AlSub: Fully Parallel Subdivision for Modeling and Rendering

DANIEL MLAKAR, Graz University of Technology
MARTIN WINTER, Graz University of Technology
HANS-PETER SEIDEL, Max Planck Institute for Informatics
MARKUS STEINBERGER, Graz University of Technology
RHALEB ZAYER, Max Planck Institute for Informatics

Mesh subdivision is a key geometric modeling task which forges smooth, supple, eye-pleasing surfaces out of coarse polygonal outlines. Since its rise to mainstream around the turn of the century, subdivision has become an unavoidable production tool. With industrial demands in sight, there has been a steady effort for developing faster and more efficient subdivision implementations. Despite the tremendous parallelism potential of subdivision algorithms, state-of-the-art implementations are only partially parallel as they are riddled by intermediate serial steps and therefore fail to unleash the compute power of massively parallel devices such as graphics processing units (GPUs).

To fully parallelize the subdivision process, we discard traditional linked list data structures in favor of a sparse matrix linear algebra formalism. Subdivision algorithms are written in the language of linear algebra with customized operators which readily demonstrate a performance edge over existing approaches. To further increase performance, we automatically identify critical matrix operations and replace them by specialized, heavily tuned GPU kernels. To substantiate the versatility of our approach we apply it to $\sqrt{3}$, Loop and Catmull-Clark subdivision schemes and show support for adaptive subdivision, semi-sharp creases, and a split evaluation scheme that separates topology and topological changes from positional updates. Our results indicate substantial performance gains over the state-of-the-art and current industry standard.

1 INTRODUCTION

Mesh subdivision is an important free-form modeling method. The input mesh undergoes a series of averaging, splitting, and fitting of surface patches. The careful design of these steps guarantees a denser and smoother output mesh enjoying highly desirable attributes, as shown in Figure 1. The process bears some similarity to early ideas in surface fitting in finite element analysis [Clough and Tocher 1965] and numerical approximation [Powell and Sabin 1977] and it has been honed for geometric modeling through the concerted effort of several pioneering researchers, *e.g.*, Chaikin [1974], Doo [1978], Doo and Sabin [1978a], and Catmull and Clark [1978]. With the adaption of mesh subdivision in animated feature films [DeRose et al. 1998], it became a standard modeling tool in the production pipeline. To date, it continues to be a highly active research topic given the ever increasing demand for realtime performance [Brainerd et al. 2016]. While not too long ago, sequential algorithms were simply running faster, as the CPU doubled its speed every other

Fig. 1. The control mesh of the ArmorGuy (courtesy of DigitalFish) subdivision model consists of 9k faces and 10k vertices and features a considerable number of creases. Using our approach, the refined mesh at level six (35M faces, 35M vertices) can be computed in 40ms without any preprocessing.

year [Moore 2000], that expectation is no longer true [Sutter 2005]. Apparently, the only available possibility to increase performance is efficient algorithmic parallelization that is capable of harnessing the compute power of modern architectures. Although mesh subdivision seems predestined for parallel execution, vectorization of the complete process remains a challenge; especially on the single instruction, multiple data (SIMD) compute model of the graphics processing unit (GPU). Up to date, the sheer processing power delivered by such hardware remains under-exploited, as even the current industry standard performs expensive serial preprocessing. To support fast adjacency queries, serial subdivision implementations have traditionally relied on mesh representations based on linked lists, *e.g.*, winged-edge representations [Baumgart 1972]. Changes to the topology in such representations requires careful pointer updates to preserve consistency. Computations and dynamic changes in the local neighborhood of mesh vertices—which are essential in subdivision—require pointer chasing, which leads do unbalanced

workload distributions and scattered memory accesses. Both hurt performance on modern parallel devices like the GPU. Although mesh representations for parallel subdivision exist, none supports efficient parallelization of the complete process. The idea of splitting a mesh into patches which can be subdivided independently seems appealing for parallelization at first. They can be refined recursively [Shiue et al. 2005] or be directly evaluated over a regular grid using precomputed basis functions [Bolz and Schröder 2002, 2003]. As these basis functions depend on the patch's topology, one or two preceeding subdivision iterations are required to enforce isolation of irregularities and decrease the number of function tables that have to be stored to a feasible amount. Also, special care has to be taken to avoid cracks in the result and get a watertight mesh. Patch-based approaches introduce significant data redundancy and computational overhead, as patches need to overlap to avoid interdependence.

Alternatively, evaluation of the refined mesh from the control mesh can be performed using pre-computed subdivision or stencil tables [Nießner et al. 2012; Pixar 2017]. This approach allows a parallel evaluation of the subdivided mesh—even if the control mesh vertices are animated—but does not solve the challenge of parallelization, as the cost of table creation amounts to a complete subdivision and is usually done on the CPU. The stencil tables need to be generated anew whenever the mesh topology changes, *i.e.*, when modeling operations are applied. This introduces unpleasant delays in the model creation process.

Due to these challenges, there are various tools and approximations used at different stages of modeling and rendering, potentially leading to different results. Furthermore, they all come with various limitations, *e.g.*, OpenSubdiv limits the number of vertices placed in regular patches to allow for efficient hardware supported rendering.

Our work asks the questions: *What if there was a unified, instantaneous, completely parallel subdivision approach that can be used for modeling and rendering? What if this method simply followed the traditional recursive manner of subdivision, naturally supporting regular and irregular patch types, creases, displacement mapping, and frequent topological changes? What if on top of all that, topological operations can be factored out for even faster evaluation during production rendering of animated subdivision meshes?*

We believe such an approach could be valuable throughout all areas using subdivision surfaces. It can be used as a drop-in replacement for OpenSubdiv, allowing the generation of subdivision surfaces instantaneously during modeling. It can be used to generate highly detailed, zoomed-in views of the limit surface without any limitations on the number of subdivision applied. It can be used to selectively subdivide parts of meshes during ray and path tracing without the need to keep preprocessed or temporary data around. It can be used to generate uniform subdivision for CAD models as preparation for simulations. Finally, such a unified approach can guarantee that the detailed view seen by a modeler matches the final production rendering geometry and much more.

In this paper we propose such an approach: With Algebra Subdivision, short *AlSub*, we recast mesh subdivision into linear algebra operations and introduce *the first complete GPU-enabled, universally applicable subdivision implementation.* This formalization allows channeling the vectorization burden to linear algebra kernels and decouples it from the intrinsic problem settings, thus, implying the versatility of our approach. Starting from a sparse matrix representation of meshes [Zayer et al. 2017], we make the following contributions:

- We show that with few linear algebra operations optimized for mesh-processing, the entire subdivision process can be described in a compact, self-contained manner suitable for execution on *massively parallel devices* like the GPU.
- We show that our sparse linear algebra formalization is *sufficiently general* to describe many existing subdivision schemes such as $\sqrt{3}$, Loop and Catmull-Clark.
- We show that the proposed approach *can be easily extended* to support additions to the standard subdivision algorithms, such as sharp and semi sharp creases, displacement mapping and subdivision of selected regions, *e.g.*, for feature adaptiveness or path tracing.
- We show that the topological operations can be separated from position refinement steps in such a linear algebra approach, allowing for even more efficient evaluation during animation of meshes.
- We show that not only efficient formulations of the algorithmic steps are of high relevance, but also data locality captured in the sparse matrix representation.

We start with a brief discussion of related work (Section 2), and then walk through the classical Catmull-Clark subdivision algorithm illustrating how relevant operations can be reformulated in the language of linear algebra (Section 3). To demonstrate the flexibility of our methodology, we show how the basic formulation can be extended to handle (semi-)sharp creases and how only regions of interest on the surface can be evaluated. We use the treatment of $\sqrt{3}$-subdivision to highlight how a reinterpretation of the scheme brings forward its algebraic nature and allows a simple implementation (Section 4). In the same spirit, we describe the sparse matrix formulations which capture the essence of the Loop scheme (Section 5). The details of the GPU implementation are then described (Section 6). In particular, we highlight how the identification of relevant connectivity relations can help tune algebra kernels. Additionally, we discuss the influence of the memory layout of the mesh matrix on performance. In Section 7, the performance of our implementation is evaluated against publicly available production and research implementations such as OpenSubdiv [Pixar 2017], the feature adaptive Catmull-Clark version of Nießner et al. [2012], and OpenMesh [Aachen-University 2017].

2 RELATED WORK

Subdivision meshes are commonly used across various fields ranging from character animation in feature film production [DeRose et al. 1998] to primitive creation for REYES-style rendering [Zhou et al. 2009], and real-time rendering [Tzeng et al. 2010].

Mesh subdivision is a refinement procedure which requires data structures capable of providing and updating connectivity information. Commonly used data structures are often variants of the *winged-edge* mesh representations [Baumgart 1972], like quad-edge [Guibas and Stolfi 1985] or half-edge [Campagna et al. 1998; Lienhardt 1994]. While they are well suited for use in the serial setting,

parallel implementations suffer from scattered memory accesses, which are particularly harmful for performance. Besides, their storage cost is a limiting factor on graphics hardware. Compressed alternative formats which were designed for GPU-rendering, like triangle stripes [Deering 1995; Hoppe 1999], do not offer complete connectivity information and are thus not suitable for subdivision.

Most recently, a compact sparse matrix mesh representation has been proposed [Zayer et al. 2017], where mesh processing operations can be expressed as sparse linear algebra and parallelized using linear algebra kernels. While the principal applicability of parallel matrix operations to mesh processing tasks has been reported, a full-fledged treatment of mesh subdivision has not been attempted. In the same spirit, the effort undertaken by Mueller-Roemer et al. [2017] for volumetric subdivision attempts to use boundary operators for boosting performance on the GPU. While these differential forms have been used earlier [Castillo et al. 2005], their storage cost and redundancies continue to limit their practical scope, especially, as data-sets with millions of elements are now mainstream.

Given the pressing need for high performance subdivision implementations, various vectorization approaches have been proposed. Shiue et al. [2005] divide the mesh into fragments which can be subdivided independently on the GPU, which reduces inter-thread communication but introduces redundant data and computations. Moreover, an initial subdivision step has to be done on the CPU. Subdivision tables have been introduced to efficiently reevaluate the refined mesh after moving control mesh vertices [Bolz and Schröder 2002]. However, the creation of such tables requires a symbolic subdivision, whose cost is similar to a full subdivision. Similarly, the pre-computed eigenstructure of the subdivision matrix can be used for direct evaluation of Catmull-Clark surfaces [Stam 1998].

