Cooperative fault detection and isolation in a surveillance sensor network: a case study

Julien Marzat, Hélène Piet-Lahanier, Sylvain Bertrand

Cooperative fault detection and isolation in a surveillance sensor network: a case study

Julien Marzat * **Hélène Piet-Lahanier** * **Sylvain Bertrand** *

* *DTIS, ONERA, Université Paris-Saclay, F-91123 Palaiseau, France*
e-mail: {firstname.lastname}@onera.fr

Abstract: This work focuses on Fault Detection and Isolation (FDI) among sensors of a surveillance network. A review of the main characteristics of faults in sensor networks and the associated diagnosis techniques is first proposed. An extensive study has then been performed on the case study of the persistent monitoring of an area by a sensor network which provides binary measurements of the occurrence of events to be detected (intrusions). The performance of a reference FDI method with and without simultaneous intrusions has been quantified through Monte Carlo simulations. The combination of static and mobile sensors has also been considered and shows a significant performance improvement for the detection of faults and intrusions in this context.

Keywords: cooperative fault detection and isolation, sensor network, intrusion detection, performance evaluation

1. INTRODUCTION

Two main applications of sensor networks are usually considered, namely environmental monitoring (continuous sensing) and event detection (surveillance, target tracking). The scenario considered in the EuroSWARM * project is the persistent monitoring and surveillance of a large area in order to protect a high value asset from intruders in a battlefield context. This paper is thus dedicated to event detection networks, whose particular features are the non-persistence of the excitation signals, limited sensor ranges, and possibly sparse coverage. Cooperative FDI can be very useful to avoid sending manned troops or unmanned vehicles from the military basis for intrusion confirmation in the case of a faulty event detection when no intrusion is present or of wrong localisation in the case of an intrusion. In the first case where no intrusion is reported, this also helps the maintenance of the network by replacing faulty nodes without the need for a systematic strategy. This would entail cost saving and security improvement.

For reliable sensor fusion and an enhanced possibility of detecting faults in a cooperative manner, the network should locally present a high density of (possibly heterogeneous) sensors. Otherwise, information gathered on possible intrusions will only be based on scarce measurements, and the lack of possibility of comparison with other sources of information makes it difficult to distinguish between true events and faults without a human intervention. Heterogeneity increases complexity, since it is hard to compare measurements between sensors of different natures and this turns out to be possible only with high-level preprocessed data.

As indicated by Ni et al. (2009), sensor network data is expected to be correlated in both the spatial domain and temporal domain and this is the main assumption at the basis of FDI methods. In surveillance networks, what is defined as a fault may also characterize an event under some conditions. The main difficulty for FDI in event detection sensor networks is thus the coupling between an event that should be detected and a fault. There are two critical situations: (1) No event is present but signals are produced by some faulty sensors and (2) An event is present but is not detected by some faulty sensors. To differentiate between a fault and an event, a training phase may be utilized to determine a model of a specific interesting event, or this could also be distinguished through the exploitation of the spatial and temporal correlations in the measurements.

In this paper, a classification of faults that can affect sensor networks and a review of the main FDI techniques are first given in Section 2. In Section 3, the design and results of extensive simulations of a reference FDI method (Choi et al. (2009)) applied on the specific case of the area monitoring scenario are provided. In particular, the performances of both FDI and intrusion detection are quantified with a Monte Carlo approach, so as to provide recommendations for the design of surveillance sensor networks. A preliminary study has also been conducted on the combination of a static network and a fleet of mobile sensors for such a mission, which shows significant improvement of the overall performance. Conclusions are discussed in Section 4.

2. LITERATURE REVIEW

2.1 Definition of faults in a sensor network

Classifications of sensor fault types in sensor networks and of associated FDI methods have been proposed in Ni et al. (2009); Muhammed and Shaikh (2017); Mahapatro and Khilar (2013); Jurdak et al. (2011); Fang and

* `https://euroswarm.wordpress.com`

Dobson (2013); Rassam et al. (2013). These references have been analyzed to propose the following synthesis. A fault is an unpermitted deviation of at least one characteristic property or parameter of the system from acceptable/usual/standard conditions. A fault may lead to a failure, which is a permanent interruption of the system ability to perform a required function under specified operating conditions. There are two classical approaches to define faults in sensor networks, which can be more or less appropriate depending on the context. Frequently there may not be a clear explanation on the cause of a fault (e.g. outliers), and hence it may be easier to describe some faults by the characteristics of the data behaviour. This is the *data-centric* view, which is the most relevant from the diagnosis point of view. The second *system-centric* approach covers physical malfunctions or degraded operational conditions on sensors and describes what type of features this could exhibit in the data collected. These two approaches are therefore complementary: data-centric faults are usually caused by system issues. Monitoring certain aspects of the hardware, such as battery life, may thus help to understand when and why a fault may occur. The data-centric perspective has been adopted in this work, i.e. how faults can be detected by investigating the data of the sensor nodes. It should although be kept in mind that some major hardware faults that are at the root of these data faults can sometimes be self-detected by a node itself (loss of connection to the network, battery end-of-life). In this case, cooperative FDI is not mandatory and the node can just be excluded from sensor fusion for event detection or target tracking.

