

Visual Sensor Networks for Infomobility

M. Magrini^a, D. Moroni^a, C. Nastasi^b, P. Pagano^b, M. Petracca^b,
G. Pieri^a, C. Salvadori^b, and O. Salvetti^a

^a *Institute of Information Science and Technologies (ISTI), Italian
National Research Council (CNR), Pisa, Italy*

^b *ReTis Laboratory (CEIICP), Scuola Superiore Sant'Anna, Pisa, Italy*

*e-mail: massimo.magrini@isti.cnr.it, davide.moroni@isti.cnr.it, christian.nastasi@sssup.it, paolo.pagano@sssup.it,
matteo.petracca@sssup.it, gabriele.pieri@isti.cnr.it, claudio.salvadori@assup.it, ovidio.salvetti@isti.cnr.it*

Abstract—The wide availability of embedded sensor platforms and low-cost cameras—together with the developments in wireless communication—make it now possible the conception of pervasive intelligent systems based on vision. Such systems may be understood as distributed and collaborative sensor networks, able to produce, aggregate and process images in order to understand the observed scene and communicate the relevant information found about it. In this paper, we investigate the peculiarities of visual sensor networks with respect to standard vision systems and we identify possible strategies to accomplish image processing and analysis tasks over them. Although the rather strong constraints in computational and transmission power of embedded platforms that may prevent the use of state of the art computer vision and pattern recognition methods, we argue that multi-node processing methods may be envisaged to decompose a complex task into a hierarchy of computationally simpler problems to be solved over the nodes of the network. These ideas are illustrated by describing an application of visual sensor network to infomobility. In particular, we consider an experimental setting in which several views of a parking lot are acquired by the sensor nodes in the network. By integrating the various views, the network is capable to provide a description of the scene in terms of the available spaces in the parking lot.

Keywords: image mining, sensor networks, infomobility, object detection, change detection.

1. INTRODUCTION

In the last years, the wide availability of embedded systems and low-cost camera sensors—together with the developments in wireless communication—has made it possible to conceive sensor-based pervasive intelligent systems centered on image data [1]. Such visual Wireless Sensor Networks (WSNs), employing a great number of low power camera nodes, may support a large class of novel vision-based applications thanks to the great informative power of imaging. For example, visual WSNs may be used to monitor in real-time the crowd in shopping mall, airports and stadiums. By mining the scene, the network may detect anomalous and potentially dangerous events [2]. Similarly, visual WSNs may be used for environmental monitoring, for the remote control of elderly patients in e-Health [3] and for human gesture analysis in ambient intelligence applications [4].

In brief, the key feature of visual WSNs is the combination of the versatility and independence from a physical infrastructure typical of general WSNs with the richness of information that can be

gained through imaging techniques, computer vision and image understanding. In this sense, a visual WSN may be understood as a distributed and collaborative sensor network, able to produce, aggregate and process images in order to mine the observed scene and communicate the relevant information found about it. In particular, each acquisition node is capable to acquire a view of the scene to be mined. Then, in cooperation with each other, the various devices in the network are able to process and aggregate the acquired views in order to predicate something about the scene.

A successful design and development of such a system cannot be achieved without suitable solutions to the involved computer vision problems. Although the computer vision problems may be still decomposed into basic computational tasks (such as feature extraction, object detection and object recognition), in the context of visual WSNs, it is not directly possible to use all the methods that have been developed to solve such tasks and that are already available in the specific literature [5]. Indeed, since WSNs usually require a large number of sensors, possibly scattered over a large area, the unit cost of each device should be small to make the technology affordable. Therefore, cost constraints limit the computational and transmission power of sensor nodes as well as the fidelity of acquired images,

with consequences on the employable computer vision algorithms. Since untethered sensors are usually supplied by batteries, it is also necessary to analyze carefully the tradeoffs between quality of processing and energy consumption, in order to avoid too frequent battery replacement.

