

A mathematical theory for mass lumping and its generalization with applications to isogeometric analysis

Yannis Voet ^{*1}, Espen Sande ^{†1}, and Annalisa Buffa ^{‡1}

¹MNS, Institute of Mathematics, École polytechnique fédérale de Lausanne, Station 8, CH-1015 Lausanne, Switzerland

December 7, 2022

Abstract

Explicit time integration schemes coupled with Galerkin discretizations of time-dependent partial differential equations require solving a linear system with the mass matrix at each time step. For applications in structural dynamics, the solution of the linear system is frequently approximated through so-called mass lumping, which consists in replacing the mass matrix by some diagonal approximation. Mass lumping has been widely used in engineering practice for decades already and has a sound mathematical theory supporting it for finite element methods using the classical Lagrange basis. However, the theory for more general basis functions is still missing. Our paper partly addresses this shortcoming. Some special and practically relevant properties of lumped mass matrices are proved and we discuss how these properties naturally extend to banded and Kronecker product matrices whose structure allows to solve linear systems very efficiently. Our theoretical results are applied to isogeometric discretizations but are not restricted to them.

Keywords: Isogeometric analysis, explicit dynamics, mass lumping, linear algebra, Loewner partial order

1 Introduction and background

Mass lumping dates from the early days of the finite element method, in the 1970s, when practitioners were facing the burden of inverting non-diagonal mass matrices in explicit time integration schemes for solving wave propagation problems. They quickly discovered that reasonably good solutions could still be recovered after replacing the mass matrix by some easily invertible (typically diagonal) ad hoc approximation, a procedure which became known as mass lumping. The mass lumping approach relied on physical intuition of mass conservation [1, 2]. Numerous strategies were proposed but all later attention in modern literature has focused on only three of them: the row-sum technique, the nodal quadrature method and a so-called “special mass lumping” scheme [2], more commonly referred to as the diagonal scaling procedure [3].

The row-sum method was among the first techniques proposed. As the name suggests, the row-sum lumped mass matrix is a diagonal matrix where the i th diagonal element is simply the sum of all entries on the i th row. Apart from being straightforward to implement, it is widely used and appreciated for the improved stability in connection with explicit time integration schemes, a property we will formally prove in this article. However, this lumping technique is prone to failure by generating singular or indefinite lumped mass matrices, leading to erroneous dynamical behaviors. Nonnegative partition of unity methods are exempt from such breakdowns as they always deliver positive definite row-sum lumped mass matrices. Therefore, the emergence of new discretization methods based on a nonnegative partition of unity, such as isogeometric analysis [4, 5], has sparked renewed interest in the row-sum technique [6, 7]. Unfortunately, the row-sum technique in isogeometric analysis leads to a loss of convergence rate for the smallest generalized eigenvalues. As a matter of fact, it was shown in [6] for 1D problems and for quadratic and cubic B-splines that the convergence rate is downgraded to quadratic order when using the row-sum lumped mass matrix. To this date, high order mass lumping in isogeometric analysis is an open research problem and has only been achieved (to our knowledge) in [8] using dual basis functions.

Another mass lumping technique specific to classical finite element methods is the nodal quadrature method. In this method, the interpolation property of Lagrange basis function is cleverly used to construct diagonal matrices by choosing quadrature nodes that coincide with the finite element nodes. Although this technique might preserve

*yannis.voet@epfl.ch

†espen.sande@epfl.ch

‡annalisa.buffa@epfl.ch

the higher order convergence of the consistent mass [9, 10], it also breaks down in case nonpositive quadrature weights are used. Positive weights can be guaranteed by either choosing specific quadrature rules and sacrificing some accuracy [3, 11] or enlarging the finite element space by introducing additional degrees of freedom [12]. Both solutions come with their own drawbacks: the first one does not preserve the rate of convergence while the second one requires solving a larger system of equations. In either case, research is hampered by tedious element specific analysis. Other authors, instead, have suggested modifying the time integration scheme to accommodate nonpositive weights. In [13], a mixed implicit-explicit time integration scheme is designed to cope with positive semidefinite (potentially singular) lumped mass matrices.

Finally, the diagonal scaling method [14] was designed to circumvent breakdowns in the row-sum and nodal quadrature methods. Positive definite lumped mass matrices are indeed guaranteed by construction. The method scales the diagonal of the consistent mass matrix, which is necessarily positive, by a positive factor that depends on the mass of the element.

Despite appealing numerical results, the mathematical theory supporting these methods is rather limited and usually restricted to specific discretization methods, orders and element types. Although a completely unified analysis is indeed probably intractable, a framework with the least restrictive assumptions is sought. For this reason, our analysis will be restricted to the row-sum technique because it allows us to draw very general conclusions. It is based on purely algebraic considerations and is therefore independent of the discretization method used. Moreover, although the three aforementioned mass lumping strategies generally yield different lumped mass matrices, they might be identical in some cases. In [3], for instance, equivalence conditions for quadrature rules have been established for the Lagrange basis. Thus, although our conclusions will generally not apply to other mass lumping strategies, they might still hold under certain circumstances.

The rest of the paper is structured as follows: in Section 2, we briefly recall the discretization of a model wave equation and the problem arising from explicit time integration methods for the semi-discrete problem. The generalized eigenvalues of the matrix pencil formed by the stiffness and mass matrices are critical for the dynamics and stability of explicit methods. Section 3 is designed to bring the reader up to date by recalling well-known theoretical results for generalized eigenproblems and paves the way for the upcoming sections. Section 4 is devoted to the theoretical analysis of perturbed generalized eigenproblems. Immediate conclusions can be drawn regarding the stability of explicit time integration methods in connection with mass lumping. Some of our results for high dimensional discretizations assume the mass matrix is expressed as a Kronecker product, a structure inherited by tensor product basis functions. This is generally not the case in practice for a variety of reasons. Therefore, Section 5 explains how the mass matrix may nevertheless be approximated by a Kronecker product. Established results can then be applied to the approximation itself. Numerical experiments are presented along the way and closely follow the theory. They serve multiple purposes by illustrating the findings, highlighting their limitations and drawing additional insight. In order to foster applications to other problems, theoretical results are stated for general matrices satisfying certain properties. Finally, Section 6 concludes our article and summarizes our findings.

2 Mass lumping in explicit dynamics

We consider as model problem the classical wave equation. Let $\Omega \subset \mathbb{R}^d$ be an open connected domain with Lipschitz boundary and let $I = [0, T]$ be the time domain with $T > 0$ denoting the final time. We look for $u: \Omega \times [0, T] \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \partial_{tt}^2 u(\mathbf{x}, t) - \Delta u(\mathbf{x}, t) &= f(\mathbf{x}, t) && \text{in } \Omega \times (0, T], \\ u(\mathbf{x}, t) &= 0 && \text{on } \partial\Omega \times (0, T], \\ u(\mathbf{x}, 0) &= u_0(\mathbf{x}) && \text{in } \Omega, \\ \partial_t u(\mathbf{x}, 0) &= v_0(\mathbf{x}) && \text{in } \Omega, \end{aligned} \tag{2.1}$$

where u_0 and v_0 are some initial displacement and velocity, respectively, and we prescribe homogeneous Dirichlet boundary conditions for simplicity. A Galerkin discretization of the spatial variables (see for instance [2, 15]) leads to solving the semi-discrete problem

$$\begin{aligned} M\ddot{\mathbf{x}}(t) + K\mathbf{x}(t) &= \mathbf{f}(t) && \text{for } t \in [0, T], \\ \mathbf{x}(0) &= \mathbf{x}_0, \\ \dot{\mathbf{x}}(0) &= \mathbf{v}_0. \end{aligned} \tag{2.2}$$

where $M, K \in \mathbb{R}^{n \times n}$ are the mass and stiffness matrices, respectively, and $\mathbf{x}(t) \in \mathbb{R}^n$ is a vector containing the coefficients of the approximate solution $u_h(\mathbf{x}, t)$ in the basis of some finite dimensional subspace. The time-dependent right hand side $\mathbf{f}(t) \in \mathbb{R}^n$ accounts for the forcing function f and potential non-homogeneous Neumann and Dirichlet

boundary conditions. Naturally, the properties of M and K depend on the discretization method. Nevertheless, they are always symmetric with M being positive definite for the standard Galerkin method. Provided Dirichlet boundary conditions are prescribed on some portion of the boundary, which is necessary for the well-posedness of the problem, K is also positive definite. Therefore, we will assume throughout the paper that both M and K are symmetric positive definite.

There exists a host of numerical methods for approximating the solution of the semi-discrete problem (2.2). The Newmark method [16], proposed in 1959 and summarized in Algorithm 2.1, has been historically used and still remains very popular today for structural problems with fast dynamic loadings such as blasts, impacts or earthquakes. The time domain $[0, T]$ is first discretized using a uniform time step $\Delta t = \frac{T}{N}$ such that $t_s = s\Delta t$, $s = 0, 1, \dots, N$ with $N \in \mathbb{N}^*$. The method then consists in approximating the solution $\mathbf{x}(t)$ as well as its first and second time derivative (interpreted as the displacement, velocity and acceleration, respectively) at the discrete times t_s . These quantities are denoted \mathbf{x}_s , \mathbf{v}_s and \mathbf{a}_s in Algorithm 2.1.

Input: Initial conditions $\mathbf{x}_0, \mathbf{v}_0$, matrices M, K and parameters β and γ

Output: Approximate displacement $\{\mathbf{x}_s\}_{s=0}^N$, velocity $\{\mathbf{v}_s\}_{s=0}^N$ and acceleration $\{\mathbf{a}_s\}_{s=0}^N$

- 1: Compute $\mathbf{a}_0 = M^{-1}(\mathbf{f}(t_0) - K\mathbf{x}_0)$
- 2: **for** $s = 0, 1, \dots, N - 1$ **do**
- 3: Compute $\tilde{\mathbf{x}}_{s+1} = \mathbf{x}_s + \Delta t\mathbf{v}_s + (\frac{1}{2} - \beta)\Delta t^2\mathbf{a}_s$
- 4: Compute $\tilde{\mathbf{v}}_{s+1} = \mathbf{v}_s + (1 - \gamma)\Delta t\mathbf{a}_s$
- 5: Compute $\mathbf{a}_{s+1} = (M + \beta\Delta t^2K)^{-1}(\mathbf{f}(t_{s+1}) - K\tilde{\mathbf{x}}_{s+1})$
- 6: Compute $\mathbf{v}_{s+1} = \tilde{\mathbf{v}}_{s+1} + \gamma\Delta t\mathbf{a}_{s+1}$
- 7: Compute $\mathbf{x}_{s+1} = \tilde{\mathbf{x}}_{s+1} + \beta\Delta t^2\mathbf{a}_{s+1}$
- 8: **end for**

Algorithm 2.1: Newmark method

Both explicit and implicit methods are recovered for specific combinations of two scalar parameters β and γ in $[0, 1]$. Explicit methods are recovered for $\beta = 0$ while implicit unconditionally stable methods arise for $\beta \geq \frac{\gamma}{2} \geq \frac{1}{4}$. Whereas explicit methods are necessarily conditionally stable, implicit methods may or may not be unconditionally stable. The central difference method for instance is an explicit second order method and its stability analysis reveals that the critical time step is given by $\Delta t_c = \frac{2}{\sqrt{\lambda_n(K, M)}}$, where $\lambda_n(K, M)$ is the largest eigenvalue of the generalized eigenvalue problem $K\mathbf{u} = \lambda M\mathbf{u}$. Unfortunately, rapidly changing boundary conditions require a sufficiently small time step for us to resolve the correct physical behavior. Thus, the unconditional stability of some implicit methods cannot be fully exploited while the computational burden prevails in Line 5. Indeed, for very large problems, solving the linear system in Line 5 at each time step becomes a major computational bottleneck. However, for explicit methods ($\beta = 0$), this linear system only involves the mass matrix. Therefore, replacing the consistent mass matrix in Lines 1 and 5 with some kind of easily invertible approximation \tilde{M} becomes very appealing. Mass lumping stems from this observation. In fact, in engineering practice all kinds of ad hoc replacements may be made involving also the stiffness matrix or the right-hand side vector but we will not consider those here.

While the computational benefits of mass lumping are obvious, the substitution of M by \tilde{M} both effects the stability of the time integration method and the accuracy of the approximate solution. Strong numerical evidence suggests that mass lumping improves the stability by increasing the critical time step [7, 8]. Verifying this claim amounts to guaranteeing that $\lambda_n(K, \tilde{M}) \leq \lambda_n(K, M)$. On the other hand, regarding accuracy, the smallest generalized eigenvalues of (K, M) usually play a prominent role although the largest ones might also sometimes be involved depending on the initial conditions and loading [17]. Thus, the smallest eigenvalues of (K, \tilde{M}) and (K, M) must be very close. Consequently, both questions of stability and accuracy can be tackled by comparing the eigenvalues and eigenspaces of some matrix pencils. The upcoming analysis requires some theoretical understanding of generalized eigenvalue problems. All necessary background knowledge will be covered in the next section.

