

# REPT: Reverse Debugging of Failures in Deployed Software

Weidong Cui<sup>1</sup>, **Xinyang Ge**<sup>1</sup>, Baris Kasikci<sup>2</sup>, Ben Niu<sup>1</sup>, Upamanyu Sharma<sup>2</sup>, Ruoyu Wang<sup>3</sup>, and Insu Yun<sup>4</sup>

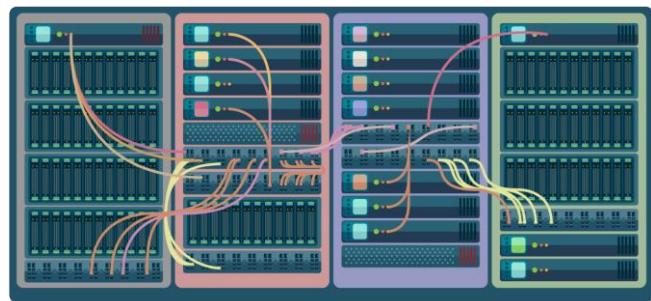
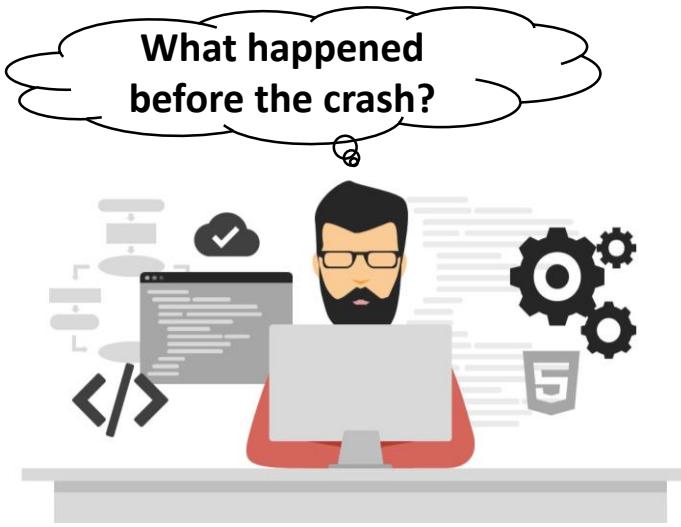
Microsoft Research<sup>1</sup>

