

ESE534: Computer Organization

Day 16: March 24, 2010
Component-Specific Mapping



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Previously


- Increasing variation with scaling (Day 8)
- Row sparing in memories (Day 12, HW6)
- Disk-drive model for mapping out defects in software (Day 12)
- PLAs for logic (last time)
- LUTs + Interconnect (Day 14)

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

2

Today

- Defect and Variation Tolerance for Compute and Interconnect
- High defect rates and extreme variation
- Nanoscale chips are like snowflakes
→ Each one is unique.
- Application mapping tailored to the component



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

3

Agenda

- nanoPLA
 - Introduction
 - Tolerating crosspoint defects
 - Tolerating extreme V_{th} variation
- FPGA
 - Tolerating defects
 - Lightweight, load-time defect avoidance

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

4


nanoPLA

Architecture

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

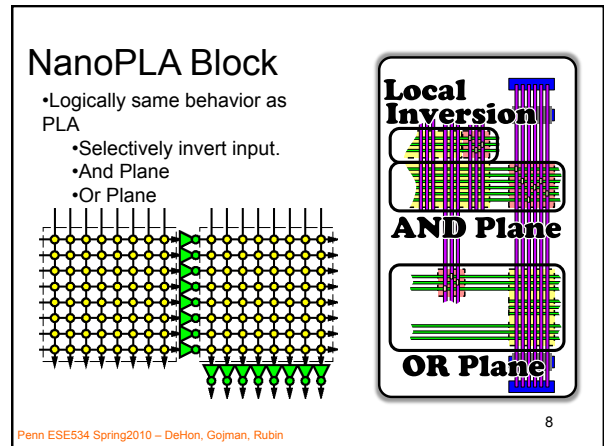
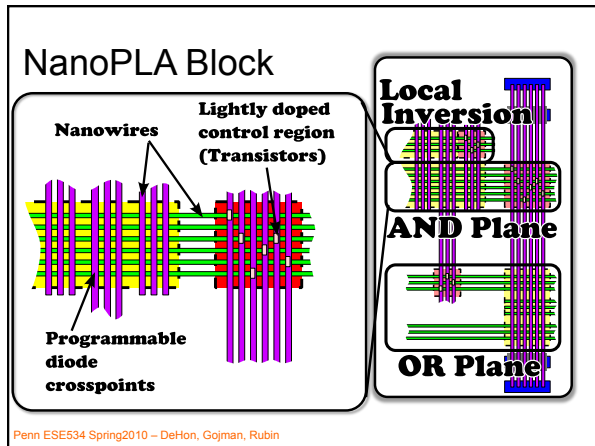
5

NanoPLA Plane

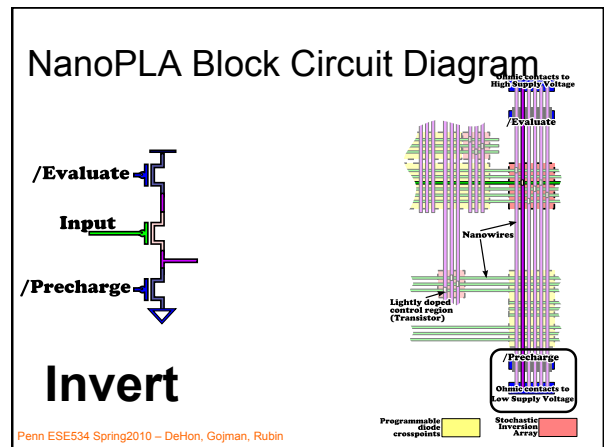
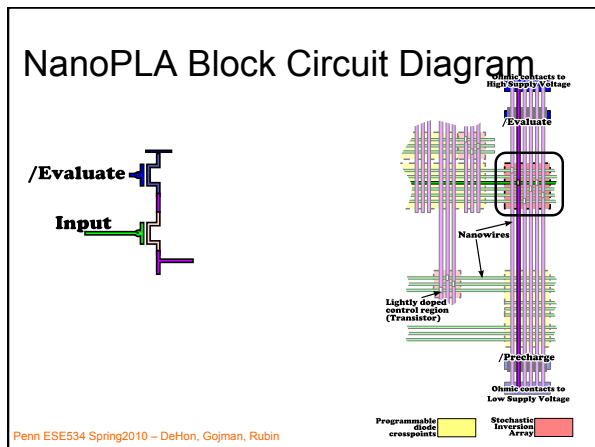
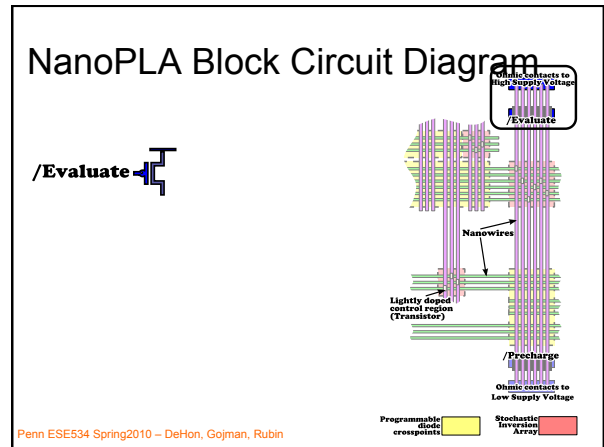
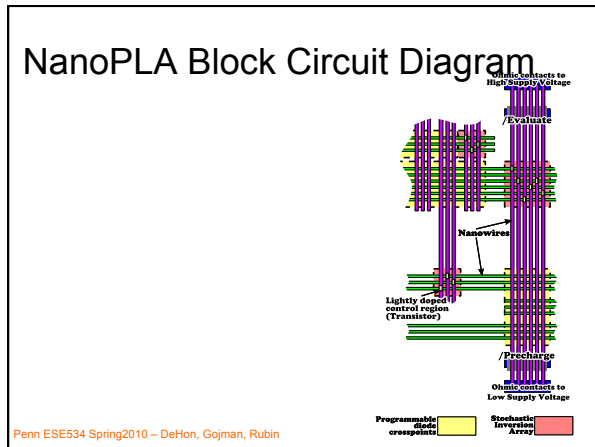


