Sound Perception in Source Process Analysis of Seismic Events

Matthew Vaughan

Advisors: Benjamin Holtzman, Felix Waldhauser, Lapo Boschi

Department of Physics

Case Western Reserve University

Cleveland, OH 44106-7079

May 2nd, 2014

Abstract:

The use of sound in seismology has been sparse, but the ability to utilize another dimension of analysis is an overlooked one for even experienced seismologists. In particular, for lesser-understood areas of seismology, such as source processes, sound may be an invaluable tool. This project involved three separate comparisons of pairs of phenomena: a landslide to an earthquake, equal magnitude earthquakes on two different faults, and clusters of repeating earthquakes. The main goal was to see in what ways sound could inform more traditional analysis of these types of events. The landslide occupied a vastly narrower frequency range than the earthquake, and attenuated quite differently due to extreme differences in source process. The data retrieved from the two different faults were damaged in such a way that the analysis must be called into question, but aspects of the sounds suggested differences in the mechanical properties of the faults. The repeating events were remarkably similar, and introduced questions regarding sound perception and how it links directly to source processes.

Introduction:

The fundamental processes underlying earthquake mechanics are poorly understood at best. Earthquakes result from the movement of tectonic plates across a fault, when friction along a patch of a fault plane causes locking along the fault. As relative motion along the fault continues, and stored energy builds, an earthquake occurs when irregularities along such a high-friction patch break at a rupture point that nucleates outward. While these basics are understood to a degree, details about the geometry and time evolution of the rupture (or the “source function”) are not well known to any extent. Since seismic data is limited to surface displacements recorded on nearby seismometers, determining the source function is a non-straightforward problem. In this project, we aimed to develop a new approach to this problem, namely using humans’ cognitive perception of motion and other features in sound to differentiate between different source functions.

Human hearing is an incredible tool that has yet to be tapped into in a significant way in seismology research. “Audification,” or the process of transforming seismic data into sound, has been used primarily as an educational prop (Peng, 2012), rather than as an added dimension to help researchers conduct analysis. Our abilities to combine spatial-temporal cues with audio cues can give more insight into problems (Holtzman, 2013) like the source function problem. The potential benefits for gaining a deeper understanding of earthquake source processes are enormous. With better understanding of earthquake fundamentals comes better earthquake prediction, on potentially significant time scales. While this is unlikely to be an immediate payoff for any single research project, any step toward this level of understanding is a powerful one.

Because traditional analysis mostly involves focal mechanisms, cross correlations between events, and looking at individual waveforms, the amount of detail may be high, but the coherence of these together is piecewise at best. Even automated, computer-powered analysis relies on impulses more than anything else, while ignoring any “noise” in the data. Yet in these segments of “noise,” we can hear minute and nuanced details that can inform us beyond traditional analyses. Our sound perception relies heavily on subconscious associations to past experiences, which allows a powerful level of pattern-matching, and creates a more universal language for seismologists and non-seismologists to engage in.

Methods:

Earthquakes primarily occupy frequency ranges outside of the range of human hearing:

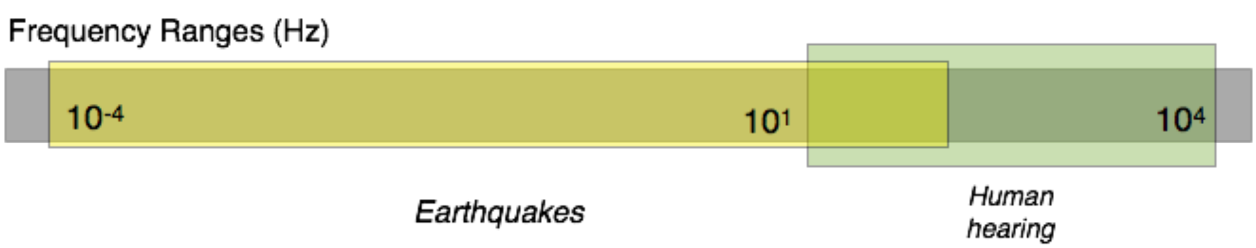

Figure 1. Human hearing occupies a range of about 20 Hz to 20 kHz. Earthquakes can generate waves with periods on the order of thousands of seconds, but can also have energy in the audible range.

The range of frequencies that humans can hear is approximately 20 Hz to 20,000 Hz, and earthquakes can generate waves with periods on the order of thousands of seconds. In addition to this, they can have energy in the audible range as well. So not only do earthquakes primarily occupy a range significantly lower than the human audible range, but they also occupy an extremely *wide* range as well. Because of these frequency limitations, we must utilize various mathematical techniques to both limit the range of frequencies we listen to at one time, and also shift those frequencies into the human audible range.

Identify and Isolate Frequency Ranges

To determine which frequencies we should try to isolate in the signal, we must determine where the most useful frequency content resides. Typically the most interesting frequency content will be where the bulk of the frequency content is represented. To determine the frequency content of the data, we performed Fast Fourier Transforms on the signals. A Fast Fourier Transform (FFT) is a fast, computational method for the decomposition of a time-domain signal into the frequency domain:

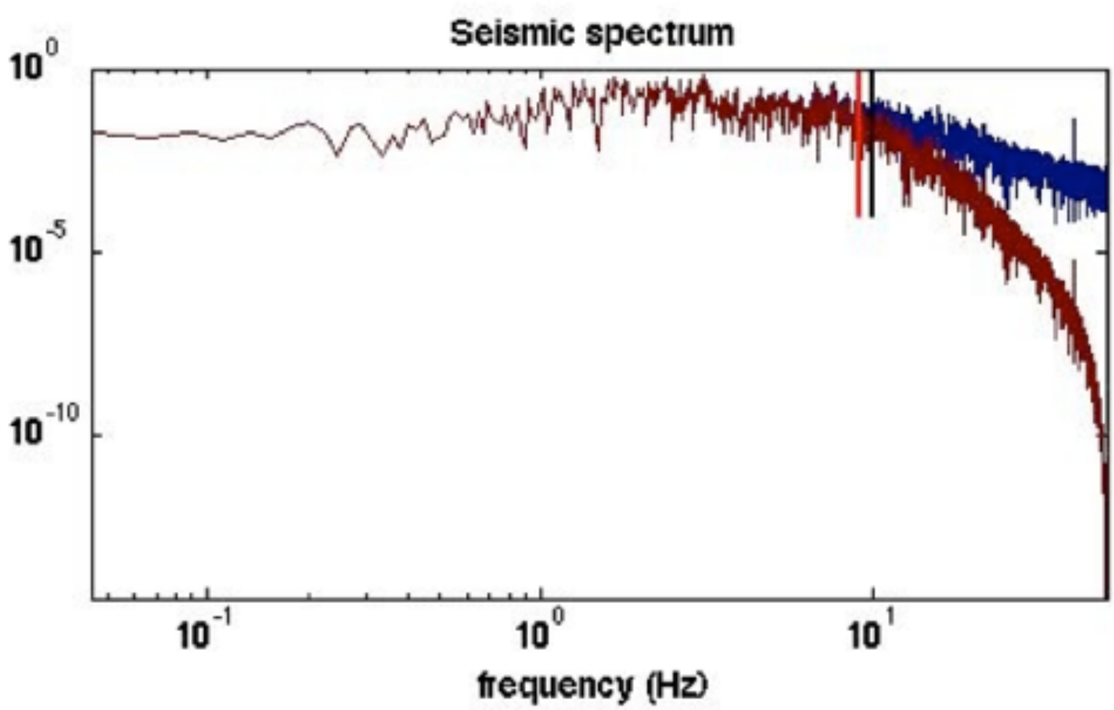

Figure 2. An FFT of a Parkfield earthquake signal (blue) and a final, filtered signal (red).

It's especially powerful and applicable here because it allows us to examine the relative magnitude of the occurrences of different frequencies in our signals. Since many earthquakes have energy across an extremely wide range of frequencies, we often cannot listen to all of it at once. Because of this, we sometimes have to determine which frequency range is most relevant to us by examining the frequency content displayed in the FFTs.

While the FFT is quite useful in that it lets us view the relative magnitudes of different frequencies for a given signal, we still must occasionally iterate on this process after having created sounds. At times, mechanical noise will exist for a given station, and will need to be to be filtered out. For this, and to isolate desired frequency ranges, we used Butterworth filtering. The Butterworth filter is a signal processing filter designed to minimize the frequency response in the band-pass. Ideally, we would not tamper with our data at all, but in some cases it is necessary to apply one or more filters. In most cases that require filtering, all that is needed is a hi-pass or lo-pass filter, but we always have the option of using a band-pass filter or some combination of those three.

