THE PROGRAMMABLE AND CONFIGURABLE LOW LEVEL VISION UNIT OF THE HERMIA MACHINE

Gaetano Gerardi[1], Giancarlo Parodi[2]

[1] Istituto di Fisica dell' Università, Via Archirafi 36 -9123 Palermo- Italy.
[2] Dip. di Ing. Biofisica ed Elettronica dell' Università, Via Opera Pia, 11a -16145 Genova -Italy

ABSTRACT

In this work the Low Level Vision Unit (LLVU) of the Heterogeneous and Reconfigurable Machine for Image Analysis (HERMIA) is described. The LLVU consists of the innovative integration of six A110 DSP (Digital Signal Processing). Two main features have been considered in the design: the generality and the extendibility; the LLVU can be fully programmable and the network interconnection of the IMS A110 can be easily reconfigured by the user. The operator can choose among different options to tailor its applications: single (8 bit) or double (16 bit) precision in data or in the MAC coefficients, producing 22 or 32 bit output, 1 or 2 (cascade or parallel) kernels with programmable width and height.

GENERAL DESCRIPTION

The HERMIA machine: HERMIA is the acronym of Heterogeneous and Reconfigurable Machine for Image Analysis. Its design has been oriented to fit, at the best, image analysis processes with the architecture. The computation paradigm of image analysis, conventionally, follows four phases: Low Level Vision, (LLV), Intermediate level Vision, (ILV), High Level Vision, (HLV) and Interpretative Analysis (IA). The full integration of all phases in a unique system is steel unsolved. In this machine the approach followed is based on the principle of specialization and co-operation between workers. The image analysis is carried out in a pipe fashion; each step is performed by a dedicated class (heterogeneity) of reconfigurable processors. Reconfigurability seems to be the key solution of vision problems. The system consists of four main units: 1) The PC (IBM compatible) performs high and interpretative analysis and is also responsible for the software management and control of the whole system. 2) A network of 16, T800, transputer realises the Intermediate Level Vision Unit (ILVU). The network, under software control, can be configured in a suitable way for the problem: pipe, mesh, ring, tree. 3) The Active Memory Unit (AMU) handles the I/O and the broadcasting of the image data. Moreover the image data can be accessed separately by all functional units by means of a run-time memory bank switch. 4) Six IMS A110 DSP realises the Low Level Vision Unit (LLVU). In this paper this unit will be described in deep. The user, interacting with the system via Menu and Icon of the MS WINDOWS 3.1 environment, has all the control of the machine. Configuration of all units can be easily performed. Complex algorithms can be build and applied to any image from the data base, stored in the system, by choosing a sequence of procedures, among a suitable image analysis library. A new high level Pictorial C language (PCIL), in which image data type and related parallel instructions are defined as Icon, makes, the user, able to realise new parallel algorithms.

THE LOW LEVEL VISION UNIT

Summary of performances: The Low-level unit should be capable of performing basic image pre-processing operation at video rate: Arithmetic and Boolean operation between images; image enhancement and restoration: point operation, local operations (convolution and deconvolution, filtering, etc.); template matching and edge detection (Robert, Sobel and other operators).

For the implementation of a such unit the Digital Signal Processor, (DSP), IMS A110 has been chosen as a building block. It consists of a configurable array of multiply-accumulators (macs). The array is arranged as 3 raw of seven macs, each raw preceded by 1120 stage programmable shift registers. A versatile postprocessing unit follows the array. A survey of the CIP is given in the appendix. In the LLVU of HERMIA we use a cascade of six IMS A110 to achieve more versatility. The user can, by software, configure their interconnection in several ways to fit its image analysis problem. All the features of the A110 are exploited. In the following we give same example of the possibilities:

- Long one dimensional up to 126 stage transversal filter
- Wide and high two dimensional filters: all filters with dimensions satisfying the rule 7X by 3Y, where X and Y are any integers satisfying $XY \leq 6$ can be produced, e.g., 7 by 3, 6, 9, 12, 15, 18 or 14 by 3, 6, 9 or 21 by 3, 6.
- Two pass filtering operations: two completely different filters can be defined (e.g., a mean operation followed by e Laplacian filter), each with dimensions satisfying the rule 7X by 3Y,

where X and Y are any integers satisfying XY ≤ 3, e.g., 7 by 3, 6, 9, or 14 by 3, or 21 by 3.

