APPLICATION OF OPTIMIZED PARALLEL PROCESSING DIGITAL COMPUTERS AND NUMERICAL APPROXIMATION METHODS TO THE ULTRA HIGH-SPEED THREE-DIMENSIONAL RECONSTRUCTION OF THE INTACT THORAX

B. K. GILBERT, L. M. KRUEGER, A. CHU and E. L. RITMAN

*Biodynamics Research Unit, Department of Physiology and Biophysics,
Mayo Foundation, Rochester, MN 55901 (U.S.A.)*

E. E. SWARTZLANDER, JR.

*Electronic Systems Division, TRW Defense and Space Systems Group, Redondo
Beach, CA 90278 (U.S.A.)*

D. E. ATKINS

*Department of Electrical and Computer Engineering, University of Michigan,
Ann Arbor, MI 48109 (U.S.A.)*

SUMMARY

In order to achieve the computational capability to carry out many thousands of cross-sectional reconstructions, necessary to support a prototype high temporal and spatial resolution cylindrical scanning multiaxial tomographic unit, a series of design, software simulation, and fabrication studies is underway to develop a special-purpose high-speed reconstruction computer. This processor will rely upon integrated circuit arithmetic components of advanced design, and highly parallel architecture to execute X-ray based transaxial reconstruction algorithms at the rate of hundreds of cross sections/sec.

SOMMAIRE

Dans le but d'atteindre la capacité de calcul nécessaire pour mener à bien plusieurs milliers de reconstitutions spatiales à partir de sections transverses, indispensables à un prototype de tomographe multiaxial, à balayage cylindrique et à haut pouvoir de résolution temporel et spatial, une série d'études de conception, de simulation

317

*logicielle et de réalisation sont en cours. Eeles permettront la mise au point d'un
ordinateur de reconstitution spécialisé, à grande vitesse. Ce processeur sera fondé
sur l'utilisation de composants arithmétiques à circuit intégré de conception récente
ainsique sur une architecture soigneusement adaptée, il permettra la mise en oeuvre
d'algorithmes de reconstitution transaxiale (utilisant les rayons X) à la vitesse de
plusieurs centaines de sections transverses par seconde.*

INTRODUCTION

Presently available commercially manufactured X-ray computer-assisted tomo-
graphy (CAT) scanners are capable of producing single cross section images of the
head, thorax, or abdomen which exhibit both high spatial and halftone resolution.
However, the operational flexibility of these devices is restricted by their inability
to record X-ray transmission data sufficient to reconstruct more than a single
10–15-mm thick cross section of tissue during a single scan, and the relatively long
duration, 10–60 sec, of each scan, which is incapable of 'freezing' the motion of
rapidly moving organs such as the beating heart and breathing lungs. Thus, the
single cross sections generated by these machines cannot provide clinical diag-
nostic or biomedical research data concerning the true anatomic shape and
dynamic changes in the structural dimensions of these important organ systems.

The development has been undertaken of an advanced cylindrical scanning
multiaxial tomography unit, the Dynamic Spatial Reconstructor, or DSR, which
will be capable of imaging the thoracic and abdominal organs with high temporal
and spatial resolution. The DSR is designed to scan over a span of 24 cm in the
cephalocaudal (z-axis) dimension of the subject, and can repeat the complete scan
procedure at 16-msec intervals, allowing the entire three-dimensional anatomic
extent, as well as the dynamic changes in shape and dimensions of moving struc-
tures such as the intact beating heart, lungs, and vascular system to be reconstructed
at a repetition rate of 60/sec (Wood, 1976; Sturm *et al.*, 1976; Sturm *et al.*, 1978;
Ruegsegger *et al.*, 1977).

Figure 1 is an artist's conception of the multiple X-ray source and detector
assemblies of a proposed Dynamic Spatial Reconstructor. The system consists of a
set of 28 independently controlled rotating-anode X-ray sources arranged around
a semicircle of radius 143 cm, and a 30-cm wide fluorescent screen bent around an
opposing semicircle of radius 58 cm. Twenty-eight sets of independently con-
trolled image intensifiers and image isocon video cameras are arranged behind the
fluorescent screen to produce 28 video images of the 30×30 cm overlapping
roentgen projection images produced on the screen.

