Sum-of-Squares Geometry Processing

ZOË MARSCHNER, PAUL ZHANG, DAVID PALMER, and JUSTIN SOLOMON,
Massachusetts Institute of Technology, USA

Self intersections in Coons patches

Continuous collision detection in NURBS patches

Bounding ellipsoids of tensor product patches

Fig. 1. We describe a framework for solving many different geometry processing problems on polynomial patches using sum-of-squares relaxation. Here, we show three of the problems we discuss in this paper: Self-Intersection, Continuous Collision Detection, and Minimal Bounding Ellipsoid. Our formulation is very general, working for any piecewise polynomial or rational patch type. We show each problem here on a model with a different patch type: cubic Coons patches from [Smirnov et al. 2020a], bicubic NURBS patches from [Trusty et al. 2021], and bicubic Bézier tensor patches.

Geometry processing presents a variety of difficult numerical problems, each seeming to require its own tailored solution. This breadth is largely due to the expansive list of *geometric primitives*, e.g., splines, triangles, and hexahedra, joined with an ever-expanding variety of *objectives* one might want to achieve with them. With the recent increase in attention toward *higher-order surfaces*, we can expect a variety of challenges porting existing solutions that work on triangle meshes to work on these more complex geometry types. In this paper, we present a framework for solving many core geometry processing problems on higher-order surfaces. We achieve this goal through sum-of-squares optimization, which transforms nonlinear polynomial optimization problems into sequences of convex problems whose complexity is captured by a single *degree* parameter. This allows us to solve a suite of problems on higher-order surfaces, such as continuous collision detection and closest point queries on curved patches, with only minor changes between formulations and geometries.

CCS Concepts: • **Computing methodologies** → **Parametric curve and surface models; Collision detection; Shape analysis.**

Additional Key Words and Phrases: parametric surfaces, collision detection, shape analysis, sum-of-squares optimization

ACM Reference Format:

Zoë Marschner, Paul Zhang, David Palmer, and Justin Solomon. 2021. Sum-of-Squares Geometry Processing. *ACM Trans. Graph.* 40, 6, Article 253 (December 2021), 13 pages. <https://doi.org/10.1145/3478513.3480551>

Authors' address: Zoë Marschner, zoem@mit.edu; Paul Zhang, [pzp1@mit.edu](mailto:pzpzp1@mit.edu); David Palmer, drp@mit.edu; Justin Solomon, jsolomon@mit.edu, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139, US.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).
0730-0301/2021/12-ART253
<https://doi.org/10.1145/3478513.3480551>

1 INTRODUCTION

Of the many geometric representations available today, *polynomial patches* are exceptionally powerful. They are a mainstay of computer-aided design and digital sculpting, where they provide piecewise-smooth parametrization. They can represent complex shapes with just a few control points. And in finite-element modeling, high degree polynomial bases can be used to construct high-fidelity solutions to partial differential equations.

Polynomial patches have recently made inroads into geometry processing problems involving PDEs and simulation [Jiang et al. 2021; Mandad and Campen 2020; Schneider et al. 2019, 2018; Trusty et al. 2021]. However, basic geometric kernels that are straightforward to implement for piecewise-linear meshes remain challenging for higher-order patches. These kernels include detecting self-intersections and collisions, computing bounding volumes, and measuring distances. Such kernels become especially important in *dynamics* problems, which in principle require continuous-time maintenance of physical feasibility.

We propose to bring the methods of *sum-of-squares programming* (SOS) to bear on this domain. The core idea of SOS programming is the replacement of polynomial positivity constraints with more computationally tractable SOS constraints of bounded degree d , which can be represented by semidefinite programs. The miracle of SOS programming comes from the Positivstellensatz [Blekherman et al. 2012], which states that for a large enough d , the globally-optimal solution to the modified problem is certifiably equivalent to the globally-optimal solution of the original problem. This certificate of correctness comes from the theory of SOS programming and is known as *exact recovery*. Thus, the SOS machinery allows us in many cases to solve seemingly non-convex optimization problems to global optimality.

A caveat to SOS relaxation is that the cost of solving the relaxed problem increases factorially with d [Blekherman et al. 2012]. While the minimal required d for exact recovery is problem-dependent and can be large in general, it is often close to the degrees of the polynomials in the original problem formulation, thereby maintaining tractability. Furthermore, we find that in practice the d required to obtain the correct solution to the original problem is often lower than the d required to obtain an exact recovery certificate. Since one often only cares about correctness of the solution, choosing the lower d allows the cost of the problem to be driven down in this case. If an algebraic proof of correctness is desired, the higher d can be chosen, at the expense of an increase in the cost of the problem.

Returning to geometry processing on polynomial patches, many geometric kernels can be formulated as polynomial optimizations. For example, surface-surface intersection (SSI) of quadratic triangles can be written as the minimization of a quartic objective function with linear inequality constraints. Applying SOS relaxation then yields a convex problem directly, and it only remains to verify that the relaxation successfully produces a solution to the original problem, i.e., that d was large enough. As illustrated by this example, the simplicity of SOS relaxation makes it readily adaptable to a wide variety of problems.

In this paper, we apply SOS methodology to several core problems in geometry processing on polynomial surfaces of varying degree. We show that a minor modification to our problem formulation allow us to support rational surfaces as well—namely, those consisting of NURBS patches. We verify the success of our SOS formulations on an exhaustive suite of test data. An overview of these results is provided in Table 1, where we show the experimentally determined minimum degree d_{correct} required to solve each problem with 100% accuracy. Finally, we apply these geometric kernels on various higher-order meshes, demonstrating the extensibility of geometry processing methods on linear meshes to higher-order surfaces. With these low-level operations out of the way, we pave the way for development of higher-level geometry processing techniques on higher-order surfaces.

2 RELATED WORK

SOS programming. Sum-of-squares (SOS) programming is a type of convex relaxation in which polynomial positivity constraints are replaced by SOS constraints. These can be transformed further into semidefinite optimization problems (SDP), which are solvable in polynomial time via interior-point methods [Alizadeh 1995; Boyd and Vandenberghe 2004; Nesterov and Nemirovskii 1994]. Modeling frameworks such as YALMIP [Löfberg 2004; Prajna et al. 2002] convert SOS formulations into SDPs, which can be solved by black-box solvers such as MOSEK and SDPT3 [MOSEK ApS 2017; Toh et al. 2001]. We provide the mathematical background relevant to this paper in Section 3. For a comprehensive review of this field, see [Blekherman et al. 2012].

SOS programming in geometry processing (GP). The use of SOS programming in geometry processing is fairly recent. Closest to our work is that of Marschner et al. [2020], where SOS programming is applied to verify injectivity of trilinear hexahedra and to repair hexahedral meshes failing injectivity. We will show that this problem

fits into our generalized framework. Prior to that, SOS programming was also used to find level set surfaces that encapsulate point clouds [Ahmadi et al. 2017].

Polynomial surfaces in GP. Polynomial surfaces have a long history in computer graphics and computer-aided design (CAD), with spline surfaces such as piecewise-bicubic Catmull-Clark surfaces [Catmull and Clark 1978] and piecewise-quartic Loop subdivision surfaces [Loop 1987] remaining popular in applications today. Recently geometry processing has seen a renaissance of polynomial meshes. Mandad and Campen [2020] use Bézier curves to compute curvilinear triangle meshes of planar domains, while Karčiauskas and Peters [2020] produce bicubic spline surfaces for locally quad-dominant meshes. Jiang et al. [2021] develop a method to generate coarse higher-order tetrahedral meshes from linear ones.

Polynomial meshes have also proven useful in solving PDEs from geometry and simulation. Catmull-Clark surfaces have been used for thin shell simulation [Wawrzinek et al. 2011]. Schneider et al. [2019, 2018] use higher-order basis functions to solve various PDEs, including linear elasticity. Cardoze et al. [2004] develop a method for maintaining quality of curved triangle meshes through fluid motion.

GP problems for higher-order surfaces. Trusty et al. [2021] develop a method for simulation of elasticity on volumes enclosed by NURBS patches without volumetric remeshing. A notable gap in these methods is lack of exact collision detection and intersection prevention. We aim to show one path toward adding these key geometry kernels via SOS programming.

Surface-surface intersection (SSI) is the problem of finding intersection curves between surfaces. SSI is frequently used in CAD, where curves are obtained by first linearizing the surfaces to find an initial intersection point, followed by stepping along the common tangent direction of the two surfaces. This method requires tuning of various tolerance parameters, as well as a sufficiently dense initial linearization. We refer the reader to [Barnhill et al. 1987] for a summary of SSI in CAD. Pekerman et al. [2008] present several algorithms for self-intersection detection and removal in curves and surfaces. Oh et al. [2012] provide a culling approach combined with Newton-Raphson iterations to compute closest point projection to freeform curves and surfaces.

3 PRELIMINARIES

For a complete review of sum-of-squares theory we refer the reader to [Blekherman et al. 2012; Parillo 2019]; Marschner et al. [2020] also present similar background in the context of a geometry processing problem. For completeness, we recall the most relevant concepts here.

