AN UPDATE ON CROSFIELD DATA COMPRESSION TECHNOLOGY

Brian Jordan*

Abstract: A brief introduction will review the various approaches to data compression which have been implemented by Crosfield Electronics Limited for the effective compression of both halftone and full color continuous tone data.

Specific high-speed hardware implementations of these data compression techniques are now installed in the field, and these will be reviewed with particular emphasis on data rate and compression ratios actually achieved in practice.

The paper will conclude with a review of the application areas for data compression specifically in the fields of data communications and a system networking.

----- o -----

In a paper I presented to the Taga Annual Technical Conference in 1983, entitled Data Impression Techniques (proceedings page 689 to 699), I described two methods of Data Compression, one being appropriate for halftone, line work and text and the other for continuous tone full color images. Crosfield Electronics has now implemented both techniques in very high performance hardware, and I should like briefly to discuss in this paper the progress that has been made and performance achieved in the initial installations.

Compression for Halftone, Line Art and Text.

The most common requirement for data compression of a mixture of prescreened graphics, line art and text is in the facsimile transmission of fully made up pages as, for

*Crosfield Electronics Limited

example, is becoming common in the newspaper industry for transmission from the central publishing unit to satellite printing plants.

The compression algorithms have been implemented initially to operate in conjunction with the Crosfield Datrax Facsimile Equipment and Direct Platemaker. The Datrax consists essentially of a reader which will allow input copy to be scanned at resolutions up to 1,200 lines per inch and on the writer to reproduce the original image to the same resolution locally or remotely on Crosfield proprietory Lasermask (TM) material or direct to a printing plate.

The compression algorithm implemented is one which uses a prediction window covering pixels which have just been read together with pixels from the previous two lines in order to attempt to predict whether the next pixel will be a black or white pixel (see Figure 1). Up to eight predictors are used in parallel, and the one producing the best result is used for the actual transmission. Run lengths of good and bad predictions are then Huffman encoded and sent to the circuits which control the communications hardware. A typical installation configuration is shown in Figure 2. The data from the reader at data rates of up to 10 Megabits per second is taken to the compression system where dedicated hardware, controlled by an Intel 8080 Microprocessor, executes the prediction algorithms, codes the data, and sends it to the Data Link Interface (D.L.I.). The Data Link Interface has the job of matching the data coming from the compressor onto the communications link. The communications link may be anything from a Low Bandwidth (e.g. 9.6 Kilobit) landline to a Wide Band satellite link. The reverse configuration at the receive site decompresses the data and passes it to the writer. Because the data is heavily encoded, any errors introduced by the transmission process would cause non-recoverable errors in the decompression system and the data link interface, therefore a cyclic redundancy check sum is applied to discrete parcels of the compressed data. The data link interface at the receive site checks these check sums and advises the transmit site over a low band with return channel whether the data has been received accurately. If it has not, then the D.L.I. will retransmit that block.

If because of Low Bandwidth communication channels or multiple retransmissions due to high communications bit error rates, the communications line cannot handle with the compression system, the data rate from the reader, then the data compression system provides a command to the reader to stop reading until the compression system has cleared down its buffers.

Compression ratios achieved through this system vary enormously according to the type of material that is being scanned. We have, however, found in practice that for the typical newspaper, an average compression ratio of 12:1 can be achieved, but the actual ratios can vary from as low as 3 or 4:1 for busy halftones to more than 30:1 for text.

The installations of this new data compression hardware include the Boston Globe where the Datrax equipment together with the Data Compression hardware is being used in daily production to enable facsimile transmission of newspaper pages from the downtown publishing building to a new satellite printing plant in Billerica. An interesting configuration is the one installed for the production of the O'Neil Data Systems 'Investors Daily'. In this application, the data compression system is driven by a Monotype raster image processor which itself is provided with text data from a Hastech front end system, and graphics data from the Datrax reader. The output of the Monotype raster image processor is compressed and transmitted via satellite from the Los Angeles area publishing facility to the East Coast of the United States.

