Agisoft Benchmarking

Sören Hese

Forschungscolloq 2019
EO Department
FSU Jena

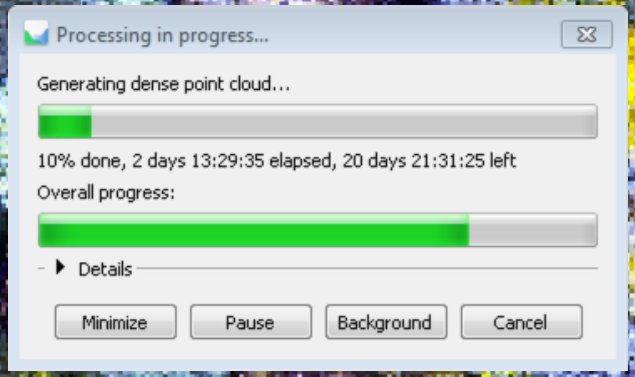

Motivation

Very long processing times with excessive image overlap
Variation 1: retro-Nadir Crown Structure Mapping with DSLR&DiGPS

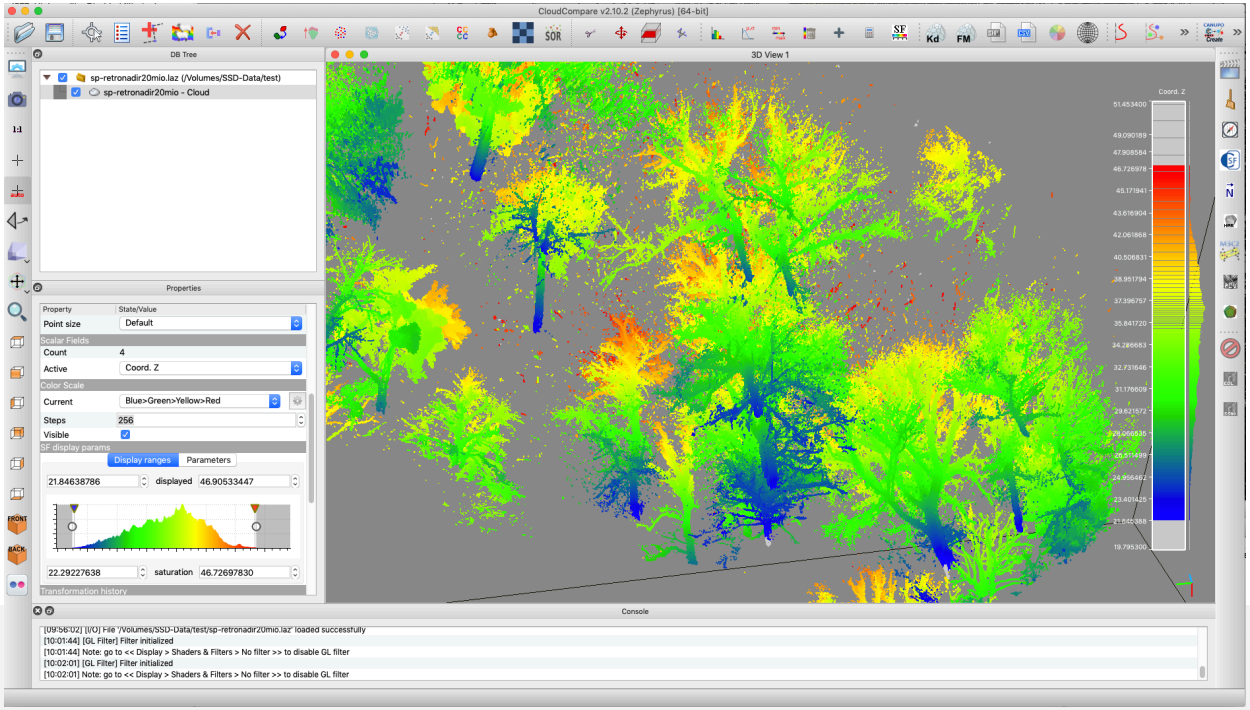

Motivation

Very long processing times with excessive overlap in Copter POI flight modi
Variation 2: POI modi with multiple altitude layers (CROWN PROJECT)

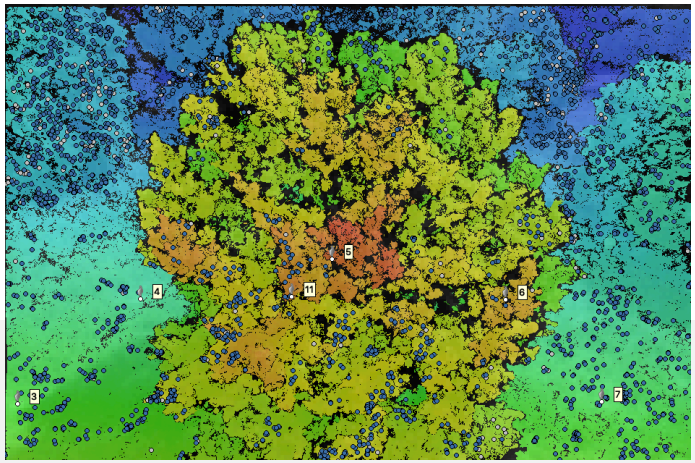

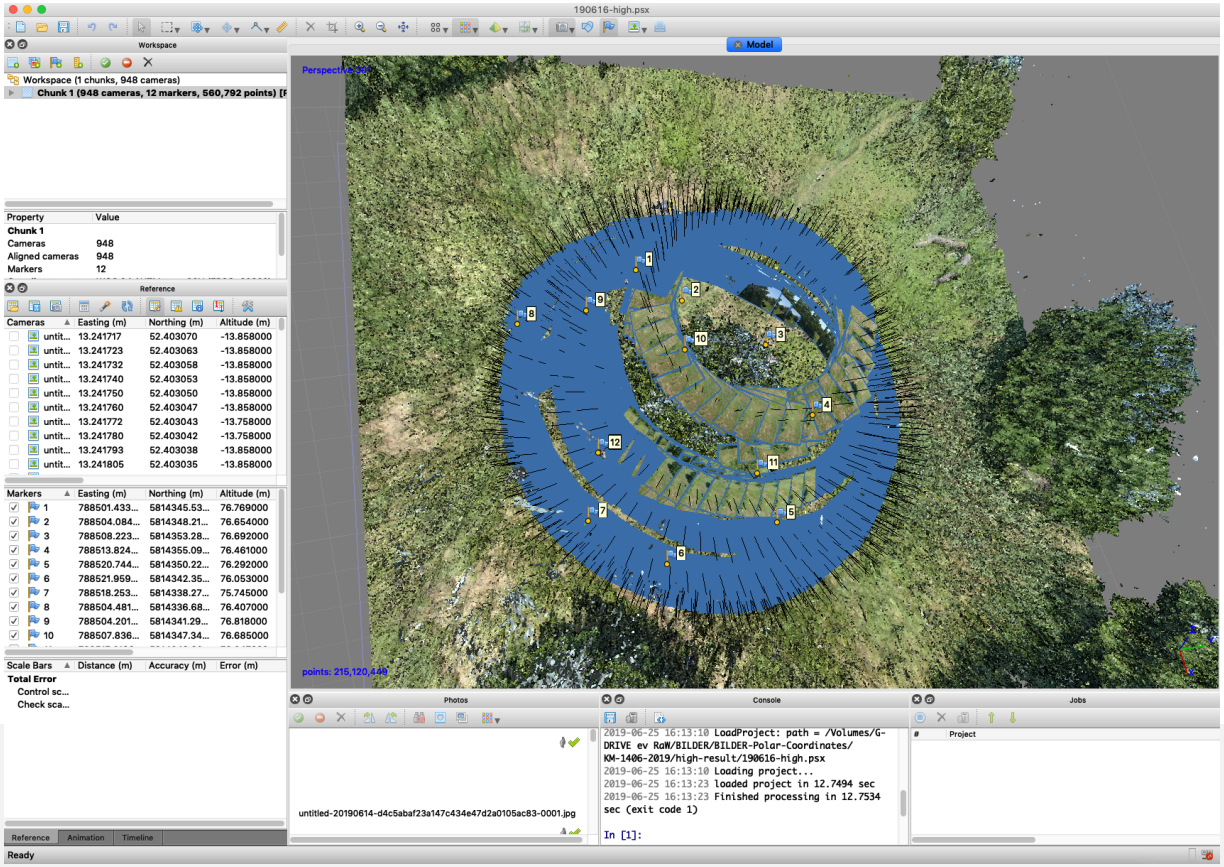

Motivation

Very long processing times with excessive Nadir data mosaic

Variation 3: 90% Overlap Mosaics – Hainich National Parc (HAINICH)

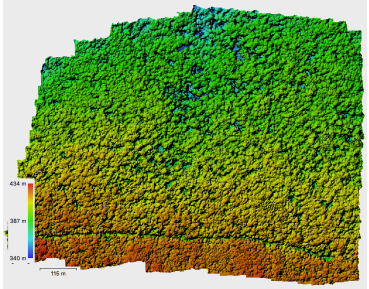

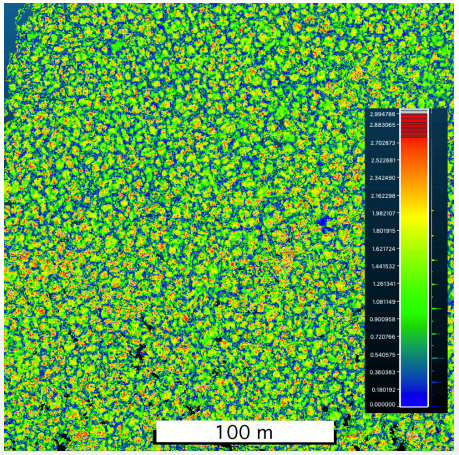

UAV BASED MULTI SEASONAL DECIDUOUS TREE SPECIES ANALYSIS IN THE HAINICH NATIONAL PARK USING MULTI TEMPORAL AND POINT CLOUD CURVATURE FEATURES

S. Hese [1]*, C. Thiel[2], A. Henkel [3]

[1] Friedrich-Schiller-Universität Jena, Institut für Geographie, Lehrstuhl für Fernerkundung, 07743 Jena, Germany - soeren.hese@uni-jena.de

[2] Deutsches Zentrum für Luft- und Raumfahrt (DLR), Institut für Datenwissenschaften, Mälzerstraße 3, 07745 Jena, Germany - christian.thiel@dlr.de

[3] Nationalpark Verwaltung Hainich, Bei der Marktkirche 9, 99947 Bad Langensalza, Germany - andreas.henkel@nnl.thueringen.de

KEY WORDS: Tree species, Crown Surface, Gaussian Curvature, Roughness, Point Cloud, UAV, Multi Seasonal

ABSTRACT:

Low cost UAV systems are a flexible and mobile platform for very detailed spatial high-resolution point cloud and surface height mapping projects. This study investigates the potential of the DJI Phantom 4 Pro 3D point clouds and derived crown surface height information in combination with RGB spectral information for mapping of deciduous tree species in the Hainich national park area. RGB image data was captured in August, early October and November 2018 to create a multi seasonal spectral dataset for a 100 ha test area. The flight campaigns were controlled from the Hainich flux tower platform in 40 m height owned and operated by University of Göttingen in the central part of the park area. Absolut georeferencing accuracy of the datasets was improved using 7 DGPS measured control points within the stand structure on small forest clearings. Image files and ground control points were processed to a dense point cloud model with 2.6 billion points (approximately 200k points per tree crown object) using the Agisoft Metashape cluster processing environment. Additionally, a digital surface model and a true ortho image mosaic with 3 cm spatial resolution was generated. For the differentiation of deciduous tree species, a reference data set with coordinates for the tree species Fagus sylvatica (beech), Fraxinus excelsior (ash), Acer pseudoplatanus (sycamore maple), Carpinus betulus (hornbeam) and dead trees and early defoliated trees was defined. The study site is however dominated by Fagus sylvatica and Fraxinus excelsior. We studied two different groups of features: tree crown surface height variability parameters using point cloud densities, point cloud height variance, local standard deviation of gaussian curvature, standard deviation of local point cloud roughness and multi temporal normalised spectral features using multi seasonal uncalibrated UAV RGB data. Analysis of feature separability showed that very high-resolution point cloud surface curvature properties with small neighbourhood radii can differentiate some tree species types but we also found multitemporal spectral ratios based on RGB data to be very successful in differentiating the main tree species.