Data-centric faults:

(1) Outlier: very common fault represented by isolated samples in the temporal or spatial sense. To model an outlier, the most common feature to consider is the distance from other readings.
(2) Spike: a rate of change much greater than expected over a short period of time (inconsistent gradient). In some cases, this may represent an event (intrusion) that should be detected.
(3) Offset (or bias): deviation in sensed data by an additive constant.
(4) Gain (or scaling): the rate of change of the sensed data is affected by a scale factor.
(5) Drift: a time-varying signal (linear or not) is added to the measurements, independently from the measured phenomenon.
(6) Stuck-at (or locked): some series of data values present no variation over a given period of time. The monitoring of gradients and variances is relevant to detect this kind of fault.
(7) High noise: unusually high noise may be due to hardware failures, battery issues or environmental disturbances.
(8) Data loss: sensed data is not received during a significant time interval.

System-centric faults:

(1) Calibration error: this is the root cause of faulty data in many cases, leading for example to the presence of offset, gain and drift.
(2) Connection or hardware failures: this causes hard faults, with loss of data.
(3) Low battery: may influence sensor readings, creating outliers, spikes or other signal distortions.
(4) Environment out-of-range: the sensitivity range of the sensor is exceeded due to current conditions. This may be due to inappropriate calibration or a wrong choice of sensor.
(5) Clipping: this saturation phenomenon can be a consequence of environment influence or misuse of the sensor, and can lead to the *stuck-at* type of fault.

The faults can be persistent or intermittent, therefore the history of events should be stored to discard sensor nodes that repeatedly present an anomalous behaviour.

2.2 Techniques for FDI in sensor networks

Fault detection is the determination of the presence of faults in a system and of their times of occurrence. It is generally followed by fault isolation to determine the type and location of the faults. Fault identification (or estimation) aims then at determining the magnitude and time-varying behaviour of the faults. The complete process is usually called either FDI or fault detection and diagnosis (FDD), the latter including identification. These tasks generally involve the generation of residuals, which are fault indicators based on discrepancies between measurements and model-based computations. Residuals should remain small as long as there is no fault, and become sufficiently large to be noticeable whenever faults occur (Marzat et al. (2012)). Emphasis will be put on cooperative processing and detection of faults affecting the collected data. Zhang et al. (2010) consider specifically detection techniques of outliers. The following generic remarks are of particular interest:

- Two variations for outlier identification exist in sensor networks. One is that each node identifies the anomalous values only depending on its historical values. The alternative is that in addition to its own historical readings, each sensor node collects readings of its neighboring nodes to collaboratively identify the anomalous values.
- Outlier detection techniques need to make use of data of neighboring nodes and spatial similarity of the sensor data. This is based on the fact that the sensor faults are likely to be spatially unrelated, while event measurements are likely to be spatially correlated.

FDI techniques can be classified in the following three main categories.

Statistical and model-based These methods represent spatial or temporal links between data via statistical models. Different types of parametric probability distributions can be employed (Gaussian, Poisson, Rayleigh, uniform), several estimation techniques have been proposed (averaging, median, maximum likelihood, Kalman filtering, Bayesian inference) and various statistical decision tests can be applied (three-sigma test, likelihood ratio). There also exist non-parametric statistical methods using other distance measures (histograms or kernel functions).
Franco et al. (2006) described a distributed FDI methodology based on a bank of Kalman filters and a Bayesian con-

sensus decision scheme. The detection procedure is based on the analysis of residuals computed as the innovations of a Kalman filter using a nominal model of the system to be monitored; once a fault is detected, a multiple model estimation scheme is called upon to achieve fault isolation. The distributed algorithm of Ji et al. (2010) computes a spatial weighted average for each node using its pre-defined set of neighbors, with a simple decision mechanism using a fixed threshold (set to zero in the case of binary data). The weights represent the confidence associated to each sensor, depending on its prior probability of becoming faulty.

Neighborhood and clustering These methods compute some measure of distance between data instances of neighbor nodes, on spatial or temporal windows. Neighbors are defined using criteria of number, distance, connection or overlapping range. The main underlying assumption is that neighboring measurements of the same phenomenon should be correlated and of very similar values. Some basic statistical quantities (similar to those from the first class of methods) can be computed using these data, and consensus mechanisms are applied (usually majority voting, possibly weighted with some prior knowledge) to identify faulty nodes, most often in a decentralized way. This category also includes nearest-neighbor clustering rules and decision processes.
Choi et al. (2009) proposed a fault detection scheme based on the comparison of sensor readings between neighbors in a homogeneous network with deterministic measurements. An iterative procedure is designed to create clusters and determine if the node is healthy. The node degree and matching threshold are adjusted at each iteration until all the nodes have been classified as healthy or faulty. A distributed fault detection scheme for sensor networks has been proposed in Chen et al. (2006). It uses local comparisons with a modified majority voting, where each sensor node makes a decision based on comparisons between its own sensing data and those of neighbors on successive time instants, while considering the confidence level of its neighbors. This requires two successive exchanges of information between neighboring nodes to obtain a local consensus. The detection criterion was improved in Jiang (2009) in the first level of exchange to avoid misclassifying sensor nodes as faulty.

