Clock Sync. and Adversarial Fault Tolerance

Christoph Lenzen – MPI for Informatics

Danny Dolev – Hebrew U. of Jerusalem

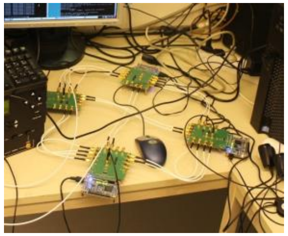

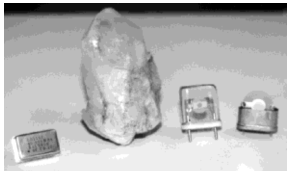

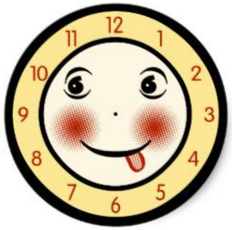

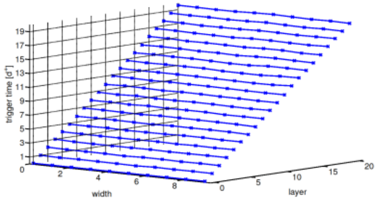

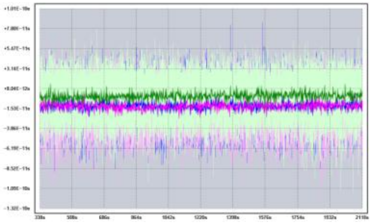

also starring: Ben Wiederhake  , Matthias Függer 

Today's Menu

1. Why does this course exist?
2. What is this course about?
3. Who are you and what do you want?
   - discussion in small groups
   - sharing your findings with everyone
4. How will we run this course?
   - your questions and input on this
5. Heads-up: What comes next?

Today’s Menu

1. **Why does this course exist?**
2. What is this course about?
3. Who are you and what do you want?
   - discussion in small groups
   - sharing your findings with everyone
4. How will we run this course?
   - your questions and input on this
5. Heads-up: What comes next?

Chips are Distributed Systems

- very large (>10^{10} transistors)

-> **fault-tolerance mandatory**

- highly concurrent/para

-> **synchron**

-

->

We should treat them as distributed systems!

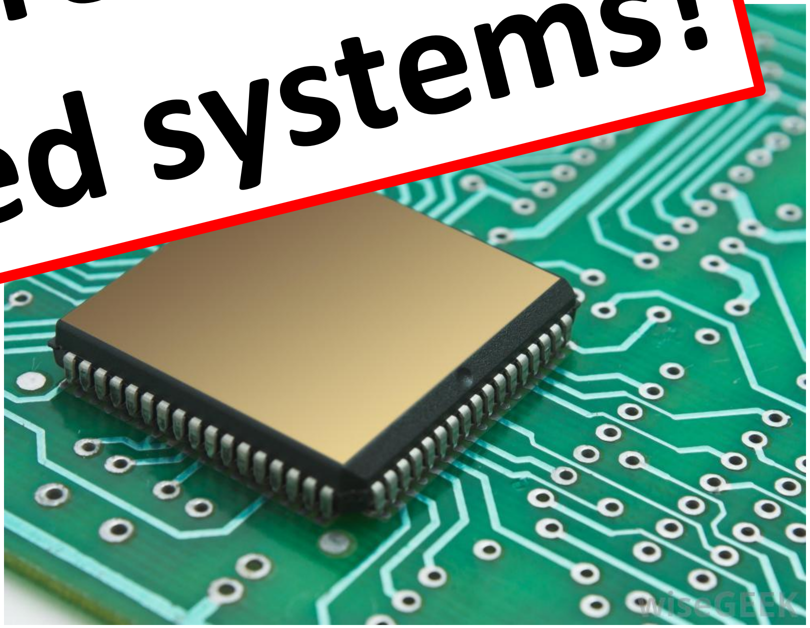

Chips are Distributed Systems

- very large (>10^{10} transistors)

-> **fault-tolerance** mandatory

- **very fast** (>10^{9} cycles/s)

-> communication “slow”

- highly concurrent/parallel

-> **synchronous** operation

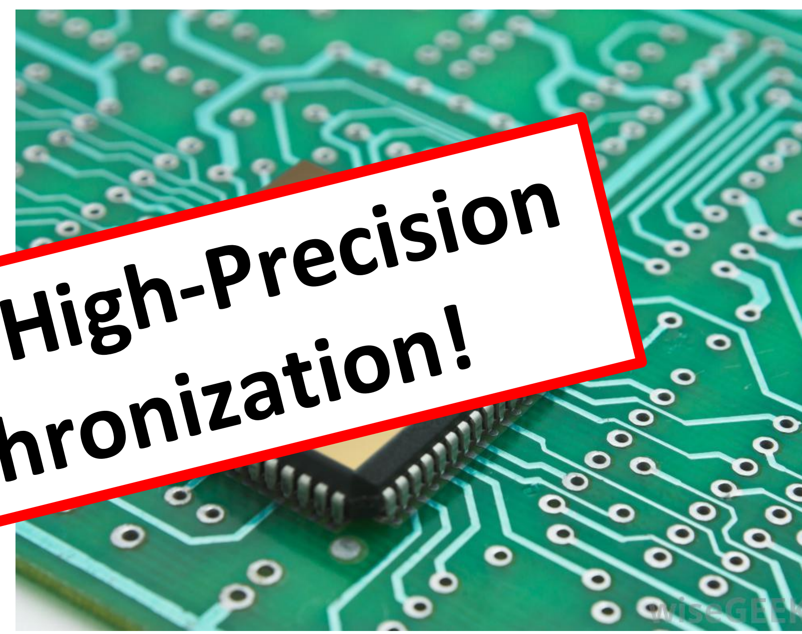

Fault-tolerant High-Precision Clock Synchronization!

