



Dedicated Optoelectronic Stochastic Parallel Processor for Real-Time Image Processing: Motion-Detection Demonstration and Design of a Hybrid Complementary-Metal-Oxide Semiconductor- Self-Electro-Optic-Device-Based Prototype

Alvaro Cassinelli, Pierre Chavel, Marc P. Y. Desmulliez

► To cite this version:

Alvaro Cassinelli, Pierre Chavel, Marc P. Y. Desmulliez. Dedicated Optoelectronic Stochastic Parallel Processor for Real-Time Image Processing: Motion-Detection Demonstration and Design of a Hybrid Complementary-Metal-Oxide Semiconductor- Self-Electro-Optic-Device-Based Prototype. *Applied optics*, 2001, 40 (35), pp.6479-6491. 10.1364/AO.40.006479 . hal-00878984

HAL Id: hal-00878984

<https://hal-iogs.archives-ouvertes.fr/hal-00878984>

Submitted on 31 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dedicated optoelectronic stochastic parallel processor for real-time image processing: motion-detection demonstration and design of a hybrid complementary-metal-oxide semiconductor–self-electro-optic-device-based prototype

Alvaro Cassinelli, Pierre Chavel, and Marc P. Y. Desmulliez

We report experimental results and performance analysis of a dedicated optoelectronic processor that implements stochastic optimization-based image-processing tasks in real time. We first show experimental results using a proof-of-principle-prototype demonstrator based on standard silicon–complementary-metal-oxide-semiconductor (CMOS) technology and liquid-crystal spatial light modulators. We then elaborate on the advantages of using a hybrid CMOS–self-electro-optic-device-based smart-pixel array to monolithically integrate photodetectors and modulators on the same chip, providing compact, high-bandwidth intrachip optoelectronic interconnects. We have modeled the operation of the monolithic processor, clearly showing system-performance improvement. © 2001 Optical Society of America

OCIS codes: 200.0200, 200.4650, 250.3140, 030.6140, 100.0100, 330.4150.

1. Introduction

Many early- and middle-vision tasks such as restoration, image segmentation, halftoning, or even motion detection can be formulated as optimization problems consisting of finding the ground states of an energy (or cost) function.¹ Stochastic optimization by simulated annealing² yields some of the best results because it is not generally constrained to quadratic or even convex cost functions. However, when implemented on conventional sequential workstations, the computational loads are too extensive for practical purposes, so that most of the proposed artificial vision systems rely—one way or another—on deterministic optimization algorithms that are at

least 100 or 200 times faster than stochastic algorithms but often lead to results further from optimality.

One can point out two different sources for such a prohibitive computational load. One obvious source is the sequential processing of pixels in a two-dimensional (2-D) image; one way of alleviating this burden might be the processing of pixels in parallel, with dedicated hardware based on optical scale parallelism—by this expression we refer to parallel image processors consisting of one processing element (PE) per pixel. This is possible, provided that some rules are observed: Briefly stated, two pixels that interact with each other should not evolve at the same time; this rule has led us to the concept of semishift invariance and coloring in optoelectronic parallel processors³ (see Fig. 1). In the case of low-level image-processing tasks, the state of each PE of the array evolves as a function of a short-range shift-invariant interconnection pattern: Therefore stochastic (as well as deterministic) algorithms can advantageously be implemented on single-instruction-multiple-data machines with a high degree of parallelism. In this context, optical convolution represents an interesting alternative to intrachip electronic shift-invariant interconnects be-

A. Cassinelli and P. Chavel (e-mail address: pierre.chavel@iota.u-psud.fr) are with the Laboratoire Charles Fabry de l'Institut d'Optique, Central National de la Recherche Scientifique, B.P. 147, 91403 Orsay Cedex, France. M. P. Y. Desmulliez is with the Department of Computing and Electrical Engineering, Heriot-Watt University, Riccarton, Edinburgh EH14 4AS, Scotland, United Kingdom.

Received 19 December 2000; revised manuscript received 16 August 2001.

0003-6935/01/356479-13\$15.00/0

© 2001 Optical Society of America




Fig. 1. PE coloring in the case of a four- or eight-nearest-neighbor-interconnection pattern. Interlaced color domains must evolve consecutively, but all pixels in one color domain can evolve in parallel; hence parallelism scales inversely with the number of colors (two and four in the left and right examples, respectively).

cause they avoid electronic layout congestion and have the benefits of an easy reconfigurability of the interconnect topology. The principles of such a mainly digital optical processor architecture, combining electronic nonlinear processing and optical linear convolution have been studied extensively in the literature under various names, including symbolic substitution, mathematical morphology, and cellular automata.^{4–6}

Whereas sequential updating of pixels is a common source of computational load shared with any deterministic algorithm, stochastic optimization algorithms suffer from the second burden of having to generate good quality, uncorrelated random numbers for every PE in the array; moreover, proper operation of the simulated annealing needs the generation to be repeated thousands of times per image under processing and even more for complex tasks. In the parallel version of the algorithm, huge amounts of (time and space uncorrelated) random numbers need to be generated to feed all the PEs of the array at the same time. Electronic (digital) random-number generators are not only time and chip-area consuming: They also suffer from the inherent pseudorandom nature of the generated sequence.⁷ On the basis of extensive previous experience,⁸ we selected differential detection of speckle as a reliable, quick, parallel, and easy way to control the generation of random numbers.

We have therefore been studying optoelectronic single-instruction-multiple-data machines that are able to minimize low-level image-energy functions at a video rate by use of the well-known simulated annealing algorithm. These devices, which we will refer to from now on as optoelectronic stochastic parallel processors (OSPPs) rely on three ingredients: (a) an optical convolution setup in charge of the reconfigurable, shift-invariant interconnection pattern; (b) an optical random-number array generator; and (c) an optoelectronic smart-pixel array (SPA) to implement the required nonlinear operation. So far, only inputs, not outputs, to our SPA have been optical, and convolution has been achieved either elec-

tronically (thanks to a hard-wired four-nearest-neighbor interconnect) or with the help of an optical convolution setup composed of an interchangeable Damman grating and a spatial-light-modulator (SLM) chip interfaced to the SPA through a—relatively—slow electronic bus; indeed, as we will show, monolithic convolution of the data using a SPA having both optical input and optical output ports would account for an enormous increase in the overall performance of the OSPP system.

The paper is organized as follows. In Section 2 we review our present setup, which is, to our knowledge, the first complete reconfigurable (nonmonolithic) OSPP prototype dedicated to low-level image processing. This proof-of-principle prototype is based on a silicon complementary-metal-oxide-semiconductor (CMOS) SPA chip. Motion detection at a video rate in a sequence of gray-level images is demonstrated. In Section 3 the flip-chip bonding technique is considered for building a hybrid Si/AsGa SPA, and figures of merit of this novel OSPP are drawn up as a function of the optoelectronic transceiver characteristics.

2. Present Demonstrator

The underlying principles of the OSPP are described extensively elsewhere.^{9,10} We report here on our latest experimental illustration, namely, the demonstration of video-rate simulated annealing for motion detection in a sequence of noisy gray-level images. The purposes of this demonstrator based on a Si-CMOS SPA are twofold: First, it provides a concrete optomechanical hardware that should help us draw realistic figures of merit while we test novel technologies for the SPA processor; second, it constitutes a proof of the good prospect of using stochastic-based algorithms at the core of more complex image-processing systems, without compromising the possibility of real-time processing—at least up to the video rate.

