

# An efficient parallelizable multigrid framework for the simulation of elastic solids

Figure 1: Lookit! Lookit!

## Abstract

TBD

**CR Categories:** K.6.1 [Management of Computing and Information Systems]: Project and People Management—Life Cycle; K.7.m [The Computing Profession]: Miscellaneous—Ethics

**Keywords:** multigrid, elasticity, parallelization

## 1 Introduction

TBD

## 2 Background

We review certain fundamental concepts of elasticity theory, initially focusing on a static linear elasticity formulation, followed by a general outline of a multigrid solver for this elliptic problem.

### 2.1 Linear elasticity

We represent the deformation of a elastic volumetric object using a *deformation function*  $\phi$  which maps any material point  $\mathbf{X}^*$  of the undeformed configuration of the object, to its position  $\mathbf{x}^*$  in the deformed configuration, i.e.  $\mathbf{x}^* = \phi(\mathbf{X}^*)$ . For simplicity, we will use the symbol  $\phi$  to denote both this deformation function, as well as the deformed position of a material point as  $\phi^* = \phi(\mathbf{X}^*)$ . A deformation of an object gives rise to elastic forces aiming to restore the object to an equilibrium configuration [Bonet and Wood 1997]. These forces are analytically given by the divergence form

$$\mathbf{f} = -\nabla^T \mathbf{P} \quad \text{or, componentwise} \quad f_i = -\sum_j \partial_j P_{ij}. \quad (1)$$

where  $\mathbf{P}$  is the first Piola-Kirchhoff stress tensor. We note that the partial derivatives in this and all subsequent formulas are taken with respect to the material (undeformed) coordinates. The stress tensor  $\mathbf{P}$  is computed from the deformation map  $\phi$ ; the analytic expression that defines the dependence of  $\mathbf{P}$  on  $\phi$ , known as the *constitutive equation*, is an intrinsic property of every elastic material.

We will henceforth adopt the common conventions of using subscripts after a comma to denote partial derivatives, and omit certain summation symbols by implicitly summing over any right-hand side indices that do not appear on the left-hand side of a given equation. Consequently, equation (1) is compactly written as  $f_i = -P_{i,j}$ . The constitutive equation of linear elasticity is

$$\mathbf{P} = 2\mu\boldsymbol{\epsilon} + \lambda\text{tr}(\boldsymbol{\epsilon})\mathbf{I} \quad \text{or} \quad P_{ij} = 2\mu\epsilon_{ij} + \lambda\epsilon_{kk}\delta_{ij} \quad (2)$$

In this equation,  $\mu$  and  $\lambda$  are the Lamé parameters of the linear material, and are computed from Young's modulus  $E$  (a measure of material stiffness) and Poisson's ratio  $\nu$  (a measure of material incompressibility) as  $\mu = E/(2+2\nu)$ ,  $\lambda = E\nu/((1+\nu)(1-2\nu))$ . Also,  $\delta_{ij}$  is the Kronecker delta,  $\boldsymbol{\epsilon}$  is the small strain tensor

$$\boldsymbol{\epsilon} = \frac{1}{2}(\mathbf{F} + \mathbf{F}^T) - \mathbf{I} \quad \text{or} \quad \epsilon_{ij} = \frac{1}{2}(\phi_{i,j} + \phi_{j,i}) - \delta_{ij} \quad (3)$$

and  $\mathbf{F}$  is the deformation gradient tensor, defined as  $F_{ij} = \phi_{i,j}$ . Using (1,2,3) we derive the differential equation of linear elasticity

$$f_i = -\mu\phi_{i,jj} - (\mu + \lambda)\phi_{j,ij} = \mathcal{L}_{ij}\phi_j \quad (4)$$

In this equation  $\mathcal{L} = -\mu\Delta\mathbf{I} - (\mu + \lambda)\nabla\nabla^T$  is the partial differential operator of linear elasticity. A static elasticity problem amounts to determining the deformation map  $\phi$  that leads to an equilibrium of the total forces on a deformable object, i.e.  $\mathcal{L}\phi + \mathbf{f}_{\text{ext}} = 0$ , where  $\mathbf{f}_{\text{ext}}$  are the external forces applied on the object. Substituting  $\mathbf{f} = -\mathbf{f}_{\text{ext}}$ , the static elasticity problem becomes equivalent to the linear partial differential equation  $\mathcal{L}\phi = \mathbf{f}$ .

### 2.2 Multigrid correction scheme

Multigrid techniques have been predominately targeted towards elliptic partial differential equations, such as the preceding formulation of elasticity. Although there is a broad gamut of multigrid cycles and schemes, the overall philosophy is well reflected in the multigrid V-cycle correction scheme which will be described in this section. Alternative schemes will be discussed in section 13.

Multigrid methods are based on the concept of a *smoother* which is a procedure resigned to smooth, and at the same time reduce the magnitude of the residual  $\mathbf{r} = \mathbf{f} - \mathcal{L}\phi$  of the differential equation, by modifying the current estimate of the unknown function  $\phi$ . For example, if the differential equation in question has been discretized into a system of linear equations, Gauss-Seidel or Jacobi iteration could be two candidates for a simple smoother. An inherent property of elliptic systems is that a smooth distribution of residuals generally implies a certain degree of smoothness in the *error*  $e = \phi - \phi_{\text{exact}}$  as well, although a careful discretization is often needed to guarantee that this property is reflected in the discrete form of the problem. Smoothers are typically simple, local and relatively inexpensive routines, which are quite efficient at eliminating high frequencies of the residual (and, as a consequence, of the error). Nevertheless, once the high frequency component of the error has been eliminated as a result of a few applications of the smoother, subsequent iterations are characterized by rapidly decelerated convergence towards the solution. Multigrid methods seek to remediate this effect of smoother *stagnation*, using the smoother as a building block to construct a solver that achieves constant rate of convergence towards the solution, regardless of the prevailing frequencies of the residual or error. This property is accomplished by observing that any lower frequency error that persists after a few smoothing iterations will appear to be higher frequency if the problem is resampled using a coarser discretization step. By transitioning to ever coarser discretizations the smoother retains its ability to make significant progress towards the solution of the problem.

The components of a multigrid algorithm are:

- The discretization of the continuous operator  $\mathcal{L}$  at a number of different resolutions, denoted as  $\mathcal{L}^h, \mathcal{L}^{2h}, \mathcal{L}^{4h}$  and so on.
- The Smoothing subroutine, defined at each resolution.
- The Prolongation and Restriction subroutines. These implement an upsampling and downsampling operation respectively, between two different levels of resolution.

- An exact solver, used for solving the partial differential equation at the coarsest discretization resolution. Any reasonable solver may be used, as the small size of the coarsest problem is expected to lead to a negligible runtime for this subroutine.

**Table 1** Multigrid Correction Scheme – V(1,1) Cycle

```

1: procedure MULTIGRID( $\phi, \mathbf{f}, L$ )  $\triangleright \phi$  is the current estimate
2:    $\mathbf{u}^h \leftarrow \phi, \mathbf{b}^h \leftarrow \mathbf{f}$   $\triangleright$  total of  $L+1$  levels
3:   for  $l = 0$  to  $L-1$  do
4:     Smooth( $\mathcal{L}^{2^l h}, \mathbf{u}^{2^l h}, \mathbf{b}^{2^l h}$ )
5:      $\mathbf{r}^{2^l h} \leftarrow \mathbf{b}^{2^l h} - \mathcal{L}^{2^l h} \mathbf{u}^{2^l h}$ 
6:      $\mathbf{b}^{2^{l+1} h} \leftarrow \text{Restrict}(\mathbf{r}^{2^l h})$ 
7:      $\mathbf{u}^{2^{l+1} h} \leftarrow 0$ 
8:   end for
9:   Solve  $\mathbf{u}^{2^L h} \leftarrow (\mathcal{L}^{2^L h})^{-1} \mathbf{b}^{2^L h}$ 
10:  for  $l = L-1$  down to  $0$  do
11:     $\mathbf{u}^{2^l h} \leftarrow \mathbf{u}^{2^{l+1} h} + \text{Prolongate}(\mathbf{u}^{2^{l+1} h})$ 
12:    Smooth( $\mathcal{L}^{2^l h}, \mathbf{u}^{2^l h}, \mathbf{b}^{2^l h}$ )
13:  end for
14:   $\phi \leftarrow \mathbf{u}^h$ 
15: end procedure

```

Table 1 gives the pseudocode for a V(1,1) cycle of the Multigrid correction scheme. Notably, this description does not presume a specific discretization, or a particular implementation of the smoothing, restriction or prolongation operators. The following sections detail our specific implementation of these components, and the factors that motivated these design decisions.

### 3 Discretization

Our method uses a staggered finite difference discretization on uniform grids. This is a familiar concept in the field of computational fluid dynamics (e.g. [Fedkiw et al. 2001]) where staggered finite difference methods are commonplace. In contrast, these formulations are less widespread in the simulation of solids, especially for computer graphics applications, where unstructured meshes coupled with finite element or mass-spring methods are more common. This trend is generally justified by the all-around geometric and algorithmic versatility of these formulations. Nevertheless, finite difference based approaches to elasticity have been investigated [Terzopoulos et al. 1987; Terzopoulos and Fleischer 1988], and a number of authors [Müller et al. 2004; James et al. 2004; Rivers and James 2007] have turned to regular grid representations for reasons of efficiency, even for discretizations other than finite differences.