3.1 Positive Polynomials and SOS Polynomials

Let $\mathbb{R}[\mathbf{u}] = \mathbb{R}[u_1, \dots, u_k]$ be the ring of real multivariate polynomials in $\mathbf{u}$. $\mathbb{R}[\mathbf{u}]_d$ denotes the subset of polynomials of degree at most d . We will use $[\mathbf{u}]_d$ to denote the basis of monomials up to degree d . Any member $f(\mathbf{u}) \in \mathbb{R}[\mathbf{u}]_d$ can be written in this basis: $f(\mathbf{u}) = [\mathbf{u}]_d^T \mathbf{f}$, where $\mathbf{f}$ denotes the vector of coefficients of monomials in f .

A special subset of $\mathbb{R}[\mathbf{u}]$ (resp. $\mathbb{R}[\mathbf{u}]_d$) is the cone of *positive polynomials* P (resp. P_d). As the name suggests, positive polynomials $f(\mathbf{u}) \in P$ satisfy $f(\mathbf{u}) \geq 0$ for all $\mathbf{u}$. Many polynomial optimization problems can naturally be written with positive polynomial constraints, and thus it is highly desirable to be able to optimize over P . Unfortunately, even membership testing in P is NP-hard in general [Blekherman et al. 2012, §3.4.3].

This leads us to the more restrictive subset Σ of *sum-of-squares (SOS) polynomials*; we use Σ_d to denote SOS polynomials of bounded degree d . Members of this set $f(\mathbf{u}) \in \Sigma$ can be decomposed into sums of squares of polynomials: $f(\mathbf{u}) = \sum_i s_i(\mathbf{u})^2$ for $s_i \in \mathbb{R}[\mathbf{u}]$. Naturally, they form a subset of P , giving us the following inclusions:

$$\Sigma_d \subset P_d \subset \mathbb{R}[\mathbf{u}]_d \quad (1)$$

Unlike P_d , Σ_d is computationally tractable—feasibility and optimization problems over Σ_d translate naturally in to SDPs. This will be made explicit in section 3.2.

3.2 SOS Optimization

Membership in Σ_d can be expressed via semidefinite programming. Using $\langle \cdot, \cdot \rangle$ to denote the Frobenius inner product, an SOS polynomial is equivalently written as

$$\begin{aligned} f(\mathbf{u}) &= \begin{pmatrix} s_1(\mathbf{u}) \\ \vdots \\ s_k(\mathbf{u}) \end{pmatrix}^\top \begin{pmatrix} s_1(\mathbf{u}) \\ \vdots \\ s_k(\mathbf{u}) \end{pmatrix} \\ &= [\mathbf{u}]_{\lfloor d/2 \rfloor}^\top \underbrace{\begin{pmatrix} s_1 & \cdots & s_k \end{pmatrix} \begin{pmatrix} s_1^\top \\ \vdots \\ s_k^\top \end{pmatrix}}_S [\mathbf{u}]_{\lfloor d/2 \rfloor} \quad (2) \\ &= \left\langle S, [\mathbf{u}]_{\lfloor d/2 \rfloor} [\mathbf{u}]_{\lfloor d/2 \rfloor}^\top \right\rangle, \end{aligned}$$

where the coefficients s_i are now encoded in the matrix S . (2) provides a linear relationship between f and S , which is positive semidefinite by construction. Indeed, we have shown that the existence of such an $S \geq 0$ is equivalent to $f \in \Sigma_d$. See Table 1 for rough dimensions of the SDP problem.

Using the above transformation to render SOS constraints into semidefinite constraints, one can perform optimization over $f(\mathbf{u}) \in \Sigma_d$ with relative ease. For this reason, it is often profitable to transform a problem involving a positivity constraint $f \in P$ to one with a set of constraints of the form $s_i \in \Sigma_d$ —resulting in an *SOS program*. A key theorem in SOS programming, the *Positivstellensatz*, explains how to effect this transformation so that the global optimum of the original problem is recovered in the limit of increasing SOS degree d [Blekherman et al. 2012]. We will use a specialized version of the Positivstellensatz below.

3.3 SOS Optimization on a Compact Domain

In the context of geometry processing, one frequently seeks to optimize a functional over a compact domain, rather than over all of $\mathbb{R}^k$. Thus, one often encounters constraints of the form $f(\mathbf{u}) \geq 0$ for $\mathbf{u} \in \mathbb{D}$ where $\mathbb{D}$ is compact. SOS programming can be extended to

Fig. 2. Quadratic triangle patch

handle such constraints. Here the key theorem is Putinar’s variant of the Positivstellensatz:

THEOREM 3.1 (PUTINAR’S POSITIVSTELLENSATZ [PUTINAR 1993]; SEE ALSO [BLEKHERMAN ET AL. 2012], THEOREM 3.138). *Let $\mathbb{D} = \{\mathbf{u} \in \mathbb{R}^k : g_i(\mathbf{u}) \geq 0\}$ be a domain with an algebraic certificate of compactness. Any polynomial $f(\mathbf{u})$ that is strictly positive on $\mathbb{D}$ admits a decomposition*

$$f(\mathbf{u}) = s_0(\mathbf{u}) + \sum_{i=1}^m s_i(\mathbf{u})g_i(\mathbf{u}), \quad (3)$$

with SOS polynomials $s_i \in \Sigma_d$ for high enough degree d .

Observe that the decomposition (3) provides a certificate of non-negativity by construction, since all s_i are nonnegative and all g_i are nonnegative on $\mathbb{D}$. For brevity we omit details of the required algebraic certificate of compactness and refer the interested reader to [Blekherman et al. 2012]. For all domains encountered in this paper, compactness certificates are readily computable. The theory of SOS programming has an elegant dual formulation that allows us to determine a sufficiently large d , but we will defer discussion of this dual to Section 5 to prioritize presenting concrete examples of what SOS programming can achieve for geometry processing.

Equality Constraints. It is often convenient to consider *semialgebraic* domains $\mathbb{D}$ defined by polynomial inequalities *and* equations. This generalization does not pose a significant problem for SOS methods, for the reasons we will outline below.

Given a domain defined by equations and inequalities $\mathbb{D} = \{\mathbf{u} : g_i(\mathbf{u}) \geq 0, h_i(\mathbf{u}) = 0\}$, a decomposition of the form

$$f(\mathbf{u}) = s_0(\mathbf{u}) + \sum_{i=1}^m s_i(\mathbf{u})g_i(\mathbf{u}) + \sum_{i=1}^n p_i(\mathbf{u})h_i(\mathbf{u}), \quad (4)$$

where the s_i are SOS and the p_i are arbitrary polynomials, certifies the nonnegativity of f on $\mathbb{D}$.

One can formally rewrite each equality constraint $h_i(\mathbf{u}) = 0$ as a pair of inequality constraints $h_i(\mathbf{u}) \geq 0$ and $-h_i(\mathbf{u}) \geq 0$, reducing $\mathbb{D}$ to the form required in Theorem 3.1. Then the theorem shows existence of a certificate of the form

$$f(\mathbf{u}) = s_0(\mathbf{u}) + \sum_{i=1}^m s_i(\mathbf{u})g_i(\mathbf{u}) + \sum_{i=1}^n (t_i^+(\mathbf{u}) - t_i^-(\mathbf{u}))h_i(\mathbf{u}), \quad (5)$$

where the s_i , t_i^+ , and t_i^- are SOS. But such a certificate is *a fortiori* of the form (4). In this way, Theorem 3.1 extends to the case of mixed constraints.

3.4 Polynomial Patches

Definition 3.1. A *polynomial patch* is a map $\mathbf{x} : \mathbb{D} \subset \mathbb{R}^k \rightarrow \mathbb{R}^n$ from a compact semialgebraic base domain $\mathbb{D} = \{\mathbf{u} \in \mathbb{R}^k : g_i(\mathbf{u}) \geq 0\}$, and such that each component of $\mathbf{x}$ is a multivariate polynomial of bounded degree $d_{\mathbf{x}}$.

In practice, the base domain is typically a canonical triangle, square, cube, or similar, and so the polynomials g_i are affine. The dimension of the patch for most geometry processing tasks is $k \in \{1, 2, 3\}$, and n controls the embedding dimension of the patch.

It is often useful to parameterize these polynomials $\mathbf{x}$ by linear combinations of a few basis functions, with coefficients that are obtained from control points $\mathbf{p}_i \in \mathbb{R}^n$. Patches of this form include Bézier and Coons patches, and they generalize the simplest polynomial patch, the *linear triangle*, which we will describe here for illustrative purposes. For the linear triangle patch, $\mathbb{D}$ is the triangle with vertices $(1, 0)$, $(0, 1)$, $(0, 0)$, $n = 3$, and the basis functions (in barycentric coordinates) are $\phi_1(\mathbf{u}) = u_1$, $\phi_2(\mathbf{u}) = u_2$, and $\phi_3(\mathbf{u}) = 1 - u_1 - u_2$. If the vertices of the linear triangle in $\mathbb{R}^3$ are $\mathbf{p}_1$, $\mathbf{p}_2$ and $\mathbf{p}_3$, then $\mathbf{x}(\mathbf{u}) = \phi_1\mathbf{p}_1 + \phi_2\mathbf{p}_2 + \phi_3\mathbf{p}_3$.

