

# Comp 590-184: Hardware Security and Side-Channels

## Lecture 21: Hardware Support for Memory Safety

April 16, 2026  
Andrew Kwong

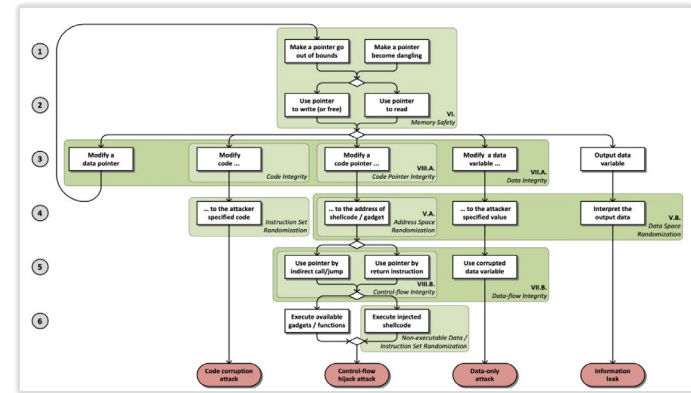


THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL

Slides adapted from  
Mengjia Yan ([shd.mit.edu](http://shd.mit.edu))

# Overview

- What can attackers do with software bugs?
  - real-world examples
- Hardware mitigations: what are the design tradeoffs?



# The Problem: Software Bugs

- Low-level Language Basics (C/C++/Assembly)
  - +Efficient, programmers have more control
  - -Programming productivity
  - -Memory Safety Bugs
    - 70% of Chrome's vulnerabilities
- Memory safe languages
  - Rust's introduction caused memory safety bugs to drop from 76% to 35% of Android vulnerabilities
- Low-level code still widely used in production systems and legacy systems
  - Operating systems, web browsers, etc.
  - Large number of CVEs every year



## The Problems of Using Pointers

---

- Pointer = Address of variables:
  - An 64-bit integer that is an index into memory, the location where a variable is stored
- **It is programmers' responsibility to do pointer check**, e.g. NULL, out-of-bound, use-after-free
- Why do Python (and other high-level programming language) not have these problems?
  - out-of-bound access => emit runtime checks
  - use-after-free => garbage collection

## Pointer example

---

```
char * buffer1 = (char*)malloc(buffer1_size);
char * buffer2 = (char*)malloc(buffer2_size);

// fill up buffer2 with 'B'
for (int i=0; i<buffer2_size; i++){
    buffer2[i] = 'B';
}

// fill up buffer1 with 'A'
for (int i=0; i<2*buffer1_size ; i++){
    buffer1[i] = 'A';
}
```

## Another Example

---


```
char * buffer1  = (char*)malloc(buffer1_size);
char * buffer2  = (char*)malloc(buffer2_size);
// fill buffer1 and buffer2 with 'A' and 'B' respectively

free (buffer1);
char *buffer3 = (char*)malloc(buffer3_size);
for (int i=0; i<buffer3_size; i++){
    buffer3[i] = 'C';
}
// print buffer1
```

# Memory Corruption Vulnerabilities

---

- Spatial safety:
  - out-of-bounds
  - Can happen on heap and stack
  
- Temporal safety:
  - Use-after-free
  - Use before initialization





From software bugs to attacks?

# Recap Stack Operations

```
int hello(){
    int a = 100;
    return a;
}

int main(){
    int a;
    int b = -3;
    int c = 12345;
    int *p = &b;
    int d = hello();
    ra → d = d+3;
    return 0;
}
```



# Stack Smash

---

```
int func (char *str) {  
    char buffer[12];  
    strcpy(buffer, str);  
    return 1;  
}  
  
int main() {  
    ...  
    func (input);  
    ...  
}
```

ra



## Shell code:

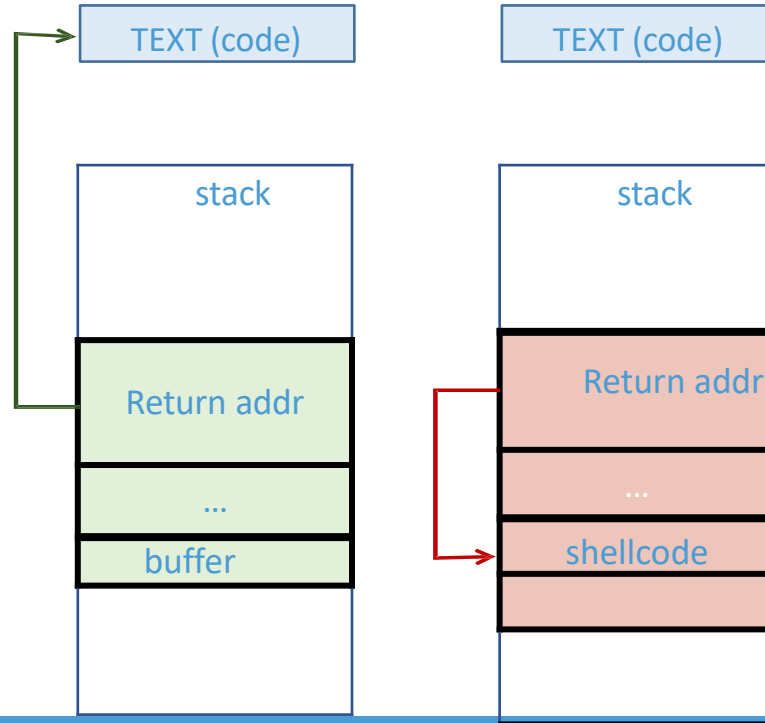
```
PUSH "/bin/sh"  
CALL system
```

## Input str:

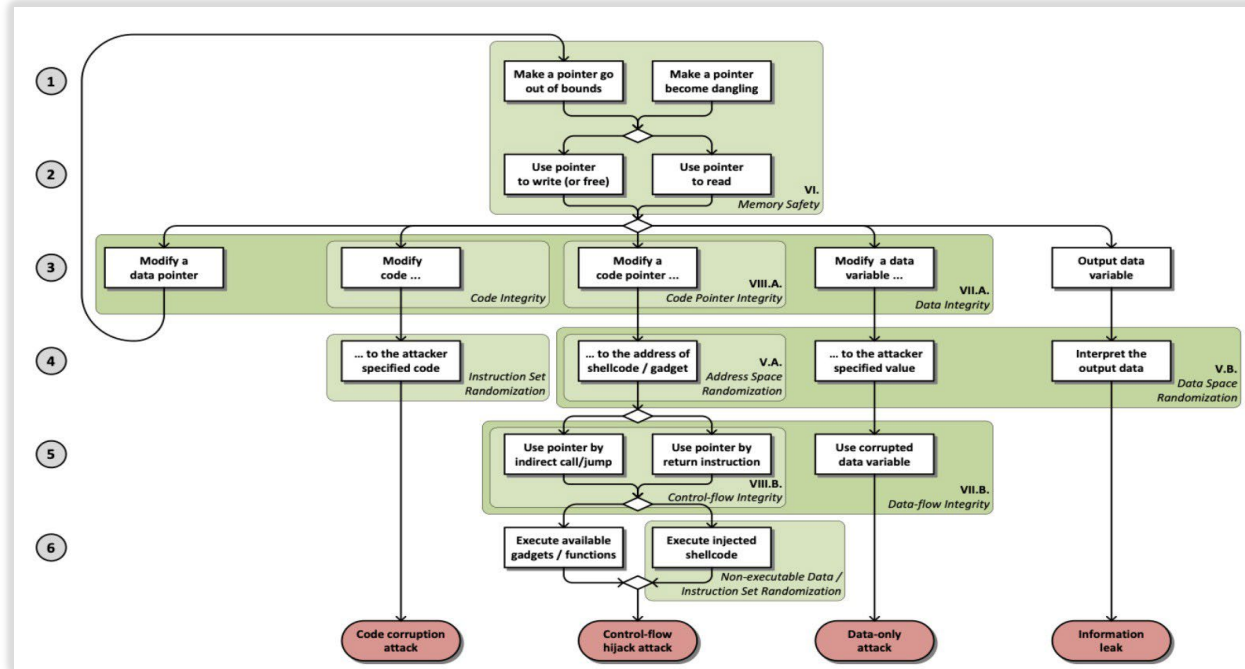
```
Shell code  
.. Some padding..  
Address of buffer
```

