

Haptic Assembly Using Skeletal Densities and Fourier Transforms*

Morad Behandish

Computational Design Laboratory
Department of Mechanical Engineering
University of Connecticut
Storrs, Connecticut 06269
Email: m.behandish@enr.uconn.edu

Horea T. Ilieş

Computational Design Laboratory
Department of Mechanical Engineering
University of Connecticut
Storrs, Connecticut 06269
Email: ilies@enr.uconn.edu

ABSTRACT

Haptic-assisted virtual assembly and prototyping has seen significant attention over the past two decades. However, in spite of the appealing prospects, its adoption has been slower than expected. We identify the main roadblocks as the inherent geometric complexities faced when assembling objects of arbitrary shape, and the computation time limitation imposed by the notorious 1 kHz haptic refresh rate. We addressed the first problem in a recent work by introducing a generic energy model for geometric guidance and constraints between features of arbitrary shape. In the present work, we address the second challenge by leveraging Fourier transforms to compute the constraint forces and torques. Our new concept of ‘geometric energy’ field is computed automatically from a cross-correlation of ‘skeletal densities’ in the frequency domain, and serves as a generalization of the manually specified virtual fixtures or heuristically identified mating constraints proposed in the literature. The formulation of the energy field as a convolution enables efficient computation using fast Fourier transforms (FFT) on the graphics processing unit (GPU). We show that our method is effective for low-clearance assembly of objects of arbitrary geometric and syntactic complexity.

1 Introduction

An integration of virtual reality (VR) tools into the modern computer-aided design (CAD) environments can facilitate an early-stage examination of a variety of product life-cycle aspects related to design, manufacturing, maintenance, service, and recycling. The so-called ‘virtual prototyping’ [2–4] results in a significant reduction of time and cost associated with ‘physical prototyping’, provides valuable insight into the functionality of the products, and facilitates the elimination of a large subset of design problems in the earlier stages of the process [5]. ‘Virtual assembly’, defined as a simulated assembly of the virtual representations of mechanical parts in an immersive 3D user interface using natural human motions, characterizes an important subset of virtual prototyping [6, 7], to which applying haptic feedback has been shown particularly beneficial in terms of task efficiency and user satisfaction [8, 9].

In the past two decades, there have been numerous studies focused on virtual assembly using a variety of visualization tools (e.g., stereoscopic displays and goggles) and tracking devices (e.g., head tracking devices and data gloves) to assist the user in the assembly tasks. More recently, an increasing number of studies have leveraged haptic devices to provide a more realistic assembly experience with force feedback, a thorough survey of which is beyond the scope of this paper. We refer the reader to [10] for an extensive review of previous studies, and to [11, 12] for recent insight on current knowledge and expected future directions in haptic assembly. Here we restrict ourselves to briefly overview the state-of-the-art practices

*A short version of this article was presented in the ASME IDETC/CIE’2015 Conferences [1].

and major recent contributions in Section 1.1, upon which we draw our newest approach for effective and efficient haptic assembly and disassembly.

Even after two decades replete with hundreds of important studies, the current models for real-time haptic-assisted assembly are still quite simplistic compared to the static and dynamic behavior of the objects in a real assembly process [12–14]. Even the most recent implementations are limited in number, complexity, and precision of assembly components. We aim to move on from few, simple, and high-clearance peg-in-hole examples to practical assembly scenarios with objects of complex geometric features and tight fits. This is hardly possible without revisiting the mathematical foundations for representing mechanical models and leveraging optimal algorithms to enable fast computations to comply with the 1 kHz requirement.

Among the major theoretical developments over the recent years in geometric computing are analytic methods [15–17]. In contrast to their combinatorial counterparts that use tools from discrete geometry, analytic techniques take advantage of functions theory, convolution algebras, and harmonic analysis to solve fundamental problems related to detecting collisions, similarity, complementarity, and symmetry. The metrics for quantifying such notions are formulated as convolutions of shape descriptor functions. There are several important advantages associated with analytic modeling, namely:

It enables formulating and solving complex problems in a uniform fashion, whose treatments with classical methods are theoretically and computationally challenging. An important example is computing Minkowski operations [18] on objects bounded by complex surfaces, which underlies swept volume generation [19] and collision detection [16], both of which are of prime importance in assembly planning (Section 2.2).

It provides insight into systematic extension of the existing solutions to more difficult problems, e.g., formulating shape complementarity metrics [20] as a generalization of the collision predicate [16] by using more sophisticated shape descriptors (Section 2.3).

It makes it possible to speed up the real-time computations to sub-millisecond times (required for haptic applications with 1 kHz frame rate) by using truncated Fourier transforms (Section 2.4).

By working in the Fourier domain, the computational resources (i.e., time and memory) that are allocated to resolving part behavior in each frame can be altered on-the-fly by low-pass filtering, regardless of syntactic complexity (e.g., polygon count).

We aim to extend the application of analytic methods—whose use has been limited for the most part to robotics applications [21, 22]—to both collision detection and geometric guidance for haptic assembly.

1.1 Related Work

Over the past few years, a ‘two-phase’ approach to haptic assembly has become popular in several different forms. The process is divided into a ‘free motion’ phase in the presence of collision-induced physical constraints, and a ‘fine insertion’ phase using pre-specified or computer-predicted geometric constraints [11]. Often the former is referred to as physically-based modeling (PBM), while the latter is called constraint-based modeling (CBM). The dynamic ‘part behavior’ (mostly for PBM) is simulated by integrating the equations of motion in real-time using physics simulation engines (PSE) [23, 24] which provide a variety of functionalities (including collision response, impact/friction mechanics, and non-smooth Lagrangian dynamics [25, 26]). The free motion phase typically relies on collision detection (CD) (reviewed in [27–29]) to constrain the motion with the so-called ‘physical constraints’ (i.e., unilateral holonomic constraints) that automatically arise from the collision response, and to compute the necessary impact/contact forces to enforce those constraints [30, 31], along with friction models [32, 33]. The fine insertion phase, on the other hand, is handled via constraint management libraries [34, 35] to implement additional ‘geometric constraints’ (i.e., bilateral degree of freedom (DOF)-limiting constraints) that are manually prescribed or heuristically predicted based on simple geometric semantics [36–38].

Among the combinatorial CD methods that are popular in haptic assembly—due to their computational efficiency in the presence of the 1 kHz requirement—are Voronoi-clipping/marching methods (e.g., V-Clip [39], SWIFT [40], and SWIFT++ [41]), oriented bounding box (OBB) tree-based methods (e.g., H-COLLIDE [42] and RAPID [43]), and volumetric enumeration methods (e.g., the well-known Voxmap PointShell (VPS) [44, 45] used in the earlier versions of SHARP [6, 7], and its improved variants [46, 47]). More efficient hierarchical data structures have been developed using spherical primitives (e.g., hierarchical bounding sphere (HBS) tree-based methods [48, 49] and inner sphere tree (IST) methods [50, 51] successfully applied to haptic rendering [52–54]). By leveraging the spherical symmetry in primitive collision predicates, sphere trees are generally more efficient alternatives to OBB trees, and represent nonuniform extensions to the uniform volumetric enumeration in VPS at every layer of the hierarchy. Although they have been shown to outperform VPS for nonconvex moving objects, their effectiveness to handle thin objects is yet to be tested [12].

In a separate line of research, analytic CD methods have been in use for decades in the robotics field for path planning in the presence of obstacles [21, 22]. Unlike combinatorial methods that search for a collision certificate point (or lack thereof) in the intersection of the objects and recover a gradient-like quantity to evaluate the collision response, analytic methods formulate the collision constraint as a convolution of the objects’ defining functions whose differentiation gives the

collision response [16], both of which convert to simple algebraic operations in the Fourier domain. However, most FFT-based convolution methods have been presented for cumulative computation of the collision predicate for all configurations, which is more than what one needs in real-time VR applications with a single pose at question during each frame [16]. In a recent work [16]—which has been significantly influential on our development in Section 2—the analytic CD was adapted to real-time applications along with explicit equations presented for gradient computations. Later, we proposed a hybrid combinatorial/analytic method in [17] that used spherical decompositions similar to [48–51] to discretize the analytic equations, and demonstrated that sub-millisecond running times can be achieved with analytic CD regardless of input shape complexity. In spite of their great promise, the nascent real-time analytic methods are yet to be tested in VR applications, and to the best of our knowledge has not been implemented into PSEs.

Regardless of which CD method is used, PSE+CD alone has been found inadequate for low-clearance haptic assembly [6, 7], partly due to the approximate nature of most CD methods (alleviated if exact boundary representations (B-reps) are used at the expense of computational performance, e.g., via the collision detection manager (CDM) module of D-Cubed in SHARP [55, 56]) and partly due to the noise in input data resulting from authentic hand vibration or added device errors [10]. This naturally led to the introduction of ad hoc solutions in a separate precision assembly phase, where a simplified set of ‘mating constraints’ are explicitly introduced in close proximity of the final intended assembly configuration. The assembly constraints can be extracted from the CAD model (e.g., from ProE models in VADE [57] and MIVAS [58]) or specified on-the-fly (e.g., via the dimensional constraint manager (DCM) module of D-Cubed in SHARP [55, 56]). Rather than using the geometric semantics of the original parts, the virtual constraint guidance (VCG) method given in [59] relied on *manually* specified abstract geometric constructs referred to as ‘virtual fixtures’ [60]—e.g., a pair of perpendicular planes intersecting at the axis of a cylindrical hole, to constrain and guide two points selected along the axis of a cylindrical peg. Alternatively, the automatic geometric constraints (AGC) method investigated in [56] and similar but independent studies in [36–38] attempted to *automatically* identify the assembly intent and associated geometric constraints by matching ‘functional surfaces’ [37]—e.g., a cylindrical surface characterized by its axis and diameter, which could be used to predict the intended mating relation and associated trajectories when a peg is brought to the proximity of a hole.

Up until recently, the two-phase framework has been widely accepted as the only promising future direction for performing haptic assembly tasks effectively [11]. Among the main difficulties with this hybrid paradigm are detecting the proximity to an insertion site, properly switching between the two phases (i.e., ‘blending’ algorithms to turn CD off and activate constraint management [56]), and handling collision events outside the insertion site when CD is switched off for insertion [12]. Furthermore, the ad hoc constraint management solutions are heavily dependent on user input based on a priori knowledge of the type of contact surfaces, which are in turn limited to simple categorized classes of geometric features—e.g., coplanarity of prismatic mates, coaxiality of cylindrical joints, and alike. In a similar classification to the one stated earlier for CD methods, the existing constraint identification algorithms can be identified as combinatorial methods that scale in running time with the complexity of the input CAD models, and look for certain heuristic indicators of a match between complementary features.