Classification and learning This set of methods builds classes from a measurement database either in a supervised way (i.e. with the help of an expert) or in a semi-supervised manner (i.e. putting in the same class elements of the database that are deemed close to one another, and relying on an expert only to label the classes). A classifier (Neural Network, SVM, Bayesian Network) is then trained with respect to these classes to assign the newly measured variables to classes that represent healthy or faulty behaviors. This category also includes spectral approaches such as Principal Component Analysis (PCA), which builds a learning model by using redundancy in the process history to predict the values of variables and generate residuals by comparing predictions to measured values. Harkat et al. (2006) used PCA to detect sensor faults and also additionally estimate their magnitude on a single node by computing the difference between a measurement of one variable and its reconstruction using the PCA projection matrix and the measurements of all the other variables. This technique was applied to supervise an air quality monitoring network.

2.3 Performance metrics

The following performance metrics are usually employed to evaluate the FDI methods:

- Detection Accuracy (DA): ratio of correctly detected faulty nodes in the total number of faulty nodes.
- False Alarm Rate (FAR): ratio of non-faulty nodes detected as faulty in the total number of fault-free nodes.
- Delay of detection (mean and variance can be computed on all nodes detected as faulty).

The trade-off between Detection Accuracy and False Alarm Rate can be represented graphically by a ROC (receiver operating characteristic) curve (see e.g. Figure 3).

2.4 Main trends and remarks

Main trends Most of the methods are based on statistical assumptions and data comparison in a neighborhood of sensor nodes, which can be easily decentralized. Decentralized approaches are either explicitly hierarchical with specific gateways and pre-defined cells in the network, or on the contrary implicitly self-organizing where each node exploits the knowledge from its neighbors without rank. The second design choice is more generic and adaptable to failures. Statistical models are very often employed, in either centralized or neighbor-based approaches. This is quite natural, regarding the usual assumptions on the spatial and temporal distributions of sensor measurements. Majority voting is probably the most employed consensus mechanism. It has the potential to enhance the detection performance from both detection accuracy and false alarm rate perspectives. It should however be noted that the performances of these techniques are worst affected by low average node degree, i.e., these techniques are topology dependent. Learning techniques are cumbersome, due to the need for a large database of data samples to train the algorithms and high computation complexity for the training step. The most used technique is to rely on a single healthy class and to detect outliers. However this strategy does not allow the identification of the type of fault.

Limitations of existing studies The majority of existing works does not take into account multivariate data and assumes the sensor data is univariate. Many of the neighbor coordination approaches only consider the spatio-temporal correlations between sensor data of neighboring nodes and ignore the dependencies among the attributes of the sensor node itself. Many of the approaches adopt a predefined threshold to detect faults. However, an optimal threshold is not always easy to determine. Rigorous analyses can be obtained under statistical assumptions. Little work has been done on diagnosing intermittent and transient faults. Most often, persistent failures are assumed. This is an important issue in practice. Little work has been done on distinguishing between events and faults. Many of existing techniques simply regard outliers as errors. Since a commonly accepted notion is that errors should be removed from the data set, important information about hidden events may be lost.

Remarks Combining static and mobile sensors for event detection is still an open issue that has not been extensively studied in the literature (a few existing references are Datta et al. (2006); Wang et al. (2007); Mohamed et al. (2013)), and collaborative fault detection in this context is all the more challenging. A simple way to address the problem would be to add (when they are available) the mobile sensor measurements along with their localization information to the static ones and use the same fusion and fault detection algorithms. Even if the dynamics of the vehicles would not be fully exploited in this framework, this would still add important information to excite the system for fault detection purposes. This is the approach evaluated in Section 3.5. The security of sensor networks (in particular robustness to attacks) is beyond the scope of this study. Some elements can be found in Chen et al. (2009), where physical protection, software and network security or secure location and data aggregation are addressed.

3. FDI IN A SURVEILLANCE SENSOR NETWORK: A CASE STUDY

In this section, the application of a FDI technique from the literature to a sensor network representative of a persistent surveillance task is analyzed through simulation results.

3.1 Scenario and simulation conditions

The simulations consider the surveillance of a sub-area of a larger protected zone. The considered dimensions are 20 m × 20 m, this metric unit is just an example, since the obtained results are not sensitive to this value and remain valid if all parameters are modified with the same scale. The sensor network is composed of a large number of cheap omnidirectional sensors with limited ranges (3 m and 4 m have been considered), which are representative of seismic sensors or magnetometers (listed as relevant for the application considered). Several measurement models of increasing complexity can be considered:

- Binary measurements (as in e.g. Arora et al. (2004)) on the full sensor range, which corresponds to a local detection by each sensor using its own information. This is a representative model for imprecise sensors (such as the seismometers and magnetometers considered).
- Position of the intruder in the sensor field of view, with possibly an associated uncertainty (which can be very large for cheap sensors). This might lead to slightly better performance than binary measurements (but this is not guaranteed).
- Spatial field mapping of the probability of presence defined at each position. This can be easily built from position measurements.

In this case study, the first measurement model has been considered for evaluating the methods. This is a representative situation in practice, and this was deemed sufficient to assess FDI performance. The surveillance mission considers intrusion detection with no prior knowledge: in the nominal case without any intruder present, a sensor should return the value 0, while it should measure 1 if an intruder is present in its field of view. No particular assumption is made on any sensor fusion algorithm for more advanced high-level recognition or combination of sensor information. The FDI algorithm will thus try to detect faulty nodes by comparing the measurements of the nodes. For this purpose, a redundancy in spatial and temporal coverage between nodes is mandatory, since this is the main assumption at the basis of all cooperative FDI approaches for sensor networks. Various percentages of faults in the network have been considered, from 0 to 50% of the nodes (randomly chosen in the simulations).

