AVS-Aware Power-Gate Sizing for Maximum Performance and Power Efficiency of Power-Constrained Processors

Abhishek Sinkar and Nam Sung Kim

January 28, 2011

Department of Electrical and Computer Engineering
University of Wisconsin - Madison

Outline

- Introduction
- Impact of process variations and PG-device size
  - Impact on VV_{DD}
  - Impact on F_{MAX} and P_{TOT}
- AVS-aware PG-device size optimization
  - Algorithm and simulation result
- Impact of WID variation on PG sizing
  - Global clocking
  - Frequency island clocking
- Experimental methodology
- Conclusion

Outline

- **Introduction**
- Impact of process variations and PG-device size
  - Impact on VV_{DD}
  - Impact on F_{MAX} and P_{TOT}
- AVS-aware PG-device size optimization
  - Algorithm and Simulation Result
- Impact of WID variation on PG sizing
  - Global clocking
  - Frequency island clocking
- Experimental methodology
- Conclusion

Introduction

Parallel processing with multiple cores

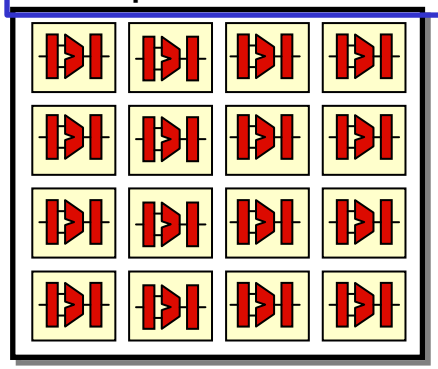

"Gigascale Integration Challenges and Opportunities", Shekhar Borkar, Intel

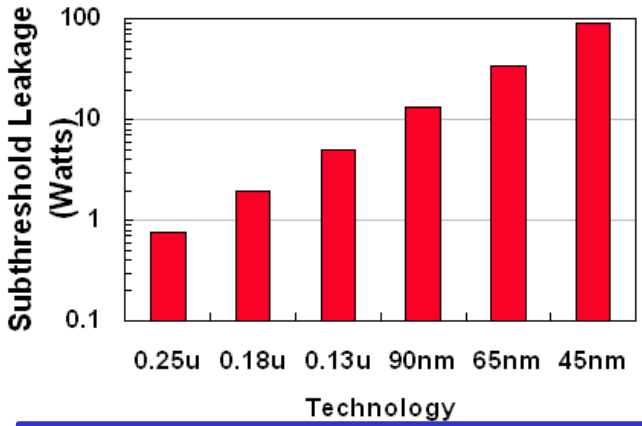

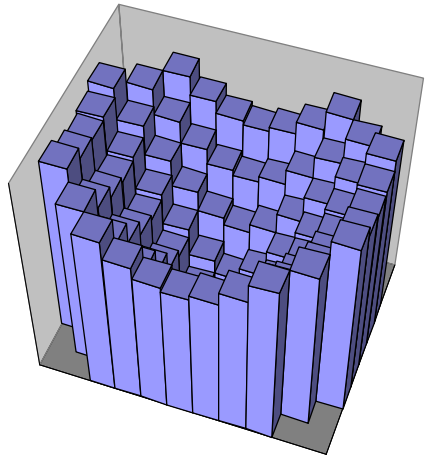

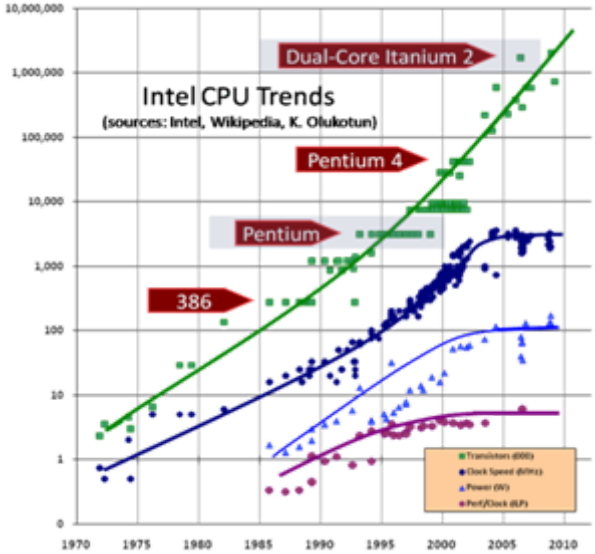

P_{LKG} increasing fraction of P_{TOT} with technology.

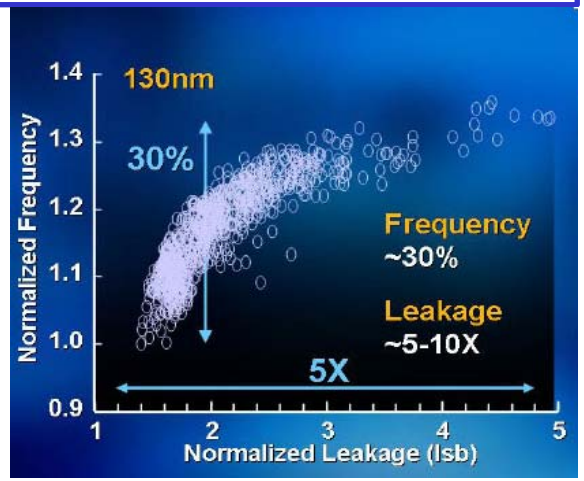

Intel

Die to Die Variation (Wafer scale)

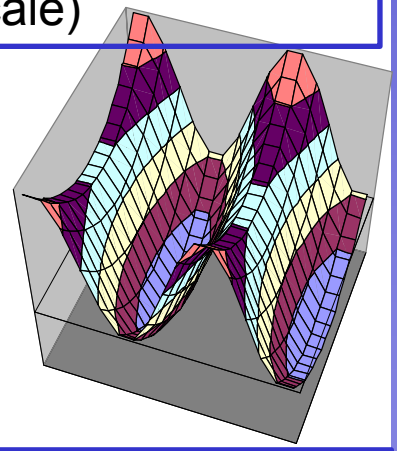

Performance and Power : increasing with technology scaling

Within Die Variation

(Die scale)

Contributions

- Analyze the impact of D2D variations on
  - Virtual rail voltage (VV_{DD})
  - Maximum operating frequency (F_{MAX})
  - Total power consumption (P_{TOT})
- Propose algorithm to find
  - Optimal PG size
  - Optimal degree of AVS
- Extend algorithm to
  - Multicore processors w/ WID variation
    - Global clocking
    - Frequency Island clocking

Outline

- Introduction
- **Impact of process variations and PG-device size**
  - Impact on VV_{DD}
  - Impact on F_{MAX} and P_{TOT}
- AVS-aware PG-device size optimization
  - Algorithm and simulation result
- Impact of WID variation on PG sizing
  - Global clocking
  - Frequency island clocking
- Experimental methodology
- Conclusion

Process Variation Impact on VV_{DD}

- VV_{DD} follows load line

  $VV_{DD} = V_{DD} - R_{PG} \times I_{TOT}(VV_{DD})$

- Fast Die
  - More I_{TOT} → Lower VV_{DD}
- Slow Die
  - Less I_{TOT} → Higher VV_{DD}
- Additionally
  - $R_{PG\text{-}SLOW} > R_{PG\text{-}FAST}$

