AHN 1-5 to Potree for online 3D viewing

Martijn Meijers

GIS-technology, Digital Technologies, Architectural Engineering + Technology,
Faculty of Architecture and the Built Environment, Delft University of Technology

Final Consortium meeting nD-PointClouds
Tuesday, November 12, 2024, 13:40–14:00 CET

Delft University of Technology
Faculty of Architecture and the Built Environment, Room B (08.BG.Oost.370)

Downloading

Getting the points...

- Overview of sheets per data set?
- Which data where?
  Scattered overview (esri, pdok, geotiles, geoforum.nl, github) - https://ahn.nl/dataroom/

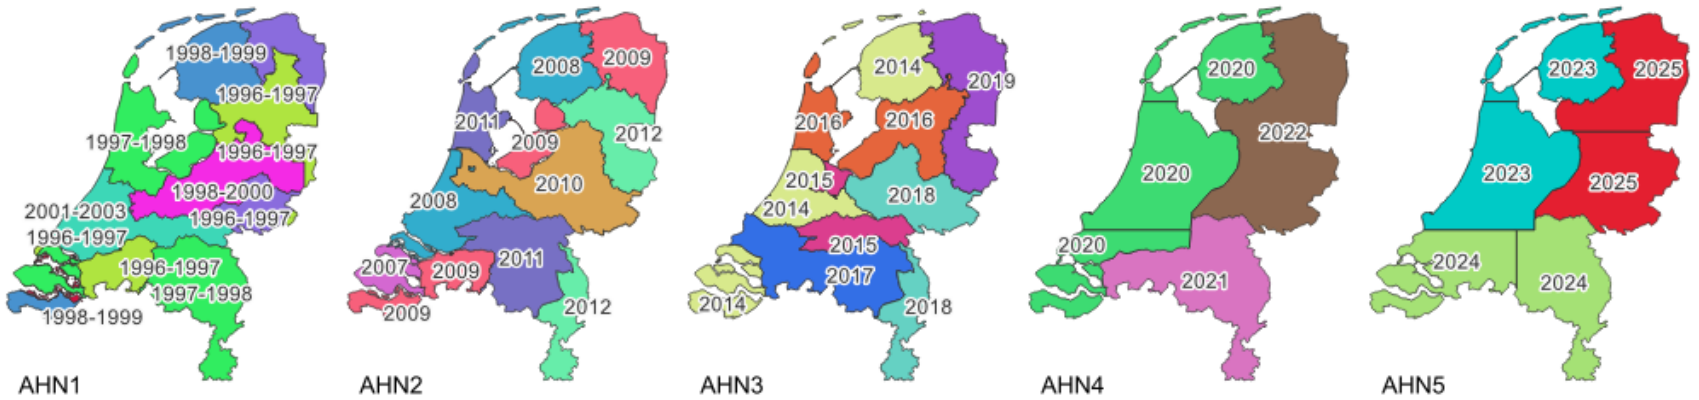

Corrupted downloads

- Download service? Checksum!

https://geoforum.nl/t/checksum-for-the-ahn-laz-downloads/4002

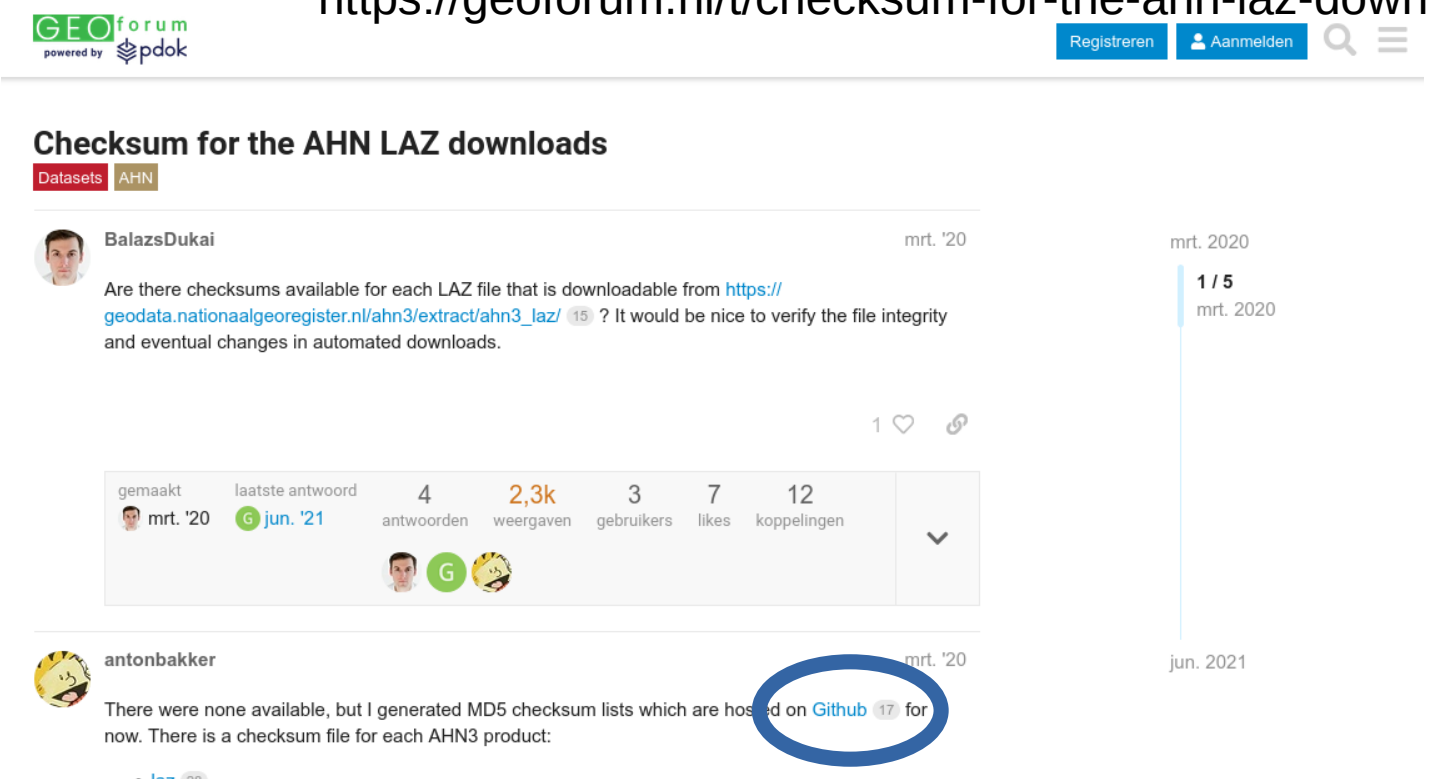

Official AHN1 ?

