JAM: A Jammed-Area Mapping Service for Sensor Networks

Anthony D. Wood, John A. Stankovic, and Sang H. Son
Department of Computer Science
University of Virginia
{*wood,stankovic,son*}*@cs.virginia.edu*

Abstract

Preventing denial-of-service attacks in wireless sensor networks is difficult primarily because of the limited resources available to network nodes and the ease with which attacks are perpetrated. Rather than jeopardize design requirements which call for simple, inexpensive, mass-producible devices, we propose a coping strategy that detects and maps jammed regions. We describe a mapping protocol for nodes that surround a jammer which allows network applications to reason about the region as an entity, rather than as a collection of broken links and congested nodes. This solution is enabled by a set of design principles: loose group semantics, eager eavesdropping, suprenacy of local information, robustness to packet loss and failure, and early use of results. Performance results show that regions can be mapped in 1 – 5 seconds, fast enough for real-time response. With a moderately connected network, the protocol is robust to failure rates as high as 25 percent.

1. Introduction

Wireless sensor networks (WSNs) are a continuation of the evolution of networks toward larger-scale, distributed computing. In contrast to mobile ad hoc networks (MANETs), which comprise the growing number of commercial handheld, cellular-aided, and laptop computing platforms, WSNs are made up of mostly small sensors with limited-resources and capabilities.

They are more likely to focus on a particular application rather than supporting general-purpose computation and communication. Environmental monitoring, battlefield intelligence, emergency response support, and real-time data fusion and collection are among their frequently cited applications [2, 12, 14, 8].

While WSNs benefit from a more cohesive design, purpose, and management than MANETs, they suffer from relative resource impoverishment. The large number of devices needed to provide fine-grained sensor coverage of areas measured in square kilometers dictates that each device be as simple and inexpensive as possible while still providing useful functionality.

Each sensor has an omnidirectional radio, small battery, one or more sensors, and may be environment and tamper-proofed. Individual sensors are not reliable due to mass manufacturing defects, harsh natural or urban deployments, and battery death. Aggregate behavior and robust algorithms provide the reliability that safety-critical applications demand.

1.1. Role of geographic location

Mobility in these networks may be defined in terms of the agents moving through them, rather than by the autonomous post-deployment movement of individual sensors. A number of localization services exist which can provide each sensor with an accurate estimate of its own location and provide directory services for others' locations [18, 4, 22, 10] without the expense of adding GPS capability to each node.

As WSNs are embedded in the environment, this geographical information is much more relevant and useful than for traditional wired networks. Here the topology is defined by physical proximity, even though real-world propagation and radio hardware produce one-way links [9], "gray-zones" [19], and other challenges to protocol design and simulation.

Many of the events and problems that occur in deployed WSNs exhibit strong spatiotemporal properties. Tracking intruders, vehicles, or animals explicitly provides information about where an entity is across time. Explosions leave behind a void in the network that is limited in area. Barriers too have physical properties and defined boundaries, whether walls, roads, or rivers. Even fires, which may eventually spread into large regions and leave behind pockets of conflagration, have boundaries that are dictated by physical processes.

1.2. A denial-of-service attack

In a network that is mostly homogeneous, that is, where there is little capability or functional differentiation except what is cooperative, distributed, or re-

dundant, deliberate attacks may also be strongly localized. One such attack, most likely to occur in a battlefield or urban warfare environment, is radio *jamming*.

The extent of the jamming is dictated by physical properties such as the available power, antenna design, obstacles and height above ground. Jamming is no different than normal radio propagation, except that it is unwanted and disruptive, creating a denial-of-service condition. It experiences the same probabilistic and transient propagation behavior noted in [9] and [19].

Sensor nodes inside a jammed region cannot effectively accomplish any aspect of their mission which depends on communication. The network at large may waste energy and cause further contention by trying to query or use the affected regions for routing messages. Unless accidental, as from a malfunctioning sensor node, whatever is causing the jamming may pose a hazard to sensor-network-supported human agents. Network-directed vehicles entering the region will be unable to communicate and may become stranded.

Defeating or avoiding jamming is a complicated game of one-upmanship, with the complexity and cost escalating with each counter-measure, counter-counter-measure, etc. Spread spectrum techniques such as frequency hopping and code-spreading [3, 24] are common defenses against both intentional jamming and high-noise environments. Other measures include antennas with steerable or adaptive nulls and multiple-element antenna arrays.

The costs and complexity of these solutions are prohibitive for WSNs, in which individual nodes must be cheap. Yet, jamming is particularly easy since many will use single-frequency communication.

