

!

Death, Taxes and Failing Chips

Chandu Visweswariah
IBM Thomas J. Watson Research Center
Yorktown Heights, NY

TAU '02

*with thanks to my colleagues and
collaborators at IBM Fishkill,
IBM Burlington, IBM Yorktown
and TU/Eindhoven*

Outline

- “The era of probabilistic design”
- Three aspects of the problem
 - modeling
 - methodology
 - analysis + synthesis
- Characteristics of a good statistical timer
- Our analysis efforts

Catastrophic vs. parametric

Chip behavior in the face of environmental and manufacturing variations

Catastrophic yield loss

Critical area
Voronoi diagrams
Redundant via insertion
Wire bending/spacing


Parametric or “circuit-limited” yield loss

Statistical timing
Yield prediction
Design centering
Design for manufacturability

Digital ASICs

Bounded vs. probabilistic analysis

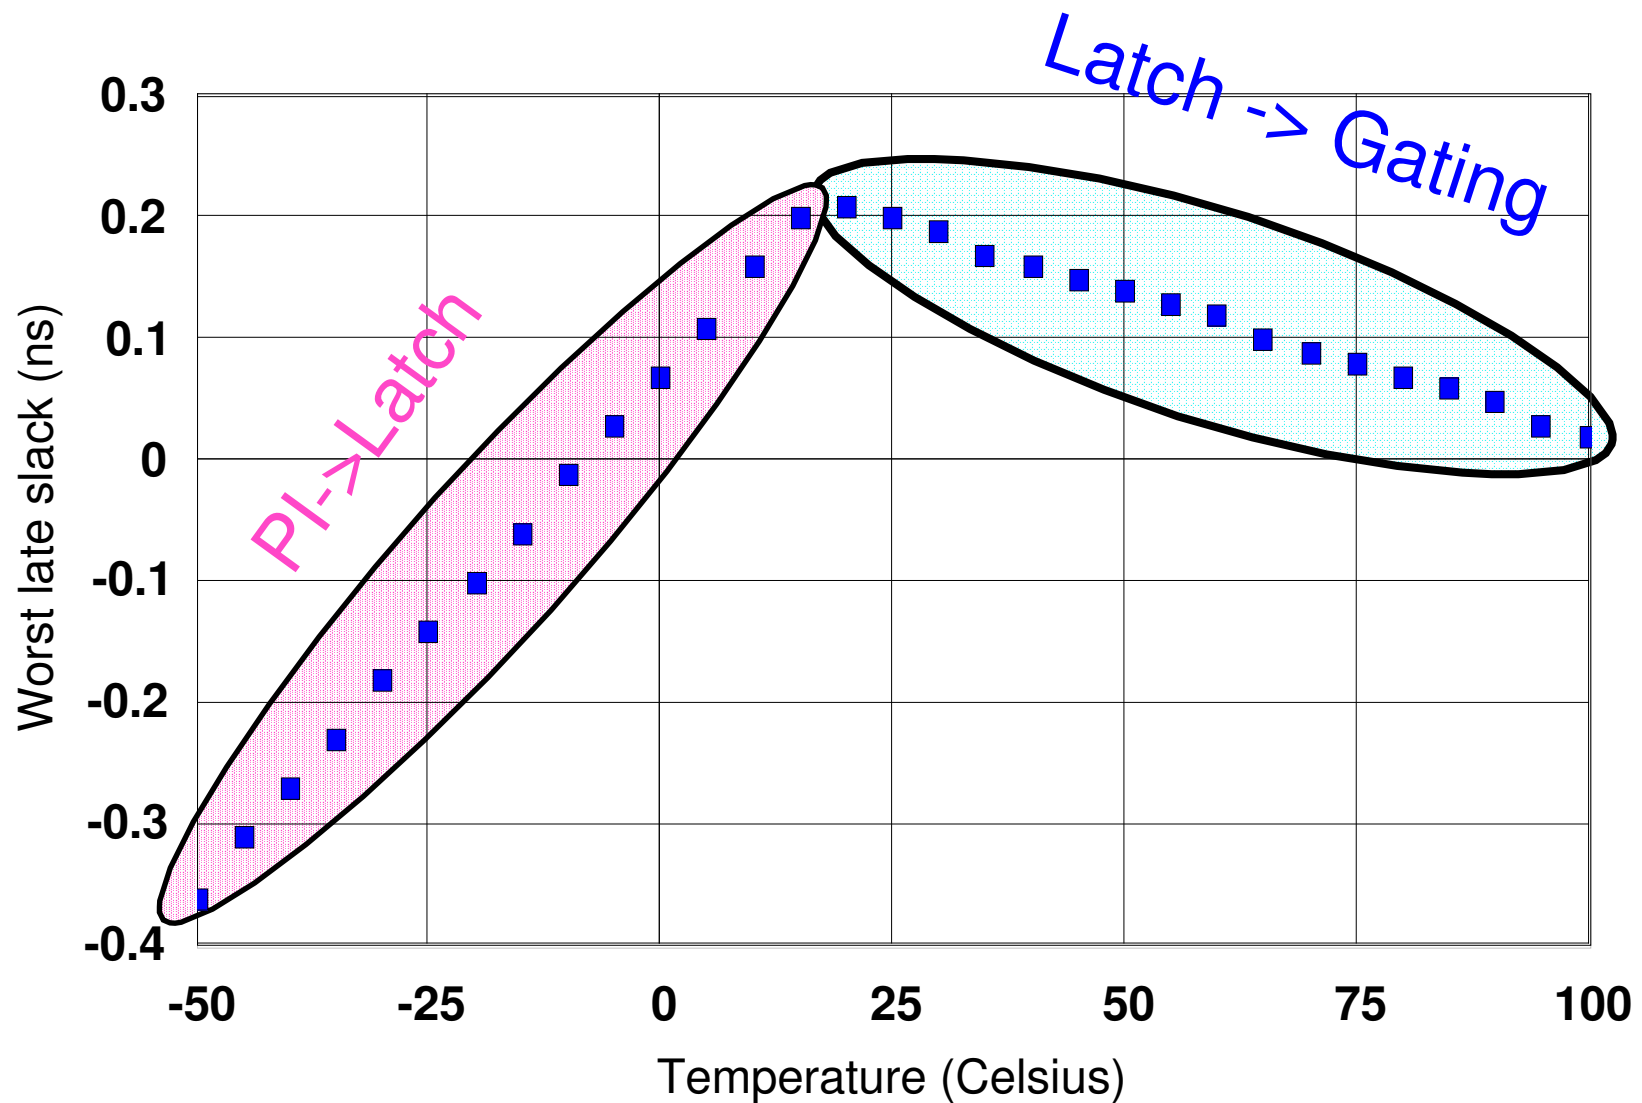
Varying temperature, mean = 25, sigma = 25



