Mapping Large Scale Finite Element Computing onto Wafer-Scale Engines

Yishuang Lin, Rongjian Liang, Yaguang Li, Hailiang Hu, Jiang Hu

Texas A&M University, College Station, Texas, USA

Speaker: Yishuang Lin

TEXAS A&M UNIVERSITY

Outline

- Overview
- Problem Background and Formulation
- Algorithm
- Experiment
- Conclusion

Outline

- **Overview**
- Problem Background and Formulation
- Algorithm
- Experiment
- Conclusion

Finite Element Method (FEM) and FE Computing

- FEM
  - Solve partial differential equations
  - Widely used in
    - Structural analysis
    - Heat transfer modeling
    - Fluid dynamics

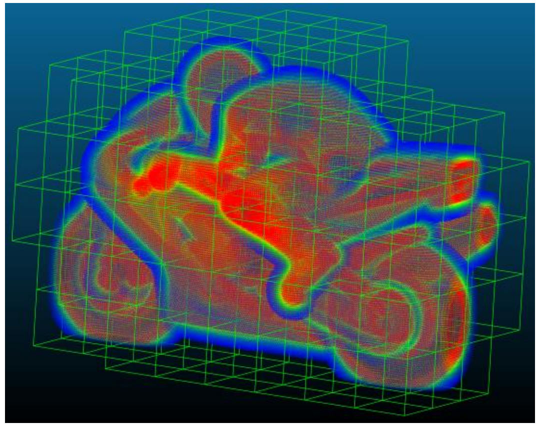

- FE Computing
  - Iterative and slow
  - Cerebras wafer-scale engine (WSE)
    - More than 800K processing elements (PE)
    - Accelerate computing

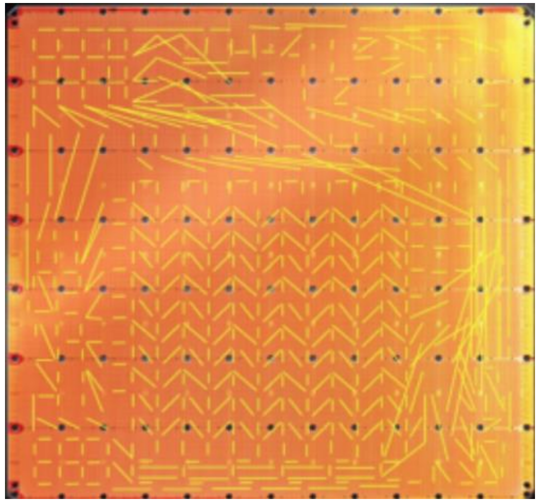

Goal: Mapping FE Computing onto PEs in WSE

- Partitioning
  - Partition object space
  - Constitute computing kernels
  - Maximize computing accuracy

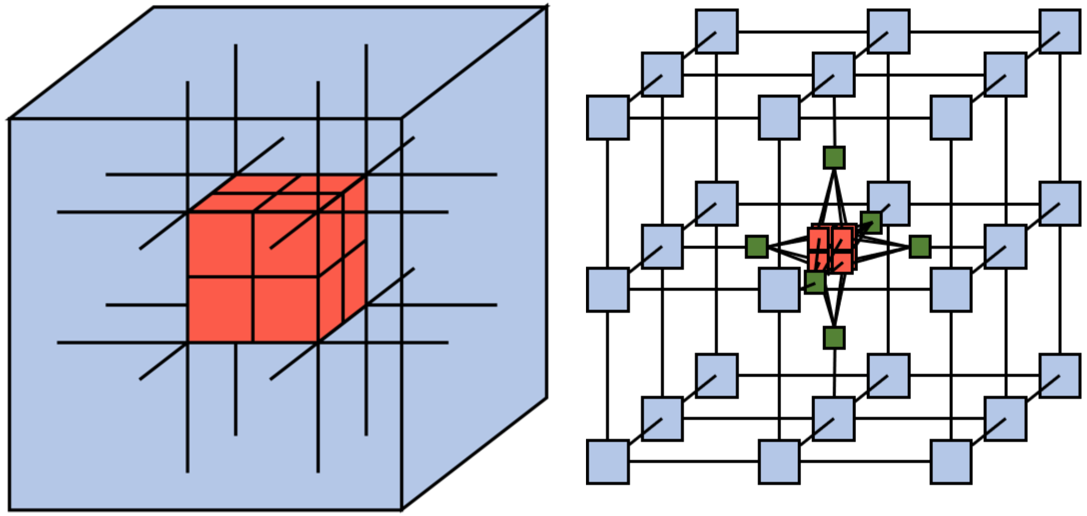

- Placement
  - Place kernel graph nodes
  - Minimize communication cost

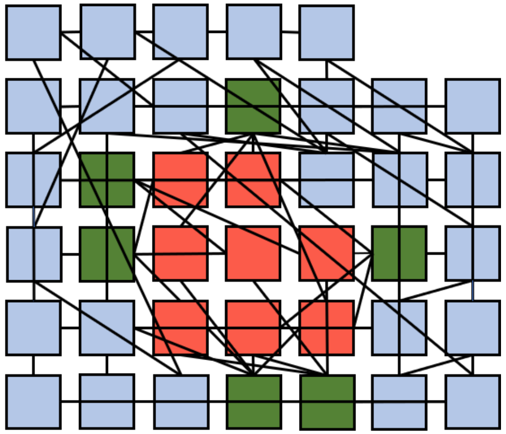

Challenges

- A new problem raised by ISPD 2021 contest
- No previous study
- Limited runtime budget
- Complicated design rules and constraints

Outline

- Overview
- **Problem Background and Formulation**
- Algorithm
- Experiment
- Conclusion

