

IEEE/ACM

2022 INTERNATIONAL
CONFERENCE ON
COMPUTER-AIDED
DESIGN

**IC
CAD**

41st Edition

X-Check: GPU-Accelerated Design Rule Checking via Parallel Sweep-line Algorithms

Zhuolun He¹, Yuzhe Ma², Bei Yu¹

¹ The Chinese University of Hong Kong

² HKUST(GZ)

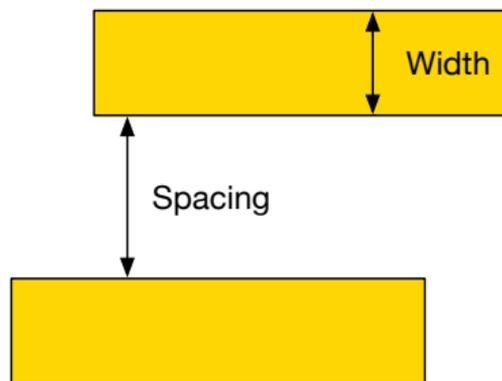
Sept. 14, 2022



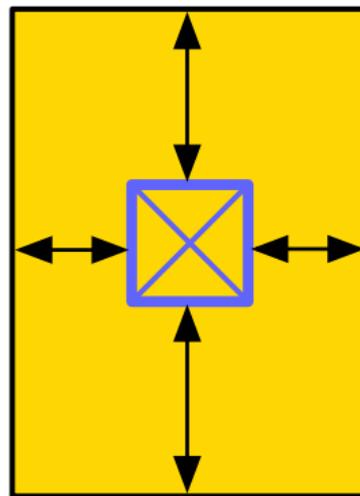
- ① Background and Motivation
- ② Algorithm: Parallel Vertical Sweeping
- ③ GPU Implementation
- ④ Experimental Results

Background and Motivation

DRC: to ensure the layout does not violate geometric constraints



(a)



(b)

Typical rules: (a) *width* and *spacing* rules in a metal layer; (b) *enclosing* rule between a metal layer and a via layer.

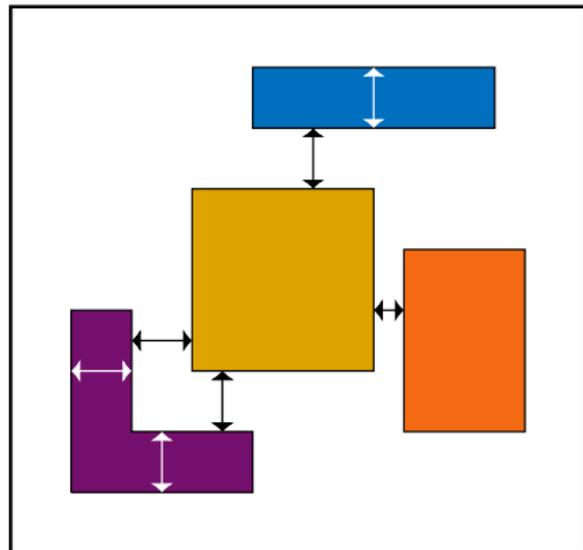
- Design rule number explosion in advanced technology
- Many classic parallel algorithms do not scale beyond a few CPU cores [G. Guo+, DAC'21]
 - data parallelism
 - task parallelism
- GPUs have demonstrated potential in EDA tool acceleration

- to cast a design automation problem into another problem solvable by current tools/infrastructure
 - DreamPlace (analytical placement → NN training) [Y. Lin+, DAC'19]
 - GATSPI (gate-level simulation → graph manipulation) [Y. Zhang+, DAC'22]
 - FastGR (batched net routing ordering → task scheduling) [S. Liu+, DATE'22]
- to design novel GPU-friendly computation kernels for some critical tasks in the design flow
 - Placement [Z. Guo+, DAC'21]
 - GAMER (maze routing) [S. Lin+, ICCAD'21]
 - STA [Z. Guo+, ICCAD'20]

Our work is closer to the second methodology.

Problem (Distance Check (informal))

- *Layout: a set of axis-parallel polygonal objects*
- *Distance rule: any two edges **must not be closer than** a predefined minimal distance*
- *Distance violation: a pair of edges in the layout that violate the distance rule*
- *Our task: report all the distance violations*




(We only consider horizontal edges.)

