Electronic Supplement to Dyadic T-mesh Subdivision

Denis Kovacs[1,2*] Justin Bisceglio[1,3†]
Denis Zorin[1‡]
[1]New York University [2]FiftyThree [3]Blue Sky Studios

1 Proof of Exhaustive Enumeration of Subdivision Stencils

We use notation introduced in the Appendix of the paper. To compute the value of P_j^1 we need to define a stencil of control points P_i^0 that influence its value. We exhaustively enumerate the 1-ring neighborhood configurations (with a suitable extension at T-vertices and T-edges) of a vertex in all possible T-mesh configurations.

Here we show that such a configuration $\mathcal{N}$ in fact covers the stencil, i.e. no P_i^0 outside $\mathcal{N}$ can influence P_j^1.

We split the parametric plane outside $\mathcal{N}$ into two parts: (1) half-slabs H_k in which a basis function B_i^0 cannot affect P_j^1 since its cross would have to intersect two edges in $\mathcal{N}$ before S_i^0 would cover S_j^1, and (2) the remaining regions (*corner zones*) C_k for which we use a Lemma from [da Veiga et al. 2012].

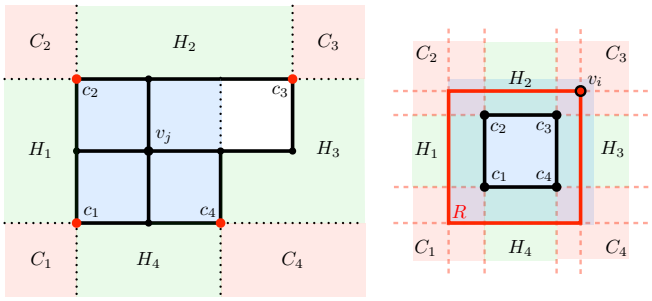

Figure 1: *Half-slab H_k (light green) and corner zones C_k (light red), corner points c_k (red), basis support S_j^1 of vertex v_j (blue). Right: minimal rectangle R that a quadrant of v_i's basis function needs to contain.*

Let

$$S_j^1 = [s_0 \dots s_1] \times [t_0 \dots t_1].$$

Half-slabs H_k are constructed for each of the four directions t,$-t$, s and $-s$. Consider a line ℓ_s in one of these directions (e.g., horizontal direction $-s$) that intersects the stencil candidate $\mathcal{N}$. Each such line is intersected by at least two edges of the stencil (this can be verified for each stencil candidate directly). Consider the set of points on ℓ_s separated from the *right* boundary of S_j^1 by two stencil edge intersections with ℓ_s. These points cannot be control points in the stencil. The union of such points for all lines ℓ_s form the half-slab H_1.

For every $\mathcal{N}$ there are four *corner points* where two H_k intersect. They bound one of the open regions C_k which we will call *corner zone*. Since the entire plane is covered by these regions:

$$\mathbb{R}^2 = \mathcal{N} \cup \bigcup_{k=1}^{t} H_k \cup C_k,$$

all that is left to do is to prove that any B_i^0 centered in any corner zone C_k can not contain S_j^1. To show this, we use *Lemma 4.2* from [da Veiga et al. 2012]. The vertex v_i is called an active T-mesh node if it is sufficiently far away from any boundary where there are enough knots to define its basis function. $\mathrm{TF}(v_i)$ is the *tiled floor* of v_i, i.e. the support of B_i^0 excluding the 5x5 grid of knot lines.

Lemma 1 *Let $\mathcal{M}$ be an AS T-mesh and v_i an active T-mesh node. Then $\mathrm{TF}(v_i)$ does not contain any T-mesh node.*

As a consequence of this Lemma, in the case of dyadic T-meshes, all T-mesh vertices v_i in a quadrant Q of a basis function B_i^0, have to be a subset of one of two possible configurations (see Figure 2).