To avoid the cost induced by exact subdivision approaches, approximation schemes have been introduced. Peters [2000] proposed an algorithm that transforms the quadrilaterals of a mesh into bicubic Nurbs patches. While the resulting surface is tangent continuous everywhere, the algorithm imposes restricting requirements on the mesh. The approach of Loop and Schaefer [2008] approximates the Catmull-Clark subdivision surface in regular regions using bicubic patches. Irregular faces still require additional computations. Approximations like the aforementioned are fast to evaluate, but along the way, desirable subdivision properties get lost and visual quality deteriorates. While regular faces can be rendered efficiently by exploiting the bicubic representation using hardware tessellation, irregular regions require recursive subdivision to reduce visual errors [Nießner et al. 2012]. Schäfer et al. [2015] took the idea one step further and enabled different subdivision depths for irregular vertices in a mesh. Brainerd et al. [2016] improved upon these results by introducing subdivision plans. Beyond classical subdivision, several extensions have been proposed to allow for meshes with boundary [Nasri 1987], sharp creases [DeRose et al. 1998], feature based adaptivity [Nießner et al. 2012], or displacement mapping [Cook 1984; Nießner and Loop 2013].

Our approach avoids the aforementioned shortcomings and requires neither preprocessing nor expensive mesh data structures. At the top level, it can be formalized mathematically in the concise language of linear algebra and hence the ensuing algorithms are easy to understand and modify without any knowledge of the underlying numerical kernels. At the lower level, our formalization reveals numerical patterns across subdivision steps which can be exploited to streamline the associated kernels and hence increase the performance yield. Our current implementation achieves considerable speedups on modern architectures and can be virtually ported to nearly any computing platform as its core is formed mainly of basic numerical algebra kernels.

As mesh subdivision requires access to neighboring primitives, a good memory layout is important for performance. Working with sparse matrix representations, insights from matrix reordering can improve the memory layout [Barnard et al. 1993; George 1971; Vo et al. 2012].

3 SPARSE LINEAR ALGEBRA FORMALIZATION

Given the generality and popularity of the Catmull-Clark subdivision scheme, we will use it to walk through the algorithmic development of our method and then briefly show how similar ideas apply to other subdivision methods. We will restrict ourselves to Loop and $\sqrt{3}$ schemes, however, the same principles extend invariably to most other existing schemes.

3.1 Catmull-Clark: Classical Formulation

The Catmull-Clark Subdivision scheme offers a generalization of bicubic patches to the irregular mesh setting [Catmull and Clark 1978]. It is fairly general, and can be applied to polygonal faces of arbitrary order. Regardless of the input, the scheme always produces quadrilaterals and the results tend to be more symmetrical when the input is a quad-mesh. A typical iteration in Catmull-Clark subdivision involves four steps. The first two append new vertices to faces and edges, respectively. The third step relocates original vertex positions so as to smooth the resulting mesh. Finally, the mesh is refined by creating new faces from each old face by connecting the old vertices with the new face and edge points.
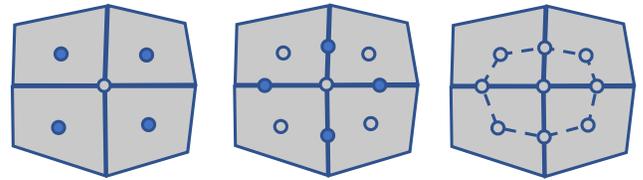
Fig. 2. The Catmull-Clark scheme inserts face-points (left), edge-points (center), and creates new faces by connecting face-points, edge-points and the original central point whose location is updated in a smoothing step (right).

Face-point calculation: For an arbitrary polygonal face i of order c_i, the position of face-point f_i is set to the barycenter of the polygon

$$f_i = \frac{1}{c_i} \sum_{j=1}^{c_i} p_j; \tag{1}$$

where p_j are the face vertices.

Edge-point calculation: For each edge $p_k p_l$, a new edge-point is introduced as the average of the endpoints p_k and p_l and the face-points f_r and f_s corresponding to the two faces bordering the edge:

$$e_{k,l} = \frac{1}{4}(p_k + p_l + f_r + f_s) \tag{2}$$

Vertex update: To produce smooth results, the original vertex locations have to be updated judiciously. In the Catmull-Clark scheme, each original vertex is set to a linear combination of its old position, the edge-mid-points of all incident edges and the surrounding face-points

$$S(p_i) = \frac{1}{n_i}\left((n_i - 3)p_i + \frac{1}{n_i}\sum_{j=1}^{n_i} f_j + \frac{2}{n_i}\sum_{j=1}^{n_i} \frac{1}{2}(p_i + p_j)\right), \tag{3}$$

where n_i is the vertex's valence, f_j are the face-points on adjacent faces and p_j the vertices in the 1-ring neighborhood of p_i.

Topology refinement: To construct the refined polygon, additional edges are created. These new edges connect the face-point to the face's edge-points. Therefore, each parent face j of order c_j is split into c_j child quadrilaterals.

Boundaries: Catmull-Clark subdivision supports meshes with boundary. Edge-points are added to mid-points of boundary edges

$$e_{i,i+1} = \frac{1}{2}(p_i + p_{i+1}) \tag{4}$$

and vertex positions p_i in the boundary polygon are updated

$$S(p_i) = \frac{3}{4}p_i + \frac{1}{8}(p_{i-1} + p_{i+1}). \tag{5}$$

3.2 Catmull-Clark: Linear Algebra Formulation

In this section, we present the higher level formalization of Catmull-Clark subdivision in the language of sparse matrix algebra. In particular, we build upon the mesh matrix $\mathcal{M}$ and some of the primitives proposed in Zayer et al. [2017]. In the matrix $\mathcal{M}$, columns correspond to faces and the row locations of their non-zeros correspond to the indices of the face vertices. The value of an entry is the position of its corresponding vertex in the cyclic order of the face. Throughout this exposition, we will make use of action maps [Zayer et al. 2017], which are helper functions called during sparse matrix-vector multiplication (SpMV) and sparse general matrix-matrix multiplication (SpGEMM) to modify classical multiplication. Results, which would have required intermediate data storage and manipulations are greedily obtained within a single multiplication pass. In this way, the overall higher level parallelism effort is channeled to linear algebra routines.

Face-point calculation: The positions of face-points added to each face coincide with its barycenter.

The required face orders can be obtained using a mapped SpMV

$$\mathbf{c} = \mathcal{M}^T \mathbf{1}; \tag{6}$$
$$\underset{val \to 1}{}$$

where $\mathbf{1}$ is a vector of ones spanning the range of the faces. The mapping below the multiplication implies that the non-zero values of $\mathcal{M}^T$ will be replaced by a 1 during multiplication. (Clearly,

multiplication by 1 can be flagged out for more efficiency, and the compressed column pointers could be used to extract the orders, however, the goal is to illustrate the concept of action maps through a simple example. Further optimizations are discussed in Section 6).

The face-points can then be obtained using the mapped SpMV

$$\mathbf{f} = \underset{val_i \to \frac{1}{c_i}}{\mathcal{M}^T \mathbf{P}}, \tag{7}$$

where $\mathbf{P}$ is the array of all point coordinates. In this case, the values read from the matrix are used as indices into a one dimensional map.

Every non-zero value val_i in $\mathcal{M}^T(i, *)$ is mapped to the reciprocal of the order of face i. For quadrilaterals, which are the most common face type in this algorithm, the SpMV simplifies to

$$\mathbf{f} = \underset{(1,2,3,4) \to \frac{1}{4}}{\mathcal{M}^T \mathbf{P}}. \tag{8}$$

Again, an action map to 1 and a row-wise division by the elements of $\mathbf{c}$ could be used but the idea is to illustrate various types of maps.

Edge-point calculation: The computation of edge points requires assigning unique indices to mesh edges. We can obtain one such enumeration from the upper (or lower) triangular part of the adjacency matrix associated with the undirected graph of the mesh. With standard sparse matrix machinery this matrix can be created, for instance, by first computing the adjacency matrix of the oriented mesh graph (mesh faces are commonly coherently oriented) and then summing it with its transpose, to account for meshes with boundaries. In view of our high performance goals, this is not a viable approach since it requires additional data creation (transpose), and more importantly, matrix assembly which is notoriously challenging on parallel platforms - especially the GPU.

With action maps this can be conveniently encoded as

$$E = \underset{\{Q_c + Q_c^{c-1}\}[\lambda]}{\mathcal{M}\mathcal{M}^T} \tag{9}$$

For the computation of E, the two circulant matrices Q_c and its power Q_c^{c-1}, where c is the face order, are combined to capture the counterclockwise and clockwise orientation inside a given face. In this context, action maps in SpGEMM can be thought of as small matrices which override the classical multiplication behavior. Whenever a collision between two entries occurs, the non-zero values are used as indices into the map.

For quads, Q_4 captures the CCW and Q_4^3 the CW adjacency.

$$Q_4 = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \end{array} \begin{array}{cccc} 1 & 2 & 3 & 4 \\ \left[\begin{array}{cccc} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{array}\right] \end{array}, \quad Q_4^3 = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \end{array} \begin{array}{cccc} 1 & 2 & 3 & 4 \\ \left[\begin{array}{cccc} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{array}\right] \end{array}; \tag{10}$$

These maps do not have to be created explicitly, as their entries can be computed on demand. This is particularly useful, when the face types vary within a mesh:

$$Q_c^r(i,j) = \begin{cases} 1 & if \quad j = ((i+r-1) \mod c) + 1 \\ 0 & else \end{cases} \tag{11}$$

The function λ is called each time a collision between elements $\mathcal{M}(i,k)$ and $\mathcal{M}^T(k,j)$ happens. It performs the map lookup and, depending on the map value, computes the result of a collision:

$$\lambda(i,j) = Q(i,j) \tag{12}$$

If the map entry is non-zero, the vertices p_i and p_j are connected to each other within a face k.

