

Architectural Design for Address Interpretation in a Laboratory Address Recognition Unit (LARU)

Jonathan J. Hull, Victor Demjanenko, Ramalignam Sridhar, Sargur N. Srihari

Center of Excellence for Document Analysis and Recognition (CEDAR)

State University of New York at Buffalo

Buffalo, New York USA

Abstract

The system architecture used for address interpretation in the CEDAR LARU is defined. Gray level address block images are input to the system and ZIP+4 Codes are output.

The objective of the CEDAR address interpretation unit (AIU) design is to provide a platform that maximizes system throughput to allow for large-scale testing of algorithmic improvements in a laboratory environment. Thus, the time to put code developed by algorithm researchers into a testable prototype is minimized. A secondary objective is to build a system that has the *capability* to reach real-time processing rates (13 mailpieces per second).

This paper outlines the basic architecture of the CEDAR AIU. Specialized image processing is performed on custom-designed hardware. This includes the reduction of image data to symbolic form, i.e., a serial representation of the contours of an address block image. All the subsequent processing of the address is performed on a bank of general purpose processors. An implementation that uses a commercially available multiprocessor board running a variant of the Unix operating system is described. This allows algorithm developers to directly compile code into the AIU test environment from development workstations. The multiprocessor also supports up to 512 megabytes of memory (DRAM) that is used to store the nationwide ZIP+4 and DSF files.

Programming Language. Programming a GPP is facilitated by general purpose programming languages, such as C, that are well-known and suited for real-time tasks.

1. Introduction

The goal of the address interpretation unit (AIU) design is to provide a development platform that will allow for the demonstration of address interpretation algorithms on large numbers of images on a regular basis. The near real-time (NRT) rate has been established at ten addresses per minute. This will allow us to run 10,000 images in about 16 hours. An additional goal is to demonstrate the feasibility of extending the NRT-AIU to process 13 addresses per second (or 77 milliseconds per address) with appropriate reasonable latency.

The NRT-AIU will facilitate testing of algorithmic improvements on large numbers of images. An example of the need for this approach is the current processing speed of the handwritten address interpretation software system. Currently, a five-digit accept rate of about 70 percent with an encode rate of ten percent is achieved at 150 seconds per address. This places a practical restriction of 500 images in an overnight run.

The rest of the paper describes the design philosophy of the NRT-AIU, the implementation of several algorithms in specialized hardware, and how the current design can be extended to a complete real-time system.

2. Design Philosophy

The design philosophy for the CEDAR AIU is illustrated in Figure 1. Image operations are separated from logical operations. Image operations are performed on special-purpose hardware and logical operations on general-purpose processors (GPPs) in a general-purpose programming language. This approach is based on both data-transfer rates required for image operations as well as speed and economy provided by GPPs.

Image Operations. Images contain large amounts of data, e.g., a typical gray-scale (8 bits per pixel) address image at 300 ppi needs 674Kbytes. Conversion from gray-scale to binary may require two passes: the first to compute a threshold by summing pixel values into a histogram, and a second to apply the threshold. Thus, over 1.3 million image accesses would be needed. Such operations are best done at high rates by special-purpose hardware. The MAXbus standard (developed by the Datacube company) is an appropriate framework for high-speed pipeline processing. The MAXbus is an image data bus with a transfer rate of 10 million pixels per second.

Logical operations involve much lower data volume than image operations. For instance, grouping connected components is an important logical operation. An image of 674Kbytes with 125 connected components would generate only 4Kbytes of connected component data (assuming 32 bytes per component).

