


Radyalis

Medical Treatment For Variability:

Lessons from Circuits Applied to Cancer Radiation Therapy

Sani R. Nassif

CEO

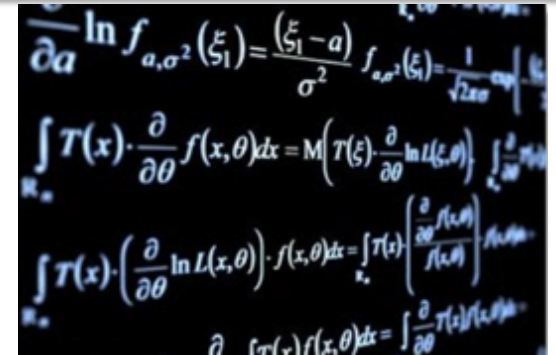
Some Context



- From 1980 to 2013, I worked at Bell Labs and IBM Research on the modeling and analysis of technology variability.
- Since 2010, I have been working on applying CAD techniques to the area of Cancer Radiation Therapy.
- As of 2014, formed “Radyalis” to put these ideas into real-world practice.
- This talk explores links between these two areas!

Do “Circuits” Matter?

- I am an Electrical Engineer.
- I have spent the last 28 years doing industrial R&D on Technology, EDA, VLSI etc...
- Observation: VLSI researchers have strong skills in and links to mathematics and physics.
 - As well as software, algorithms, modeling, etc...
 - We deal well with hierarchy and complexity.
 - And are accustomed to large-scale problems.
 - This is attributable to Moore’s Law forcing us to tackle ever larger and more complex problems.
- These foundational skills are highly portable to other areas.
- Requirements: partnership and perseverance!




$$\frac{\partial}{\partial a} \ln f_{a,\sigma^2}(\xi_1) = \frac{(\xi_1 - a)}{\sigma^2} f_{a,\sigma^2}(\xi_1) = \frac{1}{\sqrt{2\pi\sigma^2}} \left[\frac{\xi_1 - a}{\sigma^2} \right] e^{-\frac{(\xi_1 - a)^2}{2\sigma^2}}$$

$$\int \mathcal{T}(x) \cdot \frac{\partial}{\partial \theta} f(x, \theta) dx = M \left(\mathcal{T}(\xi) \cdot \frac{\partial}{\partial \theta} \ln L(\xi, \theta) \right) \int_{-\infty}^{\infty} \frac{\partial}{\partial \theta} f(x, \theta) dx$$

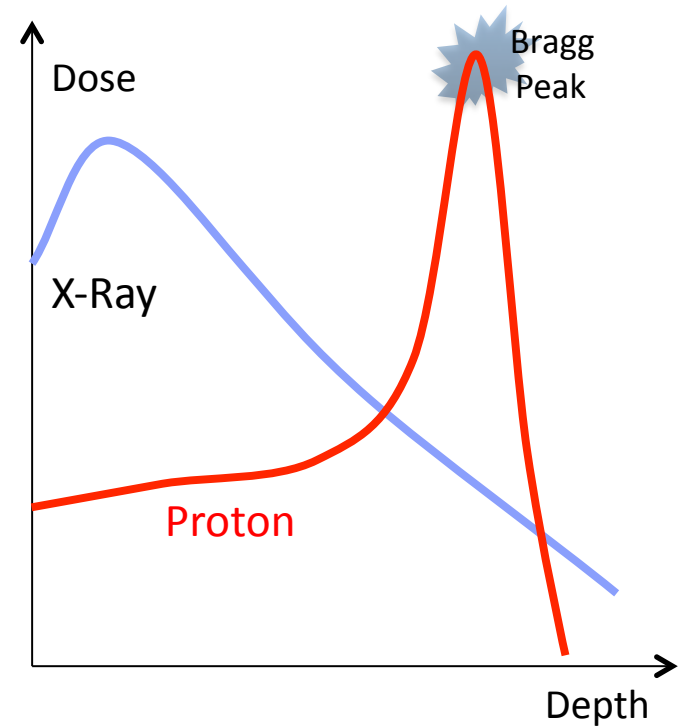
$$\int \mathcal{T}(x) \cdot \left(\frac{\partial}{\partial \theta} \ln L(x, \theta) \right) \cdot f(x, \theta) dx = \int \mathcal{T}(x) \cdot \left(\frac{\partial}{\partial \theta} \ln f(x, \theta) \right) \cdot f(x, \theta) dx$$

$$\frac{\partial}{\partial \theta} \int \mathcal{T}(x) \cdot f(x, \theta) dx = \int \frac{\partial}{\partial \theta} \mathcal{T}(x) \cdot f(x, \theta) dx$$

