

Efficient Simulation of Inextensible Cloth

Rony Goldenthal^{1,2}

David Harmon¹

Raanan Fattal³

Michel Bercovier²

Eitan Grinspun¹

¹Columbia University

²The Hebrew University of Jerusalem

³University of California, Berkeley

Abstract

Many textiles do not noticeably stretch under their own weight. Unfortunately, for better performance many cloth solvers disregard this fact. We propose a method to obtain very low strain along the warp and weft direction using Constrained Lagrangian Mechanics and a novel fast projection method. The resulting algorithm acts as a velocity filter that easily integrates into existing simulation code.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: Physically-based Modeling, Cloth simulation, Constrained Lagrangian Mechanics, Constraints, Stretching, Inextensibility, Isometry

1 Introduction

Our eyes are very sensitive to the behavior of fabrics, to the extent that we can identify the kind of fabric simply from its shape and motion [Griffiths and Kulke 2002]. One important fact is that most fabrics do not stretch under their own weight. Unfortunately, for many popular cloth solvers, a reduction of permissible stretching is synonymous with degradation in performance: for tractable simulation times one may settle for an unrealistic 10% or more strain (compare 1% and 10%, adjacent figure). Our work alleviates this problem by introducing a numerical solver that excels at time-stepping quasi-inextensible surfaces (stretching below 1%).

The solver builds on a framework of Constrained Lagrangian Mechanics (CLM) [Marsden 1999]. Warp and weft, the perpendicular sets of strands that make up a textile, are prohibited from stretching by enforcing constraint equations, not by integrating spring forces. We present numerical evidence supporting the observation that a constraint-based method is inherently well-suited to operate in the quasi-inextensible regime. In contrast, for this regime spring-based methods are known to experience a range of difficulties, leading to the adoption of various strain limiting [Provot 1995] and strain rate limiting algorithms.

We are motivated by the work of Bridson *et al.* [2002], who viewed strain limiting as one of multiple *velocity filtering* passes (another being collision handling). The velocity filter paradigm enables the design of modular systems with mix-and-match flexibility.



Figure 1: Importance of capturing inextensibility. *For efficiency, many simulation methods allow 10% or more strain, whereas many fabrics do not visibly stretch. A 1m² patch, pinned at two corners 1m apart, is allowed to relax under gravity. We compare (left to right) three simulations of progressively smaller permissible strain with an actual denim patch.*

Contributions We propose a novel CLM formulation that is implicit on the constraint gradient (§4.1). We prove that the implicit method’s nonlinear equations correspond to a minimization problem (§4.2): this result motivates a *fast projection method* for enforcing inextensibility (§4.3). We describe an implementation of fast projection as a simple and efficient velocity filter, as part of a framework that decouples timestepping, inextensibility, and collision passes (§4.4). Consequently, the fast projection method easily incorporates with a code’s existing bending, damping, and collision models, to yield accelerated performance (§5).

Before discussing these contributions, we summarize the relevant literature (§2) and describe the basic discrete cloth model (§3).

2 Related Work

For brevity, we review work on stretch resistance; for broad surveys on cloth simulation see [House and Breen 2000; Choi and Ko 2005].

The most general approach is to treat cloth as an elastic material [Terzopoulos *et al.* 1987; Breen *et al.* 1994; Eberhardt *et al.* 1996; Baraff and Witkin 1998; Choi and Ko 2002]. To reduce visible stretching, elastic models typically adopt large elastic moduli or stiff springs, degrading numerical stability [Hauth *et al.* 2003].

To address the stiffness of the resulting differential equations, Baraff and Witkin [1998] proposed implicit integration, allowing for large, stable timesteps; adaptive timestepping was required to prevent over-stretching. Eberhardt [2000] and Boxerman *et al.* [2003] adopted implicit-explicit (IMEX) formulations, which treat only a subset of forces implicitly. Our method is closely related to the IMEX approach, in the sense that stretching forces are singled out for special treatment.

These works, and many of their sequels, improved performance by allowing some perceptible stretch of the fabric. In the quasi-inextensible regime, however, implicit methods encounter numerical limitations [Volino and Magnenat-Thalmann ; Boxerman 2003; Hauth *et al.* 2003]: the condition number of the implicit system grows with the elastic material stiffness, forcing iterative solvers to perform many iterations; additionally, timestepping algorithms such as Backward Euler and BDF2 introduce undesirable numerical damping when the system is stiff [Boxerman 2003].

