


# Addressing Thermal and Power Delivery Bottlenecks in 3D Circuits


Sachin S. Sapatnekar  
ECE Department  
University of Minnesota




# 3D technologies




Wire-bonded




3D package




Face-to-face




Capacitive coupling



Inductive coupling




Through-Silicon via - bulk



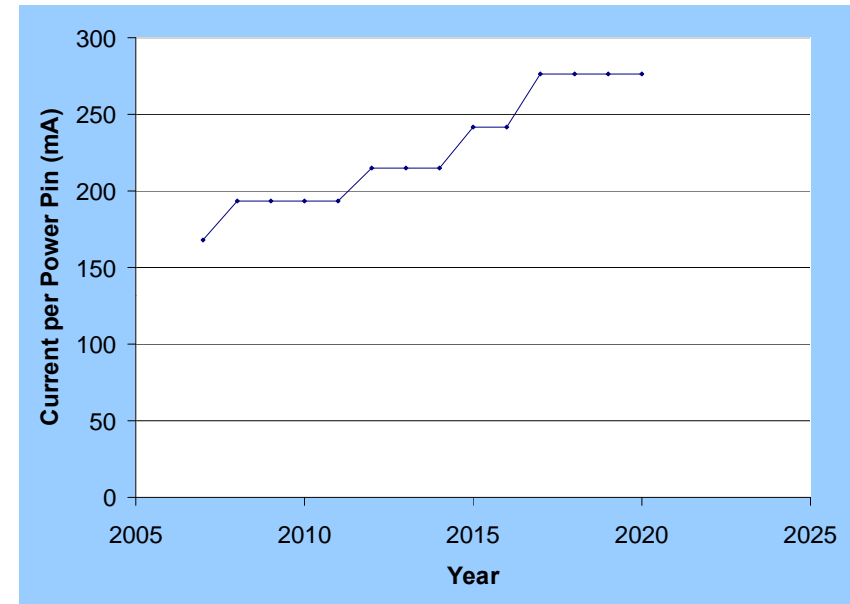
Through-Silicon via - SOI

# Thermal and Power Supply Integrity in 3D

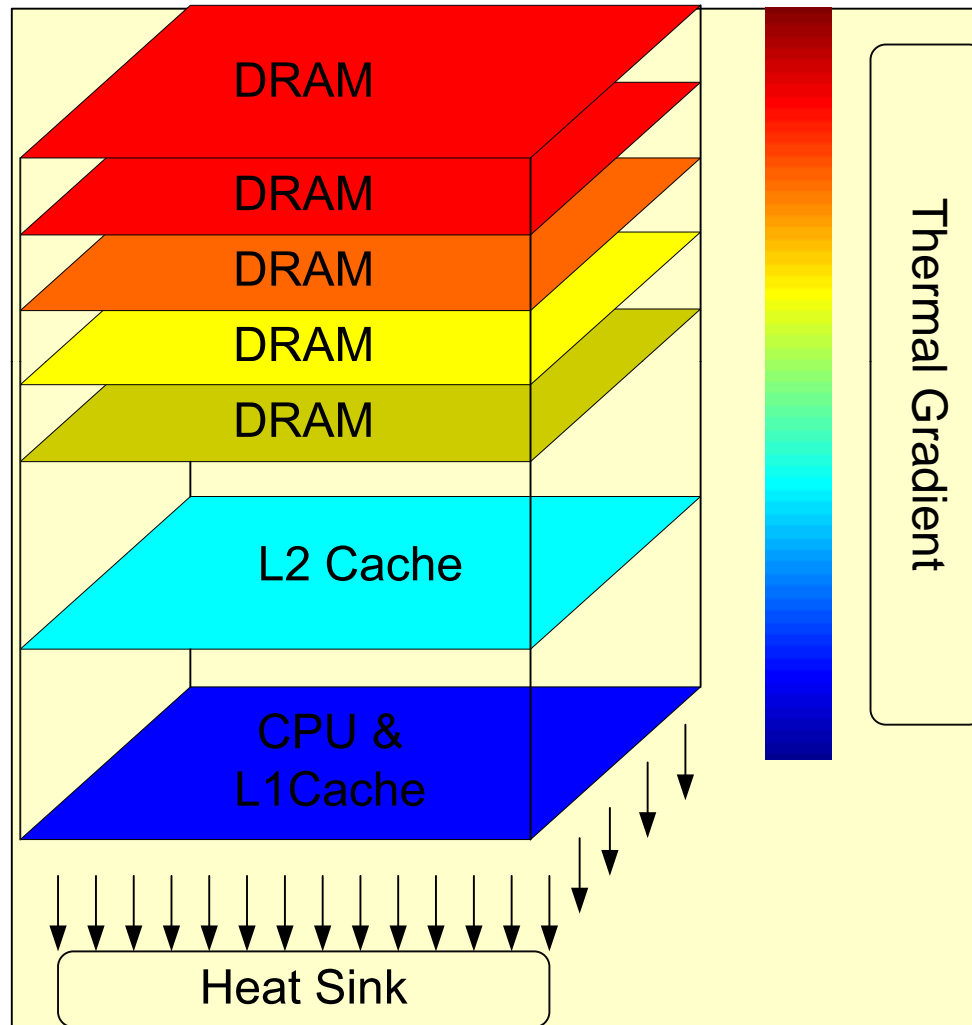
- Higher current density, faster current transients
  - Thermal management: heat sink limitations
  - Power delivery issues: via resistance, limited number of supply pins



*Current per power pin (2D) – ITRS*




# Resolving processor-memory bottlenecks




# Thermal challenges

- Each layer generates heat
- Heat sink at the end(s)
- Simple analysis
  - $\text{Power}(3\text{D})/\text{Power}(2\text{D}) = m$ 
    - $m = \# \text{ layers}$
  - Let  $R_{\text{sink}}$  = thermal resistance of heat sink
  - $T = \text{Power} \times R_{\text{sink}}$ 
    - $m$  times worse for 3D!
- And this does not account for
  - Increased effective  $R_{\text{sink}}$
  - Leakage power effects, T-leakage feedback
- **Thermal bottleneck:** a major problem for 3D
  - Impacts delays, power, reliability