3 Generalized eigenvalue problems

Given the important role generalized eigenproblems play in this work, this section provides a brief reminder of some important properties. Whereas standard eigenproblems are clearly a particular case of generalized eigenproblems, results for standard eigenproblems generally cannot be straightforwardly extended to the generalized case. Some of the differences are summarized below. Since all matrices we will be considering for our applications are symmetric, we will restrict the discussion to symmetric generalized eigenvalue problems. Let $A, B \in \mathbb{R}^{n \times n}$ be symmetric matrices. We look for pairs $(\mathbf{u}, \lambda) \in \mathbb{C}^n \times \mathbb{C}$ such that $A\mathbf{u} = \lambda B\mathbf{u}$. In analogy to standard eigenproblems, the generalized eigenvalues are the roots of the characteristic polynomial $p(\lambda) = \det(A - \lambda B)$.

Definition 3.1 (Singular matrix pencil). A matrix pencil (A, B) is called singular if $\det(A - \lambda B) = 0$ for all $\lambda \in \mathbb{C}$.

Clearly, matrix pencils formed by matrices A and B whose null space intersects nontrivially are necessarily singular. It can be achieved very easily, for example by taking

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

Moreover, contrary to standard eigenvalue problems, symmetry of the coefficient matrices is not sufficient to guarantee real eigenvalues. This point must be emphasized to avoid common misconceptions. A sufficient condition for real finite eigenvalues, which will cover all our applications, is given in the following lemma [18]. The proof is included for completeness.

Lemma 3.2. Let (A, B) be a symmetric pencil with B positive definite. Then, all generalized eigenvalues are real and there exists an invertible matrix $U \in \mathbb{R}^{n \times n}$ such that

$$U^T A U = D = \text{diag}(\lambda_1, \dots, \lambda_n), \quad U^T B U = I,$$

where D is a real diagonal matrix containing the eigenvalues.

Proof. Since B is assumed positive definite, it can be expressed using its Cholesky factorization as $B = LL^T$, where L is a lower triangular matrix. Then, the generalized eigenvalue problem $A\mathbf{u} = \lambda LL^T \mathbf{u}$ can be transformed to standard form by pre-multiplying by L^{-1} and writing $\tilde{A}\tilde{\mathbf{u}} = \lambda\tilde{\mathbf{u}}$ with $\tilde{A} = L^{-1}AL^{-T}$ and $\tilde{\mathbf{u}} = L^T \mathbf{u}$. Since \tilde{A} is symmetric, its eigenvalues are all real. Moreover, by the spectral decomposition, $\tilde{A} = \tilde{U}D\tilde{U}^T$ where D is the diagonal matrix of eigenvalues and \tilde{U} is the orthogonal matrix of eigenvectors. We now back-transform to the original problem by substituting the definition of \tilde{A} and noting that $L^{-1}BL^{-T} = I$

$$\begin{aligned} \tilde{U}^T \tilde{A} \tilde{U} &= D, & \tilde{U}^T L^{-1} A L^{-T} \tilde{U} &= D, \\ \tilde{U}^T \tilde{U} &= I, & \tilde{U}^T L^{-1} B L^{-T} \tilde{U} &= I. \end{aligned}$$

Defining $U = L^{-T} \tilde{U}$ concludes the proof. \square

More can be said if both A and B are symmetric positive definite. The following implications are a straightforward consequence of Lemma 3.2.

- If A is symmetric positive semidefinite and B is symmetric positive definite, then $\Lambda(A, B) \subset [0, +\infty)$,
- If A and B are symmetric positive definite, then $\Lambda(A, B) \subset (0, +\infty)$.

We conclude this section by introducing a few notations. $\Lambda(A, B)$ will denote the spectrum (i.e. the set of all eigenvalues) of the matrix pencil (A, B) whereas $\lambda(A, B)$ will refer to individual eigenvalues. For standard eigenvalue problems (when $B = I$), we will simply write $\lambda(A)$ instead of $\lambda(A, I)$. The matrices will be explicitly specified when referring to individual eigenvalues or to the entire spectrum if there is potential confusion. The eigenvalues will always be numbered in ascending algebraic order such that

$$\lambda_1(A, B) \leq \lambda_2(A, B) \leq \dots \leq \lambda_n(A, B).$$

4 Mathematical properties of mass lumping and its generalization

4.1 Mathematical properties of mass lumping

An important concept used throughout the analysis is the notion of equivalent pencils, which is motivated through the following lemma. Its proof is a simple exercise in [19, Chapter 15].

Lemma 4.1. Let $A, B, E, F \in \mathbb{R}^{n \times n}$ and assume E and F are invertible. Then

$$\Lambda(A, B) = \Lambda(EAF, EBF)$$

Proof. Indeed,

$$\begin{aligned} \det(EAF - \lambda EBF) &= 0 \\ \iff \det(E) \det(A - \lambda B) \det(F) &= 0 \end{aligned}$$

Since we have assumed E and F are invertible ($\det(E) \neq 0$ and $\det(F) \neq 0$), then $\det(EAF - \lambda EBF) = 0 \iff \det(A - \lambda B) = 0$. Consequently, $\Lambda(A, B) = \Lambda(EAF, EBF)$. \square

Definition 4.2 (Equivalent pencils). For matrices $A, B, E, F \in \mathbb{R}^{n \times n}$ with E, F invertible, the matrix pencils (A, B) and (EAF, EBF) are said to be *equivalent*. They are called *congruent* if $E = F^T$.

The next theorem will play a central role in our analysis. Its proof is a direct consequence of the celebrated Courant-Fischer theorem.

Theorem 4.3. Let $A, B, C \in \mathbb{R}^{n \times n}$ be symmetric positive definite matrices and let all eigenvalues be numbered in ascending algebraic order. Then

$$\lambda_k(A, C)\lambda_1(C, B) \leq \lambda_k(A, B) \leq \lambda_k(A, C)\lambda_n(C, B) \quad 1 \leq k \leq n, \quad (4.1a)$$

$$\lambda_1(A, C)\lambda_k(C, B) \leq \lambda_k(A, B) \leq \lambda_n(A, C)\lambda_k(C, B) \quad 1 \leq k \leq n. \quad (4.1b)$$

Proof. Let $\mathbf{u}_i, \mathbf{v}_i$ and \mathbf{w}_i denote the eigenvectors associated to the eigenvalues $\lambda_i(A, B)$, $\lambda_i(A, C)$ and $\lambda_i(C, B)$, respectively. We define the subspaces

$$\begin{aligned} \mathcal{U}_k &= \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_k\}, & \mathcal{U}'_k &= \text{span}\{\mathbf{u}_k, \dots, \mathbf{u}_n\}, \\ \mathcal{V}_k &= \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}, & \mathcal{V}'_k &= \text{span}\{\mathbf{v}_k, \dots, \mathbf{v}_n\}, \\ \mathcal{W}_k &= \text{span}\{\mathbf{w}_1, \dots, \mathbf{w}_k\}, & \mathcal{W}'_k &= \text{span}\{\mathbf{w}_k, \dots, \mathbf{w}_n\}. \end{aligned}$$

Since all matrices are assumed symmetric positive definite, the generalized eigenvectors of all matrix pencils are linearly independent (see Lemma 3.2). Consequently, $\dim(\mathcal{U}_k) = \dim(\mathcal{V}_k) = \dim(\mathcal{W}_k) = k$ and $\dim(\mathcal{U}'_k) = \dim(\mathcal{V}'_k) = \dim(\mathcal{W}'_k) = n - k + 1$. By the Courant-Fischer theorem (see for instance [20, Theorem 4.2.6])

$$\lambda_k(A, B) = \min_{\substack{\mathcal{S} \subseteq \mathbb{C}^n \\ \dim(\mathcal{S})=k}} \max_{\substack{\mathbf{x} \in \mathcal{S} \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T B \mathbf{x}} = \max_{\substack{\mathcal{S} \subseteq \mathbb{C}^n \\ \dim(\mathcal{S})=n-k+1}} \min_{\substack{\mathbf{x} \in \mathcal{S} \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T B \mathbf{x}} \quad 1 \leq k \leq n.$$

We already know that the minimum in the second expression is attained for $\mathcal{S} = \mathcal{U}_k$ and the maximum in the third expression is attained for $\mathcal{S} = \mathcal{U}'_k$. Moreover, we note that

$$\lambda_k(A, B) = \min_{\substack{\mathcal{S} \subseteq \mathbb{C}^n \\ \dim(\mathcal{S})=k}} \max_{\substack{\mathbf{x} \in \mathcal{S} \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T A \mathbf{x} \mathbf{x}^T C \mathbf{x}}{\mathbf{x}^T C \mathbf{x} \mathbf{x}^T B \mathbf{x}} \leq \max_{\substack{\mathcal{X}_k \subseteq \mathbb{C}^n \\ \dim(\mathcal{X}_k)=k}} \frac{\mathbf{x}^T A \mathbf{x} \mathbf{x}^T C \mathbf{x}}{\mathbf{x}^T C \mathbf{x} \mathbf{x}^T B \mathbf{x}} \leq \max_{\substack{\mathcal{X}_k \subseteq \mathbb{C}^n \\ \dim(\mathcal{X}_k)=k}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T C \mathbf{x}} \max_{\substack{\mathcal{X}'_k \subseteq \mathbb{C}^n \\ \dim(\mathcal{X}'_k)=n-k+1}} \frac{\mathbf{x}^T C \mathbf{x}}{\mathbf{x}^T B \mathbf{x}},$$

and

$$\lambda_k(A, B) = \max_{\substack{\mathcal{S} \subseteq \mathbb{C}^n \\ \dim(\mathcal{S})=n-k+1}} \min_{\substack{\mathbf{x} \in \mathcal{S} \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T A \mathbf{x} \mathbf{x}^T C \mathbf{x}}{\mathbf{x}^T C \mathbf{x} \mathbf{x}^T B \mathbf{x}} \geq \min_{\substack{\mathcal{X}'_k \subseteq \mathbb{C}^n \\ \dim(\mathcal{X}'_k)=n-k+1}} \frac{\mathbf{x}^T A \mathbf{x} \mathbf{x}^T C \mathbf{x}}{\mathbf{x}^T C \mathbf{x} \mathbf{x}^T B \mathbf{x}} \geq \min_{\substack{\mathcal{X}'_k \subseteq \mathbb{C}^n \\ \dim(\mathcal{X}'_k)=n-k+1}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T C \mathbf{x}} \min_{\substack{\mathcal{X}_k \subseteq \mathbb{C}^n \\ \dim(\mathcal{X}_k)=k}} \frac{\mathbf{x}^T C \mathbf{x}}{\mathbf{x}^T B \mathbf{x}},$$

for any given subsets $\mathcal{X}_k \subseteq \mathbb{C}^n$ and $\mathcal{X}'_k \subseteq \mathbb{C}^n$ with $\dim(\mathcal{X}_k) = k$ and $\dim(\mathcal{X}'_k) = n - k + 1$. The strategy now consists in cleverly choosing these subsets. In particular, for the upper bound, we obtain

$$\begin{aligned} \lambda_k(A, B) &\leq \lambda_k(A, C)\lambda_n(C, B) && \text{for } \mathcal{X}_k = \mathcal{V}_k, \\ \lambda_k(A, B) &\leq \lambda_n(A, C)\lambda_k(C, B) && \text{for } \mathcal{X}_k = \mathcal{W}_k, \end{aligned}$$

and for the lower bound

$$\begin{aligned} \lambda_k(A, B) &\geq \lambda_k(A, C)\lambda_1(C, B) && \text{for } \mathcal{X}'_k = \mathcal{V}'_k, \\ \lambda_k(A, B) &\geq \lambda_1(A, C)\lambda_k(C, B) && \text{for } \mathcal{X}'_k = \mathcal{W}'_k. \end{aligned}$$

□

The last theorem is useful in understanding the relationship between the eigenvalues of (K, M) and those of (K, \tilde{M}) following a worst case analysis. Provided \tilde{M} is symmetric positive definite, a straightforward application of Theorem 4.3, inequalities (4.1a), with $A = K$, $B = \tilde{M}$ and $C = M$ leads to

$$\lambda_k(M, \tilde{M}) \leq \frac{\lambda_k(K, \tilde{M})}{\lambda_k(K, M)} \leq \lambda_n(M, \tilde{M}) \quad 1 \leq k \leq n. \quad (4.2)$$

This result can also be alternatively deduced starting from Ostrowski's theorem (see for instance [20, Theorem 4.5.9]). An immediate conclusion can be drawn from these bounds: if \tilde{M} is very good preconditioner for M , then the entire spectrum is well approximated since both the lower and upper bounds are close to 1. They reveal the close connection between the quality of a preconditioner and the approximation of the spectrum. However, not all preconditioners are appropriate for our problem. Indeed, one of our main objectives when designing a preconditioner \tilde{M} is to ensure that $\lambda_k(K, \tilde{M}) \leq \lambda_k(K, M)$ for all $k = 1, \dots, n$. This requirement is sometimes easily fulfilled by simply understanding the relation between M and \tilde{M} in the Loewner partial order.

Definition 4.4 (Loewner partial order). For two symmetric matrices $A, B \in \mathbb{R}^{n \times n}$, we write $A \succeq B$ (respectively $A \succ B$) if $A - B$ is positive semidefinite (respectively positive definite).

The next lemma gathers some equivalent sufficient conditions guaranteeing that $\lambda_k(K, \tilde{M}) \leq \lambda_k(K, M)$ for all $k = 1, \dots, n$.