Given a stiff differential equation, an alternative to implicit integration is to reduce the stiff component and reformulate it as a constraint [Hairer et al. 2002]. In the smooth setting, the penalty-force and constraint-based approaches are equivalent in the limit of an infinitely stiff penalty term [Bercovier and Pat 1984]. In the discrete setting, the constraint-based approach may be implemented with various iterative or global algorithms, as surveyed below:

Iterative enforcement. Provot [1995] corrected edge lengths by iteratively displacing the incident vertices on stretched springs. While simple to implement, this approach suffers from poor convergence since each displacement may stretch other incident springs. Therefore, Provot’s method is used in cases where tight tolerances are not required, *e.g.*, [Desbrun et al. 1999; Meyer et al. 2001; Fuhrmann et al. 2003]. Bridson *et al.* [2002; 2003] used Provot’s approach in conjunction with strain rate limiting (SRL), bounding the rate of change of spring length per timestep to 10% of the current length. Müller *et al.* [2006] used a non-linear Gauss-Seidel approach to enforce inextensibility on each constraint separately. Bridson *et al.* observe that iterative strain limiting algorithms behave essentially as Jacobi or Gauss-Seidel solvers. In this light, it is not surprising that for finely-discretized quasi-inextensible fabrics, iterative constraint enforcement requires a prohibitive number of iterations (see §5).

Global enforcement. In contrast to iterative constraint enforcement, House *et al.* [1996] used Lagrange multipliers with CLM to treat stretching, and presented a hierarchical treatment of the constraint forces. The Lagrange multiplier approach alleviates the difficulties associated with poor numerical conditioning and artificial damping. House *et al.* later encountered difficulties in handling collision response within the proposed framework [2000]. By building on the velocity-filter paradigm, our method handles both inextensibility and complex collisions.

House *et al.* formulated constraints as in [Witkin et al. 1990], which is subject to numerical drift that may be exacerbated by the discontinuities introduced during collision response. Drift may be attenuated using constraint-restoring springs, but the authors reported difficulty in adjusting the spring coefficients. We postulate that one reason for their difficulties with drift was consequent to the linearization of the constraint equation, which permitted higher order errors to accumulate over time. Our method avoids a linearization of the constraint equations, and is not subject to drift.

Recently, Tsiknis [2006] proposed triangle-based strain limiting together with a global stitching step for stable constraint enforcement. Hong *et al.* [2005] used a linearized implicit formulation in order to improve stability of constrained dynamics. This allowed for larger timesteps and reduced the need for springs to maintain the cloth on the constraint manifold. Both of these approaches enforce inextensibility only for strain exceeding 10%.

In summary, when the tolerance for stretching is very small, modeling stretch response with spring-based or strain-limiting approaches is costly and even intractable; constraint-based methods present a promising alternative. The remainder of this paper discusses algorithms that excel at simulating quasi-inextensible cloth.

3 Cloth Model

Woven fabrics are not a continuous material, rather they are a complex mechanical network of interleaving yarn [Breen et al. 1994]. Since the constituent yarn is often quasi-inextensible, the material’s warp and weft directions do not stretch perceptibly. In contrast, shearing modes excite only a mechanical interaction of warp and weft and not a stretching of yarn, and therefore fabric often does

shear perceptibly. Therefore, we model shear using non-stiff stretch springs applied on both diagonals of each quad.

In imposing inextensibility on all edges of a triangle mesh, one quickly runs into parasitic stiffness in the bending modes, or *locking* [Zienkiewicz and Taylor 1989], since locally-convex regions of a triangle mesh are rigid under isometry. Instead, we consider warp-weft aligned quadrilateral meshes with a sparse number of triangles (*quad-dominant* meshes). A degree of freedom (DOF) counting argument suggests that constraining all edges of a quad mesh may circumvent the rigidification that occurs with triangle meshes: Given n vertices, we have $3n$ positional DOFs; their triangulation (resp. quadrangulation) introduces approximately $3n$ (resp. $2n$) edges, with corresponding inextensibility constraints. Subtracting constraints from positional DOFs leaves nearly zero DOFs for a triangulation. In the case of a quadrangulation, $O(n)$ DOFs remain, and we see that in a flat configuration they correspond to the normal direction at each vertex. Furthermore, under general mesh positions, the constraints are linearly independent, with a full-rank Jacobian treatable by a direct solver (§4).

We ask that a warp- or weft-aligned quad edge, $(\mathbf{p}_a, \mathbf{p}_b)$, maintain its undeformed length, l , by enforcing

$$C(\mathbf{p}_a, \mathbf{p}_b) = \|\mathbf{p}_b - \mathbf{p}_a\|^2/l - l = 0. \quad (1)$$

The solve will require the *constraint gradient*

$$\nabla_{\mathbf{p}_b} C(\mathbf{p}_a, \mathbf{p}_b) = 2(\mathbf{p}_b - \mathbf{p}_a)/l. \quad (2)$$

The complete model of in-plane deformation is compatible with an existing code’s quad- or triangle-based treatment of bending and collisions. With this simple formulation of inextensibility constraints in place, what is needed is an efficient method for enforcing constraints. In the following, we develop such a method.