Clocking VLSI Circuits

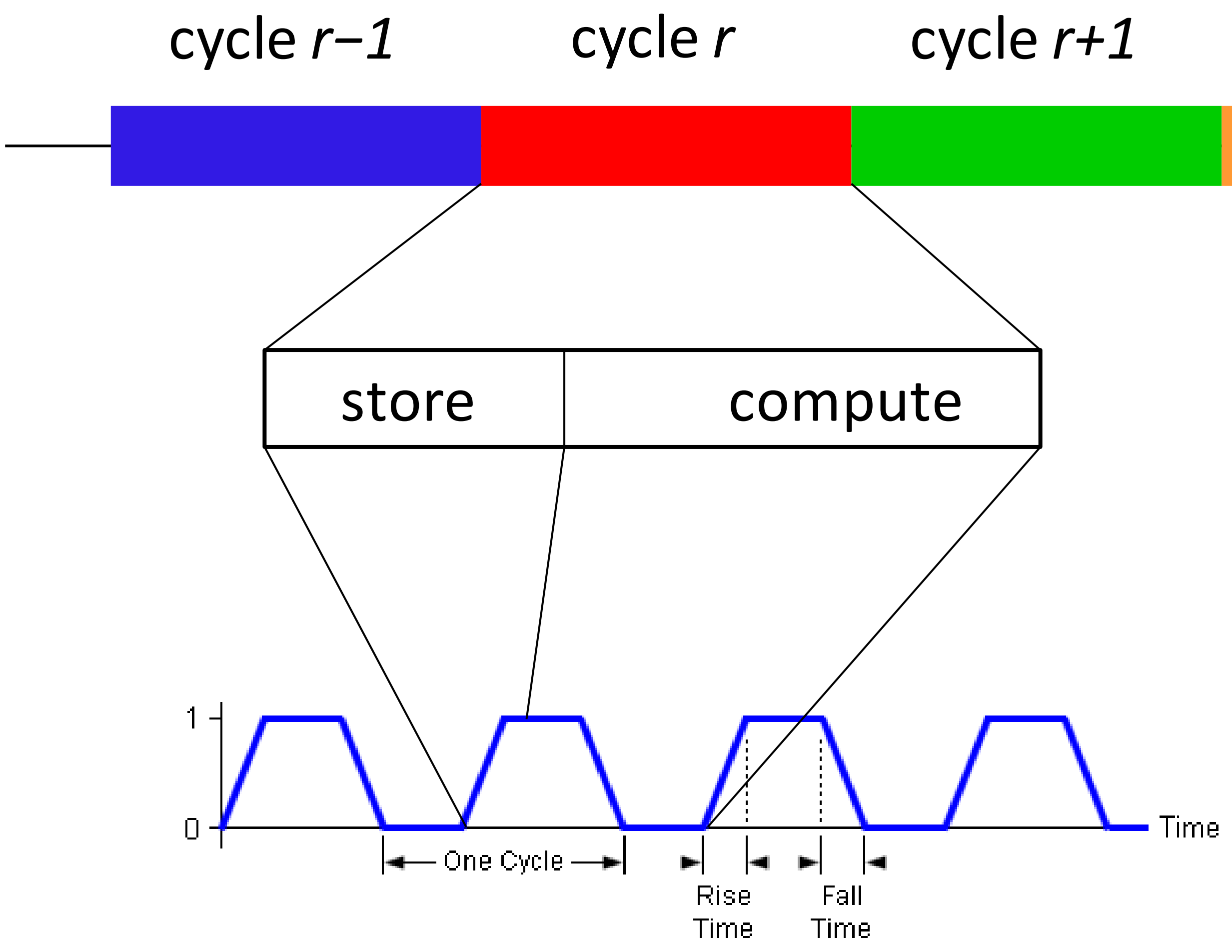

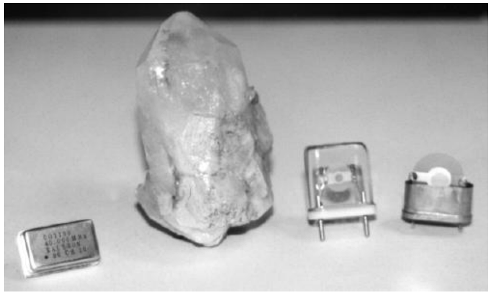

Clock Trees

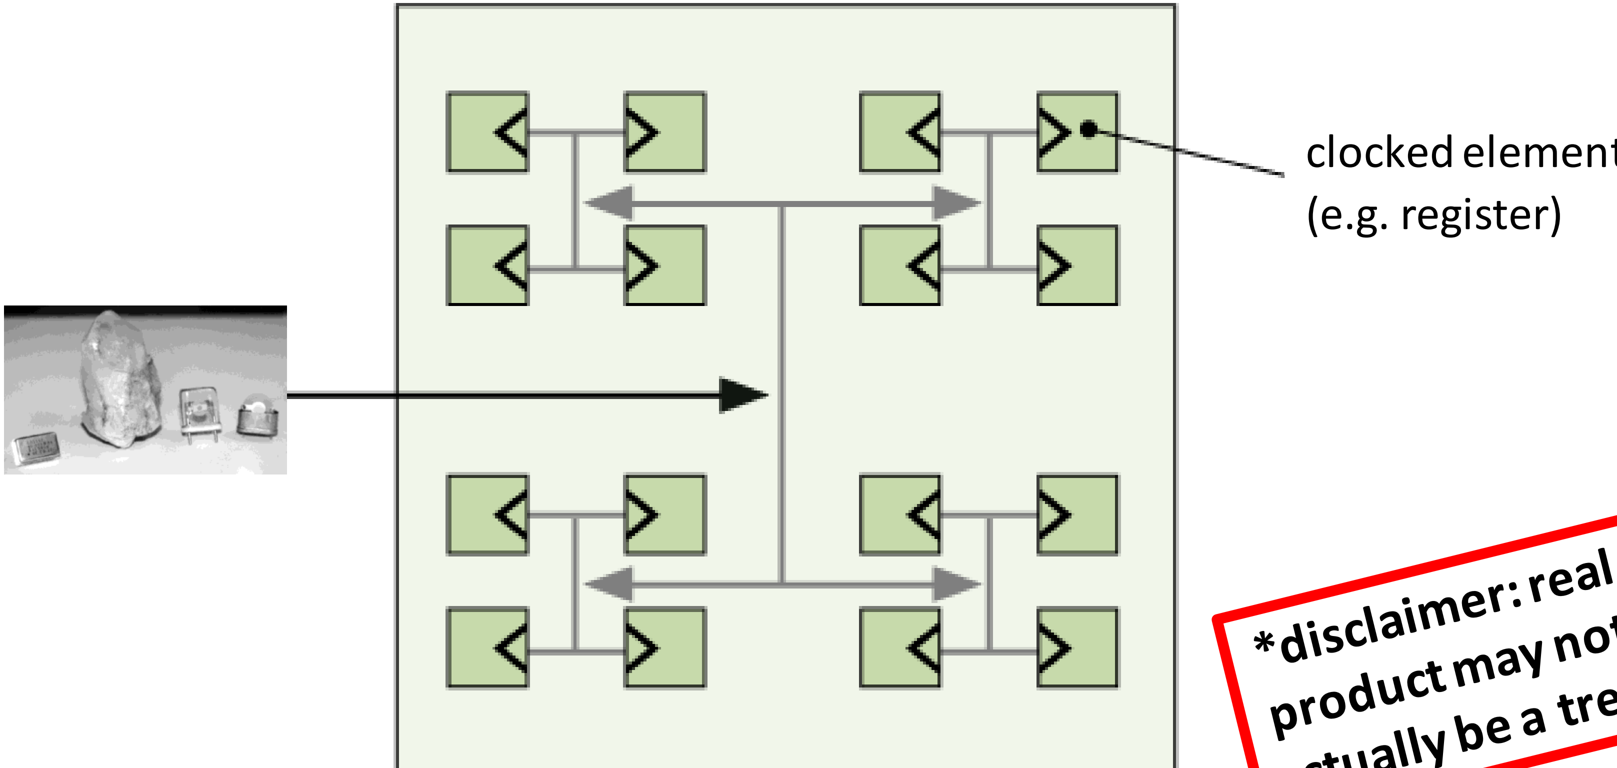

*disclaimer: real product may not actually be a tree

Distribute clock signal from single source!

+ very simple

+ self-stabilizing: recovers from **any** transient faults

+ ca. 20 ps = $2*10^{-11}$s precision (single chip)

Clock Trees: Scalability Issues

- clock tree is single point of failure

-> components **must be extremely reliable**

- tree dist./physical dist. = Ω(L) (L side length of chip)

-> max. difference of arrival times between adjacent gates grows linearly with L

