

Volcanic Earthquake Timing Using Wireless Sensor Networks

Guojin Liu^{1,2}, Rui Tan^{2,3*}, Ruogu Zhou², Guoliang Xing², Wen-Zhan Song⁴, Jonathan M. Lees⁵

¹College of Communication Engineering, Chongqing University, P.R. China

²Department of Computer Science and Engineering, Michigan State University, USA

³Advanced Digital Sciences Center, Illinois at Singapore

⁴Department of Computer Science, Georgia State University, USA

⁵Department of Geological Sciences, University of North Carolina at Chapel Hill, USA

ABSTRACT

Recent years have witnessed pilot deployments of inexpensive wireless sensor networks (WSNs) for active volcano monitoring. This paper studies the problem of picking arrival times of primary waves (i.e., P-phases) received by seismic sensors, one of the most critical tasks in volcano monitoring. Two fundamental challenges must be addressed. First, it is virtually impossible to download the real-time high-frequency seismic data to a central station for P-phase picking due to limited wireless network bandwidth. Second, accurate P-phase picking is inherently computation-intensive, and is thus prohibitive for many low-power sensor platforms. To address these challenges, we propose a new P-phase picking approach for hierarchical volcano monitoring WSNs where a large number of inexpensive sensors are used to collect fine-grained, real-time seismic signals while a small number of powerful coordinator nodes process collected data and pick accurate P-phases. We develop a suite of new in-network signal processing algorithms for accurate P-phase picking, including lightweight signal pre-processing at sensors, sensor selection at coordinators as well as signal compression and reconstruction algorithms. Testbed experiments and extensive simulations based on real data collected from a volcano show that our approach achieves accurate P-phase picking while only 16% of the sensor data are transmitted.

Categories and Subject Descriptors

C.3 [Special-purpose and Application-based Systems]: Signal processing systems; C.4 [Performance of Systems]: Measurement techniques; J.2 [Physical Sciences and Engineering]: Earth and atmospheric sciences

*The first two authors are listed in alphabetic order.

Part of this work was completed while Guojin Liu and Rui Tan were with Michigan State University as visiting scholar and postdoctoral Research Associate, respectively.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN'13, April 8–11, 2013, Philadelphia, Pennsylvania, USA.

Copyright 2013 ACM 978-1-4503-1959-1/13/04 ...\$15.00.

Keywords

Volcano monitoring, P-phase picking, hypocenter estimation, compressive sampling, wireless sensor network

1. INTRODUCTION

Volcanic eruptions have become a major hazard due to ever growing human population and urbanization around volcanoes. It is estimated that about 500 million people today live close to active volcanoes [1]. Existing volcano monitoring systems often employ broadband seismometers that collect high-fidelity seismic signals, but are expensive, bulky, and difficult to install. As a result, many of the most threatening volcanoes are monitored by fewer than 20 stations. Such poor spatial granularity limits scientists' ability to study the volcano dynamics and predict eruptions.

Recent years have witnessed pilot deployments of inexpensive wireless sensor networks (WSNs) for active volcano monitoring [30, 31, 27]. These deployments demonstrated the potential of long-term, large-area, and fine-grained volcano coverage by deploying large numbers of low-cost sensors. Significant research has been focused on improving system robustness, time synchronization, network efficiency, and communication performance issues. In previous small-scale deployments [30, 31, 27], detection and analysis of volcano activity were accomplished by transmitting raw data to a base station for centralized processing. However, as the sensor signals are sampled at high frequencies (e.g., 50 to 200 Hz), it is virtually impossible to continually collect raw, real-time data from a large-scale and dense WSN. This is due primarily to severe limitations of energy and bandwidth of current WSN platforms.

The goal of this paper is to design algorithms that can accurately determine the arrival times of primary waves (i.e., P-waves) received by seismic sensors *inside* the network, without transmitting raw measurements to the base station for centralized processing. Earthquake signal timing is a fundamental task in seismology. P-wave arrival times (i.e., P-phases) are essential information for advanced volcano monitoring tasks such as earthquake hypocenter estimation and seismic tomography [20]. Hypocenter estimation uses the P-phases of distributed sensors and a model of the P-wave propagation speed at different depths (a.k.a. velocity model) to estimate the earthquake source location. Seismic tomography updates the velocity model based on the sensors' P-phases and the associated hypocenters of earthquakes. The estimated dynamic velocity model is important for understanding the physical processes inside the volcano

conduit systems and issuing early warnings. In volcano observatories, P-phase picking is often done by visual inspection of experienced seismologists. When the volume and rate of data capture is large, however, this process is extremely labor-intensive, time-consuming, and subject to inconsistency across different examiners. In the last two decades, automated P-phase picking algorithms have been developed in seismology community for earthquake timing [14, 32, 26]. However, these algorithms are designed for powerful nodes with substantial computation, storage and power resources. It remains an open question if it is possible to implement automated *in-situ* volcanic earthquake timing in resource-constrained WSNs without transmitting a large volume of raw sensor data.

The key contribution of this paper is the development of new in-network signal processing algorithms for P-phase picking. To balance the system lifetime and network coverage, we adopt a hierarchical network architecture that consists of low-end nodes (referred to as *sensors*) and high-end nodes (referred to as *coordinators*). A large quantity of inexpensive, mote-class sensors can provide fine-grained monitoring with long lifetime, while a small number of coordinators (e.g., Imote2 and embedded PCs like Gumstix [2]) enable advanced in-network seismological signal processing. Based on this network architecture, we develop a suite of in-network P-phase picking algorithms. (1) Lightweight algorithms are designed for sensors to coarsely pick the P-phases and estimate the signal sparsity. The coarse P-phase is an important hint of the amount of new information that the sensor can contribute. The signal sparsity determines the volume of data transmission if the sensor sends its signal to a coordinator for accurate P-phase picking. (2) A sensor selection algorithm uses signal sparsities and coarse P-phases to choose a subset of the most informative sensors to transmit compressed data subject to a given upper bound on communication overhead. The bound can be set by the network designer to meet various practical system constraints, such as bandwidth limitation, energy budget, and real-time requirement. (3) A signal compression algorithm for sensors and a reconstruction algorithm for coordinators are developed based on wavelet transform and compressive sampling theory [12]. The above algorithms work collaboratively to achieve energy-efficient and accurate in-network earthquake timing. The approach presented in this paper can be extended and then applied in various monitoring applications that need accurate signal arrival times. Moreover, it has important implications to a broader class of applications that need to accurately extract features from real-time, high-frequency signals gathered by resource-constrained WSNs.

We implement and evaluate the proposed algorithms on a testbed of TelosB motes that are loaded with real seismic data collected on Mount St. Helens. The results demonstrate the feasibility of deploying our algorithms on volcano monitoring WSNs. We also conduct extensive simulations based on real data traces that contain 30 earthquakes. The results show that our algorithms can achieve accurate earthquake timing while only 16% of the sensor data are transmitted.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 states the problem and approach overview. Section 4 studies the sparsity of earthquake signal and presents the signal pre-processing algorithms at sensors. Section 5 formulates the sensor selection problem. Section 6 discusses the compression/reconstruction

algorithms. Section 7 discusses how to apply our approach to other applications. Section 8 presents the evaluation results. Section 9 concludes this paper.

2. RELATED WORK

The first field application of WSN for monitoring volcano was in 2004 [30], where four MICA2 nodes were deployed on Volcán Tungurahua, Ecuador. The system successfully collected three days of acoustic data. In 2005, the same research group deployed sixteen Tmote nodes equipped with seismic and acoustic sensors on Volcán Reventador, Ecuador, for three weeks [31]. In 2007, they deployed eight Tmote nodes on Volcán Tungurahua again and applied the Lance framework [29] to select a subset of sensors such that the total value of the raw data collected from the selected sensors is maximized subject to the network lifetime constraint. However, Lance adopts a heuristic metric to guide the sensor selection and does not apply signal compression before transmission. In the Optimized Autonomous Space In-situ Sensorweb (OASIS) project [27], twelve Imote2 nodes were deployed on Mount St. Helens in 2008. It demonstrated a long-term sustainable WSN in a challenging environment, and delivered a long-period (up to half a year), high-fidelity sensor dataset. The design of the above volcano monitoring WSNs [30, 31, 29, 27] mainly focused on the basic network services such as node sustainability, network connectivity, time synchronization and data collection. As raw sensor data were continually collected, these systems either had short lifetimes [30, 31] or had to employ heavy batteries [27]. We previously proposed a volcanic earthquake detection approach based on in-network signal processing [28]. The TelosB-based testbed experiments show that the in-network signal processing scheme reduces the node energy consumption to one sixth of the raw data collection approach. In contrast to [28], whose target was to detect the occurrence of earthquakes, this work builds on previous results by adding in the task of accurately picking P-phases after an earthquake is detected.

