

# $\mu$ RAI: Securing Embedded Systems with Return Address Integrity

Naif Saleh Almakhdhub<sup>1,4</sup>


Abraham A. Clements<sup>3</sup>

Saurabh Bagchi<sup>1</sup>

Mathias Payer<sup>2</sup>



# Current State of Security



**Target:**  
Embedded and IoT devices  
Running Microcontroller  
Systems (MCUS)



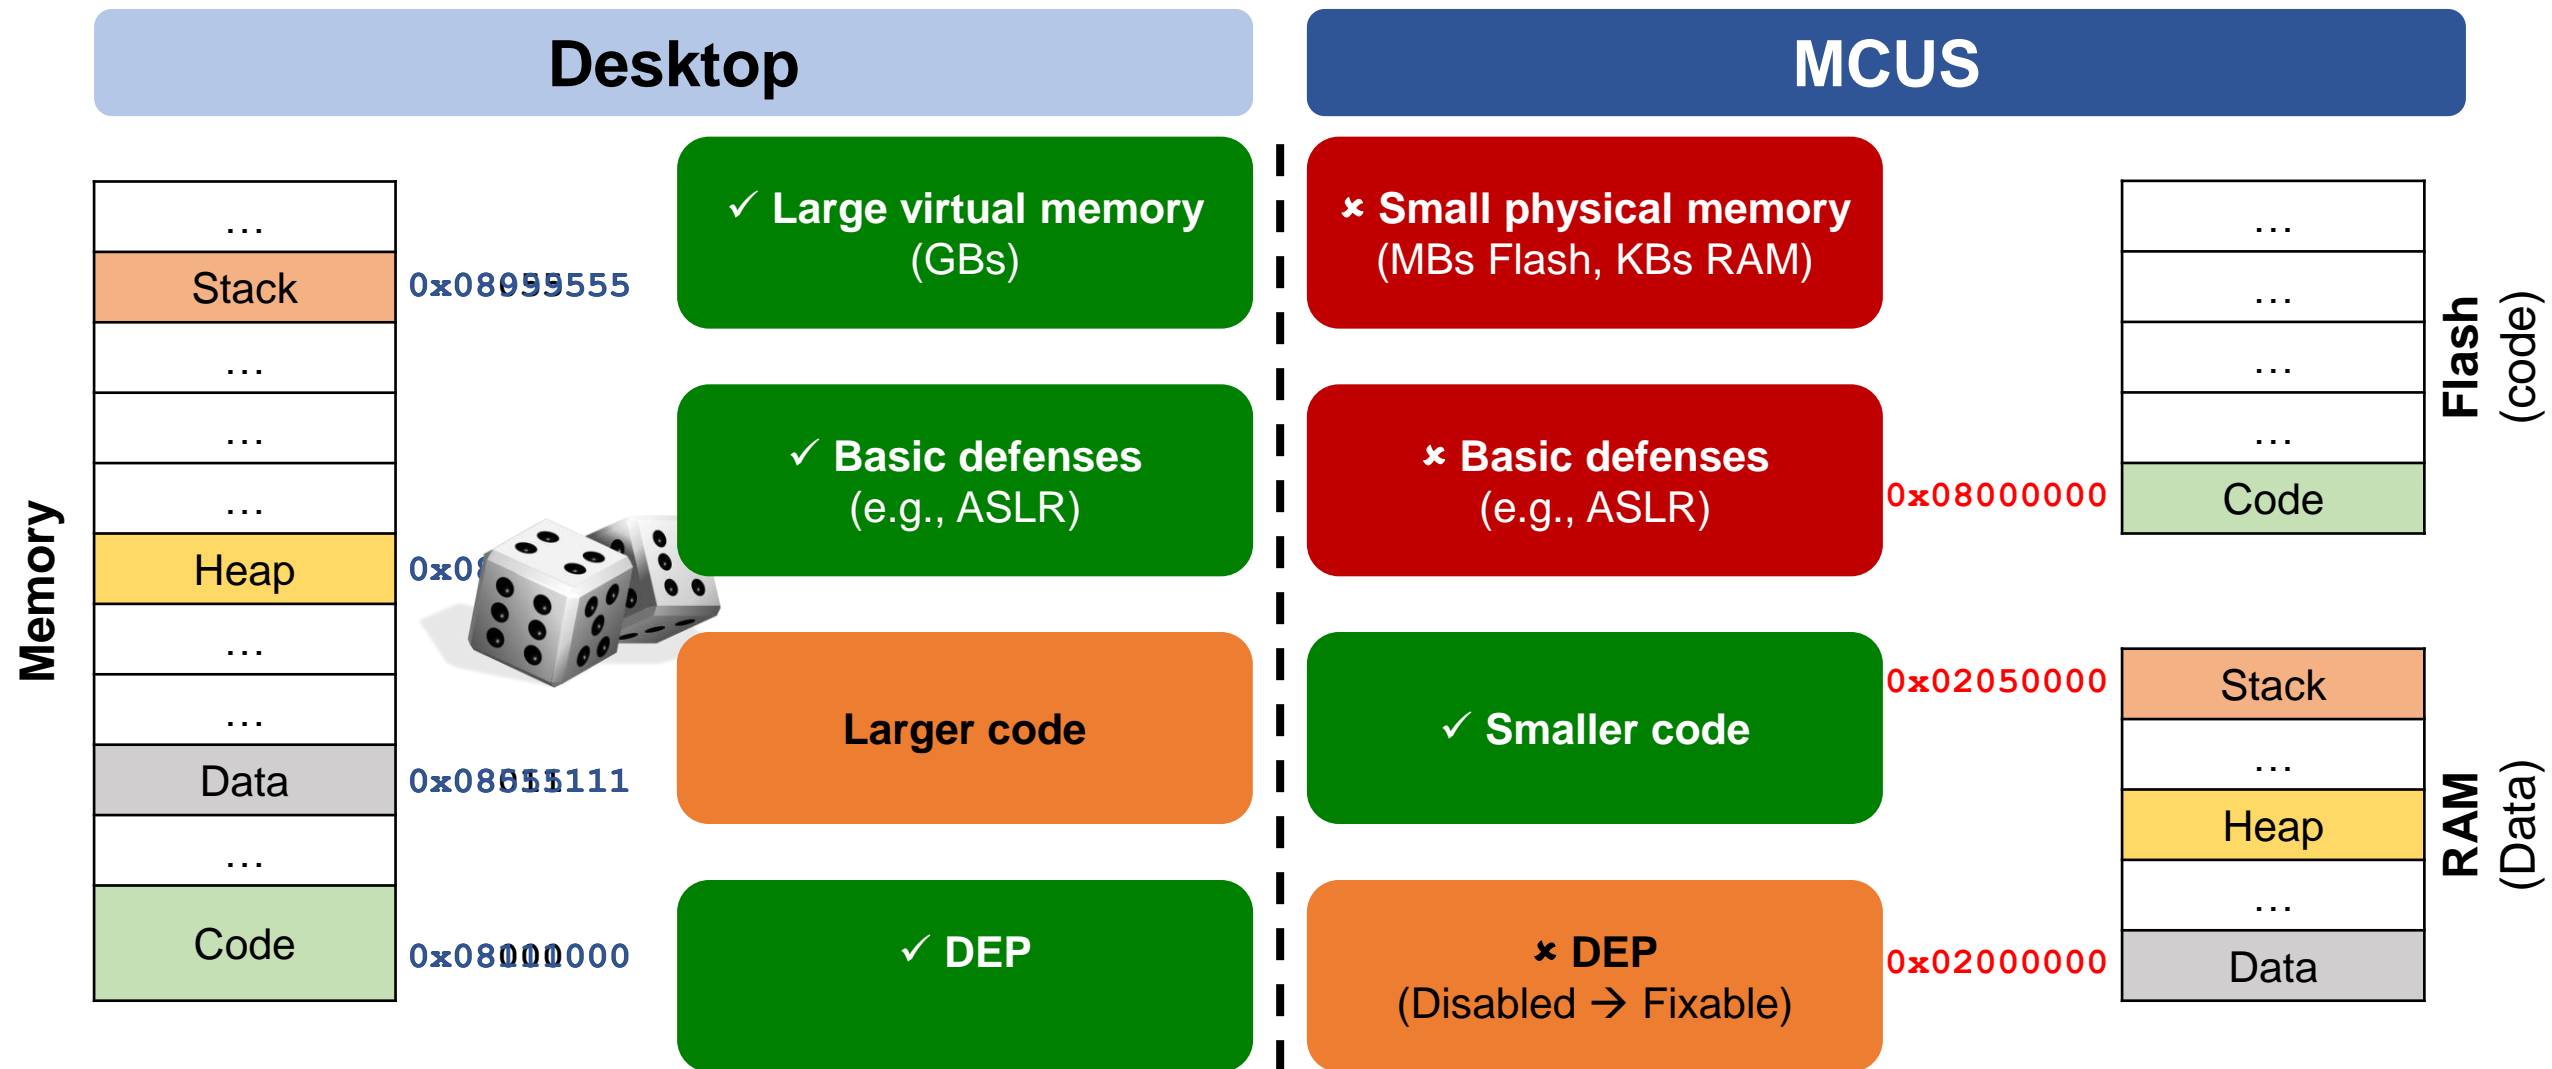
**Attack:**  
Control-flow Hijacking

[1] <https://www.wired.com/story/broadpwn-wi-fi-vulnerability-ios-android/>

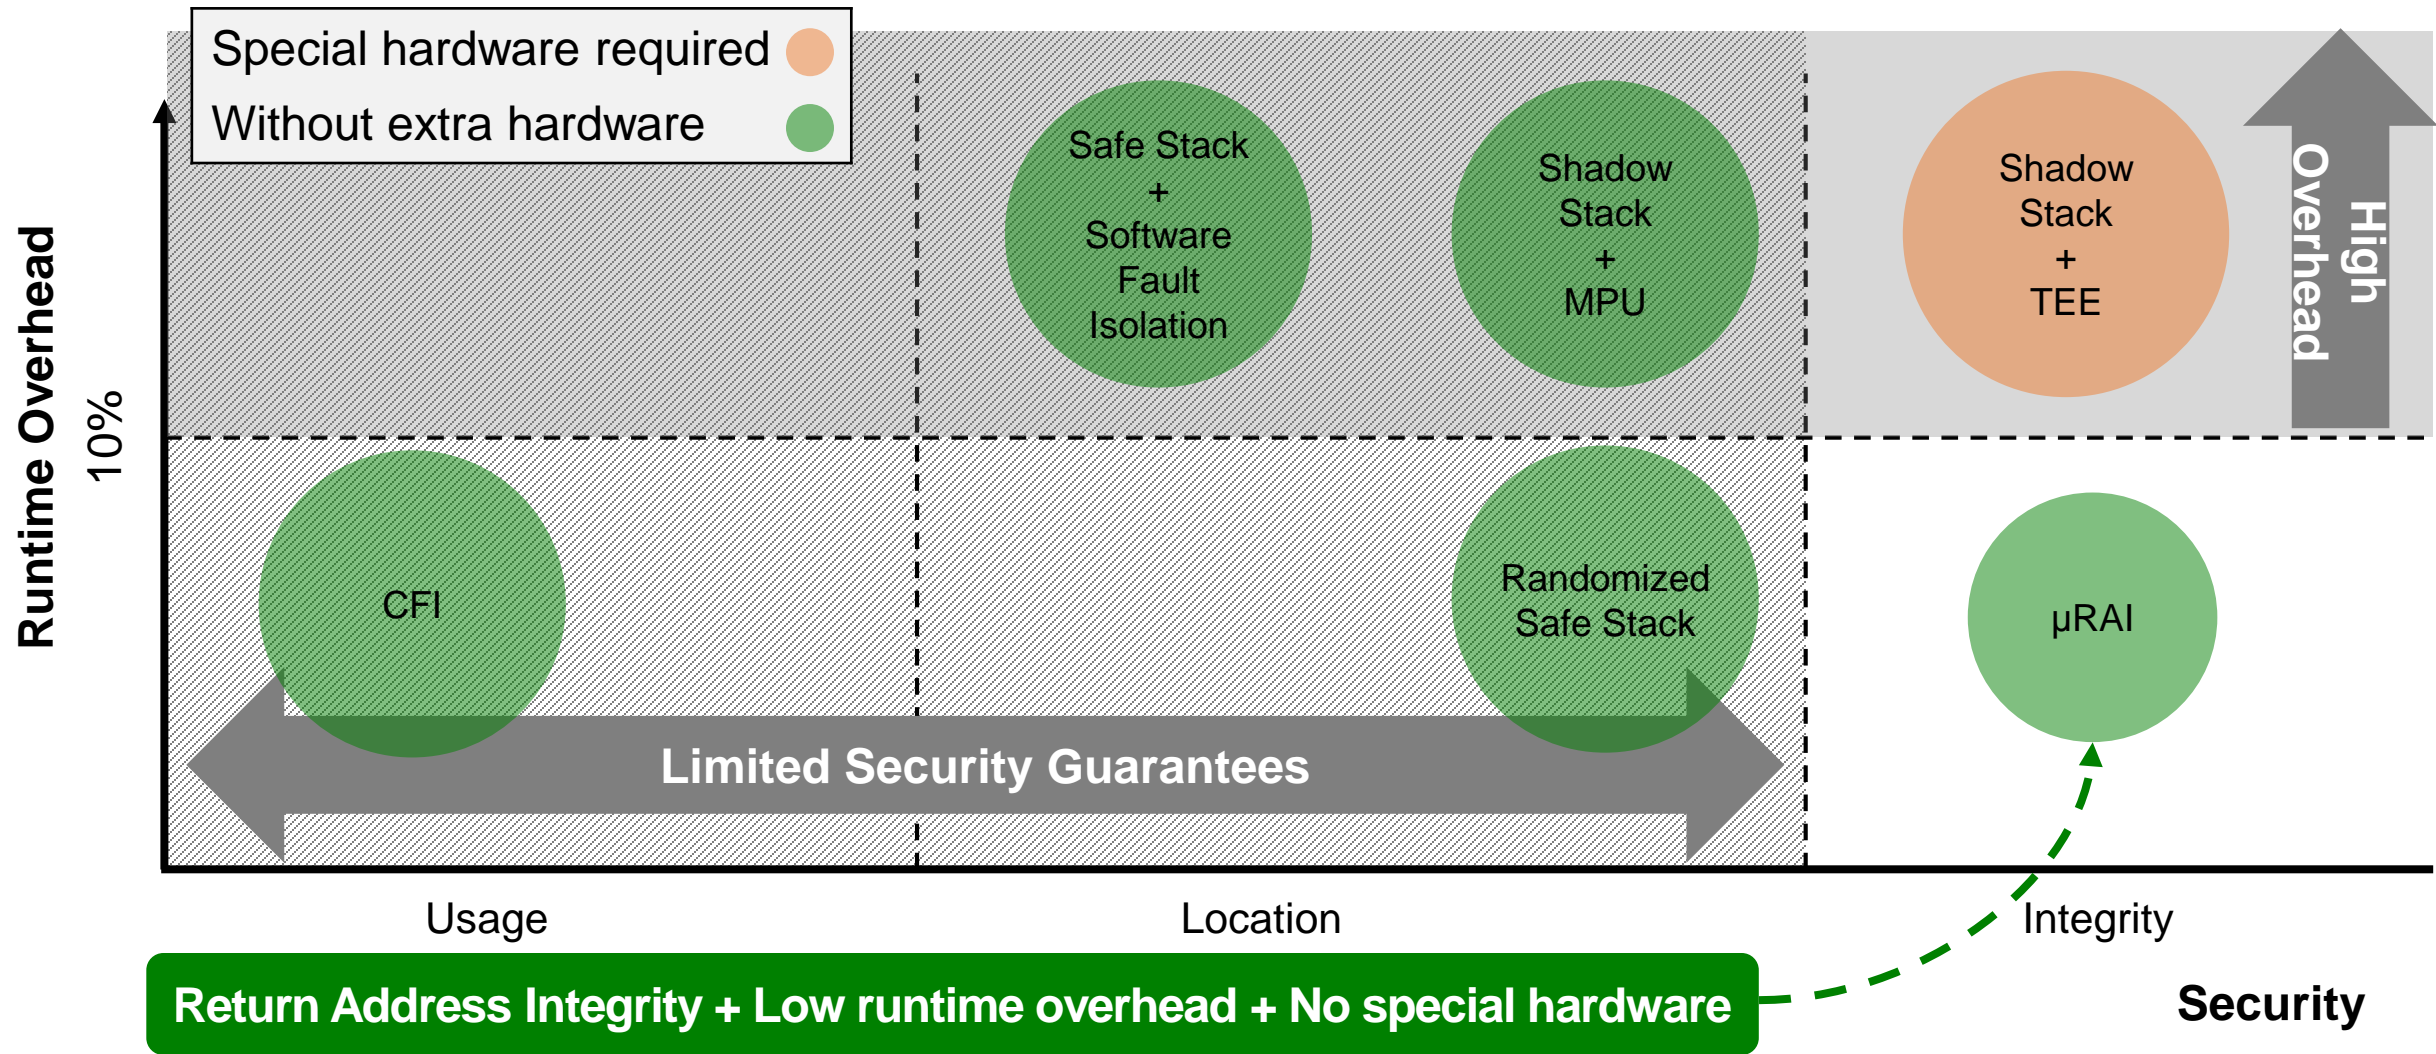
[2] <https://keenlab.tencent.com/en/2020/01/02/exploiting-wifi-stack-on-tesla-model-s/>

[3] <https://www.securityweek.com/rise-ics-malware-how-industrial-security-threats-are-becoming-more-surgical>