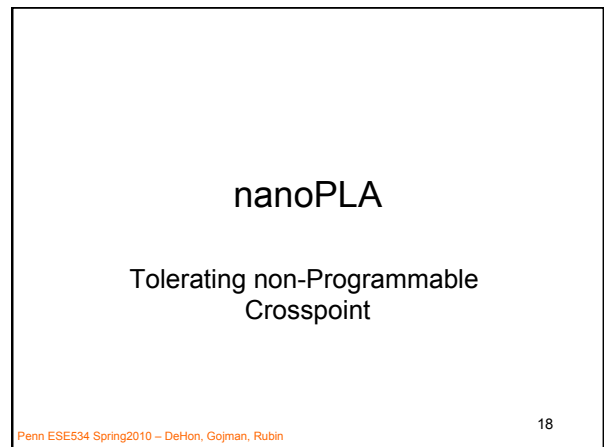
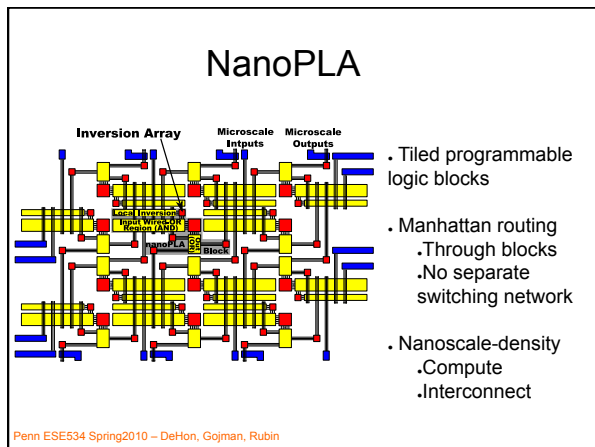
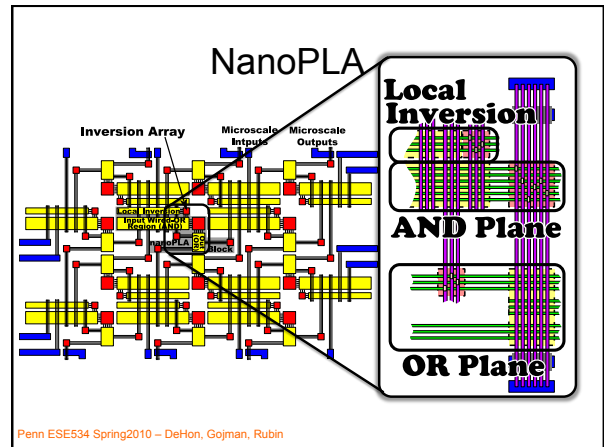
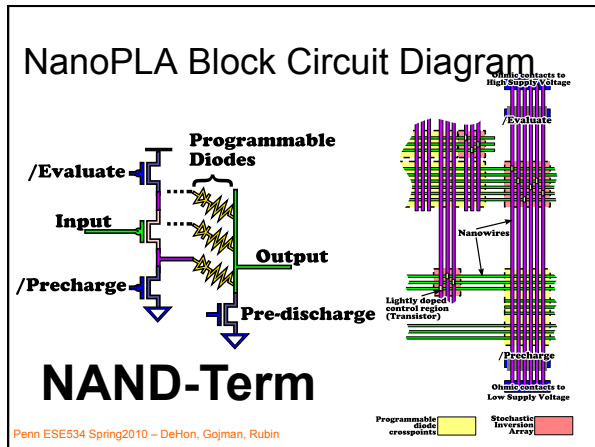
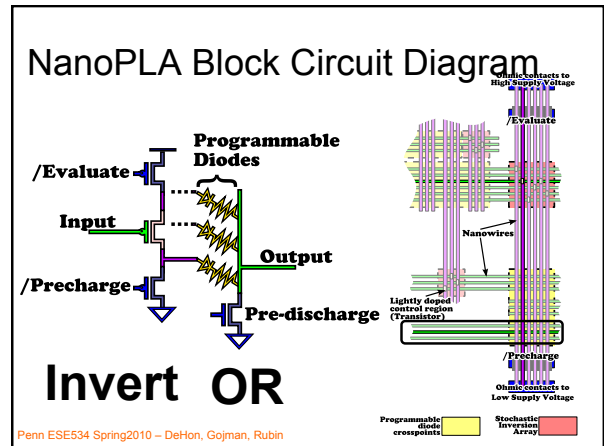
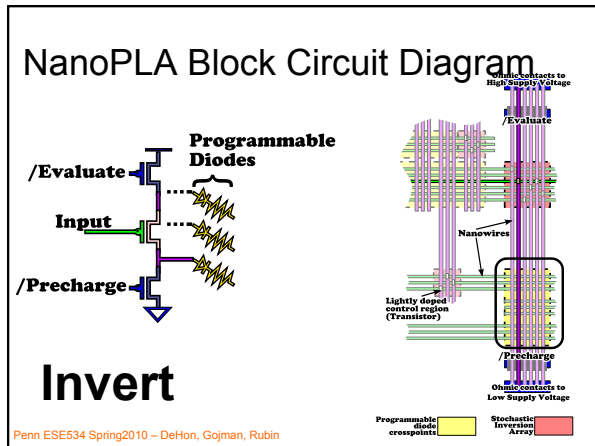
- Smallest unit of computation
- Built Bottom-Up from nanowires
- Diode-Programmable region
 - Memory cell area == Wire crossing area
 - Molecular
 - Amorphous Si
- Stochastic FET restore region
 - Core-shell nanowires with doped region

