



---

# Module 10 – Memory Safety and TEE

---

**(ECE 209) Secure and Advanced Computer Architecture**

**Nader Sehatbakhsh**

**Department of Electrical and Computer Engineering**

**University of California, Los Angeles**

# How do we enforce security in computers today?



# How do we enforce security in computers today?

## I- “Principle of Least Privilege”

*Users/Process should only have access to the data and resources needed to perform routine, authorized tasks.*



**Samueli**

School of Engineering

ECE 209 - Spring 24  
Nader Sehatbakhsh <[nsehat@ee.ucla.edu](mailto:nsehat@ee.ucla.edu)>

# How do we enforce security in computers today?

## I- “Principle of Least Privilege”

*Users/Process should only have access to the data and resources needed to perform routine, authorized tasks.*

- *This leads to “privilege separation”. Typically, “user” and “root” level access.*



# How do we enforce security in computers today?

## 2- Isolation

*A process cannot access (read or write) the memory content of any other process.*



**Samueli**

School of Engineering

ECE 209 - Spring 24  
Nader Sehatbakhsh <[nsehat@ee.ucla.edu](mailto:nsehat@ee.ucla.edu)>

# How do we enforce security in computers today?

## 3- Trusted Computing Base (TCB)

*Trust something (e.g., hardware), build everything on top/around that.*

- Only need to verify the TCB.
- Keep it simple and small so it can be easily(!) verified.



# Does it work?



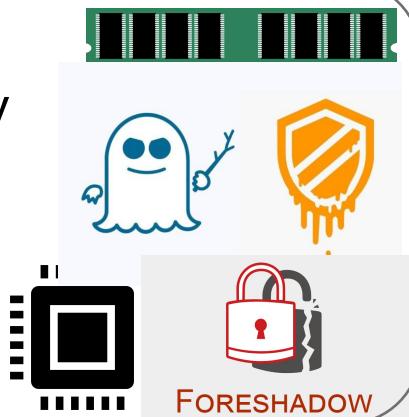
All processes were fully isolated!

# Types of Side-Channels

Digital/Microarchitectural



- cache
- memory
- u-arch
- TLB
- ...



Analog/Physical



- electromagnetic
- power
- temperature
- sound
- ...



# Is there more?

- Side-channel attacks can read/leak information but can we WRITE?



**Samueli**

School of Engineering

ECE 209 - Spring 24  
Nader Sehatbakhsh <[nsehat@ee.ucla.edu](mailto:nsehat@ee.ucla.edu)>

# Memory Safety

- Changing contents stored in memory to gain advantage!
- We have already seen an example:
  - physical attacks → rowhammer, (other physical/fault attacks)
  - but can we use other means?



# What can we do if we can modify memory?



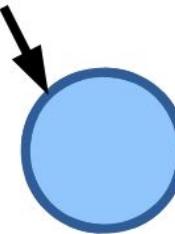
# Memory Safety

- How to modify memory? Remember that we need to somehow bypass isolation!



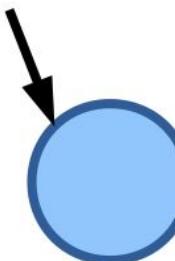
# Invalid Dereferencing

Dangling pointer:  
(temporal)

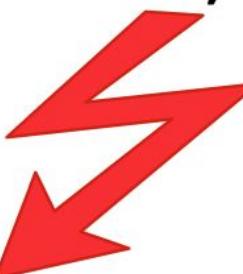


```
free(foo);  
*foo = 23;
```

Out-of-bounds pointer:  
(spatial)

