

Modeling fluid flow in the airway using CFD with a focus on fricative acoustics

Peter Anderson

Univ. of British Columbia
Dept. of Mechanical Eng.
pja00@interchange.ubc.ca

Sheldon Green

Univ. of British Columbia
Dept. of Mechanical Eng.
green@mech.ubc.ca

Sidney Fels

Univ. of British Columbia
Dept. of Electrical and Computer Eng.
ssfels@ece.ubc.ca

Abstract

A wide range of computational fluid dynamics (CFD) simulations are performed on a static geometry based on a mid-sagittal plane image of a speaker uttering the fricative /sh/ (International Phonetic Alphabet symbol: /ʃ/). The simulations include: 2D incompressible flow, 2D compressible flow, 3D compressible flow, using RANS and LES turbulence models as well as considering the case of no turbulence model. The resulting sound is recorded using direct pressure measurement and the acoustic analogy. The simulation results are compared with previous experiments for validation.

All RANS simulations are found to be inadequate due to averaging used in the model. The direct pressure measurement from 2D simulations gives unphysical results, but the acoustic analogy result is more reasonable. The 3d simulations offer the best results, yet are limited by the size and runtime of the simulation. Both direct measurement and the acoustic analogy are shown to be useful approaches in measuring the sound.

We discuss the usefulness of fluid simulations in the airway, the requirements for such simulations, and the challenges of meeting those requirements. Finally, we present suggestions for further research in this direction.

Keywords: Computational Fluid Dynamics, Large Eddy Simulation, turbulence, acoustics, acoustic analogy, speech simulation, fricative, CFD, LES, RANS.

1. Introduction

Fricatives are produced when air is channeled through a constriction in the airway, thus forming a jet, which strikes an obstacle (such as the teeth) and produces sound. The fricative may combine sound produced at the vocal chords with the sound produced at the constriction (a voiced fricative), or it may be only the sound produced at the constriction (a voiceless fricative). While a fricative can be produced with a relatively static vocal tract, in the context of normal speech

the fricative occurs between other speech sounds, and thus includes vocal tract dynamics.

Fricatives cannot be understood without understanding the fluid dynamics in the vocal tract, which are themselves not completely understood because they involve turbulent flow through a complex geometry. In addition, there is no complete theory of sound generation by turbulence. However, some general principles are known which makes this problem approachable.

When the flow passes through the constriction it greatly increases in velocity and becomes turbulent either around the constriction or when it strikes the obstacle. The unsteady turbulent flow generates sound in a number of ways. The dominant sound source typically comes from the forcing between the fluid and the obstacle it strikes (a dipole sound source). However, sound may also be generated by the fluid having an unsteady flow rate in the constriction (a monopole sound source) or by shear within the turbulence itself (a quadrupole source) [1, 2]. Once the sound is created it will propagate through the remainder of the vocal tract, being modified along the way by resonators such as the sublingual cavity, and then escape past the lips to free space where it may be detected by an ear or microphone. The location of the sound creation and the nature of the modification are still a matter of recent studies [3].

Part of the problem in understanding fricatives comes from our lack of understanding of sound. Sound is a component of fluid flow, and as such is described by the Navier-Stokes equations [4]. However, an unsteady flow also has pressure variations that respond to the changing momentum in the flow [5]. Such pressure fluctuations (often called pseudo-sound) make it difficult to separate the propagating sound field from the rest of the flow, and in fact there is no exact method known to isolate the sound component from the rest of the flow.

Despite this fundamental problem, we can still learn much about acoustics through computational aeroacoustics (CAA). In CAA, as in CFD, one simulates the Navier Stokes equations (or derived equations using simplifying assumptions), but in CAA one takes extra measures to insure that the sound field is adequately resolved and propagated. The particular challenges that CAA faces have been discussed in great detail [6, 7, 8], so here we'll consider the most relevant issues.

Resolving sound waves requires a wide scale of resolution. In speech, the frequencies 100Hz to 12000Hz are im-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.
OPAL-09, June 26-27, 2009, Vancouver, BC, CA.
Copyright remains with the author(s).

portant (but only a smaller part of that range is necessary); these have respective periods of 0.01s to 0.000083s, and respective wavelengths of 3.4m to 0.02833m in air at room temperature. The numerical method needs to adequately resolve these temporal and spatial scales. A high order finite difference scheme may be able to resolve a wavelength with 7 mesh points, but a second-order scheme common to CFD requires approximately 20 mesh points per wavelength [6], thus providing an upper limit on the highest frequency that can be resolved. Likewise, time has similar restrictions for adequate resolution. However, a finer resolution requires more RAM, more hard disk storage, and a longer simulation runtime.