The seismology community previously developed several algorithms for P-phase picking. They are typically based on the identification of changes in signal characteristics such as energy, frequency and characteristics of autoregressive models [32]. The widely adopted STA/LTA approaches [14] continuously compute the ratio of short-term average (STA) to long-term average (LTA) over a signal characteristic and raise a detection once the ratio exceeds a specified threshold. Although STA/LTA approaches are suitable for sensors with limited resources, associated accuracies are much lower than minimal requirements of volcanic earthquake timing [32]. Moreover, the heuristic STA/LTA approaches often require empirical tuning of numerous parameters, making these methods difficult to adapt to different regions or temporally changing environments. Another important category of picking algorithms is based on autoregressive (AR) models [26]. These methods pick the time instance to maximize the dissimilarity of two AR models for signals before and after the picked time instance. AR-based algorithms need few user settings and are the most accurate and robust P-phase picking algorithms to date. However, since both AR models must be constructed for each time instance, they incur high computational complexity and memory usage.

Various in-network signal processing approaches have been proposed for different applications in data-intensive WSNs.




Figure 1: The seismic signals received by three sensors when an earthquake happens on Mount St. Helens. The vertical lines represent the P-phases.

For instance, in [18], the structural damage localization task is decentralized, by pushing the feature extraction algorithms to distributed vibration sensors. VanGo [17] can calibrate the parameters of the software filters running on low-end sensors, such that uninterested high-frequency sensor data are not transmitted. However, the simple filters included in VanGo, e.g., gating, cannot meet the stringent accuracy requirements of earthquake timing.

3. PROBLEM STATEMENT AND APPROACH OVERVIEW

3.1 Design Objectives

The P-phase is the first arrival time of a P-wave of a seismic signal. Fig. 1 shows the seismic signals received by three sensors deployed on Mount St. Helens [27] along with the manually picked P-phases. It can be seen that the sensors receive different P-phases due to different signal propagation delays. P-phase variations provide critical information for volcano monitoring applications such as earthquake hypocenter estimation and seismic tomography [20]. The task of picking the P-phases of spatially distributed sensors is referred to as *volcanic earthquake timing*. When the network is dense and P-wave velocities are high, the differences between sensors' P-phases can be small, e.g., at most one second in Fig. 1. This imposes stringent accuracy requirements on volcanic earthquake timing. In this paper, we aim to develop a holistic and energy-efficient approach to accurate volcanic earthquake timing in resource-constrained WSNs. Our approach is designed to meet the following two key objectives. First, picked P-phases must achieve satisfactory precision and maximize the accuracy of earthquake hypocenter estimation that takes P-phases as inputs. Second, to achieve expected network lifetime, the volume of seismic data transmission in each timing process must meet a specified energy budget.

3.2 System Model

Hierarchical network architecture. We adopt a hierarchical network architecture that consists of *sensors* with limited resources and *coordinators* with more processing capability and higher battery capacity. Each sensor continuously samples and buffers the signal in its memory, which is consistent with the design of previous volcano monitoring WSNs [27, 31]. A considerable number of inexpensive sensors can be deployed over the volcano to provide a high level of coverage, and work with a small number of coordinators to achieve accurate P-phase picking. The adoption of this architecture is motivated by the fact that P-phase picking

Table 1: Specification of WSN platforms [3]

Node	MCU frequency (MHz)	RAM capacity (KB)	Active power (mA)	Sleep power (μ A)
MSP430-based	8-18	2-16	1.12-3.98	0.5-1.8
ATmega-based	6-16	1-8	3.12-11.0	4.2-40
Imote2	13-416	32000	≥ 31	390
BTnode	8	180	12	3000
Preon32	8-72	64	3.7-28.3	1300
SunSPOT	180	512	24	520

^a Only MCU's power is considered. The active powers of MSP430-/ATmega-based nodes are under the condition of 8MHz and $V_{cc} = 3V$.

^b 73% WSN platforms are based on MSP430 and ATmega MCUs [3]. As MSP430 is more energy-efficient, we adopt TelosB as *sensor* in this paper.

^c As Imote2 has the highest processing capability and lowest sleep power among the high-end platforms, we adopt Imote2 as *coordinator*. Sleep power is an important parameter because the nodes sleep most of the time in the absence of earthquake (cf. Section 8.3.1).

algorithms are computation-intensive and hence cannot be executed by mote-class sensors. According to Table 1, the autoregressive Akaike Information Criterion (AR-AIC) picking algorithm [26], which needs at least 52 KB RAM, can be executed on only a few powerful WSN platforms such as Imote series, BTnode, Preon32, and SunSPOT. However, the power consumption of these nodes can be up to dozens of times higher than that of mote-class platforms based on MSP430 and ATmega processors. According to our numerical study in Section 8.3.1, the hierarchical network architecture can reduce the per-node energy consumption by 68%, compared with a network composed of only powerful nodes. The hierarchical architecture thus not only allows us to increase the coverage over a volcano, but also extends the network lifetime. Such a hierarchical architecture has also been adopted in other WSN systems [15]. In this paper, we adopt TelosB as the sensor and Imote2 as the coordinator.

Sensor clustering. The network is organized into one or multiple clusters. Each cluster consists of a number of sensors and a coordinator as the *cluster head*. Our approach can be integrated with various existing clustering algorithms [8]. In Section 8.3.2, we will discuss two clustering schemes and the setting of cluster size through simulations. The rest of this paper is focused on the design of the in-network signal processing algorithms in a single cluster. As data transmissions only happen between cluster head and associated member sensors, exhaustive data collection from the whole network can be avoided.

Synchronization and earthquake onset time. All sensors are time-synchronized by on-node GPS modules [27] or an in-network synchronization service [31]. We assume that the network can detect the occurrence and onset time of earthquake. The *earthquake onset time* is a coarsely estimated time instance, typically to second precision, at which the earthquake process starts. The STA/LTA [14, 31] or Bayesian [28] methods can be used to detect the earthquake onset time. In particular, the Bayesian earthquake detection approach [28] developed in our previous work is based on in-network signal processing and decision fusion, in which local decisions of sensors are fused and onset time is estimated at the coordinator, and sent back to each sensor. Using the




Figure 2: Illustration of in-network earthquake timing.

earthquake onset time, sensors can largely narrow the range of searching for the P-phases.

3.3 Approach Overview

We propose a suite of algorithms running at the coordinator and associated sensors, which will work together to achieve the objectives discussed in Section 3.1. The operation flow of these algorithms is illustrated in Fig. 2. (1) When an earthquake is detected, each sensor chooses a segment of seismic signal around the detected earthquake onset time and applies a wavelet transform to the signal. The wavelet transform sparsifies the signal representation, reducing the volume of data transmission. (2) Based on the transformed signal, each sensor estimates the signal sparsity and executes a lightweight picking algorithm to find a preliminary P-phase. Each sensor then sends the estimated signal sparsity and the preliminary P-phase to the coordinator. (3) The coordinator selects a subset of sensors such that the expected error of earthquake hypocenter estimation, computed from the preliminary P-phases, is minimized subject to a given upper bound on communication overhead. The communication overhead can be exactly predicted from the signal sparsities of the selected sensors. (4) The selected sensors then employ compressive sampling (CS) [12] to compress the seismic signals and transmit to the coordinator. (5) Finally, the coordinator reconstructs the seismic signals and executes high-accuracy P-phase picking algorithms and possibly other advanced seismic signal analyses. In our implementation, the coordinator adopts the AR-AIC picking algorithm [26], which is widely used in seismology, although other algorithms might be used instead.

The key novelty of this paper is the efficient integration of various algorithms into a holistic approach to achieve accurate volcanic earthquake timing in resource-constrained WSNs. Our approach has the following three advantages. First, by sensor selection, the earthquake timing process has upper-bounded communication overhead. The system designer can set this bound to meet various practical system constraints such as bandwidth limitation, energy budget, and real-time requirement. Second, our approach significantly reduces the computation and communication overhead of the sensors. By employing CS algorithms based on a binary random matrix, the signal compression at sensors only involves the computation of sums. Moreover, the coordinator can determine the volume of compressed signal prior to compression, enabling efficient sensor selection and data transmission scheduling before sensors compress signals. As a result, the unselected sensors can avoid compression computation. Third, our approach allows the coordinator to

integrate a variety of centralized seismic signal analysis algorithms on the reconstructed signals, such as Fourier and polarization analyses. The coordinator can send estimated P-phases to the base station for advanced, joint hypocenter estimation across all clusters. Moreover, it can transmit the reconstructed signals to the base station for offline analysis.

