



Hierarchical NURBS in frictionless contact

İ. Temizer^{a,*}, C. Hesch^b

^a *Department of Mechanical Engineering, Bilkent University, 06800 Ankara, Turkey*

^b *Institute of Mechanics, Karlsruhe Institute of Technology, Otto-Ammann-Platz 9, Karlsruhe 76131, Germany*

Received 9 June 2015; received in revised form 24 October 2015; accepted 2 November 2015

Available online 11 November 2015

Highlights

- Novel isogeometric contact algorithms are introduced based on hierarchical NURBS.
- Hierarchical mortar discretizations are derived for frictionless contact.
- The partition of unity property is explored for the construction of the mortar basis.
- Exterior and interior point methods are applied in regularized and exact settings.
- Full or lumped recovery approach is employed to define the kinematic mortar variable.

Abstract

This work investigates mortar-based frictionless contact in the context of NURBS discretizations that are subjected to local hierarchical refinement. The investigations emphasize three sets of choices which lead to different contact algorithms that have distinct advantages and disadvantages. First, on the optimization side, both exterior and interior point methods are applied, thus spanning inexact constraint enforcement algorithms of the penalty or barrier type as well as exact ones of the primal–dual type. Second, on the discretization side, the hierarchical basis set of the mortar variables is inherited either directly from the discretization of the slave surface or after an intermediate normalization step to satisfy the partition of unity. Third, in interaction with both optimization and discretization, the kinematic mortar variable is recovered from the actual normal gap through the full or lumped solution of a linear system of equations. The implications of different choices are highlighted through benchmark problems which monitor the solution quality at the global level through the structural force evolution and at the local level through the contact pressure distribution.

© 2015 Elsevier B.V. All rights reserved.

Keywords: Contact; Isogeometric analysis; Hierarchical refinement; Mortar method

1. Introduction

Contact interfaces are often small in size in comparison to representative structural dimensions yet their modeling requires highly efficient and robust schemes for a reliable prediction of phenomena which are driven by the

* Corresponding author. Tel.: +90 312 290 3064.

E-mail address: temizer@bilkent.edu.tr (İ. Temizer).

contact conditions. The present work addresses the treatment of contact in the context of isogeometric analysis. Specifically, the aim is to explore the efficiency aspect by applying hierarchical discretization schemes that can concentrate the degrees of freedom within a zone of interest at the contact interface. Complementary investigations will additionally explore the robustness aspect by applying different optimization algorithms in the context of mortar-based discretizations.

Isogeometric analysis [1] has been established as a computational approach that has both enabling and facilitating features, for instance by allowing a direct transition from design to analysis and thereby circumventing the additional meshing stage. Contact mechanics particularly benefits from the intrinsic higher-order smoothness of isogeometric discretizations and from the underlying non-negative basis functions [2]. Almost all the studies on isogeometric contact to date have employed non-uniform rational B-splines (NURBS), mostly in a purely mechanical setting and often without friction. Among the first two studies, three-dimensional frictionless contact was studied in [3] using the penalty method together with a stable contact discretization that was developed for standard finite element discretizations [4,5]. In [6], the penalty method was used together with an intrinsically unstable Gauss-point-to-surface algorithm to investigate frictionless thermomechanical contact problems in a three-dimensional setting. Additionally, to recover stability, a NURBS-based discretization of the contact pressure was introduced in the spirit of the mortar method for two-dimensional linear elasticity. The two-dimensional extension of this mortar approach to incorporate friction was carried out using the penalty method in [7] and a three-dimensional extension was discussed in [8] where Uzawa augmentations were applied to improve the robustness of penalty regularization for both normal and tangential contact constraints. A two-dimensional mortar-based discretization scheme employing the Lagrange multiplier method, which ensures an exact enforcement of the contact constraints, was investigated in [9] in the absence friction and a three-dimensional Lagrange multiplier approach was presented in [10] within a frictional thermomechanical framework. An alternative approach to enforcing the constraints exactly is the augmented Lagrangian method [11] which was applied in a three-dimensional mortar-based frictionless setting in [12] and together with thermal contact constraints in [13]. While the formulation of [14,15] that was originally introduced for linear Lagrange discretizations is a common basis for mortar-based isogeometric contact discretizations, stable collocation-point-to-surface approaches have also been initiated in a two-dimensional frictionless setting [16].

From an optimization point of view, these studies exclusively employ exterior point methods such as the penalty, Lagrange multiplier and augmented Lagrangian methods. Recently, interior point methods [17] have also been applied to two- and three-dimensional mortar-based frictionless contact [18]. On the other hand, from a discretization point of view, all of these approaches are based on a Galerkin setting which introduces a weak form. An approach which takes the strong form as the starting point has also been introduced in a two-dimensional frictionless setting [19] based on isogeometric collocation [20].

The list of works reviewed above is not complete. In particular, applications to interface problems which benefit from similar algorithmic procedures were not mentioned, notably domain decomposition [21,22] and cohesive zones [23,24]. Similarly, no attempt is made to review the vast literature on which these studies are based with respect to the formulation and enforcement of the contact constraints as well as with respect to the description and discretization of the contact interface. The reader is referred to the cited works for specific references as well as to [25–27] for broad overviews.

The standard NURBS discretizations that have predominantly been employed in the reviewed studies only allow a global refinement of the basis structure. Among various possibilities, two widely adopted approaches enable local refinement in isogeometric analysis. One of these is T-splines [28,29]—see [30] for recent developments together with a refinement algorithm. In contact mechanics, T-splines have so far been applied only in a frictionless setting using the penalty method together within a Gauss-point-to-surface framework [31]. A major advantage with the T-spline approach is its ability to address complex topologies. Originally, T-splines were defined for bicubic surfaces [32]. However, the difficulty of constructing complex three-dimensional T-spline discretizations [33] with arbitrary orders [29] have recently been addressed. Another local refinement scheme is based on hierarchical refinement, starting with the approach of [34] that is applicable to both B-splines and NURBS, which has not been applied to contact problems so far and thus constitutes the focus of the present study. In [35], the approach in [34] has been extended to an adaptive framework which additionally ensures that the hierarchical basis satisfies the partition of unity. While these hierarchical refinement schemes can easily adopt arbitrary discretization orders, they are typically limited to a quadrilateral or hexahedral topology due to the underlying tensor-product structure. Recently, it has also become possible to overcome this restriction [36]. In some sense, ideas from T-splines and hierarchical representations may also

be combined: B-splines may be refined in a non-hierarchical fashion [37] while hierarchical refinement schemes may be extended to T-splines [30]. In this work, the NURBS-based hierarchical refinement scheme presented in [34] will be adopted, particularly in view of its algorithmic simplicity. Instead of addressing the lack of partition of unity property as in [35], its implications will first be discussed and then a straightforward normalization procedure is adopted to recover this property—see also [38] for an alternative scheme that satisfies the partition of unity by construction. Despite the special choice of the hierarchical refinement scheme, the addressed issues span a broad range of numerical concerns such that the developed contact algorithms would be beneficial for alternative refinement schemes as well.

The outline of this work is as follows. In Section 2, the hierarchical refinement scheme is described and demonstrated for bulk discretizations that automatically induce hierarchical discretizations on potential contact surfaces. The contact algorithm to be employed on such hierarchically refined surfaces is discussed in Section 3, with a particular focus on the choice of the mortar basis as well as its use in describing and enforcing the contact constraints. Extensive numerical investigations will be presented in Section 4 which range from classical benchmark examples with stationary contact interfaces, aiming to demonstrate fundamental features of various algorithmic choices, to more advanced tests with large sliding contact that focus on robustness and efficiency. Suggestions for possible future extensions and applications will be summarized in Section 5.

2. Hierarchical refinement

The process of hierarchical refinement adopted in this work has three major stages: (i) the construction of the initial coarse level geometry description, (ii) the generation of a series of fine level descriptions, and (iii) the selection of basis functions from each level. The first two stages are standard procedures in isogeometric analysis and therefore are only briefly reviewed. The last stage which generates the hierarchical basis set is addressed in detail.

2.1. Reference geometry

In order to define the univariate B-spline functions associated with the i th dimension of a patch, an open knot vector

$$\Xi^{i,0} = \{\xi_0^{i,0}, \dots, \xi_{p_i}^{i,0}, \xi_{p_i+1}^{i,0}, \dots, \xi_{n_i^0}^{i,0}, \xi_{n_i^0+1}^{i,0}, \dots, \xi_{n_i^0+p_i+1}^{i,0}\} \tag{2.1}$$

is assumed where p_i is the polynomial order of the functions and $n_i^0 + 1$ would indicate the number of control points in a one-dimensional setting. In a multi-dimensional setting, let \mathbf{X} indicate the position vector associated with the referential geometry that is described using n_X^0 control points \mathbf{X}_0^I with corresponding weights α_0^I . If the product of the univariate B-spline functions associated with a control point is denoted by $B_0^I(\xi^1, \xi^2, \xi^3)$, where ξ^i span the knot vectors, then the NURBS basis functions are defined by

$$N_0^I = \frac{\alpha_0^I B_0^I}{\sum_{J=1}^{n_X^0} \alpha_0^J B_0^J} \tag{2.2}$$

such that the referential geometry admits the *coarse level* description

$$\mathbf{X} = \sum_{I=1}^{n_X^0} N_0^I \mathbf{X}_0^I. \tag{2.3}$$

2.2. Fine level descriptions

Starting with the coarse level description, a hierarchy of *fine level* descriptions may be built as follows. Let

$$\bar{\Xi}^{i,0} = \{\bar{\xi}_0^{i,0}, \bar{\xi}_1^{i,0}, \dots, \bar{\xi}_{m_i^0-1}^{i,0}, \bar{\xi}_{m_i^0}^{i,0}\} \tag{2.4}$$

denote a modified vector which involves only unique knot entries. Here, m_i^0 indicates the number of unique knot spans. In a non-hierarchical setting, these would help construct the isogeometric elements that serve as convenient

domains of integration for evaluating the weak form, their number equal to the product $m^0 = \Pi_i m_i^0$. The next level in the hierarchy is constructed through knot refinement, specifically by halving the unique knot spans:

$$\bar{\xi}^{i,1} = \left\{ \bar{\xi}_0^{i,1} = \bar{\xi}_0^{i,0}, \bar{\xi}_1^{i,1} = \frac{1}{2}(\bar{\xi}_0^{i,0} + \bar{\xi}_1^{i,0}), \dots, \bar{\xi}_{m_i^1-1}^{i,1} = \frac{1}{2}(\bar{\xi}_{m_i^0-1}^{i,0} + \bar{\xi}_{m_i^0}^{i,0}), \bar{\xi}_{m_i^1}^{i,1} = \bar{\xi}_{m_i^0}^{i,0} \right\}. \quad (2.5)$$

This leads to a new set of weights α_1^J and control points X_1^J , computed using standard knot refinement algorithms [39], which are associated with $m^1 = \Pi_i m_i^1$ elements where $m_i^1 = 2m_i^0$. Consequently, the geometry is preserved but it admits the finer discretization $X = \sum_{I=1}^{n_x^1} N_1^I X_1^I$. By repeating this process through \mathcal{L} levels one obtains an embedded hierarchy of levels, each level $\ell \in [1, \mathcal{L}]$ offering a larger number of degrees of freedom to describe the same reference geometry:

$$X = \sum_{I=1}^{n_x^\ell} N_\ell^I X_\ell^I. \quad (2.6)$$

2.3. Hierarchical basis functions

Hierarchical refinement will be triggered element-wise [34]. At level ℓ , let S_ℓ^I denote the span of a basis function N_ℓ^I and $\mathcal{R}_\ell \subseteq [0, 1] \times [0, 1] \times [0, 1]$ denote the span of all elements which are marked for hierarchical refinement such that $\mathcal{R}_{\ell+1} \subseteq \mathcal{R}_\ell$. The algorithm for selecting basis functions from each level in order to create a hierarchical set of basis functions is demonstrated in Fig. 1 and can be summarized as follows:

- STEP 1 (*Coarse Level*): All coarse level basis functions are activated initially. Set $\ell = 0$.
- STEP 2 (*Basis Deactivation*): Deactivate all N_ℓ^I with $S_\ell^I \subseteq \mathcal{R}_\ell$.
- STEP 3 (*Basis Activation*): Activate all $N_{\ell+1}^I$ with $S_{\ell+1}^I \subseteq \mathcal{R}_\ell$.
- STEP 4 (*Level Loop*): Set $\ell = \ell + 1$ and continue with STEP 2 until all levels are checked.

The remaining n_u active basis functions will be indicated by N^I . These remaining functions are linearly independent and cover the whole patch while concentrating the degrees of freedom around the selected regions of refinement. Denoting the displacement vector by \mathbf{u} such that the current position vector is $\mathbf{x} = X + \mathbf{u}$, a subparametric hierarchical discretization can now be invoked [38]:

$$\mathbf{u} = \sum_{I=1}^{n_u} N^I \mathbf{u}^I. \quad (2.7)$$

The geometry is still described (exactly) by the initial coarse level discretization (2.3). Note that this initial discretization, and hence the ensuing hierarchical description, can employ different orders along each parametric direction.

The demonstration in Fig. 1 incorporates major aspects of the algorithm. Tracing the steps for the particular case of \mathcal{N}^1 as an example, initially all basis functions at level 0 are active. Subsequently, two of these are removed since they entirely lie in the zone of refinement towards level 1. Four basis functions from level 1 which lie within this zone are then activated. However, one is subsequently removed since it lies in the zone of refinement towards level 2. Within this zone, three basis functions from level 2 are activated, overall leading to a hierarchical discretization with eight basis functions. For the cases of \mathcal{N}^2 and \mathcal{N}^3 , one obtains nine basis functions. It is noted that the knot structure of the hierarchical basis in Fig. 1 as well as in similar subsequent figures solely serves a graphical purpose as an indication of the refinement process and does not directly induce the corresponding basis. This is evident for the case of \mathcal{N}^3 where there are nine basis functions although the knot structure can accommodate ten. In two and three dimensions, since there is no tensor product structure in a hierarchical setting, a maximum number of basis functions that can be accommodated cannot be deduced from the knot structure.

