

# Htex: Per-Halfedge Texturing for Arbitrary Mesh Topologies

WILHEM BARBIER, Unity Technologies, France

JONATHAN DUPUY, Unity Technologies, France




Fig. 1. A production asset rendered with Htex textures on an NVIDIA RTX 3090 GPU. Htex is a GPU-friendly alternative to Ptex and mesh colors that supports non-quad topologies. This frees artists from UV-authoring during asset creation, and yields high-quality seamless texture filtering. Due to its GPU-friendly nature, Htex allows for fast rendering times through look development, previzualization, and/or video-game prototyping.

We introduce per-halfedge texturing (Htex) a GPU-friendly method for texturing arbitrary polygon-meshes without an explicit parameterization. Htex builds upon the insight that halfedges encode an intrinsic triangulation for polygon meshes, where each halfedge spans a unique triangle with direct adjacency information. Rather than storing a separate texture per face of the input mesh as is done by previous parameterization-free texturing methods, Htex stores a square texture for each halfedge and its twin. We show that this simple change from face to halfedge induces two important properties for high performance parameterization-free texturing. First, Htex natively supports arbitrary polygons without requiring dedicated code for, e.g. non-quad faces. Second, Htex leads to a straightforward and efficient GPU implementation that uses only three texture-fetches per halfedge to produce continuous texturing across the entire mesh. We demonstrate the effectiveness of Htex by rendering production assets in real time.

CCS Concepts: • **Computing methodologies** → **Rendering; Texturing.**

Additional Key Words and Phrases: texturing, filtering, Ptex, GPU

Authors' addresses: [Wilhem Barbier](mailto:wilhem@wbrbr.org), Unity Technologies, France, [wilhem@wbrbr.org](mailto:wilhem@wbrbr.org); [Jonathan Dupuy](mailto:jonathan.dupuy@outlook.com), Unity Technologies, France, [jonathan.dupuy@outlook.com](mailto:jonathan.dupuy@outlook.com).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2022 Association for Computing Machinery.

2577-6193/2022/7-ART29 \$15.00

<https://doi.org/doi.org/10.1145/3543868>

### ACM Reference Format:

Wilhem Barbier and Jonathan Dupuy. 2022. Htex: Per-Halfedge Texturing for Arbitrary Mesh Topologies. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 3, Article 29 (July 2022), 14 pages. <https://doi.org/doi.org/10.1145/3543868>

## 1 INTRODUCTION

*Context.* Per-face texturing (Ptex) is a ubiquitous texturing method for 3D polygon meshes: it is used in offline production rendering as well as in content creation pipelines that depend on interactive modeling software [Burley and Laceywell 2008]. Compared to conventional texturing, Ptex completely frees artists from UV-authoring while providing seamless filtering at the same time, thus enabling faster prototyping and artistic iteration. A key issue however is that Ptex can not be fully accelerated on the GPU<sup>1</sup>. This can quickly hinder production pipelines because it becomes harder to guarantee interactive framerates as asset complexity grows through time for, e.g., modeling sessions or film pre-visualization.

*Motivation for Arbitrary Topologies.* Ptex’s GPU support is only partial in that existing implementations are restricted to quad-only meshes. While most assets tend to exhibit such topologies, many others will occasionally carry a few triangles and/or N-gons as shown in Figure 1. In theory, such non-quad faces could be removed via automated and/or manual re-meshing. However, enforcing topological edits would greatly diminish the benefits that Ptex brings in terms of rapid iteration times (since re-meshing can be as tedious as UV-authoring). We therefore argue in favor of an alternative texturing method capable of retaining the aforementioned benefits of Ptex while adding seamless support for arbitrary polygons in a GPU implementation. We introduce such a method here, which we refer to as per-halfedge texturing (Htex).

*Contributions and Outline.* We arrived at Htex by first identifying that Ptex extensively relies on topological information. Based on this observation, we looked into using halfedges as a main texturing entity (as opposed to faces in the case of Ptex). We found that halfedges encode an intrinsic quadrangulation of their polygon-mesh, which can be used for a Ptex-like implementation. We build upon these novel results to devise Htex and describe how it leads to an efficient GPU implementation. The remainder of this article is summarized in Figure 2 and organized as follows:

- In Section 3, we present the intrinsic meshes we build solely from halfedges.

<sup>1</sup>This is also true for its variant mesh colors [Yuksel et al. 2010] as we discuss in more detail in Section 2.




Fig. 2. Htex Overview. Htex builds upon (a) a halfedge mesh. We make the key observation that halfedges and halfedge/twin pairs respectively define (b) an intrinsic triangulation, and (c) an intrinsic quadrangulation of the mesh they belong to. The union of all halfedge/twin pairs and boundary halfedges thus yields a new intrinsic mesh that covers the entire surface of the original polygon mesh. Htex relies on (c) a simple parameterization for each of the intrinsic face spanned by this union and stores (d) a texture for each of them.

- In Section 4, we leverage these meshes to devise our novel texturing method, which stores a square texture for each halfedge/twin pair.
- In Section 5, we provide rendering results and performance measurements against conventional UV-mapping and a state-of-the-art Ptex GPU implementation.