- Data input precision: can be 8 or 16 bits wide with 22 or 32 bit output respectively, in twos complement format.
- Coefficient precision: 8 or 16 bits wide with 22 or 32 bit output respectively, in twos complement format.
- Processing of two different images with different kernel (satisfying the rule 7X by 3Y; XY ≤ 3) in the same time.
- Arithmetic and Boolean operation between image.
- Image size can be up to 1024 by 1024 pixels. To avoid skew in the output image, programmable shift registers, internal to the CIP, are settled, by software, to produce necessary delays.

All the configurations are implemented without speed penalty (20 MHz).

Moreover, in the LLVU (fully programmed by the ILVU), image data can come from and be sent to different unit of HERMIA. Table 1 gives all the possibilities:

Data Input from	Data Out to
PC	PC
PC	Link
PC	External
Link	PC
Link	Link
Link	External
External	PC
External	Link
External	External

Table 1

Any combination of source and destination unit can be programmed. The speed of the system depends on the option chosen. The slowest mode is obtained by using the ILVU either for send and receive image data. However, in this mode, full debug of the unit can be obtained: all the unit control registers and the output of each A110 are accessible. Medium speed is reached if we use transputer links to send and receive images' data because their intrinsic parallelism. Four transputer links are used: two of them is dedicated to the input and two to the output of the images' data. The full speed (20 MHz) is obtained if we use the external parallel bus.

All the facility offered by the backend processor of the IMS A110 are used at any level it needs. We can perform, on the output data of each mac array, various operations: edge detection enhancement, transformations and data normalisation, compensation of changing conditions, feature recognition and histogram equalisation. In particular:

- in binary image processing we can do isolated pixel removal, line linking, encoding according to connectivity, binary thinning including staircase elimination, feature growth;
- in multilevel image processing we can also do: thresholding, image contouring, static and dynamic range compression.

HARDWARE DESCRIPTION

The low-level vision unit consists of two main sections:

- processing unit
- control unit

The processing unit is shown in Fig.1. This consists of the six A110 arranged in two banks, each with three A110. The multiplexers manage the data transfer between the processors; they are the core of the interconnection network . All the configurations offered by the interconnection network are grouped in the following:

- flexible choise of the dimension of the filtering window;
- multi-step filtering;
- increasingly of the data precision;
- increasingly of the coefficient precision of the filtering window.

This way allows all the processing possibilities above listed.

The control unit manages the flow of data according to the Table 1 and controls the programming of the processors. The fully programming of the LLVU is achieved by a I/O port of the host PC.

CONCLUDING REMARKS

At the present, the complete circuit simulation of the project has given the following results:

(for frame of 512 x 512 pixels and 8 bit/pixel)

- performance of the A110 banks: 76 frames/s
- performance of the LLVU
  I/O DATA
  TRANSPUTER LINKS: 1 frame/s
  PC BUS: 2 frames/s
  EXT FAST BUS: 24 frames/s

The future implementation of the LLVU will be made with faster support components(state machines, multiplexers and buffers) embodied in high speed programmable gate arrays. In this case the processing performance could be about : 60 frames/s.

The result is very important for real-time high resolution images processing; in fact, for images of 1024 x 1024 pixels we have a performance expectation of 12 - 14 frames/s.

True colors applications (24 bit/pixel, R G B) will be possible with three LLVU boards.

For software environment, the future research lines will address three main elements: 1) software tools for debugging, 2) processing primitives for A110, 3) integration of the system in the HERMIA environment.

ACKNOWLEDGEMENTS

This work has been partially supported by National Council of Research under contract n. 91.00895.69 and n.92.01608.PF69 and C.R.R.N.S.M. (Regione Sicilia)..

REFERENCES

[1] G. Gerardi, F. Chiavetta, V. Di Gesù, D. Tegolo - Hermia: An Heteroneneus And Reconfigurable Machine For Image Analysis. MVA IAPR Workshop on Machine Vision Applications, Tokyo, pp.397-400, Nov., 28-30,1990

[2] INMOS The Digital Signal Processing Databook, July 1989.

[3] R.C. Gonzalez, P.Wintz - Digital Image Processing, Addison Wesley.

[4] R. Whitton - Cascading IMS A110s, INMOS.

[5] R. Stefanelli, A. Rosenfeld - Some Parallel Thinning Algorithms For Digital Pictures. Comm ACM 18,2.

[6] H.E.Lu, P.S.P.Wang - A Comment On Fast Parallel Algorithms For Thinning Digital Patterns. Comm ACM 29,3.

