Quality-driven Volcanic Earthquake Detection using Wireless Sensor Networks

Rui Tan[1] Guoliang Xing[1] Jinzhu Chen[1] Wen-Zhan Song[2] Renjie Huang[3]
[1]Michigan State University, USA [2]Georgia State University, USA [3]Washington State University, USA

Abstract

Volcano monitoring is of great interest to public safety and scientific explorations. However, traditional volcanic instrumentation such as broadband seismometers are expensive, power-hungry, bulky, and difficult to install. Wireless sensor networks (WSNs) offer the potential to monitor volcanoes at unprecedented spatial and temporal scales. However, current volcanic WSN systems often yield poor monitoring quality due to the limited sensing capability of low-cost sensors and unpredictable dynamics of volcanic activities. Moreover, they are designed only for short-term monitoring due to the high energy consumption of centralized data collection. In this paper, we propose a novel quality-driven *approach to achieving real-time,* in-situ*, and long-lived volcanic earthquake detection. By employing novel in-network collaborative signal processing algorithms, our approach can meet stringent requirements on sensing quality (low false alarm/missing rate and precise earthquake onset time) at low power consumption. We have implemented our algorithms in TinyOS and conducted extensive evaluation on a testbed of 24 TelosB motes as well as simulations based on real data traces collected during 5.5 months on an active volcano. We show that our approach yields near-zero false alarm/missing rate and less than one second of detection delay while achieving up to 6-fold energy reduction over the current data collection approach.*

1 Introduction

In the last two decades, volcanic eruptions have led to a death toll of over 30,000 and damage of billions of dollars [1]. The recent eruptions of Volcano Eyjafjallajökull in Iceland caused the disruption of air traffic across Europe. Traditional volcano monitoring systems often employ broadband seismometers which, although can yield high-fidelity seismic monitoring signals, are expensive, power-hungry, bulky, and difficult to install. These limitations have largely prevented them from wide deployment, even for many threatening volcanoes. For instance, Mount St. Helens, an active volcano in northwestern U.S., is currently monitored by less than 10 stations [16] providing limited coverage and coarse-grain monitoring.

The advances of wireless sensor networks (WSNs) have made it possible to greatly improve volcanic monitoring quality through numerous low-cost sensors. Moreover, WSNs enable fast ad hoc system deployment that is largely impossible in the past. Recent pilot deployments on several active volcanoes [16, 20, 21] have demonstrated the feasibility and scientific value of WSNs to volcano monitoring. However, the current efforts of these projects are mostly focused on communication and networking issues such as reliable data delivery, time synchronization, and network management. In order to detect earthquake events, sensory data are transmitted to the base station for centralized processing. However, due to the sheer amount of raw data gathered at high sampling rates, such a data collection approach leads to excessive energy consumption and short system lifetime. Moreover, it has poor timeliness due to the limited bandwidth of low-cost sensors. For instance, as shown in [21], collecting one minute of seismic data over a multi-hop link can take up to six minutes. Although data transmission can be reduced by event-triggered data collection approaches [21], the existing earthquake detection algorithms [9] are heuristic in nature and often lead to excessive event misses. For instance, only about 5% of seismic events were successfully detected in a recent WSN deployment at Volcań Reventador in northern Ecuador [21].

In this paper, we push state of the art to real-time, *in-situ*, and long-lived volcano monitoring systems with assured sensing performance. In particular, we aim to completely avoid raw data transmission by developing advanced in-network signal processing algorithms for volcanic earthquake detection. To this end, the following challenges must be addressed. First, volcanic earthquake is a sophisticated physical process featured by highly dynamical magnitude and variable source location. These unpredictable dynamics must be properly dealt with in the sensing algorithms. Second, compared with traditional expensive monitoring instruments, low-cost wireless sensors often have limited sensing capability such as low signal-to-noise ratio and narrow responsive frequency band. Therefore, they must efficiently collaborate in signal processing to achieve the stringent sensing quality requirements. Third, the computation as well as inter-node communication overhead must be minimized to improve timeliness and extend system lifetime.

We make the following major contributions in this paper:

1052-8725/10 $26.00 © 2010 IEEE
DOI 10.1109/RTSS.2010.21

IEEE computer society

- We develop a novel *quality-driven* approach to detecting volcanic earthquakes based on collaborative signal processing algorithms. Our fundamental methodology is to drive the system design based on user's requirements on system sensing quality while minimizing sensors' energy consumption.
- We develop new sensing algorithms based on the extensive analysis of real data traces collected on Mount St. Helens [16]. First, we propose a Bayesian detection algorithm based on a novel joint statistical model of seismic signal energy and frequency spectrum. Second, we develop a near-optimal sensor selection algorithm that chooses the minimum subset of informative sensors to yield system detection results. The above two algorithms enable the system to achieve satisfactory sensing quality in the presence of unpredictable dynamics of volcanic earthquakes. Moreover, they only generate light traffic from sensors to the base station and completely avoid the transmission of raw data.
- We have implemented our algorithms on a testbed of 24 TelosB motes. We conduct testbed experiments and extensive simulations based on real data traces collected by 12 nodes on Mount St. Helens [16] that contain more than 128 significant earthquake events. Experimental results show that our approach yields near-zero false alarm/missing rate and less than one second of detection delay while achieving up to 6-fold energy reduction over the current data collection approach. Moreover, our approach allows a system to configure its sensing quality under different energy budgets.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 provides an overview of our approach. Section 4 presents the earthquake detection algorithm run by sensors locally. Section 5 develops a near-optimal sensor selection algorithm. Section 6 presents implementation details and Section 7 evaluates our approach. Section 8 concludes this paper.

2 Related Work

In 2004, four MICA2 motes were deployed on Volcań Tungurahua in central Ecuador [20], which is the first mote-based volcano monitoring system. The system lived for three days and successfully collected the data of at least 9 large explosions. In 2005, the same group deployed 16 Tmote nodes equipped with seismic and acoustic sensors at Volcań Reventador in northern Ecuador for three weeks [21, 22]. The main objective of the above two deployments is to collect high-resolution/fidelity sensor data for domain scientists. A simple event-triggered data collection approach based on the STA/LTA (short-term average over long-term average) [9] earthquake detection algorithm is developed to reduce data transmission. However, this heuristic approach cannot yield provable and satisfactory detection performance. For instance, although the systems had zero false alarm rate, they suffered very low detection probabilities (about 5%) [21]. Moreover, collected data are processed in a centralized fashion leading to significant bandwidth requirement and energy consumption.