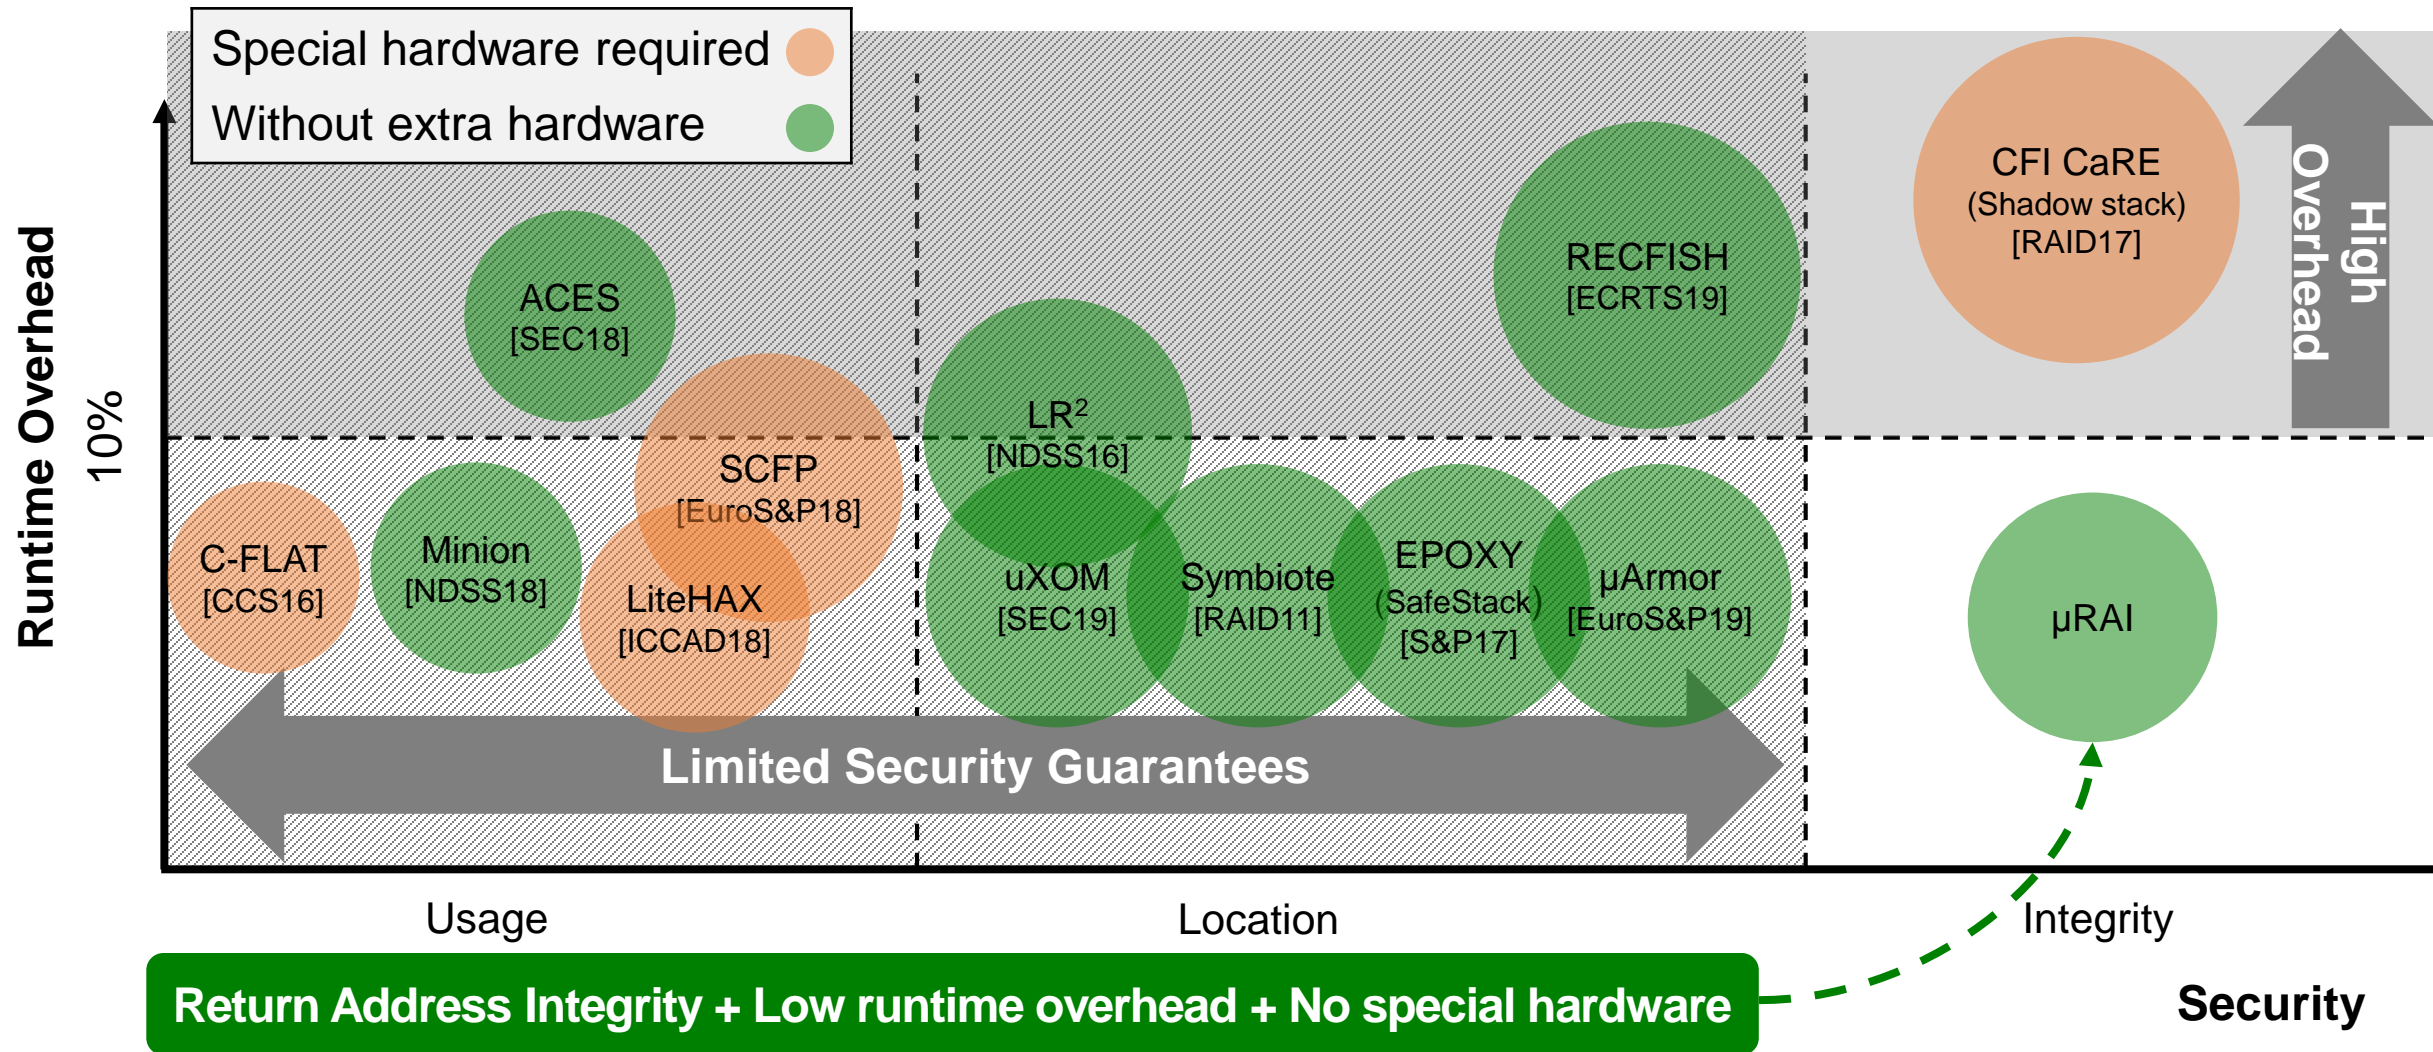
# MCUS Challenges



# MCUS Defenses for Return Addresses (Conceptual)



# MCUS Defenses for Return Addresses (Related Work)



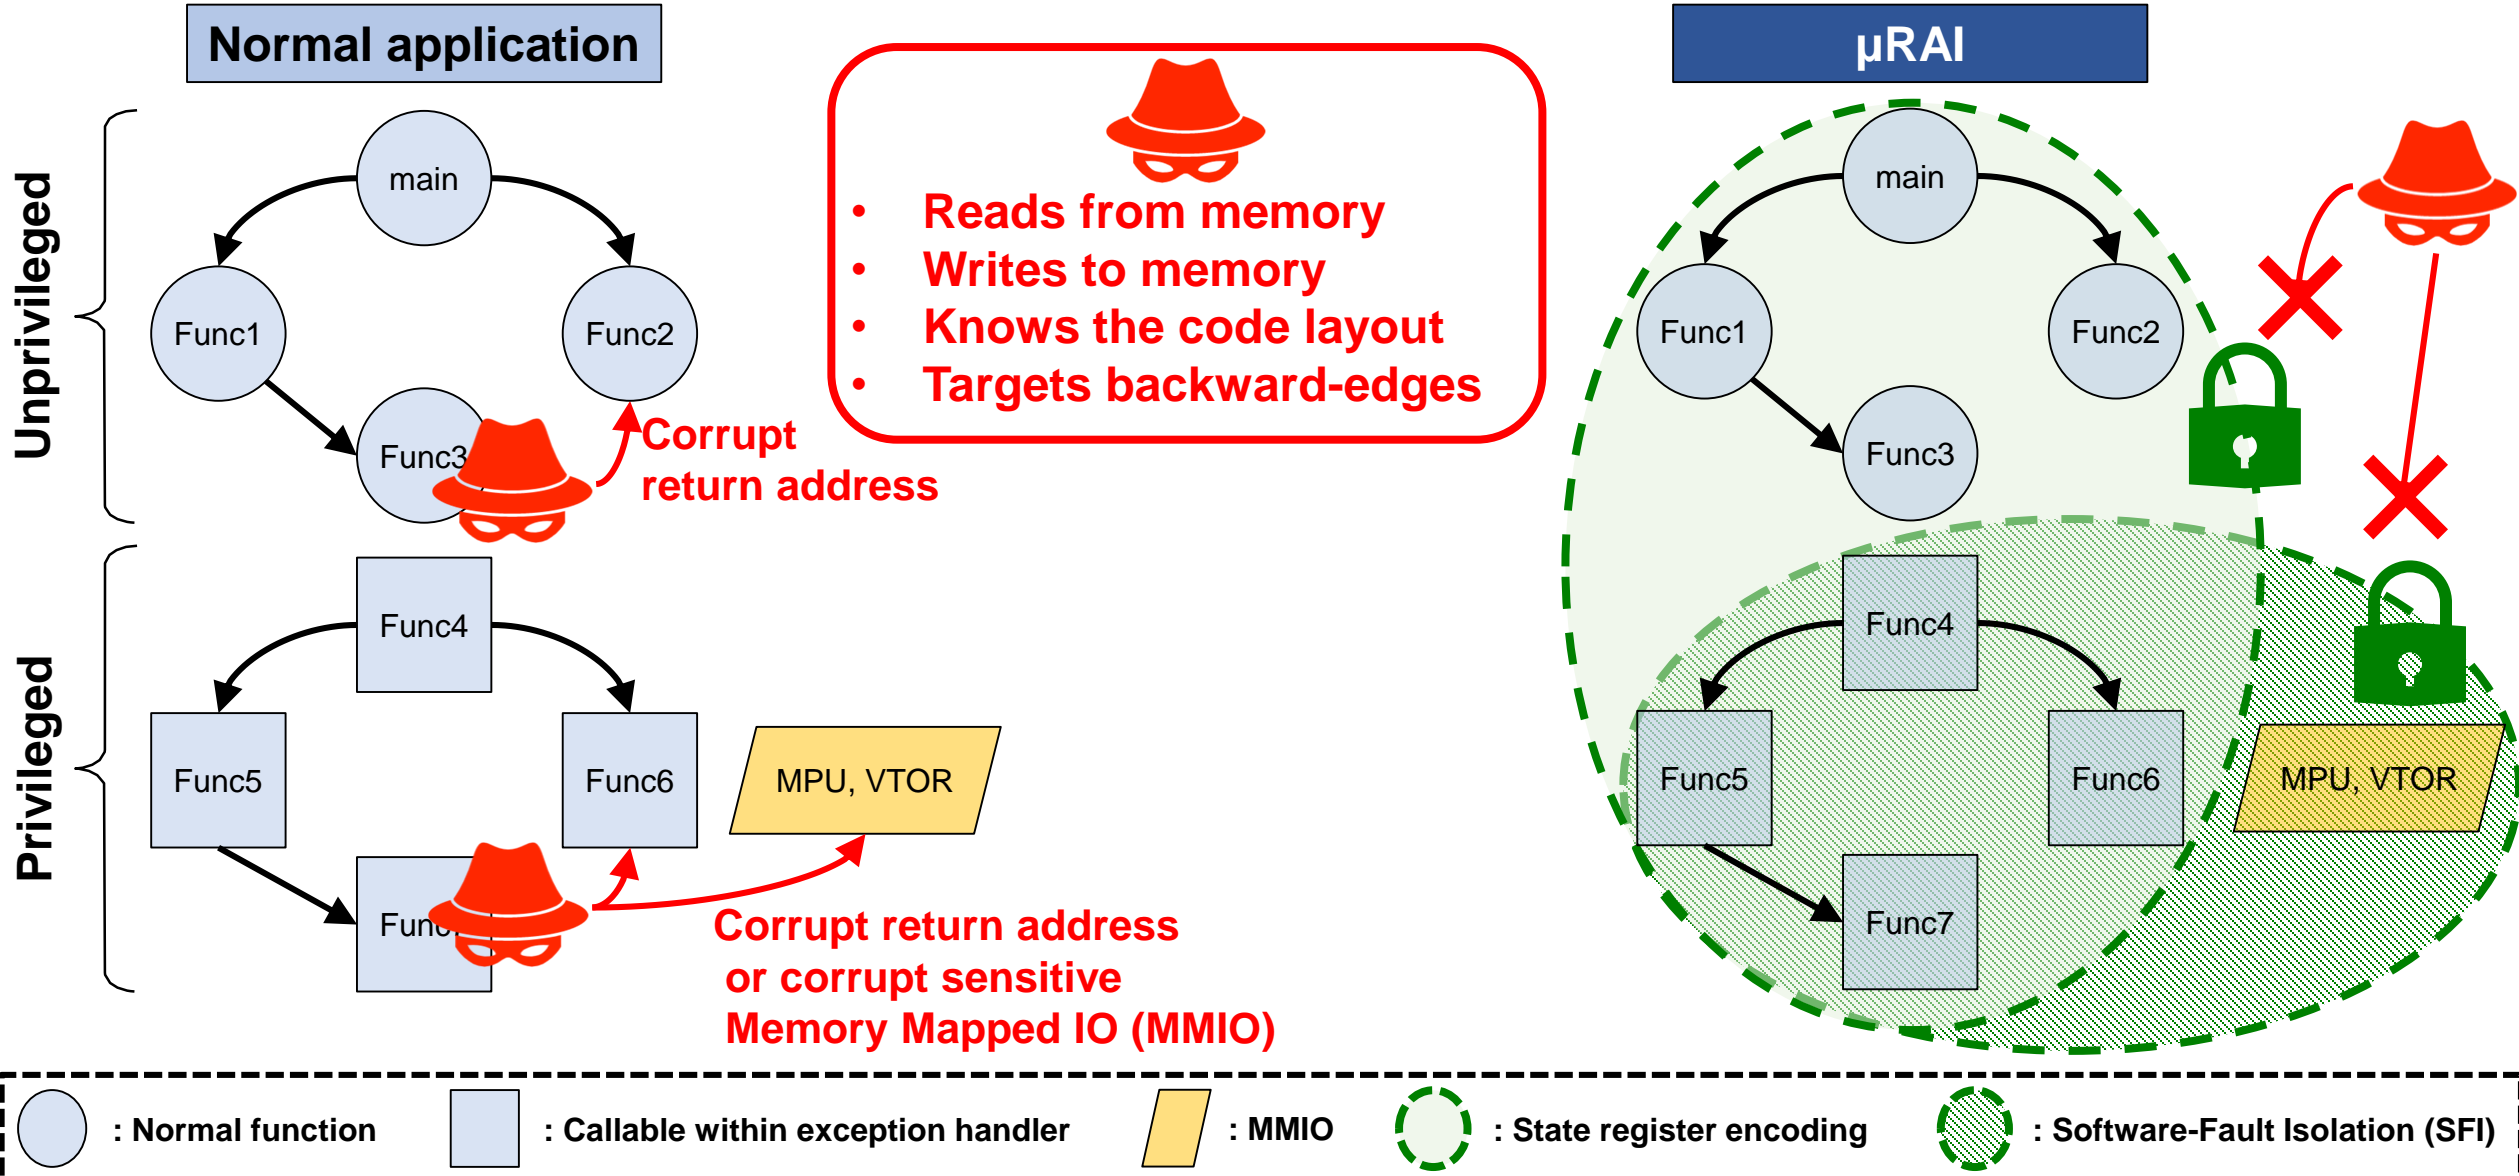
# Return Address Integrity (RAI)

- Every attack requires corrupting a return addresses by **overwriting** it

## RAI Property:

1. Ensure the return address is never writable except by an authorized instruction
2. Return addresses are never pushed to the stack or any writable memory by an adversary

# Threat Model & $\mu$ RAI Protection




# μRAI: Overview

1

Enforces the RAI property



State Register



Jump Table

Jump return\_location1

Jump return\_location2

...

