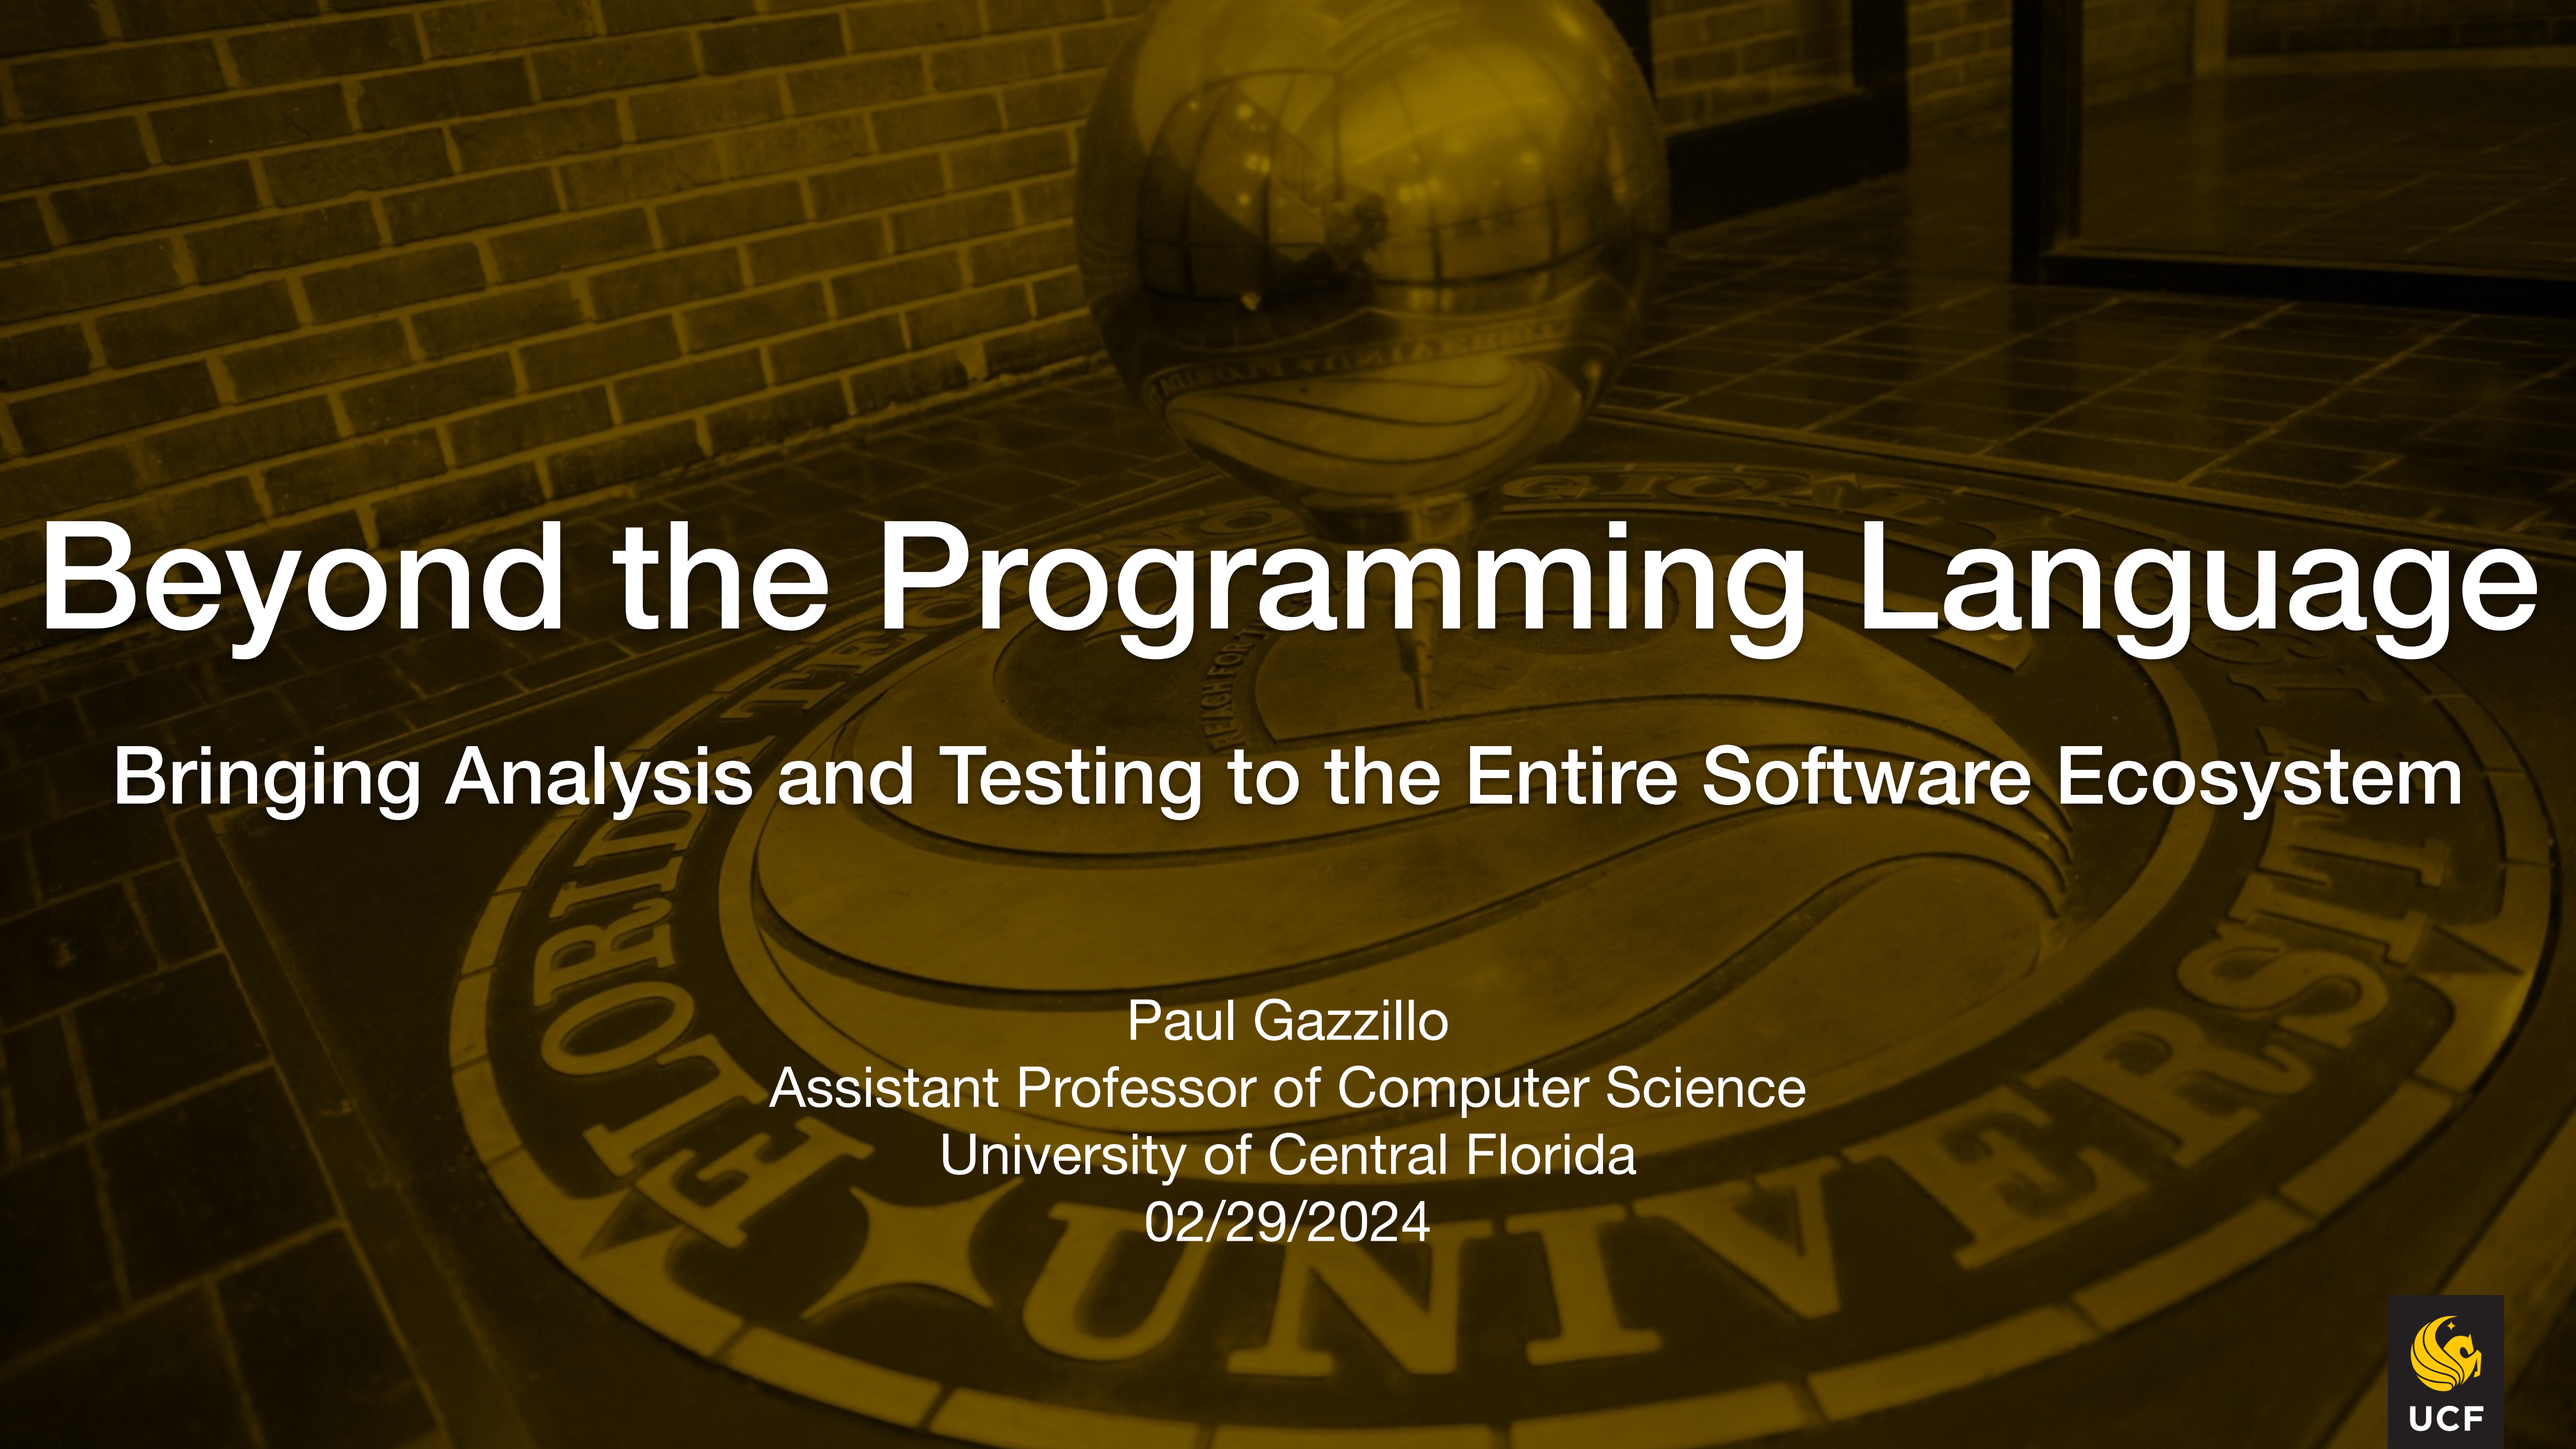

Beyond the Programming Language

Bringing Analysis and Testing to the Entire Software Ecosystem

Paul Gazzillo
Assistant Professor of Computer Science
University of Central Florida
02/29/2024

UCF

Vision

Expand the scope of software analysis beyond the programming language to the entire software ecosystem to further strengthen and secure software.