Of course, one needs a method of actually measuring the sound simulated. The most obvious way is to read the pressure fluctuations in the compressible flow at the desired location. However, when the flow velocities are well below the speed of sound, the acoustic fluctuations in the flow are less than the non-acoustic pressures by orders of magnitude, thus making it a challenge to resolve the acoustic amplitudes. Upon recovering a signal that seems to be sound, one must be aware that it may contain numerical artifacts such as those caused by improper boundary conditions, or it may be pseudo-sound. One may also use the acoustic analogy [9] to find sound at a given location. An acoustic analogy rearranges the Navier-Stokes equations to appear like the inhomogeneous wave equation, but then makes approximations to treat the two sides of that equation as independent: the classical wave equation on one side and the sound source on the other (which is composed of terms from the Navier-Stokes equations). While these approximations allow for simpler acoustic calculations, they have a dubious basis [10, 11].

The boundary conditions of a standard CFD simulation are not adequate for CAA as they cause waves to reflect back into the domain. Therefore non-reflecting boundary conditions must be developed.

One cannot ignore the demands of CFD turbulence modeling for fricative simulations. Two common approaches to turbulence modeling are models using the Reynolds-Averaged Navier-Stokes (RANS) equations or using a Large Eddy Simulation (LES). The RANS equations are time-averaged and models such as the $k - \omega$ model complete the RANS equations using known statistical properties of turbulence. A LES uses a spatial filter based on the mesh, and models the subgrid scales of turbulence while resolving the larger scales. A LES generally requires a finer mesh than RANS simulations, thus will take much longer. If one wants to completely resolve the turbulence rather than model it, one must use an exceedingly fine mesh; this is called a Direct Numerical Simulation (DNS).

Despite these difficulties and numerous other challenges that accompany CAA simulations, they may provide excellent flow and sound data, making possible analysis that will

shed light on the underlying phenomenon, and is particularly useful in cases where experiments are difficult to perform, such as the human vocal tract.

Therefore, we seek to get a better understanding of fricatives using computational fluid dynamics. In particular, we use Fluent [12], a common CFD solver, to see if we can adequately simulate the English fricative /sh/. From there we intend to draw conclusions about the theory of /sh/ and other fricatives, and also to draw conclusions regarding the type and quality of simulations needed for fricative simulations.

We don't expect Fluent to be as efficient or accurate as specialized CAA code, but that it will provide a reasonable flow simulation from which we can learn about fricatives and proper simulation methods. We expect that 2D simulations won't be adequate to capture the essence of the flow and the sound, and that 3D large eddy simulations will be needed for accurate results.

2. Methods

To investigate the capabilities of CAA, we choose to compare with the 'level 3' experimental case of /sh/ that Shadle describes [13]. The advantage of this case is that Shadle provides experimental results for a fairly simple geometry, thus we can create a comparable simulation. The disadvantage, however, is that this is a simplified geometry, and Shadle concludes that the deviation of her results from recordings of spoken /sh/ are probably due to geometrical simplifications. Thus our test case, while modeling something close to the human /sh/, is expected to sound wrong as in Shadle's experiment.

In the hopes of finding a minimal yet adequate simulation method a variety of simulations are performed. We will perform 2D and 3D simulations, in both cases using a LES and a RANS turbulence model, the $k - \omega$ SST; chosen to combine the strengths of the $k - \epsilon$ and $k - \omega$ models. The resulting 'sound' from these simulations will be recorded using the direct method (measuring pressure 20cm from the mouth) and using an acoustic analogy on a variety of source surfaces.

The simulations are performed using the commercial CFD software Fluent. The advantage of using Fluent is that it is readily accessible CFD, but the disadvantages are that it is designed to be a robust and general CFD package; consequently it is slow, and accuracy is compromised for the sake of stability, nor does it contain high-order schemes and boundary conditions needed for high-performance CAA.