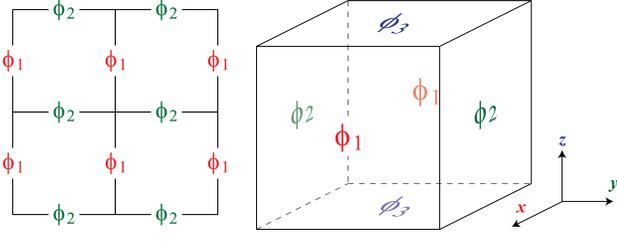
**Regular vs. Unstructured grids** Our main motivation for using a discretization based on regular grids, is avoiding the storage and access overhead of an unstructured mesh. Consider the example of a deformable model, discretized using 100K vertices. Representing the state (e.g. nodal positions) of this model would require 1.2MB of storage, using single-precision. As a rule of thumb, a tetrahedral representation using the same number of vertices would typically have more than 400K tetrahedra and require at least 6.4MB of storage to represent the mesh alone. In a parallel, streaming computing platform, this explicit representation of topology would compete with the representation of the state variables for limited cache and bandwidth resources. Additionally, unstructured meshes generally require indirect memory access or scatter-gather mechanisms which may lead to suboptimal utilization of the available bandwidth. Furthermore, uniform grids allow the use of constant stencils for the restriction and prolongation operations (and in some cases, for the discrete PDE operator  $\mathcal{L}^h$  itself), whereas this data would have to

be precomputed and streamed from memory in the case of unstructured grids. In section 12 we give further details on the working set size reductions enabled by our uniform discretization. Of course, a different measure of comparison should be established between uniform grids and unstructured, yet highly *adaptive* meshes, a topic discussed in more detail in section 13. Overall, our approach delivers very good performance at high resolution levels that compensate for the lack of a conforming or adaptive geometry.

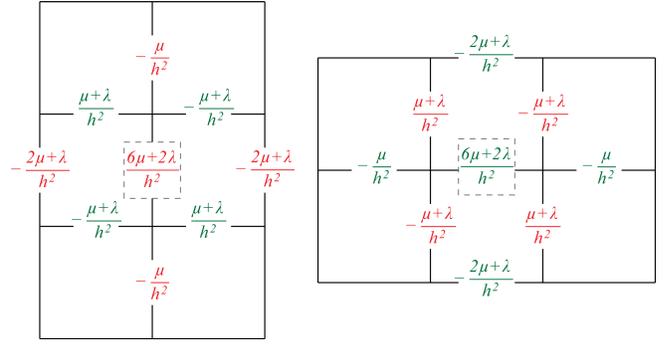
**Finite differences vs. Finite elements** In principle, either finite difference or finite element discretizations would have been viable options for our system, assuming that either method would be used with a regular grid. Our framework is based on finite differences as this method enables us to obtain a sparser, more compact discrete system of equations. This property enables computational savings and more efficient utilization of resources such as memory bandwidth and stream buffers. In 3D, for example, each of the equations of linear elasticity  $\mathcal{L}_i = -\mu \partial_{jj} - (\mu + \lambda) \partial_{ij}$  can be discretized to second-order accuracy using a stencil of 15 nonzero coefficients. For comparison, the finite element discretization of [Kazhdan and Hoppe 2008] yields 25 nonzero coefficients per equation for the much simpler 2D Poisson problem. Discretizing the 3D linear elasticity equations using trilinear finite elements and 8-point Gauss quadrature would typically require 81 nonzero coefficients per equation. Care needs to be taken however since certain useful properties of finite elements, such as symmetry and convenient treatment of certain boundaries, are not automatically guaranteed in a finite difference scheme. A significant fraction of our paper is dedicated to the implementation strategies necessary to retain these desired properties within our finite difference scheme.

**Staggered vs. Collocated grids** The deformation map  $\phi$  is a vector-valued function of the material coordinate vector  $\mathbf{X}$ . Thus  $\phi(\mathbf{X}) = (\phi_1(\mathbf{X}), \phi_2(\mathbf{X}), \phi_3(\mathbf{X}))$  where each  $\phi_i$  is a scalar-valued function. When discretizing these quantities, it would be most intuitive to use *collocated* grids, where all components of  $\phi$  are specified at the same location, for example at the nodes of a background grid. Unfortunately, for the equations of elasticity such a discretization may result in grid-scale oscillations, especially for near-incompressible materials. A comprehensive study of the causes and consequences of this behavior specifically for the equations of elasticity is beyond the scope of our current exposition. It is however qualitatively analogous to an artifact observed in the simulation of fluids with non-staggered grids, where spurious oscillations may be left over in the pressure field after a Poisson solver has been used to project a velocity field to its divergence-free component. In the context of multigrid methods, such oscillatory discretizations can be far more problematic, as they may not respect the fundamental property of elliptic PDEs that a low residual implies a smooth error, requiring more elaborate and expensive smoothers to compensate. We avoid this issue altogether by adopting a staggered discretization, which is free of this oscillatory behavior, and aligns naturally with the rest of our theoretical formulations.

Our staggered discretization is illustrated in Figure 2. A background cartesian grid serves as a reference for the placement of the unknown variables. Each component  $\phi_i$  of the deformation function  $\phi$  is stored in a separate cartesian lattice, which is offset from the nodes of the background reference grid. Specifically,  $\phi_i$  variables are stored at the centers of the background grid faces perpendicular to the cartesian axis vector  $\mathbf{e}_i$ . For example,  $\phi_1$  values are stored on grid faces perpendicular to  $\mathbf{e}_1$ , i.e. those parallel to the  $yz$ -plane. The same strategy is followed in 2D, where faces of grid cells are now identified with grid edges, thus  $\phi_1$  values are stored



**Figure 2:** Staggering of variables in 2D(left) and 3D(right). Equations  $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$  are also stored on  $\phi_1, \phi_2, \phi_3$  locations respectively.



**Figure 3:** Discrete stencils for operators  $\mathcal{L}_1$ (left) and  $\mathcal{L}_2$ (right) of the PDE system (4). The red and green nodes of the stencil correspond to  $\phi_1$  and  $\phi_2$  values respectively. The dashed square indicates the center of the stencil, where the equation is evaluated.

at the center of  $y$ -oriented edges, and  $\phi_2$  values at the center of  $x$ -oriented edges. We proceed to define the discrete approximations to first-order derivatives using central finite differences

$$\begin{aligned} D_1 u[x, y, z] &= u[x + \frac{1}{2}h_x, y, z] - u[x - \frac{1}{2}h_x, y, z] \\ D_2 u[x, y, z] &= u[x, y + \frac{1}{2}h_y, z] - u[x, y - \frac{1}{2}h_y, z] \\ D_3 u[x, y, z] &= u[x, y, z + \frac{1}{2}h_z] - u[x, y, z - \frac{1}{2}h_z] \end{aligned}$$

where  $(h_x, h_y, h_z)$  are the dimensions of the cells of the background grid. Second-order derivative stencils are simply defined as the composition of two first-order stencils, i.e.  $D_{ij} = D_i D_j$ . An implication of these definitions is that the discrete first derivative of a certain quantity will *not* be collocated with it. For example, all derivatives of the form  $D_i \phi_i$  are naturally defined at cell centers, while  $D_1 \phi_2$  is located at centers of  $z$ -oriented edges in 3D, and at grid nodes in 2D. This is not a problem, however, as all second-order derivatives are centered at the appropriate locations for a convenient discretization of (4). In particular, all stencils involved in the discretization of equation  $\mathcal{L}_i$  are naturally centered on the location of variable  $\phi_i$ . Thus, the staggering of the unknown variables implies a natural staggering of the discretized differential equations. Figure 3 illustrates this fact in 2D, where the discrete stencils for the operators  $\mathcal{L}_1$  and  $\mathcal{L}_2$  form the system (4) are shown to be naturally centered at  $\phi_1$  and  $\phi_2$  variable locations, respectively.

## 4 Construction of the Smoothing operator

The staggered discretization described in the previous section could be used essentially unmodified in the case of a relatively compressible material, i.e. with a Poisson's ratio  $\nu$  not exceeding 0.2-0.3. However, a majority of materials of interest to computer graphics, including the muscles and flesh of animated characters, are highly incompressible. A number of authors [Irving et al. 2007; English and Bridson 2007; Kaufmann et al. 2008] have discussed the challenges presented by the simulation of near-incompressible materials and proposed techniques to cope with this problem. For a multigrid solver, naive use of standard smoothers (e.g. Gauss-Seidel) in the presence of high incompressibility could lead to complications such as slow convergence or even loss of stability. We present a comprehensive and computationally inexpensive solution to this problem for linear elasticity, practically achieving performance independent of the material parameters. Moreover, in sections 7 and 8 we show how this treatment generalizes to nonlinear constitutive models.

