

ImagePlus High Performance Transaction System

by R. F. Dinan
L. D. Painter
R. R. Rodite

The need for a cost-effective method for handling, processing, storing, and retrieving transaction documents combined with the availability of hardware and software technologies capable of satisfying this need are the basis for the High Performance Transaction System discussed in this paper. The particular transactions that are the subject of this paper are bank checks, the volumes surpassing 50 billion per year and continuing to grow. Other transactions might be the handling of such documents as bill remittances, tax documents, mail-order forms, census forms, and many other similar applications. This paper discusses the system design, hardware and software architectures, and performance of the ImagePlus™ High Performance Transaction System.

This paper describes the architecture, system design implementation, and initial applications for the IBM ImagePlus™ High Performance Transaction System (HPTS) for high-volume paper applications in banking and other financial institutions. Earlier work on document processing at the IBM laboratory in Charlotte, North Carolina, from the late 1970s to 1980, resulted in some key building blocks for this system. However, recent advances in technology have allowed designs that had proved to be impractical in those earlier times.

The memory required for a single image frame might have occupied several cubic feet in 1980, but now fits on a single PC card. Very large-scale integration

(VLSI) has also brought the logic to practical size, cost, and performance levels. Technological capability now exists for high-resolution image capture at rates as high as 40 frames per second. With the promise of future development of consumer-based, high-definition digital television, continued improvements in cost and performance are projected. There is a similar promise of development improvement of mass storage media driven by optical and magnetic storage technologies.

The initiation of the transaction system development discussed in this paper was customer driven. Several major banks came to the same conclusion about the readiness of image technology to be used as a solution to their document processing needs. The Federal Reserve Bank of Boston asked the IBM Charlotte Laboratory to demonstrate the use of image in the high-speed capture, archiving, and retrieval of images of U.S. Treasury checks. That demonstration was completed successfully in 1989.

As a result and in consultation with our banking customers, an architecture was formulated; its guide-

© Copyright 1990 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

lines ensured that the image solutions would be readily assimilated into present and future operational banking environments. Keys to this end are the treatment of image data as ordinary coded data, migration in an orderly fashion from the present customer applications to image applications, and building a system core and applications on the strategic foundations associated with future processing, storage management, and communications building blocks as we could foresee them. We now present details of the HPTS architecture, system design implementation, performance, applications, and future considerations.

The need for a High Performance Transaction System (HPTS) centers around the problem of today's bank check processing environment. This is a highly automated paper-based process, with machine-readable check encoding as the method of tracking and sorting the checks at very high speeds. After two generations of hardware and software improvement, the resulting improvement of paper-handling efficiency was growing less and less cost efficient. Costs are now rising with the cost of the remaining labor content. Check volumes continue to increase, and this keeps a constant pressure on cost control and the quality of customer service. Image technology offers the promise of being the breakthrough to eliminate the high cost of paper handling, thereby reversing the upward trend in costs. A system solution is key to a successful conversion. With many large banks already relying on large MVS-based hosts to handle the check process, these banks need an uninterrupted, unfailing process even during conversion to the new system. The system architecture chosen would have to be one that provides for evolution within a normal data-processing environment.

Migration from a conventional paper-and-coded-data system to an image-and-coded-data system must be staged, starting from the present MVS-based systems. The migration, which must be carried out on several planes, is as important as creating the ultimate system design. The check processing system is actually many separate applications, and the image conversion must migrate through multiple applications at a time to maintain uninterrupted order at the bank. New applications that are not possible without image capability may be added in the future. A platform that allows convenient growth to these is required. One example of an application that was not possible prior to image processing was the one in which an image of a check is returned to the bank's customer rather than the actual paper check.

Finally, flexibility to future environmental changes in the host system is required. Check processing systems are expected to last many years and several computer generation cycles. Of high importance is the need for the architecture to be strategic in nature, so that applications remain functional, not only in today's but also in future data processing systems.

The concept that image data should be treated as much as possible like normal data is central to the HPTS architecture. Once this is achieved, most previous art practiced for today's data processing can be applied. Rather than building a completely new data-processing infrastructure for image processing, image has adapted to the present one. Carrying this concept a step further, those areas adopted are the strategic areas. Access methods, operating systems, communications, user access, programming languages, and so on all conform to that supported by Systems Application Architecture™ (SAA™). These strategic areas are in concert with future system strategies within IBM that are associated with the traditional coded data world.


