Gecko
Design for IGA-type discretization workflows

Gecko Technical Report 2

DC7 – Angelos Pagonas

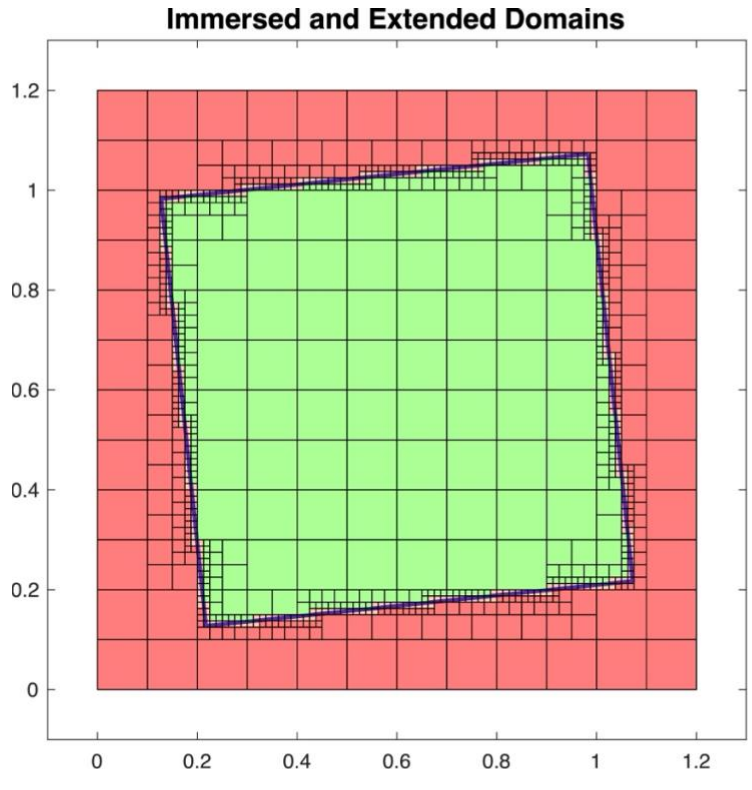

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101073106
Call: HORIZON-MSCA-2021-DN-01

Funded by the European Union

Executive summary

This report addresses the use of Isogeometric Analysis for explicit structural dynamics in immersed geometries. The combination is attractive, since the B-spline basis gives an accurate discrete spectrum and immersion avoids body-fitted meshing, but it creates a specific difficulty: the small cut elements produced by immersion result in the divergence of the largest eigenfrequency, leading to a collapse of the stable explicit time step. Mass lumping restores the time step at a cost in accuracy. This interaction between immersion and the treatment of the mass, and its consequences for accuracy and efficiency, is the subject of this work.

The work advances along two complementary directions. In the one-dimensional setting, the bar and the Bernoulli–Euler beam are examined in depth: a verified framework is put in place, the accuracy cost of mass lumping is presented and connected to the initial conditions, as well as compared with standard finite elements. In the treatment of complex immersed geometry, a controlled example shows that the mass of a cut element alone does not predict the most severe frequencies, which depend instead on the surrounding discretisation. Each direction exposes the limit of a single scalar: in one dimension the displacement error reports a loss of accuracy without explaining it, and the geometric example just given makes the same point for the prediction of frequencies. This motivates looking inside the computation rather than at its final output. First steps toward selective mass scaling, an alternative mass treatment, are also reported.

The report establishes a verified basis for the one-dimensional studies and sets the direction of the continuing work: to study immersion and mass lumping separately before their combination, and to identify the causes of these effects earlier in the computational pipeline, ahead of the final L2 error.

Table of contents

List of figures

List of tables

List of abbreviations

CAD	*Computer Aided Design*
FEM	*Finite Element Method*
IGA	*IsoGeometric Analysis*
PDE	*Partial Differential Equation*
TR1	*Technical Report 1*
nel	*Number of elements*
DOF	Degrees of freedom

Introduction

Explicit time integration is the method of choice for many transient structural-dynamics problems, and its cost is governed by the stable time step, which is in turn set by the highest eigenfrequency of the spatial discretisation. Isogeometric analysis (IGA) [Hu05] is attractive in this context: the high inter-element continuity of its spline bases yields a markedly more accurate discrete spectrum than standard finite elements [Co06], a direct advantage for dynamics. Immersed (or trimmed) discretisations are equally attractive on the geometric side, since they represent a complex domain by embedding it in a simple structured grid and resolving the geometry at the integration level [PDR07], avoiding body-fitted meshing.

Combining the two exposes a specific difficulty. Immersion produces cut elements whose physical content can be arbitrarily small, and these drive the highest eigenfrequency upward, so that the stable explicit time step collapses. Mass lumping can serve as a remedy: it bounds the eigenfrequency and restores a usable time step, but it modifies the mass operator and with it the accuracy of the solution.

It is useful to regard immersion and mass lumping as two independent research directions: each perturbs the discretisation in its own way and for its own reasons, and the two are best studied and understood separately before their combination is considered.

The report is organised along two lines: a deep one-dimensional study, and a continuation of the two-dimensional geometric study from the previous report.

That earlier report established the background reviewed above (TR1, Section 1) and introduced these one-dimensional test problems together with the initial conditions that are used again here (TR1, Section 2.1 and Section 2.2). The present report builds on that foundation: it establishes a verified one-dimensional framework in place, examines the accuracy of mass lumping within it, and extends the work along the two lines just described.

The report is organised as follows. Section 1 verifies the one-dimensional framework. Section 2 presents the accuracy of mass lumping in one dimension and Section 3 its dependence on the excitation. Section 4 takes the treatment of complex immersed geometry further. Section 5 introduces a diagnostic perspective aimed at explaining the effects rather than only reporting them, and Section 6 reports some preliminary results on selective mass scaling for iga. Sections 7 and 8 discuss the work in progress and conclusions.

1 VERIFICATION OF THE ONE-DIMENSIONAL FRAMEWORK

The accuracy results reported in the following sections rely on a verified one-dimensional framework. Verification proceeds at two levels.

First, in the boundary-fitted configuration ($\zeta=0$), where the isogeometric operators are unperturbed and closed-form references exist, the discrete spectrum is compared against analytical values, confirming that the stiffness and mass operators are assembled correctly.

Second, the time integration is checked with a discrete eigenmode as initial condition: with the period taken from the fully discrete frequency, the displacement L_2 error after one period (the accuracy metric used throughout this report) is limited to machine precision levels. This second check is not confined to the boundary-fitted case, it is carried out across the full range of immersion, so that it also certifies the assembly and marching of the cut elements. Together the two levels establish that the operators, the eigenproblem and the time integration are correct and robust before the accuracy effects of immersion and mass lumping are quantified in Section 2.

1.1 The discrete eigenspectrum

Verification of the spatial discretisation follows the path laid out by Cottrell et al. for structural vibrations [Co06]. Maximum-continuity b-splines are used throughout this report. In this section the consistent mass matrix is employed. The isogeometric eigenfrequencies are computed and compared mode by mode against the analytical k-method spectrum. For the longitudinal vibration of a fixed–fixed rod, the normalised frequency ω^h/ω is plotted against the normalised mode number n/N for degrees $p=2$ and $p=3$ (Fig. 1), with the analytical references Eqs. (32) and (33) of [Co06]. On a 2000-element mesh the computed eigenfrequencies coincide with the analytical curves across the entire spectrum and for both degrees.

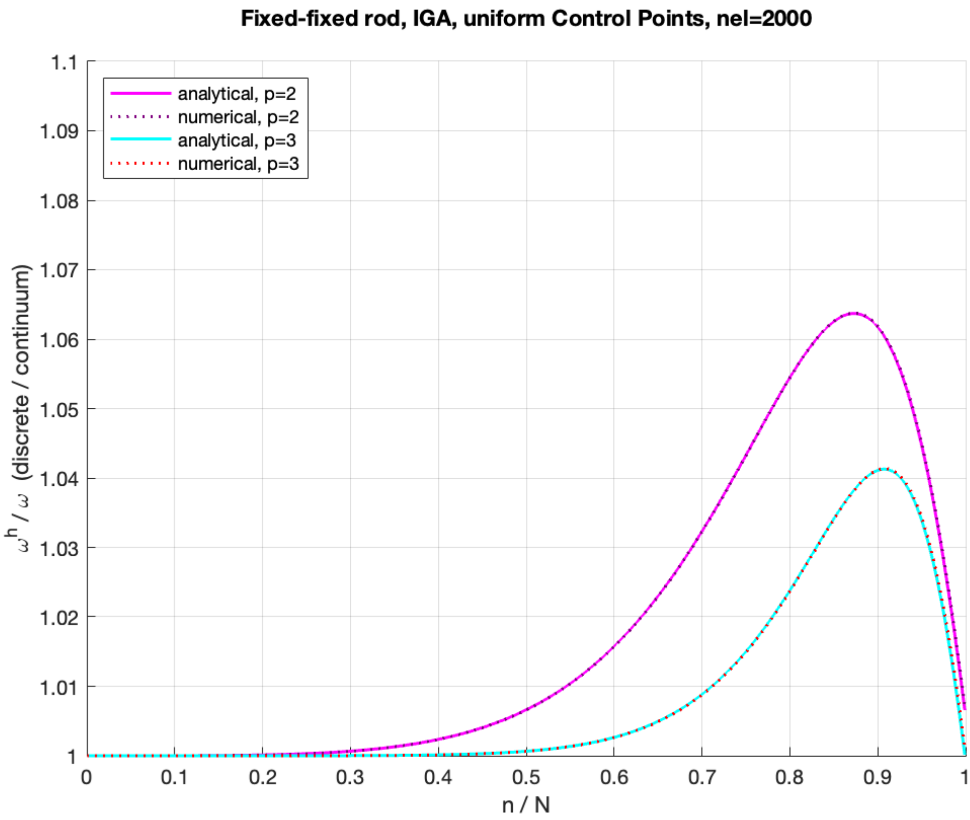

Figure 1: Discrete spectrum of the boundary-fitted fixed–fixed rod (IGA, uniform control points, consistent mass, ζ=0, nel=2000). Normalised frequency ω^h/ω against normalised mode number n/N for p=2 and p=3; solid curves are the analytical k-method spectra of [Co06], Eqs. (32) and (33), dotted curves the computed eigenfrequencies.

The same check is repeated for the bending vibration of a simply supported Bernoulli–Euler beam with cubic NURBS (p=3), whose analytical spectrum is given by Eq. (36) of [Co06] (Fig. 2). The computed spectrum again reproduces the analytical curve over the full range of modes. This agreement confirms that the operator assembly and the eigenvalue problem are correct, and that the discrete spectrum on which the time-domain study below relies is the intended one.

