

Available online at www.sciencedirect.comMICROPROCESSORS AND
MICROSYSTEMS

Microprocessors and Microsystems xxx (2006) xxx–xxx

www.elsevier.com/locate/micpro

FPGA architecture for fast parallel computation of co-occurrence matrices

D.K. Iakovidis *, D.E. Maroulis, D.G. Bariamis

Department of Informatics and Telecommunications, University of Athens, Panepistimiopolis, Ilisia, 15784 Athens, Greece

7 Abstract

8 This paper presents a novel architecture for fast parallel computation of co-occurrence matrices in high throughput image analysis
9 applications for which time performance is critical. The architecture was implemented on a Xilinx Virtex-XCV2000E-6 FPGA using
10 VHDL. The symmetry and sparseness of the co-occurrence matrices are exploited to achieve improved processing times, and smaller,
11 flexible area utilization as compared with the state of the art. The performance of the proposed architecture is evaluated using input
12 images of various dimensions, in comparison with an optimized software implementation running on a conventional general purpose
13 processor. Simulations of the architecture on contemporary FPGA devices show that it can deliver a speedup of two orders of magnitude
14 over software.

15 © 2006 Elsevier B.V. All rights reserved.

16 *Keywords:* Image analysis; Co-occurrence matrix; FPGA; Texture

18 1. Introduction

19 The co-occurrence matrix is a powerful statistical tool
20 which has proved its usefulness in a variety of image anal-
21 ysis applications, including biomedical [1,2], remote sensing
22 [3], quality control [4] and industrial defect detection sys-
23 tems [5]. It captures second-order grey-level information,
24 which is mostly related to human perception and the dis-
25 crimination of textures.

26 Although the computational complexity of the co-occur-
27 rence matrix for an image of $N \times N$ dimensions is only
28 $O(N^2)$, the processing power requirements for the compu-
29 tation of multiple co-occurrence matrices per time unit
30 can be prohibiting for the analysis of large image streams,
31 using software co-occurrence matrix implementations on
32 general purpose processors. Such demanding applications
33 include video analysis [1,6], content-based image retrieval
34 [7], real-time industrial applications [5] and high-resolution
35 multispectral image analysis [2].

Field Programmable Gate Arrays (FPGAs) are low cost, 36
reconfigurable high density gate arrays capable of perform- 37
ing many complex computations in parallel while hosted by 38
conventional computer hardware [8]. Their features enable 39
the development of a hardware system dedicated to per- 40
forming fast co-occurrence matrix computations, thus 41
meeting the requirements of real-time image analysis appli- 42
cations. On the other hand, the Very Large Scale Integra- 43
tion (VLSI) architectures could be considered as 44
competitive alternatives [9]. However, they are not recon- 45
figurable and they involve high development cost and 46
time-consuming development procedures. 47

Within the first FPGA architectures, dedicated to 48
co-occurrence matrix computations, was the one presented 49
in [5,6] for the computation of two statistical measures of 50
the co-occurrence matrix. However, these measures were 51
being approximated, without needing to compute the 52
matrix itself. In a later work, Tahir et al. [2] developed 53
an FPGA architecture for the computation of 16 co-occur- 54
rence matrices in parallel. The implementation consider- 55
ations include symmetry, but do not include sparseness. 56
As a result, a large FPGA area is utilized even for small 57
input images. 58

* Corresponding author. Tel.: +30 210 7275317; fax: +30 210 7275333.
E-mail address: rtsimage@di.uoa.gr (D.K. Iakovidis).

59 In this paper, we present a novel FPGA architecture for
60 parallel computation of 16 co-occurrence matrices that
61 exploits both their symmetry and sparseness to achieve
62 improved processing times and smaller, flexible area
63 utilization.

