UNIVERSITY OF CALIFORNIA
UC RIVERSIDE

CS133
Computational Geometry

Computational Geometry on Big Data

Big Geometric Data

Satellite Imagery

More than 17 PB

Geotagged Tweets

Billions of tweets

Check ins

Billions of check ins
Millions more every day

MapReduce

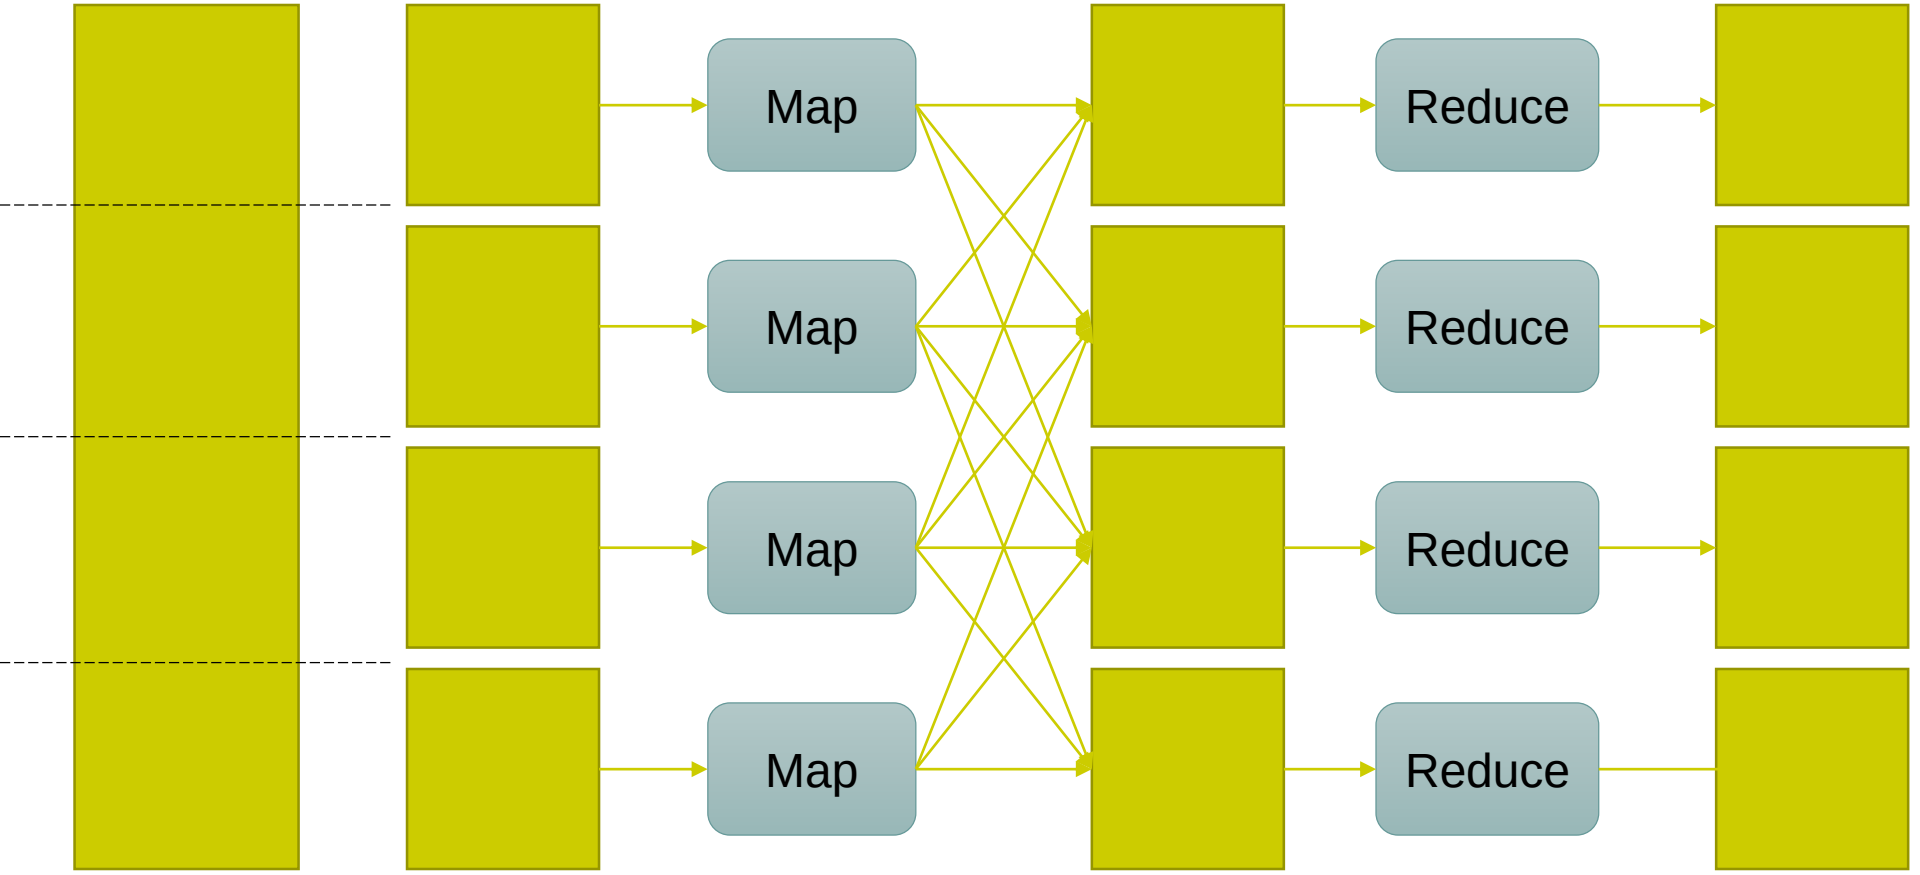
Map
Map
Map
Map
Reduce
Reduce
Reduce
Reduce
Input Big Data
Split Input
Shuffle
Output

CG Algorithms on big data

- Utilize divide and conquer algorithms
- 1. Partition the input across machines
- 2. (Optional) prune partitions that do not contribute to answer
- 3. Apply the algorithm locally in each partition
- 4. Combine the partial answers to compute the final result

Examples

- Convex hull algorithm
- Closest pair
- Farthest pair
- Voronoi diagram/Delaunay triangulation