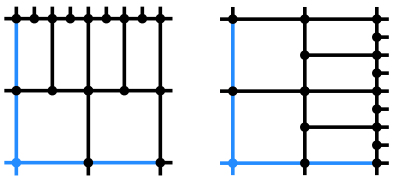

Figure 2: *Possible quadrants for DAS T-spline basis functions, visible part of basis cross marked blue*

This can be seen as follows: There can be no vertices or crossing edges on the basis cross other than the two that define the basis cross (otherwise the basis cross would be shorter), so the densest regular grid we can define on Q has 3×3 knots. W.l.o.g., let us assume there are only vertical T-edges in this grid. It contains at least one T-joint with a horizontal stem – otherwise the support of the basis function B_i^0 would be smaller. Wherever the T-joint is located, its face and edge extension together span the entire s-span of Q, making it impossible to add a vertical T-joint anywhere without violating the analysis-suitable rule that no horizontal and vertical T-joint extensions intersect.

So in one quadrant, analysis-suitable T-meshes can only have either horizontal or vertical T-joints, but not both. Since we are only considering dyadic T-meshes, there can be at most one T-joint per edge. Hence the densest Q given a fixed basis cross is defined by the cascaded T-joint pattern shown in Figure 2. It is clear that the s and t knots of any other refinement are contained in the knots of this Q.

We can hence conclude that there is one dimension along which there can be no more than 3 knots in Q.

To prove that no $v_i \in C_k$ exists such that $S_i^0 \supseteq S_j^1$, we show that the relevant quadrant Q of a basis function B_i^0 associated with such a v_i requires at least 4 knots along both dimensions.

A necessary condition for $S_i^0 \supseteq S_j^1$ is that Q contains the rectangle R spanned by v_i and the corner of S_j^1 diagonally opposite to v_i (Figure 1 right).

We then collect the knots of all vertices in R, and include the s- and t-extents of R to ensure that Q indeed contains R. Recall that the total knot count cannot exceed three in both dimensions simultaneously.

For each previously listed stencil, however, we verified by this simple counting scheme that for each stencil C_k, the number of knots in R is always ≥ 4 in each dimension. Hence, there are no $v_i \in C_k$ such that their support contains S_j^1, and there are no outside control points P_i^0 that affect P_j^1.

To verify our factorization reduces to analysis-suitable T-Splines, now all we have to do is to verify that it yields the same results we obtain with the refinement formulas in the paper for every vertex of every stencil connectivity enumerated above.

*e-mail: kovacs@cs.nyu.edu
†e-mail: justinb@blueskystudios.com
‡e-mail: dzorin@cs.nyu.edu