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin




8





Crosspoint Defects


- Crosspoint junctions may be nonprogrammable
 - E.g. the first 8x8 from Hewlett-Packard had 85% programmable crosspoints
- Preclass:
 - 100 junctions
 - 5% stuck-open rate
 - $P_{\text{xpoint}}=0.95$
 - $P_{\text{row}}?$



Penn ESE534 Spring2010 – DeHon, Gojman, Rubin

Crosspoint Defects


- Using P_{row} just calculated.
- How many perfect wires can we expect (on average) in an array of 100 wires?
- How will row sparing work here?
- What kind of overhead would row sparing require?



Penn ESE534 Spring2010 – DeHon, Gojman, Rubin

Idea 1: Don't need all junctions


- $F=a+b+e$
- Can use wire w_4
- Even though has missing crosspoints
- Preclass 2:
 - F require 3 good crosspoints
 - $P_{\text{good}}?$



Penn ESE534 Spring2010 – DeHon, Gojman, Rubin

Idea 2: Have many wires to choose from




- $F=a+b+e$
- Cannot use w_2
- Can use any of the rest
- Probability of at least one match is very high
- Preclass 3
 - Use P_{good} from Preclass 2
 - $P(1 \text{ of } 5)=?$



Penn ESE534 Spring2010 – DeHon, Gojman, Rubin

Mapping Entire Array

- Need to map all 100 functions
- Assuming $P_{1\text{-of-}5}$ from previous, what's the probability can find a match for all 100?
- In practice have 100 to choose from not 5
- What can we do if cannot find a match?





Penn ESE534 Spring2010 – DeHon, Gojman, Rubin

23

Crosspoint Defects

- Tolerate by **matching** nanowire junction programmability with pterm needs



Naeimi/DeHon, FPT2004

Design and Test of Computers, July-August 2005

Penn ESE534 Spring2010 – DeHon, Gojman, Rubin

Crosspoint Defect Rate Impact

- Toronto20 benchmark circuits
 - Typical 10k “gate” logic designs

Design	P_j			
	0.85	0.9	0.95	1
alu4	1.81	1.64	1.00	1.00
apex2	1.19	1.19	1.00	1.00
apex4	1.30	1.16	1.00	1.00
bigkey	1.00	1.00	1.00	1.00
clma	1.00	1.00	1.00	1.00
des	1.00	1.00	1.00	1.00
dsip	1.00	1.00	1.00	1.00
elliptic	1.00	1.00	1.00	1.00
ex1010	3.81	2.15	1.00	1.00
ex5p	1.00	1.00	1.00	1.00
frisc	1.00	1.00	1.00	1.00
misex3	1.31	1.31	1.00	1.00
pdc	4.75	1.79	1.00	1.00
s298	1.84	1.84	1.00	1.00
seq	1.20	1.12	1.00	1.00
spla	3.46	1.83	1.00	1.00

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Matching Lesson

- Can use almost all wires
 - @ 5% defects, need **no** extra wires
- ...despite the fact that almost all wires are imperfect
 - P_{row} from preclass 1
- Trick is to figure out where an imperfect wire is “good enough”

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

26

nanoPLA


Tolerating Extreme V_{th} Variation

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

27

Challenge


- High variation of V_{th} means some devices leak faster than others evaluate!



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

28

Preclass 4



- Sampling nanowires from a distribution
- What range of nanowire V_{th} 's must we tolerate?
- 4c: In array of 100 nanowires, what is the probability that there is one nanowire $n\sigma$ above mean and one $n\sigma$ below?
 - $\sigma = 1$?
 - $\sigma = 2$?
 - $\sigma = 3$?

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

29

Preclass


- Problem 4d: How many σ out should we expect in a collection of 100 arrays?
- Problem 5:
 - $I(V_{on}, V_{th}=V_{th}+n\sigma) = ?$
 - $I(V_{off}, V_{th}=V_{th}-n\sigma) = ?$
- What is σ/V_{th} given in example?

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

30

Challenge


- High variation of V_{th} means some devices leak faster than others evaluate!



