Predicting Black Swans for Fun and Profit

High-Dimensional Estimation of Low Failure Rates for Memory Design

Trent McConaghy
Co-founder & CTO, Solido Design Automation Inc.
www.solidodesign.com
www.trent.st

Berlin, February 2014


2008 Market Crash
BROOKLYN
BATTERY
TUNNEL
LOCAL
High Impact, Low probability
= Black Swan


NEW YORK TIMES BESTSELLER
THE
BLACK SWAN
The Impact of the
HIGHLY IMPROBABLE
Nassim Nicholas Taleb

(NNT)

Sony PS3 Black Swan

Sony PlayStation … Phenomenal Production Yield.

COMMENTS (13)

Sony Hopes to Fulfill Demand for PlayStation 4 This Year

[09/16/2013 11:10 PM]
by Anton Shilov

When Sony Corp. launched its PlayStation 3 game console seven years ago, the firm face a number of problems, one of which was high price of the system and another one was poor yield of Cell processor, the heart of the PS3. With PlayStation 4 the company made a lot of decisions differently and now it hopes that the same issues that harmed success of the PS3 will not hurt the PS4.

"It is everything we can manufacture. We want to make sure that consumers have an opportunity to buy one on November 15 and through t… holding back some inventory to make sure that people have an opportunit… on launch … manufacture, and the … production yields have been phenomenal. So this will be by a magnitude of a lot the biggest launch we have ev… America, in an interview with Fox Business, reports Games Industry web-site.

High Impact (for Sony),
Low probability *(at least Sony thought so)*

Low predictability is a common Black Swan characteristic


Q: How to handle Black Swans?
NNT: (Since hard to predict) don't try to predict.
Protect from downside, gain upside

Money,
And Moore's Law

1G phones a year

56% smartphones (↑)

- $150G now, $300G in 2015

Chips Power Smartphones

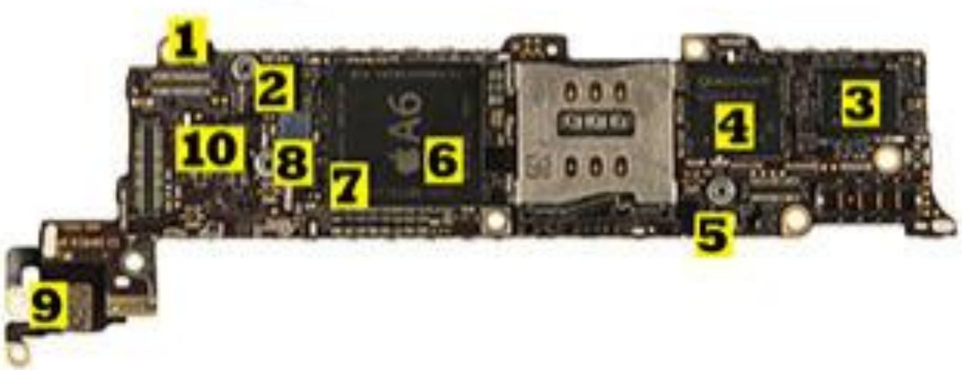

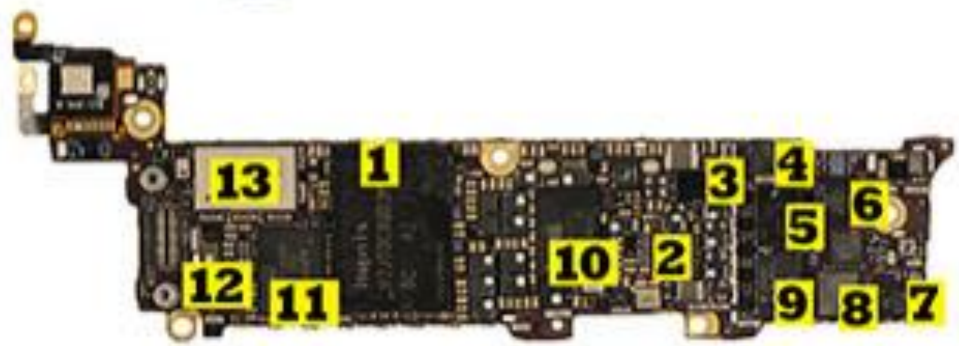

Chips Are Everywhere (not just Smartphones)

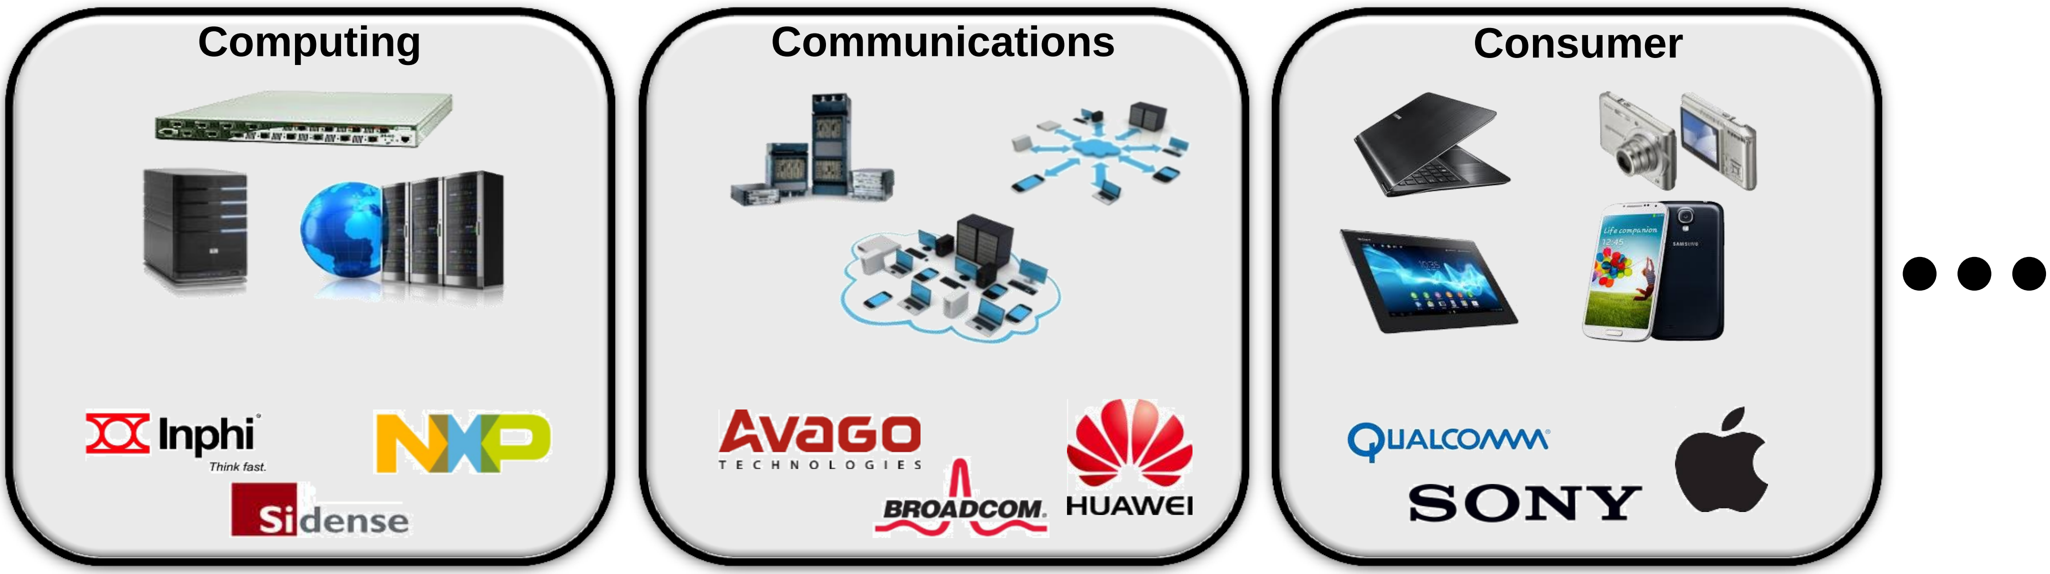

This is the $300G / year Semiconductor Industry.

But smartphones *are* the biggest driver for chips' speed, power, memory.

Moore's Law: Shrinking Transistors

For faster, lower power, and better-packed transistors
Therefore faster, lower power, and higher-memory *smartphones*