### 4.1 Augmentation and stable discretization

Simulation of near-incompressible materials is known to cause numerical complications to a range of solvers, including techniques based on finite element, finite difference or other discretizations. In particular, our PDE formulation also requires special treatment in this case. When Poisson's ratio approaches the incompressible

limit  $\nu \rightarrow 0.5$ , the Lamé parameter  $\lambda$  becomes several orders of magnitude larger than  $\mu$ . As a consequence, the dominant term of the elasticity operator  $\mathcal{L} = -\mu \Delta \mathbf{I} - (\mu + \lambda) \nabla \nabla^T$  is the rank deficient operator  $-(\mu + \lambda) \nabla \nabla^T$ ; thus  $\mathcal{L}$  becomes *near-singular*. More specifically, we see that any divergence-free field  $\phi$  will be in the nullspace of the dominant term, i.e.  $-\lambda \nabla \nabla^T \phi = 0$ . Thus, a solution to the elasticity PDE  $\mathcal{L} \phi = \mathbf{f}$  could be perturbed by a divergence-free displacement of substantial amplitude, without introducing a large residual for the differential equation. In the context of a multigrid scheme, this observation has far deeper implications other than indicating that the discretized system of equations will have a high condition number; here, the eigenspace corresponding to the smallest eigenvalues of the operator  $\mathcal{L}$  contains the entire class of divergence-free functions. These can be arbitrarily oscillatory, and lead to high-frequency errors that the multigrid method cannot smooth efficiently or correct using information from a coarser grid. Fortunately, this complication is not a result of inherently problematic material behavior, but rather an artifact of the form of the governing equations chosen to describe this physical phenomenon. Our solution is to reformulate the PDEs of elasticity into an equivalent system, which does not suffer from the near-singularity of the original differential operator. This stable differential description of near-incompressible elasticity is adapted from the theory of mixed variational formulations [Brezzi and Fortin 1991] and was demonstrated by [Gaspar et al. 2008] in simple academic problems of linear elasticity. Our work extends these formulations to nonlinear materials and domains with arbitrary boundaries.

We introduce a new auxiliary variable  $p$  (which we will refer to as the *pressure*) defined as  $p = -(\lambda/\mu) \nabla \nabla^T \phi = -(\lambda/\mu) \text{div} \phi$ . Note that, although  $p$  is a scaled divergence of the deformation field, we do not pursue or depend on any explicit associations of this variable with any "physical" pressure quantity. We can now write

$$\begin{aligned} \mathcal{L} \phi &= -\mu \Delta \phi - (\mu + \lambda) \nabla \nabla^T \phi \\ &= -\mu (\Delta \mathbf{I} + \nabla \nabla^T) \phi - \lambda \nabla (\nabla^T \phi) \\ &= -\mu (\Delta \mathbf{I} + \nabla \nabla^T) \phi + \mu \nabla p \end{aligned} \quad (5)$$

As a result, the equilibrium equation  $\mathcal{L} \phi = \mathbf{f}$  can be equivalently written as the system

$$\begin{pmatrix} -\mu (\Delta \mathbf{I} + \nabla \nabla^T) & \mu \nabla \\ \mu \nabla^T & \frac{\mu^2}{\lambda} \end{pmatrix} \begin{pmatrix} \phi \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ 0 \end{pmatrix} \quad (6)$$

The top part of system (6) follows directly from equation (5), while the bottom equation is simply a rescaled version of the definition of

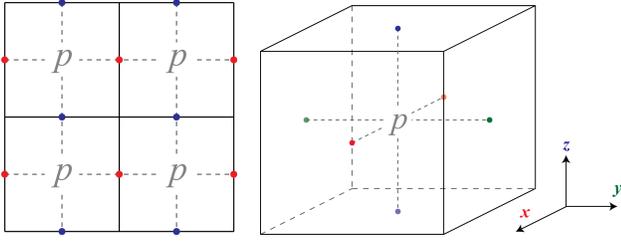


Figure 4: Placement of pressures in 2D (left) and 3D (right).

271 pressure  $p$ . Conversely, the original differential equation (4) can be  
 272 obtained from (6) by simply eliminating the pressure variable. Thus  
 273 the *augmented* differential equation system of (6) is equivalent to  
 274 the governing equations of linear elasticity (e.g. the two systems  
 275 agree in the value of  $\phi$  when solved). The important consequence  
 276 of this manipulation is that this new discretization is *stable*, in the  
 277 sense that the eigenspace corresponding to the smaller eigenvalues  
 278 of this augmented differential operator does not contain high fre-  
 279 quency deformation modes. This property can be rigorously proven  
 280 via Fourier analysis; we can verify however that as  $\lambda$  tends to infin-  
 281 ity, the term  $\mu^2/\lambda$  vanishes, and the resulting limit system is now  
 282 *non-singular*. This stability property indicates that it is possible to  
 283 smooth the error of this system efficiently with inexpensive, local  
 284 smoothers. Such a smoother is described in detail in section 4.2.

285 The newly introduced pressure variables are also discretized on an  
 286 offset cartesian lattice, with each pressure begin stored in the center  
 287 of a cell of the background lattice (see Figure 4). Pressure *equa-*  
 288 *tions*, i.e. the last row of system (6), are also centered and stored  
 289 at cell centers. As was the case with the non-augmented elastic-  
 290 ity equations, the staggering of deformation ( $\phi$ ) and pressure ( $p$ )  
 291 variables is such that all discrete first and second order differential  
 292 operators are well defined where they are naturally needed. Finally,  
 293 as a consequence of our staggered discretization, equations (4) and  
 294 (6) are equivalent at the differential *and* at the discretization level,  
 295 i.e. the discrete forms of the equations are algebraically identical,  
 296 after eliminating pressures from the augmented form.

## 297 4.2 Distributive smoothing

298 Although the augmented system (6) has the necessary stability to  
 299 admit, in principle, efficient local smoothing, this cannot be ac-  
 300 complished with a standard Gauss-Seidel or Jacobi iteration, as the  
 301 discrete augmented equations lack the formal convergence guar-  
 302 antees of these methods. First, we note that the discrete form  
 303 of system (6) is not a symmetric definite matrix; the upper left-  
 304 most block  $-\mu(\Delta\mathbf{I} + \nabla\nabla^T)$  is symmetric positive definite but the  
 305 discrete form of the off-diagonal blocks ( $\mu\nabla$  and  $\mu\nabla^T$ ) is actu-  
 306 ally *skew-symmetric* as the discrete first order operators satisfy  
 307  $D_i^T = -D_i$ . Also, negating the pressure equation to make the sys-  
 308 tem symmetric, would still render it indefinite. Secondly, in the  
 309 incompressible limit the diagonal constant term  $\mu^2/\lambda$  vanishes, so  
 310 neither Gauss-Seidel nor Jacobi methods would be usable.

311 Apart from these technical difficulties, it is generally known that  
 312 for a differential equation such as (6) exhibiting nontrivial cou-  
 313 pling between the variables  $\phi_1, \phi_2, \phi_3$  and  $p$ , a smoothing scheme  
 314 which updates several variables at once is often the optimal choice  
 315 in terms of efficiency [Trottenberg et al. 2001]. We note that this is  
 316 not the same as a more costly block smoother where a larger num-  
 317 ber of *equations* are solved simulatenously; we still process one  
 318 equation at a time, but the residual is eliminated by changing the  
 319 value of several variables, rather than just one. We adopt the *dis-*  
 320 *tributive smoothing* scheme introduced by [Gaspar et al. 2008] for

321 linear elasticity, which we later extend to nonlinear problems. Let  
 322 us redefine  $\mathcal{L}$  to denote the augmented differential operator of equa-  
 323 tion (6), and write  $\mathbf{u} = (\phi, p)$  for the augmented set of unknowns  
 324 and  $\mathbf{b} = (\mathbf{f}, 0)$  for the right-hand side vector. Thus, system (6) is  
 325 written as  $\mathcal{L}\mathbf{u} = \mathbf{b}$ . Consider the following change of variables

$$\begin{pmatrix} \phi \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{I} & -\nabla \\ \nabla^T & -2\Delta \end{pmatrix} \begin{pmatrix} \psi \\ q \end{pmatrix} \quad \text{or} \quad \mathbf{v} = \mathcal{M}\mathbf{u} \quad (7)$$

326 where  $\mathbf{v} = (\psi, q)$  is the vector of auxiliary unknown variables,  
 327 and  $\mathcal{M}$  is called the *distribution matrix*. In accordance with our  
 328 staggered formulation, the components  $\psi_1, \psi_2, \psi_3$  of the auxiliary  
 329 vector  $\psi$  will be collocated with  $\phi_1, \phi_2, \phi_3$  respectively, while  $q$   
 330 and  $p$  values are collocated as well. Using the change of variables  
 331 of equation (7), our augmented system  $\mathcal{L}\mathbf{u} = \mathbf{b}$  is equivalently  
 332 written as  $\mathcal{L}\mathcal{M}\mathbf{v} = \mathbf{b}$ . Composing the operators  $\mathcal{L}$  and  $\mathcal{M}$  yields

$$\mathcal{L}\mathcal{M} = \begin{pmatrix} -\mu\Delta\mathbf{I} & 0 \\ \mu(1 + \frac{\mu}{\lambda})\nabla^T & -\mu(1 + \frac{2\mu}{\lambda})\Delta \end{pmatrix} \quad (8)$$

333 That is, the composed system is lower triangular, and its diago-  
 334 nal elements are simply Laplacian operators. This system can be  
 335 smoothed with any scheme that works for the Poisson equation,  
 336 including the Gauss-Seidel or Jacobi methods. In fact, the entire  
 337 system can be smoothed practically with the same efficiency as the  
 338 Poisson equation, following a forward substitution approach, i.e.  
 339 we smooth all  $\psi_1$ -centered equations across the domain first, fol-  
 340 lowed by sweeps of  $\psi_2, \psi_3$ , and  $q$ -centered equations in sequence.