The key system operating points were established by identifying the significant elements in a matrix of key parameters, processing image data for those elements, and evaluating the results for adequacy to perform the applications. The following are the parameters we evaluated:

- Resolution—120, 200, 240, and 300 picture elements per inch
- Document speed—240, 345, 400, and 450 inches per second
- Compression algorithms—International Telegraph and Telephone Consultative Committee Group Three Standard and Group Four Standard (CCITT 3 and 4), Differential Pulse Code Modulation (DPCM), and Adaptive Bilevel Image Compression (ABIC)
- Scanner—Reticon Charge Coupled Photodiode (CCPD) 128 by 8 and 1024-element, low-light device
- Presentation—Black and white, gray scale, and mix of low resolution gray scale with high-resolution black and white

Description of the architecture

The High Performance Transaction System (HPTS) image record conforms to the IBM Image Object Content Architecture (IOCA). This architecture provides a standard definition for representing images

Figure 1 Image record



and the conventions to interchange images. Figure 1 illustrates the record layout used. In order to accommodate the range of future HPTS applications, several forms of object data are required in the record as illustrated by object data 1, 2, . . . n . That is, object 1 might contain the black-and-white rendition; object 2 might contain a gray-scale rendition of the same target; object 3 might contain a window, and so forth. Application-coded data are carried as a lead header, and these data have fields that allow for many applications. The IBM Mixed Object Document Content Architecture (MO:DCA) data stream is utilized to interchange data among applications and environments.

Figure 2 illustrates the applications that this architecture supports, shown in a wheel-like display. The core structure at the wheel's hub supports both present and envisioned future applications. The presentation forms of images required for each application vary with the application. Therefore, the capability is provided to have multiple forms within the record and to change those forms with time. The IOCA format accommodates varied resolutions, compressions, and other image presentation characteristics. The initial HPTS offering uses two forms of image. A black-and-white resolution image is used for applications such as unrestrained handwritten amount recognition and keying, and gray-scale medium-resolution images for viewing check endorsements. The concept of a black-and-white image is illustrated in Figure 3. All image information is depicted as either black (1) or white (0) generated by the dynamic thresholding algorithm discussed in Reference 1. The requirements for HPTS call for the black-and-white resolution to be 240 picture elements per inch.


Figure 4 illustrates an 80-pel gray-scale form. This form is used as another set of objects in the record

format. Here the image information is depicted as a 4-bit word giving 16 possible levels of gray. Both the black-and-white and gray-scale images are compressed with the same bilevel algorithm, ABIC, that normally deals with only binary, such as the high-resolution, black-and-white images. The figure illustrates how the gray-scale information is preprocessed to a form that can be conveniently handled by the bilevel compression algorithm. In this case, the darkest regions are coded as a binary 15 or binary 1111, midlevel gray as a 10 or binary 1010, and light gray as a 5 or binary 0101. When each respective ordered bit of the image is arranged specially in its own bit plane, i.e., the 20, 21, 22, and 23 bit planes, each plane can be processed separately as a bilevel image. These data planes are then serialized separately and concatenated to form one continuous bilevel form that is compressed by ABIC. The reverse process takes place when the image is reconstructed.

With the data record defined, the operating and communication environments round out the key elements of the architecture. The Systems Application Architecture (SAA) was chosen as the overall architecture. The initial operating system environments are MVS/ESA™ and OS/2®, as shown in Figure 5. These base operating system environments are utilized without modification. Core and application functions built on these bases are strategic, in that they are not only portable between SAA environments, but they also will be part of future operating systems growths and improvements. LU 6.2 was chosen for communication, thereby bringing to the base elements the IBM strategic network architecture.

Overall, the objective of bringing the HPTS image technology into the SAA technology of coded data has been accomplished. Questions of particular configurations for individualized solutions can be an-

Figure 2 High Performance Transaction System applications



swered in the following generic way. If there is a solution for coded data, that same solution applies to HPTS image data.

System design implementation

ImagePlus HPTS consists of several hardware components, and system software and application software components. The IBM approach to developing HPTS has been to provide a set of core hardware and software products upon which image applications may be built. Included is a description of the system and application design implementation. For purposes of this discussion, the system software developed in the IBM Charlotte Laboratory and to which the image applications interface is referred is called the *core software*.

