Pervasive and Mobile Computing 9 (2013) 311–323

Contents lists available at SciVerse ScienceDirect

Pervasive and Mobile Computing

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

Fast track article

Reliable networks with unreliable sensors[☆,☆☆]

Srikanth Sastry [a,*], Tsvetomira Radeva [a], Jianer Chen [b], Jennifer L. Welch [b]

[a] Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[b] Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77840, USA

ARTICLE INFO

Article history:
Available online 3 March 2012

Keywords:
Sensor network
Connectivity
Fault tolerance

ABSTRACT

Wireless sensor networks (WSNs) deployed in hostile environments suffer from a high rate of node failure. We investigate the effect of such failure rate on network connectivity. We provide a formal analysis that establishes the relationship between node density, network size, failure probability, and network connectivity. We show that large networks can maintain connectivity despite a significantly high probability of node failure. We derive mathematical functions that provide lower bounds on network connectivity in WSNs. We compute these functions for some realistic values of node reliability, area covered by the network, and node density, to show that, for instance, networks with over a million nodes can maintain connectivity with a probability exceeding 95% despite node failure probability exceeding 53%.

1. Introduction

Wireless Sensor Networks (WSNs) [1] are being used in a variety of applications ranging from volcanology [2] and habitat monitoring [3] to military surveillance [4]. Often, in such deployments, premature uncontrolled node crashes are common. The reasons for this include, but are not limited to, hostility of the environment (like extreme temperature, humidity, soil acidity, and such), node fragility (especially if the nodes are deployed from the air on to the ground), and the quality control in the manufacturing of the sensors. Consequently, crash fault tolerance becomes a necessity (not just a desirable feature) in WSNs. Typically, a sufficiently dense node distribution with redundancy in connectivity and coverage provides the necessary fault tolerance. In this paper, we analyze the connectivity fault tolerance of such large scale sensor networks and show how, despite high unreliability, flaky sensors can build robust networks.

The results in this paper address the following questions: given a static WSN deployment (of up to a few million nodes) where (a) the node density is D nodes per unit area, (b) the area of the region is Z units, and (c) each node can fail[1] with an independent and uniform probability ρ: what is the probability P that the network is connected (that is, the network is not partitioned)? What is the relationship between P, ρ, D, and Z?

Motivation. The foregoing questions are of significant practical interest. A typical specification for designing a WSN is the area of coverage, an upper bound on the (financial) cost, and guarantees on connectivity (and coverage). High reliability sensor nodes offer better guarantees on connectivity but also increase the cost. An alternative is to reduce the costs by using less reliable nodes, but the requisite guarantees on connectivity might necessitate greater node density (that is, greater number

☆ A preliminary version of this paper was presented at the 12th International Conference on Distributed Computing and Networking (ICDCN) in 2011.

☆☆ This work was supported in part by NSF grant 0964696, CCF-0830455, CCF-0937274, and CCF-0917288, and Texas Higher Education Coordinating Board grant NHARP 000512-0130-2007. This work is also partially supported by Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

* Corresponding author.
E-mail addresses: sastry@csail.mit.edu (S. Sastry), radeva@csail.mit.edu (T. Radeva), chen@cse.tamu.edu (J. Chen), welch@cse.tamu.edu (J.L. Welch).

1 A node is said to fail if it crashes prior to its intended lifetime. See Section 3 for details.

1574-1192/$ – see front matter
doi:10.1016/j.pmcj.2012.02.004

of nodes per unit area), which again increases the cost. As a network designer, it is desirable to have a function that accepts, as input, the specifications of a WSN and outputs feasible and appropriate design choices. We derive the elements of such a function and demonstrate their use.

Contribution. This paper has three main contributions. First, we formalize and prove the intuitive conjecture that as node reliability and/or node density of a WSN increases, the probability of connectivity also increases. We provide a probabilistic analysis for the relationship between node reliability (ρ), node density (D), area of the WSN region (Z), and the probability of network connectivity (P); we provide lower bounds for P as a function of ρ, D, and Z.

Second, we provide concrete lower bounds for expected connectivity probability for various reasonable values of ρ, D, and Z.

Third, we use a novel technique for network analysis which, to our knowledge, has not been utilized for wireless sensor networks before. The approach, model, and proof techniques themselves may be of independent interest.