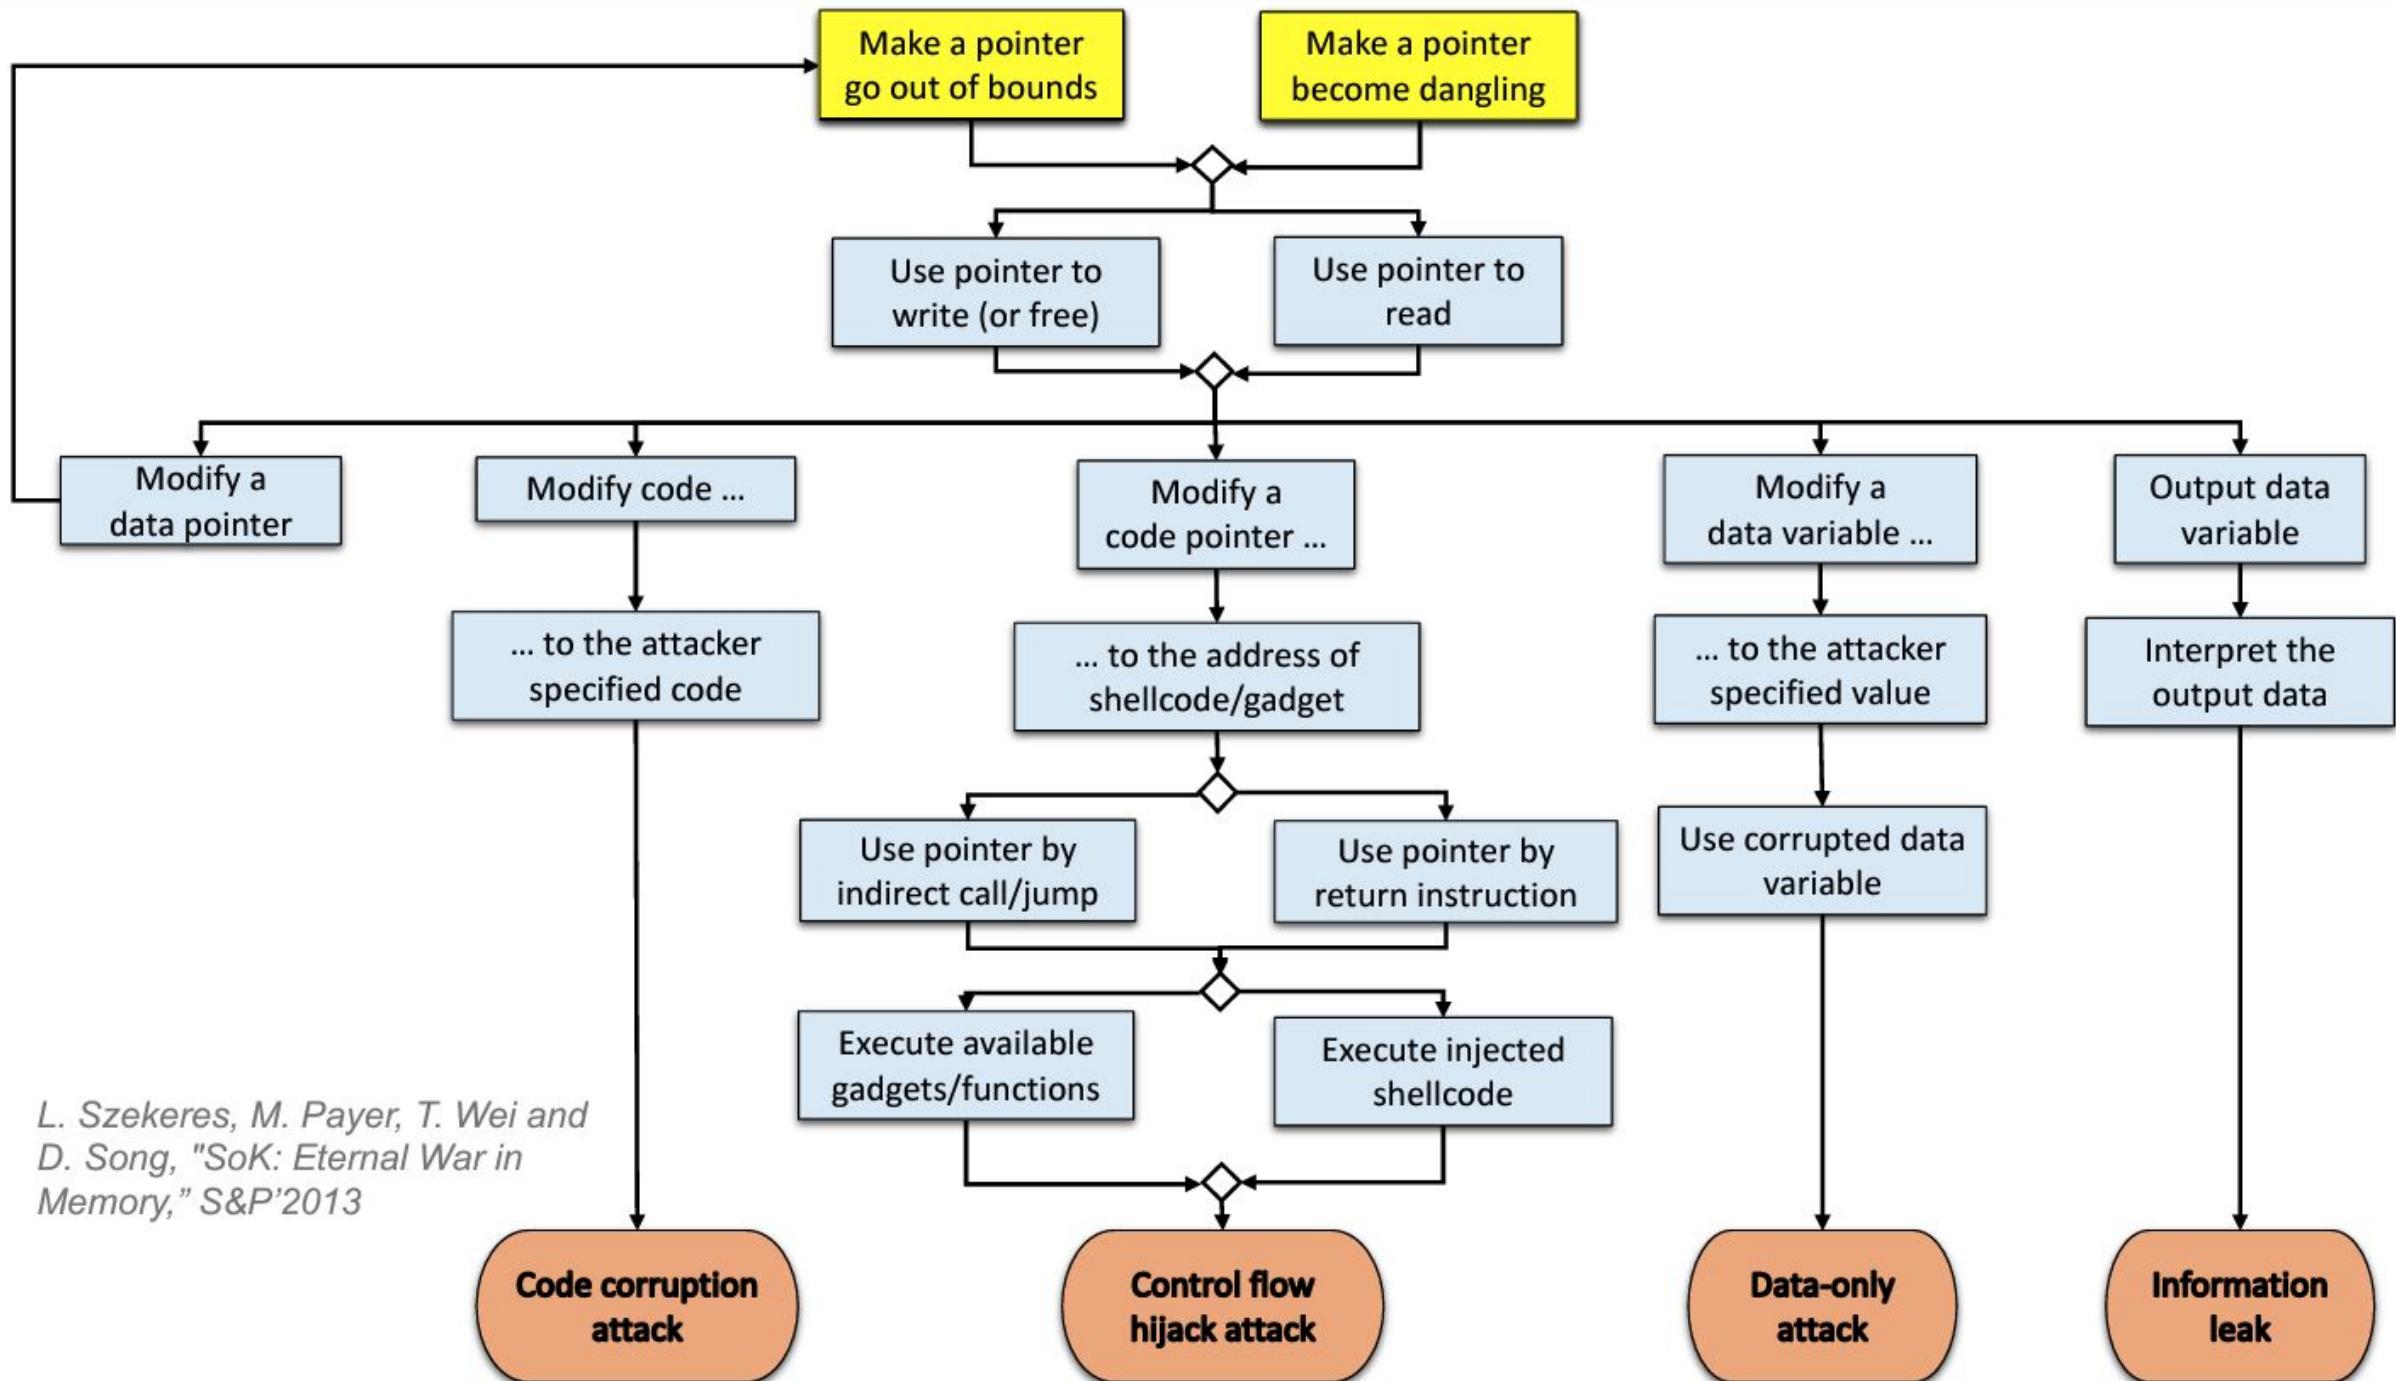


```
char foo[40];  
foo[42] = 23;
```



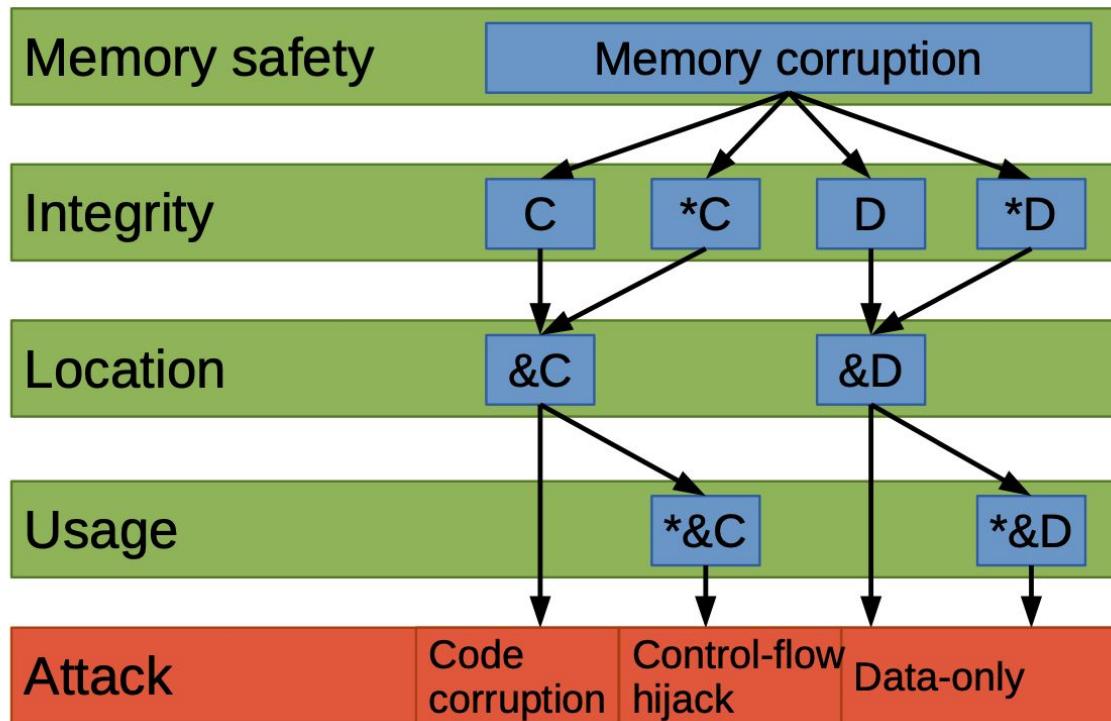
**Violation iff: pointer is read, written, or freed**

from SoK: Eternal War in Memory, Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song In: Oakland '14



L. Szekeres, M. Payer, T. Wei and  
D. Song, "SoK: Eternal War in  
Memory," S&P'2013

# What can happen?



from SoK: Eternal War in Memory, Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song In: Oakland '14

# How to defend?

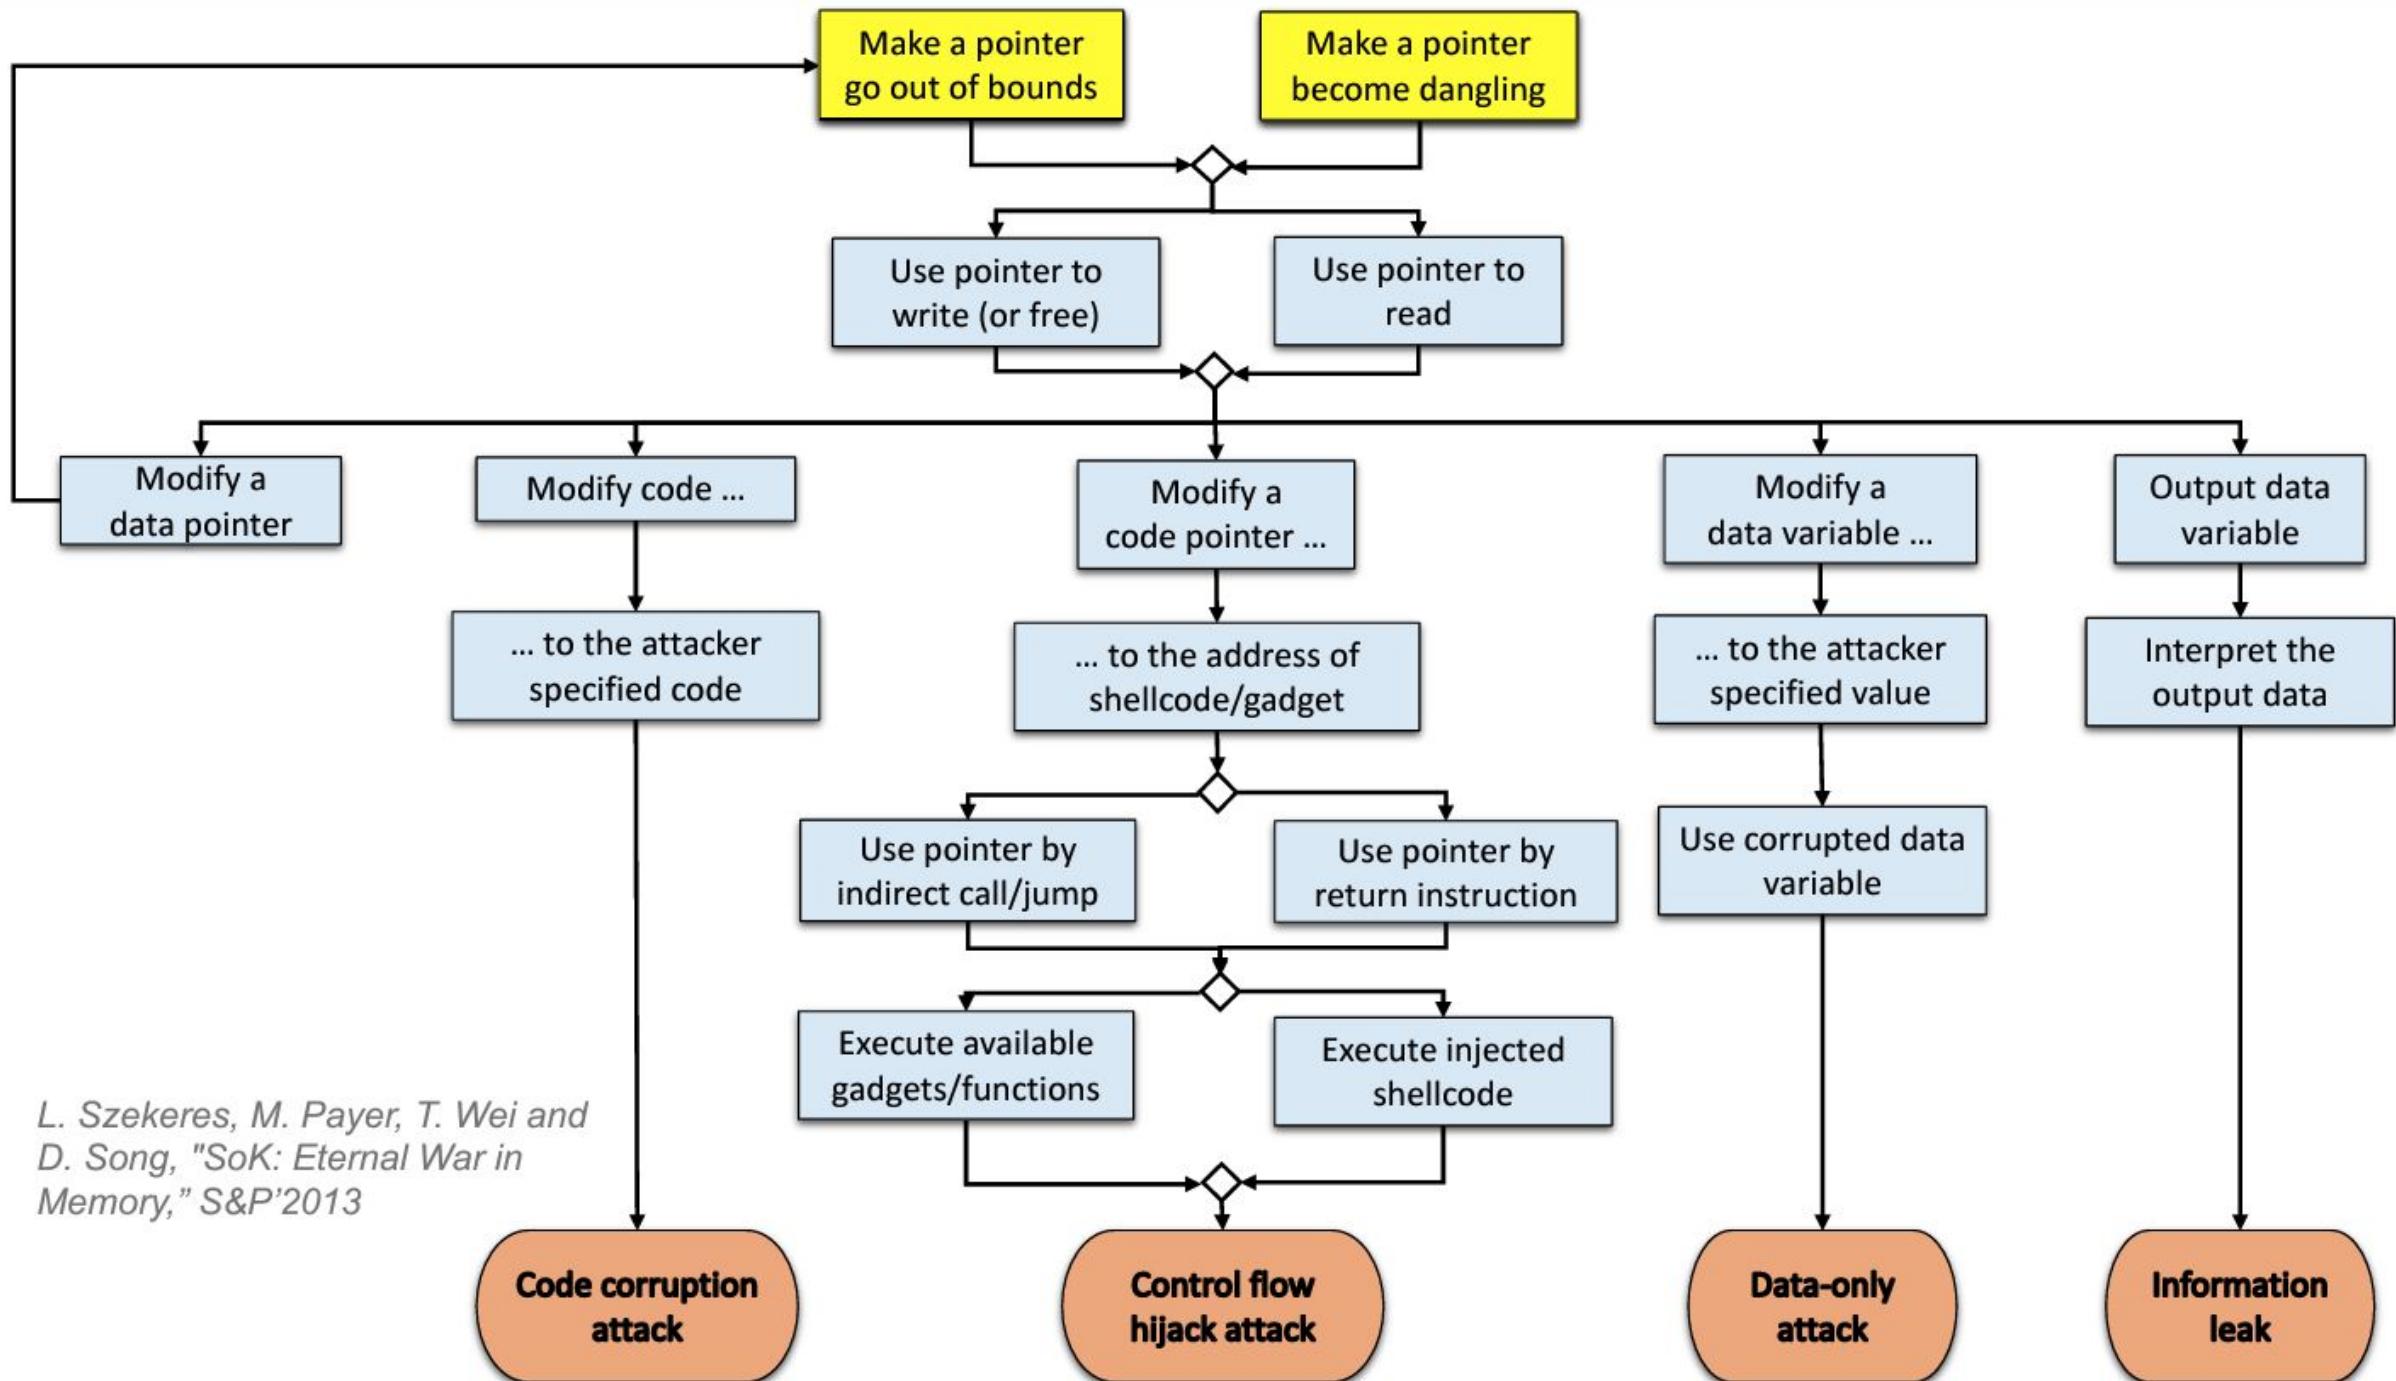
- Range of solutions
  - Finding and fixing the bug
  - Software solutions
  - Hardware solutions