In light of these considerations, performing computer vision tasks over WSNs is somewhat challenging from a technological viewpoint. However, probably, the change in perspective offered by the possibility to perform pervasive computing and by the ductility in organizing the network topology greatly compensates for the current technical limits of visual WSNs. In addition, the topic is intriguing and may open the way also to more theoretical investigations. In this paper, we investigate multi-node processing methods that may be envisaged to decompose a complex vision task into a hierarchy of computationally simpler problems, to be solved over the nodes of the network. Besides computational advantages, such hierarchical decision making approach may lead to more robust and fault-tolerant results. We illustrate these ideas by describing an application of visual sensor network to infomobility. To this end, we consider an experimental setting in which several views of a parking lot are acquired by the sensor nodes in the network. By integrating the various views, the network is capable to provide a description of the scene in terms of the available spaces in the parking lot. This paper extends and complements the conference paper [6].

2. BACKGROUND

Embedded vision platform first appeared in connection with video-surveillance and robotics. Thanks to the drop in technology costs, nowadays, their application range has become wider so as to encompass several sectors of public and private life and, in particular, infomobility. We first discuss existing embedded vision platforms and basic features of visual WSNs (Section 2.1); then we survey the specificities of performing image analysis over WSNs, introducing in particular the multi-view vision problem (Section 2.2).

2.1. Visual Sensor Networks

Following the trends in low-power processing, wireless networking and distributed sensing, visual WSNs are experiencing a period of great interest, as shown by the recent scientific production (see e.g. [7]). A visual WSN consists of tiny visual sensor nodes called camera nodes, which integrate the image sensor, the embedded processor and a wireless RF transceiver [1]. The large number of camera nodes forms a distributed system where the camera nodes are able to process image data locally (*in-node processing*) and to extract relevant information, to collaborate with other-cameras—even autonomously—on the appli-

cation specific task, and to provide the system user with information-rich descriptions of the captured scene.

In the last years, several research projects produced prototypes of embedded vision platforms which may be deployed to build a visual WSN. Among the first experiences, Panoptes project [8] aimed at developing a scalable architecture for video sensor networking applications. The key features of Panoptes sensor are a relatively low-power and high-quality video capturing device, a prioritizing buffer management algorithm to save power and a bit-mapping algorithm for the efficient querying and retrieval of video data. Nevertheless the size of the sensor, its power consumption, its relatively high computational power and storage capabilities makes Panoptes sensor more akin to smart high-level cameras than to untethered low-power low-fidelity sensors.

The Cyclops project [9] provided another representative smart camera for sensor networks. The camera nodes is equipped with a low-performance ATmega128 8-bit RISC microcontroller. From the storage memory point of view the system is very constrained, with 128 KB of FLASH program memory and only 4 KB of SRAM data memory. The CMOS sensor supports three image formats of 8-bit monochrome, 24-bit RGB color, and 16-bit YCbCr color at CIF resolution (352×288). In the Cyclops board, the camera module contains a complete image processing pipeline for performing demosaicing, image size scaling, color correction, tone correction and color space conversion.

In the MeshEye project [10], an energy-efficient smart camera mote architecture was designed, mainly with intelligent surveillance as target application. MeshEye mote has an interesting special vision system based on a stereo configuration of two low-resolution low-power cameras, coupled with a high resolution color camera. In particular, the stereo-vision system continuously determines position, range, and size of moving objects entering its fields of view. This information triggers the color camera to acquire the high-resolution image subwindow containing the object of interest, which can then be efficiently processed.

Another interesting example of low-cost embedded vision system is represented by the CMUcam3 [11], developed at the Carnegie Mellon University. More precisely, the CMUcam3 is the third generation of the CMUcam series, which has been specially-designed to provide an open-source, flexible and easy development platform with robotics and surveillance as target applications. The hardware platform is more powerful with respect to its predecessors and may be used to equip low-cost embedded system with simple vision capabilities, so as to obtain smart sensors. The hardware platform is constituted by a CMOS camera, an ARM7 processor and a slot for MMC cards. Standard

Table 1. Comparison of embedded vision platforms

Platform	Sensor	CPU	Application
Panoples [8]	External USB camera	XScale PXA255	Video streaming, surveillance
Cyclops [9]	Color CMOS 352 × 288	ATmega 128	Collaborative object tracking
Mesheye [10]	2 low-resolution sensor	ARM7	Distributed surveillance
CMUcam3 [11]	Color CMOS 352 × 288	ARM7	Local image analysis
CITRIC [12]	Color CMOS 1280 × 1024	XScale PXA255	Compression, tracking, localization

RF transceiver (e.g., TELOS mote) can be easily integrated.

