Computer Communications 75 (2016) 97–107

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Exploiting skeletonization to restore connectivity in a wireless sensor network

Yatish K. Joshi, Mohamed Younis*

Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, Baltimore, MD 21250, USA

ARTICLE INFO

Article history:
Received 20 May 2015
Revised 22 July 2015
Accepted 24 July 2015
Available online 30 July 2015

Keywords:
Topology repair
Medial axis
Straight skeleton
Fault recovery
Connectivity restoration

ABSTRACT

Recent technological advances and an increasing thrust toward automation have resulted in a rapid adoption of Wireless Sensor Networks (WSNs) as the de facto choice in monitoring and surveillance applications. Their low cost, versatility and ability to operate around the clock in inhospitable terrains without endangering human life make WSNs ideal for applications like space exploration, environmental monitoring and combat zone surveillance. In these applications WSNs are to operate autonomously for prolonged durations; thus self-healing from failures becomes a requirement to ensure robustness through sustained network connectivity. The paucity of resources makes node repositioning the method of choice to recover from failures that partition the network into numerous disjoint segments. In this paper we present a Geometric Skeleton based Reconnection approach (GSR) that exploits the shape of the deployment area in order to restore connectivity to a partitioned WSN in a distributed manner. GSR decomposes the deployment area into its corresponding two dimensional skeleton outline, along which mobile relays are populated by the surviving disjoint segments to reestablish connectivity. The performance of GSR is validated through mathematical analysis and simulation.

1. Introduction

Decreasing costs and increasing functionality of embedded computation and communication devices have made Wireless Sensor Networks (WSNs) attractive for applications that serve in harsh environments like planetary exploration, border surveillance, environmental monitoring and military surveillance. In these applications, network formation is ad-hoc in nature; for example a swarm of sensor nodes are dropped aerially in the area of interest. After deployment, these nodes are expected to discover and establish communication links with other nodes around them and form a connected network. The limited processing and communication capabilities of the deployed nodes require them to collaborate with one another to carry out application specific tasks. Thus maintaining a connected network topology is of paramount importance for the functioning of a network throughout its lifetime.

The harsh operating environment, the inherent danger in the application area, e.g., bombs in a battlefield, and limited onboard energy supply increase the probability of node failure. Some failures can cause a loss of connectivity and potentially partition the network into disjoint segments. Basically the failure of a single node can cause a network to split into disjoint blocks if such a node serves as a cut-vertex in the network topology [1]. A similar, yet more difficult scenario is when multiple collocated nodes get damaged by an external event, e.g., an explosion, flooding, sand storms, etc. Given the importance of data sharing in achieving the application goals, sustaining connectivity is critical for network operation. Therefore a network must have the ability to tolerate the occasional failure of nodes and restore connectivity without relying on external resources, e.g., remote command center, to coordinate recovery.

Tolerance of failure: Strategies for failure recovery depend on the scope of failure and the node capability [1]. The scope of failure is defined by the multiplicity of affected nodes and their location. The failure of a single node is the easiest to handle. However, the failure of multiple nodes is a major challenge, particularly when the nodes are collocated as a major void is caused and the network becomes fragmented into disjoint segments. In addition, unlike the failure of a single node or even the failure of multiple dispersed nodes, it is difficult to determine the scale of the damage if multi-collocated nodes fail. Basically, a healthy node will not be able to determine whether it lost contact with other parts of the network because a single neighbor failed, i.e., a node that acts as a cut vertex in the topology is lost, or due to the failure of multiple collocated nodes.

Tolerance mechanisms can be classified as proactive or reactive. The former is based on provisioning redundant resources at network setup in order to mitigate failure, e.g., by establishing a k-connected topology [2,3], or providing backups for faulty nodes [4,5]. Obviously

* Corresponding author. Tel.: +1 410 455 3968; fax: +1 410 455 3969.
E-mail addresses: yjoshi1@umbc.edu (Y.K. Joshi), younis@cs.umbc.edu, younis@umbc.edu (M. Younis).

http://dx.doi.org/10.1016/j.comcom.2015.07.022
0140-3664/

such methodology is not suitable for ad-hoc networks since the randomness of deployment, e.g., aerial deployment of nodes cannot guarantee node placement with the required precision and would thus require excessive resources, e.g., number of redundant nodes, to achieve k-connectivity.

Reactive strategies are based on reconfiguring the network topology to deal with failure. They restore connectivity by forming an inter-segment topology or by regrouping surviving nodes. Reactive approaches can be broken down into two main classes centralized and distributed. A centralized approach assumes knowledge of the entire network state. This global state information is utilized to optimally allocate the existing resources and coordinate recovery. Meanwhile distributed recovery schemes operate based on local state and trade off optimality in terms of desired topology features for reduced coordination overhead and responsiveness. While quite a few distributed algorithms have been proposed for tolerating single node [6,7] and non-collocated multi-node failures [8], most solutions for tolerating the failure of multiple collocated nodes are centralized and often pursue the placement of additional relays to form a connected inter-segment topology [1]. While ideal for handling simultaneous multi-node failures, a centralized approach needs to know the scope of failure, locations of surviving segments and nodes before it can begin recovery. Thus, the applicability of a centralized approach depends on the availability of external sources, i.e., satellite imagery, aerial support from drones or UAVs to collect and disseminate global network state information on demand. Such external support may not be available at all times in ad-hoc WSNs due to the harsh operating environment, resource constraints or technical difficulties. This makes distributed approaches the most suitable choice for recovery.

