




How to Use Commercial 2D IC EDA Tools to Build Commercial Quality Monolithic 3D IC Designs



Prof. Sung Kyu Lim

GTCAD Laboratory (www.gtcad.gatech.edu)

Georgia Institute of Technology

D43D Workshop, 6/26/2017, Grenoble

EDA, Your Turn Please

2/36

People say
Coolcube/3DVLSI/
M3D is cool!

My boss wants big
designs and PPA,
commercial quality.

Can you build
tools? You will sell
lots of copies.



Nope. I have not
received any
order yet.

We will build one
for you if you
pay.

Need Some Help

3/36

Fine, I will do it
myself. But...
how?

I do **not** want to (or
cannot) start from
scratch.

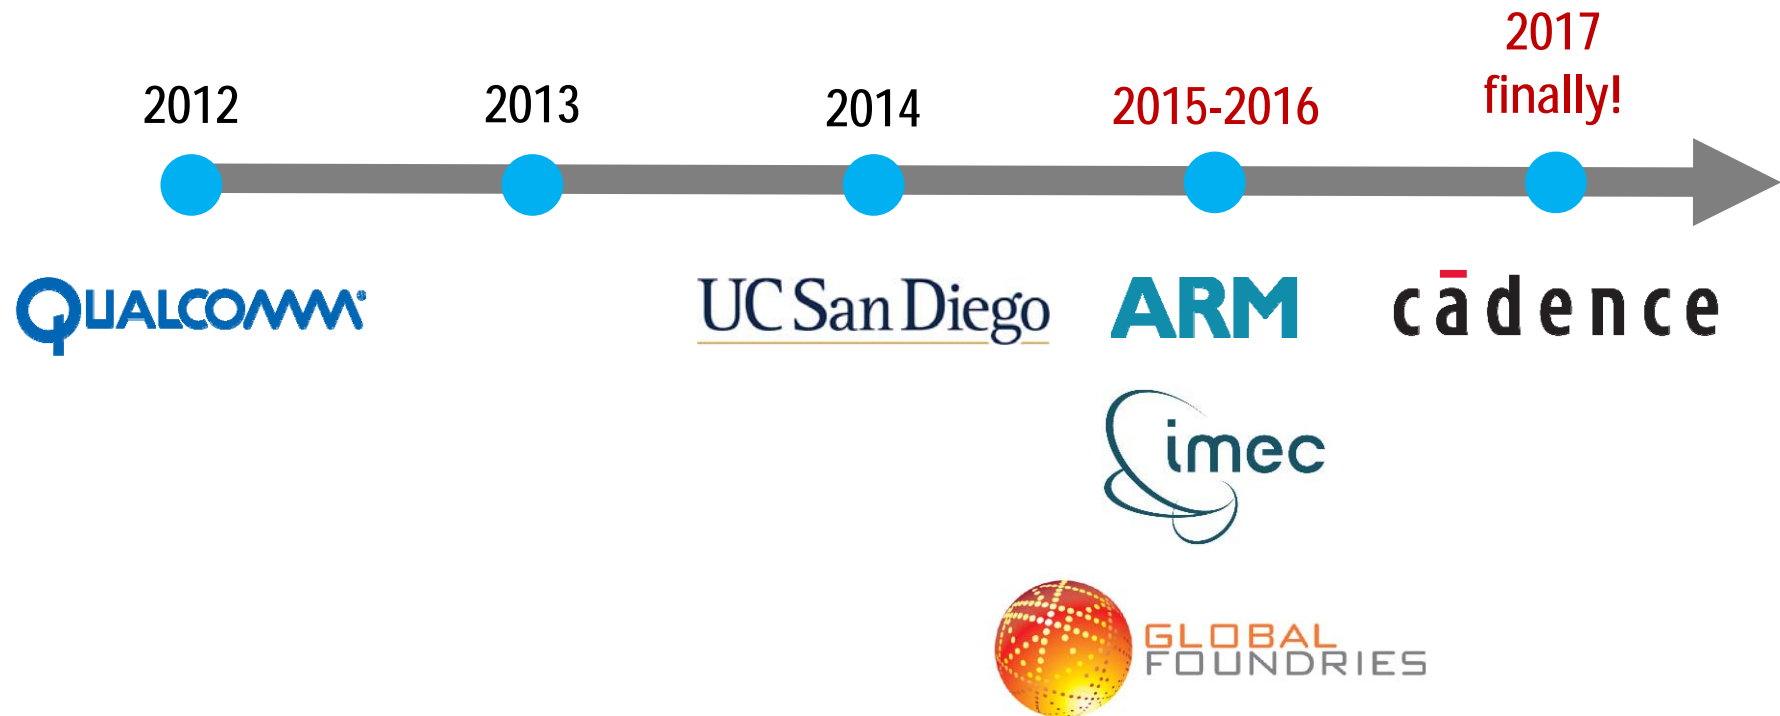
Can I recycle
commercial 2D IC
tools somehow?



A Brief History

4/36

- We went ahead ourselves (with industry partners)




- Yes, we have the tool(s) now!

We Published, Too

5/36




- Will cover the first 3 today

name	contribution	Industry collaborator	publications
Shrunk-2D	pioneer	Qualcomm	ISLPED 2014 TCAD 2016
Cascade-2D	handles arch constraints	ARM	DAC 2016 ICCAD 2016
Derate-2D	avoids shrinking	IMEC	ISLPED 2016 ICCAD 2016
TA-2D	handles inter-tier mismatch	GF	ISLPED 2016 ICCAD 2016



Shrunk-2D


The One That Started It All



Why Shrinking?

With Qualcomm

7/36



2D IC
cells fit nice




3D IC
cells overlap

Solution? Shrinking!

With Qualcomm

8/36

- Shrunk-2D flow [ISLPED'14]
 - Shrink the chip footprint
 - Shrink cell/wire dimensions (and RC) by 50%
 - Perform timing-closed 2D IC P&R as usual: **no overlap occurs!**
 - Repopulate cell/wire, tier-partition
 - Detailed routing die-by-die

