September 2022 | Aki Fujimura, CEO D2S, Inc.

O(p): GPUs, Pixels, DL, Curvy Masks & Designs

D2S Does GPU Acceleration

Curvy Shapes

Deep Learning

Pixel Domain

GPU Acceleration

GPU Accelerates Edge Manipulation, Too

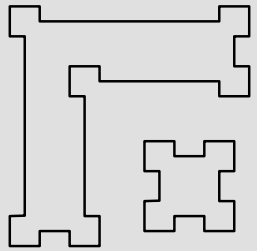

Manhattan OPC

Curvy ILT

MWCO*

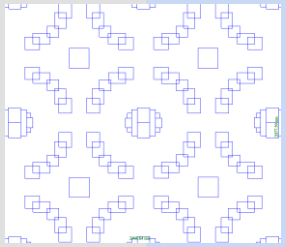

Fracturing

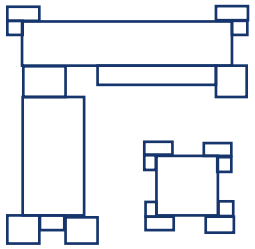

MB-MDP (overlapping)

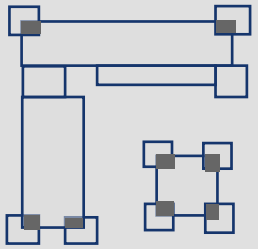

Piecewise Linear Polygons

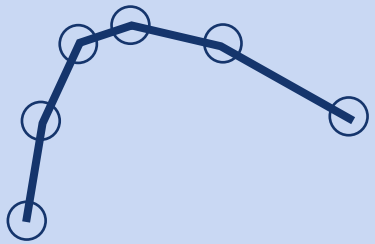

Curves (Bezier/Spline)

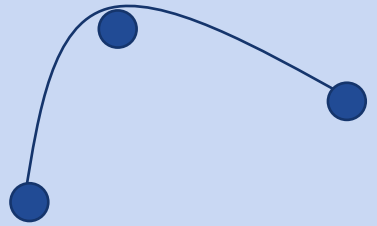

Pixels

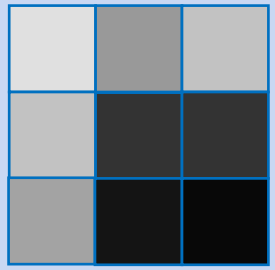

VSB

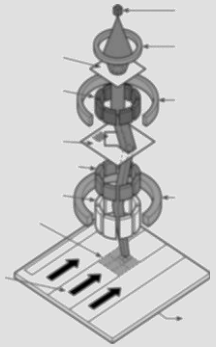

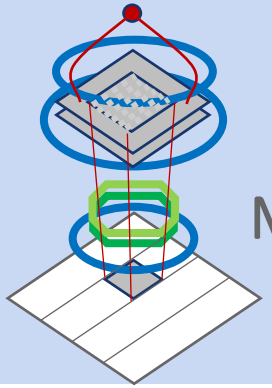

Multi-beam

*MWCO : Mask-Wafer Co-Optimization

GPU is Great for Curvy

$$O(p)$$

Because GPU is even better at pixels

GPU is SIMD and SIMD is Great at Pixels

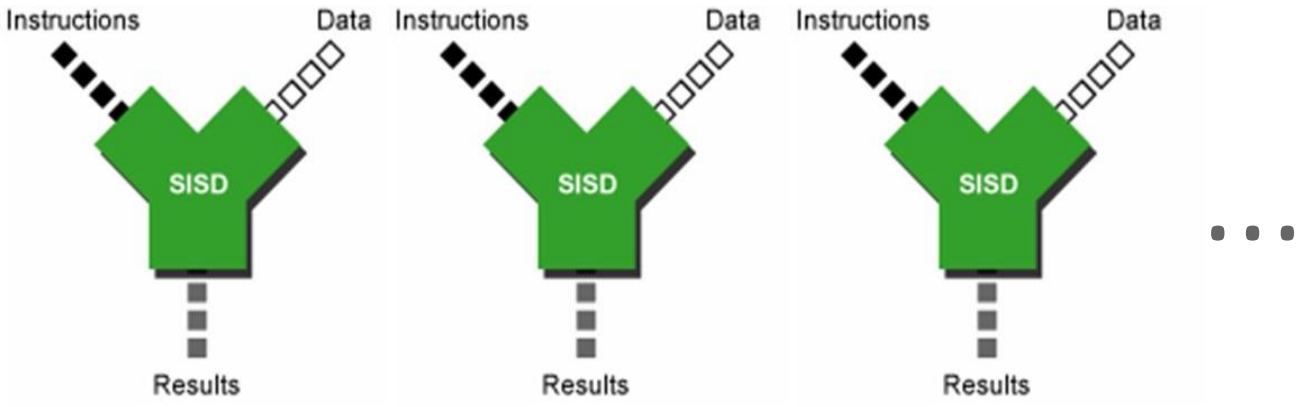

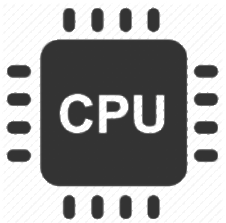

Multiple Single Instruction Single Data (SISD)

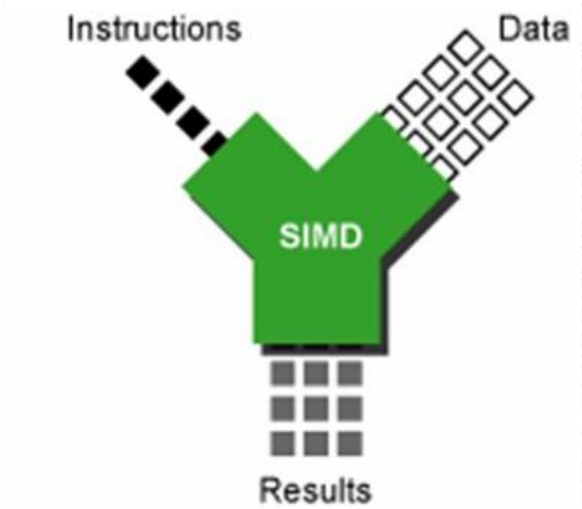

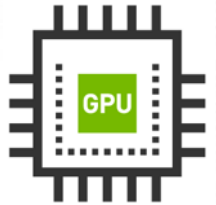

Single Instruction Multiple Data (SIMD)

