Hiding of Phase-Based Stereo Disparity for Ghost-Free Viewing Without Glasses

TAIKI FUKIAGE, TAKAHIRO KAWABE, and SHIN'YA NISHIDA, NTT Communication Science Laboratories

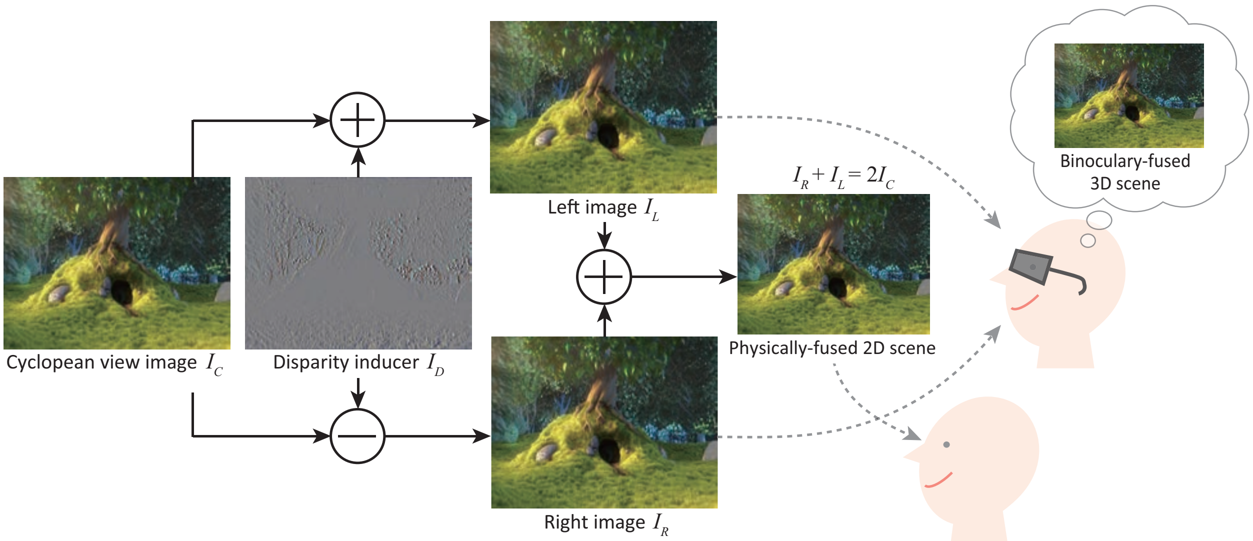

Fig. 1. A schematic of how Hidden Stereo cancels stereo ghosts for viewers without stereo glasses. Hidden Stereo generates a stereo pair by adding disparity-inducer patterns to a 2D image. The disparity-inducer patterns are identical except for the contrast polarity. Physical fusion of the stereo pair cancels out the disparity-inducer components, which makes only the original 2D image visible to viewers without stereo glasses. (The input image is taken from "Big Buck Bunny" ©by Blender Foundation, www.bigbuckbunny.org)

When a conventional stereoscopic display is viewed without stereo glasses, image blurs, or 'ghosts', are visible due to the fusion of stereo image pairs. This artifact severely degrades 2D image quality, making it difficult to simultaneously present clear 2D and 3D contents. To overcome this limitation (backward incompatibility), here we propose a novel method to synthesize ghost-free stereoscopic images. Our method gives binocular disparity to a 2D image, and drives human binocular disparity detectors, by the addition of a quadrature-phase pattern that induces spatial subband phase shifts. The disparity-inducer patterns added to the left and right images are identical except for the contrast polarity. Physical fusion of the two images cancels out the disparity-inducer components and makes only the original 2D pattern visible to viewers without glasses. Unlike previous solutions, our method perfectly excludes stereo ghosts without using special hardware. A simple algorithm can transform 3D contents from the conventional stereo format into ours. Furthermore, our method can alter the depth impression of a real object without its being noticed by naked-eye viewers by means of light projection of the disparity-inducer components onto the object's surface. Psychophysical evaluations have confirmed the practical utility of our method.

CCS Concepts: • **Computing methodologies → Perception**; *Image processing*;

Additional Key Words and Phrases: stereoscopy, spatial phase shift, backward compatible, perception

ACM Reference format:
Taiki Fukiage, Takahiro Kawabe, and Shin'ya Nishida. 2017. Hiding of Phase-Based Stereo Disparity for Ghost-Free Viewing Without Glasses. *ACM Trans. Graph.* 36, 4, Article 147 (July 2017), 17 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073672

All authors have equal contributions. This work is supported by Grants-in-Aid for Scientific Research on Innovative Areas (15H05915) from JSPS/MEXT Japan.
Author's addresses: T. Fukiage and T. Kawabe and S. Nishida, NTT Communication Science Laboratories, Nippon Telegraph and Telephone Corporation. 3-1 Morinosato Wakamiya, Atsugi, Kanagawa, 243-0198, Japan.

DOI: http://dx.doi.org/10.1145/3072959.3073672

1 INTRODUCTION

In standard 3D stereoscopic displays, an image for the left eye is superimposed on that for the right eye. For the viewers wearing shutter or polarized stereo glasses, the left and right images are separately presented to the corresponding eyes, and viewers with good stereopsis can enjoy the 3D contents with natural vergence angles. For viewers who unfortunately watch the display without stereo glasses, however, the superposition of the left and right images produces uncomfortable image blurs and bleedings, which are called 'ghosts' in this paper. Stereoscopic ghosts, caused by the binocular

Low-Cost 360 Stereo Photography and Video Capture

KEVIN MATZEN, MICHAEL F. COHEN, BRYCE EVANS, JOHANNES KOPF, and RICHARD SZELISKI, Facebook Inc.

Fig. 1. (a) our stereo rig; (b) left eye equirectangular image; (c) cross-eyed stereo pair from balcony scene.

A number of consumer-grade spherical cameras have recently appeared, enabling affordable monoscopic VR content creation in the form of full $360° \times 180°$ spherical panoramic photos and videos. While monoscopic content is certainly engaging, it fails to leverage a main aspect of VR HMDs, namely stereoscopic display. Recent stereoscopic capture rigs involve placing many cameras in a ring and synthesizing an omni-directional stereo panorama enabling a user to look around to explore the scene in stereo. In this work, we describe a method that takes images from two 360° spherical cameras and synthesizes an omni-directional stereo panorama with stereo in all directions. Our proposed method has a lower equipment cost than camera-ring alternatives, can be assembled with currently available off-the-shelf equipment, and is relatively small and light-weight compared to the alternatives. We validate our method by generating both stills and videos. We have conducted a user study to better understand what kinds of geometric processing are necessary for a pleasant viewing experience. We also discuss several algorithmic variations, each with their own time and quality trade-offs.

CCS Concepts: • **Computing methodologies → Computational photography**; **Image processing**; **Image-based rendering**; **Virtual reality**;

Additional Key Words and Phrases: image stitching, panoramas, virtual reality, stereo

ACM Reference format:
Kevin Matzen, Michael F. Cohen, Bryce Evans, Johannes Kopf, and Richard Szeliski. 2017. Low-Cost 360 Stereo Photography and Video Capture. *ACM Trans. Graph.* 36, 4, Article 148 (July 2017), 12 pages.
https://doi.org/10.1145/3072959.3073645

1 INTRODUCTION

Over the last couple of years, the general public has gained access to low-cost virtual reality head mounted displays capable of delivering immersive experiences. These devices range in price from hundreds of dollars for dedicated PC headsets down to just a few dollars for smartphone-based viewers. At the same time, fully-spherical cameras capable of capturing $360° \times 180°$ content using fisheye lenses have found an audience among photography enthusiasts as well as regular consumers. Devices such as the Ricoh Theta and the Samsung Gear 360 have enabled lightweight capture, sharing, and interactive viewing of monoscopic 360° panoramas on sites such as Facebook and YouTube. However, monoscopic panoramas fail to leverage one of the primary advantages of VR HMDs, namely their stereoscopic display.