Due to the increasing span of the basis functions, the number of basis functions activated at each refinement step has a decreasing tendency when the order increases (see level 2 in Fig. 1 from \mathcal{N}^1 to \mathcal{N}^3). The number of deactivated basis functions also decreases. In fact, the refinement of an element at a given level may not lead to the deactivation of basis functions at all (e.g. level 1 to level 2 in Fig. 1 for \mathcal{N}^2 and \mathcal{N}^3). However, overall, starting with the same number of elements at each order, the percentage increase in the number of basis functions in the transition from

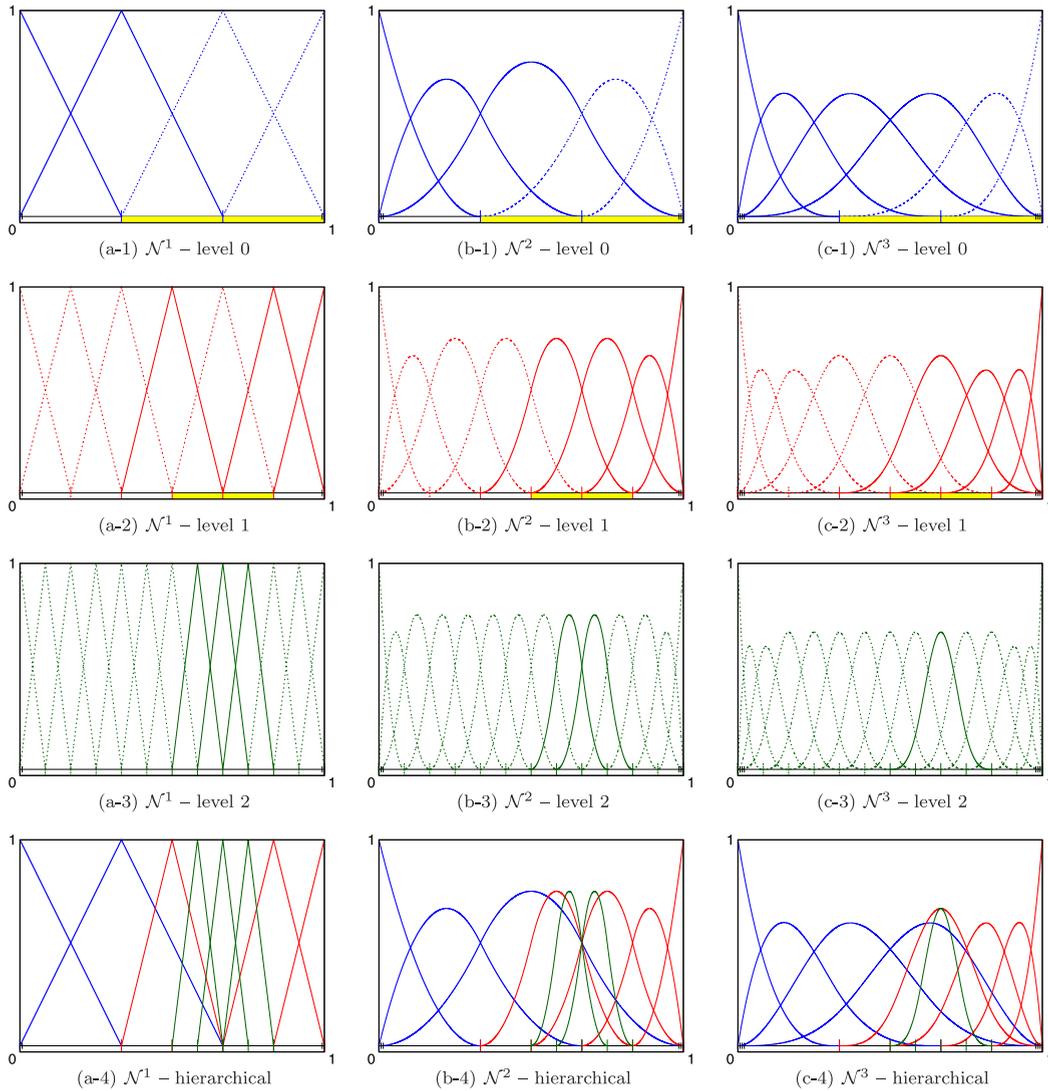


Fig. 1. The hierarchical refinement algorithm is demonstrated. The refined elements are marked with yellow and the inactive basis functions at each level are indicated with dashed lines. Here and in upcoming discussions, the notation \mathcal{N}^p refers to NURBS basis functions which emanate from B-splines of order p . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the coarse level basis set to the hierarchical one tends to decrease with increasing order. Consequently, in order to obtain a comparable increase in the numerical resolution, additional elements should be refined around the elements which indicate need for refinement. In general, the decision for hierarchical refinement may be driven through an error measure [37,36]. However, such adaptive schemes are outside the scope of the present work. Hence, in the numerical investigations, only pre-selected refinements will be employed in order to concentrate on the investigation of contact algorithms in a hierarchical setting.

The summarized hierarchical refinement strategy is reminiscent of the hanging nodes approach when first order elements are used, as depicted in Fig. 2. Indeed, there are qualitative similarities, e.g. the integration of the finite element weak form will be carried out using the hierarchical element set. Referring to Fig. 2, these are the four level 2 elements, followed by the three non-refined level 1 elements and the three level 0 elements. However, an important difference of the present hierarchical refinement algorithm from the classical hanging node approach is that the basis function activation/deactivation is not carried out locally within each element but rather globally across the whole mesh. For instance, in Fig. 2, the shape function of the central node from level 0 is active not only within the non-refined level 0 elements but across the whole domain. As a consequence, although the classical hanging nodes

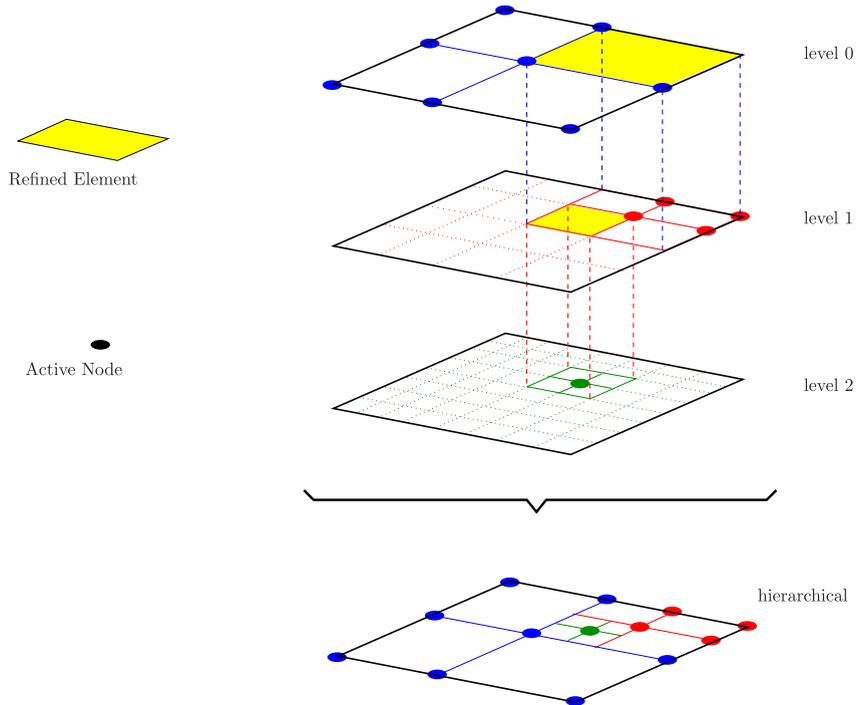


Fig. 2. The hierarchical refinement strategy is depicted with first order (Lagrange or NURBS) elements. The global shape functions which are active are indicated by their corresponding nodes. The resulting hierarchical set of elements contains four level 2, three level 1 and three level 0 elements.

approach leads to a basis which satisfies the partition of unity the present hierarchical refinement algorithm does not—see Section 3.2.

A numerical example for bulk discretization refinement at different orders is shown in Fig. 3. The hierarchical discretizations of the surfaces are directly inherited from these bulk discretizations. Indeed, the refinements demonstrated in Fig. 1 correspond to the lower surfaces in Fig. 3. The contact algorithms to be investigated next will be built on this observation. Note that only \mathbf{u}^I , not \mathbf{X}^I , is meaningful for a hierarchical basis function N^I . Nevertheless, for graphical purposes, if N^I is associated with a fine level basis function N_ℓ^J then a control point has been placed at \mathbf{X}_ℓ^J in Fig. 3. Consequently, some points are overlapping in the figure, for instance on the lower surface of \mathcal{N}^3 for $\mathcal{L} = 2$.

A hierarchical scheme may also be employed during the design process in order to endow an initially coarse geometry description with a number of fine features. In such a setting, one would first capture the base geometry exactly, excluding fine features such as a localized surface texture or a sharp protrusion. Subsequently, local hierarchical refinement would be employed to capture these features without altering the base geometry [40]. Once this is done, either an isoparametric formulation can be followed or further refinement may be carried out to improve the solution quality locally towards a subparametric formulation. In any case, in the context of isogeometric analysis, one starts from the premise that the reference geometry is exactly represented. As in standard finite element analysis, using an inexact geometry will lead to wrong stresses, and in particular to wrong contact tractions, unless refinement also updates the reference geometry.

3. Contact algorithm

3.1. Frictionless contact

The weak form of the linear momentum balance in a quasistatic setting without a body force for the frictionless contact of a slave body $\mathcal{D}^{(1)}$ and a master body $\mathcal{D}^{(2)}$ with spatial position vectors $\mathbf{x}^{(\alpha)}$ is symbolically denoted by

$$0 = \delta\mathcal{G} = \delta\mathcal{G}^o + \delta\mathcal{G}^c \tag{3.1}$$

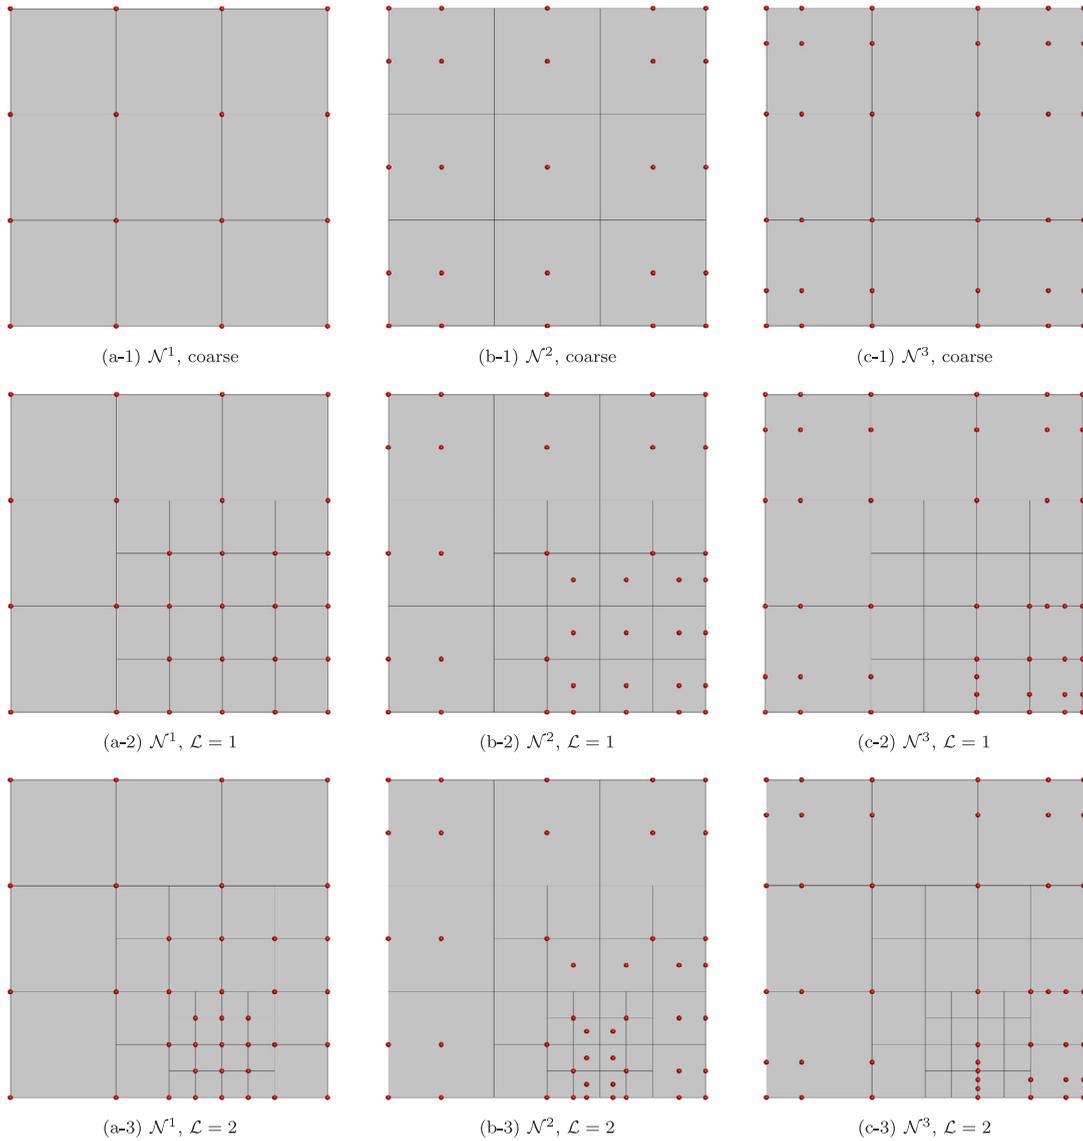


Fig. 3. Hierarchical refinement is demonstrated using \mathcal{N}^p ($p = 1, 2$ or 3) basis functions along each parametric direction, presently with unit weights (B-splines). The control points at each level are associated with the active basis functions.

where the bulk contribution is

$$\delta\mathcal{G}^o = \sum_{\alpha=1}^2 \left(\int_{\mathcal{R}_o^{(\alpha)}} \delta\mathbf{F}^{(\alpha)} \cdot \mathbf{P}^{(\alpha)} \, dV + \int_{\Gamma_o^{(\alpha)}} \delta\mathbf{x}^{(\alpha)} \cdot \widehat{\mathbf{p}}^{(\alpha)} \, dA \right) \tag{3.2}$$

and, introducing the notation $\langle \bullet \rangle = \int_{\Gamma_c^o} \bullet \, dA$, the contact contribution is

$$\delta\mathcal{G}^c = -\langle \delta g \, p \rangle. \tag{3.3}$$

In (3.2), $\mathbf{F}^{(\alpha)}$ and $\mathbf{P}^{(\alpha)}$ are the deformation gradient and the first Piola–Kirchhoff stress tensors, respectively, while $\mathcal{R}_o^{(\alpha)}$ denotes the reference configuration and $\Gamma_o^{(\alpha)} \subset \partial\mathcal{R}_o^{(\alpha)}$ denotes the portion of its boundary $\partial\mathcal{R}_o^{(\alpha)}$ with outward unit normal $\mathbf{N}^{(\alpha)}$ where the Piola traction $\mathbf{p}^{(\alpha)} = \mathbf{P}^{(\alpha)}\mathbf{N}^{(\alpha)}$ is prescribed to $\widehat{\mathbf{p}}^{(\alpha)}$. Within (3.3), $\Gamma_c^o \subset \partial\mathcal{R}_o^{(1)}$ indicates the portion of the referential slave boundary that encompasses the potential contact zone inside which the referential contact pressure p and the normal gap g should ideally satisfy $p > 0$ and $g = 0$ while outside this zone $p = 0$ and

$g < 0$. For compactness, the notation for the contact contribution is specialized such that \mathbf{x} will exclusively refer to a point on the spatial counterpart Γ^c of Γ_o^c and \mathbf{y} will indicate its closest-point projection on the master where the spatial outward unit normal is \mathbf{v} . Hence, on Γ_o^c , $\mathbf{p}^{(1)} = p \mathbf{v}$ in the absence of friction and

$$g = -(\mathbf{x} - \mathbf{y}) \cdot \mathbf{v}. \quad (3.4)$$

Denoting the referential counterparts of the position vectors \mathbf{x} and \mathbf{y} with \mathbf{X} and \mathbf{Y} , respectively, the displacement vectors

$$\mathbf{u} = \mathbf{x} - \mathbf{X}, \quad \mathbf{v} = \mathbf{y} - \mathbf{Y} \quad (3.5)$$

are introduced. The slave and master surfaces inherit their subparametric spatial surface discretizations, generally non-matching, directly from the hierarchical discretization of the volume. The coarse level shape functions of the slave/master volume will be denoted by N_0^K/M_0^L and the hierarchical ones by N^I/M^J . Without distinguishing the surface basis functions with a separate notation, the referential position vectors referring to points on the surfaces are assigned the coarse level discretizations

$$\mathbf{X} = \sum_{K=1}^{n_X^0} N_0^K \mathbf{X}_0^K, \quad \mathbf{Y} = \sum_{L=1}^{n_Y^0} M_0^L \mathbf{Y}_0^L \quad (3.6)$$

and the displacement vectors are assigned the hierarchical discretizations

$$\mathbf{u} = \sum_{I=1}^{n_u} N^I \mathbf{u}^I, \quad \mathbf{v} = \sum_{J=1}^{n_v} M^J \mathbf{v}^J. \quad (3.7)$$

The summation limits will be omitted whenever the context makes these clear. Closest-point projection onto a hierarchically refined surface for the calculation of g does not pose a significant algorithmic difficulty and hence will not be discussed.

