Fusion-based Volcanic Earthquake Detection and Timing in Wireless Sensor Networks

RUI TAN, Michigan State University
GUOLIANG XING, Michigan State University
JINZHU CHEN, Michigan State University
WEN-ZHAN SONG, Georgia State University
RENJIE HUANG, Washington State University

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. In this paper, we propose a novel *quality-driven* approach to achieving real-time, distributed, and long-lived volcanic earthquake detection and timing. By employing novel in-network collaborative signal processing algorithms, our approach can meet stringent requirements on sensing quality (low false alarm/missing rate, short detection delay, 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, less than one second of detection delay, and millisecond precision earthquake onset time while achieving up to 6-fold energy reduction over the current data collection approach.

Categories and Subject Descriptors: C.3 [**Special-purpose and Application-based Systems**]: Signal processing systems

General Terms: Design, Measurement, Performance

Additional Key Words and Phrases: Volcano monitoring, earthquake detection, data fusion, wireless sensor network

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 [Educational Broadcasting Corp. 2010]. The recent eruptions of Volcano Eyjafjallajökull in Iceland caused the disruption of air traffic across Europe [BBC News 2010]. 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

Part of this work was published in IEEE RTSS 2010.
This work is supported, in part, by the National Science Foundation under grants CNS-0954039 (CAREER), CNS-0914371, CNS-0953067, and CNS-1066391.
Author's addresses: R. Tan, G. Xing, and J. Chen, Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824, USA; W.-Z. Song, Department of Computer Science, Georgia State University, Atlanta, GA 30303, USA; R. Huang, School of Electrical Engineering and Computer Science, Washington State University, Vancouver, WA 98686, USA.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

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 [Song et al. 2009] 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 [Song et al. 2009; Werner-Allen et al. 2005; Werner-Allen et al. 2006a] 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 accurately detect and localize earthquakes, 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 reduced system lifetime. Moreover, it has poor timeliness due to the limited bandwidth of low-cost sensors. For instance, as shown in [Werner-Allen et al. 2006a], 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 [Werner-Allen et al. 2006a], the existing earthquake detection algorithms [Endo and Murray 1991] 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 [Werner-Allen et al. 2006a].

In this paper, we push state of the art to real-time, distributed, 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 and timing. To this end, the following challenges must be addressed. First, volcanic earthquake is a sophisticated physical process featured by highly dynamic 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 reduced to improve timeliness and extend system lifetime.

We make the following major contributions in this paper:

— 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 reducing sensors' energy consumption.
— We develop new sensing algorithms based on the extensive analysis of real data traces collected on Mount St. Helens [Song et al. 2009]. 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 develop a two-phase earthquake onset time estimation approach. Specifically, in the first phase, the base station identifies individual earthquakes and estimates their coarse onset times by correlating the periodical detection results. In the second phase, each sensor locally executes an existing P-phase picking algorithm [Sleeman and Van Eck 1999] and outputs onset time estimates with millisecond precision. By taking advantage of accurate occurrence detection results, such a two-phase approach avoids both the raw data transmissions and unnecessary executions of the computation-intensive P-phase picker at sensors.
— 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 [Song et al. 2009] that contain more than 128 significant earthquake events. Experimental results show that our approach yields near-zero false alarm/missing rate, less than one second of detection delay, and millisecond precision earthquake onset time 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 discusses earthquake onset time estimation. Section 7 presents implementation details and Section 8 evaluates our approach. Section 9 concludes this paper.

2. RELATED WORK

In 2004, four MICA2 motes were deployed on Volcań Tungurahua in central Ecuador [Werner-Allen et al. 2005], 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 [Werner-Allen et al. 2006a; Werner-Allen et al. 2006b]. 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) [Endo and Murray 1991] 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%) [Werner-Allen et al. 2006a]. 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 [Song et al. 2009], 15 iMote2-based nodes has been aerially 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 and timing have 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 [Endo and Murray 1991; Sleeman and Van Eck 1999; Aki and Richards 2002]. 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 [Werner-Allen et al. 2006a; Werner-Allen et al. 2006b].

Distributed signal detection and data fusion in multi-sensor systems have been extensively studied in last decades [Varshney 1996; Chair and Varshney 1986; Tsitsiklis 1993]. These studies were focused on devising the optimal decision and fusion strategies that maximize the system performance of a given network. Recent studies on data fusion in WSNs have considered the specific properties of WSNs such as sensors' spatial distribution [Niu and Varshney 2005], limited sensing/communication capability [Clouqueur et al. 2004; Niu et al. 2006], and sensor failure [Tay et al. 2008]. Our previous works [Xing et al. 2009; Tan et al. 2009] investigate the fundamental limits of spatiotemporal coverage performance of fusion-based WSNs. In practice, various data fusion schemes have been employed in surveillance WSNs [He et al. 2004; Li et al. 2002] to improve system sensing performance. However, these systems often adopted simple detection heuristics without provable sensing quality. In contrast, we aim to develop quality-driven in-network signal processing algorithms to detect and time the highly dynamic volcanic earthquakes.

3. APPROACH OVERVIEW

In this section, we provide an overview of our approach to detecting and timing volcanic earthquakes in WSNs on active volcanoes. Our approach is designed to meet three key objectives of volcano monitoring. First, the system sensing quality must satisfy the Neyman-Pearson (NP) requirement [Duda et al. 2001] 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, to meet the requirements from advanced volcano monitoring applications (e.g., earthquake source localization and seismic tomography), the per-node earthquake onset time estimation must achieve the precision in the order of sensor's sampling period. Third, the computation and communication overhead of sensors must be reduced 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 [Song et al. 2009; Werner-Allen et al. 2006a]. 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. The algorithm framework of our approach is shown in Figure 1. Each sensor detects earthquake event every detection period based on seismic frequency spectrum. To handle the earthquake dynamics such as highly dynamic magnitude and variable 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, it is desirable to obtain the node-level earthquake onset times, which can be used by the volcanologists for advanced seismic processing tasks such as earthquake source localization and seismic tomography. In our approach, the base station first identifies an individual earthquake and estimates a coarse onset time.

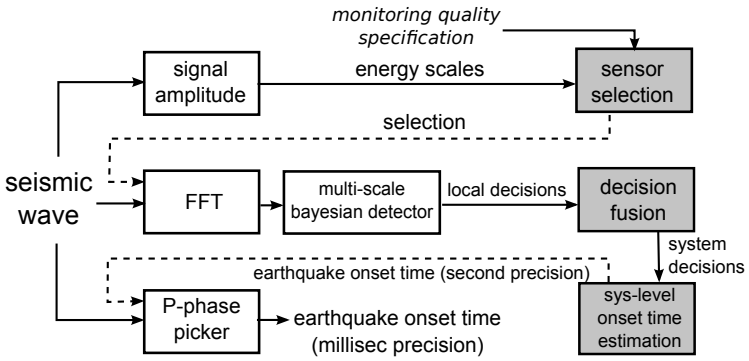

Fig. 1. Algorithm framework of volcanic earthquake detection and onset time estimation. Solid line represents data flow; dotted line represents control flow; white blocks are the components at a sensor; shadowed blocks are the components at the base station. The components *signal amplitude*, *FFT*, *mutli-scale bayesian detector* will be presented in Section 4; the components *sensor selection*, *decision fusion* will be presented in Section 5; the components *sys-level onset time estimation* and *P-phase picker* will be presented in Section 6.

