Proceedings of the 2005 IEEE
International Conference on Robotics and Automation
Barcelona, Spain, April 2005

Cooperative Fire Detection using Unmanned Aerial Vehicles*

Luis Merino, Fernando Caballero, J.R. Martínez-de Dios and Aníbal Ollero

Robotics, Computer Vision and Intelligent Control Group
University of Seville
Escuela Superior de Ingenieros, Camino de los descubrimientos s/n, 41092 Sevilla, Spain
{merino,caba,jdedios,aollero}@cartuja.us.es

***Abstract* – The paper presents a framework for cooperative fire detection by means of a fleet of heterogeneous UAVs. Computer vision techniques are used to detect and localize fires from infrared and visual images and other data provided by the cameras and other sensors on-board the UAVs. The paper deals with the techniques used to decrease the uncertainty in fire detection and increase the accuracy in fire localisation by means of the cooperation of the information provided by several UAVs. The presented methods have been developed in the COMETS multi-UAV project.**

Index Terms – Cooperative perception, data fusion, infrared images, aerial robotic vehicles.

I. INTRODUCTION

Surveillance and monitoring activities in disaster scenarios requires high mobility and perception capabilities in natural environments. Aerial robotic vehicles can be used to provide such capabilities overcoming the limitations of ground robotic vehicles to access to the appropriate locations for surveillance and monitoring. The use of a fleet of robotic aerial vehicles improves the coverage and the reliability of the mission, when comparing with the application of a single UAV. One of the main objectives of the COMETS project [1]is to design and implement a system for cooperative activities using heterogeneous UAVs.

The UAVs in COMETS are complementary platforms. Currently, both helicopters and airships have been integrated (see Fig. 1). Helicopters are suited to agile target tracking and inspection and monitoring tasks. On the other hand, airships, having much less manoeuvrability, can be used to provide global views or to act as communications relay. Furthermore, the COMETS UAVs are also heterogeneous in terms of onboard processing and sensorial capabilities. Then, some UAVs have infrared and visual video cameras while another has a still high resolution camera and a specialized fire sensor.

In order to test and validate the developments, the project includes experiments and demonstrations in forest fire applications. UAVs cooperation is very valuable in such a very challenging context. Missions involve fire alarm detection, confirmation and localization, and fire monitoring.

This paper presents the approach followed for cooperative fire detection by a fleet of heterogeneous UAVs. The objectives are to determine the position of potential fire alarms, and also to reduce the number of false alarms by means of cooperation. The paper deals mainly with fire detection algorithms using vision, alarm tracking and data fusion. Results from experiments with small controlled fires are presented (see Fig. 1).

Fig. 1 An autonomous airship and a helicopter monitoring a fire in the COMETS experiments carried out in Lousa (Portugal).

The paper is organized as follows. Section II presents the techniques for fire detection in visual and infrared images. Section III describes how to obtain an estimation of the geographical coordinates of those alarms. Section IV presents the probabilistic estimation framework used for data fusion. Section V presents the experiments and the results obtained. The conclusions and future trends in section VI complete the paper.

II. FIRE SEGMENTATION

Within the COMETS project, both infrared and visual cameras are considered as onboard sensors for the UAVs. In order to apply data fusion procedures, common features (binary images containing potential fire alarms) are extracted from both types of images.

A. Fire segmentation in infrared images

The aim of the processing of infrared images is to produce binary images containing fire alarms while discarding false alarms. Fire segmentation basically consists in thresholding since fires appear in infrared images as high intensity level regions. It should be noticed that although the temperature of fire (often over 900ºC) is much higher than the temperature of the image background, temperature-based criteria can not be used for thresholding since the measures of temperature are influenced by the emmissivity indices of the materials, which are very difficult to estimate in such an unstructured environment. The threshold selection method should take into account the particular conditions of the application such as the considerations to discard false alarms. Besides, images from miniaturized infrared cameras

* This work is supported by COMETS IST-2001-34304 and CROMAT DPI-2002-04401-C03-03