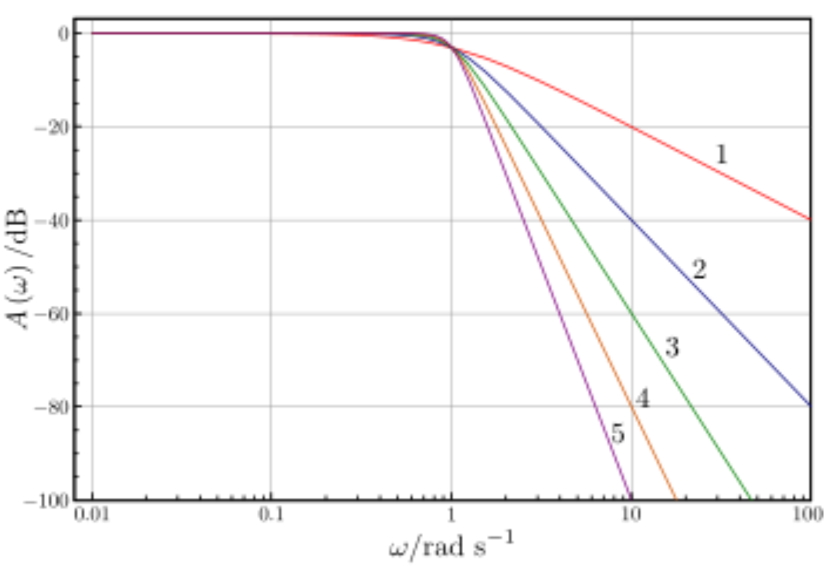

Figure 3. Plot of the gain of Butterworth lo-pass filters for a cutoff frequency of 1. Orders 1 through 5 are shown.

Frequency Shifting

At this point, the desired frequencies are isolated, but still outside the audible range. To shift the filtered signals into the audible range, we utilized a technique called "time compression." If the filtered signal comprises a frequency range from 0.1 Hz to 10 Hz, but for listening purposes we want it to occupy a range from 100 Hz to 10,000 Hz, we must shift it up by a factor of 1,000. Most of our data was initially taken at a sampling frequency of 100 Hz. So to shift the entire frequency range of the signal up by a factor of 1,000, we simply "play back" the data at a sampling rate that is 1,000 times faster than our data sampling frequency. In this example, the sampling rate for our sounds would be 100 Hz x 1,000 = 100,000 Hz. Once the signals are shifted, the actual sounds can be created. We created the sounds in MATLAB, using a function called wavwrite, which simply takes a signal and writes it into a sound file with the filename and signal as parameters. It's not always a given that the first set of filtering and shifting parameters chosen will be the best choices for a particular data set, though. We generated a number of sound files, as well as plots featuring the seismograms and FFTs of the data to attempt to find the best balance of filtering and shifting.

Spatialization and Listening

One key feature of seismic data is that especially in areas with active faults, there are often entire arrays of seismometers in the immediate area surrounding an event. This means that we can obtain data from multiple seismometer stations, and so long as we line up the data so it reflects real time, we can play back the sounds from a number of stations simultaneously. We used the software GeoMapApp to create maps of the

Figure 4. Picture of the 16-speaker array in the Columbia University sound lab.

available seismometers, and used the audio production software Reaper to arrange and spatialize the sounds we created. For cases of two stations, we used standard, stereo speakers. But for cases where we had more stations available to us, we used the 16-speaker array in the sound lab at Columbia University. The spatialization of the sound

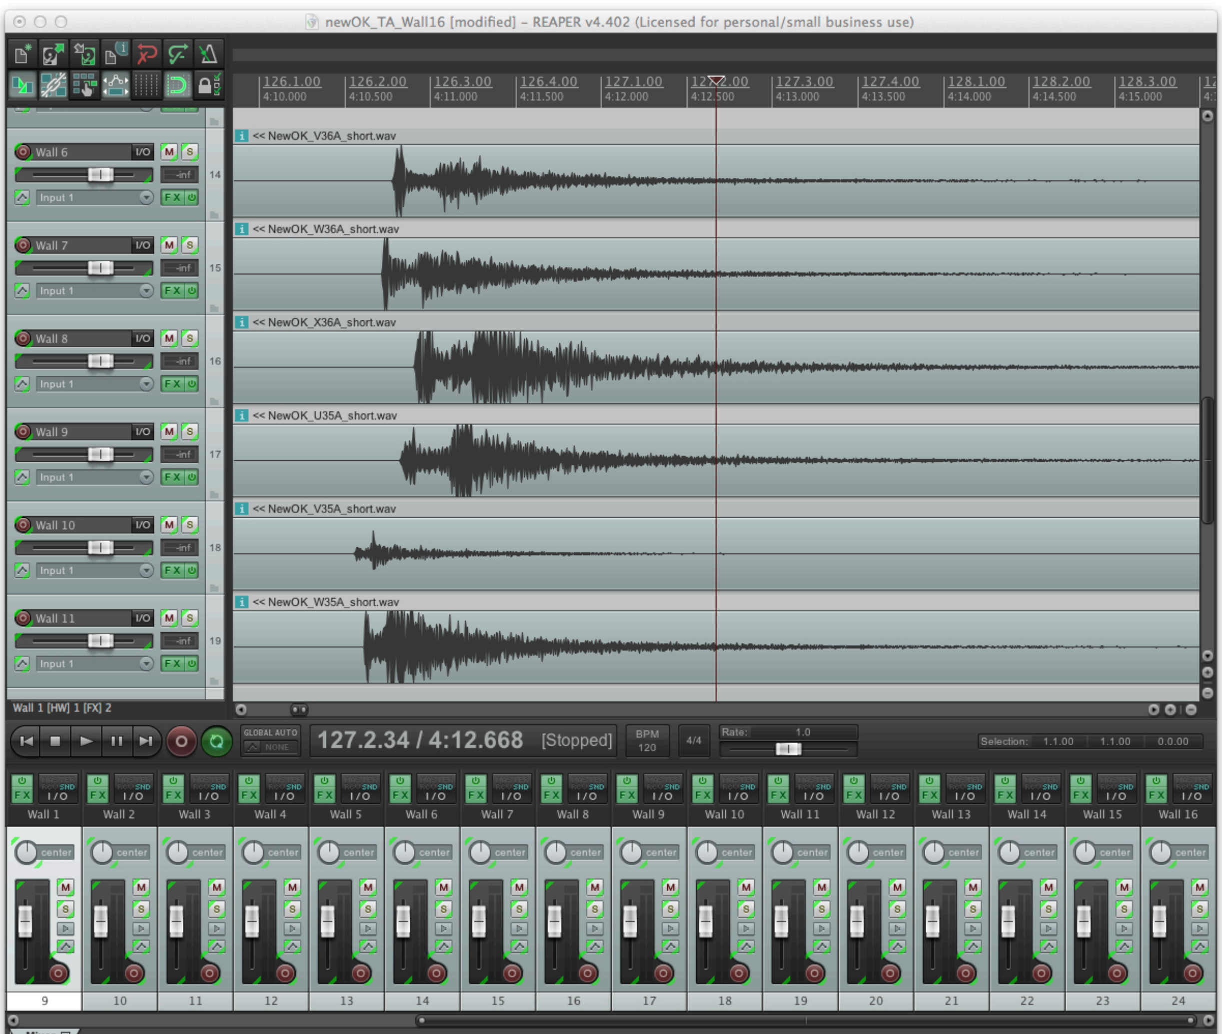

Figure 5. Reaper audio production software, with multi-channel, spatialized sounds.

allows us to create a more accurate sound-based representation of seismic events, and adds a level of nuance to the sound that cannot be gained either by analyzing waveforms visually, or by listening to a single station at a time. Features such a wave arrival time differences, frequency content differences, and attenuation differences are all audible between different stations.

However, this also adds another dimension of potential error in that how we choose stations can drastically affect the sounds. Three primary factors affect the sounds that come out of the speakers in the end: the source, the medium the seismic waves travel through, and the station recording the data. Ideally, we want to isolate the factors caused by the source and medium as much as possible, with emphasis placed on the source in this particular project. In all cases, we want to minimize the impact of the station on what we hear, and this not only means choosing high quality stations (high sampling rate, low noise), but also choosing highly similar stations when spatializing and comparing different events. In a perfect world, we would have an infinitely dense array of perfect seismometers placed around every source, but of course this is impossible and does not even estimate reality particularly well. Since we don’t have that luxury, we must take special care in selecting which seismometers we utilize for a given comparison.

