

Study of long-range infrasound propagation with high-performance numerical schemes applied to the Euler equations

O. Marsden*, L. Vayno†, C. Bogey‡, C. Bailly§
Laboratoire de Mécanique des Fluides et d'Acoustique
Ecole Centrale de Lyon & UMR CNRS 5509
36 avenue Guy de Collongue, 69134 Ecully cedex, France

1 Abstract

Numerical simulations of long-range infrasound propagation accelerated with commodity graphics hardware are reported. The test case used in this work involves a realistic atmosphere including temperature gradient effects. A low-dispersion and low-dissipation finite-difference algorithm initially developed for aeroacoustics applications is used to solve the full Navier-Stokes equations, and its performance is compared when run on a desktop processor, and when run on a mainstream graphics card.

2 Introduction

The full two-dimensional Navier-Stokes equations are solved to model the propagation of low-frequency sound waves through the atmosphere. The atmosphere is modeled from ground level to a height of 150 km. Its mean structure is constructed from an experimental temperature profile which has been fitted by a high-order polynomial function of altitude[1, 2, 3]. Finite differences based on eleven-point stencils are used to compute the derivatives involved in the Navier-Stokes equations. Time integration is performed with a six-stage optimized Runge-Kutta scheme. Additionally, a shock-capturing filtering

*assistant professor, email: olivier.marsden@ec-lyon.fr

†Masters student

‡research scientist, AIAA member

§professor at ECL & Institut Universitaire de France

technique is employed in order to handle the discontinuities that appear in the vicinity of shock waves.

The effect of source amplitude on long-distance propagated time signals is examined. Source strength appears to have very large effects on eigenrays that propagate through the thermosphere, not only in terms of amplitude, but also in terms of arrival time.

The numerical algorithm has been implemented in a functionally equivalent manner on a standard cpu with the FORTRAN language, and on a commodity graphics card with the AMD-enhanced version of the BROOK stream language. The graphics card computation is found to be approximately 18 times faster than one performed on the cpu.

3 Numerical algorithm

The full 2-D Navier-Stokes equations are solved on a regular Cartesian grid with an optimized high-fidelity numerical procedure based on explicit spatial finite differences and Runge-Kutta time integration.

Away from boundaries, spatial discretization is performed with explicit fourth-order eleven-point centred finite differences optimized to minimize dispersion for wavenumbers discretized by between 4 and 32 grid points[4]. Close to boundaries, be they solid walls or radiation conditions, optimized explicit non-centred differencing schemes are used[5]. The non-centered differencing schemes are all based on eleven-point stencils, including the one-sided stencil used for wall points. Time integration is performed with a six-step second-order optimized low-storage Runge-Kutta algorithm[4]. Characteristics regarding dispersion and dissipation for the spatial differencing schemes, filters, and the time integration scheme can be found in previous papers [4, 5]. This solver has been validated on reference viscous flow configurations and on a demanding multi-body acoustic scattering test-case, yielding results in good agreement with experimental and analytical data[6, 7].

The schemes' properties mean that the behaviour of waves discretized by at least four points per wavelength is accurately reproduced, with very low levels of dispersion and dissipation, for frequencies such that $\omega\Delta t \leq 1.25 \times \pi$. The determination of the computational time step is based on a CFL criterion, $CFL = c_{max}\Delta t/\Delta x = 0.75$, where c_{max} is the largest value of the speed of sound in the atmosphere modeled here. The ground is modeled as a non-slip boundary condition based entirely on the use of the non-centered schemes described previously. The wall-point density and pressure are updated by resolving the governing equations in which all terms containing wall velocities are eliminated, thanks to high-order non-centered differencing schemes.