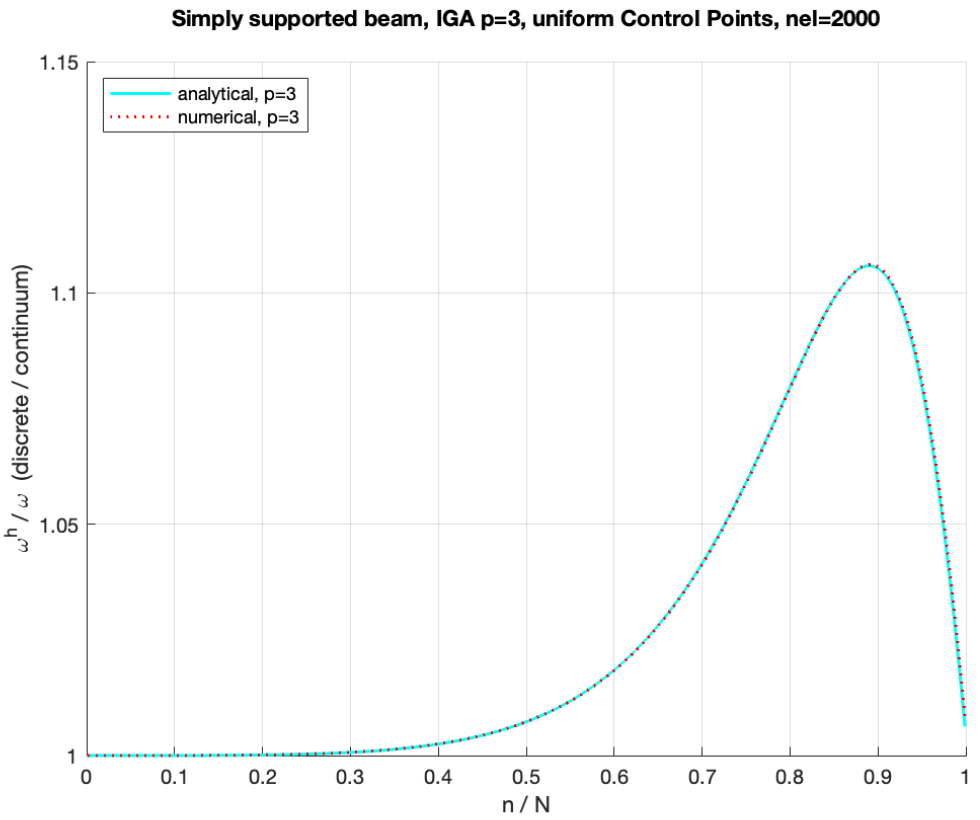

Figure 2: Discrete spectrum of the boundary-fitted simply supported Bernoulli–Euler beam (IGA, cubic NURBS p=3, uniform control points, consistent mass, ζ=0, nel=2000). Notation as in Fig. 1; analytical curve from [Co06], Eq. (36).

1.2 The time integration

The second level addresses the time integration. With the spatial operators now trusted, the remaining question is whether the explicit scheme advances them correctly. The test uses a single discrete vibration mode: when the simulation is started on one such mode, at rest, in exact arithmetic the solution retains that mode shape while only its amplitude varies, and it returns to its initial state after one full period. This holds because the mode is an eigenvector of the operators used to march it: the same stiffness and mass define both the eigenproblem and the time stepping, so the response stays within the one-dimensional subspace spanned by that mode.

The relevant subtlety is the length of the period. An explicit scheme does not advance the mode at the structure's exact frequency: discretising time introduces a small, well-understood shift, so the marched mode completes its cycle at a slightly different rate than the continuous one. If the run length were set from the exact frequency, the mode would return slightly out of phase, and that offset is itself the time-stepping error the test targets. The run is therefore timed to the scheme's own frequency with the step size chosen so that an integer number of steps completes exactly one cycle. A correct scheme then returns the mode to its starting state to within machine precision, while a larger error indicates potentially a fault somewhere in the process.

The discrepancy after one period (in the displacement L_2 norm) is measured against this discrete starting state, not the continuous analytical shape. The distinction matters: the scheme transports the discrete mode, and it is the discrete mode that should return, so comparing against the continuous eigenfunction would only re-measure the spatial error already settled by the spectrum study and could never fall to round-off. (Initial conditions built by projecting a continuum profile, where the projection error is common to both mass operators, are treated in Section 2.)

The check is repeated across a sequence of immersion configurations, from the boundary-fitted case to sliver cuts, for the rod (free–free, p=2 and p=3) and the Bernoulli–Euler beam (free–free, p=3 and p=4), with a consistent mass matrix and no stabilisation. Free ends are used here, rather than the fixed ends of the spectrum study, so that no condition need to be imposed on the immersed boundary. The rigid-body modes that the free ends introduce are identified and skipped when the elastic mode is chosen. The argument relies only on the same operators defining the mode and marching it, so it holds for every cut. A clean return therefore tests the whole immersed pipeline at once: assembly on the cut elements, the eigensolver, the timing of the period, and the marching.

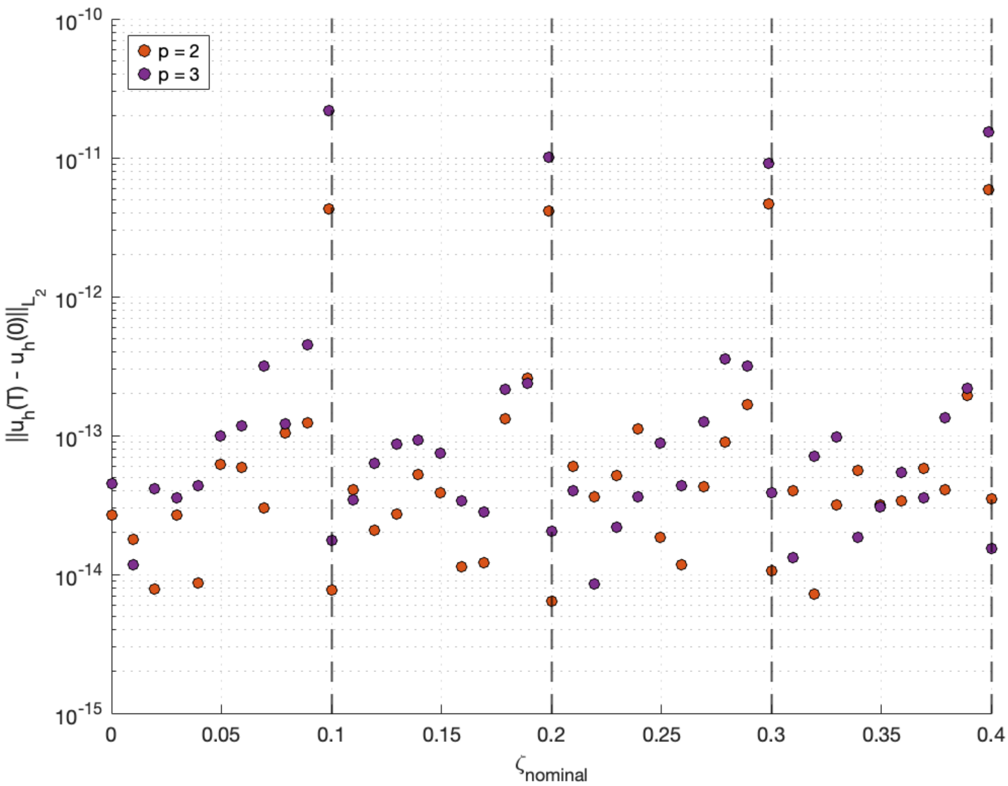

Figure 3: Time-domain self-consistency of the immersed rod (free–free, IGA, consistent mass, unstabilised, nel=100). Return-to-initial-condition L2 error $\|u_h(T)-u_h(0)\|_{L2}$ after one fully-discrete period, against the normalised fictitious-domain size $\zeta_{nominal}$, for p=2 and p=3; the initial condition is the first discrete eigenmode. Dashed lines mark the element boundaries, near which the thinnest cut elements occur.

For the rod the return-after-a-period error sits at the round-off levels of order 10^{-14}–10^{-13}, for most of the immersion scenarios and for both degrees (Fig. 3). The time integration is thereby verified, and shown to remain self-consistent under immersion. The only deviations are isolated jumps to $\approx 10^{-11}$ at the most degenerate sliver cuts, where a cut element of vanishing physical width leaves the immersed operators severely ill-conditioned.

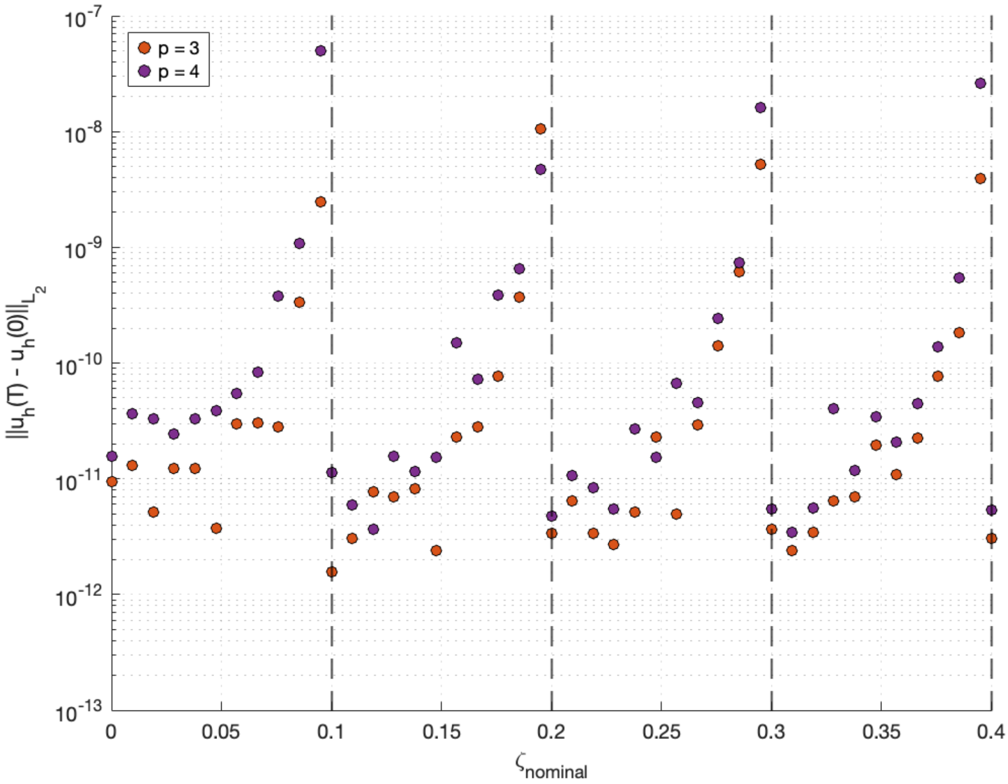

Figure 4: As Fig. 3, for the Bernoulli–Euler beam (free–free, IGA, consistent mass, unstabilised, nel=50, p=3 and p=4), first bending eigenmode. Note the higher vertical scale.

The beam behaves identically in kind (Fig. 4) but settles at a higher level of order 10^{-12}–10^{-10}, rising to $\approx 10^{-8}$ at the sliver cuts. This offset is expected and benign. The fourth-order beam is numerically far stiffer: its highest discrete frequency is much larger, reaching order 10^7 at the finest cuts, so the stable time step shrinks in proportion and a single period must be covered in up to millions of steps, over each of which a little round-off is unavoidable and accumulates. The error level therefore is affected by the stiffness and conditioning of the immersed fourth-order problem, not a flaw in the operators, the eigenmode, or the scheme.

The accuracy consequences of immersion and of mass manipulation are taken up in Section 2. Here the operators, the eigenproblem, and the time integration are established as correct and robust before any related result is presented.

