

Zeitschrift: Comtec : Informations- und Telekommunikationstechnologie = information and telecommunication technology

Herausgeber: Swisscom

Band: 77 (1999)

Heft: 10

Artikel: The image Chip

Autor: Gumm, Martin / Garina, Pierangelo

DOI: <https://doi.org/10.5169/seals-877063>

Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften auf E-Periodica. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. Das Veröffentlichen von Bildern in Print- und Online-Publikationen sowie auf Social Media-Kanälen oder Webseiten ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. [Mehr erfahren](#)

Conditions d'utilisation

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. La reproduction d'images dans des publications imprimées ou en ligne ainsi que sur des canaux de médias sociaux ou des sites web n'est autorisée qu'avec l'accord préalable des détenteurs des droits. [En savoir plus](#)

Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. Publishing images in print and online publications, as well as on social media channels or websites, is only permitted with the prior consent of the rights holders. [Find out more](#)

Download PDF: 14.08.2025

ETH-Bibliothek Zürich, E-Periodica, <https://www.e-periodica.ch>

AC 078 Atlantic:

The Image Chip

A new, versatile motion estimation chip is described which has been developed with under special consideration of concurrently developed video encoder hardware for broadcast studio environments using a new approach for minimizing cascading impairments. The chip implements an unconventional, high-performance motion estimation algorithm, which is based on tracking motion through several frames and on a genetic search strategy. The component is capable of performing a variety of different motion-estimation related tasks and offers a variety of optional enhancements to optimize and maximize the final coding quality. This functionality has been mapped into a MIMD based parallel architecture. All algorithms are executed on embedded RISC controllers, which provides a high flexibility. Several strategies for multi-chip configurations allow increase the coding quality according to system needs. The chip has been integrated in 0,35 μm CMOS technology on an area of about 60 mm².

Digital Television based on video compression according to the MPEG-2 [1] standard is about to enter broadcast studio environments. The compression of source material, editing and switching, and final assembly of

co-operation between the Integrated System Center at the EPFL in Lausanne, Switzerland, and the Centro Studi E Laboratori Telecomunicazioni (CSELT) in Torino, Italy.

The Cougar and Atlantic Approach

The architecture described in this article is the outcome of the European RACE Cougar and Acts Atlantic Projects that aim to introduce MPEG-2 digital moving picture format (CCIR601 720x576@25 Hz/720x480@30 Hz; 4:2:0 video) into

television broadcasting environments. The main goals of Cougar were three-fold: the study through computer simulation of pre-processing, motion estimation and coding techniques to optimize the quality of MPEG-2 coding within the constraints imposed by the hardware implementation [2], the design of flexible hardware to perform MPEG-2 coding for main and SNR profile at main level [3], and the identification of VLSI to implement defined key functions for the later replacement of discrete devices in prototype hardware. The Atlantic Project, [4] [5] aims to extend this work by demonstrating the use of MPEG-2 compression throughout the complete broadcast chain. In a typical studio environment, several cascaded processing steps like switching between different bit-streams, changing the bit-rate of bit-streams (transcoding), or the other editing functions (e.g. fading), commonly are executed between the compression of the source signals and the assembly of the final program stream. These manipulations require the decoding of incoming MPEG-2 signals and the re-encoding of the processed signal after the processing

MPEG-2 bit streams are required to deliver a very high final coding quality in real time if this technology is to replace today's conventional systems. Within MPEG-2 video compression, the reduction of the temporal redundancy in video sequence by motion estimation is the most computationally heavy task requiring processing power ranging from a few to dozens of giga operations per second (GOPS). Accurate and realistic motion estimation is crucial to both the overall quality of the compressed video and the achievable compression rate. This paper describes a new, flexible high performance motion estimation architecture, the so-called integrated MIMD architecture for genetic motion estimation (IMAGE) chip that was designed to perform a range of motion-estimation related tasks in existing and future high-quality video equipment for studio use. This work presents the key activity of the VLSI design group within the Acts Atlantic Project and is the outcome of the

1. Initialize the first population with the best motion vector of the previous macroblock in the same slice. Complete the population of 20 motion vectors by adding small, Laplacian distributed random values to the latter.
2. Calculate the MAE values of all motion vectors in the first population, memorize a sorted list of the nine best vectors together with their MAE value.
3. Form a new population of 20 motion vectors by applying the crossover operation to the memorized list of best vectors and adding small, Laplacian distributed random values to them.
4. Calculate for each of the 20 new motion vectors its MAE value and update the set of 9 best motion vectors if required.
5. Repeat steps 3 to 4 for the third and forth generation.
6. After the evaluation of the fourth generation, the first element of the memorized 9 best vectors represents the final best motion vector.

Fig. 1. The Modified Genetic Search Algorithm (MGS).




Fig. 2. (from [11]): Principle of vector tracing.

during which each decode-encode step introduces new coding impairments and degrades the final quality of the video signal. Therefore, technology for a digital studio must provide excellent encoding quality for the compression of the source material (even for difficult scenes like light changes or fast motion) and minimize the build-up of cascaded compression impairments.

Implementation of the encoding optimization proposed in Cougar help to provide such coding quality. To satisfy the second requirement, a special concatenation scheme [6] has been developed in Atlantic. The key idea is to preserve coding side information (i.e. motion vectors, coding decisions, etc.), transfer it to the re-encoder, and reuse it in its decision processing. This information is carried on the so-called information bus (Infobus), apart from the pure pixel data transmitted over the picture bus (Picbus). Examples for the use of such a structure for transcoding are given in [7], and for switching and editing in [8]. The infobus concept is therefore employed in the encoder architectures of the envisaged Atlantic hardware and the Cougar hardware.