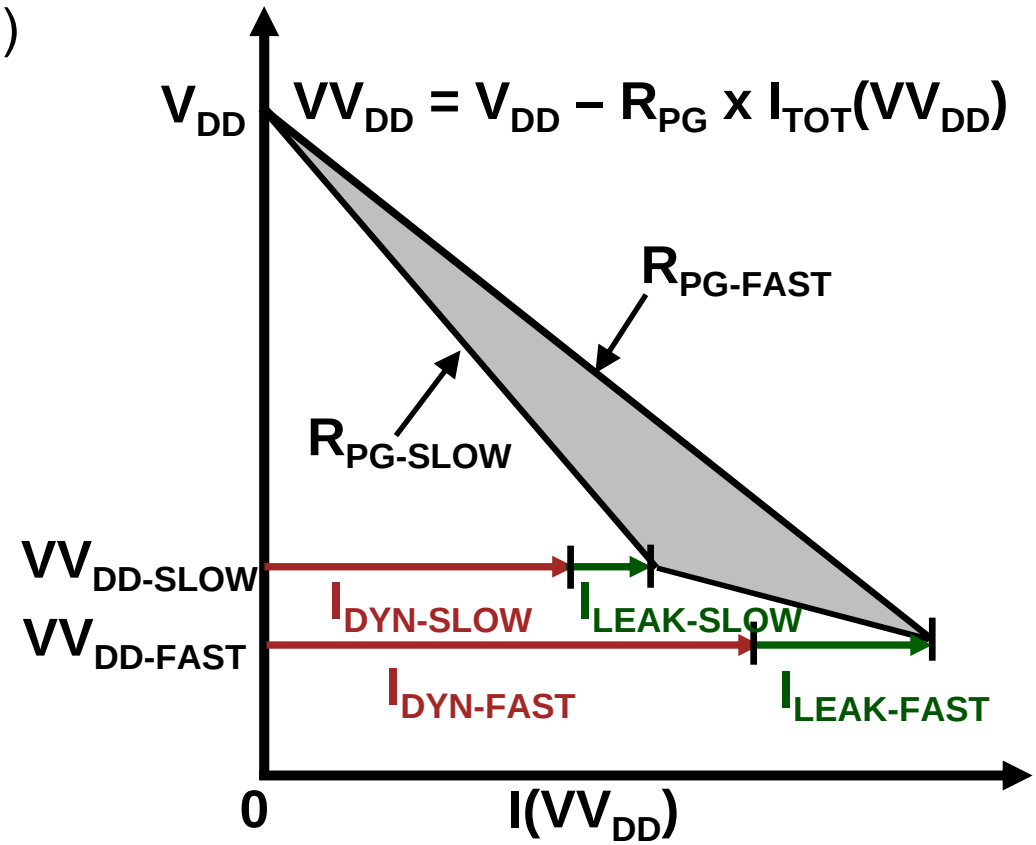

P. V. + PG Size Impact on VV_{DD}

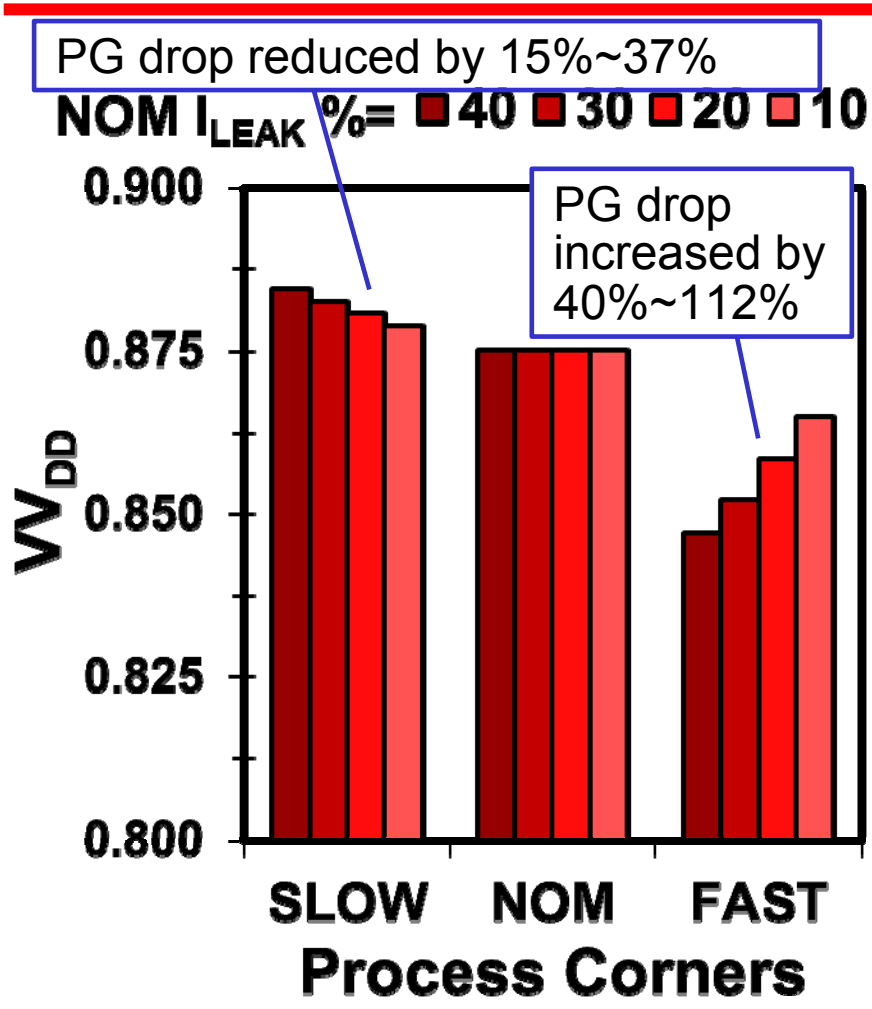

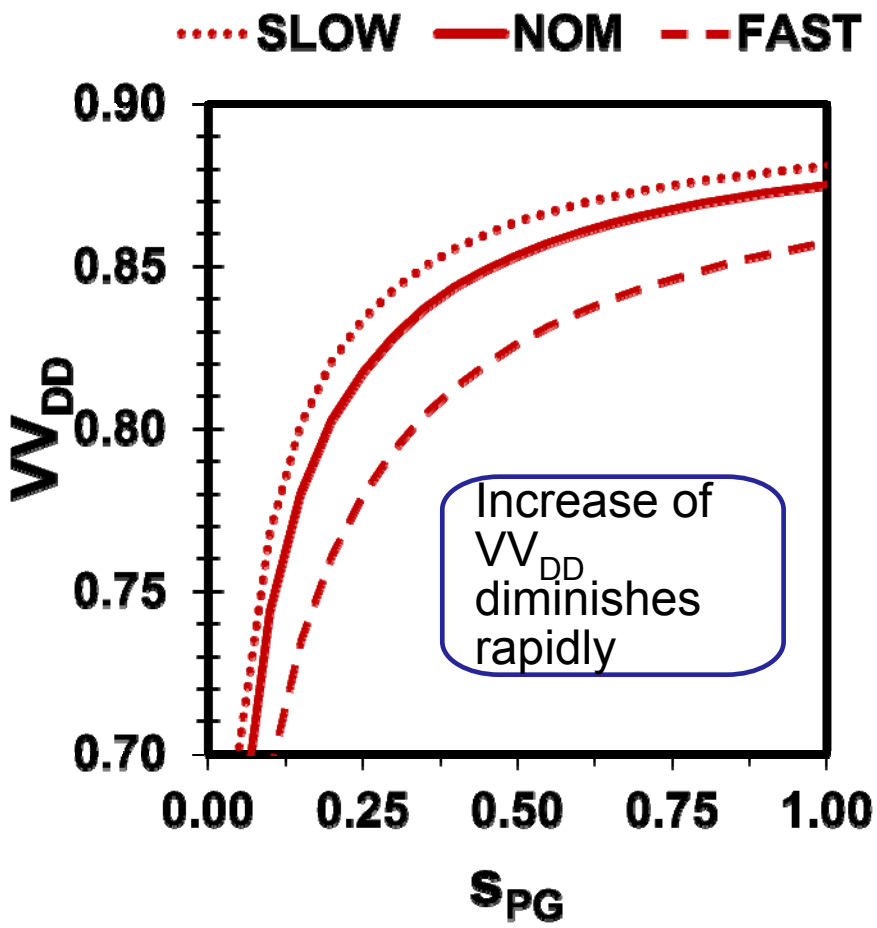

s_{PG} increases → R_{PG} decreases → VV_{DD} increases

P. V. + PG Size Impact on F_{MAX} & P_{TOT}

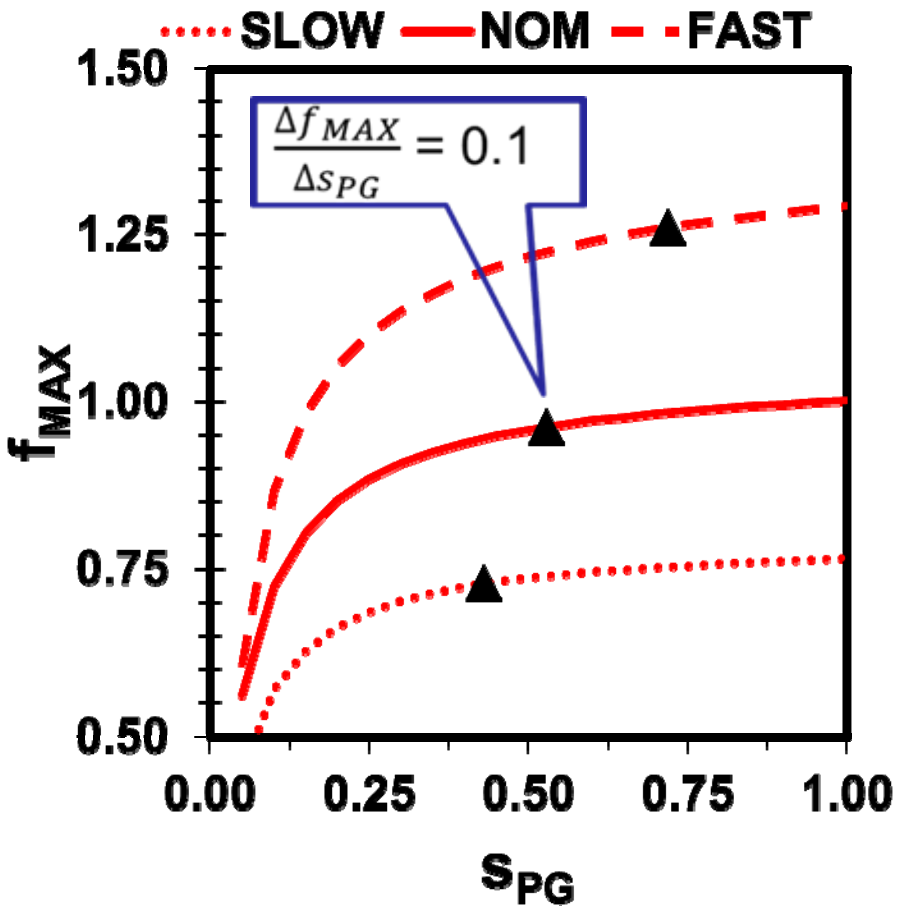

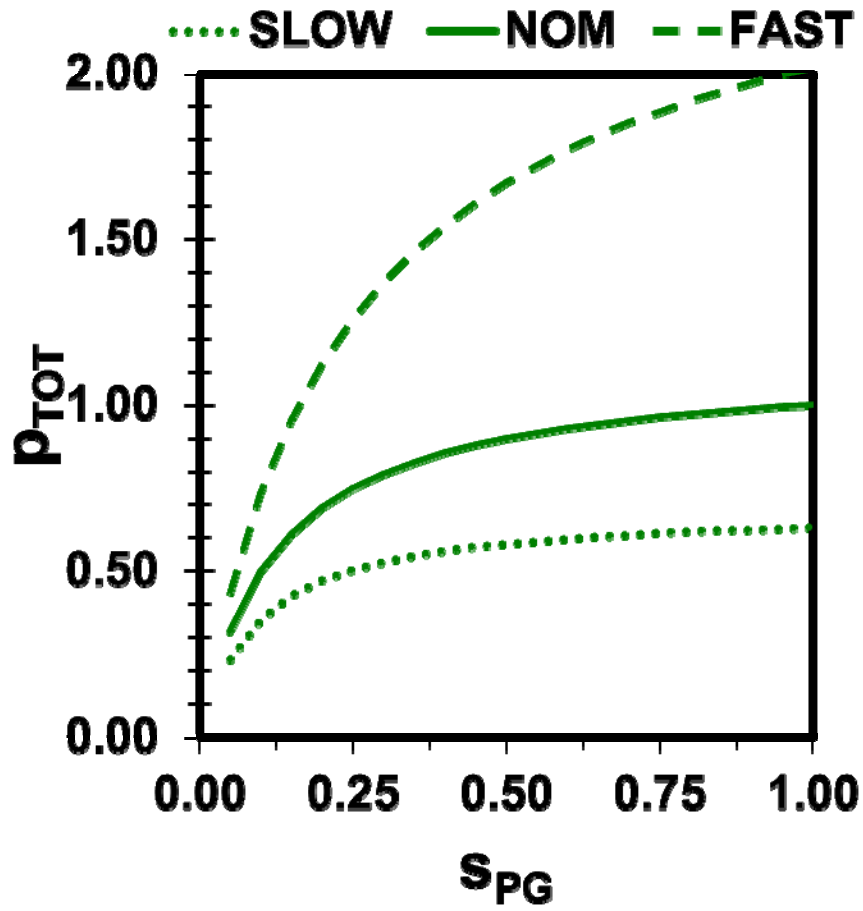

- F_{MAX} increase diminishes rapidly
- P increases faster than F_{MAX}
- Larger PG device suitable for fast die, smaller for slow die

Outline

- Introduction
- Impact of process variations and PG-device size
  - Impact on VV_{DD}
  - Impact on F_{MAX} and P_{TOT}
- **AVS-aware PG-device size optimization**
  - Algorithm and simulation result
- Impact of WID variation on PG sizing
  - Global clocking
  - Frequency island clocking
- Experimental methodology
- Conclusion