As I stated already, choosing stations that are highly similar in terms of sampling rate and data quality is important, but we can also control other factors, especially in terms of spatialization. In the case where we don't spatialize at all, and only use one station per event when comparing two events, we still must take care to pick stations that are approximately equidistant from their respective sources. Since seismic waves are quite different in the near-field and far-field, we must take distance into consideration. In the case in which we use two or more stations for each event, we must also consider the geometry (or the orientation) of the stations. This is true for similar reasons to the consideration about distance, but it is slightly nuanced. Suppose we have two events, and we'd like to use two stations for each. For Event #1, I take data from two stations that are both 20 km away from the source, and are directly across from each other (they form a straight line segment with the station). For Event #2, I take data from two stations that are also both 20 km from the corresponding source, but they form a 30 degree angle with the source. It should quickly be apparent that we are representing the seismic landscape very differently when we make choices like this, and these two sound sets would likely sound very different when spatialized. To illustrate this further, it should be noted that depending on choice of distance, it's easily possibly that these geometries would sound quite different even if Event #1 and Event #2 were actually the same event. Even with these considerations it is possible that one would encounter some problems due to the orientation of the seismometers chosen (see Appendix C for more detail).

Once the sounds were spatialized, all that was left was listening. While listening may seem like the most straightforward aspect of our methods, it should be noted that the listening occurs over many stages, and requires many repetitions. When initially looking at the frequency content and picking filtering and shifting parameters, we listen to the sounds created to determine whether these are the best parameters. When choosing stations, we must manually listen to many sets of sound data to determine if these station choices are sufficient. Even when we have finished spatializing and are doing analysis on the sounds, repetition is required, along with simultaneous viewing of the seismograms and waveforms of the events. In some cases, a discovery is made very late in the process that requires us to go back and change the filtering and shifting parameters.

Events and Comparisons:

We chose to make paired comparisons between 3 different groups of events. Because we wanted to investigate source processes, we paired up events that we expected to be quite similar, and events we expected to be quite different for different reasons. We obtained data from 3 clusters of repeating events in Parkfield, CA, along the San Andreas Fault, from a magnitude 6 earthquake and a magnitude 4 aftershock from it (also in Parkfield), from a magnitude 4 earthquake in Berkeley, CA, along the Hayward Fault, and from a landslide in the Bingham Canyon Mine in Utah. We compared the magnitude 4 earthquake from Berkeley both to the magnitude 4 aftershock from Parkfield, and the landslide from Utah. We compared each event within a repeating sequence to others in the same set.

Results:

<u>Landslide vs. Earthquake</u>

The primary feature of the landslide as compared with the earthquake was a narrow frequency range. Virtually all of the landslide's energy was between 0.2 Hz and 10 Hz, whereas the earthquake had significant energy in a range of about 0.1 Hz to about 40 Hz. However, despite this large difference in range, both events were weighted more to the lower frequencies, with peaks around 1 or 2 Hz.

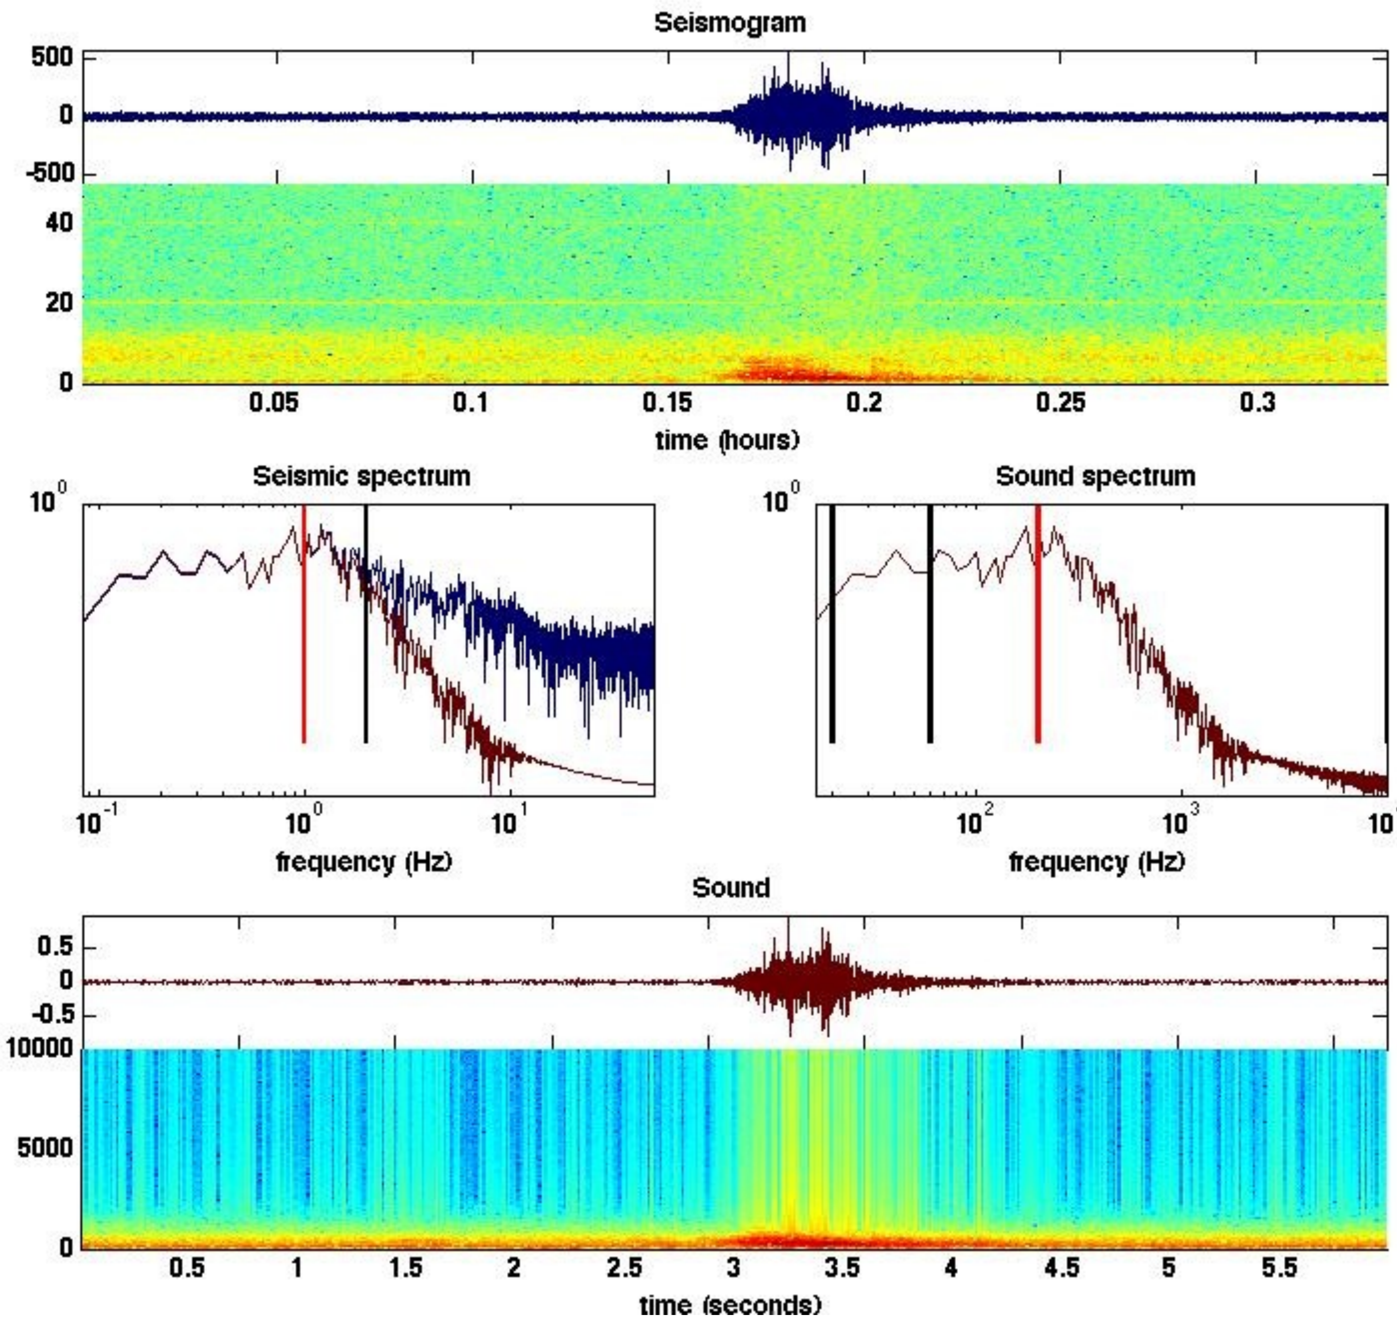

Figure 6. Plots for the Bingham Canyon Mine landslide data. The top plot is a seismogram of the signal, with a spectrogram beneath it. The middle plot on the left is the FFT of the seismic signal and the filtered seismic signal. The middle plot on the right is the FFT of the filtered, shifted signal. A lo-pass filter with a cutoff frequency of 2 Hz was applied, primarily to get rid of higher-frequency noise in the signal. The bottom plot is a seismogram and spectrogram of the sound signal.