Mask and Wafer Effects are SIMD

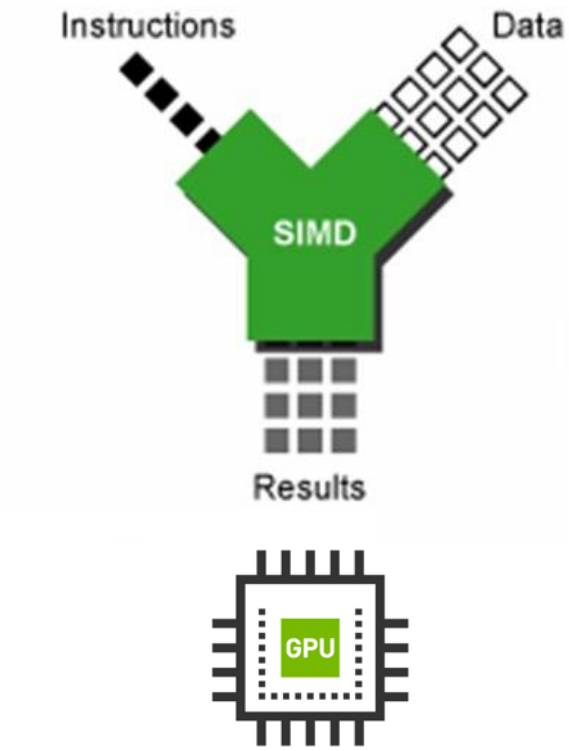

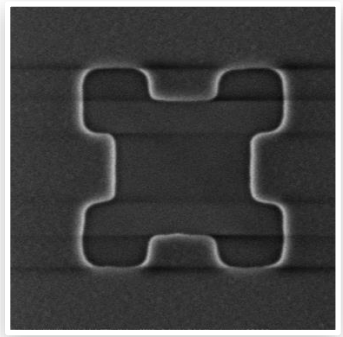

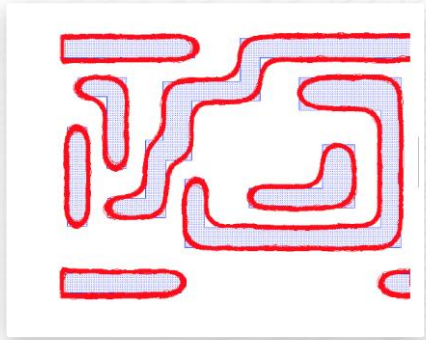

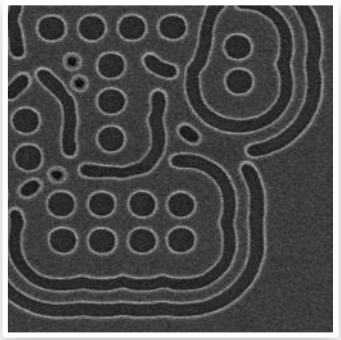

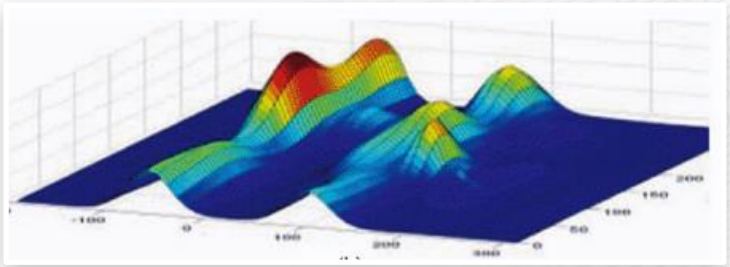

Because nature is SIMD

Image Processing is SIMD, Too

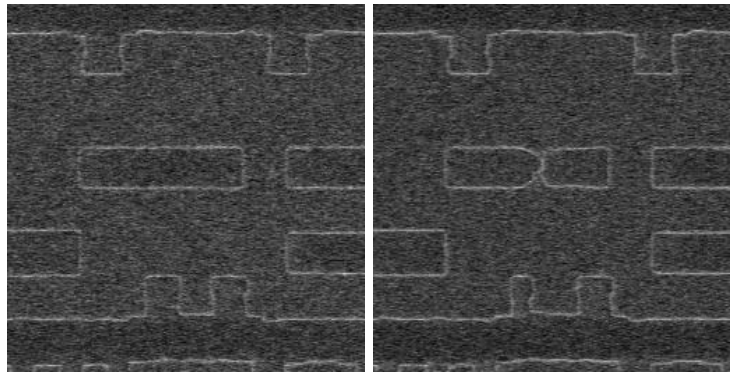

Source: NuFlare/CDLe

SEM

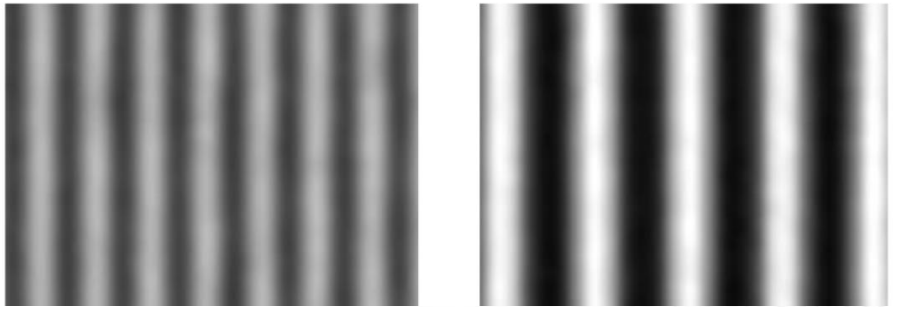

Source: Lasertec/BACUS Newsletter

Inspection

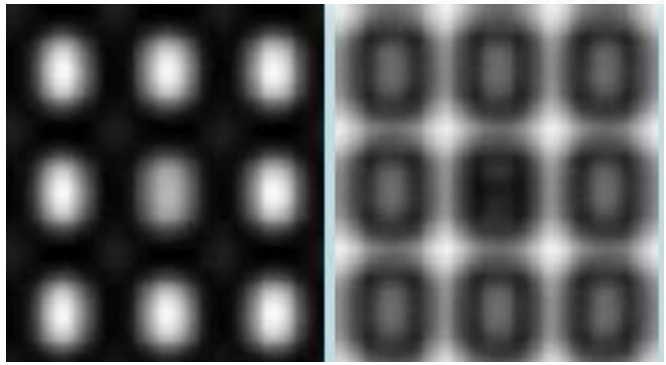

Source: L Pang, et al., "Expanding the applications of computational lithography and inspection (CLI) in mask inspection, metrology, review, and repair"

Because Images are Arrays of Pixels

Pixel-Based and Edge-Based are Duals

- Whatever you can do in one can be done in the other…given a resolution limit determined by “Nyquist”
- The only question is performance….

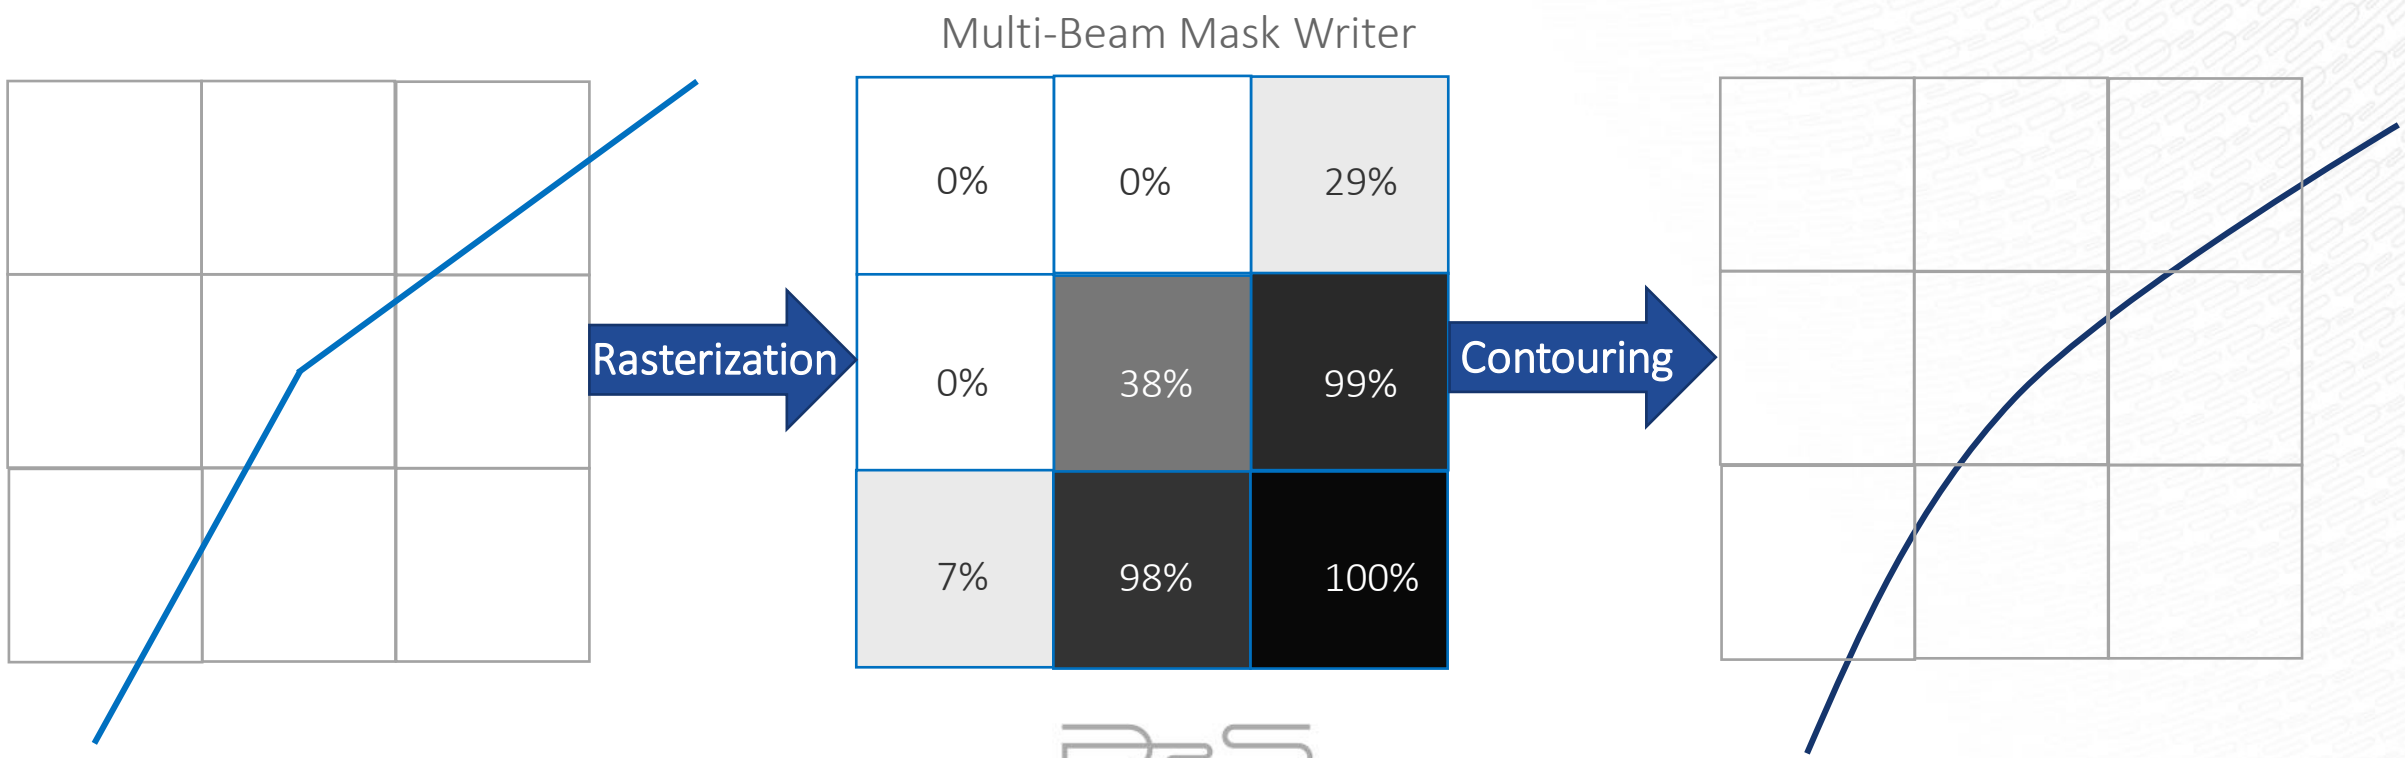

This Also Rasterizes Exactly the Same

- Rasterization is inherently a low-pass filter
- Red and blue become the same in pixels, but red uses much more data
- Curvilinear format would also become the same in pixels and thereafter
  - And represents the actual contoured shape that would be on the physical mask

