PotreeConverterMPI: An MPI implementation of PotreeConverter

Nauman Ahmed

Research Software Engineer

netherlands eScience center

netherlands eScience center

Pointcloud

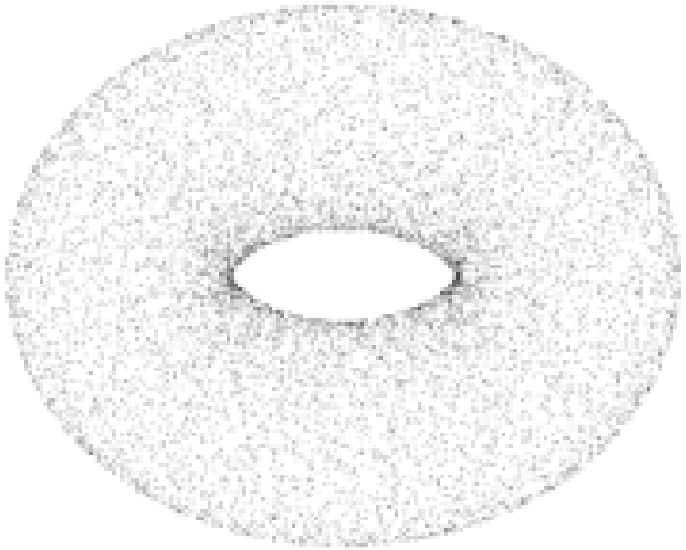

netherlands eScience center

Pointcloud format

Table 12: Point Data Record Format 3

Item	Format	Size	Required
X	long	4 bytes	*
Y	long	4 bytes	*
Z	long	4 bytes	*
Intensity	unsigned short	2 bytes	
Return Number	3 bits (bits 0, 1, 2)	3 bits	*
Number of Returns (given pulse)	3 bits (bits 3, 4, 5)	3 bits	*
Scan Direction Flag	1 bit (bit 6)	1 bit	*
Edge of Flight Line	1 bit (bit 7)	1 bit	*
Classification	unsigned char	1 byte	*
Scan Angle Rank (-90 to +90) – Left side	char	1 byte	*
User Data	unsigned char	1 byte	
Point Source ID	unsigned short	2 bytes	*
GPS Time	double	8 bytes	*
Red	unsigned short	2 bytes	*
Green	unsigned short	2 bytes	*
Blue	unsigned short	2 bytes	*

netherlands eScience center

Octree representation of pointcloud

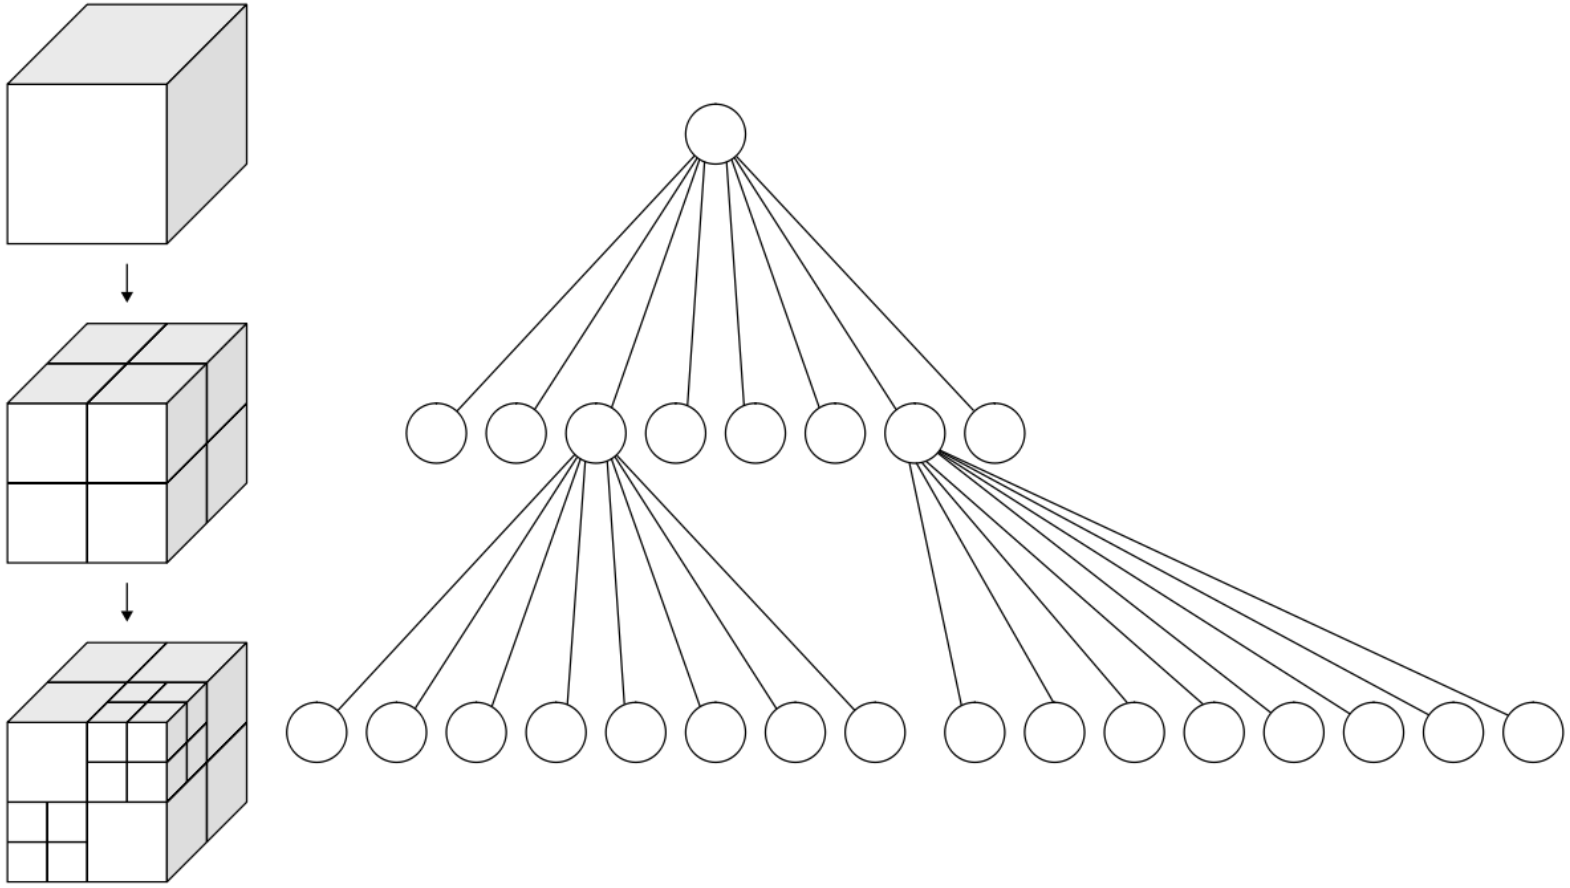

netherlands eScience center

AHN

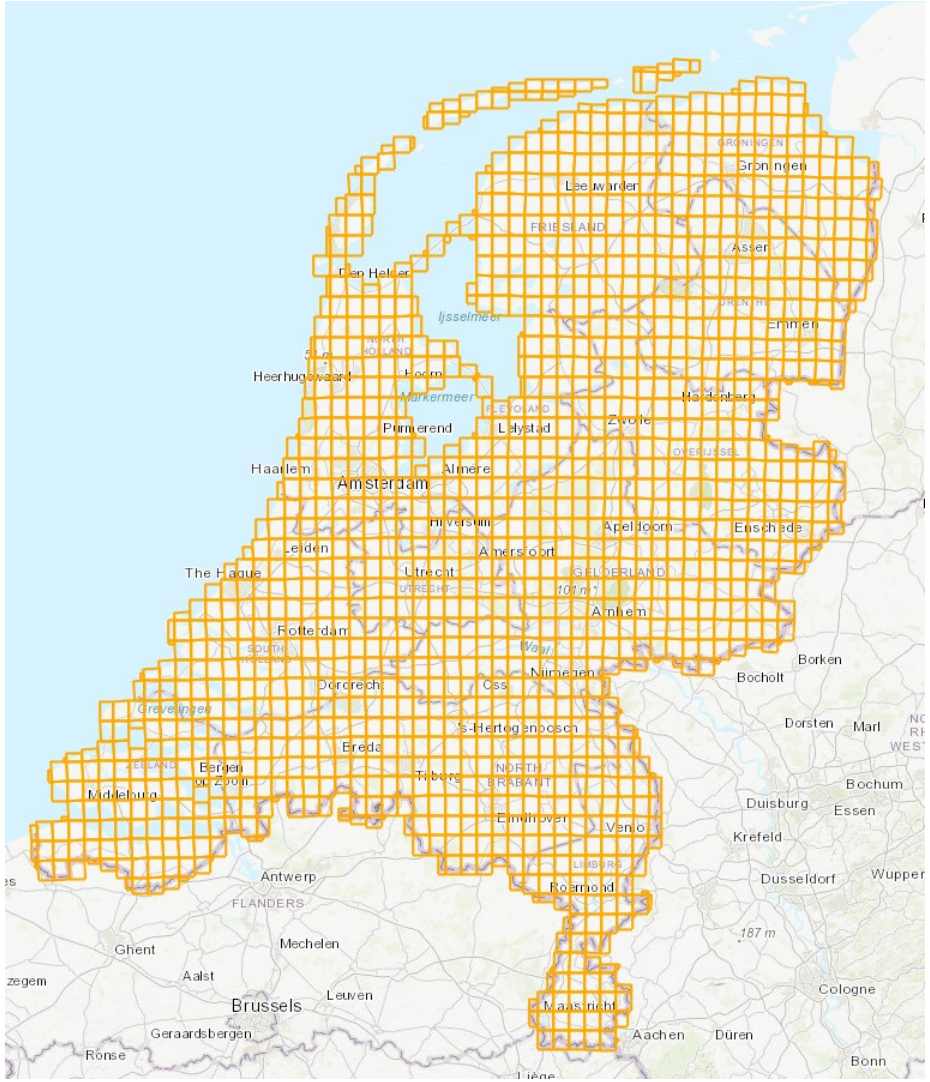

netherlands eScience center

New developments

*"**NUVIEW announced to launch 20 satellites with laser scanners to collect global elevation model (and repeat this every year),** aimed for accuracy USGS QL2 (2 pnts/m^2, 10 cm vertical accuracy). Ok, with total surface area of the Earth of about 509.600.000 km^2 of which area of land: 148.326.000 km2, that would be every year 2 x 148.326.000.000.000 = 296.652.000.000.000 points. **To compare AHN2 has about 640.000.000.000 points, So, this new dataset would be about 500 times larger."***

netherlands eScience center

Potree

LAZ/LAS

Converter

Octree

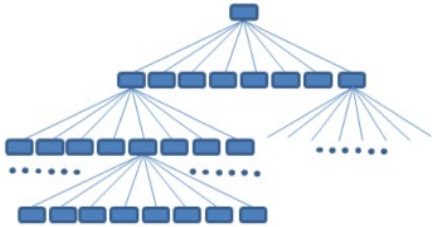

Viewer

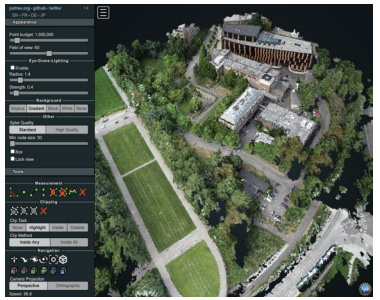

PotreeConverterMPI

Nauman Ahmed, Thijs van Lankveld, Martijn Meijers, Peter van Oosterom, Vitali Diaz, Edward Verbree

PotreeConverter on DelftBlue

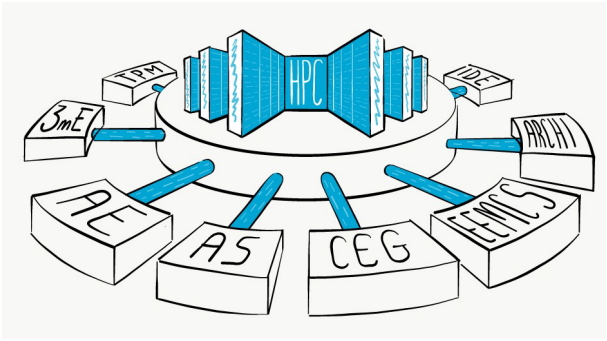

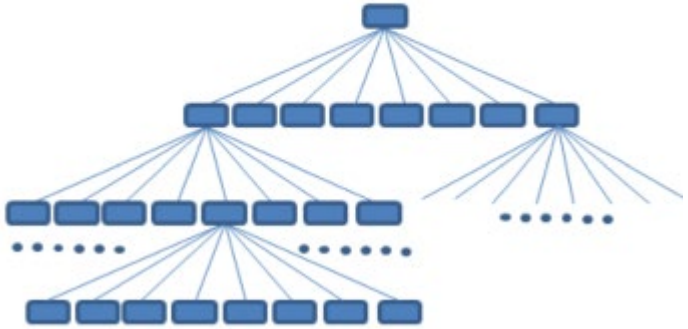

nD-PointCloud

netherlands eScience center

Problems

netherlands eScience center

1.

by Unknown Author is licensed under

netherlands eScience center

2.

