Chapter 10 Conclusions

This final Chapter is made up of the following: a summary of the Virtual Seismologist method for seismic early warning, comments on implementation issues, conclusions and other questions raised by this approach, and some possible directions for seismic early warning.

10.1 A summary of the Virtual Seismologist (VS) method for seismic early warning

The Virtual Seismologist (VS) method for seismic early warning is designed to answer the following questions: What are the most probable magnitude and location estimates given the initially sparse set of observations from an on-going earthquake rupture? How do these estimates evolve as additional data becomes available?

A Bayesian framework is used, and this is proven by the examples (Hector Mine, San Simeon, Parkfield, and Yorba Linda, California earthquakes) to be a powerful and effective approach to seismic early warning. The VS method is modeled on the thought process a human seismologist might go through as she examines incoming data from an on-going earthquake. Human seismologists are, in general, very efficient at making "back of the envelope" estimates of earthquake characteristics given the available data. One half of this "back of the envelope" thought process is a pattern-recognition problem. By visual examination of the shape and relative frequency content of the available ground motions, expert seismologists can obtain reasonable estimates of earthquake source parameters such as magnitude and distance very quickly. In the VS method, this is the role of the Bayes likelihood function; the likelihood function is a mapping of the available amplitude observations to the source parameters consistent with those observations. The ground motion models used to define the Bayes likelihood function relate 1) observed ground motion ratios to magnitude, 2) P- and S-wave amplitudes of vertical and horizontal acceleration, velocity, and filtered displacement to magnitude and distance. These ground motion models (or envelope attenuation relationships, as they have been referred to in this thesis) are new to this study. They are valid over a much larger magnitude range that traditional attenuation relationships, and are novel in the separate characterization of P-wave amplitudes. The P-wave relationships are essential for seismic early warning, since, to maximize the available warning time, it is necessary to start the early warning estimation process with the observed P-wave amplitudes. Aside from their intended application in seismic early warning, these envelope attenuation relationships are interesting because they characterize the average properties of a large suite of Southern California ground motions recorded by the California Integrated Seismic Network. (The development of these envelope attenuation relationships and their application to characterizing average ground motion properties were the subjects of Chapters 2 and 3.)

The other half of the “back of the envelope” human seismologist thought process introduces background knowledge regarding earthquake occurrence in the monitored region into the estimation process . A seismologist typically knows which stations in the network are operating or problematic, where the faults are, what types of events are expected from which faults, the level and geographic distribution of recent seismic activity, etc. Humans instinctively include this type of background information when it is necessary to make decisions or judgments with insufficient information. Similarly, in the VS method, the Bayes prior aids in transitioning the estimation/judgment-making (is it large or small, potentially damaging or not?) process from an under-determined problem, where the available information is not able to resolve the trade-offs in the earthquake source parameters, to an over-determined problem, where there is enough data to properly and uniquely constrain the source parameters. The type of information included in the Bayes prior are: station geometry, previously observed seismicity, fault locations (at least, the San Andreas fault), and the Gutenberg-Richter magnitude-frequency relationship.

The advantages of a Bayesian approach, in particular, of including prior information in the estimation process as well as the use of not-yet arrived information (Horiuchi et al., 2004; Rydelek and Pujol, 2004), are most evident for earthquakes in regions with relatively low station density. It is most impressive with the M=7.1 Hector Mine mainshock. The initial VS estimate, with the initial P detection and the 3 second P-wave amplitudes at station HEC, is an $M = 6.2 \pm 0.495$ event 2 km away from the actual epicenter; 2.5 seconds later, with no additional P arrivals, the updated VS magnitude estimate is $M = 7.2 \pm 0.415$. The spatial (within 1 km) and temporal (within 24 hours) cluster of 18 foreshocks about the mainshock is very strong and relevant prior information.