The entire assembly of 28 X-ray sources, curved fluorescent screen, and 28 video
camera chains are mounted on a cylindrical aluminum gantry which rotates at a
constant angular velocity of 90°/sec about a horizontal axis which is co-linear
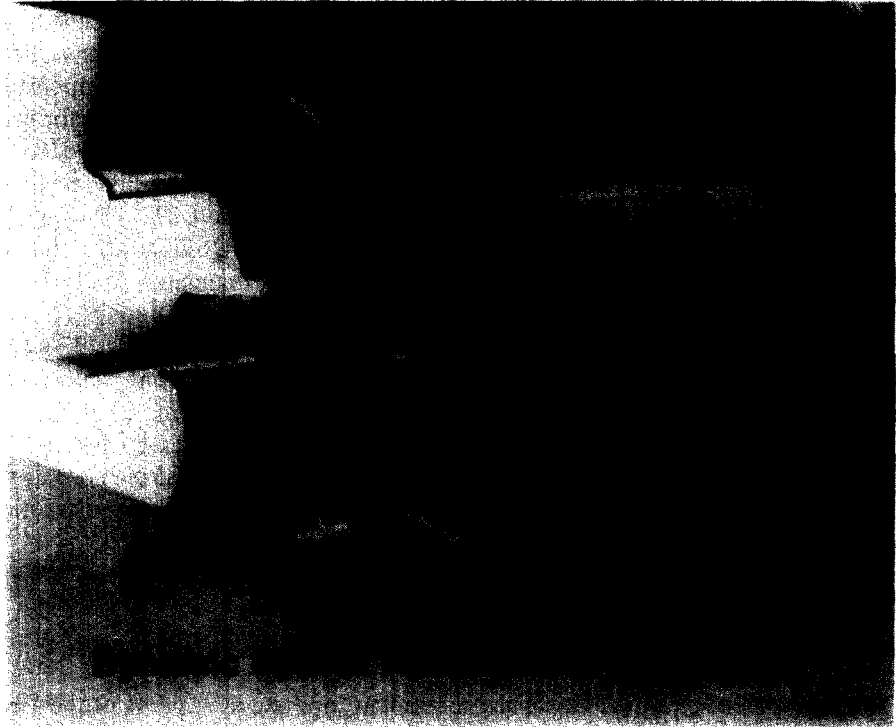
Fig. 1. Artist's conception of dynamic spatial reconstructor. Twenty-eight independently controlled rotating anode X-ray sources are positioned around lower half circumference of gantry ring. Each X-ray beam, when pulsed on, passes through subject, then creating X-ray transmission image on a portion of curved X-ray sensitive fluorescent screen. Each of 28 transmission images is scanned by 1 of 28 video camera chains mounted on upper half circumference of gantry. Cantiliver supported gantry rotates continuously; power and electronic signals pass through slip rings at far end of gantry mechanism. (Reproduced with permission from Wood *et al.* Proceedings of the Third Symposium on Coronary Heart Disease, Frankfurt, Germany (In Press)).

with the z-axis of the subject. Each X-ray source is pulsed-on sequentially to irradiate the subject for a duration of 360 μsec, while simultaneously, its associated video image intensifier and video camera are activated so as to record the image formed on the diametrically opposite portion of the curved fluorescent screen. Each scanning sequence provides a sufficient number of X-ray projections to reconstruct a three-dimensional array of X-ray density values, consisting of 240 adjacent approx. 1-mm thick cross sections of the structure under study. Since it is envisioned that the DSR will be used to investigate a variety of biomedical research and clinical diagnostic protocols, it has been provided with a variety of scanning modes allowing imaging of rapidly moving structures over a wide range of spatial and temporal resolution.

Proposed scanning protocols for diagnosis of, for example, the presence and extent of coronary artery disease (Ritman *et al.*, 1978) require the recording of

approx. 4 sec (i.e., 4 complete cardiac cycles) of projection data from each subject, creating a computational load of up to 58 000 127 × 127-element cross sections from each study. Even current state-of-the-art reconstruction processors operating at rates of approx. 1 sec/cross section would require 4 h of computation for each 1 sec of DSR operation, an impractically long duration for routine cost-effective clinical diagnostic usage for detection and study of cardiovascular abnormalities and early lung cancer for which the device is ideally suited (Wood, 1976; Ruegsegger *et al.*, 1977). Studies of the computational demands of the DSR indicate that no presently available general-purpose computers operating unaided, or augmented with commercial array processors, will be capable of fulfilling its data processing requirements at reasonable cost, since reconstruction times for each 127 × 127-element cross section must be reduced to durations not exceeding 1–10 msec/image (Gilbert *et al.*, 1978; Gilbert *et al.*, 1977; Swartzlander *et al.*, 1978; Gilbert *et al.*, 1979) if the DSR is to achieve its maximum biomedical research and clinical diagnostic potential. Thus, it will be necessary to develop special-purpose arithmetic processors specifically optimized to execute appropriate reconstruction algorithms.