2 MASS LUMPING AND ACCURACY IN ONE DIMENSION

TR1 established two results for the immersed one-dimensional rod and Bernoulli–Euler beam. With a consistent mass the largest eigenfrequency diverges as a cut element decreases, interestingly mass lumping bounds it independently of the cut size, but the time step recovered comes at a cost in accuracy (TR1, Section 1.2) [Le20, Ra24]. This section quantifies that accuracy cost for several immersion scenarios, for the rod and then the beam, on the framework verified in Section 1.

2.1 Setup and accuracy metric

The setting of [Ra24] is retained. A free–free element of length **l** is immersed in the extended domain [0, L] on a uniform background grid, and the fictitious-domain size ζ is modified from the boundary-fitted case to the deepest sliver cuts. The discretisation uses maximum-continuity b-splines, with nel=240 for the rod and nel=60 for the beam. The explicit step is $0.9 \cdot \Delta t_{crit}$.

Accuracy is measured by the displacement L_2 error after one period, $\|u_h(T) - u_h(0)\|_{L_2}$. In each case considered the structure should return to its initial state after one period, so this quantity should vanish for an exact computation (TR1, Section 1.2.2).

The initial condition is built by projecting the continuum field onto the discrete space in the L_2 sense, using the consistent mass matrix. Mass lumping is applied only to the operator that advances the dynamics. The projected initial state is then identical in the consistent- and lumped-mass runs, so its representation error enters both as the same fixed offset, and the error after one period reflects the effect of lumping on the dynamics rather than a difference in the initial condition. On every plot the critical time step Δt_{crit} is shown dashed, as the proxy for the stability gain.

2.2 Rod

2.2.1 Broadband excitation: the Gaussian pulse

For the rod under a travelling Gaussian pulse, the consistent-mass error follows the critical time step (Fig. 5a). As ζ increases the cut element fluctuates in a repeating pattern between successive element-boundary-aligned cuts, and Δt_{crit} follows along with it. Within each of these intervals the displacement error decreases together with Δt_{crit} toward the baseline set by the spatial discretisation, then recovers at the next element-boundary-aligned cut. The error stays of order 10^{-3}, set by the time step the immersion permits.

Lumping breaks this pattern (Fig. 5b). The critical time step is now bounded across the whole immersed scenarios range, which is the stability gain, but the error is held at a high, nearly cut-independent floor of order 10^{-1}. This is more than an order of magnitude above the consistent-mass error, and up to two orders above it at the deepest cuts. The accuracy carried by the consistent isogeometric spectrum (Section 1) is therefore largely lost to the row-sum lumping that secures the time step, as reported in [Ra24].

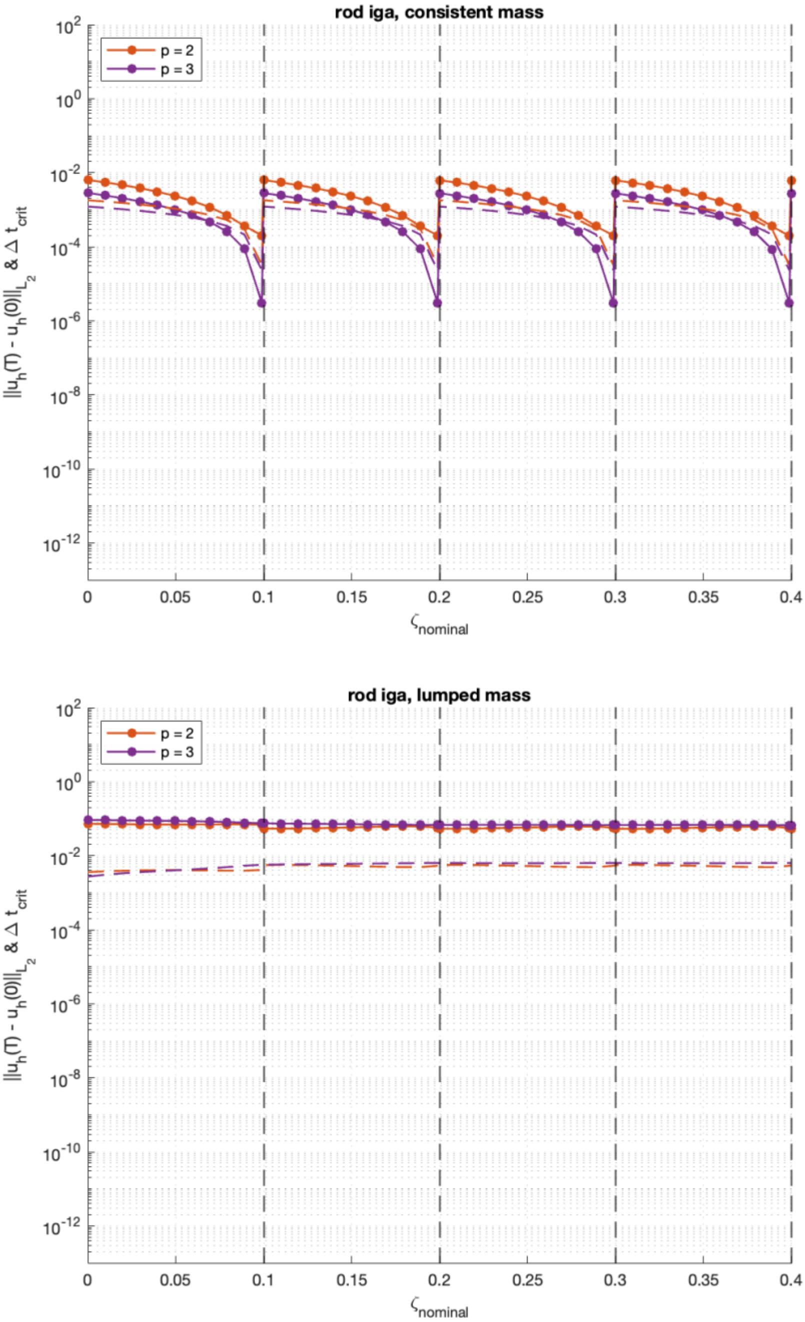

Figure 5:Accuracy cost of lumping for the immersed rod under a Gaussian pulse (IGA, free–free, nel=240, p=2 and p=3). L_2 error after one period $\|u_h(T)-u_h(0)\|_{L_2}$ (solid) and critical time step Δt_{crit} (dashed) against the fictitious-domain size $\zeta_{nominal}$, for (a) the consistent and (b) the lumped mass. Vertical lines mark the element-boundary-aligned cuts; the deepest cuts lie immediately to their left.

The same comparison against a standard finite element (C^0 Lagrange) discretisation, at approximately the same total number of DOF, clarifies the difference between the effect of the b-spline basis from that of lumping. With a consistent mass the two bases behave alike, the error follows the decreasing time step (Figs. 5a, 6a). Under lumping they differ in stability. Row-sum lumping bounds the b-spline eigenfrequency, so the isogeometric critical time step stays bounded for all immersion cases (Fig. 5b). For the Lagrange basis lumping reduces the eigenfrequency but does not bound it, and the critical time step still collapses at the small cuts (Fig. 6b). The bounded explicit time step under immersion is therefore a property of the high-continuity b-spline basis, not of lumping alone [Le20, Ra24].

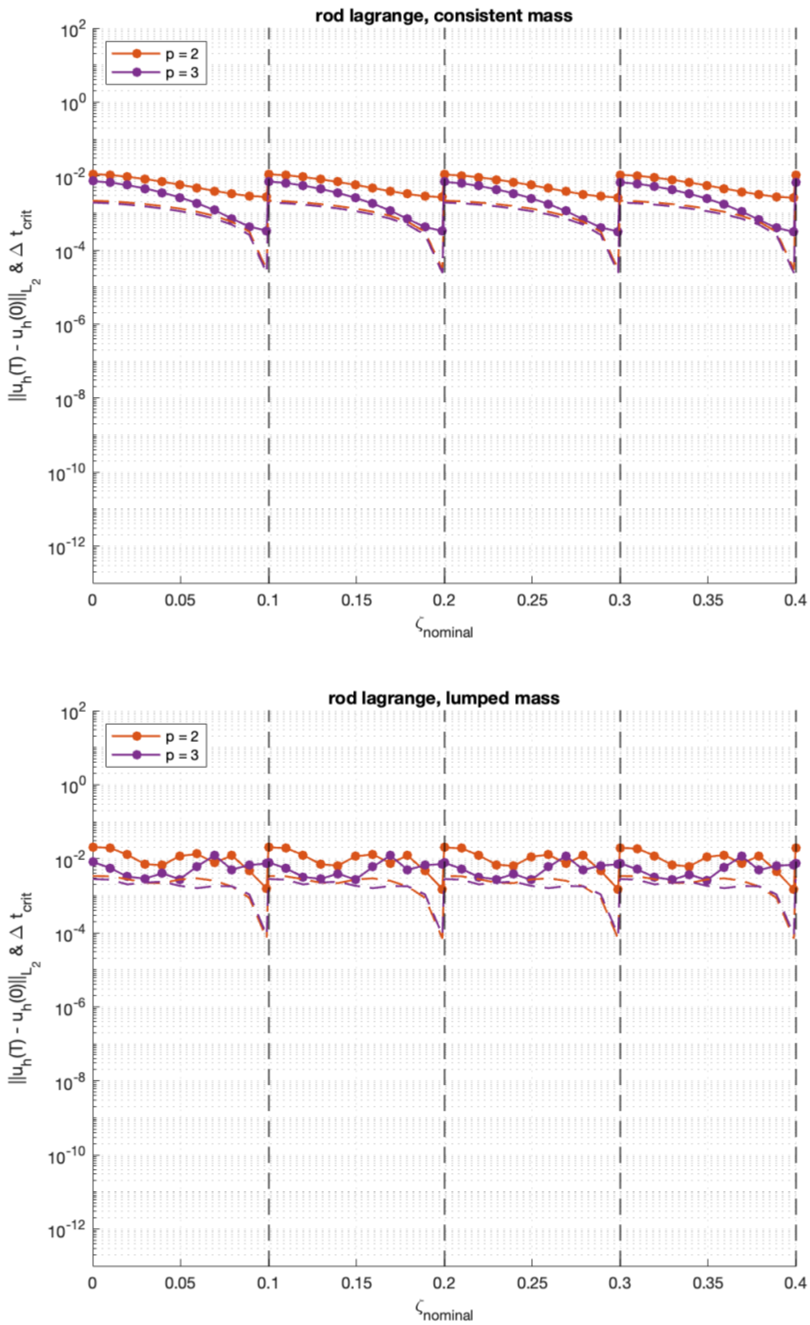

Figure 6: As Fig. 5, for a standard finite element (C^0 Lagrange) discretisation at matched total DOF (p=2: 120 elements, 241 DOF; p=3: 80 elements, 241 DOF; against 242–243 DOF for the isogeometric rod. L2 error after one period (solid) and Δ_{tcrit} (dashed) against $\zeta_{nominal}$, for (a) consistent and (b) lumped mass.)

2.2.2 Single-mode excitation

The Gaussian pulse excites a broad part of the spectrum, while a single low mode probes the opposite extreme. The initial condition here is the first non-rigid-body mode, which the isogeometric space resolves almost exactly. With a consistent mass the rod returns to its initial state at a low threshold of order 10^{-12} across the immersion cases, set by how well the discrete space represents this mode rather than by the time step. Mass lumping raises the return error to a small but distinct threshold of order 10^{-8}, nearly independent of the cut, with the critical time step bounded throughout (Fig. 7). The dependence of the lumping error on the part of the spectrum that is excited is the subject of Section 3.

