Aki Fujimura, CEO D2S, Inc.| March 22, 2026 | CSTIC Plenary Session

Changing the Face of Silicon: Curvy Manufacturing Will Lead to Curvy Design

Looking Forward to the Panel This Afternoon

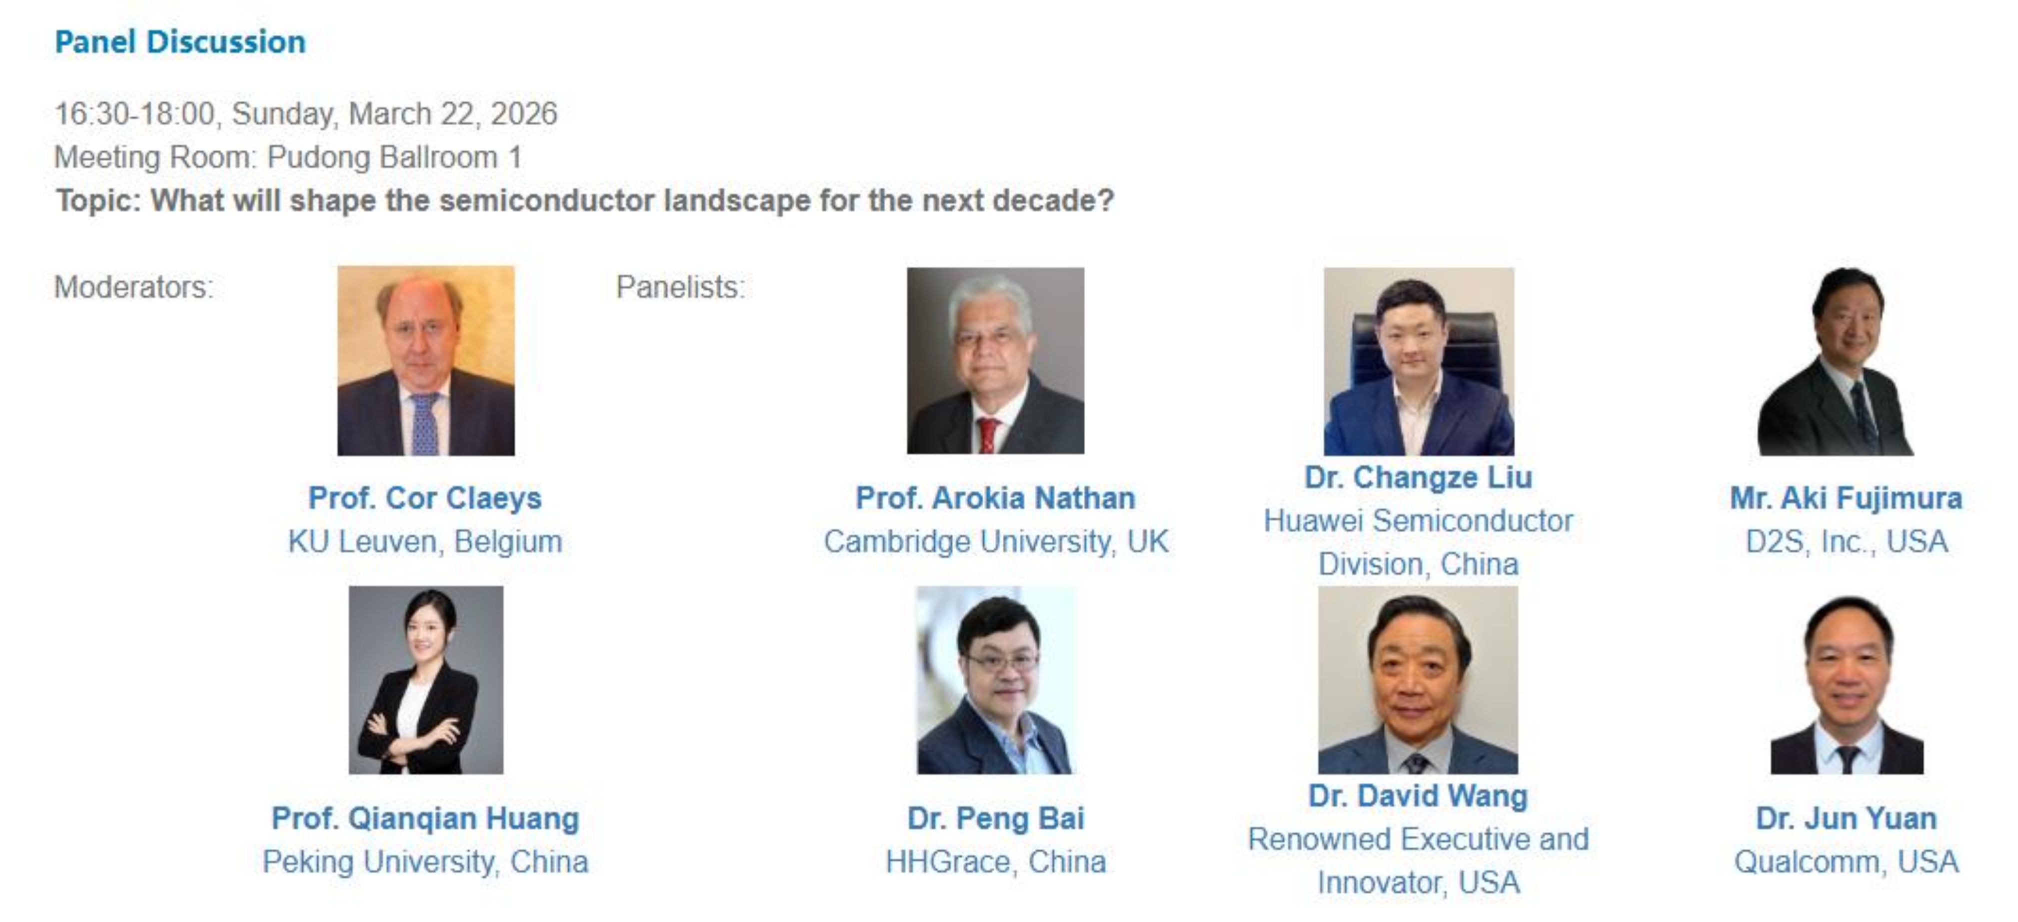

Panel Discussion

16:30-18:00, Sunday, March 22, 2026
Meeting Room: Pudong Ballroom 1
Topic: What will shape the semiconductor landscape for the next decade?

Moderators:

Prof. Cor Claeys
KU Leuven, Belgium

Prof. Qianqian Huang
Peking University, China

Panelists:

Prof. Arokia Nathan
Cambridge University, UK

Dr. Changze Liu
Huawei Semiconductor Division, China

Mr. Aki Fujimura
D2S, Inc., USA

Dr. Peng Bai
HHGrace, China

Dr. David Wang
Renowned Executive and Innovator, USA

Dr. Jun Yuan
Qualcomm, USA

TSMC’s Fab 20, where N2 is being made

Jensen Huang Said This at GTC DC 2025 in October

We have now seen that inflection point.

Accelerated computing: its moment has now arrived.

However, accelerating computing is a fundamentally different programming model.

You can’t just take a CPU software…and put it onto a GPU and have it run properly.

In fact, if you just did that, it’ll actually run slower

So, you have to reinvent new algorithms.

You have to create new libraries.

You have to, in fact, rewrite the application, which is the reason why it has taken so long.

It’s taken us nearly 30 years to get here.

But we did it one domain at a time.

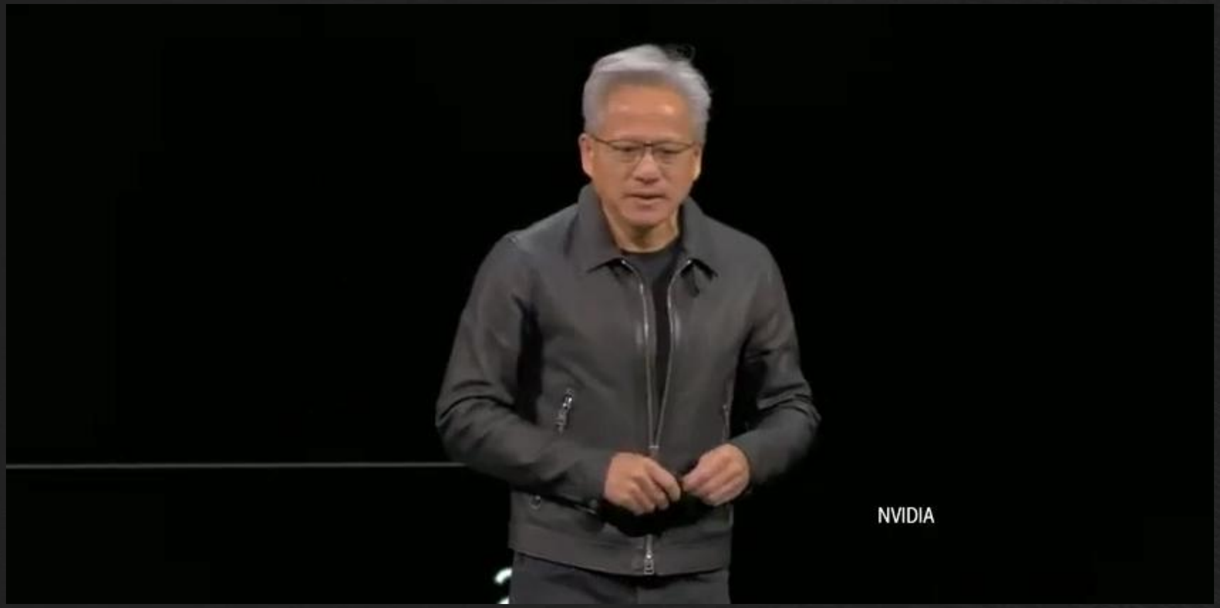

That’s Exactly What D2S Has Done in Our Domain

Reinventing new algorithms for GPU-accelerated computing

- Singular focus since 2009 : GPU-Acceleration
- GPU makes simulation-based processing super-fast
  - Like Jensen said, we reinvented the algorithms, computing in the pixel-domain, in CUDA
  - Manhattan and curvy shapes in the same time and accuracy

The Preferred GPU-Acceleration Partner for Semiconductor Manufacturing

Curvilinear Solutions in the Pixel Domain

TECHNOLOGY
GPU Acceleration

PRODUCT
TrueMask® ILT

LEARN MORE
Curvilinear

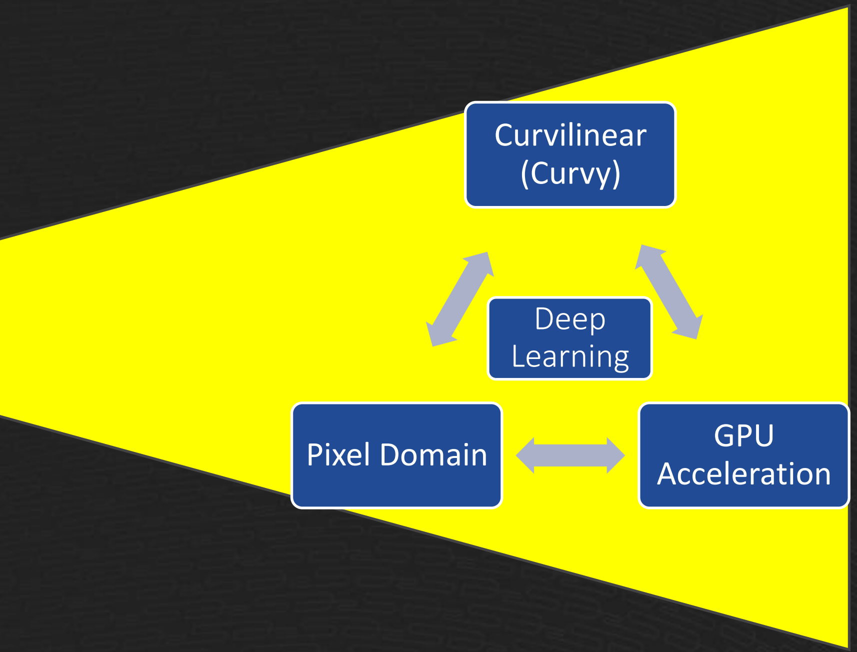

Curvilinear (Curvy)

Deep Learning

Pixel Domain

GPU Acceleration

D2S PATENTED TECHNOLOGY

21 years in Physical Design EDA, then 21 years in Manufacturing

Aki Fujimura, Chairman and CEO, D2S, Inc., SPIE Fellow and current BACUS President

1979 1982 1984 1986 1987 1991 1993 1997 1999 2002 2006 2009 2019 2026

AI Student
BS, MS @MIT
Co-op in DA
A* router
@Honeywell

@Trilogy

Founder/MTS
@Tangent

VP Eng.
Tangate
(First commercial
over-the-cell
P&R system)

@Cadence

@Pure

@Simplex

CTO@Cadence

@D2S

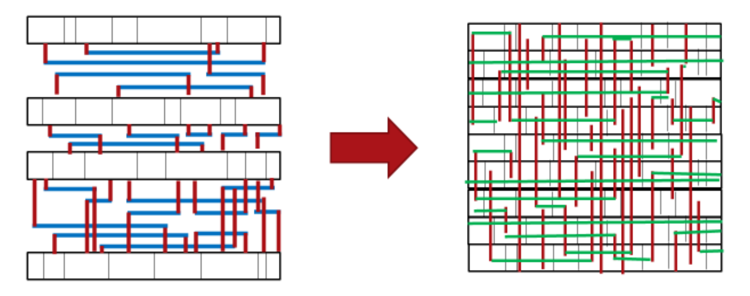

DEF, LEF, ECO, …

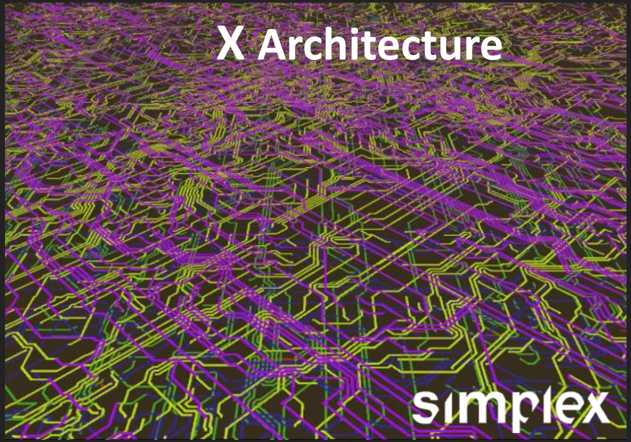

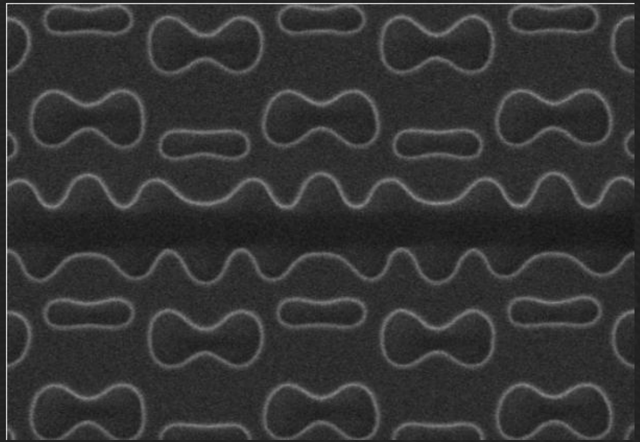

1980s was also the Decade of Automated Design

The basic scheme of automated design was set then

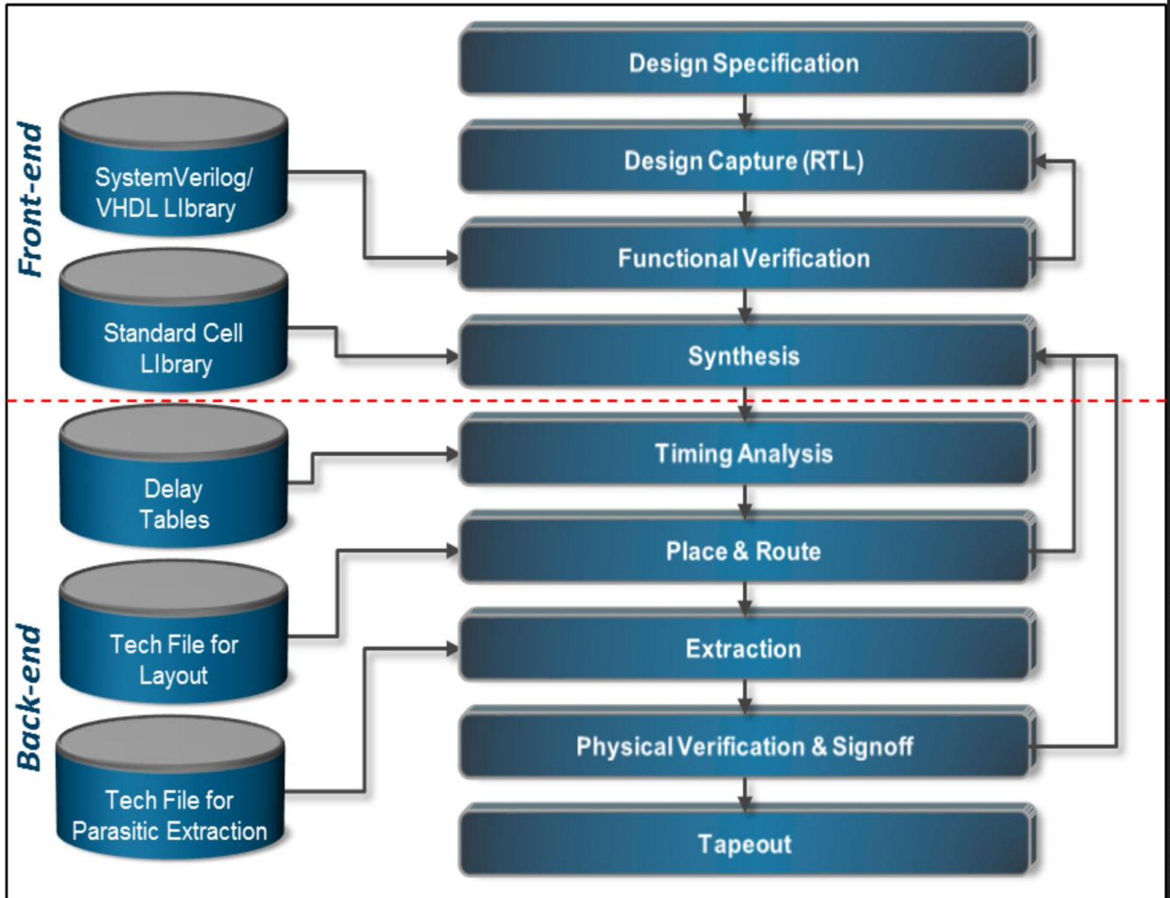

Front-end

Back-end

SystemVerilog/ VHDL LIbrary

Standard Cell LIbrary

Delay Tables

Tech File for Layout

Tech File for Parasitic Extraction

Design Specification

Design Capture (RTL)