More recently, the CITRIC platform [12] integrates in one device a camera sensor, a CPU (with frequency scalable up to 624 MHz), a 16 MB FLASH memory and a 64 MB RAM. Such a device, once equipped with a standard RF transceiver, is suitable for the development of visual WSN. The design of the CITRIC system allows to perform moderate image processing task *in-network*, that is along the nodes of the network. In this way, there are less stringent issues regarding transmission bandwidth than with respect to centralized solutions. Such results have been illustrated by 3 sample applications, namely (i) image compression, (ii) object tracking by means of background subtraction and iii) self-localization of the camera nodes in the network. Table 1 summarizes and compares the presented embedded vision platforms.

From a software perspective, it would be beneficial to implement distributed intelligent systems on top of real-time kernels and popular network protocols like IEEE 802.15.4. In this respect, TinyOS [13], the *de facto* standard Operating System (OS) for WSNs, uses a stack shared among the processes and no heap. Each instance of the tasks runs until the end of the code unless it is preempted by an Interrupt Service Routine (ISR)—an event handler activated by an event occurrence. ISRs can in turn spawn a new task or call another function or command. The usage of such a mono-stack and the lack of inter-task preemption mechanisms constrains TinyOS to schedule the tasks on the basis of the activation times, i.e., adopting a First In First Out (FIFO) policy. Other WSN OSs (e.g., ERIKA Enterprise [14]) allow task preemption and real-time priority-driven scheduling. Tasks can block on certain events and can be woken up at a certain point of time (activation). For instance, in distributed vision, events are generated at the arrival of a message from other nodes in the network, or when an image, formed in the camera sensor, is fully transferred to primary memory. Suspended tasks can be un-suspended and enqueued into the OS active list when appropriate events occur. To permit preemption some mechanisms (machine dependent) must be implemented to save the *context* of the task, e.g. registers and stack pointer, at suspension occurrence. Such mecha-

nisms permit to resume the suspended computation when the task is rescheduled. An intermediate software solution is given by Contiki [15]. This OS uses a mono-stack memory model for an event-driven kernel. The application programs are dynamically loaded at run-time. It supports a thread-like coding style (proto-threads) but enforcing a sequential flow of control; optionally multi-threading can be adopted linking to a specific library partially coded for specific platforms. TinyOS, ERIKA Enterprise, and Contiki implement (either natively or through third-party support) the IEEE 802.15.4 stack of protocols.

2.2. Image Analysis over Visual WSNs

Gathering information from a network of scattered cameras, possibly covering a large area, is a common feature of many videosurveillance and ambient intelligence systems. However, most of classical solutions are based on a centralized approach, in which image processing is accomplished in a single unit. In this respect, the need to introduce distributed intelligent system—where distribution is not merely physical but also semantic—is motivated by several requirements, namely [16]:

- *Speed*: in-network distributed processing is inherently parallel; in addition, the specialization of modules permits to reduce the computational burden in the high level decisional nodes.
- *Bandwidth*: in-node processing permits to reduce the quantity of transmitted data, by transferring only information-rich parameters about the observed scene and not the redundant image data stream.
- *Redundancy*: a distributed system may be re-configured in case of failure of some of its components, still keeping the overall functionalities.
- *Autonomy*: each of the nodes may process the images asynchronously and may react autonomously to the perceived changes in the scene.

In particular, these issues suggest to move a part of intelligence towards the camera nodes. In these nodes, artificial intelligence and computer vision algorithms are able to provide autonomy and adaptation to *internal conditions* (e.g. hardware and software failure) as well as to *external conditions* (e.g. changes in weather and lighting conditions).