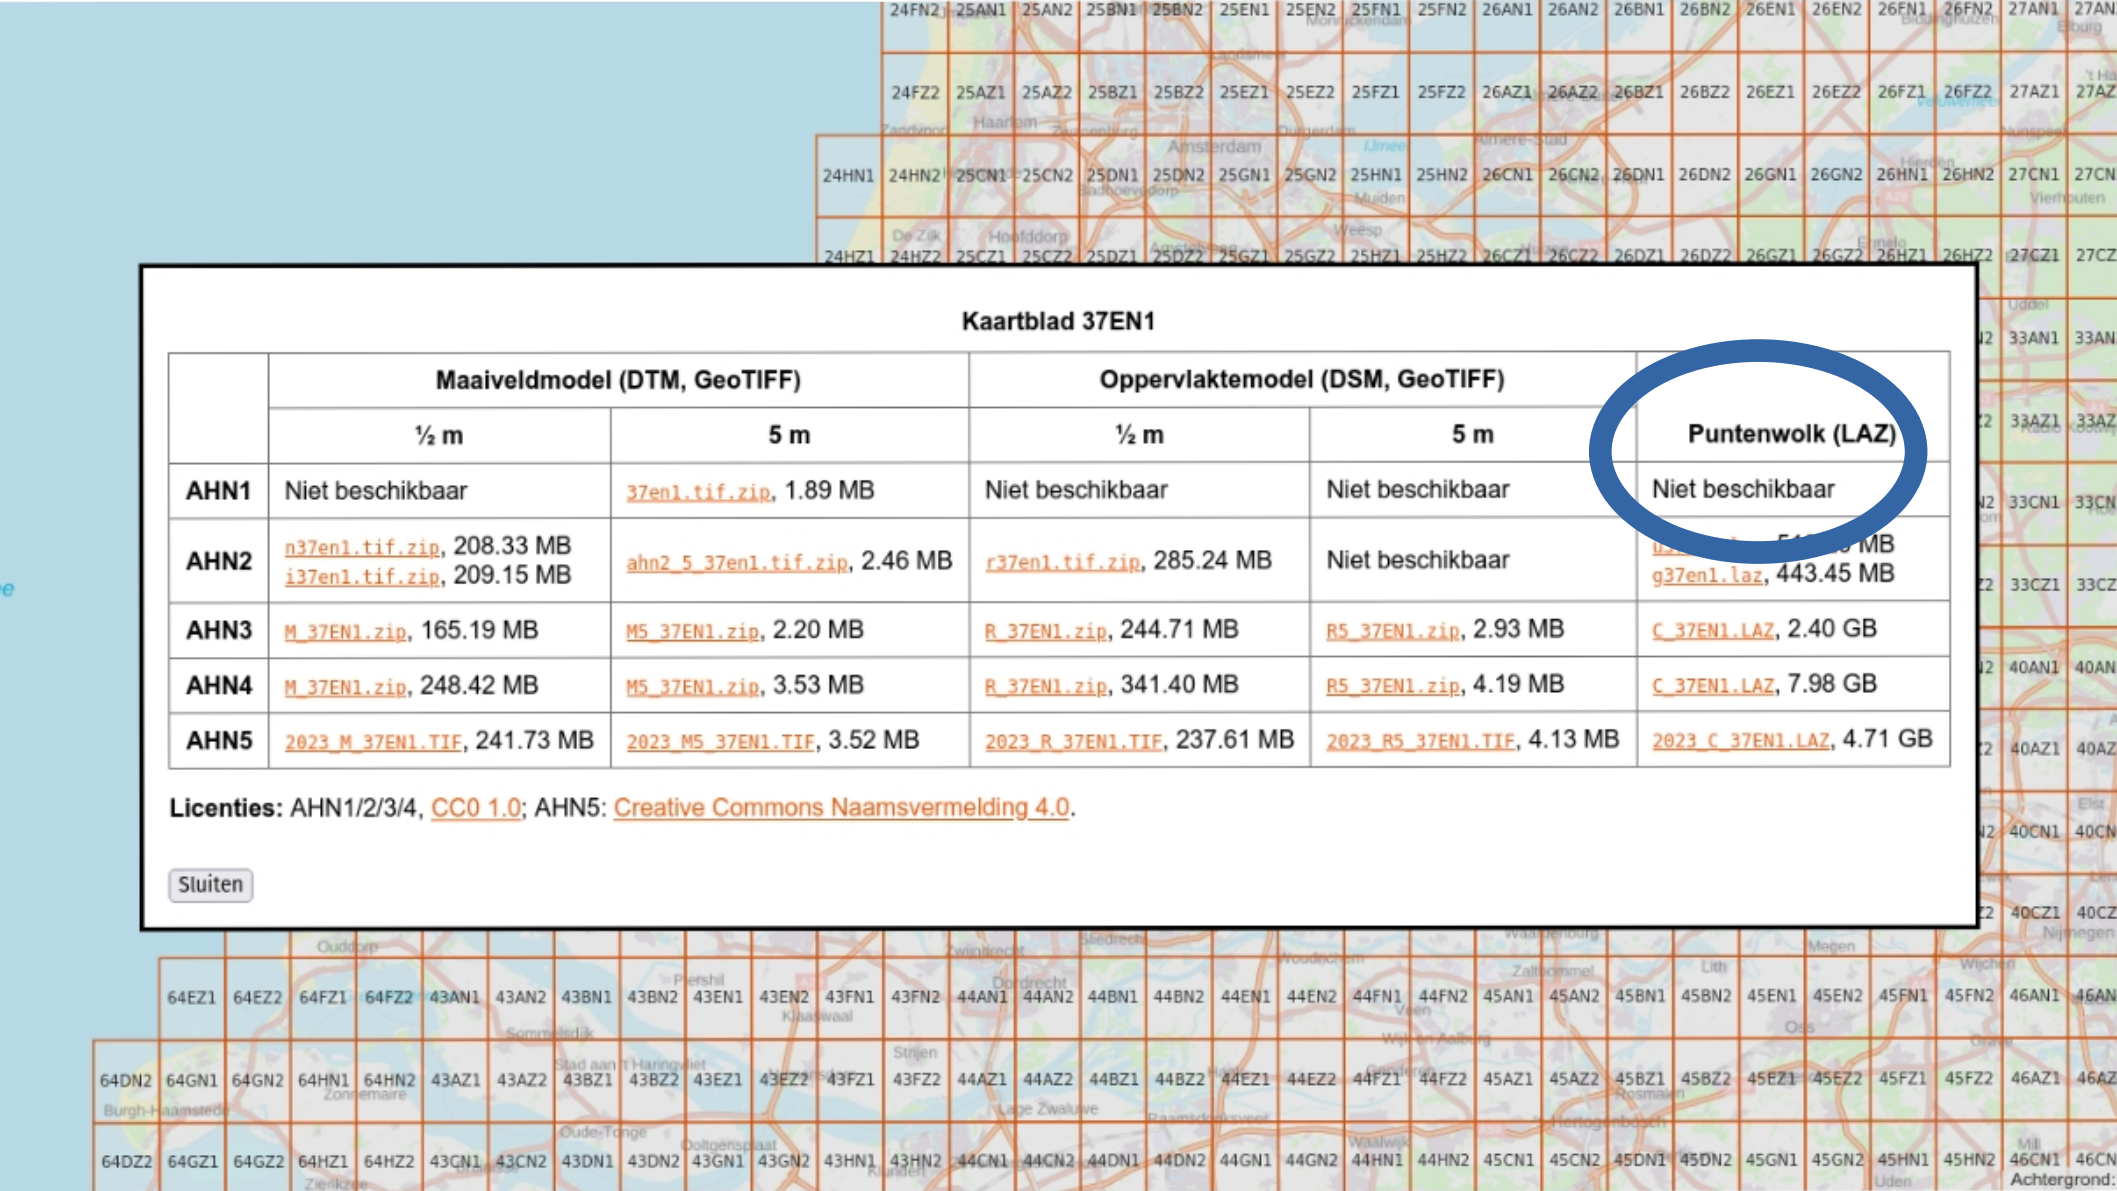

Kaartblad 37EN1

	Maaiveldmodel (DTM, GeoTIFF)		Oppervlaktemodel (DSM, GeoTIFF)		
	½ m	5 m	½ m	5 m	Puntenwolk (LAZ)
AHN1	Niet beschikbaar	37en1.tif.zip, 1.89 MB	Niet beschikbaar	Niet beschikbaar	Niet beschikbaar
AHN2	n37en1.tif.zip, 208.33 MB i37en1.tif.zip, 209.15 MB	ahn2_5_37en1.tif.zip, 2.46 MB	r37en1.tif.zip, 285.24 MB	Niet beschikbaar	[illegible] MB g37en1.laz, 443.45 MB
AHN3	M_37EN1.zip, 165.19 MB	M5_37EN1.zip, 2.20 MB	R_37EN1.zip, 244.71 MB	R5_37EN1.zip, 2.93 MB	C_37EN1.LAZ, 2.40 GB
AHN4	M_37EN1.zip, 248.42 MB	M5_37EN1.zip, 3.53 MB	R_37EN1.zip, 341.40 MB	R5_37EN1.zip, 4.19 MB	C_37EN1.LAZ, 7.98 GB
AHN5	2023_M_37EN1.TIF, 241.73 MB	2023_M5_37EN1.TIF, 3.52 MB	2023_R_37EN1.TIF, 237.61 MB	2023_R5_37EN1.TIF, 4.13 MB	2023_C_37EN1.LAZ, 4.71 GB

Licenties: AHN1/2/3/4, CC0 1.0; AHN5: Creative Commons Naamsvermelding 4.0.

Sluiten

Tile 37EN1

Actueel Hoogtebestand Nederland

Point cloud

The aerial photograph closest to the acquisition date of the AHN iteration was overlaid on the point cloud to produce a colored version for visualisation purposes. Furthermore, for each point cloud a spatial index and textual summary, including elevation histogram, are provided. Please note that the images are not from the same flight, and may be years off. Therefore, small alignment errors and differences occur between them.

Zoom in on the map to ... versions (25 tiles of 1×1.25 km). Tiled versions typically fit into memory (RAM) are easier to handle and are especially suitable for parallel processing and ... Learning.

Versi...		Point cloud		...patial index		Info		License
AHN1	G...und	37en1.laz	22.6 MiB	37e...lax	30 kiB	37en1.txt	20 kiB	Stuurgroep AHN, public domain
	Non-...ound	37en1.laz	3.5 MiB	37e...lax	37 kiB	37en1.txt	14 kiB	Stuurgroep AHN, public domain
	Me...d	37EN1.LAZ	26.1 MiB	37...N1.LAX	34 kiB	37EN1.txt	20 kiB	Stuurgroep AHN (merged product by GeoTiles)
	Colore...	Available for sub-u... only, zoom in to select.						Remixed by GeoTiles*
AHN2	Ground	g37en1.laz	143.5 MiB	g37en1.lax	298 kiB	g37en1.txt	9 kiB	Stuurgroep AHN, public domain
	Non-ground	u37en1.laz	516.2 MiB	u37en1.lax	399 kiB	u37en1.txt	31 kiB	Stuurgroep AHN, public domain
	Merged	37EN1.LAZ	959.7 MiB	37EN1.LAX	2.1 MiB	37EN1.txt	31 kiB	Stuurgroep AHN (merged product by GeoTiles)
	Colored	Available for sub-units only, zoom in to select.						Remixed by GeoTiles*
AHN3	Original	C_37EN1.LAZ	2.4 GiB	C_37EN1.LAX	2.0 MiB	C_37EN1.txt	32 kiB	Stuurgroep AHN, public domain
	Colored	Available for sub-units only, zoom in to select.						Remixed by GeoTiles*
AHN4	Original	C_37EN1.LAZ	8.0 GiB	C_37EN1.LAX	2.0 MiB	C_37EN1.txt	33 kiB	Het Waterschapshuis, public domain
	Colored	(Partially) available for sub-units.						Remixed by GeoTiles*

*The point cloud is in the public domain and the aerial photograph available under CC-BY. However, the colored point cloud is a new product created by GeoTiles.

File checksums (md5) are available for AHN1, AHN2, AHN3 and AHN4.

Raster

Processing

ronna.tudelft.nl (Dell PowerEdge R7525)

- Physical server for benchmarks, parallel computing
- Dell PowerEdge R7525 server
  - 2 x 24-core AMD EPYC 7443 processors (96 threads), 2.85GHz
  - 512 GB buffered memory (RDIMM)
  - Ubuntu operating system
- Disk storage – Direct attached
  - 15.4 Tb SSD SAS – OS
  - 19.2 Tb SSD SATA – Fast read
  - 2x 60 Tb Hard Drive SAS, 7.2K rpm in RAID-5 – Bulk storage
- Network NAS
  - 26 TB on TU Delft storage network

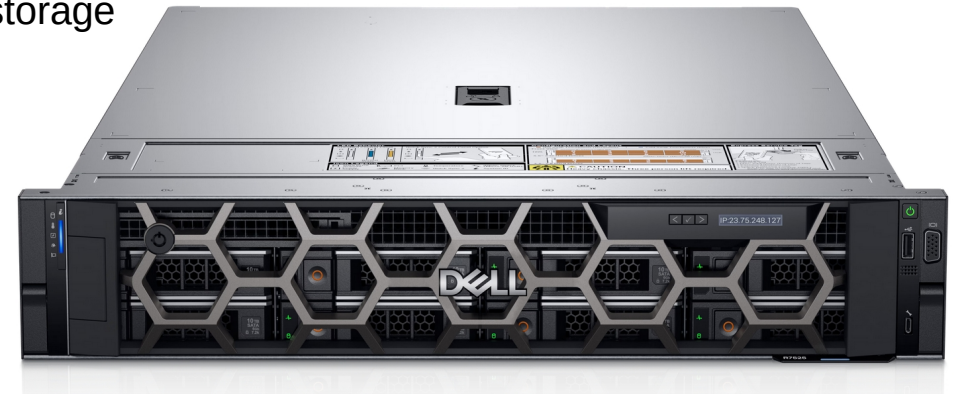

PotreeConverter

- Input: .laz files
- Output: 3 files (accessible by HTTP range requests)
  - metadata.json
  - Octree hierarchy
  - Pointcloud data with points (+attributes!) according to Octree