An image of the 2D domain can be seen in Figure 1, which is derived from Shadle's geometry. The 3D domain is extruded 25.4mm in the third dimension and narrowed at the constriction as Shadle did. The 2D domain has 71,216 cells with boundary layer cells as fine as $\Delta = 0.01$ mm and the coarsest cells being no larger than $\Delta = 2$ mm. The 3D domain has 1,262,021 cells, but computational limitations force the cells to be significantly larger, thus the finest cells

were $\Delta = 0.2\text{mm}$ at the most critical flow regions and the largest cell being no larger than $\Delta = 4\text{mm}$. The inlet is defined as a mass flow inlet (alternates were later tested in 2D—see below) with a flow rate of $0.000804[\text{kg/s}]$ in 3D and $0.0316535[\text{kg/s}]$ in 2D (that is, the 3D rate with 0.0254m divided out). To attain non-reflecting boundary conditions, a buffer zone was created inside the pressure outlet which gradually damps the waves. Such a condition can itself cause reflections if not done gradually [7], but by trial and error adequate performance was found by damping pressure according to:

$$P = P - F \cdot (P - P_0) \quad (1)$$

where

$$F = \frac{R - r}{R}, r < R \quad (2)$$

where background pressure $P_0 = 0$, r is the distance from wall, and the damper width $R = 10\text{cm}$. This buffer region was implemented in Fluent with user-defined functions. It is worth noting that Fluent does provide non-reflecting boundary conditions, but they did not work with the settings required for this simulation.

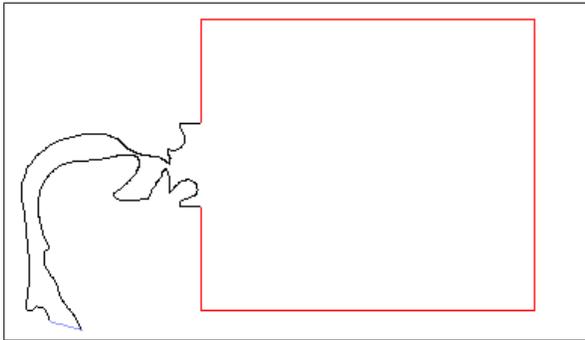


Figure 1. The 2D domain. The inlet is the paler line in the bottom left. The outlet is the rectangle enclosing the free space beyond the lips

The simulations were run with a constant time step of 0.00001s until the spectrum became fairly steady. The spatial and temporal integration schemes are all 2^{nd} order accurate. The flow is compressible, as is required to directly measure the sound waves. For all simulations the acoustic analogy data is recorded concurrently with the direct method, thus the two cases are comparable. The pressure probes are 20cm from the lips.

The sound reported by an acoustic analogy is calculated from the flow coincident with the source surface. One may consider different source surfaces, and thus investigate the contribution of each source surface upon the final sound. However, the acoustic analogy as implemented in Fluent, is only relevant for sound propagating to free space, thus an acoustic analogy result of sound created at the constriction will not consider the modifications that the sound will undergo between the constriction and when it escapes beyond

the teeth and lips. Thus we can find the sound as generated by the source surface and compare it with the sound recorded by the direct method to understand the contribution of that source surface to the final sound (Shadle did a similar study of coherence to find the important source surface). The source surfaces used are: the constriction, the cavity, the lower tooth, the upper tooth, the lower lip, and the upper lip.

The spectral analysis uses a 0.02048s hanning window (2048 data points). Ideally, these simulations would obtain about 5s of data (thus $500,000$ time steps) which allows averaging many spectra to obtain a smooth spectrum, but such a long run time isn't feasible, so we obtain an averaged-spectrum with the signal, but also apply a smoothing algorithm to supplement the averaging. Figure 2 shows an unsmoothed and unaveraged fft compared with a smoothed and averaged fft to demonstrate the ability of this method to capture the essence of the spectrum.

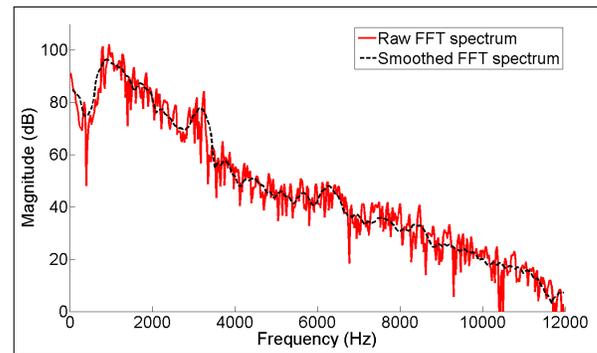


Figure 2. Spectrum processing.

Our temporal and spatial resolutions are lower than those recommended in the theory section. The time step we used only allows for 10 samples per period of a $10,000\text{Hz}$ wave, and the largest cell size in 3D only allows for 8.5 samples per wavelength of a $10,000\text{Hz}$ wave. While it is true that these resolutions are well below the desired level, in some simplified tests we found them adequate for wave propagation over short distances. Thus, this simulation should adequately resolve the desired scales, but with the warning that the higher frequencies are not as well resolved as would be hoped.