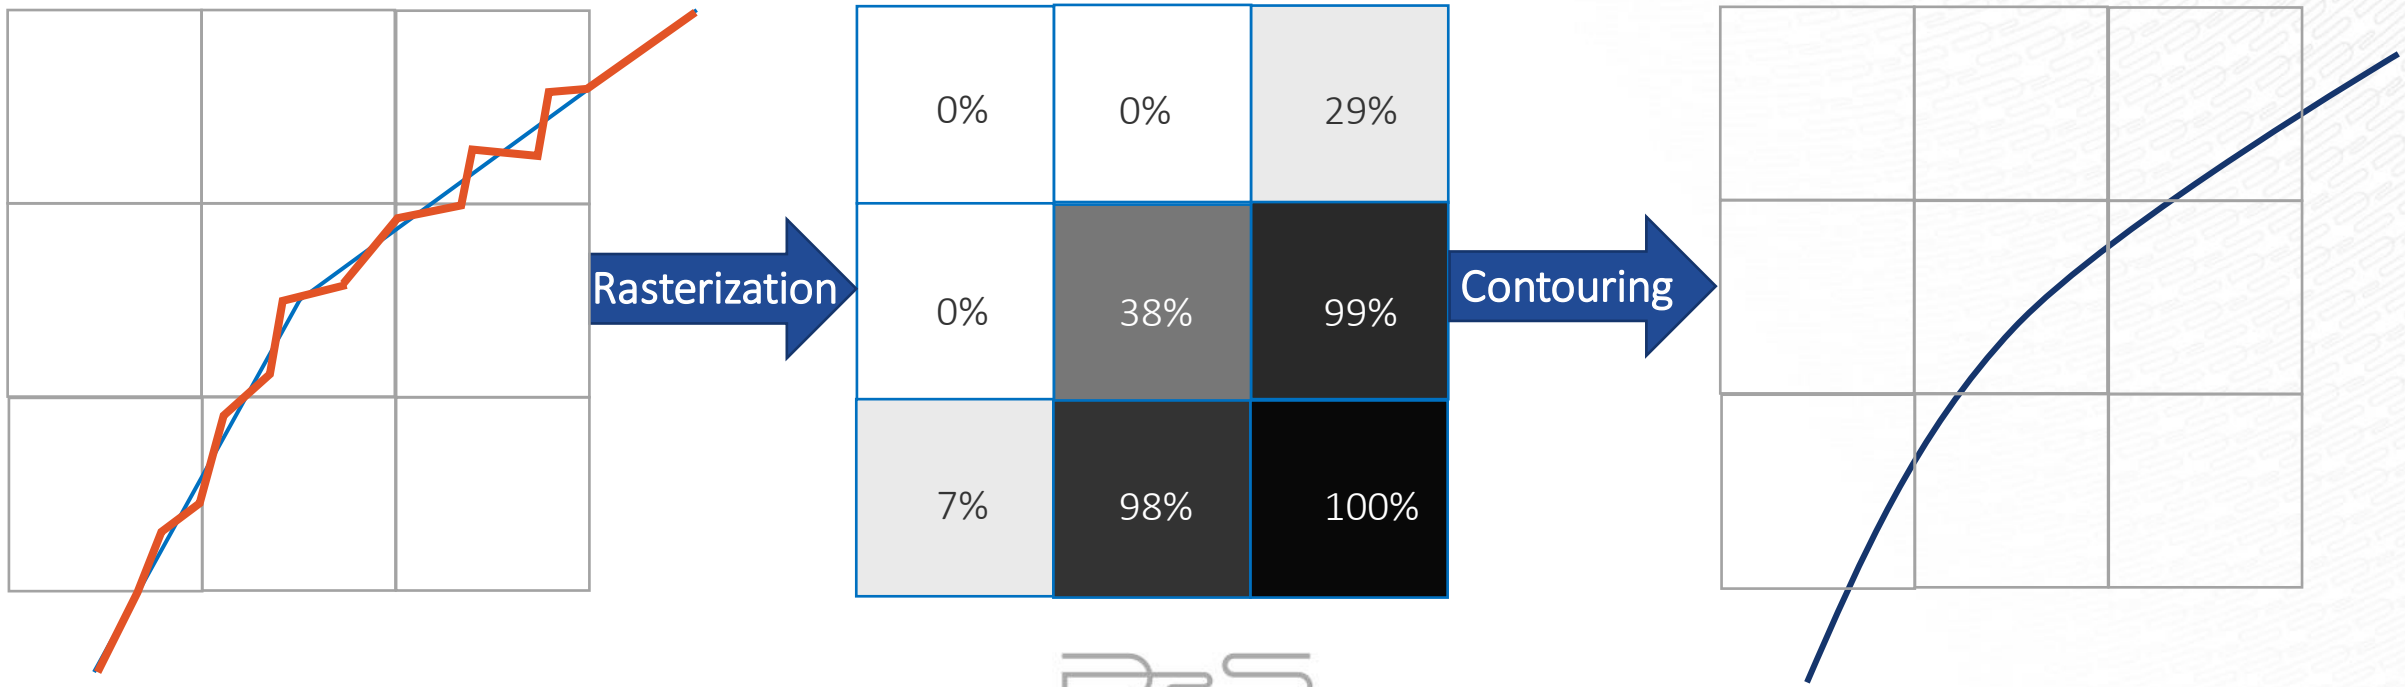

Mask-Making is Already Pixel-Based

Manhattan
OPC

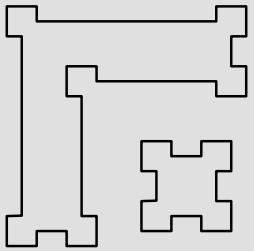

Curvy ILT
$O(p)$

MWCO*

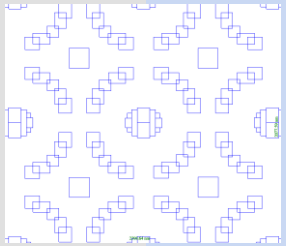

Fracturing

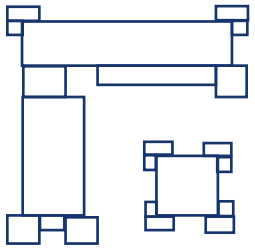

MB-MDP
(overlapping)

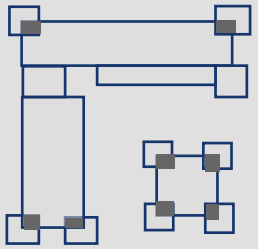

Piecewise Linear
Polygons

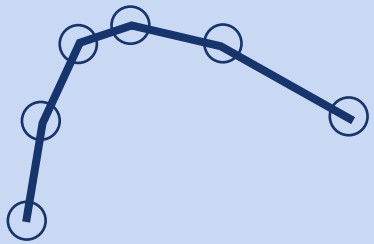

Curves
(Bezier/Spline)

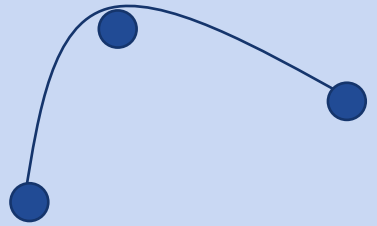

Pixels

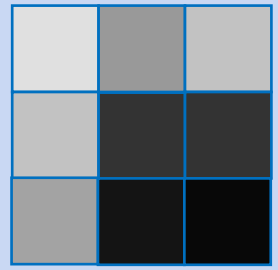

VSB

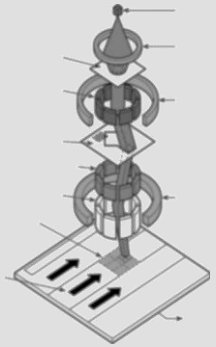

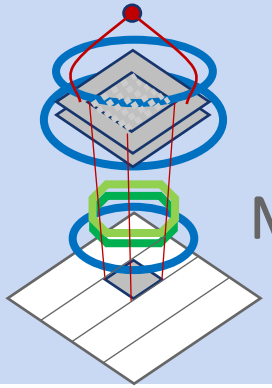

Multi-beam
$O(p)$

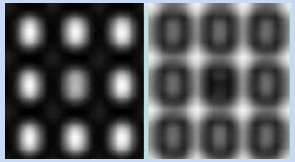

Inspection
$O(p)$

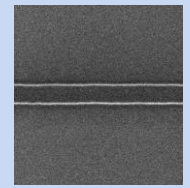

Metrology
$O(p)$

*MWCO : Mask-Wafer Co-Optimization

Datapath is Not

Manhattan
OPC

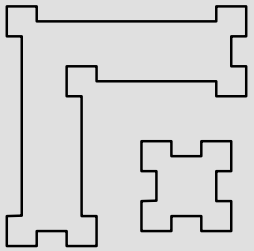

Curvy ILT
$O(p)$

MWCO*

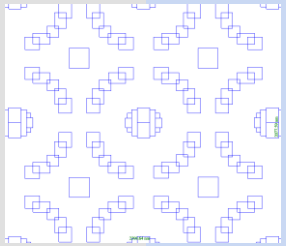

Fracturing

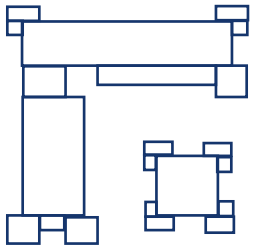

MB-MDP
(overlapping)

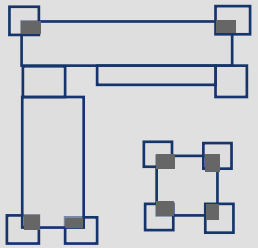

Piecewise Linear
Polygons

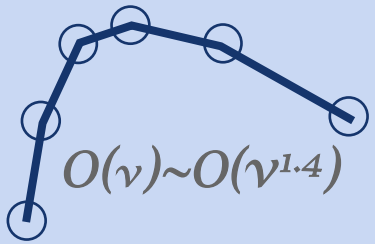

Curves
(Bezier/Spline)