# Power delivery challenges

- Each layer draws current from the power grid
- Power pins at the extreme end tier(s)
- Simple analysis
  - $\text{Current}(3\text{D})/\text{Current}(2\text{D}) = m$ 
    - $m = \# \text{ layers}$
  - Let  $R_{\text{grid}}$  = resistance of power grid
  - $V_{\text{drop}} = \text{Current} \times R_{\text{grid}}$ 
    - $m$  times worse for 3D!
- And this does not account for
  - Increased effective  $R_{\text{grid}}$
  - Leakage power effects, increased current due to T-leakage feedback
- **Power bottleneck:** a major problem for 3D
  - Impacts delays, reliability




# Thermal analysis and optimization in 3D



# Full-chip thermal analysis

- Macroscale thermal analysis for full-chip profiles
  - (as against nanoscale analysis, considering electron-phonon interactions)
- Heat generation
  - Switching gates/blocks act as heat sources
  - Time constants for heat of the order of ms or more



- Thermal equation:

$$k_t \nabla^2 T + g(x, y, z, t) = \rho c_p \frac{\partial T(x, y, z, t)}{\partial t}$$

- Boundary conditions corresponding to the ambient, heat sink, etc.
- Self-consistency: Power = f(Temperature), Temperature = g(Power)



# Thermal analysis

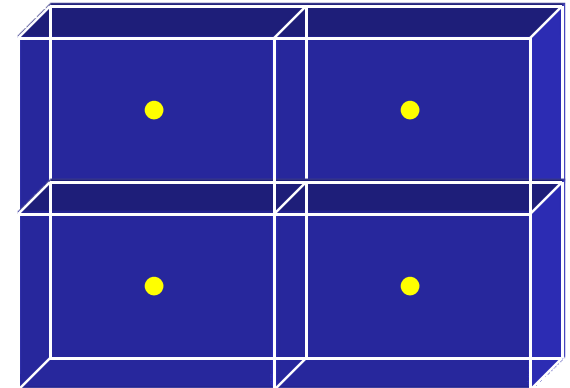
- Thermal equation: partial differential equation

$$k_t \nabla^2 T + g(x, y, z, t) = \rho c_p \frac{\partial T(x, y, z, t)}{\partial t}$$

- Boundary conditions corresponding to the ambient, heat sink, etc.
  - Self-consistency
    - Power is a function of temperature, which is a function of power!
    - Often handled using iterations
-

# The finite difference approach

- Finite difference method
  - Thermal-electrical analogy
    - Can find “thermal resistance” and “thermal capacitance” values between element nodes
  - Steady state:
$$G T = P$$
    - G is the thermal conductance matrix
    - T and P are the temperature and power density vectors
- Same structure as power grid optimization problem, which has been widely addressed in IC design. Can adapt solution techniques
  - Multigrid, etc.
  - **New** - fast random walk based solvers



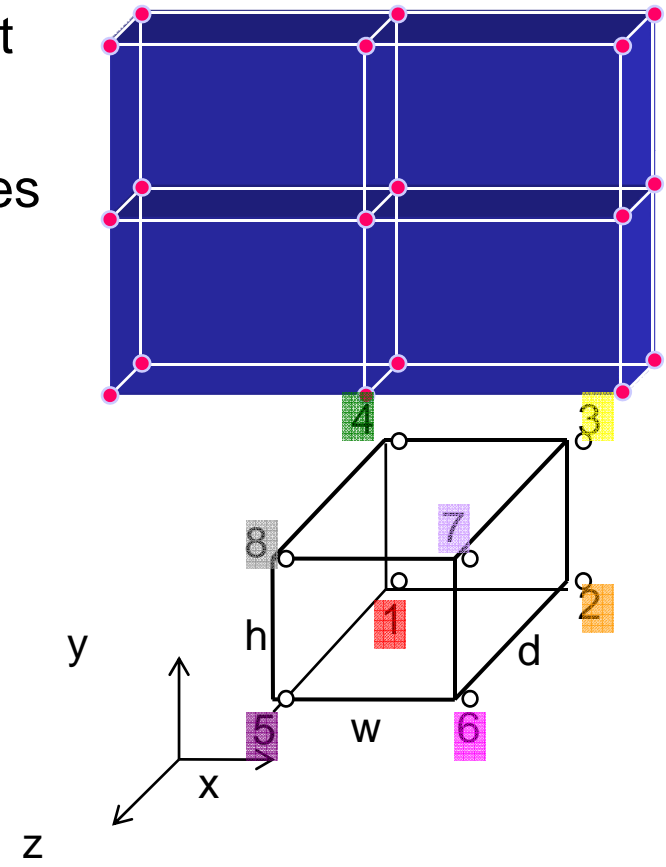
# The finite element approach

- Discretize into elements; use polynomial interpolation based on values at nodes
- Use “element stamps” and assemble these into a larger matrix
- Steady state: apply boundary conditions to get
 
$$KT = P$$
- Rectangular symmetries for on-chip geometries
  - Stamp for a hexahedral element
    - Rows and columns correspond to nodes 1 - 8

$[k] =$


+A	+B	+C	+D	+E	+F	+G	+H	1
+B	+A	+D	+C	+F	+E	+H	+G	2
+C	+D	+A	+B	+G	+H	+E	+F	3
+D	+C	+B	+A	+H	+G	+F	+E	4
+E	+F	+G	+H	+A	+B	+C	+D	5
+F	+E	+H	+G	+B	+A	+D	+C	6
+G	+H	+E	+F	+C	+D	+A	+B	7
+H	+G	+F	+E	+D	+C	+B	+A	8