VLSI Requirements for Envisaged Applications

Motion estimation is the most expensive computational part in the video encoding process. The commonly applied strategy is the matching of a block of the current picture in a defined search window of a previous (or future) picture to find the minimal prediction error. For

digital TV formats like the main level profile of MPEG-2, the requirements for both processing power and IO bandwidth are extremely high if excellent encoding quality is required as in studio environments. Therefore one main focus for VLSI activity within Atlantic was to provide a solution for a front-end high quality motion estimator that satisfies these constraints and implements the coding optimizations proposed in Cougar. This VLSI aims to replace complex, FPGA-based prototype hardware. Several observations concerning the motion estimation task had been made by simulations in Cougar. The essential requirements are the following:

- *Tracking of fast motion:* To correctly deal with sequences containing fast motion, the motion estimator must provide a search range large enough to track these motions.
- *True motion:* the motion estimator should provide motion vectors that not only minimize the prediction error but represent as much as possible true scene motion.
- *Dealing with all defined MPEG prediction modes:* Significant gains in coding quality can be achieved when cost for all prediction modes defined in MPEG-2 are calculated by the motion estimator.
- *Half-pixel accuracy:* the coding quality gain by half-pixel interpolation is well known, so this feature is a must for every motion estimator. There were further desired features for a versatile motion estimator.

- *Dual prime cost calculation:* although low-delay coding is not a primary concern in studio environments, the ability to handle the dual prime prediction mode of MPEG-2 also increases the versatility of the VLSI.

- *Search window offset:* for a scene with uniform motion (e.g. camera panning) a sophisticated encoder could analyze the global motion in a frame and offset the search window of the motion estimator at picture rate in order to bias the search in this direction.

As a first MPEG-2 encoder using Cougar technology, it was decided to use a commercially available motion estimator that has been used previously for the conversion of television standards. The so-called phase correlator produces a smooth field of «true» motion vectors that have to be converted by a processing unit into MPEG-2 compliant vectors [9]. Simulations showed that coding quality is enhanced significantly when these converted MPEG-2 motion vectors are refined to half-pixel accuracy in small zones of the reference picture (normally around the motion vector) of typically $\pm 2/\pm 1$ pixels. This refinement is a computationally heavy task as each estimated macroblock of up to 12 motion vectors must be refined in this way. The basic operation of this refinement is a full search-motion estimation in all integer positions of the refinement zone with succeeding half-pixel refinement around the found best vector. Therefore, it has been decided to include this functionality into the motion estimation VLSI to be developed.

The approach chosen for both Cougar and Atlantic encoder hardware is that the final encoder block receives all necessary coding information co-timed with the pixel data via the Infobus (this block is therefore called the DIM-coder, since most coding decisions are taken elsewhere in the system). The Infobus comprises (among others) the final motion vectors and the decision and the chosen prediction mode for the current macroblock. Choices must be made between intra-coded (not predicted), forward, backward, or bi-directional prediction. Furthermore, in field prediction, choices must be made between 16x16 and 16x8 prediction modes. This decision process has been implemented in the Cougar prototype hardware on a separate DSP board. It was desirable to incorporate this functionality into the motion estimation VLSI also.

Vector Tracing Modified Genetic Search for Motion Estimation

This section describes briefly a new motion estimation search algorithm that has been developed in Cougar ([10], [11]).

The goal was to find a strategy that satisfies on one hand the high quality requirements as described in the previous section and on the other hand has a limited complexity to allow for its hardware implementation. The final developed algorithm is a strategy which combines vector tracing techniques with a pseudo-genetic search approach.

The straight-forward approach to find the best match of a block of $n \times n$ pixels of the current picture in a search window of $M \times N$ pixels of a previous (or future) frame is an exhaustive search, i.e. calculating the MAE for all possible search locations. This method guarantees finding the minimal displacement error and hence minimize the overall PSNR. For MPEG-2 ML@MP coding, a processing power of 90 GOPS is required for a ± 32 H/ ± 32 V pixels search range that is adequate for TV studio applications. A variety of fast search algorithms have been proposed in the past to reduce the computational burden of the exhaustive search. Three strategies are primarily em-

ployed: regular reduction of the number of search positions, hierarchical pixel subsampling-based algorithms, and gradient-based algorithms. Recent reported motion estimation implementations [12] [13] [14] use a combination of the first two strategies which have the advantage of being very regular, thus easing a massive-parallel hardware implementation. On the other hand, most of these approaches suffer by tending to get stuck in a local minimum in the search window, so that several improvements like multiple-survivor strategies or optional successive up-and down-sampling have been employed to improve performance. A new search approach, the vector-tracing modified genetic search algorithm (VT-MGS), has been implemented in the image chip. The base algorithm is derived from a simple, pseudo-genetic approach, described in figure 1. Hence, terms are used, such as "population" (the set of motion vectors actually considered) and "sons" (new motion vectors generated from previous ones using a "crossover" operation). The MGS algorithm altogether uses only 80 MAE calculations to determine the motion vector for one macro block. Initialize the first population with the best motion vector

of the previous macroblock in the same slice. Complete the population of 20 motion vectors by adding small, Laplacian-distributed random values to the latter.

Calculate the MAE values of all motion vectors in the first population, memorize a sorted list of the nine best vectors together with their MAE value. Form a new population of 20 motion vectors by applying the crossover operation to the memorized list of best vectors and adding small, Laplacian distributed random values to them. Calculate for each of the 20 new motion vectors its MAE value and update the set of 9 best motion vectors if required. Repeat steps 3 to 4 for the third and fourth generation. After the evaluation of the fourth generation, the first element of the memorized 9 best vectors represents the final best motion vector.

In many cases, the motion of objects in a scene is essentially uniform for a time much longer than the temporal distance between two frames. In other words, the motion vector fields of consecutive frames are highly correlated. This characteristic can be exploited to significantly improve the results of algorithms, like MGS, by applying a technique called vec-




Fig. 3. Image chip architecture.

tor tracing that is depicted in figure 2. For the initialization vectors to estimate a macroblock in P2 from P1, the best motion vector (only the frame vector) of the same macroblock position in P1 and its eight surrounding neighbors are taken. In a similar way, the search is initialized to estimate a macroblock in B1 or B2 from P1 or P2 respectively. In combining the vector tracing technique with the MGS algorithm, both spatial and temporal correlation of the motion vector fields are exploited.

This algorithm is called VT-MGS and delivers nearly full-search PSNR on large search windows (± 75 pixels horizontally and vertically in the simulation) with only 80 MAE calculations.