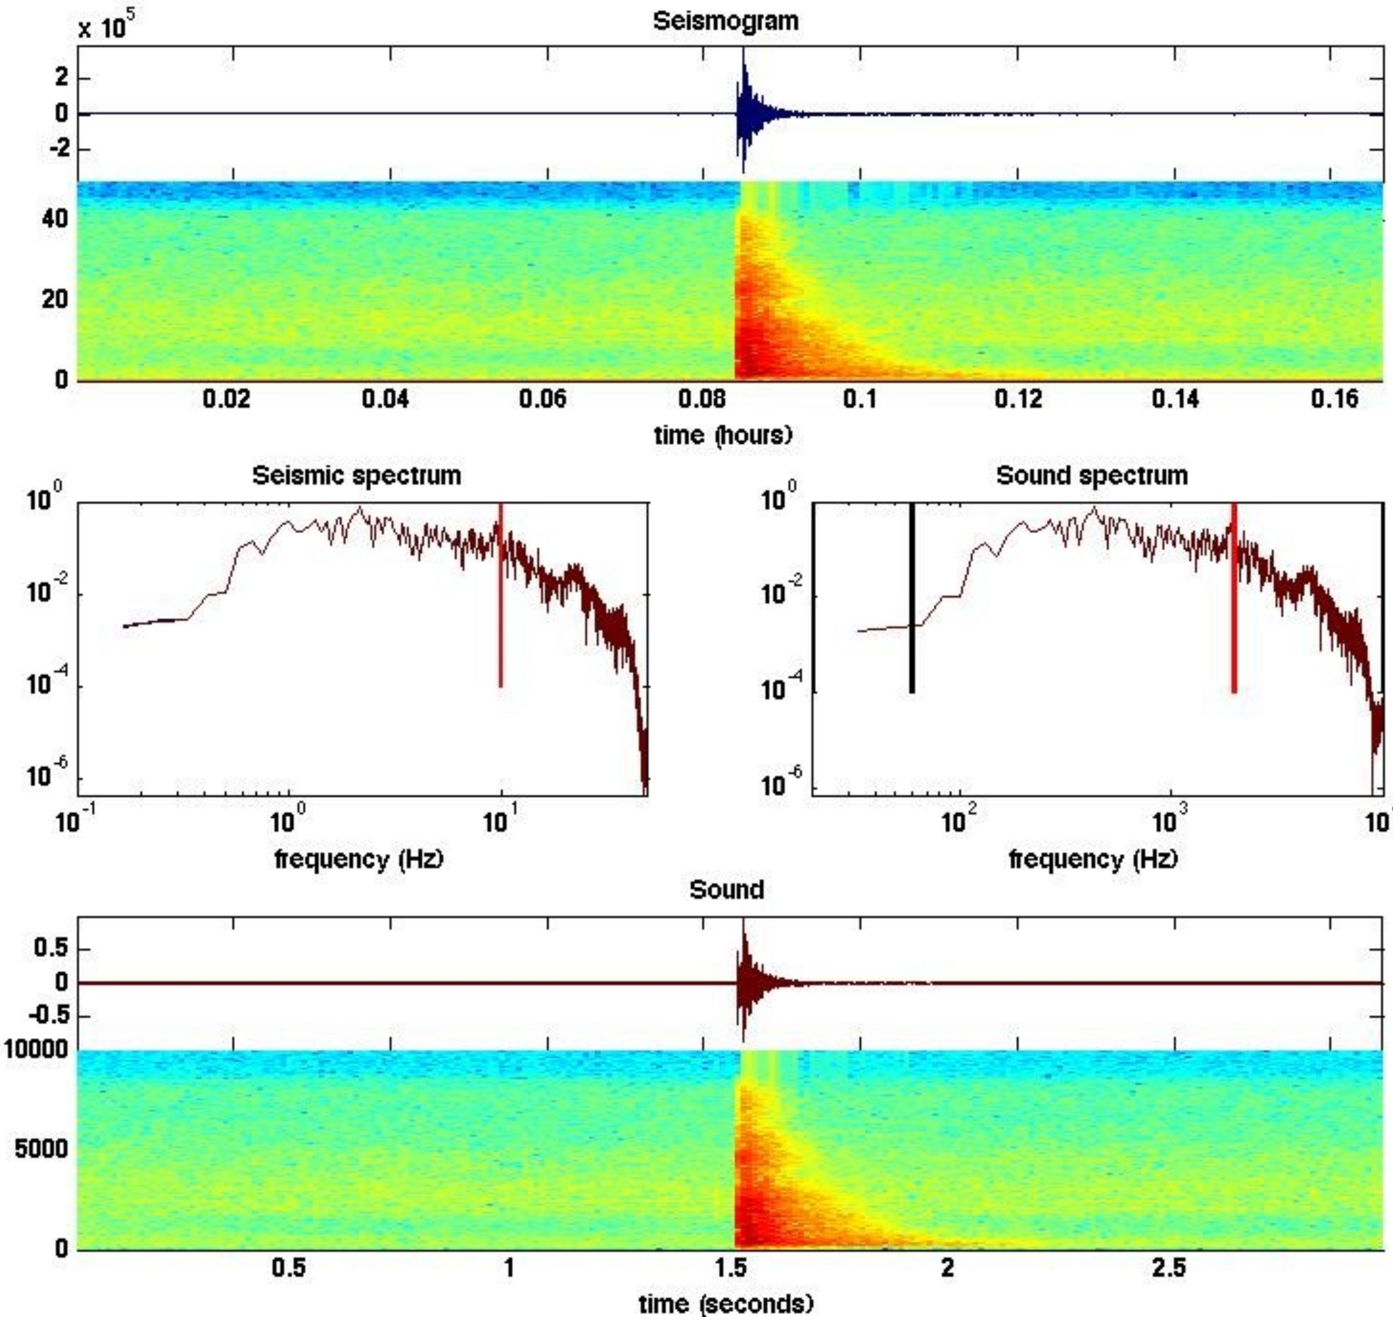

Figure 7. Plots for the M4 data from Berkeley, CA. The top plot is a seismogram of the signal, with a spectrogram beneath it. The middle plot on the left is the FFT of the seismic signal. The middle plot on the right is the FFT of the shifted signal. No filtering was used on this signal, because the data was fairly clean, with very little noise. The bottom plot is a seismogram and spectrogram of the sound signal.