We recently introduced an alternative approach in [14] (a short version of which appeared in [13]), which could be viewed as a unified analytic approach to computing predicates for collision response and geometric guidance. We defined a generic metric as a convolution that penalizes collisions and separations between objects, while it rewards shape complementarity of their mating features, resulting in an artificial geometric energy function that can be differentiated for force and torque computations. As with most analytic approaches, our method generalizes to arbitrary geometry, without any simplifying assumption on the type or complexity of the mating features. Furthermore, it completely avoids the difficulties associated with the two-phase approach, along with the need for switching/blending between two distinct modes. We demonstrated the effectiveness of the method for simple peg-in-hole benchmark examples, and tested haptic feedback in a 3 DOF setup. However, the mathematical complexity of the model in terms of volumetric integrals in the physical domain required time-consuming preprocessing steps as the geometric details of the parts and the number of DOF grew, introducing new challenges that we address in the present work by leveraging Fourier transforms.

1.2 Contributions

We propose a paradigm that streamlines assembly simulation using ideas from multivariate harmonic analysis [61]. The development expands upon our analytical formulation of the ‘geometric energy’ field in [13, 14], defined as a convolution of the skeletal density functions (SDF) [20]. The method subsumes analytic CD, and provides a generalization to analytic feature matching for geometric guidance. The SDF shape descriptors are piecewise continuous functions defined over the 3D space for each individual part, whose distributions capture the topological and geometric properties of the surface features that partake in assembly. In the context of assembly guidance, the SDF can be viewed as an implicit generalization of virtual fixtures mentioned earlier, whose computation is automated for arbitrary geometry.

The convolution in the physical space (where the part geometries reside) transfers into a pointwise multiplication of the Fourier expansion of the SDFs for the individual parts (i.e., the ‘amplitudes’ of the multi-dimensional SDF signals). Guided by this property, we show that our previous formulation leads to a straightforward mathematical relationship between the

Fourier representations of the SDF shape descriptors and the geometric energy field, which can benefit from the efficiency of the FFT algorithms [62]. Moreover, we present explicit analytic equations for computing the gradients of the convolution function (i.e., guidance forces and torques) for arbitrary spatial translations and rotations. We implement the process using optimized FFT implementation on the highly-parallel GPU architecture. We show that haptic-enabled simulation of realistic assembly scenarios with complex CAD models and low-clearance fits is made possible to an adequate fidelity with the application of GPU-accelerated FFT calls.

2 Formulation

To compute the motion trajectory of objects in a virtual assembly scene, we aim to formulate an energy field over some configuration space (often abbreviated as the C-space) of the constituent parts, that

1. penalizes collisions between the interiors of the pairs of objects, which is equivalent to an implicit representation of the so-called ‘configuration space obstacles’ [21] (i.e., ‘ridges’ of the energy terrain); and
2. rewards proper fit/contact (i.e., configurations that exhibit superior shape complementarity between assembly features), which is equivalent to a re-scoring of the feasible collision-free C-space (i.e., ‘valleys’ of the energy terrain).

A recently proposed analytic approach to CD [16] holds great promise for computing the narrowphase collision response in real-time for haptic assembly. The method defines a ‘gap function’ between the two objects as a convolution of some implicit representation of the shapes—e.g., the ‘indicator function’ (whose value is 1 for interior (and boundary) points, and 0 elsewhere) [22], distance-based depth functions, or smooth ‘bump functions’ [16]. However, even *exact* CD is not sufficient for effective low-clearance insertion, since it computes large penalties for minor collision events that can occur due to hand vibration or device encoder inaccuracies [10, 56]. Furthermore, CD alone does not accomplish the second objective stated above, as will be detailed soon. By making use of a special class of shape descriptors (i.e., the SDF [20]) that better capture the geometric information of the assembly features, we overcome both problems in a unified framework.

2.1 Preliminaries

Let us begin with a brief overview of the derivation of SDF descriptors elaborated in [20] and implemented in [14]. The different steps in the process can be identified as 1) a projection from the shape domain in the 3-space to the complex domain $\mathbb{C} \cong \mathbb{R}^2$ composed of a Cartesian product of Euclidean metrics; 2) the application of a complex kernel to the projected boundary that extracts the topological and geometric properties pertinent to assembly; and 3) a surface integral over the boundary manifold to compute the shape descriptor. Here we give a different presentation that seems backwards in those steps, but provides useful insight into the subsequent Fourier analysis.¹ Following the common conventions in solid modeling, we restrict our attention to the class of ‘well-behaved’ solid objects $\mathcal{S} \subset \mathcal{P}(\mathbb{R}^3)$,² defined as compact regular semi-analytic subsets of the Euclidean metric space with the usual topology based on the L^2 -metric. This definition is general enough to encompass all practically significant shapes, yet sufficiently specific to avoid undesirable pathological behavior at the boundary [63] and the skeleton [64]. Hereon, we use the common terms ‘solid’ and ‘r-set’ interchangeably to refer to a member of this class.

2.2 Correlation Paradigm

For a pair of arbitrary solids $S_1, S_2 \in \mathcal{S}$, the configuration space of their relative rigid motions can be characterized with the special Euclidean group $SE(3) \cong SO(3) \times T(3)$ composed of a semi-direct product of the groups of proper orthogonal rotations $\mathcal{R} \in SO(3)$ (represented by 3×3 orthogonal matrices $[\mathcal{R}]_{3 \times 3}$ with $\det(\mathcal{R}) = +1$), and translations $\mathbf{t} \in T(3) \cong \mathbb{R}^3$ (represented by arbitrary 3-vectors $[\mathbf{t}]_{3 \times 1}$), together representing all possible rigid body transformations. In a virtual assembly environment, if S_1 and S_2 represent initial instances of the rigid solid parts (e.g., at rest on the assembly table), any subsequent position and orientation of the two parts (hereon referred to as the ‘absolute configurations’) are given by $\mathcal{T}_1 S_1$ and $\mathcal{T}_2 S_2$, respectively, where $\mathcal{T}_1, \mathcal{T}_2 \in SE(3)$ are arbitrary rigid motions (represented by tuples $\mathcal{T} := (\mathcal{R}, \mathbf{t})$, or equivalently, 4×4 homogeneous matrices $[\mathcal{T}]_{4 \times 4}$). The ‘relative configuration’ of S_2 as observed from a coordinate frame fixed on S_1 is $\mathcal{T} = \mathcal{T}_1^{-1} \mathcal{T}_2 \in SE(3)$. Every such configuration can be conceptualized as a point in the 6D C-space $SE(3)$. Following the common terminology in robotics and motion planning, the 6D geometric constructs that characterize the subsets of $SE(3)$ corresponding to collisions between S_1 and $\mathcal{T} S_2$ (or equivalently between $\mathcal{T}_1 S_1$ and $\mathcal{T}_2 S_2$ in the actual assembly scene) are called C-space obstacles (or ‘C-obstacles’ for short), and their regularized complement in $SE(3)$ is called the ‘free space’.

Cross-Correlation. Let $\rho_1, \rho_2 : \mathbb{R}^3 \rightarrow \mathbb{R}$ be real-valued functions over the 3-space that implicitly define S_1 and S_2 , respectively, i.e., $S_{1,2}$ are the regularized sublevel sets corresponding to $\rho_{1,2}(\mathbf{p}) \geq 0$ (or alternatively $\rho_{1,2}(\mathbf{p}) \leq 0$). Many important

¹Although we encourage the reader to review the concepts in [14] for a more thorough perspective, it is not necessary for a comprehension of this paper, where the attempt has been to present the key concepts in a self-sufficient manner.

²The collection $\mathcal{P}(A) = \{B \mid B \subset A\}$ denotes the ‘power set’ of a set A , i.e., the set of all subsets of A .

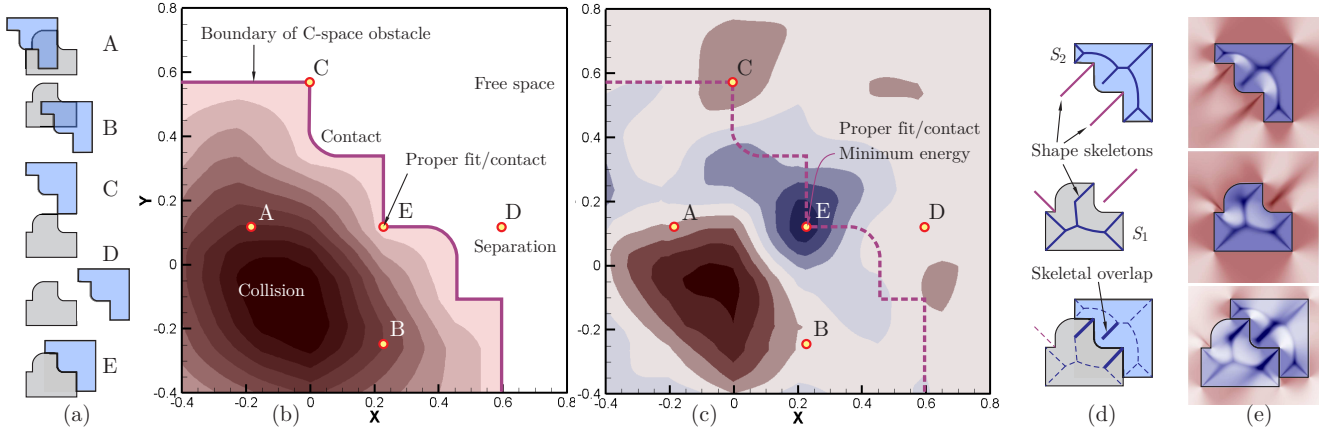


Fig. 1: Different configurations (a) are evaluated using a simple gap function (b), and the shape complementarity score function (c), based on overlapping shape skeletons (d), formulated as a cross-correlation of skeletal densities (e).

C-space problems can be formulated and solved in terms of the cross-correlation function $f_{CC} : SE(3) \rightarrow \mathbb{R}$ defined as the following volume integral:

$$f_{CC}(\mathcal{T}; S_1, S_2) = \int_{\mathbb{R}^3} \rho(\mathbf{p}; S_1) \rho(\mathcal{T}^{-1} \mathbf{p}; S_2) dV, \quad (1)$$

in which $\rho_{1,2}(\mathbf{p}) = \rho(\mathbf{p}; S_{1,2})$, where $\rho(\mathbf{p}; S)$ denotes a generic ‘defining function’ for an arbitrary r-set $S \in \mathcal{S}$; the integration variable is $\mathbf{p} = (x_1, x_2, x_3) \in \mathbb{R}^3$ and $dV = dx_1 dx_2 dx_3$ is the volume element. Note that the defining function must be invariant under rigid body transformation, hence $\rho(\mathbf{p}; \mathcal{T}S) = \rho(\mathcal{T}^{-1} \mathbf{p}; S)$ for all $S \in \mathcal{S}$ and $\mathcal{T} \in SE(3)$. Therefore, the integral in (1) overlaps the ρ -function of the two r-sets in their transformed positions and orientations. In order for $f_{CC}(\mathcal{T}; S_1, S_2)$ to remain bounded, the integrand functions need to either be compactly supported (e.g., $\rho(\mathbf{p}; S) = 0$ if $\mathbf{p} \notin S$) or approach zero as $\mathbf{p} \rightarrow \infty$ with a sufficiently rapid rate for the integral in (1) to converge.