Spatial filtering is carried to ensure stable computations. An explicit 11-point filtering stencil is designed to remove fluctuations discretized by less than four grid points per wavelength, while leaving all larger wavelengths effectively untouched[4]. As the differencing schemes used near boundaries are asymmetric, their effective wavenumbers have an imaginary part which leads to them being unstable for very high frequencies[8]. It is

therefore essential to use them in conjunction with appropriate highly selective filters, and to this end, we use the filters described in Berland *et al*[5], which also selectively damp fluctuations with fewer than four points per wavelength. Filters for grid points more than two points away from a boundary are built on eleven-point stencils, while the last and last but one point stencils up to two points away from a boundary are built respectively on four and seven points. Thus at the wall, in the x direction, the centred eleven-point filter is used, whereas in the y direction the family of non-centred filters is applied. At the lateral radiation boundaries in the x and y directions, Tam and Dong’s 2D far-field radiation condition [9] is used.

Examples of 2-D acoustic diffraction and aeroacoustic flows successfully simulated with the solver described in this work can be found in Marsden *et al*[6].

4 Non-linear shock-capturing filtering

Non-linear effects can be important in long-range atmospheric propagation, due both to the sizeable propagation distances and to large relative amplitudes which are reached in the high atmosphere. In particular, it is surmized that shock fronts may be formed in the thermosphere. This aspect poses potential problems for standard finite-differences time-domain acoustic solvers, which are not designed to cope with steep wave fronts and which can lead to Gibbs oscillations.

The CFD community has been dealing with shocked flows for a long time, and has developed a variety of different techniques to avoid numerical problems associated with the presence of shocks. Standard shock-capturing schemes developed for CFD (TVD, ENO, WENO) are however ill-suited to time-dependent problems, because they exhibit poor spectral accuracy, and tend to be excessively dissipative, particularly in the context of long-distance propagation. Hence in this work we test a non-linear filtering method designed with acoustics in mind[10].

The methodology consists in applying a second-order conservative filtering only where necessary, *i.e.* only in zones where non-linear effects are high. Understandably, much of the method’s properties come from the non-linear detection algorithm. Non-linear zones are identified thanks to a Jameson-like detector based on pressure fluctuations. The first step consists in extracting the high wavenumber components from pressure fluctuations. This is done by applying a second order filtering, as described in the following equation for grid point i , in the x direction:

$$Dp_i = (-p_{i+1} + 2p_i - p_{i-1})/4$$

and then defining the high-pass filtered squared pressure fluctuation as

$$D_i^2 = \frac{1}{2}[(Dp_i - Dp_{i-1})^2 + (Dp_i - Dp_{i+1})^2]$$

This squared pressure fluctuation is used to compute define a sensor value as

$$r_i = \frac{D_i^2}{\bar{p}_i^2} + \epsilon$$

where ϵ is a small parameter, typically 10^{-16} , whose role is to avoid numerical problems when dividing by r_i , as will be seen subsequently, and \bar{p}_i is the averaged pressure at point i .

The self-adjusting filtering strength σ_i at grid point i is computed according to

$$\sigma_i = \frac{1}{2} \left(1 - \frac{r_{th}}{r_i} + \left| 1 - \frac{r_{th}}{r_i} \right| \right)$$

where r_{th} is a threshold parameter which adjusts the extent of the zone over which filtering is applied, whose value is $r_{th} = 10^{-5}$. This filtering strength has the desired properties of being equal to zero away from shocks, where $r_i < r_{th}$, and of increasing towards a value of 1 for increasing shock intensities.

Conservative variables are filtered conservatively, *i.e.* the filtered term is computed as a difference of two fluxes, as follows:

$$U_i^f = U_i - \left(\sigma_{i+\frac{1}{2}} F_{i+\frac{1}{2}} - \sigma_{i-\frac{1}{2}} F_{i-\frac{1}{2}} \right)$$

where $\sigma_{i+\frac{1}{2}}$ is simply the average of previously calculated filtering strengths σ_i and σ_{i+1} , and $F_{i+\frac{1}{2}} = \sum_{j=1-n}^n c_j U_{i+j}$ and $F_{i-\frac{1}{2}} = \sum_{j=1-n}^n c_j U_{i+j-1}$ are the up-winded and down-winded fluxes based on a dissipative second-order filter c_j . A spectral analysis of the shock-capturing treatment is in Bogey *et al* [10].