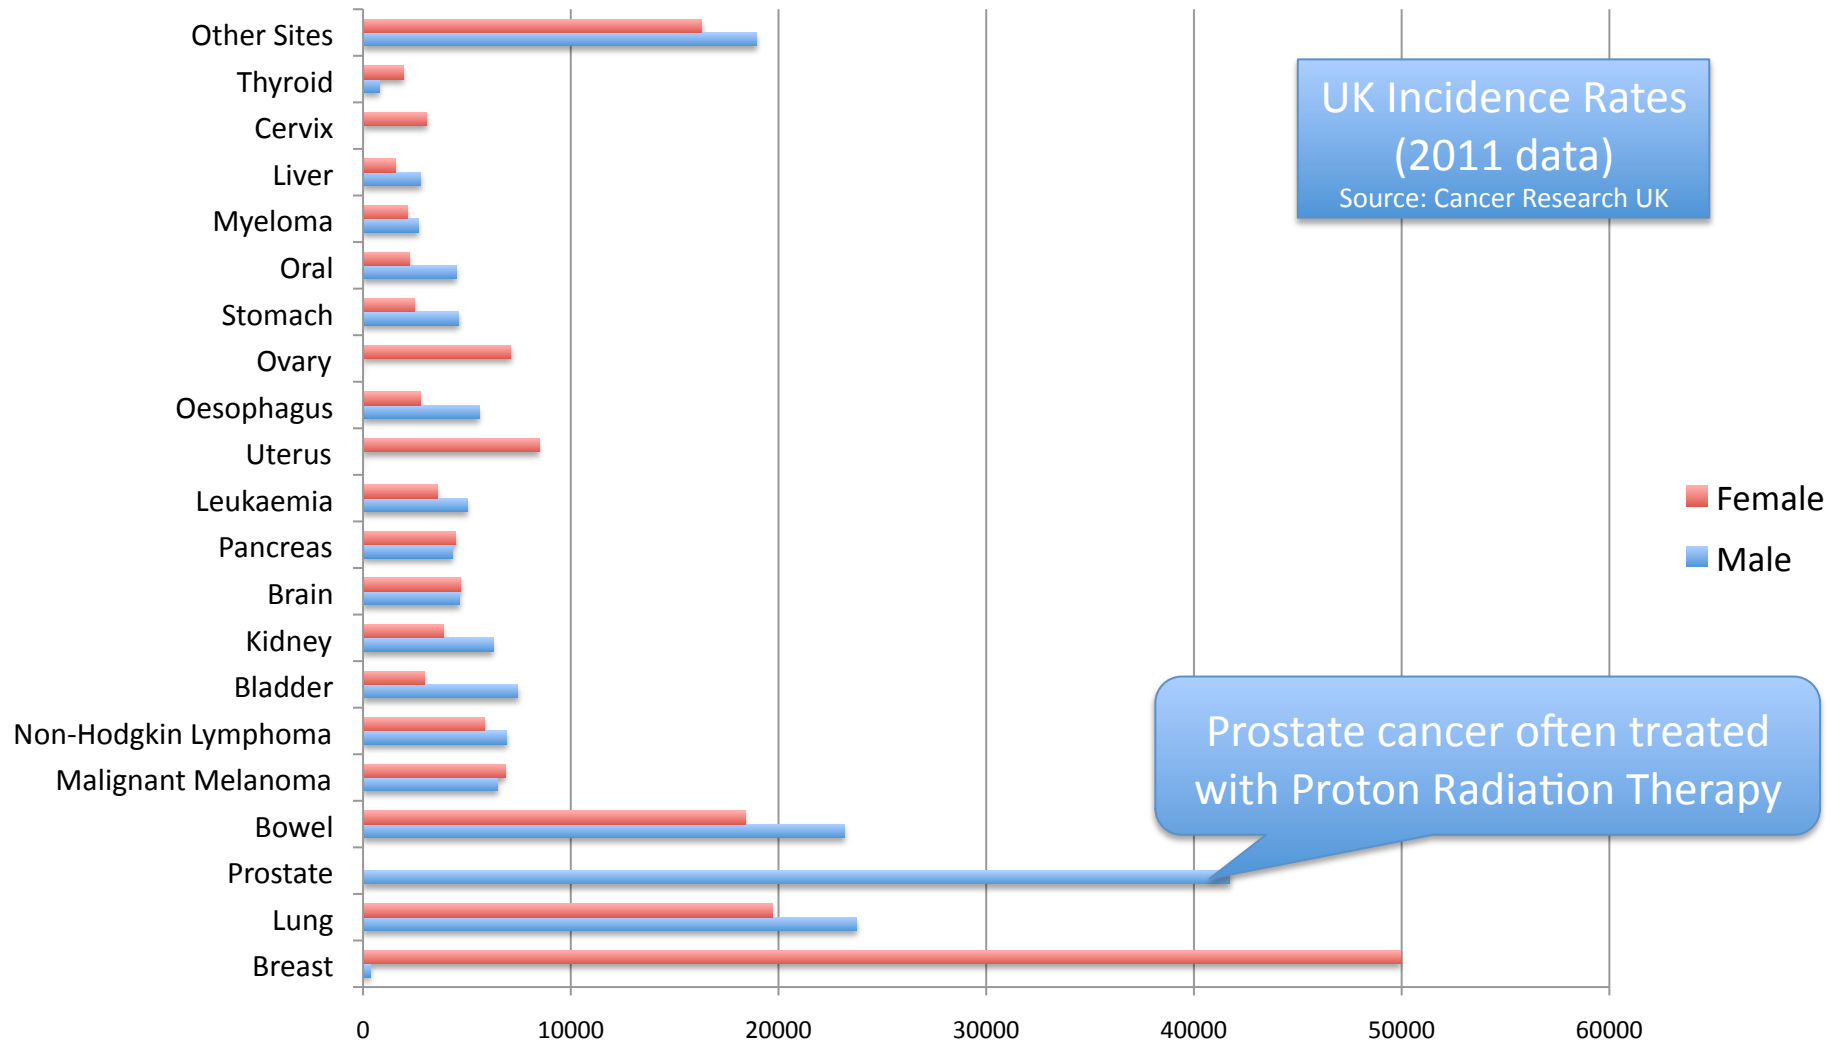


Radiation and the Human Body

- Radiation deposits ionizing energy in the atoms that make up a cell, and results in disruption of the DNA. Cells stop reproducing and eventually die...
- Two major types: X-Ray (Photon) and Proton.
 - About 60% of U.S. patients get radiation.
- Energetic Protons interact with solids in a very non-intuitive way:
 - Energetic $\approx 250\text{MeV}$ ($\frac{1}{2}c$).
 - Energy loss rate = $\delta E/\delta x \sim 1/E$.
- Leads to a strongly localized peak in energy deposited vs. depth.
 - Ideal for treating tumors deep in the body (like the Prostate, + many others)!
- Much more selective than Photons.
 - But more \$ and more difficult to plan.