Original implementation is for single compute node

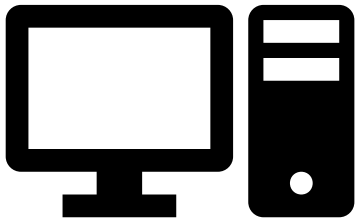

netherlands eScience center

Solutions

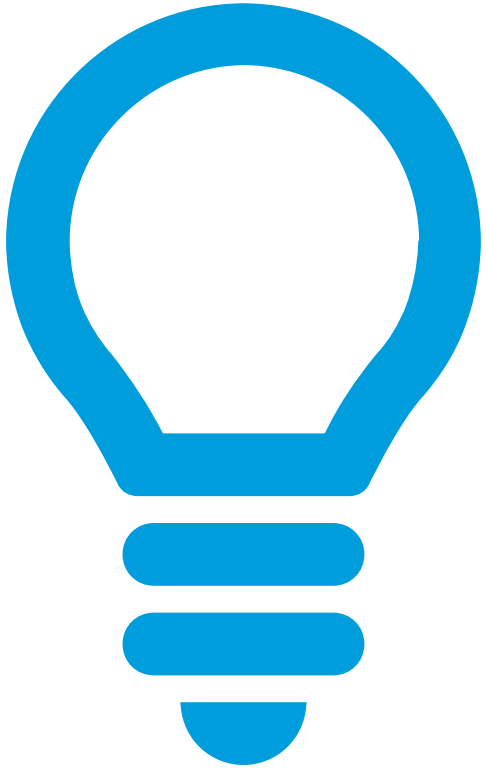

netherlands eScience center

1. Partition

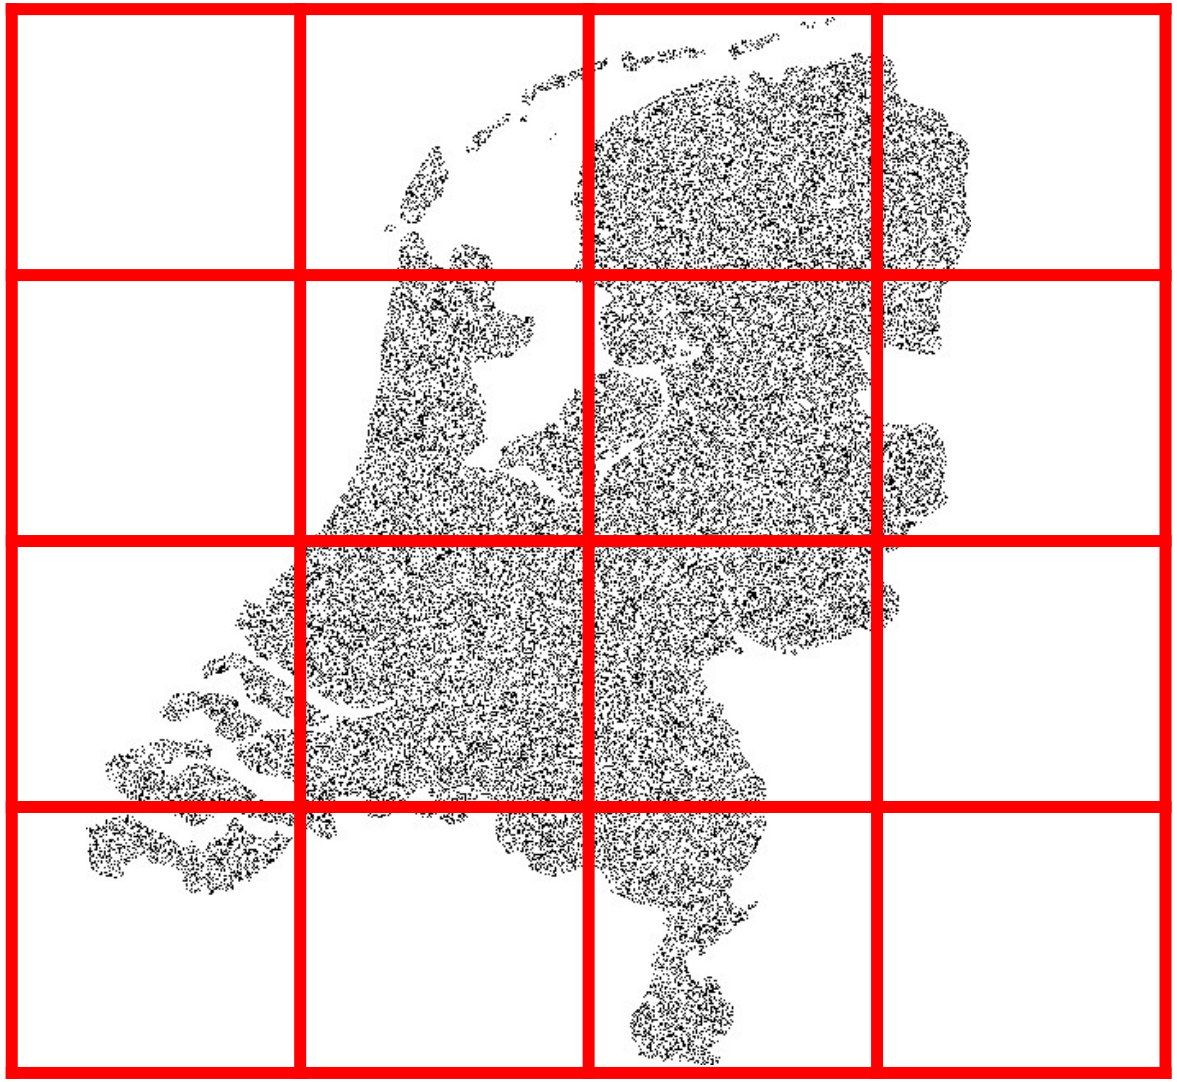

netherlands eScience center

2. MPI

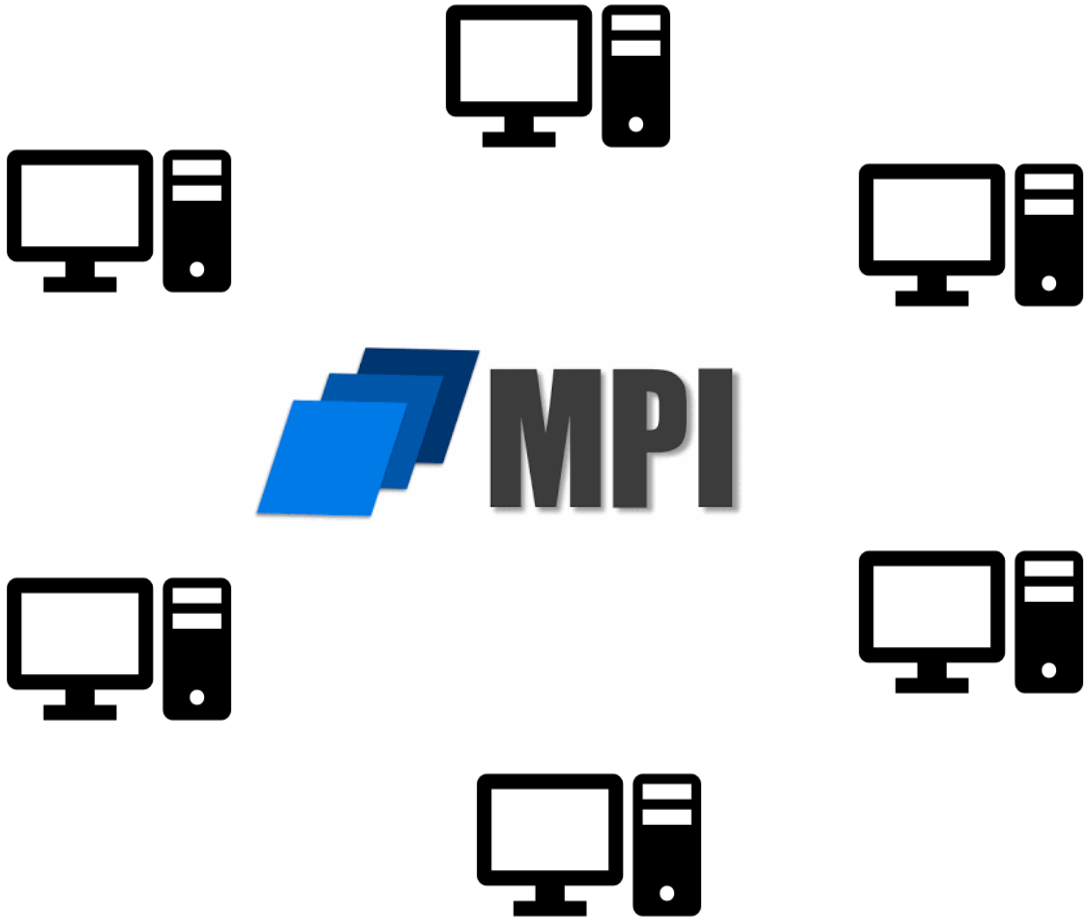

netherlands eScience center