Data Partitioning

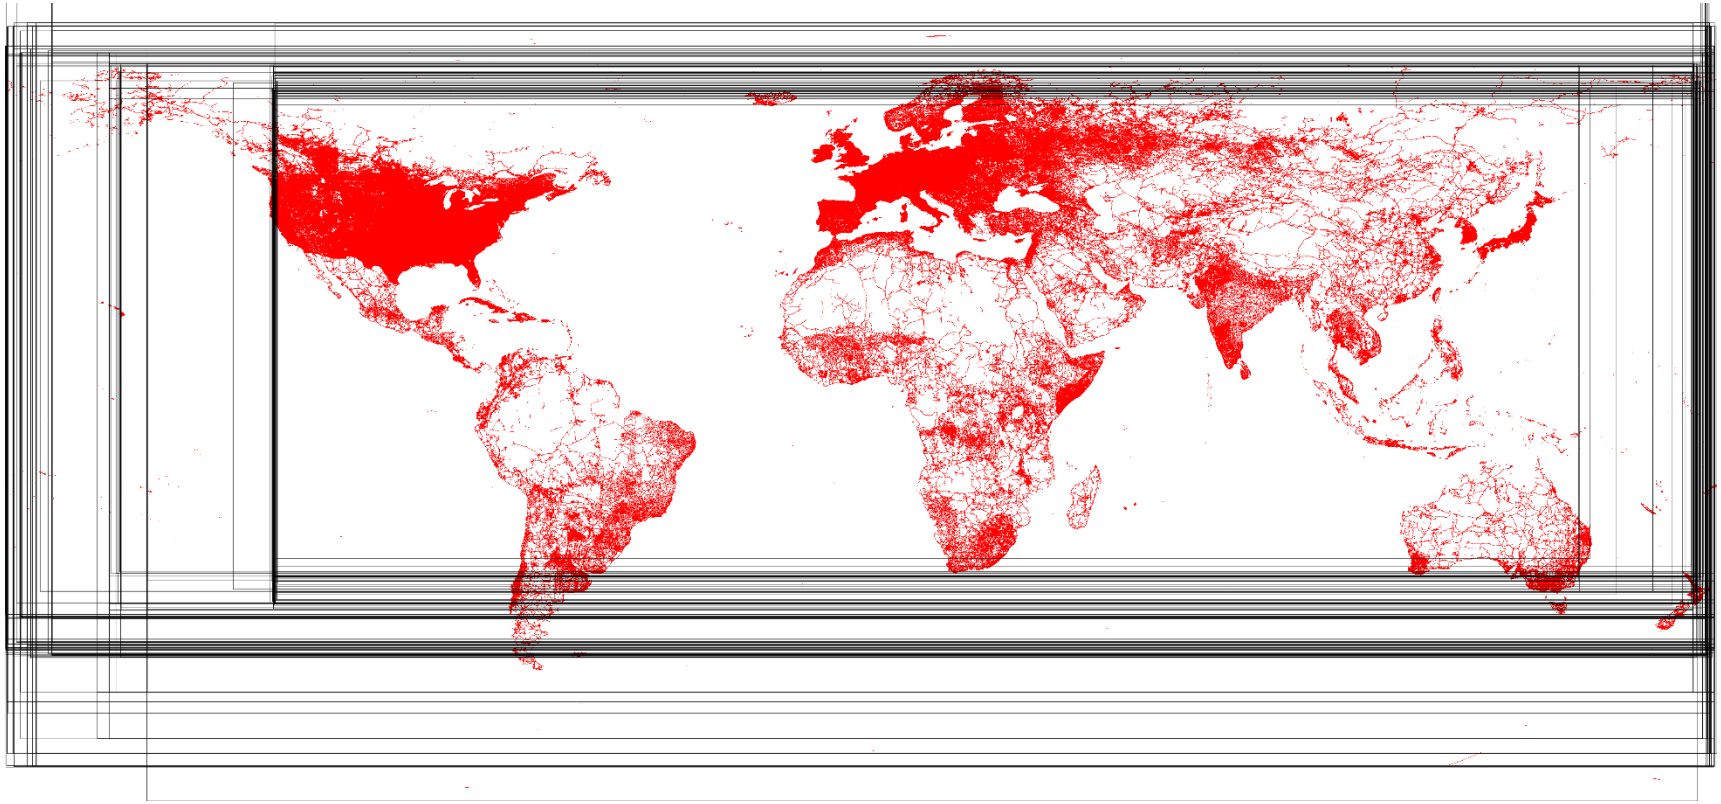

Spatial Partitioning

Skyline (Maximal Vectors)