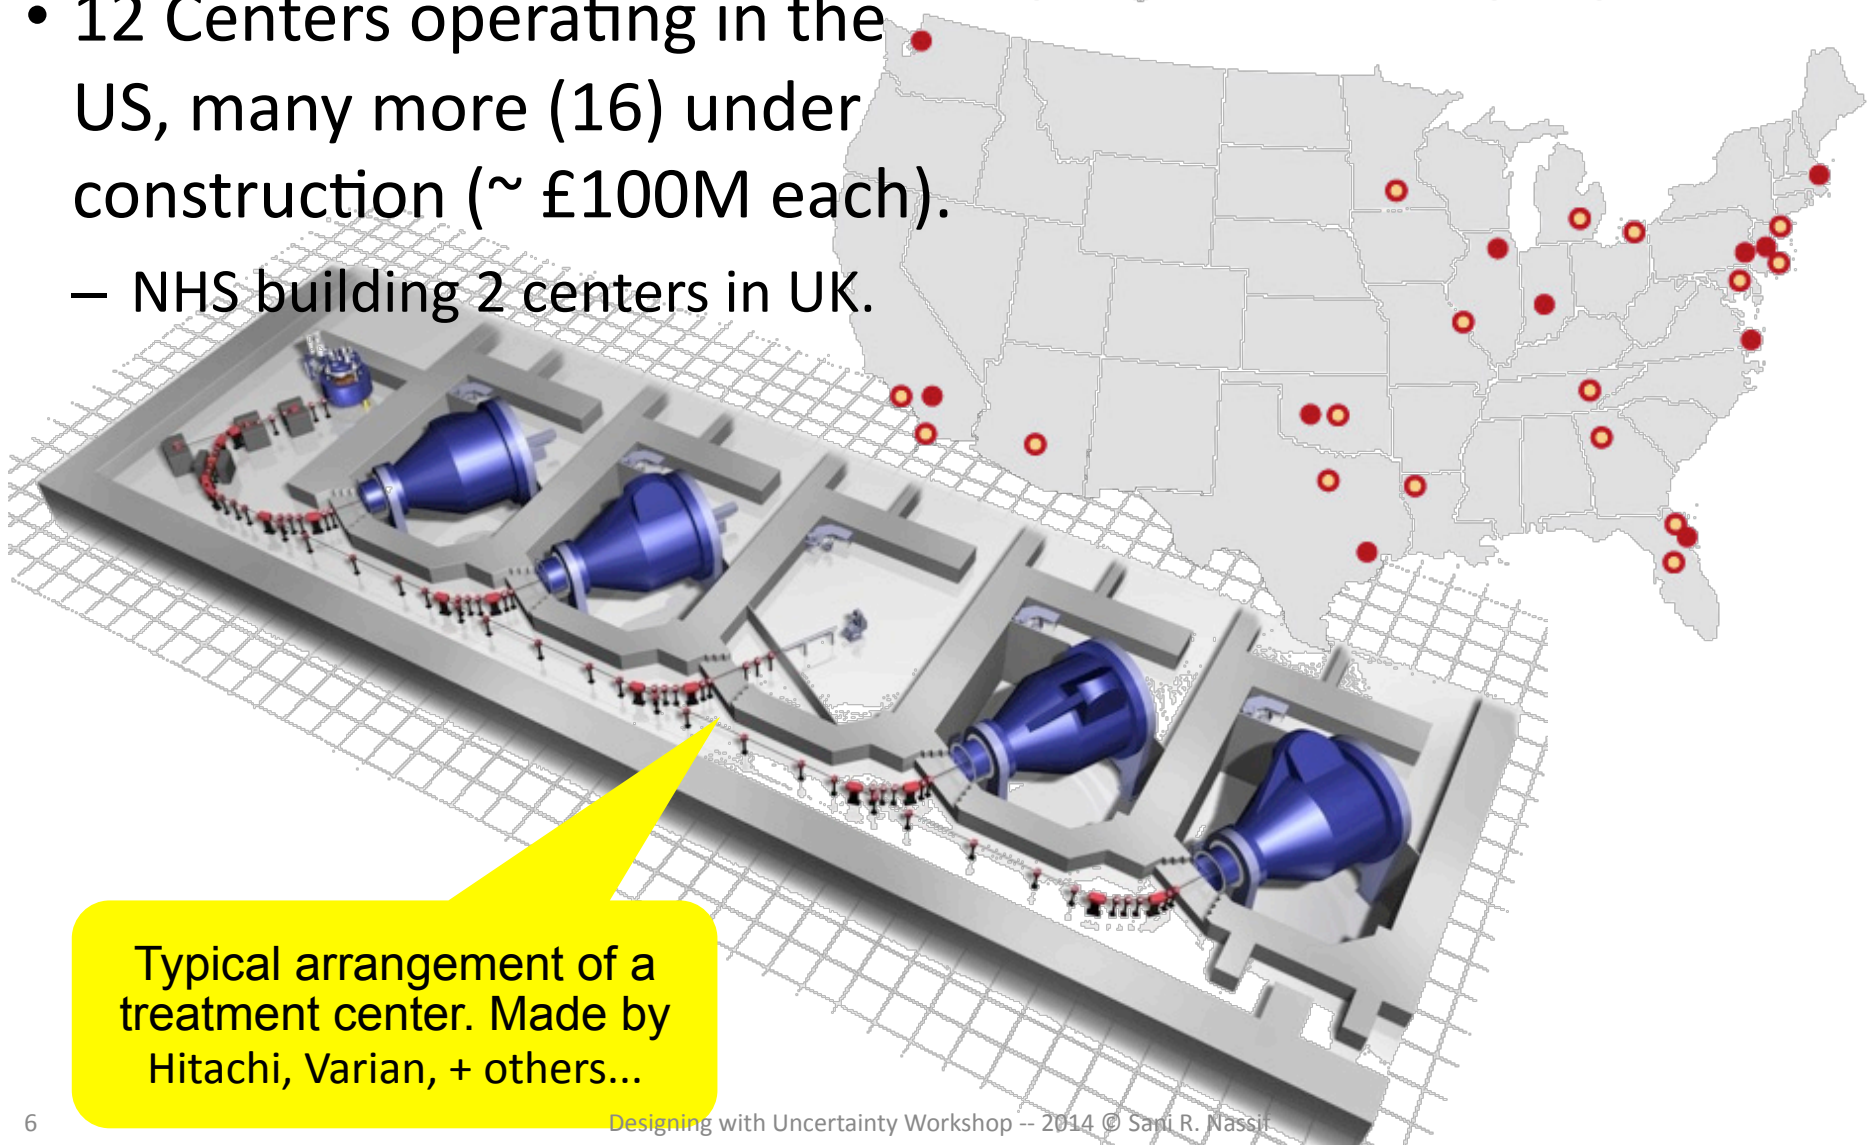
A Growing Problem Worldwide...



Radiation Treatment Centers (in the US)

- 12 Centers operating in the US, many more (16) under construction (~ £100M each).
 - NHS building 2 centers in UK.

● Operating ● Under construction/development



Typical arrangement of a treatment center. Made by Hitachi, Varian, + others...

But 1st... A Recap of Circuit Variability

- Variability and uncertainty exist at all levels of the design process, and becoming more **PERVASIVE** with scaling.
- Design variability is caused by both manufacturing fluctuations **AND** other (design related) sources.
- Variability problems are difficult to handle without adequate **MODELS** of variability and its distribution.
- Variability problems are **IMPOSSIBLE** to solve without CAD support (designers are busy enough as it is!).

("A") Taxonomy

- There are many ways to “slide and dice” variability. Some useful for process control, and others for design analysis and optimization.