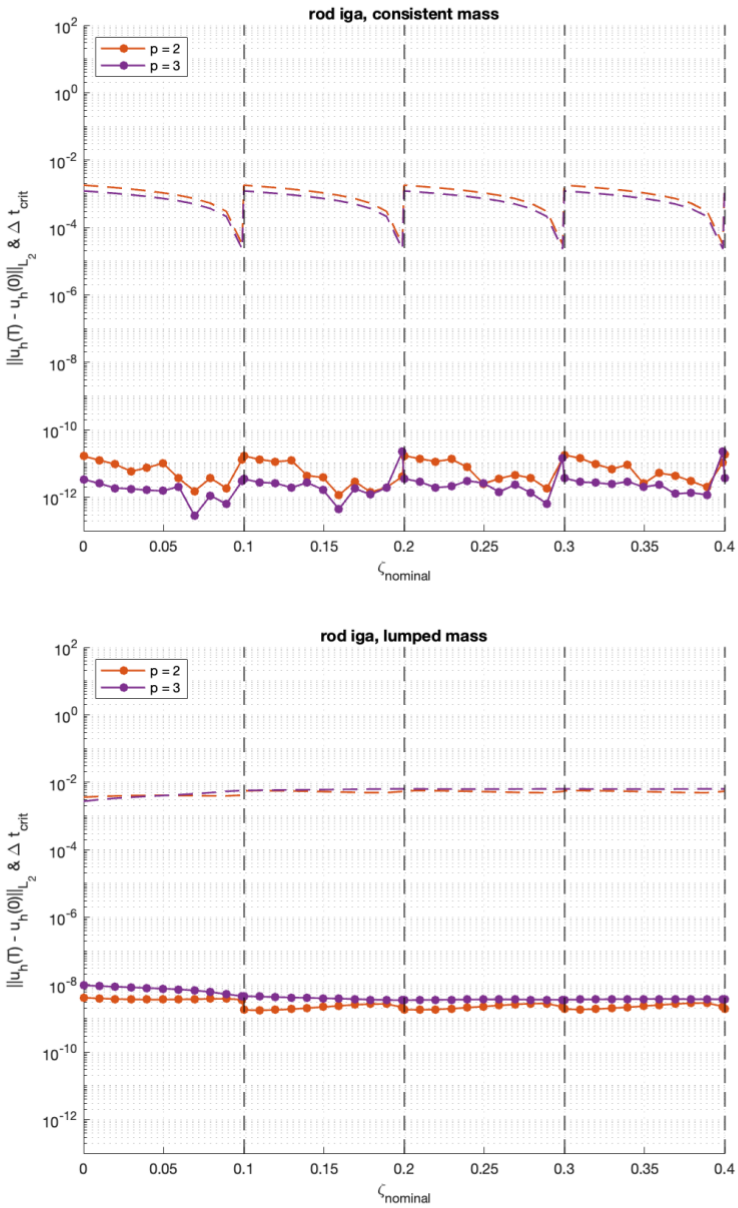

Figure 7: Immersed rod with the first continuum mode as initial condition (IGA, free–free, nel=240, p=2 and p=3). L_2 error after one period (solid) and Δt_{crit} (dashed) against $\zeta_{nominal}$, for (a) consistent and (b) lumped mass.

2.3 Bernoulli-Euler beam

The beam is excited with the same first-mode initial condition, which the spline space again resolves almost exactly (Fig. 8). With a consistent mass the error level is higher (but again cut-sensitive) compared to the rod, around the order 10^{-8} to 10^{-10}. This is consistent with the round-off accumulated over the very large number of steps that the stiff fourth-order operator requires (Section 1), not with a projection error. Lumping raises the measured error to a flat plateau of order 10^{-4}, several orders above that baseline, while bounding the critical time step. The first-mode lumping error is thus much larger for the beam ($\approx 10^{-4}$) than for the rod ($\approx 10^{-8}$).

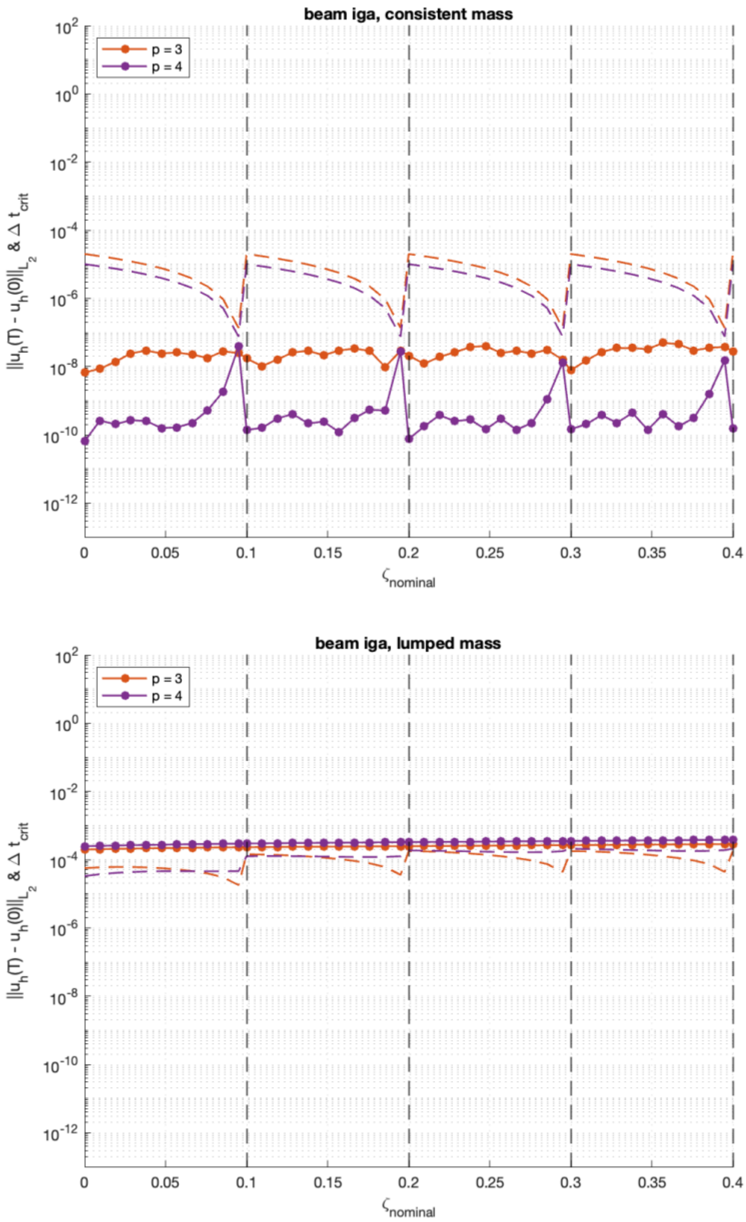

Figure 8: Immersed Bernoulli–Euler beam with the first non-rigid-body bending mode as initial condition (IGA, free–free, nel=60, p=3 and p=4). L_2 error after one period (solid) and Δt_{crit} (dashed) against $\zeta_{nominal}$, for (a) consistent and (b) lumped mass. Note the higher error floor than for the rod (Fig. 7).

3 EXCITATION DEPENDENCE OF THE ERROR

Section 2 noted that the cost of lumping a single low mode is far smaller than for a broadband pulse, and left the dependence to this section. The cost is not a fixed property of the mesh, it depends on what the initial condition excites. To make this systematic, the rod of Section 2 is driven by five initial conditions ranging from a single mode to a broadband pulse: the first, fifth and fifteenth modes, and two Gaussian pulses of different width. The mass treatment and stabilisation is varied as a sequence from the consistent mass, through α-stabilisation at several levels, to the lumped mass (Fig. 9a–f). Each panel carries the L_2 error after one period and the critical time step across the immersion cases, as in Section 2.

3.1 Consistent and lumped mass

With the consistent mass (Fig. 9a) the error already depends strongly on the excitation. The three single modes sit near their spatial-resolution thresholds and stay nearly flat across the several immersion states, rising with mode number from order 10^{-12} for the first mode to 10^{-7} for the fifteenth. The two pulses are several orders higher and follow the shrinking critical time step, as the single pulse of Section 2 did. This panel records how well the discretisation resolves each excitation before the mass is lumped or stabilisation is applied.

Lumping raises the error of every excitation (Fig. 9f), but by different relative amounts. In absolute terms the lumped error is largest for the broader-band pulse, of order 10^{-1}, and smallest for the first mode, of order 10^{-8}. Measured against the consistent baseline, the increase is largest for the single modes, about four orders of magnitude and of similar size for modes 1, 5 and 15, and smallest for the pulses, about one to two orders. To conclude which excitation is affected most by lumping therefore depends on whether the error is read in absolute or in relative terms. A single L_2 number does not separate the two contributions it contains: how difficult is to resolve the excitation and how much mass lumping is making it worse.

3.2 α-stabilization

Another approach that exists is the α-stabilisation, it keeps the consistent mass and adds artificial stiffness and mass in the fictitious domain. As mentioned already in TR1, it was introduced in [Ra24] as the intermediate solution between the consistent and lumped extremes, trading some accuracy for a bounded time step without lumping the mass. Its effect grows with the stabilisation intensity (Fig. 9b–e). A very mild intensity (Fig. 9b, $\alpha = 10^{-15}$) keeps the first mode's error below the lumped value, but the critical time step is barely changed from the unstabilised case (the dashed curve in Fig. 9b matches Fig. 9a): the stabilisation is too weak to bound the eigenfrequency, which is its purpose. Raising it until the time step is bounded (Fig. 9c, $\alpha = 10^{-11}$) affects the error levels of the single modes, the first mode already past the lumped value, while the pulses are unchanged, so the added error appears only where the consistent baseline error was low. A stronger stabilisation (Fig. 9d, $\alpha = 10^{-9}$) continues this trend, and at the strongest one (Fig. 9e, $\alpha = 10^{-6}$) makes the three single modes rise onto a common level while the pulses stay close to the consistent result.

The first mode is the limiting case. No level both keeps its error below the lumped value and bounds its time step: where the stabilisation is weak enough to stay below lumping it fails to recover the time step, and where it recovers the time step the first-mode error has already passed the lumped value. For that excitation, stabilisation offers no setting better than lumping. The distinction matters in practice. Lumping yields a diagonal mass that avoids the cost of inversion in each step, whereas stabilisation keeps the consistent mass and requires a linear solve for every step. That cost is justified only where the stabilised mass is clearly more accurate than lumping, and these results show that this may not be always the case.

As we can observe, across the five cases the error is related to which part of the spectrum is excited and by how much each mass treatment or stabilisation is altering that part. Neither the error nor a tolerable stabilisation level has identical effects from one excitation to another. Clarifying the dependence requires new metrics at the level of the spectrum: which modes a given initial condition excites, and how the mass treatment changes the modes themselves. These tools, modal participation and the modal assurance criterion, are introduced later in Section 5.