In the Optimized Autonomous Space In-situ Sensorweb (OASIS) project [16], 15 iMote2-based nodes has been aerial deployed into Mount St. Helens since July 2009. In that project, significant research efforts have been put into improving system longevity, network efficiency and performance issues. The design has successfully delivered a long-term sustainable sensor network in challenging environment, and long-period (*e.g.*, several months) valuable real-world high-fidelity volcanic sensor dataset for our research. To our best knowledge, the issue of real-time quality-ensured in-network earthquake detection has not been addressed, although the heuristic STA/LTA earthquake detection algorithm was adopted for data prioritization in their design.

There exist a vast of well-established tools and techniques for processing sensor data in seismology community [4, 9, 15]. However, most of them are designed to centrally process seismic signals collected from traditional seismological stations. Specifically, seismic data must be logged at the stations and then transmitted or manually fetched to a base station for centralized processing [21, 22].

3 Approach Overview

In this section, we provide an overview of our approach to detecting volcanic earthquakes using a WSN. Our approach is designed to meet two key objectives of volcano monitoring. First, the system sensing quality must satisfy the Neyman-Pearson (NP) requirement [8] including upper-bounded false alarm rate and lower-bounded detection probability. For instance, seismologists may request that no more than 1% of detection reports are false alarms and the system can successfully detect at least 90% earthquake events. Second, the computation and communication overhead of sensors must be minimized to improve timeliness and extend system lifetime.

We assume that the network comprises a base station and a number of sensors distributed on the volcano. In this paper, we assume that all sensors are of seismic modality, which is consistent with several first-generation volcano monitoring WSNs [16, 21]. Our approach comprises a group of detection algorithms that run at sensors and the base station, respectively. They work together to achieve the requirements on sensing quality. A system architecture of our approach is shown in Figure 1. Each sensor detects earthquake event every sampling period based on seismic frequency spectrum. To handle the earthquake dynamics such as highly dynamical magnitude and variable

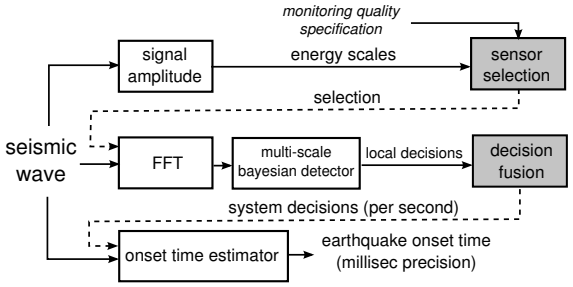

Figure 1. System architecture. White blocks are the components at a sensor; shadowed blocks are the components at the base station; solid line represents data flow; dotted line represents control flow.

source location, each sensor maintains separate statistical models of frequency spectrum for different scales of seismic signal energy received by sensor. Our study shows that the frequency-based detector typically has better detection performance when the sensor receives higher signal energy. Therefore, in our approach, the base station first selects a minimum subset of informative sensors based on the signal energies received by sensors while satisfying system sensing quality requirements. The selected sensors then compute seismic frequency spectrum using fast Fourier transform (FFT) and make local detection decisions which are then transmitted to the base station for fusion. In addition to the detection of earthquake occurrences, node-level earthquake onset time is critical for localizing earthquake source. In our approach, the base station first identifies an individual earthquake and estimates a coarse onset time. The coarse onset time is then fed back to sensors, which will pick the P-phase (*i.e.*, the arrival time of wavefront) from buffered raw seismic data using existing algorithms, *e.g.*, [15]. However, due to space limitation, we only focus on the detection of earthquake occurrences in this paper. The details of the onset time estimation are omitted and can be found in [17].

Our approach has the following advantages. First, different from existing heuristic earthquake detection algorithms such as STA/LTA, our model-driven approach can meet various sensing quality requirements including bounded false alarm rate and detection probability. Second, by employing novel in-network data fusion schemes, our approach incurs low communication overhead. Specifically, in each sampling period, only signal energy represented by an integer needs to be sent to the base station. Only when the system sensing quality meets user's requirement, local decisions made by sensors are transmitted to the base station. Third, the sensor selection algorithm allows a network to achieve desired trade-off between system sensing quality and computational overhead at sensors. In particular, based on the requirement on energy-efficiency, only a minimum number of sensors are selected to execute the computation-intensive signal processing algorithms such as FFT.

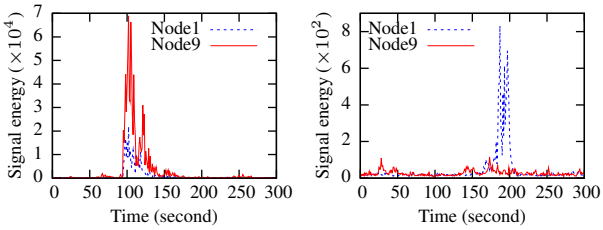

(a) Event 1 (b) Event 2

Figure 2. Seismic signal energy.

4 Local Earthquake Detection at Sensors

In this section, we design a local earthquake detection algorithm that runs at sensors locally. In order to achieve satisfactory sensing performance, the following questions must be addressed. First, what information does a sensor need to sample? Due to the resource limitation of low-cost sensors, the amount of sampled information must be minimized while critical features of earthquake should be conserved. Second, how to represent the sampled information using a sensing model? The overhead of computing and storing the model should be affordable for low-cost sensors. Third, how to accurately detect earthquakes based on the sensing model and real-time measurements? In the following, we first present a case study of sensors' measurements in earthquakes and then address the above questions.

4.1 A Case Study of Earthquake Sensing

Detecting volcanic earthquakes using low-cost accelerometers in WSN is challenging due to the dynamics of earthquake, *e.g.*, significantly variable magnitude and source location. Moreover, as seismic signal attenuates with the propagation distance, the sensors far away from the earthquake source receive weak signals and hence have lower detectabilities. Such a phenomenon is referred to as the *locality* of earthquake in this paper. In this section, we illustrate the locality of earthquake using a case study, which motivates us to propose a novel sensing model for volcanic earthquake detection.