Terminology Definition

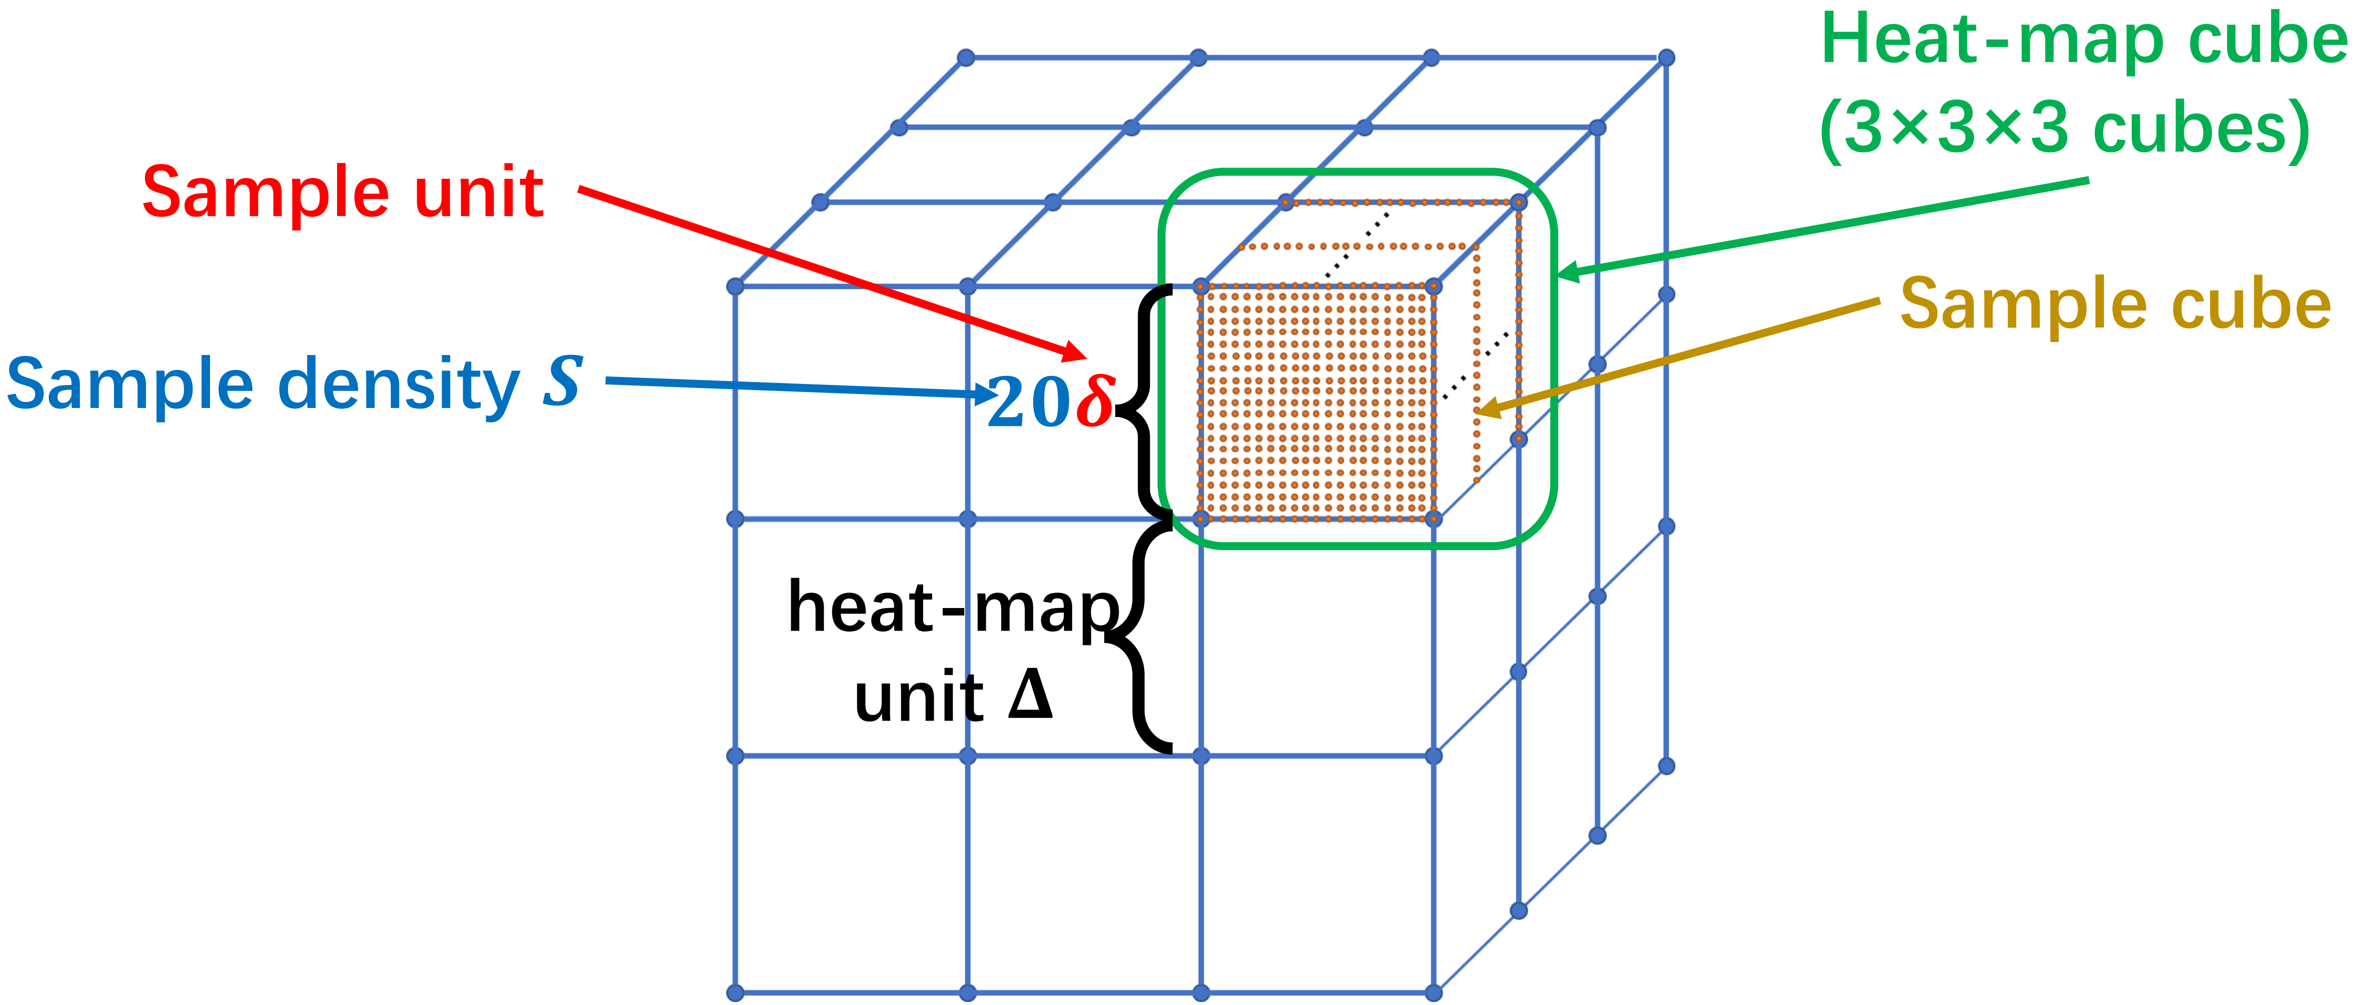

Space Partitioning and Kernel Generation

- Each heat-map cube has a target resolution $\hat{\rho}$ as specification
- Covered by contiguous and non-overlapped kernel cubes
- Computing resolution ρ defines the size of a kernel cube
  - side length= $\frac{10}{\rho}\delta$
  - $\rho = 0.5 \rightarrow \text{sidelength} = 20\delta$
  - $\rho = 1.0 \rightarrow \text{sidelength} = 10\delta$

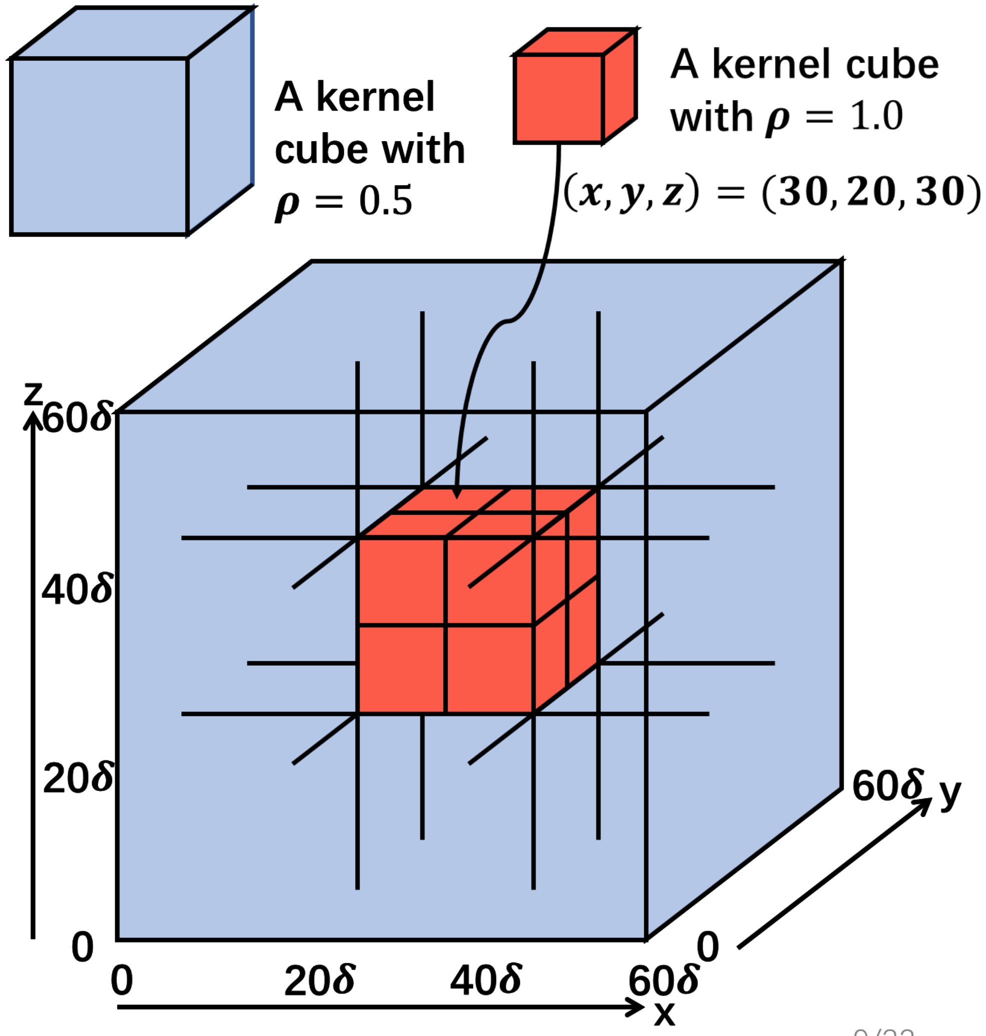

Kernel Graph

- Node set V: kernel cubes and adapters
- Edge set E
  - Two same adjacent cubes
  - An adapter and a low resolution cube
  - An adapter and a high resolution cube