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

31


NanoPLA Physical Variation



- Predominantly Random Variation
- Due to Bottom-up manufacturing process
- Variation in:
 - Dopant fluctuations
 - NanoWire alignment
 - NanoWire dimensions
 - Number of state-holding elements / junction
- Variation in V_{th} , R and C
 - Independent Gaussian distributions


Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

NanoPLA Defect Model




Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Defect Model



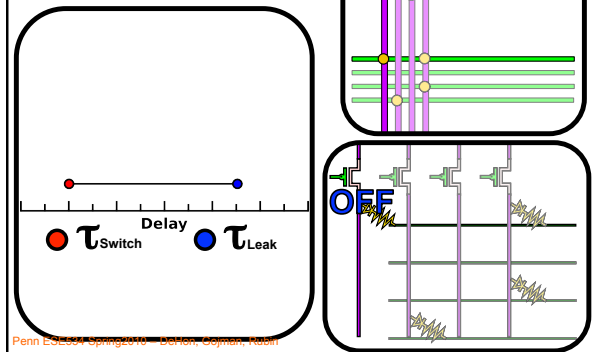
Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Defect Model

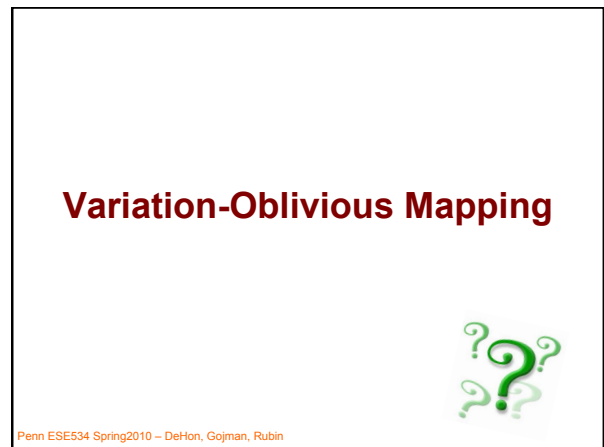
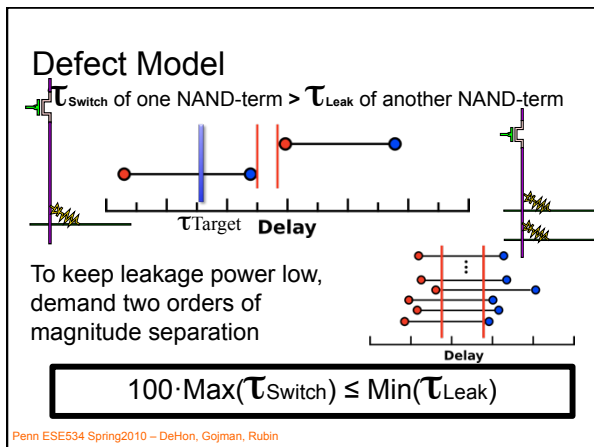
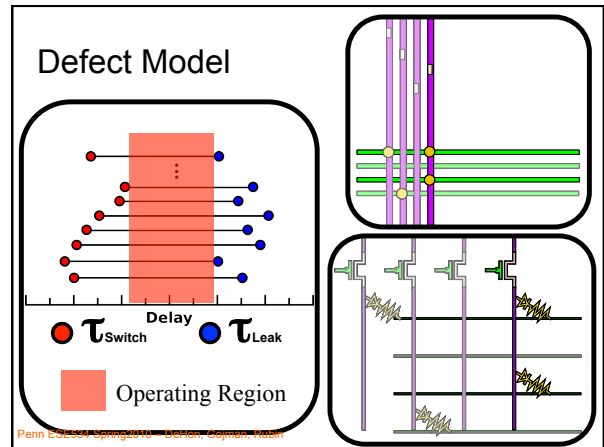
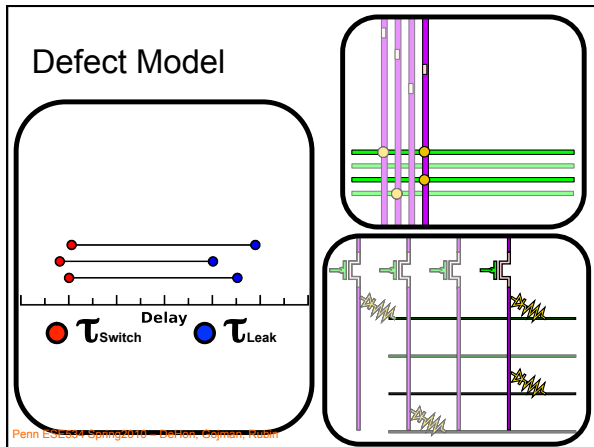
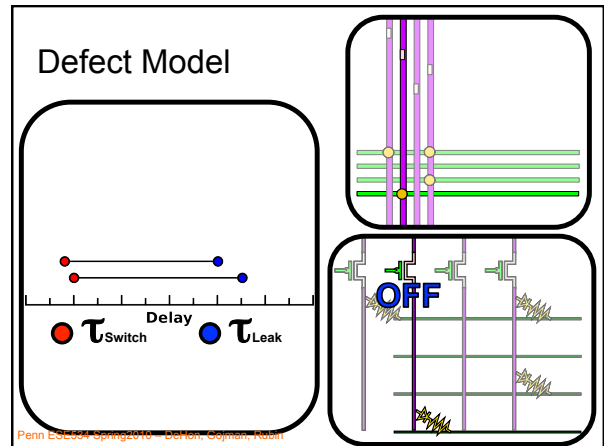
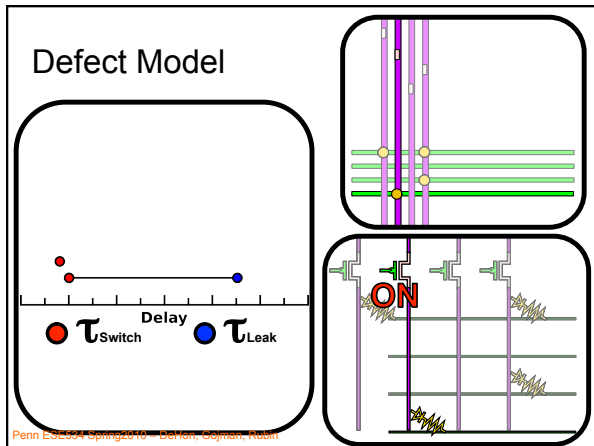


Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Defect Model




Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin



Variation-Oblivious Mapping


Perform Detailed Mapping Oblivious to the Characteristics of the underlying Resources.



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Variation-Oblivious Mapping


Perform Detailed Mapping Oblivious to the Characteristics of the underlying Resources.



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Variation-Oblivious Mapping


Perform Detailed Mapping Oblivious to the Characteristics of the underlying Resources.



Defective Mapping: $100 \cdot \text{Max}(\tau_{\text{Switch}}) \leq \text{Min}(\tau_{\text{Leak}})$ **not met**

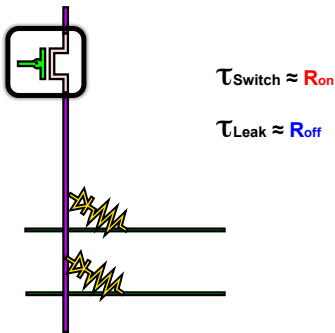
Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Timing Model



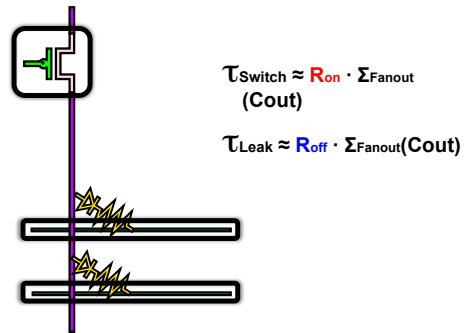
Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Timing Model









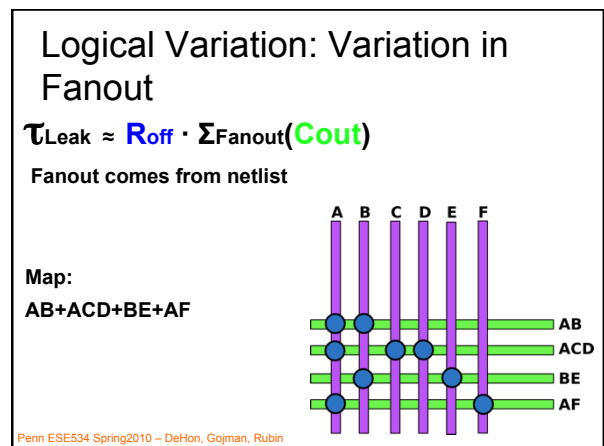
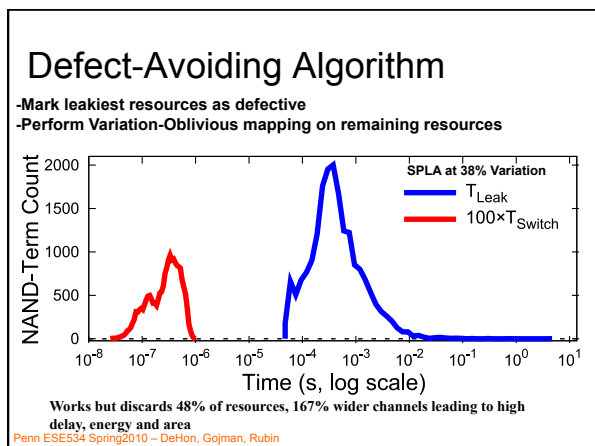
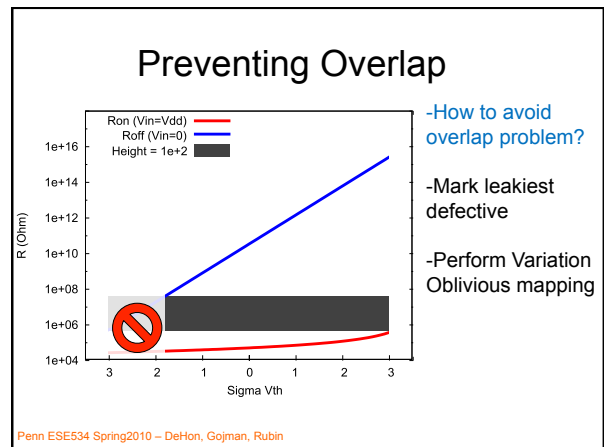