Algorithm Hardware Implications

So far, most of the reported implementations of motion estimators for MPEG1/2 or H.26x video codecs take advantage of the inherent massive parallelism of block matching and the regularity of the employed search strategies by mapping the algorithm into systolic array structures. First TV-suitable implementations could cover a full search range of ± 31 H/ ± 32 V pixel (4 chips required). [15] Newer reported motion estimators made for TV

resolution MPEG-2 encoding on main level use modified systolic array approaches that allow for the execution of combined sub-sampling and N-step-search strategies to enlarge the search range from ± 31 H/ ± 31 V ([12], [13]), ± 48 H/ ± 15 V [16], up to ± 64 H/ ± 32 V. ([14]) The most recent chips on the market ([17] [18]) use powerful full search machines with an optional large search window center offset.

A VT-MGS motion estimator cannot be based on the systolic array approach used for most other reported circuits. The complete search window must be accessible at any time during the processing, i.e. it must be stored in a fast cache memory on chip. To exploit the capabilities of the algorithm to a maximum, the search window should be as large as possible and limited by two main factors, the cache size and the memory bandwidth required to fill it, when this image has been established at 96–64 pixels (± 40 H/ ± 24 V search range).

The core of the VT-MGS algorithm is divided into two main processing steps: the generation of the next candidate vector to be tested and the evaluation of the MAE value of this vector. For the for-

mer, it is advantageous to use a programmable candidate vector generator in order to allow for the adaptation and optimization of the algorithm to the application needs. For the latter, a fast pixel processing unit is needed that is capable of deriving the search window cache address of the corresponding macroblock from the candidate motion vector, accessing the data, and calculating the corresponding MAF value. A trade-off must be found between the MAE calculation latency and the number of pixels of a macroblock that can be treated in parallel in one cycle (defined by the cache output size and pixel processor datapath width). Furthermore, its processing power must be adapted to the capability of the candidate vector generator to provide new candidates.

In contrast to systolic array approaches where parallelism is exploited on pixel level, a VT-MGS motion estimator can only be paralleled at the macroblock level because of the non-predictable search locations that are evaluated during the processing. However, paralleling can be done by implementing several processing modules onto one chip to enable the architecture to treat several macroblocks in parallel and thus to sat-

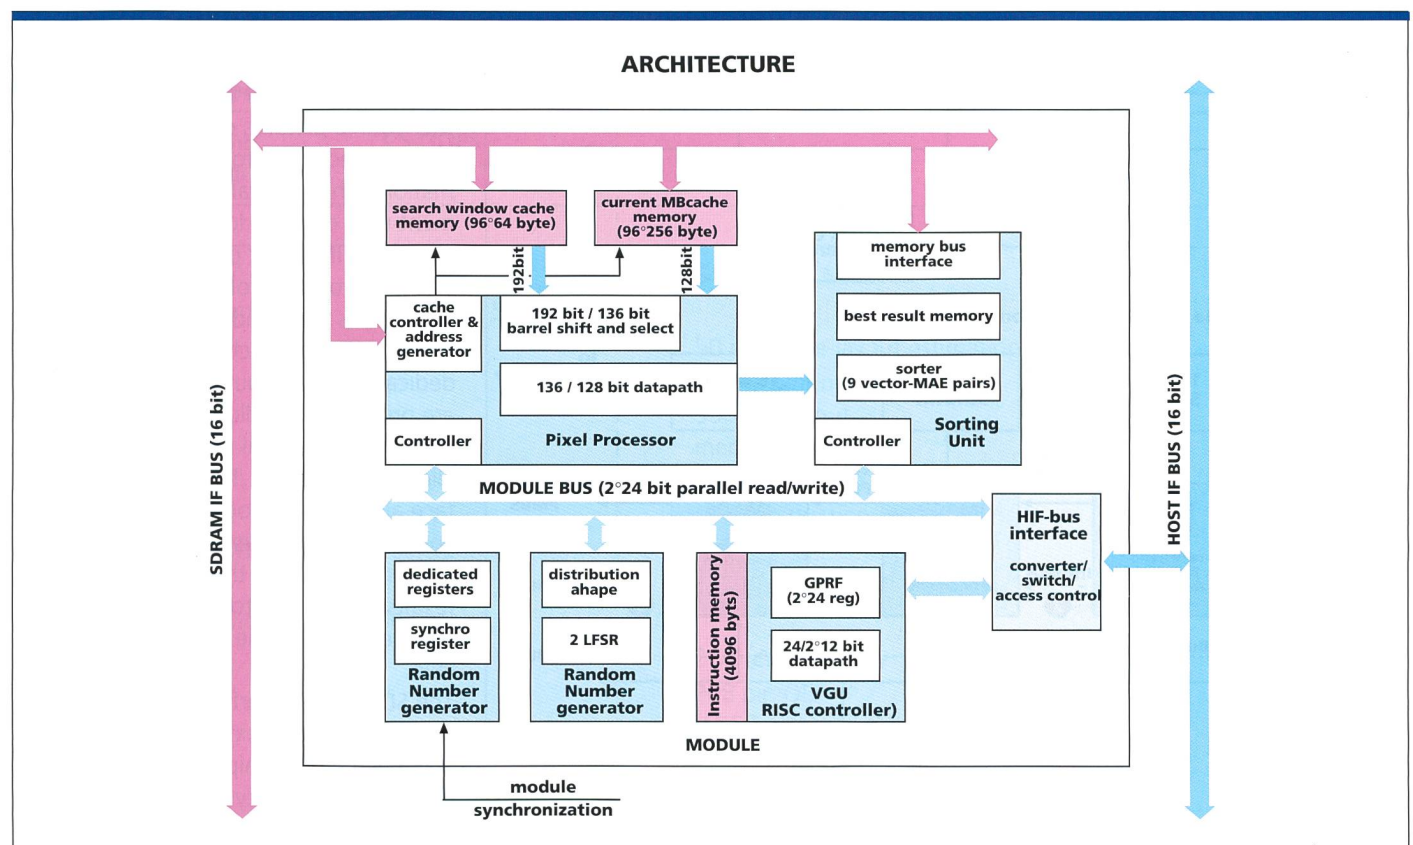


Fig. 4. Architecture of one module.

isfy given real time constraints. In turn, complexity of the overall chip-data distribution scheme is increased. One processing module will always perform either forward or backward prediction only.

