

Hardware Support for Memory Safety

Mengjia Yan

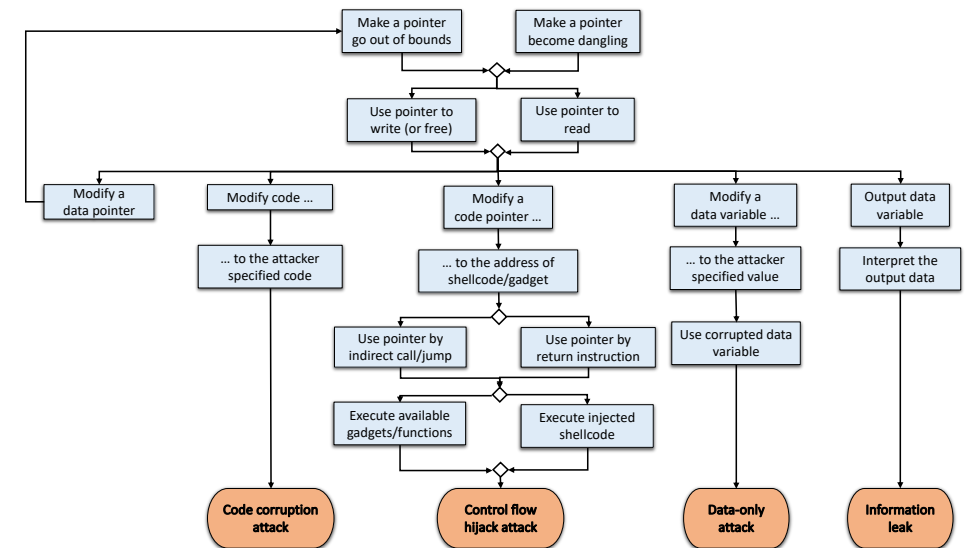
mengjia@csail.mit.edu

6.888 Secure Hardware Design



Overview

- Memory corruption problems and existing mitigations
 - SoK: Eternal War in Memory; Szekeres et al; S&P'13
- Recent progress on hardware defenses



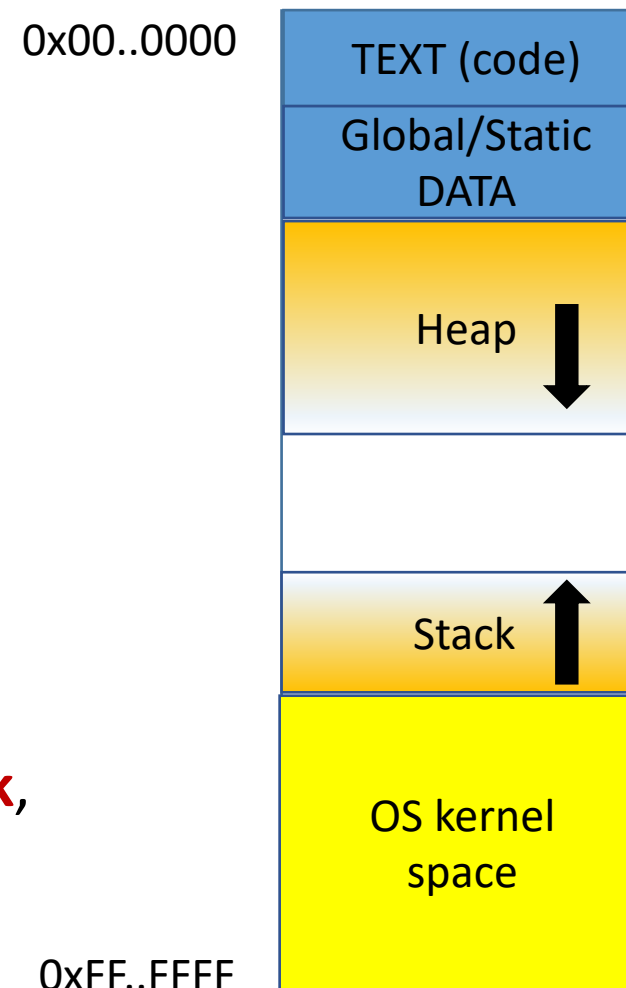
The Problem

- C/C++ is unsafe
- Everybody runs C/C++ code
- They surely have exploitable vulnerabilities



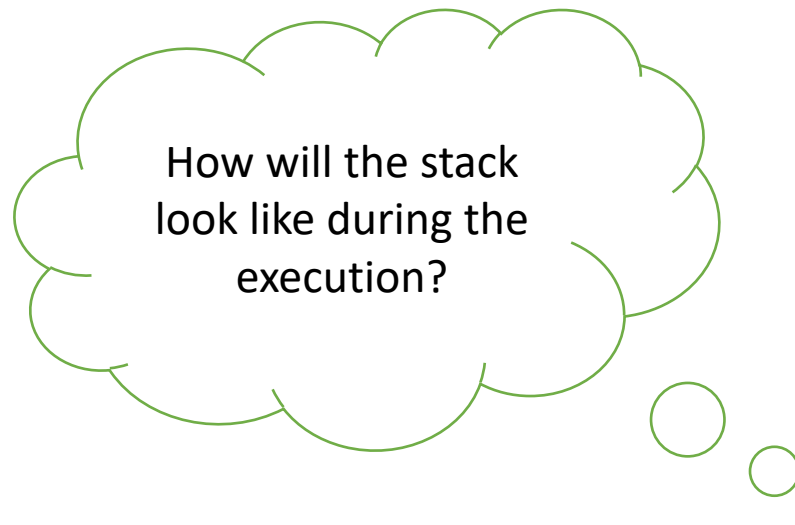
Low-level Language Basics (C/C++/Assembly)

- Programmers have more control
 - + Efficient
 - Bugs
 - Programming productivity
- **Pointers**
 - Address of variables: index of memory location where variable is stored
 - It is programmers' responsibility to do **pointer check**, e.g. NULL, out-of-bound, use-after-free



Low-level Language Basics

```
int hello() {  
    int a = 100;  
    return a;  
}  
int main() {  
    int a;  
    int b = -3;  
    int c = 12345;  
    int *p = &b;  
    int d = hello();  
    return 0;  
}
```



TEXT (code)

stack



Code Injection Attack Example

```
int func(char *str) {  
    char buffer[12];  
    strncpy(buffer, str, len(str));  
    return 1;  
}
```

```
int main() {  
    ....  
    func(input);  
    ...  
}
```

Shell code:

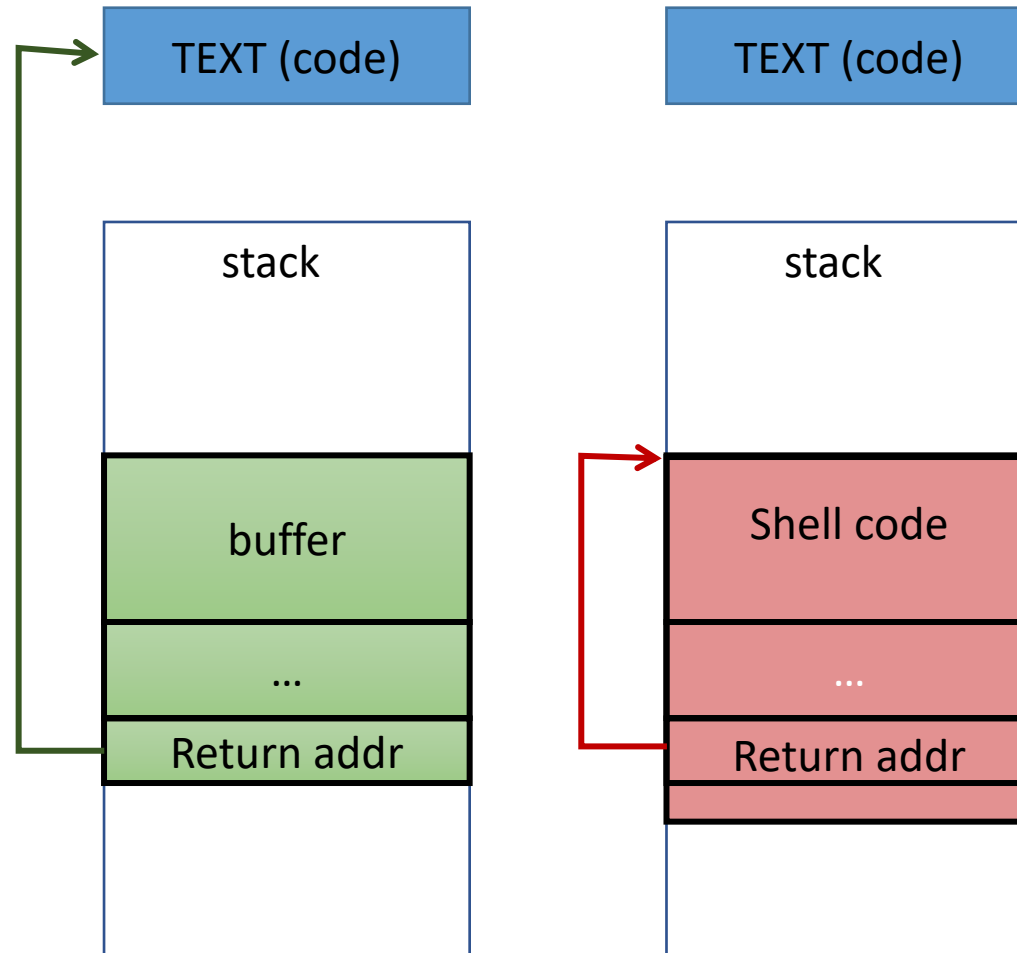
PUSH `"/bin/sh"`
CALL system

TEXT (code)