## 2 RELATED WORK

In this section, we position our work with respect to the relevant literature. Note that we do not discuss all texture mapping methods here as this is out of the scope of this paper. We refer the interested reader to the thorough survey of Yuksel et al. [2019] for an overview of such methods.

*UV Atlas.* The most ubiquitous method for texturing a polygon mesh consists in mapping it onto a square domain. In turn, this square domain is discretized into a texture called an atlas or simply UV [Maillot et al. 1993], which can be edited by artists via any 2D image-processing tool. The main issue with this approach is that the UV map generally needs to satisfy an over-constrained set of problems. Such problems include (but are not restricted to) low distortion, symmetry optimizations, seam minimization, compliance with image-processing brushes and filters, etc. This makes it difficult to devise an automatic UV generation tool that always works. Consequently, UVs require manual tweaking by artists, which is a tedious and time-consuming process as mentioned in introduction. What is worse is that, in the general case, UV-mapping systematically produces seams that lead to texture-discontinuities on the mesh surface. Recently, Liu et al. [2017] introduced an optimization that post-processes a given texture to make its data continuous at the mesh's surface. However, we have found that this optimization leaves a few discontinuities. While this is fine for most textures, it makes it impossible to use UVs for displacement maps, as shown in Figure 6. Htex is free from such limitations and allows to do displacement mapping seamlessly.

*Ptex.* Ptex [Burley and Lacey 2008] addresses the issue of quad-mesh parameterization for texturing by assigning a texture per quad. In Ptex, texture filtering is achieved per-quad by systematically sampling its associated texture as well as those opposite to its edges. This approach produces seamless filtering within each quad as well as at its boundaries.

*Ptex on the GPU.* Several GPU implementations of Ptex have been proposed. Early implementations relied on GPUs that were limited in the number of simultaneous texture-reads they could perform. To alleviate this issue, Ptex textures can be tightly packed into a larger texture [Kim et al. 2011; McDonald and Burley 2011; Nießner and Loop 2013]. Unfortunately, this makes filtering challenging as filter taps can erroneously sample neighboring data, especially when anisotropic filtering is enabled. Borders can be used to mitigate the issue, but this never completely solves the problem and can add a significant memory overhead. More recent implementations no longer attempt to minimize the number of texture samplers and rely on texture arrays of bindless textures to sample directly from the Ptex textures. For filtering this data, McDonald introduces borderless Ptex [McDonald 2013]. Htex heavily builds upon this implementation so we discuss it more thoroughly in Section 4.4. For the sake of completeness, we also mention the implementation of Toth [2013], which avoids sampling neighboring quads by discarding filter taps and relying on MSAA to reconstruct a smooth signal. This approach produces good filtering quality with high MSAA factors, but does not allow texture magnification as opposed to borderless Ptex.

*Mesh Colors.* In the case of quad-only meshes, mesh colors [Yuksel et al. 2010] provides a variant of Ptex that uses the dual parameterization of Ptex. The advantage of this approach is that it alleviates the need to sample from the neighbors of each quad, because boundary data is duplicated by construction. However, it renders its GPU implementation more difficult as the required logic for performing trilinear and anisotropic filtering no longer maps to GPU texturing units [Yuksel 2017].

In order to overcome these issues, some well-identified modifications have to be made to existing hardware [Mallett et al. 2019, 2020]. In this work, we target maximum compatibility with current hardware and therefore avoid the dual parameterization of mesh colors entirely. Nevertheless, nothing prevents Htex to be used with the dual parameterization of mesh colors.

*Issues with Non-Quads.* A common limitation of existing GPU implementations of Ptex is that they do not support non-quad faces. This is because the way non-quad faces are handled is problematic for a GPU: Ptex splits them into sub-quads that require dedicated filtering code [Ptex 2022]. While this approach works fine on CPU-based implementations, it constitutes a challenge for a GPU-based implementation due to the algorithmic branching it fundamentally incurs. We argue that the main reason for this divergent behavior is due to the fact that Ptex was originally built for quad-only meshes. Before elaborating Htex, we considered non-quad support as mandatory. As a result, Htex natively supports non-quads on the GPU. The key that makes Htex work is that it does not store a texture per face of the polygon mesh, but on an intrinsic quadrangulation. As we will show in the next section, this quadrangulation arises naturally when using a halfedge mesh and does not need to be stored explicitly in memory.

*Volumetric Texturing.* Volumetric textures are an alternative method to per-face texture mapping methods, which use the 3D positions of the mesh surface as texture coordinates [Benson and Davis 2002; DeBry et al. 2002; Dolonius et al. 2020]. The key advantages of Htex over volumetric texturing are its efficiency and ease of implementation: As we will show, Htex maps to current GPU texture filtering hardware natively and does not require any other data than a halfedge mesh, which we present in the next paragraph. In contrast, volumetric textures require 3D (rather than 2D) logic for sampling along with a sparse data-structure for storage optimizations, both of which must be implemented in software. Another important limitation of volumetric textures over Htex is their current lack of support for anisotropic filtering.