By varying the degree and number of basis functions, one can realize a variety of patch types. For example, a quadratic triangle is parameterized by six basis functions with control points $\mathbf{p}$ located at the vertices and edge midpoints. Such a quadratic triangle, along with its control points, is depicted in Figure 2. We will use n_B to denote the number of basis functions needed for a particular patch. For the purposes of this paper, we will focus our attention on quadratic and cubic triangles, quadratic Bézier curves, bicubic Bézier tensor patches, cubic Coons patches, and B-spline surfaces.

4 GEOMETRIC KERNEL PROBLEMS

In this paper, we show that a variety of geometric problems on polynomial patches can be solved with SOS programming. Our key observation is that most of these can be formulated in a very similar way. In particular, we offer the following template problem:

$$\begin{aligned} f^* &= \min_{\mathbf{u} \in \mathbb{D}} f(\mathbf{u}) \\ \mathbf{u}^* &= \arg \min_{\mathbf{u} \in \mathbb{D}} f(\mathbf{u}). \end{aligned} \quad (6)$$

where

$$\mathbb{D} = \{\mathbf{u} \in \mathbb{R}^k : g_i(\mathbf{u}) \geq 0, h_i(\mathbf{u}) = 0\}, \quad (7)$$

is a compact semialgebraic domain. In what follows, we will outline how to compute the globally optimal value f^* using the machinery of SOS programming, and we will apply this to a variety of patch problems. In Section 5, we will discuss how to compute $\mathbf{u}^*$ and determine the correct degree for the relaxation.

To apply SOS methods, we rewrite the problem in terms of an additional variable λ , which acts as a lower bound on $f(\mathbf{u})$ for $\mathbf{u} \in \mathbb{D}$. This is equivalent to asking that the polynomial $f(\mathbf{u}) - \lambda$ be positive for $\mathbf{u} \in \mathbb{D}$. Applying Theorem 3.1, such a requirement is equivalent to the SOS constraints

$$f - \lambda - \sum_i h_i p_i - \sum_i g_i s_i \in \Sigma, \quad s_i \in \Sigma, \quad p_i \in \mathbb{R}[\mathbf{u}]. \quad (8)$$

We have thus relaxed the problem (6) into the SOS form

$$\lambda^* = \left\{ \begin{array}{l} \max_{\lambda \in \mathbb{R}} \lambda \\ \text{s.t.} \quad f - \lambda - \sum_i h_i p_i - \sum_i g_i s_i \in \Sigma_d, \\ \quad s_i \in \Sigma_d \\ \quad p_i \in \mathbb{R}[\mathbf{u}]_d \end{array} \right\}. \quad (9)$$

As d increases, Theorem 3.1 guarantees that λ^* converges to the globally optimal value f^* . For high-enough d , the relaxation will be accurate to within numerical precision, and as we observe in Section 6, the convergence is dramatic at a fairly low d .

4.1 Optimization over a Polynomial Patch

The template problem (6) allows us to optimize arbitrary polynomial objective functions over a polynomial patch. One only needs to choose the g_i that encode a base domain $\mathbb{D}$ along with an appropriate objective function f . We now give several examples of objective functions of interest to geometry processing.

Closest Point (CP). The closest point problem aims to find the minimum distance between a target point $\mathbf{t}$ and a polynomial patch. Let $\mathbf{x}(\mathbf{u})$ be the shape function of the patch. Then one simply chooses the objective function $f(\mathbf{u}) = \|\mathbf{x}(\mathbf{u}) - \mathbf{t}\|_2^2$ in the template problem.

Minimal Axis Aligned Bounding Box (MBB). The axis aligned bounding box problem finds the smallest-volume AABB containing a polynomial patch $\mathbf{x}(\mathbf{u})$. The bounds can be obtained by choosing the objective function $f(\mathbf{u}) = \pm x_i(\mathbf{u})$.

Hexahedron Quality Evaluation. The hexahedral quality evaluation problem presented by Marschner et al. [2020] also fits into the template. Here $\mathbb{D}$ is a unit cube, $\mathbf{x}(\mathbf{u})$ encodes trilinear interpolation, and $f(\mathbf{u})$ is the Jacobian determinant of $\mathbf{x}(\mathbf{u})$, i.e., $f(\mathbf{u}) = \det(\nabla_{\mathbf{u}} \mathbf{x})$.

4.2 Multiple Patches in One Optimization

Often it is desirable to optimize an objective over multiple patches, or multiple copies of the same patch, simultaneously. Such problems can often be recast as optimization problems over a higher-dimensional base domain formed from the Cartesian product of the original base domains.

To wit, let polynomial patches $\mathbf{x}^1$, and $\mathbf{x}^2$ map respectively from base domains $\mathbb{D}_1 = \{\mathbf{u}^1 : g_i^1(\mathbf{u}^1) \geq 0\}$, and $\mathbb{D}_2 = \{\mathbf{u}^2 : g_i^2(\mathbf{u}^2) \geq 0\}$ into $\mathbb{R}^n$. Let

$$\begin{aligned} \mathbb{D} &= \{(\mathbf{u}^1, \mathbf{u}^2) : g_i^1(\mathbf{u}^1) \geq 0, g_i^2(\mathbf{u}^2) \geq 0, q_j(\mathbf{u}^1, \mathbf{u}^2) \geq 0\} \\ &\subseteq \mathbb{D}_1 \times \mathbb{D}_2, \end{aligned} \quad (10)$$

where the q_j are additional problem-dependent and/or symmetry-breaking constraints. One can then define a new patch $\mathbf{x}(\mathbf{u}^1, \mathbf{u}^2) = (\mathbf{x}^1(\mathbf{u}^1), \mathbf{x}^2(\mathbf{u}^2))$ and rewrite f relative to $\mathbf{x}$ and $\mathbb{D}$. For notational convenience, we will continue to use $\mathbf{u}^1$, $\mathbf{u}^2$, $\mathbf{x}^1$, and $\mathbf{x}^2$ directly.

Patch Diameter (PD). The *diameter* of a patch is the distance between the two most distant points on that patch. This can be found using the template problem with two identical base domains $\mathbb{D}_1 = \mathbb{D}_2$ and maps $\mathbf{x}^1 = \mathbf{x}^2$. Then one simply chooses the objective $f(\mathbf{u}^1, \mathbf{u}^2) = -\|\mathbf{x}^1(\mathbf{u}^1) - \mathbf{x}^2(\mathbf{u}^2)\|_2^2$.

Fig. 3. The CCD problem is demonstrated on two time-varying quadratic triangle patches over the time interval $[0, 1]$. The patches are shown at the time of first collision, at $t = 0.8377$, in the inset, with the collision point marked in red.

Surface-Surface Intersection (SSI). This problem detects when two polynomial patches intersect each other. As described previously, we use an augmented constraint set consisting of constraints from both patches. We further add the equality constraints $h(\mathbf{u}^1, \mathbf{u}^2) = \mathbf{x}^1(\mathbf{u}^1) - \mathbf{x}^2(\mathbf{u}^2) = 0$. This restricts optimization from the domain $\mathbb{D}_1 \times \mathbb{D}_2$ to its subset on which intersections occur. Feasibility of this problem determines if the two surfaces intersect. We can arbitrarily choose as an objective function $f(\mathbf{u}^1, \mathbf{u}^2) = u_1^1$.

Self-Intersection (SI). The self-intersection problem finds places where a single polynomial patch intersects itself. To fit this into the template problem, we start with the constraints of the SSI problem for two identical domains $\mathbb{D}_1 = \mathbb{D}_2$, $\mathbf{x}^1 = \mathbf{x}^2$. Then we choose the objective to be $f(\mathbf{u}^1, \mathbf{u}^2) = -\|\mathbf{u}^1 - \mathbf{u}^2\|_2^2$. This gives us points that are as far apart as possible in the base domain but map to the same embedded point in $\mathbb{R}^n$.

Continuous Collision Detection (CCD). The Continuous Collision Detection problem aims to find out if and when two patches on a time varying trajectory will intersect in a given time interval. Figure 3 shows an example of the CCD problem. Let two patches $\mathbf{x}^1, \mathbf{x}^2$ be defined by control points $\mathbf{p}_i^1, \mathbf{p}_i^2$ via $\mathbf{x}(\mathbf{u}) = \sum_i^{n_B} \mathbf{p}_i \phi_i$. Let control points have velocities v_i^1 and v_i^2 . We then define time-dependent patches by linear interpolation $\mathbf{x}(\mathbf{u}, t) = \sum_i^{n_B} (\mathbf{p}_i + \mathbf{v}_i t) \phi_i$, mapping from the augmented base domain $\mathbb{D}_1 \times \mathbb{D}_2 \times [0, t_{max}]$ to $\mathbb{R}^n$. We wish to determine if the provided velocities will result in a collision within t_{max} units of time.

As in the SSI case, we augment the constraint set with an equality constraint $h(\mathbf{u}^1, \mathbf{u}^2, t) = \mathbf{x}^1(\mathbf{u}^1, t) - \mathbf{x}^2(\mathbf{u}^2, t) = 0$. This restricts the optimization to the subset of the joint base domain on which space-time collisions occur. It just remains to choose $f(\mathbf{u}^1, \mathbf{u}^2, t) = t$ to find the earliest collision.