AVS Aware PG Size Optimization

- PG size and V_{DD} both control F_{MAX} and P_{TOT}
- ROIP^3/W (Rate Of Increase of Performance^3-per-Watt(p^3/w):
  - $p^3/w(s_{PG}) = (f_{MAX}(s_{PG}))^3 / p_{TOT}(s_{PG})$

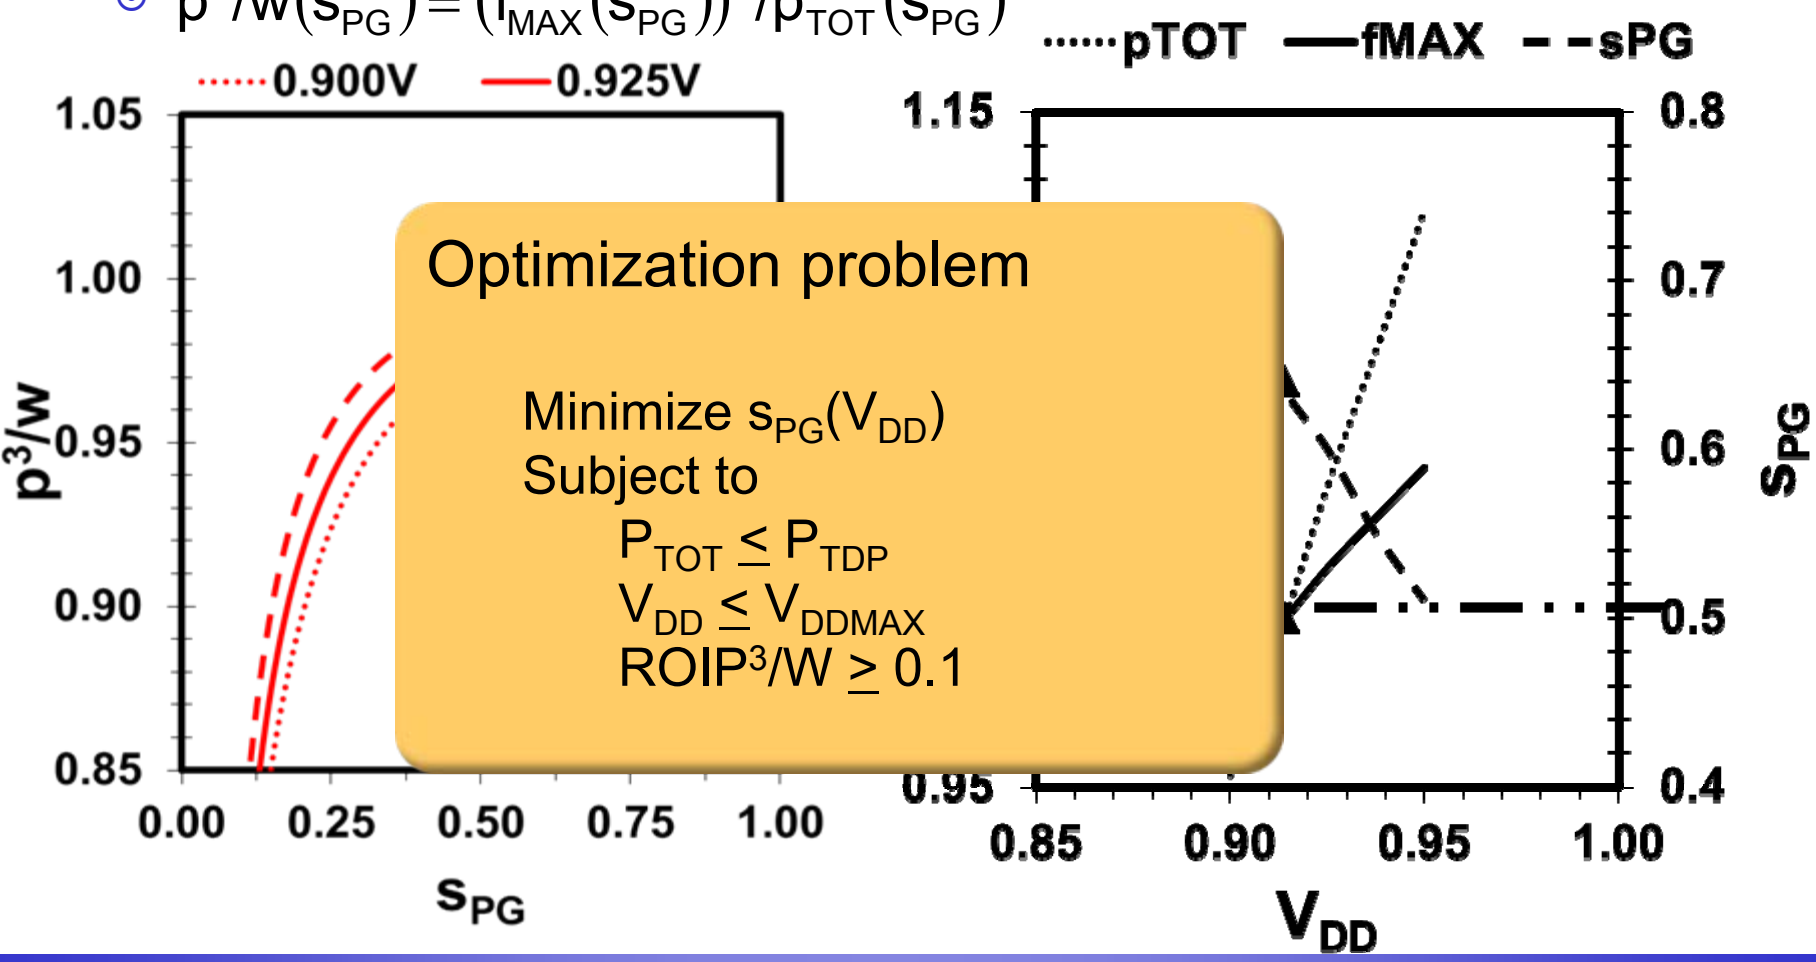

Optimization problem

Minimize $s_{PG}(V_{DD})$
Subject to
$P_{TOT} \leq P_{TDP}$
$V_{DD} \leq V_{DDMAX}$
ROIP^3/W ≥ 0.1