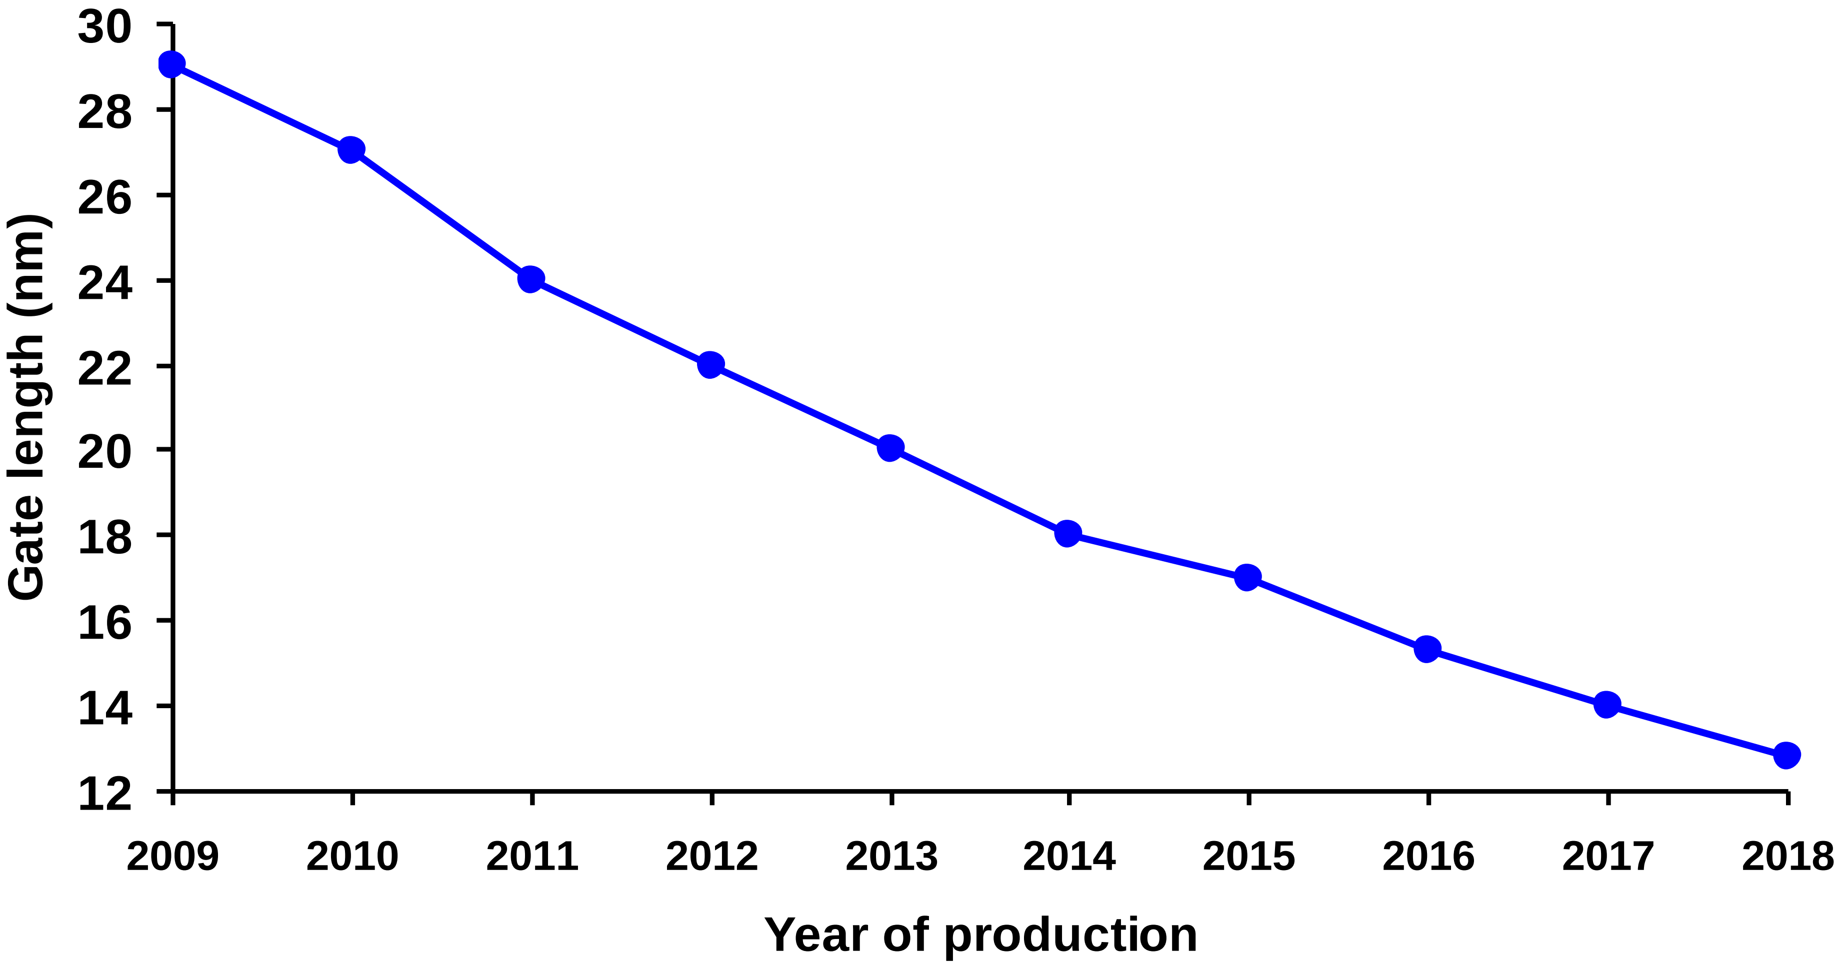

[International Technology Roadmap for Semiconductors, 2011]

Transistors are shrinking
But atoms aren't!

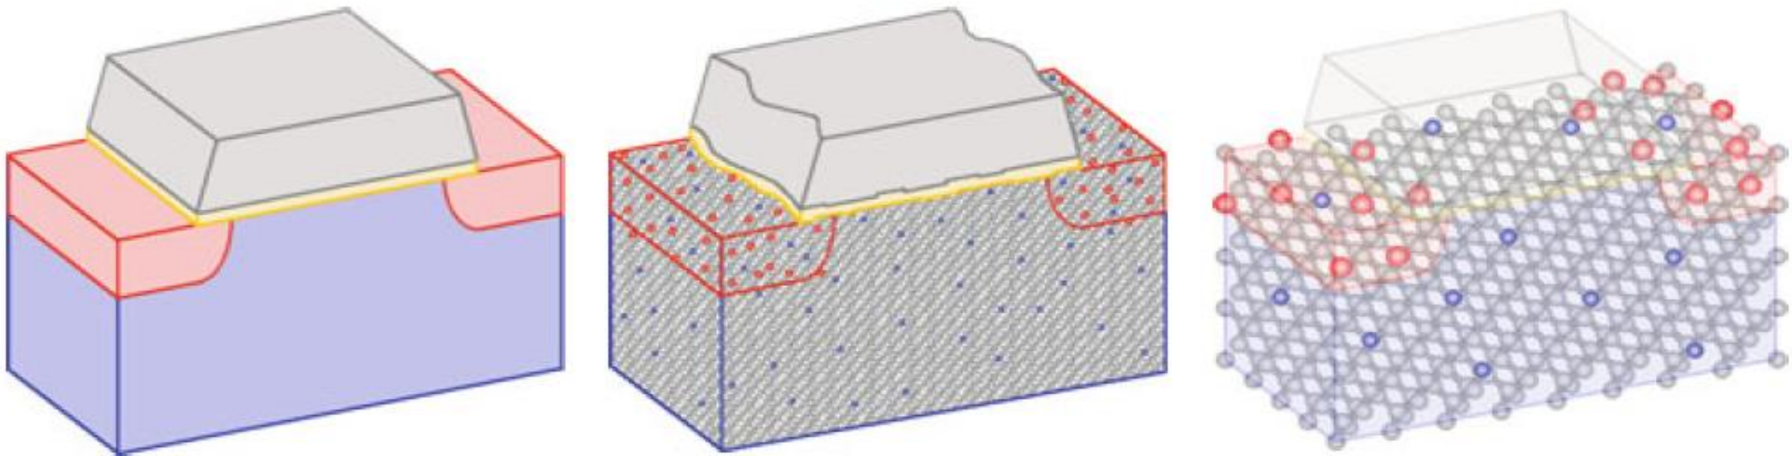

traditional 22nm sub 10 nm (2020)

A. Asenov, "Statistical Nano CMOS Variability and Its Impact on SRAM", Chapter 3, A. Singhee and R. Rutenbar, Eds., Extreme Statistics in Nanoscale Memory Design, Springer, 2010