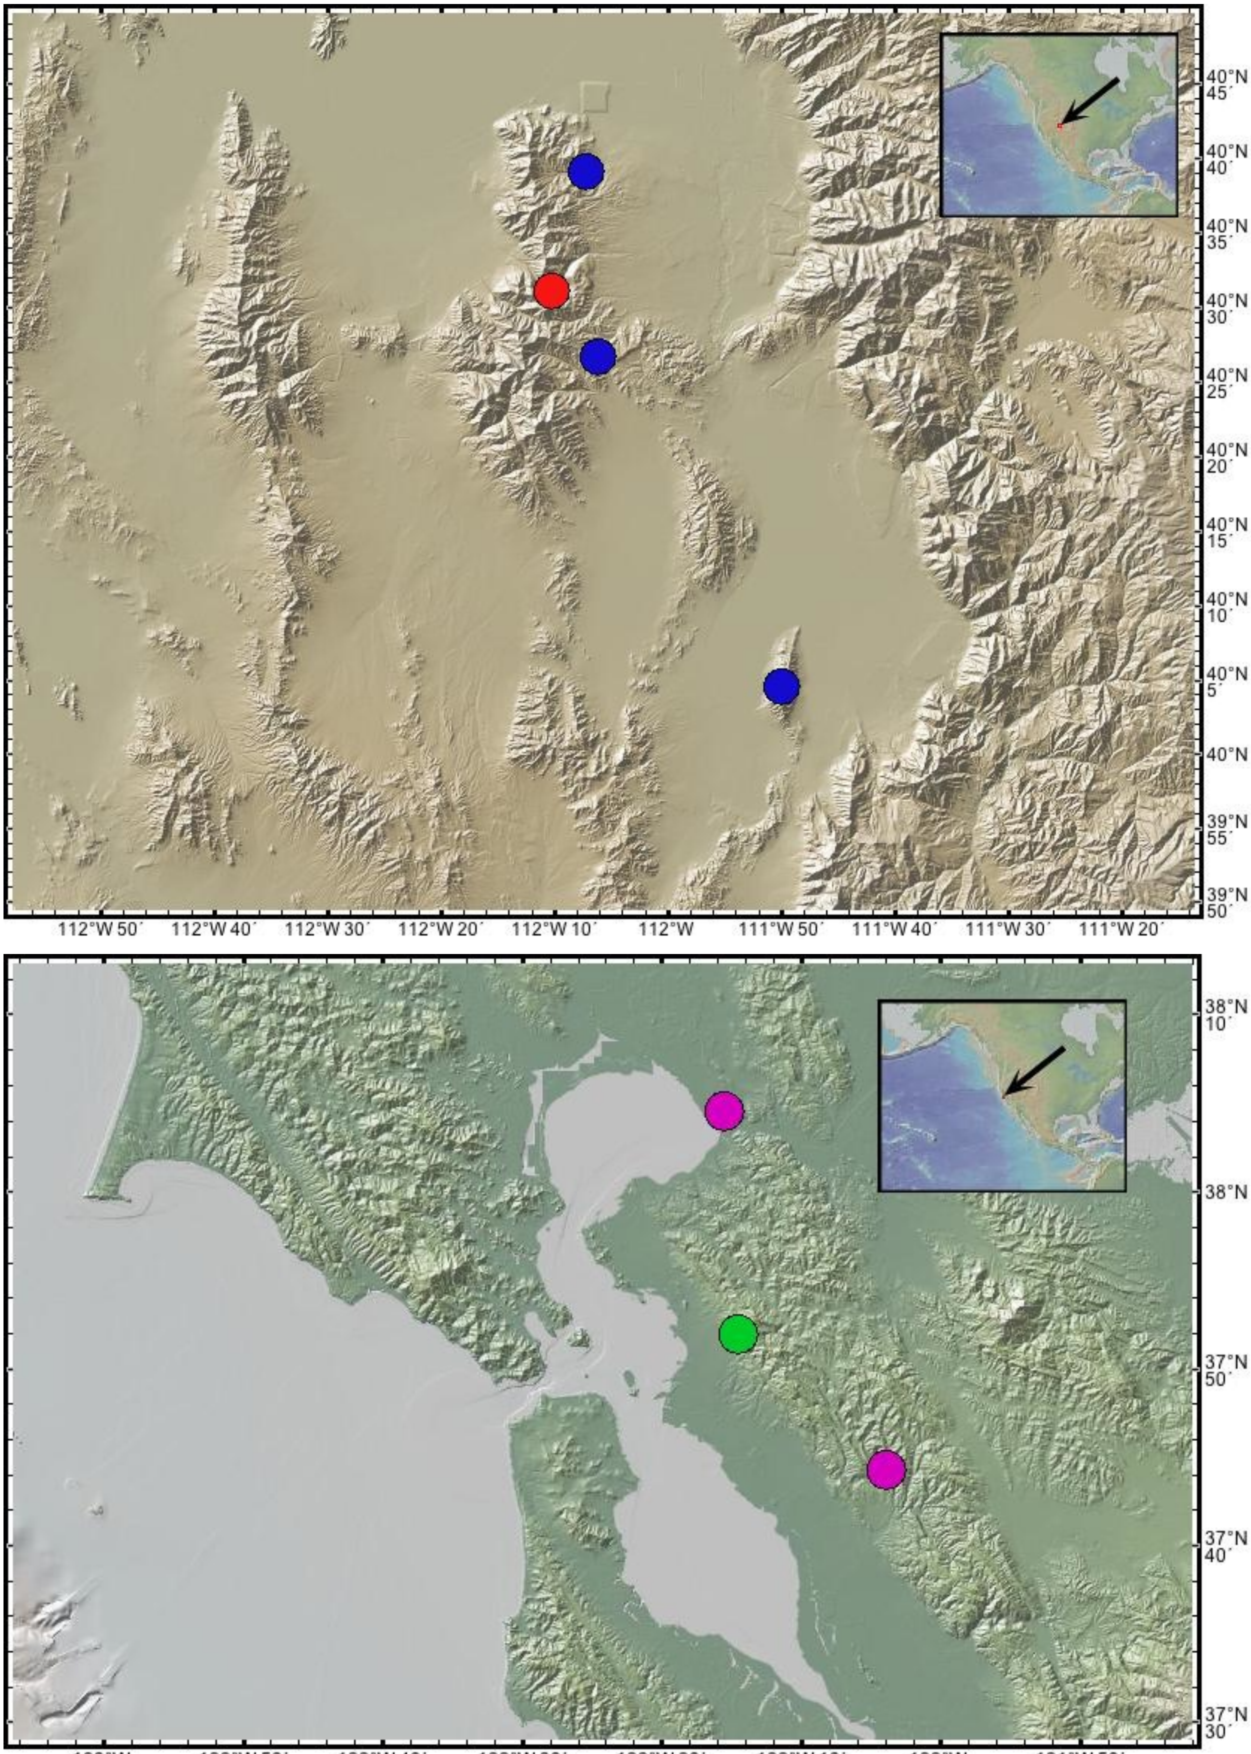

Figure 8. Maps of potential seismometer stations around the (uniquely colored) sources. Bingham Canyon landslide is on top, Berkeley M4 is on bottom. For the landslide, the two closest stations were used for comparison between the two stations used for Berkeley.

In addition to the frequency content differences, the landslide sounded different in that the entire process took longer — but the onset and attenuation were fundamentally different as well:

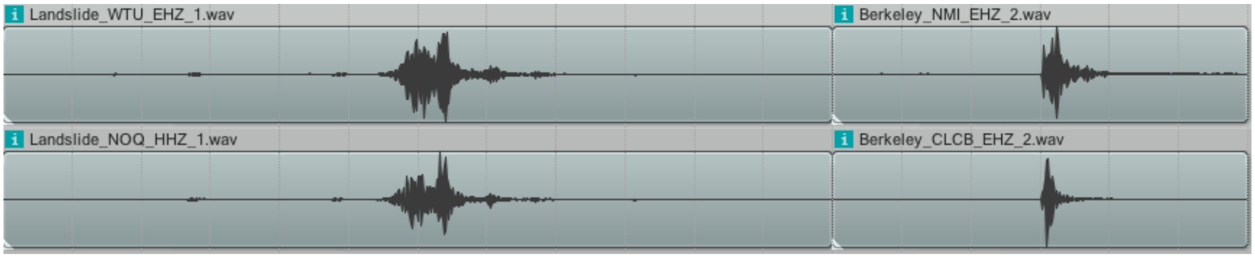

Figure 9. Waveforms for the Bingham Canyon landslide (left) and Berkeley M4 earthquake (right). Both were shifted by a factor of 200.

In a landslide, the “source” overlaps with the decay, causing a very different attenuation, most noticeable in the fact that the landslide “swells” up in volume twice. First, as it begins, and then again in the middle, before fading away. The earthquake has a clear impulse and attenuation. A landslide is also a fundamentally different process from an earthquake in that the “source” just isn’t the same. An earthquake has a well-defined source location, and we can even discuss rupture propagation. In a landslide there is no rupture at all, and the location of the “source” is more ambiguous.

Two Faults

The comparison between the Berkeley M4 and the Parkfield M4 aftershock led to the most frustration, the most profound confusion, and the least useful results. It was quickly clear that almost all of the stations we obtained data for from the Parkfield M4 aftershock were clipped. Clipping occurs when the seismometer is overdriven and becomes “saturated.” Essentially, this meant that the stations were too close to the source. Initially we thought we were going to have use different data entirely, but we ran

