The advertised title of part 3 of our ShanghaiTech Lecture is Mutual Exclusion: who's first?

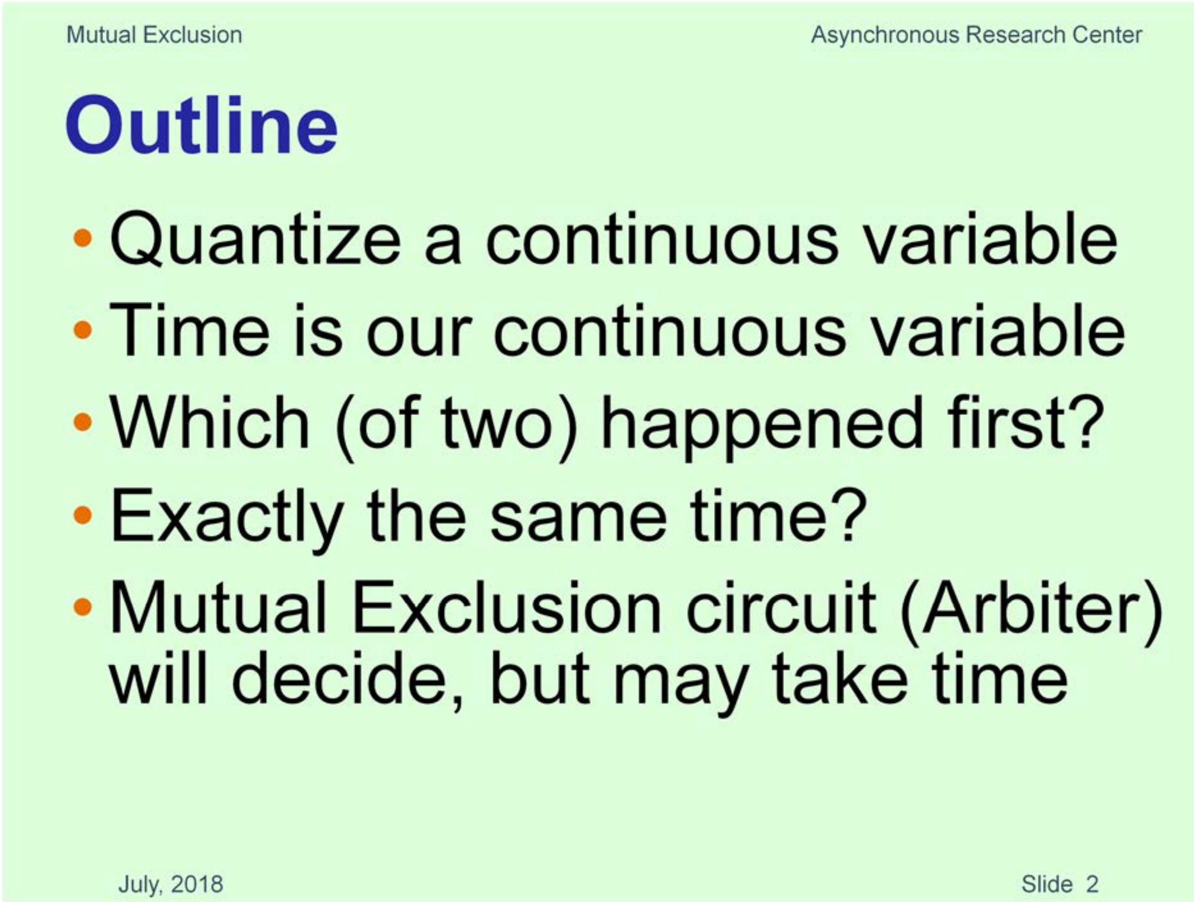

Quantizing a continuous variable
is a fundamentally hard problem.

There is always some knife-edge decision to be made.
Here we're going to deal with time as the continuous variable –
deciding which of two events came first? –
a quantized binary variable.

If the events are in different locations,
the problem is even harder, maybe impossible.

It turns out that all such decisions get harder as the
continuous variable approaches the knife edge.