4.3 Outside the Template

Rational Surfaces. A natural generalization of polynomial surfaces are *rational surfaces*—surfaces whose shape functions $\mathbf{x}$ are rational. Using control points $\mathbf{p}_i$ and basis functions ϕ_i , rational surfaces are decomposable as follows

$$x_j(\mathbf{u}) = \frac{\sum_i p_{ij} \phi_i(\mathbf{u})}{b(\mathbf{u})} = \frac{a_j(\mathbf{u})}{b(\mathbf{u})}, \quad (11)$$

for polynomials $a_j, b, \phi_i \in \mathbb{R}[\mathbf{u}]$. Straightforward application of problem formulations from Section 4.1 and Section 4.2 to rational surfaces may seem not to fit the template (6). The key observation for applying SOS programming to rational surfaces is that we simply need to clear the denominators. We demonstrate here the SSI problem on rational surfaces, with all other problems following *mutatis mutandis*. The SSI formulation applied to rational surfaces starts with

$$\begin{aligned} \min_{\mathbf{u}^1 \in \mathbb{D}_1, \mathbf{u}^2 \in \mathbb{D}_2} \quad & u_1^1 \\ \text{s.t.} \quad & \frac{\mathbf{a}^1(\mathbf{u}^1)}{b^1(\mathbf{u}^1)} = \frac{\mathbf{a}^2(\mathbf{u}^2)}{b^2(\mathbf{u}^2)} \end{aligned} \quad (12)$$

By cross-multiplying denominators, the rational constraints are transformed back into polynomial constraints.

$$\begin{aligned} \min_{\mathbf{u}^1 \in \mathbb{D}_1, \mathbf{u}^2 \in \mathbb{D}_2} \quad & u_1^1 \\ \text{s.t.} \quad & \mathbf{a}^1(\mathbf{u}^1) b^2(\mathbf{u}^2) = \mathbf{a}^2(\mathbf{u}^2) b^1(\mathbf{u}^1) \end{aligned} \quad (13)$$

Thus problems that can be solved on polynomial patches via SOS programming can be extended to rational surfaces.

Minimal Enclosing Ellipsoid (MEE). The Minimal Enclosing Ellipsoid problem finds the ellipsoid of smallest volume that fully contains a polynomial patch. Smallest ellipsoids enclosing point sets have a long history in convex optimization, where they are known as *Löwner-John ellipsoids* [Todd 2016]. The MEE of a finite collection of points or ellipsoids can be computed exactly via SDP. This involves a reduction from the problem of computing the smallest ellipsoid with arbitrary center to the case of an ellipsoid centered at the origin. We will combine this reduction with SOS methods to compute the MEE of a polynomial patch.

We can parameterize an ellipsoid in $\mathbb{R}^3$ by a positive definite matrix $A \in \mathbb{S}_{++}^3$ and center point $\mathbf{c} \in \mathbb{R}^3$:

$$\mathcal{E}(A, \mathbf{c}) = \{\mathbf{p} \in \mathbb{R}^n : (\mathbf{p} - \mathbf{c})^\top A (\mathbf{p} - \mathbf{c}) \leq 1\}, \quad (14)$$

$$\text{Vol}(\mathcal{E}(A, \mathbf{c})) \propto \det(A^{-1}) = (\det A)^{-1}.$$

Given $\mathbb{D}$ and $\mathbf{x}$ describing a 3D polynomial patch, the MEE can be computed as the solution to

$$(A^*, \mathbf{c}^*) = \left\{ \begin{array}{ll} \underset{A \in \mathbb{S}_{++}^3, \mathbf{c} \in \mathbb{R}^3}{\text{argmin}} & \text{Vol}(\mathcal{E}(A, \mathbf{c})) \\ \text{s.t.} & 1 \geq \max_{\mathbf{u} \in \mathbb{D}} (\mathbf{x}(\mathbf{u}) - \mathbf{c})^\top A (\mathbf{x}(\mathbf{u}) - \mathbf{c}) \end{array} \right\} \quad (15)$$

The problem in this form does not clearly fit into the template due to the inner maximization, which can also be viewed as a universal quantification. But as we will now show, we can still apply SOS optimization to obtain the MEE. The first step is to reduce the general MEE problem (15) to an equivalent *centered* MEE problem in one higher dimension:

$$B^* = \left\{ \begin{array}{ll} \underset{B \in \mathbb{S}_{++}^4}{\text{argmin}} & \text{Vol}(\mathcal{E}(B, 0)) \\ \text{s.t.} & 1 \geq \max_{\mathbf{u} \in \mathbb{D}} \mathbf{y}^\top(\mathbf{u}) B \mathbf{y}(\mathbf{u}) \\ & \mathbf{y}(\mathbf{u}) = \begin{bmatrix} \mathbf{x}(\mathbf{u}) \\ 1 \end{bmatrix} \end{array} \right\} \quad (16)$$

From B^* , one can recover A^* by completing the square [Todd 2016]. The constraints of (16) require that the polynomial $1 - \mathbf{y}^\top B \mathbf{y} \in \mathbb{R}[\mathbf{u}]$ be nonnegative for all $\mathbf{u} \in \mathbb{D}$. Theorem 3.1 allows us to encode this requirement with SOS constraints, resulting in the following SOS program.

$$B^* = \left\{ \begin{array}{ll} \arg \min & \log(\text{Vol}(\mathcal{E}(B, 0))) \\ \text{s.t.} & s_i \in \Sigma, \quad s_i \in \mathbb{R}[\mathbf{u}]_d \\ & 1 - \mathbf{y}^\top B \mathbf{y} - \sum_i g_i s_i \in \Sigma \\ & B \in \mathcal{S}_{++}^4 \end{array} \right\}, \quad (17)$$

where we have also convexified the objective by taking its log. Both the objective and constraints are now convex, and we can solve (17) to global optimality, and by extension the MEE problem.

Our ability to solve the MEE problem hinges on the fact that MEE has an equivalent centered formulation wherein $\mathbf{c} = 0$. Consider what would happen if we tried to relax (15) directly. We would require $1 - (\mathbf{x} - \mathbf{c})^\top A (\mathbf{x} - \mathbf{c}) \geq 0$ for $\mathbf{u} \in \mathbb{D}$. Our constraints would contain bilinear terms like $\mathbf{x}^\top A \mathbf{c}$ and linear-quadratic terms like $\mathbf{c}^\top A \mathbf{c}$, thus breaking convexity. While our solution to the MEE problem does not generalize directly to other problems outside the template, it illustrates what can and cannot be handled by SOS programming.

Minimal Surrounding Sphere (MSS). The Minimal Surrounding Sphere problem finds the sphere of smallest volume that fully contains a polynomial patch. This can be solved with a minor modification to (17). Let B_3 be the top left 3×3 block of B . We simply add the constraint that $B_3 = aI_3$, where I_3 is the 3×3 identity matrix and a is a scalar variable. This additional constraint reduces the MEE program to an MSS program without affecting convexity.

5 MOMENT RELAXATION IN THEORY AND PRACTICE

While the previous section lists a variety of problems that can be solved with SOS programming, two points remain. First, while the SOS approach of the previous section gives us a way to compute the optimal value of (6) by way of (9), we do not have a way to compute the arg min of (6). Using CP as an example, we can find the distance from a target point to its projection on a patch, but do not have the projected point. Second, we have not mentioned how to certify correctness of a solution and determine if d is large enough for Theorem 3.1 to hold. This section will address both points.

5.1 Moment Relaxation in Theory

Consider an alternative formulation of (6) as a *measure relaxation*:

$$f^* := \min_{\mu \in \mathcal{P}(\mathbb{D})} \mathbb{E}_\mu[f], \quad (18)$$

where $\mathcal{P}(\mathbb{D})$ is the space of probability measures on $\mathbb{D}$, and $\mathbb{E}_\mu[f] = \int_{\mathbb{D}} f d\mu$ denotes integration against μ . If $\mathbf{u}^* = \arg \min_{\mathbf{u} \in \mathbb{D}} f(\mathbf{u})$, then the optimum of (18) is achieved by placing all the mass of μ at $\mathbf{u}^*$ —that is, the delta measure $\delta_{\mathbf{u}^*}$ minimizes (18) [Lasserre 2001]. While this relaxation convexifies (6), optimization over the infinite-dimensional space of measures $\mathcal{P}(\mathbb{D})$ is intractable. To realize the measure relaxation computationally, μ must be represented by its moments. Using the multi-index α to index monomial exponents

Fig. 4. For the SSI problem on quadratic triangles, all intersections are correctly identified for degree $d \geq 2$, and the distance between extracted points of intersection drops to $6.7\text{e-}05$ at degree $d = 5$, even though exact recovery does not yet occur at these degrees. Triangle control points were sampled from i.i.d. standard normal distributions.

of f , the α -moment of μ is defined as $\mu_\alpha := \mathbb{E}_\mu[\mathbf{u}^\alpha]$. Let f_α be the coefficient of $\mathbf{u}^\alpha$. Then the expectation in (18) can be expressed as

$$\mathbb{E}_\mu[f] = \sum_{\alpha} f_{\alpha} \mu_{\alpha}, \quad (19)$$

a simple linear function of *finitely many* moments. It remains to constrain μ_α to be valid moments of a distribution over $\mathbb{D}$, a task fulfilled by the following theorem:

THEOREM 5.1 ([LASSERRE 2001], THEOREM 4.2). *Suppose $\mathbb{D}$ is a compact domain with algebraic certificate of compactness. Let d be large enough for $f(\mathbf{u}) - f^*$ to admit the decomposition in (3). Then the following **moment relaxation** SDP computes f^* :*

$$f^* = \left\{ \begin{array}{ll} \min_{\mu \in \mathcal{M}_d} & \mathbb{E}_\mu[f(\mathbf{u})] \\ \text{s.t.} & \mathbb{E}_\mu[q(\mathbf{u})^2] \geq 0, \quad \forall q \in \mathbb{R}[\mathbf{u}]_{\lceil d/2 \rceil} \\ & \mathbb{E}_\mu[q(\mathbf{u})^2 g_i(\mathbf{u})] \geq 0, \quad \forall q \in \mathbb{R}[\mathbf{u}]_{\lceil d/2 \rceil - w_i} \\ & \mathbb{E}_\mu[1] = 1, \end{array} \right\} \quad (20)$$

where $w_i = \lceil \deg g_i / 2 \rceil$, and $\mathcal{M}_d$ is the space of moment vectors up to degree d . Furthermore, the moment vector corresponding to $\delta_{\mathbf{u}^*}$ minimizes (20).