-> clock **frequency goes down with chip size**

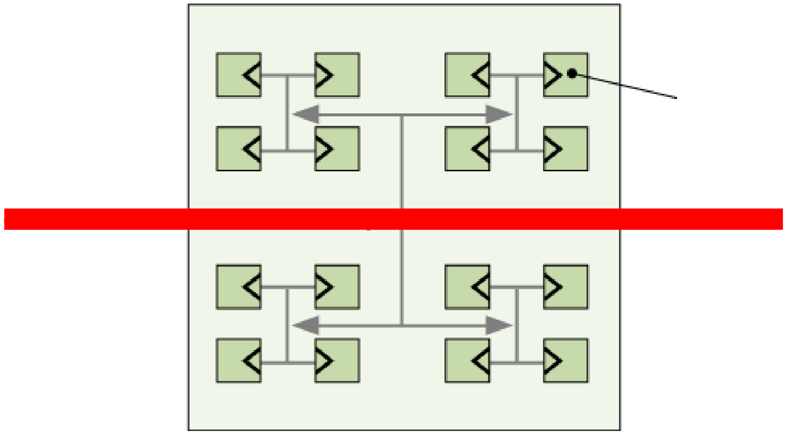

Clock Trees: Scalability Issues

- clock tree is single point of failure

-> components **must be extremely reliable**

- tree dist./physical dist. = $\Omega(L)$ (L side length of chip)

-> max. difference of arrival times between adjacent gates grows linearly with L

-> clock **frequency goes down with chip size**

- countermeasure: use higher voltage and wider wires

-> **electro-magnetic interference** causes trouble and strong currents induce large **power consumption**

GALS: Globally Sync., Locally Async.