A. Motion-Detection Algorithm

From the point of view adopted in this study, motion-detection algorithms and systems provide an output consisting of a binary map \mathbf{e} , called the label field, indicating motion or absence of it at every pixel of the input (noise-corrupted) gray-level image. There is evidence that motion detection can be achieved with a statistical regularization approach based on spatio-temporal Markov random fields.¹¹ Far from being of only academic interest, motion detection is therefore an interesting demonstration candidate for our video-rate processor.

Specifically, we started from a class of motion-detection algorithms first proposed by Lalande and Bouthemy.¹¹ The original algorithm performs an elementary preprocessing of the sequence to prepare two observed or input data fields. First, a time gradient $\mathbf{o}(s, t)$ is obtained by differentiating in time a gray-level image, where s is a 2-D pixel coordinate and t is discrete in time. Then a binary version $\hat{\mathbf{o}}(s, t)$ of the time gradient is calculated also,




Fig. 2. Underlying principles of the simplified motion-detection algorithm. The heavy circles represents two consecutive locations of a moving object. Motion information is exclusively provided to the OSPP through the observed binary field $\hat{o}(t)$.

simply with a predefined threshold. The energy function to be minimized [as a function of the 2-D label field $e(s, t)$ under estimation] incorporates both input fields as well as the previous state of the label field $e(s, t - dt)$.

In our setup the host computer is in charge of all the preprocessing. Although we started from the Lalande and Bouthemy algorithm, we found that it was necessary to introduce some simplifications to it before parallel optoelectronic implementation could be envisaged. Specifically, the need for two time-gradient input data fields $o(s, t)$ and $\hat{o}(s, t)$ was a problem. However, we concluded from simulations that it was worthwhile to investigate a model in which only the binary (thresholded) version \hat{o} of the image gradient is used as input to the algorithm. The whole idea relies on the assumption that the state of motion is likely to change only if temporal changes are notable—this is the case for a pixel lo-

cated on the edge of a moving object. Otherwise, we would prefer the label to retain the value held in the near past. Spatial regularization potentials are also defined in such a way as to favor homogeneous regions of motion. A schematic representation of the underlying principles of the simplified algorithm is shown in Fig. 2. It is appropriate here to mention the difference between motion detection and the mere subtraction of successive frames in a time sequence: A nonzero difference between two consecutive images may indicate motion, but it may as well be due to time-variable noise or a change in illumination conditions; in contrast, a null difference may be the result of subtracting consecutive pixel values at two different locations in the body of a uniformly lit moving object (see Fig. 3). The purpose of motion detection is to give a complete estimate of all pixels in the scene that are moving at a given time and follow their motion in terms of compact, but deformable, moving objects.

As expected, estimation of the current label field involves only neighborhood computations: The local energy gradient with respect to the label variable depends only on a short-range neighborhood around the pixel under consideration. Various strategies have been proposed for the order of site visiting (at random, by sequential scanning, or with a confidence ordering list to accelerate the relaxation process as in Ref. 12). However, as explained above, the parallel updating of sites happens to be the fastest strategy because the neighborhood size reduces to the four-eight nearest neighbors. Also, whereas many deterministic optimization algorithms require that much attention be paid to the label-field initialization procedure (the ultimate quality of the result strongly depends on it¹³), stochastic-based algorithms completely eliminate such a constraint. In short, all a PE has to do is

(a) Calculate the local force (i.e., the local energy gradient at a particular site s and time t with respect to the label variable), which combines one regular-




Fig. 3. Test of the simplified motion-detection algorithm implemented on the (nonmonolithic) optically interconnected OSPP prototype. A Dammann grating is in charge of an eight-nearest-neighbor-optical-interconnection pattern. This first prototype relies on a low-resolution, 24×24 -pixel large Si-CMOS chip (SPIE600). Gray levels, on a scale of 0 through 255: background = 180; object = 220. Gaussian noise, $\sigma = 10$. Object speed, $(1, \pm 2)$ pixel/frame. Size: image, 24×24 pixels; moving object, 6×5 pixels.

ization term F_S and a second (conditional) constraint-to-the-data term F_C :

$$F(s, t) = \beta_S \sum_{r \in N(s)} [2\mathbf{e}(r, t) - 1] \quad (=F_S) \\ + \beta_C [2\mathbf{e}(s, t - 1) - 1][2\hat{\mathbf{d}}(s, t) - 1] \quad (=F_C). \quad (1)$$

In Eq. 1, $N(s)$ represents the neighborhood of pixel s and β_S and β_C are empirically determined constants. The term $F_S(t)$ can be seen as a regularization field, resulting from a convolution between the current field $\mathbf{e}(t)$ and a short-range (equally weighted) convolution kernel. The term $F_C(s, t)$ behaves as a conditional constraint to the label-field data held in the near past (memory effect), provided that no significant change has been detected in the particular site [i.e., the variable $\hat{\mathbf{d}}(s, t)$ is zero]; otherwise, it tends to flip the label state, as explained above (Fig. 2). Our simulations have shown proper operation of the algorithm, provided that the ratio $\lambda = \beta_C/\beta_S$ lies between 0.4 and 1.2.

(b) Update its label $\mathbf{e}(s, t)$ according to the sigmoid probability law:

$$\text{Pr}[\mathbf{e}(s, t) = 1] = \frac{1}{1 + \exp[-F(s, t)/T]}. \quad (2)$$

Updating is performed by simply thresholding on the local force $F(s, t)$ by a standard electronic comparator; the nature of such a deterministic operation is turned probabilistic thanks to a laser speckle generator that projects a time-varying speckle light over two photodetectors attached to both positive and negative inputs of the comparator (the decision threshold is therefore randomly shaken). Characterization of these stochastic comparators (or random-number generators) is given in detail in Ref. 14.

Thanks to operations (a) and (b), the updating of a particular PE state is thereupon achieved. Updating can be done in parallel on pixels that do not directly interact with each other. Maximal sets of pixels that do not interact are said to belong to the same color domain, and colors domains are to be considered sequentially. As a result, the entire field $\mathbf{e}(t)$ is eventually updated. It should be reminded here that two neighboring (i.e., interconnected) processor cells never have the same color; therefore as the size of the neighborhood increases, so does the number of color domains (see Fig. 1). The process iterates following the annealing schedule by decreasing the speckle laser power that represents the (algorithmic) temperature (T) of the annealing.^{1,2} With this simple algorithm, one serious simulated annealing operation consists of typically 1000 updating cycles of the whole field $\mathbf{e}(t)$.