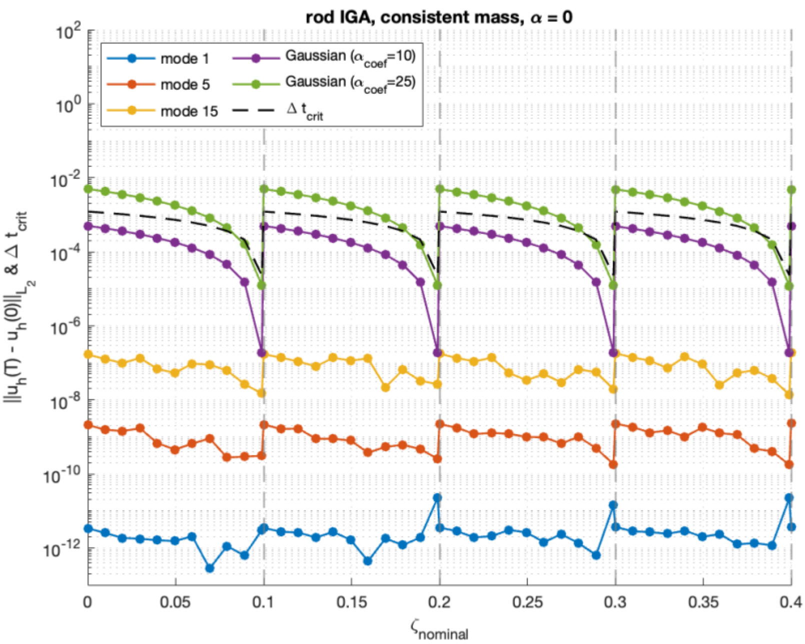

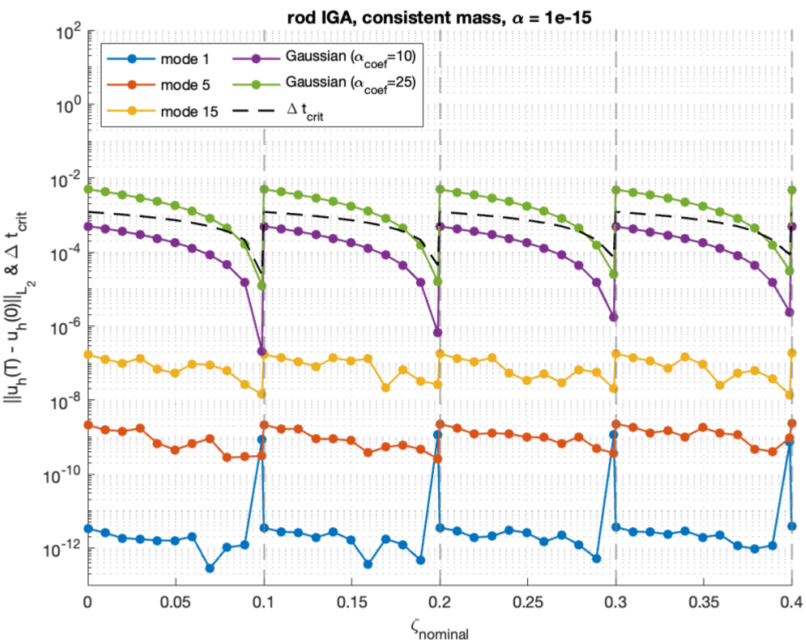

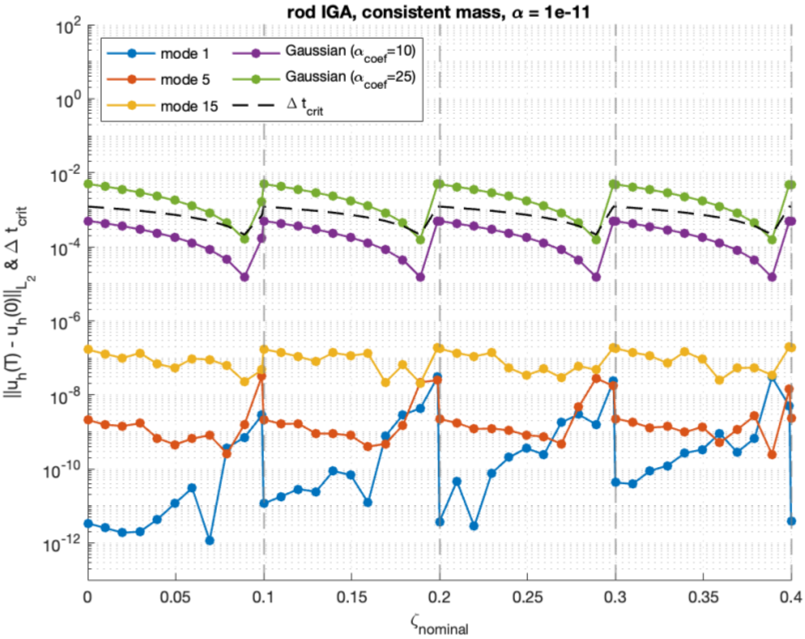
rod IGA, consistent mass, α = 1e-11
mode 1
mode 5
mode 15
Gaussian (α_{coef}=10)
Gaussian (α_{coef}=25)
Δt_{crit}
$||u_h(T) - u_h(0)||_{L_2}$ & Δt_{crit}
$\zeta_{nominal}$

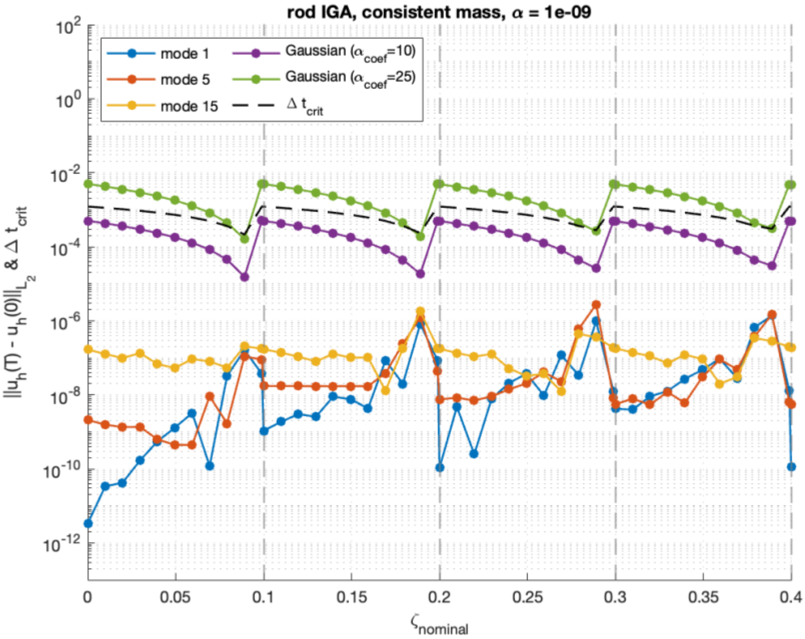
rod IGA, consistent mass, α = 1e-09
mode 1
mode 5
mode 15
Gaussian (α_{coef}=10)
Gaussian (α_{coef}=25)
Δt_{crit}
$||u_h(T) - u_h(0)||_{L_2}$ & Δt_{crit}
$\zeta_{nominal}$

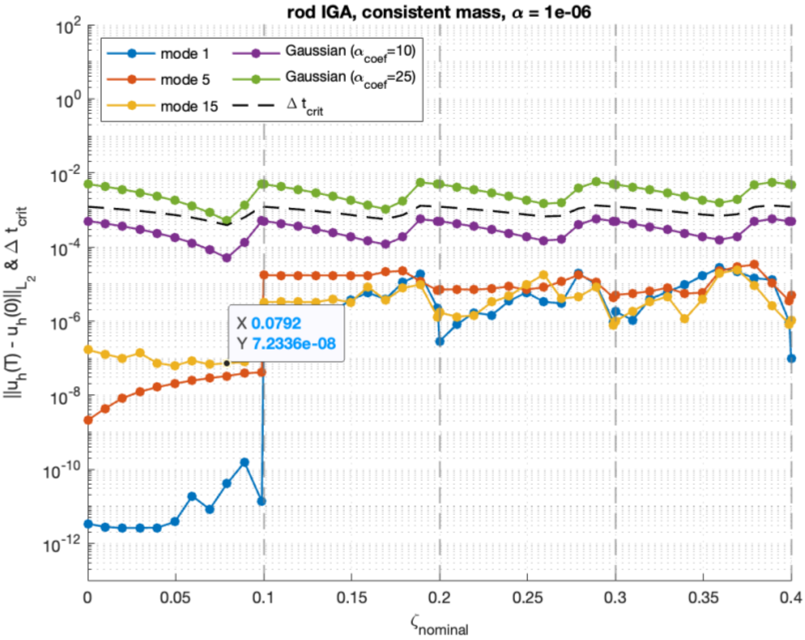

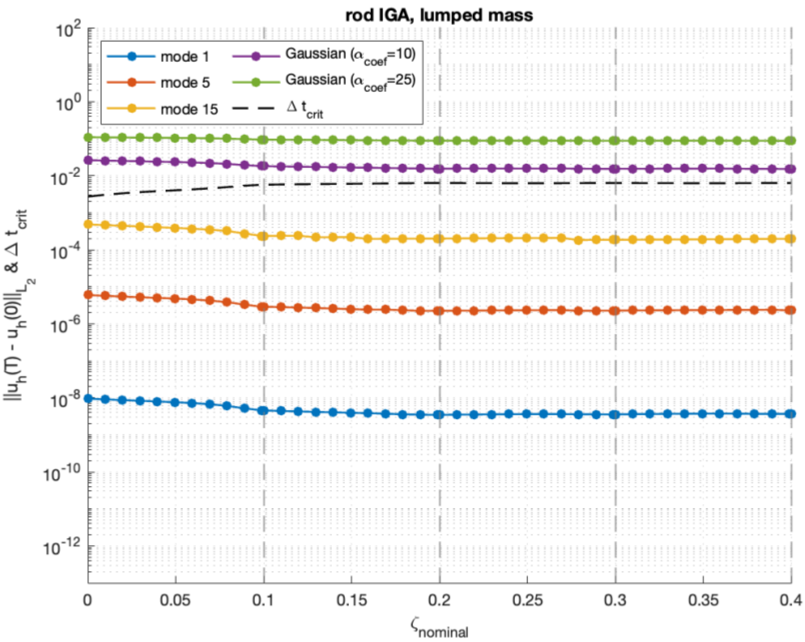

Figure 9: Excitation dependence for the immersed rod (IGA, free–free, nel=240, p=3) under five initial conditions: modes 1, 5 and 15, and two Gaussian pulses of width coefficient $\alpha_{coef} = 10$ (wider shape, narrower band) and $\alpha_{coef} = 25$ (narrower shape, broader band); α_{coef} sets the pulse sharpness and is unrelated to the stabilisation parameter. L_2 error after one period (solid, one colour per initial condition) and critical time step Δt_{crit} (dashed) against the fictitious-domain size $\zeta_{nominal}$. The mass treatment runs from (a) consistent, through α-stabilisation at (b) $\alpha = 10^{-15}$, (c) $\alpha = 10^{-11}$, (d) $\alpha = 10^{-9}$ and (e) $\alpha = 10^{-6}$, to (f) lumped

4 IMMERSION AND COMPLEX GEOMETRY

The one-dimensional studies above isolate the effect of the mass treatment on a fixed, simple geometry. The complementary path is geometric: how an immersed discretisation behaves as the cut geometry becomes more complex. This was the subject of the two-dimensional study introduced in the previous report (TR1, Section 2.3.2), in which an immersed square is integrated with a quadtree scheme that refines the cut cells to a fixed depth, here three levels. That setup is retained and not explained again, the question taken up here is what controls the spurious eigenfrequencies that small cut cells produce.