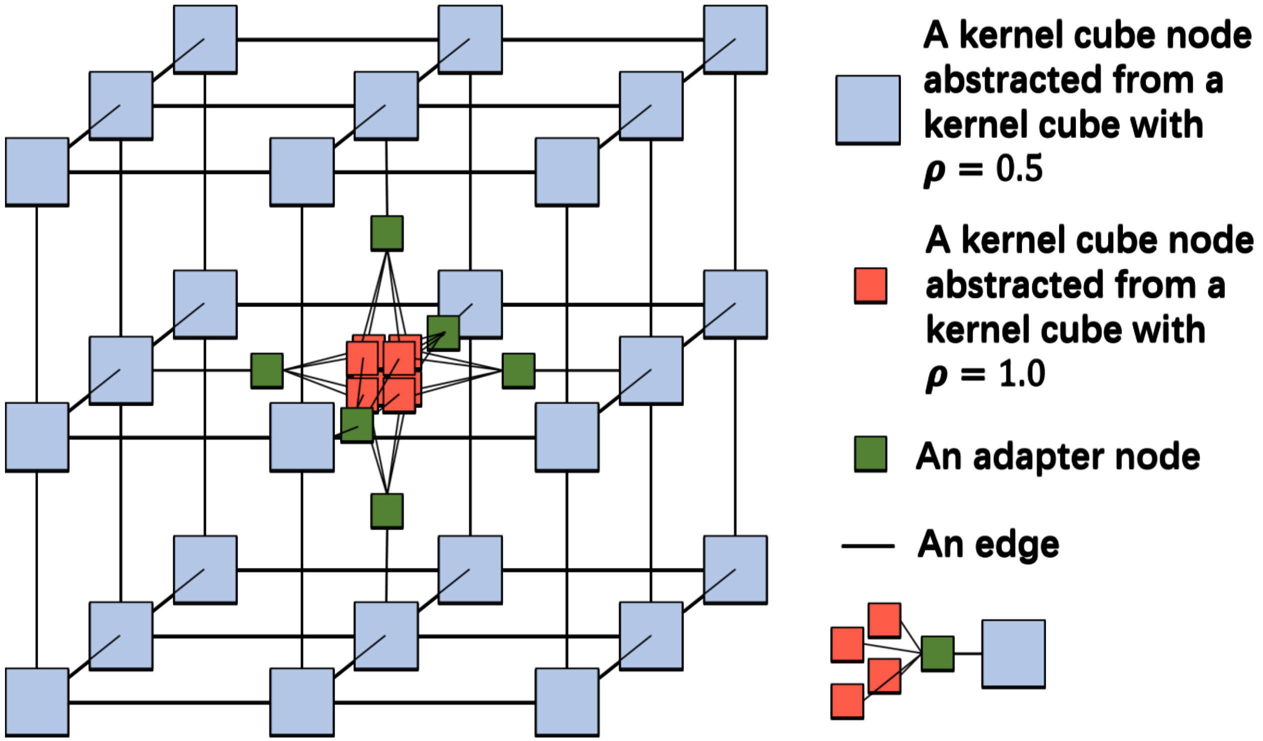

Kernel Placement

- Each node in V is mapped onto one PE
- Minimize communication cost for each edge in E

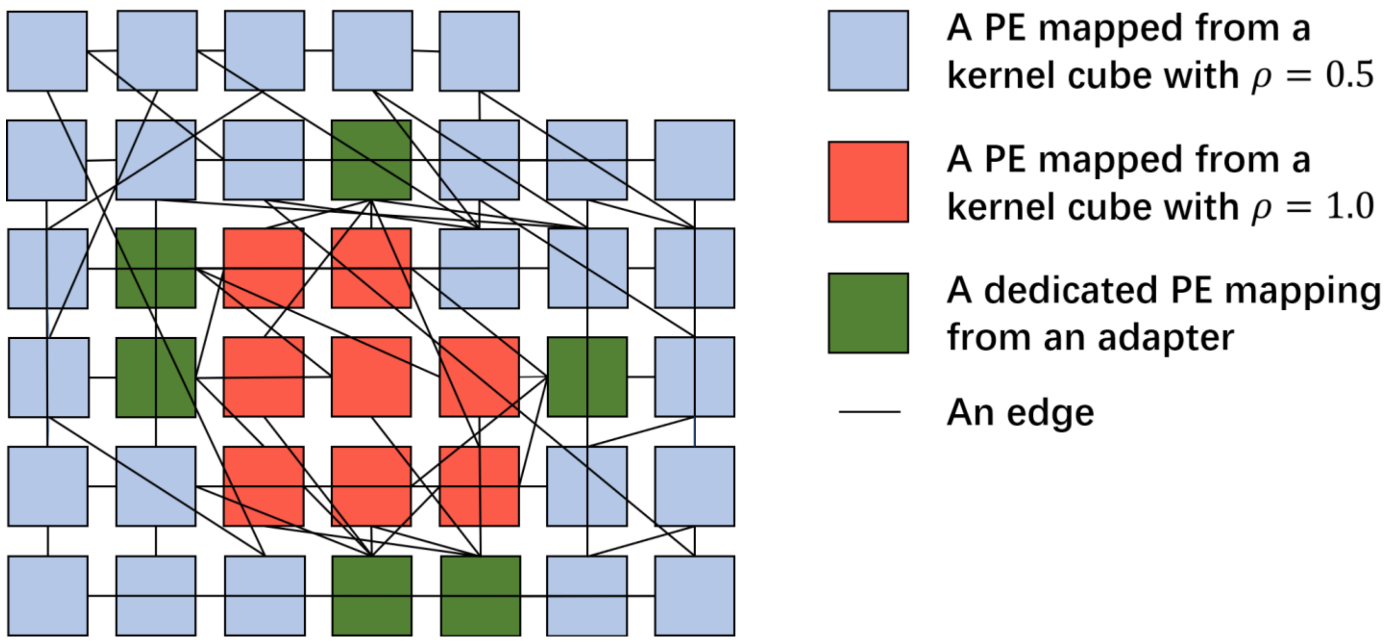

Accuracy score

Accuracy score of a kernel cube $\boldsymbol{k}$

Weighted target resolution

$$F_r(k) = \frac{\frac{\sum_{\tau \in k} \hat{\rho}(\tau)}{|k|}}{\max_{k \in K}\left(1, \max_{k \in K}\left(\max_{\tau \in k} \frac{\hat{\rho}(\tau)}{\rho(k)}\right)\right) \cdot \rho(k)}$$

Normalization factor γ

K: all kernel cubes

Overall accuracy score

$$F_r = \frac{\sum_{k \in K} F_r(k)}{M}$$

#PEs

Connectivity score

- Two connected nodes u and v are placed at (x_u, y_u) and (x_v, y_v)
- Connectivity score

$$F_w = \left(\frac{100M}{\sum_{(u,v)\in E} \left((|x_u - x_v| + |y_u - y_v|)^{1.5} \right)} \right)^{\frac{2}{3}}$$

Manhattan distance

- The overall score

$$F = \min(F_r, F_w)$$

Problem Formulation

- Input:
  - A $W \times L \times H$ target resolution matrix
  - Size of the 2D PE array
- Output:
  - Sample density S
  - Partitioning and placement solution
- Goal: Maximize overall score
- Constraints:
  - Computing resolution $\rho \in \{2^{-i} | i = 0, 1, 2\}$
  - Resolution ratio of adjacent cubes is in $\{0.5, 1.0, 2.0\}$
  - Neighbor cubes in one face have the same size
  - …