University of Michigan<sup>2</sup>

Arizona State University<sup>3</sup>

Georgia Institute of Technology<sup>4</sup>

OSDI 2018, Carlsbad, CA







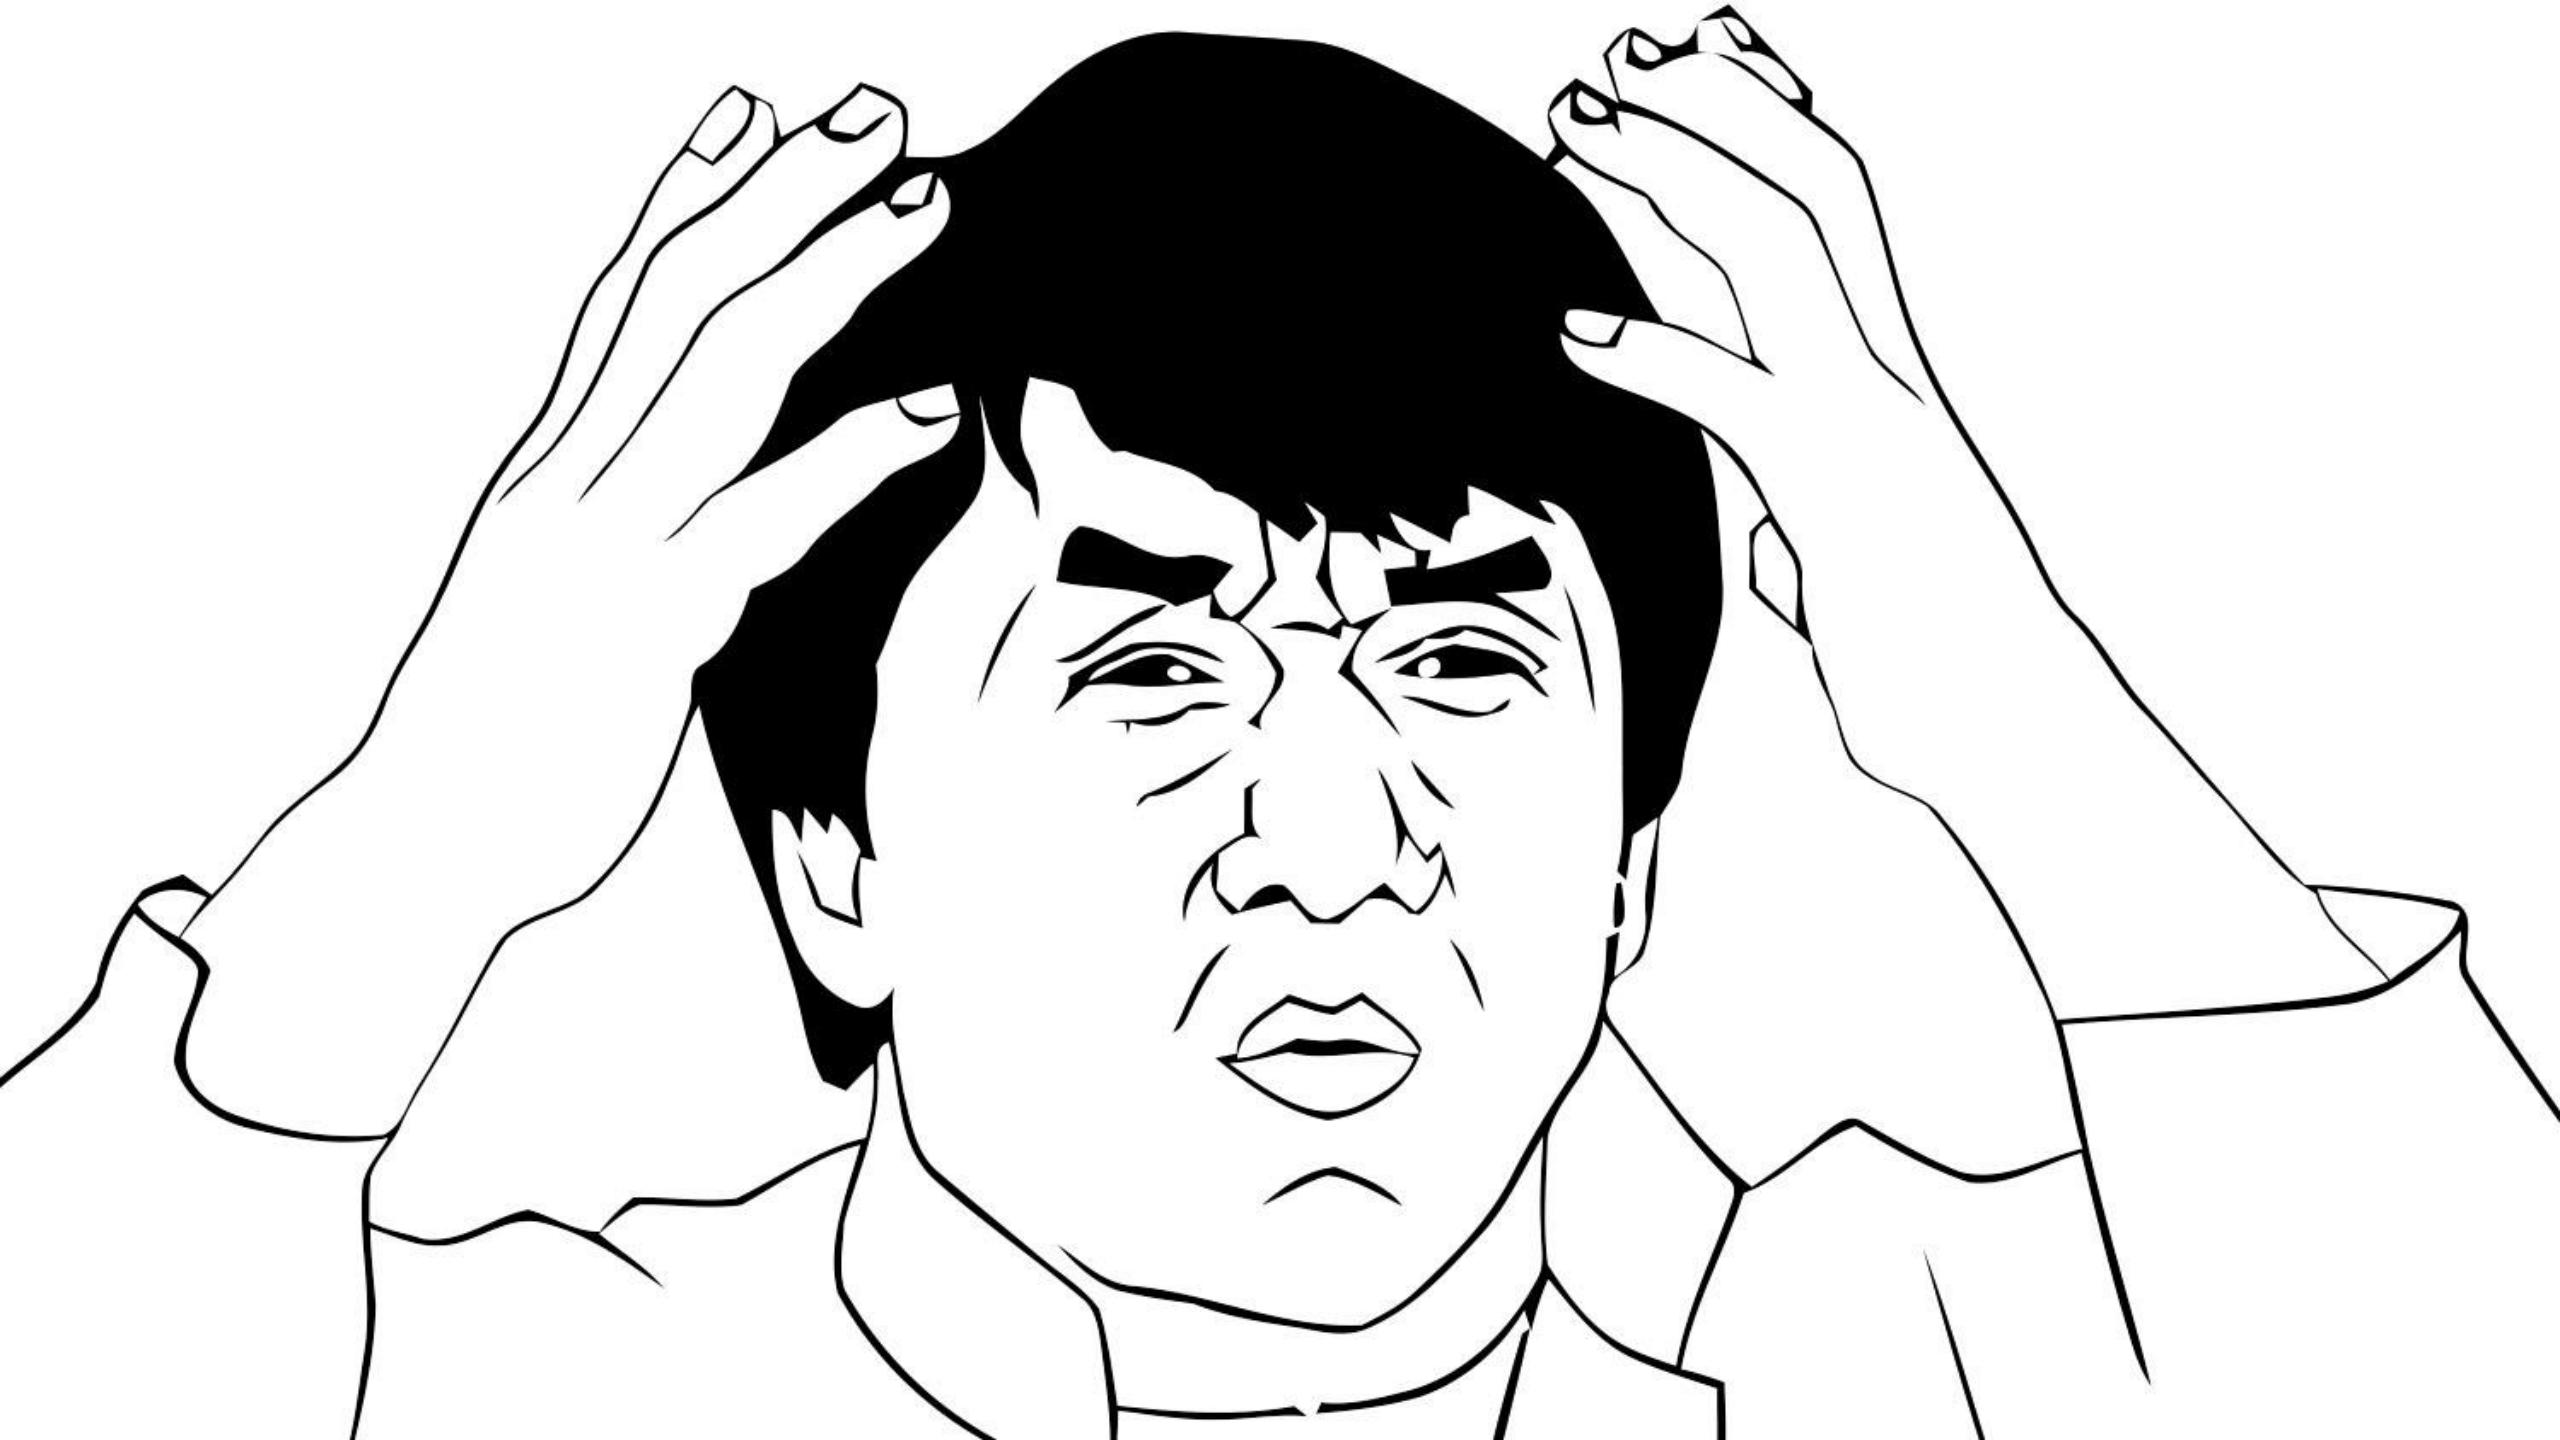


# REPT: Reverse Execution with Processor Trace

# REPT: Reverse Execution with Processor Trace

- Online hardware tracing (e.g., Intel Processor Trace)
  - Log the control flow with timestamps
  - Low runtime overhead (1 – 5%)
  - *No data!*
- Offline binary analysis
  - Recovers data flow from the control flow



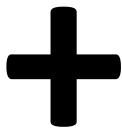


# REPT Data Recovery

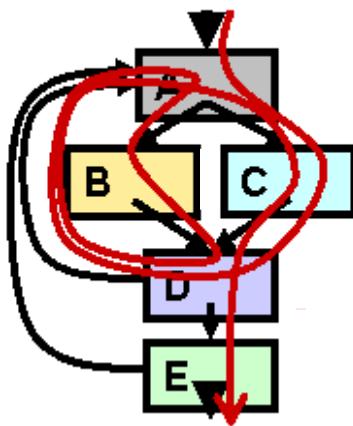
- Single-threaded execution reconstruction
- Multi-threaded execution reconstruction

## Core Dump

```
10101011010101010101011001101010010011110110110  
11010101101101010110101010101010110011010100100  
11110110110110101011010101101001010101010110110  
011010100111011011011011010101101010110101011010  
101010110011010101001111011011010101101011011010  
101101001010101010110011010010011110110110110110  
101010110101010101010101010101010101010010010111  
1011011011010101010101010101010101010101010110110  
1010010011110110110101010110101010101010101010101  
010101010110011010100100111101101101010101011011  
01010110101010101010110011010100100111101101101101  
010110101010101010101001100101010010011110110110110  
10101011010101010101010101010101010101010101010110  
