

Lesson 11  
Handling massive terrains

GE01015.2023

Hugo Ledoux

# Raster pyramids == simple & easy & used

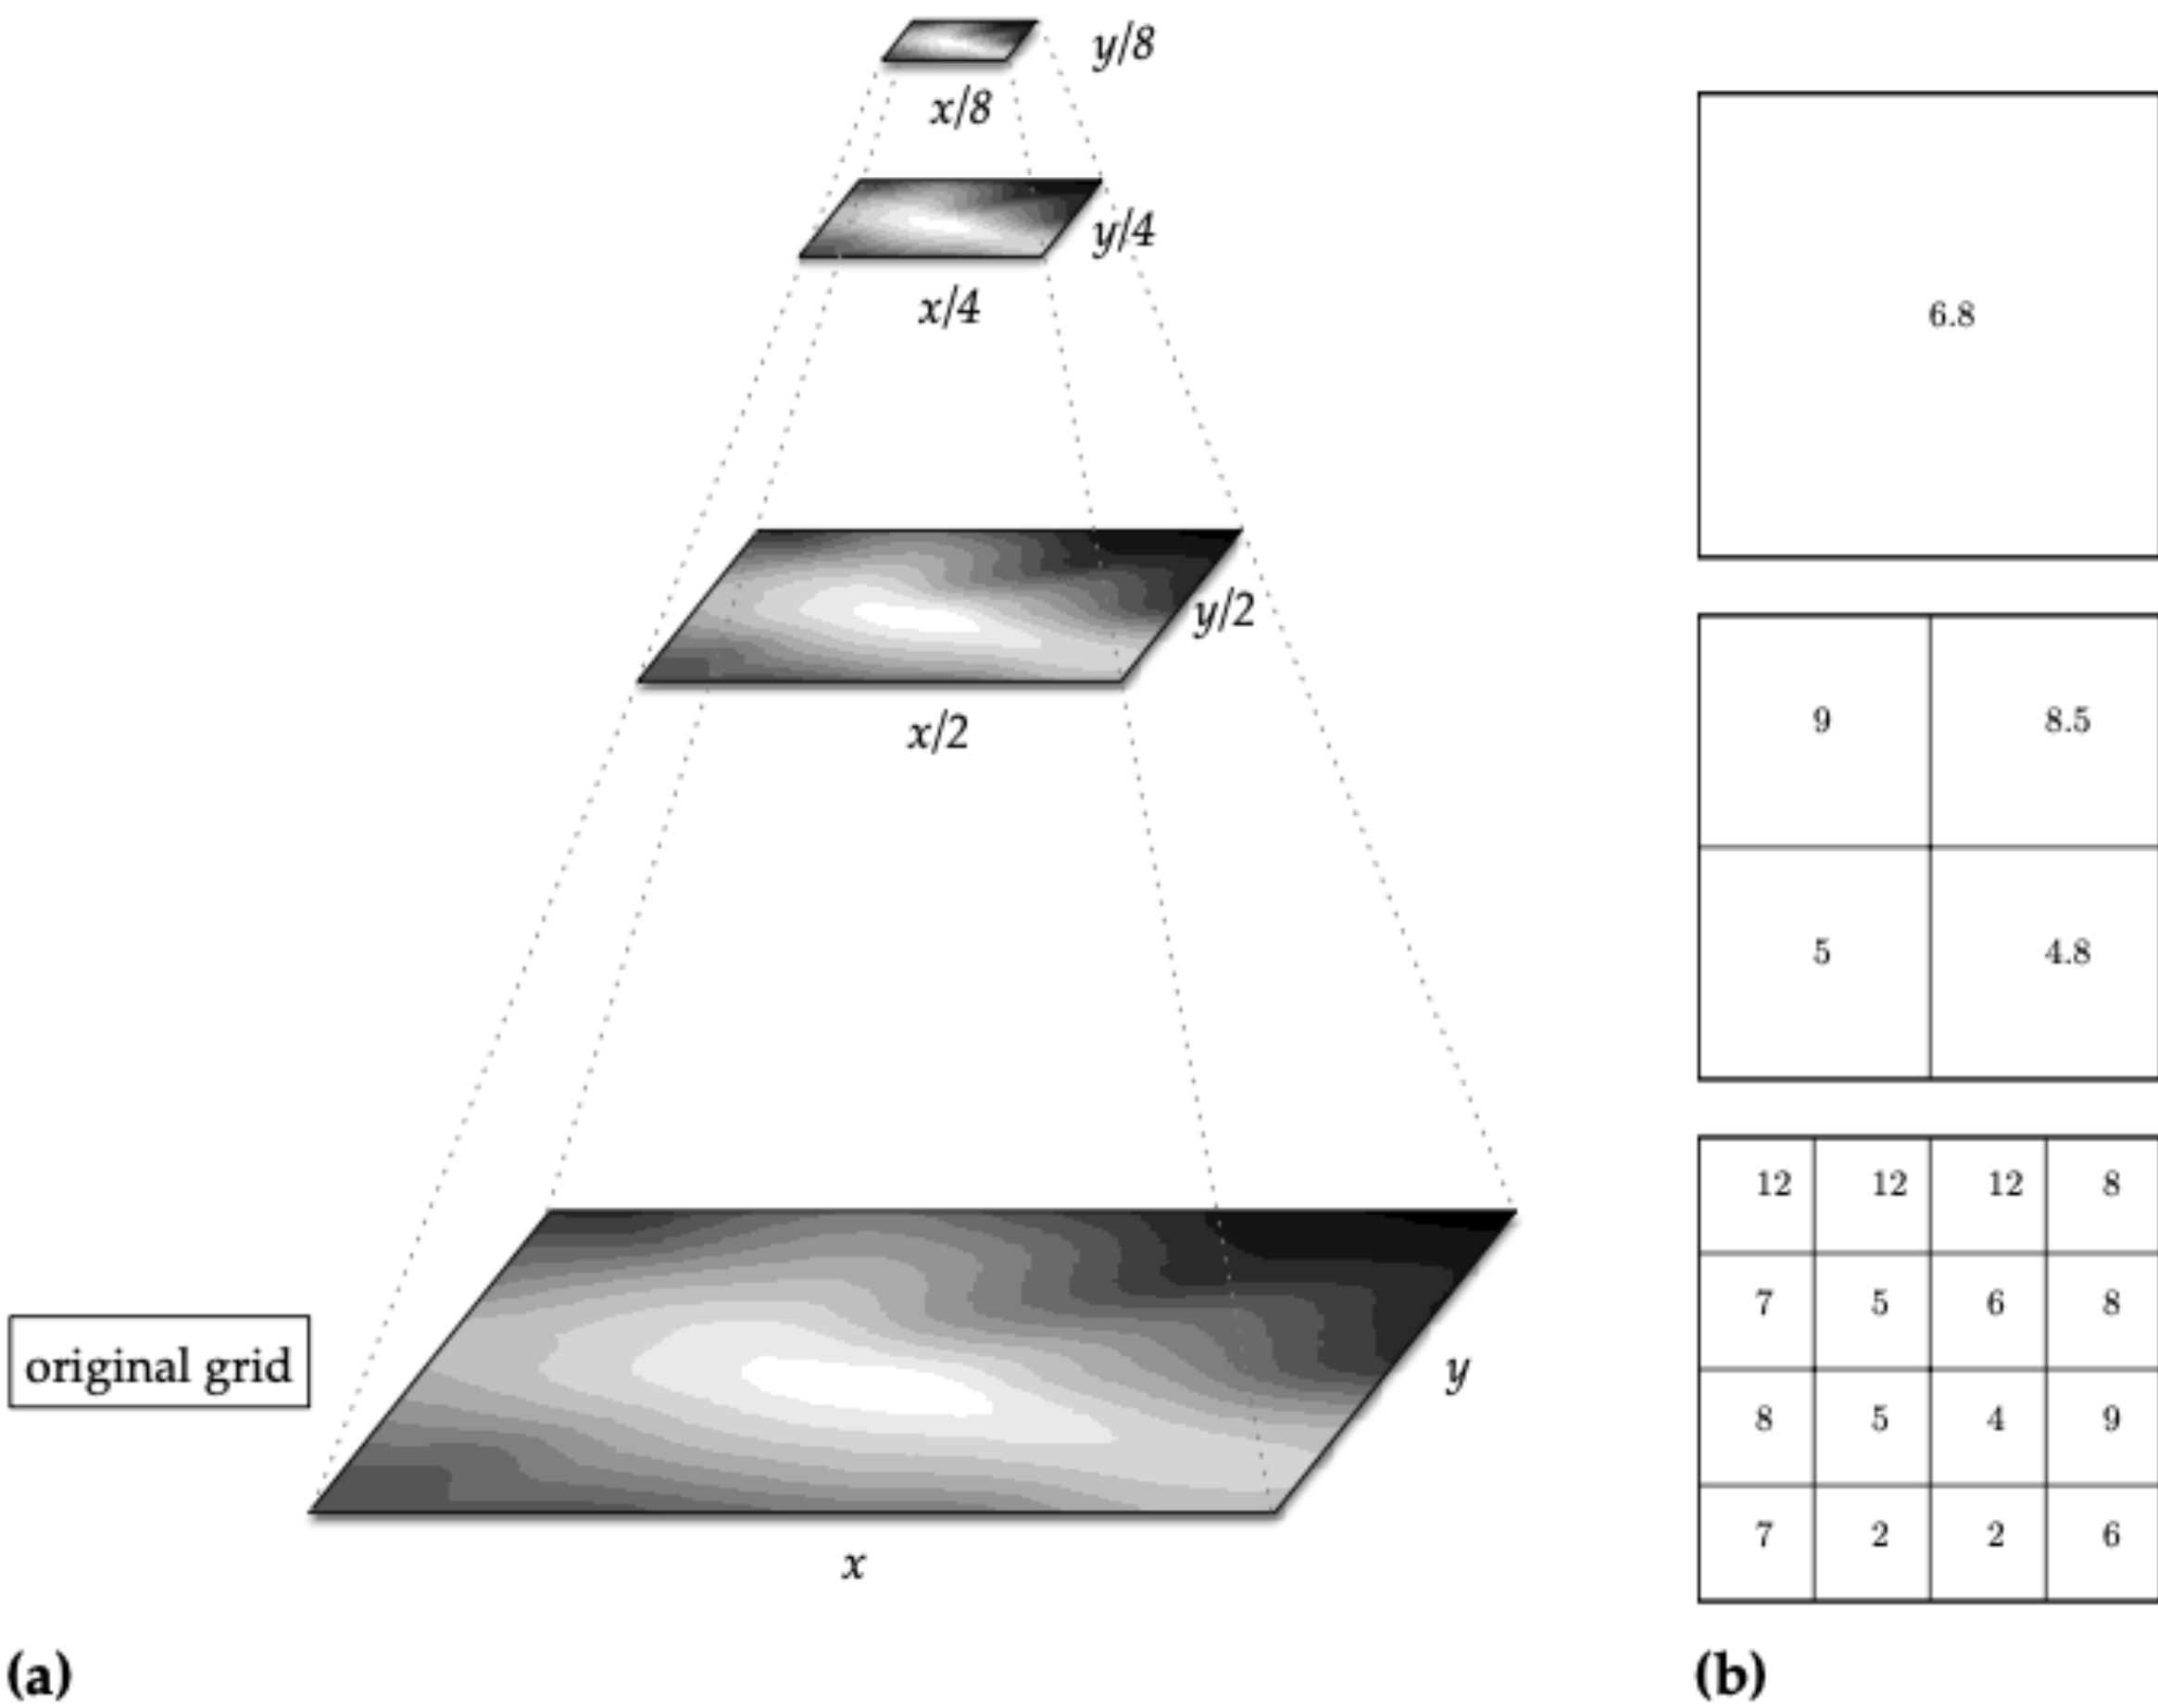
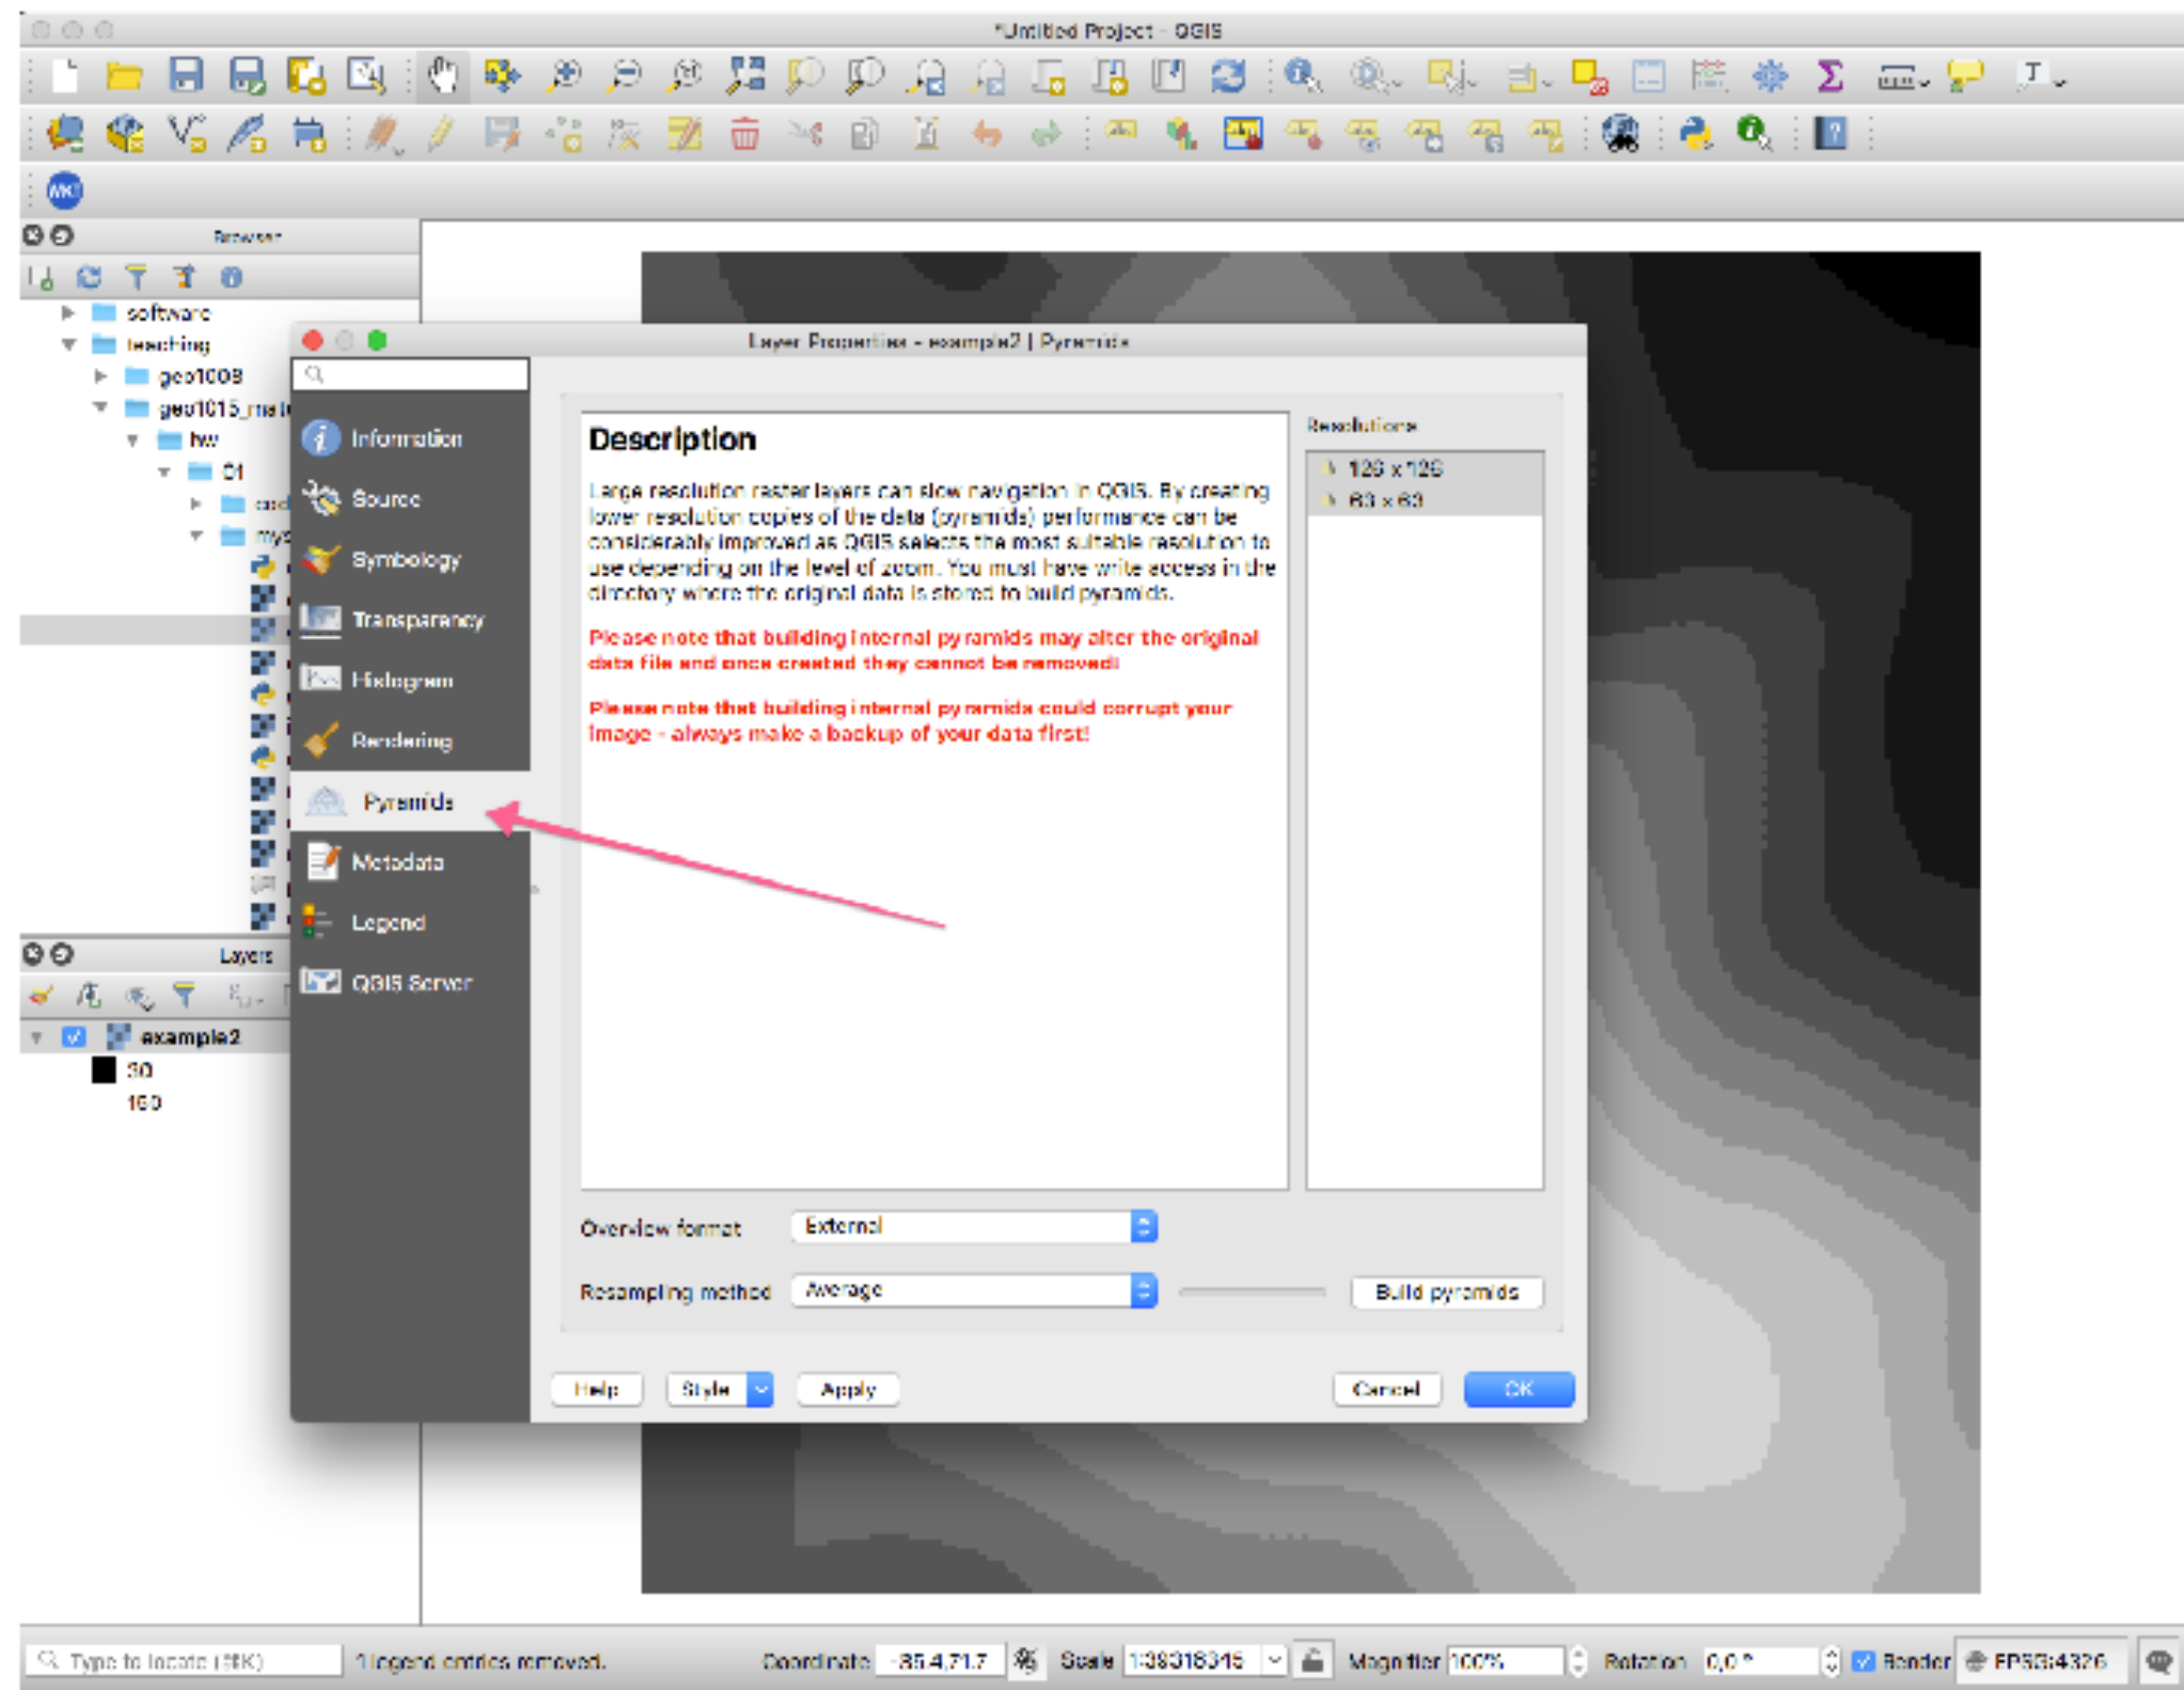


