



HAL
open science

A combinatorial model reduction method for the finite element analysis of wind instruments

Marie Jeanneteau, Paul Oumaziz, Jean-charles Passieux, Vincent Gibiat,
Jonathan Cottier

► To cite this version:

Marie Jeanneteau, Paul Oumaziz, Jean-charles Passieux, Vincent Gibiat, Jonathan Cottier. A combinatorial model reduction method for the finite element analysis of wind instruments. *International Journal for Numerical Methods in Engineering*, 2024, 10.1002/nme.7582 . hal-04677681v1

HAL Id: hal-04677681

<https://insa-toulouse.hal.science/hal-04677681v1>

Submitted on 28 Aug 2024 (v1), last revised 2 Sep 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A combinatorial model reduction method for the finite element analysis of wind instruments

M. Jeanneteau^{1,2} and P. Oumaziz¹ and J.C. Passieux¹ and V. Gibiat¹ and J. Cottier²

¹ Institut Clement Ader (ICA), Université de Toulouse, INSA-ISAE-UT3-Mines Albi, Toulouse fr
² Henri Selmer Paris

Submitted to IJNME 2024, DOI:10.1002/nme.7582

Summary A high-fidelity finite element model is proposed for the complete simulation of the time-harmonic acoustic propagation in wind instruments. The challenge is to meet the extremely high accuracy required by professional musicians, in a complex domain, for all fingerings and over a wide frequency range, within an affordable computational time. Several modelling assumptions are made to limit the numerical complexity of the problem while preserving all relevant physics. A dedicated high-performance solution strategy is also proposed, based on partitioning, condensation and model order reduction, exploiting the combinatorial nature of wind instrument fingerings. Finally, the proposed approach is applied to the simulation of an alto saxophone. An order of magnitude reduction in memory and computational cost is achieved.

keywords Wind instrument, Time-harmonic acoustic, Finite elements, Model order reduction, Static condensation

1 Introduction and motivation

Wind instruments are complex objects. These musical instruments are classically modelled as a coupled exciter-resonator system [?, ?], consisting of a mechanical oscillator - the reed or lips - and a resonant cavity - the fluid space bounded by the solid rigid structure, also known as the air column. Tone hole wind instruments are equipped with a network of chimneys or holes which the musician opens or closes with their fingers or with keys via a complex and extensive keywork. The latter can influence the vibro-acoustic behaviour of the instrument. However, these second-order exciter/structure and fluid/structure couplings are usually not taken into account. In most cases, the instrument is described only in terms of the dimensions of the air column, i.e. the bore and the holes, possibly together with their associated keys. However, an instrument characterised by a unique set of dimensional parameters presents a multitude of configurations corresponding to the many possible combinations of opening and closing of the tone holes. These configurations are called fingerings because they are modified by the position of the fingers on the tone holes or on the mechanical keyboard. The notes that can be played on the instrument depend directly on these combinations. Some notes can have more than one fingering, allowing the musician to choose the one best suited to the musical phrase. For an alto saxophone with 26 tone holes, the complex keyboard allows 50 fingerings to cover the 33 notes of the tempered scale between C#3 and A5 in English notation.

Professional musicians have high expectations of their wind instruments, both in terms of the inseparable acoustic and mechanical aspects. Their expert hearing is able to distinguish about 5 cents around the expected

note, which corresponds to a precision between 0.4 Hz and 3 Hz along the usual playing range of the same alto saxophone. This tuning accuracy of a few cents over the entire ambitus of the instrument is a major concern for the manufacturers. For instance, at Henri Selmer Paris, the development of a new model, which is still essentially experimental and empirical, is therefore the result of a costly process lasting several years.

To reduce costs and development time, industrial manufacturers want to rely on modelling or even numerical simulation of the wind instrument air column. Some acoustic properties of the instrument (tuning, homogeneity, harmonicity, ease of play or timbre, etc.) can be quantified, at least partially, from the input impedance of the resonator [?, ?]. This complex quantity that describes the acoustic response of a system at its inlet, has exact analytical Kirchhoff's solutions [?] only for ideal geometries such as cylinders or untruncated cones, which do not concern instrument factors. To access the input impedance of instruments with such complex bores, manufacturers mostly rely upon experimental [?][?], semi-experimental [?] or semi-analytical[?] approaches because their accuracy to cost is still competitive as compared to high fidelity numerical simulations. Efforts must therefore be made on both (a) physics modelling and (b) high-performance solution strategies to help finite element analysis make a breakthrough in such musical instrument design.

The accuracy of the simulations depends strongly on the modelling of the dissipative effects at the walls[?]. Within the so-called 'low reduced frequency solutions' [?] considered in musical acoustics applications, two types of approximate models are classically proposed in the literature: on the one hand, those in the line of the Zwicker-Kosten (ZK) model [?] defined in 1949 for cylinders, where dissipative effects are taken into account in the complex propagation constant; on the other hand, the ones based on viscothermal boundary layer impedances (BLI), in the line of Cremer's model [?] proposed in 1948 for an infinite plane wall, where losses are modeled by an equivalent wall admittance. The range of applicability of the two models mentioned above have been extended: approximations for non-cylindrical tubes or axisymmetric geometries have been proposed for the ZK models [?, ?], and equivalent viscothermal BLI models have been generalized to the 3D case [?, ?, ?]. These models are used both for resolution using transmission-line modelling by Transfer-Matrix Method (TMM) [?, ?], or PDEs discretization modelling like the Finite Element Method (FEM) [?, ?]. The accuracy of the simulations also depends on the radiation modelling of the tone holes [?, ?, ?], which means that potentially large areas of the surrounding air must be simulated in addition to the bore of the instrument. These additional parts of the analysis domain may more than double the total number of degrees of freedom.

Using a finite element method to solve the time-harmonic acoustics with such elaborate models on complex geometries like saxophones requires solving a series of several millions degrees of freedom (DOFs) problem with complex-possibly nonlinear-impedance conditions at the boundaries. In addition, it was shown that, in order to avoid pollution effect[?], high order finite element approximations are required [?] [?] which further increase the computational burden. To handle the arising large linear systems, iterative solvers are mainly used. However, the conditioning of such methods is highly sensitive to the wave number [?], with a deterioration of convergence for medium and high frequency ranges. In [?], the authors proposed a two-level deflation preconditioner with the complex-shifted Laplacian preconditioner associated with a GMRES solver[?] (Generalized Minimal RESidual) to shift the problematic eigenvalues far from 0. This results in a scalable method whose convergence is quite independent of the wave number. Moreover, among the large class of iterative methods, domain decomposition methods are interesting to reduce the memory consumption by subdividing the problem into subdomains. Several approaches have been developed over the last two decades: FETI-H or FETI-DPH as a FETI (Finite Element Tearing and Interconnecting) method for Helmholtz problems [?, ?], or Schwarz methods and in particular the optimised Schwarz method, which seems to be the more efficient and the most studied in the community [?, ?]. However, the choice of transmission parameters in the case of the optimised Schwarz method is not so trivial and is not yet implemented in open source scientific libraries.

All these methods help to solve a large linear system efficiently, but unfortunately the size of the linear sys-

tems to be solved is not the only problem. In fact, in the context of a large frequency sweep and because of the nonlinear nature of the chosen admittance modelling, tens of thousands of such large linear systems need to be solved for a single fingering. All fingerings considered, this represents a prohibitive amount of computing time. To overcome this difficulty, [?, ?] proposed a review of interpolatory model order reduction to reduce the number of frequency steps to be computed. The reduction methods in the literature usually rely on simplifying assumptions, such as linearity or space-frequency variable separation of operators, which may not all be compatible with the accuracy required for musical instruments. Making finite element simulations compatible with time-to-market is therefore a real challenge.

