

Postal Reply Card Processing

David C. Bartnik, Venugopal Govindaraju, Sargur N. Srihari
CEDAR, SUNY at Buffalo
UB Commons, 520 Lee Entrance
Amherst, NY 14228
bartnik@cedar.buffalo.edu

Binh C. Phan
US Postal Service
475 L'Enfant Plaza SW
Washington, DC 20060-6236
bphan@email.usps.gov

Abstract

A Reply Card Processing (RCP) Service was developed at CEDAR to demonstrate as to how the United States Postal Service could deliver business and courtesy reply cards to clients electronically instead of physically. The electronic delivery of the more than one billion cards that the USPS processes in a year was to expedite the order fulfillment cycle for participating customers, resulting in cost savings and accelerated response times.

CEDAR developed the necessary subsystems which included a Postal Image Management System. As well as providing the networking applications for ensuring the seamless transfer of card images to the Postal clients, the task involved advancing the state-of-the-art in forms processing to read the card images and to deliver the information in a form ready for database integration.

1 Introduction

Until the recent past, the United States Postal Service (USPS) has been the sole vehicle for delivering nearly all document correspondence with the exception of overnight services. With the explosion in popularity of commercial World Wide Web sites, the public can now instantly place an order for their favorite magazine, a pair of shoes, and even a box of handmade cigars. These methods for processing orders give businesses two substantial benefits: order processing costs are greatly reduced by avoiding much of the paper handling typically involved in mail orders; and higher customer satisfaction with faster processing of orders.

A popular method for initiating orders in the past has been through business reply mail, whereby the USPS only charges businesses for the postage used by customers sending in orders. Roughly one billion business reply cards and envelopes pass through post offices throughout the country each year. In order to compete with electronic order fulfillment, which continues to reduce the demand for typical

paper mail, the USPS must look at methods for providing electronic mail delivery service to its customers.

2 Reply Card Scanning

Reply cards are primarily used by companies as a convenient means for collecting orders from their customers. A customer fills out a reply card with his/her order, drops the card in a mailbox, where the USPS picks it up and delivers it to the business. The business must then organize the reply cards and enter the customer data into their order fulfillment database. This cycle of order fulfillment spans several days and incurs a significant cost in mail handling and data entry.

To design a solution to keep itself competitive with Internet and toll free phone orders, the USPS contracted the Center of Excellence for Document Analysis and Recognition (CEDAR) in the State University of New York at Buffalo in 1993 to design a low cost Reply Card Scanning (RCS) system that would convert paper business reply mail to an electronic form that could be sent directly to the mailers making use of existing postal equipment and communication networks. This system would incorporate forms processing, handwritten and machine print OCR, address interpretation, and mark sense recognition developed at CEDAR over a decade on USPS contracted projects [2, 3, 1, 4].

The RCS system incorporates existing postal equipment, but uses it in a way different from its original purpose. The mail transport typically used to read the destination address from mail pieces is used to capture the image of the information bearing (back) side of each card. These images are sent to the CEDAR Postal Image Management System (PIMS) for recognition of the customer order. The processed information is sent to a regional Electronic Post Office (EPO) where it is distributed to the destination mailers. Rejected cards are delivered to the mailers via Express Mail and handled in the standard fashion by fulfillment houses. Since almost all mail is routed through one of the major processing facilities within 24 hours of being dropped in a mailbox, processing and electronic delivery of the cards can

occur much quicker than conventional delivery.

A mailer subscribing to the RCS service would be provided with software that enables connection to the Postal Routed Network via a variable speed communication link necessary for the expected throughput. This software running on a PC/NT platform has been named the Customer Image Management System (CIMS), and provides a graphical interface to manage the transfer of data from the Postal network to the business's internal network.

Since all reply card formats are different, there needs to be a mechanism to identify each card and tell the RCS system what processing to apply to the data contained in the card. This is accomplished through the use of form definitions and a forms processing library developed specifically for the system. USPS communication to the mailers is done via the EPO. A number of EPOs are used to handle the 150 PIMS sites and the potentially unlimited number of mailers.

3 Recognition System Design

Converting data printed or written on reply cards involves applying forms processing and recognition techniques to the scanned reply card images. The PIMS recognition computer is used to process the images and send the results to the EPO. The design of the PIMS was made scalable depending on the amount of mail that a postal facility receives daily. Developed to run on a Sun/Solaris platform, the PIMS can be installed on a single processor Sun Sparc 10 for small sites or on multi-processor SparcServers which can handle many thousands of reply cards per hour.

Images received from the scanner are passed through the system as TIFF format image files. TIFF is a standard that allows multiple images to be stored within a file and provides the ability to define proprietary tags, in this case used to store RCS and other recognition information about the reply card image. The PIMS scheduler determines which of the several recognition processes running is ready for a new image. Shared memory is used to convey information between processes.