Professional content producers can deliver high-quality, stereo panoramic stills and videos using complicated and expensive hardware as well as costly processing. Some notable examples include the Google Jump and Facebook Surround 360 cameras. By combining multiple video cameras into a ring-shaped configuration, an omni-directional stereo panorama can be stitched by interpolating rays between cameras. These devices are priced such that they are accessible to professionals, but not to regular consumers.

A lower-cost approach is to have a user sweep an arc out with their smartphone, thus simulating a camera ring, which is the core principle behind the Google Cardboard Camera. However, this process can be tedious, video cannot be supported, and scenes with motion pose a challenge.

In this paper, we seek the simplest and cheapest solution to producing omni-directional stereo stills and videos with existing consumer devices. When it comes to conventional stereo photography viewed from the original viewpoint, the simplest such configuration uses two pinhole cameras displaced by a fixed baseline corresponding to a human inter-pupillary distance, with each image displayed to each eye. The question we explore is whether a two-camera configuration can be used to enable an immersive three degrees of freedom viewing experience.

A conventional camera does not satisfy the requirement that we be able to turn our heads 360°. Even an extremely wide angle lens will not allow the viewer to turn and look behind themselves. Fortunately, we can make use of two spherical cameras to satisfy this requirement. However, there are several technical challenges that prevent us from naively mapping one eye to each camera. Figure 2 illustrates several of these:

(1) There is no placement of the cameras to provide proper stereo in all directions. Any placement results in varying

0730-0301/2017/7-ART148 $15.00
https://doi.org/10.1145/3072959.3073645

Mixed-primary Factorization for Dual-frame Computational Displays

FU-CHUNG HUANG, NVIDIA
DAWID PAJĄK, Light
JONGHYUN KIM, JAN KAUTZ, and DAVID LUEBKE, NVIDIA

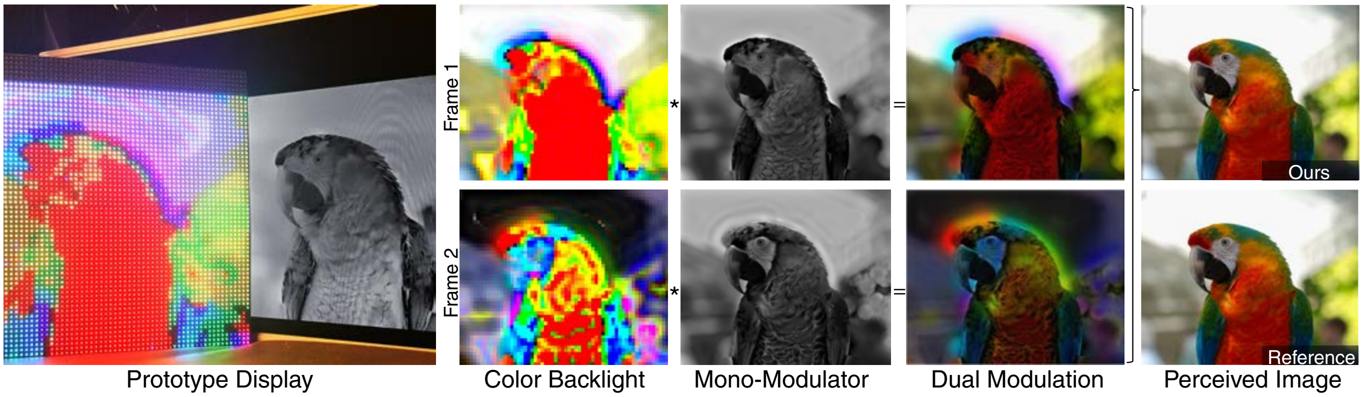

Fig. 1. The proposed display architecture. Our prototype uses a monochromatic LCD to attenuate a low-resolution color backlight generated by an LED array. Our two-frame factorization defines the color image as a pair of products of a low resolution color backlight with a high-resolution monochrome modulator. The final image (top right, simulated) is integrated through temporal (our prototype) or spatial multiplexing. Image courtesy of Derrick Coetzee.

Increasing resolution and dynamic range of digital color displays is challenging with designs confined by cost and power specifications. This necessitates modern displays to trade-off spatial and temporal resolution for color reproduction capability. In this work we explore the idea of joint hardware and algorithm design to balance such trade-offs. We introduce a system that uses content-adaptive and compressive factorizations to reproduce colors. Each target frame is factorized into two products of high-resolution monochromatic and low-resolution color images, which then get integrated through temporal or spatial multiplexing. As our framework minimizes the error in colorimetric space, the perceived color rendition is high, and thanks to GPU acceleration, the results are generated in real-time. We evaluate our system with a LCD prototype that uses LED backlight array and temporal multiplexing to reproduce color images. Our approach enables high effective resolution and dynamic range without increasing power consumption. We also demonstrate low-cost extensions to hyperspectral and light-field imaging, which are possible due to compressive nature of our system.

CCS Concepts: • **Hardware → Displays and imagers**;

Additional Key Words and Phrases: color optimization, LCD, LED array

ACM Reference format:
Fu-Chung Huang, Dawid Pająk, Jonghyun Kim, Jan Kautz, and David Luebke. 2017. Mixed-primary Factorization for Dual-frame Computational Displays. *ACM Trans. Graph.* 36, 4, Article 1 (July 2017), 13 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073654

DOI: http://dx.doi.org/10.1145/3072959.3073654

1 INTRODUCTION

Modern display design revolves around two sets of conflicting goals. On the one hand, there is a strong drive towards higher spatial resolution, dynamic range and color fidelity. This is evident in current high-definition television sets which reproduce images at 4K resolution (3840x2160) with high contrast and maximum brightness exceeding $500cd/m^2$. On the other hand, displays need to be energy efficient, low-cost and long-lasting. Systematic updates to display digital interfaces, such as VESA DisplayPort (DP), provide insight into trade-offs the industry made to satisfy the above goals. One example is Digital Stream Compression (DSC) in DP1.4. To reduce cost and power consumption of bandwidth intensive 8K (7680x4320) displays, VESA introduced a mandatory lossy compression directly in the display link. Lossy compression degrades the image quality and addresses only the transmission part of the problem: after decompression the actual data bandwidth required to drive the hardware electronics – writing the exact value of every pixel every frame – is still huge.

Most of the existing display technologies compromise some of the above goals to excel in others. The LCD design, for example, with its simple construction – a combination of a uniform backlight and a color filter array – yields low-cost at the expense of reduced spatial resolution and light efficiency. About two thirds of the light is absorbed by the color filters alone, which for HDR displays translates to a waste of hundreds of Watts. Alternative display designs have been proposed, but they all come with their own trade-offs. Field-sequential color displays are more power-efficient, but require very fast driving rates [Mori et al. 1999] which increases their cost and complexity. More recent OLED displays cannot produce high luminosity without risking early burn-out [Tsujimura 2012].

