



HAL
open science

Real-time simulation of contact and cutting of heterogeneous soft-tissues

Hadrien Courtecuisse, Jeremie Allard, Pierre Kerfriden, Stephane Pierre-Alain Bordas, Stephane Cotin, Christian Duriez

► **To cite this version:**

Hadrien Courtecuisse, Jeremie Allard, Pierre Kerfriden, Stephane Pierre-Alain Bordas, Stephane Cotin, et al.. Real-time simulation of contact and cutting of heterogeneous soft-tissues. *Medical Image Analysis*, 2014, 18 (2), pp.394-410. 10.1016/j.media.2013.11.001 . hal-01097108

HAL Id: hal-01097108

<https://inria.hal.science/hal-01097108>

Submitted on 6 Jan 2015

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

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

Published in final edited form as:

Med Image Anal. 2014 February ; 18(2): 394–410. doi:10.1016/j.media.2013.11.001.

Real-time simulation of contact and cutting of heterogeneous soft-tissues

Hadrien Courtecuisse^{a,b,c,*}, Jérémie Allard^c, Pierre Kerfriden^b, Stéphane P.A. Bordas^b, Stéphane Cotin^c, and Christian Duriez^c

^aICube / AVR, CNRS, Strasbourg, France

^bInstitute of Mechanics and Advanced Materials, Cardiff University, UK

^cSHACRA Project, INRIA, France

Abstract

This paper presents a numerical method for interactive (real-time) simulations, which considerably improves the accuracy of the response of heterogeneous soft-tissue models undergoing contact, cutting and other topological changes. We provide an integrated methodology able to deal both with the ill-conditioning issues associated with material heterogeneities, contact boundary conditions which are one of the main sources of inaccuracies, and cutting which is one of the most challenging issues in interactive simulations. Our approach is based on an implicit time integration of a non-linear finite element model. To enable real-time computations, we propose a new preconditioning technique, based on an asynchronous update at low frequency. The preconditioner is not only used to improve the computation of the deformation of the tissues, but also to simulate the contact response of homogeneous and heterogeneous bodies with the same accuracy. We also address the problem of cutting the heterogeneous structures and propose a method to update the preconditioner according to the topological modifications. Finally, we apply our approach to three challenging demonstrators: (i) a simulation of cataract surgery (ii) a simulation of laparoscopic hepatectomy (iii) a brain tumor surgery.

Keywords

Medical simulation; Contacts and collision detection; Interactive cutting; Heterogeneous structures; GPU parallelization

1 Introduction

Interactive simulation of surgical procedures could open avenues leading to improved patient care and reduced risks.¹ The main requirement to construct useful surgical simulators is to be able to reproduce the mechanical response of organs. To do so, surgical simulators involves at least four major challenges:

*Corresponding author at: ICube / AVR, Université de Strasbourg, France. hcourtecuisse@unistra.fr (H. Courtecuisse).

¹Surgical simulators are meant here as computer-based tools used to simulate a surgeon's intervention in a virtual environment.

1. The material properties of living tissues need to be characterized. Such properties are patient-specific and it is difficult predict the mechanical behavior of in vivo tissues.
2. The geometry of the organ models must be acquired from medical images, and this process is still not automatic, requiring difficult segmentation and mesh generation.
3. Significant numerical issues must be overcome, as most organs have an heterogeneous stiffness (which leads to ill-conditioned problems), are composed of multiple tissue types (which may lead to difficulties in defining boundary conditions) and are mostly incompressible (which leads to locking for most widely used finite elements).
4. The response must be computed in real-time to allow user interactions with the virtual body: pushing, prodding, palpation, needle insertion and cutting. In this sense, a simulation which would be mechanically realistic, but not interactive would not fit for our purpose.

Most of the previous works have focused on producing accurate models for the deformations of soft tissues, but real-time simulations are still usually composed of a single homogeneous organ with simple boundary conditions. Yet, we believe that it is also fundamental to take into account the deformable environment of the tissues to obtain a realistic global behavior. For instance, even if during a hepatectomy the view of the camera is focused on a small part of the liver tissues (see Fig. 1) the human body is composed of multiple organs playing an important role in the resulting deformation and motion of the liver (breathing, contact with neighboring organs, etc.). Indeed, if we look closely at the boundary conditions of the liver: the upper part is compressed by the diaphragm on which it is attached through several ligaments including the falciform ligament. Moreover, the liver is in contact with the vena cava and the stomach and linked to it by the hepatogastric ligament. Therefore, when studying the motion and deformation, the liver cannot be considered as an isolated organ.

Beyond the importance of boundary conditions, our additional working hypothesis is that realistic surgical simulation also requires to interactively compute significant topological modifications of the organs, e.g. cutting. In this case, boundary conditions are also fundamental because stiff interactions with the tool (controlled by the user) as well as complex deformable–deformable interactions between the lips of the cut must be solved simultaneously. Performing this at interactive rate is difficult, in particular for heterogeneous structures.

The contribution of this paper is to propose a consistent framework to address the aforementioned requirements of surgical simulations. This framework relies on an implicit time integration of the non-linear set of equations coming from the finite element model of the deformation. The core of the method rely on a preconditioning technique which is updated asynchronously at low frequency. This preconditioner reduces the convergence issues appearing when computing the deformation of non-linear heterogeneous structures. In addition, it provides a very good estimate of the compliance² operator associated with the coupling between the solids in contact. We also extend the method to cut these

heterogeneous organs at an interactive rate whilst resolving existing or new contact surfaces emanating from these topological modifications. Our solution is based on the Sherman Morison Formula to update the asynchronous preconditioner in case of topological modifications. Finally the generic nature of the method is demonstrated by three different types of applications: a simulation of cataract surgery, a simulation of a hepatectomy using laparoscopy procedure, and a simulation of cerebral tumor removal.

The paper is organized as follows: In Section 2, we briefly review the real-time simulation of deforming bodies. In Section 3 we define the problem at hand and provide notations and key concepts necessary to build our method. Section 4 is dedicated to the description of the asynchronous preconditioner. In Section 5 we propose to use the asynchronous preconditioner to solve the constraints associated with the multi-deformable-body contact problem under consideration. This section also provides details on the GPU implementation to achieve real-time results. In Section 6, we extend the method to handle topological modifications. In Section 7, we evaluate our method in terms of accuracy and computation time. Finally, in Section 8 we exercise the methodology in three practical problems involving complex heterogeneous structures in interaction.

2 Literature review

Biomechanical simulation with user interactions involves many challenges such as real-time computation of the deformation of soft tissues, collision detection, contact modeling, topological modification, and haptic feedback (see Nealen et al. (2006) and Payan (2012) for a broad survey).

2.1 Simulation of deformable bodies

The first methods proposed to simulate the deformation of soft tissues in real-time relied on mass-spring systems, e.g. Kühnapfel et al. (2000). Although such discrete methods are simple to implement and very fast, they are difficult to parameterize with material properties such as the Young's modulus. Moreover, they introduce anisotropy through the choice of the mesh giving rise to stability and accuracy issues (node flipping, difficulty to preserve the volume).

Finite element methods (FEM) provide high bio-mechanical realism (Zienkiewicz and Taylor, 1991), mainly because the complex non-linear behavior of soft-tissue is directly accounted for through constitutive relations. Real-time computations were first achieved for linear elastic material models (see Bro-Nielsen and Cotin (1996), Cotin et al. (1999) or James and Pai (1999)). In linear elasticity, precomputations (offline) can be used to accelerate the online simulations. However, the small strain assumption is incorrect in practice, and produces erroneous results when the solids undergo large deformations. The use of precomputed solutions for highly non-linear problems is intensively pursued, e.g. in Niroomandi et al. (2008) for hyper elasticity, and Kerfriden et al. (2011, 2012) for damage problems.

²Inverse of the stiffness.

The *co-rotational* method is an old method, which originated in continuum mechanics (Freund, 1970) and was introduced by Felippa (2000) within the field of numerical methods. In this formulation, the stiffness of each element is assumed linear within the local frame described by its rotated state, which enables to simulate geometric non-linearities (i.e. large displacements, large rotations but small strains). This enables to produce realistic simulations, while maintaining the algorithmic complexity at a minimum (Felippa and Haugen, 2005; Müller and Gross, 2004).

Later, other methods were proposed to simulate both geometric and material non-linearities. Material non-linearities are expressed through the non-linearity of the constitutive law relating the strain tensor to the stress tensor, characteristic of soft tissues. Real time models were recently proposed (Comas et al., 2008; Joldes et al., 2009; Marchesseau et al., 2010), but these models remain in general complex and expensive, and the simulation of realistic boundary conditions such as interactions between deformable organs and surgical instruments is still an issue.

2.2 Time discretization

Explicit integration schemes are widely adopted (Miller et al., 2007; Taylor et al., 2008) in surgical simulation. The main advantage is that the solution process only involves the mass matrix, which can be lumped (diagonalised). The equations of motion are thus decoupled and each degree of freedom can be solved independently, making the solution process very fast, and inherently well-suited to parallelization (Comas et al., 2008). The major drawback of explicit dynamics is the need to satisfy the Courant–Friedrichs–Lewy stability condition, which forces a strict upper bound on the time step used for integration. Explicit methods are consequently particularly well-suited to very soft tissues such as the brain (Joldes et al., 2009), but very small time steps (which prevent real-time computations) must be chosen for stiffer structures. Moreover, explicit simulations do not guarantee that, at each time step, the residual vector is minimized, and hence, that the external and internal forces balance.

These are the major reasons for developing implicit time integration techniques for real time simulations. Providing flexibility in the choice of the time step, even for very stiff objects (Baraff and Witkin, 1998), is required to achieve our aim: simulate user-controlled interactions between arbitrarily stiff anatomical structures or tools. Of course, the advantages of such methods come at the cost of having to solve a set of linear equations at each time step. Yet, this paper aims to show that implicit integration schemes can offer a reasonable tradeoff between robustness, stability, convergence and computation time, in particular when combined with a GPU implementation.

2.3 Solving the set of non-linear equations

Using implicit integration, a non-linear set of equations must be solved at each load step. This set of equations is usually solved using an iterative method based on the Newton–Raphson method which solves the set of non-linear equations through a sequence of solutions of linear equations.

The set of linear equations can either be solved by direct solvers or iterative solvers. Direct solvers provide the solution by computing the actual inverse of the system matrix (Bro-

Nielsen and Cotin, 1996), or by creating a factorization that can then be used to compute the solution (Barbič and James, 2005).

These methods are often too costly to be applied at each iteration, and are often used in combination with an approach to reduce the number of degrees of freedom of the model, either using condensation on surface nodes (Bro-Nielsen and Cotin, 1996), or reduced-coordinate models (Barbič and James, 2005). Recently, Hecht et al. (2012) proposed a method to incrementally update a sparse Cholesky factorization. They obtain fast performance by making only partial changes to the simulation's linearized system matrices, but the method is closely related to the co-rotational formulation and cannot take into account topological modifications.

The second class of methods are iterative (Saad, 1996), and start from an initial estimate and iteratively refine it to approach the exact solution. One of the most popular methods is the Conjugate Gradient (CG) algorithm. Although in theory up to n iterations are necessary to achieve convergence for n equations in the system, in practice it is possible to stop the algorithm much earlier depending on the required accuracy. Parallel implementations on CPU are now well-mastered and optimized: see for example Parker and O'Brien (2009) and Hermann et al. (2009) and start to appear on GPU (Bolz et al., 2003; Buatois et al., 2009; Allard et al., 2011). Iterative solvers are usually faster than direct methods, and require less memory storage, but they converge slowly for ill-conditioned problems, i.e. when the ratio of the largest and smallest eigenvalues is large. This is the case when solving linear systems of equations associated with the discretization of heterogeneous structures, as is the case in this paper.

Another intense area of research aims to improve the performance of the CG algorithm with the use of preconditioners to speed-up its convergence. In the context of interactive simulation, Baraff and Witkin (1998) proposed to use a diagonal inverse, often called a Jacobi preconditioner. More advanced preconditioners such as Cholesky factorizations have also been studied (Hauth et al., 2003). However, the performance improvement remains limited since the preconditioner itself is expensive to compute. Domain decomposition preconditioners are popular, for example in multi-scale parallel simulations (Dryja and Widlund, 1989) and for extended (enriched) finite element methods (Menk and Bordas, 2011). Multi-grid pre conditioners, e.g. Braess (1986) were developed for real-time simulations for example in Dick et al. (2011) and to simulate cuts in deformable objects in Dick et al. (2011) and for fracture problems in Hiriyur et al. (2012), Berger-Vergiat et al. (2012), and Gerstenberger and Tuminaro (2013).

2.4 Simulation of the interactions

A key requirement for realistic surgical simulators is to treat contact between “soft”–“soft”, “soft”–“stiff” and “stiff”–“stiff” objects. A common solution to deal with contact consists of using a penalty method, which modifies the variational principle and solves the contact condition approximately. A didactic review of constraint enforcement in a variational context is provided in Hughes (2000) and a review of error estimates associated with this enforcement of variational inequalities in finite element methods can be found in the seminal paper Brezzi et al. (1977). In penalty methods, a contact force $f = \alpha \delta n$ is added at each

contact point, where δ is a measure of the interpenetration, \mathbf{n} is the contact normal and α is a stiffness factor known as penalty coefficient. The higher α the better the constraint is satisfied. However, the higher the value of α , the worse the condition number of the system is, the higher the spurious oscillations in contact forces, and the smaller the time step. The selection of α is also problem-dependent, and depends particularly strongly on the stiffness ratio between the contacting objects. Consequently, the penalty method is limited for the applications that we have in mind, i.e. contacting heterogeneous objects.

Lagrange multipliers or augmented Lagrangian techniques (Hughes, 2000; Renard, 2013; Jean, 1999; Jourdan et al., 1998; Wriggers and Panatitopoulos, 1999) are usually preferred to penalty methods to treat contact constraints accurately and robustly. Methods used to solve contact equations with Lagrange multipliers can be classified into two categories which are numerically equivalent: *Quadratic Programming* (QP) methods and *Complementary Problem* methods that could be linear (LCP) or non-linear (NLCP).

QP methods define the constraints directly into the mechanical system. They can be used to treat the inequality of the contact (Redon et al., 2002; Pauly et al., 2004), and also to simulate friction using a discretized pyramidal cone (Kaufman et al., 2008). However, these publications address the case of rigid bodies in contact, where the number of degree of freedom (DOF) is smaller than the number of contact constraint freedoms. Indeed, the resulting number of equations with QP methods is of the same order of magnitude as the number of DOFs of the interacting objects. Therefore, these methods are difficult to adapt to the simulation of the interaction between finely meshed deformable bodies in real-time.