4 Constrained Dynamics

Given a quadrilateral mesh with n vertices and m edges, the numerical integration algorithm for constrained dynamics can be developed directly from the augmented Lagrange equation [Marsden 1999],

$$L(\mathbf{x}, \mathbf{v}) = \frac{1}{2} \mathbf{v}^T \mathbf{M} \mathbf{v} - V(\mathbf{x}) - \mathbf{C}(\mathbf{x})^T \boldsymbol{\lambda},$$

where $\mathbf{x}(t)$ is the time-varying $3n$ -vector of vertex positions, $\mathbf{v}(t) = \dot{\mathbf{x}}(t)$ is its time derivative, \mathbf{M} is the $3n \times 3n$ mass matrix, and $V(\mathbf{x})$ is the stored energy (*e.g.*, bending, shear, and gravity), $\boldsymbol{\lambda}$ is the m -vector of Lagrange multipliers $\mathbf{C}(\mathbf{x})$ is the m -vector of constraints, with the i^{th} entry corresponding to the violation of inextensibility of the i^{th} edge, as computed by (1). The corresponding Euler-Lagrange equations are

$$\mathbf{M} \dot{\mathbf{v}} = -\nabla V(\mathbf{x}) - \nabla \mathbf{C}(\mathbf{x})^T \boldsymbol{\lambda}, \quad (3)$$

$$\mathbf{C}(\mathbf{x}) = \mathbf{0}, \quad (4)$$

where $\nabla \equiv \nabla_{\mathbf{x}}$ is the gradient with respect to position, and $-\nabla V(\mathbf{x})$ is the potential force; $-\nabla \mathbf{C}(\mathbf{x})^T \boldsymbol{\lambda}$ may be viewed as the constraint-maintaining force. In the latter, the factors $-\nabla \mathbf{C}(\mathbf{x})^T$ and $\boldsymbol{\lambda}$ determine the direction and scaling for the force, respectively. $\nabla \mathbf{C}(\mathbf{x})$ is a rectangular matrix whose dimensions are $m \times 3n$.

For simulation, we must discretize (3) and (4) in time using one of various schemes, each with benefits and drawbacks. One may choose differing explicit or implicit schemes for the potential and the constraint forces (similarly, potential forces are split and separately discretized in [Ascher et al. 1997]). The discrete equations replace $\mathbf{x}(t)$ and $\mathbf{v}(t)$ with $\{\mathbf{x}^0, \mathbf{x}^1, \mathbf{x}^2, \dots\}$ and $\{\mathbf{v}^0, \mathbf{v}^1, \mathbf{v}^2, \dots\}$, where

\mathbf{x}^n and \mathbf{v}^n are the position and velocity of the mesh at time $t = nh$, and h is the size of the timestep.

One widely-used family of discretizations includes SHAKE and RATTLE, which extend the (unconstrained) Verlet scheme [Hairer et al. 2002] by considering a constraint force direction, $-\nabla\mathbf{C}(\mathbf{x})^T$, evaluated at the *beginning* of the timestep.

Unfortunately, enforcing length-preserving constraints with SHAKE fails for four common geometric configurations, which we refer to as (Q1)–(Q4) and depict in Figure 2. This figure is a reproduction from [Barth et al. 1994], which discusses these drawbacks in SHAKE but does not offer a solution. In the figure, solid and hollow dots represent edge endpoints at the start and end of the timestep, as the particles would evolve if no constraints were applied. If $-\nabla\mathbf{C}(\mathbf{x})^T$ is evaluated at the beginning of timestep positions \mathbf{x}^n , as in SHAKE, then no choice of $\boldsymbol{\lambda}^{n+1}$ satisfies $\mathbf{C}(\mathbf{x}^{n+1}) = \mathbf{0}$; numerically, for (Q2)–(Q4) this observation manifests as a singular Jacobian in Newton’s method. These four cases correspond to rapid change in edge length or orientation; in practice, they occur often.




Figure 2: Failure modes of methods using an explicit constraint direction. Reproduced from a discussion of SHAKE in [Barth et al. 1994].

4.1 Implicit constraint direction (ICD)

Consider evaluating the constraint direction, $-\nabla\mathbf{C}(\mathbf{x})^T$, at the *end* of the timestep. We observe (and prove in Appendix A) that this prevents (Q1), (Q2) and (Q4) from generating a singular matrix; (Q3) remains, but is automatically remedied by decreasing the timestep. Consider the ICD timestep, which treats potential forces explicitly¹:

$$\begin{aligned}\mathbf{v}^{n+1} &= \mathbf{v}^n - h\mathbf{M}^{-1} \left(\nabla V(\mathbf{x}^n) + \nabla\mathbf{C}(\mathbf{x}^{n+1})^T \boldsymbol{\lambda}^{n+1} \right), \\ \mathbf{x}^{n+1} &= \mathbf{x}^n + h\mathbf{v}^{n+1}, \\ \mathbf{C}(\mathbf{x}^{n+1}) &= \mathbf{0}.\end{aligned}$$