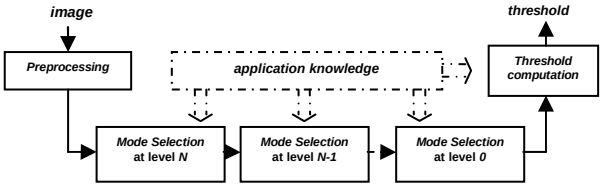

Fig. 2 Scheme of the proposed training-based threshold selection method.

are affected by blurs originated by the high frequency vibrations of the UAVs and their onboard pointing devices and the long exposure periods required by the still low sensitivity of the infrared detectors available at the current state of the technology. The used threshold selection method takes into account knowledge of the computer-vision to adapt to the specific conditions of the application.

The proposed thresholding method assumes that the intensity values of the pixels of the objects are similar within a certain region in the histogram, which will be called *histogram mode*. The method (depicted in Fig. 2) performs a coarse-to-fine selection of the mode or modes corresponding to the object of interest. It identifies the modes of interest by using multiresolution descriptions of the histogram at several levels of resolution.

Mode Selection at level l decomposes the histogram description at level l in histogram modes and selects the modes that are likely to correspond to the object of interest according to the knowledge of the application. The modes selected at level l are analyzed at the immediate finer level of resolution (level l-1). This analysis is repeated until the finest level (level 0). Finally, threshold computation is carried out by taking simple considerations on the selected histogram regions at level l=0. The knowledge of the application, which is used for the supervision of *Mode Selection*, is extracted from a set of training images and their corresponding desired threshold values, and incorporated in a fuzzy system by applying the *ANFIS* algorithm [2]. Further information of the training-based method for threshold selection can be found in [3].

The adaptation to this application was carried out by selecting training infrared images with different illumination conditions, different image backgrounds and different objects including fires and false alarms such as heat emitting sources (i.e. car engines). The algorithm was trained to detect the fires and discard the false alarms. Image blurs are assumed in the training due the selection of blurred training images. Once trained, object segmentation is carried out by a threshold-based region growing method. Fig. 3 shows an infrared image of a fire and the corresponding thresholded image.

Some heuristic simple rules are applied to discard false alarms that can not be eliminated by threshold-based segmentation. For instance, the average intensity level originated by heated objects (such as engines and chimneys which are not burning) is lower that the one produced by fires. Thus, the local contrast in case of heated objects is lower than in case of fire.

B. *Fire segmentation in visual images*

The processing of the images provided by the visual camera will also produce binary images containing potential

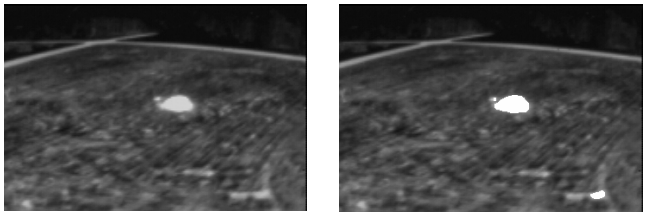

Fig. 3 Infrared fire segmentation algorithm results. Left, original image, right segmented objects.

fire alarms. The algorithm is based on the fact that visual color images of fire have high absolute values in the red component of the RGB coordinates. This property permits threshold-based criteria on the red component of the color images to segment fire images in natural scenarios.

However, not only fire gives large values in the red component. Another characteristic of fire is the ratio between the red component and the blue and green components.

Let $r(u,v)$, $g(u,v)$ and $b(u,v)$ be respectively the red, green and blue components in the pixel (u,v) of an image. Thus, the algorithm consists of two stages. First, the pixels $\mathbf{m}=[u,v]^T$ belonging to the region of the RGB space defined by

$$\left.\begin{aligned} r(u,v) - k_g \cdot g(u,v) \geq 0 \\ r(u,v) - k_b \cdot b(u,v) \geq 0 \end{aligned}\right\}, \qquad (1)$$

are selected as candidates. k_g and k_b are parameters that define the RGB region. For the pixels selected, a threshold selection based on an analysis of the histogram of the red field is used to compute the thresholds for segmenting fire. The experiments revealed that the iterative thresholding algorithm described in [4] applied over the red field provides good flame segmentation for most visual images. Fig. 4 shows some results of the algorithm.