The HPTS architecture is built upon a base of existing IBM hardware and software products that includes the following:

- System/370™ mainframe processor
- IBM 389X/XP Document Processors
- Optical and magnetic storage media
- MVS/Enterprise Systems Architecture operating system (MVS/ESA)
- System Management Software/Data Facility Product (SMS/DFP)
- Virtual storage access method (VSAM)
- DB2™ database
- Advanced function printing (AFP)
- Check Processing Control System (CPCS)
- PS/2® workstations and peripherals
- Token-Ring Network
- OS/2 Extended Edition

HPTS platform and data flow. The HPTS hardware and software system illustrations in Figures 6 and 7 provide an overview of the ImagePlus HPTS platform. A description of the flow of image and coded data through the various system components follows.

Images of documents—checks, remittances, drafts, and so forth—are captured at a rate as high as forty items per second on the IBM 389X/XP Document Processors, through the addition of the IBM 3897 Image Capture System, which contains image cameras and image processing logic. The MVS/ESA host is the image server utilizing standard DASD, tape, and optical disk via the Object Access Method (OAM). All image processing, including compression and decompression, occurs outside the host. Workstations access the image objects for both high-speed and low-speed applications with decompression and recognition occurring in special Micro Channel® cards under OS/2. Standard system printers, PS/2 workstations, displays, and scanners are all also utilized. Strategic access methods, such as Object Access Method (OAM) and virtual storage access method (VSAM), are utilized along with the strategic Check Processing Control System (CPCS). The systems shown can support a large range of customer check processing load requirements far surpassing the largest requirement today of one thousand transactions per second peak capability.

These images are transmitted to the MVS-based mainframe via a channel-to-channel (CTC) attachment and are stored in an object database via Image Host Application Services (IHAS). The coded (nonimage) data are transmitted to the mainframe via an Advanced Program-to-Program Communications (APPC) (LU 6.2) attachment. Coded data management, control, and reporting are handled by the Check Processing Control System (CPCS). After the capture of the image and the coded data from the 389X/XP, data are transmitted from the MVS-based mainframe to one of two token-ring local area network (LAN) configurations with PS/2 workstations for image application processing.

There are two paths to the LAN for data, both of which are used in the high-speed, image processor LAN configuration. This configuration supports applications, such as key entry, which require character recognition and subsecond response time. The second LAN configuration uses a single path for data traveling from the host. We now discuss these two data paths.

A channel-to-channel (CTC) connection is made between the MVS-based mainframe and the PS/2 workstations via the image processor. For high-speed, data-extensive applications, such as key entry, this path is used for images traveling from the mainframe to the LAN.

Figure 3 Black-and-white image




Figure 4 Gray-scale image





Figure 5 Operating and communication environments for the High Performance Transaction System



An IBM 3174 connection is made between the MVS-based mainframe and the PS/2 workstations using LU 6.2 protocol. In the high-speed LAN configuration (which also has a channel-to-channel image processor path), the LU 6.2 path is used to send coded data [e.g., Check Processing Control System (CPCS) Mass Data Set (MDS) records] to and from the host. In the alternate LAN configuration, which has the single LU 6.2 path, the LU 6.2 path is used to send coded data to and from the host, image data from the host to workstations, and images from workstations (where they have been captured via low-speed scanners) to the host.

As indicated in the High Performance Transaction System (HPTS) software architecture shown in Figure 7, the HPTS platform consists of the following four unique workstation types.

IBM 3898 Image Processor. This is an industrial-type PS/2 that uses the Intel 80386® chip and houses hardware for image compression, decompression, and character recognition (RECO). The processor also houses the associated device driver and image processing core software, known as Image Processing Recognition Services (IPRS). The image processor is a special-purpose workstation whose primary function is image management and distribution. There is no application software in this workstation, and, during normal operation, there is no interface to an operator.

Workstation manager. A PS/2 Model 70 or 80 houses the core software component, called Image Management Workstation Facility (IMWF), and the image application manager. The image application manager is discussed in more detail later in this paper

Figure 6 High Performance Transaction System hardware architecture





Figure 7 High Performance Transaction System software architecture



under the heading, "HPTS application architecture framework." The workstation manager is a special-purpose workstation whose primary function is application management and image tracking on the local area network (LAN). During normal operation there is an interface to an operator in support of the LAN configuration and set-up.