An important advantage of (N)LCP methods is that the number of constraint equations is proportional to the number of contacts, which is often much smaller than the number of DOF in the context of deformable models. LCP can be used to simulate frictionless contact between deformable models (Duriez et al., 2003) in real-time whereas NLCP methods can be used to simulate friction contacts with the exact friction cone (Duriez et al., 2006). The main limitation of these methods is that the solution process involves the *compliance matrix*, which is the inverse of a large system composed of the mass, the damping and the stiffness of the deformable objects. Although the evaluation of this inverse in real-time is crucial to define the boundary conditions of the deformable structures, very few methods addressed this issue. Duriez et al. (2004) proposed to precompute the *compliance matrix*, but the solution is limited to linear elastic deformation. Otaduy et al. (2009), proposed to compute the compliance matrix using additional Gauss–Seidel iterations on the deformable models, but the method was not presented in a real-time context. Saupin et al. (2008) proposed a method, named *compliance warping*, which is dedicated to co-rotational models. It consists of pre-computing the *compliance matrix* from the rest position, and updating it using a local estimation of the nodal rotations. However, this approximation can become inaccurate for large deformations, and the method is limited to relatively coarse meshes. A prediction–correction scheme is introduced in Peterlik et al. (2011) to mitigate the inaccuracies and the method extends the formulation to haptic feedback with generic constraints between the deformable models.

To our knowledge, it was never demonstrated that any of the above methods are adequate to enforce contact constraints between geometrically complex deformable objects represented by moderately fine meshes (thousands of DOFs), both accurately and at interactive rates (30–50 frames per second). This problem is particularly relevant for medical applications, since all the organs in a human body are subjected to interactions from other tissues. Resolving accurately these interactions between organs can be construed as more critical than the accuracy of the deformation of each individual organ itself. An important contribution of this article is to propose a generic solution to address this problem.

2.5 Simulation of topological modifications

When cutting through a finite element mesh, discontinuities in the displacement field must be introduced. This can be done through local re-meshing. Several approaches were proposed to maintain a relatively good mesh quality (Ganovelli and O’Sullivan, 2001; Bielser et al., 2003; Molino et al., 2007; Sifakis et al., 2007). One important difficulty is to preserve the quality and the density of the mesh during the subdivision process, in order to avoid distorted elements, which lead to convergence difficulties during the simulation. In addition, few methods were proposed to handle collision detection with the modified topology (where the quality and the density of the triangulation of the surface around the cut is difficult to control).

For the mechanical aspect, topological changes require updating of the stiffness matrix. As most recent real-time deformation methods are based on non-linear models (geometrical and/or material), topological changes add little overhead to the process. The treatment of the interactions during the cut is also problematic. Courtecuisse et al. (2010b) proposed to extend the *compliance warping* technique using the Sherman Morrison formula to update the precomputed inverse in case of topological modifications. However, numerical errors accumulate over time and the method suffers from instabilities after a large number of modifications. In this paper, we propose a solution that enables to significantly decrease the numerical issues caused by the cuts, and to preserve the accuracy of the contact response during the cut.

Alternatives to generating finite element meshes, in particular to handle topological changes are extended finite element methods (Nicolas et al., 1999). These methods allow cuts, material interfaces, and domain boundaries to be described independently of a background mesh, which may also be progressively adapted using a posteriori error estimators (Bordas and Duflot, 2007; Duflot and Bordas, 2008; Bordas et al., 2008; Ródenas et al., 2008) or local heuristics (Menk and Bordas, 2011). However, such methods are not yet developed in the real-time context, and the simulation of collisions and interactions remains unsolved, in particular because the geometry of the discontinuities is known implicitly.

3 Background

In this section we introduce some of the necessary background on which we build our method.

3.1 Deformable model and time-stepping implicit integration

Let us consider a generic dynamic deformable model. Equations used to model the dynamic behavior of bodies can be written within a synthetic formulation, given by Newton's second law:

$$\mathbb{M}(\mathbf{q})\ddot{\mathbf{q}} = \mathbb{P}(t) - \mathbb{F}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbb{H}(\mathbf{q})^T \boldsymbol{\lambda} \quad (1)$$

where $\mathbf{q} \in \mathbb{R}^n$ is the vector of generalized degrees of freedom (here, the mesh node positions), $\mathbb{M}(\mathbf{q}) : \mathbb{R}^n \mapsto \mathcal{M}^{n \times n}$ is the inertia matrix. \mathbb{F} represents internal forces applied to the simulated object depending on the current state and \mathbb{P} gathers external forces. $\mathbb{H}(\mathbf{q})^T$ is a function that gives the constraint directions depending on the position of the objects, and $\boldsymbol{\lambda}$ is the associated the vector of Lagrange multipliers containing the constraint force intensities (see details below).

$\mathbb{M}(\mathbf{q})$ and $\mathbb{F}(\mathbf{q}, \dot{\mathbf{q}})$ are derived from the physics-based deformable model. In this paper, we use the *co-rotational* formulation as a tradeoff between accuracy and computation time. In this formulation, the stiffness of the material depends on the current rotation (and thus on the current positions), which results in a geometrically non-linear elastic formulation (Felippa, 2000). After discretization, the mass matrix noted \mathbf{M} in the following, is considered as constant and lumped (we obtain a diagonal matrix).

Collision response on mechanical objects leads to discontinuities in the velocities of the colliding points. For such discontinuous events, the acceleration is not defined: the problem belongs to the field of non-smooth mechanics. To integrate the mechanics and the non-smooth events due to contact over time, we use a *time-stepping* method based on an implicit scheme: The time step is fixed and there is no limitation on the number of discontinuities that could take place during a time step (Anitescu et al., 1999), but low-order integration schemes should be used. This could lead to excessive dissipation if the time step is too large but it provides stable simulations. This is particularly relevant for interactive simulations involving contact with virtual devices controlled by an operator. With or without haptic feedback, the motion of the user will not be completely constrained and could lead to excessive energy input in the simulation. This is why the stability and the robustness of these types of simulations are crucial.

Let us consider the time interval $[t_i, t_f]$ which length is $h = t_f - t_i$. We have:

$$\begin{aligned} \mathbf{M}(\dot{\mathbf{q}}_f - \dot{\mathbf{q}}_i) &= \int_{t_i}^{t_f} (\mathbb{P}(t) - \mathbb{F}(\mathbf{q}, \dot{\mathbf{q}})) dt + h \mathbb{H}(\mathbf{q})^T \boldsymbol{\lambda} \\ \mathbf{q}_f &= \mathbf{q}_i + \int_{t_i}^{t_f} \dot{\mathbf{q}} dt \end{aligned} \quad (2)$$

To evaluate integrals $\int_{t_i}^{t_f} (\mathbb{P}(t) - \mathbb{F}(\mathbf{q}, \dot{\mathbf{q}}, t)) dt$ and $\int_{t_i}^{t_f} \dot{\mathbf{q}} dt$ we chose the following implicit Euler integration scheme:

$$\begin{aligned} \mathbf{M}(\dot{\mathbf{q}}_f - \dot{\mathbf{q}}_i) &= h(\mathbb{P}(t_f) - \mathbb{F}(\mathbf{q}_f, \dot{\mathbf{q}}_f)) + h \mathbb{H}(\mathbf{q})^T \boldsymbol{\lambda}_f \\ \mathbf{q}_f &= \mathbf{q}_i + h\dot{\mathbf{q}}_f \end{aligned} \quad (3)$$

\mathbb{F} is a non-linear function, we apply a Taylor series expansion to \mathbb{F} and make a first order approximation:

$$\mathbb{F}(\mathbf{q}_i + d\mathbf{q}, \dot{\mathbf{q}}_i + d\dot{\mathbf{q}}) = f_i + \frac{\delta \mathbb{F}}{\delta \mathbf{q}} d\mathbf{q} + \frac{\delta \mathbb{F}}{\delta \dot{\mathbf{q}}} d\dot{\mathbf{q}} \quad (4)$$

This linearization is actually the first iteration of the Newton–Raphson algorithm. This single iteration is done under the assumption of a temporal coherency of the mechanical behavior; it may lead to small numerical errors in the dynamic behavior, but these errors tend to decrease at equilibrium or with null velocity. After discretization $\mathbf{B} = \frac{\delta \mathbb{F}}{\delta \dot{\mathbf{q}}}$ and $\mathbf{K} = \frac{\delta \mathbb{F}}{\delta \mathbf{q}}$ are known respectively as the damping and stiffness matrices. Replacing (4) in (3) and using $d\mathbf{q} = \mathbf{q}_f - \mathbf{q}_i = h\dot{\mathbf{q}}_f$ and $d\dot{\mathbf{q}} = \dot{\mathbf{q}}_f - \dot{\mathbf{q}}_i$, we obtain:

$$\underbrace{(\mathbf{M} + h\mathbf{B} + h^2\mathbf{K})}_{\mathbf{A}} \underbrace{d\dot{\mathbf{q}}}_{\mathbf{x}} = \underbrace{-h^2\mathbf{K}\dot{\mathbf{q}}_i - h(f_i + \mathbf{p}_f)}_{\mathbf{b}} + h\mathbb{H}(\mathbf{q})^T \boldsymbol{\lambda}_f \quad (5)$$

where \mathbf{p}_f is the value of function \mathbb{P} at time t_f . Note that using elastic or hyperelastic FEM of deformations, matrix \mathbf{A} is symmetric definite positive. The only remaining unknown values are the Lagrange multipliers $\boldsymbol{\lambda}$ but their computation is now detailed.

3.2 Contact and friction models

Before enforcing the contact between soft tissues or with surgical instruments (rigid or deformable), one needs to detect them. In our work, we use two types of algorithms: the first is based on proximity queries,³ and provides the minimal distances between mesh (even concave meshes); the second is based on detection of volume of interpenetration (details in Allard et al. (2010)). The first algorithm have the advantage to “anticipate” the contacts before they actually appears, the second algorithm needs unnatural interpenetration between models but provide much faster results on complex meshes thanks to GPU optimizations. In the following, we consider that we use the algorithm based on proximities detection to enforce contact between point pairs at the surface of the colliding models. The formulation used for enforcing volume constraints is similar and detailed in Allard et al. (2010).

³The implementation is available on SOFA using the component *LocalMinDistance*.

Whatever method is used for detection, the collision response is based on the same model: Signorini's law. This model is defined for each *potential* contact provided by the contact detection algorithm. It expresses that there is a complementarity relation along the direction \mathbf{n} of contact⁴ between the contact force λ_n and the distance δ_n between the contacting object at the point of contact, that is:

$$\mathbf{0} \leq \delta_n \perp \lambda_n \geq \mathbf{0} \quad (6)$$

This model has several physical justifications including non interpenetration and no sticking force. Moreover the contact force vanishes if the points are not strictly in contact. Using Signorini's law, the contact space is along the normal and creates a frictionless response.

Coulomb's friction law describes the macroscopic behavior in the tangential contact space. In this law, the reaction force is included in a cone whose height and direction is given by the normal force. If the reaction force is strictly included inside the cone, objects stick together, otherwise, the reaction force is on the cone's border and objects are slipping along the tangential direction. In this last case, the friction force must be directed along the direction of motion.

$$\begin{aligned} \dot{\delta}_T = \mathbf{0} &\Rightarrow \|\lambda_T\| < \mu \|\lambda_n\| \text{ (stick)} \\ \dot{\delta}_T \neq \mathbf{0} &\Rightarrow \lambda_T = -\mu \|\lambda_n\| \frac{\dot{\delta}_T}{\|\dot{\delta}_T\|} = -\mu \|\lambda_n\| \mathbf{T} \text{ (slip)} \end{aligned} \quad (7)$$

where μ is the friction parameter, and \mathbf{T} is the direction of motion in the tangential plane to the contact normal \mathbf{n} .

During 3D slipping motion (also called *dynamic friction*), the tangential direction is unknown. We only know that the tangential force and the tangential velocity are opposite along a direction that is to be found. It creates a non-linearity in addition to the complementarity state stick/slip. Signorini's law and Coulomb's law are also valid in a *multi-contact* case. However, to solve these laws at every contact point, we have to consider the coupling that exists between these contact points. This coupling comes from the intrinsic mechanical behavior of deformable objects.

3.3 Contact mapping

From collision or proximity detection, we have a set of potential contact spots $\alpha = 1, \dots, n_c$ which are defined on the surface of the deformable bodies (triangles, lines, or points). In order to transfer the contact informations to the Degrees of Freedom of the objects, we can build a mapping function \mathbb{A} that links the positions in the contact space to the motion space (see Saupin et al. (2008) for details). For each contact point between two objects:

⁴Complementarity is noted \perp . It states that one of the two values δ_n or λ_n must be null.

$$\delta_\alpha = \mathbb{A}_\alpha(\mathbf{q}_1, t) - \mathbb{A}_\alpha(\mathbf{q}_2, t) \quad (8)$$

with $\mathbb{A}_\alpha(\mathbf{q}, t)$ the mapping function which depends on contact α and the positions \mathbf{q}_1 and \mathbf{q}_2 of the two colliding objects. To obtain a kinematic relation between the two spaces (contact, motion), we use a linearization of Eq. (8). If $\mathbb{H}_\alpha(\mathbf{q}) = \frac{\partial \mathbb{A}_\alpha}{\partial \mathbf{q}}$, we obtain, at time t for each contact:

$$\dot{\delta}_\alpha(t) = \mathbb{H}_\alpha(\mathbf{q}_1)\dot{\mathbf{q}}_1(t) - \mathbb{H}_\alpha(\mathbf{q}_2)\dot{\mathbf{q}}_2(t) \quad (9)$$

where \mathbb{H} is a non-linear function of the position \mathbf{q} . The dimensions of this jacobian matrix is the number of constraints for the rows and the size of \mathbf{q} for the columns. To simplify the solution process, we suppose that this matrix does not change during the contact response $\mathbf{H}^T \boldsymbol{\lambda}_f = \mathbb{H}(\mathbf{q})^T \boldsymbol{\lambda}_f$. In the following, this matrix is noted \mathbf{H} to emphasize that it is constant during the time step.