In this paper, a tractable 3D simulation model dedicated to the scalar Helmholtz problem is proposed, based on the FEM, using the open source DOLFINx [?] computational environment, suitable for industrial design, by saving costs on the three following aspects : (i) appropriate reduced models of viscothermal losses and radiations are used to model precisely the behaviour of the resonator. Then (ii) a condensation method and a model reduction strategy are proposed to reduce the analysis domain to the bore of the instrument only, with (non-standard) computationally affordable impedance conditions at the tone holes interfaces. The resulting moderate size linear systems are compatible with efficient sparse direct solvers. Last (iii) a high performance parallel and distributed solution strategy using interpolation and low-rank approximations is developed to rationalise the numerical cost in this multi-resolution context, i.e. with tens of thousands of linear systems with similar structure.

The purpose of our approach is to complement a 1D method like TMM with a high fidelity 3D simulation that provides additional data. First, we obtain the complete pressure field inside the instrument and it is possible to embed fine geometric details and analyze their impact on the impedance of the instrument. Second, to optimize the design, 3D analysis could be used to account for second-order effects that 1D analysis does not allow. These could include chimney orientation, geometric defects (such as orange peel inside the body), or optimizing the position and orientation of keys above the holes. It is hoped that their effect on the playability of the instrument will be accessible from the 3D simulation.

The paper is organized in the following way. The theory of the finite element methods for time-harmonic acoustics for exterior problem in unbounded domains is reviewed in the Section 2. Section 3 presents in detail the modelling strategies adopted and the corresponding High Performance Computing (HPC) and reduced numerical solution strategy. The next section presents the simulation results of a complete saxophone body study in comparison with experimental measurements; the experimental setup and specimen are briefly described first. The last section summarizes the conclusions and perspectives.

2 Problem statement

2.1 Standard formulation

Let $V \in \mathbb{R}^3$ be a solid object, origin-centered. The exterior region $R = \mathbb{R}^3 \setminus V$, with boundary $S = S_Y \cup S_p$, is defined as boundary conditions of type admittance Y and pressure source p_0 are respectively applied on S_Y and S_p . The acoustic pressure $p(x)$ for all point x of the unbounded region R is solution of the strong form of the time-harmonic problem, governed by the scalar second order differential homogeneous Helmholtz equation (x -dependency omitted for clarity and sign convention for the phase $e^{+j\omega t}$, where $j = \sqrt{-1}$):

$$\nabla^2 p + k^2 p = 0 \quad \forall x \text{ in } R, \quad (1a)$$

$$\nabla p \cdot n = -j\rho \omega Y p \quad \text{on } S_Y, \quad (1b)$$

$$p = p_0 \quad \text{on } S_p, \quad (1c)$$

$$\lim_{r \rightarrow \infty} r \left(\frac{\partial p}{\partial r} + jk p \right) = 0, \quad (1d)$$

where ρ is the density, ω the angular frequency, k the wave number and $r = \|x\|$.

A continuous approximation $\hat{p}(x)$ of $p(x)$ can be computed using the standard h -version of the finite element method (FEM). To do so, FE mesh is built on the finite computational domain $\Omega \subset R$, bounded by a truncature surface Γ ; see figure 1. Let us denote $\hat{\mathbf{p}}$ the degree-of-freedom vector gathering the values of the pressure field $\hat{p}(x)$ at the n_{DOF} nodes of this mesh.

Using the Bayliss-Gunzberger-Turkel [?] second order formulation to approximate the Sommerfeld boundary condition (1d) [?] at the spherical artificial boundary Γ , $\hat{\mathbf{p}}$ is solution of the following sparse, complex, symmetric algebraic system :

$$\mathbf{A} \hat{\mathbf{p}} = (\mathbf{K} - \omega^2 \mathbf{M} + j\omega \mathbf{C} + \mathbf{R}_\omega) \hat{\mathbf{p}} = \mathbf{0}, \quad (2)$$

where :

$$(\mathbf{K})_{ij} = \int_{\Omega} \nabla N_i \cdot \nabla N_j \, d\Omega, \quad (3a)$$

$$(\mathbf{M})_{ij} = \int_{\Omega} \frac{1}{c^2} N_i N_j \, d\Omega, \quad (3b)$$

$$(\mathbf{C})_{ij} = \int_{S_Y} \rho Y N_i N_j \, dS, \quad (3c)$$

$$(\mathbf{R}_\omega)_{ij} = \int_{\Gamma} B_1 N_i N_j \, d\Gamma + \int_{\Gamma} \frac{1}{2B_1} \nabla_{\Gamma} N_i \cdot \nabla_{\Gamma} N_j \, d\Gamma, \quad (3d)$$

where $N_i(x)$ are the C^0 continuous interpolation functions with local compact support (FE shape functions), $B_1 = jk + \frac{1}{r}$ and the angular gradient operator defined in spherical coordinate by $\nabla_{\Gamma} = \frac{1}{r} \frac{\partial}{\partial \theta} \mathbf{e}_\theta + \frac{1}{r \sin \theta} \frac{\partial}{\partial \phi} \mathbf{e}_\phi$.

The linear system (2) is then solved considering a planar acoustic source on the entry surface S_p , taken as a Dirichlet boundary condition on pressure ($p(x) = 1 \text{ Pa} \quad \forall x \in S_p$).

The impedance at the inlet section Z_{in} , ratio of mean acoustic pressure p_{in} and cross-sectional mean acoustic velocity at the input plane v_{in}^{\perp} , is finally computed by:

$$Z_{\text{in}} = \frac{p_{\text{in}}}{v_{\text{in}}^{\perp}} = \frac{\int_{S_p} \hat{p} \, d\Gamma}{\int_{S_p} \hat{v} \cdot -n \, d\Gamma} = j\rho\omega \frac{\int_{S_p} \hat{p} \, d\Gamma}{\int_{S_p} \nabla \hat{p} \cdot n \, d\Gamma}. \quad (4)$$

In the following, the ratio $Z = \frac{Z_{\text{in}}}{Z_0}$, referred to as input impedance and where $Z_0 = \rho c$ is the specific acoustic impedance of the fluid (c is the speed of sound), is preferred to facilitate comparison with experimental measurements.




Figure 1: Bounded computational domain for wind instrument acoustic problem and associated notations.

2.2 Cremer's model for thermoviscous losses

Thermal and viscous losses are complex multiphysics dissipative phenomena that are confined close to the boundary of the instrument. There are many ways to model these losses in a Helmholtz problem, the interested reader is referred to the comprehensive review of Chabassier and Thibault [?]. Among them, the so-called boundary layer impedance (BLI) is a class of methods which models dissipations as special Robin boundary conditions [?] [?] [?] [?] [?]. Without loss of generality, the equivalent wall admittance used in this paper to model the viscothermal losses at the internal surface of the instrument is the pressure-dependant admittance $Y(\omega, \nu)$ function defined in [?]:

$$Y(\omega, \nu) = \frac{1}{\rho c} \sqrt{\frac{j\omega}{c}} \left[\left(1 - \left(\frac{\nu \cdot \mathbf{n}}{\|\nu\|} \right)^2 \right) \sqrt{l_v} + (\gamma - 1) \sqrt{l_t} \right], \quad (5)$$

where $l_v = \frac{\mu}{\rho c}$ and $l_t = \frac{l_v}{\text{Pr}}$ are the viscous and thermal length, with μ the dynamic viscosity, γ the ratio of specific heats of air, Pr the Prandtl number and $\nu = -\frac{\nabla p}{j\rho\omega}$ the acoustic velocity calculated from solution at angular frequency ω . The initial formulation [?] introduces $\theta = \arccos\left(\frac{\nu \cdot \mathbf{n}}{\|\nu\|}\right)$, so called angle of incidence of the acoustic velocity to the wall, but the expression (5) is here preferred as it makes the following paragraphs easier to understand. The angle $\theta = \frac{\pi}{2}$ is known only for the cylinder (plane waves), but for any other geometry iterative resolution procedures or hypotheses to avoid it are necessary. The nonlinear problem at angular frequency ω is solved using an alternating direction fixed point algorithm. More precisely, starting from a first simulation using rigid boundary conditions ($Y = \mathbf{0} \Rightarrow \mathbf{C} = \mathbf{0}$ in (2)), one iteration is composed of the two following steps: (i) computing the wall admittance Y from the velocity field \mathbf{v} considered fixed; and then (ii) solving Problem (2) to compute the velocity field using the previous admittance Y which is considered known. These two steps are repeated until convergence. In practice, one iteration of this fixed point algorithm (involving the resolution of only two linear systems) is usually considered sufficiently accurate with minimal impact on the solution.

