

# Seismic Sound Lab: Sights, Sounds and Perception of the Earth as an Acoustic Space

Benjamin Holtzman<sup>\*1</sup>, Jason Candler<sup>2</sup>, Matthew Turk<sup>3</sup>, and Daniel Peter<sup>4</sup>

<sup>1</sup> Lamont Doherty Earth Observatory, Columbia University

<sup>2</sup> Tisch School of the Arts, New York University

<sup>3</sup> Department of Astronomy and Astrophysics, Columbia University

<sup>4</sup> Institute of Geophysics, ETH (Swiss Federal Institute of Technology) Zurich

\*benh@ldeo.columbia.edu, <http://www.seismicsoundlab.org>

**Abstract.** We construct a representation of earthquakes and global seismic waves through sound and animated images. The seismic wave field is the ensemble of elastic waves that propagate through the planet after an earthquake, emanating from the rupture on the fault. The sounds are made by time compression (i.e. speeding up) of seismic data with minimal additional processing. The animated images are renderings of numerical simulations of seismic wave propagation in the globe. Synchronized sounds and images reveal complex patterns and illustrate numerous aspects of the seismic wave field. These movies represent phenomena occurring far from the time and length scales normally accessible to us, creating a profound experience for the observer. The multi-sensory perception of these complex phenomena may also bring new insights to researchers.

**Keywords:** Audification, Sonification, Seismology, Wave field visualization

## 1 Aims

An earthquake is a minute event in the vast and slow movements of plate tectonics; it is the smallest increment of plate motion that we can experience with our unenhanced senses. Over geologic time, earthquakes are tiny, discrete events that constitute smooth, slow motion of plates. An earthquake is a rapid (seconds to minutes) release of elastic potential energy that accumulates in the two plates separated by a fault. When the fault ruptures, some fraction of that energy excites elastic waves, referred to as the seismic source. The seismic wave field is the ensemble of elastic waves that propagate through the planet after an earthquake. Only near the source can we directly feel these waves, but the wave fields from sources with magnitudes above about 4.5 can be measured globally by seismometers. The resulting seismograms are the raw data used to study both the internal structure of the planet and the nature (magnitude and orientation) of the fault rupture. The aim of this project is to develop methods for making this global wave field perceptible, as illustrated in Fig. 1.

Here, we construct a sensory experience of the global seismic wave field by making a vast shift in its time and length scales. The most basic element is to create sounds from seismograms. We shift the frequencies into an audible range by time compression of the data (as illustrated in Figure 2 and discussed in the Appendix), referred to as “audification” or “sonification” (Walker and Nees, 2006). Sounds from several seismic data sets simultaneously recorded at different locations on the globe play through speakers in the same relative positions as the seismometers to produce spatialized sound. We then produce 4D renderings of the seismic wave field from global simulations and synchronize the sound with the movies, which are played in audio-visual environments in which the listener is situated “inside” the Earth. The sounds are unmistakably natural in their richness and complexity; they are artificial in that they are generated by a simple transformation in time. We are not trying to simulate the experience of being in an earthquake; rather, as observers seek meaning in the sounds and images, they grapple with the magnitude of this shift. In a subliminal way, these sounds can bring people to realize how fragile and transient is our existence, as well as a fundamental curiosity about the planet.

Humans perceive a great deal of physical meaning through sound— for example, the motion of approaching objects, the mood of a voice, the physical character of material breaking (distinguishing snapping wood from shattering glass), the nature of motion on a frictional interface (roughness of the surface and the speed of sliding). Also, perception of motion through sound often triggers our mind to look for visual cues. This perception of motion is very sensitive to frequency (Hartmann, 1999). While our hearing has better temporal resolution, our eyes are very sensitive to spatial gradients in color from which we decipher spatial information, texture, etc. Due to these differences, “multi-sensory integration” is much more powerful than sound or image alone in eliciting in the observer a range of responses, from the instinctual to the rational.

After some historical background of signification of seismic data in Section 2, we describe how we produce sounds and images, in Sections 3 and 4 respectively. In Section 5, we discuss how we synchronize sounds and images to illustrate one fundamental aspect of the physics of the seismic wave field in the Earth. In Section 6, we elaborate on the pedagogical questions, approaches and potential. Multi-sensory perception of wave fields in the Earth provokes a wide range of questions in the observer; those questions depend on their experience. The signals in seismic waves measured at any seismometer depend on the nature of the earthquake source, the distance from the source and the structure of the Earth; questions provoked in the listener reflect any or all of these aspects.

