A Geometric-Based Approach to Fault-Tolerance in Distributed Detection Using Wireless Sensor Networks

ElMoustapha Ould-Ahmed-Vall and George F. Riley and Bonnie S. Heck
School of Electrical and Computer Engineering, Georgia Institute of Technology
Atlanta, GA 30332-0250
{eouldahm,riley,bonnie.heck}@ece.gatech.edu

***Abstract*— Distributed event detection using wireless sensor networks has received growing interest in recent years. In such applications, a large number of inexpensive and unreliable sensor nodes are distributed in a geographical region to make firm and accurate local decisions about the presence or absence of specific events based on their sensor readings. However, sensor readings can be unreliable due, for example, to noise in the sensor readings or hardware failures in the devices, and may cause nodes to make erroneous local decisions. In this work-in-progress paper, we present preliminary steps leading to a geometric based approach to fault-tolerance in distributed detection using sensor networks. The idea is to track topological inconsistencies and use them to determine the presence or not of a local detection error. The work assumes a location-aware sensor network. This geometric approach is in contrast with the existing probabilistic approaches based on special correlation.**

I. Introduction

A sensor network consists of a set of sensing elements powered by batteries and collaborating to perform sensing tasks in a given environment. It may contain one or more sink nodes (base stations) to collect sensed data and relay it to a central processing and storage system. These networks have potential for use in many military and civilian applications.

One particular application that has received a growing amount of attention in the recent years is event detection [1], [2], [3]. Nodes are tasked to determine whether a particular event of interest is occurring in their sensing range. Such an event could be, for example, a volcanic eruption at specific site [2], or the presence of a specific target [3]. An event could be detected from a high value of the sensor reading, for example. Each sensor node first determines if its sensor reading indicates the presence of an event before sending this information to its neighbors or to a sink node. However, in case of failure the sensor can produce a false positive or a false negative. That is, a high reading indicating an event occurred when it did not or a low reading indicating the absence of event when one did occur.

Event detection is commonly performed using a large number of unreliable low-cost sensor nodes. These nodes can each have a very high probability of errors (misses and false-positives). It is, therefore, important to develop fault-tolerant mechanisms that can detect detection faults and take appropriate actions.

Here, we consider fault-tolerance in the context of distributed binary detection. A node n is trying to decide whether or not a specific event is present within its coverage range. A binary variable is used to code this decision, with a value of 1 when an event is detected, and a value of 0 otherwise. In its decision scheme, the node uses the sensed data obtained by its local sensor as well as the decisions at its neighboring nodes, assuming spatial correlations.

Distributed fault-tolerance for event detection using the assumption of spatial correlation was first considered in [4]. The algorithm in [4] assumes that all nodes in the network have the same detection error probability and that this rate is known prior to the deployment. These are unrealistic assumptions. In fact, a node can become faulty with time either because of a lower energy level or because of aging or unsuitable environmental or operating conditions, thereby increasing its error probability. We can also have a heterogeneous sensor network with nodes that have different operational capabilities and accuracy levels. Moreover, the proposed algorithm in [4] is not well suited for highly localized events where the event region is very small. In fact, all nodes within a node communication range are given identical weights in the decision scheme regardless of their distances.

The work in [4] has been followed by two other publications dealing with the same problem of event region detection. In reference [5], the authors provide comments on the original paper and correct some of the mistakes in the theoretic analysis section. In [6], the authors extend the model in [4] to account for the fact that sensor errors have two different sources. An error could be noise-related or coming from a sensor fault. They also discuss the choice of the appropriate neighborhood size. However, they assume again that neighboring nodes of n at any distance have the same accuracy as estimators of the real situation at n. In such a case, the failure probability of the distributed decision scheme can be reduced by increasing the neighborhood size. Again, this is an unrealistic assumption and will introduce a large number of new errors in the case of a highly localized event. In reference [6] as in [4], it is assumed that nodes all have the same probability of failure and that this probability is known prior to the deployment.

In [7], we propose a new approach that considers the case where nodes can have different failure probability levels. This allows us to handle various types of failures including noise-related failures, biased measurement, drift over time, stuck-at failures, calibration-related failures, environment-related failures, etc. This approach can be used as a general distributed fault-tolerance mechanism for any application where nodes may have different accuracy levels. These differences can result from different locations, heterogeneous operating conditions (different sensors, different hardware conditions), different deployment times, etc.

Here, we introduce the preliminary steps to a geometric based approach to provide fault-tolerance for distributed detection. This approach contrasts with the probability-based approaches presented in the previous paragraphs. It is expected that decision errors can be detected, most of the time, by monitoring topological inconsistencies after receiving detection decisions from all the neighboring nodes.

II. Preliminary Steps Towards a Geometric Based Approach for For Fault-Tolerance in Distributed Detection

In this section, we provide preliminary observations that will be developed in future work leading to a geometric based approach to fault-tolerance in distributed detection using sensor networks. The idea is to track topological inconsistencies and use them to determine the presence or not of a local detection error. The following observations assume a location-aware sensor network where a node knows its location and can learn about the locations of its neighbors. We also assume a convex and continuous event-region.