2.3 External volume modelling

The size of the computational domain (surrounding sphere of radius r) and the mesh density selected to minimize pollution introduced by the low-order approximate local boundary conditions on Γ , leads to large sparse matrices.

Perfectly Matched Layers (PML) [?], infinite elements [?] or Higher multipole expansion orders [?] could also be used in combination with more suited ellipsoidal or cylindrical surrounding domains to reduce a bit the problem size, see [?] for a comprehensive review. In any case, the number of external DOFs (outside the instrument) exceeds the number of internal DOFs (inside), whereas the complexity of the bore designed by the maker is mostly described by the latter. Because of this last remark, the numerical strategy proposed herein focuses particularly on modelling this large external domain at a fair cost.

2.4 Problem complexity

Solving the Helmholtz problem for the $N_f \approx 50$ fingerings of the wind instrument and for $N \approx 7000$ frequencies of the whole frequency range of interest (from 20 Hz to 8000 Hz every 1 cent, with a minimal frequency step of 0.2 Hz), while satisfying the high precision on the resonance frequencies required by the factor, is thus costly, both in memory (large problem - cf Section 2.3) and in time (large number of resolutions - cf Section 2.2). Simulating the acoustic response of an instrument requires to solve almost half a million linear systems, each with several million degrees of freedom. In other words, it may take weeks to months with classic finite element analysis software to get a simulation result which explains why, as mentioned in the introduction, the experimental method is still very much used.

3 Numerical methods: reduced modelling and computational strategies

The goal of the proposed computational strategy is to achieve an order of magnitude reduction in computational time so that 3D finite element analysis becomes a competitive alternative to the manufacture of a prototype (and possibly the associated manufacturing tooling). The condensed-partitioned-reduced simulation framework proposed here reduces computational costs by combining different numerical strategies, at the physics, model and computation levels.

3.1 Incremental Cremer boundary condition

As explained in Section 2.2, given the nonlinearity of thermoviscous losses, solving the damped problem for the aforementioned N frequencies, means solving $2N \approx 14000$ problems. Given the small frequency step and the high regularity of the solution, it can reasonably be assumed that $\left(\frac{v_{\omega \cdot n}}{\|v_{\omega}\|}\right)^2 \approx \left(\frac{v_{\omega+\delta\omega \cdot n}}{\|v_{\omega+\delta\omega}\|}\right)^2$ for small $\delta\omega$. Thus, it is proposed to implement an incremental formulation \mathbf{C}_{incr} of acoustic damping operator based on the previous Cremer formulation:

$$\begin{cases} \mathbf{C}_{\text{incr}}(\omega_1)_{ij} = \int_S \rho Y(\omega_1, v_{\omega_1}^r) N_i N_j dS, \\ \mathbf{C}_{\text{incr}}(\omega_{k, k \geq 2})_{ij} = \int_S \rho Y(\omega_k, v_{\omega_{k-1}}^d) N_i N_j dS. \end{cases} \quad (6)$$

Superscripts r or d indicate whether the acoustic velocity is calculated from the rigid or damped problem. Initializing the fixed point at ω_k with the Cremer solution at angular frequency ω_{k-1} is, according to our numerical experiments, neither better nor worse than initialization with rigid condition at the right angular frequency ω_k . Solution at ω_k provides a free and relevant initialization to the Cremer iterations of section 2.2. Only one linear system has to be solved whereas the rigid condition required the resolution of an additional FE problem for each frequency. The fact that problem at frequency ω_k requires the knowledge of the solution at angular frequency ω_{k-1} introduce a causality dependence. This dependence between problems in frequency may break parallelism. It is possible to restore the parallelism by splitting the frequency range into subsets of contiguous frequencies. The first frequency of each subset is initialized using a rigid wall. Therefore, the number of linear system to solve is almost halved with this choice.

3.2 Reduced modelling adapted to the combinatorial nature of tone hole wind instrument

The goal of this section is to reduce as much as possible the analysis domain, at both spatial and frequency levels. The domain reduction proposed here is based on static condensation. Assumptions and principles of the condensed model are briefly presented in this paragraph, and all formalisation aspects are detailed in 3.2.1.

A first assumption is to consider a rigid null admittance boundary condition at any external wall (including keys). With such null-admittance condition, the exterior problem is linear and pressure independent.

Thus, the computational domain can be reduced using static condensation of the so-called external volume, at bore interfaces, so that the analysis domain is limited to the single body of the instrument. Another assumption is then to neglect the external acoustic interaction between holes, in order to replace the use of one large sphere embedding the whole instrument by several smaller independent spherical domains at the top of each hole (referred to as ‘blind’ spheres in the following). This second assumption participates in the reduction of the number of external DOFs first, but above all it makes possible to apply the static condensation of the external domains, at each hole interface, which leads to additional gains in terms of memory requirements and resolution costs. This choice is motivated by the combinatorial nature of tonehole wind instruments, for which each note of the register is defined by one or more combination of open or closed holes.

From a computational point of view, this hole-wise condensation approach, based on multiple fingering-independent spheres, has several advantages compared to a fingering-wise condensation approach based on a single fingering-dependent sphere. First, the use of ‘blind’ spheres allows to work with smaller external domains than the larger fingering-dependent sphere, and thus more easily condensed. Second, the hole-wise approach involves fewer condensations than the fingering-wise approach since the considered instruments typically have fewer holes $H \approx 25 < N_f \approx 50$ than the number of fingerings. Then, this assumption of zero-external interaction allows to define the global condensed fingering-dependent operator at the interfaces as the sum of locally small dense independent operators. Thus, instead of one large dense dependent operator, it provides a better sparsity of the final condensed system to solve.

In terms of use, this hole-wise condensation approach also allows more flexibility for design parameter studies (changing tonehole/key dimensions requires new computations only for the concerning ‘blind’ spheres), and leads to gains in data processing and storage.

The previous non-standard impedance conditions obtained by static condensation of independent ‘blind’ sphere can be computed in an offline phase. Given the frequency range of interest, this represents a expensive costs at computation and storage levels. The model reduction proposed here tackles both of these aspects, using frequency interpolation and low-rank approximation respectively.

3.2.1 Domain reduction: static condensation on ‘blind’ spheres

A naive way to model such a problem would be to consider a computational domain Ω^f as the union of two subsets $\Omega^f = B \cup S^f$, where B represents the volume of the bore (main domain of interest) and S^f a unique large surrounding sphere for each fingering f . The bore B would have a constant mesh M_B as shown in figure 4. External spherical domains S^f would generate N_f meshes M_S^f since they depend on the fingering. Such large external spheres would results in $N_f = 50$ different meshes, each with hundreds of millions of degrees of freedom.