The mining and deployment of visual information involve particular problems in computer vision, such as change detection in image sequences, object detection, object recognition, tracking, and image fusion for multi-view analysis. For each of these problems, there exists a large corpus of already implemented methods (see e.g. [17] for a survey of change detection algorithms); however most of the techniques currently available are not suitable to be used in visual WSNs, due to the high computational complexity of algorithms or to excessively demanding memory requirements. Nevertheless, some attempts to employ non-trivial image analysis methods over visual WSN have been done. For example, [18] presents a visual WSN able to support the query of a set of images in order to search for a specific object in the scene. To achieve this goal, the system uses a representation of the object given by the Scale Invariant Feature Transform (SIFT) descriptors [19]. SIFT descriptors are known to support robust identification of objects even among cluttered background and under partial occlusion situations, since the descriptors are invariant to scale, orientation, affine distortion and partially invariant to illumination changes. In particular, using SIFT descriptors allows retrieving the object of interest from the scene, no matter at which scale it is imaged.

Interesting computer algorithms are also provided on the CMUcam3 vision system [11]. Besides basic image processing filters (such as convolutions), methods for realtime tracking of blobs on the base either of color homogeneity or frame differencing are available. A customizable face detector is also included. Such detector is based on a simplified implementation of Viola-Jones detector [20], enhanced with some heuristics to further reduce the computational burden. For example, the detector does not search for faces in the regions of the image exhibiting low variance.

Camera sensors have limited fields of views and can only perceive a portion of a scene from a single viewpoint. In addition, monocular vision totally lacks 3D information. To mine the entire scene and to deal with occlusions, it is natural to equip visual WSNs with multiview capabilities [21]. Due to bandwidth and efficiency considerations, however, images cannot be routinely shared on the network, so that no dense computation of 3D properties (like disparity maps and depth) can be made.

Nevertheless, the static geometrical entities observed in the scene may be suitably codified during the setup of the acquisition system. In addition, specially-designed references may be introduced in the scene for obtaining an initial calibration of the views acquired by each camera, thus permitting to find geometrical correspondences among regions or points of interest seen by different nodes. To this end, a coordinator node, aware of the results of such calibration step, may be considered, so as to translate events from the image coordinates to physical world coordinates.

Such approach may produce more robust results as well as a richer description of the scene.

3. METHODS

In a distributed reactive application, the mission of the visual WSN is to report any relevant change in the scene, where some real-time constraints must be satisfied. To this end, one should take into account both local computation and network communication issues to bound the maximum latency in the reaction.

Considering these constraints, but in the need for an efficient mining of the visual scene under examination, we focused on two different approaches: the first one based on simpler detection methods, but more efficient from a computational and storage points of view, the latter giving better performances for more specific object identification tasks while being computationally more intensive.

Moreover, considering a more global analysis and mining of the scene viewed from several different nodes (i.e. multiview), a final processing level for the aggregation of different single in-node data, and the final result of the scene mining needs to be implemented.

In the following we will analyze the above mentioned specific methods for the detection of changes in a scene, in particular regarding the two categories: simple, but quick and efficient change detection algorithms in Section 3.1, and on the other hand, more complex and demanding algorithms for an object detection and identification in Section 3.2. Finally in Section 3.3 we will present methods for the aggregation of hierarchical identifications, made at different level of complexity and from multiple acquiring points of view, in order to be able to correctly process the set of images to detect events.

3.1. Event and Change Detection

The task of identifying changes, and in particular regions in the image under analysis, taken at different times is a widely diffused topic encompassing very different disciplines [17]. In this context, the goal is the identification of one or more regions of pixels, relative to a change in the scene, and occurring within a sensible and a priori known area.

In order to perform fast, efficient, and enough reliable methods for this change detection, the generally adopted low-level methods are based on both frame differencing, and statistical methods. This assumption can be either the quick and immediate answer to a simple problems, or a preliminary synthesis for a deeper and more effective higher level analysis. The general model used is mainly based on background subtraction, where a single frame, acquired in a controlled situation, is stored (*BG-image*) and thereafter used as the zero-level, for different types of comparison in the actual processing.




Fig. 1. Cascade of classifier for object detection in a case study of car detection.