- Select all non-dominated points

Input

Output

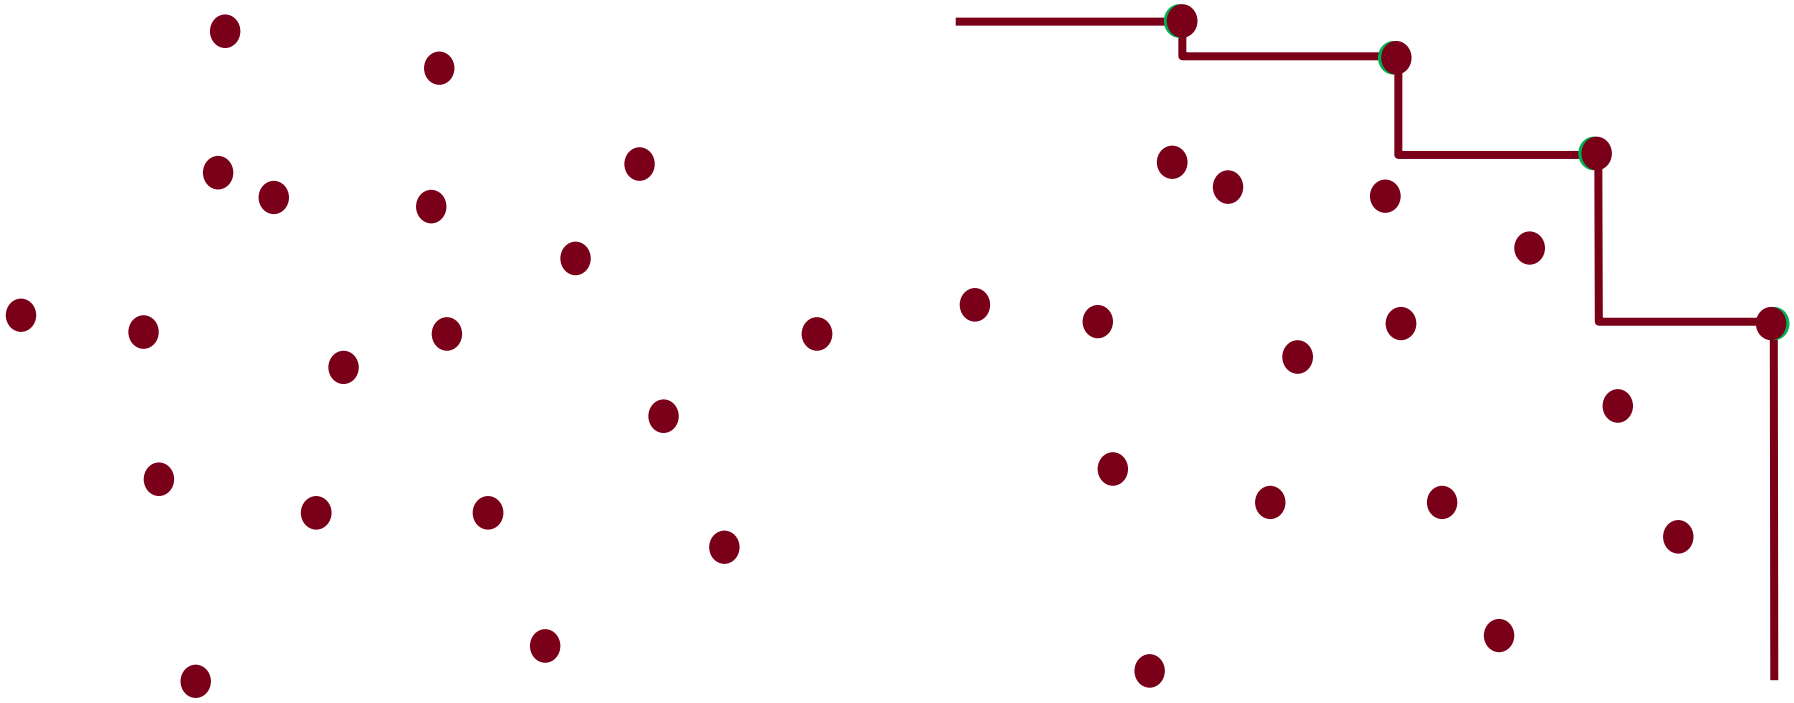

$$(x_1, y_1) \succ (x_2, y_2) \Leftrightarrow x_1 \geq x_2 \wedge y_1 \geq y_2$$

Skyline in MapReduce

Non-spatial partitioning

Spatial partitioning

①Partition

②Pruning

③Local skyline

④Global skyline

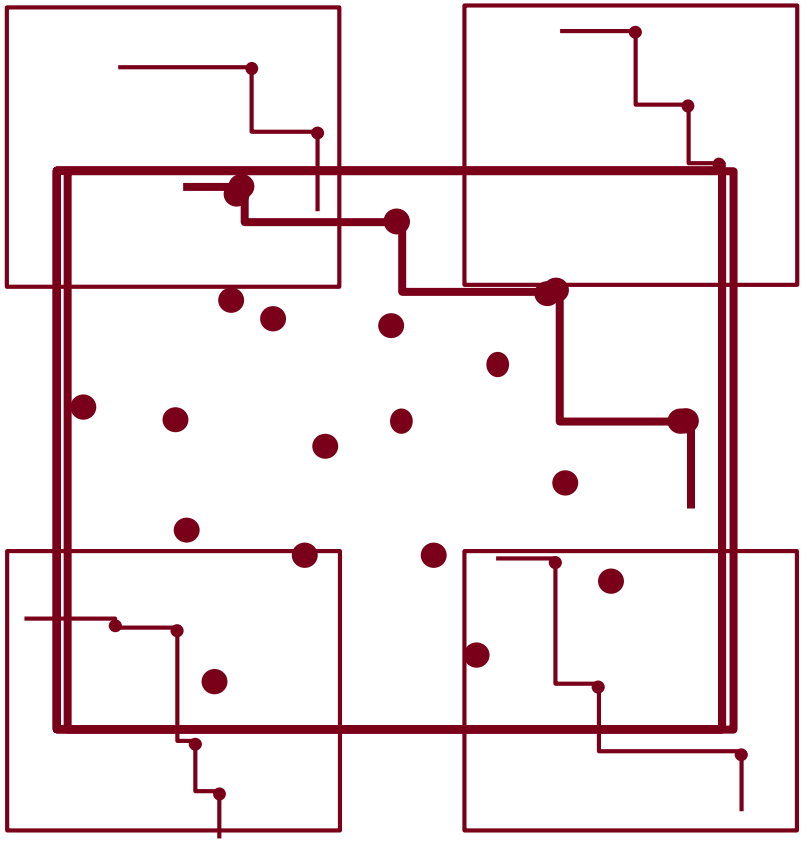

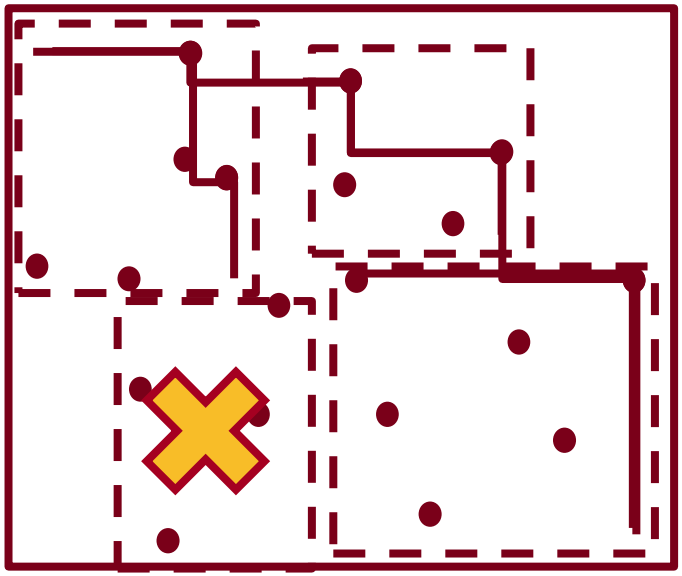

Skyline Pruning

