

# Tracing Interrupts in Embedded Software

---

Giovani Gracioli  
Federal University of Santa Catarina  
giovani@lisha.ufsc.br

Sebastian Fischmeister  
University of Waterloo  
sfischme@uwaterloo.ca



# Motivation

- **Embedded => Interrupts**  
(Especially in low power systems)

- **Embedded => Interrupts**  
(Especially in low power systems)
- **Interrupts => Complexity**  
(Non-linear program execution)

- **Embedded => Interrupts**  
(Especially in low power systems)
- **Interrupts => Complexity**  
(Non-linear program execution)
- **Complexity => Nasty bugs**  
(You only need one try, right?)

# Motivation

```
$ echo 'int main() { printf ("Hello,  
world\n"); }' | gcc -xc - -o out
```


- **Embedded => Interrupts**  
(Especially in low power systems)
- **Interrupts => Complexity**  
(Non-linear program execution)
- **Complexity => Nasty bugs**  
(You only need one try, right?)
- **Nasty bugs require good debugging support**





# Next Gen Debugging Support?



# Next Gen Debugging Support?



# Next Gen Debugging Support?





# Replay Debugging

4

Recording at run time

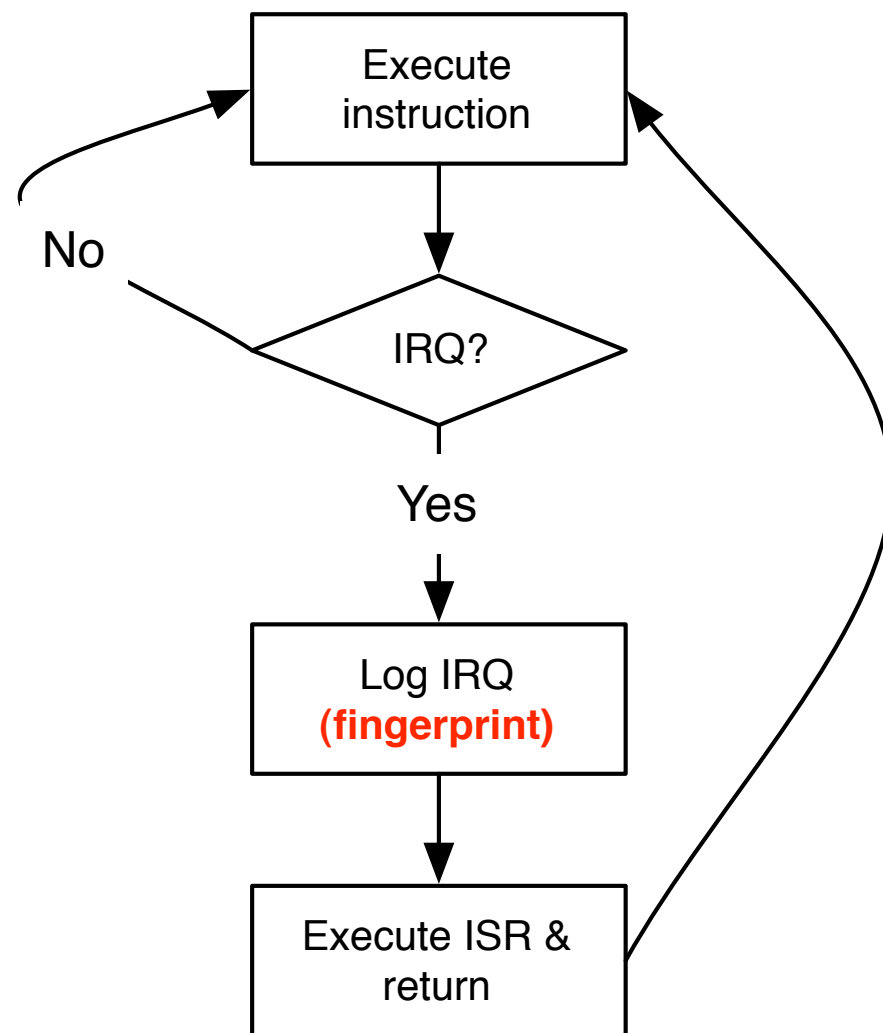
Replay offline (sim)

# Replay Debugging

4

Recording at run time

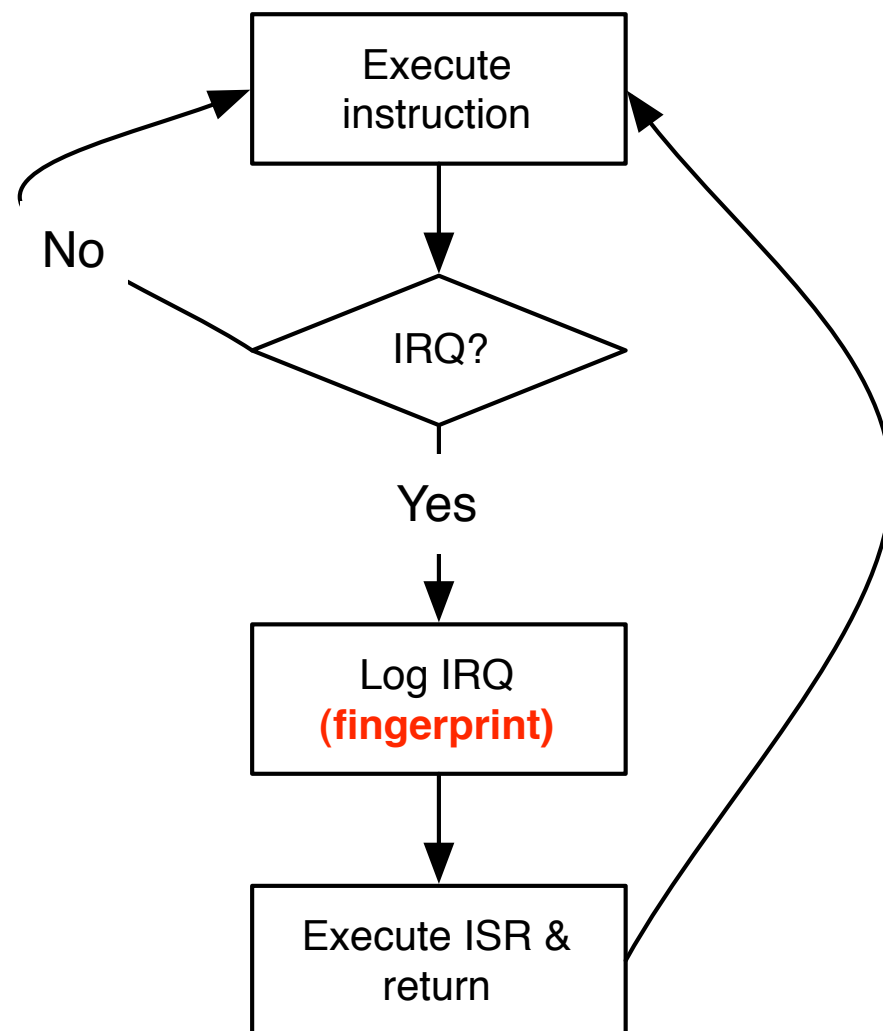
Replay offline (sim)



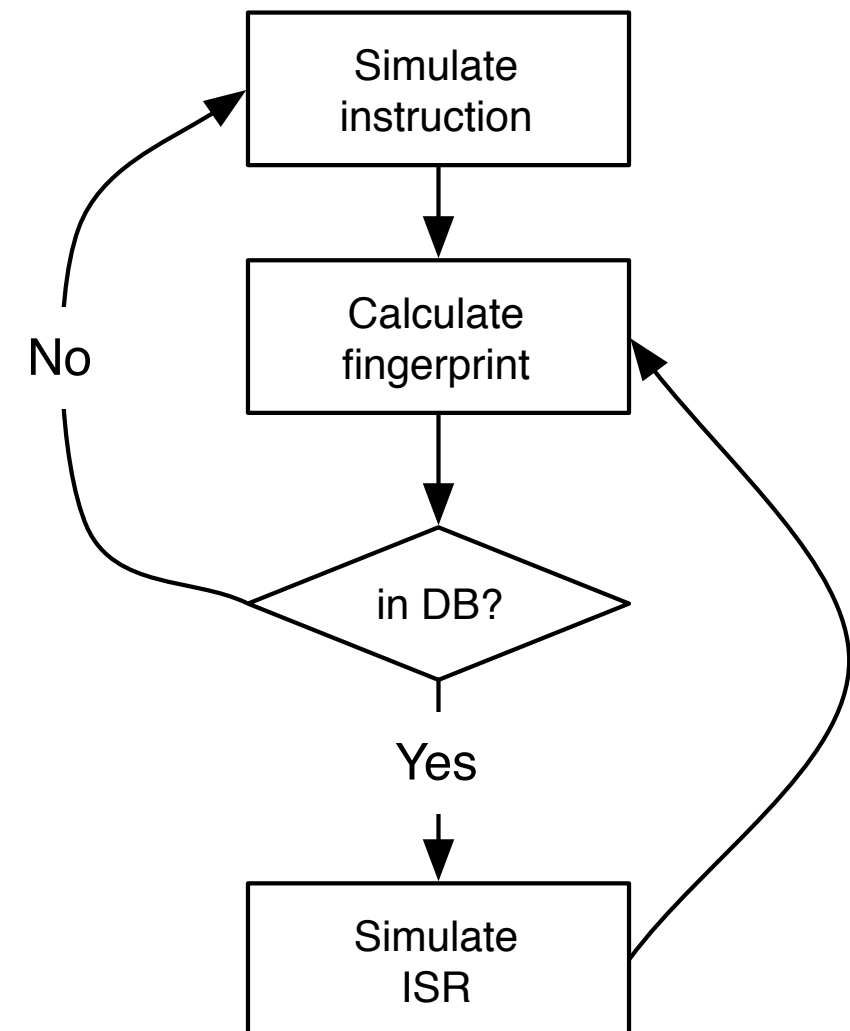
# Replay Debugging

4

## Recording at run time



## Replay offline (sim)



# Replay Debugging

Recording at run time

Replay offline (sim)

--	--


 Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

Task A




✗ Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

Task A




**X** Non-determinism happens (e.g., interrupt)



# Replay Debugging

Recording at run time

Replay offline (sim)




**✗** Non-determinism happens (e.g., interrupt)

# Replay Debugging

5

Recording at run time

Replay offline (sim)