4.1 The controlled comparison: axis aligned vs. rotated immersed square

For the axis-aligned square the answer is simple. The cut cells form along the immersed boundary in a regular pattern, and a single variable, the cut size, predicts where the largest and most dangerous eigenfrequencies appear: the smaller the physical fraction of a cell, the larger its contribution to the spectrum. Rotating the square removes this regularity. The boundary now crosses the background grid at an angle, the quadtree produces cut cells of many shapes and neighbourhoods, and the cut size is no longer the only variable, the rotation angle and the quadtree depth are needed as well.

To test whether the cut size alone still governs the outliers, two configurations are compared that are identical in every controllable respect (the same square, polynomial degree, number of elements and quadtree depth) and that contain, by construction, a tiny cut cell of the same mass. The only difference is how that cut cell connects to its neighbours (Fig. 10). If the cut mass alone defines the spectrum, both would share the same largest eigenfrequency, but as we will observe, this is not the case.

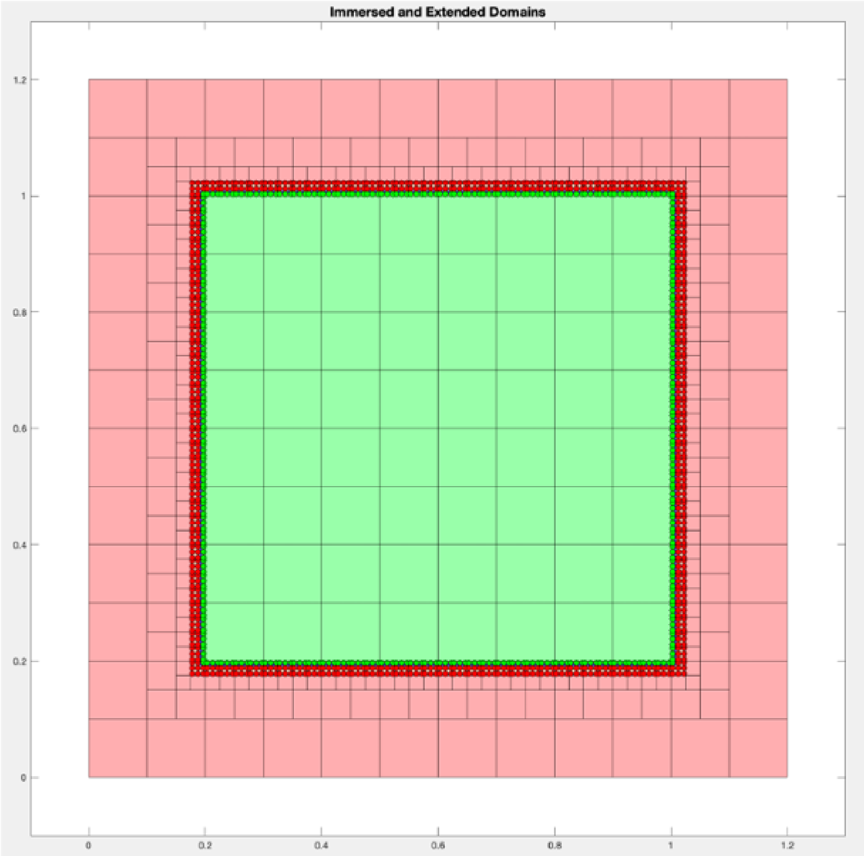
Immersed and Extended Domains
1.2
1
0.8
0.6
0.4
0.2
0
0
0.2
0.4
0.6
0.8
1
1.2

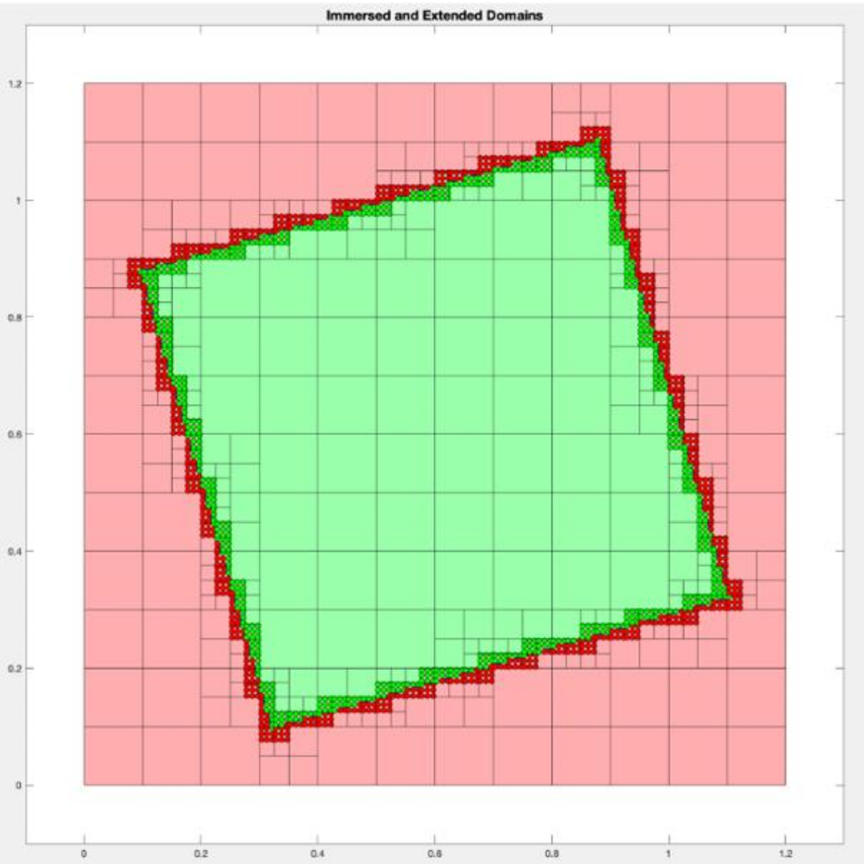
Immersed and Extended Domains
1.2
1
0.8
0.6
0.4
0.2
0
0
0.2
0.4
0.6
0.8
1
1.2

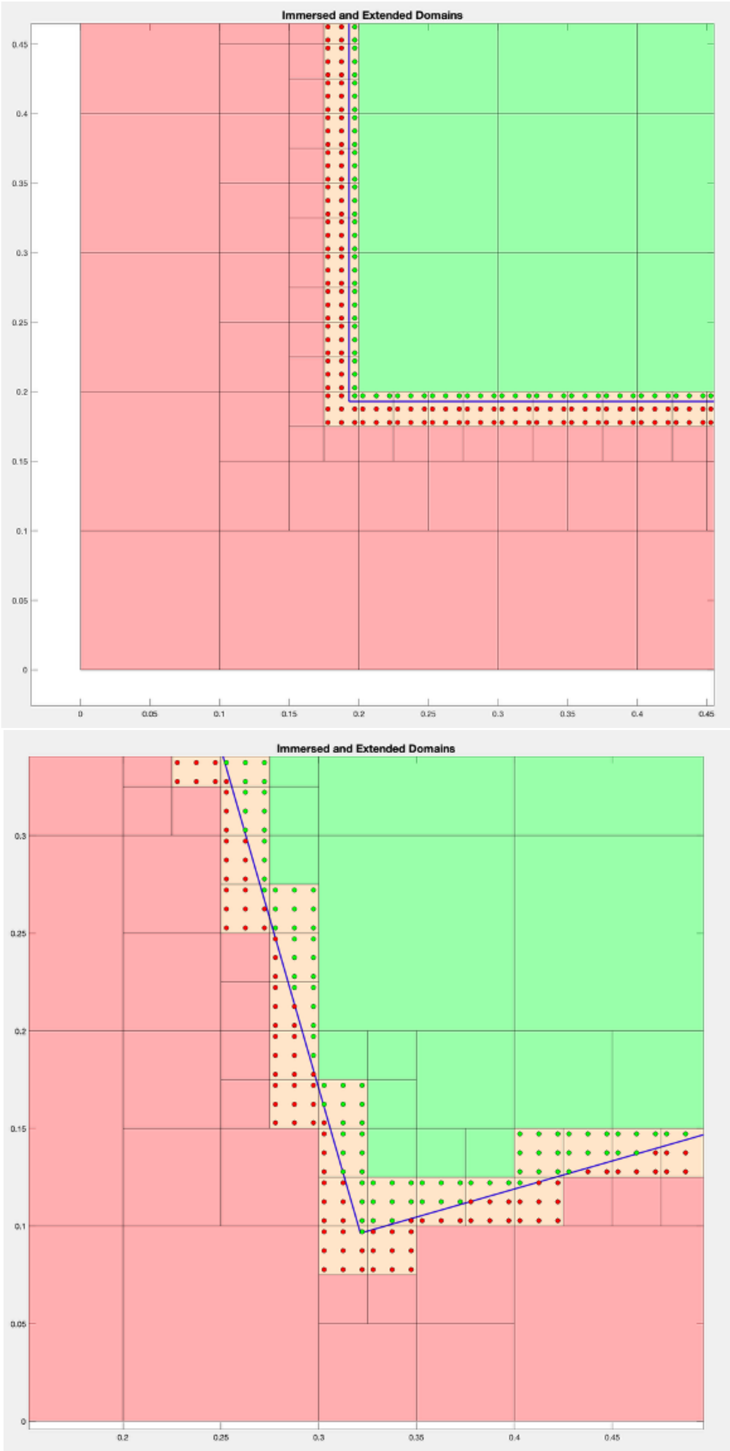

Figure 10: Controlled comparison of two immersed-square configurations integrated with a three-level quadtree, identical in square, polynomial degree, number of elements and quadtree depth, and with a tiny cut cell of identical mass by construction. (a, b) Axis-aligned square and a detail at a cut cell; (c, d) the same square rotated and a detail at a cut cell. The physical domain (green) is immersed in the extended domain (pink); cut cells along the boundary are refined by the quadtree, and quadrature points are classified as interior (green) or exterior (red). Largest discrete frequency and critical time step: aligned $\omega_{max} = 1.23\times10^{3}$, $\Delta t_{crit} = 1.463\times10^{-3}$; rotated $\omega_{max} = 8.63\times10^{9}$, $\Delta t_{crit} = 2.085\times10^{-10}$.

The aligned configuration has a largest discrete frequency of ω_{max} = 1.23×10³ and a critical time step of Δt_{crit} = 1.463×10⁻³; the rotated configuration, at the same cut mass, reaches ω_{max} = 8.63×10⁹ and Δt_{crit} = 2.085×10⁻¹⁰, six orders of magnitude apart. The cut mass is therefore not a reliable predictor of the dangerous outliers. What separates the two cases is the support of the basis functions over the cut cell and the mass of the cells they connect to. A criterion for locating problematic cells must account for this connectivity, not the cut mass alone.

4.2 Related research

The behaviour of these outlier frequencies has since been analysed in detail. Bioli and Voet [BV25] derive analytical estimates for the largest discrete frequency of lumped immersed discretisations across a range of trimming configurations, together with estimates for the smallest frequency, which has lately attracted closer scrutiny. Their analysis confirms that the outliers are governed by the trimming configuration rather than by a single geometric size.