When a simulation starts, just like a physical flow, it takes some time to reach its steady state. When the flow is unsteady as these simulated flows are, the flow will never reach a steady state, but it will reach a statistically steady state, which occurs when the long-term average flow is steady though it contains unsteady fluctuations. In this study, a statistically steady state is judged from the spectra. In Figure 3, one may see that the first spectrum (2048 samples) varies significantly from the last spectrum of the signal. One may also see the spectra that result from the averaging and smoothing (as described above), and the first spectrum that is considered to be statistically steady. Ideally the statisti-

cally steady spectrum would be taken after a longer time, but the runtime of the simulations limits this greatly.

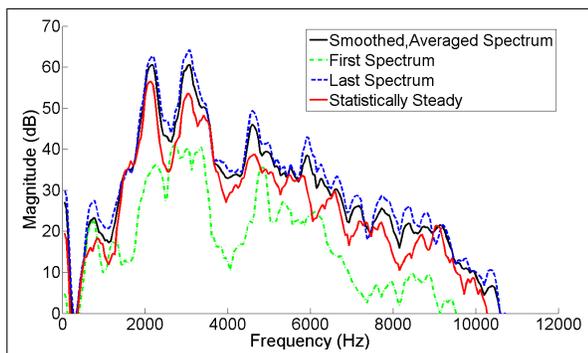


Figure 3. Determining a statistically steady simulation

3. Results

Before examining the details of the spectra from the various simulations, it is interesting to compare instances of the flow from 2D and 3D simulations. First, a pressure snapshot with the corresponding velocity snapshot from a 2D simulation is shown in Figure 4. One can observe the non-reflecting boundaries washing out the acoustic waves and the unphysically high pressures of the sound waves. The snapshots from the 3D simulations are shown in Figure 5. The jet and high-velocity flow is much more localized, and the acoustic waves have more physically realistic values.

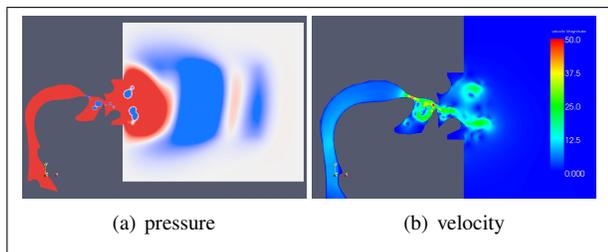


Figure 4. A 2D flow snapshot (pressure ranged from -15pa to 15pa so that acoustic waves may be seen, velocity ranged from 0 to 50m/s)

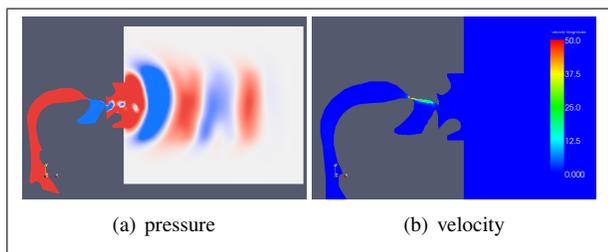


Figure 5. A 3D flow snapshot (pressure ranged from -1pa to 1pa so that acoustic waves may be seen, velocity ranged from 0 to 50m/s)

We can consider the results from 2D simulations, which are shown in Figure 6. From the first simulations, it quickly

became clear that neither the mass flow inlet nor RANS simulations gave reasonable results. Thus we focused upon using a pressure inlet at a constant 800Pa rather than a mass flow inlet, and we used either LES or no turbulence model rather than RANS. We also investigated the acoustic analogy in the case of an incompressible flow. The acoustic analogy results have the best agreement to Shadle’s data, while the direct measurements show little correlations and unreasonable amplitudes, and the simulations using RANS or a mass-flow inlet have a spectrum that doesn’t even seem to be a legitimate broadband noise signal. It is important to note that, while 800Pa is a reasonable pressure in speech, it cannot be considered an exact comparison to Shadle’s test case because pressure can vary at a mass flow inlet, and vice versa, thus one cannot define an inlet of one type that is exactly equivalent to the other. The amplitudes will be discussed later in more detail.