The seismicity-based prior does not always give the "correct" location, since not all mainshocks are preceded by foreshocks. For instance, for the M=6.5 San Simeon and M=6.0 Parkfield events, there was seismicity in the preceding 24 hours within the first triggered station's Voronoi cell that was unrelated to the mainshock. This may lead the initial location estimates to be "off". How much they are "off" is limited by the dimensions of the Voronoi cell of the first triggered station and the not-yet arrived information. This is corrected by the not-yet arrived information, usually before there are enough arrivals to uniquely locate the epicenter. For the San Simeon and Parkfield mainshocks, the VS location estimates are within 7 and 1 km, respectively, of the mainshock epicenters by the time of second VS update, 3 seconds after the initial estimate. This is part of the appeal of the Bayesian approach. For the situations when the prior seismicity is related to the on-going mainshock, the VS estimates immediately give the correct location. Jones (1984) studied 20 $M \geq 5$ mainshocks in the San Andreas system between 1966 and 1980 and found that 35% were preceded by foreshock sequences within 1 day and 5 km of the mainshocks. Abercrombie and Mori (1996) studied 59 California and Nevada $M \geq 5$ mainshocks between 1975 and 1994. On the whole, they found that 44% had foreshocks within 5 km and 30 days of the $M \geq 5$ mainshock. In all but one sequence, the foreshock activity continued to within 24 hours of the mainshock. For Southern California, using previously observed seismicity in the Bayes prior means that the initial VS location estimates (3 seconds after the initial P detection) are correct between $35-45\%$ of the time for $M \geq 5$ events. For the situations when the previously observed seismicity is unrelated to the on-going mainshock, the initial estimates are adjusted by the not-yet arrived information and corrected before there are enough arrivals to exactly locate the epicenter. Again, it should be emphasized that the Bayes prior is only important and influential when there is not enough observations available to fully constrain the magnitude and location. When sufficient observations are available, the VS estimates are determined by the data. For instance, in the Yorba Linda earthquake, the prior information is very quickly irrelevant and unnecessary; at the time of the initial VS estimate (3 seconds after the initial P detection), there are enough arrivals at the adjacent stations to locate the epicenter.

In evaluating the VS method, it is important to address separately 1) the benefits of introducing different types of prior information into the early warning estimation process and how the role of the Bayes prior evolves as additional observations become available, and 2) the effectiveness of the likelihood function. The likelihood function is how the observations come into the estimation process. In the VS method, it is based on ratios of ground motions as well as envelope attenuation relationships for vertical and horizontal P- and S-wave amplitudes. There are many other methods to translate the initially available amplitude and arrival observations to estimates of magnitude and location. The methods developed by Nakamura (1988), Allen and Kanamori (2003), Wu and Kanamori (2004a), Wu and Kanamori (2004b), Horiuchi et al. (2004), Rydelek and Pujol (2004) are some examples; a more comprehensive and up-to-date list can be found in a recent seismic early warning review paper by Kanamori (2004). Any of these methods (or possibly, combinations of different methods) can be used to define the likelihood function and placed in a Bayesian framework. For instance, the VS method incorporates the use of not-yet arrived data, as described by Horiuchi et al. (2004) and Rydelek and Pujol (2004), along with the relationships for ground motion ratios and other observed envelope amplitudes to constrain possible magnitudes and locations. The benefits of a Bayesian approach (relative to not using prior information) are most evident for earthquakes with foreshock sequences as well as events occurring in regions with low station density.

10.2 Implementation issues

The focus of this thesis was on the question: what are the most probable estimates of magnitude and location given the initially sparse set of available amplitude and arrival information at some time t shortly after the initial P detection? This work establishes that, given relevant "background" information, such as station geometry, previously observed seismicity, fault location information, the Gutenberg-Richter magnitude-frequency relationship, as well as an appropriate set of ground motion models to map the available observations to earthquake source parameters, reasonable Bayesian estimates can be made as early as 3 seconds after the initial P detection. When there is not enough information from the on-going rupture to constrain the early warning estimation problem, these Bayesian estimates are closer to the actual earthquake source parameters that estimates based on the observed amplitudes and arrivals alone.