stack

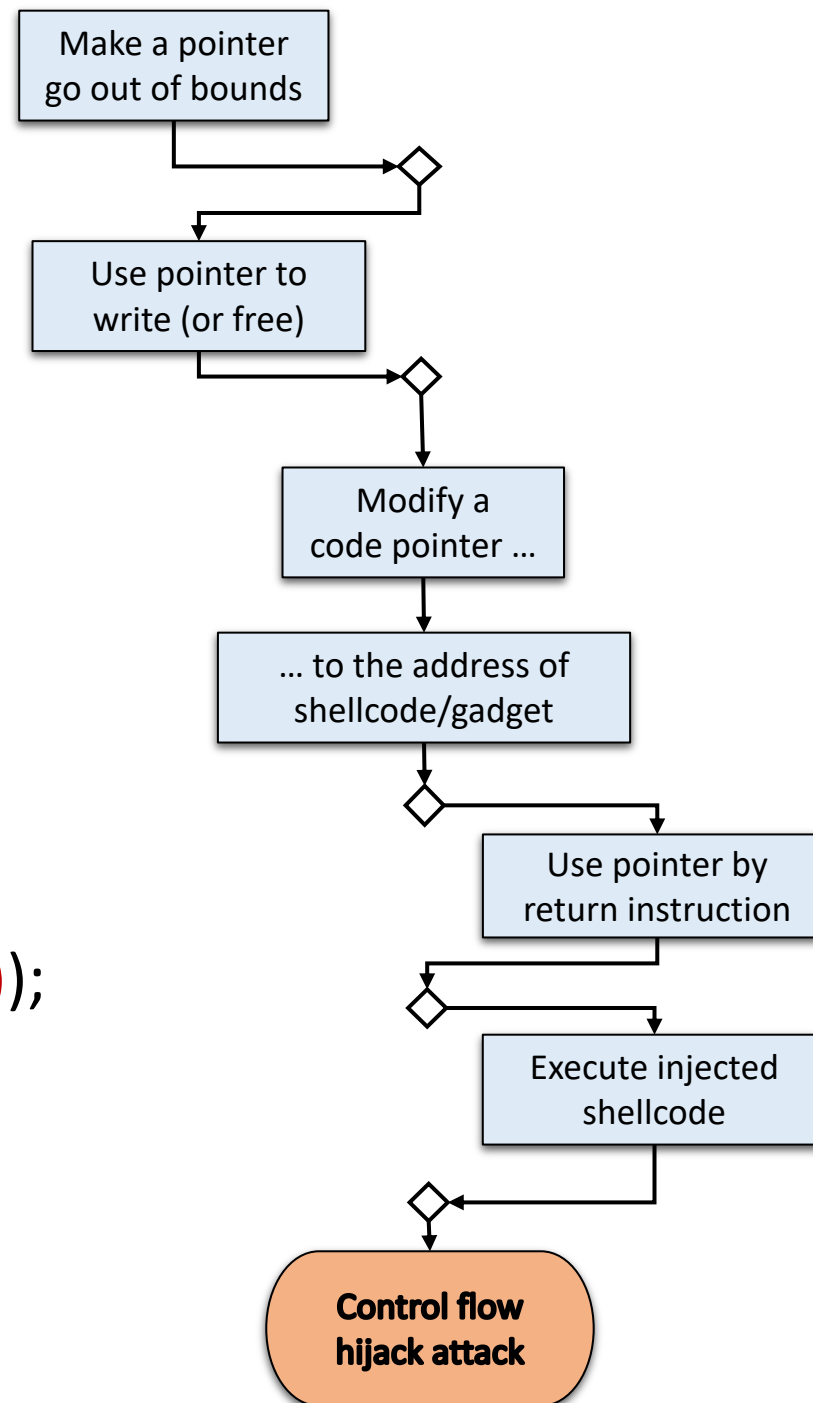


Code Injection Attack

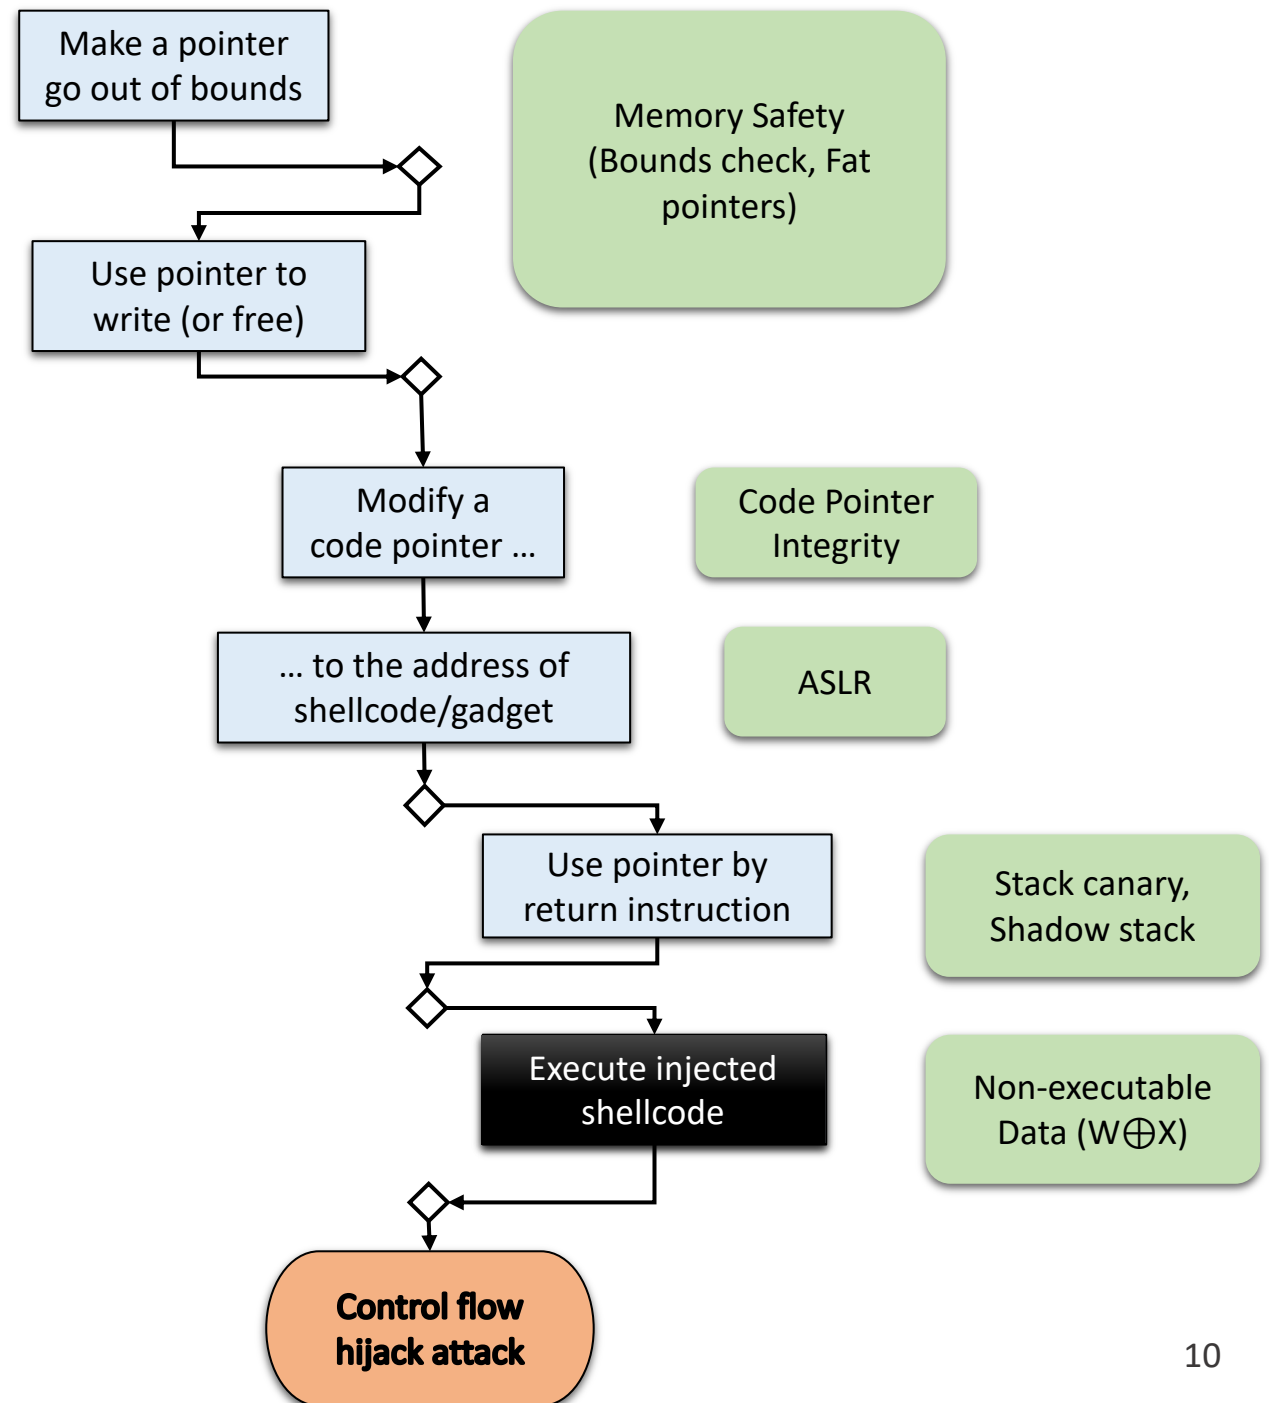


- Think about mitigations

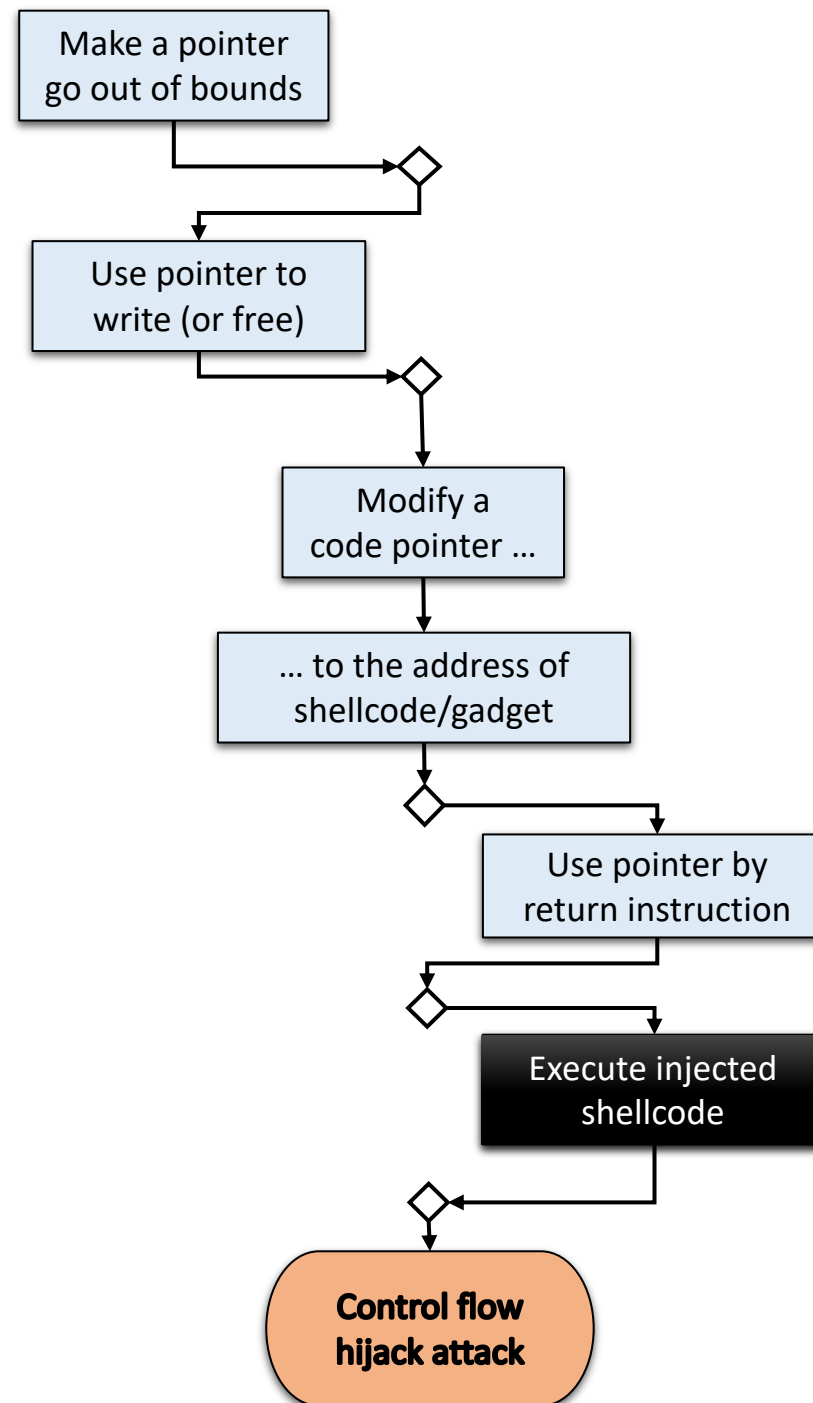
```
int func (char *str) {  
    char buffer[12];  
    strncpy(buffer, str, len(str));  
    return 1;  
}
```



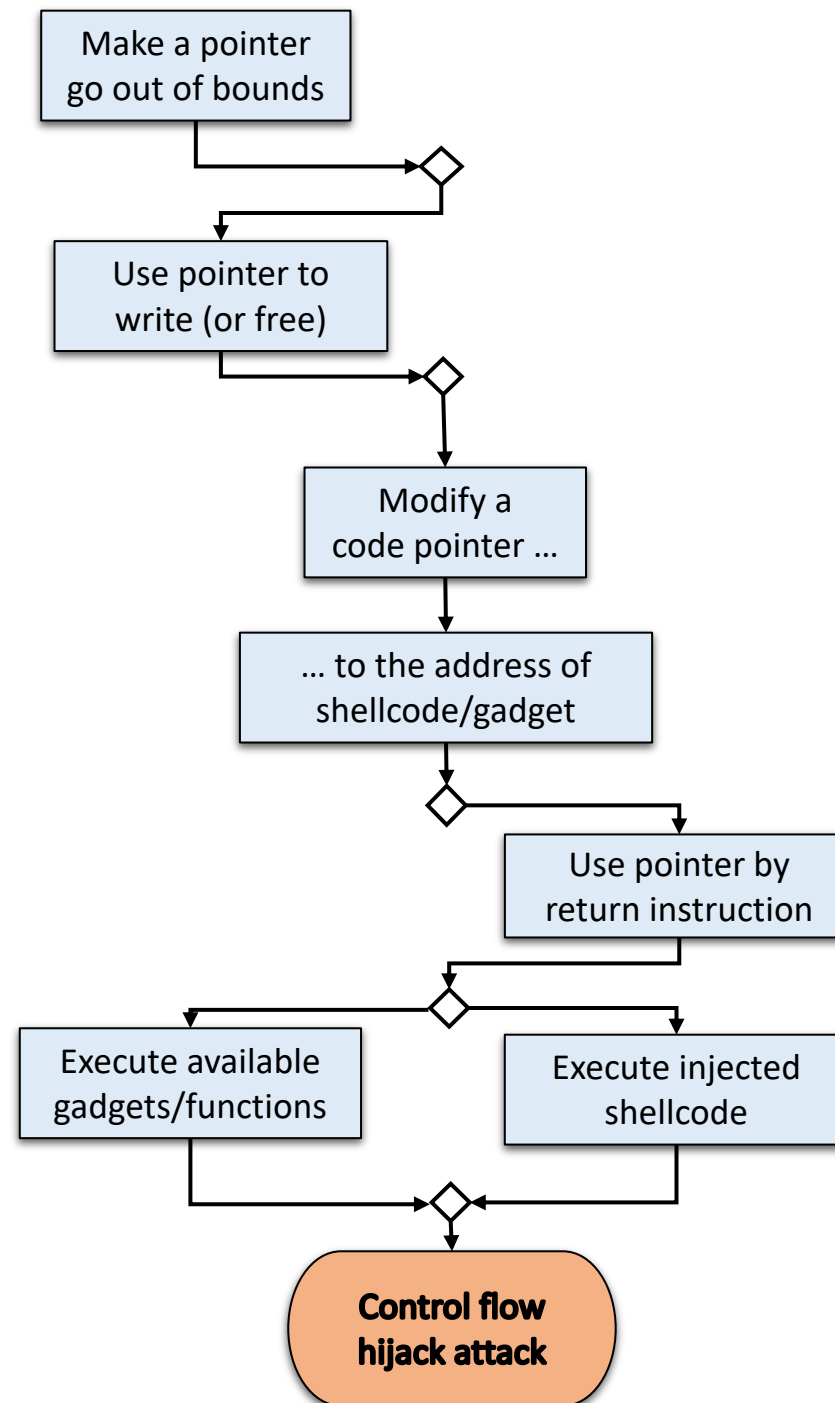
Mitigations



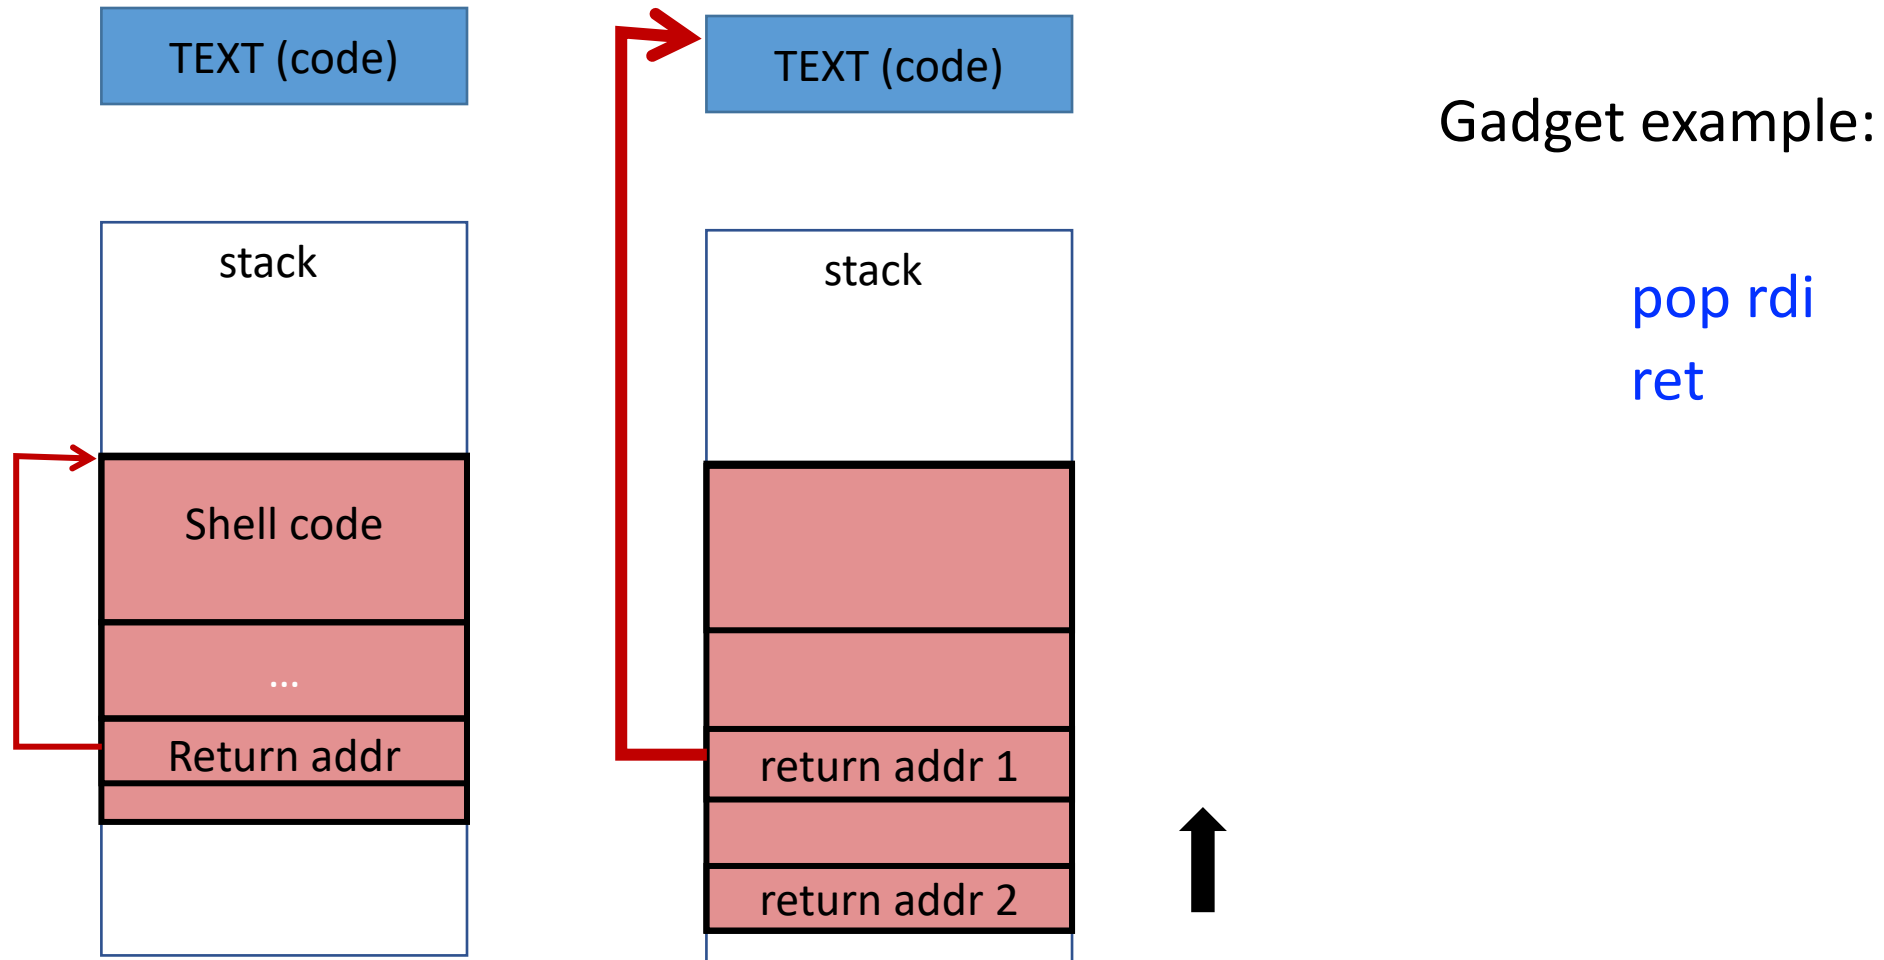
- **Think about attack variations**



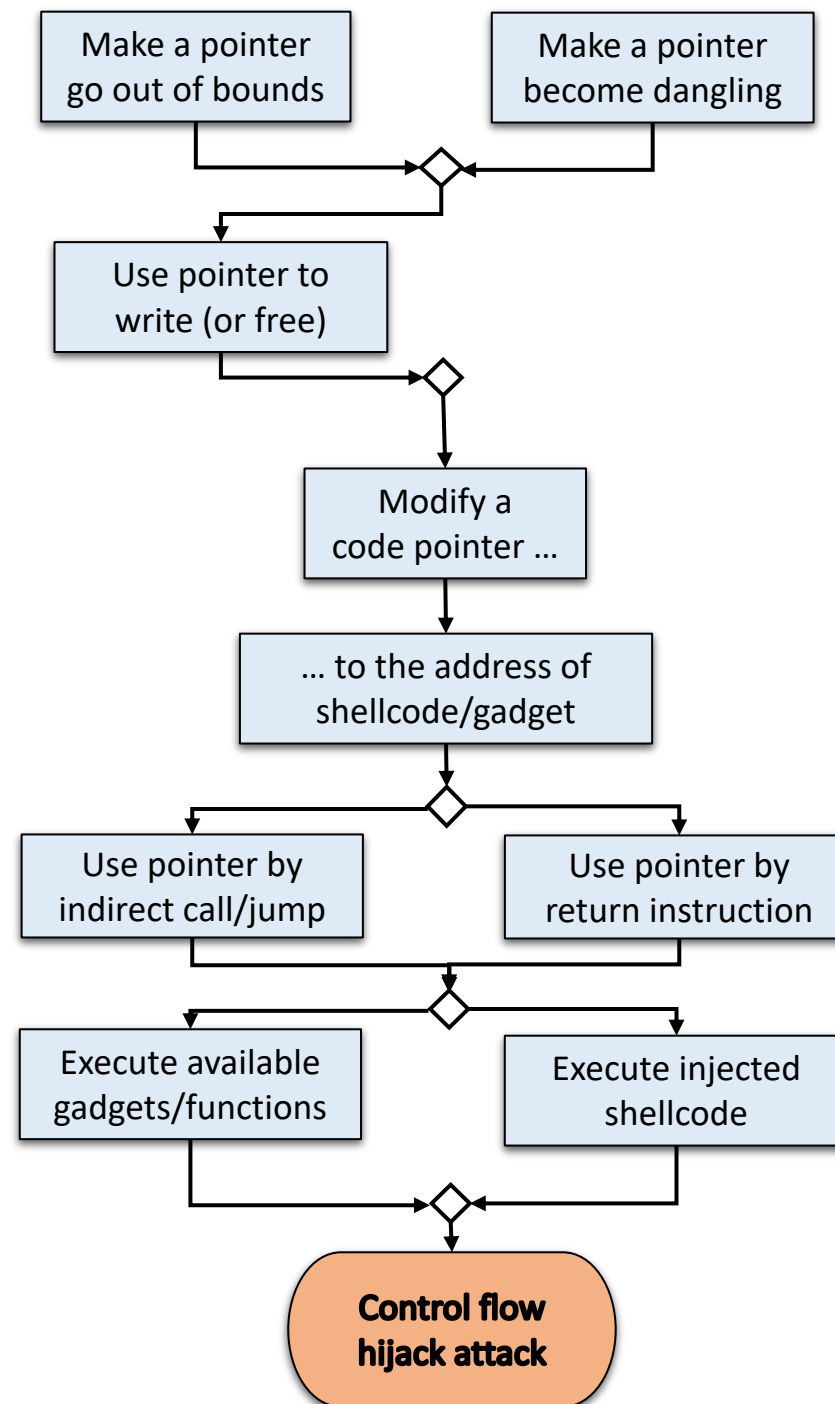
Attack Variations



Return-Oriented Programming (ROP)