3.2. Contact discretization

An additional (kinematic) contact variable γ will be introduced and, together with the (kinetic) variable p , will be associated with the hierarchical discretization of the slave surface in the spirit of the mortar method. Specifically, the number of mortar basis functions B^I is chosen to be n_u and the span of each matches N^I although the functions themselves are not necessarily the same:

$$\gamma = \sum_{I=1}^{n_u} B^I \gamma^I, \quad p = \sum_{I=1}^{n_u} B^I p^I. \quad (3.8)$$

The variable γ is weakly linked to the physical gap via the condition $\langle B^I \gamma \rangle = \langle B^I g \rangle$ which, upon defining the projected gap $\bar{g}^I = \langle B^I g \rangle$ and arranging the discrete values in vectors $\{\gamma\}$ and $\{\bar{g}\}$, leads to the formulation that recovers the kinematic contact variable from the normal gap:

$$\{\gamma\} = [\Phi]^{-1} \{\bar{g}\} \quad (\text{full recovery}). \quad (3.9)$$

Here, the components of the overlap matrix $[\Phi]$ are $\Phi^{IJ} = \langle B^I B^J \rangle$. The numerical cost of solving this linear system of equations is small in comparison to the overall cost of the problem. Nevertheless, if row-sum lumping is applied to $[\Phi]$ to obtain a diagonal matrix then the solution of a linear system of equations is not necessary:

$$w^I = \langle B^I \sum_J B^J \rangle \quad \longrightarrow \quad \gamma^I = \bar{g}^I / w^I \quad (\text{lumped recovery}). \quad (3.10)$$

Moreover, if the mortar basis satisfies the partition of unity then $w^I = \langle B^I \rangle$ such that one obtains the expression $\gamma^I = \langle B^I g \rangle / \langle B^I \rangle$ that commonly appears in mortar-based contact algorithms. Both the full and the lumped recovery approaches for γ^I will be tested. Note that the only mortar integrals that are necessary to the construction of the weak

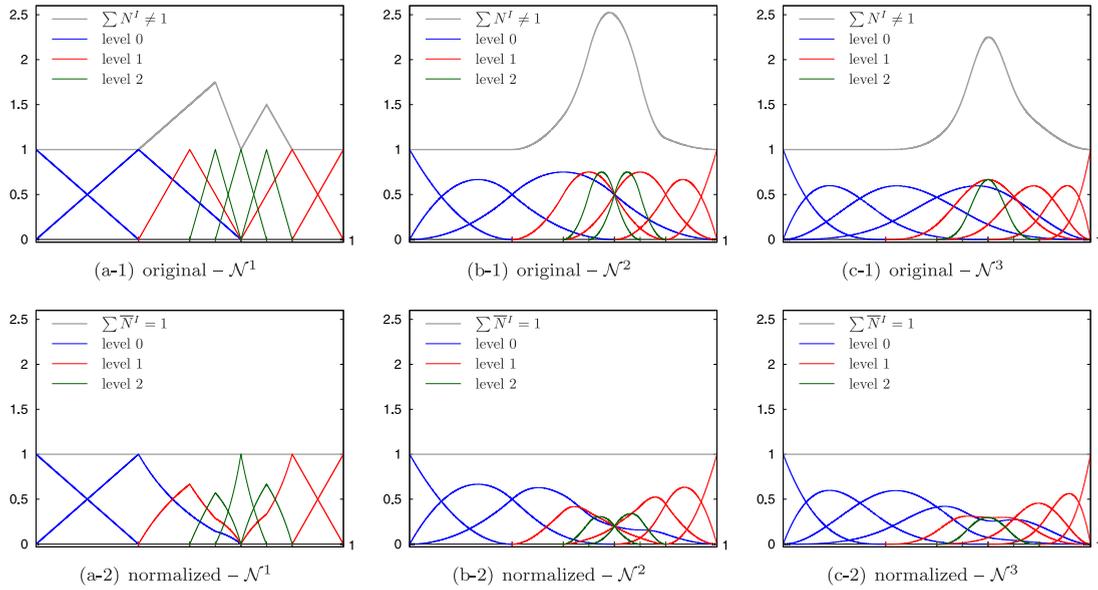


Fig. 4. The hierarchical basis sets of Fig. 1 are compared with their normalized counterparts. Note that the continuity of the normalized basis functions are now assessed with respect to the hierarchical element structure, represented by the knot structure in this one-dimensional case.

form are associated with the terms on the right-hand side of (3.9) and with the variation $\delta \bar{g}^I$, the latter appearing in (3.3) which may be restated as $\delta \mathcal{G}^c = - \sum_I \delta \bar{g}^I p^I$. Additional mortar integrals which emanate from these terms during linearization are standard [41].

In this work, two types of hierarchical mortar discretizations will be invoked. The first one inherits the hierarchical discretization of the slave surface displacement: $B^I = N^I$. However, this basis does not necessarily satisfy the partition of unity. A normalized basis

$$\bar{N}^I = \frac{N^I}{\sum_J N^J} \tag{3.11}$$

is introduced which does satisfy the partition of unity. The second choice is then $B^I = \bar{N}^I$. The numerical cost of normalization is negligible. Moreover, normalization is straightforward to implement since it is only applied to the mortar basis. Note that the bulk discretization also does not satisfy the partition of unity but this does not pose any difficulties for the purposes of the present study and therefore no modification is needed. In the numerical investigations, the non-normalized set of mortar basis functions will be denoted by \mathcal{B} and the normalized set by $\bar{\mathcal{B}}$. Normalization was discussed in [34], although it is possible to truncate the basis functions belonging to a refined level in order to ensure the partition of unity [35].

For the example of Fig. 1, the original and normalized basis functions are compared in Fig. 4. Normalized basis functions share various properties of NURBS, but not all. For instance, if the NURBS basis functions are of order p and the coarse level knot vectors have no repeated interior knots then the normalized basis functions are $p - 1$ times continuously differentiable across the hierarchical element boundaries. Also, normalization distorts initially linear basis functions, similar to the transition from unit to non-unit weights with a standard \mathcal{N}^1 discretization. However, unlike NURBS, a normalized basis function may attain multiple extrema across its span with the adopted hierarchical refinement algorithm (e.g. the last \mathcal{N}^3 basis function from level 0 in Fig. 4(c-2)).

Based on the hierarchical mortar discretization invoked above and a choice of B^I , a relation between γ^I and p^I completes the contact formulation. This relation is specified through the constraint enforcement method. Below, an exterior point method is considered, followed by an interior point method. Combinations of choices for (i) the mortar basis, (ii) the recovery approach and (iii) the method of constraint enforcement deliver classes of contact algorithms, some of which have clear advantages or disadvantages and others which perform well only for special problems.

3.3. Exterior point method

The augmented Lagrangian method is a primal–dual exterior point method that covers the penalty and Lagrange multiplier methods as special cases and requires keeping track of an active set \mathcal{A} . Presently, if $\gamma^I \geq 0$ is detected then the control point I becomes a member of \mathcal{A} for which a Lagrange multiplier λ^I is activated such that the pressure in (3.3) is formulated as

$$p^I = \lambda^I + \epsilon \gamma^I \quad (3.12)$$

where ϵ is a penalty parameter. λ^I will enforce γ^I to vanish for all $I \in \mathcal{A}$ and it should be deactivated whenever I leaves the active set, the latter detected if $p^I < 0$ at any stage during the Newton–Raphson iterations. These conditions are incorporated through a Lagrange multiplier contribution to the weak form:

$$0 = \delta \mathcal{L}^c = \sum_I \delta \lambda^I \omega \times \begin{cases} -\gamma^I & \text{if } I \in \mathcal{A} \\ \epsilon^{-1} \lambda^I & \text{if } I \notin \mathcal{A}. \end{cases} \quad (3.13)$$

Here, an additional scaling factor ω is introduced, which can be simply chosen as unity. If, instead, one chooses ω as w^I from (3.10), then the tangent matrix will be symmetric for the case of lumped recovery of γ^I . The tangent matrix will be unsymmetric in the case of full recovery. Note that although symmetry can be recovered by replacing $\omega \gamma^I$ with \bar{g}^I for $I \in \mathcal{A}$ in this case, this would be inconsistent with (3.9) as well as with (3.12), where γ^I is enforced to vanish rather than \bar{g}^I , and therefore can lead to an ill-defined problem with no solution. Using the Lagrange multiplier method, it is possible to directly enforce \bar{g}^I to zero, however this approach is intrinsically equivalent to lumping.

Overall, the full recovery approach closely resembles a recovery approach that was called *unconstrained* in [41,42], yet it is conceptually simpler and algorithmically more straightforward to implement. When all the control points on Γ_o^c are active, full and unconstrained approaches will be equivalent. Hence, it is concluded that the full recovery approach also does not display locking and therefore passes the patch test on flat or curved interfaces, which is true for the lumped approach as well. On the other hand, while no advantages of unconstrained recovery over lumping were observed in earlier studies, numerical investigations to be presented will indicate that full recovery can display numerically superior properties in certain scenarios.

The value of ϵ has no influence on the solution quality for the augmented Lagrangian method and has only a mild influence with respect to convergence. The Lagrange multiplier method is obtained by omitting γ^I in (3.12), in which case the value of ϵ in (3.13) becomes irrelevant. The penalty method is obtained as a primal exterior point method by omitting the λ^I term together with the Lagrange multiplier contribution $\delta \mathcal{L}^c$, in which case ϵ must be sufficiently large for an accurate solution without compromising convergence due to ill-conditioning. The linearization of these three possible formulations in order to determine the updates to the displacement and Lagrange multiplier degrees of freedom is largely standard and hence will not be discussed.

3.4. Interior point method

The interior point method to be employed in this work extends the presentation in [18] to the case of full recovery. Algorithmic details which are shared with this reference are omitted. In view of their considerable algorithmic similarity, this approach can easily be pursued starting from an augmented Lagrangian method implementation.

In interior point methods, the entirety of Γ_o^c is active at all times and the method is intrinsically built on the idea of maintaining a distance $g < 0$ between the surfaces. In the ideal case, based on the mortar discretization that has already been introduced, this can be readily achieved by explicitly enforcing a relation between γ^I and p^I to hold at every Newton–Raphson iteration, similar to (3.12) for the exterior point method. Specifically, a classical barrier regularization $p^I = -\mu/\gamma^I$ may be employed where μ is a barrier parameter. However, this method would have two shortcomings. The first of these is its characteristic similarity to penalty regularization: an accurate solution requires a judicious choice of μ , presently requiring a sufficiently small value without compromising convergence due to ill-conditioning. Interestingly, in sharp contrast with penalty regularization, this turns out to be a minor shortcoming such that μ can either be chosen as small as desired within a *primal* algorithm or it can be driven to zero automatically within a *primal–dual* version. The second shortcoming is major and algorithms for addressing it are central to modern interior point methods: surfaces may penetrate or touch each other during the Newton–Raphson iterations, leading to

$\gamma^I \geq 0$ and hence to unphysical or undefined p^I , yet there is no straightforward method of avoiding such occurrences by updating the displacements in a way that is applicable to a broad range of problems.

In order to address the major shortcoming, an additional contact variable s^I is introduced such that *at convergence* it satisfies $-\gamma^I = s^I > 0$, making use of (3.9), $\{s\} = -[\Phi]^{-1}\{\bar{g}\}$. This statement is equivalent to the following *gap condition* contribution to the weak form:

$$0 = \delta \mathcal{L}^c = \sum_I \delta p^I \left(\bar{g}^I + \sum_J \Phi^{IJ} s^J \right). \tag{3.14}$$

The linearization of the weak form (3.1) of the linear momentum balance involves displacements as well as Δp^I while the linearization of (3.14) entails the computation of

$$\Delta \delta \mathcal{L}^c = \sum_I \delta p^I \left(\Delta \bar{g}^I + \sum_J \Phi^{IJ} \Delta s^J \right). \tag{3.15}$$

In order to determine Δs^I , a primal algorithm can be pursued where μ is fixed at a predetermined value. *At convergence*, s^I will satisfy $p^I = \mu/s^I$. This local expression may be expressed as the *complementarity condition*

$$s^I p^I = \mu, \tag{3.16}$$

the linearization of which reads

$$s^I p^I + \Delta s^I p^I + s^I \Delta p^I = \mu \tag{3.17}$$

and hence delivers an expression for Δs^I :

$$\Delta s^I = (\mu - s^I p^I - s^I \Delta p^I)/p^I. \tag{3.18}$$

Substitution of this result into (3.15) completes the linearization of (3.14) and, together with the linearization of the linear momentum balance statement, leads to a system of equations that involves updates to the displacement and pressure degrees of freedom only, starting from an initial guess for the values of the displacements as well as for $\{s^I, p^I\}$. Note that updates to p^I determine the updates to s^I through (3.18). Since there is no active set in this interior point method, one must explicitly ensure that $\{s^I, p^I\}$ remain positive after update. This is easily achieved by scaling the updates when necessary via fraction to the boundary rules, without altering the updates to the displacements.