Corollary 4.5. Let $A, B, \tilde{B} \in \mathbb{R}^{n \times n}$ be symmetric positive definite matrices, and denote $E = \tilde{B} - B$. Then the statements

1. $E \succeq 0$,
2. $\tilde{B} \succeq B$,
3. $\Lambda(B, \tilde{B}) \subset (0, 1]$,

are all equivalent and imply that $\lambda_k(A, \tilde{B}) \leq \lambda_k(A, B)$ for all $k = 1, \dots, n$.

Proof. We first prove the equivalence between the statements:

(1) \iff (2): $\tilde{B} \succeq B \iff \tilde{B} - B = E \succeq 0$.

(2) \iff (3): Since B and \tilde{B} are symmetric positive definite, $\lambda_k(B, \tilde{B}) > 0$ for all $k = 1, \dots, n$. Moreover, $\tilde{B} \succeq B \iff \mathbf{u}^T \tilde{B} \mathbf{u} \geq \mathbf{u}^T B \mathbf{u} \forall \mathbf{u} \in \mathbb{R}^n \iff \lambda_n(B, \tilde{B}) \leq 1$ where the last equivalence follows from the Courant-Fischer theorem.

We now show that $\lambda_k(A, \tilde{B}) \leq \lambda_k(A, B)$ for all $k = 1, \dots, n$. Using for instance the third characterization combined with Theorem 4.3, inequalities (4.1a), we obtain

$$\lambda_k(A, \tilde{B}) \leq \lambda_k(A, B) \lambda_n(B, \tilde{B}) \leq \lambda_k(A, B) \quad 1 \leq k \leq n.$$

□

As we will show, these conditions are actually satisfied for an appropriate definition of mass lumping.

Definition 4.6 (Lumping operator). Let $A \in \mathbb{R}^{n \times n}$. The lumping operator $\mathcal{L}: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ is defined as

$$\mathcal{L}(A) = \text{diag}(d_1, \dots, d_n)$$

where $d_i = \sum_{j=1}^n |a_{ij}|$ for $i = 1, \dots, n$.

Remark 4.7. Definition 4.6 defines more specifically the row-sum lumping operator. Some readers might be concerned about the absolute values included in our definition. They are absolutely necessary for our results to hold as they guarantee positive definiteness of the lumped mass matrix provided A does not contain a row of zeros. For nonnegative matrices (i.e. matrices with nonnegative entries), our definition obviously coincides with the usual row-sum mass lumping. Some typical examples include isogeometric analysis [4], classical low order finite element methods and nonnegative partition of unity methods applicable to generalized or extended finite element methods [21]. Even for more classical methods, nonnegative mass matrices can still be constructed for special sets of basis functions and finite elements [22].

The next lemma is a first step in analyzing the preconditioned spectrum $(M, \mathcal{L}(M))$.

Lemma 4.8. Let $A \in \mathbb{R}^{n \times n}$ be symmetric positive definite. Then

1. $\Lambda(A, \mathcal{L}(A)) \subset (0, 1]$,
2. If A is nonnegative, then $\lambda_n(A, \mathcal{L}(A)) = 1$.

Proof. We prove each point below.

1. First note that, by construction, if A is symmetric positive definite then $\mathcal{L}(A)$ is also symmetric positive definite and consequently all eigenvalues of the matrix pencil $(A, \mathcal{L}(A))$ are strictly positive. Moreover, $\Lambda(A, \mathcal{L}(A)) = \Lambda(\mathcal{L}(A)^{-1}A)$. Denoting $\|B\|_\infty = \max_i \sum_{j=1}^n |b_{ij}|$ the infinity norm of a matrix B and applying the Gershgorin circle theorem, we obtain

$$\lambda_k(A, \mathcal{L}(A)) = \lambda_k(\mathcal{L}(A)^{-1}A) \leq \|\mathcal{L}(A)^{-1}A\|_\infty = \max_i \frac{1}{\mathcal{L}(A)_{ii}} \sum_{j=1}^n |a_{ij}| = \max_i \frac{\sum_{j=1}^n |a_{ij}|}{\sum_{j=1}^n |a_{ij}|} = 1 \quad k = 1, \dots, n,$$

showing that 1 is an upper bound on the spectrum.

2. Note that if A is nonnegative, then $(1, \mathbf{e})$ is an eigenpair of $(A, \mathcal{L}(A))$, where \mathbf{e} is the vector of all ones. Thus, the upper bound is always attained.

□

The next corollary is one of our main results.

Corollary 4.9. Let $A, B \in \mathbb{R}^{n \times n}$ be symmetric positive definite matrices and let all eigenvalues be numbered in ascending algebraic order. Then

$$\lambda_k(A, \mathcal{L}(B)) \leq \lambda_k(A, B) \quad 1 \leq k \leq n.$$

Proof. The result follows immediately from Lemma 4.8 and Corollary 4.5. □

Remark 4.10. Combining Theorem 4.3 and Lemma 4.8 also yields the upper bound

$$\lambda_k(A, B) \leq \lambda_k(\mathcal{L}(A), B) \quad 1 \leq k \leq n.$$

In words, lumping the mass matrix leads to an underestimate of the eigenvalues whereas lumping the stiffness matrix leads to an overestimate of the eigenvalues. Clearly, the above inequality is not of any practical interest for our problem and it is only mentioned for completeness.

In our quest to ensure an underestimate of the eigenvalues, we will heavily rely on the sufficient conditions stated in Corollary 4.5. In fact, the positive semi-definiteness of the error is also the main building block for ad hoc mass scaling strategies [23, 24, 25]. The objectives of mass lumping and mass scaling are similar, however, mass scaling primarily focuses on improving the stability. The issue of solving linear systems with the scaled mass matrix is usually disregarded and at best only addressed in a later design phase. On the contrary, mass lumping addresses both issues simultaneously.

Nevertheless, although the conditions stated in Corollary 4.5 are sufficient, they are not necessary, as the following example shows.

Example 4.11. Consider the symmetric positive definite matrices

$$A = \begin{pmatrix} 6 & 0 \\ 0 & 6 \end{pmatrix}, \quad B = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \quad \tilde{B} = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}.$$

The eigenvalues of (A, B) , (A, \tilde{B}) and of the error $E = \tilde{B} - B$ are

$$\begin{aligned} \lambda_1(A, B) &= 2, & \lambda_1(A, \tilde{B}) &= 2, & \lambda_1(E) &= \frac{1 - \sqrt{5}}{2} < 0, \\ \lambda_2(A, B) &= 6, & \lambda_2(A, \tilde{B}) &= 3, & \lambda_2(E) &= \frac{1 + \sqrt{5}}{2} > 0. \end{aligned}$$

Thus, the error E is indefinite and yet $\lambda_k(A, \tilde{B}) \leq \lambda_k(A, B)$ for $k = 1, 2$.

The theoretical results are now illustrated on a numerical example.

Example 4.12 (Preconditioner comparison). We illustrate the theoretical bounds with a first numerical example. We consider an isogeometric discretization of our model problem (2.1) on a stretched square already used by previous authors [26, 27]. Quadratic B-splines of maximal smoothness are used with 20 subdivisions in each parametric direction. All our experiments are done with GeoPDEs [28], an open source Matlab/Octave package for isogeometric analysis. The perturbed spectrum $\Lambda(K, \tilde{M})$ is compared for two different choices of \tilde{M} , namely the lumped mass matrix introduced in Definition 4.6 and a very good preconditioner developed by Loli et al. [27], which is provably optimal under mesh refinement. The interested reader is referred to the original article for its definition and construction. This experiment has two purposes: firstly, it illustrates the close connection between the quality of a preconditioner and the approximation of the eigenvalues. Secondly, it supports our first theoretical finding regarding mass lumping and pinpoints some characteristic features so far not captured by our bounds.

The eigenvalue ratios together with their lower and upper bounds from (4.2) are shown in Figure 4.1 for the lumped mass matrix and for the preconditioner of Loli et al. In Figure 4.2, the exact discrete spectrum is compared to the approximate spectrum obtained with our two choices of preconditioners. The excellent preconditioner of Loli et al. leads to a nearly perfect approximation of the entire spectrum while mass lumping leads to an underestimate of the eigenvalues, in agreement with Corollary 4.9.




Figure 4.1: Eigenvalue ratios $\frac{\lambda_k(K, \tilde{M})}{\lambda_k(K, M)}$ when using the lumped mass matrix and the preconditioner of Lori et al. [27]




Figure 4.2: Comparison between $\Lambda(K, M)$ and $\Lambda(K, \tilde{M})$ obtained using the lumped mass matrix and the preconditioner of Lori et al. [27]

Example 4.12 also illustrates a typical and widely appreciated feature of mass lumping: In Figure 4.1, the smallest eigenvalues are reasonably well approximated, whereas the larger eigenvalues are strongly underestimated. It explains why mass lumping does not lead to completely erroneous approximate solutions. This feature is explored more thoroughly in a second and simpler example.

Example 4.13 (*h*-refinement for 1D Laplace). We perform an *h*-refinement experiment for the isogeometric discretization of the 1D Laplace eigenvalue problem on the unit line with homogeneous Dirichlet boundary conditions. The relative eigenvalue error $\frac{\lambda_i - \tilde{\lambda}_i}{\lambda_i}$ for the first 20 discrete eigenvalues is shown in Figures 4.3 for increasingly fine meshes. The spline order is fixed at $p = 6$. As expected, for a given spline degree the relative error decreases for increasingly fine meshes but is not even across all eigenvalues: the smallest eigenvalues are always much better approximated.




Figure 4.3: Relative eigenvalue error for the first 20 discrete eigenvalues and $p = 6$

The reasons why mass lumping gives a much better approximation of the smallest eigenvalues can be intuitively explained with the perturbation theory of generalized eigenvalue problems pioneered by Stewart in the early 1970s. For generalized eigenvalue problems, the eigenvalue error is often measured in the chordal metric [18, 29, 30, 31, 32] because its ability to treat small and large eigenvalues uniformly. Error bounds are also alternatively formulated in

terms of *eigenangles*, which were first introduced by Stewart in 1979 [30] and are related to eigenvalues by

$$\theta_i = \operatorname{arccot}(\lambda_i) \quad 1 \leq i \leq n.$$

Similarly to eigenvalues, eigenangles satisfy a min-max characterization, which is the basis for deriving error bounds on eigenangles. The classical perturbation theory uniformly bounds the quantity $\sin(|\tilde{\theta}_i - \theta_i|)$ by the ratio of some error measure over the definiteness of the pencil. Due to the highly nonlinear nature of the cotangent function, large eigenvalues undergo large perturbations while small eigenvalues undergo small perturbations, which is consistent with the results in Figure 4.2. Quantitatively speaking though, the perturbations bounds from [30] seemed sharper for the smallest eigenvalues than for the largest ones in our numerical experiments. The classical perturbation theory may not be satisfactory for our problem for two main reasons. Firstly, the error E is treated as a random symmetric perturbation whereas in our problem the lumped mass matrix is constructed from the consistent mass matrix and therefore the error inherits some structure, which is ignored in the classical perturbation theory. Secondly, all error bounds we are aware of are *uniform*, meaning the error (in whichever measure is used) is bounded independently of the eigenvalue/eigenangle number similarly to Weyl's theorem for standard eigenvalue problems [20, Equation (6.3.4.1)]. Such bounds cannot explain differences in relative error as we are witnessing in Figure 4.3. Instead, focus should be directed at understanding how the error is interacting with the eigenspace. Surprisingly little attention has been paid to this problem and the few results we are aware of are limited to standard eigenvalue problems [33]. Such error bounds are not only valuable for mass lumping but also for selective mass scaling, where specific eigenvalues are selectively driven down [23]. Perturbation bounds that take into account eigenspaces are established in Theorem 4.14.