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., [Sleeman and Van Eck 1999].

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 detection 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 FFT.

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 reduced 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

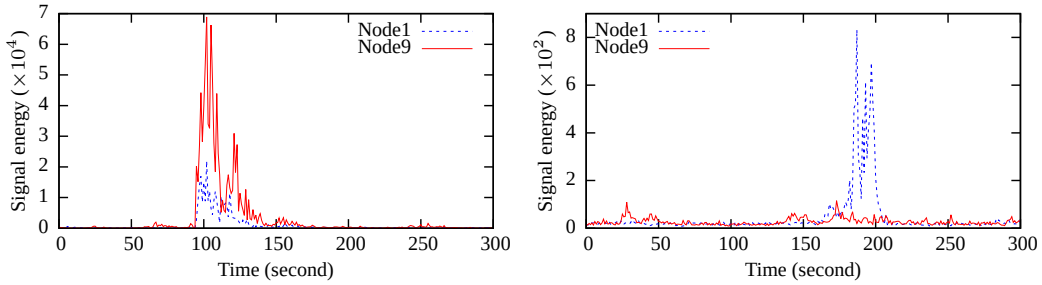

(a) Event 1 (b) Event 2

Fig. 2. Seismic signal energy.

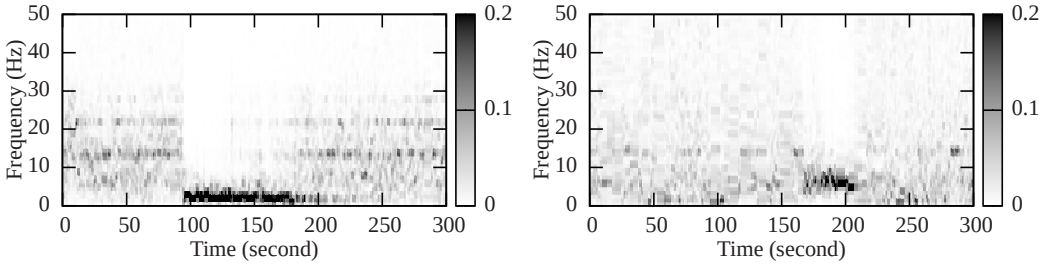

(a) Event 1 (b) Event 2

Fig. 3. Frequency spectrum of Node1.

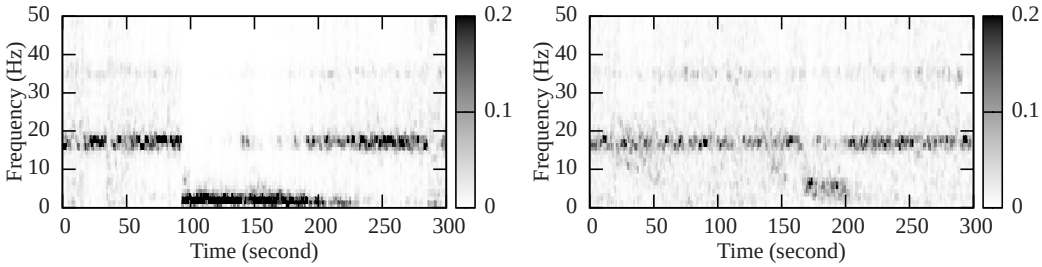

(a) Event 1 (b) Event 2

Fig. 4. Frequency spectrum of Node9.

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 [Song et al. 2009]. 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 [Niu and Varshney 2005; Clouqueur et al. 2004; Tan et al. 2010a] would not address the dynamics of volcanic earthquakes. Figures 3(a) and 3(b)

plot the frequency 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 frequency 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 propagates faster than S-wave and its frequency is typically from $1\,\text{Hz}$ to $10\,\text{Hz}$, while the slower S-wave often has a frequency of lower than $1\,\text{Hz}$ [Aki and Richards 2002]. 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 [Song et al. 2009], 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 $f\,\text{Hz}$. By applying FFT to the raw seismic data received during one second, a sensor obtains the frequency spectrum that ranges from $0\,\text{Hz}$ to $f/2\,\text{Hz}$. 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 tradeoff 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 *detection period* [Sheng and Hu 2005]. To be consistent with the frequency analysis, we let the detection 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 detection 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 lo-

cality 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 [Gutenberg and Richter 1936]. 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 [Song et al. 2009] and therefore the energy scale is from 1 to 6.

We now empirically assess the normal distribution assumption for the frequency feature using two datasets chosen from the data traces collected in the OASIS project [Song et al. 2009]. The two datasets comprise the frequency features when no earthquake happens and the earthquake energy scale p is 2, respectively. We then compute the squared Mahalanobis distance of each frequency feature $\mathbf{x}$, which is given by $(\mathbf{x} - \mathbf{m})^{\text{T}}\mathbf{C}^{-1}(\mathbf{x} - \mathbf{m})$. If $\mathbf{x}$ follows the n-dimensional normal distribution, its squared Mahalanobis distance follows the χ^2_n distribution. Therefore, we calculate the quantile-quantile (Q-Q) plot of the squared Mahalanobis distance versus the χ^2_n distribution.

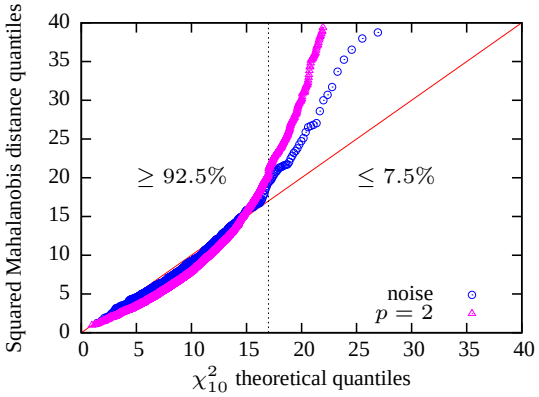

Fig. 5. The quantile-quantile plot of the squared Mahalanobis distance versus the χ^2_{10} distribution. For both cases where no earthquake happens (i.e., noise) and the earthquake energy scale is 2 (i.e., $p = 2$), more than 92.5% data points match the expected distribution.