The effect of the non-linear filtering technique described above is examined on a simple configuration of atmospheric acoustic propagation, designed to induce non-linear effects which are expected to occur in long-range propagation phenomena.

The atmosphere is designed to resemble the Misty Picture configuration[1, 2]. A semi-experimental temperature profile \bar{T} based on measurements up to an altitude of 90 km and on an empirical model above, is fitted by a tenth-order polynomial, and the resulting polynomial is used to obtain the pressure profile from the hydrostatic equilibrium equation

$$\ln(\bar{p}/p_0) = -\frac{gM}{R} \int_0^z \frac{1}{\bar{T}} dz$$

The experimentally obtained temperature profile is highly fluctuating, and exhibits negative values of the Väisälä-Brunt frequency N , defined by $N = \text{sign}(N^2) \times |N^2|^{1/2}$ where

$$N^2 = -\frac{g}{\bar{p}} \frac{d\bar{p}}{dz} - \frac{g^2}{c^2}$$

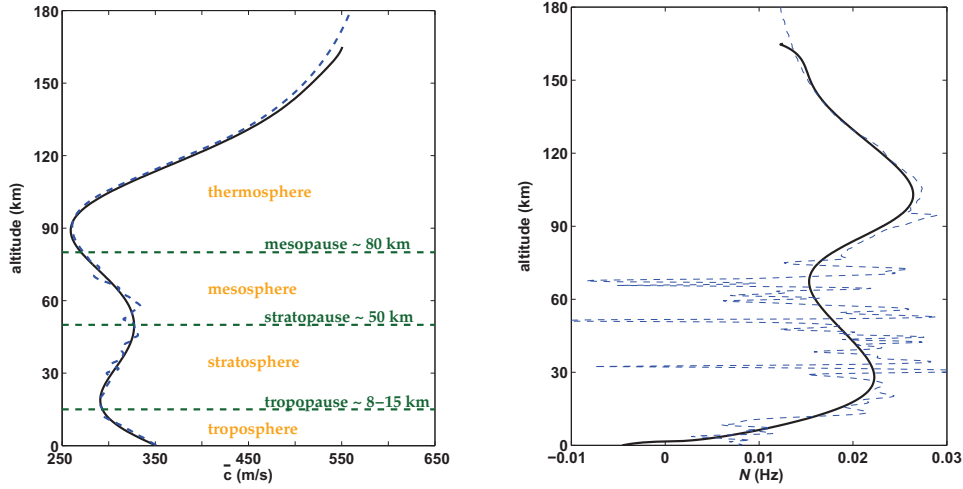


Figure 1: Sound speed and Väisälä-Brunt frequency as a function of altitude: - - - experimental/statistical data, — numerical profiles

as shown in Figure 1.

The use of a high-order polynomial allows the construction of a mean atmospheric profile which is both realistic and stable, *i.e.* $N \geq 0$. It also removes the short-wavelength fluctuations in the sound speed profile, and the acoustic scattering which is associated with them. While this removal is unphysical, it has the desirable effect of rendering acoustic results less complicated to analyse.

An evaluation of the non-linear filtering methodology is carried out by studying the vertical propagation of a low-frequency high-amplitude acoustic signal in the above-described atmospheric model. The source is described by its amplitude as a function of time

$$S(t) = \frac{1}{2} \mathcal{A} \sin(\omega_s t) \left[1 - \cos\left(\frac{\omega_s t}{2}\right) \right]$$

where $f_s = 0.1$ Hz, and by a Gaussian spatial envelope of half-width 900m. It is introduced as a forcing term on the energy equation. The source's amplitude is fixed at $\mathcal{A} = \bar{p}_{source}/30$.