10110110101010101010101010101010101010101010101010  
1100110101001001111011011010101010101010101010101  
010101010101010101010010011110110110101010101010101  
01010101010101010101010101010101010101010101010101  
01011010101010101010101010101010101010101010101010  
101011010101010101010101010101010101010101010101010  
01001111011011010101010101010101010101010101010101  
110011010100100111101101101010101010101010101010101  
011010101010101010101010101010101010101010101010101
```



## Instruction Sequence



## Execution History



How to recover  
overwritten states



lea rbx, [g]

mov rax, 1

add rax, [rbx]

mov [rbx], rax

xor rbx, rbx

lea rbx, [g]

rax=?, rbx=?, [g]=3

mov rax, 1

rax=?, rbx=?, [g]=3

add rax, [rbx]

rax=?, rbx=?, [g]=3

WRONG!

mov [rbx], rax

rax=3, rbx=?, [g]=3

→ xor rbx, rbx

rax=3, rbx=?, [g]=3

rax=3, rbx=0, [g]=3

→ lea rbx, [g]

rax=?, rbx=?, [g]=3

mov rax, 1

rax=?, rbx=g, [g]=3

add rax, [rbx]

rax=1, rbx=g, [g]=3

mov [rbx], rax

rax=3<sup>4?</sup>, rbx=g, [g]=3

xor rbx, rbx

rax=3, rbx=?, [g]=3

rax=3, rbx=0, [g]=3

lea rbx, [g]

rax=? , rbx=? , [g]=?

mov rax, 1

rax=? , rbx=g, [g]=?

add rax, [rbx]

rax=1, rbx=g, [g]=?

→ mov [rbx], rax

rax=3, rbx=g, [g]=?

xor rbx, rbx

rax=3, rbx=g, [g]=3

rax=3, rbx=0, [g]=3

lea rbx, [g]

rax=? , rbx=? , [g]=2

mov rax, 1

rax=? , rbx=g, [g]=2

add rax, [rbx]

rax=1, rbx=g, [g]=2

mov [rbx], rax

rax=3, rbx=g, [g]=?