← nodes



# Performance optimization in 3D

- Minimize power usage
- Rearrange heat sources
- Thermal vias/conduits
- Improved heat sinking




# 3D placement: One approach


Objective Function: 
$$\sum_{\text{each net } i} [WL_i + \alpha_{ILV} \cdot ILV_i] + \alpha_{TEMP} \sum_{\text{each cell } j} [R_j^{cell} P_j^{cell}]$$

- Global Placement
  - Partitioning Placement
  - Thermal Aware Net Weighting
  - Thermal Resistance Reduction Nets
- Coarse Legalization
  - Global Moves/Swaps
  - Local Moves/Swaps
  - Cell Shifting
- Detailed Legalization

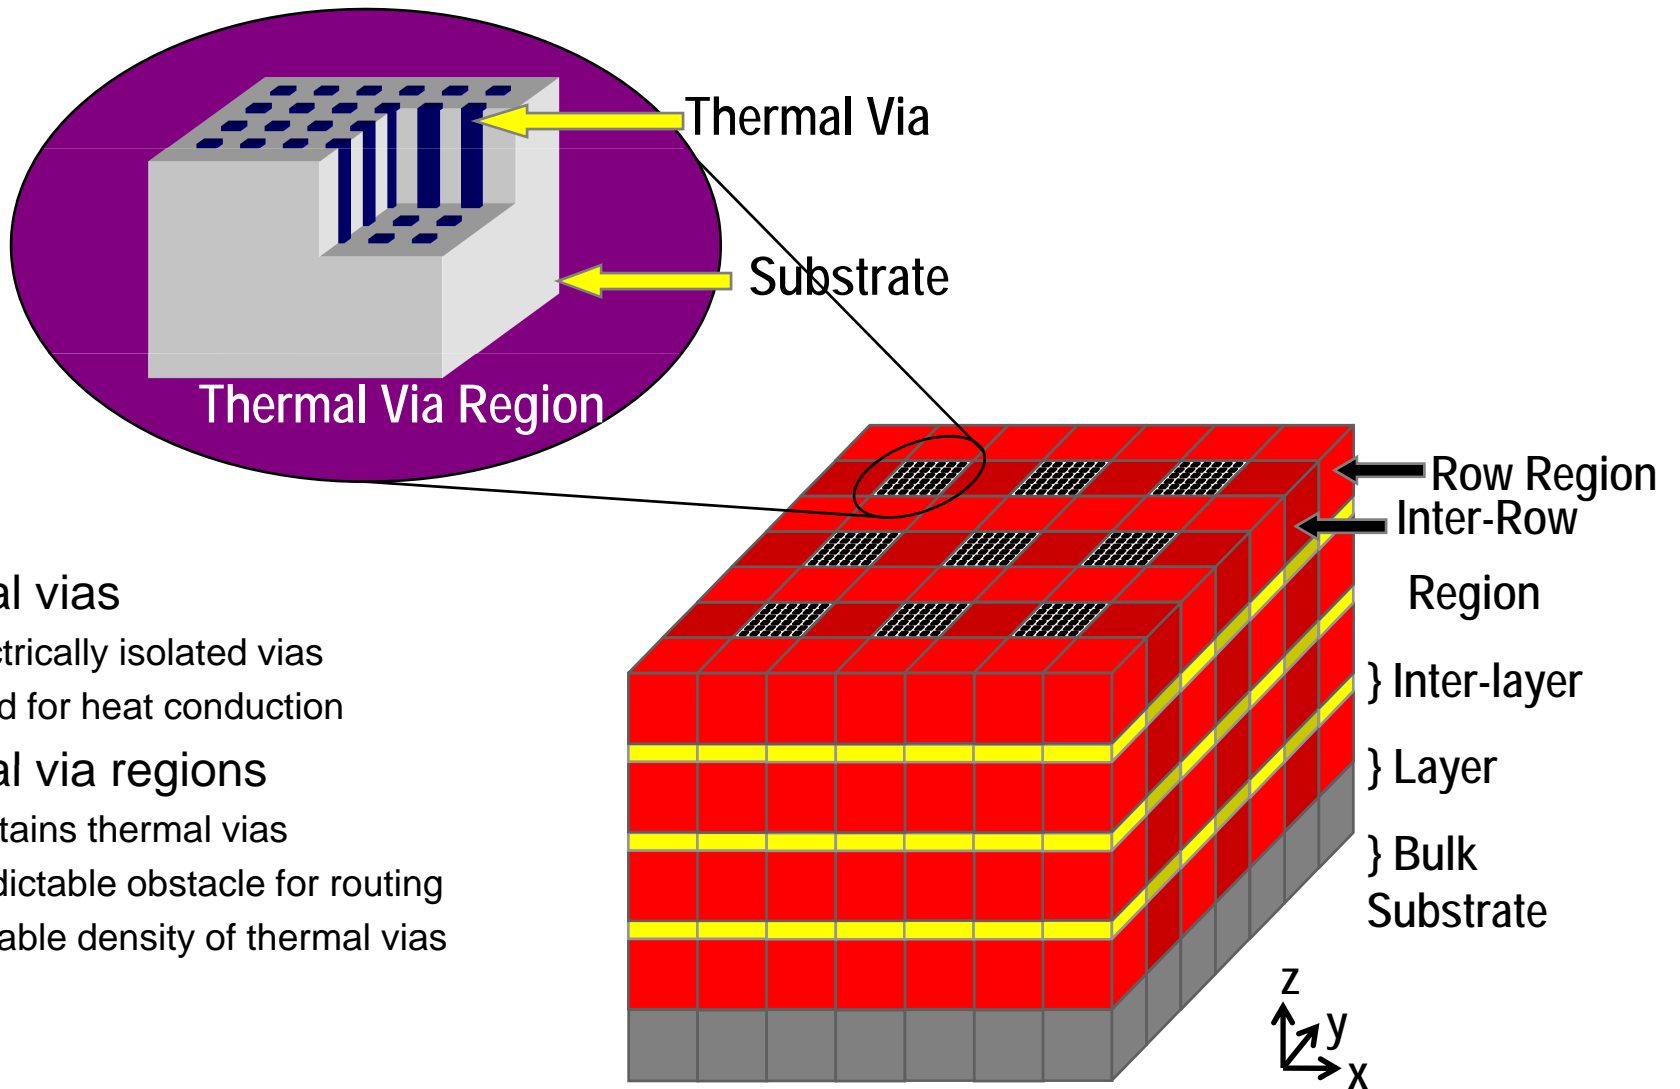
# Thermal resistance reduction nets



Runtime vs. Circuit size



# Heat removal in 3D through thermal vias




- Thermal vias
  - Electrically isolated vias
  - Used for heat conduction
- Thermal via regions
  - Contains thermal vias
  - Predictable obstacle for routing
  - Variable density of thermal vias




# Thermal via insertion


Temperature Profile: Before



Temperature Profile: After




Thermal Via Regions