CHOICE OF A SUITABLE RECONSTRUCTION ALGORITHM

The development of a high capacity processor to operate in support of the DSR required an examination of candidate reconstruction algorithms potentially suitable for execution by special-purpose processors. The by now well-known 'filtered back projection' method is considered a leading candidate in this regard. The first stage of this algorithm requires the one-dimensional linear filtration of each X-ray projection with a single, predetermined filter function. The second portion of the algorithm, the back projection of the filtered projections into the image space, actually generates the halftone values of the reconstructed cross-sectional image elements (pixels). The halftone value of each pixel is determined by selecting 1 sample, or 'ray sum', from each of the filtered projections according to a prespecified algorithm, followed by sequential addition of all selected ray sums to create the halftone value of that pixel (Herman *et al.*, 1976). The candidacy of the filtered back projection method is based on the generally high quality of the images produced by this algorithm, upon its speed in comparison to other (e.g., iterative) methods, and upon the identification of significant levels of computational parallelism within the algorithm (Gilbert *et al.*, 1979).

DEVELOPMENT OF AN ARCHITECTURAL DESIGN FOR THE PROCESSOR

Based upon an analysis of the numerical characteristics of filtered back projection,

Generation of Projection Addresses (Ray Sum Indices)
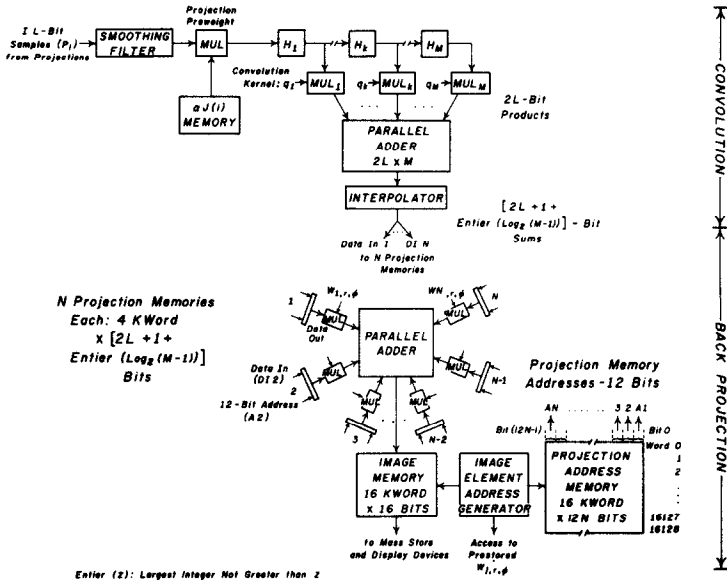& Back Projection Weights via Table Lookup

Fig. 2. Schematic of a proposed high-speed digital hardware implementation of the fan beam convolution reconstruction algorithm, employing direct convolution for linear filtration of each X-ray projection, and a memory driven parallel processing back projection method. (Reproduced with permission from Gilbert *et al.*, 1977).

a tentative design has been established for a reconstruction processor capable of exploiting the computational parallelism inherent in the algorithm. Figure 2 is a schematic diagram of a hardware processor executing the linear filtration of the individual projections, and thereafter the back projection of the filtered projections, using a parallel processing approach in which each pixel is completely back projected before computation is initiated on the next image pixel. In this diagram it is assumed, in agreement with the source-detector configuration of the DSR, that 28 projections will be back projected simultaneously. A detailed description of this design may be found in (Gilbert *et al.*, 1979).

ARCHITECTURAL DESIGN OF LINEAR FILTRATION PROCESSOR

Linear filtration of the individual projections may be carried out either by direct one-dimensional convolution in the spatial domain, or by means of the Fast

Fourier Transform (FFT) or finite field transforms (Reed *et al.*, 1977) and direct vector multiplication of the data vector and filter function vector in the spatial frequency domain (Brooks and DiChiro, 1976). From the view of hardware implementation, the convolution and frequency domain filtration techniques both have strengths and deficiencies requiring detailed examination (Swartzlander *et al.*, 1978) before the best of the available methods can be chosen. Direct convolution of real fixed precision vector operands containing 512 elements or less can be carried out entirely within the domain of real integer numbers, obviating the requirement for costly floating point hardware configurations. The increasing availability of high-speed LSI fixed precision arithmetic components of moderate cost, the conceptual simplicity of direct convolution and straightforward design of pipelined convolution processor architectures exhibiting high throughput, the freedom from large design commitments to control logic, and finally, the ease of layout, fabrication, and testing, all combine to make the direct convolution approach attractive for this particular application (Swartzlander *et al.*, 1978).