Read + eXecute

2

Protects exception handlers and privileged execution



Exception handler software-fault isolation

3

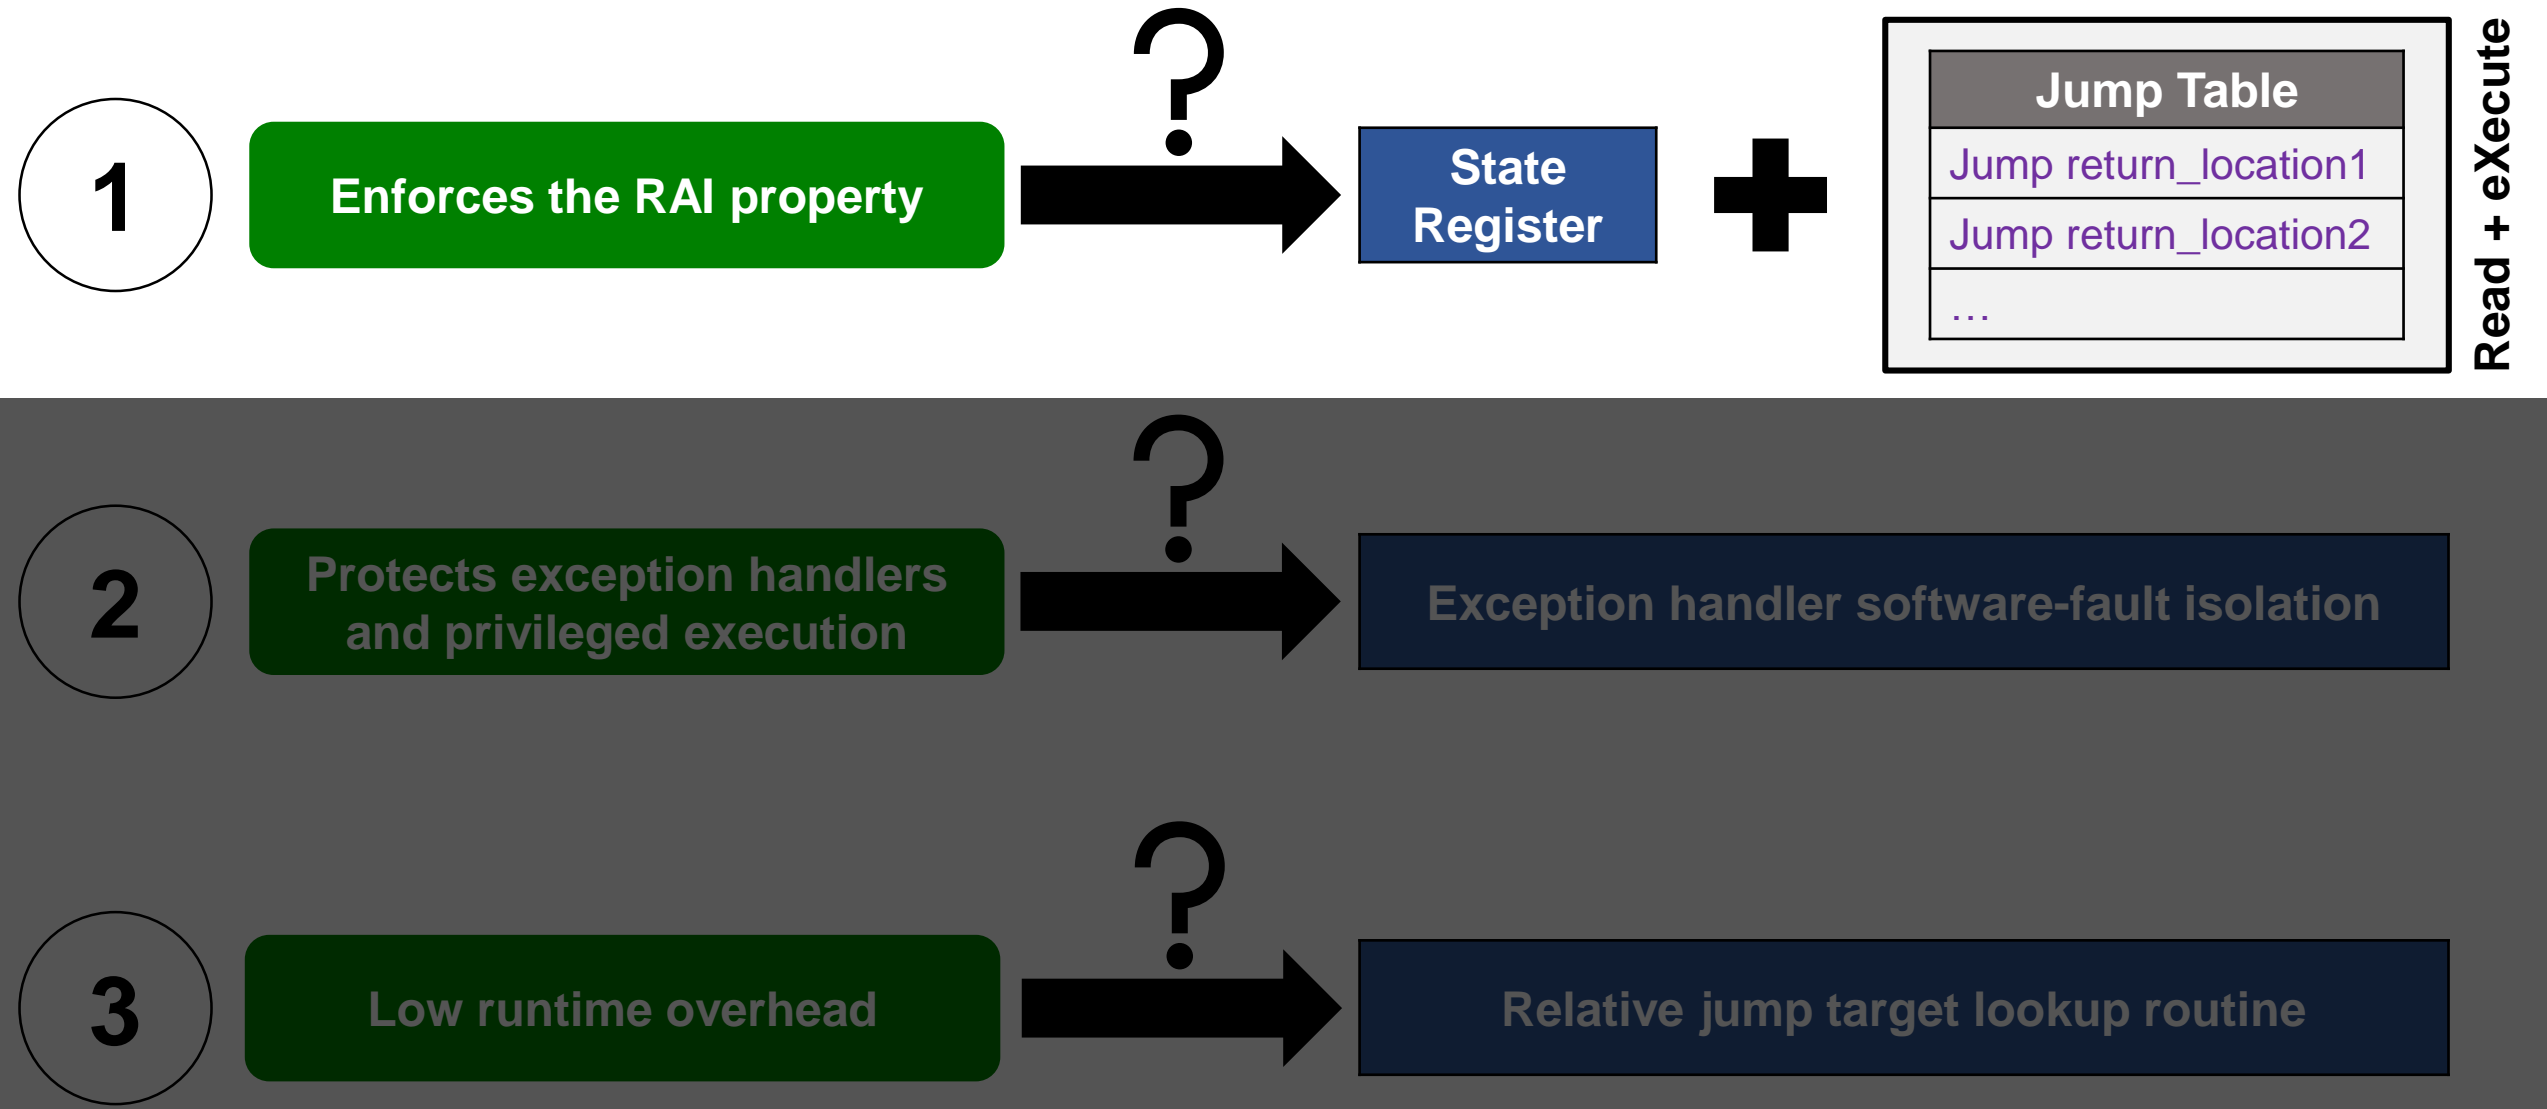
Low runtime overhead



Relative jump target lookup routine



# μRAI: Overview



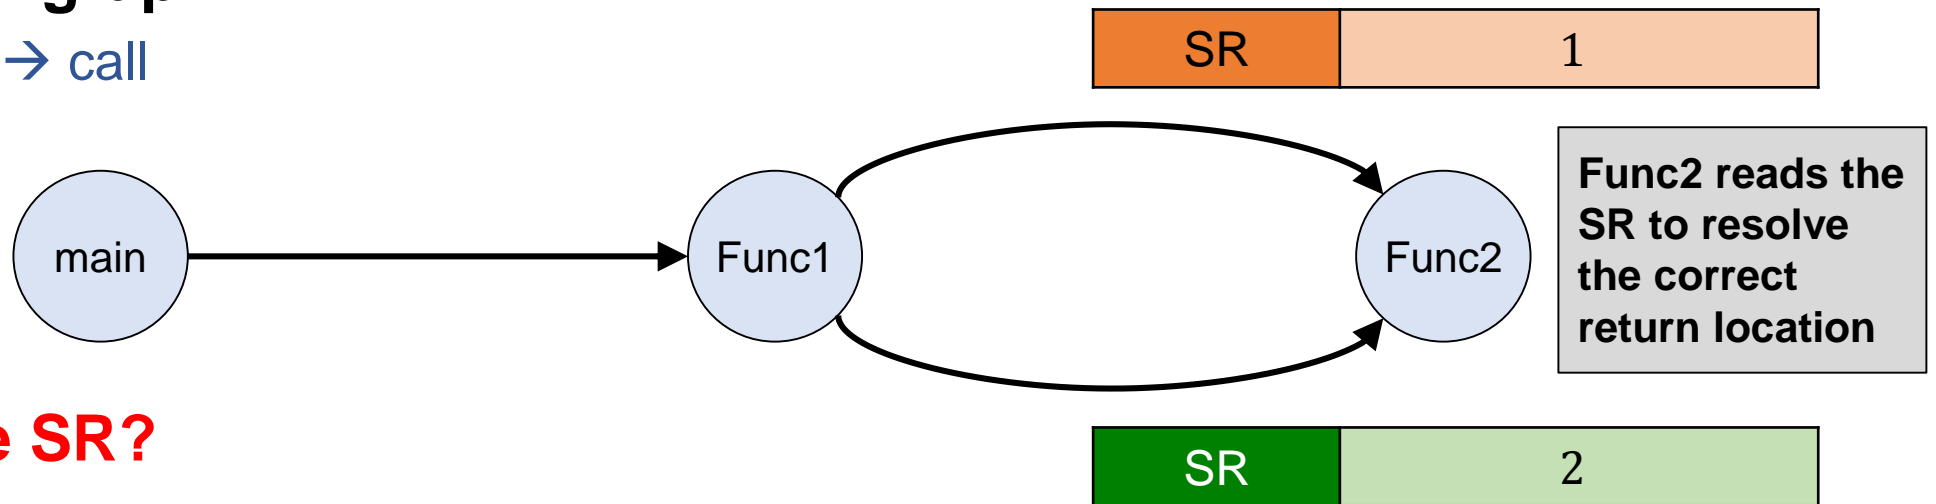
# μRAI and the State Register

- **State Register (SR):**

- Can be any general-purpose register → exclusively used by μRAI
- Never spilled → cannot be overwritten through a memory corruption
- Does not contain a return address → encoded values to resolve the return location