64 2. The co-occurrence matrix

65 The co-occurrence matrix of an $N \times N$ -pixel image I ,
66 comprises of the probabilities $P_{d,\theta}(i, j)$ of the transitions
67 from a grey-level i to a grey-level j in a given direction θ
68 at a given intersample spacing d

$$69 P_{d,\theta}(i, j) = \frac{C_{d,\theta}(i, j)}{\sum_{i=1}^{N_g} \sum_{j=1}^{N_g} C_{d,\theta}(i, j)}, \quad (1)$$

71 where $C_{d,\theta}(i, j) = \#\{(m, n), (u, v) \in N \times N: f(m, n) = j,$
72 $f(u, v) = i, |(m, n) - (u, v)| = d, \angle((m, n), (u, v)) = \theta\}$, $\#$
73 denotes the number of elements in the set, $f(m, n)$ and
74 $f(u, v)$ correspond to the grey-levels of the pixel located
75 at (m, n) and (u, v) , respectively, and N_g is the total number
76 of grey-levels in the image [11]. In accordance with [2], we
77 choose $N_g = 32$ (5-bit representation).

78 The co-occurrence matrix can be regarded symmetric
79 if the distribution between opposite directions is ignored.
80 The symmetric co-occurrence matrix is derived as $P_{d,\theta}$
81 $(i, j) = (P_{d,\theta}(i, j) + P_{d,\theta}(i, j)^T)/2$, where symbol T denotes
82 the transpose matrix. Therefore, the co-occurrence matrix
83 can be represented as a triangular structure without any
84 information loss, and θ is chosen within the range of 0°
85 to 180° . Common choices of θ include 0° , 45° , 90° and
86 135° [1,2,6,12]. Moreover, depending on the image dimensions,
87 the co-occurrence matrix can be very sparse, as the
88 number of grey-level transitions for any given distance and
89 direction, is bounded by the number of image pixels.

90 3. Architecture

91 The presented architecture was developed in Very high
92 speed integrated circuits Hardware Description Language
93 (VHDL). It was implemented on a Xilinx Virtex-
94 XCV2000E-6 FPGA, which is characterized by 80×120
95 Configurable Logic Blocks (CLBs) providing 19,200 slices
96 (1 CLB = 2 slices). The device includes 160 256×16 -bit
97 Block RAMs and can support up to 600 kbit of distributed
98 RAM. The host board, Celoxica RC-1000 has four 2 MB
99 static RAM banks. The RAM banks can be accessed by
100 the FPGA and the host computer independently, whereas
101 simultaneous access is prohibited by the board's arbitra-
102 tion and isolation circuits.

103 An overview of the proposed FPGA architecture is illus-
104 trated in Fig. 1. The FPGA includes a control unit, four
105 memory controllers (one for each memory bank) and 16
106 Co-occurrence Matrix Computation Units (CMCUs). Up
107 to four input images of N_g grey-levels can be loaded in par-
108 allel to the available RAM banks. In accordance with [2], a
109 5-bit grey-level representation was used, i.e., $N_g = 32$.
110 However, in [2] each image is loaded into a corresponding
111 RAM bank using a 5-bit per pixel representation whereas
112 in the proposed architecture a 25-bit per pixel representa-
113 tion is used. Each pixel is represented by a vector $\vec{a} = [a_p,$
114 $a_0, a_{45}, a_{90}, a_{135}]$ that comprises of five 5-bit components,
115 namely, the grey-level a_p of the pixel and the grey-levels
116 a_0, a_{45}, a_{90} and a_{135} of its neighboring pixels at 0° ,
117 45° , 90° and 135° directions.