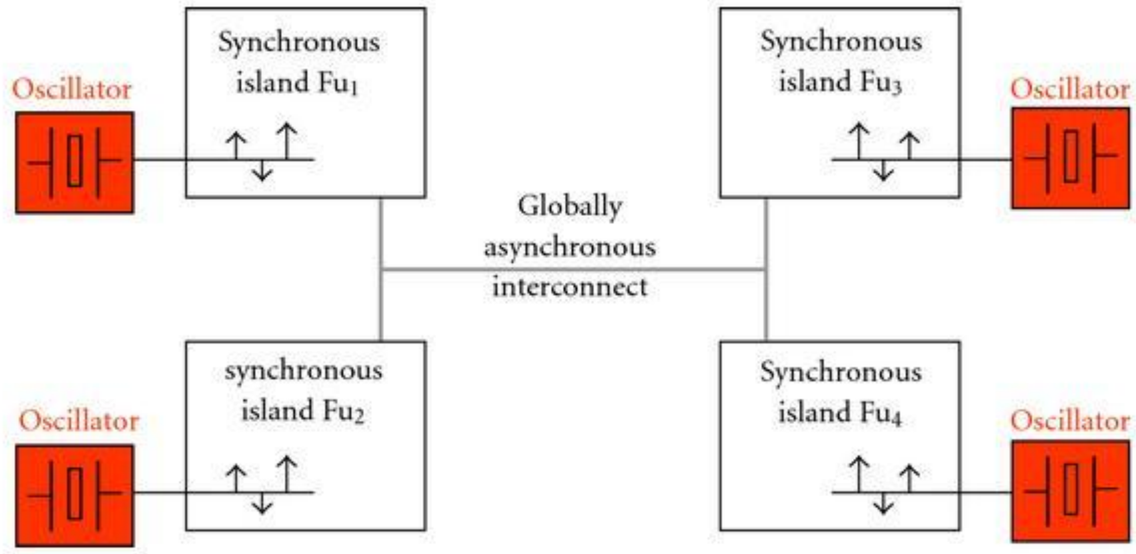

GALS: multiple separately clocked subsystems communicate asynchronously

\+ removes *some* clock tree scalability issues

\- asynchronous communication risks **metastability**

-> use of synchronizers, several clock cycles latency

What happens if we do

Computer Science

to it?

Scalable Clocking: Gradient Clock Sync

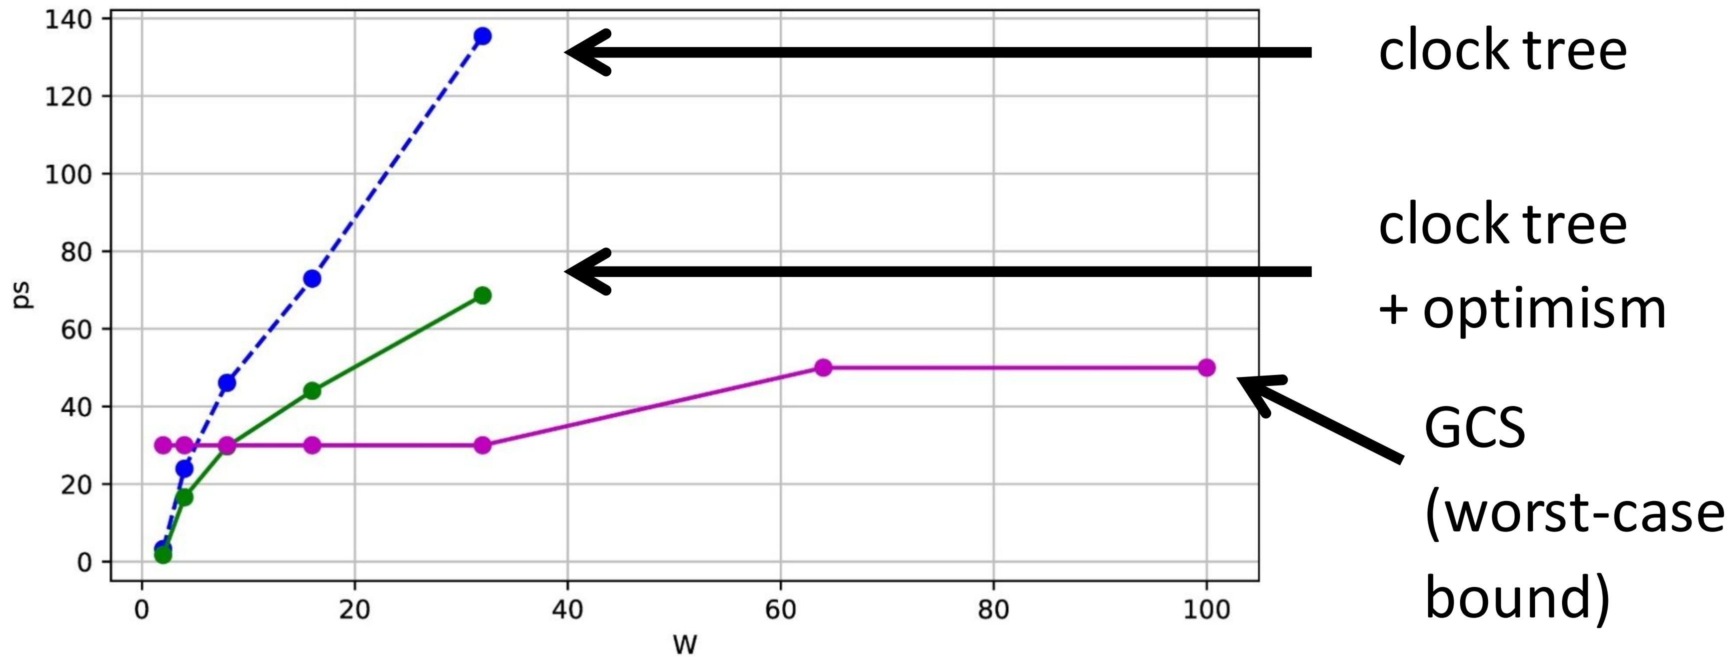

Synchronize along data flow!