Table 1. Average numbers of neighbors depending on number of nodes and sensor ranges

Number of nodes	15	20	25	30	40
Sensor range 3 m	2.43	3.50	4.57	5.64	7.79
Sensor range 4 m	4.24	5.97	7.69	9.41	12.86

A relevant operational scenario is to consider that some faulty sensors are stuck to 1, which means an undesired false detection of the presence of a target. The two cases of such a fault occurring in absence or in presence of a true event (intrusion) are considered.

(1) Persistent faults (sensor stuck to 1) in the case of no intrusion in the area.
(2) Persistent faults (sensor stuck to 1) in the case of a simultaneous intrusion in the area. Distinguishing faults from true events is known to be a difficult issue in FDI for sensor networks (Zhang et al. (2010)).

The placement of the nodes in the area has been carried out by Latin Hypercube Sampling (LHS). This method generates a near-random space-filling design of N points by dividing each axis of the space in intervals of equal length, and then by sampling randomly with a uniform distribution in only one square per line and per column. This provides a better distribution of samples in the space than grid sampling or pure direct random methods (McKay et al. (1979)), especially in high dimensional spaces. This placement does not provide the optimal coverage but still ensures a good spreading in the area with a limited computational cost (see e.g. Balesdent and Piet-Lahanier (2015) for more advanced considerations). Two sensor nodes are defined as neighbors if their fields of view intersect. Several numbers of nodes in the area have been considered so as to obtain more or less large and dense networks (see Table 1). This yields a collection of average numbers of neighbors for each node, which is one of the main parameters with a strong impact on FDI performance and reliable intrusion detection. All the results presented have been obtained via extensive Monte-Carlo simulations (1000 for each result data point) on the network sampling so as to compute metrics that are independent of the particular disposition of the network. All the cases from Table 1 have been processed in this context. The intruder has been modelled as a random walk on a duration of 200 time steps, where each subsequent step is chosen using a zero-mean Gaussian distribution with standard deviation 0.5 m. The initial position has been randomly chosen with a uniform distribution in a rectangle centered in the middle of the area and of side equal to half the length of the space.

3.2 FDI method for performance evaluation

The method described in Choi et al. (2009) has been selected for performance evaluation on the described intrusion detection scenario, since it is representative of the decentralized methods based on neighborhood and clustering that have been extensively studied and which reported the best diagnosis performances in the literature.

This approach uses an iterative procedure to create clusters and determine if the node is healthy. The node degree and matching threshold are adjusted at each iteration until all the nodes have been classified as healthy or faulty. The method is a good summary of the main features of cooperative FDI techniques: it is decentralized (each node computes its own status using only its neighbors), spatial correlation in data is exploited, a consensus is performed between neighbors and propagated inside the network, and temporal aspects are also taken into account by modifying iteratively the table of neighbors and the detection thresholds. The detailed algorithm is as follows.

For each node i, to determine its status F_i (0: good, 1: faulty, 2: undetermined) **do**

(1) **Create** the Neighbor Table (NT) and **set** F_i to 2 (undetermined)
(2) **Gather** measurements of the neighbors of i
(3) **Determine** k_i the number of matching neighbors (same measured value in the binary case)
(4) **If** $k_i > \theta_i$ **then set** F_i to 0
(5) **Repeat** l times for undetermined nodes:
If one of its matching neighbors is determined to be healthy, **then** set F_i to 0
(6) The remaining undetermined nodes are considered to be faulty ($F_i \leftarrow 1$)
(7) **Update** NT (remove faulty sensors)
(8) **Compute** threshold θ_i for the next time step

The threshold θ_i chosen as $\theta_i = \max\left(\delta, \frac{d_i}{2}\right)$ where d_i is the number of neighbors of node i and δ the minimum number of neighbors required to make a decision. Following the recommendations of Choi et al. (2009), the value of l has been fixed to 1 and the value of δ to 2.

3.3 Persistent detection of some sensor nodes with no intrusion

In this case, all the nodes should normally return the measurement 0 to indicate that no intrusion is currently in process. Therefore, faulty nodes that return 1 should be identified with the proposed method. The results of the Monte-Carlo simulations are presented in Figure 3, for all variations of parameters (sensor range, number of nodes, and percentage of faulty nodes). The perfect performance is represented by a DA indicator equal to 1 and a FAR equal to 0. It clearly appears that the most important parameter impacting FDI performance is the average number of neighbors, which is a consequence of the relation between the total number of nodes and the sensor range. Of course, for the same number of nodes there are more neighbors per node in the case of a larger range. The results show that a good FDI detection performance (DA $>$ 0.95 and FAR $<$ 0.05) can be obtained for a percentage of faulty nodes up to 25% only if the average

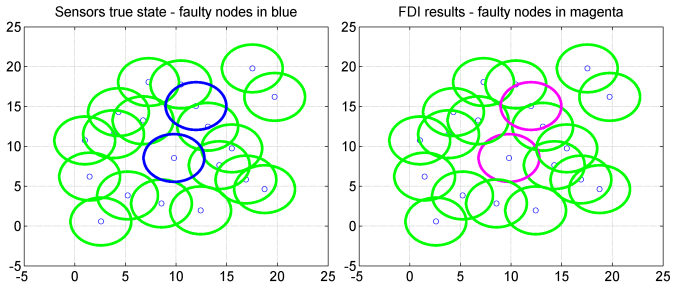

Fig. 1. Example of good FDI in a sensor network (20 nodes, 10% faults, no intrusion)