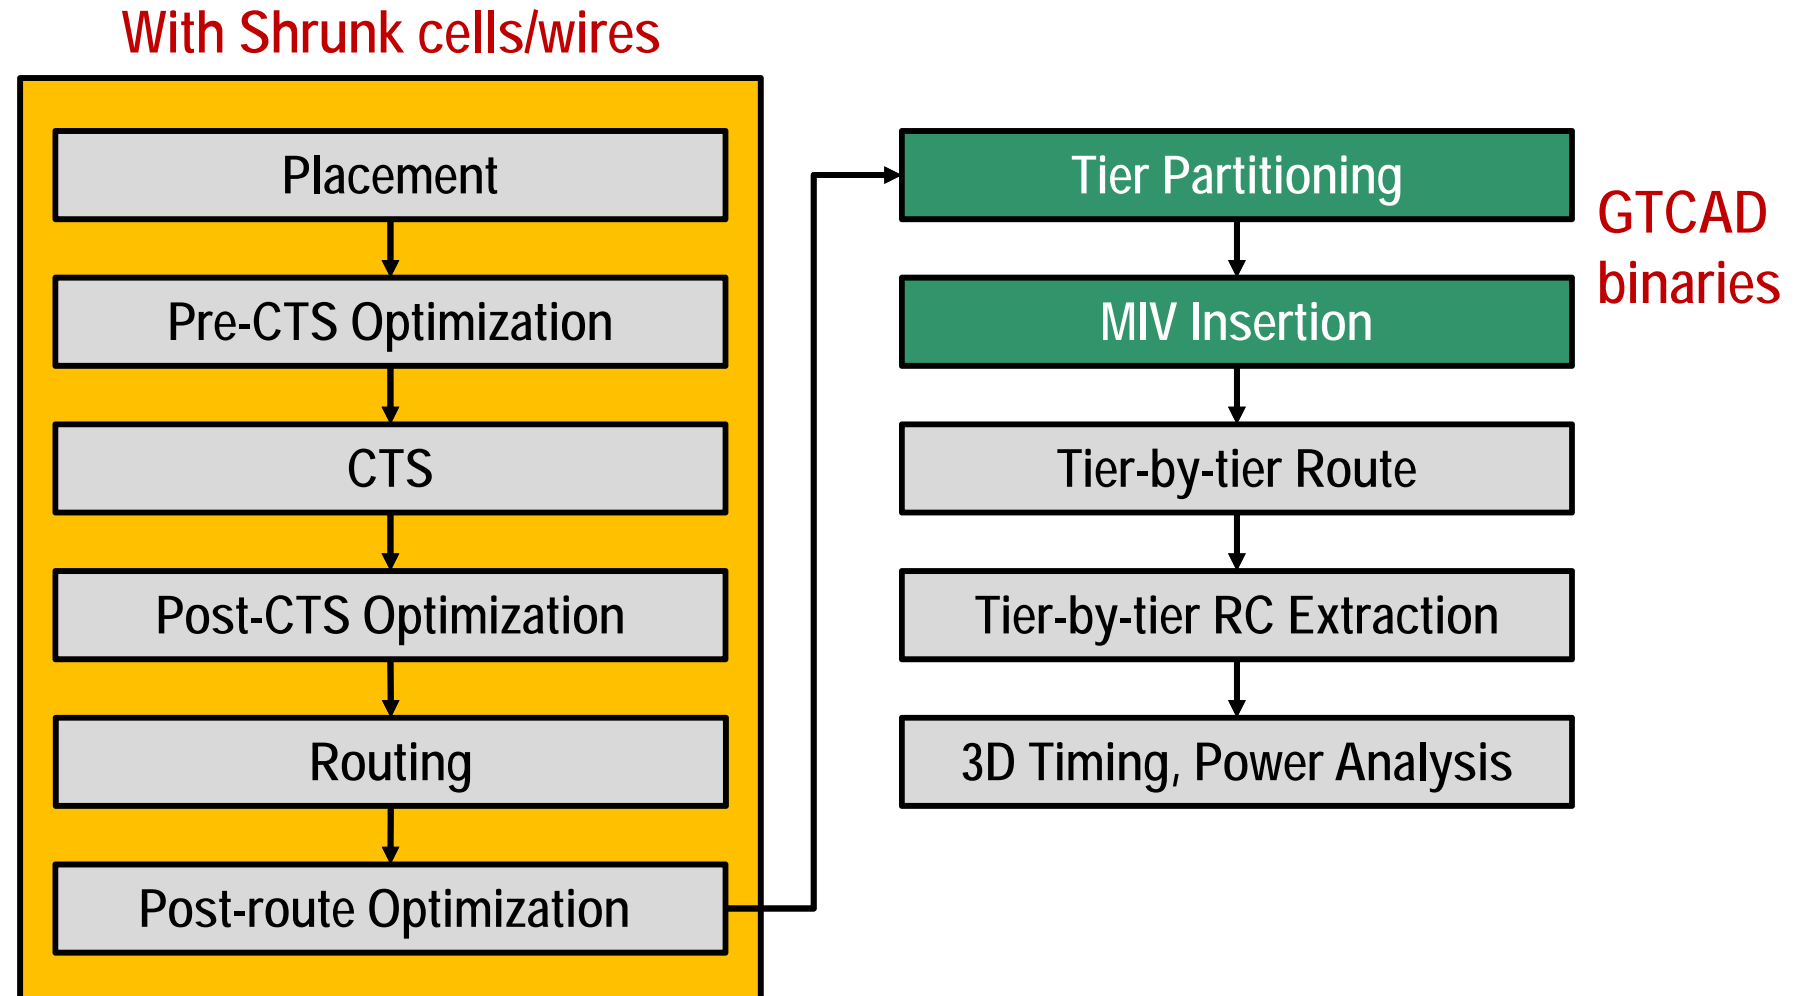


Design Flow

With Qualcomm

9/36



- Cell/wire RCs are also shrunk appropriately




Handling Memory Macros

With Qualcomm

10/36




1. Pre-Placed Memory




2. Memory Projection

reduced placement density over partial blockages



3. Shrunk 2D P&R




4. Tier Partitioning

MIV Placement




With Qualcomm

11/36

Top cells vs. Bottom cells
Differ by cell structure



Route 3D nets with Encounter




Leaf-level Clock MIV Insertion

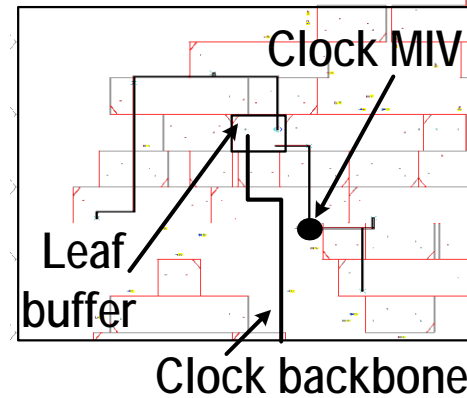
With Qualcomm

12/36


Clock back-bone on tier 0



nanoroute last-level buffer and its FFs



Leaf clock net on tier 0



Leaf clock net on tier 1


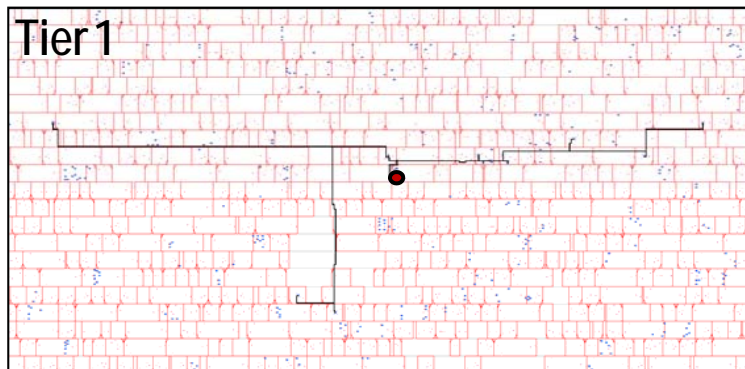
	1 MIV	Multiple MIVs
Clock WL (m)	1.03	0.80 (-21.67%)
Clock Power (mW)	68.40	48.00 (-29.82%)