Figure 12.1: (a) The pyramid for a given raster file. (b) One 4 × 4 raster downsampled twice with average-method.

# Raster pyramids == simple & easy & used



## ⚙️ How does it work in practice?



For certain GIS formats, eg GeoTIFF, the lower-resolutions rasters can be stored directly in the same file as the original raster, and this is standardised.

For other formats, if the GDAL library is used (the *de facto* open-source library for GIS images and grids), the pyramids can be stored in an auxiliary file with the extension `.ovr`, which is actually a TIFF format.

The GDAL utility `gdaladdo` [↗](#) can create automatically the pyramids for a few formats, and the downsampling method can be chosen. In QGIS, one can use `gdaladdo`, or there is also a built-in mechanism, as can be seen in Figure 12.2

**Figure 12.2:** QGIS has the option to create the pyramids automatically.

# Spatially indexing points (aka find nearest neighbours)



- Question: startinpy has a `closest_point()` function and yet it builds the DT: how does that work?



**startin** 🔄

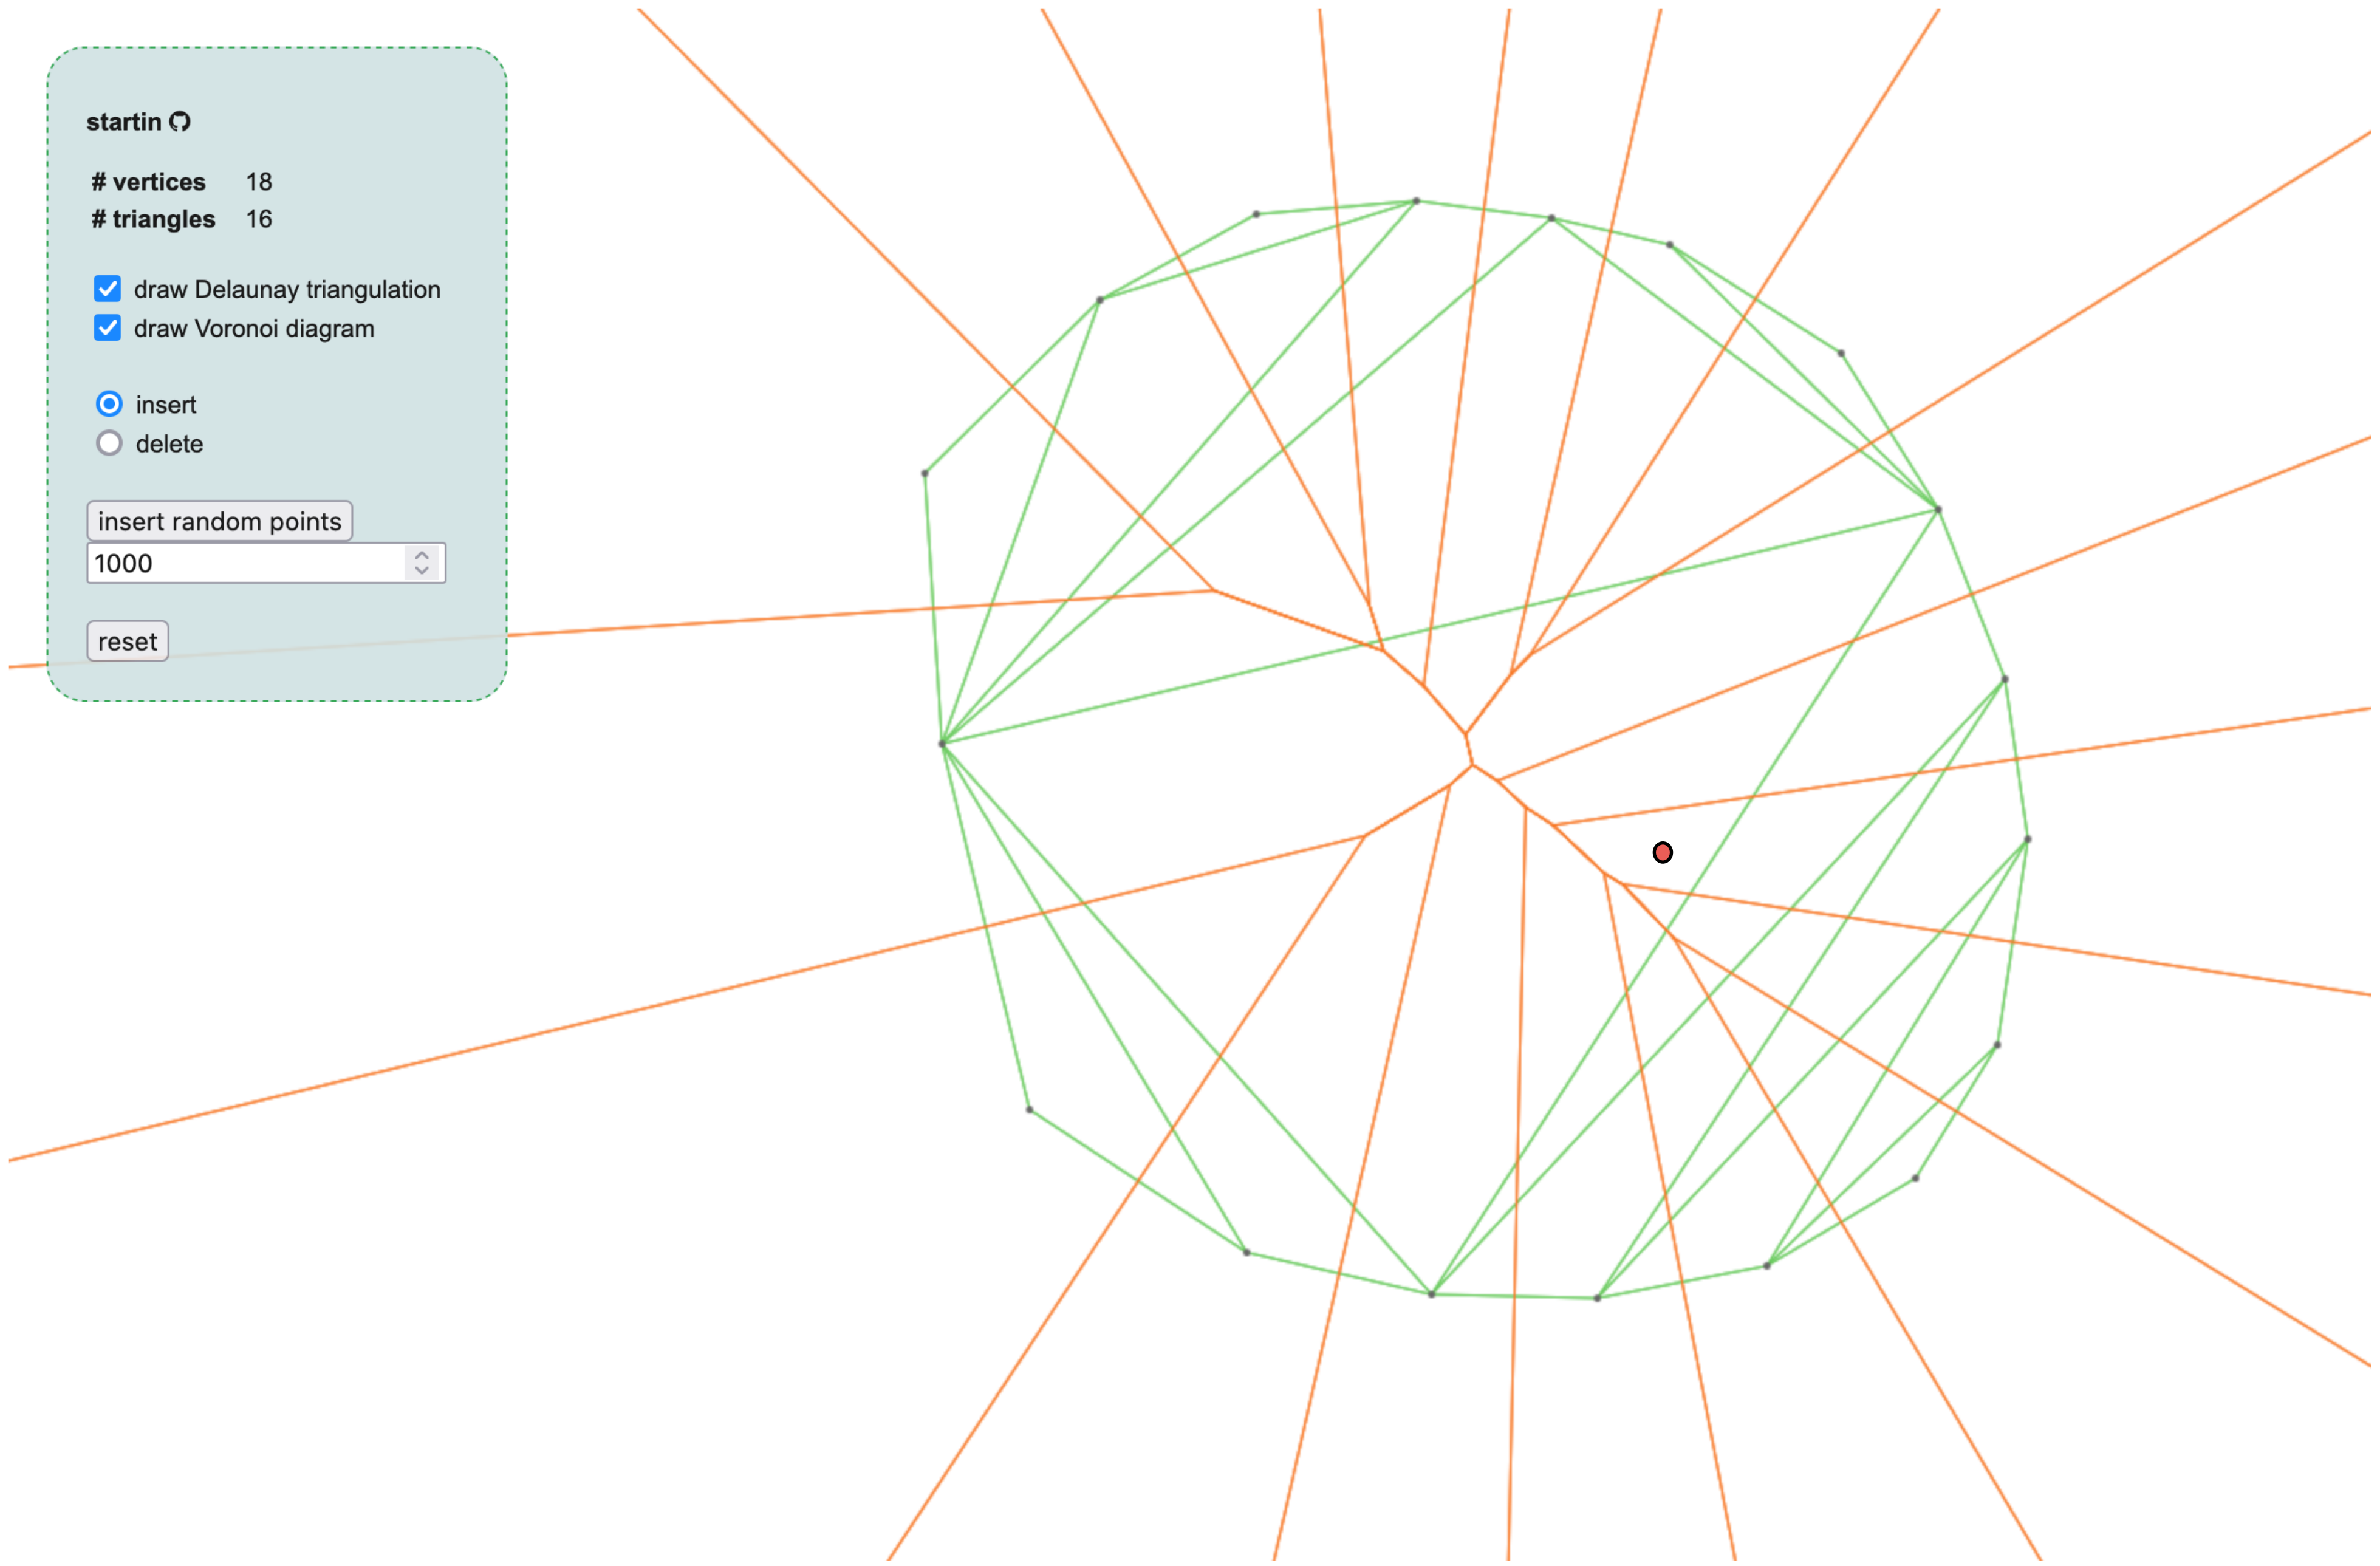
# vertices 18  
# triangles 16

draw Delaunay triangulation  
 draw Voronoi diagram

insert  
 delete

insert random points  
1000

reset





# Triangulation of billions of points? Why should you care?

- (raw points not always the best: too many, noise, doesn't fit in memory, etc)
- We want to extract iso-contours
- We want to create a grid of the area
- Noise project of RIVM: one and only one “base terrain” of NL (to replicate all computations)
- Good simplification of an area necessitate creating the TIN in the first place 🙄