PotreeConverter

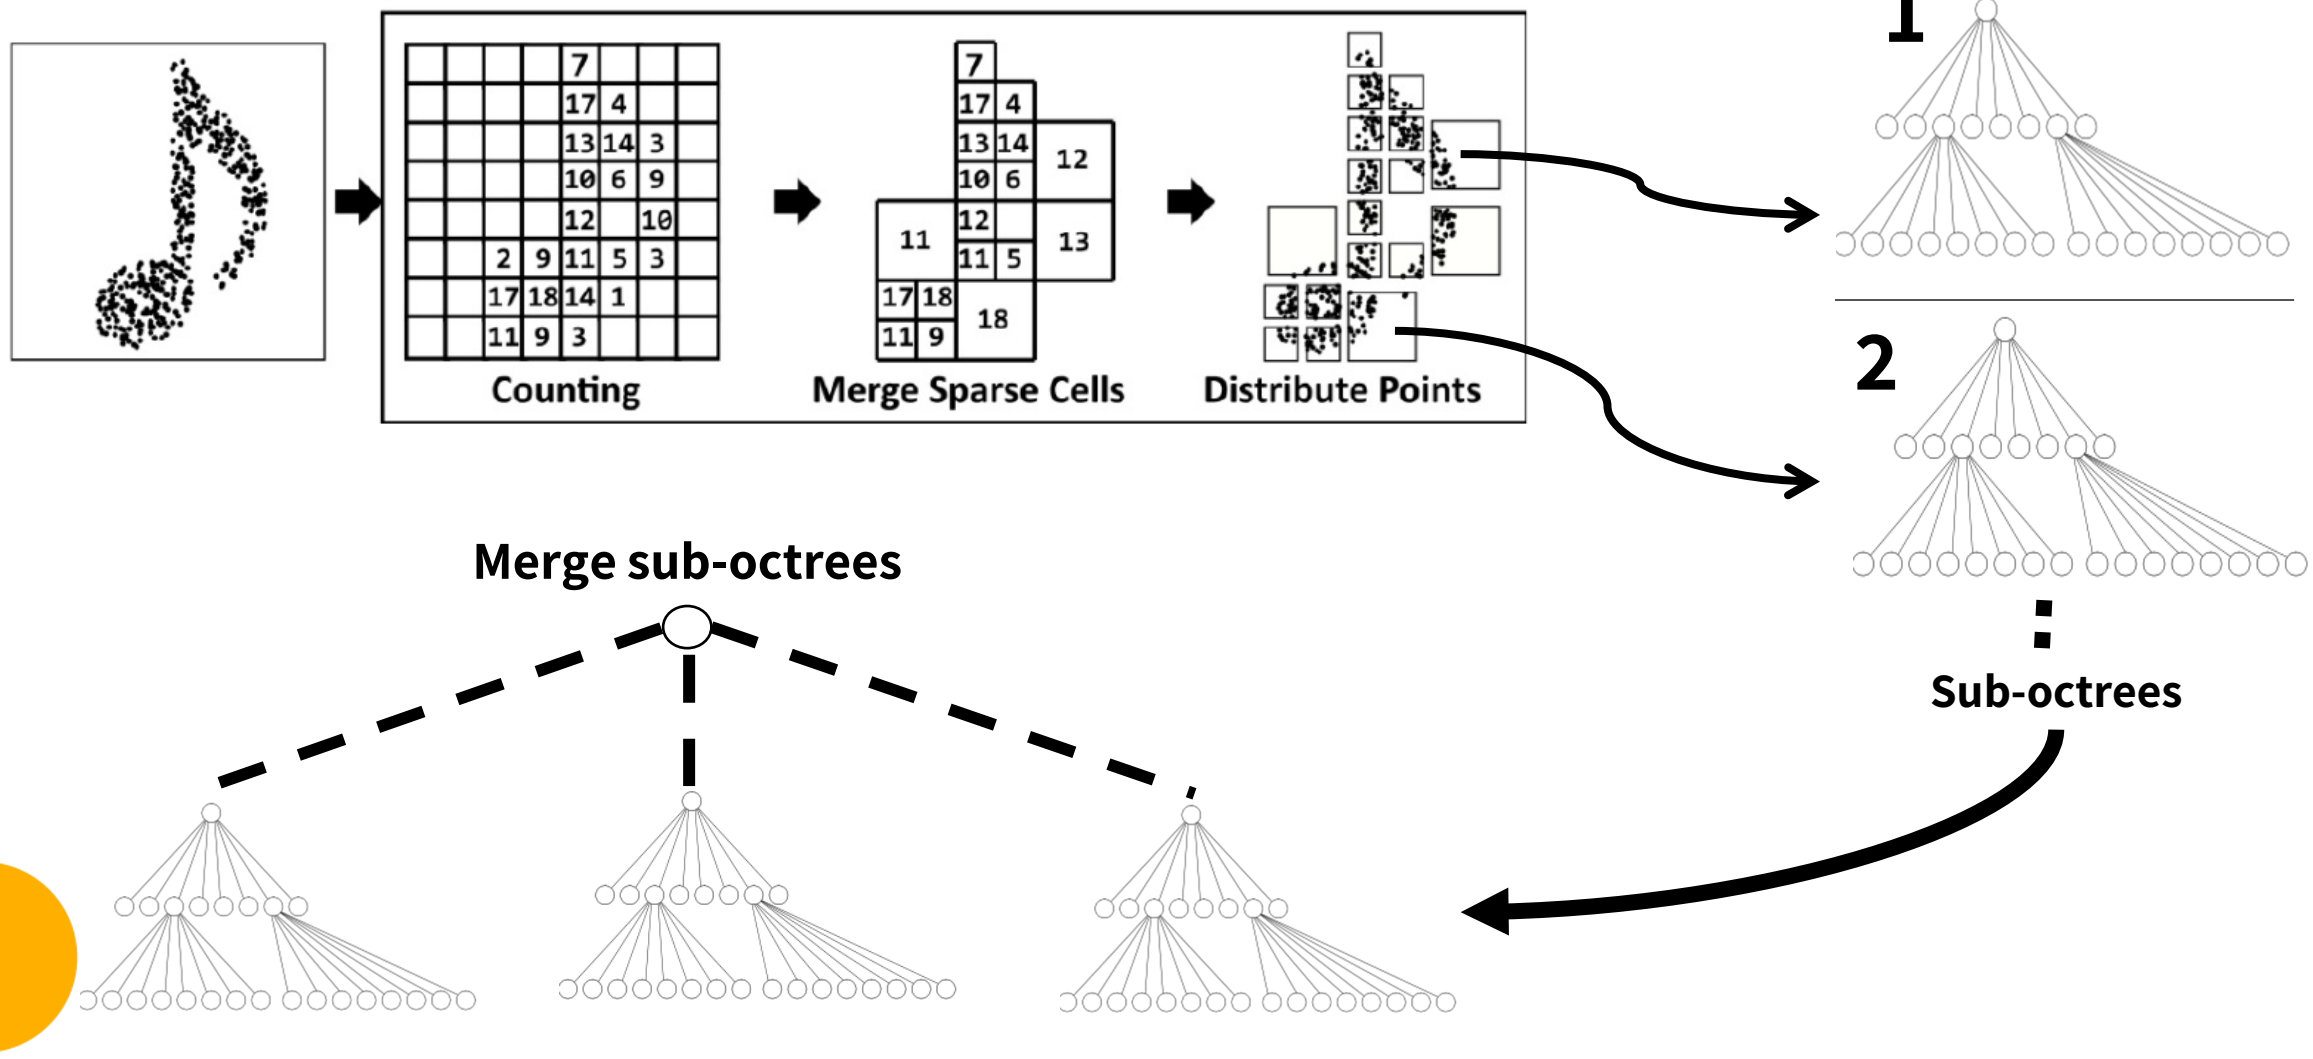

netherlands eScience center

PotreeConverterMPI Counting

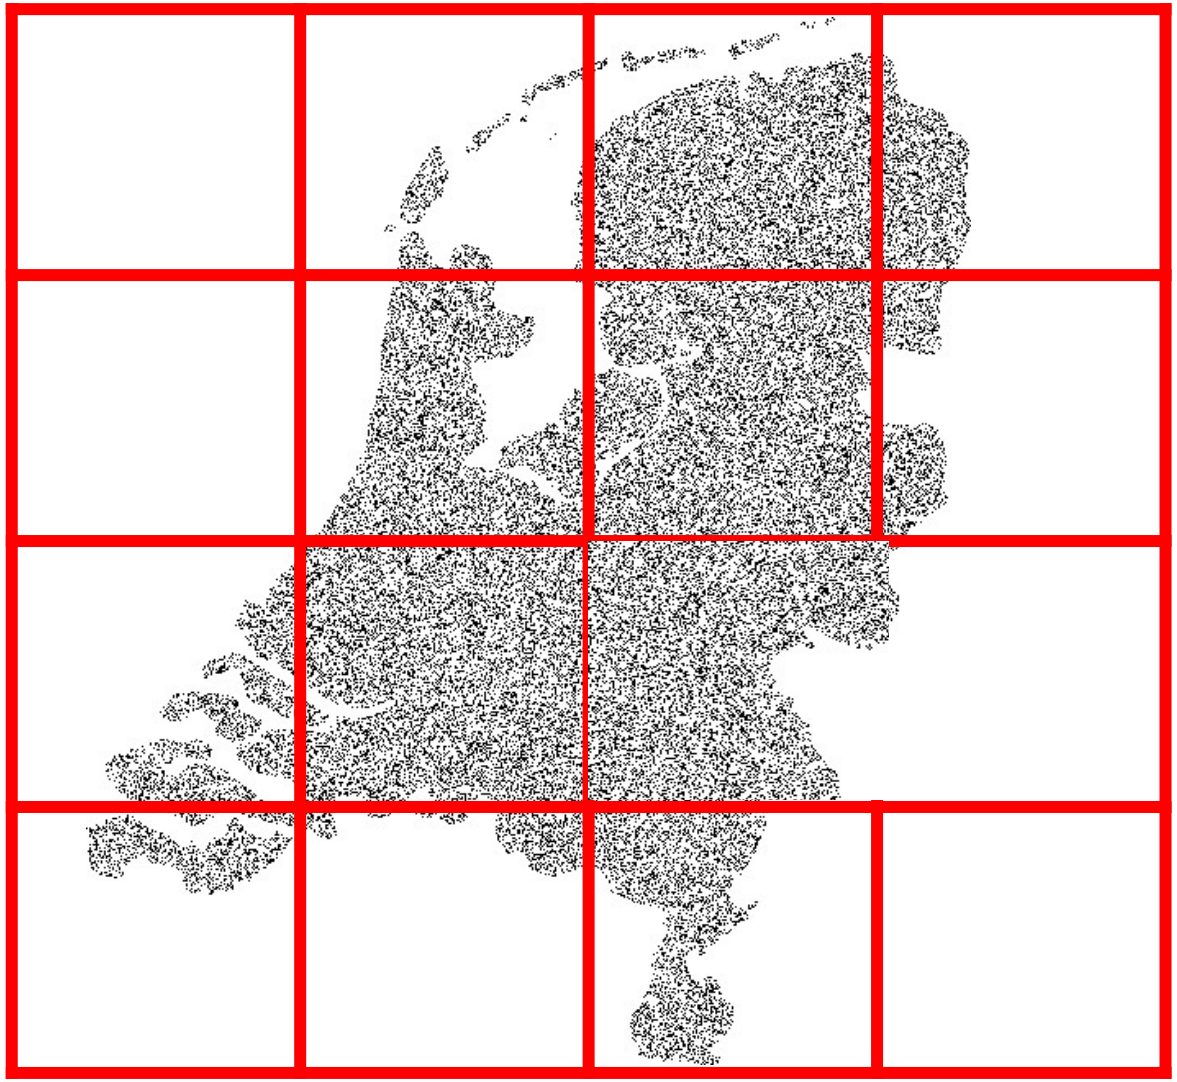

netherlands eScience center

PotreeConverterMPI Counting

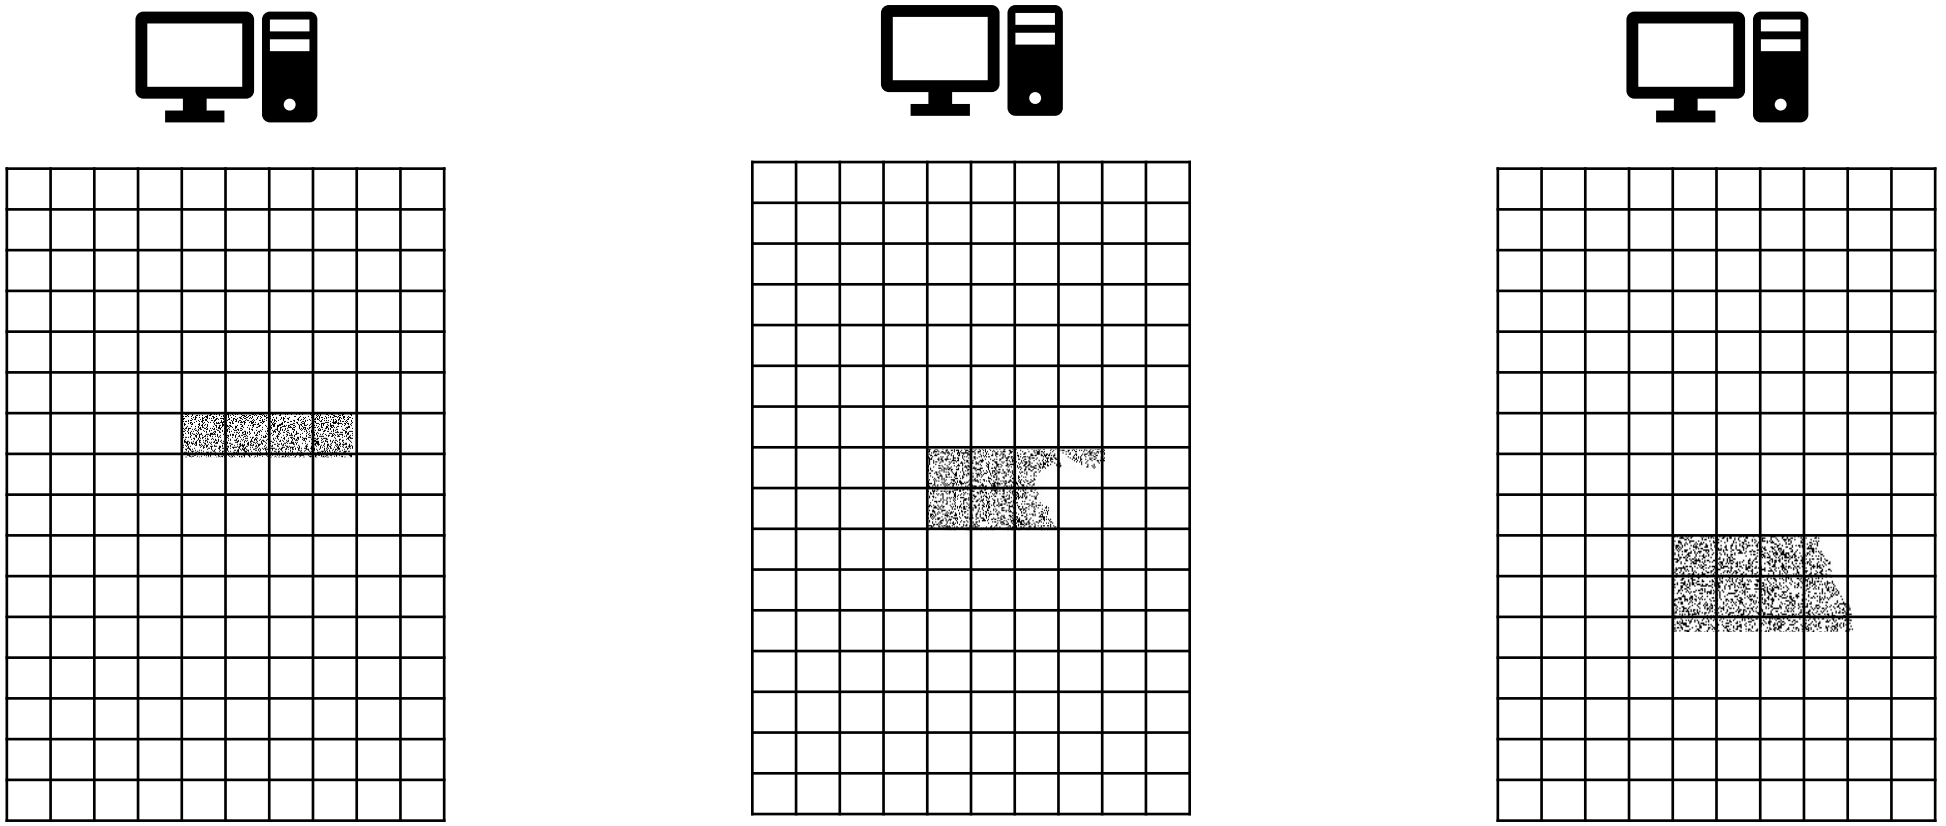

netherlands eScience center

PotreeConverterMPI Counting

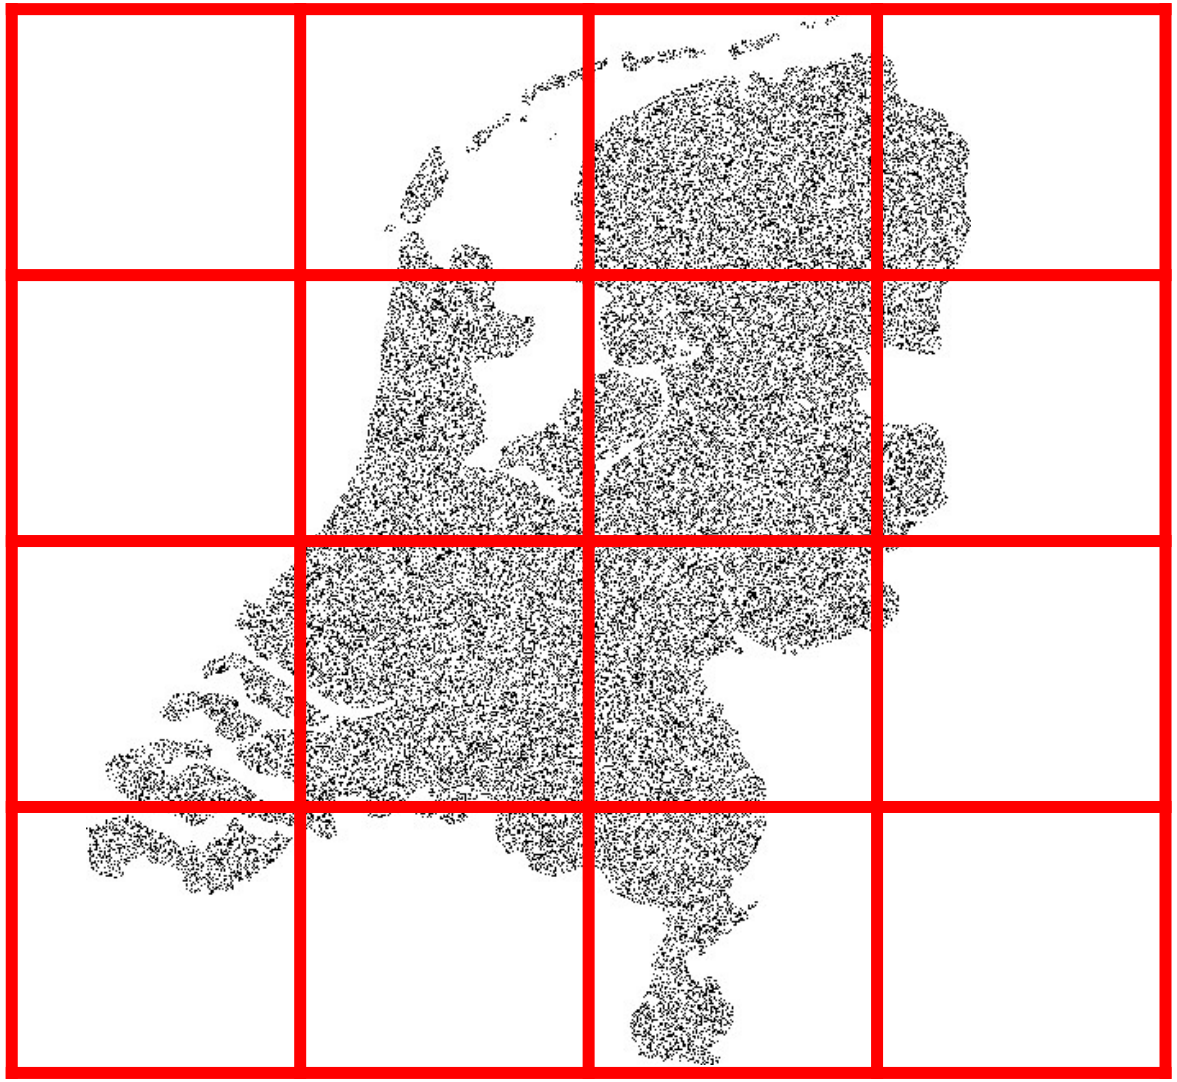