3.4 Constraint-based solution

In the following, we present how the contact laws (6) and friction laws (7) are enforced while taking into account the dynamic Eq. (5) between 2 contacting objects. To resolve these laws, we use a Lagrange Multiplier approach and a single linearization by time step. For both interacting objects we applied Eq. (5):

$$\begin{aligned} \mathbf{A}_1 \mathbf{x}_1 &= \mathbf{b}_1 + h \mathbf{H}_1^T \boldsymbol{\lambda} \\ \mathbf{A}_2 \mathbf{x}_2 &= \mathbf{b}_2 + h \mathbf{H}_2^T \boldsymbol{\lambda} \end{aligned} \quad (10)$$

In order to solve $\boldsymbol{\lambda}$ we follow the following steps.

Step 1: interacting objects are solved independently while setting $\boldsymbol{\lambda} = \mathbf{0}$. A set of independent linear systems of equations $\mathbf{A}\mathbf{x} = \mathbf{b}$ must then be solved for each object. For our purposes, this must be done in real-time and we use a GPU-based Conjugate Gradient algorithm (CG) as described in Allard et al. (2011). This method enables using meshes comprised of thousands of elements in real-time on standard architectures, but tends to converge very slowly for ill-conditioned matrices. This is particularly an issue for the simulation of heterogeneous materials with large phase contrast or for problems obtained from finite element meshes of non optimal quality (large local difference in element size). We propose in Section 4 a novel method to simulate the heterogeneities in real-time. Finally, we obtain what we call the *free motion* $\mathbf{x}_1^{\text{free}}$ and $\mathbf{x}_2^{\text{free}}$ for each object which corresponds to $d\mathbf{q}_1^{\text{free}}$ and $d\mathbf{q}_2^{\text{free}}$. After integration, we obtain $\mathbf{q}_1^{\text{free}}$ and $\mathbf{q}_2^{\text{free}}$.

Step 2: the constraint laws are linearized during the time-step. We process a proximity collision detection between the position of the objects at the previous time step, and using the linearization we obtain the free interpenetration δ^{free} and the associated \mathbf{H} , assumed constant during the time step:

$$\delta = \underbrace{\mathbb{A}_\alpha(\mathbf{q}_1^{\text{free}}) - \mathbb{A}_\alpha(\mathbf{q}_2^{\text{free}})}_{\delta^{\text{free}}} + h\mathbf{H}_1 d\dot{\mathbf{q}}_1^{\text{cor}} + h\mathbf{H}_2 d\dot{\mathbf{q}}_2^{\text{cor}} \quad (11)$$

with $d\dot{\mathbf{q}}_1^{\text{cor}}$ and $d\dot{\mathbf{q}}_2^{\text{cor}}$ being the unknown corrective motions when solving Eq. (10) with $\mathbf{b}_1 = \mathbf{b}_2 = \mathbf{0}$. When gathering Eqs. (10) and (11), we have:

$$\delta = \delta^{\text{free}} + h^2 \underbrace{[\mathbf{H}_1 \mathbf{A}_1^{-1} \mathbf{H}_1^T + \mathbf{H}_2 \mathbf{A}_2^{-1} \mathbf{H}_2^T]}_{\mathbf{w}} \lambda \quad (12)$$

With respect of the Signorini's law (Eq. (6)), this equation describes a LCP (Linear Complementarity Problem). If it is combined with Coulomb's law (Eq. (7)), we obtain a NLCP (Non-linear complementarity problem). This equation implies to evaluate the inverse of large matrices \mathbf{A}_1 and \mathbf{A}_2 (same dimension as the number of DOFs). We propose in Section 5 a novel method to obtain an approximation in real-time.

Step 3: We obtain the value of λ using a Gauss–Seidel algorithm dedicated to the NLCP created by contact and friction equations. Considering a contact α , among m instantaneous contacts, one can write the behavior of the model in contact space:

$$\underbrace{\delta_\alpha - \mathbf{W}_{\alpha\alpha} \lambda_\alpha}_{\text{unknown}} = \underbrace{\sum_{\beta=1}^{\alpha-1} \mathbf{W}_{\alpha\beta} \lambda_\beta + \sum_{\beta=\alpha+1}^m \mathbf{W}_{\alpha\beta} \lambda_\beta}_{\text{frozen}} + \delta_\alpha^{\text{free}} \quad (13)$$

where $\mathbf{W}_{\alpha,\beta}$ is a compliance matrix that models the coupling between contact points α and β . For each contact α , this method solves the contact equations by considering the others contact points ($\alpha \neq \beta$) as “frozen”. The new value of λ_α is given by solving Signorini's law and the Coulomb's law on this contact (see Duriez et al. (2006) for details of implementation and performance).

Step 4: When the value of λ is available, the corrective motion is computed:

$$\begin{aligned} \mathbf{q}_{1,t+h} &= \mathbf{q}_1^{\text{free}} + h \Delta \dot{\mathbf{q}}_1^{\text{cor}} \text{ with } \Delta \dot{\mathbf{q}}_1^{\text{cor}} = \mathbf{A}_1^{-1} \mathbf{H}_1^T \lambda \\ \mathbf{q}_{2,t+h} &= \mathbf{q}_2^{\text{free}} + h \Delta \dot{\mathbf{q}}_2^{\text{cor}} \text{ with } \Delta \dot{\mathbf{q}}_2^{\text{cor}} = \mathbf{A}_2^{-1} \mathbf{H}_2^T \lambda \end{aligned} \quad (14)$$

We finally obtain $\mathbf{q}_{1,t+h}$ and $\mathbf{q}_{2,t+h}$, the positions of object 1 and 2 that fullfils the contact and friction laws.

In this section, we identified two important difficulties to achieve real-time computations. (i) To compute the solution of the linear system defined in Eq. (5), in particular for heterogeneous objects. (ii) To compute the *compliance matrix* \mathbf{W} in Eq. (12).

These two problems are addressed in the two following sections.

4 Asynchronous preconditioner

The condition number k of the matrix \mathbf{A} , measures how much the output value of the function can change for a small change in the input argument. For heterogeneous objects or ill-structured meshes, the condition number k is often high which raises convergence issues for the conjugate gradient algorithm used to solve Eq. (5). A common technique is to use a preconditioner to reduce the condition number, ensuring a faster convergence of the algorithm. By definition, a preconditioner is an approximation of the system matrix \mathbf{A} , which is less costly to invert. Solving Eq. (5) with a preconditioner \mathbf{P} can be written:

$$\mathbf{P}^{-1}\mathbf{A}\mathbf{x} = \mathbf{P}^{-1}\mathbf{b}, \quad \text{such that } k(\mathbf{P}^{-1}\mathbf{A}) < k(\mathbf{A}) \quad (15)$$

In the real-time context, one strong limitation of this technique is the computational overhead added to the simulation: first, during the computation of the preconditioner itself, and second, during its use at each iteration of the CG (see Saad (2003) for details). Thus, the practical usefulness of preconditioners depends on the ability to strike a balance between the computational overheads and the time saved by decreasing the number of CG iterations. Several preconditioners can be used, from simple diagonal matrices (Baraff and Witkin, 1998) to precise but costly Cholesky factorizations.

We recently proposed a different approach (Courtecuisse et al., 2010a) that relies on the assumption that \mathbf{A} undergoes small perturbations between two consecutive time steps. Indeed, if $\mathbf{P}_t = \mathbf{A}_t^{-1}$ is available at a specific time t , it may remain a “good enough” approximation for the following time steps. The preconditioner can then be updated at low frequency on a dedicated CPU thread, and the last preconditioner available can be used to advance the simulation (see Fig. 2). Therefore, the overhead in computing the preconditioner is removed from the simulation loop, which allows using more advanced and computationally costly preconditioners such as a factorization of the system⁵:

$$\mathbf{P} = \mathbf{A} = \mathbf{LDL}^T \quad (16)$$

where \mathbf{D} is a diagonal matrix and \mathbf{L} is a sparse lower-triangular matrix. In our application, we rely on \mathbf{LDL}^T factorization since it produce more stable results than a Cholesky factorization. The factorization is performed by the *cs_sparse* library (Davis, 2006), using a single core on the CPU. Other libraries Toledo et al. (2003) and Schenk et al. (2008) propose

⁵Note that even if we compute an exact factorization of \mathbf{A}_{t-J} , the preconditioner remains an approximation since its computation is based on a previous configuration of the objects, and we use it with delay in the simulation.

parallel factorizations, but we found that *cs_sparse* provides sufficiently fast updates, and a sequential factorization enables to save CPU cores to compute the preconditioners of other objects in parallel.

An important advantage of this factorization is that the resulting \mathbf{L} matrix remains sparse, which makes the application of the preconditioner faster within the CG. This operation consists in solving two Sparse Triangular Systems (STS):

$$\begin{aligned} \mathbf{y} &= (\mathbf{L}^T)^{-1} \mathbf{b} \\ \mathbf{x} &= \mathbf{L}^{-1} (\mathbf{D}^{-1} \mathbf{y}) \end{aligned} \quad (17)$$

where \mathbf{L} is stored in *Compressed Row Storage* (CRS) (see Barrett et al. (1994) for details on the structure). Solving the STS is equivalent to performing a *Gauß elimination*, which is difficult to parallelize as it involves many dependencies. Therefore, the STS are solved on CPU⁶ that can take advantage of caches and of the sparsity of the matrix, to efficiently solve the system.

For large systems (see Section 7.2.2), the computation cost of the \mathbf{LDL}^T factorization can become prohibitive, and the resulting preconditioner can diverge from the actual system. However, we note that an important part of the error is associated with the rotations (Saupin et al., 2008) which can vary quickly between time steps. In order to limit the divergence of the preconditioner, we estimate the nodal⁷ rotations $\mathbf{R}_{t-h \rightarrow t}$ that were introduced since the last update of the preconditioner (i.e. between time $t-h$ and t). The most recent preconditioner \mathbf{P}_{t-h} is then rotated with the current rotation matrix $\mathbf{R}_{t-h \rightarrow t}$ as follows:

$$\mathbf{P}_r = \mathbf{R}_{t-h \rightarrow t}^T (\mathbf{L}_{t-h} \mathbf{D}_{t-h} \mathbf{L}_{t-h}^T) \mathbf{R}_{t-h \rightarrow t} \quad (18)$$

where the “rotated preconditioner” \mathbf{P}_r is less sensible to geometrical non-linearities. Finally, the method enables to simulate the deformation of homogeneous as well as heterogeneous tissues in real-time.

5 GPU-Based preconditioner for contact problems

Eq. (12) requires the computation of \mathbf{A}^{-1} , which is a large matrix (same dimension as the number of DOF) and changes at each time step. Although computing this inverse in real-time is only possible for coarse models, the resulting operator \mathbf{W} plays an important role to enforce the constraints.

⁶Using a GPU-based CG with a CPU-based preconditioner, implies to transfer the solution vector \mathbf{b} between the CPU and the GPU at each iteration of the CG. But, since the preconditioner is usually a good approximation of the actual system, only a few iterations are necessary and the cost of such transfers remains limited.

⁷Note that nodal rotations give an approximation of the co-rotational formulation where rotations are computed per elements and sum in the global stiffness matrix. $\mathbf{R}_{t-h \rightarrow t}$ is a block diagonal matrix, easy to compute and easy to invert.

5.1 Compliance and mechanical coupling

In the following, the term “*mechanical coupling*” describes the coupling between contact constraints applied on two subsets of the boundary of a deformable body. This coupling occurs through the deformation of the body itself. Indeed, even if the contact points are only defined on the boundary of the deformable bodies, they are all influenced by each other through the stiffness of the material. Consider for example Fig. 3 where, even if the deformed shapes do not show any interpenetration, the behavior computed in Fig. 3(b) is not acceptable, since stiff and soft parts deform in the same way. It shows that the contact force distribution is closely related to the underlying heterogeneity of the material, which is represented by \mathbf{W} .

In the context of explicit schemes matrix \mathbf{W} would only be built from a diagonal mass matrix. If penalty methods are used, the force distribution would mainly depend on the geometrical interpenetration, not on the inhomogeneities. In both cases, it would lead to unrealistic configurations such as in Fig. 3(b), at least during transient states. Pre-computing the inverse \mathbf{A}_0 at the initial step, and using it all along the simulation provides better results, but is limited to linear, small displacements models. Another approach adapted to the real-time context, is to use an approximation of \mathbf{W} . For instance, Saupin et al. (2008) proposed the *compliance warping* technique, that consists in updating the precomputed \mathbf{A}_0^{-1} with the nodal rotations, but this solution remains inaccurate for large deformations, and requires storing a large dense matrix which makes the method unsuitable for fine meshes.

In this paper, we propose to reuse the asynchronous preconditioner of the previous section as an approximation of the compliance operator \mathbf{W} , and we detail our GPU-Based algorithm that allows for real time computations.

5.2 Optimized preconditioner for contacts

We propose to use the asynchronous preconditioner computed in Section 4 as an approximation of \mathbf{A}^{-1} :

$$\mathbf{H}\mathbf{A}^{-1}\mathbf{H}^T \approx \mathbf{H}\mathbf{P}^{-1}\mathbf{H}^T \quad (19)$$

Indeed, since \mathbf{P} represents a close approximation of the factorization of \mathbf{A} , we propose to use it to compute \mathbf{W} in Eq. (12). For each interacting object, substituting Eq. (18) in (19) gives:

$$\mathbf{H}\mathbf{A}^{-1}\mathbf{H}^T \approx \mathbf{H}(\mathbf{R}\mathbf{L}\mathbf{D}\mathbf{L}^T\mathbf{R}^T)^{-1}\mathbf{H}^T \quad (20)$$

The above equation requires computing the product of the inverse of the preconditioner with the Jacobian of contacts \mathbf{H} , which can be achieved by computing columns independently of \mathbf{H} :

$$\mathbf{LDL}^T \mathbf{X}_i = \mathbf{H}_i \Leftrightarrow \mathbf{X}_i = (\mathbf{LDL}^T)^{-1} \mathbf{H}_i \quad (21)$$

where \mathbf{X} gives the result of the inverse of the preconditioner times the Jacobian of the contact. Therefore, we obtain \mathbf{W} within 4 steps as detail in Algorithm 1. Steps 1 and 3 are inexpensive because \mathbf{H} and \mathbf{J} are sparse matrices, \mathbf{D} is a diagonal matrix, and \mathbf{R} is a block-diagonal matrix. Step 4 involves the product of two dense matrices, which can be parallelized efficiently on GPU using the NVIDIA Corporation (2007) library. Step 2 requires the solution of a STS for each column of \mathbf{J} , composed of the lower triangular matrix \mathbf{L} . This operation remains the most expensive task and would quickly become too prohibitive if it was processed sequentially on a traditional CPU. Therefore, we propose to parallelize the computation of \mathbf{S} on GPU.