# Stack Smash / Code Injection Attack

---

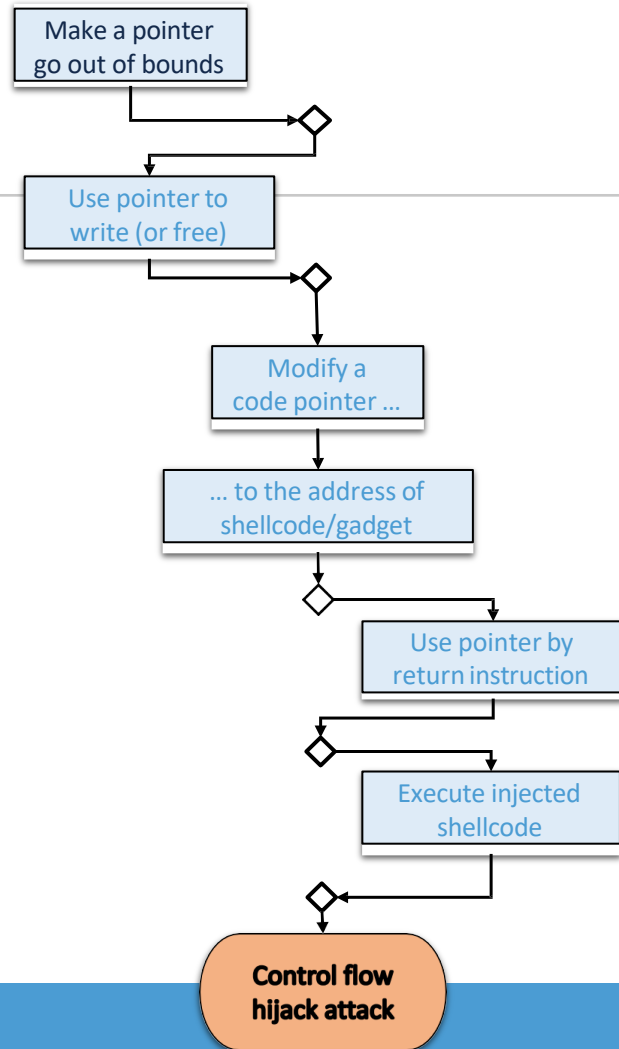


# Attack Variations

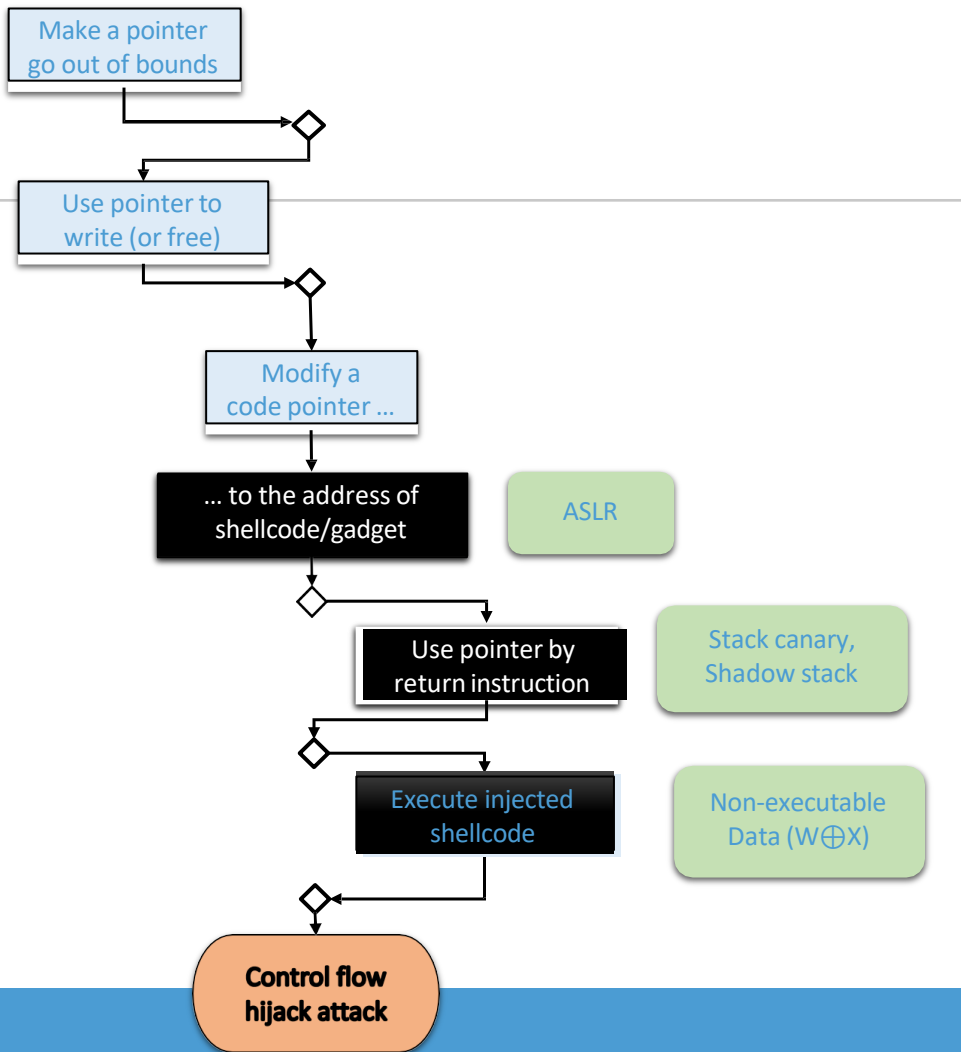
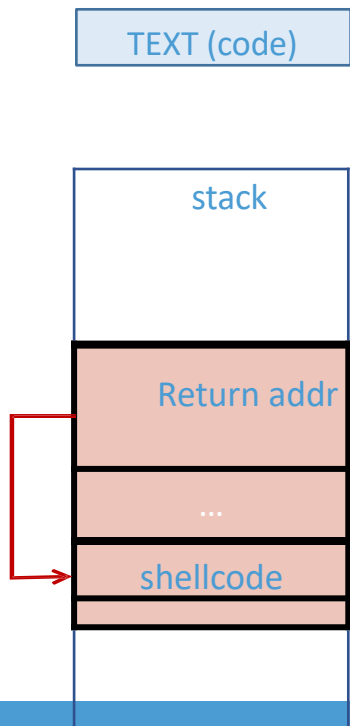