netherlands eScience center

PotreeConverterMPI Counting

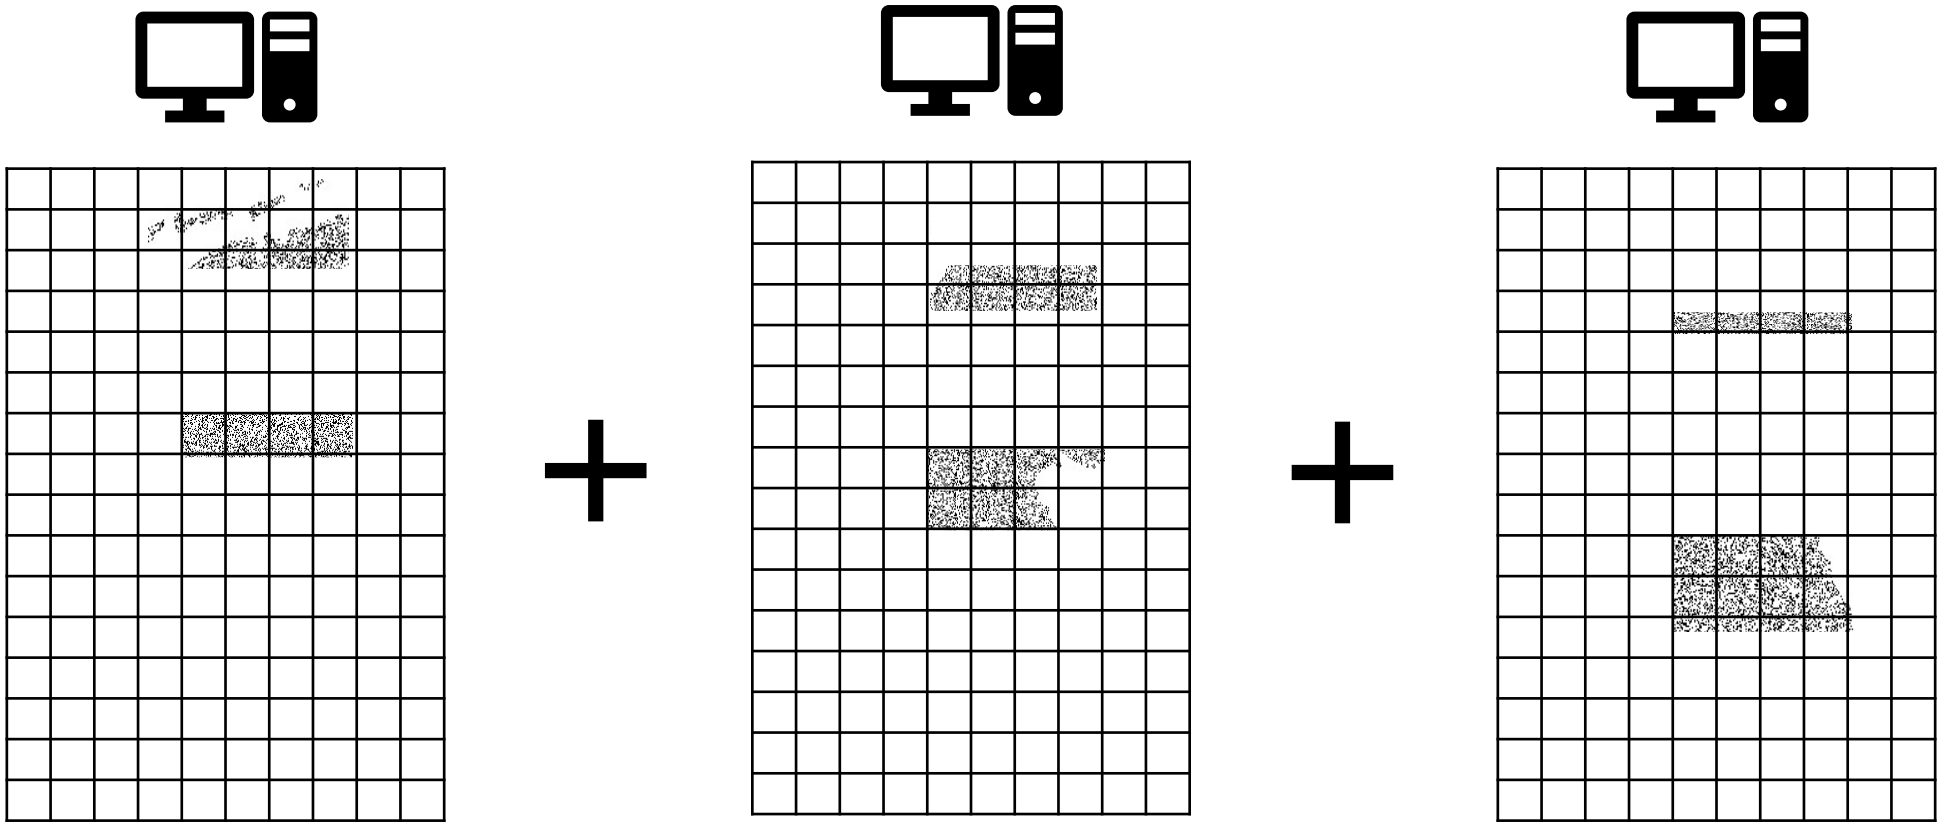

netherlands eScience center

PotreeConverterMPI Counting

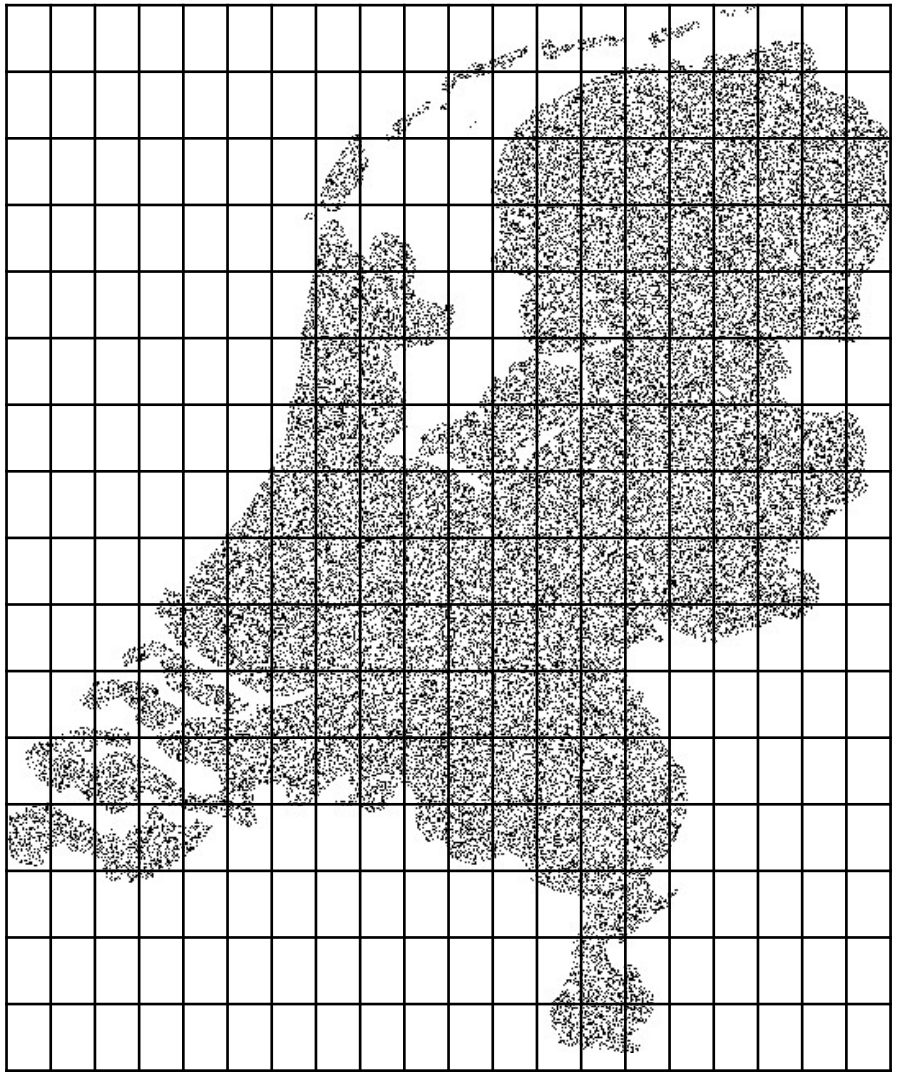

netherlands eScience center

PotreeConverterMPI erge sparse cells

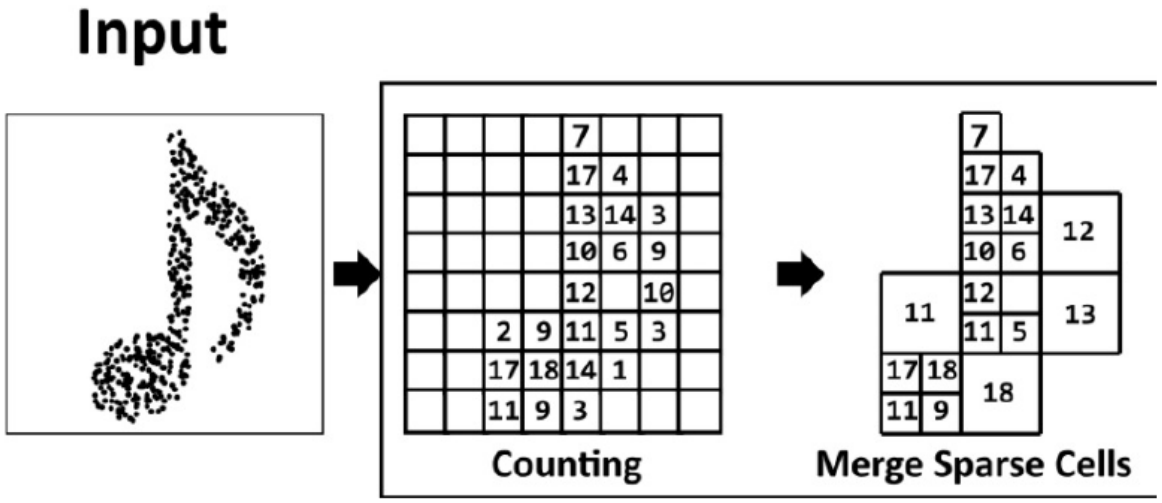

netherlands eScience center

PotreeConverterMPI Distribution

