



Distributed Surface Reconstruction

D. Marin¹ , P. Komon¹ , S. Ohrhallinger¹  and M. Wimmer¹ 

¹Institute of Visual Computing & Human-Centered Technology, TU Wien, Austria



(a) Point cloud input.



(b) Reconstructed surface, for $\delta = 1.35, l = \frac{d}{200}$.

Figure 1: Reconstructed surface of the eclepens dataset.

Abstract

Recent advancements in scanning technologies and their rise in availability have shifted the focus from reconstructing surfaces from point clouds of small areas to large, e.g., city-wide scenes, containing massive amounts of data. We adapt a surface reconstruction method to work in a distributed fashion on a high-performance cluster, reconstructing datasets with millions of vertices in seconds. We exploit the locality of the connectivity required by the reconstruction algorithm to efficiently divide-and-conquer the problem of creating triangulations from very large unstructured point clouds.

CCS Concepts

• *Computing methodologies* → *Point-based models*;

1. Introduction

Recent advances in scanning abilities, together with the increase in popularity of digital twins usage, created the tools to generate and, respectively, the need to process massive scans of real-life elements. Hence, we are now facing the question of how to quickly process amounts of data that easily exceed millions of points while still creating a faithful representation of the initial surfaces.

To answer this question, we distribute a surface reconstruction algorithm [POEM24] that filters the Delaunay triangulation by the intersection ratio of the Voronoi balls and removes long triangles, by taking advantage of this local connectivity condition. We split the input into tiles and add a parameter-based padding to each tile to create an overlap. This ensures that no surface data is missing in the final reconstruction while minimizing the amount of duplicate information and the data transfer between nodes.

2. Related Work

Numerous methods to parallelize the Delaunay triangulation [PMP14, CMYB19] have been developed recently, as it represents a first step in many surface reconstruction methods. The authors introduce various paddings to ensure that the node-computed triangulation agrees with the global one. Our method, even though it uses the Delaunay triangulation, does not require a global triangulation of the input points as it imposes a maximum edge length of triangles in the output, condition we exploit in our work.

Recently, the state-of-the-art surface reconstruction - Poisson Reconstruction [KBH06], has been adapted for out-of-core usage [KH23]. However, our method does not require any preprocessing of the point cloud or additional information such as normals.

3. Method

We adapt the recent curve and surface reconstruction algorithm Ballmerge [POEM24] to work distributedly. Its surface recon-

struction method has two variants: global and local, both starting from a Delaunay triangulation of the input. While the *Global Ballmerge* merges overlapping regions obtaining a manifold, watertight mesh, the *Local Ballmerge* filters triangles based on the following criteria: the intersection ratio between their respective Voronoi balls is less than a threshold δ , and the longest edge of the triangle is shorter than a predefined l . We will only concern ourselves with the *Local* version, as it is faster and easier to distribute. Moreover, point clouds with a magnitude of millions are usually scans of outdoor scenes for which, due to scan shadows or scan resolution, it is infeasible to enforce a manifold output.

3.1. Splitting

To be able to run the Delaunay triangulation and subsequently, the *Local Ballmerge* in parallel on each cluster node, we need to split the input data. We choose to split the input along a three-dimensional regular grid due to its simplicity and ease of parallelization. Surface reconstruction can be performed in each grid cell individually. Since triangles might be shared between cells, the reconstruction might miss important parts of the surface. To mitigate this, we introduce padding around each cell.

Since we filter triangles with edges longer than l , we pad the cells with l in the positive directions so no possibly valid triangles are missed in the triangulation. We also ensure that if points are in different cells, their Voronoi balls intersection ratio is larger than the predefined δ . To push the points sufficiently far away that their Voronoi balls are empty of samples, we additionally pad with $\frac{2l}{\sqrt{(4\delta - \delta^2)}}$ in all directions. We omit the proof due to length constraints but have verified it experimentally as well. Following the original algorithm, we use $l = \frac{d}{200}$, where d is the diagonal of the bounding box, and change their default $\delta = 1.85$ value to $\delta = 1.35$ in our experiments, as this value improved the resulted quality.