Problem (Distance Check)

Given a set \mathcal{H} of horizontal segments in \mathbb{R}^2 , report the segment pairs from \mathcal{H}^2 whose horizontal projection is nonempty, and vertical distance is smaller than δ .

Formally, we want to report:

$$\begin{aligned} & \{([l_1, r_1] \times y_1, [l_2, r_2] \times y_2) \in \mathcal{H}^2\} \\ \text{s.t. } & [l_1, r_1] \cap [l_2, r_2] \neq \emptyset, |y_1 - y_2| < \delta \end{aligned}$$



- ① Sort segment endpoints \mathcal{P} by ascending x -coordinates
- ② Initialize an empty BST \mathcal{S} (using y -coordinates as keys)
- ③ Scan endpoints from left to right
 - ① If p is the left endpoint of a segment $h = [l, r] \times y$
 - ① Range query \mathcal{S} for $[y - \delta, y + \delta]$
 - ② Report the corresponding segment pairs
 - ③ Insert h to \mathcal{S}
 - ② Otherwise (i.e., right endpoint)
 - ① Delete h from \mathcal{S}

Complexity: $O(n \log n + k)$, optimal:

- element uniqueness problem (lower bounded by $\Omega(n \log n)$) reducible to it
- we need $\Omega(k)$ time to report all the violations

Algorithm: Parallel Vertical Sweeping

Prefix Structure

$$a[] = (4, 5, 3, 6, 2, 5, 1, 1, 0)$$

Prefix sums:

$$s = (4, 9, 12, 18, 20, 25, 26, 27, 27)$$

Can we do it in parallel?

$$a[] = (4, 5, 3, 6, 2, 5, 1, 1, 0)$$

Suppose we have 3 threads.

- 1 **Batching:** each thread computes sums of 3 consecutive elements.

$$s = (?, ?, 12, ?, ?, 13, ?, ?, 2)$$

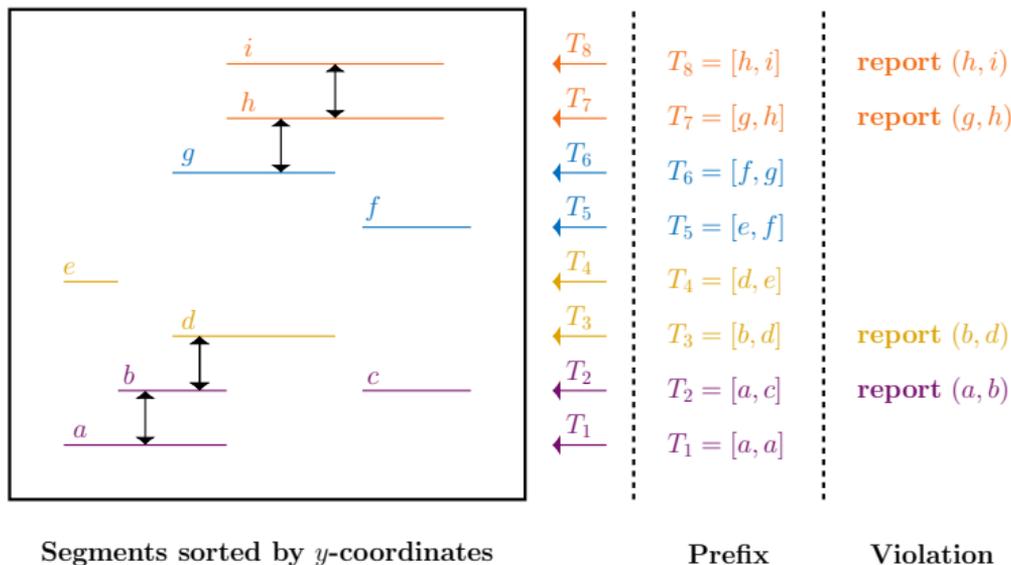
- 2 **Sweeping:** sweep the partial sums

$$s = (?, ?, 12, ?, ?, 25, ?, ?, 27)$$

- 3 **Refining:** compute other prefix sums

$$s = (4, 9, 12, 18, 20, 25, 26, 27, 27)$$

- Key idea: the prefix structure contains a set \mathcal{S} of segments that are below current segment within δ in y -direction
- Remains to check if each pair of segments overlap in the x -direction



Assume we have n elements evenly distributed to b blocks.
Let s_i be the size of the i -th prefix structure.