[7] C.M. Holt, A.Stewart, M.Clint, R.H. Perrott - An Improved Parallel Thinning Algorithm. Comm ACM 30,2.

[8] R.W. Hall - Fast Parallel Thinning Algorithms: Parallel Speed And Connectivity Preservation. Comm ACM 32,1.

[9] Z. Guo, R.W. Hall - Parallel Thinning With Two Subiteration Algorithms. Comm ACM 32,3.

APPENDIX

Inside the IMS A100 DSP: In this appendix a compressed description of the IMS A110 CIP is given. The IMS A110 consists of a configurable array of multiply-accumulators (macs). The mac array consists of three 7-stage transversal filters which can be configured either as a 21-stage linear pipeline or as 3 by 7 two-dimensional windows. The input data is 8 bit wide and it is fed to the mac array by three programmable (1120 stage) shift registers. The output of each shift register is supplied as input to one of the three 7-stage transversal filters. At each stage the input sample is multiplied by a coefficient stored in memory, and added to the output of the previous stage, delayed by one clock cycle. The output of each 7-stage mac is fed, via a delay stage to the first stage in the next transversal filter. The coefficient word width in the mac array is 8 bits wide. Two banks of coefficient are provided. At any instant one set of coefficient is in use with the mac array. The other set can be altered via the microprocessor interface. Once a new set of coefficients has been loaded, the activities of the two coefficient banks can be interchanged without interrupting the flow of data. Both data input and coefficient can be programmed independently to support twos complement or positive unsigned formats, allowing multiple devices to be used as a "slice" in higher accuracy systems. Within the mac no truncation or rounding is performed on the partial products. The mac array output is fed to backend post-processing unit. Its input stage is a programmable shifter capable of arithmetic right shift (divides) of up to 8 bits with rounding, and left shifts of up 8 bits. The output of the shifter passes into the cascade adder where it is added, along with any rounding generated by the shifter, to either the cascade input bus or a zero value. A positive or negative overflow is generated if the result of adding is out of a 22 bit value. The out of the adder can be programmed to be full wave or half wave rectified. Data are monitored by a statistics unit that can be programmed to provide: capture of the maximum or minimum value of data; overshoot or undershoot, respect to a programmed value, on the data stream counting. A LUT, consisting of 66 words by 32 bits wide, also usable as 255 bytes, permits the user to perform static or dynamic scaling, simple transformation or dynamic normalisation.