One seeming obstacle to realizing these benefits, is that we do not  
 have the auxiliary variables  $(\psi, q)$  at our disposal. As a matter of  
 fact, computing  $(\psi, q)$  from  $(\phi, p)$  would necessitate solving sys-  
 tem (7). Fortunately, such an explicit transformation is not neces-  
 sary. We start by reviewing the standard Gauss-Seidel iteration for  
 solving (or smoothing) the system  $\mathcal{L}\mathbf{u} = \mathbf{b}$ . At every step of the  
 iteration, we focus on a different equation  $\mathcal{L}_i$  (here  $i$  indicates a  
 single discrete equation, as opposed to the three coordinate compo-  
 nents of  $\mathcal{L}$ ). Each Gauss-Seidel step amounts to calculating a point-  
 wise correction to the variable  $u_i$ , collocated with the equation  $\mathcal{L}_i$ ,  
 such that the residual of  $\mathcal{L}_i$  vanishes. In more detail, we seek to  
 replace variable  $u_i$  with  $u_i + \delta$ , or equivalently  $\mathbf{u}$  with  $\mathbf{u} + \delta\mathbf{e}_i$ .  
 As a result of this value change, the residual of the equation becomes  
 $\mathbf{r} = \mathbf{b} - \mathcal{L}(\mathbf{u} + \delta\mathbf{e}_i)$ . The unknown variable  $\delta$  is determined by  
 requiring that  $r_i = \mathbf{e}_i^T \mathbf{r}$  becomes zero after this correction, thus

$$\mathbf{e}_i^T (\mathbf{b} - \mathcal{L}(\mathbf{u} + \delta\mathbf{e}_i)) = 0 \Rightarrow (\mathbf{e}_i^T \mathcal{L}\mathbf{e}_i)\delta = \mathbf{e}_i^T (\mathbf{b} - \mathcal{L}\mathbf{u})$$

341 The last equation is equivalent to  $\mathcal{L}_{ii}\delta = r_i^{\text{old}}$  or  $\delta = r_i^{\text{old}}/\mathcal{L}_{ii}$ ,  
 342 where  $\mathcal{L}_{ii}$  is the  $i$ -th diagonal element of the discrete operator and  
 343  $r_i^{\text{old}}$  denotes the  $i$ -th component of the residual vector *before*  
 344 the correction. Operating in an analogous fashion, a Gauss-Seidel step  
 345 on the *distributed* system  $\mathcal{L}\mathcal{M}\mathbf{v} = \mathbf{b}$  amounts to changing  $\psi_i$  into  
 346  $\psi_i + \delta$ , or equivalently  $\mathbf{v}$  into  $\mathbf{v} + \delta\mathbf{e}_i$ , such that the  $i$ -th residual of  
 347 the distributed equation is annihilated, as follows

$$\begin{aligned} \mathbf{e}_i^T (\mathbf{b} - \mathcal{L}\mathcal{M}(\mathbf{v} + \delta\mathbf{e}_i)) &= 0 \Rightarrow \mathbf{e}_i^T (\mathbf{b} - \mathcal{L}(\mathbf{u} + \delta\mathcal{M}\mathbf{e}_i)) = 0 \\ &\Rightarrow (\mathbf{e}_i^T \mathcal{L}\mathcal{M}\mathbf{e}_i)\delta = \mathbf{e}_i^T (\mathbf{b} - \mathcal{L}\mathbf{u}) \Rightarrow \delta = r_i^{\text{old}}/(\mathcal{L}\mathcal{M})_{ii} \end{aligned}$$

348 In this derivation we leveraged the fact that the auxiliary vector  $\mathbf{v}$   
 349 is only used in the form  $\mathcal{M}\mathbf{v}$  which is equal to the value of the origi-  
 350 nal variable  $\mathbf{u}$ . After the value of  $\delta$  has been determined,  $\mathbf{u}$  can be  
 351 updated to  $\mathbf{u} + \delta\mathcal{M}\mathbf{e}_i$ . This update can be put in a more conven-  
 352 ient form by observing that the discrete operator  $\mathcal{M}$  is symmetric  
 353 (discrete first-order derivative operators are skew-symmetric), thus  
 354  $\mathcal{M}\mathbf{e}_i$  can be taken form either the  $i$ -th column *or* row of the dis-  
 355 cretized operator. Therefore, the correction  $\mathbf{u} \leftarrow \mathbf{u} + \delta\mathcal{M}\mathbf{e}_i$  is

356 simply implemented by adding a  $\delta$ -scaled version of the finite difference stencil of the  $i$ -th row of  $\mathcal{M}$  to the current solution  $\mathbf{u}$ . Finally, the diagonal element  $e_i^T \mathcal{L} \mathcal{M} e_i$  can be precomputed as the inner product of the stencils for the  $i$ -th row of  $\mathcal{L}$  and the  $i$ -th row of  $\mathcal{M}$  (again, due to symmetry). The computational cost of distributive smoothing is comparable to that of simple Gauss-Seidel iteration, yet it allows efficient smoothing of the equations of linear elasticity, independent of Poisson's ratio. We summarize the distributive smoothing process in pseudo-code in Table 2.

**Table 2** Distributive Smoothing

```

1: procedure DISTRIBUTIVESMOOTHING( $\mathcal{L}, \mathcal{M}, \mathbf{u}, \mathbf{b}$ )
2:   for  $v$  in  $\{\phi_1, \phi_2, \phi_3, p\}$  do    ▷ Must be in this exact order
3:     for  $i$  in Lattice[ $v$ ] do        ▷  $i$  is an equation index
4:       DISTRIBUTIVESMOOTHINGSTEP( $\mathcal{L}, \mathcal{M}, \mathbf{u}, \mathbf{b}, i$ )
5:     end for
6:   end for
7: end procedure
8: procedure DISTRIBUTIVESMOOTHINGSTEP( $\mathcal{L}, \mathcal{M}, \mathbf{u}, \mathbf{b}, i$ )
9:    $r \leftarrow b_i - \mathcal{L}_i \cdot \mathbf{u}$         ▷  $\mathcal{L}_i$  is the  $i$ -th row of  $\mathcal{L}$ 
10:   $\delta \leftarrow r / (\mathcal{L} \mathcal{M})_{ii}$ 
11:   $\mathbf{u} += \delta m_i^T$                     ▷  $m_i$  is the  $i$ -th row of  $\mathcal{M}$ 
12: end procedure

```

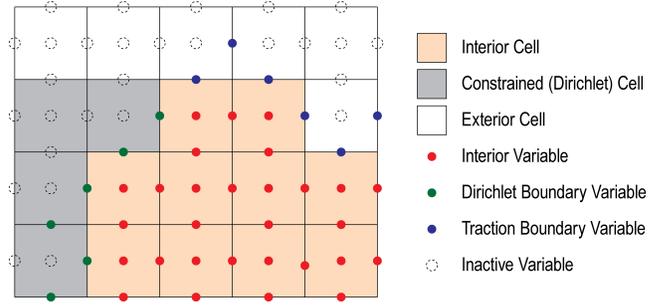
## 5 Treatment of boundaries

366 The discretization and smoothing procedures discussed in the previous sections did not address the effect of boundaries, focusing on the treatment of the interior region of the simulated deformable body. In fact, we have been able to evaluate the validity and efficiency of the preceding formulations using a periodic domain, which is devoid of any boundaries. Our findings are reported in section 12 and can serve as a theoretical upper bound of multigrid efficiency. Multigrid theory suggests that a general boundary value problem can be solved at the same efficiency as a periodic problem, at the expense of more intensive smoothing effort at the boundary. In theoretical studies of multigrid efficiency, the computational overhead of this additional boundary smoothing is often overlooked, as the cost of interior smoothing is *asymptotically* expected to dominate. Nevertheless, in our experience, practical problem sizes may never reach this asymptotic regime, and generic boundary smoothing approaches could become a performance bottleneck, even for problems with millions of degrees of freedom. In this section, we develop a boundary discretization strategy, including a novel treatment of traction boundary conditions, that facilitates the design of efficient and inexpensive boundary smoothers.

### 5.1 Domain description

387 Our geometrical description of the computational domain is based on a partitioning of the cells of the background grid. Initially, cells that have an overlap with the simulated deformable body are characterized as *interior* cells, otherwise they are designated *exterior* cells. Additionally, any cell can be user-specified to be a *constrained* (or *Dirichlet*) cell. Specifying a Dirichlet cell overrides any interior/exterior designation it may otherwise carry. Geometrically, the interface between interior and Dirichlet cells corresponds to the boundary of the computational domain where Dirichlet boundary conditions are given, while traction (or free) boundary conditions are imposed on the interface between interior and exterior cells. Intuitively, Dirichlet cells correspond to kinematically constrained parts of the object, such as the skeleton of an articulated character. This partitioning of the domain is illustrated in Figure 5.