The case study is based on the seismic data traces collected by 12 nodes in the OASIS project on Mount St. Helens [16]. We examine micro-scale *signal energy* and *frequency spectrum* which are two basic statistics computed from sensors' raw data. Figures 2(a) and 2(b) plot the signal energy received by Node1 and Node9 in two earthquake events, respectively. From the figures, we can see that Node9 receives higher signal energy than Node1 in Event 1, while Node1 receives significantly higher signal energy than Node9 in Event 2. This example shows that the signal energy received by a sensor varies significantly due to the change of the earthquake source location as well as its magnitude. Therefore, simple threshold detection approaches based on signal energy [7, 13, 18] would not address the

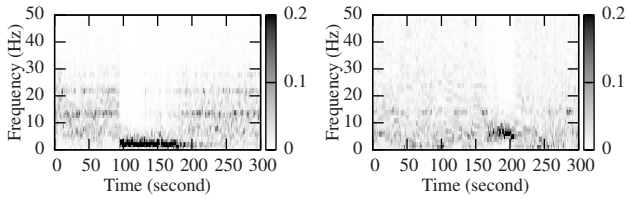

(a) Event 1 (b) Event 2

Figure 3. Frequency spectrum of Node1.

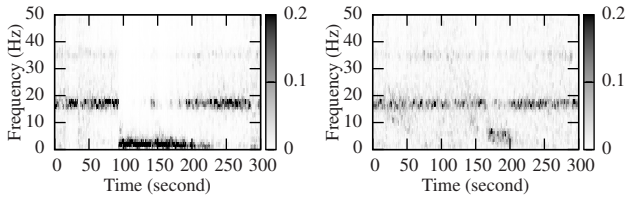

(a) Event 1 (b) Event 2

Figure 4. Frequency spectrum of Node9.

dynamics of volcanic earthquakes. Figures 3(a) and 3(b) plot the spectrum of Node1 in the two events, respectively. As the signal energy of Event 1 is much stronger than that of Event 2 (about 100 times), Node1 has significantly different frequency spectra in the two events. Specifically, the received seismic energy is mainly distributed within $[0\,\text{Hz}, 5\,\text{Hz}]$ in Event 1 and $[5\,\text{Hz}, 10\,\text{Hz}]$ in Event 2. Figures 4(a) and 4(b) plot the spectrum of Node9. We can see that Node9 has insignificant frequency feature in Event 2 due to very weak signals. Moreover, from Figures 3(a) and 4(a), we can see that Node1 and Node9 have different frequency spectra in the absence of earthquake. We can make two important observations from this case study for constructing earthquake sensing model. First, in order to achieve satisfactory sensing quality, signal energy and frequency spectrum must be jointly considered for detecting earthquakes. Second, the frequency spectra for different scales of signal energy sensed by a sensor vary considerably and hence require different mathematical representations.

4.2 Feature Extraction

To capture the significant temporal dynamics of earthquake, sensors have to perform detections at a short period, *e.g.*, per second. In the following, we discuss efficient sampling schemes to obtain both frequency spectrum and signal energy. The seismic waves emitted by an earthquake can be classified as the *primary wave* (P-wave) and *shear wave* (S-wave). The P-wave is faster than S-wave and its frequency is typically from 1 Hz to 10 Hz, while the slower S-wave often has a frequency of lower than 1 Hz [4]. Different from the high-cost broadband seismometers that are traditionally used by the seismological community, low-cost accelerometers in WSNs, *e.g.*, 1221J-002 from Silicon Designs [16], are only responsive to P-wave. As a result, the seismic energy measured by these accelerometers in the presence of earthquake is mainly distributed within $[1\,\text{Hz}, 10\,\text{Hz}]$. As shown in Section 4.1, frequency spectrum is expected to be a robust feature for detecting earthquakes using low-cost accelerometers. Suppose the sampling rate is f Hz. By applying FFT to the raw seismic data received during one second, a sensor obtains the frequency spectrum that ranges from 0 Hz to $f/2$ Hz. Each component of the spectrum represents the percentage of signal energy that is located in the corresponding frequency.

The sampling rate of accelerometers can be high (up to $400\,\text{Hz}$). In order to reduce the computation overhead of sensors, we construct feature vector from the frequency spectrum as follows. The frequency spectrum is evenly divided into n bins. Let $\mathbf{x}$ denote the feature vector at a sensor. The i^{th} component of the feature vector, *i.e.*, $\mathbf{x}[i]$, is the sum of spectrum components in the i^{th} bin. Hence, $\mathbf{x}[i]$ is the percentage of signal energy that is distributed in $(\frac{i \cdot f}{2n}\,\text{Hz}, \frac{(i+1) \cdot f}{2n}\,\text{Hz}]$, where $i = 0, 1, \ldots, n-1$. As the dimension of feature vector, *i.e.*, n, determines the computation complexity of the training and detection algorithms at sensors, n should be chosen to achieve satisfactory trade-off between detection accuracy and computation overhead.

In addition to frequency spectrum, signal energy received by sensors is also an important feature that quantifies the earthquake magnitude. The signal energy at a sensor is often estimated by the mean square of seismic intensities during a *sampling period* [14]. To be consistent with the frequency analysis, we let the sampling period be one second in this work. Let y_i denote the i^{th} seismic intensity and e denote the signal energy. For a sampling rate of f Hz, the signal energy is computed by $e = \frac{1}{f}\sum_{i=1}^{f}(y_i - \bar{y})^2$, where $\bar{y}$ is the mean of seismic intensities in a sampling period.

4.3 A Multi-scale Sensing Model

We now propose a *multi-scale* Bayesian model that jointly accounts for signal energy and frequency spectrum received by a sensor to deal with the dynamics and locality of earthquakes that are discussed in Section 4.1. In the multi-scale Bayesian model, the range of signal energy is divided into K consecutive sub-ranges, denoted by $\{R_p | p \in [1, K]\}$. Each sensor maintains $K+1$ n-dimensional normal distributions, which are denoted by $\{\mathcal{N}_p | p \in [0, K]\}$. Note that n is the dimension of the frequency feature vector. The distribution $\mathcal{N}_0$ represents the model of frequency feature vector in the absence of earthquake and $\{\mathcal{N}_p | p \in [1, K]\}$ correspond to the cases when earthquake happens and the received signal energy falls into the p^{th} energy range, *i.e.*, $e \in R_p$. Each normal distribution $\mathcal{N}_p$ is characterized by its mean vector and covariance matrix, which are denoted by $\mathbf{m}_p$ and $\mathbf{C}_p$. Specifically, $\mathbf{m}_p[i] = \mathbb{E}[\mathbf{x}[i] | e \in R_p]$ and $\mathbf{C}_p[i, j] = \text{cov}(\mathbf{x}[i] | e \in R_p, \mathbf{x}[j] | e \in R_p)$, where $\mathbf{x}[i]$ is the i^{th} component of the frequency feature vector. With the above model, the frequency spectra for different scales of

signal energy are characterized by separate normal distributions that carry sensing quality information. Such a model allows us to precisely describe sensors' performance in the presence of earthquake dynamics and locality.