We now have two SDP-based approaches to solving the same polynomial optimization problem—one derived by SOS programming, and one via moment relaxation. As the reader may have guessed, the two SDPs are dual to each other [Lasserre 2001].

If the moment vector μ solving (20) corresponds to a valid probability measure, then we know the relaxation was tight and we have achieved the global optimum of the original measure relaxation (18), and in turn of the polynomial problem (6). Verifying this when the optimal measure is a delta measure is particularly simple—this will

Table 1. Overview of SOS results on quadratic and cubic triangle patches. For each problem, we randomly generate 1000 test cases. Specifically for CCD on cubic triangles we test on a reduced set of 300 cases. We show the minimum SOS polynomial degree d_{correct} for which all problem instances are solved correctly. This degree is generally less than the degree d_{exact} required for exact recovery. For problems where it is applicable (SSI, SI, CCD), we provide the percentage of the data that intersects or collides. We also provide the number of optimization variables for each problem, which indicate the dimensions of the ultimately solved SDP problem. Additionally, we provide the median runtimes over 50 instances of each problem with standard deviations. We separate out the time spent in YALMIP converting to MOSEK’s input form, and the time spent in MOSEK actually solving the SDP problem. CCD on cubic triangles takes the longest, which is unsurprising given the large number of optimization variables. Lower-degree problems like MBB and MSS are already fast enough to be used on large meshes out of the box. *For SI, 100% exact recovery is taken out of cases where the patch did self-intersect. Exact recovery is not expected if the patch does not self-intersect.

	d_{correct}		% Correct		% Exact		% Intersects		# Vars		YALMIP Time (s)		MOSEK Time (s)	
	2	3	2	3	2	3	2	3	2	3	2	3	2	3
Triangle d_x	2	3	2	3	2	3	2	3	2	3	2	3	2	3
CP	3	5	100	100	100	100	-	-	31	64	.56±.04	.52±.04	.14±.02	.10±.08
MBB	2	4	100	100	100	100	-	-	19	46	.48±.02	.49±.02	.002±.0003	.004±.0007
PD	4	6	100	100	100	100	-	-	491	1471	.49±.02	.5±.02	.05±.009	1.24±.19
SSI	5	6	100	100	27	0	60	85	757	1261	.6±.02	.66±.08	.78±.12	6.19±1.4
SI	4	4	100	100	100*	0	51	84	491	491	.38±.02	.5±.03	.29±.06	.39±.09
CCD	5	6	100	100	1	0	82	90	2017	3697	.69±.04	1.31±.09	25.9±4.4	699±49
MSS	2	4	100	100	-	-	-	-	36	63	.41±.03	.38±.008	.006±.002	.012±.002
MEE	4	6	100	100	-	-	-	-	68	107	.42±.04	.37±.03	.13±.02	.13±.01

be the case if and only if the semidefinite matrix corresponding to the first constraint in (20) has rank one. This phenomenon is known as *exact recovery*, and it provides an *optimality certificate* verifying that d was chosen sufficiently high for the relaxed problem to solve the original problem (6). Finally, when we have exact recovery, $\mathbf{u}^*$ can be recovered as the mean (i.e., vector of first moments) of $\delta_{\mathbf{u}^*}$:

$$\mathbf{u}^* = (\mu_{\mathbf{u}_1}, \dots, \mu_{\mathbf{u}_k}) \quad (21)$$

Revisiting PD and SI. With (21) we are equipped to extract $\mathbf{u}^*$ for problems with a unique global minimum. Problems like PD, however, have multiple global minima due to the exchange symmetry $\mathbf{u}_1 \leftrightarrow \mathbf{u}_2$, which leaves the objective function unchanged. Such a symmetry means that the computed optimum μ , even if it corresponds to a valid measure, will not in general be a delta measure, making exact recovery elusive. Moreover, the mean extraction (21) will yield the Euclidean mean of optimal values, which may not itself be optimal.

To address these issues, we break the exchange symmetry by adding the generic constraint

$$g(\mathbf{u}_1, \mathbf{u}_2) = (\mathbf{u}_1 - \mathbf{u}_2) \cdot \vec{v} \geq 0, \quad (22)$$

where $\vec{v}$ is a randomly sampled unit vector. With this modification, generic uniqueness of the arg min is restored, allowing for exact recovery.

SI has the same symmetry and requires the same symmetry-breaking constraint. Furthermore, exact recovery is only possible for SI if the patch has a self-intersection. Otherwise, any pair of points $\mathbf{u}_1 = \mathbf{u}_2 \in \mathbb{D}$ are globally optimal with an optimal value of 0. Despite the possible lack of exact recovery, we address in Section 5.2 how to solve problems like SI reliably in practice.

5.2 Moment Relaxation in Practice

A general algorithm to achieve exact recovery might be to incrementally increase the SOS degree d from 1 until exact recovery is achieved. While this strategy may succeed eventually, Marschner

et al. [2020] observe that moment relaxations of a specific problem class tend to achieve exact recovery at the same degree, independent of the specific numbers involved. In particular, they observe that for a trilinear hexahedron, the hexahedron validity problem consistently achieves exact recovery with $d = 4$. Given a problem type, one might determine d empirically by generating many randomized instances of the same problem and increasing d until these generic instances can be solved to exact recovery. Then the empirically-determined degree can be applied to untested instances of the same problem.

Curiously, we observe that it is frequently possible to obtain f^* and $\mathbf{u}^* \in \mathbb{D}$ even without exact recovery. When d is smaller than required for exact recovery, the moment vector μ does not encode a delta measure. Nevertheless, we can take its mean following (21) and treat that as a proxy for $\mathbf{u}^*$. While this leaves the realm of formal SOS theory, we find in practice that this strategy is extraordinarily robust. Consider the SSI problem on quadratic triangles in Table 1. Choosing degree $d = 5$ provides us with exact recovery for only 27% of the instances of this problem. However, we can correctly detect intersection 100% of the time, and our extracted $\mathbf{u}^*$ gives us the correct intersection points 100% of the time. The same phenomenon is observed for SI and CCD—exact recovery is not required to obtain the correct solution.

Thus we distinguish the degree d_{exact} at which exact recovery is achieved from the degree d_{correct} at which a correct solution is achieved. In problems like SSI where numerical verification of intersection is easy given the candidate points of intersection, having an accurate $\mathbf{u}^*$ is just as effective a certificate of intersection as exact recovery. Since the cost of solving an SOS program is primarily dictated by d , it is extremely fortunate that for practical purposes one does not require exact recovery.

We demonstrate in Figure 4 the percentage of correctly solved SSI problems out of 200 for each of $d \in 1, 2, 3, 4, 5$ using randomly sampled quadratic triangles. Our sampling strategy is described in

Fig. 5. Conversion from (9) to functioning YALMIP code for CP. The YALMIP code matches the mathematical formulation almost line to line. After specifying the primal SOS problem, the dual moment vector and $\mathbf{u}^*$ are easy to extract. To solve other problems on different patches, only g_i , `trimapX` and f need to be changed.

Section 6.2. 128 of these problems have intersections. Out of problems with intersections, we plot the percentage of them that exhibit exact recovery for each degree. This percentage never reaches 100% and our key observation is that it does not need to! We do not require exact recovery to obtain the correct solution. We plot the percentage of problems where the SOS formulation misidentifies whether an intersection exists. At degree only $d = 2$, the percentage misidentified is already zero. Thus we are already able to determine whether a pair of quadratic triangles intersect with 100% accuracy. Finally, for problems that have intersections, using (21) we can extract the points on each surface that intersect. When the SSI problem is correctly solved, their distance must be 0. We plot the maximum distance between the extracted points over all intersecting examples at each degree. At $d = 5$, the points of intersection are identified with high precision. While these problem instances do not exhaustively cover SSI, they provide very strong empirical evidence that exact recovery is not required.

6 RESULTS AND APPLICATIONS

6.1 Implementation

Modern problem modeling languages such as YALMIP [Löfberg 2004] allow one to specify an SOS program in the concise form (9), automatically converting it to a primal-dual SDP formulation from which one can obtain the moment vector μ of (20). Figure 5 demonstrates how easy it is to specify and solve the primal and dual problems. The details of converting SOS or moment constraints into an SDP can be entirely left to the modeling language.