- **Example call graph:**

- Each edge → call

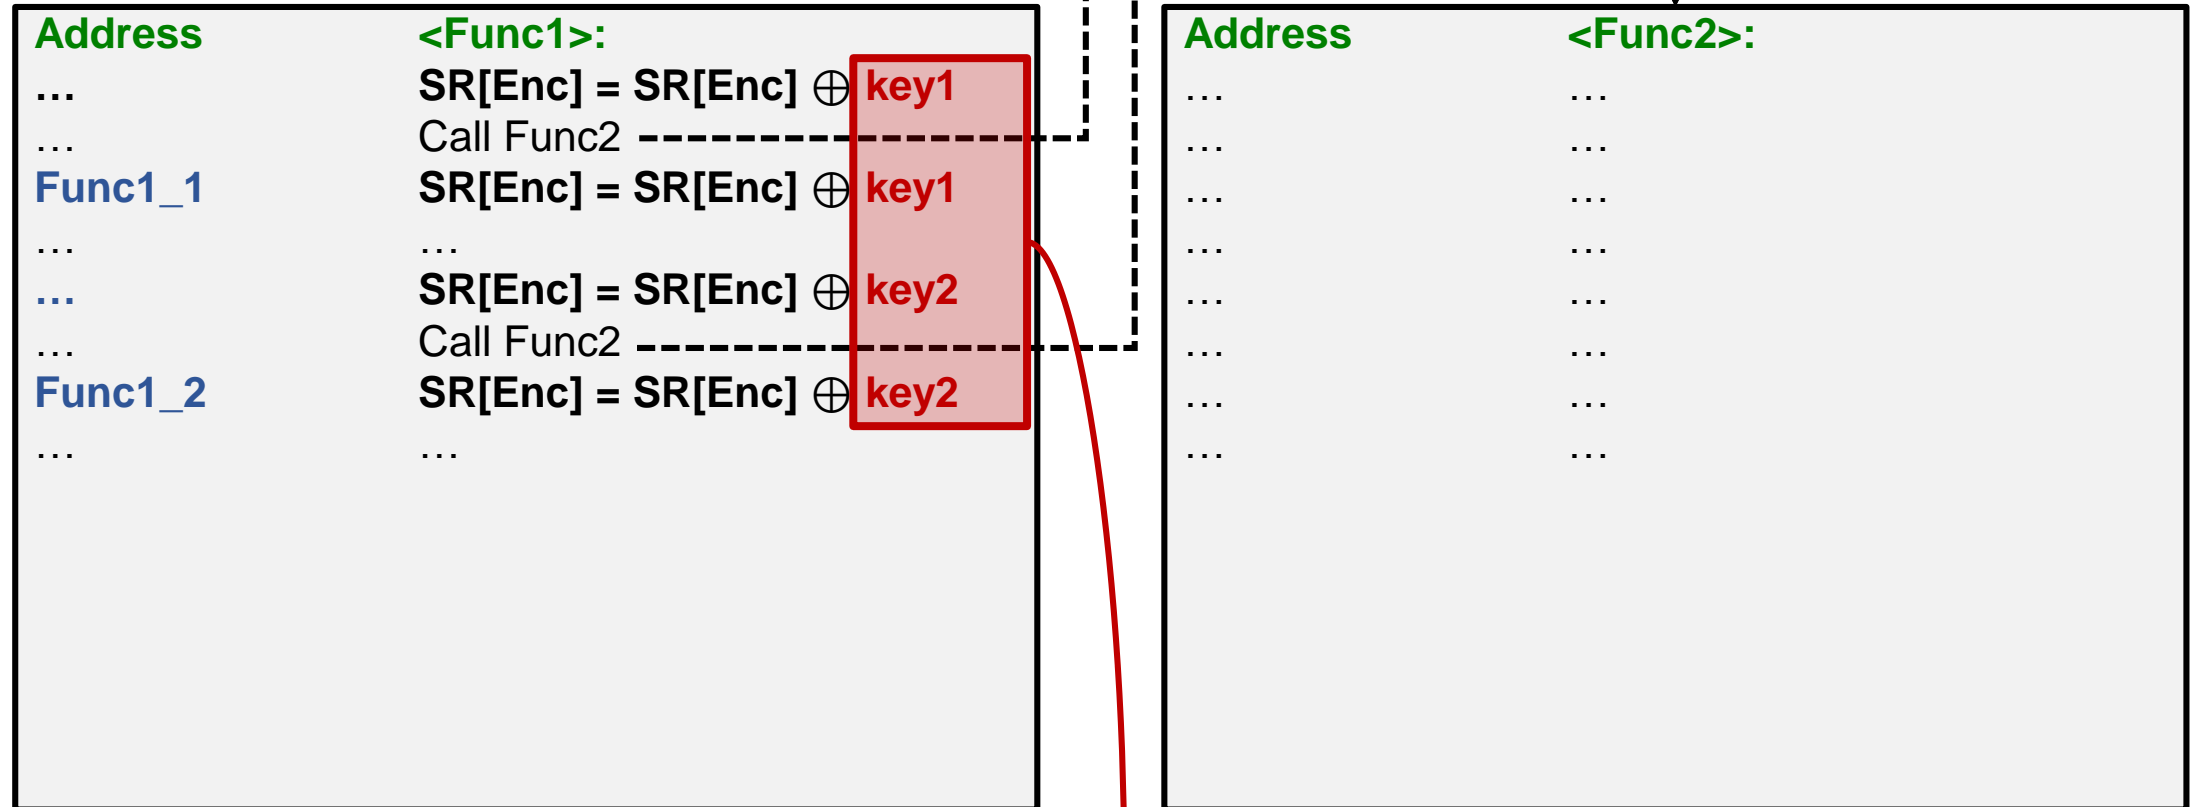


- **How encode SR?**

- **An XOR chain**

# μRAI: Terminology

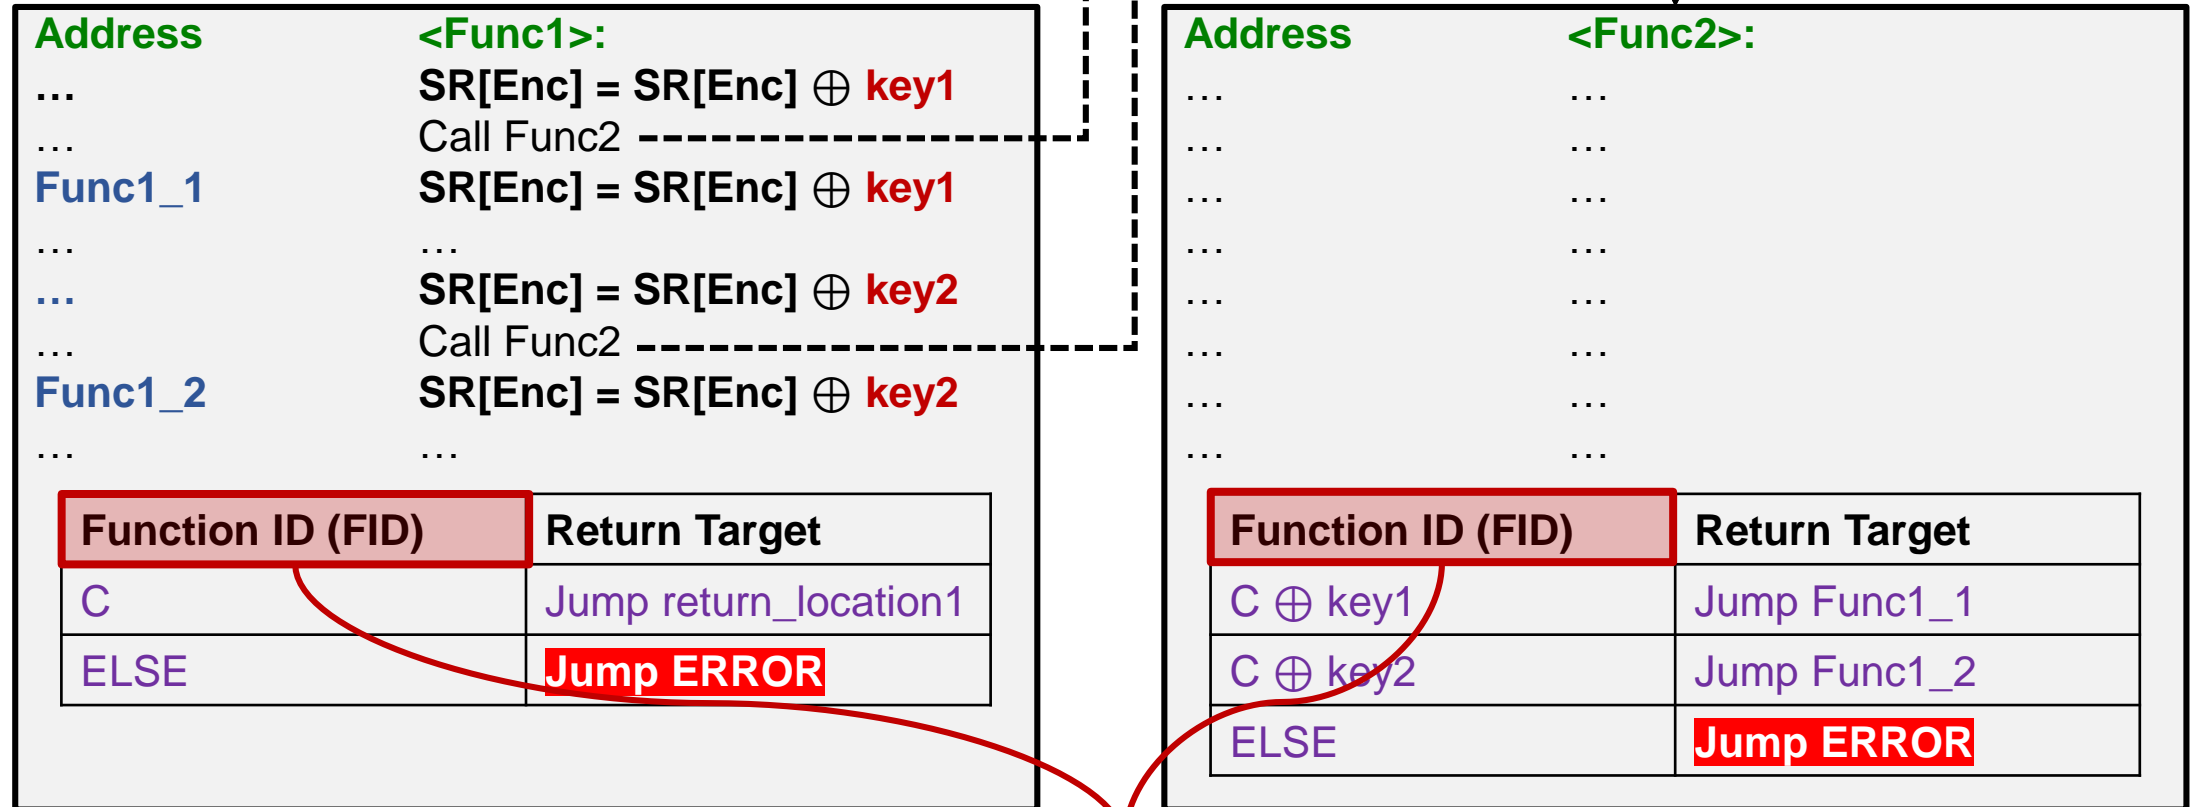
SR[Recursive]	SR[Encoded]
0	C



- **Function Keys (FKs):** Hard-coded keys used to encode the SR

# μRAI: Terminology

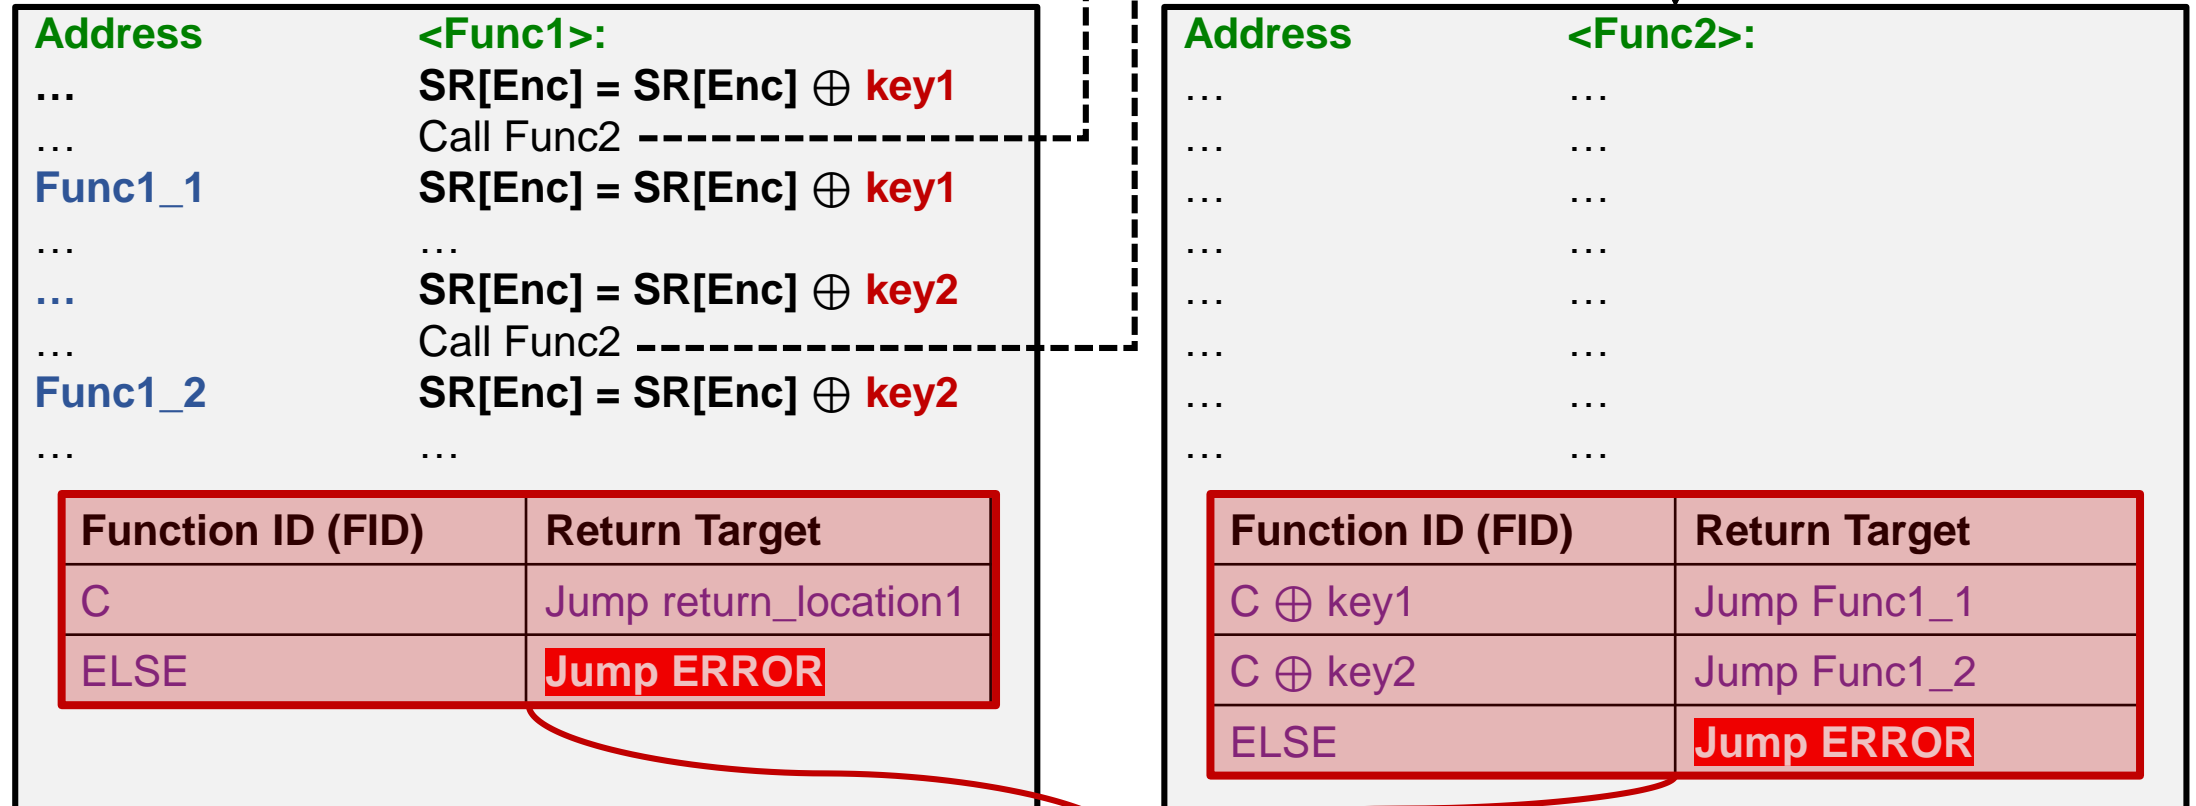
SR[Recursive]	SR[Encoded]
0	C



- **Function IDs (FIDs): Possible values of the SR for the function**

# μRAI: Terminology

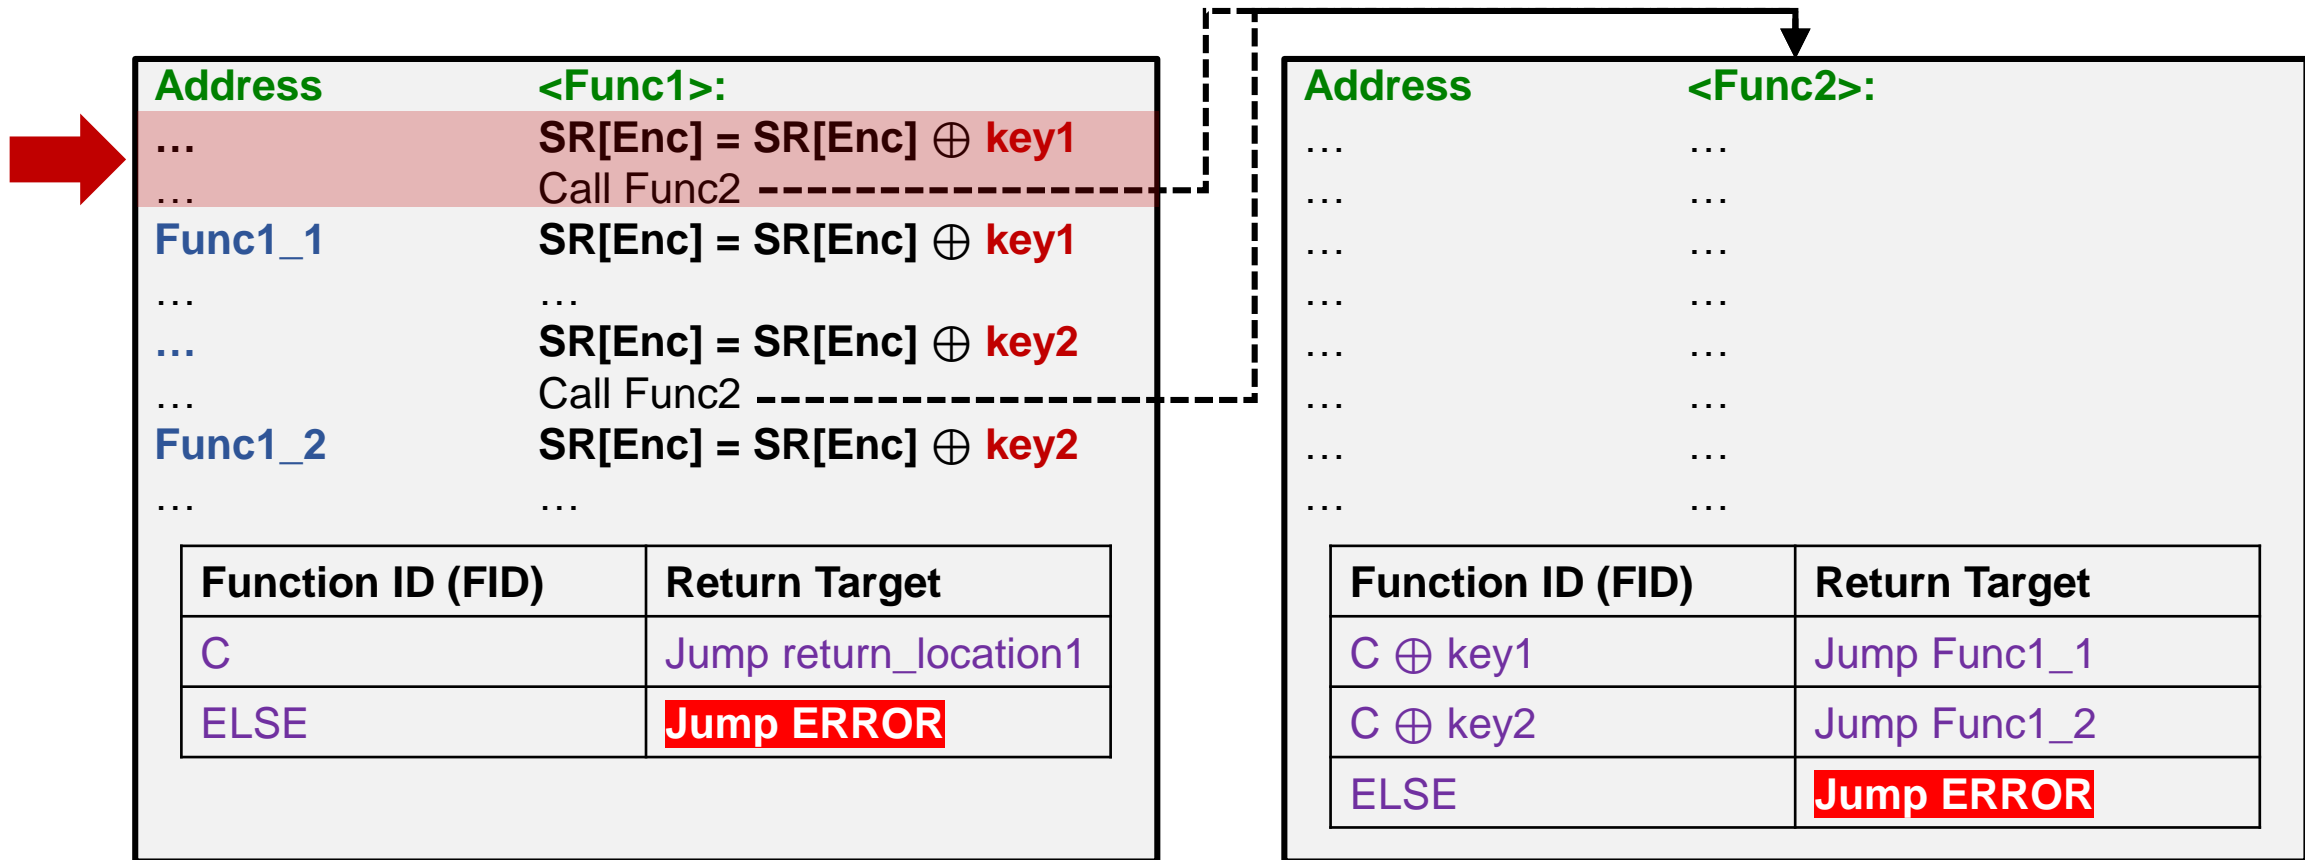