Figure 3: *DAS T-mesh face stencils.*

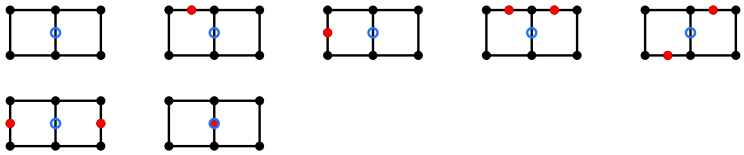

Figure 4: *DAS T-mesh edge stencils.*

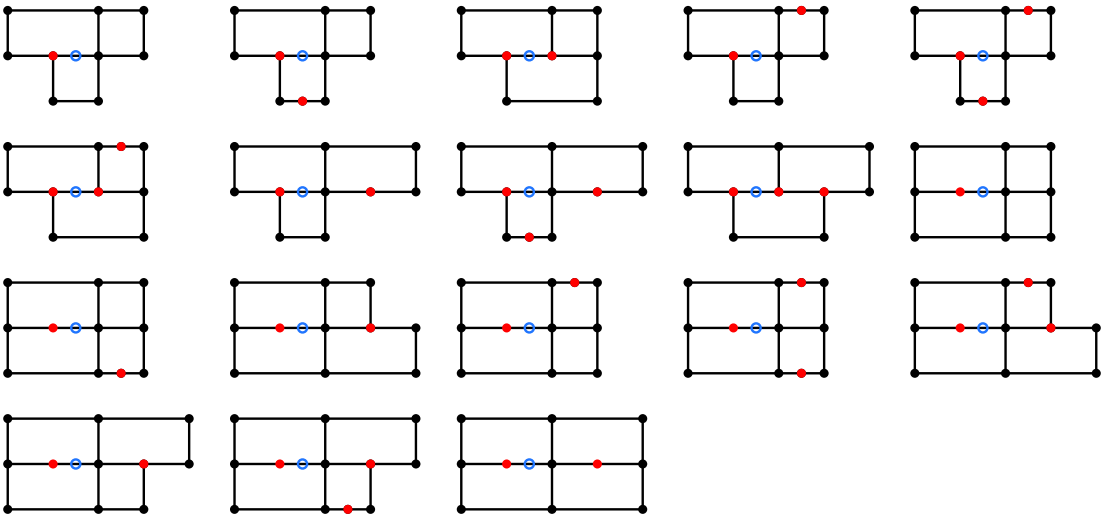

Figure 5: *DAS T-mesh T-edge stencils.*

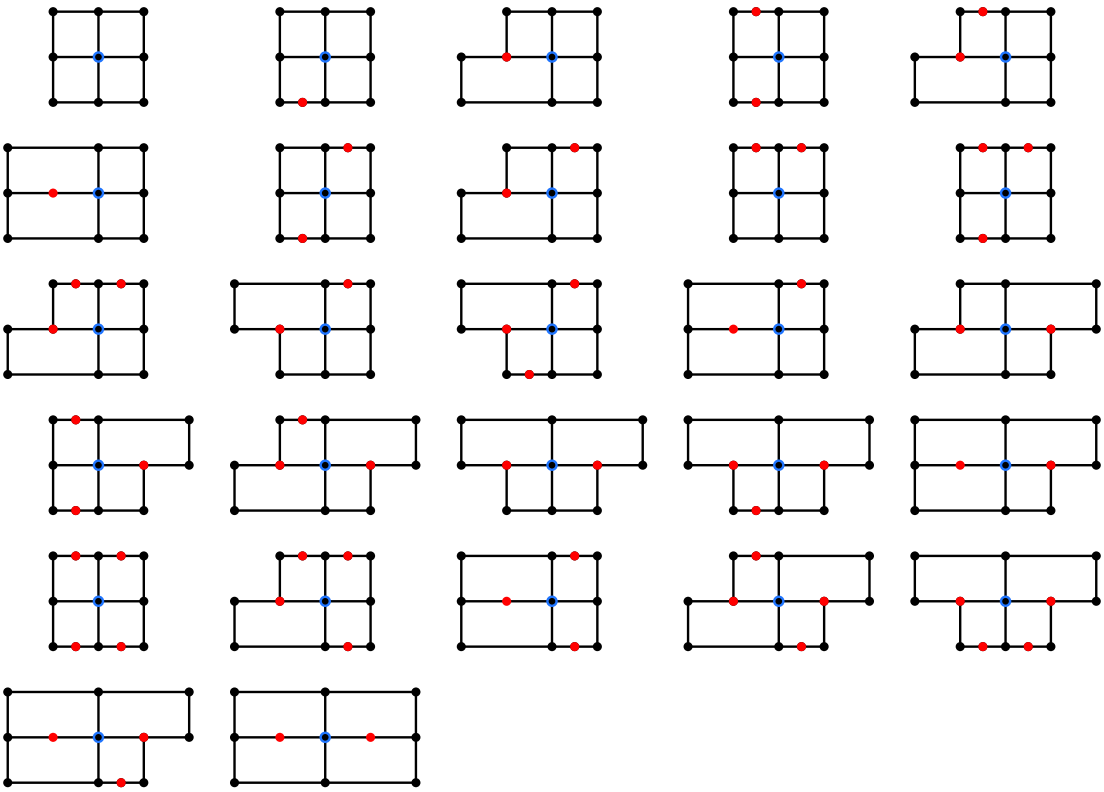

Figure 6: *DAS T-mesh vertex stencils.*