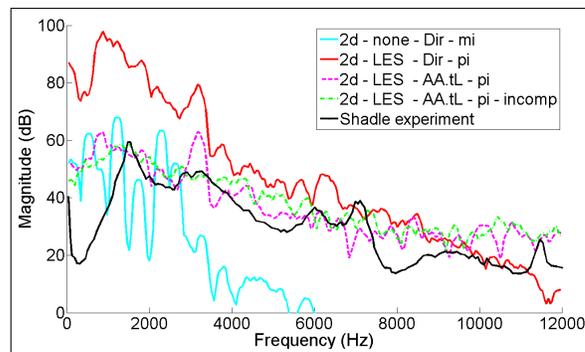


Figure 6. 2D simulations. None = no turbulence model, Dir = direct method, AA.tL = acoustic analogy from lower tooth; pi = pressure inlet; mi = mass inlet; incomp = incompressible

Next we can consider the results from the 3D simulations, as shown in Figure 7. In general, one may notice that the acoustic analogy and direct measurements match each other much better than in the 2D case, and also match Shadle’s experiments better, though they are still significantly different. One may also observe that, like the RANS in 2D simulations, the 3D RANS yields a signal that is unphysical. While the 2D results are largely limited by an unphysical geometry and 2D turbulence (which is fundamentally different from 3d turbulence), the 3D simulations are limited by the mesh being too coarse. We therefore sought to refine the mesh in the areas critical to the flow and observe how such refinements alter the spectra. This is not a proper mesh refinement study with which to observe the convergence of the numerical solutions [14] as such a grid refinement requires all parts of the mesh to be refined. However, this refinement does indicate how the spectrum changes with better flow resolution. Not surprisingly, the refined grid varies more in the high frequencies because a finer grid can better resolve smaller sound waves. These results are included in Figure 7.

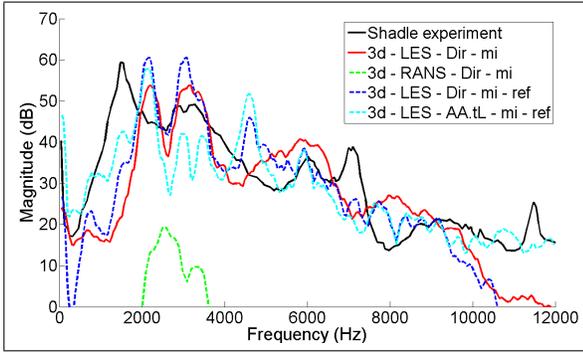


Figure 7. 3D simulations. Ref = refined

One may consider the acoustic analogy on two levels. First, one may compare the results from the AA to the results using the direct method, as a measure of error. Second, one may assume that the AA results are perfect, and use the results to study the coherence between the source surface and the final sound heard in the far field. We know this assumption isn't true, but one may still consider it to get a general feeling for the contributions of each sound source. To these ends, we present the acoustic analogy results from the same simulation but different source surfaces, and compare them with the sound recorded by the direct method, which can be seen in Figure 8.

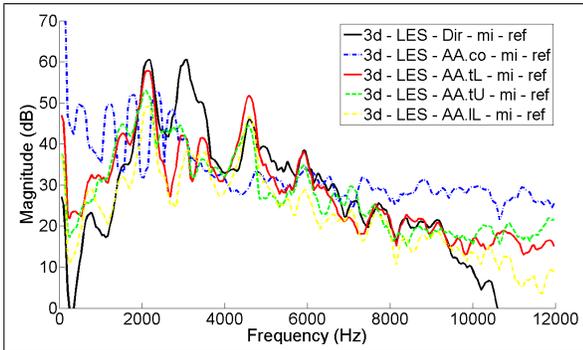


Figure 8. Acoustic analogy results. The acoustic analogy locations are: co = constriction, tL = lower tooth, tU = upper tooth, IL = lower lip

Finally, Figure 9 compares the best 2D simulation results with the best 3D simulation results, as a side by side comparison.

4. Discussion

The results from the RANS simulations and the 2D mass-flow inlet simulations appear to be unphysical upon first glance, but it is good to have an objective reason why they should be discarded. When sound propagates through the vocal tract, it encounters resonators such as the cavity below the tongue and the cavity between the lips and teeth, which will resonate with a frequency range, and thus cause a distinctive peak in the spectrum [1]. The spectra from all the

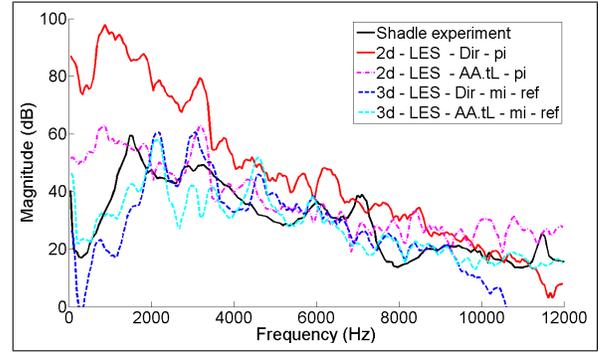


Figure 9. Comparison of best results

2D mass-flow inlet and RANS simulations don't have the characteristics of broadband noise with a few distinguished peaks which might correspond to cavities in the vocal tract, thus they are discarded as clearly violating the physics of this flow. The 2D simulations with a pressure inlet do have peaks which may represent resonance with the cavities, and the 3D results clearly do, though it is questionable how well they match the experimental results.