The upper portion of Fig. 2 demonstrates schematically the operation of a direct convolution filter. The digitized projections, each composed of J distinct samples corrected for black and white level and converted into density values, enter the reconstruction device at the upper left, where they are supplied to the processor section which carries out the filtration. The blocks designated H_i represent a series of digital registers, each connected to 1 input of a dedicated high-speed multiplication unit. Data samples are entered consecutively into register H_1, and are shifted into each successive register H_2, H_3, and so on. For the duration that a given projection data sample is in register H_i, it is presented to 1 of 2 sets of inputs on the associated multiplication unit MUL_i; the other input to the multiplier is supplied by a register which has been preset to the digital equivalent of 1 sample, a_i, of the convolution kernel. After a fixed time T_c has elapsed, each multiplier MUL_i has performed the operation $K_i P_j$; all M multipliers within a single convolution unit, operating in parallel, have formed M of these products. If the M products are summed during the next time interval T_c (by means of a parallel adder or addition tree (Swartzlander, 1973; Stenzel *et al.*, 1977)), 1 sample of the *convolved* profile will have been formed. As each projection is advanced through the series of registers, the entire convolution will be carried out within $(J+M-1)$ clock durations.

Several authors have pointed out the necessity for interpolation between samples of filtered projections recorded with a fan-shaped beam of X-rays, since an actual measured sample of the filtered data is not always available corresponding to a ray passing through the image pixel undergoing back projection (Herman *et al.*, 1976). Most reconstruction algorithms carry out these required interpolations upon demand during the back projection procedure. As an alternate approach, the interpolations could be executed unconditionally following the filtration of the projections. In Fig. 2, as the filtered data emerges from the direct convolution unit,

additional ray sums are incorporated into each projection by a hardware unit executing multiple stages of linear binary interpolations on the filtered ray sums taken pairwise. At the first stage of interpolation, a single value is computed midway between each pair of measured samples; in the second stage, 2 additional midpoints are calculated between the original pair of ray sums and the midpoint computed in the first stage. This procedure may be repeated, with 1, 3, 7, 15, or more interpolated ray sums interposed between each pair of measured ray sums when the interpolation has been completed. Interpolations may thus be executed very efficiently with a pipelined binary adder tree only, obviating the requirement for multiplication or division operations. Thereafter, only nearest-neighbor indexing of actual or 'pre-interpolated' ray sums is employed. Interpolation by this technique can be 'buried' in the transmission of the filtered projections into the high-speed memory employed in the back projection stage of the algorithm, during which the final reconstructed image is created (Gilbert et al., 1979).

OPTIMIZATION OF PROCESSOR DESIGN BY COMPUTER STUDIES OF ALGORITHM PER-
FORMANCE

Prior knowledge of the bandwidth and signal-to-noise ratio of the data vectors to be filtered may be exploited to minimize the number of elements and/or the dynamic range of the convolution kernel elements, equivalent to a decrease in the number of multiplications required or in the precision of intermediate calculations (and thereby in the amount of hardware required for filter implementation). Figures 3 and 4 demonstrate preliminary results of a series of cross section reconstructions using digitized X-ray projections of a 15 kg mongrel dog placed in an operational prototype of the DSR, using source-object-detector geometries, X-ray flux densities, and imaging system signal-to-noise ratios similar to that of the full scale DSR. A single set of projection data was repeatedly reconstructed to 127×127 image resolution using a succession of shorter, lower precision filter kernels (Fig. 3). The 'reference' image in the upper left corner was reconstructed in 48-bit floating point arithmetic with a kernel length equal to that of the projections. Little apparent degradation of image quality was observed through the first 2 stages of precision reduction, i.e., down to a 127-element kernel, but thereafter a decrease in overall image contrast appeared as a 'fogging' of the image, concomitant with a loss of spatial resolution of small image substructures.