L. Szekeres, M. Payer, T. Wei and D. Song, "SoK: Eternal War in Memory," S&P'2013

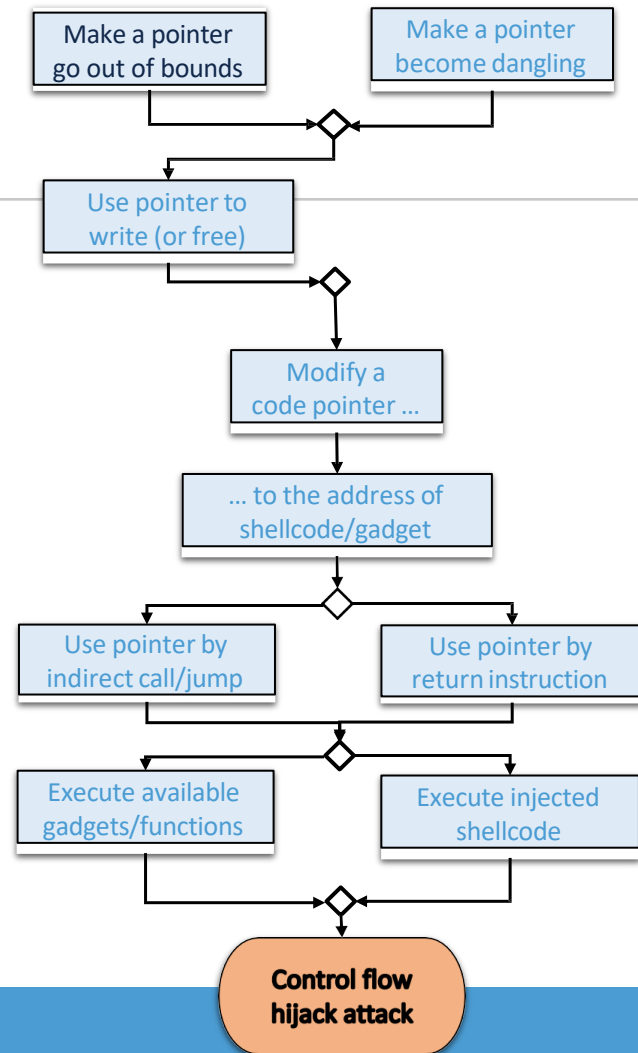
# Attack Variations



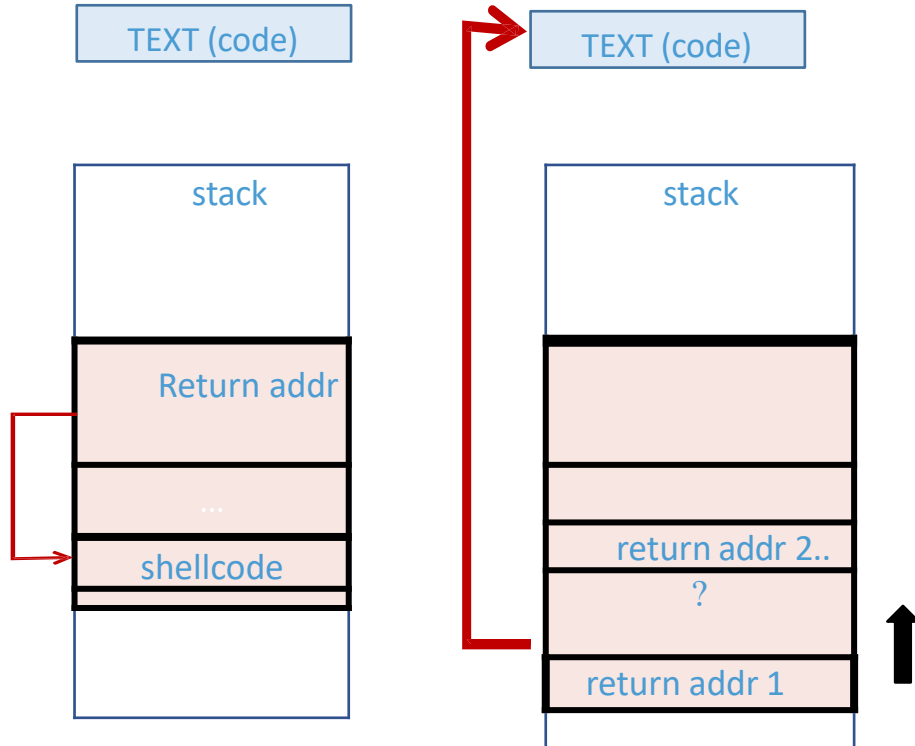
# Mitigations



# Attack Variations



# Return-Oriented Programming (ROP)

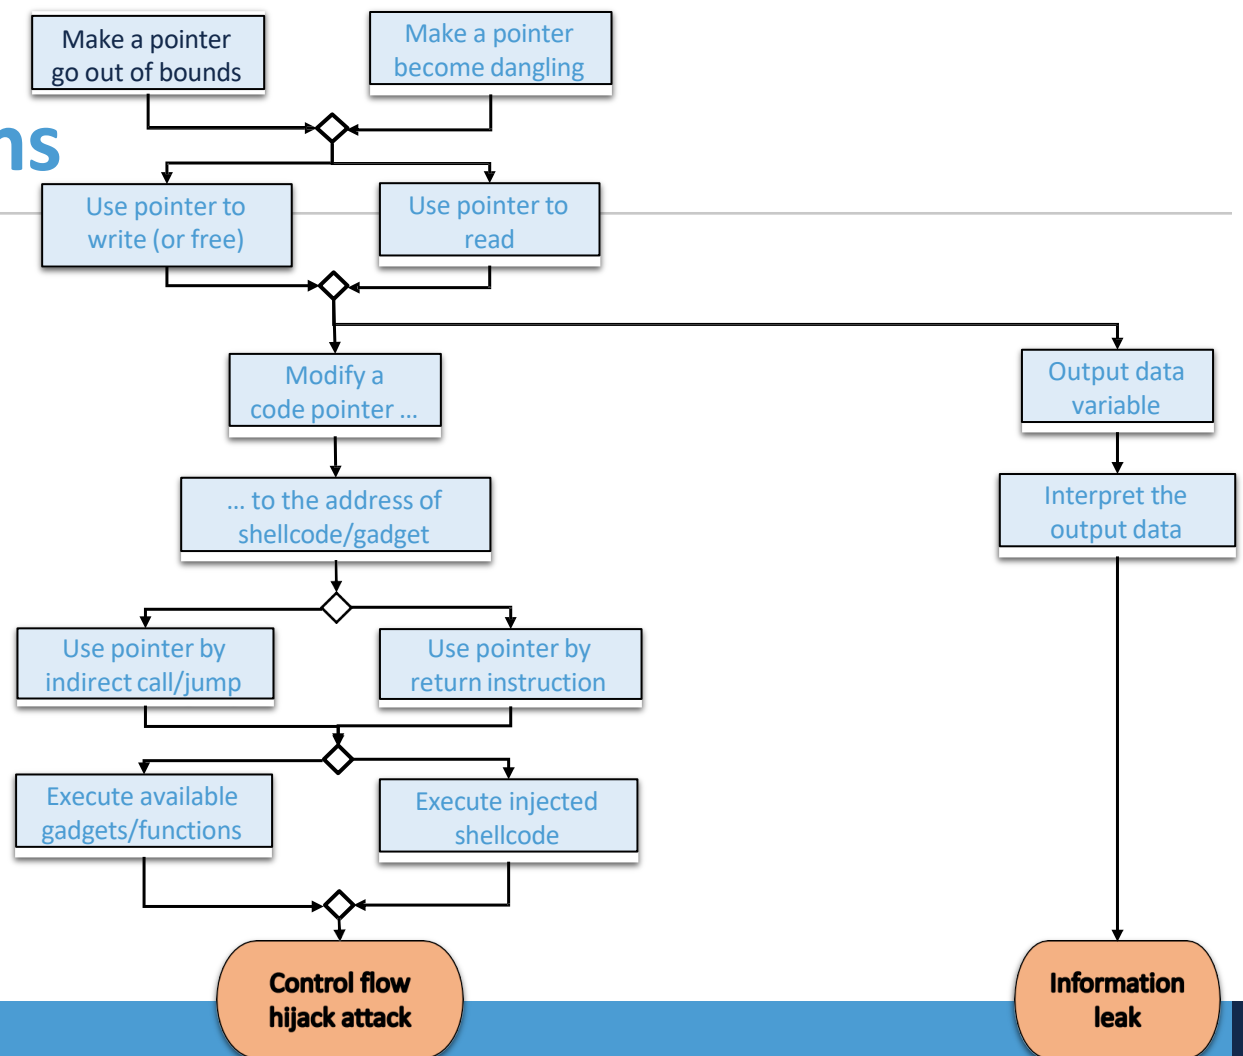


Gadget example:

```
pop rdi
```

```
ret
```

# Attack Variations

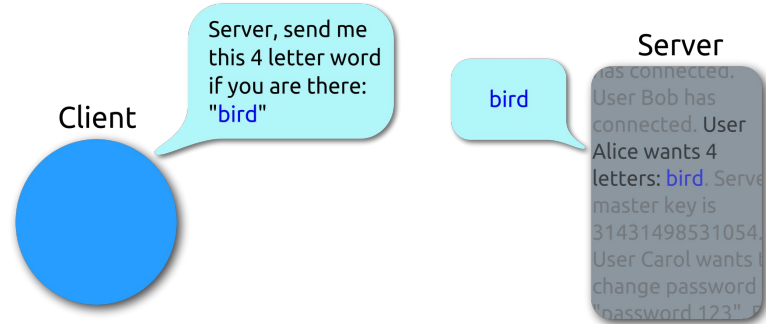


# HeartBleed Vulnerability

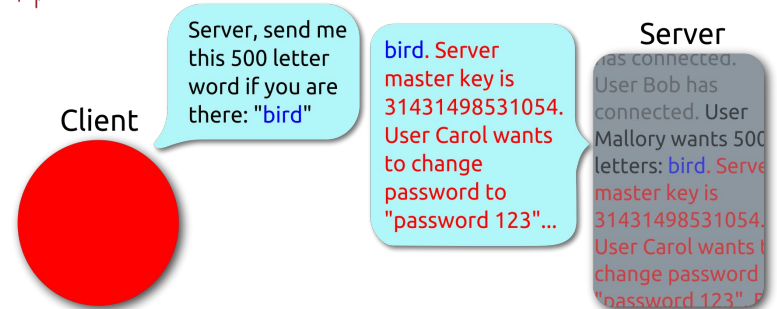
- Publicly disclosed in April 2014
- Bug in the OpenSSL cryptographic software library heartbeat extension
  - Missing a bounds check
  - `strncpy`

<https://heartbleed.com/>

