IPT-EGVE Symposium (2007)
B. Fröhlich, R. Blach, and R. van Liere (Editors)

Augmentation of Visualisation Using Sonification: A Case Study in Computational Fluid Dynamics

M. Kasakevich[1] P. Boulanger[1] W. F. Bischof[1] and M. Garcia [2]

[1]Department of Computing Science & University of Alberta, Edmonton, Alberta,
[2]Department of Mechanical Engineering & EAFIT University, Medellin, Colombia

---

Abstract

Advances in computer processing power and networking over the past few years have brought significant changes to the modeling and simulation of complex phenomena. Problems that formerly could only be tackled in batch mode, with their results being visualized afterwards, can now be monitored using graphical means while in progress. In certain cases, it is even possible to alter parameters of the computation while it is running, depending on what the scientist perceives in the current visual output. This ability to monitor and change parameters of the computational process at any time and from anywhere is called computational steering. Combining this capability with advanced multi-modal tools to explore the data produced by these systems are key to our approach. We present an advanced multi-modal interface where sonification and 3D visualization are used in a computational steering environment specialized to solve real-time Computational Fluid Dynamics (CFD) problems. More specifically, this paper describes and experimentally proves how sonification of CFD data can be used to augment 3D visualization.

Categories and Subject Descriptors (according to ACM CCS): H.5.1 [Information Interfaces and Presentation]: Audio input/output, H.5.2 [Information Interfaces and Presentation]: Auditory (non-speech) feedback, Haptic I/O, User Interfaces, Evaluation/methodology, I.3.6 [Computer Graphics]: Interaction techniques

---

1. Introduction

Advances in computer processing power and networking over the past few years have brought a significant change to the modeling and simulation of complex phenomena. Problems that formerly could only be tackled in batch mode, with their results being visualized afterwards, can now be monitored using graphical means while in progress, and, in certain cases, it is even possible to alter parameters of the computation whilst it is running, depending on what the scientist perceives in the current visual output. This ability to monitor and change parameters of the computational process at any time and from anywhere is called computational steering [BGBR06]. Combining this capability with advanced multi-modal tools to explore the data produced by those systems are key to our approach. In this paper, we present such an advanced multi-modal interface where sonification and 3D visualization are combined in a computational steering environment specialized to solve real-time Computational Fluid Dynamics (CFD) problems. More specifically, this paper describes how sonification of CFD data can be used to augment 3D visualization. We provide the results of a usability study as a proof of the concept. Figure 1 shows a general overview of the real-time CFD processing environment. This paper is not concerned with how a real-time CFD solver works [BGBR06], but instead it describes how one can convey useful information about CFD data through a combination of visual and sound modalities. Today most computational fluid dynamists use 3D visualization tools such as ParaView [Kit], AVS Express [AVS], and AMIRA [MER] for viewing the results of CFD simulations. The visualization techniques may not be sufficient because the user might miss important details: As simulation datasets become larger and have a larger dimensionality, they become increasingly difficult to visualize with simple visual encoding using color, length of arrows, icons, and others. In the scientific visualization community, this problem is referred to as the dimensionality curse. Sonification may be a possible solution to this problem because it can be used either as an alternative or as a complement to visualization, because it increases informa-

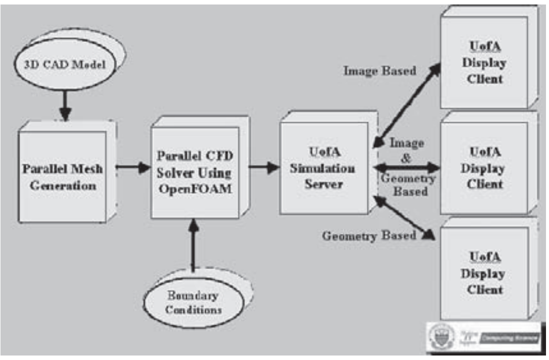

Figure 1: *Real-Time CFD Solver Architecture at the University of Alberta (UofA)*