Alternatively, it is proposed here to define the external domain as the union of a set of smaller ‘blind’ spheres S_k centered on each hole, see figure 3. Each ‘blind’ sphere shares only one interface with B at hole k , see Figure 4. Then the new mesh \tilde{M}^f of this new computational domain $\tilde{\Omega}^f$ is simply composed of the union of the bore mesh and the spheres associated only with the open holes: $\tilde{\Omega}^f = B \cup \bigcup_{k \in \mathcal{F}} S_k$. With such a choice, the construction of $N_f = 50$ very large spherical meshes with multiple interfaces to the bore was replaced by $H = 19 + 1$ small blind spheres sharing only one surface interface with the bore M_B , each sphere being activated only when the corresponding hole is open, see figure 2 for illustration. Even if the volumes of the spheres S_k overlap, their

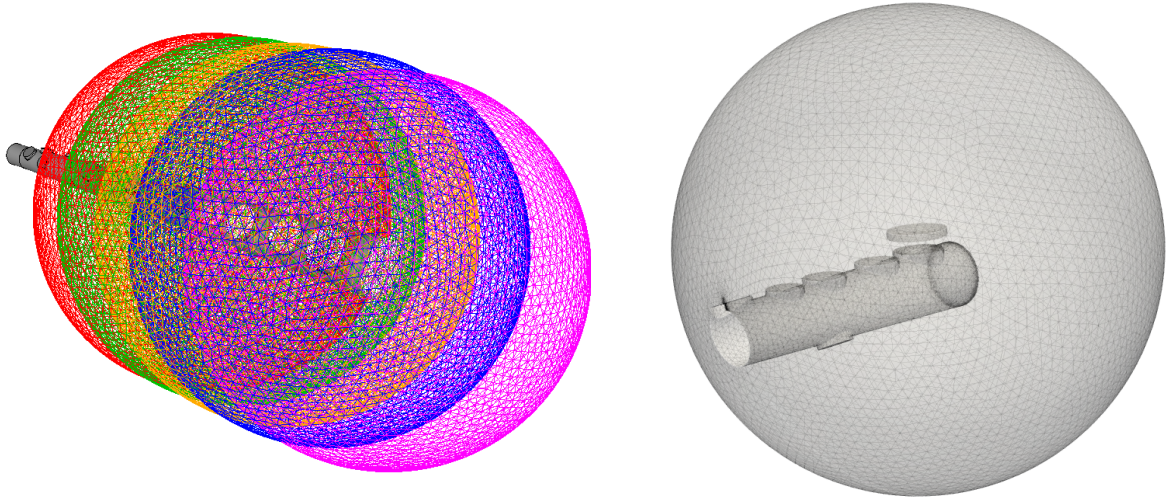


Figure 3: Mesh of G4 fingering on a alto saxophone body: (left) analysis domain made of the bore B with 5 ‘blind’ sphere meshes M_{S_k} of the exterior domain; (right) closeup on one ‘blind’ sphere mesh with one key on top of the hole, including a part of the complex external surface of the instrument.

After substitution in (7), an equivalent, but smaller, system follows:

$$\begin{bmatrix} \mathbf{A}_{B,B} & \mathbf{A}_{B,i_{f_1}} & \cdots & \mathbf{A}_{B,i_{f_H}} \\ \mathbf{A}_{i_{f_1},B} & \Sigma_{i_1,i_1} & & \\ \vdots & & \ddots & \\ \mathbf{A}_{i_{f_H},B} & & & \Sigma_{i_H,i_H} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{p}}_B \\ \hat{\mathbf{p}}_{i_1} \\ \vdots \\ \hat{\mathbf{p}}_{i_H} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_B \\ \mathbf{0}_{i_1} \\ \vdots \\ \mathbf{0}_{i_H} \end{bmatrix}, \quad (9)$$

where $\Sigma_{i_k,i_k} = \mathbf{A}_{i_k,i_k} - \mathbf{A}_{i_k,S_k} \mathbf{A}_{S_k,S_k}^{-1} \mathbf{A}_{S_k,i_k} = \mathbf{A}_{i_k,i_k}^B + (\mathbf{A}_{i_k,i_k}^{S_k} - \mathbf{A}_{i_k,S_k} \mathbf{A}_{S_k,S_k}^{-1} \mathbf{A}_{S_k,i_k}) = \mathbf{A}_{i_k,i_k}^B + \mathbf{III}_{i_k,i_k}^{S_k}$, $k \in \mathcal{F}$, with \mathbf{A}_{i_k,i_k}^B the contribution of the elements of bore mesh M_B , and $\mathbf{III}_{i_k,i_k}^{S_k}$ the Schur complement of the sphere S_k , the contribution of elements of sphere mesh M_{S_k} .

The Dirichlet boundary conditions are taken into account by substitution. This means to substitute the involved degrees of freedom $\hat{\mathbf{p}}_B$ and to modify the corresponding right-hand side $\mathbf{0}_B$.

From now on, \mathbf{S} refers to the sparse block matrix which contains the dense Schur complements of the F spheres for the considered fingering:

$$\mathbf{S} = \begin{bmatrix} \cdots & & & \\ \vdots & \mathbf{III}_{i_1,i_1}^{S_1} & & \\ & & \ddots & \\ & & & \mathbf{III}_{i_F,i_F}^{S_F} \end{bmatrix}. \quad (10)$$

Operators $\mathbf{III}_{i_k,i_k}^{S_k}$ are dense but their size is limited (hundreds of DOFs). One advantage of performing a condensation independently on each hole is that operator \mathbf{S} remain mainly sparse, since it is a block diagonal




Figure 4: Mesh M_B of the bore: (left) the whole mesh with conformal interfaces shared with ‘blind’ sphere meshes M_{S_k} for fingering G4 and (right) a closeup on the mesh refinement near the chimneys.

matrix of dense local operators. As mentioned above, with one single large sphere for the exterior domain, \mathbf{S} would be dense $(N_i \times H)^2$ vs $N_i^2 \times H$ with the proposed approach.

Let us insist that, despite the explicit computation of the local Schur complements $\mathbf{III}_{i_k, i_k}^{S_k}$ of each sphere S_k , the associated computational overhead is limited since they are constant for any fingering. They are computed offline, saved and massively reused to reconstruct \mathbf{S} for the considered fingering. This method of static condensation allows a drastic reduction in the size of the system to be solved, which is at least divided by 3. It largely compensates for the small loss of sparsity of the matrix equation (9) mentioned above and for the calculation of these operators.

Up to now, these operators are independent on the fingering but must be computed for each frequency step. In other words, one operator $\mathbf{III}_{i_k, i_k}^{S_k}$ should be computed and saved for each of the H holes and each of the $N \approx 7000$ frequency steps. Of course, more than the high computational cost of such an offline phase, the volume of data to be stored would be prohibitive.

In the following section, a dimension reduction strategy is proposed to reduce the number of effective computations (and the memory footprint) of such condensed operators by about two orders of magnitude.

3.2.2 Model reduction: construction of Schur complement matrix

As just mentioned, the computation and storage of all $H \times N$ Schur complements of the instrument would be prohibitive. In this paragraph, the computational cost is drastically reduced by combining interpolation and low rank approximation: First, given their smooth evolution in frequency, it is proposed to compute them using spline interpolation from a smaller set of linearly spaced frequencies. Second, it is shown how to reduce the storage by almost another order of magnitude using low-rank approximation based on a truncated Singular Value Decomposition (SVD).

To better illustrate the proposed reduction strategy, a running example is analysed in this section.

Let us consider a simplified instrument, tenon-body assembly of an alto saxophone (Henri Selmer Paris - SERIES II) with the single drawn out chimney about 2mm height and 24mm diameter, approximately halfway along the cone (so $H = 2$), such as external domain is modeled by small 25cm diameter spheres, centered on chimney and tube end sections. The quadratic tetrahedral meshes of the sphere contains 120k and 133k DOFs

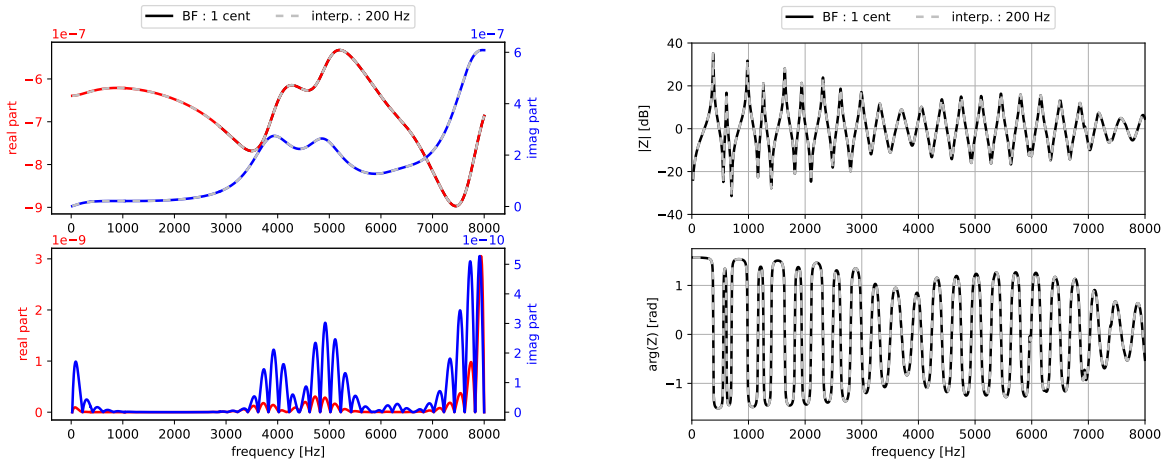
respectively, refined at holes such that there are $n_{\text{int}} = 720$ and 787 DOFs at interfaces. The bore mesh contains 250k DOFs. In this study, the radius of the sphere is deliberately small to limit computational time, without affecting the following conclusions.