[Data courtesy K. Kalafala] BC
© C. Visweswariah, 2002, do not use without permission

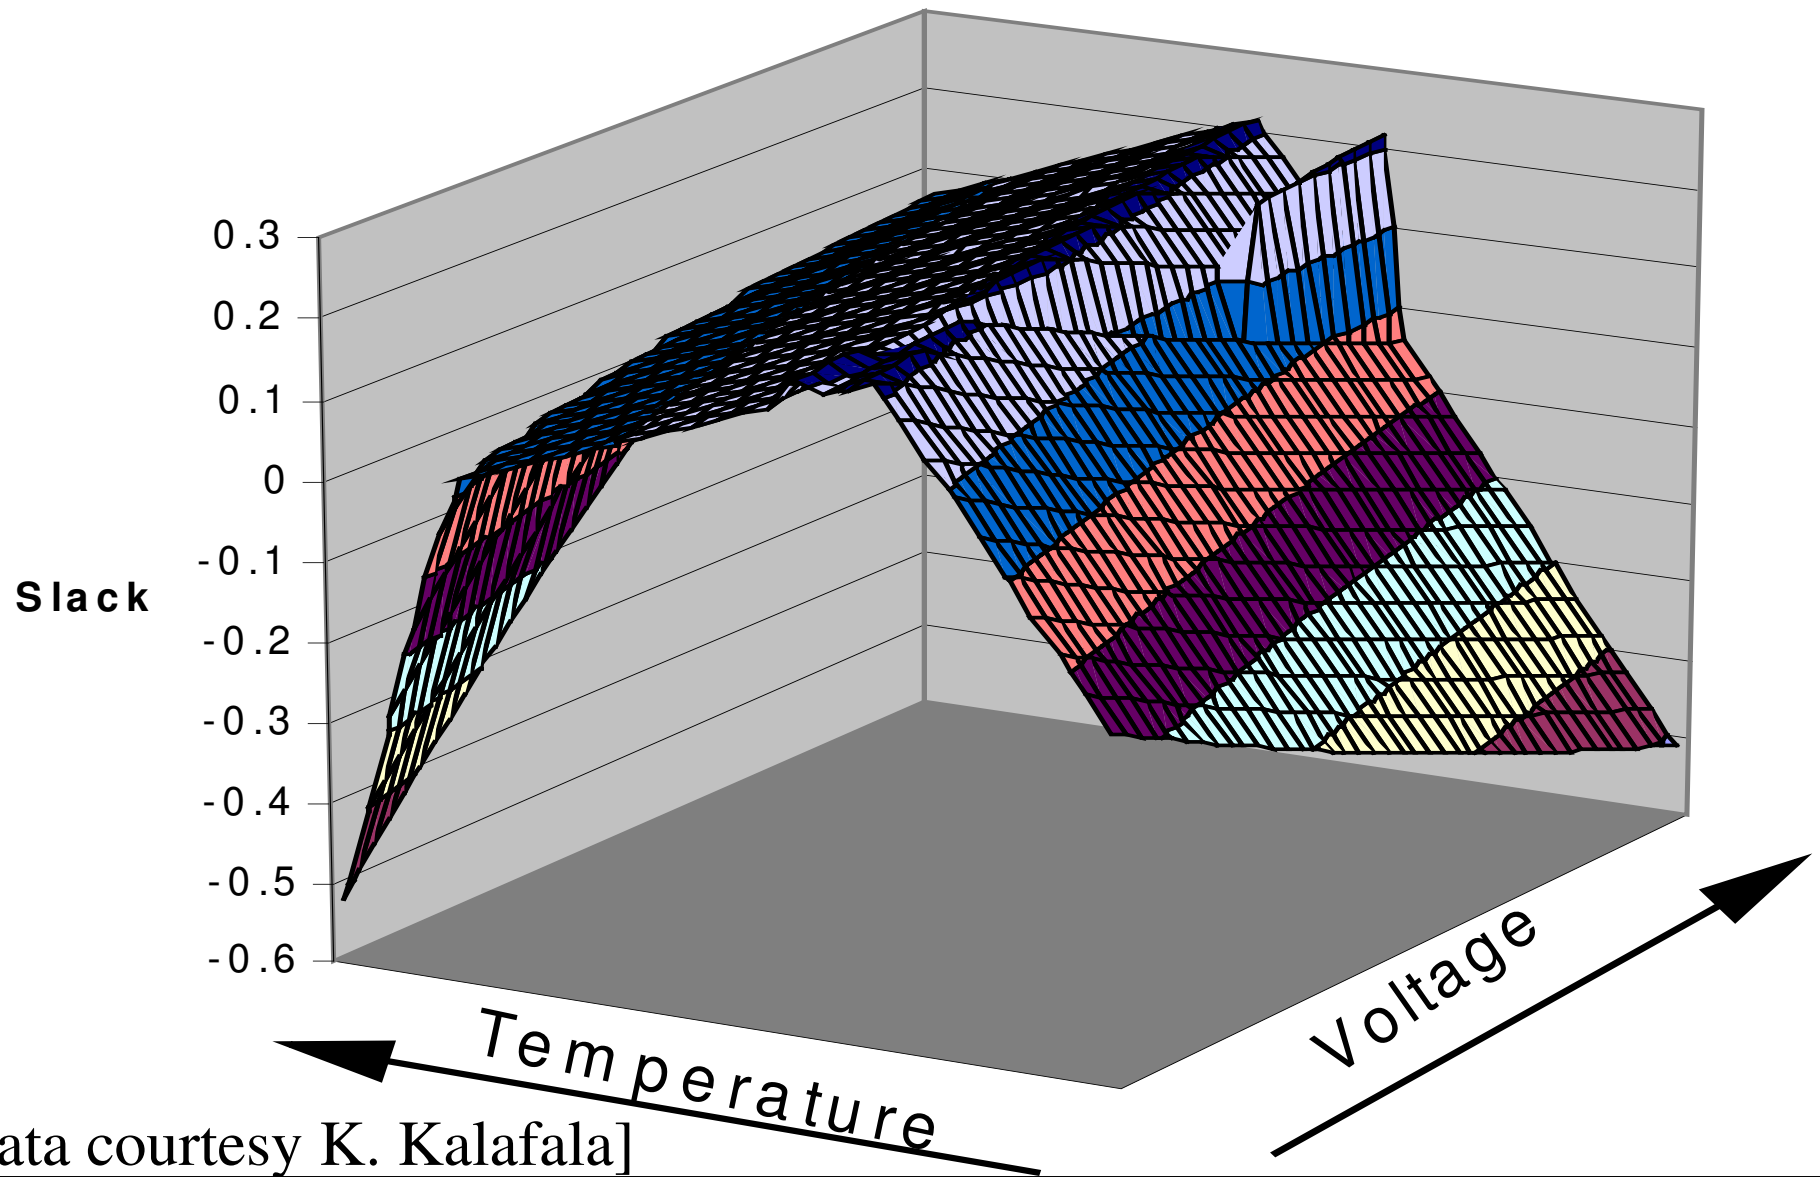
WC NOM

Bounded vs. probabilistic analysis



[Data courtesy K. Kalafala]
© C. Visweswariah, 2002, do not use without permission