One should note that it is tempting to enforce (3.16) strongly such that $p^I = \mu/s^I$ is directly employed at every iteration, eventually leading to an alternative primal algorithm which eliminates the pressure degrees of freedom. This version, however, turns out to have inferior convergence characteristics in comparison to the one summarized above. Moreover, the form presented here allows a relatively straightforward extension to the primal–dual algorithm in which μ is driven to zero automatically. In this context, it is noted that a difference between the lumped and full recovery approaches appears only in the gap condition (3.14). Although the statement of the complementarity condition in the form (3.16) also relies on a lumping argument, retaining this form is advantageous since it allows applying the primal–dual algorithm for the lumped setting without any modification—see [18] for details.

4. Numerical investigations

4.1. Remarks on setup

A series of numerical examples will be presented which concentrate on various features of different choices for the contact algorithm in the context of hierarchical discretizations. The effects or advantages of order elevation are not directly studied but orders ranging from two to six are employed in different examples to illustrate and benefit from this capability.

In all examples, an isotropic elastic material response is assumed. A Neo-Hookean material model will be employed in the case of finite deformations for which, in a d -dimensional setting, the strain energy function is

$$\mathcal{W} = \frac{A_1}{2} (\ln J)^2 + \frac{A_2}{2} (J^{-2/d} \text{tr}[\mathbf{C}] - d) \tag{4.1}$$

where A_1 and A_2 are the bulk and shear moduli, respectively, and $J = \det[\mathbf{F}]$, $\mathbf{C} = \mathbf{F}^T \mathbf{F}$. The values of the moduli are either prescribed directly or assigned indirectly through the values of the Young's modulus E and the Poisson's ratio ν .

In the examples involving incremental loading, the simulation progress will be monitored by the variable $t \in [0, T]$. When contact is initiated, the contact integrals are evaluated via Gaussian quadrature based on the (hierarchical) surface element structure of the slave (the upper body in all cases—indicated by the blue color in figures displaying the geometry). The advantages and shortcomings of such an approach for mortar-based methods have recently been assessed in [43] with respect to efficiency and accuracy. Presently, the order of integration employed ensures sufficiently accurate results. In general, for volume integrals, four integration points were employed along a parametric direction with up to \mathcal{N}^3 discretizations and six points were employed for higher order discretizations up to \mathcal{N}^6 . For contact integrals, the number was doubled with respect to these choices. The reader is referred to [21] and [10] for more accurate isogeometric mortar methods which employ segmentation-based integration schemes in the context of domain decomposition and contact, respectively. In the case of regularization, the absolute values of the penalty or barrier parameters will not be needed (except in Section 4.4.1), hence their values are only reported when necessary relative to reference values $\mu_o = 10^{-5}$ or ϵ_o , chosen to be ten times the maximum diagonal entry of the stiffness matrix associated with the bulk contribution to the weak form. For the augmented Lagrangian method, the value of ϵ in (3.12) to be reported is set in the first Newton–Raphson iteration of a load step and subsequently kept constant within that load step.

4.2. Patch test

Basic features of different choices for the contact algorithm are first demonstrated with a flat interface patch test at large deformations. Only \mathcal{N}^2 discretizations are employed, other choices displaying similar trends. Each body initially has a square geometry with dimension 1 and the initial gap between the bodies is 0.05. A displacement of 0.4 is then applied to the upper body. Presently, the bulk and shear moduli correspond to the choices of $E = 1$ and $\nu = 0.3$.

The geometry and results of the test are summarized in Figs. 5 and 6. With a uniform mesh, this patch test is passed for all combinations of different choices so that a constant pressure is obtained across the interface. With a hierarchical mesh, the test is passed for all combinations that enforce the constraints exactly, i.e. using the augmented Lagrangian method or with the primal–dual interior point method, which deliver a uniform pressure p_{exact} that results in a uniform deformation. Neither the penalty method (with small ϵ) nor the primal interior point method (with large μ) passes the test if the original non-normalized mortar basis is used, leading to a non-uniform interface deformation with excessive penetration or gap. Retaining this basis, the patch test is satisfied with full recovery if the penalty method is employed but the primal interior point method still fails. Finally, with a normalized basis the test is satisfied irrespective of other choices for the contact algorithm. However, since the test is displacement controlled, a non-zero penetration due to a small ϵ leads to a slightly lower pressure compared to p_{exact} for the penalty method and a non-zero gap due to a large μ leads to a slightly higher pressure for the primal interior point method.

The problem with the patch test arises only in the case of a non-exact enforcement of the constraints using the non-normalized mortar basis, and full recovery only helps in the case of the penalty method. The source of this observation is easily traced to the recovery statement (3.9):

- The normalized basis satisfies the partition of unity so that a constant $g = c$ not only implies a constant $\gamma = -s = c$ but also a constant $\gamma^I = -s^I = c$ irrespective of the method of recovery. Hence, $p^I = \epsilon c$ or $p^I = -\mu/c$ is obtained that leads to $p = \epsilon c$ or $p = -\mu/c$.
- The original basis does not satisfy the partition of unity such that a constant $g = c$ implies a constant $\gamma = -s = c$ only with full recovery but even then $\gamma^I = -s^I$ will not be a constant. Although p^I will not be a constant in any case, a constant p will be implied if p is directly expressible in terms of $\gamma = -s$. This is the case for the penalty method but not for the primal interior point method where, instead of the continuum expression $p = \mu/s$, the approximate discrete expression $p^I = \mu/s^I$ is employed as remarked earlier in Section 3.4.

In the limit of an exact enforcement of the contact constraints one obtains satisfactory results irrespective of the mortar basis or the recovery approach since enforcing $\gamma^I = 0$ for all I is equivalent to enforcing $g = 0$ in all cases. Moreover, in Section 4.3.3 it will be demonstrated that in this limit retaining the approximate discrete expression $p^I = \mu/s^I$ does not adversely affect the results since the local solution quality is primarily controlled by the recovery approach.

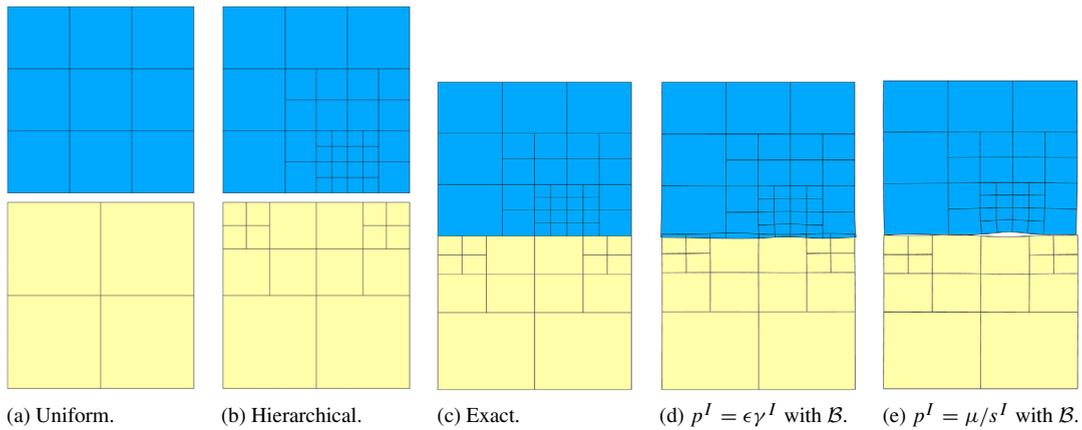


Fig. 5. The patch test geometry of Section 4.2 is summarized. The hierarchical discretization of the upper body corresponds to Fig. 3 with \mathcal{N}^2 . In (d) and (e), lumped recovery is employed.

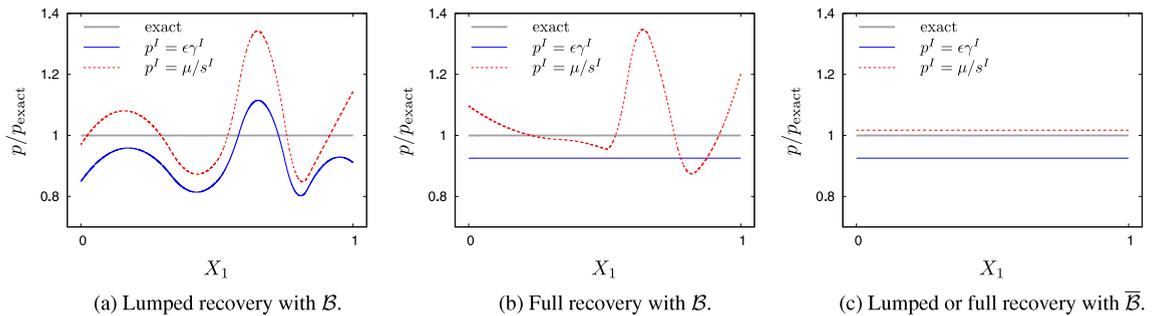


Fig. 6. The analysis of Section 4.2 is summarized, based on the geometry in Fig. 5. The exact result corresponds to an augmented Lagrangian or a primal–dual interior point method solution. Note that, here and in various upcoming examples, X_1 traces the referential position on the slave surface in the horizontal direction—see (3.6).

4.3. Hertz contact problem

In order to illustrate some of the fundamental observations regarding the impact of different choices for the contact algorithm on the local solution quality, the classical Hertz contact problem is considered. The two-dimensional geometry of the contact setup with two bodies is borrowed from [41]. To summarize, the radius of curvature for each body is 1 and each one is displaced by 5×10^{-3} towards the other from an initial configuration where the surfaces touch at a point. For each body, $E = 1$ and $\nu = 0.3$.

In all figures, the distribution of the dimensionless pressure p/p_{\max} is monitored across the dimensionless position r/r_{contact} where p_{\max} is the maximum pressure and r_{contact} is the half-width of the contact zone. The exact solution is indicated with a grey line on the background.

4.3.1. Recovery comparison

Lumping of (3.9) is an algorithmically convenient simplification that has been successfully used in various earlier studies on isogeometric contact. To briefly demonstrate an advantage of full recovery, matching non-hierarchical discretizations are assigned to the slave and the master such that the radial directions are discretized with 72 \mathcal{N}^1 elements where as the angular direction is first discretized at a very coarse resolution with 96 \mathcal{N}^3 elements and subsequently this number is doubled/quadrupled to obtain coarse/medium resolutions.

The results are summarized in Fig. 7, by default based on the augmented Lagrangian method (employing $\epsilon = \epsilon_0$ in (3.12)). At the very coarse discretization, the number of control points in the active set is very small. Nevertheless, full recovery is observed to already give a satisfactory answer that also rapidly approaches the exact solution with increasing resolution. Lumped recovery, on the other hand, delivers significantly oscillatory results at coarse resolutions yet these oscillations are not due to locking. Therefore a satisfactory solution is obtained beyond a

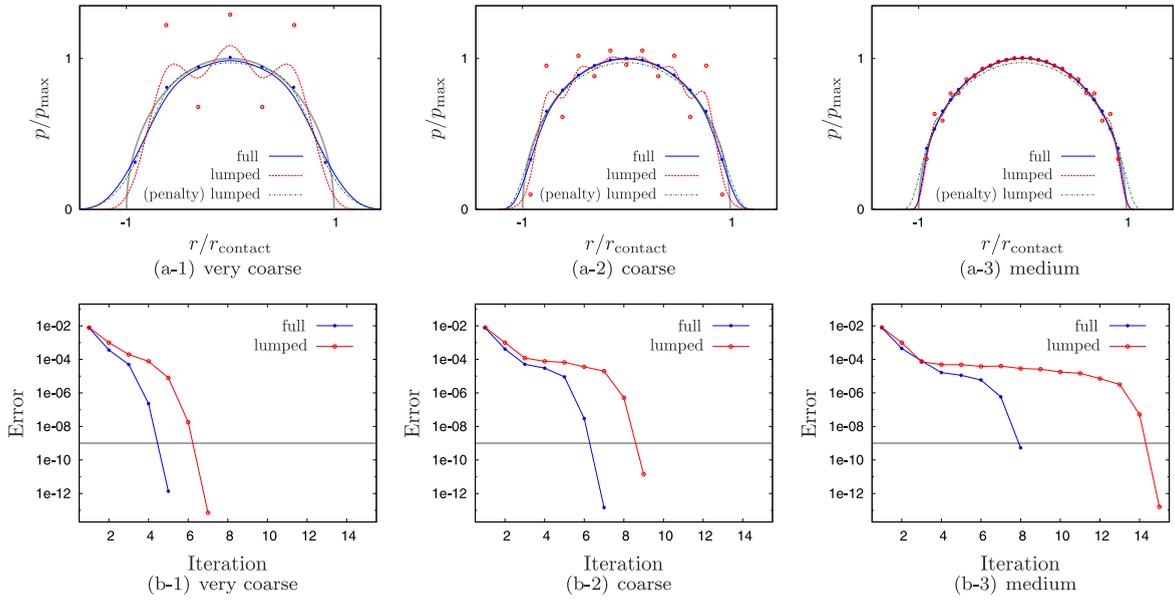


Fig. 7. The analysis of Section 4.3.1 is summarized. In the pressure plots, the points indicate the control point pressures. In the error plots, the error indicates the magnitude of the updates to the displacements scaled by the magnitude of the updated positions and the gray line indicates the convergence tolerance (10^{-9}). The penalty parameter value corresponds to $10\epsilon_\sigma$ of Fig. 9.

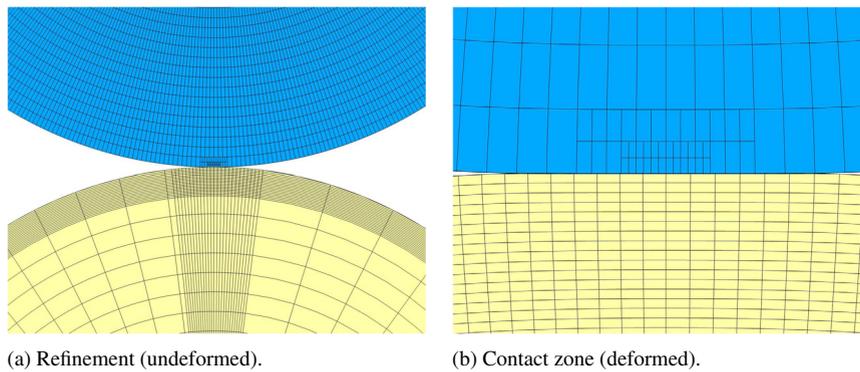


Fig. 8. The hierarchical refinement and the close-up of the contact zone are shown for the analysis of Section 4.3.2.

sufficient number of control points in the active set. It is also observed that the number of iterations to achieve a prescribed convergence tolerance is smaller for full recovery, particularly at finer discretizations. Indeed, almost invariably for all the simulations carried out in this work full recovery required less iterations than lumped recovery, specifically when the constraints are enforced exactly. Consequently, one can state that the advantage of full recovery is higher efficiency with respect to a capability of delivering: (i) a better solution at coarse resolutions and (ii) a faster convergence at fine resolutions. Note that if, instead of augmented Lagrangian, the penalty method is employed with a reasonably large but not very large penalty parameter then the oscillations observed in lumped recovery are not displayed and an acceptable solution is obtained as also shown in Fig. 7. However, this smoothing effect does not translate to an advantage of the penalty method but rather only highlights the importance of driving it to its limits to correctly assess its capabilities.