C. *Characterization of the algorithms.*

The elimination of false alarms is a main issue in fire detection. The vehicles of the fleet will cooperate to reduce the number of false alarms by means of data fusion. Thus, the probabilistic characterization of the above detection techniques is required in order to determine the fire detection capabilities of each vehicle.

Both algorithms have been analyzed with a large set of images. The algorithms are characterized by the probabilities P_D of detection and P_F of false positive outputs. These probabilities have been computed as follows:

- P_D is the ratio between the alarms correctly detected and the total number of fire alarms presented in the set of images.
- P_F is the ratio between the number of images where the algorithm detected fire incorrectly and the total number of images of the sequence.

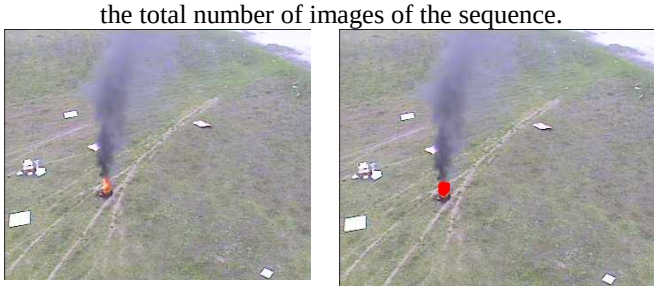

Fig. 4 Visual fire segmentation algorithm results. Left, original image, right segmented objects

TABLE I
ALGORITHMS CHARACTERISTICS

	IR	Visual
P_D	100%	90%
P_F	10%	3%

These ratios will be used to obtain a prior belief on that the object detected is a fire; and to update this belief when new measurements from the same object are obtained.

Table 1 shows the values for the algorithms used for fire segmentation in visual and infrared images (for a sequence of images different than that used during the training phase for the infrared algorithm). Actually, both algorithms can be tuned, thus leading to operating curves that relate P_D and P_F. Depending on the requirements of the mission, it could be preferred a higher detection probability or a lower false alarm ratio (for example, for alarm confirmation purposes).

D. Measurements obtained from the image processing stage.

As a result of the previous algorithms, the infrared and visual data are reduced to a common format: binary images containing potential fire alarms. The different objects presented on the image are labeled. Therefore, from each image, a set of possible alarms on the image plane is extracted, each alarm characterized by $[\mathbf{m}_a,\mathbf{C}_{ma},P_D,P_F]$, where $\mathbf{m}_a$ is the position in pixel coordinates of the object and $\mathbf{C}_{ma}$ is an estimation of the covariance matrix of the error in this position due to quantization errors. P_D and P_F are the values associated to the algorithm employed for the detection of the alarm.

III. GEOLOCATION

The sensors onboard the different UAVs (see Fig. 5) are used to compute, in a global and common coordinate frame, the position and orientation of the UAV itself and also that of the cameras that carries onboard (that will be denoted by $\mathbf{x}_{cam}$). Also, an estimation of the covariance matrix $\mathbf{C}_{cam}$ of the errors of these quantities is available. Thus, all the images gathered are labeled with these data.

If the cameras are calibrated, knowing a digital elevation map, denoted by D, it is possible to obtain the position $\mathbf{x}$ of an alarm in the common global coordinate frame, from its position on the image plane, $\mathbf{m}_a$. The position $\mathbf{x}$ is a function of the position and orientation of the camera and also depends on D:

$$\mathbf{x} = \mathbf{f}(\mathbf{m}_a, \mathbf{x}_{cam}, D) \qquad (2)$$

The function $\mathbf{f}$ encompasses the pin-hole model of the camera. Clearly, the function $\mathbf{f}$ is non-linear, and in the general case the dependence on the map D cannot be expressed analytically.