5.3 Solution of multiple STS on GPU

Implementation-wise, the simplest solution to solve the multiple STS on GPU is to use NVIDIA Corporation (2007) library. Nevertheless, this library is optimized for solving large and very sparse systems of equations (see Naumov (2011)), but the range of size of the matrices compatible with the real-time constraint is still much smaller. Therefore, to decrease the computational time of this critical step, we propose a solution to parallelize the computation of \mathbf{S} on GPU. Our solution is based on a two level parallelization strategy which is inspired from the method introduced in Courtecuisse and Allard (2009). The main difference is that the underlying \mathbf{L} matrix is stored in a sparse format, which makes it difficult to load efficiently on the GPU processors.

5.3.1 GPU-based parallelization—The multiple right-hand side vectors stored in \mathbf{J} can be computed independently from each other. Therefore, we assign the computation of each column of \mathbf{S} to an independent multiprocessor on the GPU. Each group is therefore fully processed by a single processor (see Fig. 4) which enables to use fast local synchronizations directly on the GPU. Then we use a second level of parallelism where each STS is solved with several threads. Indeed, a lot of data can potentially be treated in parallel during the solving process of each STS. This two level strategy fits the GPU architectures where local synchronizations within a group of threads are fast, whereas global synchronizations over multiple groups are much more costly.

Nevertheless, as mentioned above, solving a STS involves a number of dependencies. For instance, for the lower triangular system, the computation of the solution s_j of a given row j , requires the result of all previous solutions s_i such as $i < j$:

$$\mathbf{s}_j = \mathbf{b}_j - \sum_{i=0}^{i < j} (\mathbf{s}_i \mathbf{L}_{j,i}) \quad (22)$$

where \mathbf{b} is the solution vector and \mathbf{s} is the unknown. Therefore, each row must be processed sequentially (i.e. a synchronization of each row is necessary). However, the partial

contributions $\sum_{i=0}^{i < j} (s_i \mathbf{L}_{k,i})$ could be processed simultaneously for any row k such as $k > j$.

We propose a block-row⁸ parallelization strategy, that is inspired from the block-column scheme introduced in Courtecuisse and Allard (2009): First we accumulate the contributions of the block off-diagonal in parallel; Then we solve the block diagonal sequentially (see Fig. 5).

We use a group of $t \times t$ threads to process t rows simultaneously (each row is therefore treated by t threads in parallel). Since \mathbf{L} is sparse, it is difficult to predict if a data is located on the bloc diagonal before actually reading it in global memory. To avoid reading twice the CRS matrix, we propose to use two buffers *acc* and *diag* stored in shared memory: *acc* is used to accumulate the contributions off-diagonal whereas *diag* is used to copy in dense format the data located on the block-diagonal (see Fig. 5). A first local synchronization is then used to ensure that the t rows are fully processed by all the threads. Then, a parallel reduction (see Martin et al. (2012)) is processed to sum per row the contributions stored in *acc* (row 4 in Algorithm 2), then a second local synchronization is necessary (row 5). Finally, the block-diagonal is solved as a dense problem⁹ using *diag* in shared memory (row 5).

In our experiments we use $t = 16$ and only 3 local synchronizations are necessary to solve 16 rows with 256 GPU threads. Although this implementation may be slower than a CPU-based solver for a single STS, our GPU-based strategy enables to solve the multiple right-hand side vectors simultaneously (i.e. with the same cost as a single STS).

6 Simulation of cutting

Simulation of cutting involves two main issues: First to re-mesh the FE structure while keeping the consistency of the mesh i.e. split correctly the domain in order to be able to separate the cut parts in the simulation, and avoid degenerated elements such as sharp or thin elements. The second issue is to update adequately the mechanical properties and the equation systems of the deformable model when the mesh is cut. In this paper, we only address this second aspect and we show how to update the asynchronous preconditioner according to the topological changes.

6.1 Topological modifications, cutting

The method presented in this paper could be used with any re-meshing algorithm, as far as the modifications remain local and only affect few elements per time step. In our simulations, we use a re-meshing algorithm similar to Mor and Kanade (2000), where, rather than reconstructing the overall mesh, we incrementally update it within 3 steps: First we remove the intersected elements from the current mesh; Second we subdivide the removed elements; Third, we add back the subdivided elements.

⁸Block-row strategy instead of a block-column as in Courtecuisse and Allard (2009), because with a block-column solution, writing conflict in memory cannot be predicted due to the CRS format of \mathbf{L} .

⁹The block-diagonal is solved using a column-based strategy to avoid the need of parallel reductions. This is possible since *diag* is stored in dense format.

The subdivision process affects the *stiffness*, *mass* and *damping* matrices (see Fig. 6), and the final linear system \mathbf{A} (i.e. Eqs. (5), (16) and (20)). Nevertheless, the conjugate gradient (CG) used to solve Eq. (5) only requires performing matrix vector products. Thus, \mathbf{A} can be directly evaluated from the finite element (FE) mesh to instantly take into account the modifications in the solution. However, the preconditioner used in Eqs. (16) and (20) is updated with delay, and the modifications significantly affect its efficiency. Moreover, contrary to Eq. (16) where the preconditioned CG ensures the convergence of the system, the preconditioner is directly used to build an approximation of the contact problem in Eq. (20). Therefore, the delay of the updates can lead to instabilities and inaccuracies, in particular when treating contact with the instruments and self-collisions between different parts of the cut.

6.2 Low rank update of the asynchronous preconditioner

We propose to use the Sherman Morrison Formula (SMF) to compute the correction of the preconditioner due to the topological changes:

$$\tilde{\mathbf{P}}^{-1} = (\mathbf{P} + \mathbf{G}\mathbf{N}\mathbf{G}^T)^{-1} = \underbrace{\mathbf{P}^{-1}}_{\text{Last factorization}} - \underbrace{\mathbf{G}\mathbf{P}^{-1}(\mathbf{N}^{-1} + \mathbf{G}\mathbf{P}^{-1}\mathbf{G}^T)^{-1}\mathbf{P}^{-1}\mathbf{G}^T}_{\text{Correction due to the cut}} \quad (23)$$

where $\tilde{\mathbf{P}}$ is the modified preconditioner, \mathbf{G} is a globalization matrix which maps the rows/columns to the global system and \mathbf{N} is the perturbation of the preconditioner obtained as a difference between the modified system, and the last preconditioner:

$$\mathbf{N} = \mathbf{G}^T\tilde{\mathbf{P}}\mathbf{G} - \mathbf{G}^T\begin{pmatrix} \mathbf{P} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}\mathbf{G} \quad (24)$$

In order to keep a consistent formulation, \mathbf{P} is padded with the identity for all the added nodes during the subdivision process. This approach assumes that the added degrees of freedom were present before the cut with a unitary mass but not attached to the mechanical system, and \mathbf{N} corresponds to the correction of the padded system. \mathbf{N} is fast to compute since it only involves subtractions of small matrices associated with the nodes affected by the cut.

An important advantage of updating the preconditioner is that it helps to maintain the number of affected nodes by the perturbations at minimum. Indeed, each new factorization implicitly contains all anterior modifications to the last update of the preconditioner. The SMF correction is therefore necessary only for the perturbation that appeared since the last update, which only involves few affected nodes (see Fig. 7).

Contrary to Courtecuisse et al. (2010b), we do not store the dense inverse of \mathbf{W} . Instead we use the sparse \mathbf{LDL}^T factorization of \mathbf{A} . The SMF cannot be directly applied with such a factorization because it explicitly requires the inverse of \mathbf{W} . Thus, we proceed in two steps: First we compute the correction for the preconditioner, which is only necessary when a new topological modification is performed; Then we apply the correction until the next update of the preconditioner.

6.3 Computation of the correction

For each new topological modification, the correction is obtained by computing the two following matrices:

$$\mathbf{U} = (\mathbf{LDL}^T)^{-1} \mathbf{G}^T \quad (25)$$

$$\mathbf{Q} = (\mathbf{N}^{-1} + \mathbf{GU})^{-1} \quad (26)$$

The computation of \mathbf{U} involves the product of the inverse of the preconditioner with the globalization matrix \mathbf{G} . This implies to solve equations composed of the lower and upper triangular systems (see Algorithm 3). Since, the CRS format used to store \mathbf{L} prevents access to matrix \mathbf{L}^T by columns, and in order to use the same parallelization scheme as described in Section 5.2, the transpose matrix \mathbf{L}^T is explicitly computed by the asynchronous thread after the factorization of the preconditioner.

Eq. (26) requires to compute the inverse of two small matrices (\mathbf{N} and $\mathbf{N}^{-1} + \mathbf{GU}$) which are performed on CPU. As a difference between two mechanical matrices, \mathbf{N} is ill-conditioned. Its inverse is therefore ill-defined, but contrary to Courtecuisse et al. (2010b), our method does not lead to the accumulation of subsequent round-off errors. Indeed, as soon as a new factorization is released the perturbation is include in the new preconditioner, and \mathbf{N} is erased so that these numerical errors do not accumulate over time.

6.4 Application of the correction

Once \mathbf{U} and \mathbf{Q} are computed for a given modification, we use them to correct the preconditioner until the next update. However, the local rotations used in Eqs. (18) and (20) vary at each time step, which prevents including them in the formulation of the correction. In order to use the rotations to improve the efficacy of the preconditioner, we apply them around the corrected formulation. Substituting (23), (25), (26) in (16) gives:

$$\tilde{\mathbf{P}}^{-1} \approx \mathbf{R}(\mathbf{P}^{-1} - \mathbf{U}^T \mathbf{Q} \mathbf{U}) \mathbf{R}^T \quad (27)$$

The rotations are therefore applied around the final corrected solution. To simplify notations we do not carry those rotations in the upcoming equations. For each iteration of the preconditioned CG, the application of the preconditioner is corrected:

$$\begin{aligned} \tilde{\mathbf{P}} \mathbf{x} = \mathbf{b} &\Leftrightarrow \mathbf{x} = (\mathbf{P}^{-1} - \mathbf{U}^T \mathbf{Q} \mathbf{U}) \mathbf{b} \\ \mathbf{x} &= \underbrace{\mathbf{P}^{-1} \mathbf{b}}_{\mathbf{x}^a} - \underbrace{\mathbf{U}^T \mathbf{Q} \mathbf{U} \mathbf{b}}_{\mathbf{x}^c} \end{aligned} \quad (28)$$

where \mathbf{x}^c is the correction of the solution, which involves 3 dense matrix–vector products, and are performed on GPU using the CUBLAS library. \mathbf{x}^a is equivalent to applying the preconditioner without correction as in Eq. (16), except that if nodes are added during the subdivision process, \mathbf{P} is padded with the identity:

$$\begin{pmatrix} \mathbf{x}_m^a \\ \mathbf{x}_a^a \end{pmatrix} = \begin{pmatrix} \mathbf{P}^{-1} & 0 \\ 0 & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{b}_m \\ \mathbf{b}_a \end{pmatrix} \quad (29)$$

where subscripts m and a correspond to the modified and added nodes, respectively.

For the contact response, the corrected *Delasus operator* \mathbf{W} is obtained with:

$$\mathbf{H}\tilde{\mathbf{P}}\mathbf{H}^T = \mathbf{H}(\mathbf{P}^{-1} - \mathbf{U}^T\mathbf{Q}\mathbf{U})\mathbf{H}^T \underbrace{\mathbf{H}\mathbf{P}^{-1}\mathbf{H}^T}_{\mathbf{W}^a} - \underbrace{\mathbf{H}\mathbf{U}^T\mathbf{Q}\mathbf{U}\mathbf{H}^T}_{\mathbf{W}^c} \quad (30)$$

where \mathbf{W}^a is the asynchronous *Delasus operator*, and \mathbf{W}^c its correction. As in Eq. (29), \mathbf{H} may involve constraints on the newly created degrees of freedom, so that we pad \mathbf{P} with the identity matrix, and \mathbf{W}^a is computed by:

$$\mathbf{W}^a = \begin{pmatrix} \mathbf{H}_m\mathbf{P}^{-1}\mathbf{H}_m^T & 0 \\ 0 & \mathbf{H}_a\mathbf{H}_a^T \end{pmatrix} \quad (31)$$

$\mathbf{H}_m\mathbf{P}^{-1}\mathbf{H}_m^T$ is solved using the GPU-based algorithm introduced (Section 5.2), whereas $\mathbf{H}_a\mathbf{H}_a^T$ involves a sparse matrix product which is parallelized using the CUSPARSE library. Finally, \mathbf{W}^c is obtained by 3 small dense matrix products which are performed on GPU using the CUBLAS library.

7 Results and evaluation

In this paper, we proposed a method to significantly improve the trade-off between accuracy and computation time of an interactive simulation. In the first part of this section, we compare the accuracy of our method versus standard algorithms. In the second part of this section we measure the performance our method on academic examples. Medical applications are presented in the next section.

7.1 Accuracy and validation

7.1.1 Comparison with ABAQUS—Some comparisons of algorithms available in SOFA with analytical solutions can be found in Nesme et al. (2005) and Marchal et al. (2008). In order to provide a more thorough comparison for our method, in particular in the case of contacts, we produced a simulation involving a beam attached at one extremity and deforming under gravity until the beam contacts with a plane. The dimension of the beam is

120 × 10 × 10 mm, the mesh is composed of 10,779 linear tetrahedral elements and 2481 nodes. The constitutive law is based on small strain with large displacement (co-rotational), the mass density is set at 0:0001 Kg/mm³, the young modulus is set at 10 Mpa and the Poisson ration at 0.4.

The same simulation was computed by both SOFA and ABAQUS (Simulia, Dassault Systmes S.A.) using the same tetrahedral mesh, the same model (co-rotational for SOFA and Large Displacement Theory in ABAQUS), the same boundary conditions and assuming no friction during the contact. At equilibrium we computed the *Von Mises Stress* of the ABAQUS solution (see Fig. 8) and of our SOFA implementation (see Fig. 8(b)). Then we measured the difference of the stresses computed by the two approaches. Results are presented in Fig. 8(c) (using the same scale as the initial stresses) and in Fig. 8(d) using a normalized scale, i.e. where red parts correspond to the main difference (to emphasize areas where even a small difference exists). Although a maximum error of 20% is found near the constrained part of the beam, the average stress error is equal to 0.017 (around 1.7%), which is almost negligible. Finally, the positions of the beams match perfectly, and the maximum distance between SOFA nodes and ABAQUS nodes, is less than 0.37 mm.