tion bandwidth, and because it can help recognizing features that are not obvious from other sources. For example, the visual sense is best at processing spatial information, whereas the auditory sense is better for processing sequential or temporal information [WW86]. Furthermore, multi-modal sensory experience is known in the psychophysics literature to improve perception discrimination [WW86, SD04]. For example, global patterns could be presented through vision, while local details could be presented through sound. One can also present different data attributes to different senses, thus reducing the effect of the dimensional curse. For example, the temperature characteristics at each data vertex of a CFD solution could be presented through visual cues, pressure through haptics or force feedback, and flow or vorticity through sound. One of the goals of this paper is to describe the basic ideas of real-time CFD data sonification and its implementation in the context of a virtual wind tunnel [BGBR06]. In Section 2, we discuss the relevant literature on CFD sonification and in Section 3 the mapping functions developed for real-time CFD data. In Section 4, we describe the implementation details, and in Section 5 we provide the results of a usability study for a vortex localisation task. We then conclude by summarizing the results obtained so far.

2. Previous Work

Attempts to use sonification can be found throughout the literature. Sonification can be used for alarms, real-time feedback, exploratory data analysis, and others. Specifically to CFD, applications include testing model geometry and grid for possible discontinuities, monitoring solution the process, and exploring the CFD solution data.

An example of using sonification for monitoring CFD solution process is presented in Child's work [Chi01]. The idea behind his sonification algorithm is that one needs to be able to listen to a CFD solver progress, to determine if it is converging or if the solver needs to be stopped and restarted with different parameters. If the produced sound is converging towards a specific sound then so is the solution. The parameter-sonification mapping used in this example is fairly simple and is based on modifying sound frequency and envelope.

One of the CFD solution exploratory examples is described in Klein's work [KS04]. A sphere representing a user's head is interactively moved around the data field, and only the data samples inside that sphere affect the sound wave reaching a virtual microphone. Sonification is performed interactively and in real-time. The direction and magnitude of each vector is mapped to sound location, level and pitch. If all of the samples in the sphere are of roughly the same magnitude and direction, then a constant and smooth sound is synthesized, indicating low vorticity in the area. If flow vectors vary widely, sound appears to shift more, giving the impression of higher turbulence. The paper also indicates that the sound produced by this algorithm is very helpful and easy to understand, but this is stated without any proof or perceptual evaluation.

In general, one of the rules of data sonification is that careful analysis of specific data properties as well as of the purpose of the required sonification is very important for creating a sonification that is easy to interpret. Some interesting ideas from the literature can be used. Basic physical ideas, like the fact that the apparent magnitude of a sound varies with the inverse square of the distance, are obvious and simple. Another good idea is that of a virtual microphone that can be positioned in the flow field, and record sound sources inside a sphere of influence. One of the problems that Klein mentions in his paper is how to choose the right sphere diameter in order to preserve intelligibility. To explore this idea, one could give the user control over the sphere diameter, which could go from a single point to the whole field. A usability study can then be performed to determine the optimal radius. In many ways, Klein's [KS04] work is very close to the current project, with a number of differences. The field he is analyzing is defined on a rectilinear grid of vectors thus simplifying the connections and relative locations between data points. Choosing a more complicated field grid without such a nice structure introduces a more complicated relationship between user position and the selected data region. Further, Klein is only concerned with finding a good representation of the data in the selected region. In his system, the sphere of influence of the virtual microphone cannot be modified, making the system difficult to test in perceptual studies.

In the next sections, we not only define a mapping function we use to sonify CFD data, but also present usability analysis showing advantages of using sonification for CFD data.

3. Sound Mapping Functions for CFD

In this section, we discuss possible mapping ideas for CFD data. One of the problems of sound mapping is to determine at which scale one needs to operate. One can have an ambient, global sound associated with the whole CFD domain, or a local sound specific to a particular point, area or line of interaction.