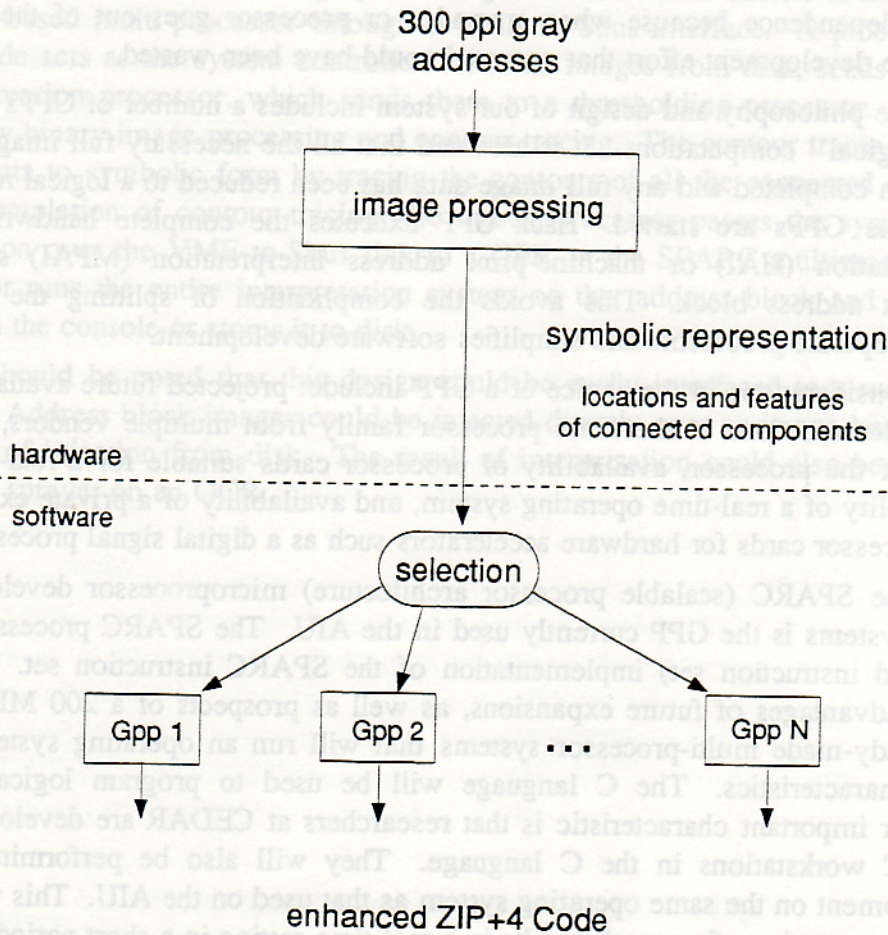


Figure 1. Design philosophy of the address interpretation unit.

GPP. A GPP offers advantages of speed, upgradability, and ease of programming for logical operations -- features not readily available in an alternate technology. Typical current micro-processors operate at rates of 28 to 100 million instructions per second (MIPS). Processors with speeds of 200 MIPS have been demonstrated and projected to be commercially available in a few years. A system with GPPs will be able to achieve a more than seven-fold speedup by merely exchanging processors: an address image requiring a half second with a 28 MIPS processor will require only 71 milliseconds with a 200 MIPS processor.

Programming Language. Programming a GPP is facilitated by general purpose programming languages, such as C, that are well-known and suited for real-time tasks.

Another advantage of C-programming is code transportability. The C code only needs to be re-compiled for a new processor. This is important since new vendors and faster processors will eventually appear on the market. It is desirable to avoid vendor or processor dependence because when a vendor or processor goes out of the market, the software development effort that occurred could have been wasted.

The philosophy and design of our system includes a number of GPPs for performing "logical" computations. It is assumed that all the necessary full image processing has been completed and any full image data has been reduced to a logical representation when the GPPs are started. Each GPP executes the complete handwritten address interpretation (HAI) or machine-print address interpretation (MPAI) system on a different address block. This avoids the complication of splitting the computation across separate processors and simplifies software development.

Considerations in the choice of a GPP include: projected future availability of faster implementations in the same processor family from multiple vendors, programmability of the processor, availability of processor cards suitable for a real-time system, availability of a real-time operating system, and availability of a private external bus on the processor cards for hardware accelerators such as a digital signal processor (DSP).

The SPARC (scalable processor architecture) microprocessor developed by Sun Microsystems is the GPP currently used in the AIU. The SPARC processor is a RISC (reduced instruction set) implementation of the SPARC instruction set. The SPARC offers advantages of future expansions, as well as prospects of a 200 MIPS processor and ready-made multi-processor systems that will run an operating system with real-time characteristics. The C language will be used to program logical operations. Another important characteristic is that researchers at CEDAR are developing code on SPARC workstations in the C language. They will also be performing their code development on the same operating system as that used on the AIU. This will allow for the demonstration of research results in a real-time setting in a short period of time.