Results of this work show that super fine very dense point cloud models and derived roughness measures of mixed forest stand surfaces hold valuable information for deciduous species discrimination and will likely also be very useful for morphological analysis of tree crown types.

1. INTRODUCTION

Small Off The Shelf (OTS) quadro copter systems provide a very cost effective, highly flexible, mobile platform for high resolution point cloud mapping. Various studies showed the full potential of these small and flexible platforms (Mumbo et al. 2017). An overview is also given in Colomina and Molina (2014). Especially in very complex 3D environments the automatic obstacle avoidance, low copter weight, long flight times and precise manoeuvring are important advantages of small OTS systems in comparison with large octocopter systems. Especially for the mapping of 2.5D structures - mapping a continuous height surface of a stand and not volume 3D structures - the airborne Lidar and the image-based UAV Structure from Motion (SfM) approach create comparable surface representation results (Thiel and Schmullius 2016). Since the development of very fast Graphic Processing Units (GPU) with multiple processing cores the computing times for image matching have decreased considerably making projects with hundreds to thousands of full-resolution 12-34 megapixel image files possible now. Studies in the past have shown that this is possible even with un-georeferenced images from various sources and with varying image quality (Snavely, Seitz and Szeliski 2008). Photogrammetric grade cameras have been used together with GPS and IMU (Inertial Measurement Unit) data to extract camera positions and orientation data and various applications with very high-resolution airborne data have been published (Oliner et al. 2006). With SfM approaches the 3D modelling is done directly based on multiple overlapping images and position/orientation data is not necessarily needed. These concepts do not require any orientation information although orientation and GPS information is an advantage for fast data processing. The applied concepts usually use features that are independent of dissimilarities between overlapping images. Namely the SIFT and the SURF operator (Bay et al. 2008, Lingua, Marenchino & Nex 2009, Lowe 2004) play a critical role for the successful introduction in various 3D point cloud mapping software solutions. First publications go back into the 1990s and earlier (Lowe 1999 and Ullman 1979). The SfM

* Corresponding author

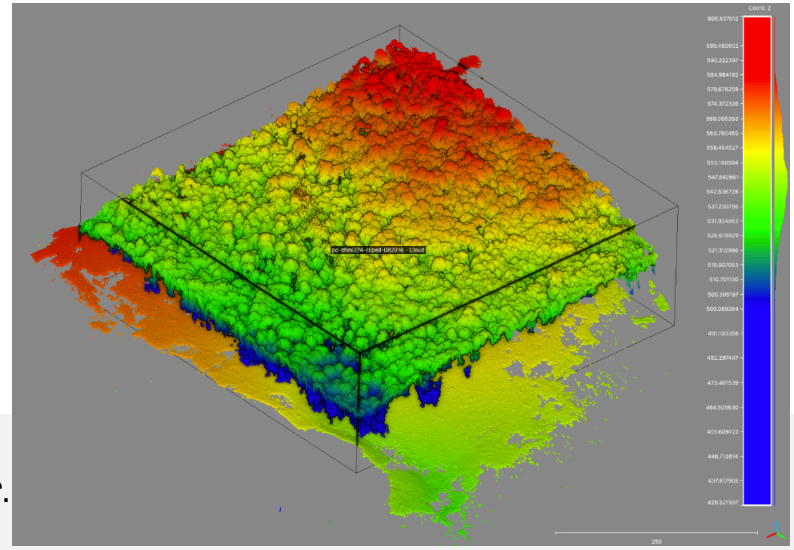

Motivation

Diverse Hardware Environment – PC – Server - HPC

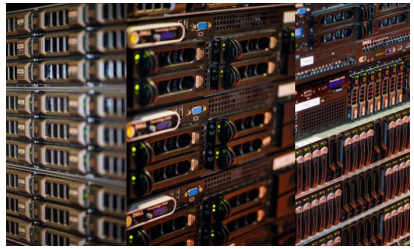

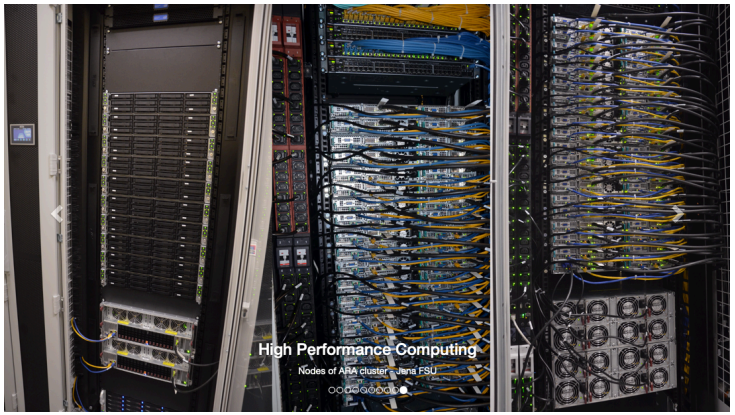

GPU PERFORMANCE COMPARISON

	P100	V100	Ratio
DL Training	10 TFLOPS	120 TFLOPS	12x
DL Inferencing	21 TFLOPS	120 TFLOPS	6x
FP64/FP32	5/10 TFLOPS	7.5/15 TFLOPS	1.5x
HBM2 Bandwidth	720 GB/s	900 GB/s	1.2x
STREAM Triad Perf	557 GB/s	855 GB/s	1.5x
NVLink Bandwidth	160 GB/s	300 GB/s	1.9x
L2 Cache	4 MB	6 MB	1.5x
L1 Caches	1.3 MB	10 MB	7.7x

Hardware Configurations:

1. MacBook Pro 2017 (16GB, Radeon 480 4GB, TB3 eGPU Radeon Vega64, Toshiba 2TB SSD, 2.9GHZ Intel Quad Core i7)
2. PC Intel Xeon 3.2GHz (X8000), 128GB RAM, Nvidia GTX1080Ti 11GB, GTX980, 500GB SSD
3. Dell Poweredge (Clustered)
    1. R710 12 Cores (ltos) noGPU (nfs/cifs RAID5 harddiscs)
    2. R720 12 Cores (noaa) noGPU (nfs/cifs RAID5 harddiscs)
    3. R730 24 Cores (earthcare) noGPU (nfs/cifs RAID5 harddiscs)
4. ARA Cluster GPU_v100 (Tesla v100) partition130, 128GB 2x Xeon 2.2Ghz each12 cores
5. ARA Cluster GPU_p100 (Tesla p100) partition129, 128GB 2x Xeon 2.2Ghz each 12 cores
6. Combination of 2. and 3.
7. (Missing: Terrasense, ARA standard_b)

Agisoft Resources Req. per Process

Agisoft Hardware Requirements per Process (GPU: CUDA /OPenCL)

- Match & Alignment (GPU/CPU) (Scale Invariant Feature Transform SIFT related Algorithm/GreedyAlgorithm+Bundle Adjust, matching and cam alignment – cam position) sparse point cloud
- Depth Maps (GPU) & Dense Cloud (GPU/CPU/SSD/RAM) Height-Field-Method with pair wise depth field computation, Fast-method with MVS – MultiView Stereo approach
- Build Mesh / Build Texture (GPU/RAM!) Poisson Reconstruction ala Meshlab, Texture Atlas Blending
- Build DEM (CPU/SSD/RAM)
- Build Orthomosaic (CPU/SSD/RAM)

Agisoft Network Processing

1. Processing nodes are registered at the coordination server
2. Client sends job to coordination server
3. Coordination Server distributes job to registered nodes.

Start a coordination server
soeren@itos:/usr/local/metashape-pro> metashape.sh —server —control 141.35.158.37 —dispatch 141.35.158.37
Start a Processing Node:
soeren@noaa:/usr/local/metashape-pro> metashape —node —dispatch 141.35.158.37