**✗** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

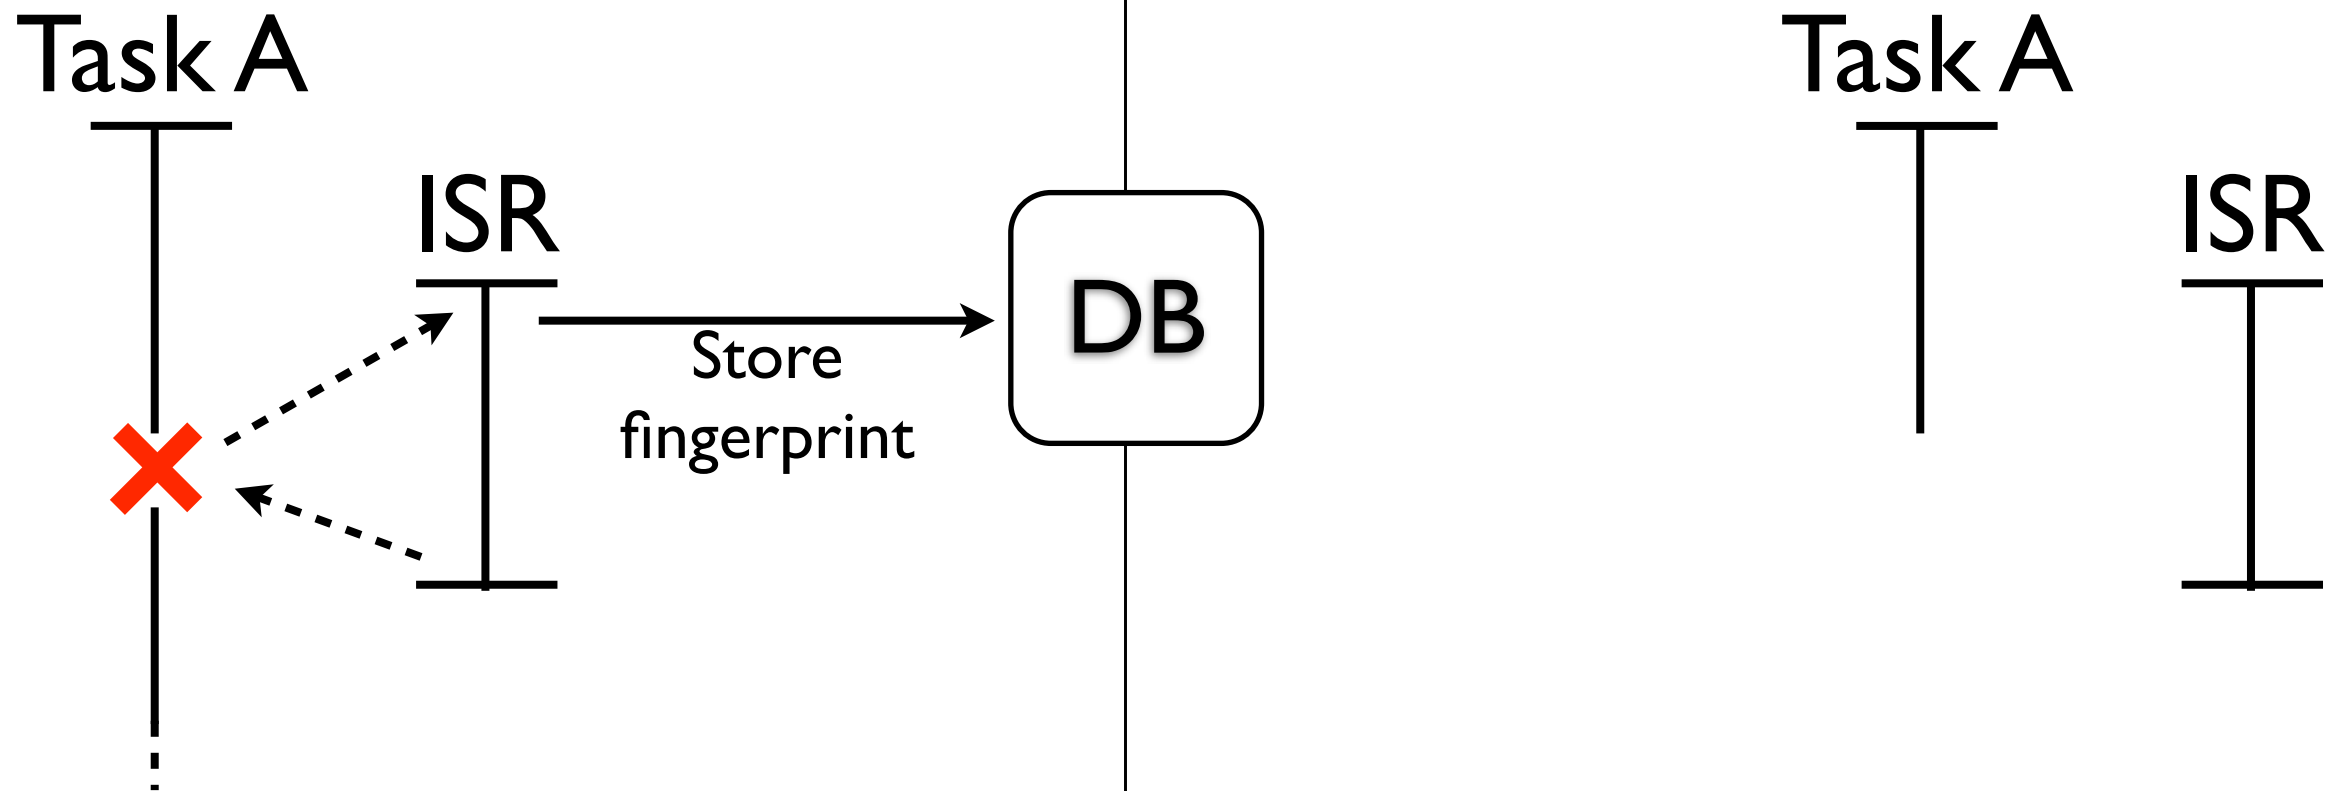


**✗** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

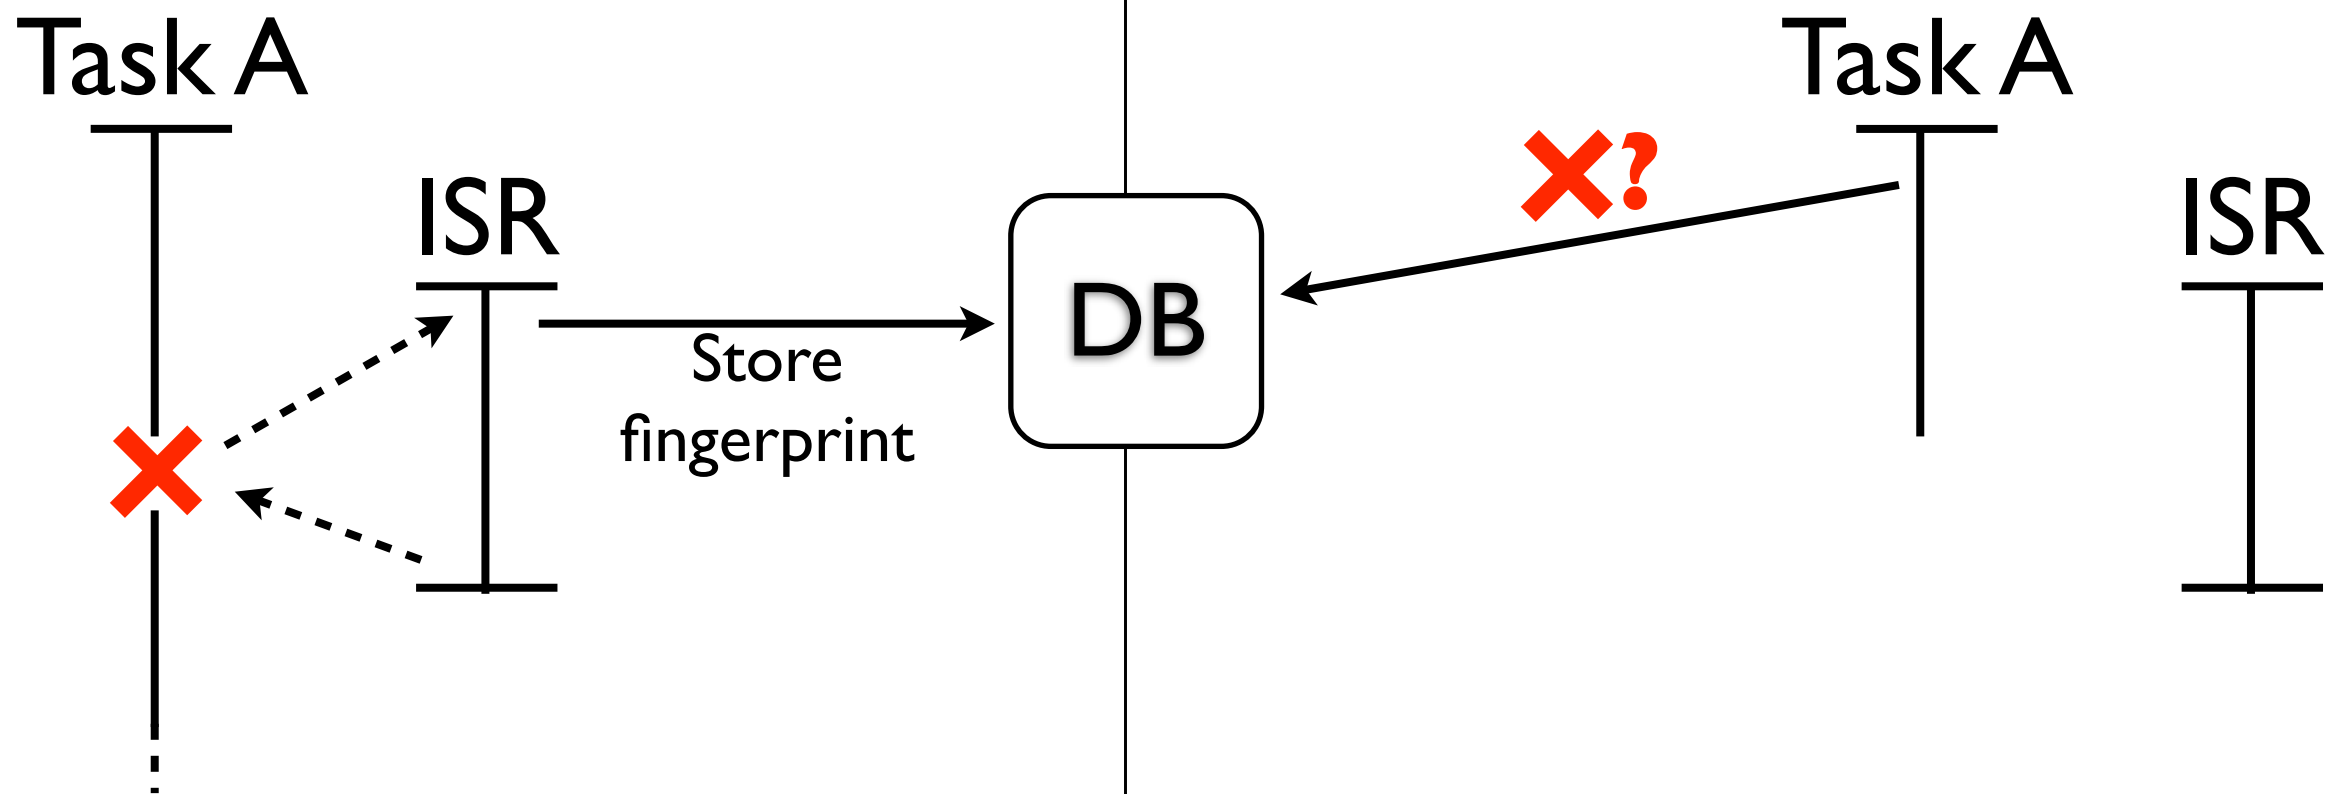


**✗** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

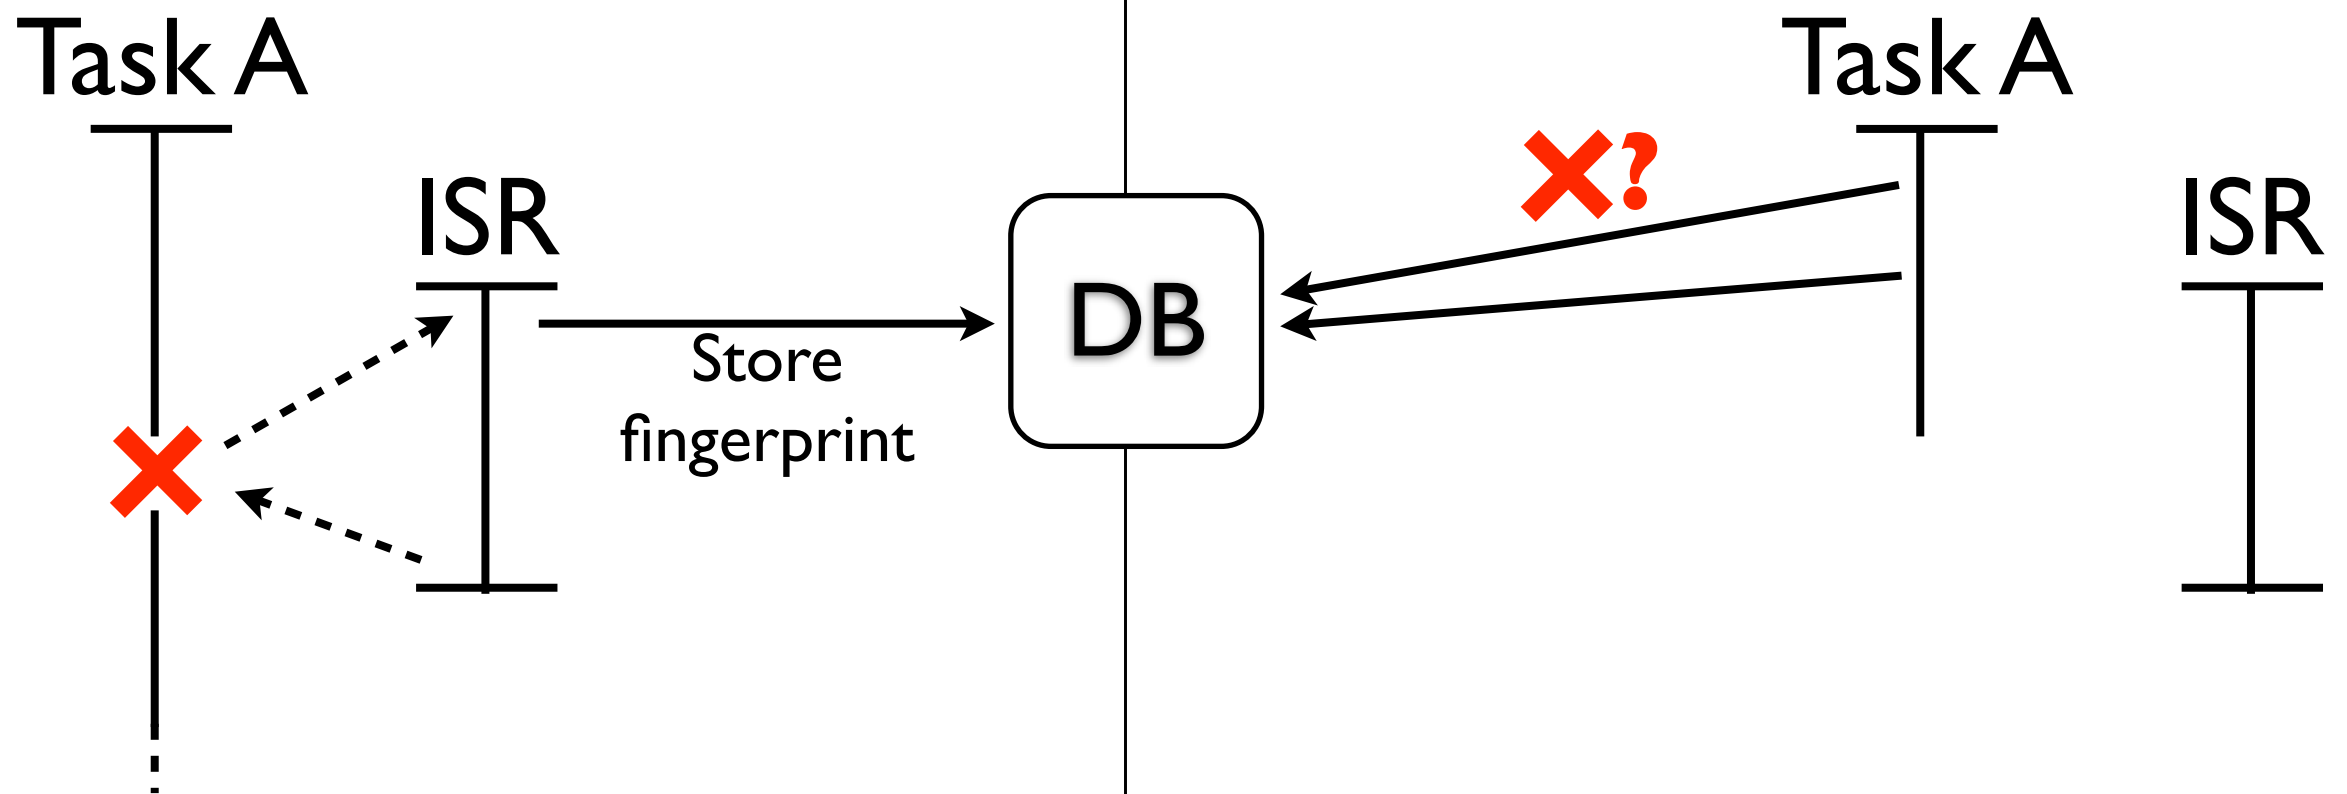


**X** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

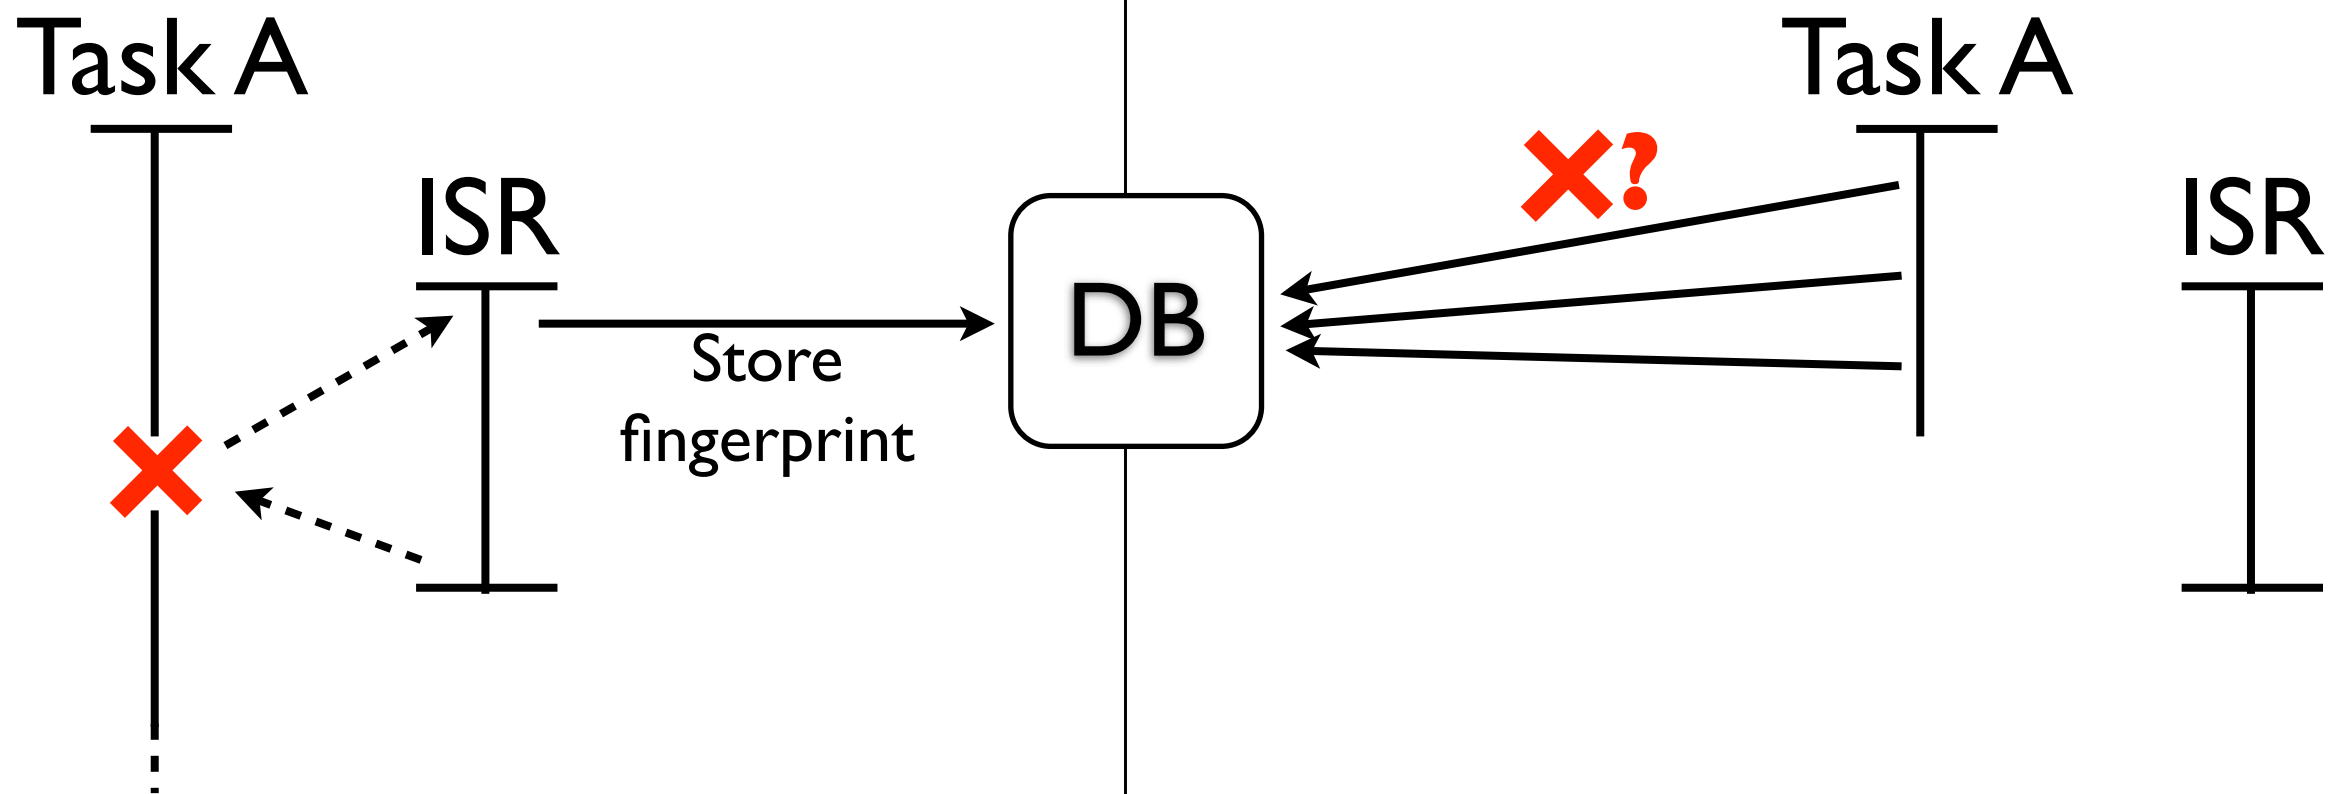


**X** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

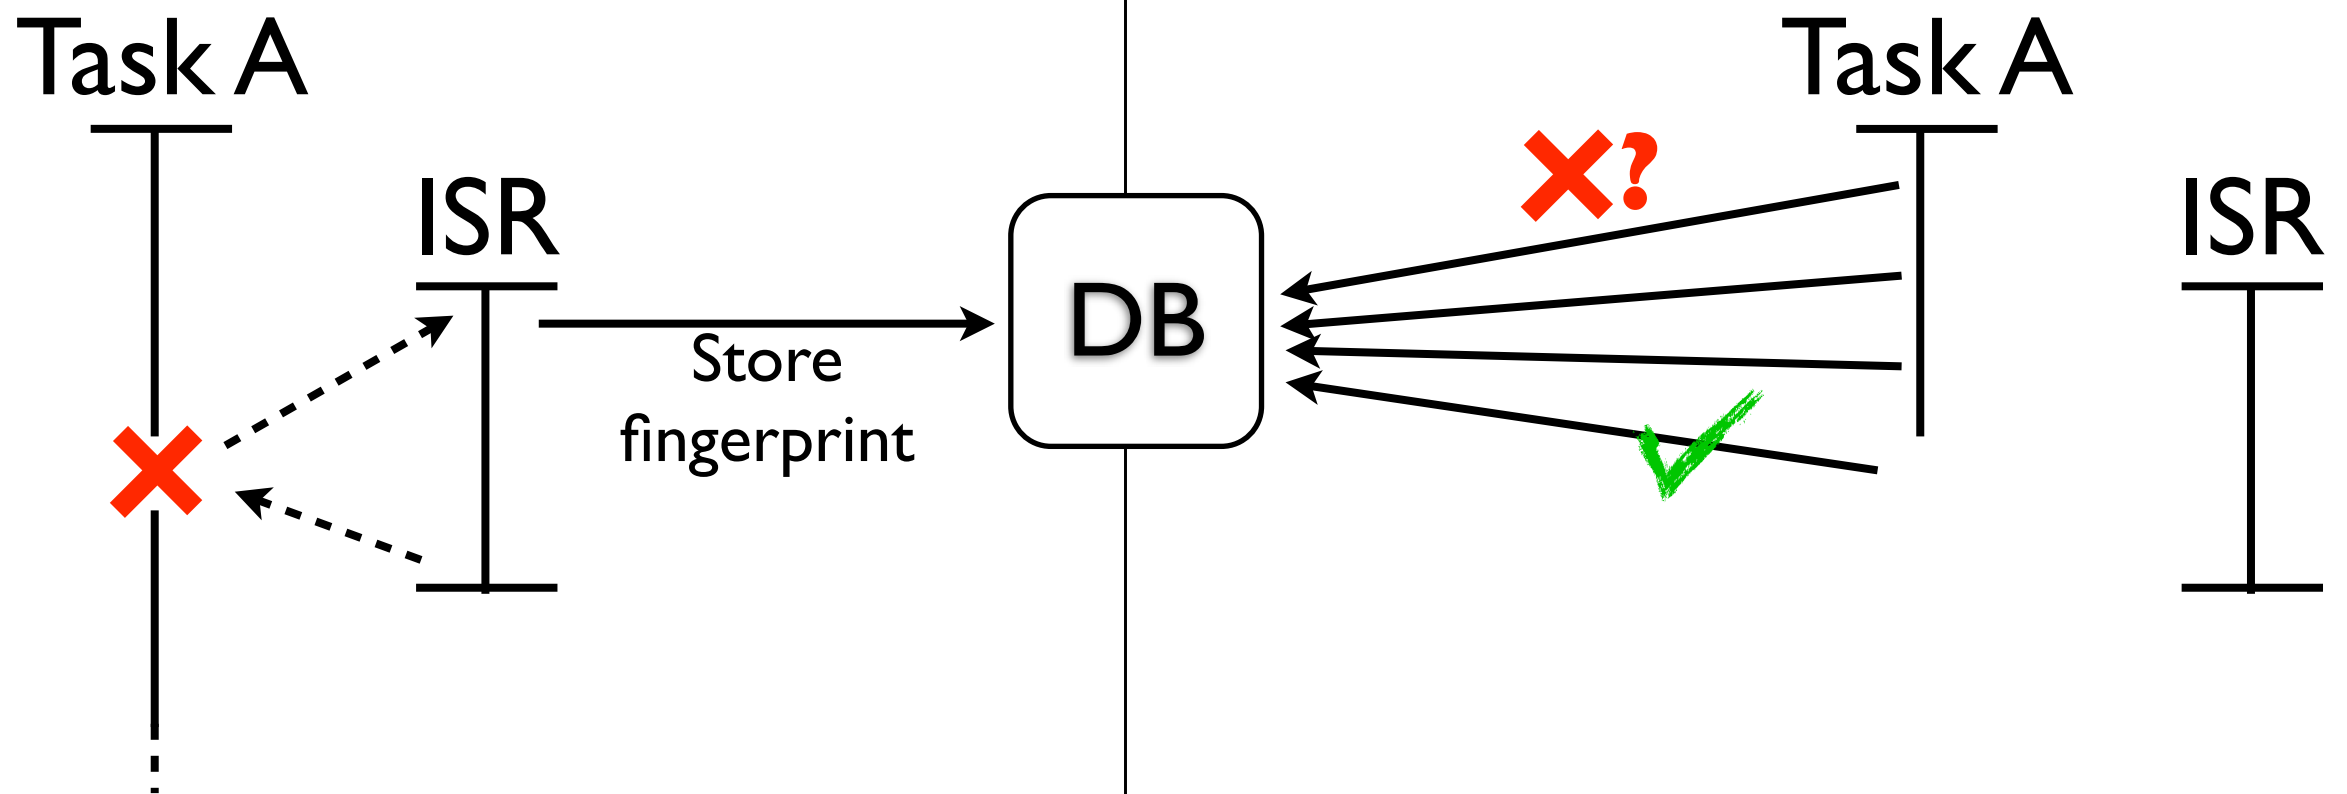


**✗ Non-determinism happens (e.g., interrupt)**

# Replay Debugging

Recording at run time

Replay offline (sim)