We now discuss how to divide the signal energy range. The range of signal intensity measured by a sensor depends on its bit-depth and calibration. Therefore, for different sensor products, the range of signal intensity varies significantly. However, through proper normalization to signal intensity, we can develop a universal scale scheme for signal energy. In this work, we employ a base-10 logarithmic scale to represent the signal energy range, which is consistent with many widely adopted earthquake magnitude scales such as the Richter magnitude scale [10]. Specifically, we let $p = \lfloor \log_{10} e \rfloor$ where e is the received signal energy. Therefore, the p^{th} energy scale range, R_p, is $[10^p, 10^{p+1})$ and p is referred to as *energy scale* hereafter. For example, the signal energy ranges from 10 to 10^6 for the data traces collected in the OASIS project [16] and therefore the energy scale is from 1 to 6.

In order to build the multi-scale Bayesian model, we need to compute the mean vector $\mathbf{m}_p$ and covariance matrix $\mathbf{C}_p$ using enough samples. As both mean and covariance can be updated efficiently with incremental algorithms when a new sample is available, the model learning can be performed on each sensor locally at low cost. Specifically, a sensor learns its sensing model as follows. When no earthquake occurs, the sensor updates the distribution $\mathcal{N}_0$ using the current extracted frequency feature vector; otherwise, it first computes the energy scale p and then updates the corresponding distribution $\mathcal{N}_p$. This model learning process can be conducted offline with data traces. Alternatively, it can be conducted online with the ground truth information from high-quality sensors. Seismological monitoring infrastructures already deployed on active volcanoes can be used to generate ground truth for training newly deployed low-cost sensors. As these infrastructures are often power-hungry, they can be turned off when the training completes.

4.4 Local Bayesian Detector

Based on the multi-scale Bayesian model presented in Section 4.3, we design a Bayesian detector for each sensor to achieve optimal local detection performance. The detector makes a decision based on both the energy scale p and frequency feature vector $\mathbf{x}$. The local decisions of sensors are then fused at base station to improve system sensing quality, which will be discussed in Section 5. A sensor makes a decision between the hypotheses that there is earthquake or not (denoted by H_p and H_0, respectively):

$$H_0 : \mathbf{x} \sim \mathcal{N}(\mathbf{m}_0, \mathbf{C}_0); \quad H_p : \mathbf{x} \sim \mathcal{N}(\mathbf{m}_p, \mathbf{C}_p).$$

Let I denote the local decision made by the sensor. Specifically, if the sensor accepts the null hypothesis H_0, $I = 0$; otherwise, $I = 1$. The detection performance is usually characterized by two metrics, namely, false alarm rate (denoted by P_F) and detection probability (denoted by P_D). P_F is the probability that the sensor decides $I = 1$ when the ground truth is H_0. P_D is the probability that the sensor decides $I = 1$ when the ground truth is H_p. Among many existing decision criteria, the minimum error rate criterion is the most widely adopted one that jointly accounts for false alarms and misses. Moreover, in contrast to other complicated decision criteria, the minimum error rate criterion has a closed-form decision function, which can largely reduce the computation overhead at sensors. Given the frequency feature $\mathbf{x}$, the decision functions for minimum error rate are [8] $g_i(\mathbf{x}) = \ln \mathbb{P}(H_i) - \frac{1}{2} \ln |\mathbf{C}_i| - \frac{1}{2} \left((\mathbf{x}-\mathbf{m}_i)^{\mathrm{T}} \mathbf{C}_i^{-1} (\mathbf{x}-\mathbf{m}_i) \right)$ for $i \in \{0, p\}$, where $\mathbb{P}(H_i)$ is the prior probability of the ground truth H_i and $|\mathbf{C}_i|$ represents the determinant of $\mathbf{C}_i$. The local detection decision I is made by $g_0(\mathbf{x}) \underset{I=1}{\overset{I=0}{\gtrless}} g_p(\mathbf{x})$. However, the matrix computations are too expensive for low-cost sensors when the dimension is high (*e.g.*, up to 10). In our approach, if sensors are trained in an online fashion as discussed in Section 4.3, each sensor transmits the mean vectors and covariance matrices to the base station, which computes the determinant and inverse of the covariance matrices and then transmits them back to sensors.

Under the above decision rule, the false alarm rate and detection probability of the sensor are given by $P_F = \int_{\mathcal{R}} \phi(\mathbf{x}|\mathbf{m}_0, \mathbf{C}_0) \mathrm{d}\mathbf{x}$ and $P_D = \int_{\mathcal{R}} \phi(\mathbf{x}|\mathbf{m}_p, \mathbf{C}_p) \mathrm{d}\mathbf{x}$, where $\mathcal{R} = \{\mathbf{x} | g_0(\mathbf{x}) < g_p(\mathbf{x})\}$ and $\phi(\mathbf{x}|\mathbf{m}_i, \mathbf{C}_i)$ is the probability distribution function (PDF) of the normal distribution $\mathcal{N}(\mathbf{m}_i, \mathbf{C}_i)$. Specifically,

$$\phi(\mathbf{x}|\mathbf{m}_i, \mathbf{C}_i) = \frac{1}{(2\pi)^{\frac{n}{2}} |\mathbf{C}_i|^{\frac{1}{2}}} \exp\left(-\frac{(\mathbf{x}-\mathbf{m}_i)^{\mathrm{T}} \mathbf{C}_i^{-1} (\mathbf{x}-\mathbf{m}_i)}{2} \right),$$

where n is the dimension of $\mathbf{x}$. We note that each pair of (H_0, H_p) where $p \in [1, K]$ gives a pair of (P_F, P_D). However, it is usually difficult to obtain the closed-form expression for the integral region $\mathcal{R}$ for computing P_F and P_D. In our approach, the base station computes the P_F and P_D for each pair of (H_0, H_p) through Monte Carlo simulation. The P_F's and P_D's for each sensor are stored at the base station, which will be used to select the most informative sensors to detect earthquakes as discussed in Section 5.

5 Dynamic Sensor Selection for Decision Fusion

As discussed in Section 4.4, sensors can yield local detection decisions by a Bayesian detector. However, the accuracy of these decisions may be poor due to the limited sensing capability of low-cost sensors. Therefore, a system-wide detection consensus is often desired for high-

fidelity volcano monitoring. In our approach, the base station generates system detection decision by fusing the local decisions from sensors. As sensors yield different sensing performances due to the dynamics and locality of volcanic earthquake as discussed in Section 4, it is desirable for the base station to select a subset of sensors with the best signal quality to achieve maximum system detection performance. Moreover, the sensor selection avoids unnecessary expensive feature extraction at the sensors with low signal quality. In this section, we first introduce the decision fusion model and analyze its performance. We then formulate the sensor selection as an optimization problem and develop a near-optimal solution.

