in a HASCI system (that's just one more way of moving a cursor around a screen and making choices). Similarly, cursor keys, control keys, joysticks, etc., are equally valid in the proper time and place. A typewriter keyboard represents merely one valid method of entering text. Others will evolve and become common, but typewriter keyboards are likely to continue in popular use for a long time.

Conclusions

The HASCI interface is by no means an end; rather it marks the beginning of an era of consumer-oriented computers.

HASCI is not intended to be a fixed thing. We hope it will evolve and improve with time. Keys will come and go, menus will change, and groups of keys will grow and shrink. We expect that computers specifically designed from the ground up to support HASCI will help to reduce substantially the overall system cost and increase system performance.

#