Agisoft Network Monitor

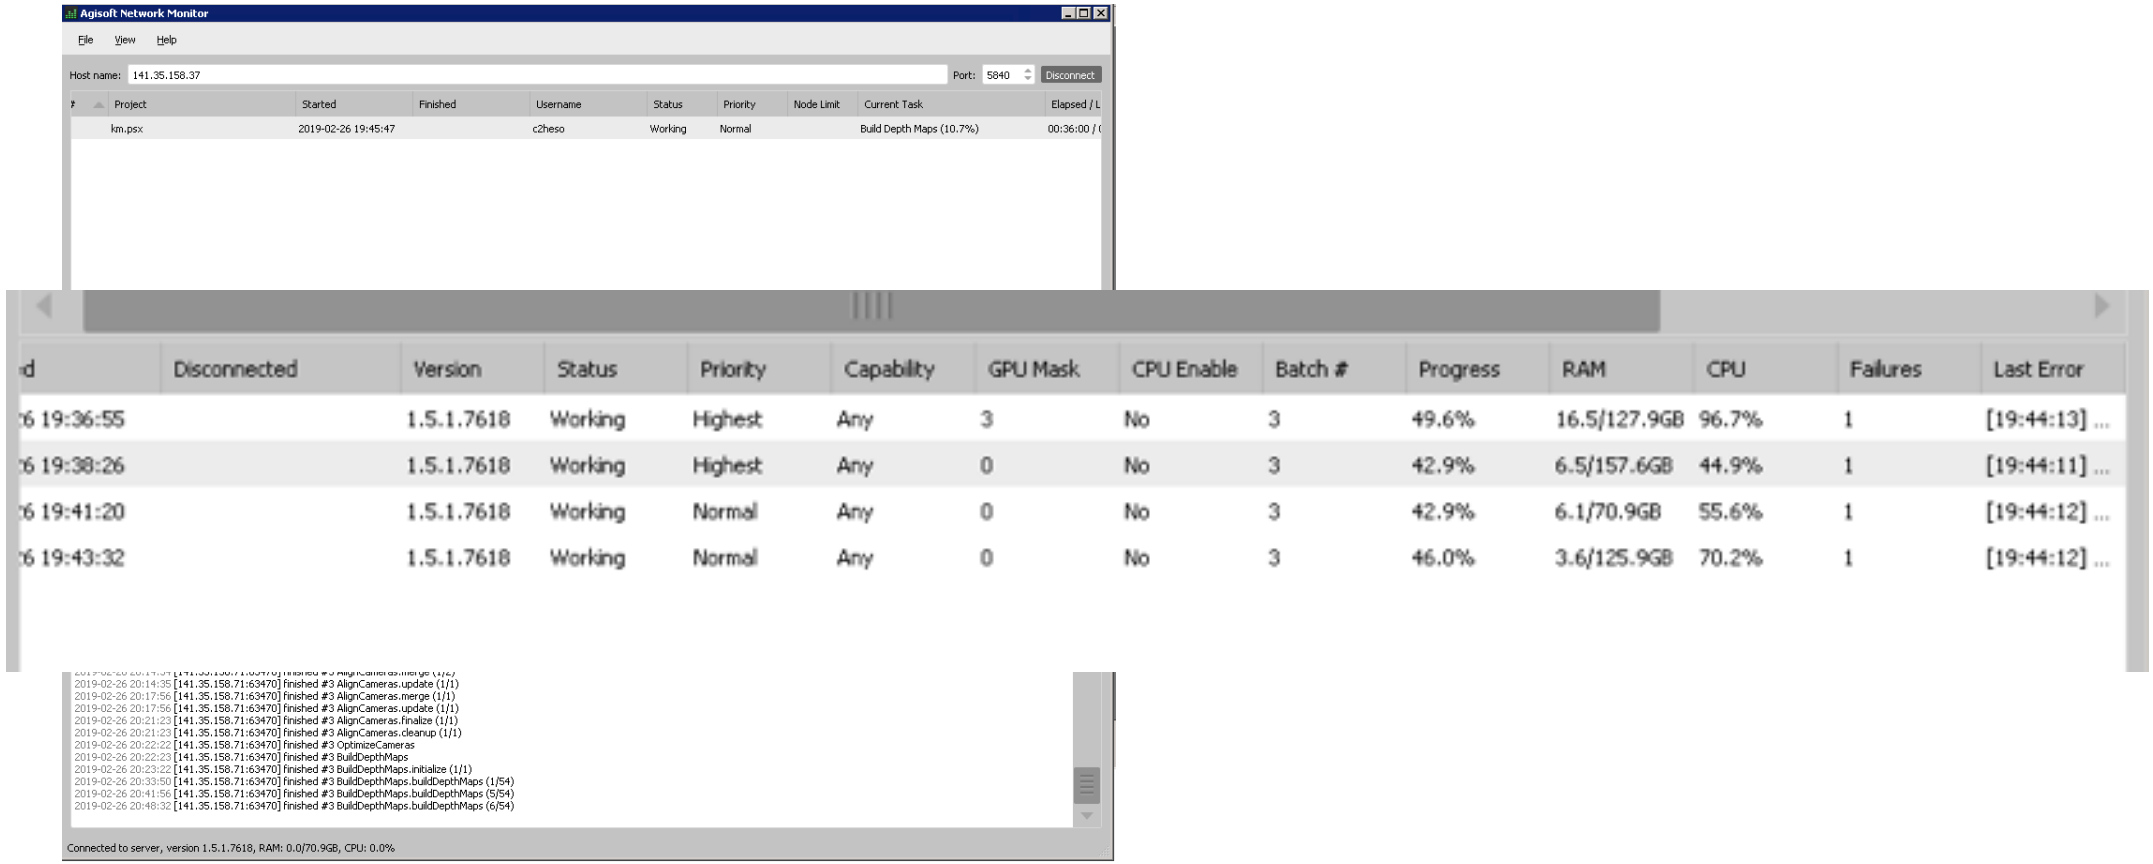

Agisoft – The Client

1. Client sends the processing job – the psx project tasks (via Metashapepro GUI or Python to the coordination server)
2. Prereq.: same Version, same working dir,

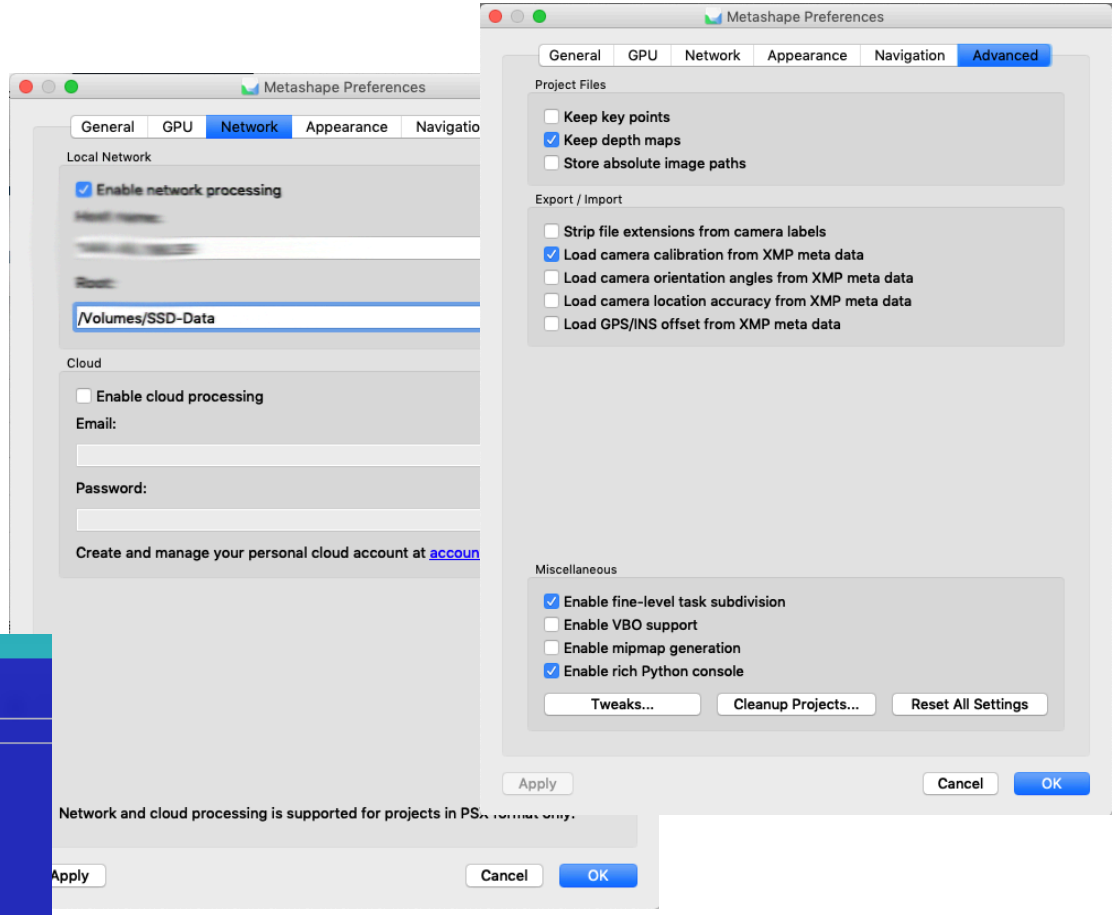

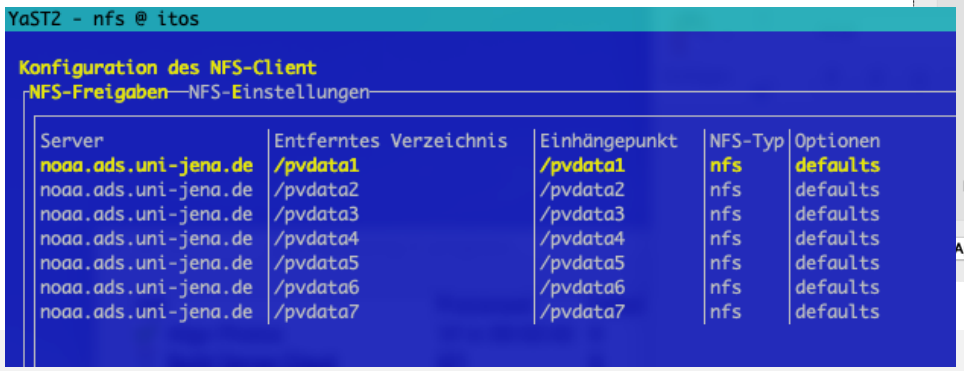

ARA Cluster

ARA Cluster (Mathematics and Computer Science Dep. - Prof. Martin Bücker) is likely the biggest cluster environment at FSU Jena with over 300 compute nodes

Tesla V100-PCIE-16GB, 80 compute units, free memory: 15703/16130 MB, compute capability 7.0, driver version: 410.79, driver/runtime CUDA: 10000/5050

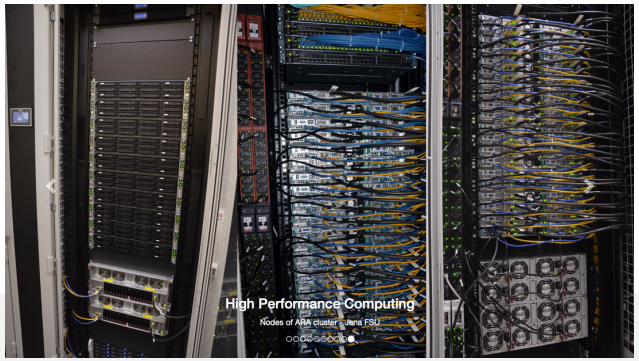


```
/home
    User's Home Directories
    NFS Filesystem
    Space available: 80 TB
    globally  accessible
```

```
/beegfs
    Scratch file system
    BeeGFS Parallel file system
    Space available: 524 TB
    2 metadata servers and 6 object storage servers,
    connected to all computer nodes through Intel Omni-Path Interconnect (HSCI)
```

```
/local
    Node-local temporary user data
    Space Available: 169 GB
    Job duration only, Files should be deleted by user job script at the end of a job
```

ARA Cluster

ARA Cluster – node configuration, compute nodes are grouped to form “Partitions”.

Partition	Nodes	Node Category	Shared	MaxTime	MaxNodesPerUser	MaxSubmitJobsPerUser
b_test	node001	Broadwell Standard	no	3:00:00		
b_standard	node[002-126,131-136]	Broadwell Standard	yes	8-08:00:00		
b_fat	node[137-140]	Broadwell Fat	yes	8-08:00:00		
s_test	node141	Skylake Standard	no	3:00:00		
s_standard	node[142-268,293-316]]	Skylake Standard	yes	8-08:00:00		
s_fat	node[269-272]	Skylake Fat	yes	8-08:00:00		
gpu_test	node127	2 NVIDIA Tesla P100	yes	01:00:00		
gpu_p100	node[128-129]	2 NVIDIA Tesla P100	yes	8-08:00:00		
gpu_v100	node130	2 NVIDIA Tesla V100	yes	8-08:00:00		
all	node[001-273]		yes	8-08:00:00		

ARA Cluster

ARA Cluster – node configuration, compute nodes are grouped to form “Partitions”.