118 All FPGA functions are coordinated by the control unit
119 which generates synchronization signals for the memory
120 controllers and the CMCUs. The control unit also
121 handles communication with the host, by exchanging
122 control and status bytes, and requesting or releasing the
123 ownership of the memory banks. Each CMCU is used

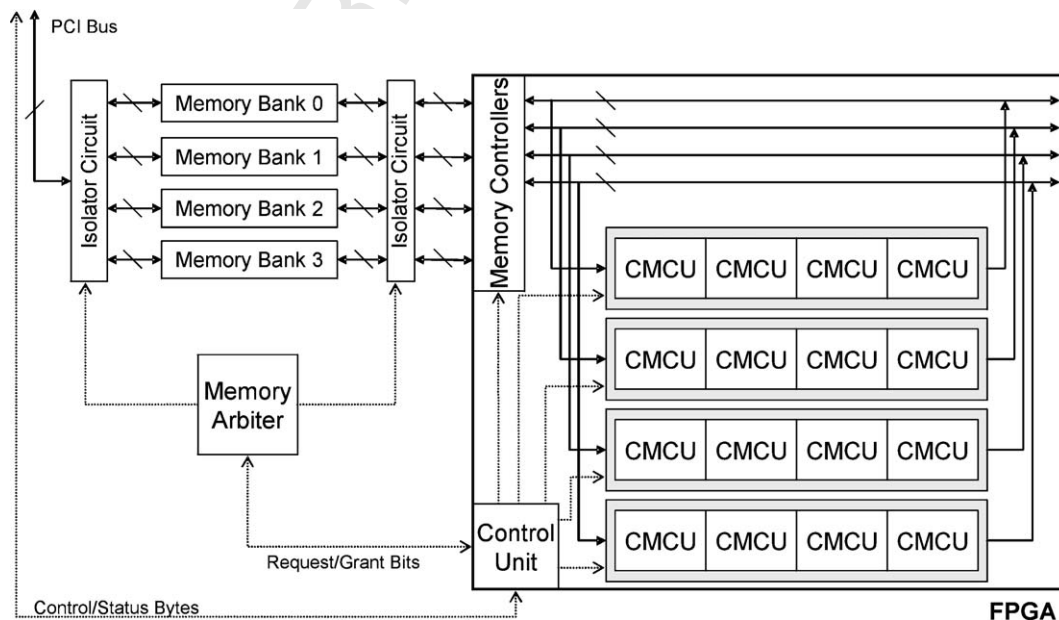


Fig. 1. Overview of the FPGA architecture.

124 for the computation of the co-occurrence matrix of an
125 image for a particular direction and distance.

126 3.1. Co-occurrence Matrix Computation Units

127 Three main objectives have been determined upon the
128 requirements of the proposed application, for the develop-
129 ment of a CMCU: (a) small FPGA area utilization to allow
130 for a potential expansion of the proposed architecture, (b)
131 high throughput of one result per cycle to achieve a high
132 per-clock performance and (c) low design complexity that
133 will contribute to achieving high operation frequency. To
134 meet these objectives we have considered various alterna-
135 tives for the implementation of the CMCUs. These include
136 the utilization of the existent FPGA BlockRAM arrays, the
137 implementation of standard sparse array structures that
138 store pairs of indices and values, and the implementation
139 of set-associative sparse arrays. The BlockRAM arrays
140 and the standard sparse array structures would not suffice
141 to meet all three objectives. The BlockRAM arrays would
142 lead to a larger area utilization compared with the sparse
143 implementations. The standard sparse arrays would result
144 in a lower throughput compared with the other two imple-
145 mentations, as the cycles needed to traverse the indices of
146 the array are proportional to its length. In comparison,
147 the set-associative arrays could be considered as a more
148 flexible alternative that can be effectively used for achieving
149 all our three objectives.

150 Fig. 2 illustrates a CMCU as implemented by means of
151 an n -way set-associative array of N_c cells and auxiliary cir-
152 cuitry which include n comparators, a n -to- $\log_2 n$ priority
153 encoder and an adder.

