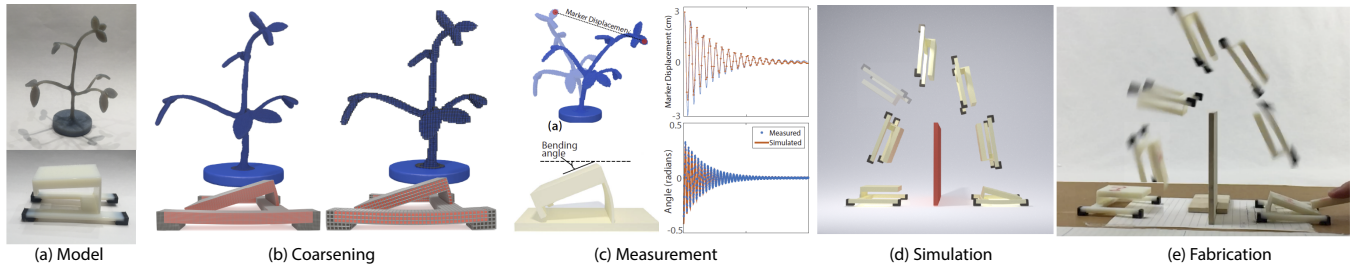


# Dynamics-Aware Numerical Coarsening for Fabrication Design

Desai Chen<sup>1,3</sup>David I.W. Levin<sup>2</sup>Wojciech Matusik<sup>1</sup>Danny M. Kaufman<sup>3</sup><sup>1</sup>MIT CSAIL<sup>2</sup>University of Toronto<sup>3</sup>Adobe Research

**Figure 1:** We introduce measurement-based, dynamics-aware coarsening (DAC) and the Boundary Balanced Impact (BBI) model - accelerating the simulation of dynamic elastica to obtain predictive and efficient accuracy required for fabrication-design testing and validation. We begin with initial models and real world fabricated materials (a). We then apply dynamics-aware, measurement-based coarsening (DAC) (b and c). The DAC model with BBI then simulates a range of designs that match the real-world dynamic behaviors of corresponding 3D-printed objects undergoing large-deformation loading, frictional contact and high-speed, transient dynamics with impact (e).

## Abstract

The realistic simulation of highly-dynamic elastic objects is important for a broad range of applications in computer graphics, engineering and computational fabrication. However, whether simulating flipping toys, jumping robots, prosthetics or quickly moving creatures, performing such simulations in the presence of contact, impact and friction is both time consuming and inaccurate. In this paper we present Dynamics-Aware Coarsening (DAC) and the Boundary Balanced Impact (BBI) model which allow for the accurate simulation of dynamic, elastic objects undergoing both large scale deformation and frictional contact, at rates up to 79 times faster than state-of-the-art methods. DAC and BBI produce simulations that are accurate and fast enough to be used (for the first time) for the computational design of 3D-printable compliant dynamic mechanisms. Thus we demonstrate the efficacy of DAC and BBI by designing and fabricating mechanisms which flip, throw and jump over and onto obstacles as requested.

## 1 Introduction

We present a pair of new methods to accurately simulate geometric and material nonlinearities subject to frictional contact, large loads and high-speed collisions at rates significantly more than an order-of-magnitude faster than previously available. Our methods combine efficiency and accuracy to enable design-for-fabrication optimization. They can be used for both fast, realistic animation and engineering analysis.

Advances in computational design, physical modeling and rapid manufacturing have enabled the fabrication of objects with customized physical properties. In computer graphics these extend from stable standing and spinning [Prévost et al. 2013; Bäcker et al. 2014] to mechanism design [Coros et al. 2013; Thomaszewski et al. 2014] and deformable character manufacture [Bickel et al. 2012; Skouras et al. 2013], to name just a few. Computational design in the presence of high-speeds, large deformation, frictional contact and impact, however, remains largely unaddressed.

Here we look towards a new generation of efficient mechanisms for practical *dynamic* function [Lipson 2014; Rus and Tolley 2015; Reis 2015; Reis et al. 2015]. In order to extend physics-driven

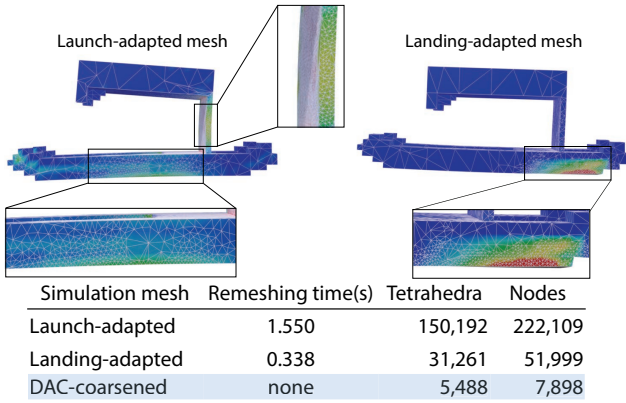
computational design to this domain, however, a bottleneck must be overcome - the physical simulation itself. Simulations must accurately replicate the behavior of elastic materials subject to high-speed, transient dynamics. Modeling these systems combines many of the remaining grand challenges in simulating elastica. Specifically we must accurately resolve nonlinear elasticity, large deformations, stiff materials, high-speed dynamics, rapid loading and unloading, frictional contact, internal friction, high-speed collisions, and rebound.

State-of-the-art FEM systems currently able to accurately match these effects are exceedingly expensive - runtimes on the order of days are standard to perform a single simulation in many cases [Belytschko et al. 2013]. Thus, while generating a single simulation for visualization or animation is already time consuming, the many simulations required during design optimization compound an already prohibitive computational burden

### 1.1 Efficiency with Accuracy

Let us explore four potential solutions for constructing fast *and* accurate simulation algorithms: (1) higher-order elements, (2) adaptive meshes, (3) reduced models, and (4) numerical coarsening. Can these methods provide the necessary efficiency to enable design-optimization while obtaining the predictive accuracy required to match fabricated results?

*Higher-order elements* offer us the opportunity to replace thin regions of our models with elements that capture higher-order deformation modes. While an attractive strategy, this poses two challenges. First, due to changing design parameters, we will need to identify suitable regions on-the-fly in order to perform this replacement. Second, coupling swapped-in higher-order elements to other element types introduces overhead. Consider, for example, replacing lower-order hexahedra with plate elements in these regions. We must then ensure continuity of displacements between plate-like portions of the design and thicker, volumetric portions. This, in turn, requires introducing difficult coupling constraints [Bergou et al. 2007; Martin et al. 2010]. Finally, even with such additional efforts, these substitutions may not always improve computational performance as high-order elements contain more DoFs than their low-order counterparts [Belytschko et al. 2013].




**Figure 2:** Comparing adaptivity and coarsening. A tetrahedral model of a jumper is adaptively remeshed to capture input stress fields from its launch and landing states. We compare the resulting element and node counts to our DAC-coarsened model. We set the minimum edge length for remeshing to one that was experimentally found to yield convergent numerical results.

Adaptive meshing allows us to reduce element counts in material regions where refinement is less critical. Let us ignore the difficulty of implementing adaptive meshing and the per time step cost to remesh. Even so, adaptive meshes are challenging in our setting. We model objects that undergo rapidly changing boundary conditions and globally varying stress fields due to contacts, loading and impacts. Adaptive meshing in our setting must then, necessarily, feature high numbers of elements to capture these details. See Figure 2. Here, using our experimentally validated, accurate element edge length as adaptive meshing threshold, our tetrahedral mesh contains 6.5X to 28X more nodal DoFs than our corresponding DAC-coarsened mesh. The DAC-coarsened mesh is likewise simpler to implement and has no per time step computational cost.

Reduced models utilizing linear modes [James and Pai 2002; Hauser et al. 2003] are widely applied to accelerate dynamic simulations. The key issue here is that linear modal models provide only a linear approximation of the deformation space leading to inaccurate linearization artifacts, such as swelling during rotation; see Figure 3. Optimized quadrature approaches, in turn, can afford efficient integration of non-linear forcing functions, but do not alleviate these artifacts [An et al. 2008] when relying on an underlying modal deformation space. Finally, nonlinear modal models [Barbič and James 2005] can alleviate some of these issues but so-far remain challenging to incorporate in the design process in comparison to their linear counterparts [Chinesta et al. 2013].

We begin by observing that numerical coarsening offers an exciting alternative for efficient yet predictive FE modeling. Coarsening methods effectively apply coarse resolution FE meshes as reduced DoF models and then seek material models that reproduce the behavior of a high-resolution FE counterpart. Analytical solutions for coarsening have been developed for linear material models (models where the stress varies linearly with strain) [Kharevych et al. 2009; Nesme et al. 2009; Torres et al. 2016]. Due to this linear assumption, and similarly to the linear modal models discussed above, we find them difficult to apply for the accurate modeling of the nonlinear materials required for 3D-printed objects. Even more recently, data-driven coarsening strategies [Chen et al. 2015] have been proposed to overcome the linear limitation of prior work. Unfortunately, while coarsening offers promise, prior work, to our knowledge, does not account dynamic effects, inertial properties, nor material damping characteristics. As we see in Figure 5, this causes



**Figure 3:** Linear modal models suffer from distortion as deformations grow. In (a) we illustrate increasing deformations, left to right, with corresponding inflation errors. Even for smaller deformations, (a) middle, the linear modal model still introduces significant distortions leading to modeling inaccuracies. In (b), left, we overlay simulations of small deformation performed respectively with the linear modal model (red) and the coarsened nonlinear FE model (grey); and, on the right, evaluate accuracy of the modal model. During simulation, even these smaller deformations introduce inaccuracies in the modal model due to element inflation; here up to 13.7%.


even the most recent, nonlinear, data-driven coarsening approaches to produce highly inaccurate dynamic simulations.