	Physical (time scale $10^1 \approx 10^9$ sec)	Environmental (time scale $10^{-9} \approx 10^1$ sec)
Spatial	$\Delta L, \Delta \mu, \Delta V_T$	Noise coupling
Temporal	NBTI, Hot- \bar{e} , electro-migration	V_{DD}, T

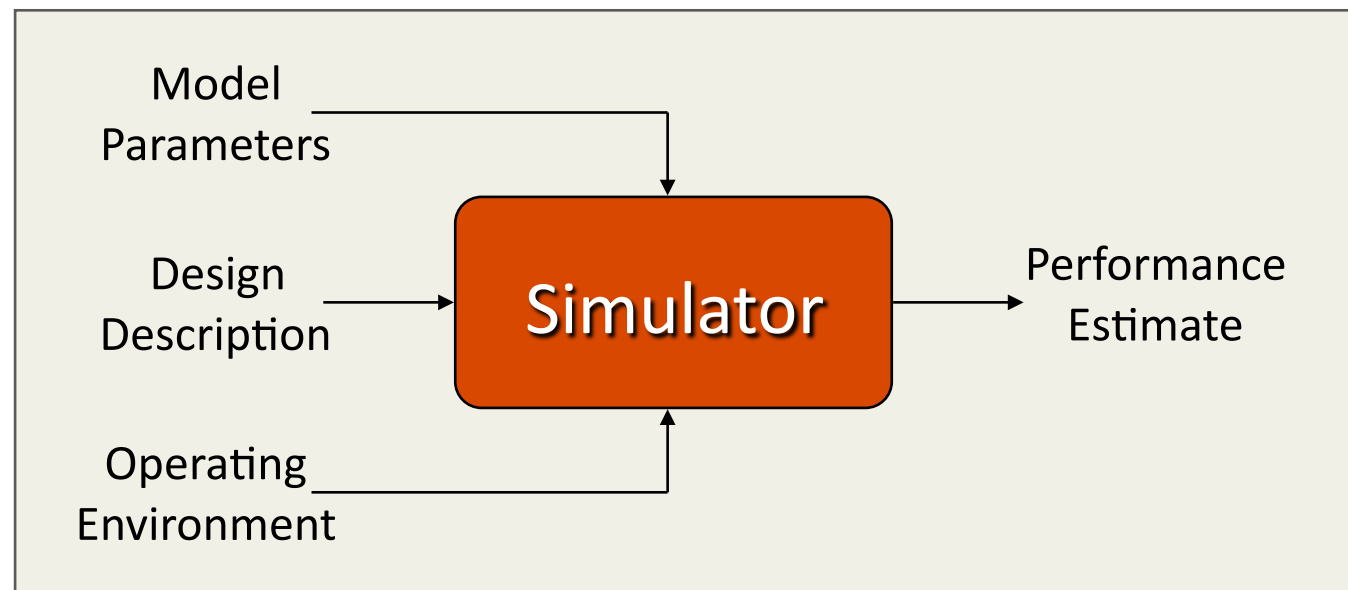
("Another") Taxonomy

- Die to die variation
 - Imposed upon design (constant regardless of design).
 - Largely Random.
 - Well modeled via worst-case files.

- Within-die variation
 - Co-generated between design & process (depend on details of the design).
 - Largely Systematic.
 - Example: Via resistance variation vs. via density.

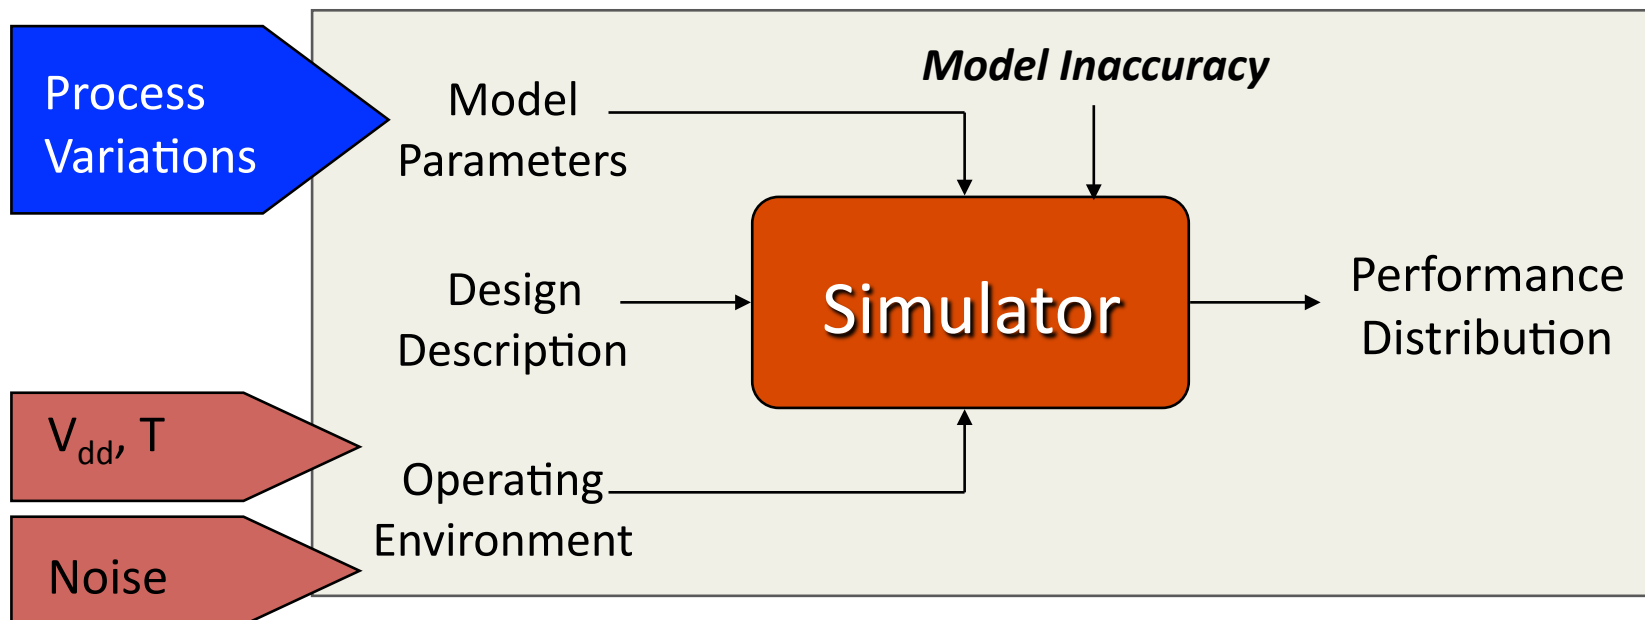
Performance Estimation

- Basic assumptions:
 - Performance measured by simulation (we hardly ever prototype).
 - Technology represented by Spice models (root of a rich hierarchy of models).
 - Parameters have distributions (with correlations + spatial and temporal dependence).
- Nominal performance estimation:



Variability Estimation

- Similar to nominal simulation, but includes variability.
- Can be as simple as Monte-Carlo analysis, or as complex as chip-level statistical static timing analysis.
 - Reality is that “worst case” analysis is the only type routinely done.