=> bound skew between **communicating** components

Fault-Tolerance

direction of propagation

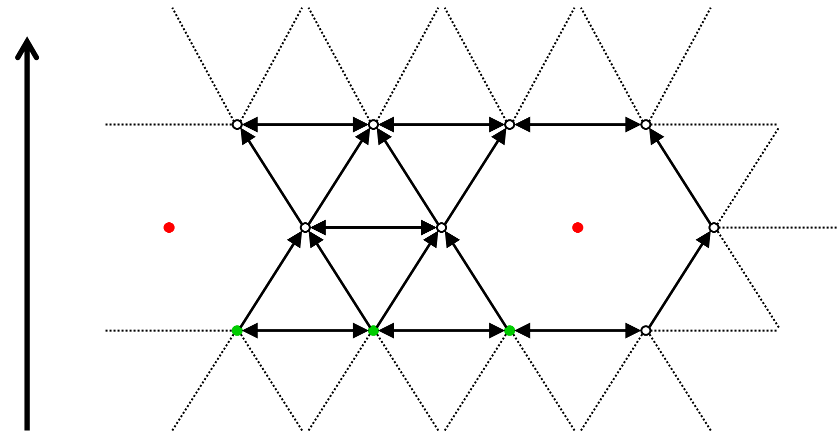

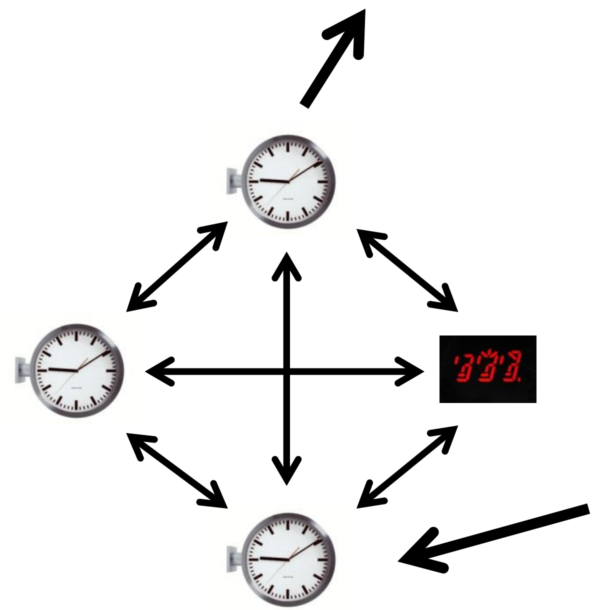

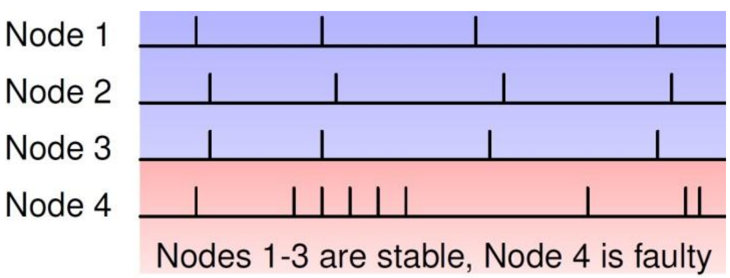

- redundancy enables tolerating (worst-case!) faults
- low-degree distribution networks needed