Bounce Maps: An Improved Restitution Model for Real-Time Rigid-Body Impact

JUI-HSIEN WANG, Stanford University
RAJSEKHAR SETALURI, Stanford University
DOUG L. JAMES, Stanford University
DINESH K. PAI, University of British Columbia

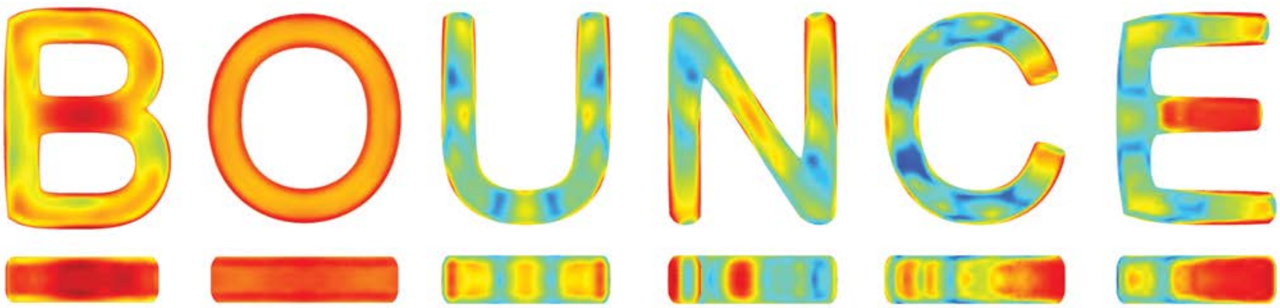

Fig. 1. **The coefficient of restitution is not constant!** Numerically computed coefficients of restitution ($blue = 0 \rightarrow red = 1$) are shown for these letters (with bottom view) and reveal significant spatial variations arising from micro-collision phenomena. Using a fast restitution analysis preprocess, we encode restitution values in Bounce Maps for fast runtime lookup during rigid-body simulation, thereby capturing natural variability in contact responses (see Figure 2). (All objects have the same physical material parameters ("steel"), and use #modes=45.)

We present a novel method to enrich standard rigid-body impact models with a spatially varying coefficient of restitution map, or *Bounce Map*. Even state-of-the art methods in computer graphics assume that for a single rigid body, post- and pre-impact dynamics are related with a single global, constant, namely the *coefficient of restitution*. We first demonstrate that this assumption is highly inaccurate, even for simple objects. We then present a technique to efficiently and automatically generate a function which maps locations on the object's surface along with impact normals, to a scalar coefficient of restitution value. Furthermore, we propose a method for two-body restitution analysis, and, based on numerical experiments, estimate a practical model for combining one-body Bounce Map values to approximate the two-body coefficient of restitution. We show that our method not only improves accuracy, but also enables visually richer rigid-body simulations.

CCS Concepts: • **Computing methodologies → Animation**; **Physical simulation**;

Additional Key Words and Phrases: Computer animation, rigid body, collision, contact, impact, chatter, coefficient of restitution, Newton's law of restitution, bounce, modal vibration.

ACM Reference format:
Jui-Hsien Wang, Rajsekhar Setaluri, Doug L. James, and Dinesh K. Pai. 2017. Bounce Maps: An Improved Restitution Model for Real-Time Rigid-Body Impact. *ACM Trans. Graph.* 36, 4, Article 150 (July 2017), 12 pages.
https://doi.org/10.1145/3072959.3073634

0730-0301/2017/7-ART150 $15.00
https://doi.org/10.1145/3072959.3073634

1 INTRODUCTION

Rigid-body impact and contact are of great importance to computer graphics. As in [Smith et al. 2012], we define contact to include resting or sliding contact, whereas impact captures instantaneous and transient collisions. Much work has been done in computer graphics to model contact and impact scenarios, as well as to develop efficient and robust techniques to handle complex collision scenarios; see Sec. 2 for a brief review.

Almost all previous work in graphics and mechanics has focused on algebraic collision laws, which *assume* a scalar coefficient of restitution, often treated as a material or object property. Indeed, most use the simplest and oldest model, Newton's Law of Restitution. Newton introduced this concept in his Principia, arguing that collision dynamics follow a simple law: the post-impact relative normal velocity v_n^+ of two objects is proportional to the pre-impact relative normal velocity v_n^-, where the ratio

$$\varepsilon := -\frac{v_n^+}{v_n^-} \in [0, 1], \tag{1}$$

is a measurable and material-dependent constant known as the *coefficient of restitution*. Thanks, in part, to Baraff's [1997] influential course notes that used ε to determine the contact impulse, this model is widely used in computer graphics.

The biggest attraction of a restitution model is, of course, its utter simplicity and efficiency. However, it has long been recognized that "rigid body" and "impact" are essentially contradictory, a meeting of an undeformable object with an impenetrable one. Such a modeling assumption would be ok (after all, many of the governing equations of classical mechanics are centuries old) if it incurred only a small or localized error, but the errors can be huge. For instance, Stoianovici

All's Well That Ends Well: Guaranteed Resolution of Simultaneous Rigid Body Impact

ETIENNE VOUGA, University of Texas at Austin
BREANNAN SMITH, Columbia University
DANNY M. KAUFMAN, Adobe Research
RASMUS TAMSTORF, Walt Disney Animation Studios
EITAN GRINSPUN, Columbia University

Iterative algorithms are frequently used to resolve simultaneous impacts between rigid bodies in physical simulations. However, these algorithms lack formal guarantees of termination, which is sometimes viewed as potentially dangerous, so failsafes are used in practical codes to prevent infinite loops. We show such steps are unnecessary. In particular, we study the broad class of such algorithms that are conservative and satisfy a minimal set of physical correctness properties, and which encompasses recent methods like Generalized Reflections as well as pairwise schemes. We fully characterize finite termination of these algorithms. The only possible failure cases can be detected, and we describe a procedure for modifying the algorithms to provably ensure termination. We also describe modifications necessary to guarantee termination in the presence of numerical error due to the use of floating-point arithmetic. Finally, we discuss the challenges dissipation introduce for finite termination, and describe how dissipation models can be incorporated while retaining the termination guarantee.

CCS Concepts: • **Computing methodologies → Physical simulation**; *Collision detection*;

Additional Key Words and Phrases: Collision response, Elastic impact, Rigid Bodies, Gauss-Seidel, Termination, Physical simulation

ACM Reference format:
Etienne Vouga, Breannan Smith, Danny M. Kaufman, Rasmus Tamstorf, and Eitan Grinspun. 2017. All's Well That Ends Well: Guaranteed Resolution of Simultaneous Rigid Body Impact. *ACM Trans. Graph.* 36, 4, Article 151 (July 2017), 19 pages.
https://doi.org/10.1145/3072959.3073689

1 INTRODUCTION

The numerical simulation of collisions between multiple objects in *simultaneous impact* is challenging. While conservation of energy and momentum completely determine the motion of a pair of colliding rigid balls, the same conservation laws are not sufficient to determine the behavior when three or more balls collide [Glocker 2004]. Additional assumptions are needed about the material properties of the objects and the way shocks propagate through them in order to make the contact problem well-posed.

Simultaneous impacts can occur in the time-continuous setting, but become even more prevalent when time is discretized, and one attempts to resolve all interferences that occur within the span of a time integration step. The problem is particularly well studied in the case of colliding rigid bodies, and it arises naturally in the simulation of granular media [Nguyen and Brogliato 2014]. This work focuses on responding to these impacts at one frozen instant in time, at the velocity level; responding to impact thus entails applying impulses to the colliding objects so that they are no longer approaching. We will not discuss the many alternative formulations based on position corrections, soft (acceleration-level) responses, etc, in detail; see any of several comprehensive surveys [Bender et al. 2014; Gilardi and Sharf 2002; Khulief 2012] for an overview.