4.3.2. Exterior point method

Hierarchical discretizations will first be investigated with the exterior point method. To concentrate on the choice of the mortar basis functions, only the slave will be discretized hierarchically and with a rather coarse resolution to emphasize various aspects (Fig. 8). Initially, $36 \mathcal{N}^1$ elements are assigned to both the master and the slave in the radial

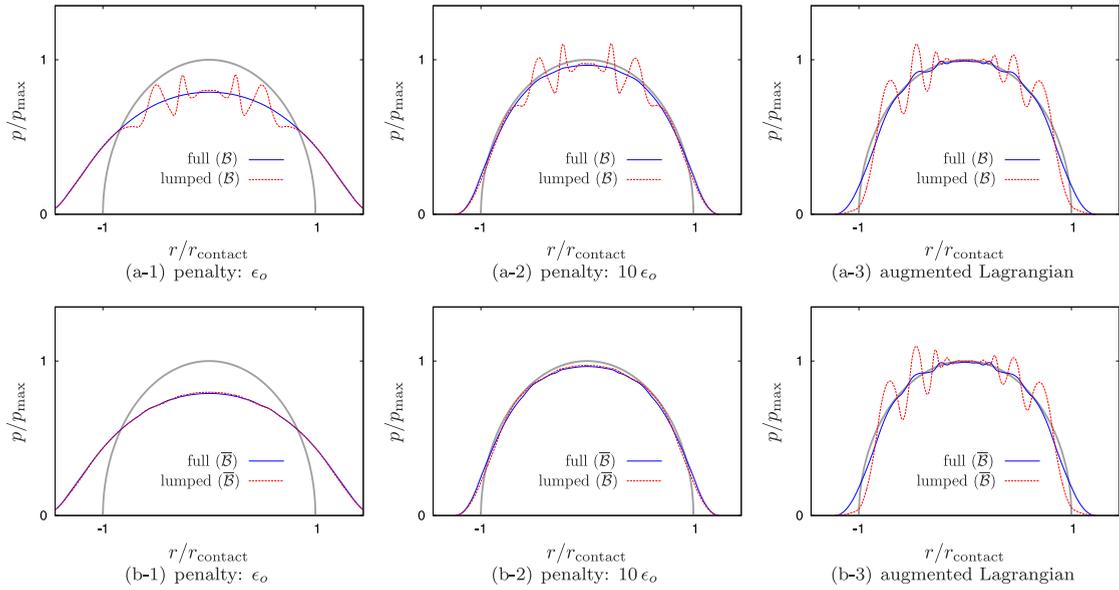


Fig. 9. The analysis of Section 4.3.2 is summarized, based on the discretization in Fig. 8 with exterior point methods for non-normalized (\mathcal{B}) and normalized ($\bar{\mathcal{B}}$) mortar basis sets.

direction and $128 \mathcal{N}^2$ elements are assigned to the slave in the angular direction whereas 36 elements are assigned to the master. Subsequently, the mesh on the master is concentrated near the contact zone by knot relocation and the mesh on the slave is refined by two levels such that the refinement zone roughly corresponds to the contact zone, overall leading to non-matching meshes at the interface. The knot relocation on the master is slightly non-symmetric, which will induce slightly non-symmetric pressure distributions.

The results summarized in Fig. 9 for the original mortar basis set is indicative of two aspects. First, unlike the uniform meshes in Fig. 7, for lumped recovery with a non-normalized mortar basis even small values of the penalty parameter lead to oscillations that correlate with the structure of hierarchical refinement. Full recovery again displays a smoothness that is also retained at the exact constraint enforcement limit of augmented Lagrangian. At this limit, lumped recovery continues to display oscillations, but this time not due to hierarchical refinement but rather due to the coarseness of the hierarchically refined mesh, similar to the case of Fig. 7. This establishes a link to the second aspect, namely that the use of a normalized mortar basis removes the oscillations observed with the penalty method for lumped recovery. In the limit of augmented Lagrangian, however, lumped recovery again displays oscillations, which verifies that these are due to the coarseness of the mesh. Overall, for this example, one can state that the effect of normalization is only important for the penalty method with lumped recovery and normalization is not necessary at the augmented Lagrangian limit. At this limit, similar to uniform meshes, full recovery again delivers superior results since the hierarchical refinement is not sufficiently fine.

Additional analysis in Fig. 10 reinforces earlier statements. Here, the hierarchical mesh in Fig. 8 is refined by uniformly subdividing each element by three along both parametric directions, thereby retaining the locations of the hierarchical transitions along the slave surface. The oscillations that were observed for the penalty method with a non-normalized mortar basis and lumping are now even further concentrated to spikes at the transition regions, thereby verifying that this combination does not benefit from refinement. In the augmented Lagrangian limit, however, the lumped recovery solution does benefit from refinement as predicted. As in earlier examples, full recovery displays a smooth result that only improves with refinement. In all cases, a difference between the normalized and non-normalized mortar basis is only apparent in the case of the penalty method with lumping and any difference is hardly noticeable for other combinations of choices even at a coarse resolution.

4.3.3. Interior point method

Fig. 11 repeats the analysis of Fig. 9 based on the interior point method. The interior point method in a primal setting resembles the penalty method and in a primal–dual setting it enforces the contact constraints exactly so that the solution obtained is the same as the augmented Lagrangian solution. For this reason, almost all the observations

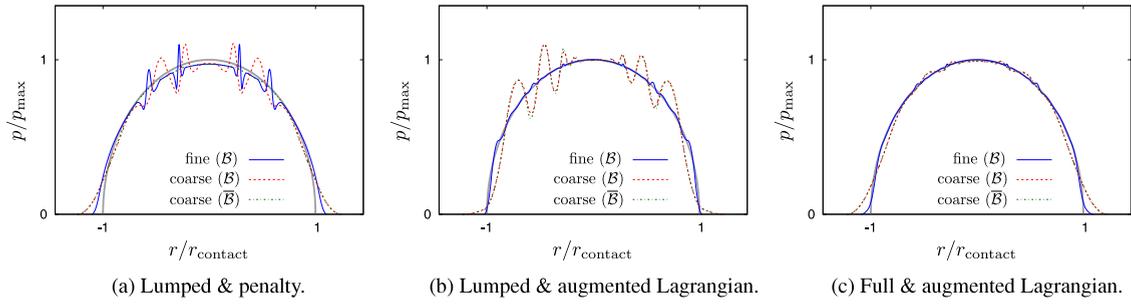


Fig. 10. Complementary analysis for Fig. 9 is presented using non-normalized (\mathcal{B}) and normalized ($\bar{\mathcal{B}}$) mortar basis sets. Here, the coarse mesh corresponds to the resolution of Fig. 8 and the fine mesh is obtained by subdividing each element by three along both parametric directions.

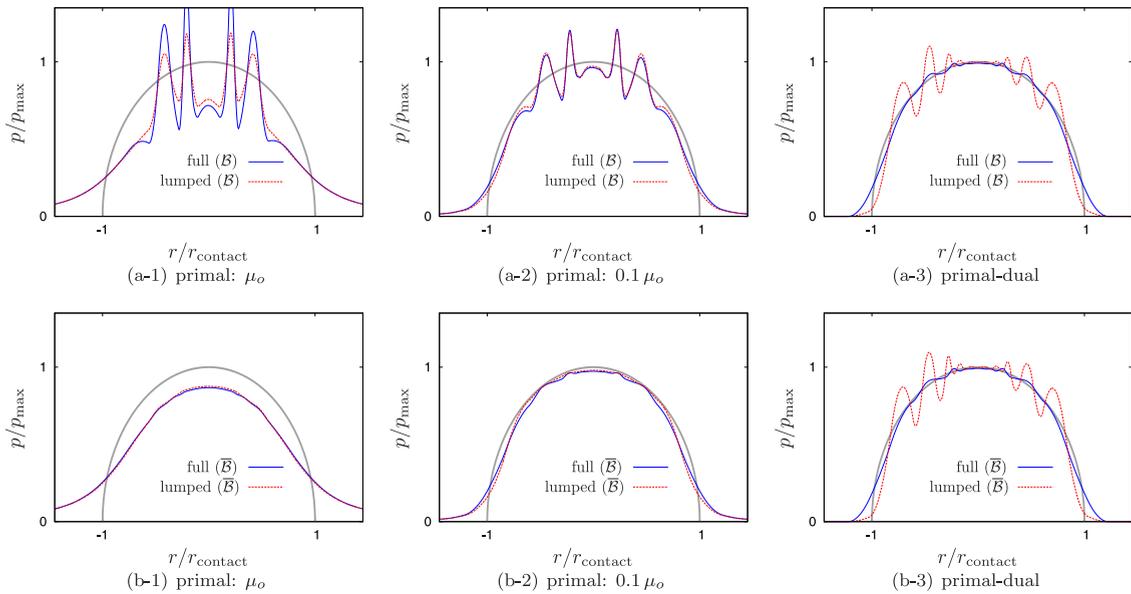


Fig. 11. The analysis of Fig. 9 is repeated with interior point methods ($\mu_o = 10^{-5}$).

stated earlier for the exterior point method apply, with one exception. For a non-normalized mortar basis, full recovery also leads to significant oscillations for the penalty method although the oscillations disappear in the primal–dual limit. This suggests that the observed issue is not a mesh resolution effect but rather due to the choice of the mortar basis functions as verified by the smooth results obtained with normalization. Overall, it can be stated that the variation in the solution quality for the Hertz problem with different combinations of choices for the contact algorithm correlates well with the preceding studies in Sections 4.2 and 4.3.1. In the remaining sections, different choices for the contact algorithms will be tested in finite deformation contact problems with evolving contact interfaces.

4.4. Ironing problem

The ironing problem is considered as a benchmark case in two- and three-dimensional settings. The geometry of the problem is directly borrowed from [8], the changes being the downward displacement of the upper body (0.3) and its stiffness (10^3 times larger than the lower one).

In this problem, the discretization of the bodies will be hierarchically refined towards the potential contact surfaces but the discretization on these surfaces will be uniform. This is achieved by refining all elements which lie adjacent to these surfaces, starting from the coarse level discretization that is visible in the corresponding figures displaying the geometry. The comparatively stiff slave body is assigned a \mathcal{N}^2 discretization in all directions while the softer master is assigned a \mathcal{N}^6 discretization in the direction of sliding in order to achieve a smoother evolution of the contact forces and a \mathcal{N}^2 discretization in the remaining directions. Due to the uniform mesh on the potential contact surfaces, there is

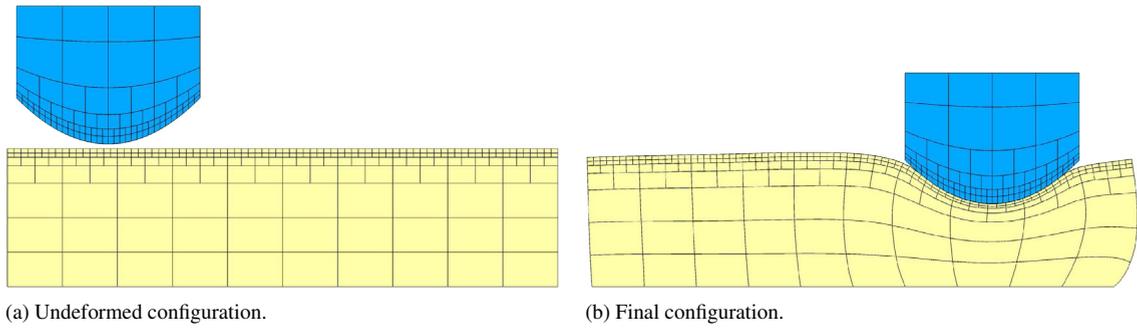


Fig. 12. The two-dimensional ironing problem of Section 4.4.1 is summarized with three levels of hierarchical refinement ($\mathcal{L} = 3$).

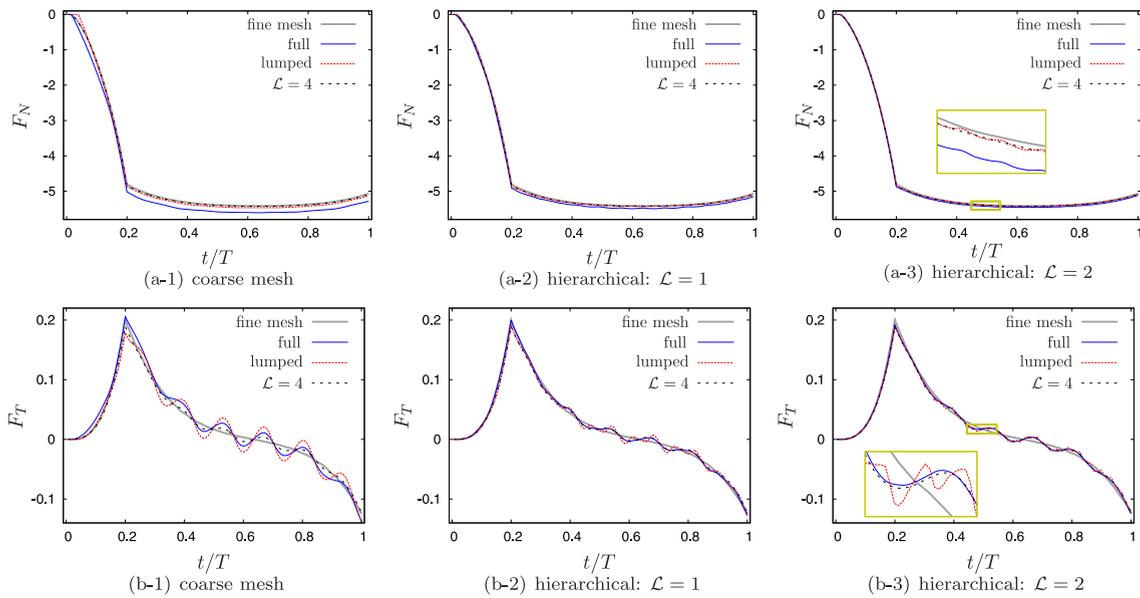


Fig. 13. The normal (F_N) and tangential (F_T) forces associated with the two-dimensional ironing problem of Fig. 12 are summarized. Note that, here and in upcoming examples, \mathcal{L} indicates the number of hierarchical levels of refinement—see Section 2.2.

no need to distinguish between normalized and non-normalized mortar basis sets in this problem. By default, an exact constraint enforcement will be pursued based on the augmented Lagrangian method (employing $\epsilon = 0.1\epsilon_o$ in (3.12)).