The future availability of faster SPARC implementations is evident. Current systems operate at 28 to 100 million instructions per second (MIPS). Boards based on a 200 MIPS chip should be available in a few years. The use of such a board would give nearly an eight times speedup over a system that used the 28 MIPS processor. Such an upgrade would be very simple: just unplug the old processor card and plug in the new one.

The advantages of cards that use the SPARC processor are clear. However, it should also be clear that if some other GPP was introduced that outperformed the SPARC, our software could be easily ported to it. Since the software is written in C, we would only need to recompile it on the new processor. Even though some operating system calls might need to be changed, there is nothing in the software that would limit it to run on the SPARC.

3. AIU System Implementation

The implementation of the AIU is summarized in Figure 2. A single VME chassis is used to hold the custom image processing hardware. This chassis is interfaced to a SPARC-based multi-processor through a VME to Sbus interface. A processor on the VME side acts as the system controller. It reads images from disk, sends them to the decompression processor, which sends them to a thresholding processor. This is followed by binary image processing and contour tracing. The contour tracer converts the image data to symbolic form by tracing the contours of all the connected components. Upon completion of contour tracing, the master processor passes the symbolic image description over the VME-to-Sbus link to a GPP on the SPARC multiprocessor. This processor runs the entire interpretation system on that address block and displays the result on the console or stores it to disk.

It should be noted that this design could be easily interfaced to a working OCR system. Address block images could be injected directly over an image bus connection in place of injection from disk. The result of interpretation could also be passed to a bar code sprayer on an OCR.