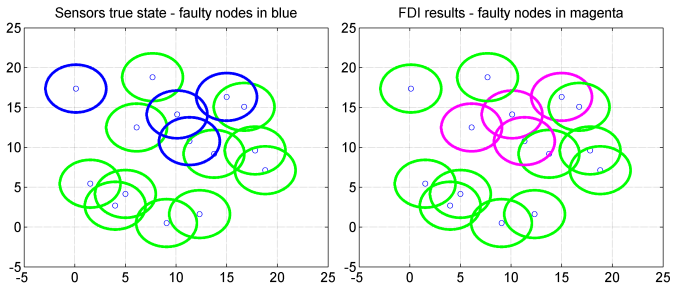

Fig. 2. Example of false alarm due to a lack of neighboring nodes (15 nodes, 25% faults, no intrusion)

number of neighbors (also known as the node degree) is larger than 4 (as a reminder, see Table 1 for the relations between those parameters). These results and the level of performance reached are consistent with those obtained in the literature with similar parameters (but not for intrusion detection scenarios, see Paradis and Han (2007); Chen et al. (2006); Ji et al. (2010)). This shows that such a method is well qualified in the targeted application, if a sufficient number of nodes can be deployed to satisfy the condition on the number of neighbors per node. For very cheap sensors, this should be possible. Moreover, the percentage of simultaneous faults should not reach very high levels since sensors can be deactivated or replaced as soon as they are detected as faulty. Figures 1 and 2 display a few examples of fault detection and identification results for different network parameters.

3.4 Persistent detection of some sensor nodes in presence of intrusion

When the same type of faults (sensor measurement persistently stuck at value 1) is considered in the presence of an intrusion, the situation is more complex and there are several possible consequences. The sensor fault can trigger false alarms in locations where there is no actual presence of an intruder and therefore result in a false detection of multiple intrusions on the overall area, which can lead to a bad use of military resources and this should be detected beforehand. If the persistent sensor fault appears where the intruder is located, it does not degrade performance of detection, but only healthy nodes can be reliably taken into account for an intrusion to be confirmed. Finally, another well-known issue with faults in sensor networks that should detect this kind of event is the difficulty to distinguish faults from intrusions, especially if there is a limited number of neighbors for each node. The simulations in the case of an intrusion provide interesting information on

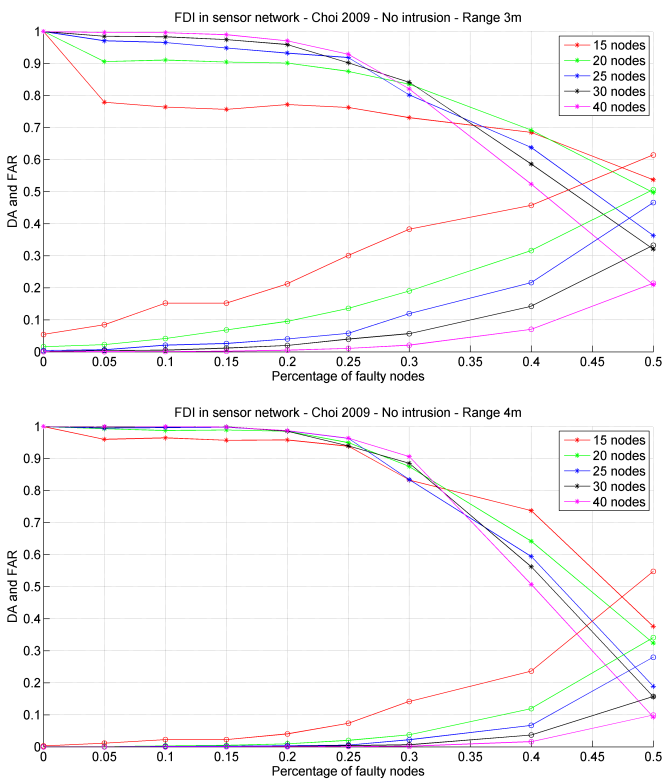

Fig. 3. Performance of FDI method with various sensor ranges and network configurations (case without intrusion, sensor ranges of 3 m (up) and 4 m (down))

all these situations. Results presented in Figure 5 show that the simultaneous occurrence of faulty nodes and intruders degrade the overall performance of FDI. Detection Accuracy slightly decreases and False Alarms are more frequent because sometimes the event is treated as a fault by the consensus mechanism, especially when there are very few neighbors per node. It is however interesting to note that the network is more robust to this kind of confusion between events and faults when the sensor range is larger, even for comparable node degrees. This is because the intrusion can be better captured by multiple neighboring nodes since their overlapping area has more chances to be larger. The FDI performance remains acceptable (DA $>$ 0.85, FAR $<$ 0.15), as long as there is still more than 4 neighbors per node.