## 3D routing with integrated thermal via insertion

- Build good heat conduction path through dielectric:
  - Thermal vias: interlayers vias dedicated to thermal conduction.
  - Thermal wires: metal wires improves lateral heat conduction.
  - Thermal vias + thermal wires → a thermal conduction network.
- Thermal wires compete with lateral signal wire routing.
- Thermal vias: large, can block lateral signal routing capacity.




[Zhang, ASPDAC06]

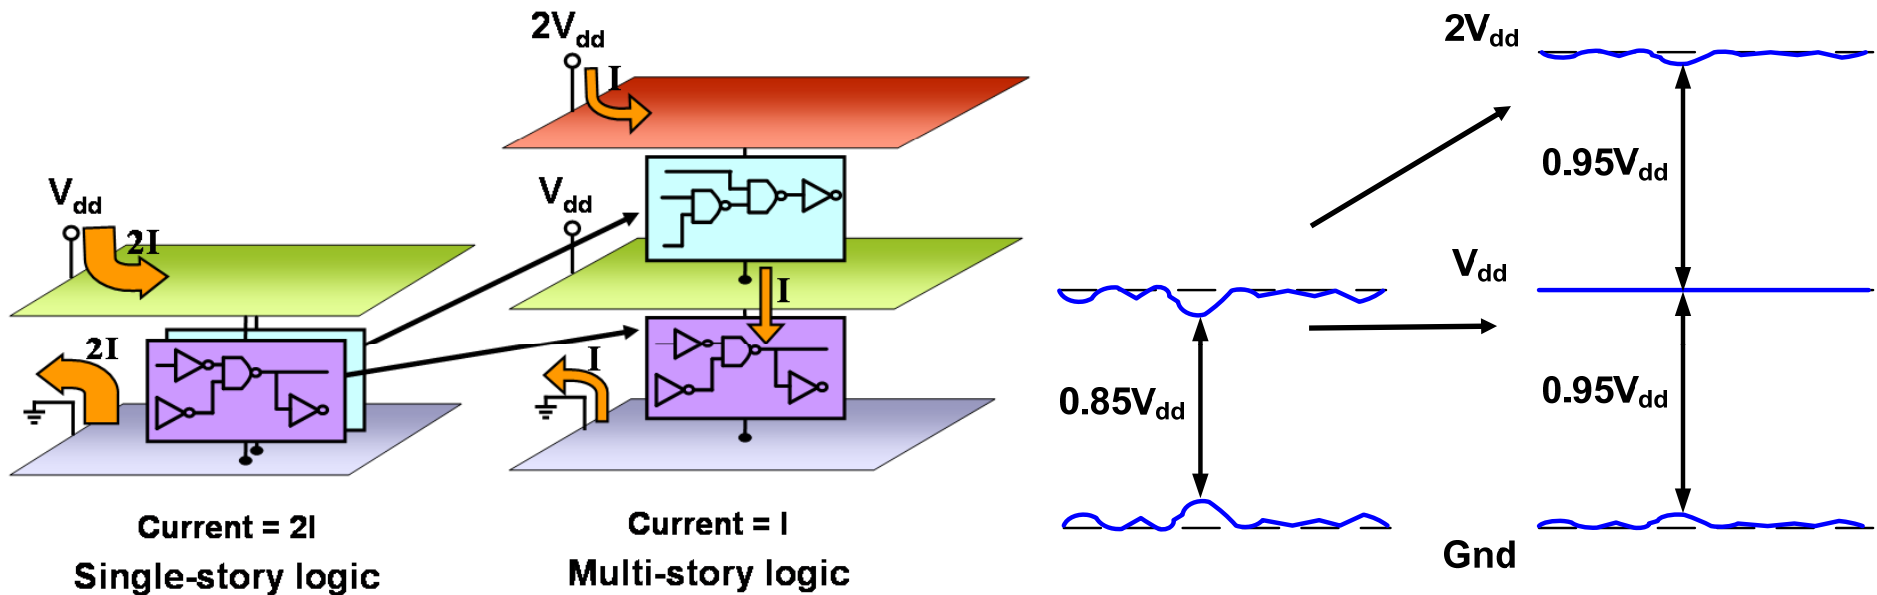
# 3D Power Delivery

# Traditional power delivery

- Requirements
  - $V_{dd}$ , GND signals should be at correct levels (low V drop)
  - Electromigration constraints
    - Current density must never exceed a specification
    - For each wire,  $I_i/w_i < J_{spec}$
  - $dl/dt$  constraints
    - Need to manage  $dl/dt$  to reduce inductive effects
- Techniques for meeting constraints
  - Widening wires
  - Using appropriate topologies
  - Adding decoupling capacitances
- Already challenged for 2D technologies
  - Reliable power delivery hard
  - Decaps get leaky
- Circuit + CAD approaches necessary