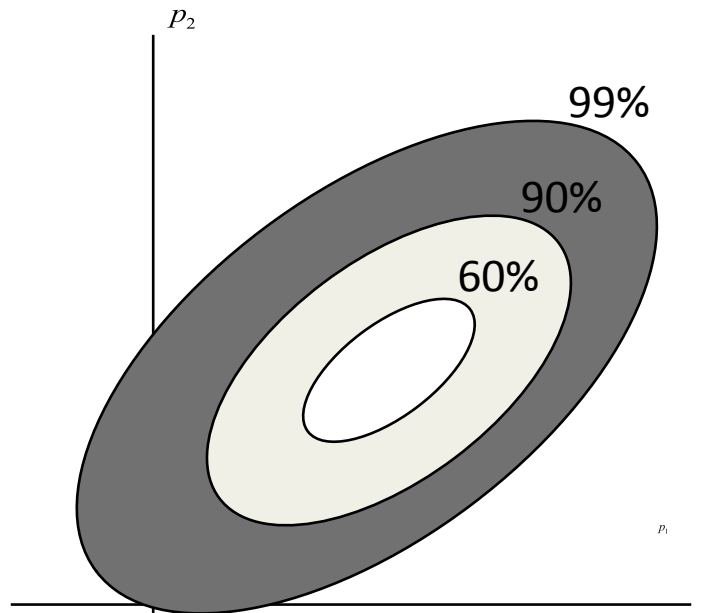
Classical Variability Analysis

- Given:
 - A design specification.
 - A technology description.
 - Target performance, ζ .
 - Determine:
 - The expected $n\sigma$ spread of ζ .
 - The distribution of ζ .
 - How to change design to improve yield.
- Worst-Case Analysis is the **only** commonly applied technique in current digital flows.
 - Worst or extreme case analysis.
 - Yield estimation.
 - Yield maximization.

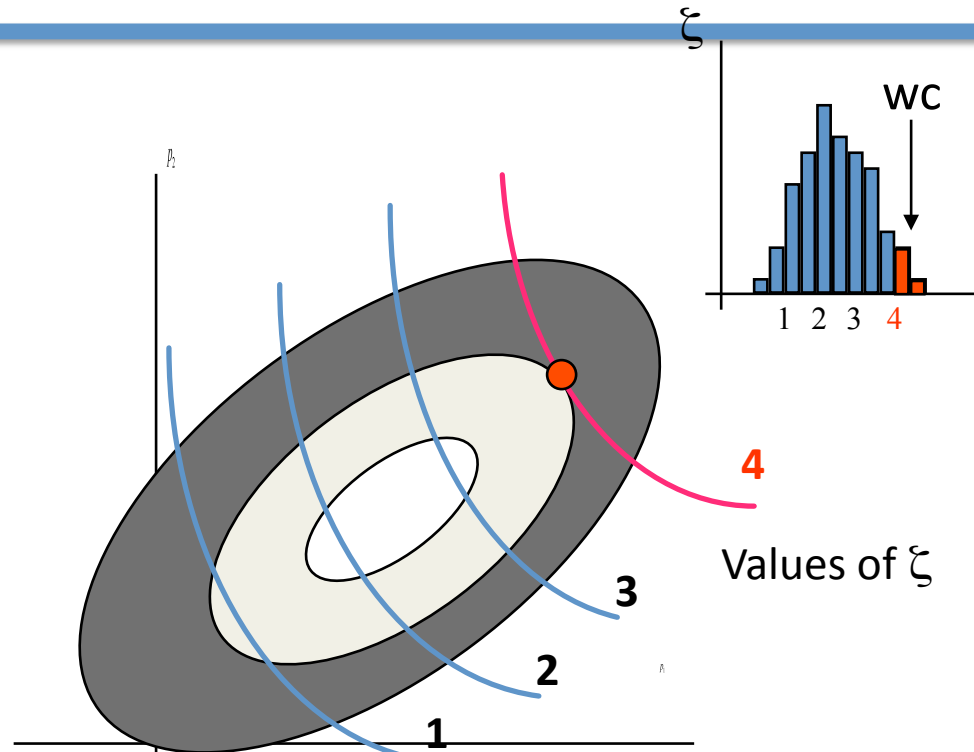
Yield estimation is used for difficult parts like SRAM.

Yield maximization is rarely used, but yield “improvement” is...