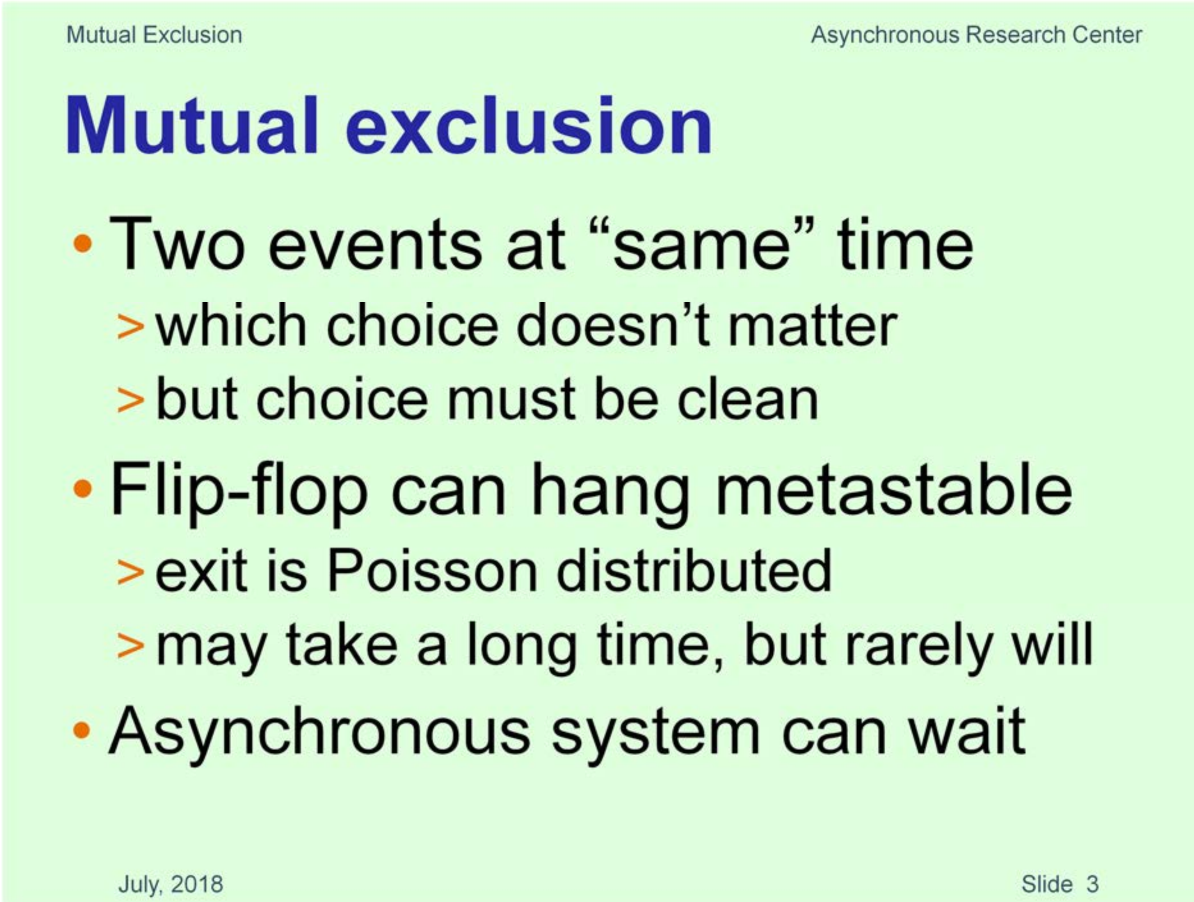

The basis of arbitration is a latch with inputs “on” and “off”.

Two cross-coupled inverting gates – each turns the other off.

Can hang “meta-stable”

Types of flip flops:

Minimum through delay – input to output.

Minimum settling time if metastable.

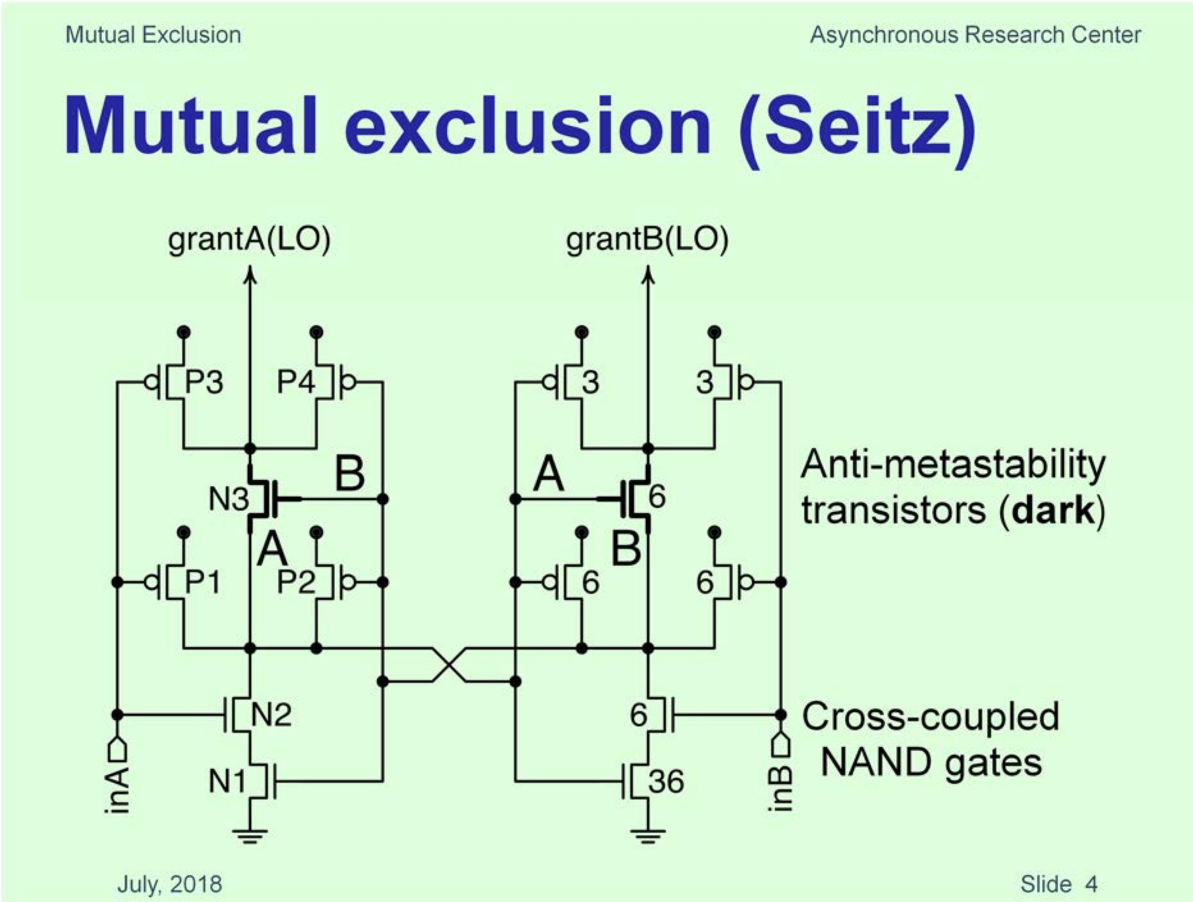

Two NAND gates at bottom: N1, N2, P1, P2 outputs A & B.