Figure 5 shows the Q-Q plots for the two datasets where the dimension is set to be 10. For both two datasets, more than 92.5% data points (on the left side of the dashed line) reside close to the diagonal line, which shows that most of the frequency features match the normal distribution. Less than 7.5% data points rest above the diagonal line, which indicates that the true distributions have long tails. In summary, the normal distribution is a good approximate distribution for describing the statistical behavior of the frequency feature.

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. Moreover, they can be turned on periodically to re-train the sensors to capture the changed signal characteristics and volcano activity pattern.

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

$$g_i(\mathbf{x}) = \ln \mathbb{P}(H_i) - \frac{1}{2} \ln |\mathbf{C}_i| - \frac{1}{2} \left((\mathbf{x} - \mathbf{m}_i)^{\mathrm{T}} \mathbf{C}_i^{-1} (\mathbf{x} - \mathbf{m}_i) \right),$$

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 $\mathbf{C}_i$ [Duda et al. 2001]. 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

$$P_F = \int_{\mathcal{R}} \phi(\mathbf{x}|\mathbf{m}_0, \mathbf{C}_0)\mathrm{d}\mathbf{x}, \qquad P_D = \int_{\mathcal{R}} \phi(\mathbf{x}|\mathbf{m}_p, \mathbf{C}_p)\mathrm{d}\mathbf{x},$$

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 of 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.

The prior probabilities (i.e., $\mathbb{P}(H_i)$) can be estimated based on historical earthquake records. When the prior probabilities are unknown, we can use the minimax error rate criterion [Duda et al. 2001], of which the error rate does not change with the prior probabilities. The minimax decision functions are $g_i(\mathbf{x}) = -\frac{1}{2}\left((\mathbf{x}-\mathbf{m}_i)^{\mathrm{T}}\mathbf{C}_i^{-1}(\mathbf{x}-\mathbf{m}_i)\right) + \delta_i$, where the constants δ_i make $P_F = P_D$. It is difficult to obtain the analytical δ_i. In practice, the base station can numerically search for the δ_i that ensures $P_F = P_D$. Moreover, the search can be sped up by leveraging the monotonic relationship between P_F and P_D.

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 qualities to achieve maximum system detection performance. Moreover, the sensor selection avoids unnecessary expensive feature extraction at the sensors with low signal qualities. 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 [Varshney 1996], *decision fusion* is preferable for WSNs due to its low communication cost [Clouqueur et al. 2004]. We use a widely adopted decision fusion model called equal gain combining (EGC) [Niu and Varshney 2005; Clouqueur et al. 2004; Tan et al. 2010a] 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. We denote α_i and β_i as the false alarm rate and detection probability of sensor i. As discussed in Section 4.4, sensor i has a pair of (α_i, β_i) for each energy scale p, where $p \in [1, K]$. 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

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

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

$$\mathbb{E}[\Lambda|H_0] = \sum_{i=1}^{n} \mathbb{E}[I_i|H_0] = \sum_{i=1}^{n} \alpha_i, \qquad \mathrm{Var}[\Lambda|H_0] = \sum_{i=1}^{n} \mathrm{Var}[I_i|H_0] = \sum_{i=1}^{n} \alpha_i - \alpha_i^2.$$

Lyapunov's central limit theorem (CLT) [Ash and Doléans-Dade 1999] 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 [Tan et al. 2010a]. 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

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

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

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. We first examine the sensing performance diversity of sensors based on data traces collected in OASIS [Song et al. 2009]. 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 [Duda et al. 2001], which is a widely adopted detectability measure, between

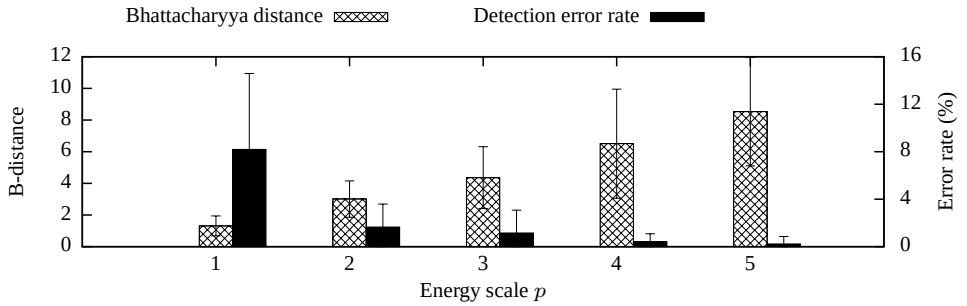

Fig. 6. Bhattacharyya distance and corresponding detection error rate versus energy scale. The results show the standard deviation over 12 sensors.

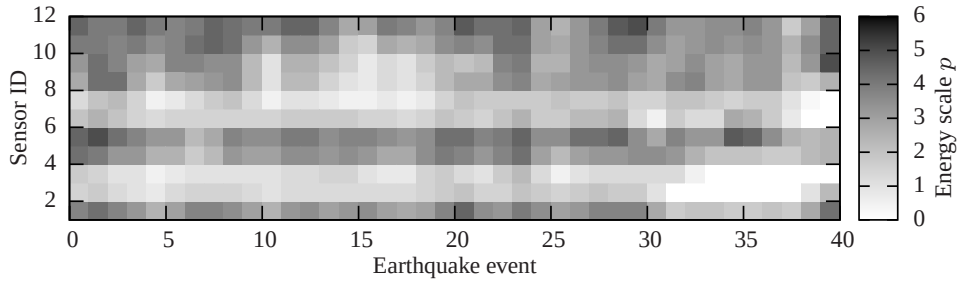

Fig. 7. Dynamics of energy scale received by sensors.

the p^{th} distribution $\mathcal{N}_p$ and the noise distribution $\mathcal{N}_0$ within its multi-scale Bayesian model. Figure 6 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 higher signal energy. Moreover, sensors show significant performance variation for the same energy scale. Figure 7 plots the maximum energy scale measured by sensors in 40 earthquake events.

We can make two important observations from Figures 6 and 7. 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 the 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 pre-computed 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 detection 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. In our scheme, the selected sensors only send binary local detection decisions to the base station. Hence, compared with the energy consumed in the feature extraction by the costly FFT, the communication cost is small. Therefore, the problem formulation in this paper is only focused on minimizing the number of selected sensors. We adopt the Neyman-Pearson

(NP) criterion [Duda et al. 2001] 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. From the NP lemma [Varshney 1996], 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. 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 [Tan et al. 2010b]. 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 detection 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. Near-optimal Solution

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 suggests a near-optimal solution with polynomial complexity.