Three computations are performed: one loosely referred to as the reference computation, performed on a finer grid ($\Delta y' = \Delta y/3$) and non-linear filtering applied with a coefficient of 0.01, one performed on the standard grid with no non-linear filtering and a second one performed on the standard grid with the non-linear filtering fully applied. Their results are compared at three different times in Figures 2 (a) (b) and (c).

These plots show vertical cuts of $p'/\sqrt{\bar{p}}$ as a function of altitude, at $x = x_{source}$. In Figure (a), first signs of wavefront steepening can be observed, due to the high source amplitude. There is no notable difference between the three sets of results, as the signal is still properly resolved by the numerical scheme on the coarse grid.




Figure 2: Computed signals at (a) $t = 10$ s, (b) $t = 30$ s and (c) $t = 70$ s.
 - - - reference computation on finer grid ($\Delta y/3$), — computation with non-linear filtering,
 — computation without non-linear filtering

In Figure (b), the reference computation (in red) shows a small leading N-wave and a well-defined higher-amplitude central N-wave. Gibbs-type oscillations can be seen in the vicinity of the shocks, but their amplitude is sufficiently small so as not to destabilize the computation. The coarse-grid solution obtained with the non-linear filtering (in blue) is in very good overall agreement with the reference solution. Both the small leading N-wave and the larger central N-wave are correctly placed, indicating that the filtering procedure does not modify shock propagation velocity. Gibbs oscillations are effectively removed. The peak amplitude of the shocks is slightly lower than that of the reference peaks, but this is an unavoidable consequence of any filtering procedure. The coarse-grid unfiltered solution (in black) is similar to the reference solution in overall shape, but exhibits strong oscillations around it. These oscillations can be interpreted as strong Gibbs oscillations, which are not properly resolved by the numerical scheme and which therefore are not propagated at the correct velocity. Strong numerical diffusion results from this behaviour.

In the final Figure (c) two well defined N-waves can be seen. The first one is the result of the fusing of the initial wave packet, while the second one results from the reflection of the source on the ground, whose positive front has had the time to catch up with the negative tail of the direct signal. The three computations show reasonable overall agreement. However, the most noticeable point regarding the unfiltered signal is that the first positive front is incorrectly placed. This can be explained by noting the strong overshoots in Figure (b) which will travel faster than the reference shock front. In any case this highly undesirable behaviour appears to be corrected by the use of the non-linear filtering procedure.




Figure 3: Pressure fluctuations p' normalized by $\sqrt{\bar{p}}$ at $t = 1320$ s. Linear ray-traced results have been superposed in gray, and eigenrays traced in black.

In summary, a self-adjusting non-linear filtering methodology has been briefly tested for long-distance acoustic propagation computations with the Euler equations. The filtering technique successfully removes Gibbs oscillations, which can be a numerical requirement for computations dealing with strong shocks, and as such can be regarded as *shock-capturing*. Moreover it enables the proper computation of shock-front velocities for relatively poorly discretized waves. This last point is essential if Euler computations are to be used to study arrival times of acoustic signals over long propagation distances.

5 Source amplitude effects on misty picture computation

The numerical algorithm described in the previous sections is used to perform a preliminary investigation of the effect of source amplitude on time signatures for a long-distance propagation example. The Misty Picture configuration is again used for this study. The previously described source is used with a central frequency of $f_s = 0.1\text{Hz}$, and is arbitrarily placed at a height of 1.5 km above the ground. The computational domain spans 450 km in the x direction and 160 km in the vertical direction, and is discretized with a spatial step of $\Delta x = \Delta y = 300\text{m}$. A CFL value of 0.55 is used, which yields a time step of $\Delta t = 0.55\Delta x / \max(c_0) = 0.3\text{s}$ based on the highest speed of sound, found in the thermosphere. Time signals are recorded at a distance of 400 km from the source, and compared as a function of source amplitude.