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Timing Model



Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin





Logical Variation: Variation in Fanout


$\tau_{Leak} \approx R_{off} \cdot \Sigma_{Fanout}(C_{out})$

A has fanout 3:
AB, ACD, AF

F has fanout 1:
AF

Even at nominal, A sees 3 times more load than F

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin



Defect-Avoiding Toy Example

Step 1 Discard Leaky Resources

Logical Fanout	R _{off} Resistance	Estimated τ_{Leak}
1	320	
10	20	
15	150	
1	300	
3	100	
	30	

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Defect-Avoiding Toy Example

Step 2 Match Functions to Resources


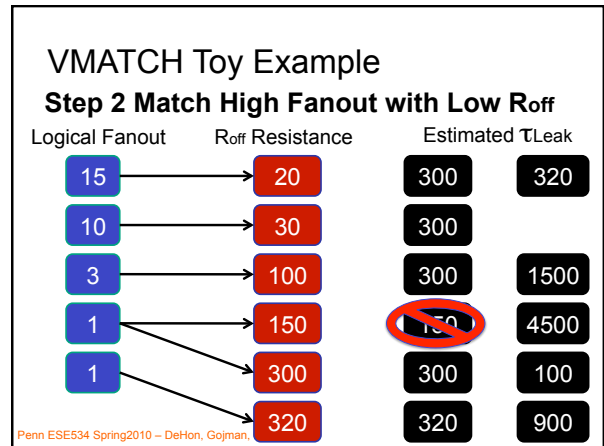
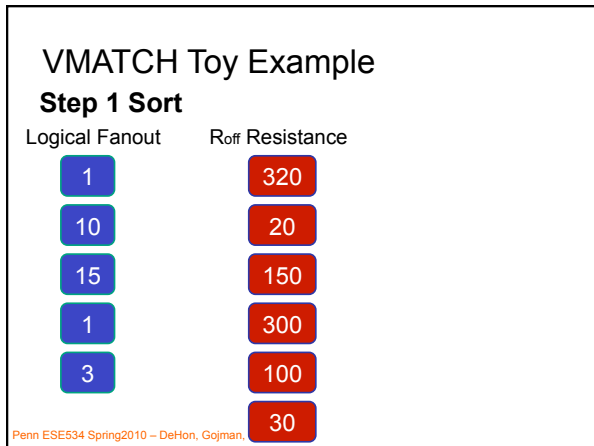
Logical Fanout	R _{off} Resistance	Estimated τ_{Leak}
1	320	320
10	20	
15	150	1500
1	300	4500
3	100	100
	30	900

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

VMATCH

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin



- ## VMATCH
- Logical Variation + Physical Variation >> Variation
 - Variation-Aware Mapping Algorithm
 - Uses logical variation (fanout variation) to counteract physical variation (R_{off} variation)
 - Matches high-fanout term with low R_{off} NAND-Term and vice versa.
- Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin








Experiments

- Simulated NanoPLA using 5nm pitch wires, 5nm transistor channel length
- Modeled variation up to ITRS predicted 38% for technology
- Implemented Variation-Oblivious, Defect-Avoiding and VMATCH Algorithms on NPR
- Routed Toronto 20 Benchmark on 100 Monte Carlo generated chips.

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin





Matching Lesson 2

- Can keep yield high
 - @ 38% variation, small amount of extra wires
- ...despite the fact that Roff varies over 12 orders of magnitude
- Trick is to match the logical variation to the physical variation to reduce both

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubín

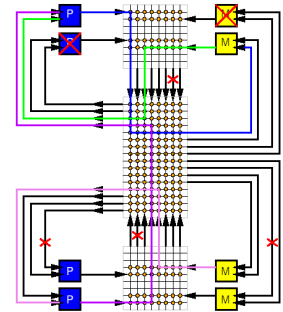
78

FPGA Defects

Full Knowledge

Problem


- At high defect rates, some resources will be unusable
- Resources
 - LUTs
 - FFs
 - Interconnection links
- Full crossbars are too expensive



Opportunity


- Avoid defective elements during mapping
 - Don't place logic in defective LUT
 - Like disk-drive model where we don't allocate a defective sector to hold data
 - Route to avoid defective interconnect

Avoid Defective LUT



[Lach et al. / FPGA 1998]

Avoid Defective Interconnect



Example

- Collect 4 numbers between 1 and 81

Defects on Mesh

- Draw on board for mapping on next slide.

85

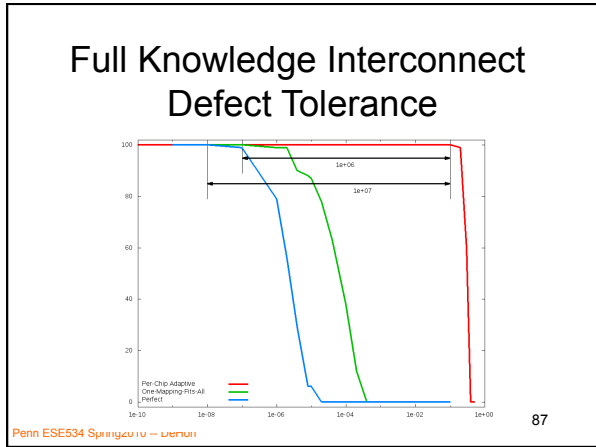
Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

Class Exercise

- Map this circuit to the defective LUT mesh, avoiding defects.