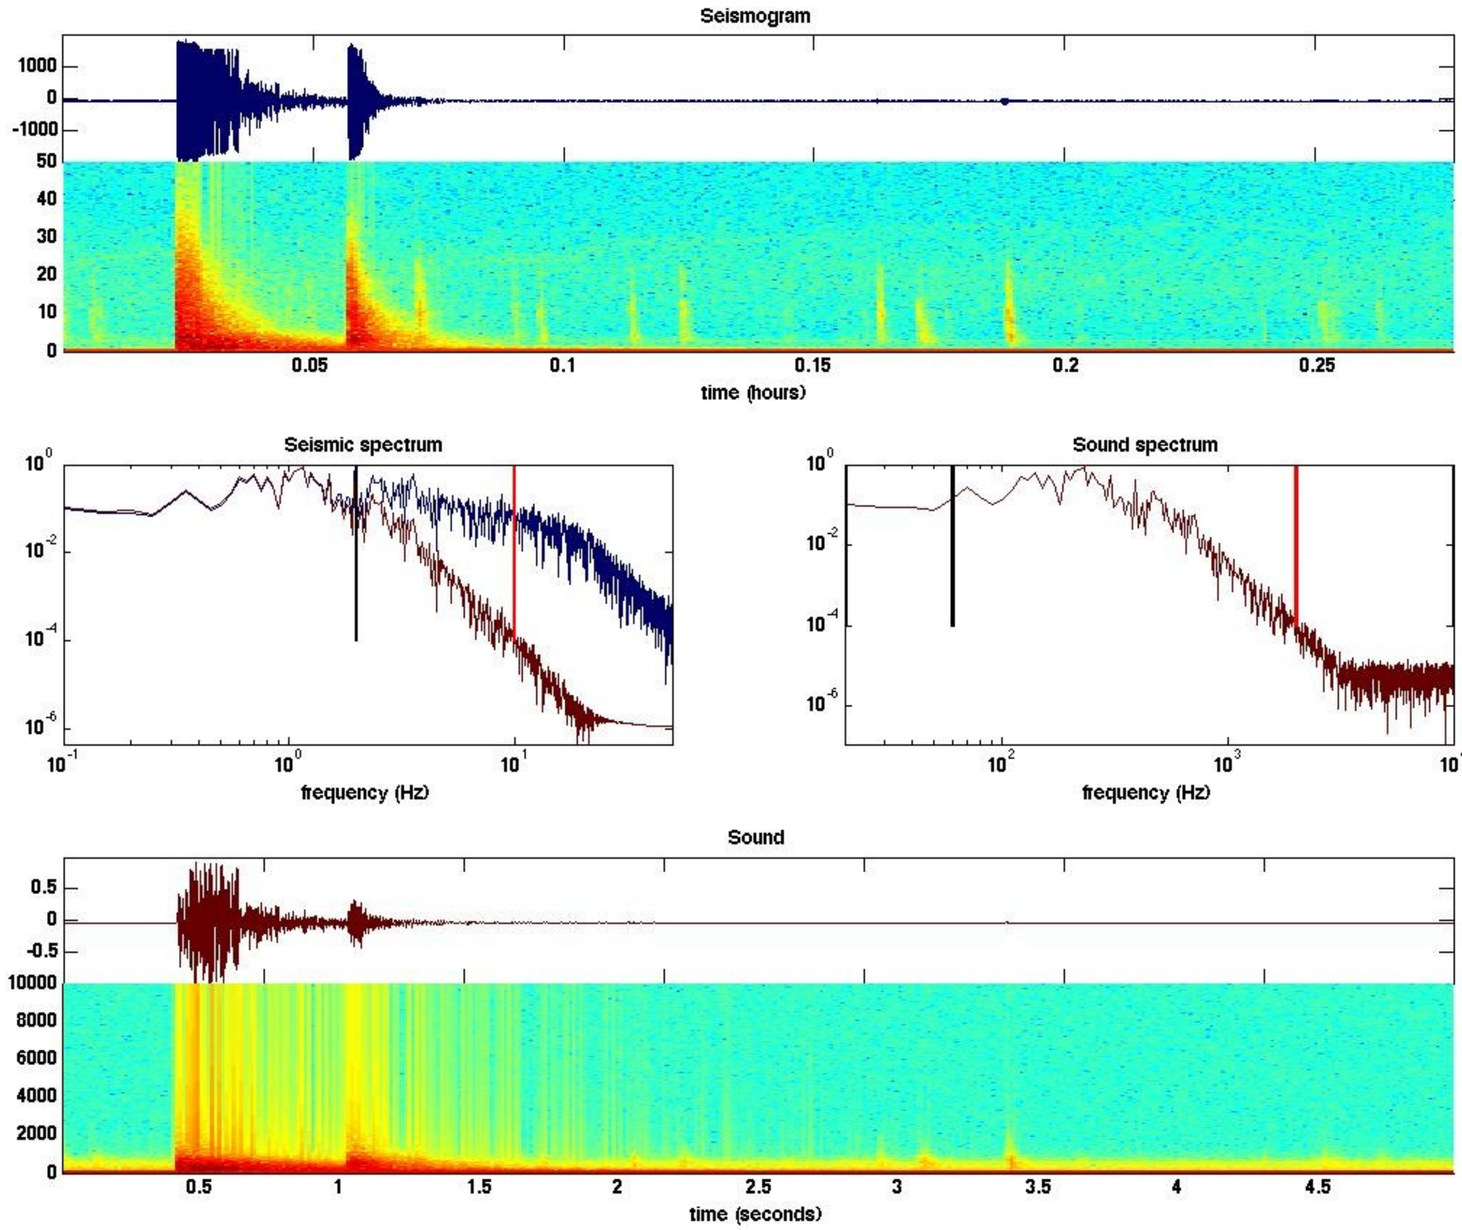

Figure 10. Plots for the M4 aftershock data from Parkfield, CA. The top plot is a seismogram of the signal, with a spectrogram beneath it. The middle plot on the left is the FFT of the seismic signal and the filtered seismic signal. The middle plot on the right is the FFT of the filtered, shifted signal. The data was shifted by a factor of 200, and a lo-pass filter with a cutoff frequency of 2 Hz was used since at the time we were using these filtering and shifting parameters for all the other data (aside from the repeating events). The bottom plot is a seismogram and spectrogram of the sound signal.

into the same clipping problem with some of our other data as well, and it appeared that one of the stations from the Parkfield M4 data didn't have clipping like the others. We decided to use this station to compare to a single station from the Berkeley M4. While

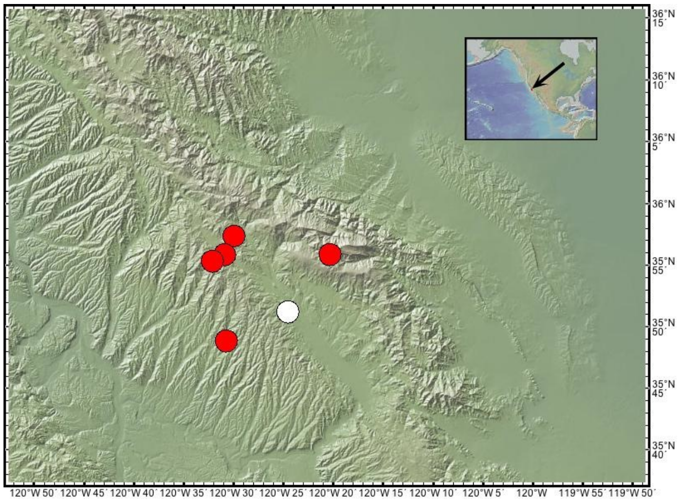

Figure 11. Map of potential seismometer stations (red) around the Parkfield M4 aftershock source location (white). See Figure 7 (bottom) for Berkeley station map.

not as powerful as creating stereo or multi-channel sounds, single-station data can still give us a lot of information. Initial listen-throughs of the data gave encouraging results. The Parkfield M4 data had interesting frequency content differences in comparison to the Berkeley M4. The Parkfield earthquake was more weighted to the lower frequencies during the impulse, while the Berkeley earthquake had a more uniform distribution in its impulse. The codas differed greatly too, with Parkfield's coda being both shorter, and fading to the lower frequencies more quickly. In the Berkeley earthquake, the coda was longer, and the middling and higher frequencies were audible for significantly longer.

We realized later that the single station that we chose (that had appeared to be free from clipping) was not as suitable for use as we thought. Due to some last minute questions regarding our filtering parameters, we chose to make new sound sets for the Parkfield aftershock and the Berkeley M4 with the same 200x shifting factor, but no filtering. The result was shocking: the unfiltered data from the *same station* we had chosen before was now hopelessly clipped. Over three quarters of the sound length were virtually useless, with a high level of distortion present. What was shocking was that judging by the FFT (Figure 9), we wouldn't have thought we were cutting out much of frequency content by applying the lo-pass filter.

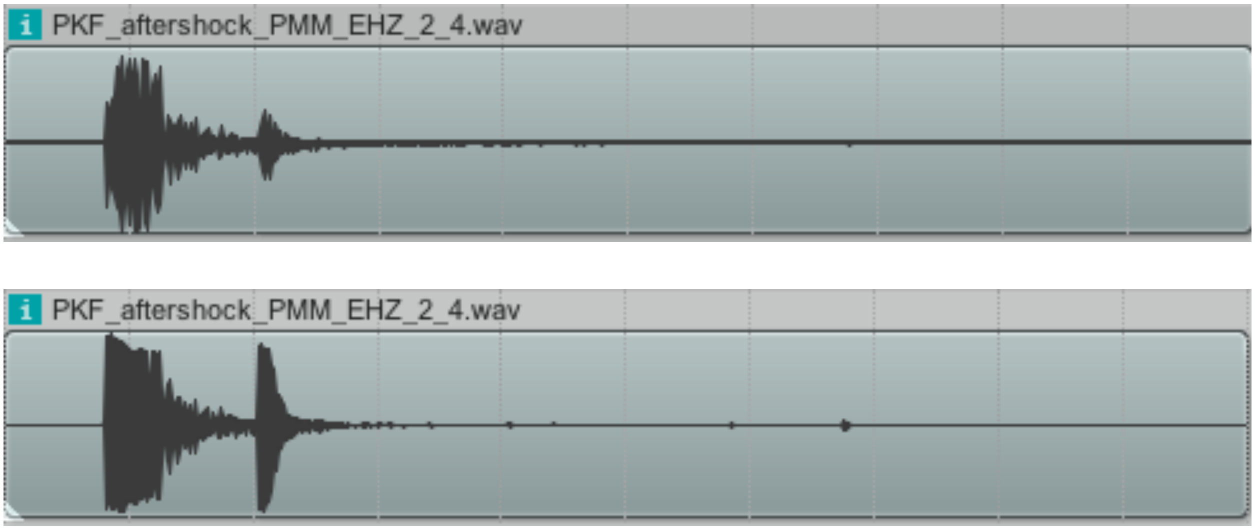

Figure 12. Waveforms for the Parkfield M4 aftershock data. The top is shifted by 200, with a lo-pass filter with a cutoff frequency of 2, while the bottom is shifted by 200, but is unfiltered. The filtered waveform, while it doesn't have a completely perfect shape, doesn't appear overly clipped. The unfiltered waveform on the other hand has the stereotypical square shape of a clipped signal.