SR[Recursive]	SR[Encoded]
0	C



- **Function Lookup Table (FLT):** List of FIDs for the function

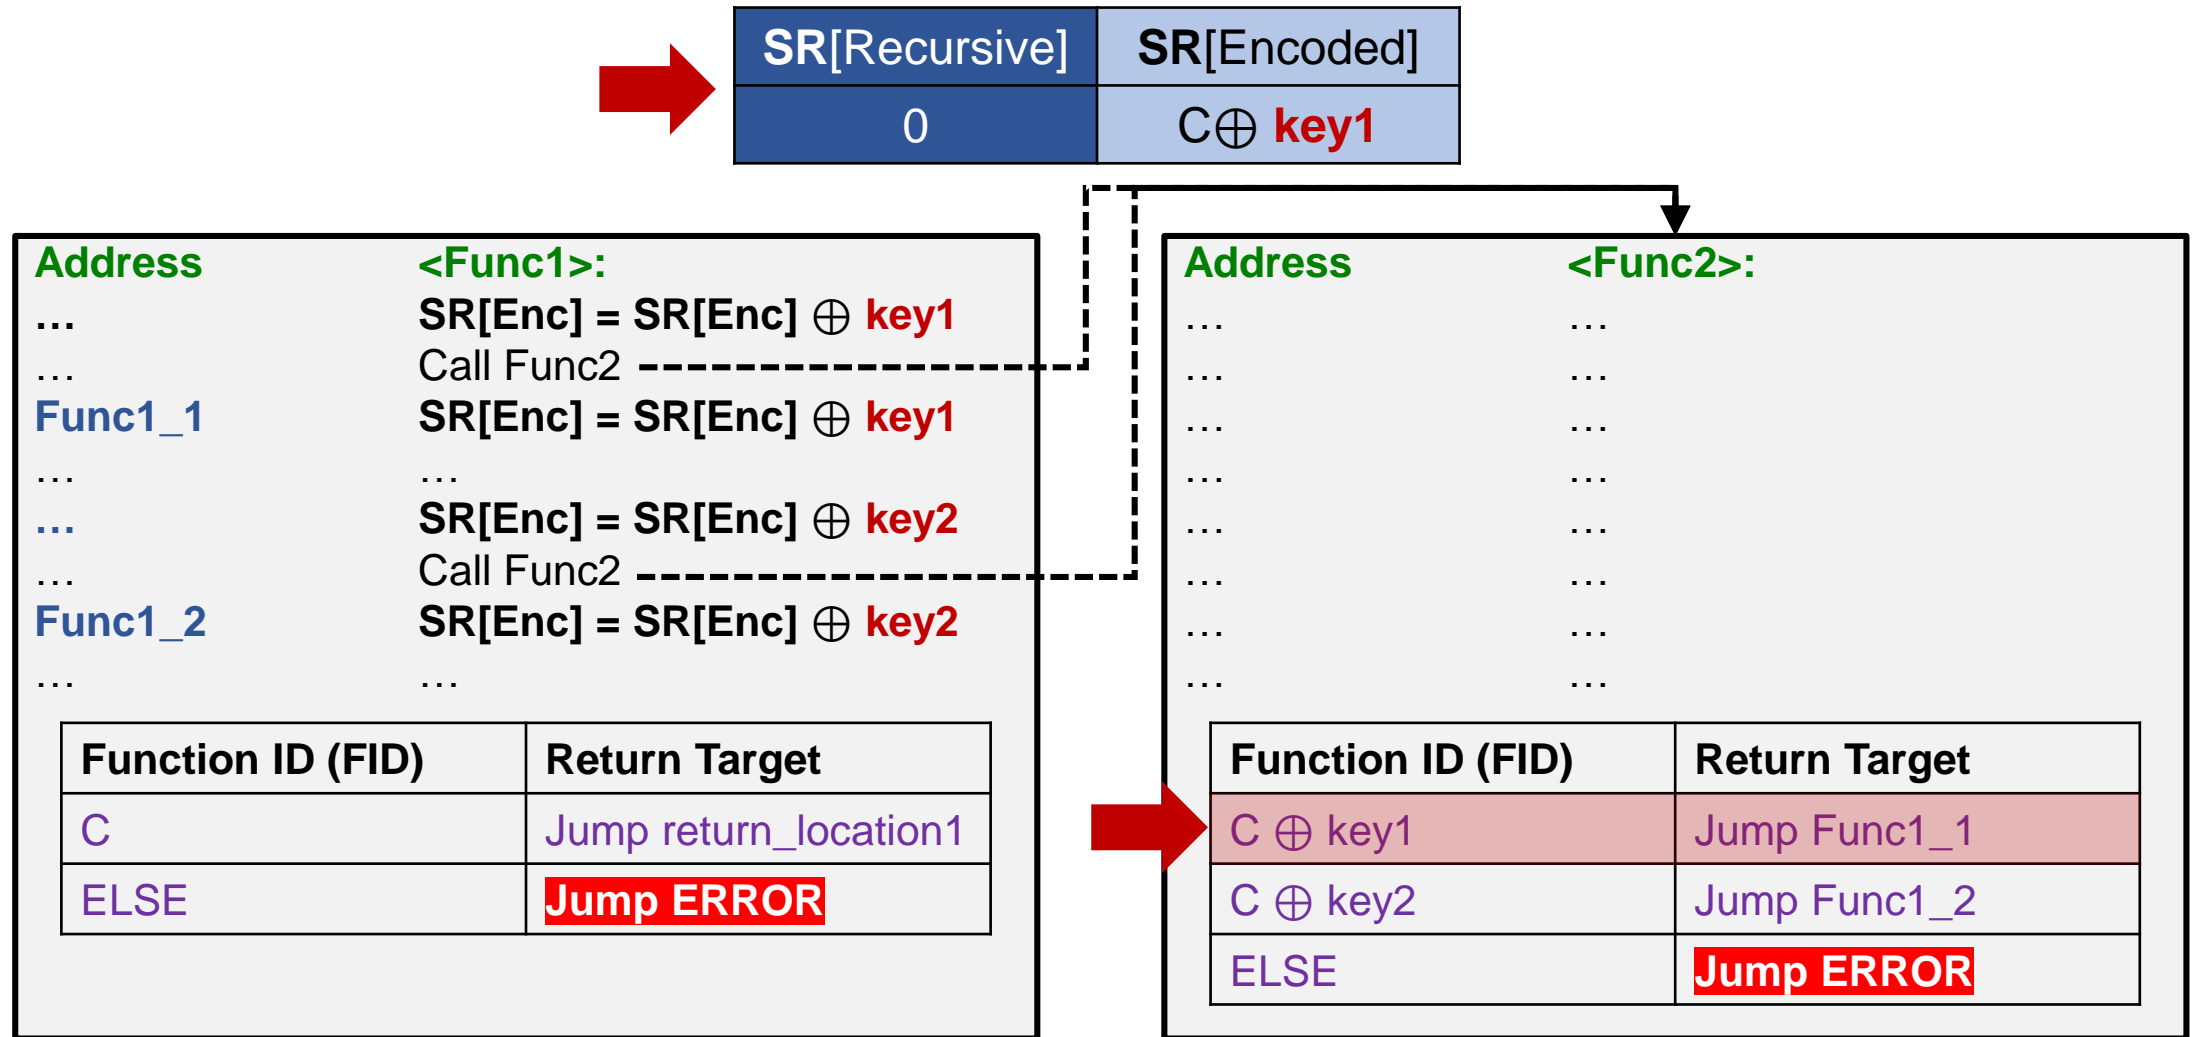
# μRAI: Transformation

SR[Recursive]	SR[Encoded]
0	C



- Encode the SR and call Func2

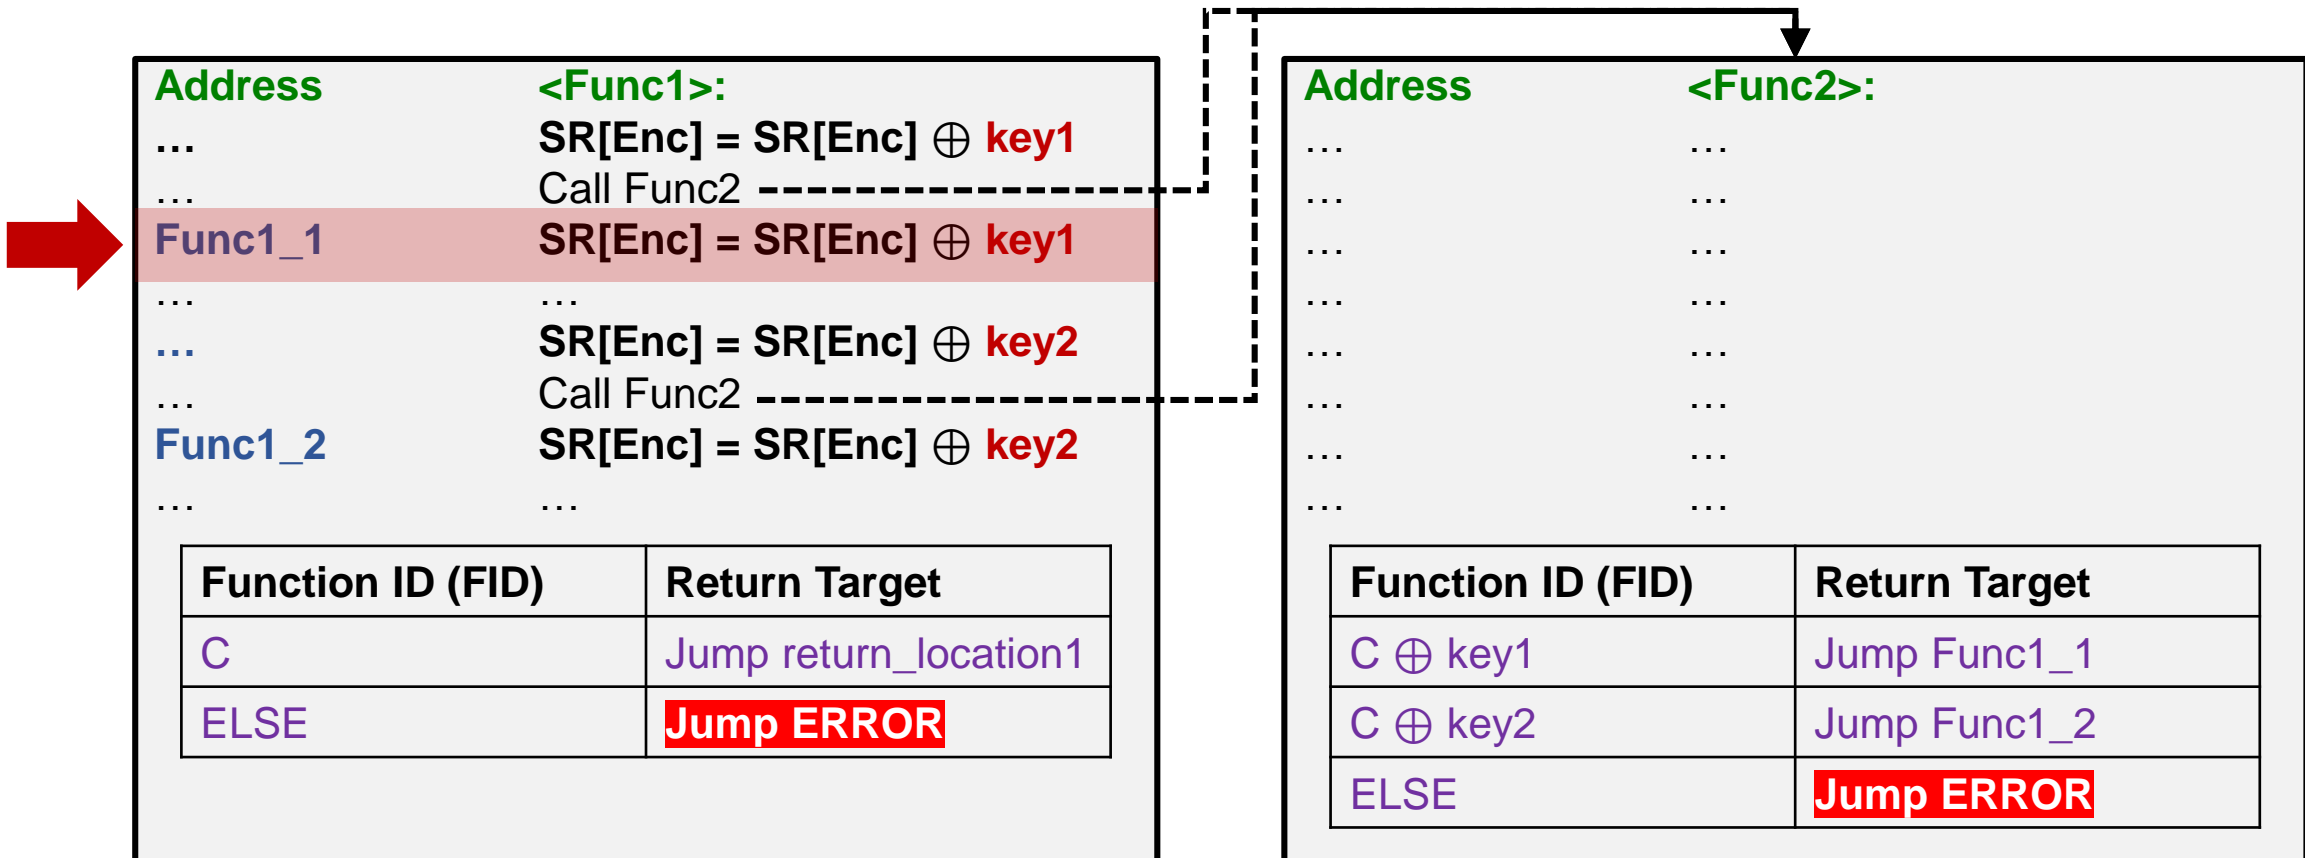
# μRAI: Transformation



- Func2 reads the SR and executes the corresponding direct jump

# μRAI: Transformation

SR[Recursive]	SR[Encoded]
0	$C \oplus \text{key1}$



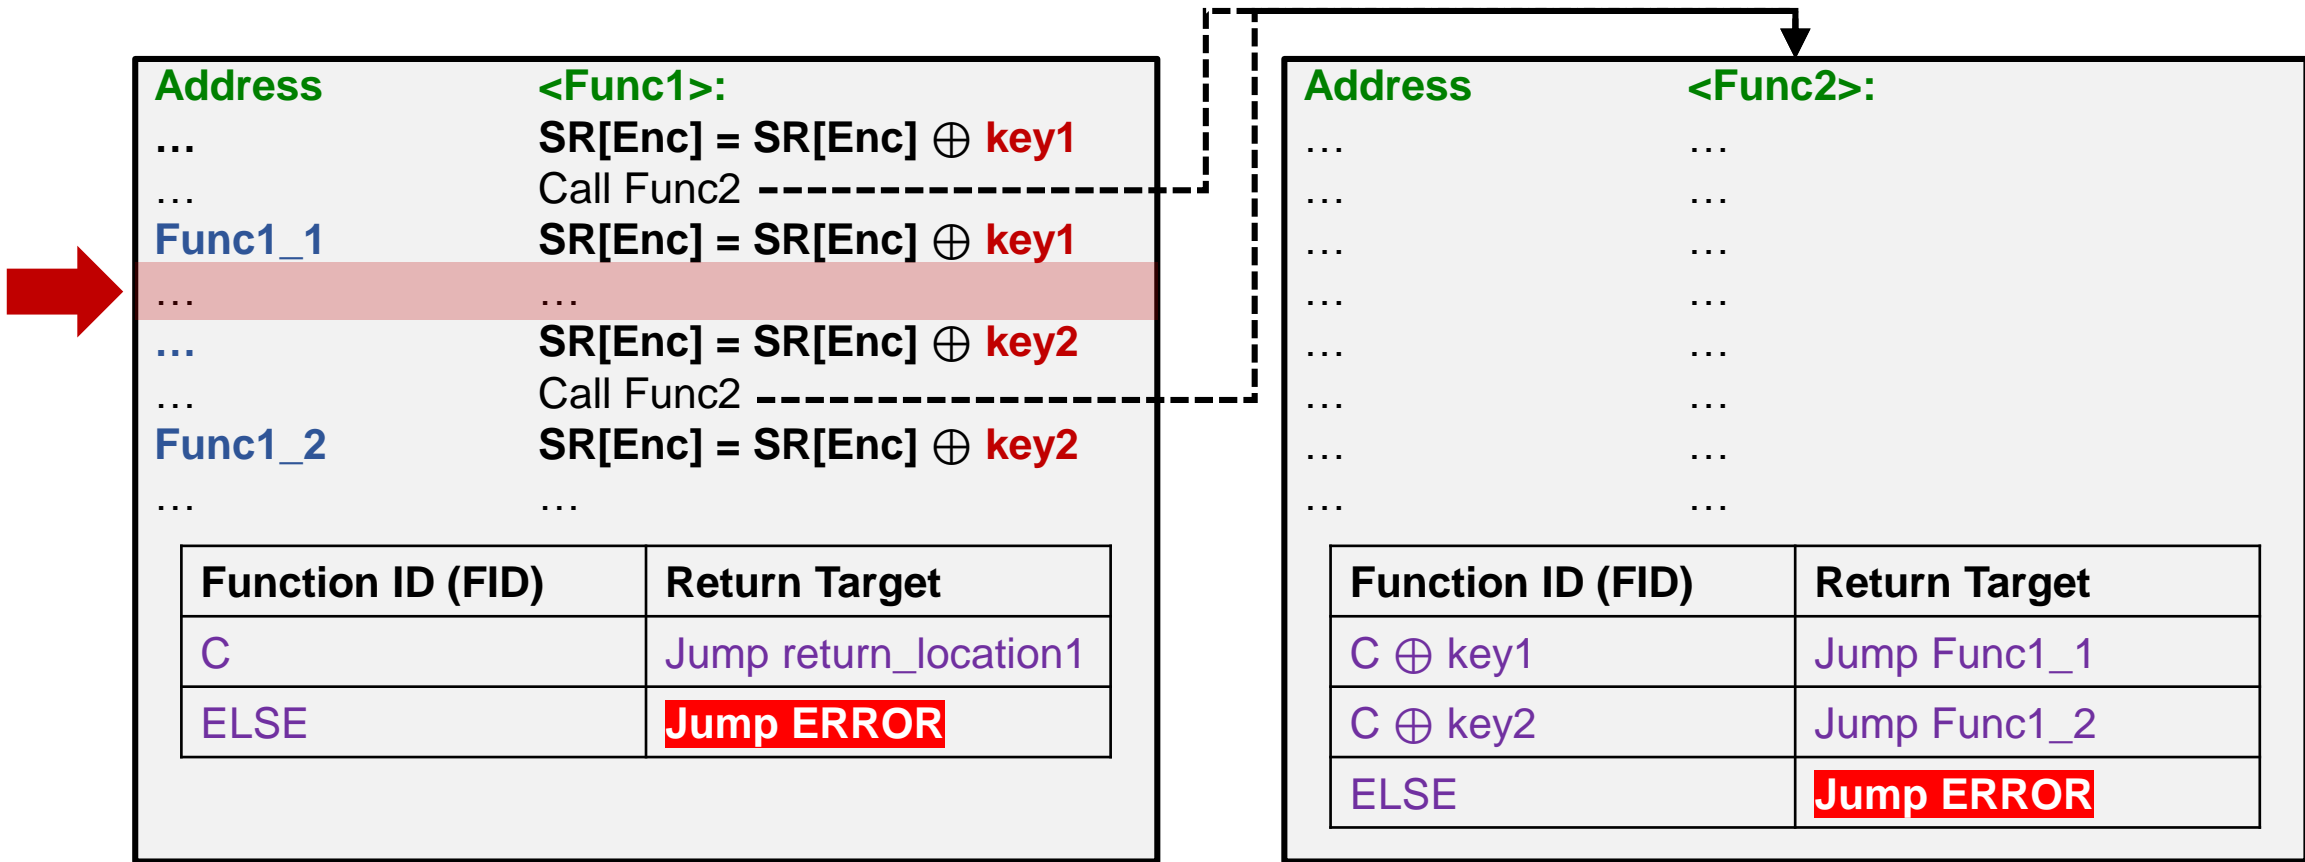
- Func2 returns correctly and the SR is decoded