154 The set-associative arrays can be utilized for efficient
155 storage and retrieval of sparse matrices, ensuring a
156 throughput of one access per cycle with a latency of
157 four cycles. An n -way set-associative array consists of
158 n independent tag arrays ($\text{tags}_0 - \text{tags}_{n-1}$) as illustrated
159 in Fig. 2. The tag-arrays are implemented in the
160 FPGA's distributed RAM and each of them consists
161 of N_c/n cells. The set-associative array uniquely maps
162 an input pair of 5-bit grey-level intensities (i, j) into
163 an address of the N_c -cell data array. The data arrays
164 are implemented using FPGA Block RAMs, each of
165 which can hold up to 256 co-occurrence matrix ele-
166 ments. The data array cells contain the number of
167 occurrences of the respective (i, j) pairs. Each of these
168 pairs are represented by a single 10-bit integer k , result-
169 ing from the concatenation of i and j . This integer can
170 be considered to consist of two parts: the first is called
171 $\text{set}(i, j)$ and comprises of the $\log_2(N_c/n)$ least significant
172 bits of k , whereas the second part is called $\text{tag}(i, j)$
173 and comprises of the $10 - \log_2(N_c/n)$ most significant bits of k .
174 The increment of a data array cell that corresponds to
175 an input pair (i, j) is implemented in four pipeline
176 stages:

177 Stage 1. The tag array cells located in the $\text{set}(i, j)$ row are
178 retrieved and stored in temporary registers.

179 Stage 2. The values of the temporary registers values are
180 compared with $\text{tag}(i, j)$.

181 a. If a match is found the column number of the
182 matching tag is written in the offset register.

183 b. If there are not any matches the $\text{tag}(i, j)$ is
184 stored in the tags array, at the first available
185 cell of the $\text{set}(i, j)$ row.

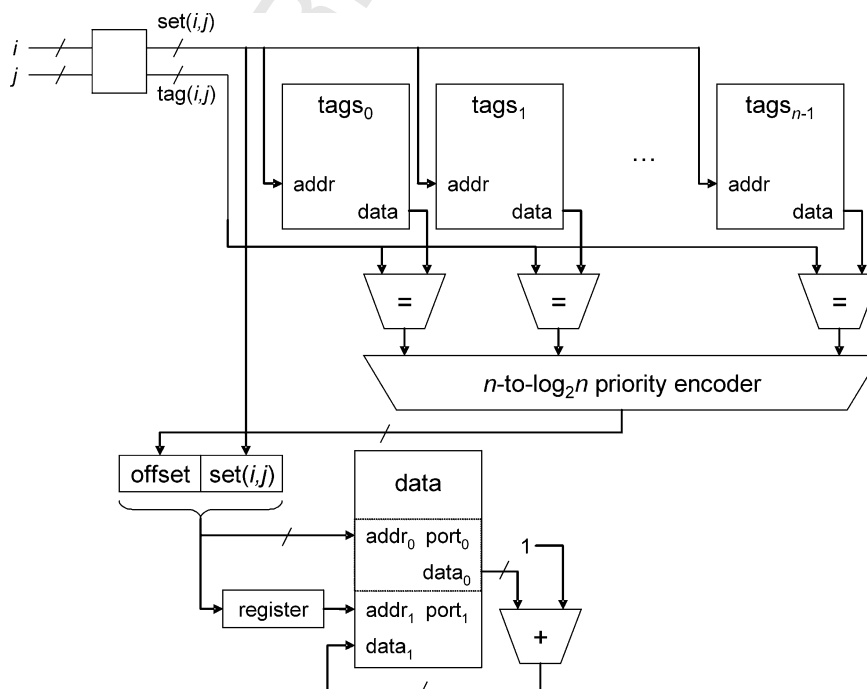


Fig. 2. The co-occurrence matrix computation unit (CMCU).

4

D.K. Iakovidis et al. / Microprocessors and Microsystems xxx (2006) xxx–xxx

186
 187 Stage 3. The contents of both the offset register and
 188 set (i, j) form an address a . The data array ele-
 189 ment stored in a is read.
 190 Stage 4. The value read in the previous cycle increases by
 191 one and it is written back to a .
 192
 193 After all input pairs are read and processed the data
 194 array will contain the co-occurrence matrix of the input
 195 image.