4. SEISMIC PRE-PROCESSING AT SENSORS

In this section, we first study the sparsity of seismic signals received by sensors. We then present a lightweight preliminary P-phase picking algorithm that is executed at sensors.

4.1 Sparsity of Volcanic Seismic Signal

In this paper, we adopt the common definition of sparsity in signal processing [12]. Let n denote signal length. Suppose Ψ is an orthonormal basis $\Psi = [\psi_1 \psi_2 \dots \psi_n] \in \mathbb{R}^{n \times n}$ where ψ_i is the i^{th} column of Ψ . A signal $\mathbf{s} \in \mathbb{R}^{n \times 1}$ in the time domain is expanded with basis Ψ as $\mathbf{s} = \Psi \mathbf{x}$, where $\mathbf{x} \in \mathbb{R}^{n \times 1}$ is the coefficient sequence of \mathbf{s} . The signal \mathbf{s} is k -sparse if the number of non-zeros in \mathbf{x} is less than or equal to k . The sparsity of signal \mathbf{s} , denoted by ρ , is defined as $\rho = k/n$. In practice, \mathbf{x} typically contains small values rather than zeros. Considering $\mathbf{x}_{(k)}$ obtained by keeping only the k largest coefficients of \mathbf{x} and setting others to zero, the corresponding signal $\mathbf{s}_{(k)}$ is $\mathbf{s}_{(k)} = \Psi \mathbf{x}_{(k)}$. The signal \mathbf{s} is k -sparse if the relative error $\frac{\|\mathbf{s} - \mathbf{s}_{(k)}\|_{\ell_2}}{\|\mathbf{s}\|_{\ell_2}}$ is smaller than a threshold, where $\|\cdot\|_{\ell_2}$ represents the ℓ_2 -norm. In this paper, the threshold is set to be 5% unless otherwise specified.

For each sensor, we choose a signal segment for 16 seconds, where 10 seconds before and 6 seconds after the earthquake onset time. Hence, $n = 16 \cdot f_s$, where f_s represents the seismic sampling rate. This setting of signal length is the minimum requirement of the AR-AIC picker [26] running at the coordinator. As the difference between the P-phases received by sensors is typically shorter than two seconds [31], this setting also ensures that all sensors' P-phases are included. The first columns of Figs. 3 and 4 show the chosen signals at Node01 and Node10 deployed on Mount St. Helens in the OASIS project [27], where f_s is 100 Hz. Vertical dashed lines represent the earthquake onset time detected by a Bayesian approach [28] and vertical red lines represent the P-phases picked by the AR-AIC picker [26]. It is clear that the P-phases are covered by the chosen signals.

4.1.1 Sparsity in Wavelet Domain

The time-domain seismic signal is often not sparse. For instance, for the signal shown in Fig. 4(d), the sparsity is 0.57. In this paper, we adopt discrete wavelet transform (DWT) with Daubechies basis to reduce signal sparsity, which produces reduced wireless data transmission. As DWT preserves time-domain characteristics, it is preferable for P-phase analysis. Moreover, the downsampling scheme of DWT allows us to develop an efficient preliminary P-phase picking algorithm in Section 4.2. The second columns of Figs. 3 and 4 show the 4-level DWT coefficients of Node01 and Node10 for two earthquakes. The vertical dotted lines represent edges between two adjacent frequency subbands in the wavelet domain. Setting the level of the DWT will be discussed in Section 4.2. Our analysis shows that the sparsity in the wavelet domain is significantly lower than that in the time domain. For instance, for the four data traces

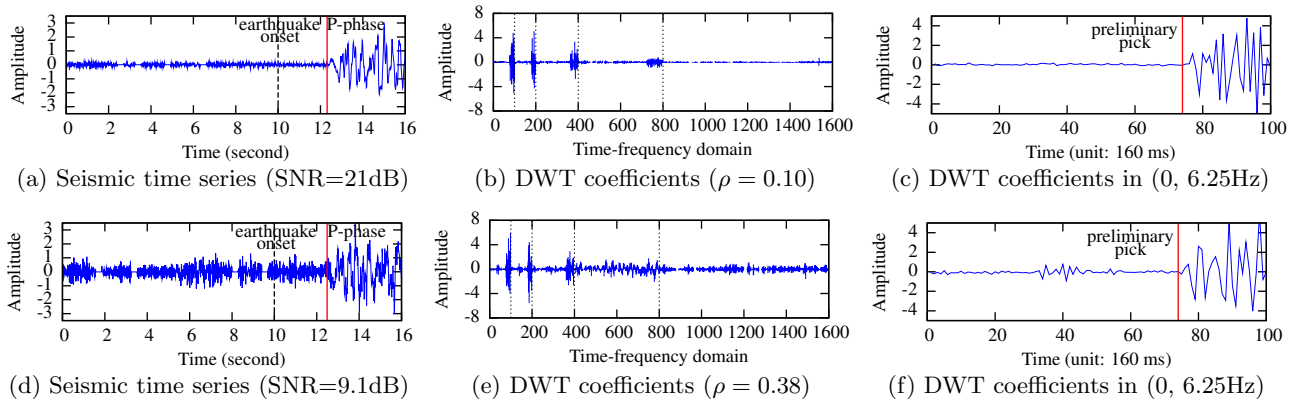


Figure 3: Earthquake01 during 12:39:23 to 12:39:39 on November 4, 2009. (a)-(c): Node01; (d)-(f): Node10.

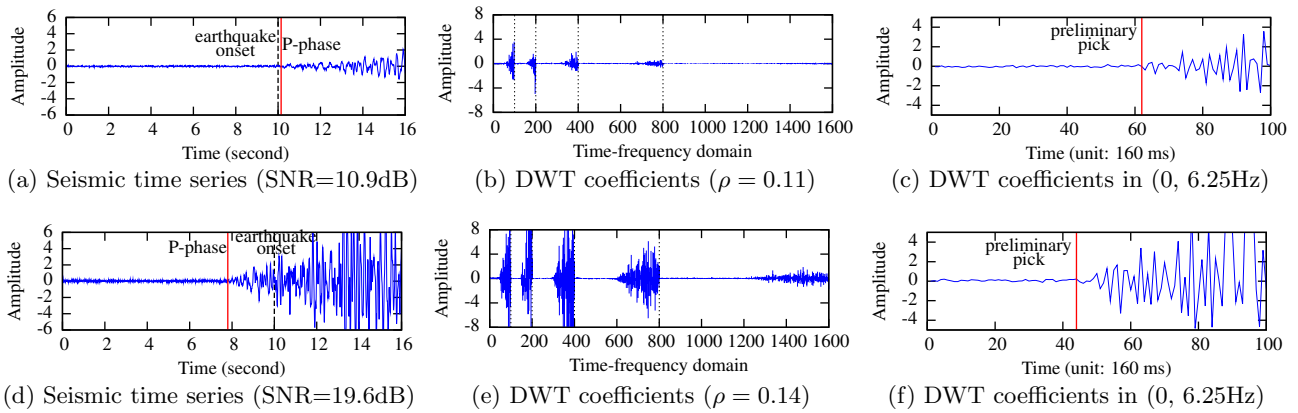


Figure 4: Earthquake02 during 00:23:58 to 00:24:14 on November 3, 2009. (a)-(c): Node01; (d)-(f): Node10.

shown in Figs. 3 and 4, the sparsity can be reduced by up to 75% using the wavelet domain.

4.1.2 Diverse Sparsity

We make the following important observations from the case study shown in Figs. 3 and 4. First, for the same earthquake, sensors receive data with different signal-to-noise ratios (SNRs), leading to different significance of P-phases. For instance, in Earthquake01 shown in Fig. 3, Node01 has a higher SNR and a more significant P-phase than Node10. As the seismic signal attenuates with propagation distance, sensors far away from the earthquake source receive weak signals, lower SNRs, and less pronounced P-phases. Second, due to highly variable event magnitude and source location, the SNR and significance of P-phase are dynamic and unpredictable. For instance, as opposed to Earthquake01, in Earthquake02 (Fig. 4), Node10 receives a much higher SNR than Node01. Third, the sparsity depends on SNR and the position of P-phase. For instance, since Node01 receives a higher SNR than Node10 in Earthquake01, the sparsity of Node01 is lower than Node10. However, although Node10 receives very high SNR in Earthquake02, its sparsity is comparable to that of Node01. This is because Node10 receives P-phase much earlier than Node01, resulting in more non-zeros in the wavelet domain. We evaluated extensively the sparsity of transformed signals based on the data traces received by 12 nodes for 30 earthquakes in the OASIS project [27]. Fig. 5 shows sparsity versus the threshold of relative error $\frac{\|s - s(k)\|_{\ell_2}}{\|s\|_{\ell_2}}$ for determining sparsity. This result vali-




Figure 5: The sparsity (with 90% confidence interval) of 360 seismic data traces received by 12 sensors.

dates our hypothesis of diverse sparsity. For instance, if the threshold is set to 5%, the sparsity ranges from 0.16 to 0.63.