The first methods are very low-level ones and are thought in such a way that it can be possible and feasible to implement them also at firmware level, so to make them real-time operative in an automatic way, examples of these can be:

- Thresholding toward the BG-image
- Contrast modification with respect to the BG-image
- Basic edge detection on the image resulting from the subtraction with the BG-image

Another class of methods is based on statistical histogram analysis. Template histograms are stored of the regions of interests in the BG-image. Then a standard distance, such as the Kullback-Leibler divergence is computed between the region of BG-image, and the region of the actual image. If such distance overcomes a fixed threshold, then a change event is detected.

3.2. Object Detection

As it is well-known, the problem of detecting classes of objects in cluttered images is challenging. Supervised learning strategies have demonstrated to provide a solution in a wide variety of real cases, but there is still a strong research activity in this field. In the context of visual WSN, the preliminary learning phase may be accomplished off-site, while only the already trained detectors needs to be ported to the network nodes.

Among machine learning methods, a common and efficient one is based on the sliding windows approach; namely rectangular subwindows of the image are tested sequentially, by applying a binary classifier able to distinguish whether they contain an instance of the object class or not. A priori knowledge about the scene or—if available—information already gathered by other nodes in the network may be employed to reduce the search space either by (a) disregarding some region in the image and (b) looking for rectangular regions within a certain scale range (e.g. rectangular regions covering less than 30% of the whole image area).

For what regards the binary classifiers itself, among various possibilities, the Viola-Jones method is particularly appealing. Indeed, such classifier is based on the use of the so-called Haar-like features, a class of features with limited flexibility but which is known to support effective learning. A Haar-like feature is essentially given by the difference between sum of pixels in rectangular blocks; as such a Haar-like feature is

computable in constant time, once the integral image has been computed (see [20] for details). Thus, having enough memory to store the integral image, the feature extraction process needs only limited computational power. Classification is then performed by applying a cascade of classifiers, as shown in Fig. 1. A candidate subwindow which fails to meet the acceptance criterion in some stage of the cascade is immediately rejected and no further processed. In this way, only *detection* should go through the entire cascade. In addition, one may train the cascade—for example using *gentle AdaBoost*—in such a way that most of the candidate subwindows that do not correspond to instances of the object class are rejected in the first stages of the cascade, with great computational advantages.

The use of a cascade of classifiers permits also to adapt the response of the detector to the particular use of its output in the network, also in a dynamical fashion, in order to properly react to changes in the internal and external conditions. First of all, the tradeoff between reliability of detections and needed computational time may be controlled by adaptive real-time requirements of the overall network. Indeed, the detector may be interrupted at an earlier stage in the cascade, thus producing a quick even though less reliable output, which may be anyhow sufficient for solving the current decision making problem. In the same way, by controlling the threshold in the last stage of the cascade, the visual WSN may dynamically select the optimal trade-off between false alarm rate and detection rate needed in a particular context.

3.3. Aggregation of Partial Results

Considering a visual WSN with multiple views of the same scene for image mining purposes, the previously described methods can be managed in a hierarchical fashion, with faster algorithms yielding a primary response to a simple change detection task, and sending their results up through the hierarchy to more performing algorithms, in case of a positive detection. These higher level algorithms, due to their greater computational requirements, are implemented in order to be called only when positive feedback is given by the lower level ones. At the highest level of the hierarchy, a final decision-maker algorithm, operates having a not so strict computational and storage constraints, that can aggregate both the results of lower




Fig. 2. Architecture of the hierarchical levels of detection, in a car detection case study.

levels processing, but also to combine the multiple views processing, each one operating with the same hierarchy (see Fig. 2).

A common hypothesis for the multiple views aggregation process works on the basis of a given knowledge base of information, relative to the approximate possible positions, or Regions of Interest, where the objects to be mined can be detected. For example the decision maker algorithm can analyze the specific outcomes from different lower-level processing in each single node, and then be able to give a final output using its knowledge and working on the basis of a weighted computation of the single in-node results.