Variation = atoms out of place

…Propagating from devices, to circuit, to yield

Process variation ↑

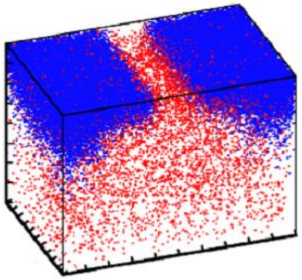

Circuit performance variation ↑

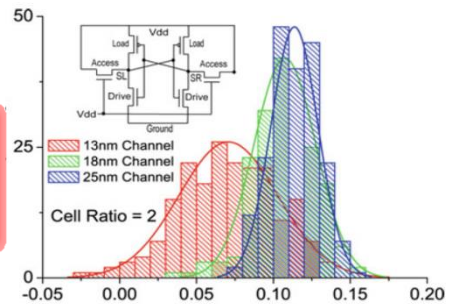

Circuit Yield ↓

Variation = atoms out of place

…Propagating from devices, to circuit, to yield… to finances

Process variation ↑

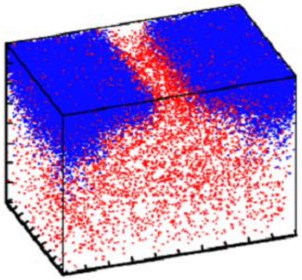

Circuit performance variation ↑

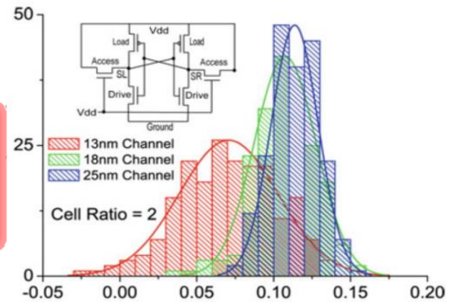

Circuit Yield ↓

Company Finances ↓

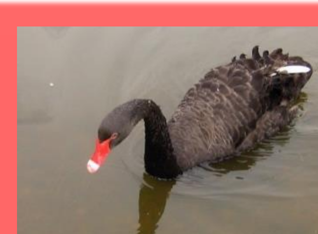

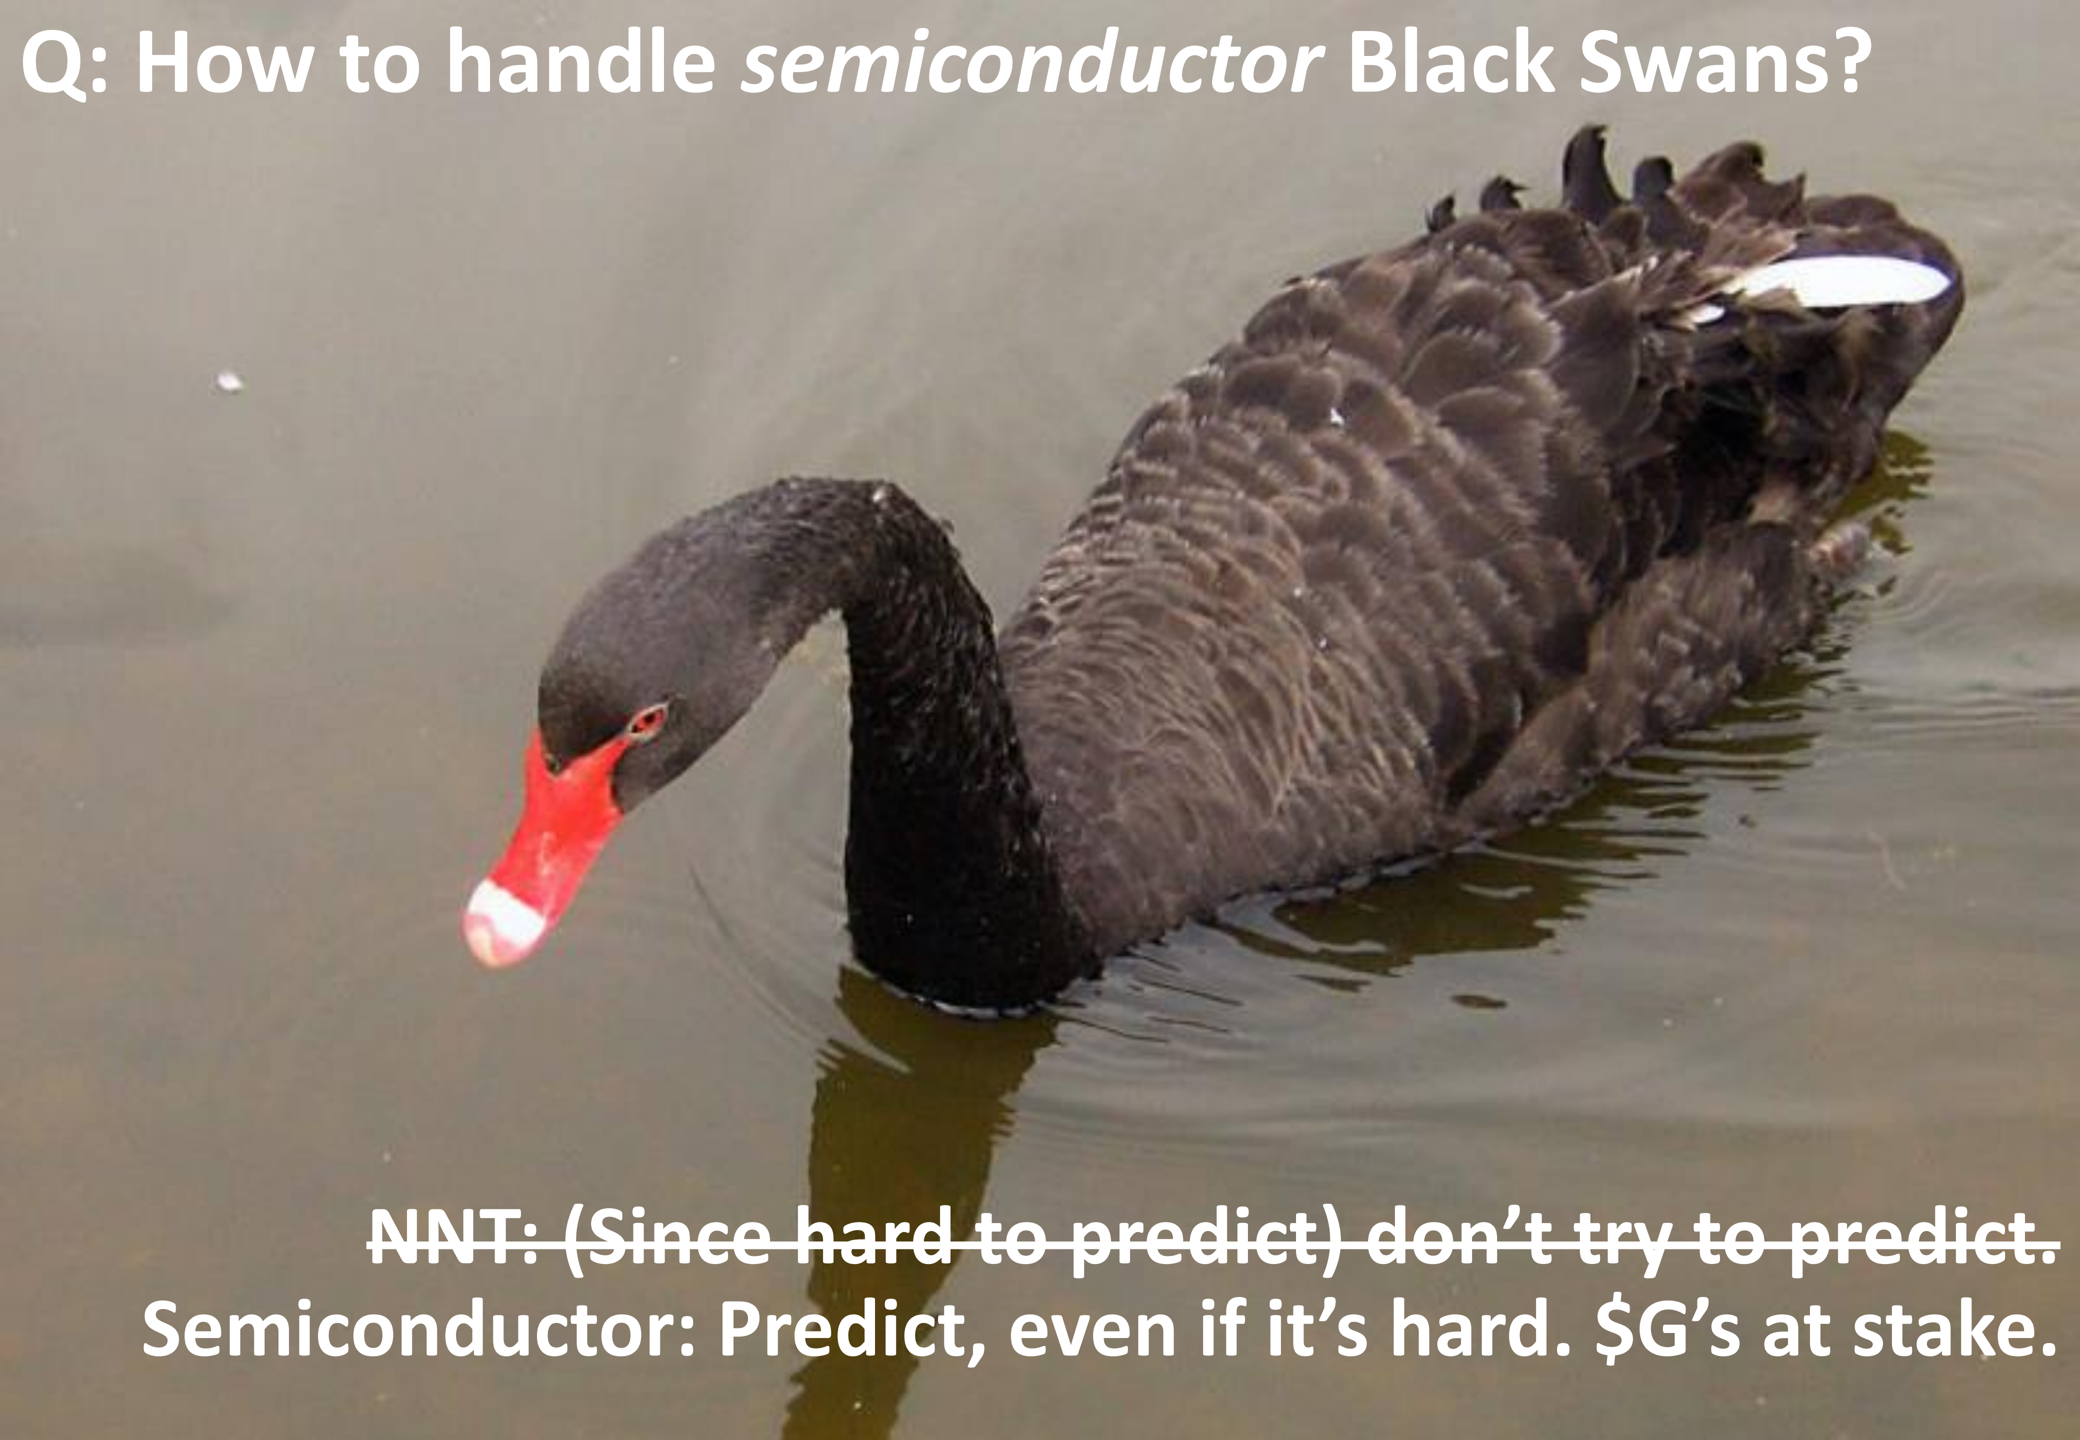
Q: How to handle *semiconductor* Black Swans?
~~NNT: (Since hard to predict) don't try to predict.~~
Semiconductor: Predict, even if it's hard. $G's at stake.

Yield Prediction

The Yield Prediction Problem

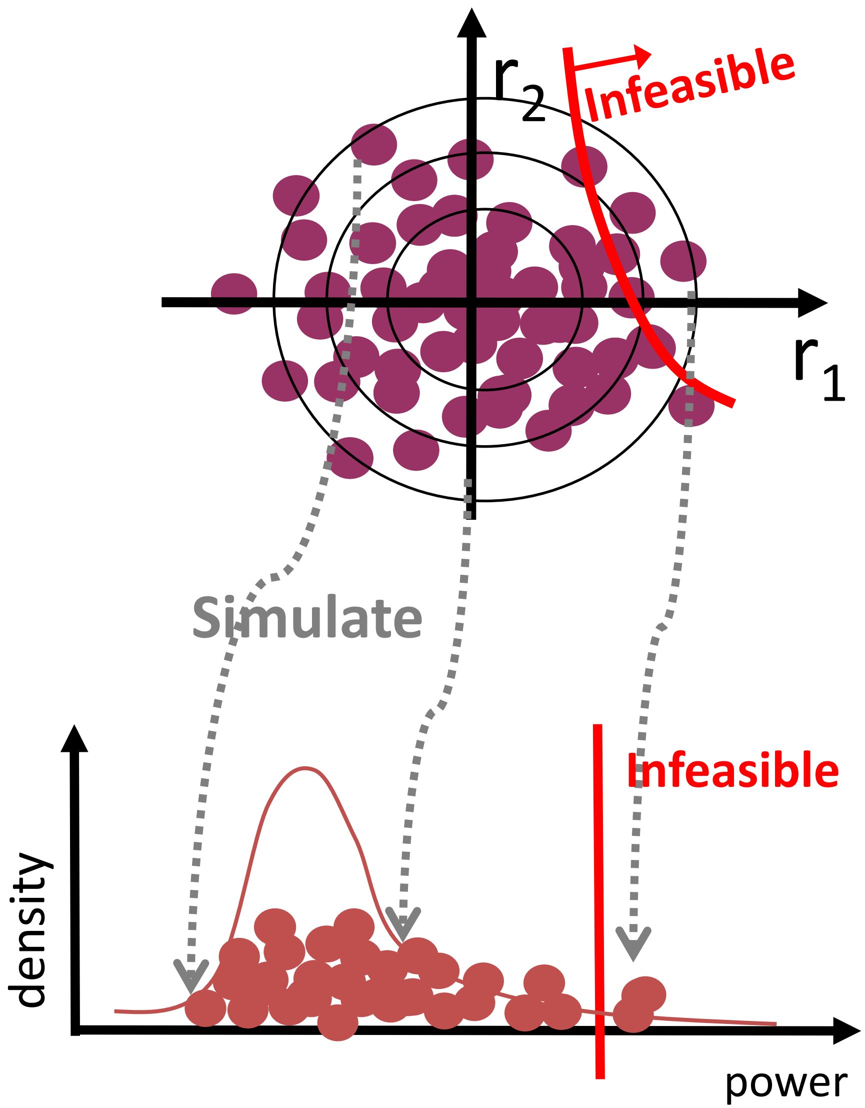

We know:

- The distribution for input variables. E.g. r1,r2 are NIID
- Given an input point, a way to estimate an output like power. E.g. simulation
- Constraint for output. E.g. power < 1 mW