196 4. Results

197 Experiments focusing to the evaluation of the time per-
 198 formance and the area utilization of the proposed architec-
 199 ture were performed using standard texture images from
 200 the Brodatz album of 16×16 , 32×32 , 64×64 , 128×128 ,
 201 256×256 and 512×512 -pixel dimensions (Fig. 3) [13].

202 Given a triangular co-occurrence matrix of $N_g = 32$, the
 203 number of pixel pairs that can be considered for its compu-
 204 tation in the case of a 16×16 -pixel input image, is smaller
 205 than the total number of co-occurrence matrix elements,
 206 and reaches the number of all image pixels. Therefore,
 207 the co-occurrence matrix will be sparse and N_c is set to a
 208 maximum possible value of $16 \times 16 = 256$. In the case of
 209 a 32×32 -pixel or a larger input image, the co-occurrence
 210 matrix is not considered sparse as the number of all possi-
 211 ble pixel pairs that can be considered for its computation is
 212 larger than the total number of its elements (i.e., 528).
 213 Therefore, N_c is set to 528. It is worth noting that the effect
 214 of sparseness in area utilization is amplified and becomes
 215 more useful as N_g increases. For example, if N_g was set
 216 at 64 or at 128 grey-levels, the co-occurrence matrix
 217 could be considered sparse for images up to 32×32 or




Fig. 3. Texture image D9 from the Brodatz album cropped at various sizes.

64 \times 64-pixel dimensions, respectively. By following a grid
 search approach for the determination of n , it was found
 that the 16-way set-associative arrays ($n = 16$) result in
 the optimal tradeoff between time performance and area
 utilization.

The proposed architecture, as implemented on the
 Xilinx Virtex-XCV2000E-6 FPGA, operates at 38.4 MHz
 and utilizes only 39% of the FPGA area for 16×16 input
 images, where the sparseness of the co-occurrence matrices
 is exploited. The use of larger input images results in
 approximately the same operating frequency reaching
 38.2 MHz and a larger area utilization of 45%. In compar-
 ison, the architecture proposed in [2] operates at 50 MHz
 and utilizes a larger area percentage (59%) on an FPGA
 of the same type, for the same N_g , regardless of the image
 dimensions. The time performance reported in [2] for the
 computation of a total of 64 co-occurrence matrices in
 512×512 -pixel 16-band multispectral images was
 $6.3 \times 10^5 \mu\text{s}$. For the same computations, the proposed
 architecture requires $28.041 \times 4 = 1.1 \times 10^5 \mu\text{s}$ (Table 1),
 which can be interpreted in approximately 500% reduction
 of the processing time. This improvement in time perfor-
 mance is mainly attributed to the use of vectors \bar{a} for the
 retrieval of five pixels in one cycle instead of the five cycles
 required in the per pixel retrieval used in [2].

Even though the implementation of the proposed archi-
 tecture was based on the Xilinx Virtex-XCV2000E-6
 FPGA, we run several simulations on state of the art
 FPGA devices, such as Virtex-XCV2000E-8 (19200 slices),
 Virtex2-XC2V6000-6 (33792 slices) and Spartan3-
 XC3S4000-5 (27648 slices). The processing times achieved
 for the computation of 16 co-occurrence matrices in hard-
 ware and software, respectively, are presented in Table 1.

Software processing times were measured using an
 MMX optimized software implementation developed in
 C programming language and executed on an Athlon
 XP2700+ processor. The optimizations were based on the
 guidelines suggested by Intel and AMD [14,15]. These
 include contiguous arrays allocation for improving CPU
 caching performance, system call overhead reduction by
 allocation of static arrays for data used iteratively within
 the program, usage of efficient C library functions such
 as `memset()` and `memcpy()`, and vectorization of several
 functions using the MMX instruction set [16]. Additional
 code fine-tuning includes code rearrangement for breaking
 dependencies in tight loops, dereferencing of commonly
 used pointers and reduction of the function call overhead
 using inline functions.