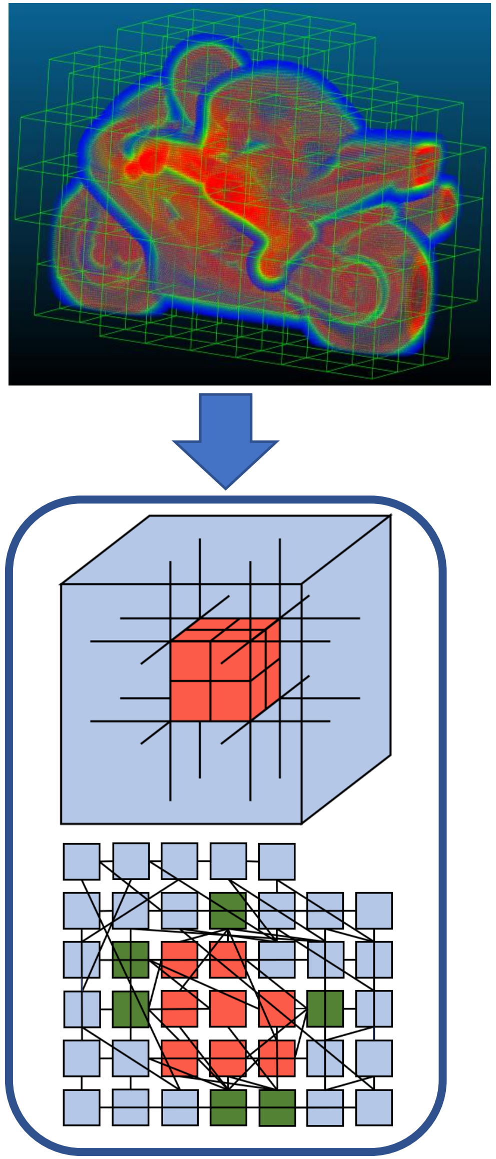

Outline

- Overview
- Problem Background and Formulation
- Algorithm
- Experiment
- Conclusion

Overview of Algorithm

- **Partitioning:**
  - Maximize accuracy score
  - Meet all constraints
- **Placement:**
  - Geometric bisection placement
    - Fast and complement ePlace
  - ePlace: Standard cell placement
  - Refine connectivity score

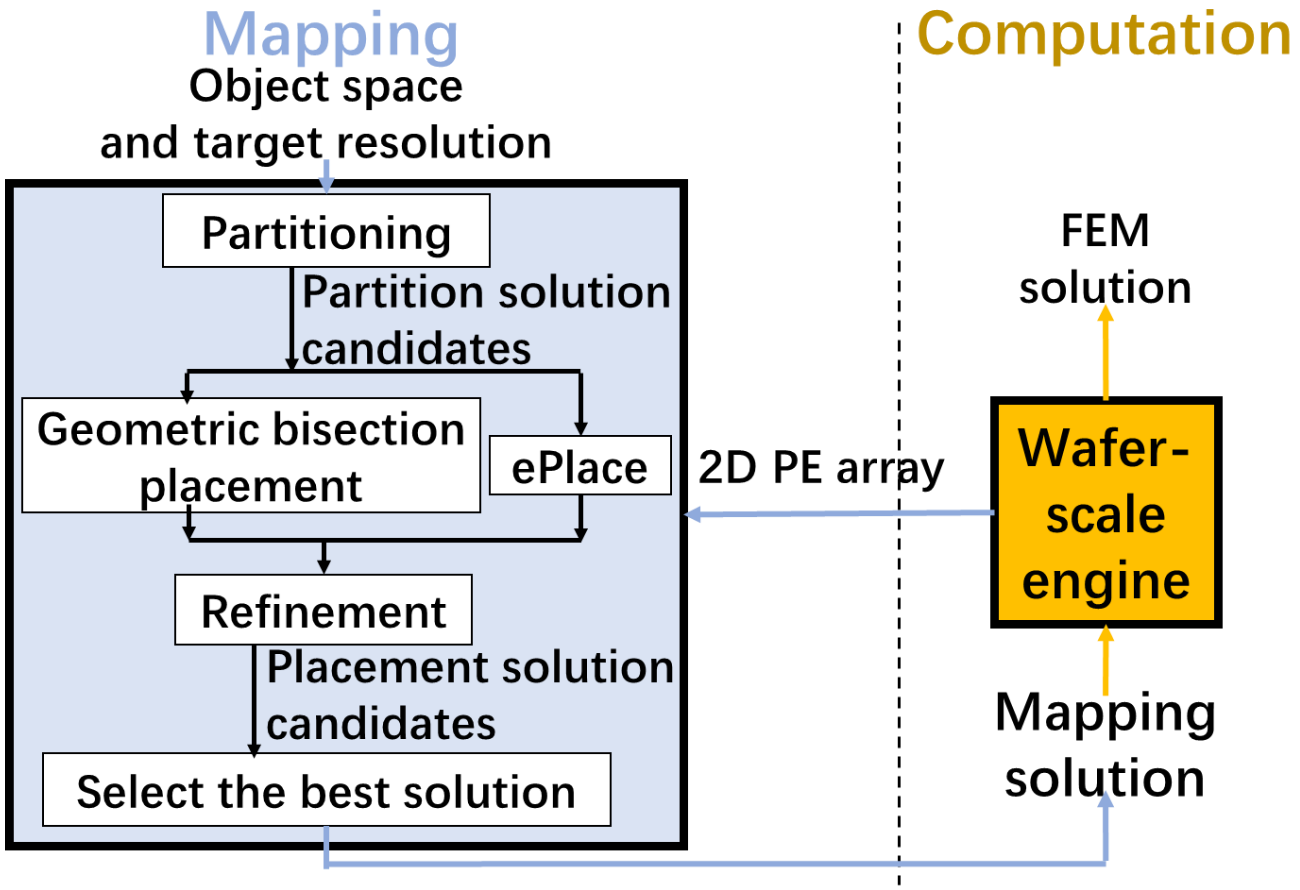

Partitioning: GEB (Greedy heuristic in Enumerated Binary search)

- Three layers:
  - Enumerate different values for sample density $\boldsymbol{S}$
  - Binary search on target normalization factor $\hat{\gamma}$
  - Greedy heuristic to generate partition candidates with specific S and $\hat{\gamma}$

Enumerate Sample Density

- Minimum sample density

Least common multiplier

$$S_{min} = \frac{\mathbf{lcm}(\mathbf{gcd}(W,L,H),10)}{\mathbf{gcd}(W,L,H)}$$

Greatest common divider

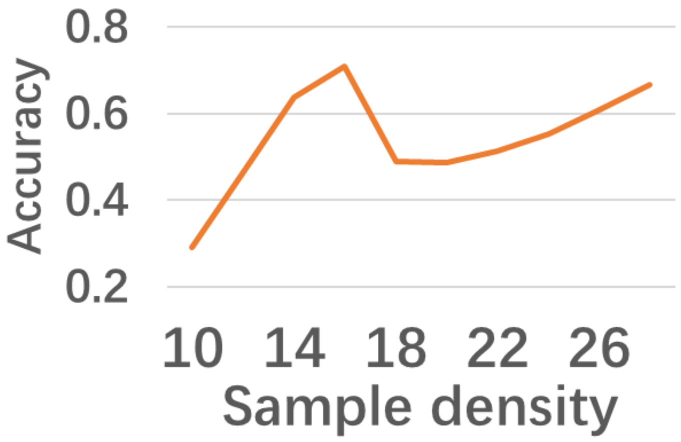

- Sample density set
  - $\mathcal{S} = \{S_{min}, 2S_{min}, 3S_{min}, \dots, S_{max}\}$
  - S_{max} is the upper bound which does not cause memory issue.