Figure 3 shows an example of a snapshot of the acoustic field normalized by the square

root of density fluctuations, at $t = 1320$ s, which corresponds to the arrival time of the shortest eigenray at the measurement location $x = 400$ km. The source amplitude in this case was $\mathcal{A} = 1 \times 10^{-3} \bar{p}_{source}$ where \bar{p}_{source} refers to the average atmospheric pressure at the source location. Linear rays are superposed in grey, and eigenrays are represented as black lines, and numbered according to their average arrival order. Strong caustics due to local variations of sound speed are clearly visible at a height of 30 km and again at a height of 110 km. An acoustic shadow zone can be observed roughly between 30 and 270 km away from the source. This view of the acoustic fluctuations illustrates well the fact that a simple initial source in a smooth atmospheric profile can yield a complex acoustic pattern downstream from the first shadow zone.

The aim of this study is to look at the effect of the source amplitude on the individual elements of the time signature. To this end, a parametric series of computations is performed for amplitudes ranging from $\mathcal{A} = 5 \times 10^{-8} \bar{p}_{source}$ to $5 \times 10^{-2} \bar{p}_{source}$ or in physical terms approximately 0.005 to 5000 Pa. The lowest source amplitude in this work is just above the threshold level of uncertainties introduced into the computation by aspects such as radiation boundary conditions and the reflecting ground condition.

The time signatures obtained over the range of source amplitudes vary considerably. This is illustrated in Figure 4 which shows the signals obtained from the highest source amplitude $\mathcal{A} = 0.05 \bar{p}_{source}$ as well as from $\mathcal{A} = 5 \times 10^{-7} \bar{p}_{source}$. The different amplitude scales (left and right) show that the ratio of source amplitude to measured signal strength at 400 km is highly different. Indeed, for the strong source this ratio is worth around 0.01, whereas for the weak source the ratio attains a value of roughly 0.2. The difference in shape between the two signals is also very noticeable. The maximum value for the strong source signal corresponds to the initial eigenray (1) arrival, in contrast with the weak source where the maximum is reached for the last (longest) eigenray (3) arrival. This variation is the result of two different causes whose effects are compounded. Firstly, the relative attenuation σ of the stratospheric eigenray (1), which is roughly constant ($\sigma = \max(p'_{eig})/\mathcal{A} \simeq 0.02 \sim 2D$ geometric attenuation) for all source amplitudes lower than about 50 Pa, diminishes for higher source amplitudes and reaches a value of $\sigma = 0.01$ for the strongest source, due presumably to a combination of dissipation in the shocked propagation and to numerics. Secondly, for low source amplitudes, an amplification mechanism appears to occur for the highest-travelling eigenray, *i.e.* the thermospheric eigenray, as will be seen subsequently. This amplification is large, as can be seen by observing the weak signal in Figure 4. In fact for source amplitudes lower than 0.5 Pa, the recorded signals have peaks that reach as much as 35% of the initial source strength.

The afore-mentioned variations are only notable for source amplitudes higher than approximately $\mathcal{A} = 5 \times 10^{-5} \bar{p}_{source}$. Figure 5 illustrates this point, by representing p^* for lower source amplitudes $\mathcal{A} < 5 \times 10^{-5} \bar{p}_{source}$. Here p^* is defined as the pressure fluctuation normalized by the amplitude of the first eigenray arrival. The signals collapse cleanly. The recorded signal is subsequently divided into three parts to examine in greater detail the effect of the source amplitude.

The first part of the signal corresponds to the arrival of the main eigenray (1) that results from refraction in the stratosphere. A linear ray-tracing algorithm applied to the chosen atmospheric model yields an arrival time of $t = 1320$ s for this eigenray. Numerically, these signals are observed to arrive between $t = 1250$ and $t = 1350$ s, as shown in Figure 6. The different signals have a similar general shape, but for sources stronger than $\mathcal{A} = 10^{-3}\bar{p}_{source}$ the positive front is observed to advance in time by up to 12 s. There is also a notable artefact at $t = 1250$ s for the strongest source, whose cause will be examined later.