Algorithm

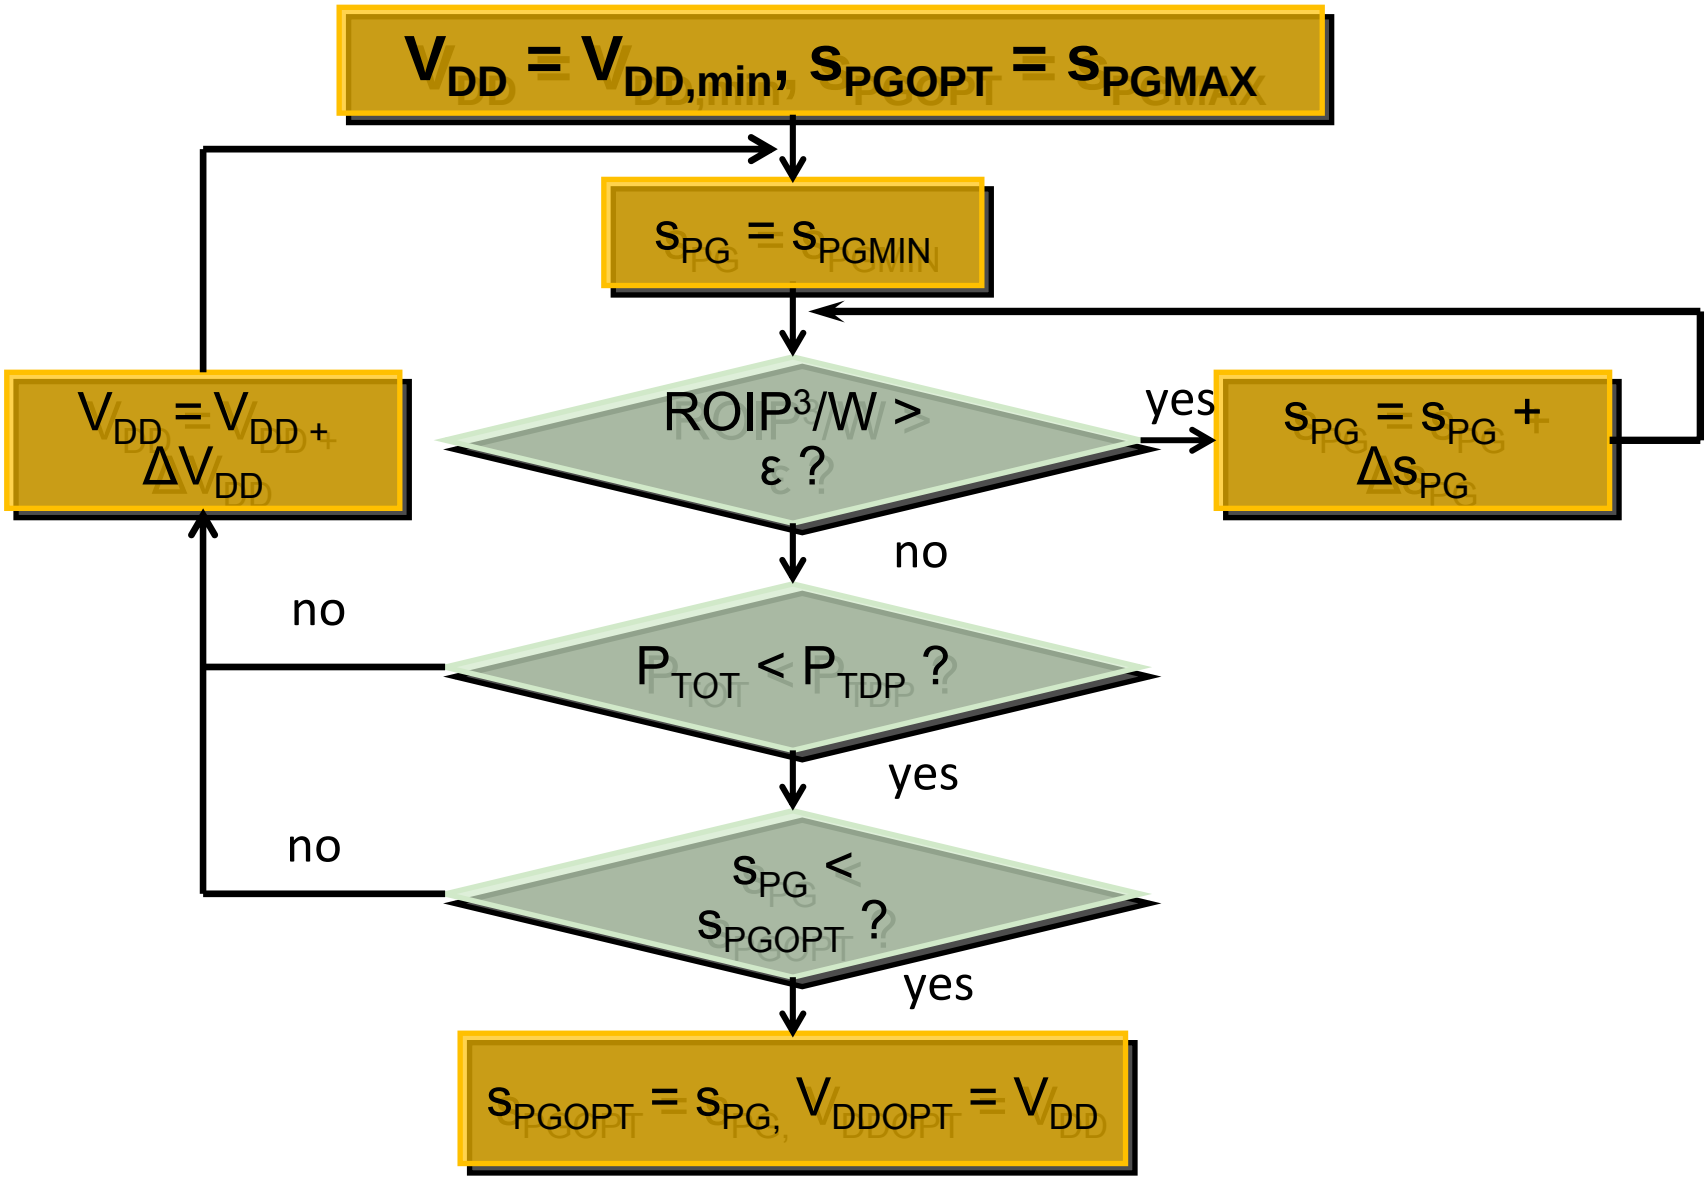
$V_{DD} = V_{DD,min}$, $s_{PGOPT} = s_{PGMAX}$
$s_{PG} = s_{PGMIN}$
$ROIP^3/W > \varepsilon$?
yes
$s_{PG} = s_{PG} + \Delta s_{PG}$
no
$P_{TOT} < P_{TDP}$?
no
$V_{DD} = V_{DD} + \Delta V_{DD}$
yes
$s_{PG} < s_{PGOPT}$?
no
yes
$s_{PGOPT} = s_{PG}$, $V_{DDOPT} = V_{DD}$

Simulation Result

Proc. Corners	Slow		Nom		Fast	
P_{TDP}	65W	70W	90W	100W	120W	130W
V_{DD} @ s_{PG}=1	0.945	0.955	0.900	0.923	0.825	0.840
s_{PGOPT}	0.755	0.715	0.640	0.515	0.568	0.455
V_{DDOPT}	0.955	0.975	0.915	0.948	0.845	0.875
f_{MAX}	0.999	0.999	0.999	~1.000	~1.000	~1.000

- 24.5%, 36% and 43.2% reduction in PG size for slow, nominal and fast corners

Outline

- Introduction
- Impact of process variations and PG-device size
  - Impact on VV_{DD}
  - Impact on F_{MAX} and P_{TOT}
- AVS-aware PG-device size optimization
  - Algorithm and simulation result
- **Impact of WID variation on PG sizing**
  - Global clocking
  - Frequency island clocking
- Experimental methodology
- Conclusion