Binary Search over Target Normalization Factor

- Target normalization factor $\hat{\gamma}$:
  The maximum γ allowed in the most inner layer greedy heuristic.
- Initial lower bound $\underline{\gamma} = 0$
- Initial upper bound

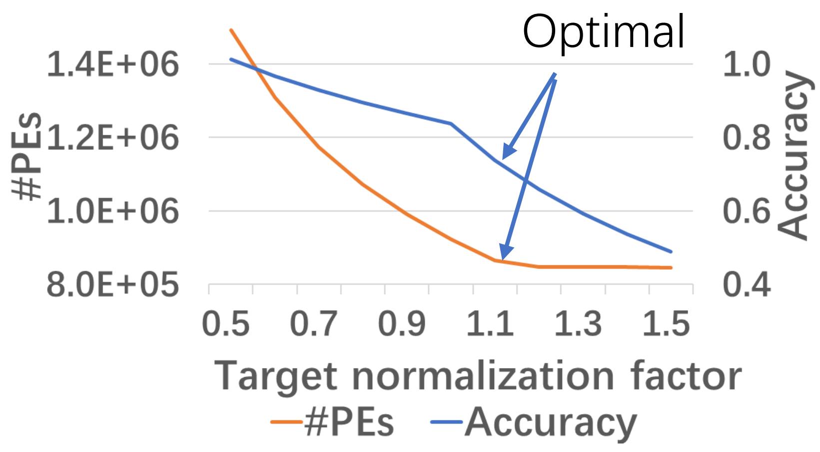

$$\overline{\gamma} = \left\lceil \left(\frac{W \times L \times H \times S^3}{1000 \times M} \right)^{\frac{1}{3}} \right\rceil$$