The above observations of dynamic, unpredictable and diverse sparsity provide important guidelines for designing volcanic earthquake timing algorithms for resource-constrained WSNs. First, due to the diversity of signal sparsity, it is desirable to collect only the most sparse seismic signals to meet a specified node energy budget keeping with the real-time requirement of data transmission. Second, as the sparsity is dynamic and unpredictable, sensors need to compute sparsity on demand when an earthquake is detected. The sparsity can then be used to predict the volume of data transmission if the coordinator requests the signal.

4.2 Preliminary P-Phase Picking at Sensors

In this section, we present a lightweight and efficient preliminary P-phase picking algorithm that runs on the sensors. Due to the downsampling scheme, the lowest frequency sub-band in the wavelet domain is a zoomed-out version of the

low-pass filtered signal. Hereafter, this subband is referred to as the *thumbnail* of the original signal in the time domain. The last columns of Figs. 3 and 4 show the thumbnails of the corresponding original signal in the first column. The thumbnails apparently preserve the shapes of the arriving P-waves. If the seismic sampling rate is f_s and the level of DWT is l , the lowest frequency subband of the wavelet domain is $[0 \text{ Hz}, \frac{f_s}{2^{l+1}} \text{ Hz}]$. By setting l such that $\frac{f_s}{2^{l+1}} \geq 5 \text{ Hz}$, the thumbnail can preserve the shape of the P-wave, which typically has a frequency lower than 5 Hz [28]. In our approach, the preliminary P-phase is picked from the thumbnail to reduce the computational complexity. However, as the time resolution of the thumbnail reduces to $(1000 \cdot 2^l)/f_s$ milliseconds, the P-phase picking error caused by the downsampling will be $(500 \cdot 2^l)/f_s$ milliseconds. For the cases shown in Figs. 3 and 4, the number of data points that a sensor needs to process is reduced from 1600 to 100, and the error caused by downsampling is 80 milliseconds. This resolution is satisfactory for the preliminary P-phase picking.

Our lightweight preliminary P-phase picking algorithm is as follows. For a candidate P-phase \hat{p} , the sensor computes the signal energies (i.e., the sample variances) of the thumbnail signals with length of two seconds before and after \hat{p} . The preliminary P-phase, denoted by p , is given by

$$p = 2^l \times \arg \max_{\hat{p} \in \text{thumbnail}} \frac{\text{signal energy after } \hat{p}}{\text{signal energy before } \hat{p}}. \quad (1)$$

Note that the scaling factor 2^l maps the pick in the thumbnail to the original time domain. The complexity of the above algorithm is $\mathcal{O}(n/2^l)$. In contrast, existing advanced picking algorithms have significantly higher complexity, e.g., $\mathcal{O}(n^3)$ for AR-AIC picker [26]. By maximizing the signal energy ratio in Eq. (1), the preliminary P-phase divides the thumbnail signal into two segments with significantly different signal energies. In the last columns of Figs. 3 and 4, the vertical red lines represent the preliminary P-phases. We can see that the preliminary P-phase picker accurately extracts the P-phases from the thumbnails. In Section 8, we will conduct extensive evaluation of the accuracy of the preliminary picker.

5. SENSOR SELECTION FOR EARTHQUAKE TIMING

In this section, we present the sensor selection algorithm that aims to maximize the accuracy of earthquake hypocenter estimation subject to a given upper bound on communication overhead. Earthquake hypocenter estimation, which takes sensors' P-phases as inputs, is the base of many advanced volcano monitoring applications such as seismic tomography [20]. The sensor selection best directs the limited network resources, e.g., bandwidth and energy, to acquire the sensor data for accurate earthquake timing.

5.1 Impact of Timing on Hypocenter Estimation

As the propagation speed of P-wave varies with the depth in earth, the earthquake hypocenter estimation is a non-linear inversion problem involving residual reduction coupled with seismic ray tracing [20]. Suppose a set of sensors, denoted by S , belongs to the cluster under consideration. Let \mathbf{z}_i and \mathbf{z}_o denote the 3-dimensional Cartesian coordinates of sensor i and the earthquake source, p_i and p_o de-

note the P-phase picked by sensor i and the earthquake time origin of the source, and \mathbf{v} denote a list of P-wave speeds at different depths. We assume that $\{\mathbf{z}_i | i \in S\}$ and \mathbf{v} are known, which can be obtained by inquiring the GPS module on sensors [27] and from existing tomographic studies [20], respectively. The \mathbf{z}_o and p_o are the unknowns to be estimated from the P-phases $\{p_i | i \in S\}$. We have

$$p_i - p_o = \tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v}) + \epsilon_i, \quad \forall i \in S \quad (2)$$

where $\tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v})$ is the P-wave travel time from the source to sensor i given the velocity model \mathbf{v} , and ϵ_i is the random error experienced by sensor i . We employ the ray tracing algorithm in the RSEIS R package [21] to calculate $\tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v})$. We assume that ϵ_i follows zero-mean normal distribution with variance ζ^2 . The variance ζ^2 captures the error of the P-phase picked from the seismic signal with respect to the true P-phase. As will be shown later, the hypocenter estimation algorithm and its accuracy analysis are independent of ζ^2 . Hence, the variance ζ^2 can be unknown to the network. The unknown p_o can be canceled out by subtracting Eq. (2) with $i = r$ from the same equation with $i \in S \setminus \{r\}$, yielding $p'_i = \tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v}) - \tau(\mathbf{z}_r, \mathbf{z}_o | \mathbf{v}) + \epsilon'_i$, where sensor r is the *reference node*, $i \in S \setminus \{r\}$, $p'_i = p_i - p_r$, $\epsilon'_i = \epsilon_i - \epsilon_r$. Note that ϵ'_i follows zero-mean normal distribution with variance $2\zeta^2$. We adopt maximum-likelihood (ML) approach to estimate \mathbf{z}_o . The ML estimate of \mathbf{z}_o , denoted by $\tilde{\mathbf{z}}_o$, is given by:

$$\tilde{\mathbf{z}}_o = \arg \min_{\mathbf{z}_o} \sum_{i \in S \setminus \{r\}} (p'_i - \tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v}) + \tau(\mathbf{z}_r, \mathbf{z}_o | \mathbf{v}))^2. \quad (3)$$

We now analyze the accuracy of $\tilde{\mathbf{z}}_o$. As there is no closed-form formula for $\tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v})$, to make the analysis tractable, we let $\tau(\mathbf{z}_i, \mathbf{z}_o | \mathbf{v}) = \|\mathbf{z}_i - \mathbf{z}_o\|_{\ell_2}/v$, where v represents the average P-wave speed. Define $\mathbf{g}_i = \frac{\mathbf{z}_r - \mathbf{z}_o}{\|\mathbf{z}_r - \mathbf{z}_o\|_{\ell_2}} - \frac{\mathbf{z}_i - \mathbf{z}_o}{\|\mathbf{z}_i - \mathbf{z}_o\|_{\ell_2}}$ and let \mathbf{G} denote the matrix composed of $\{\mathbf{g}_i | \forall i \in S \setminus \{r\}\}$ as columns. By extending the result in [13], the Fisher information matrix, denoted by \mathbf{J} , is given by $\mathbf{J} = \frac{1}{2v^2\zeta^2} \mathbf{G}\mathbf{G}^T \in \mathbb{R}^{3 \times 3}$, where the diagonal elements of \mathbf{J}^{-1} are the theoretical lower bounds for the variances of the coordinates in $\tilde{\mathbf{z}}_o$. A widely adopted error metric is $\text{tr}(\mathbf{J}^{-1}) = 2v^2\zeta^2 \text{tr}((\mathbf{G}\mathbf{G}^T)^{-1})$. As $2v^2\zeta^2$ is a scaling factor in $\text{tr}(\mathbf{J}^{-1})$, we define the error metric as

$$\mathcal{E} = \text{tr}((\mathbf{G}\mathbf{G}^T)^{-1}). \quad (4)$$

Note that \mathcal{E} depends on the true but unknown source location \mathbf{z}_o . In our approach, we replace \mathbf{z}_o in Eq. (4) with its ML estimate $\tilde{\mathbf{z}}_o$ to calculate the error metric.