Functional Verification

Synthesis

Timing Analysis

Place & Route

Extraction

Physical Verification & Signoff

Tapeout

EDA Design Flow *(simplified)*

- Fully synchronous design
  - Verilog and VHDL
  - Logic simulation and synthesis
  - Scan-based test
  - Timing analysis and assurance
- >2 Layers of Metal
  - Over-the-cell routing
  - LEF, DEF, ECO…
- Design rule checking
  - Basis for mask rule checking later
- Workstations and ISVs with TBLs*

---

* Independent software vendors with term-based licenses

We Were Computationally Limited Back Then

IBM 3090: 30 MIPS, 50 MFLOPS, 126GB disk, 128MB RAM, 100KWatts, $20M

Source: Cray

Source: IBM

We Were Computationally Limited Back Then

IBM 3090: 30 MIPS, 50 MFLOPS, 126GB disk, 128MB RAM, 100KWatts, $20M

Cray-2: 1.9GFLOPS (64bit); 500GB disk, 2GB RAM, 150KWatts, $30M

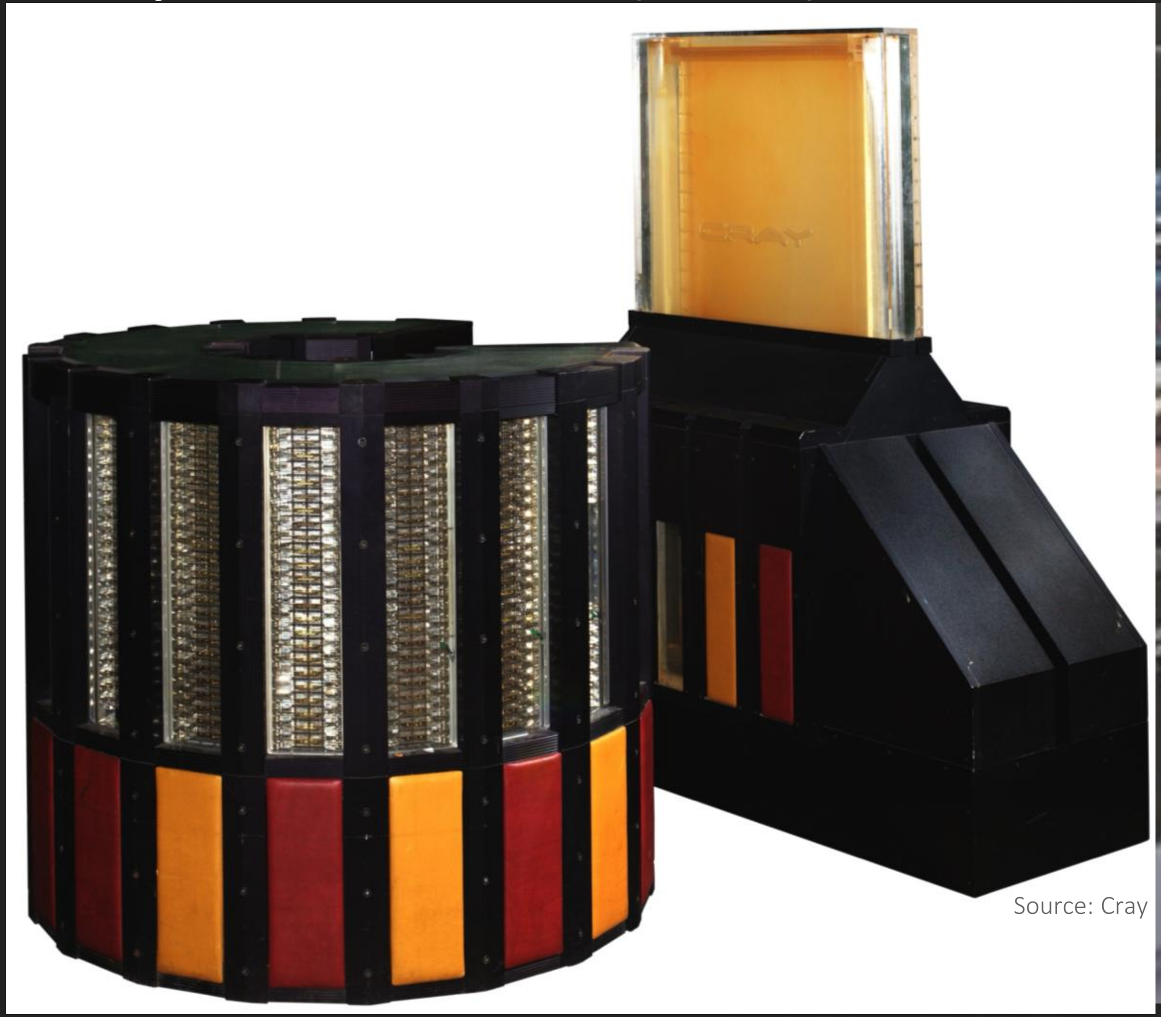

Source: Cray

Source: IBM

History is Repeating Itself

1980's

CPU

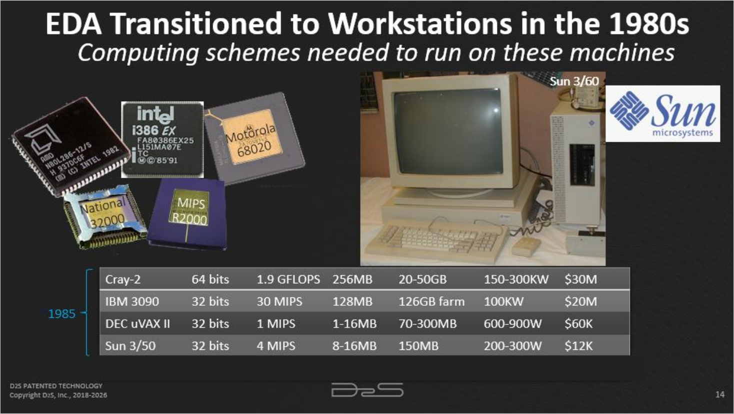

EDA Transitioned to Workstations in the 1980s

Computing schemes needed to run on these machines

1985

Cray-2	64 bits	1.9 GFLOPS	256MB	20-50GB	150-300KW	$30M
IBM 3090	32 bits	30 MIPS	128MB	126GB farm	100KW	$20M
DEC uVAX II	32 bits	1 MIPS	1-16MB	70-300MB	600-900W	$60K
Sun 3/50	32 bits	4 MIPS	8-16MB	150MB	200-300W	$12K

D2S PATENTED TECHNOLOGY
Copyright D2S, Inc., 2018-2026

14

Manhattan

Computing Limited Chip Routing to Manhattan

Needed to break away from PCB-based routing schemes

VSB

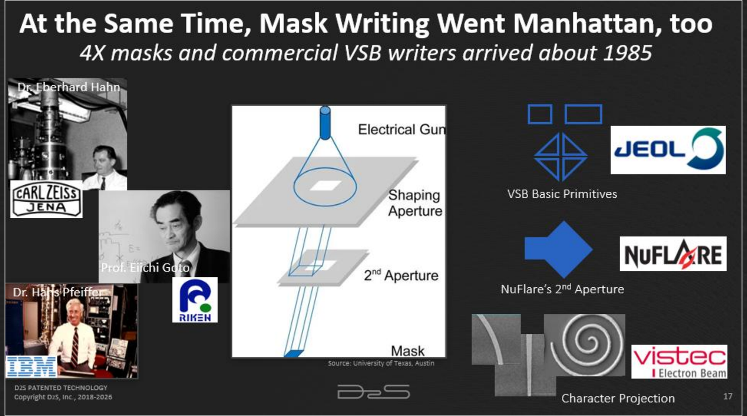

At the Same Time, Mask Writing Went Manhattan, too

4X masks and commercial VSB writers arrived about 1985

Source: University of Texas, Austin

D2S PATENTED TECHNOLOGY
Copyright D2S, Inc., 2018-2026

17

2020's

GPU

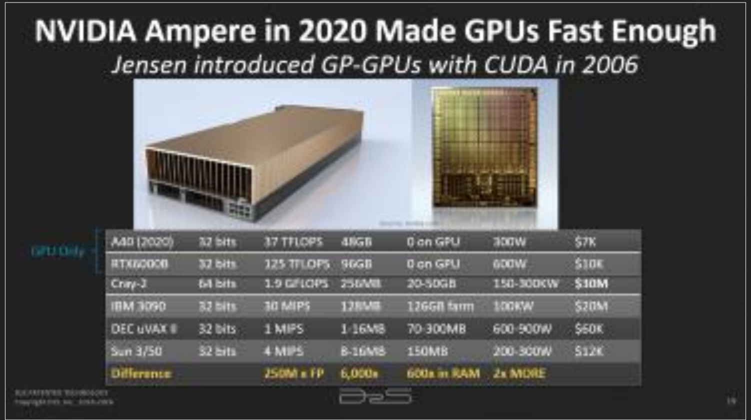

NVIDIA Ampere in 2020 Made GPUs Fast Enough

Jensen introduced GP-GPUs with CUDA in 2006

GPU Only

A40 (2020)	32 bits	37 TFLOPS	48GB	0 on GPU	300W	$7K
RTX6000B	32 bits	125 TFLOPS	96GB	0 on GPU	600W	$10K
Cray-2	64 bits	1.9 GFLOPS	256MB	20-50GB	150-300KW	$30M
IBM 3090	32 bits	30 MIPS	128MB	126GB farm	100KW	$20M
DEC uVAX II	32 bits	1 MIPS	1-16MB	70-300MB	600-900W	$60K
Sun 3/50	32 bits	4 MIPS	8-16MB	150MB	200-300W	$12K
Difference		250M x FP	6,000x	600x in RAM	2x MORE	

Curvy

Now That Computing Isn't Limited

Entirely curvilinear masks are inevitable

Multibeam

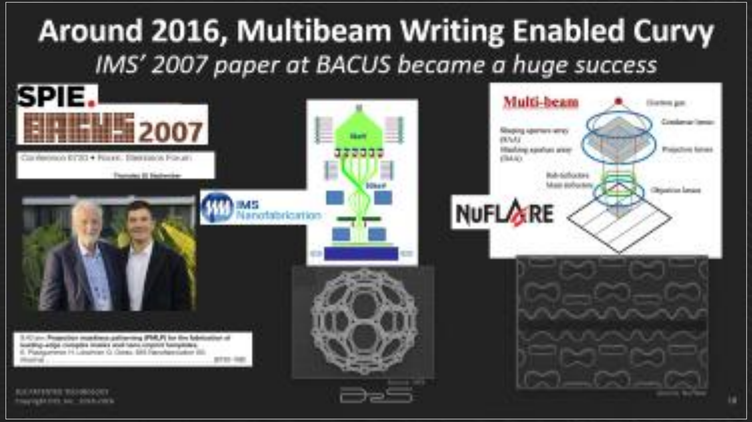

Around 2016, Multibeam Writing Enabled Curvy

IMS' 2007 paper at BACUS became a huge success

EDA Transitioned to Workstations in the 1980s

Computing schemes needed to run on these machines


Sun 3/60

~1985

Cray-2	64 bits	1.9 GFLOPS	256MB	20-50GB	150-300KW	$30M
IBM 3090	32 bits	30 MIPS	128MB	126GB farm	100KW	$20M
DEC uVAX II	32 bits	1 MIPS	1-16MB	70-300MB	600-900W	$60K
Sun 3/50	32 bits	4 MIPS	8-16MB	150MB	200-300W	$12K

EDA Transitioned to Workstations in the 1980's

We had to design to machines 4000x less than Apple Watch

~1985

Cray-2	64 bits	1.9 GFLOPS	256MB	20-50GB	150-300KW	$30M
IBM 3090	32 bits	30 MIPS	128MB	126GB farm	100KW	$20M
DEC uVAX II	32 bits	1 MIPS	1-16MB	70-300MB	600-900W	$60K
Sun 3/50	32 bits	4 MIPS	8-16MB	150MB	200-300W	$12K
Apple Watch	64 bits	2 GFLOPS	1GB	64GB	1-2W	$400

Computing Limited Chip Routing to Manhattan

Needed to break away from PCB-based routing schemes

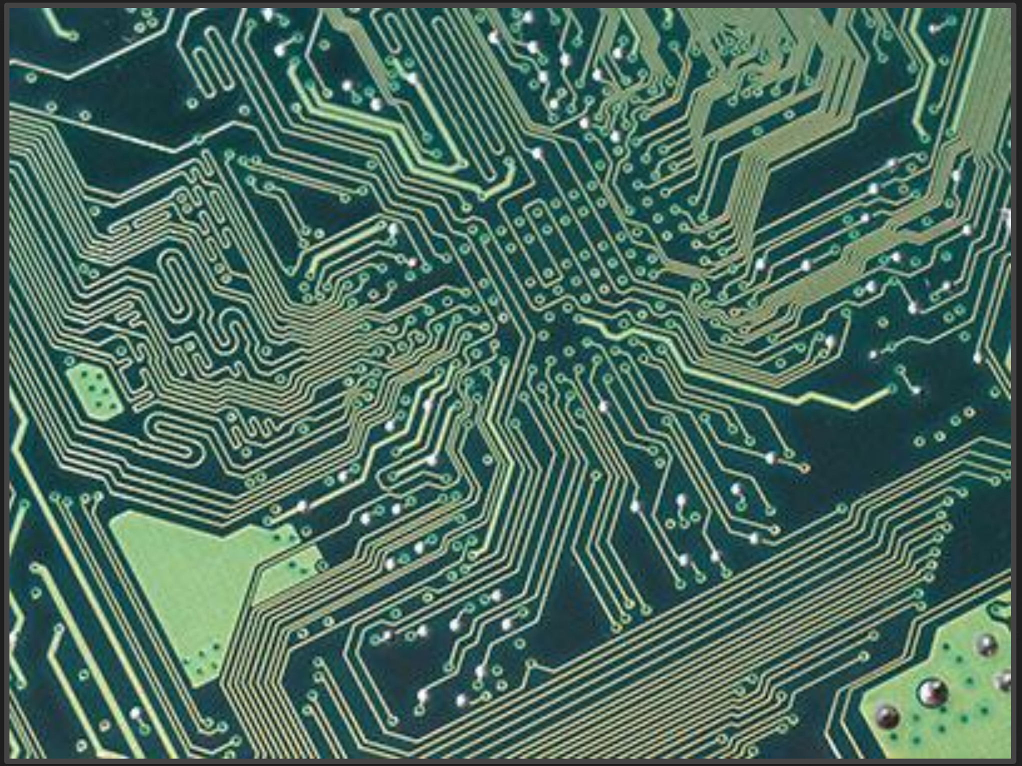

Source: ourpcb.com

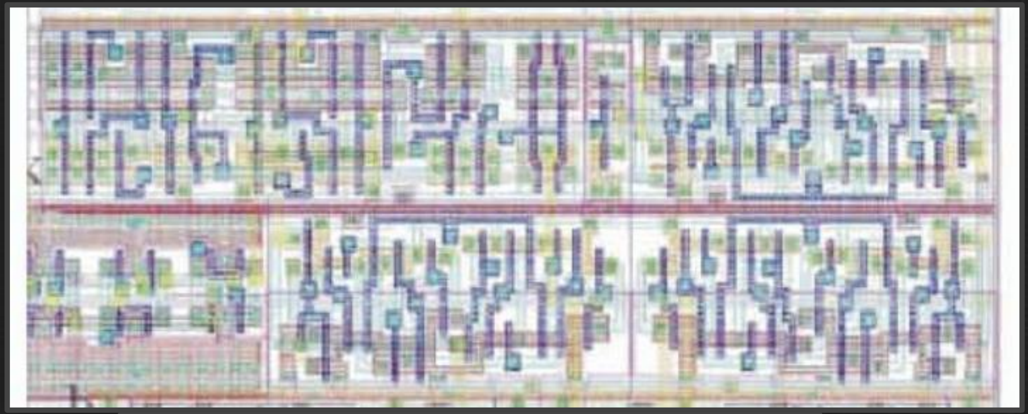

Source: Vladimir Stojanovic, course materials for 6.973 Communication System Design, Spring 2006. MIT OpenCourseWare