Define $\mathbf{x}_0^{n+1} = \mathbf{x}^n + h\mathbf{v}^n - h^2\mathbf{M}^{-1}\nabla V(\mathbf{x}^n)$, *i.e.*, \mathbf{x}_0^{n+1} is the position at the end of an unconstrained timestep. Next, eliminate \mathbf{v}^{n+1} by rewriting the above system as two equations, $\mathbf{F}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}^{n+1}) = \mathbf{0}$ and $\mathbf{C}(\mathbf{x}^{n+1}) = \mathbf{0}$, where

$$\mathbf{F}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}^{n+1}) = \mathbf{x}^{n+1} - \left(\mathbf{x}_0^{n+1} - h^2\mathbf{M}^{-1}\nabla\mathbf{C}(\mathbf{x}^{n+1})^T \boldsymbol{\lambda}^{n+1} \right).$$

$\mathbf{F}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}^{n+1})$ and $\mathbf{C}(\mathbf{x}^{n+1})$ are the residuals of the discretization of (3) and (4), respectively. In particular, \mathbf{F} measures the deviation of the trajectory away from that dictated by the governing (potential and constraint) forces; \mathbf{C} measures the deviation from the constraint manifold (in our case, the extensibility of the material). To implement ICD, we solve for the roots of \mathbf{F} and \mathbf{C} up to a desired tolerance using Newton’s method. Solving for an ICD step is

¹For an implicit treatment, write $\nabla V(\mathbf{x}^{n+1})$ in place of $\nabla V(\mathbf{x}^n)$.

Algorithm 1 Fast projection is a velocity filter that enforces constraints. It combines the robustness of using an implicit constraint direction with the efficiency of approximate manifold projection.

Input: $\tilde{\mathbf{v}}^{n+1}$ // candidate end-of-step velocity
 1: $\mathbf{x}_0^{n+1} \leftarrow \mathbf{x}^n + \tilde{\mathbf{v}}^{n+1}h$ // unconstrained timestep
 2: **while** strain exceeds threshold **do**
 3: $\delta\boldsymbol{\lambda} \leftarrow - \left[\nabla\mathbf{C}(\mathbf{x}_j^{n+1})\nabla\mathbf{C}(\mathbf{x}_j^{n+1})^T \right]^{-1} \mathbf{C}(\mathbf{x}_j^{n+1})$
 4: $\delta\mathbf{x} \leftarrow \nabla\mathbf{C}(\mathbf{x}_j^{n+1})^T \delta\boldsymbol{\lambda}$
 5: $\mathbf{x}_{j+1}^{n+1} \leftarrow \mathbf{x}_j^{n+1} + \delta\mathbf{x}$
 6: $j \leftarrow j + 1$
 7: **end while**
 8: $\mathbf{v}^{n+1} \leftarrow (\mathbf{x}^{n+1} - \mathbf{x}^n)/h$
Output: \mathbf{v}^{n+1} // constraint-enforcing velocity

costly, because there are many unknowns ($\approx 5n$), and each Newton step requires the solution of an *indefinite* linear system, whose matrix is costly to assemble. In §4.3, we develop an approximation to ICD that addresses these drawbacks without sacrificing constraint accuracy or robustness. To arrive at this fast projection method, the following section considers ICD from an alternative, geometric viewpoint.

4.2 Step and project (SAP)

Consider for a moment an alternative approach to constrained integration in two steps: (a) step forward only the potential forces to arrive at the unconstrained position, \mathbf{x}_0^{n+1} ; (b) enforce the constraints by projecting onto the constraint manifold $\mathcal{M} = \{\mathbf{x}^{n+1} | \mathbf{C}(\mathbf{x}^{n+1}) = \mathbf{0}\}$. Methods of this form are known as manifold-projection methods [Hairer et al. 2002]. To define a specific method, we must choose a projection operator. In the method we refer to as SAP, the projection of \mathbf{x}_0^{n+1} onto \mathcal{M} is the stationary point with respect to $[\mathbf{x}^{n+1}, \boldsymbol{\lambda}]^T$ of the objective function

$$W(\mathbf{x}^{n+1}, \boldsymbol{\lambda}) = \frac{1}{2} \|\mathbf{x}^{n+1} - \mathbf{x}_0^{n+1}\|^2 - \mathbf{C}(\mathbf{x}^{n+1})^T \boldsymbol{\lambda}, \quad (5)$$

where $\|\cdot\|$ is the usual Euclidean norm. Simply put, we choose the point on the constraint manifold closest to \mathbf{x}_0^{n+1} .

Theorem 1: ICD \equiv SAP

Proof: The stationary point of $W(\mathbf{x}^{n+1}, \boldsymbol{\lambda})$ satisfies

$$\begin{bmatrix} \nabla_{\mathbf{x}^{n+1}} W \\ \nabla_{\boldsymbol{\lambda}} W \end{bmatrix} = \begin{bmatrix} \mathbf{x}^{n+1} - \mathbf{x}_0^{n+1} - \nabla_{\mathbf{x}^{n+1}} \mathbf{C}(\mathbf{x}^{n+1})^T \boldsymbol{\lambda} \\ \mathbf{C}(\mathbf{x}^{n+1}) \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (6)$$

Substituting $\boldsymbol{\lambda} = h^2\mathbf{M}^{-1}\boldsymbol{\lambda}^{n+1}$ recovers ICD: $\mathbf{F}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}^{n+1}) = \mathbf{0}$ and $\mathbf{C}(\mathbf{x}^{n+1}) = \mathbf{0}$. \square

Corollary In 4.1, we interpreted the roots of \mathbf{C} and \mathbf{F} from the ICD view. We can interpret these roots from the SAP view as follows: $\mathbf{C}(\mathbf{x}^{n+1}) = \mathbf{0}$ corresponds to finding *some* point on the constraint manifold. $\mathbf{C}(\mathbf{x}^{n+1}) = \mathbf{0}$ with $\mathbf{F}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}^{n+1}) = \mathbf{0}$ corresponds to finding the *closest* point on the constraint manifold.