All batch experiments were run on an Intel i7-8700K CPU @ 3.70GHz with 16 GB of RAM. Problems were all written in MATLAB 2021a using YALMIP version 20200116 to formulate the problems

and MOSEK 9.2 to solve. Default modeling and solver parameters were maintained. Table 1 lists the median runtimes and standard deviations for 50 random instances of each problem for quadratic and cubic triangles. Time spent in YALMIP converting the problem into a form that MOSEK can handle does not vary significantly across problems or patch types. Time spent actually solving the problem in MOSEK varies much more. Clear trends are that solver time increases for each problem with SOS degree, which increases with patch degree. While CCD seems to be a significantly challenging problem, MBB is fairly easy. For most problems, the table shows that the majority of the time is spent in YALMIP. This time can be cut out entirely by formulating the problem directly as an SDP problem, or by taking advantage of YALMIP’s ability to precompile problems. We expect the runtime can be further improved by engineering tailor made optimization strategies for these problems.

6.2 SOS Problems Overview

A key parameter of our SOS formulations is the degree d . Since it controls the size of the SDP in (2), we naturally want to choose a small d . However, if d is too small, the SDP may not successfully solve the unrelaxed problem. In Table 1, we show the minimal degree d_{correct} required to solve each problem for 100% of our test cases on various triangular patches. As mentioned in Section 5.2, this is not necessarily the same degree d_{exact} needed to achieve exact recovery.

For each problem, 1000 test instances were generated, except for CCD on cubic triangles, for which we used 300 instances. For CP, MBB, PD, SI, SSI, MSS, and MEE, we randomly sampled quadratic and cubic triangles by picking all dimensions of all control points independently according to a normal distribution with variance 1. For CCD, pairs of triangles are sampled the same as before, but then shifted by $(0, 0, \frac{\sqrt{2}}{2})$ and $(0, 0, -\frac{\sqrt{2}}{2})$ respectively. Velocity vectors are

sampled the same way, but shifted by $(0, 0, -\frac{\sqrt{2}}{2})$ and $(0, 0, \frac{\sqrt{2}}{2})$ respectively. This is done to increase the number of problem instances that successfully collide without starting out in an intersecting configuration. With this sampling strategy, 82% (90%) of the sampled quadratic (cubic) triangles collide in spacetime while 22.5% (28%) of configurations start out intersecting. We verify correctness of the SOS solutions by comparing to the same problem solved on a linearization of the polynomial patch. Patches are uniformly subdivided so that each edge is split into 10 segments. We then check that solutions for the SOS solution and the linearization match up to a threshold of 10^{-2} in both the parametric and embedded spaces. In a few cases where they do not match, we increase density of the linearization until they do.

Table 1 demonstrates how challenging it can be to guess d_{correct} in advance. The SOS degrees we end up with are not clearly correlated with the degrees of the objective or constraints. The clearest trend is that d_{correct} increases with the degree of the shape function $d_{\mathbf{x}}$. There does not seem to be a monotonic relationship between d_{correct} and the degree of the objective function. For example, CP has an objective degree of $2d_{\mathbf{x}}$ while needing fairly low d_{correct} . On the other hand, CCD has an objective degree of just 1 but needs a much higher d . Even if we consider the constraint degrees as well, the degree of the CCD constraint is only $d_{\mathbf{x}} + 1$ which is still less than CP's objective degree $2d_{\mathbf{x}}$.

6.3 Closest Point

As described in 4.1, our method of formulation is applicable to the problem of finding the closest point on a polynomial patch to a target point.

Applying this method for all patches in a model lets us find the closest point on the model to a target. In the inset figure, we show this procedure applied to a collection of points surrounding the bicubic Bézier tensor patch teapot model, where $d_{\mathbf{x}} = 6$ and we chose $d = 5$. We can solve an analogous problem in one lower ambient dimension, finding the closest point to a Bézier curve in $\mathbb{R}^2$. This formulation allows for the creation of Voronoi diagrams of 2D objects formed from Bézier curves, as shown in Figure 6. In this figure, the 2D objects are letters consisting of quadratic Bézier curves, where $d_{\mathbf{x}} = 2$ and $d = 3$. We color points in the 2D domain by which curve segment they are closest to, forming curve-segment Voronoi cells.

Closest point queries allow us to build more complex operations on domains with curved boundaries. For example, the Monte Carlo walk-on-spheres procedure introduced to geometry processing in [Sawhney and Crane 2020] builds a stochastic solver for linear PDE out of closest point queries. Using the SOS formulation of CP as a building block, such a procedure can solve a PDE on a domain defined by polynomial bounding curves with exact boundary conformation, without linear remeshing. In the inset figure, we demonstrate walk-on-spheres applied to computing a

Fig. 6. Curve segment Voronoi diagrams for the glyphs H and N from [Smirnov et al. 2020b], each composed of 15 quadratic Bézier curves. On the left, we show the calculation of the closest point on each of the curve segments to the red point, with the distance to each curve represented by a circle. On the right, we show the Voronoi diagram, computed by finding the curve with the minimum distance to each point in a 500 by 500 grid around each letter.

boundary aligned cross-field. Each frame is computed by averaging complex fourth powers sampled from 100 random walks, each stopping when it arrives within 10^{-3} of the boundary. No mesh of the domain interior is required, and in principle, samples can be adapted to resolve singularities precisely.

6.4 Intersections

It is worth revisiting the intersection problems (SSI, SI, CCD) in Table 1. First, for SI on quadratic triangles, $d_{\text{correct}} = d_{\text{exact}}$. The same property does not extend to cubic triangles or SSI on quadratic triangles since the percentage that achieve exact recovery is different from the percentage that have intersection.

An interesting feature of SI is that the degree d_{correct} is the same for both quadratic and cubic triangles. This is surprising given that for other problems the cubic triangles require higher d_{correct} . We note here that CCD was almost the same in that 93% of cubic triangle CCD instances succeeded using $d = 5$ with runtime comparable to CCD for quadratic triangles. Only for the remaining percentage was $d = 6$ necessary. Due to the large time cost of CCD on cubic triangles with $d = 6$, one can imagine an optimization where $d = 5$ is used first to prune cases where a collision happens and is certifiably found. Only in the absence of this collision would one resort to $d = 6$.

In Figure 8 we demonstrate application of SI and SSI tests to airplane meshes from Smirnov et al. [2020a]. These models use cubic Coons patches for which $d_{\mathbf{x}} = 4$. The SI and SSI problems are

Fig. 7. We show the application of our CCD solution to the problem shown in (a), where the teapot moves rigidly into the elephant. The teapot and elephant models are comprised of 32 and 128 bicubic Bézier tensor patches, respectively. To accelerate the calculation of the collisions, we first use our MSS method to pre-compute a set of bounding spheres for each model, and use these to determine which pairs of patches may intersect. In (b) we color in red the bounding spheres of the elephant and bounding capsules of the teapot for which at least one intersection was found. We then solve the SOS patch CCD problem on the remaining candidate intersections. The results of this are shown in (c), where each patch is colored based on the time at which it first collides (with darker colors representing earlier times). We can also determine the exact moment in time when the two models first intersect, shown in (d).

solved with $d = 4$. We are able to find intersections for both SI and SSI and provide their exact locations. This enables users in a CAD pipeline to discover and manually fix problems in a design.

In Figure 7 we demonstrate application of the SOS relaxed CCD problem to detect collision of a rigid bicubic Bézier tensor patch teapot and elephant. Velocities were chosen to guarantee collision within 1 unit of time. This problem was solved with $d = 4$. Using CCD we are able to find the earliest instance of collision, the patches that collide, and the location of their collision. With the location of the collision we can get exact surface normals at the collision point, which can then be passed onto later steps in a simulation pipeline such as collision response.

Surface-Surface Intersections

6.5 Bounding Volumes

Bounding volumes are useful to speed up various intersection-type problems. They can be used in ray tracing to quickly detect if a ray will not intersect an object. They can also be used to quickly detect if rigid objects will not collide. The tighter a bounding volume is, the more non-collisions/intersections can be quickly pruned. On the other hand, employing a more complex bounding volume shape can result in increased computation and updating costs in cases when bounding volumes cannot be precomputed. As such, there is a tradeoff between the computation saved by having a tighter bounding volume and the cost of maintaining the bounding volume itself. SOS optimization is flexible enough to target multiple points along this tradeoff e.g. MBB, MSS, and MEE. In Figure 9 we compute the MBB, MSS, and MEE evaluated on bicubic Bézier tensor patches of the teapot mesh. As expected the total volume of MEEs is less than that of MSSs or that of MBBs. The MBB problem fits into the framework of Section 4.1, and so we can generically expect exact recovery for high enough d . The MSS and MEE problems are different in that the MSS (MEE) generically intersects the polynomial patch it bounds at multiple points. Thus we do not ever expect exact recovery. Nevertheless, we can obtain the minimal volume and the parameters of the MSS (MEE) achieving that volume from (17).