7.1.2 Preconditioner for the contact response—We now measure the error introduced when using our asynchronous preconditioner as an approximation of the compliance matrix defined in Eq. (12). We created a simulation involving an heterogeneous disk which is driven by a sphere through the center of a torus (see Fig. 9). This is a very good test for measuring both the deformation process and the contact response. We first produced a reference simulation where the exact compliance matrix is computed by inverting **A** at each time step (which is obviously not real-time), and we measured the distances between the nodal positions obtained by using an “approximated compliance matrix” and the reference simulation.

Relying on the diagonal of **A**⁻¹ to build **W (Diagonal)**, as done by Dequidt et al. (2009) for the vessel wall compliance, the lens cannot be pushed into the cavity. Indeed, the forces exerted by the ball are located at the center (stiff part), while other forces are applied by the cylinder on the periphery (soft part) and hold the object up. Without coupling, the contact forces applied onto the stiff part are not transmitted to the periphery, and the lens remains in equilibrium without being deformed. The precomputed inverse of the compliance matrix (**LDL no update**), gives better results when the object is in a similar configuration as the rest position (i.e. in the state where the precomputed inverse has been computed). However, the rigidity is not properly evaluated in case of large deformations, and a large error is introduced in the behavior.

When the preconditioner is updated periodically after a fixed number of time steps (**LDL update 20** and **LDL update 10**), the error tends to decrease with the frequency of the updates, and for the version where the preconditioner is updated every 10 time steps there is no visible error (see Fig. 9(b)). Finally, using the asynchronous version (**LDL update async**), the preconditioner is updated on average every 5 time steps and the error introduced throughout the simulation is negligible.

7.2 Performance evaluation

7.2.1 Convergence rate and computation time—We now evaluated the computational time and the convergence rate of standard solvers to solve Eq. (5). We produced a simulation where an heterogeneous beam composed of 6500 elements and 1470 nodes, is deformed under gravity (see Fig. 10). In all our experiments we set the tolerance of the preconditioned CG to 10^{-7} , and we use a time step $h = 0.02$.

Direct solvers must process a complete factorization of the matrix at each time step, to take into account the non-linearities of the model. This operation is expensive but it can be parallelized using optimized libraries such as Pardiso Schenk et al. (2008). However, due to the multiple dependencies and the relatively small size of the system, a parallel version with 2 threads (**Pardiso-2**) is only 1.5× faster than a sequential approach (**Pardiso-1**). With 4 threads (**Pardiso-4**) only 2.2× faster, and with 8 threads version (**Pardiso-8**) only 1.90× faster (and slower than **Pardiso-4**).

Standard iterative solvers (**standard**) require a large number of iterations, and thus a large computational time, due to the strong heterogeneity of the material. Indeed, although our GPU-based CG does not need to assemble the matrix and provides fast iterations, an average of 493 iterations are necessary to obtain a sufficient solution and the computation time is no longer compatible with real-time. The Jacobi preconditioner is too simple and does not manage to sufficiently reduce the number of iterations.

Pre-computing the \mathbf{LDL}^T factorization (**no update**) and using it throughout the simulation, enables to remove the overhead of the factorization while keeping a limited number of iterations. However in cases of large deformations, the actual stiffness of the material may be very different from the rest configuration, and a large number of iterations are necessary. Indeed, when the beam undergoes large deformations, we measured a maximum of 121 iterations necessary to achieve convergence, and the application time of the preconditioner was around 210 ms. Interestingly, applying the rotation around the preconditioner (*warping method*) helps to sensibly reduce the number of iteration with a limited overhead. Indeed, the unitary cost of a single iteration with the warping method is only 11% higher whereas it requires 3.5 times fewer iterations to converge.

When using our method where the \mathbf{LDL}^T factorization is updated asynchronously (**Async**), the factorization time of the preconditioner is still negligible, and the number of iterations remains very low throughout the simulation. For this preconditioner, the (*warping method*) provides only limited improvements. Indeed, the preconditioner was updated on average every 8 time steps, and the rotations between two consecutive updates remained limited. The warping heuristic is therefore particularly beneficial if the preconditioner cannot be updated sufficiently fast (for larger systems for instance). Finally, the asynchronous solution is the only method which is compatible with real-time computations. Indeed, the simulation runs at 45 FPS, and 1 s in the simulation is simulated in 1 s.

7.2.2 Factorization of the preconditioner—To test the scalability of our method, we evaluate the computational time required to factorize the mechanical matrices with dimensions up to $10,000 \times 10,000$ (which corresponds to an object with 3333 nodes with 3

dofs per node). This operation is very expensive, but since we process it asynchronously it does not impact directly the performances of the simulation.

We also evaluate the number of simulation steps necessary to update the preconditioner with a simulation running asynchronously at 25 FPS (see Fig. 11). Depending on the dimension of the matrix, the number of time steps necessary to update the factorization varies between 1 and 6. With our asynchronous version, this period can be twice larger because a first period is necessary to compute the factorization, and then it will be used until the next update of the preconditioner. However, even for a 9000×9000 matrix, the maximum delay due to the update of the preconditioner is less than 0.5 s.

7.2.3 Computation of the compliance—We now focus on the computation time (see Fig. 12) for solving a Sparse Triangular System with multiple right hand side vectors. This operation is necessary to set the contact problem in Eq. (12) and (21).

Solving the different STS on the GPU (**LDL GPU**) is much faster than solving sequentially each STS on the CPU (**LDL CPU**). Indeed, the computation time for the CPU version is linear according to the number of right-hand side vectors, whereas the GPU can process them in parallel. Therefore, for up to 78 constraints the computation time remains almost constant with our GPU implementation. Indeed, below this limit the GPU computing units are not fully utilized and the different STS are processed in parallel. Beyond this number, some GPU processors will compute several STS successively, and the computational time curve takes a staircase appearance. Nevertheless, the GPU processors are able to overlay waiting times, due to synchronizations and access in memory, with computations for another STS. Thus, solving the system for 140 constraints is only 1.8 times slower than for 70 constraints.

We also compared our GPU implementation to the NVIDIA Corporation (2007) library where an implementation of solving a STS with multiple right-hand side vector has recently been released in the procedure `cusparseScsrsm_solve`. This procedure is implemented in two steps: first an analysis of the sparsity of the matrix is processed in order to determine the dependencies, then the analysis information is used to solve the STS (see Naumov (2011) for details). As the analysis is necessary only when the preconditioner is updated, we detailed separately the cost with (**CUSPARSE + A**) and without (**CUSPARSE**) the analysis. For 78 constraints our solution is 4.2 times faster when the analysis is not required (5.72 times with the analysis), and 2 times faster for 300 constraints (2.4 times with the analysis). Although the difference may be less for larger systems, we have optimal performance for the scenarios that are compatible with real-time constraint. Indeed, for 300 constraints the computation time to build **W** with our optimized approach is around 0.04 s which is a limit for interactive rates.

7.2.4 Topological modifications—Finally, we evaluate the influence of the cut on the convergence of the standard **CG**, and using our method (**SMF**), on a simulation composed of a beam cut lengthwise and falling under gravity (see Table 7.2.3). **Build** is the time to assemble **A**: For the **CG**, the matrix is directly evaluated from the mesh structure at each iteration, whereas the preconditioner is updated on average every 4.30 simulation steps, each

update requiring to fully assemble \mathbf{A} to perform the factorization in Eq. (16). **Iterations** and **Solve** are respectively the number of iterations per time step, and the corresponding time to solve the system with a tolerance at 10^{-7} . The method significantly decreases the number of iterations, and provides an average speed-up of 2.6× compared to the **CG**.

The main overhead of the **SMF** update is the computation \mathbf{U} and \mathbf{Q} (see Fig. 13), but these operations are performed only when a new topological modification is detected (i.e. every 3.4 simulation steps on the beam example). The computation of \mathbf{U} is the most expensive, and its GPU parallelization is the key point to enable real-time computation, whereas the inversion of \mathbf{Q} is inexpensive for small perturbations but quickly becomes costly for large perturbations. In practice, the number of impacted nodes remains very small since the preconditioner is updated several times per second. Finally the application of the correction (i.e. applying \mathbf{x}^c in Eq. (28) and \mathbf{W}^c in (30)) is negligible since it represents less than 1% of the computation time of a time step.

8 Applications

We now demonstrate that our method is generic enough to address several kinds of simulation in a medical context. We use it to simulate a cataract surgery, an hepatectomy in laparoscopic surgery, and a cerebral tumor removal. We show that our method can handle the requirements of such simulations in real-time.

8.1 Application to cataract surgery

The cataract is an opacification of lens of the eye, which prevents the passage of light and results in partial or complete blindness. Millions of people are affected by this pathology, particularly in third world countries. A surgical treatment exists, which consists in extracting the diseased lens and replacing it with by an implant. The standard surgical procedure is known as *Phacoemulsification* (Tsuneoka et al., 2002), where the lens is emulsified by an ultrasonic tool. However, Phacoemulsification requires advanced technology which is not available in many countries where the prevalence of cataract is highest, and hence many patients can simply not be treated.

Another surgical procedure known as *Manual Small Incision Cataract Surgery* (MSICS) (Venkatesh et al., 2008) requires only basic technology and leads to quasi-identical results when performed by an experienced specialist. This technique requires a slightly larger incision (around 5 mm) so that the lens be extracted in a single piece. This involves that the eyeball as well as the lens must be simulated because they both undergo large deformations and high stresses. The heterogeneity of the lens must also be taken into account since it has an impact on the success of the lens extraction (the nucleus of the lens is much stiffer than its periphery).

In order to simulate the lens extraction according to the MSICS technique, the lens is modeled with 1113 nodes and 4862 tetrahedra, whereas the eye contains 1249 nodes and 3734 tetrahedra. The center of the lens is 5 times stiffer than the periphery, and the incision and meshes of the organs has been generated off-line using the The CGAL Project (2011) library. The lens is removed with the help of deformation of the eyeball, and friction with the

surgical instrument. To account for both the multiple contacts, and the heterogeneity of the lens, we used our asynchronous preconditioner to ensure the convergence of the CG.

The preconditioned CG required an average of 11.6 iterations to converge to 10^{-5} , despite strong deformations and heterogeneities. Using our method, we managed to simulate this application in real-time (see Fig. 14), while maintaining a computation speed from 18 to 25 FPS. Within a single time step, the distribution of the computation time was: 40.56% for the free motion and 44.69% for the corrective motion.

8.2 Application to liver resection

Nearly 100,000 European citizens die every year of cirrhosis of the liver or liver cancer. Surgical procedures remain the options that offer the foremost success rate against such pathologies. The ability to simulate liver resection (hepatectomy) is key for both advanced training of this complex procedure (in particular when done using a laparoscopic technique). But simulation can also have an important impact in the planning phase. Indeed, eligibility for liver surgery is based on the minimum safety liver volume remaining after resection, but this minimum value varies over time and from one patient to another according to biological and biomechanical properties of the liver. Although various preoperative planning software have been developed, they only provide the volume of the liver before and after resection. However interesting, this limited information is not sufficient to improve the rate of surgical eligibility. By combining patient-specific anatomical and biomechanical modeling, we can provide the means to rehearse for the procedure while trying to maximize the future liver remnant (FLR).

Although simulators of this procedure have been developed in the past (Bourquain et al., 2002; Lamadé et al., 2002), the originality of our approach is that our simulation is based on patient specific data. Meshes of the organs are obtained from a semi-automatic segmentation of a CT (see Soler et al. (2001) for details). We simulate 5 deformable bodies in interaction (liver, stomach, colon, intestines and diaphragm). Each organ is composed of several hundred of nodes and thousand of elements with complex shapes composed of several thousand of triangles (see Fig. 15). An important issue to produce this application concern the collision detection which is performed by the method introduced in Allard et al. (2010).

This simulation runs at a frequency of 25 FPS, including during cutting phases. The distribution of computing time in a time step is as follows: 27.77% for free movement, 11.01% for collision detection, 22.48% for the constraint motion. The preconditioned CG requires an average need of 5.82 iterations to converge. Finally, by taking into account the *mechanical coupling* between the contacts, we managed to produce a consistent haptic feedback. For instance, users can feel the stiffness of ribs behind the liver by applying contacts on the surface of the organ (see Fig. 16).

8.3 Application to brain tumor resection

The last application that uses our method is the simulation of a surgical resection of a brain tumor. Brain surgery simulation has been widely studied, and commercial products start to appear (Luciano et al., 2005; Delorme et al., 2012). A review of computer-based brain

surgery simulations can be found in Malone et al. (2010) and Alaraj et al. (2011). It is still an active research topic, for instance for the prediction of the *brain shift* (Joldes et al., 2010) or for the *Deep Brain Stimulation* (Bilger et al., 2011). The brain is a relatively soft tissue, and most of existing simulations rely on explicit integration. Nevertheless, during the surgery the accuracy of surgeon's interactions with the tissues is fundamental, as an error of few millimeters may often have dramatic consequences for the patient. The main needs of such simulation are precise modeling and topological modifications, high heterogeneities in the deformation modeling (the tumor is often stiffer than the brain tissues) and real-time computation. We show that our method meets these two issues, and may be beneficial to improve the realism of existing brain simulations.

The brain is modeled as a heterogeneous deformable body, composed of 1734 nodes and 7680 linear tetrahedral elements. The tumor is 20× stiffer than the brain. During the simulation, the preconditioner is updated every 5.6 steps, and a new topological modification appears every 5.5 simulation steps, affecting 24 nodes. A total of 553 modifications are performed, and the method remains stable with an average of 5.70 iterations to solve the linear system. The collisions and self-collisions are correctly solved while processing the modifications, and cut parts can instantaneously be separated upon contact with the instrument. Finally, we achieve between 20 and 40 FPS and the method remains interactive.

9 Conclusion

In this manuscript, we introduce a set of methods for real-time bio-mechanical simulation of soft anatomical structures, relying on an implicit time integration method. The proposed paradigm relies on an asynchronous preconditioner that is updated at low frequency, and permits to significantly reduce the number of iterations in the linear solver. It also improves the contact response process by taking into account the mechanical coupling between contact points. We also extended the approach to handle topological modifications without compromising the interactivity of the simulation. The method is particularly beneficial for heterogeneous structures. On a simple case, we verify the method by comparing it with those of a commercial software, and we demonstrate the benefits of the proposed method through applications in cataract surgery, liver surgery and neurosurgery. In the last two applications, topological changes due to cutting are also enabled with minimal impact on the computation time.