Theorem 4.14. Let (A, B) be a symmetric matrix pencil with B positive definite. Let $(\tilde{A}, \tilde{B}) = (A + E, B + F)$, where $E, F \in \mathbb{R}^{n \times n}$ are symmetric perturbations and \tilde{B} is positive definite. Let $\{(\lambda_i, \mathbf{u}_i)\}_{i=1}^n$ and $\{(\tilde{\lambda}_i, \tilde{\mathbf{u}}_i)\}_{i=1}^n$ denote the eigenpairs of (A, B) and (\tilde{A}, \tilde{B}) , respectively, where the eigenvectors \mathbf{u}_i and $\tilde{\mathbf{u}}_i$ are assumed B -orthonormal and \tilde{B} -orthonormal, respectively. Then

$$\min_j |\tilde{\lambda}_j - \lambda_i| \leq \frac{1}{\lambda_1(\tilde{B})} \left(\frac{\|E\mathbf{u}_i\|_2}{\|\mathbf{u}_i\|_2} + |\lambda_i| \frac{\|F\mathbf{u}_i\|_2}{\|\mathbf{u}_i\|_2} \right) \quad 1 \leq i \leq n, \quad (4.3a)$$

$$\min_j |\lambda_j - \tilde{\lambda}_i| \leq \frac{1}{\lambda_1(B)} \left(\frac{\|E\tilde{\mathbf{u}}_i\|_2}{\|\tilde{\mathbf{u}}_i\|_2} + |\tilde{\lambda}_i| \frac{\|F\tilde{\mathbf{u}}_i\|_2}{\|\tilde{\mathbf{u}}_i\|_2} \right) \quad 1 \leq i \leq n. \quad (4.3b)$$

Proof. We prove statement (4.3a) by first noticing that for any given vector $\mathbf{x} \in \mathbb{R}^n$ and any given scalar $\sigma \in \mathbb{R} \setminus \Lambda(\tilde{A}, \tilde{B})$, $\mathbf{x} = (\tilde{A} - \sigma\tilde{B})^{-1}(\tilde{A} - \sigma\tilde{B})\mathbf{x}$. If \tilde{B} is positive definite, there exists a matrix $\tilde{U} \in \mathbb{R}^{n \times n}$ of \tilde{B} -orthonormal eigenvectors that diagonalizes (\tilde{A}, \tilde{B}) such that $\tilde{U}^T \tilde{A} \tilde{U} = \tilde{D}$ and $\tilde{U}^T \tilde{B} \tilde{U} = I$, where \tilde{D} is the diagonal matrix of eigenvalues of (\tilde{A}, \tilde{B}) (Lemma 3.2). Then, by taking the norm, we obtain

$$\|\mathbf{x}\|_2 \leq \|(\tilde{A} - \sigma\tilde{B})^{-1}\|_2 \|(\tilde{A} - \sigma\tilde{B})\mathbf{x}\|_2 = \|\tilde{U}(\tilde{D} - \sigma I)^{-1}\tilde{U}^T\|_2 \|(\tilde{A} - \sigma\tilde{B})\mathbf{x}\|_2 \leq \|\tilde{U}\|_2^2 \frac{\|(\tilde{A} - \sigma\tilde{B})\mathbf{x}\|_2}{\min_j |\tilde{\lambda}_j - \sigma|}.$$

Noticing that $\|\tilde{U}\|_2^2 = \|\tilde{B}^{-1}\|_2 = \frac{1}{\lambda_1(\tilde{B})}$, we obtain

$$\min_j |\tilde{\lambda}_j - \sigma| \leq \frac{1}{\lambda_1(\tilde{B})} \frac{\|(\tilde{A} - \sigma\tilde{B})\mathbf{x}\|_2}{\|\mathbf{x}\|_2}.$$

In particular, the inequality trivially holds for any $\sigma \in \Lambda(\tilde{A}, \tilde{B})$. After choosing $\sigma = \lambda_i$ and $\mathbf{x} = \mathbf{u}_i$, we obtain a residual bound. Since $A\mathbf{u}_i = \lambda_i B\mathbf{u}_i$, we finally obtain

$$\min_j |\tilde{\lambda}_j - \lambda_i| \leq \frac{1}{\lambda_1(\tilde{B})} \frac{\|(\tilde{A} - \lambda_i \tilde{B})\mathbf{u}_i\|_2}{\|\mathbf{u}_i\|_2} = \frac{1}{\lambda_1(\tilde{B})} \frac{\|(\tilde{A} - A + \lambda_i(B - \tilde{B}))\mathbf{u}_i\|_2}{\|\mathbf{u}_i\|_2} \leq \frac{1}{\lambda_1(\tilde{B})} \left(\frac{\|E\mathbf{u}_i\|_2}{\|\mathbf{u}_i\|_2} + |\lambda_i| \frac{\|F\mathbf{u}_i\|_2}{\|\mathbf{u}_i\|_2} \right).$$

Statement (4.3b) follows by swapping the roles of (A, B) and (\tilde{A}, \tilde{B}) . \square

Remark 4.15. Theorem 4.14, statement (4.3a) is a generalization of a uniform bound on eigenvalues [32, Theorem 2], which appeared earlier in [34].

$$|\tilde{\lambda}_i - \lambda_i| \leq \frac{1}{\lambda_1(\tilde{B})} (\|E\|_2 + |\lambda_i| \|F\|_2) \quad 1 \leq i \leq n. \quad (4.4)$$

A direct comparison between (4.3a) and (4.4) is generally not possible since they do not bound the same quantities. Nevertheless, if the eigenvalue gap is large enough relative to the error, then the perturbed eigenvalue that attains the minimum in (4.3a) is actually the i th one [33]. In this case, Theorem 4.14 is a clear improvement. In the context of mass lumping, the error component in the eigenspaces associated to the smallest eigenvalues tends to be quite small. Moreover, for sufficiently fine meshes, the discrete eigenvalues are sufficiently close to the exact eigenvalues, which are well separated for 1D problems. Thus, Theorem 4.14 might very well directly provide an upper bound on the relative error for the smallest eigenvalues. This assertion is tested in the next example.

Example 4.16. The quality of the upper bounds given in (4.3a) and (4.4) is assessed numerically for a first order isogeometric discretization of the 1D Laplace. Figure 4.4 indicates that for the first few eigenvalues, $\tilde{\lambda}_i$ is the perturbed eigenvalue closest to λ_i . However, it is no longer the case for the upper 82 % of the spectrum. The error grows faster than the relative gap and the pairing is lost. From Figure 4.4, the reader might misleadingly think that the upper bound in (4.3a) is actually *always* an upper bound on the relative error between λ_i and $\tilde{\lambda}_i$, regardless of the eigenvalue gap. Unfortunately, other experiments showed the contrary. The results using the upper bound in (4.3b) were quite similar and are omitted.




Figure 4.4: Comparison of (4.3a) and (4.4) for the row-sum lumped mass matrix and $p = 1$

4.2 A generalization of mass lumping

For applications in explicit dynamics, the different approximation quality at different branches of the spectrum is an attractive feature of mass lumping. A good approximation of the smallest eigenvalues allows us to recover reasonably good solutions whereas a rather poor and gross underestimate of the largest eigenvalues, which usually do not contribute much to the solution, improves the stability of the time integration scheme. We will now design an entire class of preconditioners preserving those nice properties and generalizing the simple row-sum lumped mass matrix to more complex structures. A small preparatory lemma is needed before describing this class.

Lemma 4.17. Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Then

$$-\mathcal{L}(A) \preceq A \preceq \mathcal{L}(A).$$

Moreover,

1. If A does not contain a row of zeros, then $\Lambda(A, \mathcal{L}(A)) \subset [-1, 1]$,
2. If A is nonnegative, then $\dim \ker(A - \mathcal{L}(A)) \geq 1$,

Proof. Since we do not assume A to be positive definite, we run the risk of encountering singular matrix pencils (see Definition 3.1). In general, $\ker(\mathcal{L}(A)) \subseteq \ker(A)$ and $\dim \ker(\mathcal{L}(A)) > 0$ if and only if A contains a row of zeros. Thus, if A contains a row of zeros, then the null spaces of A and $\mathcal{L}(A)$ have a nontrivial intersection and the matrix pencil $(A, \mathcal{L}(A))$ is singular: generalized eigenvalues associated to an eigenvector $\mathbf{y} \in \ker(A) \cap \ker(\mathcal{L}(A))$ ($\mathbf{y} \neq \mathbf{0}$) are arbitrary. We first study the regular case by assuming A does not contain a row of zeros and later extend the argument to the more general case. If A does not contain a row of zeros, then the lumping operator still delivers a

symmetric positive definite matrix $\mathcal{L}(A)$. Thus, $\Lambda(A, \mathcal{L}(A)) \subset \mathbb{R}$. Now, by the Gershgorin circle theorem (similarly to Lemma 4.8)

$$|\lambda_k(A, \mathcal{L}(A))| = |\lambda_k(\mathcal{L}(A)^{-1}A)| \leq \|\mathcal{L}(A)^{-1}A\|_\infty = 1$$

and consequently $\Lambda(A, \mathcal{L}(A)) \subset [-1, 1]$. Moreover, since A is symmetric, the Courant-Fischer theorem still holds. Thus, $\forall \mathbf{u} \in \mathbb{R}^n, \mathbf{u} \neq \mathbf{0}$

$$-1 \leq \lambda_1(A, \mathcal{L}(A)) = \min_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T \mathcal{L}(A) \mathbf{x}} \leq \frac{\mathbf{u}^T A \mathbf{u}}{\mathbf{u}^T \mathcal{L}(A) \mathbf{u}} \leq \max_{\substack{\mathbf{x} \in \mathbb{R}^n \\ \mathbf{x} \neq \mathbf{0}}} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T \mathcal{L}(A) \mathbf{x}} = \lambda_n(A, \mathcal{L}(A)) \leq 1.$$

Consequently, $-\mathcal{L}(A) \preceq A \preceq \mathcal{L}(A)$. In particular, if A is nonnegative, $\dim \ker(A - \mathcal{L}(A)) \geq 1$ since $(A - \mathcal{L}(A))\mathbf{e} = \mathbf{0}$ where \mathbf{e} is the vector of all ones.

We now extend the proof to the more general case. Let us assume that A contains m rows of zeros. Since A is symmetric, if the i th row is zero, so is the i th column. Without loss of generality, we may assume that the rows and columns of A are ordered such that

$$A = \begin{pmatrix} B & 0 \\ 0 & 0 \end{pmatrix} \quad \mathcal{L}(A) = \begin{pmatrix} \mathcal{L}(B) & 0 \\ 0 & 0 \end{pmatrix}$$

where B is symmetric and does not contain any row of zeros. The previous proof argument can be used to show that $\Lambda(B, \mathcal{L}(B)) \subset [-1, 1]$ and $-\mathcal{L}(B) \preceq B \preceq \mathcal{L}(B)$. Consequently, $-\mathcal{L}(A) \preceq A \preceq \mathcal{L}(A)$ still holds and the zero rows only increase the dimension of $\ker(A - \mathcal{L}(A))$. If A is the zero matrix, the result trivially holds. \square

In some sense, if A is nonnegative, Lemma 4.17 states that the ‘‘upper bound’’ $A \preceq \mathcal{L}(A)$ is ‘‘attained’’ in the Loewner ordering. Namely, the error $\mathcal{L}(A) - A$ is not only positive semidefinite but also singular. For the sake of eigenvalue approximation, we are especially interested in singular errors. In fact, positive definite errors will necessarily effect the smallest generalized eigenvalues, which we of course want to avoid.

Definition 4.18 (Preconditioners P_i). Let $A \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix and consider the matrix splitting $A = D_i + R_i$ where D_i consists of all super and sub-diagonals smaller than i and R_i is the remainder. We define the sequence of preconditioners $P_i = D_i + \mathcal{L}(R_i)$ for $i = 1, \dots, n$. In particular, we observe that $P_1 = \mathcal{L}(A)$ and $P_n = A$.

We are now ready to state an important theorem generalizing Lemma 4.8.

Theorem 4.19. Let $A \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix. Then, the sequence of preconditioners $\{P_i\}_{i=1}^n$ constructed from A according to Definition 4.18 satisfies the following properties:

1. $\Lambda(A, P_i) \subset (0, 1]$ for all $i = 1, \dots, n$,
2. $\lambda_k(A, P_i) \leq \lambda_k(A, P_{i+1})$ for all $k = 1, \dots, n$ and any given $i = 1, \dots, n - 1$,
3. If A is nonnegative, then $\lambda_n(A, P_i) = 1$ for all $i = 1, \dots, n$.

Proof. We prove all properties below.

1. We first prove that P_i is symmetric positive definite for all $i = 1, \dots, n$. Symmetry is a straightforward consequence of its definition and the symmetry of A . Positive definiteness immediately follows from Lemma 4.17 and the positive definiteness of A :

$$A - P_i = R_i - \mathcal{L}(R_i) \preceq 0 \implies P_i \succeq A \succ 0.$$

From the equivalent characterizations in Corollary 4.5, $\Lambda(A, P_i) \subset (0, 1]$.

2. For the second property, we note that

$$P_{i+1} - P_i = D_{i+1} + \mathcal{L}(R_{i+1}) - (D_i + \mathcal{L}(R_i)) = D_{i+1} - D_i - \mathcal{L}(D_{i+1} - D_i) = \Delta D_i - \mathcal{L}(\Delta D_i) \preceq 0$$

where we have denoted $\Delta D_i = D_{i+1} - D_i$ and used the definition of the matrices R_i and R_{i+1} to deduce that

$$\mathcal{L}(R_{i+1}) - \mathcal{L}(R_i) = -\mathcal{L}(R_i - R_{i+1}) = -\mathcal{L}(D_{i+1} - D_i).$$

We note that ΔD_i is obviously symmetric, but it may be singular. Nevertheless, Lemma 4.17 covers for such a situation and states that $\Delta D_i - \mathcal{L}(\Delta D_i) \preceq 0$. Since P_i and P_{i+1} are symmetric positive definite and $P_i \succeq P_{i+1}$, Corollary 4.5 concludes that

$$\lambda_k(A, P_i) \leq \lambda_k(A, P_{i+1}) \quad 1 \leq k \leq n.$$

The same argument shows that more generally, $P_i \succeq P_j$ for all $j \geq i$. Thus, $\lambda_n(P_j, P_i) \leq 1$ for all $j \geq i$.

3. By the first property, $\lambda_n(A, P_i) \leq 1$. If A is nonnegative, this upper bound is attained. Indeed, if A is nonnegative, all preconditioners P_i for $i = 1, \dots, n$ are also nonnegative and all matrices have the same row-sum by construction. Thus, $A\mathbf{e} = (D_i + R_i)\mathbf{e} = (D_i + \mathcal{L}(R_i))\mathbf{e} = P_i\mathbf{e}$ where \mathbf{e} is the vector of all ones. Therefore, $(1, \mathbf{e})$ is an eigenpair of (A, P_i) and (P_j, P_i) for all $i, j = 1, \dots, n$. Consequently, $\lambda_n(A, P_i) = 1$ and $\lambda_n(P_j, P_i) = 1$ for all $j \geq i$.