- Processing = 3 phases
  - Counting
  - Distributing
  - Indexing

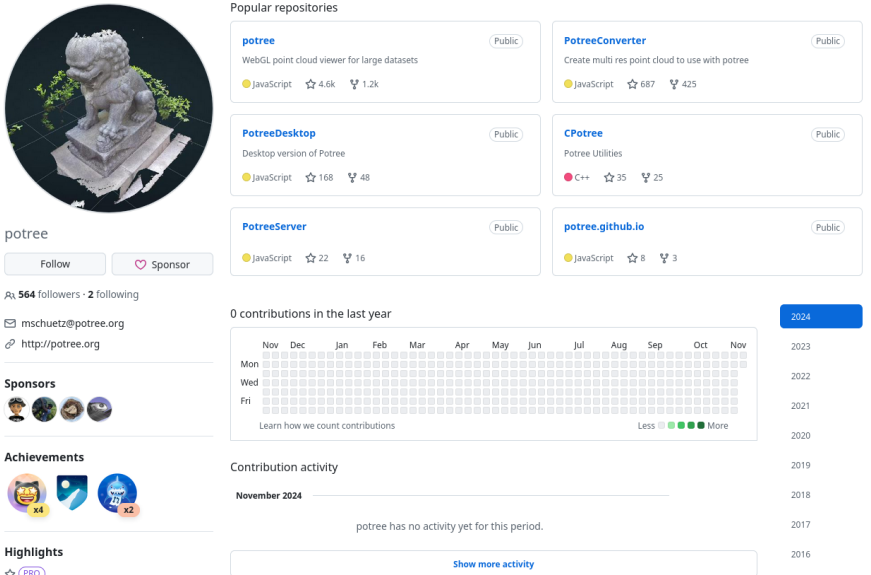

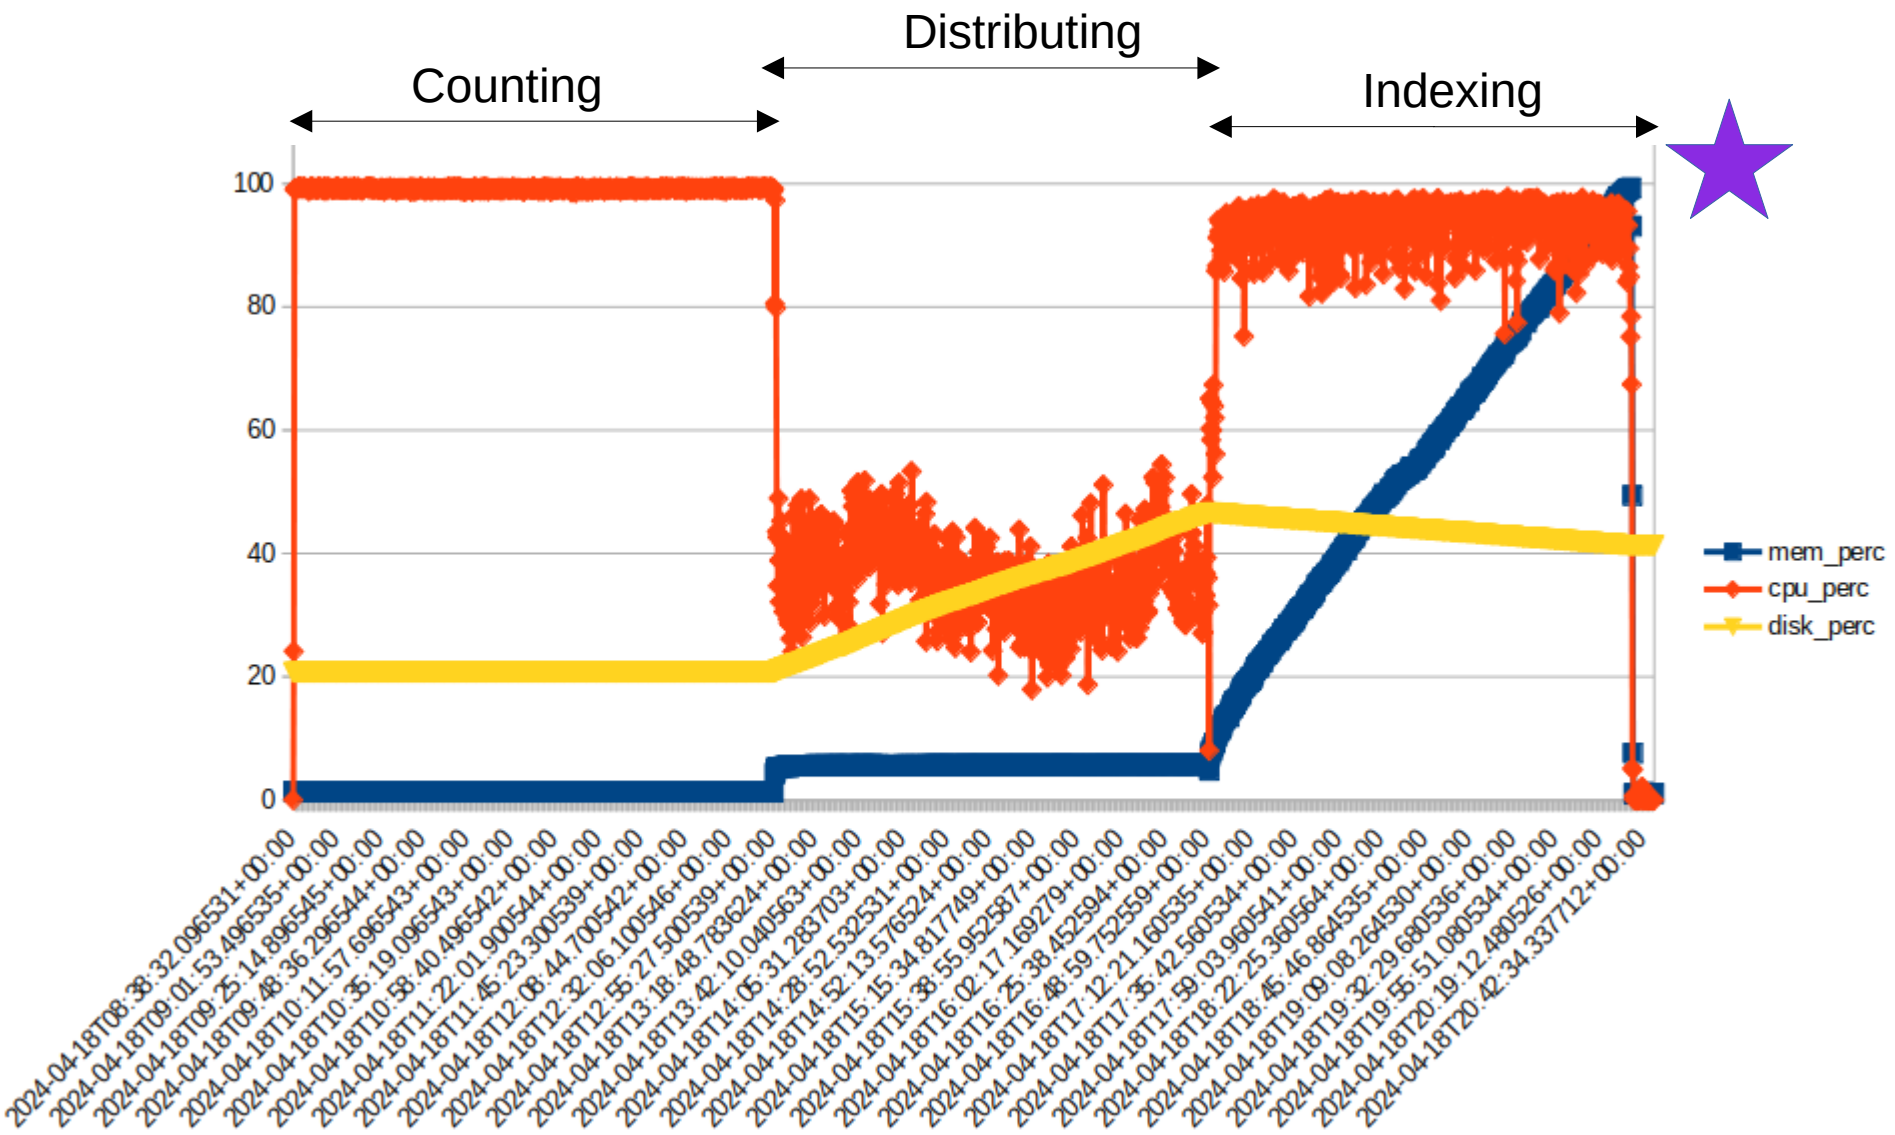
Counting
Distributing
Indexing
100
80
60
40
20
0
mem_perc
cpu_perc
disk_perc
2024-04-18T08:38:32.096531+00:00
2024-04-18T09:01:53.496535+00:00
2024-04-18T09:25:14.896545+00:00
2024-04-18T09:48:36.296544+00:00
2024-04-18T10:11:57.696543+00:00
2024-04-18T10:35:19.096543+00:00
2024-04-18T10:58:40.496542+00:00
2024-04-18T11:22:01.900544+00:00
2024-04-18T11:45:23.300539+00:00
2024-04-18T12:08:44.700542+00:00
2024-04-18T12:32:06.100546+00:00
2024-04-18T12:55:27.500539+00:00
2024-04-18T13:18:48.783624+00:00
2024-04-18T13:42:10.040563+00:00
2024-04-18T14:05:31.283703+00:00
2024-04-18T14:28:52.532531+00:00
2024-04-18T14:52:13.576524+00:00
2024-04-18T15:15:34.817749+00:00
2024-04-18T15:38:55.952587+00:00
2024-04-18T16:02:17.169279+00:00
2024-04-18T16:25:38.452594+00:00
2024-04-18T16:48:59.752559+00:00
2024-04-18T17:12:21.160535+00:00
2024-04-18T17:35:42.560534+00:00
2024-04-18T17:59:03.960541+00:00
2024-04-18T18:22:25.360564+00:00
2024-04-18T18:45:46.864535+00:00
2024-04-18T19:09:08.264530+00:00
2024-04-18T19:32:29.680536+00:00
2024-04-18T19:55:51.080534+00:00
2024-04-18T20:19:12.480526+00:00
2024-04-18T20:42:34.337712+00:00

```
Apr 18 22:34:50 ronna kernel: PotreeConverter invoked oom-killer: gfp_m

Apr 18 22:34:50 ronna kernel:  oom_kill_process.cold+0xb/0x10

Apr 18 22:34:51 ronna kernel: [  pid  ]   uid  tgid total_vm      rss p

Apr 18 22:34:51 ronna kernel: oom-kill:constraint=CONSTRAINT_NONE,nodem
user.slice/user-151276.slice/session-3188.scope,task=PotreeConverter,pi

Apr 18 22:34:51 ronna kernel: Out of memory: Killed process 390176 (Pot
shmem-rss:0kB, UID:151276 pgtables:1090144kB oom_score_adj:0

Apr 18 22:35:36 ronna kernel: oom_reaper: reaped process 390176 (Potree
```

potree / PotreeConverter

Code | Issues 134 | Pull requests 27 | Actions | Projects | Wiki | Security | Insights

Fix memory leak. #666

Open tobias93 wants to merge 1 commit into potree:develop from tobias93:develop

Conversation 2 | Commits 1 | Checks 0 | Files changed 2 +2 −4

tobias93 commented on Aug 28 — Contributor

When indexing large point clouds, we always had problems with PotreeConverter consuming enormous amounts of memory. PotreeConverter would only run with 100GiB of swap memory, if not more.

Long and thorough investigations revealed the issue: The poisson sampling sorts the points by their distance to the center of the node. This is done using std::sort with the relatively new (C++17) std::execution::par_unseq argument so the sorting happens in parallel. In Linux, the C++ standard library relies on Intel TBB for the implementation of the parallel sorting. The TBB library is where the memory is leaking.

There is not much information available online about this, but I think this is the issue that is causing the leak: https://community.intel.com/t5/Intel-oneAPI-Threading-Building/std-sort-std-execution-par-unseq-has-a-memory-leak-on-Linux/m-p/1582773

The fix for PotreeConverter is to simply not use par_unseq when sorting the points. This means that the sorting won't happen in parallel any more. However PotreeConverter already parallelizes over the chunks, so it should not be an issue.