For future work, we plan to investigate algebraic model reduction techniques (using the proper orthogonal decomposition, or multi-scale methods) to decrease the computational expense through pre computations, and to allow a finer description of organs. We also plan to investigate the use of enriched finite element methods in the real-time context to handle continuous cut with the elements. Finally, we will also focus on the estimation of the spatial and temporal discretisation error and of the model error during our simulations. An analysis of the model error will require addressing more realistic models including hyper-elastic materials, where the mechanical matrices undergo 'faster' modifications than for the co-rotational case. We will also investigate the stability of the numerical schemes for nearly incompressible materials.

Appendix A. Supplementary data

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The authors would like to thank Mouhamadou Diallo and Dr. Igor Peterlik for their help and their assistance during the comparison between SOFA and ABAQUS.

We also would like to thank the financial support of the European Research Council Starting Independent Research Grant (ERC Stg Grant agreement No. 279578) entitled Towards real time multi-scale simulation of cutting in non-linear materials with applications to surgical simulation and computer guided surgery.

References

- Alaraj A, Lemole MG, Finkle JH, Yudkowsky R, Wallace A, Luciano C, Banerjee PP, Rizzi SH, Charbel FT. Virtual reality training in neurosurgery: review of current status and future applications. *Surg Neurol Int.* 2011; 2
- Allard J, Faure F, Courtecuisse H, Falipou F, Duriez C, Kry PG. Volume contact constraints at arbitrary resolution. *ACM Trans Graph (Proc. SIGGRAPH).* 2010; 29(3)
- Allard, J, Courtecuisse, H, Faure, F. Implicit FEM solver on GPU for interactive deformation simulation GPU Computing Gems Jade Edition. Hwu, WW, editor. Elsevier; 2011. 281–294.
- Anitescu M, Potra F, Stewart D. Time-stepping for three-dimensional rigid body dynamics. *Comput Meth Appl Mech Eng.* 1999:183–197.
- Baraff, D; Witkin, A. Large steps in cloth simulation. *SIGGRAPH '98: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques; ACM; 1998. 43–54.*
- Barbič J, James DL. Real-time subspace integration for St. Venant–Kirchhoff deformable models. *ACM Trans Graph.* 2005; 24:982–990.
- Barrett, R, Berry, M, Chan, TF, Demmel, J, Donato, J, Dongarra, J, Eijkhout, V, Pozo, R, Romine, C, der Vorst, HV. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods.* second ed. SIAM; 1994.
- Berger-Vergiat L, Waisman H, Hiriyyur B, Tuminaro R, Keyes D. Inexact schwarz-algebraic multigrid preconditioners for crack problems modeled by extended finite element methods. *Int J Numer Meth Eng.* 2012; 90(3):311–328.
- Bielser, D; Glardon, P; Teschner, M; Gross, M. A state machine for real-time cutting of tetrahedral meshes. *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications PG '03; Washington, DC, USA: IEEE Computer Society; 2003. 377*
- Bilger, A; Dequidt, J; Duriez, C; Cotin, S. In: Fichtinger, G; Martel, A; Peters, T, editors. *14th International Conference on Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011, Lecture Notes in Computer Science; Toronto, Canada: Springer; 2011. 339–346.*
- Bolz J, Farmer I, Grinspun E, Schröder P. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM Trans Graph.* 2003; 22(3):917–924.
- Bordas S, Duflo M. Derivative recovery and a posteriori error estimate for extended finite elements. *Comput Meth Appl Mech Eng.* 2007; 196(35):3381–3399.
- Bordas S, Duflo M, Le P. A simple error estimator for extended finite elements. *Commun Numer Meth Eng.* 2008; 24(11):961–971.
- Bourquain H, Schenk A, Link F, Preim B, Prause G, Peitgen H. Hepavision2 – a software assistant for preoperative planning in living-related liver transplantation and oncologic liver surgery. *Can Res.* 2002:1–6.
- Braess, D. *On the combination of the multigrid method and conjugate gradients* Multigrid Methods II. Springer; 1986. 52–64.
- Brezzi F, Hager WW, Raviart P-A. Error estimates for the finite element solution of variational inequalities. *Numerische Mathematik.* 1977; 28(4):431–443.

- Bro-Nielsen M, Cotin S. Real-time volumetric deformable models for surgery simulation using finite elements and condensation. *Comput Graph Forum*. 1996; 15(3):57–66.
- Buatois L, Caumon G, Lévy B. Concurrent number cruncher – a GPU implementation of a general sparse linear solver. *Int J Parallel Emerg Distrib Syst*. 2009; 24(3):205–223.
- Comas, O; Taylor, Z; Allard, J; Ourselin, S; Cotin, S; Passenger, J. Efficient nonlinear FEM for soft tissue modelling and its GPU implementation within the open source framework SOFA. *ISBMS*; 2008. 28–39.
- Cotin S, Delingette H, Ayache N. Real-time elastic deformations of soft tissues for surgery simulation. *IEEE Trans Visual Comput Graph*. 1999; 5(1):62–73.
- Courtecuisse, H; Allard, J. Parallel dense Gauss–Seidel algorithm on many-core processors. *High Performance Computation Conference (HPCC)*; IEEE CS Press; 2009.
- Courtecuisse, H; Allard, J; Duriez, C; Cotin, S. Asynchronous preconditioners for efficient solving of non-linear deformations. *Proceedings of Virtual Reality Interaction and Physical Simulation (VRIPHYS)*; 2010a November.
- Courtecuisse H, Jung H, Allard J, Duriez C, Lee DY, Cotin S. GPU-based real-time soft tissue deformation with cutting and haptic feedback. *Progr Biophys Mol Biol*. 2010b; 103(2–3):159–168.
- Davis, TA. *CSparse*. Society for Industrial and Applied Mathematics; Philadelphia, PA: 2006.
- Delorme S, Laroche D, DiRaddo R, Del Maestro RF. Neurotouch: a physicsbased virtual simulator for cranial microneurosurgery training. *Neurosurgery*. 2012; 71:ons32–ons42.
- Dequidt, J, Duriez, C, Cotin, S, Kerrien, E. Towards interactive planning of coil embolization in brain aneurysms *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2009*. Yang, GZ, Hawkes, D, Rueckert, D, Noble, A, Taylor, C, editors. Vol. 5761. Springer; Berlin/Heidelberg: 2009. 377–385. *Lecture Notes in Computer Science*
- Dick C, Georgii J, Westermann R. A hexahedral multigrid approach for simulating cuts in deformable objects. *IEEE Trans Visual Comput Graph*. 2011; 17(11):1663–1675.
- Dick C, Georgii J, Westermann R. A real-time multigrid finite hexahedra method for elasticity simulation using cuda. *Simul Modell Pract Theory*. 2011; 19(2):801–816.
- Dryja, M, Widlund, OB. *Towards a Unified Theory of Domain Decomposition Algorithms for Elliptic Problems*. New York University, Courant Institute of Mathematical Sciences, Division of Computer Science; 1989.
- Duflot M, Bordas S. A posteriori error estimation for extended finite elements by an extended global recovery. *Int J Numer Meth Eng*. 2008; 76(8):1123–1138.
- Duriez, C; Andriot, C; Kheddar, A. Interactive haptic for virtual prototyping of deformable objects: Snap-in tasks case. *EUROHAPTICS*; Citeseer; 2003.
- Duriez, C; Andriot, C; Kheddar, A. Signorini’s contact model for deformable objects in haptic simulations. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; 2004.
- Duriez C, Dubois F, Kheddar A, Andriot C. Realistic haptic rendering of interacting deformable objects in virtual environments. *IEEE Trans Visual Comput Graph*. 2006; 12(1):36–47.
- Felippa, CA. *A Systematic Approach to the Element Independent Corotational Dynamics of Finite Elements*. Tech Rep CU-CAS-00-03. Center for Aerospace Structures; 2000.
- Felippa C, Haugen B. A unified formulation of small-strain corotational finite elements: I. theory. *Comput Meth Appl Mech Eng*. 2005; 194(21–24):2285–2335.
- Freund L. Constitutive equations for elastic–plastic materials at finite strain. *Int J Solids Struct*. 1970; 6(8):1193–1209.
- Ganovelli, F, O’Sullivan, C. Animating cuts with on-the-fly re-meshing *Eurographics 2001, Short Presentations Programme*. 2001.
- Gerstenberger A, Tuminaro RS. An algebraic multigrid approach to solve extended finite element method based fracture problems. *Int J Numer Meth Eng*. 2013; 94(3):248–2.
- Hauth M, Eitzmuß O, Straßer W. Analysis of numerical methods for the simulation of deformable models. *The Vis Comput*. 2003; 19(7–8):581–600.
- Hecht F, Lee YJ, Shewchuk JR, O’Brien JF. Updated sparse Cholesky factors for corotational elastodynamics. *ACM Trans Graph*. 2012; 31(5):X:1–X:13.

- Hermann, E; Raffin, B; Faure, F. Interactive physics simulation on multicore architectures. Proceedings of the 9th Eurographics Symposium on Parallel Graphics and Visualization (EGPGV'09); 2009 Mar.
- Hiriyur B, Tuminaro R, Waisman H, Boman E, Keyes D. A quasi-algebraic multigrid approach to fracture problems based on extended finite elements. *SIAM J Scient Comput.* 2012; 34(2):A603–A626.
- Hughes, T. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis.* Dover Publications; 2000.
- James, D; Pai, D. Artdefo: accurate real time deformable objects. Proceedings of SIGGRAPH of 26th International Conference on Computer Graphics and Interactive Techniques; ACM; 1999. 65–72.
- Jean M. The non-smooth contact dynamics method. *Comput Meth Appl Mech Eng.* 1999; 177(3–4): 235–257.
- Joldes GR, Wittek A, Miller K. Suite of finite element algorithms for accurate computation of soft tissue deformation for surgical simulation. *Med Image Anal.* 2009; 13(6):912–919. [PubMed: 19152791]
- Joldes G, Wittek A, Miller K. Real-time nonlinear finite element computations on gpu—application to neurosurgical simulation. *Comput Meth Appl Mech Eng.* 2010; 199(49):3305–3314.
- Jourdan F, Alart P, Jean M. A Gauss–Seidel like algorithm to solve frictional contact problems. *Comp Meth Appl Mech Eng.* 1998:33–47.
- Kaufman DM, Sueda S, James DL, Pai DK. Staggered projections for frictional contact in multibody systems. *ACM Trans Graph.* 2008; 27(5):1–11.
- Kerfriden P, Gosselet P, Adhikari S, Bordas S. Bridging proper orthogonal decomposition methods and augmented Newton–Krylov algorithms: an adaptive model order reduction for highly nonlinear mechanical problems. *Comput Meth Appl Mech Eng.* 2011; 200(5):850–866.
- Kerfriden P, Gourey O, Rabczuk T, Bordas S. A partitioned model order reduction approach to rationalise computational expenses in nonlinear fracture mechanics. *Comput Meth Appl Mech Eng.* 2012
- Kühnapfel U, Cakmak H, Maaß H. Endoscopic surgery training using virtual reality and deformable tissue simulation. *Comput Graph.* 2000; 24(5):671–682.
- Lamadé W, Vetter M, Hassenpflug P, Thorn M, Meinzer H-P, Herfarth C. Navigation and image-guided hbp surgery: a review and preview. *J HepatoBiliary-Pancreat Surg.* 2002; 9:592–599. DOI: 10.1007/s005340200079 [PubMed: 12541045]
- Luciano C, Banerjee P, Lemole G, Charbel F, Charbel F. Second generation haptic ventriculostomy simulator using the immersive touch system. *Stud Health Technol Inform.* 2005; 119:343.
- Malone HR, Syed ON, Downes MS, D'Ambrosio AL, Quest DO, Kaiser MG. Simulation in neurosurgery: a review of computer-based simulation environments and their surgical applications. *Neurosurgery.* 2010; 67(4):1105–1116. [PubMed: 20881575]
- Marchal, M; Allard, J; Duriez, C; Cotin, S. Towards a framework for assessing deformable models in medical simulation. Proceedings of ISBMS 2008; Springer; 2008. 176–184.
- Marchesseau, S; Heimann, T; Chatelin, S; Willinger, R; Delingette, H. Multiplicative Jacobian energy decomposition method for fast porous visco-hyperelastic soft tissue model. MICCAI'10; 2010 Sep.
- Martin, P, Ayuso, L, Torres, R, Gavilanes, A. Algorithmic strategies for optimizing the parallel reduction primitive in cudaHPCS. Smari, WW, Zeljkovic, V, editors. *IEEE*; 2012. 511–519.
- Menk A, Bordas S. Crack growth calculations in solder joints based on microstructural phenomena with x-fem. *Comput Mater Sci.* 2011; 50(3):1145–1156.
- Menk A, Bordas S. A robust preconditioning technique for the extended finite element method. *Int J Numer Meth Eng.* 2011; 85(13):1609–1632.
- Miller K, Joldes G, Lance D, Wittek A. Total lagrangian explicit dynamics finite element algorithm for computing soft tissue deformation. *Commun Numer Meth Eng.* 2007; 23(2):121–134.
- Molino N, Bao Z, Fedkiw R. A virtual node algorithm for changing mesh topology during simulation. *Proc Eurograph.* 2007:73–80.