When multiple objects are touching, resolving collisions between some of them often creates new collisions between others. For example, in a perfect head-on pool break, the cue ball first collides (only) with the ball at the tip of the pyramid, but a resolution of this collision in isolation induces new collisions against the two balls on the next level of the pyramid, and so forth. For this reason, the usual approach for solving the multiple impact problem is to adopt an iterative strategy [Han and Gilmore 1993]: a rule R is chosen for how to modify the velocities of objects to fix some subset of the collisions (we will call these rules *impact operators*). Applying the impact operator fixes some collisions while perhaps causing others; the operator is thus applied again, repeatedly, until all collisions are resolved. Algorithm 1 illustrates the general structure of algorithms that adopt this strategy; we call this structure *Gauss-Seidel-like*, by analogy to the iterative splitting method of the same name used to solve linear complementarity problems [Cottle et al. 1992; Erleben 2007]

A natural question that we must ask is ***will the iterative algorithm terminate, producing a collision-free solution in finite time?*** If not, what are the obstructions to termination? Clearly, Algorithm 1 is not guaranteed to terminate if the impact operator R is completely arbitrary. However, the rule R is necessarily restricted by physical considerations, and we can ask whether physical assumptions on R are sufficient to guarantee termination. In this setting termination has been shown for simple geometries (such as a straight line of balls of different masses, or impacts involving a limited number of objects [Jia et al. 2013]) but the general question remains open.

We thank Andy Ruina for discussions early in this project, and Eric Price for advice on proving our inexact arithmetic results. This work was supported in part by the National Science Foundation (DMS-1304211, 1117257, 1319483, 1409286, 1441889), Adobe Systems, The Walt Disney Company, Pixar, Nvidia Corporation, and Google.

0730-0301/2017/7-ART151 $15.00
https://doi.org/10.1145/3072959.3073689

Anisotropic Elastoplasticity for Cloth, Knit and Hair Frictional Contact

CHENFANFU JIANG, University of Pennsylvania and Jixie Effects
THEODORE GAST, University of California, Los Angeles and Jixie Effects
JOSEPH TERAN, University of California, Los Angeles and Jixie Effects

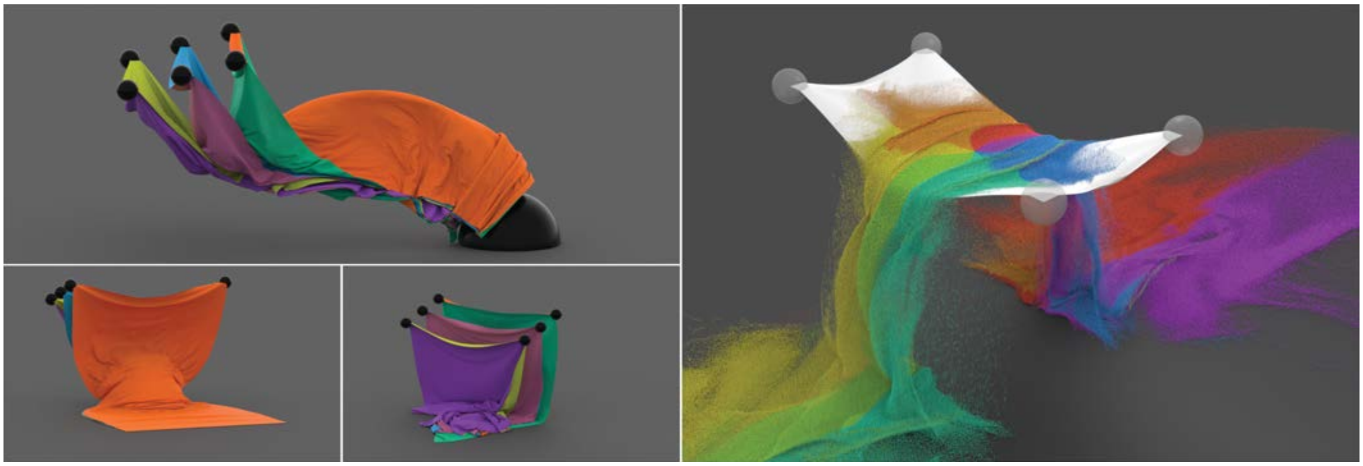

Fig. 1. The left figure shows three pieces of cloth with around 1.4M triangles pushed back and forth by a sphere, revealing intricate folds and contact. The right figure shows 7M colored sand grains coupled with elastic cloth, exhibiting beautiful flow patterns.

The typical elastic surface or curve simulation method takes a Lagrangian approach and consists of three components: time integration, collision detection and collision response. The Lagrangian view is beneficial because it naturally allows for tracking of the codimensional manifold, however collision must then be detected and resolved separately. Eulerian methods are promising alternatives because collision processing is automatic and while this is effective for volumetric objects, advection of a codimensional manifold is too inaccurate in practice. We propose a novel hybrid Lagrangian/Eulerian approach that preserves the best aspects of both views. Similar to the Drucker-Prager and Mohr-Coulomb models for granular materials, we define our collision response with a novel elastoplastic constitutive model. To achieve this, we design an anisotropic hyperelastic constitutive model that separately characterizes the response to manifold strain as well as shearing and compression in the directions orthogonal to the manifold. We discretize the model with the Material Point Method and a novel codimensional Lagrangian/Eulerian update of the deformation gradient. Collision intensive scenarios with millions of degrees of freedom require only a few minutes per frame and examples with up to one million degrees of freedom run in less than thirty seconds per frame.

CCS Concepts: • **Computing methodologies → Physical simulation**;

Additional Key Words and Phrases: MPM, elastoplasticity, friction, cloth, knit, hair

0730-0301/2017/7-ART152 $15.00
DOI: http://dx.doi.org/10.1145/3072959.3073623

ACM Reference format:
Chenfanfu Jiang, Theodore Gast, and Joseph Teran. 2017. Anisotropic Elastoplasticity for Cloth, Knit and Hair Frictional Contact. *ACM Trans. Graph.* 36, 4, Article 152 (July 2017), 14 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073623

1 INTRODUCTION

Physically based animation of elastic surfaces and curves has been an essential aspect of computer graphics for nearly three decades. Whether it be different layers of clothing in a virtual garment, individual strands in a head of hair or even yarns in a knit garment, collision and contact phenomena of these materials are essential for the richness and realism provided by physics based simulation. Unfortunately the thin nature of these materials makes collision detection and resolution challenging. This process is often the bottleneck in modern visual effects. Building on recent methods that characterize frictional contact in granular materials via elastoplasticity [Daviet and Bertails-Descoubes 2016; Klár et al. 2016; Narain et al. 2010], we present a new approach for codimensional elasticity that uses a hybrid Lagrangian/Eulerian Material Point Method (MPM) [Sulsky et al. 1994] discretization to model frictional contact with a continuum view. Unlike traditional approaches, our elastoplastic description completely characterizes all collision/contact response in the continuum and requires no separate post-processing.