4. CASE STUDY: A PERVASIVE INFRASTRUCTURE FOR URBAN MOBILITY

After having presented methods for performing image analysis over WSN, in this section we describe a real case study concerning urban mobility. To this end, we provide a brief overview of IPERMOB project (Section 4.1), we describe the embedded vision sensor prototypes to be used in the project (Section 4.2) and, finally, we report preliminary results in a particular scenario (Section 4.3).

4.1. IPERMOB Project

The Tuscany regional project IPERMOB—“A Pervasive and Heterogeneous Infrastructure to control Urban Mobility in Real-Time” [22] is exploiting an approach based on visual WSNs for building an intelligent system capable of analyzing infomobility data and of providing effective decision support to final users, e.g. citizens, parking managers and munic-

ipalities. The partners of the IPERMOB project are a balanced mix of academical and industrial partners; the work presented in this paper is mainly implemented by the Signals and Images Lab of the Institute of Information Science and Technologies of the National Research Council, and the Real Time Systems Lab of the Scuola Superiors Sant’Anna of Pisa. The IPERMOB project is developing an open, interoperable system featuring leading-edge technologies for distributed imaging and wireless telecommunication. Supported by the real-world data gathered by pervasive monitoring, IPERMOB aims to become a laboratory for the study of novel models for mobility and to represent a benchmark for technological integration. IPERMOB infrastructure is organized in 3 tiers, corresponding to (i) data acquisition and processing, (ii) Metropolitan Information Bus (MIB) and (iii) end-users applications. The data acquisition and processing tier (by using visual WSNs and Vehicular ad-hoc Networks (VANETs)) is in charge of gathering and processing information, and of transmitting aggregated data to the MIB. The MIB, acting as a middleware application, stores historical data and makes all the information available in a suitable form to the final applications.

4.2. Embedded Vision Sensor Prototypes

In IPERMOB project, visual WSNs are used to monitor parking lots and roads, in order to mine the observed scenes and predicate something about parking lot occupancy and traffic flow. To this end, a number of embedded sensor prototypes is being considered in IPERMOB. In particular, two different architectures are being explored:

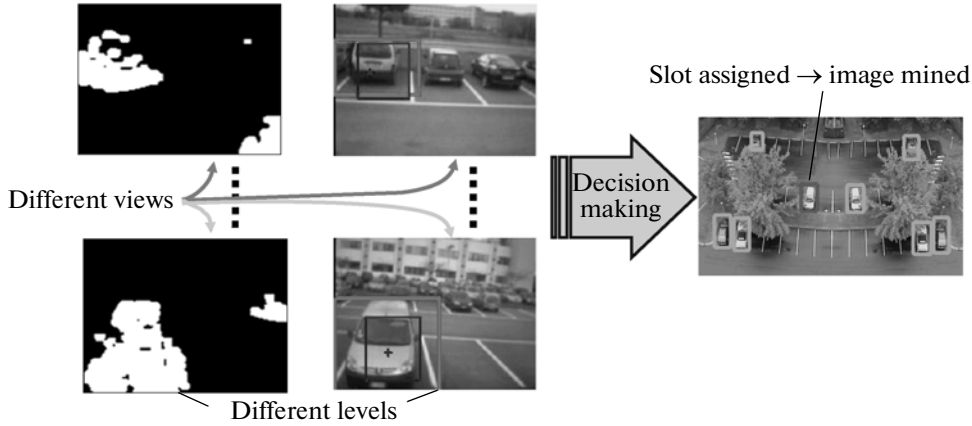


Fig. 3. The aggregation and decision-making process for WSN image mining in IPERMOB Project.

- a dual-microcontroller architecture, in which image acquisition and processing functions are separated from the network communication operations;
- a FPGA based board with dedicated hardware components for image processing and a soft microprocessor for image acquisition control and wireless communication.

The microcontroller-based board aims to keep to a minimum the power necessary to perform the required tasks, and this implies that e.g. image size and/or frame rate are somewhat limited. In this architecture, one microcontroller is dedicated to the acquisition of the image from the camera and to the bulk of the image processing operations. The other microcontroller performs the final stage of the processing and handles the radio frequency communication with the remaining nodes in the network. Extra modularity may be gained, since the image analysis algorithm may be split so as to balance the load between the two microcontrollers and minimize the amount of exchanged data.

