

Parallel Vectors Extraction using Bézier Clipping – Additional Material –

Nico Daßler¹ and Tobias Günther¹

Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany




Figure 1: Results on the BOEING dataset using Newton-Raphson only [PR99] (left) and our Bézier Clipping approach (right). Both methods use the same bicubic interpolant on cell faces. The Newton-Raphson method misses 15% of the solutions due to insufficient initial guesses.

1. Additional Performance Measurements

Table 1 and Table 2 list additional performance measurements for the independent bicubic interpolant and for the bicubic interpolant with derived acceleration. The behavior is overall the same as with the bilinear case. The component-based Bézier clipping consistently outperforms the bisection method, and the projection-based Bézier clipping consistently outperforms the component-based clipping. Due to the higher polynomial degree of the root finding problem, the bicubic interpolant with derived acceleration is the slowest, followed by bicubic interpolation. The bilinear interpolant remains to be the fastest. Across all datasets, it is by a factor of about $8.2\times$ faster than the bicubic interpolant. We also observe a slightly larger speedup when using one of our clipping solvers ($8.31\times$) compared to the bisection solver ($8.01\times$). The bicubic interpolation is about $3.47\times$ faster than the same with derived acceleration. For the solve runtime, we see an average speedup of about $7.28\times$ of the bilinear interpolant over the bicubic interpolant and a $3.05\times$ speedup of the bicubic interpolant over the bicubic interpolant with derived acceleration. The bisection solver again scales slightly worse in this scenario with a speedup of only about $2.63\times$, compared to the projection-based clipping solver with a speedup factor of about $3.76\times$.

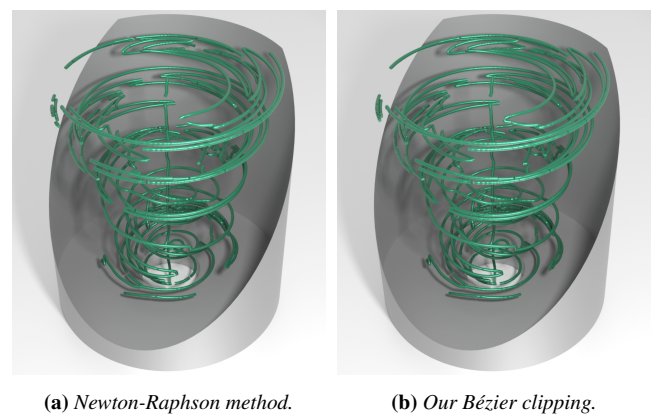


Figure 2: Results on the SWIRLING JET dataset using Newton-Raphson and our Bézier Clipping approach with bicubic interpolant and derived acceleration. The Newton-Raphson method misses 0.49% of solutions, resulting in minor visual differences.

2. Additional Comparisons

Table 1 shows that the Newton-Raphson solver misses about 15% of the solutions compared to the Bézier clipping solver on the BOE-