→ xor rbx, rbx

rax=3, rbx=g, [g]=3

rax=3, rbx=0, [g]=3

→ lea rbx, [g]

rax=? , rbx=? , [g]=2

mov rax, 1

rax=? , rbx=g, [g]=2

add rax, [rbx]

rax=1, rbx=g, [g]=2

mov [rbx], rax

rax=3, rbx=g, [g]=2

xor rbx, rbx

rax=3, rbx=g, [g]=3

rax=3, rbx=0, [g]=3

lea rbx, [g]

rax=?, rbx=?, [g]=2

mov rax, 1

rax=?, rbx=g, [g]=2

add rax, [rbx]

rax=1, rbx=g, [g]=2

→ mov [rbx], rax

rax=3, rbx=g, [g]=2

xor rbx, rbx

rax=3, rbx=g, [g]=3

rax=3, rbx=0, [g]=3

# Key Techniques

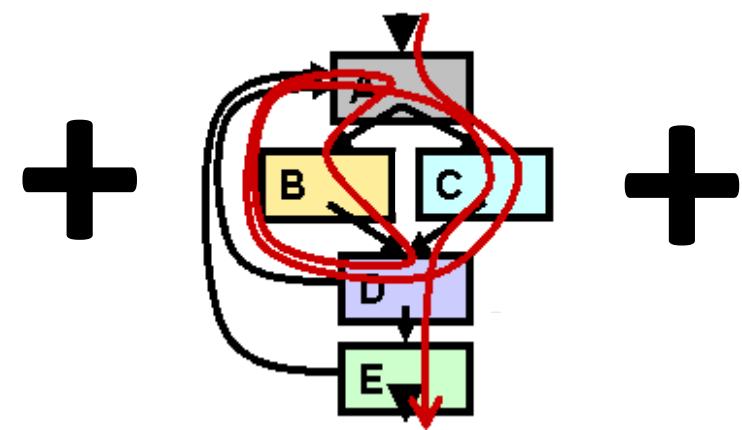
- Forward Execution
  - Recovers states before irreversible instructions
- Error Correction
  - Handles errors introduced by “missing” memory writes

# REPT Data Recovery

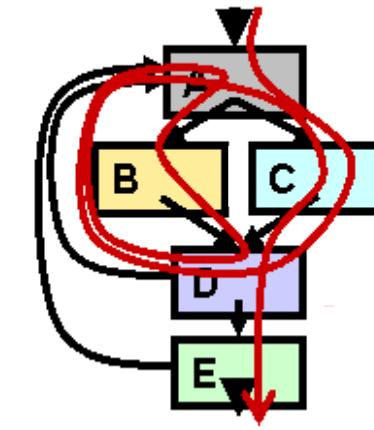
- Single-threaded execution reconstruction
- Multi-threaded execution reconstruction