Building on the promise of coarsening techniques and inspired by recent developments in frequency matching for plausible computer animation [Li et al. 2014; Wang et al. 2015], we develop a new, dynamics-aware coarsening (DAC) method that, in contrast to prior approaches, provides well-over an order-of-magnitude performance enhancement, while maintaining fabrication-level accuracy when modeling highly dynamic motions subject to frictional contact. Our method does so without complex substructuring, does not require adaptive remeshing, accounts for dynamic effects including damping, and does not introduce prohibitive linear modeling artifacts and so is applicable to a wide selection of nonlinear constitutive models for 3D-printed materials.

## 1.2 Impact Response for Elastic Materials

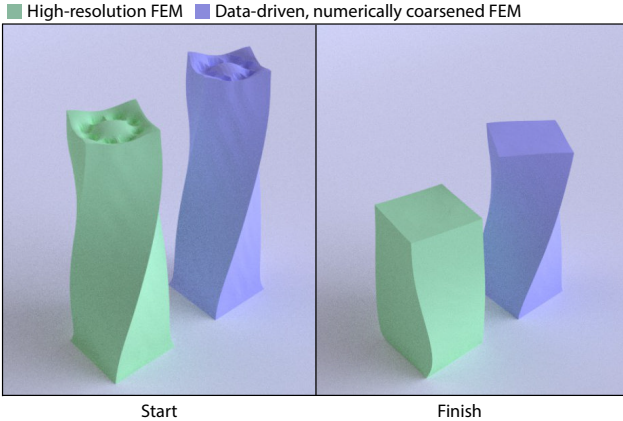
Even with a suitably accurate FE solution to model material dynamics, accurate impact response for elastic materials on collision remains highly challenging. To capture the bounces and rebounds of elastic mechanisms coming into contact with the ground - consider for example the heel strike of a sneaker - we need to get this right. State-of-the-art, implicit time-stepping methods for FEM with contact solve variational forms of time-steppers, e.g., variational Implicit Newmark, subject to additional, fully implicit contact and friction forces [Kane et al. 1999; Pandolfi et al. 2002]. These form so-called nonsmooth or complementarity integrators.

However, these complementarity integrators have two well-known flaws [Deuffhard et al. 2008]: (1) these methods can yield spurious oscillations on the contact boundaries and (2) the effective impact response of these methods is too energetic. Effectively the normal velocity on impact along the elastic boundary should be dissipated completely. Instead, complementarity methods with Newmark will generate an entirely incorrect elastic restitution; see Figure 4. To address these problems Deuffhard and colleagues [2008] introduced the now-standard DKE contact-stabilization step to filter contact response with projection. In the limit, DKE makes the impact-response model consistent, while in FE codes it is applied as an



**Figure 4:** Comparison of complementarity Newmark integration, DKE Stabilization and our BBI model for the impact resolution of an elastic, 3D-printed block. BBI closely matches the maximum rebound height achieved by the experimental result while both the Complementarity and DKE methods overestimate the rebound significantly.

effective strategy to recover from the well-known limitations of implicit integration with impact. However, it remains widely acknowledged that the right-way to accurately model high-speed collision response with implicit FE remains an open question at this time - not only in graphics - but more broadly in scientific computing as well. As an example consider Figure 4 where we see that both the complementarity model *and* the DKE filter produce different but equally incorrect predictions of the response of a stiff elastic block dropped on the floor. Here we offer a new, Boundary-Balancing Impact (BBI) model for FE that gains us accurate prediction of impact response for the stiff 3D-printed materials we focus on here; see Figure 4.



**Figure 5:** Dynamics-oblivious coarsening is significantly inaccurate for dynamic simulation. Left: two twisted elastic bars, initialized to the same configuration, are time-stepped in a dynamic simulation with an energy based, data-driven coarsening (DDFEM) model (blue) and a high-resolution FEM mesh (green). Right: small localized errors in the material model of the DDFEM simulation aggregate across the mesh over time to quickly produce large global errors when compared to the high resolution solution.

### 1.3 Summary

High-fidelity simulation methods for elastodynamics are too slow for use in fabrication design tasks while existing strategies to reduce the simulation cost of elastica (including adaptive, reduced and coarsened models) are too inaccurate and/or too expensive to employ. Finally, existing FE models for simulating elastic collision and rebound miss critical compliance coupling in the filter stage.

## 1.4 Contributions

We have exposed and analyzed the limitations of simulation methods for the predictive modeling of elastodynamics at rates sufficient for fabrication design optimization. Next, we develop our *Dynamics-Aware Coarsening* (DAC) method to address this need (§3). DAC jointly identifies *and* predictively simulates fabricated materials. To address current limitations in FE collision-response filtering we then introduce our *Boundary Balancing Impact* (BBI) model (§4). We then validate these contributions by comparing simulated results generated by DAC and BBI with real-world results experimentally obtained from a range of compliant 3D-printed jumping and throwing mechanisms that flip, throw projectiles, jump onto obstacles and jump over walls (§5).

## 2 Simulation Preliminaries

Accurate time-varying tracking of energy dissipation is key to accurate dynamic simulation [Marsden and West 2001] and so we rely on implicit Newmark time integration, a discrete variational integrator [Kane et al. 2000] with consistent high-quality energy tracking. We experimented with other integrators, including linearly implicit Newmark, Implicit Euler, and BDF2, but found them to be wanting in either stability or energetic behavior; see our Supplemental for details.

### 2.1 Discrete Model

We begin with a standard FEA material model for 3D printed materials, see e.g., Skouras et al. [2013]. To capture stiff elastic response of 3D-printed materials we use the neo-Hookean material model, augmented with Rayleigh damping to capture transient dissipation of vibrations, and discretize with second order, hexahedral finite elements. We start with a second-order, implicit Newmark time discretization

$$M\delta^{t+1} = b^t + \frac{h^2}{4}F(q^{t+1}) - \frac{h^2}{4}D(q^{t+1})v^{t+1}, \quad (1)$$

$$\begin{aligned} q^{t+1} &= q^t + \delta^{t+1}, \\ v^{t+1} &= \frac{2}{h}\delta^{t+1} - v^t, \end{aligned} \quad (2)$$

where

$$b^t = hMv^t + \frac{h^2}{4}F(q^t) - \frac{h^2}{4}D(q^t)v^t, \quad (3)$$

$h$  is the timestep size,  $F(\cdot)$  is the internal force vector,  $D(\cdot) = aM + bK(\cdot)$ , with  $a, b \in \mathbb{R}^+$  is the Rayleigh damping matrix,  $M$  is the stiffness-consistent mass-matrix [Belytschko et al. 2013],  $K(\cdot) = -\nabla F(\cdot)$  is the tangent stiffness matrix, and  $q, v \in \mathbb{R}^{3n}$  are respectively the nodal position and velocity vectors. In the absence of dissipative forces this method is symplectic and momentum preserving [Kane et al. 2000]. With dissipation we find that integration gives us accurate bookkeeping of system energy at comparable cost to implicit Euler.

### 2.2 Material Parameters

No matter how good our energy bookkeeping, the overall fidelity of our method is critically determined by the accuracy of the material parameters we select. While many material parameters are reported in the literature, there remains large and significant variation in these values across 3D-print batches, printing orientations and curings; see §5. For predictive simulation we need to identify these



**Figure 6:** Bottom: the dynamic behavior of a 3D-printed, elastic jumping mechanism in experiment. Top: our DAC-coarsening combined with our BBI impact model generate a simulation that predictively captures this experimental behavior at rates 79X faster than state-of-the-art FEM.

values. In addition, we must model dissipation requiring us to determine unreported damping properties; e.g.,  $a$  and  $b$  in the Rayleigh model. Finally, as discussed below and complicating matters even further, these material parameters are discretization dependent at non-convergent spatial resolutions. In the next section we will detail our DAC model to capture dynamic deformation, stiffness and damping at coarsened spatial resolutions.

### 2.3 Contact and Friction

For contact we need to model nonpenetration constraints and frictional contact forces that resist sliding along interfaces. Contacts are between object parts or between a part and a fixed boundary such as the ground. At each time step we apply continuous collision detection to the predicted trajectory to gather contact constraints into a contact set  $\mathcal{C}$ .