Innocent "Theory" Assumption

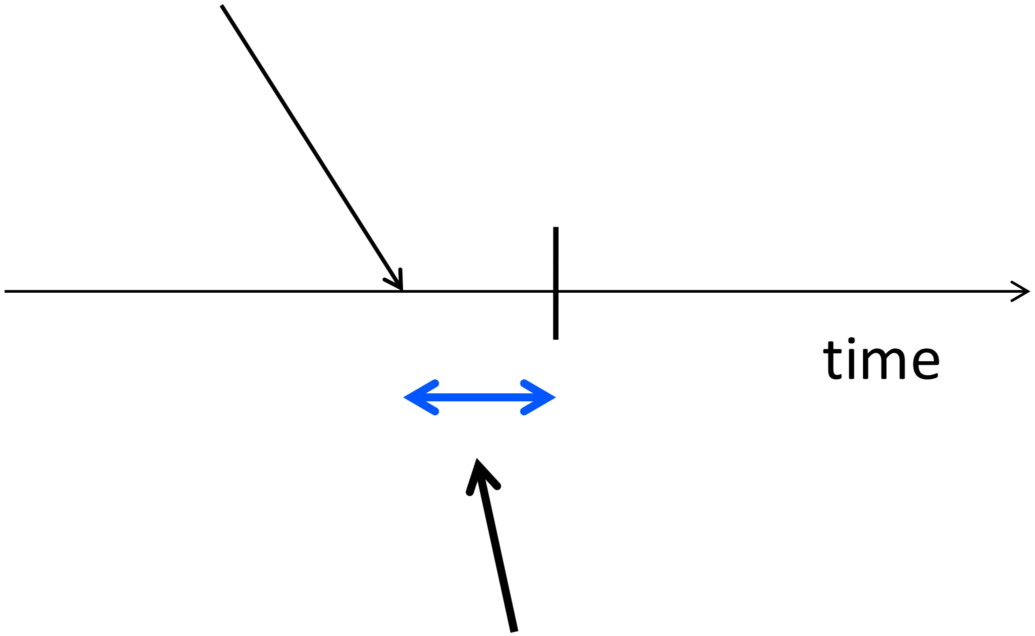

time difference can be turned into a **discrete number**

Metastability

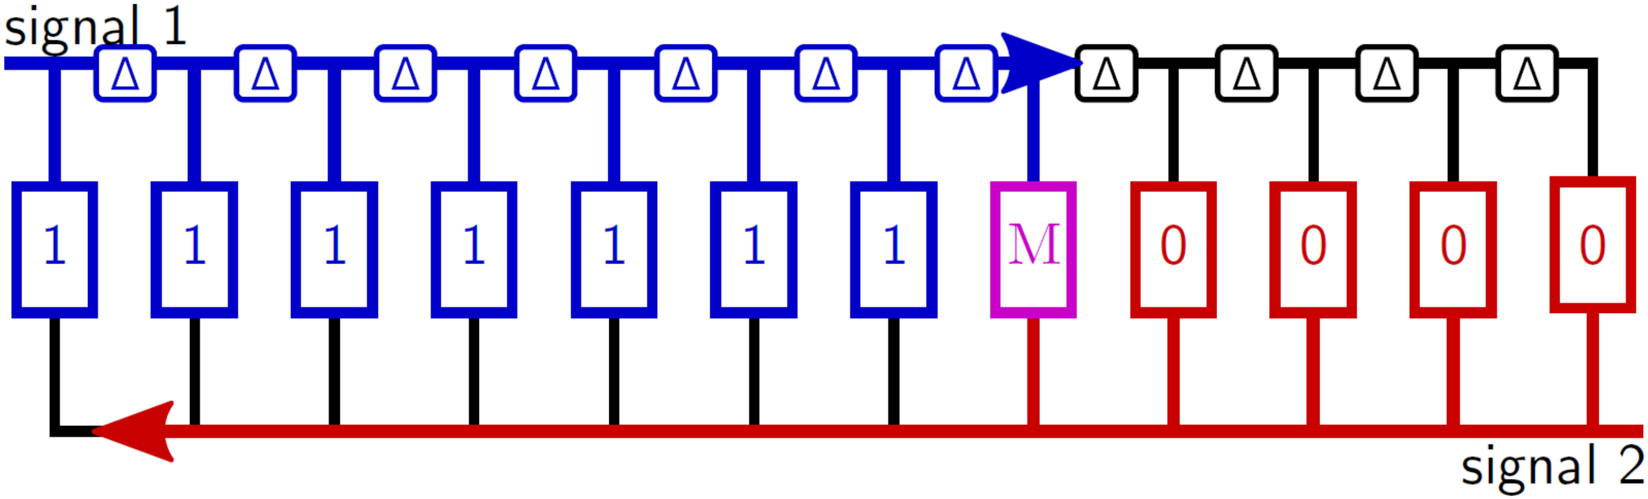

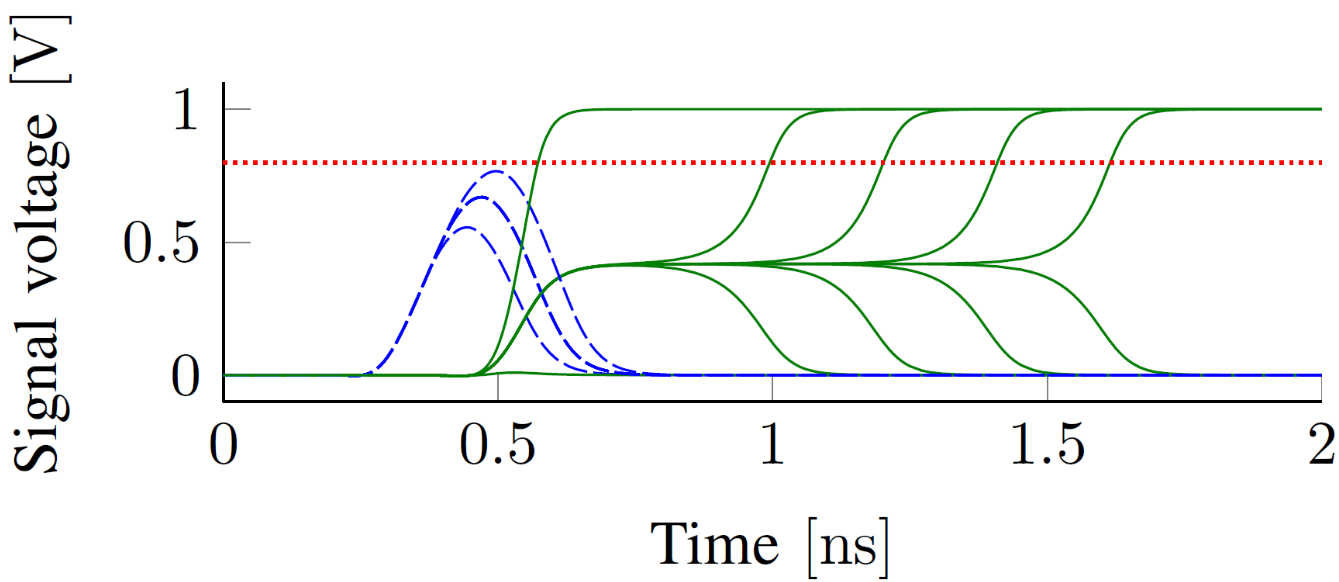

Metastability is Rare...