It is interesting to note that text data from the raster image processor is able to be compressed by a significant additional factor due to the 'regularity' of the data. When text is scanned on the laser scanner, even at the high resolutions of the Datrax, some jitter is always introduced by the reading system making the predictors in the compression system less efficient leading to lower compression ratios. Thus as the industry moves towards the 'computer-to-plate' environment a side benefit will be the more efficient compression and communication of data.

Compression of Continuous Tone Color Graphics

Again, the basic principles of the system algorithms were described in the Taga Preceedings in 1983, and I will just recap those principles here and go into a little more detail on the hardware implementation and field performance.

The data volumes involved in the raw data from color separation processes is very large. An A4 page scanned at 300 lines per inch and 300 picture elements per inch represents about 36 Megabytes for four colours and proportionately more at higher scanning resolutions. This means that a standard 2,400 foot magnetic tape depending on bit density which is used will hold two to four pages and a standard removable disk pack with 256 Megabytes formatted capacity, a maximum of six pages.

The Crosfield Data Compression System aims at an 8:1 average reduction of the data volumes saving the corresponding disk or tape storage space or transmission time with no visible distortion of the reconstituted image. The actual compression ratio will vary with the subject from about 4:1 to about 15:1.

Different applications will demand different compression ratios and will tolerate various degrees of distortion. It is obvious that the compression ratio can be increased at the expense of distortion. Facilities are available to tune the compression system to customer requirements specified at the time of placing the order. The compression ratio is not under the control of the operator, but is optimised for the stated requirements.

The speed of compression is in the order of one byte per microsecond. If the system operates in 'disk-to-disk' mode, i.e. the original image is stored on disk and after compression stored either on the same or another disk, the overall speed will be determined by the disk access and latency times. The compression system operates under control of a host computer and hence the overheads of the host computer operating system has to be taken into consideration. This implies that the actual compression time for a page or a particular image size has to be calculated for each configuration individually. The actual speed will be determined by factors other than the speed of compression.

Principles of Operation. The general scheme for data compression can be considered as a sequence of operations and as mapping, quantising and coding.

The mapping operation takes input data from the pixel domain and maps it into another domain where the subsequent operations of quantising, and coding can be used more efficiently in the sense that fewer bits are required to encode the map to data than the original format.

The quantiser rounds off the mapped data to a smaller number of permitted values.

The coder assigns a code word to each of the quantised data values.

The Crosfield Data Compression system uses adaptive transform encoding for the mapping operation using a transformation closely related to the well known Fourier Transform, and it can be shown that it is close to the optimum in respect of mean square error. The image is broken down into N x N size blocks where a transform is applied separately to each block row and column yielding a two-dimensional transform. In adaptive transform encoding each block is inspected and only those coefficients are retained which exceed a predetermined level, the others are set to zero. The retained coefficients after scaling and quantisation are encoded with a statistical coding technique. The block diagram of the system is shown in Figure 3 giving more details of the implementation.

The same steps are applied in the reverse order for restoration of the original image.

Physical Description. The Data Compression hardware consists of a 19 inch rack mountable cabinet which comprises five 15 inch by 15 inch printed circuit cards. Two further p.c.b.s which are D.E.C. standard sizes are located in the host computer. The data compression has its own D.C. power supply and needs only A.C. power. The hardware is divided into six major functional areas and organised in a manner such that each function occupies a separate board.

Installations.

The major installation sites for the data compression hardware is with Time Incorporated where it is used in conjunction with Crosfield Studio 800 series equipment to transmit full color magazine pages which have been made up electronically on the Studio 800 System to ten remote sites over a broadcast satellite transmission system. The system is now fully installed and operational and for a typical magazine page, including white space, data volumes are being reduced by an average of 16:1 over the raw data which is derived from the Studio 800 System. This is achieved without discernable reduction in the quality of the final printed image.