B. Smart-Pixel Array Hardware

Because of the general availability and maturity of silicon technology, in an earlier collaborative project with Institut d'Electronique Fondamentale, Centre National de la Recherche Scientifique, Université de

Paris-Sud,⁹ we developed a CMOS SPA, named SPIE600, consisting of 24×24 identical processor cells for the optoelectronic simulation of an Ising spin array with two phototransistors per cell serving as optical inputs for the speckle light, used to compute the updating probability of Eq. (2) by differential detection of speckle.¹⁵ Of course, because light emission by silicon is far from being well mastered, there is no optical output. The optical inputs, however, are sufficient for the Ising spin problem. Although, indeed, the optical inputs may be used for the neighborhood convolution as explained in the Introduction, a provision for electronic bipolar interaction with the four nearest neighbors has been implemented. That study, published earlier, demonstrated a video real-time massively parallel implementation of an Ising spin array.¹⁶

It appeared that the same chip could be used for demonstrating some low-level image tasks. For instance, the chip could be used for noise cleaning on binary images—a problem of only academic interest though—by simply setting all interaction coefficients of the Ising chip to unity and then optically inputting the binary image to be processed (dual rail encoded as explained in Ref. 4) onto the photodetectors on the chip. This is possible because each PE is provided with two optical inputs. This was actually done as the first full experimental demonstration of an OSPP for image processing (noise cleaning in binary images) at the video rate, exploiting the concept of semi-shift invariance by coloring.¹⁷ In the present study, optical convolution and electronic feedback have been added to the OSPP to broaden the interconnection neighborhood, even though the task demonstrated is motion detection. This will be explained in some detail below.

C. System Architecture

There is a fundamental interest in studying and evaluating the performances of an optically interconnected prototype, and the noise-cleaning OSPP demonstrator is naturally suitable for this kind of experiment. Indeed, optical inputs to the PEs can be used as channels for communicating between each other, provided that their outputs are made somehow optically available. As explained in Section 1, optical convolution is an interesting solution for the interconnection hardware because all contemplated applications of the OSPP (noise cleaning, halftoning, motion detection, etc.) need a shift-invariant interconnection pattern. Moreover, as the interconnection topology depends on each particular application, an easy way of reconfiguring it is also needed. Optical convolution also provides extension to larger convolution kernels, compromising neither the electronic complexity of each PE nor the density of the array.

With respect to the noise-cleaning demonstrator mentioned above, two major improvements were required. (1) Two SLMs were added to the system: one for optically generating the constraint-to-the-data (precalculated) binary field $F_C(t)$ of Eq. (1), and




Fig. 4. Schematic overview of the nonmonolithic OSPP demonstrator, including an optical input arm (folded), an optical convolution setup (unfolded arm), and an optical random-number generator (only the exit of the speckle fiber is represented). Polarizing beam splitters are used to increase the overall system throughput.

the other for performing the convolution in an optoelectronic feedback loop, with a convolution filter (a two-level-phase Damman grating in our experiment), so that regularization term F_S of Eq. (1) would become optically available (see Fig. 4). Of course, the optical feedback was implemented with a bulky optical setup significantly increasing the demonstrator volume. The on-chip electronic interconnects were disabled. (2) All the optical parts of the demonstrator and the associated mechanics had to be designed for easy adjustment, stable operation over a long period of time, and easy transportation, e.g., for presentation at seminars and conferences. Figure 5 represents a view of the complete setup. Details of the optomechanic implementation and testing of the nonmonolithic macro-optical prototype will be published in an article presently under preparation.

D. Experimental Results and Demonstration

The demonstration of the modified Lalande-Bouthemy algorithm was successfully completed with the OSPP prototype described above, yielding good-quality moving pixel estimates over synthetic noisy gray-level moving sequences (see Fig. 3). Whereas the video rate was compatible with the components, nonoptimal operation of the Windows computer interface operating with the driving personal computer slowed the process to approximately 5 s per full simulated annealing when the four-nearest-neighbor-interconnecting hologram is used (the array is then divided into two different color domains, see Fig. 1) and to approximately 24 s in the case of a height-nearest-neighbor-interconnecting hologram (leading to four color domains, see Fig. 1). Because our images are 24×24 pixels large, the level of parallelism is $24 \times 24/2 = 238$ in the case of the four-nearest-neighbor-interconnected array, and $24 \times 24/4 = 119$ in the case of a wider eight-nearest-neighborhood-interconnection pattern, respectively. For a fixed interconnection pattern, the level of par-




Fig. 5. View of the complete setup that uses the SPIE600 chip. It includes a CCD camera for alignment purposes and continuous monitoring of system operation. Approximate dimensions are 35 cm \times 21 cm \times 14 cm (the random-number-generator hardware is not shown in the picture).

allelism scales with the number of pixels and is restricted only by the cost of dedicated chip design and fabrication. The SPIE600 chip was designed in 1994; with presently available technology, it would be possible to design a new version of the chip with many more pixels, perhaps several hundreds, on a side, and the processor operation would be exactly as fast.

3. Design of a Hybrid CMOS-SEED-Based OSPP Prototype

If video-rate stochastic processing is to be achieved along with good results in terms of quality of the optimization (that is, at least 1000 updating cycles per annealing), then our optically interconnected SPA has to be provided with optoelectronic feedback that works at rates higher than $25 \times 1000 \times N_C$ loops/s (where N_C represents the number of color domains in the processor). Actually, our current OSPP demonstrator uses a binary ferroelectric liquid-crystal SLM chip for feedback and convolution; even if we neglect the time needed for reading and properly formatting the data flowing through the electronic loop (controlled by a 100-MHz Pentium host computer), the SLM full-frame rate is less than 2 kHz—more than 20 times under the minimum required rate in the simplest case of two-color operation of the processor array, so that only hardly 50 updating cycles could in principle be achieved at the video rate—a less than satisfactory situation. As a matter of fact, placing sources and detectors in different chips is an unsatisfactory solution because it leads to an electrical bottleneck when a full 2-D binary data field is routed from one chip to another. This is in-

deed what limits the SLM full-frame rate. We thus contemplated the design of a new version of the OSPP, relying on a unique SPA in which all PEs would be electronically isolated from each other (except for the clock and probably some additional control signals) and would include *in situ* high-bandwidth optoelectronic transducers.

A. Optoelectronic-VLSI Chip Architecture

It is appropriate at this stage to review the operation of an optically interconnected version of the OSPP. Each PE of the SPA would require the following features:

1. At present, a pair of photodetectors to implement the probabilistic updating by differential detection of speckle. These inputs are also used to calculate the local force $F_S(s, t)$ by analogous (intensity) addition of the neighborhood output optical signals. It is interesting to note that differential encoding allows robust operation under low-contrast input data conditions, which is generally the case when the signal is provided by some kind of modulator-based source;
2. A comparator (point nonlinearity) in charge of the thresholding operation that sets every PE into its new state;
3. A one-bit memory, defining the current state of the PE;
4. A dual-rail optoelectronic output to which the current state of the PE is directed so that this state can be read by the neighborhood PEs; these outputs can be either sources [light-emitting diodes (LEDs) or vertical-cavity surface-emitting lasers] or modulators such as p-i[multiple quantum well (MQW)]-n diodes.

The basic PE design then corresponds to a differential receiver enhanced with an optoelectronic transmitter device (known in the literature as a differential transceiver or two-beam transceiver).⁹ Recently there has been significant progress in the development of these receivers for data-communication and information-processing purposes, and many publications describe well-characterized and functional optoelectronic transceivers.^{18,19} The efforts in designing our elementary processor are then considerably reduced.