Codimensional elastic objects are naturally represented with a Lagrangian mesh. Such approaches were pioneered in graphics by Terzopoulos et al. [1988; 1987] and are still primarily used today. With the Lagrangian mesh model, individual particles are tracked and mesh polygons or segments are used to approximate the spatial

Phace: Physics-based Face Modeling and Animation

ALEXANDRU - EUGEN ICHIM, EPFL
PETR KADLEČEK, Charles University in Prague / University of Utah
LADISLAV KAVAN, University of Utah
MARK PAULY, EPFL

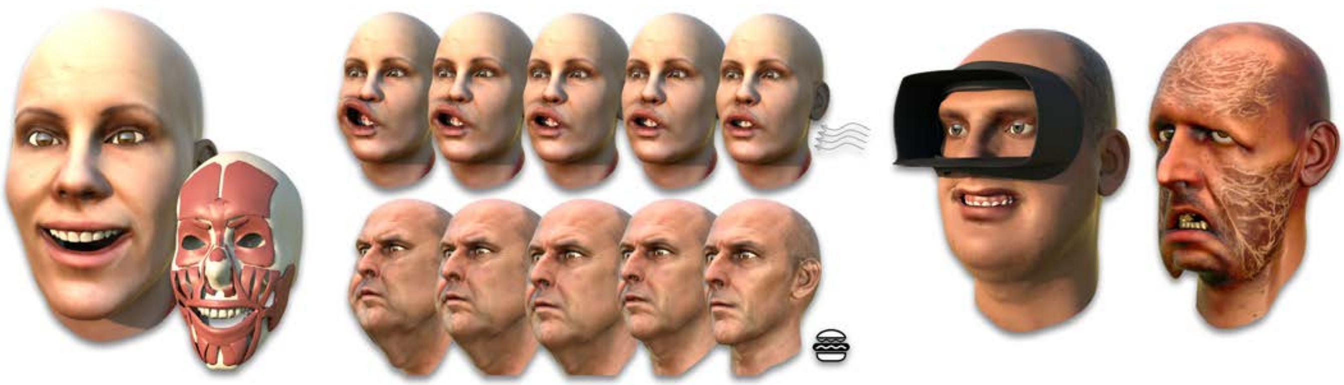

Fig. 1. Physics-based simulation facilitates a number of advanced effects for facial animation, such as applying wind forces, fattening and slimming of the face, wearing a VR headset, and even turning into a zombie.

We present a novel physics-based approach to facial animation. Contrary to commonly used generative methods, our solution computes facial expressions by minimizing a set of non-linear potential energies that model the physical interaction of passive flesh, active muscles, and rigid bone structures. By integrating collision and contact handling into the simulation, our algorithm avoids inconsistent poses commonly observed in generative methods such as blendshape rigs. A novel muscle activation model leads to a robust optimization that faithfully reproduces complex facial articulations. We show how person-specific simulation models can be built from a few expression scans with a minimal data acquisition process and an almost entirely automated processing pipeline. Our method supports temporal dynamics due to inertia or external forces, incorporates skin sliding to avoid unnatural stretching, and offers full control of the simulation parameters, which enables a variety of advanced animation effects. For example, slimming or fattening the face is achieved by simply scaling the volume of the soft tissue elements. We show a series of application demos, including artistic editing of the animation model, simulation of corrective facial surgery, or dynamic interaction with external forces and objects.

CCS Concepts: • **Computing methodologies → Physical simulation**;

Additional Key Words and Phrases: 3D avatar creation, facial animation, anatomical models, rigging

 0730-0301/2017/7-ART153 $15.00
DOI: http://dx.doi.org/10.1145/3072959.3073664

ACM Reference format:
Alexandru - Eugen Ichim, Petr Kadleček, Ladislav Kavan, and Mark Pauly. 2017. Phace: Physics-based Face Modeling and Animation. *ACM Trans. Graph.* 36, 4, Article 153 (July 2017), 14 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073664

1 INTRODUCTION

Accurate simulation of facial motion is of paramount importance in computer animation for feature films and games, but also in medical applications such as regenerative and plastic surgery. Realistic facial animation has seen significant progress in recent years, largely due to novel algorithms for face tracking and improvements in acquisition technology (Klehm et al. 2015; von der Pahlen et al. 2014).

High-end facial animations are most commonly produced using a sophisticated data capture procedure in combination with algorithmic and manual data processing. While video-realistic animations can be created in this manner, the production effort is significant and costly. A main reason is that complex physical interactions are difficult to recreate with the commonly employed reduced model representations. For example in blendshape rigs, collisions around the lip regions or inertial effects of the facial tissue are typically not accounted for. To remedy these shortcomings, artists often introduce hundreds of corrective shapes that need to be carefully sculpted and blended to achieve the desired effect in each specific animation sequence (Lewis et al. 2014).

Recent work (Barrielle et al. 2016; Ichim et al. 2016) proposes to avoid these shortcomings by augmenting the generative approach of blendshape animation with a simulation-based solution. A key

Facial Retargeting with Automatic Range of Motion Alignment

ROGER BLANCO i RIBERA*, KAIST
EDUARD ZELL*, Bielefeld University
J.P. LEWIS, Frostbite Labs and Victoria University
JUNYONG NOH†, KAIST
MARIO BOTSCH†, Bielefeld University

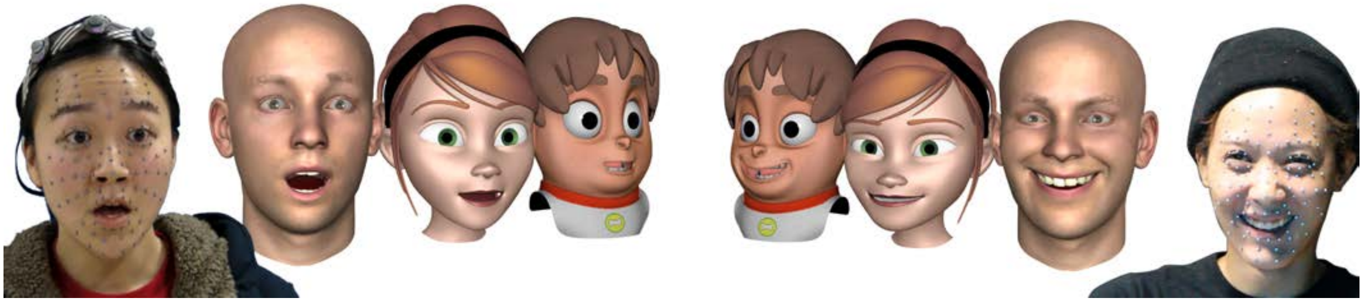

Fig. 1. Captured facial expression of two actors (*left* and *right*) with retargeting results to realistic and stylized characters. Our method automatically aligns the ranges of motion of the captured actor and the target blendshape rig, such that expressions are restored faithfully even for stylized characters.
From left to right: *The face rigs* Loki, Mery, *and* Billy, *are courtesy of* Mark Pauly, meryproject.com, *and* Jana Bergevin, *respectively. Motion capture data (*right*) courtesy of* Feel Ghood Music.

While facial capturing focuses on accurate reconstruction of an actor's performance, facial animation retargeting has the goal to transfer the animation to another character, such that the semantic meaning of the animation remains. Because of the popularity of blendshape animation, this effectively means to compute suitable blendshape weights for the given target character. Current methods either require manually created examples of matching expressions of actor and target character, or are limited to characters with similar facial proportions (i.e., realistic models). In contrast, our approach can automatically retarget facial animations from a real actor to stylized characters. We formulate the problem of transferring the blendshapes of a facial rig to an actor as a special case of manifold alignment, by exploring the similarities of the motion spaces defined by the blendshapes and by an expressive training sequence of the actor. In addition, we incorporate a simple, yet elegant facial prior based on discrete differential properties to guarantee smooth mesh deformation. Our method requires only sparse correspondences between characters and is thus suitable for retargeting marker-less and marker-based motion capture as well as animation transfer between virtual characters.

CCS Concepts: • **Computing methodologies → Animation**; *Motion capture*; *Motion processing*;

Additional Key Words and Phrases: facial animation, retargeting, blendshapes

*authors contributed equally; †corresponding author(s).