Collision Detection. This correlation paradigm is central to important applications in geometric modeling and group morphology [15, 18]. For example, if we use the indicator function $\rho(\mathbf{p}; S) = 1$ if $\mathbf{p} \in S$ and $\rho(\mathbf{p}; S) = 0$ elsewhere, the integral in (1) simply computes the volume of the intersection ($S_1 \cap \mathcal{T}S_2$) (or equivalently, the volume of the intersection ($\mathcal{T}_1 S_1 \cap \mathcal{T}_2 S_2$) of the moved pair of parts). Consequently, $f_{CC} = 0$ characterizes the collection of feasible configurations corresponding to zero intersection volume, including unassembled (no-contact) and assembled (proper point/curve/surface contact), while $f_{CC} > 0$ implicitly defines the C-space obstacles (i.e., regions of the C-space that correspond to an interpenetration of parts). The holonomic unilateral contact constraints can be implemented by penalizing the moving part’s energy function proportional to the collision volume given by the so-called gap function $f_{CC}(\mathcal{T}; S_1, S_2)$, and the collision impulse forces and torques can be obtained from differentiating (1) using Lie algebras [16], to which we will return at the end of this section. Figure 1 (b) illustrates the translational C-space landscape for a pair of 2D solids shown in panel (a), along with the colormap for the gap function. To make the illustration possible, the motion is restricted to translation only, i.e., the landscape in panel (b) is a section (corresponding to zero rotation) through the full 3D C-space. Each of the relative positions in panel (a) are represented by a point in panel (b). The gap function penalizes collision as in positions A and B ($f_{CC} > 0$), but does not differentiate point contact in C and separation in D from proper fit/contact in E ($f_{CC} = 0$). This is clearly due to the property $\rho(\mathbf{p}; S_{1,2}) = 0$ for $\mathbf{p} \notin S$ as a result of the definition.

Geometric Guidance. The question remains as how to modify this approach to 1) incorporate nonzero values to the cross-correlation function over the free space, such that it rewards proper fit/contact (e.g., position E in Fig. 1 (a)), slightly penalizes separation and insufficient contact (e.g., positions C and D), in addition to the high penalty already assigned to collision (e.g., positions A and B); and 2) provide a mechanism to adjust relaxation of the collision penalty, and control the smoothness of the transition between the free space and the obstacle space. In other words, how can we add ‘valleys’ to the free space in addition to the ‘ridges’ inside the obstacle space, and how can we control the steepness of the transition between the two? The former allows additional haptic assistance for insertion, e.g., ‘magnetic’ attraction forces and torques that guide the assembly by restricting the DOF of motion, and eventually snap the peg into the hole. The latter is important to let the user tune the flexibility of the collision response (i.e., ‘soft’ versus ‘hard’ assembly) as well as the degree of strictness with which the guidance constraints are enforced.

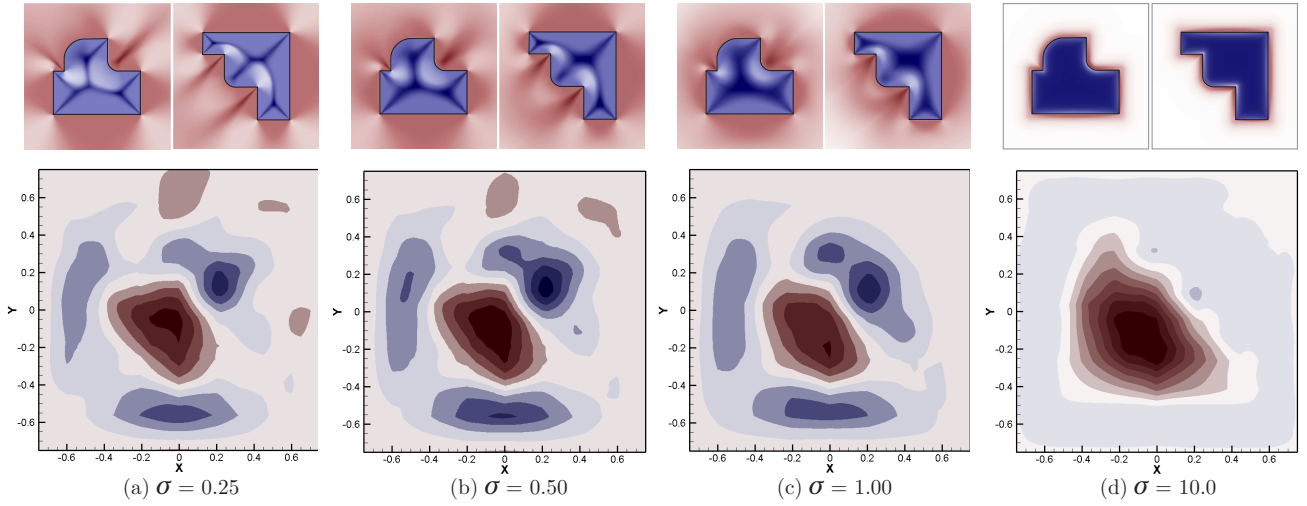


Fig. 2: The extent of geometric details captured by the skeletal density distribution is adjustable by the thickness factor σ .

2.3 Shape Descriptors

The answers to both questions lie in the development of more sophisticated shape descriptors $\rho_{1,2} = \rho(\mathbf{p}; S_{1,2})$ with $\rho(\mathbf{p}; S_{1,2}) \neq 0$ outside the object, rather than using, for example, the binary indicator functions or bump functions. We consider a special class of complex-valued functions $\rho_\sigma : (\mathbb{R}^3 - \partial S) \rightarrow \mathbb{C}$ that can be defined in the form of the following surface flux integral over the solid boundary [14, 20]:

$$\rho_\sigma(\mathbf{p}; S) = \oint_{\partial S} \phi_\sigma \left[\zeta(\mathbf{p}, \mathbf{q}; S) \right] dA_\perp, \quad (2)$$

where dA_\perp is the signed area element normal to the line that connects $\mathbf{p} \in \mathbb{R}^3$ to $\mathbf{q} \in \partial S$. As detailed in [14], $\zeta : (\mathbb{R}^3 \times \partial S) \rightarrow \mathbb{C}$ characterizes a projection of the boundary ∂S from the 3-space to the complex plane, defined as

$$\zeta(\mathbf{p}, \mathbf{q}; S) = \xi(\mathbf{p}; S) + i\eta(\mathbf{p}, \mathbf{q}), \quad (3)$$

where the real-part $\xi(\mathbf{p}; S) = \pm \min_{\mathbf{q} \in \partial S} \|\mathbf{p} - \mathbf{q}\|_2$ is the signed Euclidean distance from the boundary ∂S to the query point $\mathbf{p} \in \mathbb{R}^3$,³ and the imaginary-part $\eta(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|_2$ is simply the L^2 -distance between one particular boundary point \mathbf{q} to the point \mathbf{p} . The kernel $\phi_\sigma : (\mathbb{C} - \{0\}) \rightarrow \mathbb{C}$ is the key, whose choice characterizes the set of topological and geometric properties that are extracted from the distance distribution embedded in the ζ -map. The formulation as a surface integral characterizes a linear combination of the shares of different surface elements superimposed into $\rho_\sigma(\mathbf{p}; S)$.

It is interesting to note that the binary indicator functions used earlier to compute the collision gap function in (1) can be regarded as a special case of (2). If we let $\phi(\zeta) = -(4\pi\eta^2)^{-1}$ where $\eta = \Im\{\zeta\}$ is the imaginary-part of the argument, the combination of the integrand and normal area element in (2) gives $\phi(\zeta)dA_\perp = -d\gamma/4\pi$ where $d\gamma = dA_\perp / \|\mathbf{p} - \mathbf{q}\|_2^2$ is the signed infinitesimal spatial angle by which the query point $\mathbf{p} \in (\mathbb{R}^3 - \partial S)$ observes the area element dA at $\mathbf{q} \in \partial S$. Therefore, the surface integral in (2) computes the ‘winding number,’ which can be used for inclusion (i.e., PMC) testing [65] (i.e., the winding number is 1 in the interior, 0 in the exterior, and undefined on the boundary of S). Disregarding the boundary points—which do not contribute to the volume integral in (1)—this gives the indicator function whose convolution in (1) yields the gap function for the collision response.

Skeletal Density. We demonstrated in [20] that shape complementarity between objects of arbitrary shape can be related to the overlapping of shape skeleton branches that are generated by their assembly features—see Fig. 1 (d), for example, where the assembly features (i.e., the two sharp corners and one filleted corner fitting together) are each captured by overlapping branches of the medial axis (MA). To develop an implicit continuous function that highlights the skeletal features, we first modify the aforementioned choice of the kernel as $\phi(\zeta) \propto \eta^{-2}$ (which led to the indicator function when substituted in (2)), to $\phi_\sigma(\zeta) \propto \eta^{-2} g_\sigma(|\tan \angle \zeta| - 1)$, where $g_\sigma(x) = (\sqrt{2\pi}\sigma)^{-1} e^{-\frac{1}{2}(x/\sigma)^2}$ is the Gaussian function. This new term incorporates a

³The sign is positive if the query point is external ($\mathbf{p} \in \mathbf{c}S$) and negative if the query point is internal ($\mathbf{p} \in \mathbf{i}S$), determined by Point Membership Classification (PMC) [65].

higher contribution to the surface integral in (2) for boundary elements at which $\tan \angle \zeta \approx \pm 1$ (i.e., $\eta \approx |\xi|$) which means the boundary point $\mathbf{q} \in \partial S$ is an approximate nearest neighbor (ANN) to the query point $\mathbf{p} \in (\mathbb{R}^3 - \partial S)$. Query points near the shape skeleton have more extensive ANNs, hence receive more such contributions. As depicted in Fig. 1 (e), a continuous measure of skeletal overlap can be obtained by convolving such density functions using (1) for a given relative position and orientation of the two parts. The resulting change in the convolution function over the translational C-space is shown in Fig. 1 (c), where the energy well is located at position E. Figure 2 (a–d) illustrates the effect of changing the thickness factor $\sigma > 0$ on the geometric energy landscape. Clearly, as $\sigma \rightarrow +\infty$ the Gaussian flattens out and the density function approaches the indicator function (with the exception of a signed coefficient), and as $\sigma \rightarrow 0^+$ the high density regions further resemble the MA.

For 3D interactions in haptic assembly, we used a slightly different complex structure $\phi(\zeta) \propto \zeta^{-2} g_\sigma(|\tan \angle \zeta| - 1)$ in [14] with different signed coefficients $\pm \lambda_{1,2}$ for the interior and exterior query points, determined by the PMC function. This allowed for positive real contributions to f_{CC} in (1) due to external-internal and internal-external skeletal overlaps (i.e., rewarding ‘proper fit’), and negative real contributions due to internal-internal and external-external overlaps (i.e., penalizing ‘collision’ and ‘separation’), the relative intensities of which could be adjusted by the choice of the ‘penalty factor’ $\mathfrak{p} = \lambda_2/\lambda_1$. We refer to the resulting generic function $\rho(\mathbf{p}; S)$ as the ‘affinity function’ in this context, and to the cross-correlation function in (1) as the shape complementarity ‘score function’ (denoted f_{SC} instead of f_{CC}). The geometric energy is then defined as $E_G \propto \Re\{f_{SC}\}$ and the guidance forces and torques are obtained as its gradients. Skipping the details, here it suffices to repeat that aligning the skeletal branches (as in Fig. 1 (d)) can be thought of as a generalization of virtual fixtures [60] to arbitrary geometry, and implicit representation using SDF offers robustness and stability with respect to boundary perturbations, in contrast to the inherently unstable MA.

