Poster Abstract: *SenSlide* - A Sensor Network Based Landslide Prediction System

A. Sheth
University of Colorado, Boulder
sheth@cs.colorado.edu

K. Tejaswi, P. Mehta, C. Parekh, R. Bansal, S. Merchant, T. Singh, U. B. Desai
IIT, Mumbai

C. A, Thekkath, K. Toyama
Microsoft Research

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; D.2.8 [**Software Engineering**]: Metrics—*complexity measures, performance measures*

General Terms

Algorithms, Reliability

Keywords

Lanslide prediction, uniform wear leveling, fault tolerance

1. INTRODUCTION

Landslides are a serious geological hazard caused when masses of rock, earth, and debris flow down a steep slope during periods of intense rainfall and rapid snow melt. The western (Konkan) coast and the Himalayan region of India are subject to many such landslides every year. Landslides in these rocky regions are mainly caused by the increase in strain due to percolating rain water in rocks fissures, causing rocks to fracture and slide down the slope. According to government reports, from 1998 to 2001 alone, landslides have killed more than 500 people, disrupted the communication and transport for weeks and destroyed thousands of hectares of crop area.

Existing solutions are restricted to landslide *detection*. A trip wire is installed along the landslide prone areas, and a break in the trip wire due to the falling rocks and debris triggers an alarm. Although this is an inexpensive solution for landslide *detection*, it is ineffectual in providing warning of the impending landslide.

Typical sensors used for monitoring slope stability are multi-point bore hole extensometers, tilt sensors, displacement sensors, and volumetric soil water content sensors. These require drilling 20-30 meter holes into the surface making the installation very expensive ($\approx$ \$50 per meter) and requiring skilled labor. Furthermore, these are expensive sensors making wide scale deployment infeasible. Installing a single sensor for monitoring an entire hill side is not sufficient as the properties of the rocks change every 100-200 meters. Wiring each sensor to a central data logger is also not feasible in the rocky terrain and requires high maintenance.

In contrast to the existing approaches, we propose using a Wireless Sensor Network (WSN) of 2-axis strain gauges to predict landslides. The small sized, low cost, and wireless battery operated nodes require minimum maintenance and can easily be deployed over a wide area. Strain gauges can operate at low depths (25-30 cms) and require low excitation voltage compared to the other sensors. The orders of magnitude lower depth of operation make strain gauges ideal for deployment

The goals of *SenSlide* are quite different from previous WSN deployments we are familiar with. A majority of the WSN deployments [3, 4] are mainly data collection networks, where the sensors are periodically sampled and sensor data is collected at a central base station for *offline* analysis. At the other end of the spectrum, there have been WSN deployments, which only communicate with the base station when a *rare event* is detected [1]. In our application, we need to satisfy both ends of the spectrum. Data needs to be sampled periodically to help earth scientists gather much needed historical trend information, while ensuring that the life-time of the network is not adversely affected by frequent sampling. Thus, *SenSlide* shares features of a rare event detection network as well as a very low sampling rate data collection network.

During the system design we encountered several challenges in successfully deploying sensors for predicting landslides. These challenges are detailed in the following four sections.

Our overall system design consists of sensor nodes (mica2 motes interfaced with strain gauges) organized in a hierarchy to monitor the strain in the rocks. A subset of these are designated as *aggregators* that collect the locally smoothed sensor data and create *spatial summaries*. These aggregators communicate with the *base station* (laptops connected to the Internet) providing summary data at adaptively adjusted frequencies.

2. FILTERING OF LOCAL SENSOR DATA

The inaccurate calibration of the cheap strain gauges result in noisy sensor data. Locally sampled sensor data must

SenSys'05, November 2–4, 2005, San Diego, California, USA.
ACM 1-59593-054-X/05/0011.

be smoothed before processing to avoid false positives and negatives.

Based on field experiments using extensometers, we observe that the strain in the rocks can be modeled as a piecewise linear function of time. In the first region, which persists for a relatively long time, the strain is constant. In the second region, the strain increases linearly with time until it reaches the elastic limit of the rock.

Each sensor samples data at a fixed rate: 3 minutes based on our analysis of the field data. Each sensor calculates an exponentially weighted moving average of the samples. It communicates this average to the rest of the system at different rates, depending on the region of the strain curve at which the sensor estimates it is operating.

The slowest rate at which sensors communicate data to aggregators is once every 30 minutes. During periods of increasing strain, which the sensor detects as increasing average values, the communication rate is increased as well, up to a maximum of once every 3 minutes.