# Multi-story power supply



	1-story	2-story
<b>Current</b>	$2I$	$I$
<b>Voltage</b>	$V_{dd}$	$2V_{dd}$
<b>Power</b>	$2V_{dd} \cdot I$	$2V_{dd} \cdot I - \Delta$
<b>Noise</b>	$15\%V_{dd}$	$< 8\%V_{dd}$

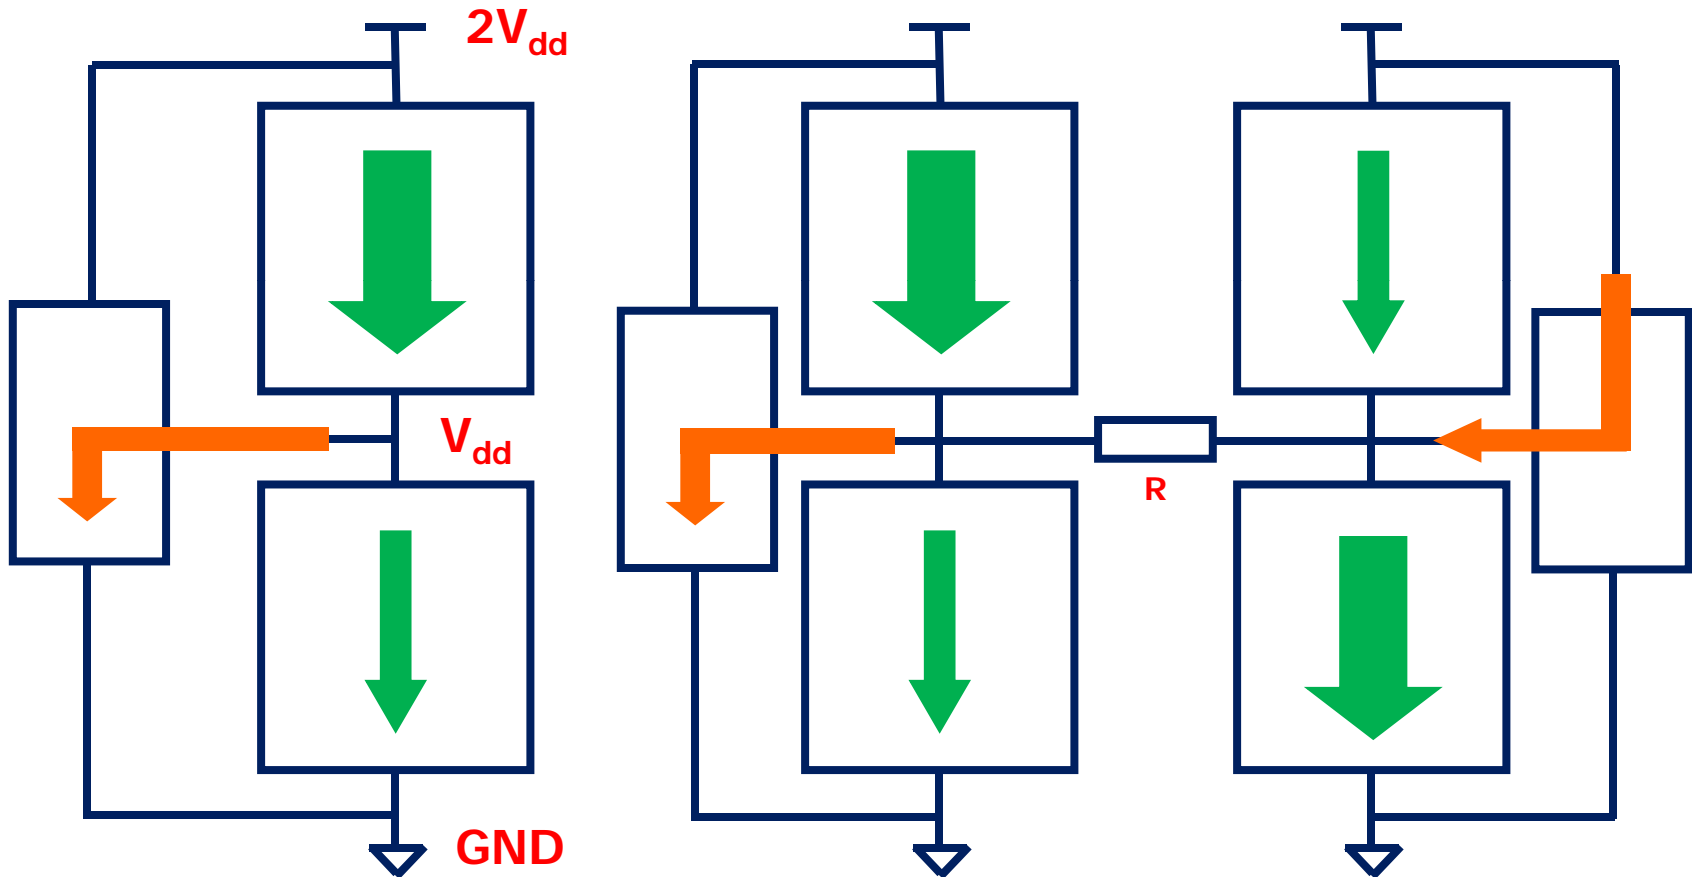
Improved supply noise due to:

- Reduced current magnitude
- Cleaner middle supply voltage


Attractive for 3D chips:

- Isolated substrate for each tier
- Chip is naturally partitioned

# CAD solutions for multi-story circuits



# Multi-story power supply: Test layout



- A test layout in MITLL's SOI process shows a 5.3% area overhead

# Overall Design Flow



**Netlist and block information**

**Floorplanning involving regular modules and regulators**

**Assigning modules using a graph partition-based algorithm**

**Module assignment**

# Estimating the wasted power



$$x_i = \begin{cases} 0 & M_i \text{ works between } 2V_{dd} \text{ and } V_{dd} \\ 1 & M_i \text{ works between } V_{dd} \text{ and } GND \end{cases}$$

$$I_R(t) = \left| \sum_{i=1}^n I_i(t) \times (1 - 2x_i) \right|$$

$$\min \overline{I_R^2(t)} = \overline{\left( \sum_{i=1}^n I_i(t) \right)^2} - 4 \sum_{i < j} \overline{I_i(t) I_j(t)} \times (x_i + x_j - 2x_i x_j)$$

$$\max S = \sum_{i < j} \overline{I_i(t) I_j(t)} (x_i + x_j - 2x_i x_j)$$


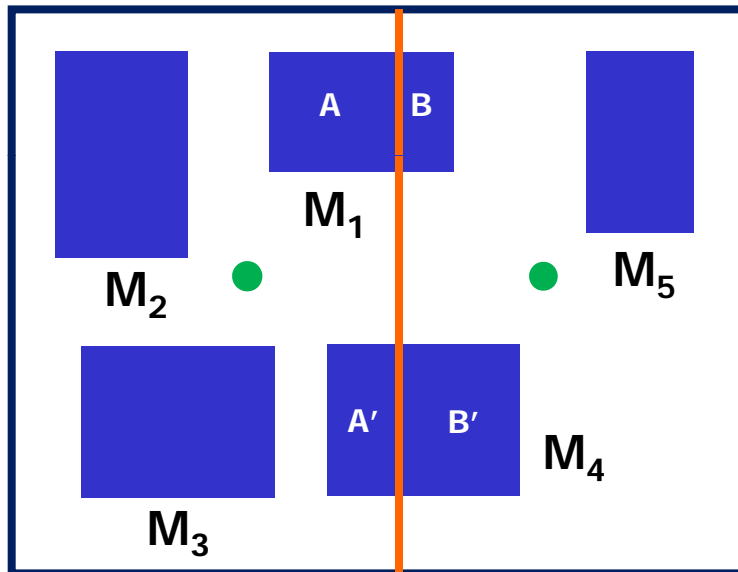
**= 0 if  $x_i = x_j$**

**= 1 if  $x_i \neq x_j$**

**Graph partitioning problem!**



# Constructing the graph



$$w(V_i, V_j) = \left( \sum_{k=1}^K \frac{S_{ik} S_{jk}}{S_i S_j} \right) \overline{I_i(t) I_j(t)}$$

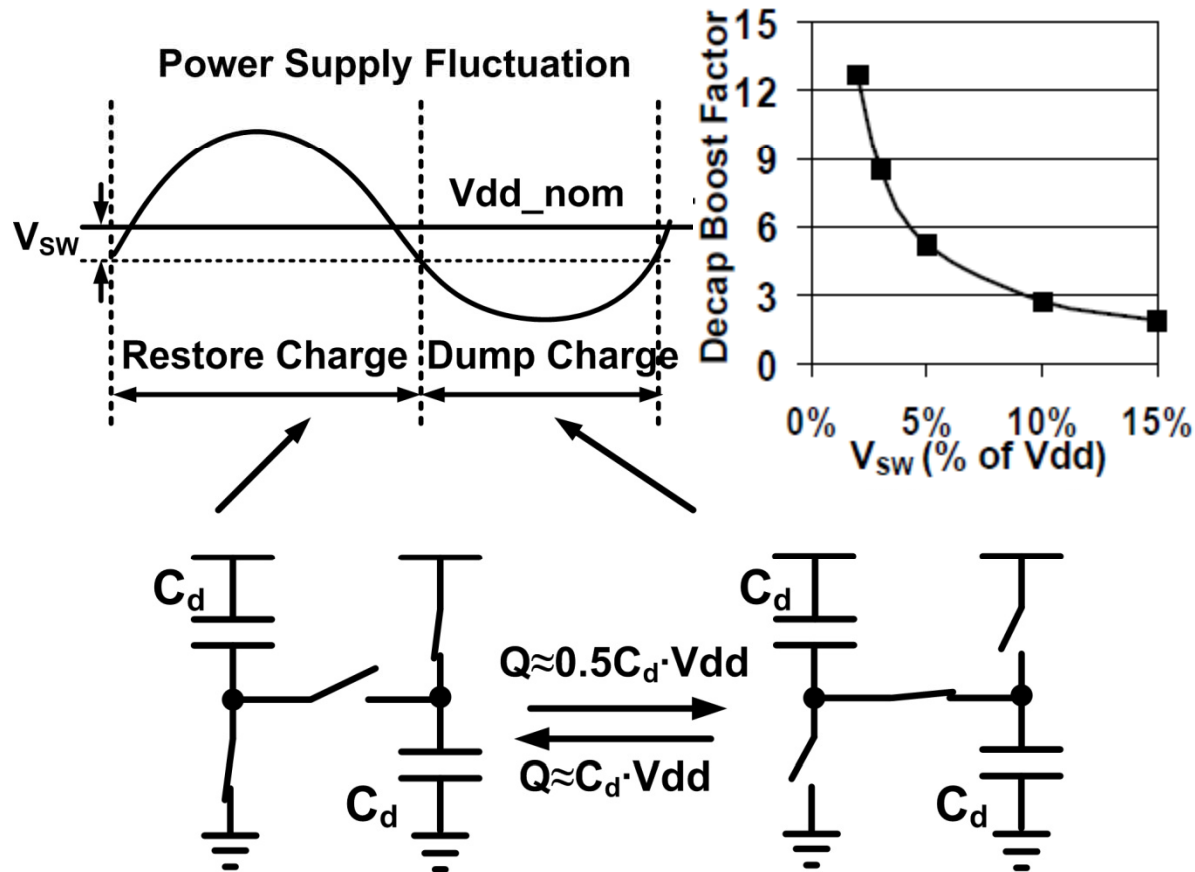
## 3D benchmarks

- Exercised on GSRC floorplanning benchmarks
- Largest floorplan has 300 modules
- Comparison with (slow)simulated annealing method