## Heartbeat – Normal usage



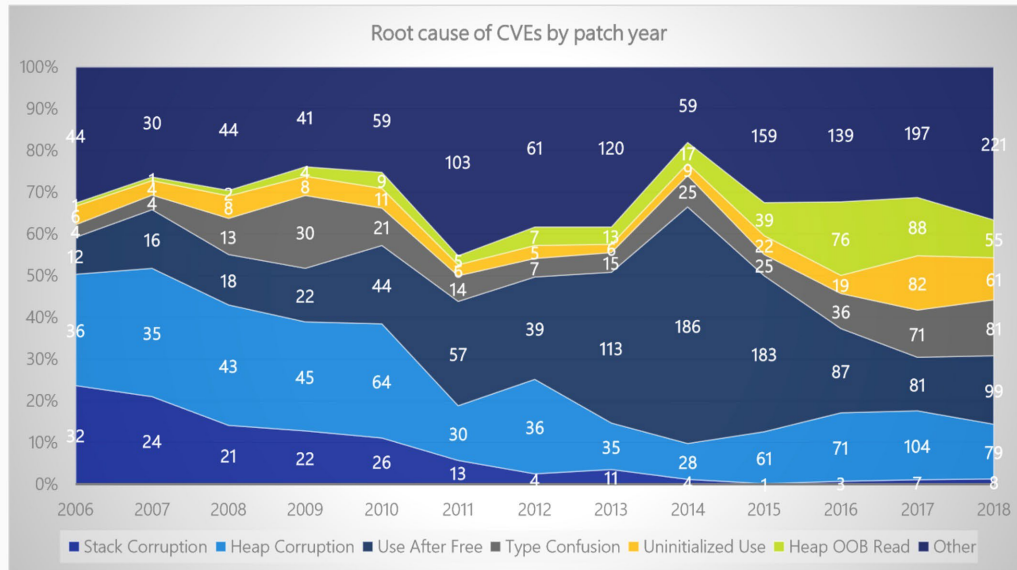
## Heartbeat – Malicious usage



# Trend reported by Microsoft

[https://github.com/microsoft/MSRC-Security-Research/tree/master/presentations/2019\\_02\\_BlueHatIL](https://github.com/microsoft/MSRC-Security-Research/tree/master/presentations/2019_02_BlueHatIL)

## Drilling down into root causes



Stack corruptions are essentially dead

Use after free spiked in 2013-2015 due to web browser UAF, but was mitigated by Mem GC

Heap out-of-bounds read, type confusion, & uninitialized use have generally increased

Spatial safety remains the most common vulnerability category (heap out-of-bounds read/write)

Top root causes since 2016:

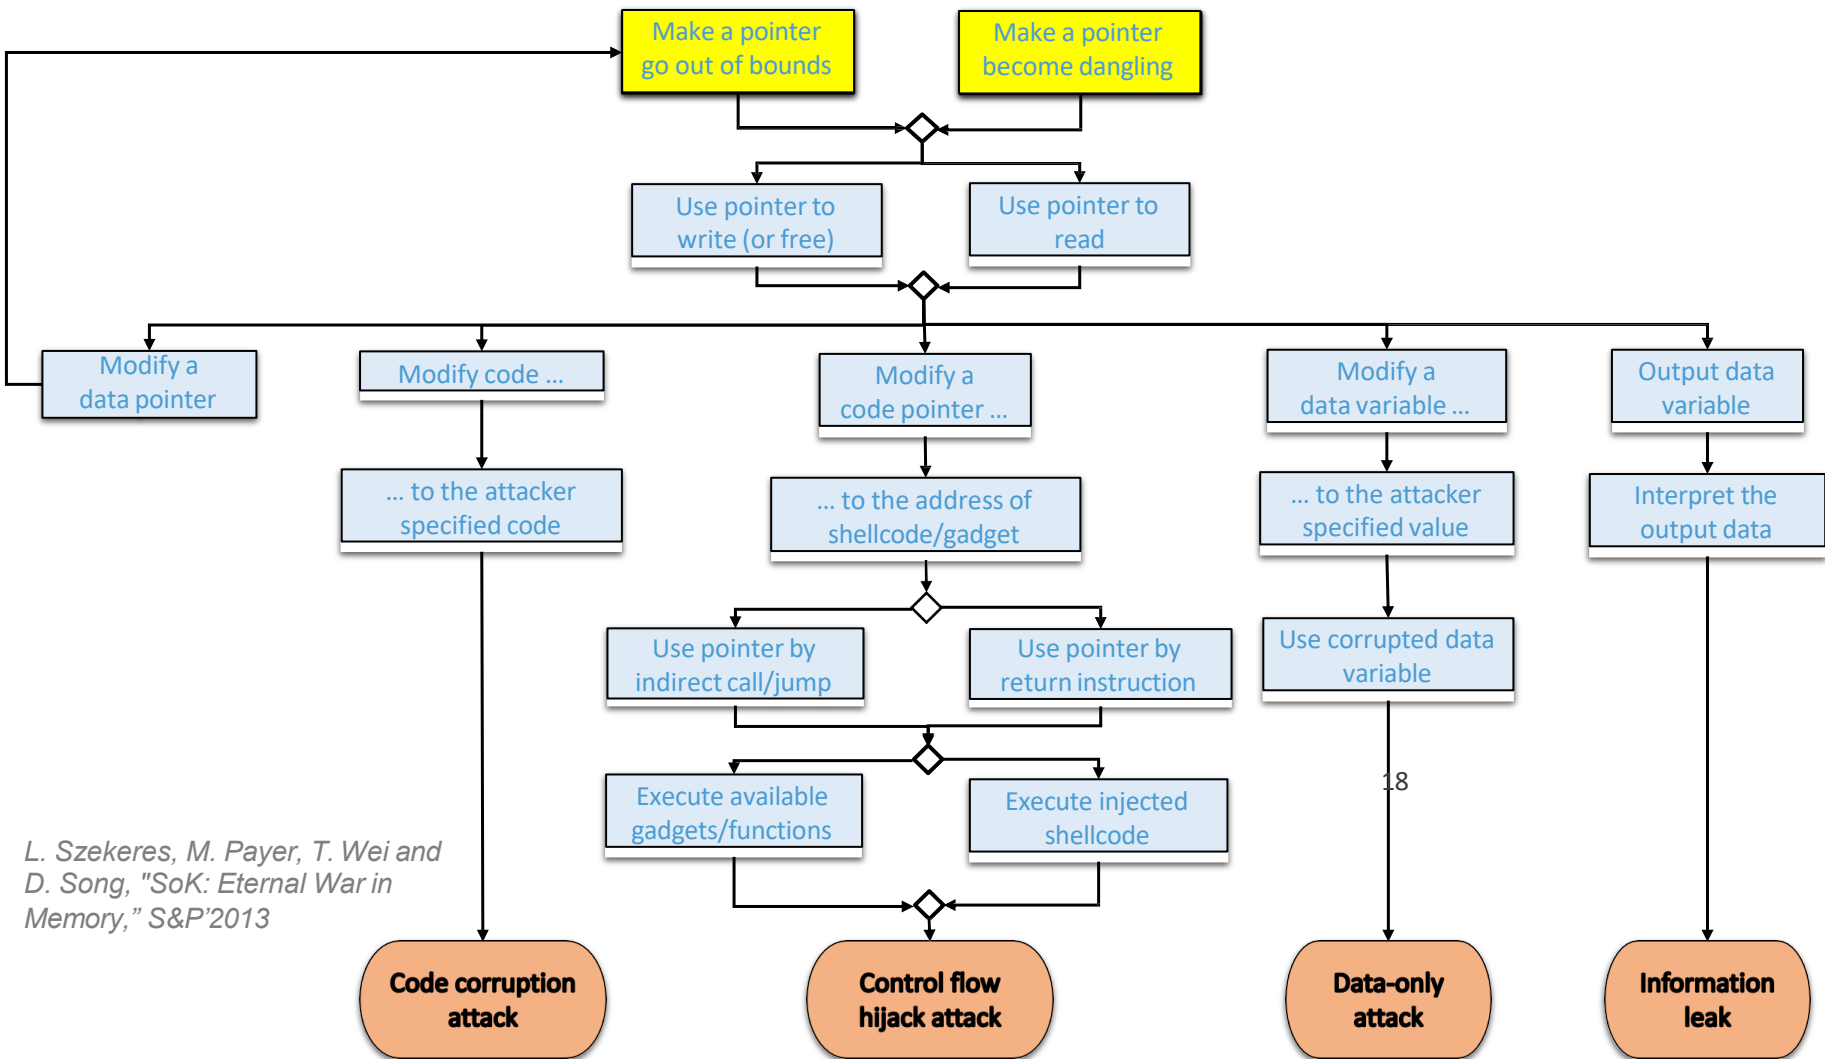
#1: heap out-of-bounds

#2: use after free

#3: type confusion

#4: uninitialized use

# Hardware Supported Mitigations



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

# Memory Safety

---

- Strongest security property that tries to address the problem at the root
  - Prevents attacker from reading or writing to memory
- Idea: include metadata and perform security checks at runtime
  - Spatial safety (bound information)
  - Temporal safety (allocation/de-allocation information)
- Software solutions
  - Problem #1: performance overhead, extra instructions to perform the check
  - Problem #2: where to store metadata? -> in shadow memory

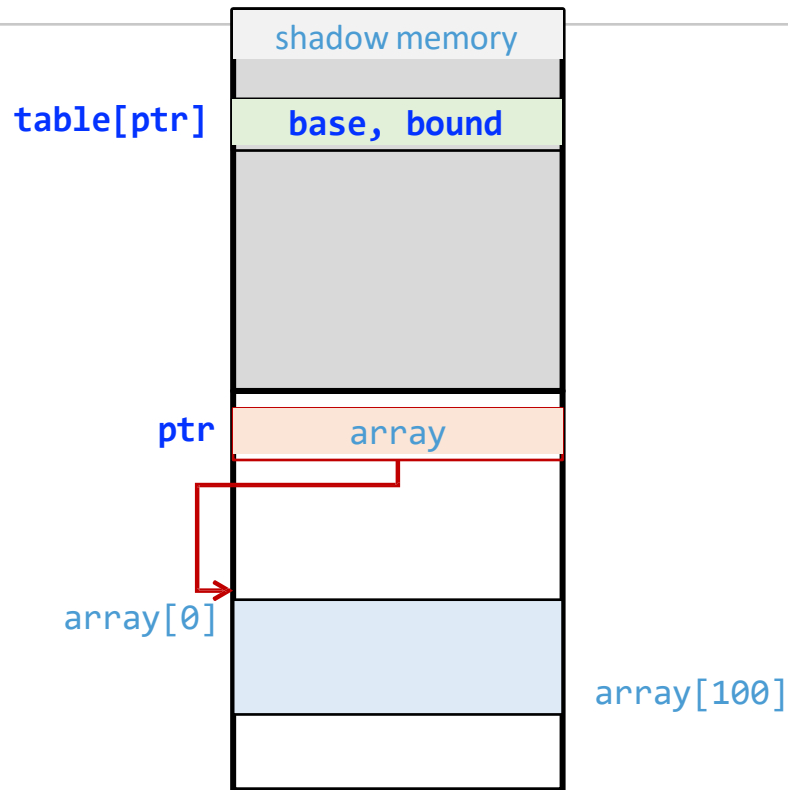
# SoftBound

## Creating a pointer:

```
int array[100];  
ptr = &array;  
ptr_base = &array[0];  
ptr_bound = &array[100];  
table[ptr] = {base, bound};
```