If jamming is only a problem in military networks, perhaps the expense of prevention can be justified. While standoff and localized jamming *are* a concern in a battlefield context, jamming may also occur in commercial and industrial networks. With the promulgation of mobile handheld devices that also operate in the unlicensed Industrial, Scientific, and Medical (ISM) band, it may even be accidental. Neither must the attack be jamming in the traditional sense; other low-energy mechanisms cause similar denials-of-service [26].

Given the frequency of denial-of-service activity on the Internet, we expect to see more of this kind of attack as WSNs become more commonly and more accessibly deployed.

Though it may not be feasible to cheaply *prevent* jamming, if the network can know the location and shape of the affected area, it can mitigate its impact. Assuming that the entire network is not affected (in which case there is no hope), it can take action to avoid the area for routing and higher-layer route planning, tune energy management protocols, and report the problem to uplink control systems.

1.3. Detection: A mapping approach

We propose a mapping service for WSNs that can provide the following benefits:

- Feedback to routing and directory services
- An effective abstraction at a higher-level than local congestion, failed neighbors, and broken routes,
- Support for avoiding the region by network-controlled vehicles, military assets, emergency personnel,
- Reports to a base-station for further jamming localization, and
- Aid to power management strategies for nodes inside and around jammed regions.

The jamming detection and mapping protocol use mostly existing data and facilities in the typical sensor communication stack, making detection and mitigation a cheaper strategy than prevention.

Generally, nodes near the border of a jammed region notify neighbors outside the region of jamming. These nodes form groups and use a lightweight, low-state management mechanism to coalesce groups and map the extent of the jammed region. Bridge members aid neighborhoods of low connectivity. An eager eavesdropping strategy provides forward and backward information diffusion among mapping members.

Contributions of this work include:

- Loose group semantics integrated with flooding and eager eavesdropping to quickly build a map of the region of interest,
- Cross-layer (MAC, routing, application) interaction to provide a useful service,
- Analysis of performance in medium-scale simulation and with failures,
- Carrier-sense defeating mechanism for high-priority message delivery, and
- Analysis of the tradeoff of time versus the amount of the region known by a portion of the group.

In the remainder of the paper we describe the Mapping Service in greater detail, including how to detect jamming and design principles employed. Then we develop evaluation criteria and show results from extensive simulation experiments. Finally, we conclude with related and future work.

2. Mapping service

Two primary components form the basis of the mapping service, shown in Figure 1: a jamming detection module, and a mapping module. Both operate on every node in the network.

The jamming detection module is responsible for monitoring the radio and medium access