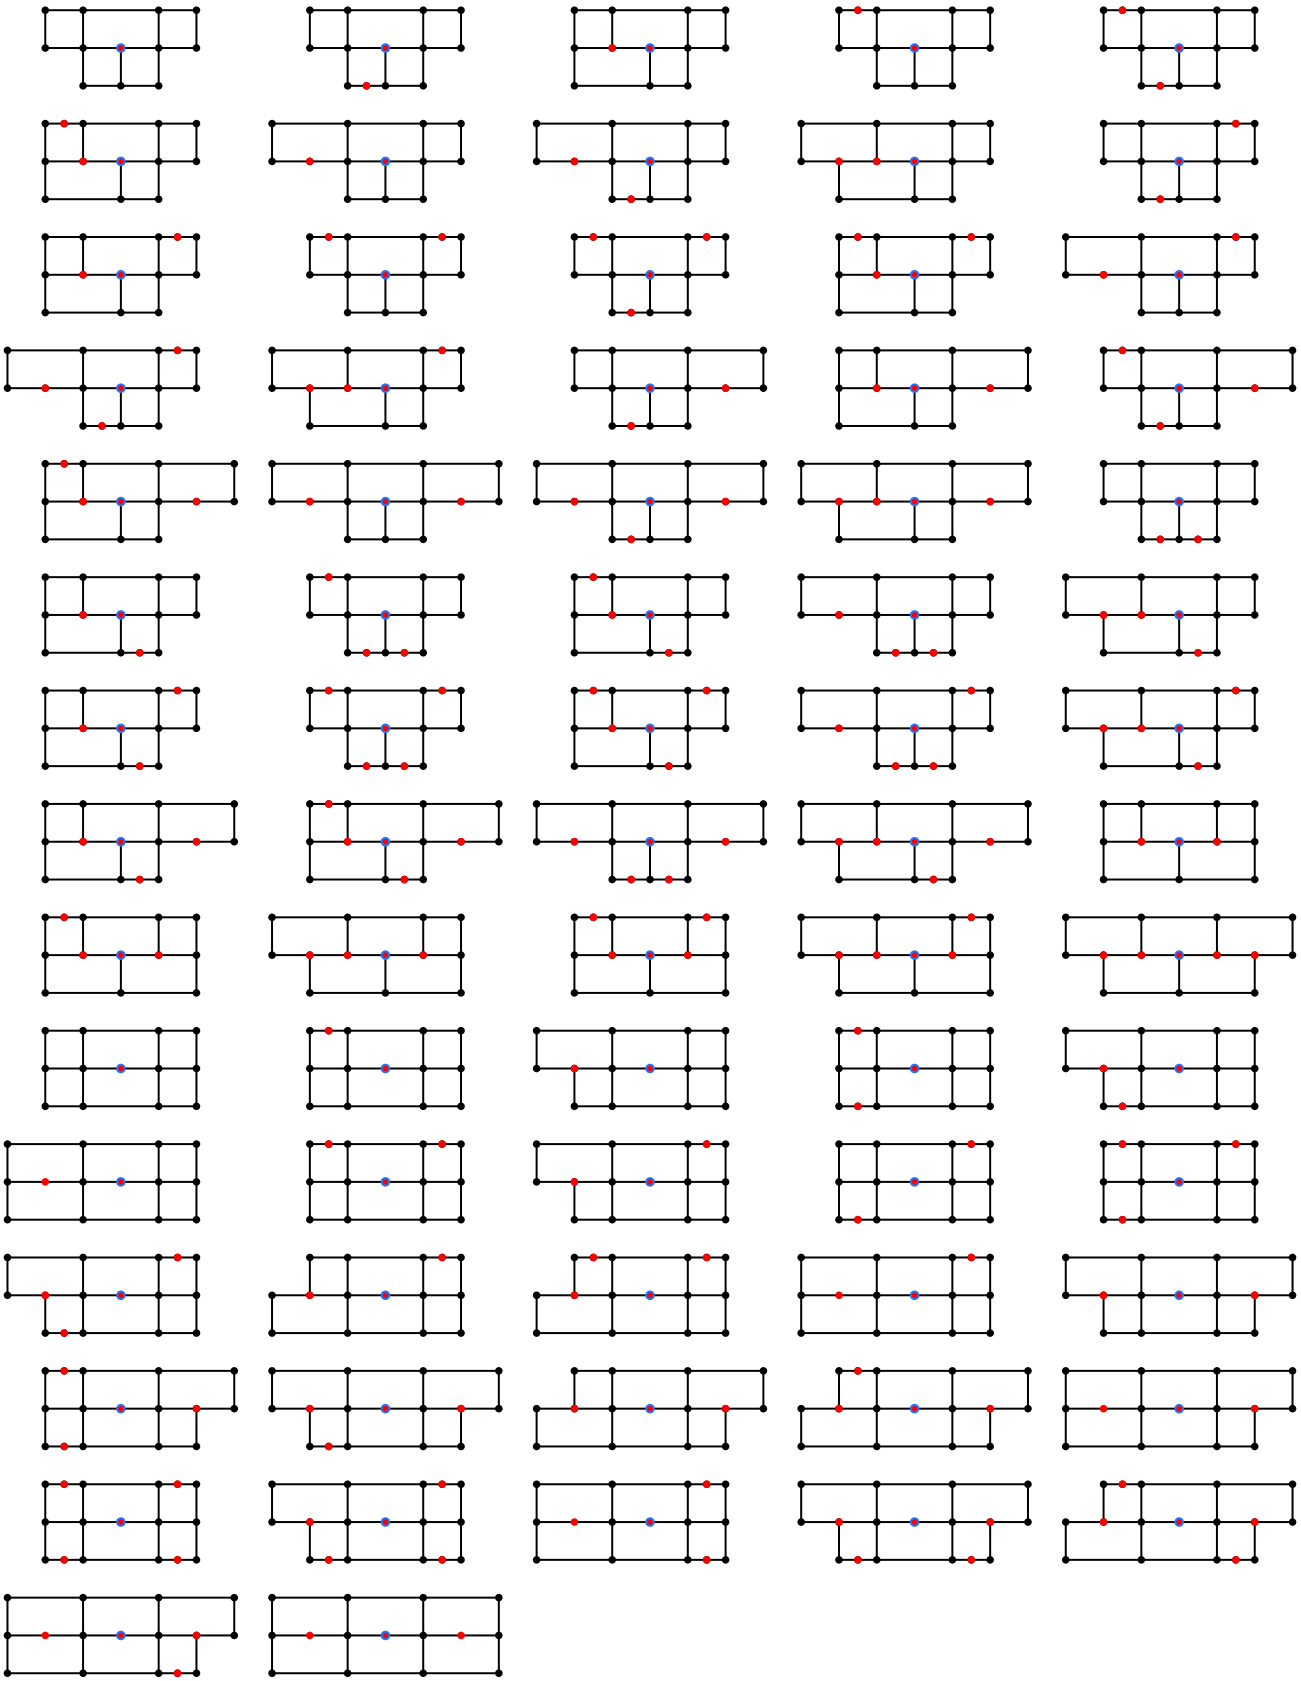

Figure 7: *DAS T-mesh T-vertex stencils.*

1.1 Characteristic Maps and Tangent Plane Analysis

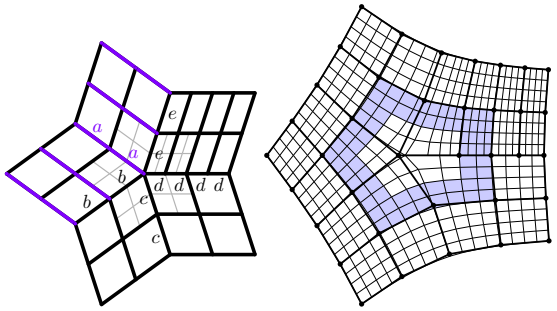

Figure 8: *Left: limit topology around an extraordinary vertex. Right: ring of Bezier patches extracted from limit stencil*

A complete analysis of tangent plane continuity at extraordinary vertices is relatively complex, due to the large number of configurations that need to be considered. However, with one additional assumption, a finite enumeration for moderate vertex valences is possible.