We first analyze the condition for the NP criterion. We assume that n is large enough such that the detection performance approximations made in Section 5.1 are accurate. We will discuss how to deal with the inaccuracy caused by small n in Section 5.4. According to the NP lemma [Varshney 1996], 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

$$\eta = \sum_{i=1}^{n} \alpha_i + Q^{-1}(\alpha) \cdot \sqrt{\sum_{i=1}^{n} \alpha_i - \alpha_i^2}, \tag{3}$$

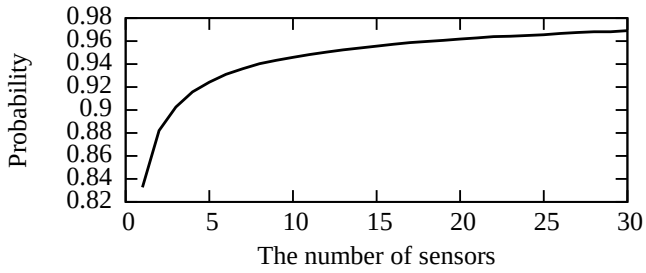

Fig. 8. The probability that f increases with $\sum_{i=1}^{n}(\alpha_i - \beta_i)$ versus the number of sensors.

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

$$f = \frac{Q^{-1}(\alpha)\sqrt{\sum_{i=1}^{n}\alpha_i - \alpha_i^2} + \sum_{i=1}^{n}(\alpha_i - \beta_i)}{\sqrt{\sum_{i=1}^{n}\beta_i - \beta_i^2}}. \quad (4)$$

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 false alarm rate and miss rate of sensor i is given by α_i and $1 - \beta_i$, respectively. Therefore, the error rate of sensor i is $\frac{\alpha_i + 1 - \beta_i}{2}$ and the average error rate over all sensors is $\frac{1}{n} \cdot \sum_{i=1}^{n} \frac{\alpha_i + 1 - \beta_i}{2}$. Hence, if n is fixed, $\sum_{i=1}^{n}(\alpha_i - \beta_i)$ is a metric for quantifying the system detection performance. The advantage of using this metric is that it is a linear combination of sensors' local performance metrics, which enables us to develop efficient sensor selection algorithms. We now investigate the monotonic relationship between the function f and $\sum_{i=1}^{n}(\alpha_i - \beta_i)$. We conduct Monte Carlo simulation to evaluate the probability that f increases with $\sum_{i=1}^{n}(\alpha_i - \beta_i)$. For each trial of the Monte Carlo simulation, two points are randomly and uniformly sampled from the $2n$-dimensional space $\{\alpha_i, \beta_i | i \in [1, n], \alpha_i \in (0, 1), \beta_i \in (0, 1)\}$ to evaluate the increasing relationship between f and $\sum_{i=1}^{n}(\alpha_i - \beta_i)$. After a large number of trials, the probability for the increasing relationship is shown in Fig. 8. We can see that the increasing relationship holds with a probability of at least 90% when the number of sensors exceeds 3. Therefore, in practice, we can find the near-optimal solution to the sub-problem by selecting n sensors with minimized $\sum_{i=1}^{n}(\alpha_i - \beta_i)$, which can be easily solved by sorting sensors ascendingly according to the value of $(\alpha_i - \beta_i)$.

5.4. Dynamic Sensor Selection Algorithm

In this section, we develop a dynamic sensor selection algorithm to solve Problem 1 based on the analysis in Section 5.3. Before presenting the algorithm, we first discuss how to handle the inaccuracy of the normal approximations made in Section 5.1. Re-

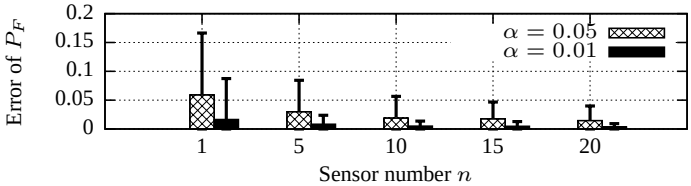

Fig. 9. Error of P_F due to approximation (99.5% confidence level).

Table I. An example of sensor detection performance look-up table ($K = 5$).

Sensor ID	1				···	N			
Energy scale p_i	1	2	···	K	···	1	2	···	K
(α_i, β_i) (%)	(6.9, 96.8)	(5.9, 98.1)	···	(0.008, 99.9)	···	(7.3, 95.3)	(4.8, 97.8)	···	(0.0036, 99.9)

Algorithm 1 Dynamic sensor selection algorithm

Input: (1) sensor detection performance look-up table
(2) energy scales measured by sensors, $\{p_i | i \in [1, N]\}$
(3) system detection performance requirements $\{\alpha, \beta\}$
Output: minimum subset S, detection threshold η

```
query the look-up table using {p_i | i ∈ [1,N]} to find {α_i, β_i | i ∈ [1,N]}
sort sensors according to (α_i − β_i) ascendingly
for n = 1 to N do
  if n < n_s then
    for all subset S with size of n do
      compute the PMF of Λ using (1)
      if exists η such that system P_F ≤ α and P_D ≥ β then
        return S and η
      end if
    end for
  else
    S = { top n sensors }
    compute f with S using (4)
    if Q(f) ≥ β then
      compute η with S using (3)
      return S and η
    end if
  end if
end for
exit with no solution
```

garding the divide-and-conquer strategy proposed in Section 5.3, when n is small, we compute the PMF of Λ (given by (1)) and then search for the optimal detection threshold. A question is when we should switch to the normal approximations. We propose a numerical approach to determine the switching point for n, denoted by n_s. We investigate the impact of n on the accuracy of system false alarm rate. Specifically, we first compute the detection threshold η using (3) with randomly generated local false alarm rates, i.e., α_i. We then compute the true P_F with the threshold η using (1). Figure 9 plots the error between the requested P_F (i.e., α) and the true P_F. For instance, if $\alpha = 0.05$ and $n = 15$, the maximum error of P_F is about 0.05 and hence the true P_F is within $[\alpha - 0.05, \alpha + 0.05]$. The figure also shows that if more stringent requirement is imposed, i.e., smaller α, the error decreases accordingly. We can evaluate the impact of n on the accuracy of detection probability as well. With such an approach, we can choose the switching point n_s to achieve desired accuracy.

We now present the dynamic sensor selection algorithm to solve Problem 1. As discussed in Section 4.4, in the training phase, the base station computes and stores sensors' false alarm rates and detection probabilities, which can be organized into a look-up table. For instance, Table I shows the look-up table constructed for the 12 sensors deployed in the OASIS project [Song et al. 2009]. With the sensor detection performance look-up table, we develop the dynamic sensor selection algorithm, which is 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 with no solution (Line 20). As a result, no sensor will be selected and hence the frequency feature extraction by the costly FFT can be avoided.