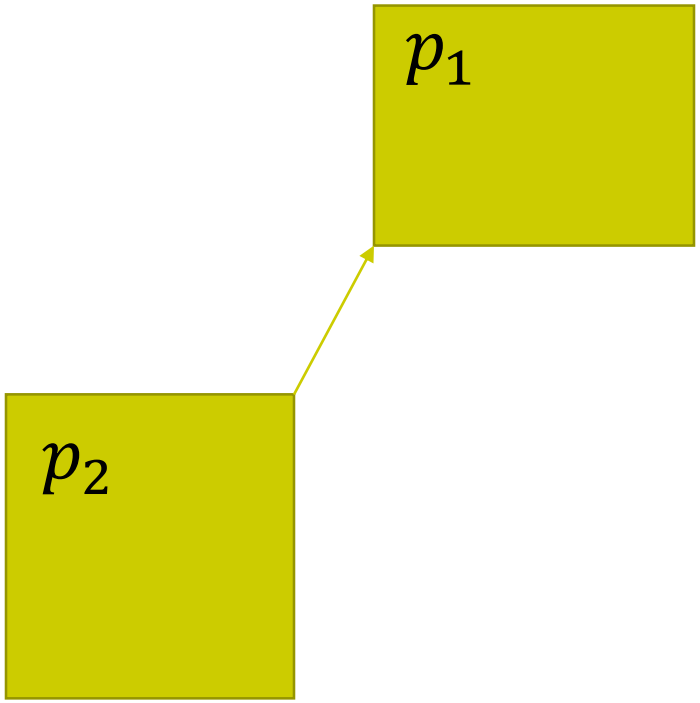

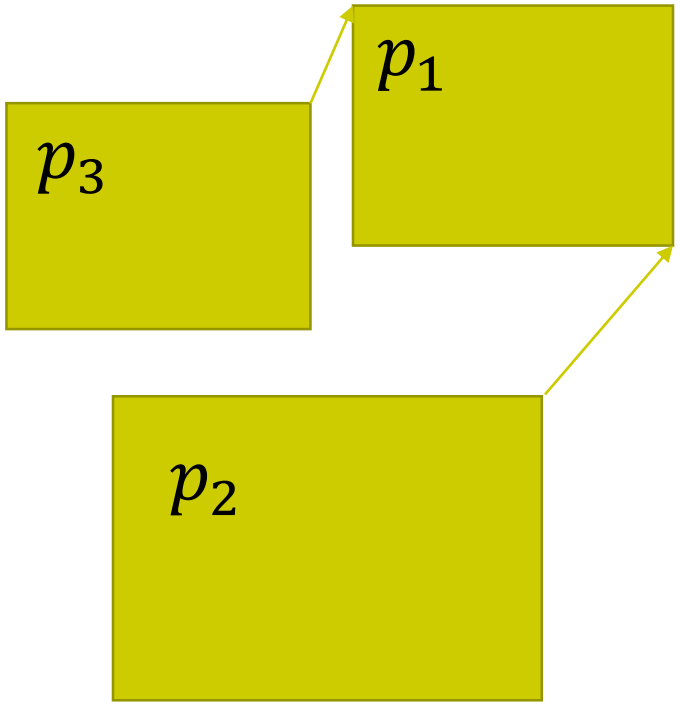

Partition domination rules

$(p_1.xmin, p_1.ymin) \succ (p_2.xmax, p_2.ymax)$

$(p_1.xmin, p_1.ymax) \succ (p_3.xmax, p_3.ymax)$

$(p_1.xmax, p_1.ymin) \succ (p_2.xmax, p_2.ymax)$

Convex Hull

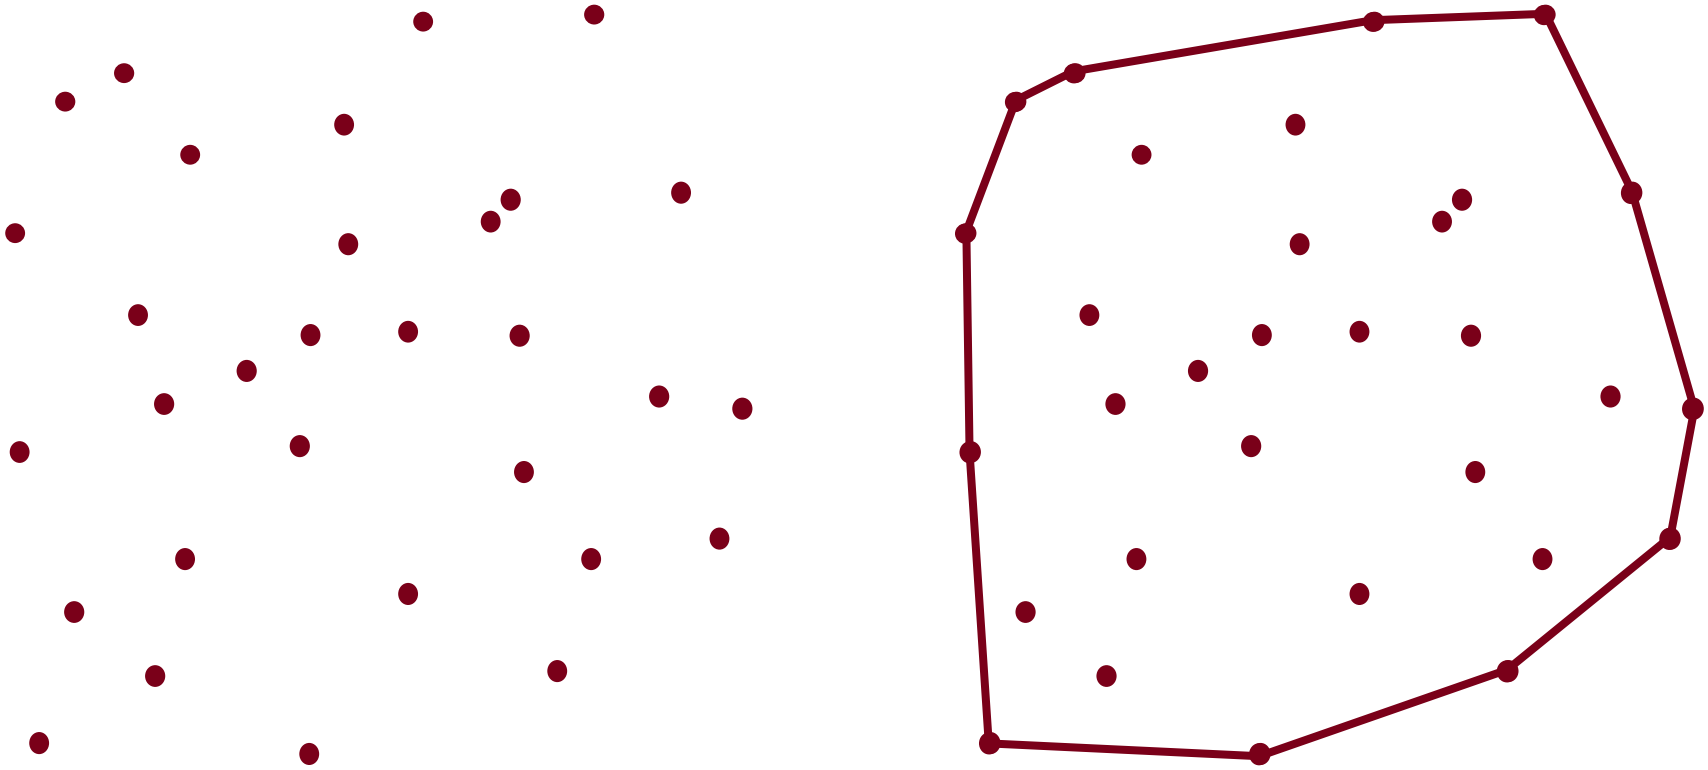

Convex Hull in MapReduce

①Partition

②Pruning

③Local hull

④Global hull