## Check a pointer:

```
{base, bound} = table[ptr];  
if (base > ptr || bound < ptr)  
    go to err;  
*ptr = 0xFF;
```



# SoftBound

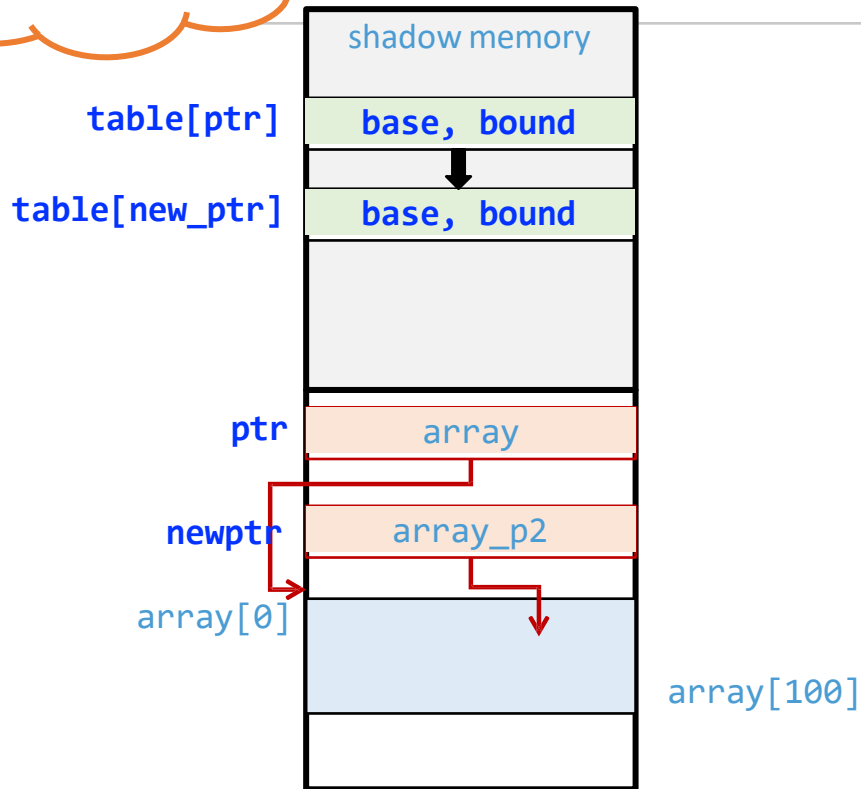
Compare number of  
memory accesses?

## Creating a pointer:

```
int array[100];  
ptr = &array;  
ptr_base = &array[0];  
ptr_bound = &array[100];  
table[ptr]={base, bound};
```

## Pointer arithmetic:

```
int* array_p2 = &array[10];  
newptr_base = table[ptr].base;  
newptr_bound = table[ptr].bound;  
table[newptr]={base, bound};
```



## HW Support for Memory Safety

---

- A lot of work in both academia and industry
  - Various design trade-offs
- Intel MPX (Memory Protection Extensions)
  - announced in 2013, produced in 2015 (Skylake), discontinued in 2019
- Arm MTE (Memory Tagging Extensions)
  - Introduced with ARM v9

# Intel MPX (Memory Protection Extension)

## 4 bounds registers (bnd0-3)

- `Bndmk`: create base and bound metadata
- `Bndldx/bndstx`: load/store metadata from/to bound tables
- `Bndcl/bndcu`: check pointer with lower and upper bounds

### Original Program

```
p=malloc(16);  
... // p = p + 4;  
*p = 'a';
```

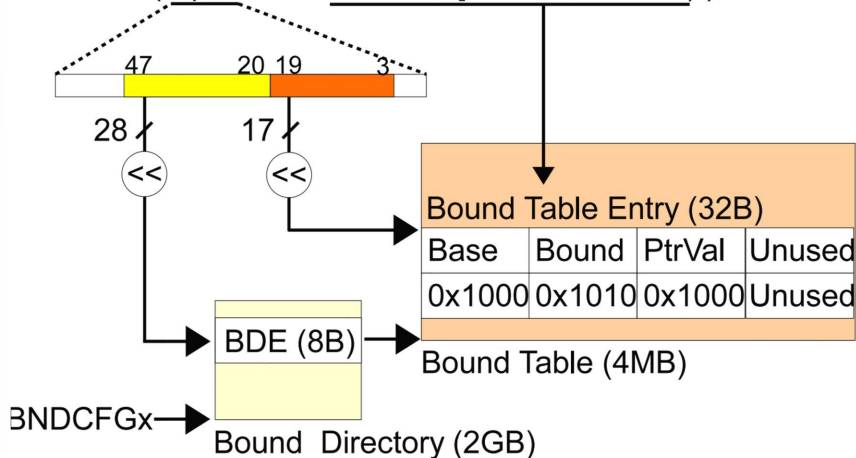
### Instrumented Program

```
p=malloc(16);  
bnd0 = bndmk(p,16);  
bndstx (&p,p,bnd0);  
... // p = p + 4;  
bnd1 = bndldx(&p,p);  
bndcl (&p, bnd1);  
bndcu (&p, bnd1);  
*p = 'a';
```

Store the metadata in  
a two-level table  
in hardware



```
bndstx ( Ptr Addr., Ptr Val., [Base, Bound] );  
bndstx ( &p, 0x1000, [0x1000, 0x1010] );
```



# Analysis of Intel MPX

---

## Performance and cost:

- + Reduce number of instructions, and reduce register pressure
- + No branch instructions, doesn't pollute the branch predictor
- High overhead: Not any faster than software!
  - 50% overhead average case, 300% overhead worst case
- + Two-level page table organization is more area-efficient
- High overhead: loading/storing bounds registers involves two-level table lookup



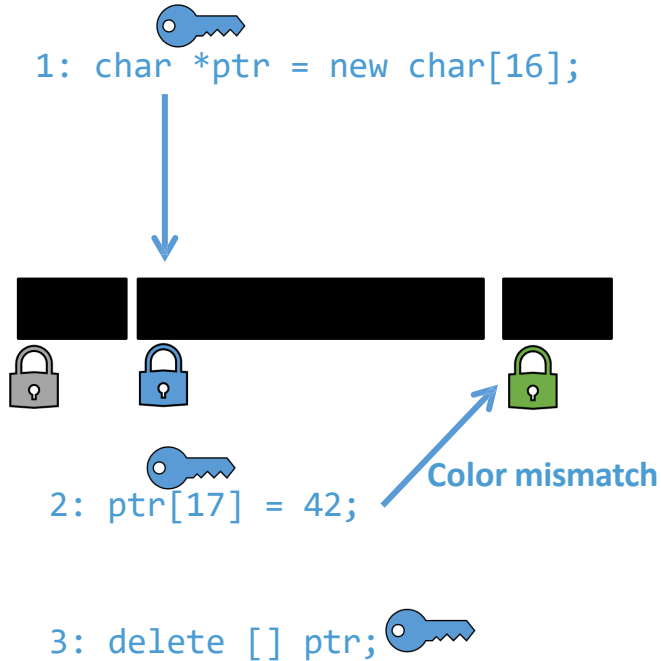
## Compatibility:

- Not straightforward about how to extend the scheme to support temporal safety
- Does not support multithreading transparently
  - Causes data races in legacy programs
- All the code need to be rewritten, otherwise either security breaks or correct code broken

*Intel MPX Explained: A Cross-layer Analysis of the Intel MPX System Stack; OLEKSENKO et al; SIGMETRICS'18*



# ARM MTE (Memory Tagging Extension)




- The concept of keys and locks
- Memory locations are tagged by adding **four bits** of metadata to each **16 bytes** of physical memory

# ARM MTE (Memory Tagging Extension)



```
1: char *ptr = new char[16];
```



```
2: ptr[17] = 42;
```

```
3: delete [] ptr; 
```

- The concept of keys and locks
- Memory locations are tagged by adding **four bits** of metadata to each **16** bytes of physical memory

# Analysis of ARM MTE



```
1: char *ptr = new char[16];
```



```
2: ptr[17] = 42;
```

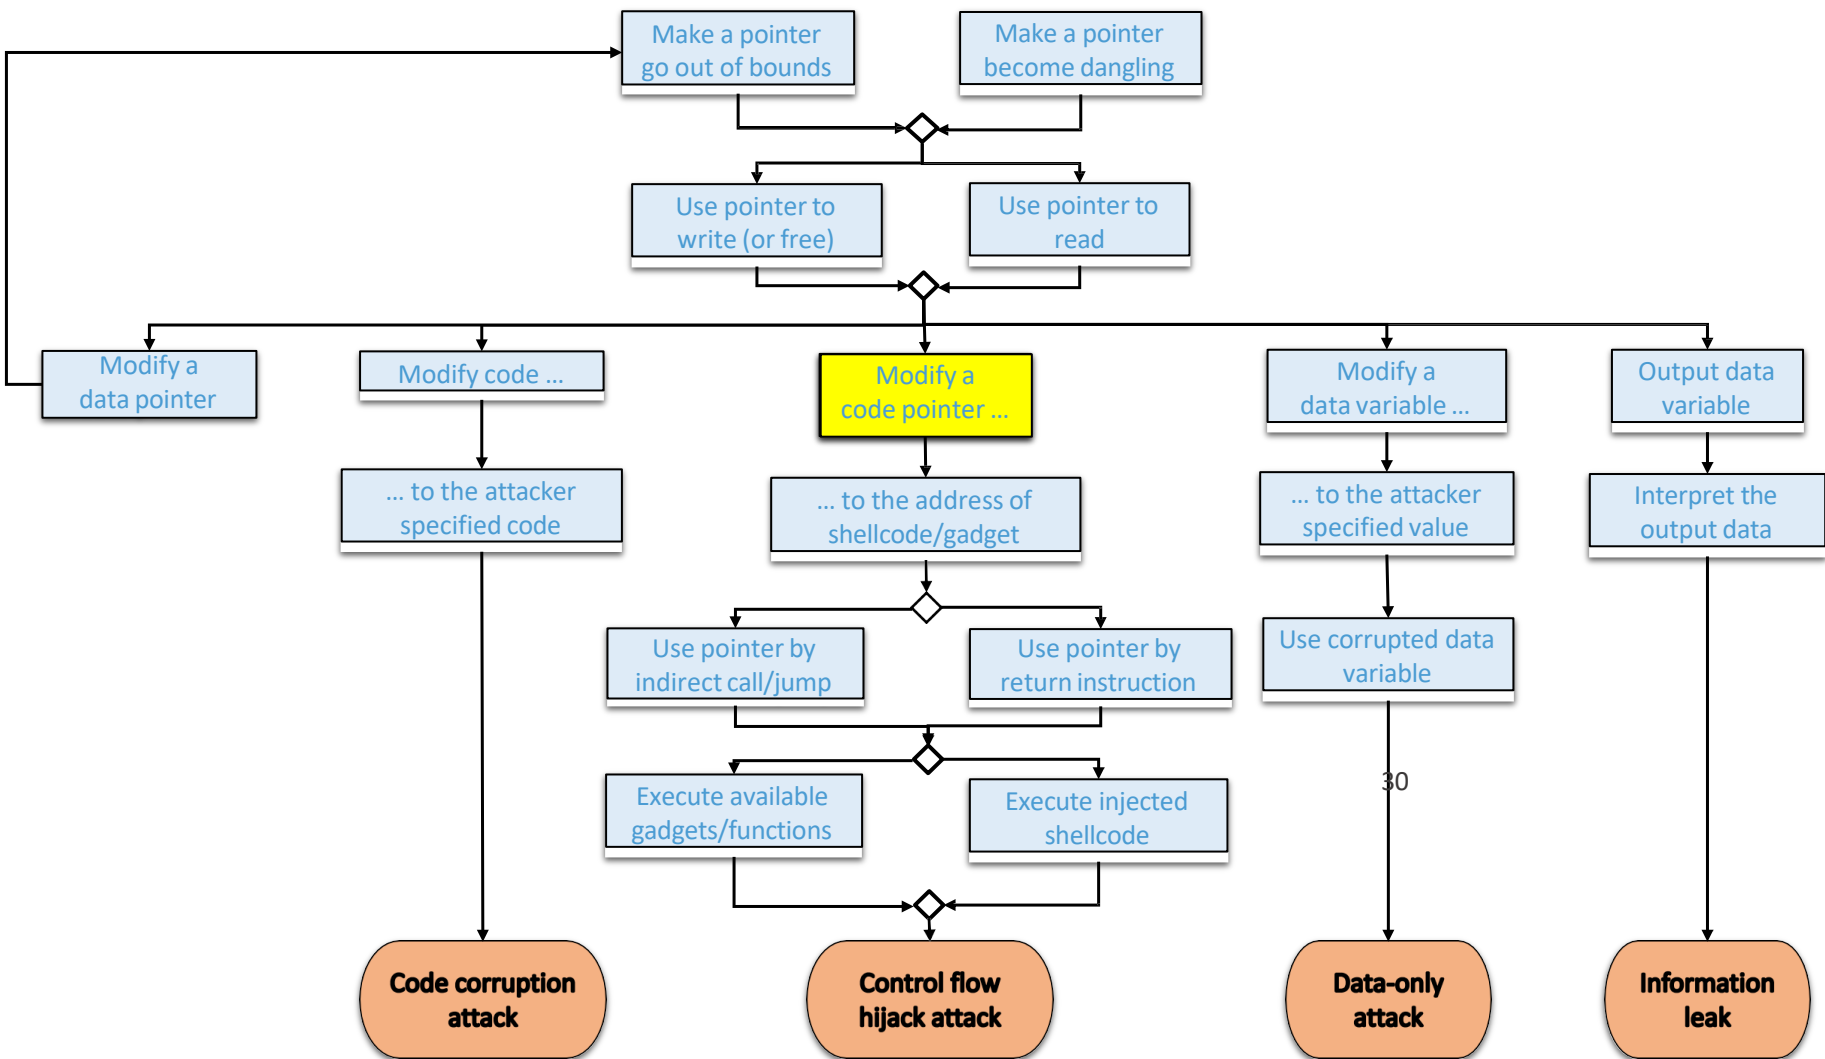
```
3: delete [] ptr;
```

- Where to store tags (key and lock)?
  - Pointer tag is stored in top 4 unused bits inside the pointer (no extra register needed)
  - Physical memory tag is stored in hardware (new hardware needed for both DRAM and cache)
- Limited tag bits
  - Cannot ensure two allocations have different colors
  - But can ensure that the tags of sequential allocations are always different