*Halfedge Mesh Representation.* Htex builds upon a halfedge mesh. A halfedge mesh [Botsch et al. 2010; Kettner 1999; Weiler 1985] decomposes a polygon mesh into two sets: a set of vertex points and a set of halfedges. The vertex points encode the positional information for the mesh, while the halfedges—from which edges and faces emanate naturally—encode its topology. The halfedge are solely characterized by their operators TWIN, NEXT, PREV, VERT, EDGE, and FACE. In our implementation, we rely on a generalization of directed edges [Campagna et al. 1998] as our halfedge data-structure, which is illustrated in Figure 3 along with the aforementioned halfedge operators. Note that this same data-structure was used recently in the context of GPU-based subdivision [Dupuy and Vanhoy 2021].

### 3 INTRINSIC MESHING WITH HALFEDGES

In this section, we describe the key insight that Htex builds upon. Specifically, we show that halfedges encode an intrinsic triangulation of the polygon mesh they belong to (Section 3.1). We then show how to further pair such triangles to form an intrinsic quadrangulation over the polygon mesh (Section 3.2).

#### 3.1 Halfedges as an Intrinsic Triangulation

The intrinsic triangulation we introduce is shown in Figure 4 (a,b). Its properties are straightforward:

*A Triangle for Each Halfedge.* Each halfedge maps to exactly one triangle. For instance:

- the blue halfedge  $h_7$  in Figure 4 (a) maps to the blue triangle  $t_7$  in Figure 4 (b).
- the orange halfedge  $h_4$  in Figure 4 (a), maps to the orange triangle  $t_4$  in Figure 4 (b).

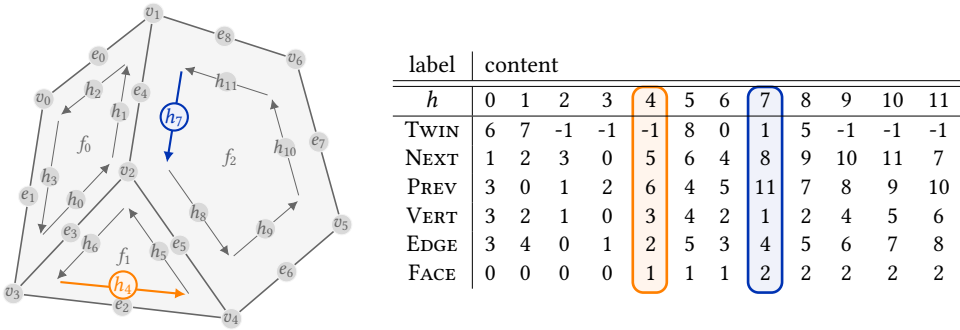


Fig. 3. Halfedge mesh representation used for the implementation of Htex. The blue and orange halfedges are highlighted as illustrative examples of a regular and border halfedge, respectively.

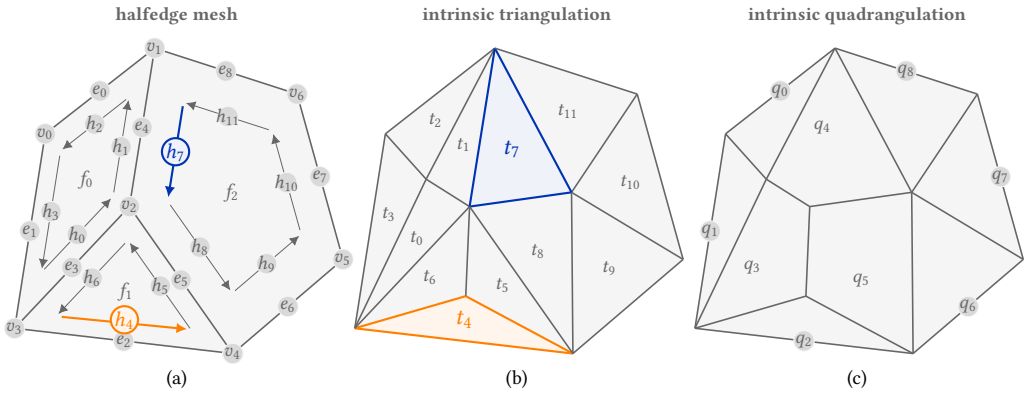


Fig. 4. Intrinsic triangulation and quadrangulation with halfedges.

*Triangle Vertices.* We retrieve the vertices of each triangle thanks to the halfedge operators NEXT and PREV. For instance, the vertices of the blue triangle  $t_7$  in Figure 4 (b) are

$$\begin{aligned} \text{the vertex } v_0 &:= \text{VERT}(\text{NEXT}(h_7)) = \text{VERT}(h_8), \\ \text{the average } v_1 &:= \frac{1}{5} (\text{VERT}(h_7) + \text{VERT}(h_8) + \dots + \text{VERT}(h_{11})), \\ \text{and the vertex } v_2 &:= \text{VERT}(h_7). \end{aligned}$$

Algorithm 1 provides pseudocode for computing these vertices.

*Triangle Adjacency.* We retrieve adjacent triangles thanks to the halfedge operators NEXT, PREV, and TWIN. For instance:

- the triangles adjacent to that of the blue triangle  $t_7$  in Figure 4 (b) are those formed by halfedges  $\text{PREV}(h_7) = h_{11}$ ,  $\text{NEXT}(h_7) = h_8$ , and  $\text{TWIN}(h_7) = h_1$ .
- the triangles adjacent to that of the orange triangle  $t_4$  in Figure 4 (b) are those formed by halfedges  $\text{PREV}(h_4) = h_6$ ,  $\text{NEXT}(h_4) = h_5$ , and  $\text{TWIN}(h_4) = -1$ , since  $h_4$  is a boundary halfedge.