To simplify the following, for each such contact  $k \in \mathcal{C}$ , the relative acceleration between material points  $\mathbf{x}_i$  and  $\mathbf{x}_j$  (at contact  $k$ ) can be expressed via the map  $\Gamma_k : \dot{\mathbf{q}} \rightarrow \dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j$ . See our Supplement for details on construction of  $\Gamma_k$ . If  $\mathbf{y} \in \mathbb{R}^3$  is a force applied to point  $\mathbf{x}_i$ , and an equal but opposite force is applied to point  $\mathbf{x}_j$ , then  $\Gamma_k^T \mathbf{y}$  is the resulting generalized force applied to the contacting system.

In turn, points in contact apply an equal and opposite force along their shared, unit-length normal  $\mathbf{n}_k \in \mathbb{R}^3$ . In global coordinates this is equivalent to applying a force of magnitude  $\bar{\alpha}_k \in \mathbb{R}^+$  along a generalized normal

$$\mathbf{n}_k = \Gamma_k^T \mathbf{n}_k \in \mathbb{R}^{3n}, \quad (4)$$

to the system. The subspace of generalized normal directions

$$\mathbf{N} = (\mathbf{n}_1 \dots \mathbf{n}_{|\mathcal{C}|}) \quad (5)$$

then forms a basis for contact forces. Concatenating the corresponding force magnitudes in  $\alpha = (\bar{\alpha}_1, \dots, \bar{\alpha}_{|\mathcal{C}|})^T$ , the total contact force applied in the system is then  $\mathbf{N}\alpha$ .

Friction forces lie in the tangent plane orthogonal to the contact normal. At each contact  $k$  we sample an orthogonal pair of unit length vectors from the tangent plane. The  $3 \times 2$  matrix composed

column-wise of these samples is given by  $\mathbf{T}_k$  so that a friction force,  $\mathbf{f}_k \in \mathbb{R}^3$ , applied at a contact  $k$ , lies in the span of  $\mathbf{T}_k$  with  $\mathbf{f}_k = \mathbf{T}_k \bar{\beta}_k$ , where each  $\bar{\beta}_k \in \mathbb{R}^2$  gives the frictional response coefficients at contact  $k$ .

The total friction force applied to the system at each contact  $k$  must be equal and opposite and is  $\mathbf{f}_k = \Gamma_k^T \mathbf{T}_k \bar{\beta}_k$ . The generalized basis for a friction force at contact  $k$  is then

$$\mathbf{T}_k = \Gamma_k^T \mathbf{T}_k \in \mathbb{R}^{3n \times 2}. \quad (6)$$

We build the corresponding subspace of generalized tangent directions,

$$\mathbf{T} = (\mathbf{T}_1 \dots \mathbf{T}_{|\mathcal{C}|}) \quad (7)$$

and form the corresponding vector of frictional force coefficients as  $\beta = (\bar{\beta}_1^T, \dots, \bar{\beta}_{|\mathcal{C}|}^T)^T$ . The total friction force on the system is then  $\mathbf{T}\beta$ .

Contact and friction forces can be inexpensively modeled explicitly [Belytschko et al. 2013] but this introduces instabilities and nonphysical oscillations on boundaries even at small time step for stiffer materials [Deuffhard et al. 2008]. Thus FEM state-of-the-art generally turns to implicit time-integration for efficient contact force modeling.


Kane et al. [1999] proposed the now standard nonsmooth-Newmark method for contact modeling. This is a fully implicit time-stepping model that couples frictional contact with internal energies and forcing in each solve.

$$\mathbf{M}\delta^{t+1} = \mathbf{b}^t + \frac{h^2}{4}\mathbf{F}(\mathbf{q}^{t+1}) - \frac{h^2}{4}\mathbf{D}(\mathbf{q}^{t+1})\mathbf{v}^{t+1} + \frac{h^2}{2}\mathbf{N}\alpha + \frac{h^2}{2}\mathbf{T}\beta. \quad (8)$$

Note that here, unlike internal forces, contact forces are evaluated solely at the time step endpoint to ensure dissipation. Next, to fully define the time stepper, consistency conditions are required for contact and friction forces.

Enforcing the complementarity model for contact requires contact forces to balance along boundaries

$$0 \leq \alpha \perp \mathbf{N}^T \delta^{t+1} \geq 0, \quad (9)$$



**Figure 7:** Dynamics-aware coarsening (DAC) coarsens meshes to capture deformation with calibrated stiffness. We observe FE meshes capture accurate deformation modes up to a quite coarse resolution. However, these same coarse meshes suffer from numerical stiffening (see Figure 8). We design coarse mesh FE models by matching both significant deformation modes (above) and their captured material response (see Figures 11 and 21) to obtain efficient and predictive coarse-mesh FEM simulations of dynamics.

Friction in turn is modeled with the *Maximal Dissipation Principle*, requiring friction to maximize the rate of negative work done at each contact,  $-\mathbf{f}_k^T \dot{\mathbf{x}}_k$ . The total dissipation performed by friction is then

$$\sum_{k \in \mathcal{C}} (-\mathbf{f}_k^T \dot{\mathbf{x}}_k) = - \left[ \sum_{k \in \mathcal{C}} \mathbf{f}_k^T \Gamma_k \right] \dot{\mathbf{q}} = -\beta^T \mathbf{T}^T \dot{\mathbf{q}}. \quad (10)$$

Then, maximizing the Coulomb-constrained dissipation simultaneously at *all* contact points, with the implicit Newmark discretization [Pandolfi et al. 2002], gives us the final condition for our numerical integration


$$\min_{\beta} \{ \beta^T \mathbf{T}^T \left( \frac{2}{h} \delta^{t+1} - v^t \right) : \mu_k \bar{\alpha}_k \geq \|\bar{\beta}_k\|, \forall k \in \mathcal{C} \}, \quad (11)$$

where  $\mu_k$  is the local friction coefficient at contact  $k$ .

Taken together stepping this implicit *complementarity* Newmark integrator with (3), (8), (9), (11) and (2) provides predictive simulation at slower speeds of contact. For higher speed contacts and impacts, the remaining challenge lies in stabilizing contact stresses, velocities and displacements [Deuffhard et al. 2008] to ensure that simulated objects’ bounces and rebounds consistently match with their real-world counterparts. In Section 4 we will address this issue with a new impact model for FE simulation. First, however, we address our fundamental scaling problem: how can we gain accurate dynamic simulation without being bottlenecked by systems too large to solve quickly per time step?

### 3 Dynamics-Aware Coarsening

We couple numerical coarsening with parameter acquisition. Our DAC method computes *numerical* stiffness parameters for the nonlinearly elastic, Neo-Hookean material model, and damping parameters for the Rayleigh damping model so that, when applied to the coarsened simulation mesh, the dynamic behavior of a high-resolution simulation is preserved.



**Figure 8:** Static Deformation Test. We apply an identical load to three meshes with the same model geometry. With the same material parameters, a high-resolution mesh (left) is effectively 2.5x softer than the corresponding coarse mesh (right). Applying our captured numerical Young’s modulus to the coarse mesh (middle) regains the correct deformation of the original, high-resolution mesh on the left.

Because the goal of DAC is to replicate the dynamic behavior of a high-resolution simulation, we focus on creating a coarse mesh which captures both the large-scale deformation modes, and the corresponding natural frequencies of the high-resolution mesh. We do this in two stages. First, we produce a coarse hexahedral mesh that can replicate the large scale deformation modes of our high-resolution mesh. Second, we compute material and damping parameters that yield matching fundamental frequencies for each mode shape on this coarse mesh.

#### 3.1 Geometric Coarsening

DAC uses an iterative procedure to create the coarse mesh while maintaining mode shapes. We initialize our mesh to a coarse hexahedral discretization of the starting geometry,  $\mathbf{q}^0$ , and then subdivide recursively until we reach a convergent mesh resolution. We then solve the generalized mass-PCA system

$$\mathbf{K}(\mathbf{q}^0) \mathbf{q} = \lambda \mathbf{M} \mathbf{q} \quad (12)$$


for the dominant shape modes of the convergent system and then coarsen via bisection with mass-PCA until we reach a maximally coarse mesh that matches the dominant four shape modes to tolerance. We use a relative geometric difference of 5% (Hausdorff distance) as our tolerance threshold. Typically this is a short validation step as even the coarsest meshes generally satisfy this criteria (Figure 7).

#### 3.2 Material Parameter Fitting

Our geometric coarsening ensures that our DAC mesh captures significant deformation modes of our design accurately. However, when simulated, these same coarse meshes suffer from *numerical stiffening* – an increase in effective stiffness and damping as a consequence of decreased mesh resolution; see Figure 8. This leads to unacceptably inaccurate simulated trajectories no matter how we simulate this system. Regaining predictive stiffness and damping by refining the discretization would take us back to intractable mesh sizes; see Figure 2. Rather than refine, we keep our coarse mesh (as we have already ensured that our geometry is resolved there) and instead calibrate its frequency spectrum to directly match experiment.

To do this, we rescale our coarse model stiffness so that its fundamental frequency matches an observed frequency. As we will see below, this simple analysis sufficiently recovers effective stiffness and damping to regain a predictive nonlinear simulation with our coarse mesh.