Image Architecture Description

Global Architecture

The top view of the IMAGE architecture with its main building blocks is depicted in figure 3. It consists of the two processing modules (MODULE1 and MODULE2), the global chip controller (GLOC), the SDRAM memory interface (SDRAM_IF), the Infobus manager (IM), the video input (VID_IN1, VID_IN2) and output (VID_OUT) buffers, the host interface (HOST_IF), and the combined JTAG/general test interface (TEST_IF, described in section IX: Test Approach). Two main

buses, the memory bus and the host interface bus, connect these blocks. The memory bus is controlled by the GLOC and is used for all the transfers between the video input and output buffers, the SDRAM_IF and the memory input ports of the two modules. The arbitrated host interface bus is used by the two RISC controllers of the modules during processing to access the shared RAM data, the Infobus manager buffers and registers, and the global controller registers. Furthermore, the board controller CPU uses the bus during the chip setup phase. The two processing modules work in parallel and independently according to the MIMD principle, i.e. they execute their own program on private data. Although in general the two programs will be identical for the two modules, they are executed at different times.

Acknowledgements

The authors would like to express their acknowledgement to the members of the design team at CSLT and EPFI who contributed to this work: Mareo Gandini, Friederich Mombers, Andrea Finotello, Ilario Remi, Mauro Marchisio, Stephanie Dogimont, and Alessandro Torielli. This work was done within the framework of the ACTS Atlantic project AC078 (Advanced Television at Low Bitrates & Network Transmission over Integrated Communication Systems), and funded in part by the European Union. The authors would like to thank the members of the consortium for their indispensable input and the permission to publish this paper.

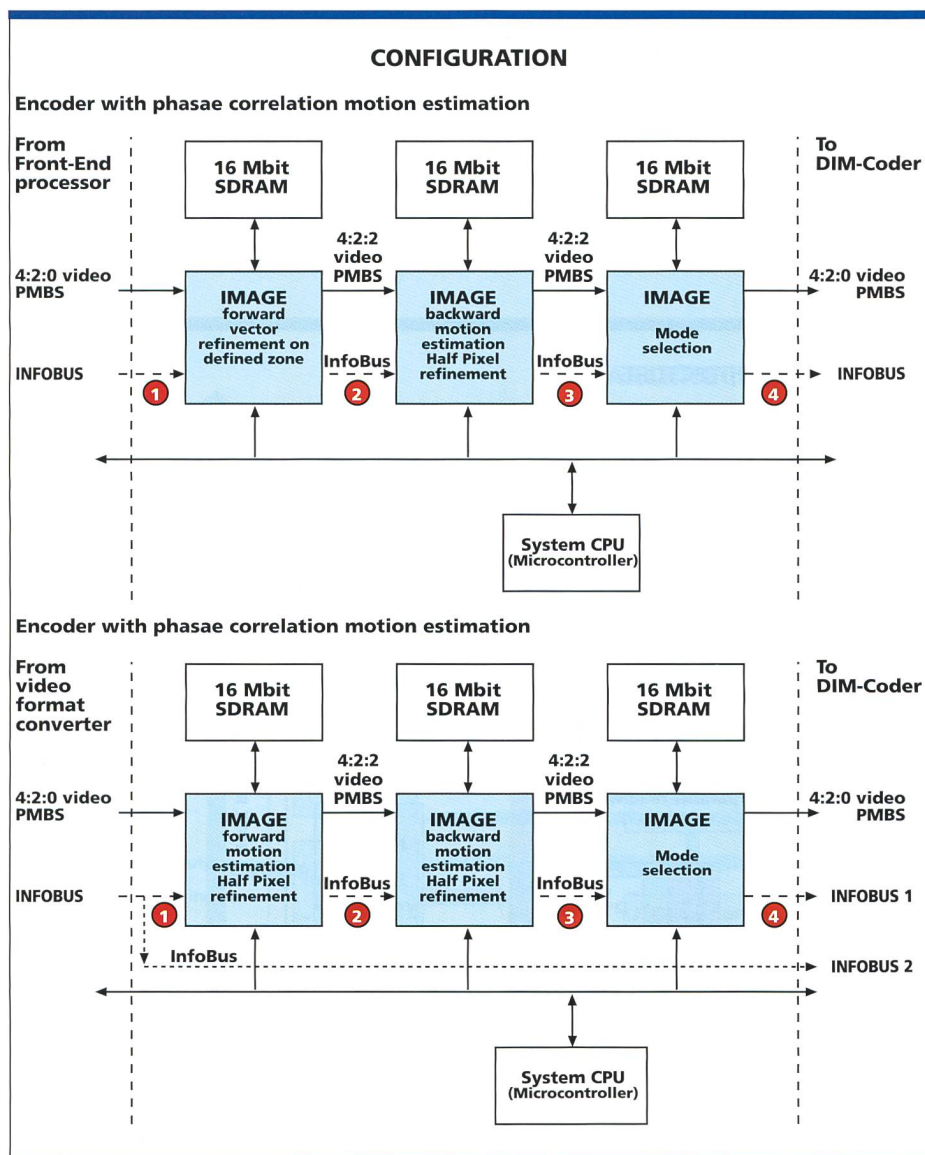


Fig. 5. Encoder configuration examples using the image chip.

The GLOC is responsible for the activation of all data transfers onto the memory bus, the synchronization of the chip with the system, and the scheduling of all processing on the chip. A micro-programmable sequencer initiates all block transfers on the memory bus. The user can change the GLOC micro-program to increase the flexibility of the general processing.

The VID_IN1, VID_IN2 and VID_OUT ports consist of double-buffered memories to fulfil the real time constraints. Either source pictures or decoded pictures can be used to perform the motion estimation, while typically the vector refinement uses only the locally decoded pictures. The IM block performs extraction and insertion of information and calculation results from and into the information bus. For multi-chip configurations, dedicated serial input and output links are provided that allow for the fast transmission of temporal results between parallel working chips. This chip link interface also contains logic that compares incoming temporal results with processing results provided by the chip. The better results (motion vector-MAE value pairs) are then sent via the link to the next chip.