Note their cross-coupling.

Dark anti-metastability transistors N3.

Permit output only when A and B differ

by more than the transistor threshold.

WHY NOT JUST KICK IT TO THE RIGHT ?

Mutual Exclusion Asynchronous Research Center

Continental divide

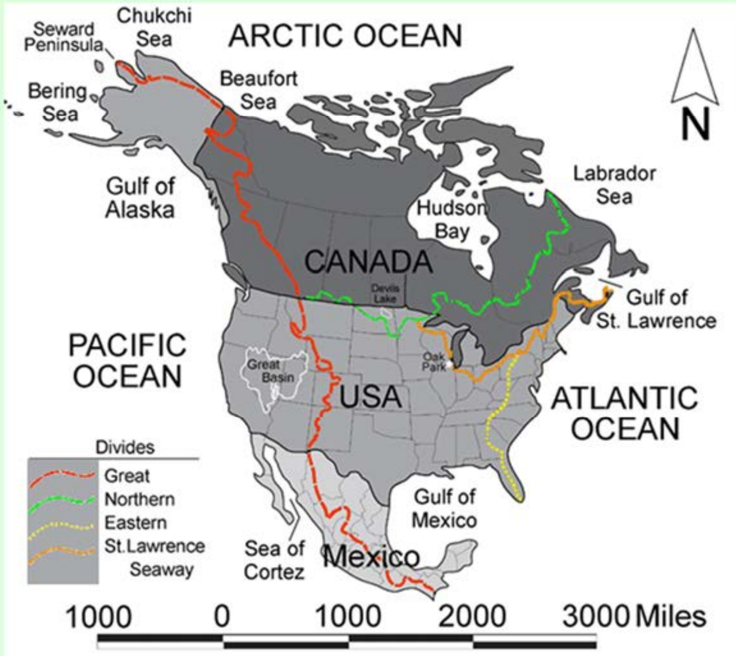

July, 2018 Slide 5

Given two stable states (Atlantic and Pacific ocean)
and a mountain range in between . . .
There must be a metastable place between them.
We call that place the “continental divide.”

You can move the continental divide, but you can’t eliminate it.
You can detect when metastability ends, but not when it starts.

Mutual Exclusion Asynchronous Research Center

Continental divide

July, 2018 Slide 6

Here’s a photograph I took exactly at the continental divide.
A “Pass” is a lower place in the ridge
where it’s easier to cross the mountains.
Most passes have names.
In the USA, most passes are marked like this.
Rabbit Ears Pass is west of Denver on US highway 40.

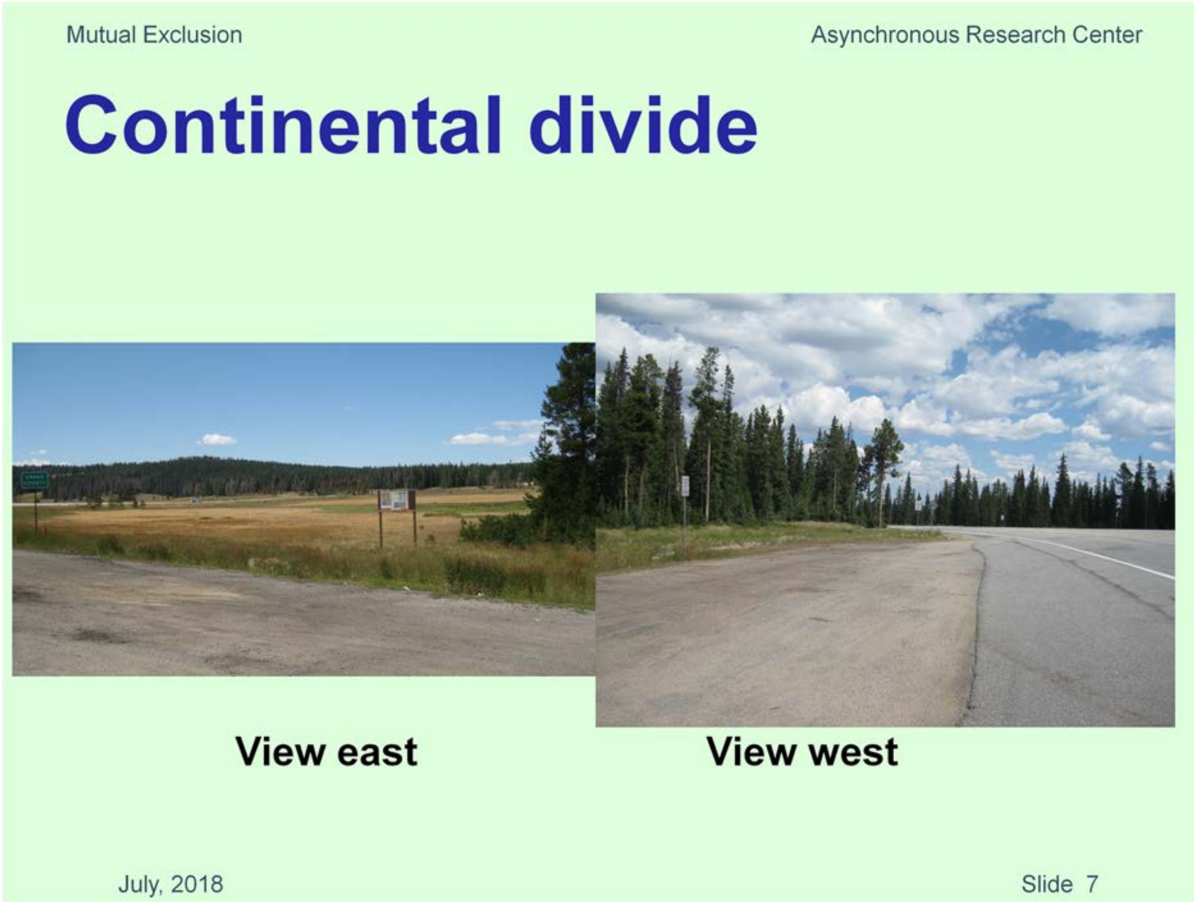

Here are views looking East and West from Rabbit Ears Pass.