These effects are depicted more graphically in Fig. 4, in which reconstructions whose projections were filtered with successively shorter kernels appear above dot and bar graphs of the gray values of pixels along 1 row (Row 63) of each image. The vertical gray bars represent the magnitudes of the pixel halftone values taken from the reference image (the upper lefthand image of Fig. 3); the white dots represent the halftone values from the pixels in the same positions (i.e., Row 63)
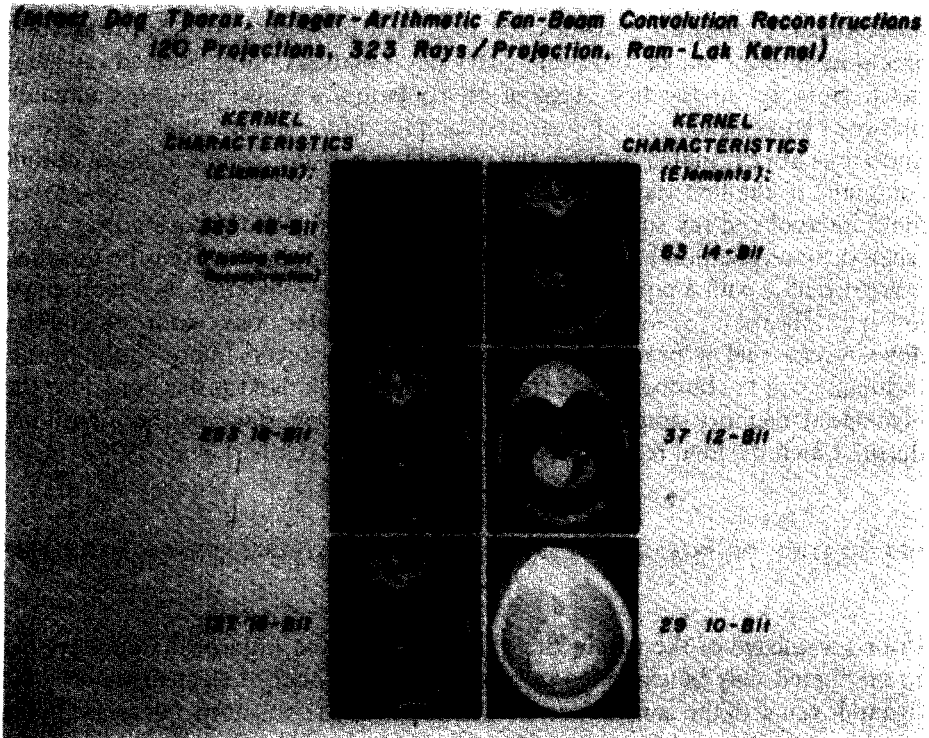
Fig. 3. Effects on reconstructed image quality resulting from successively decreasing both number and dynamic range of elements in convolution kernel. Upper left image, reconstructed in floating point arithmetic, is highest resolution image expected from the given projection data. Decreasing length of kernel and precision of kernel elements causes an average gray scale shift of entire image, or fogging effect, coupled with an actual loss of contrast and spatial resolution of small substructures.

of the image arrays reconstructed with the shortened kernels. Comparisons of the successive sets of gray bars and white dots demonstrate that the use of truncated kernels results in a halftone offset in the image and a decrease in spatial and contrast resolution which gradually erodes the spatial resolution of substructural features. A first order correction to the image, the subtraction of a fixed halftone offset from every image element, does not compensate for the progressive loss in image structural detail associated with the kernel truncation. On the basis of such simulation studies, it has been tentatively concluded that a filter kernel with 16–17 bit dynamic range is sufficient to reproduce the anatomic structures of interest which can be recovered from projection data recorded with the DSR imaging system. Continuing investigation of kernel characteristics is directed toward the optimization of kernel filtration capabilities while minimizing the number of kernel elements.
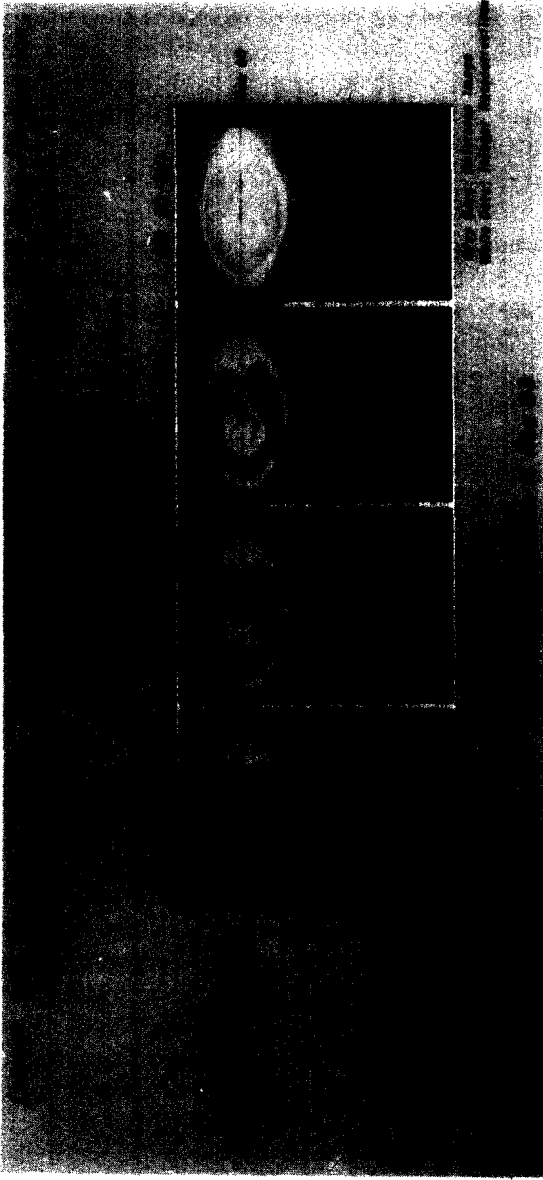
Fig. 4. Graphical illustration of the effect of decreasing element precision and number of elements in convolution kernel from 127 element kernel with 16-bit element precision, down to 29 elements with 10-bit precision. Gray bars in lower panels depict halftone values from row 63 of upper leftmost reference image in Fig. 3. White dots, representing halftone values of images reconstructed with successively more abbreviated kernels, show a progressive gray scale offset and loss of spatial resolution of small substructures within the reconstruction.