The errors on the position and orientation of the camera and those of the position of the alarm on the image plane will translate into uncertainties in the final position $\mathbf{x}$ (see scheme of Fig. 6). It is very interesting to have and maintain an estimation of the uncertainty in the position of the alarms. Thus, the uncertainties are propagated to obtain an estimation of the covariance matrix of the errors in $\mathbf{x}$ by using the so-called Unscented Transform [5][6].

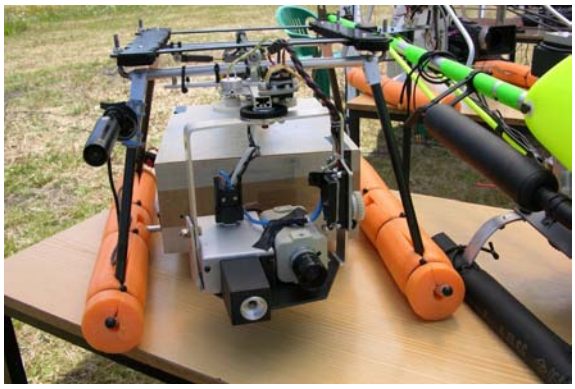

Fig. 5 Each UAV carries a electronic box and sensors to compute the position and orientation of the onboard cameras.

The unscented transform is chosen because allows considering a more general class of functions than the usual first order expansion (in this case the dependency of $\mathbf{f}$ in (2) respect to D cannot be expressed analytically). Also the estimated covariance matrix is more accurate than that obtained by means of a Taylor expansion [5].

Then, each UAV will compute the position of the detected alarms on the image plane by means of this geolocation procedure. As a result, each UAV will provide measures of the form $[\mathbf{x}_m,\mathbf{C}_m,P_D,P_F]$, where $\mathbf{x}_m$ is the measured position of the alarm in the common coordinate frame and $\mathbf{C}_m$ the estimated covariance of the errors on this position.

IV. COOPERATIVE ALARM DETECTION AND LOCALIZATION

The objective of the system is, from the measurements provided by the vehicles of the fleet, to cooperatively estimate the geographical position of potential fire alarms while trying to reduce the number of false alarms. Thus, the state of an alarm i is defined by a continuous random variable $\mathbf{x}_{ai}$ (the position of the alarm) and a Bernoulli binary random variable $h_i \in \{0,1\}$ (h_i=1 means that the alarm is a fire).

Then, at time k the current information about the set of potential fire alarms is defined by $A(k) = \{[\bar{\mathbf{x}}_{ai}(k), \mathbf{C}_{ai}(k), p_i(k)], i = 1,\ldots,L(k)\}$, where $\bar{\mathbf{x}}_{ai}(k)$ is the estimated position for alarm i at time k, $\mathbf{C}_{ai}(k)$ is the estimated covariance matrix of the errors in $\mathbf{x}_{ai}$ and $p_i(k)$ is the estimated probability for this alarm to be a fire (h_i=1).

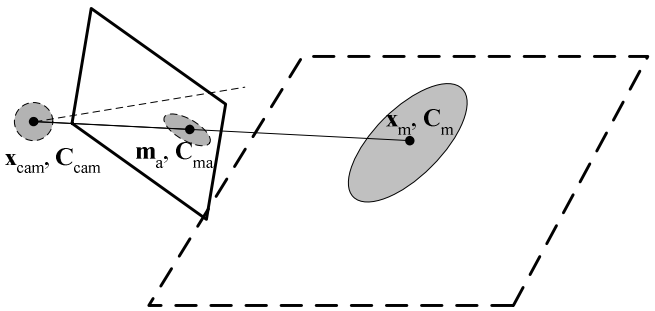

Fig. 6 Scheme of the uncertainties propagation during the geolocation process. The errors on the estimated camera position and orientation and those on the estimated position of the alarm on the image plane propagate through the projection model into errors in the final estimated geographical position of the alarms.

