Chapter 1 Introduction

If earthquakes developed more slowly (like hurricanes), it would be possible to optimize our response to earthquake emergencies; earthquakes would be less devastating. There would be time for expert seismologists to examine the incoming data and advise the affected regions of the appropriate response. By examining the overall shape and relative frequency content of the available ground motions, and with their knowledge of factors such as where the faults are, what size events they are capable of generating, which seismic stations are operating, and recent levels of seismic activity, etc, experienced seismologists can typically make fairly accurate judgments regarding whether the developing earthquake will be potentially damaging or not. They would be able to determine which areas will be significantly affected, and provide these areas with estimates of when the ground motions will arrive and what levels of shaking to expect. There would be time for emergency response agencies, local governments, and the media to coordinate their efforts to inform, advise, and direct the response of the affected communities to the approaching earthquake. In this scenario of hurricane-like earthquakes, by the time the damaging ground motions arrived at a given area, vulnerable structures would have been evacuated, trains and traffic stopped, airport activity suspended or diverted, data saved, computers and other equipment shut down, hazardous materials secured, and emergency services aware of which regions will be hardest hit and thus where to focus their resources (Goltz, 2002).

Unfortunately, earthquakes develop over short timescales; the time between the initiation of the rupture and the arrival of the damaging ground motions at a given site is on the order of seconds. It is not possible for a human to process incoming information and make the appropriate judgments (is it large or small, which regions will be affected, what levels of shaking are expected and when will the shaking commence, etc.) in a rational manner given so little time. If we wish to have a few seconds of warning between detecting the earthquake rupture and the commencement of strong shaking at a given site, the estimation/decision-making function of the human seismologist must be automated. This problem of providing affected regions on the order of seconds of warning of approaching strong ground shaking is known as seismic early warning and is the topic of this thesis.

1.1 Background on seismic early warning

While earthquakes occur much too fast relative to the speed at which humans are capable of processing information, in another sense, seismic early warning is possible because damaging earthquake ground motions propagate slowly relative to the speed at which information travels. The large amplitude ground motions from an earthquake (and hence the most damaging) are typically from the S-wave and later arriving surface waves, which travel at about 3.5 km/s; electronically transmitted information travels at about 300,000 km/s. If the initiating earthquake can be detected, and its characteristics (is it large or small?) identified early enough in the rupture process, information about the imminent strong ground motions can arrive at regions further from the source region on the order of seconds before the start of strong shaking.

Kanamori (2004) provides a recent review and history of research efforts in seismic early warning. The first printed reference to seismic early warning dates back to 1868, in an editorial by J.D. Cooper of the San Francisco Daily Evening Bulletin written after an $M = 7$ earthquake on the Hayward fault (Nakamura and Tucker, 1988; Kanamori, 2004). Cooper proposed the installation of an array of seismic detectors from 10 to 100 kilometers away from San Francisco. When large ground motions triggered the array, a signal would be telegraphed to San Francisco and would automatically ring a bell hung high in City Hall. Cooper's article in 1868 listed a number of concepts that still define the boundaries for the seismic early warning problem in modern times: 1) the need for a dense seismic network, 2) the longest warning times would be available from large distant earthquakes; there would be little or no warning from that were too close to close to the target warning region, and 3) the need to automate the system. Cooper recognized that human operators "might not always retain the presence of mind enough to telegraph at the moment or might sound the alarm too often" (Nakamura and Tucker, 1988).

The model for a modern seismic computerized alert network (SCAN) was proposed by Heaton (1985). In Heaton's model, ground motions recorded by a dense array of broadband, high dynamic range seismometers would be digitally telemetered to a central processing facility. The central processing site would take the input ground motions and calculate estimates of earthquake size and location. These estimates, along with their reliabilities, would be sent to system subscribers. Computers at subscriber sites would combine the available estimates (and their reliabilities) along with the subscriber location and site data to calculate when and what level of ground motions to expect, as well as what actions to take. The appropriate actions and what level of ground motions and reliability would be required to initiate the actions would vary depending on the subscribers. For example, there is little cost or inconvenience associated with stopping elevators or ensuring that gates at fire stations are open. These actions could be initiated even when the reliabilities of the SCAN estimates are low (and there is still a possibility that it is a false alarm). However, false alarms would be more costly for tasks such as diverting airport traffic or initiating shut-down procedures at a nuclear plants. These actions would only be initiated if the SCAN estimates indicated that there was a high reliability that ground motions exceeding some critical level were imminent. Again, the reliability and the ground motion level required to initiate action would vary from user to user. The SCAN would send continuous improved updates of earthquake size, location, and estimate reliabilities as the ground motions propagated to more seismic stations in the network.