You can move the continental divide,
for example by moving some dirt with a shovel,
but you can’t eliminate the divide.

You can detect when metastability ends,
but not when it starts.

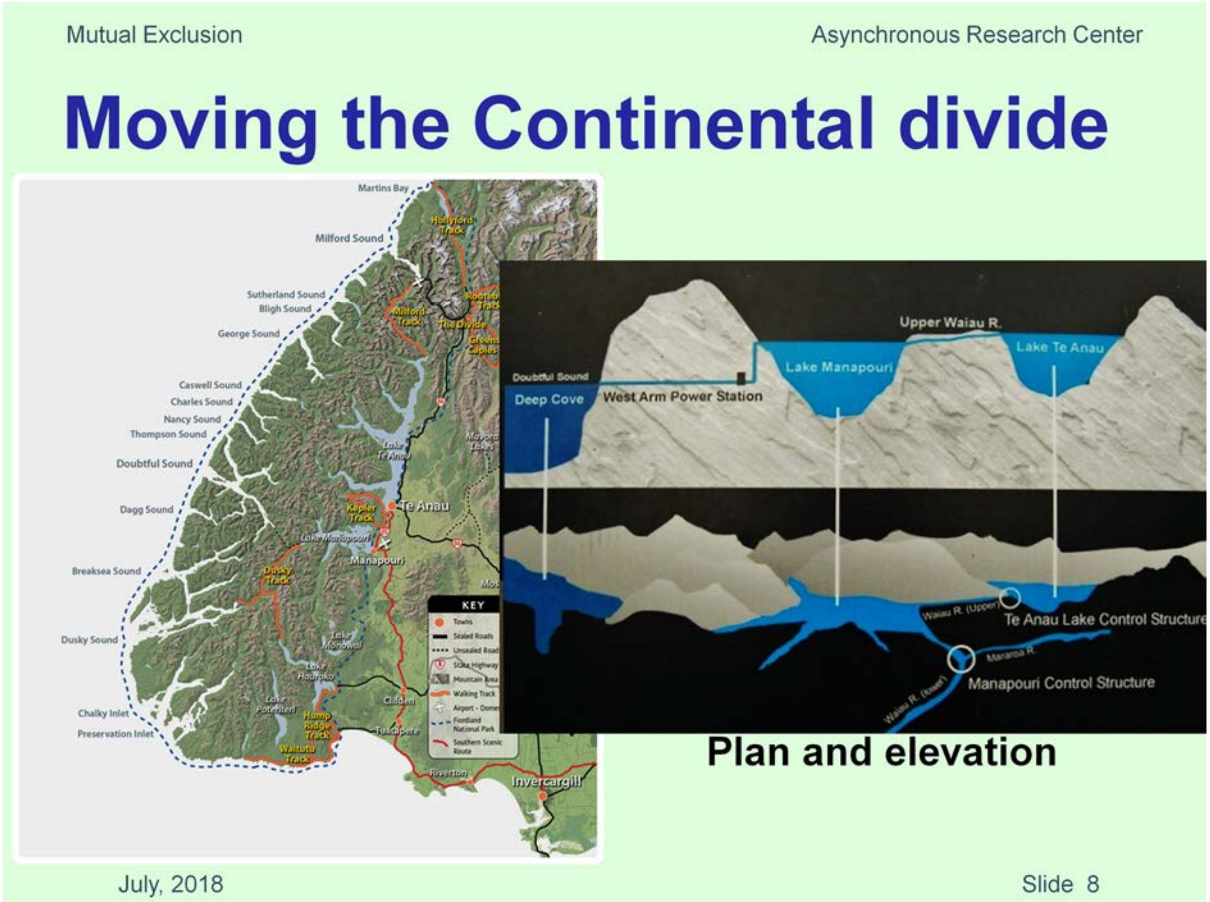

Here’s an example of moving a continental divide.

The “Southern Alps” are high mountains in New Zealand.

The highest is Mt Cook at 3724 meters.

Lake TeAnau and Lake Manapouri drained into the
South Pacific Ocean to the East of New Zealand.

They were on the **East** side of the continental divide.

A proposed power dam would have spoiled the natural beauty.

Instead, they built a tunnel from the lakes to the Tasman Sea.

Now water flows West instead of East.

The lakes moved to the **West** side of the continental divide.

A model of the underground power plant with transparent dirt shows the tunnels that provide access.

The right side shows the underground generators.

I went to visit and stood where the picture was taken.

The 50 Hz generators hum.

Tom Chaney in St Louis built this device to show metastablity.
There's an extendable piece of coaxial cable – the silver tube.
It works like a trombone to change the length of the cable.

Each pulse from the pulse generator splits.
One side sets a flip-flop through the adjustable trombone.
The other side goes directly to reset the same flip-flop.
You can adjust the trombone to control which arrives first.

In 1973 IEEE published the classic Chaney and Molnar
Anomalous Behavior of Synchronizer and Arbiter Circuits.