4.3 Fast projection method

Starting with the unconstrained position, \mathbf{x}_0^{n+1} , we propose to find a *close*, but not necessarily closest, point on the constraint manifold. To find this nearby point, fast projection starts at \mathbf{x}_0^{n+1} , and

takes a sequence of steps, \mathbf{x}_j^{n+1} , $j = 1, 2, \dots$, toward the constraint manifold, with each step minimizing the distance of \mathbf{x}_{j+1}^{n+1} to its predecessor, \mathbf{x}_j^{n+1} . In comparison, to would solve SAP we minimize (5) using Newton’s method: each iteration would improve upon a guess for *the* shortest step from \mathbf{x}_0^{n+1} onto the constraint manifold.

A step of fast projection Formally, each step of fast projection, $\mathbf{x}_{j+1}^{n+1} = \mathbf{x}_j^{n+1} + \delta\mathbf{x}$, minimizes the objective function

$$\tilde{W}(\delta\mathbf{x}, \delta\boldsymbol{\lambda}) = \frac{1}{2} \|\delta\mathbf{x}\|^2 - \mathbf{C}(\mathbf{x}_j^{n+1} + \delta\mathbf{x})^T \delta\boldsymbol{\lambda}, \quad (7)$$

with respect to the step increment, $\delta\mathbf{x}$, and an auxiliary variable, $\delta\boldsymbol{\lambda}$. Using $\mathbf{C}(\mathbf{x}_j^{n+1} + \delta\mathbf{x}) \approx \mathbf{C}(\mathbf{x}_j^{n+1}) + \nabla\mathbf{C}(\mathbf{x}_j^{n+1})\delta\mathbf{x}$, we obtain a quadratic objective function whose stationary equations with respect to $\delta\mathbf{x}$ and $\delta\boldsymbol{\lambda}$ are

$$\begin{aligned} \delta\mathbf{x} &= \nabla\mathbf{C}(\mathbf{x}_j^{n+1})^T \delta\boldsymbol{\lambda}, \\ \nabla\mathbf{C}(\mathbf{x}_j^{n+1})\delta\mathbf{x} &= -\mathbf{C}(\mathbf{x}_j^{n+1}). \end{aligned}$$

We eliminate $\delta\mathbf{x}$, and solve a linear system in $\delta\boldsymbol{\lambda}$:

$$\left(\nabla\mathbf{C}(\mathbf{x}_j^{n+1})\nabla\mathbf{C}(\mathbf{x}_j^{n+1})^T\right) \delta\boldsymbol{\lambda} = -\mathbf{C}(\mathbf{x}_j^{n+1}). \quad (8)$$

Finally, we compute the increment $\delta\mathbf{x} = \nabla\mathbf{C}(\mathbf{x}_j^{n+1})^T \delta\boldsymbol{\lambda}$.

As with ICD/SAP, a fast projection step requires a linear solve. However, fast projection’s system, (8), is smaller ($\approx 2n \times 2n$ compared to $\approx 5n \times 5n$), positive definite (compared to indefinite) and sparser. As a result it is considerably cheaper to evaluate, assemble, and solve than its ICD/SAP counterpart.

Fast projection algorithm We repeatedly take fast projection steps until the maximal strain is below a threshold, *i.e.*, the constraint may be satisfied up to a given tolerance. This process is summarized in Algorithm 1.

Fast projection finds a manifold point, \mathbf{x}^{n+1} , that is close, but not closest, to the unconstrained point, \mathbf{x}_0^{n+1} . Referring to the Corollary, we conclude that fast projection exactly solves $\mathbf{C} = 0$ while it approximates $\mathbf{F} = 0$.

One important question is whether the fast projection’s error in \mathbf{F} is acceptable. Compare a sequence of fast projection iterations to ICD/SAP’s sequence of Newton iterations. The first iteration of these methods is identical. At the end of this first iteration, $\mathbf{F}, \mathbf{C} \in O(h^2)$. Additional fast projection iterations seek $\mathbf{C} \rightarrow 0$, and since $\mathbf{C} \in O(h^2)$, increments in \mathbf{x} are $O(h^2)$, therefore \mathbf{F} remains in $O(h^2)$. Observe that $\mathbf{F} \in O(h^2)$ is considered acceptable in many contexts, *e.g.*, [Baraff and Witkin 1998; Choi and Ko 2002] halt the Newton process after a single iteration. As recorded in the adjacent figure of a typical run, the first iteration of the fast projection method eliminates first-order error. The remaining iterations perturb \mathbf{F} only to higher-order (often decreasing the error further).