So far, only the performance of the FDI algorithm when there are at the same time faults and actual intrusion has been discussed. However, another important operational consideration in this scenario is whether the correct detection of the intrusion can still be achieved even in the presence of such faults. The intrusion is determined to be correctly identified if a healthy node, i.e. one that has not been classified as faulty by the FDI algorithm, detects the target at least once during the time interval studied. The rates of correct intrusion detection for the same simulations have been computed and are reported in Figure 6. In the simulation conditions considered, a meaningful result is that the rate of correct intrusion detection increases with the percentage of faults. This is again mainly because in the case of only a few faults and a reduced number of nodes, events are sometimes treated as

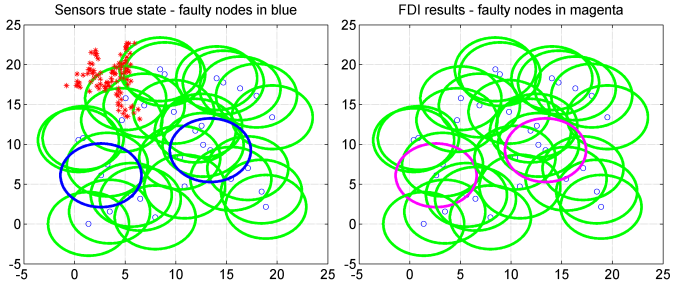

Fig. 4. Illustration of FDI results in presence of intrusion (30 nodes with 5% faulty, intruder trajectory in red)

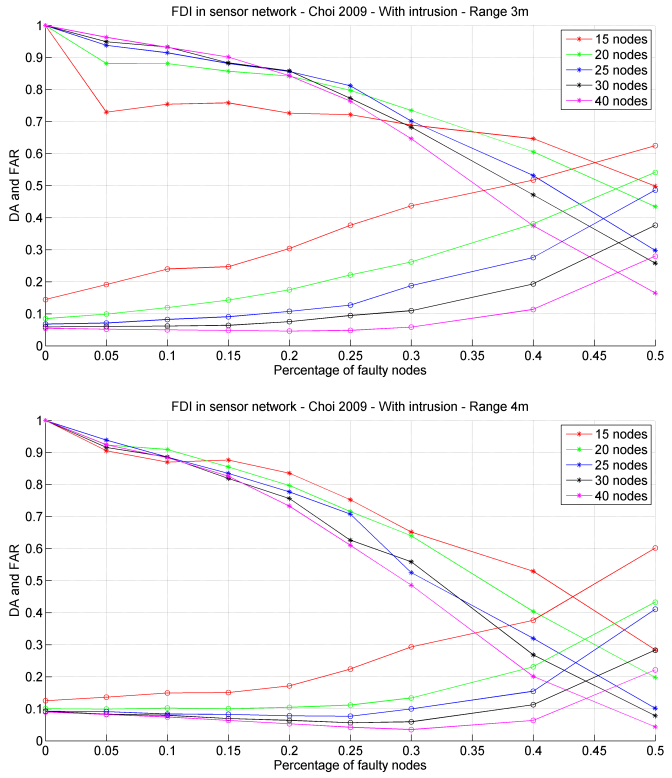

Fig. 5. Performance of FDI method with various sensor ranges and network configurations (with intrusion)

faults (in the specific case of the *stuck-at* fault considered here) and if the percentage of faults increases then some nodes are not detected as faulty and can contribute to detecting the event. When the number of nodes increases (and thus the number of neighbors), the performance of both fault detection and correct intrusion detection is significantly improved. More than 7 neighbors per node are thus required here to guarantee a sufficient level of correct intrusion detection (at least 0.8 to 0.9).

3.5 Network of static and mobile sensors

Another possible improvement to the network is to consider a fleet of mobile vehicles, equipped with the same small-range sensors (4 m range) than the static network, which follows a predefined trajectory covering the entire monitored area. These sensors communicate with the static ones (no loss or delay assumed) and can thus help to detect faults and events on the monitored area. We consider the case of a small fleet of 6 sensors, which are