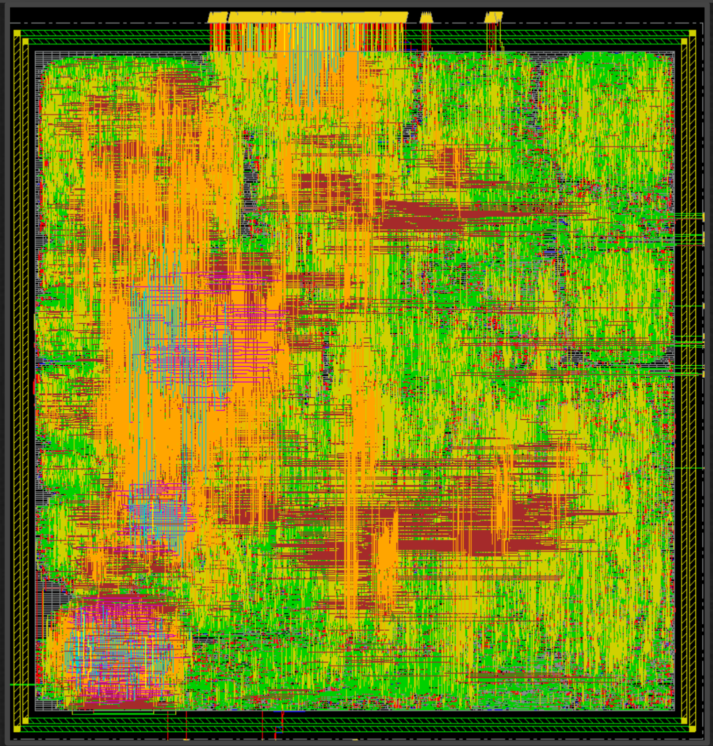

Source: University of Bremen

At the Same Time, Mask Writing Went Manhattan, too

4X masks and commercial VSB writers arrived about 1985

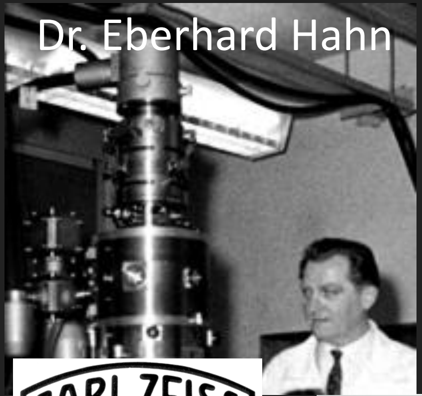

Dr. Eberhard Hahn

Prof. Eiichi Goto

Dr. Hans Pfeiffer

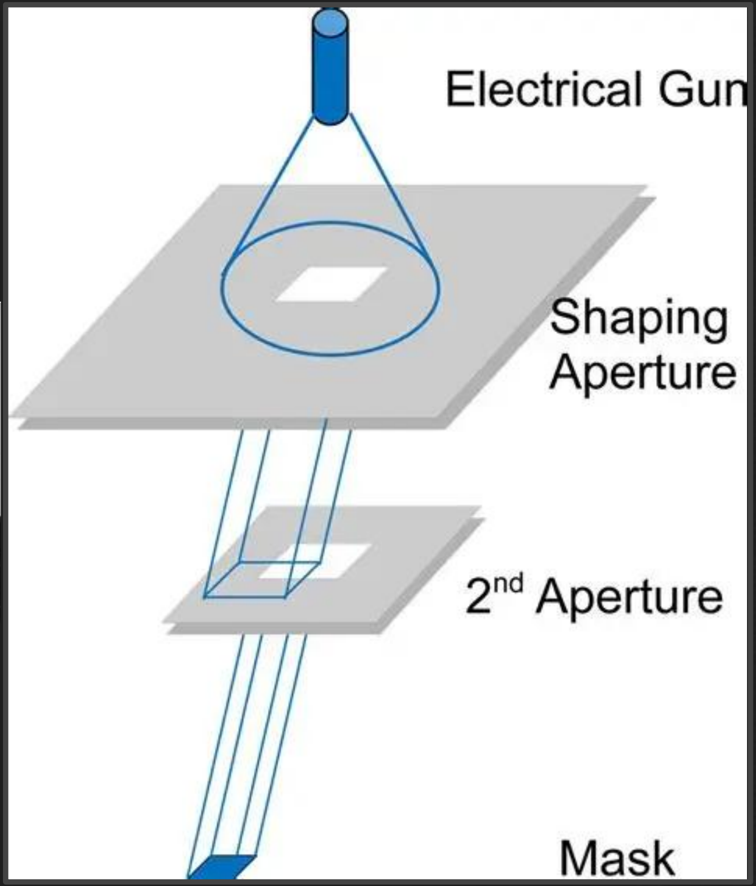

Source: University of Texas, Austin

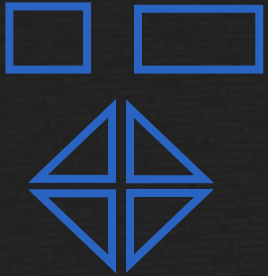

VSB Basic Primitives

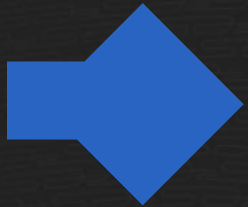

NuFlare's 2nd Aperture

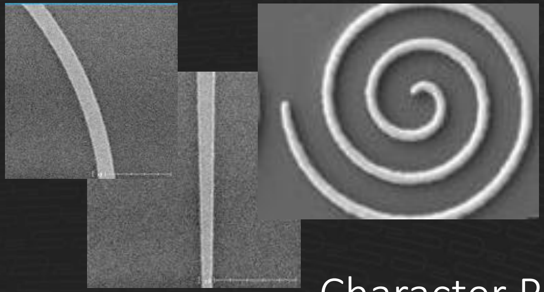

Character Projection

Around 2016, Multibeam Writing Enabled Curvy

IMS' 2007 paper at BACUS became a huge success

SPIE.

BACUS 2007

Conference 6730 • Room: Steinbeck Forum

Thursday 20 September

9:40 am: **Projection maskless patterning (PMLP) for the fabrication of leading-edge complex masks and nano-imprint templates,**
E. Platzgummer, H. Löschner, G. Gross, IMS Nanofabrication AG (Austria) . . . . [6730-108]

IMS
Nanofabrication

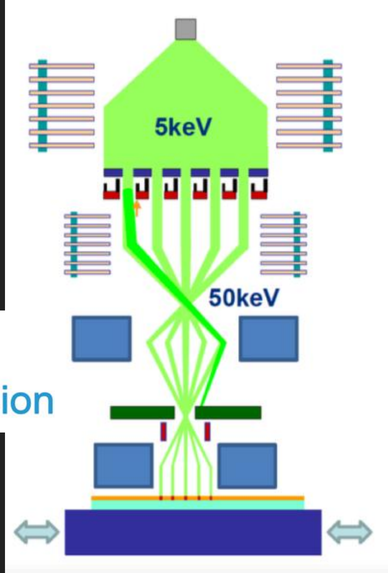

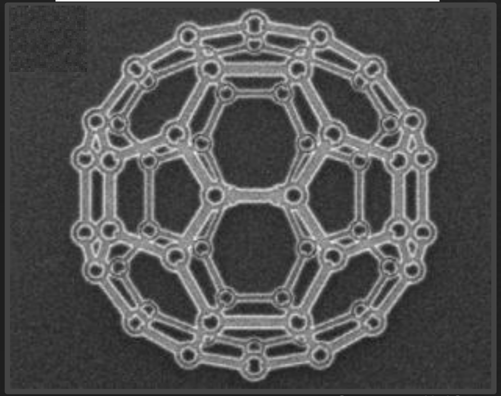

Source: IMS

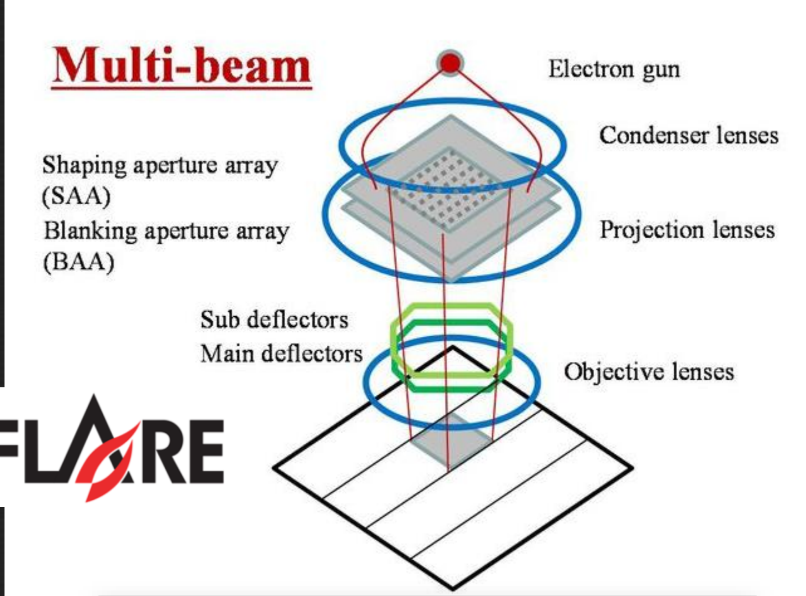

NUFLARE

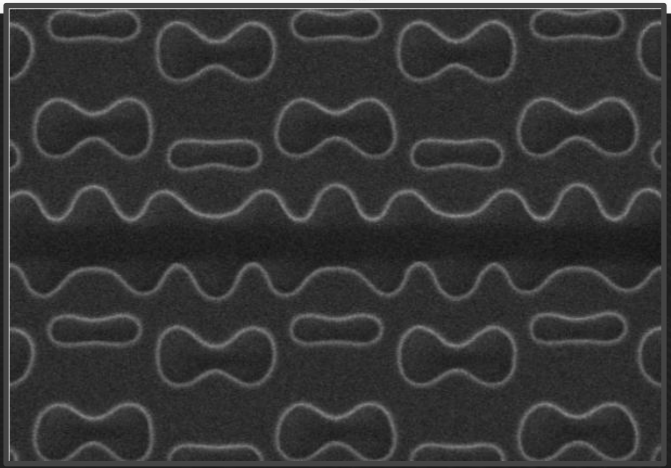

Source: NuFlare

NVIDIA Ampere in 2020 Made GPUs Fast Enough

Jensen introduced GP-GPUs with CUDA in 2006

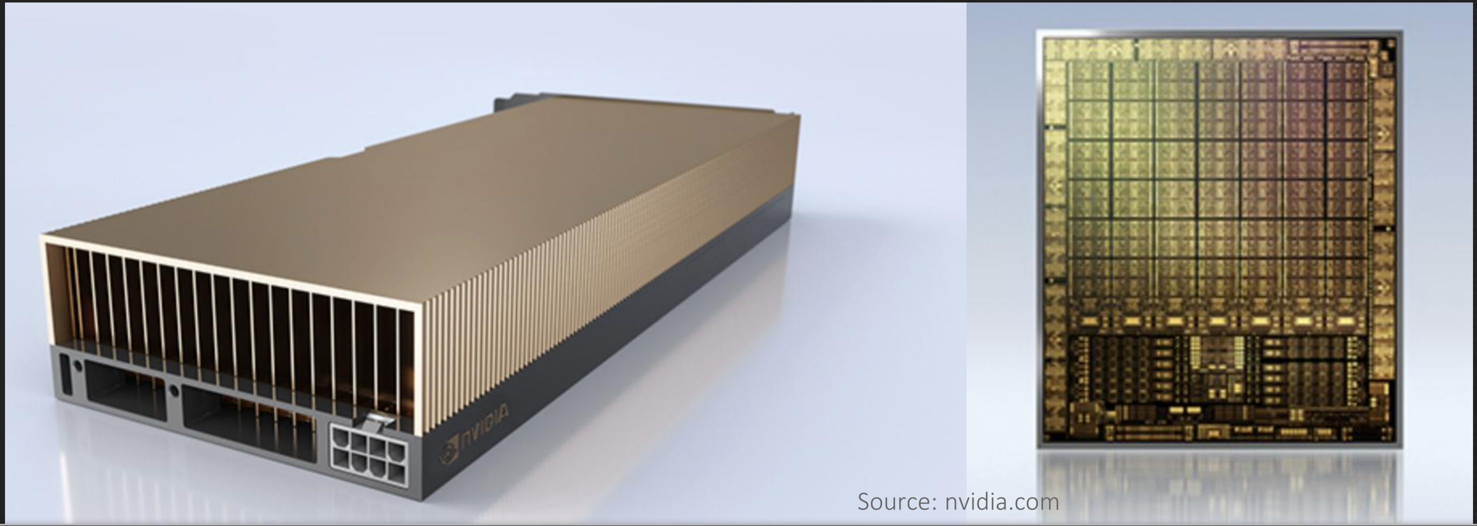

Source: nvidia.com

GPU Only (A40, RTX6000B)

A40 (2020)	32 bits	37 TFLOPS	48GB	0 on GPU	300W	$7K
RTX6000B	32 bits	125 TFLOPS	96GB	0 on GPU	600W	$10K
Cray-2	64 bits	1.9 GFLOPS	256MB	20-50GB	150-300KW	**$30M**
IBM 3090	32 bits	30 MIPS	128MB	126GB farm	100KW	$20M
DEC uVAX II	32 bits	1 MIPS	1-16MB	70-300MB	600-900W	$60K
Sun 3/50	32 bits	4 MIPS	8-16MB	150MB	200-300W	$12K
Difference		**250M x FP**	**6,000x**	**600x in RAM**	**2x MORE**	