# Hardware Support for Memory Safety

- Idea: include metadata and perform security checks at runtime
  - Spatial safety (bound information)
  - Temporal safety (allocation/de-allocation information)





L. Szekeres, M. Payer, T. Wei and  
D. Song, "SoK: Eternal War in  
Memory," S&P'2013

# Main Ideas

- Control-flow integrity (CFI) → check all transitions and allow only those that are valid.
- Two main questions:
  - How to find/know all the valid transitions? (why)
  - How to check and enforce this in runtime?



```

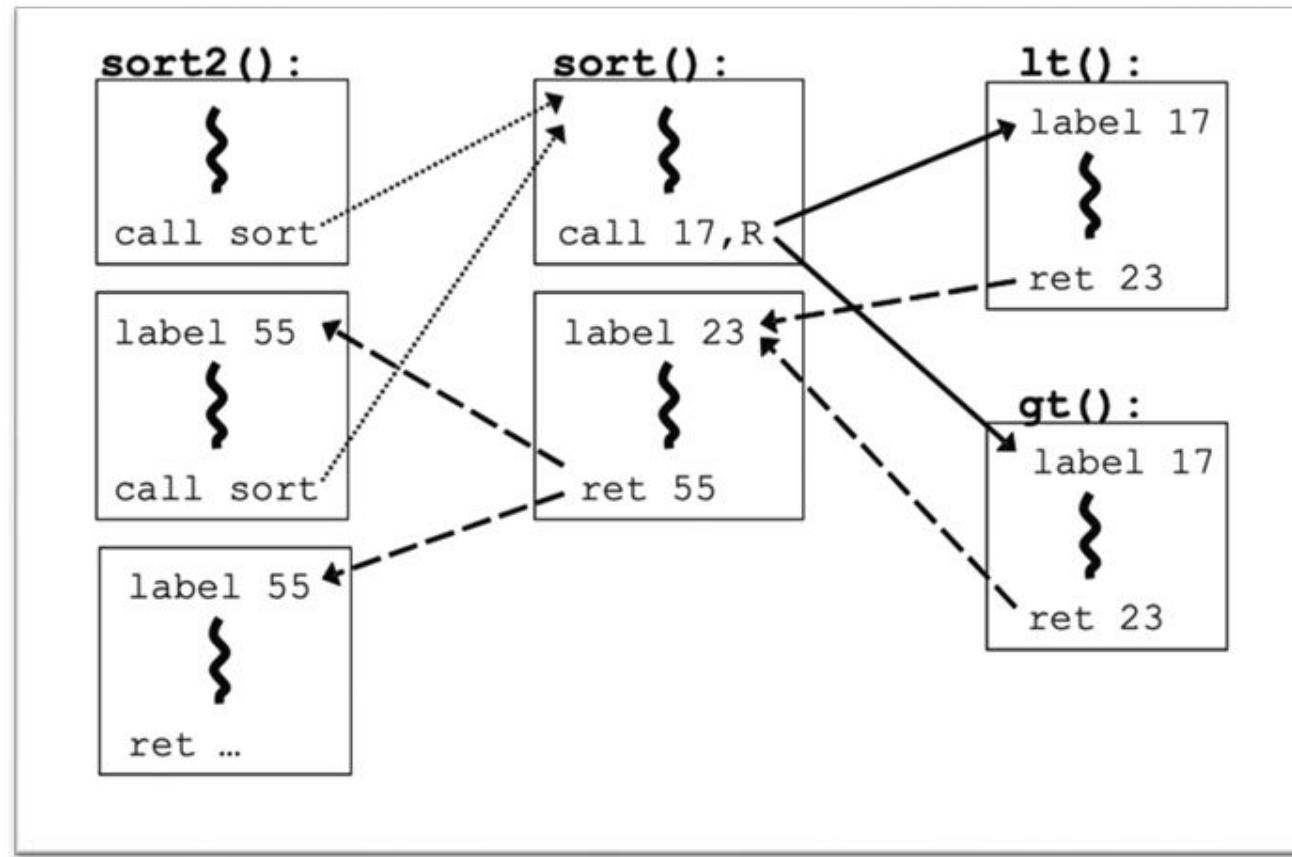
bool lt(int x, int y) {
    return x < y;
}

bool gt(int x, int y) {
    return x > y;
}

sort2(int a[], int b[], int len)
{
    sort( a, len, lt );
    sort( b, len, gt );
}

sort(int x[], int len, fun_ptr)
{
    for(int i=0; ....)
        for (int j=i; ....)
            if (fun_ptr(x[i], x[j]))
                ... //swap x[i] and x[j]
}

```



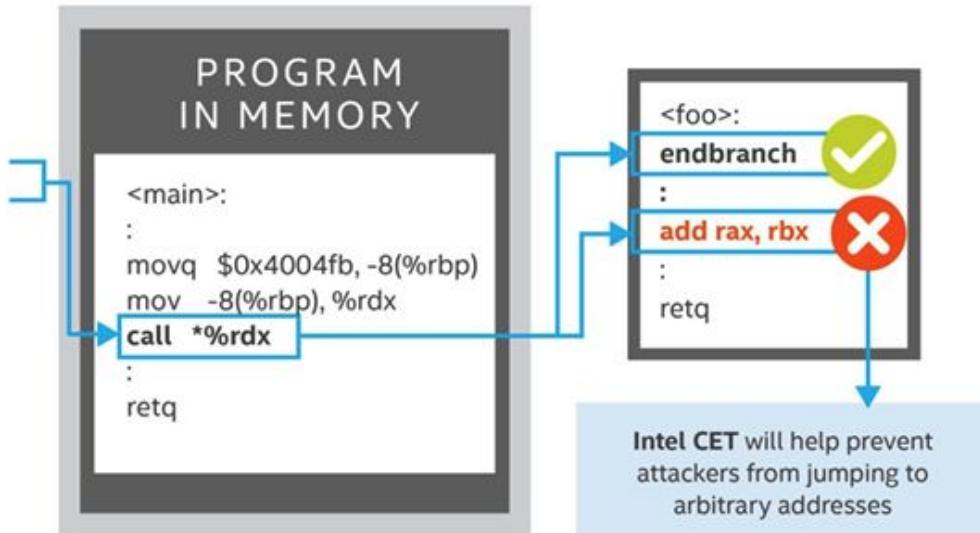
*Control-Flow Integrity Principles, Implementations, and Applications;  
Mart'ın Abadi, et al. CCS'05*

# Intel® Control-Flow Enforcement Technology (Intel CET)

$$\text{INTEL CET} = \text{INDIRECT BRANCH TRACKING (IBT)} + \text{SHADOW STACK (SS)}$$

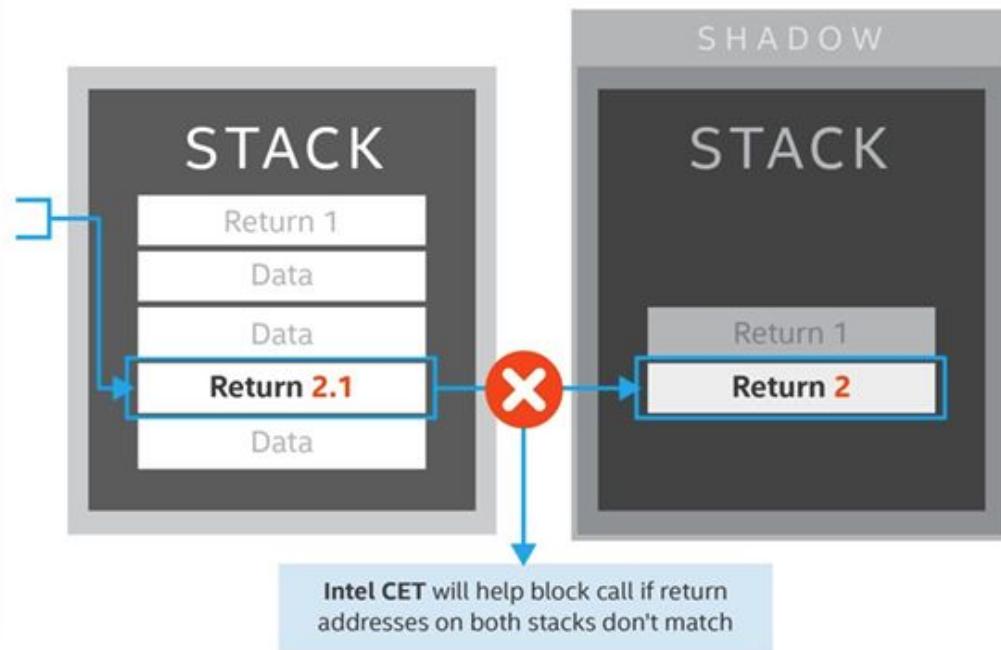
## INDIRECT BRANCH TRACKING (IBT)

IBT delivers indirect branch protection to defend against jump/call oriented programming (JOP/COP) attack methods.

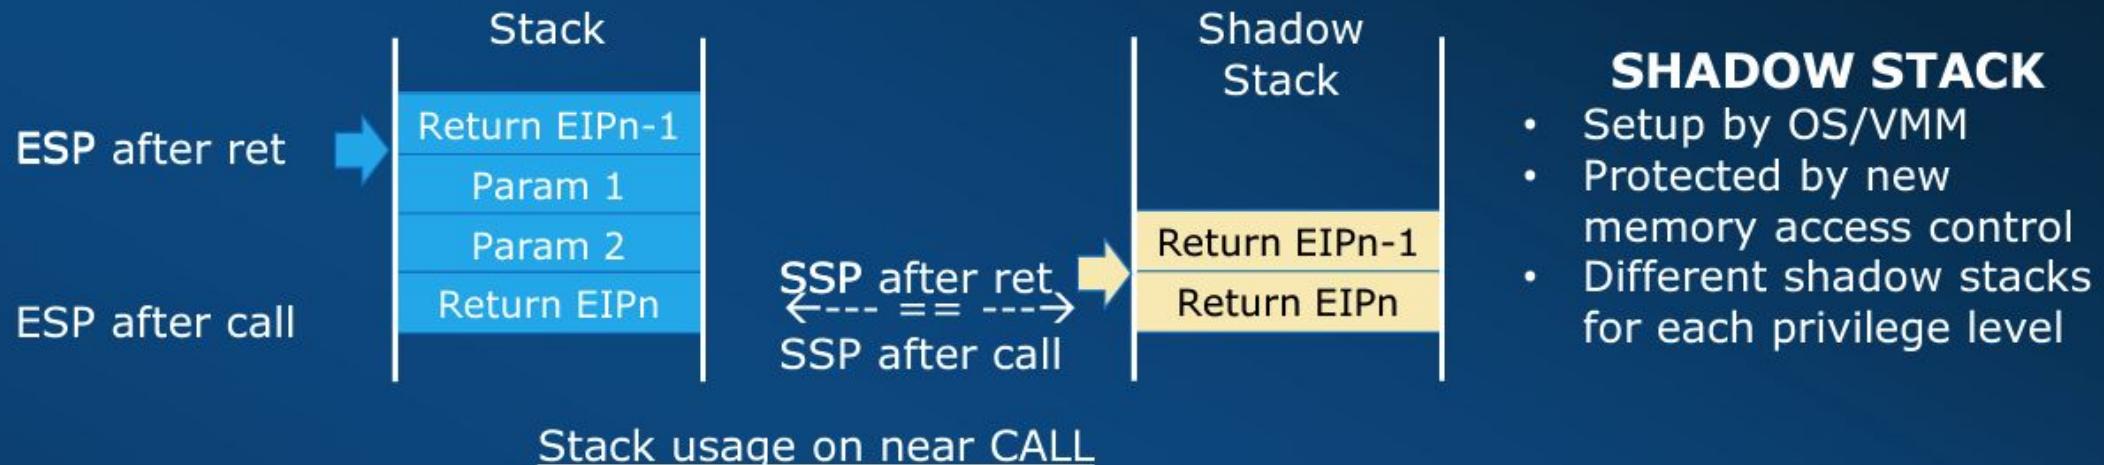


## SHADOW STACK (SS)

SS delivers return address protection to defend against return-oriented programming (ROP) attack methods.