Organization. The rest of this paper is organized as follows: the related work is described next in Section 2. The system model assumptions are discussed in Section 3. The methodology includes tiling the plane with regular hexagons. The analysis and results in this paper use a topological object called a *level-z polyhex* that is derived from a regular hexagon. The level-z polyhex is introduced in Section 3. Section 4 introduces the notion of *level-z connectedness* of an arbitrary WSN region. Section 5 uses this notion of level-z connectedness to formally establish the relationship between P, ρ, D, and Z. Finally, Section 6 provides lower bounds on connectivity for various values of ρ, D, and Z.

2. Related work

There is a significant body of work on topological issues associated with WSNs [5]. These issues are discussed in the context of coverage [6], connectivity [7], and routing [8].

The results in [7] focus on characterizing the fault tolerance of sensor networks by establishing the k-connectivity of a WSN. However, such characterization results in a poor lower bound of $k-1$ on the fault tolerance simply because the failure of some specific sets of k nodes partitions the network. Unfortunately, this is an extreme case behavior, and it fails to characterize the expected probability of network partitioning in practical deployments.

The results in [9–13] establish and explore the relationship between coverage and connectivity. The results in [11,12] show that in large sensor networks if the communication radius r_c is at least twice the coverage radius r_s, then complete coverage of a convex area implies connectivity among the working set of nodes. In [10], Bai et al. explore the relationship between coverage and connectivity if r_c/r_s is less than 2; they establish optimal coverage and connectivity in regular patterns including square grids and hexagonal lattice. However, maintaining connectivity in such scenarios requires deployment of additional sensors in periodic 'strips' across the region. The ratio r_c/r_s is weakened further in [9] to show that if $r_c/r_s = 1$ then, even if each node is highly unreliable, for large networks in a square region we can still maintain connectivity with coverage; however, as node failure probability increases, connectivity does not imply coverage. Ammari et al., extend these results in [13] with a focus on k-coverage: they show that if $r_c/r_s = 1$ in a k-covered homogeneous WSN, then the network fault tolerance is given by $4r_c(r_c + r_s)k/r_s^2 - 1$ as long as the entire neighborhood of any sensor does not fail at the same time. Another related result is [14] which shows that if a uniform random deployment of sensors in a WSN covers an entire area and $1 \leq r_c/r_s \leq 2$, then the probability of maintaining connectivity approaches 1 as r_c/r_s approaches 2.

A closely related work [15] explores the relationship among node density, transmission range, and k-connectivity in WSNs where the nodes are distributed uniformly at random. However, the results in [15] are applicable only for circular regions and they do not consider node failures in their analysis.

Our work differs from the works cited above in three aspects: (a) we focus exclusively on maintaining connectivity (and we are agnostic to coverage), (b) while the results in [9–12] apply to specific deployment patterns or shape of a region, our results and methodology can be applied to any arbitrary region and any homogeneous deployment, and (c) our analysis is probabilistic insofar as node crashes are assumed to be independent random events and we assess the probability of maintaining connectivity despite such crashes; we focus on the probability of network connectivity in the average case instead of the worst case.

Apart from static analysis of coverage and connectivity, a significant body of work focuses on dynamic protocols to maintain coverage and/or connectivity by coordinating nodes to remain either active or asleep. Such schedules extend the lifetime of a network and reduce the overall power consumption. Examples of such protocols include AFECA [16], Naps [17], GAF [18], Span [19], ASCENT [20], PEAS [21], and partial clustering [22]. Our work, although related to the above, is orthogonal. Each of the above protocols behaves correctly only when the node distribution in the WSN is "adequately redundant"; we provide a quantitative measure of such a "adequate redundancy" by providing lower bounds on the probability of connectivity. Moreover, our results are obtained through mathematical analysis, instead of simulation and experimentation. Therefore, unlike simulation-based results, which are sensitive to the fidelity of the simulated runs to the real-world behavior, our results are robust and applicable to all homogeneous WSN deployments.

We assume that the node distribution in a WSN is homogeneous and the region is tiled by hexagons such that nodes in a given hexagon can communicate with the nodes in neighboring hexagons. Although a similar analysis could be done by tiling the region with squares (or any other polygon that tiles the plane), we chose a hexagonal tessellation for the following reason. Traditionally, the communication and sensing range of wireless sensors is approximated by a circle, and among the set of regular polygons that tile a plane, hexagons are the closest approximation to a circle.