WID Variation

Systematic L variation

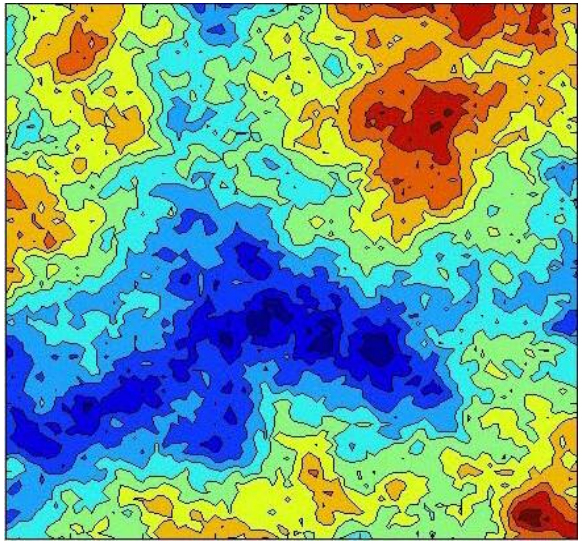

Normalized $f(V_{DD})$, $I(V_{DD})$

1.03 1.15	1.06 1.38	1.01 1.04	1.00 1.00
LLC		LLC	
1.04 1.20	1.08 1.54	1.01 1.20	1.01 1.29
1.05 1.61	1.11 1.82	1.12 1.80	1.08 1.44
LLC		LLC	
1.07 1.58	1.05 1.38	1.05 1.27	1.00 1.14

12% Faster

82% Leakier

S. Sarangi et al. IEEE Trans. On Semiconductor Manufacturing, vol. 21, no. 1, pp. 3~13, Feb. 2008.

- WID spatially correlated
- Result in C2C F_{MAX} and I_{LEAK} variation
- As # of cores increases
  - Relative variation among cores more significant

Global Clocking

- Limits F_{MAX} of multicore processor to that of slowest core

- Total power consumption of a N core processor
  - $$P_{TOT} = \sum_{i=1}^{N} \left(F_{MAX,j}\left(VV_{DD,j}\right) \times C_{EFF} \times VV_{DD,i} + I_{LEAK,i}\left(VV_{DD,i}\right) \right) \times VV_{DD,i}$$
  - where N = no. of cores
    j = index of slowest core on die
    C_{eff} = effective switched capacitance per core

Global Clocking

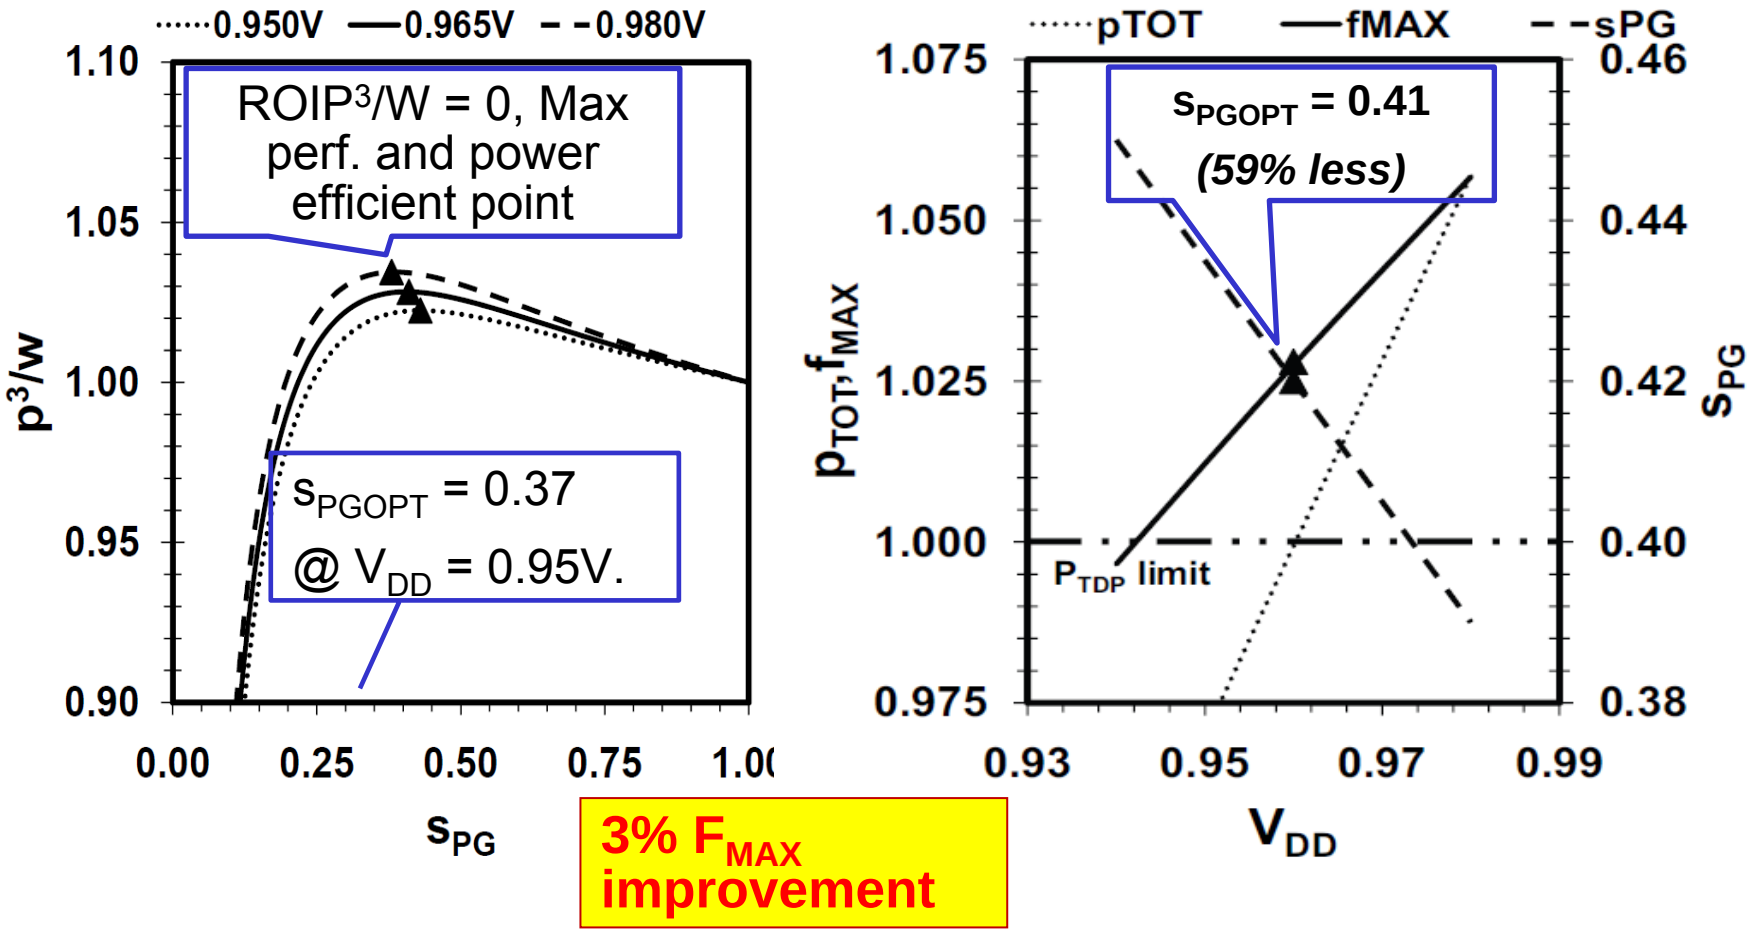

As PG size increases, faster(leakier) cores increase I_{LEAK}, F_{MAX} is fixed by slowest core.