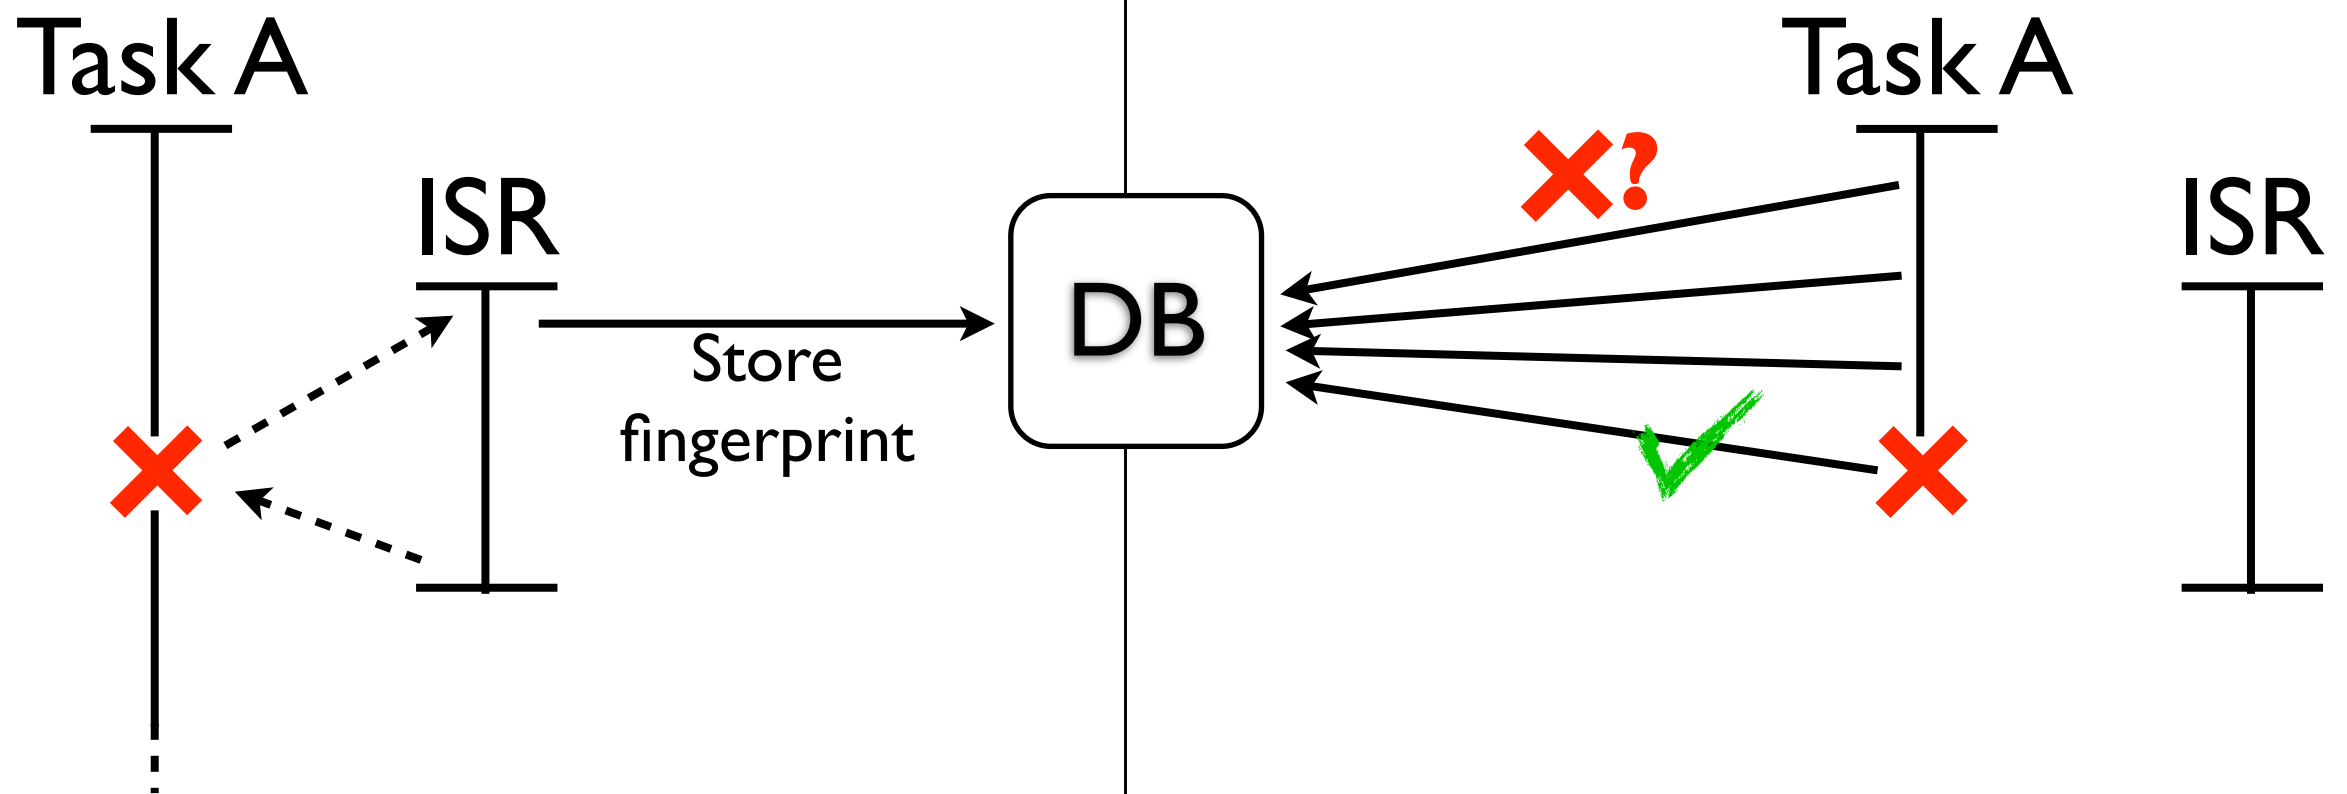
**✗ Non-determinism happens (e.g., interrupt)**



# Replay Debugging

Recording at run time

Replay offline (sim)

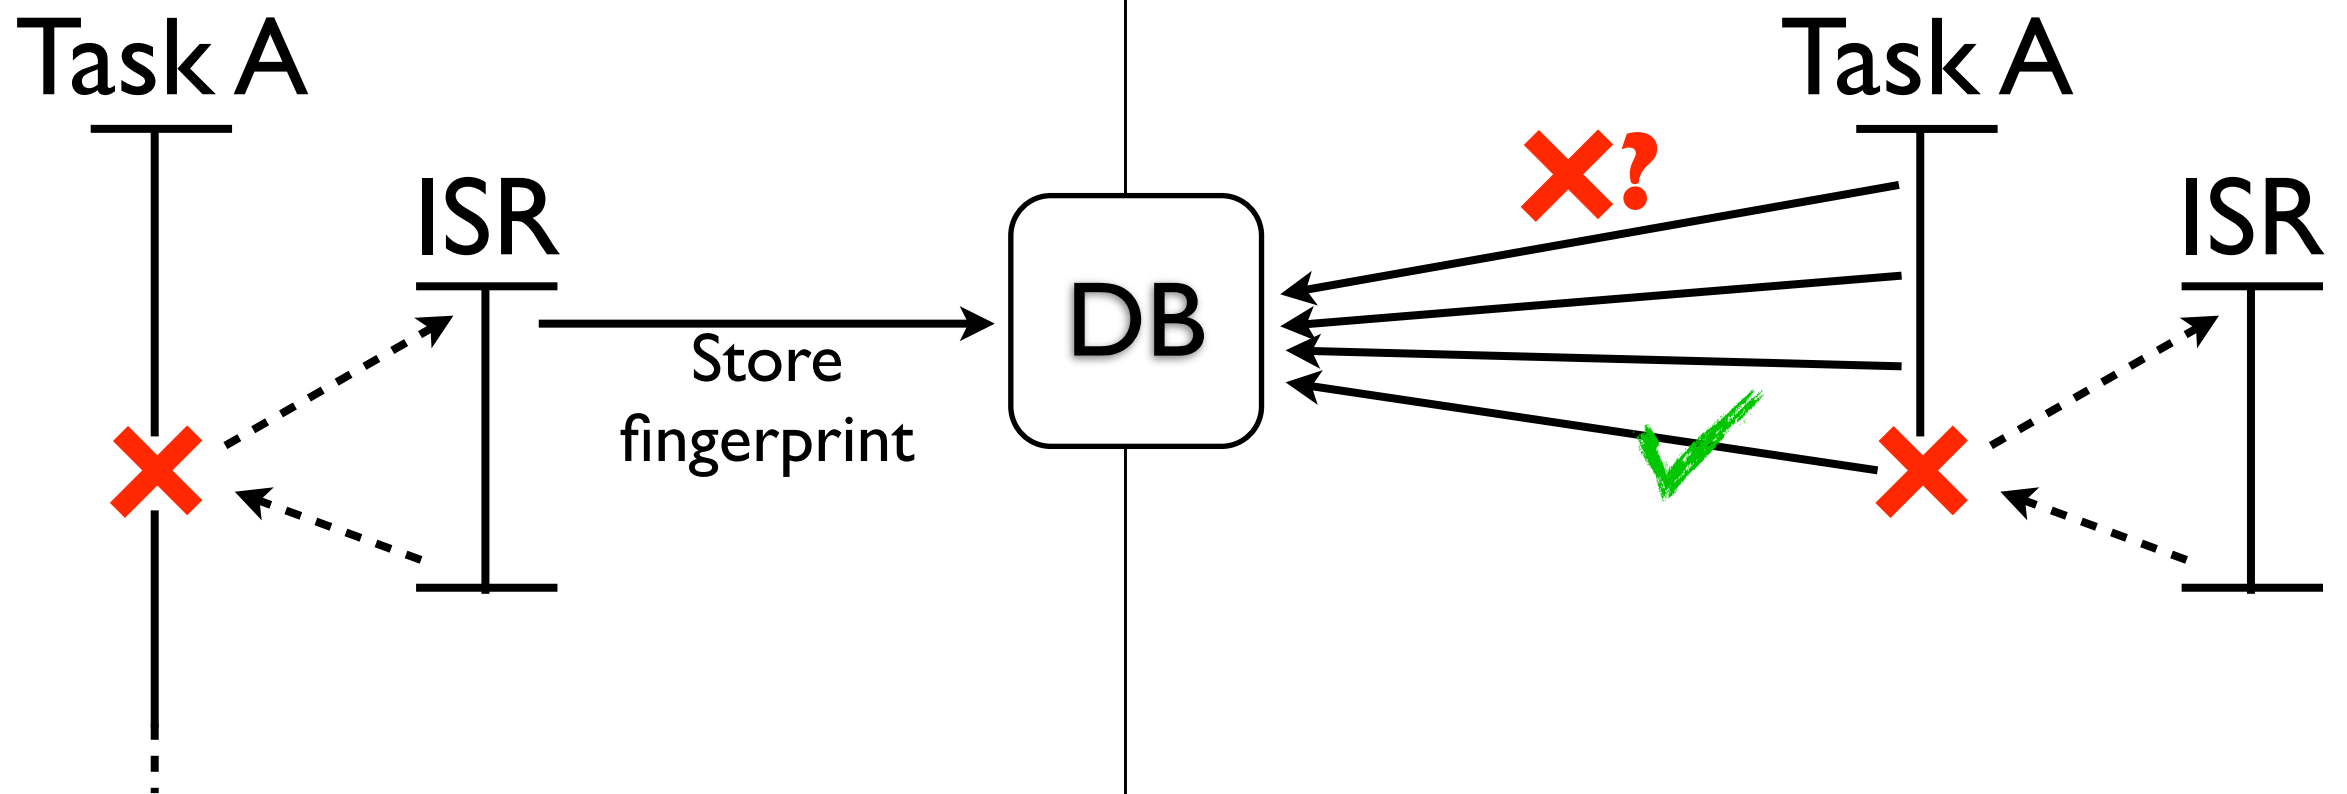


**✗ Non-determinism happens (e.g., interrupt)**

# Replay Debugging

Recording at run time

Replay offline (sim)