The theoretical error metric given by Eq. (4) is the same for different P-phase pickers that yield zero-mean errors with respect to the true P-phase. As will be shown in Section 8, the preliminary P-phase picker has zero-mean error with respect to the AR-AIC picker that has near zero-mean error (100 ms with respect to the manual picks [26]). Hence, the error metric calculated from the preliminary P-phases is a good estimate of the error metric calculated from the P-phases picked by AR-AIC at the coordinator.

5.2 Dynamic Sensor Selection Problem

Our study in Section 4 shows that sensors have diverse signal sparsity. As a result, the volume of data transmission varies significantly across different sensors. The coordinator requests the compressed signals from a subset of sensors

to minimize hypocenter estimation error subject to a given upper bound on communication cost. We make the following assumptions. First, the volume of compressed signal is given by $m(\rho_i)$, where ρ_i is the sparsity of sensor i . The expression of $m(\rho_i)$ will be given in Section 6. Second, the communication cost of a data unit from sensor i to the coordinator is c_i , which is referred to as *unit communication cost*. The sensor selection problem is formulated as follows:

Sensor Selection Problem. *When an earthquake is detected, given the sparsity of all sensors $\{\rho_i|\forall i\}$ and the unit communication costs $\{c_i|\forall i\}$, find a subset of sensors S such that the error metric \mathcal{E} given by Eq. (4) is minimized, subject to $\sum_{i \in S} c_i \cdot m(\rho_i) \leq C$.*

In the above problem, C is the upper bound on the total communication cost in each earthquake timing process. By properly setting the unit communication costs, the upper bound C can represent different costs, e.g., the number of transmitted packets, the energy consumed in an earthquake timing process, or the latency of the data collection. Moreover, c_i can incorporate the residual battery energy such that the solution can balance sensors' energy consumption for multiple rounds of earthquake timing. For instance, by defining c_i as the reciprocal of sensor i 's residual energy, the most informative sensors with more residual energies and less transmission volume will be selected.

In our approach, the coordinator first solves Eq. (3) using the Nelder-Mead algorithm [23]. For any candidate sensor subset, we consistently use $\tilde{\mathbf{z}}_o$ to compute \mathcal{E} . As \mathcal{E} is a non-linear and non-convex function, it is difficult to solve the sensor selection problem in polynomial complexity. In our implementation (cf. Section 8.1.1), the execution time of the Nelder-Mead algorithm on Imote2 is around 4 seconds. A brute-force search takes 0.08 and 8.2 seconds when the cluster size is 10 and 16, respectively. Note that our numerical study in Section 8.3.2 shows that the gain of hypocenter estimation performance rapidly diminishes after the cluster size is greater than 15. Therefore, the computation overhead of the brute-force search is acceptable without sacrificing too much hypocenter estimation accuracy due to the setting of cluster size. In Section 5.3, we propose an approximate sensor selection algorithm that can scale with the cluster size but will sacrifice hypocenter estimation accuracy.

If the coordinator is equipped with a seismometer to sample the seismic signal, it can be always selected to improve the hypocenter estimation accuracy. Moreover, the P-phase picked from the coordinator's signal can be used as a reference to identify wrong preliminary P-phases sent from the sensors as well as wrong P-phases picked from the reconstructed signals at the coordinator.

5.3 Approximate Sensor Selection Algorithm

In this section, we propose a new heuristic metric that allows us to develop an efficient sensor selection algorithm. The metric is defined as

$$\mathcal{V} = \sum_{i \in S} \frac{1}{(p_i - \tilde{p}_o - \tau(\mathbf{z}_i, \tilde{\mathbf{z}}_o|\mathbf{v}))^2}, \quad (5)$$

where S is the subset of selected sensors and \tilde{p}_o is the ML estimate of p_o . Specifically, $\tilde{p}_o = \frac{\sum_{i \in S} p_i - \tau(\mathbf{z}_i, \tilde{\mathbf{z}}_o|\mathbf{v})}{N}$, where N is the number of sensors in the cluster. The denominator in Eq. (5) is the squared error in P-phase. The sensor selection problem is to select a subset of sensors S to maximize \mathcal{V} sub-

ject to the constraint $\sum_{i \in S} c_i \cdot m(\rho_i) \leq C$. This problem is a 0-1 knapsack problem, which can be solved optimally in pseudo-polynomial time. Eq. (5) is a specialization of the heuristic metric adopted in Lance [29] that defines the total value of selected sensors as the sum of the values of individual sensors. A key difference is that Lance does not consider signal compression. The evaluation results in Section 8.2.2 show that the solution given by this approximate algorithm approaches to the optimal solution described in Section 5.2 when the constraint C becomes larger.

6. COMPRESSIVE SAMPLING FOR EARTHQUAKE TIMING

This section presents our approach of compressing and collecting the seismic signals from the selected sensors based on compressive sampling (CS) [12]. We first briefly review the CS theory. Let $\mathbf{y} \in \mathbb{R}^{n \times 1}$ denote the compressed signal and $\mathbf{A} \in \mathbb{R}^{m \times n}$ denote the random projection matrix, where $m < n$. The compression is expressed as $\mathbf{y} = \mathbf{A}\mathbf{x}$, where \mathbf{x} is a vector of wavelet coefficients of the original signal. Note that the typical use of CS is to apply the combined transform and random projection (i.e., $\mathbf{A}\Psi^{-1}$) to the time-domain signal \mathbf{s} . However, in our approach, these two steps are separated to efficiently estimate the preliminary P-phase in the lowest subband of \mathbf{x} and the sparsity. These two numbers are important inputs to the sensor selection algorithms. If \mathbf{x} is k -sparse and \mathbf{A} complies with the *restricted isometry property* (RIP) of order k , the original signal \mathbf{s} can be exactly reconstructed from \mathbf{y} [12]. The wavelet transform of the reconstructed signal, denoted by $\tilde{\mathbf{x}}$, is given by $\tilde{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_{\ell_1}$ subject to $\mathbf{y} = \mathbf{A}\mathbf{x}$. The above optimization can be solved by various algorithms such as the iterative hard thresholding method [11]. With $\tilde{\mathbf{x}}$, the reconstructed seismic signal is given by $\Psi\tilde{\mathbf{x}}$.

We now discuss the design of CS for earthquake timing. We adopt the binary random projection matrix [9] that is promising for the implementation on resource-constrained sensors. Specifically, only the positions of '1's need to be stored and the multiplication $\mathbf{A}\mathbf{x}$ is simply the sum of the elements of \mathbf{x} at these positions. The binary random matrix complies with RIP of order k if $m \geq h \cdot k \cdot \log(n/k)$, where h is an unknown constant [9]. From the results shown in Fig. 5, the sparsity ρ of volcanic seismic signal typically ranges from 0.1 to 0.6. Hence, $\log(n/k) = \log(1/\rho)$ ranges from $\log(1.67)$ to $\log(10)$. We define $\eta = \log(10) \cdot h$. If $m \geq \eta \cdot \rho \cdot n$, the RIP condition must be satisfied. Therefore, we let

$$m(\rho) = \eta \cdot \rho \cdot n. \quad (6)$$

Many studies have reported that $\eta = 4$ is a safe setting that ensures satisfactory reconstruction [12]. However, as the sparsity ρ estimated in Section 4.1 does not follow the strict definition of sparsity (i.e., the ratio of non-zeros), the setting of $\eta = 4$ might be overly conservative for earthquake timing, which may result in excessive data transmission. In Section 8, we evaluate in detail the impact of η on the quality of seismic signal reconstruction as well as the P-phase picking. The results show that the setting of $\eta = 1.5$ can lead to a good trade-off between the volume of data transmission and the P-phase picking error introduced by reconstruction. In practice, the setting of η can also be determined based on the seismic data obtained in offline earthquake shaking table experiments. We note that the CS-compressed signal can be

further compressed by other data compression algorithms if more computation resource is available.

7. DISCUSSION

The approach presented in this paper can be applied to a broader class of sensor network applications where sensors sample the physical phenomena at high frequencies and extract signal features from the samples. Many signal feature extraction algorithms are not affordable for resource-constrained sensors because of either the large volume of data or high complexity of the algorithms. Therefore, it is desirable to select a subset of most contributory sensors to transmit their compressed data to a more powerful node for feature extraction.