2.4 Fourier Convolution

As a result of the definition, the SDF defined in (2) is invariant under rigid body motion. This means that for a rigid solid instance $S \in \mathcal{S}$ and a transformation $\mathcal{T} \in \text{SE}(3)$, computing the ρ -function at $\mathbf{p} \in \mathbb{R}^3$ for the moved instance $\mathcal{T}S$ is equivalent to querying the precomputed SDF for the original instance at points displaced with the inverse transformation $\mathcal{T}^{-1} \in \text{SE}(3)$, i.e., $\rho_\sigma(\mathbf{p}; \mathcal{T}S) = \rho_\sigma(\mathcal{T}^{-1}\mathbf{p}; S)$. Therefore, for a pair of solids $S_1, S_2 \in \mathcal{S}$ transformed with $\mathcal{T}_1, \mathcal{T}_2 \in \text{SE}(3)$, respectively, the score function per unit volume (whose real-part is the ‘geometric energy density’) at a query point $\mathbf{p}' \in \mathbb{R}^3 - (\partial(\mathcal{T}_1 S) \cup \partial(\mathcal{T}_2 S_2))$ can be computed from a product of the SDFs (measuring the skeletal overlap):

$$\left. \frac{df_{SC}}{dV} \right|_{\mathbf{p}'} = \rho_\sigma(\mathbf{p}'; \mathcal{T}_1 S_1) \rho_\sigma(\mathbf{p}'; \mathcal{T}_2 S_2) \quad (4)$$

$$= \rho_\sigma(\mathbf{p}; S_1) \rho_\sigma(\mathcal{T}_2^{-1} \mathcal{T}_1 \mathbf{p}; S_2). \quad (5)$$

where $\mathcal{T} = \mathcal{T}_1^{-1} \mathcal{T}_2 \in \text{SE}(3)$ is the motion of S_2 as observed from a coordinate frame attached to S_1 (compare with (1)), and $\mathbf{p} = \mathcal{T}_1^{-1} \mathbf{p}'$ is the new coordinates of the query point measured with respect to that frame. The overlap density df_{SC}/dV is then accumulated over the 3-space to obtain the score in (1).

Motion Decomposition. The only time-dependent variable in (1) is the 6D relative motion $\mathcal{T} \in \text{SE}(3)$. To simplify the subsequent development, let us decompose the motion into the translational component $\mathbf{t} \in \text{T}(3)$ described by a 3-tuple $(t_1, t_2, t_3) \in \mathbb{R}^3$, and the rotational component $\mathcal{R} \in \text{SO}(3)$ represented by a 3×3 proper orthogonal matrix $[\mathcal{R}]_{3 \times 3}$ —‘orthogonal meaning $\mathcal{R}^{-1} = \mathcal{R}^T$ and ‘proper’ meaning $\det(\mathcal{R}) = +1$. As a result of the product structure $\text{SE}(3) = \text{SO}(3) \times \text{T}(3)$, the transformation sequence applies as $\mathcal{T}\mathbf{p} = (\mathcal{R}\mathbf{p}) + \mathbf{t}$ hence $\mathcal{T}^{-1}\mathbf{p} = \mathcal{R}^T(\mathbf{p} - \mathbf{t})$. Noting the rotational invariance property of the SDF, it is easy to see that for all $S \in \mathcal{S}$,

$$\rho_\sigma(\mathcal{T}^{-1}\mathbf{p}; S) = \rho_\sigma(\mathcal{R}^T(\mathbf{p} - \mathbf{t}); S) = \rho_\sigma(\mathbf{p} - \mathbf{t}; \mathcal{R}S). \quad (6)$$

Using the notation $\rho_{1,2}(\mathbf{p}) = \rho_\sigma(\mathbf{p}; S_{1,2})$ and defining the reflection of a generic function as $\tilde{f}(\mathbf{p}) = f(-\mathbf{p})$, the integral in (1) becomes a 3D convolution over a section through the C-space corresponding to the fixed rotation $\mathcal{R} \in \text{SO}(3)$:

$$\begin{aligned} f_{SC}((\mathcal{R}, \mathbf{t}); S_1, S_2) &= \int_{\mathbb{R}^3} \rho_1(\mathbf{p}) \tilde{\rho}_2(\mathcal{R}^T(\mathbf{t} - \mathbf{p})) dV \\ &= (\rho_1 * (\tilde{\rho}_2 \circ \mathcal{R}^T))(\mathbf{t}). \end{aligned} \quad (7)$$

Fourier Transforms. Using the orthonormal Fourier basis of the form $e^{2\pi i(\omega \cdot \mathbf{p})}$ where $\mathbf{p}, \omega \in \mathbb{R}^3$, one could decompose a function $f : \mathbb{R}^3 \rightarrow \mathbb{C}$ (defined over the 3D physical domain) into its components captured by the function $\hat{f} : \mathbb{R}^3 \rightarrow \mathbb{C}$ (defined over the 3D frequency domain). The forward Fourier transform $\hat{f} = \mathcal{F}\{f\}$ is thus defined as the inner product

$$\hat{f}(\omega) = \left\langle f, e^{+2\pi i(\omega \cdot \mathbf{p})} \right\rangle = \int_{\mathbb{R}^3} f(\mathbf{p}) e^{-2\pi i(\omega \cdot \mathbf{p})} dV, \quad (8)$$

and the inverse Fourier transform $f = \mathcal{F}^{-1}\{\hat{f}\}$ is defined as follows, to retrieve the function as a superposition of its orthogonal components:

$$f(\mathbf{p}) = \left\langle \hat{f}, e^{-2\pi i(\omega \cdot \mathbf{p})} \right\rangle = \int_{\mathbb{R}^3} \hat{f}(\omega) e^{+2\pi i(\omega \cdot \mathbf{p})} dV. \quad (9)$$

The integrals in (8) and (9) are sometimes referred to as the continuous Fourier Transform (CFT), whose discredited form for computation purposes is given in Section 3.2. The *convolution theorem* states that the Fourier transform of a convolution is the pointwise product of Fourier transforms [61], hence the integral in (7) converts in the frequency domain to the simple product $\hat{f}_{\text{SC}}(\omega) = \mathcal{F}\{f_{\text{SC}}\} = \mathcal{F}\{\rho_1\} \mathcal{F}\{\tilde{\rho}_2 \circ \mathcal{R}^T\}$. As a direct consequence of the definition in (8), the rotation and reflection commute with the Fourier transform (noting that $\omega \cdot (\mathcal{R}^T \mathbf{p}) = (\mathcal{R} \omega) \cdot \mathbf{p}$ and $\omega \cdot (-\mathbf{p}) = (-\omega) \cdot \mathbf{p}$ in (8)), hence $\mathcal{F}\{\tilde{\rho}_2 \circ \mathcal{R}^T\} = \mathcal{F}\{\tilde{\rho}_2\} \circ \mathcal{R}^T$ in which $\mathcal{F}\{\tilde{\rho}_2\} = \tilde{\mathcal{F}}\{\rho_2\}$. The score function in the physical domain can thus be computed by applying an inverse transform as

$$f_{\text{SC}}((\mathcal{R}, \mathbf{t}); \mathcal{S}_1, \mathcal{S}_2) = \mathcal{F}^{-1} \left\{ \mathcal{F}\{\rho_1\} (\mathcal{F}\{\tilde{\rho}_2\} \circ \mathcal{R}^T) \right\}. \quad (10)$$

At a first glance, this might appear as aggravating the computational burden, by requiring the evaluation of 3 volume integrals rather than the one in (7). However, (8) and (9) can be computed efficiently using the FFT algorithm [62], as will be demonstrated in Section 3.

Convolution Gradients. The forces and torques can be obtained from the gradient of the geometric energy field $E_G : \mathbb{R}^3 \rightarrow \mathbb{R}$, which is proportional to the real-part of the shape complementarity score function $f_{\text{SC}} : \mathbb{R}^3 \rightarrow \mathbb{C}$, i.e., its rate of change with respect to infinitesimal translations and rotations. This is possible by differentiating either the physical domain formulation (7) (detailed in [14]) or the frequency domain formulation in (10). Another important advantage of working in the Fourier domain is that the translational differentiation is replaced with a multiplier:

$$\left\langle \frac{df_{\text{SC}}}{d\mathbf{t}}, \mathbf{e} \right\rangle = (2\pi i) \mathcal{F}^{-1} \left\{ (\mathbf{e} \cdot \omega) \mathcal{F}\{\rho_1\} (\mathcal{F}\{\tilde{\rho}_2\} \circ \mathcal{R}^T) \right\}, \quad (11)$$

where $\mathbf{e} \in \mathbb{R}^3$ represents any direction in the vector space $\mathbb{T}(3) \cong \mathbb{R}^3$, along which the differentiation occurs. The rotational differentiation, on the other hand, is more involved since $\text{SO}(3)$ is not a vector space and cannot be globally parameterized by a single continuous 3D grid. To obtain a local parametrization, the tangent direction at $\mathcal{R} \in \text{SO}(3)$ is obtained as $\mathcal{R}\Omega$ where $\Omega \in \mathfrak{so}(3)$ can be represented by a 3×3 skew-symmetric matrix $[\Omega]_{3 \times 3}$, and $\mathfrak{so}(3)$ denotes the Lie algebra for $\text{SO}(3)$, which is a vector space tangent to $\text{SO}(3)$ at the identity rotation [16]. Without getting into much detail, we present the rotational gradient as

$$\left\langle \frac{df_{\text{SC}}}{d\mathcal{R}}, \mathbf{e} \right\rangle = (2\pi i) \mathcal{F}^{-1} \left\{ (\mathcal{R}^T \Omega \omega) \cdot \mathcal{F}\{\rho_1\} (\mathcal{F}\{\tilde{\rho}_2^*\} \circ \mathcal{R}^T) \right\}, \quad (12)$$

where $\mathbf{e} \in \mathbb{R}^3$ is the dual vector of $\Omega \in \mathfrak{so}(3)$, which means $\Omega \omega = \mathbf{e} \times \omega$. $\rho_2^* : \mathbb{R}^3 \rightarrow \mathbb{C}^3$ is a vector function defined as $\rho_2^*(\mathbf{p}) = +\rho_2(\mathbf{p})\mathbf{p}$ (hence $\tilde{\rho}_2^*(\mathbf{p}) = -\tilde{\rho}_2(\mathbf{p})\mathbf{p}$), whose Fourier transform $\mathcal{F}\{\tilde{\rho}_2^*\}$ can be obtained by componentwise FFTs of its complex components projected on the 3 coordinate axes. See [16] for a more in-depth treatment of rotational gradients.