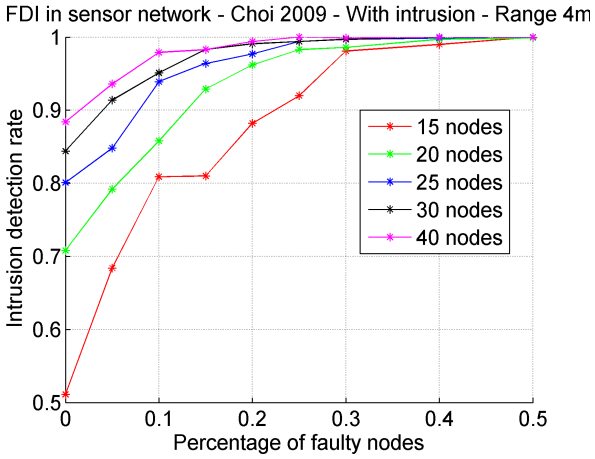

Fig. 6. Performance of intrusion detection in presence of faults

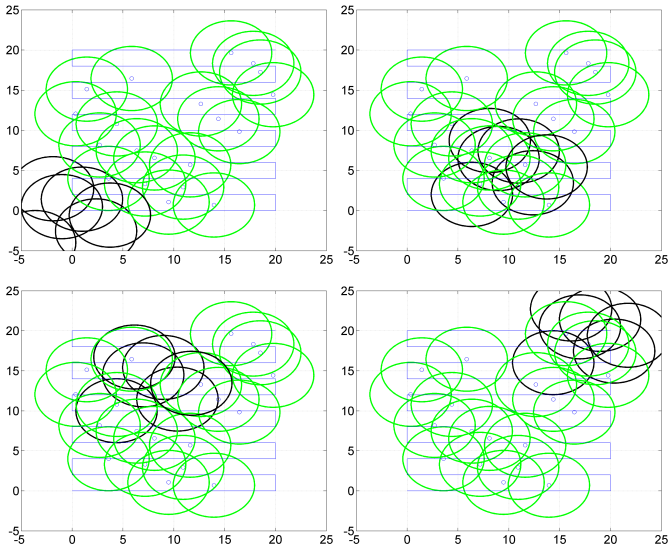

Fig. 7. Collaboration between the static sensor network (in green) and a swarm of vehicles (in black) following a reference trajectory (in blue)

non-faulty for the entire duration of the mission but which could be detected as faulty by the FDI algorithm. Figure 7 presents successive screenshots of the mission. Note that in the result figures, the number of nodes is the one of the static network.

FDI performances with simultaneous intrusion are presented in Figure 8 and Figure 9. The detection accuracy remains at a similar level than the original static network, but false alarms have been drastically reduced thanks to the mobile fleet. Regarding intrusion detection, the performance of the combined mobile and static network is excellent. The impact of employing fleets of mobile sensors together with static sensors is thus beneficial in terms of intrusion detection even in the presence of faulty sensors. What is even more interesting is that high level of performances are reached with a smaller number of static nodes (15 are sufficient here, in spite of the corresponding small static node degree). The addition of mobile sensors is thus an efficient way to increase dynamically the network node degree.

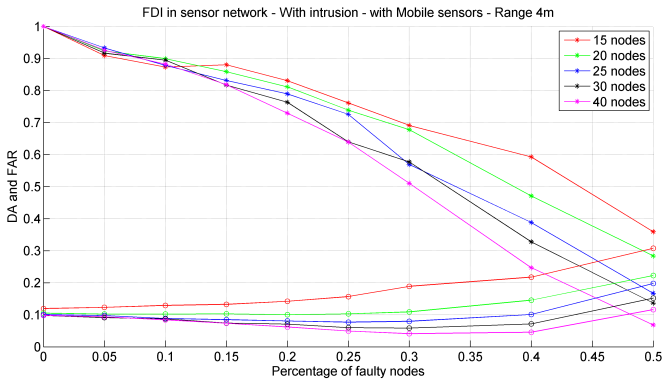

Fig. 8. FDI results in the sensor network with the help of a fleet of mobile sensors (with intrusion)

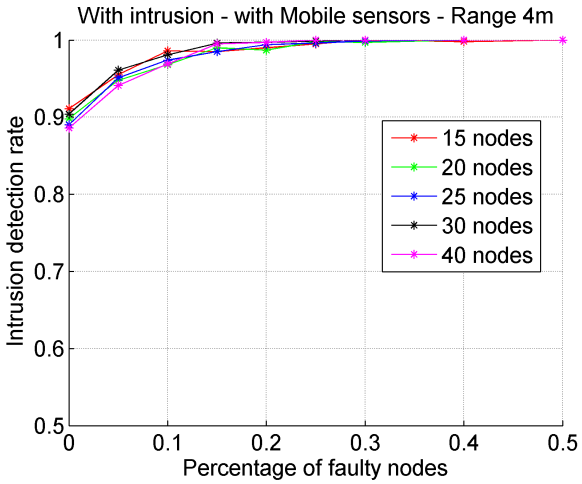

Fig. 9. Network combining static and mobile sensors - Performances of intrusion detection (to be compared with Figure 6)

4. CONCLUSIONS

This paper has presented a summary of the state of the art on cooperative fault detection and isolation (FDI) in sensor networks and original numerical evaluations. The performances of a reference method have been quantified on a simplified surveillance scenario. Extensive simulations have been conducted to assess the reachable performances of both FDI and intrusion detection in many different cases: static sensors with different node degrees (number of neighbors for each node) and a combination of mobile and static sensors. These analyses yield the following global remarks.

To implement cooperative FDI, there is a need for redundancy in the sensor network: at least 2 to 3 neighbors are required for voting schemes and much more to be able to distinguish between faults and events (i.e. the intrusion that should be detected). It has been shown that a good level of performance (high detection accuracy and few false alarms) can be attained with more than 4 neighbors per node, without simultaneous intrusion in the monitored area. Existing methods are efficient if the percentage of faulty nodes in the network is limited (less than 25%). This is a reasonable assumption since confirmed faulty

nodes can be deactivated or replaced through systematic maintenance.

To obtain a good performance of correct intrusion detection, with simultaneous faults, there should be more than 7 neighbors per node. Otherwise, if there are a lot of faults in the network, they will be detected but a true intrusion could also be classified as a fault. It should be noted that these results have been obtained without considering advanced sensor fusion techniques, which might improve fault and event detection performances and thus reduce the necessary minimum node degree. Performances of both FDI and correct intrusion detection are improved by considering both static and mobile sensors in the process, this type of architecture should thus be further investigated.

This study has obtained progress beyond the state of the art by applying a reference method for FDI in a sensor network to the specific case of intrusion detection in a surveillance mission, and by analyzing simultaneously the performance of these two distinct capabilities. Moreover, this has been extended by combining the static network with a fleet of mobile vehicles and evaluating the resulting increased performance.