## 2 Background and Development of the Project

The earliest example (to our knowledge) of earthquake sonification is a record from 1953, called “Out of this World” (Road Recordings, Cook Laboratories), recorded by Prof. Hugo Benioff of Caltech, brought to our attention by Douglas Kahn (Kahn, 2013). He ran seismic data from 2” magnetic tape directly to vinyl

print, accelerating the tape playback so that the true frequencies of the seismic signal would be shifted into our range of hearing. A concise description of the frequency shifting process in an analog system is found in the liner notes: *“It is as though we were to place the needle of a phonograph cartridge in contact with the bedrock of the earth in Pasadena... and we “listen” to the movement of the Earth beneath a stable mass, or pendulum, which is the seismometer... But even during nearby earthquakes, the rate of motion is so slow in terms of cycles per second that the taped signals cannot be heard as sound... It is somewhat as though we played a 33-1/3 rpm record at 12,000 rpm,”* Sheridan Speeth at Bell Labs used this technique to distinguish between seismic signals from bomb tests and earthquakes, in the early days of nuclear test monitoring (Speeth, 1961; Kahn, 2013). Based on his experiments with sonification of active-source seismic data (explosions, not earthquakes as the source of elastic waves) in the 1970’s, David Simpson wrote *“Broadband recordings of earthquakes...provide a rich variety of waveforms and spectral content: from the short duration, high frequency and impulsive ‘pops’; to the extended and highly dispersed rumbles of distant teleseisms; to the rhythmic beat of volcanic tremor; to the continuous rattling of non-volcanic tremor; and the highly regular tones of Earth’s free oscillations”* (Simpson et al., 2009). A small number of people have continued to develop these efforts with scientific and artistic aims. Those we are aware of include, in some roughly chronological order, David Simpson (Simpson et al., 2009), Florian Dombois (Dombois, 2002), John Bullitt (Baker, 2008)<sup>5</sup>, Andy Michael, Xigang Peng, Debi Kilb (Peng et al., 2012). Our impression is that many of these people (including us) had the initial raw idea to time-compress seismograms that felt completely original and obvious, then later discovered that others had preceded them with the same excitement.

We began in 2005 with a low budget system of 8 self-powered speakers and an audio interface, interested in using the spatialized sound to understand the nature of earthquakes, wave propagation through the planet as an acoustic space. In subsequent years, our large step forward was to couple the sounds with animations of the seismic wave field from both real data and simulations. Seeing the wave field while hearing the sounds enables an immediate comprehension of the meaning of the sounds that is a huge leap from the sound alone, as discussed above. Our first effort to couple the sound and image was to add sound to the animation of smoothed seismic data across Japan from the 2007 Chuetsu-oki (Niigata) earthquake (Magnitude 6.8)<sup>6</sup>, using the method of Furumura (2003). This kind of wave field visualization is only possible for very dense arrays of seismometers, where the station spacing determines the minimum wavelength that can be seen and the areal coverage of the area determines the maximum wavelength. With the advent of the USArray program and its transportable array (TA), such images have become possible for long period waves across the continental US. These “ground motion visualizations” (GMVs) were developed by Charles Ammon and Robert Woodward at IRIS (Integrated Research Institu-

<sup>5</sup> <http://www.jtbullitt.com/earthsound>

<sup>6</sup> <http://www.eri.u-tokyo.ac.jp/furumura/lp/lp.html>

tions for Seismology <sup>7</sup>). The TA contains more than 400 broadband seismometers with an aperture of more than 1000 kilometers, with station spacing of 70 km, such that there is a sampling of the wavelengths across the seismic spectrum that can be seen in the data. We currently synchronize multi-channel sound to these GMVs, but that project will be described in future work.

Subsequently, we synchronized sounds to visualizations of simulations of the global seismic wave field, which is the focus of this paper, as illustrated in Fig. 1. The simulations are generated using SPECFEM (Komatitsch and Tromp, 2002; Komatitsch et al., 2002; Komatitsch and Tromp, 1999), a numerical (spectral element) model that calculates the elastic wave field in the whole Earth resulting from an earthquake. Candler and Holtzman began a collaboration with Turk, in the context of a grant from the National Science Foundation (see Acknowledgements) to develop this material for the Hayden Planetarium at the American Museum of Natural History in New York City, with its 24-channel sound system and a hemispheric dome.

### 3 Sound Production

Here, we first describe the methodology for making a sound using data from a single seismometer and then for an array of seismometers for spatialized sound. At present, most of the following processing is done in MATLAB, unless otherwise noted.