□

Lemma 4.17 states that $\mathcal{L}(A) \succeq A$. Increasing the bandwidth of the preconditioner allows us to find increasingly better preconditioners that are “closer” to A in the Loewner ordering:

$$\mathcal{L}(A) = P_1 \succeq P_2 \succeq \dots \succeq P_{n-1} \succeq P_n = A.$$

Theorem 4.19 then simply states that $\Lambda(A, P_i)$ converges monotonically to 1 for increasing values of i . In particular, the theorem can be used to construct a sequence of preconditioners for the mass matrix and the first element of this sequence coincides with the usual row-sum lumped mass matrix. Since $\Lambda(M, P_i)$ converges monotonically to 1 for increasing values of i , $\Lambda(K, P_i)$ is expected to yield an increasingly good approximation of $\Lambda(K, M)$. Indeed, since $P_i \succeq P_{i+1}$, Corollary 4.5 yields $\lambda_k(K, P_i) \leq \lambda_k(K, P_{i+1})$. By repeating the same argument for all $i = 1, \dots, n - 1$, we obtain

$$\lambda_k(K, \mathcal{L}(M)) = \lambda_k(K, P_1) \leq \lambda_k(K, P_2) \leq \dots \leq \lambda_k(K, P_{n-1}) \leq \lambda_k(K, P_n) = \lambda_k(K, M) \quad 1 \leq k \leq n.$$

In other words, $\Lambda(K, P_i)$ converges monotonically from below to $\Lambda(K, M)$ for increasing values of i . In practice, we are only interested in small values of i such that the preconditioners P_i have narrow bandwidths. Linear systems of size n with banded symmetric positive definite matrices of bandwidth $b \ll n$ can be solved in $O(nb^2)$ floating point operations (flops), which represents a considerable saving over the $O(n^3)$ complexity for general matrices using standard Gaussian elimination [35, 36]. The tridiagonal case is particularly appealing since linear systems can be solved in $O(n)$ flops using the Thomas algorithm. In general, choosing the right preconditioner is a compromise between accuracy and computational cost.

Example 4.20 (Lumped mass preconditioners P_i - 1D). We consider an isogeometric discretization of the 1D Laplace equation on $\Omega = (0, 1)$ with B-splines of order $p = 3$ and $p = 5$ with C^{p-1} continuity and 400 subdivisions. For a p th order discretization with C^{p-1} continuity, the mass and stiffness matrices are banded matrices of bandwidth p . Therefore, they have $2p + 1$ bands. We construct the first three lumped mass preconditioners of the sequence $(P_1, P_2$ and $P_3)$. The preconditioner P_i is a banded matrix of bandwidth $i - 1$ and has $2i - 1$ bands. Thus, P_1 , P_2 and P_3 are diagonal, tridiagonal and pentadiagonal matrices, respectively. None of these preconditioners are able to represent the mass matrices exactly. Already for $p = 3$, the mass matrix has 7 bands. However, the further the bands are from the diagonal, the smaller the magnitude of their entries. The spectrum of (K, M) and (K, P_i) for $i = 1, 2, 3$ are compared in Figures 4.5 and 4.6 for $p = 3$ and $p = 5$, respectively. The approximation quality improves tremendously already for small values of i . For $p = 3$, the preconditioner P_3 leads to a nearly perfect approximation. We notice that $\Lambda(K, P_i)$ converges monotonically from below to $\Lambda(K, M)$, in perfect agreement with our theoretical results. Alert readers will also notice the “spikes” appearing towards the end of the spectrum. The largest eigenvalues of (K, M) grossly overestimate those of the continuous operator and for this reason were coined *outliers* in [6] where they were first discussed. Unfortunately, in our case, $\Lambda(K, P_i)$ approximates the entire spectrum of (K, M) , including the outliers. Fixing this issue will be the object of future work.




Figure 4.5: Comparison of $\Lambda(K, M)$ and $\Lambda(K, P_i)$ for $i = 1, 2, 3$ and $p = 3$




Figure 4.6: Comparison of $\Lambda(K, M)$ and $\Lambda(K, P_i)$ for $i = 1, 2, 3$ and $p = 5$

The performance of these preconditioners in approximating the spectrum of (K, M) is directly related to the clustering on the eigenvalues of (M, P_i) around 1. The eigenvalues of the matrix pencils (M, P_i) for $i = 1, 2, 3$ are shown in Figures 4.7 and 4.8 for $p = 3$ and $p = 5$, respectively. As predicted theoretically, the spectrum of the preconditioned matrices converges monotonically to 1.




Figure 4.7: $\Lambda(M, P_i)$ for $i = 1, 2, 3$ and $p = 3$




Figure 4.8: $\Lambda(M, P_i)$ for $i = 1, 2, 3$ and $p = 5$

Remark 4.21. Theorem 4.19 can also be used to construct a sequence of preconditioners for the stiffness matrix K . It can be easily shown that in this case the approximate spectrum converges monotonically from above to $\Lambda(K, M)$ for increasing values of i . However, the main computational bottleneck in explicit dynamics consists in solving linear systems with the mass matrix, not the stiffness matrix. Therefore, we will not explore this option any further.

Theorem 4.19 is based on purely algebraic considerations on the entries of the matrices. Thus, it holds regardless of how complicated the physical domain is or how large the dimension of the problem becomes. However, for multidimensional problems the mass matrix no longer has a banded structure. Therefore, if banded preconditioners are constructed, the preconditioned spectrum will most likely converge very slowly to 1. This foreboding is confirmed in the next two-dimensional example.

Example 4.22 (Lumped mass preconditioners P_i - 2D). We consider the 2D Laplace model problem on the unit square $\Omega = (0, 1)^2$ discretized with maximally smooth B-splines of order $p = 3$ and $p = 5$ and 20 subdivisions in each parametric direction. As shown in Figures 4.9 and 4.10, increasing the bandwidth of the preconditioners in this case barely improves the approximation of the eigenvalues. This poor approximation is closely linked to the poor clustering of the eigenvalues of (M, P_i) around 1. As a matter of fact, Figures 4.11 and 4.12 reveal that only very few eigenvalues of the preconditioned matrix are close to 1 whereas the vast majority are clustered around zero.




Figure 4.9: Comparison of $\Lambda(K, M)$ and $\Lambda(K, P_i)$ for $i = 1, 2, 3$ and $p = 3$




Figure 4.10: Comparison of $\Lambda(K, M)$ and $\Lambda(K, P_i)$ for $i = 1, 2, 3$ and $p = 5$




Figure 4.11: $\Lambda(M, P_i)$ for $i = 1, 2, 3$ and $p = 3$




Figure 4.12: $\Lambda(M, P_i)$ for $i = 1, 2, 3$ and $p = 5$

It is well-known in preconditioning practice that good preconditioners must preserve the essential structure of the matrices [37, 38]. Therefore, it would be wise to further extend the family of lumped mass preconditioners to better account for the structure. For tensor product basis functions, there exists a natural generalization to multidimensional problems.

4.3 Multidimensional problems

We first discuss the case of a scalar PDE in 2D. The extension to 3D is straightforward and vector problems will be discussed at the end of this section. If the physical domain is a square, the density function is separable and tensor product basis functions are used, the mass matrix is expressed as $M = M_1 \otimes M_2$ where M_1 and M_2 are mass

matrices for univariate problems (see for instance [26]). We then consider the family of lumped mass preconditioners given by $P_{ij} = P_{1,i} \otimes P_{2,j}$ where $P_{1,i}$ and $P_{2,j}$ are banded preconditioners defined in Definition 4.18 for some indices $1 \leq i, j \leq n$. This definition is the natural extension to two-dimensional problems, as shown in the next theorem.

Definition 4.23 (Preconditioners P_{ij}). Let $A = A_1 \otimes A_2$ where $A_1, A_2 \in \mathbb{R}^{n \times n}$ are symmetric positive definite matrices. We define the family of preconditioners $P_{ij} = P_{1,i} \otimes P_{2,j}$, where $P_{1,i}$ and $P_{2,j}$ are preconditioners constructed from A_1 and A_2 , respectively, for some indices $1 \leq i, j \leq n$ according to Definition 4.18. In particular, we observe that $P_{nn} = A$.

Theorem 4.24. Let $A = A_1 \otimes A_2$ where $A_1, A_2 \in \mathbb{R}^{n \times n}$ are symmetric positive definite matrices. Then, the family of preconditioners $\{P_{ij}\}_{i,j=1}^n$ constructed from A according to Definition 4.23 satisfies the following properties:

1. $\Lambda(A, P_{ij}) \subset (0, 1]$ for all $i, j = 1, \dots, n$,
2. $\lambda_k(A, P_{sq}) \leq \lambda_k(A, P_{ij})$ for all $k = 1, \dots, n^2$ if $i \geq s$ and $j \geq q$,
3. If A_1 and A_2 are nonnegative, then $\lambda_n(A, P_{ij}) = 1$ for all $i, j = 1, \dots, n$.

Proof. We prove all properties below using the results of Theorem 4.19. Let us first define the product of two sets $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathbb{C}$ as $\mathcal{S}_1 \mathcal{S}_2 = \{s_1 s_2 : s_1 \in \mathcal{S}_1, s_2 \in \mathcal{S}_2\}$.

1. We will first prove that A and P_{ij} are symmetric positive definite. Our arguments are based on some basic properties of the Kronecker product. The reader may for instance refer to [39, Chapter 4] for their proof. For matrices $B, C \in \mathbb{R}^{n \times n}$
 - $\Lambda(B \otimes C) = \Lambda(B)\Lambda(C)$,
 - $(B \otimes C)^T = B^T \otimes C^T$,
 - If B and C are invertible, $(B \otimes C)^{-1} = B^{-1} \otimes C^{-1}$.

By the first and second property, since A_1, A_2 and $P_{1,i}, P_{2,j}$ are symmetric positive definite for all $i, j = 1, \dots, n$, then so are $A_1 \otimes A_2$ and $P_{1,i} \otimes P_{2,j}$. Therefore, $\lambda_k(A, P_{ij}) > 0$ for all $k = 1, \dots, n$. Then, using Lemma 4.1 and the properties of the Kronecker product

$$\Lambda(A, P_{ij}) = \Lambda(A_1 \otimes A_2, P_{1,i} \otimes P_{2,j}) = \Lambda(P_{1,i}^{-1} A_1 \otimes P_{2,j}^{-1} A_2) = \Lambda(A_1, P_{1,i}) \Lambda(A_2, P_{2,j}).$$

Thus, using the first property of Theorem 4.19, $\lambda_n(A, P_{ij}) = \lambda_n(A_1, P_{1,i}) \lambda_n(A_2, P_{2,j}) \leq 1$ and $\Lambda(A, P_{ij}) \subset (0, 1]$.

2. Using Theorem 4.3, inequality (4.1a), we obtain

$$\lambda_k(A, P_{sq}) \leq \lambda_k(A, P_{ij}) \lambda_n(P_{ij}, P_{sq}) \quad 1 \leq k \leq n^2.$$

Using again Lemma 4.1 and the properties of the Kronecker product, $\lambda_n(P_{ij}, P_{sq}) = \lambda_n(P_i, P_s) \lambda_n(P_j, P_q)$. Finally, when proving the third property of Theorem 4.19, we showed that $\lambda_n(P_j, P_i) \leq 1$ for all $j \geq i$. Therefore, $\lambda_n(P_{ij}, P_{sq}) \leq 1$ if $i \geq s$ and $j \geq q$ and the result follows.

3. Moreover, if A_1 and A_2 are nonnegative, so are the preconditioners $P_{1,i}$ and $P_{2,j}$ for all $i, j = 1, \dots, n$ constructed from A_1 and A_2 and Theorem 4.19 states that $\lambda_n(A_1, P_{1,i}) = 1$ and $\lambda_n(A_2, P_{2,j}) = 1$ for all $i, j = 1, \dots, n$. Therefore, $\lambda_n(A, P_{ij}) = \lambda_n(A_1, P_{1,i}) \lambda_n(A_2, P_{2,j}) = 1$.

□

Remark 4.25. The previous theorem is extended straightforwardly to matrices $A = A_1 \otimes A_2 \otimes A_3$ and the associated family of preconditioners $P_{ijk} = P_{1,i} \otimes P_{2,j} \otimes P_{3,k}$ where $P_{1,i}, P_{2,j}$ and $P_{3,k}$ are constructed from A_1, A_2 and A_3 , respectively, for some indices $1 \leq i, j, k \leq n$ according to Definition 4.18.

Linear systems with a Kronecker product structure can be solved very efficiently and especially if the factor matrices themselves have good properties. It is the main motivation for designing Kronecker product preconditioners [26, 27, 40]. In the 2D case, if all factor matrices have size n and c_1 and c_2 denote the cost of solving a single linear system with $P_{1,i}$ and $P_{2,j}$, respectively, then the cost of solving a linear system with P_{ij} is $O(n(c_1 + c_2))$. Thus, the preconditioner of largest bandwidth dictates the overall cost. If $c_1 \approx c_2 \approx n$, this cost amounts to $O(n^2)$, which is linear in the size of the system. The analysis can be easily generalized to the 3D case.