4.4.1. Two-dimensional case

The two-dimensional test and the corresponding forces measured on the upper surface of the slave body are summarized in Figs. 12 and 13. Here, the fine mesh solution employs uniform refinement at the resolution of level 4 and delivers a converged force curve for which lumped and full recovery results are virtually overlapping, hence are not distinguished. The visible offset from the fine mesh solution for the normal force as well as the oscillation in the tangential force rapidly diminishes with hierarchical refinement. After four levels of hierarchical refinement, the forces still do not overlap with the fine mesh solution since refinement is highly localized near the contact surfaces and hence the solution is still significantly influenced by the coarse elements away from the contact zone. However, due to the fine meshes on the surfaces, the difference between full and lumped recovery diminishes with increasing refinement and after four levels of hierarchical refinement the two curves cannot be distinguished. Moreover, it is observed that full recovery displays smoother tangential contact forces at coarse resolutions, which correlates with earlier observations regarding smoother pressure distributions. On the other hand, it can also be observed that full recovery always overestimates the magnitude of the normal force in comparison to lumped recovery. This point will be revisited in Section 4.5.

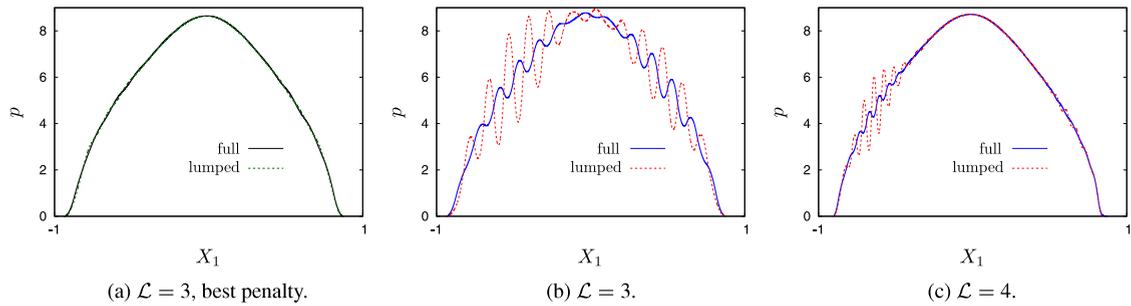


Fig. 14. The (non-Hertzian) pressure distributions for the analysis of Fig. 13 are shown at $t/T = 1$. The best penalty solution is obtained with $\epsilon = 10^4$.

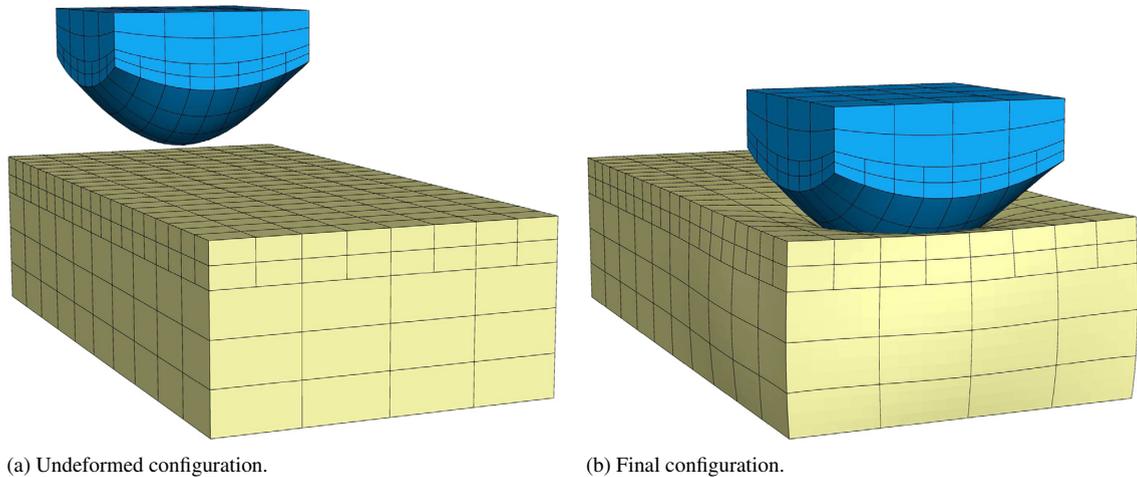


Fig. 15. The three-dimensional ironing problem of Section 4.4.2 is summarized with $\mathcal{L} = 1$.

The forces in Fig. 13 are virtually overlapping for the augmented Lagrangian method and the primal interior point with $\mu = 10^{-5}$, hence are not distinguished. In comparison to global measures such as forces, local measures such as the pressure distribution are more sensitive to the choice of regularization parameters. The pressure distribution in the final configuration is compared in Fig. 14 for different contact algorithms. In these cases, a value of $\mu = 10^{-8}$ also delivers virtually overlapping pressure distributions with augmented Lagrangian. Again, full recovery delivers a less oscillatory result than lumped recovery and, although the forces do not significantly improve beyond $\mathcal{L} = 3$, the oscillations in the pressure significantly diminish with increasing hierarchical refinement. The best penalty regularization result with respect to the augmented Lagrangian limit for this example was obtained with $\epsilon = 10^4$, which is similar in quality to the choice of $\mu = 10^{-4}$. At this value, there remains a difference with the force predictions of the augmented Lagrangian method (not shown) yet larger choices of ϵ prohibit convergence despite aggressive load step refinement. This is a major advantage of the primal interior method where μ can be chosen to be even smaller than 10^{-8} without compromising numerical performance—see also [18] for extensive discussions in the context of non-hierarchical discretizations. The smoothing effect in the pressure distribution that was earlier discussed in Section 4.3 is observable in this case as well at this insufficiently large value of ϵ .

4.4.2. Three-dimensional case

In order to demonstrate analysis capability, the ironing problem is revisited in a three-dimensional setting (Fig. 15). The results are summarized in Fig. 16. Here, the reference fine mesh solution employs: 16/3 elements in the horizontal/vertical directions of the slave, 60 elements in the sliding directions of the master and 9 elements in the remaining directions. The full recovery solution is chosen as the reference fine mesh solution in the coarse mesh and hierarchical solution figures.

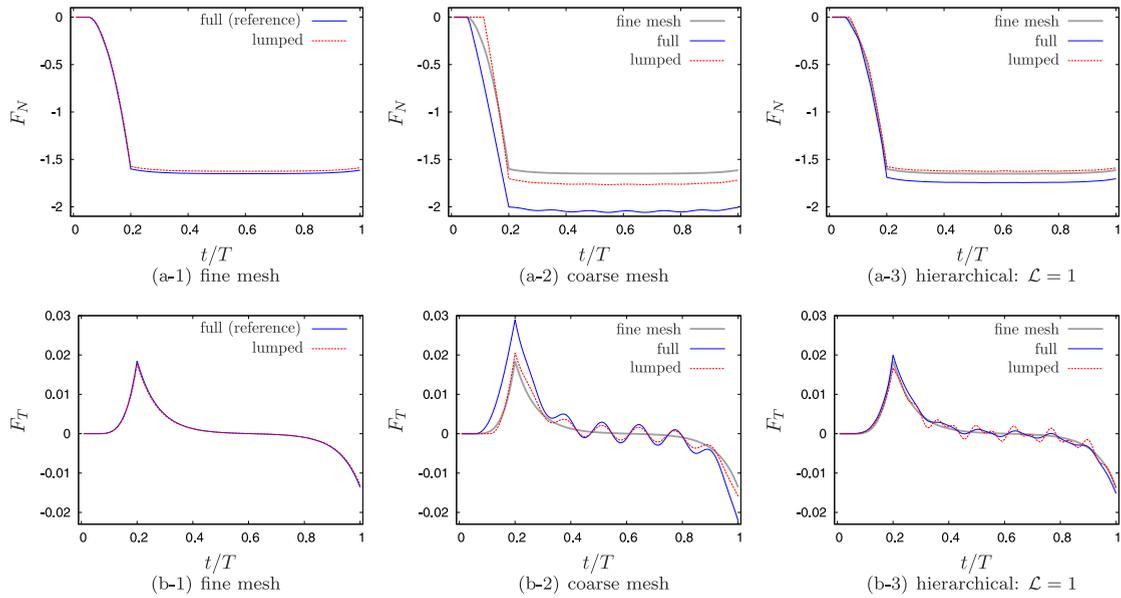


Fig. 16. The normal (F_N) and tangential (F_T) forces associated with the three-dimensional ironing problem of Fig. 15 are summarized.

On the finest uniform mesh that could be employed for this setup, full and lumped recovery solutions slightly differ from each other which suggests the need for further mesh refinement. Nevertheless, the full recovery solution is chosen as a sufficiently accurate reference solution. On the coarse mesh, both normal and tangential forces significantly differ from the reference solution. As in the two-dimensional case, full recovery overestimates the magnitude of the normal force at such coarse meshes. Only with a single level of refinement, the normal force predictions improve significantly, lumped recovery delivering a solution that is very close to the reference solution. For the tangential force, full recovery delivers a solution that again rapidly achieves smoothness in comparison to lumped recovery. Despite the superior performance of full recovery in most cases tested so far, the ironing problem suggests that this approach may be inferior to lumped recovery at considerably coarse meshes. This point is investigated next in detail.

4.5. Shell problem

In the following two- and three-dimensional test cases, thick-walled shell-like bodies are considered. The setup of this benchmark problem in three dimensions is borrowed from [41] and the two-dimensional case that is aimed to highlight potential shortcomings of full recovery is a direct simplification of this setup. The bodies are discretized with \mathcal{N}^2 elements in all directions for both cases. In order to concentrate on the comparison of the two recovery approaches, only the augmented Lagrangian method will be employed (with $\epsilon = 0.1\epsilon_0$ in (3.12)). In agreement with the discussions of Section 4.3, the forces measured with lumped recovery in these test cases differ only slightly for normalized and non-normalized mortar basis sets due to the exact enforcement of the constraints. Moreover, the difference vanishes with increasing mesh resolution. Therefore, the choice of the mortar basis will not be distinguished in the presentation of lumped recovery results with the exception of Fig. 19.

4.5.1. Two-dimensional case

In the two-dimensional case (Fig. 17), the flat slave body is held stationary by fixing displacements at the left and right edges while the shell-like master body is displaced upwards towards the slave by displacement control in a similar fashion. Due to the symmetry of the setup, only the normal force F_N exerted on the slave by the master will be monitored. In the case of uniform meshes, the accuracy of the solution in this problem with respect to F_N is solely driven by the number of elements in the tangential direction. A mesh with a single element in the normal direction and 32 elements in the tangential direction already delivers a converged force profile that is identical for both full and lumped recovery. This solution is chosen as the reference. In subsequent discretizations, the master body is assigned the reference resolution while the slave is first coarsely discretized with 4 elements and subsequently elements adjacent to the potential contact surface is refined through one or two levels.

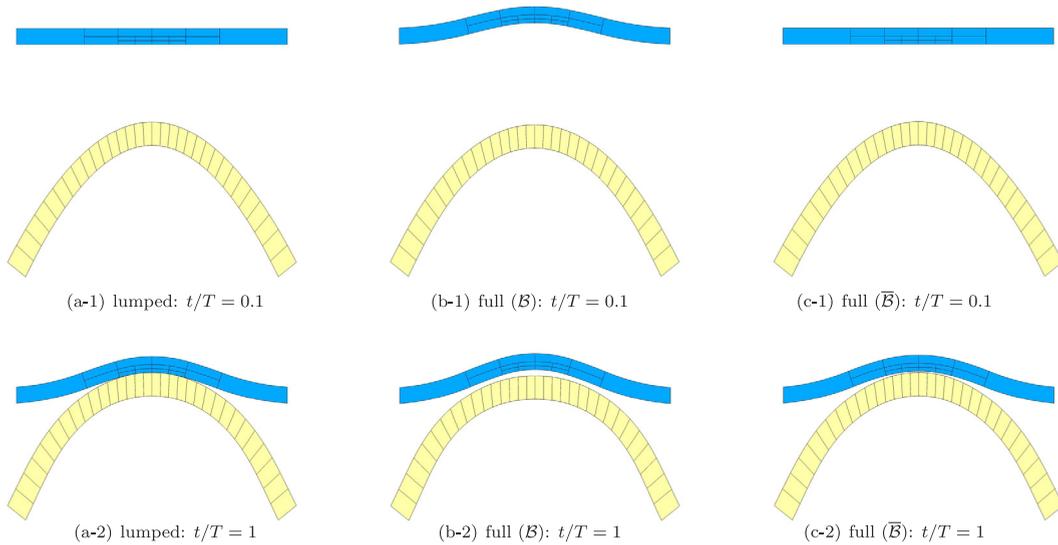


Fig. 17. The two-dimensional problem of Section 4.5 is summarized with $\mathcal{L} = 2$.

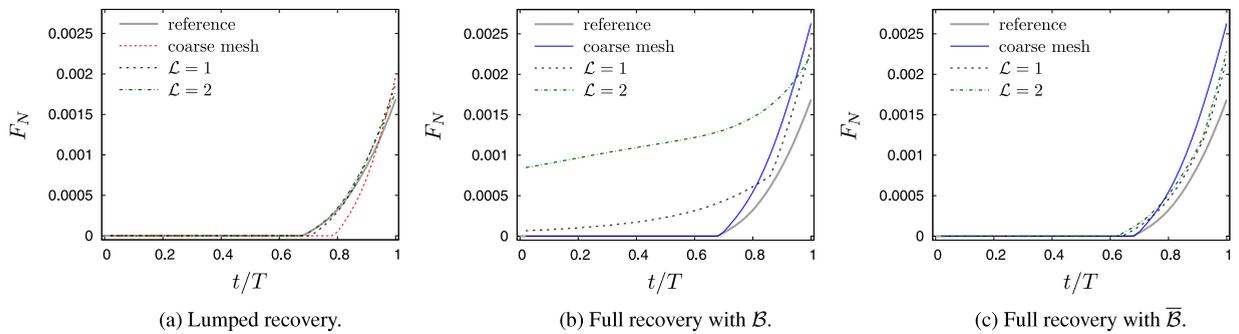


Fig. 18. The normal force (F_N) evolutions corresponding to the cases in Fig. 17.

The results corresponding to the cases in Fig. 17 are summarized in Fig. 18. In the lumped recovery case, a single level of refinement significantly improves the solution but the second localized refinement does not. This is expected since the mesh away from the contact zone is still significantly coarse. When full recovery is employed with the *non-normalized* mortar basis, a significantly different picture emerges where the solution quality deteriorates with increasing levels of refinement. By switching to a *normalized* basis, one obtains solutions that are comparable to the lumped recovery results. Fig. 17 reflects this behavior, where deformation occurs with non-normalized full recovery as soon as the simulation starts despite the large distance between the bodies. Additionally, a gap is observed at the contact interface of the final configuration with normalized full recovery which reflects as an over-estimation of the force magnitude in comparison to the reference solution, similar to an earlier observation in Section 4.4.1.