6. EARTHQUAKE ONSET TIME ESTIMATION

In addition to the accurate detection of earthquake occurrences, another import requirement of volcano monitoring is to identify individual earthquakes as well as estimate their onset times and durations [Werner-Allen et al. 2006a; Aki and Richards 2002; Sleeman and Van Eck 1999]. In particular, the fine-grained per-node earthquake onset times are critical for advanced volcano monitoring applications, such as earthquake source localization and seismic tomography. Traditionally, the volcanologists often manually pick the per-node onset times from the collected raw seismic data, which is an extremely labor-intensive process. In this paper, we aim to develop an in-network earthquake onset time estimation approach. The automatically picked onset times significantly narrow the manual search regions for more accurate onset times. Moreover, only a small amount of raw seismic data around the automatically picked onset times need to be sent to the base station when further manual picks are still needed. In this section, we develop a two-phase approach to estimating earthquake onset time with millisecond precision. Specifically, in the first phase, the base station correlates the per-second detection results yielded by the decision fusion process to identify the onset time and duration of individual earthquakes, which is referred to as *system-level* onset time estimation. In the second phase, given the system-level onset time, each sensor locally executes an existing P-phase picking algorithm [Sleeman and Van Eck 1999] and outputs onset time estimate with improved accuracy, which is referred to as *node-level* onset time estimation. Note that in seismology, P-phase refers to the arrival time of P-wave front. By taking advantage of accurate occurrence detection results, the computation-intensive P-phase picker is executed at sensors only when an earthquake is detected with high confidence (i.e., low false alarm rate and high detection probability). Moreover, the system-level onset time estimation narrows the range of time, over which the sensors have to run the P-phase picker to achieve node-level fine-grained onset time estimation.

6.1. System-level Onset Time Estimation

In this section, we discuss how to temporally correlate the per-second detection results yielded by the decision fusion process to identify individual earthquake events that are described by the earthquake onset time and duration. Although the system detection performance can be improved by fusing sensors' local detection results, the possibilities of false alarms and misses cannot be completely eliminated. In particular, they must be properly dealt with in order to precisely estimate the onset time of earthquake events. Figure 10(a) shows the per-second decision sequence yielded by the decision fusion process. A key observation from real measurements is that a true earthquake event often generates clustered detection results and hence isolated positive or negative decisions are likely false alarms and misses. Based on this observation, we propose to use mathematical morphology [Soille 2003] to identify individual events in the presence of random detection errors. Mathematical morphology is a widely adopted tool to identify geometrical structures in image processing.

In the morphological processing, by applying the *opening* operator [Soille 2003] on the decision sequence, multiple continuous false alarms can be eliminated if the number of these false alarms is less than the operator diameter. Moreover, by applying the *closing* operator [Soille 2003], multiple continuous misses can be restored as successful detections if the number of these misses is less than the operator diameter.

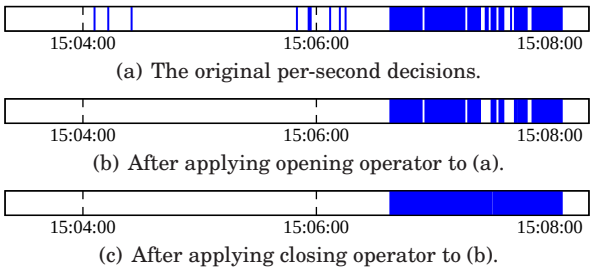

(a) The original per-second decisions.

(b) After applying opening operator to (a).

(c) After applying closing operator to (b).

Fig. 10. Morphological processing of an earthquake event that occurred on Mt. St. Helens, October 31, 2009.

We have implemented the two morphological operators and applied the combination of them to the per-second decision sequence. In this work, the diameter of the opening operator is set to be 3. Under such a setting, we ignore any detected event that lasts for shorter than 3 seconds. As shown in Figure 10(b), several isolated false alarms before the earthquake event are eliminated by the opening operator. The diameter of the closing operator should be set according to the earthquake *recurrence interval* [Aki and Richards 2002] of the volcano. In this work, we set 11 and the closing operator shows satisfactory performance. As shown in Figure 10(c), the closing operator can join the fragmented positive decisions to yield a single earthquake event. Given the morphologically processed detection decisions, it is easy to determine the earthquake onset time and duration.

6.2. Node-level Onset Time Estimation

The granularity of the system-level onset time estimation is equal to the detection period, which is one second in our system. However, localizing earthquake source requires the onset times estimated by spatially distributed nodes with millisecond precision [Aki and Richards 2002]. We now discuss our approach to node-level onset time estimation with improved accuracy. Once the morphological filter presented in Section 6.1 yields a system-level onset time estimate, the base station notifies the sensors to perform node-level onset time estimation by running the automatic P-phase picking algorithm proposed in [Sleeman and Van Eck 1999]. Given a coarse estimate of earthquake onset time, the algorithm constructs two regression models for the buffered seismic data before and after the coarse onset time, respectively. The algorithm then picks a millisecond onset time that has maximum likelihood to match the regression models. It is important to note that the algorithm relies on correct onset time estimate with the granularity of one second, which is guaranteed by the other modules in our approach.

7. 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 [Song et al. 2009] 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: 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 [Song et al. 2009]. 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 detection period.

When the detection period is set to be one second, the sampling rate is consistent with previous deployments [Song et al. 2009; Werner-Allen et al. 2006a].

Seismic processing: We use the `KissFFT` [Borgerding 2010] 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. We modified the P-phase picking software developed in [Sleeman and Van Eck 1999] to run in TinyOS. The picker requires 6.5KB ROM and 12.5KB RAM. In this work, we only evaluated the performance of the P-phase picker in TinyOS simulator (i.e., TOSSIM) [Levis et al. 2003].

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 detection 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 detection periods. When a sensor has received all packets from its children for the current detection period, it sends out the aggregated packets for the current and previous detection 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 earthquake detection.

8. PERFORMANCE EVALUATION

We conduct testbed experiments and extensive simulations based on real data traces collected by 12 nodes in the OASIS project [Song et al. 2009]. 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 lasts for 10 minutes and contains one or more earthquake events. In Section 8.1, we present the experimental results on energy usage and communication performance using a testbed of 24 TelosB motes. In Section 8.2, we present the simulation results on earthquake detection and onset time estimation in TOSSIM.

8.1. Testbed Experiments

8.1.1. 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 Node 01 in the deployment. The STA/LTA threshold is set to be 2, which is suggested by the volcanologists at U.S. Geological Survey [Song et al. 2009]. 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 [Song et al. 2009] 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 [Endo and Murray 1991]. A node sends its detection result to the base station only when the decision is positive. If more than

30% nodes make positive decisions, the base station makes a positive system detection decision. Note that these settings are consistent with the detection approach in [Werner-Allen et al. 2006a]. (3) In the *Chair-Varshney* approach, each node performs FFT and makes a local detection decision every detection period. The base station fuses the local decisions by the Chair-Varshney's rule [Chair and Varshney 1986] that is the optimal decision fusion model. Specifically, the test statistic is $\Lambda = \sum_{i=1}^{n} \log \frac{\beta_i(1-\alpha_i)}{\alpha_i(1-\beta_i)} \cdot I_i$, where I_i is the local detection decision of sensor 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. The one-hop topology allows us to evaluate the basic performance of various approaches without the impact of topology-dependent factors such as extra computation and communication overhead caused by packet aggregation in multi-hop networks. As real deployments [Song et al. 2009; Werner-Allen et al. 2006a] typically adopt multi-hop tree topology, we also aim to evaluate the impact of multi-hop tree topology on the proposed approach. In the 3-hop tree network experiment, we will evaluate the communication performance of our approach under various timeliness requirements. Note that the setting of hop count in our evaluation is comparable with the settings in several real deployments, e.g., six hops in [Werner-Allen et al. 2006a] and five hops in [Song et al. 2009]. We also note that, compared with traditional data collection approaches, the advantages of our fusion-based approach are likely more evident on larger networks.