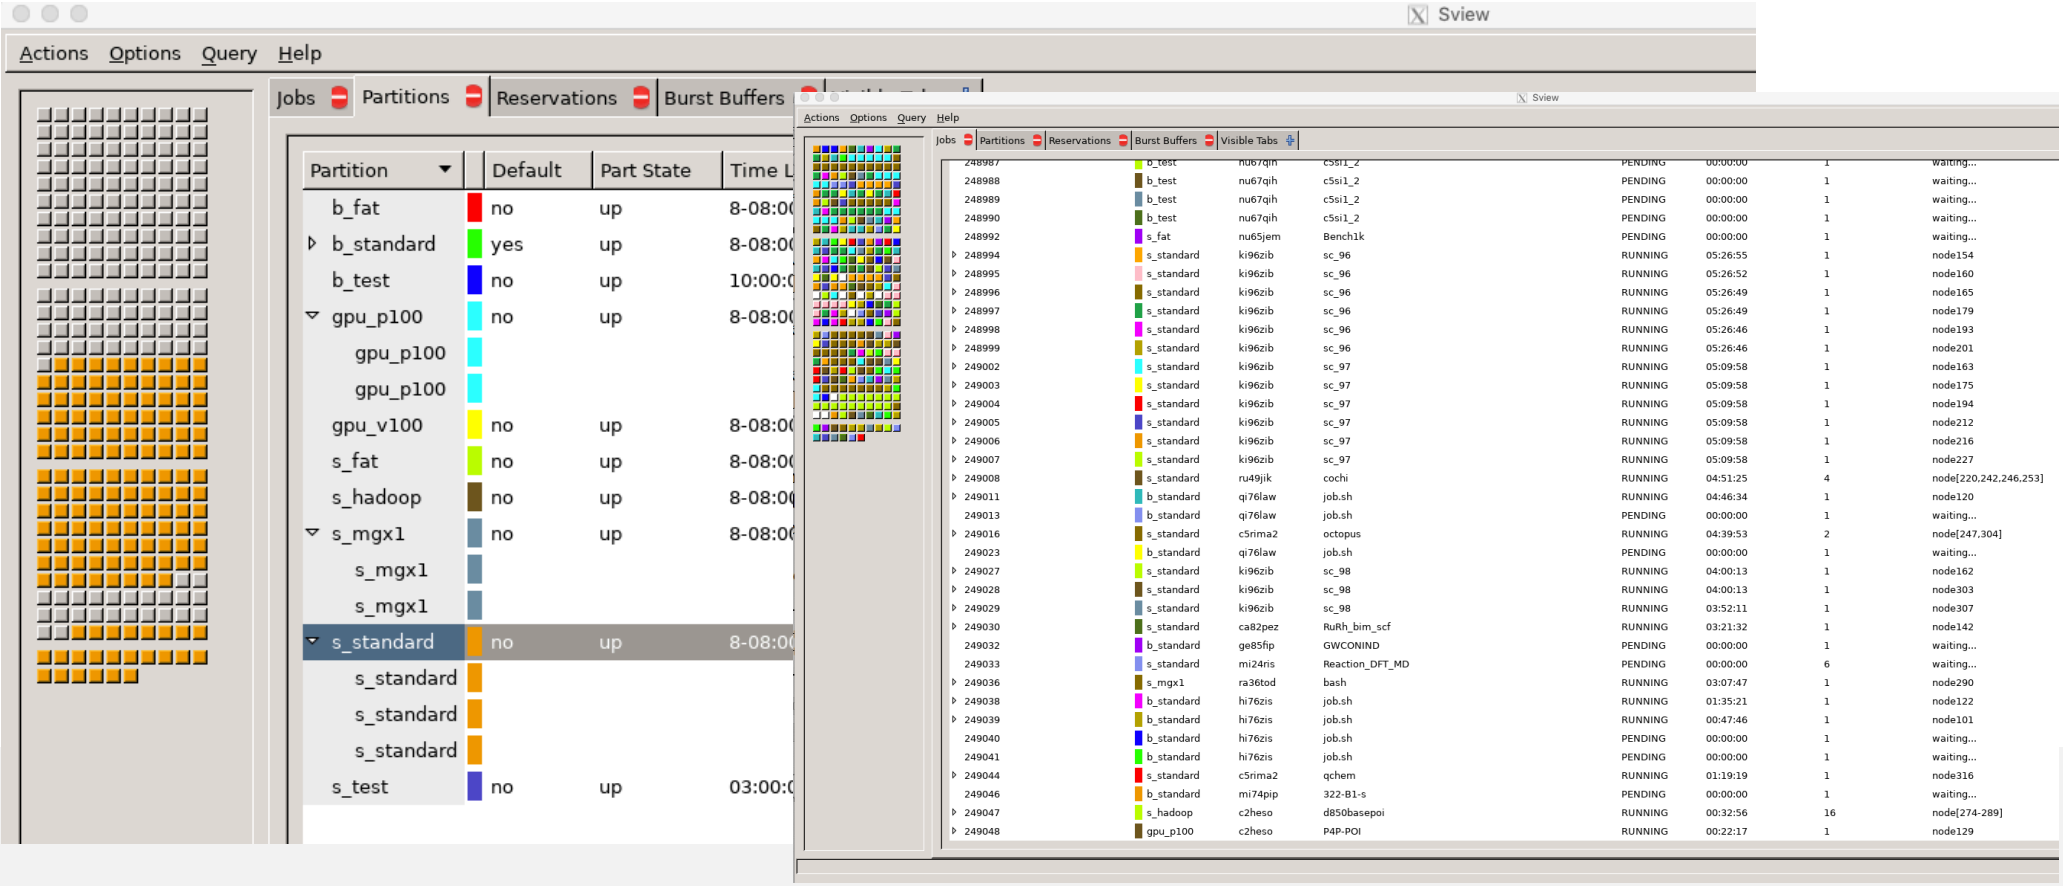

ARA Cluster

[Slurm Workload Manager - Overview](https://slurm.schedmd.com/overview.html)
https://slurm.schedmd.com/overview.html
https://ara-wiki.rz.uni-jena.de/Hauptseite#Job_parameters
Jobs on ARA run, by default, in shared-mode, meaning that multiple jobs can run on the same compute nodes. Sometimes, this behaviour is not desired (e.g. for benchmarking purposes), in which case it can be turned off by specifying the SLURM parameter: –exclusive.

SLURM-Batch.sh — Processi

```
#!/bin/sh
#SBATCH --nodes=1
#SBATCH --job-name=hainichmed1118
#SBATCH --partition=gpu_v100
#SBATCH --gres=gpu:2
#SBATCH module load apps/metashape/1.5.1
#SBATCH module load tools/python/3.6
#SBATCH module load nvidia/cuda/9
#SBATCH --begin=now+10
# time $(i=0; while (( i < 9999999 )); do (( i ++ )); done)
export DISPLAY=localhost:12.0 && /cluster/apps/metashape/1.5.1/metashape.sh -platform offscreen -r metashape.py
```

Agisoft Processing Tweaks

#agisoft Processing Tweaks:

- main/dense_cloud_max_neighbors = 60 reducing the amount of overlapping images in dense point cloud processing
- reduction of image number from POI flights
- area/volume reduction in GUI
- Increase RAM