In order to investigate the source of the observations stated above, the evolution of the kinematic contact variable γ , which is the detection of the normal gap g by the mortar discretization, is depicted in Fig. 19. Since lumping is a strong simplification of the recovery statement (3.9), the resulting γ profile is not a faithful reproduction of the actual gap g although the quality significantly improves with normalization. Nevertheless, irrespective of the mortar basis and the resolution, lumped recovery will not indicate contact before there is penetration ($g < 0$) at the potential contact interface since γ^I is proportional to the projected gap $\bar{g}^I = \langle B^I g \rangle$ and B^I are non-negative. On the other hand, full recovery is intrinsically based on a faithful reproduction of g through the direct solution of (3.9). Hence, with or without normalization, it delivers a γ profile that conforms very accurately to g before contact is initiated. However, contact is initiated based on the values of γ^I in all cases. Consequently, two shortcomings arise with full recovery. In order to discuss these shortcomings, the structure of hierarchical refinement of the mortar basis is depicted in Fig. 20.

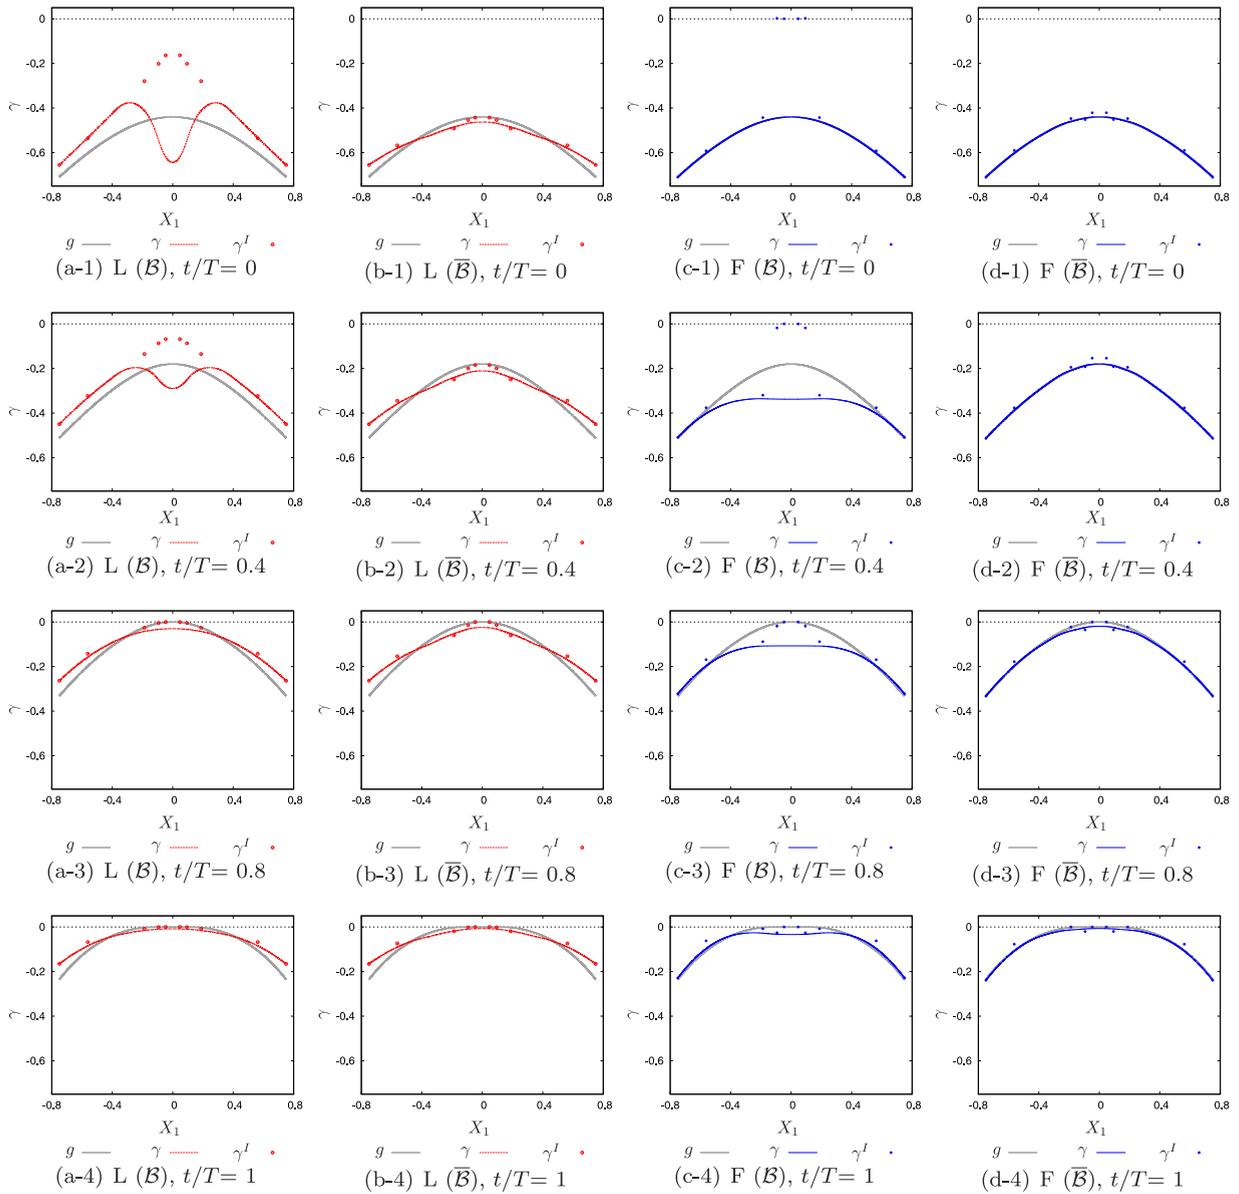


Fig. 19. For the cases in Fig. 17, instances from the variation of the recovered normal gap (γ) profiles are summarized for full (F) and lumped (L) recovery approaches with the non-normalized (\mathcal{B}) and normalized ($\bar{\mathcal{B}}$) mortar basis sets. The actual normal gap (g) profile corresponds to the γ profile of the reference mesh solution.

Clearly, refinement at the coarsest level does not remove any basis functions in this example. By making use of Eqs. (3.6)–(3.8) may be expanded in this case explicitly as

$$\gamma = \sum_{I=1}^{n_u} B^I \gamma^I \stackrel{!}{=} \sum_{K=1}^{n_x^0} B_0^K \gamma_0^K + \sum_{\ell=1}^2 \sum_{\alpha=1}^{n_\ell^\ell} B_\ell^\alpha \gamma_\ell^\alpha \quad (4.2)$$

where the first set is the coarse level contribution to the hierarchical discretization of γ and the remaining are associated with level 1 and 2 contributions such that $n_x^0 + \sum_{\ell} n_u^\ell = n_u$. The following observations may then be stated:

1. *Non-normalized basis ($B^I = N^I$):* The geometry is exactly represented by the coarse level basis functions N_0^K . If g could also be exactly represented by N_0^K , then γ_ℓ^α would vanish. This is indeed the case, for instance, if the

master body is also flat. When surfaces are curved, g depends on the discretization of both bodies and hence does not in general admit a representation based on N_0^K alone. This is the case for the present problem and some of the resulting non-zero γ_ℓ^α which essentially hierarchically refine the gap representation are in fact positive from the outset (Fig. 19(c-1)), thus leading to an immediate generation of a spurious force despite the distance between the bodies.

2. *Normalized basis ($B^I = \bar{N}^I$):* The normalized forms \bar{N}_0^K cannot exactly represent the geometry. Indeed, even if the master body is also flat such that the gap is a constant at $g = c$, due to the partition of unity one necessarily obtains $\gamma^I = c$ for all I . Consequently, a spurious force is not generated. However, since NURBS discretizations are not interpolatory, γ^I can lie far outside the γ profile, in particular at coarse discretizations (Fig. 19(d)). Therefore, contact can be initiated before penetration occurs, thereby leading to an over-estimation of the force magnitude due to displacement control, which was also observed in Section 4.4.

The observation regarding the first shortcoming suggests that whenever refinement at a given level removes basis functions then this major shortcoming may be alleviated. This indeed has been verified in various test discretizations based on the present setup (not shown) and was also the case in the hierarchical discretizations employed in Sections 4.3.2 and 4.3.3 of the Hertz contact problem. However, the adopted hierarchical refinement algorithm does not guarantee removal in general and therefore such an analysis is not pursued further. One may speculate that a discretization of the pressure in a fashion similar to that of the surface through coarse (as in (3.6)) and hierarchical (as in (3.7)) contributions may resolve this issue. Another possibility would be to adopt an alternative hierarchical refinement algorithm that guarantees removal of basis functions from the refined zone (which would still require addressing the second shortcoming). The investigation of either possibility is a major task that shall not be undertaken in the present work.

The second shortcoming is in fact not special to hierarchical discretizations. For a non-hierarchical discretization, full recovery is similar in nature to checking contact with a node-to-surface (NTS) type formulation based on the projection of spatial control point positions onto the master surface. Despite this similarity, it is expected to perform significantly better than NTS due to the underlying mortar structure. Such a comparison is beyond the scope of this work. It is sufficient to note that either method would predict contact initiation more accurately with an increasing number of degrees of freedom on the potential contact surface since the control degrees of freedom would become increasingly interpolatory. It is highlighted that the size of the contact zone may be small—only the number of degrees of freedom will control the quality of contact detection. In the light of experience gained through the simulations carried out in this work, it can be stated that this shortcoming is minor: it is significant only at considerably coarse discretizations and rapidly diminishes with the refinement of the coarse level mesh or with a uniformly increasing refinement on the potential contact surface. This was indeed the case with the Hertz contact and ironing problems discussed earlier. Hence, one might even state that there is no shortcoming—at coarse resolutions lumping allows significant local penetration whereas full recovery leaves a gap. However, in the context of hierarchical discretizations, it is additionally important to investigate how much a contact algorithm benefits from local refinement. This last point is investigated next.

4.5.2. Three-dimensional case

In the three-dimensional setup with two thick-walled shells in sliding contact, the bodies are initially discretized coarsely with 6 elements along the tangential directions and a single element in the thickness direction. Subsequently, a hierarchical refinement concentrated in the contact zone is carried out through two levels to obtain the discretization depicted in Fig. 21. The reference solution is based on a uniformly fine mesh at the level 2 resolution, namely 24 elements along the tangential directions and 4 elements in the thickness direction. This mesh delivers a sufficiently accurate but not converged results, hence each recovery method is explicitly compared with its corresponding fine mesh solution. The presented results are reliably obtained using the augmented Lagrangian method or the primal interior point with a sufficiently small μ (10^{-5} and beyond). In view of the fact that the use of a non-normalized basis requires significant care with full recovery (Section 4.5.1), only normalized basis functions have been considered in this remaining example. Lumped recovery is only very slightly sensitive to normalization as in the two-dimensional case.

The normal and tangential force evolutions are summarized in Fig. 22. In the case of full recovery, detection of contact is based on the values of γ^I which are obtained through the solution of a linear system of equations that couples all degrees of freedom. Hence, even if refinement is localized its effect is distributed over the entire potential

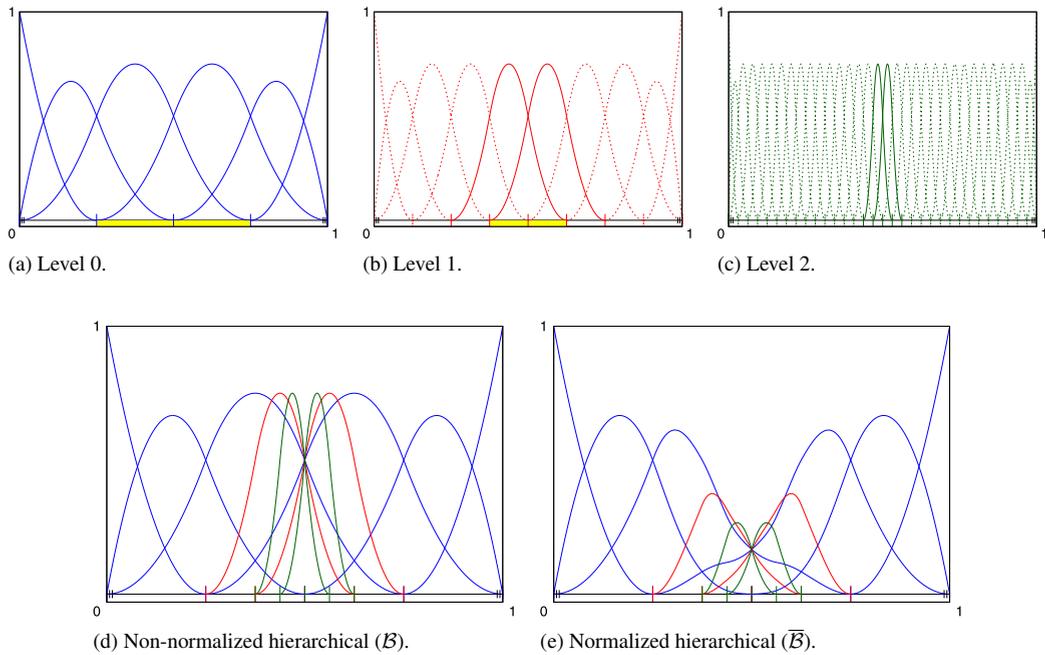


Fig. 20. The refinement at each level, the active basis functions and the hierarchical mortar basis in non-normalized (\mathcal{B}) and normalized forms ($\bar{\mathcal{B}}$) are depicted for the hierarchical \mathcal{N}^2 discretization of the slave body in Fig. 17.

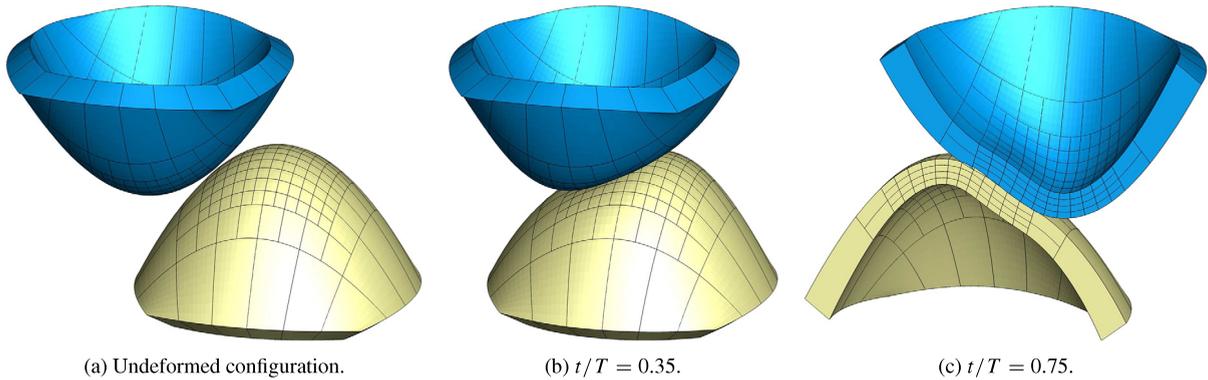


Fig. 21. Simulation instances from the analysis of Section 4.5.2 with a hierarchical discretization that is localized in the contact zone.

contact surface. On the other hand, each degree of freedom detects contact independently in the case of lumped recovery. Consequently, it is structurally more amenable to benefiting from local hierarchical refinement. Indeed, the solution improves with both approaches after refinement yet lumped recovery delivers a solution that is remarkably close to the fine mesh solution while full recovery is still influenced by regions of coarse discretization far from the contact zone.