Single vs Multiple MIV/F2F Insertion

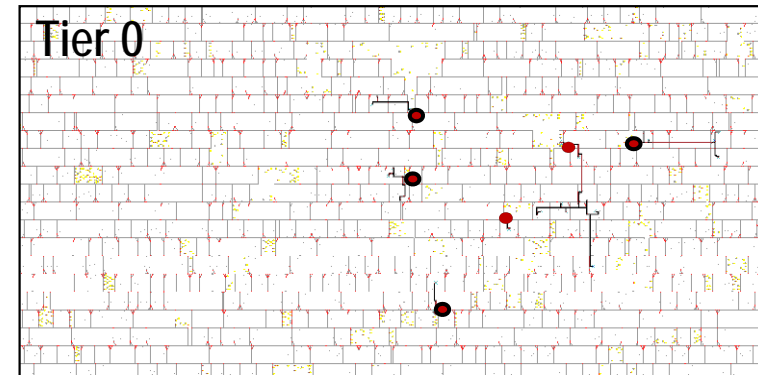

With Qualcomm

13/36

Single MIV per net



Multiple MIVs per net



	1 MIV	Multiple MIVs
#MIV	106k	235k (+120.44%)
Total WL (m)	15.61	14.29 (-8.43%)





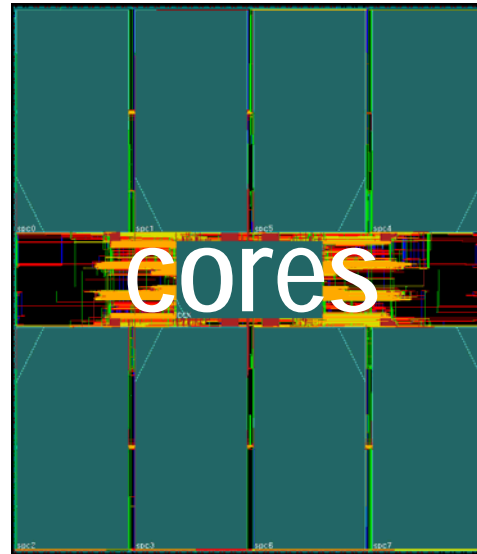
Commercial-Grade 8-Core Designs

With Qualcomm

14/36




870MHz
ST28nm (FDSOI)




MIV Maps

15/36



Logic + Memory
#MIV = 4,205




Folded
#MIV = 838,360

Detailed Comparisons

16/36





- PDK: ST28nm FDSOI

T2	2D	3D core/cache	diff	3D folding	diff
Footprint (mm ²)	15.6	7.8	-50%	7.8	-50%
Si area	15.6	15.6	0%	15.6	0%
WL (m)	99.4	95.2	-4.2%	76.58	-23.0%
# Cells	2.62M	2.58M	-1.28%	2.47M	-5.41%
# Buffers	0.53M	0.50M	-5.99%	0.45M	-16.02%
# HVT cells	83.34%	85.94%		88.63%	
Total power (W)	5.70	5.61	-1.5%	5.03	-11.8%
Cell (W)	2.94	2.89	-1.7%	2.76	-6.1%
Net (W)	2.74	2.70	-1.5%	2.26	-17.5%
Leakage (W)	0.016	0.014	-12.5%	0.010	-37.5%



Cascade-2D

Architects Called For It



Shrinking Causes Issues

With ARM

18/36

I want block A
and B to be on
top of each other
in my M3D
design.

Can you handle
that?



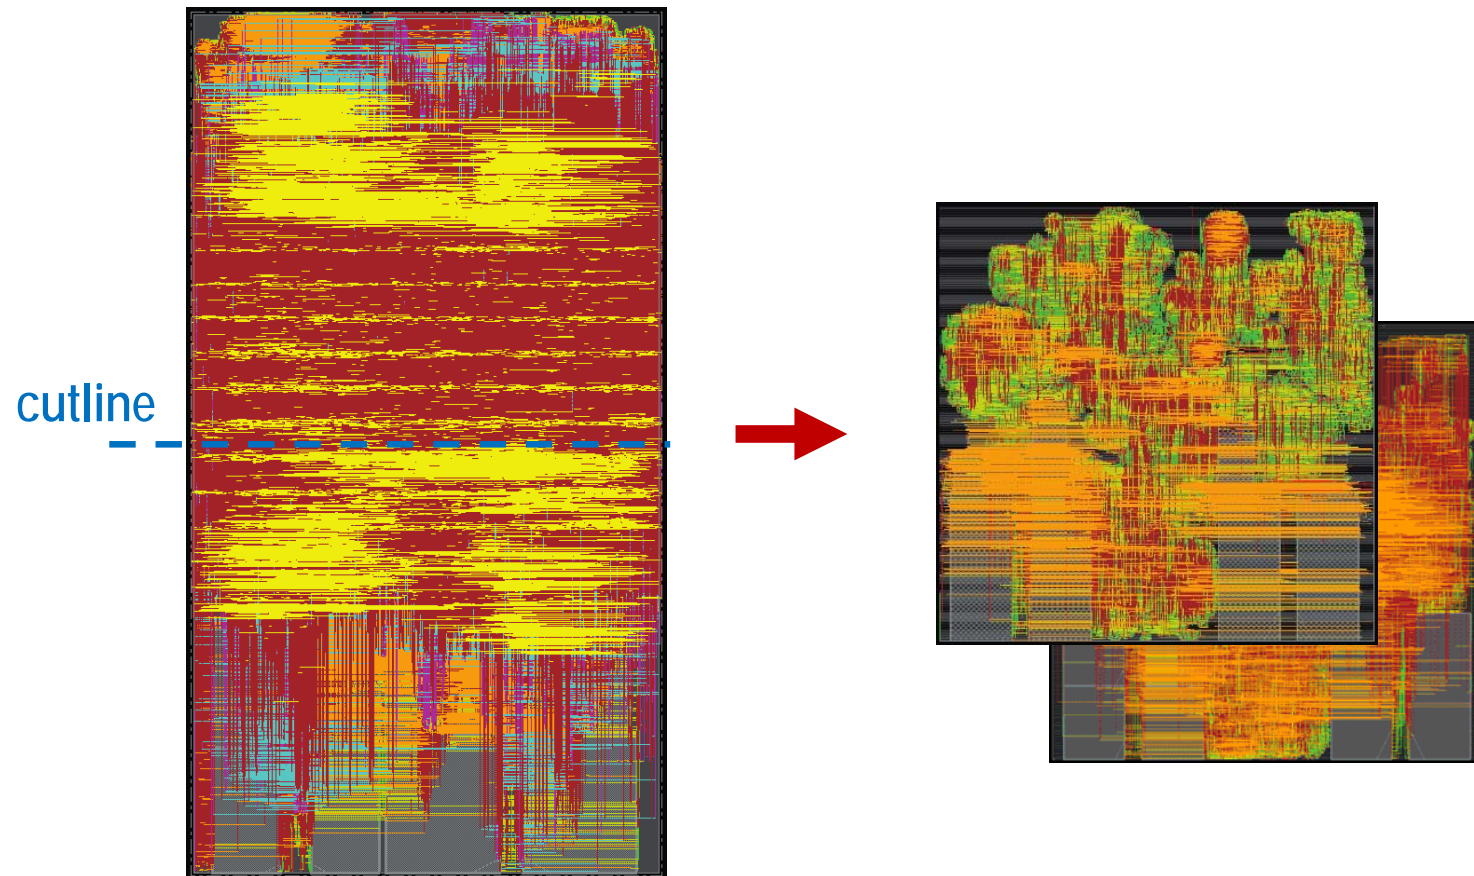
Not now. Gimme
some time.