To make the coloring of the SPA easily reconfigurable, we can add additional features:

5. A three-bit memory to encode the corresponding PE color (five colors are enough for the twelve-nearest-neighborhood-interconnected SPA) along with a binary counter, so that only one common clock signal (ticking the binary counter) would be needed to control the sequential updating of color domains (this signal may be provided optically as well). Also, the whole optical architecture of the OSPP demonstrator can be greatly simplified (saving one SLM and its corresponding optical arm) with a final improvement;
6. An additional one-bit memory for storing the constraint-to-the-data term $F_C(s, t)$ during the whole




Fig. 6. Block diagram of a single OSPP smart pixel (\mathbf{a}_1 , \mathbf{a}_2 , \mathbf{a}_3 , \mathbf{e} , and \mathbf{b} are one-bit memory registers, loaded in parallel over the rows).

simulated annealing process. Sequential loading of color numbers as well as the precalculated binary constraint-to-the-data field $F_C(t)$ can be done in parallel over the rows of the array before the simulated annealing iterations are started. The technique has been tested on the present chip SPIE600: A full loading of the chip registers can be done in less than 3 μ s, which would hardly affect the 40 ms needed to perform simulated annealing at the video rate. The schematic of a single PE with one dual-rail optical input and one dual-rail optical output is shown in Fig. 6. The optical architecture of the optically interconnected OSPP is shown in Fig. 7. A holographic array illuminator creates an array of reading beams to be projected onto the PE modulators. Thanks to a cube beam splitter and a computer-generated hologram, the resulting reflected signals (from the modulators in the ON state) are convoluted and instantaneously back projected to the photodiodes of the SPA (a process that we referred to earlier as monolithic feedback). Relative weighing of each




Fig. 7. Optical architecture of the monolithic OSPP. An array illuminator is used to generate a reading bunch of beams to be projected onto the PE modulators. The laser source operates in a quasi-cw mode. PBSC, polarizing beam-splitter cube; CGH, computer-generated hologram; AI, array illuminator; SPA, smart-pixel array.

Table 1. OSPP Performance As a Function of the SPA Technology^a

Technology	Assumed Response Time	Number of Simulated Annealing cycles per Second ($N_C = 4$)
The photthyristor array that we tested	26 ms	0.01
SPIE600 (nonmonolithic demonstrator: CMOS chip and liquid-crystal SLM)	400 μ s	0.6
S-SEED (MQW)	50 μ s	5
FET-SEED (MQW)	10 ns	25,000
LED(MQW)-MSM	4 ns	62,500
CMOS-SEED (MQW)	1.6 ns	150,000
VCSEL(MQW)/MSM	1.6 ns	150,000

^aPerformance in terms of simulated annealing cycles per second by use of different technological approaches (a single simulated annealing operation represents 1000 updating cycles per color domain; $N_C = 4$ is the number of color domains). There is a gap between electronic enhanced detectors [hybrid CMOS-SEED or field-effect-transmitter (FET)-SEED devices] and simple symmetric-SEED³⁶ (S-SEED) devices. MSM: metal-semiconductor-metal detector; VCSEL: vertical-cavity surface-emitting laser. The photthyristor device that we tested was a preliminary test device and is believed to be far from the limit of that technology.

contribution F_S and F_C to the local force of Eq. (1) is done by adjustment of the power of the optical source with respect to a fixed electrical input to the chip (the “constraint term β_C ” input line of Fig. 6).

B. Optoelectronic Transducers

In comparing different input–output devices, one has to consider several competing parameters such as power consumption and dissipation, sensitivity, response time, and driver layout area. Better optoelectronic transducers do not only lead to improved performances of the whole system but they also potentially broaden the range of applications for a given optical power (higher sensibility allows wider convolution kernels without additional optical power requirement). Liquid-crystal modulator-based smart pixels are limited to response times in the microsecond regime; as far as response time is concerned, there is a gap between these devices and semiconductor receiver–transmitter-based smart pixels (typically hundreds of megahertz). Although performances of the latter may seem to exceed our needs (response times of approximately 0.1 μ s would just be sufficient), we are obliged to consider semiconductor devices because liquid-crystal microsecond regimes are unacceptable. Table 1 summarizes other possible technologies, along with the corresponding maximum (estimated) simulated annealing rates, strictly on the basis of response time; Subsection 3.C will have a much more detailed investigation of one of the technologies. Our selection of the technology was based on data directly available to us on two types of SPA:

- AsGa photthyristor array chips and hybrid CMOS–self-electro-optic-effect-device (SEED) chips,
- AsGa MQW-SEED²⁰ flip chip bonded to standard Si-CMOS circuits.

1. Photthyristor Pairs

At first sight, an array of AsGa differential photthyristor pairs²¹ seems to fulfill all the requirements (i.e., differential detection and one-bit memory,

thanks to optical bistability) without the help of any additional electronics.²² High sensitivities of 2.6 fJ at 720 nm have been reported, and optical cascaded transmission between arrays has been demonstrated at rates as fast as 50 MHz.²³ We were fortunate to be able to test one integrated circuit composed of 32 strongly positively doped–weakly negatively doped–weakly positively doped–strongly negatively doped photthyristor pairs made available to us by IMEC.²⁴ The matrix was laid out in a checkerboard fashion, with one separate voltage supply for each interlaced matrix (bipartitioning of the processor is hard wired, as in SPIE600, and is therefore automatically ready for two-color operation). Unfortunately, our measurements on this particular device indicated poor overall performances: The average detection threshold turned out close to 1.2 pJ at 783 nm (far from the femtojoule regime), and the average total power emitted per site (at 860 nm) was approximately 31 nW. Our photthyristor array has three major drawbacks: (a) an important mismatch between pairs (four elements were useless in our circuit); (b) low sensibility (approximately 100 times lower) at its own emission wavelength, severely compromising the all-optical interconnection issue; and (c) important divergence (LED-Lambertian source), which leads to cross talk and to insertion loss and imposes the use of additional imaging optics to concentrate their energy (microlens array, for instance). It is possible to switch a photthyristor by use of the light emitted by its neighbors, but, even considering null optical losses, the feedback loop would be very slow: $31 \text{ nW}/(100 \times 1.2 \text{ pJ}) = 26 \text{ ms}$ with the devices that we tested. Hence no more than one serious simulated annealing process (composed of 1000 updating cycles) would be possible every 2.6 s, which does not meet the requirement for a video real-time OSSP demonstrator. We understand from the latest publications on this device (Ref. 22) that better performance was achieved from later devices, but unfortunately, to our knowledge, such devices are no longer being developed.

2. Hybrid CMOS-SEED Technology

Rather than compete with CMOS technology, hybrid CMOS-SEED technology²⁵ seeks to complement it in the interconnection domain by providing additional high-density, high-bandwidth optical inputs and outputs to the existing electronic circuitry. Performances of the flip-chip-bonded GaAs-MQW modulator–detectors depend strongly on the design of its driving (CMOS) circuit. Operation of a two-beam, differential receiver–transmitter circuit at a rate of 620 Mbit/s and a sensitivity of 30 fJ per beam by use of transimpedance front-end amplifiers has been reported.¹⁸ An almost 100% device yield for large arrays (68 × 68) has also been demonstrated.¹⁹ Operating such devices at normal incidence with collimated light simplifies the imaging system (divergence is due only to diffraction). In short, the CMOS-SEED technology currently available offers fairly good performance while considerably reducing the effort in the design of our PE because well-characterized and tested standard transceiver cells are available from many university research groups.