From the geometric coarsening step above, we retain the numerical eigenvalues,  $\lambda_i^0$ , of our coarse mesh with corresponding defor-



**Figure 9:** Top: tracking the oscillations of a 3D-printed calibration rig allows us to measure mesh-dependent stiffness and damping parameters. Bottom: the resulting DAC-simulated frames at corresponding times for visual comparison to the captured motion.

mation modes  $m_i$  approximated linearly by the damped harmonic oscillator [Shabana 2012]

$$\ddot{m}_i = -(a + b\lambda_i)\dot{m}_i - \lambda_i m_i, \quad (13)$$

or equivalently

$$m_i(t) = A_i \exp(-d_i t) \sin(2\pi f_i t + \theta_i),$$

$$f_i = \frac{1}{2\pi} \sqrt{\lambda_i - \left(\frac{a + b\lambda_i}{2}\right)^2}, \quad (14)$$

$$d_i = \frac{1}{2}(a + b\lambda_i),$$

where  $a$  and  $b$  are Rayleigh Damping parameters and  $\lambda_i$  is the  $i^{\text{th}}$  eigenvalue associated with the  $i^{\text{th}}$  deformation mode.

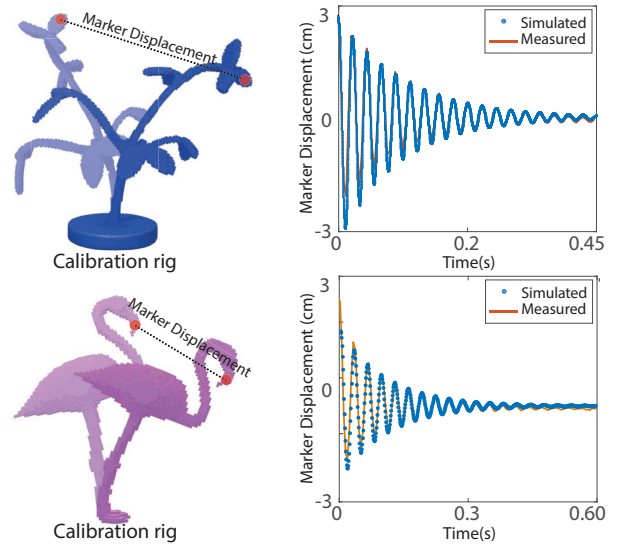
We then 3D-print calibration rigs with tracker markings, see e.g., Figure 9, and capture high-speed video (240 fps) of the rigs vibrating. We extract a tracked trajectory  $m_t$  of the marker motion from the video capture and solve the inverse harmonics problem [Mandelsham and Taylor 1997] to find the printed beam’s frequency,  $f_t$ , and damping,  $d_t$ , parameters. Setting the tracked  $f_t$  and  $d_t$  in (14) and  $a = 0$ , based on our observation of minimal effective mass-damping, simultaneously retrieves the captured target eigenvalue  $\lambda_t$  and the unknown and unreported stiffness damping parameter  $b$  required for dynamic simulation.

With our tracked  $\lambda_t$  in hand, our final step is to map the initial material Young’s modulus,  $E_m$ , that we use to compute  $\lambda_i^0$  (we set  $E_m = 1$  throughout) to a new, numerical modulus value,  $E_n$ . We seek an  $E_n$  that will match the numerical stiffness response of the simulated coarse FE mesh to the captured material response. We do this with a simple argument of fixed proportionality between the principle eigenvalues and the moduli by setting

$$E_n \leftarrow \frac{\lambda_t}{\lambda_1^0} E_m. \quad (15)$$

As validation we confirm that our coarsened FE simulations, initialized to the starting calibration pose, with Young’s modulus set to  $E_n$ , and stiffness damping set to the measured damping parameter, match both high-resolution simulation and the tracked calibration rig, up to viscosity—which we so far find unnecessary to model; see Figure 21.

Note that, in our experiments, increasing the number of mode shapes which must fall below our error threshold of 5% Hausdorff distance does not improve simulation accuracy. Figure 11 shows a comparison of DAC meshes created using 4 and 10 modes for



**Figure 10:** Measured and simulated vibrations of the 3D-printed plant (digital material) and 3D-printed flamingo (Rigur RGD450) models (magnify the plots for details of fits).

the error threshold. The resulting simulations are indistinguishable from each other and both match measured experimental data equally well.

In practice we observe that DAC captures object motions containing large contributions from a number of modes. To understand why we see that DAC scaling corrects frequencies of modes well beyond the first, so that, for example, for our plant and walker models, DAC reduces average frequency error across the first 10 modes from 27.1% to 2.4%.


Figure 10 shows an example of DAC coarsening applied to plant and flamingo models. With DAC we accelerate the simulation of the time varying deformation of our plant object by 25X while achieving good approximation to the high-resolution simulation mesh; see Figure 12. Here Figure 12 distinguishes between convergent FEM discretizations and accurate FEM discretizations for meshes used in our examples. Convergent discretizations are ones for which the spatial resolution is high enough so that the modal frequencies of the mesh have converged to their final values (changing by less than 1% with respect to previous subdivision). Accurate discretizations are ones for which the resolution of the FE mesh is such that the modal frequencies are within 5% of convergent values. In this paper all results are with respect to the coarser, accurate FE meshes. Even compared to these we achieve speedups of up to 79X. We also note that, if measurement data is not available, DAC calibration can be carried out using high-resolution simulation data.

## 4 Boundary Balancing Impact Model

With our DAC discretization in place, we will now derive a new, Boundary-Balancing Impact (BBI) model for FE that gains us accurate prediction of impact-response for the 3D-printed elastic materials we focus on in this work.

### 4.1 Complementarity Integration Revisited

The nonsmooth Newmark complementarity integrator we reviewed in Section 2 has several well-known flaws [Deuffhard et al. 2008] that we illustrate next. In Figure 14 right, we drop a 3D-printed



**Figure 11:** DAC coarsening comparison for increasing the number of fitted modes. Our default setting requires less than 5% relative Hausdorff distance up to the 4<sup>th</sup> mode and gives a coarsened FE mesh with  $dx = 0.6\text{mm}$ . If we increase the number of modes for our DAC fit to less than 5% relative Hausdorff for up to the 10<sup>th</sup> mode shape, we instead obtain a mesh with  $dx = 0.3\text{mm}$ . Note that the resultant simulations for these two DAC meshes are indistinguishable (magnify the plot for detail of fit).


block from height of 3 cm onto a flat surface. As it is both stiff and highly damped it lands without perceptible rebound. Yet, when we simulate the same drop of the block with the Newmark complementarity integrator the block rebounds up to a height of 0.89 cm; see Figure 14 left. Errors on this scale are unacceptable in a fabrication design process where they can make the difference between success and failure - see Figure 13.

What is going wrong in these examples? The Newmark discrete velocity update step in (2) gives rise to an arbitrary, undesirable (and for elastic materials) generally incorrect choice of restitution. Consider the impact of our material at contact with a normal  $n$ . Here the complementarity constraints ensure that the new displacement  $\delta^{t+1}$  along this normal are zero so that  $n^T \delta^{t+1} = 0$ . However, although this nicely satisfies position constraints, upon substitution we see that  $n^T v^{t+1} = -n^T v^t$  so that Newmark gives an incorrect, fully elastic (coefficient of restitution = 1) effective impact response. Yet the impact response for elastica along an impact boundary should instead be inelastic with the normal velocity on impact along the elastic boundary dissipated completely [Doyen et al. 2011]. In turn this results in much too large rebounds upon impact as we observe in Figures 13 and 14. A related error for complementarity integrators manifests in commonly observed spurious oscillations in positions and tractions along contact boundaries. These oscillations are the combined result of instabilities in contact stresses, velocities and displacements. Notably, both of these issues arise with arbitrary impact geometries, not just in the planar example discussed here.

### 4.2 DKE Contact Stabilization


To address these widely reported problems Deuffhard and colleagues [2008] introduced a contact stabilization step to filter contact response with projection. They observe that contact forces acting on the material boundary should be balanced and so proposed a now-standard FEM contact-stabilization filter (DKE) that applies an L2-projection that zeroes out normal displacements along the boundary of materials at contact interfaces.

This projection on displacement is performed at the end of each time step, after the position update (8) has been solved. This ensures a correct *inelastic* response at the contact boundary and is



Simulation	Model	dx(mm)	Timing avg/step(s)	Memory (G)	Elements
FEM	Plant	0.100	1876.3	48.5	1,004.6K
FEM	Plant	0.150	226.8	11.9	297.9K
DAC	Plant	0.600	9.2	0.4	16.1K
FEM	Jumper	0.150	—	379.7	5,450.0K
FEM	Jumper	0.375	791.0	18.0	346.3K
DAC	Jumper	1.500	10.4	0.3	5.4K

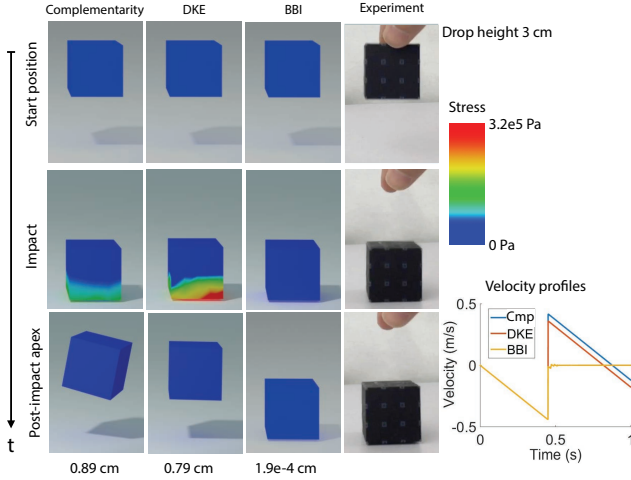
**Figure 12:** Statistics for our DAC simulations compared with two choices of accurate FE: a convergent FE model, and a validated accurate FE model (validated as matching experimental behavior and modal frequencies are within 5% of convergent values). For each simulation we report the model used (plant and jumper models), the mesh element size, the average wall-clock time spent per dynamic time step, memory usage, and the number of elements in the simulated mesh. Timings were recorded on an Intel Xeon E5-2666 v3, 2.9Ghz with 4 CPU threads.