Remark 4.26. Definition 4.23 can be generalized to matrices $A = S^T(A_1 \otimes A_2)S$ where $A_1, A_2 \in \mathbb{R}^{n \times n}$ are symmetric positive matrices and S is an invertible matrix. The preconditioners are then naturally defined as $P_{ij} = S^T(P_{1,i} \otimes P_{2,j})S$. However, such preconditioners are only useful if linear systems with S can be solved easily.

Theorem 4.24 significantly increases the possibilities for preconditioning the mass matrix. It is worthwhile noting that different preconditioners may have the same sparsity pattern.

Example 4.27 (Lumped mass preconditioners - 2D). We consider exactly the same problem as in Example 4.22: the 2D Laplace model problem on the unit square $\Omega = (0,1)^2$ discretized with maximally smooth B-splines of order $p = 3$ and $p = 5$ and 20 subdivisions in each parametric direction. For this simple model problem, the mass matrix is expressed as a Kronecker product $M = M_1 \otimes M_2$, with $M_1 = M_2$ because of the same discretization parameters in each direction. We now compare the preconditioners P_i for $i = 1, 2, 3$ from the previous example to the preconditioners P_{ii} constructed following Definition 4.23. For $i = 1$, the two preconditioning strategies coincide. Hence, Figures 4.13 and 4.14 show the comparison for $i = 2$ and $i = 3$, respectively and $p = 3$. The same comparison for $p = 5$ is shown in Figures 4.15 and 4.16. In both cases, the preconditioners P_{ii} lead to a significant improvement. However, it is not a fair comparison since they also usually have many more nonzero entries. Thus, let us select two different preconditioners with exactly the same sparsity pattern and the same number of nonzeros. For instance, P_2 and P_{12} are both tridiagonal. Yet, Figure 4.17 reveals that the preconditioner P_{12} performs much better.




Figure 4.13: Comparison of $\Lambda(K, P_2)$ and $\Lambda(K, P_{22})$ for $p = 3$




Figure 4.14: Comparison of $\Lambda(K, P_3)$ and $\Lambda(K, P_{33})$ for $p = 3$




Figure 4.15: Comparison of $\Lambda(K, P_2)$ and $\Lambda(K, P_{22})$ for $p = 5$




Figure 4.16: Comparison of $\Lambda(K, P_3)$ and $\Lambda(K, P_{33})$ for $p = 5$




Figure 4.17: Comparison of $\Lambda(K, M)$, $\Lambda(K, P_2)$ and $\Lambda(K, P_{12})$ for $p = 3$

Example 4.28 (Accuracy test - 1D and 2D). As we have discussed in Section 1, the row-sum lumped mass matrix limits the convergence rate of the smallest eigenvalues to quadratic order (at least for 1D problems and for quadratic and cubic B-spline discretizations [6]). This property is now investigated numerically for our new class of lumped mass preconditioners for both 1D and 2D model problems. We consider a cubic isogeometric discretization of the 1D and 2D Laplace on the unit line and unit square, respectively, with mixed boundary conditions. The smallest eigenfrequency $\omega_1 = \sqrt{\lambda_1}$ for these model problems is given by

$$\omega_1 = \frac{\pi}{2} \text{ (1D)}, \quad \omega_1 = \frac{\pi}{\sqrt{2}} \text{ (2D)}.$$

We report the relative error $\frac{\omega_1 - \omega_{h,1}}{\omega_1}$ in Figures 4.18 and 4.19 for our 1D and 2D problems, respectively. Unfortunately, our class of lumped mass preconditioners does not seem to improve the convergence rate but may improve the constant by several orders of magnitude. This fact is interesting given how close P_3 (or P_{33}) is to the consistent mass matrix M . Small differences are enough to ruin the convergence rate, which we expect is due to the consistency error of the method. The same trend was also observed for other types of boundary conditions.




Figure 4.18: Relative error in the smallest eigenfrequency for a cubic ($p = 3$) discretization of the 1D Laplace on the unit line with mixed boundary conditions




Figure 4.19: Relative error in the smallest eigenfrequency for a cubic ($p = 3$) discretization of the 2D Laplace on the unit square with mixed boundary conditions

Remark 4.29 (Vector PDEs). Vector PDEs, such as those modeling the elastic deformation of structures, again require a slight adjustment to the definition of the preconditioner due to the specific structure of the mass matrix.

If M_S denotes the mass matrix of a multidimensional scalar PDE, and if the same approximation spaces are used for each component of the approximate solution of the vector PDE, then the mass matrix of a vector problem M_V is either $I_d \otimes M_S$ or $M_S \otimes I_d$ depending on the ordering of the basis functions [41]. With the first definition M_V is a block diagonal matrix. This ordering is used in GeoPDEs for instance [28]. In order to preserve the structure, the preconditioner is naturally defined as $P_V = I_d \otimes P_S$ or $P_V = P_S \otimes I_d$. Again, due to the Kronecker product structure, all the properties of P_S carry over to P_V . Numerical testing for elasticity problems produced figures similar to those reported in Example 4.27 and are omitted.

5 Nearest Kronecker product approximation

As we have seen in Section 4.3, the performance of the lumped mass preconditioners can be drastically improved for tensor product basis functions if the mass matrix is expressed as a Kronecker product. Unfortunately, this structure cannot be guaranteed for most problems of practical interest. In several cases, the entries of the mass matrix are computed from

$$M_{ij} = \int_{\hat{\Omega}} \hat{\rho}(\hat{\mathbf{x}}) \hat{B}_i(\hat{\mathbf{x}}) \hat{B}_j(\hat{\mathbf{x}}) |\det(J_F(\hat{\mathbf{x}}))| d\hat{\Omega} \quad i, j = 1, \dots, n. \quad (5.1)$$

The integration is done in the parametric domain $\hat{\Omega}$ where $\hat{B}_i(\hat{\mathbf{x}})$ are tensor product basis functions defined in the parametric domain (for instance B-splines in isogeometric analysis), $\hat{\rho}(\hat{\mathbf{x}}) = \rho(F(\hat{\mathbf{x}}))$ is the density function, $F: \hat{\Omega} \rightarrow \Omega$ is the mapping from parametric to physical domain and $J_F(\hat{\mathbf{x}})$ is its Jacobian matrix. The mass matrix loses its Kronecker product structure when the integrand is no longer a separable function. A nontrivial density function, geometry or the use of NURBS basis functions are some of the causes. Nevertheless, for tensor product basis functions, the mass (and stiffness) matrix can be generally very well approximated by a sum of Kronecker products, a fact which has already been exploited for designing fast assemblage algorithms [26, 42, 43, 44, 45]. For 2D problems, we are interested in finding the best factor matrices B and C such that $M \approx \tilde{M} = B \otimes C$. Several different approaches have been proposed in the literature. The one proposed in [42] and later simplified in [43], is tailored to PDE problems and attempts to directly approximate all non-separable functions in the integral (5.1) by a sum of separable functions. In other words, it finds a low-rank approximation for the integrand itself. Another approach, used in [44] and based on earlier work from [46, 47], is more general and does not require any knowledge of the origin of the matrix. Of course, the first approach seems better suited to our problem. However, its implementation, although simplified, remains technical. The implementation of the second approach, in comparison, is rather straightforward. The method relies entirely on standard matrix operations. For this reason, we will assess if the easiest method can deliver useful results. We will initially restrict the discussion to 2D problems by closely following the presentation in [46]. The extension to 3D problems will be discussed subsequently.

5.1 Approximation for 2D problems

Given a matrix $M \in \mathbb{R}^{n \times n}$ with $n = n_1 n_2$, our goal is to find the best factor matrices $B \in \mathbb{R}^{n_1 \times n_1}$ and $C \in \mathbb{R}^{n_2 \times n_2}$ such that $\phi_M(B, C) = \|M - B \otimes C\|_F$ is minimized. This problem was first investigated by Van Loan and Pitsianis [46]. They observed that since both the Kronecker product $B \otimes C$ and $\text{vec}(B) \text{vec}(C)^T$ form all the products $b_{ij} c_{kl}$ for $i, j = 1, \dots, n_1$ and $k, l = 1, \dots, n_2$ but at different locations, there must clearly exist a linear mapping $\mathcal{R}: \mathbb{R}^{n_1 n_2 \times n_1 n_2} \rightarrow \mathbb{R}^{n_1^2 \times n_2^2}$ such that $\mathcal{R}(B \otimes C) = \text{vec}(B) \text{vec}(C)^T$. This mapping can be defined explicitly. Consider the following block matrix where $A_{ij} \in \mathbb{R}^{n_2 \times n_2}$ for $i, j = 1, \dots, n_1$ and define the mapping as

$$A = \begin{pmatrix} A_{1,1} & \cdots & A_{1,n_1} \\ \vdots & \ddots & \vdots \\ A_{n_1,1} & \cdots & A_{n_1,n_1} \end{pmatrix} \quad \mathcal{R}(A) = \begin{pmatrix} \text{vec}(A_{1,1})^T \\ \text{vec}(A_{2,1})^T \\ \vdots \\ \text{vec}(A_{n_1,n_1})^T \end{pmatrix}.$$

Then, by construction, for a matrix $A = B \otimes C$,

$$B \otimes C = \begin{pmatrix} b_{1,1} C & \cdots & b_{1,n_1} C \\ \vdots & \ddots & \vdots \\ b_{n_1,1} C & \cdots & b_{n_1,n_1} C \end{pmatrix} \quad \mathcal{R}(B \otimes C) = \begin{pmatrix} b_{1,1} \text{vec}(C)^T \\ b_{2,1} \text{vec}(C)^T \\ \vdots \\ b_{n_1,n_1} \text{vec}(C)^T \end{pmatrix} = \text{vec}(B) \text{vec}(C)^T.$$

More generally, since the mapping \mathcal{R} inherits the linearity of the vectorization operator, it transforms a sum of r Kronecker products into a rank r matrix. A matrix is said to have Kronecker rank r if it can be expressed as the sum of r Kronecker products. Thus, A has Kronecker rank r if and only if $\mathcal{R}(A)$ has (matrix) rank r . The mapping

\mathcal{R} was originally called *rearrangement* in [46] but later authors have also called it *reordering*. In any case, we must stress that $\mathcal{R}(A)$ is generally not a permutation of the rows and columns of A or any change of basis.

The minimization of $\|M - B \otimes C\|_F$ in the Frobenius norm allows to rewrite the problem in a more familiar way. The Frobenius norm is defined as the square root of the sum of squares of the matrix entries and does not depend on the location of these entries. Therefore, the mapping \mathcal{R} can be applied to a matrix without affecting its Frobenius norm and the minimization problem becomes

$$\min \phi_M(B, C) = \min \|M - B \otimes C\|_F = \min \|\mathcal{R}(M) - \mathcal{R}(B \otimes C)\|_F = \min \|\mathcal{R}(M) - \text{vec}(B) \text{vec}(C)^T\|_F.$$

Thus, finding the best factor matrices B and C is equivalent to finding the best rank-1 approximation of $\mathcal{R}(M)$. The transformation from M to $\mathcal{R}(M)$ requires knowing the block-partitioning of the matrix M , which in the context of isogeometric analysis is always clear from the underlying discretization. The next theorem summarizes a very useful result in this context.

Theorem 5.1 ([46, Theorems 5.1, 5.3 and 5.8]). Let $M \in \mathbb{R}^{n \times n}$ be a block-banded, nonnegative and symmetric positive definite matrix. Then, there exists banded, nonnegative and symmetric positive definite factor matrices B and C such that $\phi_M(B, C) = \|M - B \otimes C\|_F$ is minimized.

Thus, the properties of M are inherited by its approximation $B \otimes C$. Once the Kronecker product approximation has been computed, the lumped mass preconditioners can be constructed from the factor matrices B and C following Theorem 4.24. The overall construction consists in a two-level approximation of the mass matrix. Algorithm 5.1 is a practical algorithm for computing the preconditioner. For simplicity, we will assume that the preconditioners of the factor matrices have equal bandwidth.

Input: Mass matrix M

Output: Preconditioner $\tilde{P}_{ii} = \tilde{P}_{1,i} \otimes \tilde{P}_{2,i}$

- 1: Compute $\mathcal{R}(M)$
- 2: Compute a rank-1 approximation of $\mathcal{R}(M)$ such that $\mathcal{R}(M) \approx \mathbf{b}\mathbf{c}^T$
- 3: Reshape the vectors \mathbf{b} and \mathbf{c} to matrices B and C , respectively, and set $\tilde{M} = B \otimes C$
- 4: Construct the preconditioners $\tilde{P}_{1,i}$ and $\tilde{P}_{2,i}$ from B and C , respectively, following Definition 4.18
- 5: Set $\tilde{P}_{ii} = \tilde{P}_{1,i} \otimes \tilde{P}_{2,i}$

Algorithm 5.1: Lumped mass and Kronecker product approximation

More specifically, a practical implementation in Line 2 would extract a dense submatrix from $\mathcal{R}(M)$ containing all its nonzero entries and then compute a low-rank approximation of that submatrix only. The truncated singular value decomposition is an obvious choice for computing a low-rank approximation. However, there exists a host of other black-box algorithms for computing good but not optimal low-rank approximations. Some of them do not require the matrix to be explicitly available. It was the main motivation in [44] for using Adaptive Cross Approximation (ACA). In our context, it means we could potentially compute a good Kronecker product approximation of the mass matrix without ever assembling it. In Line 5, the preconditioner \tilde{P}_{ii} is not necessarily formed explicitly. Only the factors $\tilde{P}_{1,i}$ and $\tilde{P}_{2,i}$ are stored.