The 3D translational and rotational gradient vectors can be computed in a componentwise fashion by substituting for the base vectors $\mathbf{e} \in \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ one at a time in (11) and (12). The complete 6D gradient $\nabla f_{\text{SC}} : \text{SE}(3) \rightarrow \mathbb{C}^6$ is defined as $\nabla f_{\text{SC}} = (df_{\text{SC}}/d\mathbf{t}, df_{\text{SC}}/d\mathcal{R})$, and the geometric force and torque are obtained as $(\mathbf{F}_G, \mathbf{T}_G) = -\nabla E_G \propto \Re\{\nabla f_{\text{SC}}\}$ as detailed in [14].

3 Implementation

In this section we present the numerical algorithms that carry out a discrete approximation of the SDF integrals and the subsequent Fourier transforms for a particularly simple representation (namely, triangular mesh B-reps). We also present a straightforward complexity analysis of each step, and refer the reader to [20] for more details.

3.1 Precomputations

The following summarizes the offline SDF precomputations for mesh representations in both physical and frequency domains, the former being described in more detail in [14].

Representation. The analytic definition of the SDF in (2) does not impose any restriction on the representation scheme, as long as it satisfies the informational completeness requirement [66], and more specifically, supports distance and inclusion queries for arbitrary query points in the 3D space. Nevertheless, the numerical computation of the surface integral in (2) lends itself well to B-reps, and particularly to triangular mesh approximations. For mesh generation from STEP models (exported from CAD software of choice), we use the NETGEN library [67]. We obtain the unsigned distance field using Havoc3D [68], which computes the rasterized Voronoi diagram using the OpenGL rendering pipeline and depth-buffer. To compute the PMC and correct the sign of the distance function, on the other hand, we take advantage of the winding number approach in [65], which can be thought of as a special case of (2)—i.e., one with a simple inverse-square kernel $\phi(\xi + i\eta) = (4\pi\eta^2)^{-1}$ as explained earlier in Section 2.3—which can be computed using the same subroutines that compute the SDF with arbitrary kernels.

Preprocessing. The sequence of steps can be summarized as 1) generating a mesh from the CAD model using NETGEN [67]; 2) computing the unsigned distance function using Havoc3D [68]; 3) evaluating the PMC function [65] (to correct the distance signs) using (2) with an inverse-square kernel; and 4) evaluating the SDF descriptor [20] using (2) with a combined Gaussian and inverse-square kernel. The last two steps can be implemented by approximating the integral in (2) as a discrete Riemann sum over the triangles. In order to ensure numerical stability one must guarantee an upperbound on the error, which is not possible by assigning a lumped weight to each triangle in the sum. This is due to the inverse-square term in the kernel, resulting in large errors as the query point gets closer to a particular triangle. To overcome this difficulty, the algorithm carries out adaptive recursive subdivisions of the triangles with a spatial angle-based threshold, i.e., until the triangle is subdivided to small enough pieces each observed from the query point by a small spatial angle $\delta\gamma = \delta A_{\perp} / (4\pi\eta^2)$. More implementation details for SDF precomputation in the physical domain are given in [14], and will not be repeated here. It suffices to mention that for a mesh $\Delta_n(S) = \bigcup_{j=1}^n \delta_j$ that approximates the boundary ∂S with the complex of n faces δ_j ($1 \leq j \leq n$) and $O(n)$ vertices and edges, and a uniformly sampled 3D grid of m query points $G_m(S)$ that contains the bounding box of S , all steps take $O(mn)$ sequential time and $O(m+n)$ memory space. Furthermore, the summation can be parallelized in a trivial manner by assigning disjoint subsets of the grid to different processors.

We implement (2) in parallel for the multi-core central processing units (CPU) using the Boost C++ libraries [69] for multi-threading,⁴ and obtain significant speed-ups that scale almost linearly with the number of dedicated cores. In addition, the 3D grid structure maps properly to the single-instruction multiple-thread (SIMT) execution model of the modern many-core GPUs, allowing us to further speed up the process using NVIDIA’s compute-unified device architecture (CUDA).

3.2 FFT Computations

By sampling the physical and frequency domain data over the 3D grids $G_m(S)$ and $\hat{G}_m(S)$, respectively, the volume integrals in (8) and (9) that define the forward and inverse CFTs, respectively, can be approximated by the sums

$$\hat{f}_k \approx \sum_{i=1}^m f_i e^{-2\pi i(\omega_k \cdot \mathbf{p}_i)} \delta V, \quad 1 \leq k \leq m, \quad (13)$$

$$f_i \approx \sum_{k=1}^m \hat{f}_k e^{+2\pi i(\omega_k \cdot \mathbf{p}_i)} \delta V, \quad 1 \leq i \leq m, \quad (14)$$

where $\mathbf{p}_i \in G_m(S)$ and $\omega_k \in \hat{G}_m(S)$ are uniformly sampled physical and frequency nodes at which the function values $f_i := f(\mathbf{p}_i)$ and $\hat{f}_k := \hat{f}(\omega_k)$ are stored, respectively. With the exception of a constant factor (depending on the chosen conventions), these sums define the discrete Fourier transform (DFT) whose cumulative computation for all grid nodes takes $O(m^2)$ basic operations using the cascade method. However, the same computation can be carried out in $O(m \log m)$ steps using the radix-2 FFT algorithm [62].

⁴In the earlier implementation reported in [14] we used OpenMP for the task of CPU multi-threading. A significant improvement in performance was observed by reimplementing the same steps using the Boost C++ [69] ‘thread’ library.

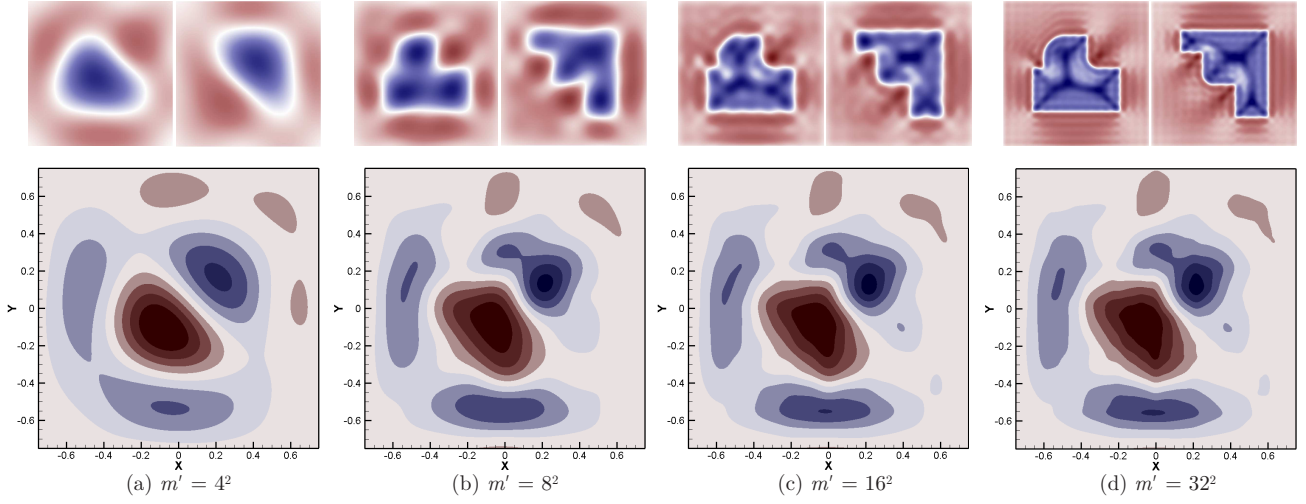


Fig. 3: Frequency domain representation allows for a systematic means of successive approximation of the energy field.

A key observation is that the frequency domain representations of the SDFs, namely $\hat{\rho}_1 = \mathcal{F}\{\rho_1\}$, $\hat{\rho}_2 = \mathcal{F}\{\rho_2\}$, and $\hat{\rho}_2^* = \mathcal{F}\{\rho_2^*\}$ depend on part geometries alone (and not on the instantaneous assembly configuration), hence can be pre-computed offline prior to the virtual assembly session. For a pair of parts, computing the forward FFT in (13) to obtain $\hat{\rho}_{1,k}, \hat{\rho}_{1,k}, \hat{\rho}_{2,k}^*$ ($1 \leq k \leq m$) from $\rho_{1,i}, \rho_{2,i}, \rho_{2,i}^*$ ($1 \leq i \leq m$) takes $O(m \log m)$ per part.

The subsequent computation of the convolutions in (10), (11), and (12) (for geometric energy, force, and torque evaluations) in real-time for a particular relative orientation $\mathcal{R} \in \text{SO}(3)$ of the parts takes place entirely in the frequency domain. The sequence of operations is 1) interpolating $(\hat{\rho}_2 \circ \mathcal{R}^T)$ and $(\hat{\rho}_2^* \circ \mathcal{R}^T)$ over a rotated grid $\mathcal{R}^T \hat{G}_m(S_2)$, followed by a reflection, which takes $O(m)$ basic trilinear interpolation steps; 2) a pointwise multiplication of the interpolated data with $\hat{\rho}_1$ data over $\hat{G}_m(S_1)$, which also takes $O(m)$ basic steps; and 3) an inverse FFT along with applying the proper coefficients (e.g., 2π for force/torque computations), which takes $O(m \log m)$ steps and yields the convolution results over a grid of translations. Therefore, the total arithmetic complexity of the process (for a single rotation and all translations of interest) is $O(m \log m)$.

We use the FFTW library [70] for the CPU sequential implementation and NVIDIA’s cuFFT(W) for the GPU parallel implementation of the FFT, the running times of both to be presented in Section 4.2 for comparison.

3.3 Low-Pass Filtering

The additional important advantage of this representation is that one could decide to keep only a small subset of $m' \ll m$ frequency domain data nodes (i.e., the dominant modes) for computing the Fourier convolutions in (10), (11), and (12) in real-time, in a trade-off between the desired accuracy and available computational power. This results in a reduction of the real-time process complexity to $O(m' \log m')$ operations per frame, which is practically almost $O(1)$. In other words, the interpolation (of the rotated and reflected data), the pointwise multiplication, and the inverse FFT steps are carried out for a significantly smaller sample of low-frequency grid nodes, a process referred to as ‘low-pass filtering’. As we will demonstrate in Section 4.2, the pointwise multiplication step is the bottleneck and dictates an upperbound on the choice of m' in order to stay within a prespecified time allocated to guidance force and torque computations at each frame, hence to achieve the desired frame rate of 1 kHz.

It is worthwhile noting that the inverse FFT cumulatively computes the correlations for a range of m translations corresponding to a single rotation. This is a nonoptimal approach from a theoretical point of view, since we only need the results for a single configuration, i.e., the instantaneous relative translation and rotation of the objects $(\mathcal{R}, \mathbf{t}) \in \text{SE}(3)$ at the current simulation frame. Therefore, $O(m' \log m')$ of the FFT can be reduced to $O(m')$ of a cascade partial summation of (14) for a single translation index i . However, this does not yield a significant performance gain in practice, where the running time is governed by the notably slower pointwise multiplication step. Besides, parallel implementation of a cascade sum (especially on the GPU) for small m' does not necessarily outperform its FFT counterpart.