The SDRAM_IF interface block is able to control up to two external 16 Mbit SDRAM memory chips running at full chip clock frequency and using a 16-bit data bus. If no picture reordering is required (i.e. for one phase vector tracing), one SDRAM is sufficient to store 3

frames, along with motion vectors and MAE results. The interface allows for flexible access of arbitrary regions within the stored frames, a feature required for vector post-processing. The SDRAM_IF is activated by GLOC instructions and then handles autonomously the complete required operation (e.g. loading a block of initialization motion vectors, a zone extraction, etc.).

A 16-bit host interface (HOST_IF) is provided to connect the board controlling CPU to the chip. The interface allows for the seamless connection to a variety of common DSPs and micro-controllers for both, synchronous and asynchronous modes.

Processing Module

The processing module is the central part of the image chip where all calculations take place. Its 24-bit architecture is depicted in figure 4. The module is divided into the following blocks which are described in detail below: a RISC controller or vector generation unit (VGU) for algorithm execution and overall module control with its private instruction memory, the pixel processor (PP), the pixel caches for search window and current MB, the sorting unit (SU), the random number generator (RU), and a set of dedicated registers. All units are connected through

a 24-bit high-speed bus that allows for parallel read and write operations by the RISC controller. The memory bus interfacing is done by the SU and the cache controller of the pixel processor whereas the interfacing to the shared host-interface bus is handled by a special block (HIFB_IF) that includes a crosspoint switch, data converter, and a bus access controller. During most of processing, the VGU controller is directly connected through this block to the module bus only when the VGU accesses data from outside the module and the HTFB_IF block requests the shared bus. Furthermore, it is possible to access all of registers of the module address space via the chip's host-interface. This is foreseen for use only during the chip's initialization phase to fill the RISC's instruction memory or for debugging purposes during normal chip processing.

The main task of the VGU is to generate a new candidate motion vector within the tight timing constraints given by MSE calculation time of the PP (16 clock cycles). The VGU base concept is a simplified standard pipelined RISC architecture, restricted to integer operations (no multiplication/division). Two different data types are to be treated: signed 2x12-bit (vector components) and unsigned 24-bit (MAE values and general data)

integer values. The VGU is a four-stage pipelined, 24-bit RISC architecture that executes one instruction per cycle. Standard RISC instructions have been combined with specialized instructions to speed-up algorithm execution time. A special synchronization instruction eases software polling. Most instructions exist for both data types and for register-register and register-immediate addressing. A subword parallelism approach allows for the treatment of two vector components in one cycle.

The 3-pipeline stage pixel processor (PP) is the central calculation unit of every module and is capable of calculating one MAE value at a random position in the search window every 16 cycles by using a 128-bit datapath. It consists of four main parts: the pixel search window and current MB caches with their cache controller, a barrel shifter and select block, the datapath, and the controller. The PP controller receives the candidate motion vector (2 vectors in dual prime case) and the operation instruction from the VGU. The PP controller communicates the end of the commanded operation through a status register that is polled by the VGU. The SU block consists of three main parts: sorter, best-result storage, and memory bus interface. The sorter represents the implementation of the best parent vector

Image processing power

Pixel processing power	3,5 GOPS	
RISC processor power		108 MIPS

Image chip data

Operating frequency		54 MHz/18 MHz
Implementation Technology	5 layer 0,35 μ m HCMOS6	
Estimated chip area		50 mm ² (pre-layout estimation)
Transistors count		1,5 M with memory/ 500 K without memory (pre-layout estimation)
VGU program size		1024 instructions each
Global sequencer microcode	max 32 instructions	
On-chip RAM		- 25 Kbytes
Max Memory Bandwidth	864 Mbit/s	
Power consumption		2 W (estimated)
Package		208 pin QFP (136 functional pins)
Design style		VHDL description + synthesis + standard cell layout

Tab. 1. Image chip features.