# Core Dump

## **Instruction Sequence #1**



## Instruction Sequence #2

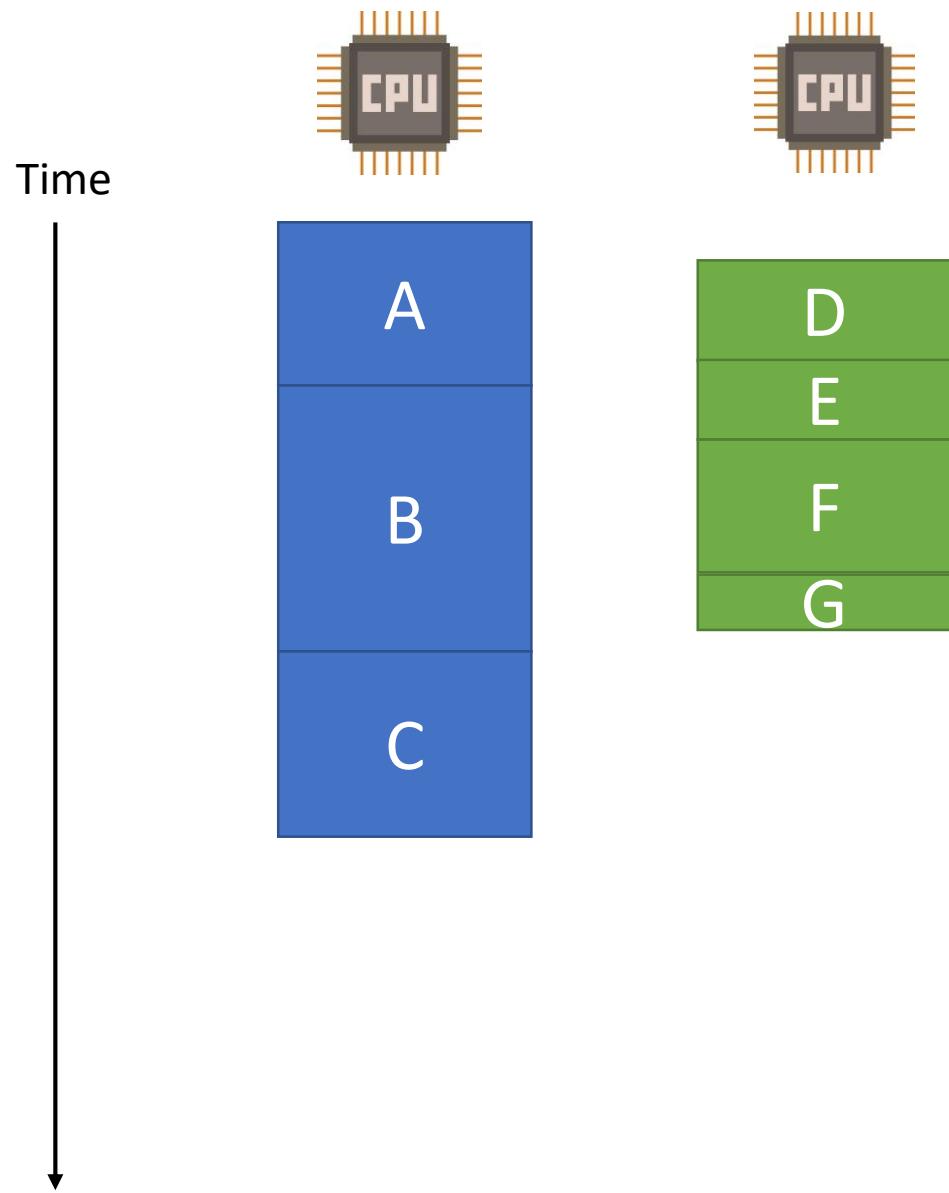


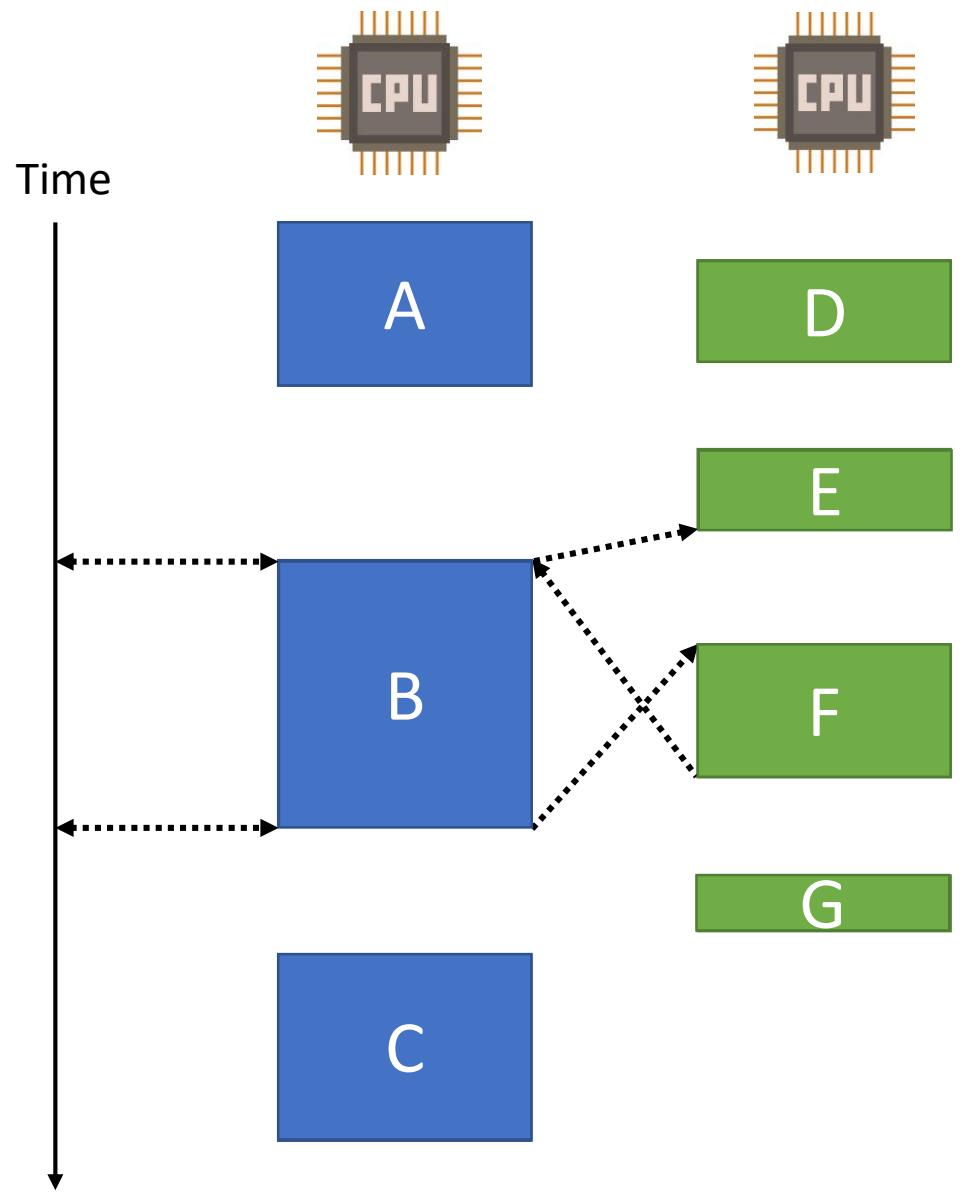