Image entry workstation. A ps/2 Model 55SX is the image entry workstation with a medium-resolution display (IBM 8503) that is used for high-speed, image-assisted data entry. This workstation houses the core

software component, called image entry workstation facility, which includes the applicable workstation and communication service routine subcomponents of the Image Management Workstation Facility. The workstation also houses image application software. The image entry workstation is a typical workstation whose primary function is to serve as the interface between the image application software and the operator.

Image extended workstation. This workstation is a ps/2 Model 60 or higher with a high-resolution dis-

play (IBM 8514) that is used for such extended image application functions as image compression and decompression at the workstation using Document Storage System (DSS) card, image printing on an IBM 4019 workstation printer, and optical disk medium storage and retrieval. This workstation is also used for low-speed image capture on such scanners as the IBM 3119. Coded data capture is also supported via a code-line reader. The image extended workstation houses the core software component, called Image Extended Workstation Facility, which includes the appropriate workstation and communication service routine subcomponents of the Image Management Workstation Facility. It also houses image application software. The image extended workstation is a typical workstation whose primary function is to serve as the interface between the image application software and an operator.

Upon completion of image application processing on the LAN, all coded data updates and newly captured images are transmitted back to the host processor via the LU 6.2 (3174) connection. The host image applications use the services of Check Processing Control System (CPCS) to process and store the coded data updates and the services of IHAS to process, store, and print images.

HPTS system software building blocks. The system or core software components upon which the image applications are being built are the following:

- Check Processing Control System (CPCS)
- MVS Host Support Program (HSP)
- IBM 3890/XP MVS Support
- Image Application Library Services (ALS)
- Image Host Application Services (IHAS)
- Image Workstation Application Services (IWAS)
- IP Recognition Services (IPRS)

Check Processing Control System (CPCS). The primary functions include coded data management, control, and reporting. This product is enhanced to support image by providing such functions as priority processing (which allows a customer to direct the sequence in which the application tasks are processed), unencoded item capture, electronic item capture, and system/workflow management.

Although CPCS is a key component of HPTS, it is important to note that IBM implementation of an open-ended architecture provides the means whereby a non-CPCS customer may install the other components of HPTS and bridge them to other existing item processing systems.

MVS Host Support Program (HSP). This core software component provides several functions, including simulation tools for IBM 389X/XP Document Processors and MVS driver code for 389X/XP Document Processors. The primary function used in support of image is the Advanced Program-to-Program Communications (APPC) software support for the LU 6.2 (3174) coded data and low-speed image path between the MVS host and the image workstations on the LAN.

3890/XP MVS Support. This product provides several functions to support document processing applications and is a prerequisite for CPCS and HPTS products. The 3890/XP MVS Support provides macros with which a user may write sort programs for 389X/XP Document Processors, and it also provides the queued sequential access method (QSAM) modules needed to support channel-attached document processors. This support further provides a simulator that allows the user to capture data from disk, then pass the data through the user sort programs and create CPCS records without the use of a document processor.

Image Host Application Services (IHAS). IHAS is a feature of the Image Application Library Services (ALS) that is a set of mainframe services used by all applications to access images in the image database, to access coded data results from image compression, decompression, and character recognition (RECO), and to access image document window coordinate and processing parameters. IHAS also provides facilities for printing images on high-speed mainframe printers (such as the 3835) and for transmitting images to workstations on a LAN.

Image Workstation Application Services (IWAS). IWAS is a feature of the Image Application Library Services (ALS) that is a set of workstation and LAN services that support image processing. In contrast with IHAS, whose services are required in total by all HPTS applications, IWAS has been divided into three subcomponents. The purpose for this division is to provide customers with a means of configuring workstations and LANs specific to the applications they have chosen to implement.

The Image Management Workstation Facility (IMWF) is a set of local area network (LAN) communication services used by the HPTS applications to route images and coded data to and from the mainframe and to other workstations on the LAN.

The *Image Extended Workstation Facility* (IXWF) is a set of device support and image enhancement routines used by HPTS applications that require advanced image workstation capabilities. These services are used by the applications to display (on an 8514) and print (on a 4019) gray-scale and bilevel images, compress and decompress images at the workstation, and interactively enhance the image as it is displayed. IXWF also provides device support for low-speed image capture and magnetic ink character recognition (MICR) code line data capture.