# μRAI: Transformation

➔

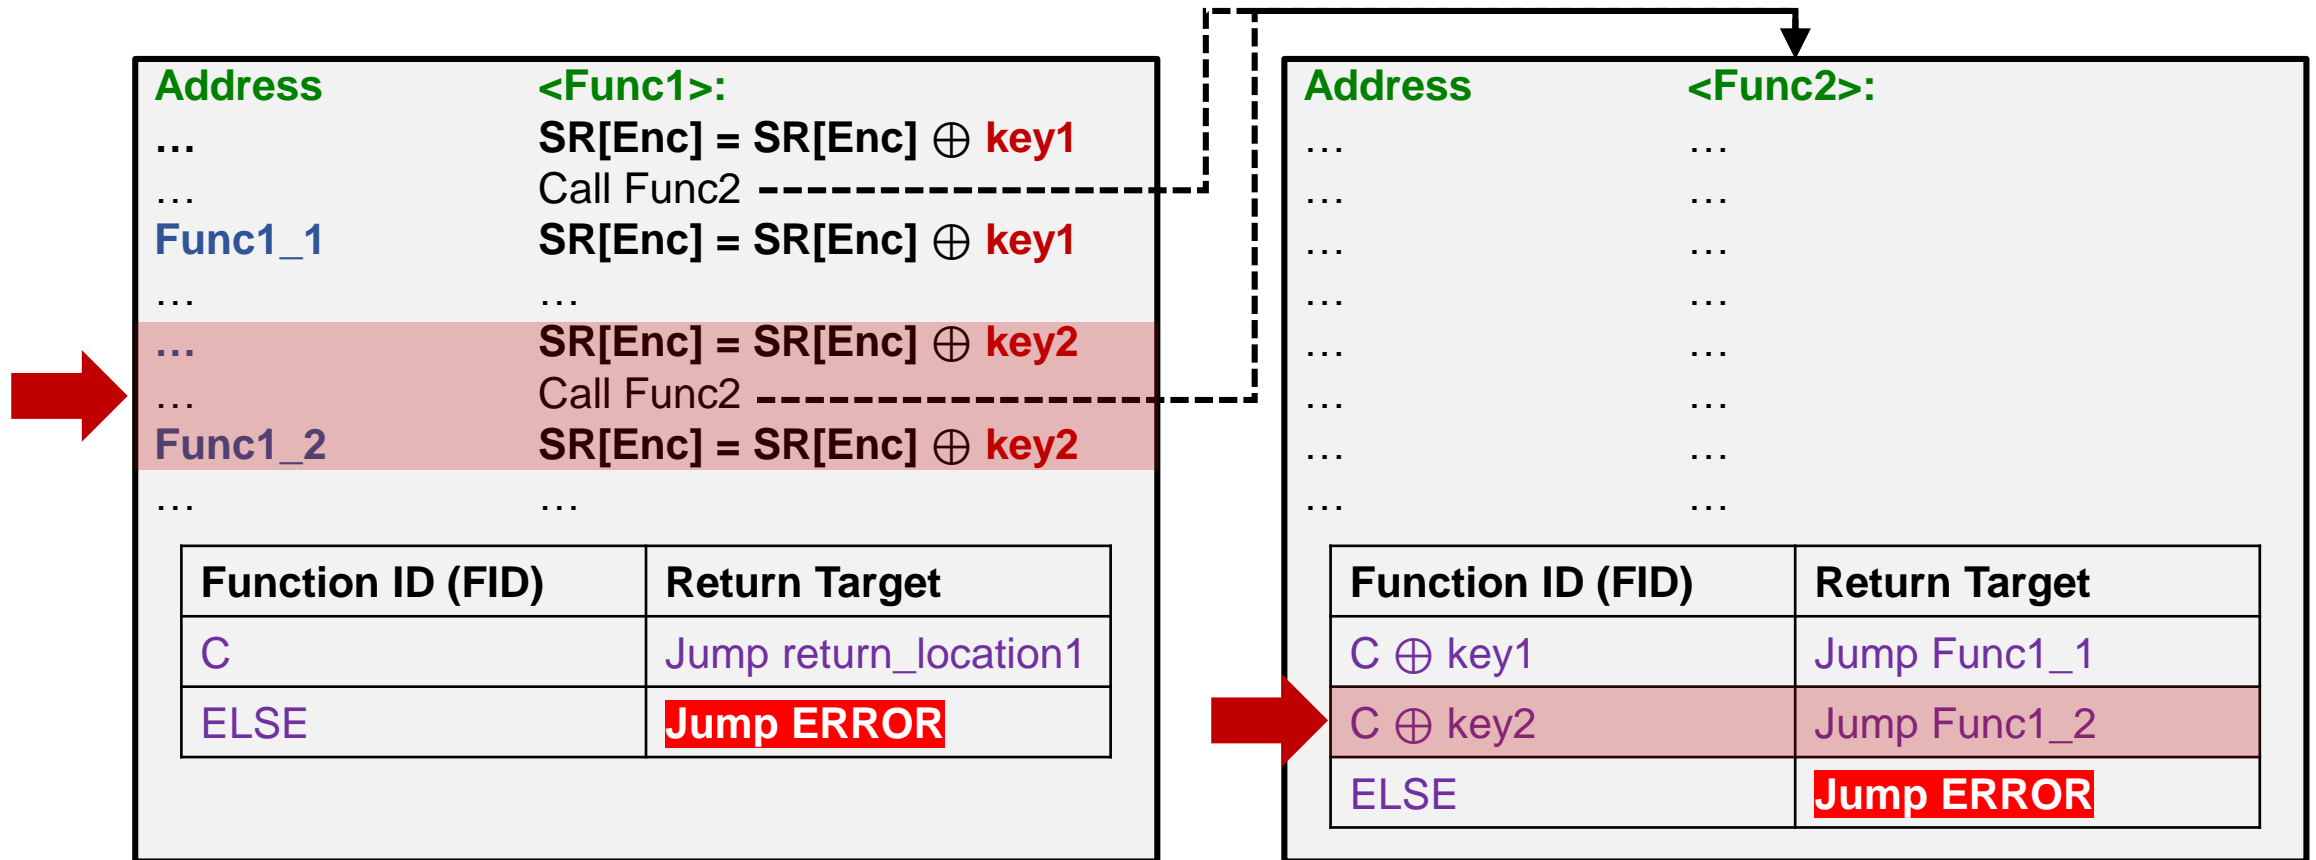
SR[Recursive]	SR[Encoded]
0	C



- The previous SR value is restored


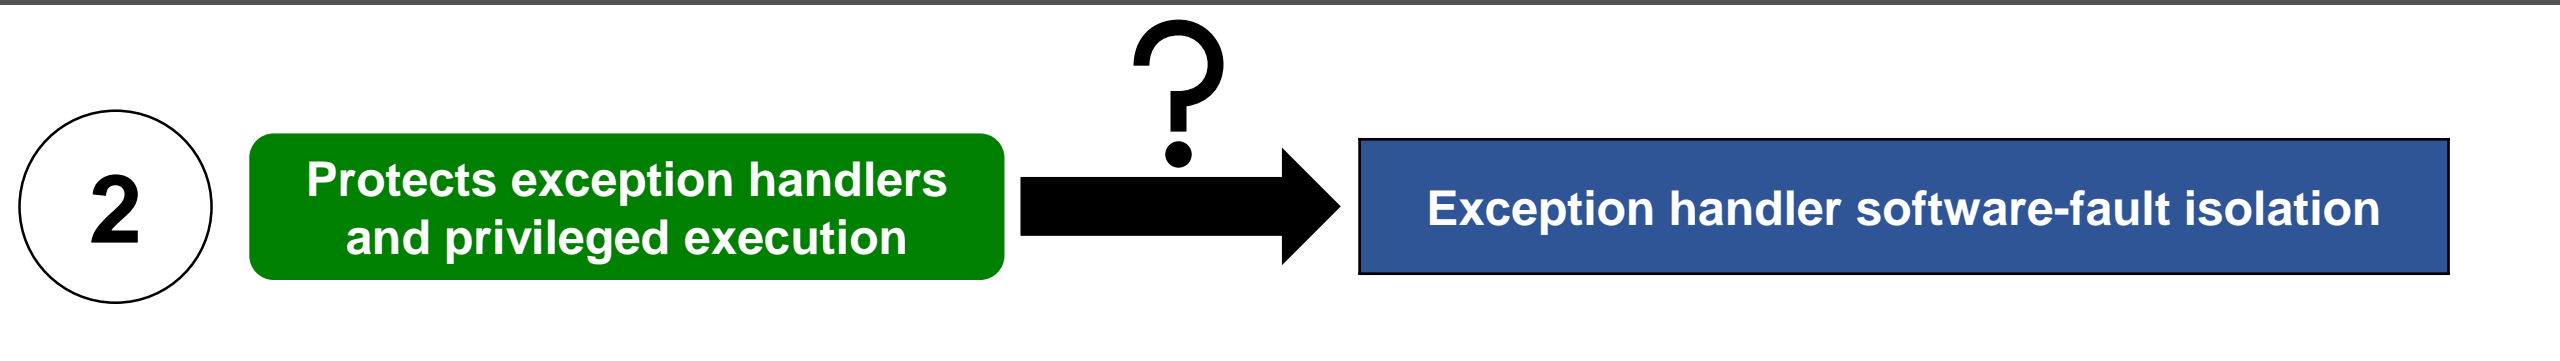
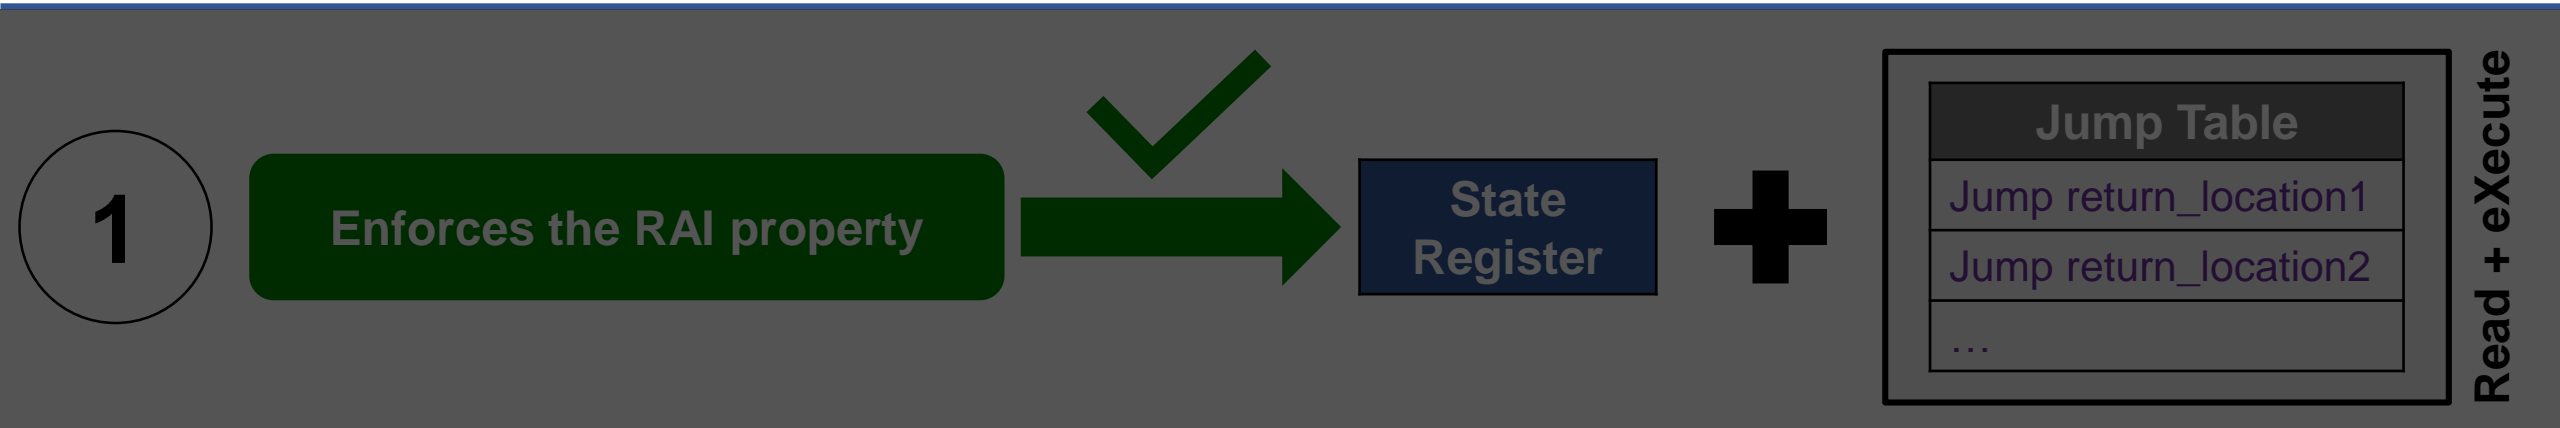
# μRAI: Transformation

SR[Recursive]	SR[Encoded]
0	C



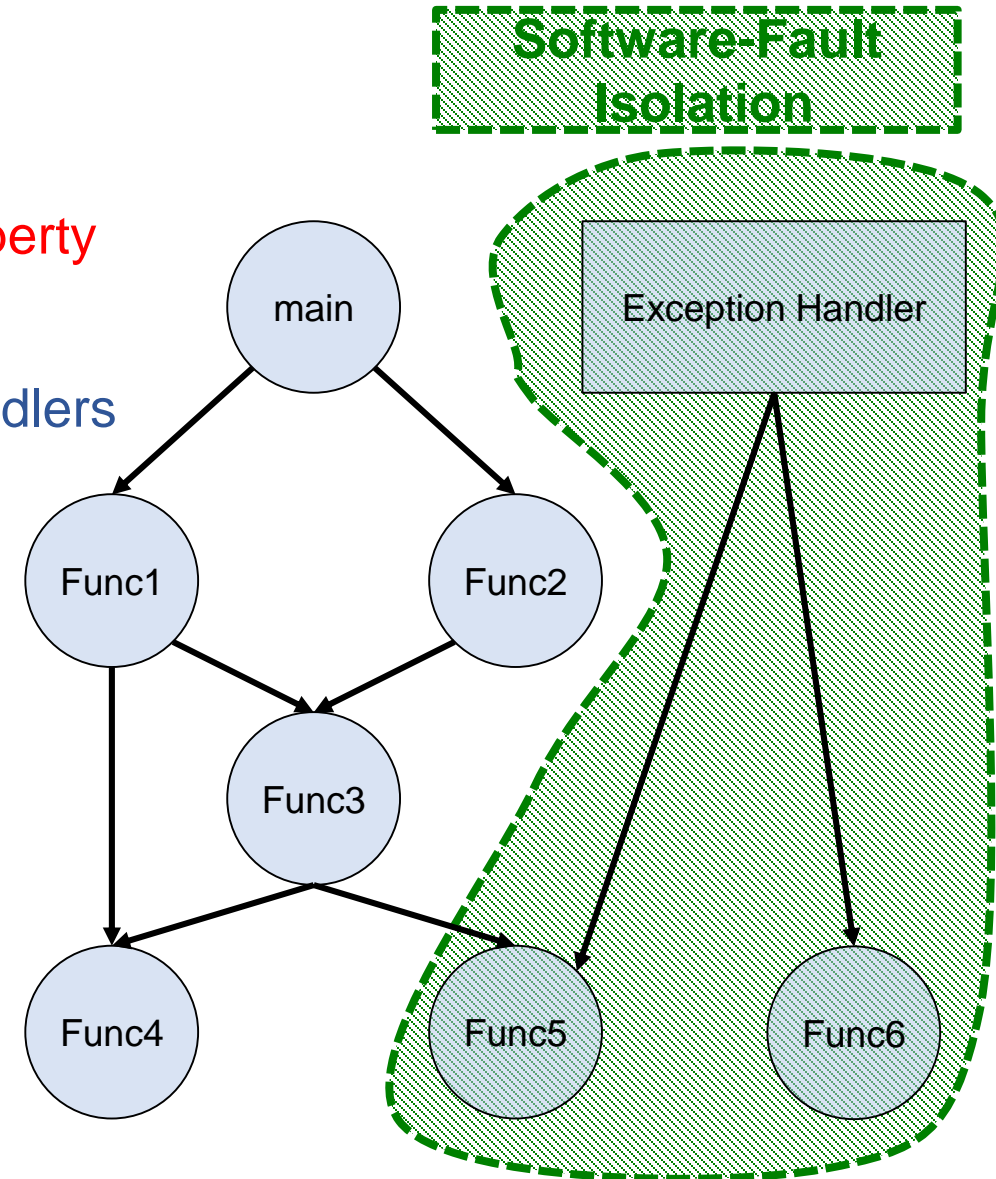

- The same happens for other calls. Func1 can then return correctly

# μRAI: Overview

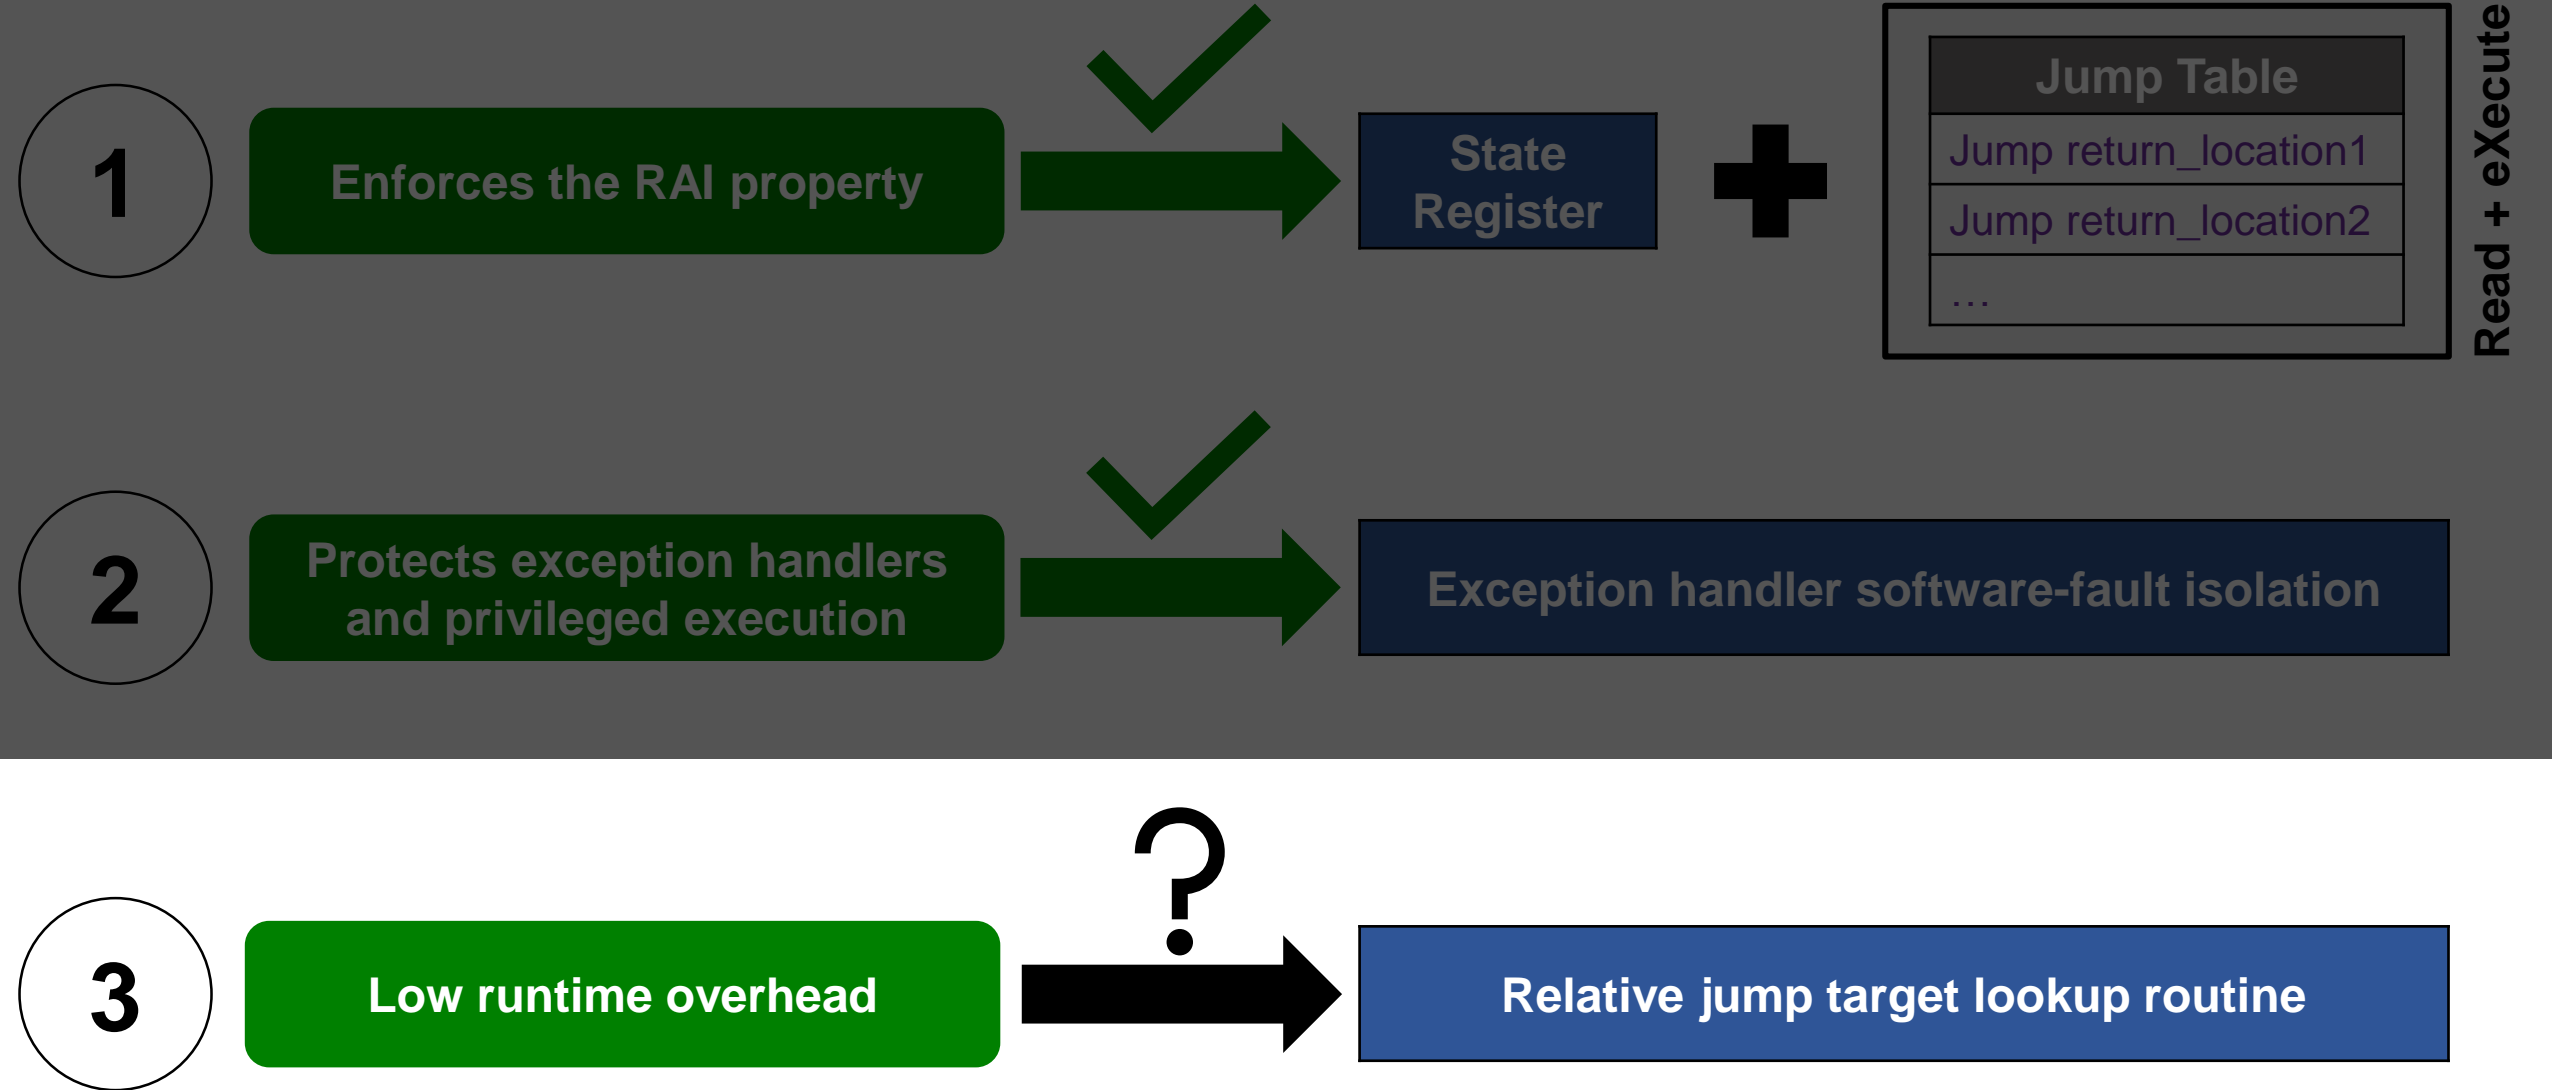


# μRAI: Enforce RAI for exception handlers

- **Exception handlers execute with privileges**
  - Can disable the MPU → enable code injection
  - Can corrupt exception stack frame → break RAI property
- **Solution:**
  - Apply SFI only to functions callable by exception handlers
  - Limit SFI overhead compared to full-SFI