Bounded vs. probabilistic analysis



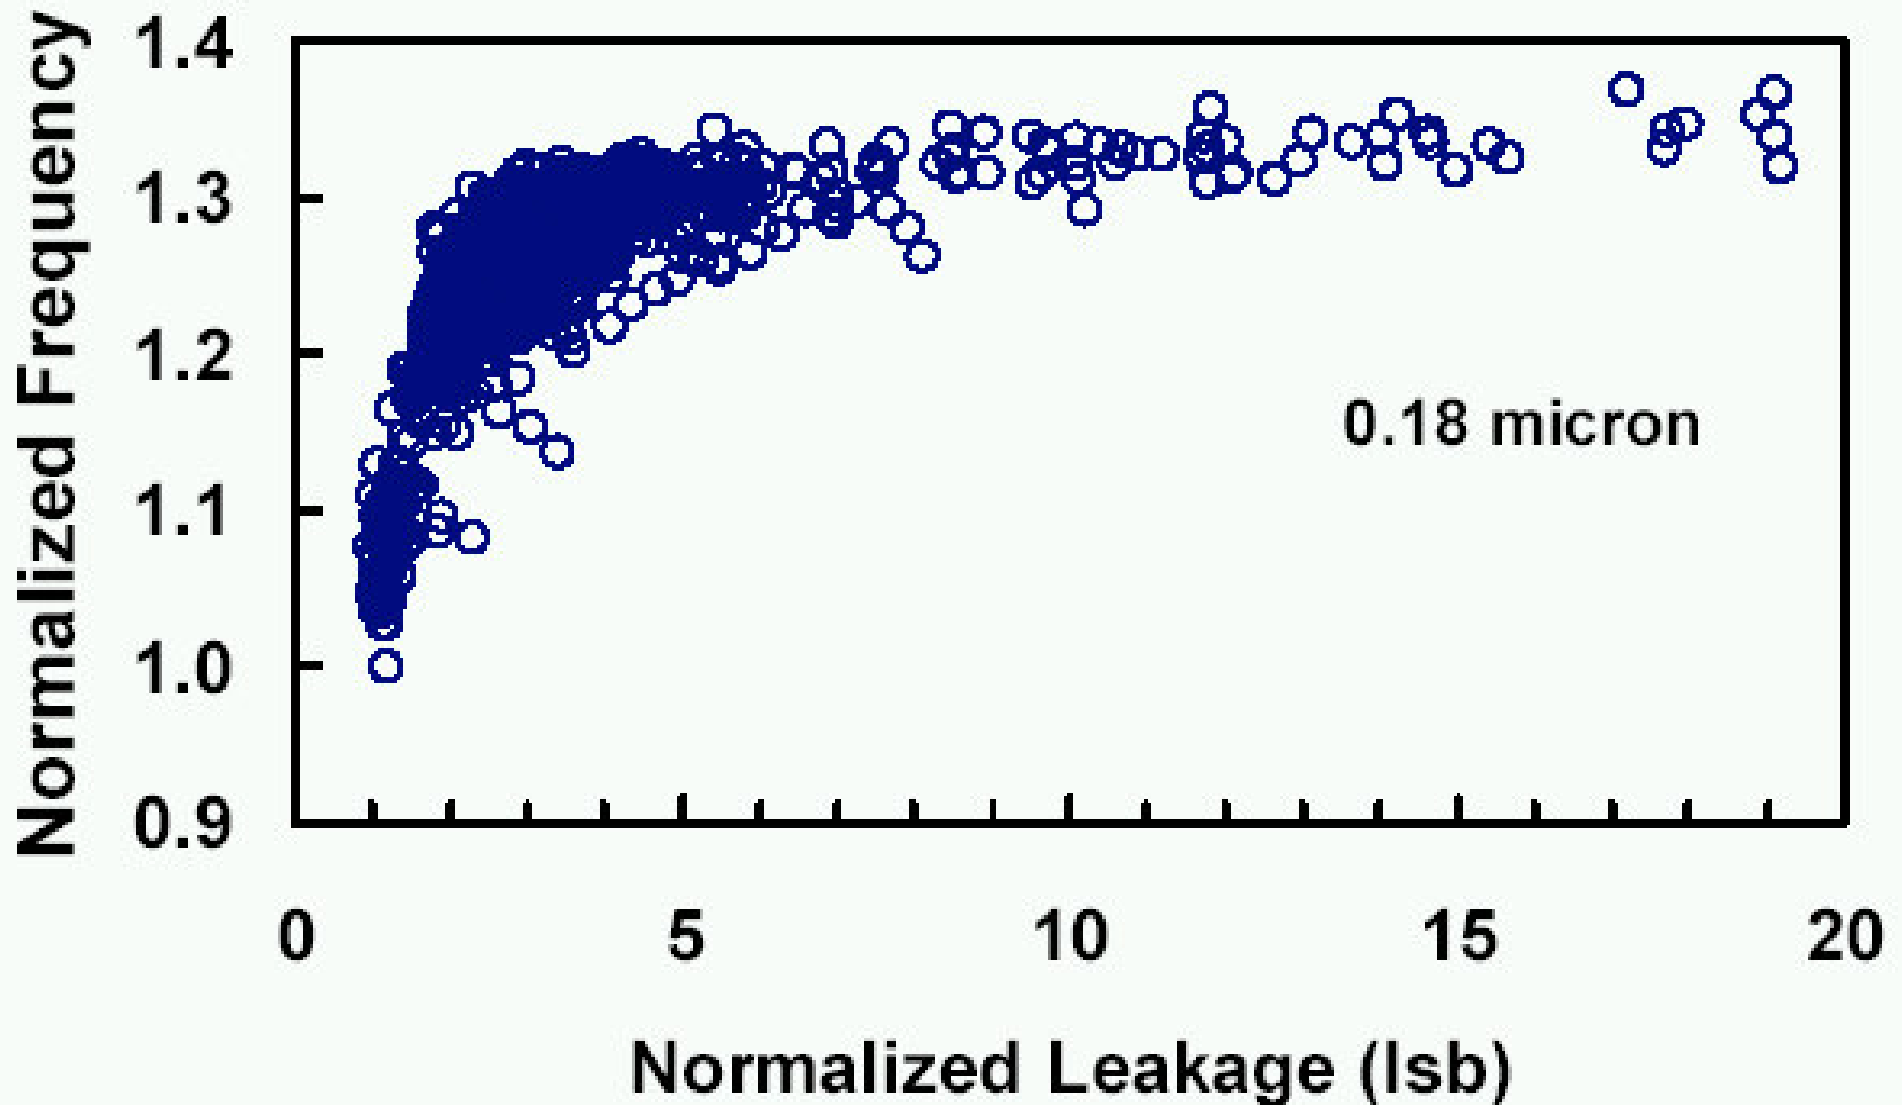
[Data courtesy K. Kalafala]