The *Image Entry Workstation Facility* (IEWF) is a set of device-support routines that is a subset of those provided in the Image Extended Workstation Facility (IXWF). These services are used by HPTS applications (such as key entry) that require base image functions (such as displaying all or portions of a bilevel image) in support of image-assisted data entry. The IEWF routines support the display of images on medium- or high-resolution 8503 or 8514 displays.

Image Processing Recognition Services (IPRS). This core software component supports image processing on the LAN for applications that require high-speed image decompression and numeric hand-printed character recognition and machine-print recognition (RECO). IPRS also contains memory management and image distribution functions that send images to the workstations on the LAN for validation and viewing of the RECO results.

Performance considerations

More than 50 billion checks are processed annually in the U.S., most of which are processed against very tight deadlines to minimize the amount of dollar float charged to the financial institution. Banks have been very creative to date by implementing controls that flag high-dollar items for special handling, by chartering airplanes to transport the checks, and by automating through the use of high-speed document processing equipment for transporting 40 documents per second. Images that are of sufficient quality to perform the applications illustrated in Figure 2 must be captured at a minimum of 240 pels per inch. This requirement is driven by the need to perform the recognition of unrestrained numeric hand printing as an assist to the keying application.

The second requirement when viewing the captured image is the ability to recognize the overprinting that results from endorsements placed on a docu-

ment by machine, hand stamp, or writing. It is necessary to distinguish overprinting from printing on the base document itself. Gray scale is required to provide this viewing depth. The capture operating point selected to meet this requirement is 240 pel and 256 levels of gray, which is reduced to 80 pel and 16 levels during processing. The 240 pel and 256 levels are specified for the initial capture to allow for intelligent processing and commonality with the need for a 240 pels-per-inch black-and-white image. The check sizes range from 6 by 2.75 inches to 8.75 by 3.67 inches, which results in 950 000 to 1 850 000 bytes of information generated for each side of the document. Because the system captures both sides of each document, the bandwidth of the data path from the scanner to the image capture processor must handle 170 megabytes per second. The goal of the system design is to minimize the amount of data, with a specific goal to fit within a 3.0 megabyte channel when transmitting the processed data to the host. Several processing means (filtering, separate black-and-white and gray-scale renditions, reduced resolution and gray-scale depth, elimination of extraneous information, and the application of a compression algorithm) are used to reduce the data. The result is images that average 42 000 bytes each, well within the capability of the 3.0 megabyte channel.

The images are stored on DASD on a real-time basis using IHAS to provide the fast storage and retrieval times required by the applications. Longer-term storage is accomplished by transferring the data after processing to other media, such as tape or optical disk. A single reader/sorter typically processes 400 000 documents per day, requiring 17 gigabytes of storage if all data are retained. Several reader/sorters may reside in a single processing location. The CPU requirements vary, depending on the current applications. A black-and-white front-image-only entry system that performs keying, balancing, data preparation, work in progress (WIP) and statements requires 1.4 million instructions per record per reader/sorter.

The image processor contains a high-capacity data link that is a channel-to-channel configuration and is capable of decompressing, windowing, and performing recognition results at a rate of 20 document images per second. The system allows for recognition assistance via other documents that normally accompany a deposit, such as deposit slips, adding machine tapes, and electronic teller input. For example, in the case of adding machine tapes, the tape image is

entered into the system via a low-speed scanner, and recognition is performed on the machine printed characters. The results of this process are correlated with the accompanying documents for confirmation of the character recognition performed on the documents. Unrecognized documents are routed to the keying workstations via the 16 Mb LAN for processing by an operator. The image processor supports an average of ten keying workstations, each capable of processing 4000 images per hour.

The image qualification process is a substantial improvement over conventional processing. System productivity improvement results from the fewer items that have to be keyed and from the faster operator keying rate that results from image. The results of the keying and recognition are compiled and validated automatically at the workstation manager.

For applications not requiring the high-speed image processor, decompression is performed in firmware at the workstation. These functions are generally of a document-research type that do not place a heavy load on the system because of the comparatively long time it takes the operator to analyze the image before an action is taken. The need for exception handling of a document or report is met by the use of a low-speed scanner and/or a printer attached to a workstation. Multiple scanners accommodate a wide variety of document sizes from adding machine tapes (2.5 inches in width and lengths of up to 5 feet) to sheet paper of 8.5 by 13 inches. Scanning time depends on the particular scanner being used and the mode of operation selected (black and white or gray scale).