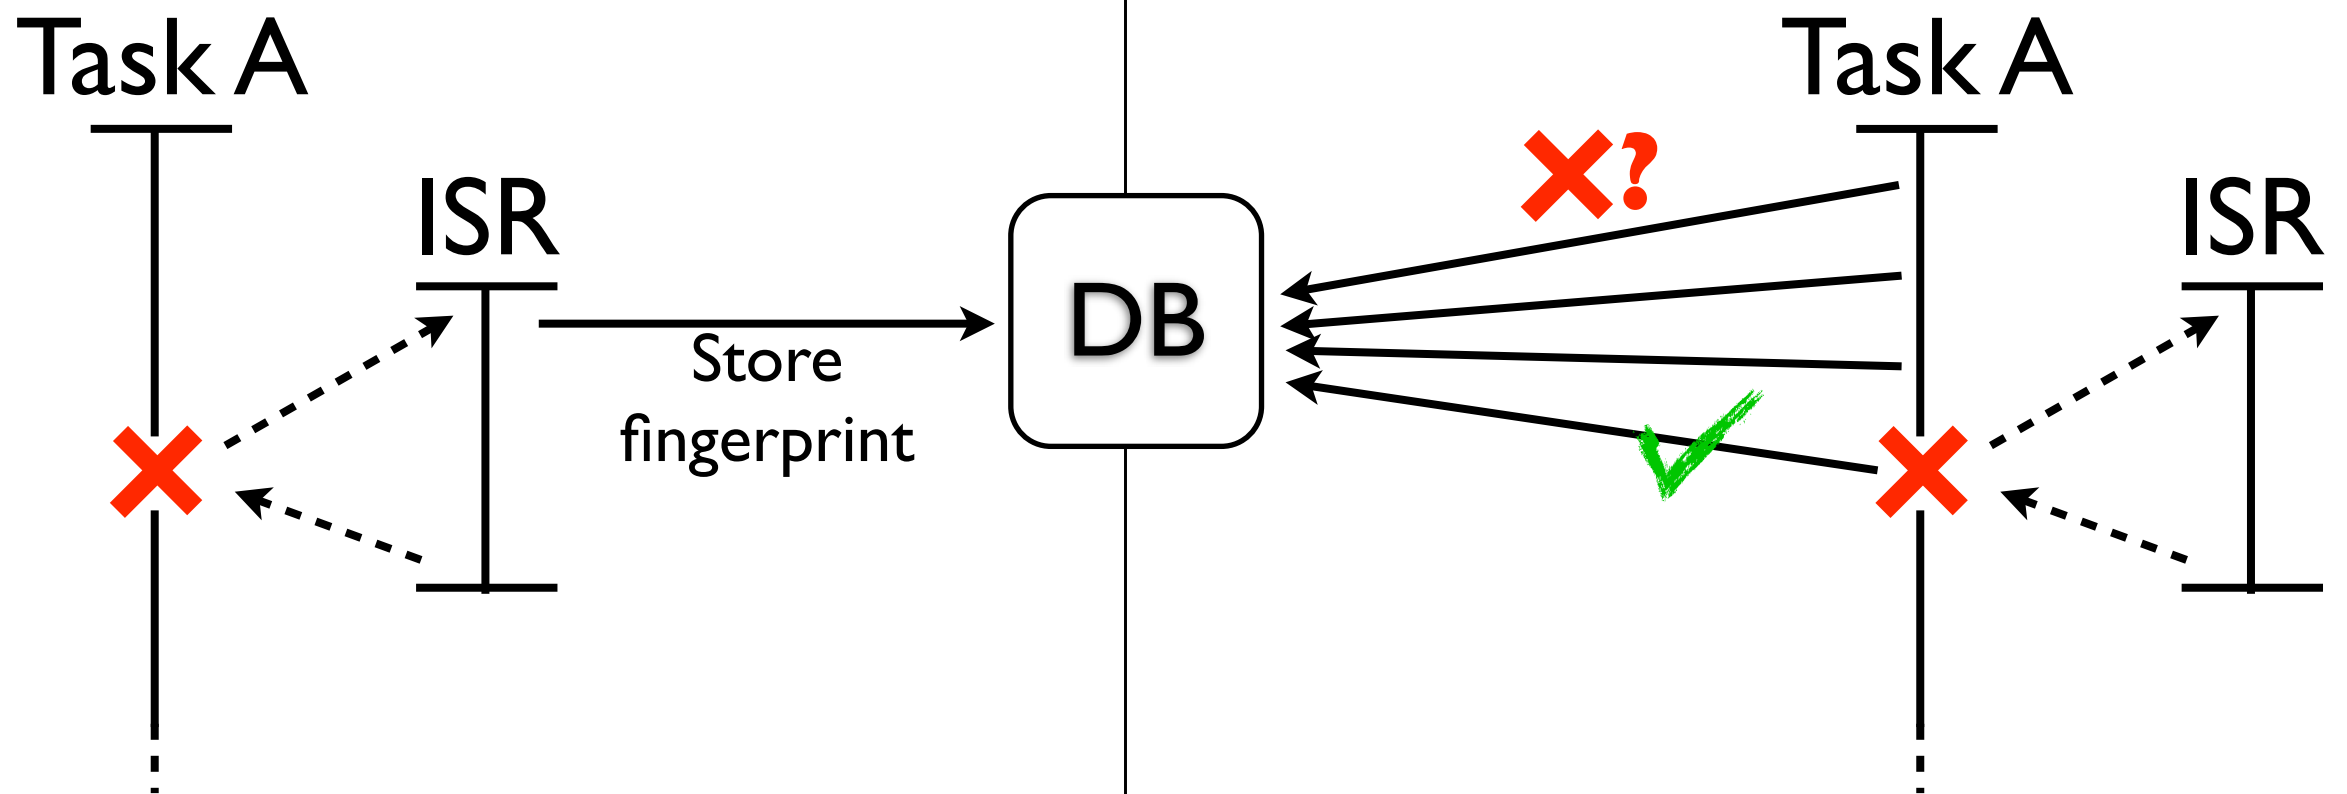


**X** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

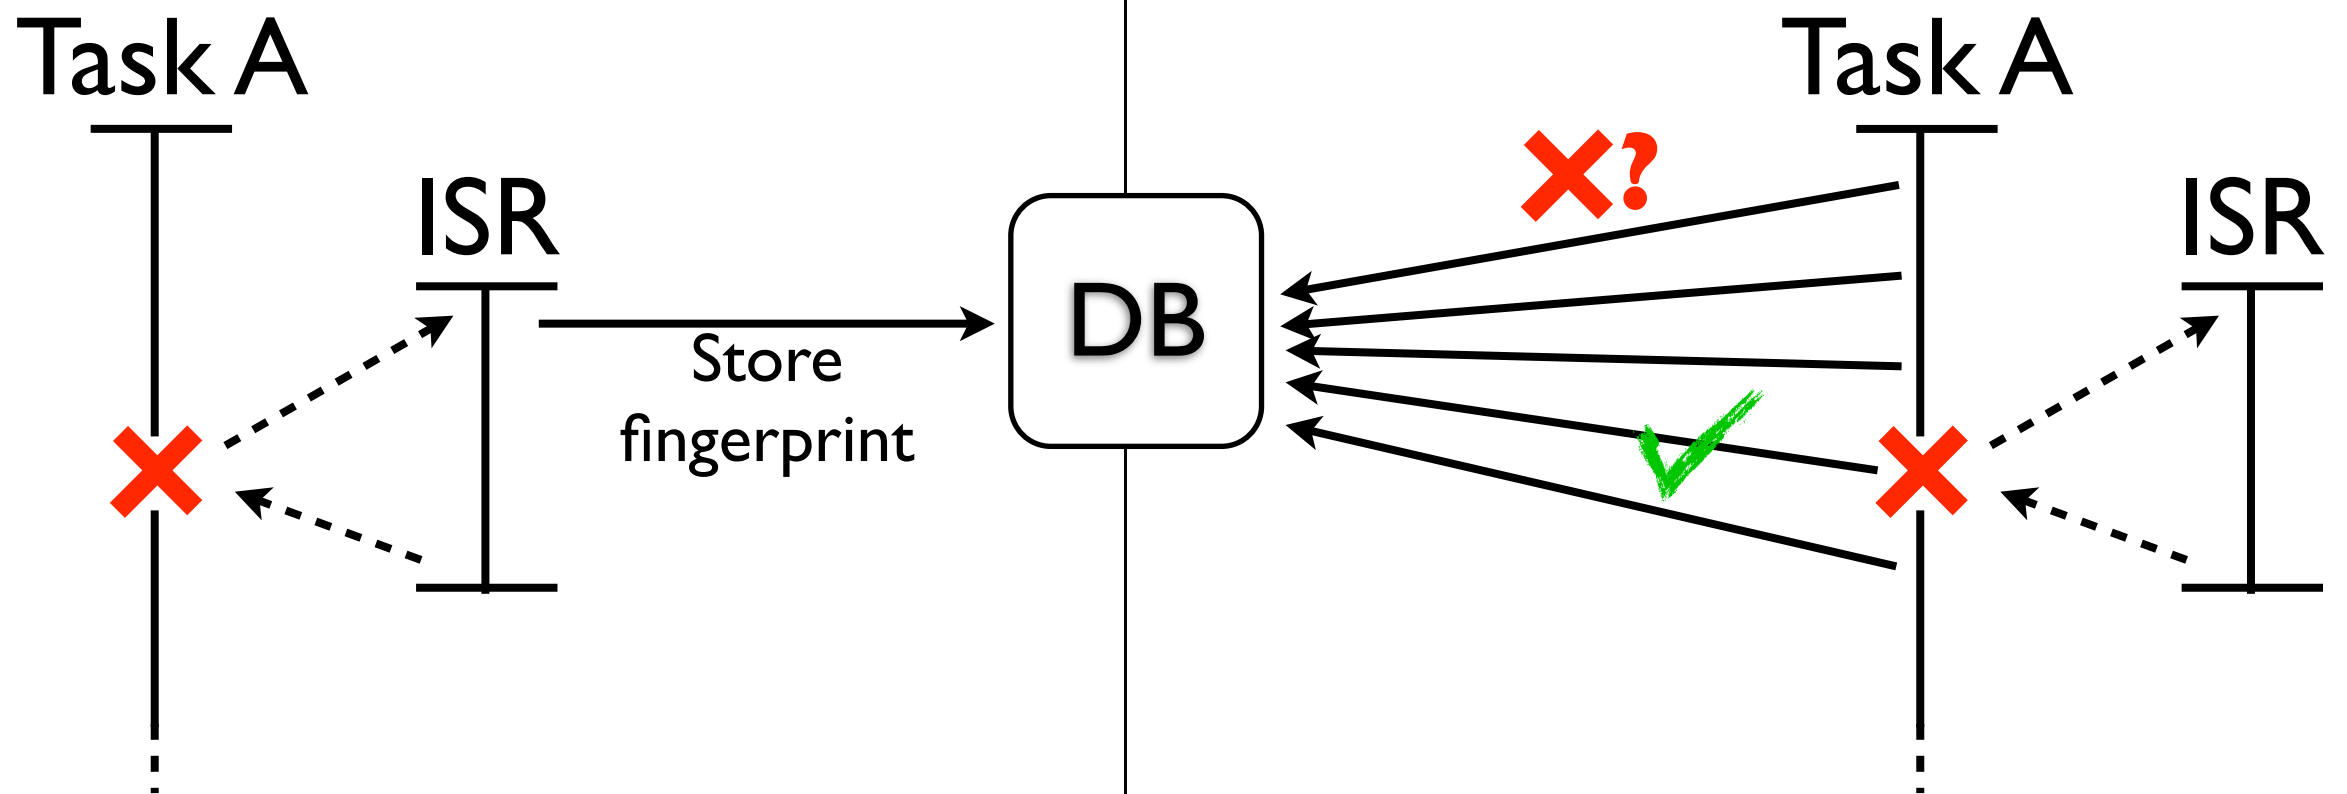


**X** Non-determinism happens (e.g., interrupt)

# Replay Debugging

Recording at run time

Replay offline (sim)

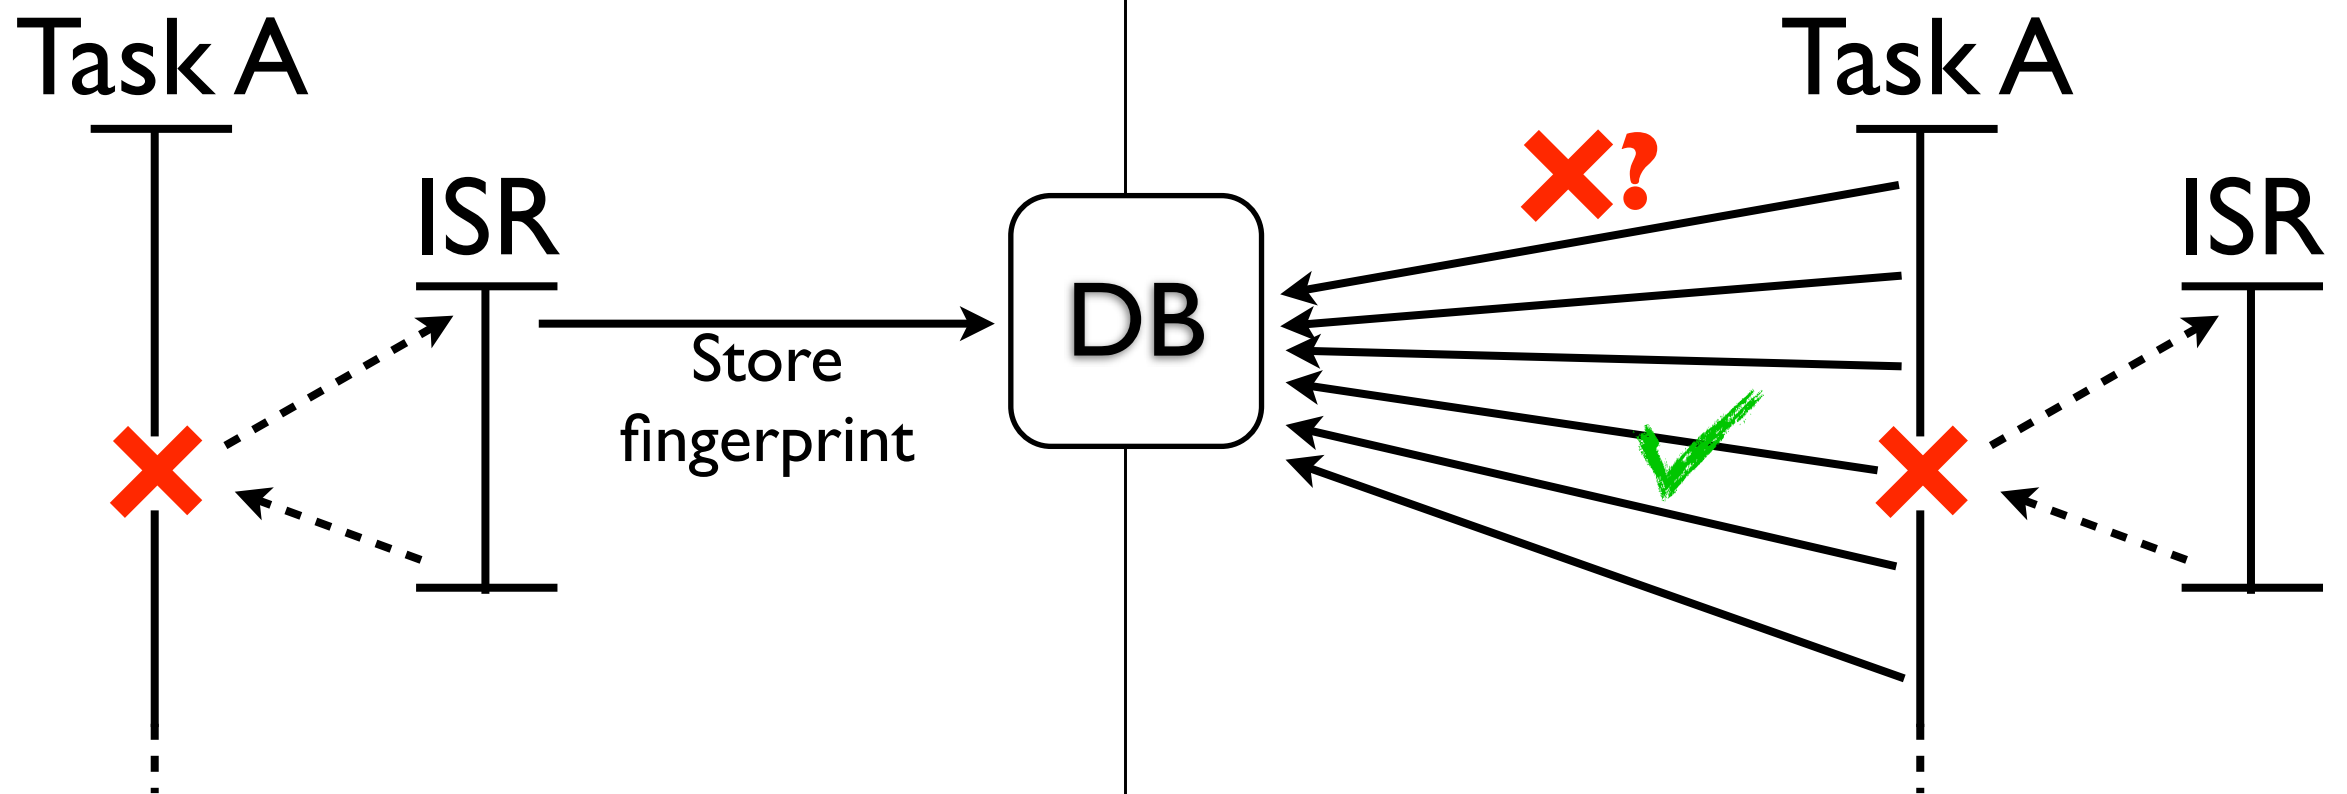


**✗ Non-determinism happens (e.g., interrupt)**

# Replay Debugging

Recording at run time

Replay offline (sim)



**✗** Non-determinism happens (e.g., interrupt)


# Replay Debugging

- **Benefits**
  - **Communicate and document bugs**  
(Show the trace)
  - **Go forward and backwards in time**  
(Thoroughly analyze how the bug happened)
  - **Retest specific scenario for fixes**  
(Exercise the system with the specific trace)

# Recording Phase

7


First try: record program counter



**✗** Non-determinism happens (e.g., interrupt)

# Recording Phase

First try: record program counter




**X** Non-determinism happens (e.g., interrupt)



# Recording Phase


First try: record program counter



**✗** Non-determinism happens (e.g., interrupt)

# Recording Phase

First try: record program counter




Second try: record program counter  
and state information

✗ Non-determinism happens (e.g., interrupt)

# Recording Phase

First try: record program counter



Second try: record program counter  
and state information



✗ Non-determinism happens (e.g., interrupt)

# Fingerprinting

9

## Compress system state

**System state** 512+ bits

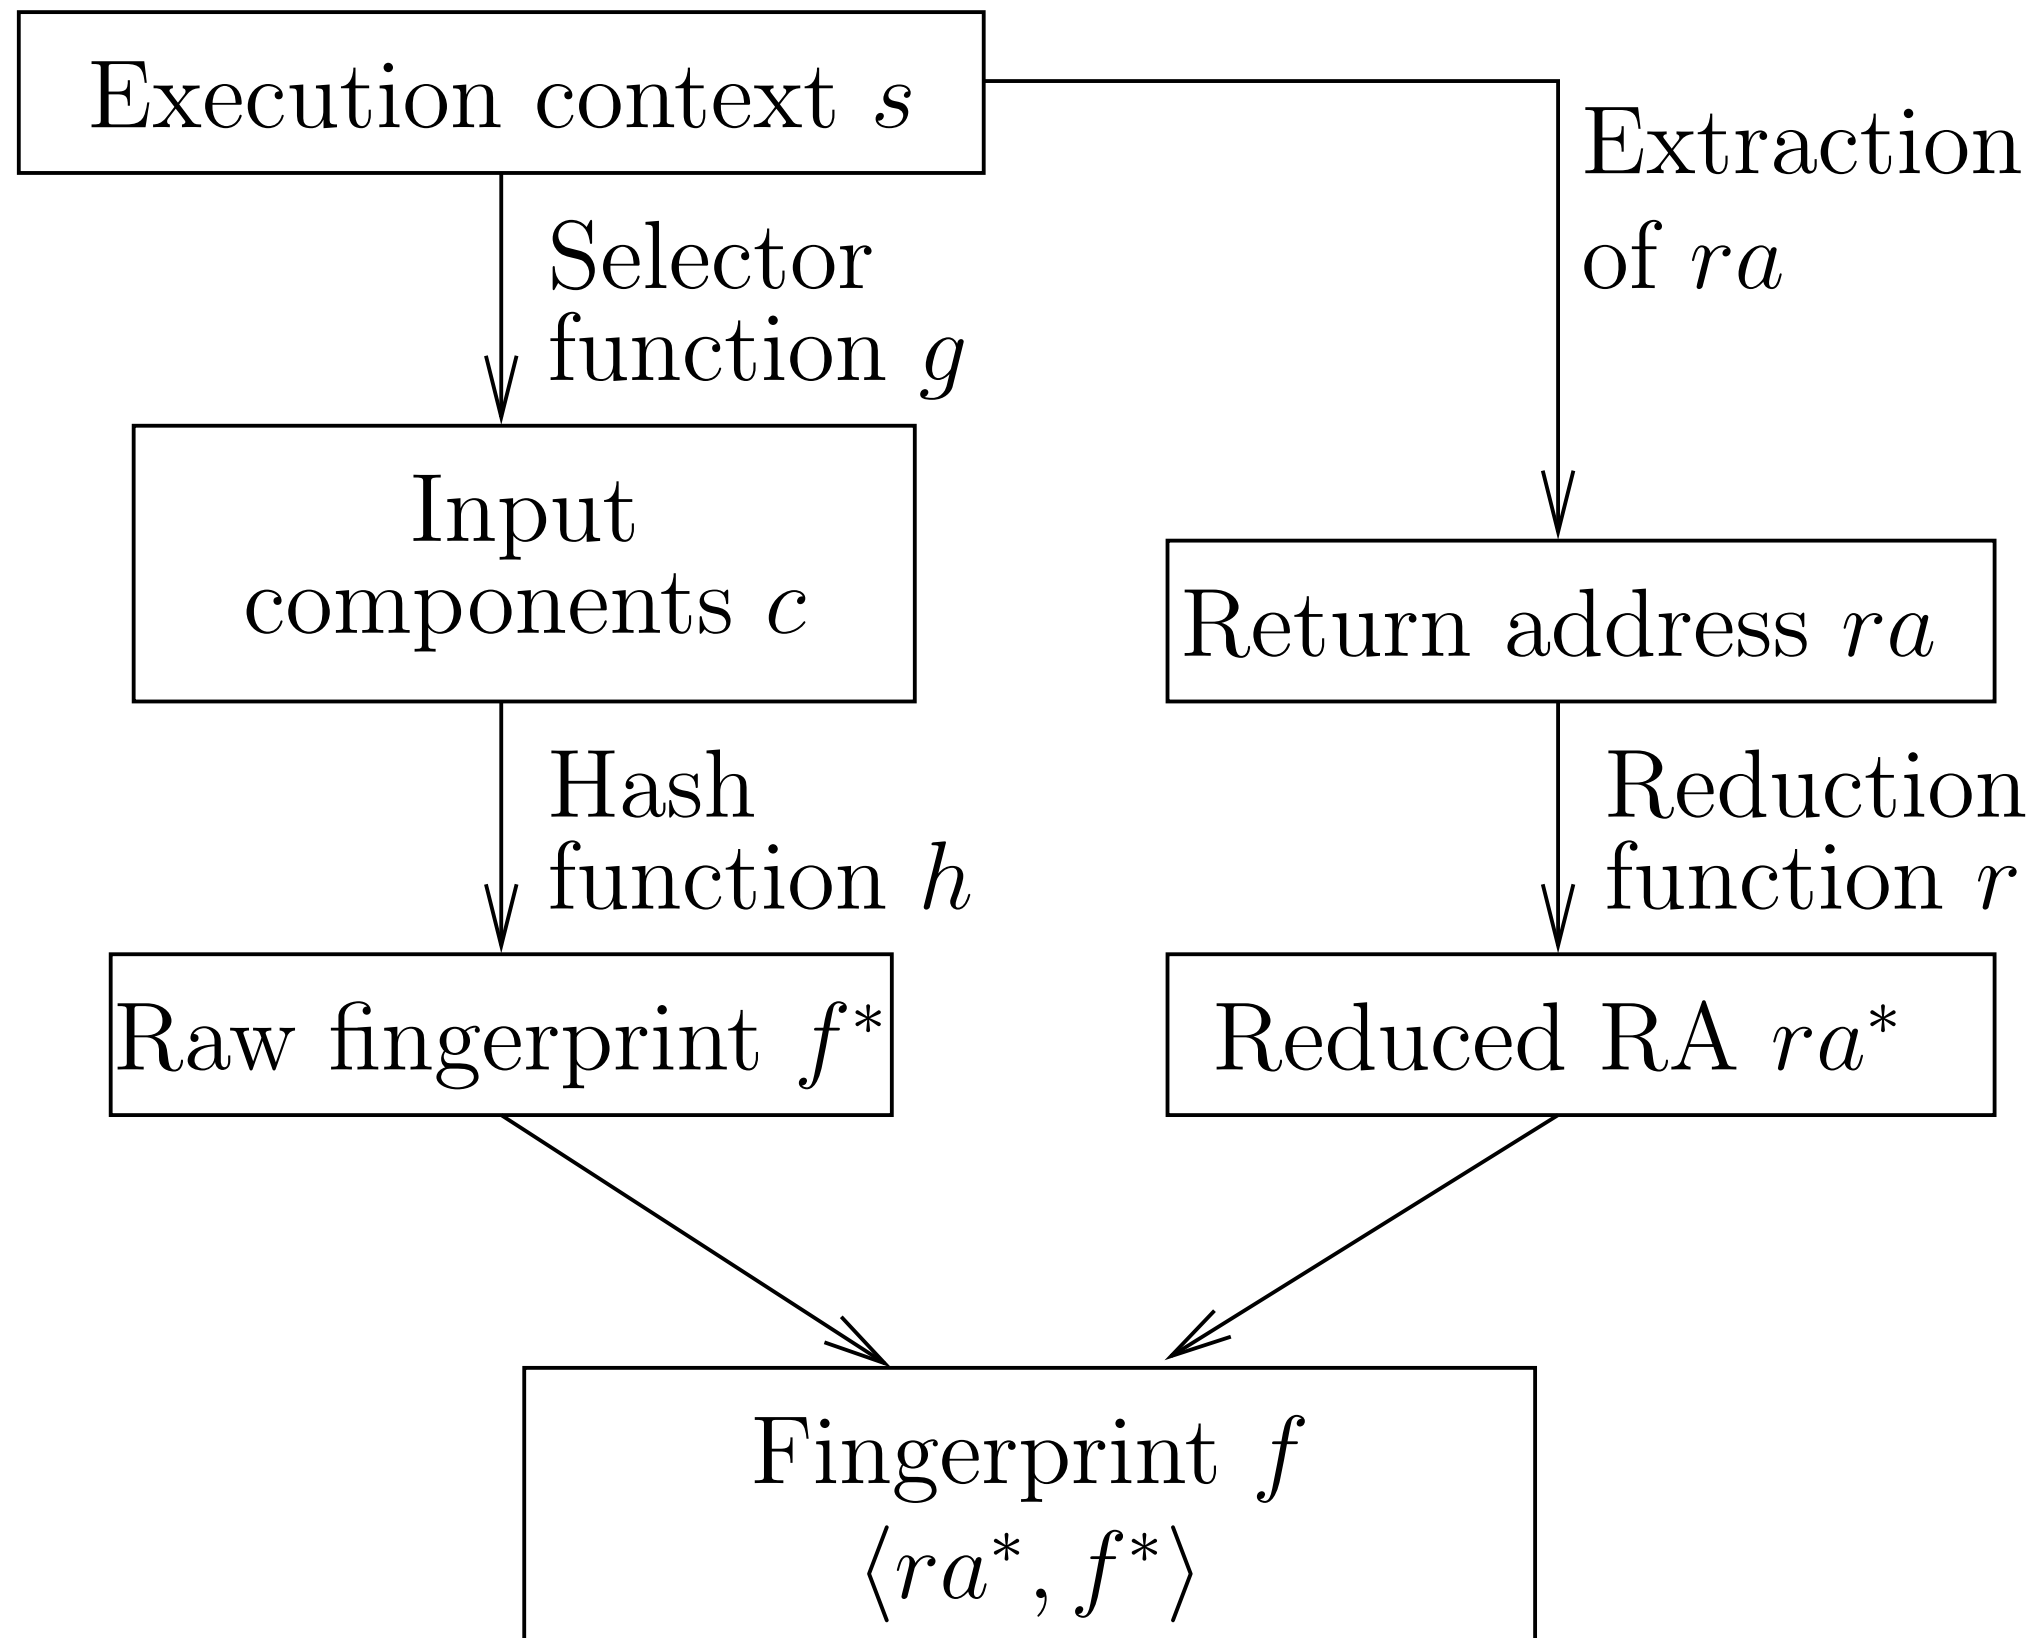


**Fingerprint** <32 bits

01001000100100100101110

# System Model

10



# Caveats

- False input duplicates (*selected bad subset*)

$$s \neq s' \wedge g(s) = g(s')$$

- False loop positives (*falsely believe IRQ happened here*)

$$s \neq s' \wedge s.ra = s'.ra \text{ but also } (h \cdot g)(s) = (h \cdot g)(s')$$