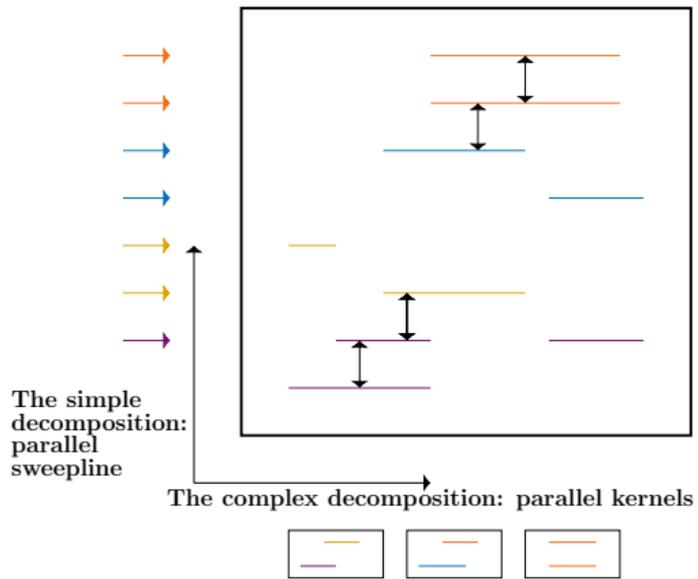
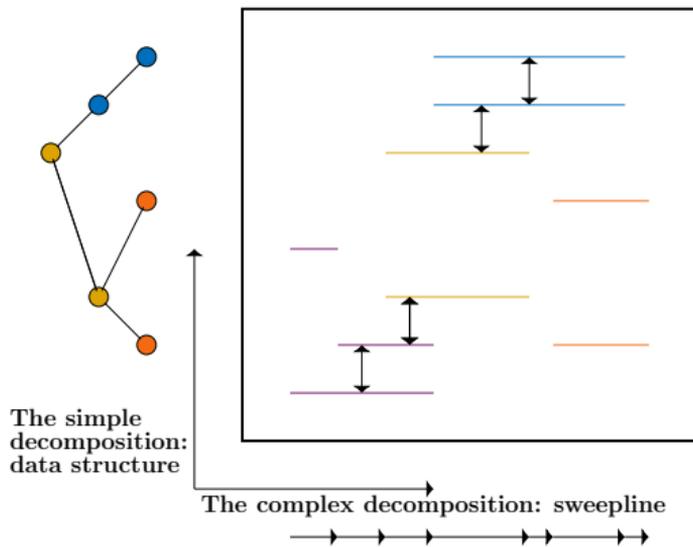
- 1 Batching: b binary search, $O(\log(n/b))$ depth, $O(b \log(n/b))$ work
- 2 Sweeping: $\sum_{k=1}^b O(\log(s_{(k-1)n/b} + n/b))$ work and depth
- 3 Refining: building the i -th prefix structure takes $O(\log s_{i-1})$ time. Total work $\sum_{k=1}^n O(\log(s_{k-1}))$, depth $\max_k \sum_{i=1}^{n/b} O(\log(s_{(k-1)n/b+i-1}))$.

Note that $s_i = O(i)$.

The worse case: $O(n \log n)$ work and $O((b + n/b) \log n)$ depth.

When $b = \Theta(\sqrt{n})$, the depth is $O(\sqrt{n} \log n)$.