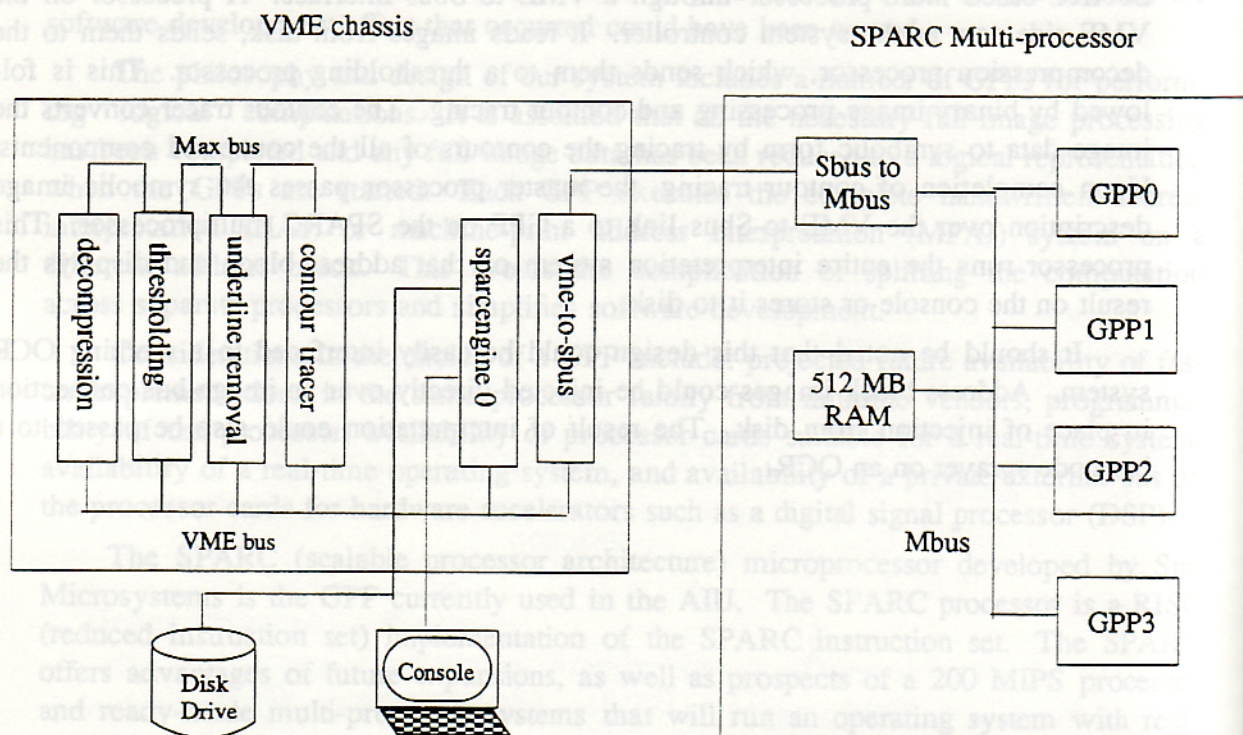


Figure 2. Address interpretation unit implementation.

The feasibility of this design is determined in part by the transfer rates needed for the near real-time system (six seconds per image). An analysis of 500 300-ppi (pixels per inch) handwritten address block images shows that an average compressed 8-bit gray scale image contains 80k bytes. The transfer of this data from disk over the VME bus to the decompression processor is easily accommodated since a transfer rate of only about 16k bytes per second would be needed. The output of the decompression processor will then be passed over the MAXbus. Since the average uncompressed image contains 674Kbytes, a transfer rate of only about 112K bytes per second will be needed. This is easily accommodated on the MAXbus. The same transfer rate will be needed from the thresholding hardware to the underline removal processor and from the underline removal processor to the contour tracer.

The projected data output rate requirements of the contour tracer were also calculated from the 500 images. About 162K bytes are needed to describe the contours of the average image. This implies a data-transfer requirement of 27K bytes per second

on the VME-to-Sbus link to process images at the rate of six seconds per piece. This is easily accommodated with standard commercially available hardware.

The remainder of this section of the paper describes the implementation of several of the image operations: decompression, thresholding, underline removal, and contour tracing. These operations have been demonstrated to be the minimum essential image processing algorithms. However, this does not preclude the development and integration of other image processing operations. The architecture is flexible in this respect and allows for such upward compatibility.

Decompression Processor

Several gray scale image compression techniques have been implemented on commercial chips. Among these are Group 4 (G4) fax and JPEG¹. We have also considered a decompression processor strictly for binary images. This is a valuable addition to the AIU design since greater compression ratios can be achieved for binary than for gray scale images. This allows for the storage of many more address images in a given amount of disk space.

A binary image decompression processor for facsimile images would receive input over the VME bus and output address block images over the MAX bus. Effective real-time rates of 13 pieces per second could be achieved. An important characteristic of this portion of the system is that it retains compatibility with some of the OCR systems currently used by USPS that can record compressed images in facsimile format from live mail. Thus, we could demonstrate our algorithmic developments on the same image databases that would be processed on OCRs.

Thresholding

The thresholding algorithm currently used in our software HAI and MPAI systems is based on an analysis of the gray level histogram of an image. This technique has been used with over 4000 gray-scale address images scanned at 300 ppi and found over 98 percent effective at calculating the correct threshold. Therefore, this algorithm is being used in the AIU.

The algorithm assumes that it is given an essentially bi-level image. This is a very common characteristic of address block images. The algorithm determines a single fixed threshold for the input image from the histogram of its gray-scales. The threshold value is statistically determined to minimize the within-class variance of the "black" and "white" portions of the histogram and maximize the between-class variance of these groups.

¹ For the past few years, a standardization effort known by the acronym JPEG, for Joint Photographic Experts Group has been working towards the first international digital image compression standard for both gray-scale and color. The "joint" refers to collaboration between CCITT and ISO.

The algorithm contains three steps:

1. calculate gray-scale histogram,
2. calculate fixed threshold from the histogram,
3. apply threshold and generate binary image.

We have already implemented this algorithm using off-the-shelf image processing boards. We have also designed a dedicated hardware solution.

The hardware architecture of the thresholding processor uses the LSI Logic histogram/Hough transform processor (L64250). The chip computes the histogram required by the algorithm. A micro-processor reads the contents of the histogram and calculates the global threshold. The threshold is applied to the output channel. The micro-processor is programmed in a high level language. This allows for maximum flexibility in algorithm development. Experimentation with different methods for histogram computation is possible by modifications to the code that runs on the micro-processor.