- False RA positives (*falsely believe IRQ happened here*)

$$s \neq s' \wedge s.ra \neq s'.ra \text{ but also } r(s.ra) = r(s'.ra) \wedge (h \cdot g)(s) = (h \cdot g)(s')$$

# Caveats

- False input duplicates (*selected bad subset*)

$$s \neq s' \wedge g(s) = g(s')$$

- False loop positives (*falsely believe IRQ happened here*)

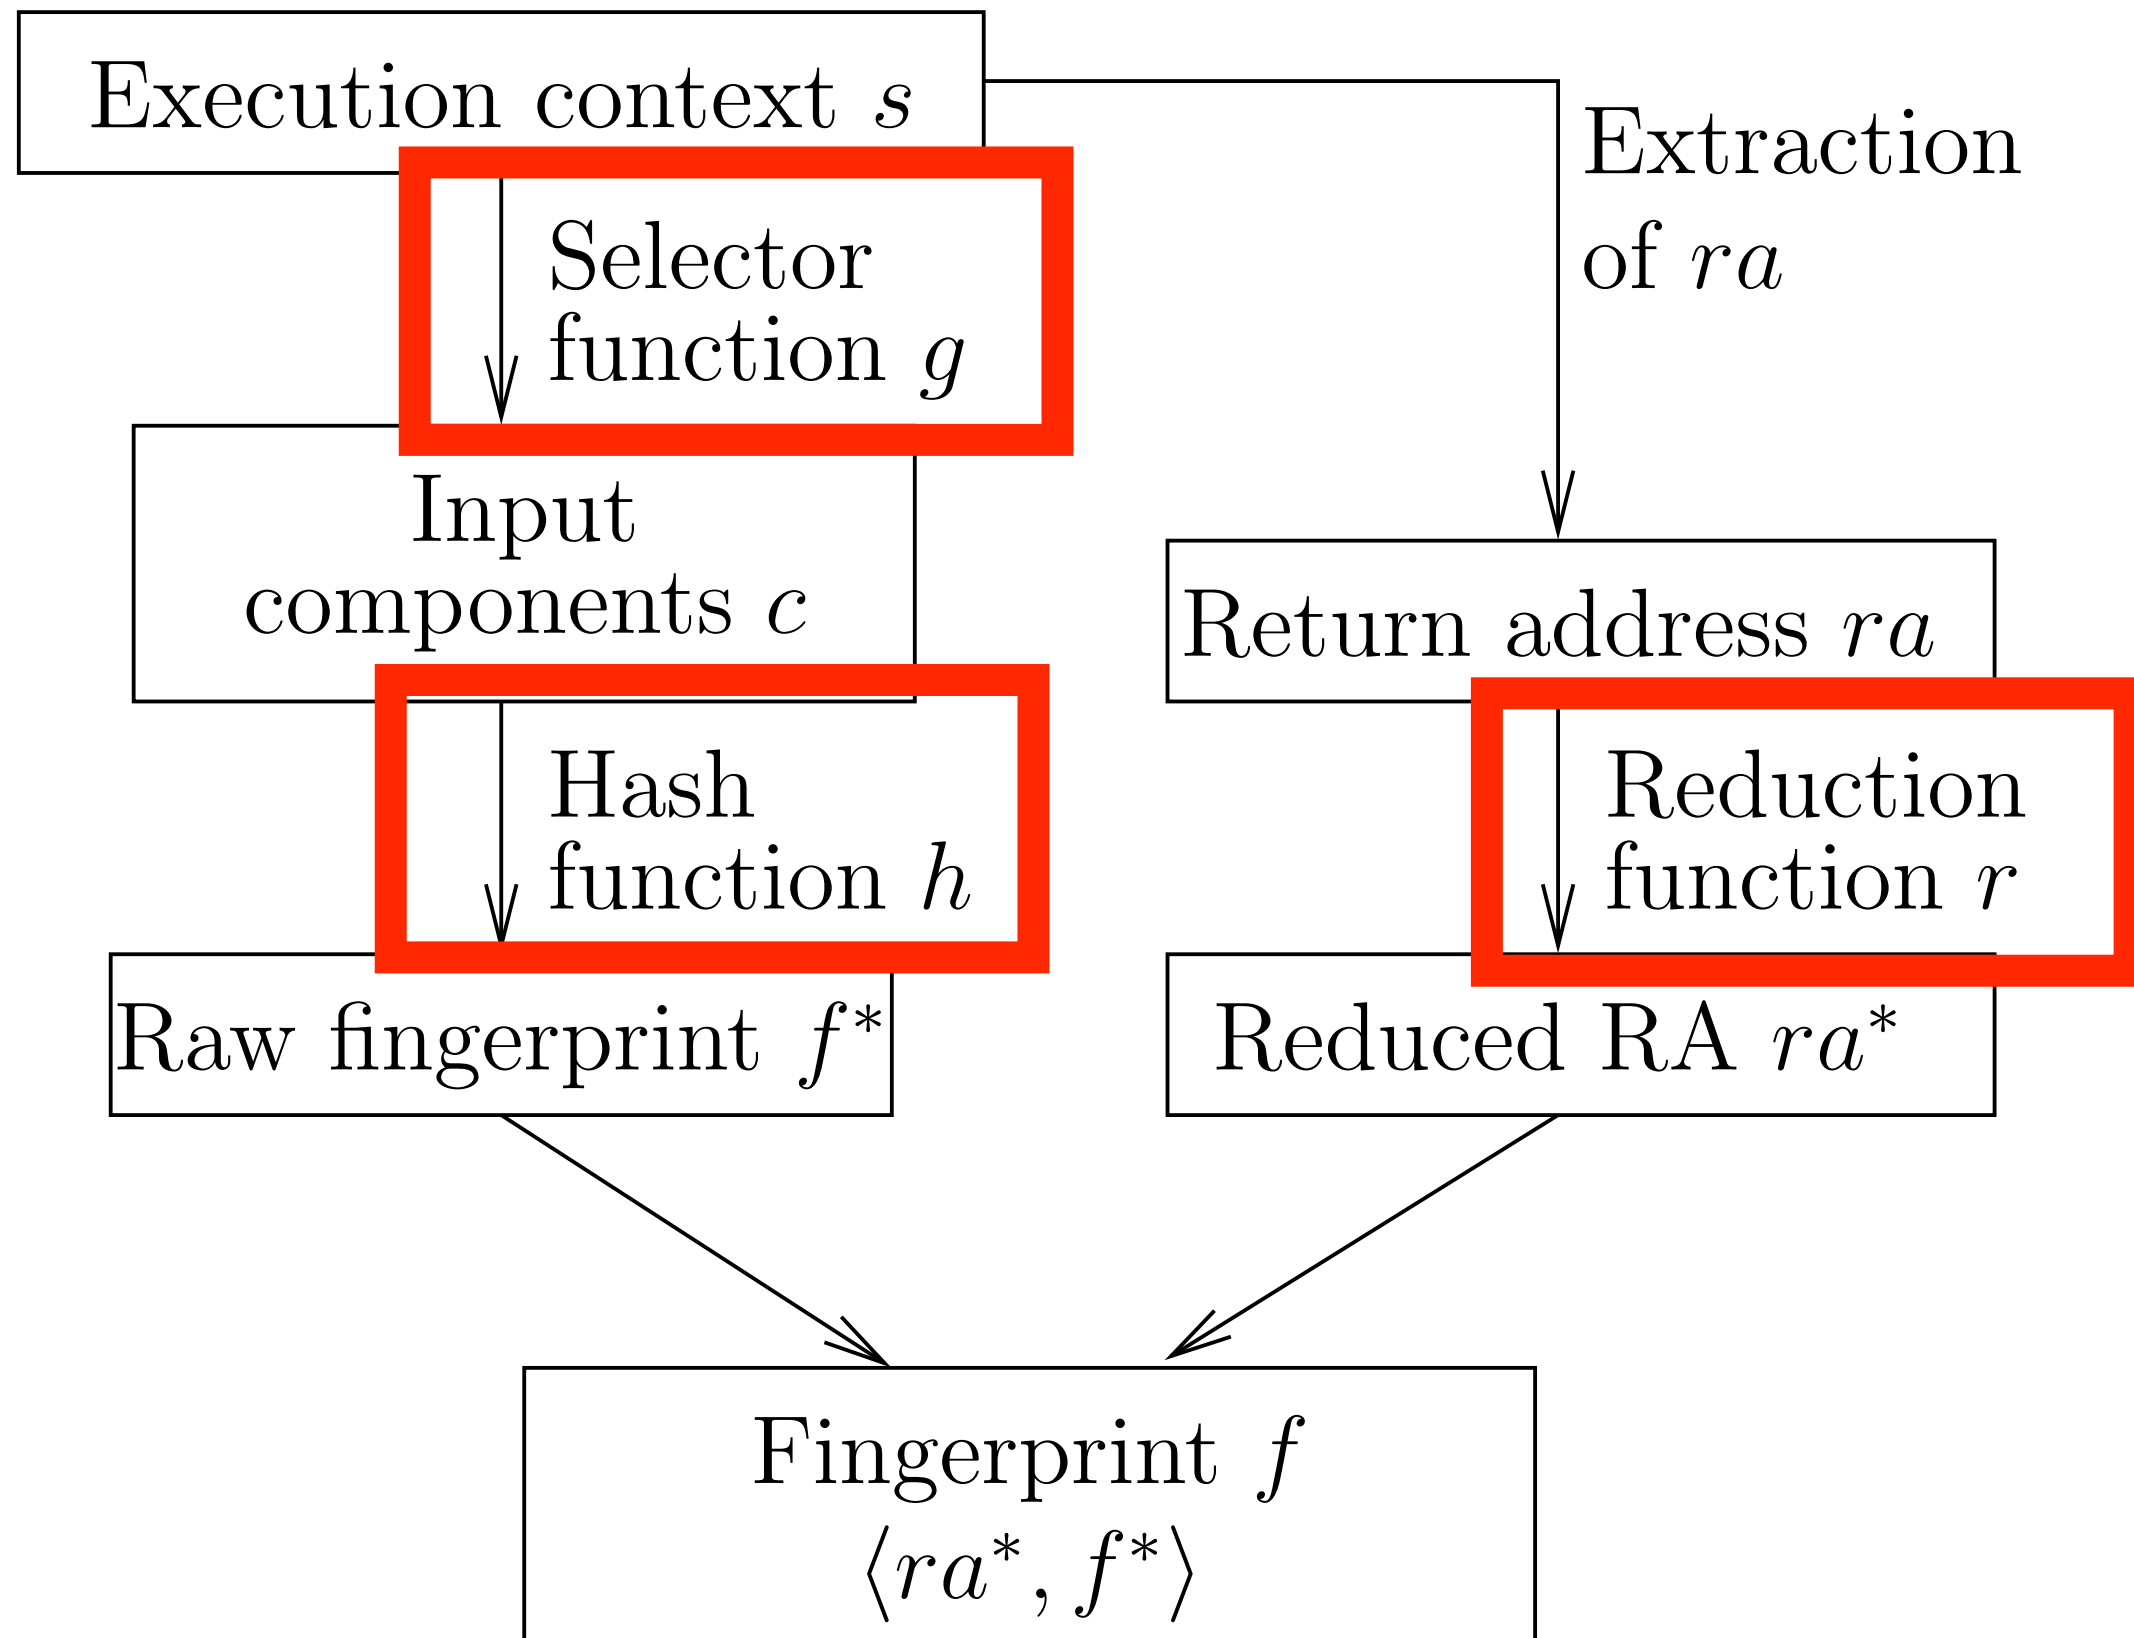
$$s \neq s' \wedge s.ra = s'.ra \text{ but also } (h \cdot g)(s) = (h \cdot g)(s')$$

- False RA positives (*falsely believe IRQ happened here*)

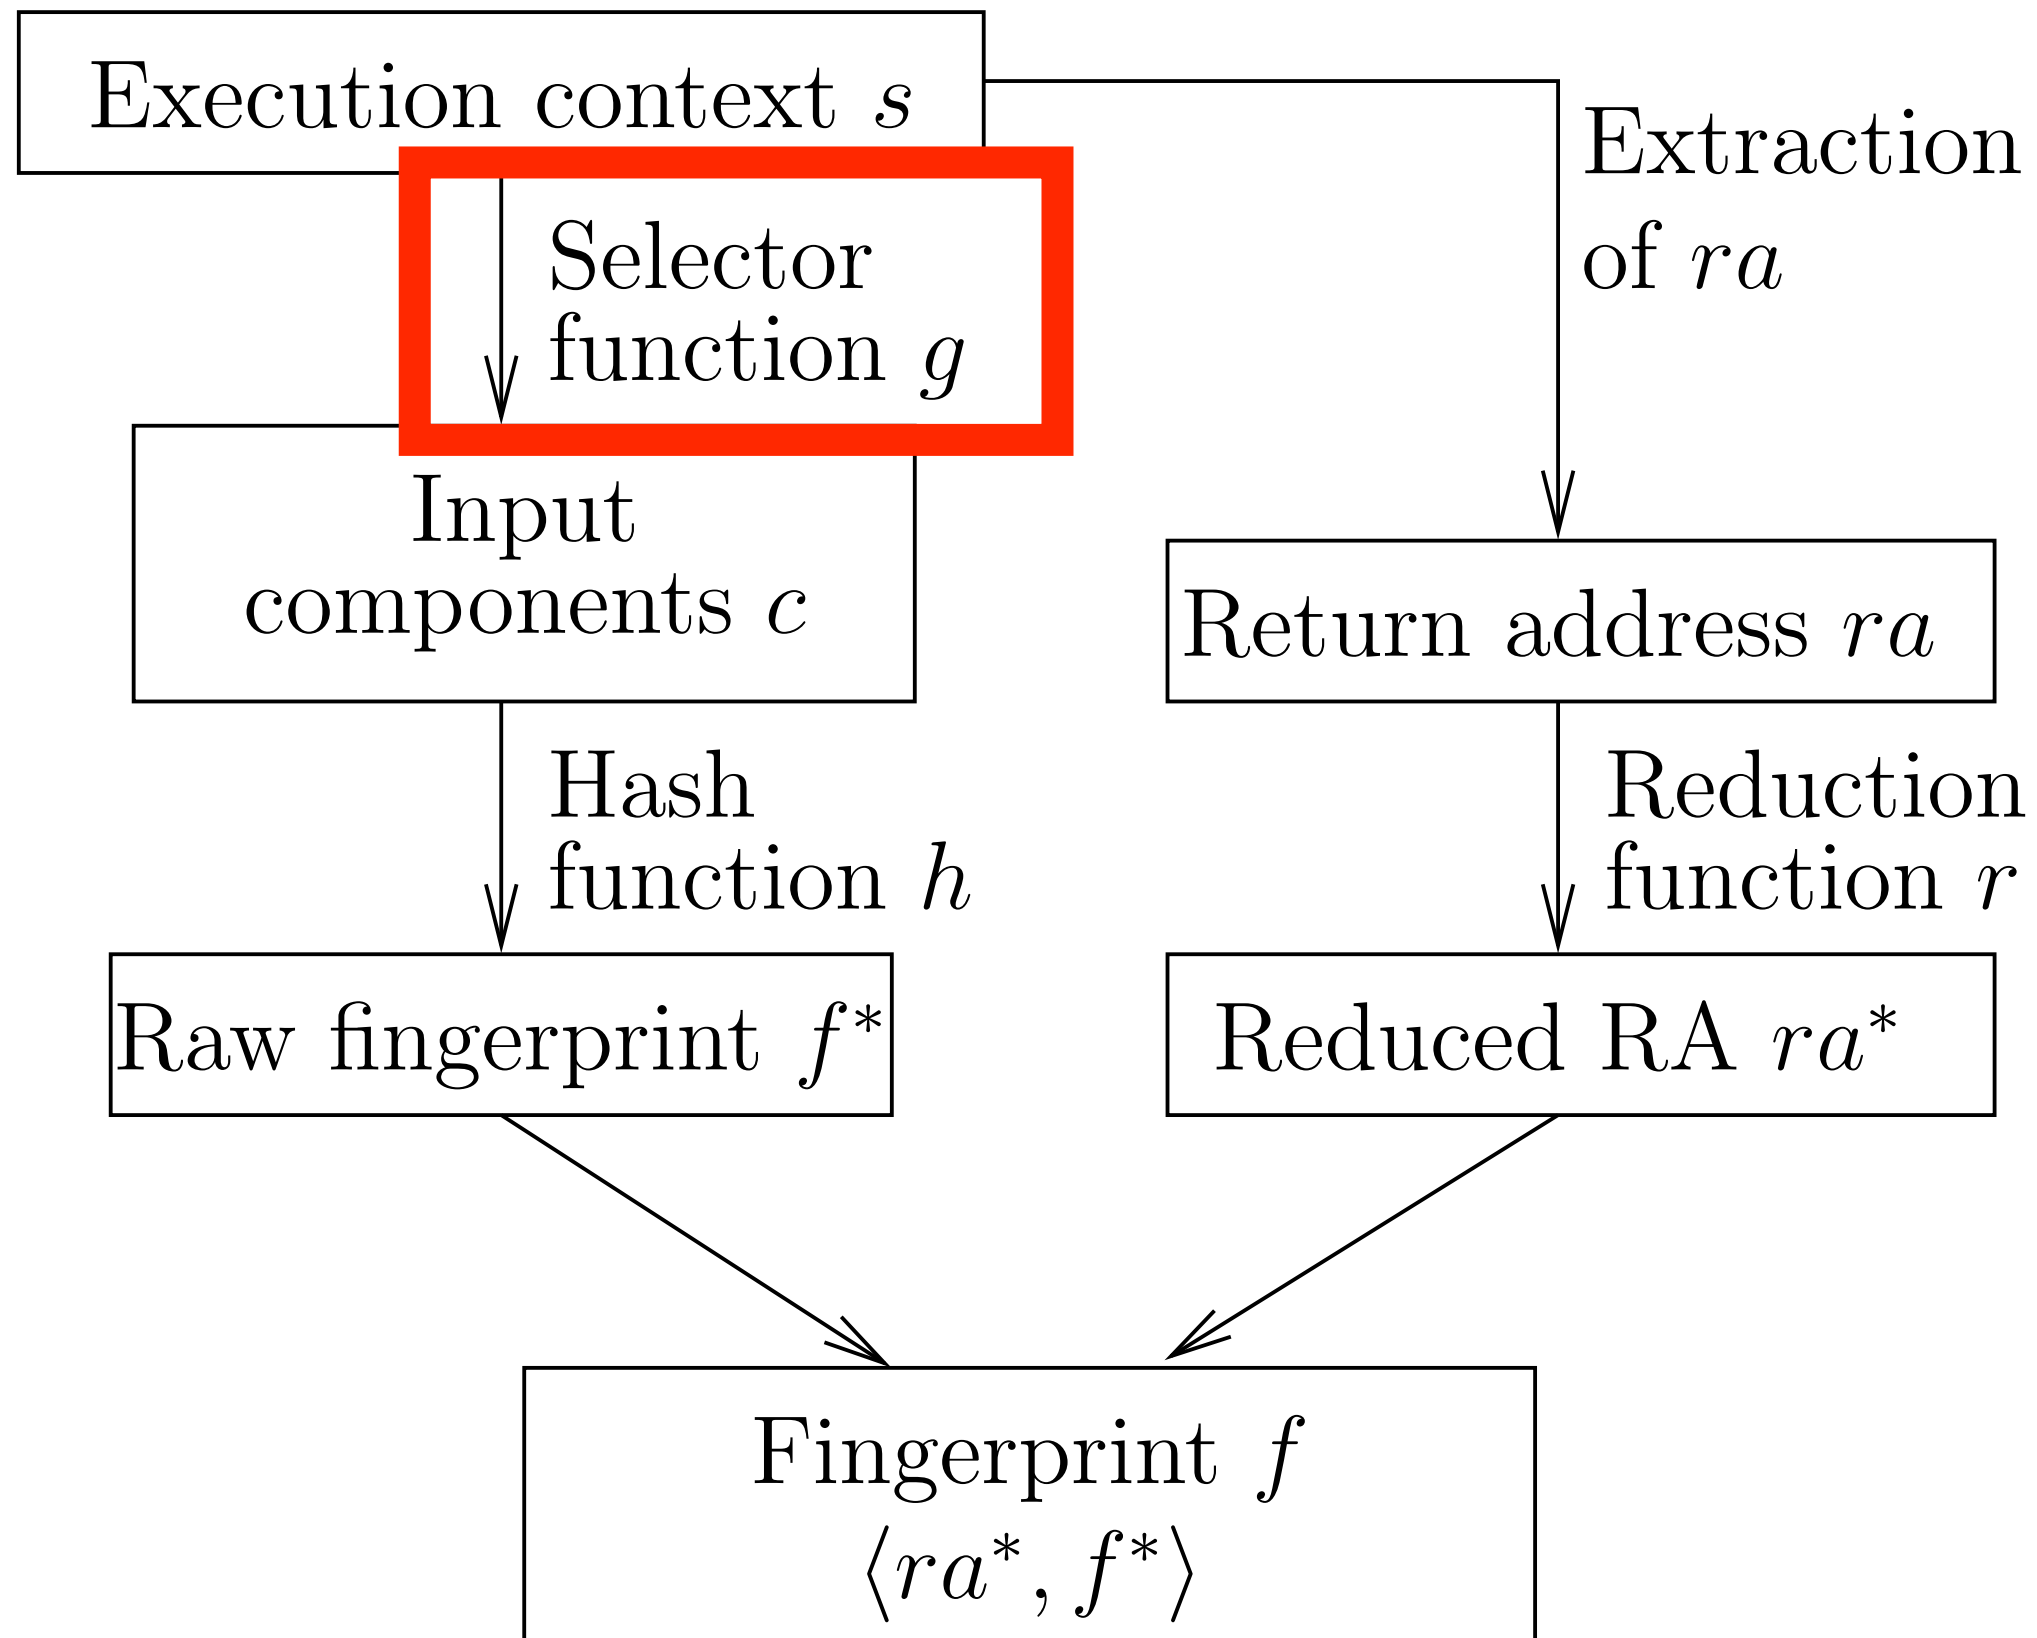
$$s \neq s' \wedge s.ra \neq s'.ra \text{ but also } r(s.ra) = r(s'.ra) \wedge (h \cdot g)(s) = (h \cdot g)(s')$$