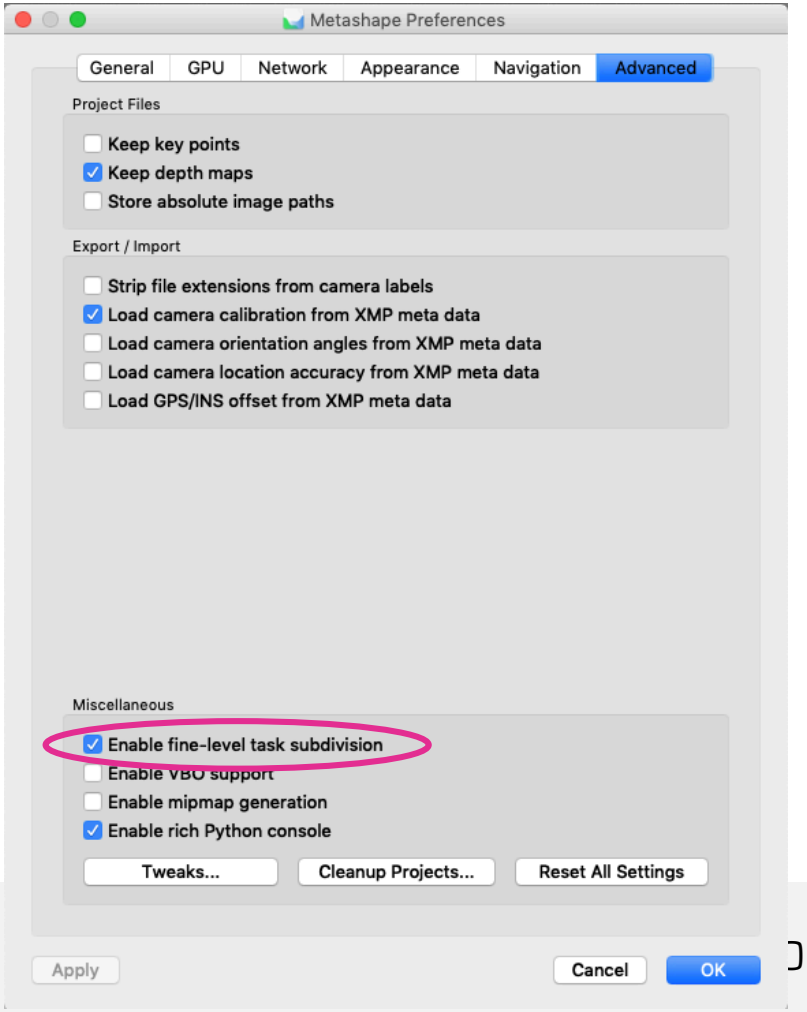

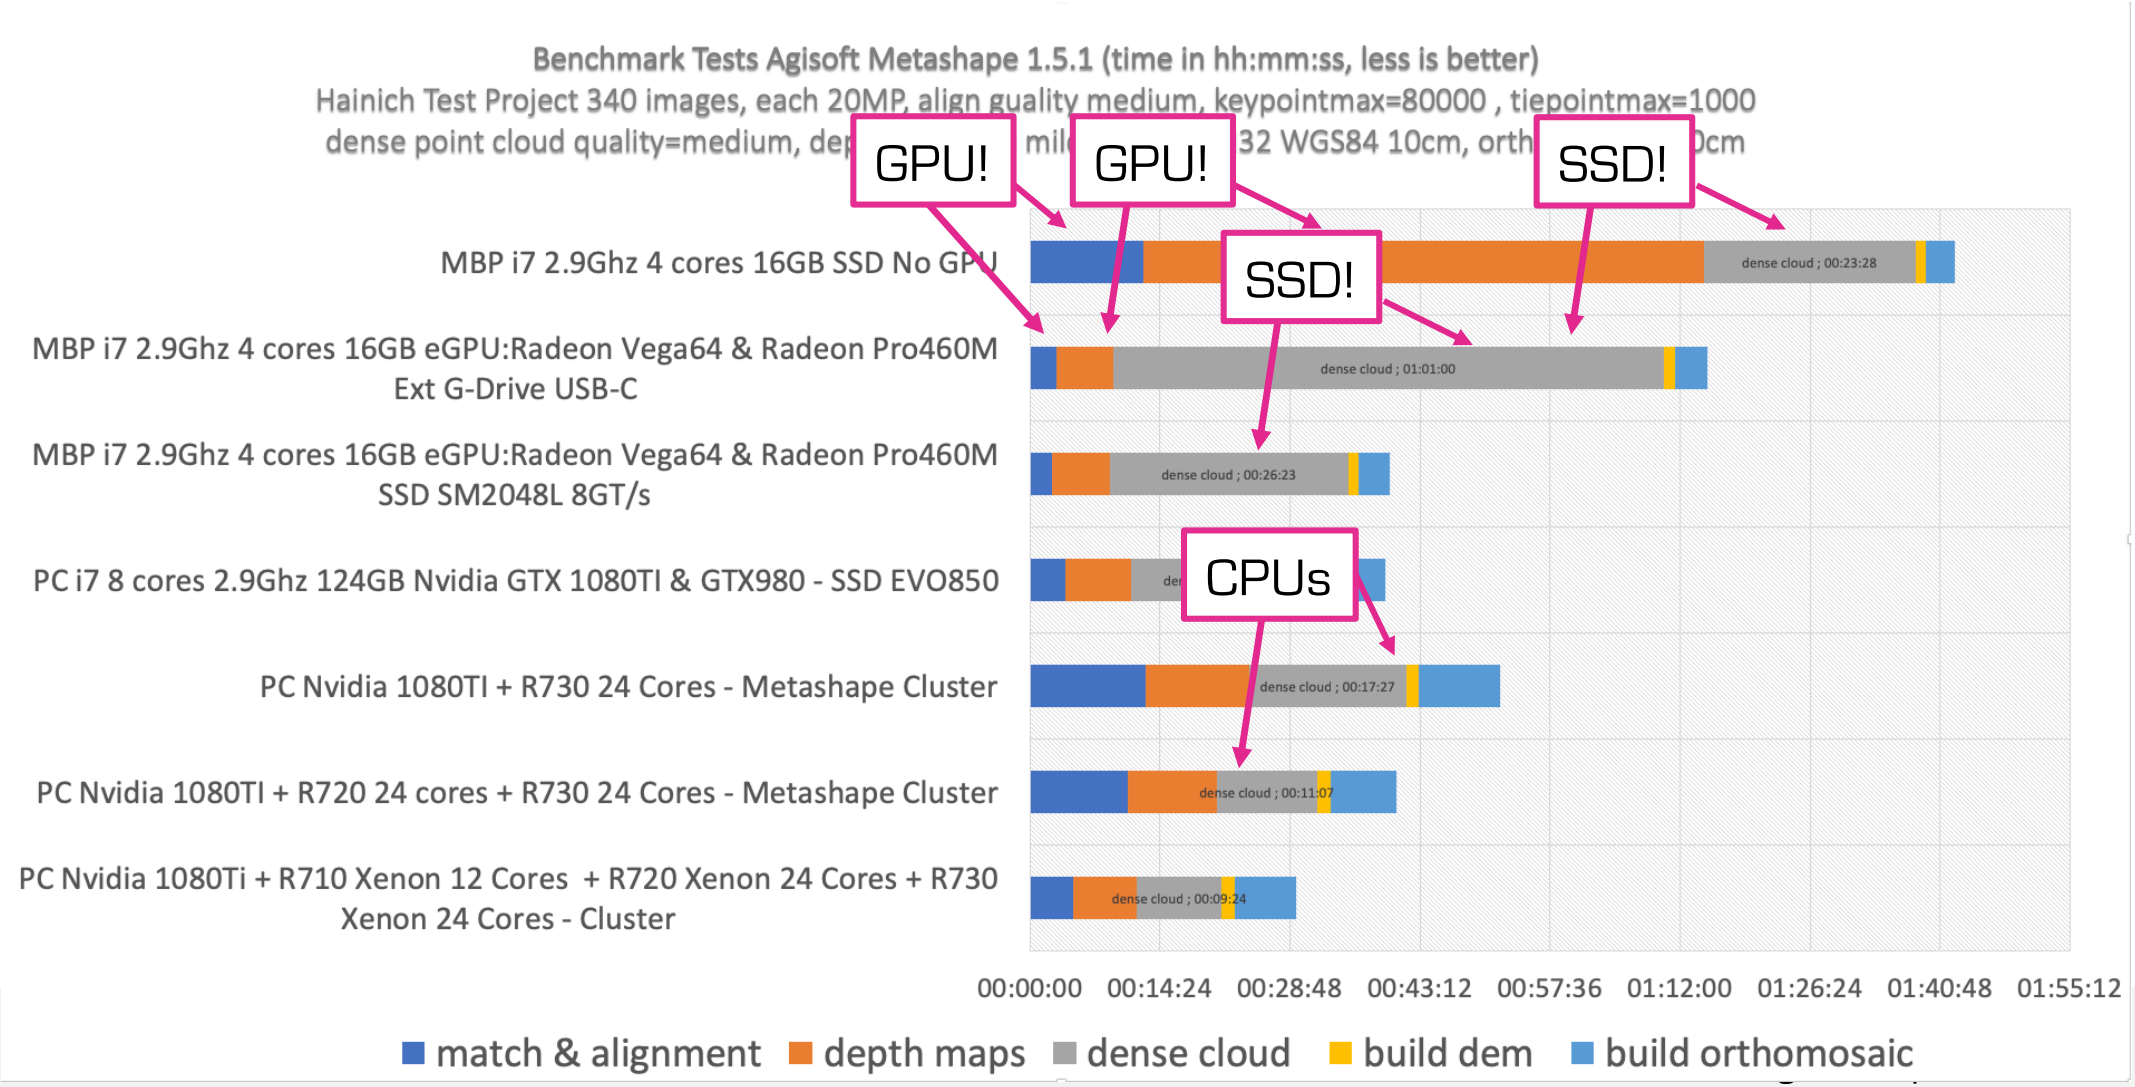
Benchmark Tests Agisoft Metashape 1.5.1 (time in hh:mm:ss, less is better)
Hainich Test Project 340 images, each 20MP, align guality medium, keypointmax=80000 , tiepointmax=1000
dense point cloud quality=medium, dep mil 32 WGS84 10cm, orth 0cm
GPU!
GPU!
SSD!
SSD!
CPUs
MBP i7 2.9Ghz 4 cores 16GB SSD No GPU
MBP i7 2.9Ghz 4 cores 16GB eGPU:Radeon Vega64 & Radeon Pro460M Ext G-Drive USB-C
MBP i7 2.9Ghz 4 cores 16GB eGPU:Radeon Vega64 & Radeon Pro460M SSD SM2048L 8GT/s
PC i7 8 cores 2.9Ghz 124GB Nvidia GTX 1080TI & GTX980 - SSD EVO850
PC Nvidia 1080TI + R730 24 Cores - Metashape Cluster
PC Nvidia 1080TI + R720 24 cores + R730 24 Cores - Metashape Cluster
PC Nvidia 1080Ti + R710 Xenon 12 Cores + R720 Xenon 24 Cores + R730 Xenon 24 Cores - Cluster
dense cloud ; 00:23:28
dense cloud ; 01:01:00
dense cloud ; 00:26:23
dense cloud ; 00:17:27
dense cloud ; 00:11:07
dense cloud ; 00:09:24
00:00:00
00:14:24
00:28:48
00:43:12
00:57:36
01:12:00
01:26:24
01:40:48
01:55:12
match & alignment
depth maps
dense cloud
build dem
build orthomosaic