Recall that for a mesh a number of knot values can be chosen independently, with the rest determined by compatibility conditions. Specifically, if we assume that all independent knot intervals are set to 1, then a finite (although very large) enumeration of cases of self-reproducing connectivities near extraordinary vertex is possible.

To determine the limit behavior at an extraordinary vertex, we can assume that enough subdivision steps have occured that the topology around the vertex is *self-similar*, i.e. the control mesh of the set of patches around the extraordinary vertex is the same at all subdivision levels. To characterize these topologies, we define a *spoke* to be the edge of the unsubdivided mesh (and all T-joints are regular) incident at the extraordinary vertex of interest. Then self-similar configurations are characterized by the following conditions:

- there is a single extraordinary vertex in the control mesh, and it is not a T-joint;
- knot intervals on spokes are equal;
- there are only T-joints along spokes. A row of faces along a spoke either all have T-joints on the spoke or none of them do.

These conditions allows us to characterize a configuration by a small number of parameters (Fig. 8): 1) the valence, 2) the knot interval of the form $1/2^i$ for each spoke (all other knot intervals are determined by compatibility constraints, and the intervals on adjacent spokes cannot differ by more than a factor of two) 3) whether there are T-joints on a spoke and to which side their stem is pointing.

The two-ring control mesh for the central ring of patches is obtained by taking the vertices of the union of quads forming 2×2 grids in each of k sectors for a vertex of valence k. We note that scaling all knots by the same amount does not change local surface behavior, so one of the knot intervals in the self-similar control mesh can be always chosen to be 1, and the rest set with respect to it.

We enumerate possible configurations by going over all combinations of parameter values and checking the analysis-suitable conditions. Of course the number of configurations grows exponentially, so the method is practical only for sufficiently low valences (up to $n = 9$).

We use the standard approach to verifying C^1 continuity for spline-based schemes [Reif 1995; Peters and Reif 1998].

First, we construct the subdivision matrix mapping the control points of the two-ring to the points of the two-ring on the next refinement level, and compute its subdominant eigenvalues and eigenvectors x^ℓ, $\ell = 1, 2$ with components x^ℓ_i. The two-dimensional control mesh with control points (x^1_i, x^2_i) define the control mesh for the *characteristic map* from the plane to the plane. Nonvanishing Jacobians and bijectivity of the characteristic map are sufficient for C^1-continuity. The characteristic map is also self-similar (i.e., its values on a nested sequence of ring domains are obtained by scaling), so it is sufficient to examine it on a single ring domain. The ring domain is obtained as a set of patches forming outer rings after two subdivision steps (Fig. 8). As there are no extraordinary vertices in the control meshes of these patches, all subdivision rules affecting the limit surface on these patches are just analysis-suitable T-spline rules, and patches are polynomial.

For each patch, nonegativity of the Jacobian can be verified explicitly, by computing the Jacobian as a polynomial and converting it to the Bezier form. Positivity of Bezier coefficients of the Jacobian is sufficient. Finally, global bijectivity can be inferred from local bijectivity by simple winding number tests as shown in [Zorin 2000].