To apply our approach, a lightweight algorithm should be available to compute a coarse estimate of the feature, and a closed-form expression or heuristic metric is then used to predict the quality of upper-layer application based on the coarsely estimated features. In particular, our approach can be applied to structural damage localization [18] and most applications based on time difference of arrival (TDOA). As the P-phase picking addressed in this paper is a critical component in a class of TDOA-based applications such as acoustic event localization, our approach can be easily applied to these applications. We now briefly discuss how to extend our approach to the structural damage localization based on the natural frequencies received by distributed vibration sensors [18]. The natural frequency identification algorithm involves high-order curve fitting, and hence can be computationally prohibitive for low-end sensors due to the lack of floating point arithmetic support. To apply our approach, low-end sensors can use simple peak detectors [10] to coarsely estimate the natural frequencies, and the coordinator can use the Damage Localization Assurance Criterion [18] to guide the sensor selection.

8. PERFORMANCE EVALUATION

In this section, we conduct testbed experiments and extensive simulations based on real data traces collected by 12 nodes on Mount St. Helens in the OASIS project [27]. Our system implementation and testbed experiments verify the feasibility of the proposed signal processing algorithms on low-end sensor platforms. The trace-driven simulations extensively evaluate the performance of our approach. We finally conduct two numerical studies to evaluate the energy efficiency of the hierarchical network architecture and the impact of sensor clustering on earthquake timing.

8.1 Testbed Experiments

8.1.1 System Implementation

Sensors: Our system implementation is based on TelosB motes. Similar mote-class sensor platforms were also used in previous volcano monitoring systems [30, 31, 29, 28]. We implement all the four seismic processing algorithms, i.e., DWT, sparsity estimation, preliminary P-phase picker and CS in TinyOS 2.1. We conducted extensive code optimization on all the signal processing algorithms. First, we adopt fixed point arithmetic, which can speed up the decimal computation up to 10 fold on TelosB with respect to default floating point arithmetic. Second, we maintain a single input/output data buffer for the four pipelined algorithms and wire the output of each algorithm back to the




Figure 6: Sensors' workloads in an earthquake.




Figure 7: Distribution of preliminary picking error.

buffer. This pipeline implementation significantly reduces RAM usage. Our current implementation of CS uses predefined binary random matrices for sensors, which avoids the overhead of online matrix generation each timing process. In future work, we will explore efficient methods to generate the same binary random matrix on sensor and coordinator without incurring high communication costs. A possible solution is to use a common seed to generate the same projection matrix on both the sensor and coordinator. To improve the realism of the experiments, we reserve 320 KB on the mote's flash and load it with real seismic data trace collected by the OASIS system [27]. A mote acquires 100 seismic samples from flash every second, which is consistent with the sampling rate in OASIS. Our implementation uses 21 KB ROM and 8 KB RAM.

Coordinator: We use a laptop computer to simulate the coordinator and implement all its algorithms in ANSI C. The ANSI C implementation can be easily ported to embedded computing platforms such as Imote2. To evaluate the computational overhead of these algorithms, we cross-compile the programs and run them in the SimIt-ARM 3.0 [5], which simulates the XScale processor on Imote2.

8.1.2 Experiment Results

We evaluate the computation and storage overhead of the algorithms running on sensors in a testbed of 12 TelosB motes, loaded with the real data traces sampled by 12 nodes of the OASIS system [27] in an earthquake. Fig. 6 shows the execution times of various signal processing algorithms at different sensors during an earthquake event. It is clear that the end-to-end execution time does not exceed 3 seconds, which introduces moderate workload and energy consumption to the sensors. Moreover, our implementation of CS is very efficient and most computation overhead is due to the DWT. The variation of execution time is mainly caused by sparsity estimation and CS. In the sparsity estimation algorithm, the wavelet coefficients are sorted using quick sort, which has a variable execution time. As seismic signals at sensors have different sparsities, sensors have different numbers of rows in the project matrix \mathbf{A} , leading to a variable execution time of CS. Nonetheless, as the variation is less than one second, the computation overhead is relatively evenly distributed among the sensors.

8.2 Trace-Driven Simulations

Our simulations use a data set collected by 12 Imote2-based nodes on Mount St. Helens in the OASIS project [27], which spans 5.5 months and comprises 30 significant earthquakes. In this section, we simulate a cluster of 12




Figure 8: P-phase picking error and relative reconstruction error vs. the coefficient η .




Figure 9: The error metric and the number of selected sensors vs. the allowed number of packets.

sensors, which exactly correspond to the 12 nodes in the OASIS project [27]. Note that the cluster size of 12 is a reasonable setting that will be evaluated in Section 8.3.2. For each earthquake, we use a Bayesian approach [28] to detect the onset time based on 10 minutes of data traces. In our simulations, we also use the locations of OASIS nodes and a velocity model \mathbf{v} (cf. Section 5.1) obtained in a tomographic study of Mount St. Helens [20].

8.2.1 Accuracy of P-phase Picking

We first evaluate the accuracy of the preliminary P-phase picker described in Section 4.2. The error is defined as the absolute difference between the preliminary P-phase and the P-phase picked by the AR-AIC picker [26] on the original seismic signal. Fig. 7 shows the distribution of preliminary picking error based on 100 sensor data traces. The mean error is 8.5 milliseconds. Therefore, the preliminary P-phase picker can be approximated as a zero-mean error picker with respect to the AR-AIC picker. The standard deviation is about 400 milliseconds. To evaluate the effectiveness of the preliminary picker, we also calculate the error of the earthquake onset time with respect to the P-phase picked by the AR-AIC picker. The mean and standard deviation of the error of earthquake onset time are 280 and 1310 milliseconds, respectively. Therefore, compared with the earthquake onset time, the results of the preliminary picker are more concentrated on the true P-phases.

The coefficient η in Eq. (6) is an important coefficient for CS. We now evaluate the impact of η on the quality of seismic signal reconstruction as well as the P-phase picking at the coordinator. The relative reconstruction error is calculated as $\|\tilde{\mathbf{s}} - \mathbf{s}\|_{\ell_2} / \|\mathbf{s}\|_{\ell_2}$ where \mathbf{s} and $\tilde{\mathbf{s}}$ are the original and reconstructed signals. The picking error is calculated as the difference between the P-phases picked by the AR-AIC picker on \mathbf{s} and $\tilde{\mathbf{s}}$. Fig. 8 shows the standard deviation of picking error as well as the relative reconstruction error versus the coefficient η based on 100 data traces. We can see that the standard deviation of picking error dramatically drops when η increases from 0.75 to 1.5 and becomes flat after 1.5. Therefore, $\eta = 1.5$ is a proper setting to achieve the satisfactory gain of P-phase picking accuracy to the data transmission volume. When $\eta \geq 1.5$, the mean picking error is within $[-15 \text{ ms}, 15 \text{ ms}]$. Therefore, it is shorter than 1.5 sampling periods given that the sampling rate is 100 Hz. This time error can be translated to an error distance of 75 to 120 meters based on the P-wave speed (5 to 8 km/s). As the precision for distance in volcano models (e.g., \mathbf{v} in

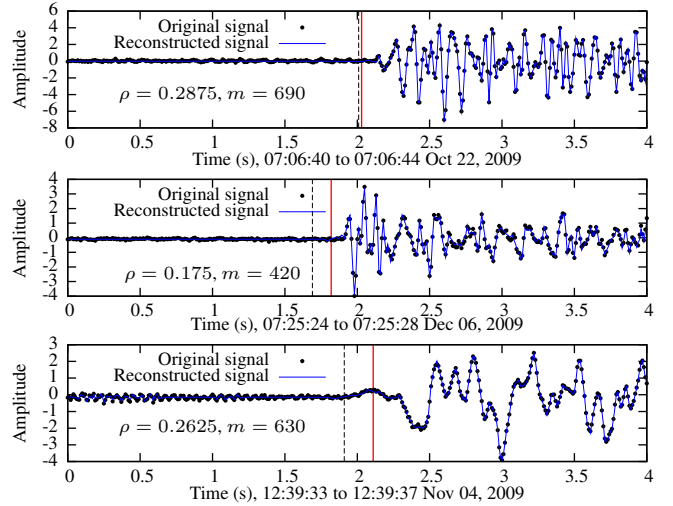


Figure 10: Original and reconstructed signals with picks. Vertical dashed/solid lines represent the picks by AR-AIC algorithm on the original/reconstructed signals.

Section 5.1) is typically in the order of kilometers [20], the error of 15 ms is small and can be safely approximated as zero-mean error. With $\eta = 1.5$, Fig. 10 shows the original and reconstructed signals received by a sensor for three earthquakes. It is apparent that the signals are accurately reconstructed and P-phases are well preserved.