FABRICATION AND OPERATIONAL VERIFICATION OF A PROTOTYPE ADVANCED TECH-
NOLOGY HIGH-SPEED CONVOLUTION UNIT

Optimization of designs for special-purpose reconstruction processors are dependent upon the availability of high-speed components, particularly adders, multipliers, and memories whose high integration levels allow minimization of parts count and maintenance of very high throughput. Preliminary fabrication studies have been initiated of hardware processors exploiting such components which are capable of carrying out single 127×127 fan beam convolution reconstructions within 1–10 msec. Using monolithic 8×8-bit emitter coupled logic (ECL) two's complement multipliers of advanced design, several of which can be configured to form products of 8M-bit by 8N-bit operands (each component exhibits worst-case propagation delays of 25 nsec (Fletcher, 1976), a small two-section subunit of the direct convolution processor of Fig. 2 has been fabricated. Two of these multiplier components, labeled 'TRW', are visible in Fig. 5, as are the commercial ECL components comprising the shift registers (the H_i of Fig. 2) and adders which support the operation of the multipliers. Figures 6 and 7 display digital and waveform data, respectively, recorded from this prototype device, which during each clock cycle T_c forms the entire sum of products:

$$a_{i+1}x_{i+1} + a_ix_i = y_i, \qquad i = 1, \ldots, n \qquad (1)$$

with each successive clock cycle repeating the operation for a new (incremented) value of the subscript i. Figure 6 depicts the input data to, and output data from the convolver for 16 complete sums of products, that is, for $i = 1, 2 \cdots 16$. Each row of ones and zeros represents the input and output data corresponding to a complete evaluation of Eqn. 1 for a single value of i. All data bits were monitored simultaneously by a logic state analyzer, the CRT screen of which was photographed for presentation in Fig. 6. The 32-bit capacity of the logic state analyzer allowed only the 4 least significant bits (LSB) of the four 8-bit input operands, a_i, a_{i+1}, x_i, and x_{i+1}, and the 16 LSB of the resultant sum of products, y_i, to be displayed. The superscript numbers in Figs. 6 and 7 refer to specific data bit positions within the processor, which was operated initially at a system clock rate of 12 MHz, subsequently increased to 33 MHz. Two vertically positioned time axes appear at the right margin of Fig. 6, corresponding to the 12 MHz and 33 MHz system clock rates. Although the logic state analyzer verified correct operation of the convolution processor at the 12 MHz clock rate, operational constraints of the analyzer precluded use at the 33 MHz rate; correct operation of the convolution unit at the higher rate was therefore verified for each bit position with a wideband oscilloscope. The 33 MHz time scale in Fig. 6 is included for reference, indicating that each evaluation of Eqn. 1 for successive values of the subscript i was completed within approx. 30 nsec, i.e., 16 evaluations of Eqn. 1 were carried out within 480 nsec. Figure 7 depicts the actual signal waveforms displayed on an oscilloscope
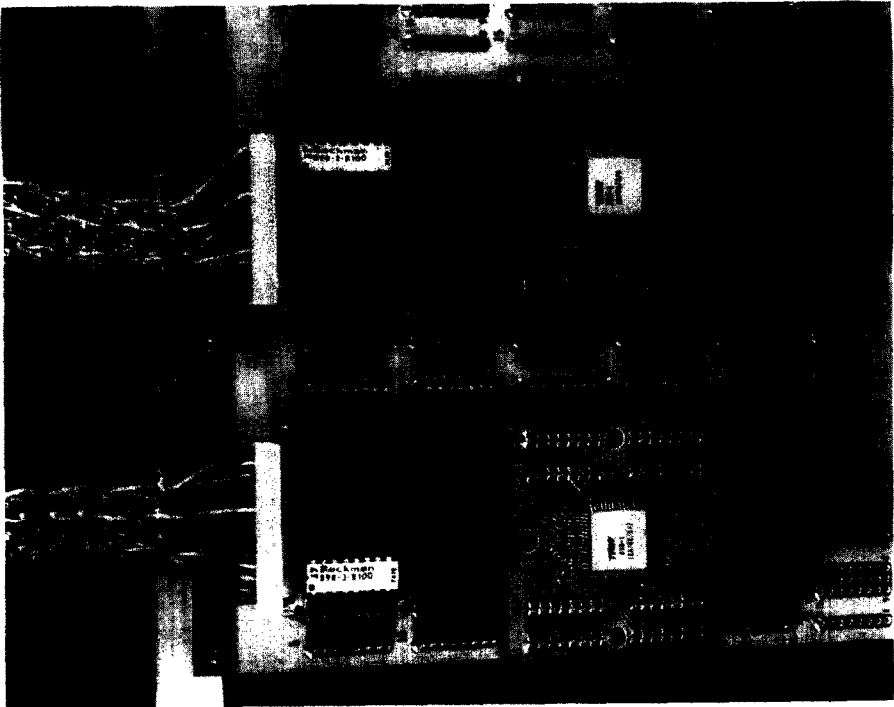
Fig. 5. Prototype 2 section high-speed direct convolution unit employing specially fabricated ECL 8 × 8-bit two's complement multipliers (marked 'TRW') in conjunction with commercially available adders and registers. (Reproduced with permission from Gilbert et al., 1977.)