We tested the fix with a point cloud consisting of approximately 100GB of uncompressed LAS files. In the current version of PotreeConverter you can see the memory climbing up to ~60GB during the indexing phase, while it stays below 12GB with the fix:

Old:

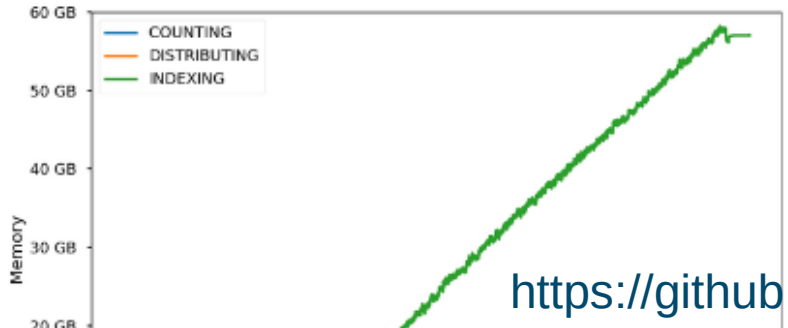

Reviewers
No reviews
Still in progress? Learn about draft PRs

Assignees
No one assigned

Labels
None yet

Projects
None yet

Milestone
No milestone

Development
Successfully merging this pull request may close these issues.
None yet

Notifications — Customize
Unsubscribe
You're receiving notifications because you commented.

3 participants

Roeland (GeoDelta), thanks for the tip!

https://github.com/potree/PotreeConverter/pull/666

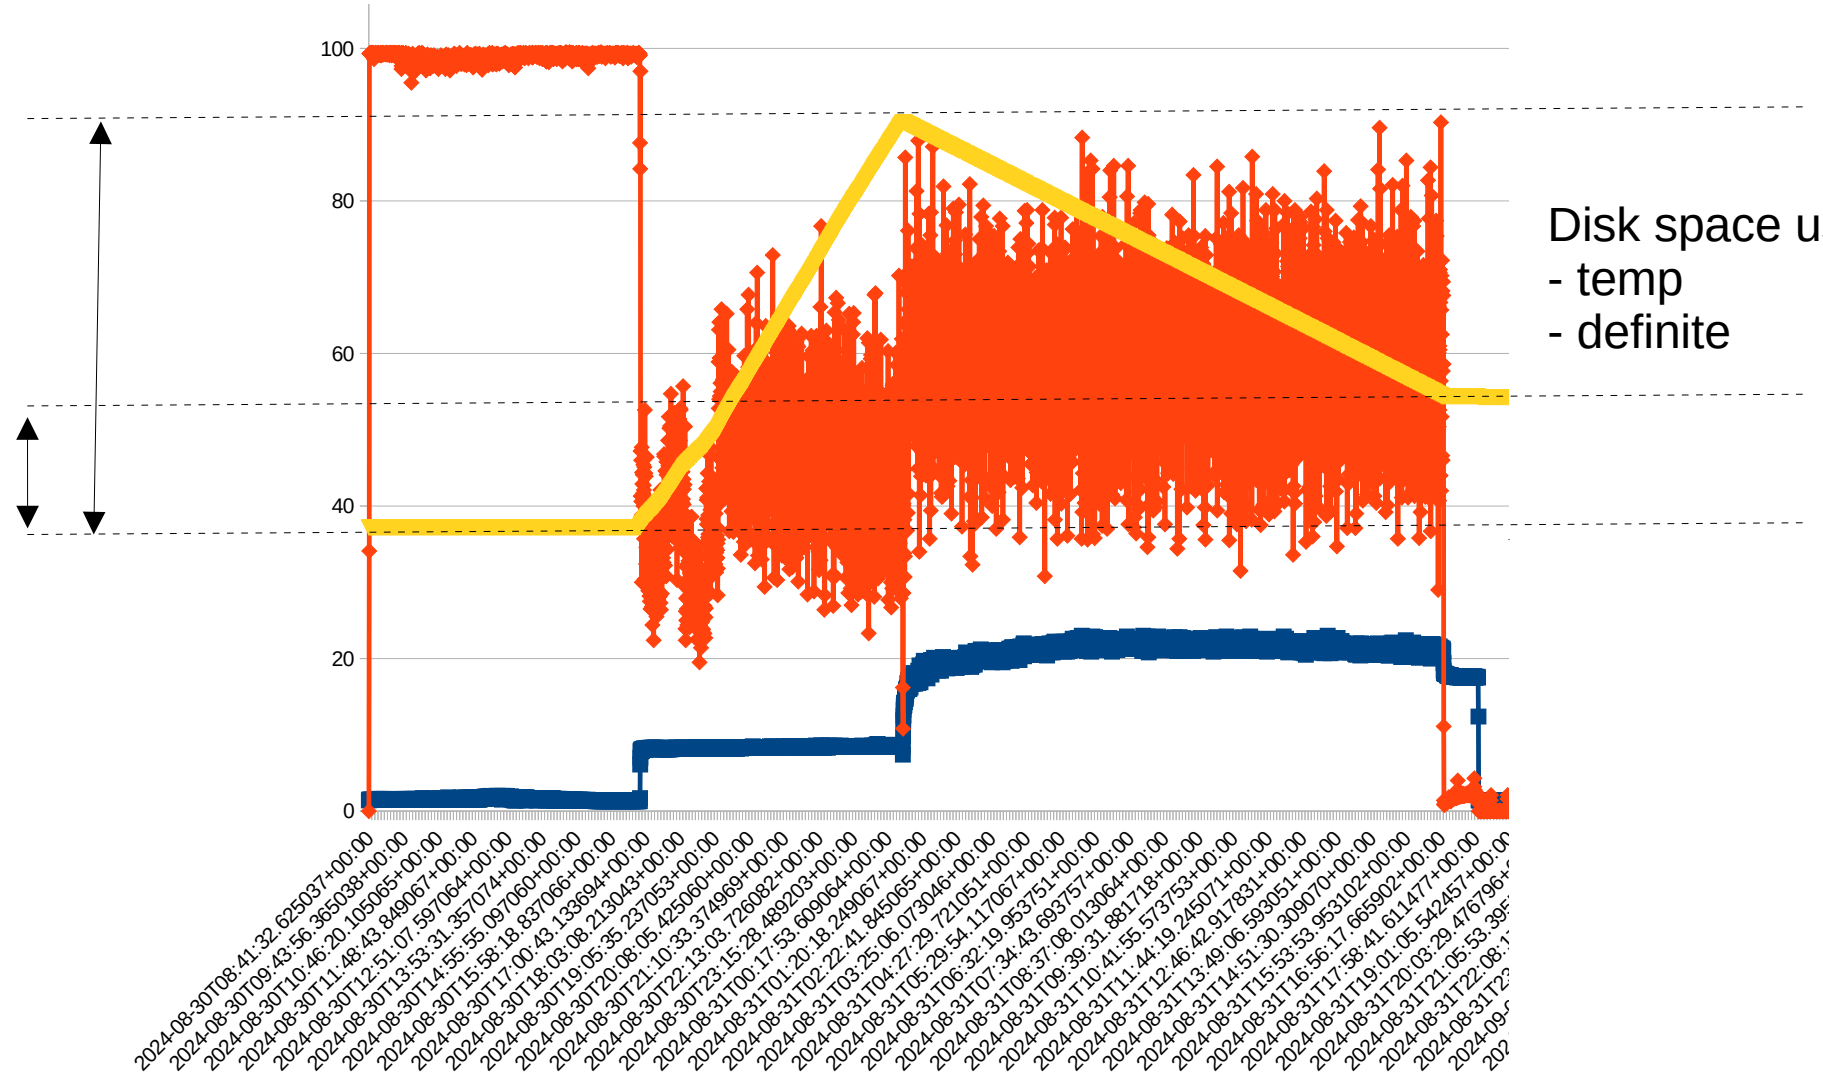

Disk space usage: Huge!
- temp
- definite

potree conversion on ronna

name	#points	#files in	Input size (Gb)	Output size (Gb)	Duration (hours)
AHN1	11.984.853.767	1358	33.1	44.2	0.2
AHN3	557.925.797.136	1374	2390.7	5035.8	12.3
AHN4	947.364.043.509	1381	6145.4	9931.8	33.4
AHN5 (2023 part)	289.944.615.278	499	1394.3	2695.9	8.6

AHN2 converted before (previous project)
Duration conversion on ronna, multi-user