Non-spatial partitioning

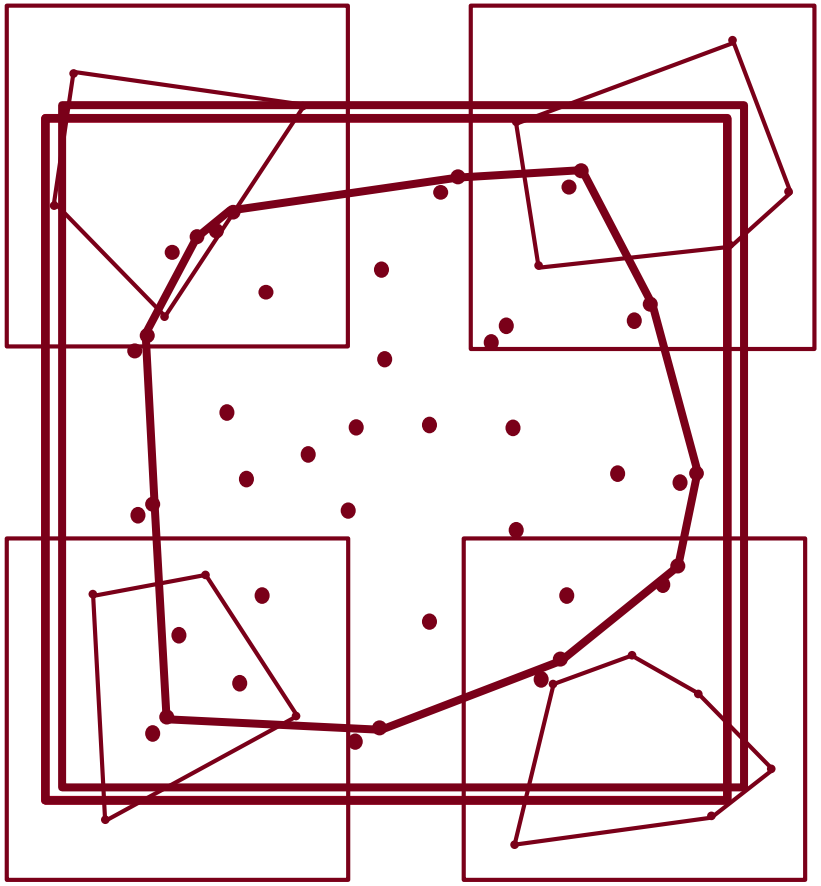

Spatial partitioning

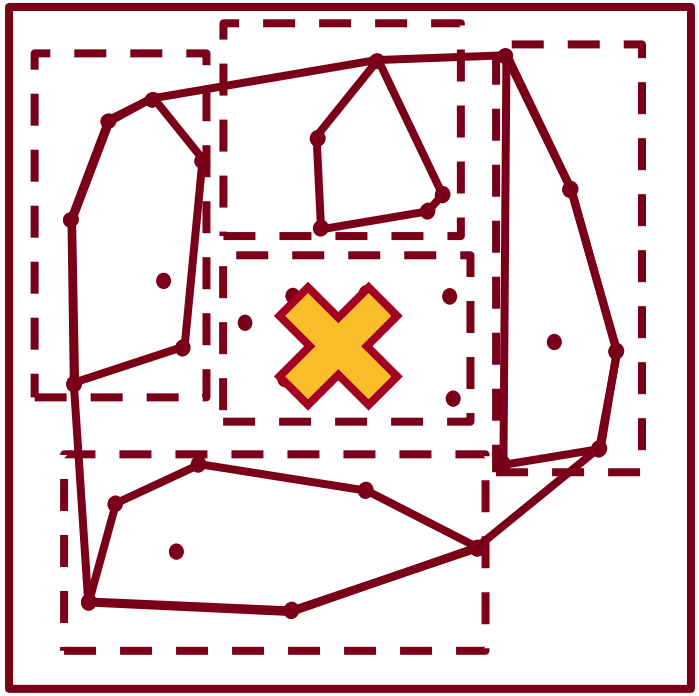

Pruning

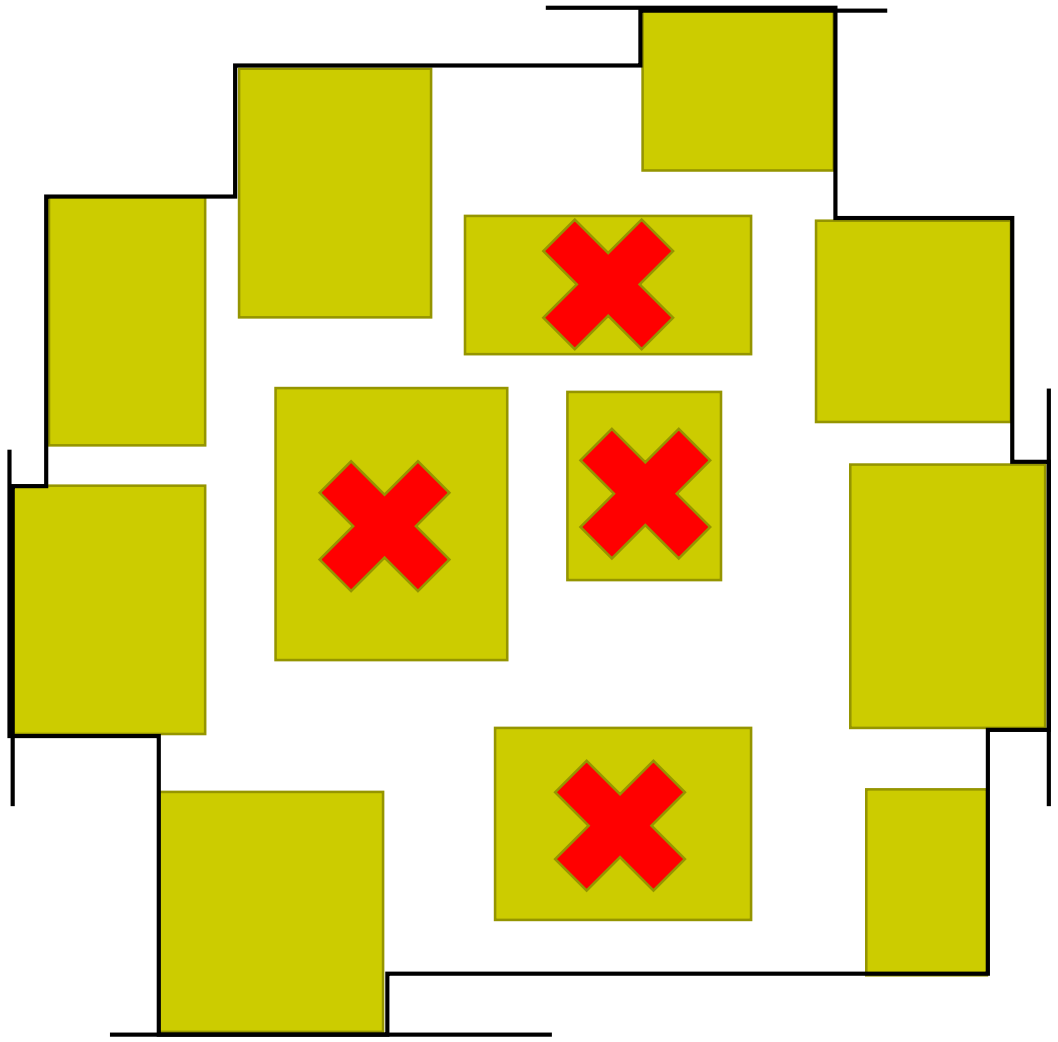

The intersection of the four skyline pruning rules with all directions

Closest Pair

Find the pair of points that have the shortest Euclidean distance

Input

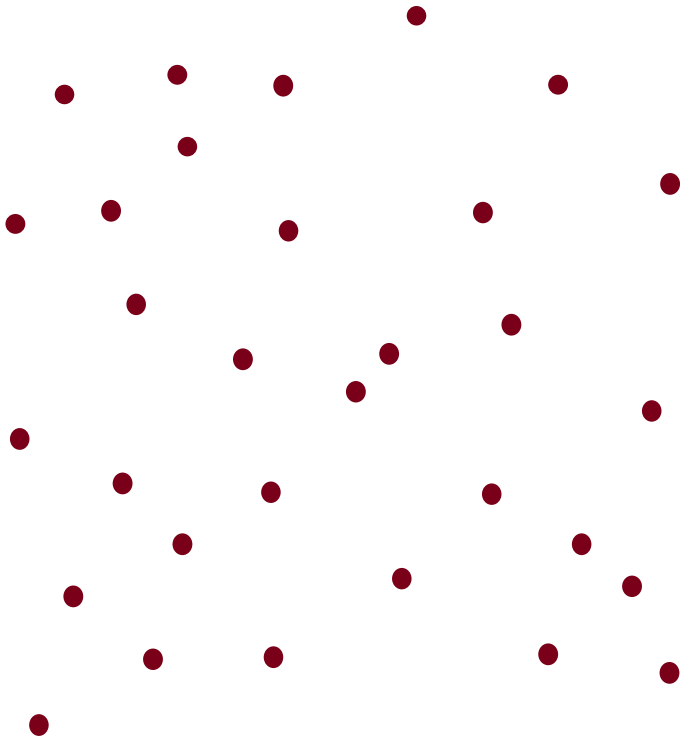

Output