# Shadow Stack Operation

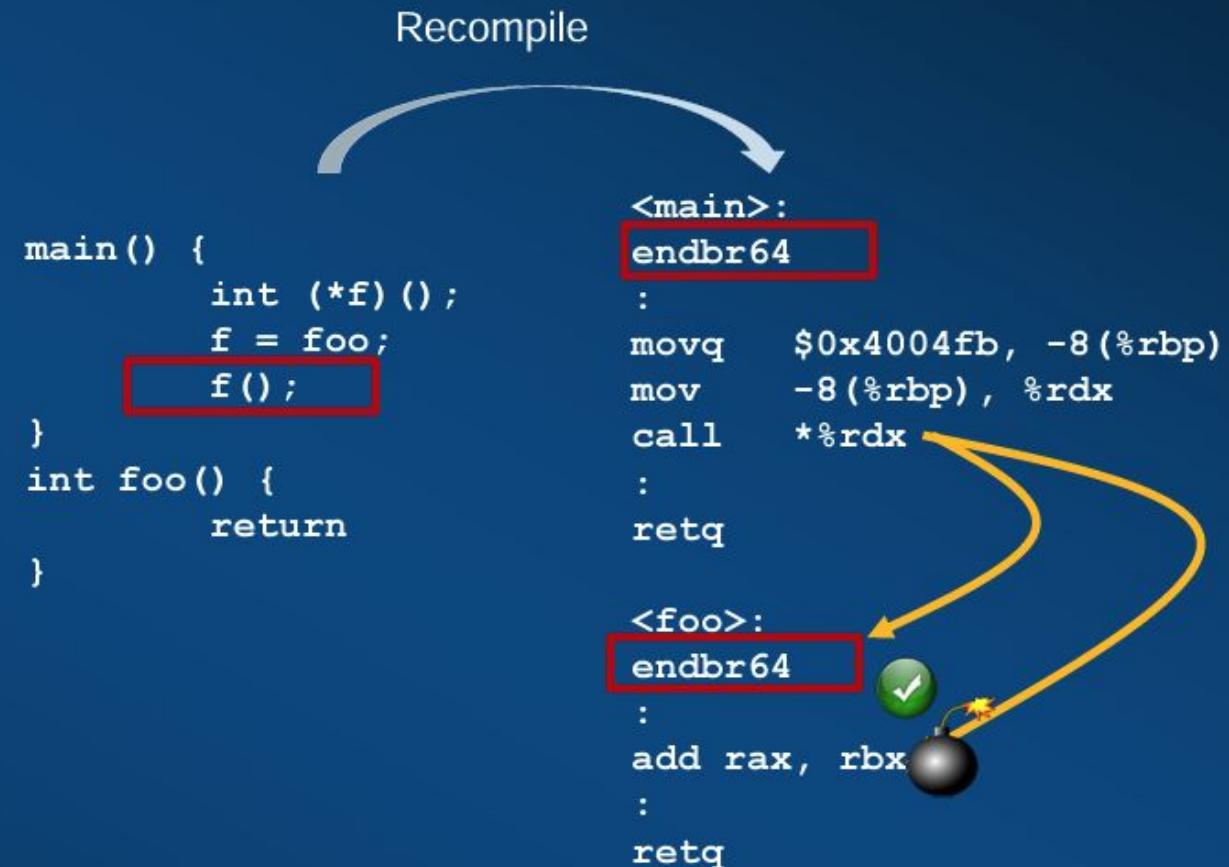


- Call
  - pushes return address on both stacks
- No parameters passing on shadow stack
- Return
  - pops return address from both stacks
  - Controlflow Protection (#CP) exception in case the two return addresses don't match

Keeps stack ABI intact – no changes to data stack layout

# ENDBRANCH

- New Instruction to mark legal targets of indirect jumps
- Added by the compiler
- Decodes as “NOP” on legacy processors
- An indirect jump to a target not marked by ENDBR signals an exception



# Main Ideas

- Control-flow integrity (CFI) → check all transitions and allow only those that are valid.
- Two main questions:
  - How to find/know all the valid transitions? (why)
  - How to check and enforce this in runtime?



# Main Ideas

- Memory tagging → Augment every pointer with its tag. Only legal instructions can update the tag!

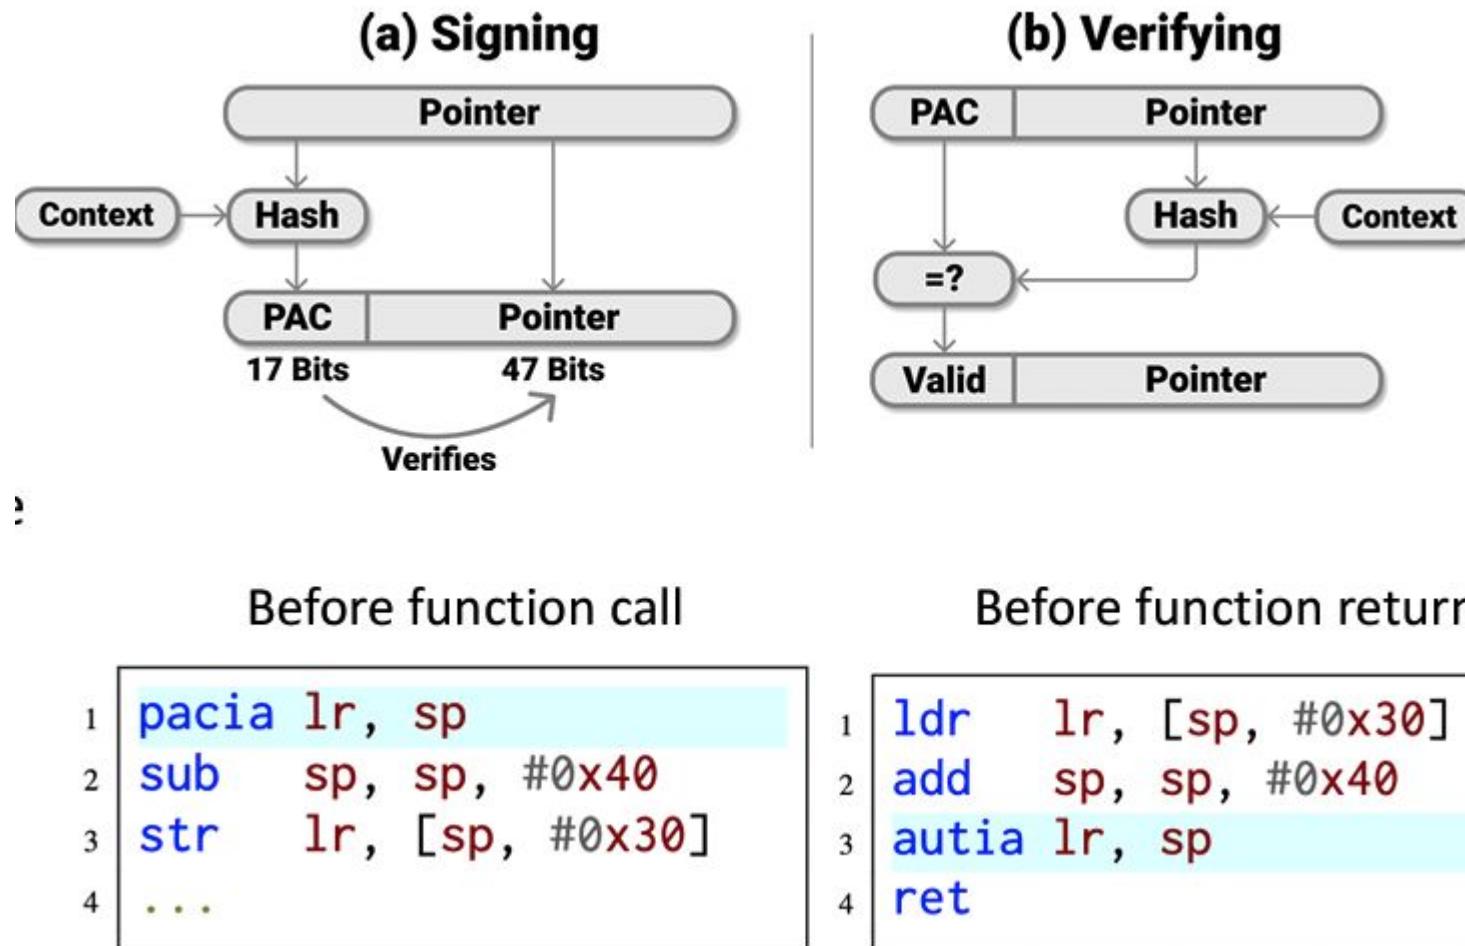


# Main Ideas

- Memory tagging → Augment every pointer with its tag. Only legal instructions can update the tag!
- Further expand this to every memory content → data, pointer, and code



# ARM PAC



# Challenges

- How to store tags? How to move them? How much is the slowdown?
- What MAC to use? What to do if they don't match?
- Does this always work?



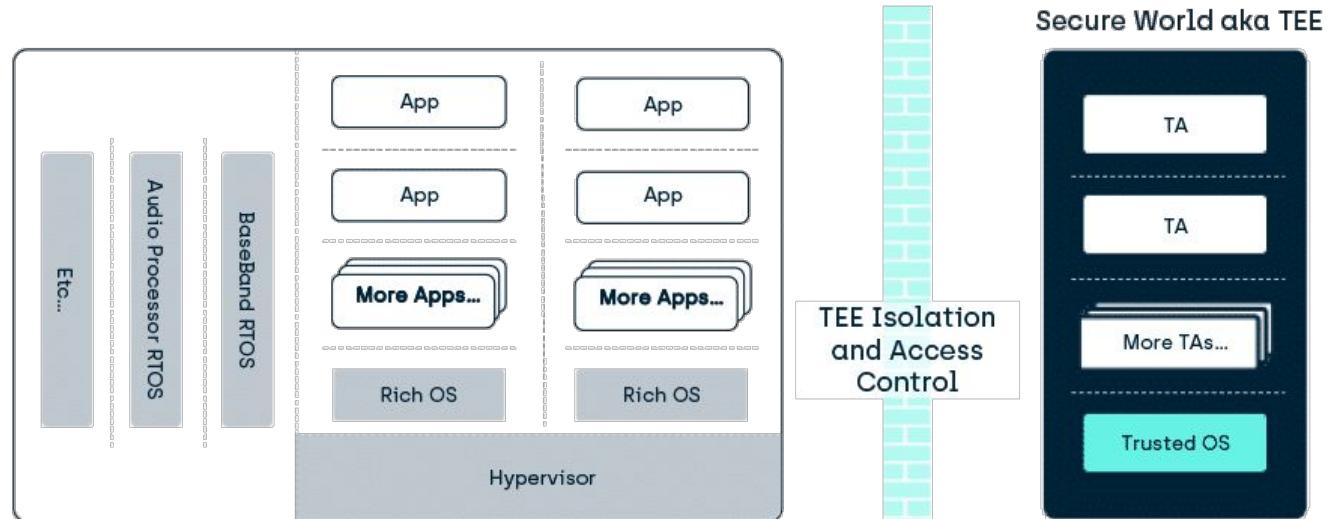
# Trusted Execution Environment



# TEE

- Why do we need them?

## 1. Taking isolation to next level



Note that this is not just isolating activity on the application cores  
TrustZone also isolates the Secure World (TEE) from all the other bus masters!

Isolation controlled ONLY  
by the TEE designers.  
Designers (and hackers) of  
the rest of the system  
cannot override this!

From  
<https://www.trustonic.com/technical-articles/what-is-a-trusted-execution-environment-tee/>

# Why we need more isolation?

- Larger the code, harder to remove bugs.
- Side channels.
- Different vendors/designers which we don't trust.
- More sharing.