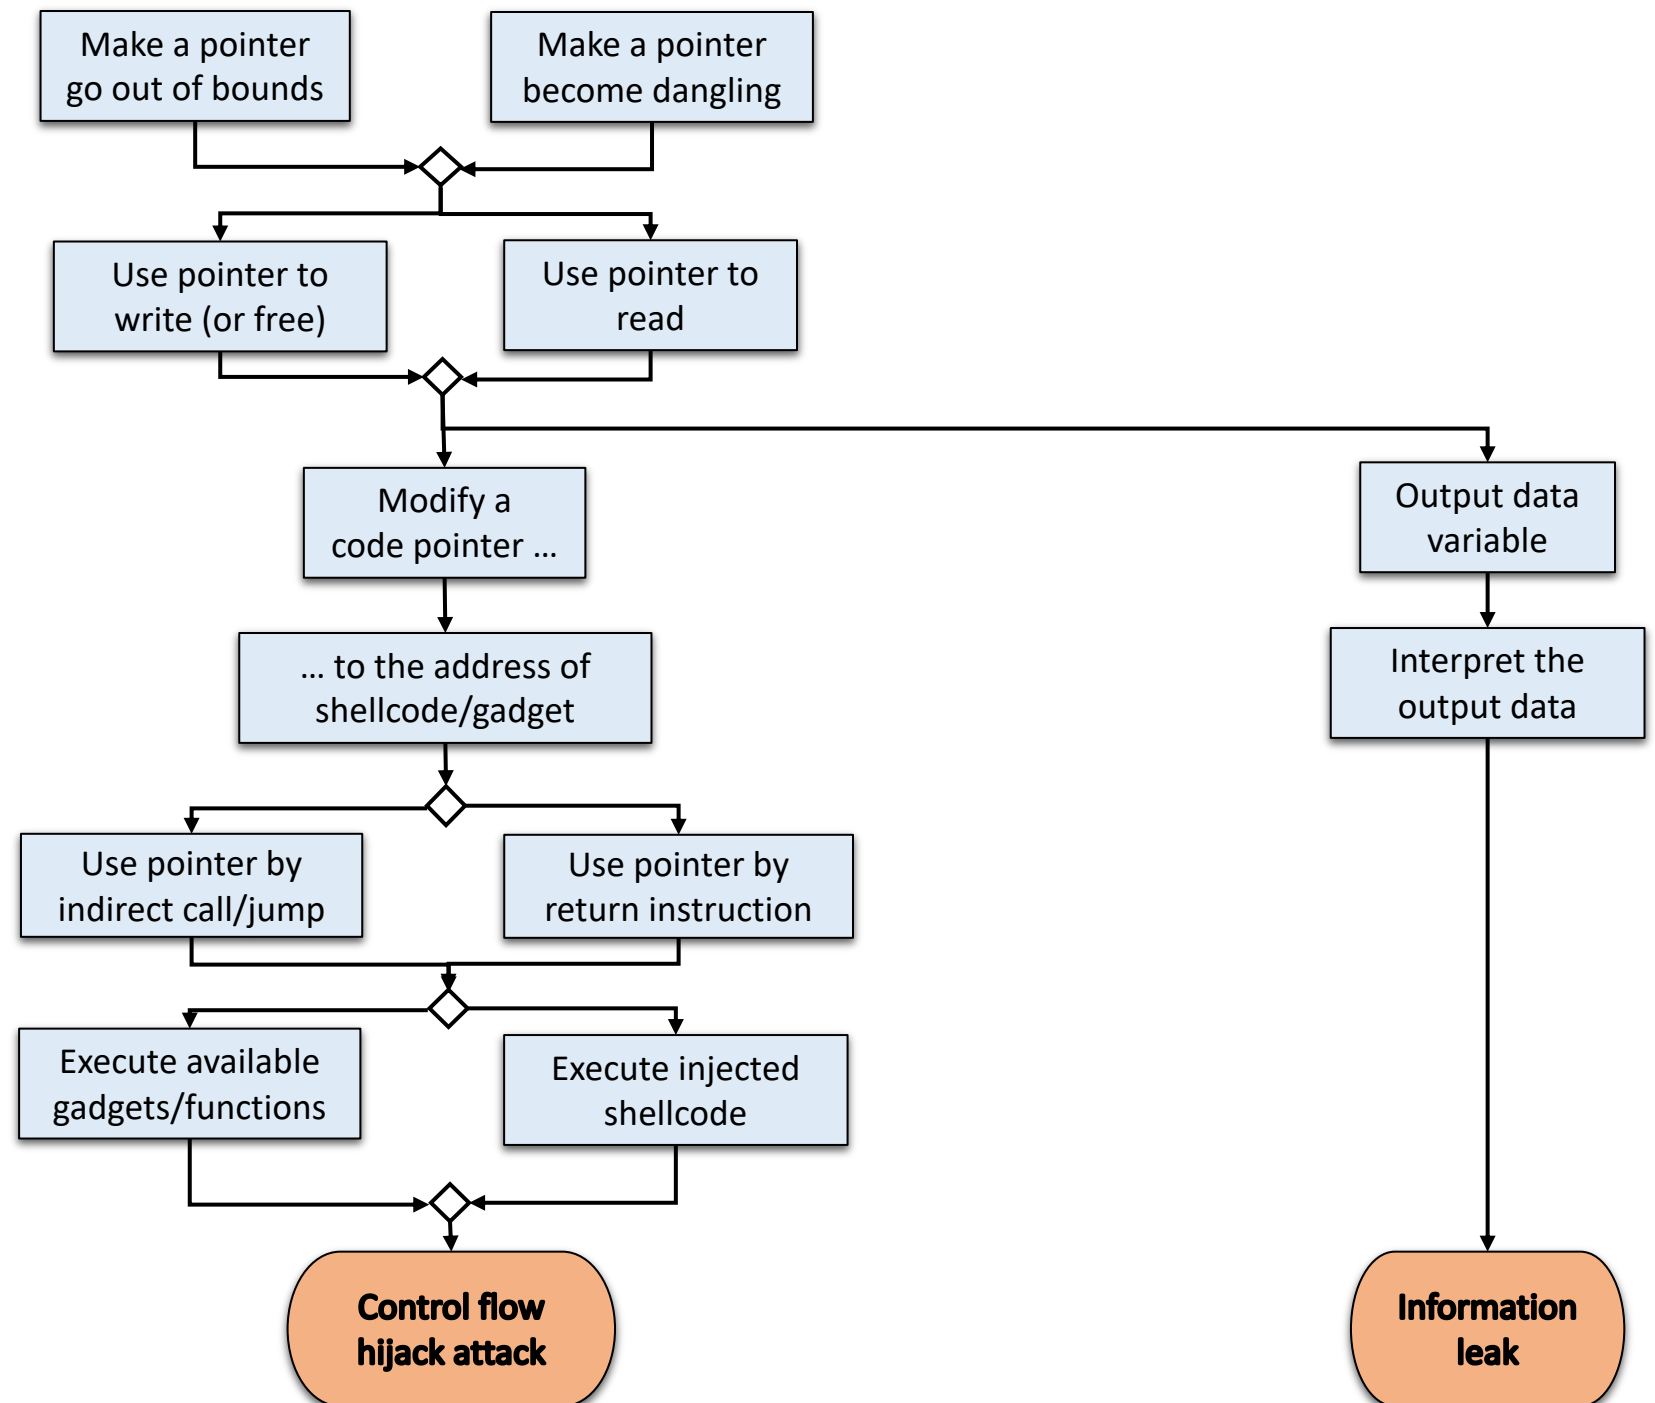
Attack Variations



Use-After-Free

- Example
- How to check?

- **More Variations**

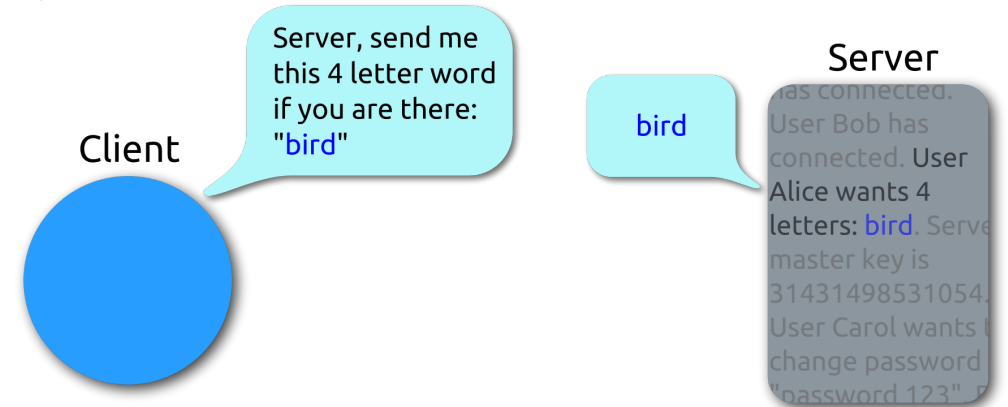


HeartBleed Vulnerability

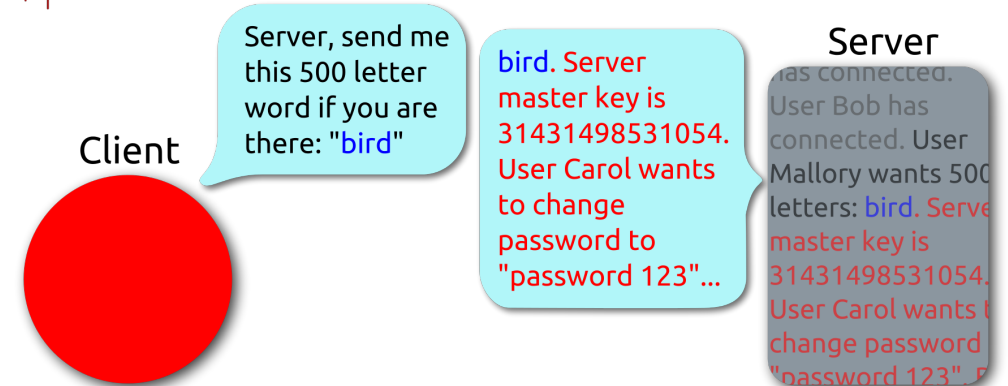
- Publicly disclosed in April 2014
- Missing a bound check
- Bug in the OpenSSL cryptographic software library heartbeat extension