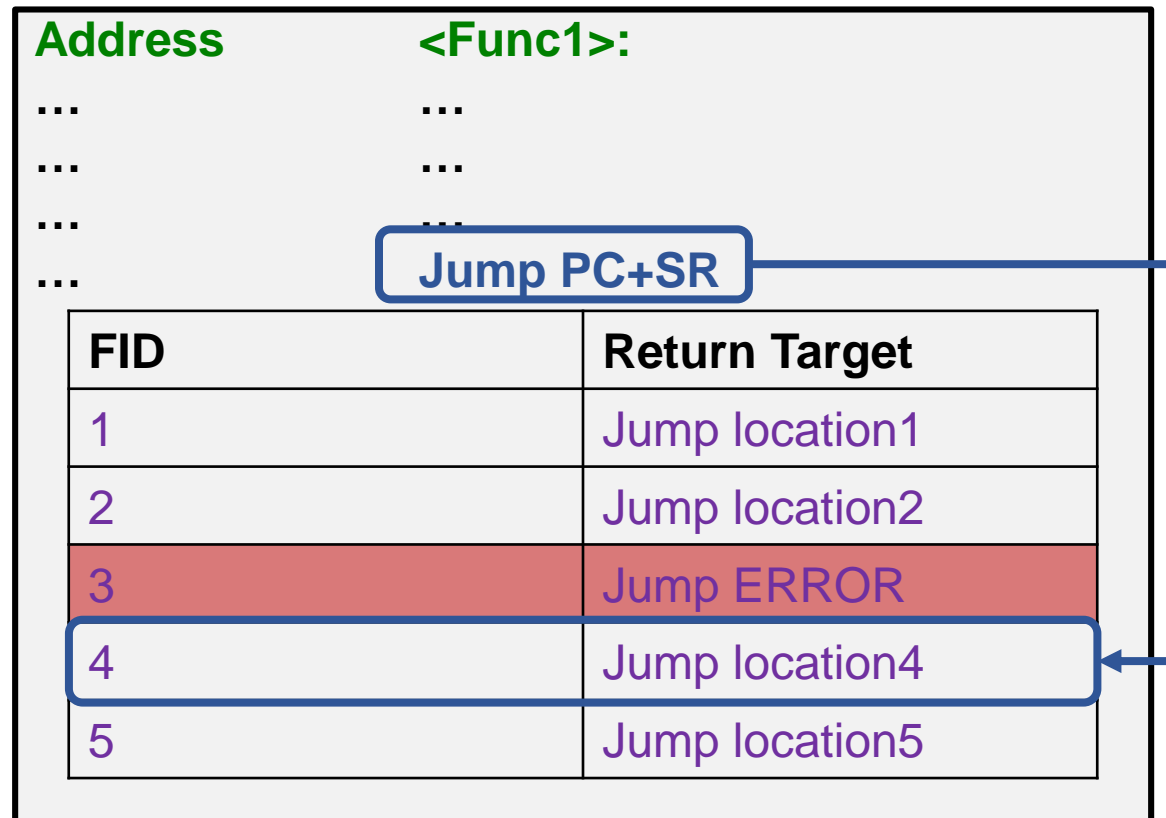


# μRAI: Overview



# Target Lookup Routine (TLR)

- How can we find the correct direct jump in the FLT efficiently?
  - Use a relative jump before the FLT
  - Resolve the correct return location efficiently regardless of FLT size



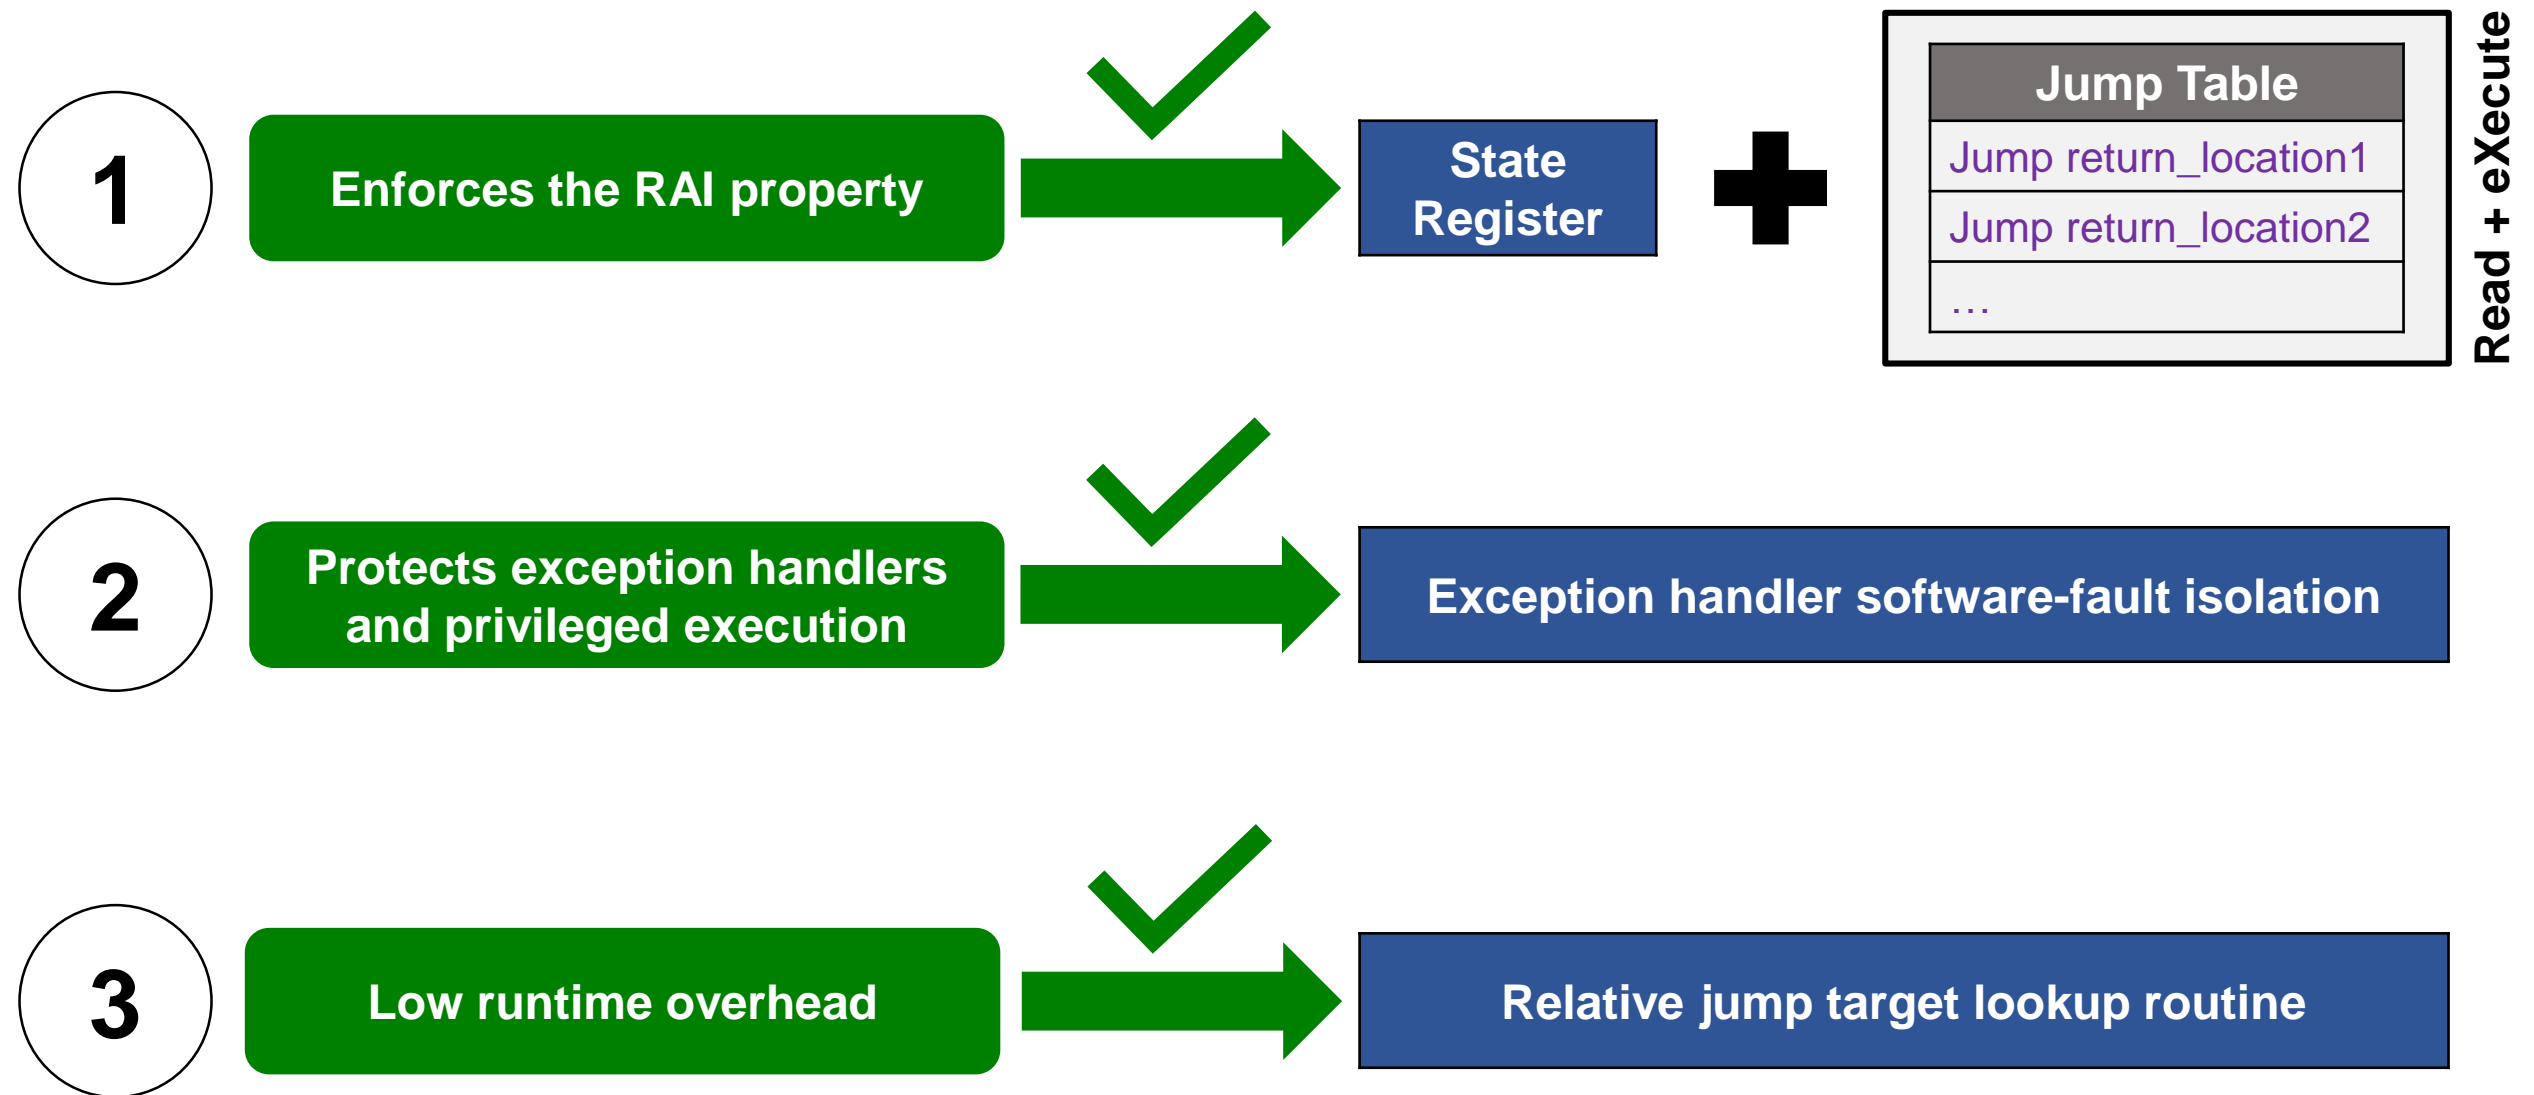
Assume the correct return location is **location4**

Comparing all FID can be slow!

Use SR as an index of a jump table

Align the FLT  
→ make SR = 4

# μRAI: Overview



# Evaluation

- **Five MCUS applications on Cortex-M4:**
  - PinLock
  - FatFs\_uSD
  - FatFs\_RAM
  - LCD\_uSD
  - Animation
- **CoreMark benchmark[1]**
  - Standard MCUS performance benchmark

[1] EEMBC, “Coremark - industry-standard benchmarks for embedded systems,” <http://www.eembc.org/coremark>