**Figure 13:** Impact model testing with a 3D-printed jumper dropped onto the ground. We show overlaid frames captured, left to right, from an experiment (left) and two corresponding simulations (at  $h = 10^{-4}\text{s}$ ) that respectively apply DKE (middle) and BBI (right) impact models. In experiment the fabricated jumper lands, bounces up and then rests upright on its feet (left). However, simulation with the DKE model (middle) rebounds too high, flips over and so incorrectly predicts that the jumper will fail by landing on its back after collision (middle). With our BBI model (right), our simulation qualitatively predicts the experimentally determined landing behavior for this design.

effective in producing desirably stabilized contact tractions [Deuffhard et al. 2008; Krause and Walloth 2012]. Nevertheless, when

we compare DKE against experiment (Figures 13, 14, and 16), we see that the DKE projection likewise introduces unacceptably large rebound errors when compared with real-world results.



**Figure 14:** Impact validation test. A 3D-printed, stiff elastic block is dropped face-first on the ground. Left-to-right we compare the simulated results of three FE impact models - complementarity, DKE and our BBI model - with experimental results. We show configuration, at start, impact, and post impact maximum rebound height with details on stress distribution at impact, apex height reached, and (inset) the velocity profile for each simulation. As the effective restitution of elastica varies with angle of impact, see below in Figure 15 for a comparable oblique drop experiment.



**Figure 15:** Comparison of our BBI model to experiment for a 3D-printed cube dropped from an oblique initial orientation.

### 4.3 Boundary Balancing Impact

To better understand why DKE projection does such a poor job of resolving impacts in our elastic materials, let us consider again our simple dropped block example in Figure 14. We observe that by projecting out just the normal component of displacement along the boundary the DKE method artificially *concentrates* high stresses along the elements just inside the boundary. This can be seen in the impact row of Figure 14. These concentrated stresses effectively load the near-boundary layers which then spring back, introducing a much too large response as seen in the final row of Figure 14. The

problem here is that the L2-projection applied is material-oblivious and yet material properties clearly mediate impact response. Compliance distributes contact stresses quickly through an elastic material while, in damped materials, internal friction rapidly attenuates the response.

With these observations in mind we define a new Boundary Balancing Impact (BBI) model that can effectively impose boundary force-balance in a material-aware fashion. Starting with our base time integrator we define a compliant, discretization- and material-aware metric for projection below. The resulting stabilizing impact model better duplicates results in our design applications and experiments. Upon completing each time step from  $t - 1$  to  $t$  we replace the Newmark update in (2) with a compliant boundary projection update

$$\begin{aligned}
 c &= 2\delta^{t+1} - hv^t, \\
 A &= M + \frac{h^2}{4}K(q^{t+1}) + \frac{h}{2}D(q^{t+1}), \\
 d^* &= \operatorname{argmin}_d \left\{ \frac{1}{2} \|d - c\|_A^2 : N^T d \geq 0 \right\}, \\
 v^{t+1} &= \frac{1}{h} d^*, \\
 q^{t+1} &= q^t + \delta^{t+1}.
 \end{aligned} \tag{16}$$

This new impact model projects the explicit predicted velocity displacement,  $c$ , to the nearest set satisfying force balance on the boundary with respect to the local approximation of both material *stiffness* and *damping*. We find that this effectively distributes the contact-stabilized displacement across the material from the boundary layers. The effect of impact is communicated to the material interior while still ensuring that impacts are correctly inelastic along the active contact boundary. This leads to accurate predictions of real-world bouncing and rebounds. In our simple drop test we see in the impact row of Figure 14 that, at impact, response for this damped material has been correctly dissipated and the resulting normal displacement closely matches real world results in the apex row of Figure 14.

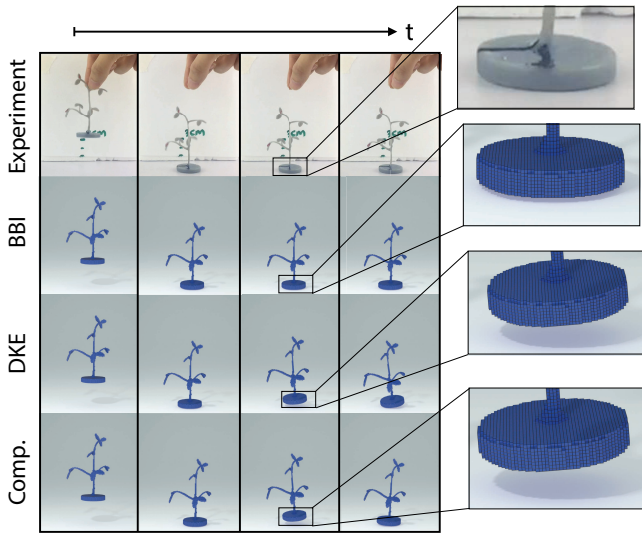
As we consider a range of impact angles, as well as impact with more complex geometries, multi-material 3D-prints and self contact we see that BBI still consistently and more accurately captures impact-response behavior. See Figures 13-17. These examples demonstrate the range and complexity of responses obtained by BBI coupling impact to stiff elastic materials with large deformations as well as objects composed of multiple materials undergoing both self-contact and stiction.

## 5 Evaluation

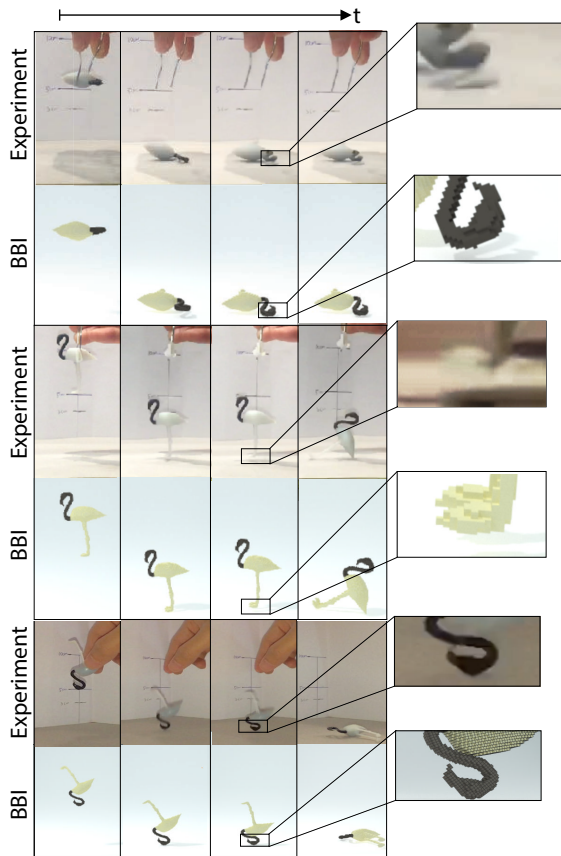
To test our proposed algorithms we compare simulated results generated by DAC and BBI with real-world results experimentally obtained from a range of compliant 3D-printed jumping and throwing mechanisms. For each mechanism we begin with a dynamic, time-varying goal, e.g., for our *jump-over* goal: “when pressed and released, jump over a given wall and land upright on the far side” (see Figure 25).

Simulations suitable for fabrication design must accurately predict both when a design fails and when it succeeds, thus we require real world examples of each. Below we outline our approach to creating these mechanism examples for testing, discuss our identification results, and review our implementation. We then detail our experiments comparing DAC and BBI’s simulated predictions for each design task below in §5.4 and validate outcomes of the simulations against repeated user trials in a study discussed in §5.5.





**Figure 16:** A drop test of the 3D-printed plant model (digital material). We compare simulated results from Complementarity, DKE and our BBI model with experiment. BBI solely reproduces the observed impact behavior.



**Figure 17:** Drop tests of a multi-material 3D-printed flamingo model (Rigur and TangoBlackPlus) from a variety of orientations. While effective restitution of elastica varies with angle of impact, in all cases our BBI simulations agree with experiment.