Figure 4 (a–d) illustrates the successive Fourier approximations for the 2D skeletal densities in Fig. 2 (b) with $\sigma = 0.5$. A grid size of $m = 512^2 = 262,144$ nodes is used, resulting in the same number of frequency domain amplitudes. In this case, the evolution of the score function is not substantial for $m' > 16^2 = 256$ (less than 0.1% of m). However, shapes with higher geometric detail require a larger number of modes to capture the smallest features. It is very important to note that the input model complexity (e.g., the number of triangles $n_{1,2}$ per mesh representation of the solids $S_{1,2} \in \mathcal{S}$) is irrelevant here, and the suitable value of m' is determined from the desired fidelity with which the small geometric features are captured in the output energy field. This indifference to the syntactic representation complexity is a significant advantage of our method

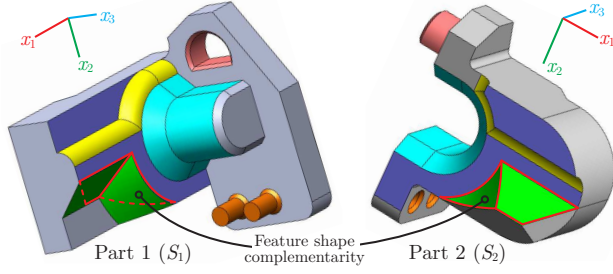


Fig. 4: A non-trivial, zero-clearance assembly pair.

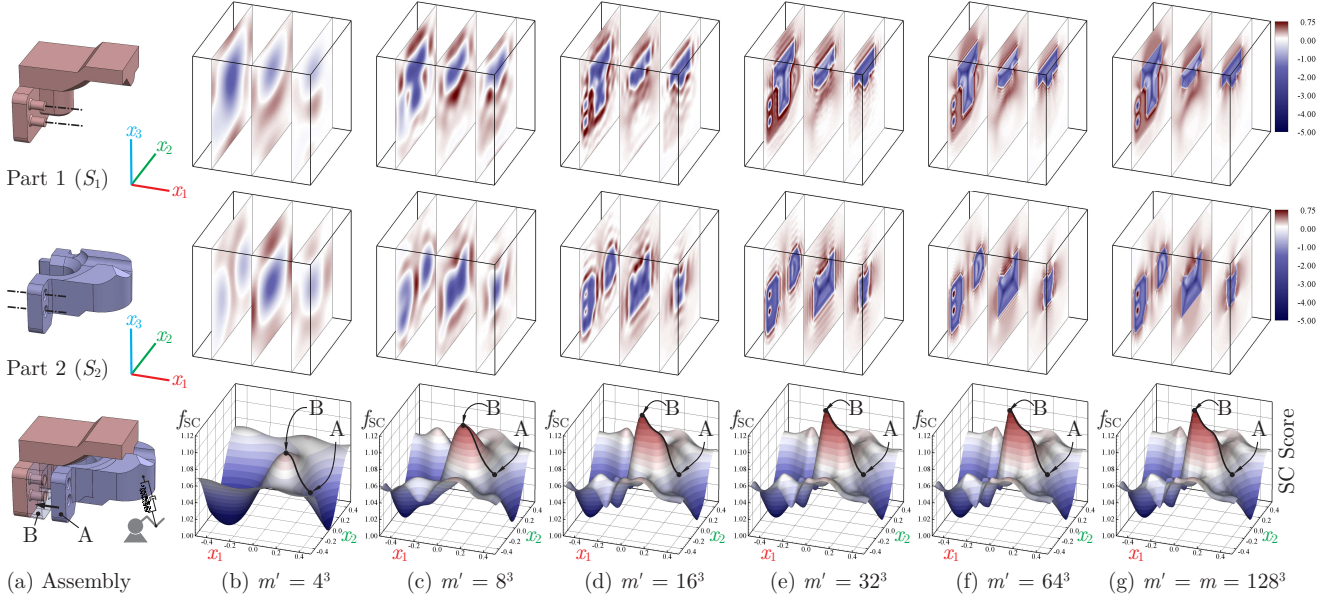


Fig. 5: The effect of FFT filtering on part SDFs (top) and score variations versus biaxial relative translation (bottom).

over most other collision detection and constraint management algorithms, whose running times depend on (and scale with) the representation size, resulting in a failure to handle large mesh sizes in real-time due to the high frame rate limitation.

4 Results & Discussion

In [14], we tested the SDF descriptors for haptic-assisted assembly of simple peg-in-hole geometries. The examples provided important insight into the properties of SDF and its effectiveness as an automated generalization of virtual fixtures for arbitrary geometry. However, the numerically intensive cascade computation of the integral in (1) required substantial preprocessing time. Here we use the FFT convolution technique, using optimized CPU- and GPU-accelerated implementations. To demonstrate the practicality of the method, we use the pair of 3D assembly parts shown in Fig. 4. The solids in this example are made of semi-algebraic r-sets with only planar, cylindrical, spherical, and toroidal surface patches, which forms a small subset of the general semi-algebraic class (i.e., solids bounded by polynomial surfaces of arbitrary degrees) and even more general semi-analytic class covered by our formulation. Nevertheless, an automatic identification of the correspondence between the mating features (depicted with different colors) is not trivial from an algorithmic perspective—e.g., recognition and matching of the partially complementary features connected by a curve. Furthermore, there are 3 pairs of pegs and holes with zero clearance, making this example sufficiently challenging.

4.1 Offline Preprocessing

The running times for the offline steps made of precomputing the part SDFs in the physical domain and their forward FFT into the frequency domain are plotted in Figs. 6 and 7.

Precomputing SDFs. The performance of the CPU- and GPU-accelerated SDF computation for parts in Fig. 4 is presented in Fig. 6 (a, b) for different mesh and sample sizes. The NVIDIA Tesla K20c GPU (2,496 CUDA cores, 5GB device memory)

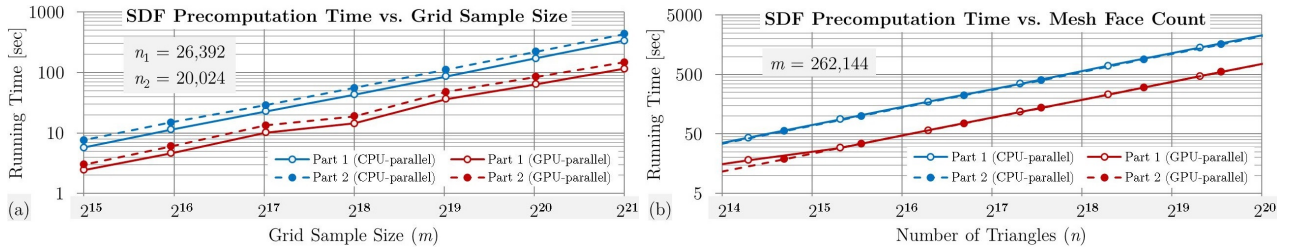


Fig. 6: CPU vs. GPU performances for SDF computation.

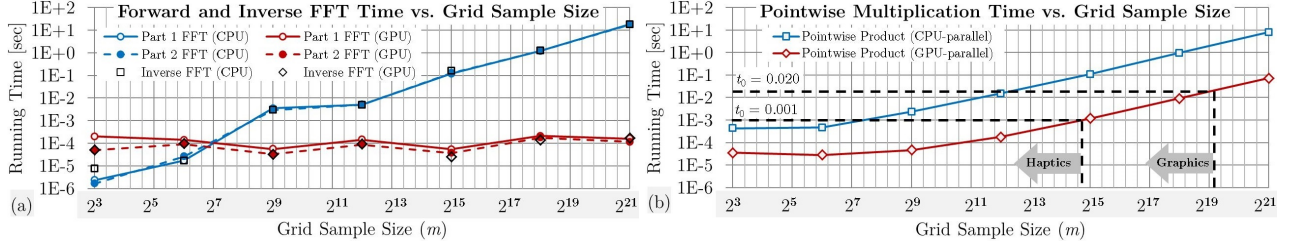


Fig. 7: CPU vs. GPU performances for FFT convolution.

outperforms the Intel Xeon E5-2687W CPU (32 cores, 3.10 GHz clock-rate, 64GB host memory) by average speed-up factors of $2.5 - 3.0 \times$.⁵ For this particular example, a sample size of $m = 128^3 = 2,097,152$ and $n_{1,2} \approx 2 \times 10^5$ is adequate to capture the geometric details within the SDF field, which takes about 2–3 minutes per part to precompute the SDF offline. The resulting SDFs of the parts in Fig. 4 (using $\sigma := 0.5$ and $p = \lambda_2/\lambda_1 := 3$) are plotted for their imaginary parts in Fig. 5 (g). The correspondence between high density regions (e.g., along the axes of cylindrical features or along the bisectors of corners of the boundary), which are analytic generalizations of virtual fixtures, is apparent.

Forward FFT of SDFs. Figure 7 (a) presents the performance of the forward FFTs versus the sample size to map the part SDFs into the frequency domain. It is important to note that the forward FFT of part SDFs is also an offline preprocessing step, which takes negligible time compared to the previous steps (typically in the sub-millisecond range).

4.2 Real-Time Processing

The combination of Fourier amplitude product and inverse FFT cumulatively produces the geometric energy response for all relevant transformational configurations, which can take up to 0.1 second for large sample sizes, as depicted in Fig. 7 (b). However, during haptic assembly we ideally have less than 1 millisecond, but we also need to evaluate only a single configuration at any instant of time during the virtual assembly; namely, the one corresponding to the instantaneous physical positioning of the objects. Two different strategies can be used to adapt this method to real-time applications:

Restricted Motion DOF. It is not possible in practice to precompute the energy field for a 6D grid of all possible motions. However, for most assembly scenarios the motion during the insertion phase is constrained to one or two DOF. For example, if the rotational space is limited to a finite number of permissible relative orientations, the inverse FFT for 3D translational motion can be precomputed and stored for each orientation, and queried rapidly during motion. This approach allows for computing the guidance force feedback to full accuracy, but goes against the philosophy of avoiding the multi-phase approach and manual specifications, from which we set off to pursue this method.

Truncated Inverse FFT. The major benefit of working in the frequency domain is the systematic means it provides to trade off the accuracy of physical domain representation with computation time. This provides a chance for real-time computation of translational convolution for arbitrary rotations given at any instant of assembly simulation. If t_0 is the amount of dedicated computation time available at each frame (e.g., ideally $t_0 \leq 0.020$ seconds for 50 Hz graphic rendering, and $t_0 \leq 0.001$ seconds for 1 kHz haptic feedback), one can always choose the maximal sample size $m_0 \ll m$ whose processing time is approximately t_0 . Hence real-time computing is contingent upon keeping only $m' \leq m_0$ dominant modes of the SDFs in the frequency domain to approximate the convolution function (see Fig. 4).

⁵This is less than expected, because the GPU implementation suffers from extra overhead due to data transfer between host and device memories, and is suboptimal in performing conditional instructions. Much better speed-ups up to two orders of magnitude are obtained in the convolution step, as depicted next.