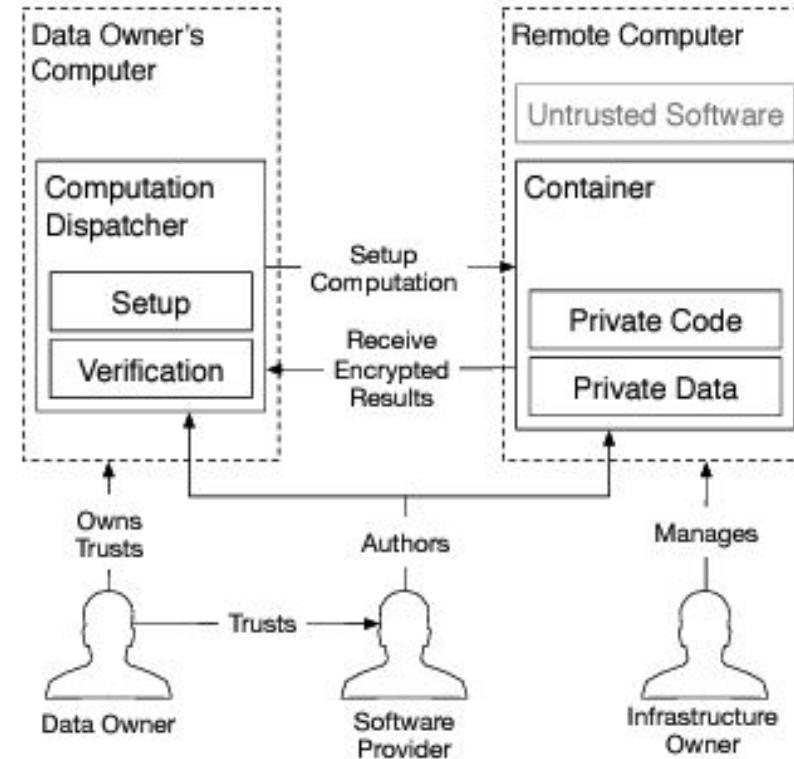


# TEE

- Why do we need them?

1. Taking isolation to next level

2. Remote computing



From SGX Explained.

# Why we need remote computing?

- We won't trust software but hardware and cryptographic methods.
- The trust issue is either due to bugs or just having a third-party and/or regulations.

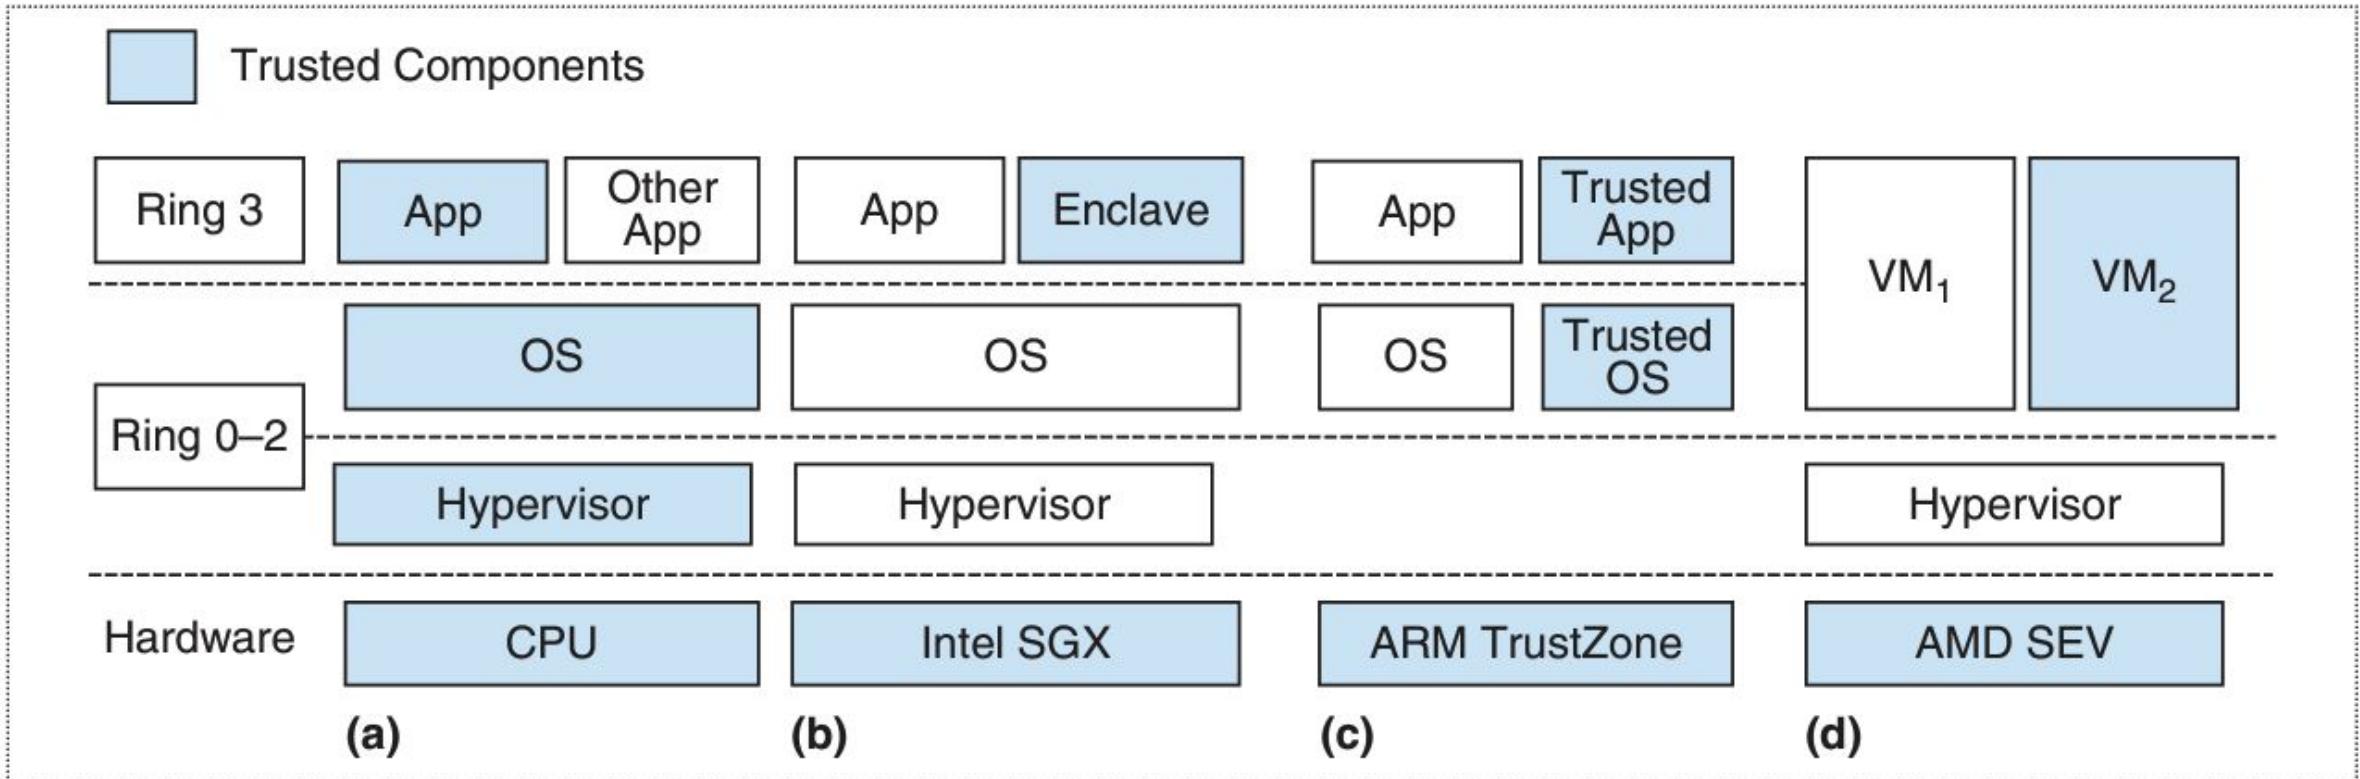


**Samueli**

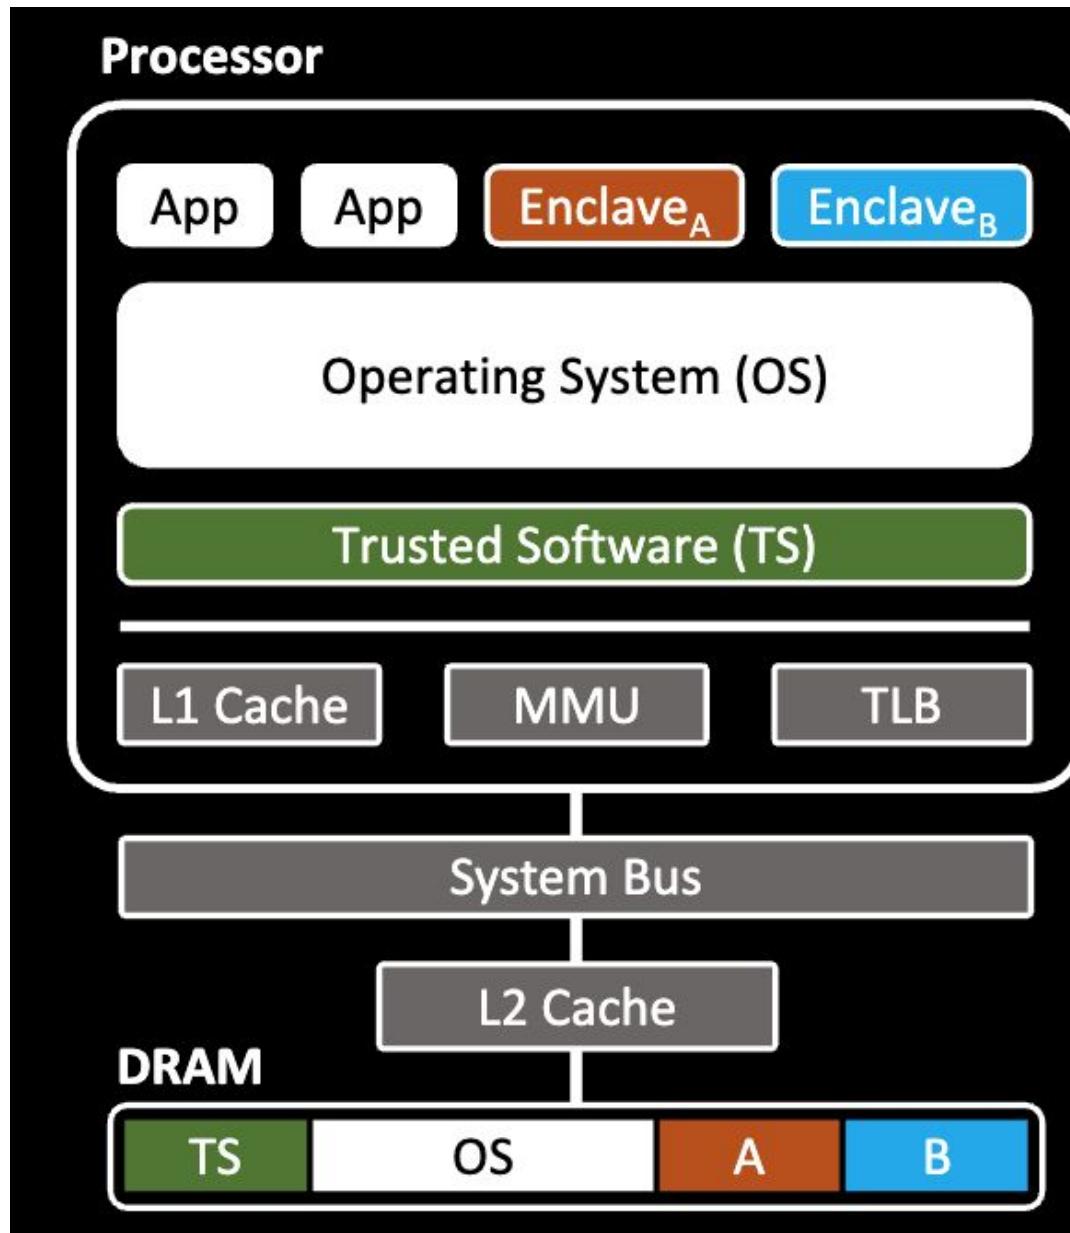
School of Engineering

ECE 209 - Spring 24  
Nader Sehatbakhsh <[nsehat@ee.ucla.edu](mailto:nsehat@ee.ucla.edu)>