© C. visweswaran, 2002, do not use without permission

Bounded vs. probabilistic analysis



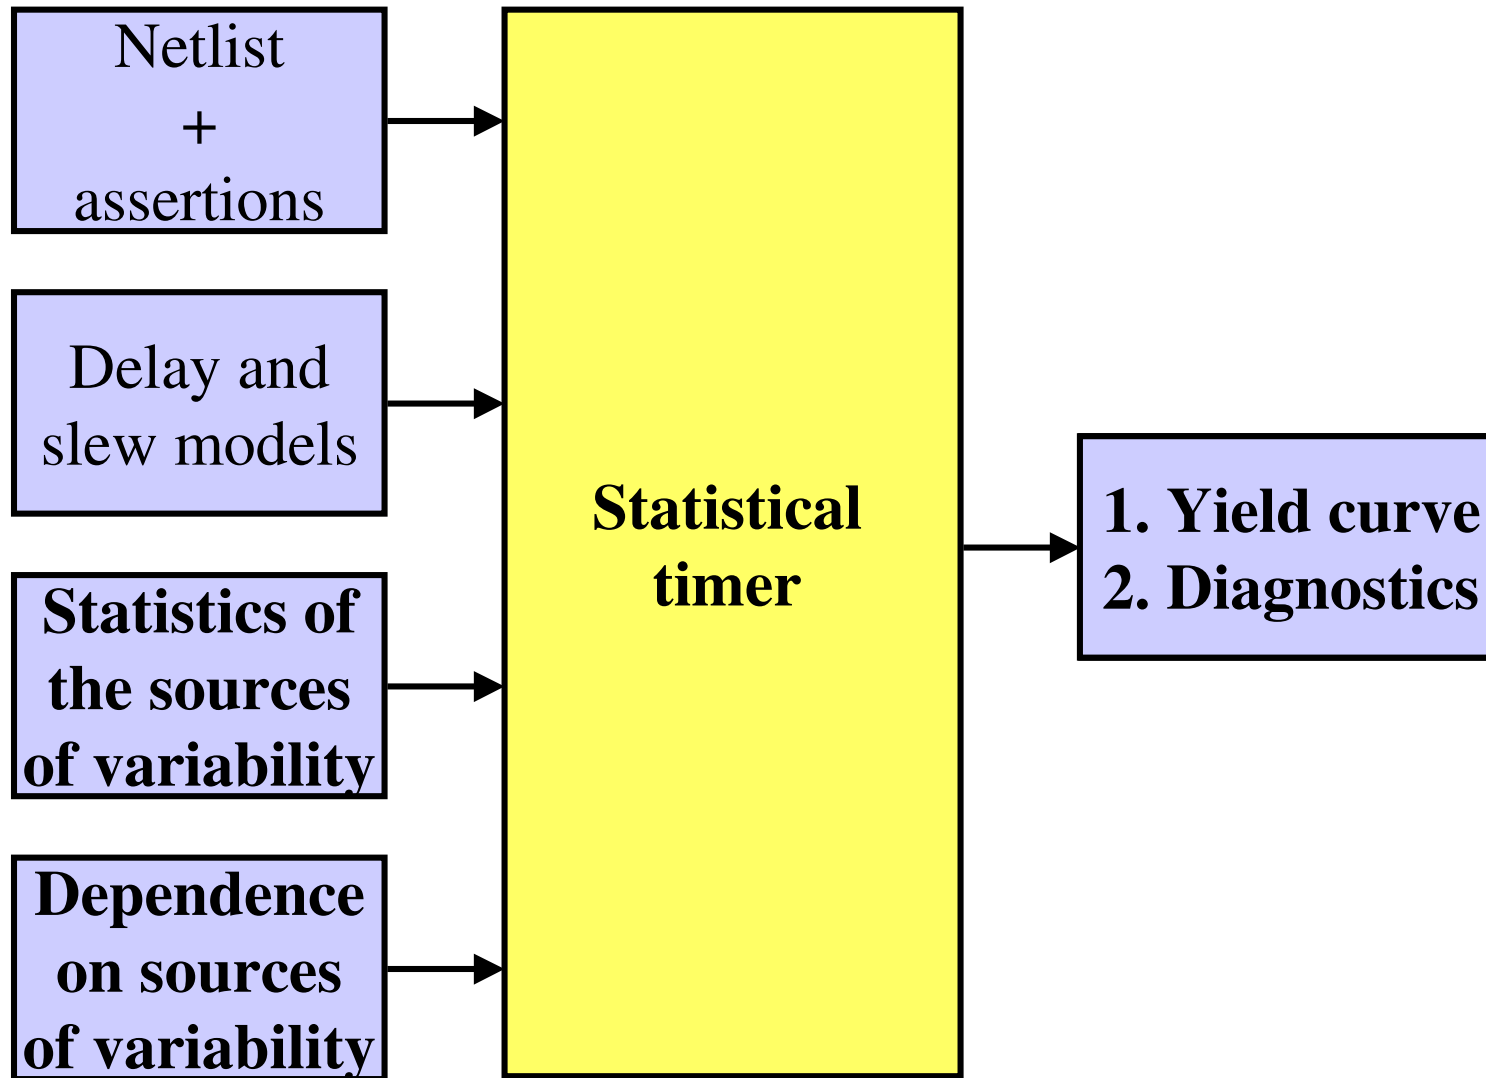
The era of probabilistic design




[T. Karnik, S. Borkar, V. De, ICCAD 2002]