Dataset	Algorithm	Runtime (Serial)	Runtime (Parallel)	Solve Runtime	# Duplicates	# Solutions	\varnothing Sol. Depth	\varnothing Clip Size
DELTA WING	Bisection	736.70 \pm 6.61	50.68 \pm 0.27	104.12 \pm 0.19	924	2398	45.78	0.500
	Clipping (C)	672.01 \pm 16.91	47.49 \pm 1.24	49.66 \pm 0.31	95	2398	17.04	0.723
	Clipping (P)	672.87 \pm 0.19	43.74 \pm 0.95	18.15 \pm 0.03	0	2398	7.92	0.866
	Newton-Raphson	654.80 \pm 2.46	41.83 \pm 0.49	54.83 \pm 0.18	0	2299	—	—
BORROMEAN	Bisection	1766.03 \pm 5.00	115.25 \pm 0.28	351.34 \pm 0.56	3862	7002	42.95	0.500
	Clipping (C)	1574.61 \pm 3.93	99.90 \pm 0.18	156.21 \pm 0.38	555	7002	18.39	0.692
	Clipping (P)	1455.86 \pm 5.06	92.55 \pm 0.21	40.45 \pm 0.08	0	7002	6.39	0.953
	Newton-Raphson	1551.81 \pm 3.48	95.67 \pm 0.38	124.05 \pm 0.24	0	6973	—	—
SWIRLING JET	Bisection	8016.07 \pm 18.11	523.40 \pm 0.87	1446.89 \pm 5.81	13345	11765	68.90	0.500
	Clipping (C)	7028.38 \pm 26.62	452.12 \pm 0.83	499.11 \pm 2.11	2585	11765	31.21	0.639
	Clipping (P)	6618.40 \pm 7.68	423.61 \pm 0.74	78.78 \pm 0.14	88	11765	7.27	0.957
	Newton-Raphson	6897.20 \pm 21.22	434.04 \pm 1.44	326.08 \pm 0.22	0	11708	—	—
BOEING	Bisection	3860.64 \pm 24.13	294.56 \pm 0.98	1977.25 \pm 5.22	23740	29511	49.25	0.500
	Clipping (C)	2780.34 \pm 71.16	188.00 \pm 0.82	973.74 \pm 4.37	2973	29511	25.41	0.635
	Clipping (P)	2246.04 \pm 23.74	150.55 \pm 0.71	330.83 \pm 0.49	0	29511	10.6	0.752
	Newton-Raphson	2874.58 \pm 4.56	176.61 \pm 0.35	1134.34 \pm 1.47	0	24932	—	—
NINE CL	Bisection	2867.64 \pm 21.73	190.55 \pm 0.35	35.08 \pm 0.19	0	1625	46.85	0.500
	Clipping (C)	2840.56 \pm 1.85	187.36 \pm 0.24	8.20 \pm 0.08	0	1625	5.23	0.996
	Clipping (P)	2822.20 \pm 11.60	185.98 \pm 0.46	6.73 \pm 0.01	0	1625	5.23	0.995
	Newton-Raphson	2898.47 \pm 8.16	188.16 \pm 0.59	3.18 \pm 0.02	0	1625	—	—

Table 1: Metrics recorded for bicubic interpolation on different datasets using the bisection method, our component-wise clipping (C), our projection-based clipping (P), and the Newton-Raphson method. All runtimes are measured and averaged for 10 runs, listing mean and standard deviation.

Dataset	Algorithm	Runtime (Serial)	Runtime (Parallel)	Solve Runtime	# Duplicates	# Solutions	\varnothing Sol. Depth	\varnothing Clip Size
DELTA WING	Bisection	1020.94 \pm 2.19	68.44 \pm 0.17	324.44 \pm 1.69	2194	3693	40.05	0.500
	Clipping (C)	883.40 \pm 3.57	56.34 \pm 0.22	183.33 \pm 0.25	188	3693	19.36	0.652
	Clipping (P)	778.63 \pm 0.26	49.59 \pm 0.18	64.65 \pm 0.14	0	3693	9.14	0.785
	Newton-Raphson	805.31 \pm 0.23	50.46 \pm 0.13	145.37 \pm 0.07	0	2928	—	—
BORROMEAN	Bisection	6862.38 \pm 7.41	463.37 \pm 1.07	596.15 \pm 0.96	5685	7281	36.15	0.500
	Clipping (C)	6593.85 \pm 3.08	442.83 \pm 1.10	310.25 \pm 1.09	599	7281	16.52	0.686
	Clipping (P)	6357.05 \pm 4.69	431.21 \pm 1.16	86.73 \pm 0.12	0	7281	6.44	0.926
	Newton-Raphson	6484.48 \pm 3.39	438.23 \pm 1.34	248.00 \pm 0.35	0	7090	—	—
SWIRLING JET	Bisection	31801.73 \pm 34.19	2166.73 \pm 5.21	2313.39 \pm 14.15	12674	12640	55.32	0.500
	Clipping (C)	30396.87 \pm 10.00	2068.10 \pm 5.66	926.76 \pm 2.47	1942	12640	26.71	0.635
	Clipping (P)	29689.55 \pm 21.24	2034.15 \pm 5.85	195.51 \pm 0.27	438	12640	7.53	0.883
	Newton-Raphson	29953.97 \pm 6.88	2052.56 \pm 6.12	656.98 \pm 0.88	0	12134	—	—
BOEING	Bisection	12089.95 \pm 53.28	831.96 \pm 1.91	10158.63 \pm 55.16	48344	55133	65.12	0.500
	Clipping (C)	6367.85 \pm 16.18	391.01 \pm 2.15	4490.11 \pm 8.13	6818	55133	32.22	0.640
	Clipping (P)	4839.47 \pm 7.64	682.23 \pm 6.17	2911.22 \pm 9.76	8579	55133	13.19	0.649
	Newton-Raphson	4252.16 \pm 22.44	253.75 \pm 0.88	2494.52 \pm 11.98	0	36025	—	—
NINE CL	Bisection	14747.24 \pm 33.03	1016.16 \pm 3.65	55.95 \pm 0.39	0	1625	34.62	0.500
	Clipping (C)	14675.09 \pm 2.36	1010.99 \pm 3.32	14.79 \pm 0.09	0	1625	4.31	0.994
	Clipping (P)	14664.73 \pm 11.16	1014.50 \pm 3.35	12.23 \pm 0.02	0	1625	4.62	0.993
	Newton-Raphson	14639.45 \pm 7.57	1014.59 \pm 3.19	3.83 \pm 0.03	0	1625	—	—