## 5.1 Creating Mechanism Examples

We focus our tests here on jumping-related mechanisms. The analysis and design of jumping mechanisms is an increasingly active [Churaman et al. 2011; Bergbreiter and Pister 2007; Bingham et al. 2014; Bergbreiter 2008; Jung et al. 2014; Koh et al. 2013; Koh et al. 2015; Vella 2015; Li et al. 2015], challenging and practical domain that incorporates high-speed transient dynamics of stiff elastic materials undergoing impact and so is an ideal test case for DAC and BBI. Research in jumping mechanisms has focused on efficient energy transfer into jump height, see e.g., Noh et al. [2012] with aligned research on a range of approaches for controlled jumping [Loepfe et al. 2015; Bartlett et al. 2015; Li et al. 2015]. In all cases, to our knowledge, mechanisms have been developed via costly, manual iterations of hands-on experiment, re-design, fabrication and one-off simulations [Cho et al. 2009; Bartlett et al. 2015], while even the stable landing of dynamic jumps has remained highly challenging [Jung et al. 2015].

For each of our jumping and throwing goals (see §5.4 below) we first attempted to manually create successful designs. We obtained designs that came close to ideal, but, consistent with the above cited literature, we were unable to hand tune these mechanisms to fully satisfy design goals. E.g., for the jump-over goal we found a design that often cleared the wall but did not land upright; see Figure 25, left. These mechanisms are our *initial* designs.

Presuming our initial designs are potentially close to successful designs, we perform local optimization over a pair of key design parameters, here generally length and height (see Figure 19), to seek a nearby solution. For each mechanism we pose its design goal as an objective and a set of constraints; these functions are detailed in Figure 18. To evaluate these objectives and constraints we apply DAC and BBI simulation at each design sample queried by the optimizer. Statistics for the numbers of simulation samples, iterations and timings for performing these local optimizations are summarized in Figure 18.

For all these examples we found that a relatively small number of iterations were needed to find successful designs; this suggests that our initial, hand-tuned designs are close to solutions in a basin. However, practical design optimization would demand a global optimization strategy to resolve non-convexity. In such general cases initial designs can be expected to start far from optima while the search space is often large - local optimization would be insufficient.

The mechanisms found by this process are our *final* designs. We compare the trajectories and outcomes predicted by BBI and DAC against the real world initial and final mechanisms: validation results are detailed for each example below in §5.4, while our user study results are presented in §5.5. The uncut video footage of all experiments are available online [Chen et al. 2017].


## 5.2 Implementation

All results are computed on an Amazon EC2 compute-optimized instance with 4 CPU threads (Intel Xeon E5-2666 v3, 2.9Ghz), while all mechanisms were printed on a Objet500.

While DAC and BBI give us orders of magnitude speedups for predictive simulation of deforming dynamics, our experiments (see Supplemental) show that dynamic FE simulation is unnecessary when modeling initial loading as well as during some portions of free-flight. With careful book-keeping and mapping of state simpler and more efficient models can be employed during these phases to gain further speedup.

Design Task	Objectives	Constraints	Iterations	Simulations	Time(m)
Flipper	$\theta_c^2$	—	5	44	242
Catapult	$\ x_c - x_{target}\ $	—	5	39	350
Jump Onto	$\theta_c^2$	$x_{wall} - x_c \leq 0,$ $y_{wall} - y_c \leq 0$	5	47	376
Jump Over	$\theta_c^2$	$x_{wall} - x_c \leq 0,$ $y_{wall} - y_{highest} \leq 0$	7	56	415

**Figure 18:** Design optimization statistics. For each dynamic design optimization we report the design task objectives and constraints applied, the number of optimization iterations performed, the total number of simulations performed for each design and finally the total wall-clock time (minutes) spent in design optimization. Sampled frames for each simulated and corresponding fabricated mechanism designs are given in Figures 22–26 and design parameters optimized over are summarized in Figure 19. Here  $\theta_c$  and  $x_c$  represent jumper angle of rotation and catapult projectile center-of-mass position at landing while  $x_{target}$  is the desired projectile target.



**Figure 19:** Design geometries and parameters.

When initially loading mechanisms, e.g., when pressing down a jumper in Figure 6, we observe that the process is effectively quasistatic and so simulate with an efficient quasistatic solver detailed in our Supplemental. Upon completion of the initial loading we map state to our full dynamic solver with DAC and BBI.


We also track the time-varying elastic potential energy stored in simulated compliant objects. When damping causes this internal potential to fall to zero, e.g., during portions of free-flight, we switch from our full DAC discretization to a rigid body discretization. We use DMV [Moser and Veselov 1991], an efficient energy-momentum preserving rigid-body integrator, to then time step the system in SE(3) coordinates until the next collision is reached, at which point we map rigid-body state back to the DAC model to capture the new deformation dynamics at impact. See our Supplemental for details on this process.

### 5.3 Identification

Material (model)	Young's Modulus (GPa)			Damping
	Reported	Identified	Numerical	
Rigur, vertical (jumper)	$1.9 \pm 0.2$	1.66	0.65	$1.64E-04$
Rigur, horizontal (jumper)	$1.9 \pm 0.2$	2.06	0.81	$1.40E-04$
DM4825, horizontal (plant)	$1.2 \pm 0.3$	1.57	0.72	$2.55E-04$
TangoBlackPlus (jumper)	$(2 \pm 1)E-04$	$6.32E-04$	$4.625E-04$	$9.10E-02$

**Figure 20:** Material parameters identified by our fine and coarse matching against calibration. Left to right we list previously reported Young's moduli for vertically and horizontally oriented prints compared with our identified moduli. We then report the matched numerical moduli we use for each for our DAC models and finally, list previously unreported damping parameters we identify and use in our simulations.

We compute each DAC model's coarse mesh resolution and material parameters using the measurement procedure described above




**Figure 21:** Measured and simulated vibrations of a 3D-printed jumper (Rigur) model (magnify the plot for details of fit).

in §3; see also Figure 21. We model heterogeneous materials by computing numerical Young's moduli for soft (TangoBlackPlus) and rigid (Rigur RGD450) materials and coarsening as much as possible while retaining a single material per element. A major benefit of our coarsening scheme is that it identifies both actual and numerical moduli and damping parameters of real-world materials. Figure 20 details these parameters. The damping parameters we present here have not, to our knowledge, been previously identified in the literature. On the other hand, in our experiments and simulations we find that the previously identified Poisson's ratio for our printed materials [Major et al. 2011], at 0.45, is effective, and keep it fixed at this value for all examples reported here.

### 5.4 Experiments

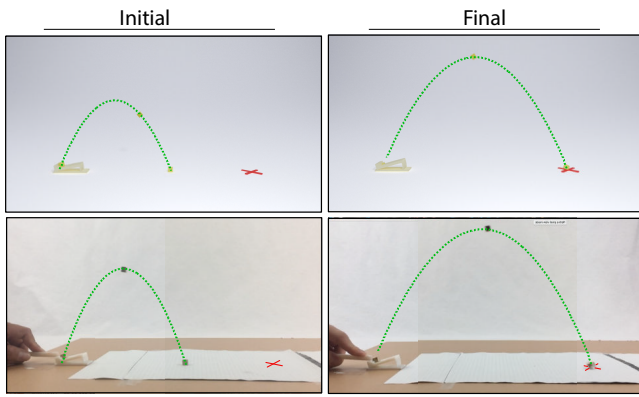
In this section we detail the individual design examples. We visually compare the trajectory behavior of both the real world fabricated initial and final mechanisms with corresponding DAC and BBI simulation results at the same parameters. In the following section we then discuss the results of the user study we perform to evaluate the results of initial and final fabricated mechanisms over repeated user trials comparing against DAC and BBI predicted simulation outcomes.



**Figure 22:** A comparison of experimental (bottom) and DAC/BBI simulated (top) results for initial (left) and final (right) designs of a flipping mechanism.

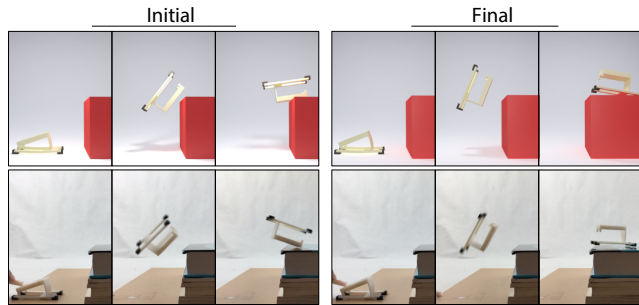
**Flipper** Our first design example is a simple forward flipping jumper. We begin with the geometry in Figure 22 left, load the jumping model by pressing down and then releasing. The design goal is a shape that, upon release, jumps forward into the air, flips and then lands stably on its feet, see e.g., Figure 6. See Figure 22 comparing simulated and experimental trajectories for both initial and final design samples.

**Catapult** Our next design example is a catapult mechanism. By adding a firing basket to the above flipper geometry, fixing the



**Figure 23:** A comparison of experimental (bottom) and DAC/BBI simulated (top) results for initial (left) and final (right) designs of the catapult mechanism.

mechanism base to the ground, and then adding a projectile cube in the basket to the design model we obtain a catapult for throwing metal cubes at targets. This system requires modeling the sliding contact and impact between the cube and basket. The design goal here is to find a catapult geometry that, under loading produces the correct combination of launch position and release velocity for the catapult arm and block so that the block hits a predetermined target. See Figure 23 comparing simulated and experimental trajectories for both initial and final design samples.