The last part of the signal is due to the arrival of the low thermospheric eigenray, referenced as (3) in Figure 3. Signals associated to this event are observed to arrive between $t = 1350$ s and $t = 1600$ s. Two different views of this period, shown in Figures 7 (a) and (b), are used to examine this arrival. Representation (a) shows pressure signals scaled by the first eigenray amplitude. The amplification factor mentioned previously is clearly visible for the weakest sources, and it can be seen to diminish progressively with increasing source strength. For the lowest amplitudes, the signal shows no sign of having been shocked, and it has kept its initial duration of 20 s. For higher amplitudes, the amplification progressively disappears, rendering the representation (a) less useful for examining the time signatures. Instead, the signals are scaled by the amplitude of the eigenray (3) arrival, and plotted in Figure 7 (b). Here the effects of the source strength on the waveforms is clearly visible. Significant steepening occurs for sources stronger than $\mathcal{A} = 1 \times 10^{-5}\bar{p}_{source}$. The arrival also undergoes significant widening, with the highest amplitude arrival spanning around 110 s to be compared to the initial 20 s for the emitted signal and for the weaker arrivals. For the strongest sources, clear shock fronts delimit the overall arrival, and form a “U-wave” shape characteristic of N-waves having traversed a caustic. It should be noted that of the three eigenrays, the low thermospheric (3) eigenray is subjected to the strongest caustic, at an altitude of 120 km. This explains why the (3) eigenray arrivals are significantly more distorted than the others.

Finally, we consider the signal due to the arrival of the high thermospheric eigenray, referenced as (2) in Figure 3. Its behaviour is the most variable of the three eigenrays. For low source amplitudes, a very clean wavepacket is observed to arrive around $t = 1480$ s, as shown by Figure 8. Interestingly, this wavepacket is twice as long as the source signal, lasting 40 s instead of the source’s 20, and its shape is very well matched by $\sin(\omega_s t) \left[1 - \cos\left(\frac{\omega_s t}{4}\right)\right]$. As the source amplitude is increased, the wavepacket is seen to distort progressively towards N-waves, and the amplitude relative to the stratospheric arrival is observed to decrease. This arrival is also stronger than that of the stratospheric eigenray (1). The most spectacular effect of increasing source amplitude is seen on the arrival time. For $\mathcal{A} > 5 \times 10^{-5}\bar{p}_{source}$, the arrival time progressively diminishes from 1480 to 1260 s, which corresponds to the initial oscillation visible on the strongest source signal in Figure 6.

Figure 9 represents $p'/\sqrt{\bar{p}}$ over the whole computational domain for source amplitudes between $A = 5 \times 10^{-5}\bar{p}_{source}$ to $A = 5 \times 10^{-2}\bar{p}_{source}$, at $t = 1320$ s, which corresponds to




Figure 4: Pressure fluctuations for source amplitudes $\mathcal{A} = 5 \times 10^{-8} \bar{p}_{source}$ (left axis) and $\mathcal{A} = 5 \times 10^{-2} \bar{p}_{source}$ (right axis).




Figure 5: Pressure fluctuations for source amplitudes $\mathcal{A} < 5 \times 10^{-5} \bar{p}_{source}$ normalized by amplitude of initial eigenray (1) arrival. Vertical line represents arrival time of eigenray (1) obtained by linear ray tracing.

the arrival time of the initial eigenray (1). It gives a visual illustration of the movement of some parts of the pressure field compared to others, as well as an indication of their relative strength.

To conclude, time signatures recorded a long distance away from an acoustic source appear to vary in a highly non-linear fashion as a function of source amplitude. As expected, the highest travelling waves are subject to the strongest non-linear effects. These effects include both relative variations of amplitude between early and late arrivals, and also substantial changes in arrival times. In the equations used for this work, no account is taken of atmospheric absorption due to molecular relaxation. Relaxation effects are expected to be strongest in the high atmosphere, *i.e.* where non-linear propagation effects are also maximal. Hence quantitative aspects described in this work are likely to be different once relaxational effects are included.




Figure 6: Normalized pressure fluctuations due to the stratospheric eigenray (1). Colour scheme ranges from black for weakest source to orange for strongest source.