HPTS image application software

For purposes of discussion, image processing applications have been divided into the two groups, front-end subsystem and back-end subsystem. The following is a general description of each of the subsystems and a discussion of the application architecture framework within which all the HPTS image applications must fit.

Front-end processing subsystem. The purpose of the applications in the front-end subsystem is to perform all the required processing on a group of items to facilitate movement of these items to the next bank or processing center. The applications that are collectively referred to as the front-end subsystem are the following.

Data preparation (data prep) application. The data preparation application functions are performed on the image extended workstation between the time of receipt of items at the bank processing center and the time of capture and sorting on the 389X/XP. The purpose of the data preparation application is to provide for conditioning of the items and to accept item control and source input information. The item conditioning process prepares the items for high-speed processing on the 389X/XP and includes the removal of paper clips, orientation of documents, insertion of control documents, and the placement of documents into trays or boxes for transportation to the 389X/XP. The control and source data input process includes the ability to do the following: (1) capture images of adding machine tapes or cash letters (using a workstation scanner) to provide assistance data for downstream RECO processing balancing; (2) capture control numbers from MICR-encoded control documents (using a 4706-type document code-line reader); and (3) weigh the groups of items (using a PS/2-attached weight scale). Additional information that can be captured in the data prep application is optionally key-entered by the data preparation operator. Examples of the data that can be entered by the operator are source of the items, time and date of receipt, type of work, and the availability schedule to be assigned to the items.

Key entry. The key entry application functions are performed on either the image entry workstation (for amount key entry) or the image extended workstation (for reject processing) between the time of capture and sorting on the 389X/XP and processing by the balancing application. The purpose of the key entry application is to support high-speed keying from image of data associated with the items. The two categories or classes of check items processed by this application are MICR reject items from the 389X/XP and *unqualified items*. The term *unqualified items* means those items whose amounts must be MICR-encoded on the bottom of the check.

Balancing. The balancing application functions are performed on the image extended workstation. In a typical work flow, the balancing application is invoked for a unit of work (or group of items) upon completion of capture and sorting on the 389X/XP and processing by the key entry application. The purpose of the balancing application, through the use of images and nonimage data, is to provide the necessary tools to a balancing operator to balance the dollar amount of the documents processed in HPTS to a known control total.

Work in progress (WIP). The WIP application functions are performed on a standard PS/2 similar to the image entry workstation. Unlike other image applications for which there is a logical flow of items from one to the next application, WIP is always running and is not a specific step within an application sequence. The WIP application is primarily concerned with the current state of operation (or work in progress) within HPTS. The purpose of WIP is to provide HPTS with production control, performance measurement, modeling, and forecasting information upon which bank personnel can make scheduling, personnel assignment, or other types of decisions.

Back-end processing subsystem. The back-end processing subsystem consists primarily of applications associated with handling the commercial and retail accounts (savings, checking, etc.) of the bank itself. The applications that are included in the back end are the following.

Image statements. The purpose of the image-statement application is to provide the commercial and retail customers of the bank with the opportunity to receive printed copies of images as part of their account statement, rather than the physical documents themselves. The mainstream functions of the statement application are performed on the host on statement-cycle day. Images that were stored on magnetic or optical media on the day of capture are retrieved. The image processor is used to decompress and window the images, with the results printed on a mainframe printer, using advanced function printing (AFP) facilities. The workstation components of the application are performed on the image extended workstation and provide a mechanism for handling exception statements and for servicing customer requests for interim account statements.

Exception processing. The purpose of the exception-processing application is to handle posting system exceptions. The functions of this application are performed on the image extended workstation and include the handling of overdrafts, stop payments, referrals, closed accounts, and signature verification. The use of on-line images results in an operator throughput improvement over existing paper-handling exception processing systems.

Return-item processing. The return-item processing application is the means by which a bank handles its incoming and outgoing returns. The image extended workstation provides image enhancement capabilities that assist in determining the source of a deposit

item. These capabilities include contrast and pseudocolor enhancement of the endorsements. As with the exception-processing application, the return-item application results in improved processing times through the use of images rather than physical paper handling.