Yet it was right in front of us — the unfiltered signal was very clearly clipped. It was apparent both in the waveform, and even more so in the sound (the clipped sound is heavily distorted, and doesn't sound anything like an earthquake normally would).

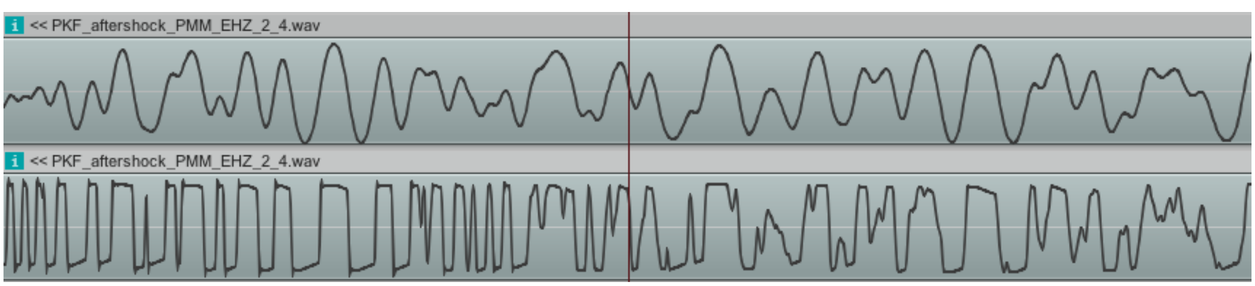

Figure 13. Same waveforms as in figure 11 (above). Zoomed in such that the "square-wave" style of the clipping is quite apparent in the unfiltered (bottom) waveform. There is still no visibly (or audibly) obvious clipping in the filtered (top) signal.

Even knowing this, it didn't seem like the filtered signal was clipped at all — or at least nowhere near as badly as the unfiltered signal. It was unclear whether this would have any effect on how relevant our results from that station would be though, so we must state that all the results from this comparison must be taken with a grain (or a few cupfuls) of salt.

Repeaters

The Parkfield repeaters were the biggest point of interest coming into this project, since these were the ones with the most definitive traditional analyses. According to standard analysis techniques, these earthquakes are identical for all intents and purposes. Their locations, magnitudes, frequency makeups, durations, and overall waveform characteristics are extremely similar. While the repeating events were

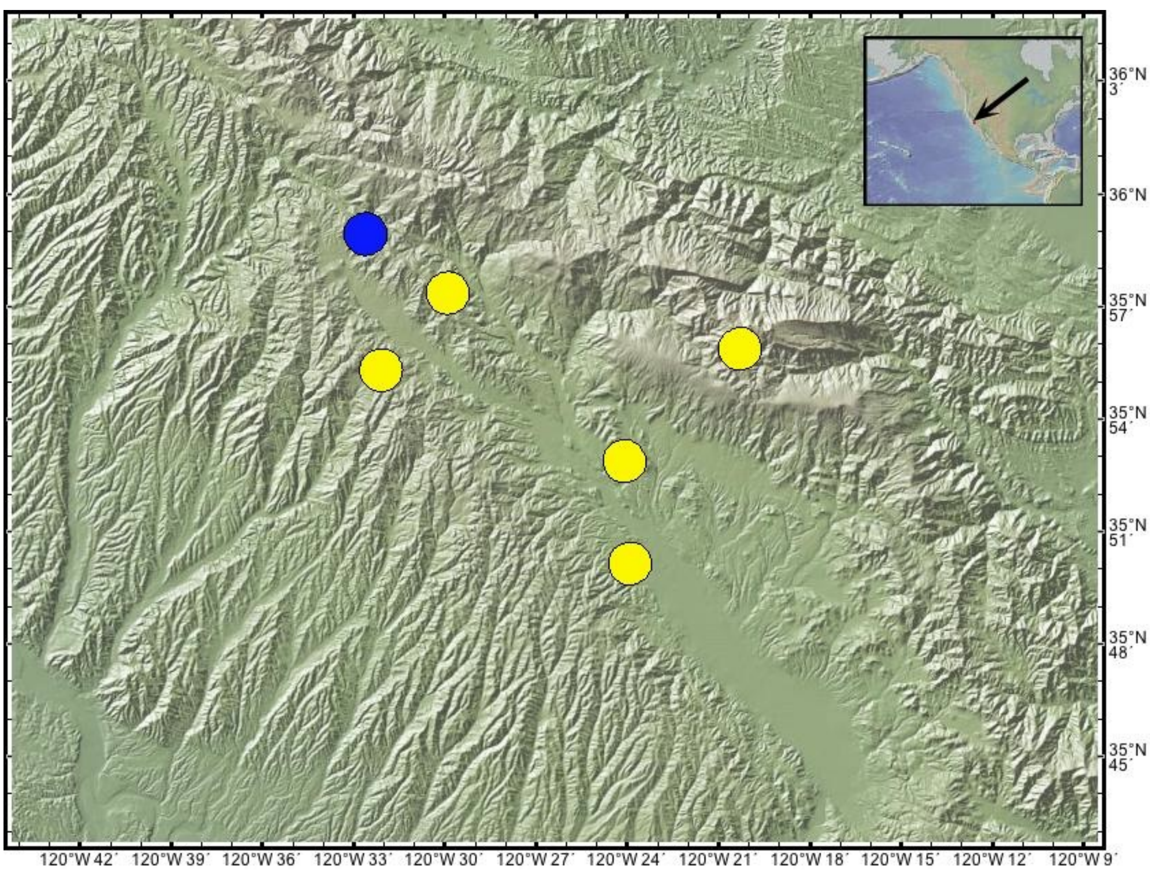

Figure 14. Map of the available stations (yellow) and the source (blue) for the Parkfield repeating events.

probably the most striking to listen to, there is less to say about them because of how remarkably consistent they were. The sounds were uncannily similar in every respect: frequency content and distribution through time, duration, and characteristics of the spatialization. However, there was an unexpected feature to the sounds. There was a sense of coherent motion across the array, that was far too slow and pronounced to be accounted for by difference in wave arrival time. Wave arrival time differences, while audible in a 200x shifted signal, are not large enough to account for the relatively slow-and-steady motion heard across the array. Additionally, the motion was too smooth and uniform to be a result of reflections or reverberations. While disconcerting, this sensation of motion across the array was also remarkably similar between events, making it all the more conspicuous — without detracting from the consistency among all the events.

Discussion and Questions:

It's challenging to discuss the differences between an earthquake and a landslide, because the sounds reflect the reality: they are entirely different types of processes. The landslide has swells and a flow to it that just isn't present in an earthquake. The earthquake sounds more like a single clap of thunder; an impact followed by a train of reverberation. The point about a landslide not having a rupture is relevant as well — it's difficult to compare two events that are so different at their core. In addition to the question of landslide magnitudes (see Appendix D), the most important question that comes out of these results is: how can we conceptualize comparisons between these two phenomena? In particular, how can we conceptualize the idea of a “source” for a landslide, when it seems like any description would be inherently ambiguous?

For the comparison of the San Andreas Fault and the Hayward Fault M4 earthquakes, it’s unclear whether we can really utilize our findings. When we used the filtered signal from the Parkfield aftershock on the San Andreas Fault, it seemed that we were picking up some relevant and encouraging differences. For one thing, it is thought that the San Andreas Fault tends to build stress and then have a stress drop “all at once,” where most of the energy is released in a single earthquake, and then the fault stops shaking. The Hayward Fault is considered to be “looser,” such that when an earthquake occurs, there is more extraneous shaking, triggered earthquakes, and other phenomena. This seemed to be reflected in the attenuation differences of the two events. The frequency content differences were not expected per se, and could have opened a door to some discussion. However, the presence of clipping in the higher frequencies draws into question the validity of such a discussion.