Frequency Island Clocking

- Each core runs at its own max. frequency
- Total power consumption of a N core processor
  - $P_{TOT} = \sum_{i=1}^{N} \left(F_{MAX,i}\left(VV_{DD,i}\right) \times C_{EFF} \times VV_{DD,i} + I_{LEAK,i}\left(VV_{DD,i}\right) \right) \times VV_{DD,i}$
- For compute-bound workloads with sufficient # of threads
  - $\text{Performance(Throughput)} = \left[\sum_{1}^{N} F_{MAX,i} \right] / N$

Throughput Experiment with FI Clk.

As # of cores increases, C2C F_{MAX} variations increase leading to higher $F_{MAX,avg}$

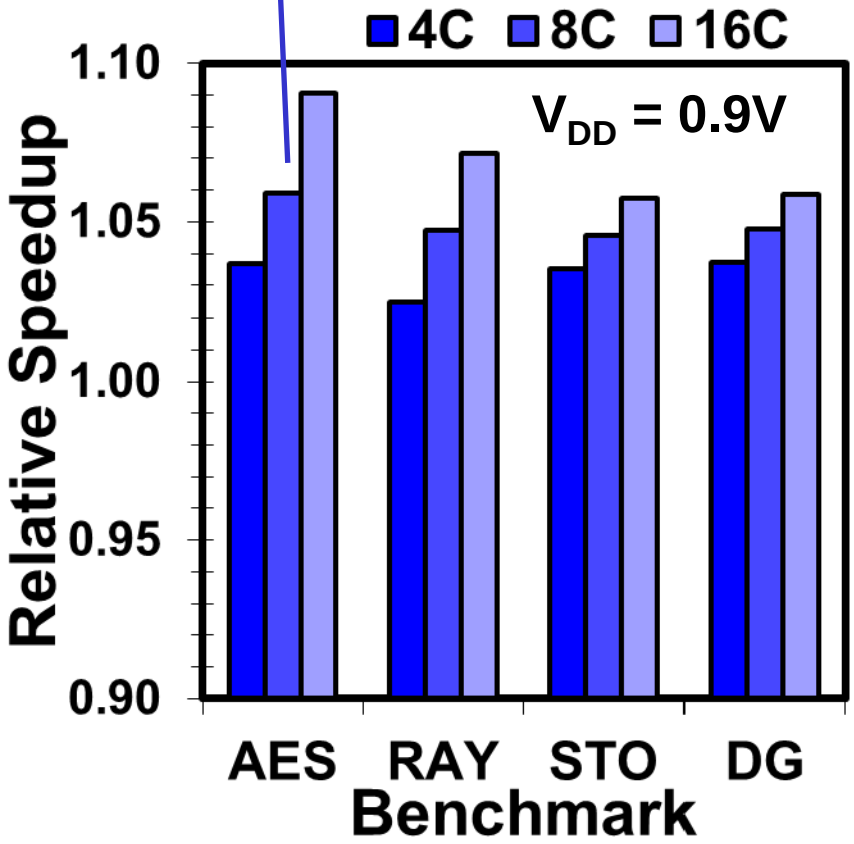

N = 16

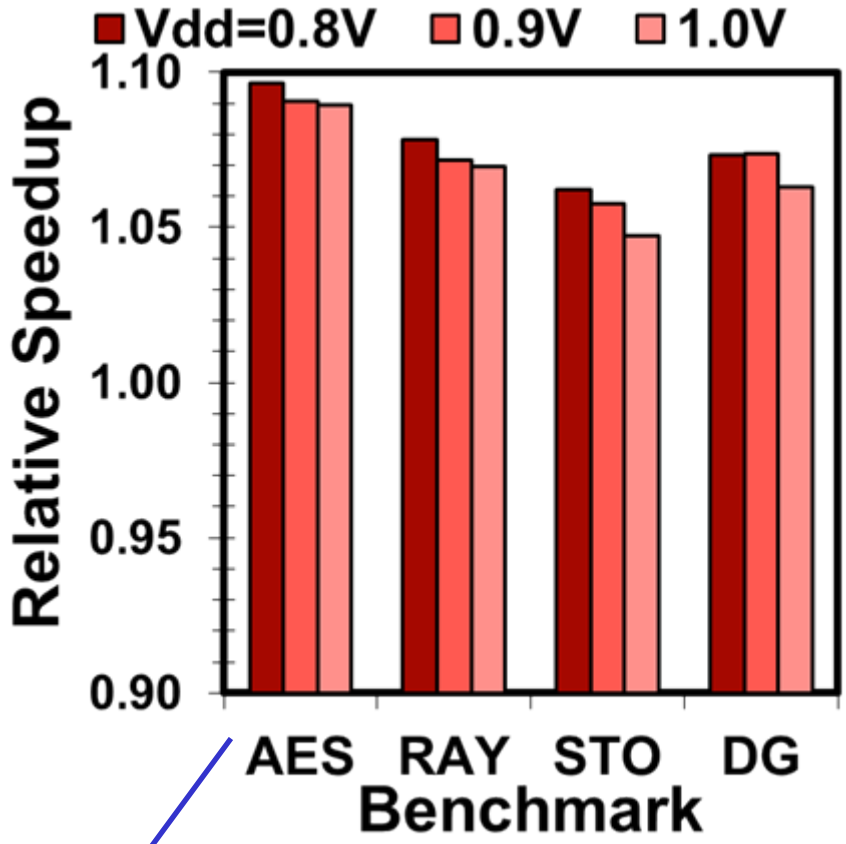

Compute bound workloads with large # of threads, future RMS applications

Speedup degrades as V_{DD}(thus F_{MAXavg}) increases due to limited main memory resource

PG Sizing with FI clocking and WID Var.