Table 2: Metrics recorded for bicubic interpolation with derived acceleration to compare the performance on different datasets using the bisection method, our component-wise clipping (C), our projection-based clipping (P), and the Newton-Raphson method. All runtimes are measured and averaged for 10 runs, listing mean and standard deviation.

ING dataset. Figure 1 shows the extracted feature curves for this experiment using the Newton-Raphson method on the left and our projection-based Bézier clipping solver on the right. Both methods used the same bicubic interpolant on cell faces. The missing solutions are directly reflected in the rendered results. This suggests that a Newton-Raphson solver is not sufficient for this dataset and interpolant combination. In contrast, the Newton-Raphson solver was better suited for the SWIRLING JET dataset, when using bicubic interpolation and derived acceleration. Table 2 shows that only about 0.49% of the solutions are missing, resulting in minor visual differences in the extracted feature curves, which can be seen in

Figure 2. However, for both datasets the clipping solver is still preferred due to its faster solve runtime, achieving a speedup factor of about 3.4 \times on the BOEING dataset and a speedup factor of about 3.3 \times on the SWIRLING JET.

3. Cross Product Components of Two Bilinear Fields

In the special case of $\mathbf{v}(\mathbf{x})$ and $\mathbf{w}(\mathbf{x})$ being both bilinearly interpolated across a cell face, the biquadratic cross product can be computed symbolically in Bézier form. We list the control points of the

z -component of the cross product surface $\mathbf{c}(u, v)$ explicitly:

$$\mathbf{c}_{00}^{[z]} = \mathbf{w}_{00}^{[x]} \cdot \mathbf{v}_{00}^{[y]} - \mathbf{w}_{00}^{[y]} \cdot \mathbf{v}_{00}^{[x]} \quad (1)$$

$$\mathbf{c}_{01}^{[z]} = \frac{1}{2} \left(\mathbf{w}_{00}^{[x]} \cdot \mathbf{v}_{01}^{[y]} + \mathbf{w}_{01}^{[x]} \cdot \mathbf{v}_{00}^{[y]} - \mathbf{w}_{00}^{[y]} \cdot \mathbf{v}_{01}^{[x]} - \mathbf{w}_{01}^{[y]} \cdot \mathbf{v}_{00}^{[x]} \right) \quad (2)$$

$$\mathbf{c}_{02}^{[z]} = \mathbf{w}_{01}^{[x]} \cdot \mathbf{v}_{01}^{[y]} - \mathbf{w}_{01}^{[y]} \cdot \mathbf{v}_{01}^{[x]} \quad (3)$$

$$\mathbf{c}_{10}^{[z]} = \frac{1}{2} \left(\mathbf{w}_{00}^{[x]} \cdot \mathbf{v}_{10}^{[y]} + \mathbf{w}_{10}^{[x]} \cdot \mathbf{v}_{00}^{[y]} - \mathbf{w}_{00}^{[y]} \cdot \mathbf{v}_{10}^{[x]} - \mathbf{w}_{10}^{[y]} \cdot \mathbf{v}_{00}^{[x]} \right) \quad (4)$$

$$\mathbf{c}_{11}^{[z]} = \frac{1}{4} \left(\mathbf{w}_{00}^{[x]} \cdot \mathbf{v}_{11}^{[y]} + \mathbf{w}_{01}^{[x]} \cdot \mathbf{v}_{10}^{[y]} + \mathbf{w}_{10}^{[x]} \cdot \mathbf{v}_{01}^{[y]} + \mathbf{w}_{11}^{[x]} \cdot \mathbf{v}_{00}^{[y]} \right. \quad (5)$$