D2S PATENTED TECHNOLOGY

Now That Computing Isn’t Limited

Entirely curvilinear masks are inevitable

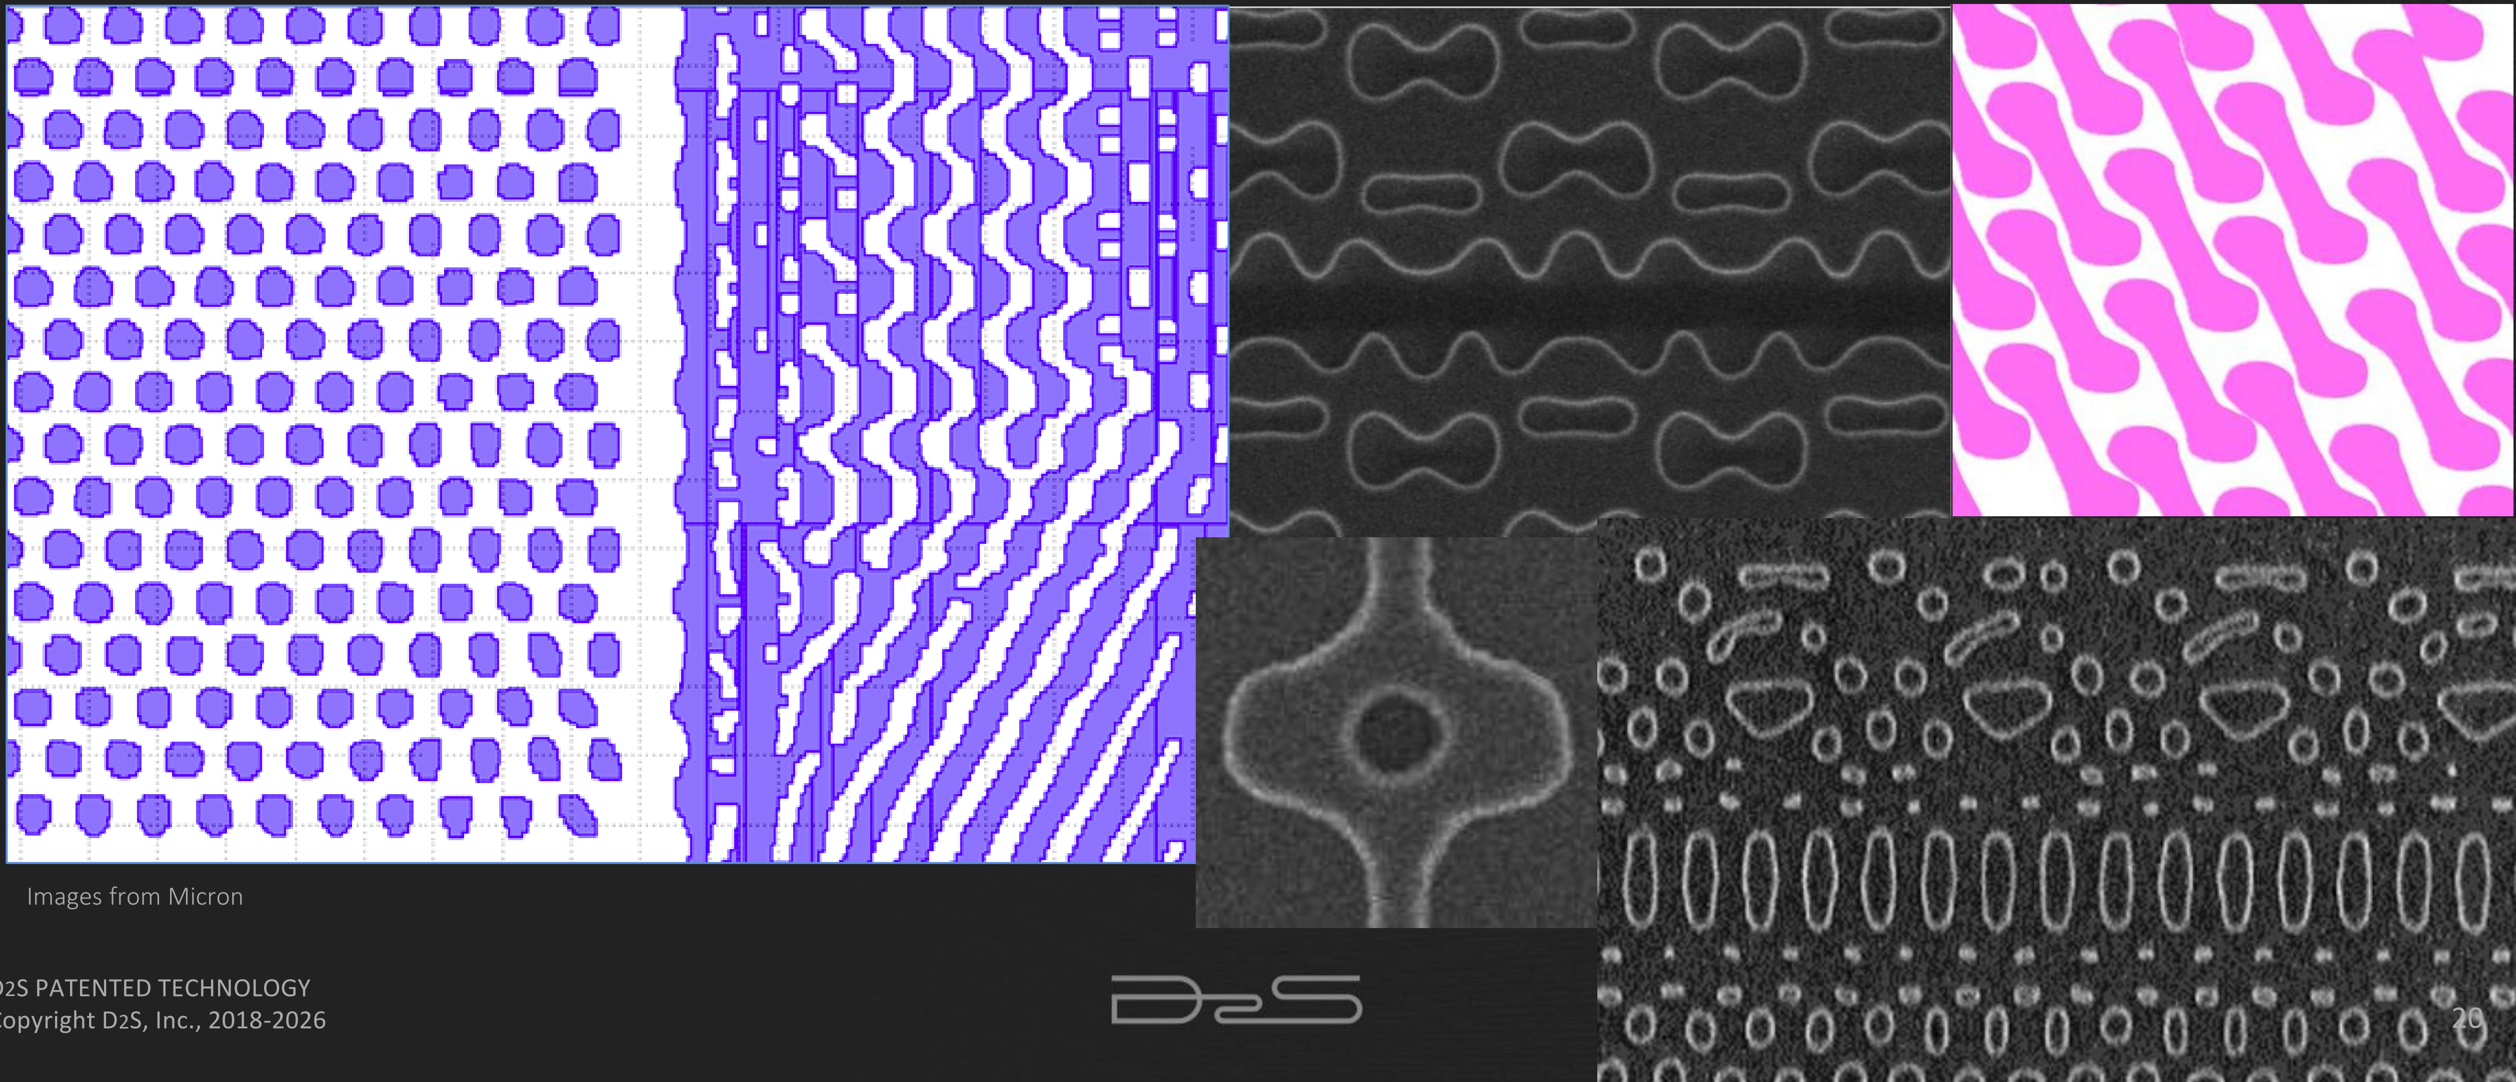

Images from Micron

Only Entirely Curvilinear Masks are Manufacturable

But not all curvy masks are entirely manufacturable masks

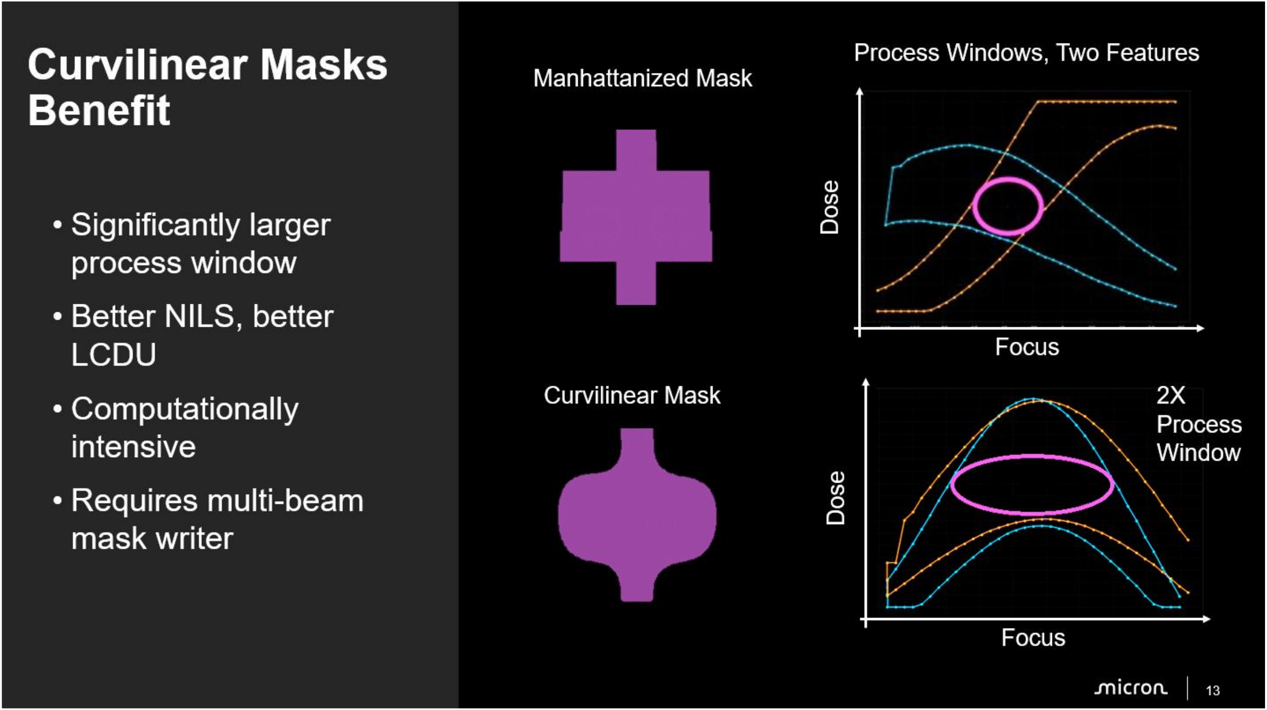

Curvilinear Masks Benefit

- Significantly larger process window
- Better NILS, better LCDU
- Computationally intensive
- Requires multi-beam mask writer

Source: Micron

- Four reasons for curvy masks
  1. 2X Process Window
  2. Manufacturable masks: Ask for what you can get, and get what you asked for
  3. Manufacturable masks: Best area uniformity on mask for best wafer CDU
  4. Enables Curvy Designs

D2S PATENTED TECHNOLOGY

Wafer Designs are Made So Only 1D CD is Important

Design is resilient to misalignment and corner rounding

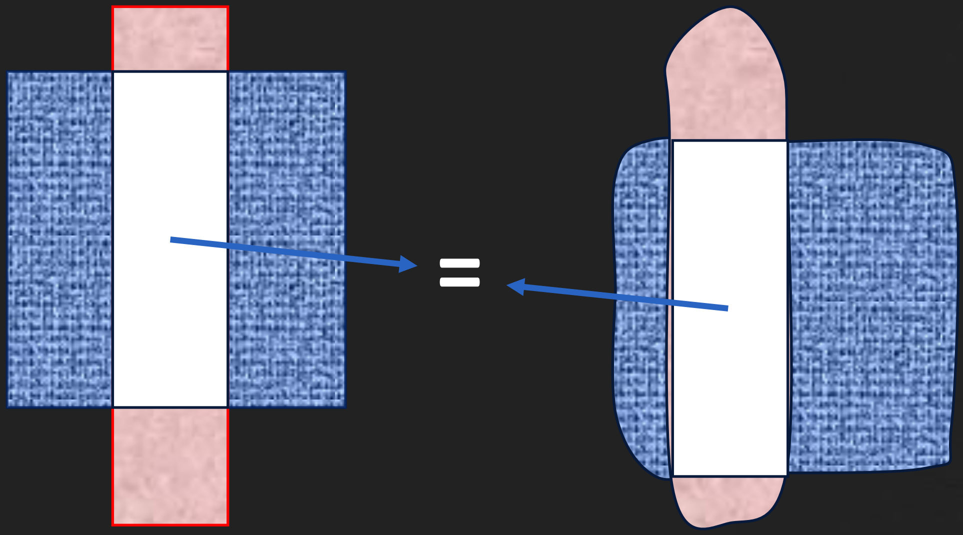

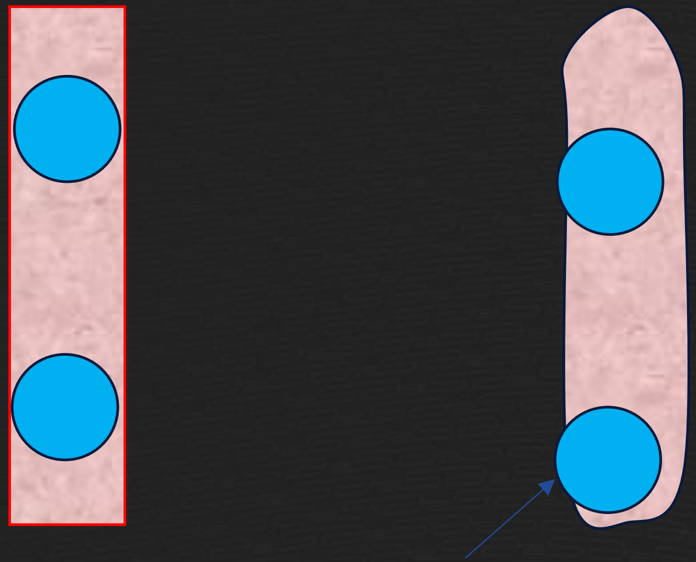

This is bad, but it’s a compromise

But that’s not true for mask…On mask, 2D area uniformity is critical

Wafer Accuracy is About Area Accuracy on Mask

So smaller shapes we need to handle now can't be doing 1D on mask

This is already true for all decorated OPC masks before Curvy...