Unique indices for edges can easily be generated by enumerating the non-zeros in the upper triangular part of the matrix E.

To complete the computation of edge-points, faces adjacent to a given edge are required. For this purpose, a secondary matrix F can be used. This matrix has the same sparsity pattern as the adjacency matrix of the oriented graph of the mesh but each non-zero entry i, j stores the index of the face containing the edge $p_i p_j$. It can be similarly constructed by matrix multiplication such that whenever the action map returns a non-zero for a collision between elements $\mathcal{M}(i,k)$ and $\mathcal{M}^T(k,j)$, the face index k is stored in the result.

$$F = \underset{\{Q_n\}[\gamma]}{\mathcal{M}\mathcal{M}^T} \tag{13}$$

with the function

$$\gamma(i,j,k) = \begin{cases} k & if \quad Q = 1 \\ 0 & else \end{cases} \tag{14}$$

Hence, for each edge $p_i p_j$ in the mesh, its unique edge index is known from E and the two adjacent faces are $F(i,j)$ and $F(j,i)$. The edge-point position can then be computed.

Vertex update: The position update in Equation 3 can be conveniently rewritten as

$$S(p_i) = \left(1 - \frac{2}{n_i}\right)p_i + \frac{1}{n_i^2}\sum_{j=1}^{n_i}p_j + \frac{1}{n_i^2}\sum_{j=1}^{n_i}f_j, \tag{15}$$

such that the update can be split into three summands. Vertex valencies can be obtained globally as the vector

$$\mathbf{n} = \underset{val \to 1}{\mathcal{M}\mathbf{1}}. \tag{16}$$

The first term involves only the original position and can be calculated in the customary ways.

The second term sums the 1-ring neighborhood of the vertex. This is done using the matrix F, which has the same sparsity pattern as the vertex-vertex adjacency matrix without diagonal, in the index mapped SpMV

$$\underset{val_i \to \frac{1}{n_i^2}}{F\mathbf{P}}. \tag{17}$$

The last term sums the face-points on faces adjacent to the vertex and is computed via

$$\underset{val_i \to \frac{1}{n_i^2}}{\mathcal{M}\mathbf{f}}. \tag{18}$$

Topology refinement: Each face r of the control mesh contributes c_r child quadrilaterals to the refined mesh. A new face consists of one vertex of the parent, its face-point and two edge-points. To add a face to a mesh matrix $\mathcal{M}$, the column representing the new polygon can simply be appended to the matrix.

For a column $\mathcal{M}(*, r)$ in the control mesh matrix, c_r columns are added to the refined mesh matrix where c_r is the order of the face. For each of the new columns, four indices have to be determined. The indices of original vertices are already known, and the index of a face-point on f_r is $|v| + r$, where $|v|$ is the number of mesh vertices. The indices of the two edges can be fetched using the matrix E and incremented by $|v| + |f|$, where $|f|$ is the number of mesh faces, to get their edge-points' indices.

Boundaries: In practice, meshes often feature boundaries, which need to be treated using specialized subdivision rules. The boundary vertices can be conveniently identified from E as entries which have a value of 1, as they are encountered only once.

Boundary meshes are handled in a build and repair fashion. The individual steps are carried out as usual and external vertices are identified and repaired.

3.3 Catmull-Clark: Modes of Operation

Our approach can be used in one of two different ways, depending on the requirements of the specific practical application: dynamic and static topology of the subdivision meshes.

Static topology is common, *e.g.*, in production rendering applications, where only vertex attributes, *e.g.* positions, change over time but the mesh connectivity is persistent. Subdivision algorithms make heavy use of adjacency information. The fact that this information can be prepared upfront and does not have to be re-computed in every frame, reduces the overall production time. To accomplish this goal in our approach, we factor all computations dealing with mesh connectivity information into a *build step*, that only has to be done once upfront before the mesh is subdivided for the first time. In this preprocessing phase, all information required to evaluate refined vertex data from coarse control mesh data is computed and stored on a per subdivision level basis. Each of these levels contains the topology of that level and all information needed to subdivide vertex data from one level to the next.

The *evaluation step* is straight forward: For each iteration of subdivision, the information generated in the build step is used to calculate the refined vertex data of the next depth from the coarser data as described in Section 3.2.

The build step is concerned with the computation of matrices E and F and the creation of the refined mesh matrix up to the defined maximum level. This information is later used in the evaluation step to calculate face orders, vertex valencies, face-points, edge-points and vertex updates. Of course the data precomputed in the build step has to be kept around for the evaluation, which exactly reflects the increased memory footprint if the two steps are separated. Clearly, it is possible to move face and vertex order computations into preprocessing resulting in slightly increased build times and memory footprint but higher evaluation performance.

Dynamic topology is ubiquitous in 3D modeling and CAD applications during the content creation process. Faces, vertices and edges are frequently added, modified and removed which poses a great challenge to many existing approaches that rely on expensive preprocessing, as it has to be repeated on every topological update. This fact has led to the use of different subdivision approaches for model preview and rendering causing discrepancies between the
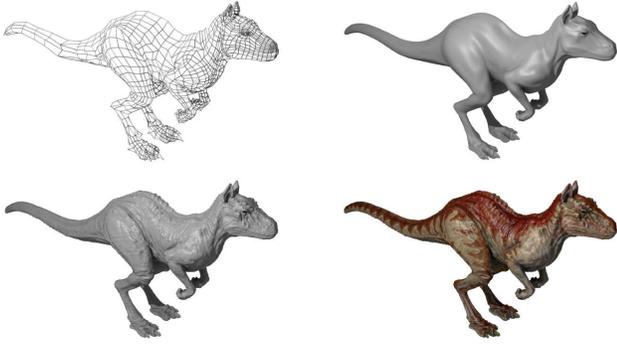
Fig. 3. AlSparse is capable of subdividing a coarse control mesh (top, left) instantaneously to a dense and smooth refined mesh (top, right). Displacement (bottom, left) and texture (bottom, right) mapping are possible naturally without any additional effort.

two images. Due to the efficiency of our complete approach, we can avoid any preprocessing and compute both the build and evaluation step instantaneously avoiding any additional memory costs.

3.4 Catmull-Clark: Selective and Feature Adaptive Subdivision

In this section we will discuss how selected regions in the mesh can be subdivided individually which might be used in several applications. It might for example improve flexibility in path and ray tracing approaches were spatially coherent ray packets are gathered and evaluated, as our approach is not limited to work on a per patch basis, but is capable of evaluating arbitrary portions of the mesh while eliminating data redundancy between adjacent faces.

Another setting where only portions of the mesh have to be subdivided recursively is feature adaptive subdivision, which we will use to demonstrate our approach. As the subdivision steps progress, the faces of the resulting mesh can be regarded as an arrangement of regular lattice, except for a limited number of locations where the lattice regularity is perturbed. These locations are called extraordinary vertices and their number does not increase throughout the subdivision process [Doo and Sabin 1978b]. To take advantage of hardware tessellation and still get an exact result, the idea emerged to subdivide regions around irregular vertices and features such as creases and corners manually and build bicubic patches for regular regions [Nießner et al. 2012].

Using our matrix formalism, traditional mesh data structure are not required for identifying sub-meshes that need to be refined around extraordinary vertices. These are vertices whose valency is different than four and can be identified from vertex orders calculated in Equation 16. The regions surrounding them can be obtained by propagation using the mesh matrix $\mathcal{M}$. Starting with a vector $\mathbf{x_0}$ spanning the number of vertices and initialized to 1 at extraordinary vertices and to 0 elsewhere.

A propagation iteration is carried out in two steps. First, the neighboring faces are determined as the non-zeros of the resulting vector

$$\mathbf{q_i} = \mathcal{M}^T \mathbf{x_i} \qquad (19)$$
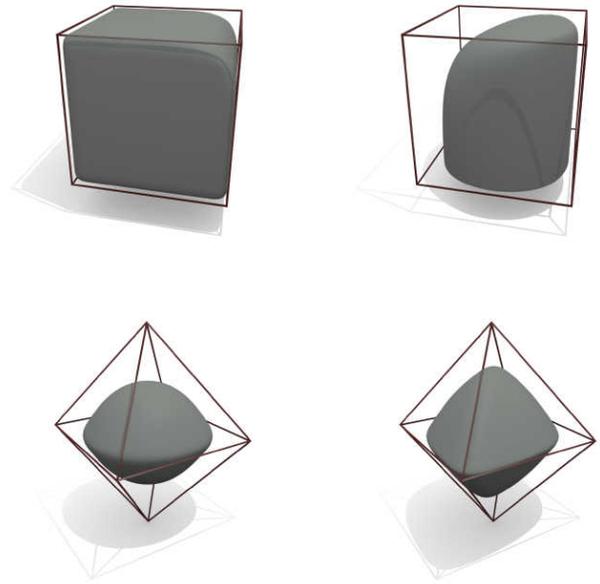
Fig. 4. Due to its linear algebra formulation, our approach naturally supports extensions, such as semi-sharp and infinitely sharp creases. Top: a cube with semi-sharp and infinitely sharp creases. Bottom: octahedron with smooth and semi-sharp creases.

and their vertices can be revealed as the non-zero entries resulting from the product

$$\mathbf{x_{i+1}} = \mathcal{M}\mathbf{q_i}. \qquad (20)$$

This in fact reflects that the adjacency matrix can be obtained from the mapped mesh matrix product and the power of the adjacency matrix reflects the neighborhood order around a vertex.

In our current implementation, a 3-ring propagation is used to conform with the output OpenSubdiv provides after each step.

Displacement mapping and hierarchical edits [Forsey and Bartels 1988] are also straight forward in our approach regardless, as we have access to the refined mesh after each iteration and can therefore simply use the refined data and uv-coordinates to arbitrarily modify the vertices. An example displacement and texture mapped subdivision model can be seen in Figure 3.