We want to know:

- Yield = volume under PDF that meets output constraint
- Alternatively, Pfail = 1 – yield

Example approach: Monte Carlo (shown)

- Draw N points in input space
- Simulate each
- Est. yield = (# meet constraint) / N

High-Sigma
Yield Prediction

Memory Bitcell Failures are Highly Improbable (1/1G)
How can we estimate yield? Try Monte Carlo…

10K MC samples:

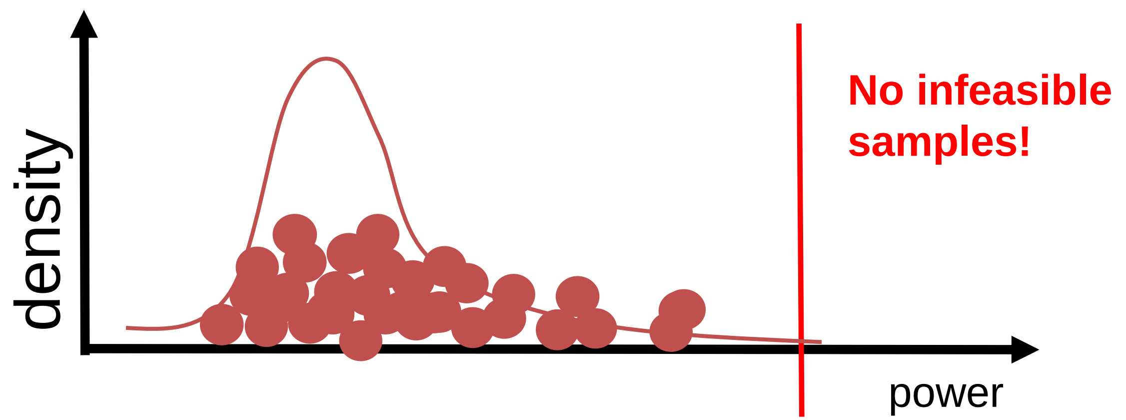

5G MC samples:

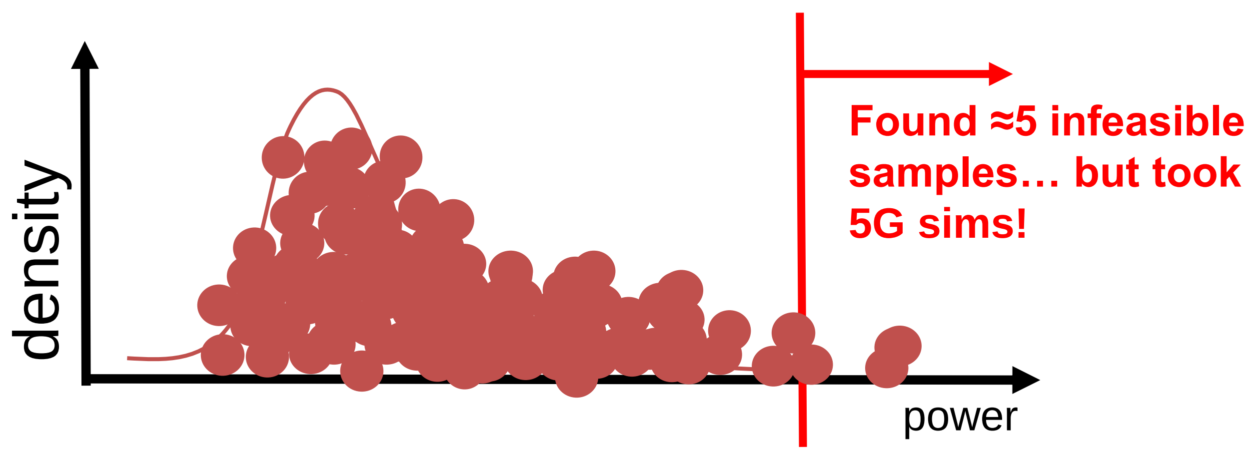

Let's Survey Approaches…

Approach: 10K Monte Carlo Samples

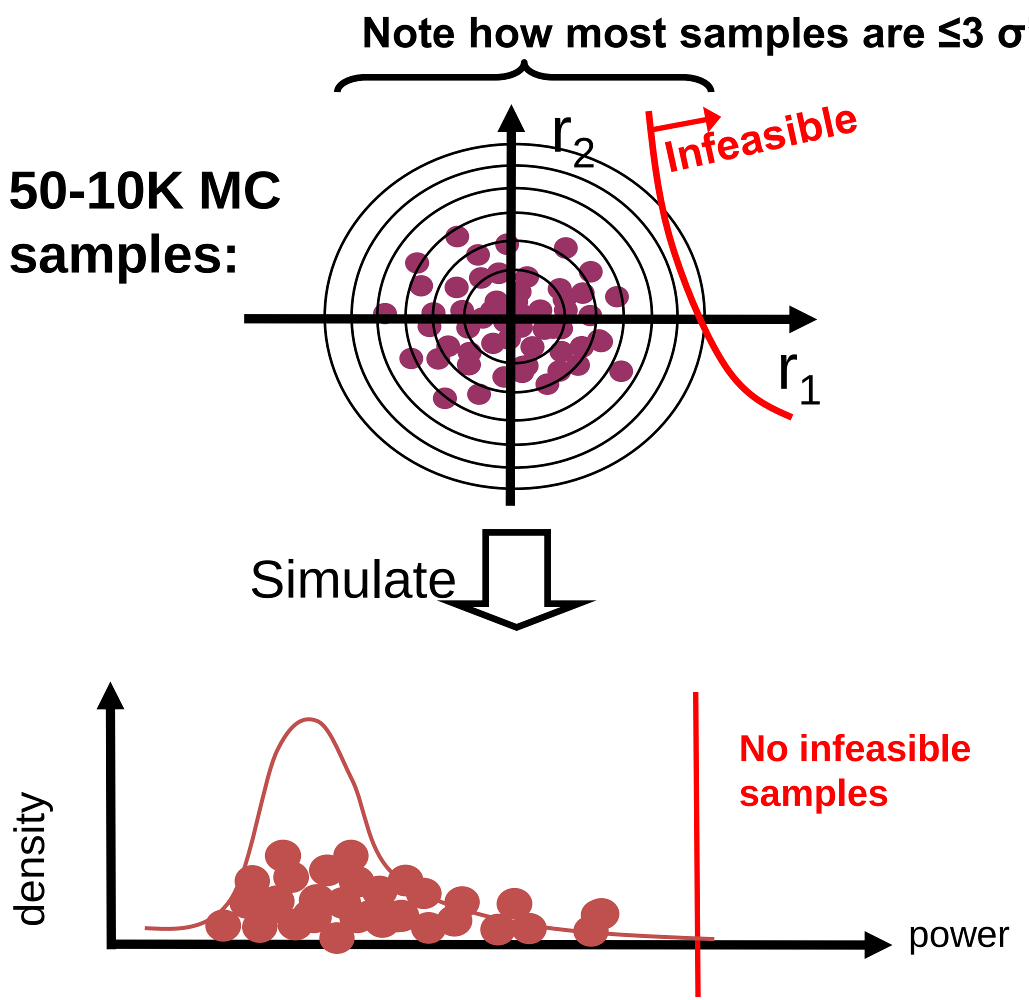

	10K Monte Carlo Samples
Fast	Yes
Accurate	**No**
Scalable	Yes
Verifiable	Yes

Approach: 5G MC Samples

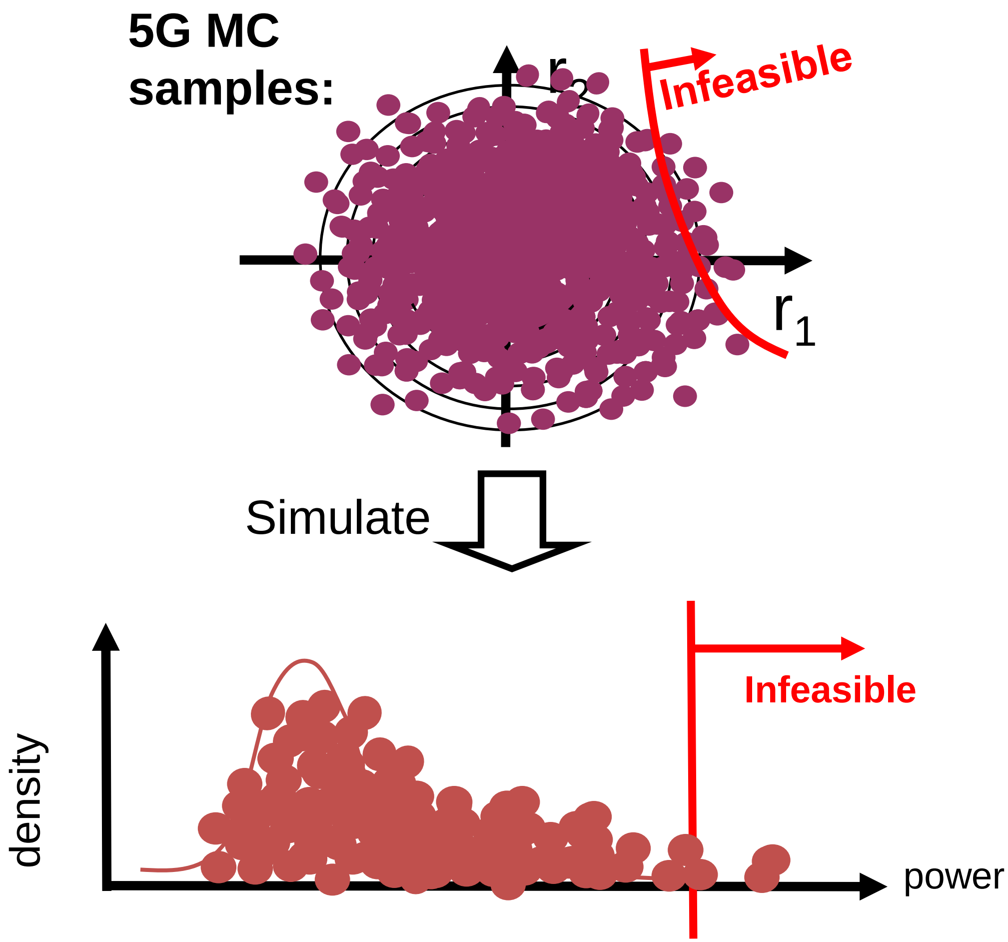

	10K MC Samples	5G MC Samples
Fast	Yes	**No**
Accurate	**No**	Yes
Scalable	Yes	Yes
Verifiable	Yes	Yes