- Mor, AB; Kanade, T. Modifying soft tissue models: progressive cutting with minimal new element creation. *Proceedings of MICCAI 2000*; 2000. 598–607.
- Müller, M; Gross, M. Interactive virtual materials. *GI '04: Proceedings of Graphics Interface 2004*; Waterloo, Ontario, Canada: Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo; 2004. 239–246.
- Naumov M. Incomplete-lu and cholesky preconditioned iterative methods using cusparse and cublas. 2011
- Nealen A, Mueller M, Keiser R, Boxerman E, Carlson M. Physically based deformable models in computer graphics. *Comput Graph Forum*. 2006;25.
- Nesme M, Marchal M, Promayon E, Chabanas M, Payan Y, Faure F. Physically realistic interactive simulation for biological soft tissues. *Recent Res Develop Biomech*. 2005; 2:1–22.
- Nicolas M, Dolbow J, Belytschko T. A finite element method for crack growth without remeshing. *Int J Numer Meth Eng*. 1999; 46:131–150.
- Niroomandi S, Alfaro I, Cueto E, Chinesta F. Real-time deformable models of non-linear tissues by model reduction techniques. *Comput Meth Programs Biomed*. 2008; 91(3):223–231.
- NVIDIA Corporation. CUBLAS Library. 2007.
- NVIDIA Corporation. CUSPARSE Library. 2007.
- Otaduy MA, Tamstorf R, Steinemann D, Gross M. Implicit contact handling for deformable objects. *Comput Graph Forum (Proc Eurograph)*. 2009; 28(2):559–568.
- Parker, EG; O'Brien, JF. Real-time deformation and fracture in a game environment. *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*; 2009 Aug. 156–166.
- Pauly, M; Pai, DK; Guibas, LJ. Quasi-rigid objects in contact. *SCA '04 Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*; 2004. 109–119.
- Payan, Y. *Soft Tissue Biomechanical Modeling for Computer Assisted Surgery*. Studies in Mechanobiology Tissue Engineering and Biomaterials. Springer; Berlin: 2012.
- Peterlik I, Nouicer M, Duriez C, Cotin S, Kheddar A. Constraint-based haptic rendering of multirate compliant mechanisms. *IEEE Trans Haptics*. 2011; 4(3):175–187. [PubMed: 26963485]
- Redon, S; Kheddar, A; Coquillart, S. Gauss' least constraints principle and rigid body simulations. *IEEE; IEEE International Conference on Proceedings of Robotics and Automation, 2002, ICRA'02*; 2002. 517–522.
- Renard Y. Generalized Newton's methods for the approximation and resolution of frictional contact problems in elasticity. *Comput Meth Appl Mech Eng*. 2013; 256:38–55.
- Ródenas J, González-Estrada O, Tarancón J, Fuenmayor F. A recovery-type error estimator for the extended finite element method based on singular + smooth stress field splitting. *Int J Numer Meth Eng*. 2008; 76(4):545–571.
- Saad, Y. *Iterative Methods for Sparse Linear Systems*. Vol. 620. PWS publishing company; Boston: 1996.
- Saad, Y. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics; Philadelphia, PA, USA: 2003.
- Saupin, G, Duriez, C, Cotin, S, Grisoni, L. Efficient contact modeling using compliance warping. *Computer Graphics International*; 2008.
- Schenk O, Bollhöfer M, Römer RA. On large scale diagonalization techniques for the Anderson model of localization. *SIAM Rev*. 2008; 50(1):91–112.
- Sifakis E, Der KG, Fedkiw R. Arbitrary cutting of deformable tetrahedralized objects. *Proc Eurograph*. 2007:73–80.
- Soler L, Delingette H, Malandain G, Montagnat J, Ayache N, Koehl C, Dourthe O, Malassagne B, Smith M, Mutter D, Marescaux J. Fully automatic anatomical, pathological, and functional segmentation from CT scans for hepatic surgery. *Comput Aided Surg*. 2001; 6(3):131–142. [PubMed: 11747131]
- Taylor, Z; Comas, O; Cheng, M; Passenger, J; Hawkes, D; Atkinson, D; Ourselin, S. Modelling anisotropic viscoelasticity for real-time soft tissue simulation. *Proceedings of MICCAI 2008*; Sep, 2008 703–710.
- The CGAL Project. *CGAL User and Reference Manual*. 3.8 Edition. CGAL Editorial Board; 2011.

- Toledo S, Chen D, Rotkin V. Taucs: A Library of Sparse Linear Solvers Version 2.2. 2003
- Tsuneoka H, Shiba T, Takahashi Y, et al. Ultrasonic phacoemulsification using a 1.4 mm incision: clinical results. *J Catar Refract Surg.* 2002; 28(1):81–86.
- Venkatesh R, Tan C, Singh G, Veena K, Krishnan K, Ravindran R. Safety and efficacy of manual small incision cataract surgery for brunescant and black cataracts. *Eye.* 2008; 23(5):1155–1157. [PubMed: 18566610]
- Wriggers, P, Panatiotopoulos, P. *New Developments in Contact Problems.* Springer; 1999. 384
- Zienkiewicz, O, Taylor, R. *The Finite Element Method.* forth ed. Vol. 1. McGraw-Hill; 1991.

Algorithm 1

Optimized algorithm to accumulate the contributions of a deformable on \mathbf{W} using the asynchronous preconditioner (i.e $\mathbf{W} = \mathbf{W} + \mathbf{H}(\mathbf{R} \mathbf{L} \mathbf{D} \mathbf{L}^T \mathbf{R}^T)^{-1} \mathbf{H}^T$).

- | | | |
|----|---|---|
| 1. | $\mathbf{J} = \mathbf{H} \mathbf{R}$ | <i>(rotate constraints)</i> |
| 2. | $\mathbf{S} = \mathbf{L}^{-1} \mathbf{J}^T$ | <i>(solve a STS for each column of \mathbf{J})</i> |
| 3. | $\mathbf{T} = \mathbf{D}^{-1} \mathbf{S}$ | <i>(apply the diagonal)</i> |
| 4. | $\mathbf{W} = \mathbf{W} + \mathbf{S}^T \mathbf{T}$ | <i>(sum the contributions)</i> |

Algorithm 2

Algorithm used to solve a Sparse Triangular System
with multiple right-hand side vectors on GPU.

```
1  bx = 0 repeat
2      accumulate_contributions(acc,diag) //seeFig. 5;
3      local_synchronization;
4       $cont[ty] = \sum_{i=0}^t acc[ty][i]$  //Parallel reduction;
5      local_synchronization;
6      solve_bloc_diagonal(cont,diag) //see Courtecuisse and allard (2009);
7      local_synchronization;
8      bx = bx + t //We treat the next t rows in parallel;
9  until bx < dim ;
```

Algorithm 3

Computation of \mathbf{U} in the correction of the SMF. Step 1 and 3 can be imply to solve respectively the lower and upper STS for each column of \mathbf{G} .

-
1. $\mathbf{S} = \mathbf{L}^{-1} \mathbf{G}^T$ *(solve the lower STS)*
 2. $\mathbf{T} = \mathbf{D}^{-1} \mathbf{S}$ *(apply the diagonal)*
 3. $\mathbf{U} = \mathbf{L}^{-T} \mathbf{T}$ *(solve the upper STS)*
-




Fig. 1.

Hepatectomy in laparoscopy: (left) camera view during a hepatic surgery and (right) constraints from the environment.




Fig. 2.

The preconditioner is updated asynchronously within a dedicated CPU thread. We use the last preconditioner available to advance the simulation so that the simulation never needs to wait for the computation of the current preconditioner to be complete.




Fig. 3. Contact force distribution in different scenarii and using different approximation of the *mechanical coupling*. (a) Is homogeneous, (b) is heterogeneous and contacts are solved without coupling, and (c) is heterogeneous and contacts are solved with coupling.

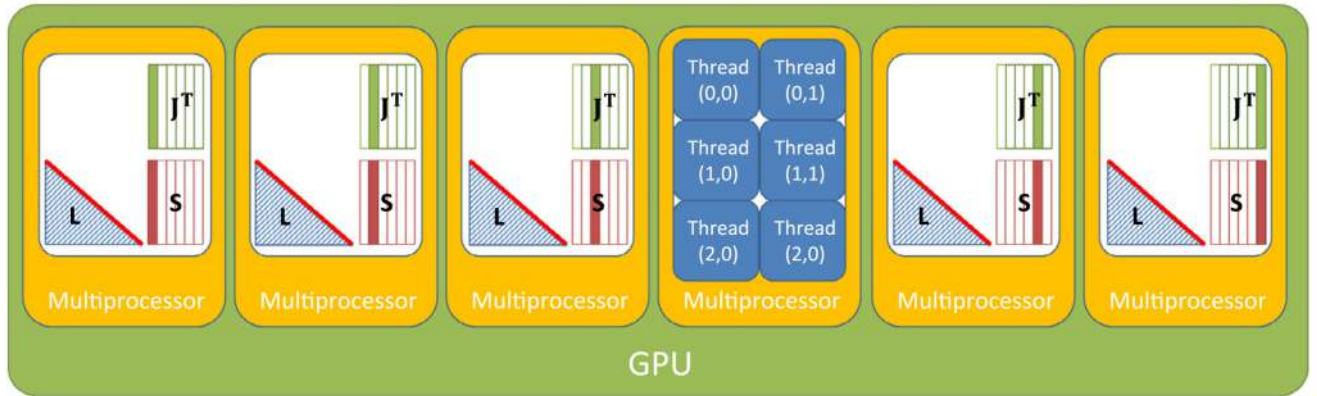


Fig. 4. First level of parallelism achieved for solving a Sparse Triangular System with multiple right hand side vector on GPU.

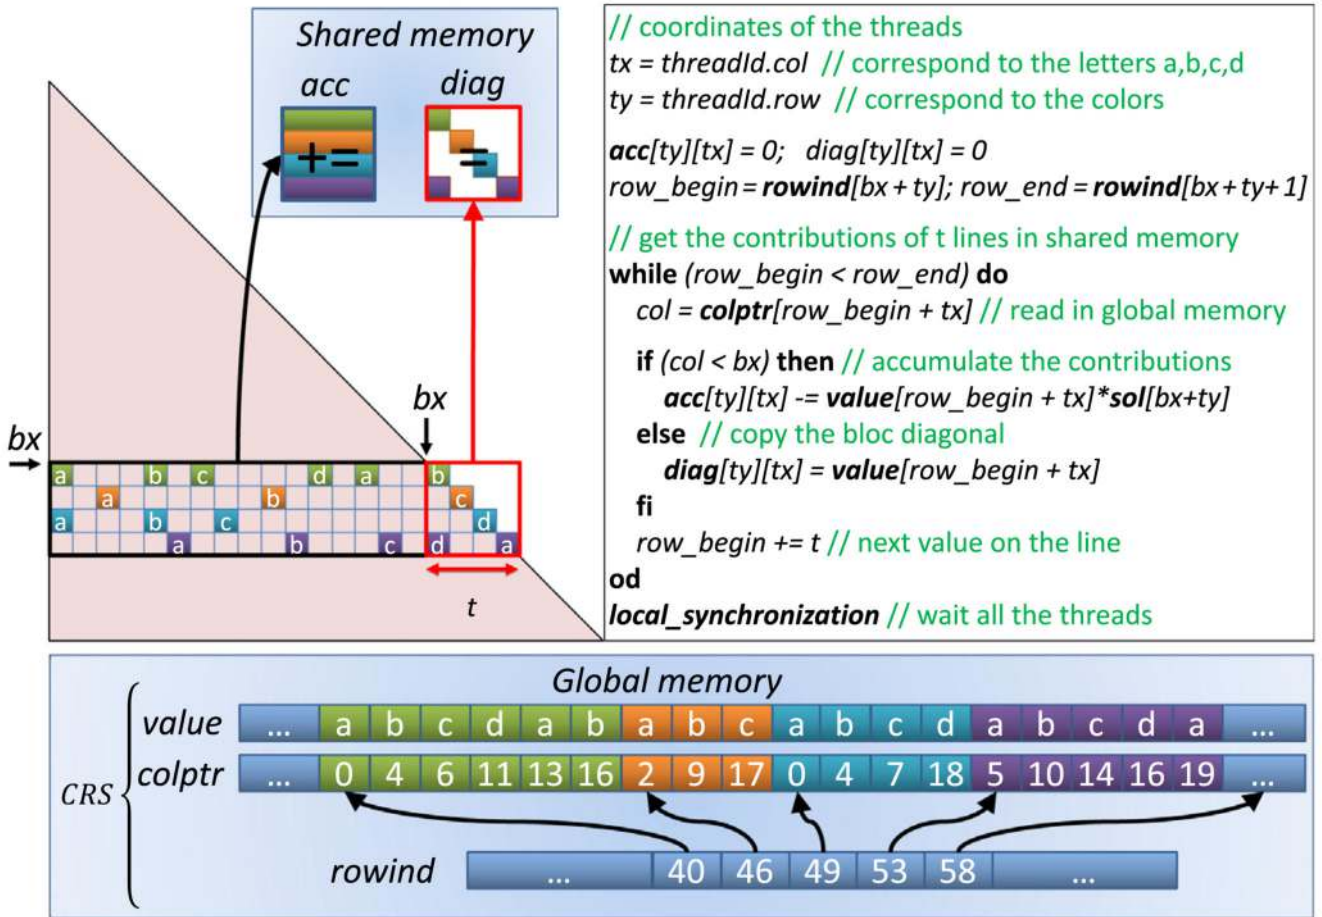


Fig. 5. Parallel accumulation of the contributions for solving a STS described by the CRS matrix. $t \times t$ threads (illustrated here with $t = 4$) are used such that t rows are processed simultaneously (colors). Each being accumulated by t threads in parallel (letters a, b, c, d). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

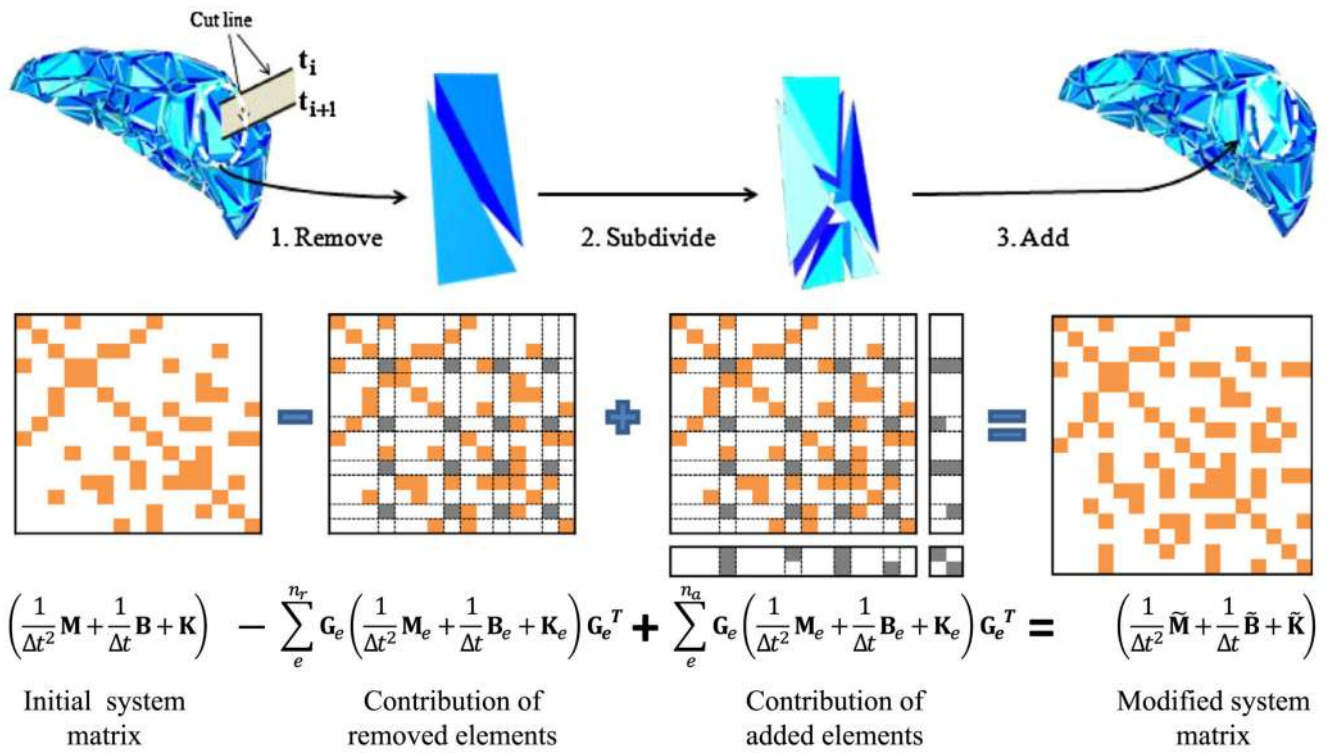


Fig. 6. Incremental update of the mesh structure for the cut.