## Execution History

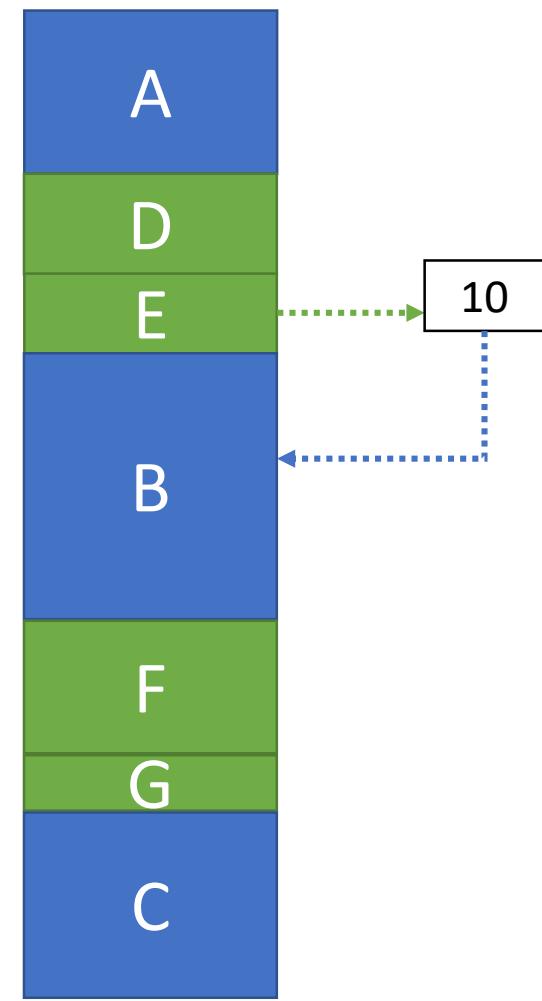
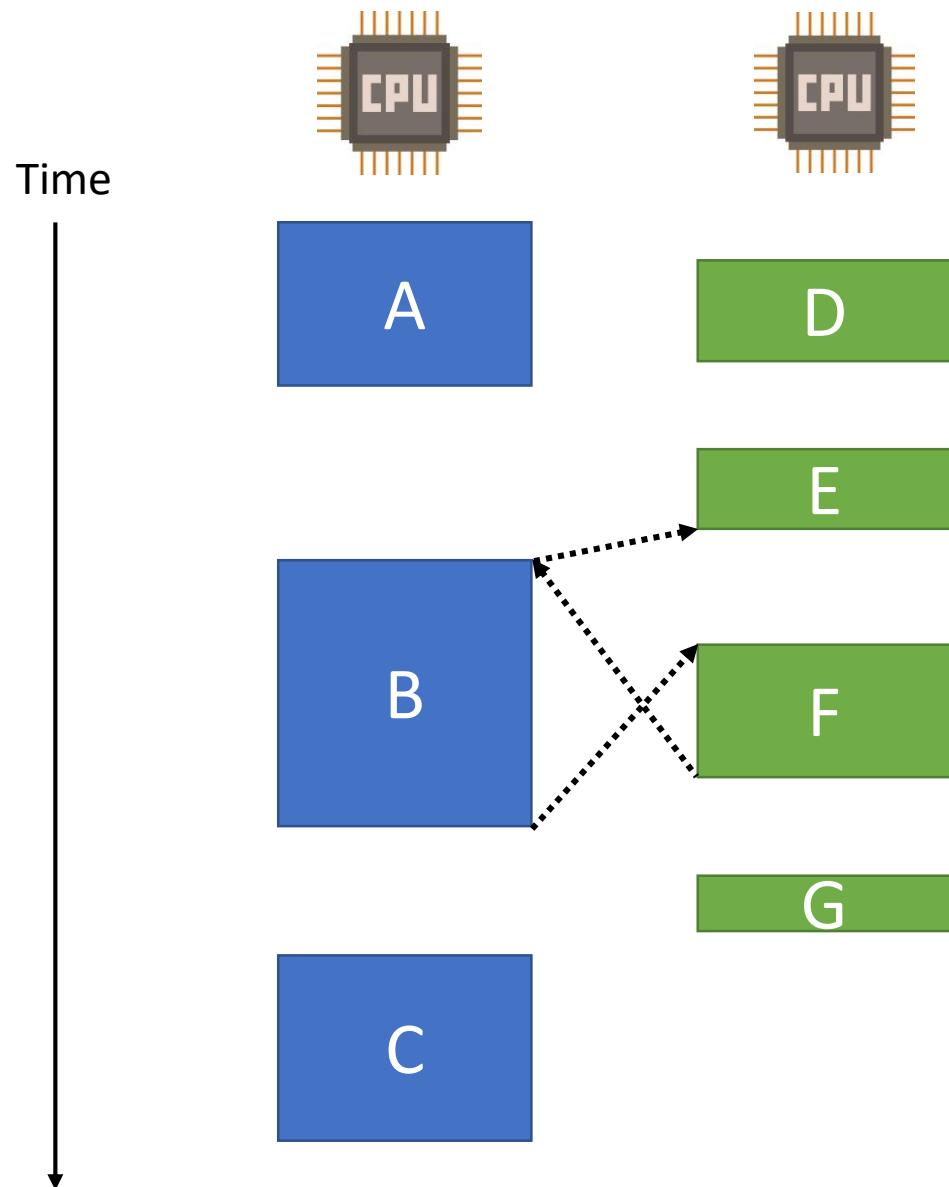
= ?

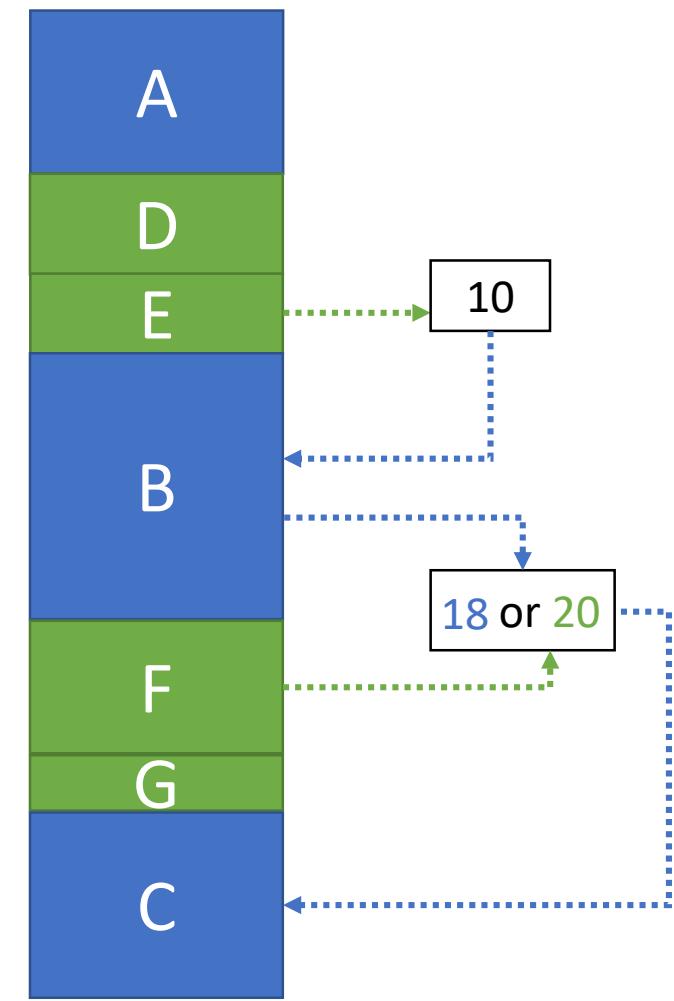
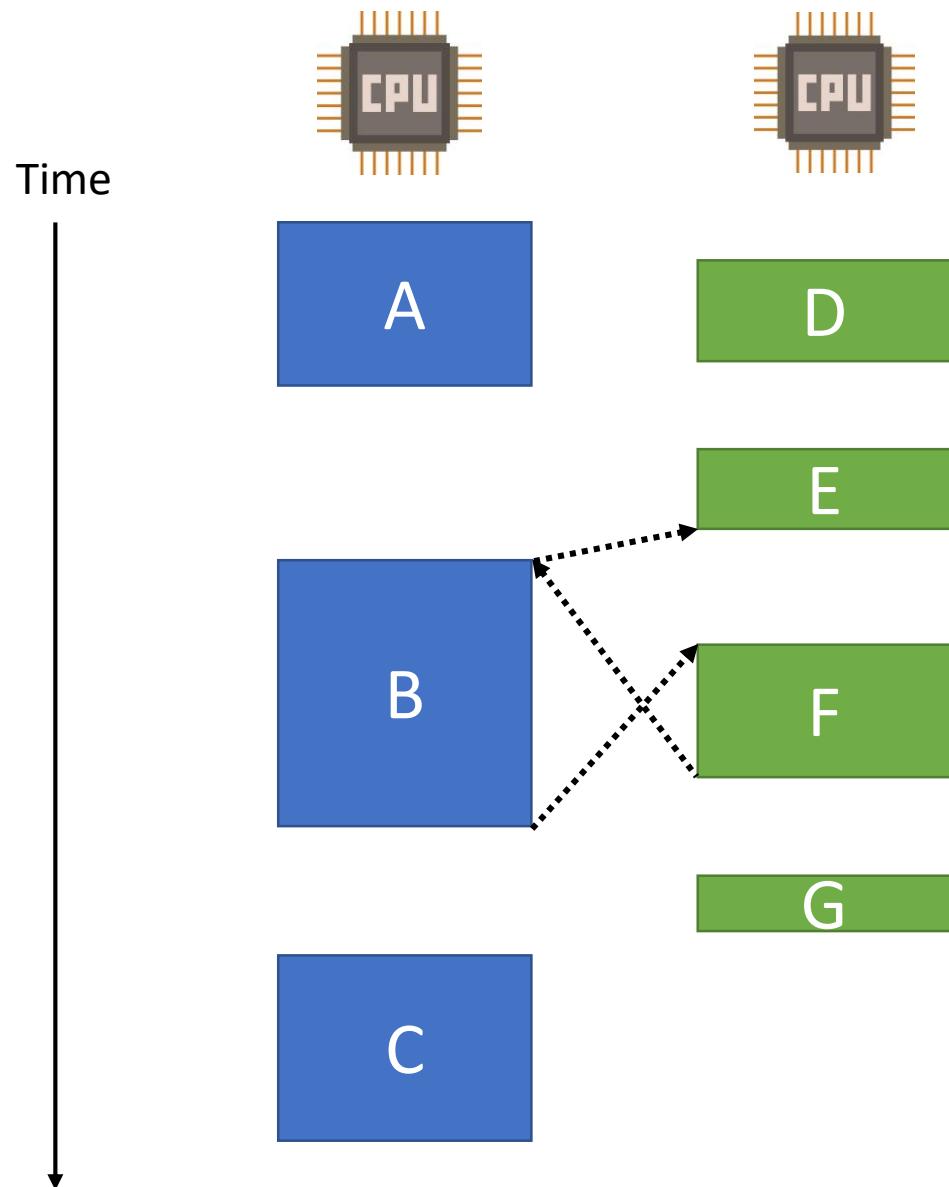