401 This definition provides an intuitive way to specify the degrees of



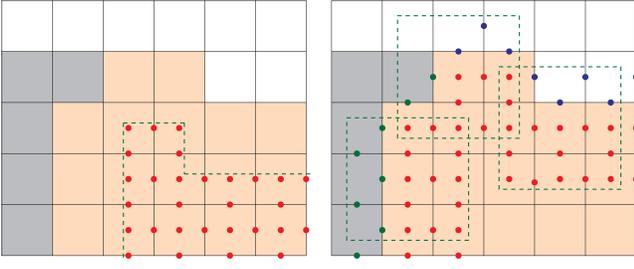
**Figure 5:** Classification of cells and variables near the boundary

402 freedom of our problem and their associated equations. Any of the variables  $\phi_1, \phi_2, \phi_3$  or  $p$  located strictly inside the interior region (i.e. either on an interior cell center, or on the face between two interior cells) is designated an *interior* variable. For every interior variable we also introduce in our system the discrete equation of (6) centered at the same location. Any variable which is included in the discrete stencil of an interior equation, but is not interior itself, is a *boundary* variable; such a variable will be further designated a Dirichlet boundary variable if it touches a Dirichlet cell (either inside or on the boundary of one), otherwise it is designated a traction boundary variable. Variables not appearing in the stencil of any interior equation are labeled *inactive* and can generally be ignored.

### 5.2 A general-purpose box smoother

415 We first describe a general-purpose treatment of the boundary equations and variables. In order to allow for a unique solution to our discrete problem, one additional equation must be provided for each boundary variable. Note that, as a result of our domain description, there will be no boundary *pressure* variables, i.e. no exterior pressure ever appears on the discrete stencil of an interior equation. For every Dirichlet boundary variable  $\phi_i$  we need to specify a Dirichlet condition of the form  $\phi_i(\mathbf{X}^*) = \phi_i^*$ , while every traction boundary variable  $\phi_i$  will be naturally matched with a traction condition  $e_i^T \mathbf{P}(\mathbf{X}^*) \mathbf{N} = t^*$ , where  $\mathbf{N}$  is the normal vector to the object surface ( $t^*=0$  would be the zero-traction or *free* boundary condition). The point  $\mathbf{X}^*$  used in these boundary equations need not coincide with the location of the boundary variables, allowing the flexibility to use a boundary condition defined on the object surface even when the associated boundary variable is offset from it due to staggering. As a consequence, averaging (for Dirichlet equations) or one-sided finite differences (for traction equations) may be needed in the discrete formulation of these boundary equations. In general, any first-order accurate (or better) finite difference approximation for either type of boundary condition is acceptable for our purposes.

435 Although a well-posed system can be constructed as described, the distributive smoothing scheme cannot be used in the immediate vicinity of the boundary. In that region, the distributive correction for certain variables extends outside the domain, affecting boundary and even inactive variables. Moreover, the boundary equations themselves need to be smoothed, and that cannot be accomplished by simply substituting a different smoother for the specific equations where distributive smoothing is not applicable. In such situations a *box smoother* is a broadly applicable solution. This process amounts to collectively solving a number of equations in a rectangular box, simultaneously adjusting the values of all variables within that region. More formally, we compute a collective correction  $\delta = (\delta_{i_1}, \dots, \delta_{i_N})$  such that the  $N$  equations  $i_1, \dots, i_N$  will be simultaneously satisfied after the correction has been applied. The correction vector can be obtained as the solution of the equation



**Figure 6:** Left: Extent of distributive smoothing (interior region), Right: Boundary region with some boxes used by the box smoother.

found the Kaczmarz smoother to converge extremely slowly (a fact well documented in the literature) and consequently requiring a very large number of iterations, making it a very sub-optimal solution. Our proposed solution stems from a novel perspective of the constitutive equations *and* the boundary conditions that results in a well conditioned, symmetric positive definite boundary system.

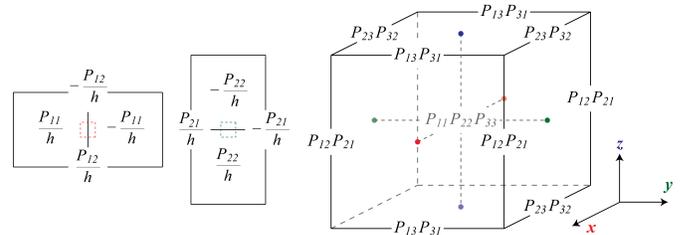
First, we revisit the constitutive equation of linear elasticity (2). The scalar coefficient  $\text{tr}(\epsilon)$  appearing in equation (3) is equivalently written as  $\text{tr}(\epsilon) = \sum_i \epsilon_{ii} = \sum_i \phi_{i,i} - d$ , where  $d = \text{tr}(\mathbf{I})$  equals the number of spatial dimensions. Similarly, the last equation of system (6) is equivalent to  $-(\mu/\lambda)p = \nabla^T \phi = \sum_i \phi_{i,i}$ . Thus, we have  $\text{tr}(\epsilon) = -(\mu/\lambda)p - d$ , and equation (2) becomes

$$\mathbf{P} = \mu(\mathbf{F} + \mathbf{F}^T) - \mu p \mathbf{I} - (2\mu + d\lambda)\mathbf{I} \quad (9)$$

The difference between equations (2) and (9) is that the original definition of stress is physically valid for any given deformation field  $\phi$  while the formulation of equation (9) will correspond to the real value of stress only when the augmented system (6) is solved exactly. We can verify that the *position* equations  $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$  of system (9) are equivalent to the divergence form  $\mathcal{L}_i u = -\partial_j P_{ij}$ , where  $\mathbf{P}$  is now given by the new definition of equation (9). The discrete stencils for these equations can be constructed as a two-step process. First, we construct a finite difference stencil for the expression  $-\partial_j P_{ij}$ , treating every value  $P_{ij}$  appearing in this stencil as a separate variable (see Figure 7, left). As a second step, finite difference approximations are substituted in place of the  $P_{ij}$  values.

For interior equations, this process yields exactly the same results as the direct discretization of system (6). Certain interior equations near the boundary, however, are special in the sense that their stencil extends onto *boundary* variables. For those equations, we turn our attention to the stress variables appearing in their discrete divergence form. Such a stress variable  $P_{ij}$  will be characterized as *interior* if the discrete stencil for  $P_{ij}$  uses only interior or Dirichlet variables, and *exterior* if its stencil uses at least one traction boundary variable (see Figure 8, left). In our proposed formulation, any exterior stresses will not be evaluated by means of a finite difference stencil; instead a specific value will be substituted for them, using an appropriate traction boundary condition. More specifically:

- Stress variables of the form  $P_{ij}$  ( $i \neq j$ ) are centered on grid edges in 3D (see Figure 7, right) and on grid nodes in 2D. This stress variable appeared in the finite difference approximation of the term  $-\partial_j P_{ij}$  in equation  $\mathcal{L}_i$ . Let  $\mathbf{X}^*$  be the location where equation  $\mathcal{L}_i$  is centered. The stress variable  $P_{ij}$  is located one half of a grid cell away from  $\mathbf{X}^*$ , *along the direction*  $\mathbf{e}_j$ . Without loss of generality, assume  $P_{ij}$  is located at  $\mathbf{X}^* + \frac{h_j}{2} \mathbf{e}_j$ .  $P_{ij}$  neighbors exactly four cells; out of those, the two centered at  $\mathbf{X}^* \pm \frac{h_i}{2} \mathbf{e}_i$  are *interior* cells, since we assumed that  $\mathcal{L}_i$  was an interior equation. The two other neighbor cells of  $P_{ij}$  are centered at  $\mathbf{X}^* \pm \frac{h_i}{2} \mathbf{e}_i + h_j \mathbf{e}_j$ . We



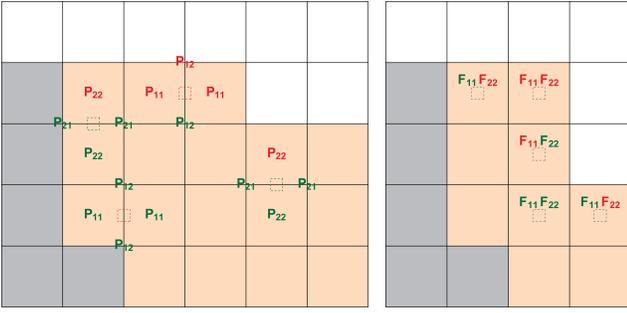
**Figure 7:** Left: Equations  $\mathcal{L}_1, \mathcal{L}_2$  expressed as divergence stencils. Right: Placement of the components of stress tensor  $\mathbf{P}$  in 3D.

