DYNAMIC SENSOR COALITION FORMATION TO ASSIST THE DISTRIBUTED TRACKING OF TARGETS: APPLICATION TO WIDE-AREA SURVEILLANCE

M Briers[1], S R Maskell (QinetiQ Ltd, Malvern)
S Reece, S J Roberts (University of Oxford)
V D Dang, A Rogers, R K Dash, N R Jennings (University of Southampton)

[1] Corresponding author: mbriers@QinetiQ.com

Keywords

Distributed Information Fusion, Automated Sensor Coalition Formation, Covariance Inflation, Collaborative Sensing Systems, Wide Area Surveillance

Abstract

The protection of infrastructure and facilities within the UK is of prime importance in the current environment where terrorist threats are present. Surveillance of large areas within such facilities is a complex, manpower intensive and demanding task. To reduce the demands on manpower, new systems will need to be developed that use a mixed sensor suite associated with access to databases containing historical data and known threats. This requires fusion of mixed type data from disparate sources. The methods used for the fusion process, and the location of the fusion process, will be dependent on the data, sensor or database. The communication requirements will also be of paramount importance within the monitoring network. As computers increase in performance and reduce in cost and power consumption, there is a growing trend for more processing to be carried out locally. This raises issues of compatibility, timeliness, global awareness of the situation and distributed versus centralised control of the system. This paper presents a generic solution to the wide-area surveillance problem through the application and combination of Covariance Inflation (a distributed fusion mathematical framework that circumvents problems with data incest) with agent-based technologies (allowing the dynamic formation of sensor coalitions) to track, and potentially risk assess, targets within the region of interest. A discussion will be provided into the distributed detection and tracking of an intruding vehicle at a commercial airport to place the seemingly abstract technology into context.

1 Introduction

There is an increasing need to develop *automated* systems that can *assist* security officials with the protection of UK infrastructure. The deployment of networked sensing technology for Wide Area Surveillance (WAS) is seen as an effective mechanism for combatting (unwanted) intrusion. In this paper, we describe a novel mathematical framework that is robust to communication failures, is able to integrate with legacy systems, and has the ability to be human-interactive.

It is assumed that the area of interest (AOI) will be monitored by a heterogeneous mix of sensors, each with a differing modality and therefore ability to detect (and subsequently form a track on) a target. The nature of the distributed fusion process within the WAS system considered here refers to distributing the processing whilst creating redundancy in the communications process - not ensuring that every sensor has knowledge of all of targets within the AOI (such inference is commonly referred to as decentralised processing). Each sensor (or fusion node) is required to maintain a description (in the form of the posterior probability distribution) of a certain phenomena of interest, which may differ over different sensor types. The aim is to present a centralised fused picture of the potential threats within the wide area to the security office. It is also of paramount importance that one utilises the "human in the loop". To this end, the framework described allows operators to provide pertinent prior information. One such example would be the specification by a human operator of a perceived threat level for the targets within the AOI. Sensors are then able to form dynamically changing coalitions, thus improving the threat assessment carried out by (automatically) focusing resources on targets of high interest.

Whilst much related research has been focused upon the use of specific sensor types for WAS this paper attempts to address the problem of hierarchical distributed (multi-)sensor fusion with automated sensor coalition formation. To allow the integration of the technology with existing legacy systems, it is important that the mathematical exposition be made in full generality (removing the dependence upon specific sensor types), and that human-level interaction be possible. Much of the technical material appearing in this paper is therefore based upon the generic concepts of Bayesian (statistical) information fusion and agent-based optimisation. The novel combination of both technologies generates a system that improves legacy systems and thus increases the ability to guard against potential threats.

The format of the paper is as follows: in the following section we provide some technical detail on the theoretical concepts utilised in this publication. Then, in section 3, some initial results are described that attempt to demonstrate the mode of operation of an automated WAS system. Finally, in section 4, conclusions and directions of future work are described.

2 Technical Formalism

In this section we provide an overview of the technical concepts to be utilised within the WAS system being developed. Since we advocate such a system for a plethora of application domains with differing sensor types, no specific idiosyncracies are discussed and so the material is presented in full generality.

2.1 Consistent distributed processing

There are many sophisticated (individual) sensing technologies that are able to potentially provide a probabilistic description of either the presence, location and/or identification of a target within the its (unique) field-of-view. Many legacy security systems used in practice will be able to provide such information. It is therefore necessary to focus research on architectures that allow one to fuse such information, whilst ensuring a certain amount of communication channel redundancy in the system. Due to the inherent uncertainty in any system, it is extremely important that one manages uncertainty correctly.

2.1.1 Covariance inflation

Using the Bayesian framework we are able to quantify the degree of uncertainty in our estimate of the phenomena of interest through the posterior probability distribution. However, when fusing information from multiple sensors it is important that a certain amount of communication redundancy be built into the system to allow for sensor malfunctions and/or the creation of ad-hoc sensor networks (that is, the rapid deployment/removal of sensors). To ensure robustness in such a system, we allow for a communication sensor hierarchy, whereby (the same) information travels to its destination by more than one route. Such reuse of information at a higher level in the hierarchy does cause problems; the system becomes unintentionally overconfident. One such solution to this problem would be to introduce a protocol that meant that information was only used once. This would severely damage the generality of the system and would be difficult to integrate with existing sensors. We thus have to devise a mathematical mechanism for dealing with unknown correlations in the data.