3. ROUTING AND UNIFORM WEAR LEVELING

Convergecasting all sensor data to a central base station leads to non-uniform wear leveling of nodes in the network, and failure of all links one-hop away from the base station would partition the entire network. To avoid this, *SenSlide* designates multiple aggregator nodes that filter data and transmit to the base station only a summary of the data. Non-aggregator nodes transmit averages of local strain sensor data to the closest aggregator.

Data is routed within the network using Beacon Vector Routing (BVR), a scalable point-to-point routing protocol [2]. Periodically, each sensor node transmits a *node status* message to its closest base station containing its ID, its BVR coordinate vector, its energy level, and its neighbor list. The periodicity of these messages is much lower than the minimum sensor data communication rate.

A designated *leader* base station combines status information received at multiple base stations and designates certain nodes as aggregators. Aggregator nodes are selected based on the k-means clustering algorithm. The BVR coordinate vector, which is based on radio connectivity, is used to calculate the "k" aggregator nodes in the network. Our algorithm accounts for the constraints that (a) aggregators are not too close to the base stations nor at the extremities of the network, and (b) aggregators have sufficient energy to survive long enough to detect a rare event, and (c) they are not geographically co-located so as to survive environmental hazards that affect a large region of the network.

To avoid hot spots in the network, aggregator nodes are not static and are re-assigned by the leader base station. The aggregator selection algorithm is triggered when the energy level at the aggregator drops below E_{thresh} or when an aggregator is considered to have failed.

4. SPATIAL SUMMARY OF SENSOR DATA

Even though data from individual sensors is smoothed, single sensor observations are insufficient to predict a landslide. Aggregator nodes should summarize spatial sensor data for accurate prediction. This is achieved by curve fitting, where the aggregator only computes the coefficients of the curve that approximate the smoothed sensor data and transmits these coefficients to the base station. The base station can reconstruct the distribution of strain over the entire region by receiving only the coefficients from individual aggregator nodes. Any significant deviation in the coefficients triggers an alarm, which adaptively increases the sampling rate of the network. We are in the process of carrying out experiments to select appropriate curve fitting techniques based on distributed strain sensor data collected from different rocks.

5. FAULT TOLERANCE

SenSlide achieves fault tolerance by introducing redundancy at various levels of the system and detects failure using low overhead system heartbeats.

Sensor node and link failures are detected by exploiting the underlying broadcast nature of the wireless medium by using symmetric links for communication. Only nodes that are a single hop away from an aggregator request an explicit acknowledgement.

Aggregator failure is detected by monitoring the loss rate at base stations since aggregators transmit periodic summary data to the closest base station. An aggregator failure triggers the aggregator node selection algorithm. Changes in aggregator nodes and other control messages are broadcast by the leader base station to the entire network.

Base station failures are detected by having the multiple base stations monitor the liveness of each other. A leader base station is elected via a fault-tolerant distributed leader election algorithm running on the base stations. Periodic data received at a base station (aggregator summaries as well as periodic node status messages) is synchronously replicated on other base stations for fault-tolerance.

Notice that by using the replicated node status information, we can reconstruct the aggregator selection in the presence of base station failure. Based on the summary data from the aggregators we can reconstruct the distribution of strain over the entire network at the base station.

6. STATUS AND FUTURE WORK

We have completed the design of *SenSlide* and initial simulations of the aggregator selection algorithm. The 2-axis strain gauges have been interfaced with the mica2 motes, and are being used to collect extensive strain data for different rock types. This data will be used to select suitable curve fitting algorithms for spatial summarization. We plan to incorporate this data into network simulations to better understand the behavior of the system at scale. We also intend to evaluate our prototype in an indoor testbed and subsequently deploy it in the field.

7. REFERENCES

[1] A. Arora et al. A line in the sand: a wireless sensor network for target detection, classification, and tracking. *Comput. Networks*, 46(5):605–634, 2004.

[2] R. Fonseca et al. Beacon vector routing: Scalable point-to-point in wireless sensornets. In *Proceedings of the 2nd Symposium on Networked Systems Design and Implementation (NSDI)*, May 2005.

[3] R. Szewczyk et al. Lessons from a sensor network expedition. In *Proceedings of the First European Workshop on Sensor Networks (EWSN)*, Jan. 2004.

[4] N. Xu, et al. A wireless sensor network for structural monitoring. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 13–24, New York, NY, USA, 2004. ACM Press.