In two dimensions the configurations multiply quickly, through cut size, rotation and quadtree depth, and a complex geometry can place many cells in unfavourable neighbourhoods at once. Mass lumping applied on top of this does not simply bound the spectrum. Voet et al. [VSB25] show that lumping can damage the accuracy of the low frequencies and modes and induce spurious oscillations, and proposed a stabilisation based on polynomial extensions that restores accuracy comparable to a boundary-fitted discretisation. Guarino et al. [Gu25] extended the same treatment to plate and shell problems. These treatments remove the damage that is specific to immersion. When the mass is still lumped, the boundary-fitted lumped accuracy remains, but Sections 2 and 3 show that this is already limited.

Immersion and mass lumping clearly have different roles. The first changes the cut geometry, the support of the basis functions and the connectivity of the cells. The second changes the mass operator used to advance the solution in time. Their effects could compound, but each is best understood alone before their combination is studied.

5 RECAP AND DIAGNOSTICS

The two studies above reach similar limitations for different reasons. The excitation study (Section 3) shows that the displacement error depends on what is excited by the initial condition. The immersion study (Section 4) shows that the dangerous eigenfrequencies are not predicted by the cut mass alone. These two point to different limitations appearing when some single scalar measures (L2 error, sliver cut mass) are not adequate to explain or predict what is needed.

Moving from reporting the symptom to explaining it requires looking inside the computation rather than at its final output. To assist with this task, two additional metrics are introduced here.

The computation can be viewed as a chain. First, the stiffness and mass operators are assembled, then their generalised eigenproblem gives the discrete spectrum, an initial condition selects which part of that spectrum is excited, and the time integrator advances the result. The displacement error is located at the end of this chain, but the two diagnostic metrics below inspect earlier links, the spectrum and the excitation, where the cause of the error can be seen before it reaches the output.

5.1 Modal Assurance Criterion (MAC)

The first metric demonstrates what mass lumping does to the modes themselves. The modal assurance criterion compares each consistent-mass mode shape with each lumped-mass mode shape: a bright diagonal means the lumped modes still match their consistent counterparts, and off-diagonal entries mean the shapes have changed or the modes have reordered. For a standard finite element discretisation the diagonal stays bright across the whole spectrum, so lumping leaves the mode shapes essentially in the same shape (Fig. 11b).

For the isogeometric discretisation the diagonal is bright over the lower part of the spectrum but breaks down over the upper part, where the lumped modes no longer correspond one to one with the consistent ones and begin to mix and reorder (Fig. 11a). Lumping therefore disturbs the isogeometric modes in a way it does not disturb the standard finite element modes, and it does so in the upper part of the spectrum.

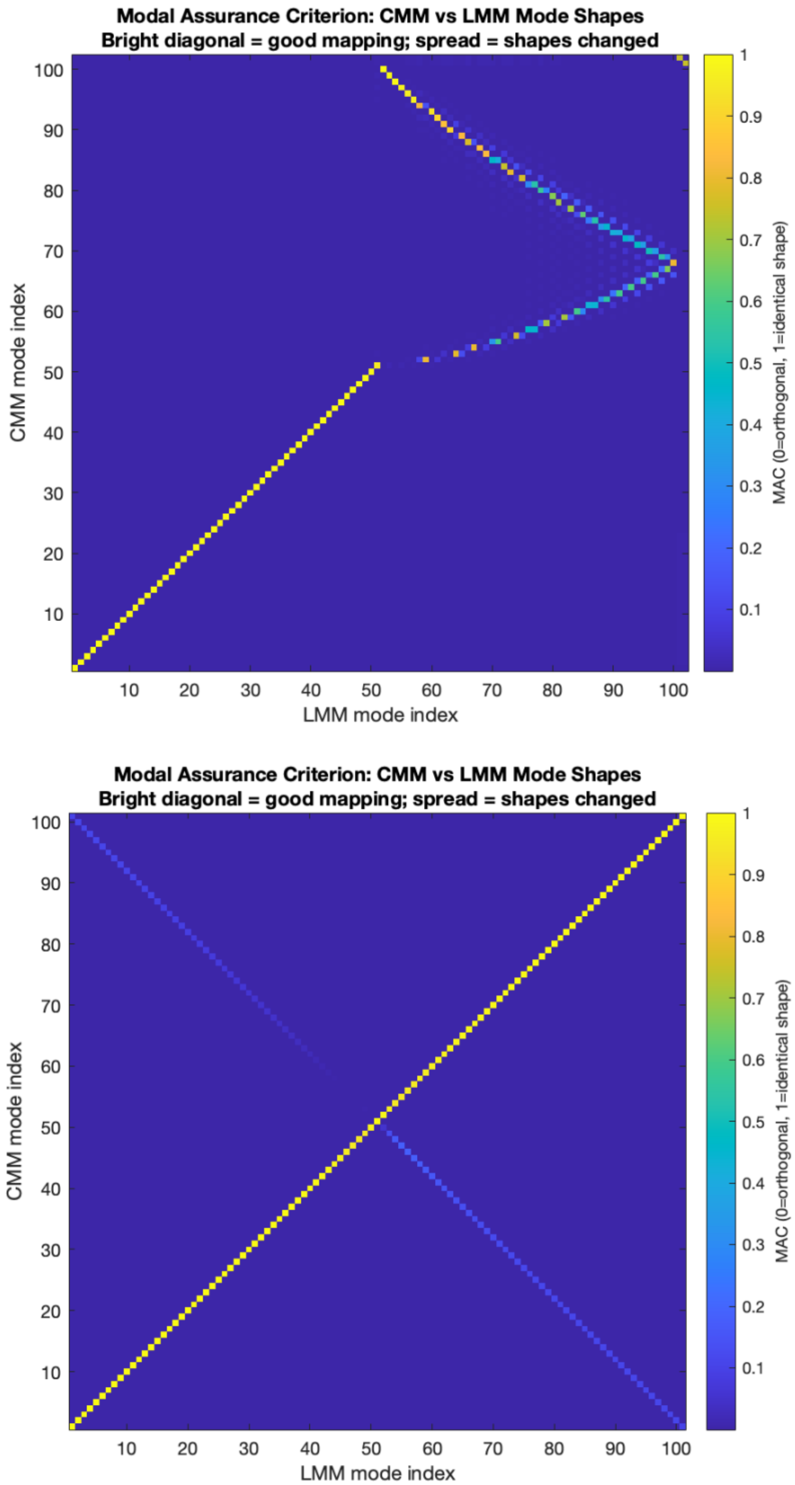

Figure 11: Modal assurance criterion between consistent-mass and lumped-mass mode shapes for the boundary-fitted rod, p=2 (consistent-mass index against lumped-mass index; 1 = identical shape, 0 = orthogonal). (a) IGA and (b) standard FE discretisation. The standard FE diagonal remains bright across the spectrum; the IGA diagonal is bright over the lower modes but does not hold over the upper modes, where the lumped modes reorder and mix.

5.2 Modal Participation

The second metric indicates which modes a given initial condition actually excites. The spectrum describes what the system can do, the initial condition decides which part of it is excited. Modal participation measures this directly, as the fraction of the initial energy carried by each discrete mode.

For the Gaussian pulse on the rod the energy is concentrated in a band of low-to-middle modes, rising to a peak and then decaying toward the higher modes (Fig. 12). On the coarser mesh the participation is still significant near the top of the spectrum, so the mesh does not fully resolve the pulse (Fig. 12a). On the finer mesh the same band is reproduced but now it is fully resolved and thus decays to zero well below the top of the spectrum, and the higher modes carry no energy (Fig. 12b). The content an initial condition excites is therefore a quantity that can be read directly, and refinement shows exactly how the discretisation resolves it.

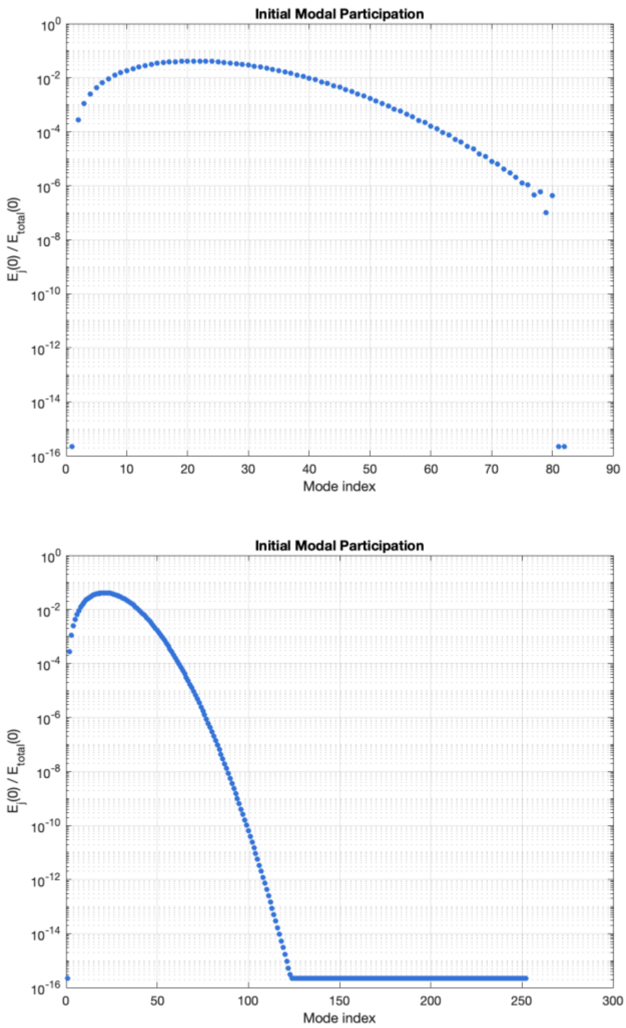

Figure 12: Initial modal participation for a Gaussian pulse on the boundary-fitted rod (IGA, consistent mass): the fraction of the initial energy $E_j(0)/E_{total}(0)$ carried by each discrete mode, for (a) number of elements = 80 and (b) number of elements = 250. The same low-to-middle band is excited in both; on the coarse mesh the participation remains noticeable near the top of the spectrum, while on the fine mesh it decays to zero well below it.

6 SELECTIVE MASS SCALING: FIRST STEPS

Another way to treat the mass operator is selective mass scaling. In addition to lumping the mass matrix, it modifies it further so that the highest eigenfrequencies are lowered, which increases the stable time step, while the lowest modes and the rigid-body motion are left as nearly unchanged as possible. The cost is that the modified mass is no longer diagonal, so each explicit step requires the solution of a linear system.

The formulation used here is Method II of Olovsson et al. [OSU05], which modifies each element mass matrix so that the translational rigid-body modes, and hence linear momentum, are preserved while the total mass is unchanged. The technique is established in explicit finite element practice and is available in commercial software such as LS-DYNA, where lumped mass is the explicit default and selective mass scaling is offered as an option that enlarges the time step at an added cost per step.

A more recent formulation preserves both linear and angular momentum while keeping the mass diagonal [Kr24].

The aim here is to present the first steps towards selective mass scaling for isogeometric discretisations.

6.1 Validation at p=1

Method II was implemented for the one-dimensional rod, with a single parameter β setting the amount of scaling: β = 0 leaves the mass unscaled, and larger β scales the upper modes more strongly. At the lowest polynomial degree, p = 1, where the b-spline and Lagrange bases coincide, the implementation reproduces the spectral effect of Method II. The frequency ratio relative to the lumped mass stays near one over the low modes and reaches a plateau over the upper modes, lower for larger β (Fig. 13), matching the behaviour reported for Method II [OSU05, Fig. 3]. A separate check confirmed that translational rigid-body motion is preserved during explicit integration after the selective mass scaling is applied.