For a global sonification, every node contributes to the sonification process, either uniformly or dependent on the distance to a virtual microphone in 3D space. For local sonification, only the field values interpolated at the virtual microphone position are used to synthesize the sound. In CFD, the field value at time-step t at the microphone position $r_m = (x_m, y_m, z_m)^T$ is defined by a tensor $\mathbf{p}_m(t) = (\rho(t), p(t), T(t), v(t), \Omega(t))^T$ where $\rho(t)$ is the fluid density, $p(t)$ the pressure, $T(t)$ the temperature, $v(t) = (v_x(t), v_y(t), v_z(t))^T$ the fluid speed, and $\Omega(t)$ is the vorticity of the fluid, which is proportional to the rotational speed. In this scheme, the field tensor values $\mathbf{p}_m(t)$ are interpolated for sonification by first determining which grid cell the virtual microphone is located in and then, using Schaeffer's interpolation scheme, the flow parameters $\mathbf{p}_m(t)$ at that position are computed using the following equation:

$$\mathbf{p}_m(t) = \frac{\sum\limits_{ForAllVertices} \mathbf{p}_n(t)/||r_m - r_n||^2}{\sum\limits_{ForAllVertices} 1/||r_m - r_n||^2} \qquad (1)$$

Following this interpolation process, noise is then shaped and modified in amplitude and frequency to simulate a wind effect. White noise is filtered through a band-pass filter with fixed bandwidth, were the central frequency is linearly mapped to the field velocity modulus $||v(t)||$ at a given point. The amplitude of the sound is calculated from both velocity value $||v(t)||$ and the angle α between a virtual pointer and the field velocity vector at a given point by the simple relationship $F_1(||v(t)||) \times F_2(\alpha)$. This mapping makes the sonification synthesis sensitive to the amplitude and orientation of the flow relative to the probe. The scaling function F_1 first transforms the modulus of the speed vector by a simple non-linear mapping $||v(t)||^{5/3}$ based on psychophysical considerations [Yos00, GGF89], and the value of this mapping is normalized to a value in the interval [0,1]. Similarly, the function F_2 scales the angle α using the same non-linear relation and then is normalized to a value in the interval [0.5, 1] to ensure that we hear the simulated wind even if we are not facing it directly.

An interesting extension of this sonification algorithm is to expand the interpolation function to nodes inside an influence radius R where the flow field is interpolated using Schaeffer's interpolating function. In this case, nodes in an area around the virtual microphone contribute to the sonification. By interactively changing the radius of interaction, the user can change between local and global sonification.

Further modifications of the mapping algorithm can be introduced: Instead of mapping an increase in velocity to an increase in amplitude, one can be map it to decrease in amplitude, or one can relate amplitude to the angle α only, while keeping central frequency dependent on velocity only, thus creating a frequency-dependent mapping rather than an amplitude-dependent mapping.

Finally, one can view the grid nodes within a radius R of the virtual microphone position r_m as virtual sound sources, located at a distance $d = ||r_m - r_n||$ from the microphone. The contributions of each virtual source are then added using the familiar $1/d^2$ law for sound propagation. As in the other interpolation schemes, the radius of influence could be modified to change the extent of the sensitivity of the virtual microphone. One could also use different attenuation laws that amplify certain preferred orientations in the flow field. For any of these sonification schemes, one will have to do a perceptual analysis to determine the efficiency of each mapping function.

4. Multi-Modal Interface Implementation

In this section, we discuss the implementation details of the proposed multi-modal interface. The system has been implemented on a dual Xeon PC under Windows XP operating system using Visual Studio .NET programming environment and Max/MSP [Cyc04] graphical programming environment. There are two main and one subordinate threads that run simultaneously. The main (sonification and visualization) threads are completely independent of each other and depend on the subordinate (haptic) thread to receive position and orientation of the virtual pointer as well as the radius of the interaction space. Since the main purpose of this interface is to sonify CFD data, the sonification thread is only dependent on the pointer position received from the haptic thread. The visualization thread is also dependent on the haptic thread for visual interaction with the CFD data through the virtual pointer. If no other threads are not running, the visualization thread can be used, without interaction, for viewing the data. In this case, the CFD data can be viewed from different viewpoints without interactive sonification.

The data is produced by the simulation server as a series of time-steps, where the data structure stays the same, but the attribute values at each nodes change. The data structure includes the coordinates and connections between the data nodes and vertices of a CFD mesh. Visualization of the CFD data is done using OpenGL Performer [SGI91] from geometry produced by the VTK [Kit] library. The advantage of OpenGL Performer is that, for a given visualization task, the display modality can be easily changed, for example, from a computer screen to a stereo cave display.