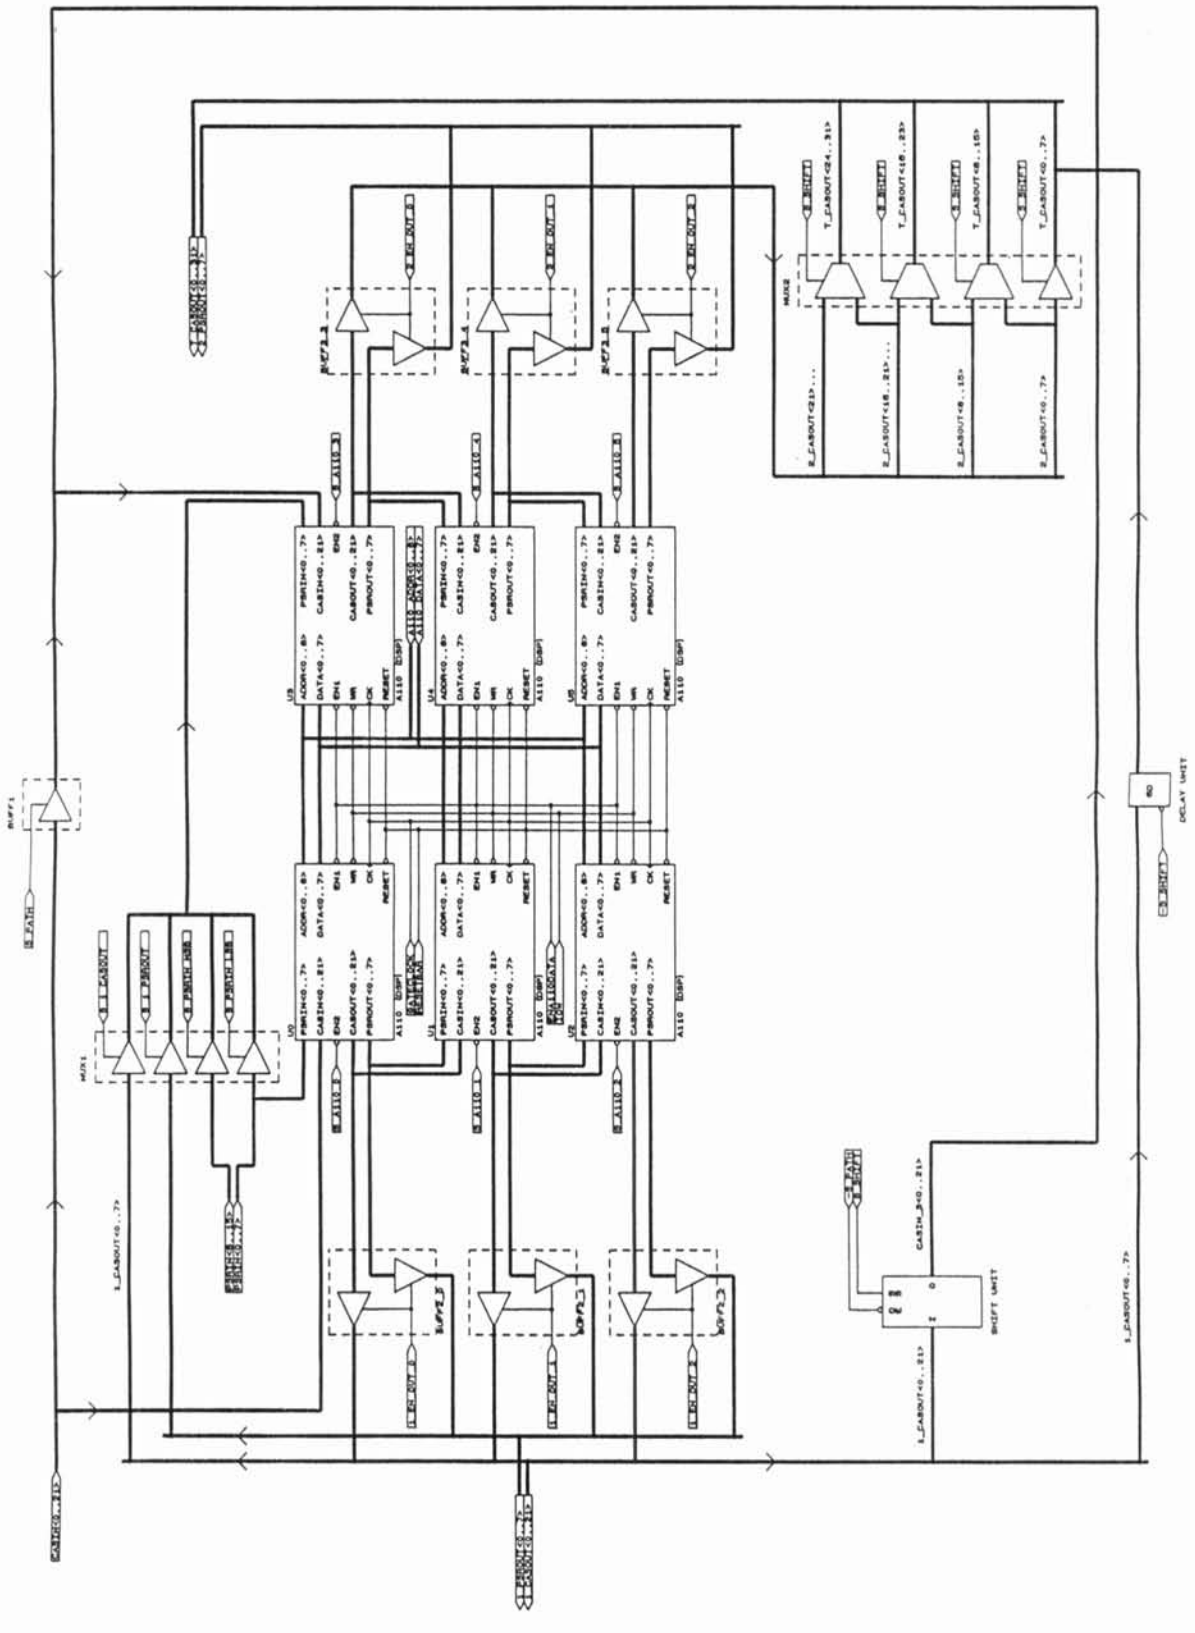

Fig.1 - Processing unit