Cascade-2D

With ARM

19/36

- Cut-and-slide [ICCAD'16]
 - Still uses 2D IC P&R tool



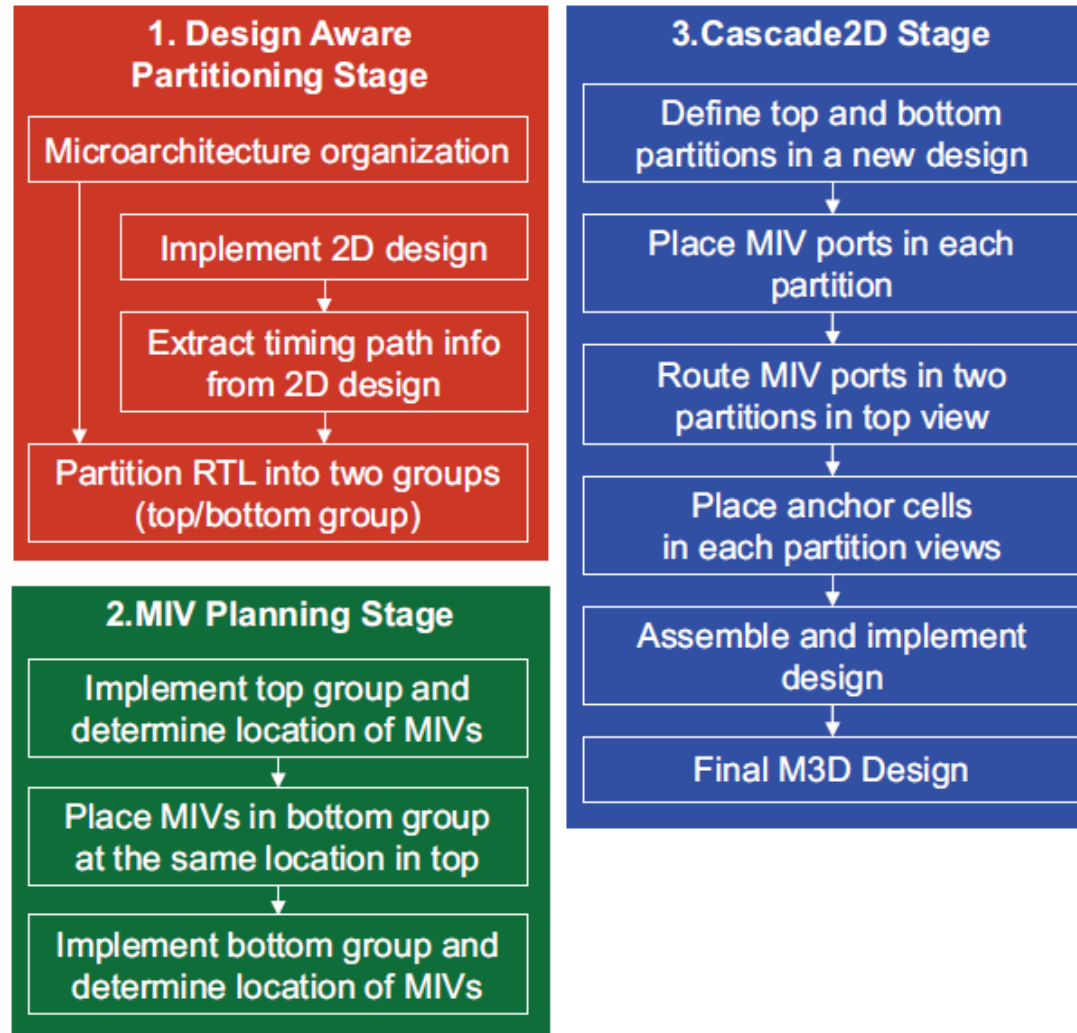
Key Issue: Handling 3D Connections

With ARM

20/36

1. Tier partitioning first

2. MIVs placed next

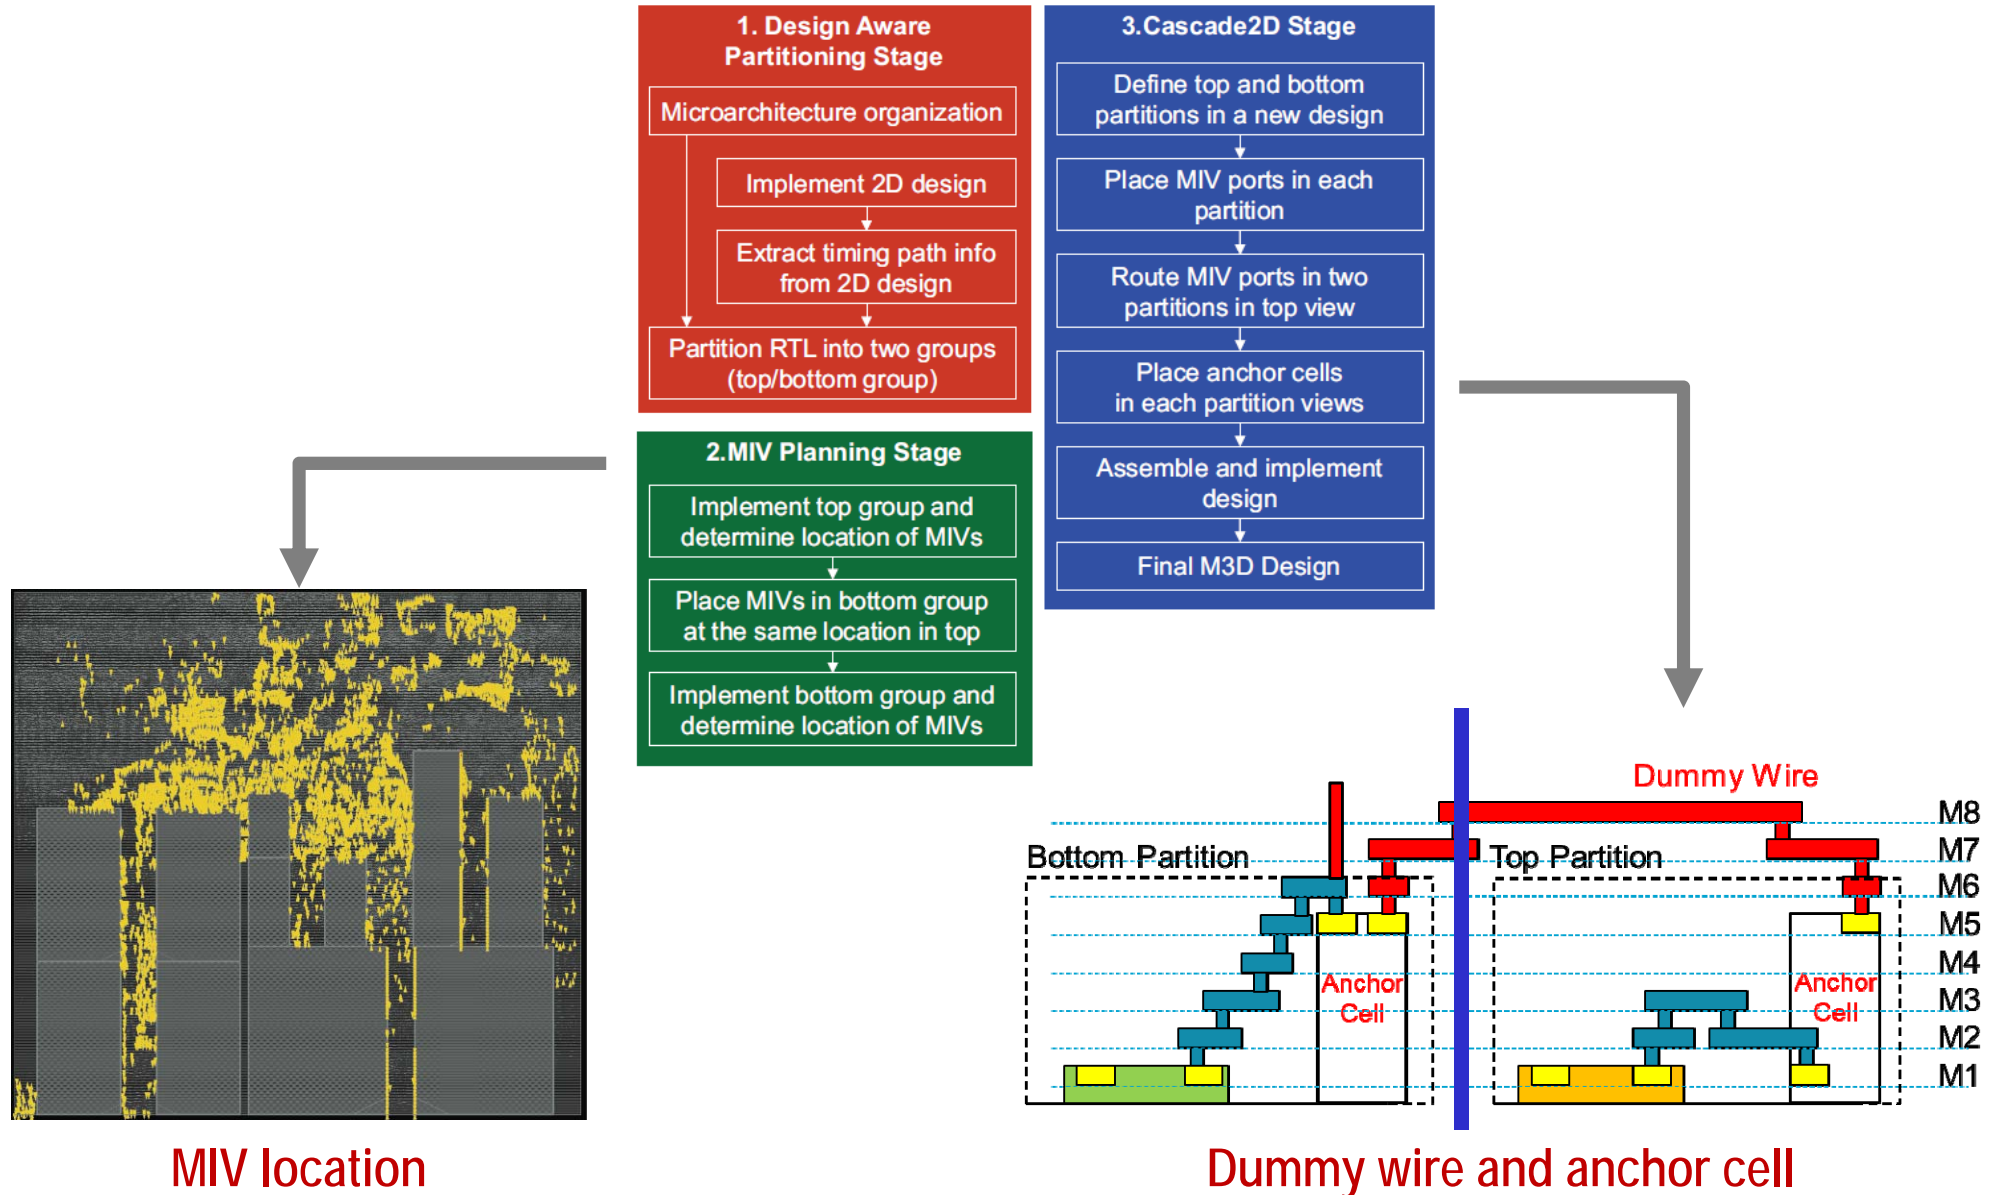


3. Full-chip P&R

Key Issue: Handling 3D Connections

With ARM


21/36



Details

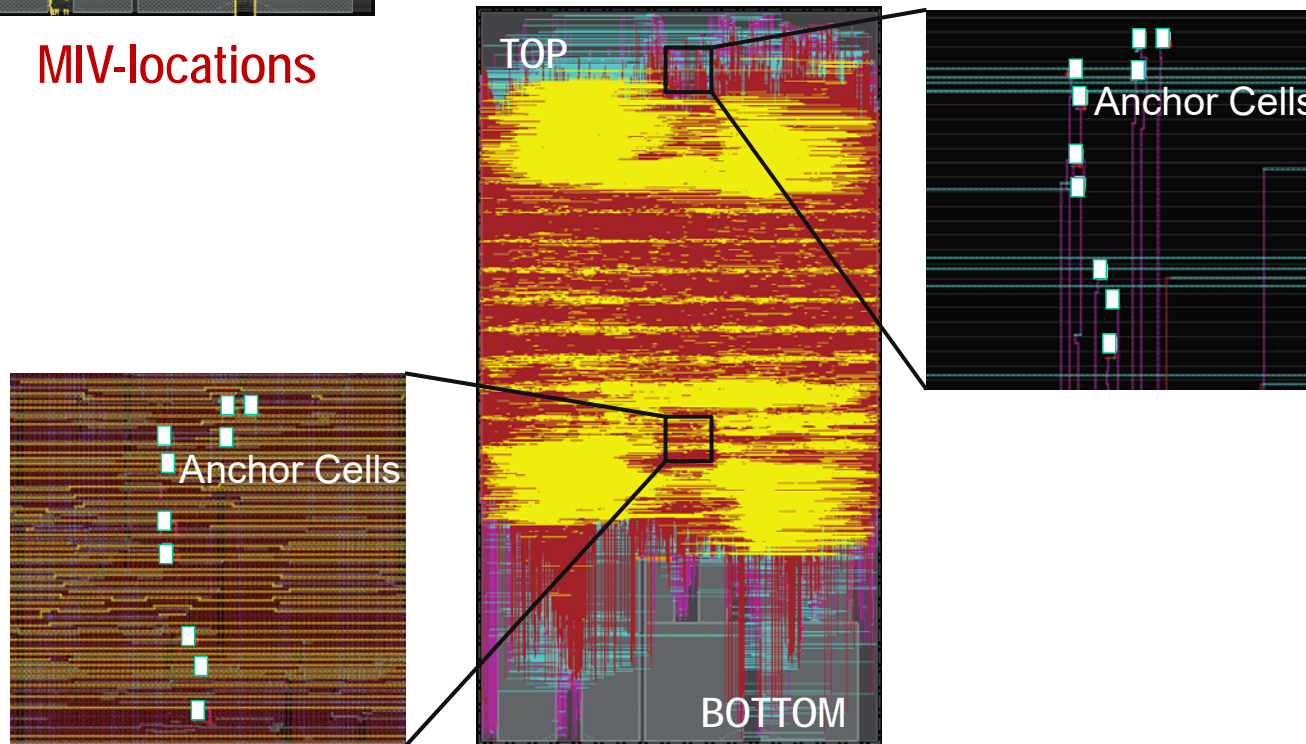
With ARM

22/36




MIV-locations

Dummy wires and anchor cells



Cascade-2D Design




Floorplanning Constraint Works


With ARM

23/36

- Handled during stage 3 of C2D
 - Works without fences