Calls for good design of fingerprint

# Designing the Fingerprint







# Selector Function

# Selector Function

- **Many elements to choose from:** data registers, control registers, GPIO, IO status, general SDRAM, stack, heap, ...

# Selector Function

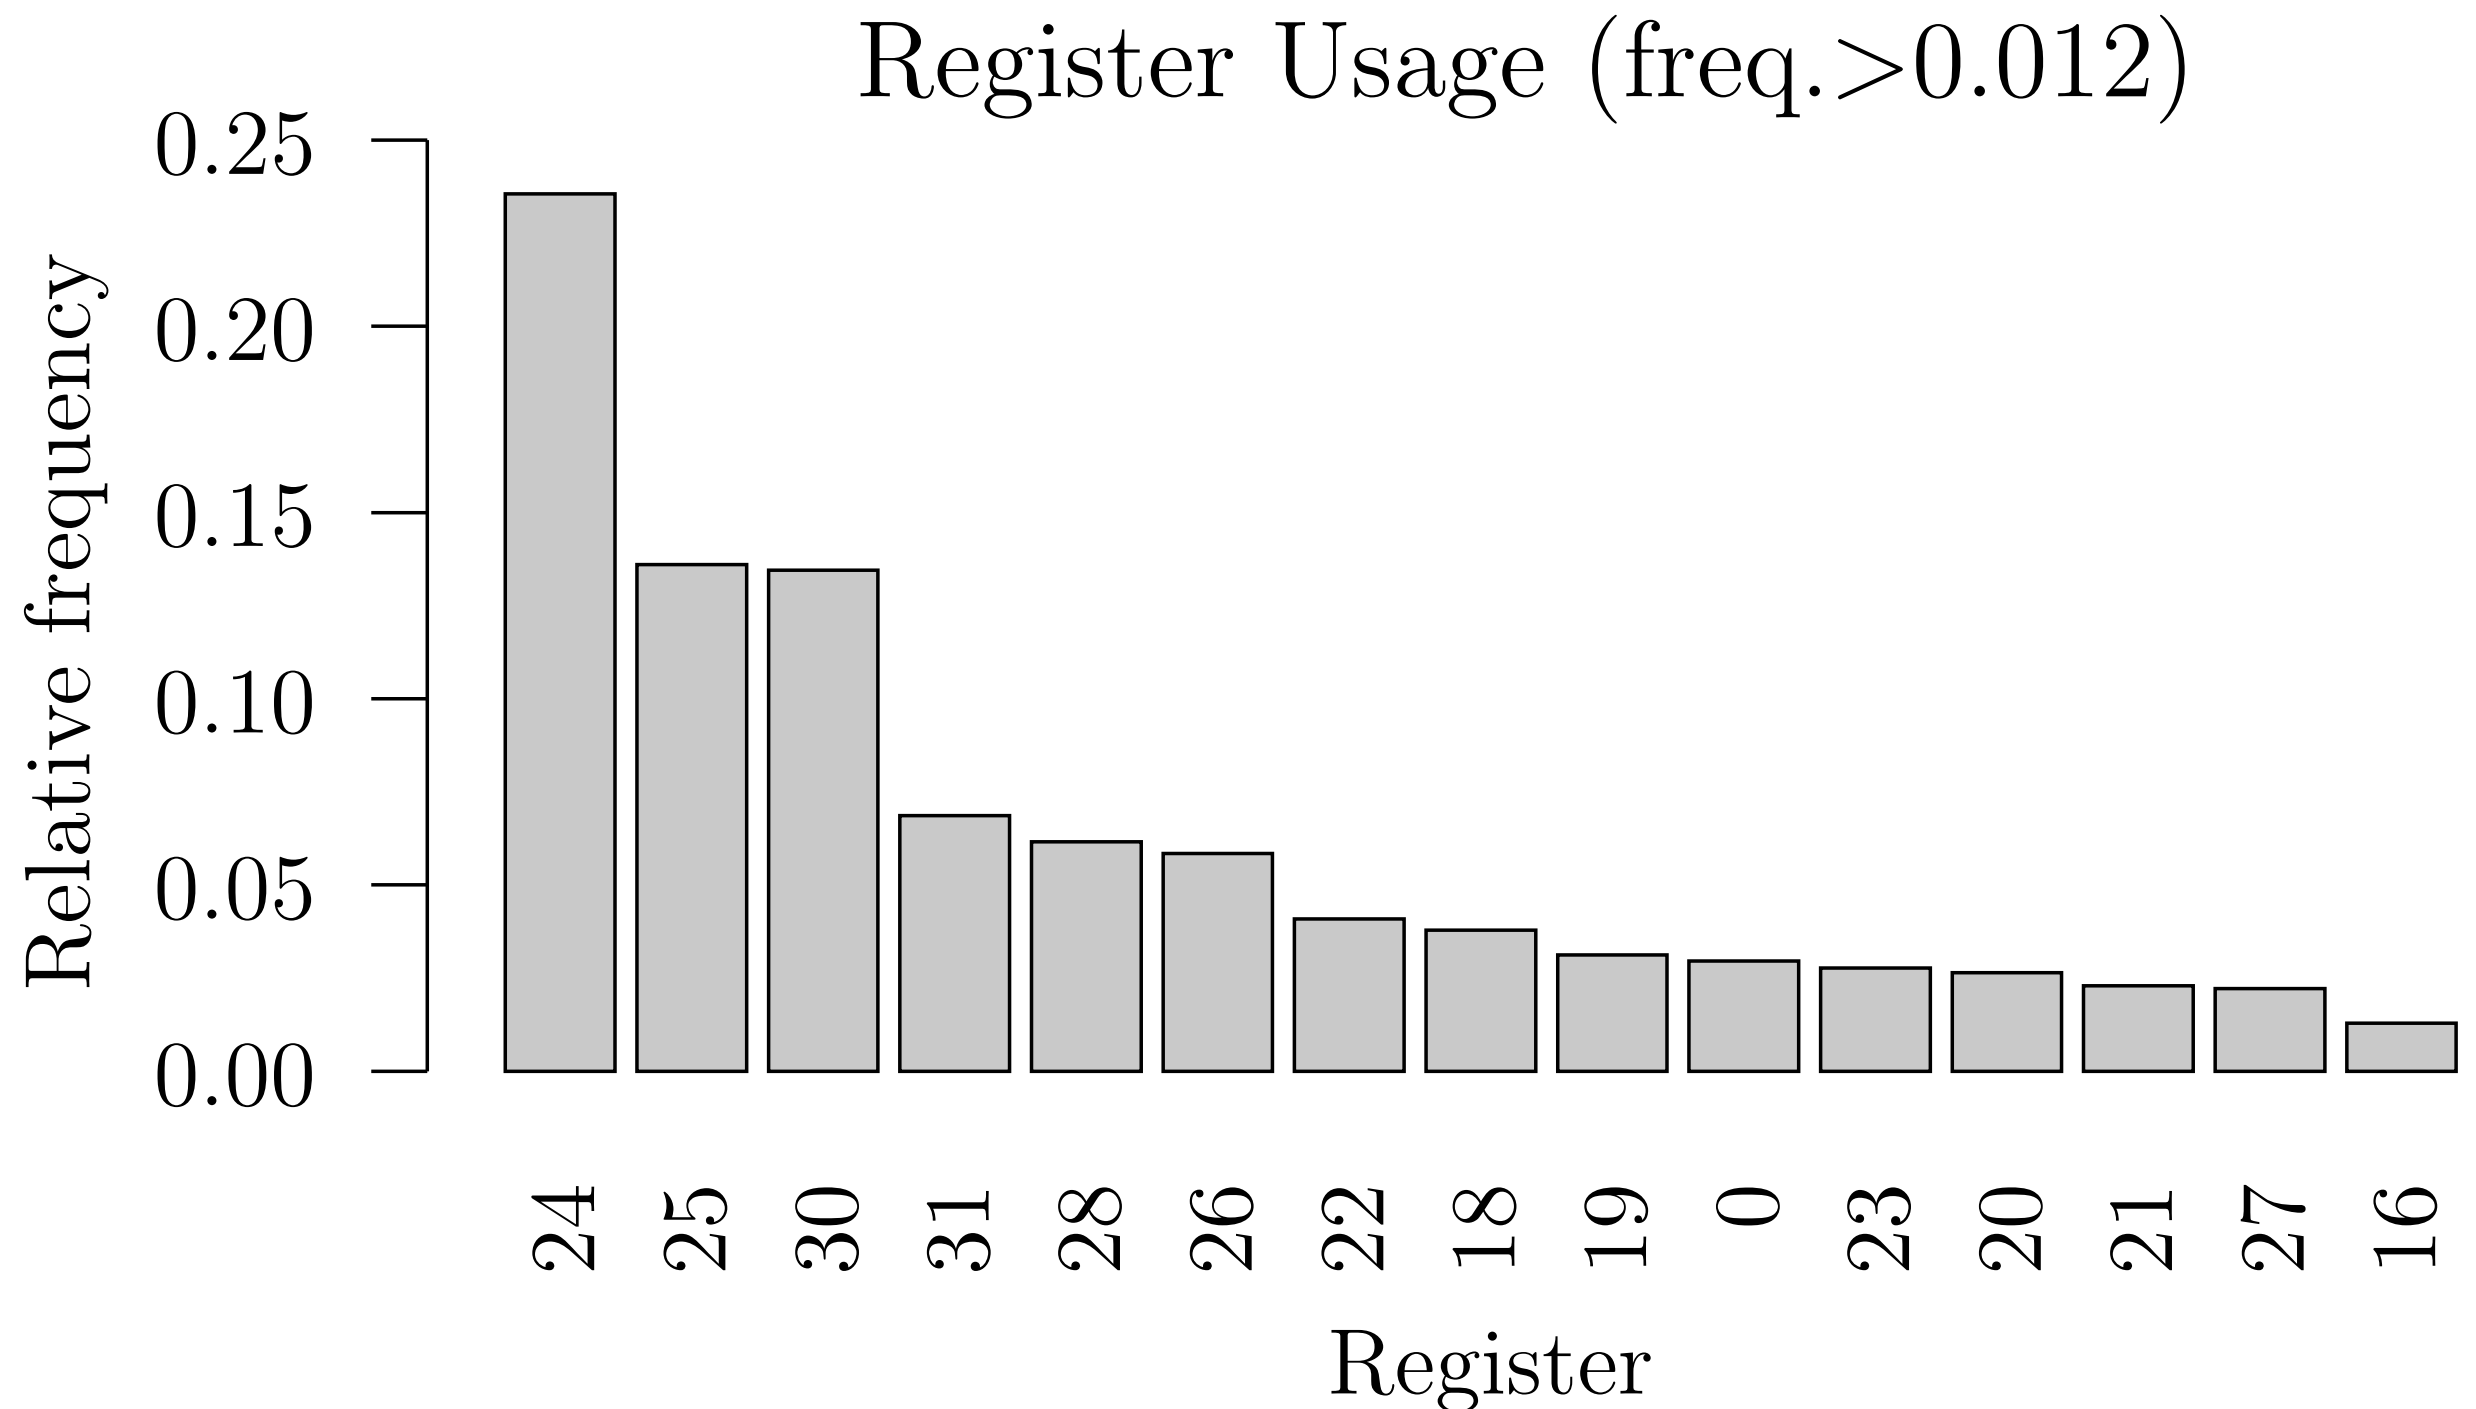
- **Many elements to choose from:** data registers, control registers, GPIO, IO status, general SDRAM, stack, heap, ...
- **Can't take all:**
  - Too much computation overhead
  - Not all are used equally often