As it will be seen, this probability $p_i(k)$ will be calculated and updated from the probabilities P_D and P_F associated to the fire detection algorithms. $L(k)$ is the number of alarms at time k. The objective is to iteratively re-estimate the state of the alarms as new measurements arrive from the different UAVs.

Fig. 7 illustrates the procedure that consists of the following stages:

A. Prediction stage

At time $k+1$, the state for the alarms can be predicted by using a motion model. Through this paper, the position of the alarms are assumed to be static. This is realistic for the considered scenario, and also allows the measurements to be fused with arbitrary latency [7]. More complicated models can be considered for other kind of alarms, although the synchronization issues have to be taken into account.

B. Data association

At time $k+1$ the different UAVs provide a set of measurements,

$$M(k+1)=\left\{\left\lfloor \mathbf{x}_{mj}(k+1), \mathbf{C}_{mj}(k+1), P_{Dj}, P_{Fj} \right\rfloor, \; j=1,\ldots,N(k+1)\right\} \quad (3)$$

where $N(k+1)$ is the total number of measurements provided by the UAVs at time $k+1$.These measurements should be associated to the current set of alarms (see Fig. 7, center). A gated nearest neighbor technique is used [8]. The measurement j is associated with alarm i if

$$d_{ij}=(\mathbf{x}_{ai}-\mathbf{x}_{mj})^T \mathbf{S}_{ij}^{-1}(\mathbf{x}_{ai}-\mathbf{x}_{mj}) \le d, \quad (4)$$

where $\mathbf{S}_{ij}=\mathbf{C}_{ai}+\mathbf{C}_{mj}$ and d is chosen from the χ^2 distribution [8][9].

Additional constraints are applied. For the measurements given by one UAV, if there are some that accomplish (4), then the nearest (that of minimum d_{ij}) is chosen as the correct association. Of course, different measurements from *different* UAVs can be associated to the same alarm.

C. Update stage

Assume that the measurements $\{j=1,\ldots,M_i\}$ have been associated to alarm i. Then, the estimated position $\bar{\mathbf{x}}_{ai}(k)$ and covariance matrix $\mathbf{C}_{ai}(k)$ are updated following equations

$$\mathbf{C}_{ai}^{-1}(k+1)=\mathbf{C}_{ai}^{-1}(k)+\mathbf{C}_{m1}^{-1}(k+1)+\cdots+\mathbf{C}_{mM_i}^{-1}(k+1). \quad (5)$$

$$\bar{\mathbf{x}}_{ai}(k+1)=\mathbf{C}_{ai}(k+1)\left[\mathbf{C}_{ai}^{-1}(k)\bar{\mathbf{x}}_{ai}(k)+ +\mathbf{C}_{m1}^{-1}(k+1)\mathbf{x}_{m1}(k+1)+\cdots+\mathbf{C}_{mM_i}^{-1}(k+1)\mathbf{x}_{mM_i}(k+1)\right] \quad (6)$$

that corresponds to the information from the Kalman filter [7]. This leads to a reduction on the uncertainties on the position of the alarm, as is illustrated in Fig. 7, right.

The probability $p_i(k)$ of being a fire alarm is updated following the Bayes rule:

$$p_i(k+1)=\frac{(P_{D1}\cdots P_{DM_i})p_i(k)}{(P_{D1}\cdots P_{DM_i})p_i(k)+(P_{F1}\cdots P_{FM_i})[1-p_i(k)]}. \quad (7)$$

D. Insertion of new alarms.

If a measure j has not been associated with any of the alarms, then a new track is initialized (see Fig. 7, right).

The prior belief of this new track $p_{new}(0)$ is computed as:

$$p_{new}(0)=\frac{P_{Dj}}{P_{Dj}+P_{Fj}}, \quad (8)$$

that supposes an initial probability 0.5 of having fire in this position.