Solving this problem with a brute force (BF) approach would consist in computing, for the H spheres S_k and for the N frequencies ω_l of interest, the $n_{\text{int}} \times n_{\text{int}}$ dense Schur complements matrices $\mathbf{III}^{S_k}(\omega_l)$. Let us now define \mathbf{U}^{S_k} the dense $m \times N$ matrix, the so-called condensation matrix of sphere S_k , such that each column contains the upper triangle matrix coefficients of Schur complements of sphere S_k computed at ω_l ($m = \frac{n_{\text{int}}(n_{\text{int}}+1)}{2}$) taking advantage of the symmetry of \mathbf{III}):

$$(\mathbf{U}^{S_k})_{(i-1)(n_{\text{int}}-\frac{i}{2})+j,k} = (\mathbf{III}^{S_k}(\omega_l))_{i,j}, \text{ with } \begin{cases} 1 \leq i \leq n_{\text{int}} \\ i \leq j \leq n_{\text{int}} \\ 1 \leq l \leq N \end{cases} . \quad (11)$$

The frequency behaviour of the matrix coefficients of the Schur complement matrices is analysed by studying the different rows of \mathbf{U}^{S_k} . Given the smooth evolution of coefficient values with respect to the frequency (see Figure 5a), a first gain on the computational level is obtained by restricting the computation of the Schur complements to a reduced number $n_S \ll N$ of regularly distributed sampling frequencies (Δf_S step) and then storing them in smaller $m \times n_S$ matrix $\hat{\mathbf{U}}^{S_k}$ ($1 \leq l \leq n_S$ in eq. (11)). From these sampling values, it is proposed to approximate the full set of N using piecewise quintic spline interpolation of $\hat{\mathbf{U}}^{S_k}$ matrix coefficients.

This approach has a minimal impact on the resulting input impedance estimation as illustrated in figure 5b. More precisely, let $\tilde{\mathbf{U}}^{S_k}$ be $m \times N$ matrix obtained from interpolation of $\hat{\mathbf{U}}^{S_k}$ on the whole frequency range



(a) Frequency behaviour of a randomly chosen matrix coefficient: upper graph shows real (red) and imaginary (blue) parts with respect to the frequency; lower graph shows absolute error between the approaches.

(b) Input impedance: upper and lower graphs show magnitude and phase respectively with respect to the frequency.

Figure 5: Comparison of Brute Force (BF) approach (solid line) with $N = 7000$ points, one every cent, and proposed interpolation (dashed line) with a $n_S = 41$ sampling points, namely, one every 200 Hz.

of interest, Z the input impedance computed using the BF approach, and \tilde{Z} computed based on $\tilde{\mathbf{U}}$ using the interpolation approach. Let finally $\epsilon_{\mathbf{X},\mathbf{Y}}$ be L^2 relative error matrix or vector norms such as:

$$\epsilon_{\mathbf{X},\mathbf{Y}} = \frac{\|\mathbf{X} - \mathbf{Y}\|_2}{\|\mathbf{X}\|_2} \quad (12)$$

The evolution of the interpolation error on the condensation matrix $\epsilon_{\mathbf{U},\tilde{\mathbf{U}}}$ and on the impedance $\epsilon_{\mathbf{Z},\tilde{\mathbf{Z}}}$ as a function of the number of sampling points n_s is given in Figure 7.

The relative error $\epsilon_{\mathbf{Z},\tilde{\mathbf{Z}}}$ is less than or about 10^{-6} for $n_s \geq 81$, so computing the Schur complements every $\Delta f_S = 100$ Hz seems sufficient for correct input impedance estimation.

The question now is how to store the condensation matrix $\hat{\mathbf{U}}^{S_k}$ at the right cost. The singular value decomposition (SVD) of $\hat{\mathbf{U}}^{S_k}$ leads to the following factorization:

$$\hat{\mathbf{U}}^{S_k} = \Psi \Lambda \Phi^*, \quad (13)$$

where Ψ is the $m \times m$ matrix of orthonormal matricial functions, Φ the matrix $n_s \times n_s$ of orthonormal frequency functions, with Φ^* its adjoint matrix, and Λ the matrix $m \times n_s$ null except for the $\inf(m, n_s)$ diagonal coefficients which contain the descending order singular values λ_i .




Figure 6: Singular values of matrix $\hat{\mathbf{U}}$ for $n_s = 81$.

Singular values decrease rapidly (see figure 6), so we prefer to use a truncated SVD to construct $\hat{\mathbf{U}}_r^{S_k} \approx \hat{\mathbf{U}}^{S_k}$ an approximation at rank r , such that $\hat{\mathbf{U}}_r^{S_k} = \Psi_r \Lambda_r \Phi_r^*$, with Ψ_r and Φ_r^* the r first columns and respectively rows of Ψ and Φ^* , and Λ_r diagonal matrix which the r first singular values of Λ . Indeed, the Eckart-Young theorem states that the truncated matrix $\hat{\mathbf{U}}_r^{S_k}$ is the closest possible matrix of rank r to the original matrix $\hat{\mathbf{U}}^{S_k}$ according to the Frobenius norm. The theorem also states that this distance is equal to the sum of the squared truncated singular values $\sum_{i=r+1}^m \lambda_i^2$. In other words, if the singular values decrease rapidly, it means that it is possible to consider low rank (few tens in this study) approximations of $\hat{\mathbf{U}}^{S_k}$ without sacrificing accuracy.

Finally, noting $\mathbf{L}_r = \Psi_r \Lambda_r$, we can write: $\hat{\mathbf{U}}_r^{S_k} = \mathbf{L}_r \Phi_r^*$ with \mathbf{L}_r of size $m \times r$ and Φ_r^* of size $r \times n_s$. Storing the last two matrices \mathbf{L}_r and Φ_r^* involves $r \times m + r \times n_s$ data, compared with $m \times n_s$ data from $\hat{\mathbf{U}}_r^{S_k}$; the smaller the

rank r , the more effective the memory saving is.

The previous reduction technique offers a final computational cost saving when it comes to interpolation of the Schur matrices in the frequency domain: instead of interpolating and evaluating the Schur complements from the $\hat{\mathbf{U}}^{S^k}$, it is possible with this technique to interpolate them on the r rank truncature only, i.e. directly on Φ_r^* . Since, as shown in the examples, $r \ll n_s$, it is far less computationally demanding.

In order to quantify the impact of this reduction technique on the quality of the simulation results, let us finally note $\tilde{\mathbf{U}}_r^{S^k}$ be $m \times N$ matrix obtained from interpolation of $\hat{\mathbf{U}}_r^{S^k}$ on the frequency range and $\tilde{\mathbf{Z}}_r$ computed with the interpolation approach on $\tilde{\mathbf{U}}_r$. The effect of this approximation on condensation matrix $\epsilon_{\mathbf{U}, \tilde{\mathbf{U}}_r}$ and on the quality of the impedance estimation $\epsilon_{\mathbf{Z}, \tilde{\mathbf{Z}}_r}$ is analysed for the smallest number of sampling frequencies in Figure 7, where blue and red coloured points correspond to these last relative errors for ranks 10 and 20 respectively. Using a too small truncation ratio $\frac{r}{n_s}$ has a visible impact on relative error on matrix $\epsilon_{\mathbf{U}, \tilde{\mathbf{U}}_r}$, then on resulting input impedance. Approximation for the rank $r = 10$ leads to an increase of almost an order of magnitude in the relative error on condensation matrices for cases $n_s = 81$ and 101 (where $\frac{r}{n_s} = 10\%$ and 12% respectively), thus relative errors greater than 10^{-5} on the impedance; and is not observed for case $n_s = 41$ of larger truncation ratio 24%. When increasing the rank to $r = 20$, the errors introduced by the truncation on either the condensation matrices or the impedance become negligible, or say, largely below the other sources of modelling approximations.