86

Penn ESE534 Spring2010 -- DeHon, Gojman, Rubin

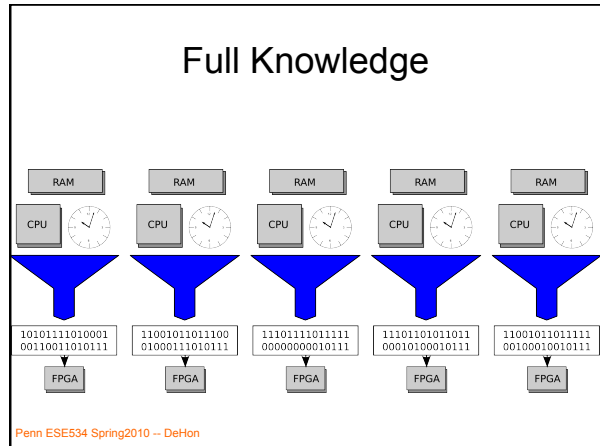
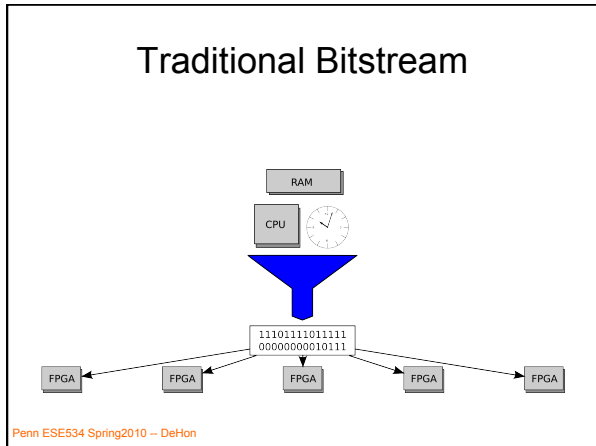


FPGA Routing Defects

Pre-computed Alternatives
Lightweight Defect Tolerance
Choose Your own Adventure

88

Penn ESE534 Spring2010 -- DeHon



Choose Your own Adventure

- Bitstream with alternatives
 - Still just one bitstream

Penn ESE534 Spring2010 -- DeHon

Choose Your own Adventure

- Bitstream with alternatives
 - Still just one

Penn ESE534 Spring2010 -- DeHon

Loading

Examine path

- Configure
- Test
 - Good: skip to next net
 - Bad: grab next path

Penn ESE534 Spring2010 -- DeHon

Loading

Examine path

- **Configure**
- Test
 - Good: skip to next net
 - Bad: grab next path

Penn ESE534 Spring2010 -- DeHon

Loading

Examine path

- Configure
- **Test**
 - Good: skip to next net
 - Bad: grab next path

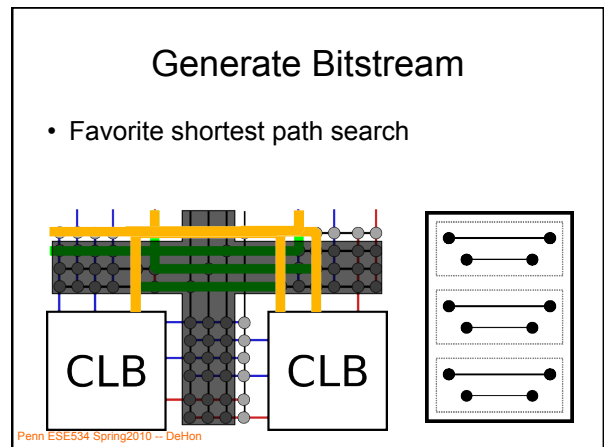
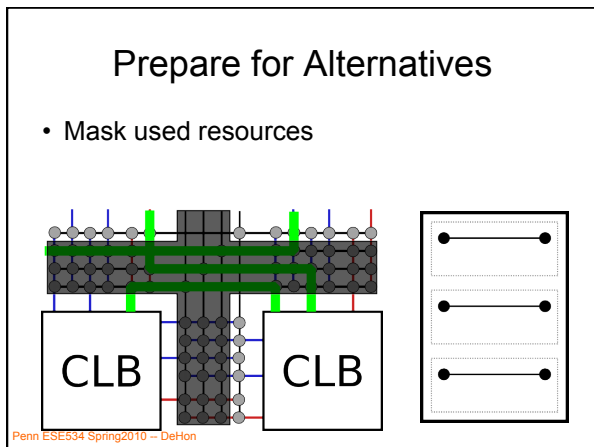
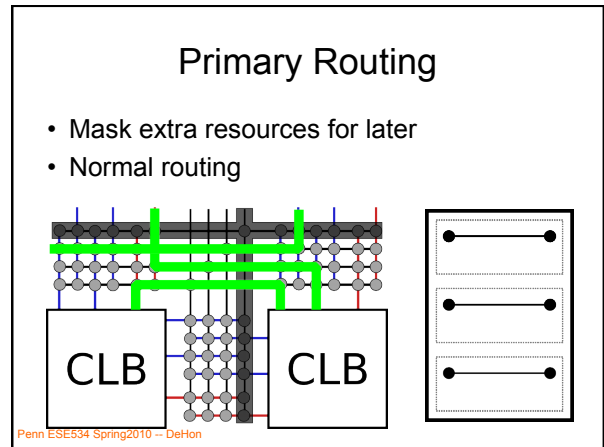
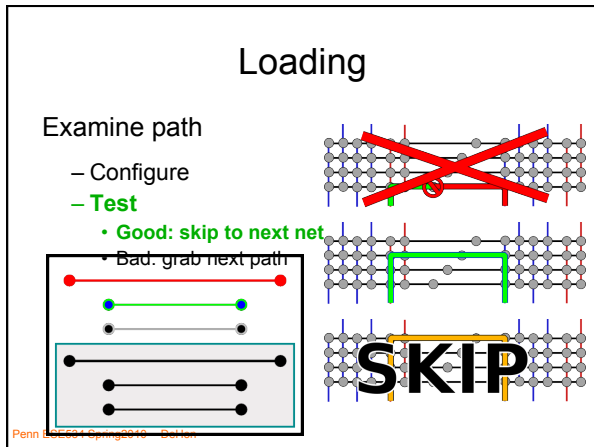
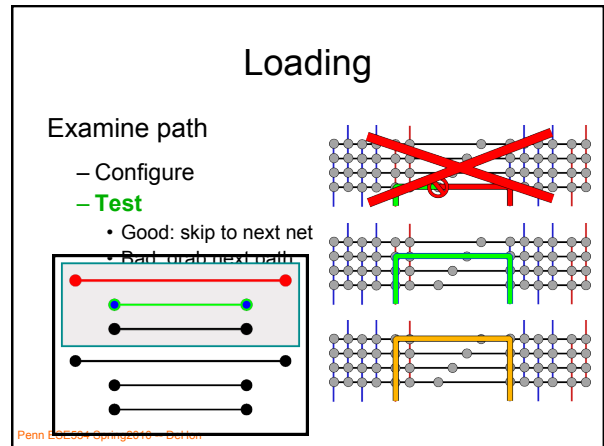
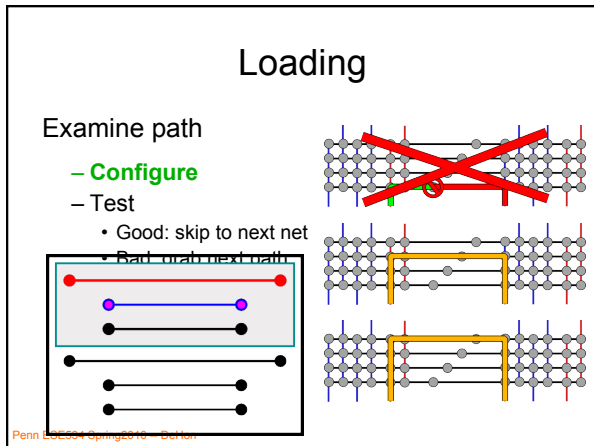
Penn ESE534 Spring2010 -- DeHon