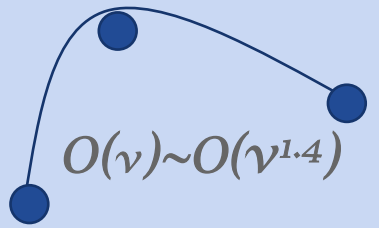

Pixels

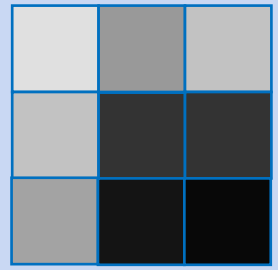

VSB

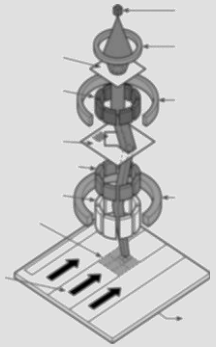

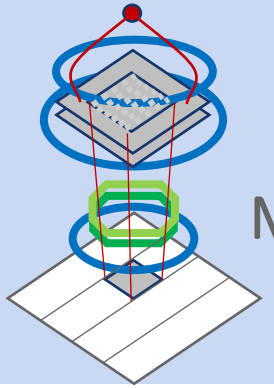

Multi-beam
$O(p)$

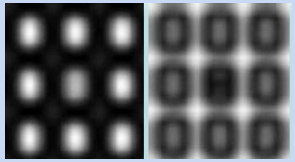

Inspection
$O(p)$

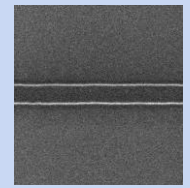

Metrology
$O(p)$

*MWCO : Mask-Wafer Co-Optimization

D2S is Adding a Pixel-Based Datapath

Manhattan OPC

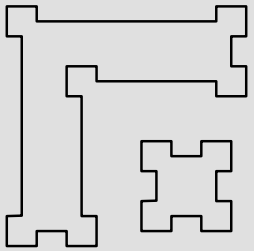

Curvy ILT

$O(p)$

MWCO*

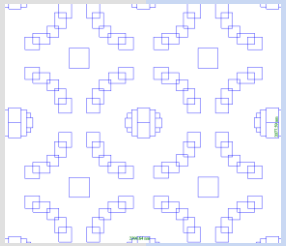

Fracturing

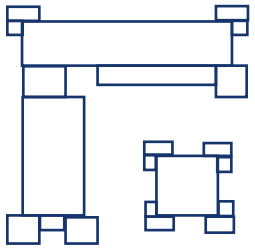

MB-MDP
(overlapping)

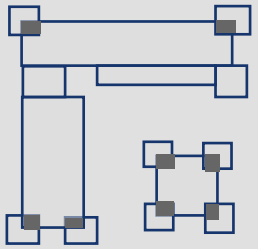

Piecewise Linear
Polygons

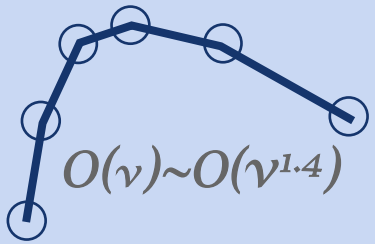

Curves
(Bezier/Spline)

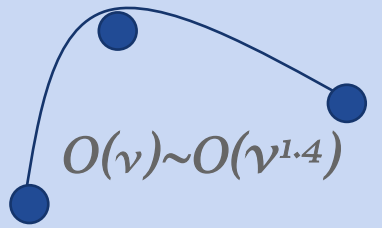

$O(p)$
Pixels

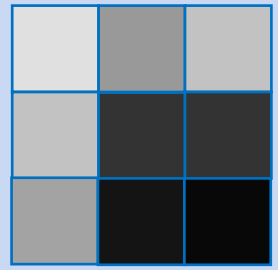

VSB

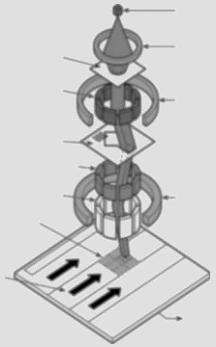

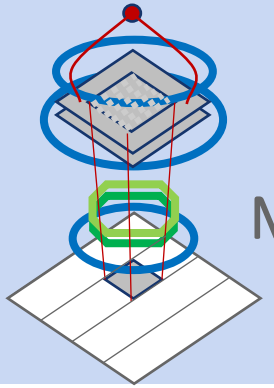

Multi-beam

$O(p)$

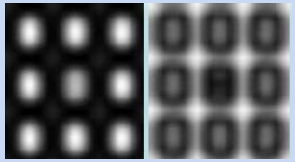

Inspection

$O(p)$

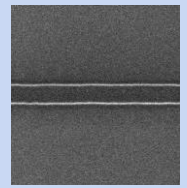

Metrology

$O(p)$

*MWCO : Mask-Wafer Co-Optimization

Edge or Pixel, Curvy Improves Manufacturing

Simple example : Biasing

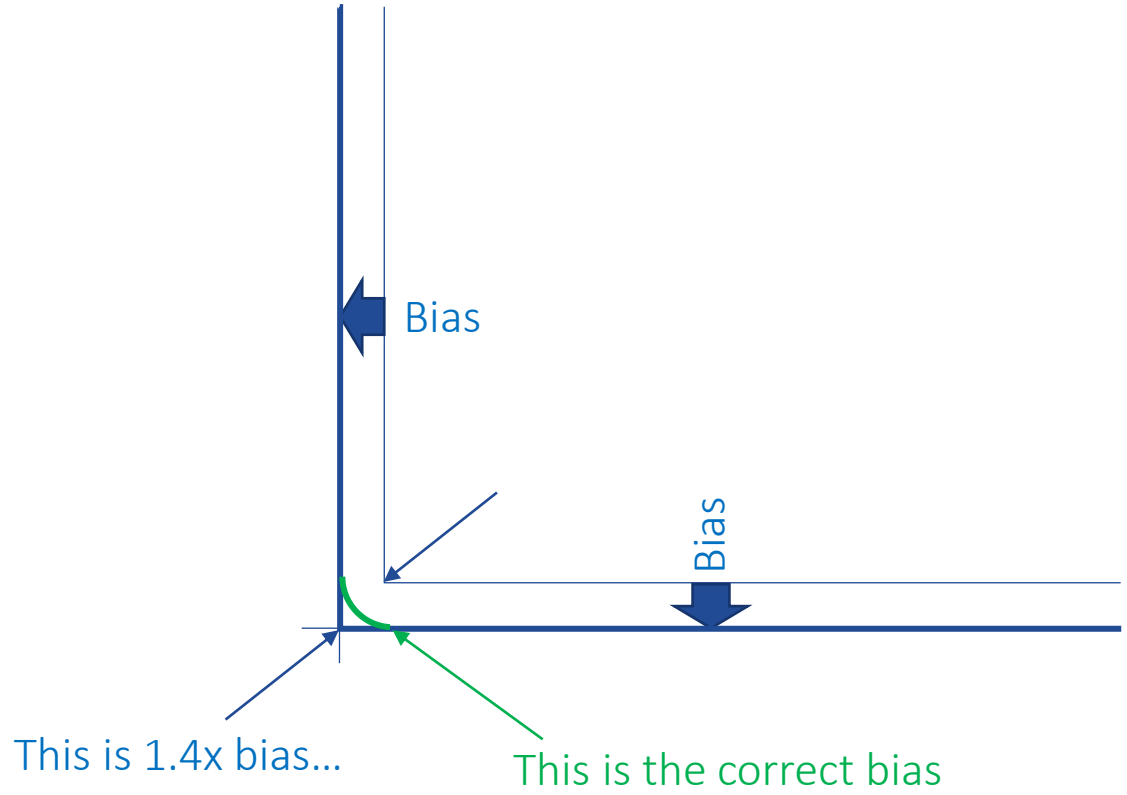

For traditional Manhattan edge-based CAD, it is bad to have bias create curvy shapes. So this is understandable but incorrect. Etching, for example, doesn't work that way.