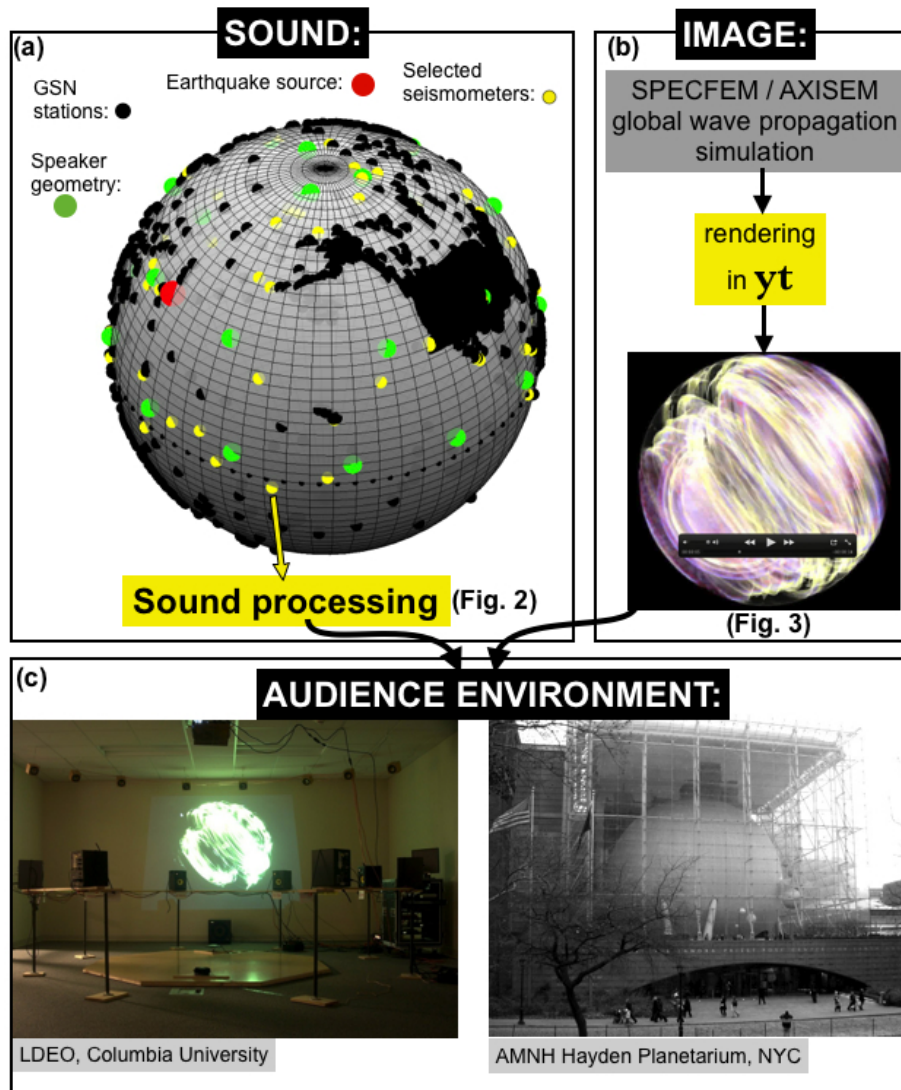
#### 3.1 For a single seismometer

For a chosen event, the data is downloaded from IRIS (using the package “Standing Order for Data” (SOD) <sup>8</sup>, with de-trending performed in SAC (Seismic Analysis Code). The signal in SAC format is read into MATLAB (Fig. 2a). We choose initial values of  $f_e$  (a reference frequency in the seismic signal, where the subscript  $e$  refers to “earth”, from some part of the spectrum that we want to shift to the center of the sonic band) and  $f_s$  (the frequency in the sonic range to which we would like to shift  $f_e$ , as illustrated in Fig. 2b and c), listen and repeat, until achieving the desired sound and time-compression. We also specify what time derivative we would like to listen to (displacement, velocity or acceleration). The sounds are sharper and some volume compression occurs with each successive time derivative. We then design a filter (low-pass, high-pass or band-pass Butterworth), using MATLAB functions, as illustrated in Fig. 2d. We then sweeten and apply additional cleaning/filtering/compression utilizing iZotope RX and Reaper audio software. Alternative methods are under development using python and SoX <sup>9</sup>.

<sup>7</sup> [www.iris.edu](http://www.iris.edu)

<sup>8</sup> <http://www.seis.sc.edu/sod/>

<sup>9</sup> <http://sox.sourceforge.net>



**Fig. 1.** (a) Sound processing: for a given earthquake, we obtain a list of active seismic stations on the globe (black dots), decide on the speaker geometry (green dots) and the relation of it to the earthquake source (red dot), and then find the nearest stations to the speaker locations (yellow dots), download the data for those stations from IRIS, and then run them through our sound generating/processing programs. (b) Image generation: for that given event, we run a simulation either in SPECFEM3D or AXISEM, render the graphics in yt and then synchronize with the sounds. (c) We have a small multi-channel (16) sound system with a projector at LDEO, but are developing this material for the Hayden Planetarium in 2014.