4.4 Implementation

We implemented fast projection as a velocity filter, enabling easy integration into our existing cloth simulation system; refer to Algorithm 1. Step 3 requires solving a sparse symmetric positive definite





Figure 3: Performance of fast projection vs. implicit springs. For a 1D chain simulated in MATLAB, we plot the computation time of one simulated second, as a function (a) of permissible strain (log-log plot for 80 vertices), and (b) of discretization resolution (linear plot for 1% permissible strain).

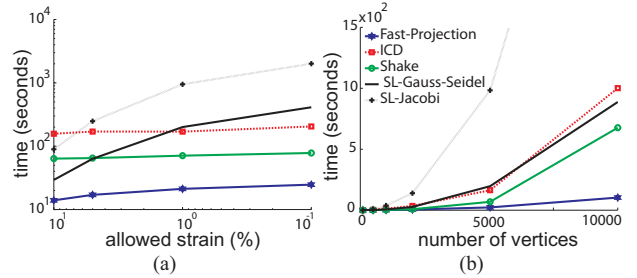


Figure 4: Performance of several constraint-enforcing methods. For a 2D cloth, simulated in C++, we plot the computation time of one simulated second, as a function (a) of permissible strain (log-log plot for 5041 vertices), and (b) of discretization resolution (linear plot for 1% permissible strain).

linear system; we use the PARDISO [Schenk and Gärtner 2006] solver. Each row of $\nabla\mathbf{C}(\mathbf{x}_j^{n+1})$ corresponds to one edge, and is computed using (2). The right-hand side, $\mathbf{C}(\mathbf{x}_j^{n+1})$, is given by (1). Note that Algorithm 1 is derived from the SAP point of view, therefore the terms h^2 and M^{-1} are omitted, incorporating them in, as in ICD, should merely scale the Lagrange multipliers. We tested both approaches, using a direct linear solver, both methods exhibited similar performance.

5 Results

One-dimensional chain Our first experiment compared the performance of fast projection against an implicit treatment of stiff springs. We observed the scaling of computational cost as a function of (a) permissible strain and (b) mesh resolution.

The physical setup consisted of a chain pinned at the top node and released to free fall under gravity. The simple 1D chain resists stretching, but not bending.

In this didactic example, timings refer to MATLAB’s (sparse) direct solver. Our method shows asymptotically better performance as permissible strain vanishes (see Figure 3a). Likewise, our algorithm enjoys favorable performance as mesh resolution increases (see Figure 3b). Using 80 vertices and 1% strain, the fast projection method achieved a 25× speedup. Note that there exists considerable difficulty in setting spring coefficients *a priori* to satisfy a given strain limit. For settings more pragmatic than a simple chain, such as the following draping experiment, we were unable (despite considerable effort) to set spring coefficients that achieve a pre-




Figure 5: Qualitative visual comparison. Snapshot of a cloth draped using (a) fast projection and (b) implicit constraint direction.

scribed small strain. This explains why spring methods are often treated with strain-limiting procedures.

Draping cloth The next experiment compared fast projection, ICD, SHAKE, and the strain limiting approach. Our goal was to evaluate how the spatial discretization and permissible strain affect performance of these four algorithms. The setup consisted of draping a cloth over a polygonal model of a sphere. We measured strain before the collision reaction pass.

For the strain limiting algorithms, we iterate until strain is in the permissible range. Both Jacobi and Gauss-Siedel approaches were tested; with Gauss-Siedel, a random permutation was applied in order to reduce bias resulting from the particular edge ordering. The SHAKE method used the acceleration suggested in [Barth et al. 1994] to rebuild the matrix once per step or when it fails to converge. As a consequence, the algorithm required extremely small timesteps to converge, but each timestep was computationally inexpensive, as matrix re-assembly and re-factoring was infrequent. ICD was able to use larger timesteps than the SHAKE method and still converge, however, since each timestep is substantially more expensive than a SHAKE step, the overall time is higher. Figure 4a shows timing comparison of these methods, and Figure 4b compares performance as the stiffness is increased for a ≈ 5000 vertex cloth. All CLM methods show similar good scaling performance, asymptotically better than the strain limiting approach, with the fast projection being the fastest. As the resolution is refined and strain of 1% is allowed (Figure 4b), the fast projection method outperforms the other methods.

Figure 5 shows the same frame from simulations performed using the fast projection and ICD method, with qualitatively comparable results. Figures 6 and 7 show still frames from more complex simulations demonstrating that fast projection is capable of producing complex, realistic simulations of cloth. Simulations were run on one CPU of a 2.66GHz Intel Core 2 Duo.