# TEE Trust Boundaries



# Options



From CURE: A Security Architecture with Customizable and Resilient Enclaves

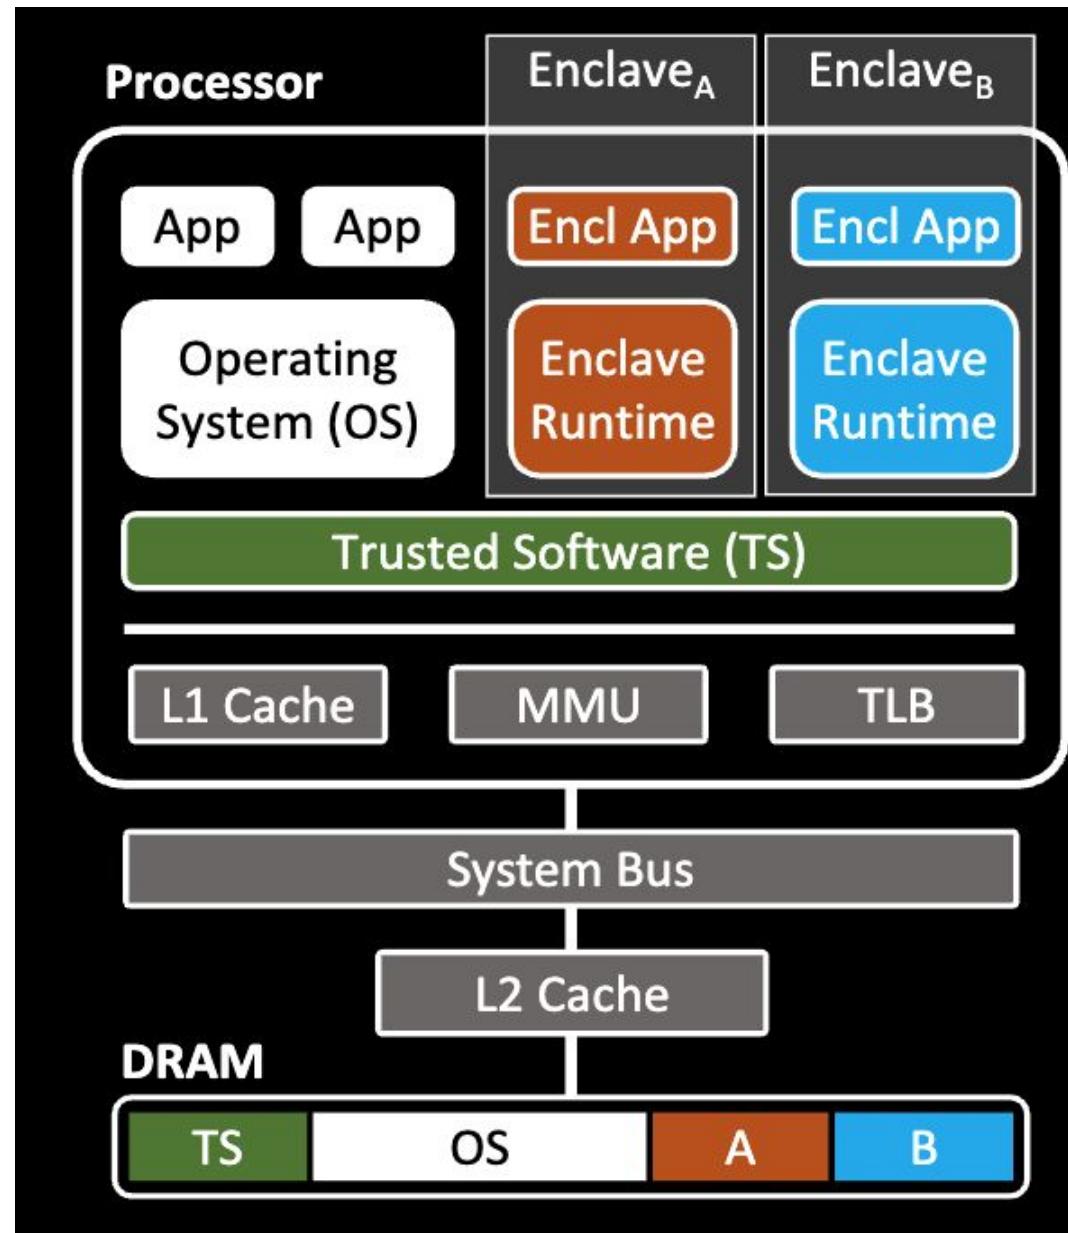


**Samueli**

School of Engineering

ECE 209 - Spring 24  
Nader Sehatbakhsh <[nsehat@ee.ucla.edu](mailto:nsehat@ee.ucla.edu)>

# Options



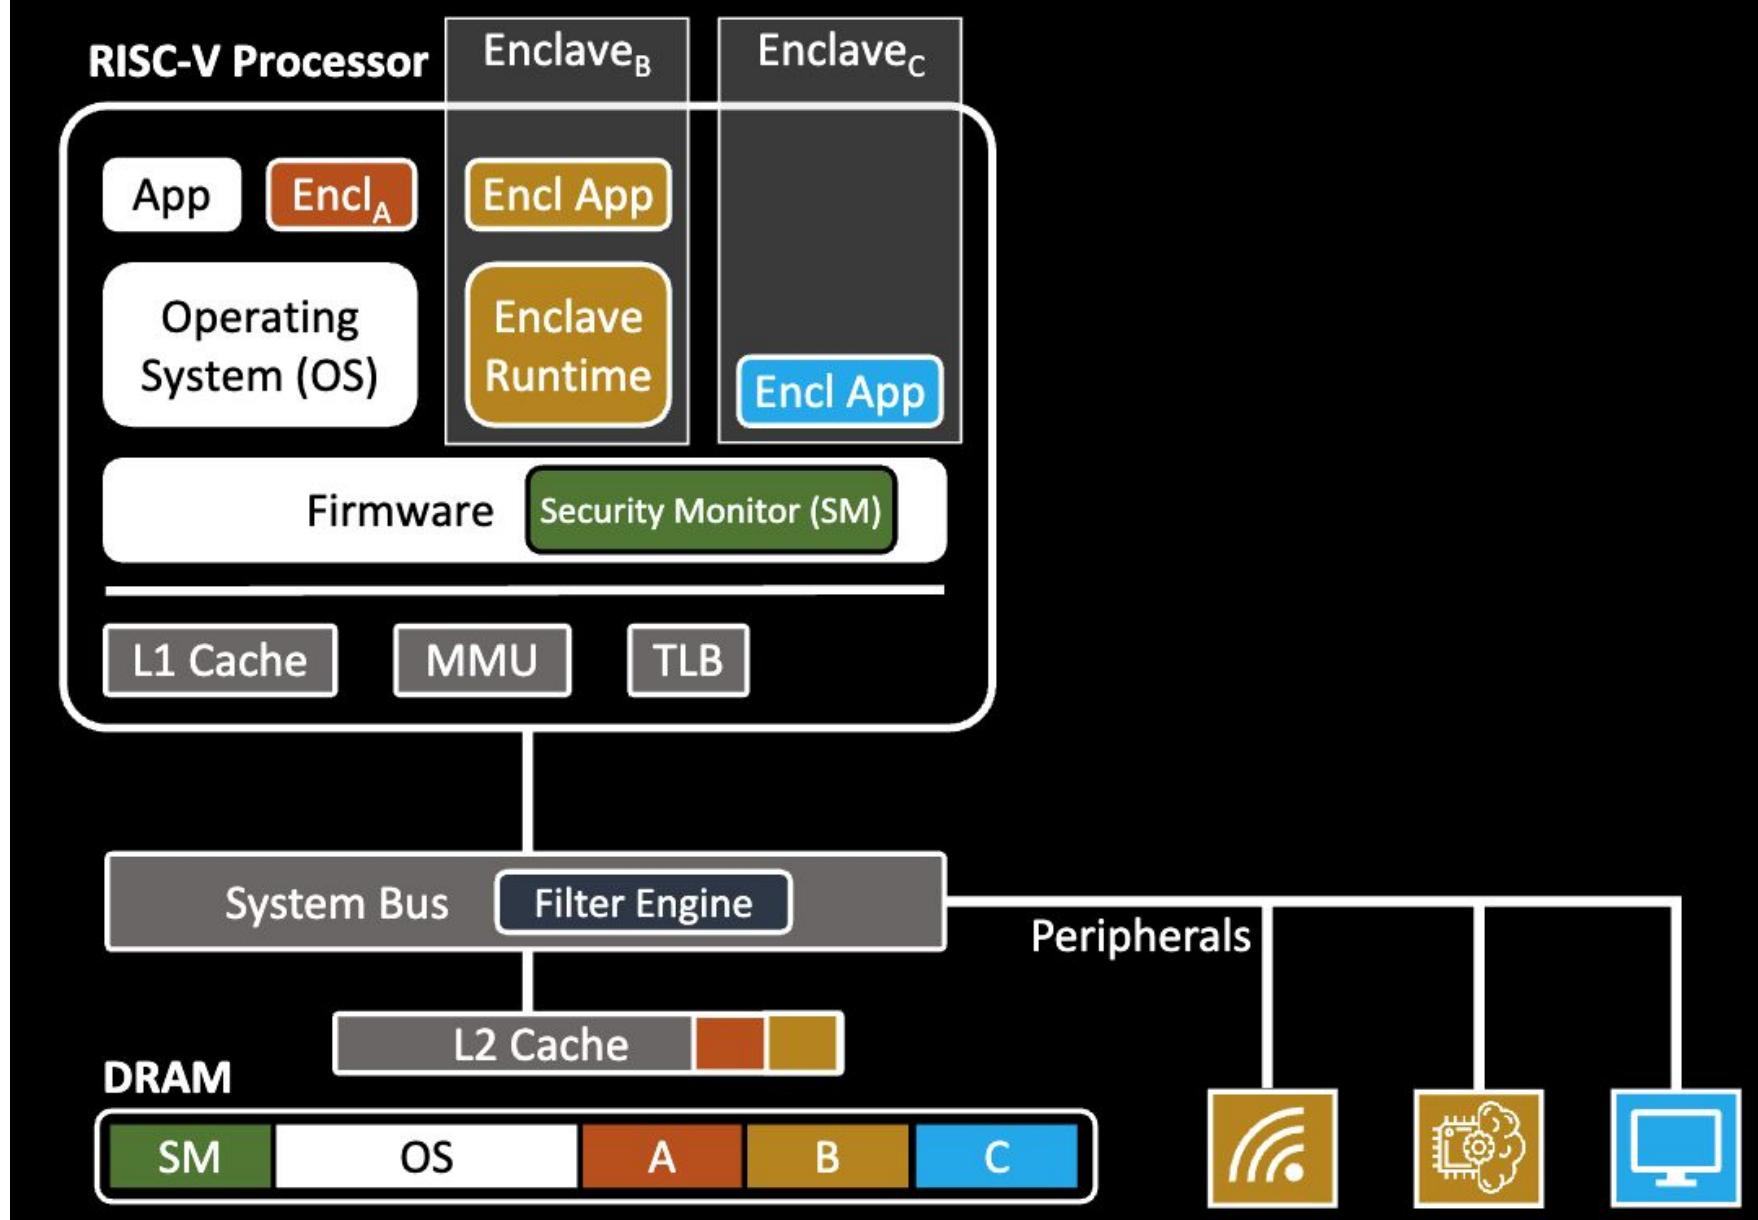
From CURE: A Security Architecture with  
Customizable and Resilient Enclaves