3.5 Catmull-Clark: Creases

Sharp and semi-sharp creases have become indispensable in subdivision surface modeling to describe piecewise smooth and tightly curved surfaces respectively [DeRose et al. 1998], as shown in Figure 4. Without that feature, supporting topology has to be added to the mesh to achieve a similar visual appearance, which may lead to substantially increased polygon counts in the refined mesh. Creases are edges that are tagged by a (not necessarily) integer sharpness value. In each subdivision iteration an edge-point is added to each crease edge at a position dependent on the crease's sharpness value and crease vertices are updated using a special set of rules depending on the number and average sharpness value of incident crease edges. As the general computation of creases is beyond the scope of this paper we refer the reader to DeRose et al. [1998] for a detailed

description. To support creases, we use a sparse symmetric crease matrix C of size $n_v \times n_v$. The entry $C(i,j) = \sigma_{ij}$ holds the sharpness value of the crease between vertices i and j. To calculate the position of crease vertices and edge points, the crease valency $\mathbf{k}$ (number of creases incident to a crease vertex)

$$\mathbf{k} = \underset{val \to 1}{C\mathbf{1}} \tag{21}$$

and the vertex sharpness $\mathbf{s}$ (average over all incident crease sharpnesses)

$$\mathbf{s} = \underset{val_i \to \frac{val_i}{k_i}}{C\mathbf{1}} \tag{22}$$

need to be determined, which we complete using a single mapped SpMV. In our approach we are handling creases in a similar fashion as boundaries—we first subdivide the mesh as if it was smooth and subsequently correct all vertices corresponding to creases. With the computed vectors $\mathbf{k}$ and $\mathbf{s}$ and the already available adjacency information in E, the correction of crease vertices can simply be carried out in parallel using the rules provided by DeRose et al. [1998]. After each iteration of subdivision a new crease matrix has to be created, that holds the updated sharpness values for the subdivided creases. As a crease may turn into a smooth edge after a certain number of subdivisions, we perform the crease matrix creation in two steps:

(a) To *determine the sparsity pattern* we calculate the sharpness value for a new crease in the refined mesh from the sharpness values of the parent crease and their neighbors using Chaikin's edge subdivision algorithm [Chaikin 1974] and reducing the resulting sharpness by one to account for the performed subdivision step [DeRose et al. 1998]

$$\sigma_{ij} = \max\{\tfrac{1}{4}\left(\sigma_i + 3\sigma_j\right) - 1, 0\} \tag{23}$$

$$\sigma_{jk} = \max\{\tfrac{1}{4}\left(3\sigma_j + \sigma_k\right) - 1, 0\} \tag{24}$$

where σ_i, σ_j and σ_j are sharpness values of three adjacent parent creases i, j and k. σ_{ij} and σ_{jk} are the sharpness values of the two child creases of j. Counting the number of resulting non-zero sharpnesses in each column of the new crease matrix and performing a parallel scan over these counts results in the column pointer of the new crease matrix and the total number of non-zero entries in the new C which can be used to allocate row index and value arrays.

(b) To *populate the sparsity pattern* a similar kernel as in the first step is used, but instead of counting the non-zeros, the updated sharpness values and the corresponding row indices are written to the sparsity pattern of the new crease matrix.

4 √3-SUBDIVISION

The √3-subdivision scheme is specialized for triangle meshes and is based on a uniform split operator which introduces a new vertex for every triangle of the input mesh [Kobbelt 2000]. It defines a natural stationary subdivision scheme with stencils of minimum size and maximum symmetry.

The subdivision process involves two major steps. The first one inserts a new vertex f_i at the center of every triangle i.

Each new vertex is then connected to the vertices of its master triangle and an edge flip is then applied to the original edges, see Figure 5. In the second step, the positions of the old vertices are updated using the following smoothing rule

$$S(p_i) = (1 - \alpha_i)p_i + \frac{\alpha_i}{n_i} \sum_{1}^{n_i} p_j \tag{25}$$

where n_i is the valence of vertex p_i and α_n is obtained by analyzing the eigen-structure of the subdivision matrix:

$$\alpha_i = \frac{4 - 2\cos(\frac{2\pi}{n_i})}{9}. \tag{26}$$

Clearly the topological operations involved in this scheme anticipate an edge-based mesh representation and all the implementations we are aware of rely on the half-edge data structure.
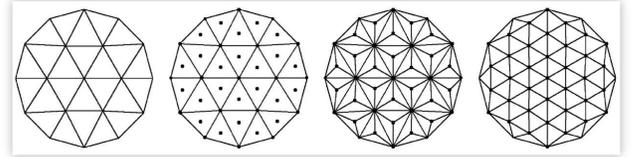
Fig. 5. Original description of the √3-subdivision scheme. First a new vertex is inserted at every face of the given mesh. Second, an edge flip applied to the original mesh edges yields the final result, which is a 30 degree rotated regular mesh. Applying this scheme twice leads to a 1-to-9 refinement of the original mesh. Original image from [Kobbelt 2000], copyright ACM.

In order to adapt this subdivision scheme to our matrix algebra framework, we reinterpret the whole process in a slightly different manner. By reasoning only on triangles as detailed in Figure 6, the topological operations get simplified and the subdivision scheme can be easily abstracted using sparse matrix algebra. In fact, we need only a good bookkeeping of triangle-triangles adjacency to obtain new triangulations and update vertex positions. Please note, that the boundary can be treated by using adequate smoothing [Kobbelt 2000] using similar ideas to the outline given earlier for the Catmull-Clark scheme, but we omit it here to keep the presentation succinct.
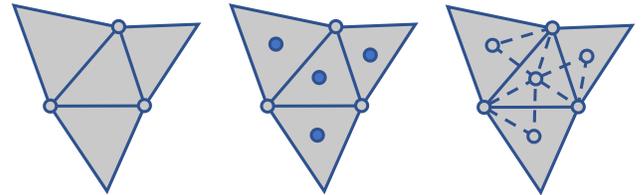
Fig. 6. After inserting the new vertices (blue), each triangle contributes three new triangles to the refined mesh which are obtained by joining its vertices to their respective right and left neighboring new vertices.

New vertex points. A new vertex is added to each triangle's barycenter. The average of triangle vertices can be calculated using the mapped SpMV

$$\mathbf{f} = \underset{(1,2,3) \to \frac{1}{3}}{\mathcal{M}^T \mathbf{P}} \tag{27}$$

Required adjacency information. The $\sqrt{3}$ scheme adds a vertex to each triangle and connects it to the new vertices on the three neighboring triangles. To find these neighbors efficiently, we can again use the oriented graph adjacency matrix to store the index of the adjacent face to any given edge, as in Equation 13.

Vertex update. The second term in the smoothing step can also be performed with action maps

$$\underset{val_i \to \frac{\alpha_i}{n_i}}{F\mathbf{P}} \tag{28}$$

Topology refinement. For a mesh with n_v vertices, each vertex of a given triangle (p_k, p_l, p_m) with index i, contributes a new triangle to the refined mesh. For instance, vertex p_k contributes the triangle consisting of p_k itself, the face-point f_i which can be conveniently indexed by $i + |v|$ and the barycenter on an adjacent triangle, which then takes index $F(l, k) + |v|$. The mesh matrix of the refined mesh can be efficiently created in parallel by appending new columns.

5 LOOP SUBDIVISION

This scheme is another triangle mesh subdivision method which was introduced by Loop [1987]. It refines a mesh by inserting new edge-points as described in Figure 7-left. For each triangle, these points can be used to perform a split into four new triangles. The original vertex positions are then smoothed using local weighted averaging as summarized in Figure 7-right. The weighted average in the smoothing step is based on convergence consideration and is defined as

$$S(p_i) = (1 - n_i \beta_i) p_i + \beta_i \sum_1^{n_i} p_j, \tag{29}$$

where

$$\beta_i = \frac{1}{n_i} \cdot \left(\frac{5}{8} - \left(\frac{3}{8} + \frac{1}{4} \cdot \cos\left(\frac{2\pi}{n_i}\right) \right)^2 \right). \tag{30}$$
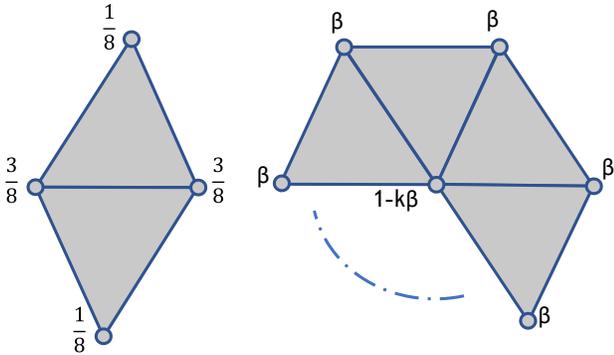
Fig. 7. For each edge, the Loop scheme inserts a new vertex as a weighted sum of the vertices of the adjacent triangles (left). In the smoothing step, original positions are updated using a β-weighted combination of the their neighbors (right).

In the following, we briefly describe the algebraic machinery we use to capture the topological modifications intrinsic to this scheme.

New vertex points. For a given edge, the vertex insertion requires the edge vertices and the vertices opposite to the edge. We can gather this information by using the adjacency matrix of the directed graph of the mesh and for each edge store the index of the remaining triangle vertex as the non-zero value.

$$\underset{\{Q_3\}[\lambda]}{G = \mathcal{M}\mathcal{M}^T} \tag{31}$$

$$\lambda(i, j) = \begin{cases} k & if \quad Q_3 = 1 \\ 0 & else \end{cases} \tag{32}$$

where k is the vertex opposite to edge $p_i p_j$.

Unique edge indices can be obtained from this matrix by summing it with its transpose to obtain a matrix E and incrementally assigning indices to the non-zeros of the upper triangular part of E similarly to how it is done in the context of Catmull-Clark subdivision. The new vertex locations can then be obtained by looking up the unique edge indices and for each edge $p_i p_j$, obtaining the opposite vertices as $G(i, j)$ and $G(j, i)$ and performing the summation as given in Figure 7-left.

Vertex update. The second term in Equation 29 can be computed using the mapped SpMV below

$$\underset{val_i \to \beta_i}{G\mathbf{P}}, \tag{33}$$

where the action maps substitutes values in row i by β_i.

Topology refinement. For each triangle (p_k, p_l, p_m) in the control mesh, three new triangles of the refined face are simple arrangements of an original vertex and two new edge-points. The fourth triangle is only composed of the three new edge-points. The original vertices' indices are p_k, p_l, p_m and unique edge indices can be obtained from the upper triangular part of E, as done in the Catmull-Clark scheme. With this information, the refined mesh matrix can be constructed efficiently in parallel.