Layer	WastedPower / UsefulPower (%)		Maximum IR Noise (mV)		Runtime (sec)	
	Partition-Based	Annealing	Partition-Based	Annealing	Partition-Based	Annealing
n100Layer0	3.3	3.1	52.8	62.0	0.03	80
n100Layer1	3.1	3.8	28.9	42.5	0.02	80
n100Layer2	3.7	5.7	45.4	54.6	0.02	80
n200Layer0	8.7	6.4	55.2	88.4	0.31	157
n200Layer1	5.6	6.4	62.1	64.4	0.16	160
n200Layer2	5.6	7.1	77.4	52.7	0.18	165
n300Layer0	4.7	4.5	61.1	56.0	1.83	235
n300Layer1	6.3	6.3	33.4	36.8	0.69	236
n300Layer2	5.4	4.6	46.5	39.5	0.77	236

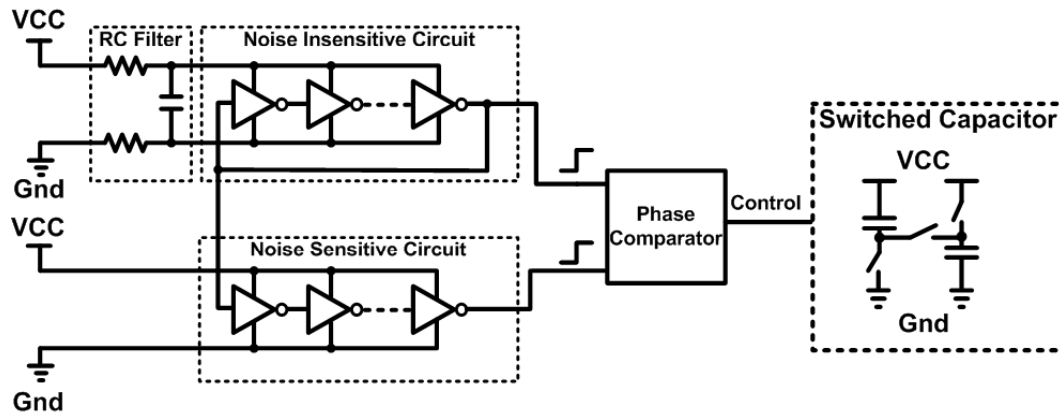
**Runtime Comparison:  $> 10^3$  x speedup over SA**

# Switched decaps for active noise cancellation

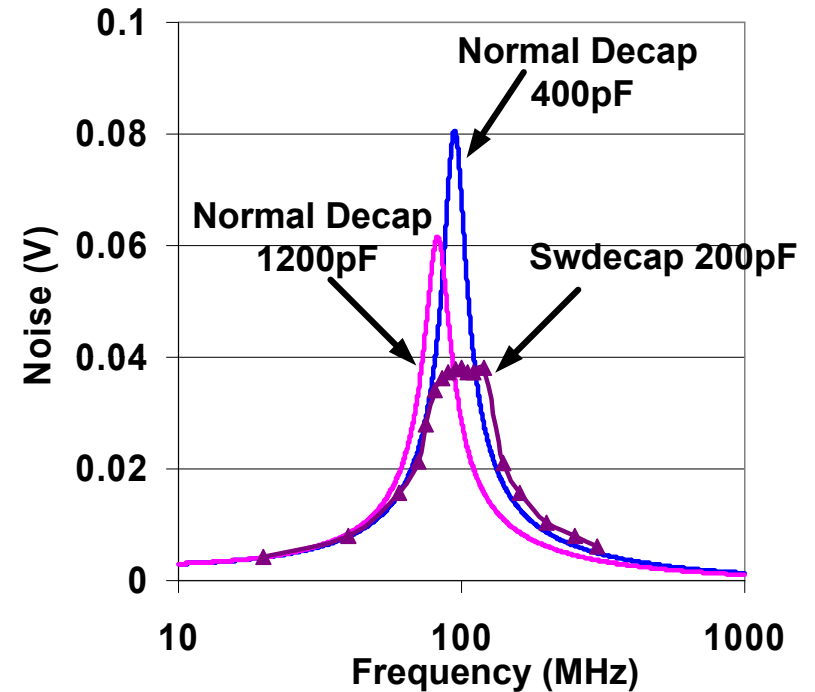


- Charge provided by switched decap ( $=0.5C \cdot V_{dd} + C \Delta V_{dd}/2$ ) much larger than that of a conv. decap ( $=2C \cdot \Delta V_{dd}$ )
- For a supply noise ( $\Delta V_{dd}$ ) of 5%, effective decap value is boosted by 7.5X