Control Meshes of Characteristic Maps for Dyadic T-meshes

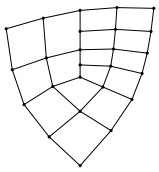

Figure 9: *DAS T-mesh valence 3 characteristic map.*

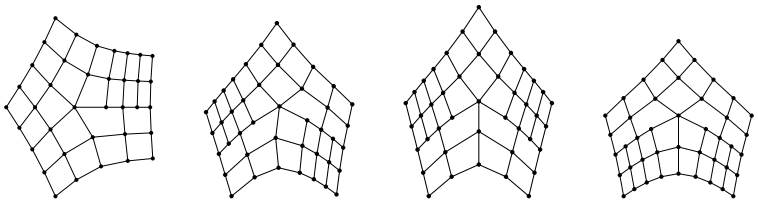

Figure 10: *DAS T-mesh valence 5 characteristic maps.*

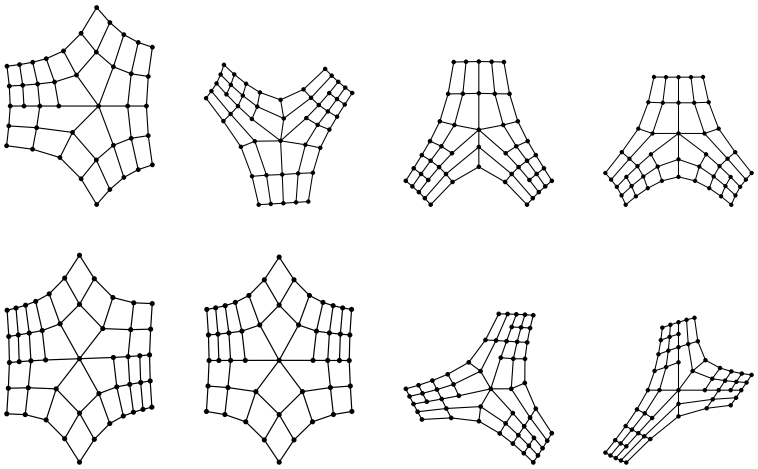

Figure 11: *DAS T-mesh valence 6 characteristic maps.*

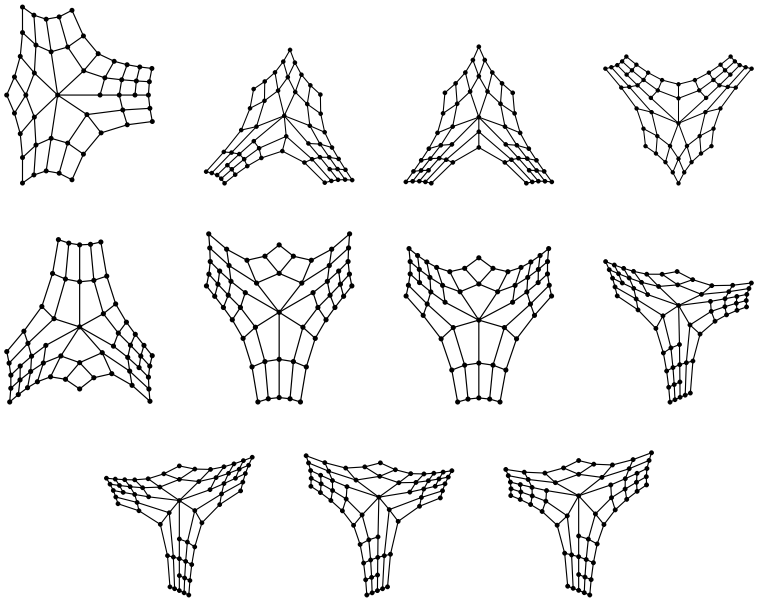

Figure 12: *DAS T-mesh valence 7 characteristic maps.*

References

DA VEIGA, L. B., BUFFA, A., CHO, D., AND SANGALLI, G. 2012. Analysis-suitable t-splines are dual-compatible. *Computer Methods in Applied Mechanics and Engineering 249–252*, 0, 42 – 51. Higher Order Finite Element and Isogeometric Methods.

PETERS, J., AND REIF, U. 1998. Analysis of algorithms generalizing b-spline subdivision. *SIAM Journal on Numerical Analysis 35*, 2, 728–748.

REIF, U. 1995. A unified approach to subdivision algorithms near extraordinary vertices. *Computer Aided Geometric Design 12*, 2, 153–174.

ZORIN, D. 2000. A method for analysis of c 1-continuity of subdivision surfaces. *SIAM Journal on Numerical Analysis 37*, 5, 1677–1708.