3.2. Distributed Version

The regularly subdivided tiles are assigned and distributed to nodes. Each node performs surface reconstruction on its assigned tiles by locally executing *Local Ballmerge*. As points obtained from real-world scans usually are not distributed uniformly, the number of points within tiles may vary greatly. To balance out the load, we assign tiles to nodes using longest-processing-time-first list scheduling, based on the number of points in the tiles, which is an efficient and good approximation of the optimal solution to the load-balancing problem. Merging is trivially done by taking the union of all local results. Because of the padding, the overall result is guaranteed to be the same as when running *Local Ballmerge* on the entire input. Expected run time complexity for points sampled on a surface is $\mathcal{O}(n/c + (n \log n)/p)$ in the best case of uniform input point distribution, worst is $\mathcal{O}(n/c + n \log n)$ as in the original, for n points, p reconstruction nodes and c cores for splitting.

4. Results

Splitting is done on the GPU using CUDA on a single node, while the reconstruction is distributed via MPI to up to 32 CPU nodes. We have evaluated our implementation on the [VSC3+](#) HPC cluster,

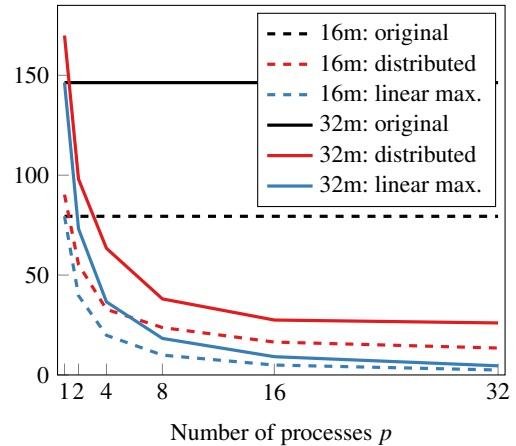


Figure 2: Run time comparison between original and distributed version using datasets with 16 and 32 million points on 1-32 nodes.

varying input sizes and numbers of nodes. We used truncated versions of an aerial photogrammetry scan to simulate increasing input size - Figure 1. Increasing the number of nodes with a fixed input size and subdivision shows timings close to the linear speedup (theoretical maximum) - Figure 2. By using 16 nodes and scaling up the input size, the run times of both original and distributed versions increase according to their expected behavior of $\mathcal{O}(n \log n)$. However, the distributed version exhibits much slower growth, peaking at a speedup factor of 5 for 32 million points and 32 nodes due to data duplication. While varying the number of nodes for a fixed number of points, finer subdivisions allow for better work distribution, resulting in potentially lower run times. However, choosing finer subdivisions results in significantly more duplicated points, eventually canceling out any speed gains from adding more nodes.

5. Conclusion and Future Work

Our method adapts a surface reconstruction algorithm to work distributedly, with minimal data duplication, achieving speedups close to linear for various configurations. We plan to investigate other data-splitting approaches, as well as the possibility of parallelizing *Local Ballmerge* on a node level. Moreover, we aim to evaluate our method in more scanner scenarios and compare it to recent work.

References

- [CMYB19] CARAFFA L., MEMARI P., YIRCI M., BRÉDIF M.: Tile & merge: Distributed delaunay triangulations for cloud computing. In *2019 IEEE International Conference on Big Data* (2019), pp. 1613–1618. 1
- [KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proceedings of SGP* (2006), SGP '06, p. 61–70. 1
- [KH23] KAZHDAN M., HOPPE H.: Distributed poisson surface reconstruction. In *Computer Graphics Forum* (2023), vol. 42, p. e14925. 1
- [PMP14] PETERKA T., MOROZOV D., PHILLIPS C.: High-performance computation of distributed-memory parallel 3d voronoi and delaunay tessellation. In *SC '14* (2014), pp. 997–1007. 1
- [POEM24] PARAKKAT A. D., OHRHALLINGER S., EISEMANN E., MEMARI P.: Ballmerge: High-quality fast surface reconstruction via voronoi balls. In *Computer Graphics Forum* (2024). 1