Clearly, we can expect the approximation error of M by $B \otimes C$ to be related to the low-rank approximation error of $\mathcal{R}(M)$ by $\text{vec}(B) \text{vec}(C)^T$. The next argument attempts to formalize this intuition. Assuming $\mathcal{R}(M)$ has rank r , its reduced singular value decomposition (SVD) is

$$\mathcal{R}(M) = U_r \Sigma_{rr} V_r^T = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + U_{r-1} \Sigma_{r-1, r-1} V_{r-1}^T$$

where $U_r = [\mathbf{u}_1, \dots, \mathbf{u}_r]$ and $V_r = [\mathbf{v}_1, \dots, \mathbf{v}_r]$ have orthonormal columns and are formed by the left and right singular vectors, respectively, and $\Sigma_{rr} = \text{diag}(\sigma_1, \dots, \sigma_r)$ is the diagonal matrix of decreasingly ordered singular values. In Algorithm 5.1, Line 2, we may set for instance $\mathbf{b} = \sqrt{\sigma_1} \mathbf{u}_1$ and $\mathbf{c} = \sqrt{\sigma_1} \mathbf{v}_1$. Then, by construction, the Frobenius norm of the error $E = \tilde{M} - M = B \otimes C - M$ is given by

$$\|E\|_F = \|\mathcal{R}(E)\|_F = \|\mathcal{R}(M) - \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T\|_F = \|U_{r-1} \Sigma_{r-1, r-1} V_{r-1}^T\|_F = \sqrt{\sum_{i=2}^r \sigma_i^2}.$$

Our ability to approximate the spectrum of M also depends on the low-rank approximation error of $\mathcal{R}(M)$. This fact is a direct consequence of [20, Corollary 6.3.8], applied to the symmetric matrices M and \tilde{M}

$$\sum_{i=1}^n (\lambda_i(M) - \lambda_i(\tilde{M}))^2 \leq \|E\|_F^2 = \sum_{i=2}^r \sigma_i^2.$$

The error we are committing will surely affect all eigenvalues. However, our main concern when approximating M by \tilde{M} is to guarantee that $\lambda_k(K, \tilde{M})$ remains close to $\lambda_k(K, M)$ for the smallest eigenvalues while underestimating the largest eigenvalues. We have already noted in Corollary 4.5 that if the error $E = \tilde{M} - M$ is symmetric positive semidefinite, then $\lambda_k(K, \tilde{M}) \leq \lambda_k(K, M)$. Unfortunately, this property is not satisfied by the nearest Kronecker product preconditioner, the error matrix E being in general indefinite, and the eigenvalues of (K, \tilde{M}) might be greater than those of (K, M) . However, as we have seen in Example 4.11, the positive semidefiniteness of the error only guarantees an underestimate of the eigenvalues, and it is not a necessary condition. Moreover, our strategy relies on a two-level approximation of M . It is first approximated by a Kronecker structured matrix \tilde{M} and then \tilde{M} is approximated by a lumped mass preconditioner \tilde{P}_{ii} . The relation between $\lambda_k(K, \tilde{M})$ and $\lambda_k(K, \tilde{P}_{ii})$ is well understood thanks to Theorem 4.24, guaranteeing that $\lambda_k(K, \tilde{P}_{ii}) \leq \lambda_k(K, \tilde{M})$. The problem of understanding the relation between $\lambda_k(K, M)$ and $\lambda_k(K, \tilde{M})$ once again amounts to assessing the quality of the preconditioner \tilde{M} . From (4.2),

$$\lambda_1(M, \tilde{M}) \leq \frac{\lambda_k(K, \tilde{M})}{\lambda_k(K, M)} \leq \lambda_n(M, \tilde{M}).$$

Although $\lambda_n(M, \tilde{M})$ may be greater than 1 due to the indefiniteness of the error, we will show numerically that it remains very close to 1. The use of Kronecker product approximations as preconditioners is not a new idea. It has already been explored for various applications including image processing [48], Markov chains [49], stochastic Galerkin finite element discretizations [50] and isogeometric analysis [26, 27, 40]. However, only few authors have undertaken the challenge of obtaining bounds for the preconditioned spectrum and most attempts are application specific. The following theorem is a variation of a general result presented in [50, Lemma 5.3] obtained by exhibiting the singular values and giving a more explicit bound on the condition number of the preconditioned matrix.

Theorem 5.2. Let $M = \sum_{i=1}^r B_i \otimes C_i$ be a symmetric positive definite matrix of Kronecker rank r where each factor matrices B_i, C_i are symmetric for $i = 1, \dots, r$ and let $\mathcal{R}(M) = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$ be the reduced singular value decomposition of the reordered matrix. Let \tilde{M} be the nearest Kronecker product approximation. Then, assuming that $\delta = \sum_{i=2}^r \frac{\sigma_i}{\sigma_1} \max_k |\lambda_k(U_i, U_1)| \max_k |\lambda_k(V_i, V_1)| < 1$ where $\mathbf{u}_i = \text{vec}(U_i)$ and $\mathbf{v}_i = \text{vec}(V_i)$, it holds

$$\kappa(\tilde{M}^{-1/2} M \tilde{M}^{-1/2}) \leq \frac{1 + \delta}{1 - \delta}$$

where κ denotes the spectral condition number.

Proof. By reverting the ordering, M is expressed as a sum of r Kronecker products

$$\mathcal{R}(M) = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T \iff M = \sum_{i=1}^r \sigma_i (U_i \otimes V_i)$$

where $\mathbf{u}_i = \text{vec}(U_i)$ and $\mathbf{v}_i = \text{vec}(V_i)$ for $i = 1, \dots, r$. Similarly, by construction, $\mathcal{R}(\tilde{M}) = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$ and thus $\tilde{M} = \sigma_1 (U_1 \otimes V_1)$. The symmetry and positive definiteness of U_1 and V_1 has already been established in [46, Theorem 5.8]. We now establish symmetry of all U_i and V_i for $i = 1, \dots, r$ using proof arguments similar to those presented in [49, Theorem 4.1]. By the singular value decomposition, $\mathcal{R}(M) \mathbf{v}_k = \sigma_k \mathbf{u}_k$ and $\mathcal{R}(M)^T \mathbf{u}_k = \sigma_k \mathbf{v}_k$ for $k = 1, \dots, r$. Since $\mathcal{R}(M) = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T = \sum_{i=1}^r \mathbf{b}_i \mathbf{c}_i^T$ with $\mathbf{b}_i = \text{vec}(B_i)$ and $\mathbf{c}_i = \text{vec}(C_i)$, we obtain

$$\begin{aligned} \mathbf{u}_k &= \frac{1}{\sigma_k} \mathcal{R}(M) \mathbf{v}_k = \frac{1}{\sigma_k} \sum_{i=1}^r \mathbf{b}_i \mathbf{c}_i^T \mathbf{v}_k = \sum_{i=1}^r \alpha_{ik} \mathbf{b}_i, \\ \mathbf{v}_k &= \frac{1}{\sigma_k} \mathcal{R}(M)^T \mathbf{u}_k = \frac{1}{\sigma_k} \sum_{i=1}^r \mathbf{c}_i \mathbf{b}_i^T \mathbf{u}_k = \sum_{i=1}^r \beta_{ik} \mathbf{c}_i, \end{aligned}$$

with $\alpha_{ik} = \sigma_k^{-1} \mathbf{c}_i^T \mathbf{v}_k$ and $\beta_{ik} = \sigma_k^{-1} \mathbf{b}_i^T \mathbf{u}_k$. By reshaping \mathbf{u}_k and \mathbf{v}_k to matrices, we obtain

$$U_k = \sum_{i=1}^r \alpha_{ik} B_i, \quad V_k = \sum_{i=1}^r \beta_{ik} C_i.$$

Since all B_i and C_i are symmetric by assumption, so are U_k and V_k for all $k = 1, \dots, r$. We now establish the

bound on the condition number by noting that

$$\begin{aligned}
|\lambda_k(M, \tilde{M}) - 1| &\leq \|\tilde{M}^{-1/2} M \tilde{M}^{-1/2} - I\|_2 = \left\| \sum_{i=2}^r \frac{\sigma_i}{\sigma_1} (U_1^{-1/2} U_i U_1^{-1/2} \otimes V_1^{-1/2} V_i V_1^{-1/2}) \right\|_2 \\
&\leq \sum_{i=2}^r \frac{\sigma_i}{\sigma_1} \|U_1^{-1/2} U_i U_1^{-1/2}\|_2 \|V_1^{-1/2} V_i V_1^{-1/2}\|_2 \\
&= \sum_{i=2}^r \frac{\sigma_i}{\sigma_1} \max_k |\lambda_k(U_i, U_1)| \max_k |\lambda_k(V_i, V_1)| = \delta.
\end{aligned}$$

Assuming that $\delta < 1$, since $\tilde{M}^{-1/2} M \tilde{M}^{-1/2}$ is symmetric positive definite,

$$\kappa(\tilde{M}^{-1/2} M \tilde{M}^{-1/2}) = \frac{\lambda_n(M, \tilde{M})}{\lambda_1(M, \tilde{M})} \leq \frac{1 + \delta}{1 - \delta}.$$

□

The previous theorem is not entirely satisfactory since δ still depends on the unknown quantities $\max_k |\lambda_k(U_i, U_1)|$ and $\max_k |\lambda_k(V_i, V_1)|$, which we were unable to bound in any convenient way. Nevertheless, our numerical experiments indicated that these quantities were $O(1)$. Most importantly, δ depends on the singular value ratios $\frac{\sigma_i}{\sigma_1}$ for $i = 2, \dots, r$. Therefore, a large singular value gap ($\sigma_2 \ll \sigma_1$) is important for constructing a good preconditioner and the smallest singular values will not contribute much to the sum. The quality of the approximation is now assessed numerically on a few examples.

Example 5.3 (Low-rank approximation). We consider two 2D examples from isogeometric analysis on nontrivial geometries. A brief description is given below.

- **Example 1:** The domain is a stretched square already used by several other authors in [26, 27] (Figure 5.1).
- **Example 2:** The domain is a perturbed triangle featuring a re-entrant corner (Figure 5.3).

We consider the same discretization parameters in all examples, namely cubic B-splines of maximal smoothness and 20 subdivisions in each parametric direction. The factor matrices B and C are obtained by truncating the singular value decomposition of a dense submatrix extracted from $\mathcal{R}(M)$ and containing all its nonzero entries. The rank of $\mathcal{R}(M)$ is exactly 2 for Example 1 since $|\det(J_F(\tilde{\mathbf{x}}))|$ is the sum of two separable functions for the stretched square domain. The second example is more challenging with $\mathcal{R}(M)$ having numerical rank 22. For each case, the discrete eigenvalues using the consistent mass matrix are compared to the approximate eigenvalues using the nearest Kronecker product (NKP) preconditioner with and without mass lumping. The curve using \tilde{P}_{33} nearly overlaps with the one for the NKP preconditioner and is omitted.




Figure 5.1: Stretched square domain

Figure 5.2: Comparison of $\Lambda(K, M)$ and $\Lambda(K, \tilde{M})$ for the nearest Kronecker product (NKP) preconditioner and $\Lambda(K, \tilde{P}_{ii})$ for $i = 1, 2, 3$




Figure 5.3: Domain with a re-entrant corner

The condition number of the preconditioned mass matrix with the NKP preconditioner is 2.18 and 39.1 for Examples 1 and 2, respectively. The assumptions of Theorem 5.2 are satisfied for the first example and the upper bound on the condition number is 2.73. The assumptions are not satisfied for the second example due to the much slower decay of the singular values. In Figure 5.2, we notice that the largest eigenvalues of (K, \tilde{M}) are actually slightly greater than those of (K, M) . This situation cannot be prevented but our second level approximation involving mass lumping usually fixes it. On the contrary, in Figure 5.4, the largest eigenvalues of (K, M) are greatly underestimated when using the NKP preconditioner. This situation is very favorable and was generally experienced in our experiments with complicated problems. Any mass lumping done subsequently will bring down the eigenvalues further. However, the bandwidth must still be large enough for the smallest eigenvalues to be well approximated. The selective damping of the higher modes will be addressed in future work.