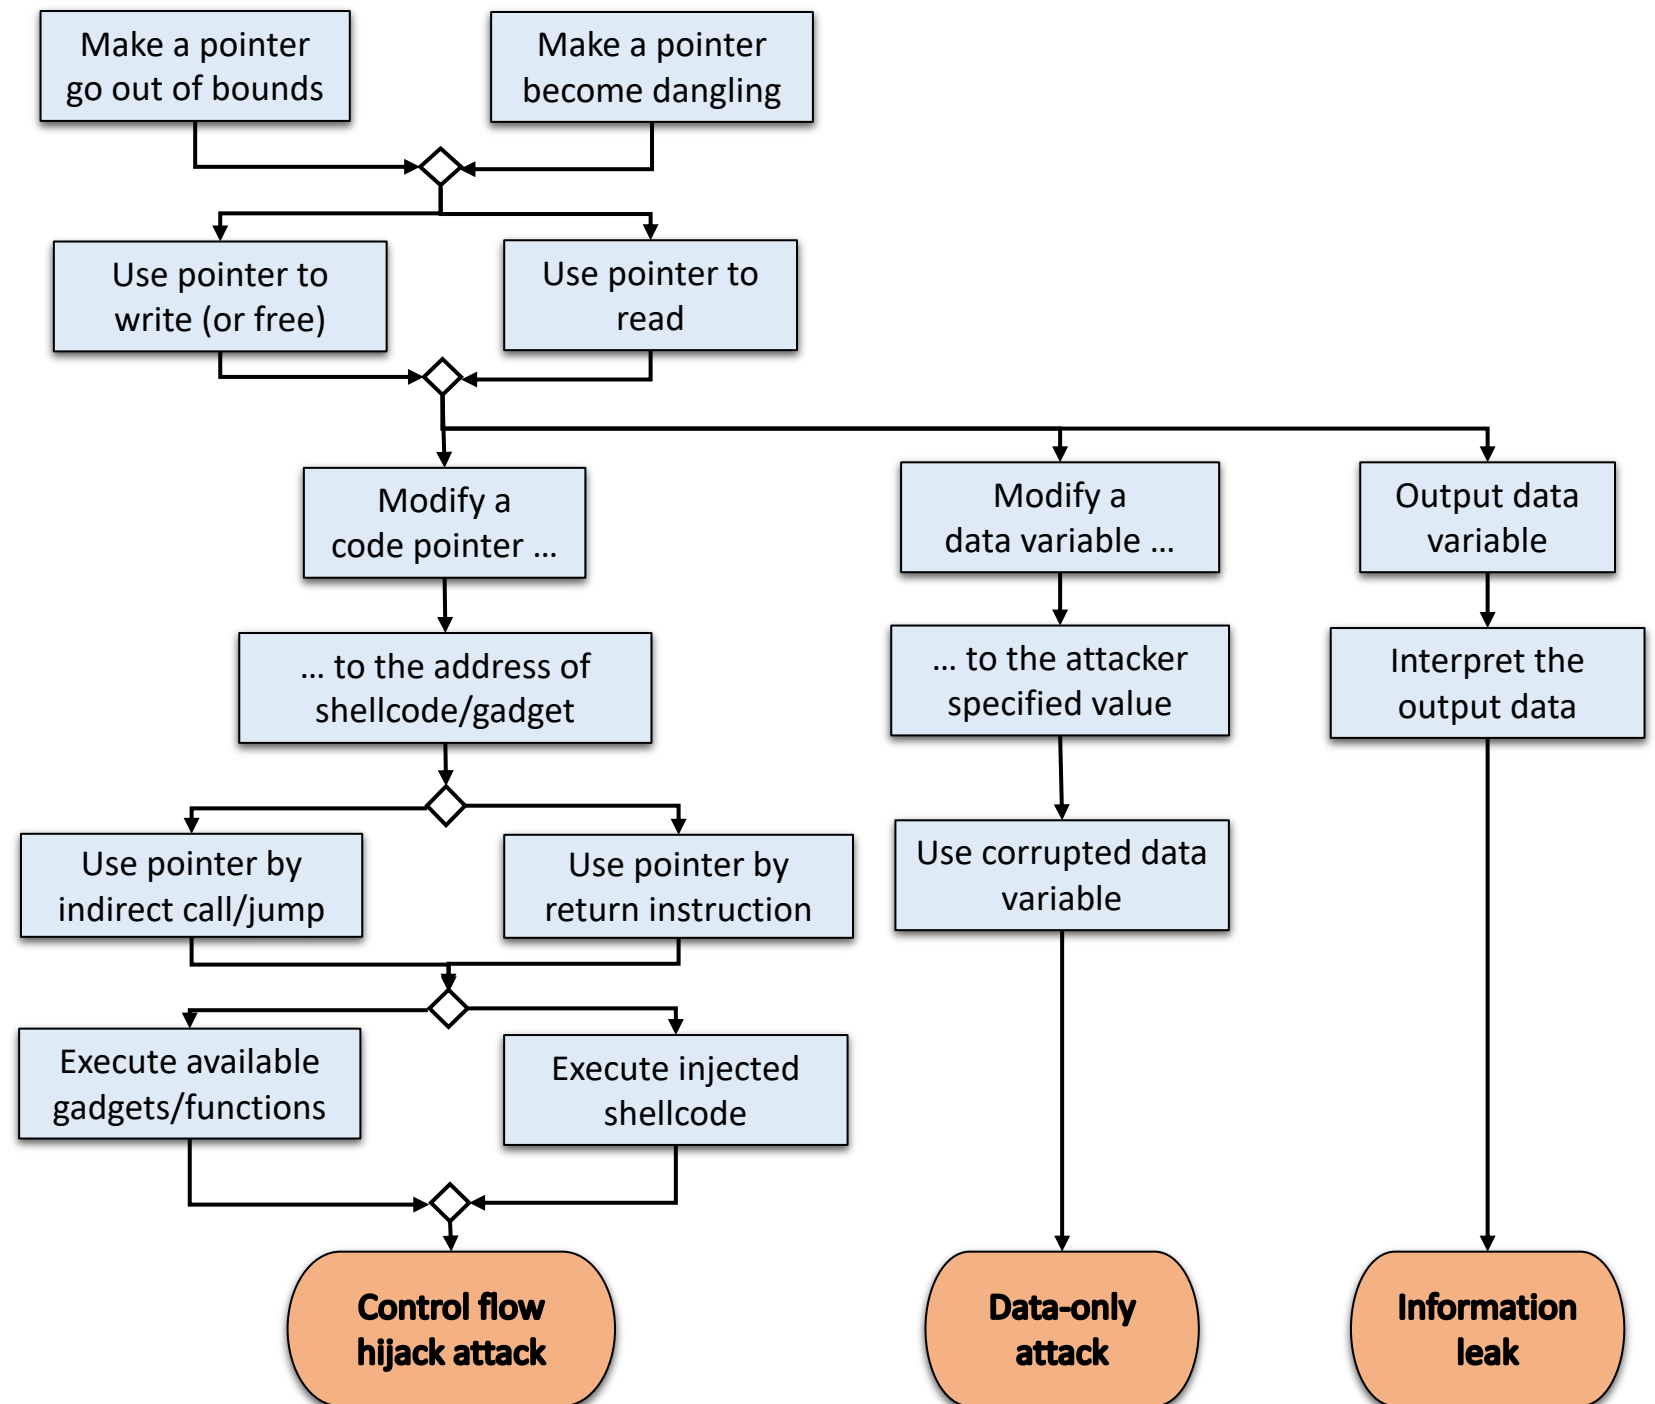
<https://heartbleed.com/>

Heartbeat – Normal usage



Heartbeat – Malicious usage





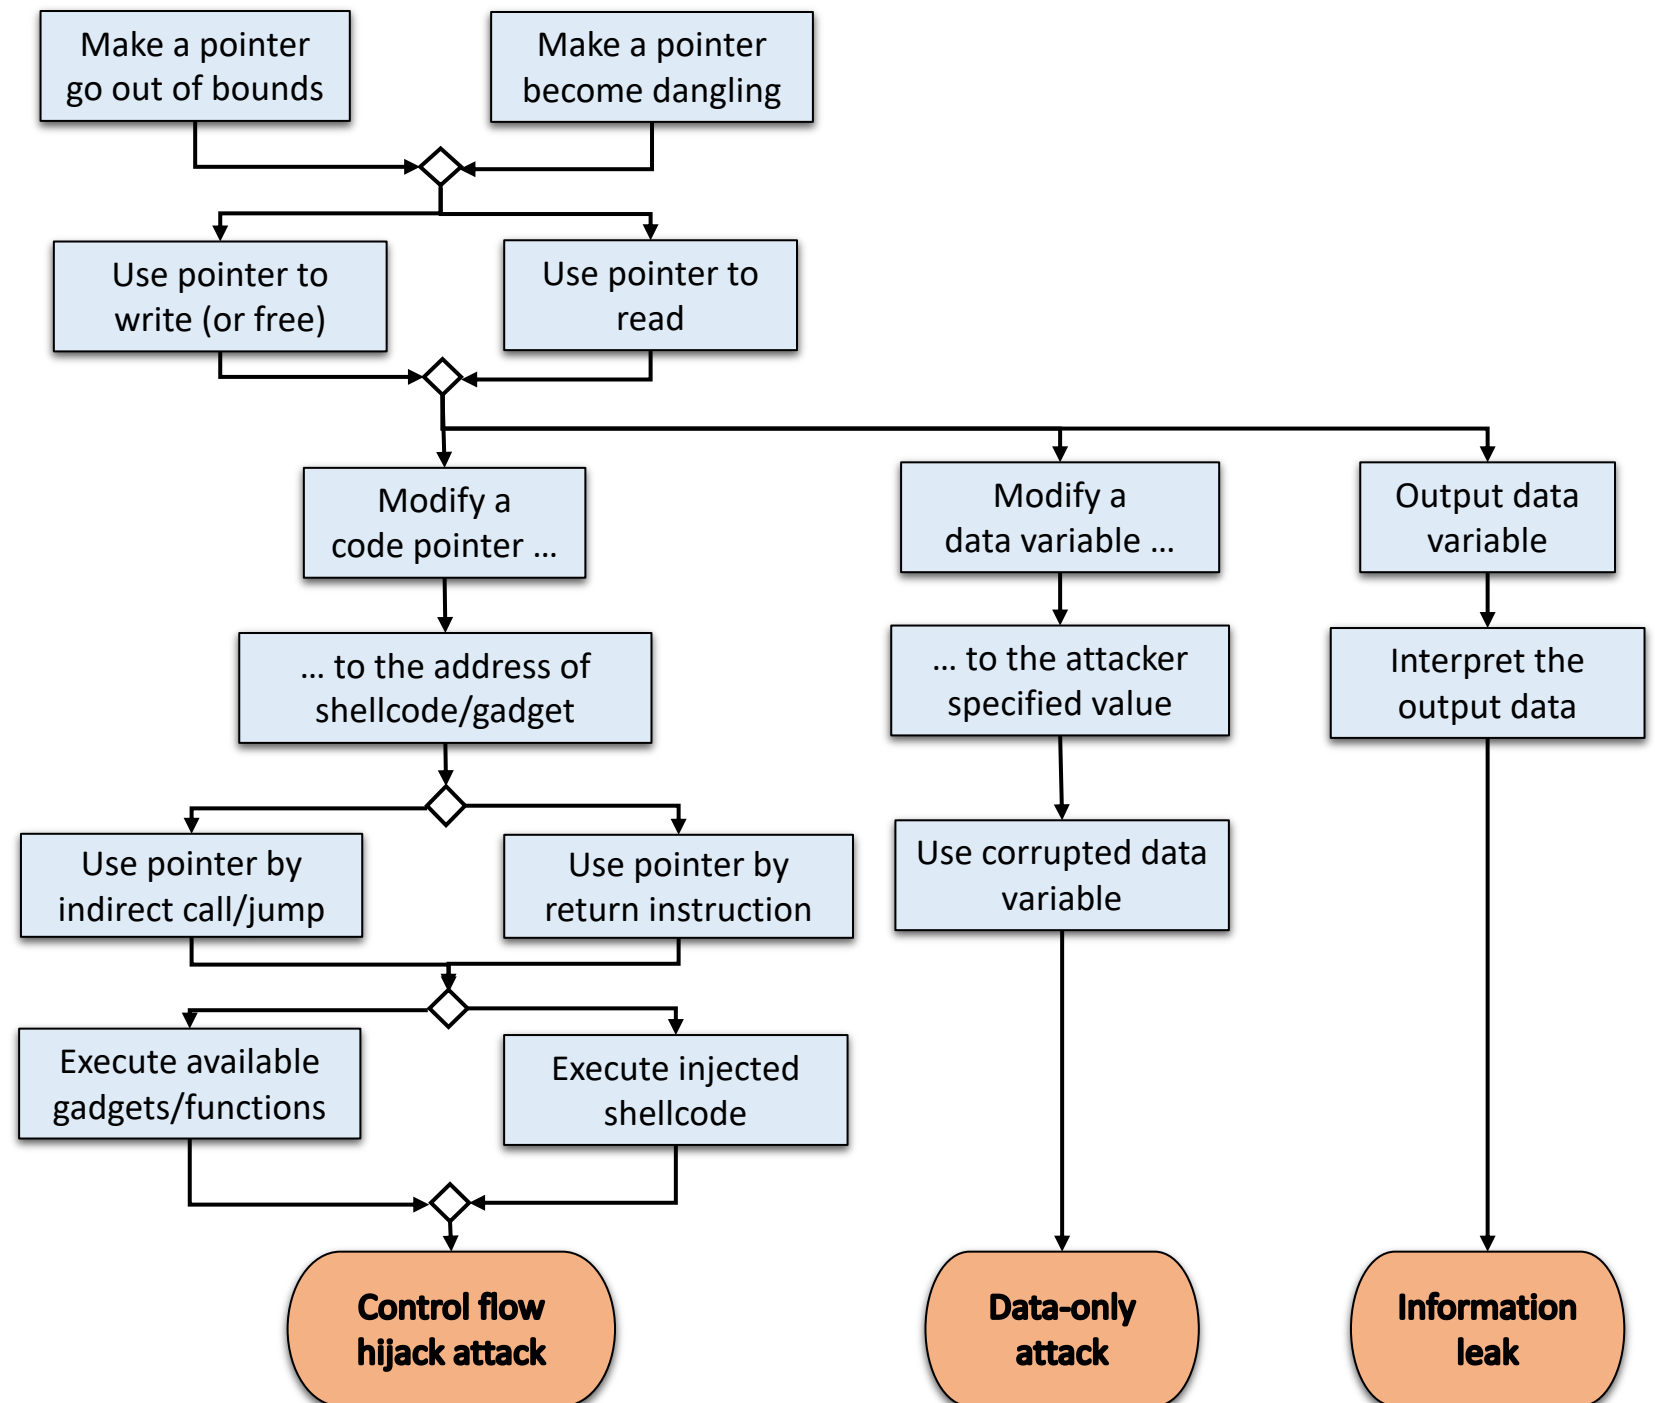
Data-only attack example:

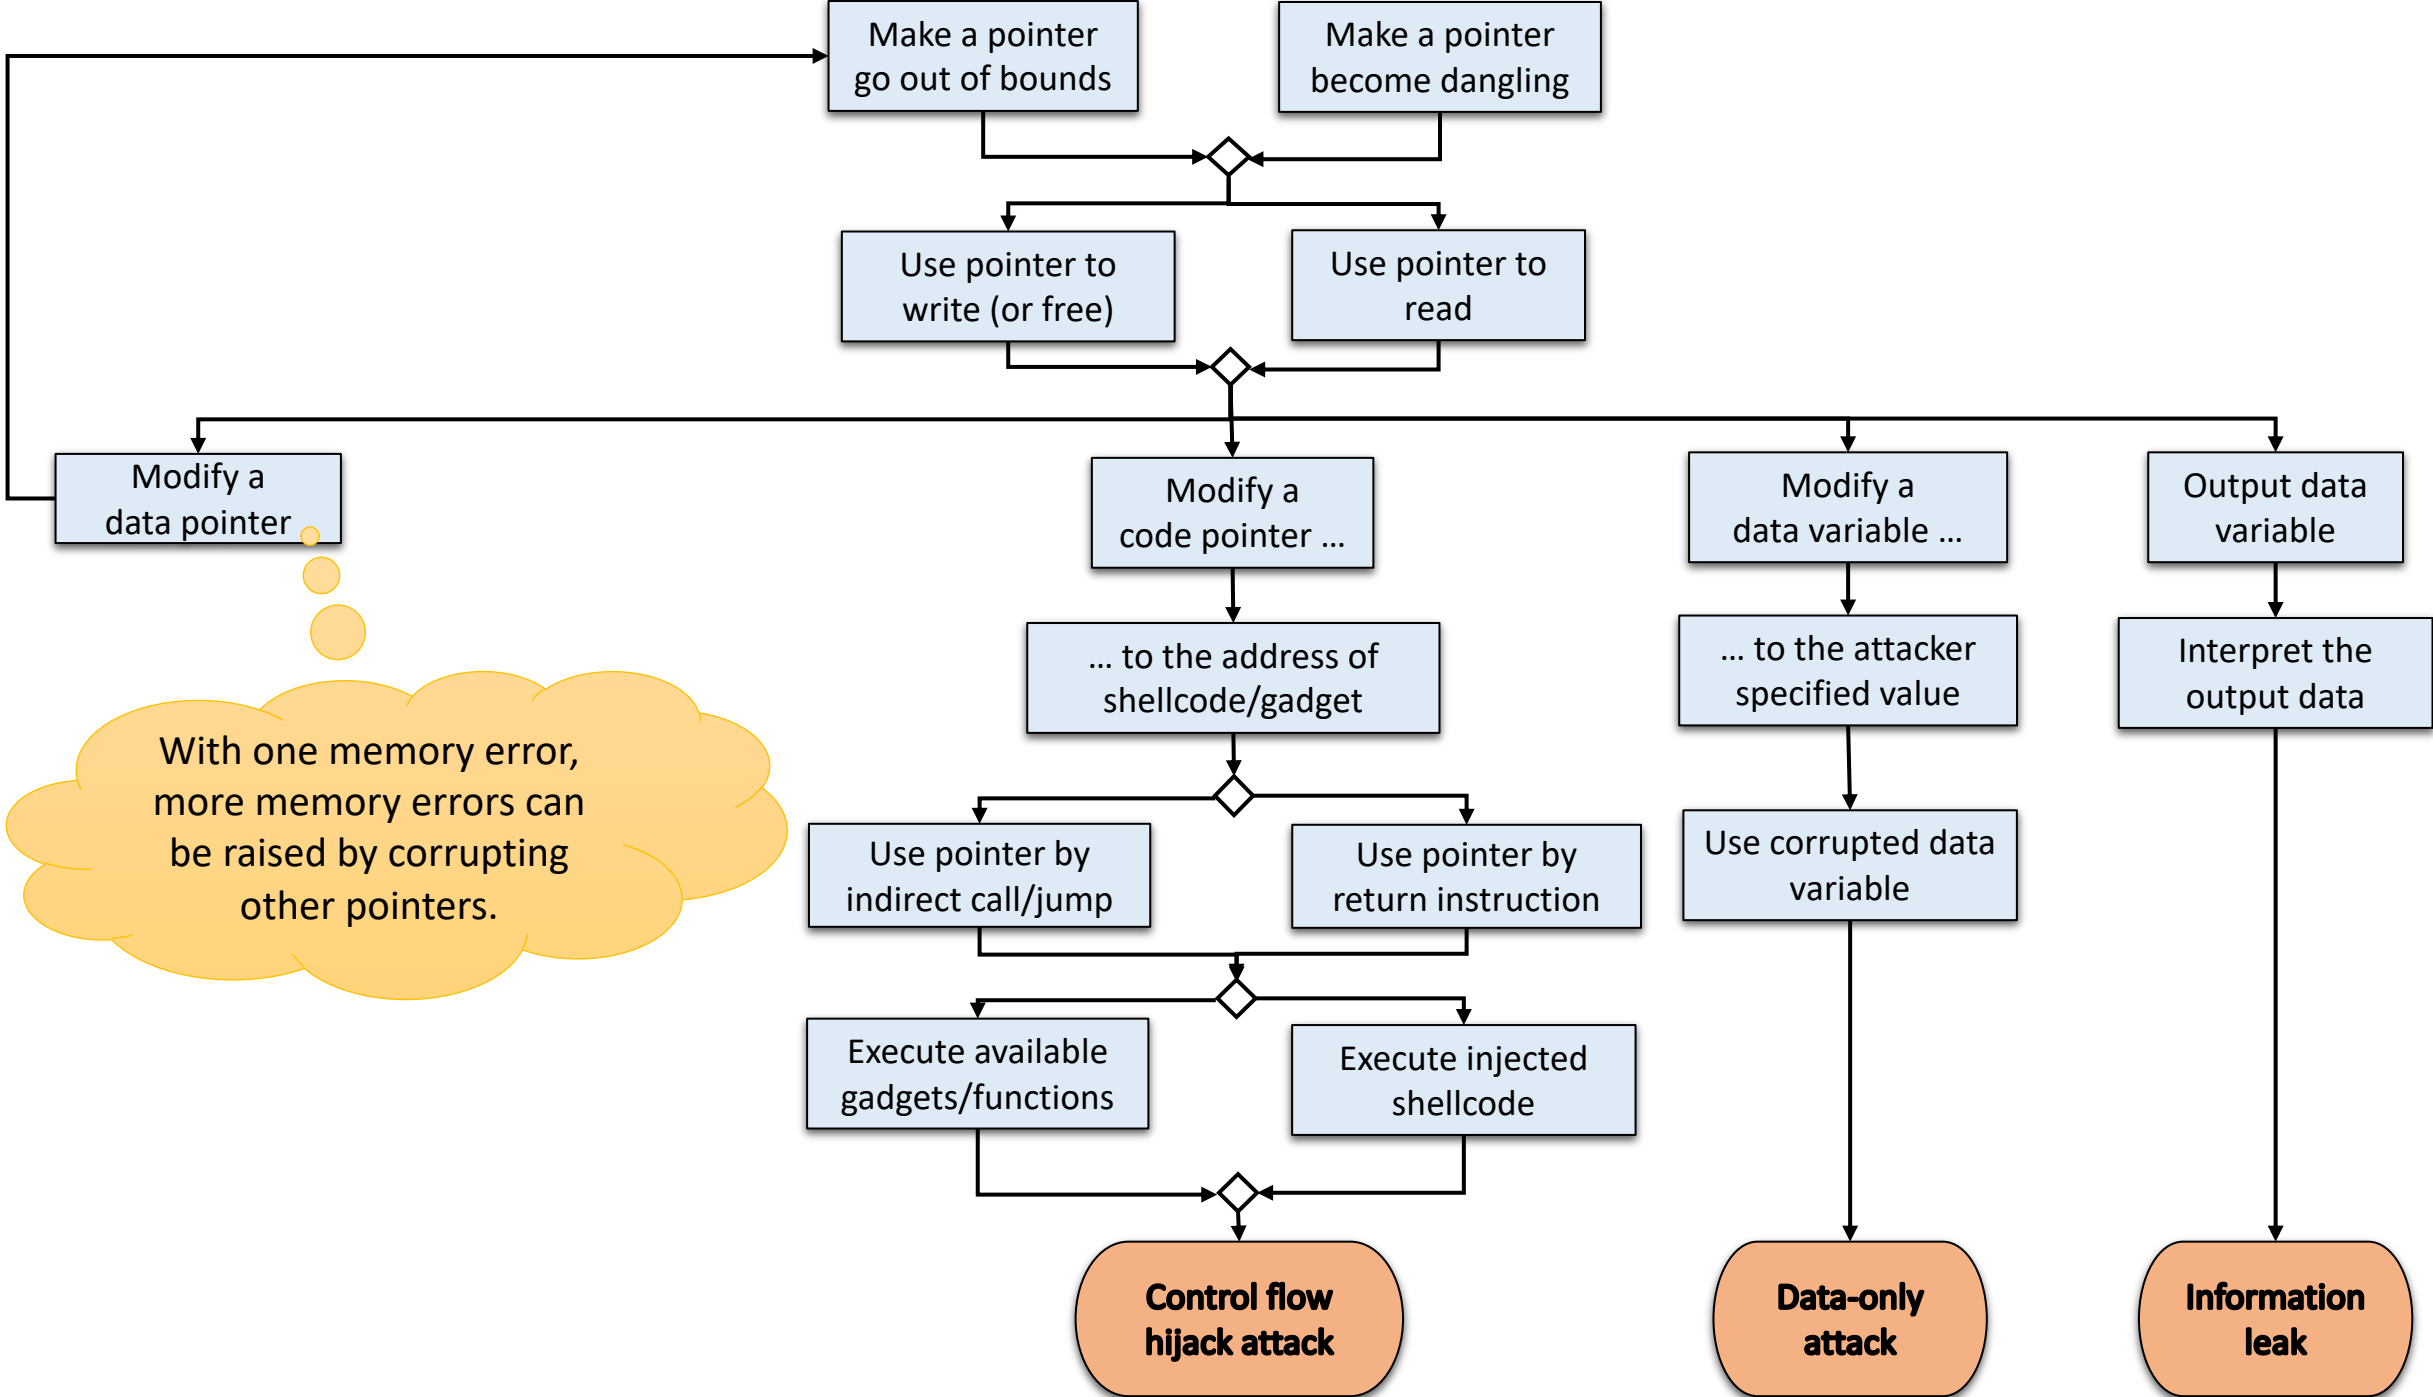
```
bool isAdmin = false;
```

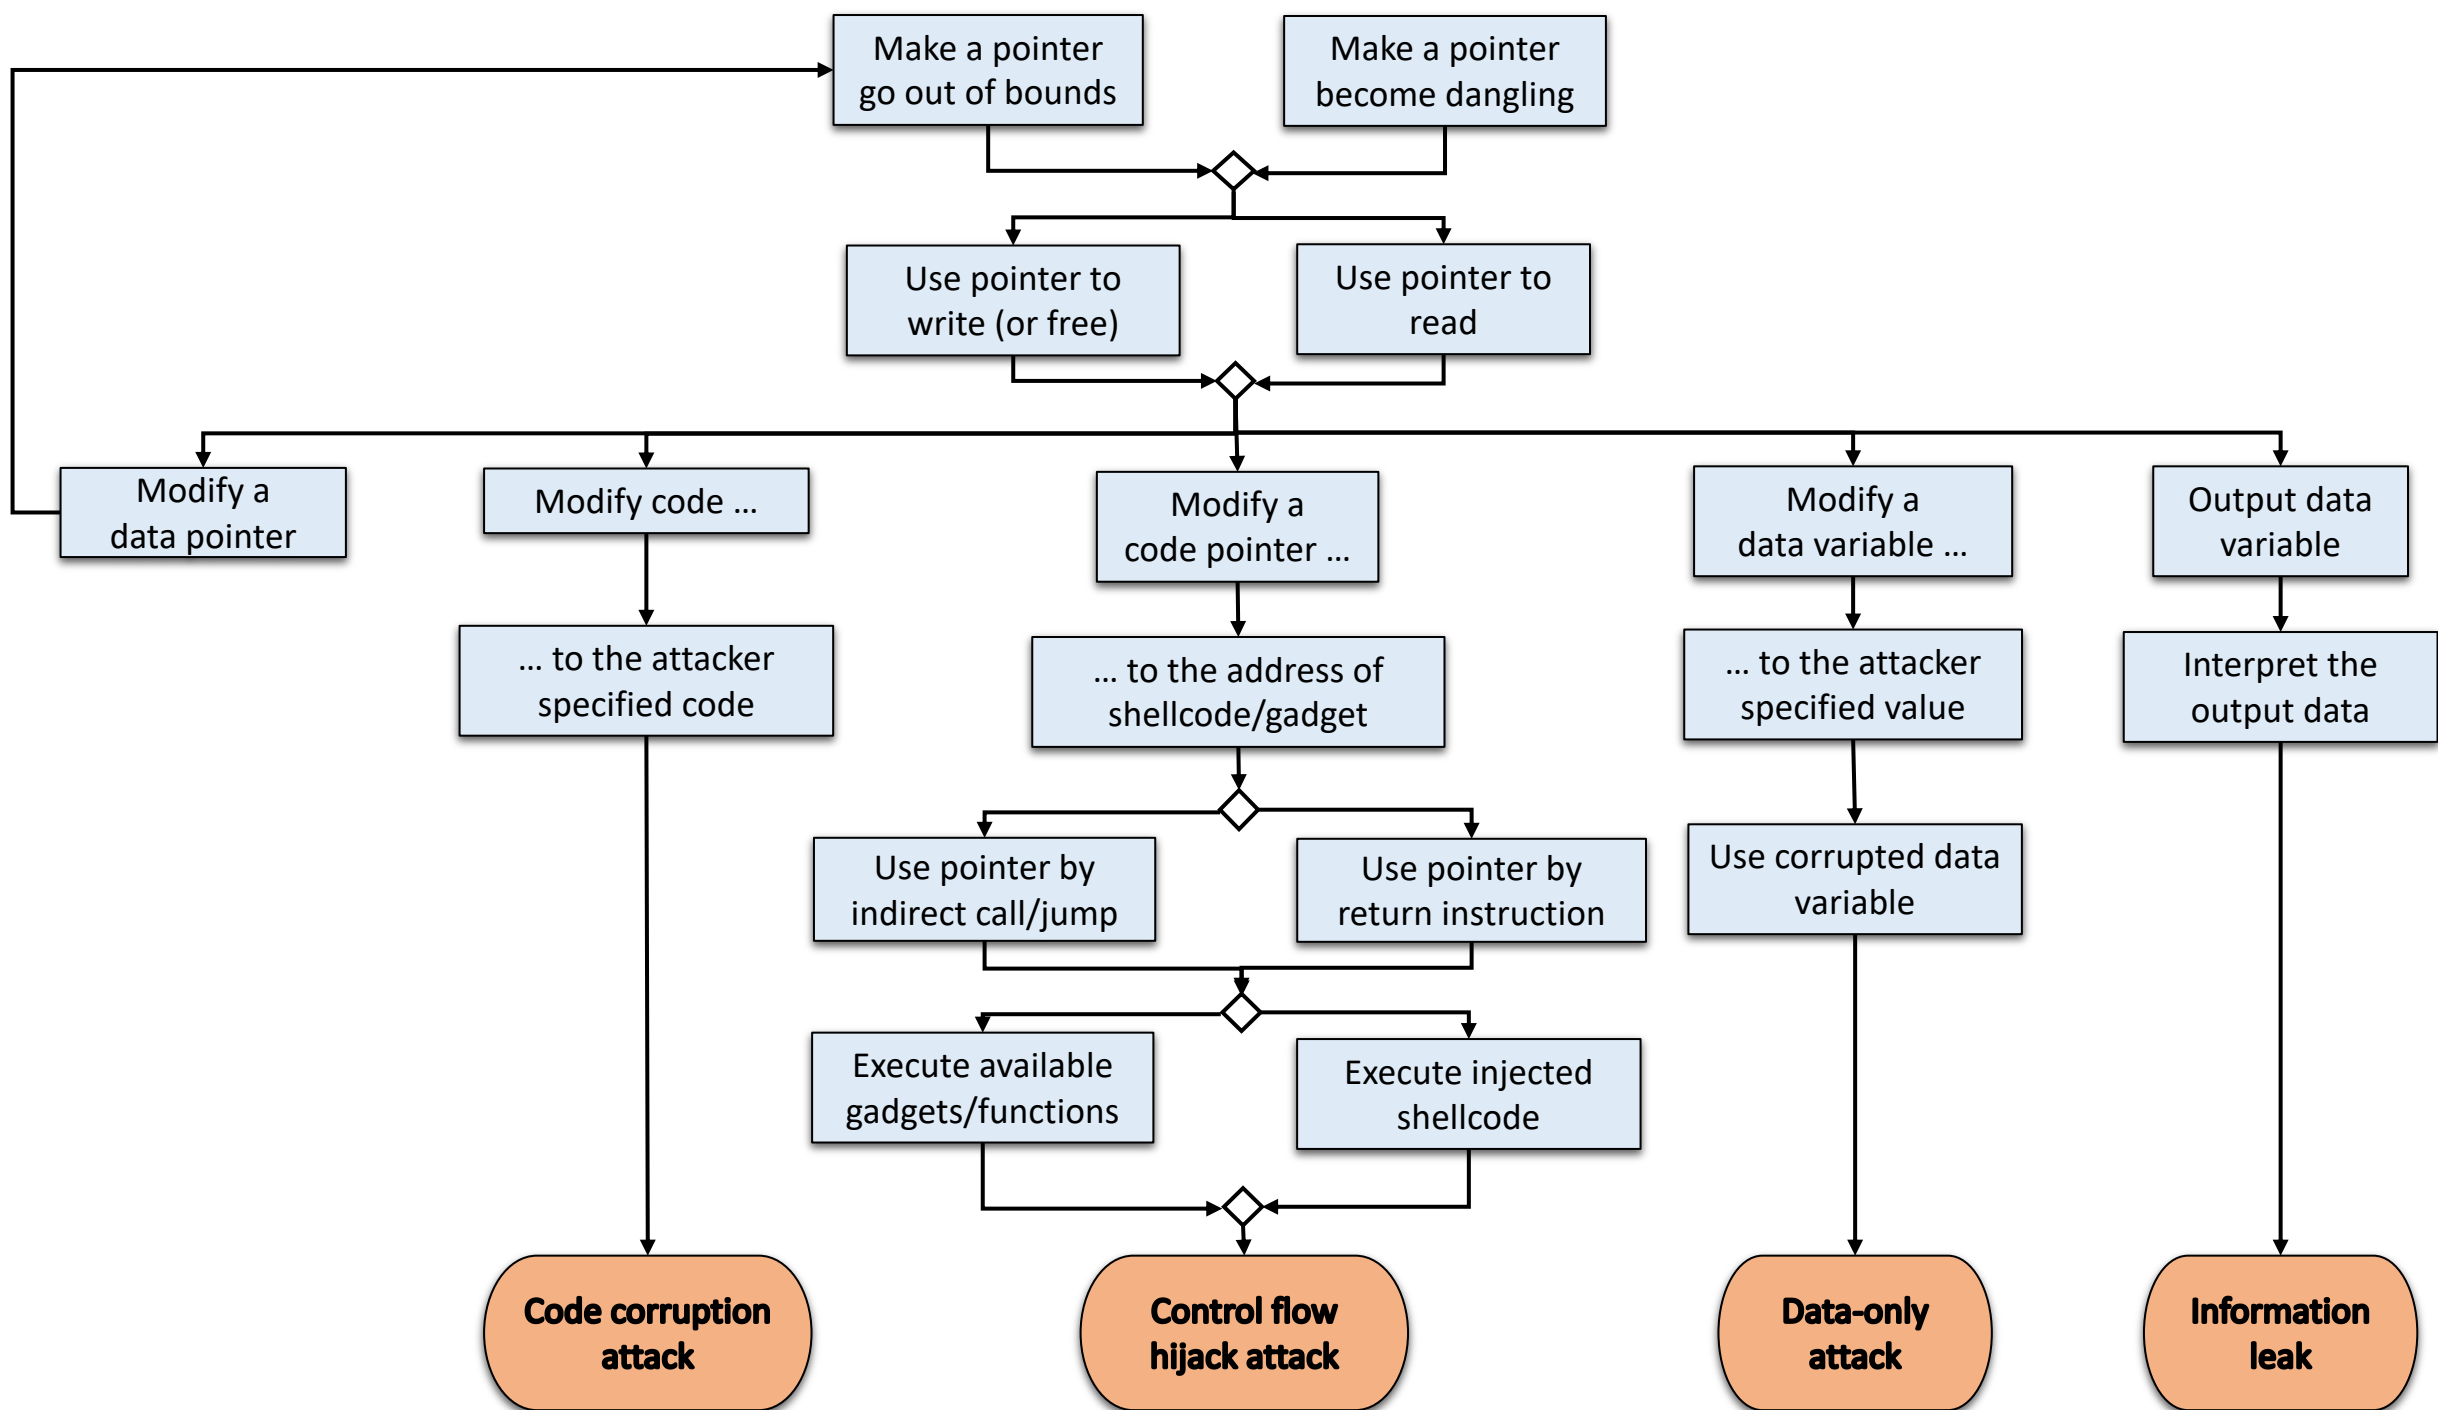
```
...
```

```
if (isAdmin)
```

```
// do privileged operations
```