from several circuit nodes within the functioning convolution unit, including the bit 2 and bit 15 outputs of 1 of the 2 multipliers, and the bit 15 output of the 16-bit adder. The 33 MHz system clock waveform is the lowermost trace. Examination of this otherwise classical digital data presentation, particularly the scale of the time axis, reveals the tremendous speed at which the unit functions. In this example, the transitions of the output bits 2 and 15 from the multiplier occur in 7 and 10 nsec, respectively, after the application of the input data operands; the outputs from the adder occur approx. 5 nsec thereafter (Fletcher, 1976).

The convolution unit design described above is expandable to 200 or more sections without significant difficulty and without sacrifices in operand resolution or processing speed. For example, assuming that a 256-element kernel is to be convolved with a 256-element data vector, an entire convolution executed by a single direct convolution unit implemented in ECL and operating at a system clock rate of 20 MHz could be completed within 26 μsec, at a computational rate of 5×10^9 multiplications and 5×10^9 additions/sec; 28 X-ray projections for a device such as the DSR could be filtered within 716 μsec.
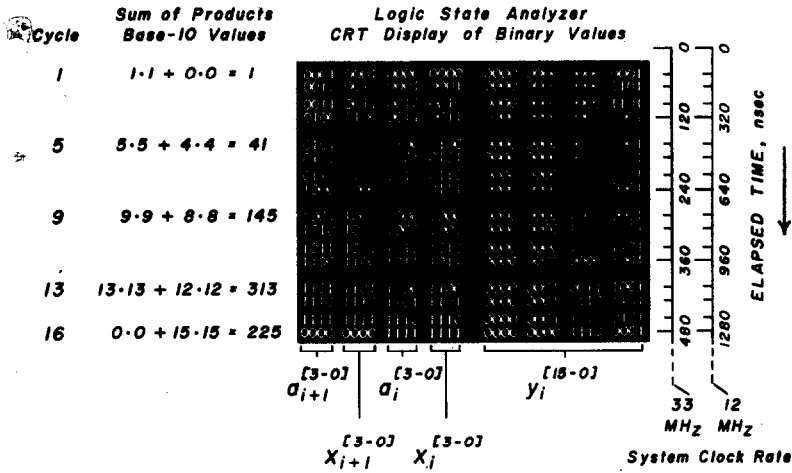
Fig. 6. Logic state analyzer display of digital data from prototype 2 section high-speed direct convolution processor of Fig. 5. Superscript numbers below the binary display refer to specific bit positions or groups of bits within the convolution processor. In this figure, processor evaluates the equation