6 Discussion

Our experiments have focused on measuring the performance of enforcing inextensibility using CLM compared to strain limiting and stiff springs. In addition to the direct benefit of fast projection on computation times, further benefits can be reaped from the resulting inextensibility. For example, the work of Bergou *et al.* [2006] assumes inextensibility in order to accelerate bending computation. In adopting the velocity-filtering viewpoint, we have gained speed, simplicity, and software modularity—all key to a practical and maintainable implementation. This has, however, come at a



Figure 6: Inextensibility and dynamics. Inextensibility ensures that the tight-fitting pants do not drop past the dancer’s narrow waist. Using fast projection, an implicit treatment of shear and bending, and a mesh with 10600 vertices, the average simulation time per (30Hz) frame was 9 seconds.

theoretical cost: there is no longer an efficient way to perfectly enforce both ideal inextensibility and ideal collision handling, since one must execute one filter before the other, and both ideals correspond to sharp constraints. To enforce both perfectly would require combining them in a single pass, an elegant and exciting prospect from the standpoint of theory, but one which is likely to introduce considerable complexity and convergence challenges. Practically, we observe that this drawback does not cause artifacts in our simulation, for several reasons: first, we execute collision-handling last, to avoid glaring collision artifacts, yet we assert that empirically our strain remains negligible, as required. Second, unlike constraint-enforcement approaches such as [Witkin et al. 1990], the inextensibility filter does not assume that the constraint is maintained at the beginning of the timestep and errors are not accumulated during the simulation.

Conclusion Despite the fact that the most common fabrics do not visibly stretch when draped over the body, the trend in our community is to favor stretching formulations based on penalty-springs. The consequent numerical difficulties are then addressed by a combination of (a) considerably relaxing realism, allowing strain $\approx 10\%$, and (b) simple iterative strain and strain-rate algorithms with poor convergence behavior. With Constrained Lagrangian Mechanics as our alternative point of departure, we have demonstrated a straightforward filter, with good convergence behavior, for enforcing inextensibility. We provide one immediate and pragmatic approach to fast and realistic fabric simulation using CLM, and we hope that it will spur a renaissance of activity along this direction.

Acknowledgments We thank OptiTex for providing the 3D garment geometry as well as the animated figurines. We are grateful to David Ismailov for setting up the cloth models, and Ruzz Oved and Yaniv Goral for lighting and shading our scenes. This work has benefited from the valuable insights of Jerrold E. Marsden, Ari Stern, and Max Wardetzky, and from the generous support of the NSF (MSPA 0528402, CSR 0614770, CAREER 0643268), Autodesk, mental images, NVIDIA, and Elsevier.



Figure 7: Enforcing inextensibility using fast projection yields lively motion with detailed wrinkles and folds. *Frames from ballet and runway sequences simulated using fast projection. The elastic term was integrated implicitly (top) and explicitly (bottom), respectively. The cloth contains 8325 (top) and 10688 (bottom) vertices, with average simulation time per (30Hz) frame of 5.2 and 7.8 seconds, respectively.*

References

- ASCHER, U. M., RUUTH, S. J., AND SPITERI, R. J. 1997. Implicit–explicit Runge–Kutta methods for time-dependent partial differential equations. *Applied Numerical Mathematics: Transactions of IMACS* 25, 2–3, 151–167.
- BARAFF, D., AND WITKIN, A. 1998. Large steps in cloth simulation. In *Proceedings of SIGGRAPH 98*, ACM Press / ACM SIGGRAPH, New York, NY, USA, 43–54.
- BARTH, E., KUCZERA, K., LEIMKUEHLER, B., AND SKEEL, R. 1994. Algorithms for Constrained Molecular Dynamics. March.
- BERCOVIER, M., AND PAT, T. 1984. A C^0 finite element method for the analysis of inextensible pipe lines. *Computers and Structures* 18, 6, 1019–1023.
- BERGOU, M., WARDETZKY, M., HARMON, D., ZORIN, D., AND GRINSPUN, E. 2006. A quadratic bending model for inextensible surfaces. In *Fourth Eurographics Symposium on Geometry Processing*, 227–230.
- BOXERMAN, E. 2003. *Speeding up cloth simulation*. Master’s thesis, University of British Columbia.
- BREEN, D. E., HOUSE, D. H., AND WOZNY, M. J. 1994. Predicting the drape of woven cloth using interacting particles. In *Proceedings of ACM SIGGRAPH 1994*, ACM Press/ACM SIGGRAPH, New York, NY, USA, 365–372.
- BRIDSON, R., FEDKIW, R. P., AND ANDERSON, J. 2002. Robust treatment of collisions, contact, and friction for cloth animation. *ACM Transactions on Graphics* 21, 3 (July), 594–603.
- BRIDSON, R., MARINO, S., AND FEDKIW, R. 2003. Simulation of clothing with folds and wrinkles. In *Symposium on Computer animation*, 28–36.
- CHOI, K.-J., AND KO, H.-S. 2002. Stable but responsive cloth. *ACM Transactions on Graphics* 21, 3, 604–611.
- CHOI, K.-J., AND KO, H.-S. 2005. Research problems in clothing simulation. *Computer-Aided Design* 37, 6, 585–592.
- DESBRUN, M., SCHRÖDER, P., AND BARR, A. 1999. Interactive animation of structured deformable objects. In *Graphics Interface ’99*, 1–8.
- EBERHARDT, B., WEBER, A., AND STRASSER, W. 1996. A fast, flexible, particle-system model for cloth draping. *IEEE Comput. Graph. Appl.* 16, 5, 52–59.