DOI: http://dx.doi.org/10.1145/3072959.3073674

ACM Reference format:
Roger Blanco i Ribera, Eduard Zell, J.P. Lewis, Junyong Noh, and Mario Botsch. 2017. Facial Retargeting with Automatic Range of Motion Alignment. *ACM Trans. Graph.* 36, 4, Article 154 (July 2017), 12 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073674

1 INTRODUCTION

Facial animation retargeting addresses the general problem of animation transfer between virtual characters, with the transfer of performance capture to virtual characters being the main application. Recent developments in vision- and depth-sensor-based facial motion capture [Cao et al. 2014; Ichim et al. 2015; Li et al. 2013; Thies et al. 2016; Weise et al. 2011] made accurate captures of an actor, traditionally limited to big film or game studios, affordable to a much broader audience. Current real-time capture systems typically adapt a realistic generic blendshape model to the actor. Since the modified and the original character have semantically equivalent blendshapes, the captured actor performance is then transferred between the characters by directly mapping the blendshape weights. The special case of equivalent blendshapes between two characters is often named *parallel parametrization* in retargeting context.

In practice, it is uncommon to encounter facial rigs with a complete set of semantically equivalent blendshapes. Creating facial rigs for animation is time consuming and requires highly skilled artists. Therefore, a rig is carefully designed to fit the animation needs, only modeling the necessary expressions. In addition, expressive digital characters are often stylized and exaggerate the facial proportions of humans. An effective retargeting method must either transfer animation from facial motion capture markers to a blendshape rig or between faces with different blendshape sets. Several

Example-Based Synthesis of Stylized Facial Animations

JAKUB FIŠER, Czech Technical University in Prague, Faculty of Electrical Engineering
ONDŘEJ JAMRIŠKA, Czech Technical University in Prague, Faculty of Electrical Engineering
DAVID SIMONS, Adobe Research
ELI SHECHTMAN, Adobe Research
JINGWAN LU, Adobe Research
PAUL ASENTE, Adobe Research
MICHAL LUKÁČ, Adobe Research
DANIEL SÝKORA, Czech Technical University in Prague, Faculty of Electrical Engineering

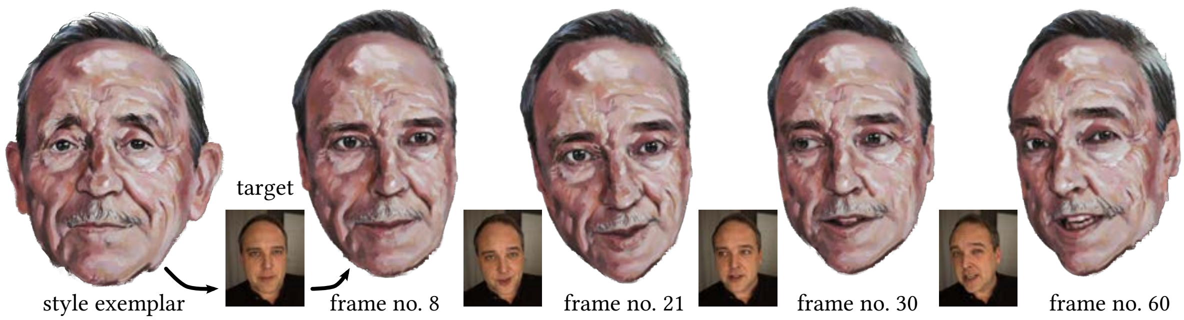

Fig. 1. The painterly style from an exemplar painting (far left) is transferred to a target video sequence (bottom). Exemplar painting: © *Graciela Bombalova-Bogra*. Video sequence: © *Ted Forbes* via YouTube.

We introduce a novel approach to example-based stylization of portrait videos that preserves both the subject's identity and the visual richness of the input style exemplar. Unlike the current state-of-the-art based on neural style transfer [Selim et al. 2016], our method performs non-parametric texture synthesis that retains more of the local textural details of the artistic exemplar and does not suffer from image warping artifacts caused by aligning the style exemplar with the target face. Our method allows the creation of videos with less than full temporal coherence [Ruder et al. 2016]. By introducing a controllable amount of temporal dynamics, it more closely approximates the appearance of real hand-painted animation in which every frame was created independently. We demonstrate the practical utility of the proposed solution on a variety of style exemplars and target videos.

CCS Concepts: • **Computing methodologies → Non-photorealistic rendering**; *Image processing*;

Additional Key Words and Phrases: Style Transfer, Hand-Drawn Animation

ACM Reference format:
Jakub Fišer, Ondřej Jamriška, David Simons, Eli Shechtman, Jingwan Lu, Paul Asente, Michal Lukáč, and Daniel Sýkora. 2017. Example-Based Synthesis of Stylized Facial Animations. *ACM Trans. Graph.* 36, 4, Article 155 (July 2017), 11 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073660

DOI: http://dx.doi.org/10.1145/3072959.3073660

1 INTRODUCTION

Recently, neural-based style transfer has become extremely popular thanks to the seminal work of Gatys et al. [2016] and its numerous publicly available implementations like *DeepArt* and *Prisma*. Selim et al. [2016] extended this technique to provide better results when stylizing head portraits. In their system, additional spatial constraints improve the resemblance between the stylized portrait and its real counterpart. They align the style image to the target photo and compute a set of gain maps to modify the response of the neural network in order to suppress the local differences in appearance.

Although their neural-based style transfer produces impressive results on various styles, it has one key limitation. For styles that contain rich textural information, the method tends to distort local visual features. In some cases, the overall appearance of the synthesized output becomes notably different from the original style exemplar (see Fig. 2, top row). This issue stems from the original method of Gatys et al. being based on a variant of parametric texture synthesis [Portilla and Simoncelli 2000], which is known to produce such artifacts [Efros and Freeman 2001]. Fišer at al. [2016] demonstrated that non-parametric texture synthesis can alleviate this issue, but it is not clear how to apply their analogy-based style transfer technique, designed for 3D rendering, to portrait stylization.

Another issue with Selim et al.'s approach is that it requires perfect alignment (warping) of the source style with the target photo. When the facial proportions of the stylized portrait differ considerably

A Data-driven Approach to Four-view Image-based Hair Modeling

MENG ZHANG, MENGLEI CHAI, HONGZHI WU, HAO YANG, and KUN ZHOU,
State Key Lab of CAD&CG, Zhejiang University

Fig. 1. Given four input hair images taken from the front, back, left and right views, our method computes a complete strand-level 3D hair model that closely resembles the hair in all images.

We introduce a novel four-view image-based hair modeling method. Given four hair images taken from the front, back, left and right views as input, we first estimate the rough 3D shape of the hair observed in the input using a predefined database of 3D hair models, then synthesize a hair texture on the surface of the shape, from which the hair growing direction information is calculated and used to construct a 3D direction field in the hair volume. Finally, we grow hair strands from the scalp, following the direction field, to produce the 3D hair model, which closely resembles the hair in all input images. Our method does not require that all input images are from the same hair, enabling an effective way to create compelling hair models from images of considerably different hairstyles at different views. We demonstrate the efficacy of our method using a wide range of examples.

CCS Concepts: • **Computing methodologies → Shape modeling**;

Additional Key Words and Phrases: hair modeling, image-based modeling, patch-based texture synthesis

ACM Reference format:
Meng Zhang, Menglei Chai, Hongzhi Wu, Hao Yang, and Kun Zhou. 2017. A Data-driven Approach to Four-view Image-based Hair Modeling. *ACM Trans. Graph.* 36, 4, Article 156 (July 2017), 11 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073627

1 INTRODUCTION

Hair plays a crucial role in producing realistically looking characters in computer-generated imagery. The generation of compelling 3D hairs, however, requires considerable efforts, due to the intricate hair structures and the wide variety of real-world hairstyles. It usually takes a few days of manual work by digital artists to create realistic 3D hair models in the entertainment industry, using dedicated hair design tools.