$\mathcal{L}^* \delta = \mathbf{r}^*$ , where  $\mathcal{L}^*$  is the  $N \times N$  submatrix of  $\mathcal{L}$  corresponding to the rows and columns indexed  $i_1$  through  $i_N$ , and  $\mathbf{r}^*$  contains the corresponding  $N$  entries of the residual vector before the correction. Our complete smoothing subroutine starts with a boundary box smoothing sweep, proceeds with a sweep of interior distributive smoothing and finishes with a last boundary pass. The criterion for performing distributive smoothing on a certain interior equation is that the discrete stencil of the corresponding *distributed* equation, i.e. the  $i$ -th equation of the composed system  $(\mathcal{LM})\mathbf{v} = \mathbf{b}$  needs to contain only interior variables, as illustrated in Figure 6 (left). During the boundary sweep, we collectively solve all equations in overlapping boxes that are two grid cells wide, and centered at the centers of all the outermost layer of interior cells, as seen in Figure 6 (right). The local system  $\mathcal{L}^* \delta = \mathbf{r}^*$  can be pre-factorized using a pivoted LU decomposition and solved by means of forward and back substitution. In our experiments the box smoother performed very well, generally allowing the entire multigrid scheme to converge at the interior efficiency (i.e. achieving the convergence rates observed for a periodic problem).

### 5.3 A fast symmetric Gauss-Seidel smoother

Although the box smoother described in section 5.2 was effective in achieving a good convergence rate, the runtime overhead associated with it was enormous. For a problem with 32K vertices, the execution time of boundary smoothing was about 60 – 80 times longer than that of the interior distributive smoothing. Asymptotically, the cost of boundary smoothing is  $O(N^{2/3})$ , while the distributive smoothing has an asymptotic  $O(N)$  complexity, where  $N$  is the number of vertices in the simulation. These asymptotics, however, will not materialize into a tangible benefit for realistic problem sizes. Even for models with 2M vertices, boundary smoothing would still require more than 15 times the runtime of interior smoothing. Furthermore, when using a direct solver for box smoothing in a parallel machine we would waste substantial memory bandwidth to streaming the precomputed LU factors, and tolerating on-the-fly factorization would likely be an even costlier alternative.

We propose a novel formulation that enables equation-by-equation smoothing that is both efficient and inexpensive. The main obstacle to designing efficient equation-by-equation boundary smoothing schemes, is that the discrete system of equations near the boundary lacks properties such as symmetry, definiteness or diagonal dominance. In particular, loss of definiteness is predominantly a side-effect of the augmented discretization (6) the distributive smoother owes its efficiency to. Additionally, even natural discretizations of the boundary conditions (especially for the traction boundary) can easily result in loss of symmetry, even in our non-augmented system (4). An alternative local smoother would be the Kaczmarz method [Trottenberg et al. 2001], which does not require symmetry or definiteness of the boundary system; we have nevertheless



**Figure 8:** *Left: Stress variables used in the divergence form of certain interior equations. Boundary stress variables are colored red, interior stresses are green. All boundary stresses can be set to a specific value using a traction condition from a nearby boundary. Right: Interior and boundary gradients used in pressure equations.*

545 can verify that if those two cells were interior or Dirichlet,  $P_{ij}$   
 546 would have been an *interior* stress, i.e. its finite difference stencil  
 547 would have only included interior or Dirichlet variables. Thus, if  $P_{ij}$   
 548 is an exterior stress, one of the cells centered at  $\mathbf{X}^* \pm \frac{h_i}{2}\mathbf{e}_i + h_j\mathbf{e}_j$   
 549 must be exterior. This means that  $P_{ij}$  is incident on a traction boundary  
 550 face perpendicular to the direction  $\mathbf{e}_j$ , thus the traction condition  $\mathbf{P}\mathbf{e}_j = \mathbf{t}$   
 551 associated with this part of the boundary specifies a value  $P_{ij} = t_i$  for  
 552 this component of the stress. In the common case of a free boundary the  
 553 traction value is simply zero, giving  $P_{ij} = 0$ .  
 554

- 555 • Stress variables of the form  $P_{ii}$  are located at cell centers, and appear  
 556 in the finite difference approximation of  $-\partial_i P_{ii}$  in equation  $\mathcal{L}_i$ . Similar  
 557 to the previous case,  $P_{ii}$  is located one half grid away from the location  
 558  $\mathbf{X}^*$  of  $\mathcal{L}_i$  along the direction  $\mathbf{e}_i$ . Without loss of generality, assume  
 559  $P_{ij}$  is located at  $\mathbf{X}^* + \frac{h_i}{2}\mathbf{e}_i$ . From (9) we have  $P_{ii} = 2\mu\partial_i\phi_i - \mu p - (2\lambda + d\mu)$ ,  
 560 thus the stencil for  $P_{ii}$  uses variables  $\phi_i(\mathbf{X}^*)$ ,  $p(\mathbf{X}^* + \frac{h_i}{2}\mathbf{e}_i)$   
 561 which are both *interior* (since  $\mathcal{L}_i$  is interior) and the one additional  
 562 variable  $\phi_i(\mathbf{X}^* + h_i\mathbf{e}_i)$ .  $P_{ii}$  would be an exterior stress only if  
 563  $\phi_i(\mathbf{X}^* + h_i\mathbf{e}_i)$  was an exterior variable; in this case  $P_{ii}$  would have  
 564 been “near” (specifically half a cell away from) a traction boundary face  
 565 normal to  $\mathbf{e}_i$ . Once again, we will use the traction condition associated  
 566 with this boundary to set  $P_{ii} = t_i$  (or  $P_{ii} = 0$  for a free boundary). The  
 567 subtlety of this formulation is that the stress variable  $P_{ii}$  is not located  
 568 exactly on the boundary; nevertheless the discrete stencil for  $P_{ii}$  is  
 569 still a valid *first-order* approximation of the continuous quantity  $P_{ii}$   
 570 at the boundary. At the worst case, our solver merely loses second-order  
 571 accuracy, which is a widely acceptable compromise for computer graphics  
 572 purposes, and creates hardly any visible artifacts in our experience.  
 573  
 574  
 575

576 In summary, we have justified that all *exterior* stress variables can be  
 577 eliminated (and replaced with known constants) from the divergence form  
 578 of interior position equations. A similar treatment is performed on the  
 579 discretization of the pressure equation  $\mathcal{L}_p = \mu \sum_i F_{ii} + \frac{\mu^2}{\lambda} p$ . Similar to  
 580 stresses, the deformation gradients  $F_{ii}$  are also characterized as interior  
 581 or exterior, based on whether they touch traction boundary variables. Since  
 582  $P_{ii} = 2\mu F_{ii} - \mu p - (2\lambda + d\mu)$ , we observe that  $F_{ii}$  is exterior if  
 583 and only if the stress  $P_{ii}$  is exterior (see figure 8, right). For such  
 584 exterior gradients or stresses we can use the traction condition  $P_{ii} = t_i$   
 585 to eliminate  $F_{ii}$  from the pressure equation. This is accomplished by  
 586 replacing  $\mathcal{L}_p \leftarrow \mathcal{L}_p - \frac{1}{2}(P_{ii} - t_i)$  for every exterior gradient  
 587  $F_{ii}$ . If more than one gradients are exterior, we can annihilate all

589 of them from the stencil for  $\mathcal{L}_p$  in a similar fashion. Note that in  
 590 the process of annihilating these exterior gradients, we modify the  
 591 original coefficient  $\mu^2/\lambda$  of the pressure variable at the center of  
 592  $\mathcal{L}_p$ , yet the modified coefficient will retain a positive sign regardless  
 593 of how many gradients are eliminated.

Our manipulations effectively remove all traction boundary variables from the discretization of the interior equations. For every Dirichlet boundary variable, we assume a Dirichlet condition of the form  $\phi_i = c_i$  is provided. Thus, we can substitute a given value for every Dirichlet variable in the stencil of every interior equation that uses it. As a result, our overall discrete system can be written as  $\mathcal{L}^* \mathbf{u}^* = \mathbf{b} - \mathbf{b}^D = \mathbf{b}^*$ , where  $\mathbf{u}^*$  only contains interior variables, and  $\mathbf{b}^D$  results from moving the known Dirichlet variables to the right-hand side. The discrete system matrix  $\mathcal{L}^*$  has as many rows and columns as interior variables, and will differ from  $\mathcal{L}$  near the boundaries, as it incorporates the effect of the boundary conditions. An analysis of our formulation can verify that  $\mathcal{L}^*$  has the form

$$\mathcal{L}^* = \begin{pmatrix} \mathbf{L}_\phi & \mathbf{G} \\ -\mathbf{G}^T & \mathbf{D}_p \end{pmatrix}$$

In this formulation  $\mathbf{L}_\phi$  is symmetric, positive definite, and  $\mathbf{D}_p$  is a diagonal matrix with positive diagonal elements. The skew-symmetry of the off-diagonal blocks is anticipated, since  $\mathbf{G}$  originates from a discretization of first-order derivatives. As a final step, we define the *substitution* matrix  $\mathcal{U}$

$$\mathcal{U} = \begin{pmatrix} \mathbf{I} & -\mathbf{G}\mathbf{D}_p^{-1} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$$

594 and use it to pre-multiply our equation as

$$\mathcal{U}\mathcal{L}^*\mathbf{u}^* = \begin{pmatrix} \mathbf{L}_\phi + \mathbf{G}\mathbf{D}_p^{-1}\mathbf{G}^T & \mathbf{0} \\ -\mathbf{G}^T & \mathbf{D}_p \end{pmatrix} \mathbf{u}^* = \mathcal{U}\mathbf{b}^* \quad (10)$$