Subsection 3.C is devoted to establishing a theoretical upper bound for the frequency of operation of an OSPP demonstrator that uses a CMOS-SEED SPA device. Two different receiver front-end amplifier circuits will be considered: transimpedance front-end amplifiers and charge-sense amplifiers. The methodology used in the present study is strongly inspired by earlier works of one of the authors.²⁶

C. Optical and Electronic Power Budget

The purpose of this section is to derive an upper bound for the operation frequency of the monolithic OSPP demonstrator (i.e., its feedback rate), depending on the total optical and electrical power available on the system and the maximum thermal dissipation of the SPA.

1. Optical Power Limitations

We will consider quasi-cw mode of operation of the external optical laser source (the reading laser), and we will also assume a maximum power of 1 W. The photodetectors used to convert the optical signal into current or voltage variations are p-i(MQW)-n photodiodes, having at least (an average) responsivity of $S \approx 0.5$ A/W. The voltage output from the electronic circuitry modulates the device into high R_{ON} (typically 70%) or low $R_{OFF} = R_{ON}/C$ reflectivity levels (C is the contrast ratio of the modulator—typically $C \approx 2:1$).

To estimate the overall system performance (i.e., maximum frequency of operation of the feedback loop), we have to determine the lossiest optical path from the source to the detectors of the SPA. Optical losses have been extrapolated on the realistic basis of our previous prototype demonstrator described in Section 2. Maximal losses occur on the route going from the laser source of total power P_{tot} to a particular modulator—e.g., the plus modulator of PE(j),

$$p_m(j) = \eta_{S-PE} P_{tot}, \quad (3)$$




Fig. 8. Light from the (global) laser source to a particular p-i-n photodetector—let us say of PE(i)—goes first through the array illuminator (AI), reflects on every neighboring modulator—i.e., modulators of PE(j), with $j \in N(i)$ —and is diffracted by the computer-generated hologram (CGH) before arriving at the destination.

and then from there to the plus photodiode of PE(i) (a neighboring PE, see Fig. 8), leading, if PE(j) was in the ON state, to the optical power,

$$p_d^{ON}(i) = \eta_{PE-PE} R_{ON} p_m, \quad (4)$$

and if PE(j) was in the OFF state,

$$p_s^{OFF}(i) = \frac{p_d^{ON}(i)}{C} = \eta_{PE-PE} \frac{R_{ON}}{C} p_m. \quad (5)$$

The efficiency of the path from the laser source to a particular PE modulator, η_{S-PE} , depends on the efficiency and the fanning out of the array illuminator. We make a safe assumption of an efficiency of 40% for the array illuminator in the case of a Dammann grating (and perhaps as much as 70% with Talbot holograms²⁷). We then have $\eta_{S-PE} = 40\% / (2 \times N_{PE})$, where N_{PE} is the number of PEs in the SPA and the factor 2 accounts for the two photodiodes per PE. The fan-out efficiency of the interconnecting hologram, η_{PE-PE} , depends on the neighborhood under consideration. The two-level-phase Dammann grating used in our previous demonstrator has shown 16%, 9%, and 7% efficiency, respectively, in the case of a 4-, 8-, and 12-nearest-neighbor-interconnection pattern (leading to a total efficiency of 64%, 72%, and 84%, respectively). The minimum optical power that has to be detected by PE(i) (e.g., when a unique neighboring PE flips its state) is given by

$$\begin{aligned} p_{min} &= 2(p_d^{ON} - p_d^{OFF}) = 2p_d^{ON} \left(1 - \frac{1}{C}\right) \\ &= 2\eta_{S-PE}\eta_{PE-PE}R_{ON} \left(1 - \frac{1}{C}\right) P_{tot}. \end{aligned} \quad (6)$$

For a total optical power $P_{tot} = 1$ W and a state-of-the-art 32 × 16-pixel large array, p_{min} ranges from 52 μ W in the case of a simple mesh-interconnected ar-

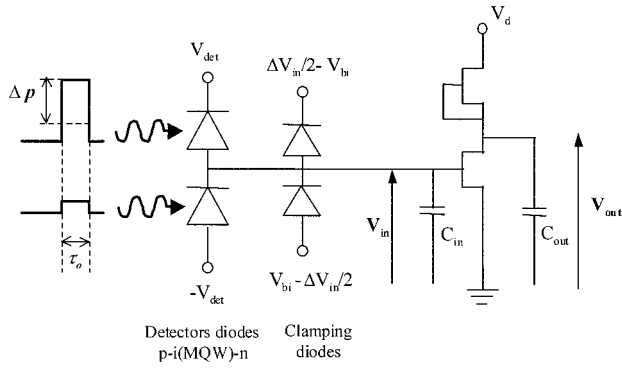


Fig. 9. Single-stage diode-clamped receiver. Pulse duration τ_o is 3 or 6 ns, for a minimum (detectable) differential optical power Δp_{\min} of 52 or 24 μW , respectively.

ray, 29 μW in the case of a 8-nearest-neighbor interconnection, and 24 μW in the case of a wider 12-nearest-neighbor-interconnect topology.

The optoelectronic conversion time T_{IN} (from optical signal to logic-level voltage) depends on the front-end amplifier in use. We have considered here a diode-clamped receiver front-end amplifier. In this case, the energy carried by the optical signal of power p_{\min} has to be sufficient to produce a small voltage swing ΔV_{in} about the threshold voltage of the input transistor of the diode-clamped receiver (see Fig. 9). We consider here a clamped-diode single-stage voltage amplifier with zero output conductance as in Ref. 28. To be detected, the pulse duration τ_o times its power p_{\min} must satisfy the following condition,

$$p_{\min} \tau_o \geq \Delta E_{\min} = \frac{1}{S} C_{\text{in}} \Delta V_{\text{in}}, \quad (7)$$

where C_{in} is the total output capacitance formed by the detectors, the bonding pad, and the input transistor gate ($C_{\text{in}} \approx 300$ fF). If relation (7) is satisfied, then the sum of the conversion and amplification times is given by²⁹

$$T_{\text{IN}} = \frac{3}{4} \frac{C_{\text{in}} \Delta V_{\text{in}}}{S p_{\min}} + 2 C_{\text{out}} \frac{\Delta V_{\text{log}}}{g_m \Delta V_{\text{in}}}, \quad (8)$$

where C_{out} (≈ 120 fF) is the output capacitance of the amplification stage, g_m is its (constant) transconductance ($\approx 10^{-3}$ S), and ΔV_{log} is the typical logic-level voltage swing (≈ 1 V).

In the quasi-cw mode of operation (the duration of the reading pulse is $\tau_o = T_{\text{conv}}$) and during the updating cycle of 1 PE, the output modulators are driven to a high- or low-reflectivity state prior to the illumination by the reading beams (see the chronogram in Fig. 10), so that the output switching time is simply given by

$$T_{\text{OUT}} = 2 \frac{C_{\text{ex}} V_o}{\Delta I_{\text{trans}}}, \quad (9)$$

where V_o is the voltage applied to the modulator (8V), C_{ex} is the output capacitance of the smart pixel, and




Fig. 10. Chronogram of two neighboring PEs in the case of a two-color array and corresponding processing latency times.