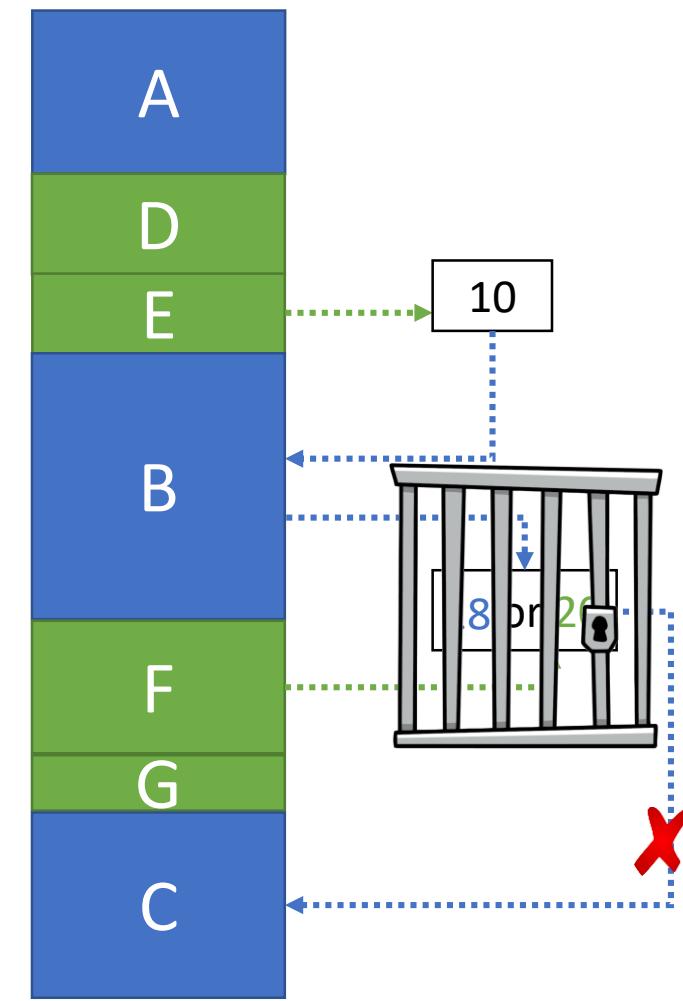
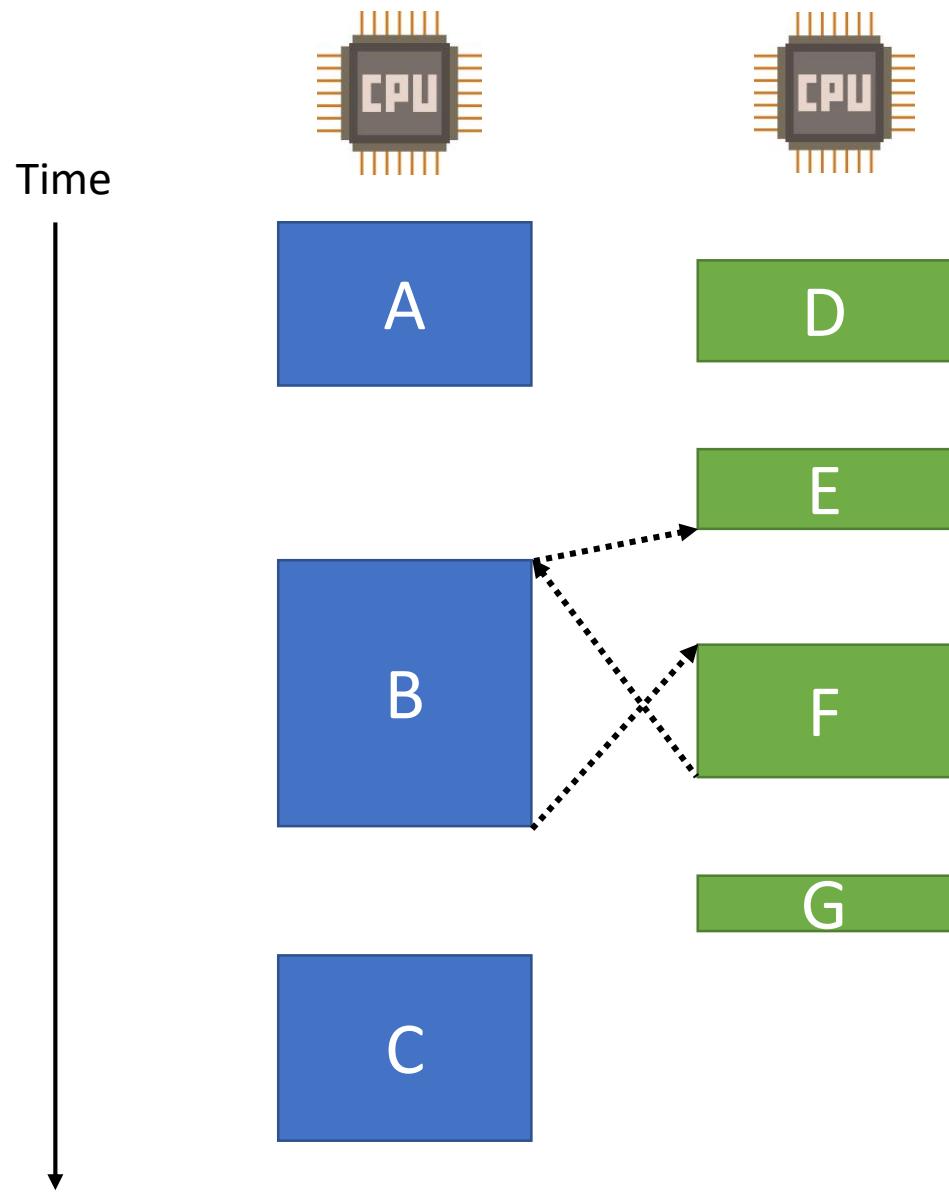
## How to determine the thread interleavings?