595 Equation (10) is the basis of our boundary smoother. The top left  
 596 block  $\mathbf{L}_\phi + \mathbf{G}\mathbf{D}_p^{-1}\mathbf{G}^T$  is a symmetric and positive definite matrix and  
 597 can be smoothed via Gauss-Seidel iteration. One may recognize that  
 598 this matrix is qualitatively similar to the discretization of our non-  
 599 augmented system (4) and, in fact, the two are identical away from  
 600 boundaries. In section 4.1 we discussed the problematic conditioning of  
 601 this matrix in the near-incompressible regime. In our boundary smoother,  
 602 however, this formulation is only used for a very narrow region around  
 603 the boundary, specifically the same interior equations affected by our  
 604 previous box smoother formulation (depicted as red interior nodes in  
 605 Figure 6, right). The boundary and interior regions are smoothed in  
 606 separate sweeps; during the sweep of the boundary smoother, all  
 607 interior variables not being smoothed are effectively treated as Dirichlet  
 608 values. The fact that the boundary smoother is confined in a narrow  
 609 region between boundary conditions has a strong stabilizing effect,  
 610 as compared to its use for interior smoothing. In practice, we found  
 611 that 2 Gauss-Seidel boundary sweeps for every sweep of the distributive  
 612 interior smoother are sufficient for Poisson’s ratio up to  $\nu = .45$ ,  
 613 while 3-4 Gauss-Seidel sweeps will accommodate values as high as  
 614  $\nu = .495$ .

The last step in the boundary smoothing process is the treatment of the pressure and boundary variables, which were not included in the Gauss-Seidel update. Since the lower right block of equation (10) is diagonal, all pressure equations can be satisfied exactly via a simple Gauss-Seidel sweep (which is essentially equivalent to forward-substitution). Note that this update of pressures only needs to occur once *after* the position variables have been smoothed by a given number of sweeps. Lastly, the boundary traction variables can be updated using the traction conditions  $P_{ij} = t_i$  that have been imposed. For simplicity we will assume that no interior cell

625 is trapped between two exterior cells which are adjacent on two  
 626 opposing faces, i.e. the active domain does not contain a one-cell  
 627 wide isthmus between traction boundaries. This assumption is only  
 628 required of the finest level of discretization in the multigrid scheme  
 629 and can be procedurally accomplished by minimal thickening of the  
 630 active domain. In this case all external diagonal stress components  
 631  $P_{ii}$  have a stencil that touches only a single boundary variable, lo-  
 632 cated at the surface of the active domain. Thus, if the equation  
 633  $P_{ii} = t_i$  is stored at the location of this specific boundary variable,  
 634 a Gauss-Seidel step will compute the exact value of that boundary  
 635 variable. Similarly, all off-diagonal stress components  $P_{ij}$  ( $i \neq j$ )  
 636 will contain only a single boundary variable, located *outside* the  
 637 active domain, whose value has not been already determined, and  
 638 we store the boundary condition  $P_{ij} = t_i$  on the location of that  
 639 variable. In summary, after the smoothing of interior position vari-  
 640 ables, all our smoother has to do to is perform Gauss-Seidel relax-  
 641 ation on pressure equations, diagonal stress equations  $P_{ii} = t_i$  and  
 642 off-diagonal stress equations  $P_{ij} = t_i$ , in this specific order. At the  
 643 end of the process, all boundary equations will be satisfied *exactly*  
 644 (i.e. they will have zero residuals), which we exploit next.

## 645 6 Construction of the Transfer operators

646 We designed the Restriction ( $\mathcal{R}$ ) and Prolongation ( $\mathcal{P}$ ) operators  
 647 employed by the algorithm of Table 1 aiming to keep implementa-  
 648 tion as inexpensive as possible, while conforming to the textbook  
 649 accuracy requirements for full multigrid efficiency (see [Trotten-  
 650 berg et al. 2001]). We define the following 1D averaging operators

$$\begin{aligned} \mathcal{B}^1 u[x] &= \frac{1}{2} u[x - \frac{h}{2}] + \frac{1}{2} u[x + \frac{h}{2}] \\ \mathcal{B}^2 u[x] &= \frac{1}{4} u[x - h] + \frac{1}{2} u[x] + \frac{1}{4} u[x + h] \\ \mathcal{B}^3 u[x] &= \frac{1}{8} u[x - \frac{3h}{2}] + \frac{3}{8} u[x + \frac{h}{2}] + \frac{3}{8} u[x - \frac{h}{2}] + \frac{1}{8} u[x + \frac{3h}{2}] \end{aligned}$$

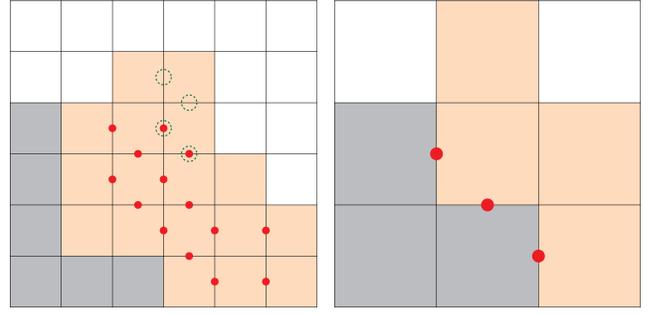
651 The restriction and prolongation operators will be defined as tensor  
 652 product stencils of the preceding 1D operators as

$$\begin{aligned} \mathcal{R}_1 &= \mathcal{B}^2 \otimes \mathcal{B}^1 \otimes \mathcal{B}^1 & \mathcal{P}_1^T &= 8 \mathcal{B}^2 \otimes \mathcal{B}^3 \otimes \mathcal{B}^3 \\ \mathcal{R}_2 &= \mathcal{B}^1 \otimes \mathcal{B}^2 \otimes \mathcal{B}^1 & \mathcal{P}_2^T &= 8 \mathcal{B}^3 \otimes \mathcal{B}^2 \otimes \mathcal{B}^3 \\ \mathcal{R}_3 &= \mathcal{B}^1 \otimes \mathcal{B}^1 \otimes \mathcal{B}^2 & \mathcal{P}_3^T &= 8 \mathcal{B}^3 \otimes \mathcal{B}^3 \otimes \mathcal{B}^2 \\ \mathcal{R}_p &= \mathcal{B}^1 \otimes \mathcal{B}^1 \otimes \mathcal{B}^1 & \mathcal{P}_p^T &= 8 \mathcal{B}^1 \otimes \mathcal{B}^1 \otimes \mathcal{B}^1 \end{aligned}$$

653 where  $\mathcal{R}_i, \mathcal{P}_i$  are the restriction and prolongation operators used for  
 654 variable  $u_i$ , respectively. We opted to define the prolongation oper-  
 655 ator in terms of its transpose, since  $\mathcal{P}_i^T$  has the same stencil every-  
 656 where, while  $\mathcal{P}_i$  is composed of several different stencils. We can  
 657 easily verify that  $\mathcal{P}$  simply corresponds to trilinear interpolation.

658 Our domain description for the finest grid was based on a parti-  
 659 tioning of the cells into interior, exterior and Dirichlet. The coarse  
 660 background grid is derived by the natural 8-to-1 coarsening of the  
 661 cartesian background lattice. Furthermore, a coarse cell is design-  
 662 ated a Dirichlet cell if *any* of its eight fine sub-cells is Dirichlet.  
 663 If any of the fine sub-cells are interior and none is Dirichlet, the  
 664 coarse cell will be considered interior. Otherwise, the coarse cell is  
 665 exterior. Thus, the coarse active domain is geometrically a super-  
 666 set of the fine domain, while its Dirichlet parts are more extensive.  
 667 Despite this geometrical discrepancy, which is in any case no larger  
 668 than the grid size, we were still able to retain the interior efficiency  
 669 of the multigrid scheme, by taking a few simple additional steps.

670 In our treatment of boundary equations in section 5.3 we effectively  
 671 forced all boundary conditions to be satisfied *exactly* after every ap-  
 672 plication of the smoother, by absorbing boundary equations into the  
 673 interior system. In general, if a smoother by design leaves a residual



674 **Figure 9:** Fine grid domain description (left) and its coarse grid  
 675 form (right). Red dots indicate (a) coarse Dirichlet equations that  
 676 were interior in the fine grid, and (b) the fine grid equations they  
 677 restrict residuals from. Green circles indicate fine interior variables  
 678 that prolongate their correction from boundary coarse variables.

674 on the boundary equations, this residual has to be restricted (sepa-  
 675 rately from interior residuals) onto the boundary equations of the  
 676 fine grid. In our case all boundary residuals in the fine grid are  
 677 zero, thus all coarse boundary equations will be *homogeneous*; for  
 678 Dirichlet equations they will have the form  $u_i^{2h} = 0$  (i.e. the coarse  
 679 grid incurs no correction), while traction equations will be of the  
 680 form  $\hat{P}_{ij}^{2h} = 0$ , where  $\hat{\mathbf{P}}$  is the homogeneous part of  $\mathbf{P}$  (i.e. we  
 681 omit any constant terms). We also note that, due to the possible  
 682 geometrical change of the Dirichlet region, certain coarse Dirichlet  
 683 equations will be centered on locations that were interior in the fine  
 684 grid (shown as red dots in Figure 9, right). The fine grid interior  
 685 equations (red dots in Figure 9, left) that would restrict their resid-  
 686 uals onto these (now Dirichlet) coarse locations, will not have their  
 687 residuals well represented on the coarse grid. We compensate for  
 688 this inaccuracy by performing an extra 2-3 sweeps of our boundary  
 689 Gauss-Seidel smoother over these equations, driving their residuals  
 690 very close to zero, just before the Restriction operation takes place.