- Decompose a problem by the 'simple' direction for parallelism, and leave the 'complex' work to each individual processor.
- The emphasis is different from the sequential version: we use sweepline to deal with the hard direction and maintain the easy direction for efficient query
- In the distance check case: horizontal is the hard direction (2 endpoints per segment, no total order)



GPU Implementation

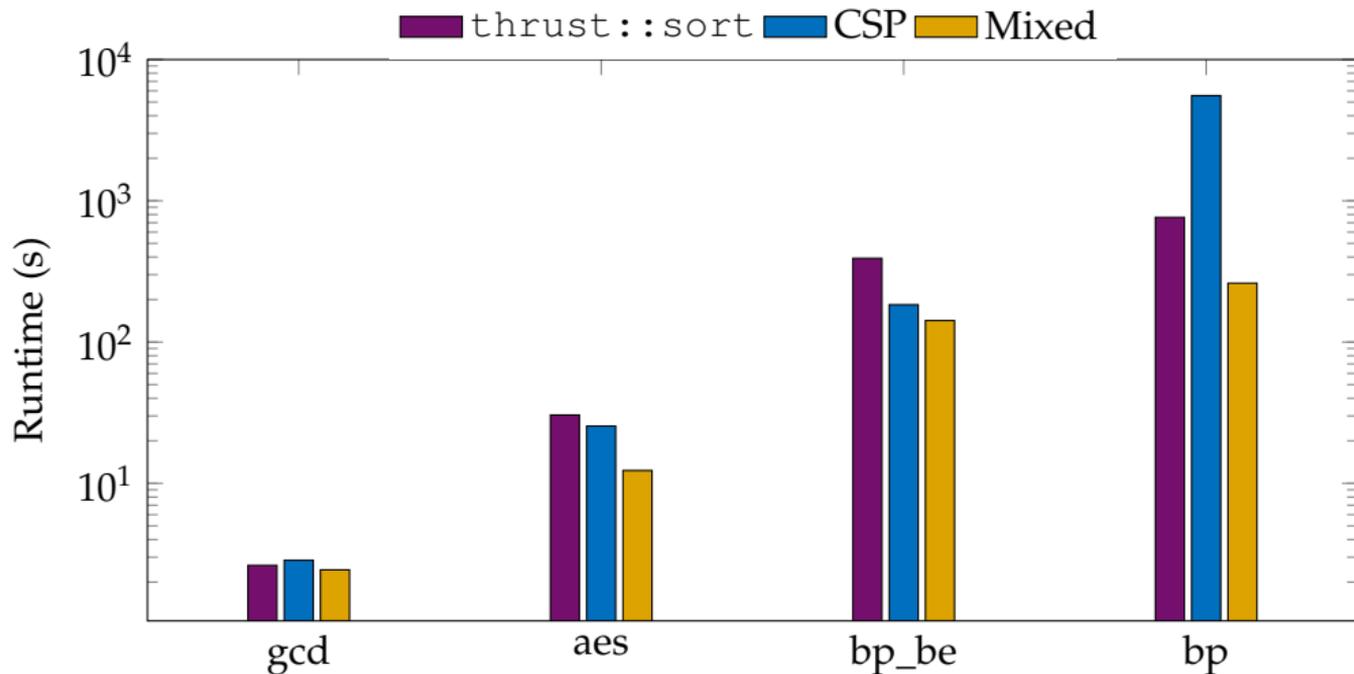
- The sweepline framework is divide-and-conquer (GPU-friendly)
- dynamic algorithm selection: don't invoke GPU if not necessary
- kernel granularity
 - tile-wise
 - polygon-wise
 - per prefix structure
 - per check
- Sorting?