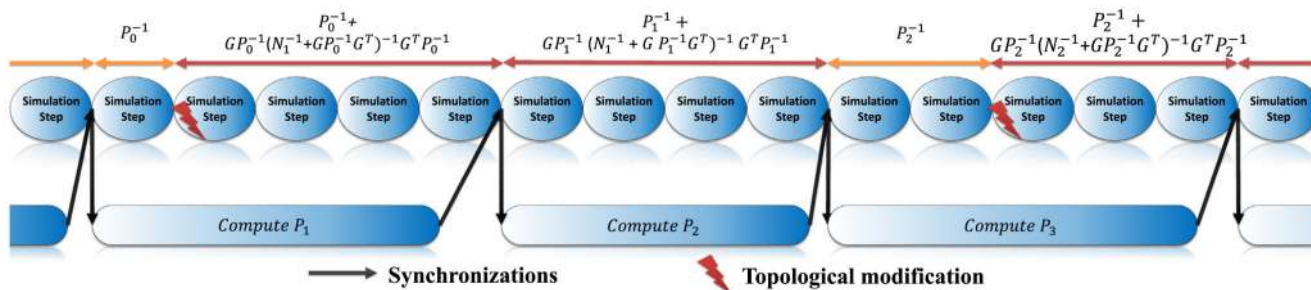





Fig. 7. Correction of the preconditioner during topological modifications. When a modification is performed on the mesh, we first compute the correction of the current factorization. Then we compute the correction of the preconditioner which was being calculated at the time of the cut. After two consecutive updates without topological modification, the preconditioner does not need any additional correction.




(a) Simulation with ABAQUS



(b) Simulation with SOFA.



(c) Difference of stresses between ABAQUS and SOFA



(d) Scaled difference of stresses between ABAQUS and SOFA

Fig. 8. Comparison of our SOFA implementation of a co-rotational model in contact with the ABAQUS solution.

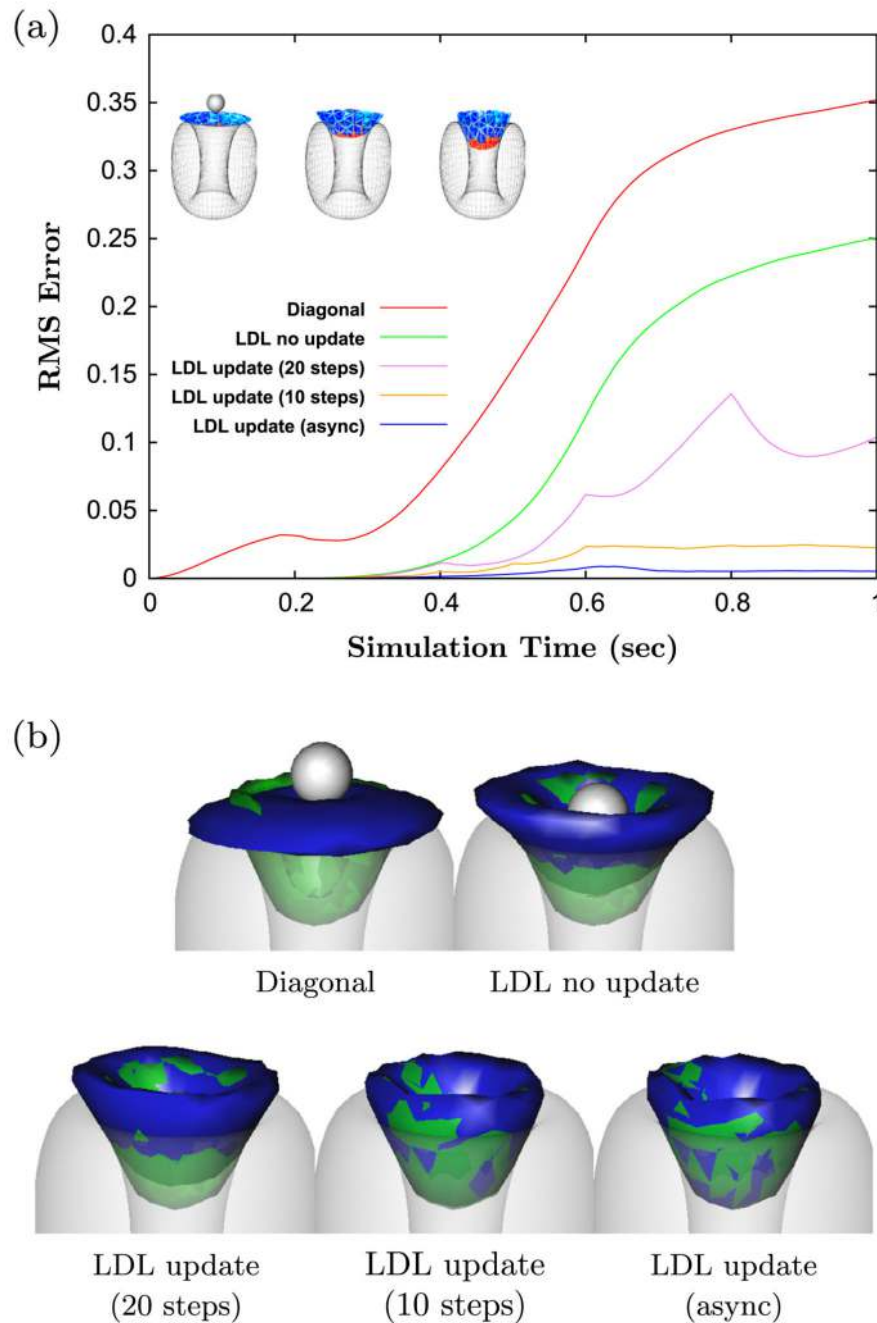


Fig. 9. Distance of nodal positions by using different approximations of the compliance matrix compared to a reference simulation. The reference simulation is obtained by computing the compliance matrix as the exact inverse of the system matrix every time step. (a) Root Mean Square error of the nodal position (compared to the reference simulation) by using different approximation as compliance matrix. (b) Simulated soft disk (blue) with reference (green) after 0.8 s of simulation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

	Method	Iterations	Computational time (ms)		
			Inverse	Solving	Total
Direct Solvers	Pardiso-1	1	78.01	1.99	80.37
	Pardiso-2	1	50.81	1.27	52.42
	Pardiso-4	1	35.34	1.05	36.80
	Pardiso-8	1	40.89	1.70	43.05
Iterative Solvers	Standard	CG	493.66	0.01	64.41
		Jacobi	314.28	14.20	53.43
	No update	LDL^T	59.11	0.02	87.86
		$\text{LDL}^T + \text{warp}$	15.12	0.24	23.74
	Async	LDL^T	9.16	0.03	14.44
		$\text{LDL}^T + \text{warp}$	6.85	0.29	12.56




Fig. 10.

Simulation of an heterogeneous deformable beam falling under gravity and average computational time for 200 simulation steps and convergence rate for different preconditioners. Red parts are 50× stiffer than bleu parts. **Inverse** corresponds to the inversion of the diagonal matrix for the Jacobi preconditioner, and to the factorization of the system for LDL^T preconditioners. **Solving** is the time taken to solve the system. **Total** is the total time of a single time step.




Fig. 11. Time in (ms) to factorize the system depending on the dimension on the system. **Simulation steps** gives the number of simulation steps that are necessary to update the asynchronous preconditioner with a simulation running at 25 FPS.

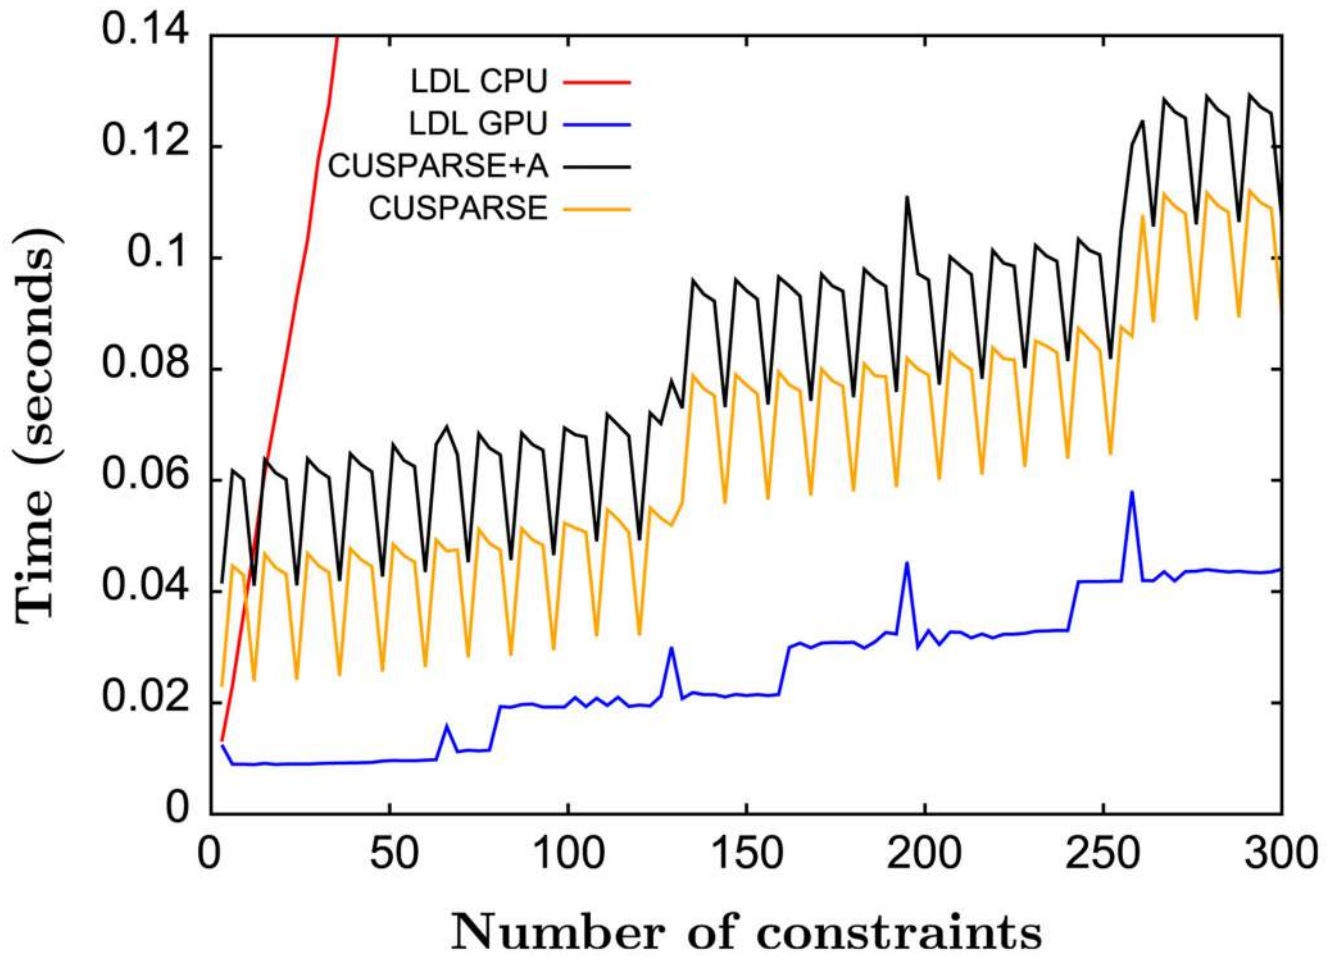


Fig. 12.
Computation time for solving a STS with multiple right-hand side vectors.

Nodes	Method	Affected Nodes	Build (ms)	Iterations	Solve (ms)	Total (ms)
1200	CG	NA	0.01	489.40	49.50	50.99
	SMF	46.39	7.30	6.59	10.41	20.25
1600	CG	NA	0.01	584	60.71	62.47
	SMF	52.50	9.00	7.26	15.70	27.57
2000	CG	NA	0.01	687.29	76.29	78.38
	SMF	54.75	10.24	7.43	21	34.40




Fig. 13.

Performances and convergence comparison for the standard CG versus our method based on the Sherman Morrison Formula. The main overheads for the computation of the correction using the SMF are reported on the graph (right). Matrices U and Q are computed for every new topological modification which correspond to the **Build** step in the table.





Fig. 14.
Simulation of the lens extraction using MSICS technique.



	Nodes	Elements	Triangles
Liver	506	1607	3680
Stomach	306	756	594
Colon	347	819	702
Diaphragm	223	582	2162
Intestine	166	486	6460

Fig. 15. Simulation of an hepatectomy with haptic feedback, and associated dataset used. (left) FE mesh used, (middle) boundary conditions from the collision detection and (right) simulation of cutting.



Fig. 16.
Real-time simulation of a brain tumor resection.