# Idea 1

**Pick the most frequently used ones**


# Idea 1

## Pick the most frequently used ones






# Does it work?

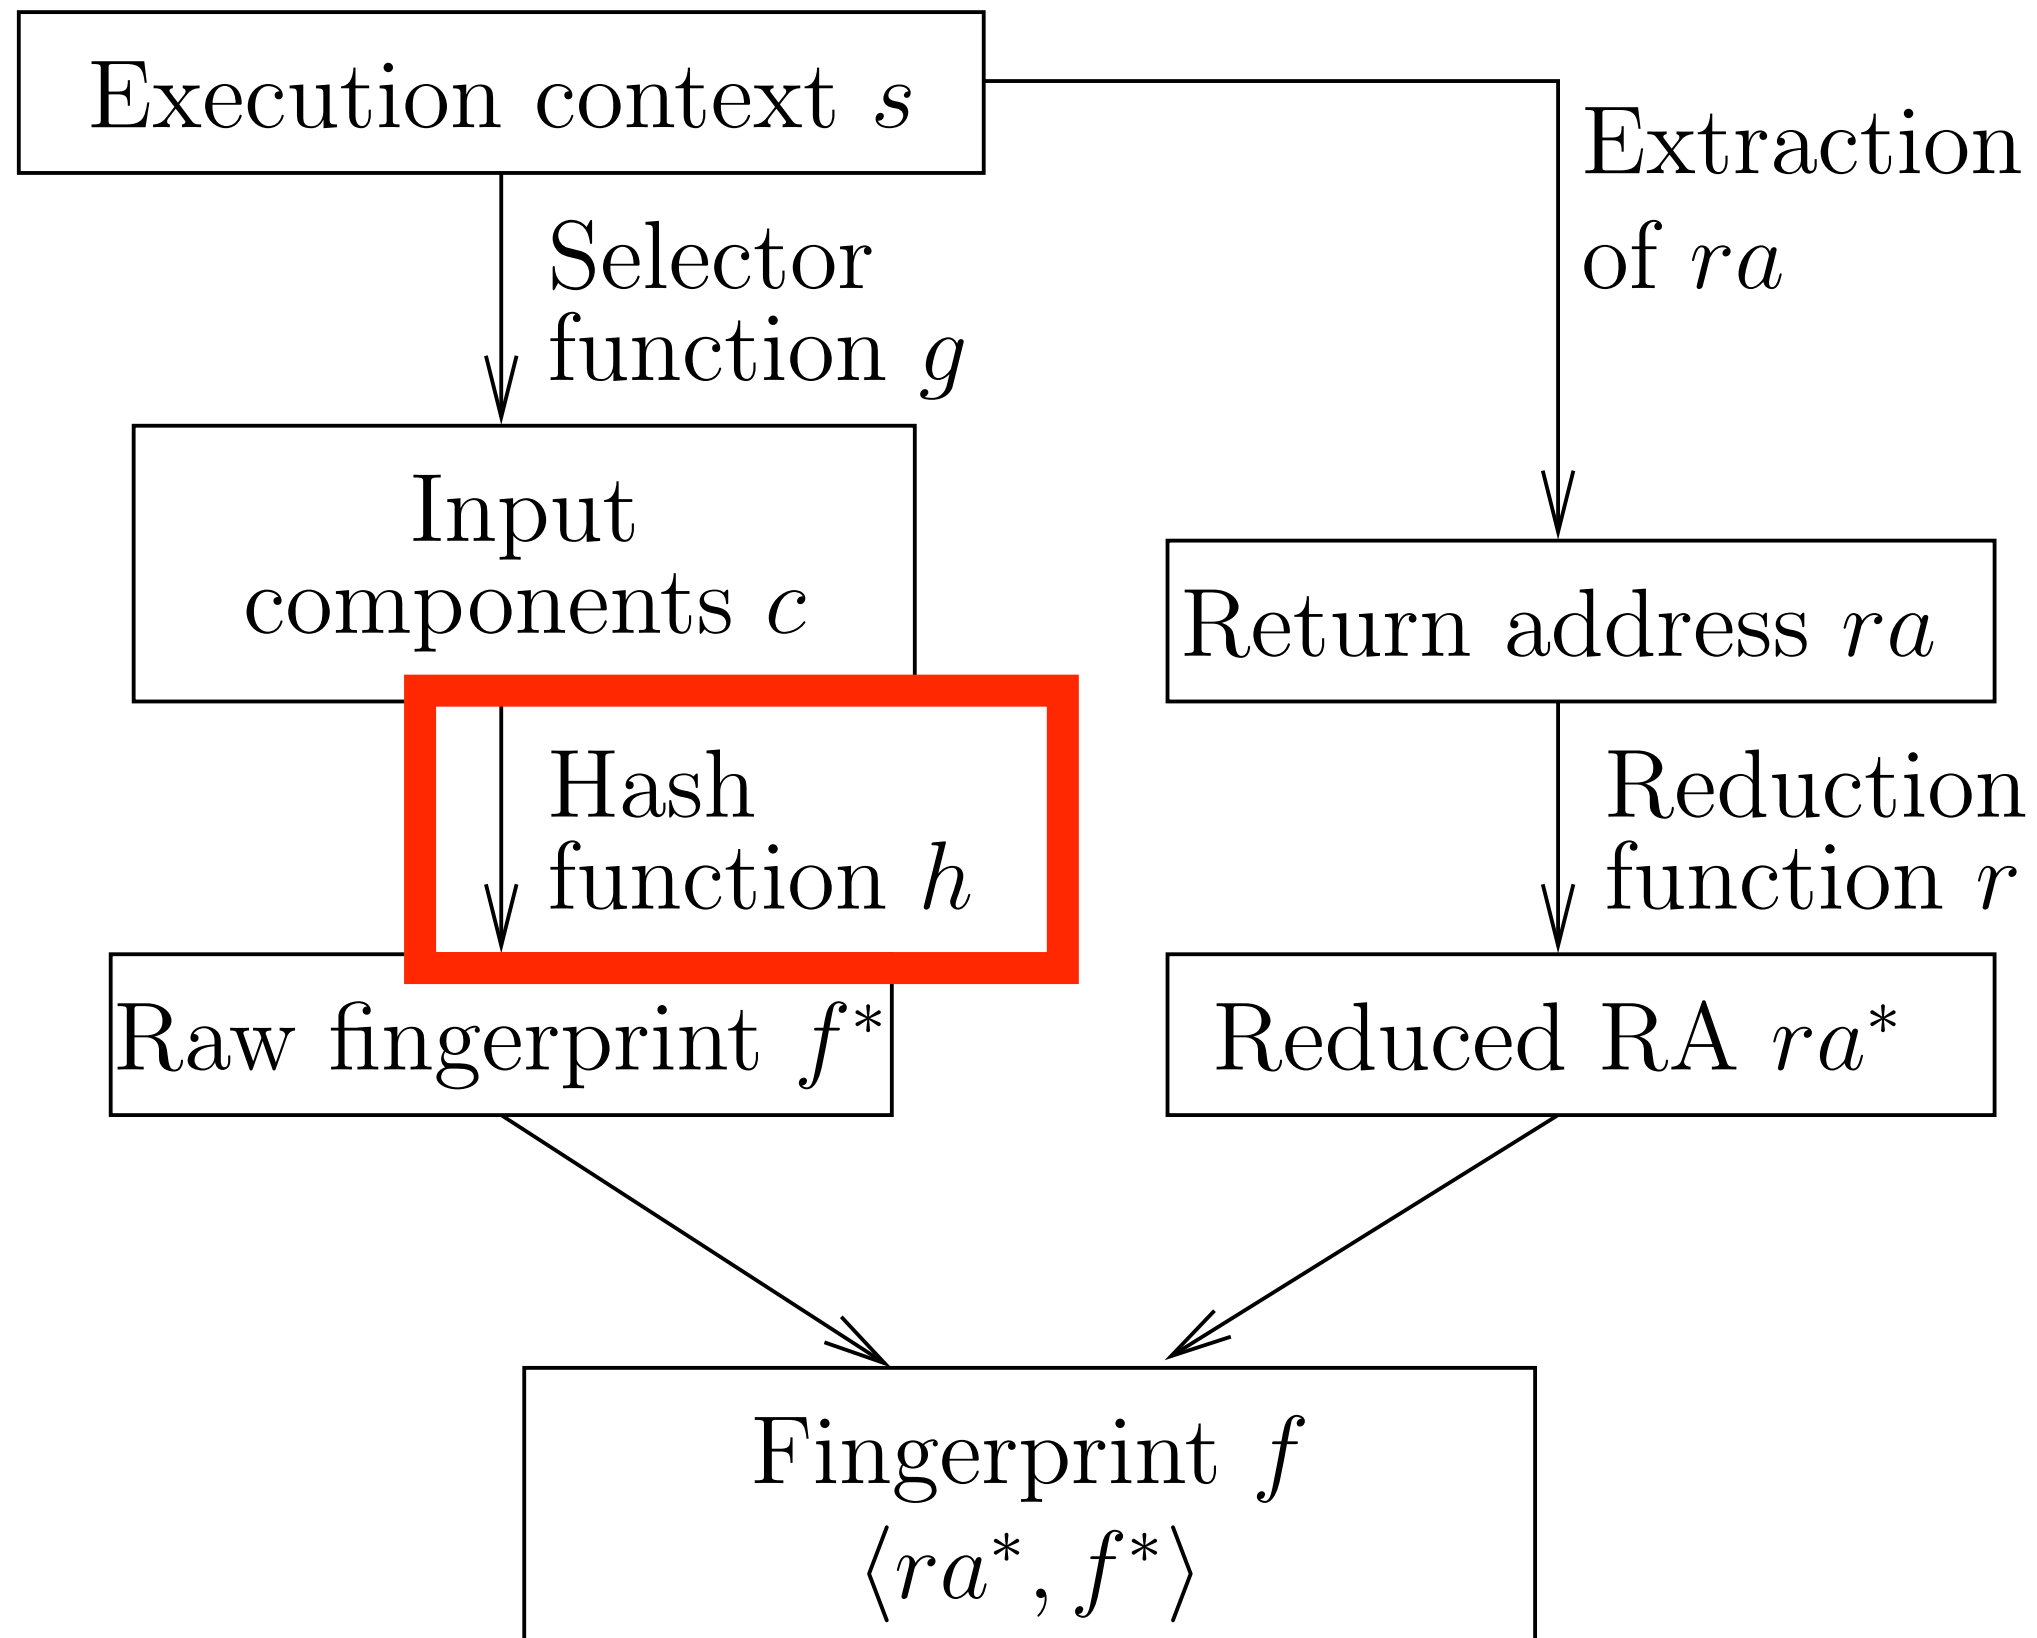
## Yes, quite well.



# Does it work?









# All Perform Equally Badly

19




# Observation

## Return addresses tend to cluster



# Why?

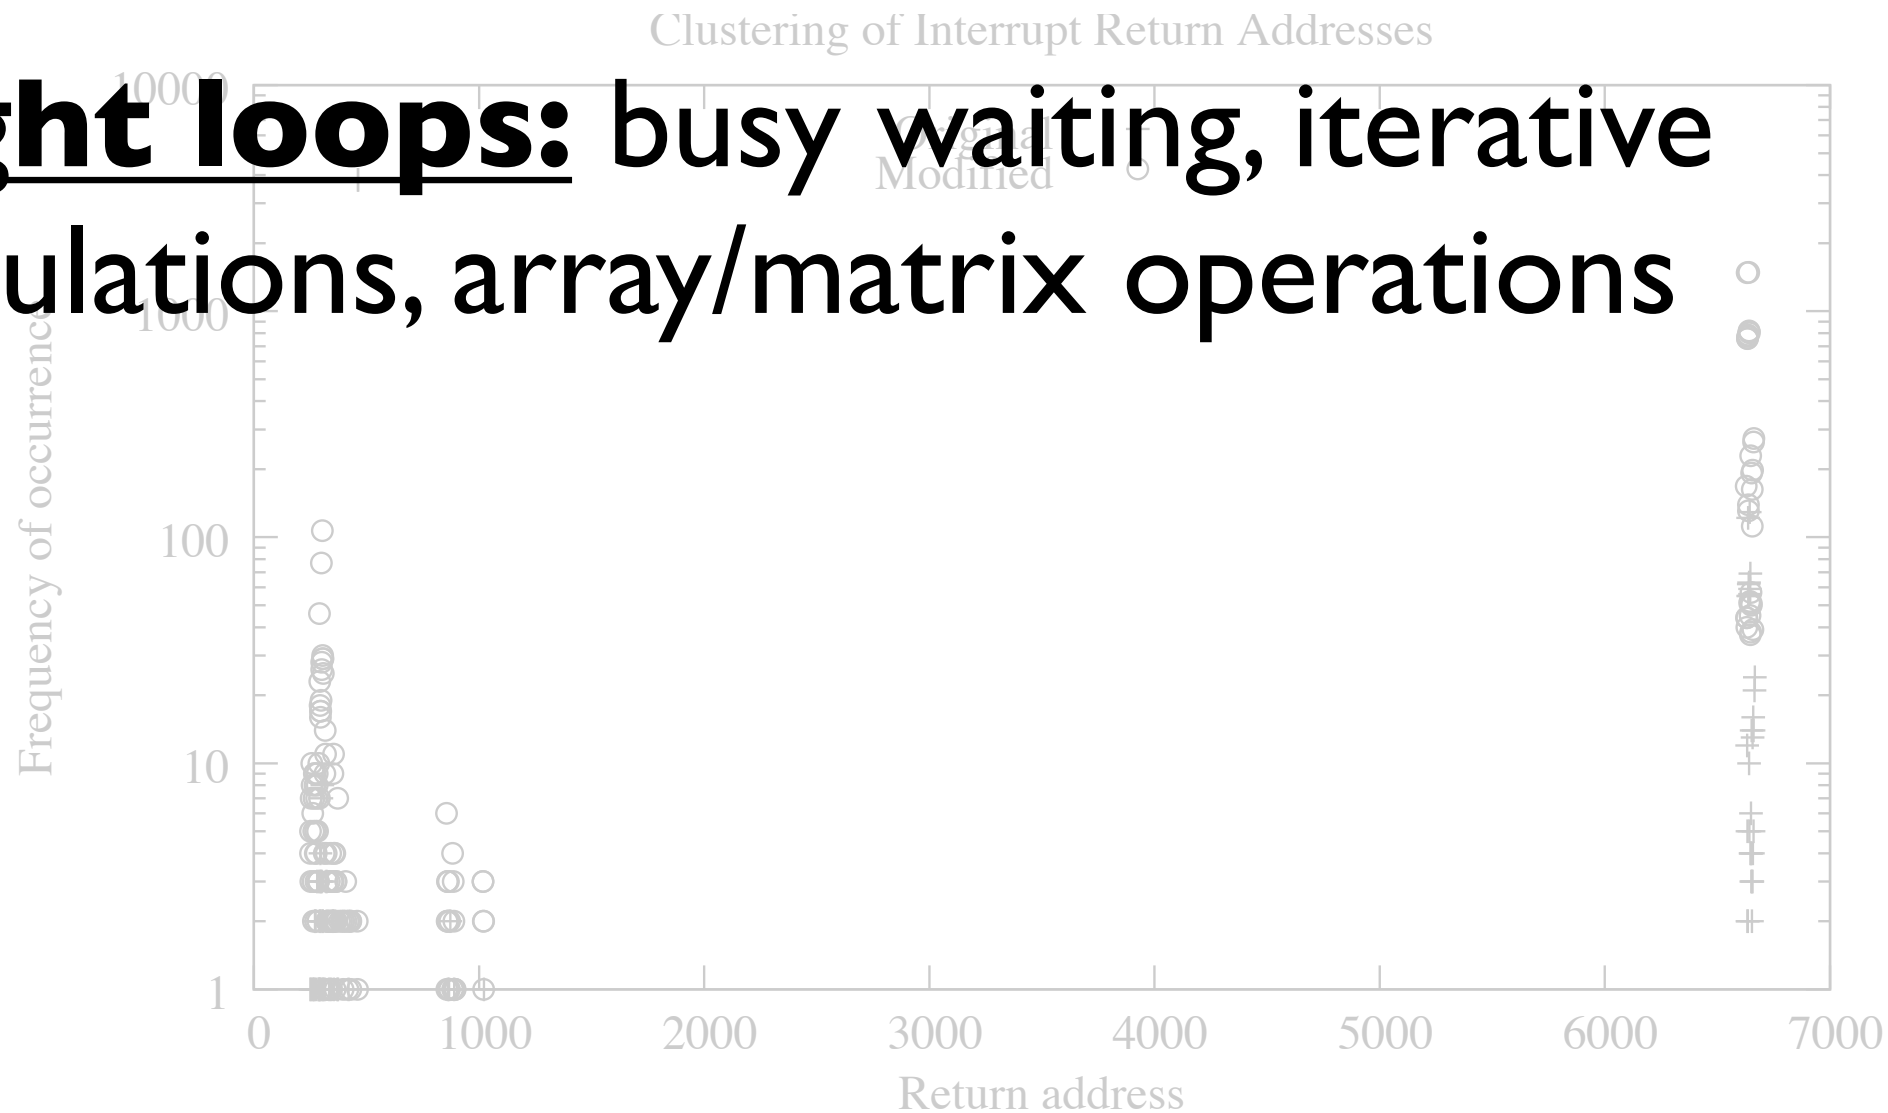
## Return addresses tend to cluster



# Why?

## Return addresses tend to cluster

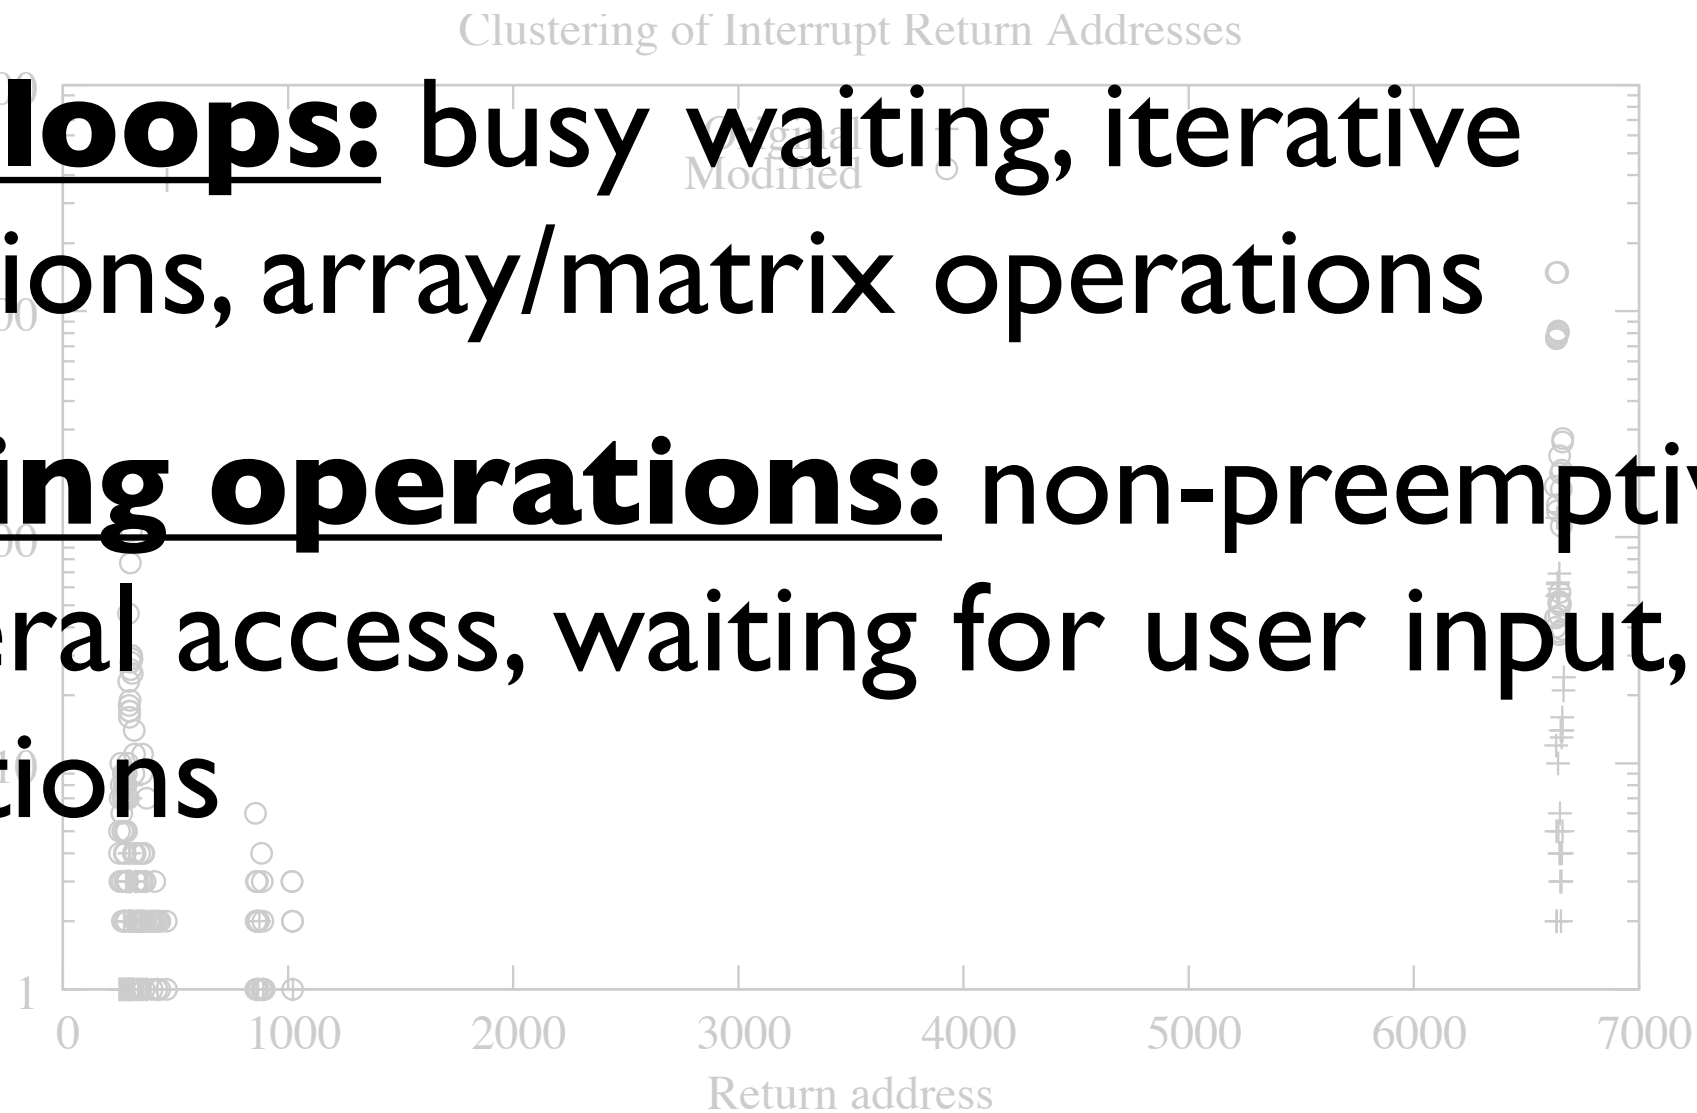
- **Tight loops:** busy waiting, iterative calculations, array/matrix operations