Figure 7: Pressure fluctuations due to low thermospheric eigenray (3), normalized (a) by stratospheric eigenray (1) amplitude and (b) by low thermospheric eigenray (3) amplitude. Colour scheme ranges from black for weakest source to orange for strongest source.




Figure 8: Pressure fluctuations due to high thermospheric eigenray (2), normalized by stratospheric eigenray (1) amplitude. Colour scheme ranges from black for weakest source to orange for strongest source.




Figure 9: Normalized pressure fluctuations represented at $t = 1320$ s, for source strengths ranging from $A = 5 \times 10^{-5} \bar{p}_{source}$ (top left) to $A = 5 \times 10^{-2} \bar{p}_{source}$ (bottom right).

6 Acceleration with graphics hardware

In the course of this work, the Navier-Stokes solver was implemented in a stream-processing language in order to evaluate the potential both of the language and of graphics hardware for this kind of algorithm. This has become possible since the advent of double-precision floating point operations in graphics hardware in 2008.

Graphics cards (GPUs) has been capable of performing useful computations for the past 6 years, since the introduction of programmable graphics pipelines. This reusing of GPUs for purposes other than graphics is referred to as *GPGPU*, or General Purpose computations on GPUs. The reader is referred to [11] for an overview of past and current trends in GPGPU.

Over the last few years a number of open portable languages have been developed for GPGPU programming. Unfortunately, none of them currently allows double precision programming in a portable fashion, and a vendor-specific language must be used. This limitation will probably disappear in the near future, possibly with the advent of OpenCL, a computational counterpart to OpenGL which is currently in the process of being certified.

In this study ATI hardware was used, and accordingly programming was performed using the BROOK+ language which exposes ATI double-precision floating-point operations to the programmer. As such, this work shows that graphics cards are highly amenable to high-precision Navier-Stokes solvers, but makes no effort to compare hardware from different vendors.

The BROOK+ programming language used in this is derived from the BROOK language originally developed at Stanford university[12]. It is designed as an extension to the C language, with specific constructs for describing data-parallel operations. The general computational model is referred to as *streaming*, and is a form of *SIMD* or Single Instruction Multiple Data programming. It is ideally suited to describing array problems that require identical treatment everywhere, which is the case for Euler equations solvers away from boundary conditions. However, near boundaries different conditions or equations are required and thus these zones do not equate well with the streaming model. An illustration of this is given by the number of lines of code in our BROOK+ based solver: only 500 lines are required for the central grid zone, while more than 3000 lines are needed to describe the boundary conditions. The BROOK+ language is currently available in a beta or testing version, and as such has a number of more or less frustrating quirks and bugs. These bugs will presumably disappear in time as the compiler and development environment matures.

For comparison, the Navier-Stokes solver has been implemented in an algorithmically equivalent fashion both in BROOK+ and in Fortran90. The results obtained on the GPU are not identical to those obtained with Fortran, due to slightly different conventions in rounding modes or other hardware aspects such as fused floating point operations. Fortunately, no systematic bias is induced by these differences. This last point was established by comparing both Fortran and GPU results to analytical solutions of propagating

acoustic pulses.

For the grid size used in this work, the optimized Fortran code ran at roughly 1.4 GFLOPs on a single core of a recent processor. The BROOK+ code running on an ATI HD3870 GPU ran around 18 times faster, yielding an effective computational rate of around 25 GFLOPs. This acceleration is both highly appreciable, since it was obtained at relatively little cost, and also a little frustrating, since it seems to be substantially lower than the maximum rated hardware speed of around 100 GFLOPs in double precision for our GPU. A preliminary performance appraisal suggests that memory access is the main bottleneck, but this point would need a more quantitative examination.