Underline Removal

We have considered several algorithms to remove underlines from address images. The use of any of these techniques is motivated by the presence of machine-printed underlines in addresses and because they interfere with the location and recognition of the ZIP Code and other information [1]. Indeed, it was shown (on 500 handwritten address images) that removing underlines can add up to five points to the ZIP Code recognition rate.

We have considered several alternative hardware implementations for underline removal. The most promising method we have experimented with is based on the Hough transform. The algorithm works by calculating the Hough transform at several angles around 0 degrees (the resolution is a parameter). Peaks in the transform space are located and taken to indicate the presence of an underline in the image. The underline is removed by changing pixels at the corresponding locations in the image to white.

It is feasible to implement the Hough transform in hardware. A chip is available that computes the Hough transform (LSI Logic L64250) directly at a single angle. An input image could be simultaneously sent to a number of Hough transform chips in parallel. Each chip would calculate the Hough transform at one angle. The specific angles would be chosen to represent the range of skew in over 99 percent of the mailpieces encountered in a typical mailstream. Underlines would be detected by software running on a dedicated micro-processor that locates peaks above a threshold at a particular angle. The corresponding underlines would then be located and removed. The skew of an address block can also be determined by locating peaks in the Hough transform output. The skew can be corrected by rotating the image.

Since the Hough transform chips would operate in parallel, the time required to process an image should be proportional to image size (assuming 10MHz data transfer rates on the MAXbus). Thus, real-time performance is expected.

Contour Tracer

The transformation from the image domain to the symbolic domain is performed with a board based on the LSI Logic L64290 chip. This chip is given a binary image of up to 1024 by 1024 pixels and locates the contours of the objects in the image. It returns the sequence of X,Y coordinates and discrete curvature values for each contour. The bounding box, area, and perimeter of each contour are also returned. Figure 3 shows an example of the input and output of the chip. A binary address block such as the one in Figure 3(a) is input to the chip. It computes the contours shown in Figure 3(b). The output representation for these contours is shown in Figure 3(c). The data for the outer and inner contours of "D" are illustrated.

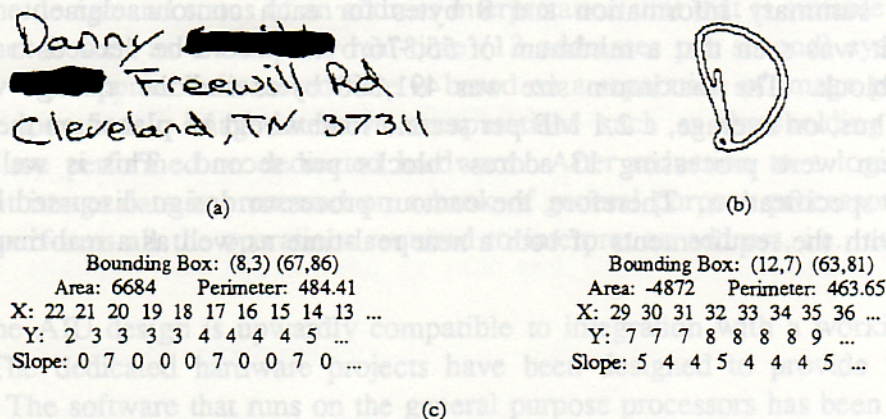


Figure 3. Example of contour tracer input and output.

The image size constraint of the chip (1k by 1k maximum) is not a difficult problem to handle. An analysis of 500 handwritten address block images that were scanned at 300 ppi showed that all of the images had less than 1024 rows and only four of the images had more than 2048 columns. In each of these four cases, all of the ZIP Code was contained in the rightmost 2048 columns. Because of this analysis, the contour tracing system was built to process images with a maximum size of 1024 rows and 2048 columns. Only the rightmost 2048 columns will be processed in images that are larger than this. This will be done in the NRT-AIU by splitting an input image at a seam that is 1024 columns from the right end of an image and sending both portions through the chip consecutively. The two outputs will be "stitched" later on by software. In a real-time system, two identical contour tracers could easily be used with each operation on its own 1Kx1K path. The same stitching technique will then be used on their outputs.