AHN-potree conversion statistics: In-/output size

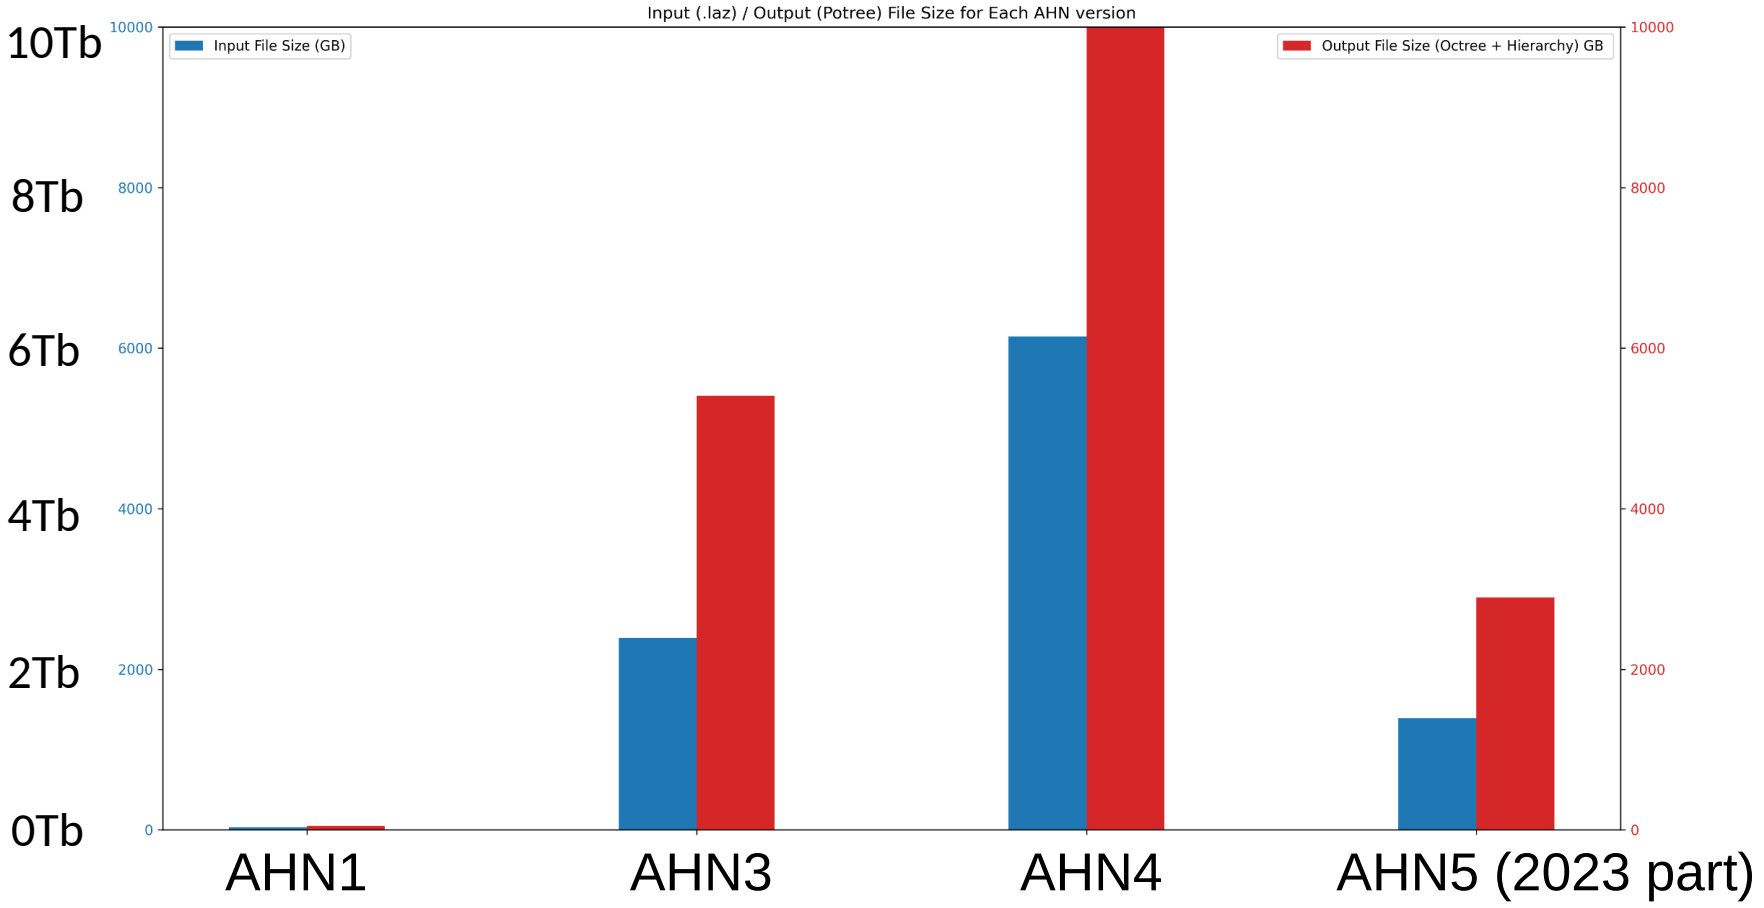

AHN-potree conversion: input size vs duration

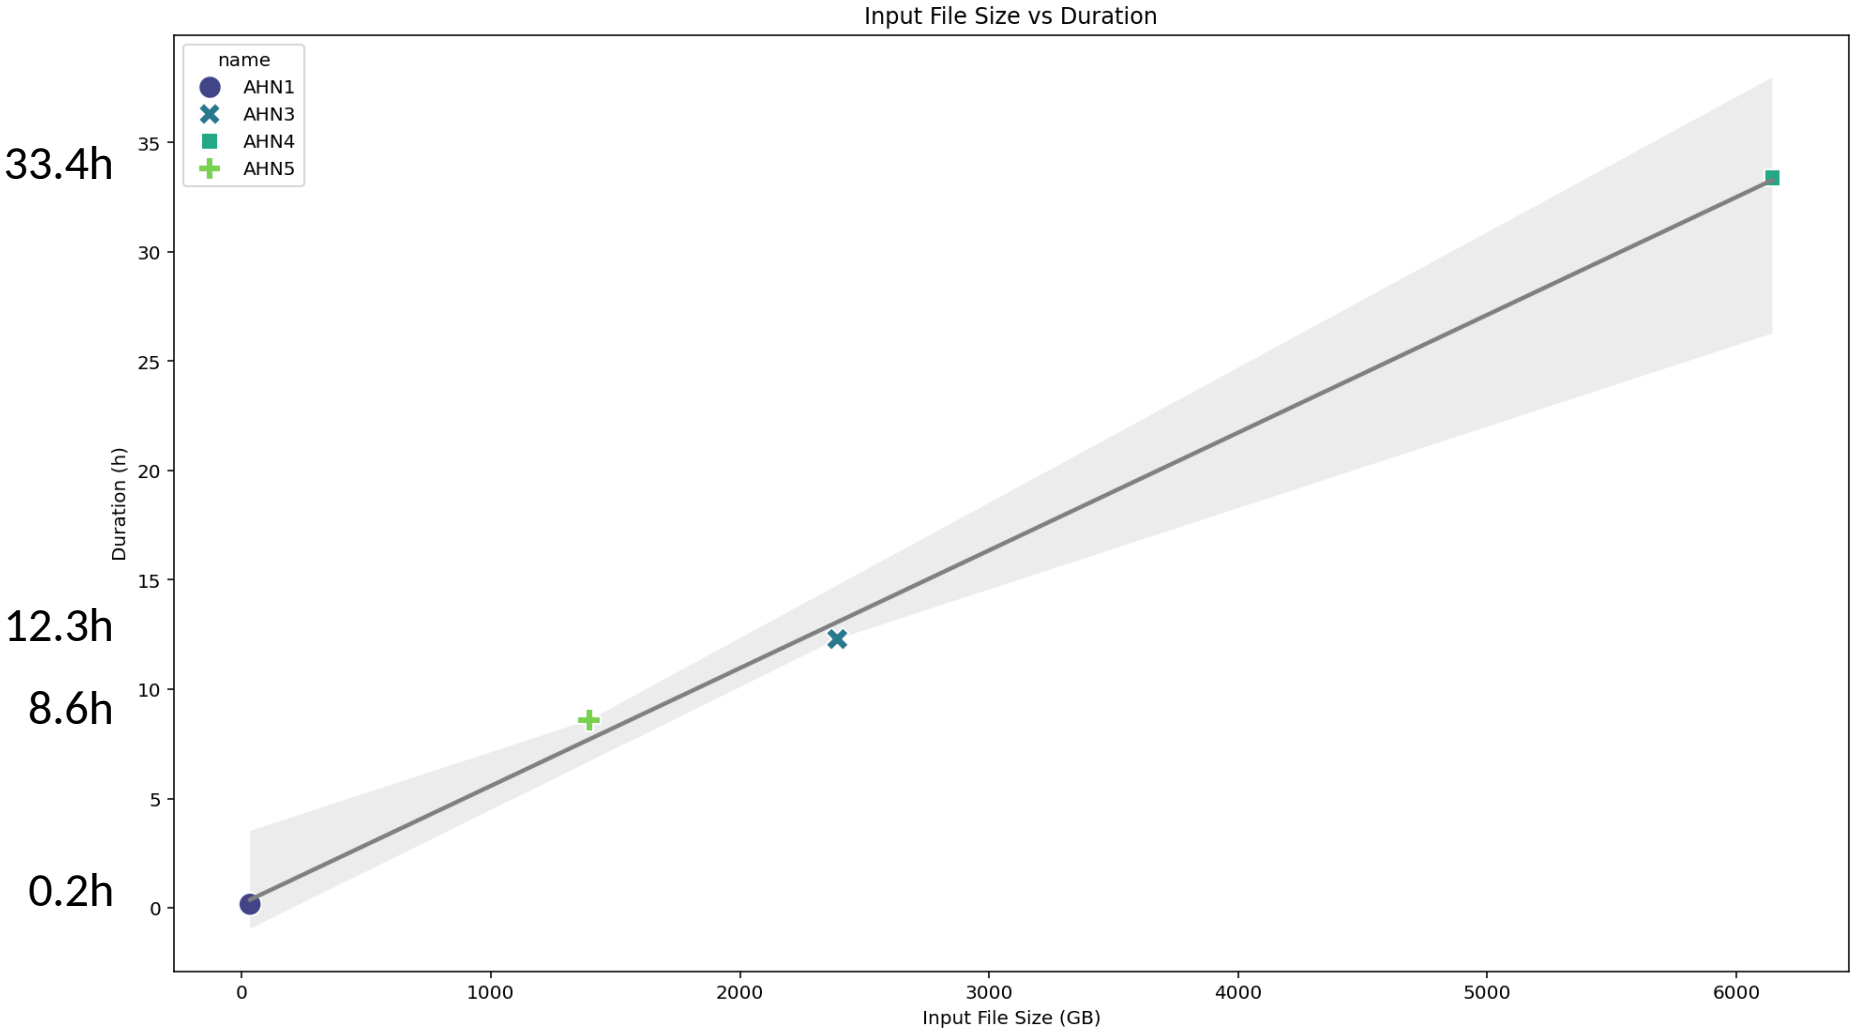

Viewing

AHN-1 + AHN-2 + AHN-3 + AHN-4 + AHN-5

http://viewer.pointclouds.nl all* AHN versions in potree

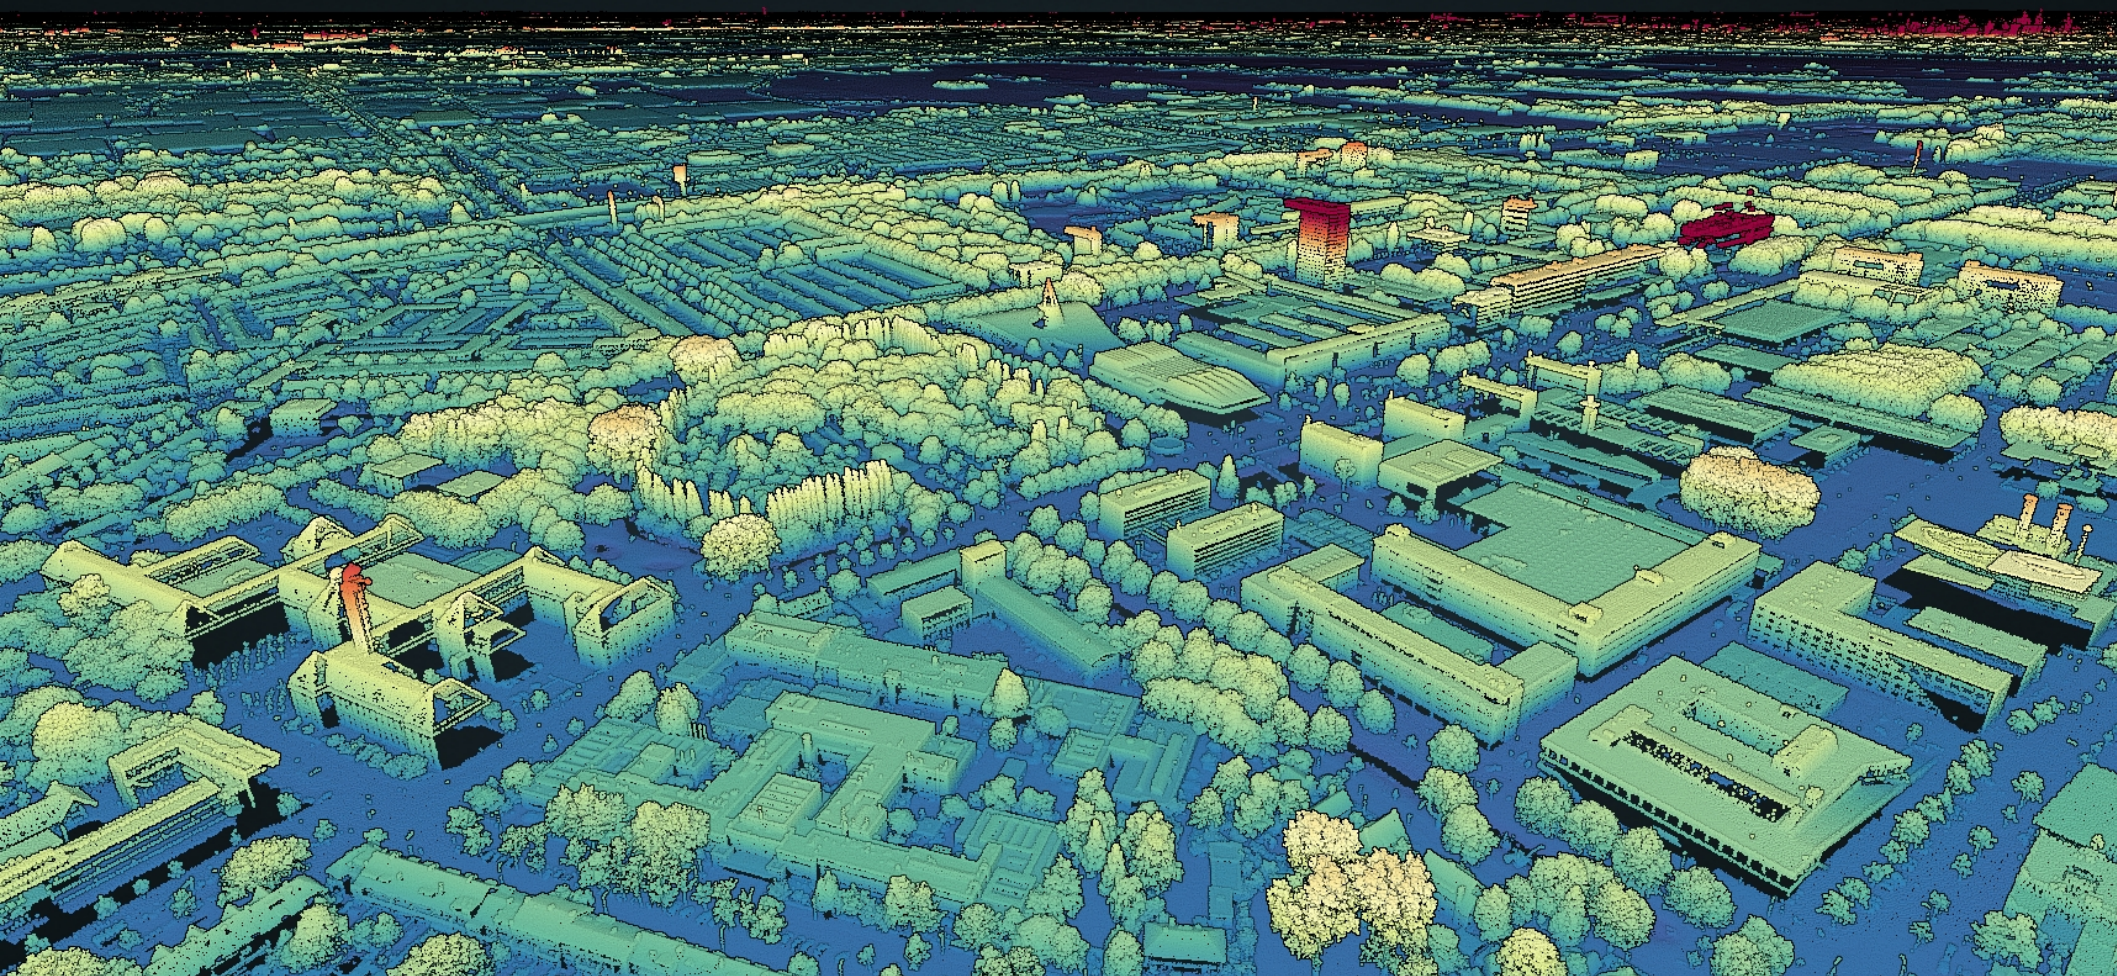