**Figure 24:** A comparison of experimental (bottom) and DAC/BBI simulated (top) results for initial (left) and final (right) designs for a jumper mechanism to jump onto a platform and land upright.


**Jumping onto obstacles** Here we consider a design example where the goal is to find a geometry for a 3D-printed jumper mechanism that, upon release, jumps forward and upwards into the air, flips (possibly multiple times) and then lands stably upon its feet on top of a flat obstacle, see e.g., Figure 24. See Figure 24 comparing simulated and experimental trajectories for both initial and final design samples.

**Jumping over obstacles** In this example, the goal is to find a geometry for a 3D-printed mechanism that, upon release, jumps forward and upwards high enough into the air to clear a wall, flip (possibly multiple times) over it and then land stably on the other side. See Figure 25 comparing simulated and experimental trajectories for both initial and final design samples.

**Flipper variations** We additionally evaluate six further design example comparisons between fabricated mechanisms and DAC



**Figure 25:** A comparison of experimental (bottom) and DAC/BBI simulated (top) results for initial (left) and final (right) designs for a jumper mechanism to jump over a wall of specified height and land upright.




**Figure 26:** A comparison of experimental (bottom) and DAC/BBI simulated (top) results for six more design samples: initial (left) and final (right) designs of variations on the flipping mechanism (“Normal”).

and BBI simulated trajectories over a range of flipper mechanism variations. In Figure 26 we compare initial and final designs created by three variations away from the base initial (“normal”) flipper mechanism design.

## 5.5 User Study Results

Here we present the results of a user study evaluating the goal outcomes of initial and final fabricated mechanisms over repeated user

trials as compared against outcomes predicted by DAC and BBI.



**Figure 27:** Statistics summarizing our user study comparing results from trials with our initial and final catapult mechanisms. We report the distance to target per trial. The final catapult design dramatically outperforms its initial counterpart, consistently, across all users, coming close to matching DAC/BBI predicted outcomes.

We asked five users to perform twenty attempts each with both the initial and final versions of each of the above mechanisms; the flipper, catapult, jump onto and jump over. For the flipper, jump onto and jump over mechanisms we count the number of successful attempts for each user; see Figure 28. Here we define goal success as satisfying the objectives and constraints posed by each design as described above; e.g. flipping, clearing all obstacles, and landing feet down on the desired area (Figures 22, 24 and 25). In Figure 28 we summarize the total number of successes for each user per optimized and unoptimized design as well as the aggregate totals. For the optimized and unoptimized catapult designs we report the mean distance to the target in millimeters and the standard deviation for each user (Figure 27).


In order to consistently apply the loading force to each mechanism, users are instructed to fully load each mechanisms by pressing the top until it makes contact with the bottom. We ask users to apply the loading force with a 3D printed bar at a designated loading point marked on each mechanism with permanent marker. The user then deformed the mechanism to the loaded state and released the load with a sliding motion, similar to the launching motion in Tiddlywinks. This ensures that contact break between the stick and the mechanism is close-to instantaneous and consistent. The same criteria are used in simulation to obtain comparable loading. The Objet500 we employ in all our fabrication examples works at 85 microns precision while our design parameter changes are on the order of millimeters and are thus be reliably manufactured. Additionally to further minimize variability in experimental conditions we use the same printer for all examples and always print in the top left area of the build tray. We orient models consistently and perform experiments within two weeks of printing to avoid long-term material degradation - an interesting topic for future modeling and design research.

Throughout the study, mechanism goal outcomes consistently matched those predicted by their corresponding DAC/BBI simulations. Final, optimized, mechanisms were much more reliable than their initial counterparts - under simulation these were successful. In terms of jumping (flipping, over, onto) tasks final designs completed more than 85% of their attempts for every task - close to matching DAC/BBI predicted success. while the initial designs successfully completed the simplest flipping task in 3% of all attempts (Figure 28) and had zero successes for the jump-onto and jump-over tasks - matching DAC/BBI predicted failure. This large difference validates that designed mechanism successes are quite reliably modeled by DAC/BBI. For instance, our catapult design achieves a 10 $\times$  reduction in mean error with respect to target distance for all users. Note that here there is a small increase in standard deviation due to the fact that our optimized design of the catapult necessarily throws the cube much further, amplifying

any variation errors in initial targeting (Figure 27). For supporting evidence of each mechanism’s experimental behavior please see our supplemental materials which include uncut videos [Chen et al. 2017] of all user studies.

## 6 Discussion and Conclusions

Beyond the direct application to fast and accurate dynamic simulation in contact-rich environments (see e.g., Figure 29), we see this work as a first step towards making the design optimization of functional dynamics broadly possible for real-world applications. What we show is that a principled approach to numerical coarsening, which accounts not just for quantities such as potential energy, but also mode shapes and damping, can allow for highly accurate dynamics simulations to be carried out with exceedingly low resolution meshes, leading to large gains in performance. However, while our method is the first to perform dynamics-aware coarsening for nonlinear viscoelastic materials, we are limited to material models parameterized by Young’s Modulus. Extending this technique to more general polynomial models would broaden its applicability.



**Figure 29:** A simulated deformable passive walking mechanism. In this preliminary work we take parameters captured from silicon to create “minion” walkers. We use six design parameters: three for length, height and depth of body; and three for length, height and depth of legs. We start with the closer walker as an initial, hand-tuned design. We then locally optimize from this design for a faster walking speed. The final design is farther from the viewer and obtains an emergent gallop to accelerate at the beginning of the motion and so, on the right, quickly pulls ahead of our initial design; see our supplemental video.


We note that prior to our introduction of DAC, design optimization of real-world transient dynamics was computationally intractable. DAC and BBI in the fabrication-centric dynamic simulation framework presented here begin to open exciting areas of future work: from the design of prosthetics and robotics, to the creation of highly dynamic toys, animatronic characters and beyond. As a next critical step on this path we look forward to the integration of DAC and BBI within a design optimization tool for fabrication.

## Acknowledgements

We thank Eitan Grinspun, Uri Ascher and Etienne Vouga for valuable discussion; Gaurav Bharaj for user study assistance; Dan Ramirez for video voice overs; and David Salesin for support throughout. This work was supported in part by the NSF (Grant CCF 1138967).

## References

- AN, S. S., KIM, T., AND JAMES, D. L. 2008. Optimizing cubature for efficient integration of subspace deformations. *ACM Trans. Graph.* 27, 5 (Dec.), 165:1–165:10.



**Figure 28:** Statistics summarizing our user studies comparing results from trials (left to right) with our our initial and final flipper, jump-onto, and jump-over mechanisms. In orange bars we report the number of successes per user, while in grey we report total successes per mechanism across all trials. Note that initial designs for both the jump-onto and jump-over have no successes at all, while, across all three mechanisms, final designs dramatically outperform their initial counterparts. This success is consistent across all users - closely matching DAC/BBI predicted outcomes.