Overall Results - sub divided

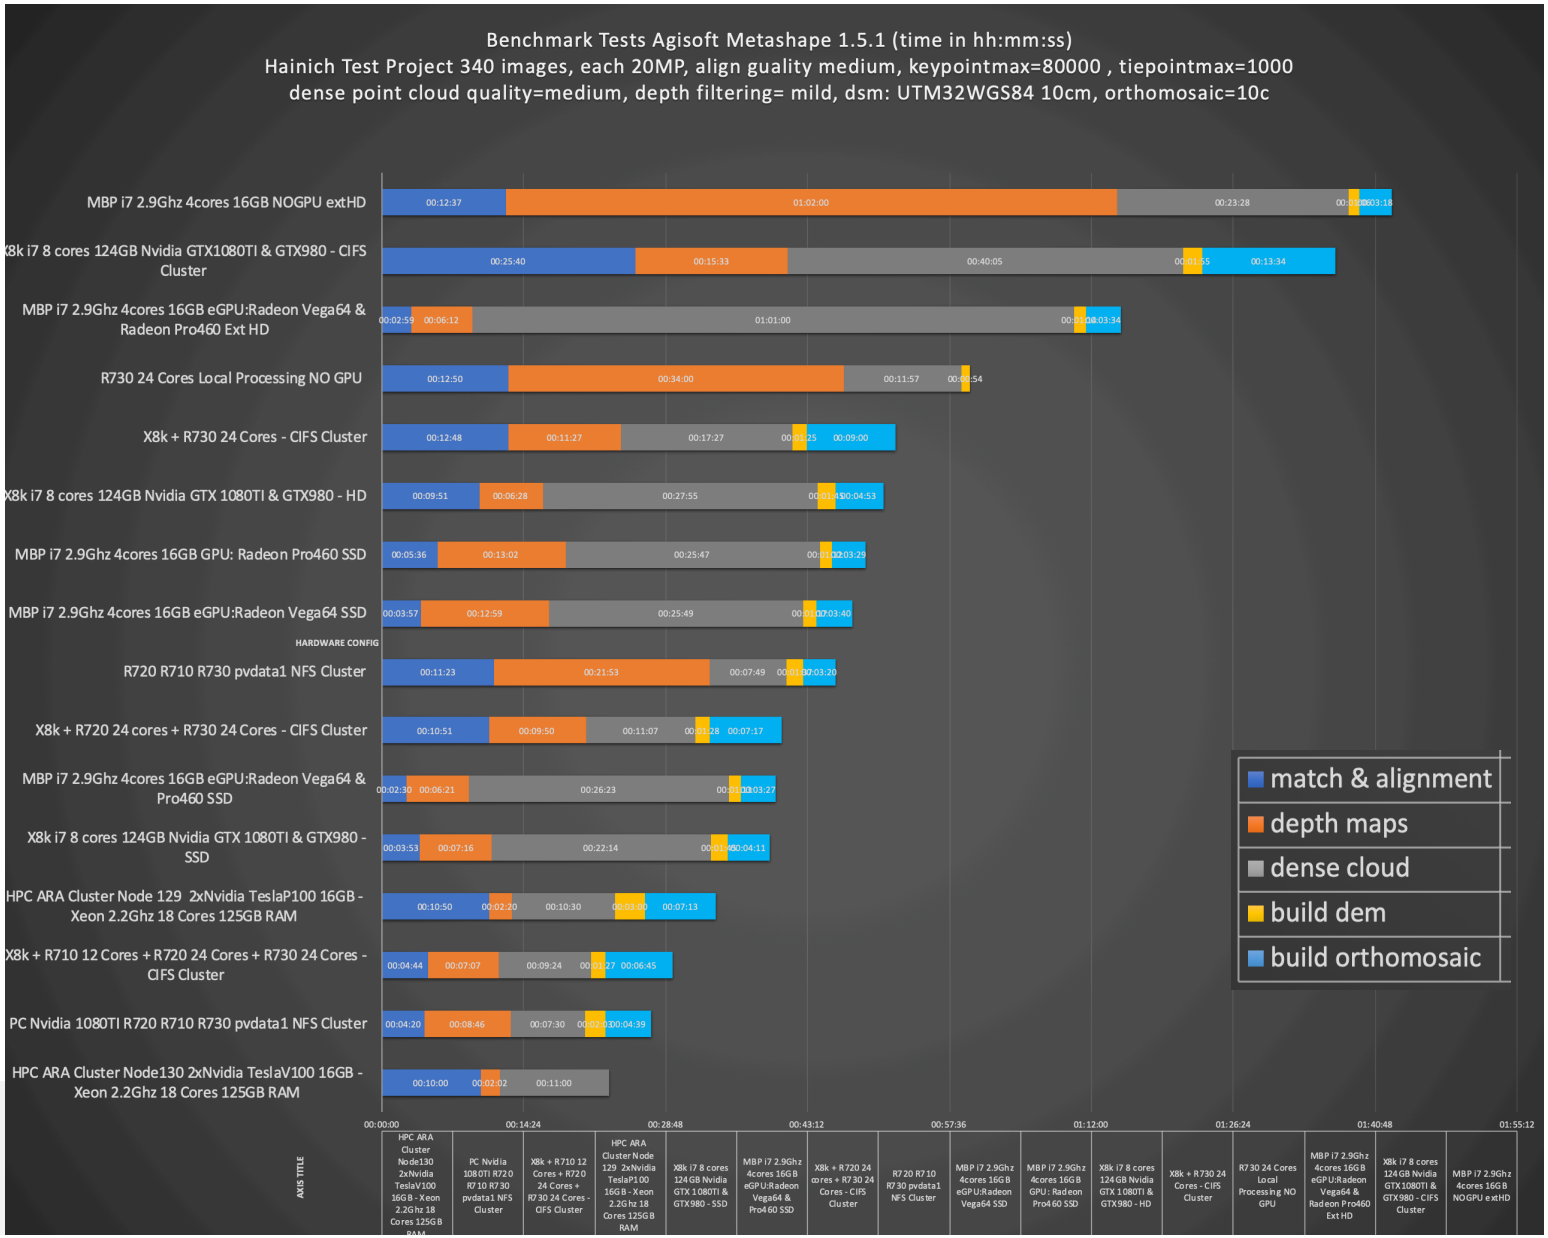

Match & Align

CIFS

SSD

SSD

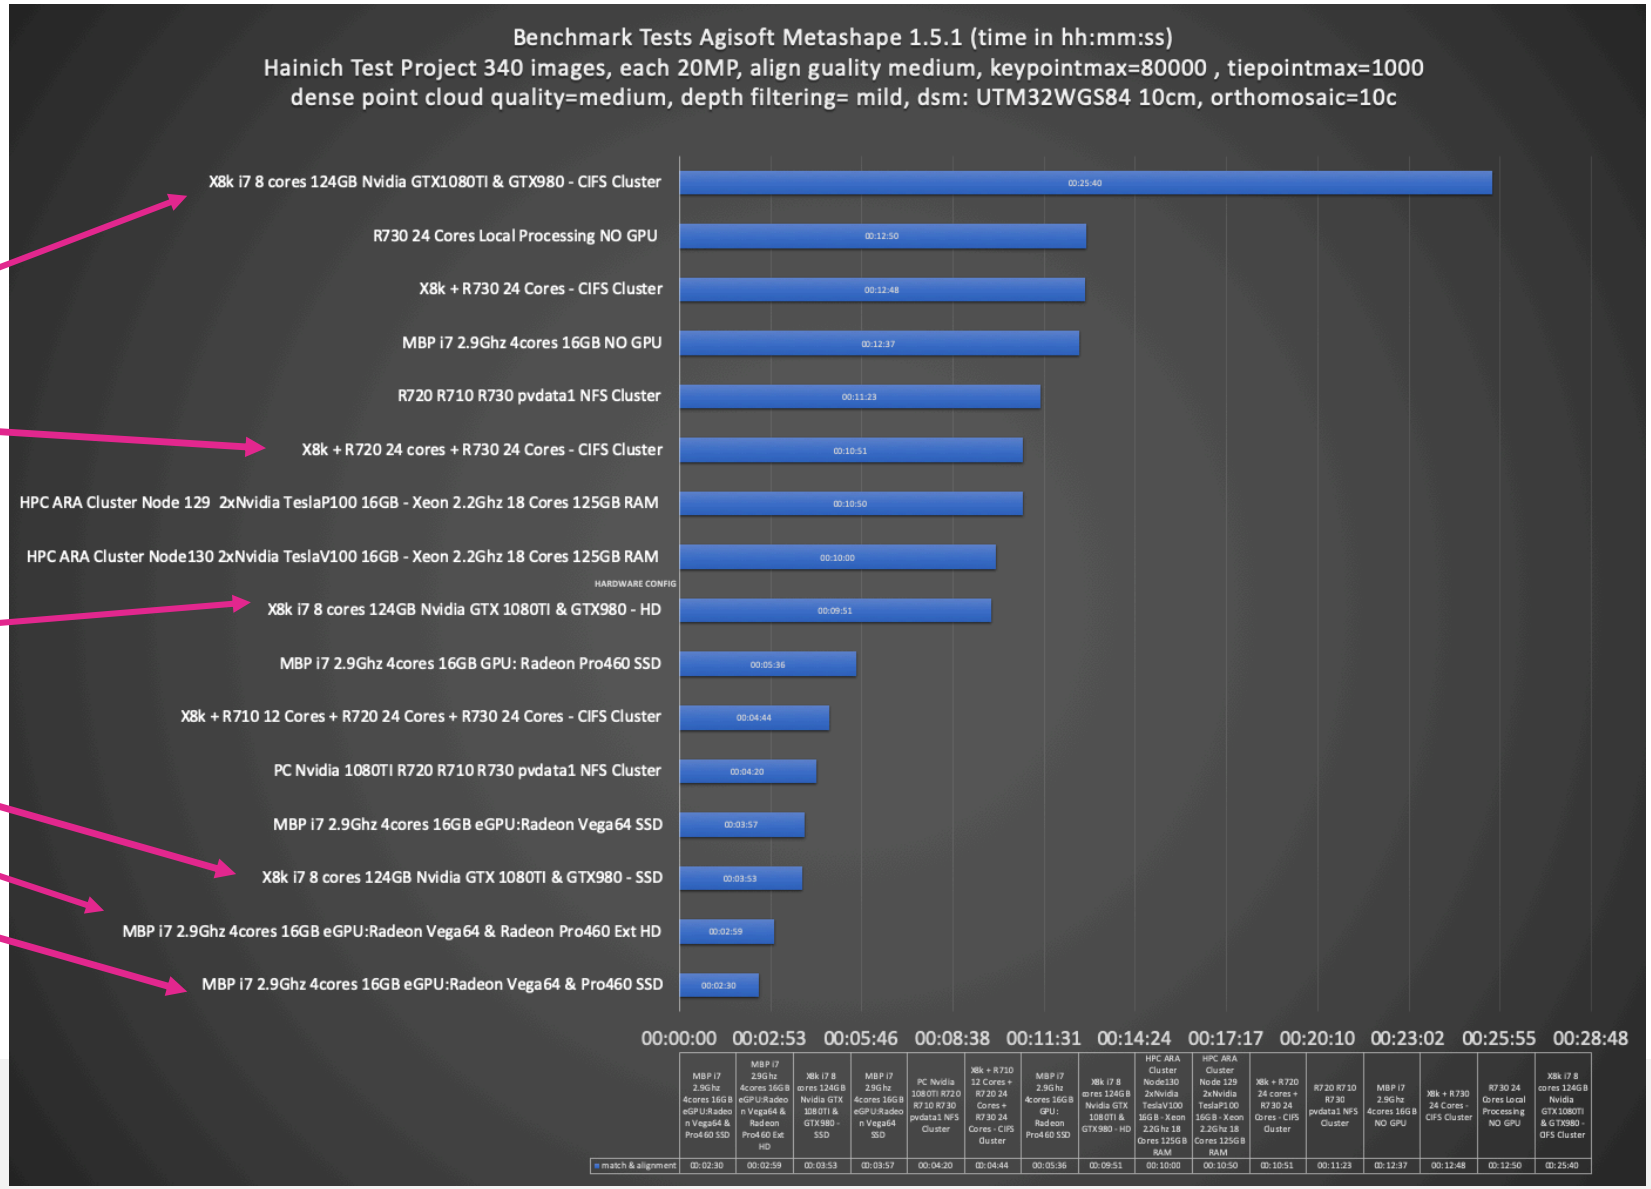

Depth Maps (Advantage Tesla)