### 3.2 For multiple seismometers and spatialized sound

For a single earthquake, we generate sounds from multiple seismometers to convey the entire wave field and the motion of seismic waves through and around the Earth. As illustrated in Fig. 1a, for a chosen speaker geometry, we build a 3D image of the locations of the speakers on a sphere in MATLAB (green dots). For a given earthquake, we download a list of seismometers that were active at that time (for whatever spatial scale we are interested in, but in this case, global), from <http://global.shakemovie.princeton.edu/> (black dots). We then orient the Earth (or the earthquake source and the array of seismometers) relative to the speaker locations and find the nearest seismometers to each speaker (yellow dots). We run the multi-channel sounds and synchronize with the animation in Reaper, which is also capable of varying time-compression of audio and video interactively. Reaper actions can be scripted using `python` and driven externally using Open Sound Control, forming the platform for installations and exhibits.

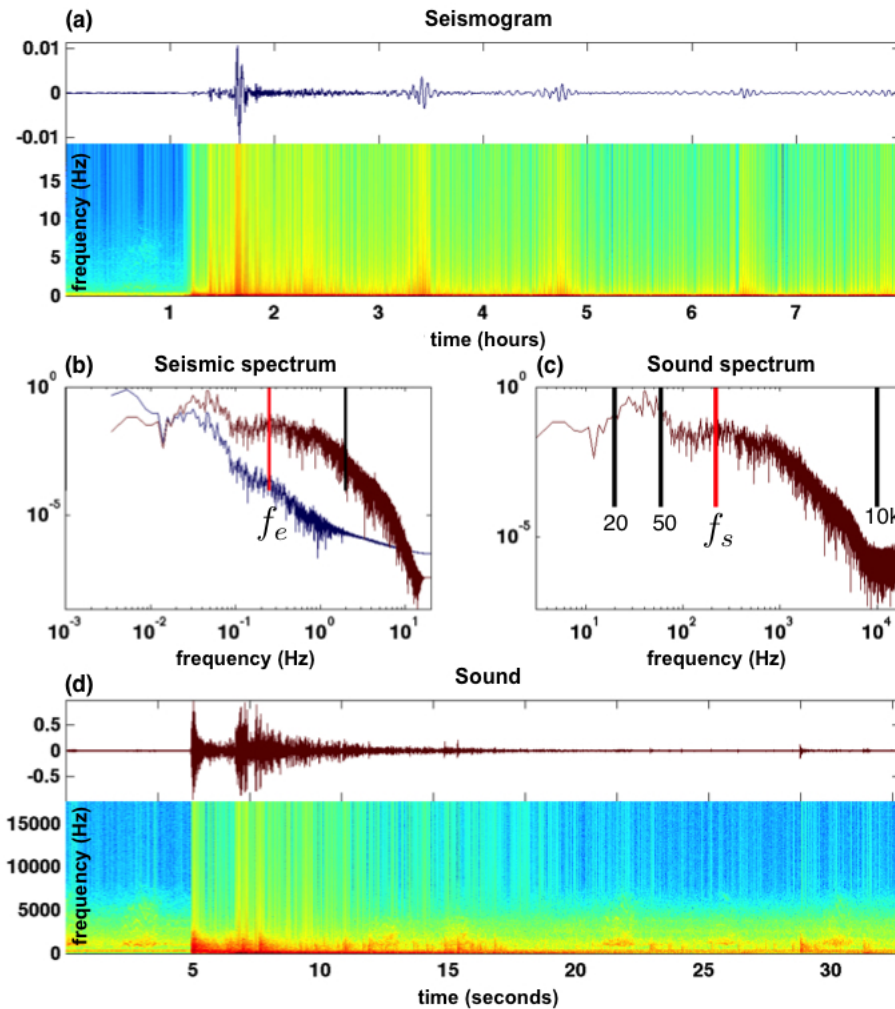
Our current speaker geometries include a flexible array of 24 speakers in the Seismic Sound Lab<sup>10</sup> at the Lamont Doherty Earth Observatory (LDEO) and the 24-speaker dome system in the Hayden Planetarium. At LDEO, the 24 channel array is comprised of an 8-channel ring of self-powered monitors and a 16 channel array of mid-range satellite speakers powered by car-audio amplifiers, with a subwoofer and a infra-speaker platform on the floor. The audience can sit inside the ring (on the infraspeaker) or outside the ring, illustrated in Fig. 1c. The infra-speaker floor adds a great deal of dynamic range to the experience, as listeners often are not aware that their perception of the sound has changed from the sonic to the sub-sonic as the seismic waves disperse and decrease in frequency. At the Hayden Planetarium, the 24-channel speaker array resides behind the hemispheric screen, as shown in Fig.1c (12 at the horizon, 8 at the tropic, 3 at the apex and one subwoofer track).