At the very least, this comparison brings up some questions regarding the use of sound in seismology. First, to what extent is clipping a factor in modern seismometers, and what might be some methods for combatting it? This would allow a better use of seismometers in the near-field, especially for larger earthquakes. Second, why were we able to “filter out” the clipping? This makes sense on one level: clipping causes the seismometer to sustain its maximum power output for some amount of time, which would severely distort the higher frequencies. However, what does this really mean for the lower frequencies? Are they mostly unaffected, as it appeared in the waveform and in the sound? If so, this would imply that we could use clipped sounds so long as we were only interested in the lower frequencies, but it also seems too easy. It seems unlikely that the lower frequencies would be completely unaffected, and the effects would obviously be unpredictable (given that we didn't really notice them at first for the single seismometer). Lastly, how does over-driving a seismometer affect its long-term usability as a seismic measuring tool? It seems plausible that some damage could come to the devices, if not after one such occasion, then after many. Since data quality is such a concern, and stations must often be rejected due to bad quality, this should be an on-going consideration.

The repeating events had a lot to offer, especially since they were the only comparison in this project that we were able to spatialize more fully than in stereo. We used up to 5 stations at one time, which added a smoothness and coherence to the data that was not present in the other two comparisons made for this project. While we expected that these events would match each other quite well, it was still astounding to actually hear. Without watching the waveforms during playback, it was impossible to tell which event was being played. This was especially impressive given the presence of a slow sense motion across the array. We already established in the Results section that the motion cannot be accounted for by differences in wave arrival time or by the presence of reflections. The motion was too slow, and too obviously two-dimensional across the plane of the arrays. This begs the question, what caused this perception? More generally, what leads to a perception of motion when listening to spatialized sound, and how does this perception line up with the physical source processes? In what ways do the choices surrounding accurate station-speaker symmetrical geometries affect these perceptions?

The final clear-cut, if straightforward, conclusion that came out of this project was that the spatialization of the sounds is an incredibly powerful factor. The comparison between the Bingham Canyon landslide and the Berkeley M4 earthquake was undoubtedly interesting and eye-opening in a number of ways, but the Parkfield repeaters data allowed a much deeper view into the events. The difference between listening to a single station and listening to stations in two dimensions is akin to the difference between seeing in black-and-white and seeing in color. The texture it adds, along with the sheer amount of data that can be analyzed at once is simply staggering. Going forward, maintaining a focus on the spatialization will be crucial in using sound to analyze source processes and everything else in seismology.

Appendix:

A. Magnitude Scaling and Normalization

An issue beyond the scope of this project is that of magnitude range. Earthquakes have a wide range of possible magnitudes, translating to an enormous variation in energy release. A difference of magnitude of 1 means a 10 times difference in shaking amplitude, or approximately a 32 times difference in energy release. What this means is that listening to earthquakes of different magnitudes in the same file is quite difficult, because the decibel range that humans can hear is fairly limited in relation.

Partly because of this, we must normalize the amplitudes of all data used. When this happens, we set the maximum amplitude moment in the sound (corresponding to the largest magnitude event) to a set volume in the sound. This means that every sound file has the same maximum volume, so that when we play multiple sound files at once, the scaling is correct within one file, but not between files (unless they had the same largest amplitude).

One potential solution involves concatenation of all signals that we wish to listen to at one time, and normalizing the resulting signal so that everything is normalized in the same way. Then we could "chop up" the sounds to allow us to rearrange them as usual. We would then have to utilize compression techniques to shrink the decibel range of the sounds.

B. Compression

Somewhat related to the issue of magnitudes and amplitude normalization (Appendix A) is the audio tool of compression. Compression is an audio engineering effect that allows us to increase the volume of quiet sounds and decrease the volume of loud sounds so that they occupy a smaller decibel range. While we did not have the time and tools to implement compression for this project, it is highly recommended, since earthquakes have such a wide range on the energy spectrum, and often many of the interesting details in the sound happen in the lower-decibel range. By utilizing compression, one can examine wider decibel range, and in particular bring out details of the attenuation that otherwise would be lost due to the normalization to a high-energy impulse.

At the same time, however, it is unclear whether this removes some level of “authenticity” in representing the earthquake. It’s true that finding detail in the sound is important, and perhaps being able to hear on a log-scale of decibels would be most useful for examining earthquakes, but it’s unclear whether compression creates a sound environment that is more realistic or less so.

C. Spatialization and Seismometer Geometry

Picking station geometry can be challenging, especially when the arrays we are utilizing are not always particularly dense. There are also a number of ways in which two spatialized sets of stations chosen for two different events could appear very valid, but leave you with misleading results. Let’s say you have stations that are only on the north west side of a source, and have the same hold true for the other source, with the same distances and relative locations. While the visual appearance would suggest that they are exactly comparable, it’s possible that the focal mechanisms of the two events would make it such that the arrays were in different orientations relative to the focal mechanisms. Similarly, since the fault plane and the array plane actually exist in three dimensions, and are not necessarily perpendicular to each other, the same problem could theoretically exist even with a symmetrical, fairly dense array that surrounded each source at every angle. This is why having as-dense an array as possible is ideal, as you can eliminate most of the concerns regarding two-dimensional rotations of a focal mechanism. This also suggests that it will usually be strictly better to have an array that surrounds a source, rather than one that occupies one side or a smaller angular chunk.

D. Landslide Magnitudes

When we compared the earthquakes from Berkeley and Parkfield, we deliberately chose events of similar magnitude. The reasoning behind this is that we want to control as many of the factors that we're not specifically analyzing, and make them close to identical between the two events. However, for a landslide, "controlling" the magnitude is tricky. For an earthquake, the magnitude calculation is based primarily on the maximum surface-wave amplitude present in the signal. Due to the the mechanics of a landslide, this maximum surface-wave amplitude may be unrealistically large, and over-represent its size. For this reason, our choice of which earthquake to compare to the landslide was somewhat arbitrary, and any future comparisons of this nature should involve an attempt at creating a less arbitrary method of event-pairing.

E. Signal Length

One issue that inevitably came up in every comparison in this project was a question of signal/sound length. The most notable change in our data after all manipulation was complete (aside from the frequency changes) was the time length. Because we must change the sampling rate by the same factor that we shift the frequency range, we often end up with sounds that are very short — often on the order of a few seconds long for individual events — and therefore difficult to analyze using human audio perception. The nuances in the sound can be difficult to pick up and articulate even when the sounds are not particularly complex. When there are many layered features to the sound though (like in a multi-channel, spatialized case), it becomes an increasing struggle to keep track of everything, even with multiple repetitions. Ideally, one would lengthen these types of signals using signal stretching, and without misrepresenting the signal itself.

References:

Holtzman, B., Candler, J., Turk, M., 2013, Seismic Sound Lab: Sights, Sounds and Perception of the Earth as an Acoustic Space, accepted to Lecture Notes in Computer Science: Proceedings of the Computer Music Multidisciplinary Research conference, Marseille, 2013.

Peng, Z., Aiken, C., Kilb, D., Shelly, D.R., Enescu, B., 2012, Listening to the 2011 Magnitude 9.0 Tohoku-Oki, Japan, Earthquake, *Seismological Research Letters*, *Vol. 83 No. 2,* p. 287-293.

Benioff, H., Ewing, M., Press, F., 1953, Sound Waves in the Atmosphere Generated by a Small Earthquake, *Proceedings of the National Academy of Sciences of the United States of America v.37(9),* p.600-603.

Waldhauser, F. "Streaks, Multiplets, and Holes: High-resolution Spatio-temporal Behavior of Parkfield Seismicity." Geophysical Research Letters 31.18 (2004): n. pag. Print.

Beeler, N. M. "A Simple Stick-Slip and Creep-Slip Model for Repeating Earthquakes and Its Implication for Microearthquakes at Parkfield." Bulletin of the Seismological Society of America 91.6 (2001): 1797-804. Print.

Vidale, J. E., W. L. Elisworth, A. Cole, and C. Marone. "Variations in Rupture Process with Recurrence Interval in a Repeated Small Earthquake." Nature 368.6472 (1994): 624-26. Print.

Sammis, C. G. "Repeating Earthquakes as Low-Stress-Drop Events at a Border between Locked and Creeping Fault Patches." Bulletin of the Seismological Society of America 91.3 (2001): 532-37. Print.

Acknowledgements:

I would like to thank Professors Ken Singer and Ken Kowalski for their diligence in running Senior Project Seminar. I would also like to thank Professor Harsh Mathur for his insightful questions and encouraging interest in this project. His comments always helped me see my project from an outside perspective, and consider what kind of an impact my work could have beyond Senior Project.

I would like to thank each of my advisors — Ben Holtzman, Felix Waldhauser, and Lapo Boschi — for their enthusiasm and support throughout the ups and downs of this project. I would like to thank Felix Waldhauser in particular for his invaluable data contributions, and his unworried thoughts and suggestions when things went wrong.

I would especially like to thank Ben Holtzman for teaching me, guiding me, and forcing me to think both more analytically and creatively than I otherwise could. Without him, this project would not have been possible.