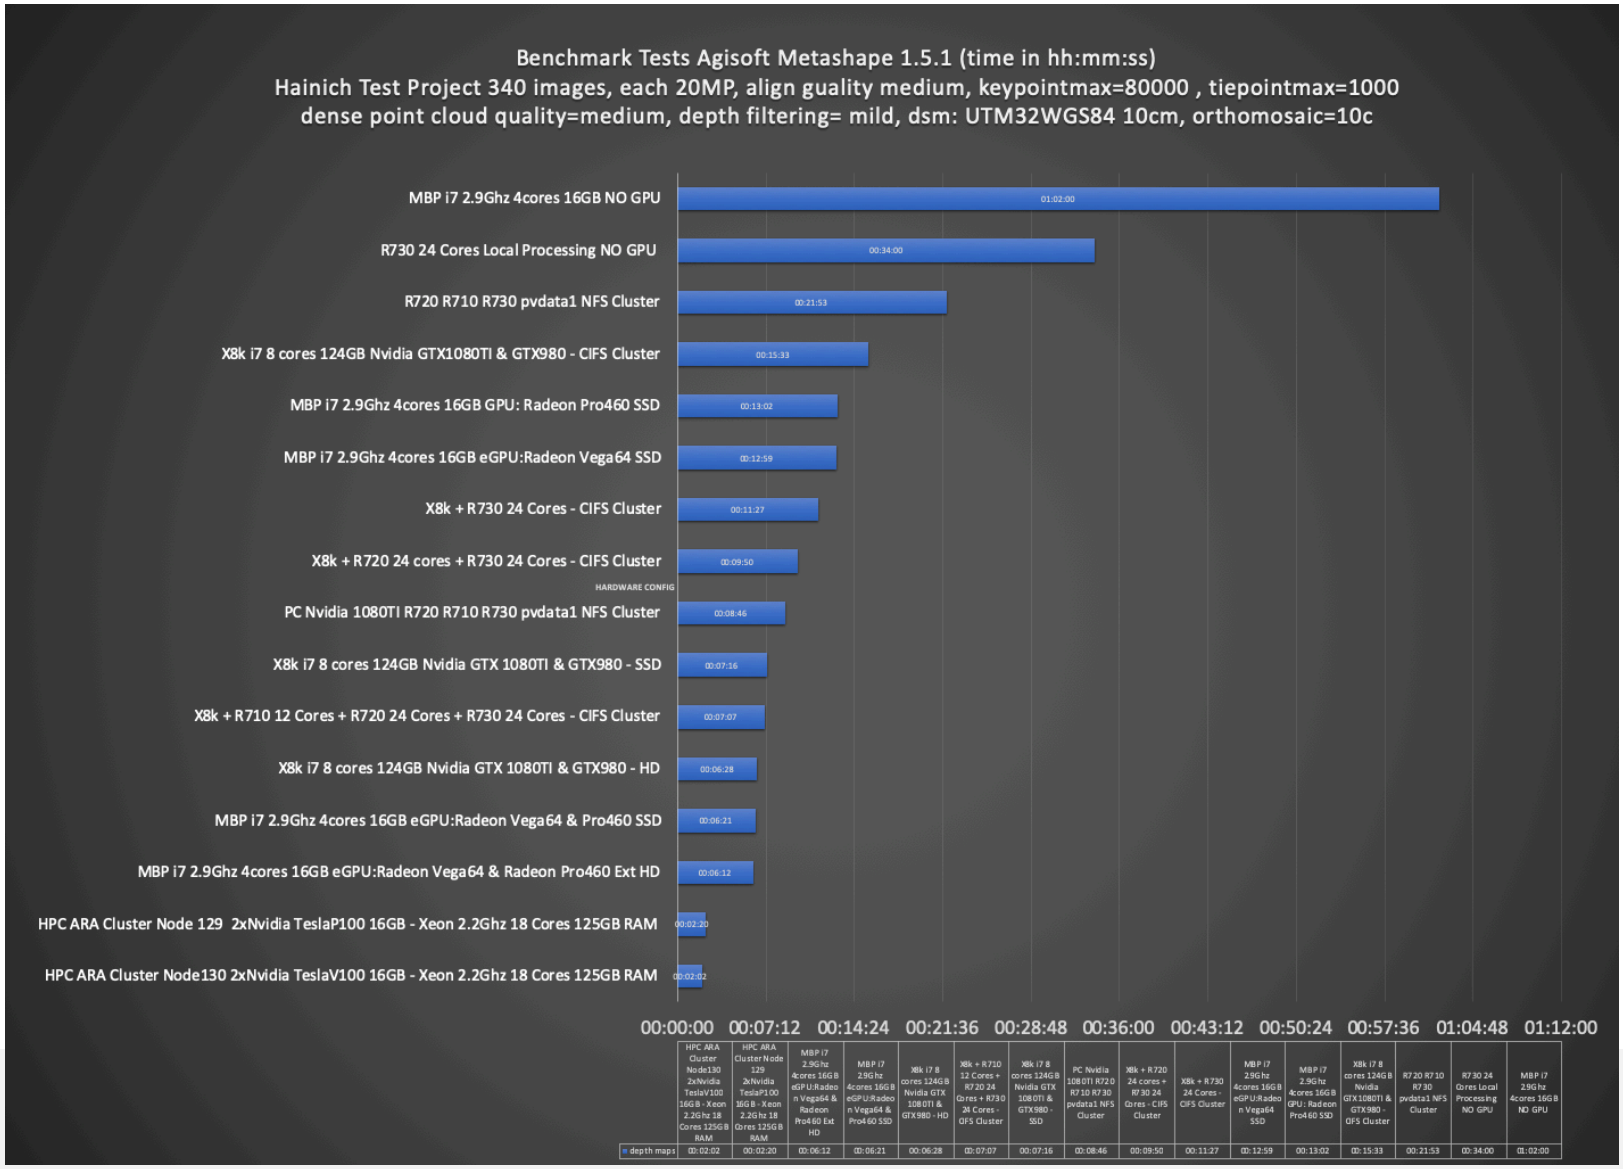

Dense Cloud

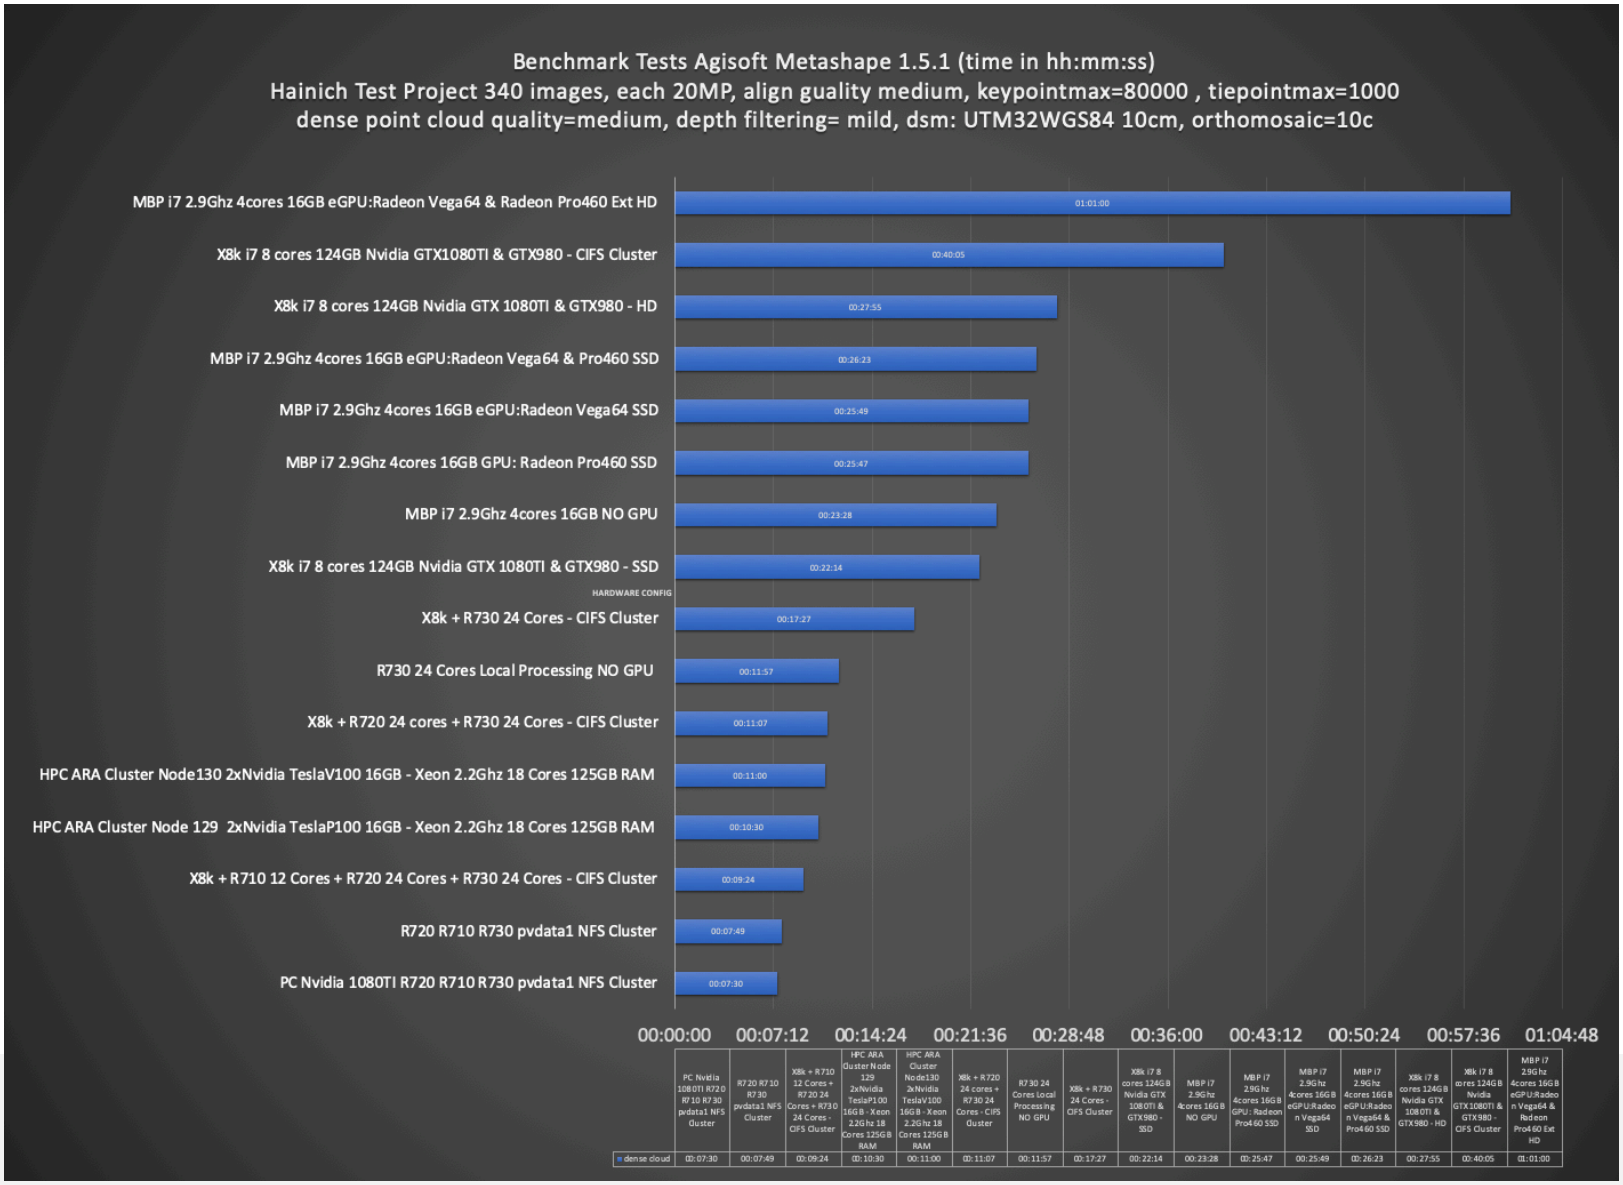
Benchmark Tests Agisoft Metashape 1.5.1 (time in hh:mm:ss)
Hainich Test Project 340 images, each 20MP, align guality medium, keypointmax=80000 , tiepointmax=1000
dense point cloud quality=medium, depth filtering= mild, dsm: UTM32WGS84 10cm, orthomosaic=10c
MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 & Radeon Pro460 Ext HD
01:01:00
X8k i7 8 cores 124GB Nvidia GTX1080TI & GTX980 - CIFS Cluster
00:40:05
X8k i7 8 cores 124GB Nvidia GTX 1080TI & GTX980 - HD
00:27:55
MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 & Pro460 SSD
00:26:23
MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 SSD
00:25:49
MBP i7 2.9Ghz 4cores 16GB GPU: Radeon Pro460 SSD
00:25:47
MBP i7 2.9Ghz 4cores 16GB NO GPU
00:23:28
X8k i7 8 cores 124GB Nvidia GTX 1080TI & GTX980 - SSD
00:22:14
HARDWARE CONFIG
X8k + R730 24 Cores - CIFS Cluster
00:17:27
R730 24 Cores Local Processing NO GPU
00:11:57
X8k + R720 24 cores + R730 24 Cores - CIFS Cluster
00:11:07
HPC ARA Cluster Node130 2xNvidia TeslaV100 16GB - Xeon 2.2Ghz 18 Cores 125GB RAM
00:11:00
HPC ARA Cluster Node 129 2xNvidia TeslaP100 16GB - Xeon 2.2Ghz 18 Cores 125GB RAM
00:10:30
X8k + R710 12 Cores + R720 24 Cores + R730 24 Cores - CIFS Cluster
00:09:24
R720 R710 R730 pvdata1 NFS Cluster
00:07:49
PC Nvidia 1080TI R720 R710 R730 pvdata1 NFS Cluster
00:07:30
00:00:00 00:07:12 00:14:24 00:21:36 00:28:48 00:36:00 00:43:12 00:50:24 00:57:36 01:04:48
PC Nvidia 1080TI R720 R710 R730 pvdata1 NFS Cluster
R720 R710 R730 pvdata1 NFS Cluster
X8k + R710 12 Cores + R720 24 Cores + R730 24 Cores - CIFS Cluster
HPC ARA Cluster Node 129 2xNvidia TeslaP100 16GB - Xeon 2.2Ghz 18 Cores 125GB RAM
HPC ARA Cluster Node130 2xNvidia TeslaV100 16GB - Xeon 2.2Ghz 18 Cores 125GB RAM
X8k + R720 24 cores + R730 24 Cores - CIFS Cluster
R730 24 Cores Local Processing NO GPU
X8k + R730 24 Cores - CIFS Cluster
X8k i7 8 cores 124GB Nvidia GTX 1080TI & GTX980 - SSD
MBP i7 2.9Ghz 4cores 16GB NO GPU
MBP i7 2.9Ghz 4cores 16GB GPU: Radeon Pro460 SSD
MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 SSD
MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 & Pro460 SSD
X8k i7 8 cores 124GB Nvidia GTX 1080TI & GTX980 - HD
X8k i7 8 cores 124GB Nvidia GTX1080TI & GTX980 - CIFS Cluster
MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 & Radeon Pro460 Ext HD
dense cloud
00:07:30 00:07:49 00:09:24 00:10:30 00:11:00 00:11:07 00:11:57 00:17:27 00:22:14 00:23:28 00:25:47 00:25:49 00:26:23 00:27:55 00:40:05 01:01:00

Overall

(only applies to low res setup - would look very different with quality high/ultra!)