ACKNOWLEDGEMENTS

This work has received funding from the European Union under the Pilot Project on Defence Research managed by the European Defence Agency under grant agreement No PP-15-INR-01 05 EuroSWARM ("Unmanned Heterogeneous Swarm of Sensor Platforms"). This work reflects only the author's views. The European Commission and European Defence Agency are not responsible for any use that may be made of the information it contains (Art. II.9 of the Model Grant Agreement).

REFERENCES

Arora, A., Dutta, P., Bapat, S., Kulathumani, V., Zhang, H., Naik, V., Mittal, V., Cao, H., Demirbas, M., and Gouda, M. (2004). A line in the sand: a wireless sensor network for target detection, classification, and tracking. *Computer Networks*, 46(5), 605–634.

Balesdent, M. and Piet-Lahanier, H. (2015). A multi-level optimization approach for the planning of heterogeneous sensor networks. In *Modelling, Computation and Optimization in Information Systems and Management Sciences*, 221–233. Springer.

Chen, J., Kher, S., and Somani, A. (2006). Distributed fault detection of wireless sensor networks. In *Proceedings of the Workshop on Dependability issues in wireless ad hoc networks and sensor networks*, 65–72.

Chen, X., Makki, K., Yen, K., and Pissinou, N. (2009). Sensor network security: a survey. *IEEE Communications Surveys & Tutorials*, 11(2).

Choi, J.Y., Yim, S.J., Huh, Y.J., and Choi, Y.H. (2009). A distributed adaptive scheme for detecting faults in wireless sensor networks. *WSEAS Transactions on Communications*, 8(2), 269–278.

Datta, S., Klinowski, C., Rudafshani, M., and Khaleque, S. (2006). Distributed localization in static and mobile sensor networks. In *Proceedings of the IEEE International Conference onWireless and Mobile Computing, Networking and Communications*, 69–76.

Fang, L. and Dobson, S. (2013). In-network sensor data modelling methods for fault detection. In *International Joint Conference on Ambient Intelligence*, 176–189.

Franco, E., Olfati-Saber, R., Parisini, T., and Polycarpou, M.M. (2006). Distributed fault diagnosis using sensor networks and consensus-based filters. In *Proceedings of the 45th IEEE Conference on Decision and Control, San Diego, CA, USA*, 386–391.

Harkat, M.F., Mourot, G., and Ragot, J. (2006). An improved PCA scheme for sensor FDI: Application to an air quality monitoring network. *Journal of Process Control*, 16(6), 625–634.

Ji, S., Yuan, S.F., Ma, T.H., and Tan, C. (2010). Distributed fault detection for wireless sensor based on weighted average. In *Proceedings of the Second International Conference on Networks Security Wireless Communications and Trusted Computing*, volume 1, 57–60.

Jiang, P. (2009). A new method for node fault detection in wireless sensor networks. *Sensors*, 9(2), 1282–1294.

Jurdak, R., Wang, X.R., Obst, O., and Valencia, P. (2011). Wireless sensor network anomalies: Diagnosis and detection strategies. In *Intelligence-Based Systems Engineering*, 309–325.

Mahapatro, A. and Khilar, P.M. (2013). Fault diagnosis in wireless sensor networks: A survey. *IEEE Communications Surveys & Tutorials*, 15(4), 2000–2026.

Marzat, J., Piet-Lahanier, H., Damongeot, F., and Walter, E. (2012). Model-based fault diagnosis for aerospace systems: a survey. *Proceedings of the IMechE, Part G: Journal of Aerospace Engineering*, 226(10), 1329–1360.

McKay, M.D., Beckman, R.J., and Conover, W.J. (1979). Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2), 239–245.

Mohamed, N., Al-Jaroodi, J., Jawhar, I., and Eid, A. (2013). Using mobile sensors to enhance coverage in linear wireless sensor networks. In *Proceedings of the 12th IEEE International Symposium on Network Computing and Applications*, 1–6.

Muhammed, T. and Shaikh, R.A. (2017). An analysis of fault detection strategies in wireless sensor networks. *Journal of Network and Computer Applications*, 78, 267–287.

Ni, K., Ramanathan, N., Chehade, M.N.H., Balzano, L., Nair, S., Zahedi, S., Kohler, E., Pottie, G., Hansen, M., and Srivastava, M. (2009). Sensor network data fault types. *ACM Transactions on Sensor Networks*, 5(3).

Paradis, L. and Han, Q. (2007). A survey of fault management in wireless sensor networks. *Journal of Network and systems management*, 15(2), 171–190.

Rassam, M.A., Zainal, A., and Maarof, M.A. (2013). Advancements of data anomaly detection research in wireless sensor networks: a survey and open issues. *Sensors*, 13(8), 10087–10122.

Wang, D., Liu, J., and Zhang, Q. (2007). Probabilistic field coverage using a hybrid network of static and mobile sensors. In *Proceedings of the Fifteenth IEEE International Workshop onQuality of Service*, 56–64.

Zhang, Y., Meratnia, N., and Havinga, P. (2010). Outlier detection techniques for wireless sensor networks: A survey. *IEEE Communications Surveys & Tutorials*, 12(2), 159–170.