Figure 1.1 shows a schematic of Heaton's seismic computerized alert network. The UrEDAS (Urgent Earthquake Detection and Alarm System) system used by Japan Railway (Nakamura, 1988) is perhaps the most famous example of a seismic computerized alert network, where the estimates from the central processing site are used to automatically initiate damage-mitigating actions, in this case, stopping high speed trains. In general, most recent research has been focused on the problem of reliably determining whether the event will be damaging or not from the data available from the initial part of the rupture (task of the central processing facility in Figure 1.1); the source estimation and user response problems have typically been dealt with separately. The need to move from this separatist view to a more integrated approach considering both the source estimation problem and the user response within a unified framework will be discussed.

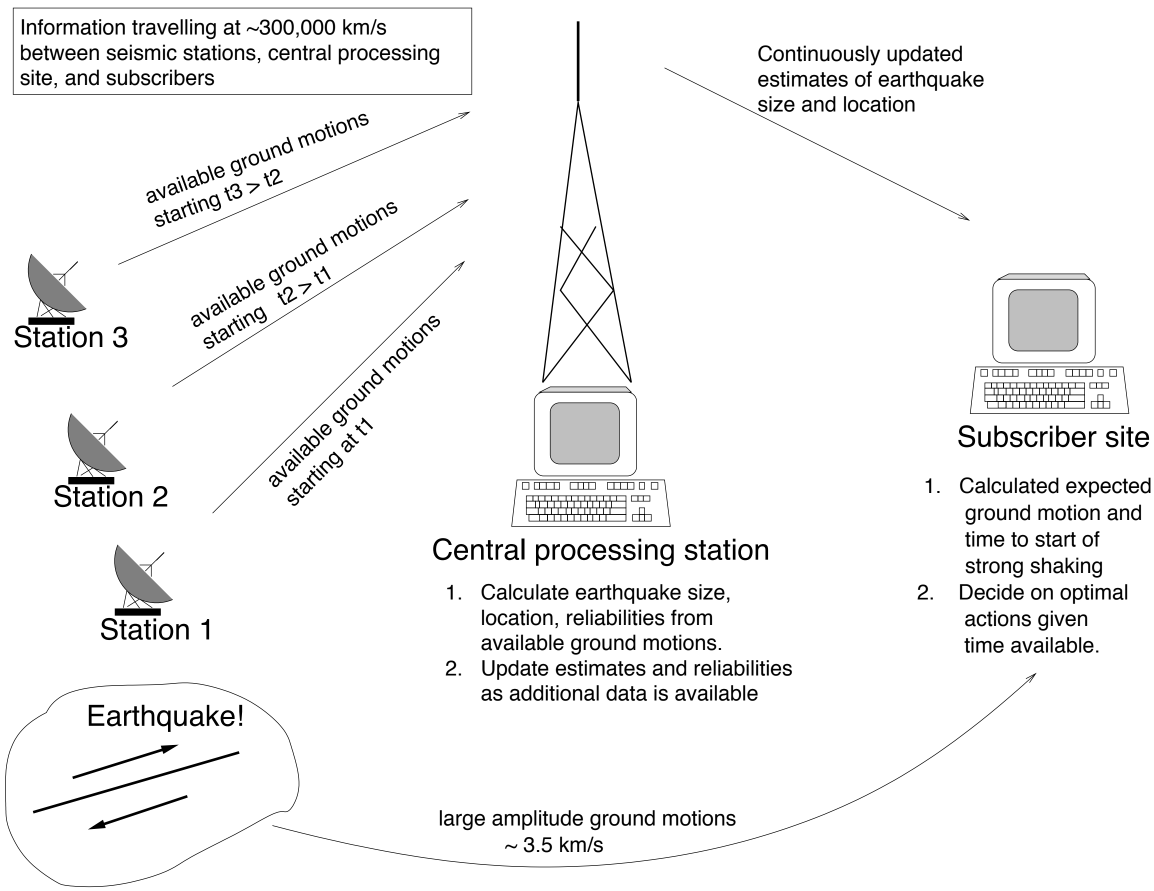

Figure 1.1: Schematic for a seismic computerized alert network (SCAN), adapted from Heaton (1985). At the time of this writing, most research effort has been focused on the tasks of the central processing station - that is, in providing reliable estimates of magnitude and location (or of expected ground motions) given the initially available data. Aside from UrEDAS (Nakamura, 1988), a general framework to guide different types of subscribers in deciding on the optimal action given the available warning time has yet to be proposed.