6 OPTIMIZATION OF ALGEBRAIC OPERATIONS

The higher level formalization discussed in Section 3 can be easily implemented by minor adjustment to standard sparse matrix algebra kernels. This already yields better performance on a variety of basic mesh processing tasks as reported by Zayer et al. [2017]. In our view, the algebraic operations can be further optimized on the lower level to unleash the computational power embedded in modern hardware. In fact, our approach uses SpMVs and SpGEMMs not for computations on arbitrary matrices but on mesh and adjacency matrices. This enables us to exploit the particular computational patterns found in those matrices and streamline them through efficient and highly optimized GPU kernels.

6.1 Reduced Mesh Matrix

In our general implementation we use the Compressed Sparse Column (CSC) matrix format, which is comprised of three arrays. The first two hold row indices and values of non-zero entries. The column pointer contains an index to the start of each column in the first two arrays [Saad 1994]. Whenever the mesh matrix $\mathcal{M}$ represents a homogeneous mesh, such as a triangular or quadrilateral mesh, the column pointer can be dropped, because the face orders are

consistent and the start index in the row index and value array is known for each face. Reordering the row index-value pairs, such that the values are sorted in each column also renders the value array unnecessary, because the cyclic order of vertices in a face is implicitly given by the order of their appearance in the row indices array. The memory requirement of the reduced mesh matrix is therefore equal to that of a face table of the mesh. An evaluation of this reduction on basic matrix vector multiplication has been done before [Zayer et al. 2017].

Some subdivision implementations are designed to work for triangular meshes exclusively, *e.g.*, $\sqrt{3}$ and Loop. While the Catmull-Clark algorithm accepts polygonal meshes as input, it produces only quadrilaterals. Therefore, this optimization can be applied frequently in subdivision implementations and it cuts down data creation as well as memory consumption and expensive memory accesses.

6.2 Implicit mapped sparse matrix-matrix multiplication

Mapped multiplications of the form

$$A = \underset{\{Q\}[\alpha]}{\mathcal{M}\mathcal{M}^T} \tag{34}$$

are extensively used in our mathematical formalization for capturing various connectivity information. Tailoring high performance algorithms for sparse generalized matrix-matrix multiplication (SpGEMM) is one of the most fascinating and challenging tasks in modern numerical kernel development. Despite the steady improvements reported every year, the cost of this operation is still relatively high, especially, as general purpose algorithms have to address intricate issues pertaining to memory allocation and load balancing for the arbitrary case. Therefore, it is worthwhile to avoid explicit multiplication wherever possible. A close examination of what happens during a multiplications as in Equation 34 reveals that the result can be directly created from $\mathcal{M}$.

During the sparse matrix-matrix multiplication $\mathcal{M}\mathcal{M}^T$, each row $r_i = \mathcal{M}(i, *)$ is multiplied with each column $c_j = \mathcal{M}^T(*, j)$. Both vectors, r_i and c_j, encode one vertex each and have non-zero entries at the positions corresponding to their surrounding faces. During multiplication, a collision between two entries $\mathcal{M}(i, k)$ and $\mathcal{M}^T(k, j)$ happens if both are non-zero, meaning that vertices i and j share a face k. Clearly, vertices not part of the same face will never induce a collision. An action lookup corresponds to a check whether two given vertices of a face are in a specific relation. In case of a parameterized action map, a collision invokes the function α.

A mapped SpGEMM as in Equation 34, would usually require multiplication of each row vector in $\mathcal{M}$ with all n_v columns of $\mathcal{M}^T$. As collisions can only happen locally within a face the number of columns c_j that have to be multiplied with a certain row r_i is reduced from n_v to the valency of vertex i. Still, there might be collisions between two entries where the map lookup and therefore also the invocation of α result in a zero return value which do not contribute to the corresponding entry in the result matrix. Switching from a vertex- to a face-centric viewpoint allows us to eliminate invocations to α that will for sure result in a zero return, namely those, where the map lookup returns zero. As action maps encode relations between positions of vertices within the cyclic order of a face, the function α only needs to be called with the j-th vertex if and only if $Q(i, j) \neq 0$. Therefore, $Q(i, *)$ is an evaluation pattern for the i-th vertex in each face, which determines the invocations of α that have a non-zero return. This means the mapped multiplication can be carried out implicitly, by working in parallel on the row indices of $\mathcal{M}$ and invoking α for each pair of vertices in each face that fulfill the relation encoded by Q.

Before the actual multiplication can be carried out in the way described above, the number of non-zero values in each column of the result needs to be determined, to be able to allocate sufficient memory for the arrays of its CSC representation. This can be done using a preceding symbolic pass, similar to general SpGEMM algorithms. In parallel for each entry in the row indices of $\mathcal{M}$, we determine the number of *local* per-face non-zero α invocations for the vertex, by counting the non-zeros in the map row corresponding to the vertex's position in the cyclic order the current face. The *global* number of non-zero invocations for each vertex is accumulated in an array, which then corresponds to the number of non-zero entries in the vertex's column of the result. A simple parallel scan (cumulative sum) over that array gives the column pointer and the number of non-zero values of the resulting matrix. Now, the row index and value arrays can be allocated and subsequently filled during the evaluation pass. It is worth noting that this step can be skipped if each row of the map has the same number of non-zero entries $z_{r_i} = z_r$. Then the number of non-zero values of the vertex's column in the result, is independent of its position in the cyclic order of adjacent faces. If this is the case, the invocations on each vertex can be directly calculated as a multiple of the vertex order $z_r \mathbf{n}$.

6.3 Specialized SpMVs

Certain patterns in SpMVs formulated in the higher level linear algebra can be transformed to specifically tailored and more efficient GPU kernels. Basis for all eventual SpMV optimizations is a naïve GPU implementation [Steinberger et al. 2016]. Like classical serial SpMV, it does not require explicit transpose computation for performing transposed matrix-vector multiplication. This is particularly important since we make extensive use of the mesh matrix and its transpose in SpMVs and calculating the transpose of a sparse matrix is relatively expensive.

Direct mapped SpMV. In the mapped SpMV for CSC matrices, multiple threads collaborate to calculate a single element of the result vector. Parallelization is done over the elements of the vector. A thread reads a single entry of the input vector and multiplies it with the mapped non-zero elements of the corresponding column. These intermediate products are directly accumulated in the result vector using an atomic addition operation to avoid race conditions.

Transpose mapped SpMV. In the transpose mapped SpMV for CSC matrices, a single thread is responsible for a single output element, which eliminates the need for atomic operations. Each thread iterates over the non-zero elements of its column, uses the map to substitute them and multiplies each mapped value with the corresponding vector element. This means, that in contrast do the direct mapped SpMV, a thread has to read multiple elements from the input vector.

Specializations and Memory Optimizations. Depending on the different input parameters to the mapped SpMVs we can take several optimization steps to produce more efficient GPU kernels. We distinguish between matrix, vector and map-based optimizations.

If the *input matrix* is in reduced form, every column has the same number of non-zero entries, which renders the loop over each column obsolete and it can be unrolled. Therefore, the column pointer is not required to perform the multiplication. Value arrays can also be omitted, because row indices in each column are sorted in a reduced mesh matrix to reflect the cyclic order of the face. In both SpMV versions each thread works on one column of the matrix. If the mesh matrix represents a quadrilateral mesh, as it is very common in the Catmull-Clark scheme, each column has exactly four entries. Instead of performing four individual memory accesses while reading the row indices, a single 128-bit request can be issued, reducing the number of reads by a factor of four.

The content of the *input vector* may enable further optimizations. To count elements, such as vertices in each face or faces adjacent to each vertex, the mesh matrix is multiplied with a one vector. In this case the reads of vector elements is obsolete and can be omitted, as the linear combination reduces to a simple sum. In many cases a vector of positions is used in a mapped SpMV with the mesh matrix, to average over local neighborhoods in the mesh. As every input position consists of multiple components the number of threads can be increased such that the multiplication is carried out on a per component level. Without loss of generality, consider the case of averaging the vertex positions for each face, *e.g.*, when calculating face-points in the Catmull-Clark scheme. Each position consists of four components and each column in M has four non-zero entries. In this case an SpMV kernel can be constructed that is launched with 16 threads per face, each responsible for a single component of one vertex position. Each group of sixteen consecutive threads can then calculate the mapped multiplication of a single column. In the general case their intermediate products are then combined in the result using atomic addition operations. As it is known that each output component will only depend on intermediate products of four vertices, efficient SIMD level communication primitives (shuffle instructions on NVIDIA hardware) can be used to combine the results and the result can be written by a single thread without the need for atomics.

Certain properties of the *map* might be exploited to optimize SpMV kernels. If the map is a constant function, as it is often the case for averaging over a local neighborhood, the value of the map can simply reside in shared or constant memory or even in a register to eliminate frequent map lookups. In the non-transposed case, maps that output the same value for each entry in a face can be handled similarly, as each column is only used by a single thread.

Fusion. Kernel fusion is an important paradigm in parallel computing, as it enables to reduce kernel launch overheads and costly memory loads and stores by merging kernels that have overlapping inputs or data dependencies. Whenever two operations in the high level linear algebra formulation require the same input vectors, and the number of threads required for both computations agree, the two generated kernels can be merged, such that data is not required to be loaded multiple times. The input to the fused kernel is then
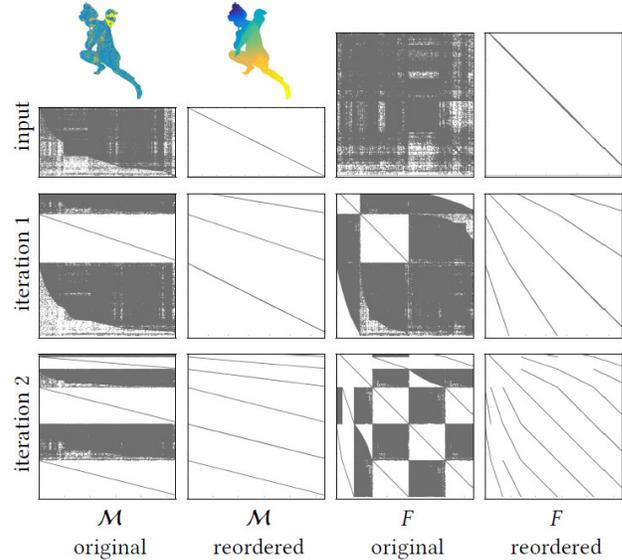
Fig. 8. Evolution of the mesh matrix M of the original and RCM-reordered Angel model (first and second columns respectively) throughout two Catmul-Clark iterations. Color-coded geometric layouts of both orderings is shown on top. The evolution of the respective F matrices is shown in the third and fourth columns.

the union of the sets of inputs to the two individual kernels. The same is true for the output variables.