**Samueli**

School of Engineering

# Options



From CURE: A Security Architecture with Customizable and Resilient Enclaves



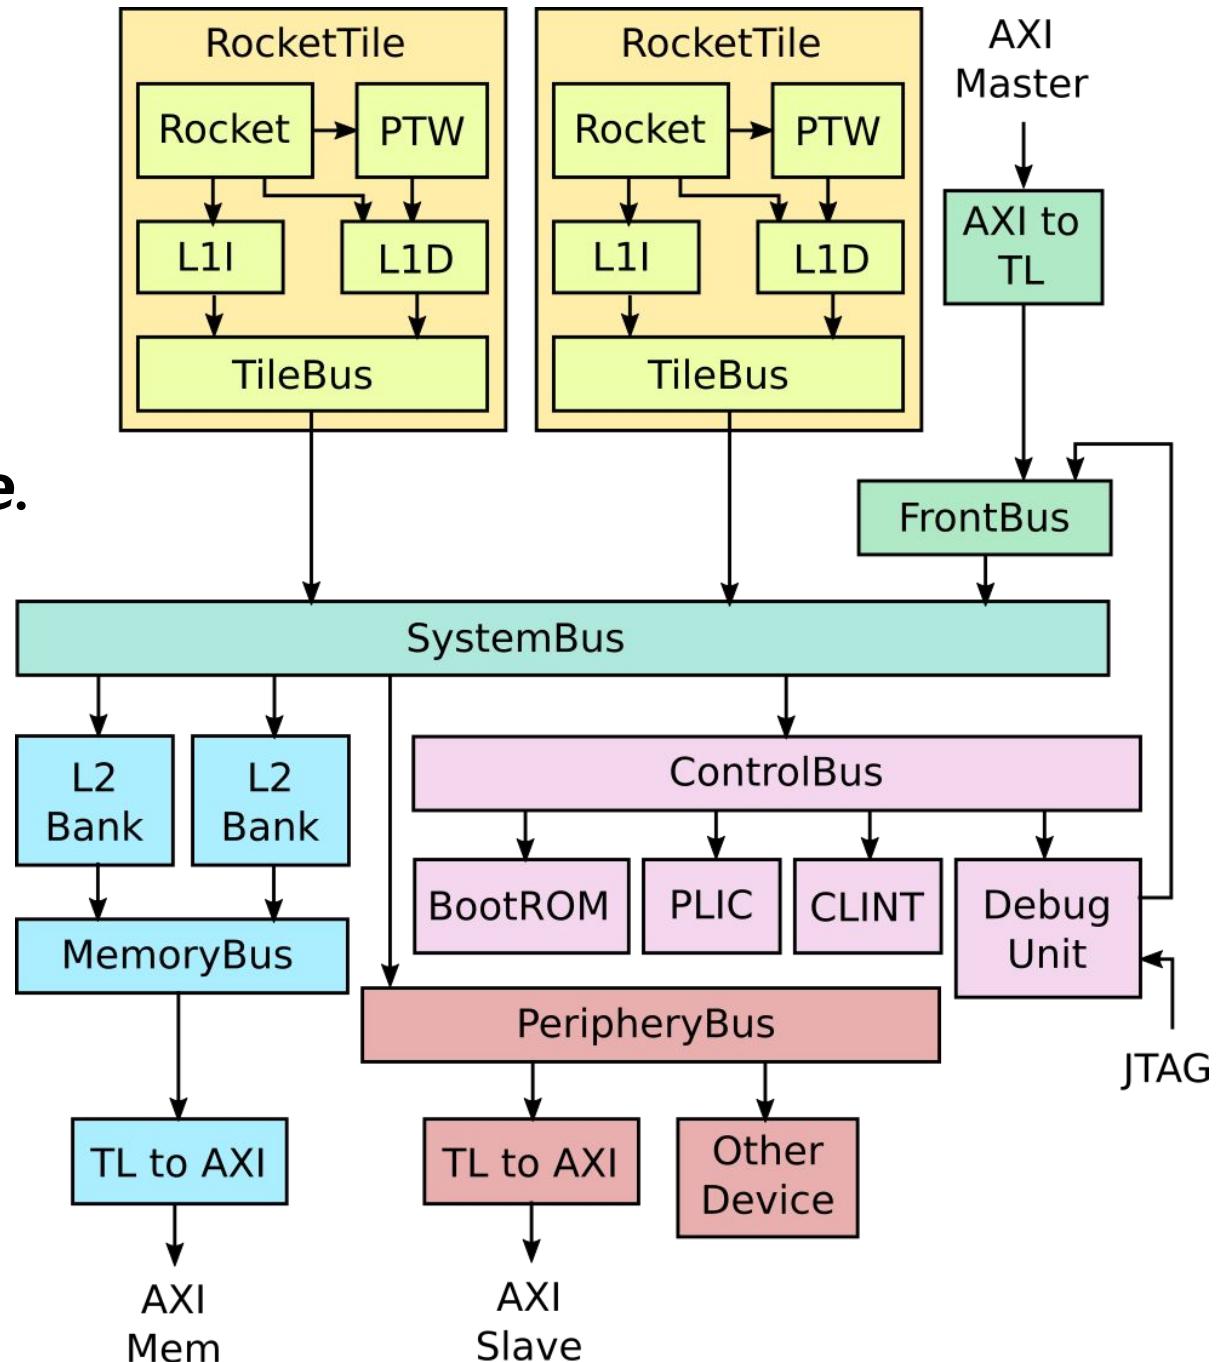
**Samueli**

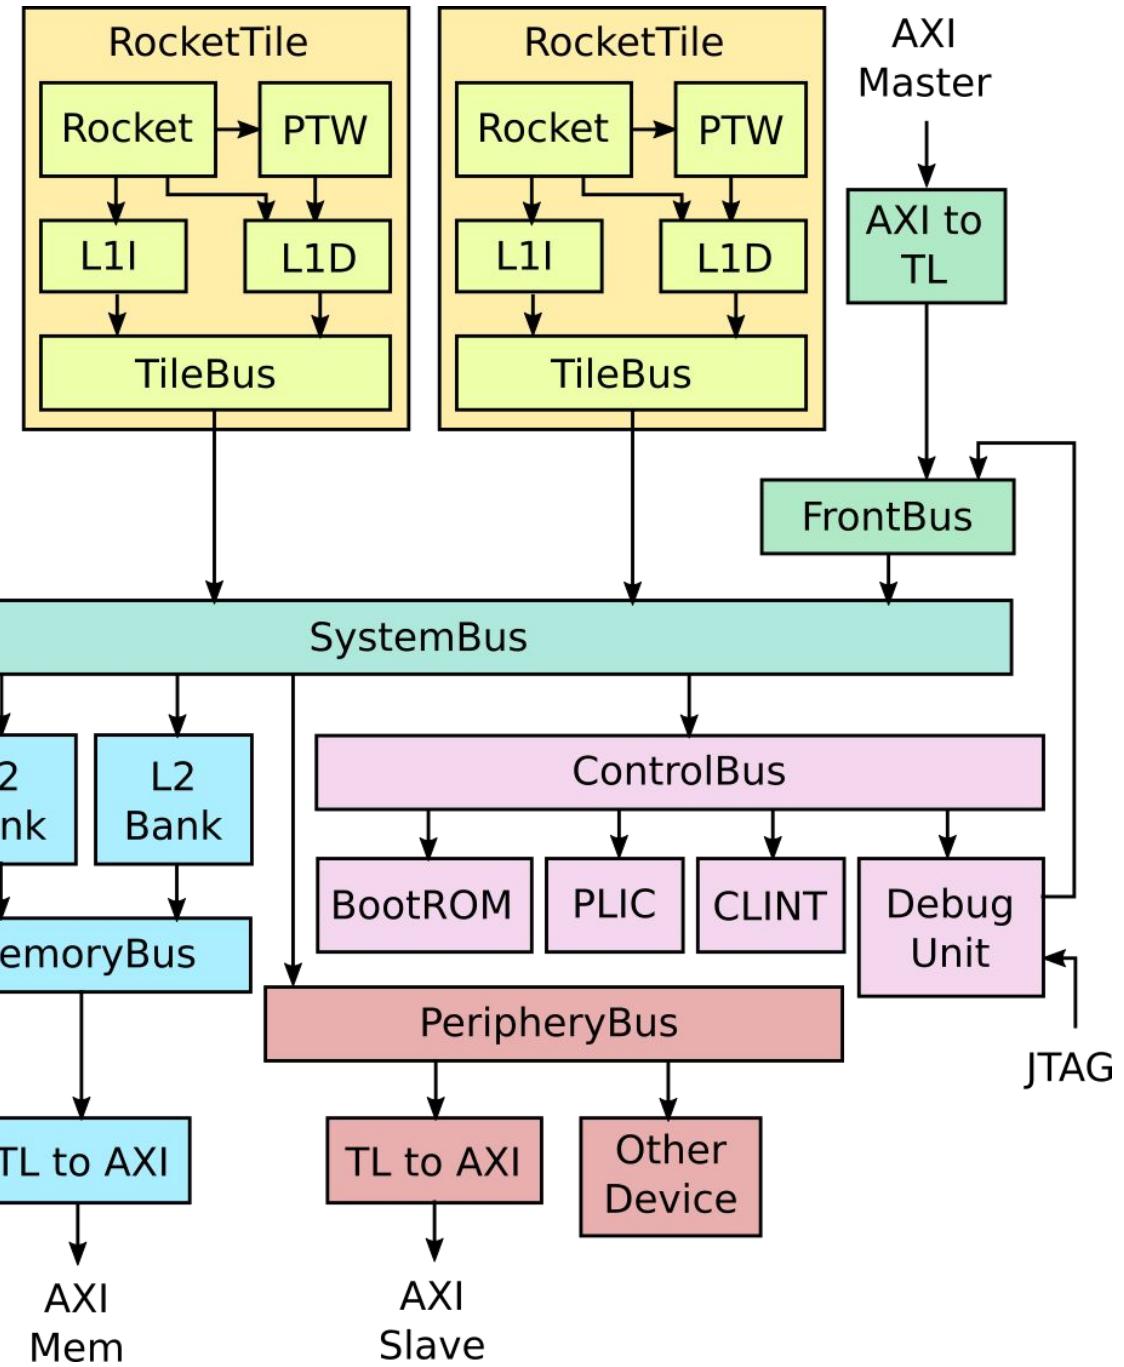
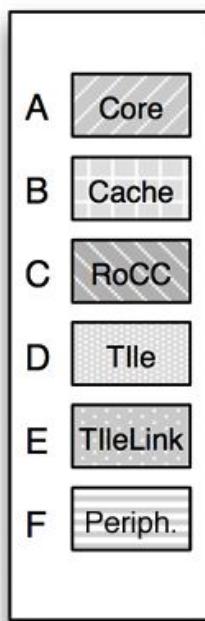
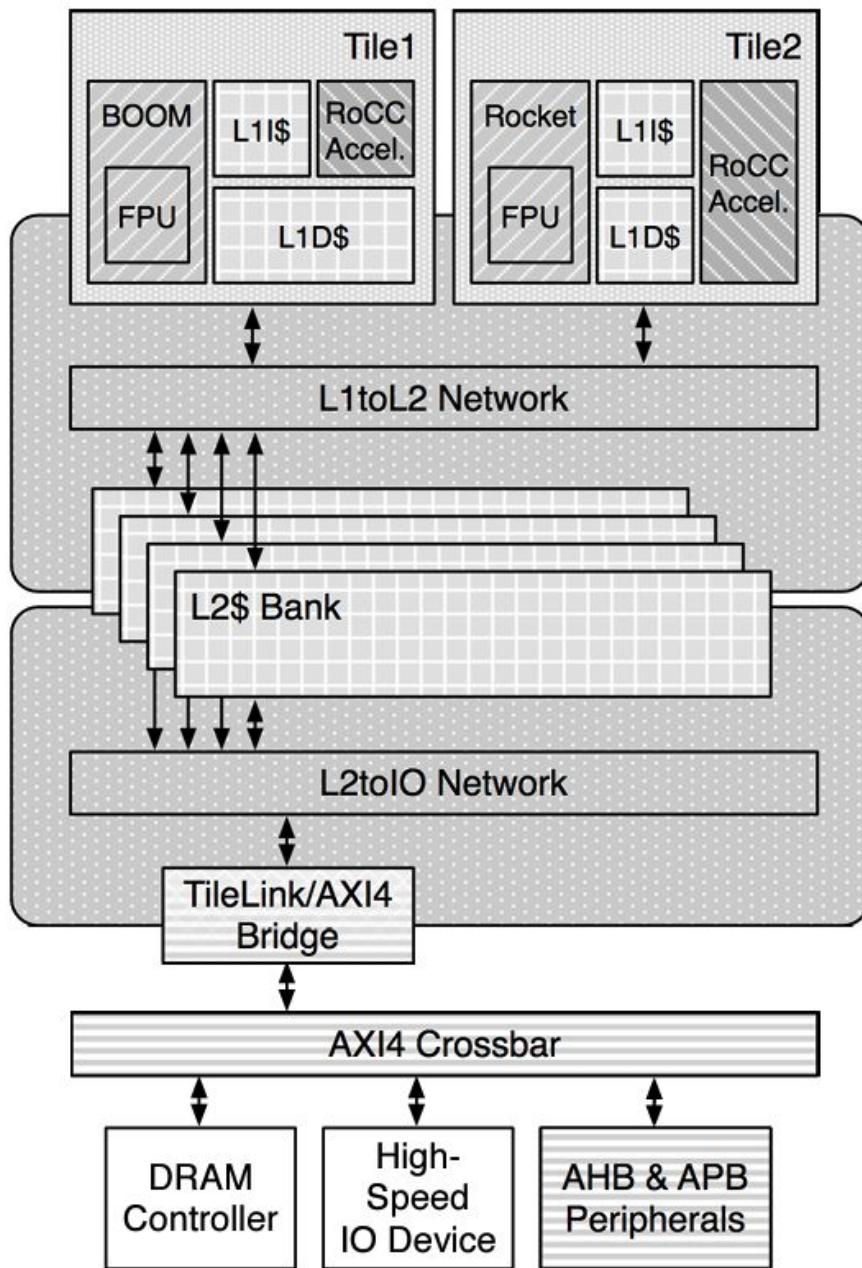
School of Engineering