Our overall impression of the GPGPU experience was positive. It appears to be a very promising technique for the future, particularly since GPUs are increasing in power more rapidly than are CPUs. The hardware is capable of performing computations in double precision, and its programmability is such that there is no particular limit in the type of algorithms that can be programmed. However, both programming languages and graphics hardware are changing very fast, and the field is far from being mature. This point should be kept in mind before embarking in large-scale software development on graphics cards.

7 Conclusions

A computational solver based on the resolution of the full Navier-Stokes equations with high-order space and time methods is presented. A non-linear filtering technique suitable for long-range acoustic propagation is tested, and shown to yield significantly improved waveforms for a simple example of non-linear propagation.

The solver is used to examine the effect of source strength on long-distance propagation. For low source amplitudes, time signatures collapse cleanly, and the longest eigenray provides the strongest arrival. For higher source amplitudes, the relative level of arrivals due to the different eigenrays is strongly modified, to the point where the strongest arrival is due to the shortest eigenray for the most powerful sources. Arrival times of high-atmosphere eigenrays are also strongly modified.

Finally, the algorithm used in this work has been implemented in a programming language destined to run on graphics hardware. An acceleration factor of 18 was obtained compared to optimized Fortran code running on a recent processor. Such graphics hardware shows great potential for future large-scale computations requiring vast computational resources.

Acknowledgments

This work was undertaken as part of a collaboration with the DASE department of the CEA. The authors would like to thank Olaf Gainville for his help in providing tabulated data on the Misty Picture experiment.

References

- [1] O. Gainville, P.-F. Piserchia, P. Blanc-Benon, and J Scott. Ray tracing for long range atmospheric propagation of infrasound. *12th AIAA/CEAS Aeroacoustics Conference AIAAP*, (2006-2451), 2006.
- [2] O. Gainville. *Modélisation de la propagation atmosphérique des ondes infrasonores par une méthode de tracé de rayons non linéaire*. PhD thesis, EC-Lyon, 2008.
- [3] L. C. Sutherland and H. E. Bass. Atmospheric absorption in the atmosphere up to 160 km. *JASA*, 120(5):2895, 2004. See also Erratum. *J. Acoust. Soc. Am.*, **120**(5), 2985 (2006).
- [4] C. Bogey and C. Bailly. A family of low dispersive and low dissipative explicit schemes for noise computations. *Journal of Computational Physics*, 194(1):194–214, 2004.
- [5] J. Berland, C. Bogey, O. Marsden, and C. Bailly. High-order, low dispersive and low dissipative explicit schemes for multiple-scale and boundary problems. *Journal of Computational Physics*, 224:637–662, 2007.
- [6] O. Marsden, C. Bogey, and C. Bailly. High-order curvilinear simulations of flows around non-cartesian bodies. *Journal of Computational Acoustics*, 13(4), 2005.
- [7] O. Marsden, C. Bogey, and C. Bailly. Direct noise computation of the turbulent flow around a zero-incidence airfoil. *AIAA Journal*, (4):874–883, 2008.
- [8] D. Gaitonde and M. Visbal. Padé-type higher-order boundary filters for the Navier-Stokes equations. *AIAA Journal*, 38(11):2103–2112, november 2000.
- [9] C.K.W. Tam and Z. Dong. Radiation and outflow boundary conditions for direct computation of acoustic and flow disturbances in a nonuniform mean flow. *Journal of Computational Acoustics*, 4(2):175–201, 1996.
- [10] C. Bogey, N. De Cacqueray, and C. Bailly. A shock-capturing methodology based on adaptive spatial filtering for high-order non-linear computations. *Journal of Computational Physics*, 2008. doi:10.1016/j.jcp.2008.10.042.

- [11] J. Owens, D. Luebke, N. Govindaraju, J. Krüger, A. Lefohn, and T. Purcell. A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum*, 26(1):80–113, March 2007.
- [12] Ian Buck, Tim Foley, Daniel Horn, Jeremy Sugerman, Kayvon Fatahalian, Mike Houston, and Pat Hanrahan. Brook for gpus: stream computing on graphics hardware. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*, pages 777–786, New York, NY, USA, 2004. ACM.