# Partitioning the area into sub-tiles and triangulate each? Easy in theory...





# What I describe today is basically this paper

## Streaming Computation of Delaunay Triangulations

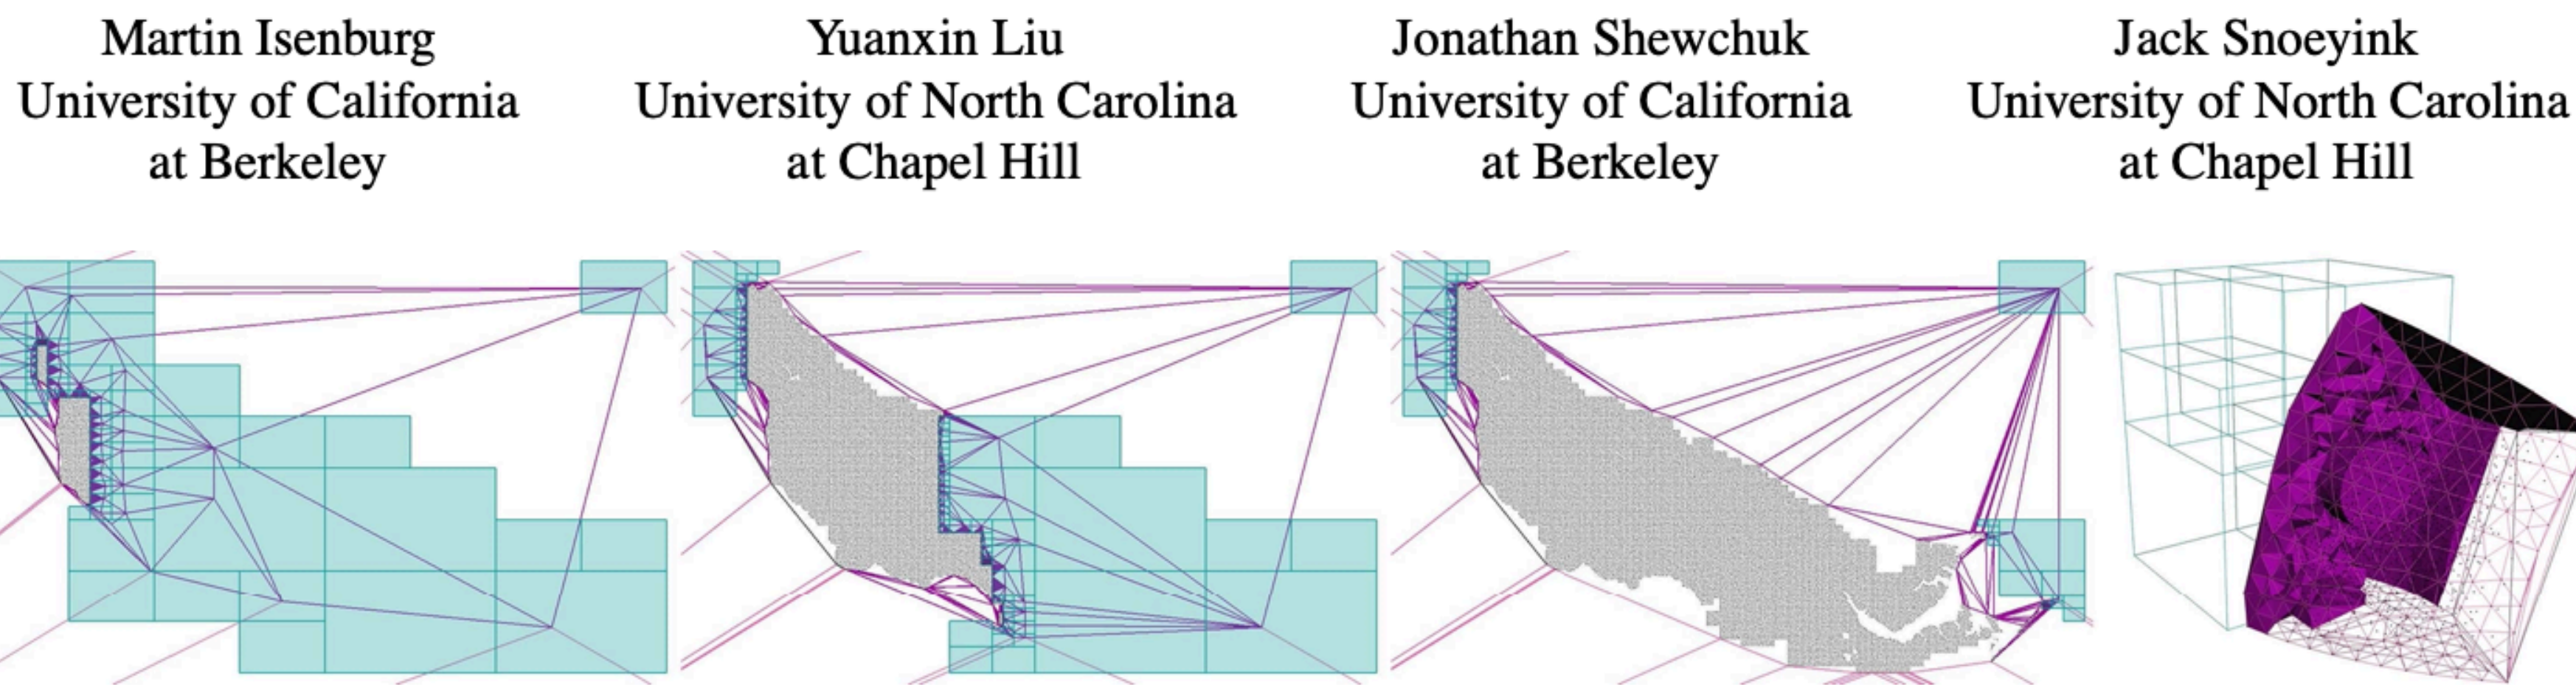


Figure 1: Streaming computation of Delaunay triangulations in 2D (Neuse River) and 3D. Blue quadrants or octants are unfinalized space where future points will arrive. Purple triangles and tetrahedra are in memory. Black points and their triangles and tetrahedra have already been written to disk or piped to the next application.

### Abstract

We show how to greatly accelerate algorithms that compute Delaunay triangulations of huge, well-distributed point sets in 2D and 3D by exploiting the natural spatial coherence in a stream of points. We achieve large performance gains by introducing *spatial finalization* into point streams: we partition space into regions, and augment a stream of input points with finalization tags that indicate when a point is the last in its region. By extending an incremental algorithm for Delaunay triangulation to use finalization tags and produce streaming mesh output, we compute a billion-triangle terrain representation for the Neuse River system from 11.2 GB of LIDAR

Delaunay triangulator, by Agarwal, Arge, and Yi [2005]; see Section 6. We also construct a nine-billion-triangle, 152 GB triangulation in under seven hours using 166 MB of main memory.

A *streaming* computation makes a small number of sequential passes over a data file (ideally, one pass), and processes the data using a memory buffer whose size is a fraction of the stream length. We have implemented two- and three-dimensional triangulators that read streams of points as input, and produce Delaunay triangulations in streaming mesh formats. The memory footprint of the 2D triangulator is typically less than 0.5% of the output mesh size (sometimes much less). The memory footprint of the 3D triangulation



# Unix pipes for streaming

136

12 Handling and processing massive terrains


## ⚙ Streaming is realised with Unix pipes

The key to implementing streaming of geometries is to use Unix pipes (also called *pipelines*).

Pipelines were designed by Douglas McIlroy at Bell Labs during the development of Unix, and they allow to chain several processes together. The output of a process becomes the input of the next one, and so on (the data flowing through the process is the *stream*). Given 2 processes, the 2nd one can usually start before the 1st one has finished processing all the data.

In Unix, the pipe operator is the vertical line “|”, and several commands can be chained with it: “cmd1 | cmd2 | cmd3”. A simple example would be “ls -l | grep json | wc -l” which would:

1. list all the files in the current directory (one file name per line);
2. send this to the operator *grep* which would discard all lines not having the keyword “json”;
3. send this to the operator “wc -l” which counts the number of line.







WIKIPEDIA  
The Free Encyclopedia

- Main page
- Contents
- Current events
- Random article
- About Wikipedia
- Contact us
- Donate

Contribute

- Help
- Learn to edit
- Community portal
- Recent changes
- Upload file

Tools

- What links here
- Related changes
- Special pages
- Permanent link
- Page information
- Cite this page
- Wikidata item

Print/export

- Download as PDF
- Printable version

Languages

- Deutsch
- Français
- Italiano
- 日本語
- Português
- Русский

Not logged in [Talk](#) [Contributions](#) [Create account](#) [Log in](#)

Article

[Talk](#)

Read

[Edit](#)

[View history](#)

Search Wikipedia

# Wavefront .obj file

From Wikipedia, the free encyclopedia

*For other uses, see [Obj \(disambiguation\)](#).*

**OBJ** (or **.OBJ**) is a geometry definition file format first developed by [Wavefront Technologies](#) for its [Advanced Visualizer](#) animation package. The file format is open and has been adopted by other 3D graphics application vendors.

The OBJ file format is a simple data-format that represents 3D geometry alone — namely, the position of each [vertex](#), the [UV position](#) of each texture coordinate vertex, [vertex normals](#), and the faces that make each polygon defined as a list of vertices, and texture vertices. Vertices are stored in a counter-clockwise order by default, making explicit declaration of face normals unnecessary. OBJ coordinates have no units, but OBJ files can contain scale information in a human readable comment line.

## OBJ geometry format


<b>Filename extension</b>	.obj
<b>Internet media type</b>	model/obj <sup>[1]</sup>
<b>Developed by</b>	<a href="#">Wavefront Technologies</a>
<b>Type of format</b>	3D model format

### Contents [\[hide\]](#)



- File format
  - Geometric vertex
  - Parameter space vertices
  - Face elements
    - Vertex indices
    - Vertex texture coordinate indices
    - Vertex normal indices
    - Vertex normal indices without texture coordinate indices
    - Line elements
  - Other geometry formats
  - Reference materials
  - Relative and absolute indices
- Material template library
  - Basic materials
  - Texture maps
  - Texture options
  - Vendor specific alterations
    - Physically-based Rendering
- See also



If all in memory it's "easy"