Recently, significant research effort has been devoted to hair digitalization, to help reduce the laborious work in the hair modeling process. Multi-view hair modeling [Hu et al. 2014; Luo et al. 2013; Paris et al. 2008] can generate impressive reconstructions of challenging, real-world hairstyles from dozens of images under different views, usually captured in a controlled environment. They often require complex acquisition setups, which are not easily accessible to average users. On the other hand, single-view hair modeling approaches [Chai et al. 2016, 2012; Hu et al. 2015] only take a single hair image as input, and produce plausible 3D hair models by using various forms of priors. As there is no input information about the hair at other views, the modeling result may not match the reality at views distant from the input one. Yu et al. [2014] propose a hybrid image-CAD system to create a 3D hair model from two or three images. Their focus is on visually pleasing results, not close visual matches to the input images.

In this paper, we introduce a novel four-view image-based hair modeling method, to fill in the gap between existing work on multi-view and single-view hair modeling. Given four images taken from the front, back, left and right views as input, our pipeline generates a complete strand-level 3D hair model that closely resembles the hair in all input images, with the help of a small amount of user interaction. We start by estimating the rough 3D shape of the hair observed in the input images, then synthesize a hair texture on the surface of the shape, from which the hair growing direction information on the hair shape surface is calculated and diffused to construct a 3D direction field in the hair volume. Finally, we grow hair strands from the scalp, following the direction field, to produce the 3D hair model.

In designing the above pipeline, we must address one challenging issue – how to generate a 3D hair model that matches *all* four input views with high fidelity. Unlike previous multi-view modeling methods which calculate accurate correspondences between a set

This work is supported by the National Science Foundation China (U1609215), the National Program for Special Support of Eminent Professionals of China, and the Fundamental Research Funds for the Central Universities (2017XZZX009-03). Corresponding authors: Hongzhi Wu (hongzhi.wu@gmail.com), Kun Zhou (kunzhou@acm.org).

DOI: http://dx.doi.org/10.1145/3072959.3073627

Interactive Design Space Exploration and Optimization for CAD Models

ADRIANA SCHULZ, JIE XU, and BO ZHU, Massachusetts Institute of Technology
CHANGXI ZHENG and EITAN GRINSPUN, Columbia University
WOJCIECH MATUSIK, Massachusetts Institute of Technology

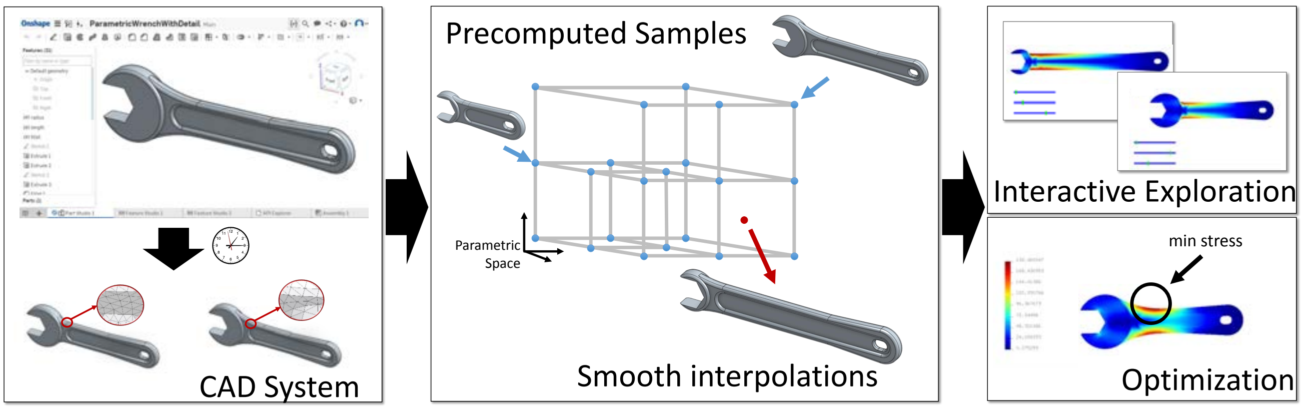

Fig. 1. Our method harnesses data from CAD systems, which are parametric from construction and capture the engineer's design intent, but require long regeneration times and output meshes with different combinatorics. We sample the parametric space in an adaptive grid and propose techniques to smoothly interpolate this data. We show how this can be used for shape optimization and to drive interactive exploration tools that allow designers to visualize the shape space while geometry and physical properties are updated in real time.

Computer Aided Design (CAD) is a multi-billion dollar industry used by almost every mechanical engineer in the world to create practically every existing manufactured shape. CAD models are not only widely available but also extremely useful in the growing field of fabrication-oriented design because they are parametric by construction and capture the engineer's design intent, including manufacturability. Harnessing this data, however, is challenging, because generating the geometry for a given parameter value requires time-consuming computations. Furthermore, the resulting meshes have different combinatorics, making the mesh data inherently discontinuous with respect to parameter adjustments. In our work, we address these challenges and develop tools that allow interactive exploration and optimization of parametric CAD data. To achieve interactive rates, we use precomputation on an adaptively sampled grid and propose a novel scheme for interpolating in this domain where each sample is a mesh with different combinatorics. Specifically, we extract partial correspondences from CAD representations for local mesh morphing and propose a novel interpolation method for adaptive grids that is both *continuous/smooth* and *local* (i.e., the influence of each sample is constrained to the local regions where mesh morphing can be computed). We show examples of how our method can be used to interactively visualize and optimize objects with a variety of physical properties.

This work is supported by the National Science Foundation, under grant CMMI-1547154, grant CCF-1138967, grant CMMI-1644558 and grant IIS-1409286.
Author's addresses: A. Schulz, 77 Massachusetts Ave, MIT 32-D414, Cambridge, MA 02139; email: adschulz@mit.edu.

0730-0301/2017/7-ART157 $15.00
https://doi.org/10.1145/3072959.3073688

CCS Concepts: • **Computing methodologies → Shape modeling**; *Shape analysis*; *Modeling and simulation*;

Additional Key Words and Phrases: CAD, parametric shapes, simulation, precomputations, interpolation

ACM Reference format:
Adriana Schulz, Jie Xu, Bo Zhu, Changxi Zheng, Eitan Grinspun, and Wojciech Matusik. 2017. Interactive Design Space Exploration and Optimization for CAD Models. *ACM Trans. Graph.* 36, 4, Article 157 (July 2017), 14 pages.
https://doi.org/10.1145/3072959.3073688

1 INTRODUCTION

Computational design typically involves fast exploration of shape variations or shape optimization. This field makes extensive use of parametric models: consistent families of geometries, each described by a given point in parameter space. Parametric shapes are advantageous because they constrain the manipulations to structure-preserving (and potentially fabrication-aware) variations and also reduce the search space, making computations more efficient. These techniques have been applied to many types of parametric shapes, from mechanical objets to garments, and masonry. So far, however, not much progress has been made in efficient exploration and optimization of the world's preeminent source of parametric data: CAD.

Practically every man-made shape that exists has once been designed by a mechanical engineer using a parametric CAD tool (e.g., SolidWorks, OpenScad, Creo, Onshape), and these models are available on extensive repositories (e.g., GrabCAD and Thingiverse, Onshape's public database). CAD shapes have the advantage of being parametric from construction, and the parameters exposed by expert engineers contain specific design intent, including manufacturing

Lightweight Structure Design Under Force Location Uncertainty

ERVA ULU, Carnegie Mellon University
JAMES MCCANN, Carnegie Mellon University
LEVENT BURAK KARA, Carnegie Mellon University

Fig. 1. We present a method for lightweight structure design for scenarios where external forces may contact an object at a multitude of locations that are unknown a priori. For a given surface mesh (grey), we design the interior material distribution such that the final object can withstand all external force combinations capped by a budget. The red volume represents the carved out material, while the remaining solid is shown in clear. Notice the dark material concentration on the fragile regions of the optimum result in the backlit image. The cut-out shows the corresponding interior structure of the 3D printed optimum.