Even a Simple Bias Operation is Better with Curvy

Rectilinear bias is off by 40% on corners

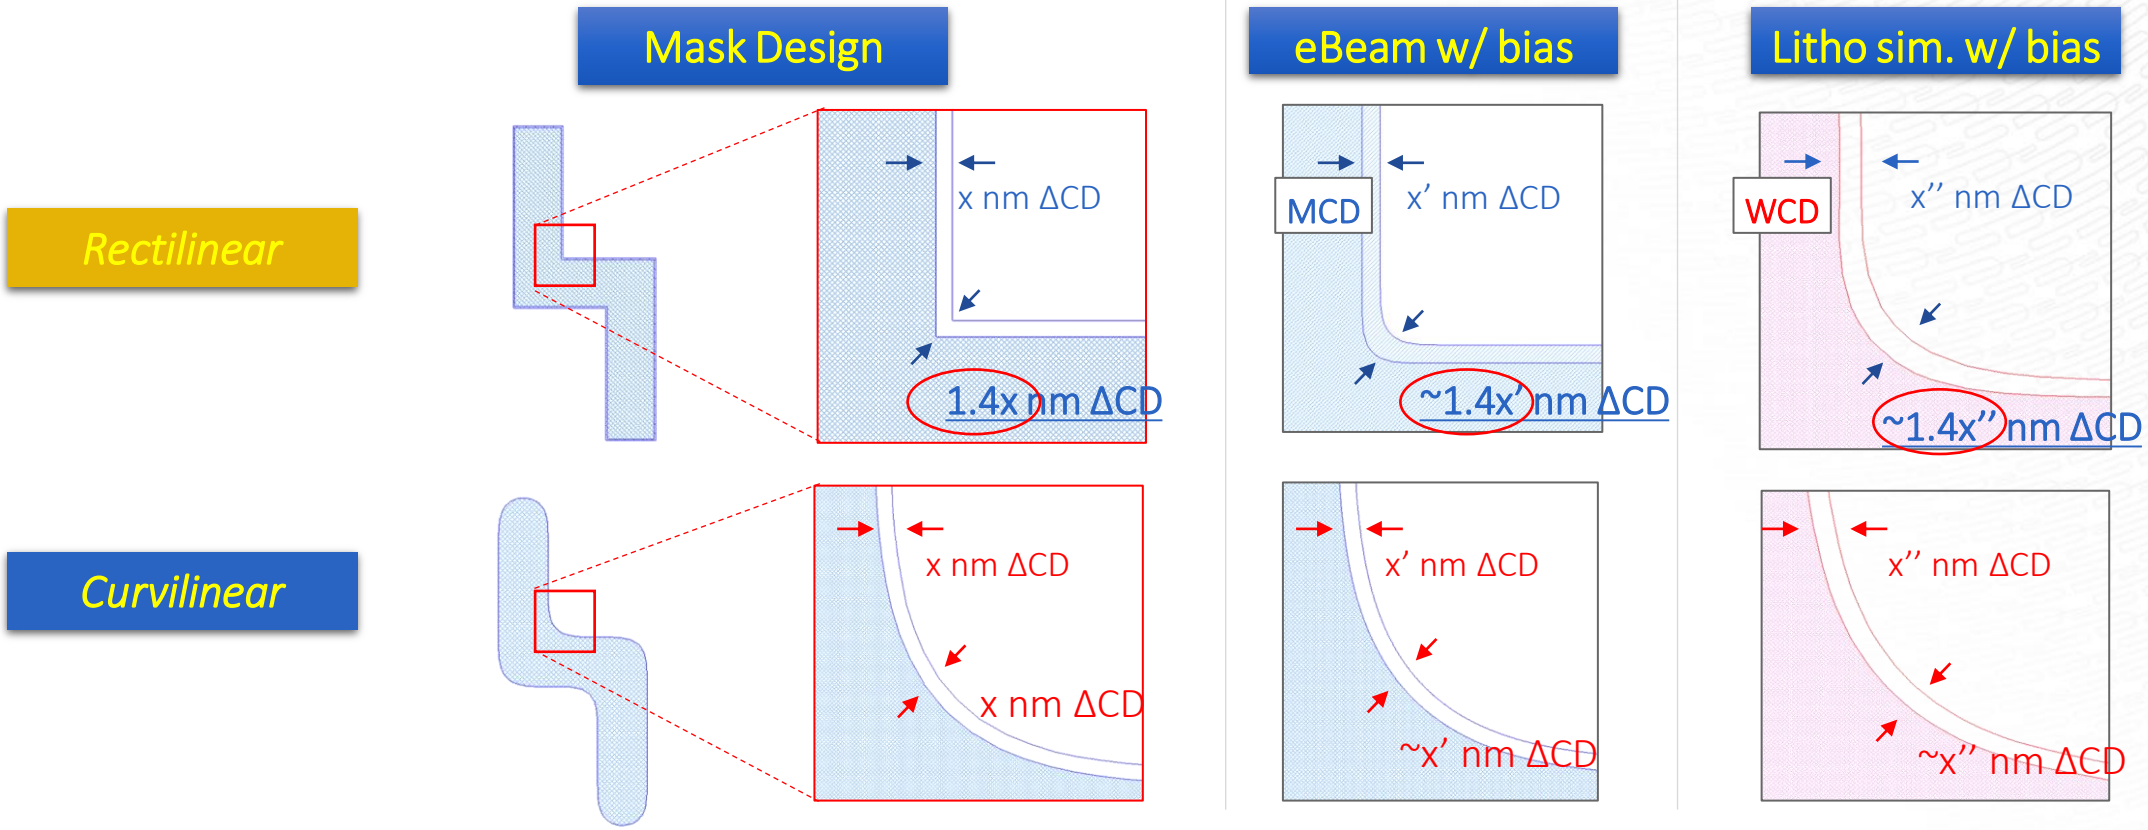

Even Worse is Manufacturing Variation

Dose margin is bad on 90º corners

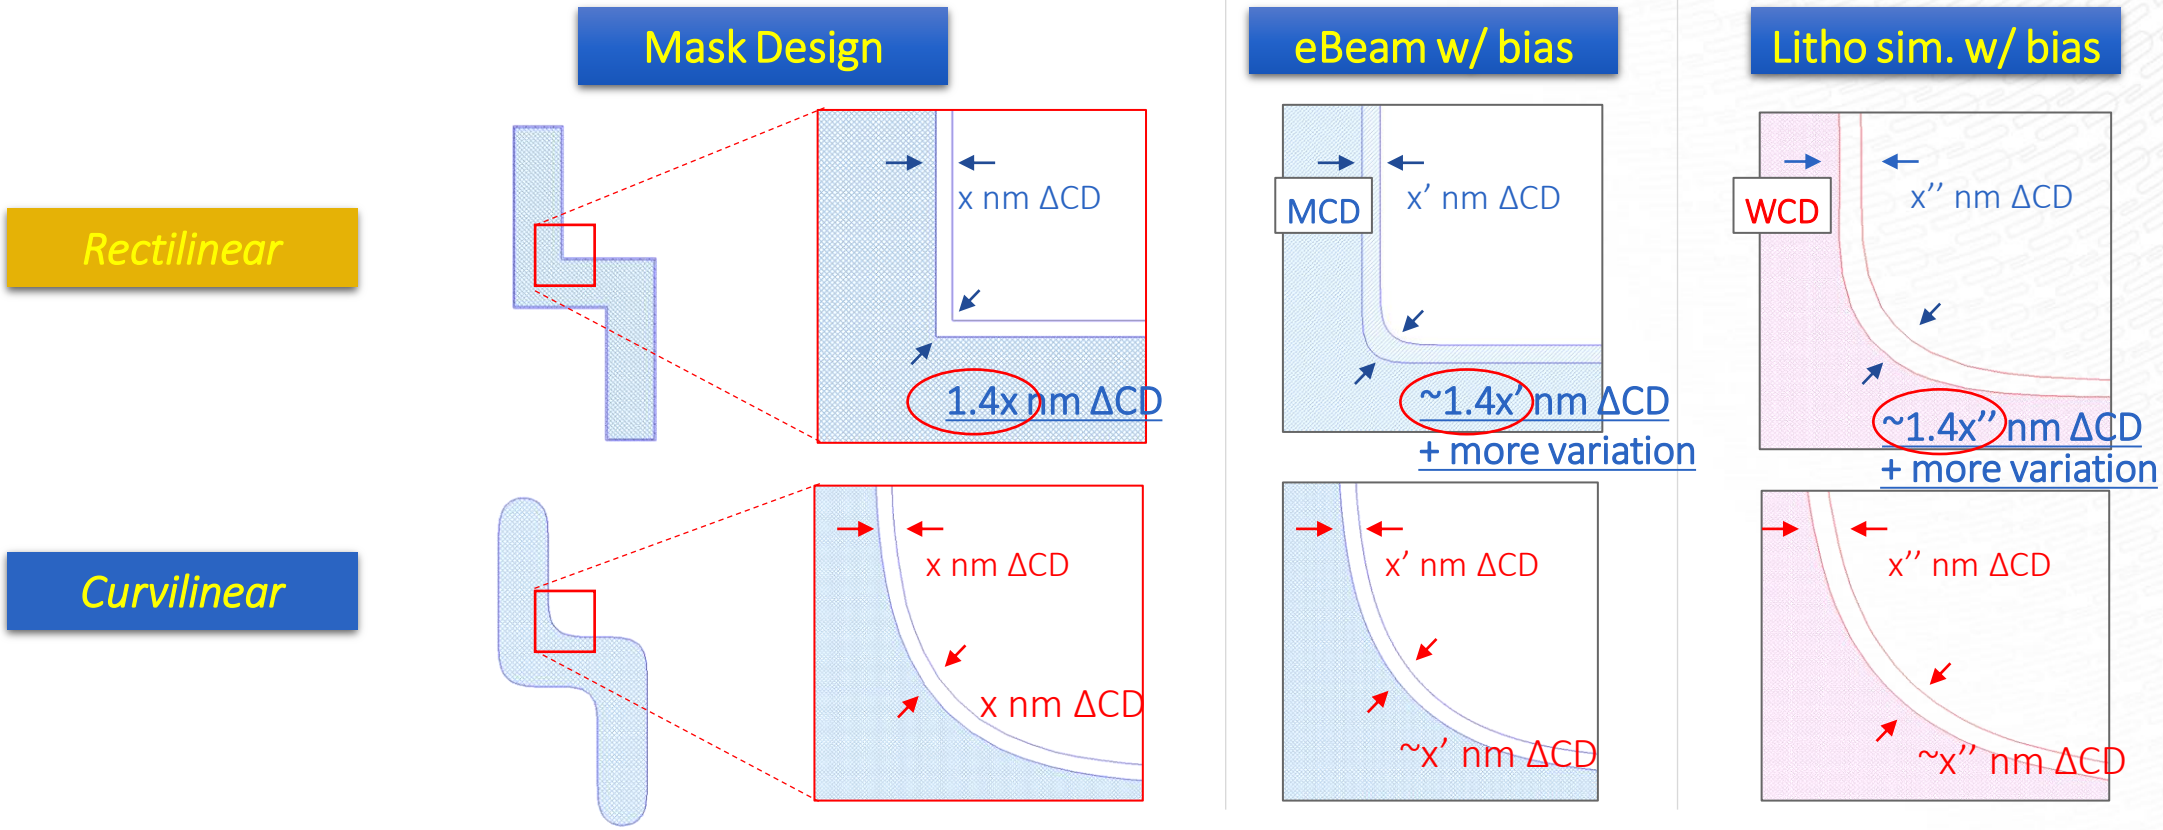

Manufacturable Shapes are More Reliably Manufacturable