The results reveal the superior performance of the hard-
 ware implementations of the proposed architecture over
 the software implementation. The speedup factors achieved
 in hardware vary depending on the FPGA model used. The
 minimum speedup is approximately 20 in the case of
 XCV2000E-6 for 512×512 images, whereas it exceeds
 100 in the case of XC2V6000-6 for 16×16 images. The var-
 iance in speedup is mainly attributed to the different fre-
 quencies of the various FPGA models and does not

Table 1

Processing times (μs) achieved for various input image dimensions using various FPGA devices, and software

Implementation	Frequency (MHz)	Image dimensions (pixels)					
		16 \times 16	32 \times 32	64 \times 64	128 \times 128	256 \times 256	512 \times 512
<i>Processing times (μs)</i>							
Hardware							
XCV2000E-6	38	30	113	442	1756	7013	28,041
XCV2000E-8	51	22	83	323	1283	5123	20,483
XC3S4000-5	72	15	59	230	915	3653	14,606
XC2V6000-6	83	13	51	198	788	3149	12,590
Software							
Athlon XP 2700+	2167	1371	3247	10,018	36,320	143,600	562,080

275 correlate with the sparseness of the co-occurrence matrix,
 276 which mainly affects the area utilization. In Table 1 it can
 277 be observed that the increase in processing times as the
 278 image dimensions increase by two is not exactly divided
 279 by four, as it would have been expected by the quadrupli-
 280 cation of the image pixels. This is explained by the constant
 281 time period spent for resetting the FPGA circuit.

282 5. Conclusions

283 We presented a novel FPGA architecture which is
 284 capable of performing fast parallel co-occurrence matrix
 285 computations in grey-level images. It performs better
 286 than the state of the art FPGA architecture presented
 287 in [2]. The proposed architecture and the architecture
 288 in [2] have two main differences pinpointed to the input
 289 data format and the co-occurrence matrix representation.
 290 The vector representation of the input image pixels and
 291 the use of set-associative arrays for the sparse representa-
 292 tion of the co-occurrence matrix result in a higher time
 293 performance and smaller area utilization respectively.
 294 Its advantageous time performance compared with the
 295 architecture in [2] and with an optimized software imple-
 296 mentation for general purpose processors, makes it
 297 appealing for use in high throughput applications. More-
 298 over, the smaller FPGA area it utilizes, allows for the
 299 exploitation of the remaining area for other tasks, such
 300 as the computation of co-occurrence matrix features
 301 [11], or the computation of more co-occurrence matrices
 302 in parallel, if the host board is equipped with more
 303 RAM banks.

304 It is worth noting that the computation of co-occurrence
 305 matrices in conjunction with feature extraction in the same
 306 FPGA design still remains a challenge. In [2], two different
 307 FPGA designs, one for the computation of co-occurrence
 308 matrices and one for the feature extraction, are inter-
 309 changeably configured on a single FPGA.

310 Within our future perspectives are the extension of the
 311 current architecture for efficient on-chip extraction of mul-
 312 tiple textural features from grey-level and colour images, in
 313 the same FPGA design, and its integration in a complete,
 314 hardware/software system with real-time video analysis
 315 capabilities.

6. Uncited reference

Ref. [10].


Acknowledgments

This research was funded by the Operational Program for
 Education and Vocational Training (EPEAEK II) under the
 framework of the project “Pythagoras – Support of Univer-
 sity Research Groups” co-funded by 75% from the Europe-
 an Social Fund and by 25% from national funds.