8.1.2. 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 [Song et al. 2009]. We first evaluate the timeliness of our approach. As discussed in Section 6.2, accurate per-second detection is a prerequisite of the millisecond onset time estimation. Hence, one second is the delay bound in each detection period. 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.

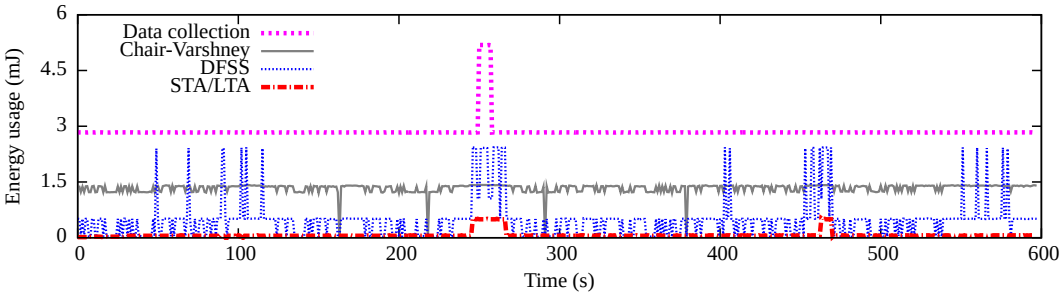

Fig. 11. Energy consumption of Node 11.

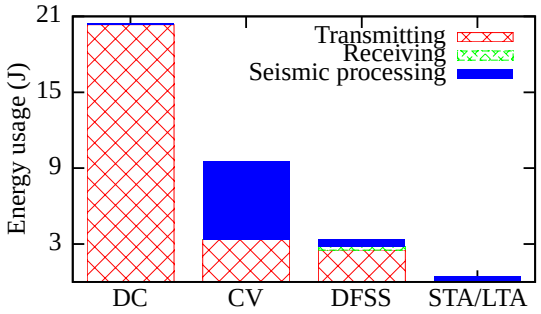

Fig. 12. Total energy consumption of 12 nodes for 10 minutes. DC represents *data collection*, CV represents *Chair-Varshney*.

The energy consumption is then estimated based on the measured current usage of processor and transceiver [Moteiv Corp. 2004]. Figure 11 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. *Chair-Varshney* consumes more energy than *DFSS* but less than *data collection*. *STA/LTA* consumes the least energy. However, *STA/LTA* has a false alarm at around the 460^{th} second. Figure 12 shows the corresponding breakdown of energy consumption. We can see that *data collection* consumes a large amount of energy in transmitting raw data. As *Chair-Varshney* performs FFT on every node all the time, it consumes a significant amount of energy in seismic processing. *STA/LTA* consumes a little energy in executing the STA/LTA algorithm on the seismic data. Moreover, as only positive local detection decisions are transmitted, *STA/LTA* significantly reduces the energy consumed in data transmission. However, as discussed in Section 2, the heuristic STA/LTA algorithm cannot yield provable detection performance. In practice, the STA/LTA algorithm is best used to trigger data collection. Therefore, the STA/LTA algorithm often adopts conservative settings (e.g., low threshold) in order not to miss any earthquake event even though a lot of false alarms will be raised. Suppose two carbon-zinc AA batteries are used, which have a total of $4680\,\text{J}$ of energy storage [AllAboutBatteries.com 2011]. The projected lifetime of a node is 19 days and 3.9 months for *data collection* and our approach, respectively.

8.1.3. 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

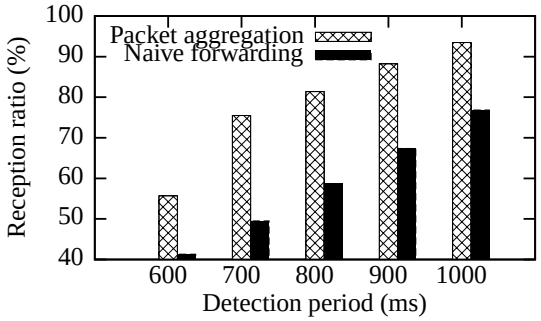

Fig. 13. Reception ratio of energy scale information at the base station in a 3-hop network.

packet immediately without aggregation. Figure 13 plots the reception ratio of energy scale information at the base station versus detection period. Due to limited wireless bandwidth, we observe low reception ratios when the detection period is shorter than 600 ms. However, our approach can reach a reception ratio of 93.5% when the detection period is one second, which is consistent with the setting in real deployments [Song et al. 2009; Werner-Allen et al. 2006a]. In contrast, *naive forwarding* only achieves a reception ratio of 77%.

8.2. Trace-driven Simulations

In addition to the testbed experiments, we also conduct simulations in TOSSIM [Levis et al. 2003] based on real data traces. The trace-driven simulations allow us to extensively evaluate the detection and timing 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 a long period of time (based on the data traces that span 5.5 months) and evaluate the configurability of our approach with respect to system sensing qualities such as false alarm rate. Second, we evaluate the performance of our two-phase earthquake onset time estimation approach.

8.2.1. Detection Performance and Configurability. Figure 14 plots the false alarm rate of the per-second system detection results without morphological processing versus the requested false alarm rate. Note that the false alarm rate is calculated by the ratio of the number of positive per-second system decisions to the number of detection periods when no earthquake happens. 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 14.

Figure 15 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.

Figure 16 plots the number of false alarms and misses of earthquake events after the morphological processing. Note that the earthquake event is described by the onset time and duration. We can conclude that our approach generates fewer detection errors than *STA*/*LTA* and has comparable detection performance with *Chair-Varshney*. However, as shown in Figures 11 and 12, *Chair-Varshney* consumes signif-

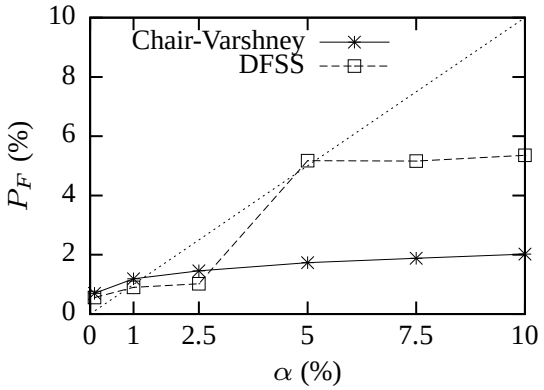

Fig. 14. System P_F versus requested P_F.