Approach: MC + Extrapolate Tail

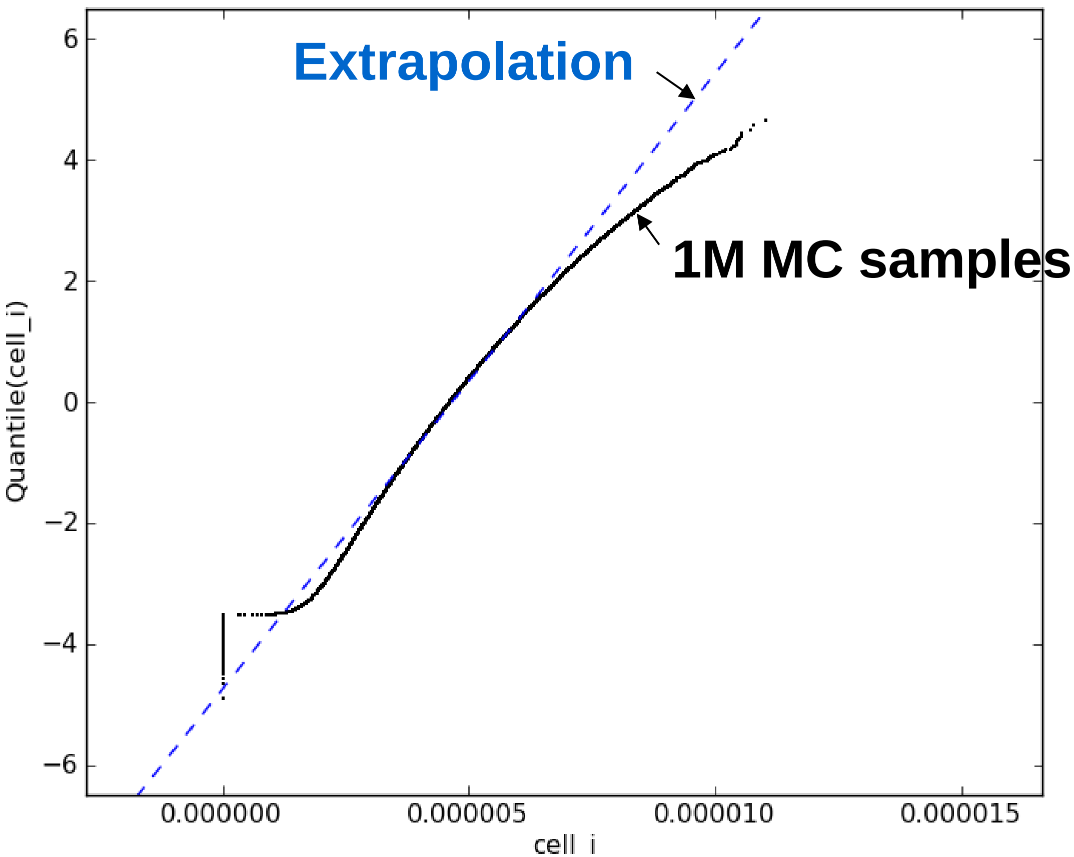

		MC + Extrapolate Tail
Fast	..	**No**
Accurate	..	**No**
Scalable	..	Yes
Verifiable	..	**No**

Approach: Importance Sampling

Step 1: Find highest-probability regions of process space that cause infeas.
Step 2: Draw points from a distribution that is *biased to those regions*

Prob. Failure = weighted sum of failing points
Yield = (1 – Prob. Failure)

Step 1:

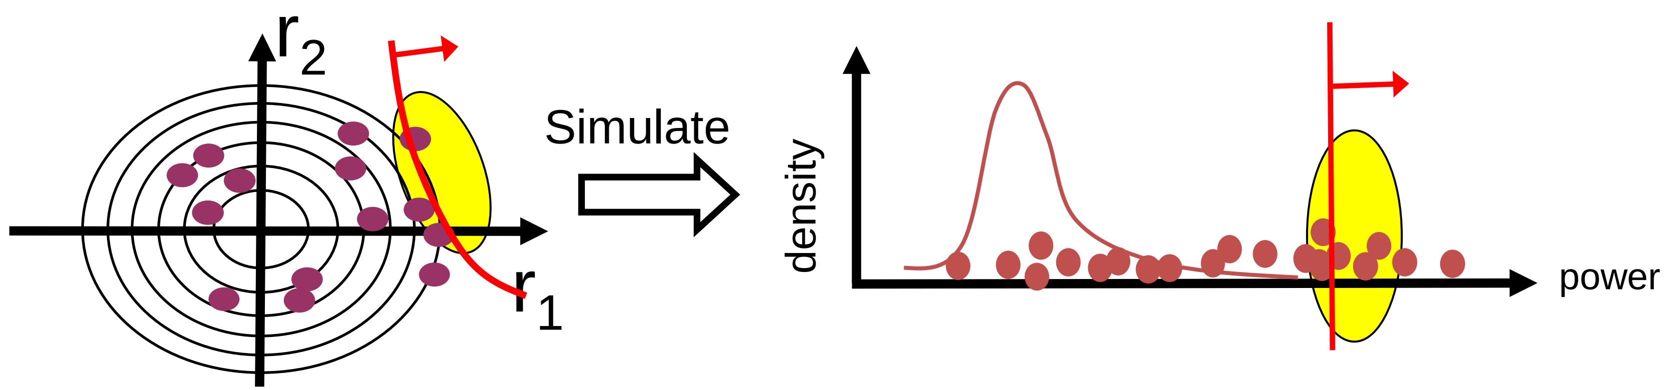

Step 2:

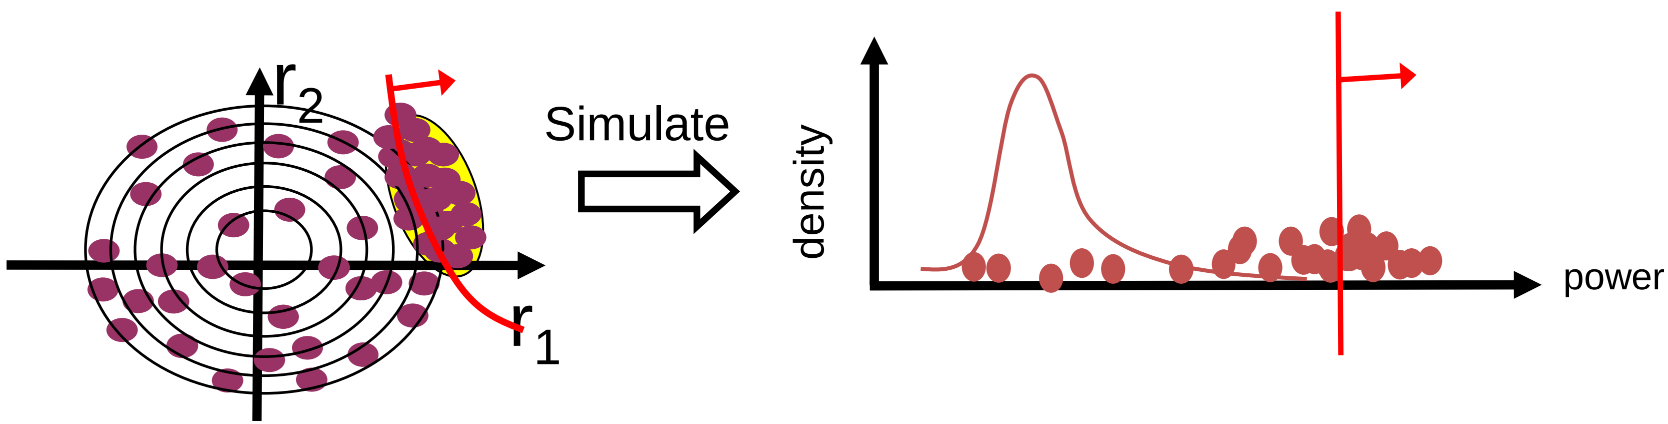

Approach: Importance Sampling

Bitcell test problem:

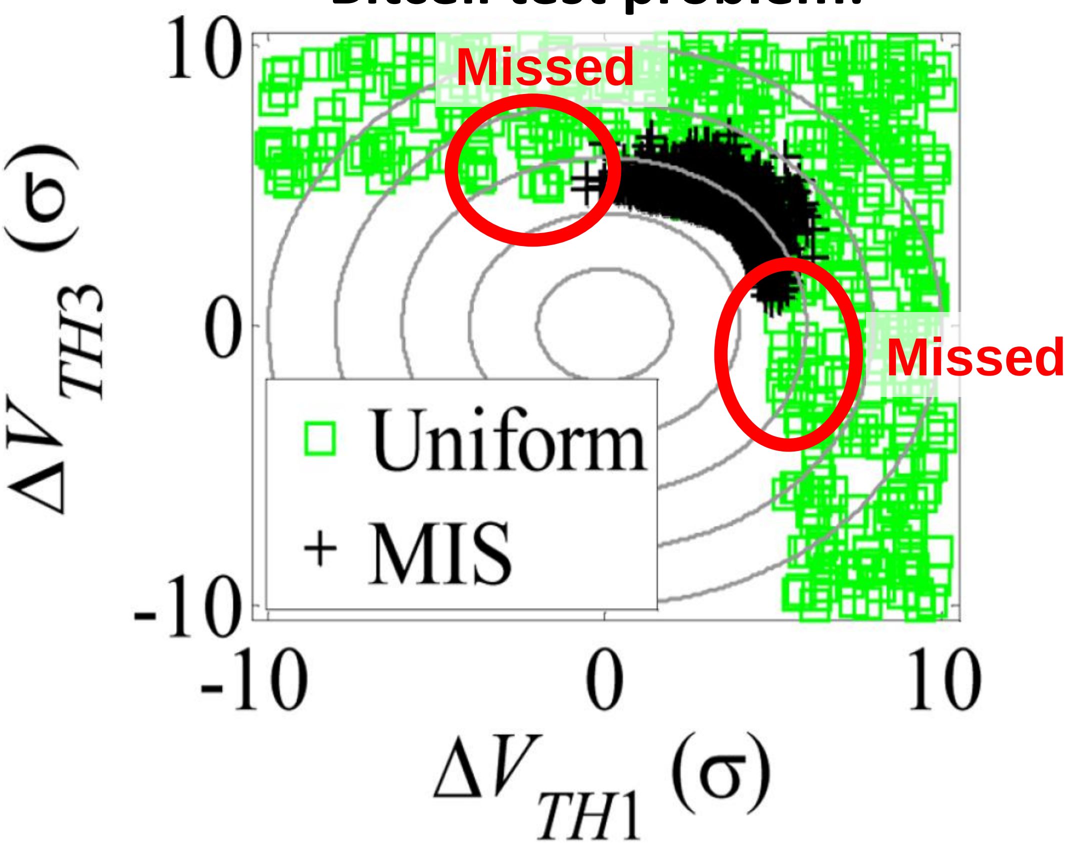

		Importance Sampling
Fast	..	**Yes**
Accurate	..	**No**
Scalable	..	**Sometimes**
Verifiable	..	**No**