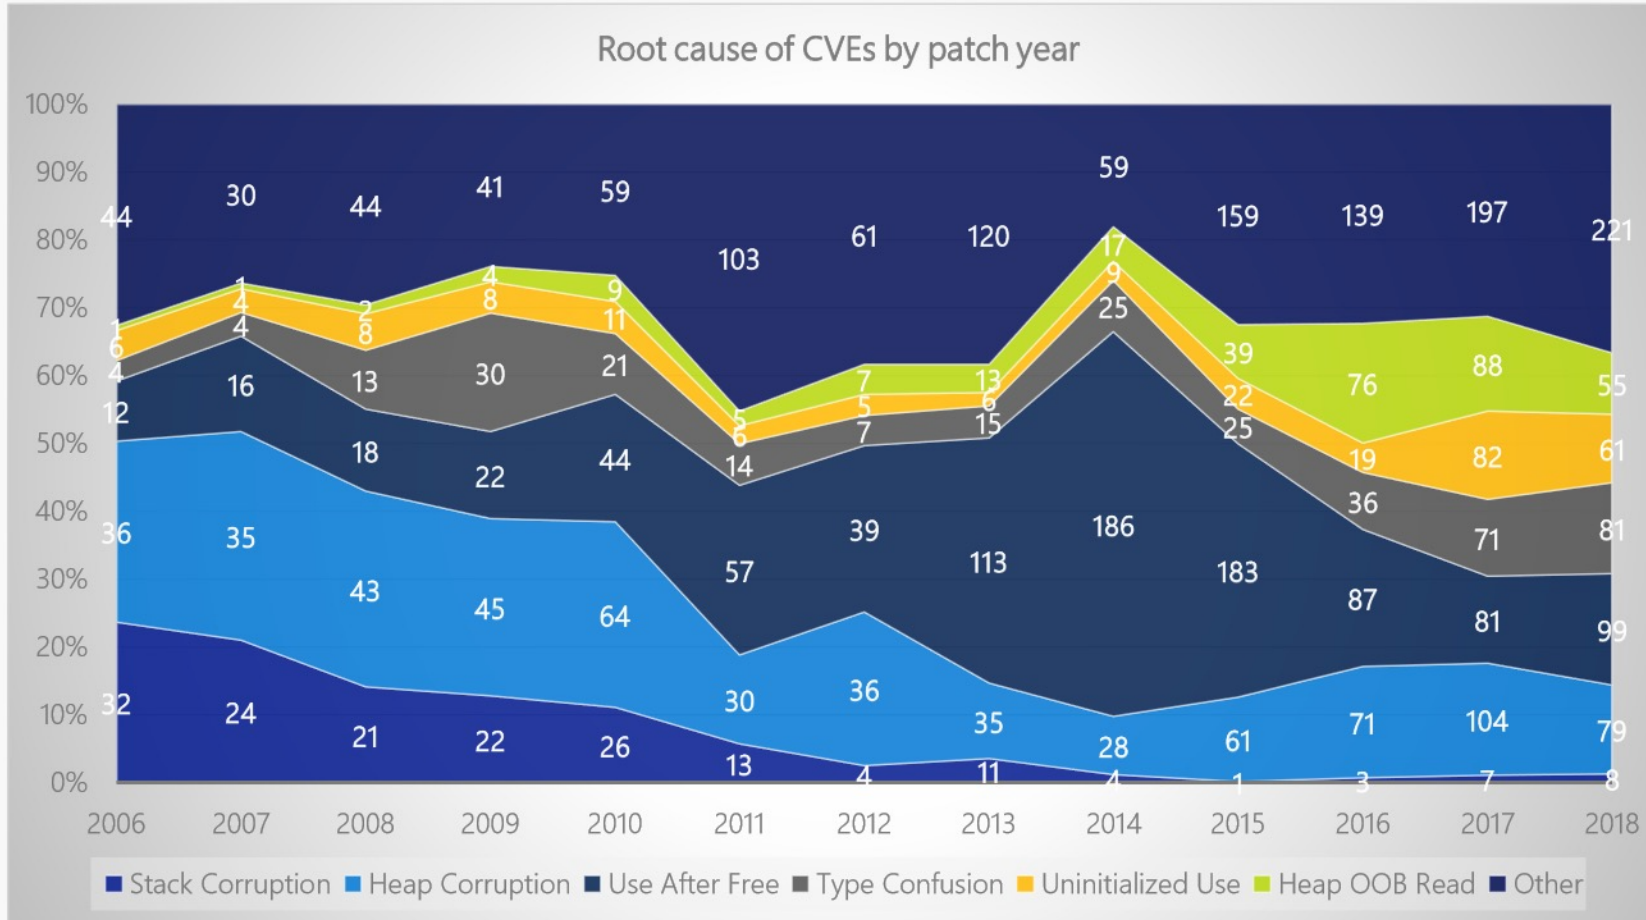




Drilling down into root causes

Trend reported by Microsoft

https://github.com/microsoft/MSRC-Security-Research/tree/master/presentations/2019_02_BlueHatIL



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

Why not High-level Language?

- Benefits:
 - Easier to program
 - Simpler concurrency with GC
 - Prevents classes of kernel bugs
- Downsides (performance):
 - Safety tax: Bounds, cast, null-pointer checks
 - Garbage collection: CPU and memory overhead, pause time
 - Feasibility?

The benefits and costs of writing a POSIX kernel in a high-level language; Cutler et al (OSDI'18)

BISCUIT: new x86-64 Go kernel

No fundamental challenges due to HLL

But many implementation puzzles

- Interrupts
- Kernel threads are lightweight
- Runtime on bare-metal
- ...

Why not Rust (no GC)?

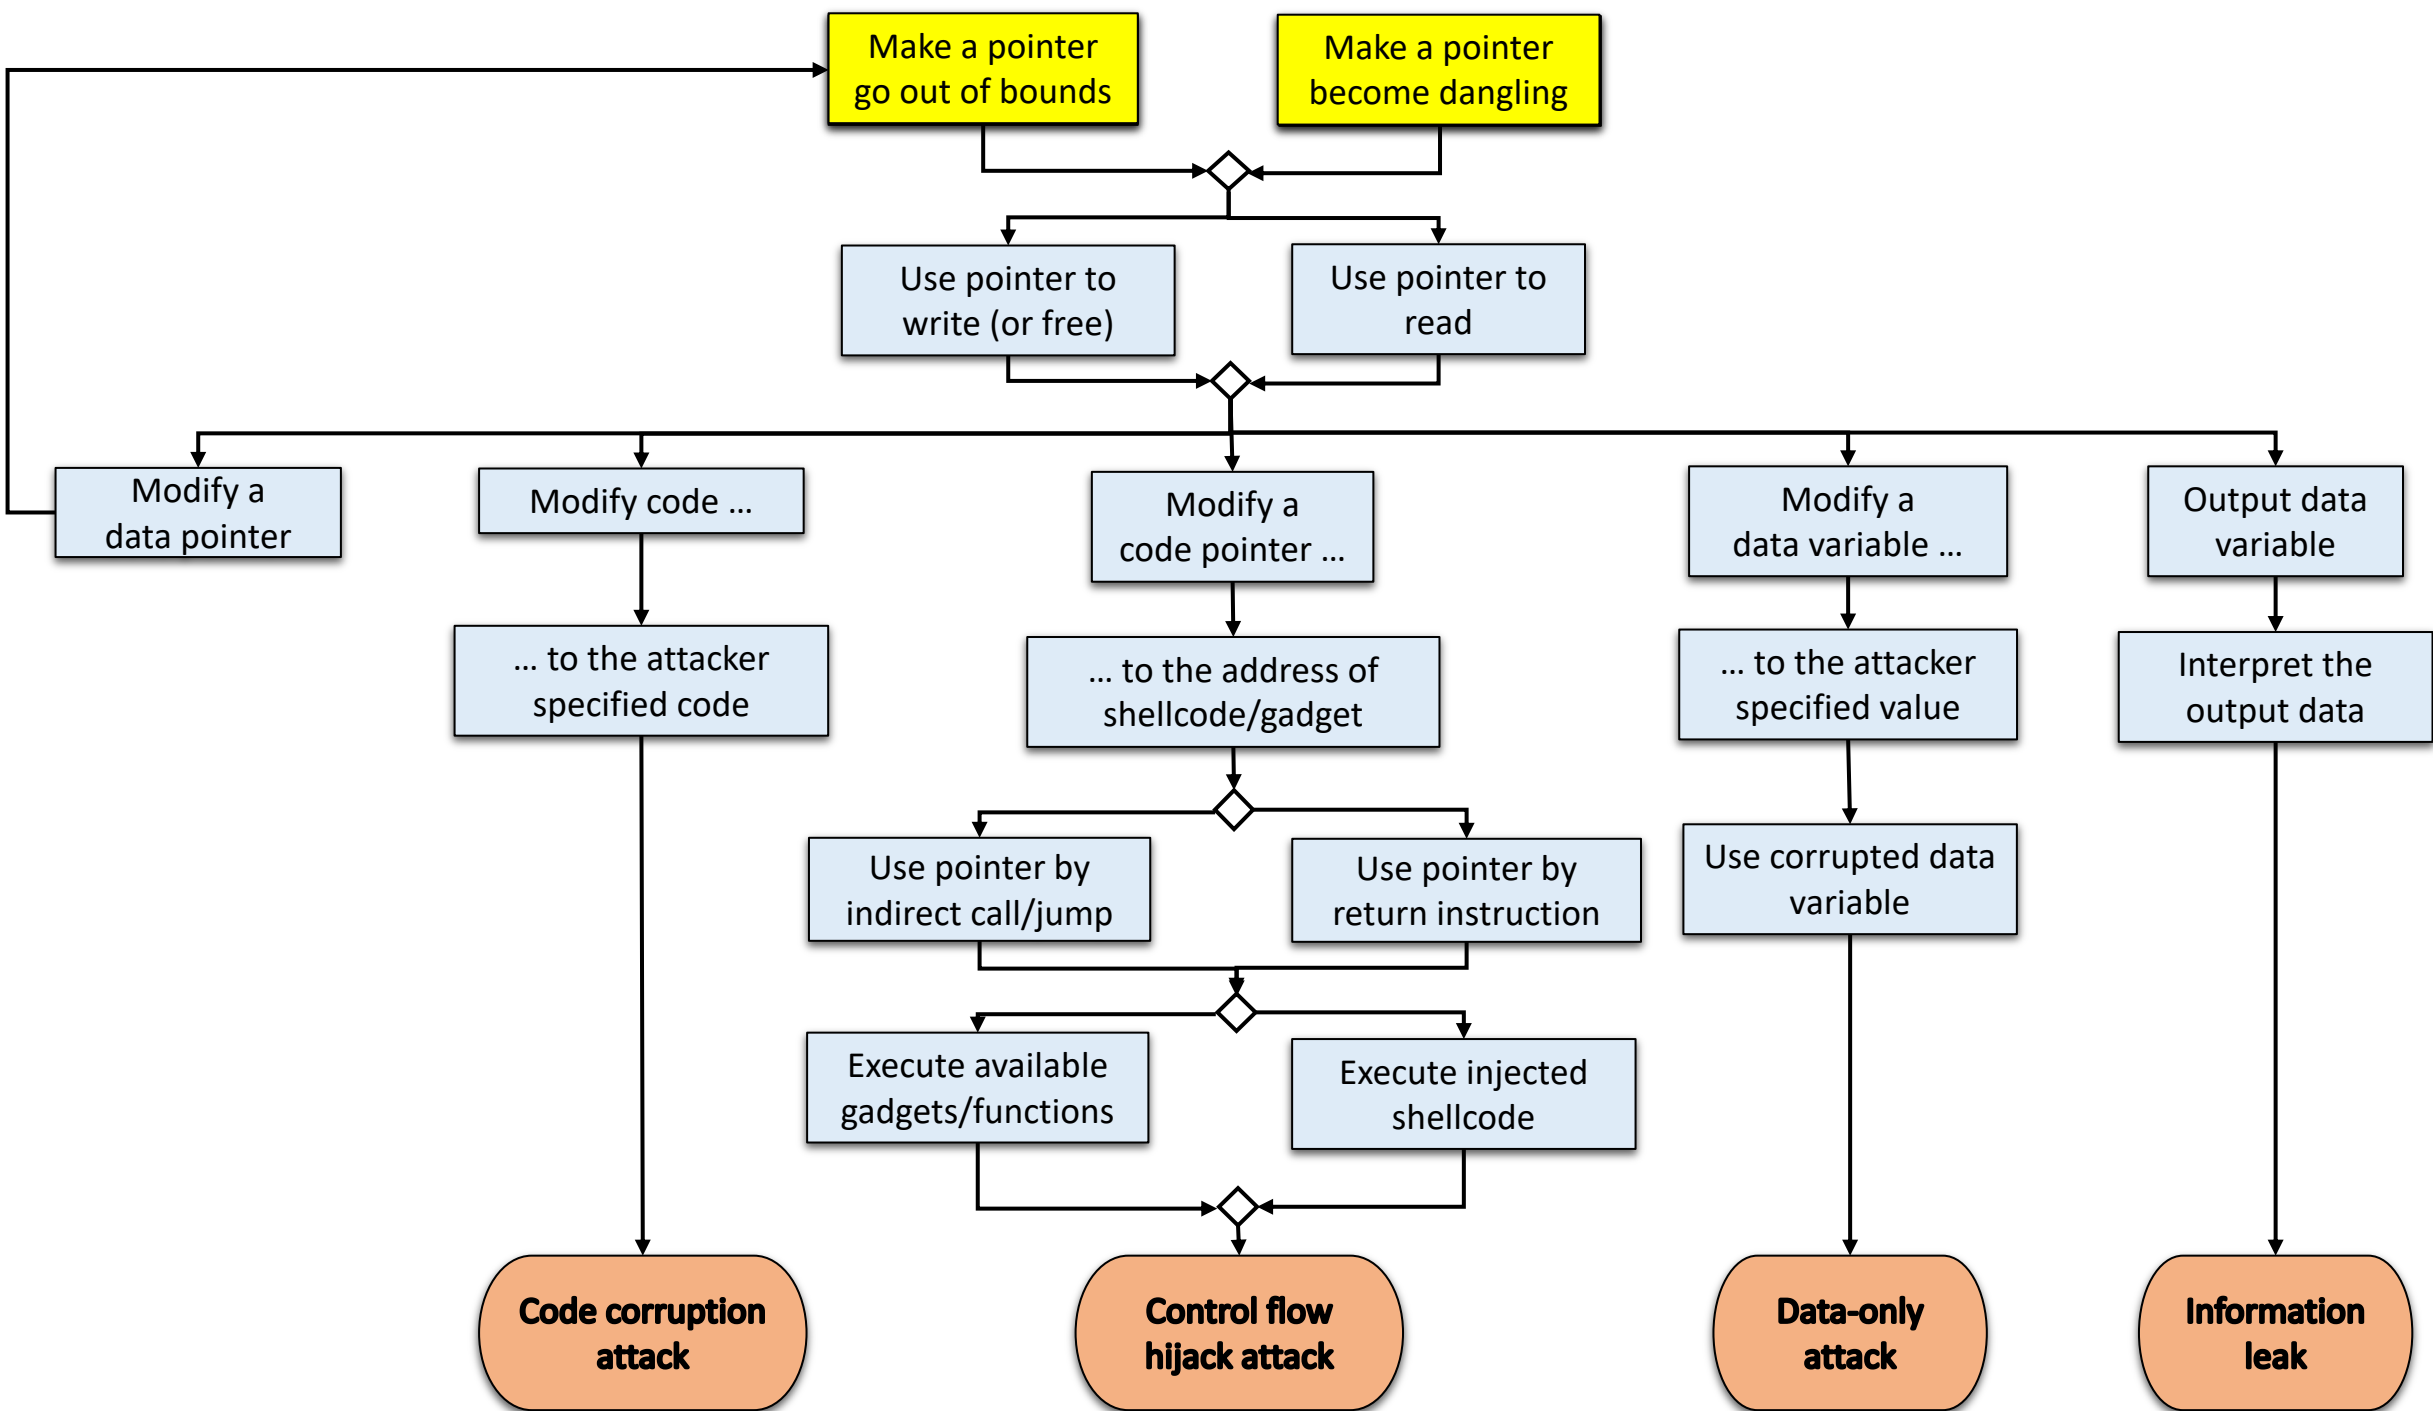
Rust compiler analyzes the program to partially automate freeing of memory. This approach can make sharing data among multiple threads or closures awkward

Surprising puzzle: heap exhaustion

HW Support for Memory Safety

- Spatial safety (bound information)
- Temporal safety (allocation/de-allocation information)
- Low-level reference monitor
 - SW approach: add checks → performance overhead (e.g., SoftBound)
 - Execution time: Extra instructions to perform the check
 - Memory: Maintain extra meta data (in shadow memory)

SoftBound: Highly Compatible and Complete Spatial Memory Safety for C; Nagarakatte et al; PLDI'09



Intel MPX (Memory Protection Extension)

4 bound registers (bnd0-3)

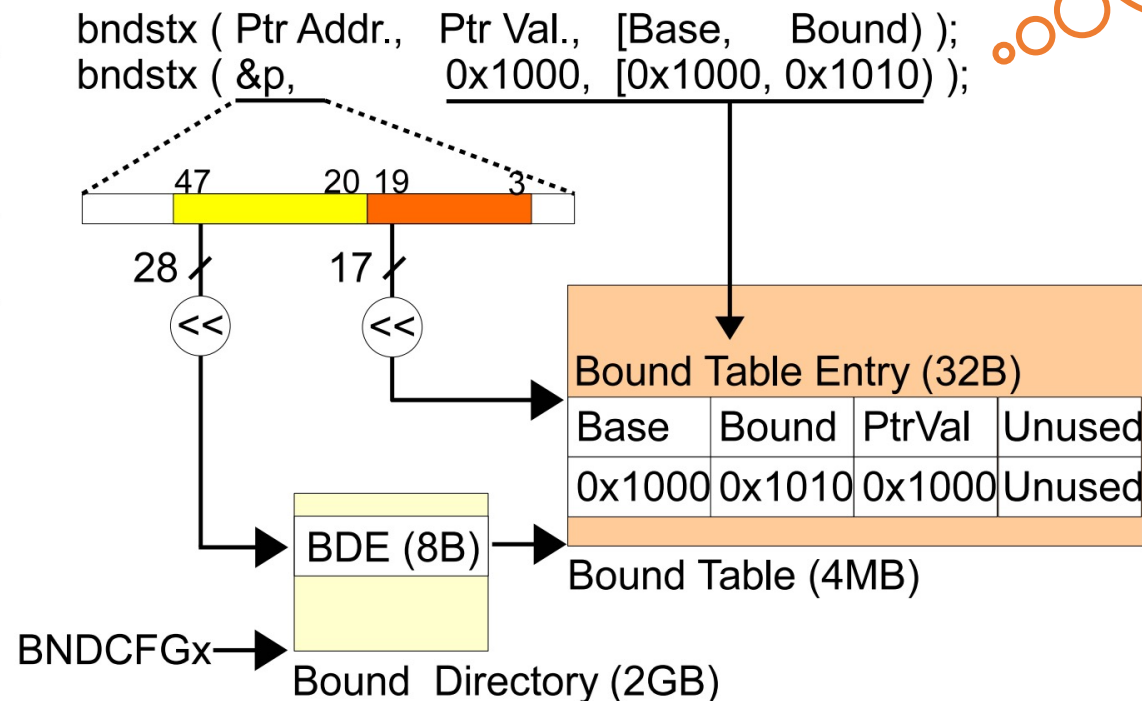
- **Bndmk**: create base and bound metadata
- **Bndldx/bndstx**: load/store metadata from/to bound tables
- **Bndcl/bndcu**: check pointer with lower and upper bounds