E. Not detected alarms.

Finally, it is checked if the alarms are in the field of view of the UAVs that have not provide any measurement for them. If so, for each UAV with fire detection capabilities defined by P_D and P_F that does not provide measures the probability is updated by using (9):

$$p_i(k+1)=\frac{(1-P_D)p_i(k)}{(1-P_D)p_i(k)+(1-P_F)[1-p_i(k)]}. \quad (9)$$

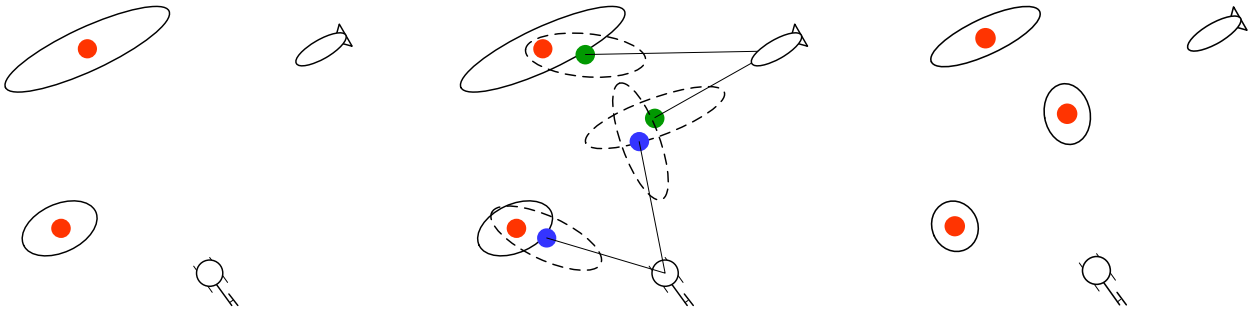

Fig. 7 Alarms tracking scheme. Left: previous detected alarms and their uncertainties (presented as ellipses). Center: The UAVs gather new measurements. These measurements are associated to the current tracks. Right: the update stage reduces the uncertainties of the tracks with new measurements. New tracks are added.

V. Experimental results

A. *Description of the experiments*

One of the key issues of the COMETS project is the demonstration. The project will be demonstrated in fire detection and monitoring activities, and also for terrain mapping missions.

Several experiments with controlled small fires have been carried out at Lousa, Portugal during the years 2003 and 2004 for integration and testing purposes. In this section, a brief explanation of the experiments is presented.

In the experiments three vehicles are considered (one autonomous helicopter, one autonomous blimp and a teleoperated helicopter). As it was mentioned in section III, the UAVs are equipped with sensors and hardware to obtain the position and orientation of their onboard cameras. A common coordinate frame is set based on the coordinates given by the GPS (all UAVs have one GPS on board).

For geolocation purposes, a digital elevation map of the zone is available.

The Lousa airfield is also a training site for firemen in the region of Coimbra, in Portugal. The fire brigades are in charge of organizing small controlled fires that are being used for fire detection and monitoring demonstration purposes.

The controlled fires used in the fire detection tests are originated by the burning of small shrubs (see Fig. 8).

B. *Experimental results.*

This section presents the results obtained during a fire detection test. Two UAVs are involved: one carries a visual video camera and the other an infrared camera.

The algorithms of section II are used for fire segmentation. The probabilities P_D and P_F of Table I apply.

In the system currently implemented, the measurements obtained by the different UAVs are received in a central processing node where the data association and tracking is carried out.

Fig. 9 shows a typical snapshot of the alarm tracking procedure. Two alarms (labeled by 0 and 3) of the current list of tracked fire alarms can be seen. In this case, two new measurements from different UAVs are received. These measurements are associated to alarm 3, and lead to a reduction in the uncertainties on the position of the alarm, and also on the belief of being an actual fire alarm. On the other hand, alarm 0 is a false alarm that is in the field of view one of the UAVs but it has not been detected.