References

- [1] S.A. Karkanis, D.K. Iakovidis, D.E. Maroulis, D.A. Karras, M. Tzivras, Computer aided tumor detection in endoscopic video using color wavelet features, *IEEE Trans. Inf. Technol. Biomed.* 7 (2003) 141–152.
- [2] M.A. Tahir, A. Bouridane, F. Kurugollu, An FPGA based coprocessor for GLCM and haralick texture features and their application in prostate cancer classification, *Analog Integr. Circ. Signal Process.* 43 (2005) 205–215.
- [3] A. Baraldi, F. Parmiggiani, An investigation of the textural characteristics associated with gray level cooccurrence matrix statistical parameters, *IEEE Trans. Geosci. Remote Sens.* 33 (2) (1995) 293–304.
- [4] K. Shiranita, T. Miyajima, R. Takiyama, Determination of meat quality by texture analysis, *Pattern Recogn. Lett.* 19 (1998) 1319–1324.
- [5] J. Iivarinen, K. Heikkinen, J. Rauhamaa, P. Vuorimaa, A. Visa, A defect detection scheme for web surface inspection, *Int. J. Pattern Recogn. Artif. Intell.* (2000) 735–755.
- [6] D.K. Iakovidis, D.E. Maroulis, S.A. Karkanis, I.N. Flaounas, Color texture recognition in video sequences using wavelet covariance features and support vector machines, in: *Proceedings of 29th EURO-MICRO*, September 2003, Antalya, Turkey, 2003, pp. 199–204.
- [7] C.-H. Wei, C.-T. Li, R. Wilson, A content-based approach to medical image database retrieval, in: Z. Ma (Ed.), *Database Modeling for Industrial Data Management: Emerging Technologies and Applications*, Idea Group Publishing, 2005.
- [8] T.A. York, Survey of field programmable logic devices, *Microprocess. Microsyst.* 17 (7) (1993) 371–381.
- [9] M. Ba, D. Degrugillier, C. Berrou, Digital VLSI using parallel architecture for co-occurrence matrix determination, in: *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, vol. 4, 1989, pp. 2556–2559.
- [10] K. Heikkinen, P. Vuorimaa, Computation of two texture features in hardware, in: *Proceedings of the Tenth International Conference on Image Analysis and Processing*, September 1999, Venice, Italy, 1999, pp. 125–129.

- 359 [11] R.M. Haralick, K. Shanmugam, I. Dinstein, Textural features for
360 image classification, *IEEE Trans. Syst. Man Cybern.* 3 (1973) 610–621.
361 [12] S. Theodoridis, K. Koutroumbas, *Pattern Recognition*, Academic
362 Press, San Diego, 1999.
363 [13] P. Brodatz, *Textures: A Photographic Album for Artists and*
364 *Designers*, Dover Publications, New York, 1966.
365 [14] IA-32 Intel Architecture Optimization Reference Manual, Intel
366 Corp., 2004.
367 [15] Athlon Processor x86 Code Optimization Guide, AMD Inc., 2002.
368 [16] Intel Pentium 4 Processor Optimization Reference Manual, Intel
369 Corporation, 1999–2000.



Dimitris K. Iakovidis received his B.Sc. degree in Physics from the University of Athens, Greece. In April 2001, he received his M.Sc. degree in Cybernetics and in February 2004 his Ph.D. degree in Computer Science from the Department of Informatics and Telecommunications, University of Athens, Greece. Currently he is working as a Research Fellow in the same Department and he has co-authored more than 30 papers on image analysis, systems, and biomedical applications. Also he is a regular reviewer for many international journals. His research interests include image analysis, system development, pattern recognition and bioinformatics.



Dimitris E. Maroulis received the B.Sc. degree in Physics, the M.Sc. degree in radioelectricity, the M.Sc. in electronic automation and the Ph.D. degree in Computer Science, all from the University of Athens, Greece, in 1973, 1977, 1980 and 1990, respectively. In 1979, he was appointed Assistant in the Department of Physics, in 1991 he was elected Lecturer and in 1994 he was elected Assistant Professor, in the Department of Informatics of the same university. He is currently working in the above Department in teaching and research activities, including Projects with European Community. His main areas of activity include data acquisition systems, real-time systems, signal processing and biomedical systems.



Dimitris G. Bariamis is a student of the Department of Informatics and Telecommunications, pursuing a Ph.D. degree in hardware architecture design. His research interests include FPGA design, and software programming and optimization techniques.

385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409

UNCORRECTED