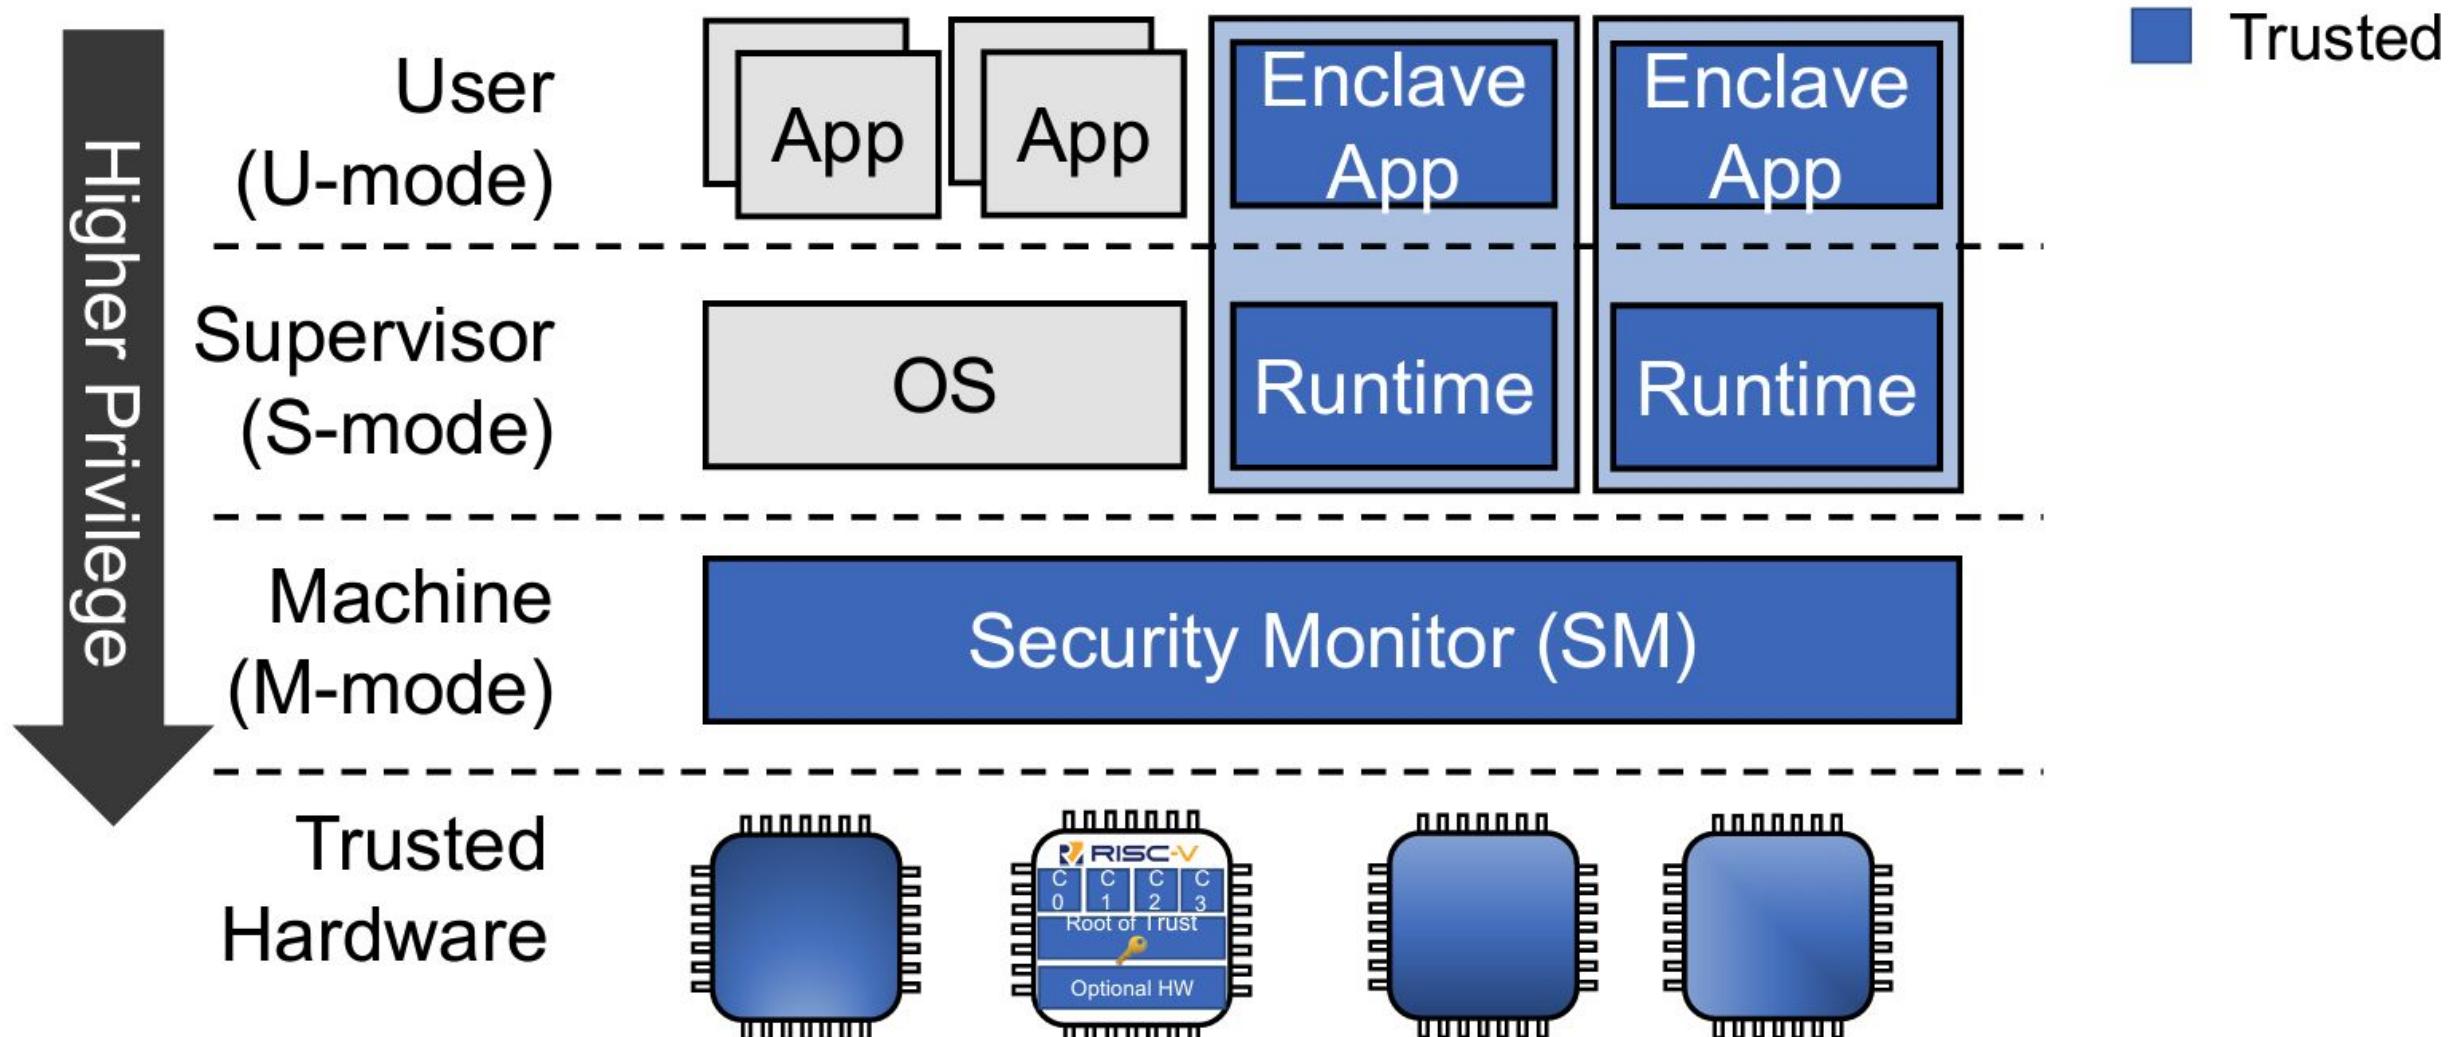
ECE 209 - Spring 24  
Nader Sehatbakhsh <[nsehat@ee.ucla.edu](mailto:nsehat@ee.ucla.edu)>

# Microarch

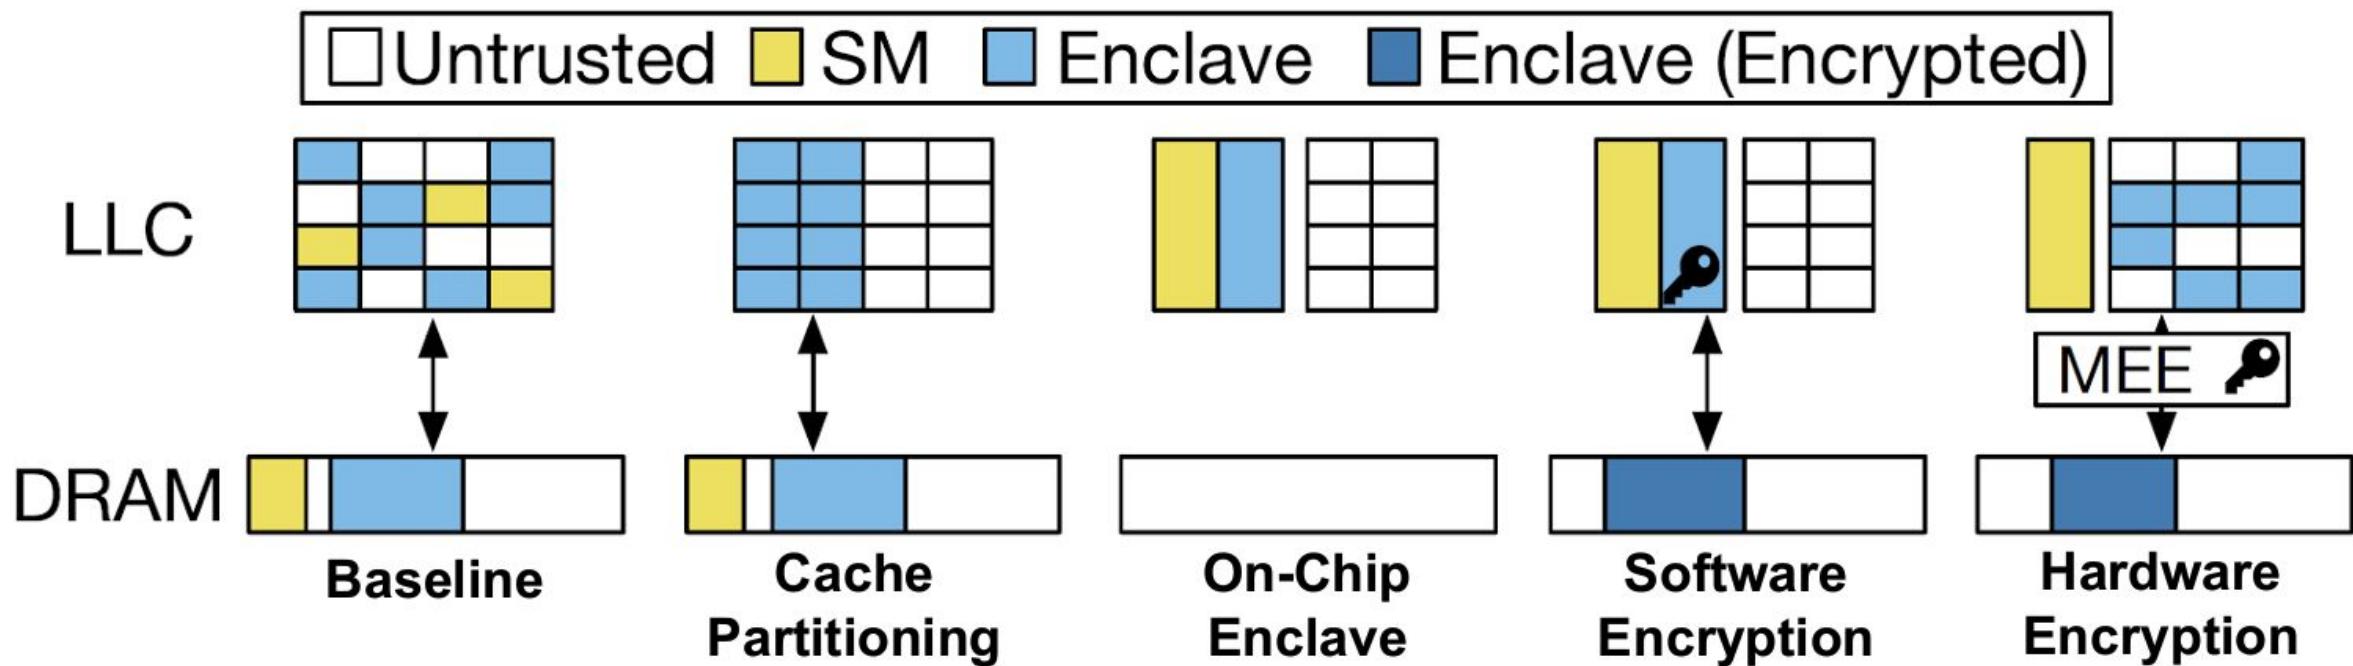
- How to architect a TEE?
- We are focusing on **Keystone**.







# How to protect memory?



# What are the challenges?

- First, is this really secure?
  - Not really, there are still many side channel attacks because TEE is not truly isolated!



# What are the challenges?

- First, is this really secure?
  - Not really, there are still many side channel attacks because TEE is not truly isolated!
- Second, how to balance performance and security?
  - Challenges: true isolation reduces utilization, parallelism, etc.

# What are the challenges?

- First, is this really secure?
  - Not really, there are still many side channel attacks because TEE is not truly isolated!
- Second, how to balance performance and security?
  - Challenges: true isolation reduces utilization, parallelism, etc.
- Third, how to handle security for other components?

# Summary

- In addition to hardware concerns, side channel, transient, and faults, computing systems are vulnerable to software attacks due to bugs in software.
- There are two main solutions for software vulnerabilities:
  - Either fix them using a plethora of solutions (hardware, software, etc.)
  - Or use TEEs to isolate the environment.
    - TEEs have additional bonuses including remote computing and/or protection against faults, side-channel, and/or transient attacks.

# End of Presentation



# Acknowledgement

- This course is partly inspired by the following courses created by my colleagues:
  - 6.888 MIT (Yan)
  - CS 598 UIUC (Fletcher)
  - CS 7290 Georgia Tech (Kim)
  - ECE 752 Wisconsin (Lipasti)
  - ECE 7103A Georgia Tech (Qureshi)

