# Last time

Memory safety

Memory-safe language concepts

Safe unsafe languages?

2/12

# So... perfection?

## Write all software in a memory-safe language?

## **TCB considerations**

## Memory safety in compiled languages

- 1. Compiler-added run-time safety checks
- 2. Limited unsafety
- 3. Continued dangers of native instructions

3/12

High-level language interpreters have to be written in something. You might be able to write a lot of a Java interpreter in Java, but at the lowest levels you will find lots and lots of C++ code. At the lowest levels of the C standard library, you will find \_\_\_\_\_\_, sometimes

Languages like \_\_\_\_\_ and \_\_\_\_\_ claim to provide memory safety, but they are compiled languages. How is this possible?

The compiler can add extra code to check some accesses at run time. For example, if you are indexing within an array, the compiler can implicitly add code such as if  $0 \le i \le n$ .

Languages that aspire to "systems programming" (i.e., things that have to be aware of or manipulate the lowest-level primitives such as hardware registers) have to allow for unsafe operations. There is no memory-safe way to perform arbitrary register, memory or I/O operations, so these kinds of languages have to provide some way to break abstraction layers. C code can include assembly via the **asm** keyword. Rust code can explicitly violate memory safety guarantees if it uses the **unsafe** keyword.

Even with those checks, however, if you load someone else's native instructions and execute them, !

# Safe compiled code?

## What is a language?

#### Software

AddressSanitizer, CCured, Cyclone, "fat pointers", Go, Rust, ...

#### Hardware:

Arm MTE, CHERI, Hardbound, MPX, segmentation, Watchdog, ...

4/12

When we think of a language, we typically think about \_\_\_\_\_\_ and the \_\_\_\_\_\_ for writing it. However, in addition to \_\_\_\_\_\_, we also have \_\_\_\_\_\_\_ that are defined by language specifications and — crucially —\_\_\_\_\_\_. If we take this expanded view of what makes a language, we can see a number of approaches applied in various places that can be used to improve the security of compiled code, too.

#### Software

AddressSanitizer (and other "sanitizers" like Thread Sanitizer and the Undefined Behaviour Sanitizer) can help spot memory errors during testing that might otherwise have gone unnoticed. CCured is an example of an approach that uses static analysis to figure out how pointers in a C program are "meant" to be used and dynamic analysis to ensure that they are, in fact, used that way. Cyclone is a C dialect with better memory safety properties than vanilla C, which it is designed to be compatible with (or at least easy to adapt from). Newer languages like Go and Rust have more expressive type systems that make it possible to write memory-safe code even in highperformance compiled languages with limited run-time checking.

#### Hardware

Arm MTE has been adopted by Android to detect memory safety violations at run time. Hardbound, MPX and Watchdog attempt to provide various forms of hardware memory safety enforcement. CHERI is a designed-for-security instruction set extension for ARM and MIPS that is just about to ship its first hardware prototypes; it has the potential to change by allowing high-level object accesses to be precisely enforced by hardware.

# Today

## Finding memory safety violations

- testing
- formal methods
- hybrid approaches

5/12

# Testing

# Default approach... if we're lucky!

## Code coverage

- what's that?
- gcov, llvm-cov, SanitizerCoverage...

## Limitations

6/12

Hopefully software developers are employing	g practices such as	
Testing software helps us find defects, but it	s more valuable than just	that: designing software to be
testable also tends to encourage	via	with
, all of which	also help with software qu	ality. So, testing is good in
more than one way, and we know we ought	to do it! Let's be honest, t	hough: we don't always do
what we know we ought to do. And even wh	nen we do,	?
We'll have a little demo involving code cover	rage data, how it's generat	ed and how it's used.
Testing is great, but ultimately testing can no	ever	: it can only