To sum up, $20m + 20n_s$ are thus enough to store the Schur complement data, compared with $n_{\text{int}}^2 N$, i.e. a reduction of almost 3 orders of magnitude in the number of data to be stored.




Figure 7: Relative errors on condensation matrices $\epsilon_{\mathbf{U}, \tilde{\mathbf{U}}}$ (light gray dashed or dotted lines) and on resulting input impedance vector $\epsilon_{\mathbf{Z}, \tilde{\mathbf{Z}}}$ (dark gray solid line) with respect to the interpolation number of frequency sampling points. The effect of truncation is shown by blue and red coloured markers, for $r = 10$ or $r = 20$ respectively.

3.3 Parallelization strategy

All codes presented in this paper are implemented with the open-source computational environment DOLFINx of FEniCS Project. Message Passing Interface (MPI) protocol [?] is used to perform parallelization at two levels: first, dividing the global communicator into S sub-communicators, it allows to perform simultaneous calculations for several sphere meshes M_{S_k} (Schur Complement computations), or on same body mesh M_B but for several frequency range (input impedance computations); second, it performs parallel solving of the previous linear systems (PETSc objects) on each sub-communicator.

A first step consists in building the H interpolation functions $\phi_r^{*S_k}(\omega)$ for all holes of the instrument. The DOLFINx mesh of the bore volume B is then built on a single CPU 0 of each sub-communicator, and thus the sequential sparse AIJ PETSc matrix \mathbf{K}^0 , \mathbf{M}^0 and \mathbf{C}^0 . Finally, the sequential sparse AIJ Schur complement matrix \mathbf{S}^0 is constructed from interpolation functions $\phi_r^{*S_{k,k \in \mathcal{F}}}$ for the considered fingering, at the frequency of interest. Parallel AIJ PETSc matrices (\cdot^{\parallel}) of the previous sequential operators are finally obtained by scattering on the (sub-)communicator. A schematic diagram of the parallel algorithm is shown in Figure 8. We used MUMPS' parallel sparse direct solver [?] for the linear system resolution.




Figure 8: Sketch of the overall algorithm and strategy for the parallel implementation.

4 Application to the analysis of a real instrument

4.1 Specimen

The specimen used in this study was an alto saxophone body with 19 toneholes, so $H = 20$, visible in figure 9. This part of the instrument was specifically chosen to validate the proposed reduced modelling strategies, as it is the part of the instrument with the lowest dimensional and geometric tolerances (max. 0.1mm, on chimney radius). No keypad were mounted on the instrument, so the chimneys were closed by brass discs and the sealing was guaranteed by a neoprene glue joint.

To ensure planar acoustic wave on the measure plan, the cylindrical tenon, normally welded to the neck, was added in the entrance of the body.

4.2 Experimental setup

The impedance measurement method used in this paper is based on the Two Microphones - Three Calibrations (TMTC) method [?]. Four microphones were used (see Figure 9), that is, 3 pairs, to avoid uncertainties when the ratio of the sensors spacing to the half-wavelength is an integer.




Figure 9: Experimental set-up in the semi-anechoic chamber at Henri Selmer Paris acoustic lab. The studied specimen is an alto saxophone body (Henri Selmer Paris - SERIES III) represented here in a configuration where all the holes are open

Only one calibration was performed for a closed (infinite impedance) impedance head configuration. Analytical formulations were used for the last two calibrations (null and unit input impedance). Controlled temperature of the semi-anechoic chamber was measured at 20.4°C during the calibration and a maximum temperature deviation of 0.2°C was allowed with the measurements.

A logarithmic chirp of 5 s was used, with a sampling frequency of 96 kHz and a 0.2 Hz frequency resolution of the measurements on the wide 20 Hz to 8000 Hz. One measurement consisted in five acquisitions, from which the mean input impedance was computed. Between all acquisition of a same measurement, the instrument position was unchanged. Conversely, between measurements, the instrument was removed from the set up, then reassembled. Finally, for each fingering, 5 measurements (each one consisting of the average of 5 chirps) were made in order to compute the average impedance, thus reducing the uncertainties associated with the operator.

4.3 Simulation

Results obtained on the specimen described above are presented in this section. Seven different configurations were studied, corresponding to the standard fingering for alto first register notes D \sharp 4, E4, F4, G4, A4, A \sharp 4 and B4 of the English naming convention.

Fingering	D#4	E4	F4	G4	A4	A#4 _p	B4
H	1	2	3	5	6	7	8
resonance 1	-10.6 ± 1.0	-3.8 ± 2.3	-6.4 ± 0.9	-2.0 ± 2.3	-11.3 ± 3.6	-5.3 ± 1.2	-1.2 ± 0.7
resonance 2	-6.2 ± 0.8	1.4 ± 0.6	-0.1 ± 0.8	3.9 ± 1.1	-5.3 ± 1.3	-7.1 ± 0.5	15.01 ± 0.9
resonance 3	-5.0 ± 0.6	-0.6 ± 0.7	1.1 ± 0.9	6.3 ± 2.2	34.3 ± 1.8	25.9 ± 1.7	-20.0 ± 0.9
resonance 4	-5.7 ± 0.6	-1.9 ± 0.8	0.5 ± 0.9	-	3.9 ± 1.7	15.5 ± 1.1	14.3 ± 2.2
resonance 5	-5.1 ± 0.5	-6.7 ± 0.9	-7.4 ± 1.1	-3.6 ± 0.9	-0.2 ± 1.78	-7.0 ± 2.0	41.2 ± 2.2
resonance 6	-6.6 ± 0.7	-10.8 ± 1.2	-0.7 ± 0.9	-14.7 ± 2.7	-31.6 ± 1.2	4.8 ± 3.0	-7.9 ± 1.9

Table 1: Discrepancy (in cents) between the predicted and measured resonance frequencies for the 6 first resonances of 7 different configurations. Colour legend: discrepancy lower than 5 cents in green, 10 cents in yellow, 20 cents in orange and red above. White cells correspond to large differences attributed to experimental biases.

Pre-preprocessing stages, CAD then meshing aspects, were fully automated and performed under the free software salome-meca [?]. The selected geometrical parameterization required the least knowledge of keyword: a normal offset of the key with respect to the chimney was imposed, so that only the distance between the section of the tonehole and the assumed cylindrical keypad, is needed. Of no consequence for this study, as no key was mounted on the body. Thus, both resolutions stages, calculation of the Schur complements and the computation on the condensed domain, were also automated.

Quadratic tetrahedral meshes of the spheres and internal body used for this study were built with a maximum cell size of 4.25 mm ($< \lambda/10$ at 8 kHz), 3.5 mm face size at walls and refined at tone hole so that there are at least 100 DOFs on the circumference, with a growth rate of 30% and 5 segments per radius at the cone-chimney junction, see Figure 4. Sampling frequency step Δf_s was set to 100 Hz ($n_s = 81$) for the computation and the truncature rank chosen was 20 the storage of approximated Schur complement matrices. Such mesh parameters led to systems of only 800 kDOFs for the condensed body.

All calculations were performed on a (Intel(R) Xeon(R) Gold 6230 CPU @ 2.10 GHz) machine, with 2 sockets of 20 CPU - 40 threads each, and 126 Go total RAM. The global communicator was divided into five parallel 16-threads sub-communicators, on which the 7070 steps of the frequency range were split.