# Supply noise cancellation: Results

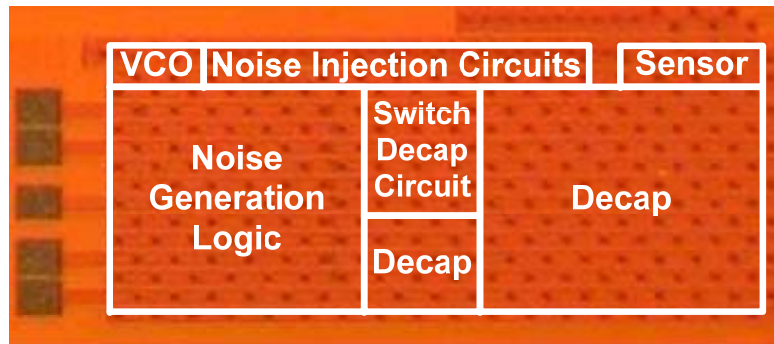


**Digital Switched Decap Circuit**

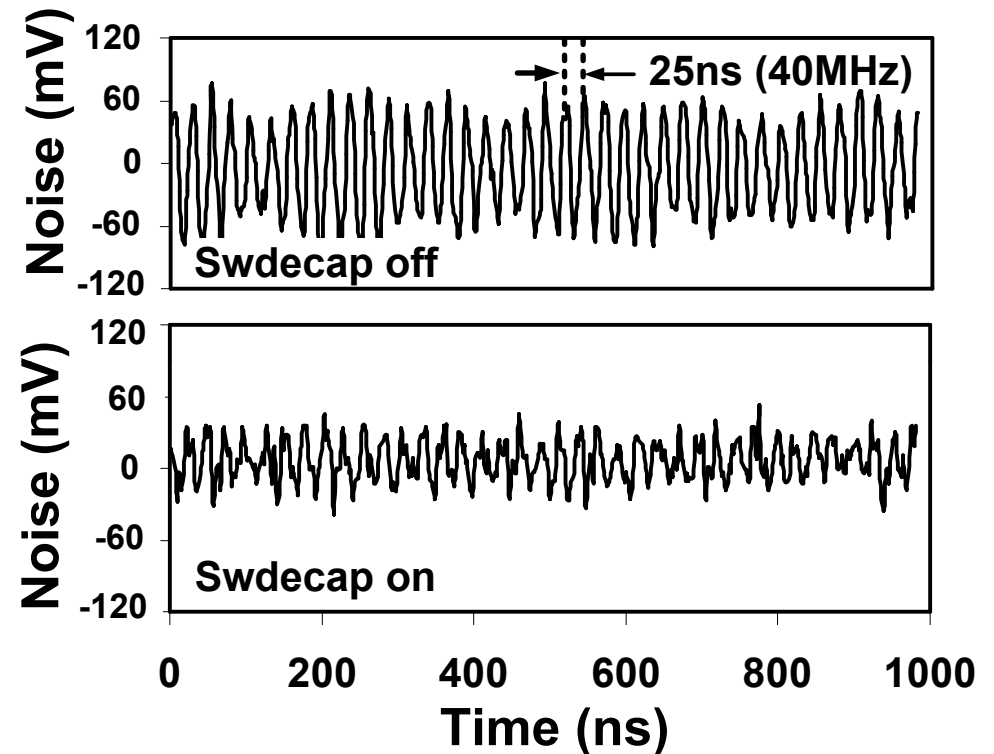


- 200pF switched decap has lower noise than 1200pF conventional decap
- 5–11X boost over passive decaps depending on supply noise magnitude

# Proof of concept: Switched decap test chip



<b>Technology</b>	0.13 $\mu$ m CMOS
<b>Quiescent Current</b>	0.54mA
<b>Regulation Freq.</b>	10MHz-300MHz
<b>Regulator Area (w/o decap)</b>	100 $\mu$ m $\times$ 70 $\mu$ m
<b>Regulator Area (w/ 300pF decap)</b>	190 $\mu$ m $\times$ 220 $\mu$ m
<b>Total Die Area</b>	0.9mm $\times$ 1.8mm



- **2.2-9.8dB reduction of the 40MHz resonant noise using 100-300pF switched decaps**


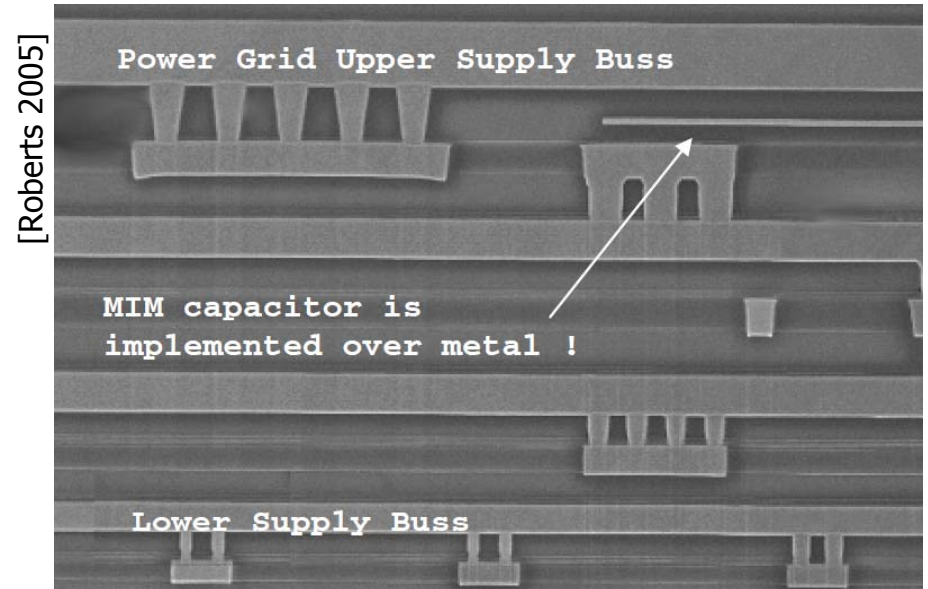
# Comparison with passive damping

<b>Swdecap Value</b>	<b>Resonant Suppression</b>	<b>Equivalent Passive Decap</b>	<b>Decap Boost</b>
100pF	2.2dB	500pF	5X
200pF	5.5dB	1500pF	7.5X
300pF	9.8dB	3500pF	11X

# MIM decaps

- Capacitance density\*
  - CMOS –  $17.3 \text{ fF}/\mu\text{m}^2$  at 90nm
  - MIM –  $8.0 \text{ fF}/\mu\text{m}^2$
- Leakage density\*
  - CMOS –  $1.45\text{e-}4 \text{ A}/\text{cm}^2$
  - MIM –  $3.2\text{e-}8 \text{ A}/\text{cm}^2$
- Congestion
  - MIM – routing blockage
- Described in paper 2D-4, ASPDAC09

\* Numbers deduced from Roberts et al., IEDM05 and PTM simulations



# Conclusion

- Power, thermal issues are major bottlenecks for 3D integration
  - The root cause of both is closely related
- Solutions can come through low power design, physical design, and novel circuit techniques