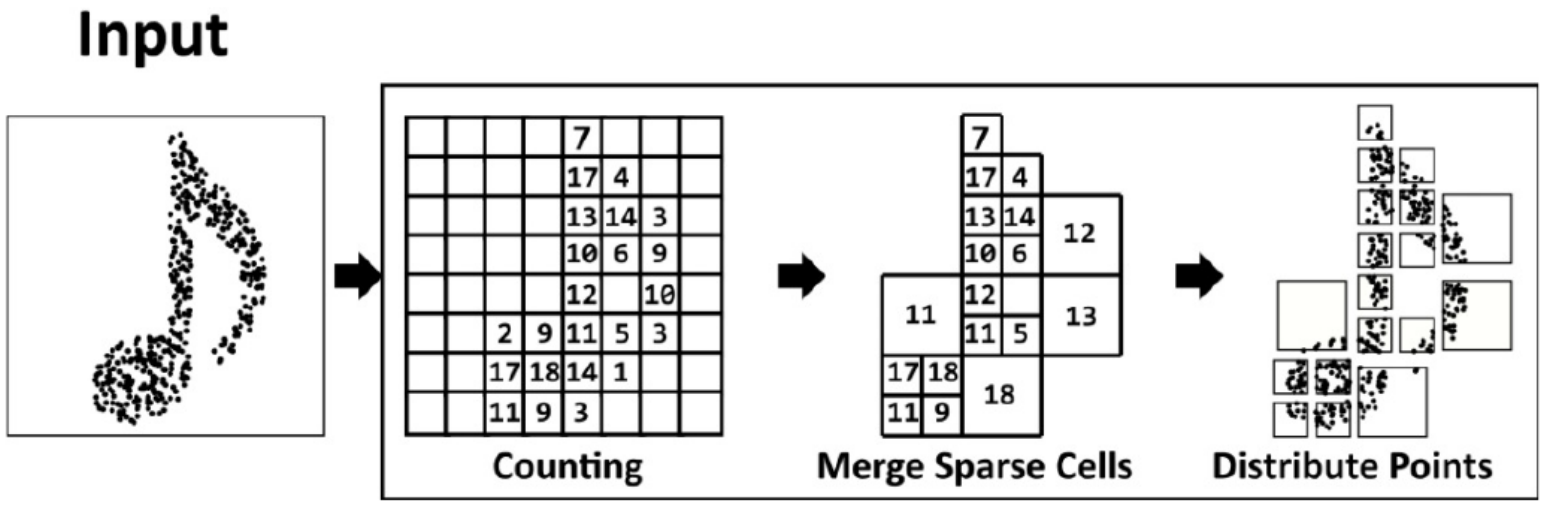

netherlands eScience center

PotreeConverterMPI Distribution

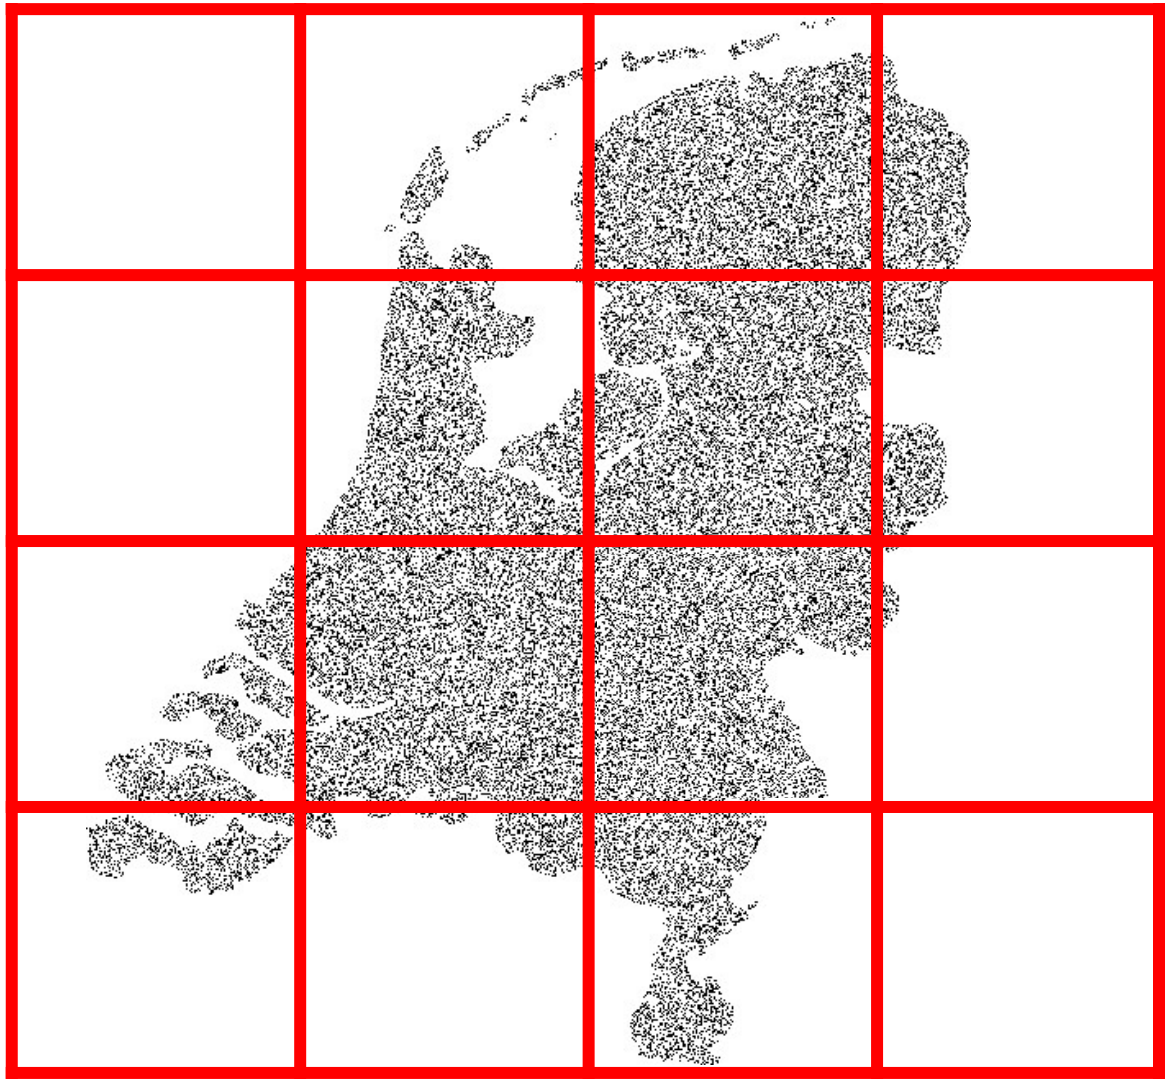

netherlands eScience center

PotreeConverterMPI Indexing

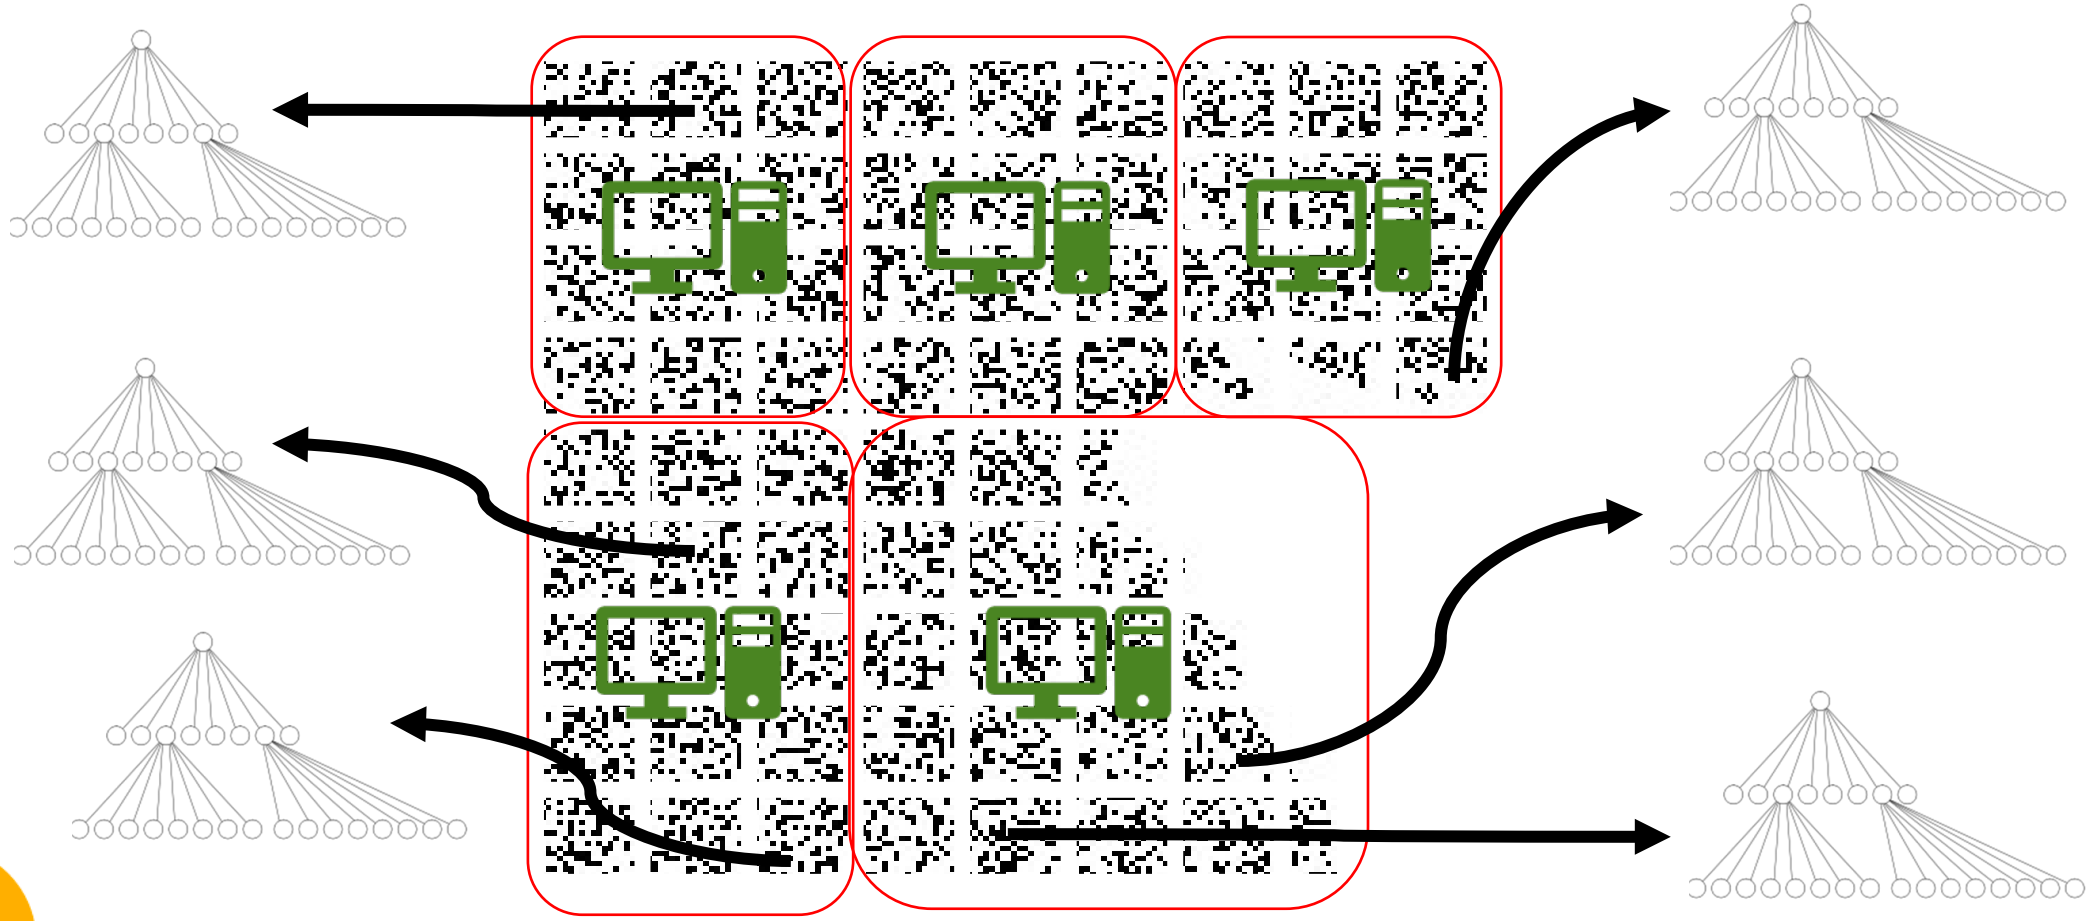

netherlands eScience center

PotreeConverterMPI Distribution & Indexing

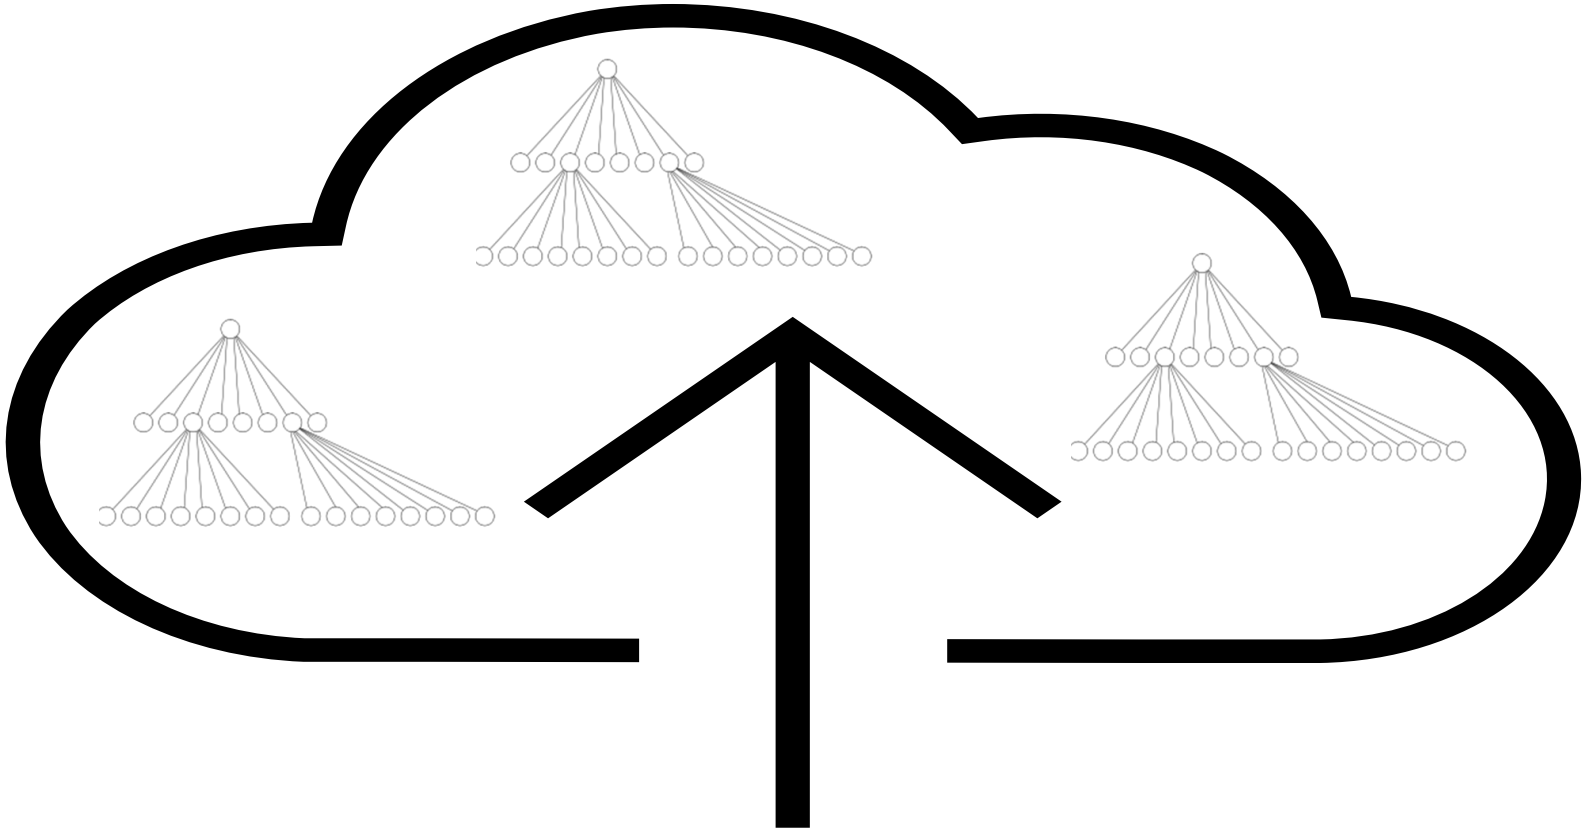