$$\left. - \mathbf{w}_{00}^{[y]} \cdot \mathbf{v}_{11}^{[x]} - \mathbf{w}_{01}^{[y]} \cdot \mathbf{v}_{10}^{[x]} - \mathbf{w}_{10}^{[y]} \cdot \mathbf{v}_{01}^{[x]} - \mathbf{w}_{11}^{[y]} \cdot \mathbf{v}_{00}^{[x]} \right) \quad (6)$$

$$\mathbf{c}_{12}^{[z]} = \frac{1}{2} \left(\mathbf{w}_{01}^{[x]} \cdot \mathbf{v}_{11}^{[y]} + \mathbf{w}_{11}^{[x]} \cdot \mathbf{v}_{01}^{[y]} - \mathbf{w}_{01}^{[y]} \cdot \mathbf{v}_{11}^{[x]} - \mathbf{w}_{11}^{[y]} \cdot \mathbf{v}_{01}^{[x]} \right) \quad (7)$$

$$\mathbf{c}_{20}^{[z]} = \mathbf{w}_{10}^{[x]} \cdot \mathbf{v}_{10}^{[y]} - \mathbf{w}_{10}^{[y]} \cdot \mathbf{v}_{10}^{[x]} \quad (8)$$

$$\mathbf{c}_{21}^{[z]} = \frac{1}{2} \left(\mathbf{w}_{10}^{[x]} \cdot \mathbf{v}_{11}^{[y]} + \mathbf{w}_{11}^{[x]} \cdot \mathbf{v}_{10}^{[y]} - \mathbf{w}_{10}^{[y]} \cdot \mathbf{v}_{11}^{[x]} - \mathbf{w}_{11}^{[y]} \cdot \mathbf{v}_{10}^{[x]} \right) \quad (9)$$

$$\mathbf{c}_{22}^{[z]} = \mathbf{w}_{11}^{[x]} \cdot \mathbf{v}_{11}^{[y]} - \mathbf{w}_{11}^{[y]} \cdot \mathbf{v}_{11}^{[x]} \quad (10)$$

Analogously, the x -components of the cross product $\mathbf{c}_{ij}^{[x]}$ are found by selecting the vector components $[z]$ and $[x]$ of \mathbf{v} and \mathbf{w} , and the y -components $\mathbf{c}_{ij}^{[y]}$ of the cross product use the components $[y]$ and $[z]$ of \mathbf{v} and \mathbf{w} . In the main paper, we derived the cross product generally for input fields of arbitrary Bézier degree, which includes the biquadratic cross product as a special case.

4. Interpolant Comparison

Our Bézier clipping based parallel vectors solver is applicable to higher-order polynomial interpolants. To examine the differences between a piecewise linear [PR99], a bilinear [BRG21], a bicubic interpolant (Section 3.2 in main paper), and the derived acceleration (Section 3.3 in main paper), we introduce an analytic vector field with known ground truth corelines. The NINECL data set, is a 3D steady vector field $\mathbf{v}(x, y, z)$ defined in the domain $[-2, 2]^3$:

$$\mathbf{v}(x, y, z) = \begin{pmatrix} -y(1-y)(1+y) \\ x(1-x)(1+x) \\ 1 \end{pmatrix}, \quad (11)$$

The flow contains vertical corelines for $\mathbf{v} \parallel \mathbf{a}$, including five vortices at $(0,0)$ and $(\pm 1, \pm 1)$, as well as four bifurcation lines at $(0, \pm 1)$ and $(\pm 1, 0)$. Further, there are four degenerated lines at $(\pm \frac{1}{\sqrt{3}}, \pm \frac{1}{\sqrt{3}})$ at which one eigenvalue vanishes to zero. We compare the accuracy of the different interpolants by measuring the residual error quantitatively and by comparing the resulting corelines visually. For this, we densely sample the analytic field onto a 128^3 grid and then progressively downsample to a resolution of 16^3 . With this, we can analyze how many sample points the different interpolants require in order to accurately capture the features of the underlying analytic field. Figure 3 compares the residual error for all interpolants. It shows that the derived acceleration interpolant outperforms the others significantly, suggesting that this interpolation method is better at representing the data of the underlying vector field. With the derived acceleration, a resolution of 64^3 is sufficient to achieve a very high degree of accuracy. However, this accuracy comes with a large drop in performance, which