2D IC design



M3D design (A should be on top of B)


Which Node Is Best for M3D?

With ARM


24/36

- Performed frequency sweeps across three technology nodes
 - Design: commercial in-order 32-bit AP
 - Technology: foundry 28nm, 14/16nm, and predictive 7nm


	28nm	14/16nm	7nm
Transistor type	Planar	FinFET	FinFET
Supply Voltage	0.9V	0.8V	0.7V
Contact Poly Pitch	110-120nm	78-90nm	50nm
M1 Pitch	90nm	64nm	36nm




28nm 2D




28nm M3D




14/16nm 2D



14/16nm M3D



7nm 2D




7nm M3D

Cascade-2D Results



With ARM

25/36


- Ourperforms S2D



Power saving over 2D






Cell area saving over 2D



Derated-2D

Shrinking Not Necessary



Shrinking Causes Issues

With IMEC

27/36

I tried your S2D
on my 10nm
designs. It asks
for 7nm license.

I did, but now I
get tons of DRC
errors!



Good. You gotta
pay me.


Good. You gotta
buy my 7nm
cells.

Solution? Placement Projection!


With IMEC

28/36


- Project 2D placement onto 3D IC footprint



Derated-2D
439um x 437um
(cells and interconnects are
not shrunk)



Derated-2D
Placement
(cells and interconnects are
not shrunk)




Placement projection
= $0.7 \times x/y\text{-coordinates}$
(will have lots of overlap)
310um x 309um