Mutual Exclusion Asynchronous Research Center

Scope trace

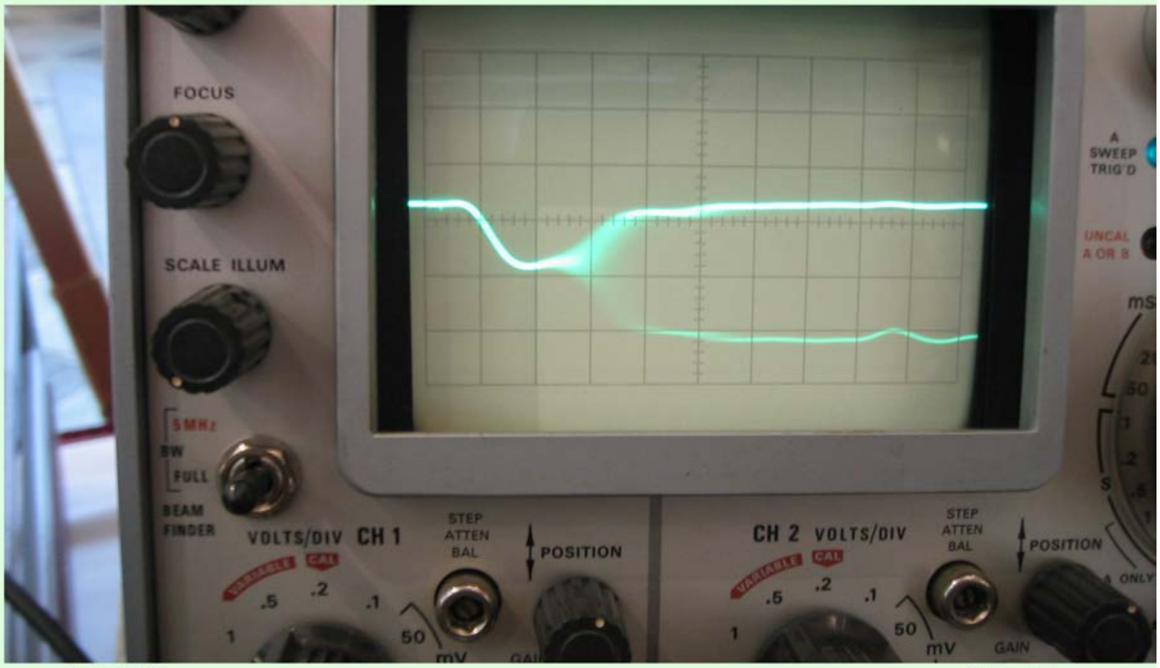

July, 2018 Slide 11

Here you see the result.

The Flip Flop starts out HI, but then falls to metastable,
the flat place in the middle.

Sometimes it goes back to the upper voltage,
less often it goes to the lower voltage.

How long it takes to leave metastability varies,
making the trace fuzzy near the end of metastability.

Changing the length of the trombone by as little as 1 mm
changes the picture a lot.

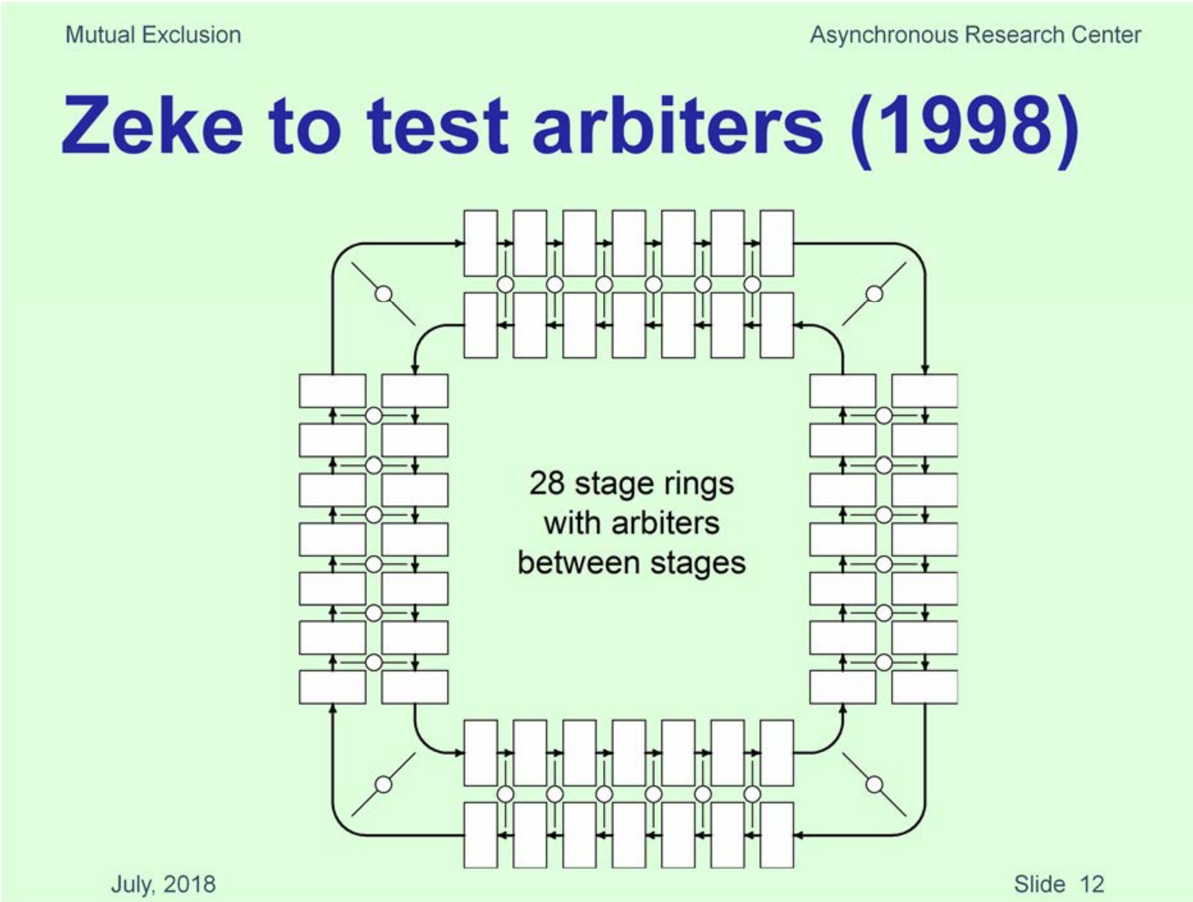

We built two chips to test the Seitz arbiter.