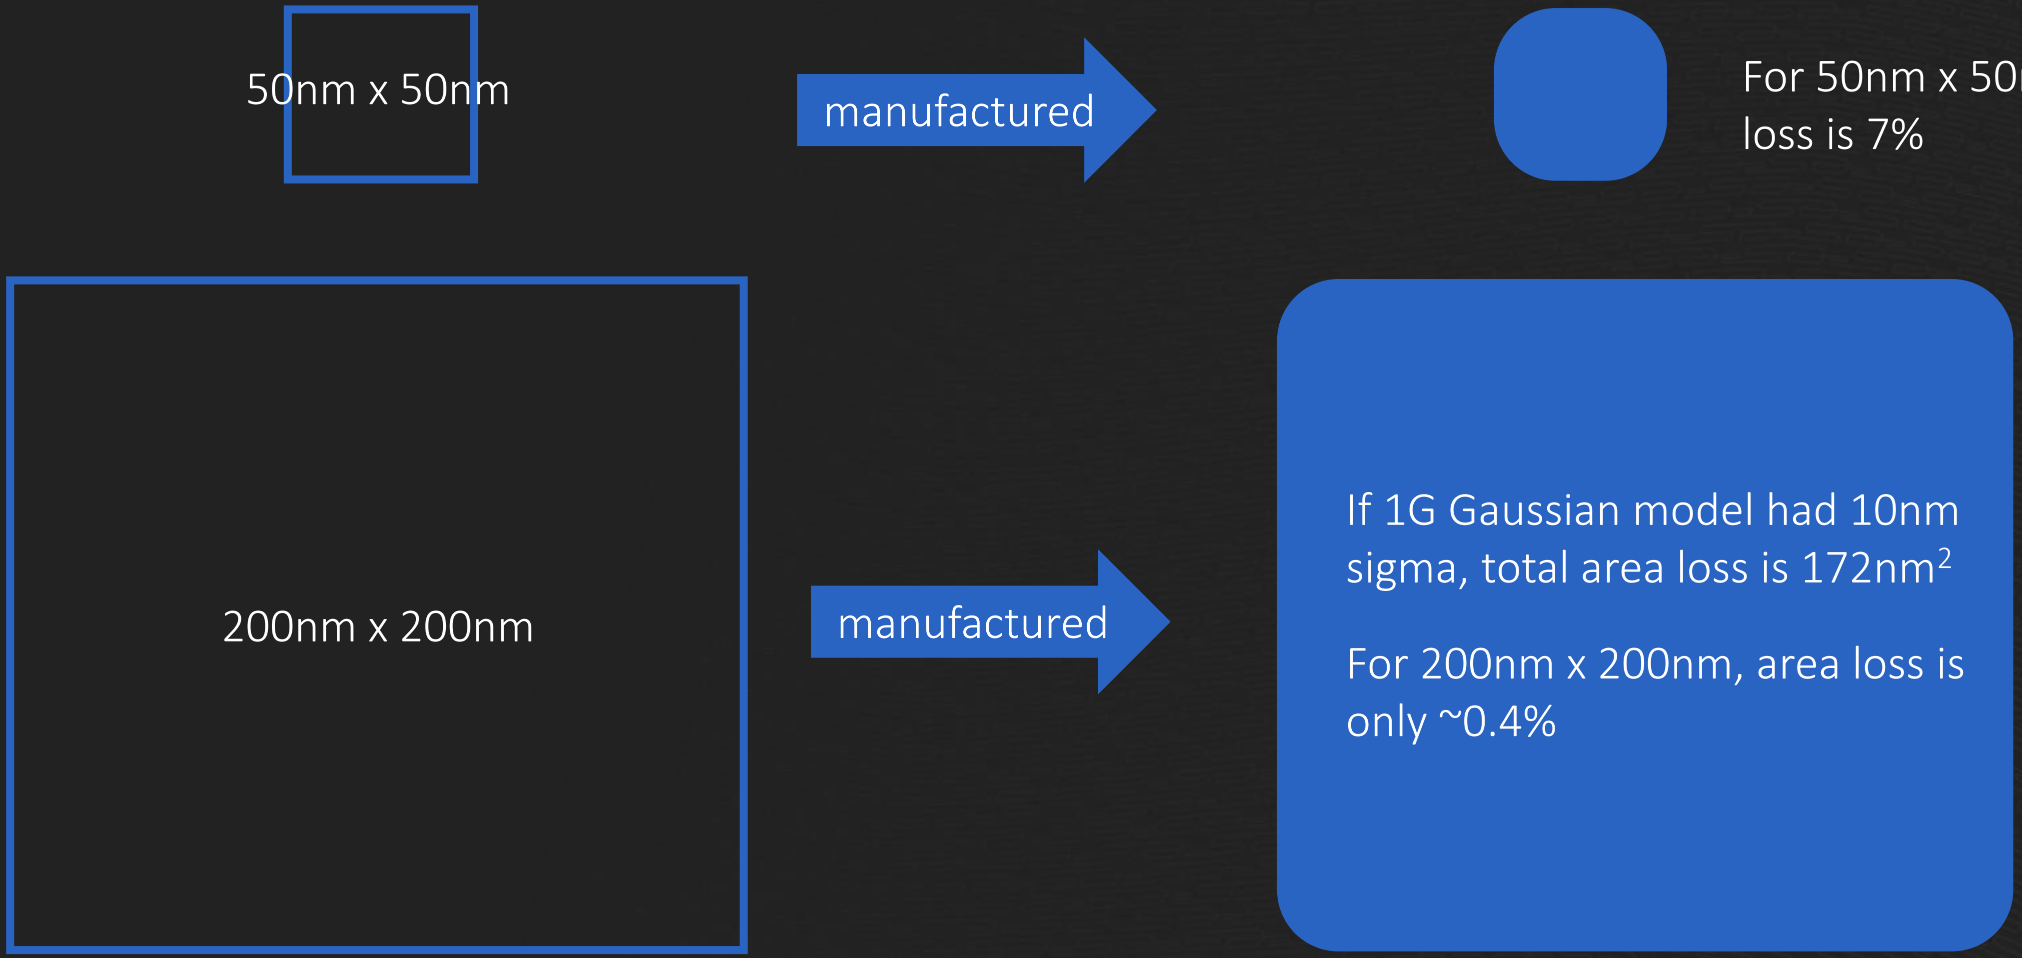

50nm x 50nm

manufactured

For 50nm x 50nm, area loss is 7%

200nm x 200nm

manufactured

If 1G Gaussian model had 10nm sigma, total area loss is 172nm^2

For 200nm x 200nm, area loss is only ~0.4%

Unmanufacturable Shapes Always Have Bad Dose Margin

Instance-to-instance mask area uniformity is critical for wafer CDU

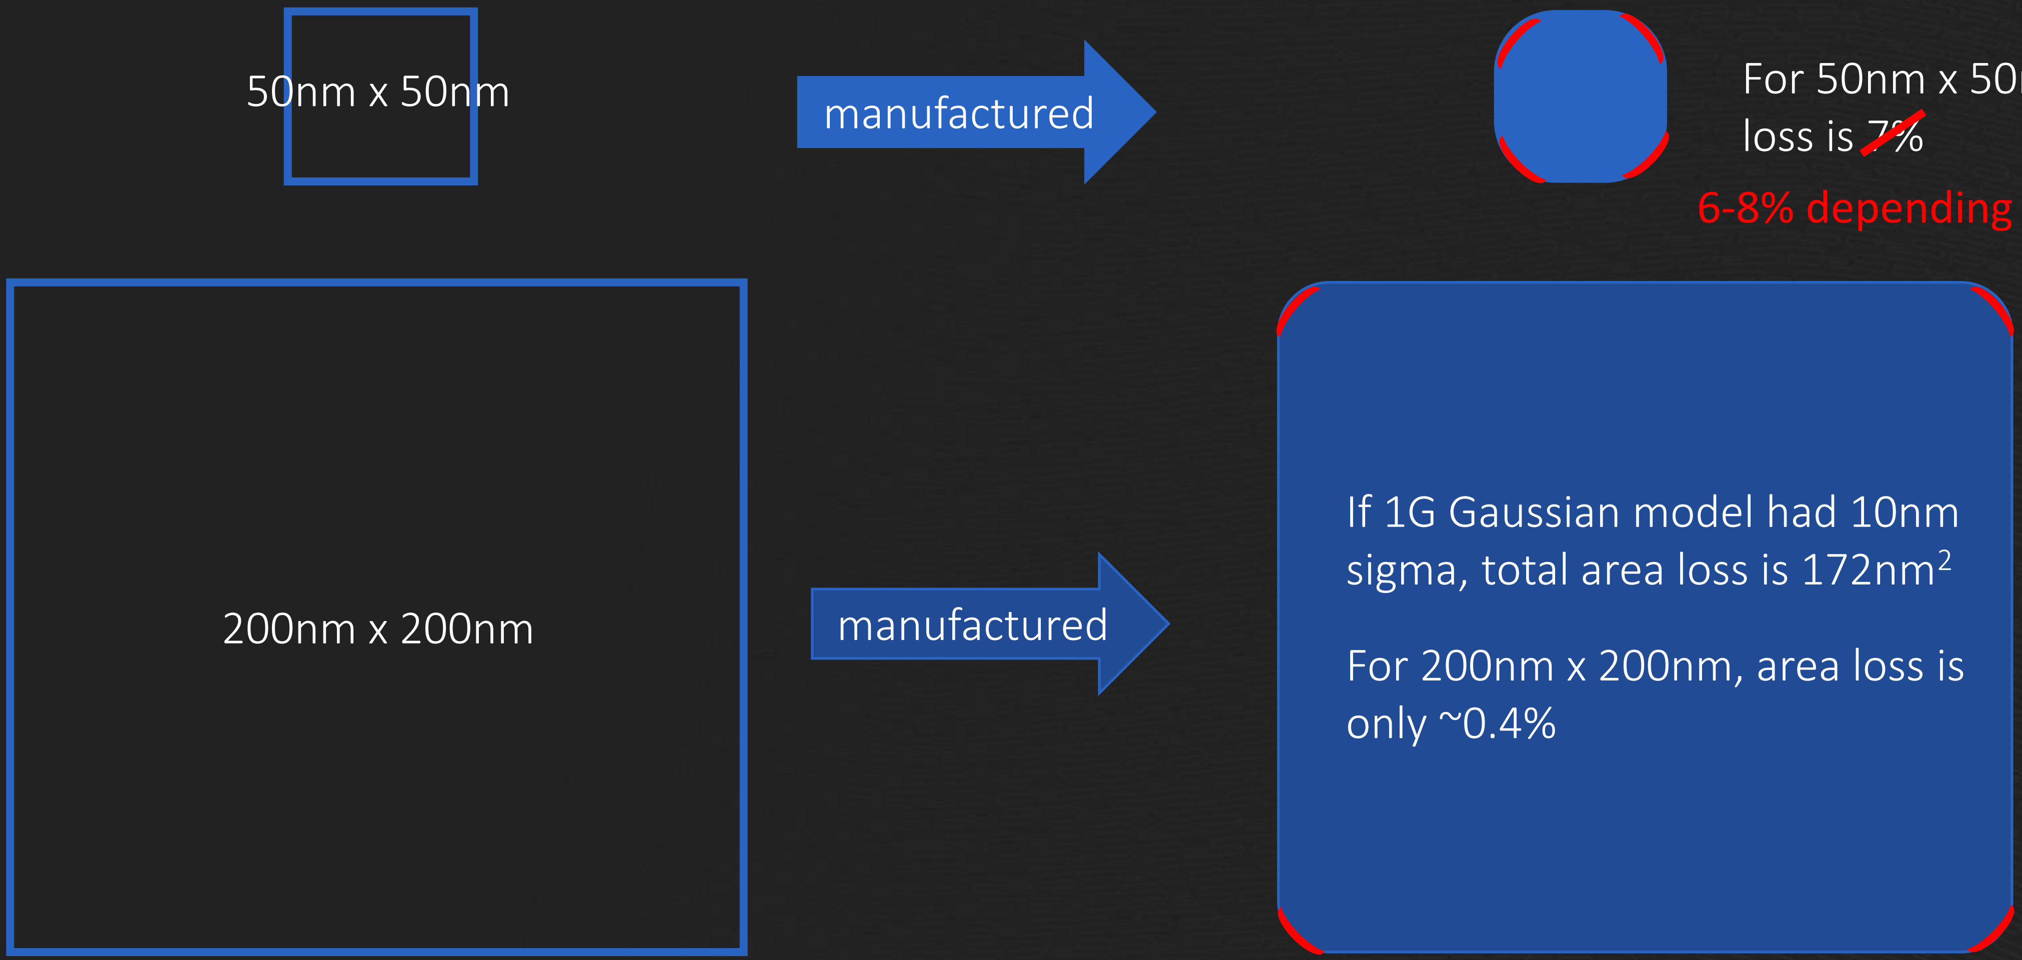

50nm x 50nm

manufactured

Poor dose margin on corners

For 50nm x 50nm, area loss is ~~7%~~

6-8% depending on instance

200nm x 200nm

manufactured

If 1G Gaussian model had 10nm sigma, total area loss is 172nm^2

For 200nm x 200nm, area loss is only ~0.4%

Much Better to Ask for What You Can Get

Instance-to-instance variation is critical for manufacturing

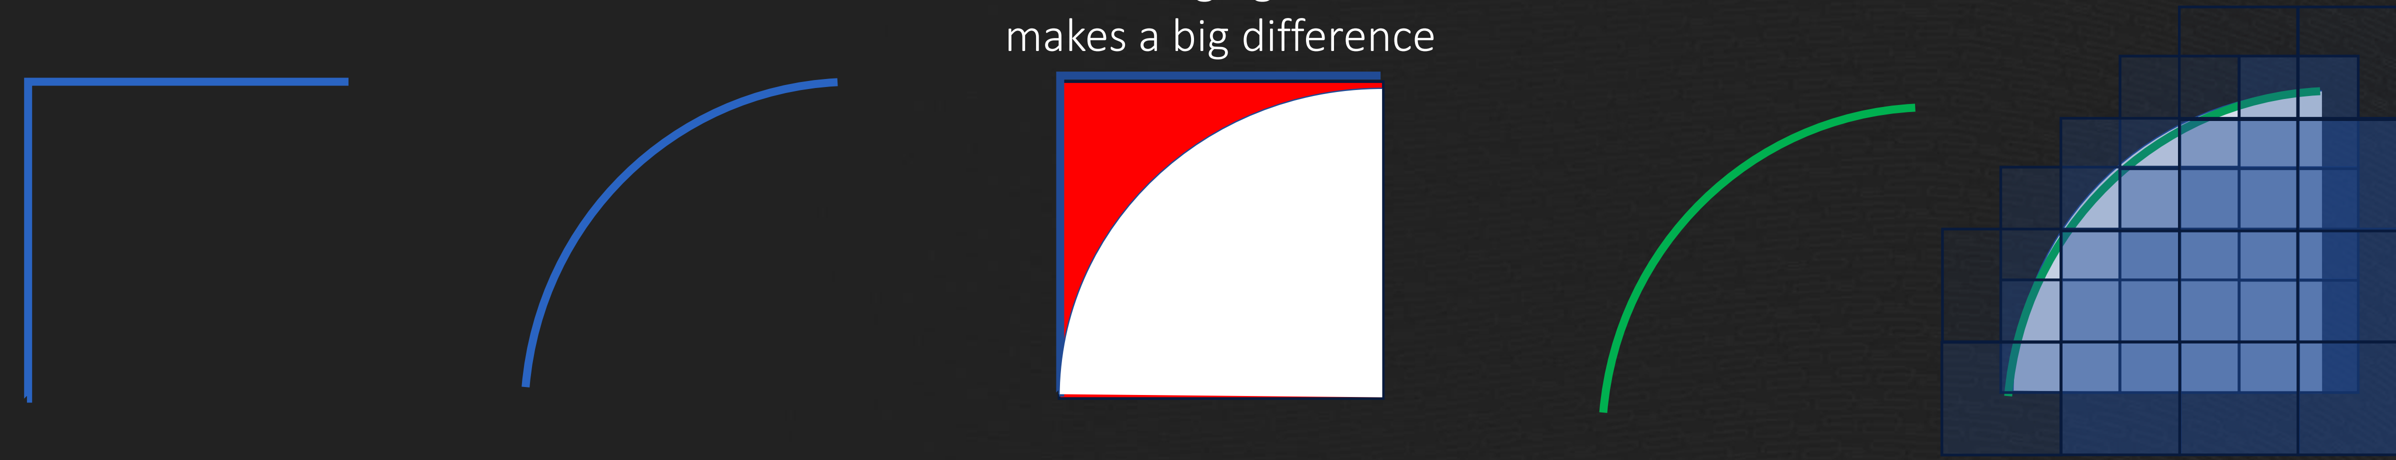

Trying to manufacture this

and getting this

has more variation
because the dose of this
area changing +/-x%
makes a big difference

Best to ask for this

and get this

If you ask for what you can get…

You get better mask area uniformity
and therefore, better wafer CDU

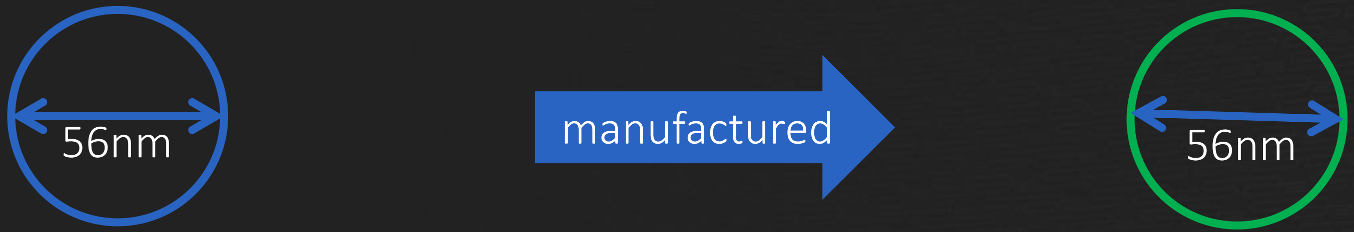

56nm

manufactured

56nm

Multi-Beam or VSB, Design What Can be Manufactured so you can Manufacture it Exactly and get Less Variation

Multi-Beam Writing Needs Pixel-Level Dose Correction (PLDC)

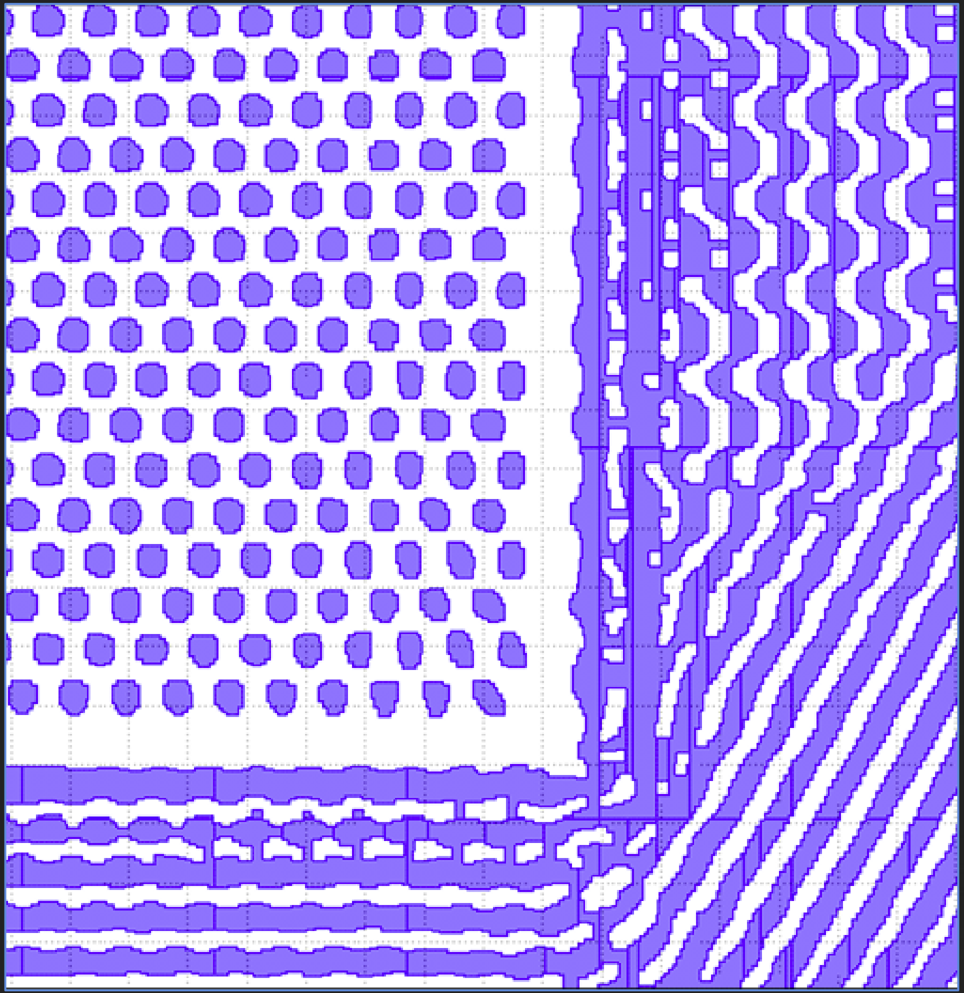

Ezequiel Russel, Micron, et al, BACUS 2023

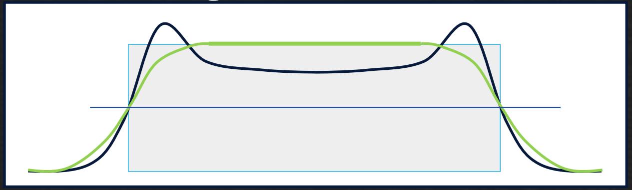

- 11% - 22% better LCDU*
- 20%+ better GCDU**
- Linearity correction including variable etch bias
- Full-reticle curvilinear or all Manhattan = same compute time