5.1 Decision Fusion Model

As one of basic data fusion schemes [19], *decision fusion* is preferable for WSNs due to its low communication cost [7]. We use a widely adopted decision fusion model called equal gain combining (EGC) [7,13,18] that fuses sensors' local decisions with equal weight. Suppose there are n sensors taking part in the fusion and let I_i denote the local decision of sensor i. The EGC compares the test statistic $\Lambda = \sum_{i=1}^{n} I_i$ against a threshold denoted by η. If Λ exceeds η, the base station decides that an earthquake has occurred; otherwise, it makes a negative decision.

We now analyze the system detection performance of the EGC fusion model. In the absence of earthquake, the local decision of sensor i, $I_i|H_0$, follows the Bernoulli distribution with α_i as success probability. As sensors have different false alarm rates, the test statistic $\Lambda|H_0$ follows a generalized Binomial distribution. The probability mass function (PMF) of $\Lambda|H_0$ is given by

$$\mathbb{P}(\Lambda = \lambda|H_0) = \sum_{||S||=\lambda, \forall S} \prod_{i\in S} \alpha_i \prod_{j\in S^{\mathrm{C}}} (1-\alpha_j), \quad (1)$$

where S is any subset of sensors with size of λ and S^{C} represents the complement of S. Hence, the cumulative distribution function (CDF), denoted by $F_{\Lambda|H_0}(x)$, is given by $F_{\Lambda|H_0}(x) = \sum_{\lambda=0}^{\lfloor x \rfloor} \mathbb{P}(\Lambda = \lambda|H_0)$. Therefore, the system false alarm rate can be computed as $P_F = 1 - F_{\Lambda|H_0}(\eta)$. Similarly, the system detection probability can be computed as $P_D = 1 - F_{\Lambda|H_1}(\eta)$. Note that replacing α_i in (1) with β_i yields the PMF of $\Lambda|H_1$. However, computing the CDF of Λ has a complexity of $O(2^n)$ and hence is infeasible when the number of fused sensors is large.

We now propose approximate formulae for the system detection performance of the EGC fusion model when the number of fused sensors is large. As sensors independently make local decisions, the mean and variance of $\Lambda|H_0$ are given by $\mathbb{E}[\Lambda|H_0] = \sum_{i=1}^{n} \mathbb{E}[I_i|H_0] = \sum_{i=1}^{n} \alpha_i$, $\mathrm{Var}[\Lambda|H_0] = \sum_{i=1}^{n} \mathrm{Var}[I_i|H_0] = \sum_{i=1}^{n} \alpha_i - \alpha_i^2$. Lyapunov's central limit theorem (CLT) [5] is a CLT variant for independent but non-identically distributed variables. We have proved the Lyapunov condition for a sequence of Bernoulli random variables in [18]. Therefore, according to Lyapunov's CLT, $\Lambda|H_0$ follows the normal distribution when n is large, *i.e.*, $\Lambda|H_0 \sim \mathcal{N}\left(\sum_{i=1}^{n} \alpha_i, \sum_{i=1}^{n} \alpha_i - \alpha_i^2\right)$. Similarly, $\Lambda|H_1 \sim \mathcal{N}\left(\sum_{i=1}^{n} \beta_i, \sum_{i=1}^{n} \beta_i - \beta_i^2\right)$. Hence, the system false alarm rate and detection probability are given by

$$P_F \simeq Q\left(\frac{\eta - \sum_{i=1}^{n} \alpha_i}{\sqrt{\sum_{i=1}^{n} \alpha_i - \alpha_i^2}}\right),\ P_D \simeq Q\left(\frac{\eta - \sum_{i=1}^{n} \beta_i}{\sqrt{\sum_{i=1}^{n} \beta_i - \beta_i^2}}\right),$$

where $Q(\cdot)$ is the Q-function of the standard normal distribution, *i.e.*, $Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} \mathrm{e}^{-t^2/2} \mathrm{d}t$.

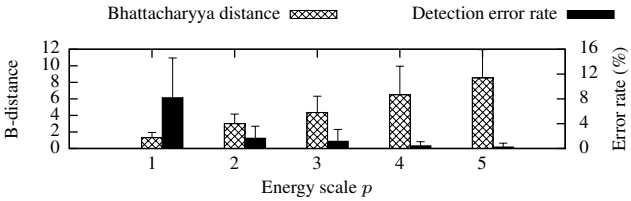

Figure 5. Bhattacharyya distance and corresponding detection error rate versus energy scale. The results show the standard deviation over 12 sensors.

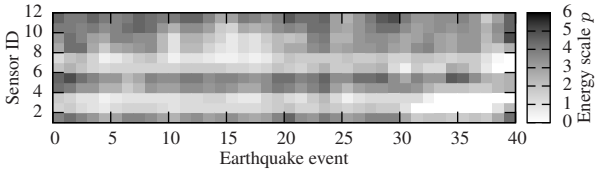

Figure 6. Dynamics of energy scale received by sensors.

5.2 Dynamic Sensor Selection Problem

The case study in Sectin 4.1 shows that a sensor exhibits different frequency patterns for different energy scales. Moreover, sensors receive significantly different energy scales due to the locality of earthquake. Our objective is to select a subset of sensors with the best signal quality to maximize system detection performance. To this end, we first examine the sensing performance diversity of sensors based on data traces collected in OASIS [16]. The result motivates us to formulate a dynamic sensor selection problem to achieve satisfactory trade-off between system detection performance and computation overhead at sensors. For each sensor, we compute the Bhattacharyya distance [8], which is a widely adopted detectability measure, between the p^{th} distribution $\mathcal{N}_p$ and the noise distribution $\mathcal{N}_0$ within its multi-scale Bayesian model. Figure 5 plots the error bars of Bhattacharyya distance and the corresponding detection error rate versus the energy scale p. We can see that the frequency-based detector has better performance when

a sensor receives stronger signal energy. Moreover, sensors show significant performance variance for the same energy scale. Figure 6 plots the maximum energy scale measured by sensors in 40 earthquake events.