- As #. of cores/die ↑
  - V_{DDOPT} ↓
  - s_{PGOPT} ↑
- In FI clocking faster(and leakier) cores increase $F_{MAX,AVG}$
- Larger PG increases I_{LEAK} in fast cores; so V_{DDOPT} must reduce to satisfy P_{TDP}

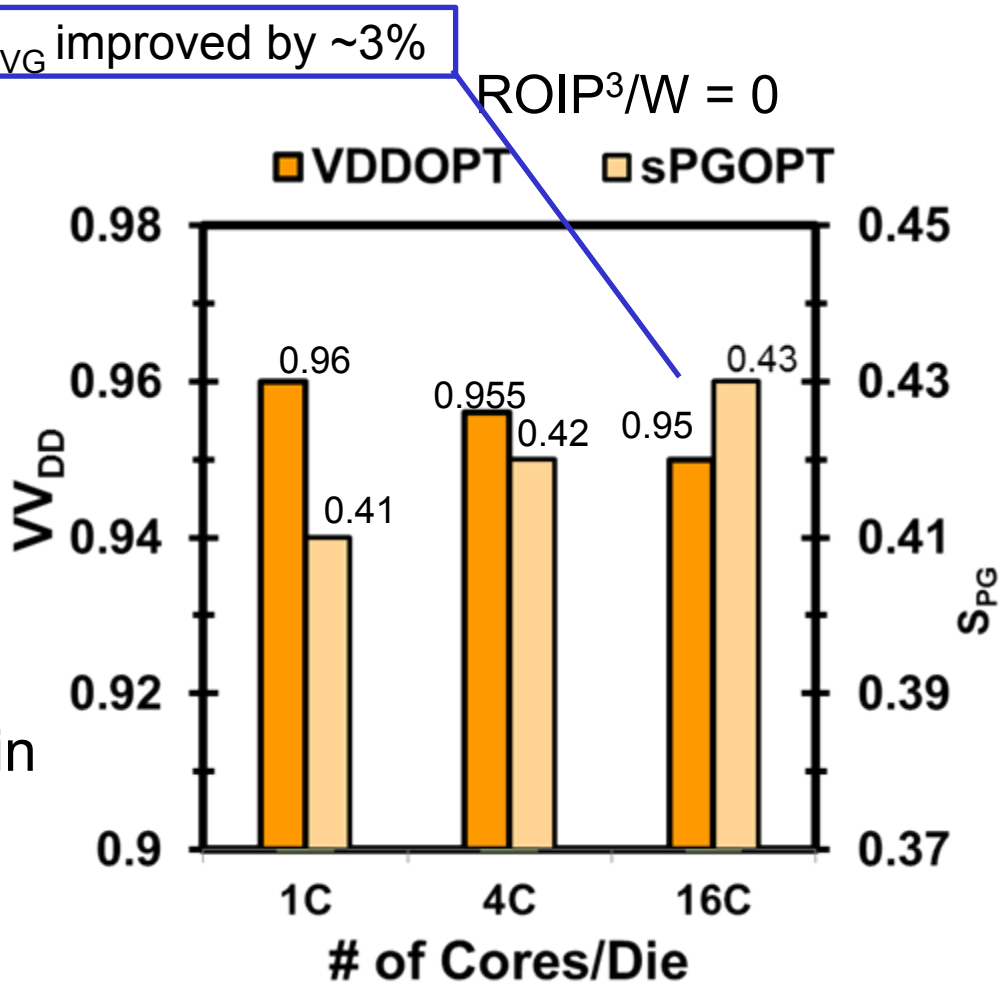

Use $F_{MAX,AVG}$ to compute P^3/W constraint

Outline

- Introduction
- Impact of process variations and PG-device size
  - Impact on VV_{DD}
  - Impact on F_{MAX} and P_{TOT}
- AVS-aware PG-device size optimization
  - Algorithm and simulation result
- Impact of WID variation on PG sizing
  - Global clocking
  - Frequency island clocking
- **Experimental Methodology**
- Conclusion

Experimental Methodology

- For frequency and leakage modeling with power gates
  - V_{th} and L_{eff} WID spatial and D2D variation map*
    - WID variation : Correlation coefficient = 0.5
    - D2D variation : $\sigma_{V_{th}}^{sys} = 6.4\%$

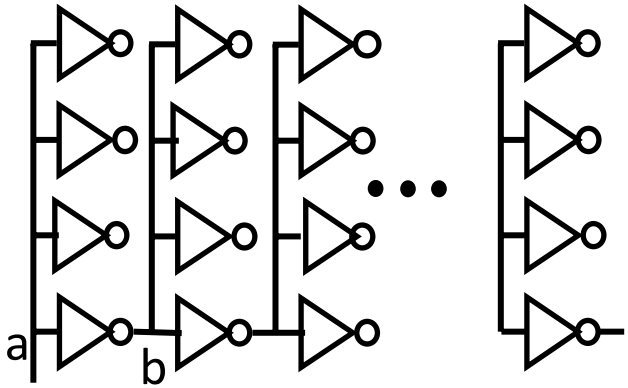

24 FO4 INV chain for measuring $f(V_{DD})$
32nm PTM SPICE model

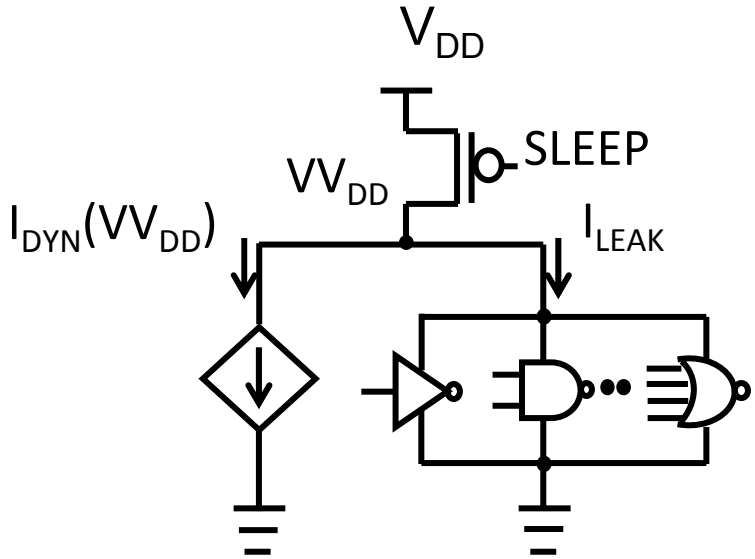

Dummy gates for measuring $I(V_{DD})$

Effective Widths

50% INV, 30% NAND,20% NOR

*S. Sarangi et al. IEEE Trans. On Semiconductor Manufacturing, vol. 21, no. 1, pp. 3~13, Feb. 2008.

Experimental Methodology

- Power and thermal constraint
  - P_{TDP} at $V_{DD,TDP}$ = 90 W (at the nominal corner)
  - T_{jmax} = 100 ° C
- Performance simulation with GPGPU-Sim
  - Simulator modified to support FI clocking

# of SM Cores	4/8/16	Shared Mem/SM	16KB
SIMD Width/SM	1/4/8	# of Mem Ch.	4
# of Threads /SM	1024	BW/ Mem Ch.	8B/Cycle
3 of CTAs/SM	8	DRAM Rq. Queue	16
# of Registers/SM	16384	Mem Controller	FR-FCFS
Constant and Texture Cache Sizes	8KB, 2-Way, 64B Line	GDDR3 Mem# of SM Cores . tCL/tRP/tRAS	10/10/35/25

Conclusion

- Effect of PG-device sizing on F_{MAX} and P_{TOT}
  - F_{MAX} and P_{TOT} both increase, P_{TOT} increases faster than $F_{MAX.}$
  - Rate of increase diminishes quickly for slow die than for fast die
- Reduction in PG-device size
  - D2D variation :
    - 24.5%(slow), 36%(nominal) and 43.2%(fast)
  - WID variation
    - Global Clk. **59%**
    - FI Clk. **58%** (4 cores), **57%** (16 cores)
- F_{MAX} penalty
  - D2D variation : negligible
  - WID variation : improved F_{MAX} by ~3%
- As #. of cores increases
  - Opt. PG size increases
  - Opt. V_{DD} for AVS decreases