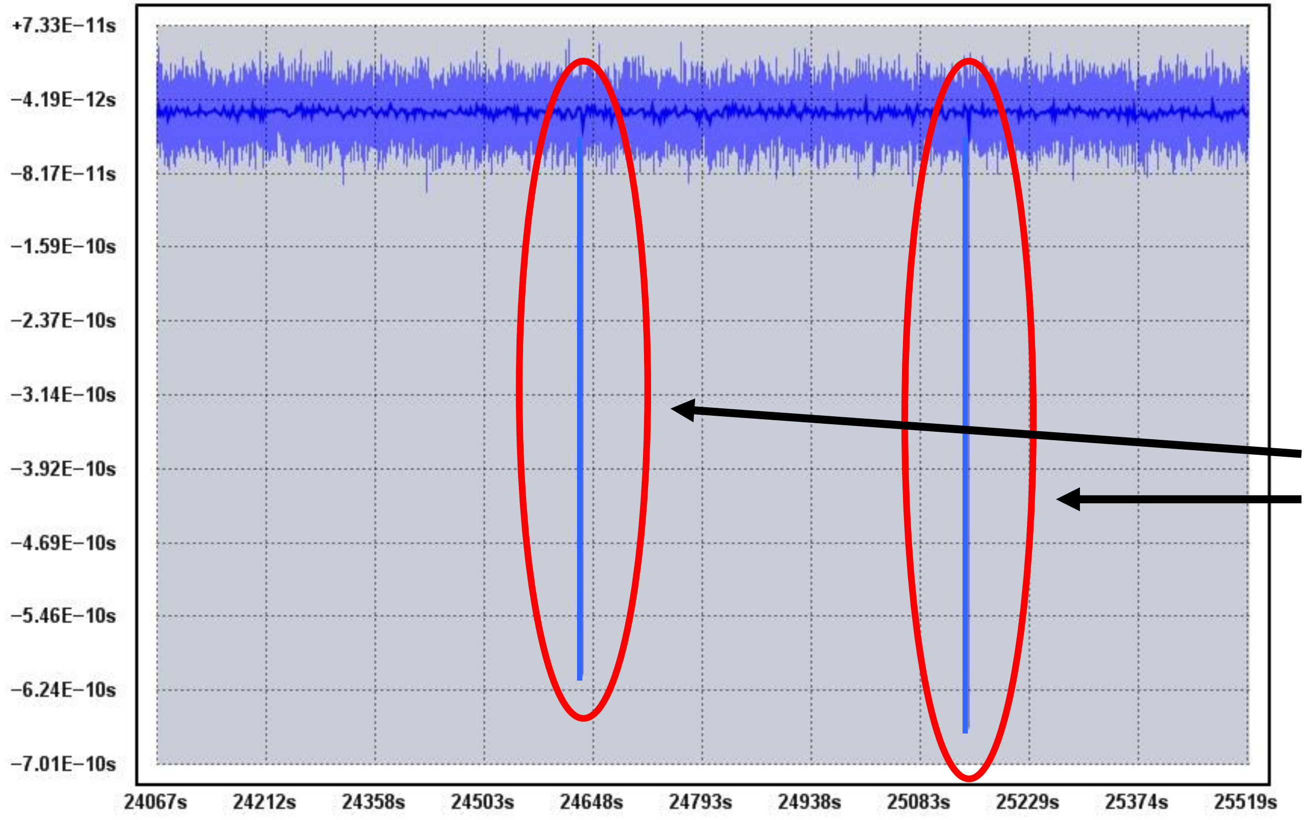

measurement equipment metastable

...**unless** your system runs at **GHz speeds!**

A “CS” Approach to Metastability

AND	0	1
0	0	0
1	0	1

→

AND$_M$	0	1	M
0	0	0	0
1	0	1	M
M	0	M	M

- What can be computed “with” metastable inputs?
- What is the complexity of such circuits?
- Can we avoid synchronizers (and their latency)?

This, and more...

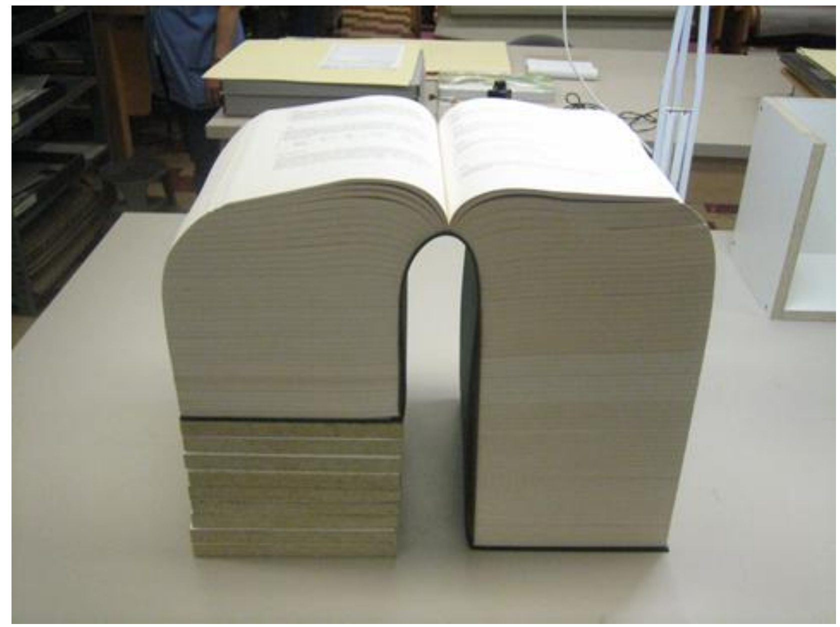

...is to become a book!