# Formal methods

## Modeling programs

# Proving properties of programs\* and compilers†



## **Curry-Howard Correspondence**

* Can prove more than you'd think SIGNPLAN Notices 42(1), 2007 (DOI Communications of the ACM 54(5), † e.g., Kästner et al., CompCert: Pro-	c, e.g., Cook et al., Proving that programs eventually do someth I: 10.1145/1190215.1190257) and Cook et al., Proving Program , 2011, (DOI: 10.1145/1941487.1941509). actical experience on integrating and qualifying a formally veri production of formation and contents and provide the second se	ing good, <i>ACM</i> Termination, tied optimizing
compiler , ERTS 2018: Embedded F	Real Time Software and Systems, 2018.	8/12
Constructing a model of a high	n-level functional program is one thing: after all, func	tional
programming languages take th	heir cues from mathematics! Constructing a model of	a low-level
program in a language with lot	and pointer-based	,
however, is something else.		
People do, however, use formal	l methods in real (albiet someone size-limited) system	s. The seL4
microkernel has been formally	verified by first verifying a model, then generating co-	de <i>from</i> that
model, then proving that the g	enerated code corresponds <i>to</i> the model. That's not qu	uite proving
properties about C code, but it	t's close!	
In addition to proving property	ies about software in its source code, one practical (th	ough
computationally-challenging) a	approach is to prove properties about software in its	-
form. This is u	useful for the artifact that ultimately matters most:	
To e	do this, you need a formal model not of your software	e, but of the
hardware it will execute on.		
People also use	to prove that the output of a compile	r really
matches its inputs. This is mean	nt to address the problem raised in Ken Thompson's	famous Turing
Award lecture, Reflections on t	rusting trust. COMPCERT is such a compiler altho	ough every new
release "fixes a few bugs"??		
If you like monads, \$\pi\$-calcu	ılus or intuitionistic modalities or just	
! Why are functio	nal programs easier to verify? Because the language	
Tł	he more sophisticated the type system, the fewer prog	rams you can
	e programs you <i>can't</i> write are the	
write with the caveat that the		ones. For a lot

# Hybrid approaches

## Concolic testing:

- model checking
- concolic execution

## Fuzzing



9/12

*Concolic* is a portmanteau of \_\_\_\_\_\_ and \_\_\_\_\_. It aims to be practical (the concrete execution part, not just all theorems) while also gaining some of the generality of symbolic approaches. *Model checking* is heavily used by hardware engineers to verify (not just \_\_\_\_\_\_) things like state machines.
Concolic execution can be applied to real software using tools like KLEE. In such testing (for which we'll have a brief demo if there's time), we can run real test cases while telling the tool (KLEE in this case) that it should treat some values as \_\_\_\_\_\_, i.e., effectively try \_\_\_\_\_\_\_ instead of one specific one. This can help us spot tricky corner cases that might escape our testing! *Fuzzing* is the process of running software with inputs that we mutate a little bit every time we run it. This isn't strictly as powerful as concolic execution, but it benefits from being highly practical: we can fuzz software \_\_\_\_\_\_\_ for concolic execution.

# Fuzzing

## **Black-box fuzzing**

#### **Glass-box fuzzing**

DART, SAGE, AFL, LibFuzzer OSS-Fuzz and its trophies

10/12

*Black-box fuzzing* mutates input data with no knowledge of the target program's internals. This can find bugs, but it's not as powerful as *glass-box* fuzzing.

In glass-box fuzzing, we allow the compiler to instrument programs, much like

. This information is then used to decide whether we've explored the control-flow graph of a particular function, etc., "enough". It can also be used to decide whether two fuzz-created crashes are !

There are lots of tools for doing glass-box fuzzing. We'll have a demo (time permitting) of AFL: American Fuzzy Lop (named for a species of fuzzy rabbit).

The large-scale OSS-Fuzz project applies fuzzing to a number of critical, widely-used open-source software projects. It includes components with fun names like ClusterFuzz, and it also has quite a few trophies in large, complex projects.

# Summary

Testing

Formal methods

Hybrid approaches

# ... to *mitigate* risk

11/12

All of this work — with the exception of formal proof, which is limited in scope — serves to *mitigate* the risks of imperfect software.