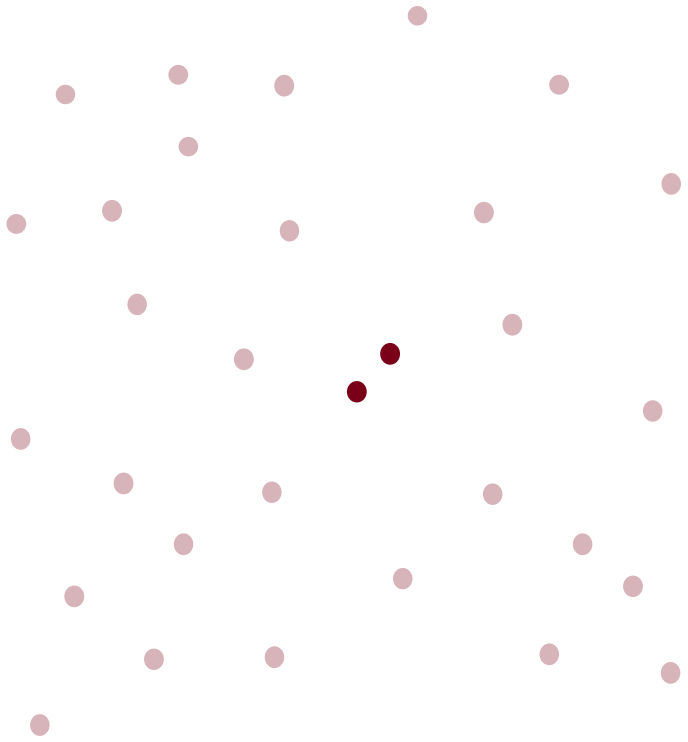

Closest Pair in MapReduce

UCR

Non-spatial partitioning

Spatial partitioning

①Partition

② Local closest pair

③ Global closest pair

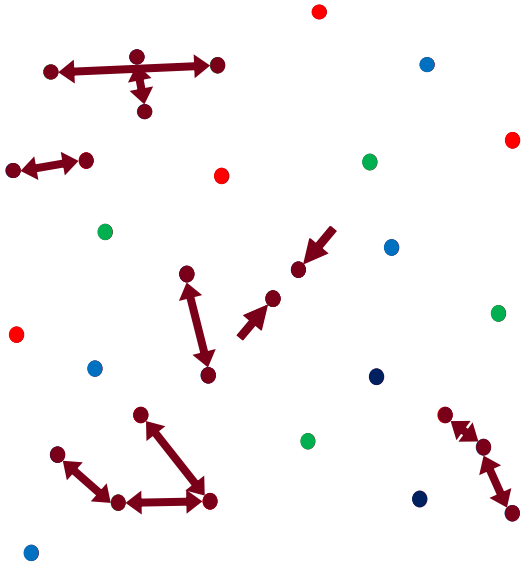

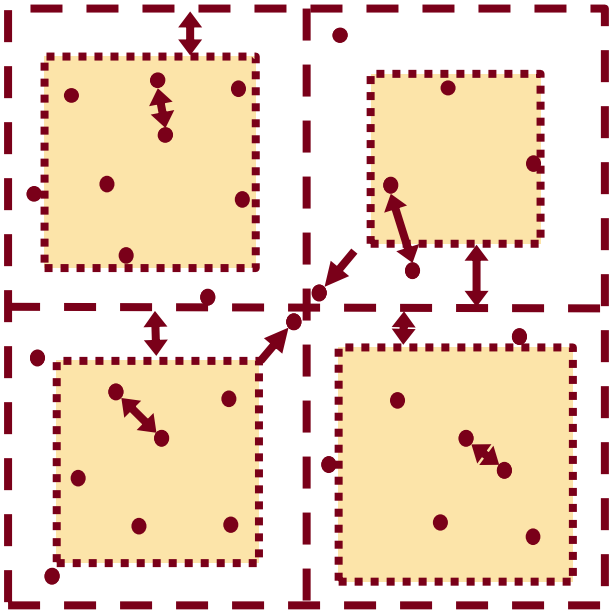

Farthest Pair

Find the pair of points that have the largest Euclidean distance

Input

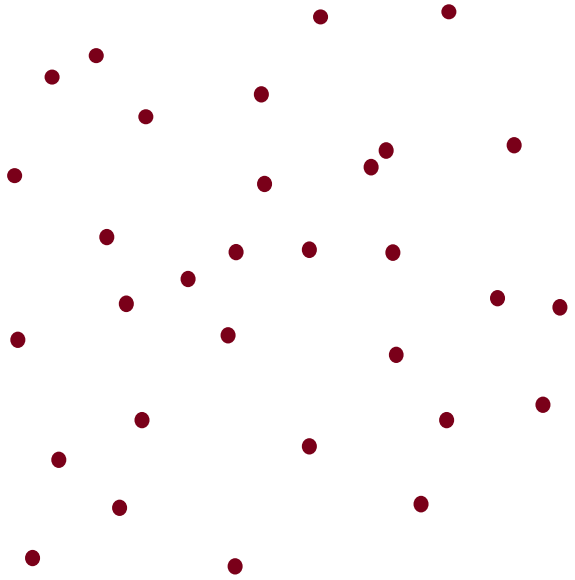