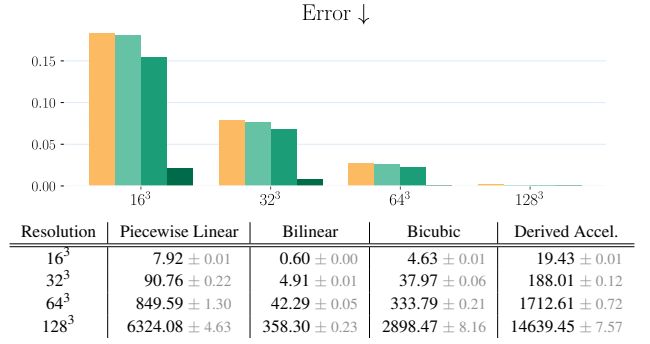


Figure 3: Top: Comparison of cross product residuals (lower is better) for: piecewise linear (●), bilinear clipping (●), bicubic clipping (●), and bicubic clipping with derived acceleration (●). Bottom: Comparison of the total runtime (serial) in milliseconds.




Figure 4: Comparison of interpolants on the NINE CL dataset downsampled to a 16^3 grid. From left to right: bilinear (●), bicubic (●) and derived acceleration (●) are shown. Ground truth lines are shown in white, and for vortex cores swirling streamlines are visualized. The interpolant with derived acceleration reproduces the expected lines better than the other two approaches.

is visible in the table below the plot. Derived acceleration is approximately $4.9\times$ slower than the bicubic interpolant and $38.0\times$ slower than the bilinear interpolant. It is also approximately $2.2\times$ slower than the piecewise linear interpolator. Fig. 4 shows the extracted corelines for resolution 16^3 alongside the ground truth.

5. Pseudocodes

The Bézier-based formulation includes tensor contractions, which we in part accelerate by precomputing weight tables that store the binomial coefficient fractions. For reproducibility of those algorithms, we list here the pseudocodes of three important steps. First, Alg. 1 lists the computation of control points for the cross product of two given vector fields in Bézier form, where the degree of both vector fields can be arbitrary.

Second, Alg. 2 describes how a $4 \times 4 \times 4$ block of values $\mathbf{p}_{i,j,k}$ is tricubically interpolated, such that the inner $2 \times 2 \times 2$ block is interpolated with C^1 continuity. For this, the 4×4 matrix \mathbf{M} is applied successively in each dimension, which converts from the values to the Bézier form of the cubic Hermite interpolant.

Lastly, Alg. 3 lists how the Bézier control points of the acceleration field are computed from the Bézier form of both the velocity field $\mathbf{v}_{i,j,k}$ and the Jacobian field $\mathbf{J}_{i,j,k}$. Again, the binomial coefficient weights are precomputed to avoid redundant computations.

Algorithm 1 Computation of Cross Product Control Points $\mathbf{c}_{i,j}$

```

1: Precompute weights:
2: for  $i' = 0$  to  $n_v$  do
3:   for  $j' = 0$  to  $m_v$  do
4:     for  $i'' = 0$  to  $n_w$  do
5:       for  $j'' = 0$  to  $m_w$  do
6:          $\omega_{i',j',i'',j''} \leftarrow \frac{\binom{n_v}{i'} \binom{m_v}{j'}}{\binom{n_c}{i}} \cdot \frac{\binom{m_w}{j''} \binom{m_w}{j''}}{\binom{m_c}{j}}, i \leftarrow i' + i'', j \leftarrow j' + j''$ 
7: Compute control points:
8: for  $i = 0$  to  $n_c$  do
9:   for  $j = 0$  to  $m_c$  do
10:     $\mathbf{c}_{i,j} \leftarrow \mathbf{0}$ 
11:    for all  $i', i''$  with  $i' + i'' = i$  do
12:      for all  $j', j''$  with  $j' + j'' = j$  do
13:         $\mathbf{c}_{i,j} \leftarrow \mathbf{c}_{i,j} + \omega_{i',j',i'',j''} \cdot (\mathbf{v}_{i',j'} \times \mathbf{w}_{i'',j''})$ 
14: return  $\{\mathbf{c}_{i,j}\}_{0 \leq i \leq n_c, 0 \leq j \leq m_c}$ 