Kanamori (2004) classifies early warning approaches as either regional or site specific. In the regional or traditional seismological approach, magnitude and location are first estimated from the available data; these estimates are used to predict ground motions at other sites. This approach is employed in Japan, Taiwan, and Mexico. In on-site or site-specific approaches, available ground motions at a given site are used to predict the later-arriving (and larger) ground motions at that and other sites. Work on site-specific approaches has been done by Kanamori (2004), Wu and Kanamori (2004a), Wu and Kanamori (2004b), Leach and Dowla (1996), and Bose et al. (2003).

Existing paradigms for seismic early warning all base their estimation on the early portion of seismograms available from the initiating rupture. Most methods share the characteristic that the initial estimates have large uncertainties or have trade-offs between parameters; additional data decreases these uncertainties or resolves these trade-offs at the cost of reducing the available warning times. The large uncertainties of the initial estimates should not be surprising; in most cases, the problem of locating the earthquake and determining its magnitude is not yet uniquely determined given the initially available data. (For instance, it is not possible to uniquely locate an earthquake and determine its magnitude from the first few seconds of P wave amplitudes at a single station via the traditional seismological method.) That almost all methods share this characteristic of large initial uncertainties is also expected; rules of logic state that if there are several different ways of using the same information, the same conclusions will be reached irrespective of the analysis-path chosen (Sivia, 1996). This implies that a paradigm-shift is necessary if we wish to make substantial gains over the current state of the art in seismic early warning. Improving the reliabilities or resolving the trade-offs inherent in the initial estimates requires additional information. Since waiting for the availability of future observations is not a satisfactory solution (as this decreases the available warning time), perhaps we need to look to the past as our source of information. The introduction of prior information into the seismic early warning problem is a central theme of this thesis, and is facilitated by a Bayesian approach.

1.2 Virtual Seismologist (VS) method for seismic early warning

The Virtual Seismologist (VS) method is a Bayesian approach for seismic early warning designed for network-wide deployment on modern digital seismic networks such as the Southern California Seismic Network (SCSN). It is designed for regions where the interlacing of active fault systems with areas of high population density mean that the seismic hazard and risk is geographically distributed, such as the case for California and Japan; it can also be applied to the more straightforward situation where there is a well-defined target warning area and the source of damaging earthquakes is relatively fixed, as is the case for Bucharest and Mexico City. The VS method is modeled on "back of the envelope" methods used by human seismologists to quickly distinguish between small and large events. Expert seismologists can typically determine earthquake characteristics from quick visual examination of overall shapes and relative frequency content of available ground motions. One component of the VS method is a new way to estimate magnitude from ratios of available ground motion amplitudes as early as 3 seconds after the initial P wave detection. The use of prior information differentiates the VS method from other proposed paradigms for seismic early warning. Any approach to estimate the size and location of an earthquake or the expected ground motions at a given site from the typically sparse set of available observations immediately after the initial P wave detection from an on-going rupture will have large uncertainties or trade-offs between the estimated parameters. The trade-offs are eventually resolved and the uncertainties reduced by additional data. However, waiting for additional observations also reduces the available warning time. The VS method uses prior information, such as the state of health of the seismic network, fault locations, and previously observed seismic activity, to resolve the trade-offs in magnitude and location estimates left unresolved by the limited available observations in the initial stages of the earthquake rupture, similar to how a human seismologist might incorporat such information when making judgments with a sparse set of observations.

1.3 Bayes' theorem in qualitative terms and a road map

Our approach to seismic early warning consisted of asking the following questions: given the available observations at some time t, what are the most probable magnitude and location estimates? How do these estimates evolve as additional data become available?

As mentioned previously, the Virtual Seismologist (VS) method is a Bayesian approach. According to Bayes' theorem, the best (assuming best = most probable) magnitude and location estimates at any given time t is a weighted combination of 2 terms: the likelihood and the prior. The likelihood function is a mapping between the available observations to estimates of the parameters of interest. In the case of seismic early warning, it relates the available ground motion amplitudes to estimates of magnitude and location. Since the likelihood depends only on the observed seismograms, there are large uncertainties or trade-offs between parameters when there are not enough data available to uniquely determine magnitude and location. The prior is a statement of our beliefs regarding earthquake occurrence before examining the available observations. In the VS method, it is a combination of disparate types of information that are potentially relevant to the early warning estimation process. The types of prior information evaluated in this thesis are station geometry, previously observed seismicity, and fault locations. Bayes' theorem, how to formulate the likelihood function and the Bayesian prior, and how to apply these to seismic early warning are discussed in depth in Chapter 4.