Statistical Performance Basics



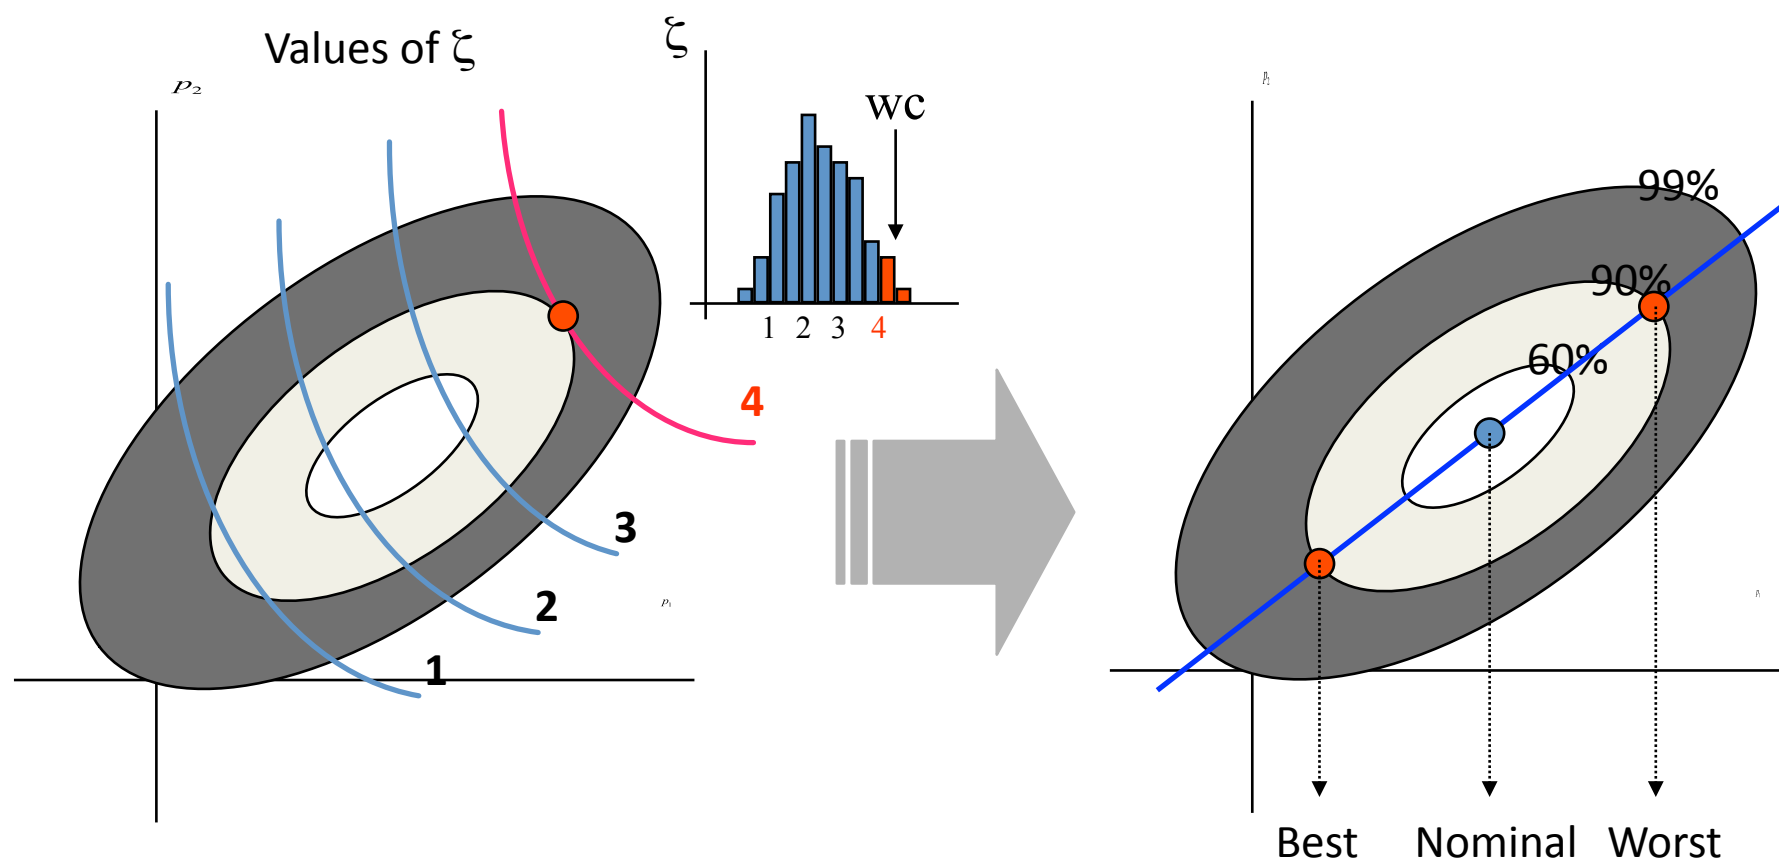
Distribution of
Model Parameters



Contours of
Performance Values

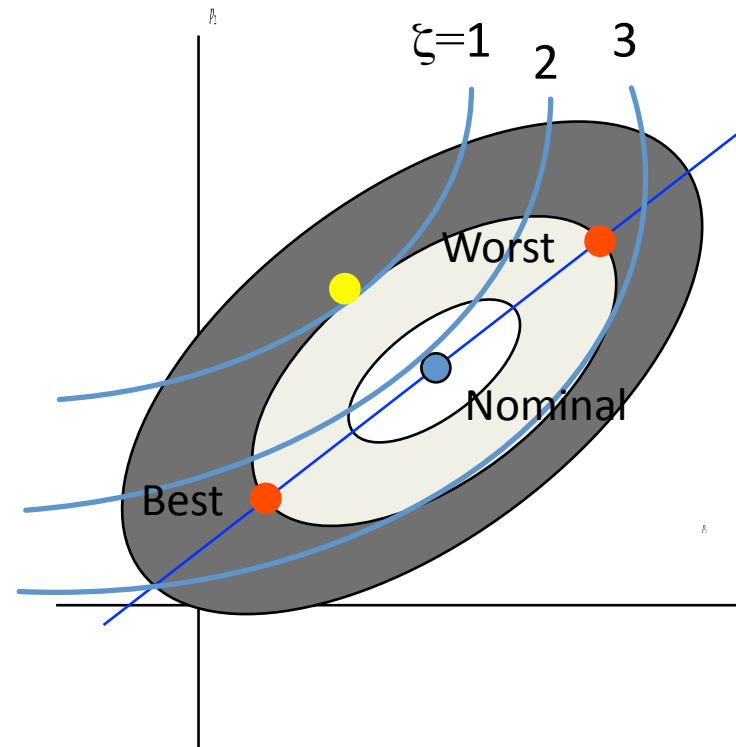
Worst Case Analysis (WCA)

- Perform analysis on one or more typical circuits and apply resulting parameters to all circuits.



Problems with WCA

- Assumes *identical* sensitivities to parameter variations for *all* performances of *all* circuits.
- Performances which exhibit different sensitivities will produce *inaccurate* worst cases.
- Current design methodologies do not assess such sensitivities, and therefore are open to significant errors when WCA is applied “blindly”.