Because sonification is the main emphasis of this project, only a basic visualization interface was developed. Visualization consists of displaying arrows at each node of the fluid field (see Figure 2), a representation of the virtual pointer, and the (influence) interaction sphere in the field (see Figure 3). Direction, size and color of each arrow correspond to

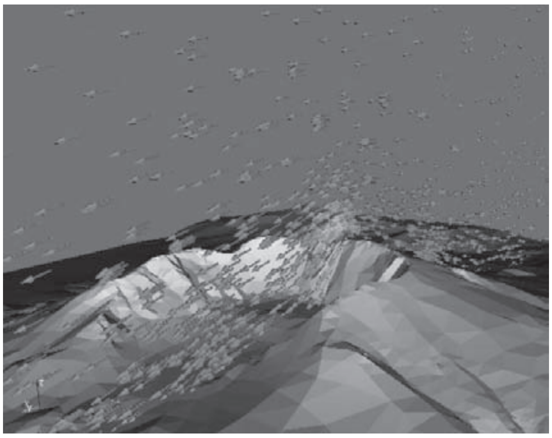

Figure 2: *CFD visualization of the airflow over Mount-Saint Helens.*

the velocity vector values at that node. A virtual pointer is displayed as a white arrow whose position and direction corresponds to the relative position and direction of the haptic device. The influence radius is represented as a sphere located at the end of the virtual pointer, with a diameter specified by the haptic device. Finally, the visual fluid field can be rotated using the mouse to look at it from different viewpoints.

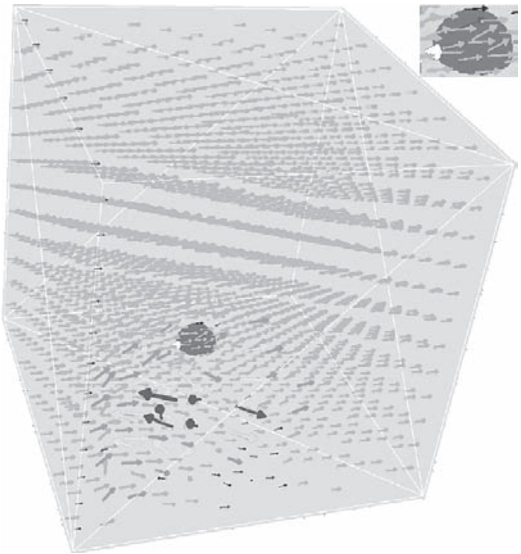

Figure 3: *Example field used in the usability study described in Section 5, and a zoomed in region showing the 3D pointer with an interaction space sphere.*

The haptic device is used to provide feedback on the position of the virtual pointer within the fluid field data allowing users to navigate only inside a given 3D field. The haptic device could also be used as a force-feedback interface, producing a force that is proportional to the flow density and its direction.

Both, the main (visual and auditory) threads and the secondary (haptic) thread connect independently to the simulation server to receive the dataset. Connection is done using Quanta libraries [EVL02], an advanced communication package used for data exchange with the simulation server. The simulation server is in charge of distributing, in real-time, the CFD time-steps computed on a remote high performance computer. After receiving a predefined number of time-steps, both programs disconnect from the simulation server and start the rendering process.

Max/MSP libraries are used to produce a sound for the given dataset. Max/MSP provides a nice graphical programming environment for sound manipulations (see Figure 4). A main thread is written as an object for the Max/MSP interface, which can then connect to other Max/MSP objects to produce a sophisticated sonification program synthesizing sounds.

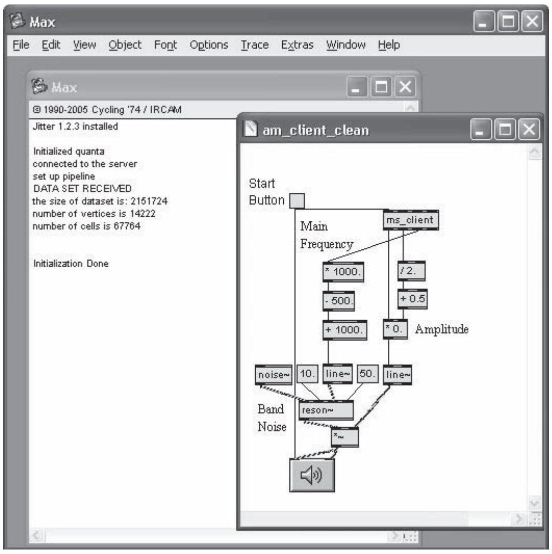

```
© 1990-2005 Cycling '74 / IRCAM
Jitter 1.2.3 installed

Initialized quanta
connected to the server
set up pipeline
DATA SET RECEIVED
the size of dataset is: 2151724
number of vertices is 14222
number of cells is 67764


Initialization Done
```