The first of these was Zeke with two rings of 28 stages each.

Data elements flow clockwise in the outer ring,
counterclockwise in the inner ring.

Each data element carries a “meeting count”.

Whenever two elements meet in any stage,

the stage increment both their meeting counts.

Arbiters between the stages ensure

that only one element at a time may pass from stage to stage.

Data elements must never pass each other without counting.

Zeke ran at about a billion moves per second,

so in a very short time it tested a lot of “who’s first” decisions.

None ever failed.

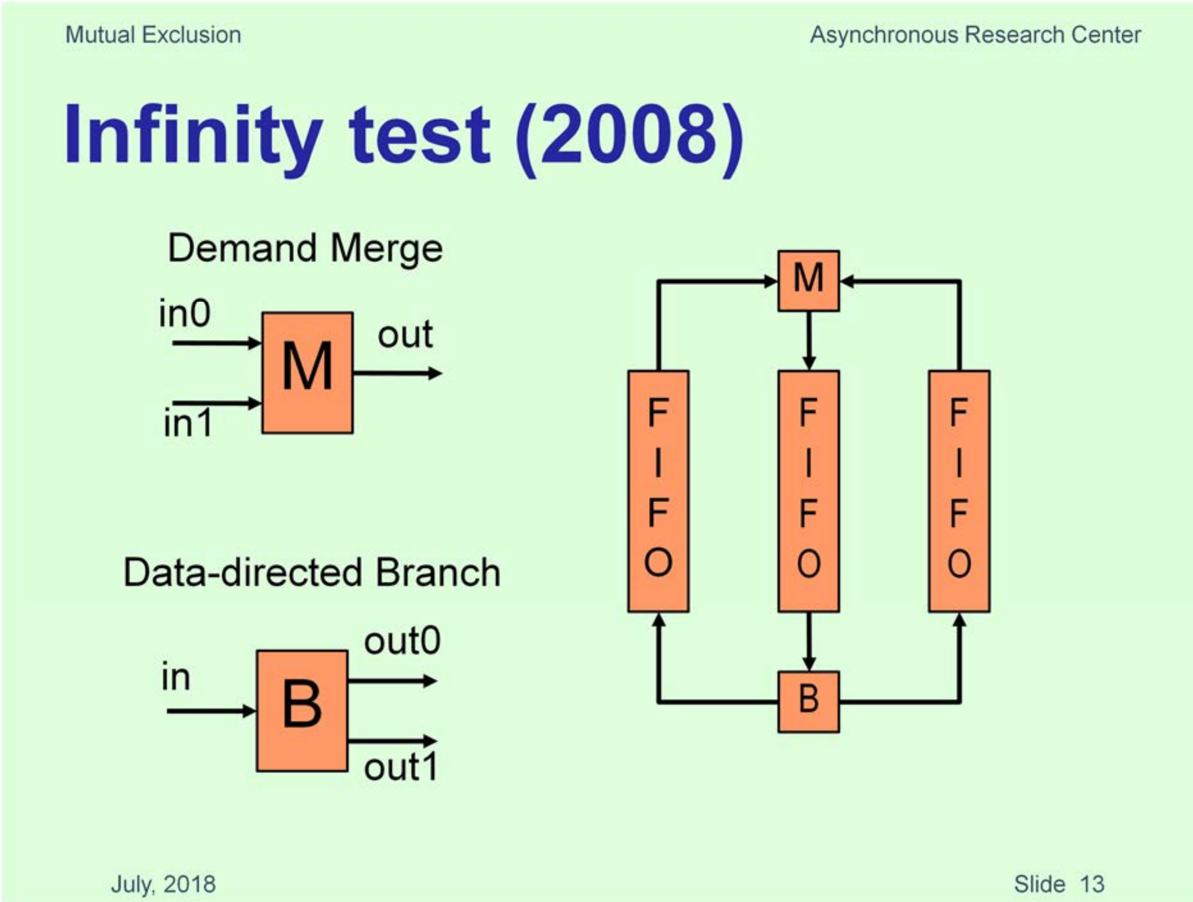

The second experiment we built was called “Infinity”.

You have to squint to see this as an infinity symbol.

Infinity has two rings of 100 stages each.

They share a section that is 50 stages long.

Each data element carries a left/right bit.

The Data-directed branch (B) tests the left/right bit of each data element to send it in that direction.

The Demand Merge (M) accepts data from either input: first-come-first-served.

Infinity works like a charm.