References

- [1] ISO/EC 13818-2: 1996 Information technology – Generic coding of moving pictures and associated audio information: Video.
- [2] P. Tudor, P. Brightwell, Cougar: Software Model Architecture for Hardware-Specific Optimization Experiments Proc. of the Hamlet RACE2110 workshop, pp. 85–94, February 27–28, 1996, Rennes.
- [3] M. Knee, P. Tudor, et al., Cougar: Hardware and Software for SNR MPEG-2 Codecs Proc. of the European Conf. on Multimedia Applications, Services and Techniques (ECMAST), pp. 243–253, Louvain-la-Neuve, May 1996.
- [4] Atlantic web homepage.
- [5] N. Wells, The Atlantic project: Models for Program Production and Distribution Proc. of the European Conf. on Multimedia Applications, Services and Techniques (ECMAST), pp. 243–253, Louvain-la-Neuve, May 1996.
- [6] M. Knee, N. Wells, Seamless Concatenation – A 21st Century Dream proc. of the Montreux 20th International Television Symposium, June 1997.
- [7] P. Tudor, O. Werner, Real Time Transcoding of MPEG-2 Video Bitstreams accepted for Int. Broadcasting Convention, Amsterdam, September 12–16, 1997.
- [8] P. Brightwell, S. Dancer, M. Knee, Flexible Switching and Editing of MPEG-2 Video Bitstreams, accepted for Int. Broadcasting Convention, Amsterdam, September 12–16, 1997.
- [9] P. Brightwell, T. Graham, Better Motion estimation techniques for MPEG-2 Coding, Proc. of the Hamlet RACE2110 workshop, pp. 29–37, February 27–28, 1996, Rennes.
- [10] M. Mattavelli, D. Nicoulaz: A Low Complexity Motion Estimation Algorithm for MPEG-2 Encoding Proc. of the Hamlet RACE2110 workshop, February 27–28, 1996, Rennes.
- [11] M. Mattavelli: Motion Analysis and Estimation: From Ill-posed Discrete Inverse Linear problems to MPEG-2 Coding, Parenthesis no. 1595, chapter 7, EPFL, 1997.
- [12] K. Suguri, T. Minami et al., A Real-time Motion Estimation and Compensation LSI with Wide Search Range for MPEG-2 Video Encoding IEEE Journal of Solid States circuits, Vol. 31, No 11, pp.1733–1740, November 1996.
- [13] H. Lin, A. Anesko et al.: A 14 GOPS Programmable Motion Estimator for H.26x Video Coding, Digest of Techn. Papers, IEEE Int. Solid State Circuits Conf., p. 246, February 1996.
- [14] R. Pacalet, A. Lafage et al. (ENST, Philips), A Real Time MPEG-2 MP@ML Motion-Estimator Chipset, Digest of Techn. Papers, IEEE Int. Solid State Circuits Conf., pp. 260, February 1997.
- [15] K. Ishihara, S. Masuda et al., A Half-Pel Precision MPEG-2 Motion Estimation Processor with Concurrent Three Vector Search, Digest of Techn. Papers, IEEE Int. Solid State Circuits Conf, pp.288–289, February 1995.
- [16] M. Mizuno, Y. Ooi et al. (NEC), A 1.5 W Single-Chip MPEG-2 MP@ML Encoder with Low-Power Motion Estimation and Clocking, Digest of Techn. Papers, IEEE Int. Solid State Circuits Conf., pp. 256, February 1997.
- [17] E. Ogura, M. Takashima et al., A 1.2 W Single-Chip MPEG-2 MP@ML Video Encoder LSI including Wide Search range Motion Estimation and 81MOPS Controller, Digest of techn. papers, IEEE Int. Solid State Circuits Conf., pp. 32, February 1998.
- [18] E. Miyagoshi, T. Araki et al., A 100 mm² 0,95W Single-Chip MPEG-2 MP@ML Video Encoder with a 128 GOPS Motion Estimator and a Multi-tasking RISC-Type Controller, Digest of techn. papers, IEEE Int. Solid State Circuits Conf., pp. 30, February 1998.
- [19] M. Gumm, M. Brochi et al.: A High Fault-Coverage Design-For-Testability Approach for a MIMD Based Multimedia Processor, Europ. Design & Test Conf, User Forum Volume, p. 91, Paris, March 97.
- [20] IEEE Standard Test Access Port and Boundary Scan Architecture, IEEE standard 1149.1a, 1993.

list necessary for the pseudo-genetic search. The sorter is capable of inserting a new vector/MAE-value pair at its proper position in a list within one clock cycle, depending on the MAE value.

The RU block provides random numbers with a distribution that approximates the Laplacian-shaped random number distribution, on which the original VT-MGS algorithm is based. Programmable standard derivations of the distribution for both, the x and the y components of the random vector are provided.

Image Application Examples

Figure 5 shows application examples of the image chip for different encoder architectures. In figure 5a, the candidate motion vectors is sent via the Infobus along with their corresponding refinement zone information, (1) co-timed with the input video data to the image chips. The first chip does the forward vector refinement and puts the results (e.g. 6 forward motion vectors) into the dedicated slots of the Infobus (2). The backward vector refinement chip receives these results and also puts its processing results (e.g. 6 backward motion vectors) on the Infobus (3) that is then sent to the last chip, which performs the prediction selection. This chip inserts the finally selected prediction mode and related information in the dedicated slots of the output Infobus (4), which contains then all necessary information for the final encoder circuitry. In figure 5b, the image chip is used primarily as a motion estimator. Unlike the former case, the Infobus input (1) supplied to the forward motion estimation chip is optional (guide vectors could be sent). They perform the traced genetic motion estimation with half-pixel refinement of the results. The transfer of preliminary and final results via the Infobus is similar to the first example.

Multi-Chip Configurations

The image architecture is scalable to allow for the enlargement of the search window or the increase of processing power. In figure 6, an example of a multi-chip configuration is depicted where 4 chips are used for each prediction direction to increase the overall processing power (i.e. the full search range) in vector refinement. For each direction, the chips are connected in a chain. The chips 1 + 2 (5 + 6) each