© C. Visweswariah, 2002, do not use without permission

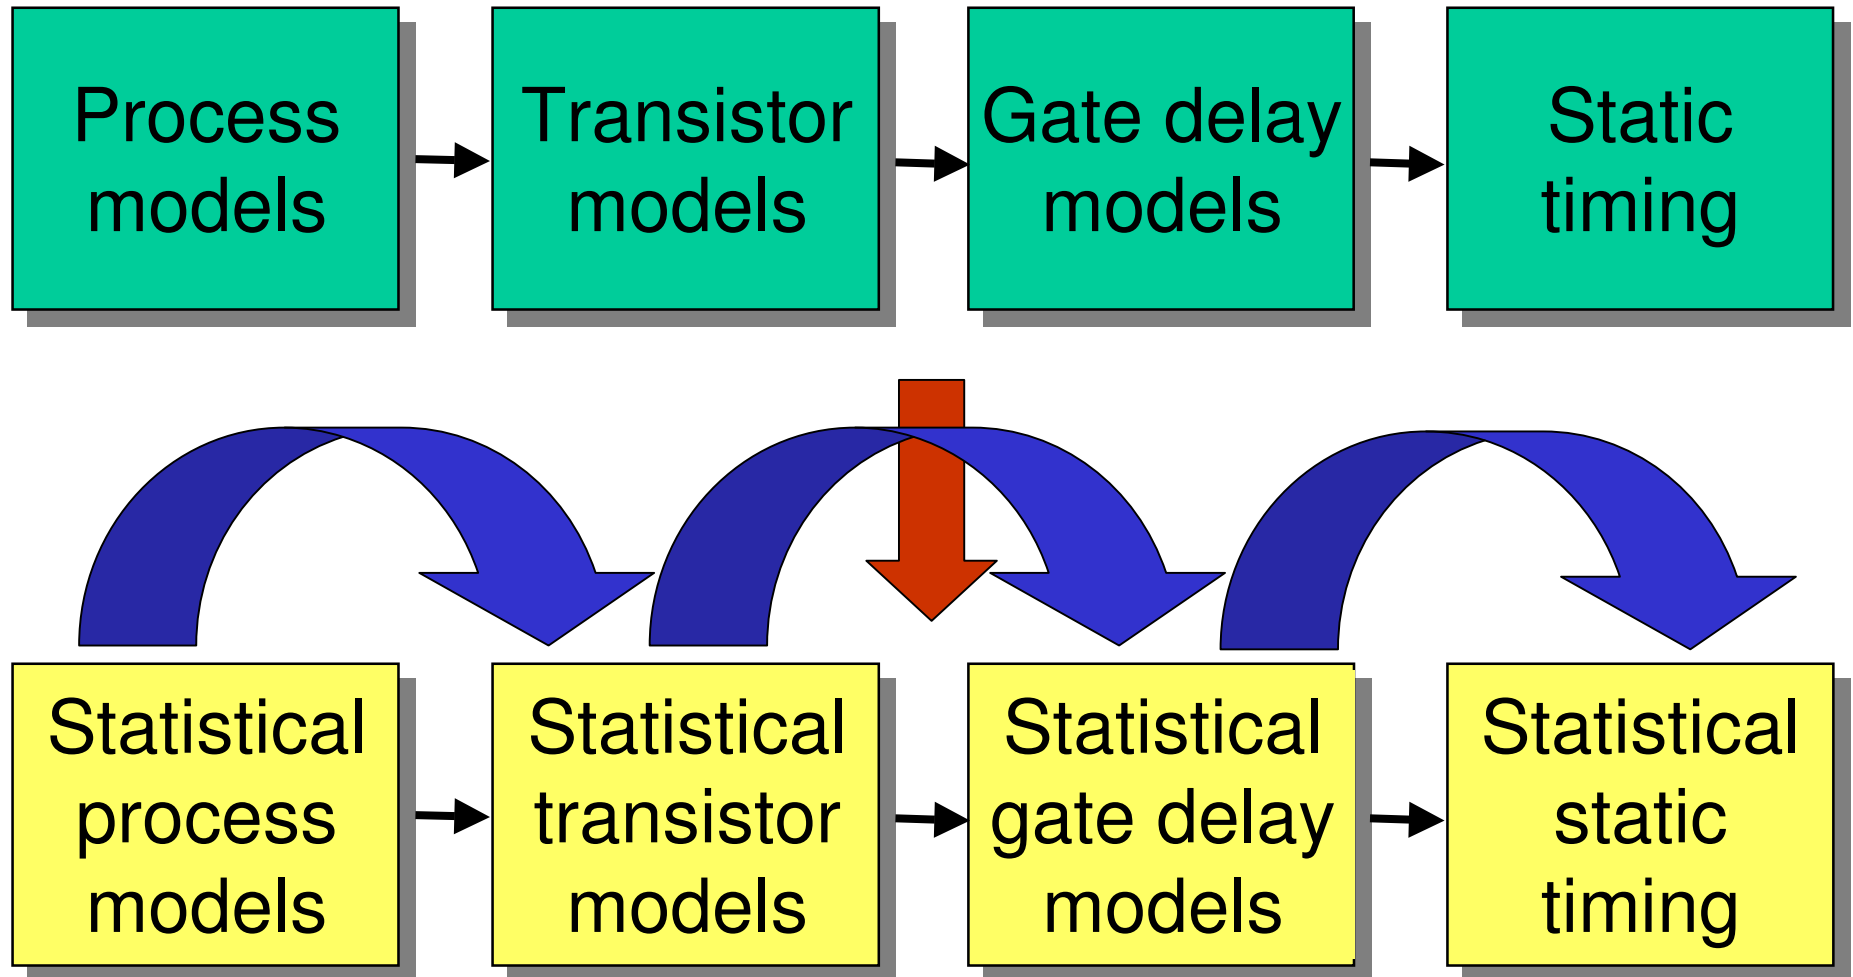
Statistical timer



The full picture



Methodology



Methodology issues

<i>ASIC</i>	<i>Microprocessor</i>
<i>No at-speed test, often no AC test</i>	Sorted
<i>Large, flat</i>	Hierarchical
<i>Library-based</i>	Custom circuits and library-based
<i>Focus on worst-case timing</i>	Focus on nominal (and best case!) timing