netherlands eScience center

Merge sub-octrees

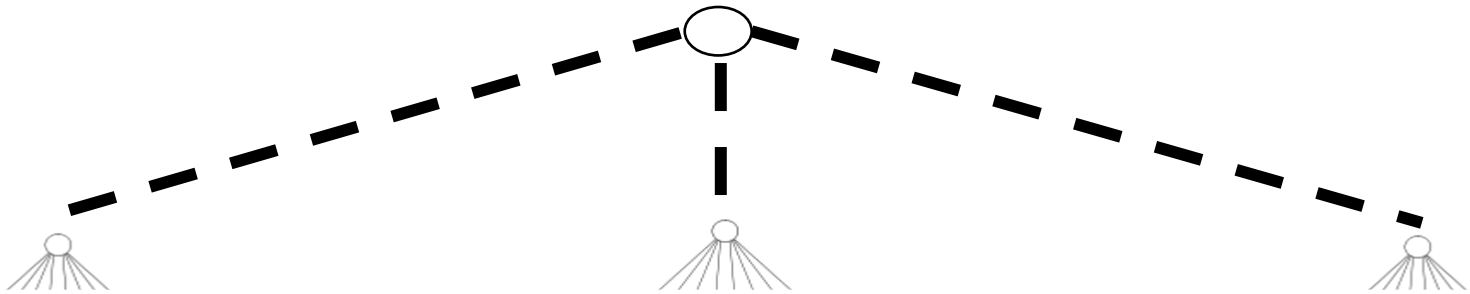

netherlands eScience center

Partition loader/unloader script

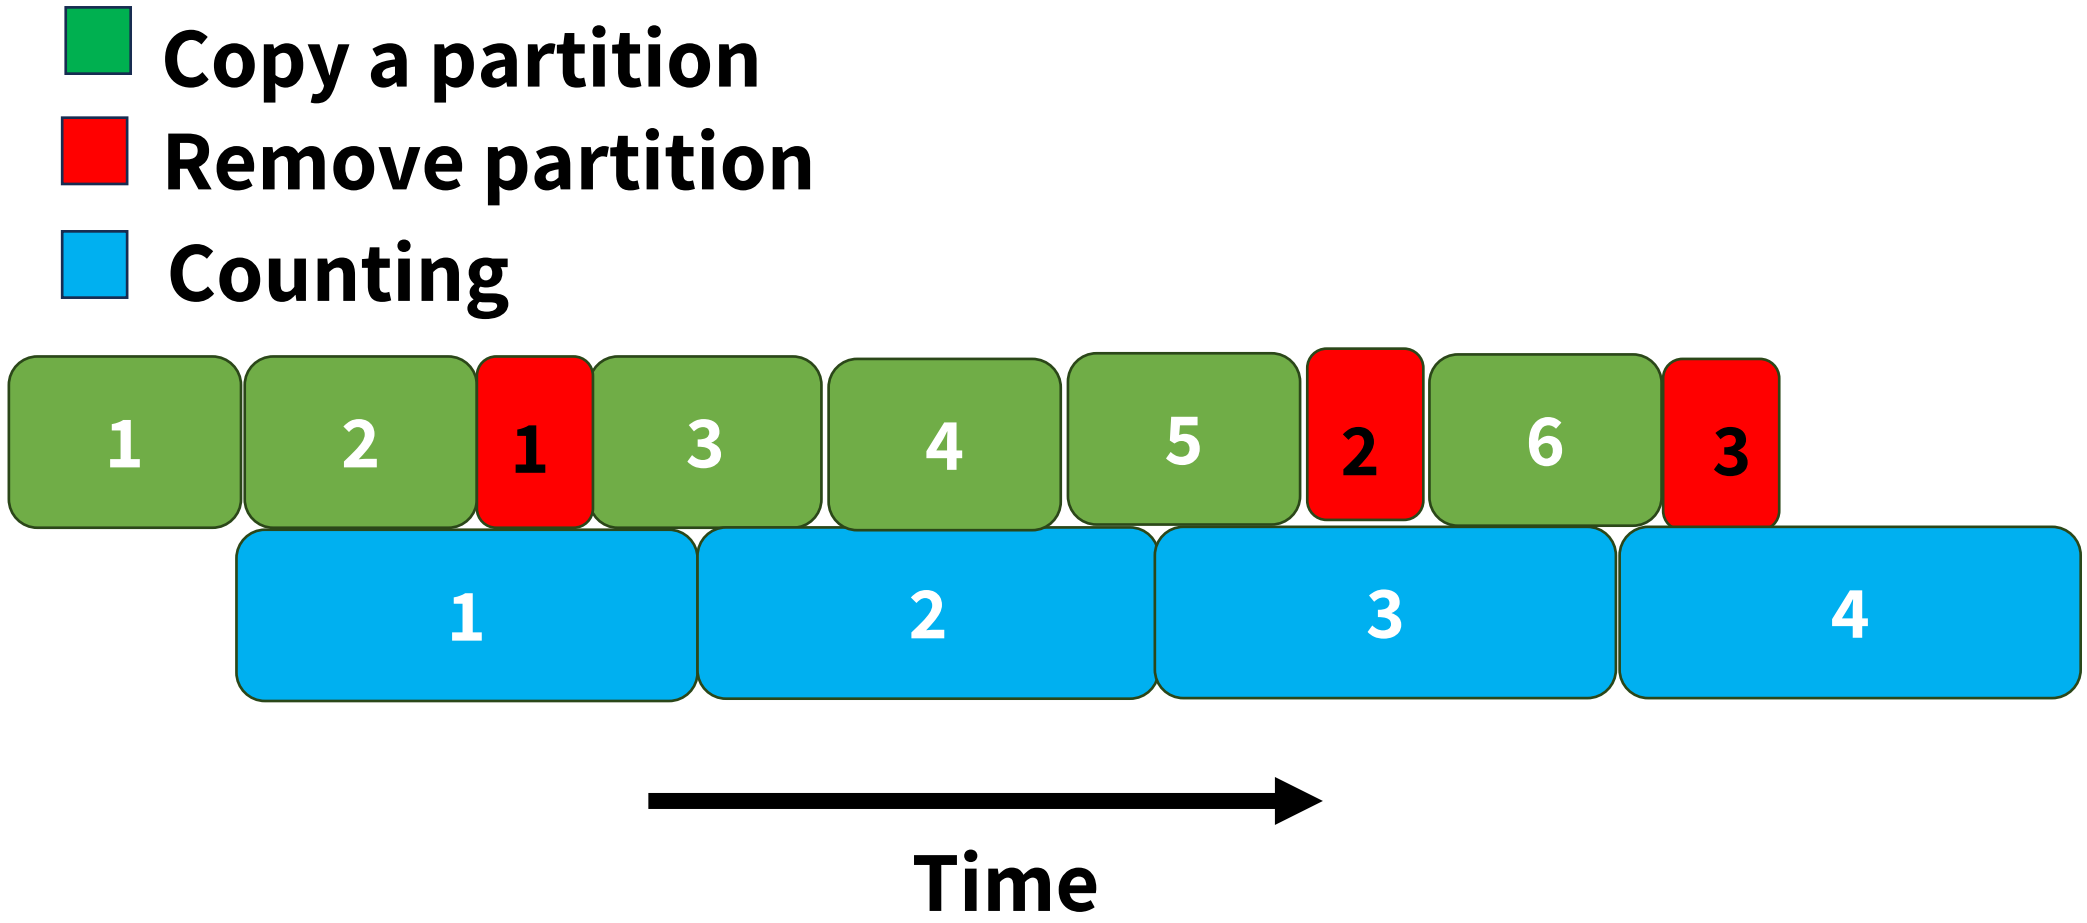

netherlands eScience center

Partition loader/unloader script

- Copy a partition
- Remove partition
- Move partial output to “BIG” storage
- Distribution and Indexing