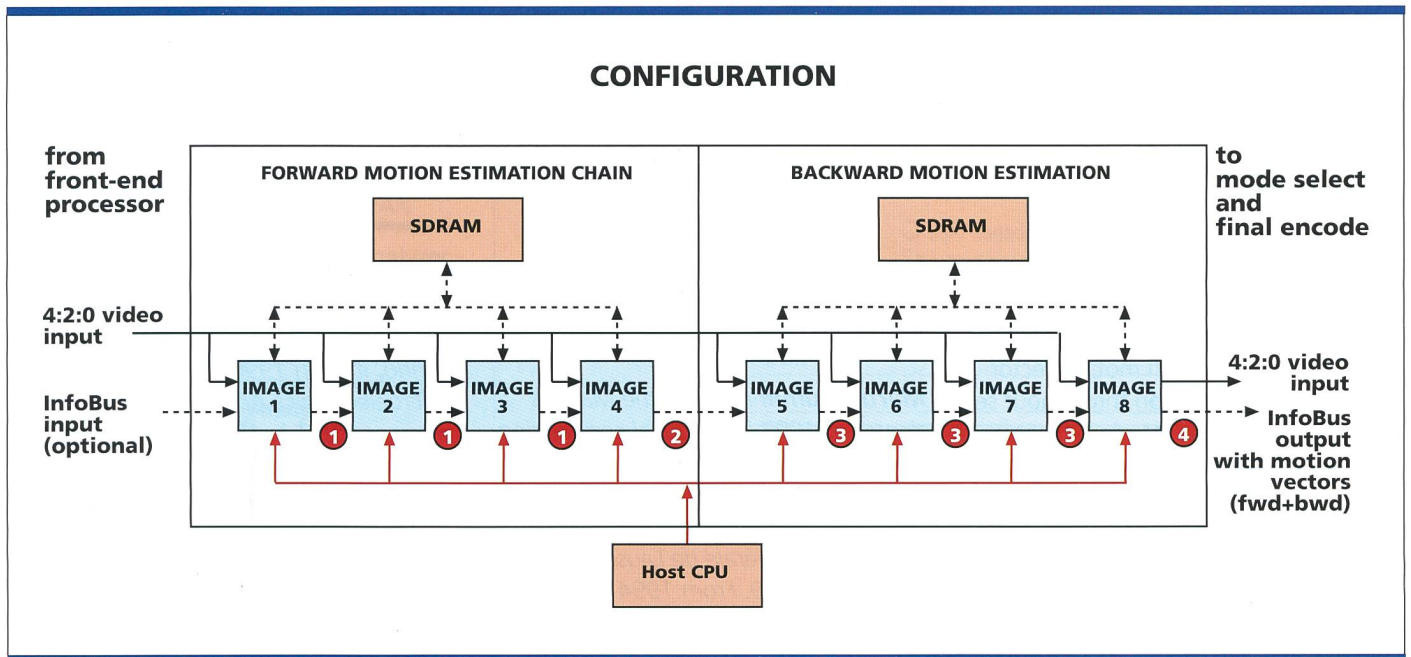


Fig. 6. Image multi-chip configuration example (increase of the search window to 192×128 pixel). Tab. 1. Image chip features.

operate on one 16×16 -line candidate vectors, the chips 3 + 4 (7 + 8) each on two 16×8 -line candidate vectors. Results are passed via a dedicated link to the successor (1) until they reach the last chip of a chain that puts the complete set of forward prediction motion vectors in their dedicated slots into the Infobus (2). The vectors are then sent to the second chain of chips, which are responsible for the backward prediction. The last chip of each chain also controls the common SDRAM block. The backward chain works in the same way; the results are passed through the chain (3) and finally put in the corresponding Infobus slots (4). A configuration similar to the one described here can be used to increase the search window by the factor four (± 80 H/ ± 48 V) if required.

Test Approach

A thorough analysis of the design-for-testability (DFT) aspects has been undertaken for the entire chip. [19] A three-fold DFT approach was implemented, consisting of the combination of partial scan design, built-in self-test (BIST) techniques (for memory and large combinatorial blocks like datapaths), and their integration into a JTAG [20] controlled hierarchical test structure. Most of the 43 distinct physical memory blocks on the chip can be tested by two configurable array built-in self-test (ABIST) blocks that execute a standard memory test algorithm. The two VGUs in the

modules are used in test mode as local test controllers. They are capable of performing several different self-tests in the module. The chip's JTAG test interface provides control over all chip self test functions and access to most chip registers for debugging besides the standard functions.

Implementation Results

The chip has been described in VHDL and synthesized onto a standard cell library of a 5-metal layer 0.35 μm CMOS process. Most of the chip's core logic operates at a clock frequency of 54 MHz, and only the video and Infobus ports operate at 18 MHz to interface seamlessly

Zusammenfassung

Der Image-Chip

Der Artikel beschreibt einen neuartigen, vielseitigen Motion Estimation Chip, der mit Hilfe einer neuen Methode die Qualitätsverluste in kaskadierten Dekoder-/Encoder-Systemen minimiert. Der Chip wurde unter besonderer Berücksichtigung einer gleichzeitig entwickelten Hardware zur Videokodierung für Fernsehstudios konzipiert. Er realisiert einen unkonventionellen, leistungsstarken Motion-Estimation-Algorithmus, der die Bewegung über mehrere Videobilder zurückverfolgt und auf einer genetischen Suchstrategie basiert. Der Chip kann eine Vielzahl verschiedener, mit der Bewegungsanalyse in Zusammenhang stehender Aufgaben erfüllen und bietet eine Reihe von Optionen zur Erreichung einer optimalen Kodierqualität. Diese Funktionalität wurde in eine parallele, auf dem MIMD-Prinzip basierende Architektur integriert. Sämtliche Algorithmen werden von integrierten RISC-Prozessoren ausgeführt, was eine hohe Flexibilität ermöglicht. Dank mehrerer Strategien für Multi-Chip-Konfigurationen ist es möglich, die Kodierqualität je nach Systembedürfnissen zu erhöhen. Der Chip wurde in 0,35- μm -CMOS-Technologie auf einer Fläche von etwa 60 mm^2 realisiert.

with the video system. Special attention was paid to keeping the chip's power consumption low by using gated clock strategies and low power RAMs provided by the manufacturer. Table 1 summarizes the relevant technical data for image.

Summary of Results

The image chip implements the new, powerful VT-MCS motion-estimation algorithm that provides near full-search quality and delivers smooth motion vector fields by tracking motion throughout several frames. Besides standard features such as half-pixel accuracy and the calculation of all MPEG-2 defined prediction costs, several enhancements such as DC-removed search, search window offsetting or biasing of matching results towards the (0,0) vector are available. The chip requires very limited interaction with a board controller, mainly for chip setup and for uploading some information that changes the picture rate. All algorithm execution is software based providing high flexibility to fine-tune the search algorithm to the needs of a single application. Due to the flexible concept, required vector refinement and prediction selection functions can be executed by the chip, with various optional enhancements, just by loading new VGU assem-

bler programs and another GLOC micro-code. An integrated SDRAM interface guarantees a high frame memory bandwidth and a generic host interface allows for the use of different types of board controllers. An integrated test interface

for controlling different self-tests and for debugging is provided. The chip is scalable for the use in different systems, from a low-delay coder to a powerful and high quality studio encoder. 4

Martin Gumm received the Masters degree in 1993 in electrical engineering from the Technical University of Darmstadt, Germany, and the Ph.D. degree in technical sciences from the Ecole Polytechnique Fédérale de Lausanne, Switzerland, in 1999. From 1994 to 1995 he worked as a research and teaching assistant at the University of Stuttgart. In 1996 he joined the multimedia group of the Integrated Systems Center at EPFL where was responsible for the design group involved in the IMAGE design. His interests are VLSI architectures for multimedia applications and real-time systems, design-for-testability, and high-level HDL-based design. E-mail: martin.gumm@epfl.ch

Pierangelo Garino received the degree in Computer Science at the University of Torino, Italy, in 1990. He joined the Microelectronics Applications Department at CSELT in 1983. Since then he has been involved in different issues concerning digital circuits, including design methodologies as well as the design and test of VLSI circuits. He participated to several international co-operative research projects, dealing with the audio-video coding techniques and the design of ASIC components for multimedia applications. He was the coordinator of the design of the IMAGE circuit at CSELT. His research interests are in System on Chip design, including HW/SW code-sign/co-verification techniques and low power design, related to multimedia and wireless communication systems. E-mail: Pierangelo.Garino@CSELT.IT