The system utilises a T-2 Satellite link at 1.544 Megabytes per second. For an average magazine page compression takes approximately 2 minutes per page, transmission approximately 1 minute 15 seconds per page and decompression 2 minutes per page and therefore a four colour page can be compressed, transmitted, and decompressed in less than 6 minutes with a throughput of four full color pages of one page every two minutes.

Using Data Compression techniques the throughput of pages over this relatively wide band communications link is in the region of thirty full color pages per hour and a bottleneck can build up at the output scanner on the remote sites. Data volume reduction using the Data Compression System has made it possible to substantially enhance the throughput of the output recording device.

In a normal magazine close the color graphics elements normally close well ahead of the text, and can be transmitted out to the remote sites well ahead of press time. However, a large number of pages of text close within the final few hours before the press start time. This text has to be merged with the black printer of the graphics transmitted to the remote sites and output on the output scanner. Using the normal output formats of the recorder this makes for a rather inefficient throughput as, either all of the color has to be held back until the text is available for that page which produces bottlenecking during the critical text close or the three colours have to be sent ahead of time and then single black printers with text retransmitted when the text becomes available. In order to overcome this problem it is advantageous to transmit pages in batches of three imposed so that they can be output simultaneously three around the drum on the expose recorder. As shown in Figure 5 this is a much more efficient utilisation of the output scanner allowing three 3 colour pages to be output on one piece of film on a 640M output recorder and nine pages of black printer plus text. However, the overheads involved in assembling three pages together in uncompressed form would produce file sizes of 100 Megabytes approximately per three pages and would also take an extensive period of time. Crosfield has developed for Time Incorporated a mini imposition system which allows the pages to be imposed to groups of three at the central production site after compression but before transmission.

Hence, as the color graphics pages are closed they are compressed and imposed to three page impositions ready for transmission to the remote site. At the remote site, three of the four colors are output for each page as shown in Figure 5. Meanwhile the compressed graphics pages are retained at the central production facility. As text pages become available the imposition system imposes the compressed text data, selects the appropriate three graphics pages and re-imposes those pages allowing both text and graphics to be transmitted to the remote site where the black printers plus text are finally output. Because we are in a compressed domain the imposition of three full color pages takes no more than two minutes. This has allowed the throughput using one output recorder to go up as far as eighteen pages per hour.

The above example is just one indication of the substantial improvement in processing times which can be achieved using compressed data. Now that Crosfield Electronics has a proven high speed, high quality color data compression system, we can proceed to the goal outlined in the Taga 1983 paper of achieving a true networked electronic pre-press system using compressed data throughout the system substantially reducing process times, archiving times and eliminating the need for removable disks.