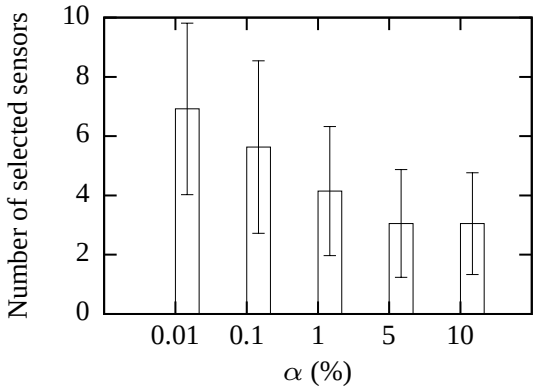

Fig. 15. The number of selected sensors versus requested P_F.

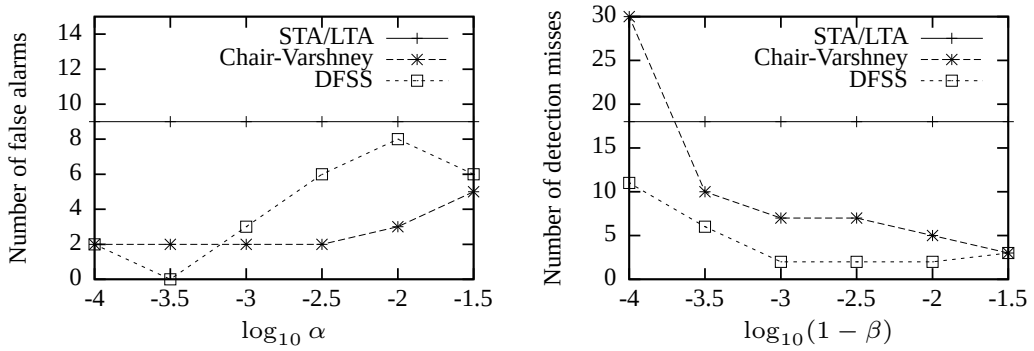

(a) False alarms (b) Detection misses

Fig. 16. Number of detection errors after morphological processing. The test data lasts for 19 hours and includes 139 earthquake events.

icantly more energy. When the detection performance requirement is extremely high (e.g., $\beta = 1 - 10^{-4}$), the sensor selection algorithm exits with no solution and hence no detection is made as discussed in Section 5.4. As a result, the system misses more events as shown in Figure 16(b). Therefore, in practice, the requirement on detection probability should be set in the achievable range to avoid the saturation.

8.2.2. Onset Time Estimation Performance. In this section, we evaluate the performance of our two-phase onset time estimation approach. In the simulations, when the base station yields a system-level earthquake onset time, the raw seismic data that are 10 seconds before and 6 seconds after the system-level onset time are fed into the P-phase picking algorithm [Sleeman and Van Eck 1999] at each sensor. Figure 17 shows the node-level onset times (represented by vertical lines) estimated by four sensors in a typical earthquake event. The Y-axis of the figures represents the slant distance from the vent. The event shows a suddenly strong shake at its start and the sensor closest to the vent (i.e., Node 04) receives the signal before others. This indicates that the event is a typical explosive earthquake at the vent.

We manually check the node-level earthquake onset time estimation results of 112 earthquakes that are detected by our *DFSS* approach (after excluding false alarms). For each earthquake, we first manually pick the onset times of all sensors from their seismic waves. If the difference between the output of the P-phase picking algorithm

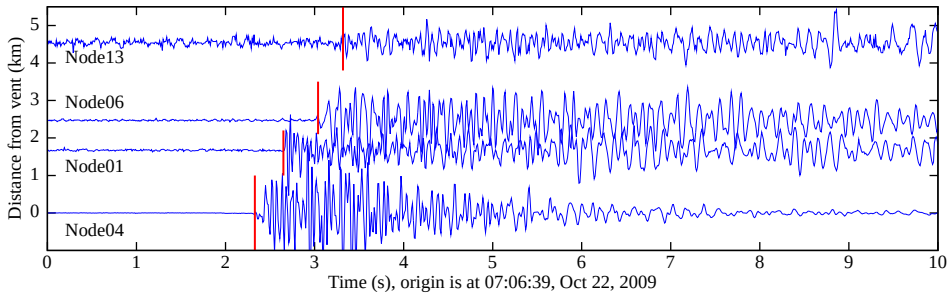

Fig. 17. An example of node-level onset time estimation. System-level onset time is at the 3^{th} second.

[Sleeman and Van Eck 1999] and the manual pick is less than one second, the onset time estimate is regarded to be a *good estimate*. Our analysis shows that 68.8% of node-level onset time estimates are good estimates. We summarize the reasons for the remaining 31.2% inaccurate node-level onset time estimates as follows. First, ideally, only a subset of sensors receiving prominent earthquake signals should perform the node-level onset time estimation. However, all sensors in the simulations are required to perform the estimation when an earthquake is detected. As a result, the sensors receiving insignificant earthquake signals yield inaccurate onset time estimates. As the P-phase picking algorithm [Sleeman and Van Eck 1999] operates in time domain, the sensor selection approach described in Section 5, which selects sensors with prominent features in frequency domain, cannot be readily applied to select the sensors for node-level onset time estimation. In our future work, we will explore efficient sensor selection approaches for node-level onset time estimation. Second, the dataset used in our simulations does not contain the complete data due to packet losses and communication errors in data collection, which undermines the performance of the P-phase picking algorithm. In a real deployment, this issue can be avoided since each sensor locally performs the P-phase picking algorithm on the buffered data. In summary, the simulation results demonstrate the feasibility of the proposed two-phase onset time estimation approach. We note that the performance of the two-phase onset time estimation approach can be improved by leveraging more accurate P-phase picking algorithms when they are available.

Finally, we manually choose 63 earthquake events and evaluate the node-level onset time estimation results. Figure 18 plots the average delay of P-phase arrival with respect to Node 01 which is a sensor on the middle of the volcano. We can see from the figure that, on average, the sensors on the middle of the volcano first receive the P-wave before the sensors that are close to or far away from the vent. The time difference is within 180 milliseconds. Moreover, the sensors that have equal distance to the

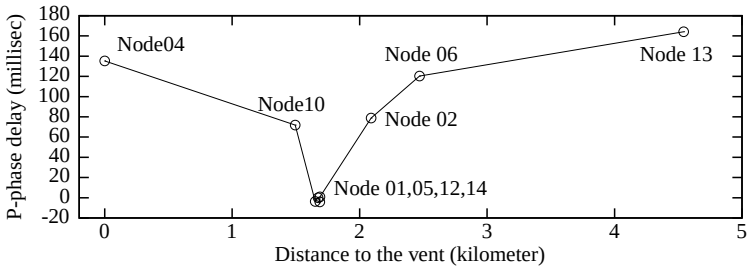

Fig. 18. Delay of P-phase arrival (with respect to Node 01) versus distance to the vent. The results are the average over 63 earthquake events.

vent (i.e., Node 01, 05, 12 and 14) have consistent P-phase arrivals. The above results are consistent with the findings from an earlier study [Werner-Allen et al. 2006a] in which the coarse onset times were identified manually to assist the P-phase picking algorithm.