Two commonly used parallelizable sorting algorithms

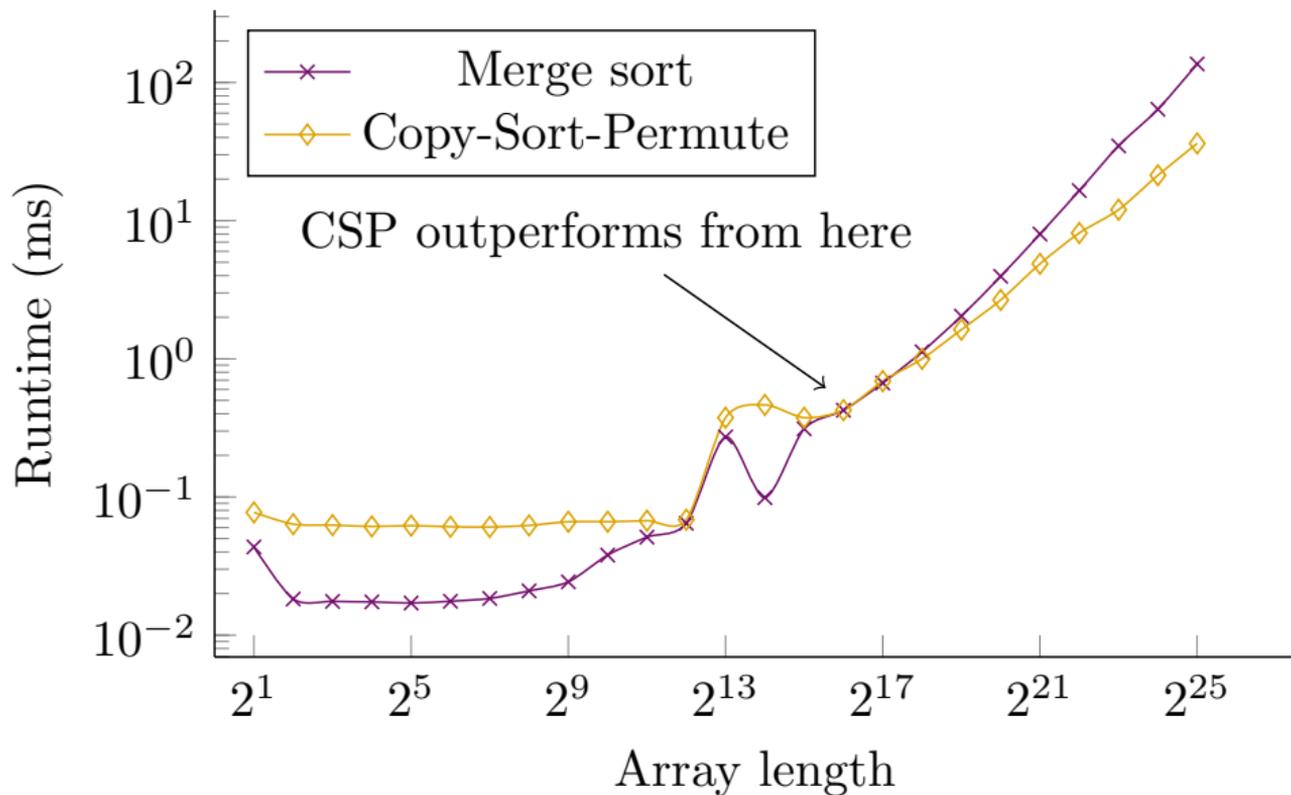
- Merge sort
 - comparison-based
 - e.g., when you pass a *comparison function object* as an argument to `thrust::sort`
- Radix sort
 - non comparison-based
 - works for numeric data types (e.g., `int`) and default comparators

```
1 // Assume we want to sort array by S::key.
2 // n is the length of the array.
3 // effectively equivalent to thrust::sort(array, array+n);
4 template <typename S>
5 void sort_long_arrays(S *array, int n) {
6     int *keys;    // the buffer for keys
7     int *indices; // the buffer for indices
8     S *tmp;      // the buffer for permutation
9
10    // step 0: properly allocate the buffers
11    cudaMallocManaged(...)...
```

```
1  // step 1: Copy
2  for (int i = 0; i < n; ++i) {
3      keys[i] = array[i].key;
4      indices[i] = i;
5  }
6  // step 2: Sort
7  thrust::sort_by_key(keys, keys+n, indices);
8  // step 3: Permute
9  thrust::copy_n(
10     thrust::make_permutation_iterator(
11         array, indices),
12     n, tmp);
13  thrust::copy_n(tmp, n, array);
14 }
```



Runtime of enclosing check on Metal 1 in log scale.