## 4 Image Production

In the last decade, a major advance in theoretical seismology has come from the development of the spectral element method (SEM) for simulating the wave fields that emanate from an earthquake. SEM combines the geometric flexibility of the finite element method with the temporal and spatial resolution of the spectral methods (Komatitsch and Tromp, 1999, 2002; Komatitsch et al., 2002). At the global scale, the SPECFEM3D code, implementing SEM, can model wave fields with a broad range of frequencies, for realistic crustal structure (e.g. CRUST2.0), limited only by the computational resources available. The movies and synthetic seismograms can be downloaded several hours after any event for which there is a centroid moment tensor produced (CMT,  $M > 5.5$ , <http://www.globalcmt.org>). We are also currently using AXISEM, a quasi-3D code based on the same numerical methods, but with higher symmetry that allows for higher frequency calculations (Nissen-Meyer et al., 2007, 2014)<sup>11</sup>.

<sup>10</sup> [www.seismicsoundlab.org](http://www.seismicsoundlab.org)

<sup>11</sup> <http://www.seg.ethz.ch/software/axisem>



**Fig. 2.** Our standard plots for sound generation and analysis. This data is from the 2011 Tohoku, Japan, M9 earthquake. **(a)** Seismic waveform and its spectrogram. **(b)** Fourier transform of the signal (blue) and filtered signal (red), in this case a high pass filter. The red line marks the reference frequency ( $f_e$ ) used for calculating the time-compression factor (appendix 1). **(c)** FFT of the sound signal ( $f_e$  has been shifted to  $f_s$ ). The black lines mark the ranges of our sound systems (20-50 for bass and 50-10k Hz for mid-upper ranges). **(d)** Sound waveform and its spectrogram. Note the difference between the original seismogram and the sound signal due to the high pass filter; the large pulses, which are the surface waves, are absent in the sound, and much more high frequency detail is apparent in the “coda”. Without the high-pass filter, the surface waves dominate the sound.

Visualizing the output of these simulations allows us to see the wave field in motion. A “membrane” rendering the wave field at the Earth’s surface is automatically constructed as part of the Shakemovie <sup>12</sup> output (Tromp et al., 2010). The wave field can also be rendered on a 2D cross section through the Earth, on, for example, a great circle section <sup>13</sup>. Combined cross section and spherical membrane renderings give a good sense of the 3D wave field and are computationally much cheaper. A beautiful volumetric rendering with great pedagogical value was made for the 1994 Bolivia M8.2 deep focus ( $\sim 630$  km) earthquake (Johnson et al., 2006). To render such complex volumetric forms in a meaningful (and inexpensive) way is an important visualization challenge and major aim of this project.

The `python` environment “`yt`” (Turk et al., 2011) is designed to process, analyze and visualize volumetric data in astrophysical simulations, but in the context of this project is being adapted to spherical geometries relevant to seismology. To visualize this data, techniques used in the visualization of astrophysical phenomena such as star formation and galaxy evolution were applied to the seismic wave fronts. Data was loaded into `yt` in a regular format and decomposed into multiple patches, each of which was visualized in parallel before a final composition step was performed. The visualization was created using a volume rendering technique, wherein an image plane traversed the volume and at each step in the traversal the individual pixels of the image were evolved according to the radiative transfer equation, accumulating from absorption and attenuating due to emission. Colormaps are constructed using Gaussian and linear functions of RGB values and are mapped to the amplitude of net displacement in each voxel and time step. In a given filter, the color shows the radiative value (local emission at each point) and the curvature of the top of the filter shows the alpha value, that describes the transparency (alpha=1 is completely opaque, alpha=0 is completely transparent). The combination of the color map and the alpha function is called the “transfer function”, as illustrated in Fig. 3. This approach results in a smooth highlighting of specific displacement values throughout the volume, illustrating the global wave field for one time step. The example snapshots shown in Fig. 3 were generated from a SPECFEM simulation of the 2011 Tohoku Magnitude 9 earthquake (at [www.seismicsoundlab.org](http://www.seismicsoundlab.org)), discussed further below. Graphics are rendered for flat screen and 180-degree fisheye projections.


## 5 Synchronization of Sound and Image

Our ongoing challenge is to synchronize the natural sounds and synthetic images into a meaningful cinematic object, in order to convey physical aspects of the seismic wave field. The most fundamental character of the Earth as an acoustic space is that its spherical (or near spherical) form with a liquid outer core controls the wave propagation (Shearer, 2009). The waves whose propagation is always controlled by or “aware of” the free surface of the sphere are called surface

<sup>12</sup> <http://global.shakemovie.princeton.edu>

<sup>13</sup> <http://seis.earth.ox.ac.uk/axisem/>





**Fig. 3.** Three time steps of yt-rendered animations, (a) Small displacement transfer function (top) corresponds predominantly to body waves (high frequency). (b) Large displacement transfer function (top) corresponds predominantly to the surface waves (low frequency). This simulation is of the first 4 hours after the 2011 Tohoku, Japan, Mag. 9.1 earthquake.