ΔI_{trans} is the current flowing to the output capacitance. Typically, $T_{\text{OUT}} \approx 1$ ps.

The total on-chip processing time T_{PE} is given by

$$T_{\text{PE}} = T_{\text{IN}} + T_{\text{elec}} + T_{\text{OUT}}, \quad (10)$$

where T_{elec} is the electronic processing time determined by a simulation by Simulation Packages for Integrated Circuits in Electronics (SPICE) (typically tens of nanoseconds for CMOS technology). The clock period must exceed the total system processing time,

$$T_{\text{PE}} + T_{\text{flight}} \leq T_{\text{clock}} = \frac{1}{F_{\text{clock}}}, \quad (11)$$

where T_{flight} is the time of flight of the optical signals (of the order of 1 ns in our architecture). The updating rate of each color domain (i.e., the feedback rate) is simply $F_C = F_{\text{clock}}/N_C$. An upper bound for the frequency of operation is then given by the following relation,

$$P_{\text{opt}} = \frac{3}{8} \frac{\Delta E_{\min}}{\eta_{\text{S-PE}} \eta_{\text{PE-PE}} R_{\text{ON}}} \left(\frac{1}{F_c} - \frac{N_C}{F_{\text{MAX}}} \right)^{-1} \leq P_{\text{tot}}, \quad (12)$$

where

$$F_{\text{MAX}} = \frac{1}{T_{\text{amp}} + T_{\text{elec}} + T_{\text{OUT}} + T_{\text{vol}}} \approx 83 \text{ MHz} \quad (13)$$

is an upper bound for the clock frequency, independent of the total optical power available on the system and depending exclusively on the electronic circuit and the optical architecture characteristics.

2. Electrical Power Budget and Thermal Limitations

We calculated the electrical and thermal budget, and we found that for a reasonable heat-removal capability of the array Q of the order of 10 W/cm^2 and a




Fig. 11. Optical power requirement versus clock frequency. Solid curve, the 12-nearest-neighbor-optimally-interconnected array; dashed curve, 8-nearest-neighbor-optimally-interconnected array; dotted curve, 4-nearest-neighbor-optimally-interconnected array.

usual electrical power supply of 5 W, the 16×32 SPA under study is only optical power limited.

D. System Performance

Figure 11 shows the laser power consumption for a 32×16 array-based monolithic OSPP, as a function of the frequency (clock) operation, in the case of three different interconnection patterns. Table 2 summarizes the overall performances of the proposed compact prototype, in terms of simulated annealing operations per second. Although the estimation of 150,000 annealing cycles/s (Table 1) was too optimistic, the performance of the new OSPP remains at least 3 orders of magnitude above the video-rate regime. It is important to say here that with our current speckle-generator configuration (with a rotating diffuser) as many as 100,000 spatiotemporal uncorrelated speckle fields can be generated per second, which is large enough for a video-rate serious simulated annealing but is too short if 1000 simulated annealings are required per second. However, it has been shown that with use of an alternative speckle-generator hardware (based on cavitation noise), it is possible to overcome this limitation.¹⁴

To our knowledge, dedicated electronic video-

Table 2. Hybrid 32×16 OSPP Expected Performances^b

Neighborhood	F_{clock} ($P_{\text{tot}} = 1 \text{ W}$) (MHz)	CDR (Mbits/s)	S.A./s
4 (OSPP ₄)	73.3	36.7	36,700
8 (OSPP ₈)	66.8	16.7	16,700
12 (OSPP ₁₂)	64.1	12.8	12,800

^bExpected performances for the CMOS-SEED, OSPP 32×16 pixels. (CDR, smart-pixel optical channel data rate; S.A., simulated annealing).

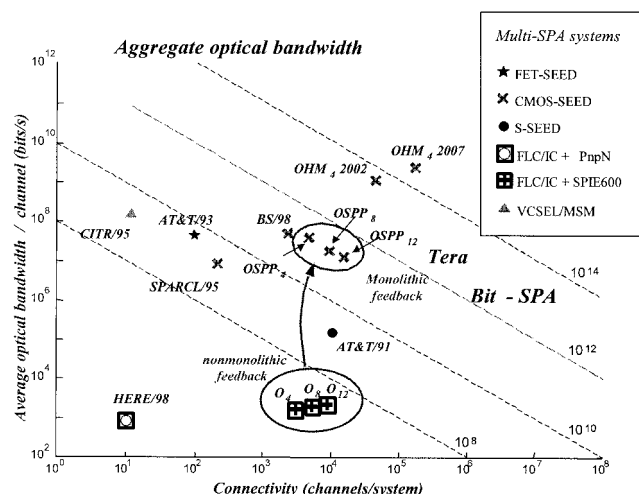


Fig. 12. Optical bandwidth improvement of the OSPP demonstrator and comparison with other SPA-based prototypes (CITR/95,³² HERE/98,³³ SPARCL/95,³⁴ BS/98,³⁵ AT&T/91,³⁶ and AT&T/93³⁷). Optical hardware module 2000/2007 represents perspectives for the CMOS-SEED systems, on the basis of foreseeable improvements in SPA hybrid CMOS-SEED technology³⁸ and by use of a compact, rugged optical hardware module as described in Ref. 39. We assumed a 5-SPA system with $2 \times$ diffractive optical elements that interconnect adjacent arrays (an architecture equivalent to that of the AT&T/93 demonstrator). FLC/IC, integrated-circuit-ferroelectric-liquid-crystal modulator; MSM, metal-semiconductor-metal detector; S-SEED, symmetric-SEED³⁷; FET-SEED, field-effect-transistor-SEED.

processing hardware does not perform better (in terms of processing rate) at statistical regularization-based (Markov random field-based) motion-detection tasks. For the sake of comparison, a general-purpose vision system (powerful and versatile, yet rather expensive and cumbersome), the TRANSVISION system³⁰ (1991), is able to process as many as 2 images/s when configured for Markov random field-based motion detection, and a dedicated (thus cheaper) digital-signal-processing-based hardware, the PCMARKOV³¹ personal computer compatible card designed exclusively for motion detection (1996), is capable of a maximum processing rate of 3 images/s. These systems are able to process standard resolution (256×256 pixels) images instead of hyper-low resolution (24×24 or 32×16 pixels) images, but, as explained before, this is just a different issue, concerning the technologically achievable SPA density. The OSPP is a parallel processor; hence there is no trade-off between SPA density and frequency of operation—the latter depends solely on the PE optical input-output bandwidth. Figure 12 shows the expected aggregate optical bandwidth (or aggregate capacity) achieved with a 32×16 hybrid CMOS-SEED-based OSPP prototype (the aggregate optical bandwidth of the system is defined as the average input-output bits per second flowing within the free-space 2-D optical bus during the annealing process). Figure 12 also draws a comparison between our OSPP prototype and other recent optoelectronic system demonstrators (photonic switching systems, free-




Fig. 13. Interconnection-density improvements for the monolithic OSPP demonstrator.

space optical backplanes, etc.) that use one or more SPAs. Indeed, aggregate capacity can be seen as a general system-performance indicator, resulting—in the case of the OSPP—from the (feedback rate) \times (size of the image) product, divided by the number of color regions of the semishift-invariant single-instruction-multiple-data processor. The feedback rate (or clock frequency F_{clock}) divided by the number of colors N_C corresponds to the individual optical bandwidth of the PE input–output differential channel (which strongly depends on the transducer technology). Note that the connectivity of the system refers to the number of optical input–output channels present on the system.