Note that by construction, all triangles have exactly three neighbors except for those lying at the mesh boundary, which have exactly two. In the latter case, the neighbors are always given by the boundary halfedge's PREV and NEXT operators.

---

**Algorithm 1** Halfedge to triangle
 

---

```

1: function INTRINSICTRIANGLEVERTICES(halfedgeID: integer)
2:   nextID  $\leftarrow$  NEXT(halfedgeID)
3:    $v_0 \leftarrow$  VERT(nextID)
4:    $v_2 \leftarrow$  VERT(halfedgeID)
5:    $v_1 \leftarrow v_2$ 
6:    $n \leftarrow 1$ 
7:    $h \leftarrow$  nextID

8:   while  $h \neq$  halfedgeID do
9:      $v_1 \leftarrow v_1 +$  VERT( $h$ )
10:     $n \leftarrow n + 1$ 
11:     $h \leftarrow$  NEXT( $h$ )
12:   end while

13:   $v_1 \leftarrow v_1/n$ 

14:  return  $\langle v_0, v_1, v_2 \rangle$ 
15: end function

```

---

### 3.2 Building an Intrinsic Quadrangulation with Triangle Pairs

We now describe an intrinsic quadrangulation that builds upon the triangulation described in the previous subsection. We start by considering the case of boundary-free polygon-meshes, and then generalize our result for boundaries.

*Quads as Halfedge/Twin Pairs.* In the case of a boundary-free polygon mesh, we have  $E = 2H$ , where  $E$  and  $H$  respectively denote its number of edges and halfedges. This suggests that we can build a quadrangulation of  $E$  quads by creating unique pairs of adjacent triangles. In practice, we form an intrinsic quad by pairing the two triangles formed by a halfedge and its twin. The advantage of this approach is that it creates a bijection between the edges of the polygon mesh and its intrinsic quadrangulation. A practical consequence of this bijection is that each halfedge then maps to the intrinsic quad it spans using its EDGE operator. Our approach is illustrated in Figure 4, where we pair the blue triangle  $t_7$  with its twin  $t_1$  in (4, a) to yield the quad  $q_4$  in (4, c). From halfedge perspective, the halfedge pair  $h_1$  and  $h_7$  form the edge  $\text{EDGE}(h_7) = \text{EDGE}(h_1) = e_4$ , which maps to quad  $q_4$ . Notice that we purposely use consistent subscripts between the edges  $e$  in (4, a) and the quads  $q$  in (4, c) to emphasize their bijective relation.

*Boundaries.* We now generalize our intrinsic quadrangulation in the case where boundaries are present in the polygon-mesh. In such configurations, we still form a quad for each non-boundary edge. This leaves isolated triangles at the boundary as shown in Figure 4. We convert these isolated triangles into quads by pairing the boundary halfedges that span them with themselves. Intuitively, this is equivalent to placing a virtual halfedge on the other side of the boundary. We illustrate this approach using Figure 4: we build the quad  $q_2$  in (4, c) by pairing the orange halfedge  $h_4$  in (4, a) with itself. Again, we emphasize the bijection between edges  $e$  in (4, a) and quads  $q$  in (4, c) by using consistent subscripts. Our construction thus forms  $E \in [H, 2H)$  intrinsic quads for any

non-manifold mesh. Furthermore, each halfedge still maps to the intrinsic quad it spans via its `EDGE` operator. For instance, the orange boundary halfedge  $h_4$  in (4, a) maps to the quad  $q_2$  in (4, c) via its edge  $e_2 = \text{EDGE}(h_4)$ .

## 4 PER-HALFEDGE TEXTURING

In this section, we leverage the intrinsic mesh we derived in Section 3 to devise Htex. We start by a high-level description of Htex (Section 4.1). Next, we build our mesh parameterization (Section 4.2). We then further describe how we store and filter texture data (Section 4.3). Finally, we describe our GPU implementation, which is based on OpenGL (Section 4.4).

### 4.1 High-Level Description

Htex stores a separate texture per intrinsic quad of the halfedge mesh, each of which can be independently sized. Since we are now dealing with quad-only topologies, we can use the same approach as Ptex for texture filtering. However, Htex offers a few additional advantages. For filtering, we look-up adjacent quads similarly to Ptex. Since all adjacency information is provided directly by the halfedges of the mesh, Htex only stores texel data without any further adjacent information. For rendering, we draw the intrinsic triangulation, i.e., a triangle for each halfedge rather than the actual polygons of the mesh. Note that our triangulation does not alter the shape of the original polygonal mesh in any way due to its construction. We then perform texture lookup using halfedge operators, which tell us which textures to lookup in memory for seamless filtering.