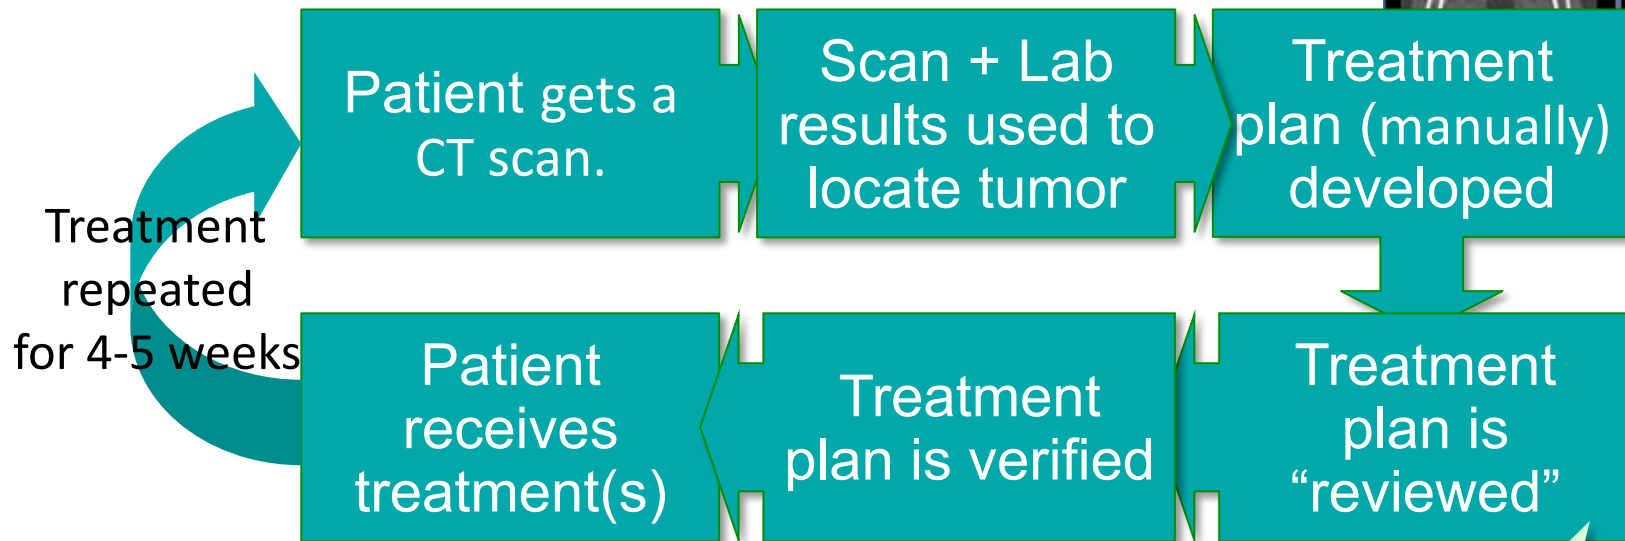





Variability in Today's (digital) Design Flow

- A deep stack with a physical “bottom” and a rather abstract “top” (and interspersed with magic).

Chip Characterization (Delay, Power)	Excel
Core Characterization (Delay, Power)	Excel + STA
Unit Characterization (Delay, Power)	Cleverness + STA
Macro Characterization (Delay, Power)	STA
Cell Library Characterization (Delay/Slew vs. input/load)	Spice
Corner Generation (Spice Models for Slow/Fast etc...)	 Magic + Spice
Technology Characterization (Spice Models)	TCAD + Spice


From Circuits to Cancer?



Treatment repeated for 4-5 weeks

Delay: Many days!

Verification is done by applying treatment to a sensor in a water bath.



60% of cost is in "staff" because treatment requires Oncologists + PhD Physicists!


Plan?: Select Best Beams for Patient

- Many degrees of freedom available for treatment.
 - Beam Energy, Direction, Treatment Time, Beam Shape ...
- A very large search space with millions of options & constraints.
 - Impact of each beam requires a large (~hours) Monte-Carlo to evaluate.
 - Current approaches use simple analytical models instead...
- Typical computational cost for optimization: 3 days.
 - Experts “gave up”, or scaled problem down drastically.
- But... VLSI engineers are not scared of solving large optimization problems.



Formulating the Problem: Objective


- The physical domain is naturally discretized because scan data has limited spatial resolution, so...
 - We create a discrete version of the problem working on each individual “voxel”.
 - Tumor voxel □: high dose.
 - Sensitive Organ voxel □: low dose.
- Leads to a linear program.
 - 10cm cube with 1mm voxels.
 - $100 \times 100 \times 100 = 1\text{M}$ voxels.
 - 1M dose constraints: $\text{Dose}(x,y,z) \leq / \geq C(x,y,z)$.
- VLSI folk love linear programs...



Enlarged voxels for illustration only


Formulating the Problem: Variables

- The radiation delivery “gantry” has limited setting resolution, so...
 - Beam angles and energies can be considered to be discrete.
 - We can enumerate all possibilities, free variable is beam time, T .
 - Number of possible beams $\approx 1M$.
- For each beam, we determine the dose to each voxel per unit time.
 - Total Dose(x,y,z) = $\sum T_i$ Dose $_i$ (x,y,z).
- So the linear program is 1M variables and 1M constraints.