AHN-1 + AHN-2 + AHN-3 + AHN-4 + AHN-5

All potree visualizations, e.g. show Number of returns: 1-7

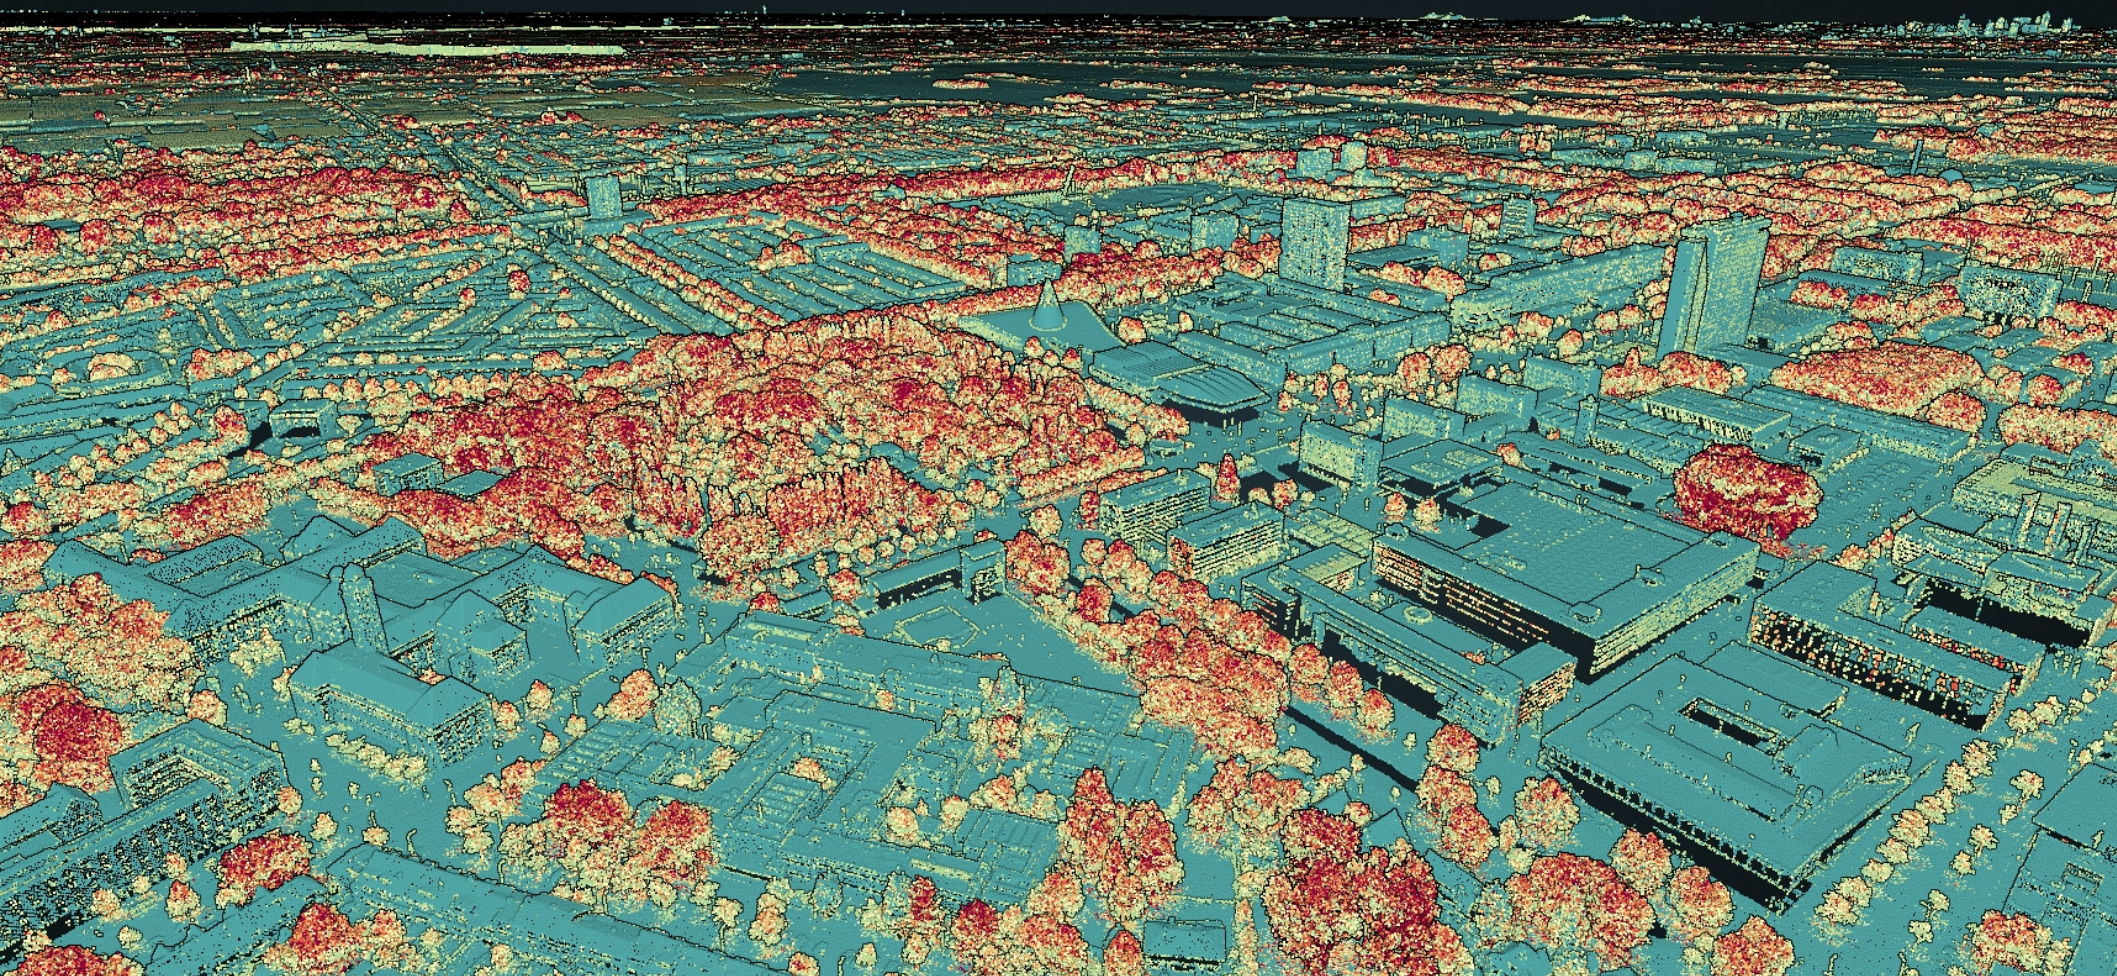


AHN-1 + AHN-2 + AHN-3 + AHN-4 + AHN-5
... or classification (given data set)

AHN-1 + AHN-2 + AHN-3 + AHN-4 + AHN-5

Tools: height, area, volume, angles, profile, ...