### 4.2 Parameterization

Htex stores a texture for each of our intrinsic quad and we derive here the parameterization that tells how to map texels onto the actual surface of the polygon mesh. To this end, we use the fact that, by construction, each quad can be triangulated into the two intrinsic triangles formed by a halfedge and its twin (or itself for a boundary halfedge). We therefore parameterize each quad as the union of these two intrinsic triangles as follows:

*Intrinsic Triangle Parameterization.* We start with a parameterization for our intrinsic triangles using barycentric-coordinate space  $(u, v, w)$ . We configure this space such that:

- (i)  $(u, v, w) = (1, 0, 0)$  maps to the vertex point  $v_2 = \text{VERT}(h)$ , and
- (ii)  $(u, v, w) = (0, 1, 0)$  maps to the vertex point  $v_0 = \text{VERT}(\text{NEXT}(h))$ ,

where  $h$  denotes the halfedge that spans the intrinsic triangle.

*Retrieving the Quad Texture Coordinates.* We define our quad parameterization as a split into that of the two intrinsic triangles that compose it along the diagonal  $(1,0) - (0,1)$  in  $uv$ -space. We determine whether to map a triangle on the upper or lower part of this split using the fact that we represent halfedges as integers: We map the halfedge with largest integer value to the lower left part of the quad. Our approach is illustrated in Figure 2 (c), where each frame is located at the origin of the quad's parametric domain. Thanks to this construction, we can map the parameterization of an intrinsic triangle into that of a quad. Algorithm 2 provides pseudocode for computing such a mapping given a halfedge and the barycentric coordinates of its associated intrinsic triangle.

*Rendering the Intrinsic Triangulation.* Our parameterization allows to render a polygon-mesh using its intrinsic triangulation. In order to texture these triangles, we only need the halfedge that spans it as it provides all the operators to retrieve its associated texel data.

---

**Algorithm 2** Intrinsic Triangle to Texture Coordinates
 

---

```

1: function TRIANGLETOQUADUV( halfedgeID: integer,  $\langle u, v \rangle$ : barycentric coordinates )
2:   if halfedgeID > TWIN(halfedgeID) then
3:     return  $\langle u, v \rangle$ 
4:   else
5:     return  $\langle 1 - u, 1 - v \rangle$ 
6:   end if
7: end function

```

---

### 4.3 Storage and Filtering

*Texture Data.* We store texture data as regular textures, just like Ptex. In the case of boundary halfedges, our construction only fills half of the texture, effectively resulting in wasted space at mesh boundaries. Fortunately, boundaries are very sparse within typical polygon meshes. As such, we did not make any effort to avoid this waste. Rather, we mirror the texel content of the existing halfedge as shown in Figure 2 (d) (see the orange inset), which is useful for the seamless filtering approach we introduce next. In terms of texture count, Htex stores exactly  $E$  textures, where  $E$  denotes the number of edges on a mesh. Compared to Ptex in the case of a quad-only mesh, Htex thus stores a factor of  $E/F$  more textures, where  $F$  denotes the number of quads in the mesh. For non-quad faces, Htex stores as many textures as Ptex.

*Filtering.* Our filtering approach works with standard GPU texture filtering units. In essence, it is similar to that of Ptex as it requires fetching the texel content of neighboring faces. An important distinction however is that, since we render the intrinsic triangulation rather than the actual polygon mesh, we only need to deal with 3 neighbors (rather than 4 for Ptex). As explained in Section 3, we retrieve these neighbors using the halfedge operators NEXT, PREV, and TWIN. Furthermore, since we store halfedges and their twins together in the same texture, sampling from the intrinsic quad already takes care of one of the neighbors. This leaves two extra lookups for PREV and NEXT. Htex thus provides the same filtering capabilities as Ptex but at the cost of 3 texture fetches rather than 5.

### 4.4 GPU Implementation

Our GPU implementation of Htex is very similar to borderless Ptex [McDonald 2013]. We provide more details in the following paragraphs.

*Texture and Sampler Initialization.* We store each texture in a separate texture object that we manipulate using the bindless texture extension GL\_ARB\_bindless\_texture. We sample these textures in the same fashion as borderless Ptex, which requires two main steps:

- (1) set the sampler of each texture to CLAMP\_TO\_BORDER with a border color of zero.
- (2) add an extra border-normalization channel per texture resolution with all texels set to one.

*Rendering.* During rendering, we render each intrinsic triangle associated with a halfedge of the polygon mesh. For texture sampling, we proceed as shown in Algorithm 3: Each triangle accumulates the sample of its texture with those associated with the halfedges NEXT and PREV. For each of these three texture fetches, we also fetch the border-normalization channel and use it as normalization. An important step of Algorithm 3 is to convert the barycentric coordinates of the current triangles to its two neighbors. Borderless Ptex accomplishes this by precomputing 16 transformation matrices. In our case, our intrinsic construction makes the transformation obvious and simply consist in reflecting the coordinates (see lines 7 and 9 in Algorithm 3).

*Corner Pre-processing.* As borderless Ptex, Htex produces continuous filtering everywhere except at corner texels. In practice, this discontinuity is un-noticeable for most texture data. A notable exception is displacement mapping as the discontinuity results in cracks. For this case we follow previous methods [McDonald and Burley 2011; Nießner and Loop 2013] and pre-process texel corners to enforce that they all have the same value. Note that this produces continuous filtering whenever texture have the same resolution. In the case of varying texture resolutions, we have observed that a few cracks may still appear due to precision issues within GPU texture samplers.