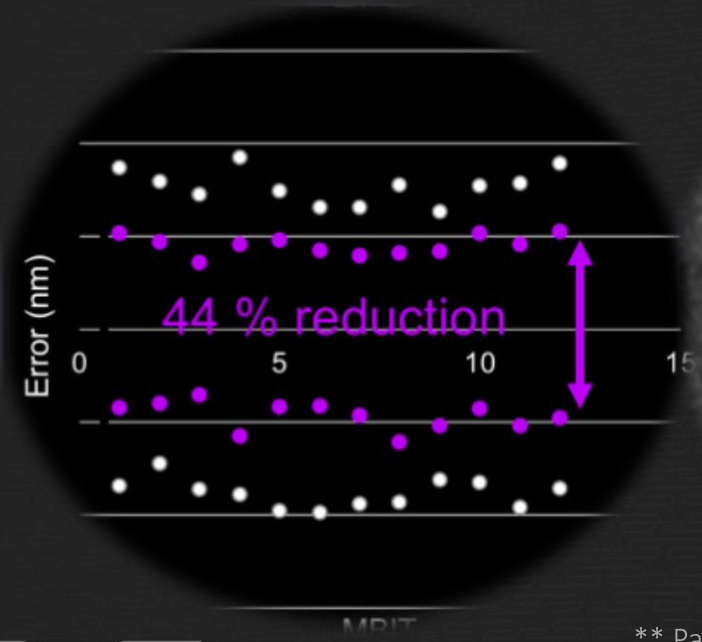

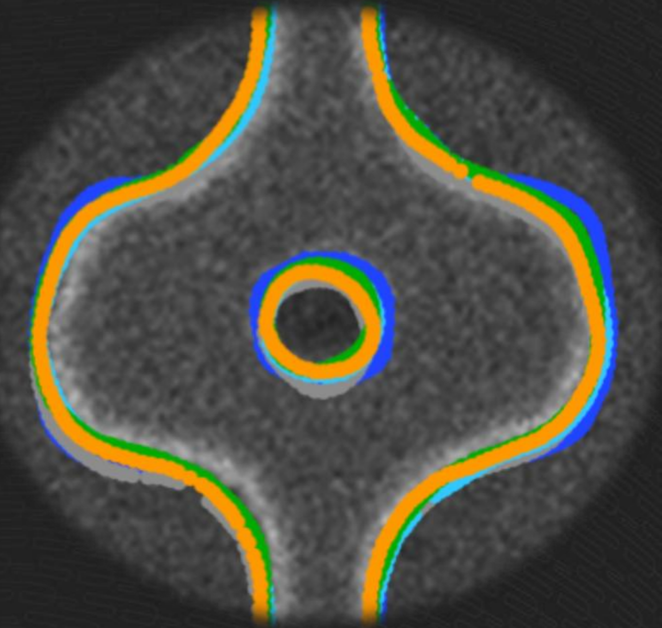

** Paris Spinelli, Micron, et al, Photomask Japan, 2025

* Matsumoto, M., et al, “Mask performance improvement by pixel level dose correction,” Proc. SPIE 13216, 132160H (2024)

Geometric Effects are Corrected by nm-Bias
Dose Effects are Corrected by Dose-Bias

Beam Blur + +

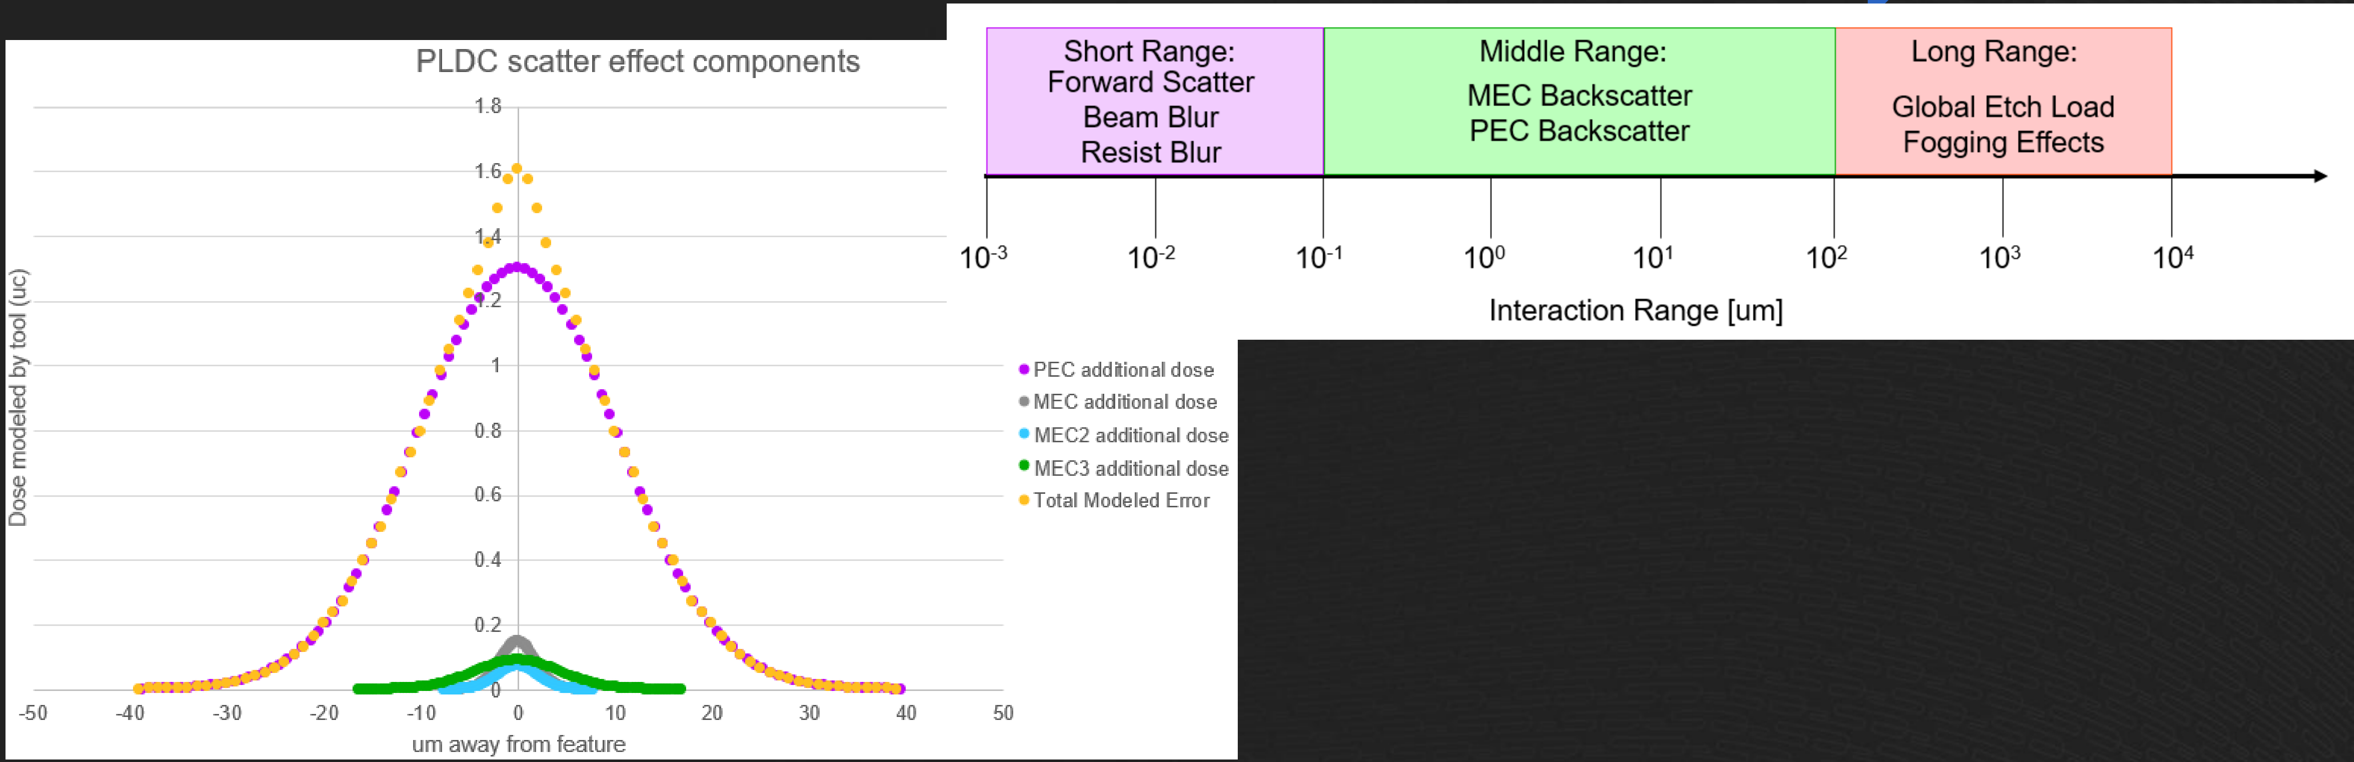

Geometric Effects are Corrected by nm-Bias
Dose Effects are Corrected by Dose-Bias


Beam Blur + Edge Dose Enhancement +

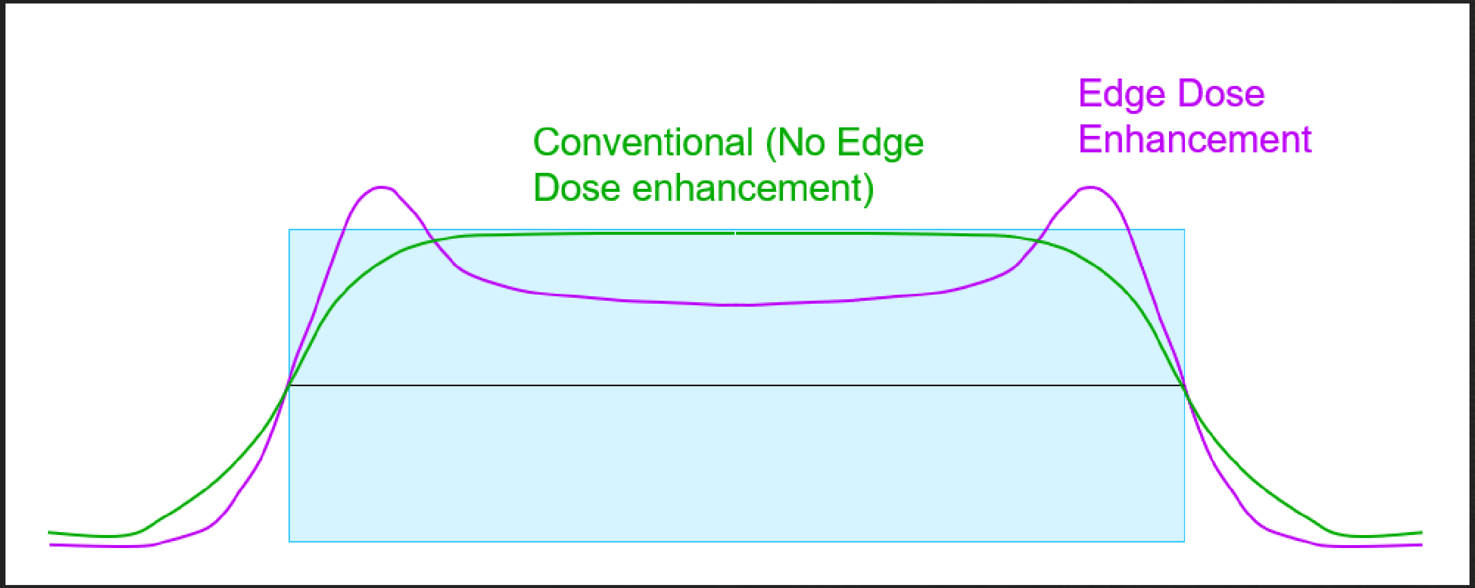

Geometric Effects are Corrected by nm-Bias
Dose Effects are Corrected by Dose-Bias