Treats

We intend to treat you to the second ≈33.33% of its contents!

Today's Menu

1. Why does this course exist?
2. **What is this course about?**
3. Who are you and what do you want?
   - discussion in small groups
   - sharing your findings with everyone
4. How will we run this course?
   - your questions and input on this
5. Heads-up: What comes next?

Outlook

winter 2020/21: clocking in the past &
    future from 40’s to 40’s

this course: fault-tolerant clocking
    Byzantine faults & self-stabilization

winter 2021/22: handling metastability
    going beyond synchronizers

Outlook

winter 2020/21: clocking in the past &
future from 40’s to 40’s

this course: fault-tolerant clocking
Byzantine faults & self-stabilization

winter 2021/22: handling metastability
going beyond synchronizers

Warning: Contents May Advance Quickly

lectures	content
2	model & getting our feet wet
3-5	limits on Byzantine fault-tolerance
6-8	optimal skew under Byzantine faults
9-11	low-degree clock distribution networks
12-13	self-stabilization and recovery
14-16	opt. skew with Byzantines & self-stabilization
17-19	consensus
20-22	pulse synchronization from consensus
23-24	synchronous counting
25-27	low-degree gradient clock distribution
28	summary & feeling good about ourselves

Today's Menu

1. Why does this course exist?
2. What is this course about?
3. **Who are you and what do you want?**
   - introduce yourself
   - what you are attending this course for
4. How will we run this course?
   - your questions and input on this
5. Heads-up: What comes next?

Now

~15 min. in breakout room (no recording):

+ implicit soundcheck for everyone
+ introductions
+ what would you like to take away from this course
+ questions

Today's Menu

1. Why does this course exist?
2. What is this course about?
3. Who are you and what do you want?
   - discussion in small groups
   - sharing your findings with everyone
4. **How will we run this course?**
   - your questions and input on this
5. Heads-up: What comes next?

Our Expectations

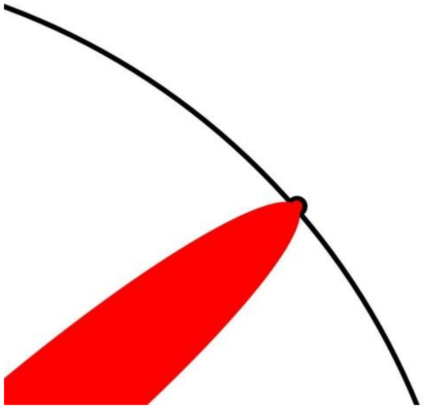

+

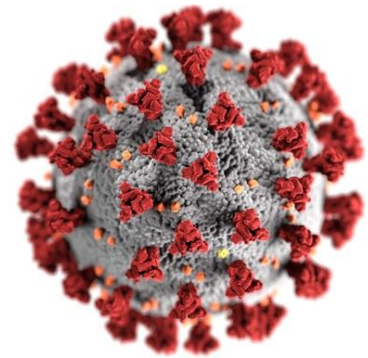

matt.might.net/articles/phd-school-in-pictures/

+

=

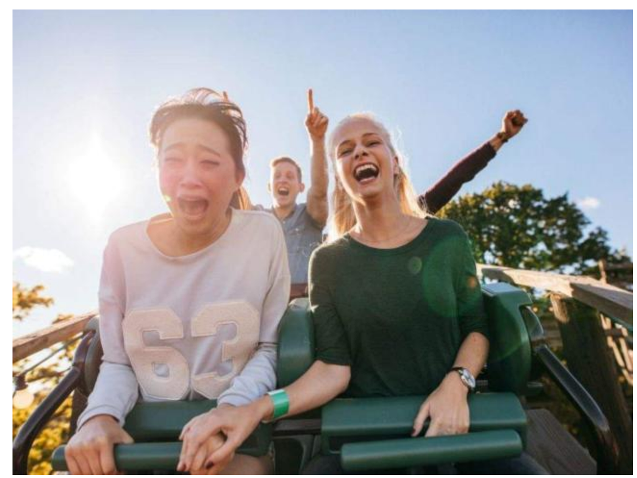

Our Expectations of You

1. For each topic (i.e., 2-3 lectures), study the reading assignment.
2. Write a **short summary of the topic**, including your thoughts and questions. **25% grade contribution**
3. Attend* the sessions:
   + brief intro/overview by the lecturer
   + discuss and/or exercise in breakout room
   + **25% grade contribution from participation**
4. After the lecture period is over, write a **report on handcrafted questions one of the topics.**

   50% grade contribution

***Recordings!** Contact us in case of privacy concerns!

Questions?

Today's Menu

1. Why does this course exist?
2. What is this course about?
3. Who are you and what do you want?
   - discussion in small groups
   - sharing your findings with everyone
4. How will we run this course?
   - your questions and input on this
5. **Heads-up: What comes next?**

Schedule for the next 7 Days

1. Read the 3-page summary of motivation and model **by tomorrow.**
2. Write an email to the mailing list. Any **questions** on the summary are **highly encouraged!**
3. I‘ll present the model and setting in depth on Monday (second opportunity for questions).
4. Study and summarize the reading assigment, handing it in **before the lecture** on Wednesday!
5. On Wednesday, **Danny takes over** for the first chapter.

See You on Monday!

Bring a Question!