Under these assumptions, the following observations can be made:

1) If all neighbors of a node n are detecting an event and n is not detecting this event, then there must be a detection error at one or several neighbors of n. For example in figure 1, when all nodes except n detect the event, this could indicate that the node n is faulty while its neighbors are not. This observation comes from the fact that because of convexity the event cannot be present at all neighbors of n and not at n itself. This observation can be used to allow nodes to detect their failures with a high probability.
2) If an event is detected by a node n and all of its neighbors except one interior neighbor, there must be a local detection error at one or several neighbors of n. This situation occurs, for example, if n and all its neighbors except node 8 in figure 1 detect the event. This observation is again due to the assumptions of convexity and continuity of the event region. It can be used by the node n to detect errors at its neighbors.
3) It can be proven that if 3 nodes $n1$, $n2$ and $n3$, as in figure 2 detect the presence an event, then any node in the interior triangle ($X1X2X3$ must detect this event. In a similar way, if these three nodes do not detect an event, a node in the interior triangle should not detect it. The interior triangle is formed as follows. We define the point $Oj, j \in \{1, 2, 3\}$ as the location of node j. The circle C_j is the circle centered at Oj with a radius equal to the sensing range. We define the point $X_{ij}, i, j \in \{1, 2, 3\}$ as the intersection the circle C_i and the line $OiOj$. There two such points on each circle as in the figure 2. The point $Xj, j \in \{1, 2, 3\}$ is defined as the intersection of the two tangents of the circle C_j at the two points $X_{ji}, i \in \{1, 2, 3\} \setminus j$. The interior triangle is the triangle defined by the points $X1X2X3$. It can be proven that if the three nodes $nj, j \in \{1, 2, 3\}$ have the same local detection decision, then any node inside this triangle must have this same decision when no error is present. It is of importance that this result does not require that the three trusted nodes be within a minimum distance of each other in the case of a continuous event-region. This observation can be used to place a number of "trusted" nodes in strategic locations in the deployment region so that every node in the network can detect when it is faulty. We should note here that if a regular (non-trusted) node is surrounded by three trusted nodes, it can be viewed as a trusted node after verifying its decision by comparing it with the ones of the three trusted neighbors.

These observations will be used as a starting point to develop a more rigorous geometric approach to fault-tolerance in distributed detection applications of sensor networks. We will also investigate the implication of relaxing the convexity assumption.

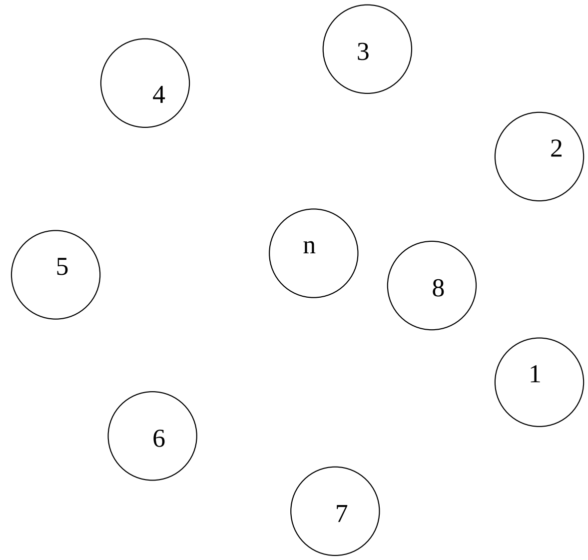

Fig. 1. All nodes except n or 8 detect the event

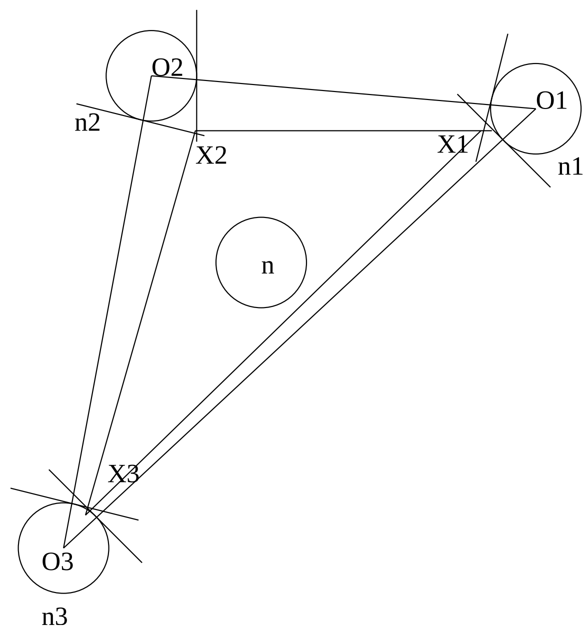

Fig. 2. Node n in the interior triangle of three trusted nodes

Acknowledgment

This work is supported in part by NSF under contract numbers ANI-9977544, ANI-0136969, ANI-0240477, ECS-0225417, CCR 0209179.

References

[1] J. Chamberland and V. V. Veeravalli, "Distributed detection in sensor networks," *IEEE on Signal Processing, Vol.51, No.2*, 2003.

[2] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh, "Monitoring volcanic eruptions with a wireless sensor network," in *Proceedings of European Workshop on Sensor Networks (EWSN'05)*, 2005.

[3] R. R. Brooks, C. Griffin, and D. Friedlander, "Self-organized distributed sensor network entity tracking," *International Journal of High Performance Computing Applications, Vol. 16, No. 3*, 2002.

[4] B. Krishnamachari and S. Iyengar, "Distributed bayesian algorithms for fault-tolerant event region detection in wireless sensor networks," *IEEE Transactions on Computers, Vol.53, No.3*, March 2004.

[5] Q. Chen, K.-Y. Lam, and P. Fan, "Comments on "distributed bayesian algorithms for fault-tolerant event region detection in wireless sensor networks"," *IEEE Transactions on Computers, Vol.54, No.9*, September 2005.

[6] X. Luo, M. Dong, and Y. Huang, "On distributed fault-tolerant detection in wireless sensor networks," *IEEE Transactions on Neural Networks, to appear*, 2005.

[7] E. Ould-Ahmed-Vall, G. F. Riley, and B. S. Heck, "A distributed fault-tolerant algorithm for event detection using heterogeneous wireless sensor networks," in *Proccedings of the 45th IEEE Conference on Decision and Control (CDC'06), under review*, 2006.