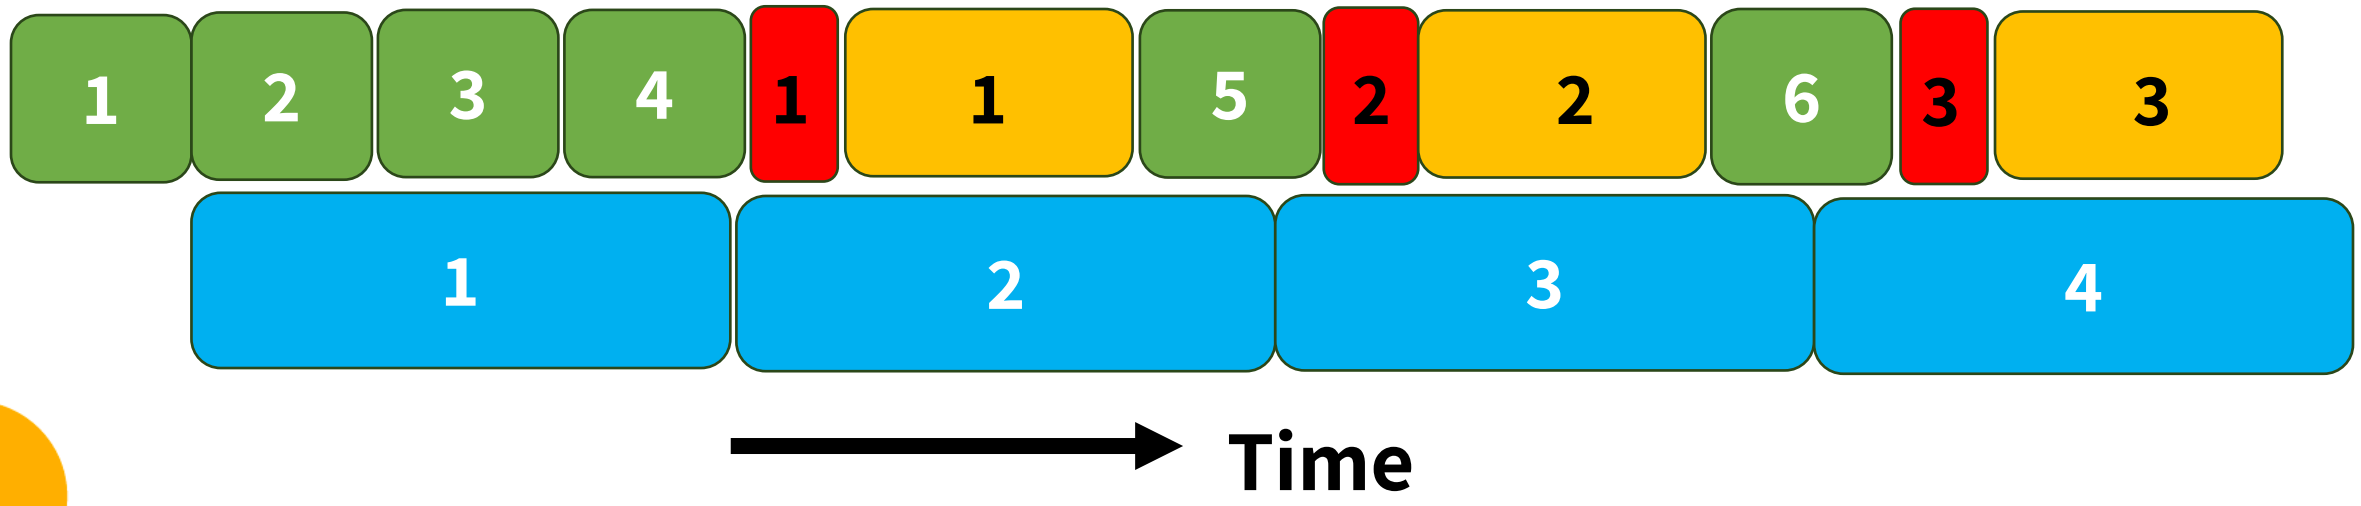

netherlands eScience center

Performance

- Converted AHN3 on 20 nodes of Delftblue. 48 threads per node.

Node Category	Number	Cores	CPU/GPU	RAM	SSD
Compute 1	218	48	2x Intel XEON E5-6248R 24C 3.0GHz	185 GB	150 GB
Compute 2	90	64	2x Intel Xeon E5-6448Y 32C 2.1GHz	250 GB	360 GB
Fat type-a	6	48	2x Intel XEON E5-6248R 24C 3.0GHz	750 GB	150 GB
Fat type-b	4	48	2x Intel XEON E5-6248R 24C 3.0GHz	1500 GB	150 GB

netherlands eScience center

Performance

- ========== STATS ==========

#points: 557'915'228'142

#input files: 1'370

sampling method: poisson

input file size: 14'548.8GB

duration: 150'155.039s (41 hours)

throughput (points/s) 3.7M

netherlands eScience center

Performance

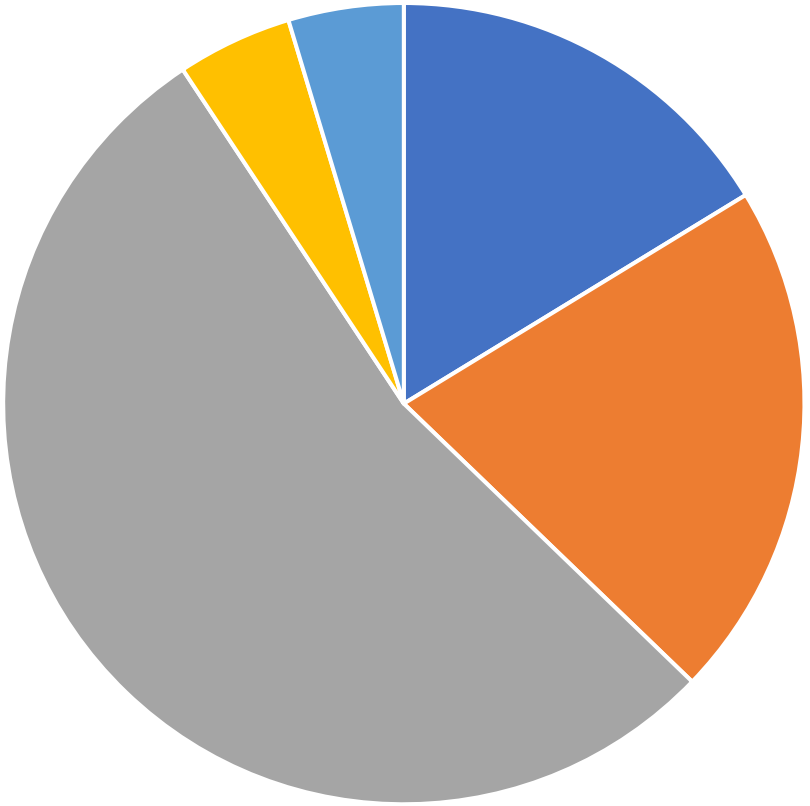

Counting (16%) Distribution(18.6%) Indexing(55%) Output(4.5%) Merging(4.5%)

netherlands eScience center

Performance

- Perfect overlap in copying/removing data and processing in counting phase

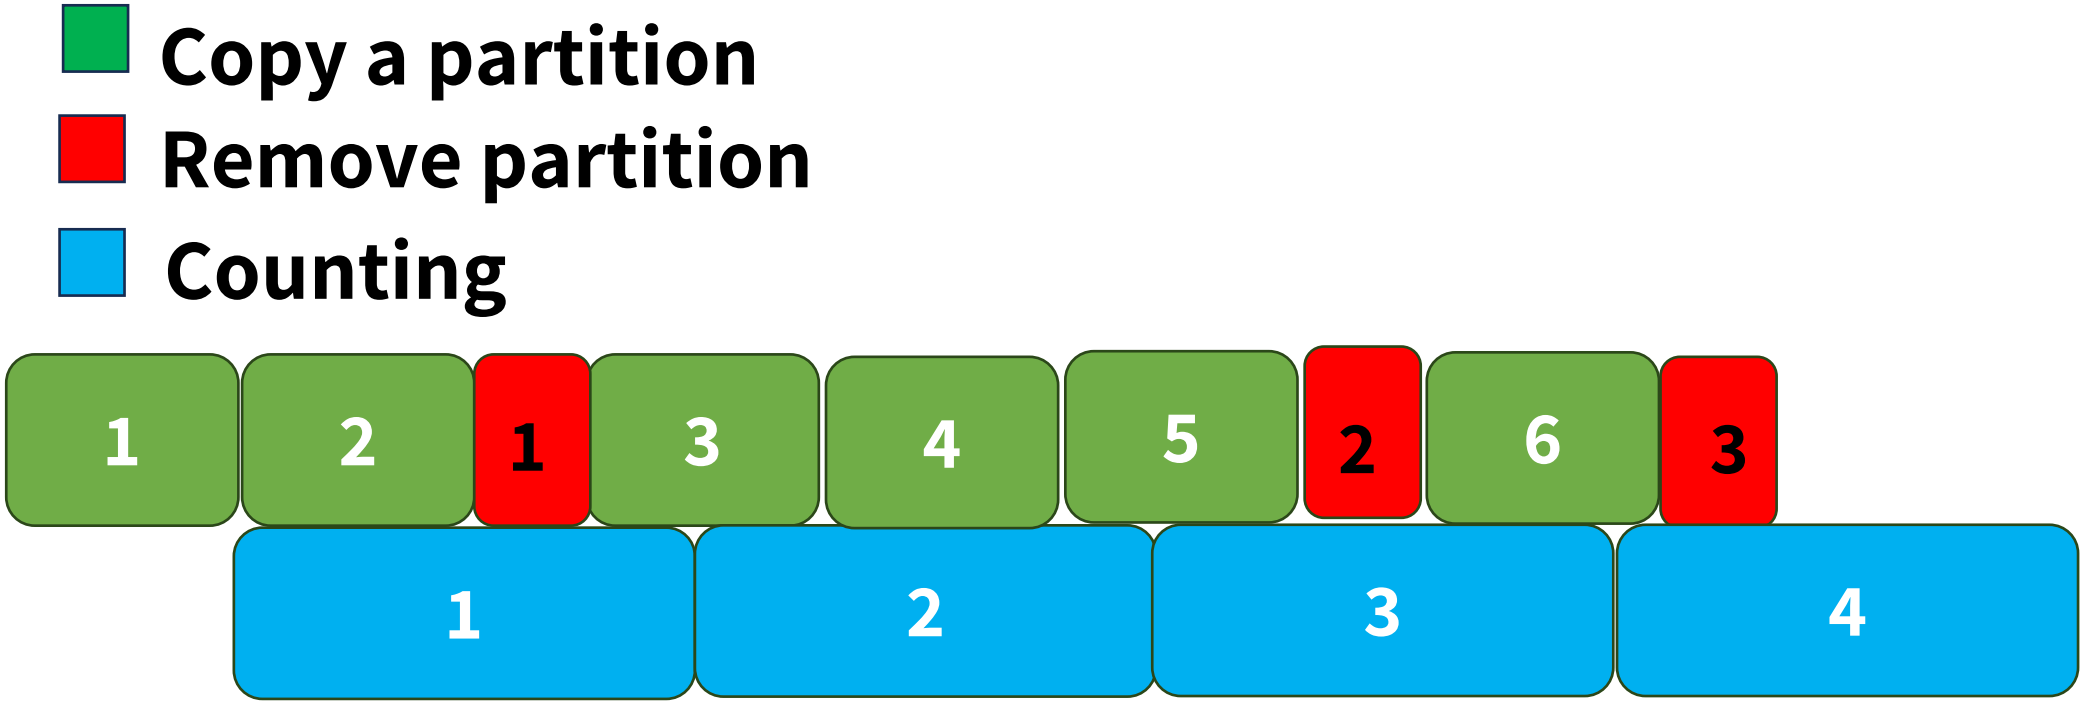

netherlands eScience center

Performance

- Perfect overlap in copying/removing data and processing in distribution and indexing phase.