8.2.2 Effectiveness of Sensor Selection

In our simulations, each sensor directly communicates with the coordinator. Each packet carries total ten 4-byte data points. Therefore, by setting the unit communication cost $c_i = 1/10$, the upper bound of communication cost C characterizes the number of packet transmissions that are allowed in an earthquake timing process. Fig. 9 plots the error metric \mathcal{E} (given by Eq. (4)) of the optimal (cf. Section 5.2) and approximate (cf. Section 5.3) sensor selections versus C in an earthquake. The figure also illustrates the number of selected sensors in the optimal solution. Note that the number of selected sensors in the approximate solution is at most one more than that of the optimal solution. From the figure, we can see that if more packet transmissions are allowed, the coordinator will select more sensors to collect data from them. Consistent with intuition, the error metric decreases with the number of packet transmissions. When $C \in [170, 250]$, a total of four sensors are selected. However, the selected four sensors can be different. When more packets are allowed, sensors with higher ρ 's, though more contributory to hypocenter estimation, will be selected. We can see that the error metric for the optimal solution becomes flat when more packets are allowed. This result can be exploited to reduce the communication cost without sacrificing too much hypocenter estimation accuracy. When C is lower than 300, the approximate solution has much worse performance than the optimal solution. However, the approximate solution approaches the optimal solution when C is greater than 500.

8.2.3 Impact of Random Packet Loss

As volcano monitoring WSNs are deployed in harsh environments, sensors are subject to unreliable communication

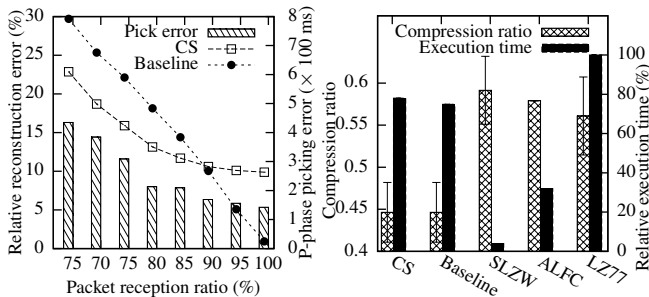


Figure 11: Impact of packet loss on the reconstruction and P-phase picking.

links [27]. We evaluate the impact of random packet loss on our earthquake timing approach. In the simulations, we assume that each link from sensor to coordinator has the same packet reception ratio (PRR). The coordinator can detect lost packets from the sequence numbers in the received packets. When the coordinator reconstructs the signal, it only uses the rows in the projection matrix \mathbf{A} that correspond to the received data points. Therefore, the effect of packet loss is similar to choosing a smaller m in CS. We compare our CS-based approach with a baseline that implements a lossy compression scheme. The baseline transmits the largest coefficients together with their indexes in the wavelet domain. The number of transmitted coefficients is chosen to make sure that the baseline produces the same number of packets as our approach. The curves in Fig. 11 plot the relative reconstruction errors of our CS-based approach and the baseline versus PRR. When no packet loss happens, the baseline outperforms our approach. This result is consistent with previous studies on CS [16]. When the PRR is lower than 90%, our approach outperforms the baseline. It has been observed in previous deployments [30, 27] that the PRR varies with time due to changing environment and can be lower than 90%. Therefore, we can switch between the baseline and CS according to recent PRRs (e.g., measured in the per-second earthquake onset time detection [28]). The histograms in Fig. 11 plot the average P-phase picking error of our CS-based approach. We can see that the reconstruction is resilient to packet loss when the PRR is no lower than 80%. Error correction mechanisms such as Forward Error Correction can be integrated with the baseline to improve its resilience to packet loss. However, they can increase both computation and communication overhead. A comprehensive comparison that accounts for error correction mechanisms is left for our future work.

8.2.4 Compression Efficiency

We now compare our CS-based approach with several baselines in terms of compression ratio and execution time. In addition to the lossy baseline approach used in Section 8.2.3, we adopt the following three lossless baseline algorithms: (1) SLZW [25], a lossless compression algorithm designed for WSNs; (2) ALFC [19], a real-time predictive lossless compression algorithm developed for OASIS [27]; (3) Lempel-Ziv coding (LZ77), a widely employed scheme used in traditional data-collection-based volcano monitoring systems. Compression ratio is defined as the ratio of compressed size

Figure 12: Compression ratios and relative execution times of various schemes in 30 earthquakes.

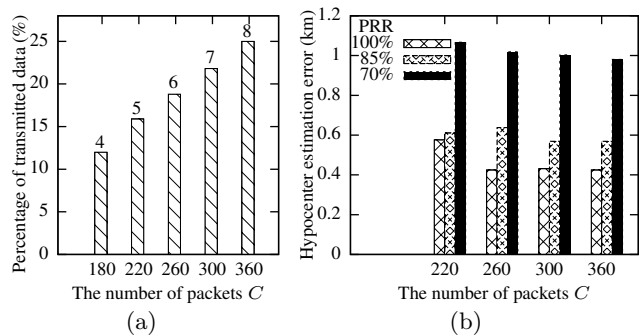


Figure 13: Hypocenter estimation results for an earthquake at 16:56:47 Nov 03 2009: (a) The percentage of transmitted data versus the allowed number of packets C ; (b) Hypocenter estimation error versus C under various PRRs.

to the original size. Fig. 12 plots the compression ratios and relative execution times of various approaches. The relative execution time is calculated with respect to LZ77. Our approach and the lossy baseline approach have comparable compression efficiency. Our approach saves more than 10% data transmission volume compared with the lossless baselines. It is faster than LZ77 but slower than SLZW and ALFC. However, none of the lossless baseline algorithms can predict the exact volume of compressed signal prior to compression. Therefore, they do not allow effective scheduling of data transmission.

8.2.5 Earthquake Hypocenter Estimation

As the data transmission scheduling is guided by the accuracy of hypocenter estimation, the final set of simulations evaluate the impact of our timing approach on the hypocenter estimation. We first use the P-phases picked by AR-AIC on the original signals of all 12 sensors to localize the earthquake source. This source location is regarded here to be the groundtruth location. We then localize the earthquake source based on the P-phases obtained in our timing approach. The hypocenter estimation error is the Euclidean distance from the groundtruth location. Fig. 13(a) plots the percentage of transmitted data in our timing approach with respect to the total volume of raw data at sensors. The number over each bar is the number of selected sensors. We observe that when C ranges from 180 to 360, only 12% to 25% of the sensor data are transmitted. Fig. 13(b) plots the corresponding hypocenter estimation errors under different PRRs. Note that when $C = 180$, the hypocenter estimation error is around 9 km (not shown in Fig. 13(b)). Consistent with intuition, the hypocenter estimation error decreases with C and PRR. From the two figures, by setting $C = 220$, the hypocenter estimation error is below 1 km, a common result in volcano seismology [20], at the expense of only 16% data transmission.

8.3 Impact of Network Architecture and Clustering

In order to choose the right hardware platform and network organization, a sensor network designer must carefully consider the trade-offs between many factors, including hardware availability, energy consumption, sensing coverage,

system delay, and etc. In this work, we adopt a hierarchical network architecture where mote-class sensors sample fine-grained signals and powerful coordinators run computation-intensive seismological algorithms. However, one may argue that a network composed of only powerful nodes such as Imote2 is more desirable in data-intensive applications like volcano monitoring. In contrast to mote-class nodes, these nodes can collect and directly process the seismic signals for various advanced monitoring tasks, reducing the energy cost of communication with the cluster heads. In this section, we quantitatively study the impact of network architecture and sensor clustering on the energy efficiency and performance of earthquake timing under realistic settings.

8.3.1 Energy Efficiency under Different Network Architectures

This numerical study compares the per-node energy consumption under the *hierarchical* and *non-hierarchical* network architectures discussed in Section 3.1. In the non-hierarchical network, a cluster is composed of only high-end sensors and each sensor runs the AR-AIC algorithm. In the hierarchical network, we consider our CS-based approach proposed and a centralized approach. To simplify the analysis, we assume a 1-hop star topology centered at the coordinator. For our approach, we assume all sensors are always selected. For the centralized approach, each sensor transmits a segment of raw seismic signal (cf. Section 4.1) to the coordinator. Since earthquakes are usually rare events, the network must perform earthquake detection most of the time in order to capture these events. We should thus also model the energy consumed in earthquake detection. Assume each sensor detects earthquake every second using some detection algorithm. The sensors send detection results to the cluster head, which then fuses the results to make the final detection decision, subsequently sends the earthquake onset time back to the sensors in case of positive decision. This is a common detection approach adopted in previous volcano monitoring systems [31, 28]. Due to space limitation, the details of the energy consumption modeling are omitted here and can be found in [22]. We compare the energies consumed in computation and communication by a sensor per day under the two network architectures. Fig. 14 shows the map and contours of ratio of energies under the hierarchical (CS-based approach) and non-hierarchical networks. Note that the execution time of the detection algorithm on TelosB (i.e., X -axis of Fig. 14) varies, depending on the detection algorithm. For instance, the STA/LTA-based and Bayesian detection algorithms require around 10ms and 100ms on TelosB [28], respectively. The execution time on Imote2 is scaled accordingly. From Fig. 14, we observe that the hierarchical network consumes much less energy in computation and communication than the non-hierarchical network under a range of settings. This is true because, primarily, when the sensors sleep most of the time in the absence of earthquake, the sleep power of Imote2 is at least 18 times of that of TelosB [6, 4].¹ After the current draw of the sensor circuit is taken into consideration [27], the sensor’s projected lifetimes over two D-cell batteries are about 6 and 2 months, re-

¹We assume that the sensors can sleep with the help of Direct Memory Access controller in signal sampling [24]. The results become more favorable to the hierarchical architecture if the sensors stay in idle state instead of sleep, because the energy ratio increases from 18 to at least 20 [6, 4].