4.4 Validation of the numerical approach

Two types of output of this high fidelity FE model are analyzed in this section: the classical quantities of interest, namely the input impedance, but also the pressure fields, which are additional outputs from the model, that provide a wealth of other information, such as the precise location of pressure nodes and anti-nodes, which could be invaluable for the design of the chimney and keys. In this paragraph, the first quantity is used to validate the model by comparison with experiments ; the second provides a better understanding of the conditions under which the proposed modelling strategies are most effective.

Figure 10 presents a comparison of the measured and predicted input impedance curves for three different configurations corresponding to the D#4, F4 and A4 fingerings (opened holes of each configuration are shown with green interfaces). An extremely good agreement can be observed between the experimental measurement (dashed blue line) and the numerical solution computed with the proposed model (solid red line), both on magnitude and phase. However, such a visual match is not sufficient to validate a model for musical acoustics given the high sensitivity of musicians mentioned in the introduction. To go further in the analysis of the accuracy of the prediction, table 1 summarizes the discrepancy between simulations and measurements (in cents) for the first 6 resonances for 8 different fingerings. Cell colours are green, yellow or orange depending on whether the minimum deviation is less than 5, 10 or 20 cents respectively, and red above (peak 3 for G4 fingering is empty because of experimental issues - no phase sign change and no clear maximum for the peak).

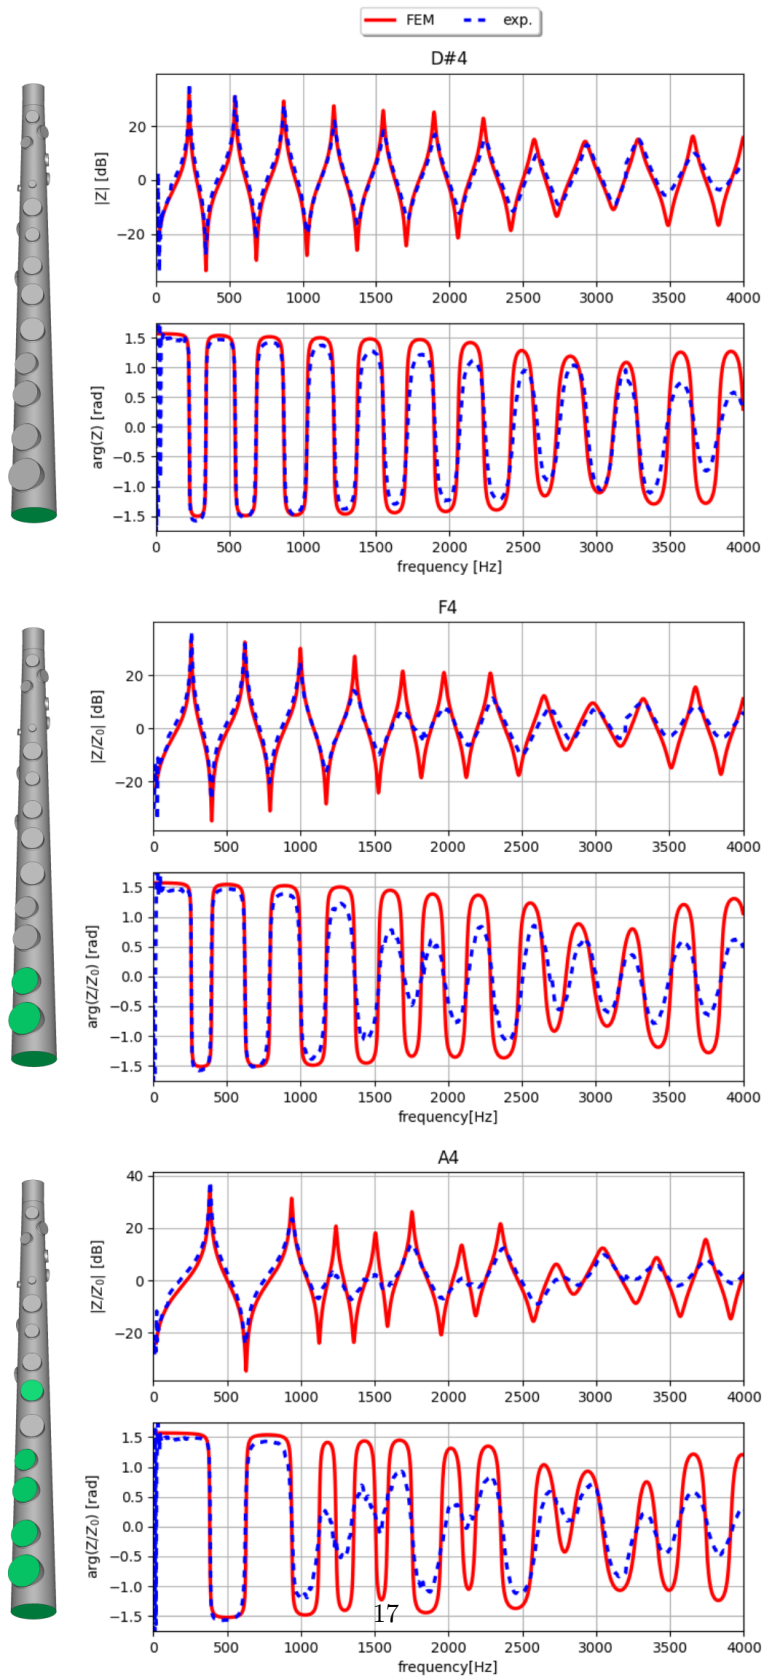


Figure 10: Comparison of the input impedance predicted by the proposed modelling strategy (solid red) and measured experimentally (dashed blue) for the 3 configurations: D#4 (top), F4 (middle) and A4 (bottom). On the left side, the corresponding fingerings (open toneholes in green) and on the right the magnitude (top) and phase (bottom)

Remark. Estimating the resonant frequency from the measured impedance was most often done by locating where the phase was zero. Sometimes the phase simply did not change sign and in these situations the maximum amplitude was used (italic values in Tab. 1). This explains why, to minimize the addition of experimental and post-processing uncertainties, 25 measurements (with 5 bench assembly and disassembly) were performed for each fingering as explained in section 4.2.

It can be seen that 75% of resonance frequencies are correctly estimated with minimum deviations of less than 10 cents, and even 51% very correctly within 5 cents. For all configurations, the position and amplitude of the first 2 peaks are always correctly estimated (see Tab.1). For subsequent resonances, dissipation is sometimes misestimated by the model, which leads to higher amplitudes of input impedance peaks, in the high frequency range but also in the middle range as visible on peaks 5 and 6 of configuration F4, with deviations of 15.7 ± 0.0 and 12.9 ± 0.0 dB respectively, without affecting the estimation of resonance frequencies for this case.

For configurations A4, A#4 and B4, significant deviations were observed from the measurement at very specific resonances. The acoustic pressure fields (real part) of these configurations for different resonances (largest cent errors identified by circles) are shown in Figure 11. The arrows indicate the opened tone holes for each of the fingerings.

Tone holes suspected to explain these discrepancies were designated by circled arrows. Non-negligible acoustic pressure was observed at the second opened chimney (counting from input section). The relatively high pressure on these open holes is explained because of the presence of anti-node of high pressure located at the neighbouring closed chimney. More precisely, for peaks 3 and 6 of configuration A4, the average pressure at the tone hole section is about 0.5 Pa for opened chimneys on either side of the closed hole (1 Pa) of this cross-fingering. Conversely, for B4, a significant depression was observed at the second open tone hole.

These higher discrepancies could be explained by the ‘blind’ sphere strategy, which neglects external interactions between holes. This phenomenon is known to occur more likely in the presence of large diameter short height chimneys close to each other [?], see first two opened tone holes of configuration A#4 for instance. If this intuition is confirmed by future investigations, a computationally still efficient alternative could consist in considering a larger blind sphere with two interacting holes in it.