Tier Partitioning

With IMEC

29/36

- Bin-based FM mincut partitioning




Tier partitioning result

Top / Bottom


Partitioning bin size = 8.8um x 8.8um

Local area skew < 5%

310um x 309um



Top die
310um x 309um
Cell = 51,548
Cell area = 63,433 um²

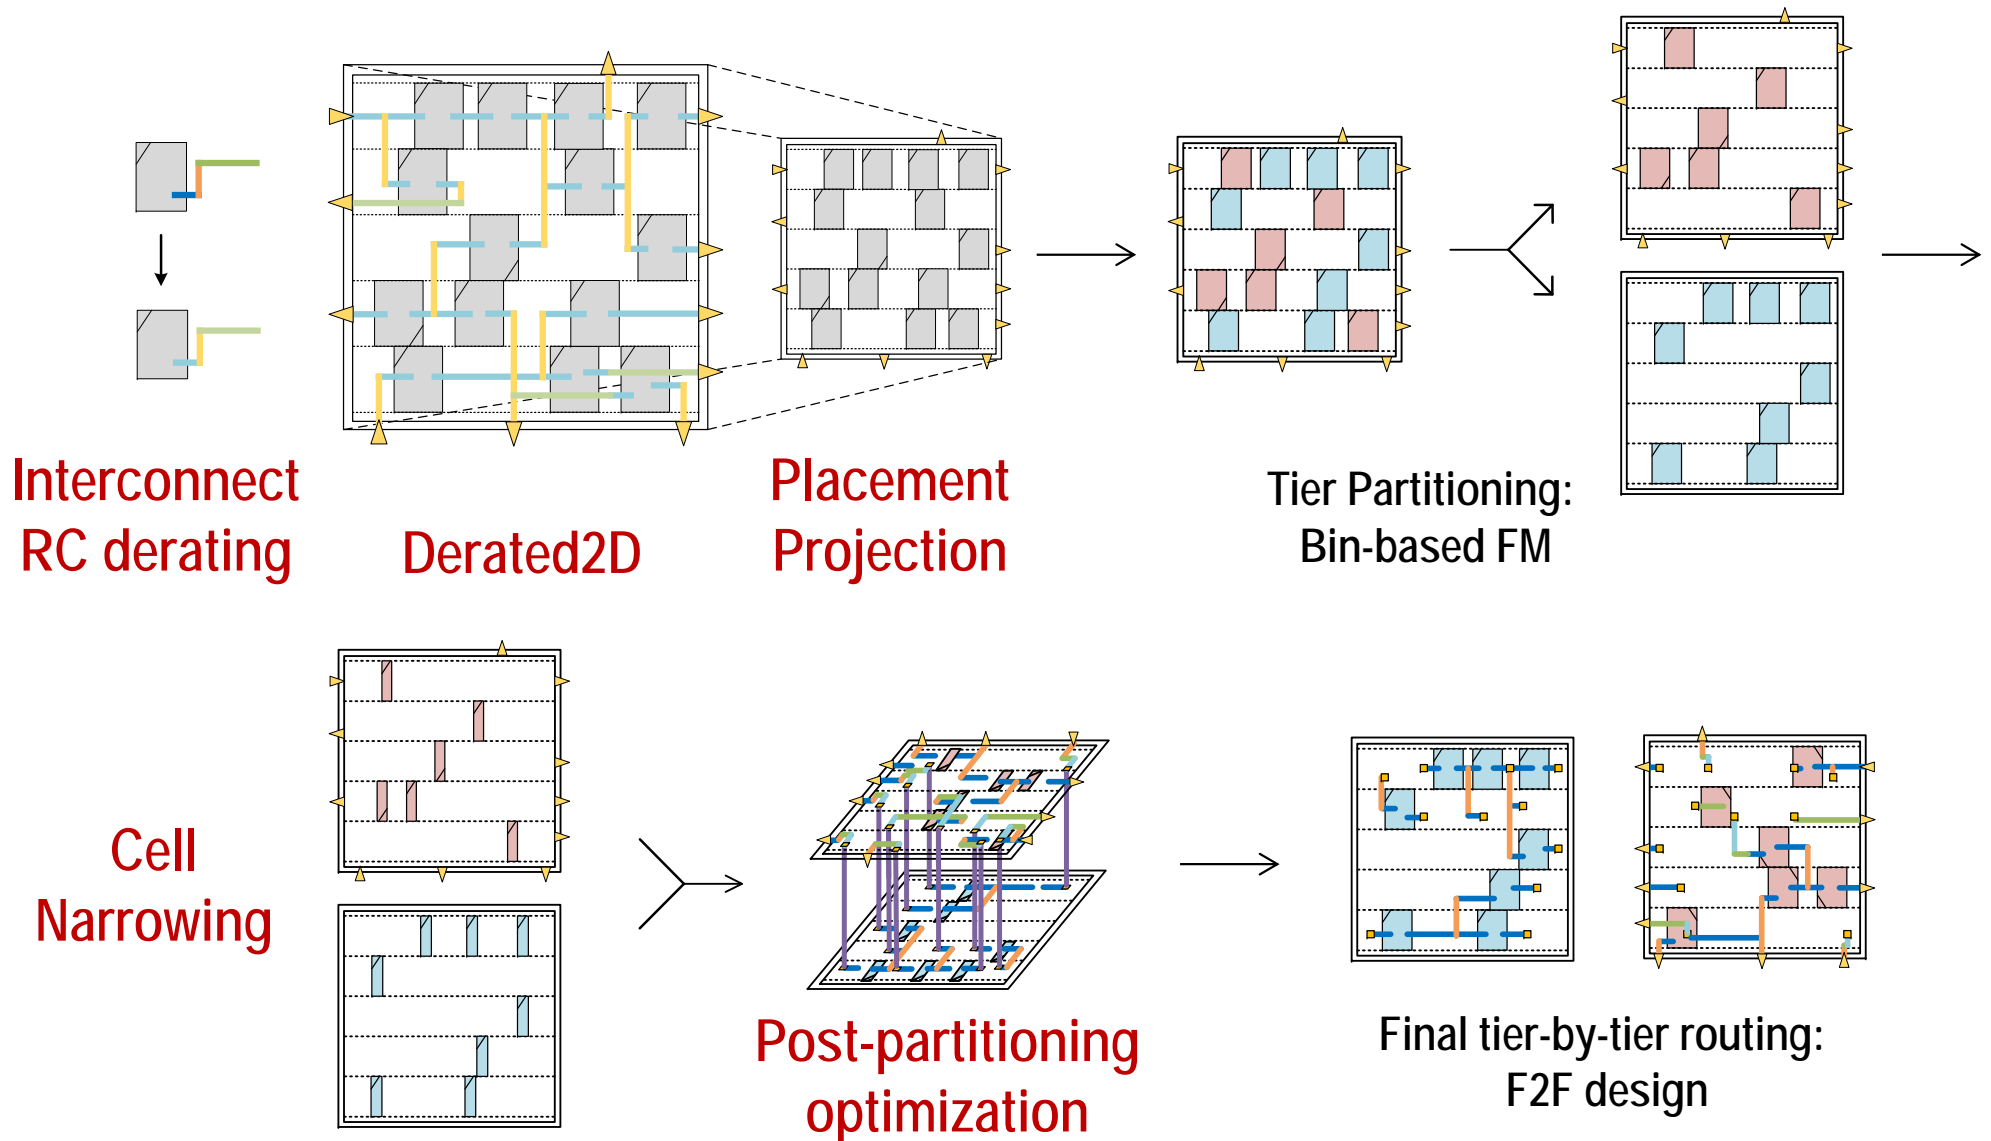


Bottom die
310um x 309um
Cell = 68,762
Cell area = 66,344 um²

Overall Design Flow

With IMEC

30/36




Post-partitioning Optimization


With IMEC

31/36

- We still need to use a 2D IC optimizer



F2F stack-up view (M3D similar)