Recognition of the customer data entered on the reply card is handled through an elaborate forms processing system [2]. In designing the system, the developers needed a way to uniquely identify each card through the images captured from the mailstream. The goal was to identify 98% of the reply cards passing through the system, with accuracy constrained to one error for every 10,000 cards. It was determined that common form identification techniques such as horizontal and vertical line matching, or component analysis would not work due to the subtle differences between cards.

4 Forms Processing

4.1 Registration and Identification

The solution proved to be simple and agreeable to mailers as well. A 12 digit form identification number printed in a 12 point OCRA font is used to uniquely identify the mailer and the ID of that mailer's reply card format. To aid form registration, an 'L' bracket registration mark surrounded the form ID. The two components are placed at the bottom-rightmost corner of the card. Once the registration mark is located, skew is calculated by determining the average midpoint at several locations on the horizontal line. The location of the registration mark in the input image is compared to the location stored in the form definition to give the registration offset. The offset combined with the skew angle allows the forms processing library to accurately locate components such as lines, tags, and fields in the image.

After registration, the form identification number is read and validated against the current database of active cards maintained by the FDS site. A neural network based machine print OCRA character recognizer is used for form identification [4]. If the form identification number validation is successful, the form definition is loaded into the PIMS current form definition buffer. A cache of form definitions is used by the PIMS in an attempt to minimize the loading of form definitions as different reply cards pass through the system. Since thousands of card formats may be active in the system at any one time, it is not efficient to load all form definitions into memory at once. It is likely that a PIMS site generally sees the same n card formats during processing on a given day. If the form definition buffer is slightly larger than n , little disk access will be required during processing. This buffer is flushed periodically as new form definitions are updated from the FDS site.

4.2 Attributes

A reply card's form definition contains all the information required to process the card. The information is stored in the form of attributes ordered hierarchically starting with global card attributes. To aid the monitoring or reply card processing, the image of the reply card used to create the form definition is compressed and stored in the definition file. Again, the TIFF file format is used for form definition storage since user data can easily be added to the image data, and end users are likely to have utilities that can manipulate the files if necessary. The complete reply card is thought of as a form, while zones hold the regions of interest that will be processed individually. Zones are used to group fields, allowing attributes to be defined across a group. Fields must be contained within a zone, even if this means each zone will contain only one field. Using this

hierarchy, nearly any reply card or document form can be represented and processed by the PIMS recognition system.

In order for the RCS system to process and recognize the customer data entered on each reply card, it must first know what to expect to see on the card. This information is included in the reply card form definition in the manner of field attributes. Field attributes tell the recognition system the type of data contained within the field and what processing to apply to the field. Field types that require recognition can also describe a set of constraints to apply to the recognition of the field. Constraints are necessary to limit the number of errors passed downstream by the recognition modules. A mailer obviously knows the range of customer identification numbers that appear in their customer database, so the form definitions for this mailer's reply cards can include a check to verify that all recognized customer numbers exist in their database. Invalid customer numbers can be rejected and processed manually to see the reason for the discrepancy. Other field recognition constraints include the number of characters allowed, expected font for machine print data, and checkdigit algorithm to apply. An address interpretation attribute requests that the system compute a postal ZIP+4 Code for the customer address read from each card. The system can be expanded to include an unlimited amount of field types and attributes through the use of user defined types that are created with the form definition graphical editor.

4.3 Form Removal

Once the reply card image has been identified and its form definition loaded, all preprinted background lines, tags, and graphics must be removed leaving just the customer entered data. Lines consist of horizontal and vertical strokes making up underlines and tables, while tags represent everything else, such as text and graphics. The goal of line removal is to erase all pixels that belong to the line without damaging strokes passing through the line. The PIMS recognition processing locates lines within a field and removes the horizontal and vertical strokes of the lines in its first pass. Any strokes broken by the line removal are then repaired in a second pass. Line adjacency graph (LAG) representations of the zone image components are used to perform these operations quickly.

5 Recognition of Field Data

Mailers have several options regarding the processing of customer data on participating reply cards. Some may wish to only receive the full image of the scanned cards and apply their own image based data entry or OCR systems. It is likely that most mailers would request full processing of the customer fields on the cards, requiring form removal and

recognition. Text fields make up the bulk of the fields processed, including address and order fields. Mark sense such as checkboxes are also common and often can be grouped together to adhere to a set of rules. An example of this would ask the customer to check only one box in a list of choices. If more than one box appears to be checked, the recognition results are suspect.

The recognition system collects field data as specified in the reply card form definition. As each field is encountered its image is extracted, the appropriate recognizer is called with the field recognition attributes, and any constraints specified in the attributes are applied. Stray marks entering a field from outside the field boundaries are removed by first determining average size of the field components and then removing those that touch a boundary, are smaller than average, and enter a certain percentage into the field.