Beam Blur + Edge Dose Enhancement + Variable Etch Bias

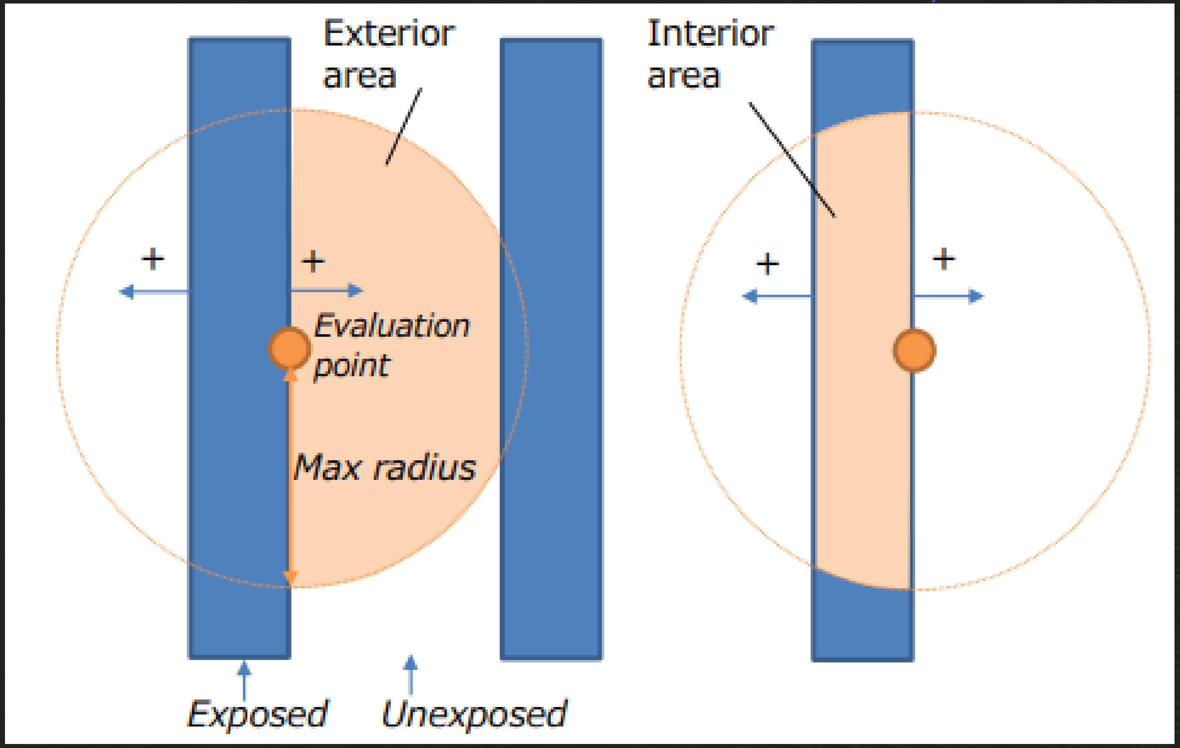

D2S PATENTED TECHNOLOGY

PLDC Improves Uniformity and Linearity

For masks written on multi-beam mask writers

Beam Blur + Edge Dose Enhancement + Variable Etch Bias

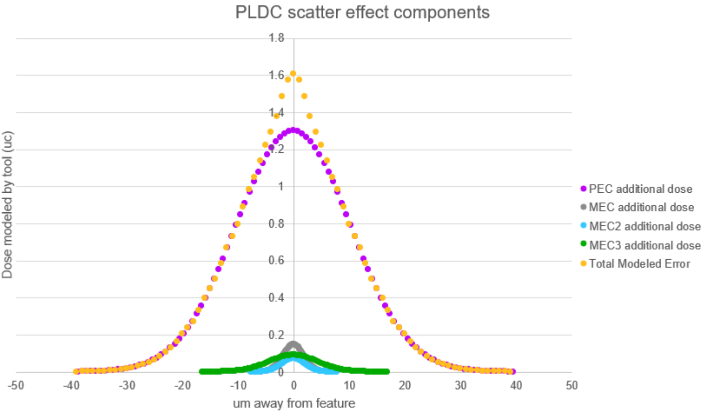

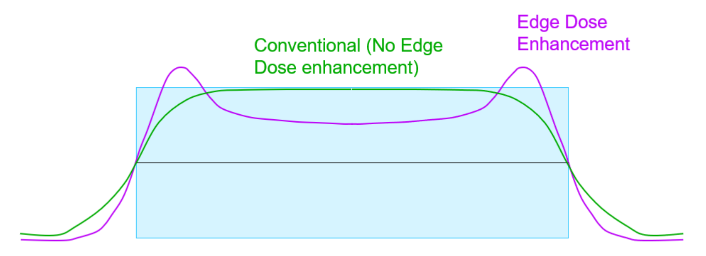

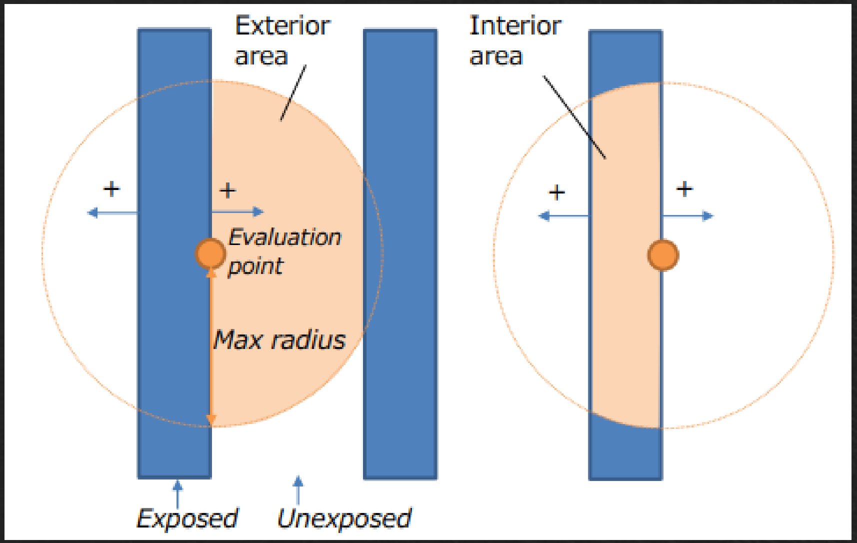

Sum of All Corrections Shows 44% EPE Improvement

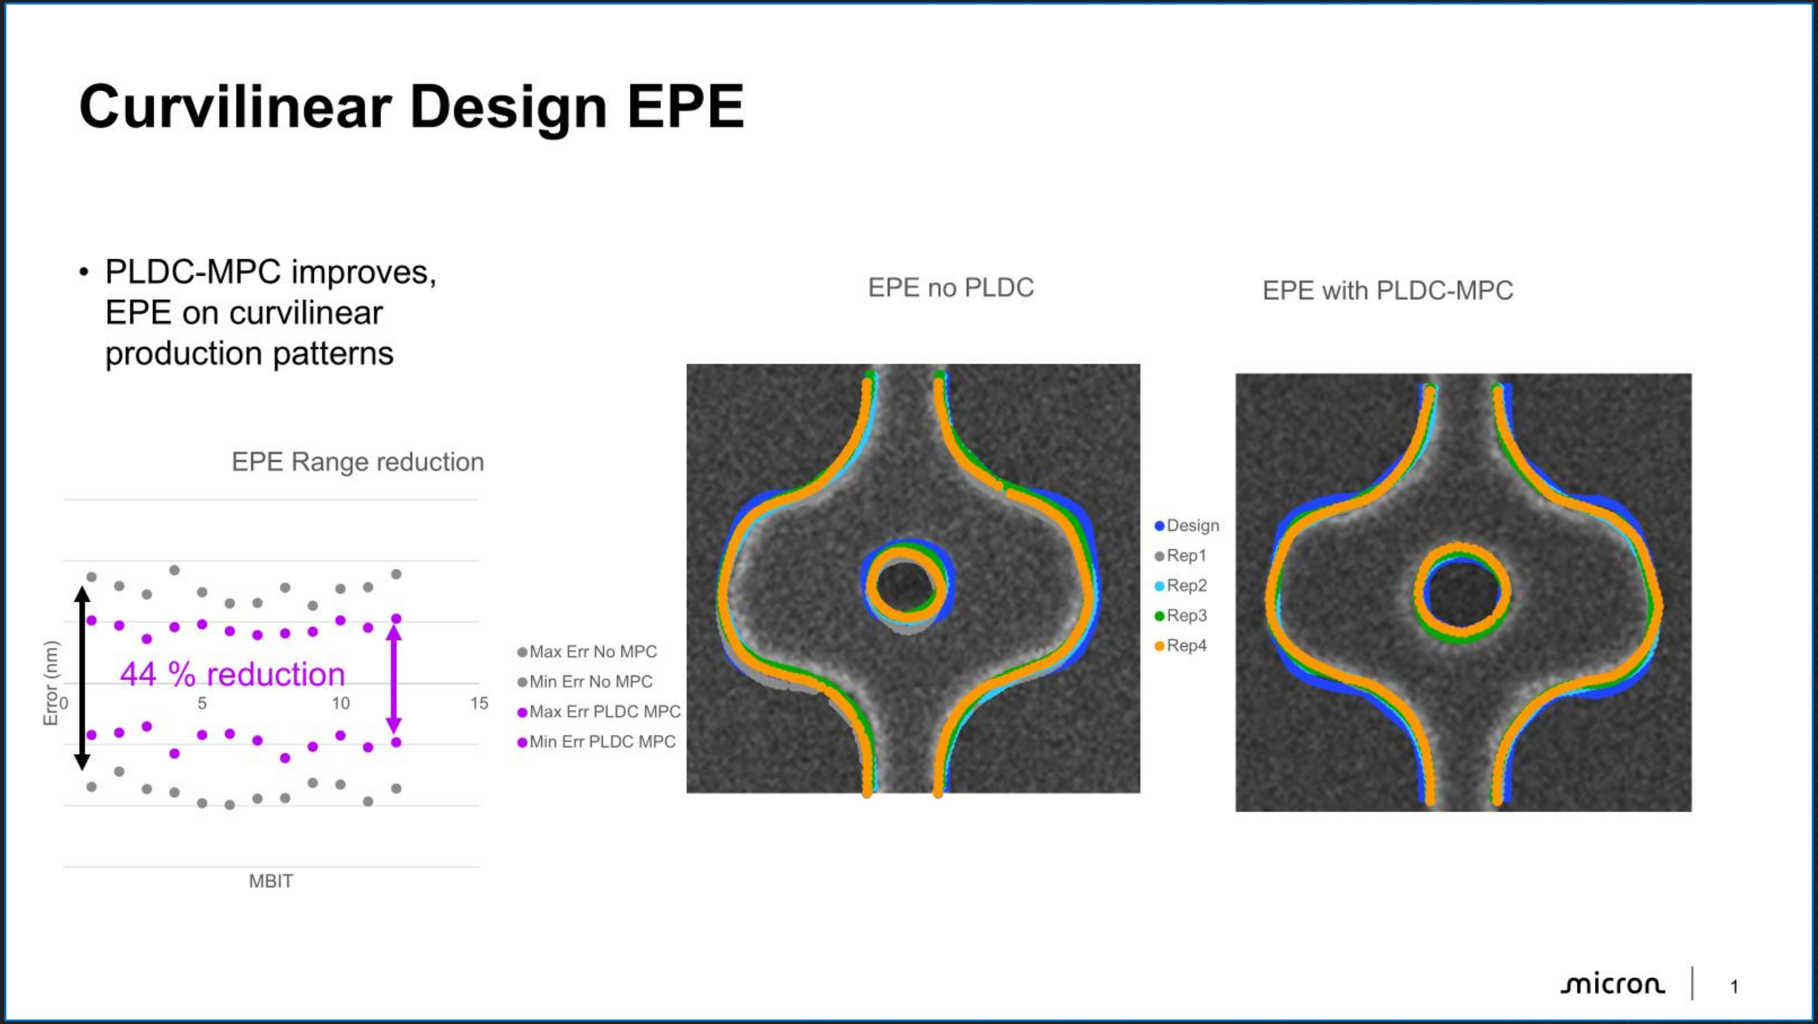

Curvilinear Design EPE

- PLDC-MPC improves, EPE on curvilinear production patterns

But How Do You Do 193i Curvilinear Without Multi-Beam?

Mask Wafer Co-Optimization (MWCO)

Optimizing for Wafer Reduces Shot Count

MWCO moves VSB shot edges to get 2X process windows

More on ILT and MWCO from Leo Pang at 13:40 today

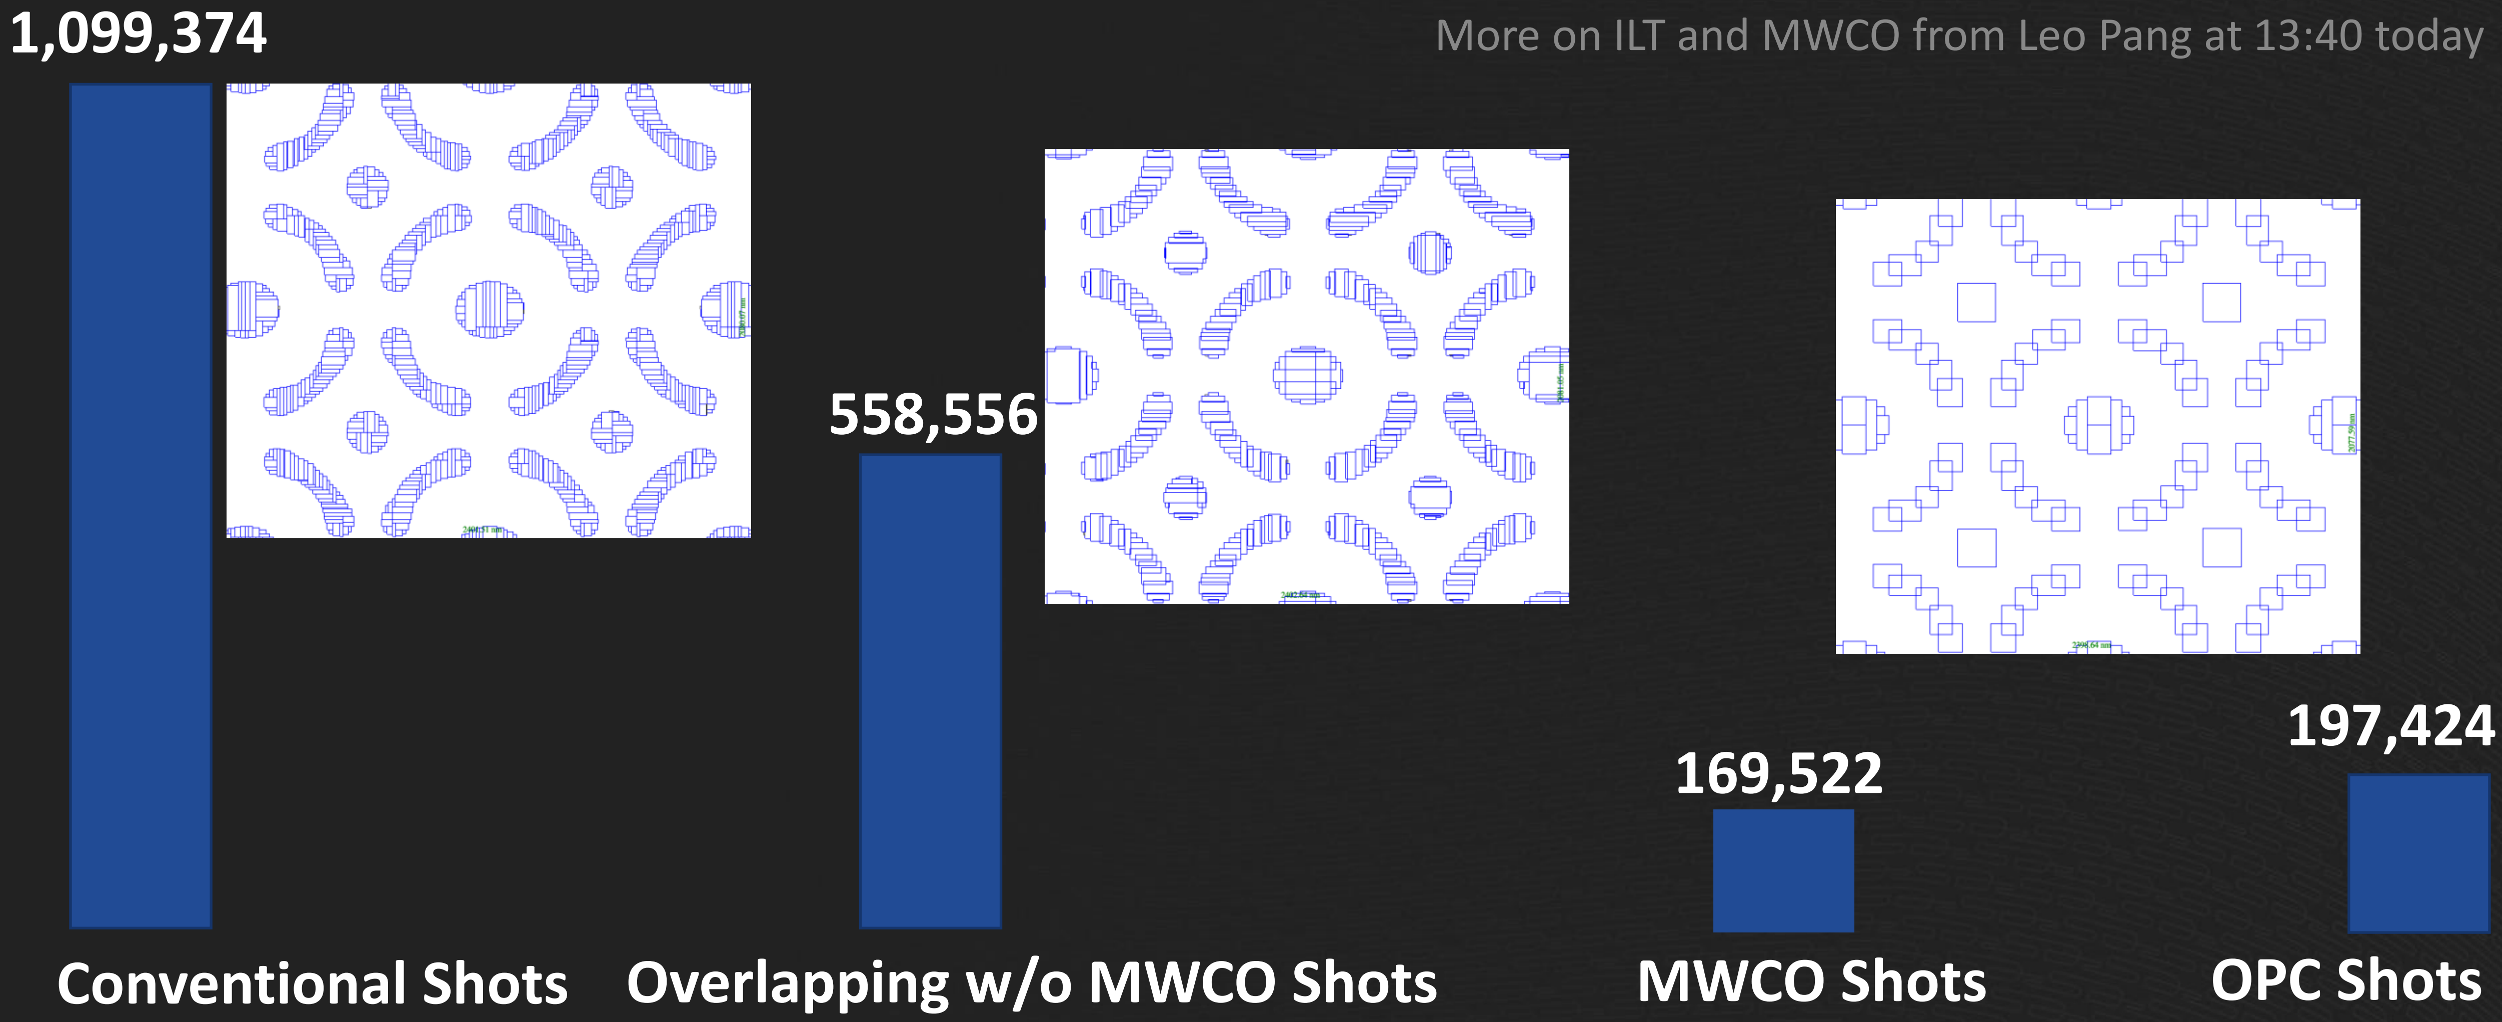

D2S PATENTED TECHNOLOGY

Ask for What You Can Get

And Get What You Asked For

This should guide what we mean by "curvilinear"

Now that Curvy Mask Manufacturing is Possible
Curvy Design is Enabled

Same pixel-based GPU technology built for manufacturing is what's needed in EDA to enable Curvy Design