Figure 7 presents the running times of the convolution step, composed of CPU- and GPU-parallel pointwise multiplication of the mapped SDFs (plotted in panel (b)), followed by an inverse FFT to obtain the geometric energy field in the physical domain (plotted in panel (a)). Here the GPU implementation of both steps makes a crucial difference, as the inverse FFT time on the GPU stays significantly below the 1 millisecond threshold even for sample sizes as large as $m \sim 10^6 - 10^7$. The pointwise multiplication is the bottleneck, whose running time determines the upperbound m_0 on the size of truncated Fourier expansions. For this step, CPU parallelization speeds up the multiplication process by about $25\times$ over the sequential implementation, and as depicted in Fig. 7 the GPU implementation enhances it by an additional factor of about $100\times$ (hence an overall speed-up as much as $2,500\times$). The difference grows with the sample size, which indicates the scalability of the process.

Effects of Filtering. Figure 5 (b–f) shows the effects of successive FFT filtering described in Section 3.3 for different numbers of retained dominant modes $m' \leq m$ on the SDF (only imaginary part plotted on top rows). It also plots the score function on the bottom row, over a 2D section corresponding to a biaxial relative translation along the x_1x_2 -plane through the 6D convolution. As more frequency domain data is kept, the geometric details (e.g., pertaining to the small pairs of cylindrical pegs and holes depicted by their axis lines in panel (a)) start to emerge in the SDF shortly after $m'/m \approx 0.2\%$ in panel (d). An important observation is that the maximum score (i.e., minimum energy) configuration (denoted by B) does not change much even with very few number of frequencies in panels (b) and (c). However, the slopes and curvatures of the energy profile characterizing the forces/torques and the stiffness of combined physical and geometric constraints do change significantly. For example, for the uniaxial motion from A to B, filtering with $m'/m < 1\%$ results in a relaxed collision response and geometric guidance along BA with ‘soft’ snapping at B, as a result of faded geometric details. By increasing the precision with $m'/m \geq 1\%$, a brisker response is imposed by larger transverse slopes along the AB trajectory with ‘hard’ snapping at B due to sharper curvature. However, the changes are insignificant after $m'/m > 2\%$ with $\sigma = 0.5$, which enables speed-ups of two orders of magnitude by disposing of 98% of the frequency data. One needs to use smaller σ factors to capture more geometric details thus more meaningful frequencies, and larger $p = \lambda_2/\lambda_1$ to impart stronger collision response compared to collision-free geometric guidance.

4.3 Haptic Experiment

We test the method to assemble the pair of parts in Fig. 4, keeping Part 1 stationary and manipulating Part 2 using a SensAble® PHANTOM® Omni® device (with 6 DOF input, 3 DOF output). Because the device is not capable of torque feedback, we choose to restrict the motion to the translational space, due to the observation that force response is insufficient to create a satisfactory user experience when dealing with rotational constraints. We precompute the SDFs with $\sigma = 0.5$ for fine triangular meshes of $n_1 = 422,272$ and $n_2 = 320,384$, and a grid sample size of $m = 128^3 = 2,097,152$. The confinement of motion to translation only allows for exact precomputation of the convolution. However, after transferring to the Fourier domain, we keep only $m' = 16^3 = 4,096$ (less than 0.2% of m) dominant modes of the SDFs and zero-pad the rest. Figure 8 is the time plot of the haptic performance over a short assembly session, where the user explores colliding the objects from different sides and finally assembles them into the obvious configuration. The haptic performance is consistent, always maintaining the servo-loop rate of $1.00_{-0.15}^{+0.24}$ kHz not only during free motion, but also during collision and final insertion.

5 Conclusion

Haptic-enabled assembly planning has been restrained for a long time from achieving its full potential, due to the seemingly contradictory requirements of handling high geometric complexity and maintaining a response rate as fast as 1 kHz. We expanded upon our generic method in [13, 14] to develop a unified force model for collision response and geometric guidance that applies to arbitrary geometry, by formulating the interaction energies as a convolution of the so-called SDF descriptors. Although the generality of the method inevitably imparts additional computational complexity, we demonstrated that the guidance forces and torques can be efficiently computed in the Fourier domain, where the convolution converts to pointwise product of SDF amplitudes. We showed that the SDF shape descriptors and their Fourier expansions can be computed in a preprocessing step that takes a few minutes per part for reasonable mesh and sample sizes. For real-time computations, we showed that a very small subset of the dominant frequency domain data can be used to compute the pointwise multiplication followed by an inverse FFT. Our results confirm that such low-pass filtering of the SDF information, together with the computational power offered by the modern GPUs, enable fast evaluation of the geometric energies, forces, and torques within the available 1 millisecond time frame during haptic assembly, with very little compromise in accuracy. Unlike the existing approaches to collision detection or constraint management which are restricted by topological complexity (e.g., connectivity and number of holes), geometric complexity (e.g., convexity and type of surfaces), or syntactic complexity (e.g., number of triangles or voxels), our method does not impose any such restriction. Instead, it allows for a systematic trade-off between achieved fidelity and computational efficiency regardless of the input size or complexity.

The outcome of this research is a powerful paradigm that streamlines haptic assembly using spectral analysis of shape

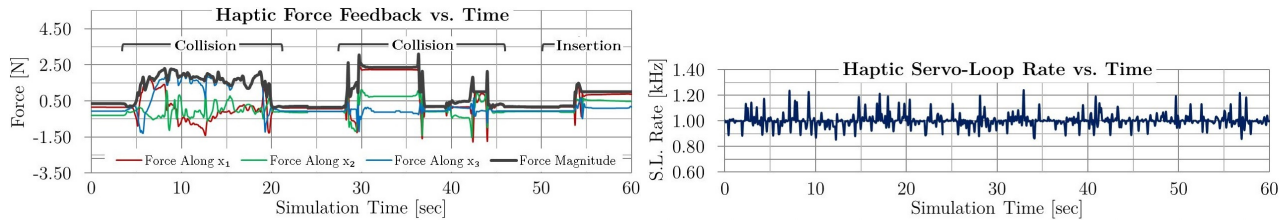


Fig. 8: Performance of a haptic assembly simulation.

descriptors in the Fourier domain, and opens up new promising theoretical and computational directions for VR researchers and haptics software developers.

6 Acknowledgment

This work was supported in part by the National Science Foundation grants CMMI-1200089, CMMI-0927105, and CNS-0927105.

References

- [1] Behandish, M., and Ilieş, H. T., 2015. “Haptic assembly using skeletal densities and Fourier transforms”. In Proceedings of the 2015 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE’2015), no. DETC2015-47923, American Society of Mechanical Engineers (ASME), pp. V01BT02A060:1–11.
- [2] Bullinger, H. J., Breining, R., and Bauer, W., 1999. “Virtual prototyping—state of the art in product design”. In 26th International Conference on Computers and Industrial Engineering, pp. 103–107.
- [3] Wang, G. G., 2002. “Definition and review of virtual prototyping”. *Journal of Computing and Information Science in Engineering*, **2**(3), pp. 232–236.
- [4] Deviprasad, T., and Kesavadas, T., 2003. “Virtual prototyping of assembly components using process modeling”. *Journal of Manufacturing Systems*, **22**(1), pp. 16–27.
- [5] Bordegoni, M., Colombo, G., and Formentini, L., 2006. “Haptic technologies for the conceptual and validation phases of product design”. *Computers and Graphics*, **30**(3), pp. 377–390.
- [6] Seth, A., Su, H. J., and Vance, J. M., 2006. “SHARP: A system for haptic assembly and realistic prototyping”. In Proceedings of the 2006 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE’2006), no. DETC2006-99476, American Society of Mechanical Engineers (ASME), pp. 905–912.
- [7] Seth, A., Su, H. J., and Vance, J. M., 2008. “Development of a dual-handed haptic assembly system: SHARP”. *Journal of Computing and Information Science in Engineering (JCISE)*, **8**(4), pp. 1–8.
- [8] Gomes de Sa, A., and Zachmann, G., 1999. “Virtual reality as a tool for verification of assembly and maintenance processes”. *Computers and Graphics*, **23**(3), pp. 389–403.
- [9] Volkov, S., and Vance, J. M., 2001. “Effectiveness of haptic sensation for the evaluation of virtual prototypes”. *Journal of Computing and Information Science in Engineering*, **1**(2), pp. 123–128.
- [10] Seth, A., Vance, J. M., and Oliver, J. H., 2011. “Virtual reality for assembly methods prototyping: A review”. *Virtual Reality*, **15**(1), pp. 5–20.
- [11] Vance, J. M., and Dumont, G., 2011. “A conceptual framework to support natural interaction for virtual assembly tasks”. In Proceedings of the 2011 ASME World Conference on Innovative Virtual Reality, no. WINVR2011-5570, American Society of Mechanical Engineers (ASME), pp. 273–278.
- [12] Perret, J., Kneschke, C., Vance, J. M., and Dumont, G., 2013. “Interactive assembly simulation with haptic feedback”. *Assembly Automation*, **33**(3), pp. 214–220.
- [13] Behandish, M., and Ilieş, H. T., 2014. “Peg-in-hole revisited: A generic force model for haptic assembly”. In Proceedings of the 2014 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE’2014), no. DETC2014-35290, American Society of Mechanical Engineers (ASME), pp. V01BT02A041:1–10.
- [14] Behandish, M., and Ilieş, H. T., 2015. “Peg-in-hole revisited: A generic force model for haptic assembly”. *Journal of Computing and Information Science in Engineering (JCISE)*, **15**(4), pp. 041004:1–11.
- [15] Lysenko, M., Nelaturi, S., and Shapiro, V., 2010. “Group morphology with convolution algebras”. In Proceedings of the 14th ACM Symposium on Solid and Physical Modeling (SPM’2010), pp. 11–22.
- [16] Lysenko, M., 2013. “Fourier collision detection”. *International Journal of Robotics Research*, **32**(4), pp. 483–503.