Both 2D and 3D simulations measure the sound at the same location, and Shadle's data is scaled for distance, but the amplitudes should be viewed with some caution. The 3D mass flow inlet was designed to match Shadle's 670 [cm³/s] volume flow rate by assuming incompressibility at the inlet. The 2D mass flow inlet was scaled to match the 3D rate by dividing the third dimension out (2.54cm deep), but most simulations used a pressure inlet instead. Also, some ambiguity also comes from the constriction. Shadle formed a narrow constriction by filling the third-dimension with clay. The constriction shape of the 3D geometry was estimated from Shadle's description, but the 2D simulation cannot include the narrowing in the third dimension at the constriction. As a consequence the 3D constriction area is about 0.4% of the inlet area, while the 2D constriction area is about 9.6% of the inlet area (in 2D, it is really a length rather than an area). Thus the velocity increase at the constriction is not expected to be the same, and there remains some ambiguity between the 3D simulation geometry constriction shape and the experiment. As one might expect, the velocity at the 3D constriction was observed to be much higher than at the 2D constriction, yet the sound from the 2D simulations has much higher amplitudes. This is attributed to the inability of the 2D equations to describe the energy dissipation that occurs in 3D turbulence [15], and thus is a fundamental shortcoming of 2D simulations.

One issue that is difficult for simulations to handle is the material properties. Shadle's experiments used plexiglass and clay to form the constriction, while simulations using basic wall boundary conditions will treat all walls as an acoustically hard surface. This may have caused discrepancies between these results and Shadle's, making this validation less certain. However, this will be a bigger concern

when trying to simulate a true fricative, because the flesh walls of a true vocal tract will increase the bandwidth of the cavities and cause energy losses as a function of the frequency [1]. To simulate this properly would require specialized boundary conditions at all of the walls.

From Figure 8 one may observe that the acoustic analogy using either the upper tooth, lower tooth, or the lower lip as the source surface matches the direct recordings quite well. From this one might draw some useful conclusions. First, the sound is very close to its final form at the teeth and the acoustic analogy is able to capture this. Second, because the acoustic analogy can replicate the sound from direct measurement (in 3D simulations), there is little need to extend the domain far beyond the lips. One can consider a fictitious source surface just outside of the lips and propagate the sound to the far field. This allows the domain and nonreflecting boundaries to be much smaller.

Though the acoustic analogy does fit the direct measurement quite well, there are two exceptions worth noting. First, from about 9500Hz and above the direct method spectrum drops in amplitude while the acoustic analogy stays roughly the same. This is quite likely an indication that the mesh and time step were too coarse to adequately resolve those frequencies, thus one should trust the acoustic analogy results more. This failure in the higher frequencies was forecast and discussed in Methods. Second, there is a distinct peak at 3000Hz which the direct method picked up but none of the acoustic analogy surfaces recorded. This peak may represent a quadrupole sound source, which is sound created by stresses within turbulence rather than sound created by turbulence interacting with a surface. The AA source surfaces will not account for this quadrupole noise because they are calculated from an impermeable source surface [12]. However, in such a flow the quadrupole contributions are expected to be small, and it may be that this peak is a numerical artifact, such as domain resonance.

In this study the recordings were taken at 20cm, but for the sake of investigating the far field in the simulation results, Figure 10 shows a measurement at 10cm compared with the measurement at 20cm. One may note that from about 4500Hz and above, the two spectra stay close to parallel, an indication that at 10cm from the mouth, this frequency and higher ones can be considered in the far field. The frequency 4500Hz has wavelength $\lambda = 7.5$ cm, thus this is slightly more than one wavelength distant from the source, but still not an unreasonable estimate of the far field. If treated as a point source, the decibel amplitude of sound pressure should decrease 6dB for each doubling of distance from the source. In Figure 10, the difference is around 7dB which is still reasonably close to the expected value.