Electronic kill. The electronic-kill application provides a mechanism for banks to process incoming and outgoing electronic media. This application runs on the host and is the key application required for interbank truncation of physical documents.

HPTS application architecture framework

The ImagePlus HPTS application architecture is based on the Systems Application Architecture (SAA) concept of cooperative processing. As such, each HPTS application consists of at least two parts: the host application (which runs on the host processor) and the workstation application (which runs on one of the two workstation types, depending upon the application).

The host application interfaces with the Check Processing Control System (CPCS) to retrieve coded data and processing instructions. The host application also interfaces with IHAS to retrieve images. An Advanced Program-to-Program Communication (APPC) conversation between the host application and the LAN-based application provides the host application with the ability to perform background tasks associated with assembling data for processing on the LAN and sending those data to the LAN. This cooperative processing architecture also allows the customer to maintain all the databases (both image and nonimage) that are required to run the item-processing operation on a centralized mainframe computer, while distributing the sophisticated operator-interface functions to OS/2 workstations to take advantage of the SAA Common User Access (CUA) facilities.

The workstation application interfaces with the operator. The functions performed by the workstation application are assisted by workstation-communication or device-support services provided by either or both the Image Entry Workstation Facility and Image Extended Workstation Facility core software components.

Applications such as key entry are of a highly interactive nature and require that some of the workstation application functions be offloaded to a separate

processor on the LAN. This processor is called the workstation manager, and the application software that runs in the workstation manager station is referred to as the application manager.

The application manager monitors the processing status on the LAN. Through the use of the Image Management Workstation Facility core software routines, the application manager (1) determines which images and coded data are available for processing by the workstation application; (2) determines the coded data that have been processed and sends the data back to the host application; and (3) initiates requests to the host application for more coded and image data, when processing on the LAN is nearing completion. The relationship between the image applications and the core image software is shown in Figure 7. This application architecture framework, together with the previously described portions of the HPTS architecture, form the new image platform for document processing in financial institutions for the 1990s.

Concluding remarks

The IBM ImagePlus High Performance Transaction System (HPTS) was architected and designed with the assistance of the IBM Research Division and six IBM customers in the financial services industry. The primary functions of HPTS are high-speed image capture, image database management, numeric hand print and machine print character recognition, and image enhancement at a workstation.

The initial implementation of HPTS is aimed at check-processing customers. However, all the functions provided by the system are required to support high-speed image and transaction processing applications, which extend well beyond check processing. HPTS should be readily extendable to other applications that today involve high-speed processing and manual handling (data entry) from paper documents such as, remittances, tax documents, mail-order forms, census data, standardized tests, manufacturing receiving and distribution functions, and many more.

Standard system architecture components (such as MVS, OAM, VSAM, DB2, OS/2, image scanners, and displays) were chosen with these needs in mind. The implementation of Systems Application Architecture (SAA), and particularly cooperative processing, APPC, and Common User Access (CUA), in conjunction with image processing technology makes HPTS

a viable solution for the 1990s and beyond for many high-speed image and transaction processing needs.

The authors have attempted to describe the results of a project that spanned over four years in time and several hundred person years of effort. The application addressed by this system, that of check processing, had previously been automated via magnetic character recognition and standard data processing hardware and software. The migration to an automated image platform based on standard strategic hardware and software was both natural and required by the massive increase in the use of bank checks.

Future expansions of this technology will be driven by the customers' business and operational needs. We have laid the foundation for these extensions that we expect to continue to follow the strategic system direction discussed in this paper. We believe that the large object content, high transaction rates, and rich image processing requirements will begin setting the pace for the data processing strategies.

ImagePlus, Systems Application Architecture, SAA, MVS/ESA, System/370, and DB2 are trademarks, and OS/2, PS/2, and Micro Channel are registered trademarks, of International Business Machines Corporation.

Intel and 80386 are registered trademarks of the Intel Corporation.

Acknowledgment

The authors want to thank W. Leighton Carmichael, Director of Document and Image Processing, for his executive leadership throughout the project.

Cited reference

1. J. M. White and G. D. Rohrer, "Image Thresholding for Optical Character Recognition and Other Applications Requiring Character Image Extraction," *IBM Journal of Research and Development* 27, No. 4, 400-411 (1983).