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';
```



Analysis of Intel MPX

Intel MPX is **impractical** for fine-grained memory safety

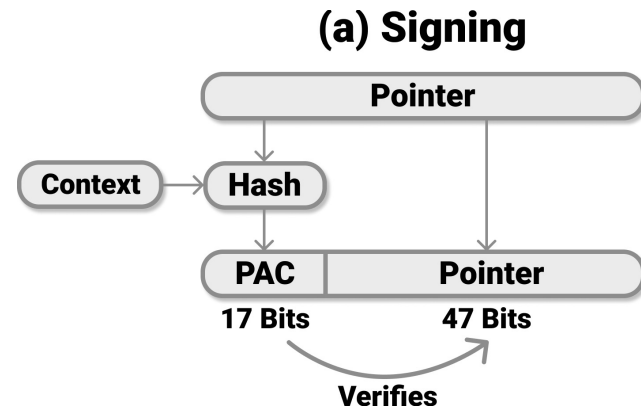


- High overheads
 - Check is sequential
 - loading/storing bounds registers involves two-level address translation
- Does not provide temporal safety
- Does not support multithreading transparently
- Meltdown? [Bound Range Exceeded \(#BR\) hardware exception](#)

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

ARM PA (Pointer Authentication)

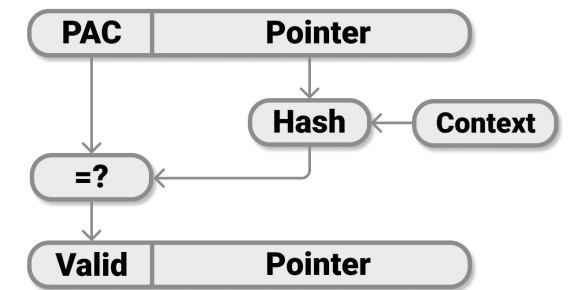
- Widely used in Apple processors
- Motivation:
 - 64-bit pointer, but 48-bit virtual address space
 - Unused high bits
- Hash:
 - A tweakable message authentication code (MAC)
 - ARM calls it **PAC (pointer authentication code)**
- Context:
 - secret key
 - salt (could be the stack pointer)



Before function call

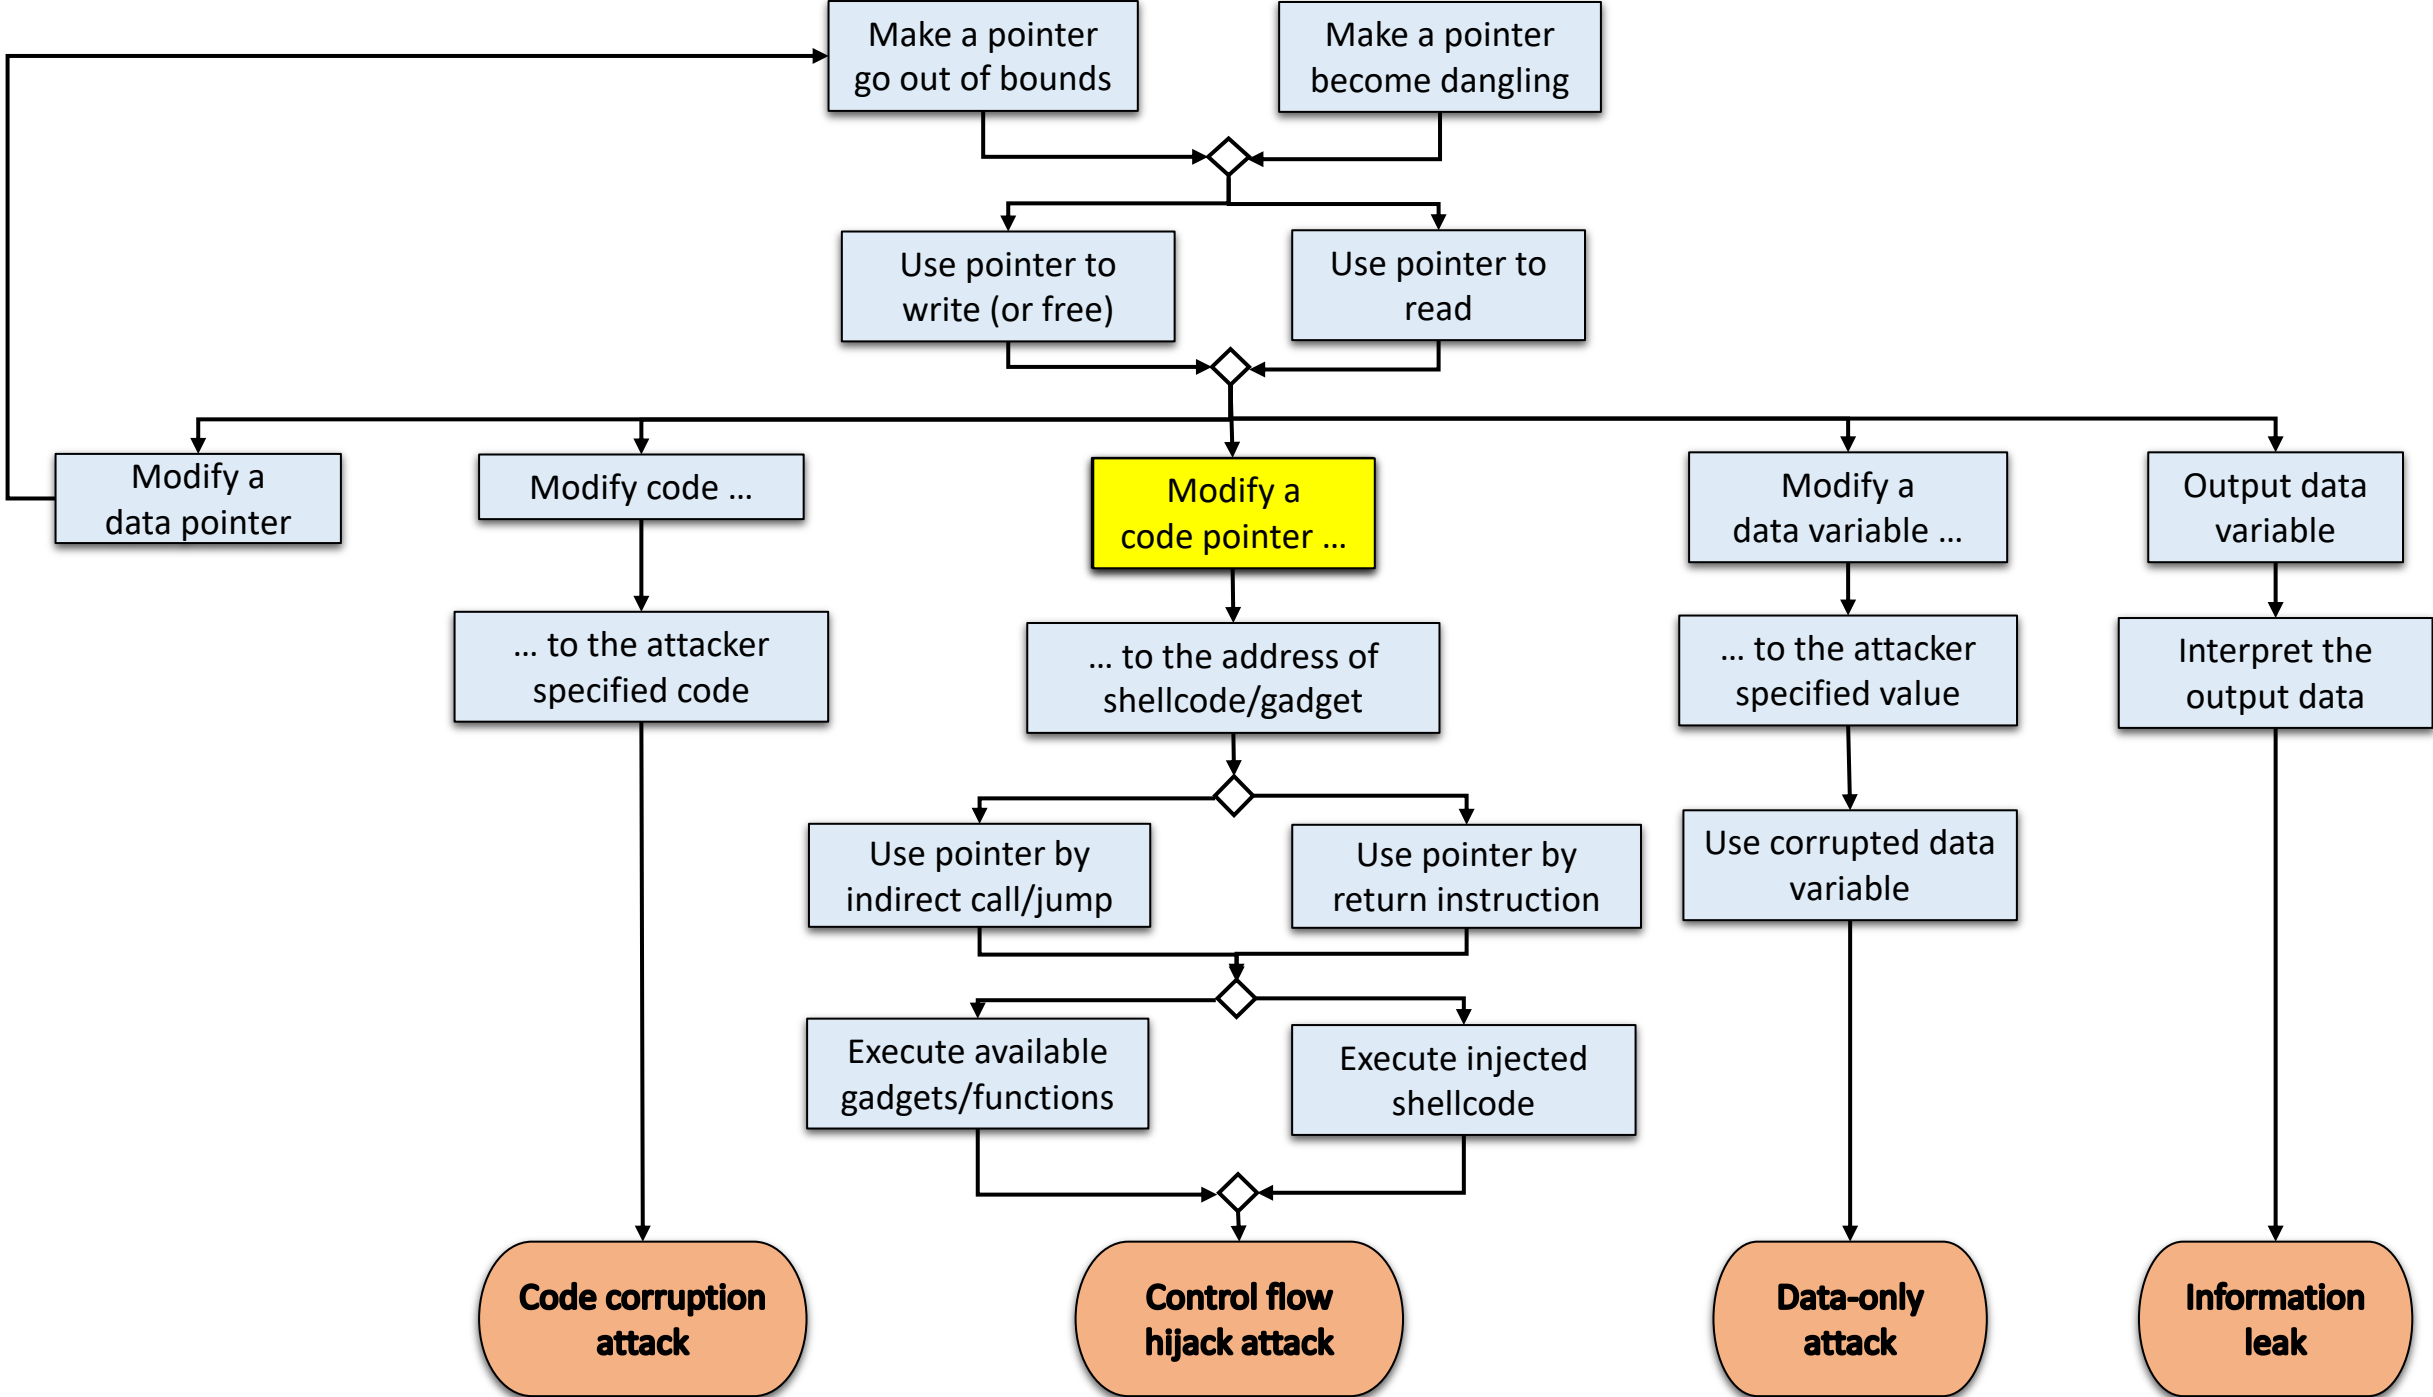
```
1 pacia lr, sp
2 sub sp, sp, #0x40
3 str lr, [sp, #0x30]
4 ...
```

(b) Verifying



Before function return

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



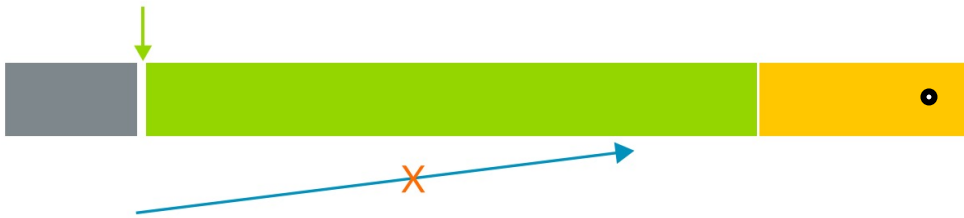
ARM MTE/Intel MPK

```
char *ptr = new char [16]; // memory colored
```



```
ptr[17] = 42; // color mismatch -> overflow
```

```
delete [] ptr; // memory re-colored on free
```



```
ptr[10] = 10; // color mismatch -> use-after-free
```

Armv8.5-A Memory Tagging Extension White paper
<https://security.googleblog.com/2019/08/adopting-arm-memory-tagging-extension.html>

- 2019, Google announced that it is adopting Arm's MTE in Android
- Memory locations are tagged by adding **four bits** of metadata to each **16** bytes of physical memory
- Where to store tags?
 - Pointer tag is stored in top byte of the pointer
 - Physical memory tag is stored in hardware

Any problems?

ARM MTE/Intel MPK