Another case where fusion might be advantageous, is when the output of the first kernel is the input to the second one. In that case the data does not have to go through global memory from the first to the second kernel but can directly be used in the same kernel after it was computed. If the data that causes the dependency is not needed in any further computation it is not even necessary to store it in global memory at all. Especially for subdivision, where different output data is generated from shared input data, *e.g.*, edge points and face points, fusion can greatly reduce memory access, and thus improve performance, especially for memory bound kernels.

6.4 Mesh reordering

Fast mesh querying is key to any efficient, high performance subdivision implementation. While this may be sufficient in theory, non-algorithmic factors such as memory access and cache effects are crucial for algorithmic performance in practice. Data layout in memory directly affects access patterns and therefore ensuring the locality of such patterns would allow taking advantage of caching mechanisms. In this way, global reads and writes, which are known to cause performance deterioration, especially on modern GPUs, can be reduced.

In our context, this translates to ensuring that primitives which are topologically close in the mesh reside close in memory as well. The mesh memory layout is reflected in the sparsity pattern of the mesh matrix, and locality can be enforced by clustering the nonzero elements close to the diagonal. The closeness to the diagonal can be measured in terms of bandwidth and wavefront [Davis 2006].

For a rectangular sparse matrix $\mathbf{A}$, the bandwidth of a row k, denoted by β_{r_k}, is the maximum width between its non-zero entries. The *row bandwidth* of the matrix $\mathbf{A}$ is then defined as

$$\beta_r = \max_k \left\{ \beta_{r_k} \right\}. \tag{35}$$

We can define the *column bandwidth* β_c accordingly. Then, the bandwidth of a matrix is the maximum of its row and column bandwidth. When the matrix is symmetric, both are equal and we refer to them indifferently as the matrix bandwidth. With respect to the mesh matrix, the row bandwidth defines the vertex bandwidth which we denote β_v. Similarly, the face bandwidth β_f corresponds to the column bandwidth.

A column j is *active* in row i if $j \geq i$, and there is a non-zero entry in that column in any row with index $k \leq i$. Let c_i denote the number of active columns in row i. In plain English, this is the number of columns that have nonzero entries on both sides of row i. The matrix row wavefront ω_r is the maximum over all c_i. Similarly, we can define the matrix column wavefront ω_c. With respect to the mesh matrix, these correspond to the vertex wavefront ω_v and face wavefront ω_f. Since a column, or a row cannot be active more than its width, it follows that $\omega_v \leq \beta_v$ and $\omega_v \leq \beta_v$.

For bandwidth reduction, the reverse Cuthill McKee (RCM) algorithm is fairly well known to produce an inexpensive and low bandwidth reordering [Cuthill and McKee 1969; George 1971]. The original RCM only works on square symmetric matrices. To reorder (usually non-square) mesh matrices, the RCM algorithm can be applied to the graph Laplacian of the mesh. The acquired permutation is then applied to the rows of M. The columns are sorted by minimum row index of their non-zero entries in ascending order.

As shown in Figure 8, there is no reason to expect mesh creators to deliver coherently ordered meshes. Therefore, a reordering of the input can be beneficial in any subsequent mesh processing operations. As subdivision dynamically changes the mesh after each step it is expected that the quality of the mesh ordering deteriorates as it gets refined. The cost of a full RCM reordering after each iteration is simply too prohibitive. We observed however, that the way we append new vertices and faces to the original mesh helps preserve a tightly aligned sparsity pattern of the mesh matrix and other intermediate matrices we use, in particular F as shown in Figure 8. Please note that by construction, the matrix E has a quite similar sparsity patterns to F. For the model shown in Figure 8 the mesh layout metrics in the original mesh read $\beta_f = 236\,320$, $\omega_f = 155\,119$, $\beta_v = 473\,753$, $\omega_v = 178\,423$. After reordering they drop significantly to $\beta_f = 14\,299$, $\omega_f = 624$, $\beta_v = 28\,643$, $\omega_v = 1\,252$. This implies that significant performance gains can be expected. Empirical evaluation of the gains induced by mesh reordering are discussed in Section 7.

7 RESULTS

This section provides an evaluation of our approach. In particular, we distinguish the performance of simply formulating the algorithms in the language of sparse matrix algebra (*AlSub SpLA*) and the performance resulting from using the optimized kernels (*AlSub opt.*). Comparisons are made to the current industry standard, Open-Subdiv, which splits subdivision into three steps. First, a symbolic subdivision up to a given maximum level is performed to create the

Fig. 9. Some of the meshes used in our experiments. Neptune, girl, Eulaema Bee (courtesy of The Smithsonian Institution), hat, car (courtesy of Yasutoshi Mori) and ArmorGuy (courtesy of DigitalFish).

refined topology, which is then used in a second step to precompute the stencil tables for face-points, edge-points and vertex update. In our evaluation we summarize these first two steps as *build*. The stencil tables are then used to perform the evaluation of refined vertex data (*eval*). While the evaluation is performed in parallel on the GPU, the two preprocessing steps are done on the CPU. To compare our $\sqrt{3}$ implementation we use OpenMesh which is a serial mesh processing framework, widely used in academics.

Evaluation Setup. All tests are performed on an Intel Core i7-7700 with 32GB of RAM and an Nvidia GTX 1080 Ti. The provided measurements are the sum of all kernel timings required for the subdivision, averaged over several runs. The input models are unaltered and thus have not been reordered unless specifically marked differently.

7.1 $\sqrt{3}$ performance

To evaluate the performance of $\sqrt{3}$ we use differently sized models and different subdivision depths, listed in Table 1. As OpenSubdiv lacks support for this scheme and we are not aware of a different GPU implementation for $\sqrt{3}$, we compare our approach to Open-Mesh. Here we evaluate what we call the "modeling" use case, where

	mesh	c_f	c_v	n_i	r_f	r_v
$\sqrt{3}$	fox	622	313	6	453.4k	226.7k
	girl_bust	61.3k	30.7k	6	44.7M	22.4M
	goblet	1.0k	520	6	729.0k	364.5k
	Hhomer	10.2k	5.1k	6	7.4M	3.7M
	star	10.4k	5.2k	6	7.6M	3.8M
	bee	16.9M	8.5M	1	50.8M	25.4M
	neptune	4.0M	2.0M	2	36.1M	18.0M

Table 1. $\sqrt{3}$-subdivision test meshes: Number of faces c_f and vertices c_v of the control meshes as well as the applied number of iterations n_i and faces r_f and vertices r_v in the refined mesh.

	mesh	c_f	c_v	n_i	r_f	r_v
Loop	archer_t	3.2k	1.6k	6	13.1M	6.5M
	hat_t	8.8k	4.4k	6	36.2M	18.1M
	goblet	1.0k	520	6	4.1M	2.0M
	Hhomer	10.2k	5.1k	6	41.8M	20.9M
	phil_t	6.1k	3.1k	6	24.9M	12.5M
	star	10.4k	5.2k	6	42.5M	21.3M
	bee	16.9M	8.5M	1	67.8M	33.9M
	neptune	4.0M	2.0M	2	64.1M	32.1M

Table 2. Loop test meshes: Number of faces c_f and vertices c_v of the control meshes well as the applied number of iterations n_i and faces r_f and vertices r_v in the refined mesh.
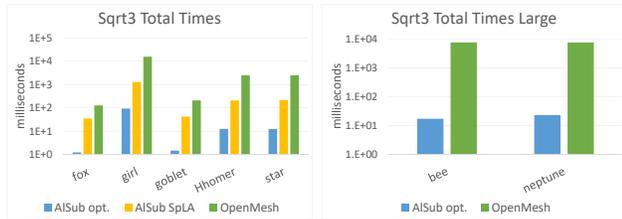
Fig. 10. $\sqrt{3}$ subdivision: Performance comparison of our unoptimized (AlSub SpLA) and optimized (AlSub opt.) implementations to OpenMesh.

Fig. 11. Loop subdivision: Peak GPU memory requirements comparison between our approaches and OpenSubdiv. No preprocessing in AlSub - full subdivision in each iteration.

Fig. 12. Loop subdivision: Comparison of AlSub and OpenSubdiv. Topology refinement is performed in each iteration. Evaluation times are given for OpenSubdiv.

no preprocessing is performed and topology as well as vertex data are subdivided from to input geometry to a given subdivision depth.

Figure 10 shows the achieved performance of our unoptimized and optimized approach compared to OpenMesh. While it is clear, that a parallel GPU implementation is capable of outperforming a serial CPU approach (up to two orders of magnitude), it is interesting to compare our SpLA version to its optimized counterpart. Throughout all test cases, we achieve a performance gain of 10× or more with our optimizations. Especially when starting with a smaller input model, the AlSub opt. pulls away further, which we attribute mostly due the involved SpGEMM operations which show a certain overhead independent of the input size. This overhead also reflects in the temporary memory requirements, which prohibit very large meshes (bee or neptune) to complete with our unoptimized version. AlSub opt. handles these cases without trouble.

7.2 Loop performance

To evaluate the performance of our linear algebra implementations of Loop subdivision we used the meshes in Table 2. We compared the time our approach needs for the complete subdivision of the mesh (topology and vertex positions) to the build and evaluation performance of OpenSubdiv, again considering a "modeling" use case where the input topology changes.