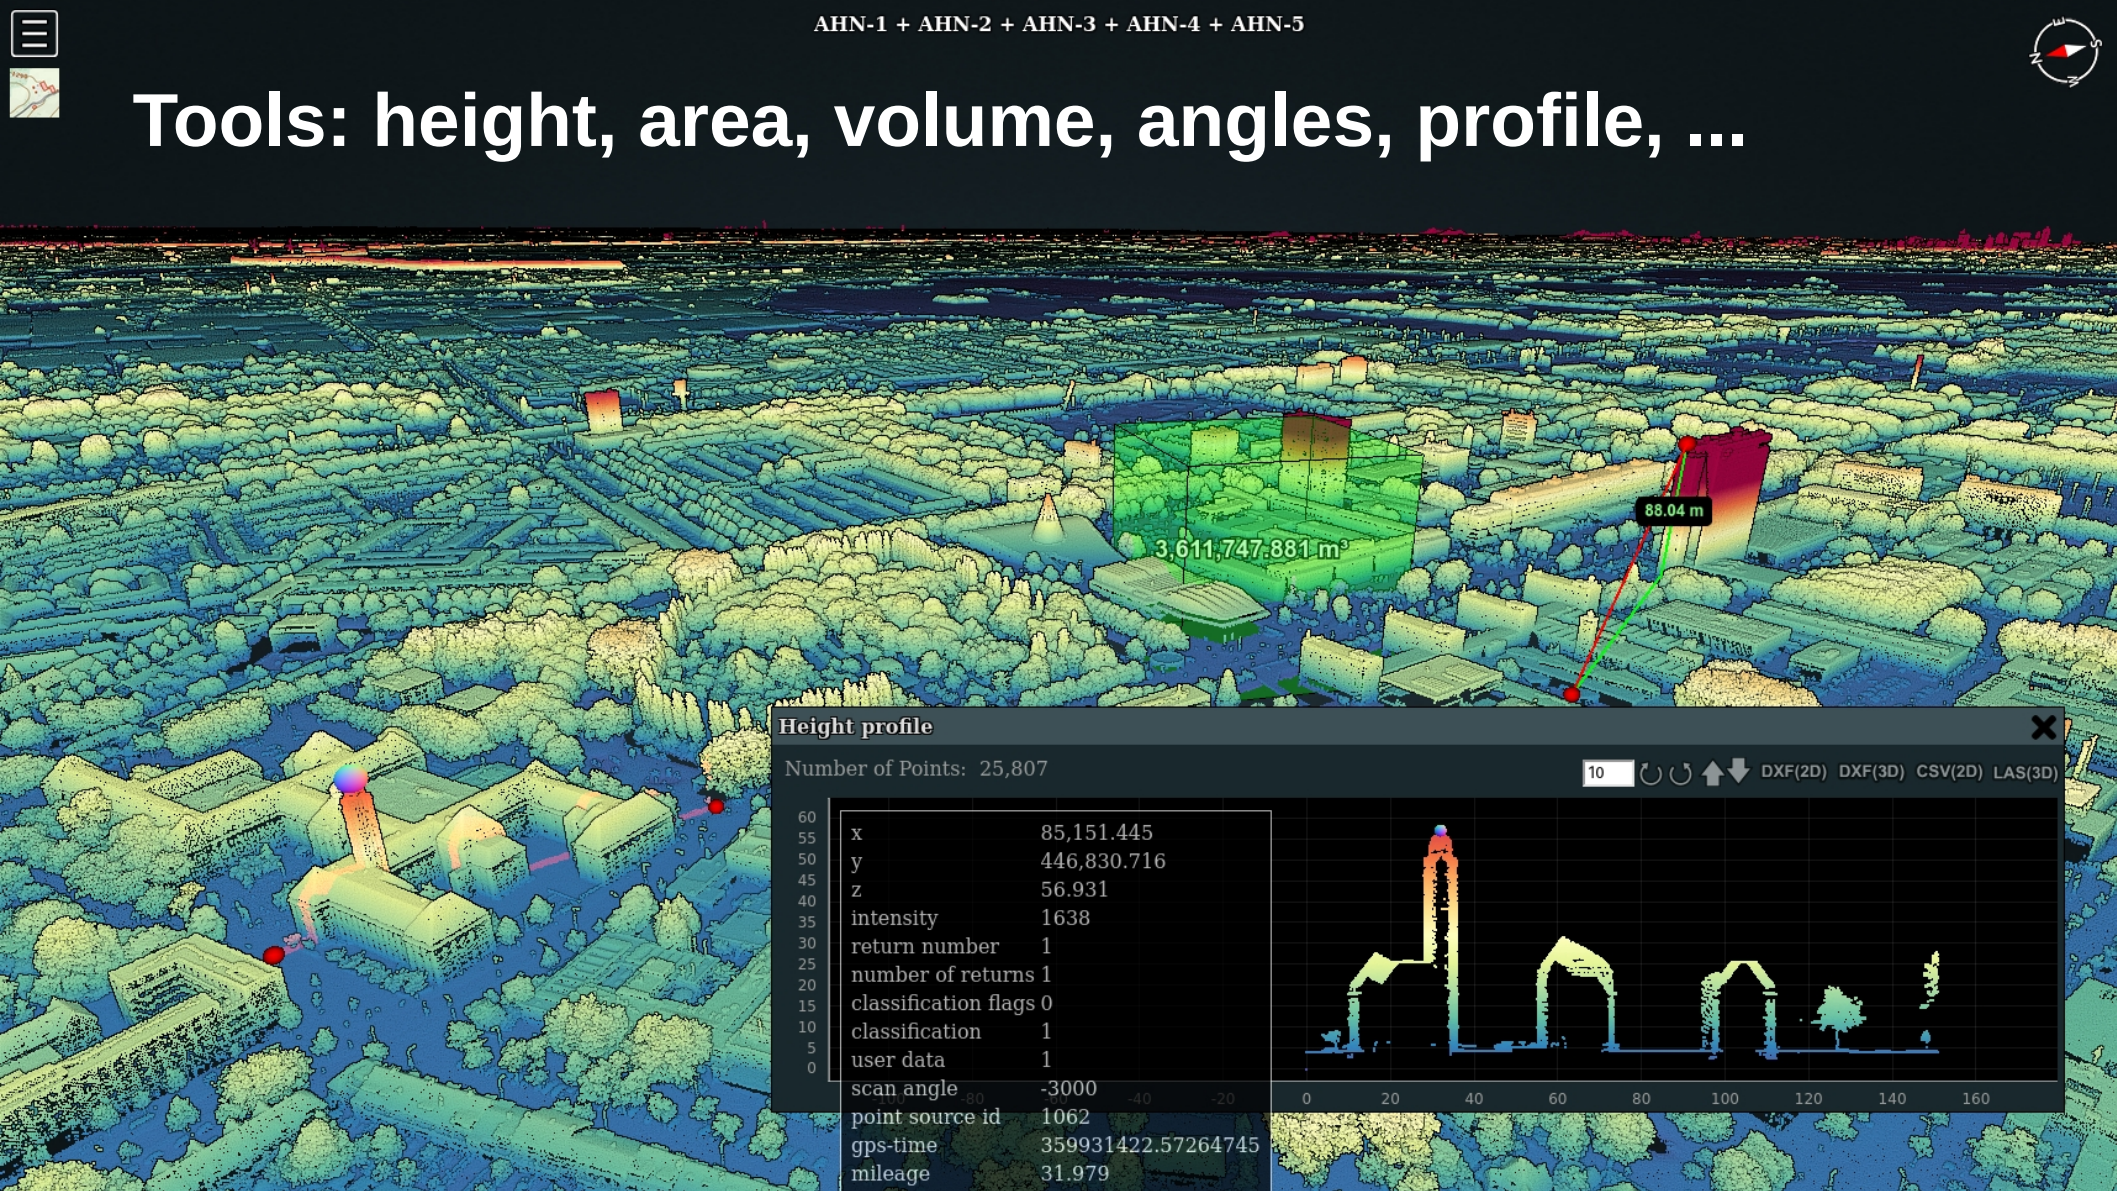

Visual change detection (AHN2 → 5)