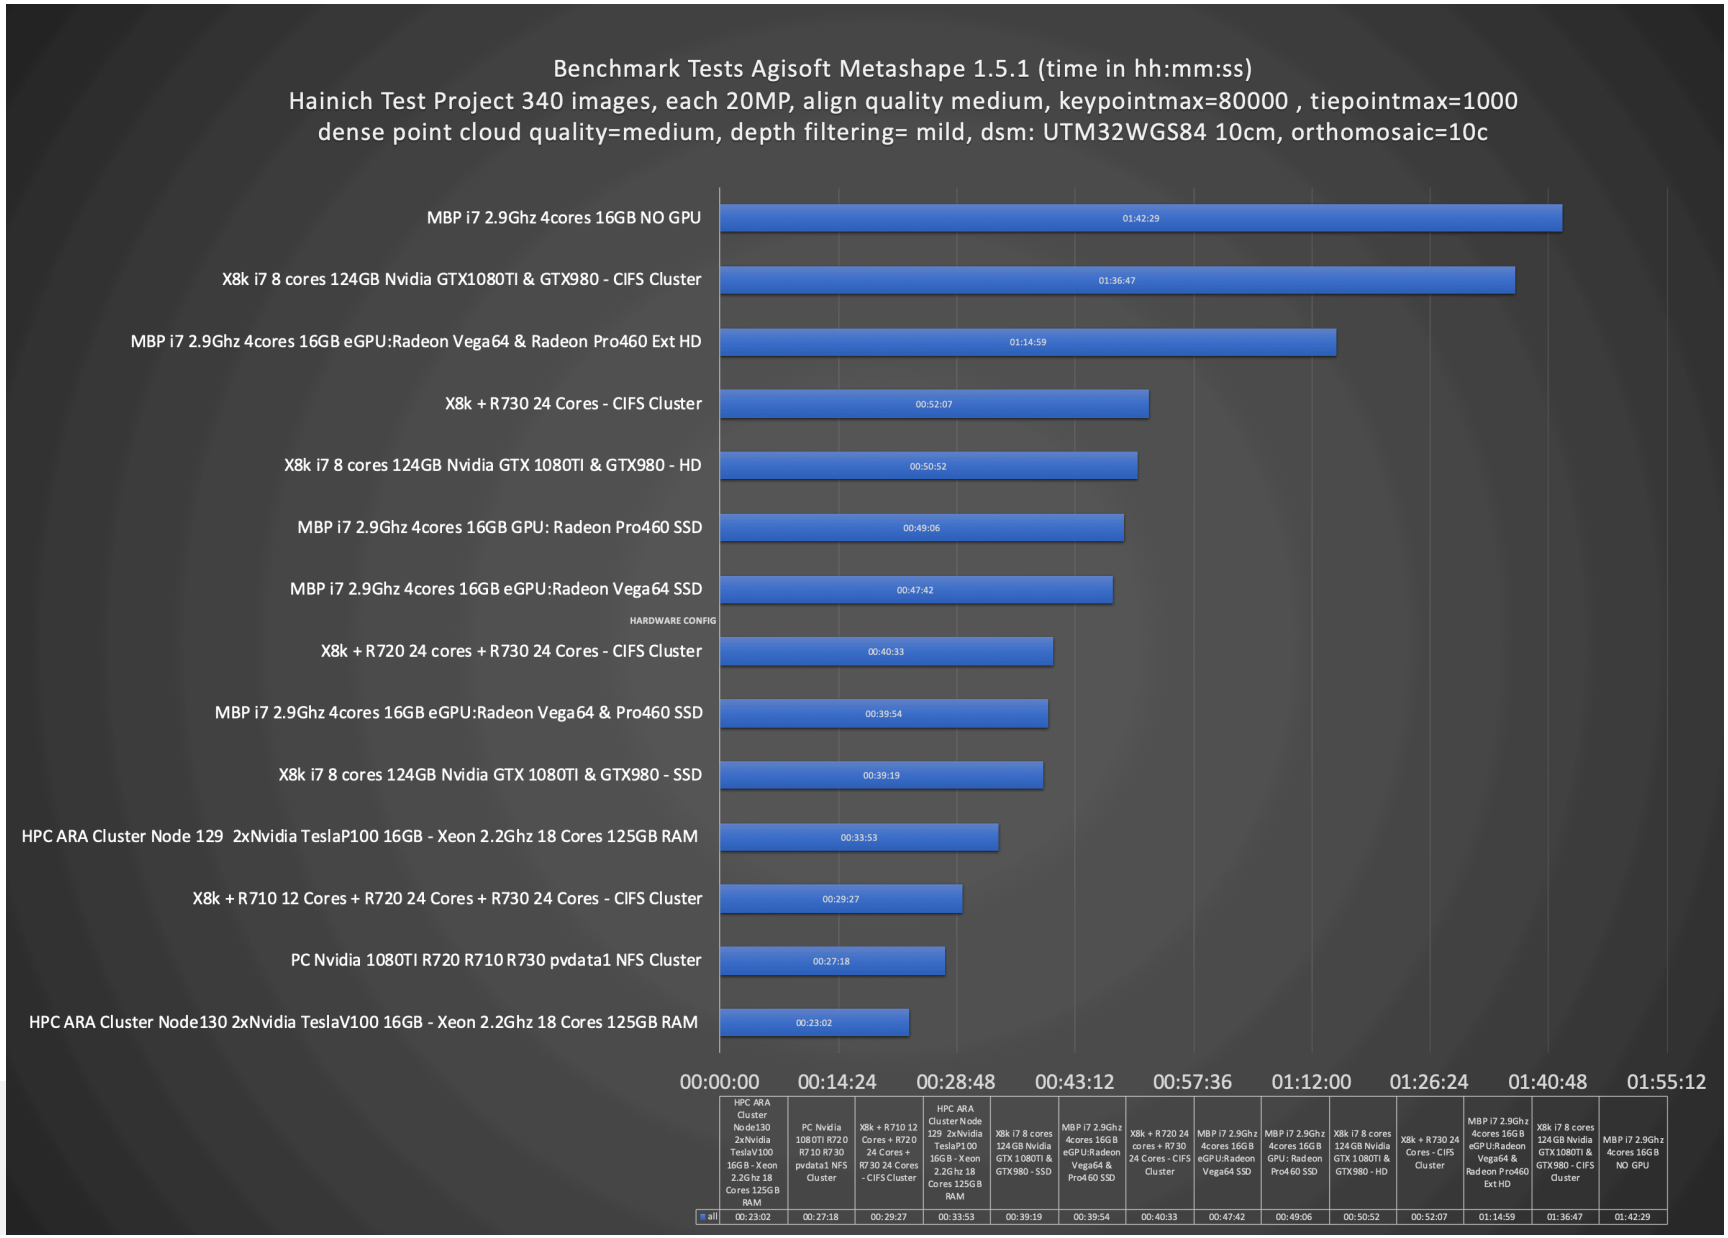

	HPC ARA Cluster Node130 2xNvidia TeslaV100 16GB - Xeon 2.2Ghz 18 Cores 125GB RAM	PC Nvidia 1080TI R720 R710 R730 pvdata1 NFS Cluster	X8k + R710 12 Cores + R720 24 Cores + R730 24 Cores - CIFS Cluster	HPC ARA Cluster Node 129 2xNvidia TeslaP100 16GB - Xeon 2.2Ghz 18 Cores 125GB RAM	X8k i7 8 cores 124GB Nvidia GTX 1080TI & GTX980 - SSD	MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 & Pro460 SSD	X8k + R720 24 cores + R730 24 Cores - CIFS Cluster	MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 SSD	MBP i7 2.9Ghz 4cores 16GB GPU: Radeon Pro460 SSD	X8k i7 8 cores 124GB Nvidia GTX 1080TI & GTX980 - HD	X8k + R730 24 Cores - CIFS Cluster	MBP i7 2.9Ghz 4cores 16GB eGPU:Radeon Vega64 & Radeon Pro460 Ext HD	X8k i7 8 cores 124GB Nvidia GTX1080TI & GTX980 - CIFS Cluster	MBP i7 2.9Ghz 4cores 16GB NO GPU
all	00:23:02	00:27:18	00:29:27	00:33:53	00:39:19	00:39:54	00:40:33	00:47:42	00:49:06	00:50:52	00:52:07	01:14:59	01:36:47	01:42:29

Overall

- NVIDIA Volta architecture disappointing (AI / ML optimized)
- Strong GPU performance however helps some processes (align/match/depth maps)
- Speedy SSDs seem to accelerate the full workflow
- Good results with AMD Vega64 card
- Network speed (cifs/nfs file systems) seem to be bottle neck for cluster processing
- Dense point cloud calculation: workflow bottle neck (high and ultra settings – full res settings
- Good MBP results only valid for very small projects! (16GB RAM limit (as of 2019))

More Agisoft Benchmarking

https://www.pugetsystems.com/labs/articles/Agisoft-PhotoScan-Multi-Core-Performance-709/
https://www.agisoft.com/forum/index.php?topic=1868.0
https://www.anandtech.com/bench/CPU/1057

Copter POI Crown Point Cloud (DGPS registered) vers. TLS Crown Point Cloud (DGPS registered)
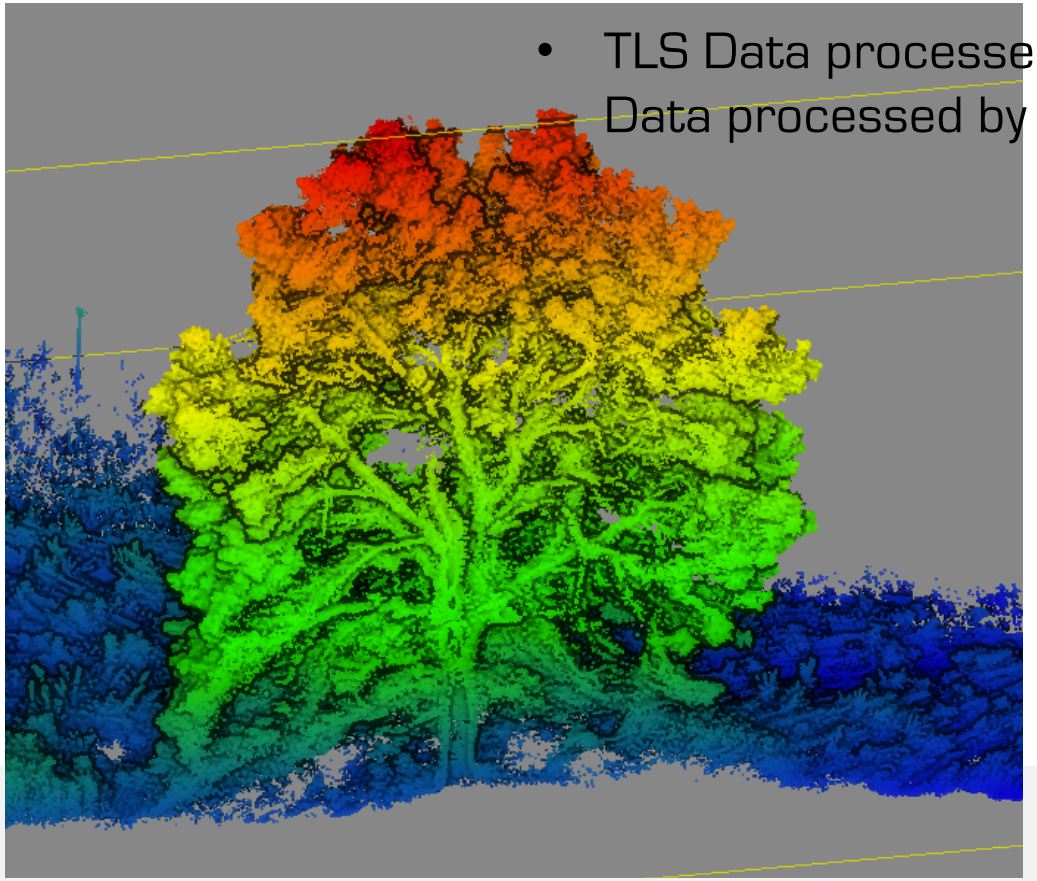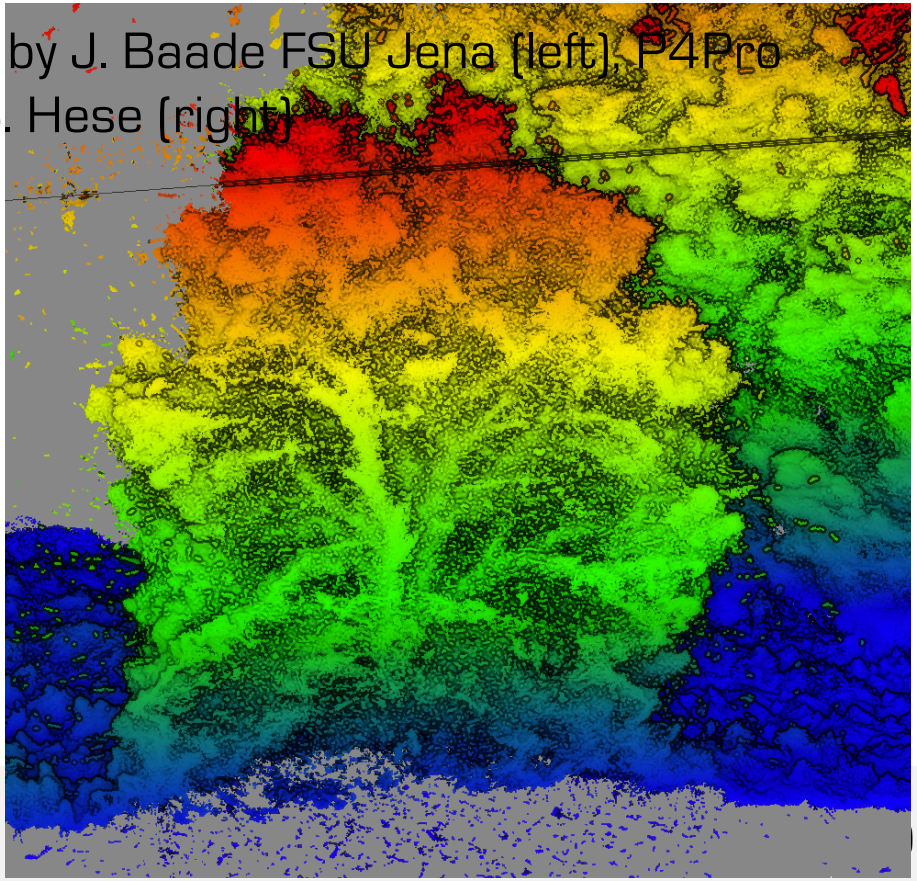
- TLS Data processed by J. Baade FSU Jena (left), P4Pro Data processed by S. Hese (right)