Overlapped Top/Bottom cell placement




Optimization engine will legalize the overlap:
placement is **DAMAGED!**

Enabling Post-partitioning Optimization

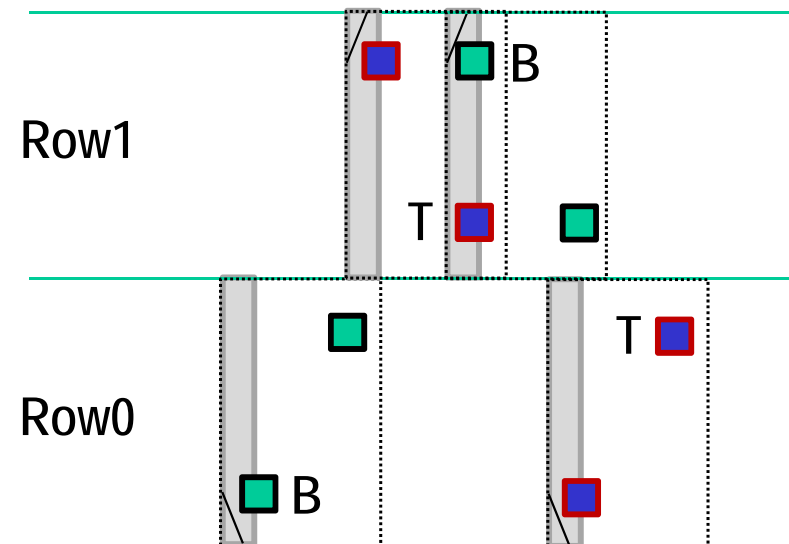
With IMEC

32/36

- Idea: cell narrowing (site-sized MACRO LEF)
 - To **temporarily** remove overlap just to do timing closure



Pins are fine: no overlap
Cells are not fine: overlap




Pins are **not** overlapping
Cells are **not** overlapping
Optimization works
And placement is not damaged


Details

With IMEC

33/36



cell narrowing




after optimization

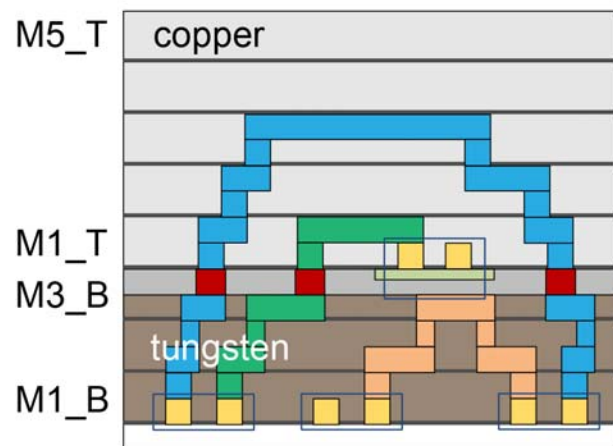
Fighting Bottom-Tier Degradation

With IMEC

34/36



(a) Shrunk 2D



(b) Derated 2D

Idea: Use top tier metals (= faster) for routing bottom gates

metric	S2D	D2D
Top placed, top routed	17,432	17,410
Top placed, both routed	0	22
Bot placed, bot routed	22,280	19,072
Bot placed, both routed	0	3,208
Both placed, both routed	19,984	19,984
Ave top tier WL (um/net)	5.40	6.85
Ave bot tier WL (um/net)	3.50	2.64
Fmax (GHz)	0.68	0.75

LDPC designed with IMEC N7

AES (Pin-cap Dominated)

With IMEC

35/36

- Used ST28nm FDSOI

	2D	S2D	S2D – F2F	D2D	D2D – F2F No opt	D2D – F2F Post-Part opt
Footprint (um ²)	251001	120408 (-52%)	120408 (-52%)	251001	120408 (-52%)	120408 (-52%)
WL (m)	2.021	1.485 (-27%)	1.676 (-17%)	1.979 (-2%)	1.581 (-22%)	1.596 (-21%)
F2F via#	-	-	50947	-	43413	75837
Cell#	123214	122418 (-1%)	122418 (-1%)	121143 (-2%)	121143 (-2%)	121373 (-1%)
Buffer#	22134	21414 (-3%)	21414 (-3%)	19785 (-11%)	19785 (-11%)	20015 (-10%)
Ave Buf cap (fF)	3.24	3.22 (-1%)	3.22 (-1%)	3.10 (-4%)	3.10 (-4%)	3.05 (-10%)
WNS (ns)	-0.012	-0.017	-0.048	-0.011	-0.008	0.006
TNS (ns)	-0.338	-0.607	-11.399	-0.173	-0.018	0.000
Wire cap (pF)	240.5	199.8 (-17%)	203.8 (-15%)	162.6 (-32%)	190.3 (-21%)	197.1 (-18%)
Pin cap (pF)	472.8	457.4 (-3%)	457.4 (-3%)	435.9 (-8%)	435.9 (-8%)	405.2 (-14%)
Switching (mW)	129.8	119.3 (-8%)	120.3 (-7%)	109.2 (-16%)	114.3 (-12%)	109.6 (-16%)
Internal (mW)	87.0	83.9 (-4%)	83.8 (-4%)	79.9 (-8%)	80.2 (-8%)	73.5 (-16%)
Total power (mW)	217.1	203.5 (-6%)	204.4 (-6%)	189.4 (-13%)	194.8 (-10%)	183.3 (-16%)

LDPC (Wire-cap Dominated)

With IMEC

36/36

- Used ST28nm FDSOI

	2D	S2D	S2D – F2F	D2D	D2D – F2F No opt	D2D – F2F Post-Part opt
Footprint (um ²)	92129	43688 (-53%)	43688 (-53%)	92129	43688 (-53%)	43688 (-53%)
WL (m)	1.661	1.124 (-32%)	1.199 (-28%)	1.618 (-3%)	1.206 (-27%)	1.248 (-25%)
F2F via#	-	-	19131	-	19068	30871
Cell#	46585	45571 (-2%)	45571 (-2%)	44802 (-4%)	44802 (-4%)	46955 (+1%)
Buffer#	12331	11639 (-6%)	11639 (-6%)	11191 (-9%)	11191 (-9%)	13344 (+8%)
Ave Buf cap (fF)	4.85	2.09 (-57%)	2.09 (-57%)	2.38 (-51%)	2.38 (-51%)	2.00 (-59%)
WNS (ns)	-0.0322	-0.0292	-0.0026	-0.0131	-0.0533	-0.0301
TNS (ns)	-5.9264	-0.6927	-0.0026	-0.2286	-6.4526	-0.2506
Wire cap (pF)	215.2	163.3 (-24%)	155.5 (-28%)	144.8 (-33%)	149.9 (-30%)	162.9 (-24%)
Pin cap (pF)	193.9	179.1 (-8%)	179.1 (-8%)	169.0 (-13%)	169.0 (-13%)	158.0 (-19%)
Switching (mW)	118.1	99.2 (-16%)	96.4 (-18%)	91.4 (-23%)	92.4 (-22%)	93.2 (-21%)
Internal (mW)	68.2	60.1 (-12%)	59.2 (-13%)	54.3 (-20%)	54.0 (-21%)	52.7 (-23%)
Total power (mW)	186.6	159.6 (-14%)	155.8 (-17%)	145.9 (-22%)	146.6 (-21%)	146.2 (-22%)