The speed of the contour tracer is compatible with a full real-time implementation that operates at 13 addresses per second. The execution time of the contour tracer chip for a typical image is less than the product of the image size and the clock rate. In a full real-time implementation, we would use a two board implementation in which each board processes its own 1Kx1K patch of the address block in parallel. Since the average 300 ppi address block image contains 674K pixels, 67 milliseconds would be needed to transfer the image to both contour tracers over the 10MHz MAXbus and then each contour tracer chip would require only 15,000,000 seconds or 23 milliseconds to process the image. This is well within the 77 milliseconds (13 address per second rate) provided for the pipeline processing of an address. Also, it provides a safety factor in case sections of the image are traced at a slower than typical time.

The output requirements of the chip were estimated from a set of 500 handwritten address block images that had been scanned at 300 ppi. Each component will need 16 bytes of summary information and 8 bytes for each contour element. In the 500 images, it was seen that a minimum of 55,376 bytes would be needed to represent an address block. The maximum size was 491,368 bytes and the average was 162,584 bytes. Thus, on average, a 2.1 MB per second load would be placed on the VMEbus if the system were processing 13 address blocks per second. This is well within the VMEbus specification. Therefore, the contour processor design discussed here is compatible with the requirements of both a near real-time as well as a real-time implementation.

We have considered several alternative hardware implementations for underline removal. The algorithm is based on the Hough transform. The algorithm works by calculating the Hough transform at angles around 0 degrees (the resolution is a parameter). Peaks in the transform are located and taken to indicate the presence of an underline in the image. The image size of the chip (4Kx4K) is a maximum of 16Kx16K. An analysis of 500 handwritten address block images that were scanned at 300 ppi showed that all of the images had less than 1024 rows and only four of the images had more than 2048 columns. In each of these four cases, all of the contour elements were contained in the remaining 2048 columns. Because of this analysis, the contour processing system was built to process images with a maximum size of 1024 rows and 2048 columns. Only the remaining 2048 columns will be processed in images that are larger than this. This will be done in the VRT-AIU by splitting an input image at a 1024 column from the right end of an image and sending both portions to the chip consecutively. The two outputs will be stitched back on by software. In a real-time system, two identical contour tracers could easily be used in operation on its own 1Kx1K patch. The same stitching technique will then be used to process an image should be proportional to image size (assuming 10MHz data transfer rates on the MAXbus). Thus, real-time performance is expected.

4. Extension to Real-Time

The basic design of the NRT-AIU is upwardly compatible with an implementation that runs at 13 pieces per second (or 0.077 seconds per piece). Each of the image processing operations can be made to run at the required rates. In some cases, duplication of selected hardware would be needed. However, most importantly, there is no inherent theoretical limit on bus transfer capacity that would be exceeded by increasing the image throughput rate. It is also clear that the capability will exist to increase the speed of the software that runs on the GPPs to the 13 piece per second rate. The speed improvement will come from improvements to the software as well as improvements in the speed of the processors themselves.

5. Summary and Conclusions

The design and status of an address interpretation unit that is suitable for both near real-time as well as a complete real-time (13 addresses per second) system was discussed. The approach discussed here is based on a separation of image processing and logical operations. Image-intensive computations such as thresholding and contour tracing are performed on dedicated hardware. After reduction to a logical form, the resultant image data is processed on a bank of general purpose processors. Each processor performs all the operations required to interpret an address, i.e., assign a ZIP+4 Code.

The AIU design is upwardly compatible to integration with a working OCR system. The dedicated hardware projects have been designed to provide the necessary speed. The software that runs on the general purpose processors has been demonstrated to have complete real-time capability. Future software work will include the transfer of performance improvements from research to the AIU system. The AIU will be used as a platform to demonstrate the latest research results on large numbers of images.

References

1. *Advanced Research in Handwritten Address ZIP Code Recognition: Interim Report*, SUNY at Buffalo, December, 1990..