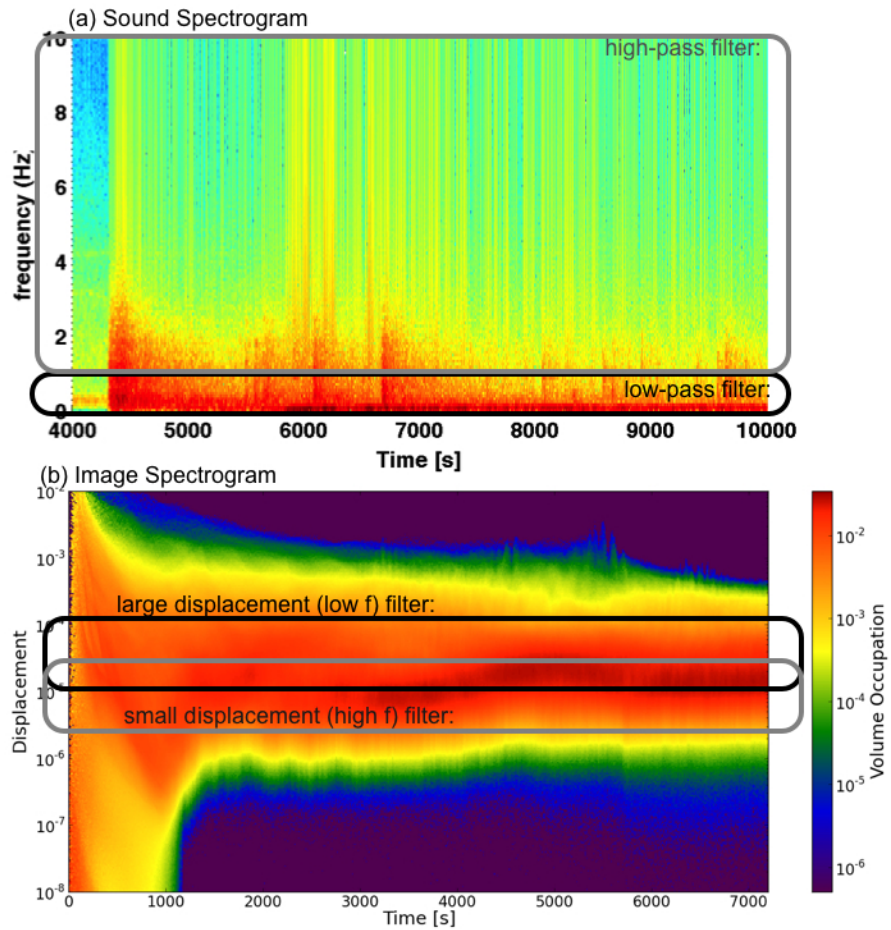
waves. Those that propagate through the interior as wave fronts are called body waves. Surface waves are relatively large amplitude and low frequency, while body waves have smaller amplitudes and higher frequencies. Physically, however, they are continuous parts of the same wave field; surface waves are essentially sums of body waves. However, because they can be identified in seismograms as discrete phases and are thus separable, seismologists analyze them using different methods. Thus, we want to visually and sonically distinguish between surface waves and body waves, to explore this aspect of the physics.

To demonstrate this difference, using the data and the simulation from the Tohoku earthquake, our initial efforts involve filtering both the images and the sounds in parallel, as summarized in Figs. 3 and 4. To make the sounds, we run a low pass and high pass filter on the seismic data, above and below about  $1.0 - 0.5$  Hz, as illustrated in Fig 4a. The surface waves clearly propagate as a wave packet, and the coherent motion is clear when listening in spatialized sound environments.

To make the images, we apply different transfer functions centered on different bands of displacement amplitude, mapping amplitude to color and transparency. The larger displacement amplitudes correspond to the lower frequency surface waves, as shown in Fig. 3, top row. The surface wave transfer function is modified from one designed to render the appearance of a plasma. The smaller displacement amplitudes correspond to the higher frequency body waves, as shown in Fig. 3, bottom row. The body wave transfer function is designed to render semi-transparent, semi-reflective sheets moving through space (a bit like a jellyfish). The wave fields, when separated like this, look very different. Correspondingly, the sounds are very different. Furthermore, the movies with sounds play back at different speeds. The surface wave movies actually have to be shifted more to be in the audible range and thus play back faster than the body wave movies. The synchronization is tight; the sounds correspond to events in the wave motion, and the meaning of the two aspects of the wave field becomes clear. However, much potential for improvement remains in the development of quantitative relationships between the audio and image filters, such that we can explore the behavior of narrower frequency (band-pass) aspects of the wave field.