```

Algorithm 2 Computation of Bézier Control Points $\mathbf{b}_{i,j,k}$

```

1: Contraction along x-axis:
2: for  $i = 0$  to 3 do
3:   for  $j = 0$  to 3 do
4:     for  $k = 0$  to 3 do
5:        $\mathbf{t}_{i,j,k}^x \leftarrow \sum_{a=0}^3 \mathbf{M}_{i,a} \cdot \mathbf{p}_{a,j,k}$ 
6: Contraction along y-axis:
7: for  $i = 0$  to 3 do
8:   for  $j = 0$  to 3 do
9:     for  $k = 0$  to 3 do
10:       $\mathbf{t}_{i,j,k}^{xy} \leftarrow \sum_{b=0}^3 \mathbf{M}_{j,b} \cdot \mathbf{t}_{i,b,k}^x$ 
11: Contraction along z-axis:
12: for  $i = 0$  to 3 do
13:   for  $j = 0$  to 3 do
14:     for  $k = 0$  to 3 do
15:        $\mathbf{b}_{i,j,k} \leftarrow \sum_{c=0}^3 \mathbf{M}_{k,c} \cdot \mathbf{t}_{i,j,c}^{xy}$ 
16: return  $\{\mathbf{b}_{i,j,k}\}_{0 \leq i,j,k \leq 3}$ 

```

Algorithm 3 Computation of Acceleration Control Points $\mathbf{a}_{i,j,k}$

```

1: Precompute weights:
2: for  $i' = 0$  to  $n_1$  do
3:   for  $i'' = 0$  to  $n_1$  do
4:      $\omega_{i',i'',i}^u \leftarrow \frac{\binom{n_1}{i'} \binom{n_1}{i''}}{\binom{2n_1}{i}}, i = i' + i''$ 
5: for  $j' = 0$  to  $n_2$  do
6:   for  $j'' = 0$  to  $n_2$  do
7:      $\omega_{j',j'',j}^v \leftarrow \frac{\binom{n_2}{j'} \binom{n_2}{j''}}{\binom{2n_2}{j}}, j = j' + j''$ 
8: for  $k' = 0$  to  $n_3$  do
9:   for  $k'' = 0$  to  $n_3$  do
10:     $\omega_{k',k'',k}^w \leftarrow \frac{\binom{n_3}{k'} \binom{n_3}{k''}}{\binom{2n_3}{k}}, k = k' + k''$ 
11: Compute acceleration control points:
12: for  $i = 0$  to  $2n_1$  do
13:   for  $j = 0$  to  $2n_2$  do
14:     for  $k = 0$  to  $2n_3$  do
15:        $\mathbf{a}_{i,j,k} \leftarrow \mathbf{0}$ 
16:       for all  $i', i''$  with  $i' + i'' = i$  do
17:         for all  $j', j''$  with  $j' + j'' = j$  do
18:           for all  $k', k''$  with  $k' + k'' = k$  do
19:              $w \leftarrow \omega_{i',i'',i}^u \cdot \omega_{j',j'',j}^v \cdot \omega_{k',k'',k}^w$ 
20:              $\mathbf{a}_{i,j,k} \leftarrow \mathbf{a}_{i,j,k} + w \cdot (\mathbf{J}_{i',j',k'}^u \cdot \mathbf{v}_{i'',j'',k''}^v)$ 
21: return  $\{\mathbf{a}_{i,j,k}\}_{0 \leq i \leq 2n_1, 0 \leq j \leq 2n_2, 0 \leq k \leq 2n_3}$ 

```

A C++ reference implementation is available at [DG26].

References

- [BRG21] BAEZA ROJO I., GÜNTHER T.: Coreline criteria for inertial particle motion. In *Topological Methods in Data Analysis and Visualization IV*. Springer, 2021, pp. 133–157. doi:10.1007/978-3-030-83500-2_8. 3
- [DG26] DASSLER N., GÜNTHER T.: Reference implementation for parallel vectors extraction using Bézier clipping. <https://github.com/fau-vc/pvBezier>, 2026. Accessed: 2026-03-04. 4
- [PR99] PEIKERT R., ROTH M.: The "Parallel Vectors" operator – a vector field visualization primitive. In *Proc. IEEE Visualization* (San Francisco, CA, USA, 1999), IEEE, pp. 263–270. doi:10.1109/VISUAL.1999.809896. 1, 3