In all the methods used, the $k - \omega$ SST model failed to find a reasonable spectrum. Also, an initial simulation with the $k - \epsilon$ turbulence model showed a similar behavior to the $k - \omega$, which speaks against the usefulness of RANS

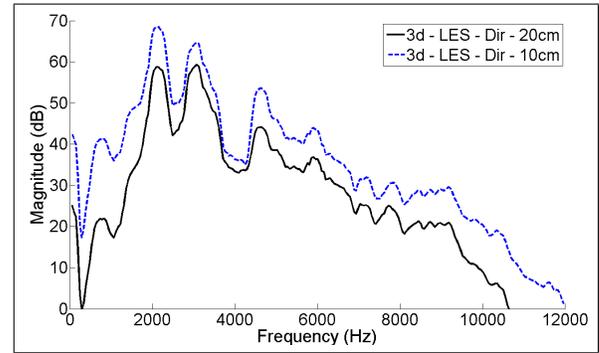


Figure 10. Investigation of far field location

turbulence models in CAA simulations. This isn't surprising because RANS equations are time-averaged and thus smear out acoustic signals.

To give an idea of simulation runtimes, the 3D RANS simulation took 338s per time step and the LES took 304s per time step on the same mesh (both parallel processing on 3 cores). The 2D LES on a single core took 15.9s per time step. Running a 2D simulation without a turbulence model offered a large increase in speed, while changing the flow from compressible to incompressible offered a smaller speed increase.

In 2D we ran simulations using the LES turbulence model, and with no turbulence model. A comparison of these simulations is shown in Figure 11. Using no turbulence model is presumably a DNS, which requires a very fine mesh. While the 2D mesh was not fine enough for a DNS, it is worthwhile to note that the spectra between these two simulations are very similar. This implies that the subgrid turbulence model and the wall model in the 2D LES have a minimal contribution to this flow, and one may consider not running a turbulence model at all for significant speedup.

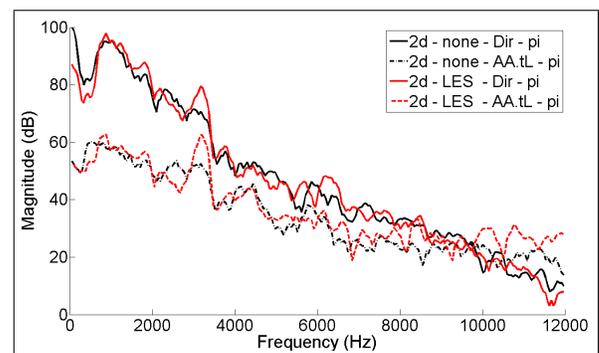


Figure 11. Comparison of LES with no turbulence model

While these simulations are primarily compared to Shadle's experiments, we can still make a statement concerning 2D simulations of the true geometry. Because the 2D geometry is derived from a mid-sagittal X-ray of the vocal tract, and because a 2D simulation can never include the true 3-dimensionality of the vocal tract, we may consider the 2D simulation geometry as good as it can get. Thus it is reason-

able to compare these results not just with Shadle’s experiment, but with a true /sh/, which is shown in Figure 12. The simulation results bear little resemblance to the spoken /sh/; however, one might also note how little resemblance Shadle’s spoken /sh/ has to the experimental /sh/ and in comparison with Fant’s /sh/ (as given in [16]).

One challenge is knowing how to quantitatively compare simulation results to an experiment or to real life, especially considering the great variation that is seen in speech. Before we can judge how well a simulation has performed, we must know what the defining characteristics of /sh/ are that we seek to reproduce, which is a challenge that extends beyond comparing spectra. If one can attain successful fricative simulations on a static geometry, then another great challenge is simulating the dynamics of speech.

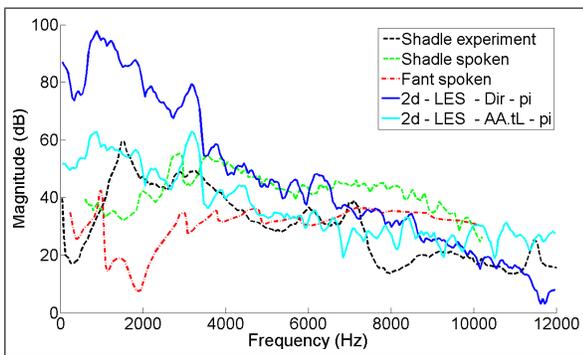


Figure 12. Comparison of 2D simulation with spoken /sh/

5. Conclusion

While these simulations don’t provide as close a fit with Shadle’s experimental data as hoped, numerous observations have been made concerning the strengths and weaknesses of the simulations which may be applied in future attempts to model fricatives computationally. First, they demonstrate the inability of the $k - \omega$ sst model (and quite likely all RANS models) to find a reasonable spectrum. Second, they demonstrate the superiority of 3D simulations to find a physically reasonable spectrum. Third, they demonstrate how suitable non-reflecting boundaries may be created in Fluent. They also show that an acoustic analogy may offer reasonable results, and can likely be used to simplify the computational domain in future simulations.