Restore 3D ratio to 1D

Number of sample cubes

Number of points all PEs can compute

Greedy Heuristic to Generate Partition Candidates

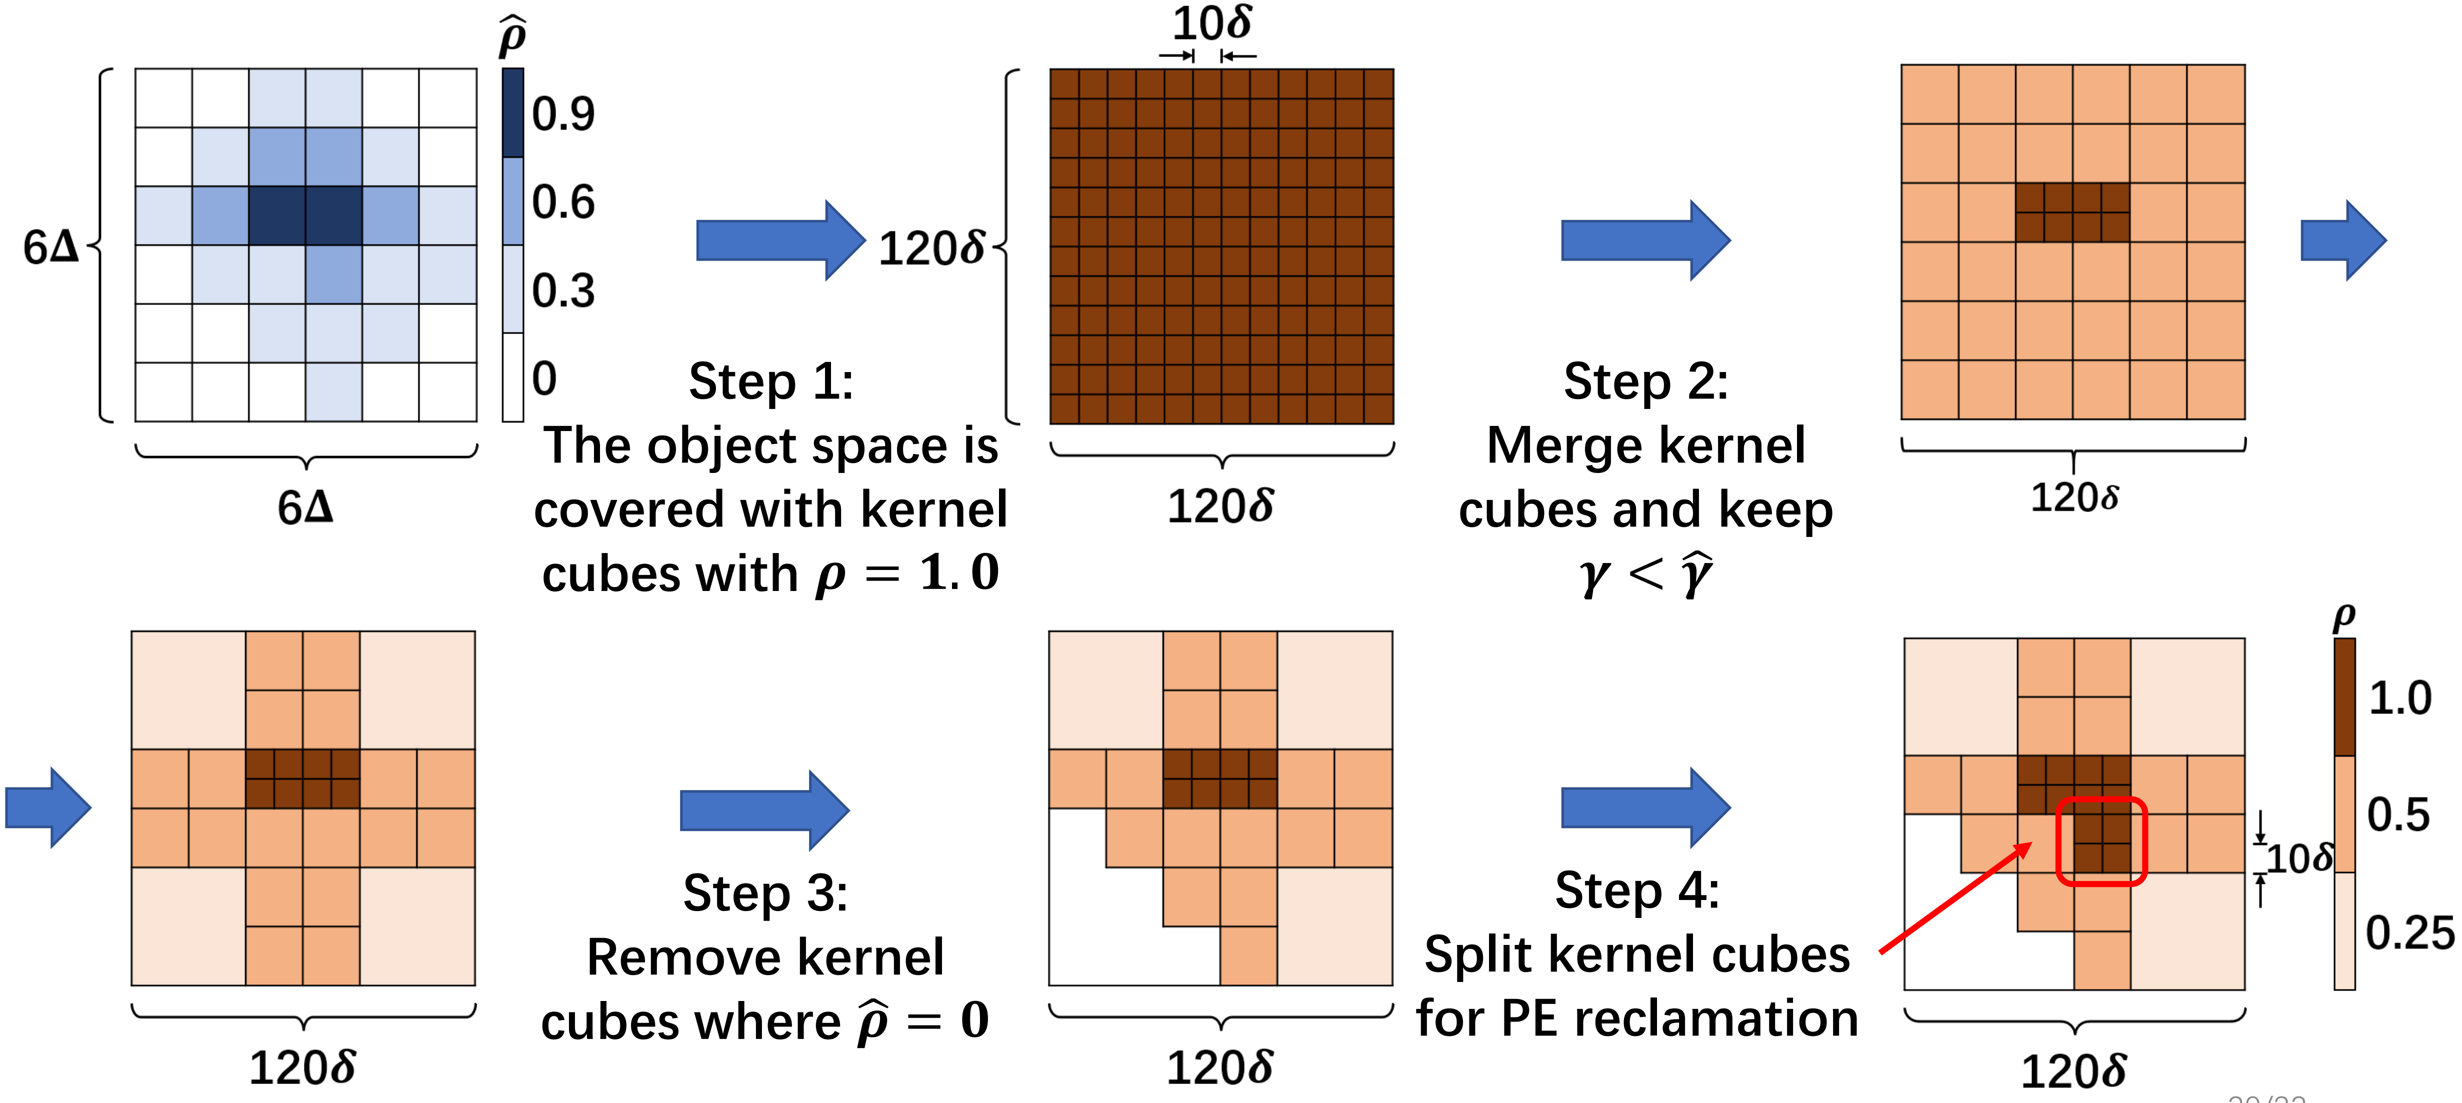

Placement – Geometric Bisection Placement

- **Input:**
  - A set of kernel cubes and adapters
  - Coordinate of the corresponding placement area
  - Size of the corresponding placement area
  - Bisection direction
- **Output: Placement coordinates**
- **Recursion stop criteria: Size of set small then threshold, go to node pillar placement**