- BÄCHER, M., WHITING, E., BICKEL, B., AND SORKINE-HORNUNG, O. 2014. Spin-it: Optimizing moment of inertia for spinnable objects. *ACM Trans. Graph.* 33, 4 (July), 96:1–96:10.
- BARBIČ, J., AND JAMES, D. L. 2005. Real-time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph.* 24, 3 (July), 982–990.
- BARTLETT, N. W., TOLLEY, M. T., OVERVELDE, J. T. B., WEAVER, J. C., MOSADEGH, B., BERTOLDI, K., WHITESIDES, G. M., AND WOOD, R. J. 2015. SOFT ROBOTICS. A 3D-printed, functionally graded soft robot powered by combustion. *Science* 349, 6244 (July), 161–165.
- BELYTSCHKO, T., LIU, W. K., MORAN, B., AND ELKHODARY, K. 2013. *Nonlinear Finite Elements for Continua and Structures*. John Wiley & Sons, Nov.
- BERGBREITER, S., AND PISTER, K. S. 2007. Design of an autonomous jumping microrobot. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, IEEE, 447–453.
- BERGBREITER, S. 2008. Effective and efficient locomotion for millimeter-sized microrobots. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 4030–4035.
- BERGOU, M., MATHUR, S., WARDETZKY, M., AND GRINSPUN, E. 2007. Tracks: Toward directable thin shells. *ACM Trans. Graph.* 26, 3 (July), 50:1–50:10.
- BICKEL, B., KAUFMANN, P., SKOURAS, M., THOMASZEWSKI, B., BRADLEY, D., BEELER, T., JACKSON, P., MARSCHNER, S., MATUSIK, W., AND GROSS, M. 2012. Physical face cloning. *ACM Trans. Graph.* 31, 4 (July), 118:1–118:10.
- BINGHAM, J. T., LEE, J., HAKSAR, R. N., UEDA, J., AND LIU, C. K. 2014. Orienting in mid-air through configuration changes to achieve a rolling landing for reducing impact after a fall. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 3610–3617.
- CHEN, D., LEVIN, D. I. W., SUEDA, S., AND MATUSIK, W. 2015. Data-driven finite elements for geometry and material design. *ACM Trans. Graph.* 34, 4 (July), 74:1–74:10.
- CHEN, D., LEVIN, D., MATUSIK, W., AND KAUFMAN, D. M. 2017. *Uncut user study videos*. [https://www.dropbox.com/sh/5quhghwg6vjjuz/AAAuBS8ilZ2Tsv52C\\_YQYrma](https://www.dropbox.com/sh/5quhghwg6vjjuz/AAAuBS8ilZ2Tsv52C_YQYrma).
- CHINESTA, F., LEYGUE, A., BORDEU, F., AGUADO, J. V., CUETO, E., GONZÁLEZ, D., ALFARO, I., AMMAR, A., AND HUERTA, A. 2013. Pgd-based computational vademecum for efficient design, optimization and control. *Archives of Computational Methods in Engineering* 20, 1, 31–59.
- CHO, K.-J., KOH, J.-S., KIM, S., CHU, W.-S., HONG, Y., AND AHN, S.-H. 2009. Review of manufacturing processes for soft biomimetic robots. *International Journal of Precision Engineering and Manufacturing* 10, 3, 171–181.
- CHURAMAN, W. A., GERRATT, A. P., AND BERGBREITER, S. 2011. First leaps toward jumping microrobots. *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011)*, 1680–1686.
- COROS, S., THOMASZEWSKI, B., NORIS, G., SUEDA, S., FORBERG, M., SUMNER, R. W., MATUSIK, W., AND BICKEL, B. 2013. Computational design of mechanical characters. *ACM Trans. Graph.* 32, 4 (July), 83:1–83:12.
- DEUFLHARD, P., KRAUSE, R., AND ERTEL, S. 2008. A contact-stabilized Newmark method for dynamical contact problems. *International Journal for Numerical Methods in Engineering* 73, 9, 1274–1290.
- DOYEN, D., ERN, A., AND PIPERNO, S. 2011. Time-Integration Schemes for the Finite Element Dynamic Signorini Problem. *SIAM Journal on Scientific Computing* 33, 1 (Jan.), 223–249.
- HAUSER, K. K., SHEN, C., AND O’BRIEN, J. F. 2003. Interactive deformation using modal analysis with constraints. In *Graphics Interface*, CIPS, Canadian Human-Computer Communication Society, 247–256.
- JAMES, D. L., AND PAI, D. K. 2002. Dyr: Dynamic response textures for real time deformation simulation with graphics hardware. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH ’02, 582–585.
- JUNG, G.-P., KIM, J.-S., KOH, J.-S., JUNG, S.-P., AND CHO, K.-J. 2014. Role of compliant leg in the flea-inspired jumping mechanism. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014)*, IEEE, 315–320.
- JUNG, S.-P., JUNG, G.-P., KOH, J.-S., LEE, D.-Y., AND CHO, K.-J. 2015. Fabrication of Composite and Sheet Metal Laminated Bistable Jumping Mechanism. *Journal of Mechanisms and Robotics* 7, 2 (Feb.), 021010.

- KANE, C., REPETTO, E. A., ORTIZ, M., AND MARSDEN, J. E. 1999. Finite element analysis of nonsmooth contact. *Computer Methods in Applied Mechanics and Engineering* 180, 1-2, 1–26.
- KANE, C., MARSDEN, J. E., ORTIZ, M., AND WEST, M. 2000. Variational integrators and the Newmark algorithm for conservative and dissipative mechanical systems. *International Journal for Numerical Methods in Engineering* 49, 10, 1295–1325.
- KHAREVYCH, L., MULLEN, P., OWHADI, H., AND DESBRUN, M. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph.* 28, 3 (July), 51:1–51:8.
- KOH, J.-S., PIL JUNG, S., WOOD, R. J., AND CHO, K.-J. 2013. A jumping robotic insect based on a torque reversal catapult mechanism. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2013)*, IEEE, 3796–3801.
- KOH, J.-S., YANG, E., JUNG, G.-P., JUNG, S.-P., SON, J. H., LEE, S.-I., JABLONSKI, P. G., WOOD, R. J., KIM, H.-Y., AND CHO, K.-J. 2015. BIOMECHANICS. Jumping on water: Surface tension-dominated jumping of water striders and robotic insects. *Science* 349, 6247 (July), 517–521.
- KRAUSE, R., AND WALLOTH, M. 2012. Presentation and comparison of selected algorithms for dynamic contact based on the Newmark scheme. *Applied Numerical Mathematics* 62, 10 (Oct.), 1393–1410.
- LI, S., HUANG, J., DE GOES, F., JIN, X., BAO, H., AND DESBRUN, M. 2014. Space-time editing of elastic motion through material optimization and reduction. *ACM Trans. Graph.* 33, 4 (July), 108:1–108:10.
- LI, S., KATZSCHMANN, R., AND RUS, D. 2015. A soft cube capable of controllable continuous jumping. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 1712–1717.
- LIPSON, H. 2014. Challenges and Opportunities for Design, Simulation, and Fabrication of Soft Robots. *Soft Robotics* 1, 1 (Mar.), 21–27.
- LOEPFE, M., SCHUMACHER, C. M., LUSTENBERGER, U. B., AND STARK, W. J. 2015. An untethered, jumping roly-poly soft robot driven by combustion. *Soft Robotics* 2, 1, 33–41.
- MAJOR, Z., REITER, M., HEMMETER, E., AND HIPTMAIR, F. 2011. Combination of novel virtual and real prototyping methods in a rapid product development methodology. *Polimeri* 32.
- MANDELSHTAM, V. A., AND TAYLOR, H. S. 1997. Harmonic inversion of time signals and its applications. *The Journal of Chemical Physics* 107, 17 (Nov.), 6756–6769.
- MARSDEN, J. E., AND WEST, M. 2001. Discrete mechanics and variational integrators. *Acta Numerica* 2001 10 (May), 357–514.
- MARTIN, S., KAUFMANN, P., BOTSCH, M., GRINSPUN, E., AND GROSS, M. 2010. Unified simulation of elastic rods, shells, and solids. *ACM Trans. Graph.* 29, 4 (July), 39:1–39:10.
- MOSER, J., AND VESELOV, A. P. 1991. Discrete versions of some classical integrable systems and factorization of matrix polynomials. *Communications in Mathematical Physics* 139, 2, 217–243.
- NESME, M., KRY, P. G., JEŘÁBKOVÁ, L., AND FAURE, F. 2009. Preserving topology and elasticity for embedded deformable models. *ACM Trans. Graph.* 28, 3 (July), 52:1–52:9.
- NOH, M., KIM, S.-W., AN, S., KOH, J.-S., AND CHO, K.-J. 2012. Flea-Inspired Catapult Mechanism for Miniature Jumping Robots. *IEEE Transactions on Robotics* 28, 5, 1007–1018.
- PANDOLFI, A., KANE, C., MARSDEN, J. E., AND ORTIZ, M. 2002. Time-discretized variational formulation of non-smooth frictional contact. *International Journal for Numerical Methods in Engineering* 53, 8 (Mar.), 1801–1829.
- PRÉVOST, R., WHITING, E., LEFEBVRE, S., AND SORKINE-HORNUNG, O. 2013. Make it stand: Balancing shapes for 3D fabrication. *ACM Trans. Graph.* 32, 4 (July), 81:1–81:10.
- REIS, P. M., JAEGER, H. M., AND VAN HECKE, M. 2015. Designer Matter: A perspective. *Extreme Mechanics Letters* 5 (Dec.), 25–29.
- REIS, P. M. 2015. A Perspective on the Revival of Structural (In) Stability With Novel Opportunities for Function: From Buckli-phobia to Buckli-philia. *Journal of Applied Mechanics* 82, 11, 111001.
- RUS, D., AND TOLLEY, M. T. 2015. Design, fabrication and control of soft robots. *Nature* 521, 7553 (May), 467–475.
- SHABANA, A. A. 2012. *Theory of Vibration*. Volume I: An Introduction. Springer Science & Business Media, New York, NY, Dec.
- SKOURAS, M., THOMASZEWSKI, B., COROS, S., BICKEL, B., AND GROSS, M. 2013. Computational design of actuated deformable characters. *ACM Trans. Graph.* 32, 4 (July), 82:1–82:10.
- THOMASZEWSKI, B., COROS, S., GAUGE, D., MEGARO, V., GRINSPUN, E., AND GROSS, M. 2014. Computational design of linkage-based characters. *ACM Transactions on Graphics* 33, 4 (July), 1–9.
- TORRES, R., RODRÍGUEZ, A., ESPADERO, J. M., AND OTADUY, M. A. 2016. High-resolution interaction with corotational coarsening models. *ACM Trans. Graph.* 35, 6 (Nov.), 211:1–211:11.
- VELLA, D. 2015. ROBOTICS. Two leaps forward for robot locomotion. *Science* 349, 6247 (July), 472–473.
- WANG, B., WU, L., YIN, K., ASCHER, U., LIU, L., AND HUANG, H. 2015. Deformation capture and modeling of soft objects. *ACM Trans. Graph.* 34, 4 (July), 94:1–94:12.