The VS method operates on envelopes of various channels (horizontal and vertical acceleration, velocity, and filtered displacement) of ground motion. Ground motion envelopes are defined as the maximum absolute value of ground motion on a given channel over a one second window. It is this particular data stream that is transmitted from the seismic stations to the central processing facility of the seismic network in closest to real-time, and it thus a logical choice for seismic early warning.

To define the Bayesian likelihood function, ground motion models relating the observed ground motion envelopes to magnitude and epicentral distance (or location) are required. I developed a paramterization that decomposes observed ground motion envelopes into P-wave, S-wave, and ambient noise envelopes. I applied this envelope parameterization to a large suite of observed ground motion envelopes recorded within 200 km of Southern California earthquakes in the magnitude range $2 \leq M \leq 7.3$. Chapter 2 describes this envelope parameterization and the development of envelope attenuation relationships for various channels of ground motion. Given a magnitude and distance, envelope attenuation relationships can be used to predict the expected ground motion envelope as a function of time.

The envelope attenuation relationships discussed in Chapter 2 describe how various channels of Southern California ground motion envelopes vary over a wide range of magnitude, distance, site condition, and frequencies. In Chapter 3, I use these envelope attenuation relationships to examine some average properties of ground motion. Some of the topics to which these envelope attenuation relationships are applied are: the difference between rock and soil sites, the magnitude and distance scaling of P- and S-waves, and the difference between horizontal and vertical ground motions.

Chapter 4 discusses Bayes' theorem and its application to seismic early warning. The focus is on 1) how to use ground motion ratios to estimate magnitude and 2) how to define the Bayesian likelihood function using the ground motion ratios and the envelope attenuation relationships developed and discussed in Chapters 2 and 3. The Bayesian prior is also discussed, although this concept is most effectively explained by example.

Chapters 5 through 8 describe the application of the VS method to the incoming ground motion envelopes from actual Southern California earthquakes. The events used are the 16 October 1999 $M = 7.1$ Hector Mine, the 22 December 2003 $M = 6$ San Simeon, the 28 September 2004 $M = 6$ Parkfield, and the 3 September 2002 $M = 4.75$ Yorba Linda events. These examples illustrate the importance of including prior information in the seismic early warning estimation process. Aside from its primary intended application for seismic early warning, the VS method can also be used as part of routine seismic network operations to provide a check to arrival-based location estimates. The concept of amplitude-based locations was introduced by Kanamori (1993). The traditional seismological approach to locating earthquakes is based on finding the location and origin time that minimizes the residuals between observed and calculated arrival times. This timing-based approach is the most precise way to locating earthquakes. However, they are vulnerable to large errors. Given the available observations at some large t (say 80 seconds) after the earthquake origin time, the VS method can be used to provide amplitude-based location estimates. Amplitude-based location estimates are not as precise as those based on arrival times. In particular, such amplitude-based locations are sensitive to non-uniform azimuthal distribution of stations and unaccounted site effects. However, they are completely independent of timing and are very robust. Robust amplitude-based estimates can be used as a check on the arrival-based locations. If the arrival-based locations are consistent with the amplitude-based estimates, then the arrival-based estimates are most likely correct. However, if there is a large discrepancy between the arrival and amplitude-based locations, then perhaps the arrival-based location need to be reviewed. Arrival-based locations will be discussed for each of the example events.