5. Conclusion

In this work, hierarchical local refinement has been applied to frictionless contact problems using mortar-based discretizations in the context of isogeometric analysis with NURBS. The developments were carried out along three branches of investigations. First, two different numerical representations of the actual normal gap across the contact interface were introduced, namely a full recovery approach that is based on the solution of a linear system of equations in addition to a lumped recovery approach as its simplification. Second, two choices of hierarchical mortar basis sets were considered, one based on the hierarchical surface discretization of the slave body which does not satisfy the

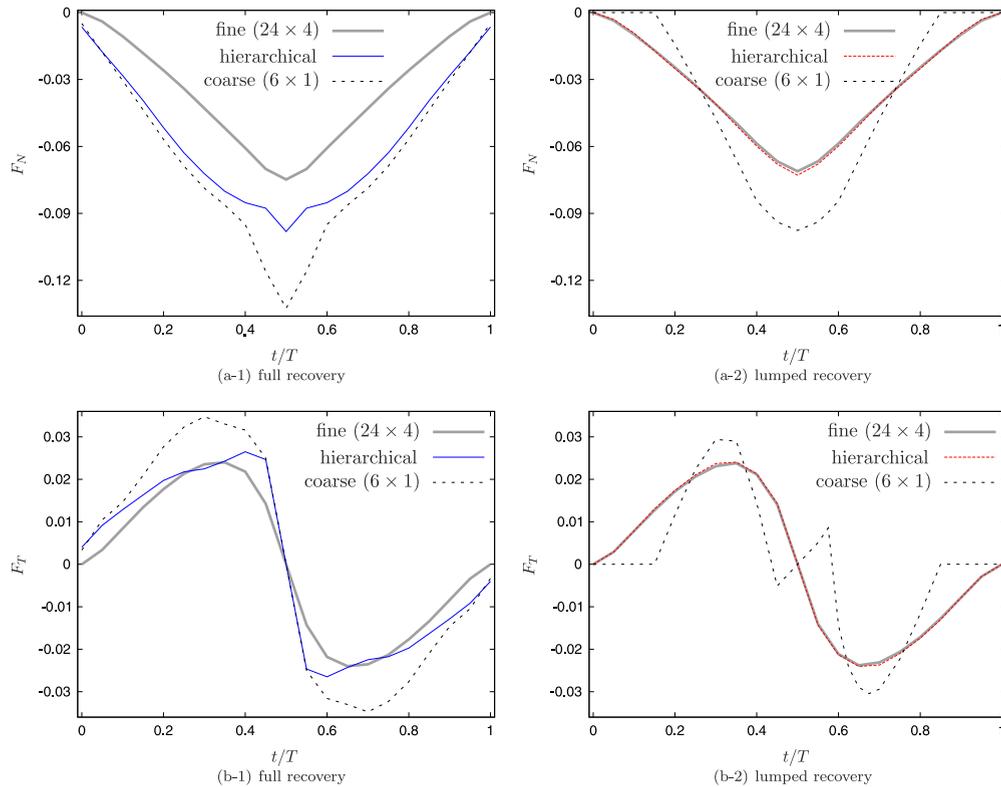


Fig. 22. The normal (F_N) and tangential (F_T) forces associated with the three-dimensional shell problem of Fig. 21 are summarized. The hierarchical solution locally refines the coarse mesh towards the fine mesh through two levels within the contact zone.

partition of unity and another based on the normalization of this set. Third, both exterior and interior point methods were applied for the enforcement of the contact constraints, including penalty and barrier type regularizations as inexact approaches as well as primal–dual algorithms which are exact. The investigations carried out within this third branch were complementary to both the first, in the sense that existing methods for lumped recovery were extended to the case of full recovery, as well as to the second, since the choice of the mortar basis set strongly influences the local and global solution quality. Overall, the results indicate the following:

- If hierarchical refinement is carried out towards uniform surface meshes then full recovery can display smoother force variations at the global level and more accurate pressure distributions at the local level.
- For local hierarchical refinement, full recovery should be accompanied by a normalized mortar basis but exact and inexact constraint enforcement methods are equally applicable. Lumped recovery, on the other hand, should preferably be accompanied by an exact enforcement method, but it may be employed with or without normalization.
- On coarse meshes, full and lumped recovery approaches both have their advantages and disadvantages. However, lumped recovery is more amenable to benefiting from local hierarchical refinement within the potential surfaces of contact such that the forces improve more rapidly with increasing levels of refinement in comparison to full recovery.

The investigations which have been carried out are comprehensive yet by no means exhaustive. Among other possibilities, the consideration of different hierarchical refinement algorithms, possibly in an adaptive fashion, and the incorporation of friction within a thermomechanical setting stand out as two immediate paths for further research. Problems which share similarities in the algorithmic basis for the treatment of interfaces, such as domain decomposition and cohesive zones, may also benefit from such investigations. Finally, it would be of interest to consider numerically more efficient alternatives to mortar schemes, such as point-to-surface or collocation type algorithms, for problems where computation speed is of major concern. Numerical efficiency may be particularly important for the dynamic contact of complex three-dimensional bodies, for which it may also be advantageous to

overcome the quadrilateral or hexahedral topology restriction as in [36] within the volume as well as on potential contact surfaces, the latter requiring possible modifications to the method of constructing the mortar basis. Such investigations will contribute to the efforts towards the numerically efficient and robust modeling of interfaces.

Acknowledgments

The first author acknowledges support by the European Commission under the project *MultiscaleFSI* (Grant No. PCIG10-GA-2011-303577) and also would like to thank M.A. Scott for very helpful early discussions on hierarchical refinement with NURBS.

References

- [1] J.A. Cottrell, T.J.R. Hughes, Y. Bazilevs, *Isogeometric Analysis*, Wiley, 2009.
- [2] L. De Lorenzis, P. Wriggers, T.J.R. Hughes, Isogeometric contact: a review, *GAMM-Mitt.* 37 (2014) 85–123.
- [3] J. Lu, I sogeometric contact analysis: Geometric basis and formulation of frictionless contact, *Comput. Methods Appl. Mech. Engrg.* 200 (2011) 726–741.
- [4] P. Papadopoulos, J.M. Solberg, A Lagrange multiplier method for the finite element solution of frictionless contact problems, *Math. Comput. Model.* 28 (1998) 373–384.
- [5] R.E. Jones, P. Papadopoulos, A novel three-dimensional contact finite element based on smooth pressure interpolations, *Internat. J. Numer. Methods Engrg.* 51 (2001) 791–811.
- [6] İ. Temizer, P. Wriggers, T.J.R. Hughes, Contact treatment in isogeometric analysis with NURBS, *Comput. Methods Appl. Mech. Engrg.* 200 (2011) 1100–1112.
- [7] L. De Lorenzis, İ. Temizer, P. Wriggers, G. Zavarise, A large deformation frictional contact formulation using NURBS-based isogeometric analysis, *Internat. J. Numer. Methods Engrg.* 87 (2011) 1278–1300.
- [8] İ. Temizer, P. Wriggers, T.J.R. Hughes, Three-dimensional mortar-based frictional contact treatment in isogeometric analysis with NURBS, *Comput. Methods Appl. Mech. Engrg.* 209–212 (2012) 115–128.
- [9] J.-Y. Kim, S.-K. Youn, Isogeometric contact analysis using mortar method, *Internat. J. Numer. Methods Engrg.* 89 (2012) 1559–1581.
- [10] M. Dittmann, M. Franke, İ. Temizer, C. Hesch, Isogeometric analysis and thermomechanical mortar contact algorithms, *Comput. Methods Appl. Mech. Engrg.* 274 (2014) 192–212.
- [11] P. Alart, A. Curnier, A mixed formulation for frictional contact problems prone to Newton like solution methods, *Comput. Methods Appl. Mech. Engrg.* 92 (1991) 353–375.
- [12] L. De Lorenzis, P. Wriggers, G. Zavarise, A mortar formulation for 3D large deformation contact using NURBS-based isogeometric analysis and the augmented Lagrangian method, *Comput. Mech.* 49 (2012) 1–20.
- [13] İ. Temizer, Multiscale thermomechanical contact: computational homogenization with isogeometric analysis, *Internat. J. Numer. Methods Engrg.* 97 (2014) 582–607.
- [14] M.A. Puso, T.A. Laursen, A mortar segment-to-segment contact method for large deformations, *Comput. Methods Appl. Mech. Engrg.* 193 (2004) 601–629.
- [15] M.A. Puso, T.A. Laursen, A mortar segment-to-segment frictional contact method for large deformations, *Comput. Methods Appl. Mech. Engrg.* 193 (2004) 4891–4913.
- [16] M.E. Matzen, T. Cichosz, M. Bischoff, A point to segment contact formulation for isogeometric, NURBS based finite elements, *Comput. Methods Appl. Mech. Engrg.* 255 (2013) 27–39.
- [17] M.H. Wright, The interior-point revolution in optimization: history, recent developments, and lasting consequences, *Bull. Amer. Math. Soc. (N.S.)* 42 (2005) 39–56.
- [18] İ. Temizer, M.M. Abdalla, Z. Gürdal, An interior point method for isogeometric contact, *Comput. Methods Appl. Mech. Engrg.* 276 (2014) 589–611.
- [19] L. De Lorenzis, J.A. Evans, T.J.R. Hughes, A. Reali, Isogeometric collocation: Neumann boundary conditions and contact, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 21–54.
- [20] F. Auricchio, L. Beirão Da Veiga, T.J.R. Hughes, A. Reali, G. Sangalli, Isogeometric collocation methods, *Math. Models Methods Appl. Sci.* 20 (2010) 2075–2107.
- [21] C. Hesch, P. Betsch, Isogeometric analysis and domain decomposition methods, *Comput. Methods Appl. Mech. Engrg.* 213–216 (2012) 104–112.
- [22] A. Apostolatos, R. Schmidt, R. Wüchner, K.-U. Bletzinger, A Nitsche-type formulation and comparison of the most common domain decomposition methods in isogeometric analysis, *Internat. J. Numer. Methods Engrg.* 97 (2014) 473–504.
- [23] C.V. Verhoosel, M.A. Scott, R. de Borst, T.J.R. Hughes, An isogeometric approach to cohesive zone modeling, *Internat. J. Numer. Methods Engrg.* 87 (2011) 336–360.
- [24] R. Dimitri, L. De Lorenzis, P. Wriggers, G. Zavarise, NURBS- and T-spline-based isogeometric cohesive zone modeling of interface debonding, *Comput. Mech.* 54 (2014) 369–388.
- [25] T.A. Laursen, *Computational Contact and Impact Mechanics*, first ed., Springer, Berlin, Heidelberg, New York, 2003, corr. second printing.
- [26] P. Wriggers, *Computational Contact Mechanics*, second ed., Springer, Berlin, Heidelberg, New York, 2006.
- [27] B. Wohlmuth, Variationally consistent discretization schemes and numerical algorithms for contact problems, *Acta Numer.* 20 (2011) 569–734.

- [28] M.R. Dörfel, B. Jüttler, B. Simeon, Adaptive isogeometric analysis by local h-refinement with T-splines, *Comput. Methods Appl. Mech. Engrg.* 199 (2010) 264–275.
- [29] Y. Bazilevs, V.M. Calo, J.A. Cottrell, J.A. Evans, T.J.R. Hughes, S. Lipton, M.A. Scott, T.W. Sederberg, Isogeometric analysis using T-splines, *Comput. Methods Appl. Mech. Engrg.* 199 (2010) 229–263.
- [30] E.J. Evans, M.A. Scott, X. Li, D.C. Thomas, hierarchical T-splines: analysis-suitability, Bézier extraction, and application as an adaptive basis for isogeometric analysis, *Comput. Methods Appl. Mech. Engrg.* 284 (2015) 1–20.
- [31] R. Dimitri, L. De Lorenzis, M.A. Scott, P. Wriggers, R.L. Taylor, G. Zavarise, Isogeometric large deformation frictionless contact using T-splines, *Comput. Methods Appl. Mech. Engrg.* 269 (2014) 394–414.
- [32] T.W. Sederberg, J. Zheng, A. Bakenov, A. Nasri, T-splines and T-NURCCSs, *ACM Trans. Graph.* 22 (2003) 477–484.
- [33] Y. Zhang, W. Wang, T.J.R. Hughes, Conformal solid T-spline construction from boundary T-spline representations, *Comput. Mech.* 51 (2013) 1051–1059.
- [34] A.-V. Vuong, C. Giannelli, B. Jüttler, B. Simeon, A hierarchical approach to adaptive local refinement in isogeometric analysis, *Comput. Methods Appl. Mech. Engrg.* 200 (2011) 3554–3567.
- [35] C. Giannelli, B. Jüttler, H. Speleers, THB-splines: The truncated basis for hierarchical splines, *Comput. Aided Geom. Design* 29 (2012) 485–498.
- [36] M.A. Scott, D.C. Thomas, E.J. Evans, Isogeometric spline forests, *Comput. Methods Appl. Mech. Engrg.* 269 (2014) 222–264.
- [37] K.A. Johannessen, T. Kvamsdal, T. Dokken, Isogeometric analysis using LB B-splines, *Comput. Methods Appl. Mech. Engrg.* 269 (2014) 471–514.
- [38] D. Schillinger, L. Dedè, M.A. Scott, J.A. Evans, M.J. Borden, E. Rank, T.J.R. Hughes, An isogeometric design-through-analysis methodology based on adaptive hierarchical refinement of NURBS, immersed boundary methods and T-spline CAD surfaces, *Comput. Methods Appl. Mech. Engrg.* 249–252 (2012) 116–150.
- [39] L. Piegl, W. Tiller, *The NURBS Book*, second ed., Springer, Berlin, Heidelberg, New York, 1996.
- [40] S. Hahmann, G.-P. Bonneau, A. Yvart, Hierarchical triangular splines, *ACM Trans. Graph.* 24 (2005) 1374–1391.
- [41] İ. Temizer, A mixed formulation of mortar-based frictionless contact, *Comput. Methods Appl. Mech. Engrg.* 223–224 (2012) 173–185.
- [42] İ. Temizer, A mixed formulation of mortar-based contact with friction, *Comput. Methods Appl. Mech. Engrg.* 255 (2013) 183–195.
- [43] P. Farah, A. Popp, W.A. Wall, Segment-based vs. element-based integration for mortar methods in computational contact mechanics, *Comput. Mech.* 55 (2015) 209–228.