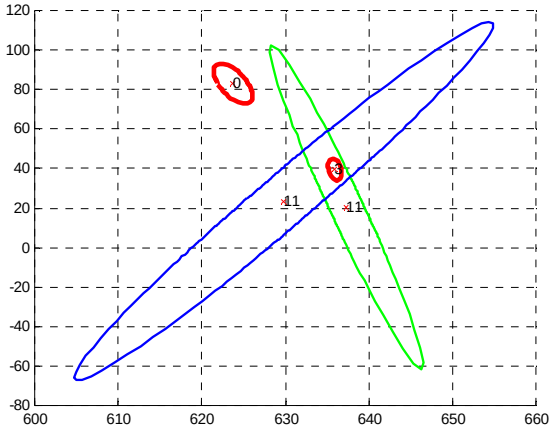

Fig. 9 A snapshot of the tracking process during the experiments. Two alarms 0 and 3 are being tracked. Two new measurements arrive (labeled as 11) and are associated to alarm 3 (uncertainties magnified).

Fig. 10 shows the evolution on the standard deviation for the position of the alarms 0 (dashed) and 3 (dotted). The initial values depends on the accuracy of the position estimation capabilities of the particular UAV, and also strongly on the orientation of the camera due to the non-linearity of the projection. The initial estimation of the covariance of the errors on the sensors is quite conservative. As new measurements arrive the uncertainty is reduced.

Fig. 11 shows also the evolution on the probability p of being fire for each alarm. This probability evolves following (7), (8) and (9). It can be observed that alarm 0 (dashed plot) is detected at time 20, and also at time 21 (in this case by the same UAV), increasing the probability p. However, this alarm is later not detected being in the field of view of the sensor of the other UAV. As a result, its probability of being a fire alarm is decreased (in fact, this was a false alarm). On the other hand, the dotted plot shows the evolution of the probability for alarm 3, an actual fire alarm.

Fig. 8 Image from on UAV of one of the fire detection experiments.

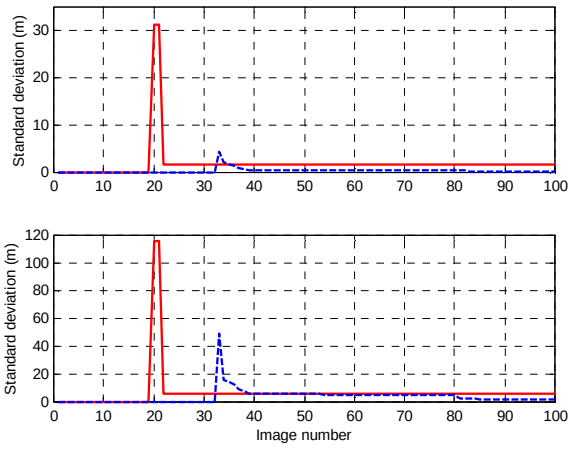

Fig. 10 Evolution of the standard deviation of the error in X and Y coordinates for the detected alarms.

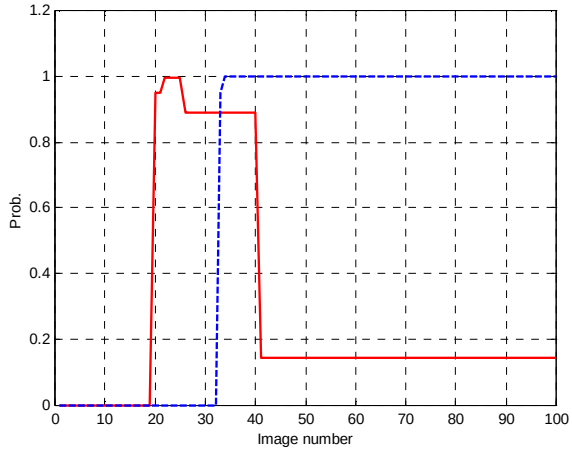

Fig. 11 Evolution of the belief on being a fire alarm.