Self-Intersections

Fig. 8. Our method for finding self-intersections and surface-surface intersections applied to an invalid airplane mesh, obtained from the pipeline of [Smirnov et al. 2020a]. In the top row, we show on the left all the patches that intersect with another patch in red. On the right, two of those surface-surface intersections are shown, with a red point marking the intersection point calculated by our method. In the bottom right, all the patches that include self-intersections are shown in red. Two of these self-intersections are shown on the left, with the intersection points plotted in red.

Fig. 9. In the top row, we show a comparison of the MBB, MSS, and MEE for the bicubic Bézier tensor patch teapot. The total volume of bounding elements is computed for each type; ellipsoids provide the tightest bounding volume. The bottom row shows a comparison of the three different bounding types for an individual patch.

Fig. 10. The MBB and SSI problems solved on generic NURBS patches. On the left, two bounding boxes are drawn around the blue and green patches. On the right, the red point indicates where these patches intersect.

In Figure 7, we demonstrate the MSS on rigid bicubic Bézier tensor patches as a way to shorten computation time of CCD problems between meshes. Instead of running CCD between all pairs of patches, we can drastically cut down computation time by precomputing the MSS for each patch. It is straightforward to check if two moving spheres will collide or not, allowing us to reduce the number of CCD computations required.

6.6 Rational Surfaces

We demonstrate application of SSI to NURBS in Figure 11 on gears obtained from [Koch et al. 2019]. Intersecting teeth are easily detected and highlighted. We also demonstrate CCD in Figure 1 on the castle from [Trusty et al. 2021] to find which patches of the rocket will first collide with patches of the castle. These collisions could then be fed back into their elasticity simulation to render colliding elastic objects without need for linearization or volumetric re-meshing.

While handling of rational surfaces allows us to process generic NURBS surfaces, we find that many of the models in [Koch et al.

Fig. 11. This helical gear, from the ABC dataset [Koch et al. 2019], is modeled with NURBS surfaces defining its teeth, shown in blue in the left figure. To determine whether the two gears intersect, we can apply our SSI formulation to the NURBS surfaces in the two gears. To speed up the runtime, we first compute the bounding boxes of each patch with our MBB method and use these to identify patches which may intersect. The patches in the two gears that intersect are shown in red on the right.

2019] and [Trusty et al. 2021] have mainly patches with constant denominators $b = 1$, making them simply B-Splines. The primary exceptions in these datasets are patches that comprise spherical or cylindrical features. In Figure 10, we show MBB and SSI on generic rational surfaces without constant denominator computed using SOS programming.

7 DISCUSSION AND OPEN PROBLEMS

This paper presents the use of SOS programming tools to solve a variety of problems in geometry processing. The miraculous success of SOS for geometry processing—beyond the bounds specified by the SOS theory—leaves many open questions for theory to answer. We document these questions below to inspire additional research developing SOS theory relevant to our new problem domain.

7.1 Advancing SOS Methodology

Exact Recovery. A surprising aspect of our results is the robustness of the obtained solutions even when they do not come with exact recovery. This gap between theory and practice suggests formulating a modified notion of exact recovery where atomic measures are not required. It also suggests that there could be alternative formulations for the SSI, SI and CCD problems where exact recovery is found at the same degree that the correct solution is obtained.

For example, we relied heavily on Putinar’s Positivstellensatz Theorem 3.1 for our SOS formulations in this paper. While this form of the Positivstellensatz provides a concise form for certificates for compact domains, lower-degree certificates can sometimes be found by employing the full Positivstellensatz, which incorporates products of constraints [Blekherman et al. 2012].

Hardware. Despite the flexibility of SOS programming, the runtime cost of solving SDPs can be substantial. Similar to how the availability of high-performance GPUs catalyzed deep learning, we expect that design of specialized SDP hardware would enable the use of SOS programming in many more contexts where runtime was previously too large. At the least, individual kernel problems do not interact and so can be parallelized.

7.2 Exact Recovery in SI and CSI

Self Intersection with Tolerance. One might also look for a version of SI that generically has a unique optimum when no intersection is present. Recall that for SI we chose the objective $f(\mathbf{u}) = -\|\mathbf{u}_1 - \mathbf{u}_2\|_2^2$ to maximize distance in $\mathbb{D}$ when the SI problem only really cares about $\mathbf{u}_1 \neq \mathbf{u}_2$. This suggests an alternative formulation with constraint $g(\mathbf{u}_1, \mathbf{u}_2) = \|\mathbf{u}_1 - \mathbf{u}_2\|_2^2 - \epsilon \geq 0$, thus freeing the objective slot to be filled with something generic like $f(\mathbf{u}_1, \mathbf{u}_2) = (\mathbf{u}_1 - \mathbf{u}_2) \cdot \vec{v}$ for a random unit vector $\vec{v}$. Now feasibility of the problem decides if there is or is not an intersection, while exact recovery can be expected in either case. The cost of this formulation is of course the addition of the tuning parameter ϵ , and one might be left wondering if an undetectable self-intersection could have occurred in less than ϵ distance in $\mathbb{D}$.

Continuous Self-Intersection (CSI). A secondary benefit of reformulating SI with tolerance ϵ is that it would allow us to solve the CSI problem. Given control points and velocities of a single polynomial patch, one looks for the earliest time that a self-intersection occurs. Now that the objective slot is vacated, it can be used to minimize t similarly to CCD. The same drawback applies of course— ϵ must be chosen in advance.

7.3 Degenerate Instances and Generalized Exact Recovery

Our discussion of exact recovery has focused on the case where the optimal moments correspond to a delta measure, i.e., where the moment matrix has rank one. For some problems, optimal measures generically have more than one support point; this is the case for MSS and MEE, for which optimal measures are supported on points where the patch contacts the surface of the bounding sphere or ellipsoid. PD and SI would exhibit similar behavior if not for our symmetry-breaking modification. Other problems, such as CP, have *degenerate instances*, i.e., problem instances for which there are multiple global optima to the polynomial optimization problem. For these instances, optimal solutions to the measure relaxation are mixtures of these global optima.

In such cases, one might still be interested in extracting the support points. While more complicated than in the rank-one case, it is still possible to extract an *atomic measure*—a mixture of deltas—corresponding to a moment matrix when it satisfies the so-called *flat extension* property. This property holds when increasing the degree of the moment relaxation does not increase the rank of the moment matrix. When flat extension holds, one can apply a procedure based on diagonalization of a set of commuting multiplication operators to extract the measure support points [Blekherman et al. 2012, §3.5.6]. It would be interesting to test whether our MEE and MSS relaxations satisfy the flat extension property, which would serve as a certificate of generalized exact recovery and allow for extraction of the support points.

7.4 Optimization with Quantifiers

Certain problems are difficult to manipulate into the form of the template problem due to the presence of *quantifiers*. Quantifiers occur when the problem is more naturally formulated as a minimax or nested optimization problem. An example of this is MEE (15). While we managed to manipulate the MEE problem into a convex

form, there are similar problems with quantifiers for which convex relaxations are still unknown. One example is the hexahedron repair problem presented in [Marschner et al. 2020], where they aim to minimally perturb a trilinear map such that the map becomes injective. Their alternating optimization solution, while effective in practice, comes with no global optimality guarantees.

Continuous Collision Response (CCR). Another example of an optimization with quantifiers is CCR: given velocities $\hat{\mathbf{v}}_i^1$ and $\hat{\mathbf{v}}_i^2$ for two patches $\mathbf{x}_1$ and $\mathbf{x}_2$, find new velocities close to the input such that the new velocities do not result in collision in time t_{max} . Using the same notation as from CCD, we can write the CCR problem as the following.

$$v^* = \left\{ \begin{array}{l} \underset{\mathbf{v}_i^1, \mathbf{v}_i^2 \in \mathbb{R}^3}{\operatorname{argmin}} \quad \sum_{i=1}^{n_B} \|\mathbf{v}_i^1 - \hat{\mathbf{v}}_i^1\|^2 + \|\mathbf{v}_i^2 - \hat{\mathbf{v}}_i^2\|^2 \\ \text{s.t.} \quad \epsilon \leq \left\{ \begin{array}{l} \min_{\mathbf{u}^1, \mathbf{u}^2, t} \quad \|\mathbf{x}^1(\mathbf{u}^1, t) - \mathbf{x}^2(\mathbf{u}^2, t)\|_2^2 \\ \text{s.t.} \quad t \in [0, t_{max}] \\ \mathbf{u}^1 \in \mathbb{D}_1 \\ \mathbf{u}^2 \in \mathbb{D}_2 \end{array} \right. \end{array} \right\} \quad (23)$$

It could be that a clever manipulation of (23) convexifies it just as we did for MEE.

8 CONCLUSION

Many problems in geometry processing suffer from nonlinearity, forcing users to cope with only locally optimal solutions. This nonlinearity makes it challenging to leave the realm of piecewise linear geometry, where at least low-level operations such as closest-point queries, injectivity testing, intersection, and collision detection can be solved with confidence, a solid base facilitating development of more complex algorithms.

With our SOS framework, these and other low level operations are readily extended to higher-order surfaces. SOS geometry processing alleviates concerns, for example, that curved surface continuous collision detection might miss collisions due to linearization error, leading to unrealizable states. Similarly, SOS-based closest projection does not suffer from local optima that could affect processes down the line. SOS programming transforms these and other problems on a huge variety of curved patch types into problems where a user can confidently certify a global optimum.