Definition of yield


- Risk management
 - with PSROs (Performance-Sensitive Ring Oscillators) and appropriate sign-off criteria
 - at multiple levels
 - with/without AC or at-speed test
- Environmental vs. manufacturing variations
 - require 100% yield in environmental window
 - guaranteed 100% yield in the manufacturing window is overkill

Other methodology implications

- Number of timing runs is excessive
 - early and late mode
 - LCD (Linear Combination of Delay) or “interval delay” to model ACLV
 - CPPR (Common Path Pessimism Removal)
 - NBTI (Negative Bias Temperature Instability)
 - BEOL variations
 - coupling noise

What is required is a reduction in the number of timing runs. **Opportunity!** New analysis modes into the current methodology!

Design methods



- Examples:
 - adaptive body bias
 - mixing of logic families

Modeling

- What are the sources of variation that really matter?
 - mathematical vs. empirical answers
- What are the means, deviations and correlations of the sources of variation?
- What is the dependence of the delay and slew of each edge of the timing graph to each source of variation? Is this computed during the library characterization?
- What about custom circuits?

Analysis wish list


**W
i
s
h

l
i
s
t


t
o
p

1
0**

Number 1: path sharing



Number 2: clock correlation



**W
i
s
h

l
i
s
t


t
o
p

1
0**


Number 2: clock correlation

- Importance of correlations
 - consider a circuit with 50K latches, each with a setup and hold test, each of which has a 99.99% probability of being met
 - if all tests are perfectly correlated, yield=99.99%
 - if all tests are perfectly independent, yield is 0.005%
 - the truth is closer to the perfectly correlated case!

Number 3: global correlation



Number 3: global correlation



W Number 4: bounded vs. statistical

i

s

h

/

i

s

t

t

o

p


1

0

- Bounded
 - input vectors
 - environmental variables
 - PLL jitter
- Statistical
 - manufacturing parameters
 - coupling noise?
- Should be easy to switch between columns
- Large vs. small number of random variables


W i s h l i s t t o p 1 0

Number 5: slew/load dependence



**W
i
s
h
l
i
s
t
t
o
p
1
0**

Number 6: deterministic vs. random ACV




[From M. Orshansky, L. Milor, P. Chen, K. Keutzer, C. Hu, ICCAD 2000]

© C. Visweswarian, 2002, do not use without permission

**W
i
s
h
l
i
s
t
t
o
p
1
0**

Number 7: the tail matters!



- Avoid pessimism
- Capture correlations

**W
i
s
h

l
i
s
t

t
o
p

1
0**

Numbers 8 and 9

- Number 8
 - fit well with rest of existing methodology
 - reduce number of timing runs required
- Number 9
 - provide diagnostics