Approach: Worst-Case Distances (WCD)

Bitcell test problem:

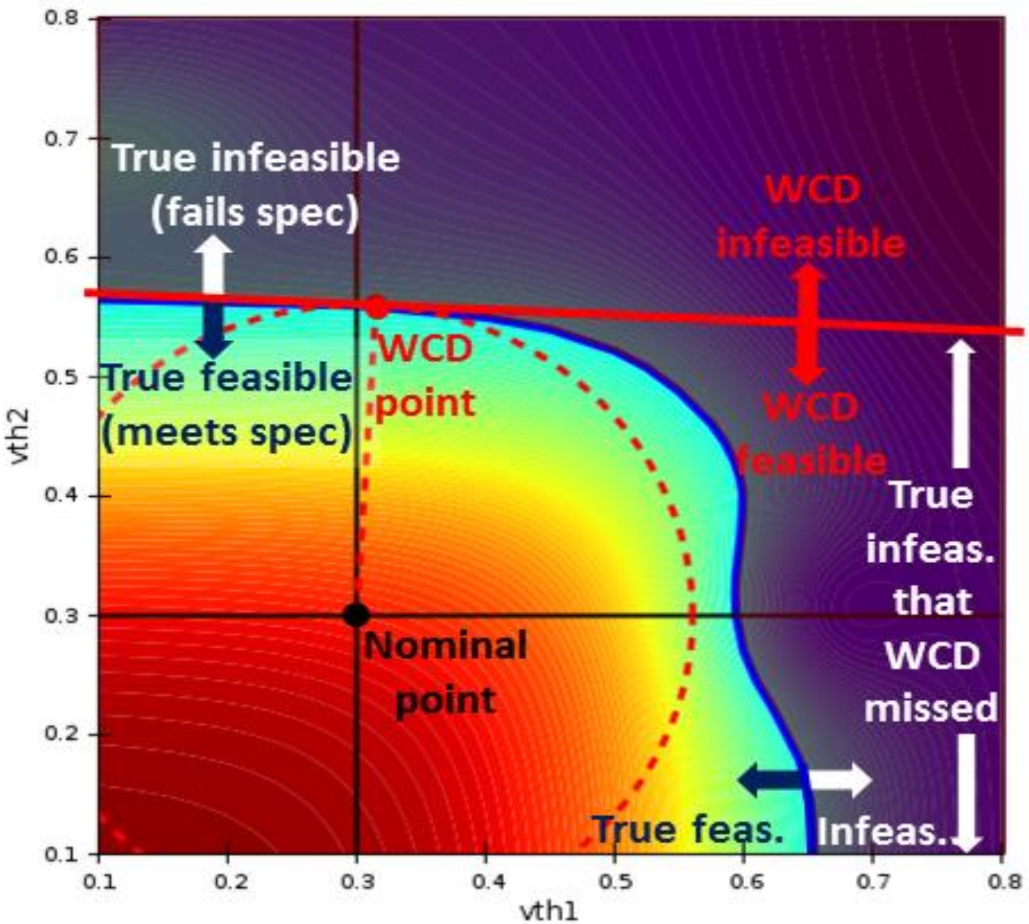

		WCD
Fast	..	**Yes**
Accurate	..	**No**
Scalable	..	**Sometimes**
Verifiable	..	**No**

High-Sigma Yield Prediction:
A New Approach

Rethinking the problem

Hint: Benefit of MC: computational complexity is *independent of dimensionality* (pretty amazing!)

Hint: In other domains, can simplify the problem by adding more information about the problem

Hint: Adding info “problem is 1/1G failures” means we have to consider “just” 10^{10} MC samples. $10^{10} \ll 10^{150}$

Idea: Recast the yield estimation problem into a *ranking* problem: of the 10^{10} MC samples, find the 10-50 that are highest-rank. (≈web search problem)

Approach: High-Sigma Monte Carlo (HSMC)

1. Generate (but don't simulate) 5G Monte Carlo samples
2. Order the 5G samples:
   - Choose initial samples, & simulate
   - Build model that predicts order
   - Order the samples using model
3. Simulate in that order
   - Until tail found / all failures found

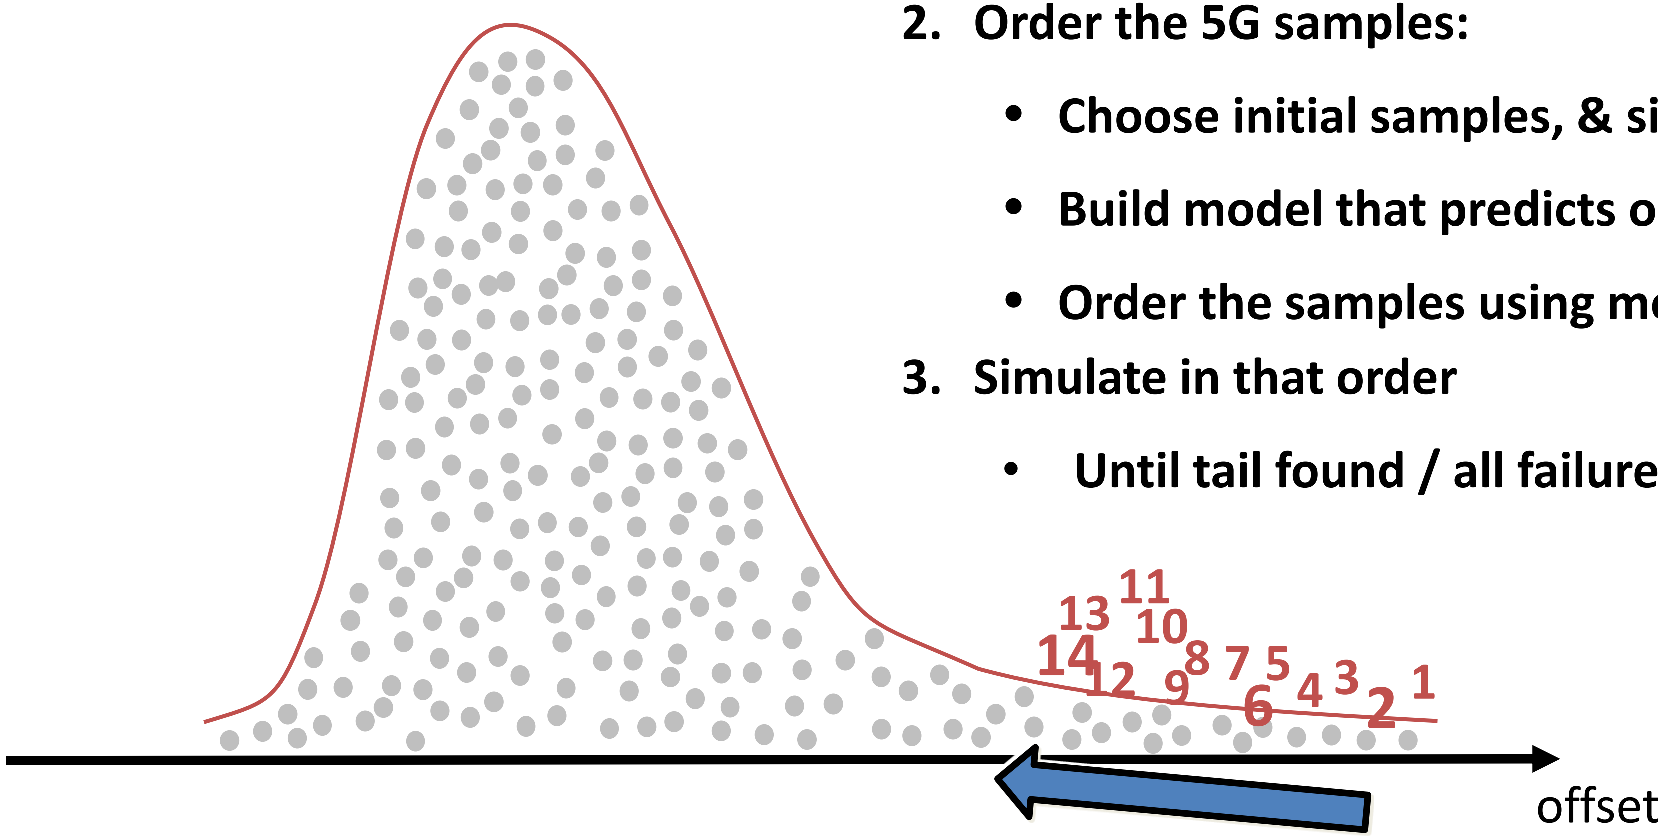

HSMC: Algorithm Details

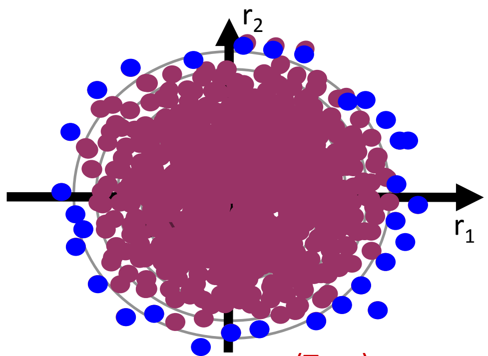

- Generate (but don't simulate) 5G Monte Carlo samples
  - Choose initial samples = the 1000 generated samples that are farthest from nominal

Contour lines are model- predicted output values.