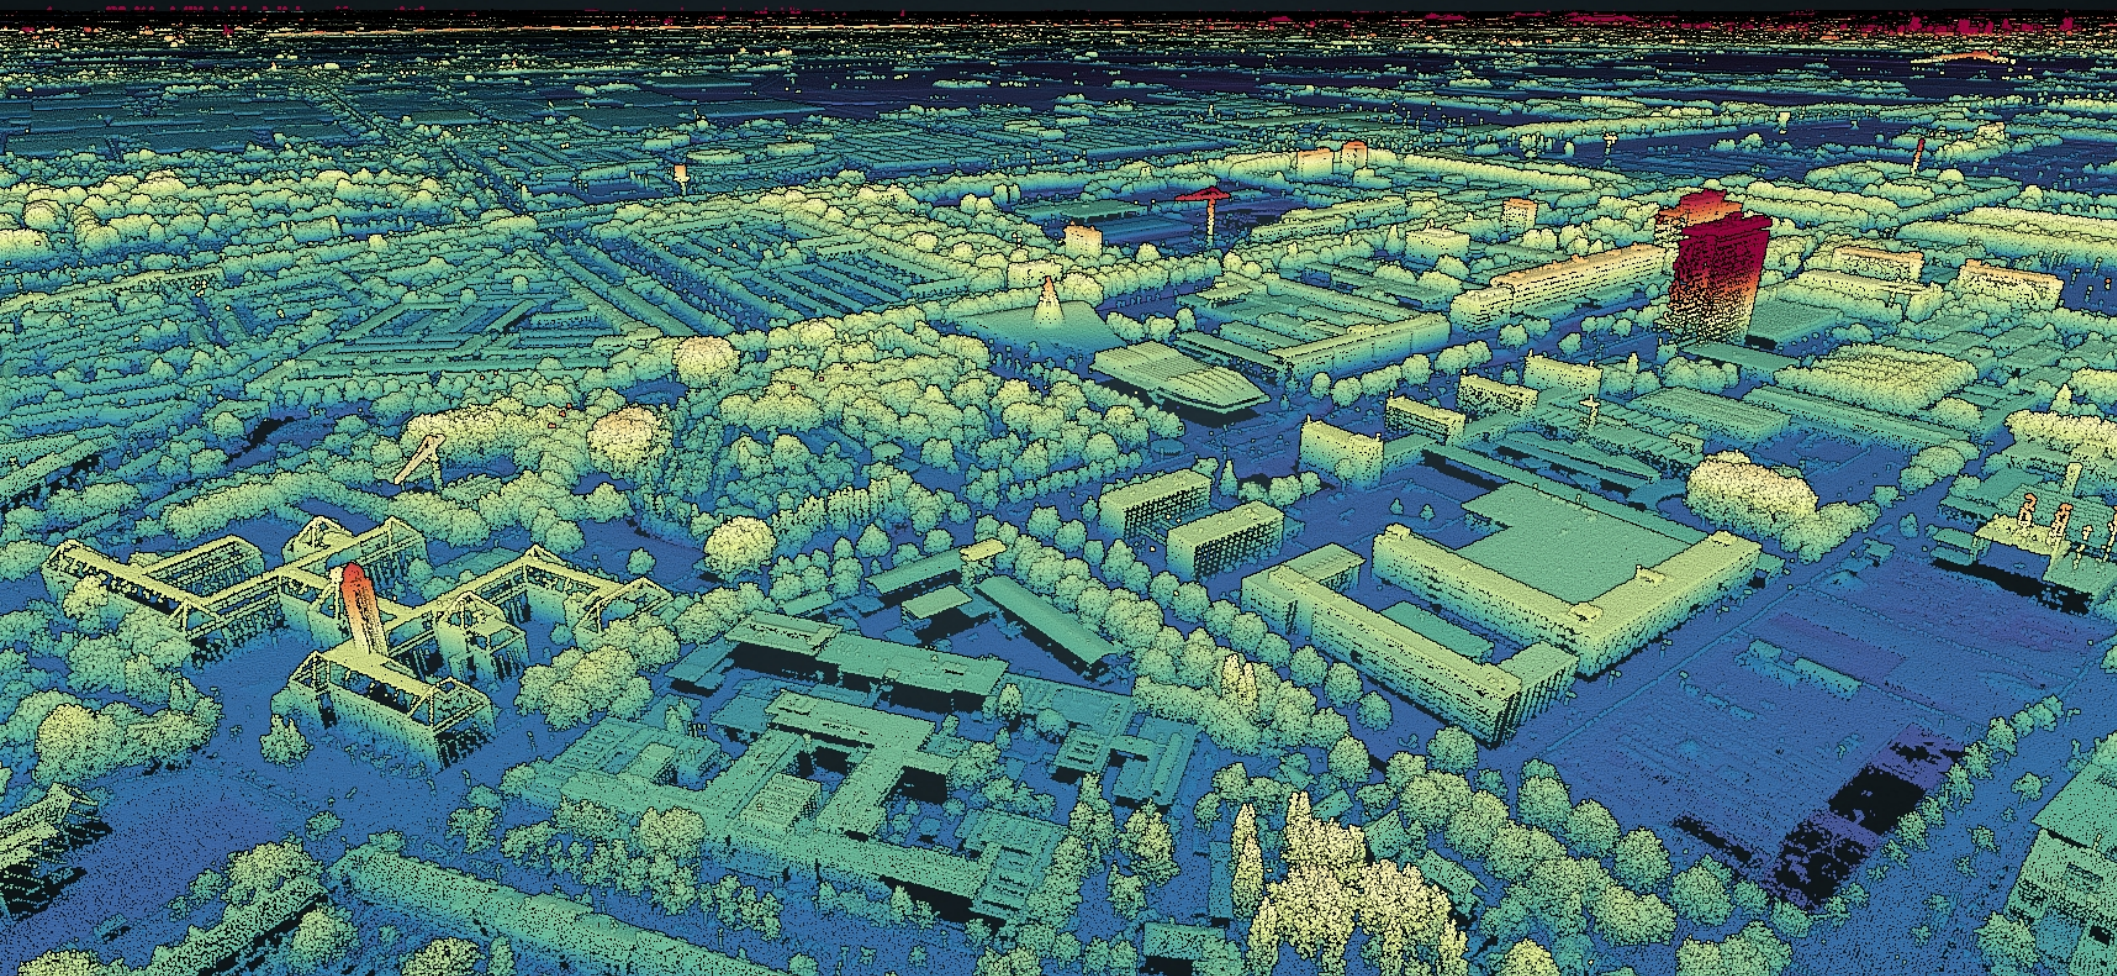

Visual change detection (AHN2 → 5)

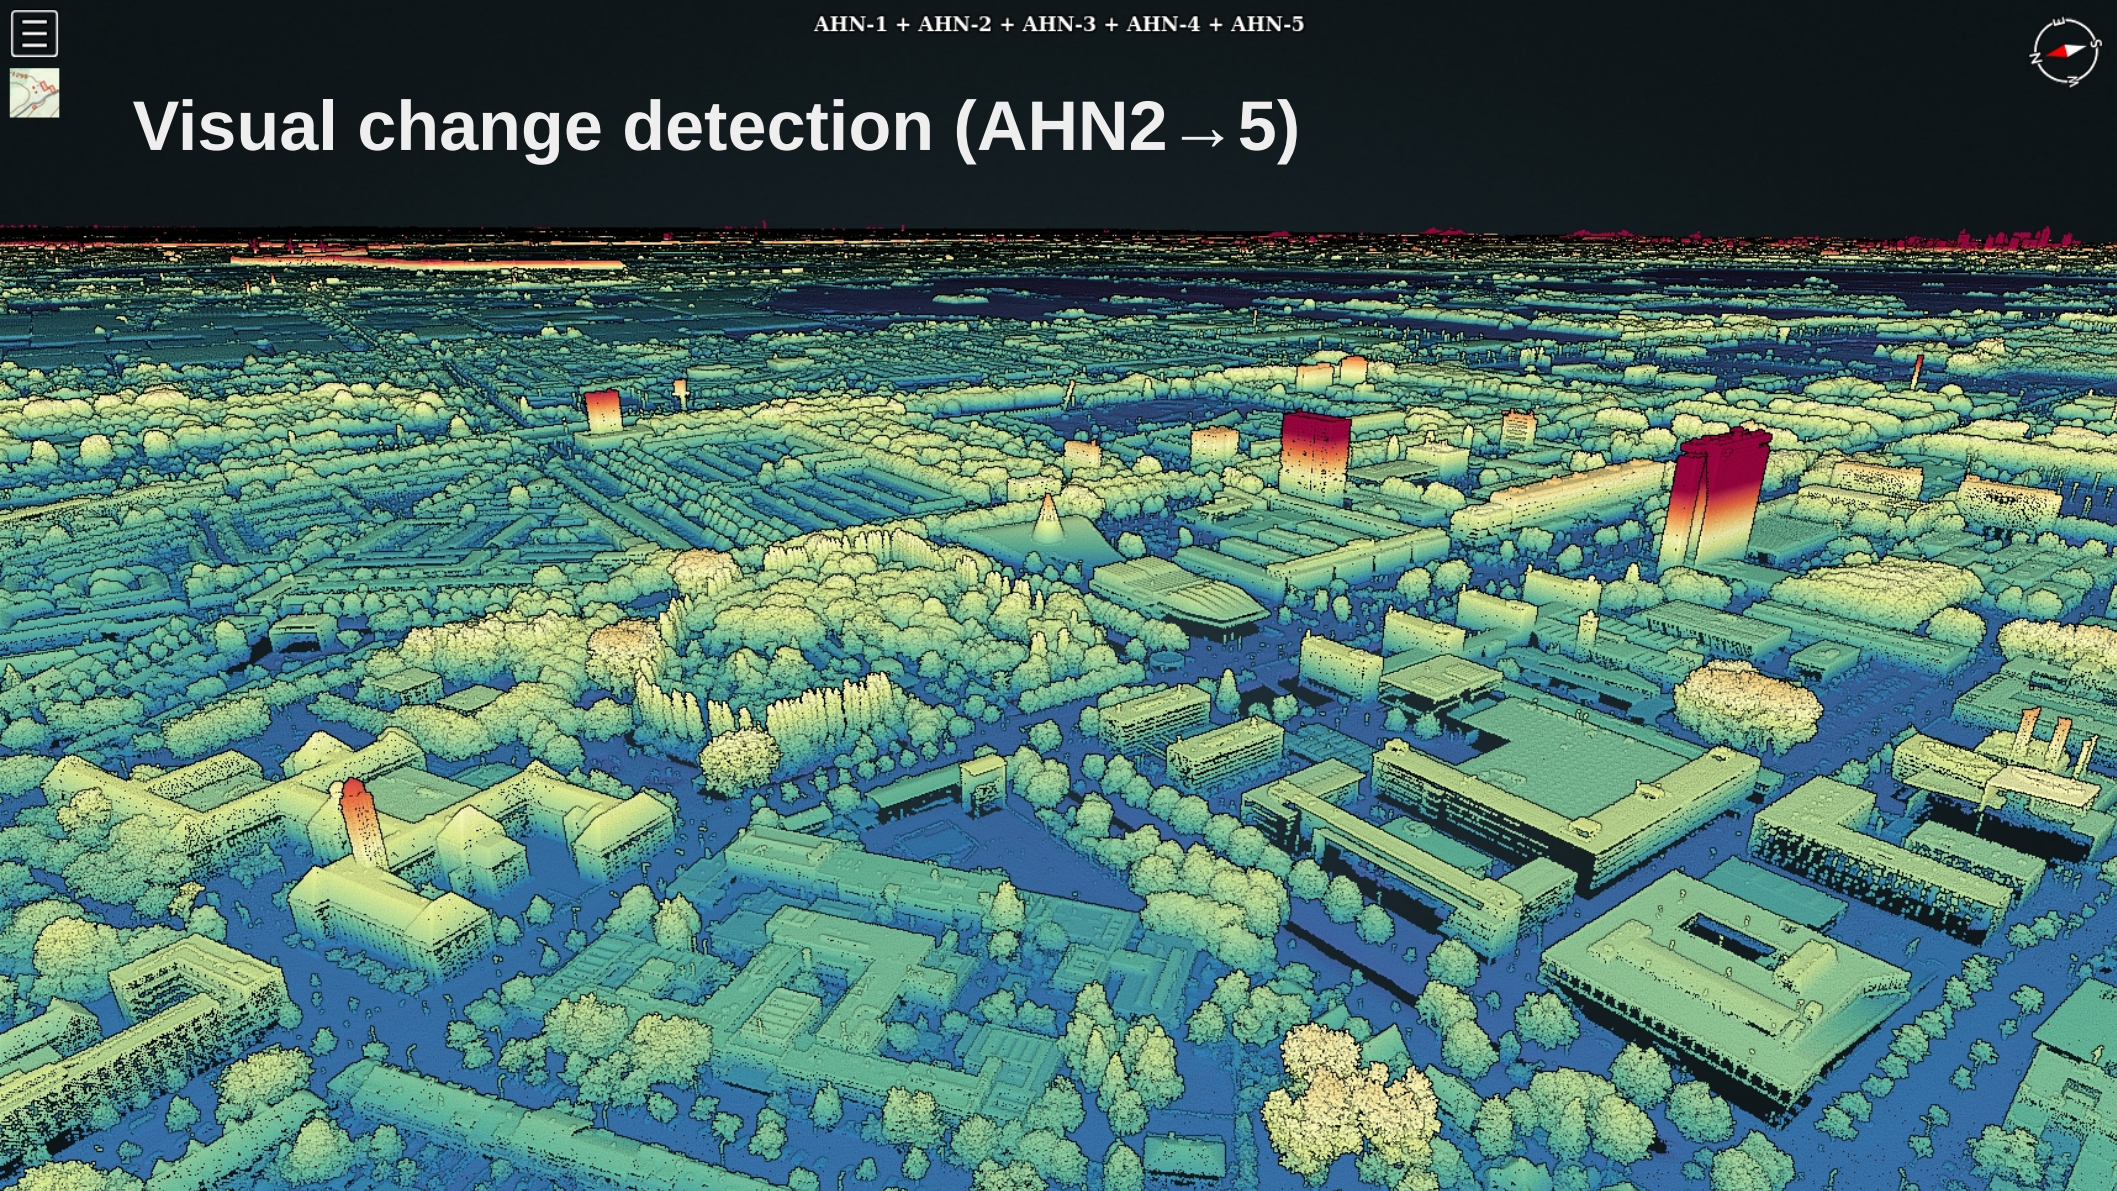

Concluding remarks

- Conversion challenges
- Visualization of + Direct Interaction with AHN data sets
- Governance (‘keep it running’)
  - *AHN5 part ‘24 + ‘25 (once ready) to be added

Questions?

dr.ir. Martijn Meijers

GIS-technology, Digital Technologies, Architectural Engineering + Technology,
Faculty of Architecture and the Built Environment, Delft University of Technology

tel: (+31) 15 278 56 42
mailto:b.m.meijers@tudelft.nl
https://www.tudelft.nl/en/staff/b.m.meijers/

TU Delft GDMC nD-PointCloud netherlands eScience center