---

**Algorithm 3** Htex Sampling on the GPU
 

---

```

1: function HTEXTURE(halfedgeID: integer,  $\langle u, v \rangle$ : barycentric coordinates)
2:   nextID  $\leftarrow$  NEXT(halfedgeID)
3:   prevID  $\leftarrow$  PREV(halfedgeID)
4:    $c = \langle c_0, \dots, c_n \rangle \leftarrow 0$  ▷ we have  $n$  channels plus the border-normalization channel
5:    $\langle x, y \rangle \leftarrow$  TRIANGLETOQUADUV(halfedgeID,  $\langle u, v \rangle$ )
6:    $c \leftarrow c +$  TEXTURE(EDGE(halfedgeID),  $\langle x, y \rangle$ )
7:    $\langle x, y \rangle \leftarrow$  TRIANGLETOQUADUV(nextID,  $\langle v, -u \rangle$ )
8:    $c \leftarrow c +$  TEXTURE(EDGE(nextID),  $\langle x, y \rangle$ )
9:    $\langle x, y \rangle \leftarrow$  TRIANGLETOQUADUV(prevID,  $\langle -v, u \rangle$ )
10:   $c \leftarrow c +$  TEXTURE(EDGE(prevID),  $\langle x, y \rangle$ )
11:  return  $\langle \frac{c_0}{c_n}, \dots, \frac{c_{n-1}}{c_n} \rangle$ 
12: end function

```

---

## 5 RESULTS AND EVALUATION

In this section, we evaluate the performances of Htex. We first showcase some renderings to evaluate the filtering quality and scalability of Htex (Section 5.1). We then provide some performance measurements and compare them against borderless Ptex and classic UV-mapping (Section 5.2). Finally, we discuss some limitations of Htex (Section 5.3).

### 5.1 Rendering Results

*Anisotropic Filtering.* Htex maps seamlessly to the GPU's hardware accelerated anisotropic filter. We demonstrate this in Figure 5, where we render an Htex asset with varying anisotropic filtering factors. We generated this asset by taking direct inspiration from the original Ptex article, see [Burley and Lacewell 2008], Figure (10). To do so, we applied the same texture with tri-planar mapping on the original mesh, which is available on the Ptex website. Notice how anisotropic filtering improves the quality of the texture at grazing angles. To further demonstrate Htex's filtering quality, we provide an animated rendering of the same asset in our supplemental video.

*Displacement Mapping.* Htex allows for crack-free displacement maps as illustrated in Figure 6. We emphasize that crack-free displacement is impossible to achieve with UV-maps as seams will systematically result in texture discontinuities. Furthermore, we mention that the asset used for this particular example is similar in nature to the one shown in Figure 1: it is composed of mostly quads (1494 in total), but also a few triangles (38 in total). Htex supports it effortlessly. Note that we further demonstrate the benefits of Htex with respect to UV-maps in the context of GPU-accelerated displacement mapping in our supplemental video.

*Scalability.* To demonstrate the scalability of Htex, we render production assets in Figure 1 and Figure 7. The asset from Figure 1 consists of 10634 quads, 436 triangles, and 562 boundary edges. The asset from Figure 7 is a quad-only Ptex asset with 5812 faces and 572 boundary edges. To

produce the rendering in Figure 7, we resampled the Ptex textures into Htex ones, which consist of displacement and albedo data. We provide an animated zoom-in sequence in our supplemental video that relies on adaptive tessellation to render the asset in real time. Note that we displace the mesh along the normals of the input mesh, rather than from those of the limit surface of Catmull-Clark subdivision. We plan to add subdivision in our implementation, but note that this is orthogonal to the problem of texturing, which Htex addresses.

## 5.2 Performances

*Methodology.* We position Htex’s performances against our implementation of borderless Ptex [McDonald 2013] and conventional UVs. Our performance measurements correspond to the median rendering timing over 100 successive frames rendered at  $1920 \times 1080$  resolution. Our GPU rendering pipeline uses tessellation shaders and we measure its performances with and without sampling a displacement map. The fragment shader simply samples an extra RGB texture. The rendering involves a single asset, which we draw so that it occupies the entire framebuffer. Our test machine has an NVIDIA RTX 3090 GPU and an Intel i9-10980XE CPU with 128GiB of RAM.

Table 1. Htex performance comparison against UVs on an NVIDIA RTX 3090.

	RGB only	RGB + disp
UVs	0.40 ms	0.59 ms
Htex	0.42 ms	0.68 ms
speedup	$\times 0.95$	$\times 0.87$

*Comparison Against UV-maps.* We report our timings for Htex’s performance against UVs in Table 1. We performed our measurement using the asset from Figure 1 for which we also have a UV map. As demonstrated by the reported numbers, Htex’s overhead is almost negligible when displacement is disabled. As for when displacement is enabled, Htex’s overhead becomes slightly more important. We attribute this increased difference to the fact that texture fetches during vertex processing do not necessarily pipeline as well as in fragment shaders and cache misses are more frequent. We finally mention that the displaced, UV-based configuration is purely synthetic as UVs will systematically produce cracks on the final surface. As such, these numbers should rather be seen as a baseline.

Table 2. Htex performance comparison against borderless Ptex [McDonald 2013] on an NVIDIA RTX 3090.