One benefit to mailers from using the RCS system is automatic address correction and postal ZIP+4 Code generation using the latest USPS delivery point directories for customer address fields. The recognition system can perform address interpretation on machine print and handwritten fields [3, 1]. With reply cards, the exact location of the address is known when the card is identified, completely removing the cost and errors incurred with typical address block location. The print type and, in many cases of machine print, the font of the text is known for the reply card. Nearly all handwritten address fields are separated by line and address words, to facilitate more efficient data entry. Breaking the address apart up front gives the system valuable recognition clues and reduces the effort required by parsing routines.

Typically, address fields are separated into 3 lines: customer name, street line, and city/state/ZIP line. Two common variations to this are the addition of an apartment field after the street line, and the breaking up the last line into separate city, state, and ZIP Code fields. Since the location of each address word field is stored in the reply card form definition, the address interpretation routines now can process the address lines, one at a time, starting with recognition of the ZIP Code. The recognized ZIP Code is then used to generate a lexicon of city/state pairs. Using the segmented city/state image and the new lexicon, a word recognizer is called to verify the validity of the ZIP Code recognition results. The recognition confidences of the ZIP Code and city/state images are thresholded to reduce recognition errors caused by poor quality images.

The second phase of address interpretation, looks at the delivery, or street, line to determine the ZIP Code's four digit add-on. This is accomplished in a similar manner to ZIP Code verification, but in the case of street line recognition, the word recognizer directs the system to the value of the add-on associated with the recognized street word. Given a ZIP Code and street number, CEDAR designed query in-

terfaces that work with the USPS delivery point file directory produce a lexicon of streets valid for the combination. The street word image is segmented from the street line, and the word recognizer is applied to the image given the street lexicon. The result is not only the name of the street that matched the image, but also the four digit add-on that can be returned for the complete address. In the case of post office box addresses, the system detects 'PO BOX'-like keywords and performs digit recognition of the box number. In all cases, the recognition results are verified against the USPS delivery point directory to reduce errors.

6 RCS Data at Mailer Sites

After processing at the scanning site is complete, the data produced by the system is sent to the mailers through the nationwide Postal Routed Network. Well known FTP protocols are used to pass information between subsystems, making integration and testing reliable. Since FTP can add significant overhead to the transmission of large number of small files, the RCS result files are first packed into large bundles and transferred. This allows throughput to approach 100% of the available bandwidth. Error checking provided by FTP also ensures no loss of data due to breaks in transfer.

Once a mailer's bundle of RCS result files is received at the mailer's CIMS site, the file is unbundled and sorted by reply card type. The system separates the results based on the outcome of recognition processing, pulling aside cards in need of manual verification. If a mailer has an image based reentry system already in place, it can be plugged into the data flow path of the CIMS, substituting the reject reentry interface provided with the RCS system. The data can then be automatically prepared for loading into the customer order database.

7 System Testing and Results

The RCS system was developed to capture images of reply cards, recognize the data contained within, and send the results downstream, replacing much of the paper handling typically involved. Tests involving the processing of USPS generated Change of Address cards were conducted to evaluate the capability of the system with documents other than standard reply cards. Drop out ink was added to the designs of this card, reducing the forms processing costs incurred with the traditional format. Initial testing of hand print Change of Address cards showed resolution of addresses to a five digit ZIP Code for nearly 70% of cards processed, with 48% resolved to the street or post office box level. Further modifications of the system to take more advantage of the layout of these cards would likely yield even greater performance.

Over the course of two years, the CEDAR RCS system has participated in multiple field tests involving several large volume mailers, which included Columbia House, Rodale Press, Fingerhut, Corporation, and Neodata Corporation. Throughput of the system easily met the requirement of 10,000 pieces per hour on a Sun SparcCenter 1000 for the standard reply card formats. An average form identification and registration rate of 98% was recorded with no cases of misidentification. Although the variety of card formats involved in the tests was small, it is expected that this form identification performance would hold close to these results with much larger sets. As important as recognition performance, overcoming the mis-sorts caused by physical sorting devices at the postal processing sites can prove frustrating. Feeding cards at the rate of 40,000 per hour, sorting equipment tuned to the general envelope mailstream can have difficulty with thin paper reply card stock.

References

- [1] P. B. Cullen, R. S. Curtis, P. D. Filipski, J. J. Hull, and S. N. Srihari. machine print address interpretation utilizing contextual information. pages 97–105, London, England, May 1994.
- [2] D. Niyogi, S. N. Srihari, and V. Govindaraju. Analysis of printed forms. In *Handbook of Character Recognition and Document Image Analysis*, chapter 18. World Scientific Publishing Co., Singapore, 1997.
- [3] S. N. Srihari and E. J. Keubert. Integration of handwritten address interpretation technology into the united states postal service remote computer reader system. Ulm, Germany, August 1997. Fourth International Conference on Document Analysis and Recognition (ICDAR'97).
- [4] S. N. Srihari, Y.-C. Shin, V. Ramanaprasad, and D.-S. Lee. A system to read names and addresses on tax forms. pages 5–10, Montreal, Canada, October 1994. International Conference on Document Analysis and Recognition.