Figure 14: Ratio of energy consumed by a sensor in the hierarchical (CS-based) and non-hierarchical networks.




Figure 15: Hypocenter estimation error metric ϵ versus cluster size under two clustering schemes.

spectively, under the hierarchical (CS-based approach) and non-hierarchical architectures, if a STA/LTA-based earthquake detection algorithm is adopted and 100 positive detection decisions are made by the detection algorithm per day. Moreover, our CS-based approach can increase lifetime by 7% and 12% compared with the centralized approach if 100 and 200 positive detection decisions are made per day, respectively. Note that the network will make more than 200 positive detection decisions per day if its false alarm rate is no lower than 3%, which is common for earthquake detection algorithms [28].

8.3.2 Impact of Sensor Clustering

This numerical study evaluates the impact of sensor clustering on earthquake timing. A hundred sensors are randomly deployed over a $6 \times 6 \text{ km}^2$ square region. We assume that the earthquake occurs at 10 km beneath the center of the region. We consider the following two clustering schemes. (1) *Geographic clustering*: A cluster head is randomly selected from the network and the sensors that are geographically closest to it are its members. This approach is similar to a class of sensor clustering algorithms based on sensor locations [8]. (2) *Random clustering*: Sensors are randomly selected from the network to form a cluster. Although this scheme is not practical, it gives the upper bound on the hypocenter estimation accuracy because sensors are most scattered. Fig. 15 shows the hypocenter estimation error metric (given by Eq. (4)) of a cluster averaged over many runs versus the cluster size under the two schemes. We observe that, for both schemes, the hypocenter estimation error has a sharp drop initially and then becomes flat when the cluster size increases. This result implies that adding a sensor becomes less beneficial for a larger cluster. From the figure, a setting of around 15 for cluster size is preferable. Although this numerical study is based on simplified assumptions, it provides insights into the impact of sensor clustering on earthquake hypocenter estimation. In practice, similar numerical studies, which integrate available geographical information such as the volcano surface altitude data, can be conducted to guide the sensor clustering as well as the setting of cluster size.

9. CONCLUSION AND FUTURE WORK

This paper presents a holistic and energy-efficient approach to accurate volcanic earthquake timing in WSNs. We develop a suite of in-network seismic signal processing algo-

rithms that collaboratively pick the arrival times of seismic primary waves received by sensors. A dynamic sensor selection problem is formulated to maximize the performance of earthquake hypocenter estimation subject to a given upper bound on communication overhead. We further develop the signal compression and reconstruction algorithms based on compressive sampling theory. Testbed experiments and extensive simulations based on real data traces collected on an active volcano demonstrate the effectiveness of our approach.

In this paper, we use an XScale processor simulator on a laptop to simulate the coordinator (cf. Section 8.1.1). In our future work, we plan to use Imote2 to extensively evaluate the computation and communication overhead of coordinator's algorithms, which allows us to accurately study the trade-off between energy consumption of coordinator and lifetime extension of sensors. The results can guide the choices of batteries for both coordinator and sensors. Moreover, we plan to deploy and further evaluate our approach in a real volcano monitoring WSN system [7].

10. ACKNOWLEDGMENTS

The authors thank our shepherd Dr. John Stankovic and the anonymous reviewers for providing valuable feedbacks to this work. This work was supported in part by U.S. National Science Foundation under grants OIA-1125163, CNS-0954039 (CAREER), CNS-1218475, OIA-1125165, CNS-1066391, in part by National Natural Science Foundation of China under grant 61171089, and in part by the Fundamental Research Funds for the Central Universities under grant CD-ZJR10160005.

11. REFERENCES

- [1] AlertNet. <http://bit.ly/f9JhLc>.
- [2] Gumstix. <http://www.gumstix.com>.
- [3] List of wireless sensor nodes. <http://bit.ly/TfLEom>.
- [4] Power modes and energy consumption for the imote2 sensor node. <http://bit.ly/THlmRz>.
- [5] SimIt-ARM. <http://bit.ly/T44mj1>.
- [6] TelosB datasheet. <http://bit.ly/Psjj2S>.
- [7] VolcanoSRI project. <http://sensornet.cse.msu.edu>.
- [8] A. Abbasi and M. Younis. A survey on clustering algorithms for wireless sensor networks. *Computer communications*, 30(14):2826–2841, 2007.
- [9] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *Annu. Allerton Conf. Commun., Control, and Comput.*, 2008.
- [10] F. Blais and M. Rioux. Real-time numerical peak detector. *Signal Processing*, 11(2):145–155, 1986.
- [11] T. Blumensath and M. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3), 2009.
- [12] E. Candès and M. Wakin. An introduction to compressive sampling. *IEEE Signal Process. Mag.*, 25(2), 2008.
- [13] Y. Chan and K. Ho. A simple and efficient estimator for hyperbolic location. *IEEE Trans. Signal Process.*, 42(8), 1994.
- [14] E. Endo and T. Murray. Real-time seismic amplitude measurement (RSAM): a volcano monitoring and prediction tool. *Bulletin of Volcanology*, 53(7), 1991.
- [15] O. Gnawali, B. Greenstein, K.-Y. Jang, A. Joki, J. Paek, M. Vieira, D. Estrin, R. Govindan, and E. Kohler. The TENET Architecture for Tiered Sensor Networks. In *SenSys*, 2006.
- [16] V. Goyal, A. Fletcher, and S. Rangan. Compressive sampling and lossy compression. *IEEE Signal Process. Mag.*, 25(2), 2008.
- [17] B. Greenstein, C. Mar, A. Pesterev, S. Farshchi, E. Kohler, J. Judy, and D. Estrin. Capturing high-frequency phenomena using a bandwidth-limited sensor network. In *SenSys*, 2006.
- [18] G. Hackmann, F. Sun, N. Castaneda, C. Lu, and S. Dyke. A holistic approach to decentralized structural damage localization using wireless sensor networks. In *RTSS*, 2008.
- [19] A. Kiely, M. Xu, W. Song, R. Huang, and B. Shirazi. Adaptive linear filtering compression on realtime sensor networks. In *PerCom*, 2009.
- [20] J. Lees and R. Crosson. Tomographic inversion for three-dimensional velocity structure at mount st. helens using earthquake data. *J. Geophysical Research*, 94(B5), 1989.
- [21] J. M. Lees. RSEIS: Seismic time series analysis tools. <http://bit.ly/Qcj60K>.
- [22] G. Liu, R. Tan, R. Zhou, G. Xing, W.-Z. Song, and J. M. Lees. Volcanic earthquake timing in wireless sensor networks. Technical Report MSU-CSE-12-8, CSE Dept, Michigan State University, 2012.
- [23] J. Nelder and R. Mead. A simplex method for function minimization. *The computer journal*, 7(4), 1965.
- [24] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *IPSN*, 2005.
- [25] C. Sadler and M. Martonosi. Data compression algorithms for energy-constrained devices in delay tolerant networks. In *SenSys*, 2006.
- [26] 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.
- [27] 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.
- [28] R. Tan, G. Xing, J. Chen, W. Song, and R. Huang. Quality-driven volcanic earthquake detection using wireless sensor networks. In *RTSS*, 2010.
- [29] G. Werner-Allen, S. Dawson-Haggerty, and M. Welsh. Lance: optimizing high-resolution signal collection in wireless sensor networks. In *SenSys*, 2008.
- [30] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh. Monitoring volcanic eruptions with a wireless sensor network. In *EWSN*, 2005.
- [31] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *OSDI*, 2006.
- [32] M. Withers, R. Aster, C. Young, J. Beiriger, M. Harris, S. Moore, and J. Trujillo. A comparison of select trigger algorithms for automated global seismic phase and event detection. *Bulletin of the Seismological Society of America*, 88(1), 1998.