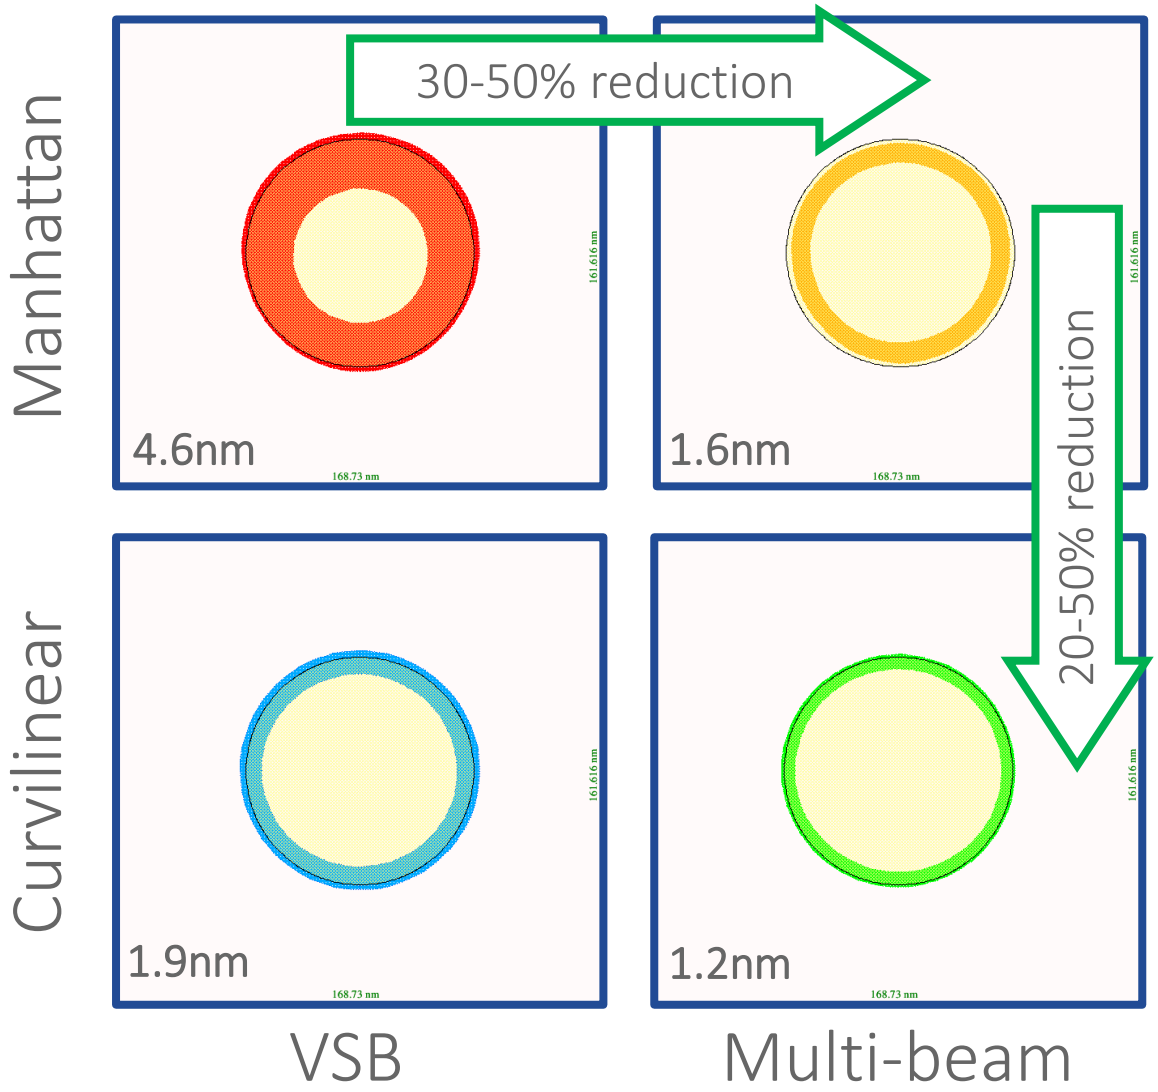

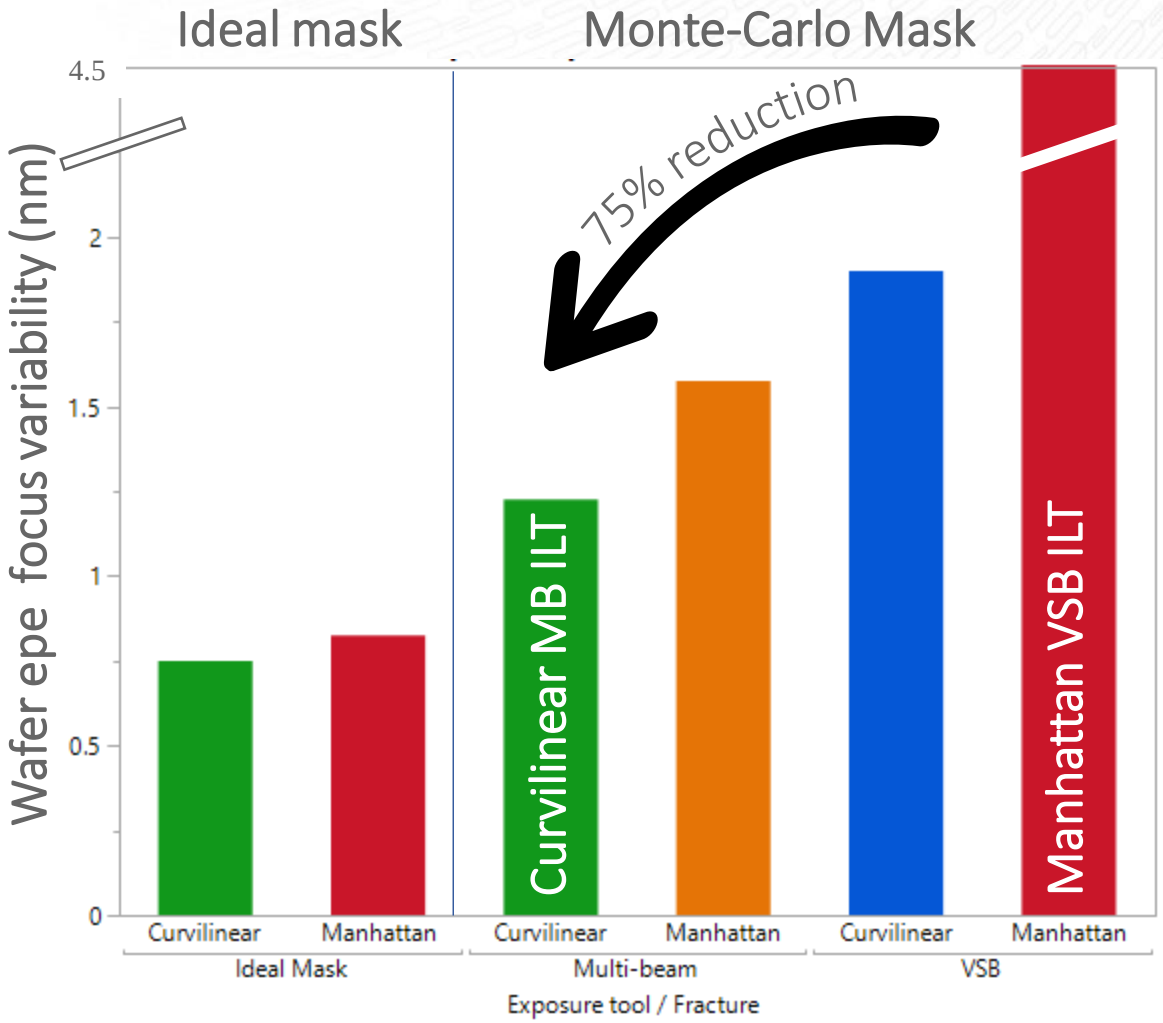

Ryan Pearman, Jeff Ungar, Nagesh Shirali, Abhishek Shendre, Mariusz Niewczas, Leo Pang, and Aki Fujimura "How curvilinear mask patterning will enhance the EUV process window: a study using rigorous wafer+mask dual simulation", Proc. SPIE 11178, (2019)