Figure 4: *Simple program in Max/MSP environment.*

Because all threads have a copy of the dataset and directly receive virtual pointer position from the haptic thread, they are completely independent from each other in the data processing pipeline. At any given moment, the sonification thread reads the 3D pointer position and the interaction space radius from the haptic tread. Depending on these values, it calculates which mesh nodes are the closest and then interpolates the values of the flow properties that need to be rendered at the pointer position. The 3D pointer orientation is also received from the haptic thread and used to calculate the mapping function described in Section 3. The mapping produces a bandpass-noise with central frequency ranging between 500Hz and 1500Hz dependent on velocity value at the pointer, and amplitude dependent on the velocity and the angle between the velocity vector and the pointer. For frequency-dependent mapping, the main frequency ranges between 1000Hz and 4000Hz, and amplitude depends on the angle only.

5. Usability Analysis

The goal of the usability analysis presented here was to examine whether and how one can improve the localization of vortex centers in virtual wind tunnel applications using a combination of auditory and visual feedbacks. To this purpose we investigated the efficiency of vortex localization using a purely visual interface, a purely auditory interface, and a multimodal visual-auditory interface.

Each participant was presented with a set of fluid fields with a vortex at a random location and was required to locate each of the vortices as quickly and as accurately as possible (see Figure 3). Each trial began with the probe at a random location. The participant had to move the probe to the vortex location using a haptic input device (Phantom Omni) and press a button when the probe was believed to be at the correct position. In conditions with visual feedback, the flow field was displayed on a screen with the location of the vortex being marked by a large red arrow (see Figure 3). This display could be rotated using a mouse and thus viewed from different viewpoints. In conditions with auditory feedback, a sound presented with earphones provided feedback regarding the distance between the probe and the vortex location, as described in detail below.

Each participant was tested with each of the three feedback conditions (Visual-Only, Auditory-Only, Visual+Auditory). In conditions with auditory feedback, sound was generated in one of three ways: In condition Positive-Amplitude, sound amplitude increased the closer the probe was to the vortex, in condition Negative-Amplitude, sound amplitude decreased the closer the probe was to the vortex, and in condition Frequency, sound amplitude remained constant but center frequency increases the closer the probe was to the vortex. Each participant was randomly assigned to one of the Sound-Modulation conditions. A total of 30 participants were tested, and each participant completed 36 trials each.

Three measures were used to evaluate the efficiency of vortex localization: Time, the time (in mins) it took the participant to move the probe from the start location to the vortex location; path length, the arc length of the path the probe was moved from the start location to the vortex location; and localization accuracty, the distance between the probe and the vortex center when the participant indicated, using a button press, that the vortex center had been reached.

An outlier analysis was performed and outlier data points were discarded. Only 1% of the data were considered to be outliers, while the rest 99% of the data was included in the final analysis. The distributions of the dependent measures path length, time and localization accuracy were skewed, so a normalizing transform (log) was applied to each dependent measure before the statistical analysis.

An analysis of the Feedback condition showed that participants' performance was worst for the Auditory-only condition, better for the Visual-only condition, and best for the Visual+Auditory condition. This was true for path length [F(2,58)=327.97, p<0.001], time [F(2,58)=21.23, p<0.001], and for localization accuracy [F(2,58)=216.61, p<0.001]. A further analysis showed that this effect was mostly due to a much inferior performance in the Auditory-only condition.

In the following we analyze the Visual-only and the Visual+Auditory conditions in more detail. Participants took less time to locate the vortices in the Visual+Auditory condition than in the Visual-Only condition [F(1,29)=16.99, p<0.001], path length was shorter [F(1,29)=7.48, p<0.05], but localization accuracy was not affected [F(1,29)=0.00, p>0.1]. This is shown in Figure 5. These results indicate that participants were faster in locating the goal position and explored less space using a multi-modal (Visual+Auditory feedback) system.