Distributed reactive strategies utilize the surviving nodes to recover from node failures. Most published recovery schemes in this category, e.g. [6–8], can deal with only a single node failure or multiple non-collocated failures as they rely on the neighbors of a failed node to restore connectivity by utilizing 1-hop or 2-hop information to either move inwards in the direction of failure until connectivity is restored or by moving a redundant node to the failed node's location. These schemes however do not scale if the failure spans multiple collocated nodes since the surviving nodes do not have enough information to determine the scope of failure and appropriately plan the recovery. Distributed approaches like DarDs [9] and DORMS [10] handle recovery from multiple collocated node failures by provisioning a common meeting point before failure that serves a point of convergence for all disjoint segments and is used to restore initial connectivity.

Contribution: We present a novel distributed Geometric Skeleton based Reconnection (GSR) approach that restores connectivity after the failure of multiple collocated nodes. Our approach is motivated by the fact that the WSN topology affects its operation and in practice it is influenced by the terrain and environmental factors at time of deployment. Unlike other published schemes we aim to factor in the pre-failure network topology in our recovery strategy and provision a plan that can be implemented after failure. We argue that the use of a common meeting point imposes increased overhead and may slow down convergence. Using a geometric skeleton will enable efficient handling of failure in any part of the network by allowing surviving nodes from the individual segments to reach each other faster and facilitate localized tolerance of failures that are far from any common meeting point. It is worth noting that GSR may be applied in some applications, which can be viewed as mobile ad-hoc networks (MANET). An example of that includes a networked set of robots that participate in a combat or reconnaissance mission. We argue, however, that most MANETs involve coordination among the nodes to deal with broken links caused by node mobility and not due to the failure of multiple collocated nodes.

Fig. 1 illustrates our strategy. Given an ad-hoc network Fig. 1(a) we utilize the WSN boundary to decompose the network into its constituent geometric skeleton Fig. 1(b). A geometric skeleton is a descriptor that decomposes a shape into its essential support structure, e.g., bone structure in humans, based on how components are connected. This skeleton is stored by nodes within the network and serves as a backbone, along which mobile nodes can be deployed by disjoint segments after failure, Fig. 1(c), in order to find other survivors and reestablish network connectivity, as shown in Fig. 1(d). Thus by exploiting the pre-failure network topology, GSR provisions a recovery plan that can be independently implemented by the surviving segments to restore connectivity in the network while reducing the recovery overhead. We consider two types of skeletons, namely the straight skeleton [12,13] and the medial axis [14] described in later sections, categorize their impact on recovery, highlight their differences and provide guidelines for choosing the best skeleton for GSR. GSR is validated through extensive simulation experiments and is shown to outperform competing schemes both in terms of the travel distance overhead and the number of involved nodes. The paper is organized as follows. The next section sets GSR apart from existing solutions in the literature. GSR consists of two phases; the first takes place before a failure takes place and is described in Section 3; the second phase is for restoring connectivity in reaction to failure and is detailed in Section 4. The performance of GSR is analyzed in Section 5 and is validated through simulation in Section 6. Finally, the paper is concluded in Section 7.

2. Related work

As pointed out earlier, strategies for tolerating node failure can be classified based on the scope of failure into single node and multi-node failures. The latter can be further categorized into collocated and non-collocated failures. To tolerate a single node failure, most distributed approaches in the literature pursue relocation of nodes and rely on local neighborhood information stored by neighbors of the failed node to initiate recovery [6,7]. When multiple non-collated nodes fail, these techniques may cause resource conflicts, e.g., engage a node in more than one failure recovery. Some approaches, e.g. [8], avoid resource conflicts by synchronizing the various recovery actions. However even these approaches cannot be scaled to handle multiple collocated node failures, since a node would have to maintain state information spanning many hops, in fact it should have the entire network state available in order to avoid conflicts. Thus, they are deemed ineffective as the messaging and storage overhead required to maintain a multi-hop network state increases exponentially with network size. Given the focus of the paper, the remaining part of this section covers tolerance of multiple collocated node failures. A survey and detailed analysis for the recovery schemes for single and multiple non-collocated node failures can be found in [1].

Strategies for repairing a network topology after multiple collocated nodes failure can be classified into centralized [16–18] and distributed [9,10,19]. Centralized schemes utilize relays to form a connected inter-segment topology and re-establish communication paths between the disjoint segments. These relays may be new nodes or simply existing nodes whose repositioning does not seriously impact the intra-segment topology. Since the entire network state is factored in, centralized recovery schemes provide the best solution in terms of metrics like the number of relays deployed and the total distance traveled by them during the recovery process, if existing relays relocate as part of the solution. However, centralized schemes cannot be applied in all scenarios since their applicability is dependent on the entire network state being known after failure, which may not be feasible as pointed out earlier.

Distributed schemes are the solutions of choice for autonomously-operated ad-hoc WSNs, e.g., those serving remote or inhospitable areas such as space exploration or military reconnaissance. The general methodology in this case is to utilize mobile nodes that exist in the network. Published schemes can be classified based on the node mix