Exploration of Three Dimensional Vector Field Data Using Sound

Eric Klein* Oliver Staadt† Bernd Hamann‡

Abstract

We describe and analyze a working implementation of a new technique for creating wind-like noises from three dimensional vector fields. This process can also be referred to as sonification. This technique allows the user to use 3D sound algorithms to map vectors in a listener's local neighborhood into smooth wind-like sound (aerodynamic sound). The three types of information provided by this technique are flow direction, flow velocity, and flow turbulence (vorticity). The end result is a system that helps the listener achieve an intuitive understanding of the dataset by using their sense of sound.

1 Introduction

The concept of using sound in computer simulations to enhance the user's understanding of their environment is not new. In fact, it has been used in computer gaming technology since at least the 1980s. It is however not often used in the scientific visualization community. As the size of scientific data sets is increasing rapidly, it is becoming harder and harder to work with this data due to its size and complexity. Often when using purely visual investigative techniques the data sets are so massive that the result is an overload of the user's visual system. Combining visualization techniques with sonification has the potential of offloading some of the data interpretation to an often underused sense. An ideal place to introduce sound is in 3D vector fields. These vector fields are usually very dense volumes of data (often 256x256x256 voxels per timestep or more), sometimes with many timesteps, and even using current techniques for vector field visualization such as particles, streamlines, streamers, streamtubes [1] and LIC (Line Integral Convolution) [2], it is sometimes difficult to see the flow pattern in such a system. Another thing that makes vector fields prime candidates for sonification is that they almost always represent some kind of flow, whether it is a fluid, a gas or a plasma. The idea of a flow combines very naturally with the concept of wind, and most people have an intuitive sense of what wind should sound like like for a flow (such as in a wind tunnel air flow).

Others have noticed this opportunity before, and have exploited it to sonify vector field data. There are two distinct approaches that have been successfully employed for sonification of vector field data. Eckel et. al. [3] have used their cyberstage for exploring vector fields, and sampling either points or whole paths within the data to produce an aerodynamic sound from that point (or path). It is unclear from the description in their paper precisely what information the aerodynamic sounds conveys, but presumably it conveys at least velocity information about the flow. Since the cyberstage has a complex positional audio system, it would also make sense if the wind noise also conveys local flow direction. A second approach comes from Volpe and Glinert [4] in which the vector field is displayed using streamtubes. The user explores the data along the streamtubes, and a continuously changing pitch is heard to represent the changing vorticity along the streamtube.

Our technique is most akin to Eckel's [3]. In our technique, points in the listener's local region are sampled and converted to wind sounds. For a given point, the wind sound coveys two pieces of information: how fast the data flow is at that point, and which direction it is flowing. Our approach adds vorticity data to this. This is accomplished by randomly sampling different points within the listener's local region several times a second, and smoothly interpolating between them. This approach works well in practice and conveys to the listener the data from both Volpe's technique [3] and Eckel's technique [4].

2 Spatial Sound

Our approach is based on 3D sound spatialization [5] [6]. Sound spatialization algorithms all convert a particular sound that lies somewhere in space around the listener's head in such a way that the user can better perceive exactly where that sound is in space. This is accomplished using either multiple speakers or a head related transfer function (HRTF) [7] in concert with a set of headphones. HRTFs are a large and active area of research, and many groups (such as the user interface lab at UC Davis [8]) are working on developing and improving HRTFs.

HRTFs have developed to such a point that it is often possible using headphones to fairly accurately identify which direction a sound is coming from. When head tracking is added and the user is able to move his or her head to help locate the sound, these systems become extremely useful. In this paper we describe a technique to map the data in the listener's local region to a position in relation to the listener's head. Once that is done, the user can listen to the data.

3 Mapping a Vector to Sound

The first step in our technique is to be able to map a single data point to a single sound location. This can be achieved in a 3D vector field since each data point is a vector which gives us both a direction and a magnitude. This is accomplished with the following eqation:

$$P = T \times \left(-\frac{V}{|V|} \times \frac{1}{\sqrt{|V|}} \right)$$

where P is the point in relation to the listener's head, V is a vector from the vector field, and T is the transform from the world coordinate space to the listener's coordinate space. The $\frac{-V}{|V|}$ term gives a unit vector in the direction from the listener to the sound (in world space), and the 1 over square root of the vector magnitude term gives us the scaling factor. This makes sense because sound often follows a $\frac{1}{distance^2}$ attenuation pattern. Essentially that means that the longer the vector is, the louder the sound should be, and thus the closer it should be to the listener. It would also be possible to increase pitch as well as volume with increasing vector magnitude (in contrast to Volpe's [4] method, which uses

---

[*] elklein@ucdavis.edu

[†] staadt@cs.ucdavis.edu

[‡] hamann@cs.ucdavis.edu

Data Feature	Sound Property
Vector Direction	Sound location
Vector Magnitude	Distance between sound and listener. Also controls sample rate.
Vorticity (turbulence)	Gustiness of wind. How wildly the wind sound shifts around, and how fast.