Manufacturable Shapes are More Reliably Manufacturable

On Wafer, too

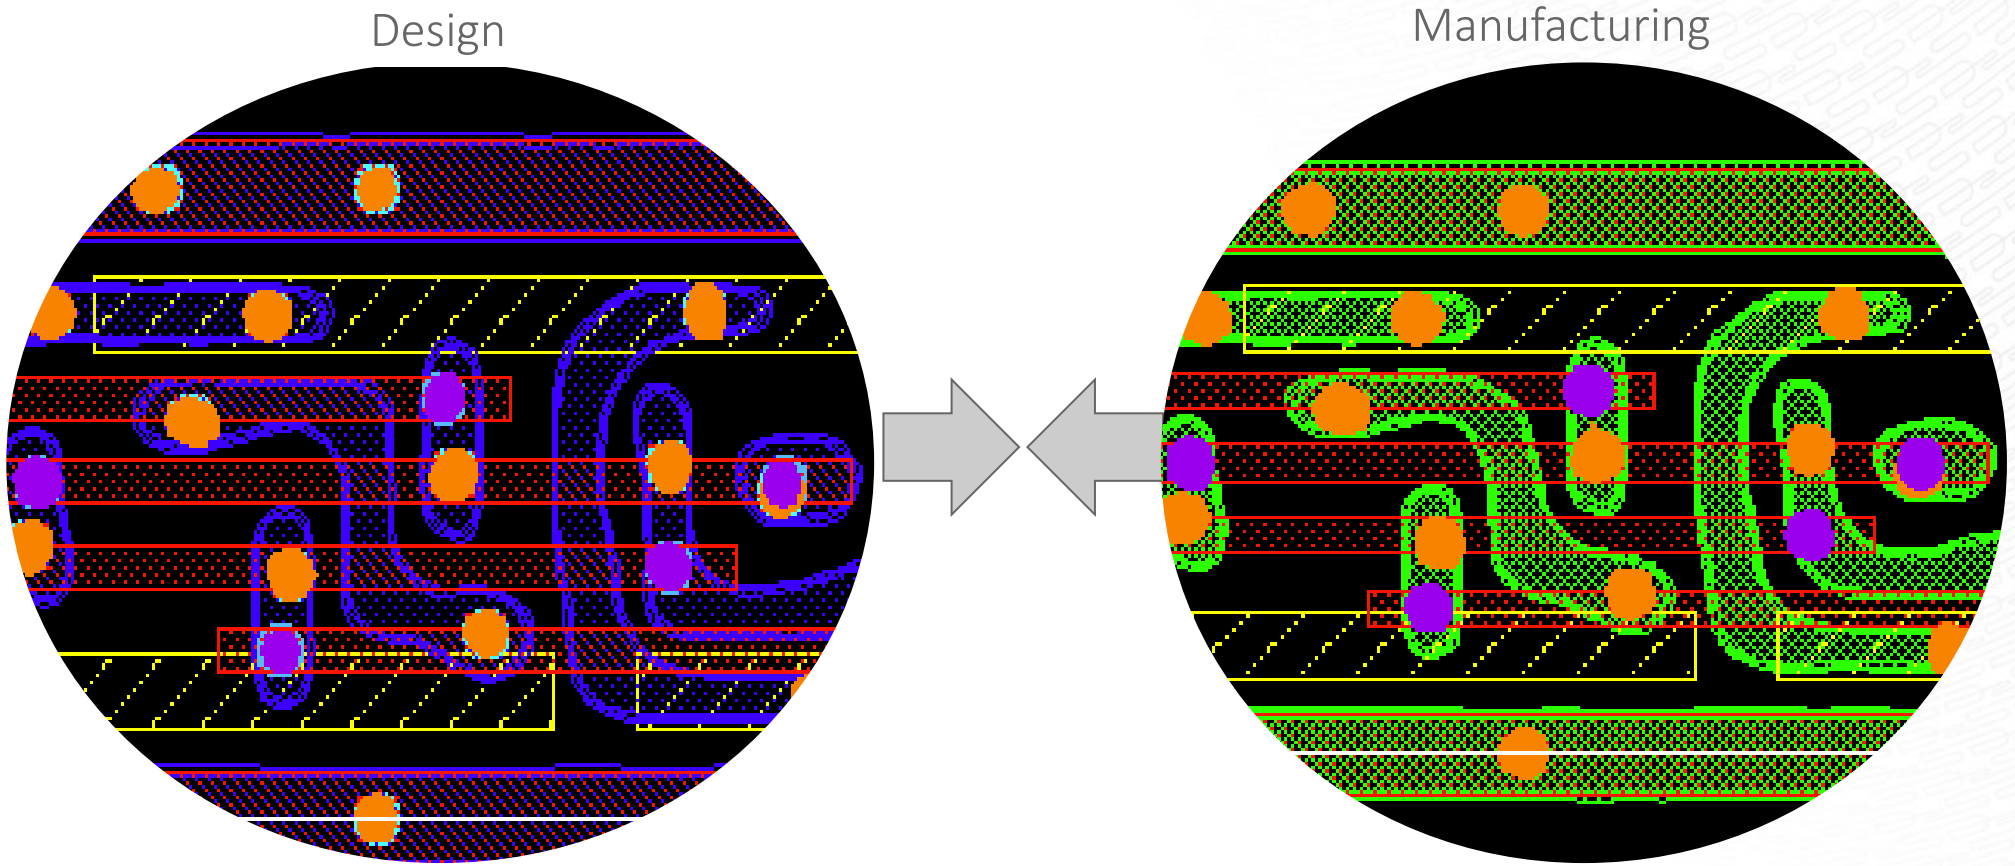

First thing ILT does is to compute manufacturable curvy targets anyway

I Presented This at Design Automation Conference

"I was in EDA physical design since 1979 so I know it's important for you to know:
Curvilinear Designs are Now Manufacturable.
In fact, More Reliably Manufacturable."

Is it Time to Break the Manhattan Assumption?

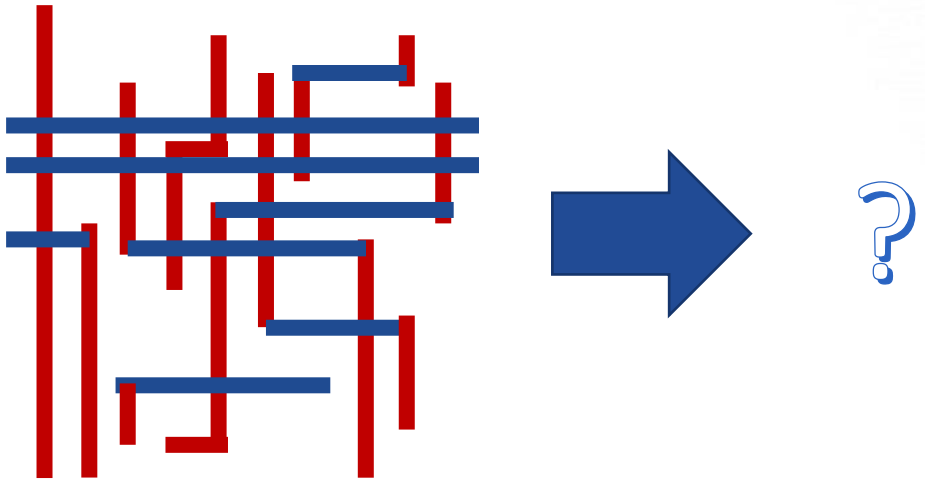

Most chips are interconnect-limited; Reducing vias will reduce routing congestion

Curvy Designs are Better for Designers, Too

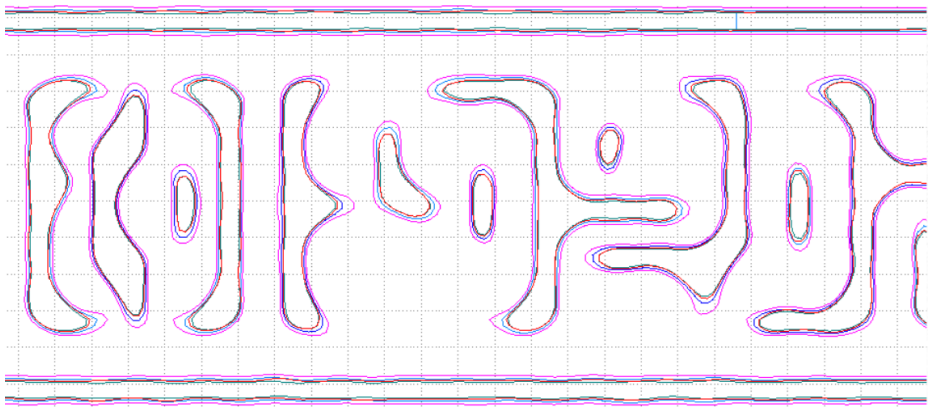

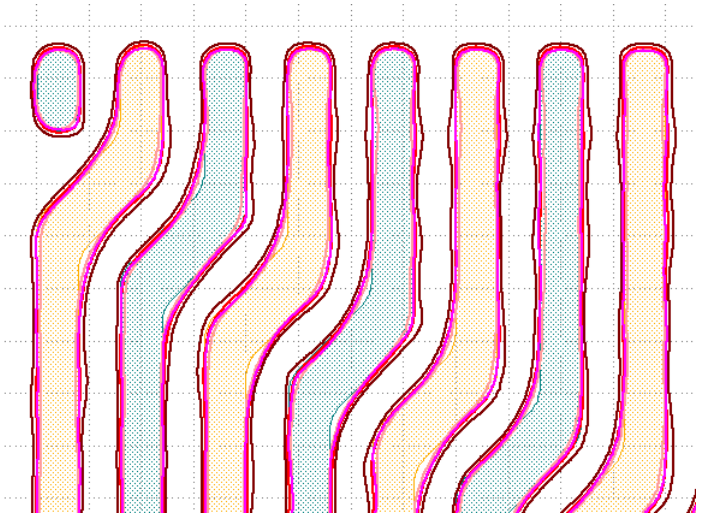

- Manufacturable Curvy Designs
- Improves, all at the same time
  - Yield
  - Power
  - Performance
  - Area
- The barriers are:
  - ~~VSB mask writing~~
  - EDA infrastructure

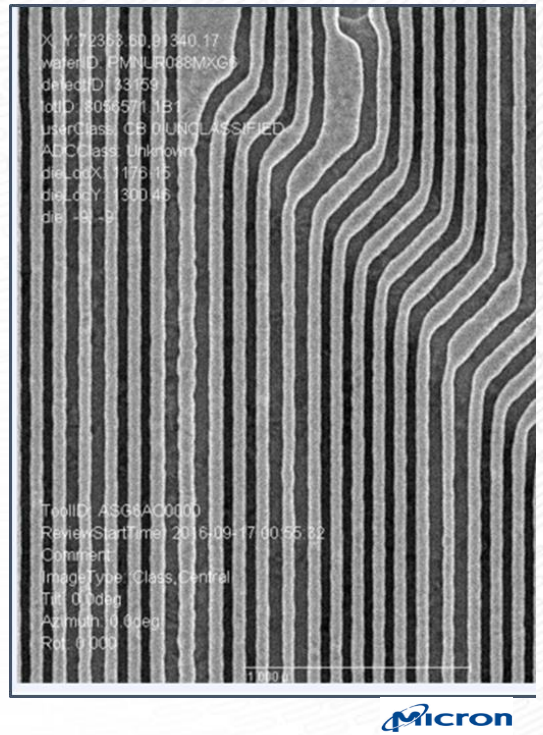

General Perception: “Everything has to Change”