# Why?

## Return addresses tend to cluster

- **Tight loops:** busy waiting, iterative calculations, array/matrix operations
- **Blocking operations:** non-preemptive peripheral access, waiting for user input, costly instructions

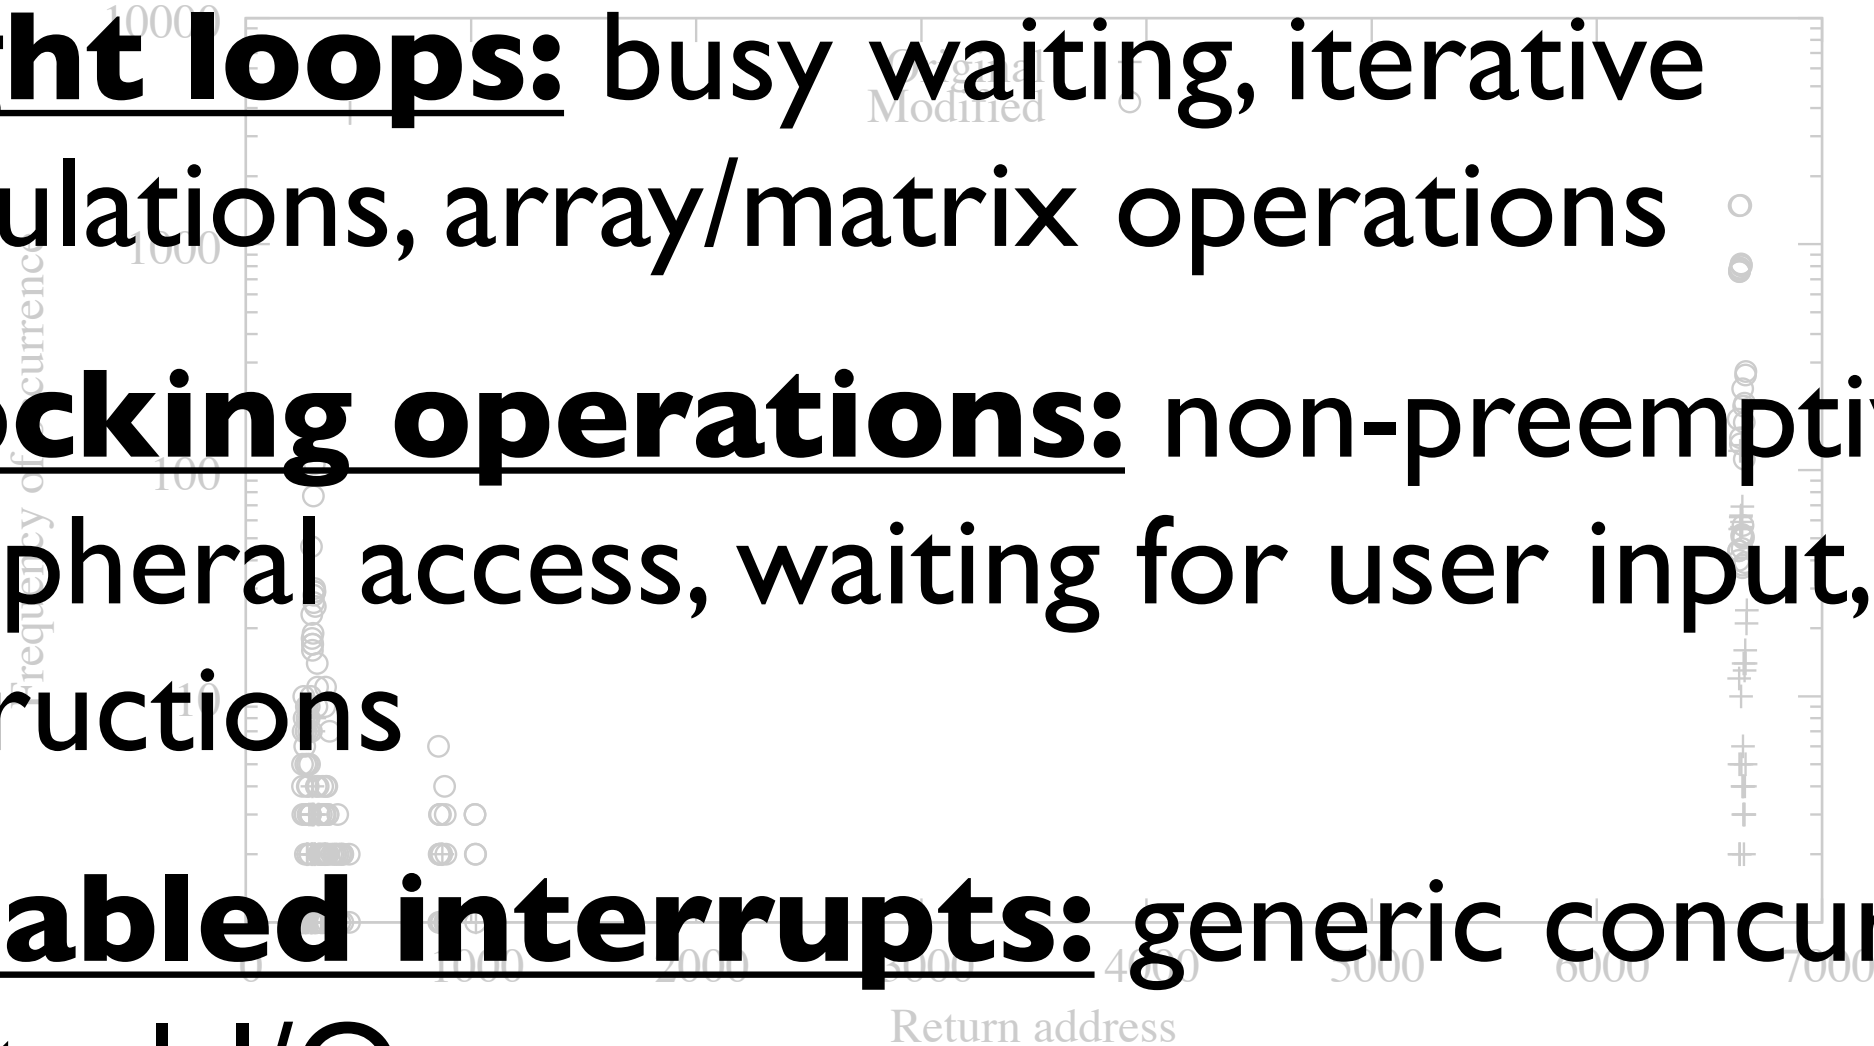


# Why?

## Return addresses tend to cluster

Clustering of Interrupt Return Addresses

- **Tight loops:** busy waiting, iterative calculations, array/matrix operations
- **Blocking operations:** non-preemptive peripheral access, waiting for user input, costly instructions
- **Disabled interrupts:** generic concurrency control, I/O access



# Why?

## Return addresses tend to cluster

Clustering of Interrupt Return Addresses

- **Tight loops:** busy waiting, iterative calculations, array/matrix operations
- **Blocking operations:** non-preemptive peripheral access, waiting for user input, costly instructions
- **Disabled interrupts:** generic concurrency control, I/O access

Context specific hashing => RR hashing



# Idea 2: Round-Robin Hashing<sup>22</sup>

# Idea 2: Round-Robin Hashing<sup>22</sup>

- Use **different hash functions** to minimize collisions

# Idea 2: Round-Robin Hashing<sup>22</sup>


- Use **different hash functions** to minimize collisions
- Insert **markers** in the code for different hash functions

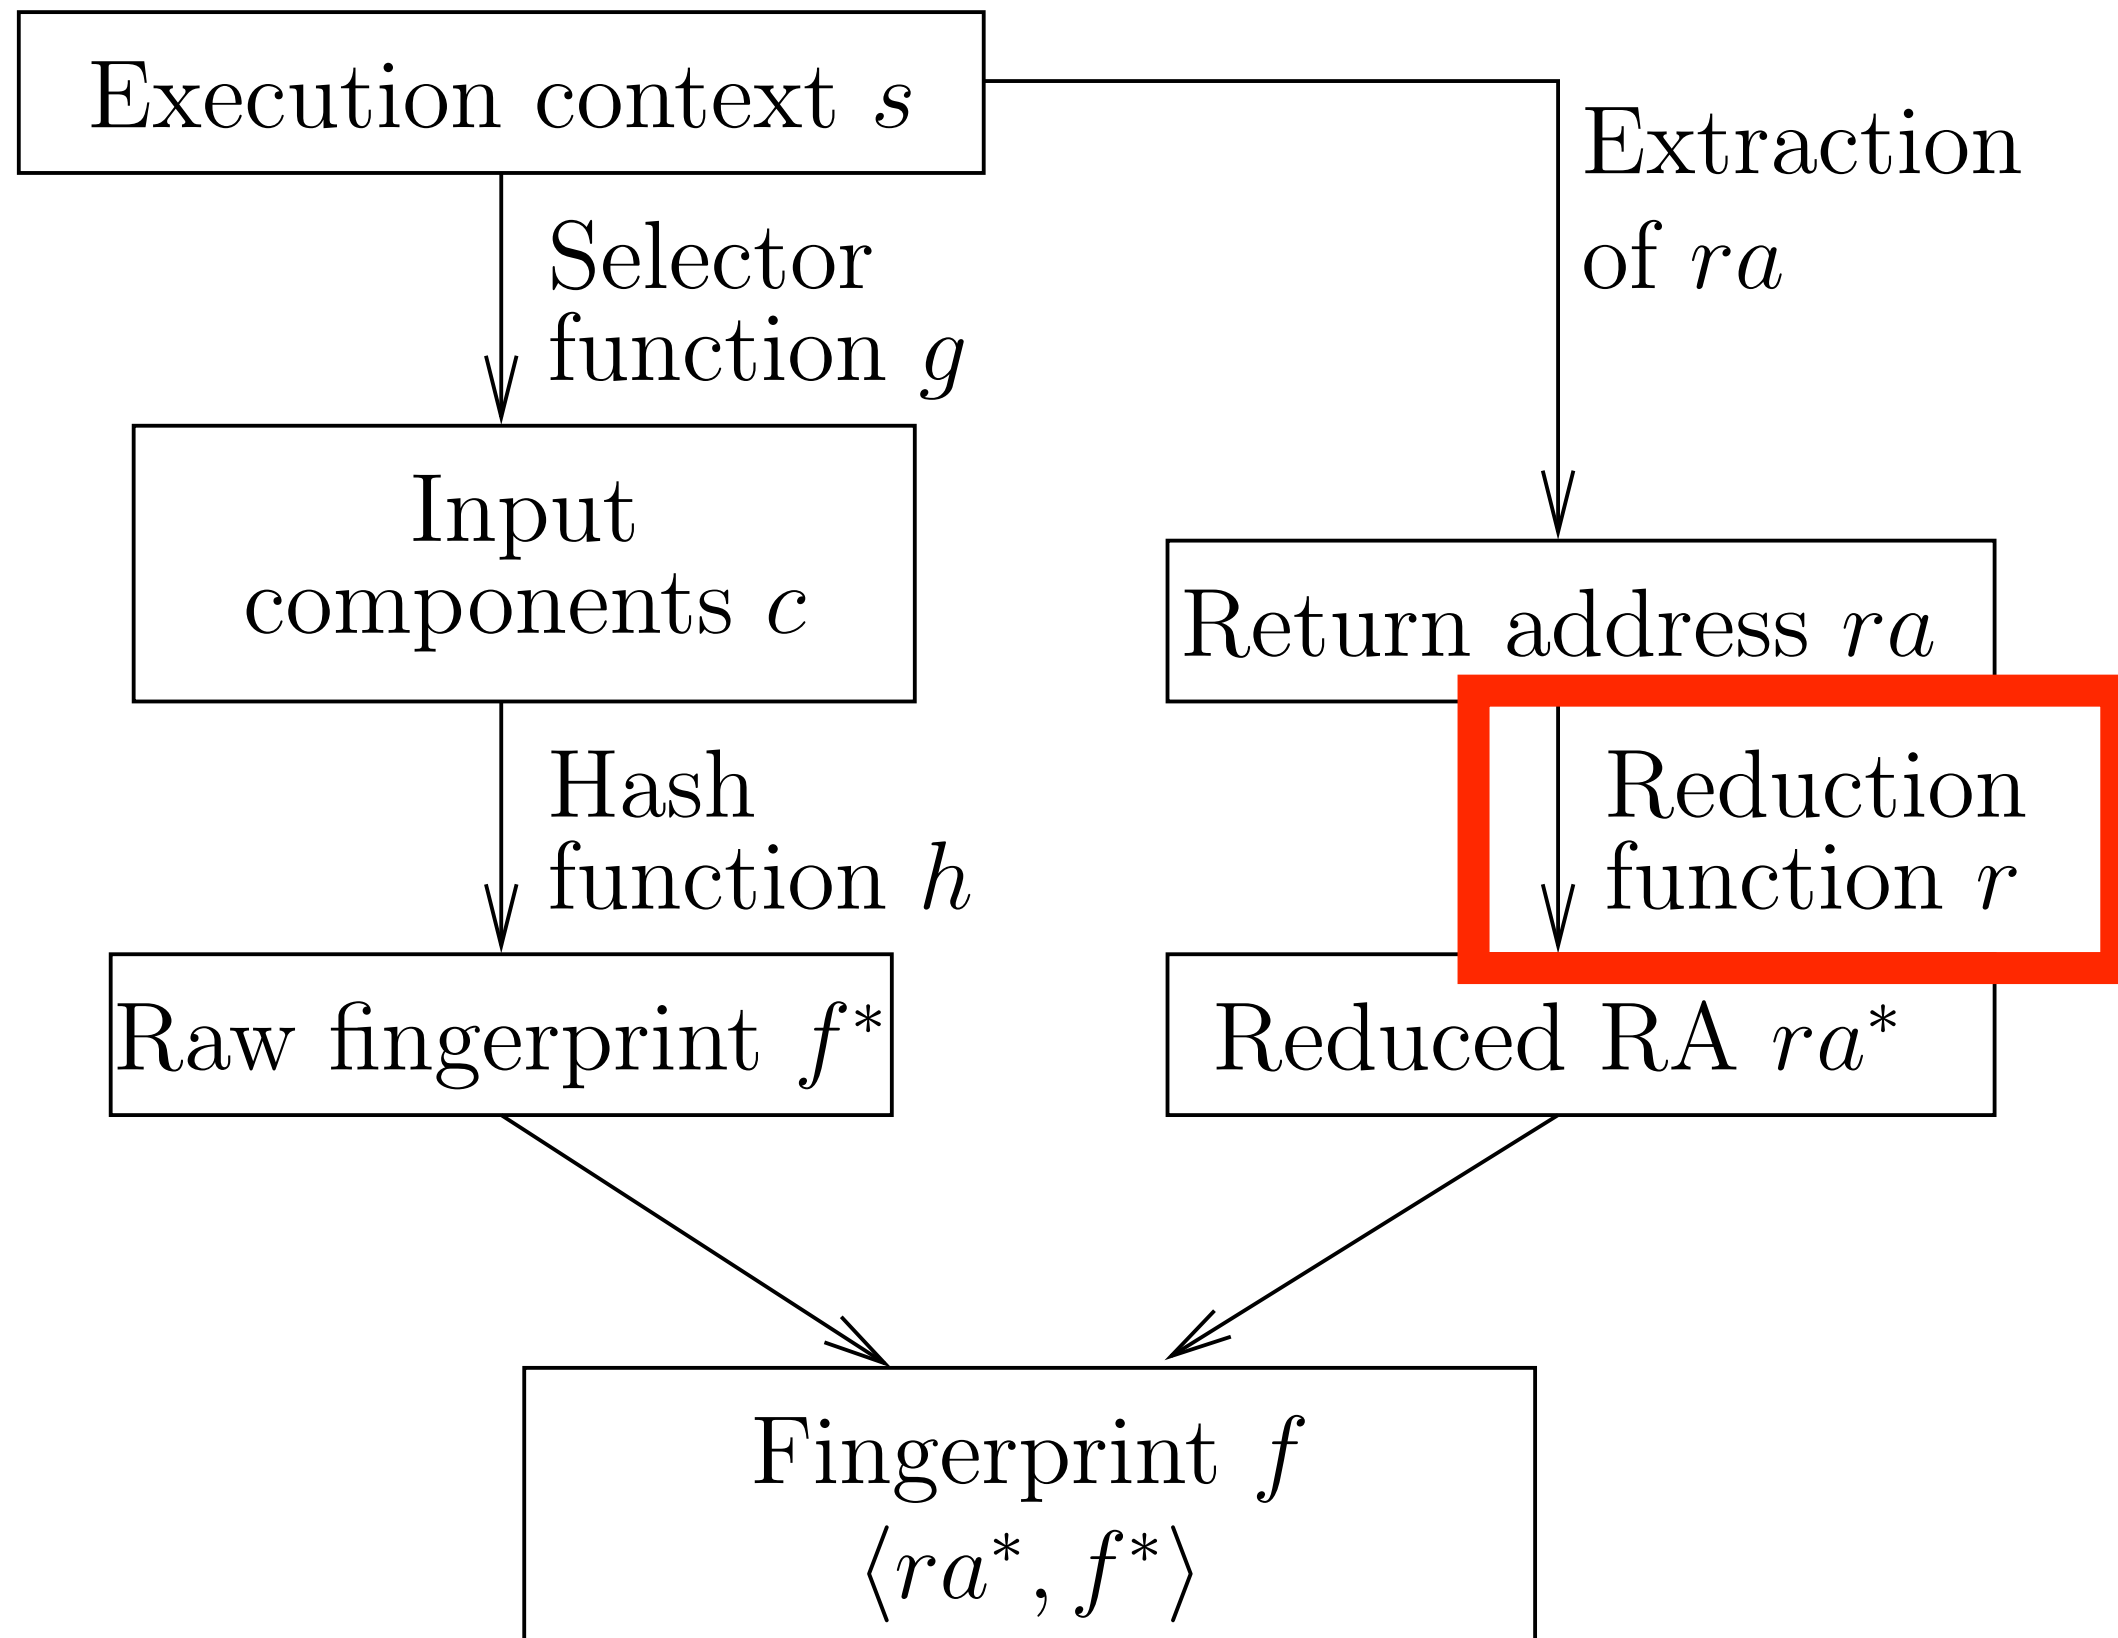
# Idea 2: Round-Robin Hashing <sup>22</sup>

- Use **different hash functions** to minimize collisions
- Insert **markers** in the code for different hash functions
- In the replay use the correct hash function for according to the marker

# Does it work?


## RR hashing works surprisingly well





# How best to use the FP bits?


- Give all to the hash value? *(minimize hash collisions)*
- Give all to the RA? *(minimize false RA positives)*



# Engineering the Fingerprint

26

Trade-off Between  $ra^*$  and  $f^*$  Width





# Conclusions

- Debugging embedded systems is painful  
=> Replay debugging can help [ST08]
- Frequency-based selection function ✓
- Round robin hashing ✓
- Tradeoff engineering for RA vs hash ✓
- Future work: Lots, see discussion section

# Questions?

Work has been sponsored by:

Graduate Students' Exchange Program (GSEP) funded by Foreign Affairs and International Trade Canada (DFAIT),  
Ontario Research Fund Research Excellence (ORF-RE) &  
CIMIT under U.S.Army Medical Research Acquisition Activity Cooperative Agreement W81XWH-07-2-0011