$$a_{i+1}x_{i+1} + a_i x_i = y_i, \ i = 1, 2 \ldots 16.$$
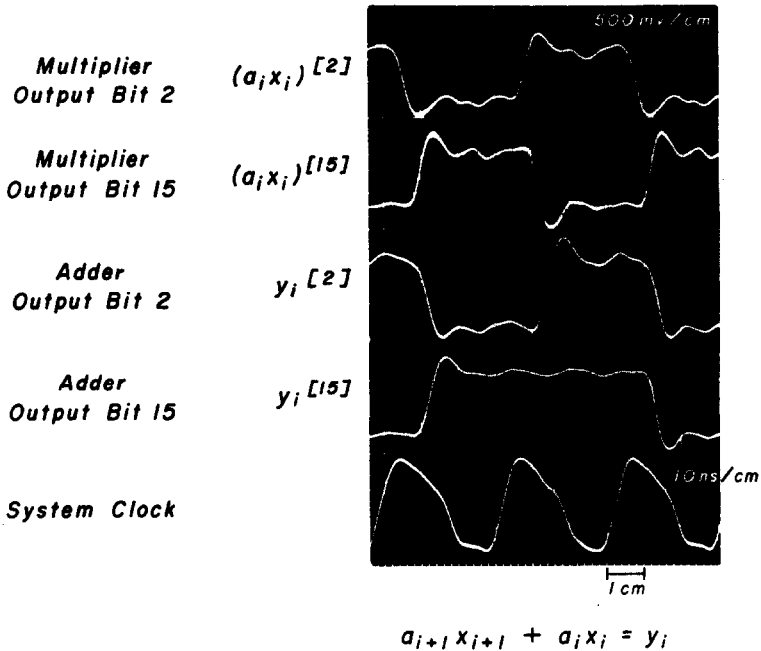

$$a_{i+1}x_{i+1} + a_i x_i = y_i$$

Fig. 7. Waveform data recorded with wideband oscilloscope from selected circuit nodes within the prototype ECL high-speed direct digital convolution unit of Fig. 5. Superscript numbers refer to specific bit positions or groups of bits within the convolution processor. System clock rate, lowermost trace, is 33 MHz.

ACKNOWLEDGEMENTS

REFERENCES

BROOKS, R. A. and DiCHIRO, G., *Phys. Med. Biol.*, **21** (1976) pp. 689–732.

FLETCHER, L. A., Proceedings of the 1976 Government Microcircuits Applications Conference (GOMAC), U.S. Army Electronics Command, Fort Monmouth, New Jersey, 1976, pp. 362–365.

GILBERT, B. K., BEISTAD, R. D., SWARTZLANDER, E. E., JR., KRUEGER, L. M., CHU, A., BREUER, D. R. and RITMAN, E. L., *Proc. of the San Diego Biomed. Symp.*, **16** (1977) pp. 159–168.

GILBERT, B. K., CHU, A., ATKINS, D. E., SWARTZLANDER, E. E., JR. and RITMAN, E. L., *Comput. Biomed. Res.*, **12**(1) (1979) pp. 17–38.

GILBERT, B. K., CHU, A. and RITMAN, E. L., in *Roentgen video techniques for dynamic studies of structure and function*, P. Heintzen and J. H. Bursch, (Eds.), Georg Thieme Verlag, Stuttgart, Germany, 1978, pp. 300–312.

HERMAN, G. T., LAKSHMINARAYANAN, A. V., NAPARSTEK, A., RITMAN, E. L., ROBB, R. A. and WOOD, E. H., in *Medical data processing*, M. LAUDET *et al.*, (Eds.), Taylor and Francis, Ltd., London, 1976, pp. 581–598.

REED, I. S., KWOH, Y. S., TRUONG, T. K. and HALL, E. L., *IEEE Trans. Nucl. Sci.*, **NS24**(1) (1977) pp. 843–849.

RITMAN, E. L., STURM, R. E., ROBB, R. A. and WOOD, E. H., in *Roentgen video techniques for dynamic studies of structure and function*, P. Heintzen and J. H. Bursch, (Eds.), Georg Thieme Verlag, Stuttgart, Germany, 1978, pp. 271–284.

RUEGSEGGER, P. E., RITMAN, E. L. and WOOD, E. H., *Proc. of the San Diego Biomed. Symp.*, **16** (1977) pp. 143–157.

STENZEL, W. J., KUBITZ, W. J. and GARCIA, G. H., *IEEE Trans. Comput.*, **C26**(10) (1977) pp. 948–957.

STURM, R. E., RITMAN, E. L., JOHNSON, S. A., WONDROW, M. A., ERDMAN, D. I. and WOOD, E. H., *Proc. of the San Diego Biomed. Symp.*, **15** (1976) pp. 181–188.

STURM, R. E., RITMAN, E. L. and WOOD, E. H., in *Roentgen video techniques for dynamic studies of structure and function*, P. Heintzen and J. H. Bursch, (Eds.), Georg Thieme Verlag, Stuttgart, Germany, 1978, pp. 261–270.

SWARTZLANDER, E. E., JR., IEEE Trans. Comput., **C22**(11) (1973) pp. 1021–1024.

SWARTZLANDER, E. E., JR., GILBERT, B. K. and REED, I. S., *IEEE Trans. Comput.*, **C27**(1) (1978) pp. 21–31.

WOOD, E. H., *Chest*, **69** (1976) pp. 394–408.