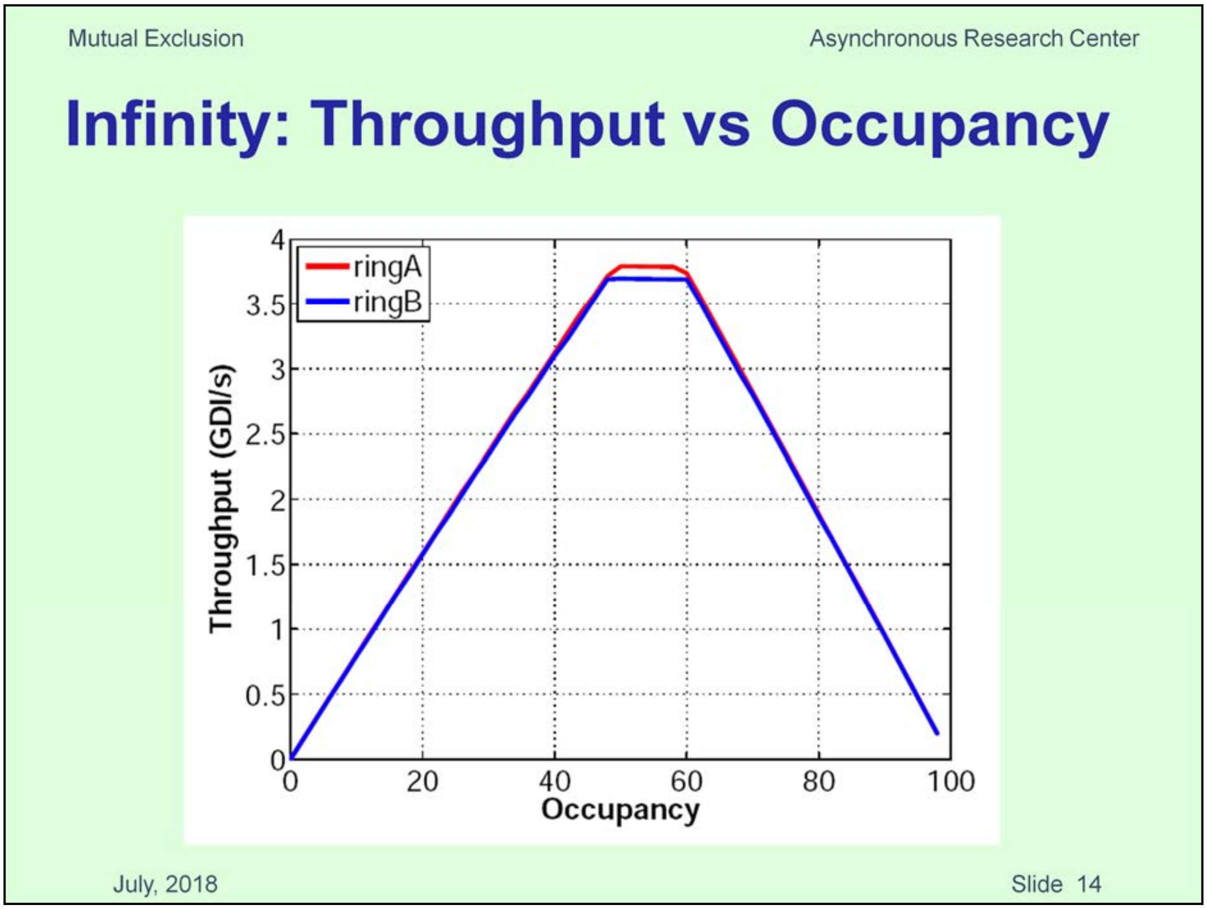

This is a “canopy graph” for the Infinity chip.

Vertical axis is throughput.

Horizontal axis is number of data items in ring.

Entirely full or entirely empty gives no throughput.

Maximum throughput somewhere in between.

Infinity uses 90 nm technology and runs at ~3.7 GDI/sec

Graphs are red for one ring blue for the other.

I don’t know why the shapes differ slightly at the top.

The flat top means there’s a slow stage. Which one?

It’s NOT the arbiter!

Long wires between some stages cause excess delay.