- EBERHARDT, B., ETZMUSS, O., AND HAUTH, M. 2000. Implicit-explicit schemes for fast animation with particle systems 137–154.
- FUHRMANN, A., GROSS, C., AND LUCKAS, V. 2003. Interactive animation of cloth including self collision detection. In *WSCG '03*, 141–148.
- GRIFFITHS, P., AND KULKE, T. 2002. Clothing movement—visual sensory evaluation and its correlation to fabric properties. *Journal of sensory studies* 17, 3, 229–255.
- HAIRER, E., LUBICH, C., AND WANNER, G. 2002. *Geometric Numerical Integration*. No. 31 in Springer Series in Computational Mathematics. Springer-Verlag.
- HAUTH, M., ETZMUSS, O., AND STRASSER, W. 2003. Analysis of numerical methods for the simulation of deformable models. *The Visual Computer* 19, 7-8, 581–600.
- HONG, M., CHOI, M.-H., JUNG, S., WELCH, S., AND TRAPP, J. 2005. Effective constrained dynamic simulation using implicit constraint enforcement. In *International Conference on Robotics and Automation*, 4520–4525.
- HOUSE, D. H., AND BREEN, D. E., Eds. 2000. *Cloth modeling and animation*. A. K. Peters, Ltd., Natick, MA, USA.
- HOUSE, D. H., DEVAUL, R. W., AND BREEN, D. E. 1996. Towards simulating cloth dynamics using interacting particles. *International Journal of Clothing Science and Technology* 8, 3, 75–94.
- MARSDEN, J. 1999. *Introduction to Mechanics and Symmetry*. Springer.
- MEYER, M., DEBUNNE, G., DESBRUN, M., AND BARR, A. H. 2001. Interactive animation of cloth-like objects in virtual reality. *The Journal of Visualization and Computer Animation* 12, 1 (Feb.), 1–12.
- MÜLLER, M., HEIDELBERGER, B., HENNIX, M., AND RATCLIFF, J. 2006. Position based dynamics. In *Proceedings of Virtual Reality Interactions and Physical Simulation (VRIPHYS)*, C. Mendoza and I. Navazo, Eds., 71–80.
- PROVOT, X. 1995. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Graphics Interface*, 147–154.
- SCHENK, O., AND GÄRTNER, K. 2006. On fast factorization pivoting methods for sparse symmetric indefinite systems. *Elec. Trans. Numer. Anal* 23, 158–179.
- TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically deformable models. In *Computer Graphics (Proceedings of ACM SIGGRAPH 87)*, ACM Press, New York, NY, USA, 205–214.
- TSIKNIS, K. D. 2006. *Better cloth through unbiased strain limiting and physics-aware subdivision*. Master’s thesis, The University of British Columbia.
- VOLINO, P., AND MAGNENAT-THALMANN, N. Comparing efficiency of integration methods for cloth simulation. *Computer Graphics International*.
- WITKIN, A., GLEICHER, M., AND WELCH, W. 1990. Interactive dynamics. *Computer Graphics (Proceedings of ACM SIGGRAPH 90)* 24, 2, 11–21.
- ZIENKIEWICZ, O. C., AND TAYLOR, R. C. 1989. *The finite element method*. McGraw Hill. 2.

Appendix A

We briefly explain why ICD and fast projection (FP) are not troubled by configurations (Q1), (Q2), and (Q4), and are resilient to (Q3). Facts about the behavior of SHAKE are taken from [Barth et al. 1994]. **Q1:** SHAKE’s force $\nabla C(\mathbf{x}^n)^T \boldsymbol{\lambda}^{n+1}$ cannot reduce the single edge’s length back to l ; our force $\nabla C(\mathbf{x}_j^{n+1})^T \boldsymbol{\lambda}^{n+1}$ can reduce that edge’s length back to l . **Q2:** here $\nabla C(\mathbf{x}_j^{n+1})$ and $\nabla C(\mathbf{x}^n)^T$ are both full-rank, yet SHAKE fails since $\nabla C(\mathbf{x}_j^{n+1}) M^{-1} \nabla C(\mathbf{x}^n)^T$ is singular; ICD and FP use $\nabla C(\mathbf{x}_j^{n+1}) D \nabla C(\mathbf{x}_j^{n+1})^T$, where D is symmetric full-rank, and this product is not singular. **Q4:** here $\nabla C(\mathbf{x}^n)$ is rank-deficient, so SHAKE fails; ICD and FP do not use $\nabla C(\mathbf{x}^n)$. Regarding **Q3**, ICD and FP may fail if $\nabla C(\mathbf{x}_j^{n+1})$ is rank-deficient; for sufficiently small timestep, h , this case is always avoidable.