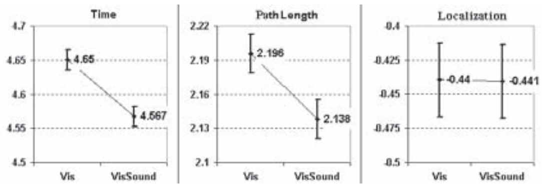

Figure 5: *The results of multi-modal interface vs. visual only interface.*

An analysis of auditory feedback type for Visual+Auditory condition showed that participants' performance improved in some cases. Participants took less time to locate the vortices in the Positive-Amplitude and Negative-Amplitude conditions [F(2,357)=10.05, p<0.001], Localization accuracy was better for the Positive-Amplitude condition [F(2,357)=6.59, p<0.01], but path length was not affected [F(2,357)=2.85, p>0.05]. These results indicate that Amplitude conditions helps in navigating to the goal the fastest, while Positive-Amplitude helps to locate the goal with least error.

Overall, the results indicate that multi-modal feedback system shows an improvement in localization over the visual-only system in most criteria.

6. Conclusions

In this paper, we present results on a multi-modal interface for a virtual wind tunnel. We implemented several mapping algorithms that create very promising results for CFD sonification. These different sonification mappings were then studied experimentally showing several advantages of the multi-modal interface over pure visualization. The experimental study also helped us determine how some mappings are better at helping users to perform analysis of the fluid in a vortex localization task.

We are planning to explore the use of multiple speakers around the user head to provide information on the sound directions. As demonstrated by other researchers, this will give the user a better feeling of immersion into the simulation, giving him/her a more natural representation of the fluid field direction.

We plan to explore other fluid-field sound renderings with this versatile architecture, including sonification along path lines, streak lines, streamlines, and stream tubes.

References

[AVS] AVS ADVANCED VISUAL SYSTEMS. AVS Express, 1991. http://www.avs.com.

[BGBR06] P. BOULANGER, M. GARCIA, C. BADKE, J. RYAN. An Advanced Collaborative Infrastructure for the Real-Time Computational Steering of Large CFD Simulations. *European Conference on Computational Fluid Dynamics*, TU Delft, Netherlands, Sept. 2006.

[Chi01] E. CHILDS. The sonification of numerical fluid flow simulations. *In Proceedings of the 2001 International Conference on Auditory Display*, pp 44−49, Jul.29 − Aug.1 2001.

[Cyc04] CYCLING ’74. Max/MSP, 2004. http://www.cycling74.com/.

[EVL02] EVL ELECTRONIC VISUALIZATION LABORATORY. QUANTA, July 2002. http://www.evl.uic.edu/cavern/quanta/index.html.

[GGF89] G.A. GESCHEIDER, W.L. GULICK, R.D. FRISINA. *Hearing: Physiological Acoustics, Neural Coding, and Psychoacoustics*. Oxford: Oxford University Press, 1989.

[Kit] KITWARE. VTK. http://www.vtk.org/. [Kit] KITWARE. ParaView. http://www.paraview.org/.

[KS04] E. KLEIN, O.G. STAADT. Sonification of three-dimensional vector fields. *In Proceedings of the SCS High Performance Computing Symposium* pp 8, 2004.

[MER] MERCURY COMPUTER SYSTEMS INC. AMIRA. http://www.amiravis.com/.

[SD04] C. SPENCE, J. DRIVER. *Understanding intersensory integration*. Oxford: Oxford University Press, 2004.

[SGI91] SILICON GRAPHICS, INC. (SGI). OpenGL Performer, 1991. http://oss.sgi.com/projects/performer/.

[WW86] R.B. WELCH, D.H. WARREN. Intersensory interactions. In K. R. Boff, L. Kaufman, J. P. Thomas, editors, *Handbook of Perception and Human Performance: Vol 1. Sensory Processes and Perception*, pp 25−1 − 25−36. New Cork: Wiley, 1986.

[Yos00] W.A. YOST. *Fundamentals of Hearing: An Introduction, Forth Edition*. New York: Academic Press, 2000.