Conventional

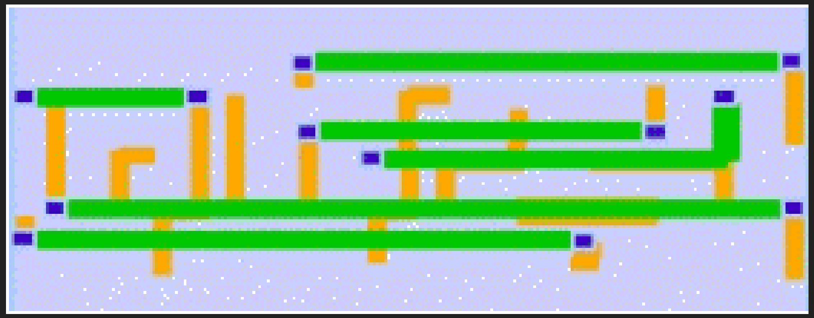

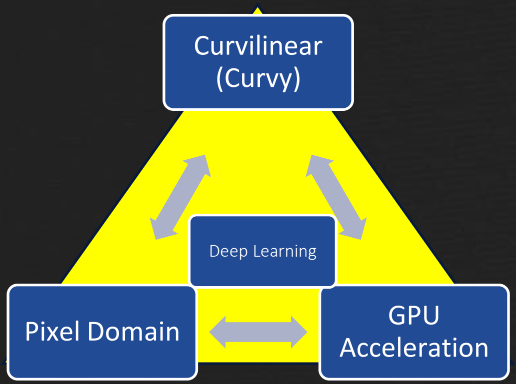

Curvilinear (Curvy)

Deep Learning

Pixel Domain

GPU Acceleration

Curvy

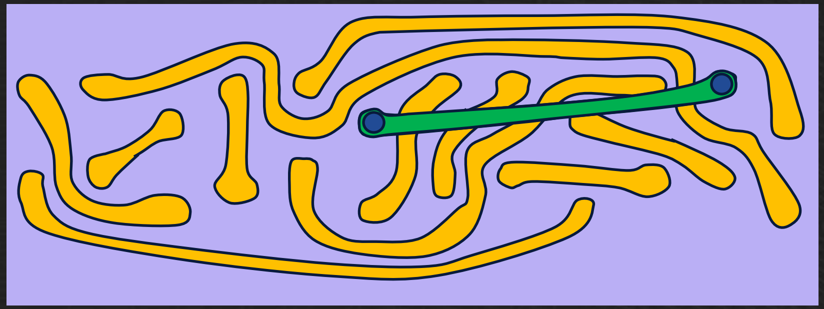

This is Possible with VSB Today

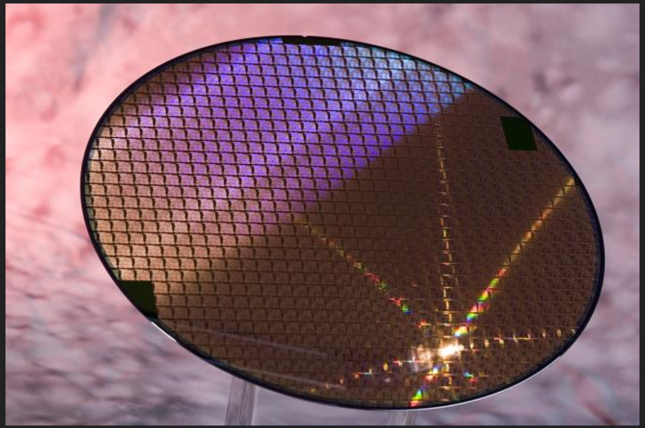

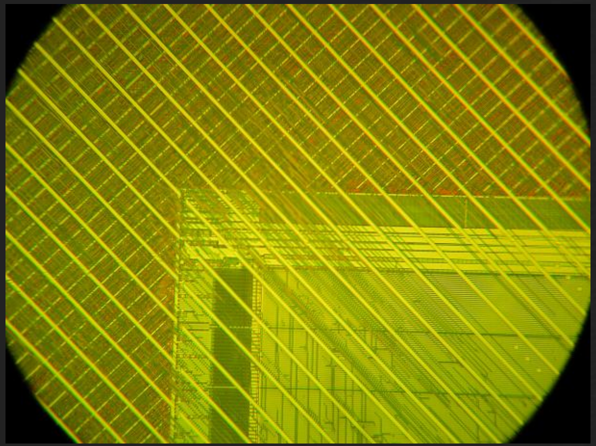

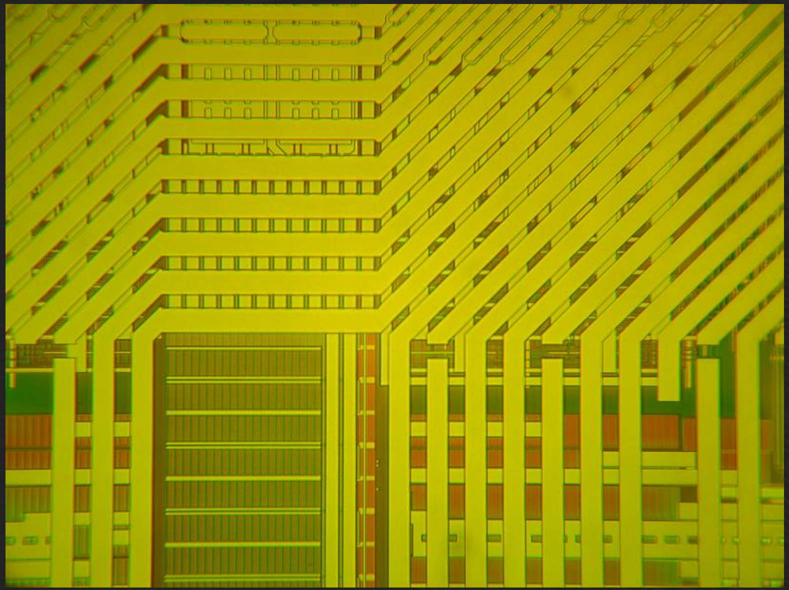

30% Wire Length Reduction from X By Many Customers

Curvy design is a superset of X-architecture

ATi

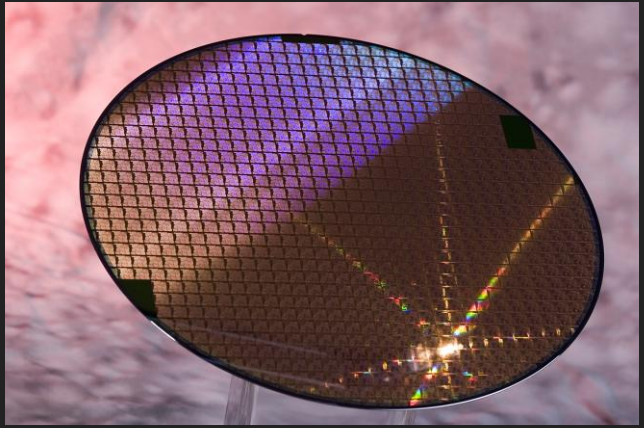

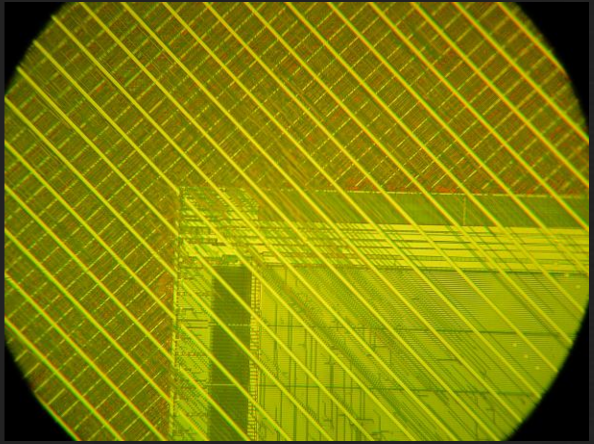

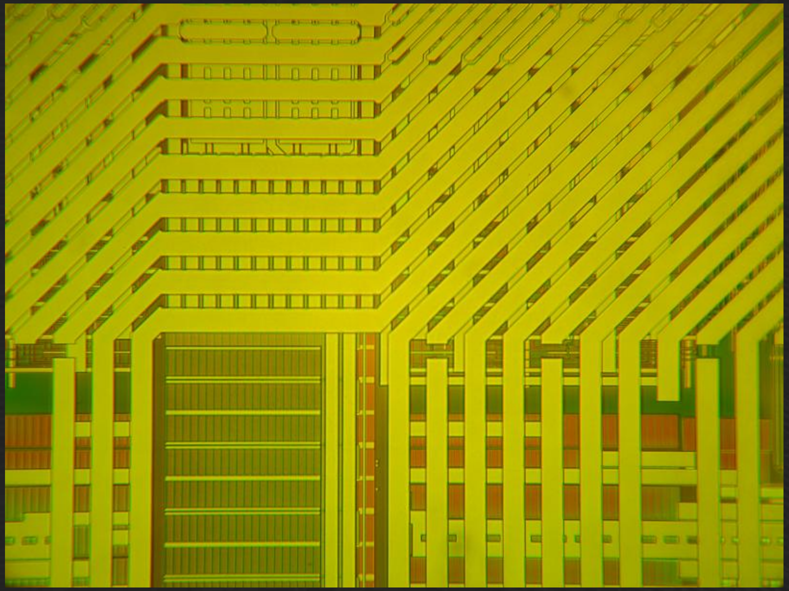

20% Faster Processing Time + 10% Area Reduction

斜め配線が可能な設計手法「Ｘ　Ａｒｃｈｉｔｅｃｔｕｒｅ」による初のＬＳＩ設計について

2002年2月6日

従来手法の設計に比べ処

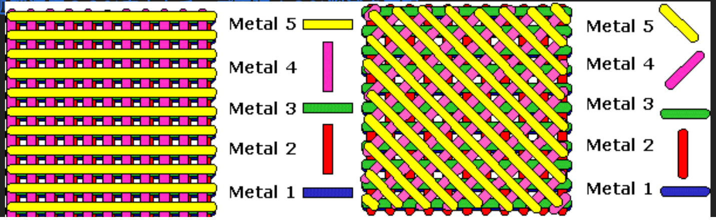

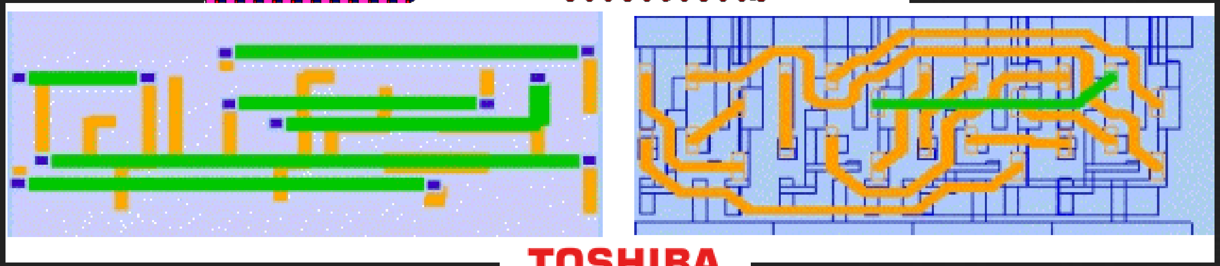

TOSHIBA

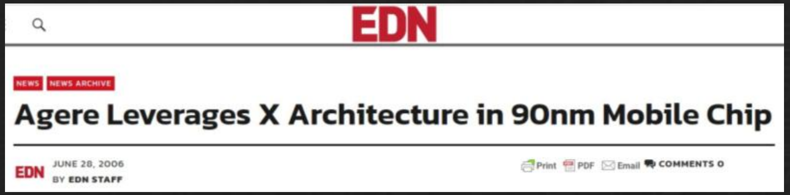

EDN

NEWS NEWS ARCHIVE

Agere Leverages X Architecture in 90nm Mobile Chip

EDN JUNE 28, 2006
BY EDN STAFF

Print PDF Email COMMENTS 0

X Facts from 16 production chips

- -30% wirelength
- -40% vias
- +20% performance
- -10% area
- Reduce 1 layer pair
- Faster design cycles

D2S

Ryoung-han (Ryan) Kim

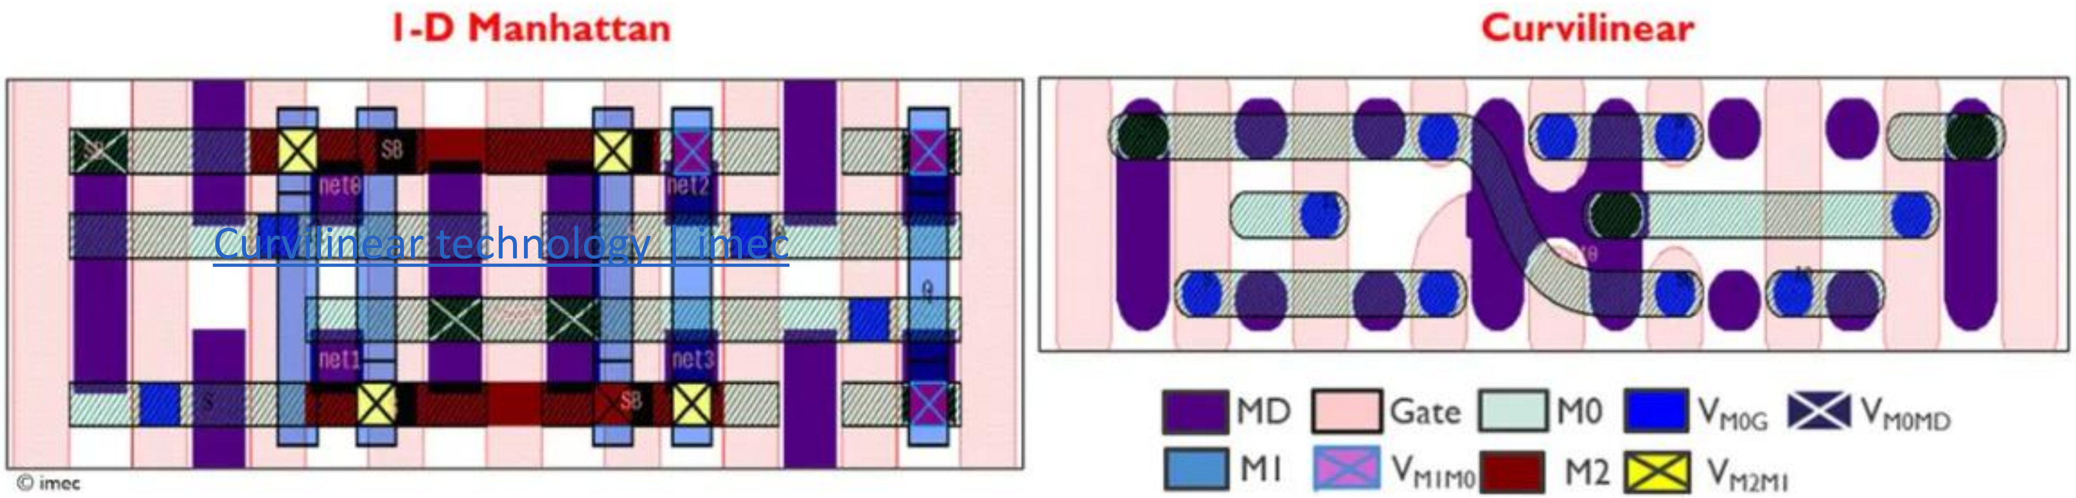

Curvilinear technology | imec

Figure 5 – Source/drain contact and gate routing using (left) 1-D Manhattan and (right) curvilinear designs. Curvilinear design completes the design within M0 while Manhattan design requires M2.

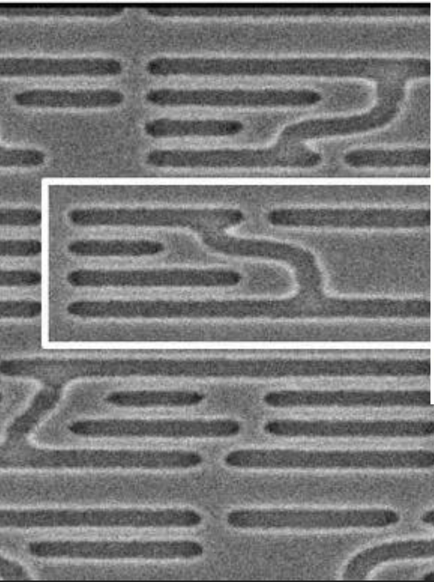

www.imec-int.com/en/articles/curvilinear-technology-game-changer-logic-technology-roadmap

Steve Teig, Amazon’s VP and Distinguished Engineer, designer of Perceive low-power AI chip, says:

“Curvy is a transformative P&R technology”

- Simultaneously achieve the following benefits:

Goal	Performance	Total Power		Area	Wirelength	Vias	Perf / W
Balance	+25%	-20-25%		-10%	-30-45%	-40-50%	+50-60%
	+30-35% for chiplets	-15% dynamic	-25-30% leakage	-20-25% for chiplets			

- Individually achieve the following benefits:

Goal	Performance	Total Power	Area	Perf / W
Max Performance	+40-43%	+10-15%	-5-7%	+25-30%
Max Power Savings	0%	-40-45%	-8-10%	+40-45%
Max Area Reduction	+15-18%	-15-20%	-10-14%	+35-40%
Max Efficiency	+20-25%	-25-30%	-8-10%	+65-75%

- Improved yield and reliability: 40% fewer vias, lower current density, reduced electromigration
- Improved timing closure

Mask or Wafer, Design What Can be Manufactured so you can Manufacture it Exactly and get Less Variation

D2S