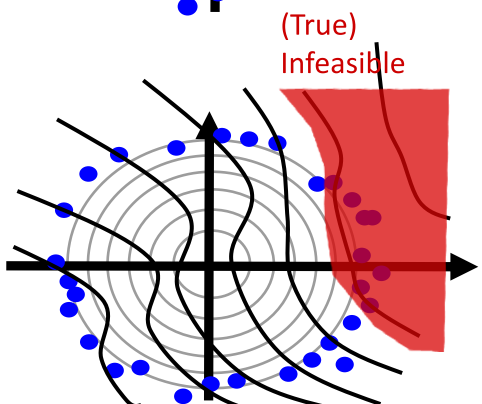

- Simulate initial samples
- Build model that predicts order, mapping process points → simulated output values

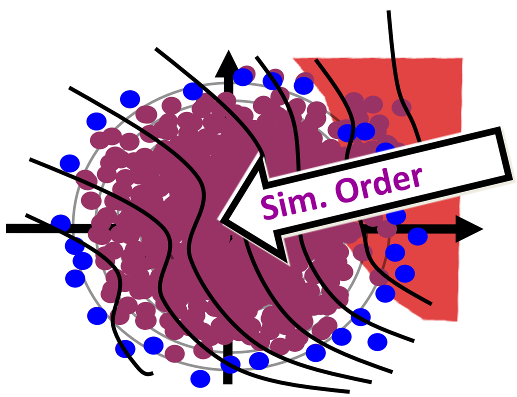

- Order the samples using model. Simulate the points in order, until all failures found or max # sims hit

(PS Parallelize everything: generating samples, simulating, model-building, sorting 5G samples)

Details of HSMC Modeling

- 1K to 40K training points, 60 to 30K variables
- For real-valued outputs, use regression:
  - basis functions are 2nd order polynomial & hockey-stick
  - elastic-net regularized formulation
  - solve with coordinate descent
  - "Fast Function Extraction" (FFX)

- For binary-valued outputs: use two-class classification:
  - basis functions like regression
  - logistic regression formulation
  - solve with stochastic gradient descent

HSMC on Example Test Problem I
(delta bit voltage on 6T SRAM bitcell)

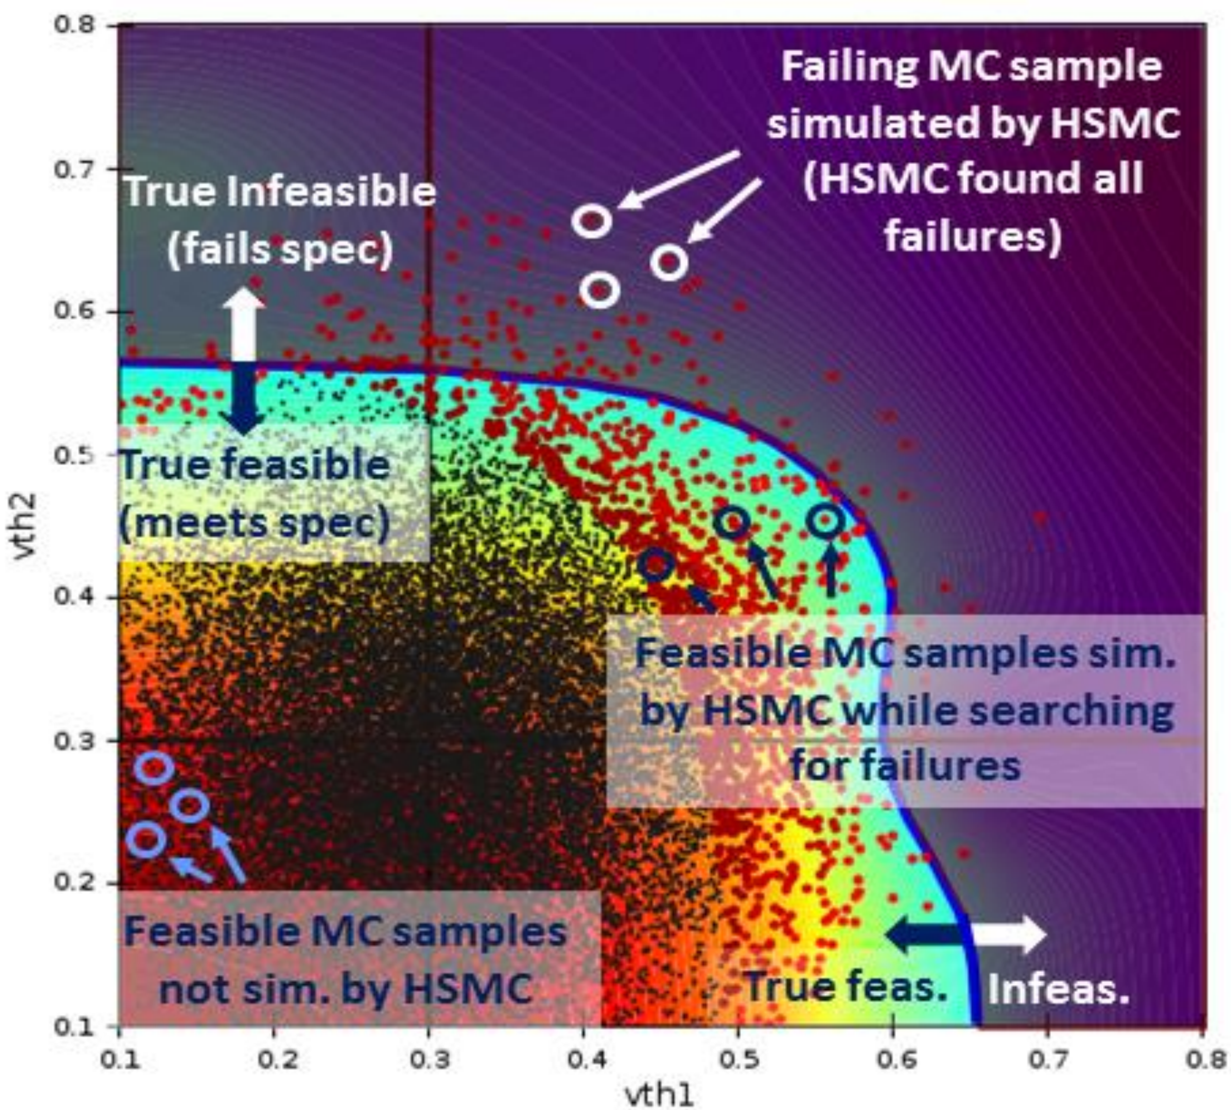

HSMC in Action: NQ Plot of Bitcell cell_i

(HSMC worked off 100M generated)