Copy a partition

Remove partition

Move partial output to “BIG” storage

Distribution and Indexing

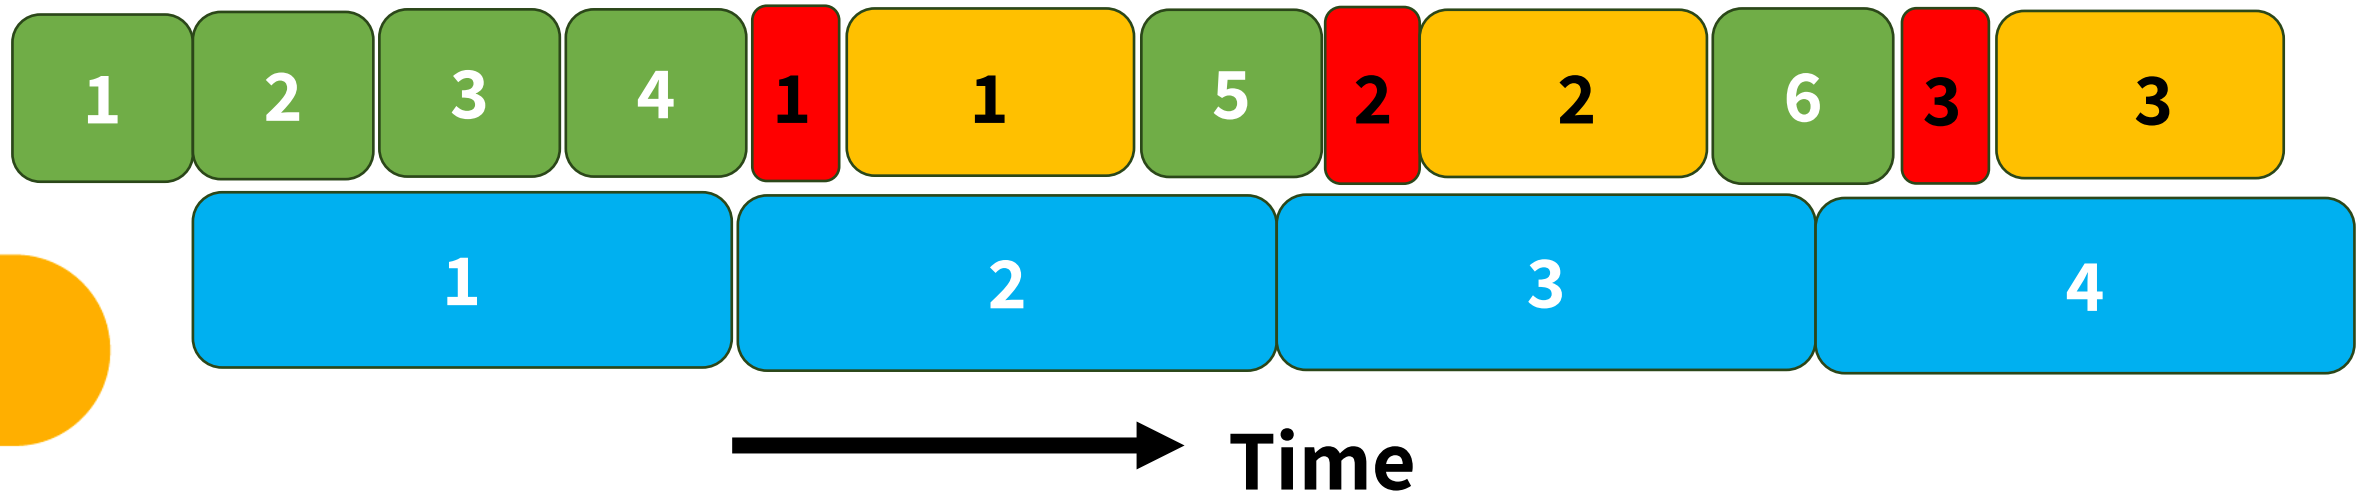

netherlands eScience center

Limitations

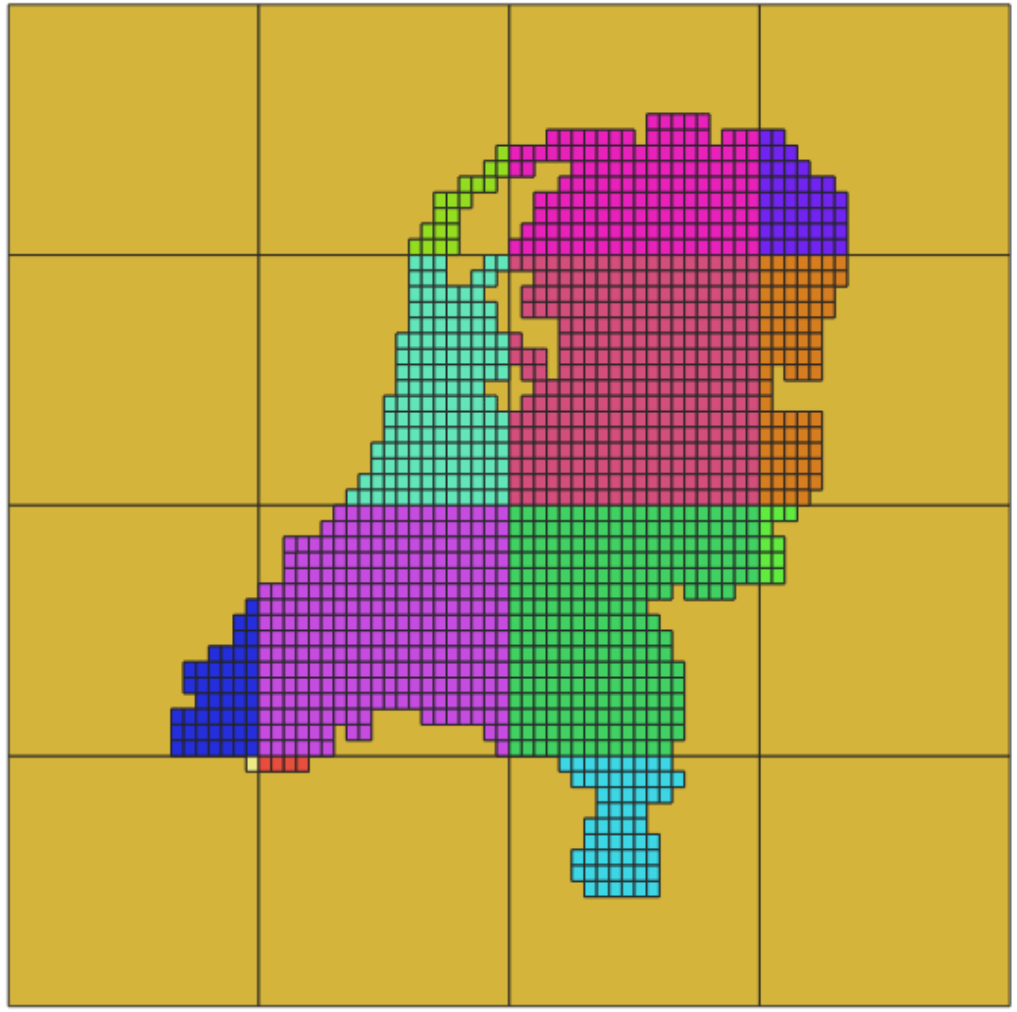

netherlands eScience center

Limitations

- The implementation is 6-7x slower than ideal.
  - Investigating. May be due to shared scratch space. Copying the data to local SSD per node may help

netherlands eScience center

Further Improvments

- Copy a partition
- Remove partition
- Move partial output to “BIG” storage
- **Counting, Distribution and Indexing**

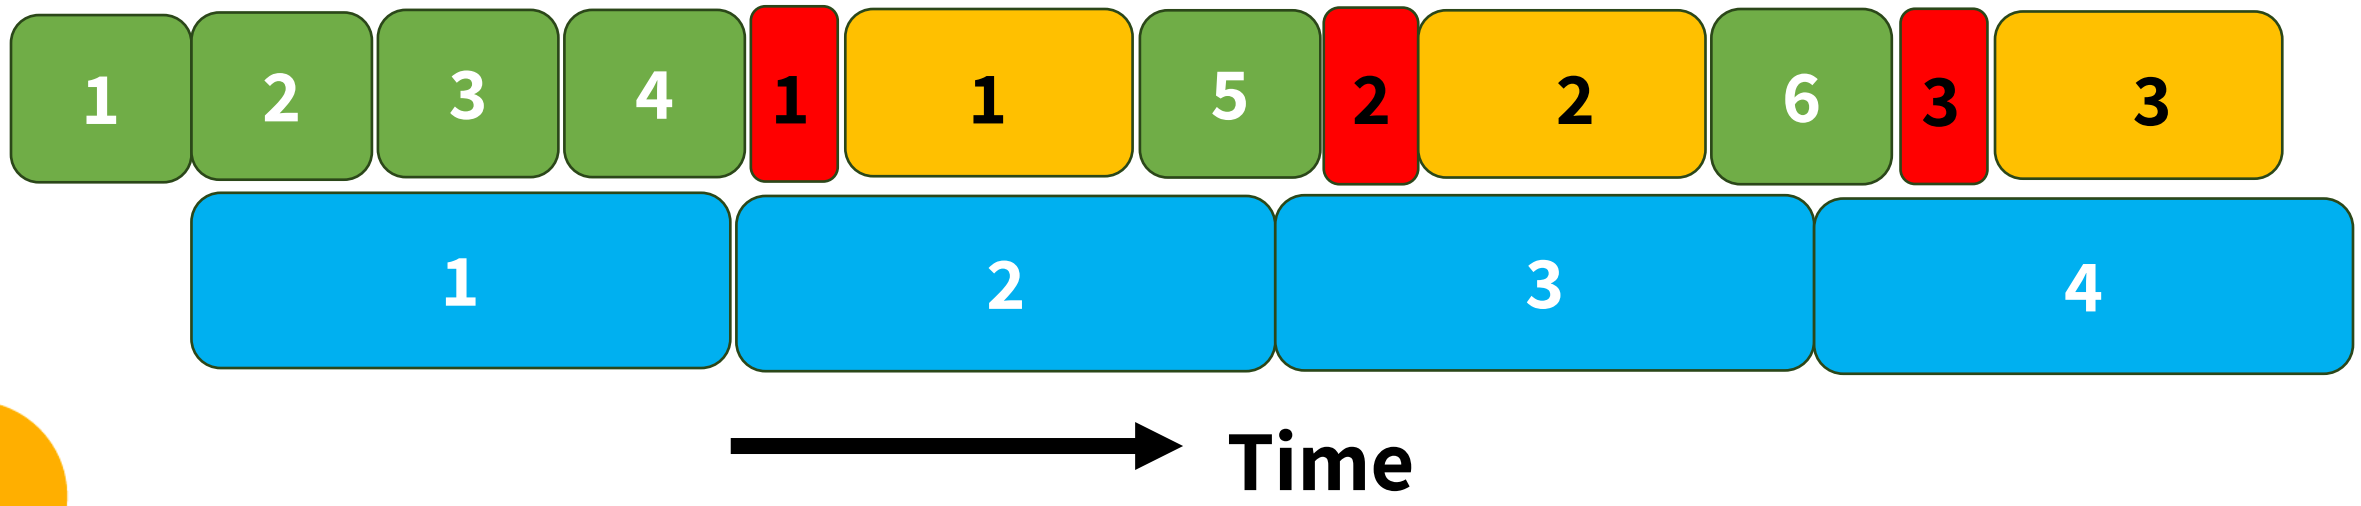

netherlands eScience center

Target

"An implementation that works on a laptop to a cluster irrespective of the input data size. "

netherlands eScience center

Questions and Suggestions?

www.eScienceCenter.nl

Naumanahmedraja@gmail.com