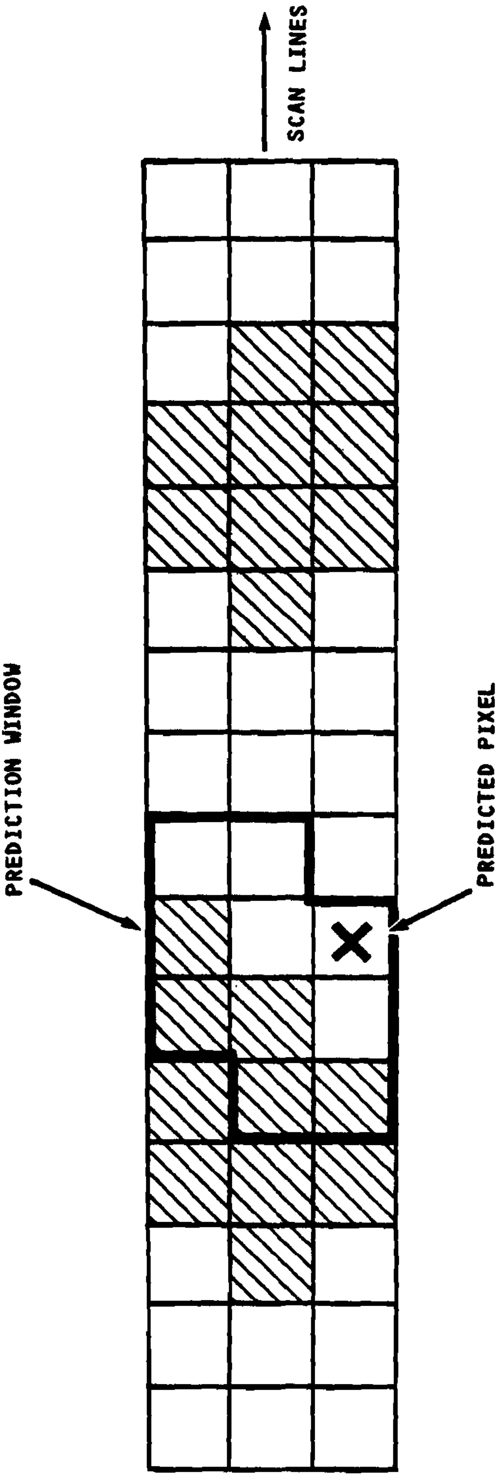

FIG. 1

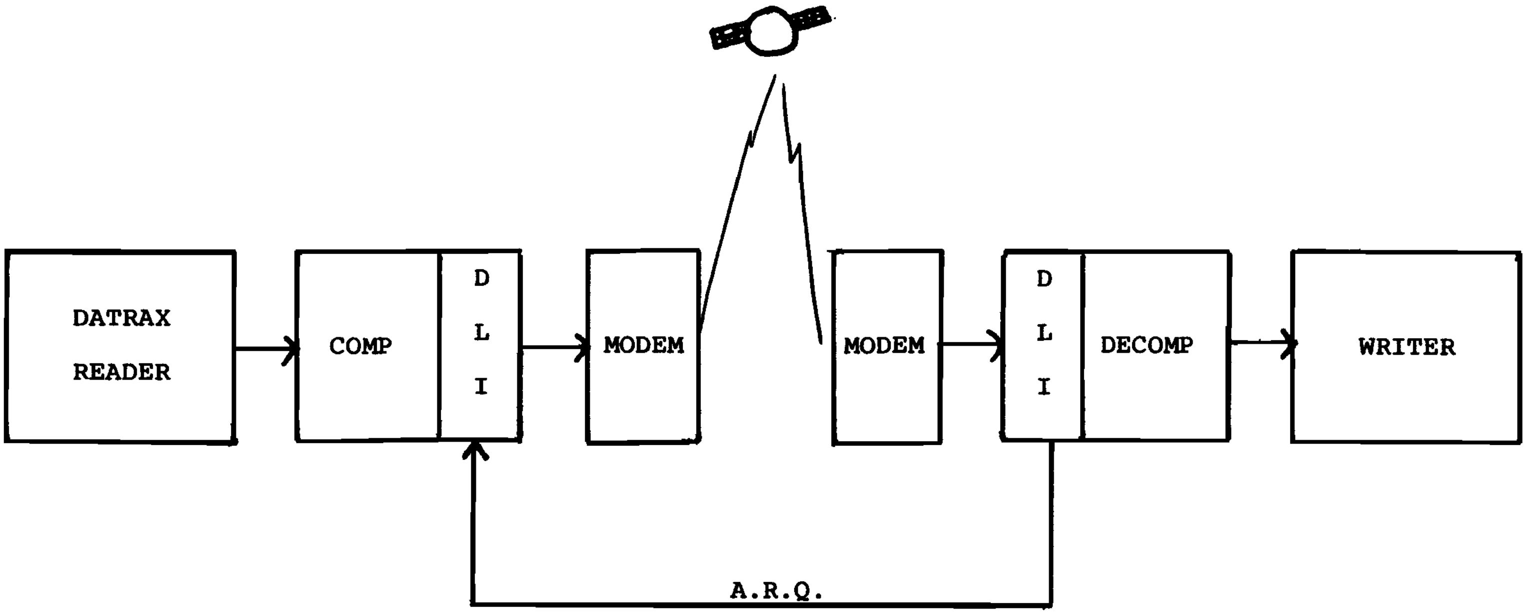

FIGURE 2

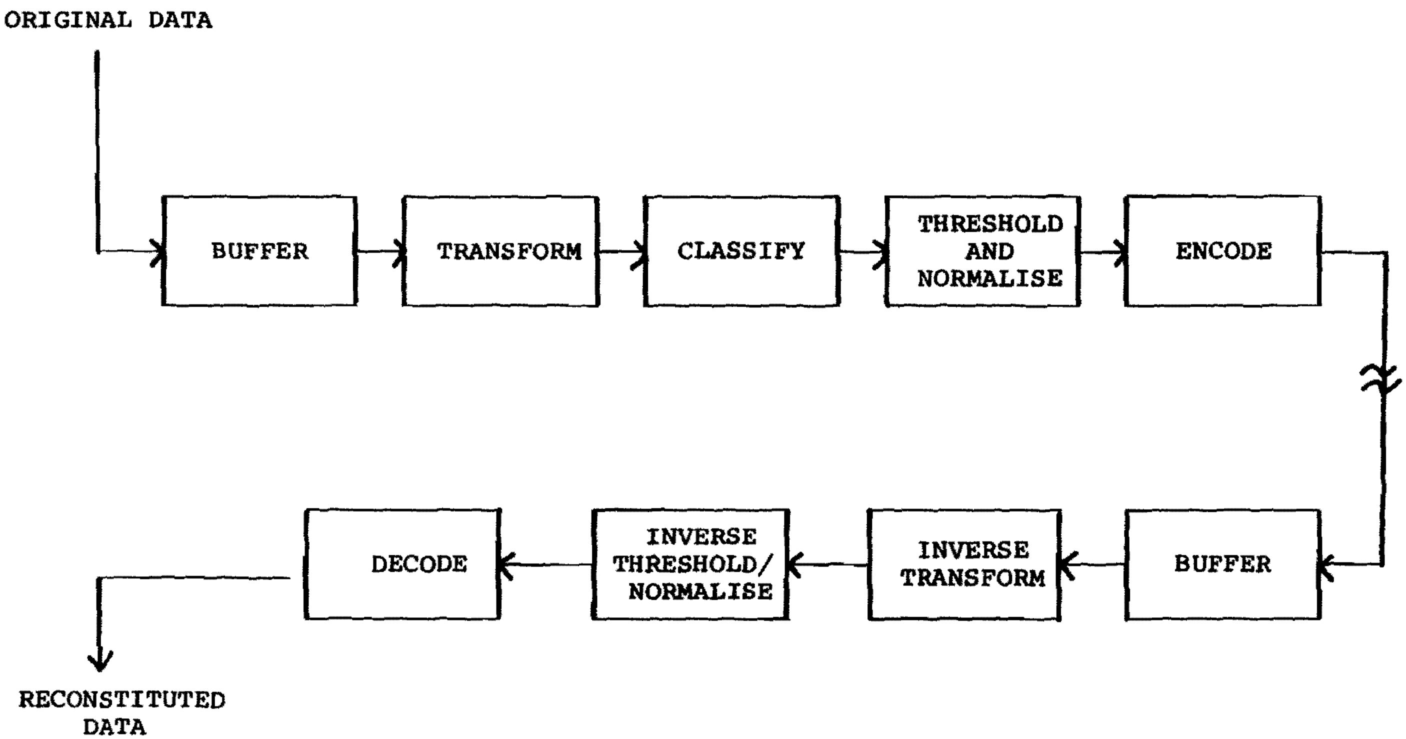
ORIGINAL DATA
BUFFER
TRANSFORM
CLASSIFY
THRESHOLD AND NORMALISE
ENCODE
BUFFER
INVERSE TRANSFORM
INVERSE THRESHOLD/ NORMALISE
DECODE
RECONSTITUTED DATA