While developing a working prototype of the VS method was beyond the scope of this work, there are implementation issues I would like to address.

10.2.1 Processing time

In illustrating the VS performance for various Southern California mainshocks, the question of processing time is not directly addressed. The focus was on "what constraints on the source parameters can be extracted from the available information", and not on "how fast these constraints can be extracted". In practice, of course, the question of how fast these estimates can be extracted is crucial to successful seismic early warning.

Maximizing the Bayes posterior yields the VS estimates for the most probable magnitude and location estimates given the available data. It is equivalent to minimizing a function over a 3 dimensional space. Since the observations are assumed to be independent, the problem reduces to minimizing a residual sum of squares. The calculations necessary for the VS method are ideal for parallel implementation; the value of the Bayes posterior at any given coordinate in magnitude-location space is independent of the values of its neighbors. This means that the required communication between processor nodes will be relatively small, which bodes well for the processing time. Of course, the more processors involved, the faster the processing time. Even if the costs of the hardware necessary to implement the VS method might be prohibitive today, with the steady increase of processing power per dollar in successive generations of computers, it may be that the VS method will be very affordable to implement in the next few years. It is encouraging that Horiuchi et al. (2004) report that their real-time earthquake information system (REIS) operates on 2 Linux computers, and can determine earthquake hypocenters from 5 or less arrivals in about 0.1 seconds.

Aside from parallel implementation, using "informed" sampling could potentially reduce the processing time. The VS estimation problem is a non-linear inverse problem; in general, the possible approaches to non-linear inverse problems are classified as either gradient-based or direct-search methods. A variant of the Neighborhood Algorithm, which is a direct-search method used extensively in this thesis, could be developed, where the sampling is initially concentrated about the regions where the location prior is peaked (for example, at locations of previously observed seismicity within the first triggered station's Voronoi cell). The Neighborhood Algorithm lends itself well to parallel implementation. In fact, message passing interfaces (MPIs) were added to the Neighborhood Algorithm by Sambridge (2003) to allow it to run on Beowulf clusters.

10.2.2 Station geometry information

The station geometry information, in the form of Voronoi cells and the probability density functions of epicentral distances consistent with the given Voronoi cells, should be updated whenever there is a change in the operating stations in the network, for example, whenever a new station is put on-line or taken off-line. These calculations take only a few seconds, and would ideally become part of routine network processing.

These calculations are truly prior information, in that they available even before the earthquake initiates. The Voronoi calculations are not meant to be done in real-time as part of the early warning estimation process; however, this information should be ready to use once the initial P wave is detected.

The sensitivity of the VS estimates to inaccuracies in the station geometry information - for instance, assuming a station is operational, when it is in fact not, or that the event is closest to the station with the first reported P detection (it might be possible that the station that was actually closest did not send data) - was not dealt with in this study. However, these are important operational questions to address.

The method of not-arrived data is used by Horiuchi et al. (2004) to remove erroneous arrival times in their real-time calculations of hypocenter locations. There is a potential to also use this concept to determine, in real-time, whether stations are operational or not, and thus determine whether the station geometry information is accurate. For instance, given an initial configuration of stations assumed to be operational, the distance between various stations is known. Consider a station (let's call it A) with n Voronoi edges. If station A has the first P detection, then the time interval between the P arrival at one of the n surrounding stations and the initial arrival at A must be less than or equal to the distance between these stations divided by an average P wave velocity. If there is no reported arrival at an adjacent station after (distance between stations) $\times \alpha$ seconds (where α is an average P-wave velocity), then there is a possibility of a problem at the adjacent station.

Similarly, the incoming amplitude observations can be used to remove erroneous signals from the estimate. The magnitude estimates from the vertical ground motion ratios are typically within 1 magnitude unit above or below the actual magnitude. If a particular station gives a ratio-based magnitude estimate that differs by more than 2 magnitude units from the mean of the available ratio-based estimates, it is possible that that particular station is having problems.