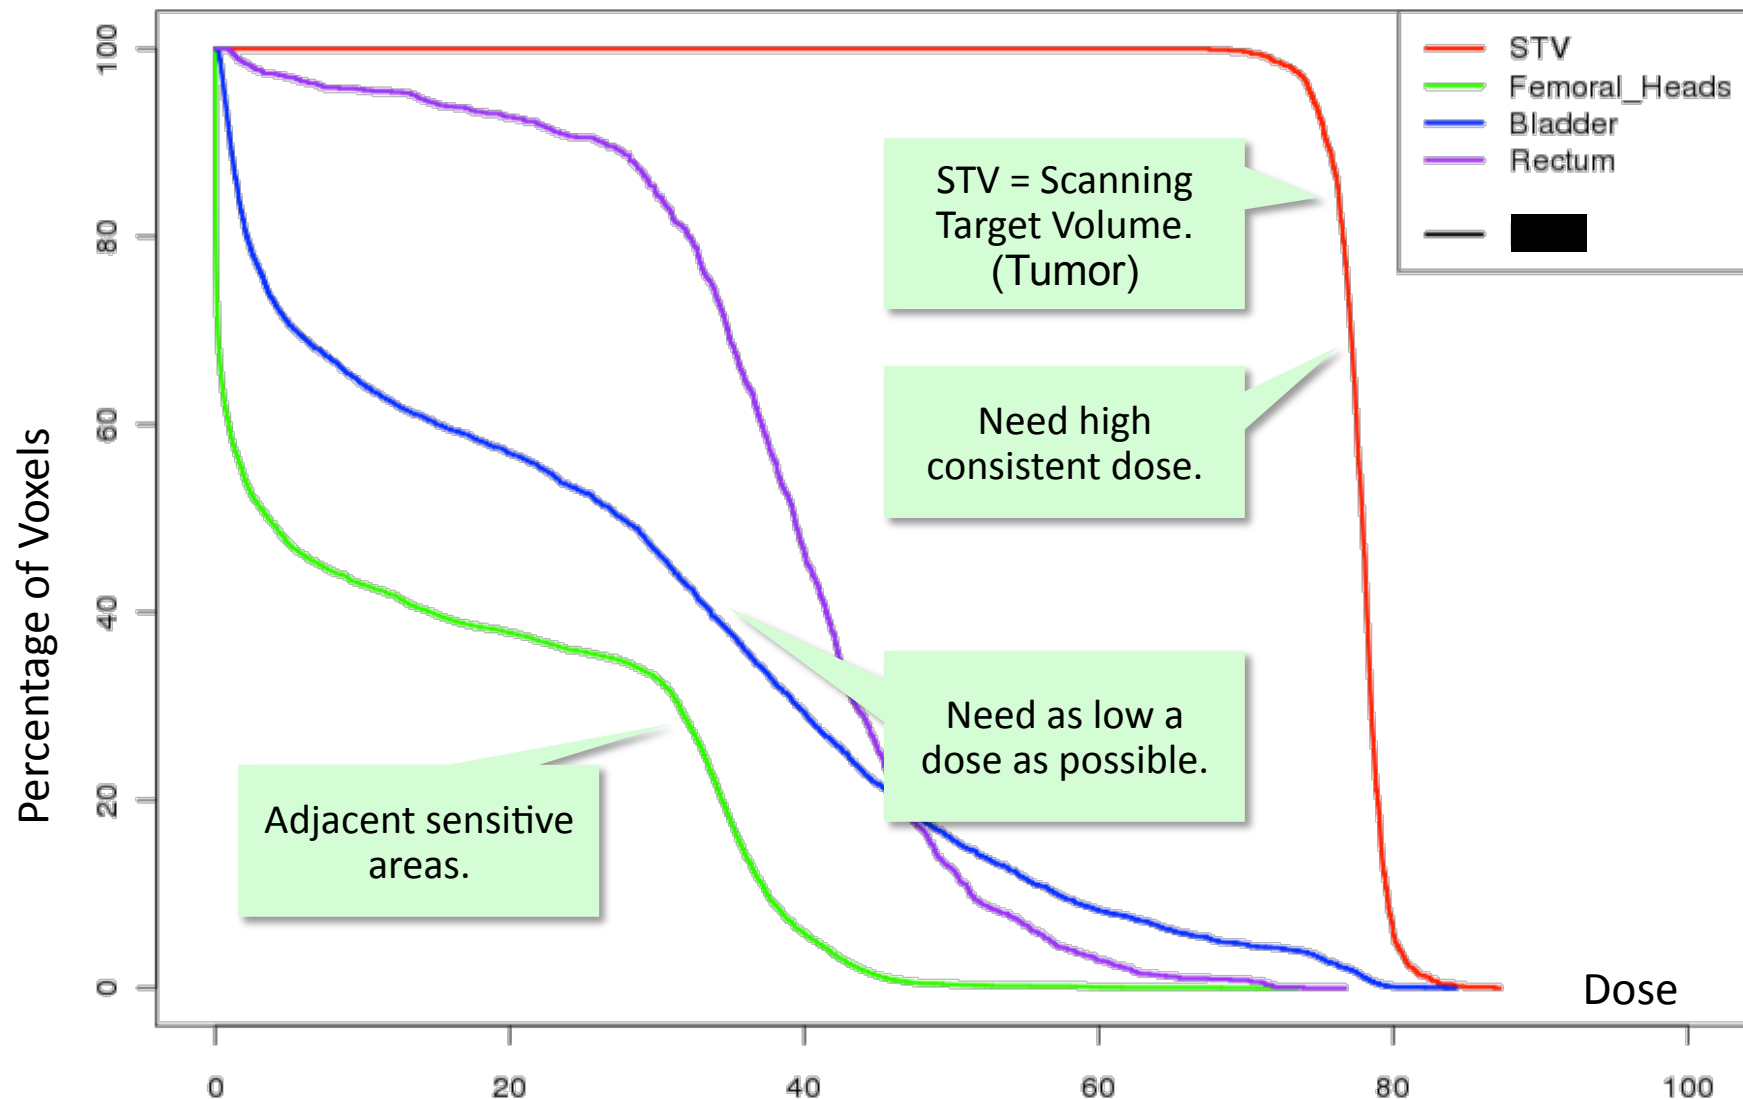


Plan Results: Dose Volume Histogram

- Summarizes optimization result, treatment plan outcome.



Example Results (Prostate Case)...



Where is the Variability?

Positional

- Is the tumor where we think it is.
- Is the patient where she/he is supposed to be.
- Is the tumor moving (because patient is breathing).

Physical

- Are we sure of the composition of each voxel.
- Did our Monte-Carlo particle simulation converge.
- Did we model all the physics necessary.

Biological

- Is the tumor where we think it is.
- Will this tumor react to this amount of dose.
- Will the surroundings react to this amount of dose.

HARD...

WCA of Positional Variability


- Considering a specific voxel at (x,y,z) , the updated LP constraint would be:

$$\sum T_i \text{Dose}_i(x+\Delta X, y+\Delta Y, z+\Delta Z) \leq C(x,y,z)$$

- Just like for circuits:

- Many performances (voxels vs. delays).
- Few stochastic variables with known distributions (ΔX vs. ΔV_T).
- Computable sensitivities ($\delta\text{Dose}/\delta\Delta X$).

- We are developing an analytical optimizer capable of handling positional uncertainty!



The Road Forward

- In this new area, embarking on applying the knowledge learned from VLSI EDA to a very different topic.
- Many of the problems are quite similar...
 - Or perhaps can be cast to use the same underlying mathematical formulations, concepts and resultant solution algorithms.
- As VLSI Research (and funding) slows down with the end of Moore's law and the absence of a economically useful Post-Silicon technology, this phenomena is going to become more common!

From VLSI to Cancer?



- I began my EE education in 1976, and have worked in the area of Design Automation for VLSI since that time.
- I am tired of having to explain what I do to people.
- With some hard work, and an open and collaborative attitude, it is possible to make contributions to other areas!
- VLSI and Design Automation are quite broad subjects, good preparation for working on other problems.
- I can even explain what I do to my Mother now!