We can make two important observations from Figures 5 and 6. First, for a particular event, sensors have different detection performances due to different received energy scales. As a result, sensors with poor sensing performances should be excluded from participating in system decision fusion. Moreover, if a sensor has sufficient sensing performance for system decision fusion, it must make local decisions by costly FFT to extract frequency features. Therefore, it is desirable to select the minimum subset of informative sensors to fuse their decisions. Second, each sensor has unpredictable signal energy pattern due to the stochastic nature of earthquake magnitude and source location. Although the optimal sensor selection can be precomputed for all possible combinations of sensors' energy scales, both the time and storage complexities are exponential, *i.e.*, $O(K^N)$, where K is the number of energy scales and N is the total number of sensors. Therefore, the sensors that have the best sensing performances must be dynamically selected in each sampling period.

We now formally formulate the sensor selection problem. We aim to select the minimum number of sensors being involved in the feature extraction and decision fusion processes, subject to bounded system detection performance. We adopt the Neyman-Pearson (NP) criterion [8] for characterizing system detection performance, *i.e.*, we allow users to specify the upper and lower bounds on system false alarm rate and detection probability, respectively. NP criterion is useful when the two types of errors, *i.e.*, false alarms and misses, need separate considerations. There exists a fundamental trade-off between the two metrics for any detection system, *i.e.*, higher detection probability is always achieved at the price of higher false alarm rate [19]. Depending on the characteristics of volcanoes to be monitored, seismologists may have different requirements on false alarm rate and detection probability. For instance, for an active volcano with frequent tiny earthquakes, it is desirable to reduce false alarms to avoid excessive sensor energy consumption and prolong system lifetime. On the other hand, for a dormant volcano, it is more critical to detect every important earthquake event while a higher false alarm rate can be tolerated. Note that our approach can be extended to address other performance metrics such as error rate that jointly accounts for false alarms and misses. Due to space limitation, the extension is omitted and can be found in [17]. Based on the fusion model in Section 5.1, the sensor selection problem is formally formulated as follows:

Problem 1. *Given the local false alarm rates and detection probabilities of all sensors,* i.e.*,* $\{\alpha_i, \beta_i | i \in [1, N]\}$*, to find a subset of sensors,* S*, and the decision fusion threshold at the base station,* η*, such that* $||S||$ *is minimized, subject to that the system false alarm rate is upper-bounded by* α *and the system detection probability is lower-bounded by* β*.*

The brutal-force solution, *i.e.*, iterating all possible subsets of sensors, has an exponential complexity of $O(2^N)$. As the dynamic sensor selection is conducted every sampling period (one second in our system), such a complexity would impede the system timeliness. In the rest of this section, we first reduce the complexity of Problem 1 with approximations and then develop a near-optimal sensor selection algorithm with polynomial complexity.

5.3 Dynamic Sensor Selection Algorithm

We adopt a divide-and-conquer strategy to solve Problem 1. The sub-problem of Problem 1 is to select n sensors out of the total N sensors such that the system detection performance is optimized. By iterating n from 1 to N, Problem 1 is solved once the optimal solution of the sub-problem satisfies the detection performance requirement. The brutal-force search for the optimal solution of the sub-problem has a complexity of $O\left(\binom{N}{n}\right)$. The following analysis shows that the sub-problem can be reduced to a sorting problem with polynomial complexity.

We first analyze the condition for the NP criterion. Due to the fundamental trade-off between false alarm rate and detection probability [19], P_D is maximized when P_F is set to its upper bound. Therefore, by letting $P_F = \alpha$, the detection threshold at the base station is

$$\eta = \sum_{i=1}^{n} \alpha_i + Q^{-1}(\alpha) \cdot \sqrt{\sum_{i=1}^{n} \alpha_i - \alpha_i^2}, \tag{2}$$

where $Q^{-1}(\cdot)$ is the inverse function of $Q(\cdot)$. Hence, the system detection probability is $P_D = Q(f)$, where

$$f = \frac{Q^{-1}(\alpha)\sqrt{\sum_{i=1}^{n} \alpha_i - \alpha_i^2} + \sum_{i=1}^{n}(\alpha_i - \beta_i)}{\sqrt{\sum_{i=1}^{n} \beta_i - \beta_i^2}}. \tag{3}$$

As $Q(\cdot)$ is a decreasing function, P_D is maximized if f is minimized. Therefore, the sub-problem is equivalent to minimizing f.

However, the function f has a complex non-linear relationship with each sensor's detection performance represented by α_i and β_i. We now propose a linear approximation to f. The Monte Carlo simulations show that f increases with $\sum_{i=1}^{n}(\alpha_i - \beta_i)$ with high probability ($\geq 95\%$). Due to the space limitation, the details of the simulations are omitted and can be found in [17]. Therefore, the sub-problem is reduced to selecting n sensor to minimize $\sum_{i=1}^{n}(\alpha_i - \beta_i)$, which can be easily solved by sorting sensors ascendingly according to the value of $(\alpha_i - \beta_i)$.

Based on the above analysis, we now develop a dynamic sensor selection algorithm to solve Problem 1, which is

Algorithm 1 Dynamic sensor selection algorithm

Input: local P_F's and P_D's $\{\alpha_i, \beta_i | i \in [1, N]\}$, system performance requirements $\{\alpha, \beta\}$
Output: minimum subset S, detection threshold η
1: sort sensors according to $(\alpha_i - \beta_i)$ ascendingly
2: **for** $n = 1$ to N **do**
3: $S = \{$ top n sensors $\}$
4: compute f with S using (3)
5: **if** $Q(f) \geq \beta$ **then**
6: compute η with S using (2)
7: **return** S and η
8: **end if**
9: **end for**
10: exit with no solution

listed in Algorithm 1. With the solution given by Algorithm 1, the base station will request the selected sensors to perform FFT and make their local decisions. Finally, the base station compares the sum of local decisions against the detection threshold η to make a system detection decision. In the absence of earthquake, Algorithm 1 would exit without a solution (Line 10). As a result, no sensor will be selected and hence costly seismic processing algorithms such as FFT can be avoided.

6 Implementation

We have implemented the proposed detection algorithms in TinyOS 2.1.0 on TelosB platform and conducted testbed experiments in laboratory. In the future work, we plan to deploy our implementation on the OASIS system [16] that is currently monitoring Mount St. Helens. Our implementation uses 45.3KB ROM and 9.5KB RAM when a sensor buffers 8 seconds of raw data for earthquake onset time estimation. Several important implementation details are presented as follows.