10.2.3 Seismicity, fault locations, and foreshock/aftershock statistics

As mentioned numerous times, the manner in which previously observed seismicity and fault locations are included in the Bayes prior are very ad hoc. However, the intention in this thesis was to show that this type of information can be very useful in seismic early warning, rather than to lay out the exact details of how to do this. Fortunately, this problem has already been thoroughly addressed by researchers working on foreshock/aftershock statistics (Reasenberg and Jones, 1989) and short-term, seismicity-based forecasts, such as the STEP (Gerstenberger et al., 2003) model. In the ad hoc approach in this thesis, the previously observed seismicity was used primarily as a constraint on location; in short-term earthquake forecasts, the magnitudes and time elapsed since these events also affect the forecasted earthquake probabilities. For instance, there were two events within the Voronoi cell of station PKD, which had the first P detection from the San Simeon mainshock. One was an M=1.88 event on the San Andreas about 22 hours before the San Simeon event, the other an M=1.22 event 8.5 km away from the eventual epicenter 11 hours before the mainshock. From the ad hoc method in this thesis, the San Andreas event is given more weight, due to its proximity to the San Andreas. However, proper use of foreshock/aftershock statistics would probably weigh the M=1.22 event more, not because it is closer to the mainshock epicenter (since this is not known), but because it had a slightly larger magnitude and was more recent. In practice, the output of these short-term, seismicity-based forecasts should be used to quantify how the recent seismic activity affects relative earthquake probabilities.

In the examples in this thesis, the San Andreas fault trace was the only fault included in the Bayes prior (for the San Simeon and Parkfield events). For Hector Mine, no fault information was included, even though the Landers fault trace is close by and there are a number of faults associated with the Eastern California Shear Zone in the epicentral region. The national hazard maps (Frankel et al., 1997) should be used to decide which faults should be included in the Bayes prior, and how much weight they should be assigned relative to each other. Typically, when both short-term (from recent seismicity) and long-term (from fault locations) information are available, the short-term forecasts dominate the hazard calculations.

10.2.4 Communicating source estimates and uncertainties

In regions where station density is high, such as the epicentral region of the Yorba Linda earthquake (Chapter 8), the Voronoi cells are small enough such that the errors is assuming that the epicenter is located at the first triggered station are small (on the order of 10 km). In addition, the prior information is unnecessary soon after the initial P detection because there is quickly enough information to uniquely constrain the epicentral location. For events in densely instrumented regions of the network, the marginal pdfs of the magnitude and location estimates can be adequately described by a 3 Gaussian functions. If the seismic network were broadcasting the source estimates, it would be relatively straightforward to describe and broadcast these estimates to users.

Having high enough station density (such that inter-station distances are on the order of acceptable errors in epicentral location estimates) saves a great deal of complexity in terms of the information that needs to be broadcast, assuming that the seismic network performs the source estimation calculations and broadcasts warning information to users. However, seismic networks typically have a finite extent, and will thus have outer edges. Also, uniformly dense station coverage is a rare and expensive thing, and it is common to have some regions in the interior of the network with relatively low station density. It is important for early warning systems to be able to handle events occurring at the outer edges of the network or in regions with sparse instrumental coverage. For events occurring in such regions, such as the Hector Mine, San Simeon, and Parkfield earthquakes, the possible errors in assuming that the epicenter is located at the first station can be large (in the hundred km range). The Voronoi boundaries, not-arrived data, previously observed seismicity, and fault location information allow most probable source estimates to be obtained. However, the uncertainties on such estimates which rely on these types of prior information are not simple to describe. (See the location estimates 3 seconds after the initial P detection in Chapters 5, 6, and 7.) If the seismic network is charged performing the source estimation calculations and broadcasting warnings, it may be necessary to transmit the full 3-d Bayes posterior pdf. This may require a prohibitive amount of bandwidth.