## 6 Emerging Scientific and Pedagogical Questions

Who is the audience for these movies and what do we want them to understand? We largely ignore this question when designing the sounds and images, but give it full importance when designing presentations and exhibits. We design the material to be as rich in visual and auditory information as possible. In the spirit of Frank Oppenheimer and the Exploratorium, our belief is that the material should contain as much of the richness of the process itself as possible; that richness is what attracts peoples innate curiosity and will be understood in different terms at every stage of experience. Filtering of information to isolate certain patterns (for example, as a function of frequency) should be done late in



**Fig. 4.** (a) Sound spectrogram with approximate filter bands superimposed. (b) Image spectrogram showing displacement on the y-axis and the volume occupied by pixels with that displacement value, represented by the color spectrum. These filtering and rendering methods are works in progress.

the production, ideally by the observer as part of the process of understanding through experimentation.

There is large diversity in people’s sonic and visual perception, as well as ranges of experience in perceiving and interpreting physical behavior. The movies may provoke very similar questions, but the language used to articulate those questions will be very different for a 5-year old than for a professional seismologist. In past presentations, people have used interesting physical analogies to describe the sounds, including: “whales singing in the distance through the deep ocean”, “hands clapping once in an empty warehouse”, “a loose piece of sheet metal flapping in the wind”, “a tree bough cracking in a cold winter night”, “a bowling ball hitting the dirt with a thud”, “a trailer truck bouncing over a bump in the highway as you cling to the underside of its chassis”. These analogies speak to the detailed information we associate with sound on the acoustic properties of different spaces and the physical processes that produce sound, and also to the diversity of sounds in seismic data.

As discussed above, the phenomena in the sounds can be separated into (1) the physical characteristics of the rupture process, (2) the mechanical properties of the rock volumes that the waves are passing through and (3) the geometry or acoustic space— the internal structure of the Earth. The spatialization is important for conveying relative location of events, depth relative to the surface, and motion of wave fronts. In our demonstrations, we try to isolate these effects by a comparative approach. For example, to convey the concept of magnitude, we listen to two earthquakes with different magnitudes as recorded at one seismometer. To convey the concept of the Earth as an acoustic space and compare different paths through it, we listen to one earthquake recorded at different seismometers on the globe. For Seismologists, our intent is that the multi-sensory presentation of seismic data and simulations enables the recognition of patterns that would be missed only through visual inspection seismograms as waveforms and subsequent signal processing; sound may help us identify patterns and events in the waveforms that can then be further isolated by iterative signal processing and listening. Work towards this aim is ongoing.

## 7 Conclusions

The synchronized visual and auditory rendering of the seismic wave field in the Earth is allowing us to experience with our senses an otherwise imperceptible natural phenomenon. The novelty and information-richness of these movies has enormous educational value, but also is beginning to provide new questions and insights for researchers. We have much work to do in the technical aspects of the visualization and sonification and their union, that will improve the perceptibility of patterns in the data.

**Acknowledgements:** We have the good fortune of working with Pritwiraj Moulik, Anna Foster, Jin Ge, Yang Zha and Pei-ying Lin at LDEO and Lapo Boschi at Univ. Paris VI. They have taught us a great deal of Seismology

and contributed generously to many aspects of the project. In addition to co-author Daniel Peter, Vala Hjorleifsdottir, Brian Savage, Tarje Nissen-Meyer and Jeroen Tromp have all contributed to bringing SPECSEM and AXISEM into this project. Art Lerner-Lam gave us the initial financial support and scientific encouragement when we began this project for the LDEO Open House in 2006. Matthew Vaughan has made invaluable contributions to production of movies and sounds. David Simpson, Douglas Repetto, George Lewis and Dan Ellis have provided encouragement and assistance on computer music/sound aspects. Denton Ebel, Carter Emmart, Rosamond Kinzler and Peter Kelemen have made possible bringing this project to the Hayden Planetarium. This work is directly supported by NSF grant EAR-1147763, “Collaborative Research: Immersive Audio–Visualization of Seismic Wave Fields in the Earth (Earthscope Education & Outreach)”.

## A Scaling frequency and duration

In the process of shifting the frequency of a seismic signal, the number of samples (or data points) in the waveform signal ( $n$ ) does not change. All that changes is the time interval assigned between each sample,  $dt$ , where the sampling frequency,  $f_{Sam} = 1/dt$ . Broadband seismometers generally have sampling rates of 1, 20 or 40 Hz. For sound recording a typical sampling frequency is 44.1 kHz. In the context of frequency shifting, consider an arbitrary reference frequency  $f_{Ref}$  such that  $f_{Ref}/f_{Sam} < 1$ , because  $f_{Ref}$  must exist in the signal. When considering frequency shifting in which the number of samples does not change, this ratio must be equal before and after frequency shifting. In the problem at hand, we refer to the original sampling rate of the seismic data as  $f_{Sam}^e$  (for “earth”), with reference frequency  $f_{Ref}^e$ , and the shifting sampling rate and reference frequency as  $f_{Sam}^s$  and  $f_{Ref}^s$ , (for “sound”) respectively, such that