Experimental Results

- Implemented in C++ and CUDA
- Integrated into KLayout¹ (version 0.26.6)
 - Baseline: KLayout DRC Engine (8 threads)
- Test cases synthesized from OpenROAD²
- Environment:
 - Intel Xeon 2.90 GHz Linux machine with 128 GB RAM
 - One NVIDIA GeForce RTX 3090 GPU
 - NVCC 11.4, GNU GCC 10.3

¹<https://klayout.de>

²<https://github.com/The-OpenROAD-Project>

Design	Layer	#Tiles	#Polygons	#Edges	#Edge/Polygon	Width Check Time (s)		
						KLayout	X-Check	Speedup
gcd	Metal1	1	391	24440	62.5	<0.1	0.1	-
	Metal2	1	1229	4916	4.0	<0.1	<0.1	-
aes	Metal1	16	17739	2059906	116.1	2.9	3.0	0.97×
	Metal2	16	76007	304028	4.0	0.2	0.1	-
bp_be	Metal1	56	34747	27245522	784.1	21.9	19.3	1.13×
	Metal2	56	393834	1575336	4.0	0.4	0.4	-
bp	Metal1	144	107706	52595418	488.3	38.9	33.0	1.18×
	Metal2	144	833588	3334352	4.0	0.9	0.9	-
Average								1.09×

Design	Layer	Enclosing Check			Space Check		
		KLayout	X-Check	Speedup	KLayout	X-Check	Speedup
gcd	Metal1	38.4	2.4	16.00×	12.6	2.4	5.25×
	Metal2	2.5	2.5	1.00×	6.4	2.4	2.67×
aes	Metal1	15470.4	12.3	1257.76×	4493.8	67.5	66.57×
	Metal2	2227.0	14.5	153.59×	2778.5	9.9	280.66×
bp_be	Metal1	66194.6	128.6	514.73×	6718.7	123.7	54.31×
	Metal2	3089.2	147.4	20.96×	4171.5	16.6	251.30×
bp	Metal1	98370.4	235.3	418.06×	14019.7	233.4	60.07×
	Metal2	3958.7	276.6	14.41×	5164.4	65.9	78.37×
Average				61.36×			45.00×

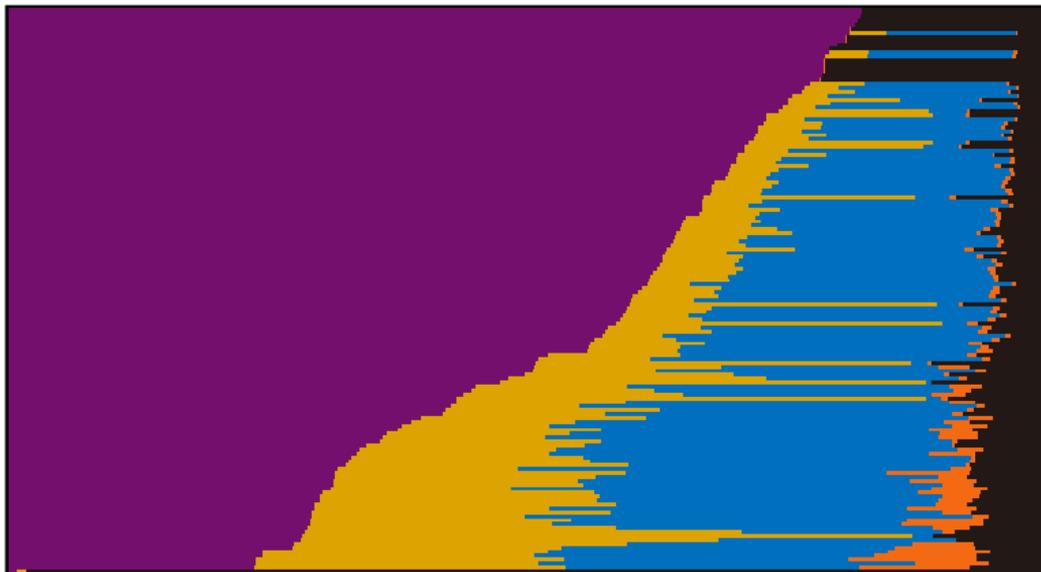


(a) KLayout



(b) X-Check

Each horizontal bar is for one thread. The purple and gold portions are for the *merge* and the *check* stages, respectively.



Each horizontal bar is for one thread. The **purple** portion is for *merge*, **gold** for *sort*, **blue** for *prefix build*, **orange** for *violation report*, and **black** for the rest, respectively.

- Parallel sweepline algorithm for DRC
- GPU implementation considerations
- Integration into an end-to-end flow
- Future work
 - Parallelize/Accelerate the merge stage
 - GPU infrastructure: associative data structures and thread-safe solution

THANK YOU!