One such technique is known as *Covariance Inflation*[1] (CInf). CInf is a novel tool for generating an outer bound for the family of covariance matrices consistent with known diagonal terms and possibly unknown but bounded off-diagonal terms. The motivation for an outer bound follows that of covariance intersection [2] which states that a conservative estimate is better than an over-confident estimate. Both the Kalman filter and covariance intersection are special cases of the covariance inflation approach to estimate fusion. In order to use CInf within loopy networks it is necessary to determine the absolute bound on the correlation between two estimates prior to fusion. In order to calculate this bound, each agent must maintain knowledge of the proportion of its own state estimate which is potentially shared with other agents in the system. This is a measure of the projection of an agent's state estimate vector onto the vector space comrpising all other agents state estimates. The absolute value of this projection is called the agent's *coupling scalar*. By combining coupling scalars associated with different state estimates an agent is able to determine a bound on the correlation between these estimates. Thus, an agent is able to fuse two potentially correlated estimates using CInf and the coupling scalars.

Compared to CI, CInf places very little extra burden on the multi-agent system communication and computation load and is thus an efficient and effective means for performing decentralised data fusion within arbitrary and loopy multi-agent networks. The reader is invited to study [1] for details of the CInf approach to data fusion.

2.1.2 Sensor teamwork

As well as providing communication redundancy in the system, another potential improvement is through the creation of dynamically evolving sensor coalitions. This is where sensors "join forces" to perform a particular given task, with the coalition being able to better solve the task than any individual sensor in isolation.

2.2 Dynamic coalition formation

Let n be the number of sensors in the system and I be the set of sensors: $I = \{1, 2, ...n\}$. Let k_i be the number of states that sensor i has (1 sleeping state and $k_i - 1$ sensing states). When a sensor i is in different sensing states, it can see different subsets of the set of targets. When a sensor is in sensing mode, it incurs an additional cost c_i.

Let m be the number of targets and T be the set of targets: $T = \{t_1, t_2, ...t_m\}$. When a coalition $C \subseteq I$ tracks a target t_k it gains a value $V(C, t_k)$ that equals the value of the fused information about target t_k, denoted $Inf(t_k)$: $V(C, t_k) = Inf(t_k)$. The problem is then for the agents to organise into a set of coalitions $\{C_1, C_2, \ldots, C_p\}$ such that the system welfare:

$$S = \sum_{j=1}^{m} V(C_j) - \sum_{i \text{ is in sensing mode}} c_i$$

is maximised.

Specifically, let $visibility(i, s_i, t_j)$ be $true$ if sensor i in state s_i can see t_j. The problem is then for each sensor i to select its state s_i such that the system welfare:

$$S = \sum_{j=1}^{m} V(\{i | visibility(i, s_i, t_j) = true\}, t_j) - \sum_{i \text{ is in sensing mode}} c_i$$

is maximised.

Now, the value of a coalition is not known before the coalition actually carries out the sensing task. As a result, we use a predicted value for the coalition which is achieved via the rough measurement carried out by the regional radars.

Note that this is a dynamic setting, specifically, the coalitional values can change (as the targets can move around) and the set of sensors can change over time (new sensors may be added into the system and existing sensors may stop to function or be removed from the system).

3 Application

In this section we describe the operation of the system within the context of the protection of individuals, assets and property on the airport airside environment. Future publications will provide a more detailed analysis of data taken from a series of trials.

3.1 Hypothesised mode of operation

An intruder is detected by the perimeter nodes, automatically creating an object of high importance within the system. Such an object therefore demands sensor resources to be focused upon it; coalitions are automatically determined using the mechanism previously described. At the same time an airport fire engine is performing foreign object detection on the runway. This has been designated as a low priority target by the human operator, ensuring that the coalitions are sensibly formed to ensure track consistency of the unwanted intruder. All of the communication between nodes is being performed via multiple routes and the data incest issues are resolved by the use of the Covariance Inflation mechanism.

4 Conclusions

In this paper we have described the (initial and) novel combination of Covariance Inflation (a distributed fusion mathematical framework that circumvents problems with data incest) with agent-based technologies (allowing the dynamic formation of sensor coalitions) to track, and potentially risk assess, targets within the region of interest. A technical description of the underpinning technology has been provided together with its initial application to the problem of wide area surveillance.

Future work is anticipated to portray the benefits of employing such an approach over existing security solutions by the deployment of the framework in differing scenarios. Preliminary results show that this WAS system outperforms individual legacy systems, thus improving the ability of security services to successfully protect infrastructure against unwanted attack.

Acknowledgements

This research was undertaken as part of the ARGUS II DARP (Defense and Aerospace Research Partnership). This is a collaborative project involving BAE SYSTEMS, QinetiQ, Rolls-Royce, Oxford University and Southampton University, and is funded by the industrial partners together with the DTI, MOD and EPSRC.

References

[1] S Reece and S Roberts. Robust, multi-vehicle mapping. In *Proceedings of the Eighth International Conference on Information Fusion (Fusion'05), Philadelphia, Pennsylvania*, 2005.

[2] S J Julier and J K Uhlmann. A non-divergent estimation algorithm in the presence of unknown correlations. In *Proceedings of the 1997 American Control Conference (ACC'97), Albuquerque, New Mexico*, pages 2369–2373, 1997.