In the FPGA-based board the power consumption constraint has been slightly relaxed to explore a solution with more flexibility and processing capabilities. This board features plenty of RAM space and may store several images and intermediate processing results. Images are fed from the camera with a custom component written in Verilog, suited to work at the maximum resolution and speed permitted by the camera (640×480 , 30 fps). The FPGA hosts also a soft microprocessor in order to be programmed using the same code developed for the dual microcontroller unit—written in C—although custom Verilog components can be developed to speed up the vision algorithms.

4.3. The Parking Lot Scenario

In this section, we consider the parking lot scenario, i.e. a case study at the Pisa International Airport

for the landside vehicle flow, in which a set of camera nodes with partially overlapping field of views is in charge of observing and estimating dynamic real-time traffic related information, in particular regarding traffic flow and availability of parking places. It is assumed that each camera knows the geometry of the parking spaces that it is in charge of monitoring. In addition, we assume that a coordinator node—on which the final *decision-maker* algorithms described in Section 3.3 are hosted—knows the full geometry of the parking lot as well as the calibration parameters of the involved cameras, in order to properly aggregate their beliefs.

Figure 3 shows the workflow in the real case study. In a typical scenario, some of the camera nodes in the first level of the hierarchy detects a change in the scene (by using the methods reported in Section 3.1) and triggers object detection methods in the remaining camera nodes that are in charge of monitoring that area. The aggregation and decision making process are finally performed by the coordinator node.

For what regards object detection, several cascade of classifiers have been trained, each one specialized in detecting particular view of cars (e.g. front, rear and side views). A large set of labeled acquisition has been made of the real case study and used for training purposes. We notice that first stages in the cascade have low computational complexity but reasonable performance (that is almost 100% detection rate but even 50% false alarm rate). Composition of the stages entails then high performance of the entire cascade, since both overall detection rate and overall false alarm rate are just the product of the detection rate and false alarm rate of the single stages. For example, using $N = 10$ stages in the cascade, each one with detection rate $v = 99.5\%$ and false alarm rate $\mu = 50\%$, one gets an overall detection rate $v_{\text{global}} = v^N \approx 95\%$ and false alarm rate $\mu_{\text{global}} = \mu^N \approx 0.1\%$.



Fig. 4. Car detection and analysis of parking lot occupancy status.




Fig. 5. Detection of vehicles on a road for traffic flow.

Figures 4 and 5 show examples of two processed image sequences in the framework of the project, respectively for parking availability and traffic flow.

5. CONCLUSION AND FURTHER WORK

In this paper, strategies for tackling the image mining problem over visual sensor networks have been presented, taking into account both the structural limits of embedded platforms and the possibilities gained by performing pervasive computing and by the very flexible organization of the network topology.

A multi-node processing method has been introduced to decompose a complex computer vision task into a hierarchy of computationally simpler problems to be solved over the nodes of the network. Besides computational advantages, such hierarchical decision making approach appears to be more robust and fault-tolerant.

Such ideas and preliminary results are illustrated by means of an application scenario regarding infomobility, which is currently in development within the IPERMOB project. Set-up and commissioning of the visual sensor network in the project testbed will start in fall 2010. It is expected that the testbed will give precise information whether the presented approach is mature for technological transfer.

6. ACKNOWLEDGMENTS

This work has been partially supported by POR CREO Tuscany Project “IPERMOB”—A Pervasive and Heterogeneous Infrastructure to control Urban Mobility in Real-Time (Regional Decree no. 360, 01/02/2010). We would like to thank Dr. Mario De Pascale (IT Manager at Societa Aeroporto Toscano S.p.A.) for his kind support and for providing access to the testing area.