	RGB only	RGB + disp
Ptex	0.9 ms	2.4 ms
Htex	0.77 ms	1.7 ms
speedup	$\times 1.17$	$\times 1.41$

*Comparison Against Borderless Ptex.* We report our timings for Htex’s performance against Ptex in Table 2. We performed our measurement using the asset from Figure 7. As demonstrated by the reported numbers, Htex provides a solid speedup over our borderless Ptex implementation. This is due to the fact that Htex requires less texture taps and memory accesses as we do not require fetching adjacency information or UV transformation matrices. We also re-emphasize that our implementation supports arbitrary polygons rather than just quads in the case of borderless Ptex.




Fig. 5. Hardware anisotropic filtering with Htex.




Fig. 6. Displacement mapping with Htex. Htex allows for artifact-free displacement mapping on the GPU while UV maps will systematically produce cracks (highlighted in red in insets) at UV-seams. Cracks can be greatly reduced by the method of Liu et al. [2017] although never entirely (bottom inset).

*Memory Consumption.* The memory consumption of Htex is a function of the number of texels stored on the surface of the polygon mesh. Just like borderless Ptex, we also require an extra border-normalization channel per texture resolution for filtering. We mention that, in practice, the overhead incurred by this extra channel quickly becomes negligible as the number of channels increases. Moreover, we believe that the need for an extra border-normalization channel could be removed entirely by having GPU texture samplers return the blending value of each border. It remains to be discussed whether this is a reasonable possibility.

### 5.3 Limitations

*Subpixel-Face-Filtering.* Fundamentally, Htex remains a per-face texturing solution. As such, it shares the same limitations as other per-face methods like Ptex and mesh colors. One such limitation is the fact that both texture and geometry alias at the same resolution. This means that Htex is unable to filter texture data below the resolution of the intrinsic quadrangulation of its input mesh. We illustrate this issue in Figure 8, where we render a highly tessellated sphere with Htex. Notice the aliasing at the silhouette of the sphere, which is especially visible for high-frequency textures.

*Low-Quality Meshes.* Another limitation common to per-face texturing methods is that parameterization quality directly depends on that of the input polygon-mesh. If the input polygon-mesh




Fig. 7. Ptex production asset (Ptex T-rex model ©Walt Disney Animation Studios.) rendered with Htex and parameterization comparison.




Fig. 8. Htex texture aliasing. Htex aliases at the same resolution as its intrinsic quadrangulation.

carries, e.g., highly elongated faces, then so will the intrinsic quadrangulation and hence our parameterization will produce highly anisotropic texels. Note that per-face texturing methods can somewhat minimize this effect through their support for arbitrary-sized textures, although this will hardly be as good as having a high-quality mesh from the start.

*Non-Manifold Meshes.* Since Htex builds upon halfedge meshes, it is fundamentally limited to manifold meshes because halfedges can not describe non-manifold topologies. Hence, non-manifold meshes can not be used with Htex. We believe that a way to extend support to non-manifold topologies would be to rely on a mesh-matrix data-structure [DiCarlo et al. 2014; Mahmoud et al. 2021; Zayer et al. 2017], but we have not looked into this.

## 6 CONCLUSION

Htex is a novel parameterization-free texturing method that can be implemented on modern GPU. The key advantage of Htex over Ptex and mesh colors is that it supports non-quad topologies seamlessly thanks to the powerful properties of halfedges. While Htex is probably too costly to be used for an entire video game, we believe it could be used for assets that can benefit from crack-free displacement mapping with tessellation shaders. In future work, we would like to combine our method with halfedge-based subdivision algorithms [Dupuy and Vanhoy 2021] to render Ptex-based assets more faithfully within a unified, halfedge based, framework.

## ACKNOWLEDGMENTS

The T-Rex asset from Figure 1 was modeled by Olivier Roos for Studio Manette. We thank Ignacio Llamas and Ignacio Castaño for sharing with us the asset from Figure 6. Finally, we thank Brent Burley and Disney for giving access to the T-rex model from Figure 7.