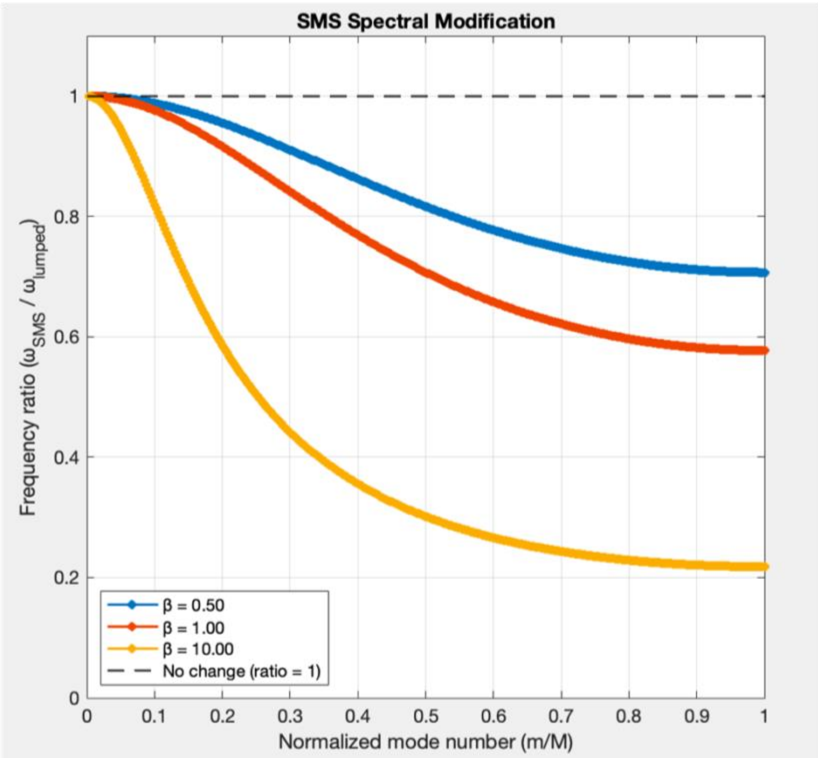

Figure 13: Spectral effect of Method II selective mass scaling on the rod at p = 1, where the b-spline and Lagrange bases coincide (401 DOF). Frequency ratio $\omega_{SMS}/\omega_{lumped}$ against normalised mode number for β = 0.5, 1 and 10. The low modes stay near the lumped values (ratio ≈ 1) and the upper modes are scaled down, more strongly for larger β.

6.2 Rod IGA

In this section the same approach is applied to the isogeometric rod at p = 2, the scaling lowers the upper part of the spectrum, including the outliers, while leaving the low modes in place (Fig. 14). The effect is of the same kind in the boundary-fitted (Fig. 14a) and immersed (Fig. 14b) cases.

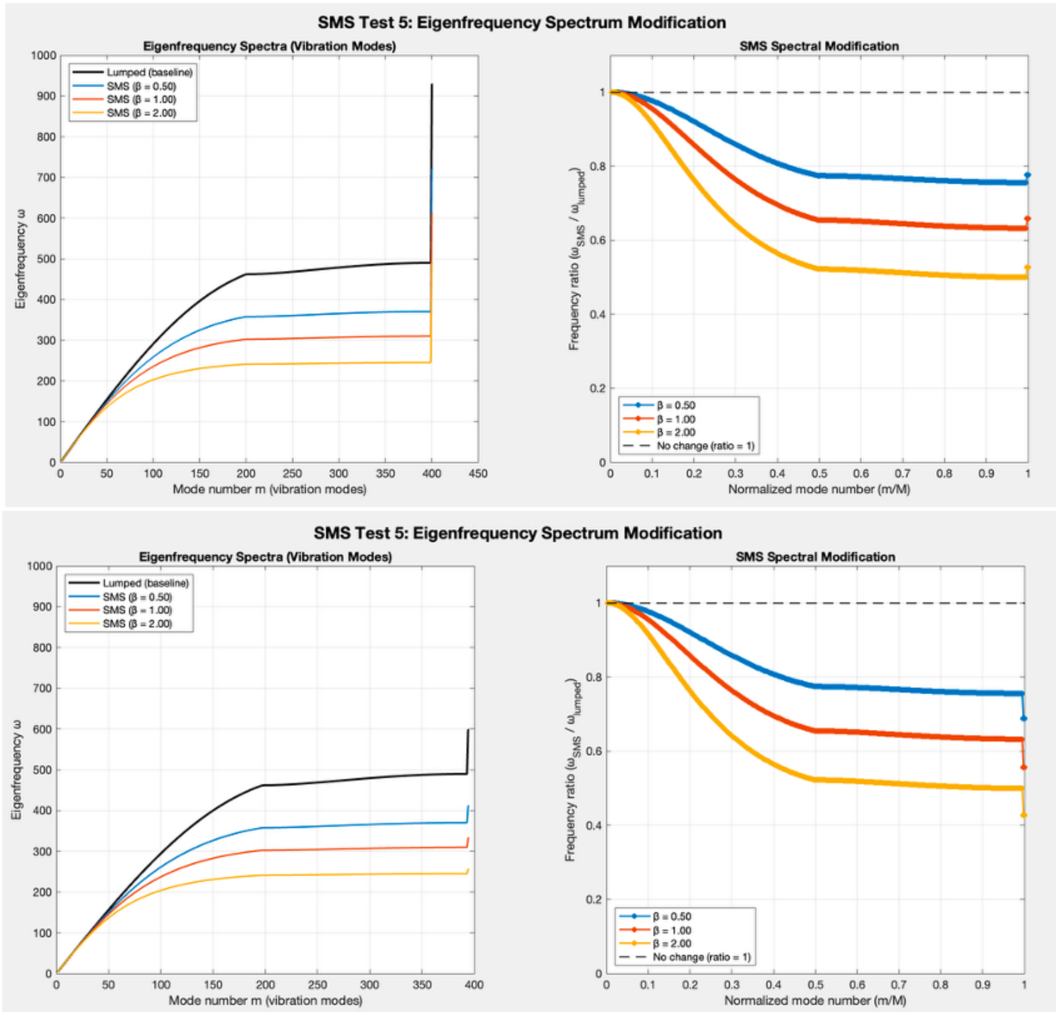

Figure 14: Selective mass scaling on the isogeometric rod (p = 2, 402 DOF), for (a) the boundary-fitted and (b) the immersed case. Each panel pair shows the eigenfrequency spectra (left) and the frequency ratio $\omega_{SMS}/\omega_{lumped}$ against normalised mode number (right), for β = 0.5, 1 and 2, with the row-sum lumped spectrum as baseline. In both cases the upper part of the spectrum, including the outliers of the lumped spectrum, is scaled down while the low modes are left in place.

These are just some preliminary results on this new topic. The idea of selective mass scaling does not imply that the time step could keep increasing without bound: beyond a certain level the highest frequency, and with it the critical time step, saturates, so further scaling adds cost through the per-step solve without enabling a larger time step. The outliers of the isogeometric spectrum enable this saturation earlier.

A time-domain accuracy study, and the balance between the larger time step and the per-step solver cost, are left for further work.

7 WORK IN PROGRESS AND FUTURE WORK

The new metrics introduced in Section 5 are the start of a broader strategy. The main idea is to follow the solution from the spectrum to the time-domain error, for the consistent and the lumped mass on a common base, and track with qualitative and quantitative metrics the accuracy decay. An earlier available (in the computational pipeline) insight could provide a deeper root-cause explanation of the observed phenomena.

Since explicit dynamics ultimately runs on an explicit solver, the value of any mass treatment is a balance of accuracy against cost and this could vary a lot depending on the application, the specific needs and details. A cost-versus-accuracy analysis is therefore an interesting study, together with convergence studies and related measures, that would turn the trade-offs described in this report into directly comparable statements across methods.

Regarding mass treatment/manipulations schemes, the selective mass scaling of Section 6 will be carried beyond its spectral first steps to a time-domain accuracy study and to the balance between the enlarged time step and the per-step solver cost. More broadly, the diagnostics motivate the development of mass-treatment strategies better matched to the isogeometric spectrum than simplistic row-sum lumping, which is an active direction of ongoing research.

8 CONCLUSIONS

This report established a verified one-dimensional framework in place and used it to set out the accuracy cost of mass lumping and its dependence on the excitation. Additionally, investigated the treatment of complex immersed geometry further, began to look beyond the symptoms of these effects toward their origin, and reported first steps toward selective mass scaling. In both settings a single scalar measure is inadequate, though for different reasons: the displacement error reports the loss of accuracy from mass lumping without being able to explain it, and the cut-element mass does not predict the most severe frequencies under immersion.

This sets the direction of the work that follows. Immersion and mass lumping alter the discretisation in different ways, each should be analysed independently before their combination is considered. At the same time, the aim is to obtain understanding at an earlier stage of the computational pipeline, closer to the formation of the operators, so that the root-cause of the observed effects can be identified before it reaches the time-domain solution.

9 REFERENCES

1. [Hu05] T.J.R. Hughes, J. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry, and mesh refinement. Computer Methods in Applied Mechanics and Engineering, 2005, 194, pp.4135–4195.
2. [Co06] J. Cottrell, R. Reali, Y. Bazilevs, T.J.R. Hughes. Isogeometric analysis of structural vibrations. Computer Methods in Applied Mechanics and Engineering, 2006, 195, pp.5257-5296.
3. [PDR07] J. Parvizian, A. Düster, E. Rank. Finite cell method: h- and p-extension for embedded domain problems in solid mechanics. Computational Mechanics, 2007, 41, pp. 121–133.
4. [Le20] Leidinger L., Explicit Isogeometric B-Rep Analysis for Nonlinear Dynamic Crash Simulations. PhD Dissertation, Technical University of Munich, 2020
5. [Ra24] L. Radtke, M. Torre, T. J. Hughes, A. Düster, G. Sangalli, and A. Reali, “An analysis of high order fem and iga for explicit dynamics: mass lumping and immersed boundaries,” International Journal for Numerical Methods in Engineering, vol. 125, no. 16, p. e7499, 2024
6. [BV25] I. Bioli, Y. Voet. A theoretical study on the effect of mass lumping on the discrete frequencies in immersogeometric analysis. Journal of Scientific Computing, 2025, 104, art. 103.
7. [VSB25] Y. Voet, E. Sande, A. Buffa. Mass lumping and stabilization for immersogeometric analysis. arXiv:2502.00452, 2025.
8. [Gu25] G. Guarino, Y. Voet, P. Antolin, A. Buffa. Stabilization techniques for immersogeometric analysis of plate and shell problems in explicit dynamics. Journal of Sound and Vibration, vol. 626, 2025.
9. [OSU05] L. Olovsson, K. Simonsson, M. Unosson. Selective mass scaling for explicit finite element analyses. International Journal for Numerical Methods in Engineering, 2005, 63(10), pp. 1436–1445.
10. [Kr24] L.-M. Krauß, R. Thierer, M. Bischoff, B. Oesterle. Intrinsically selective mass scaling with hierarchic plate formulations. Computer Methods in Applied Mechanics and Engineering, 2024, 432, Part B.