It would be more efficient for the network to provide users with the likelihood function, which can usually be well described by Gaussians. (It is the prior information that makes the Bayes posterior highly irregular and hence non-Gaussian.) Users would have their systems set up to regularly (perhaps once a day) download the appropriate prior information (short-term seismicity-based forecasts from the USGS, station health from the network). In the event of an earthquake, users would perform the final step in the source estimation (combining the likelihood and the prior) and ground motion prediction calculations on site. The network's role would be 1) to develop algorithms and software to determine the likelihood function, 2) develop the software to combine the likelihood function and prior information and distribute this to users, and 3) educate users on how to optimally use the Bayesian source estimates. The users would 1) install the network-provided software on-site, 2) regularly download updated short-term earthquake forecasts (STEP) from USGS and state-of-health reports from the network, 3) develop algorithms to predict the ground motions at their site, and 4) determine, based on the relative cost of missed and false alarms for their particular application, whether to base their decisions on estimates with or without the Gutenberg-Richter relationship, and the necessary thresholds for predicted ground motion levels and uncertainties. The network will have to educate the users about the implications of their choices for these user-specific parameters. How much computational resources to allot to the calculations would be at the discretion of the user; the network could provide suggested minimum processor specifications.

The system described requires a level of coordination between the network and the users that is not consistent with the traditional divide in seismic early warning research between the source estimation (presumably calculated and broadcast by the network) and how the user might respond to early warning information. In general, there is a need to move towards a more integrated view of seismic early warning, (as just discussed in the previous Chapter), where source estimation and user response are considered as two interacting and interrelated parts of a single problem.

10.3 A unified earthquake information system: a possibility for the future

From Kanamori (2004), real-time seismology is the practice in which seismic data are collected and analyzed quickly *after* a significant event, so that the results can be used for early warning and post-earthquake response. The various time scales involved in applications and products of real-time seismology (early warning, rapid magnitude and location, ShakeMaps, real-time loss estimates) and the various needs of the users of this information require that estimates be made available as quickly as possible. There are contrasting needs for speed and accuracy of information. In general, the earliest estimates (of source parameters, distribution of shaking, etc) will be based on a minimum amount of observations, and will thus have large uncertainties. It is at this early stage of the estimation process that a Bayesian approach is most beneficial. However, these benefits are directly related to how appropriate the prior information is to the process being studied.

A Bayesian approach to seismic early warning, such as the Virtual Seismologist method, provides a strong link between earthquake forecasting and real-time seismology. Research in earthquake forecasts is in general focused on developing models to describe earthquake occurrence in a region and quantify the implications of these models on seismic hazard. Earthquake forecasts provide the ideal priors for the VS method. At the moment, short-term seismicity-based forecasts are the most developed. However, as the general state of knowledge develops, perhaps it will be possible to accurately model the spatial and temporal effects of other factors, such as Coulomb stress transfer, the role of fluids in faults, various triggering mechanisms (static, dynamic, tidal, SmS, etc.), or perhaps we will find that it is impossible to predict with high precision the effects of such factors. As earthquake forecast models improve, so will the quality of the Bayes prior for seismic early warning. The better and more relevant the prior information is, the more reliable the initial estimates for seismic early warning. Improvements in the reliability of seismic early warning will encourage more applications to be developed that make use of such warnings. Successful use of these warnings will help our communities better react to earthquake emergencies and mitigate the damage caused by earthquakes. It is an appealing vision for the future to be able to develop a unified earthquake information system covering various time scales - from long-term hazard (years), to short-term forecasts (days), to seismic early warning (seconds), rapid magnitude and locations, to ShakeMap, to predicted damage distribution. Such a unified system would help our communities take the appropriate precautions and actions over various time scales before, during, and after the ground has stopped shaking. A Bayesian approach to seismic early warning is an important step towards such a system.