**W
i
s
h

l
i
s
t

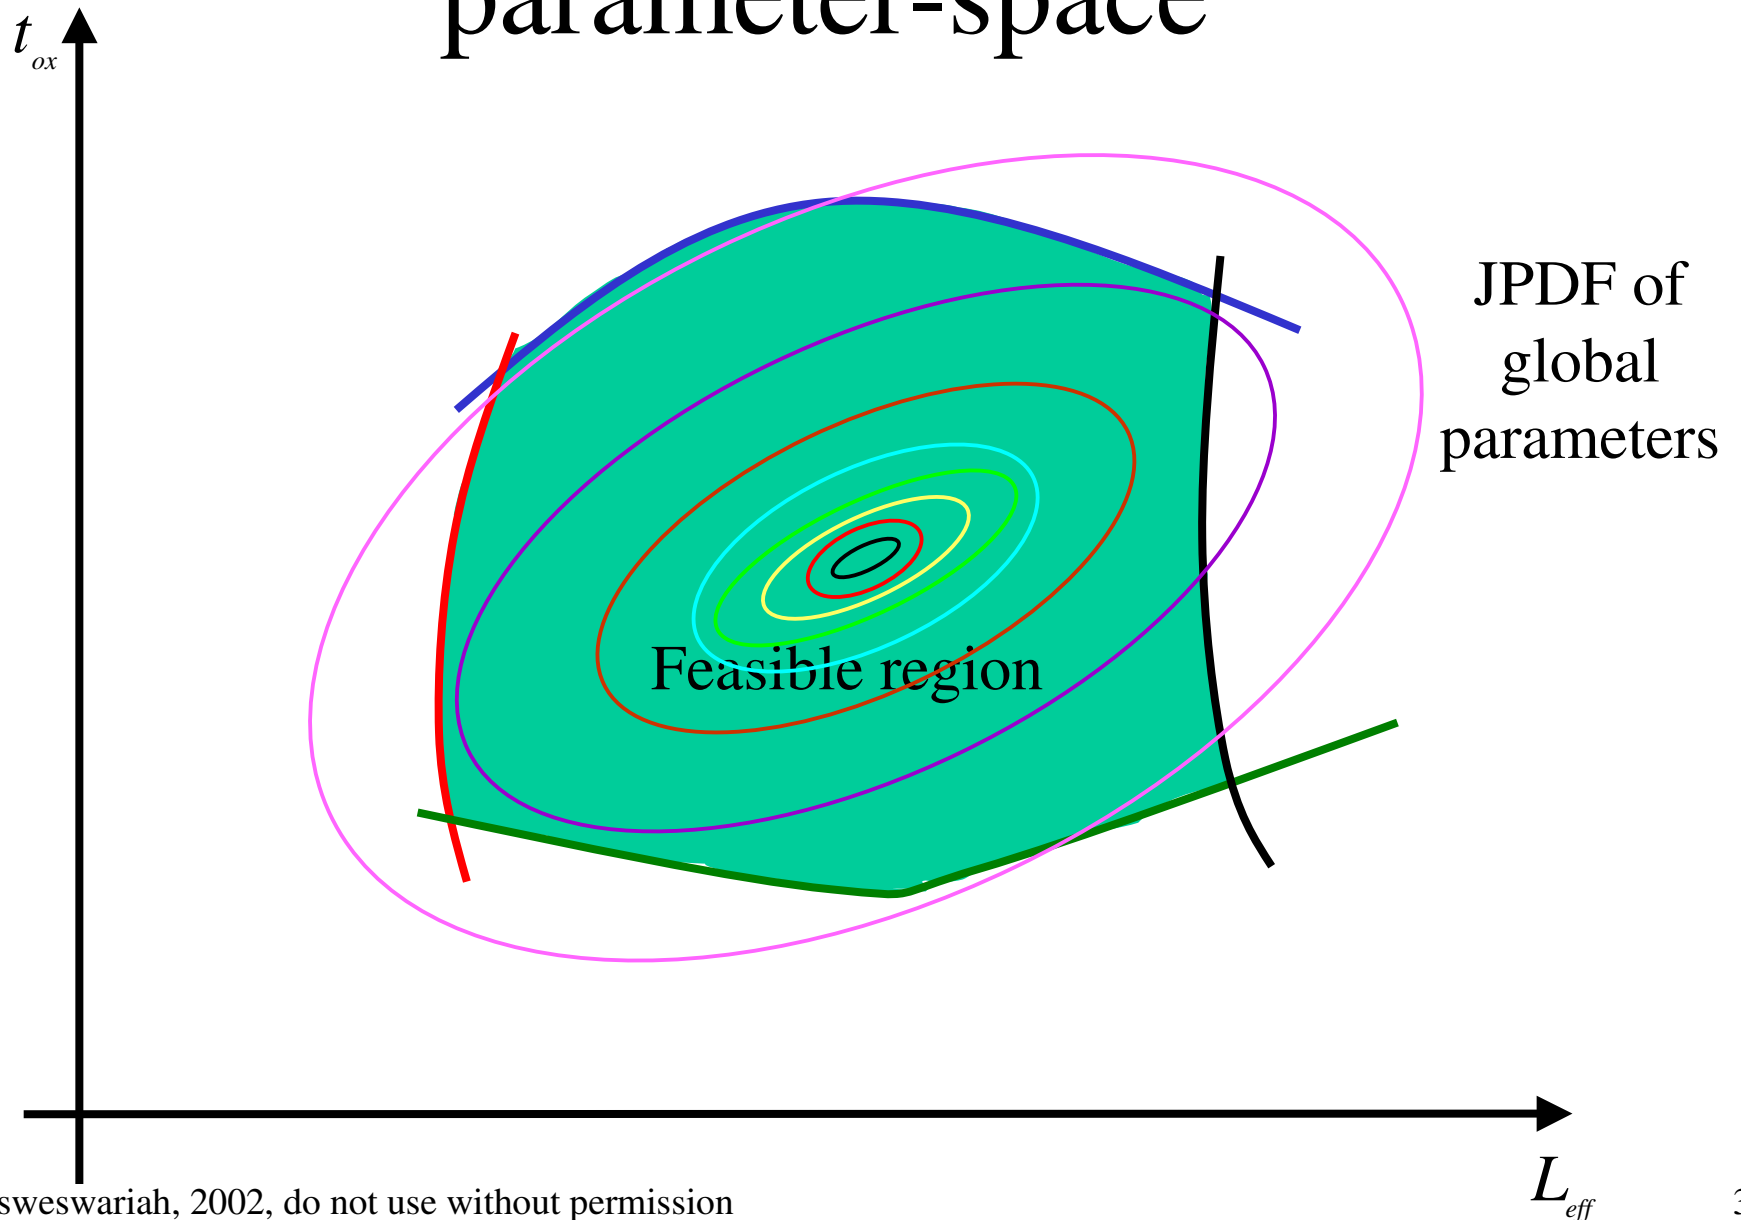
t
o
p

1
0**

Number 10: flexibility

<i>Quick and dirty</i>	Slow and accurate
<i>For optimization</i>	For sign-off
<i>Incremental</i>	Not incremental
<i>Usually block-based, performance-space methods</i>	Usually path-based, parameter-space methods

Performance-space vs. parameter-space



Our analysis efforts

- Three slides deleted
- See DAC '03 submission for details
- Example: reduced run time from 68 hours for repeated EinsTimer runs on a 200K gate ASIC to about 15 minutes

Conclusions

- Brave old world of probabilistic design
- Statistical considerations must influence all stages of design
- Comprehensive solution required encompassing methodology, modeling, analysis, synthesis, test, design methods
- The computation will not prove to be the hard part; if nothing else, Monte Carlo with intelligent sampling will come to the rescue