We introduce a lightweight structure optimization approach for problems in which there is uncertainty in the force locations. Such uncertainty may arise due to force contact locations that change during use or are simply unknown a priori. Given an input 3D model, regions on its boundary where arbitrary normal forces may make contact, and a total force-magnitude budget, our algorithm generates a minimum weight 3D structure that withstands any force configuration capped by the budget. Our approach works by repeatedly finding the most critical force configuration and altering the internal structure accordingly. A key issue, however, is that the critical force configuration changes as the structure evolves, resulting in a significant computational challenge. To address this, we propose an efficient critical instant analysis approach. Combined with a reduced order formulation, our method provides a practical solution to the structural optimization problem. We demonstrate our method on a variety of models and validate it with mechanical tests.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; *Mesh models*; • **Applied computing** → *Computer-aided design*;

Additional Key Words and Phrases: structural analysis, structural optimization, digital fabrication

ACM Reference format:
Erva Ulu, James McCann, and Levent Burak Kara. 2017. Lightweight Structure Design Under Force Location Uncertainty. *ACM Trans. Graph.* 36, 4, Article 158 (July 2017), 13 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073626

DOI: http://dx.doi.org/10.1145/3072959.3073626

1 INTRODUCTION

With the emergence of additive fabrication technologies, structural optimization and lightweighting methods have become increasingly ubiquitous in shape design [Christiansen et al. 2015; Lu et al. 2014; Stava et al. 2012; Wang et al. 2013]. In many such methods, a common approach is to model the external forces as known and fixed quantities. In many real world applications, however, the external forces' contact locations and magnitudes may exhibit significant variations during the use of the object. In such cases, existing techniques are either not directly applicable, or require the designer to make overly conservative simplifications to account for the uncertainty in the force configurations [Choi and Park 2002].

We propose a new method for designing minimum weight objects when there exists uncertainty in the external force locations. Such uncertainties may arise in various contexts such as (*i*) multiple force configuration problems where the object experiences a large set of known force configurations such as those arising in machinery, (*ii*) unknown force configuration problems where the location of the contact points may change nondeterministically such as consumer products that are handled in a multitude of ways, or (*iii*) moving contact problems where a contact force travels on the boundary of an object; such as automated fiber placement manufacturing or cam-follower mechanisms.

Our approach takes as input (1) a 3D shape represented by its boundary surface mesh, (2) a user-specified *contact region*; a subset of the boundary where external forces may make contact, and (3) a *force-budget*; a maximum cap on the total summed magnitude of the external forces at any given time instance, and produces a minimum weight 3D structure that withstands any force configuration capped by the budget (Figure 1).

Design and Volume Optimization of Space Structures

CAIGUI JIANG, KAUST and Max Planck Institute for Informatics
CHENGCHENG TANG, KAUST and Stanford University
HANS-PETER SEIDEL, Max Planck Institute for Informatics
PETER WONKA, KAUST

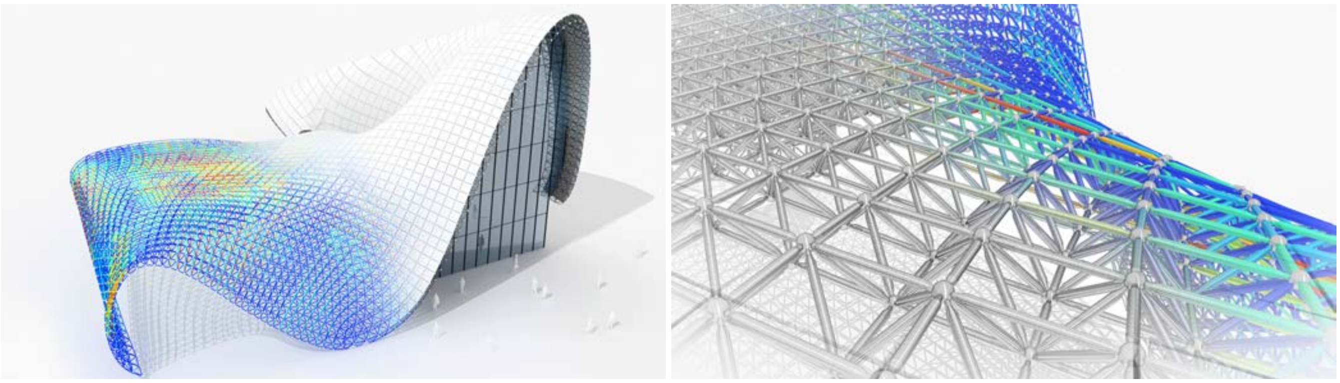

Fig. 1. Left: A statically sound space structure designed and optimized with our framework, motivated by the real architectural project shown in Figure 2. Right: The space structure is constructed with six types of customized beams to minimize the total volume of the material used for beams while maintaining moderate manufacturing complexity. Here, a hotter color indicates a larger beam cross-section area. Our framework automatically determines the optimal cross-section areas of the six types of beams as well as the assignment of beam types.

We study the design and optimization of statically sound and materially efficient space structures constructed by connected beams. We propose a systematic computational framework for the design of space structures that incorporates static soundness, approximation of reference surfaces, boundary alignment, and geometric regularity. To tackle this challenging problem, we first jointly optimize node positions and connectivity through a nonlinear continuous optimization algorithm. Next, with fixed nodes and connectivity, we formulate the assignment of beam cross sections as a mixed-integer programming problem with a bilinear objective function and quadratic constraints. We solve this problem with a novel and practical alternating direction method based on linear programming relaxation. The capability and efficiency of the algorithms and the computational framework are validated by a variety of examples and comparisons.

CCS Concepts: • **Computing methodologies → Shape modeling**;

Additional Key Words and Phrases: Architectural geometry, static equilibrium, structural optimization, optimization

ACM Reference format:
Caigui Jiang, Chengcheng Tang, Hans-Peter Seidel, and Peter Wonka. 2017. Design and Volume Optimization of Space Structures. *ACM Trans. Graph.* 36, 4, Article 159 (July 2017), 14 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073619

DOI: http://dx.doi.org/10.1145/3072959.3073619

1 INTRODUCTION

Space structures, also called space frames or space frame structures, are elegant and materially efficient *truss*-like structures consisting of beams (*two-force members*) connected at nodes. Space structures are desirable and often necessary in industrial design and architectural construction. The design and optimization of space structures present many challenges, especially for designs with complex geometries [Freund 2004].

In industrial design, space structures have been widely used for structures that should be both lightweight and statically sound, such as bikes, cars, or airplanes. A major advantage of space structures is their static soundness with limited material usage. In many situations, space structures also enable simple manufacturing processes by assembling or welding beams or bars. Moreover, the design space for space structures is rich, allowing for the possibility of the emergence of elegant structures that support designs with highly customized shapes. Similar ideas have also been extended to 3D printing and personalized design and fabrication.

In architectural construction, space structures often serve as statically sound supportive structures that approximate intended shapes or that underlie desired freeform surfaces. In many prominent projects, they are not visible. However, they are essential for the realization of the architectural design. For example, the Heydar Aliyev Cultural Center in Azerbaijan designed by Zaha Hadid has an underlying space structure as shown in Figure 2.

Space structures can also remain visible. Aesthetic considerations, such as simplicity and regularity, may influence the choice to make space structure visible. Visible space structures are not limited to