9. CONCLUSION

WSNs have been increasingly deployed for monitoring active volcanoes. This paper presents a quality-driven approach to detecting and timing highly dynamic volcanic earthquakes based on in-network collaborative signal processing. In particular, we aim to reduce 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.

Acknowledgments

We hereby thank Geoffrey Werner Challen and other anonymous reviewers for reviewing this paper and providing constructive comments.

REFERENCES

AKI, K. AND RICHARDS, P. 2002. *Quantitative seismology*. University Science Books.

ALLABOUTBATTERIES.COM. 2011. http://www.allaboutbatteries.com/Energy-tables.html.

ASH, R. B. AND DOLÉANS-DADE, C. A. 1999. *Probability & Measure Theory* 2nd Ed. A Harcourt Science and Technology Company.

BBC NEWS. 2010. Volcano erupts in south iceland. http://news.bbc.co.uk/2/hi/8578576.stm.

BORGERDING, M. 2010. KissFFT project. http://sourceforge.net/projects/kissfft/.

CHAIR, Z. AND VARSHNEY, P. 1986. Optimal data fusion in multiple sensor detection systems. *IEEE Transactions on Aerospace and Electronic Systems AES-22,* 1, 98–101.

CLOUQUEUR, T., SALUJA, K. K., AND RAMANATHAN, P. 2004. Fault tolerance in collaborative sensor networks for target detection. *IEEE Transactions on Computers 53,* 3, 320–333.

DUDA, R., HART, P., AND STORK, D. 2001. *Pattern Classification*. Wiley.

EDUCATIONAL BROADCASTING CORP. 2010. Forces of the wild. http://www.pbs.org/wnet/nature/forces/lava.html.

ENDO, E. AND MURRAY, T. 1991. Real-time seismic amplitude measurement (RSAM): a volcano monitoring and prediction tool. *Bulletin of Volcanology 53,* 7, 533–545.

GUTENBERG, B. AND RICHTER, C. F. 1936. Magnitude and energy of earthquakes. *Science 83,* 2147, 183–185.

HE, T., KRISHNAMURTHY, S., STANKOVIC, J. A., ABDELZAHER, T., LUO, L., STOLERU, R., YAN, T., GU, L., HUI, J., AND KROGH, B. 2004. Energy-efficient surveillance system using wireless sensor networks. In *The 2nd International Conference on Mobile Systems, Applications, and Services (MobiSys)*. 270–283.

LEVIS, P., LEE, N., WELSH, M., AND CULLER, D. 2003. TOSSIM: Accurate and scalable simulation of entire TinyOS applications. In *The 1st ACM Internatinal Conference on Embedded Networked Sensor Systems (SenSys)*. 126–137.

LI, D., WONG, K., HU, Y.-H., AND SAYEED, A. 2002. Detection, classification and tracking of targets in distributed sensor networks. *IEEE Signal Processing Magazine 19,* 2, 17–29.

MOTEIV CORP. 2004. Telos (rev b): Preliminary datasheet.

NIU, R., CHEN, B., AND VARSHNEY, P. 2006. Fusion of decisions transmitted over Rayleigh fading channels in wireless sensor networks. *IEEE Transactions on Signal Processing 54,* 3, 1018–1027.

NIU, R. AND VARSHNEY, P. K. 2005. Distributed detection and fusion in a large wireless sensor network of random size. *EURASIP Journal on Wireless Communications and Networking* 4, 462–472.

SHENG, X. AND HU, Y. 2005. Maximum likelihood multiple-source localization using acoustic energy measurements with wireless sensor networks. *IEEE Transactions on Signal Processing 53,* 1, 44–53.

SLEEMAN, R. AND VAN ECK, T. 1999. Robust automatic P-phase picking: an on-line implementation in the analysis of broadband seismogram recordings. *Physics of the Earth and Planetary Interiors 113*, 265–275.

SOILLE, P. 2003. *Morphological image analysis: principles and applications*. Springer-Verlag.

SONG, W., HUANG, R., XU, M., MA, A., SHIRAZI, B., AND LAHUSEN, R. 2009. Air-dropped sensor network for real-time high-fidelity volcano monitoring. In *The 7th Annual International Conference on Mobile Systems, Applications and Services (MobiSys)*. 305–318.

TAN, R., XING, G., CHEN, J., SONG, W., AND HUANG, R. 2010b. Quality-driven volcanic earthquake detection using wireless sensor networks. Tech. Rep. MSU-CSE-10-23, Computer Science and Engineering, Michigan State University.

TAN, R., XING, G., LIU, B., AND WANG, J. 2009. Impact of data fusion on real-time detection in sensor networks. In *The 30th IEEE Real-Time Systems Symposium (RTSS)*. 323–332.

TAN, R., XING, G., WANG, J., AND SO, H. C. 2010a. Exploiting reactive mobility for collaborative target detection in wireless sensor networks. *IEEE Transactions on Mobile Computing 9,* 3, 317–332.

TAY, W., TSITSIKLIS, J., AND WIN, M. 2008. On the Impact of Node Failures and Unreliable Communications in Dense Sensor Networks. *IEEE Transactions on Signal Processing 56,* 6, 2535–2546.

TSITSIKLIS, J. N. 1993. *Advances in Signal Processing*. JAI Press, Chapter Decentralized Detection, 297–344.

VARSHNEY, P. K. 1996. *Distributed Detection and Data Fusion*. Springer.

WERNER-ALLEN, G., JOHNSON, J., RUIZ, M., LEES, J., AND WELSH, M. 2005. Monitoring volcanic eruptions with a wireless sensor network. In *The 2nd European Workshop on Wireless Sensor Networks (EWSN)*. 108–120.

WERNER-ALLEN, G., LORINCZ, K., JOHNSON, J., LEES, J., AND WELSH, M. 2006a. Fidelity and yield in a volcano monitoring sensor network. In *The 7th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 381–396.

WERNER-ALLEN, G., LORINCZ, K., RUIZ, M., MARCILLO, O., JOHNSON, J., LEES, J., AND WELSH, M. 2006b. Deploying a wireless sensor network on an active volcano. *IEEE Internet Computing 10,* 2, 18–25.

XING, G., TAN, R., LIU, B., WANG, J., JIA, X., AND YI, C.-W. 2009. Data fusion improves the coverage of wireless sensor networks. In *The 15th Annual International Conference on Mobile Computing and Networking (MobiCom)*. 157–168.