General references

- R. B. Arps, T. K. Truong, D. J. Lu, R. C. Pasco, and T. D. Friedman, "A Multi-Purpose VLSI Chip for Adaptive Data Compression of Bilevel Images," *IBM Journal of Research and Development* 32, No. 6, 775-795 (1988).
- Check Processing Control System General Information Manual*, GH20-1008, IBM Corporation; available through IBM branch offices.
- Image Object Content Architecture Reference*, SC31-6805-0, IBM Corporation; available through IBM branch offices.

ImagePlus High Performance Transaction System, General Information Manual, GC31-2706-0, IBM Corporation; available through IBM branch offices.

Mixed Object Document Content Architecture Reference, SC31-6802-0, IBM Corporation; available through IBM branch offices.

W. B. Pennebaker, J. L. Mitchell, G. G. Langdon, Jr., and R. B. Arps, "An Overview of the Basic Principles of the Q-Coder Adaptive Binary Arithmetic Coder," *IBM Journal of Research and Development* 32, No. 6, 717-726 (1988).

Systems Application Architecture, Common Communications Support Summary, GC31-6810-0, IBM Corporation; available through IBM branch offices.

Raymond F. Dinan *IBM Services Sector Division, 1001 W. T. Harris Boulevard, Charlotte, North Carolina 28257.* Mr. Dinan is presently responsible for technical management of high-performance transaction systems at the IBM laboratory in Charlotte, North Carolina. From 1970 to 1985 he had management responsibilities on several development projects for high-speed document processors. He directed development efforts on the IBM 3890 Document Processor and the follow-on product engineering and feature development activities. Development of the IBM 3890 Magnetic Ink Character Recognition system and the Document Coding Feature Recognition were also his responsibility. Other development activity included work in the areas of ink-jet printing, ultrasonics, high-speed paper handling, and coding. From 1963 to 1970 Mr. Dinan worked on the development of earlier-technology IBM check-handling and processing machines. Technical responsibilities included character recognition systems development for the E13B and CMC-7 fonts. Mr. Dinan has been recognized for his technical and managerial contributions by several IBM awards. He has published several technical articles and invention disclosures. Mr. Dinan received a B.S. in electrical engineering from the Newark College of Engineering, Newark, New Jersey in 1962. (The Newark College of Engineering has been renamed the New Jersey Institute of Technology.) Mr. Dinan received an M.S. degree in computer science in 1972 from the State University of New York at Binghamton.

Lynn D. Painter *IBM Services Sector Division, 1001 W. T. Harris Boulevard, Charlotte, North Carolina 28257.* Lynn Painter received her B.S. degree in information systems from Virginia Commonwealth University, Richmond, Virginia, in 1982. Upon graduation, she joined IBM in Charlotte, North Carolina, as an applications programmer, developing software solutions for IBM's internal accounting and financial systems. In 1985, Ms. Painter participated in a workstation printer marketing support program with the IBM National Distribution and Information Products Divisions. In 1986, after returning from her marketing assignment, she joined the document processing systems organization as the image software manager. Since that time, she has managed the development of the image system and application software. This includes managing more than one hundred programmers and managing the partnership development project for the High Performance Transaction System with six IBM customers in the financial services industry. In 1988, Ms. Painter received her M.B.A. degree from the University of North Carolina in Charlotte, North Carolina.

Robert R. Rodite *IBM Services Sector Division, 1001 W. T. Harris Boulevard, Charlotte, North Carolina 28257.* Mr. Rodite received his B.S. degree in electrical engineering from Lafayette College, Easton, Pennsylvania, in 1964 and his M.S. degree in electrical engineering from the California Institute of Technology, Pasadena, California, in 1965. He joined IBM in Endicott, New York, in 1965. Prior to 1981, he was responsible for several electronic technology projects in IBM. From 1981 to 1985, Mr. Rodite had IBM program management responsibility for high-speed document processors, including the IBM 3890. During this period, his organization developed new models of the IBM 3890, including development of technologies for image digitalization, ink-jet printing, and real-time processing. In 1985, he initiated the image system project. Mr. Rodite has received IBM Invention Awards in recognition of his patent activity, and he has published several technical papers since joining IBM. Mr. Rodite is currently managing the overall development of document processing image systems at the IBM laboratory in Charlotte, North Carolina.

Reprint Order No. G321-5410.