# Insertion deletes all the triangles whose circumcircle contains $p$






# Streaming DT architecture

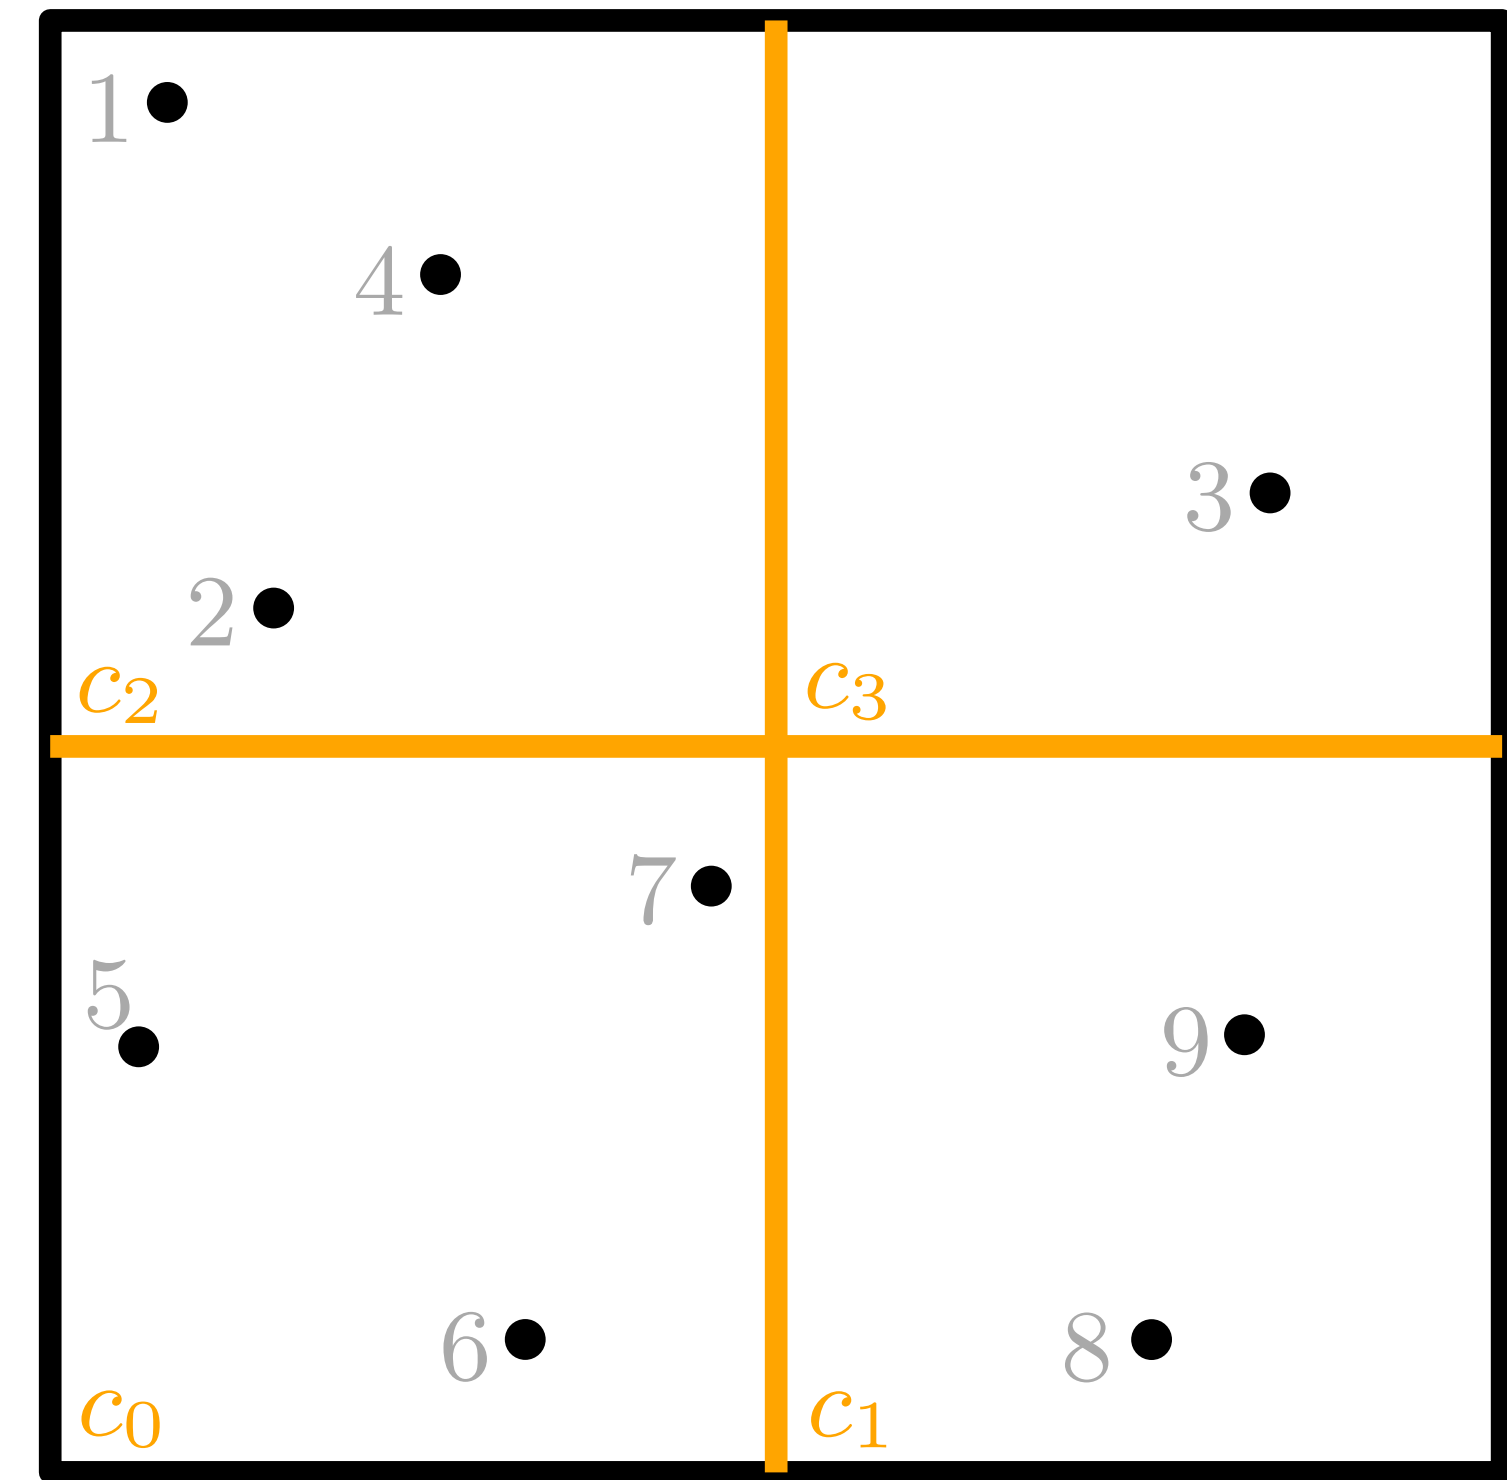


# Finaliser based on quadtree



1. x y z
2. x y z
3. x y z
4. x y z
5. x y z
6. x y z
7. x y z
8. x y z
9. x y z


9 points and a "normal" stream



1. x y z
2. x y z
3. x y z
- finalise  $c_3$
4. x y z
- finalise  $c_2$
5. x y z
6. x y z
7. x y z
- finalise  $c_0$
8. x y z
9. x y z
- finalise  $c_1$

9 points and a stream with finalisation tags





blue cell = not finalised cell

Yellow tr = in memory

White tr = written to stream

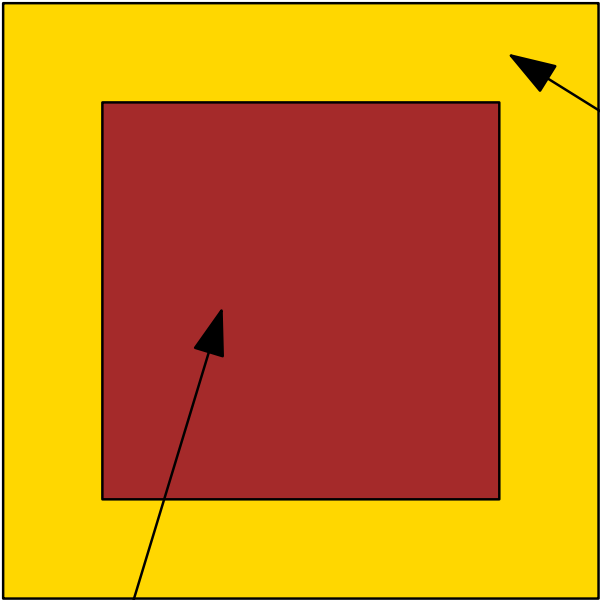

“a correlation between the proximity in space of geometric entities and the proximity of their representations in [the file]”