Figure 13 shows improvement on the prototype compactness. The system interconnection density refers to the total optical interconnection length (i.e., the collected length, taking into account every optical path within the system) divided by the total volume of the demonstrator. For an optoelectronic system based on free-space 2-D interconnected arrays, there is a theoretical upper bound for the interconnection density, given by $D = L/V = 2\pi/\lambda^2$, where λ is the wavelength of light used within the system.⁴⁰ For instance, if we consider the former standard 850-nm wavelength for optical communications, the interconnection density is approximately 8700 km/cm³—at least 7 orders of magnitude above that of present demonstrators. Nevertheless, the goal is that the final size of the whole OSPP module be compatible with standard multichip electronic boards, which is likely to be a realistic aim considering that the random-number rotating diffuser does not need to be more bulky than a conventional heat-removal motorized fan.

4. Conclusion

We have described the first reconfigurable optoelectronic stochastic parallel processor (OSPP) dedicated to real-time (low-level) image tasks. Design and

testing of a nonmonolithic free-space interconnected prototype (a SLM is required in addition to the silicon SPA) have been completed; we successfully demonstrated application of this system to motion detection in gray-level (synthetic) image sequences. Hardware reconfiguration of the OSPP involves rearrangement of the shift-invariant intra-SPA interconnection topology; this was achieved in our first prototype by simple replacement of the optical filter used within the nonmonolithic convolution setup. In fact, this (bulky) optomechanical hardware simulates intra-chip interconnects only through a relatively slow electronic feedback loop, which is responsible for the relatively poor overall system performance. Theoretical modeling shows that the feedback loop can be notably improved by the use of a hybrid CMOS-SEED-based SPA (capable of both optical input and optical output of signals) and a monolithic convolution setup. Indeed, even for the contemplated slower configuration ($N_C = 5$), as many as 12,800 images can be processed per second, outperforming most dedicated electronic image-processing systems—even when these rely on fast deterministic optimization algorithms. We also pointed out that the integration of a complete OSPP module in a conventional electronic board can be considered a realistic goal.

With this first complete OSPP demonstrator we hope to have drawn some of the attention we think this approach deserves to the field of dedicated hardware for video real-time image processing. Theoretical modeling of a hybrid CMOS-SEED-based OSPP should also pave the way for the implementation of more compact and powerful prototypes to come. This study represents one of the possible pathways for the introduction of novel optoelectronic devices into dedicated image-processing systems; nevertheless, application of the OSPP to neural-networklike signal-processing systems can be contemplated if non-shift-invariant and reconfigurable intrachip interconnects were somehow made available for use within the SPA—an issue far beyond the scope of this paper but one that is worth investigating.^{41–43}

Early stages of this study were supported in part by the European Commission under contract ERBCI1*CT93-0004. The main part of the presented study was achieved during the thesis research of A. Cassinelli at the Institut d'Optique. Philippe Lalanne and Donald Prévost actively participated in early stages of this study. The essential role of Eric Belhaire, Francis Devos, Antoine Dupret, Patrick Garda, and Jean-Claude Rodier in designing and testing the SPIE600 chip is gratefully acknowledged. We also thank André Villing for his contribution to the design of the demonstrator electronic driving circuits.