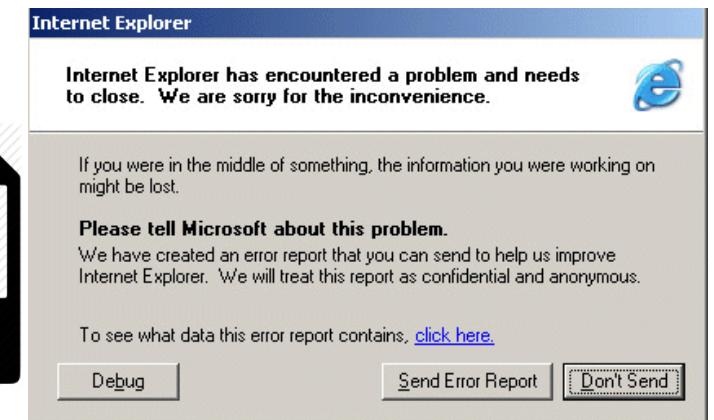
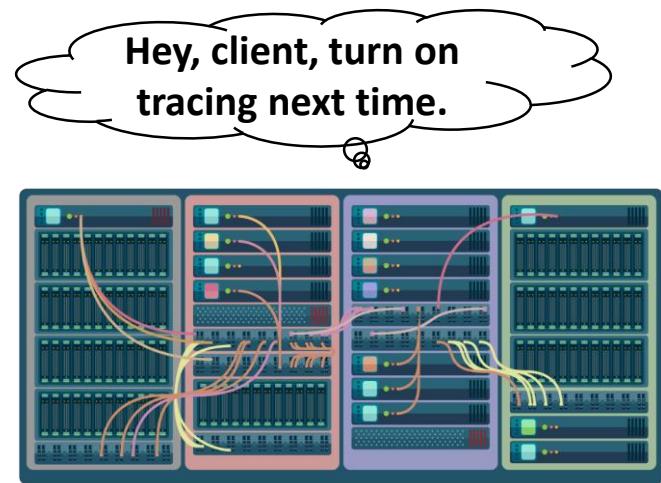




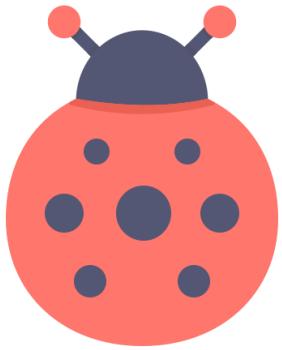
# Key Techniques

- Hardware Timestamps
  - Constructs a partial order
- Concurrent memory write detection
  - Constrains their usage to avoid propagating a wrong value

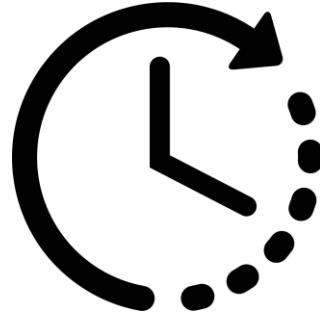
With REPT, ...



# Demo



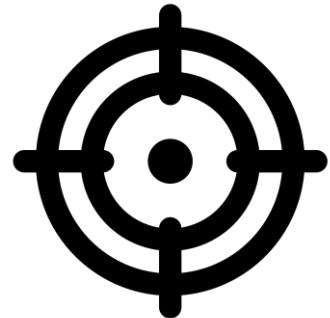
16 bugs



1-5% overhead



14 bugs



92% accuracy

# Conclusion

- Debugging production failures is important but hard
- REPT is a practical reverse debugging solution for production failures
  - Online hardware tracing to log the control flow with timestamps
  - Offline binary analysis to recover the data flow with high accuracy
- REPT has been deployed on Microsoft Windows