Geometric Bisection Placement - Recursive Bisection

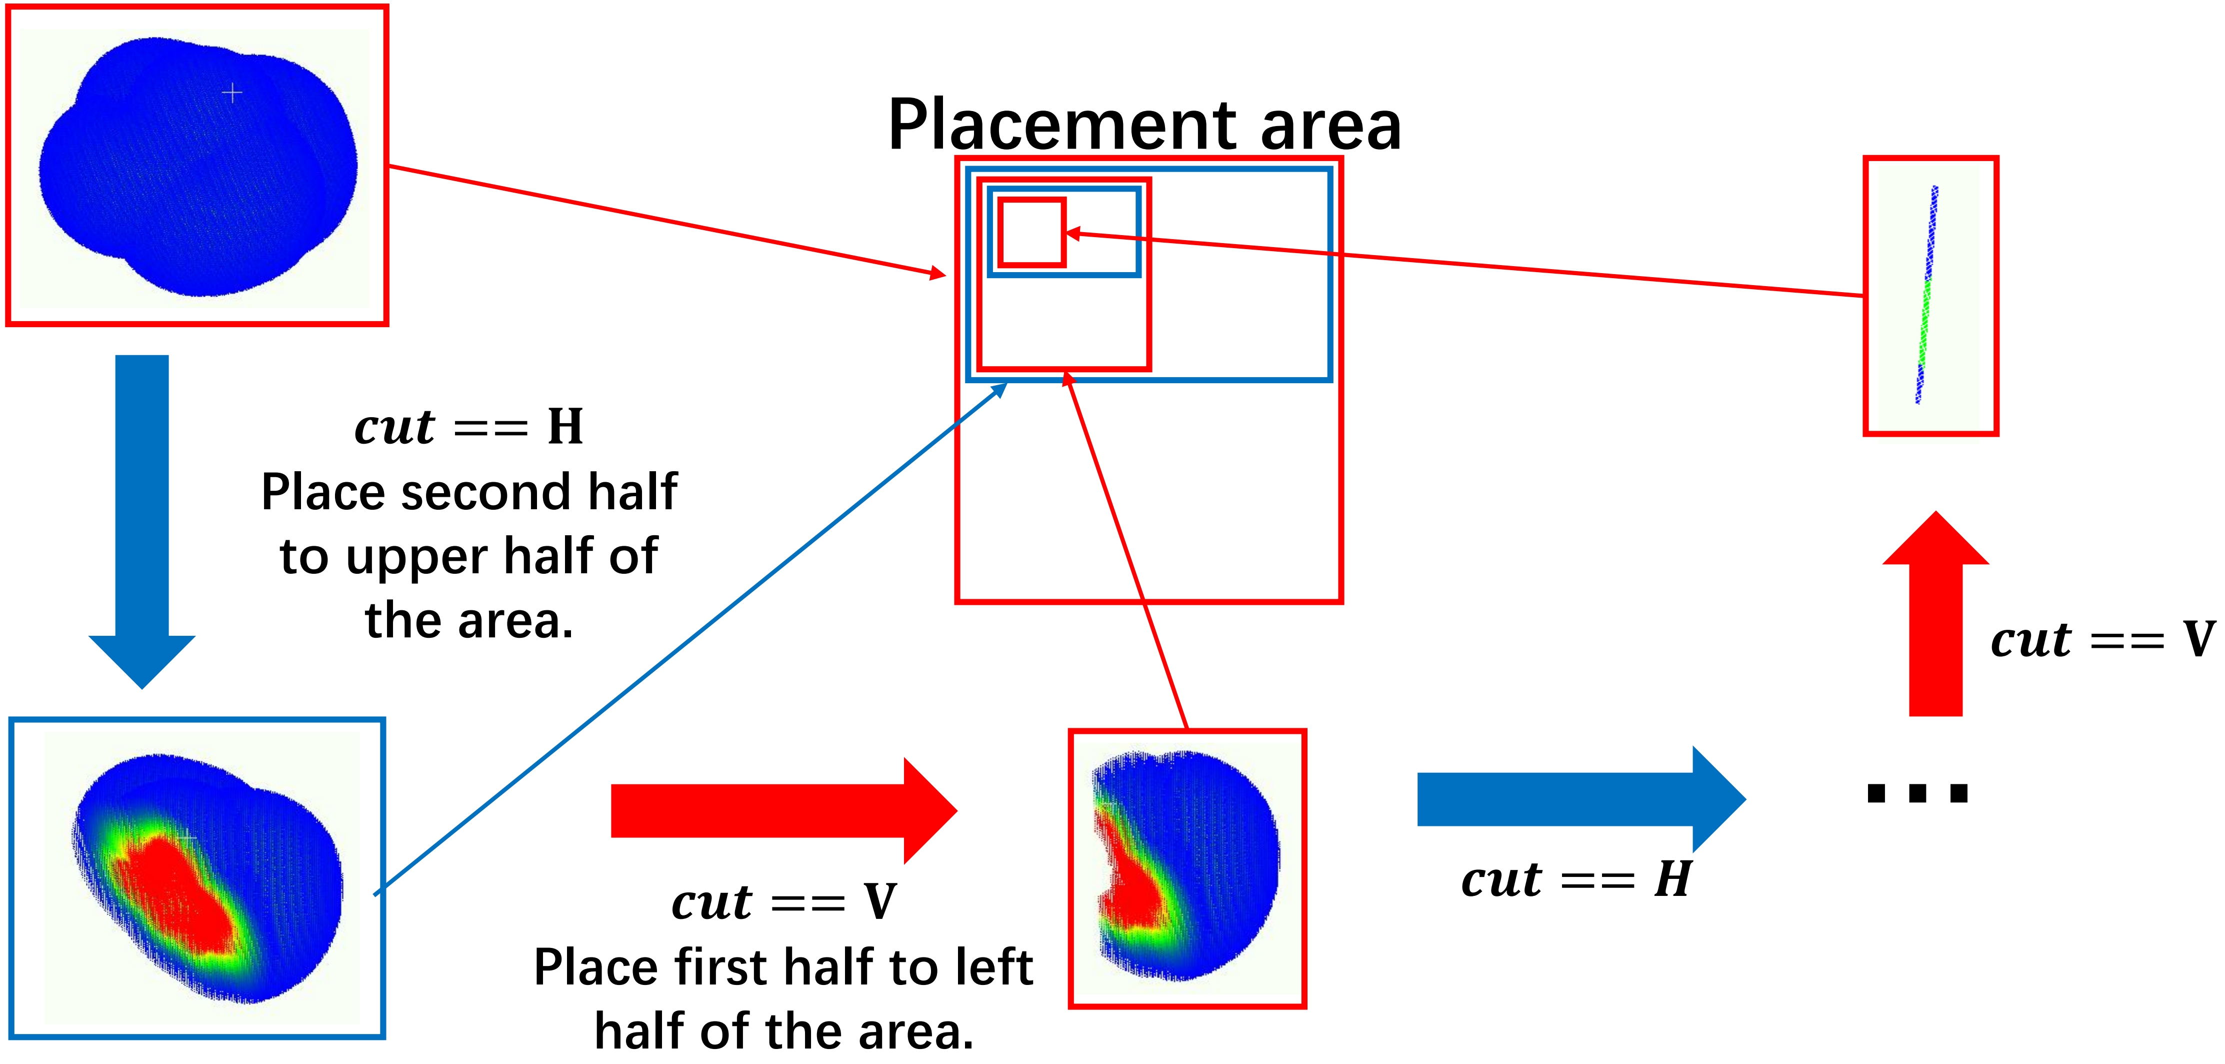

Geometric Bisection Placement – Node Pillar Placement

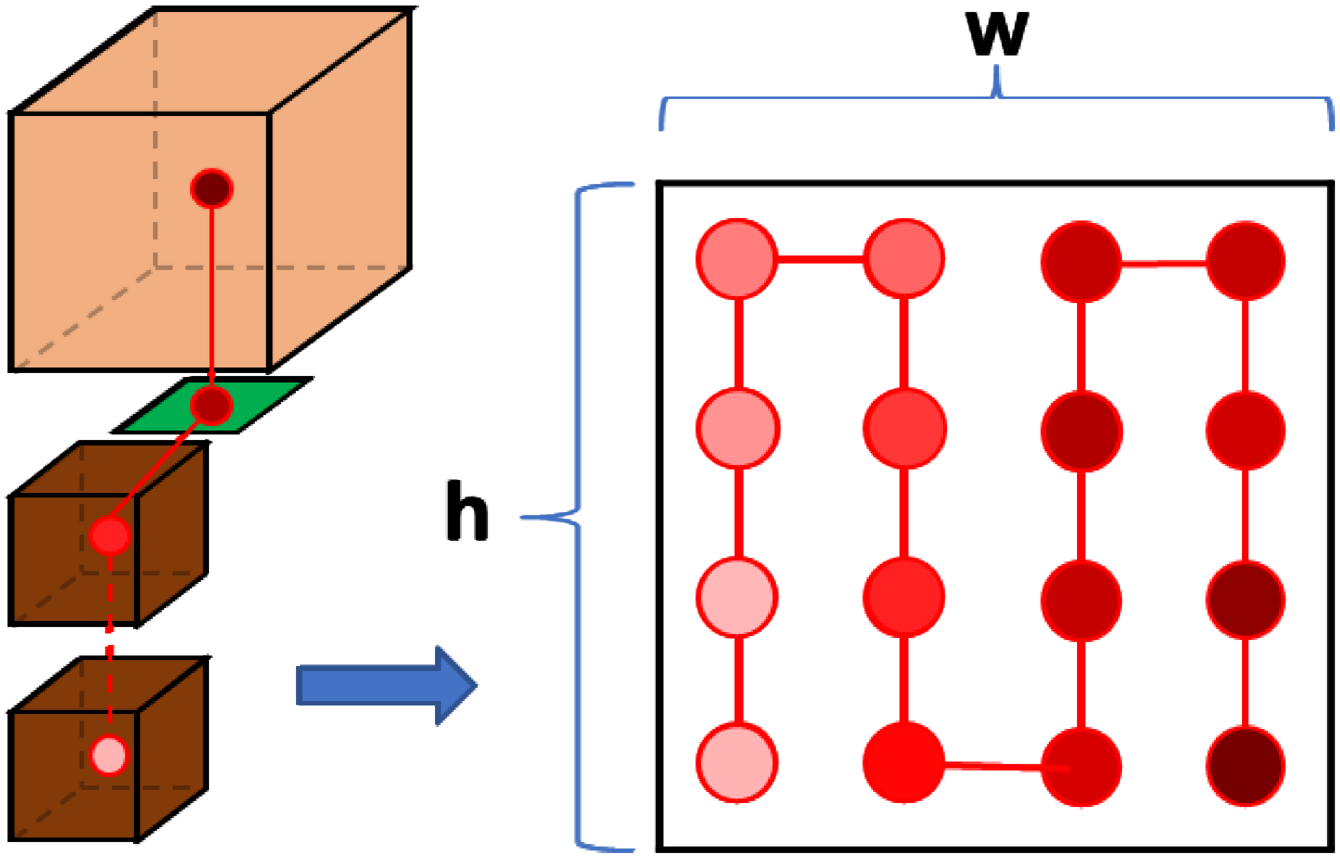

Placement - ePlace

- ePlace is used in global placement
- Polar is used for legalization after global placement

J. Lu et al., "ePlace-MS: Electrostatics-based placement for mixed-size circuits," IEEE TCAD, vol. 34, no. 5, pp. 685–698, 2015.
T. Lin et al., "Polar: Placement based on novel rough legalization and refinement," in Proc. ICCAD, 2013, pp. 357–362.

Placement - Refinement

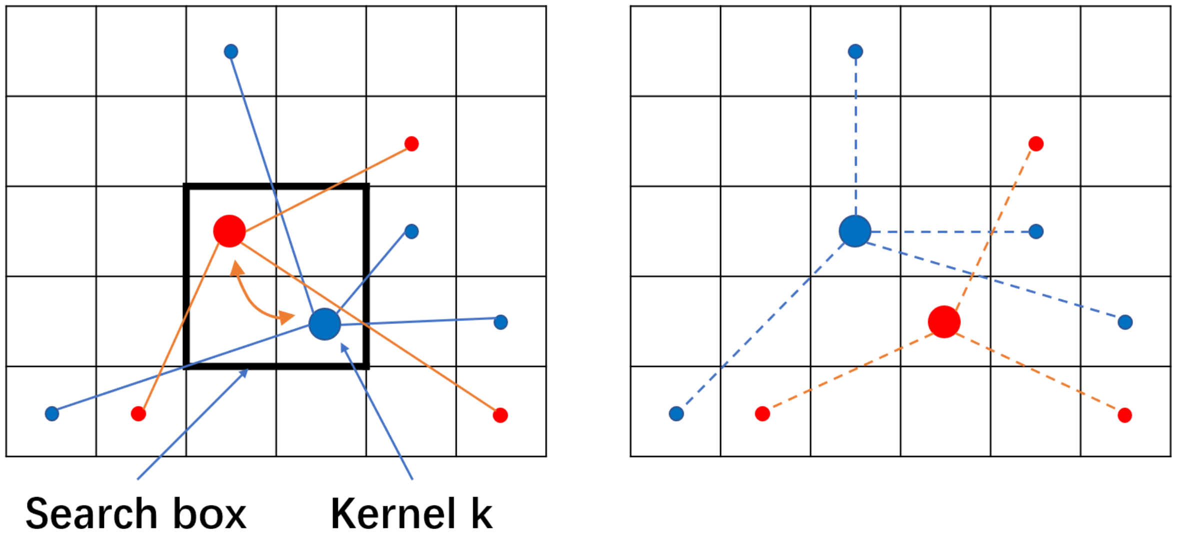

- Find a search box
- Swap with other nodes
- Accept the swap with the best score

Outline

- Overview
- Problem Background and Formulation
- Algorithm
- **Experiment**
- Conclusion

Test Cases

Test case	Volume
Bullet	119 × 119 × 119
Flange	207 × 207 × 207
Propeller Tip	115 × 115 × 115
Motorbike	203 × 203 × 203