Simpler 1D calculations were performed using the open source OpenWInD software [?]. Infinite flanged or unflanged tonehole radiation models were considered for the TMM computations, in order to compare the ranges obtained with the proposed reduced FE model. Taking into account 3D effects seems necessary as part of a predictive approach as shown in graph in Figure 12 which compares the deviations in cents between the FE model (coloured histogram) and the ranges obtained from TMM computations (black lines delimiting the values obtained with the different tonehole radiation models). The reduced 3D FEM approach gives a more accurate prediction for 83% of the peaks compared to the TMM, with a mean deviation of 10.4 cents compared with 28.6 cents respectively.

Apart from these specific situations, static condensation is an effective modelling strategy. Indeed it has made it possible to reduce the number of DOFs by at least a factor of three compared to a standard approach with the use of blind spheres (compared to the case with a unique large surrounding sphere the gain is even more important), while still meeting musicians’ accuracy criteria in terms of tuning. On the Table 2 we present a comparison between our different methods and a standard FEM about computational data like the number of linear problem to solve and its size in term of degrees of freedom, the estimated CPU time and the required storage. It shows that with a limited storage (about 1 Go) we can reach a speed-up of almost 60 compared to a standard FEM approach using a unique large surrounding sphere.

Moreover considering that the computational resources are limited (for example a fixed number of CPU), reducing the size of the linear problem is even more interesting than just solving the linear problems quickly. In fact, since it takes fewer CPUs to solve a linear problem, we can use many more to perform frequency-domain parallelism and win on both ends.

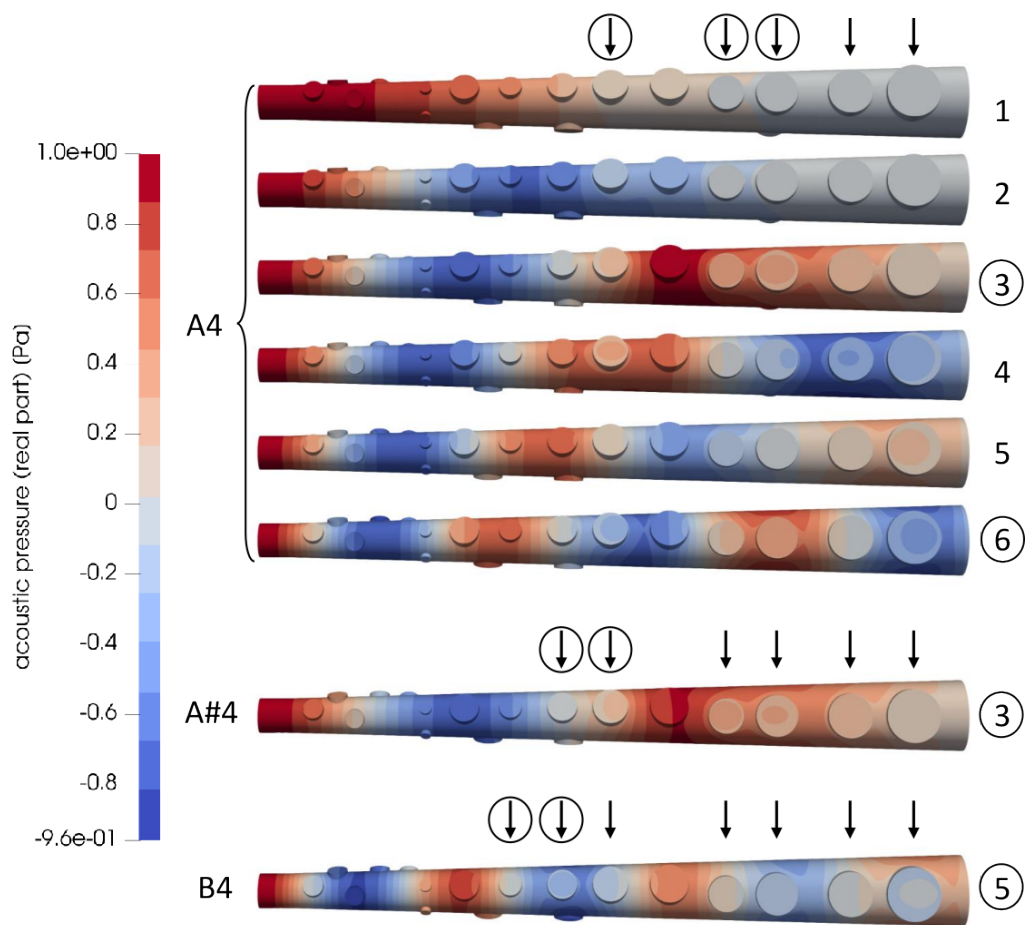


Figure 11: Predicted pressure fields (real part) in the instrument body for different resonances of configurations A4, A#4 and B4. The arrows indicate the opened tone holes for each of the fingerings. Circles indicate configurations with larger cent discrepancies and corresponding suspected tone holes.




Figure 12: Accuracy on the resonance frequencies prediction of the proposed reduced FEM modelling strategy (coloured bars) and standard TMM (black ranges) for the 7 configurations.

	nDOF (per pb)	Number of problems	Estimated CPU per fingering	Storage
Standard FEM	$\sim 10^8$	$2N_f N = 700k$	~ 1 month	-
+ incr. Cremer BLI	$\sim 10^8$	$N_f N = 350k$	~ 15 days	-
+ condensation	$8 \cdot 10^5$	$N_f N = 350k$	2.5 days	$mNH = 560$ Go
+ freq. parallelism	$8 \cdot 10^5$	$N_f N = 350k$	12 h	$mNH = 560$ Go
+ interpolation	$8 \cdot 10^5$	$N_f N = 350k$	12 h	$mn_s H = 6.48$ Go
+ reduction	$8 \cdot 10^5$	$N_f N = 350k$	12 h	$(m + n_s)rH = 1.6$ Go

Table 2: Illustration of the computer requirements to solve the different numerical approaches. The numbers are computed with $N_f = 50$ fingerings, $N = 7000$ frequency steps, $n_s = 81$ sampling points, $n_{\text{int}} = 700$ interface DOFs thus $m = 250k$ independant schur complement terms, reduction rank $r = 20$ and $H = 20$ spheres because of the 19 tone holes. There are $S = 5$ sub-communicators and 16-thread per sub-communicator. The CPU time for one problem on the bore only was $t_{\text{bore}} = 30$ s. To give an order of magnitude, we estimated the size of the FE problem including the bore and one large sphere modelling the exterior domain to $n_{\text{DOF}} \approx 10$ millions of DOFs and an associated CPU time of $t_{\text{full}} = 3$ min (just a coarse estimation as this problem was actually not solved).

5 Conclusions and perspectives

In this article, we present a 3D-FEM approach to solve the 3D scalar Helmholtz problem, dedicated to wind instruments. Computational costs, both in terms of number of resolutions (3.1), and memory, were reduced by proposing an external domain condensation technique. It is based on independent blind spheres (3.2.1), associated with a reduction strategy of the Schur complements estimation (3.2.2) combining interpolation and low-rank approximation. These combined strategies make it possible to gain more than an order of magnitude in input impedance calculation time, while maintaining high predictive accuracy in estimating playable notes (75% of the first 6 resonances of the 7 fingerings studied show a deviation of less than 10 cents from experiment). If the input impedance is the primary quantity of interest for the manufacturer, an added benefit of this new approach is the ability to visualize and quantitatively analyse pressure and velocity fields on complex 3D geometries, allowing for more detailed consideration of design parameters. For example it is now possible to consider conducting sensitivity studies that take into account more geometric details, such as fillets or chamfers.

Although the method presented in this article drastically reduces the computation cost of high fidelity acoustic FE models of wind instruments, it remains computationally significant. For instance it is still too large to consider using it in an automatic design optimisation algorithm. However, there is still room for improvement. A better parallelism management of the meshes and the assembly of the finite elements operators could be performed. Moreover, given the regularity of the solution fields for this musical application, other reduction approaches at the numerical method level itself could be considered: the use of Splines [?] in substitution to classic finite elements to reduce the approximation subspace to a smaller set of regular functions might be a good candidate. An adaptive strategy in the frequency domain could also be devised to limit the computational effort around the the resonances only, which would even further reduce the simulation times.