691 In section 5.3 we stated that the interior domain must not have a  
 692 one-cell wide isthmus, where an interior cell is neighbored by two  
 693 exterior cells across two opposing faces, in order for the traction  
 694 boundary variables to be exactly computable using a sequence of  
 695 Gauss-Seidel updates. Although this property can be procedurally  
 696 enforced in the fine grid, the coarse grid could easily violate it,  
 697 as illustrated in Figure 9. Instead of attempting to compute these  
 698 coarse boundary variables in some more elaborate way, we simply  
 699 elect to not use them. This is possible, since the formulation of  
 700 section 5.3 defines a form of the equations that completely excludes  
 701 the traction boundary variables. The consequence of this approach  
 702 is that certain fine grid variables (depicted as green circles in Figure  
 703 9, left) cannot get an accurate correction, since it would have to  
 704 be prolonged from an omitted coarse boundary variable. Again,  
 705 we found that performing 2-3 Gauss-Seidel smoother sweeps over  
 706 these equations immediately after the prolongation of the correction  
 707 will be enough to compensate for this inaccurate correction.

## 708 7 Co-rotational linear elasticity

709 Our discussion so far was based on the constitutive equations of  
 710 linear elasticity. This model allowed us to detail the theoretical and  
 711 engineering subtleties associated with designing a highly efficient  
 712 multigrid scheme, and provides excellent conditions for the imple-  
 713 mentation of real-time multigrid solvers for models with hundreds  
 714 of thousands of degrees of freedom. The physical validity, however,  
 715 of linear elasticity is primarily limited to scenarios of moderate de-  
 716 formation. In the large deformation regime, and in the presence

of large rotational deformations in particular, the linear elasticity model develops artifacts such as volumetric distortions in parts of the domain with large rotations. Interestingly, our experiments with linear elasticity seemed to indicate that visually plausible deformations were attainable even for certain examples with higher deformation, where conventional wisdom would lead us to expect more visible artifacts; we traced this effect to our ability to use a high Poisson’s ratio which helped with volume conservation in the moderate deformation regime. Nevertheless, it is easy to demonstrate cases where linear elasticity produces visibly nonphysical results, and where a nonlinear treatment of deformation is required.

Our first extension is the co-rotational linear elasticity model, which has been used in slightly different forms by a number of authors in computer graphics [Müller et al. 2002; Hauth and Strasser 2004; Müller and Gross 2004], and has also been employed in the context of a finite element based multigrid framework by [Georgii and Westermann 2006]. The co-rotational formulation extracts the rotational component of the local deformation at a specific part of the domain by computing the polar decomposition of the deformation gradient tensor  $\mathbf{F} = \mathbf{R}\mathbf{S}$  into the rotation  $\mathbf{R}$  and the symmetric tensor  $\mathbf{S}$ . The stress is then computed as  $\mathbf{P} = \mathbf{R}\mathbf{P}_L(\mathbf{S})$ , where  $\mathbf{P}_L$  denotes the stress of a *linear* material, as described in equation (2). Thus, the co-rotational formulation computes stress by applying the constitutive equation of linear elasticity in a frame of reference that is rotated with the material deformation, and subsequently rotating the result back to unrotated coordinates. After the necessary substitutions, the constitutive equation takes the following form

$$\begin{aligned} \mathbf{P} &= 2\mu(\mathbf{F}-\mathbf{R}) + \lambda\text{tr}(\mathbf{R}^T\mathbf{F}-\mathbf{I})\mathbf{R} \\ &= 2\mu\mathbf{F} + \lambda\text{tr}(\mathbf{R}^T\mathbf{F})\mathbf{R} - (2\mu + d\lambda)\mathbf{R} \\ &= 2\mu\mathbf{F} - \mu p\mathbf{R} - (2\mu + d\lambda)\mathbf{R} \end{aligned} \quad (11)$$

where the last form of the stress in equation (11) results from introducing an auxiliary pressure variable  $p = -(\lambda/\mu)\text{tr}(\mathbf{R}^T\mathbf{F})$  in a fashion similar to the augmentation used for the linear elasticity problem in section 4.1. As before, the augmented position equations are defined as  $-\partial_j P_{ij} = f_i$ . After combining with the equations defining the pressures and rearranging some terms we get

$$\begin{pmatrix} -2\mu\Delta\mathbf{I} & \mu(\nabla^T\mathbf{R}^T)^T \\ \mu(\mathbf{R}\nabla)^T & \frac{\mu^2}{\lambda} \end{pmatrix} \begin{pmatrix} \phi \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f} + (2\mu + d\lambda)\nabla \cdot \mathbf{R} \\ 0 \end{pmatrix} \quad (12)$$

which we use as the augmented PDE of co-rotational linear elasticity. The notation for the off-diagonal blocks of the matrix in equation (12) were used to indicate whether the operators  $\nabla$ ,  $\nabla^T$  operate or not on the rotation matrix  $\mathbf{R}$ . In index form, these operators equal  $[\mu(\nabla^T\mathbf{R}^T)^T]_i = \mu\partial_j R_{ij}$ , and  $[\mu(\mathbf{R}\nabla)^T]_i = \mu R_{ij}\partial_j$  respectively. *Discuss quasi-linear form here*

## 8 Neo-hookean elasticity

TBD

## 9 Dynamics

TBD

## 10 Collisions

TBD

## 11 Parallelization and Optimization

TBD

## 12 Results

TBD

## 13 Discussion and Limitations

TBD

## 14 Conclusion

TBD

## References

- BONET, J., AND WOOD, R. 1997. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge University Press.
- BREZZI, F., AND FORTIN, M. 1991. *Mixed and hybrid finite element methods*. Springer: Berlin.
- ENGLISH, E., AND BRIDSON, R. 2007. Animating developable surfaces using nonconforming elements. *ACM Trans. Graph. (SIGGRAPH Proc.)* 26, 3.
- FEDKIW, R., STAM, J., AND JENSEN, H. 2001. Visual simulation of smoke. In *Proc. of ACM SIGGRAPH 2001*, 15–22.
- GASPAR, F., GRACIA, J., LISBONA, F., AND OOSTERLEE, C. 2008. Distributive smoothers in multigrid for problems with dominating grad-div operators. *Numerical Linear Algebra with Applications* 15, 8, 661–683.
- GEORGII, J., AND WESTERMANN, R. 2006. A multigrid framework for real-time simulation of deformable bodies. *Computers and Graphics* 30, 3, 408 – 415.
- HAUTH, M., AND STRASSER, W. 2004. Corotational Simulation of Deformable Solids. In *In Proc. of WSCG*, 137–145.
- IRVING, G., SCHROEDER, C., AND FEDKIW, R. 2007. Volume conserving finite element simulations of deformable models. *ACM Trans. Graph. (SIGGRAPH Proc.)* 26, 3.
- JAMES, D., BARBIC, J., AND TWIGG, C. 2004. Squashing cubes: Automating deformable model construction for graphics. In *SIGGRAPH 2004 Sketches & Applications*, ACM Press.
- KAUFMANN, P., MARTIN, S., BOTSCH, M., AND GROSS, M. 2008. Flexible Simulation of Deformable Models Using Discontinuous Galerkin FEM. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*.
- KAZHDAN, M., AND HOPPE, H. 2008. Streaming Multigrid for Gradient-Domain Operations on Large Images. *ACM Transactions on Graphics* 27, 3.
- MÜLLER, M., AND GROSS, M. 2004. Interactive virtual materials. In *Graph. Interface*, 239–246.
- MÜLLER, M., DORSEY, J., MCMILLAN, L., JAGNOW, R., AND CUTLER, B. 2002. Stable real-time deformations. In *ACM SIGGRAPH Symp. on Comput. Anim.*, 49–54.
- MÜLLER, M., TESCHNER, M., AND GROSS, M. 2004. Physically-based simulation of objects represented by surface meshes. In *Proc. Comput. Graph. Int.*, 156–165.
- RIVERS, A., AND JAMES, D. 2007. FastLSM: fast lattice shape matching for robust real-time deformation. *ACM Transactions on Graphics (TOG)* 26, 3.

- 813 TERZOPOULOS, D., AND FLEISCHER, K. 1988. Deformable mod-  
814 els. *The Visual Computer* 4, 6, 306–331.
- 815 TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K.  
816 1987. Elastically deformable models. *Computer Graphics (Proc.*  
817 *SIGGRAPH 87)* 21, 4, 205–214.
- 818 TROTTEMBERG, U., OOSTERLEE, C., AND SCHULLER, A. 2001.  
819 *Multigrid*. San Diego: Academic Press.