$$\frac{f_{Ref}^e}{f_{Sam}^e} = \frac{f_{Ref}^s}{f_{Sam}^s} \quad (1)$$

As illustrated in Fig. 2, we look at the Fourier transform (FFT) of the original signal, choose a reference frequency based on what part of the signal spectrum that we want to hear (e.g. 1 Hz for body waves), and then choose a reference frequency to shift that value to (e.g. 220 Hz, towards the low end of our hearing). We then re-arrange Eqn. 1 to determine the new sampling rate:

$$f_{Sam}^s = \frac{f_{Ref}^s}{f_{Ref}^e} f_{Sam}^e \quad (2)$$

which is entered as an argument into the the “wavwrite” function in MATLAB.

Similarly, duration is  $t = n \cdot dt$  where  $n$  is the total number of samples, and  $dt$  is the time step in seconds between each data point or sample. Since  $n$  is constant for the original data and the sound ( $n_e = n_s$ ), we can write  $\frac{t_e}{dt_e} = \frac{t_s}{dt_s}$ .

This is usefully re-arranged to

$$t_s = \frac{f_{Sam}^e}{f_{Sam}^s} t_e, \quad (3)$$

which is useful for synchronizing the sounds with animations.

## Bibliography

- Baker, B. (2008). The internal ‘orchestra’ of the Earth. *The Boston Globe*.
- Dombois, F. (2002). Auditory seismology: on free oscillations, focal mechanisms, explosions and synthetic seismograms. *Proceedings of the 8th International Conference on Auditory Display*.
- Furumura, T. (2003). Visualization of 3D Wave Propagation from the 2000 Tottori-ken Seibu, Japan, Earthquake: Observation and Numerical Simulation. *Bulletin of the Seismological Society of America*, 93(2):870–881.
- Hartmann, W. M. (1999). How We Localize Sound. *Physics Today*, 52(11):24–29.
- Johnson, A., Leigh, J., Morin, P., and van Keken, P. (2006). Geowall: Stereoscopic visualization for geoscience research and education. *IEEE Computer Graphics*, 26:10–14.
- Kahn, D. (2013). *Earth Sound Earth Signal*. University of California Press.
- Komatitsch, D., Ritsema, J., and Tromp, J. (2002). The spectral-element method, beowulf computing, and global seismology. *Science*, 298(5599):1737–1742.
- Komatitsch, D. and Tromp, J. (1999). Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophysical Journal International*, 139(3):806–822.
- Komatitsch, D. and Tromp, J. (2002). Spectral-element simulations of global seismic wave propagation - I. Validation. *Geophysical Journal International*, 149(2):390–412.
- Nissen-Meyer, T., Fournier, A., and Dahlen, F. A. (2007). A 2-D spectral-element method for computing spherical-earth seismograms—I. Moment-tensor source. *Geophysical Journal International*, 168:1067–1093.
- Nissen-Meyer, T., van Driel, M., Sthler, S., Hosseini, K., Hempel, S., Auer, L., and Fournier, A. (2014). AxiSEM: broadband 3-D seismic wavefields in axisymmetric media. *Solid Earth Discussions*, 6:265–319.
- Peng, Z., Aiken, C., Kilb, D., Shelly, D. R., and Enescu, B. (2012). Listening to the 2011 Magnitude 9.0 Tohoku-Oki, Japan, Earthquake. *Seismological Research Letters*, 83(2):287–293.
- Shearer, P. (2009). *Introduction to Seismology (Second Edition)*. Cambridge University Press.
- Simpson, D., Peng, Z., and Kilb, D. (2009). Sonification of earthquake data: From wiggles to pops, booms and rumbles. *AGU Fall Meeting . . . .*
- Speeth, S. D. (1961). Seismometer Sounds. *The Journal of the Acoustical Society of America*, 33(7):909–916.
- Tromp, J., Komatitsch, D., Hjörleifsdóttir, V., Liu, Q., Zhu, H., Peter, D., Bozdog, E., McRitchie, D., Friberg, P., Trabant, C., and Hutko, A. (2010). Near real-time simulations of global CMT earthquakes. *Geophysical Journal International*, 183(1):381–389.

- Turk, M. J., Smith, B. D., Oishi, J. S., Skory, S., Skillman, S. W., Abel, T., and Norman, M. L. (2011). yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *Astrophysical Journal Supplements*, 192:9.
- Walker, B. and Nees, M. (2006). Theory of Sonification. *Principles of Sonification: An Introduction to Auditory Display and Sonification*, Ch. 2:1–32.