Table 1: Vector Field Data Conveyed Using Sound

pitch to represent vorticity changes), although this isn't done currently. Lastly, we need to map from the world coordinate space to the listener's coordinate space, in order to account for the listener's current orientation in space.

4 Mapping from a Region to Sound Vorticity

Now that we have the ability to convert a single location in our dataset into a sound, we still have one very important piece of information that we wish to convey: vorticity. Since vorticity is the measure of change over a region (it can also be thought of as turbulence), this makes it necessary for us to consider more than just a single point in order to sonify vorticity.

Our approach to this problem is to sample in a uniformly random way a bounding volume around the listener's head. These samples are taken periodically based on the magnitude of each sample (the higher the magnitude, the sooner the next sample will be taken) and varies between one and ten times a second. Each sample is then mapped to a position around the user's head. In order to ensure that the sound doesn't jump discontinuously between each sample time, it is necessary to smoothly interpolate between sample positions at intermediate time steps. This is done by creating a curved line (a Hermite curve) that connects each set of sample points. This line has the property that if two sample points are the same distance from the listener (but possibly on opposite sides) the line will keep approximately this distance the entire time. Since we smoothly interpolate along this line, the sound appears to move smoothly and continuously across all samples.

This technique has the advantage of moving the sound very little in areas of low vorticity and very greatly in areas of high vorticity, much as actual wind would do. The end result is that in high vorticity areas, you get a rapidly shifting, gusty wind noise, and in low vorticity areas, you get a slowly moving smooth wind noise.

5 Results

Using our technique the listener gets a very good sense of the data flow in their neighborhood, and by moving around within the vector field, the listener is able to gain a good intuitive understanding, using only sound, of the general flow properties of the vector field. Table 1 illustrates the various information conveyed to the listener from these wind sounds.

We have found though that this technique is best used alongside traditional visual techniques. On its own, this technique would require a lot of work to fully understand a vector field. However, when used in conjunction with other techniques, it helps the listener to achieve a better grasp of the data. We found that the greatest utility for this technique is the information that it provides on vorticity data. It is fairly easy to visually provide the direction and magnitude information about a given location, but it is much harder to show vorticity visually. With the continually shifting wind sounds (and the way in which they shift) it is very easy and natural to understand the vorticity.

This project is especially useful for situations where there is other information, such as the interior or exterior view of a vehicle, that the user must explore with their eyes, and yet flow (air flow or waterflow around the vehicle for example) information must still be conveyed. It could also be useful as an aid to accessability, in the case where blind users may need to explore the data.

6 Future Work

There are several areas where this technique can be further improved. An interesting possibility is to use spatial sound to help navigate to critical points within the data. Calculating the direction of curvature for a streamline at a given point and using that as the sound direction could do this. Then using the local rate of curvature we could derive a pitch and/or volume. This information would combine to give the listener a guess at a direction to navigate in, and approximately how close they are to the critical point. The sampling scheme discussed earlier could further be employed (but with averaging as well as interpolation) to smooth out any anomalous guesses. Another possibility is to fully synthesize the sound based on attributes from the data using a technique similar to [9], rather than simply modifying an original static sound sample.

References

[1] Dave Darmofal and Robert Haimes. Visualization of 3-d vector fields: Variations on a stream. `http://citeseer.nj.nec.com/2303.html`.

[2] B. Cabral and L. Leedom. Imaging vector fields using line integral convolution. In *Computer Graphics Proceedings, Anual Conference Series*, pages 263–270, 1993.

[3] G. Eckel et. al. A spatial auditory display for the cyberstage. *International Conference on Auditory Display*, 1998.

[4] C. Volpe and E. Glinert. Auralization of streamline vorticity in computational fluid dynamics data. *IEEE Visualization*, pages 51–57, 1997.

[5] Microsoft Corporation. Direct sound api for 3d sound. `http://msdn.microsoft.com/library/default.asp?url=/library/en-us/direct%x9_m/directx/sound/directsound.asp`.

[6] Creative Labs Corporation. Openal library for 3d sound. `http://www.openal.org/`.

[7] V. R. Algazi, R. O. Duda, R. Duraiswami, N. A. Gumerov, and Z. Tang. Approximating the head-related transfer function using simple geometric models of the head and torso. *J. Acoust. Soc. Am.*, 112:2053–2064, Nov. 2002.

[8] UC Davis CIPIC Interface Laboratory. `http://interface.cipic.ucdavis.edu/`.

[9] Y. Dobashi, T. Yamamoto, and T. Nishita. Real-time rendering of aerodynamic sound using sound textures based on computational fluid dynamics. In *ACM Trans. On Graphics (Proceedings of SIGGRAPH 2003)*, pages 732–740, 2003.