Data acquisition and seismic processing: To improve the realism of testbed experiments, we create a volume of 320KB on mote's flash and load it with the seismic data traces collected in OASIS [16]. We implement a nesC module that provides the standard `ReadStream` interface to read seismic data from flash to simulate data acquisition in real deployments. A node acquires 100 seismic intensities every sampling period. When the sampling period is set to be one second, the sampling rate is consistent with previous deployments [16,21]. We use the `KissFFT` [3] library to compute the frequency spectrum of seismic signals. In particular, we use the fixed-point FFT routines that are suitable for the 16-bit processor on TelosB mote.

Networking: Sensors are organized into a multi-hop tree rooted at the base station. In order to achieve timeliness, sensors are scheduled in a TDMA fashion. Specifically, a sensor reserves 250 ms for the FFT and Bayesian detector in each sampling period. The remaining time is divided into a number of slots, which are distributed among sensors for transmitting energy scales and local decisions. In order to reduce transmissions, the packets are aggregated along the routing path to the base station. For instance, when a non-leaf node has received all the energy scales from its children, it aggregates them together with its own into a single packet before forwarding. In our implementation, an energy scale entry is 1 byte where node ID uses 5 bits and energy scale uses 3 bits. Moreover, to improve reliability, a sensor buffers energy scale or decision packets from its children for at most 8 sampling periods. When a sensor has received all packets from its children for the current sampling period, it sends out the aggregated packets for the current and previous sampling periods. The sensor selection and decision fusion algorithms presented in Section 5 are implemented in Java on a desktop computer that serves as the base station. The sensor selection algorithm typically takes 10 ms to 20 ms, and hence has little impact on the timeliness of event detection.

7 Performance Evaluation

We conduct testbed experiments as well as extensive simulations based on real data traces collected by 12 nodes in the OASIS project [16]. The data set used in our evaluation spans 5.5 months (from October 1, 2009 to March 15, 2010) and comprises 128 manually selected segments. Each segment is 10 minutes and contains one or more significant earthquake events. In Section 7.1, we present the experimental results on energy usage and communication performance using a testbed of 24 TelosB motes. In Section 7.2, we present the simulation results on detection performance in TOSSIM.

7.1 Testbed Experiments

7.1.1 Methodology

The multi-scale Gaussian model of each sensor is trained offline using randomly selected 64 data segments. The ground truth information regarding the presence of earthquake event is generated by the STA/LTA algorithm using the data traces of Node01 in the deployment. The STA/LTA threshold is set to be 2, which is suggested by the volcanologists at U.S. Geological Survey [16]. We note that the STA/LTA algorithm can yield detection errors.

In this section, our approach is referred to as *decision fusion with sensor selection* (DFSS). We compare our approach with the following three baseline approaches. (1) In the *data collection* approach, each node transmits compressed raw data to the base station. We adopt incremental encoding to compress raw data, which can achieve 4-fold data volume reduction for 32-bit seismic signal in the absence of earthquake. Note that the OASIS system [16] currently adopts data collection and analyzes collected data offline at servers. (2) In the *STA/LTA* approach, each node makes local detection decision by the STA/LTA algorithm [9]. If more than 30% nodes make positive decisions, the base station first waits 30 seconds and then downloads one

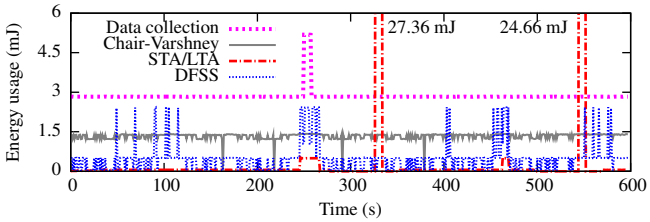

Figure 7. Energy consumption of Node 11.

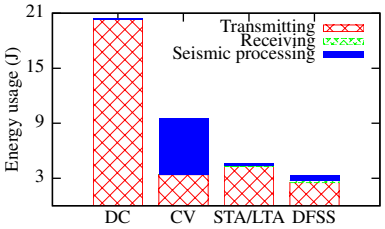

Figure 8. Total energy consumption of 12 nodes for 10 minutes. DC represents data collection, CV represents Chair-Varshney.

minute of compressed raw data from all nodes. Note that these settings are consistent with the detection approach in [21]. (3) In the *Chair-Varshney* approach, each node performs FFT and makes a local detection decision every sampling period. The base station fuses the local decisions by the Chair-Varshney's rule [6] that is the optimal decision fusion model. Specifically, the test statistic is $\Lambda = \sum_{i=1}^{n} \log \frac{\beta_i(1-\alpha_i)}{\alpha_i(1-\beta_i)} \cdot I_i$. As the Chair-Varshney's rule inherently accounts for the diversity of sensors' sensing qualities by weighting their local decisions, it is unnecessary to perform sensor selection. However, the Chair-Varshney's rule has no closed-form formula for its detection performance. Hence, we use a brutal-force approach to compute the CDF of Λ and find the detection threshold that satisfies detection performance requirements. Note that the brutal-force algorithm runs at the base station. The following experiments are conducted in two network topologies: an one-hop network composed of 12 TelosB motes and a 3-hop network composed of 24 TelosB motes.

7.1.2 Timeliness and Energy Consumption

In this section, 12 TelosB motes are organized into an one-hop network and each one corresponds to a node in OASIS [16]. We first evaluate the timeliness of our approach. As the network makes a detection decision per second, the delay bound of the system is one second. The average time of each component of the system is as follows: computing an energy scale for one second of seismic data take 6.7 ms; transmitting a TinyOS message with default size takes 9 ms; FFT and the local Bayesian detector take 164.7 ms. Therefore, our approach can achieve satisfactory timeliness on low-cost sensors with limited computational capability.

We now evaluate the energy consumption of various approaches. We measure the execution time of seismic processing and count the transmitted and received packets. The energy consumption is then estimated based on the measured current usage of processor and transceiver [12]. Figure 7 shows the energy consumption trace of Node 11 for 10 minutes. There is a significant earthquake event from the 245^{th} to 265^{th} second. As the byte length of encoded raw data increases in the presence of event, *data collection* has a spike during the earthquake. *STA/LTA* yields a high spike after the event, as it transmits compressed data at high speed after a detection. Note that *STA/LTA* has a false alarm at around the 550^{th} second. Figure 8 shows the corresponding breakdown of energy consumption. We can see that *Chair-Varshney* consumes a significant amount of energy in seismic processing, as it performs FFT on every node all the time. Suppose two carbon-zinc AA batteries are used, which have a total of 4680 J of energy storage [2]. The projected lifetime of a node is 19 days and 3.9 months for *data collection* and our approach, respectively.