5.2 Approximation for 3D problems

The natural analog of the minimization problem for 3D discretizations is

$$\min \phi_M(B, C, D) = \min \|M - B \otimes C \otimes D\|_F$$

where $M \in \mathbb{R}^{n \times n}$ with $n = n_1 n_2 n_3$ and the factor matrices $B \in \mathbb{R}^{n_1 \times n_1}$, $C \in \mathbb{R}^{n_2 \times n_2}$ and $D \in \mathbb{R}^{n_3 \times n_3}$. This extension was not originally investigated by Van Loan and Pitsianis [46] but was later analyzed in [49]. The definition of the mapping \mathcal{R} must be adjusted such that $\mathcal{R}(B \otimes C \otimes D) = \text{vec}(B) \circ \text{vec}(C) \circ \text{vec}(D)$. More generally, it transforms a Kronecker product of p factor matrices to a p th order rank-1 tensor obtained by taking the outer product of the vectorization of the factor matrices. The symbol \circ is used to denote the outer product. Note that this definition is consistent with the one previously introduced for the 2D case since $\mathbf{b} \circ \mathbf{c} = \mathbf{bc}^T$. The minimization problem then translates to finding the best rank-1 tensor approximation of $\mathcal{R}(M)$. Unfortunately, contrary to the matrix case, we cannot hope to compute this *best* rank-1 approximation and we must content ourselves with *good* rank-1 approximations such as those computed with the high order SVD (HOSVD) [51]. Another major shortcoming of the extension to higher dimensions is the lack of theoretical results related to the computed factor matrices. Some proof arguments used in [46] for the two dimensional case cannot be straightforwardly extended to higher dimensions. On the other hand, Algorithm 5.1 can be straightforwardly adapted to higher dimensional problems. The computational results were quite similar to those reported in Example 5.3 and are omitted.

6 Conclusion

Mass lumping is used on a daily basis for industrial applications in structural dynamics and consists in replacing the consistent mass matrix with some kind of easily invertible (typically diagonal) approximation. In this article,




Figure 5.4: Comparison of $\Lambda(K, M)$ and $\Lambda(K, \tilde{M})$ for the nearest Kronecker product (NKP) preconditioner and $\Lambda(K, \tilde{P}_{ii})$ for $i = 1, 2, 3$

we have unraveled some attractive mathematical features of mass lumping partly responsible for its success. Our results hold under rather broad assumptions, indicating that some properties are not specific to any discretization method. More generally, we have highlighted the close connection between the quality of a preconditioner and the approximation of the eigenvalues. Based on this insight, we have generalized the concept of mass lumping and its properties to matrices with more complex structures including Kronecker products, whose structure can be used to solve linear systems very efficiently. We have then proposed a two-level approximation of the mass matrix in case it is not already expressed as a Kronecker product. Numerical experiments have highlighted the practical usefulness of this approximation in the context of isogeometric analysis by increasing the critical time step of explicit time integration schemes. However, our generalization of mass lumping improves the approximation of all discrete eigenvalues, including the outliers. Outlier eigenvalues are a distinctive feature of maximally smooth isogeometric discretizations. Apart from being completely meaningless, they significantly constrain the critical time step of explicit time integration schemes. Understanding and removing these outliers is therefore of high interest and several follow-up papers have been written on this topic; see e.g., [17, 45, 52, 53, 54]. Future work will attempt to selectively approximate the accurate low-frequency part of the spectrum while undermining the outliers. One possible research direction consists in combining the mass lumping techniques presented here with the optimal spline spaces in [55] which were shown to be outlier-free in [53] and [54].

Acknowledgments

The second author kindly acknowledges the support of the SNSF through the project “Smoothness in Higher Order Numerical Methods” n. TMPFP2_209868 (SNSF Swiss Postdoctoral Fellowship 2021).

The third author kindly acknowledges the support of SNSF through the project “Design- through-Analysis (of PDEs): the litmus test” n. 40B2-0 187094 (BRIDGE Discovery 2019).

References

- [1] O. C. Zienkiewicz, R. L. Taylor, J. Z. Zhu, *The finite element method: its basis and fundamentals*, Elsevier, 2005.
- [2] T. J. Hughes, *The finite element method: linear static and dynamic finite element analysis*, Courier Corporation, 2012.
- [3] S. Duczek, H. Gravenkamp, Mass lumping techniques in the spectral element method: On the equivalence of the row-sum, nodal quadrature, and diagonal scaling methods, *Computer Methods in Applied Mechanics and Engineering* 353 (2019) 516–569.
- [4] T. J. Hughes, J. A. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, *Computer methods in applied mechanics and engineering* 194 (39-41) (2005) 4135–4195.
- [5] J. A. Cottrell, T. J. Hughes, Y. Bazilevs, *Isogeometric analysis: toward integration of CAD and FEA*, John Wiley & Sons, 2009.
- [6] J. A. Cottrell, A. Reali, Y. Bazilevs, T. J. Hughes, Isogeometric analysis of structural vibrations, *Computer methods in applied mechanics and engineering* 195 (41-43) (2006) 5257–5296.
- [7] L. Leidinger, *Explicit Isogeometric B-Rep Analysis for Nonlinear Dynamic Crash Simulations*, Ph.D. thesis, Technische Universität München (2020).
- [8] C. Anitescu, C. Nguyen, T. Rabczuk, X. Zhuang, Isogeometric analysis for explicit elastodynamics using a dual-basis diagonal mass formulation, *Computer Methods in Applied Mechanics and Engineering* 346 (2019) 574–591.
- [9] I. Fried, D. S. Malkus, Finite element mass matrix lumping by numerical integration with no convergence rate loss, *International Journal of Solids and Structures* 11 (4) (1975) 461–466.
- [10] G. Cohen, P. Joly, N. Tordjman, Higher-order finite elements with mass-lumping for the 1D wave equation, *Finite elements in analysis and design* 16 (3-4) (1994) 329–336.
- [11] M. Duruffé, P. Grob, P. Joly, Influence of Gauss and Gauss-Lobatto quadrature rules on the accuracy of a quadrilateral finite element method in the time domain, *Numerical Methods for Partial Differential Equations: An International Journal* 25 (3) (2009) 526–551.

- [12] G. Cohen, P. Joly, J. E. Roberts, N. Tordjman, Higher order triangular finite elements with mass lumping for the wave equation, *SIAM Journal on Numerical Analysis* 38 (6) (2001) 2047–2078.
- [13] D. S. Malkus, M. E. Plesha, Zero and negative masses in finite element vibration and transient analysis, *Computer methods in applied mechanics and engineering* 59 (3) (1986) 281–306.
- [14] E. Hinton, T. Rock, O. Zienkiewicz, A note on mass lumping and related processes in the finite element method, *Earthquake Engineering & Structural Dynamics* 4 (3) (1976) 245–249.
- [15] A. Quarteroni, *Numerical models for differential problems*, Vol. 2, Springer, 2009.
- [16] N. M. Newmark, A method of computation for structural dynamics, *Journal of the engineering mechanics division* 85 (3) (1959) 67–94.
- [17] T. J. Hughes, J. A. Evans, A. Reali, Finite element and NURBS approximations of eigenvalue, boundary-value, and initial-value problems, *Computer Methods in Applied Mechanics and Engineering* 272 (2014) 290–320.
- [18] G. Stewart, J. Sun, *Matrix Perturbation Theory*, Computer Science and Scientific Computing, ACADEMIC Press, INC, 1990.
- [19] B. N. Parlett, *The symmetric eigenvalue problem*, SIAM, 1998.
- [20] R. A. Horn, C. R. Johnson, *Matrix analysis*, Cambridge university press, 2012.
- [21] M. A. Schweitzer, Variational mass lumping in the partition of unity method, *SIAM Journal on Scientific Computing* 35 (2) (2013) A1073–A1097.
- [22] Y. Yang, H. Zheng, M. Sivaselvan, A rigorous and unified mass lumping scheme for higher-order elements, *Computer Methods in Applied Mechanics and Engineering* 319 (2017) 491–514.
- [23] L. Olovsson, K. Simonsson, M. Unosson, Selective mass scaling for explicit finite element analyses, *International Journal for Numerical Methods in Engineering* 63 (10) (2005) 1436–1445.
- [24] S. K. Stoter, T.-H. Nguyen, R. R. Hiemstra, D. Schillinger, Variationally consistent mass scaling for explicit time-integration schemes of lower-and higher-order finite element methods, *arXiv preprint arXiv:2201.10475* (2022).
- [25] A. Tkachuk, M. Bischoff, Variational methods for selective mass scaling, *Computational Mechanics* 52 (3) (2013) 563–570.
- [26] L. Gao, V. M. Calo, Fast isogeometric solvers for explicit dynamics, *Computer Methods in Applied Mechanics and Engineering* 274 (2014) 19–41.
- [27] G. Loli, G. Sangalli, M. Tani, Easy and efficient preconditioning of the isogeometric mass matrix, *Computers & Mathematics with Applications* (2021).
- [28] R. Vázquez, A new design for the implementation of isogeometric analysis in Octave and Matlab: GeoPDEs 3.0, *Computers & Mathematics with Applications* 72 (3) (2016) 523–554.
- [29] G. W. Stewart, Gershgorin theory for the generalized eigenvalue problem $Ax=\lambda Bx$, *Mathematics of Computation* (1975) 600–606.
- [30] G. W. Stewart, Perturbation bounds for the definite generalized eigenvalue problem, *Linear algebra and its applications* 23 (1979) 69–85.
- [31] J.-G. Sun, A note on Stewart’s theorem for definite matrix pairs, *Linear Algebra and its Applications* 48 (1982) 331–339.
- [32] C. R. Crawford, A stable generalized eigenvalue problem, *SIAM Journal on Numerical Analysis* 13 (6) (1976) 854–860.
- [33] I. C. Ipsen, B. Nadler, Refined perturbation bounds for eigenvalues of Hermitian and non-Hermitian matrices, *SIAM Journal on Matrix Analysis and Applications* 31 (1) (2009) 40–53.
- [34] C. R. Crawford, *The numerical solution of the generalized eigenvalue problem*, Ph.D. thesis, University of Michigan (1970).

- [35] G. H. Golub, C. F. Van Loan, *Matrix computations*, JHU press, 2013.
- [36] A. Quarteroni, R. Sacco, F. Saleri, *Numerical mathematics*, Vol. 37, Springer Science & Business Media, 2010.
- [37] A. J. Wathen, Preconditioning, *Acta Numerica* 24 (2015) 329–376.
- [38] J. W. Pearson, J. Pestana, Preconditioners for Krylov subspace methods: An overview, *GAMM-Mitteilungen* 43 (4) (2020) e202000015.
- [39] R. A. Horn, C. R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, 1991.
- [40] G. Sangalli, M. Tani, Isogeometric preconditioners based on fast solvers for the Sylvester equation, *SIAM Journal on Scientific Computing* 38 (6) (2016) A3644–A3671.
- [41] Y. Voet, *Nonlinear Finite Elements in Dynamics*, Tech. rep., École polytechnique fédérale de Lausanne (2020).
- [42] A. Mantzaflaris, B. Jüttler, B. N. Khoromskij, U. Langer, Low rank tensor methods in Galerkin-based isogeometric analysis, *Computer Methods in Applied Mechanics and Engineering* 316 (2017) 1062–1085.
- [43] F. Scholz, A. Mantzaflaris, B. Jüttler, Partial tensor decomposition for decoupling isogeometric Galerkin discretizations, *Computer Methods in Applied Mechanics and Engineering* 336 (2018) 485–506.
- [44] C. Hofreither, A black-box low-rank approximation algorithm for fast matrix assembly in isogeometric analysis, *Computer Methods in Applied Mechanics and Engineering* 333 (2018) 311–330.
- [45] J. Chan, J. A. Evans, Multi-patch discontinuous Galerkin isogeometric analysis for wave propagation: Explicit time-stepping and efficient mass matrix inversion, *Computer Methods in Applied Mechanics and Engineering* 333 (2018) 22–54.
- [46] C. F. Van Loan, N. Pitsianis, Approximation with Kronecker products, in: *Linear algebra for large scale and real-time applications*, Springer, 1993, pp. 293–314.
- [47] N. P. Pitsianis, *The Kronecker product in approximation and fast transform generation*, Ph.D. thesis, Cornell University (1997).
- [48] J. G. Nagy, M. E. Kilmer, Kronecker product approximation for preconditioning in three-dimensional imaging applications, *IEEE Transactions on Image Processing* 15 (3) (2006) 604–613.
- [49] A. N. Langville, W. J. Stewart, A Kronecker product approximate preconditioner for SANs, *Numerical Linear Algebra with Applications* 11 (8-9) (2004) 723–752.
- [50] E. Ullmann, A Kronecker product preconditioner for stochastic Galerkin finite element discretizations, *SIAM Journal on Scientific Computing* 32 (2) (2010) 923–946.
- [51] L. De Lathauwer, B. De Moor, J. Vandewalle, A multilinear singular value decomposition, *SIAM journal on Matrix Analysis and Applications* 21 (4) (2000) 1253–1278.
- [52] D. Gallistl, P. Huber, D. Peterseim, On the stability of the Rayleigh–Ritz method for eigenvalues, *Numerische Mathematik* 137 (2) (2017) 339–351.
- [53] E. Sande, C. Manni, H. Speleers, Sharp error estimates for spline approximation: Explicit constants, n-widths, and eigenfunction convergence, *Mathematical Models and Methods in Applied Sciences* 29 (06) (2019) 1175–1205.
- [54] C. Manni, E. Sande, H. Speleers, Application of optimal spline subspaces for the removal of spurious outliers in isogeometric discretizations, *Computer Methods in Applied Mechanics and Engineering* 389 (2022) 114260.
- [55] M. S. Floater, E. Sande, Optimal Spline Spaces for L2 n-Width Problems with Boundary Conditions, *Constructive Approximation* 50 (1) (2019) 1–18.