Output

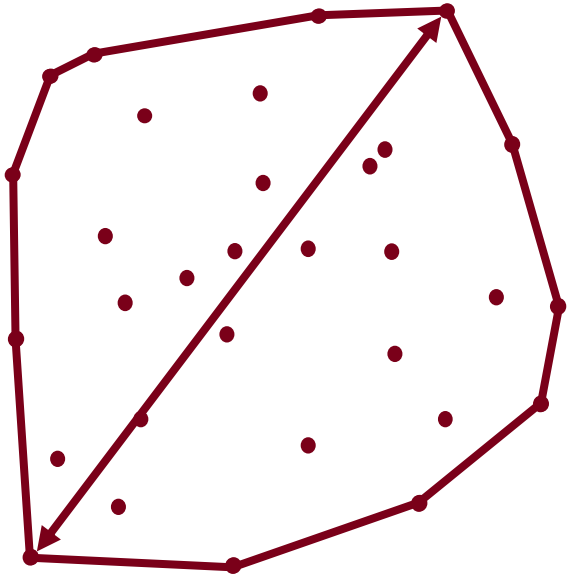

Farthest Pair in MapReduce

①Partition

②Pruning

③Local farthest pair

④Global farthest pair

Non-spatial partitioning

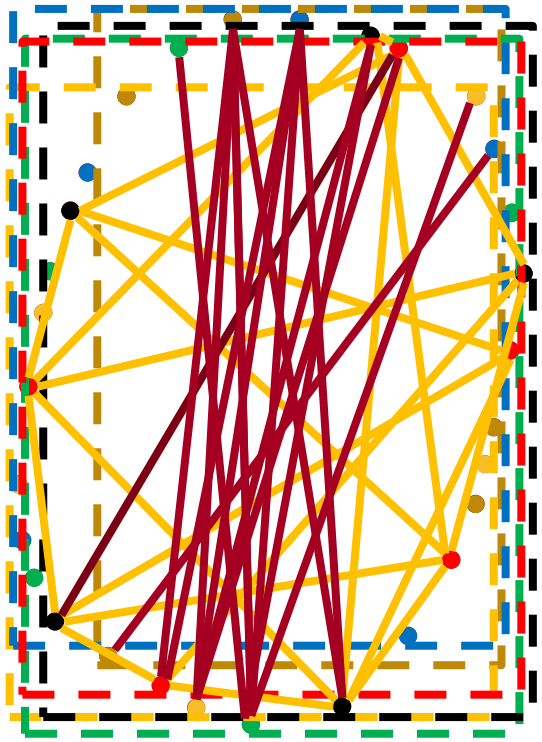

Spatial partitioning

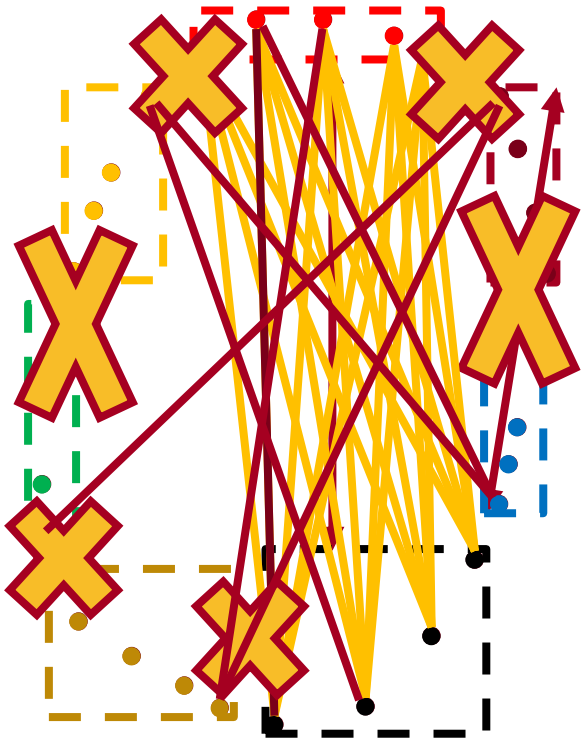

Voronoi Diagram

- **Partitioning**
- Local VD
- Pruning
- Vertical Merge
- Pruning
- Horizontal Merge
- Final output