```
char *ptr = new char [16]; // memory colored
```



```
ptr[17] = 42; // color mismatch -> overflow
```

```
delete [] ptr; // memory re-colored on free
```



```
ptr[10] = 10; // color mismatch -> use-after-free
```

Armv8.5-A Memory Tagging Extension White paper
<https://security.googleblog.com/2019/08/adopting-arm-memory-tagging-extension.html>

- 2019, Google announced that it is adopting Arm's MTE in Android
- Memory locations are tagged by adding **four bits** of metadata to each **16** bytes of physical memory
- Where to store tags?
 - Pointer tag is stored in top byte of the pointer
 - Physical memory tag is stored in hardware
- Limited tag bits
 - Cannot ensure two allocations have different colors
 - But can ensure that the tags of sequential allocations are always different

Intel® Control-Flow Enforcement Technology (Intel CET)

INTEL
CET

=

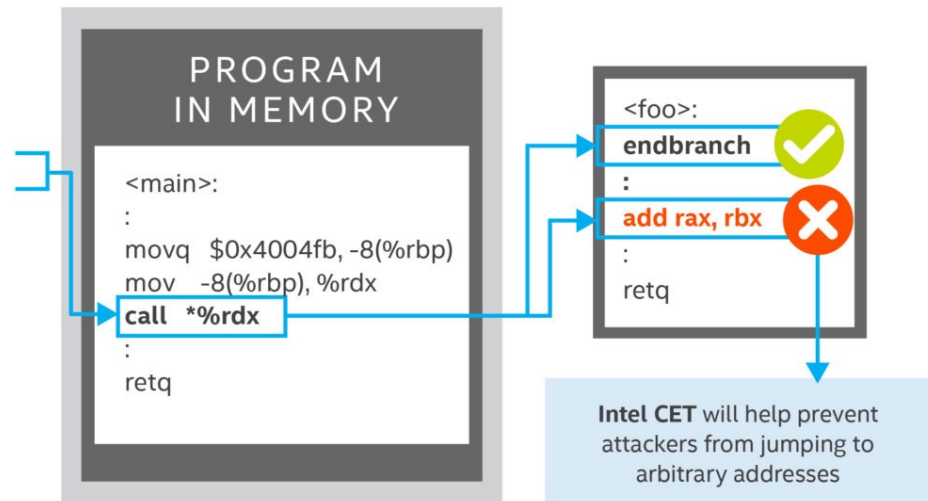
INDIRECT BRANCH
TRACKING (IBT)

+

SHADOW
STACK (SS)

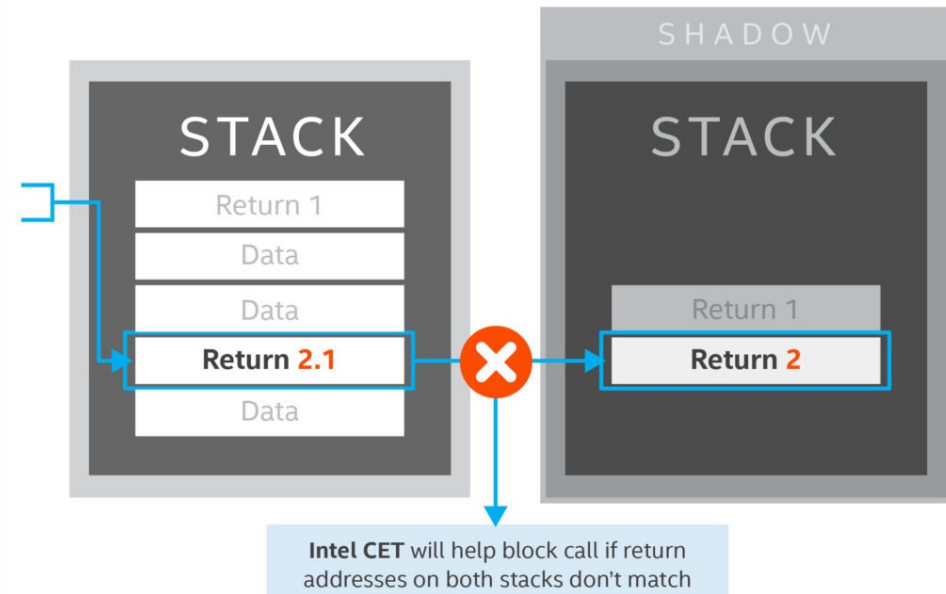
INDIRECT BRANCH TRACKING (IBT)

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.



100

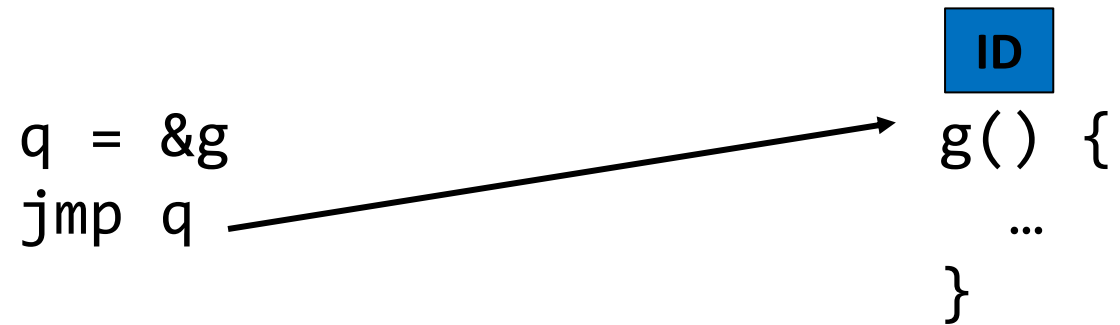
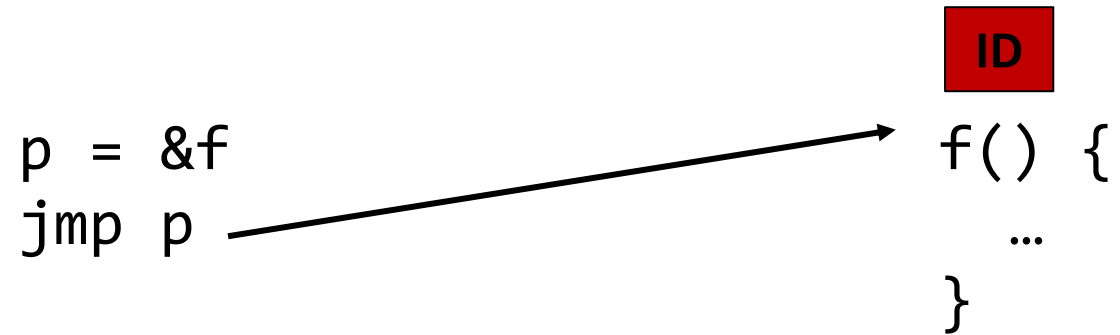
```
p = &f
jmp p
```

f() {
...
}

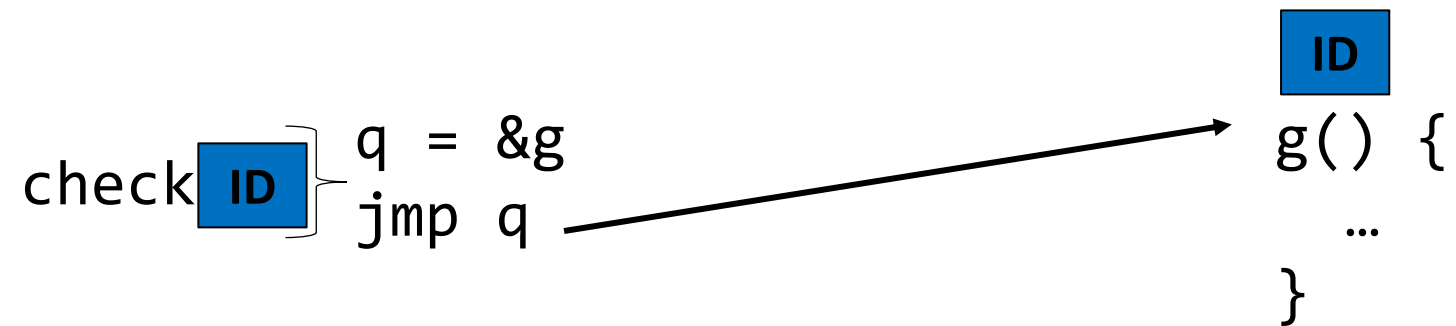
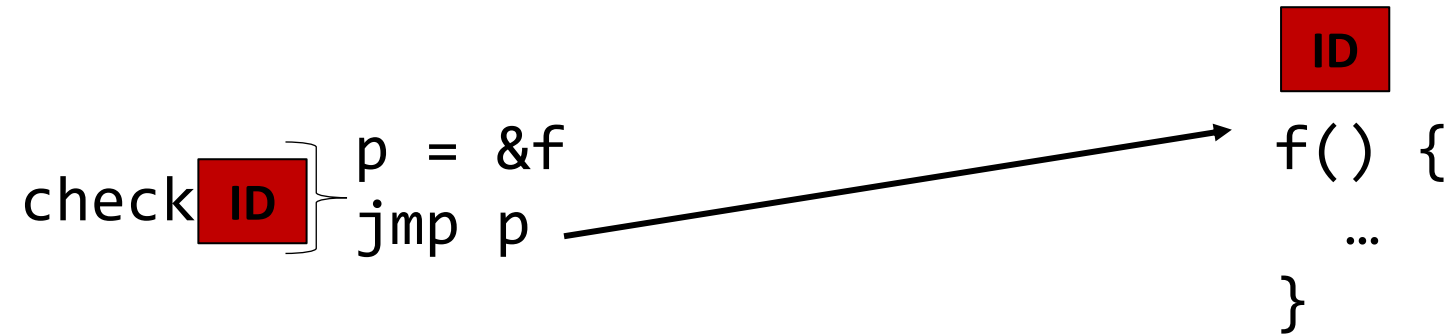
```
q = &g
jmp q
```

g() {
...
}

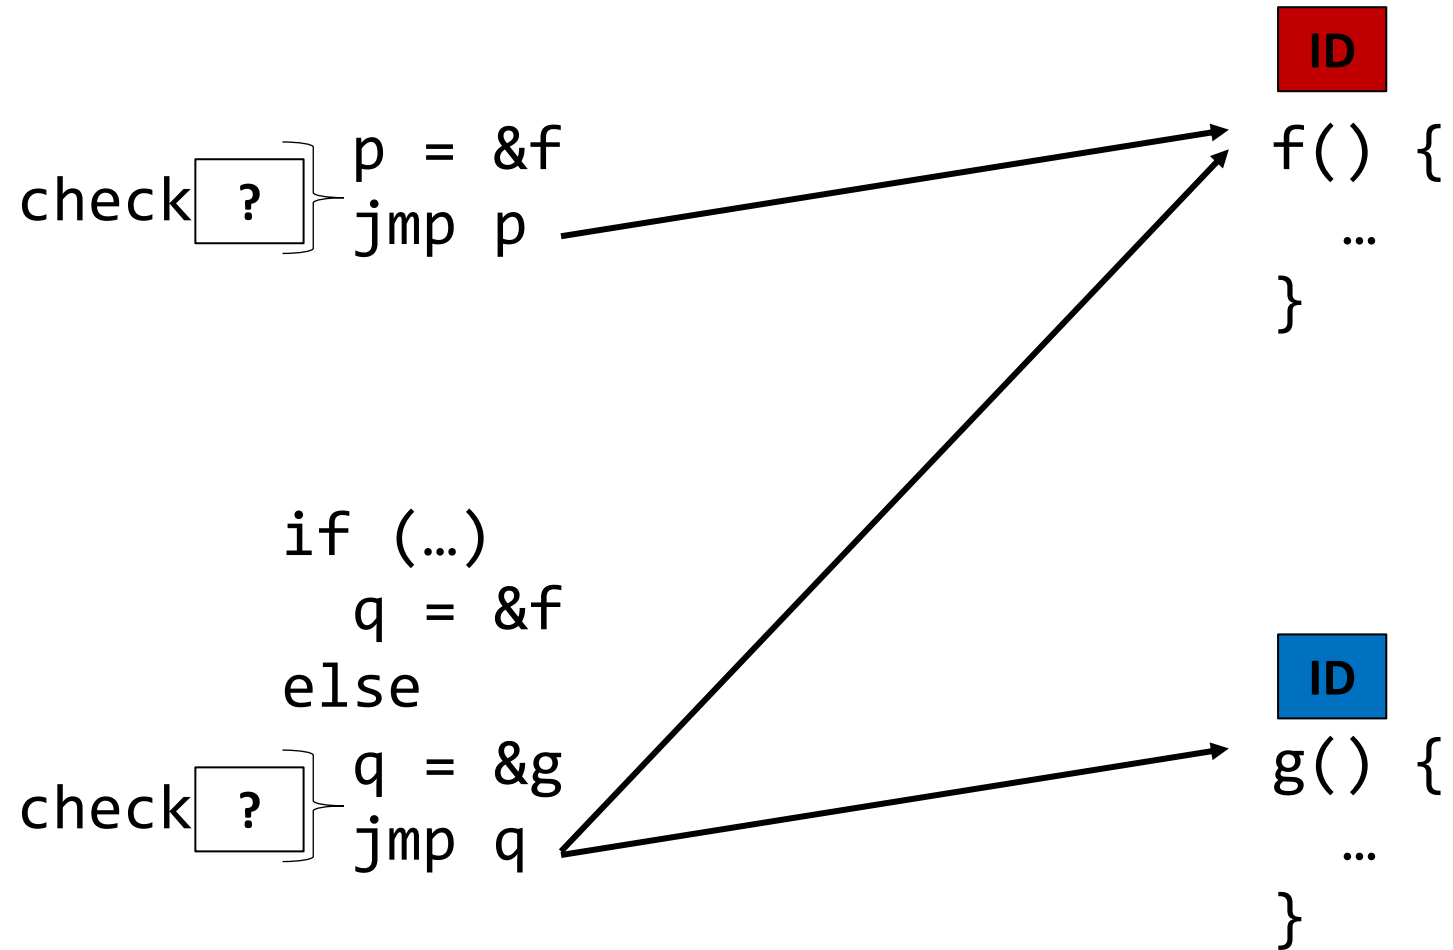
Control Flow Integrity (CFI)



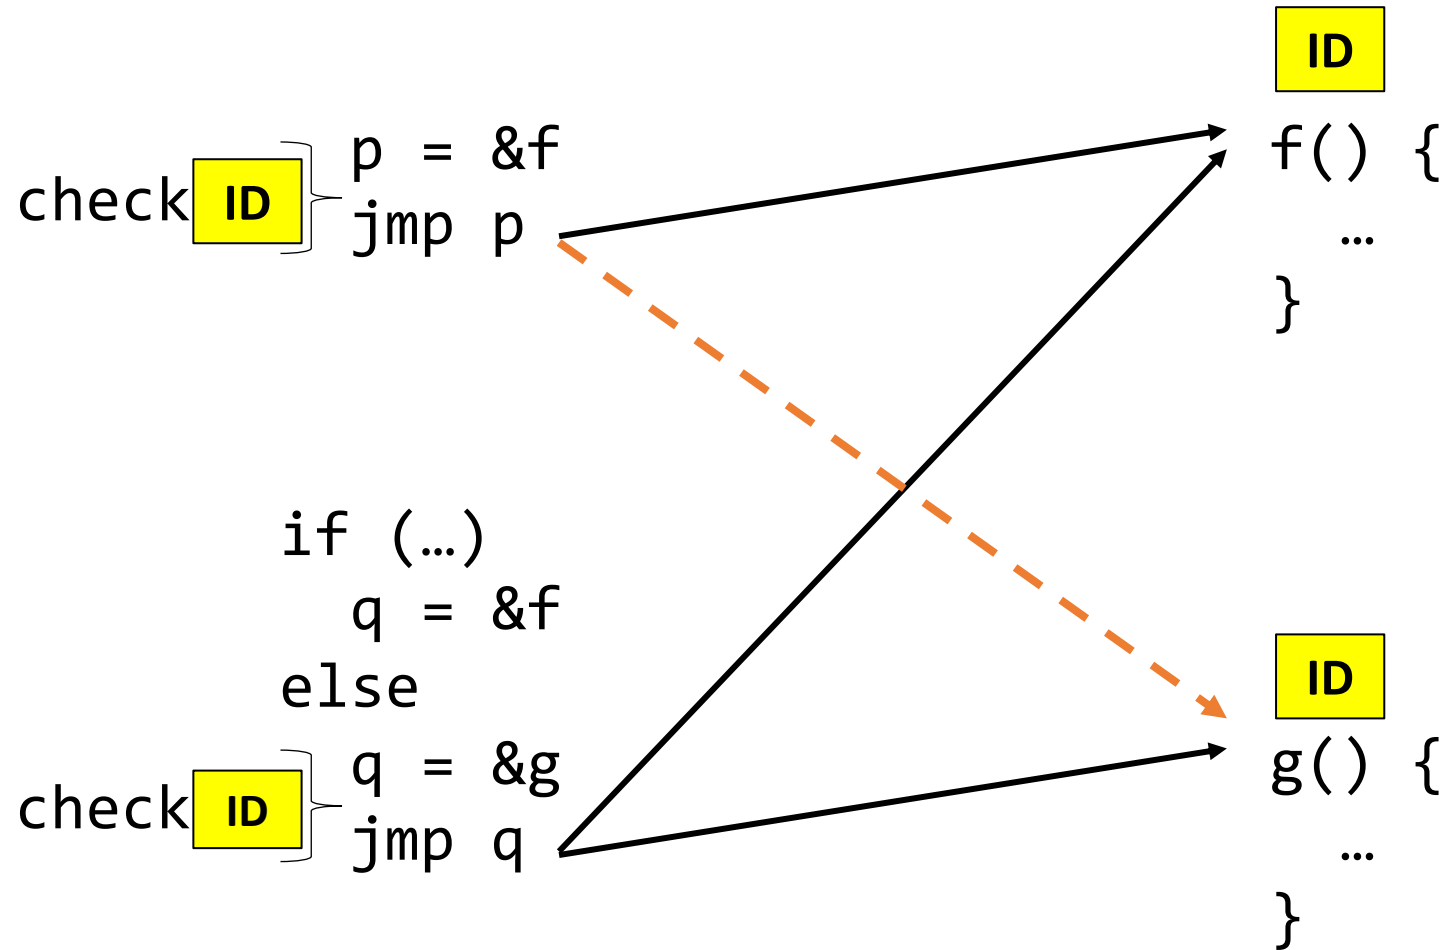
Control Flow Integrity (CFI)



Control Flow Integrity (CFI)



Over-approximation Problem



Summary

- Memory corruption problems: An eternal war
- Attack variations and mitigations
- Recent hardware support