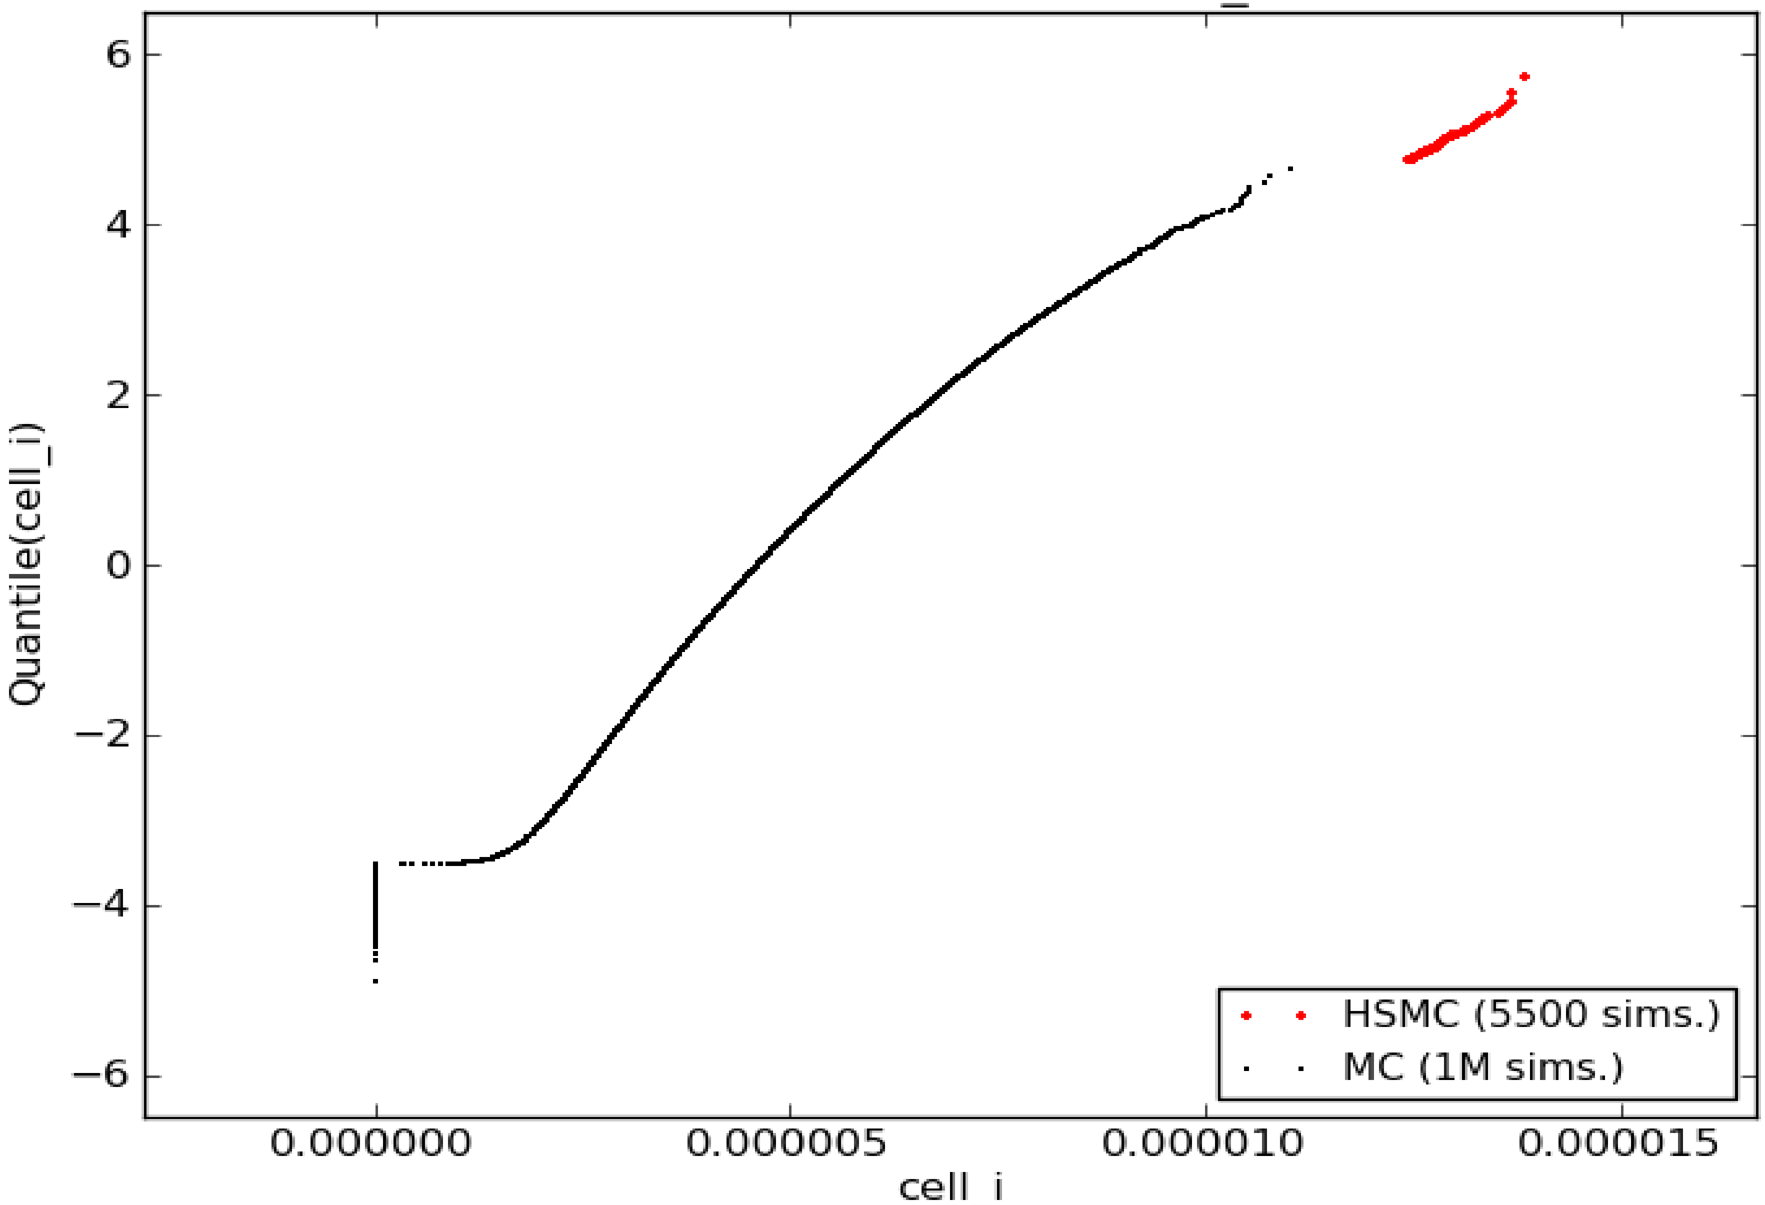

HSMC On Binary-Valued Outputs
Example on 30K-Variable Memory Column

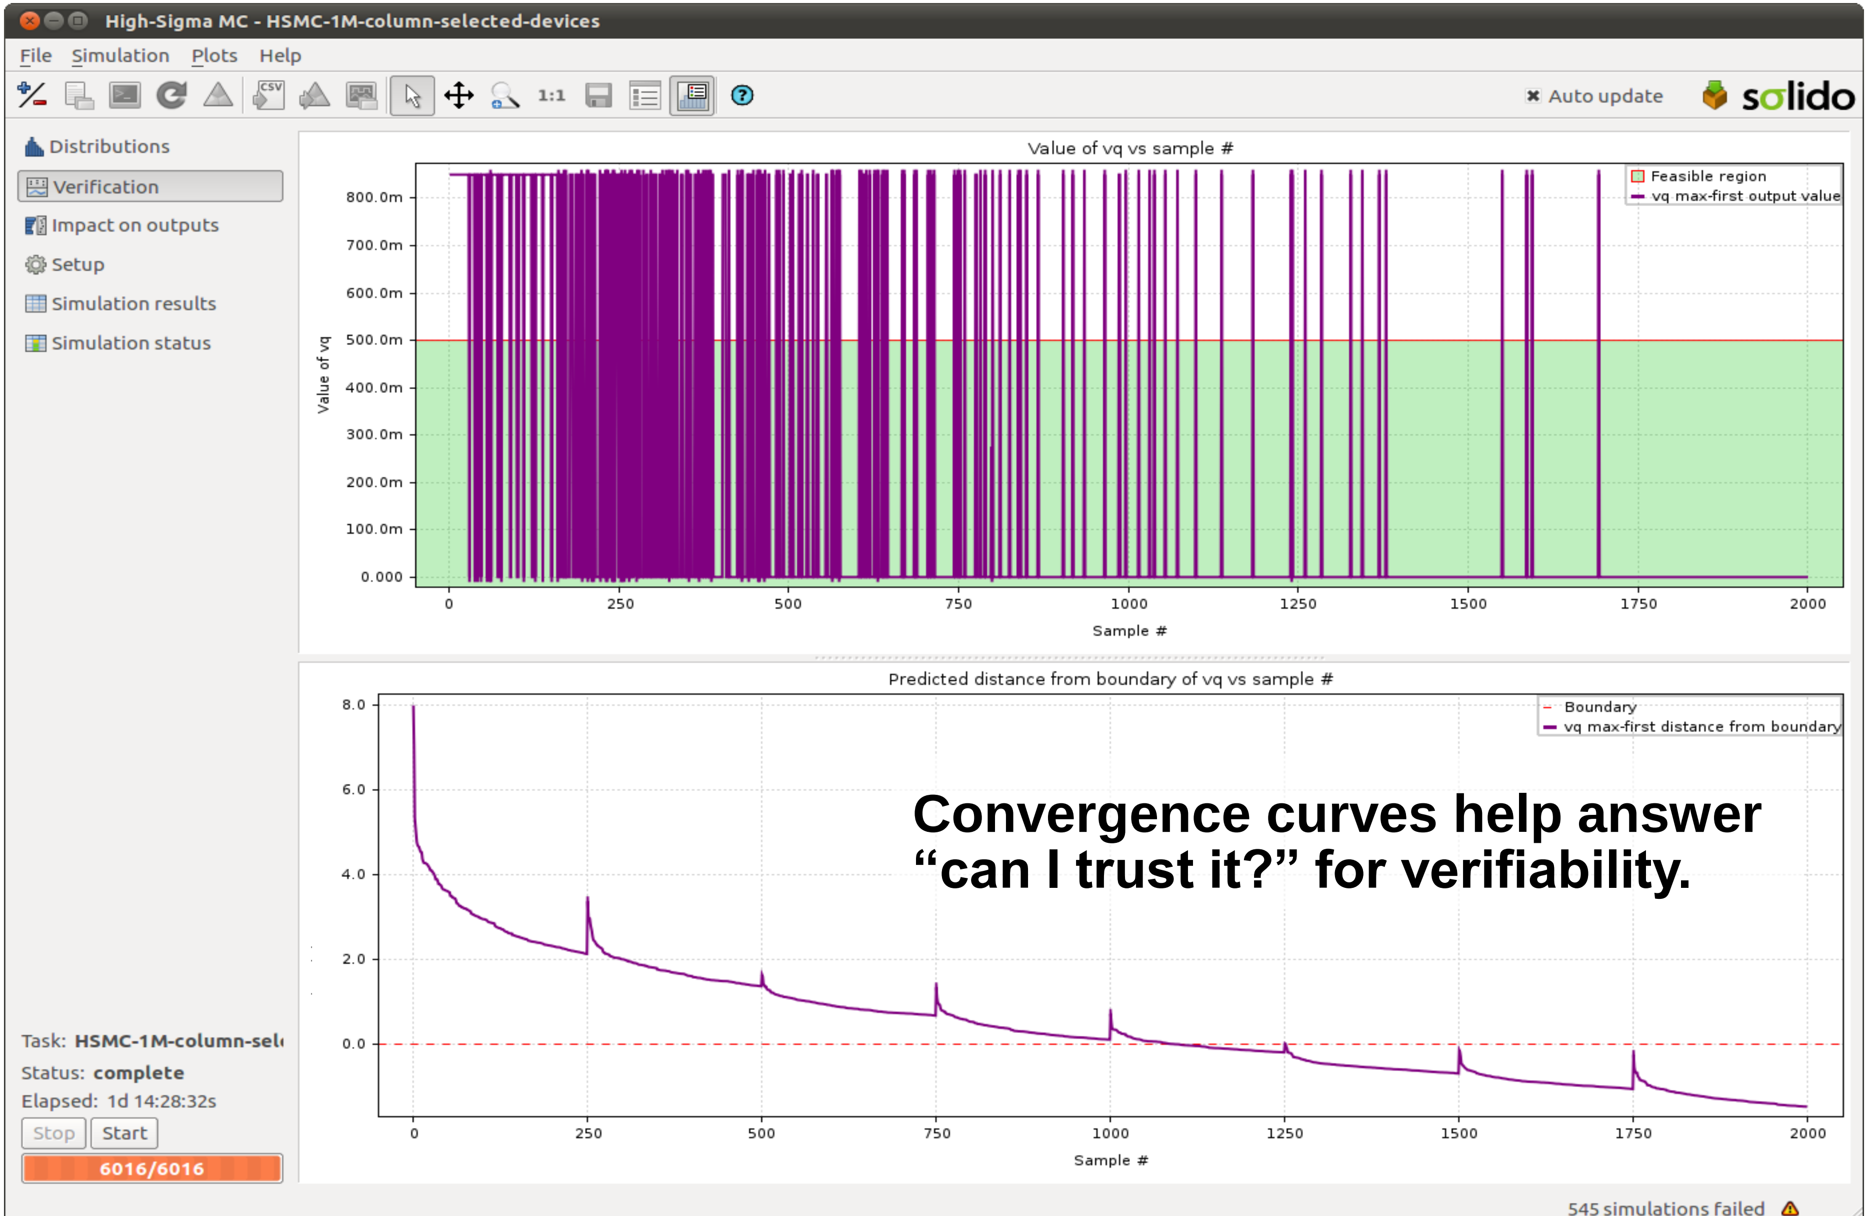

Summary of Approaches To Yield Prediction

	10K MC	5G MC	MC + Extrap.	Importance Sampling	Worst-Case Distances	HSMC
Fast	Yes	No	No	Yes	Yes	Yes
Accurate	No	Yes	No	No	No	Yes
Scalable	Yes	Yes	Yes	Sometimes	Sometimes	Yes
Verifiable	Yes	Yes	No	No	No	Yes

Sony's doing Yield Prediction. PS4 != Black Swan

COMMENTS (13)

Sony PlayStation 4 Enters Mass Production, Has Phenomenal Production Yield.

Sony Hopes to Fulfill Demand for PlayStation 4 This Year

[09/16/2013 11:10 PM]
by Anton Shilov

When Sony Corp. launched its PlayStation 3 game console seven years ago, the firm face a number of problems, one of which was high price of the system and another one was poor yield of Cell processor, the heart of the PS3. With PlayStation 4 the company made a lot of decisions differently and now it hopes that the same issues that harmed success of the PS3 will not hurt the PS4.

"It is everything we can manufacture. We want to make sure that consumers have an opportunity to buy one on November 15 and through the holidays, so we are

Semiconductor companies now routinely use HSMC to predict their black swans

QUALCOMM®

NVIDIA®

• • •

and many more

Conclusion

- **Black swan events are very rare and high impact**
- **A popular belief is that they are unpredictable.**
  - **Advice: "Don't try to predict. Just protect from downside, and stay open for upside."**
- ***So what* if something's hard? That's a challenge to algorithms people :)**
- **The specific problem to solve the high-sigma yield prediction.**
- **Many approaches (extrapolate, importance sampling, WCD).**
- **A new approach, High-Sigma Monte Carlo, *ranks* 5G points with the help of machine learning.**
- ***HSMC slays semiconductor black swans***