Actually: Only Routing plus Performance Improvements

Custom Design

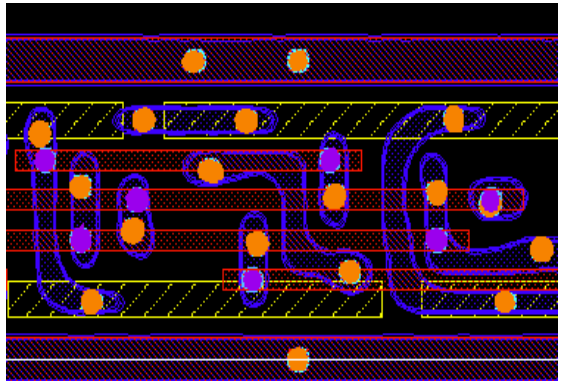

DRC/LVS

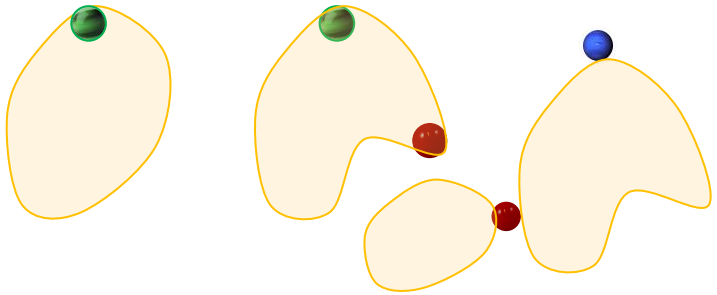

Parasitic Extraction

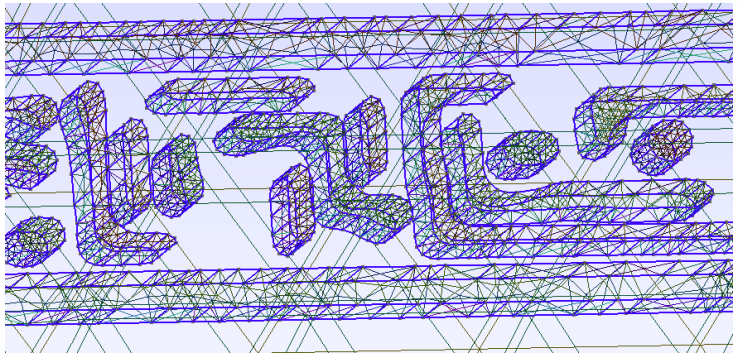

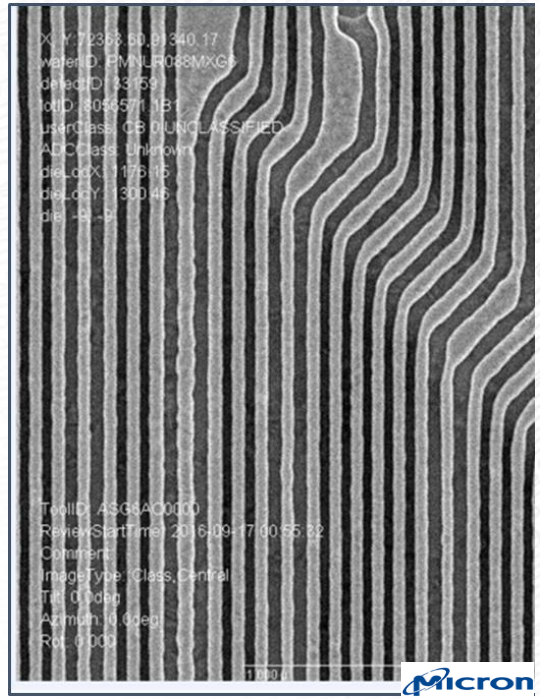

Routing

Curvy-Pixel-GPU Can Work for Them, Too

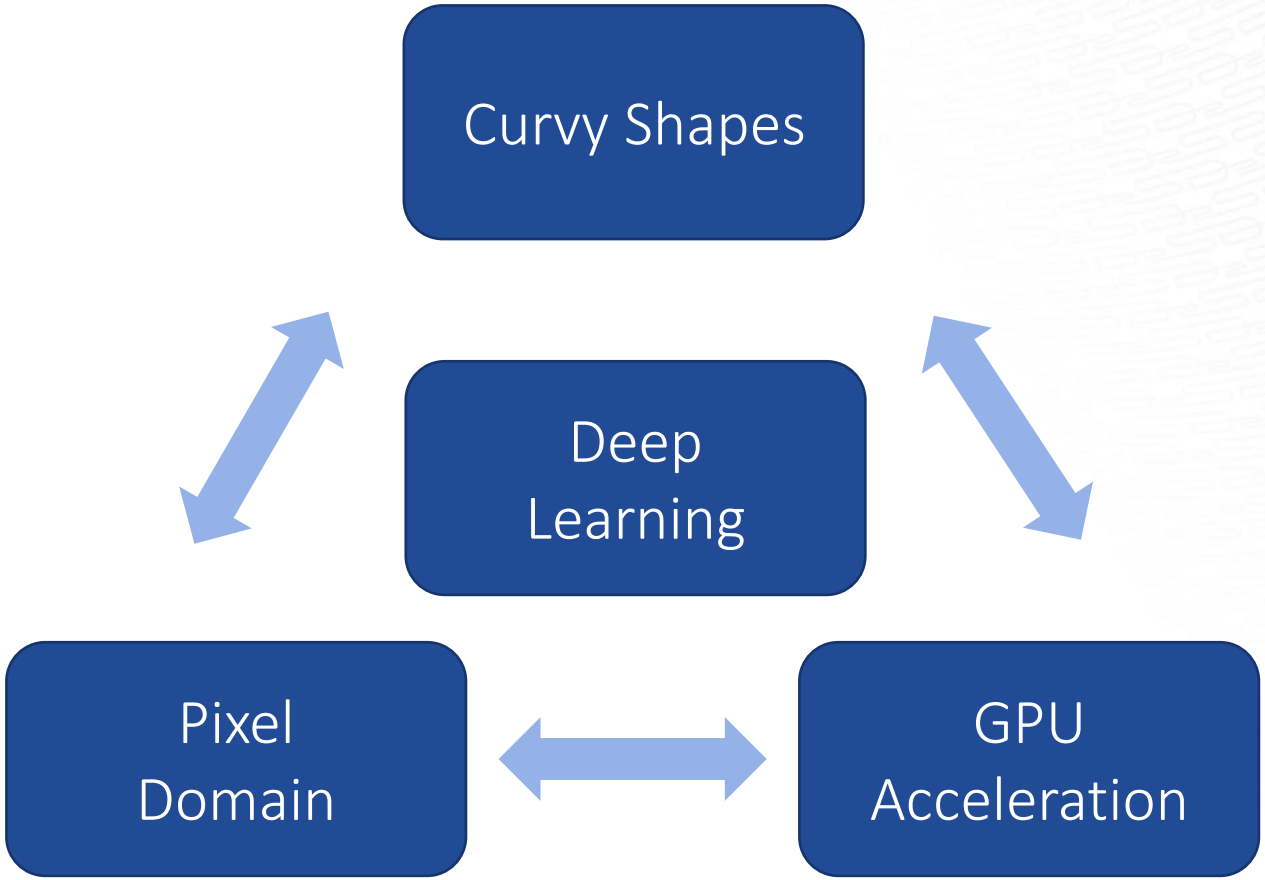

Today's GPU Workstation = 8,000 Cray-2s

60,000,000x Price Performance

It's time to rethink EDA

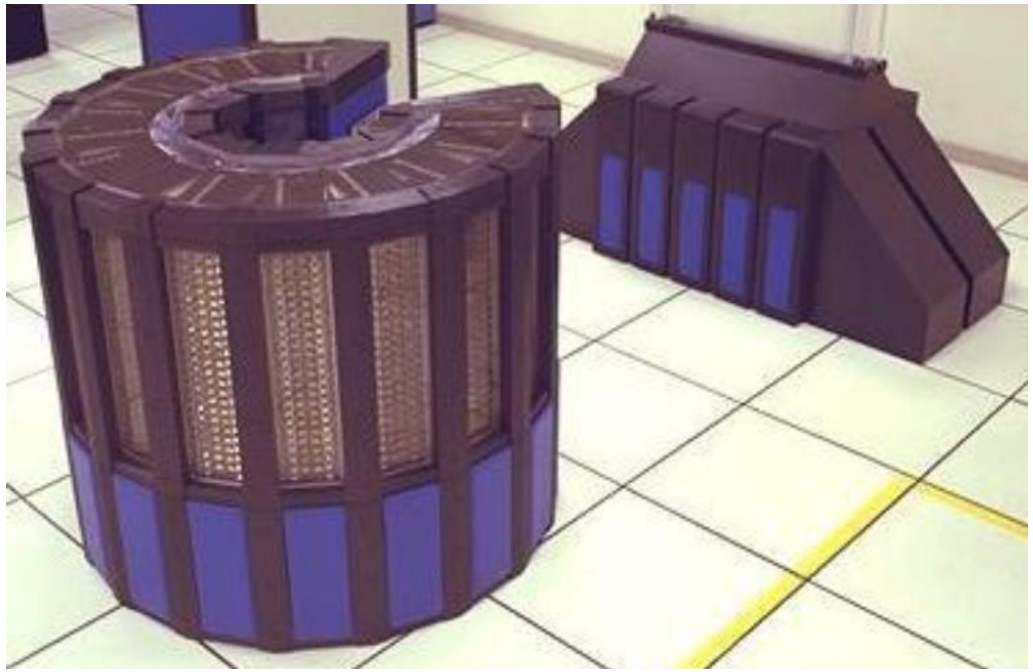

Cray-2 (1985)
1.9 GFLOPS w/500MB @ $15M

nVIDIA RTX 3090 Ti (2021)
15,300 GFLOPS w/24GB @ $2,000

Today's GPU Workstation = ~~8,000~~ 16,000? Cray-2s

150,000,000x? ~~60,000,000~~x Price Performance

It's definitely time to rethink EDA

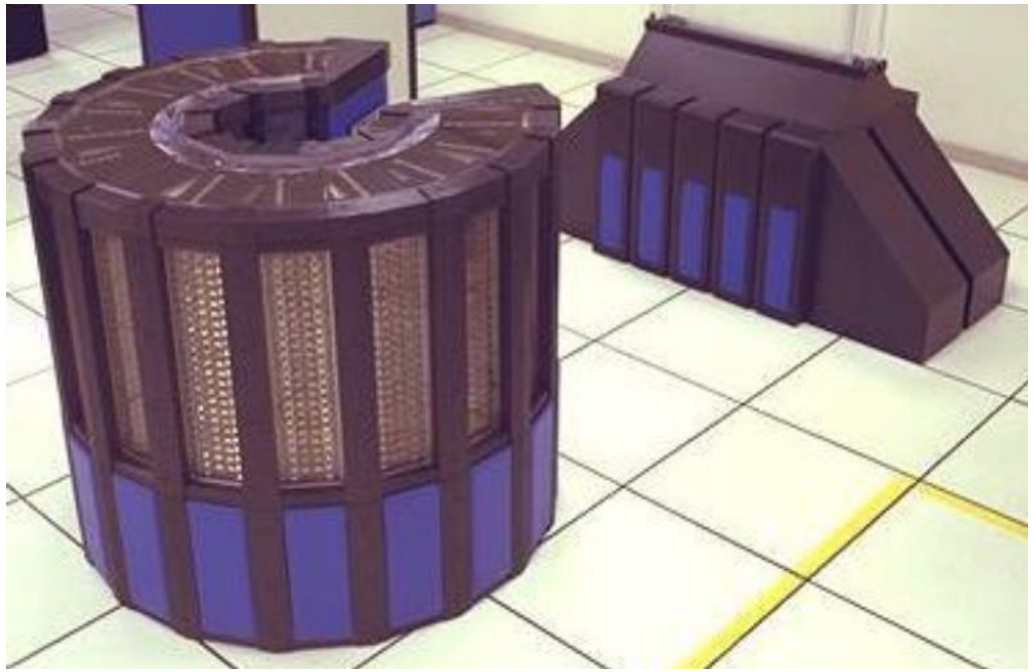

Cray-2 (1985)
1.9 GFLOPS w/500MB @ $15M

Announced Last Tuesday