SOS programming suggests a broadly-applicable framework for development of geometry processing algorithms. To fill in the details of this new approach, Section 7 presents a variety of concrete problems for future research at the intersection of SOS programming and geometry processing. At a coarser level, our work serves as a stepping stone toward design of flexible simulation tools for curved surface representations. Recent work by Trusty et al. [2021] also progresses in this direction but lacks modules for exactly resolving intersection and collision. Our tools enable confident collision detection operations analogous to the ones applied to piecewise-linear surfaces and can readily be incorporated into simulation algorithms. This opens the way toward collision response tools for a variety of geometric representations, including the velocity filter in Section 7.4.

SOS geometry processing also supports interactions between multiple element types. While many geometry processing and simulation tools expect a unified set of, e.g., only triangle or hex elements, one can use SOS tools to formulate—for example—colliding a NURBS surface with a quadratic Bézier triangle mesh. Another interesting extension would be to NURBS patches with trimming curves, which could be implemented with more general domain constraint polynomials g_i along with the denominator-clearing method from Section 4.3. One only needs to describe the element types as polynomials to the SOS framework. Beyond developing numerical tools, a programming language-inspired approach to geometry processing might consider what level of abstraction a language needs to expose this flexibility to a user.

These opportunities for further development aside, our broad SOS framework and specific model problems are already beneficial to geometry processing. These problems would otherwise each require their own solutions, including heuristics for number of initial points, density of linearization, tolerance for intersection, gradient step size, and many other nonlinear optimization parameters, all of which must be tuned per surface type. SOS instead provides a single unified framework for common objectives across the most popular surface representations.

ACKNOWLEDGMENTS

The authors would like to thank Diego Cifuentes for much help understanding SOS theory and Misha Bessmeltsev, Etienne Vouga, Ilya Baran, and Paul Stallings for valuable suggestions and discussion. We would also like to thank Dmitriy Smirnov and Ty Trusty for providing several models used in the figures.

Paul Zhang acknowledges the generous support of the DOE Computational Science Graduate Fellowship. David Palmer acknowledges the generous support of the Hertz Fellowship and MathWorks Fellowship.

The MIT Geometric Data Processing group acknowledges the generous support of ARO grant W911NF2010168, of AFOSR award FA9550-19-1-031, of NSF grants IIS-1838071 and CHS-1955697, from the CSAIL Systems that Learn program, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, from a gift from Adobe Systems, from an MIT.nano Immersion Lab/NC-SOFT Gaming Program seed grant, and from the Skoltech-MIT Next Generation Program.

REFERENCES

- Amir Ahmadi, Georgina Hall, Ameesh Makadia, and Vikas Sindhwani. 2017. Geometry of 3D Environments and Sum of Squares Polynomials. <https://doi.org/10.15607/RSS.2017.XIII.071>
- Farid Alizadeh. 1995. Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM Journal on Optimization* 5, 1 (1995), 13–51. <https://doi.org/10.1137/0805002>
- Robert E. Barnhill, Gerald Farin, M. Jordan, and Bruce R. Piper. 1987. Surface/surface intersection. *Computer Aided Geometric Design* 4, 1-2 (1987), 3–16. [https://doi.org/10.1016/0167-8396\(87\)90020-3](https://doi.org/10.1016/0167-8396(87)90020-3)
- Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas. 2012. *Semidefinite Optimization and Convex Algebraic Geometry*. SIAM. <https://doi.org/10.1137/1.9781611972290>
- Stephen Boyd and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511804441>
- David Cardoze, Alexandre Cunha, Gary L. Miller, Todd Phillips, and Noel Walkington. 2004. A bezier-based approach to unstructured moving meshes. In *Proceedings of the twentieth annual symposium on Computational geometry*. 310–319. <https://doi.org/10.1145/997817.997864>
- Edwin Catmull and James Clark. 1978. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-aided design* 10, 6 (1978), 350–355. <https://doi.org/10.1145/280811.280992>
- Zhongshi Jiang, Ziyi Zhang, Yixin Hu, Teseo Schneider, Denis Zorin, and Daniele Panozzo. 2021. Bijective and Coarse High-Order Tetrahedral Meshes. *ACM Transactions on Graphics* (2021). <https://doi.org/10.1145/3450626.3459840>
- Kęstutis Karčiauskas and Jörg Peters. 2020. Low degree splines for locally quad-dominant meshes. *Computer Aided Geometric Design* 83 (2020), 101934. <https://doi.org/10.1016/j.cagd.2020.101934>
- Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. 2019. ABC: A big cad model dataset for geometric deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9601–9611.
- Jean B. Lasserre. 2001. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization* 11, 3 (2001), 796–817. <https://doi.org/10.1137/S1052623400366802>
- Johan Löfberg. 2004. YALMIP: A Toolbox for Modeling and Optimization in MATLAB. In *In Proceedings of the CACSD Conference*. Taipei, Taiwan. <https://doi.org/10.1109/CACSD.2004.1393890>
- Charles Loop. 1987. Smooth subdivision surfaces based on triangles. *Master's thesis, University of Utah, Department of Mathematics* (1987).
- Manish Mandat and Marcel Campen. 2020. Bézier guarding: precise higher-order meshing of curved 2D domains. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 103–1. <https://doi.org/10.1145/3386569.3392372>
- Zoë Marschner, David Palmer, Paul Zhang, and Justin Solomon. 2020. Hexahedral Mesh Repair via Sum-of-Squares Relaxation. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 133–147. <https://doi.org/doi.org/10.1111/cgf.14074>
- MOSEK ApS. 2017. *The MOSEK optimization toolbox for MATLAB manual. Version 8.1*. <http://docs.mosek.com/8.1/toolbox/index.html>
- Yurii Nesterov and Arkadii Nemirovskii. 1994. *Interior-Point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611970791>
- Young-Taek Oh, Yong-Joon Kim, Jieun Lee, Myung-Soo Kim, and Gershon Elber. 2012. Efficient point-projection to freeform curves and surfaces. *Computer Aided Geometric Design* 29, 5 (2012), 242–254. <https://doi.org/10.1016/j.cagd.2011.04.002>
- Pablo A. Parrilo. 2019. Algebraic Techniques and Semidefinite Optimization. <https://learning-modules.mit.edu/materials/index.html?uuiid=course/6/sp19/6.256#materials>
- Diana Pekerman, Gershon Elber, and Myung-Soo Kim. 2008. Self-intersection detection and elimination in freeform curves and surfaces. *Computer-Aided Design* 40, 2 (2008), 150–159. <https://doi.org/10.1016/j.cad.2007.10.004>
- Stephen Prajna, Antonis Papachristodoulou, and Pablo A. Parrilo. 2002. Introducing SOSTOOLS: A general purpose sum of squares programming solver. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, Vol. 1. IEEE, 741–746.
- Mihai Putinar. 1993. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal* 42, 3 (1993), 969–984.
- Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Trans. Graph.* 39, 4 (2020). <https://doi.org/10.1145/3386569.3392374>
- Teseo Schneider, Jérémie Dumas, Xifeng Gao, Mario Botsch, Daniele Panozzo, and Denis Zorin. 2019. Poly-Spline Finite-Element Method. *ACM Trans. Graph.* 38, 3, Article 19 (March 2019), 16 pages. <https://doi.org/10.1145/3313797>
- Teseo Schneider, Yixin Hu, Jérémie Dumas, Xifeng Gao, Daniele Panozzo, and Denis Zorin. 2018. Decoupling Simulation Accuracy from Mesh Quality. *ACM Transactions on Graphics* 37, 6 (10 2018). <https://doi.org/10.1145/3272127.3275067>
- Dmitriy Smirnov, Mikhail Bessmeltsev, and Justin Solomon. 2020a. Learning Manifold Patch-Based Representations of Man-Made Shapes, In *International Conference on Learning Representations*. *arXiv preprint arXiv:1906.12337*.
- Dmitriy Smirnov, Matthew Fisher, Vladimir G. Kim, Richard Zhang, and Justin Solomon. 2020b. Deep parametric shape predictions using distance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 561–570. <https://doi.org/10.1109/CVPR42600.2020.00064>
- Michael J Todd. 2016. *Minimum-volume ellipsoids: Theory and algorithms*. SIAM. <https://doi.org/10.1137/1.9781611974386>
- Kim-Chuan Toh, Michael J. Todd, and Reha H. Tütüncü. 2001. SDPT3—a Matlab software package for semidefinite-quadratic-linear programming, version 3.0. *Web page* <http://www.math.nus.edu.sg/mattohkc/sdpt3.html> (2001). <https://doi.org/10.1080/10556789908805762>
- Ty Trusty, Honglin Chen, and David I.W. Levin. 2021. The Shape Matching Element Method: Direct Animation of Curved Surface Models. *ACM Transactions on Graphics* (2021). <https://doi.org/10.1145/3450626.3459772>
- Anna Wawrzinek, Klaus Hildebrandt, and Konrad Polthier. 2011. Koiter's Thin Shells on Catmull-Clark Limit Surfaces. In *Vision, Modeling, and Visualization (2011)*, Peter Eisert, Joachim Hornegger, and Konrad Polthier (Eds.). The Eurographics Association. <https://doi.org/10.2312/PE/VMV/VMV11/113-120>