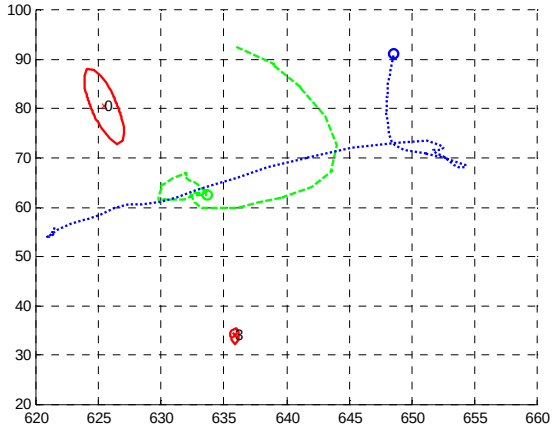

Fig. 12 Trajectories and final position of the detected alarms. Axes in meters.

As ground truth, the position of the fire alarms were recorded by using a GPS. Fig. 12 shows the trajectories of the UAVs and the position of the detected alarms. The accuracy on the estimated absolute position is within 5 meters in coordinates X and Y.

VI. Conclusions and Future Trends

A. Conclusions.

The paper has presented the multi-UAV framework used in the COMETS project for cooperative fire detection. It considers cameras of different modalities. Algorithms for fire detection in infrared and visual cameras have been presented.

The measures resulting of applying these fire detection algorithms are integrated into a cooperative state estimation procedure, in which the position of the fire alarms and the nature of the alarm (false alarm or real fire) are iteratively re-estimated as new measurements come from the different vehicles.

Experimental results of actual flights with small controlled fires are presented.

B. Future trends.

Data association is a key problem. Here a simple nearest neighbor strategy is considered. However, in complex scenario, this can lead to bad association. Thus, more complex techniques such as multiple hypothesis tracking [10] will be considered.

In the paper, although of different modalities, only vision related results have been presented. However, one of the UAVs of the fleet will carry in the close future a different kind of sensor that will be integrated in the fire detection system. The sensor consists of a photo detector that responses to a UV radiation characteristic of fire, and thus, the sensor is able to detect the presence of fire within its field of view. The sensor gives a scalar measure, but it is not possible to discriminate if one measure corresponds to a close small fire or to a farther and bigger fire.

For this kind of sensor, a grid-based localization technique is more suitable for fire localization. The objective will be to fuse the data gathered by this sensor with the data obtained from the images to increase the reliability.

Acknowledgment

The authors thank the cooperation of the partners of the COMETS project, as well as other members of the COMETS and CROMAT team at Seville.

References

[1] http://www.comets-uavs.org. COMETS Project web site.
[2] J. S. R. Jang, "ANFIS: Adaptive-Network-based Fuzzy Inference Systems", *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 23, no. 3, pp. 665-685, May 1993.
[3] J.R. Martinez-de Dios and A. Ollero, "A new training-based approach for robust thresholding", in *Proceedings of the World Automation Congress WAC'2004*, 28 June - 1 July, 2004, Seville, Spain.
[4] T.W: Ridler and S. Calvard. "Picture thresholding using an iterative selection method". *IEEE Transactions in Systems, Man and Cybernetics*, Vol. 8, pp. 630-632, 1978.
[5] S. Julier and J. Uhlmann, "A new extension of the kalman filter to nonlinear systems," in *Proceedings of the 11th Int. Symp. on Aerospace/Defence Sensing, Simulation and Controls*, 1997
[6] T. Schmitt, R. Hanek, M. Beetz, S. Buck and B. Radig. "Cooperative probabilistic state estimation for vision-based autonomous mobile robots". *IEEE Transactions of Robotics and Automation*, Vol. 18, pp. 670-684, October 2002.
[7] M. Ridley, E. Nettleton, S. Sukkarieh and H. Durrant-Whyte. "Tracking in decentralised air-ground sensor networks". *Proceedings of the 5th International Conference in Information Fusion, ICIF 2002*. Vol. 1, pp. 616-623, July 2002.
[8] H.J.S. Feder, J.J. Leonard and C.M. Smith. "Adaptive mobile robot navigation and mapping". *International Journal of Robotics Research*. Vol. 18, no. 7, pp. 650-668.
[9] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.
[10] D.B. Reid. "An algorithm for tracking multiple targets". *IEEE Transactions on Automatic Control*. Vol. 24, pp. 843-854. December 1979.