# Spatial coherence

AHN3 - 37EN1


order in file

- 0
- 120M
- 240M
- 360M
- 480M




last point position

first point position


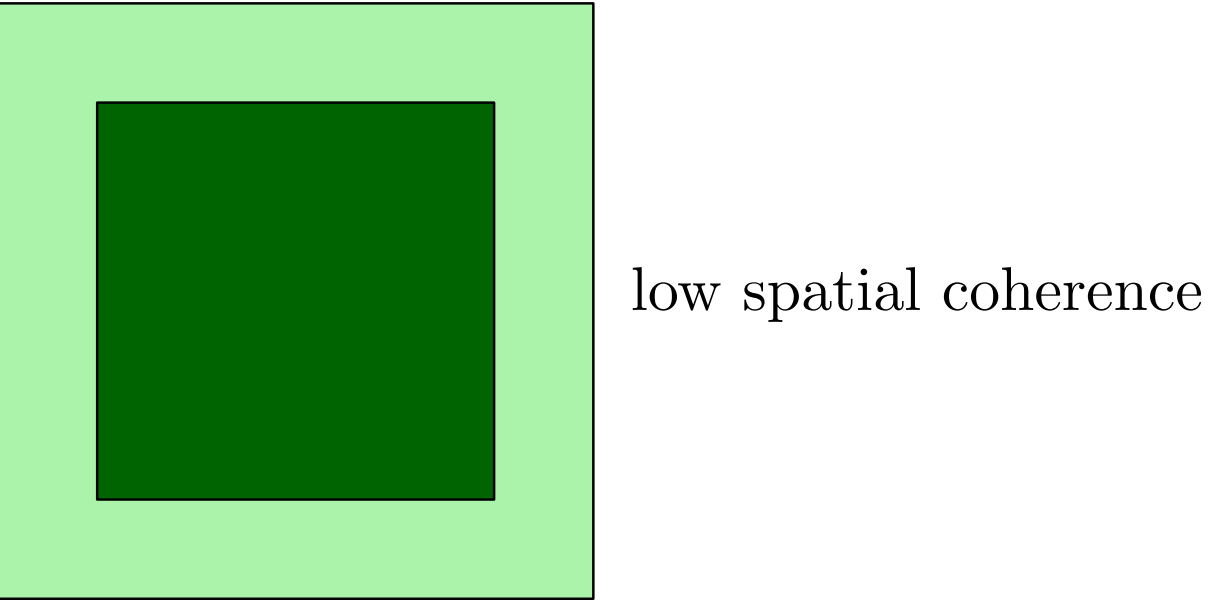
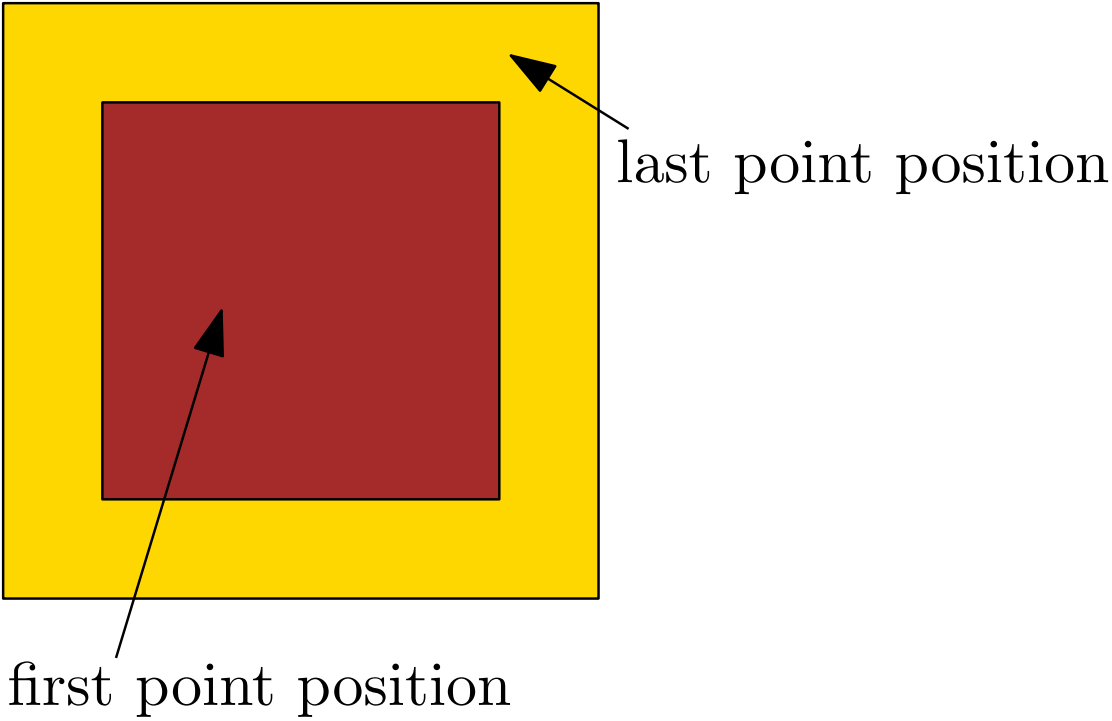




low spatial coherence




high spatial coherence

# Spatial coherence






# Spatial coherence



(a) 6 million point Baisman Run dataset in Broadmoor, Maryland



(b) 500 million point Neuse river basin dataset in North Carolina

Figure 2.4: Spatial coherence as inherent property of real-world datasets. Adapted from [Isenburg et al. \[2006b, p. 3\]](#).



(geo1015) ~/projects/sst git:(develop) (5m 1.94s)



```
./target/release/sstfin -vv /Users/hugo/data/ahn3/c_37en1_big_delft.laz 20 | ./target/release/sstdt2  
-vv > /dev/null
```

```
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.4  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell 43--74 finalised (3775 vertices)  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.4  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell 43--75 finalised (4637 vertices)  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0, 3, 0, 3] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0, 3, 0] finalised  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.4  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell 43--76 finalised (5503 vertices)  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell 43--77 finalised (10477 vertices)  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0, 3, 1, 2] finalised  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z INFO sstfin Third pass ❏  
2023-01-10T16:24:16Z INFO sstfin ✔  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell 43--78 finalised (7309 vertices)  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.3  
2023-01-10T16:24:16Z WARN sstdt2::triangulator walk.4  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell 43--79 finalised (1294 vertices)  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0, 3, 1, 3] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0, 3, 1] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0, 3] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2, 0] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1, 2] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[1] finalised  
2023-01-10T16:24:16Z INFO sstdt2::triangulator Cell qtc[] finalised  
2023-01-10T16:24:17Z INFO sstdt2::triangulator Finalise the quadtree root cell  
2023-01-10T16:24:17Z INFO sstdt2 dt.number_of_vertices() = 0  
2023-01-10T16:24:17Z INFO sstdt2 max # points in DT during process: 242733  
2023-01-10T16:24:17Z INFO sstdt2 max # triangles in DT during process: 300567  
2023-01-10T16:24:17Z INFO sstdt2 ✔
```



# Maarten de Jong MSc Geomatics thesis on the topic

MSc thesis in Geomatics

## Simplification of Massive TINs with the Streaming Geometries Paradigm

Maarten de Jong  
2021





# Some questions

- How to create a 50cm grid with TIN linear interpolation for 1B points?
- If we want to use Laplace interpolation, is it possible?