REFERENCES

1. S. Soro and W. Heinzelmann, "A Survey of Visual Sensor Networks," *Adv. Multimedia Article ID 640386* (2009).
2. A. Adam, E. Rivlin, I. Shimshoni, and D. Reinitz, "Robust Real-Time Unusual Event Detection Using Multiple Fixed-Location Motions," *IEEE Trans. PAMI* **30**, 555–560 (2008).
3. S. Colantonio, D. Conforti, M. Martinelli, D. Moroni, F. Salvetti, and A. Sciacqua, "An Intelligent and Integrated Platform for Supporting the Management of Chronic Heart Failure Patients," in *Proc. Computers in Cardiology* (Bologna, 2008), pp. 897–900.
4. O. Salvetti, E. A. Cetin, and E. Pauwels, Special Issue on Human-Activity Analysis in Multimedia Data. *Eurasip J. Adv. Signal Proc.* (2008).
5. P. Pagano, F. Piga, G. Lipari, and Y. Liang, "Visual Tracking Using Sensor Networks," in *Proc. 2nd Int. Conf. Simulation Tools and Techniques, ICST* (Rome, 2009), pp. 1–10.
6. M. Magrini, D. Moroni, C. Nastasi, P. Pagano, M. Petracca, G. Pieri, C. Salvadori, and O. Salvetti, "Image Mining for Informability," in *3rd Int. Workshop on Image Mining Theory and Applications* (INSTICC Press, Angers, 2010), pp. 35–44.
7. D. Kundur, C. Y. Lin, and C. S. Lu, "Visual Sensor Networks," *EURASIP J. Adv. Signal Proc. Article ID 21515* (2007).
8. W. Feng, B. Code, E. C. Kaiser, M. Shea, W. Chang Feng, and L. Bavoli, "Panoptes" Scalable Low-Power Video Sensor Networking Technologies," in *Proc. ACM Multimedia* (Berkeley, CA, 2003), 562–571.
9. M. H. Rahimi, R. Baer, O. I. Iroezzi, J. C. García, J. Warrior, D. Estrin, and M. B. Srivastava, "Cyclopes: in situ Image Sensing and Interpretation in Wireless Sensor Networks," in *Proc. SenSys* (San Diego, 2005), pp. 192–194.
10. S. Hengstler, et al., "Mesheye: a Hybrid-Resolution Smart Camera Mote," in *Proc. ACM'2007* (Seoul, 2007), pp. 360–369.
11. A. Rowe, et al., "CMUcam3: An Open Programmable Embedded Vision Sensor," Technical Report CMU-RI-TR-07-13 (Robotics Institute, Pittsburgh, PA, 2007).
12. P. Chen, et al., "Critic: A Low-Bandwidth Wireless Camera Network Platform," in *Distributed Smart Cameras* (2008), pp. 1–10.
13. TinyOS: Open-Source operation System, Available from: <http://www.tinyos.net>
14. <http://erika.tuxfamily.org>
15. Contiki: The Operating System for Embedded Smart Objects – the Internet of Things, Available from <http://www.sics.se/contiki>
16. P. Remagnino, A. O. Shihab, and G. A. Jones, "Distributed Intelligence for Multi-Camera Visual Surveillance," *Pattern Recogn.* **37**, 675–689 (2004).
17. R. J. Radke, S. Andra, O. Al-Kofani, and B. Roysam, "Image Change Detection Algorithms: a Systematic Survey," *IEEE Trans. Image Proc.* **14**, 294–307 (2005).
18. T. Yan, D. Ganesan, and R. Marumatha, "Distributed Image Search in Camera Sensor Networks," in *Proc. ACM SenSys* (Raleigh, NC, 2008), pp. 155–168.
19. D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *Int. J. Comput. Vision* **60**, 91–110 (2004).
20. P. A. Viola and M. J. Jones, "Robust Real-Time Face Detection," *Int. J. Comput. Vision* **57**, 137–154 (2004).
21. P. Pagano, F. Piga, and Y. Liang, "Real-Time Multi-View Vision Systems Using WSNs," in *Proc. ACM Symp. Appl. Comp.* (Honolulu, 2009), pp. 2191–2196.
22. IPERMOV: A Pervasive and Heterogeneous Infrastructure to Control Urban Mobility in Real-Time, Available from: <http://www.ipermob.org>