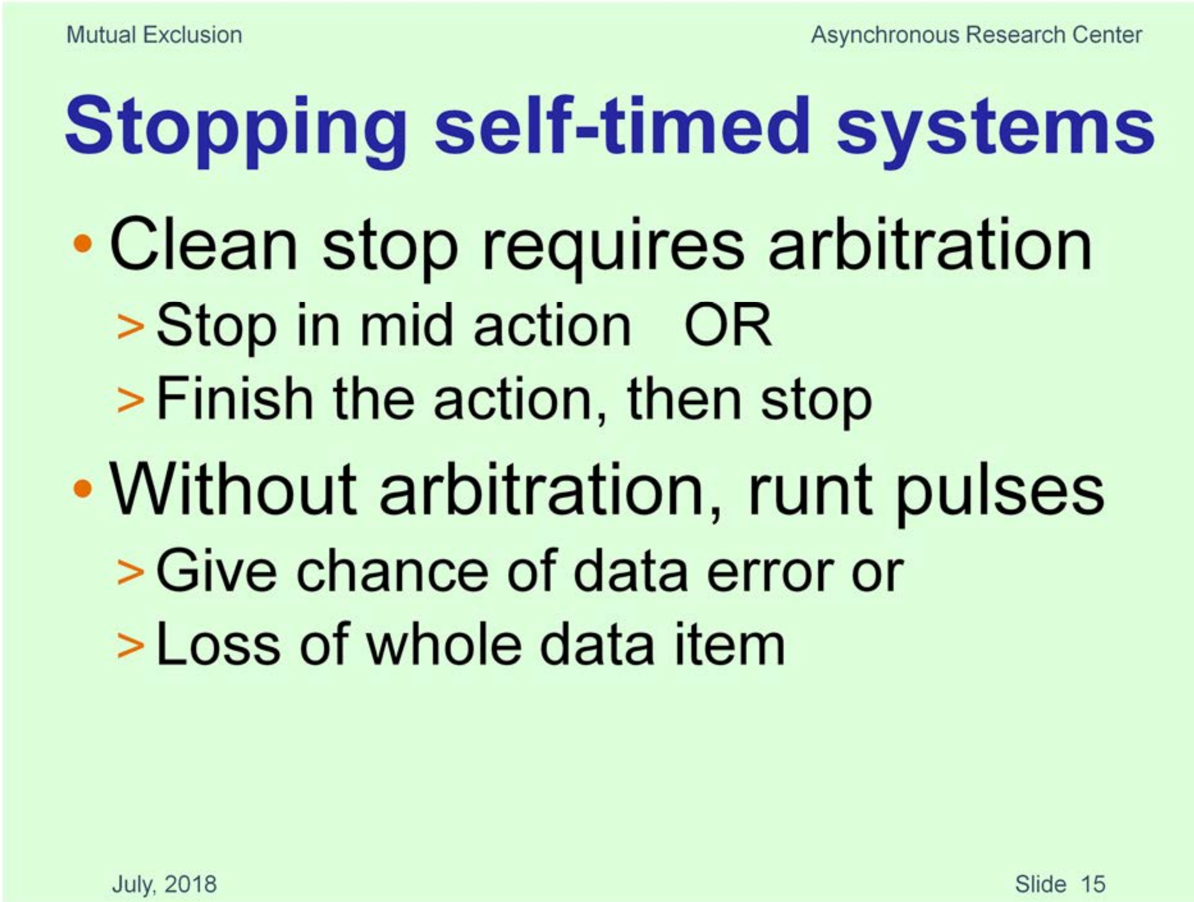

How to stop a self-timed system.

Suppose the stop signal comes just as an action starts?

Each action must decide between:
STOP vs FINISH PRESENT ACTION

We stopped our early experiments by blocking some stage.

Those experiments made a mistake in about 2% of stops.

They would corrupt data, or lose an entire data item.

The problem is that a sudden stop can create “runt” pulses.

A child skipping rope cannot stop in mid skip!

A swinging child cannot stop in mid swing!

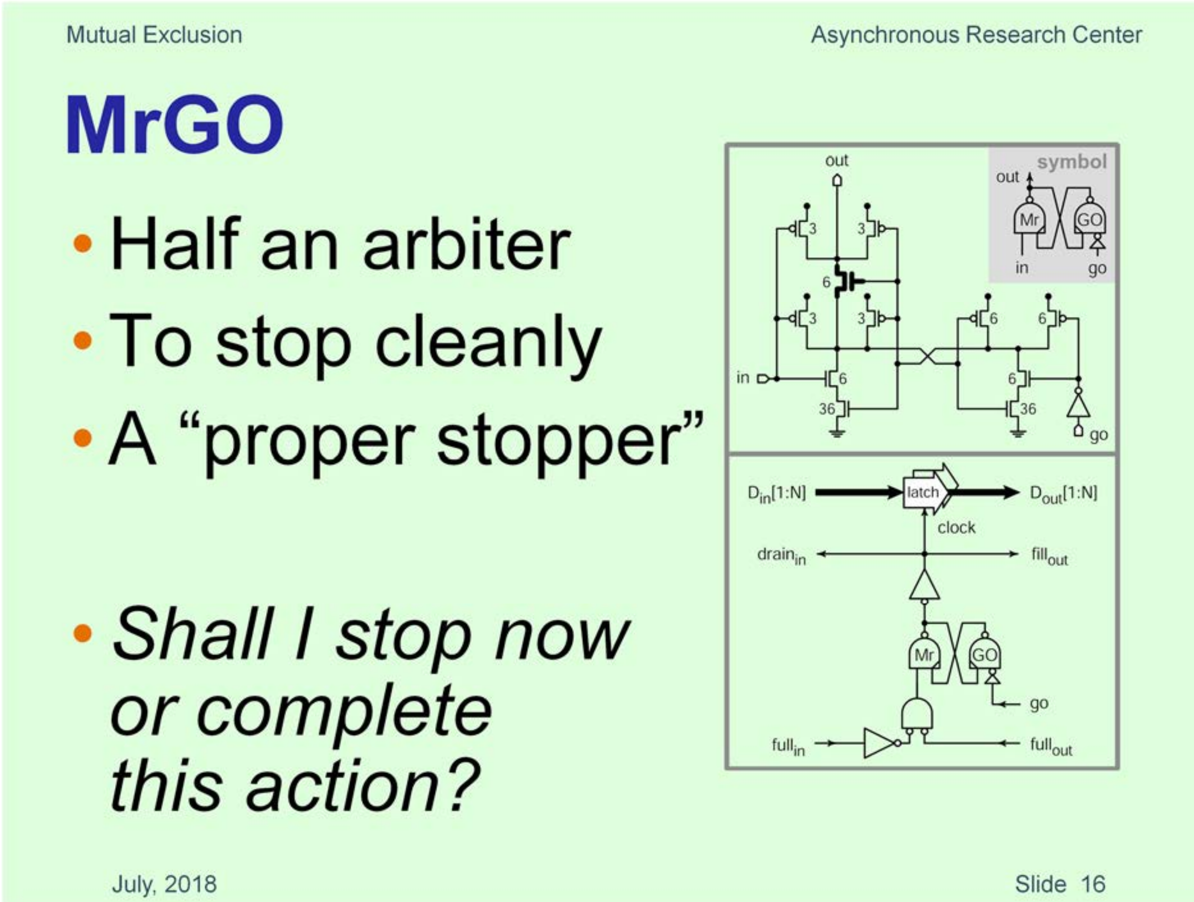

For reasons Marly will explain,
we now include MrGO in every one of our Joints.

MrGO contains a part of a Seitz arbiter.

The arbiter decides cleanly whether:
to stop immediately or
to complete this cycle and then stop.

The MrGO circuits can stop the system anywhere,
at any time, without error, even while it is running.

Discussion

I have come to the end of my talk.

I’m now open for questions and discussion.