References

1. S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721–741 (1984).

2. S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimisation by simulated annealing," *Science* **220**, 671–680 (1982).
3. P. Chavel, A. Cassinelli, and I. Glaser, "Optoelectronic cellular automata for video real time vision," in *Optics in Computing 2000*, R. A. Lessard and T. Galstian, Proc. SPIE **4089**, 374–381 (2000).
4. A. Huang, "Parallel algorithms for optical digital computers," in *Proceedings of the Tenth International Optical Computing Conference*, S. Horwitz, ed. (IEEE Computer Society and International Optical Conference, New York, 1983), pp. 13–17.
5. J. Taboury, J. M. Wang, P. Chavel, F. Devos, and P. Garda, "Optical cellular processor architecture. 1. Principles," *Appl. Opt.* **9**, 1643–1650 (1988).
6. J. Taboury, J. M. Wang, P. Chavel, and F. Devos, "Optical cellular processor architecture. 2: Illustration and system considerations," *Appl. Opt.* **28**, 3138–3147 (1989).
7. D. E. Knuth, "Random numbers," in *The Art of Computer Programming*, (Addison-Wesley, Reading, Mass., 1969), Vol. 2, Chap. 3, pp. 9–22.
8. Ph. Lalanne, J. C. Rodier, P. Chavel, E. Belhaire, and P. Garda, "Optoelectronic devices for Boltzmann machines and simulated annealing," *Opt. Eng.* **32**, 1904–1914 (1993).
9. A. Dupret, E. Belhaire, J. C. Rodier, Ph. Lalanne, D. Prévost, P. Garda, and P. Chavel, "An optoelectronic CMOS circuit implementing a simulated annealing algorithm," *IEEE J. Solid-State Circuits* **31**, 1046–1050 (1996).
10. P. Chavel, Ph. Lalanne, and J. C. Rodier, "Optoelectronic stochastic processor arrays: demonstration of video rate simulated annealing noise cleaning operation," in *Proceedings of the Conference on Massively Parallel Processing Using Optical Interconnections (MPPOI)* (IEEE Computer Society, Los Alamitos, Calif., 1996), 158–167.
11. P. Lalande and P. Bouthemy, "A statistical approach to the detection and tracking of moving objects in an image sequence," *Signal. Proc.* **5**, 947–950 (1990).
12. P. Lalande, "Détection du mouvement dans les séquences d'images selon une approche markovienne; application à la robotique sous-marine," Ph.D. dissertation (Université de Rennes I, Rennes, France, 1990).
13. A. Caplier, "Modèles markoviens de détection du mouvement dans les séquences d'images. Approches spatio-temporelle et mises en œuvre temps réel," Ph.D. dissertation (Institut National Polytechnique de Grenoble, Grenoble, France, 1995).
14. Ph. Lalanne, E. Belhaire, J. C. Rodier, A. Dupret, P. Garda, and P. Chavel, "Gaussian random number generation by differential detection of speckles," *Opt. Eng.* **34**, 1835–1837 (1995).
15. Ph. Lalanne, D. Prévost, G. Prémont, and P. Chavel, "Optoelectronic implementation of stochastic artificial retinas," *Ann. Phys.* **24**, 125–152 (1999).
16. P. Chavel and Ph. Lalanne, "On parallel algorithms for optical image processors," in *Optical Computing: Proceedings of the International Conference, Heriot-Watt University*, Edinburgh, UK, August 22–25 1994, B. S. Wherrett, ed. (Institute of Physics, London, 1995), pp. 11–16.
17. A. Cassinelli, P. Lalanne, P. Chavel, and I. Glaser, "Demonstration of video-rate optoelectronic parallel processors for noise cleaning in binary images by simulated annealing," in *Optical Computing '98*, P. H. Chavel, D. A. Miller, and H. Thienpont, Proc. SPIE **3490**, 163–166, Bellingham, (1998).
18. T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, and L. M. F. Chirovsky, "Optical receivers for optoelectronic VLSI," *IEEE J. Sel. Top. Quantum Electron.* **2**, 106–116 (1996).
19. A. V. Krishnamoorthy and K. W. Goosen, "Progress in optoelectronic-VLSI smart pixel technology based on GaAs/AlGaAs MQW modulators," *Int. J. Optoelectron.* **11**, 181–198 (1997).
20. D. A. B. Miller, "Quantum well self-electrooptic devices," *Opt. Quantum Electron.* **22**, 61–98 (1990).
21. M. Kuijk, P. Heremans, R. Vounckx, and G. Borghs, "The double heterostructure optical thyristor in optical information processing applications," *Int. J. Opt. Comput.* **2**, 433–436 (1991).
22. G. Prémont, "Photothyristors PnpN et recuit simulé optoelectronique pour la vision bas-niveau," Thèse de l'université de Paris-Sud, Orsay, France (1999).
23. B. Knupfer, M. Kuijk, P. Heremans, R. Vounckx, and G. Borghs, "Cascadable differential PnpN optoelectronic switch operating at 50 Mbit/s with ultrahigh optical input sensitivity," *Electron. Lett.* **31**, 485–486 (1995).
24. P. Heremans, M. Kuijk, R. Vounckx, and G. Borghs, "Fast and sensitive two-terminal double-heterojunction optical thyristors," *Microelectron. Eng.* **19**, 49–52 (1992).
25. K. W. Goosen, J. A. Walker, L. A. D'Asaro, S. P. Hui, and B. Tseng, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.* **7**, 360–362 (1995).
26. M. P. Y. Desmulliez and B. S. Wherrett, "Trends in hybrid-VLSI optoelectronics for optical information processing," in *Optical Interconnections and Computing*, M. Lapp, A. Hasegawa, B. J. Snavey, H. Starle, A. C. Tam, and T. Wilson, eds. (Oxford U. Press, Oxford Series of Optical and Imaging Sciences, London) (to be published).
27. I. Seyd-Darwish, "Etude et réalisation d'un automate cellulaire optoelectronique parallèle," Ph.D. dissertation (de l'Université de Paris Sud, Orsay, France, 1991).
28. M. P. Y. Desmulliez, B. S. Wherrett, A. J. Waddie, J. F. Snowden, and J. A. B. Dines, "Performance analysis of self-electro-optic-effect-device-based (SEED-based) smart-pixel arrays used in data sorting," *Appl. Opt.* **35**, 6397–6415 (1996).
29. A. L. Lentine, L. M. Chirovsky, and T. K. Woodward, "Optical energy considerations for diode-clamped smart pixel optical receivers," *IEEE J. Quantum Electron.* **30**, 1167–1171 (1994).
30. J. P. Dérutin, B. Besserer, T. Tixier, and A. Kikeli, "A parallel vision machine: transvision," in *Proceedings of Computer Architecture for Machine Perception*, B. Zavidovique and P. L. Wendel, eds. (Paris, 16–18 Dec. 1991), pp. 241–251.
31. C. Dumontier, "Etude et mise en œuvre temps réel d'un algorithme de détection du mouvement par approche markovienne," Ph.D. dissertation (Institut National Polytechnique de Grenoble, Grenoble, France, 1996).
32. D. V. Plant, B. Robertson, H. S. Hinton, M. H. Ayliffe, G. C. Boisset, W. Hsiao, D. Kabal, N. H. Kim, Y. S. Liu, M. R. Otazo, D. Pavlasek, A. Z. Shang, J. Simmons, K. Song, D. A. Thompson, and W. M. Robertson, "4 × 4 vertical-cavity surface-emitting laser (VCSEL) and metal-semiconductor-metal (MSM) optical backplane demonstrator system," *Appl. Opt.* **35**, 6365–6368 (1996).
33. R. Buczynski, R. Ortega, T. Szoplik, R. Vounckx, P. Heremans, I. Veretennicoff, and H. Thienpont, "Fast optical thresholding with an array of optical thyristor differential pairs," *J. Opt. A Pure Appl. Opt.* **1**, 267–279 (1999).
34. J. M. Wu, Ch. B. Kuznia, B. Hoanca, Ch. H. Chen, and A. A. Sawchuk, "Demonstration and architectural analysis of complementary metal-oxide semiconductor/multiple-quantum-well smart-pixel array cellular logic processors for single-instruction multiple-data parallel-pipeline processing," *Appl. Opt.* **38**, 2270–2281 (1999).
35. A. C. Walker, T.-Y. Yang, J. Gourlay, J. A. B. Dines, M. G. Forbes, S. M. Prince, D. A. Baillie, D. T. Neilson, R. Williams, L. C. Wilkinson, G. R. Smith, M. P. Y. Desmulliez, G. S. Buller, M. R. Taghizadeh, A. Waddie, I. Underwood, C. R. Stanley, F. Pottier, B. Vögele, and W. Sibbett, "Optoelectronic systems based on InGaAs-complementary-metal-oxide-semiconductor smart-pixel arrays and free-space optical interconnects," *Appl. Opt.* **37**, 2822–2830 (1998).

36. F. B. McCormick, T. J. Cloonan, F. A. P. Tooley, A. L. Lentine, J. M. Sasian, J. L. Brubaker, R. L. Morrison, S. L. Walker, R. J. Crisci, R. A. Novotny, S. J. Hinterlong, H. S. Hinton, and E. Kerbis, "A six-stage digital free-space optical switching network using S-SEEDs," *App. Opt.* **32**, 5153–5171 (1993).
37. F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, R. J. Crisci, R. A. Novotny, H. S. Hinton, and E. Kerbis, "Five-stage free-space optical switching network with field-effect transistor self-electro-optic-effect-device smart-pixel arrays," *App. Opt.* **33**, 1601–1618 (1994).
38. Advanced research initiative in microelectronics, IST programme for Future and Emerging Technologies: Technology roadmap for Optoelectronic Interconnects for Integrated Circuits (MEL-ARI OPTO), published by the European Commission with no author name, (September 1999). <http://www.cordis.lu/esprit/src/melari.htm>
39. M. H. Ayliffe, D. Kabal, F. Lacroix, E. Bernier, P. Khurana, A. G. Kirk, F. A. P. Tooley, and D. V. Plant, "Electrical, thermal and optomechanical packaging of large 2-D optoelectronic device arrays for free-space optical interconnects," *J. Opt. A Pure Appl. Opt.* **1**, 267–271 (1999).
40. H. M. Ozatkas and J. W. Goodman, "Lower bound for the communication volume required for optically interconnected array of points," *J. Opt. Soc. Am. A* **7**, 2100–2106 (1990).
41. S. Lin, A. Grot, J. Luo, and D. Psaltis, "GaAs optoelectronic neuron arrays," *Appl. Opt.* **32**, 1275–1288 (1992).
42. I. Bar-Tana, J. P. Sharpe, D. J. McKnight, and K. M. Johnson, "Smart-pixel spatial light modulator for incorporation in an optoelectronic neural network," *Opt. Lett.* **20**, 303–305 (1995).
43. A. Kirk, T. Tabata, and M. Ishikawa, "Design of an optoelectronic cellular processing system with a reconfigurable holographic interconnect," *Appl. Opt.* **33**, 1629–1639 (1994).