FIGURE 3

FIG. 4

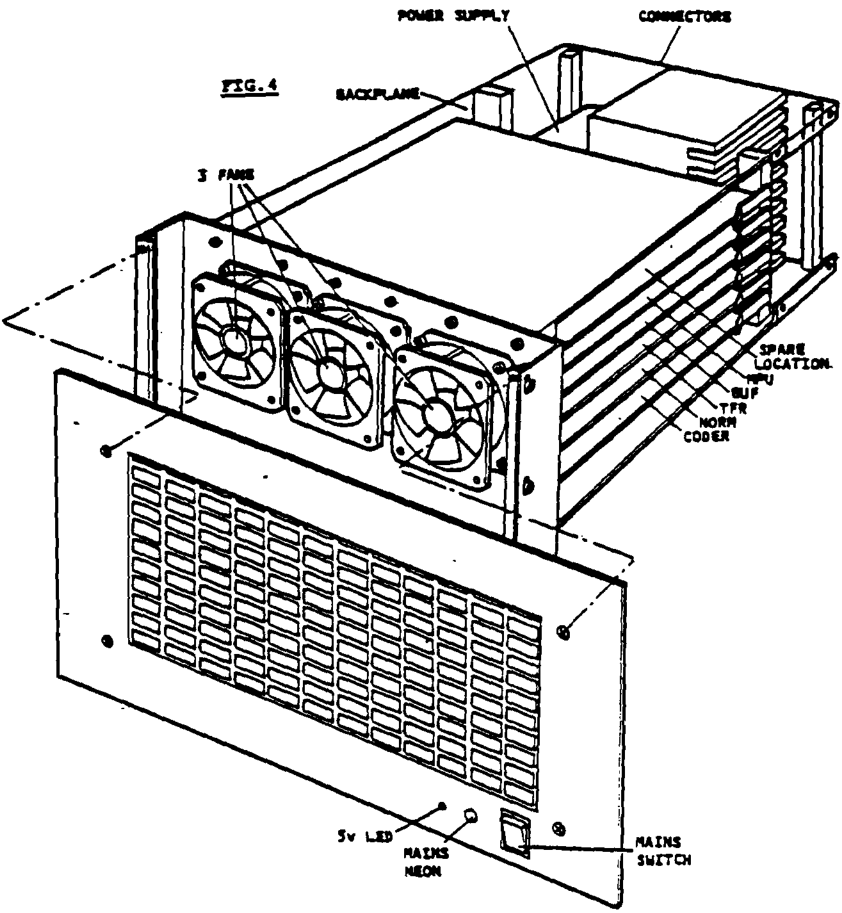
POWER SUPPLY
CONNECTORS
BACKPLANE
3 FANS
SPARE LOCATION
MPU
BUF
TFR
NORM
CODER
5v LED
MAINS NEON
MAINS SWITCH

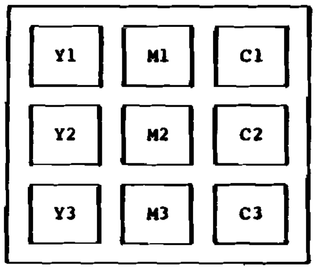

THREE COLOUR

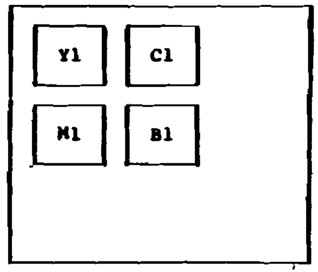

WITHOUT IMPOSITION

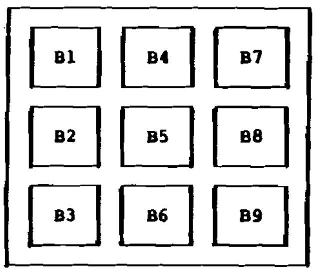

BLACK PRINTER WITH TEXT

FIGURE 5