## REFERENCES

- David Benson and Joel Davis. 2002. Octree Textures. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques* (San Antonio, Texas) (SIGGRAPH '02). Association for Computing Machinery, New York, NY, USA, 785–790. <https://doi.org/10.1145/566570.566652>
- Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Lévy. 2010. *Polygon Mesh Processing*. AK Peters / CRC Press, 250 pages. <https://hal.inria.fr/inria-00538098>
- Brent Burley and Dylan Lacewell. 2008. Ptex: Per-Face Texture Mapping for Production Rendering. *Computer Graphics Forum* 27, 4 (2008), 1155–1164. <https://doi.org/10.1111/j.1467-8659.2008.01253.x>
- Swen Campagna, Leif Kobbelt, and Hans-Peter Seidel. 1998. Directed Edges—A Scalable Representation for Triangle Meshes. *Journal of Graphics Tools* 3, 4 (1998), 1–11. <https://doi.org/10.1080/10867651.1998.10487494>
- David (grue) DeBry, Jonathan Gibbs, Deborah DeLeon Petty, and Nate Robins. 2002. Painting and Rendering Textures on Unparameterized Models. *ACM Trans. Graph.* 21, 3 (jul 2002), 763–768. <https://doi.org/10.1145/566654.566649>
- Antonio DiCarlo, Alberto Paoluzzi, and Vadim Shapiro. 2014. Linear algebraic representation for topological structures. *Computer-Aided Design* 46 (2014), 269–274. <https://doi.org/10.1016/j.cad.2013.08.044> 2013 SIAM Conference on Geometric and Physical Modeling.
- D. Dolonius, E. Sintorn, and U. Assarsson. 2020. UV-free Texturing using Sparse Voxel DAGs. *Computer Graphics Forum* 39, 2 (2020), 121–132. <https://doi.org/10.1111/cgf.13917> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13917>
- Jonathan Dupuy and Kenneth Vanhoy. 2021. A Halfedge Refinement Rule for Parallel Catmull-Clark Subdivision. *Computer Graphics Forum* (2021). <https://doi.org/10.1111/cgf.14381>
- Lutz Kettner. 1999. Using generic programming for designing a data structure for polyhedral surfaces. *Computational Geometry* 13, 1 (1999), 65 – 90. [https://doi.org/10.1016/S0925-7721\(99\)00007-3](https://doi.org/10.1016/S0925-7721(99)00007-3)
- Sujeong Kim, Karl Hillesland, and Justin Hensley. 2011. A Space-Efficient and Hardware-Friendly Implementation of Ptex. In *SIGGRAPH Asia 2011 Sketches* (Hong Kong, China) (SA '11). Association for Computing Machinery, New York, NY, USA, Article 31, 2 pages. <https://doi.org/10.1145/2077378.2077417>
- Songrun Liu, Zachary Ferguson, Alec Jacobson, and Yotam Gingold. 2017. Seamless: Seam erasure and seam-aware decoupling of shape from mesh resolution. *ACM Transactions on Graphics (TOG)* 36, 6, Article 216 (Nov. 2017), 15 pages. <https://doi.org/10.1145/3130800.3130897>
- Ahmed H. Mahmoud, Serban D. Porumbescu, and John D. Owens. 2021. RXMesh: A GPU Mesh Data Structure. *ACM Trans. Graph.* 40, 4, Article 104 (jul 2021), 16 pages. <https://doi.org/10.1145/3450626.3459748>
- Jérôme Maillot, Hussein Yahia, and Anne Verroust. 1993. Interactive Texture Mapping. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (Anaheim, CA) (SIGGRAPH '93). Association for Computing Machinery, New York, NY, USA, 27–34. <https://doi.org/10.1145/166117.166120>
- Ian Mallett, Larry Seiler, and Cem Yuksel. 2019. Patch Textures: Hardware Implementation of Mesh Colors. In *High-Performance Graphics (HPG 2019)* (Strasbourg, France). The Eurographics Association. <https://doi.org/10.2312/hpg.20191194>
- Ian Mallett, Larry Seiler, and Cem Yuksel. 2020. Patch Textures: Hardware Support for Mesh Colors. *IEEE Transactions on Visualization and Computer Graphics* (2020), 12 pages. <https://doi.org/10.1109/TVCG.2020.3039777>
- John McDonald. 2013. Eliminating Texture Waste: Borderless Realtime Ptex. Retrieved April 30, 2022 from <https://www.gdcvault.com/play/1017757/Eliminating-Texture-Waste-Borderless-Realtime>
- John McDonald and Brent Burley. 2011. Per-Face Texture Mapping for Real-Time Rendering. In *ACM SIGGRAPH 2011 Talks* (SIGGRAPH '11). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/2037826.2037840>
- Matthias Nießner and Charles Loop. 2013. Analytic Displacement Mapping Using Hardware Tessellation. *ACM Trans. Graph.* 32, 3, Article 26 (jul 2013), 9 pages. <https://doi.org/10.1145/2487228.2487234>
- Ptex. 2022. *Ptex documentation: Adjacency data*. Retrieved April 30, 2022 from <https://ptex.us/adjdata.html>
- Robert Toth. 2013. Avoiding Texture Seams by Discarding Filter Taps. *Journal of Computer Graphics Techniques (JCGT)* 2, 2 (6 December 2013), 91–104. <http://jcg.t.org/published/0002/02/07/>
- K. Weiler. 1985. Edge-Based Data Structures for Solid Modeling in Curved-Surface Environments. *IEEE Computer Graphics and Applications* 5, 1 (1985), 21–40. <https://doi.org/10.1109/MCG.1985.276271>
- Cem Yuksel. 2017. Mesh Color Textures. In *High-Performance Graphics (HPG 2017)* (Los Angeles, CA). ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3105762.3105780>

- Cem Yuksel, John Keyser, and Donald H. House. 2010. Mesh colors. *ACM Transactions on Graphics* 29, 2, Article 15 (2010), 11 pages. <https://doi.org/10.1145/1731047.1731053>
- Cem Yuksel, Sylvain Lefebvre, and Marco Tarini. 2019. Rethinking Texture Mapping. *Computer Graphics Forum (Proceedings of Eurographics 2019)* 38, 2 (2019), 535–551. <https://doi.org/10.1111/cgf.13656>
- Rhaleb Zayer, Markus Steinberger, and Hans-Peter Seidel. 2017. A GPU-Adapted Structure for Unstructured Grids. *Comput. Graph. Forum* 36, 2 (may 2017), 495–507. <https://doi.org/10.1111/cgf.13144>