1.4 Integrating the source estimation and user response problems in seismic early warning

Like most research on seismic early warning, this thesis was initially focused on only the source estimation problem. The approach adopted was to use Bayes' theorem to address the question: given the available observations from a rupturing earthquake, what are the *most probable* estimates of magnitude and location? In general, a Bayesian approach is well-suited for solving problems in real-time seismology. Real-time problems are typically characterized by the conflicting requirements of timeliness and reliability of estimates. Typically, the initial estimates will be made when there are less than adequate observations to fully resolve the parameters of interest. The influence of the Bayes prior is strongest in such situations; its influence diminishes with additional observations. Station geometry, previously observed seismicity, and fault locations are useful in resolving the trade-offs in the initial estimates and are included in the Bayes prior. Because we were looking for the *most probable* magnitude and location, it also seemed natural to include the Gutenberg-Richter magnitude-frequency relationship in the prior. The Gutenberg-Richter relationship says quite firmly that smaller magnitude earthquakes occur more frequently than larger events. In the case studies applying the Virtual Seismologist (VS) method to Southern California earthquake datasets (Chapters 5-8), VS estimates with and without the Gutenberg-Richter (G-R) in the prior were tracked. The different priors meant that there were differences in the initial estimates; these eventually converged once sufficient observations were available. Whenever there were trade-offs in the source parameters, estimates *without* the G-R in the prior were consistently closer to the eventual magnitude than those *with* the G-R. This led to the question of what information should be broadcast to the user. Initially, from the user perspective, it seemed that estimates that were closer to the actual magnitude would be better. This implied that choice of prior, in particular, the inclusion of the Gutenberg-Richter relationship, was inappropriate. However, the question of most probable magnitude cannot be divorced from the Gutenberg-Richter relationship, which has been observed to hold in general worldwide.

This apparent inconsistency is resolved by explicitly accounting for user considerations. How would a user/subscriber to an early warning system optimally respond to the warning information? Do these considerations dictate what type of information should be broadcast, and ultimately, how to address the source estimation problem in the first place? Chapter 9 describes how a user with information about the relative costs of missed warnings and false alarms for a given application might apply simple decision theory in responding to early warning information. Among the conclusions are: 1) if the cost of false alarms is relatively high, the Gutenberg-Richter should be included - either in the source estimation or user decision process, and, 2) if the cost of missed warnings is high, the user will have to live with false alarms. Perhaps the most important lesson from this thought exercise is that traditional separation of the source estimation and user response problems in earthquake early warning is an artificial divide. These are two sides of the same coin. A more integrated approach that recognizes the role of the user decision-making process in formulating the source estimation problem is necessary. The source estimation and the user response should be treated as two parts of a single problem, as opposed to two somewhat related but separate problems.

1.5 Communicating the earthquake source estimates and their uncertainties

Ultimately, the source estimates and their associated uncertainties will be used to calculate predicted ground motions at subscriber sites. The uncertainties in the source estimates dictate the uncertainties in the predicted ground motions, which, as discussed in Chapter 9, play a role in determining the optimal user response. It is therefore not sufficient to transmit just the most probable source estimates. The uncertainties are equally important.

A high enough density of stations reduces the complexity of information necessary to broadcast. For events in the densely instrumented regions of the network -such as the epicentral region of the Yorba Linda earthquake (Chapter 8)- the station spacing is close enough that the location of the first triggered station is a fairly good estimate of epicentral location. The marginal probability density functions (pdfs) of the source estimates (magnitude and epicentral location) can be reasonably described by Gaussian pdfs. Each estimate update would involve 6 numbers - 3 means, and 3 standard deviations. This is a reasonable amount of information to transmit to users.

This is not the case for earthquakes occurring along the periphery of the network, or in regions where station density is low. A seismic network will always have outer boundaries, and more often than not, will have regions in its interior where inter-station distances are large. In these regions, the possible errors from assuming the epicenter is located at the first triggered station can be greater than a hundred kilometers. It is for events in such regions that a Bayesian approach, in particular, the use of prior information, is most beneficial. Unfortunately, the use of previously observed seismicity, fault locations, and station geometry to constrain the location estimates make the marginal pdfs highly non-Gaussian. One possibility is to transmit the portion of the Bayes posterior pdf (whose maximum corresponds to the most probable source estimates) within the first triggered station's Voronoi cell. The users could estimate the uncertainties in the predicted ground motions from the uncertainties in the source estimates. Transmitting portions of a multi-dimensional pdf will take a large amount of bandwidth. Alternatively, a more efficient scheme would be for the network to provide the likelihood function (which can always be adequately described by Gaussians) to the users, and for the users have the available the appropriate prior information (short-term earthquake forecasts and national hazard maps). The final step in the source estimation (combining the likelihood and prior) and the ground motion prediction calculations would be calculated by the users on site. The role of the network would be to develop and maintain the source estimation software, to broadcast the likelihood function to users, and to educate the users on how to optimally use the early warning information to guide decisions/actions in the few seconds between the initiation of earthquake rupture and the onset of damaging ground motions at the user site. This issue of how to communicate the source estimates is another indication of the need for a more integrated approach to seismic early warning.