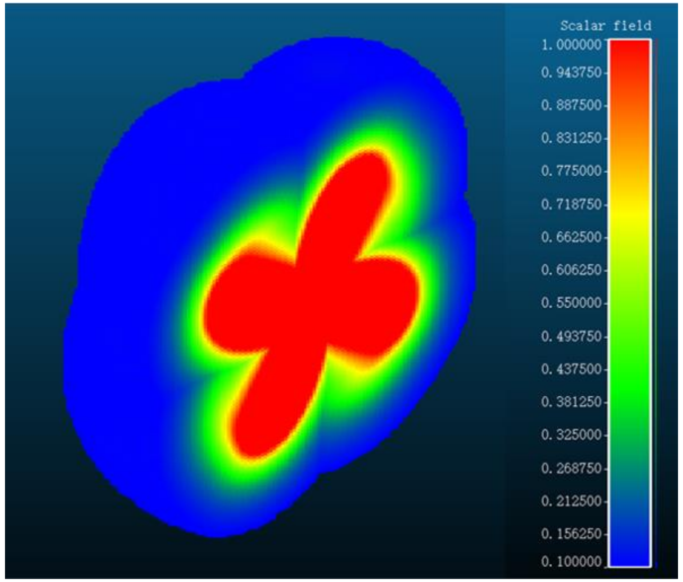

Target resolution of Propeller Tip

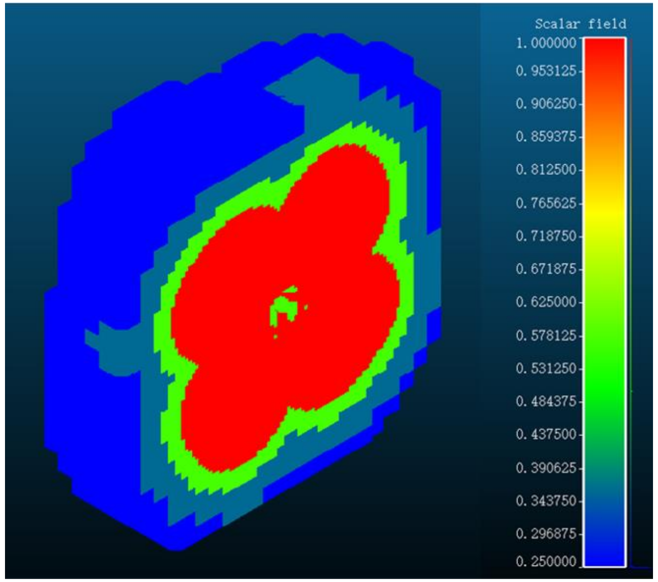

Computing resolution of Propeller Tip

Score Comparison with Other Teams and Baseline

Design	2nd Place Team			3rd Place Team			Naïve			Ours		
	Accuracy	Connectivity	Score	Accuracy	Connectivity	Score	Accuracy	Connectivity	Score	Accuracy	Connectivity	Score
Bullet	0.5194	0.6771	0.5194	0.4756	0.5063	0.4756	0.4564	0.6345	0.4564	0.6890	0.7952	0.6890
Flange	0.4689	0.6994	0.4689	0.3883	0.7599	0.3883	0.3909	0.6653	0.3909	0.4741	0.6898	0.4741
Propeller Tip	0.7089	0.8102	0.7089	0.7088	0.4905	0.4905	0.4938	0.7887	0.4938	0.7089	0.7409	0.7089
Motorbike	0.6670	0.8764	0.6670	0.6516	0.6489	0.6489	0.5195	0.7406	0.5195	0.6647	0.9577	0.6647
Total			2.3642			2.0033			1.8606			2.5367

- 7% better on overall score than the 2nd place team
- 27% better on overall score than the 3rd place team

Runtime Analysis

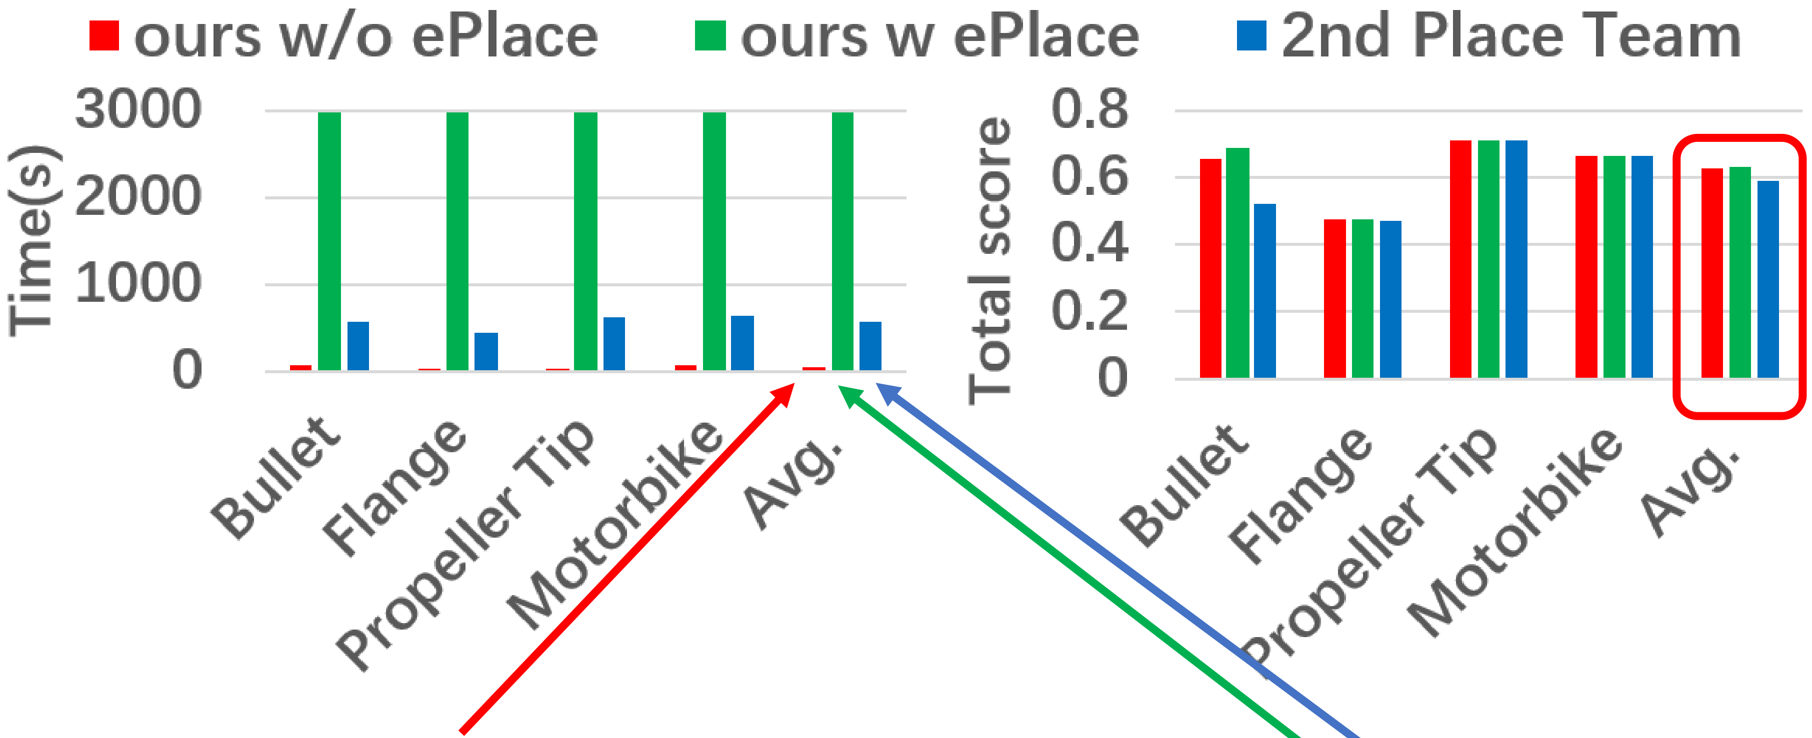

- Our method w/o ePlace is 60× faster than our method with ePlace
- Our method w/o ePlace is 11× faster than the 2nd place team

Outline

- Overview
- Problem Background and Formulation
- Algorithm
- Experiment
- **Conclusion**

Conclusion

- A solution to a new problem brought by ISPD 2021 contest
- Map finite element computing onto a wafer scale engine
- Partitioning: a greedy heuristic in enumerated binary search technique
- Placement: geometric bisection placement and ePlace followed by refinement

Q&A

Thanks and Questions?