Timing results are shown in Figure 11. AlSub performs the full subdivision without any preprocessing. Again, our optimizations increase performance by about one order of magnitude. For Open-Subdiv we distinguish between the individual timings of preprocessing (*build*) and evaluation (*eval*). Considering the time it takes from a topology changing modeling operation and the final subdivided mesh, our approach clearly outperforms OpenSubdiv by multiple orders of magnitude due to its sequential preprocessing. However, more interestingly, the time needed for AlSub's *complete*

subdivision is only slightly higher than OpenSubdiv's evaluation step. Which is a surprising result, considering that OpenSubdiv's evaluation step is provided by an optimized GPU kernel, which essentially only weighs input vertices by predefined weights, and AlSub performs the complete topological refinement and generates all output vertices. This fact is especially true for very large meshes (bee, Neptune). Figure 12 shows the memory requirements for the tested cases. Again, our optimized version requires less memory than the SpLA version. As we complete the entire subdivision process, we do not keep any additional preprocessed data around, which explains the memory overhead of OpenSubdiv. Internally AlSub stores a vertex position as four float values while OpenSubdiv uses three. As the bee model is only subdivided once, AlSub can not compensate this overhead with the lower memory requirements of our preprocessed data, which is why OpenSubdiv has a lower peak memory consumption in this special case.

	mesh	c_f	c_v	n_i	r_f	r_v
	ArmorGuy	8.6k	10.0k	6	35.2M	35.3M
	hat	4.4k	4.4k	6	18.1M	18.1M
	coat	5.6k	5.7k	6	22.8M	22.8M
Catmull-Clark	bike	53.9k	54.3k	6	13.4M	13.5M
	car	149.7k	164.9k	6	38.5M	38.7M
	dress	2.3k	2.4k	6	9.2M	9.2M
	bee	16.9M	8.5M	1	50.8M	50.8M
	neptune	4.0M	2.0M	2	48.1M	48.1M

Table 3. Catmull-Clark test meshes: Number of faces c_f and vertices c_v of the control meshes well as the applied number of iterations n_i and faces r_f and vertices r_v in the refined mesh.
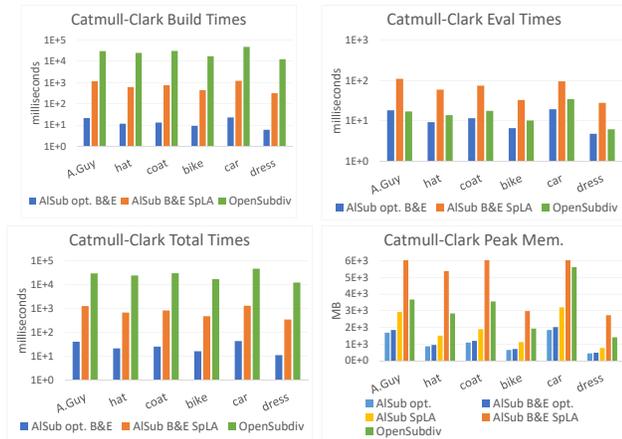
Fig. 13. Catmull-Clark subdivision: Comparison of the individual steps (build & eval) of AlSub and OpenSubdiv for small to medium sized meshes including peak memory consumption for all approaches.

7.3 Catmull-Clark performance

AlSub's Catmull-Clark subdivision performance is evaluated on a variety of differently sized meshes which are listed in Table 3.

The evaluation results for Catmull-Clark subdivision are given in Figures 13 and 14, for which we now also distinguish between the "modeling" and "rendering" use case:

Modeling. This specific use case is again a representative for scenarios in which the mesh topology changes frequently. This means that any eventually preprocessed data concerning the topology of the mesh has to be re-computed. This scenario can be seen in Figures 13 and 14 bottom left. Again, we observe similar behavior as before with AlSub opt. being more than one order of magnitude faster than AlSub SpLA, which is about one order of magnitude faster than OpenSubdiv when performing preprocessing and evaluation. Note that this kind of delay is only introduced by OpenSubdiv if a topology-changing modeling operation is carried out. Otherwise, the simple GPU evaluation is sufficient. Nevertheless, this delay might still yield unpleasant behavior during modeling.

Rendering. In contrast to "modeling", topology is considered static in "rendering", which is the intended use case of OpenSubdiv. In this case, information required during subdivision that only depends on the topology can be precomputed and stored for later use. The
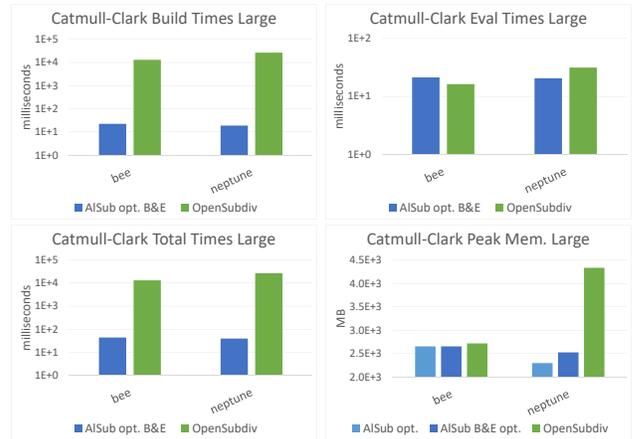
Fig. 14. Catmull-Clark subdivision: Comparison of the individual steps (build & eval) of AlSub and OpenSubdiv for large meshes including peak memory consumption for all approaches.

evaluation stage uses this information to subdivide the vertex data in every render, *e.g.*, when replaying an animation. Relying on AlSub's split into *build* and *eval* stages, similar optimizations are possible in our approach. The top rows of Figures 13 and 14 shows that when splitting our approach into these two phases, AlSub opt. achieves nearly one order of magnitude performance gain over AlSub SpLA in both *build* and *eval*. Furthermore, performing the *build* stage complete on the GPU, yields significant performance gains over OpenSubdivs build stage. Even more important for this use case is that AlSub opt. also outperforms OpenSubdiv in *eval* with a small margin. Note that this fact does not hold for the Armor Guy model with its high number of creases, for which our approach performs additional steps to 'fix' the geometry.

Considering the memory requirements, it becomes visible that there is a small overhead when splitting AlSub into *build* and *eval* phases, as additional data needs to be kept for *eval*. However, due to our efficient sparse matrix formulations and matrix optimizations, the memory required for AlSub is significantly below OpenSubdiv, up to two orders of magnitude in some cases. Note that our approach seems to be more efficient, when multiple iterations of subdivision can be performed. For a single iteration, such as for *bee*, the evaluation performance as well as the memory requirements are similar to OpenSubdiv.

Considering the sum of all these results, AlSub seems to be a suitable drop-in replacement for OpenSubdiv in the modeling and rendering use case, virtually removing preprocessing costs and significantly reducing memory requirements. With its negligible *build* cost AlSub opt. build+eval can even be used in a modeling context to further optimize rendering when no topological changes are required.

Adaptive Subdivision. To show that our approach could be used in a setting where only certain regions of the mesh have to be subdivided, we compare to the feature adaptive Catmull-Clark implementation of OpenSubdiv, which is based on the approach proposed by Nießner et al. [2012]. Here, only regions around irregularities
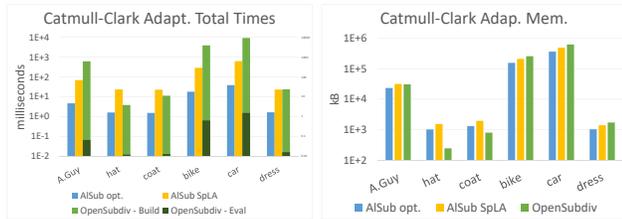
Fig. 15. Catmull-Clark adaptive results.

have to be subdivided. For regular mesh regions patches are built which can be evaluated using hardware tessellation. While this is only one specific use case of selective subdivision, there are many more, as *e.g.*, in path tracing where spatially close rays are gathered and evaluated in a batch.

Figure 15 compares performance and required peak memory of our approaches with those of OpenSubdiv. For OpenSubdiv we distinguish again between build and eval times. Our timings are the total times including topology and vertex position refinement. In all cases we only measure the time it takes to subdivide the irregular regions and discard regular patches, which could be evaluated using hardware tessellation. It is interesting to see that the evaluation performance is a lot higher for OpenSubdiv if the mesh has a very small number of faces to subdivide (hat, coat, dress). As the number of irregularities increases (ArmorGuy, bike, car), the margin to AlSub gets significantly smaller. This can be attributed to the fact, that when subdividing just a small number of faces, the overhead of our approach gets too large to compensate as we can by far not fully utilize the GPU.

7.4 Mesh reordering

To highlight the effect of mesh reordering, we compared the performance of different subdivision implementations on meshes in their original ordering and after reordering using the RCM method. From the results in Figure 16, it seems that our optimized kernels benefits most from better memory layouts.

Surprisingly, reordering can increase performance of up to $5 - 8\times$ in our optimized version for models which show a bad input data layout like the Beetle model. While the speedups of other approaches are also significant, the relative speedup of AlSub opt. is on average slightly higher than for the other approaches. We attribute the lower gains of SpLA to our optimized versions mainly operating on the input data directly due to kernel fusion while the SpLA version creates additional data structures, for which the data layout does not change significantly and always reduces performances compared to the optimized version. OpenSubdiv gains its speedups mainly in the eval step, while AlSub improves performance equally among build and eval, indicating that the memory access patterns is more important on the GPU than the CPU.

Given the significant speedups which can be gained from reordering, it seems natural to attempt to find a fast reordering which can be used after each iteration to consolidate the memory layout. Our attempts in this direction suggest that is a challenging problem since any gains get outweighed by the cost of reordering itself. Therefore, for scenarios such as production rendering where the topology does not change, it would be worthwhile to have reorderings precomputed for every few subdivision steps and deployed during batch processing tasks.

8 CONCLUSION

In this paper, we proposed a full fledged treatment of parallel mesh subdivision using linear algebra primitives and showed thereby, that efficient algorithmic solutions can be written in a higher level language which is accessible to all practitioners, as it only requires basic mathematics. Unlike traditional approaches, where bookkeeping stalls performance and impedes vectorization, our sleek mesh representation is an integral part of the solution and naturally evolves throughout the subdivision steps. While a direct implementation of this formulation already indicates high performance and opens the door for fine grained parallelism, the formulation itself sheds lights on operations which can be further optimized at a lower level, increasing performance by about one order of magnitude.