# Security Evaluation Using PinLock: Unlock The Lock


## Buffer overflow



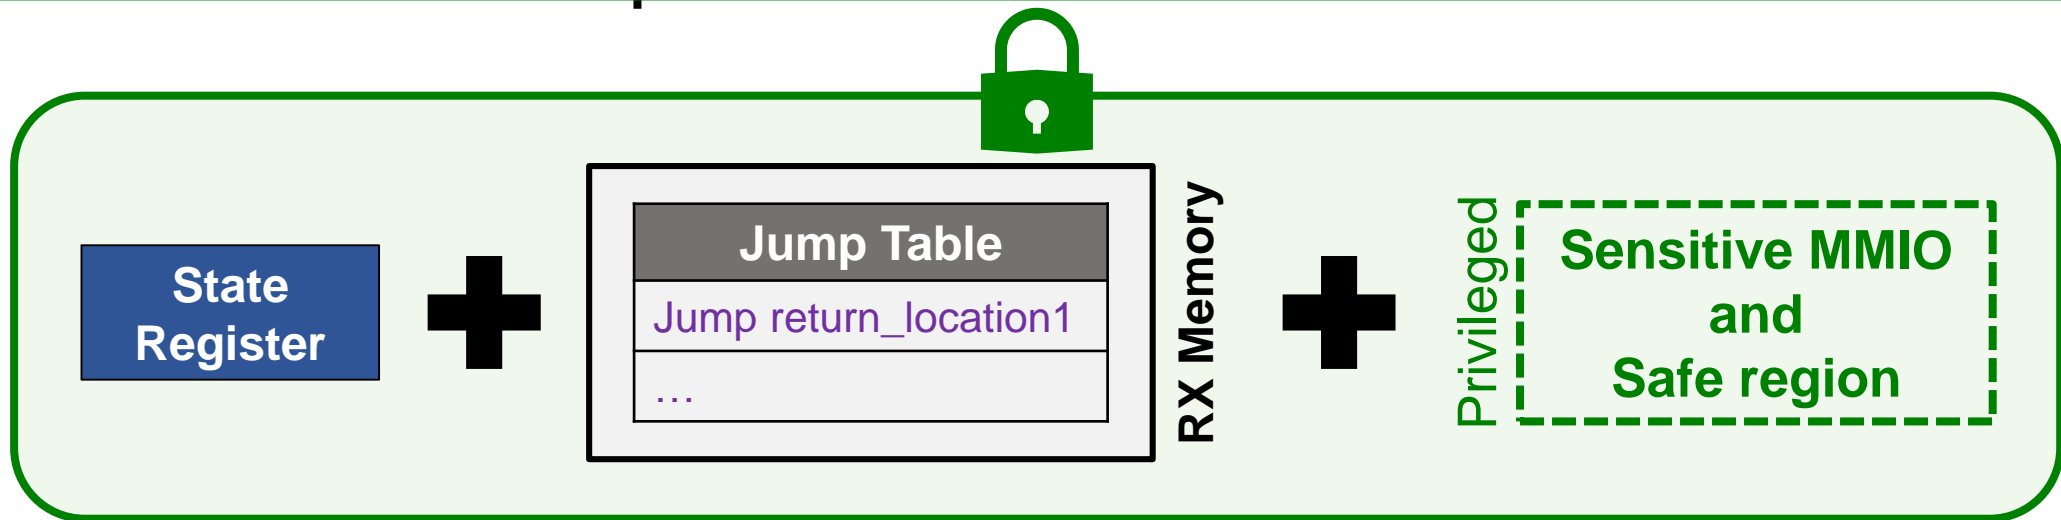
## Arbitrary write



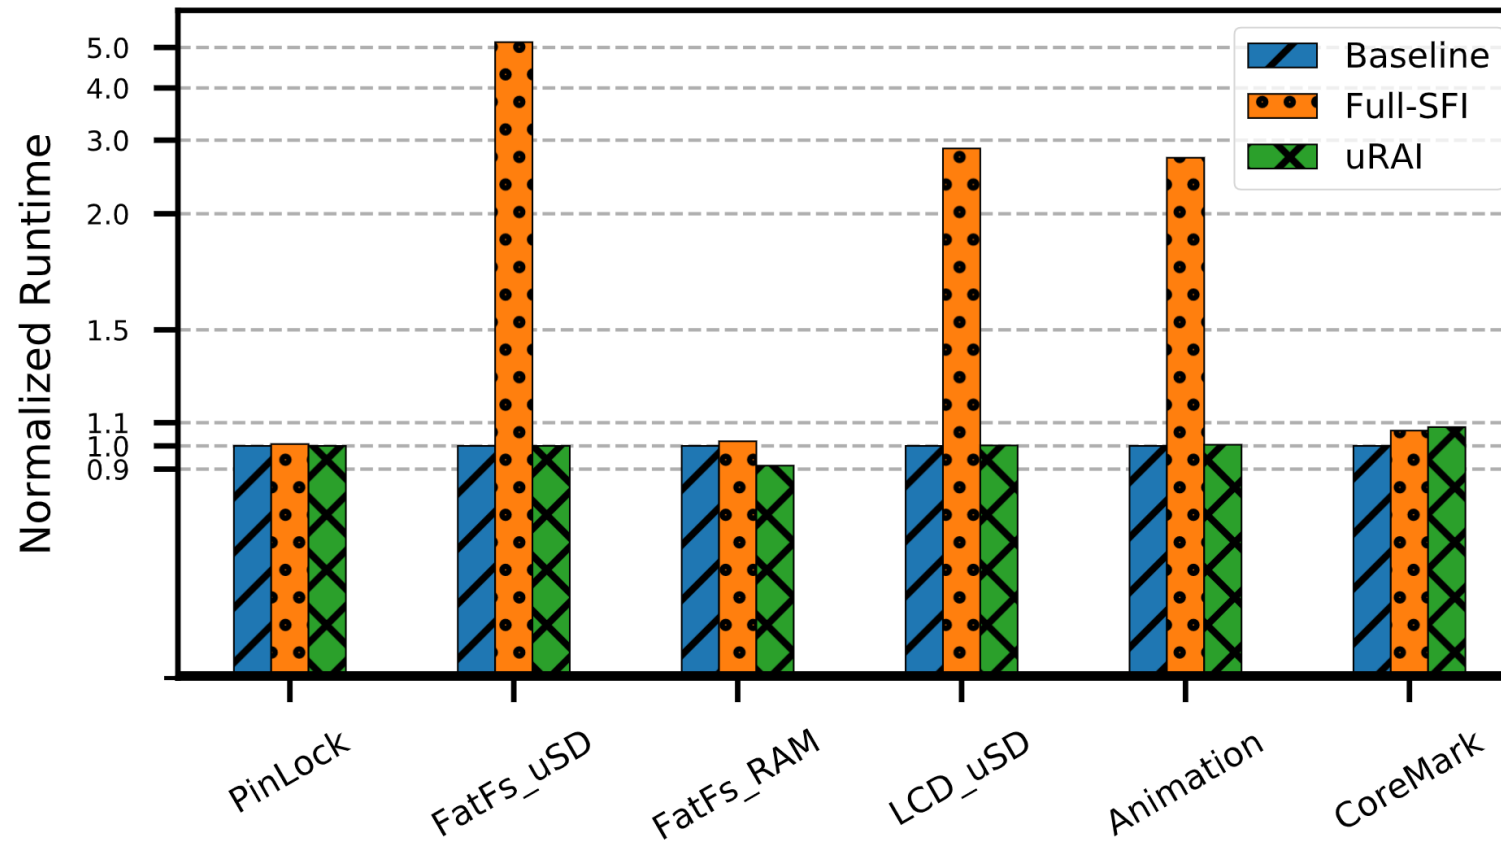
## Stack pivot



✓  $\mu$ RAI prevents all control-flow hijacking attack scenarios targeting return addresses



# Performance results



- Requiring full-SFI results in high overhead → **average of 130.5%**
- $\mu$ RAI results in low overhead → **average of 0.1%**

# μRAI: Conclusion

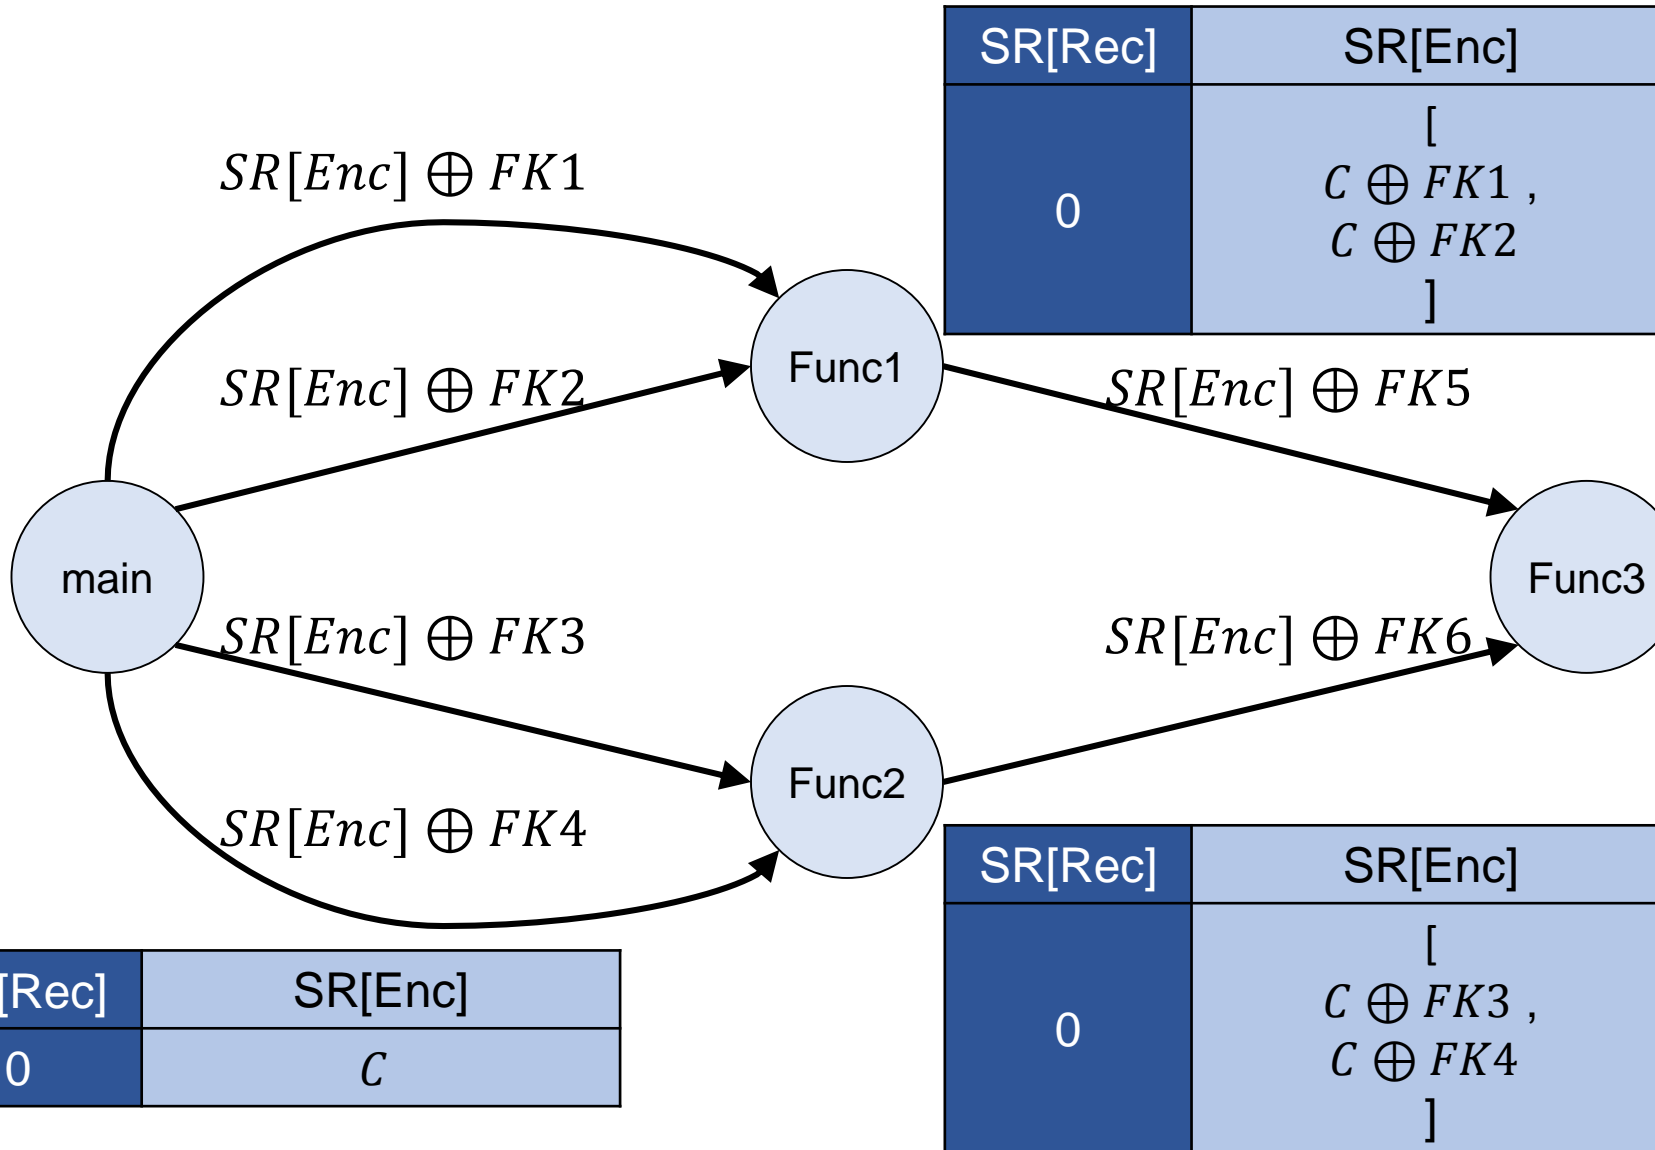
- **Control-flow hijacking on MCUS is a threat**
- **μRAI secures MCUS against control-flow hijacking**
  - Enforces the RAI property for MCUS → protects backward edges
  - Complemented with type-based CFI → end-to-end code pointer protection
- **Presents a portable encoding scheme**
  - Does not require special hardware features (only a register and an MPU)
  - Applicable to other systems
- **Low runtime overhead**

<https://github.com/embedded-sec/uRAI>

---

# Backup Slides

# Challenge: Path Explosion



- Func3 has 2 call sites
- FLT is size is 4!

SR[Rec]	SR[Enc]
0	[ C ⊕ FK1 ⊕ FK5, C ⊕ FK2 ⊕ FK5, C ⊕ FK3 ⊕ FK6, C ⊕ FK4 ⊕ FK6 ]

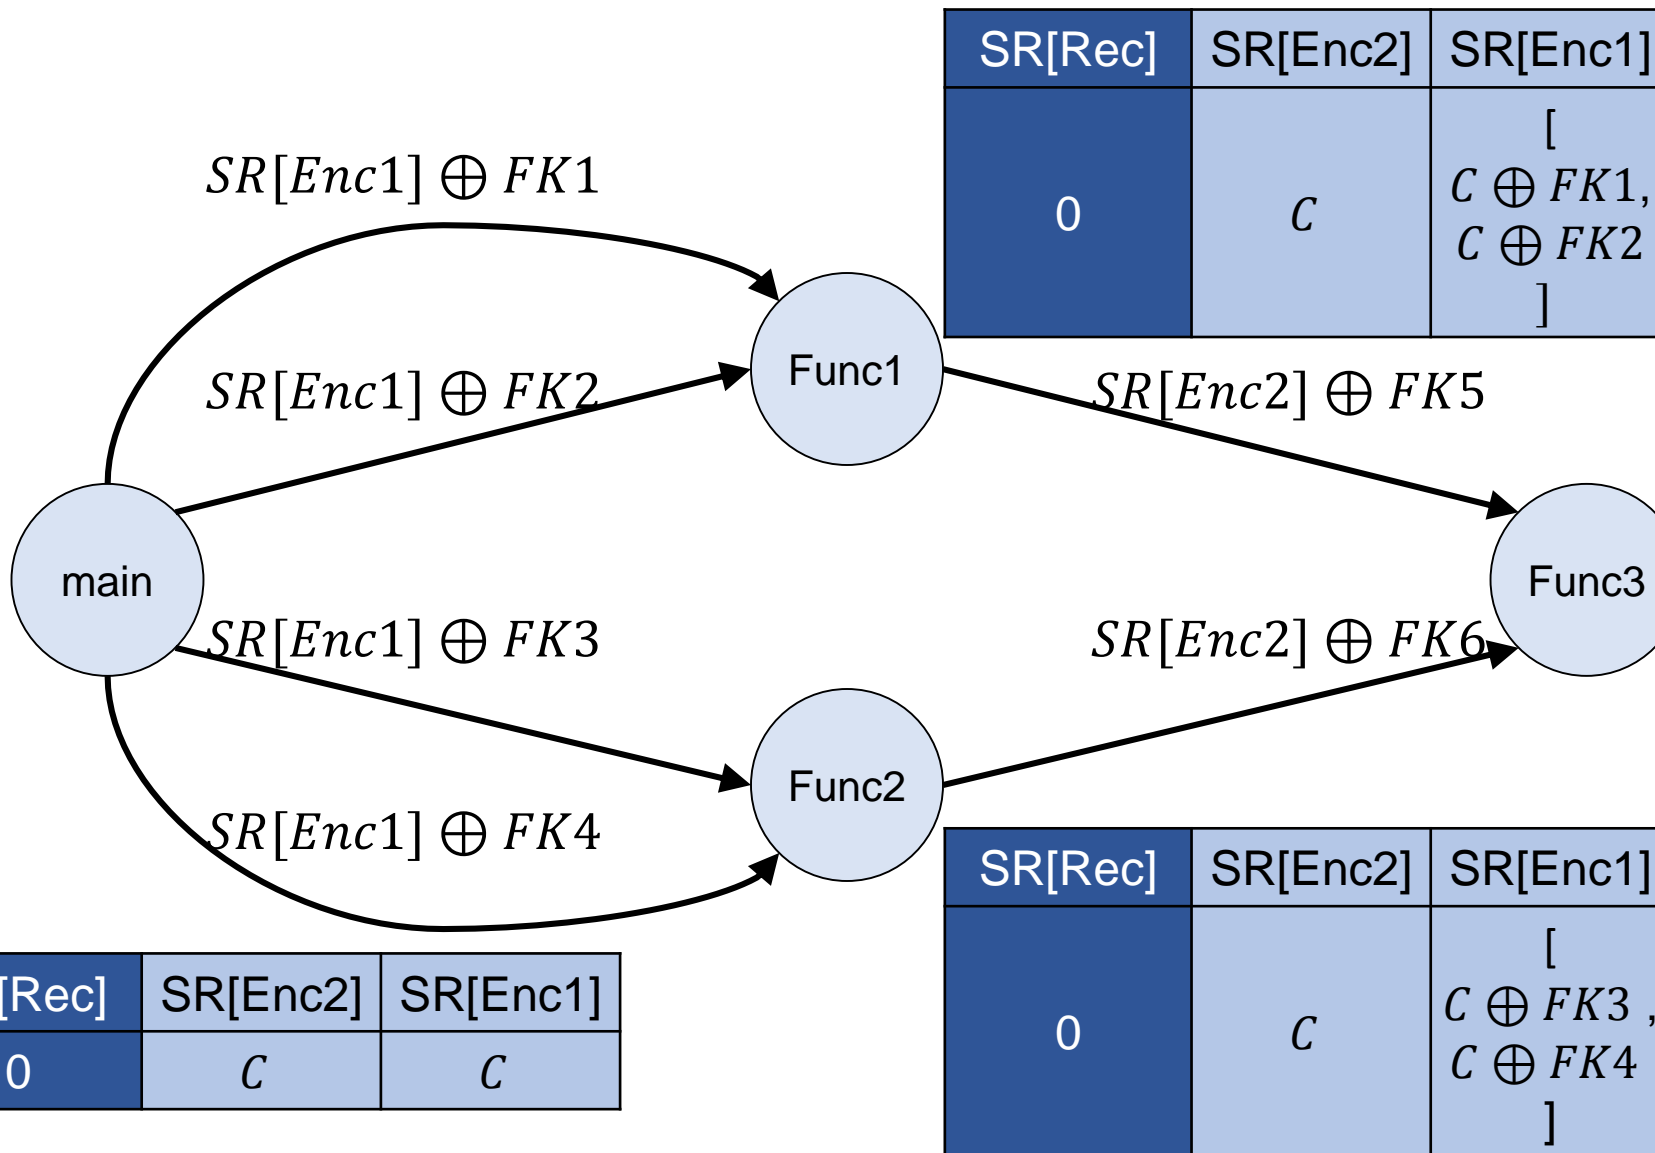
# Path Explosion Solution: Segmentation

- State register segmentation

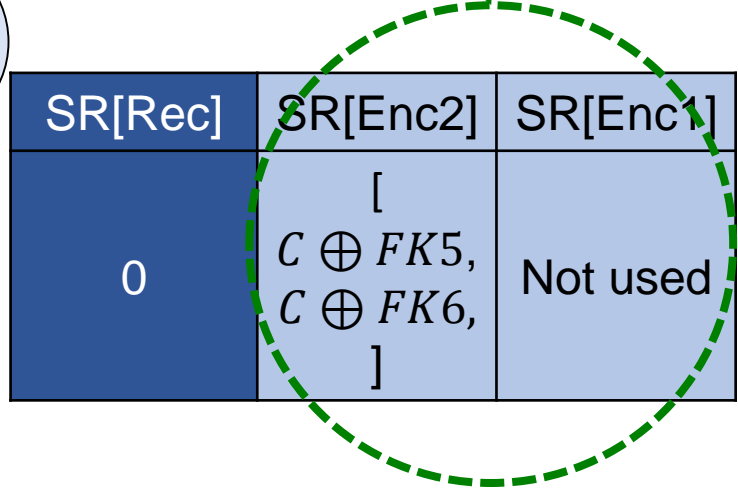
Recursion counter (Higher N bits)			Encoded value (Lower 32-N bits)			
Segment 1	...	Segment K	Segment M	...	Segment 2	Segment 1

- Functions only use the bits in their assigned segment.

# Path Explosion Solution: Segmentation




• Segmentation reduces FLT from 4 to 2



# μRAI: Scalability

- What if no more values can be found for the SR?
- Solution:
  - Partition the call graph if no solution is found
  - Entering a new partition  
→ Save and reset the SR to a privileged safe region
  - Returning to a previous partition  
→ Restore the SR





# Store Instructions Protected with EH-SFI


App	# of Store instruction			
	Static	Total	(Static/Total)%	Dynamic
PinLock	56	516	10.9	7
FatFs_uSD	99	1,802	5.5	906K
FatFs_RAM	7	1,116	0.6	7
LCD_uSD	99	2,814	3.5	48K
Animation	99	2,760	3.6	66K
CoreMark	56	1,024	5.5	7

# SR Layout: Recursion

- The SR has two parts:

- ENC: Encoded value
- REC: Recursion counter

• Cannot use XOR with recursion  
 • Collision occurs with existing values Func1  
 →  $SR \oplus ANY\ KEY \oplus ANY\ KEY = SR$



SR[Rec]	SR[Enc]
0	C

SR[Rec]	SR[Enc]
0	$C \oplus key1$