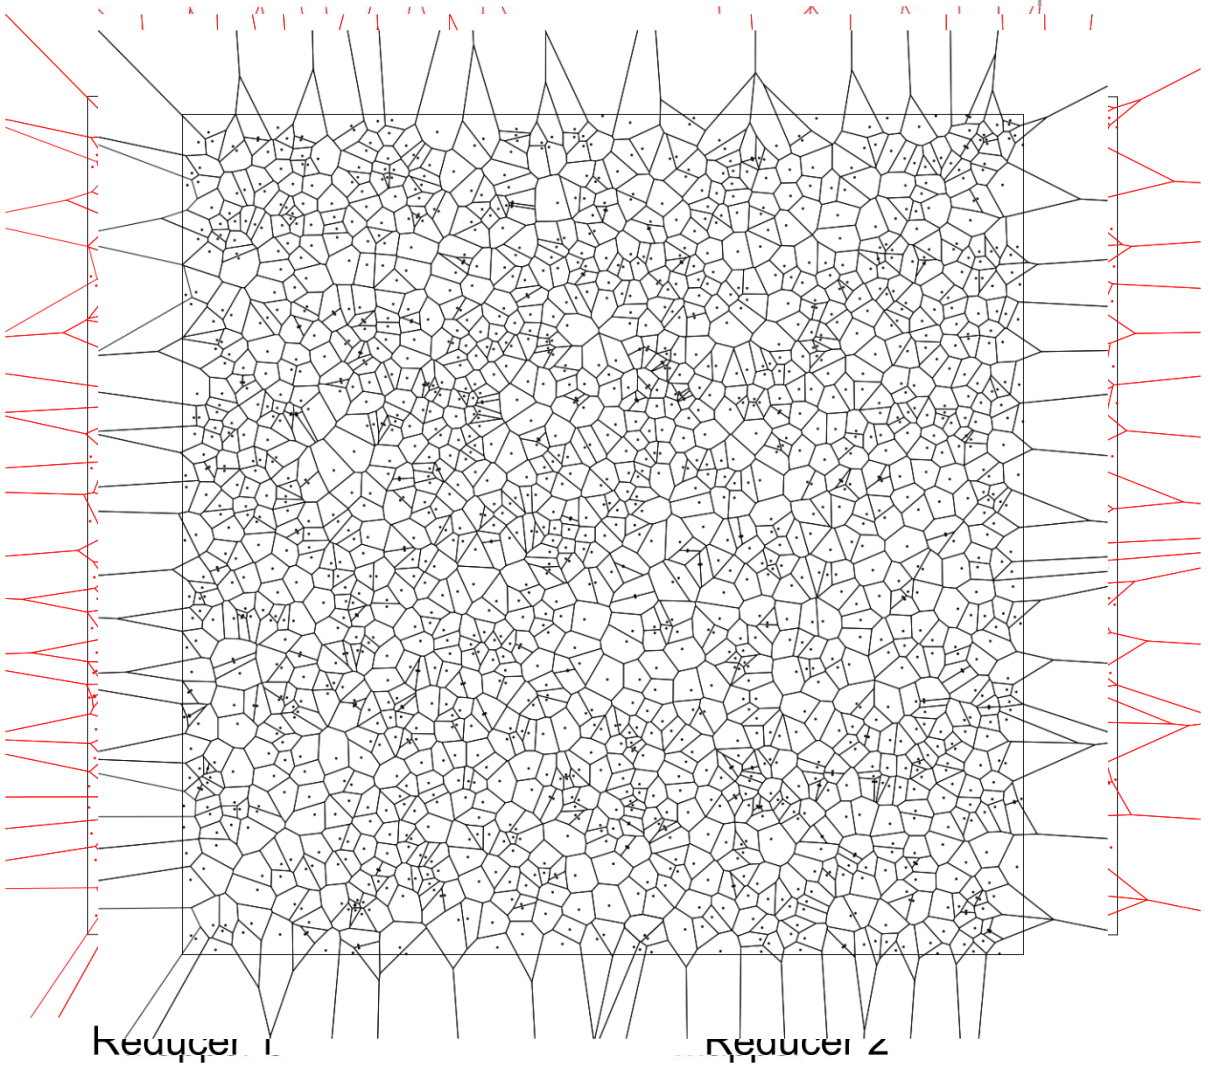

Voronoi Diagram Pruning

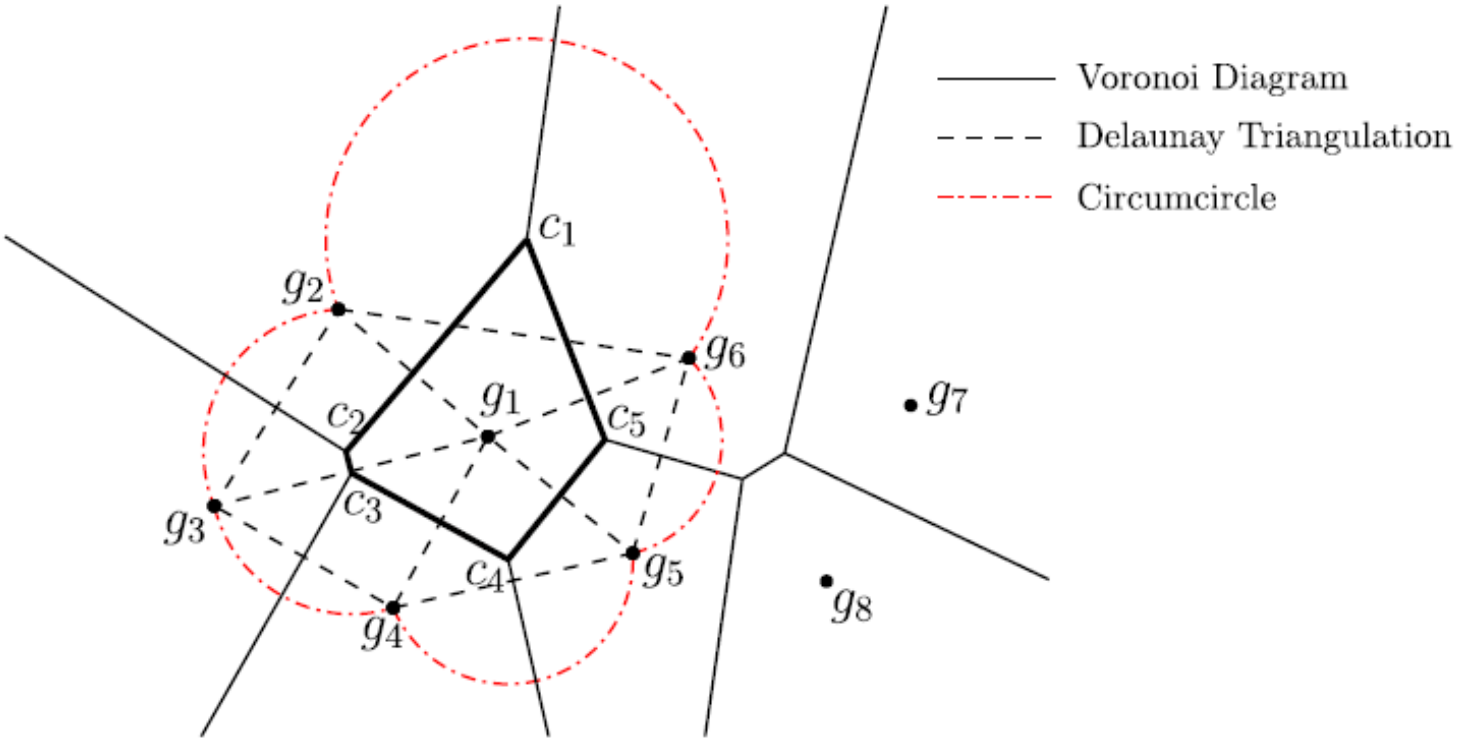

Conclusion

- Computational geometry algorithms can be parallelized
- Both non-spatial and spatial partitioning can be used
- Spatial partitioning enables some pruning techniques
- This method applies to several computational geometry algorithms