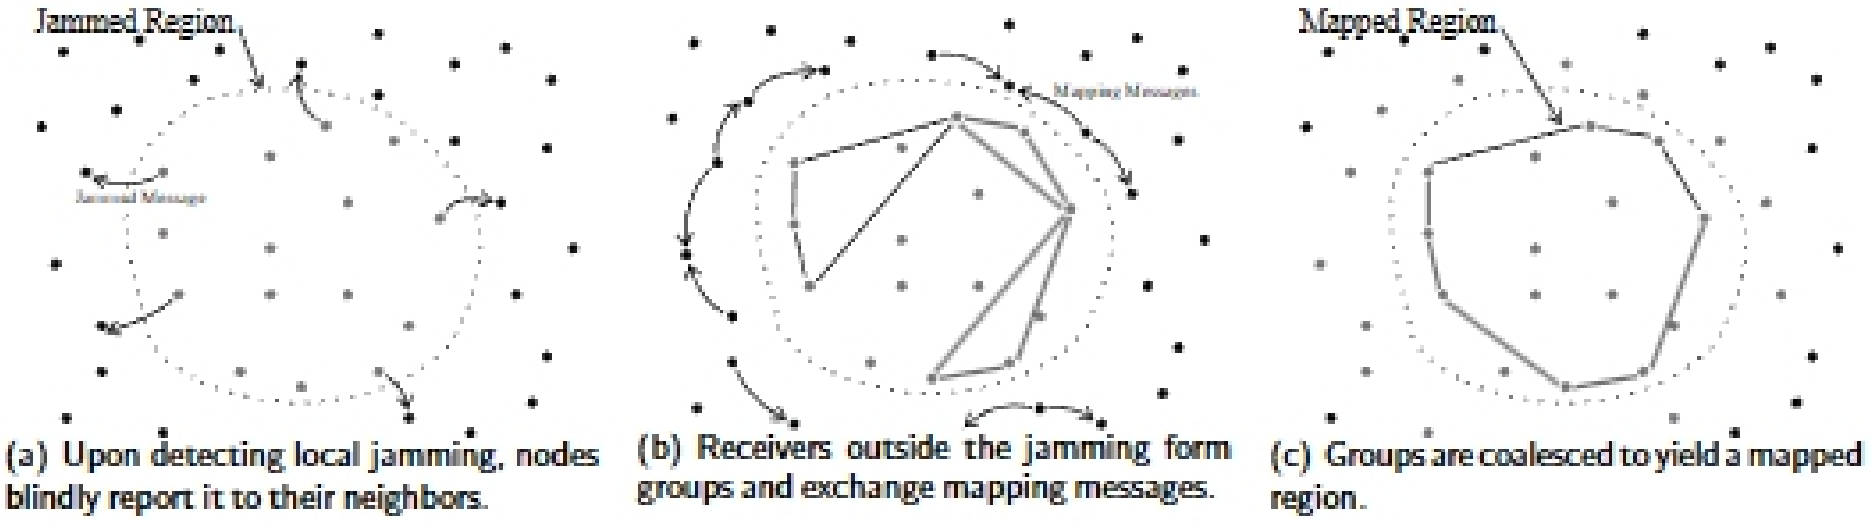

(a) Upon detecting local jamming, nodes blindly report it to their neighbors.

(b) Receivers outside the jamming form groups and exchange mapping messages.

(c) Groups are coalesced to yield a mapped region.

Figure 2. Overview of nodes collaboratively mapping a jammed region in the network

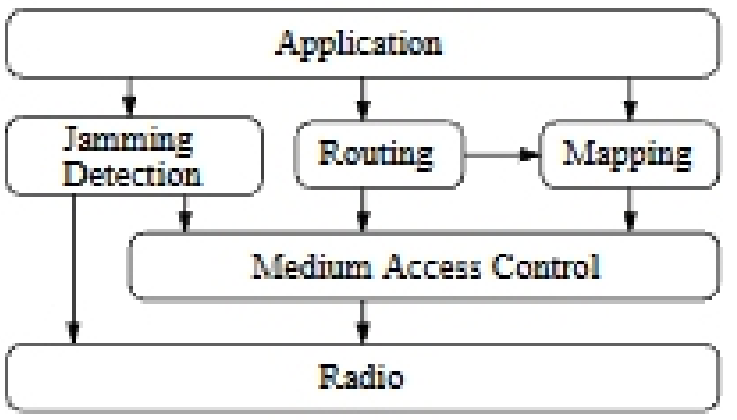

Figure 1. Architectural diagram of a mapping service. The arrows indicate a "uses" relation. Jamming Detection and Mapping interact remotely, that is, between nodes using the radio.

control (MAC) layers and applying heuristics to determine that the node is jammed. As described in Section 2.1, when it determines that the local node is most likely jammed, it sends a message to its neighbors by overriding the carrier-sense multiple access (CSMA) limitation usually enforced by the MAC, shown in Figure 2(a). It alerts the application layer, which can apply power management strategies to help the node outlast the jamming.

Mapping is initiated by the neighbors of jammed nodes who receive the jamming notifications. Each receiver forms a group, explicitly adding nearby jammed nodes as jammed members; the receiver itself becomes a mapping member. Figure 2(b) shows mapping messages, which contain information about the local group, being exchanged between neighbors. Neighboring groups are coalesced and eventually most or all of the mapping members know about the jammed region, as shown in Figure 2(c).

Details of the mapping protocol are in Section 2.2.

When the jammer(s) move or simply stop the attack, the jammed nodes recover and send notifications to their neighbors informing them of this change. The mapping members change the status of the formerly-jammed nodes and send messages to update the group. When a mapping member knows of no neighboring nodes that remain jammed, it retires from the group.

2.1. Jamming detection

Jamming is interfering with the ability of an adversary to communicate. A receiver may recognize known types of jamming by their unique energy patterns. However, this approach requires digital signal processing (DSP) capabilities and a library of patterns that may not be available except in military deployments.

We apply heuristics to determine whether the current node is experiencing non-transient interference that might be called jamming. In fact, we may expand our definition of jamming to include any kind of denial-of-service condition in which the *utility* of the communication channel drops below a certain threshold. This allows us to broaden our jamming model to include mobility, pulse jamming, and even link-layer jamming [15, 26], all based on their impact on the local ability to communicate.

The idea is that below this utility threshold, we are unable to communicate effectively enough for long enough to accomplish our objectives. This is necessarily dependent on the purpose of the sensor network.

Factors which impact this utility metric:

- Repeated inability to access wireless channel
- Bad framing
- Checksum failures
- Illegal values for addresses or other fields
- Protocol violations (e.g., missing ACKs)
- Excessive received signal level
- Low signal-to-noise ratio
- Repeated collisions
- Duration of condition

These data may be obtained from the local radio hardware, MAC layer, network layer, or other available