

fine cells containing more edges than a threshold, what created a nested grid. This refinement could be recursively repeated in a process similar to a quad-tree creation but, according to our preliminary experiments, for the overlay problem a good performance is obtained by using only one or two levels.

Since uniform grids performed well in EPUGO, we intend to extend this strategy to 3D. The idea is to create a 3D grid and insert in the grid cells the triangles intersecting them.

5. 2D INTERSECTION

We have developed the algorithm EPUGO for exactly overlaying maps. All coordinates are represented with GMP rationals. Special cases are treated handled using SoS.

Since the map is represented as a unstructured set of edges that contain information about the neighbor faces, the overlay operation is performed by computing the intersection between the edges and classifying them. Algorithm 1 summarizes the overlay process.

Algorithm 1 Computes the overlay of two maps A and B

- 1: Create the uniform grid
 - 2: Compute the intersection between all edges of the maps
 - 3: Locate all vertices of map A in map B and vice-versa
 - 4: Classify the edges and create the resulting map
-

Initially, the uniform grid is created. This process is easily parallelized over the edges. If the algorithm is configured to use more than one level in the grid, the cells containing too many edges are refined creating a nested grid in them.

After creating the index, the intersections between edges in each grid cell are computed. For each cell c , the intersection between pairs of edges from both maps in c are computed. Again, this step is highly parallelizable over the grid cells.

Then, we compute the face from one map in which each vertex from the other map is. This process is performed by creating, for each vertex v , a semi-infinite vertical ray starting in v and determining what is the lowest edge e directly above v . The face containing v is one of the two faces adjacent to e . If there is no edge above v , v is in the outside face. This process also uses a uniform grid to index the edges that need to be tested for intersection with the vertical rays. Furthermore, this step is highly parallelizable over the vertices being located.

Finally, in the next step the output edges are computed. If an input edge e from map A does not intersect any edge from the map B , it will be an output edge if and only if it is inside a face f (that is not the exterior face) of B (this can be determined by checking in what face of B one of its two vertices is) and its left and right faces will be the intersection between f and the faces from B that are around e .

If e intersects n edges from the other map, e is divided in $n+1$ sub-edges and each sub-edge e_i is classified using a strategy similar to the one used when an edge does not intersect other edges: e_i will be an output edge iff it is inside a face of the other map and the new faces to its left and right sides will be determined based on the face of the other map where e_i is and on the faces of e 's map that are around e .

Since there is no data dependency between the edges, they can be easily processed in parallel.

It is worth mentioning that preliminary experiments have showed that simply using rationals in an implementation created for floating-point numbers isn't fast. For example, in our implementation we avoid recreating temporary variables in expressions since the creation of rational numbers usually lead to memory allocations on the heap, what is significantly slow when these expressions are evaluated in parallel. Thus, EPUGO needed to be carefully implemented to be efficient.

The current parallel implementation is very efficient and our experiments showed that its performance is comparable to the performance of the overlay module present in GRASS GIS. For

example, in our largest experiment using maps with 20 and 30 million edges, EPUGO was 50% faster than GRASS GIS that, even though is sequential, does not use exact arithmetic.

6. 3D INTERSECTION

We intend to extend the ideas used in EPUGO to compute the intersection between two triangulations. The initial steps of the new method, named 3D-EPUGO, were already implemented and will be described in this section.

Initially, a 3D uniform grid is created and the triangles are inserted into the grid cells they intersect. For performance purposes, the axis-aligned bounding boxes ($AABBs$) of the triangles are computed in parallel and, then, the triangles are inserted in the grid cells intersecting their $AABBs$.

We intend to use the uniform grid as a filtering structure to accelerate the computation of the intersection, that is, the presence of a triangle in a grid cell c that does not actually intersect c does not affect the correctness of the algorithm. Thus, the choice of inserting the triangles into the grid basing on their $AABBs$ simplifies the creation of the uniform grid while it adds an overhead for the steps of the algorithm that relies in the grid to index the data. As a future work we intend to study these overheads in order to choose the grid creation strategy that leads to the best performance.

The most critical step of the overlay is computing the intersection between the triangles. This step was implemented using a strategy similar to the one used in EPUGO: for each grid cell, the intersections between pairs of triangles from the two triangulations are computed. For efficiency, the pairs of triangles are intersected using the algorithm presented by Möller [6], that uses several techniques to avoid unnecessary computations.

The next steps (computing and classifying the output triangles basing on the intersections computed in the previous step) are still under development and, thus, they will be implemented as future work. Furthermore, the current version of this algorithm does not handle the special cases and, thus, using *SoS* to handle these cases is also object of future work.

6.1 Conclusions

This paper presented a proposal for the development of 3D-EPUGO, an efficient algorithm for exactly computing the overlay of triangulations. 3D-EPUGO will use arbitrary-precision rational numbers to represent and process the spatial data.

As a proof of concept, we have developed EPUGO, a 2D version of 3D-EPUGO. Our preliminary experiments have showed that, due to the techniques we used, EPUGO was able to efficiently compute overlays. Thus, we expect that 3D-EPUGO will be also efficient.

This research was partially supported by NSF grant IIS-1117277 and by CAPES (Ciencia sem Fronteiras).

7. REFERENCES

- [1] F. Feito and C. Ogayar et al. Fast and accurate evaluation of regularized boolean operations on triangulated solids. *Computer-Aided Design*, 2013.
- [2] W. R. Franklin and N. Chandrasekhar et al. Efficiency of uniform grids for intersection detection on serial and parallel machines. In *Proc. Comp. Graphics Int.* 1988.
- [3] J. D. Hobby. Practical segment intersection with finite precision output. *Comput. Geom.*, 13(4):199–214, 1999.
- [4] C. M. Hoffman. The problems of accuracy and robustness in geometric computation. *Computer*, 22(3):31–40, 1989.
- [5] C. Li. *Exact geometric computation: theory and applications*. PhD thesis, Department of Computer Science, Courant Institute - NYU, 2001.
- [6] T. Möller. A fast triangle-triangle intersection test. *Journal of graphics tools*, 2(2):25–30, 1997.
- [7] J. Yongbin and W. Liguang et al. Boolean operations on polygonal meshes using obb trees. In *ESIAT*, 2009.