Papierdünne Chips im Kommen

Handys und Palmtops sind die Treiber für die Entwicklung immer dünnerer Chipgehäuse. Das setzt aber auch voraus, dass die Chips selbst immer dünner werden. Die dünnsten heute verfügbaren Wafer sind ohnehin nur noch 0,15 mm «dick». Nach Angaben der japanischen Wirtschaftszeitschrift «Nikkei Sangyo» hat Tokyo Seimitsu Geräte und Verfahren entwickelt, um Wafer bis auf eine Dicke von nur 30 µm chemisch-mechanisch abzuschleifen. Dabei werden besondere Vorkehrungen getroffen, um ein Brechen der Chips durch die Vielzahl von kleinen Vertiefungen auf der Waferrückseite (Durchmesser um die 1 µm) zu verhindern: Sie werden chemisch aufgefüllt. Das Unternehmen glaubt, dass sich solche ultradünnen Chips für die Personenidentifizierung einsetzen lassen oder direkt in Papier oder Bekleidung für Kassenterminals eingebaut werden können.

Japans Postministerium will nicht kleckern, sondern klotzen

Um gleich 25% will das japanische Postministerium seinen Etat für das Finanzjahr 2000 (beginnend am 1. 4. 2000) erhöhen – das entspricht rund 1,2 Mia. US-\$ mehr als im laufenden Jahr. Ein Viertel dieser Mehrausgaben soll in das Lieblingsprojekt des Ministeriums gehen: Die Transformation Japans zu einer «Information Society». Dazu gehören der Ausbau von internetbasierten Dienstleistungen, Entwicklungsarbeiten für «Smart Gates» (also intelligente Zugangsverfahren), die digitalen Hörfunk- und Fernsehdienste sowie für die bereits mehrfach kolportierte «Stratosphäre Wireless Platform» (eine Initiative für Satellitensysteme der nächsten Generation). Ob die durch Bankenskandale und zahlreiche wirtschaftliche Ankurbelungsprogramme geschwächte Finanzlage des Staates eine solch kräftige Erhöhung hergibt, wird sich wohl erst zeigen müssen.

Superschnelles Bildverarbeitungssystem

Hamamatsu und die Universität Tokyo haben zusammen ein Bildverarbeitungssystem entwickelt, das Bildaufnahme und Bildverarbeitung in 1/500 Sekunde abwickeln kann. Das System arbeitet ähnlich wie ein Parallelrechner: Alle 128x128 Pixel haben ihre eigenen Bildprozessoren, die alle gleichzeitig arbeiten. Weiter wird die jüngste hoch empfindliche Bildaufnahmetechnik von Hamamatsu eingesetzt. Die Partner glauben, dass sich das System auf vier Chips unterbringen lässt und für deutlich weniger als 1000 US-\$ im nächsten Jahr an den Markt gebracht werden kann.

Hamamatsu Photonics K.K.
325-6, Sunayama-cho
Hamamatsu-shi
Shizuoka 430
Japan
Tel. +81-534-52 2141
Fax +81-534-52 2139

LAN • WAN ATM • ISDN

Messtechnik aus einer Hand



RADCOM

- ATM/LAN/WAN und ISDN Protokoll-Analyzer
- SDH-Analyzer
- H.323 Analyse
- SS7 Jitter
- Ausbaubar & mobil
- Internetworking Analysen



SHOMITI

- 10/100MB/s LAN-Tester
- GigaEthernet-Analyzer
- Analyzer-Switches
- TokenRing-Analyzer
- «State of the Art»
- Erweitert Notebook zum Analyzer

Apptitude

Application Performance?



- Applikations-Analyse Tools
- LAN & WAN Probes
- Health-Software



- SS7/ISDN/PCS/GSM Analyzer/Simulator
- Bulk Call Simulator

FE Frederick Engineering

- preiswerte ISDN/WAN Analyzer
- BERT-Tester
- FrameRelay Analyzer/Simulator

Mess-Lösungen für Netzwerke

Hinterbergstr. 9
6330 Cham
Switzerland
Tel: 041-748 60 10
Fax 041-748 60 11
info@emitec.ch
www.emitec.ch



EMITEC



*Wir kennen,
was wir vermitteln!*

Sema Sprachreisen
Karstgässchen 4
8201 Schaffhausen
Tel. 052 625 68 25, Fax. 052 624 06 32
www.semasprachreisen.ch

Swisscom-Mitarbeiterabbatt
5 %

- | | |
|---|---|
| <input type="checkbox"/> USA/ Kanada | <input type="checkbox"/> England/ Irland |
| <input type="checkbox"/> Australien/ Neuseeland | <input type="checkbox"/> Frankr./ Italien/ Costa Rica |

Name: _____

Strasse: _____

Plz/Ort: _____

Tel. _____

comtec



Skalpelli ...

Zange ...

Domino Gigabit!

**Platz 1 für den
Gigabit-Analysator im
Network Computing
Magazine**

Die Kommunikations-Verbindungen sind die Nervenstränge Ihres Unternehmens, sie müssen den ständig steigenden Belastungen gewachsen sein. Höhere Übertragungsraten und eine wesentlich intensivere Nutzung verlangen ein völlig neues Netz-Management. Die Gigabit-Technik ist heute bereits Realität. Im Notfall muss eine schnelle Diagnose und Problemlösung her. Wavetek Wandel Goltermann hat die Lösung:

DominoGigabit, ein Internetwork-Analysator mit hervorragenden Leistungsparametern, einzigartiger Messgenauigkeit und grosser Mobilität. Damit können Sie sicher sein, alles Menschenmögliche für Ihr Kommunikationsnetz getan zu haben.

Rufen Sie uns an: **+41 31 996 44 11** oder besuchen Sie uns auf unserer Web-Seite: **www.wwg solutions.com**



WAVETEK
WANDEL
GOLTERMANN
Communications Test Solutions