- [17] Behandish, M., and Ilies, H. T., 2015. "Analytic methods for geometric modeling via spherical decomposition". *Computer-Aided Design, Special Issue on the 2015 SIAM Conference on Geometric and Physical Modeling (GD/SPM'2015)*, **70**, pp. 100–115. 2015 SIAM/ACM Symposium on Solid and Physical Modling (GD/SPM'2015).
- [18] Roerdink, J. B. T. M., 2000. "Group morphology". *Pattern Recognition*, **33**(6), pp. 877–895.
- [19] Nelaturi, S., and Shapiro, V., 2011. "Configuration products and quotients in geometric modeling". *Computer-Aided Design*, **43**(7), pp. 781–794.
- [20] Behandish, M., and Ilies, H. T., 2014. Shape complementarity analysis for objects of arbitrary shape. Tech. Rep. CDL-TR-14-01, University of Connecticut.
- [21] Lozano-Perez, T., 1983. "Spatial planning: A configuration space approach". *IEEE Transactions on Computers*, **C-32**(2), pp. 108–120.
- [22] Kavraki, L. E., 1995. "Computation of configuration-space obstacles using the fast Fourier transform". *IEEE Transactions on Robotics and Automation*, **11**(3), pp. 408–413.
- [23] Chan, L. S. H., and Choi, K. S., 2009. "Integrating PhysX and OpenHaptics: Efficient force feedback generation using physics engine and haptic devices". In Proceedings of the 2009 Joint Conferences on Pervasive Computing (JCPC'2009), Institute of Electrical and Electronics Engineers (IEEE), pp. 853–858.
- [24] Sagardia, M., Stouraitis, T., and Silva, J. L. e., 2014. "A new fast and robust collision detection and force computation algorithm applied to the physics engine bullet: Method, integration, and evaluation". In Proceedings of the 2014 Conference and Exhibition of the European Association of Virtual and Augmented Reality (EuroVR'2014).
- [25] Renouf, M., Acary, V., and Dumont, G., 2005. "3D frictional contact and impact multibody dynamics. a comparison of algorithms suitable for real-time applications". In ECCOMAS Thematic Conference on Mutlibody Dynamics.
- [26] Tching, L., and Dumont, G., 2008. "Haptic simulations based on non-smooth dynamics for rigid-bodies". In Proceedings of the 15th ACM Symposium on Virtual Reality Software and Technology (VRST'2008), Association for Computing Machinery (ACM), pp. 87–90.
- [27] Lin, M., and Gottschalk, S., 1998. "Collision detection between geometric models: A survey". In Proceedings of the 1998 IMA Conference on Mathematics of Surfaces, Vol. 1, pp. 37–56.
- [28] Jimenez, P., Thomas, F., and Torras, 2001. "3D collision detection: A survey". *Computers and Graphics*, **25**(2), pp. 269–285.
- [29] Kockara, S., Halic, T., Iqbal, K., Bayrak, C., and Rowe, R., 2007. "Collision detection: A survey". In Proceedings of the 2007 IEEE International Conference on Systems, Man and Cybernetics (ISIC'2007), Institute of Electrical and Electronics Engineers (IEEE), pp. 4046–4051.
- [30] Hasegawa, S., and Fujii, N., 2003. "Real-time rigid body simulation based on volumetric penalty method". In Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Institute of Electrical and Electronics Engineers (IEEE), pp. 326–332.
- [31] Hasegawa, S., and Sato, M., 2004. "Real-time rigid body simulation for haptic interactions based on contact volume of polygonal objects". *Computer Graphics Forum*, **23**(3), pp. 529–538.
- [32] Lötstedt, P., 1984. "Numerical simulation of time-dependent contact and friction problems in rigid body mechanics". *SIAM Journal on Scientific and Statistical Computing*, **5**(2), pp. 370–393.
- [33] Stewart, D. E., 2000. "Rigid-body dynamics with friction and impact". *SIAM Review*, **42**(1), pp. 3–39.
- [34] Marcelino, L., Murray, N., and Fernando, T., 2003. "A constraint manager to support virtual maintainability". *Computers and Graphics*, **27**(1), pp. 19–26.
- [35] Murray, N., and Fernando, T., 2004. "An immersive assembly and maintenance simulation environment". In Proceedings of the 2004 IEEE International Symposium on Distributed Simulation and Real-Time Applications (DS-RT'2004), Institute of Electrical and Electronics Engineers (IEEE), pp. 159–166.
- [36] Jacob, R., Mitrouchev, P., and Leon, J. C., 2008. "Contact identification for assembly-disassembly simulation with a haptic device". *The Visual Computer*, **24**(11), pp. 973–979.
- [37] Jacob, R., Mitrouchev, P., and Leon, J. C., 2011. "Assembly simulation incorporating component mobility modelling based on functional surfaces". *International Journal on Interactive Design and Manufacturing*, **5**(2), pp. 119–132.
- [38] Boussuge, F., Léon, J. C., Hahmann, S., and Fine, L., 2012. "An analysis of DMU transformation requirements for structural assembly simulations". In Proceedings of the 8th International Conference on Engineering Computational Technology.
- [39] Mirtich, B., 1998. "V-Clip: Fast and robust polyhedral collision detection". *ACM Transactions on Graphics (TOG)*, **17**(3), pp. 177–208.
- [40] Ehmann, S. A., and Lin, M. C., 2000. "Accelerated proximity queries between convex polyhedra by multi-level Voronoi marching". In Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2000), Vol. 3, Institute of Electrical and Electronics Engineers (IEEE).
- [41] Ehmann, S. A., and Lin, M. C., 2001. "Accurate and fast proximity queries between polyhedra using convex surface decomposition". *Computer Graphics Forum*, **20**(3), pp. 500–511.
- [42] Gregory, A., Lin, M. C., Gottschalk, S., and Taylor, R., 1999. "A framework for fast and accurate collision detection

- for haptic interaction”. In ACM SIGGRAPH 2005 Courses, no. 34, Association for Computing Machinery (ACM), pp. 38–45.
- [43] Gottschalk, S., Lin, M. C., and Manocha, D., 1996. “OBBTree: A hierarchical structure for rapid interference detection”. In Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, pp. 171–180.
- [44] McNeely, W. A., Puterbaugh, K. D., and Troy, J. J., 2005. “Six degree-of-freedom haptic rendering using voxel sampling”. In Proceedings of the ACM SIGGRAPH’2005 Courses, no. 42, Association for Computing Machinery (ACM).
- [45] McNeely, W. A., Puterbaugh, K. D., and Troy, J. J., 2006. “Voxel-based 6-DOF haptic rendering improvements”. *Haptics-e*, **3**(7).
- [46] Barbič, J., and James, D., 2007. “Time-critical distributed contact for 6-DOF haptic rendering of adaptively sampled reduced deformable models”. In Proceedings of the ACM SIGGRAPH’2007/Eurographics Symposium on Computer Animation, Association for Computing Machinery (ACM), pp. 171–180.
- [47] Sagardia, M., Hulin, T., Preusche, C., and Hirzinger, 2008. “Improvements of the Voxmap-PointShell algorithm—fast generation of haptic data-structures”. In 53rd Internationales Wissenschaftliches Kolloquium Technische Universität Ilmenau.
- [48] Hubbard, P. M., 1996. “Approximating polyhedra with spheres for time-critical collision detection”. *ACM Transactions on Graphics (TOG)*, **15**(3), pp. 179–210.
- [49] Bradshaw, G., and O’Sullivan, C., 2004. “Adaptive medial-axis approximation for sphere-tree construction”. *ACM Transactions on Graphics (TOG)*, **23**(1), pp. 1–26.
- [50] Weller, R., and Zachmann, G., 2009. “Inner sphere trees for proximity and penetration queries”. In Proceedings of the 2009 Robotics: Science and Systems Conference (RSS’2009), Vol. 2.
- [51] Weller, R., and Zachmann, G., 2011. “Inner sphere trees and their application to collision detection”. In *Virtual Realities*, G. Brunnett, S. Coquillart, and G. Welch, eds. Springer Vienna, pp. 181–201.
- [52] Ruffaldi, E., Morris, D., Barbagli, F., Salisbury, K., and Bergamasco, M., 2008. “Voxel-based haptic rendering using implicit sphere trees”. In Proceedings of the 16th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Institute of Electrical and Electronics Engineers (IEEE), pp. 319–325.
- [53] Weller, R., and Zachmann, G., 2009. “A unified approach for physically-based simulations and haptic rendering”. In Proceedings of the ACM SIGGRAPH’2009 Symposium on Video Games, Association for Computing Machinery (ACM), pp. 151–159.
- [54] Weller, R., and Zachmann, G., 2009. “Stable 6-DOF haptic rendering with inner sphere trees”. In Proceedings of the 2009 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE’2009), no. DETC2009-87209, American Society of Mechanical Engineers (ASME), pp. 1591–1600.
- [55] Seth, A., Vance, J. M., and Oliver, J. H., 2007. “Combining geometric constraints with physics modeling for virtual assembly using SHARP”. In Proceedings of the 2007 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE’2007), no. DETC2007-34681, American Society of Mechanical Engineers (ASME), pp. 1045–1055.
- [56] Seth, A., Vance, J. M., and Oliver, J. H., 2010. “Combining dynamic modeling with geometric constraint management to support low clearance virtual manual assembly”. *Journal of Mechanical Design*, **132**(8), pp. 081002:1–7.
- [57] Jayaram, S., Jayaram, U., Wang, Y., and Tirumali, H., Lyons, K., and Hart, P., 1999. “VADE: A virtual assembly design environment”. *Journal of Computer Graphics and Applications*, **19**(6), pp. 44–50.
- [58] Wan, H., Gao, S., Peng, Q., Dai, G., and Zhang, F., 2004. “MIVAS: A multi-modal immersive virtual assembly system”. In Proceedings of the 2004 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE’2004), no. DETC2004-57660, American Society of Mechanical Engineers (ASME), pp. 113–122.
- [59] Tching, L., Dumont, G., and Perret, J., 2010. “Interactive simulation of CAD models assemblies using virtual constraint guidance”. *International Journal on Interactive Design and Manufacturing*, **4**(2), pp. 95–102.
- [60] Rosenberg, L. B., 1993. “Virtual fixtures: Perceptual tools for telerobotic manipulation”. In Proceedings of the 1993 IEEE Virtual Reality Annual International Symposium, Institute of Electrical and Electronics Engineers (IEEE), pp. 76–82.
- [61] Katznelson, Y., 2004. *An Introduction to Harmonic Analysis*, 3rd edition ed. Cambridge University Press.
- [62] Cooley, J. W., T. J., 1965. “An algorithm for the machine calculation of complex Fourier series”. *Mathematics of Computation*, **19**(90), pp. 297–301.
- [63] Requicha, A. G., 1980. Mathematical models of rigid solid objects. Production Automation Project, Technical Memo. No. 28, University of Rochester.
- [64] Chazal, F., and Soufflet, R., 2004. “Stability and finiteness properties of medial axis and skeleton”. *Journal of Dynamical and Control Systems*, **10**(2), pp. 149–170.
- [65] Klein, F., 2009. “A new approach to point membership classification in B-rep solids”. In *Mathematics of Surfaces XIII*.

Springer Berlin Heidelberg, pp. 235–250.

- [66] Requicha, A. G., 1980. Representations of rigid solid objects. Production Automation Project, Technical Memo. No. 29, University of Rochester.
- [67] Schöberl, J., 1997. “NETGEN: An advancing front 2D/3D-mesh generator based on abstract rules”. *Computing and Visualization in Science*, **1**(1), pp. 41–52.
- [68] Hoff III, K. E., Culver, T., Keyser, J., Lin, M., and Manocha, D., 1999. “Fast computation of generalized Voronoi diagrams using graphics hardware”. In Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, Association for Computing Machinery (ACM), pp. 277–286.
- [69] Schling, B., 2011. *The Boost C++ Libraries*. XML Press.
- [70] Frigo, M., and Johnson, S. G., 2005. “The design and implementation of FFTW3”. *Proceedings of the IEEE, Special issue on Program Generation, Optimization, and Platform Adaptation*, **93**(2), pp. 216–231.