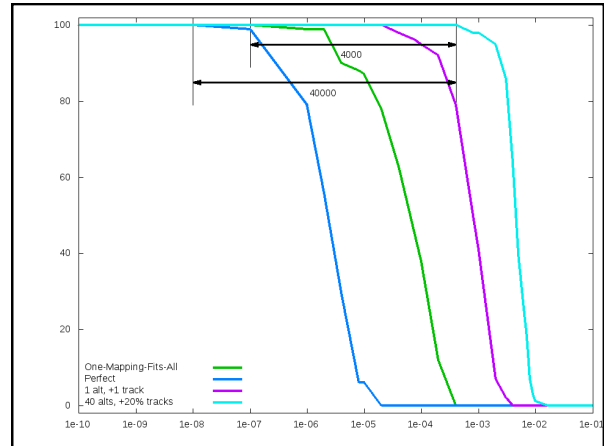
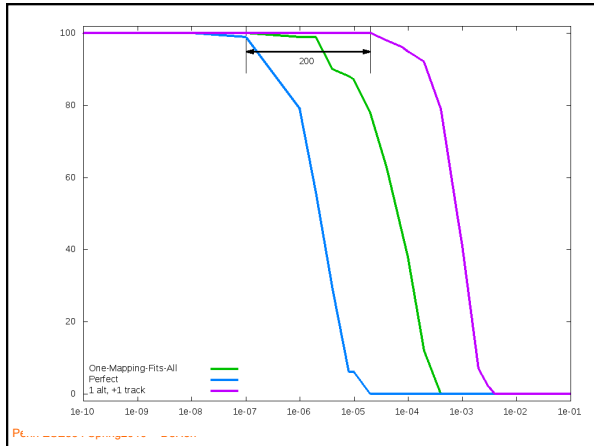
Loading

Examine path

- Configure
- **Test**
 - Good: skip to next net
 - Bad: grab next path

Penn ESE534 Spring2010 -- DeHon





Admin

- Homework 7 due Monday
- No office hours today (André away)
- Reading for Monday on web
 - Classic paper on Rent's Rule

106

Big Ideas [MSB Ideas]

- Nanoscale Integrated Circuits are like snowflakes – each unique
- Cannot expect perfect ICs
 - ...or even perfect building blocks
- But even defective components are often “good enough”
- Post-fabrication matching allows us to assign functions based on capabilities

107

Big Ideas [MSB-1 Ideas]

- Matching in nanoPLA
 - 5% crosspoint non-programmable defects
 - $\sigma=38\%$ variation in V_{th}
- Defect-aware routing in FPGAs
 - 10% switch and segment defects with full knowledge
 - 1% with pre-computed alternatives

108