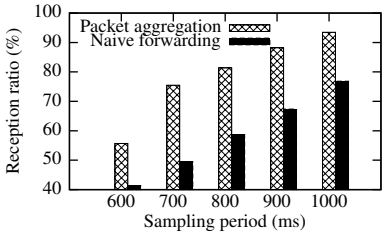

Figure 9. Reception ratio of energy scale information at the base station in a 3-hop network.

7.1.3 Communication Performance

We now evaluate the communication performance of our approach in a 3-hop network composed of 24 TelosB motes. We adopt the *naive forwarding* as the baseline approach, where an intermediate node forwards a received packet immediately without aggregation. Figure 9 plots the reception ratio of energy scale information at the base station versus sampling period. Due to limited wireless bandwidth, we observe low reception ratios when the sampling period is shorter than 600 ms. However, our approach can reach a reception ratio of 93.5% when the sampling period is one second which is consistent with the setting in real deployments [16, 21]. In contrast, *naive forwarding* only achieves a reception ratio of 77%.

7.2 Trace-driven Simulations

In addition to the testbed experiments, we also conduct simulations in TOSSIM [11] based on real data traces. The trace-driven simulations allow us to extensively evaluate the detection performance under a wide range of settings. Our evaluation is mainly focused on two aspects. First, we examines the detection performance of various approaches in

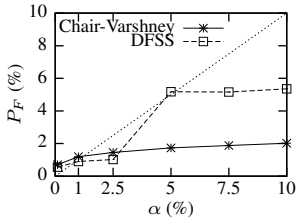

Figure 10. System P_F versus requested P_F.

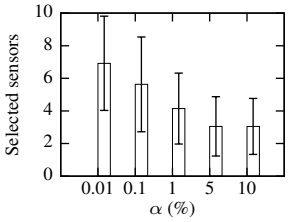

Figure 11. The number of selected sensors versus requested P_F.

a long period of time (based on the data traces that span 5.5 months). Second, we evaluate the configurability of our approach with respect to system sensing qualities such as false alarm rate.

Figure 10 plots the false alarm rate of the per-second system detection results versus the requested false alarm rate. We can see that when the requested P_F is greater than 5%, the measured P_F of our approach flats out, as Algorithm 1 can find a solution with minimum size of two sensors mostly. Moreover, as all sensors are always involved in the fusion process, *Chair-Varshney* has poor configurability as shown in Figure 10.

Figure 11 plots the number of selected sensors versus the requested false alarm rate. The error bar shows one standard deviation over 139 earthquake events. When lower performance requirement is imposed (*i.e.*, greater α), fewer sensors will be selected, which means less energy consumption. This result shows that our approach yields interesting trade-off between energy consumption and detection performance.

8 Conclusion

WSNs have been increasingly deployed for monitoring active volcanoes. This paper presents a quality-driven approach to detecting highly dynamical volcanic earthquakes based on in-network collaborative signal processing. In particular, we aim to minimize sensors' energy consumption subject to sensing quality requirements. Our approach is evaluated through testbed experiments and extensive simulations based on real data traces collected on Mount St. Helens. The results show that our approach can significantly reduce energy consumption compared with state-of-the-art approaches while providing assured system sensing quality.

Acknowledgment

This work is supported, in part, by the National Science Foundation under a CAREER grant CNS-0954039.

References

[1] `http://www.pbs.org/wnet/nature/forces/lava.html`.

[2] `http://www.allaboutbatteries.com/Energy-tables.html`.

[3] KissFFT, 2010. `http://sourceforge.net/projects/kissfft/`.

[4] K. Aki and P. Richards. *Quantitative seismology*. University Science Books, 2002.

[5] R. B. Ash and C. A. Doléans-Dade. *Probability & Measure Theory*. A Harcourt Science and Technology Company, 2nd edition, 1999.

[6] Z. Chair and P. Varshney. Optimal data fusion in multiple sensor detection systems. *IEEE Trans. Aerospace Electron. Syst.*, Jan. 1990.

[7] T. Clouqueur, K. K. Saluja, and P. Ramanathan. Fault tolerance in collaborative sensor networks for target detection. *IEEE Trans. Comput.*, 53(3), 2004.

[8] R. Duda, P. Hart, and D. Stork. *Pattern Classification*. Wiley, 2001.

[9] E. Endo and T. Murray. Real-time seismic amplitude measurement (RSAM): a volcano monitoring and prediction tool. *Bulletin of Volcanology*, 53(7), 1991.

[10] B. Gutenberg and C. F. Richter. Magnitude and energy of earthquakes. *Science*, 83(2147), 1936.

[11] P. Levis, N. Lee, M. Welsh, and D. Culler. TOSSIM: Accurate and scalable simulation of entire TinyOS applications. In *SenSys*, 2003.

[12] Moteiv. Telos (rev b): Preliminary datasheet. 2004.

[13] R. Niu and P. K. Varshney. Distributed detection and fusion in a large wireless sensor network of random size. *EURASIP J. Wireless Communications and Networking*, (4), 2005.

[14] X. Sheng and Y. Hu. Maximum likelihood multiple-source localization using acoustic energy measurements with wireless sensor networks. *IEEE Trans. Signal Process.*, 53(1), 2005.

[15] R. Sleeman and T. Van Eck. Robust automatic P-phase picking: an on-line implementation in the analysis of broadband seismogram recordings. *Physics of the Earth and Planetary Interiors*, 113, 1999.

[16] W. Song, R. Huang, M. Xu, A. Ma, B. Shirazi, and R. LaHusen. Air-dropped sensor network for real-time high-fidelity volcano monitoring. In *MobiSys*, 2009.

[17] R. Tan, G. Xing, J. Chen, W. Song, and R. Huang. Quality-driven volcanic earthquake detection using wireless sensor networks. Technical Report MSU-CSE-10-23, Computer Science and Engineering, Michigan State University, 2010.

[18] R. Tan, G. Xing, J. Wang, and H. C. So. Exploiting reactive mobility for collaborative target detection in wireless sensor networks. *IEEE Trans. Mobile Comput.*, 9(3), 2010.

[19] P. K. Varshney. *Distributed Detection and Data Fusion*. Springer, 1996.

[20] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh. Monitoring volcanic eruptions with a wireless sensor network. In *EWSN*, 2005.

[21] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *OSDI*, 2006.

[22] G. Werner-Allen, K. Lorincz, M. Ruiz, O. Marcillo, J. Johnson, J. Lees, and M. Welsh. Deploying a wireless sensor network on an active volcano. *IEEE Internet Computing*, 10(2), 2006.