From the observations of these simulations, we can make recommendations for future fricative simulations. First, a 3D geometry should be used, but the domain can be significantly truncated. The crucial flow features occur at the constriction in the vocal tract, thus one might start the domain further up the from the vocal chords allowing just enough distance between the inlet and the constriction for the flow to fully develop. The domain can also be truncated beyond the lips. Soon after the lips the sound can be considered to come from a simple sound source and can be propagated

to a further distance using a theoretical approach. Truncating the domain will save many mesh cells; however, those savings should be used toward obtaining a better flow resolution around the constriction and the teeth. Because the constriction creates a strong jet down the mid-sagittal plane, the highest flow gradients and important flow features occur here. Thus one should concentrate more cells in the mid-plane of the domain. The wall should be meshed in much finer detail, preferably enough to resolve the boundary layer without a wall function (see [17, 12] for further discussion). Ideally, such a simulation would be done with specialized CAA code. Non-reflecting inlet and outlet boundaries using a sophisticated method such as those discussed in [6, 7] should be implemented, and will be of smaller computational expense than the large damper used in this study.

In such a simulation, one might consider numerous AA surfaces. Rather than including the whole tooth or lip as a source surface, one might divide these surfaces into small sections to investigate the dipole source locations in finer detail. Also, it would be helpful to place a permeable source surface in front of the lips to account for quadrupole sources. A carefully developed simulation with these characteristics should improve upon the simulations presented in this study, and is a recommended next step.

6. Acknowledgements

We would like to thank the Artisynth project for support, and Donald Derrick for his discussions.

References

- [1] K. Stevens, *Acoustic Phonetics*. Cambridge: The MIT Press, 2000.
- [2] M. J. Lighthill, “The Bakerian Lecture, 1961. Sound Generated Aerodynamically,” *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 267, no. 1329, pp. 147–182, 1962.
- [3] M. S. Howe and R. S. McGowan, “Aeroacoustics of [s],” *Proceedings of the Royal Society a-Mathematical Physical and Engineering Sciences*, vol. 461, no. 2056, pp. 1005–1028–, 2005.
- [4] D. G. Crighton, “Acoustics As A Branch Of Fluid-Mechanics,” *Journal of Fluid Mechanics*, vol. 106, no. MAY, pp. 261–298–, 1981.
- [5] J. E. F. Williams, “Hydrodynamic Noise,” *Annual Review of Fluid Mechanics*, vol. 1, pp. 197–&–, 1969.
- [6] C. K. W. Tam, “Computational aeroacoustics: An overview of computational challenges and applications,” *International Journal of Computational Fluid Dynamics*, vol. 18, no. 6, pp. 547–567–, 2004.
- [7] T. Colonius and S. K. Lele, “Computational aeroacoustics: progress on nonlinear problems of sound generation,” *Progress in Aerospace Sciences*, vol. 40, no. 6, pp. 345–416–, 2004.
- [8] M. Wang, J. B. Freund, and S. K. Lele, “Computational prediction of flow-generated sound,” *Annual Review of Fluid Mechanics*, vol. 38, pp. 483–512–, 2006.

- [9] M. J. Lighthill, "On Sound Generated Aerodynamically. I. General Theory," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 211, no. 1107, pp. 564–587–, 1952.
- [10] A. T. Fedorchenko, "On some fundamental flaws in present aeroacoustic theory," *Journal of Sound and Vibration*, vol. 232, no. 4, pp. 719–782–, 2000.
- [11] C. K. W. Tam, "Computational aeroacoustics examples showing the failure of the acoustic analogy theory to identify the correct noise sources," *Journal of Computational Acoustics*, vol. 10, no. 4, pp. 387–405–, 2002.
- [12] Fluent, *Fluent 6.2 Documentation*. Fluent, 2005.
- [13] C. H. Shadle, "Articulatory-Acoustic Relationships In Fricative Consonants," in *Speech Production and Speech Modelling*, pp. 187–209–, Netherlands: Kluwer Academic, 1990.
- [14] P. J. Roache, "Quantification of uncertainty in computational fluid dynamics," *Annual Review of Fluid Mechanics*, vol. 29, pp. 123–160, 1997.
- [15] P. K. Kundu and I. M. Cohen, *Fluid Mechanics*, vol. 2. San Diego: Academic Press, 2002.
- [16] G. Fant, *Acoustic Theory of Speech Production*, vol. 2. The Hague: Mouton, 1970.
- [17] S. B. Pope, *Turbulent Flows*. Cambridge University Press, 2000.