# Analysis of ARM MTE

---

- **Security:**
  - - Coarse-grained spatial safety. Non-sequential violation is detected probabilistically
  - + Can support temporal safety similar to spatial safety
- **Performance and other overhead:**
  - + Storage overhead is ok
    - 4 bits per 16 bytes
  - + Performance overhead is low, mostly lies in the allocation and free time, since need to modify tags in bulk
- **Compatibility:**
  - + To protect heap, modify libraries to do malloc and free; modify OS to trap on invalid pointer. No extensive code rewritten needed.



# Control-flow Integrity

---

- To maintain code pointer integrity
- Naive idea:
  - Make pointer immutable (read-only)
    - Only work for global offset table
- How about other pointers?
  - Return address?
  - Programmer-defined function pointers

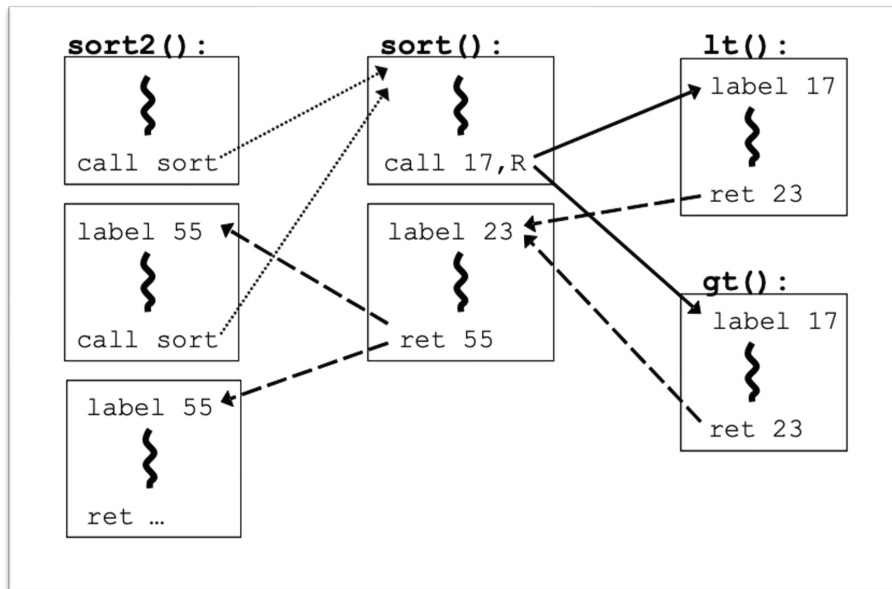
# Control Flow Integrity (CFI)

```
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;*  
Martín Abadi, et al. CCS'05

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

INTEL  
CET

=


INDIRECT BRANCH  
TRACKING (IBT)

+

SHADOW  
STACK (SS)

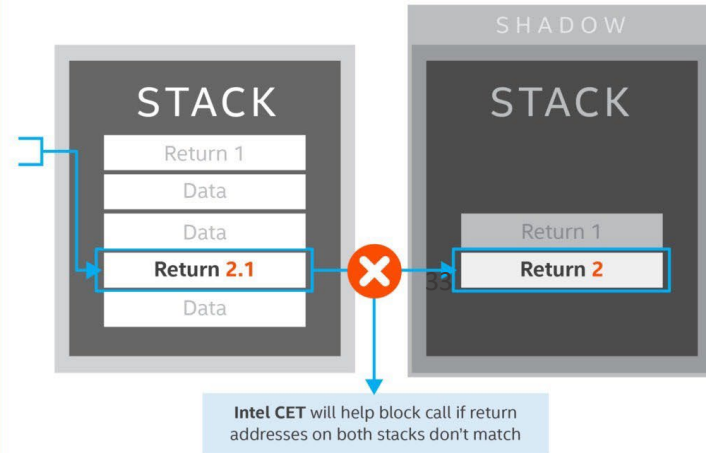
## INDIRECT BRANCH TRACKING (IBT)

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



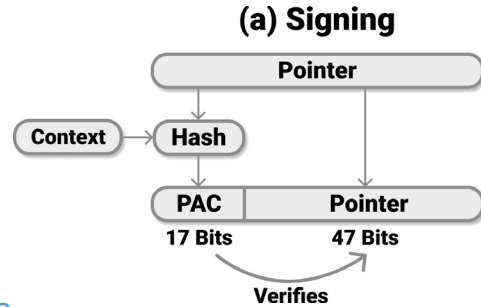
## SHADOW STACK (SS)

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



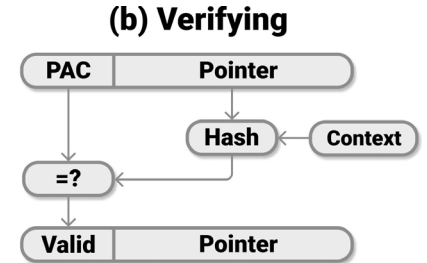
# ARM PA (Pointer Authentication)

- Introduced in 2017, widely used in Apple processors
- 64-bit pointer, but 48-bit virtual address space
  - Unused high bits
- Hash:
  - A tweakable message authentication code (MAC)
  - ARM calls it **PAC (pointer authentication code)**
- Context:
  - secret key
  - salt (could be the stack pointer)



Function prologue

```
1 pacia lr, sp
2 sub sp, sp, #0x40
3 str lr, [sp, #0x30]
4 ...
```

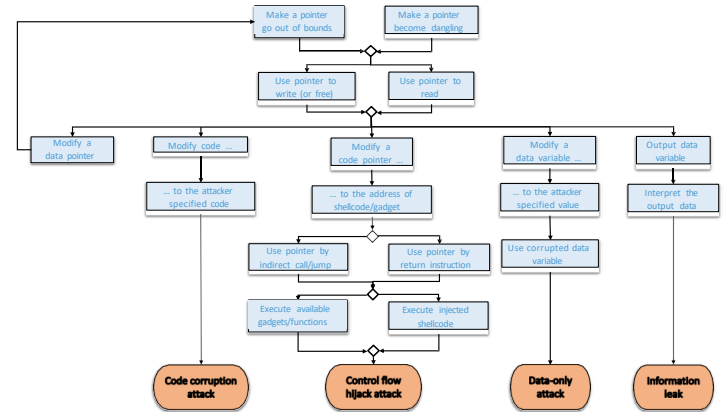


Function epilogue

```
1 ldr lr, [sp, #0x30]
2 add sp, sp, #0x40
3 autia lr, sp
4 ret
```

# Summary

- Memory corruption problems: An eternal war
- Attack variations and mitigations
- Trade-offs in hardware support widely used on real systems





THE UNIVERSITY  
*of* NORTH CAROLINA  
*at* CHAPEL HILL