Our evaluation shows that our full subdivision approach significantly outperforms other approach in scenarios where an input mesh must be subdivided once. Splitting the subdivision into a preprocessing and evaluation step for cases where the topology of models does not change, makes it a direct competitor for OpenSubdiv. Performing the preprocessing step on the GPU shows extreme speedups compared to OpenSubdiv's CPU preprocessing. Additionally, our evaluation step is also faster than OpenSubdiv, while operating at significantly less memory. Furthermore, we have shown that our approach can handle widely used extensions to the standard algorithms, such as creases and boundaries as well as adaptive subdivision efficiently.

REFERENCES

RWTH Aachen-University. 2017. OpenMesh.

Stephen T. Barnard, Alex Pothen, and Horst D. Simon. 1993. A Spectral Algorithm for Envelope Reduction of Sparse Matrices. In *Proceedings of the 1993 ACM/IEEE Conference on Supercomputing (Supercomputing '93)*. ACM, New York, NY, USA, 493–502.

Bruce G. Baumgart. 1972. *Winged Edge Polyhedron Representation.* Technical Report. Stanford, CA, USA.

Jeffrey Bolz and Peter Schröder. 2002. Rapid Evaluation of Catmull-Clark Subdivision Surfaces. In *Proc. Web3D '02*. ACM, 11–17.

Jeffrey Bolz and Peter Schröder. 2003. Evaluation of Subdivision Surfaces on Programmable Graphics Hardware. (Mar 2003), 4.

Wade Brainerd, Tim Foley, Manuel Kraemer, Henry Moreton, and Matthias Nießner. 2016. Efficient GPU Rendering of Subdivision Surfaces Using Adaptive Quadtrees. *ACM Trans. Graph.* 35, 4, Article 113 (July 2016), 12 pages.

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.

Paul Castillo, Robert Rieben, and Daniel White. 2005. FEMSTER: An Object-oriented Class Library of High-order Discrete Differential Forms. *ACM Trans. Math. Softw.* 31, 4 (Dec. 2005), 425–457.

Edwin Catmull and James Clark. 1978. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design* 10, 6 (1978), 350 – 355.

George M. Chaikin. 1974. An algorithm for high speed curve generation. *Computer Graphics and Image Processing* 3 (1974), 346–349.

Ray W. Clough and J. L. Tocher. 1965. Finite Element Stiffness Matrices for Analysis of Plates in Bending. In *Proceedings of Conference on Matrix Methods in Structural Mechanics*. Wright-Patterson A. F. B., Ohio. Air Force Institute of Technology.

Robert L. Cook. 1984. Shade Trees. *SIGGRAPH Comput. Graph.* 18, 3 (Jan. 1984), 223–231.

Elizabeth Cuthill and James McKee. 1969. Reducing the Bandwidth of Sparse Symmetric Matrices. In *Proceedings of the 1969 24th National Conference (ACM '69)*. ACM, New York, NY, USA, 157–172.

Timothy Davis. 2006. *Direct Methods for Sparse Linear Systems.* Society for Industrial and Applied Mathematics.

Michael Deering. 1995. Geometry Compression. In *Proc. SIGGRAPH '95*. ACM, 13–20.
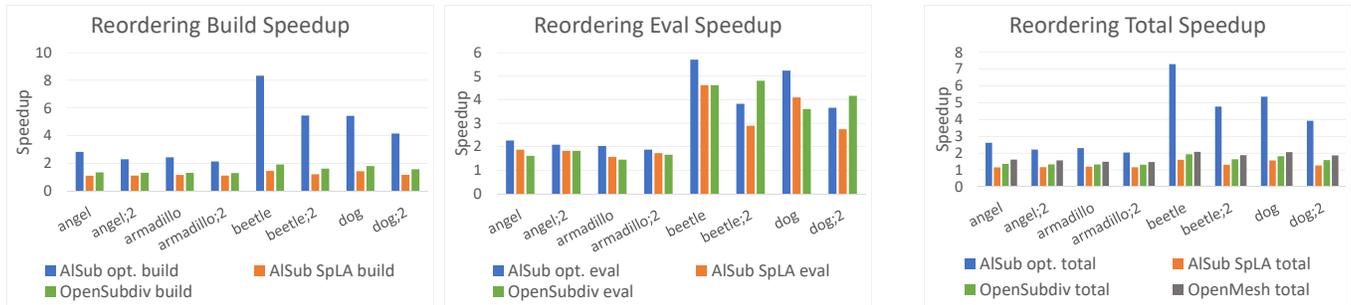
Fig. 16. Reordering results as relative speed up to the non-reordered meshes over one Catmull-Clark iteration (no postfix) and two iterations (';2' postfix).

Tony DeRose, Michael Kass, and Tien Truong. 1998. Subdivision Surfaces in Character Animation. In *Proc. SIGGRAPH '98*. ACM, 85–94.

Daniel Doo. 1978. A subdivision algorithm for smoothing down irregularly shaped polyhedrons. In *Proced. Int'l Conf. Ineractive Techniques in Computer Aided Design*. 157–165. Bologna, Italy, IEEE Computer Soc.

Daniel Doo and Malcolm Sabin. 1978a. Behaviour of Recursive Division Surfaces Near Extraordinary Points. *Computer-Aided Design* 10 (Sept. 1978), 356–360.

Daniel Doo and Malcolm Sabin. 1978b. Behaviour of recursive division surfaces near extraordinary points. *Computer-Aided Design* 10, 6 (1978), 356 – 360.

David R. Forsey and Richard H. Bartels. 1988. Hierarchical B-spline Refinement. In *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '88)*. ACM, New York, NY, USA, 205–212.

John A. George. 1971. *Computer Implementation of the Finite Element Method.* Ph.D. Dissertation. Stanford, CA, USA. AAI7205916.

Leonidas Guibas and Jorge Stolfi. 1985. Primitives for the Manipulation of General Subdivisions and the Computation of Voronoi. *ACM Trans. Graph.* 4, 2 (April 1985), 74–123.

Hugues Hoppe. 1999. Optimization of Mesh Locality for Transparent Vertex Caching. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 269–276.

Leif Kobbelt. 2000. $\sqrt{3}$-subdivision. In *Proc. SIGGRAPH '00*. ACM, 103–112.

Pascal Lienhardt. 1994. N-dimensional generalized combinatorial maps and cellular quasi-manifolds. *International Journal of Computational Geometry & Applications* 4, 03 (1994), 275–324.

Charles Loop. 1987. Smooth subdivision surfaces based on triangles. (1987).

Charles Loop and Scott Schaefer. 2008. Approximating Catmull-Clark Subdivision Surfaces with Bicubic Patches. *ACM Trans. Graph.* 27, 1, Article 8 (March 2008), 11 pages.

Gordon E. Moore. 2000. Readings in Computer Architecture. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Chapter Cramming More Components Onto Integrated Circuits, 56–59.

Johannes S. Mueller-Roemer, Christian Altenhofen, and André Stork. 2017. Ternary Sparse Matrix Representation for Volumetric Mesh Subdivision and Processing on GPUs. *Computer Graphics Forum* 36, 5 (2017), 59–69.

Ahmad H. Nasri. 1987. Polyhedral Subdivision Methods for Free-form Surfaces. *ACM Trans. Graph.* 6, 1 (Jan. 1987), 29–73.

Matthias Nießner and Charles Loop. 2013. Analytic Displacement Mapping Using Hardware Tessellation. *ACM Trans. Graph.* 32, 3, Article 26 (July 2013), 9 pages.

Matthias Nießner, Charles Loop, Mark Meyer, and Tony DeRose. 2012. Feature-adaptive GPU Rendering of Catmull-Clark Subdivision Surfaces. *ACM Trans. Graph.* 31, 1, Article 6 (Feb. 2012), 11 pages.

Jörg Peters. 2000. Patching Catmull-Clark Meshes. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 255–258.

Graphics Technologies Pixar. 2017. OpenSubdiv.

Michael J. D. Powell and Malcolm Sabin. 1977. Piecewise Quadratic Approximations on Triangles. *ACM Trans. Math. Softw.* 3, 4 (Dec. 1977), 316–325.

Youcef Saad. 1994. *SPARSKIT: A Basic Tool Kit for Sparse Matrix Computations.* Technical Report. Computer Science Department, University of Minnesota, Minneapolis, MN 55455.

Henry Schäfer, Jens Raab, Benjamin Keinert, Mark Meyer, Marc Stamminger, and Matthias Nießner. 2015. Dynamic Feature-adaptive Subdivision. In *Proc. of the 19th Symposium on Interactive 3D Graphics and Games (i3D '15)*. ACM, 31–38.

Le-Jeng Shiue, Ian Jones, and Jörg Peters. 2005. A Realtime GPU Subdivision Kernel. In *ACM SIGGRAPH 2005 Papers (SIGGRAPH '05)*. ACM, New York, NY, USA, 1010–1015.

Jos Stam. 1998. Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values. In *Proc. SIGGRAPH '98*. ACM, 395–404.

Markus Steinberger, Andreas Derler, Rhaleb Zayer, and Hans-Peter Seidel. 2016. How naive is naive SpMV on the GPU?. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*. 1–8.

Herb Sutter. 2005. The free lunch is over: A fundamental turn toward concurrency in software. *Dr. Dobb's Journal* 30, 3 (2005).

Stanley Tzeng, Anjul Patney, and John D. Owens. 2010. Task Management for Irregular-parallel Workloads on the GPU. In *Proceedings of the Conference on High Performance Graphics (HPG '10)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 29–37.

Huy T. Vo, Claudio T. Silva, Luiz F. Scheidegger, and Valerio Pascucci. 2012. Simple and Efficient Mesh Layout with Space-Filling Curves. *Journal of Graphics Tools* 16, 1 (2012), 25–39.

Rhaleb Zayer, Markus Steinberger, and Hans-Peter Seidel. 2017. A GPU-Adapted Structure for Unstructured Grids. *Computer Graphics Forum* (2017).

Kun Zhou, Qiming Hou, Zhong Ren, Minmin Gong, Xin Sun, and Baining Guo. 2009. RenderAnts: Interactive Reyes Rendering on GPUs. In *Proc. SIGGRAPH Asia 2009*. ACM, Article 155, 11 pages.