Memory Safety Dominates Exploits

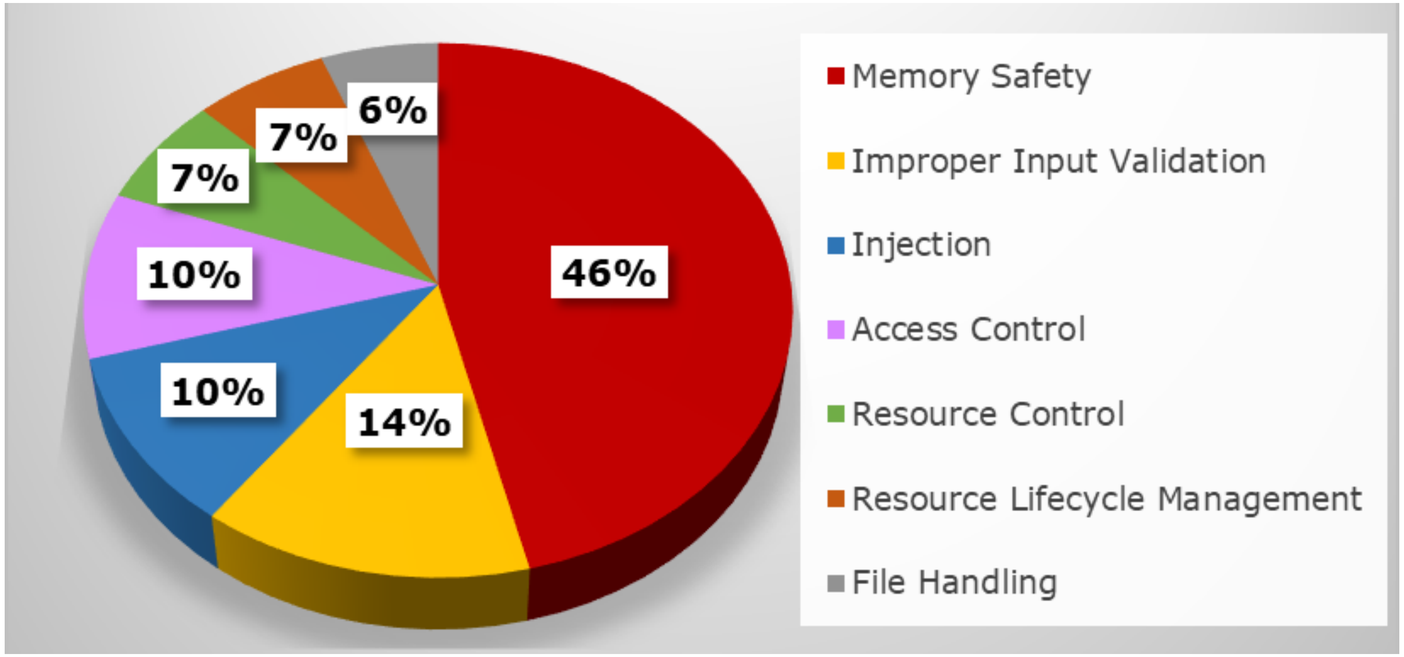

Source: [2023 CWE Top 10 KEV Weaknesses List Insights]

C/C++ Is the Origin

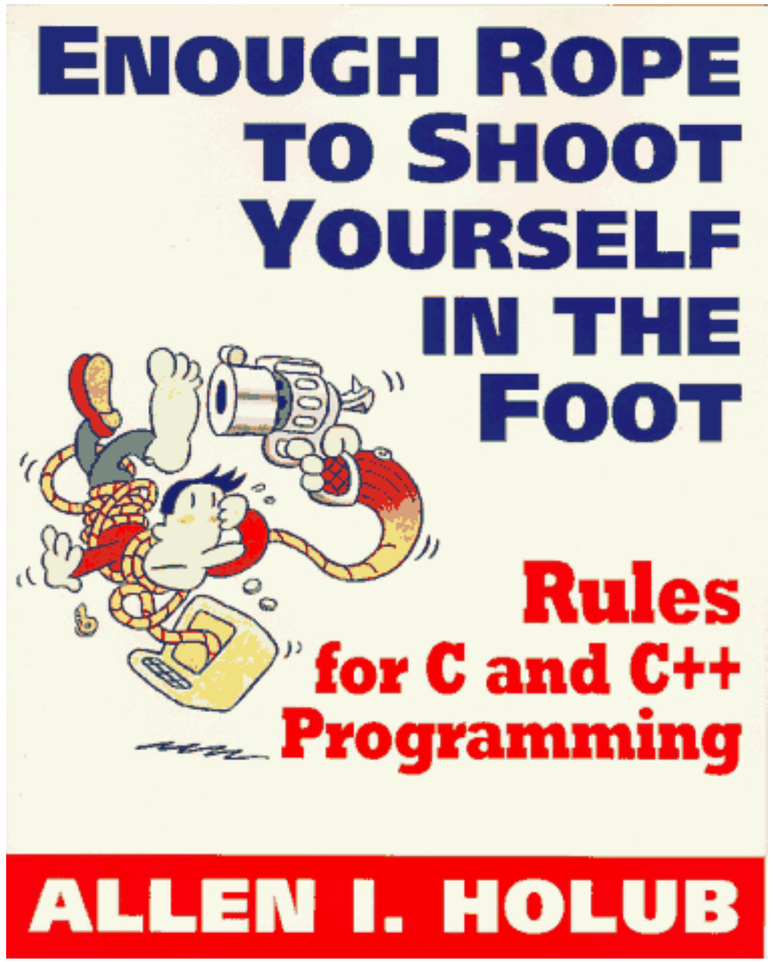

Memory Safe Programming is Solved

CYCLONE

Checked C

UCF

Just a Matter of Time

THE WHITE HOUSE

PRESS RELEASE: Future Software Should Be Memory Safe

FEBRUARY 26, 2024

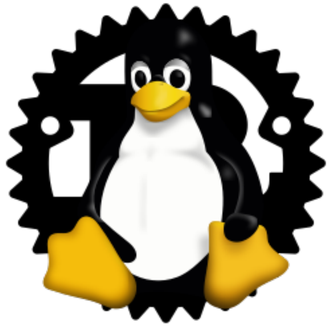

Rust for Linux

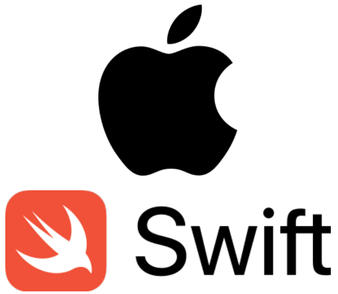

Swift for Apple

What's Left After Memory-Based Exploits?

Memory bugs happen here

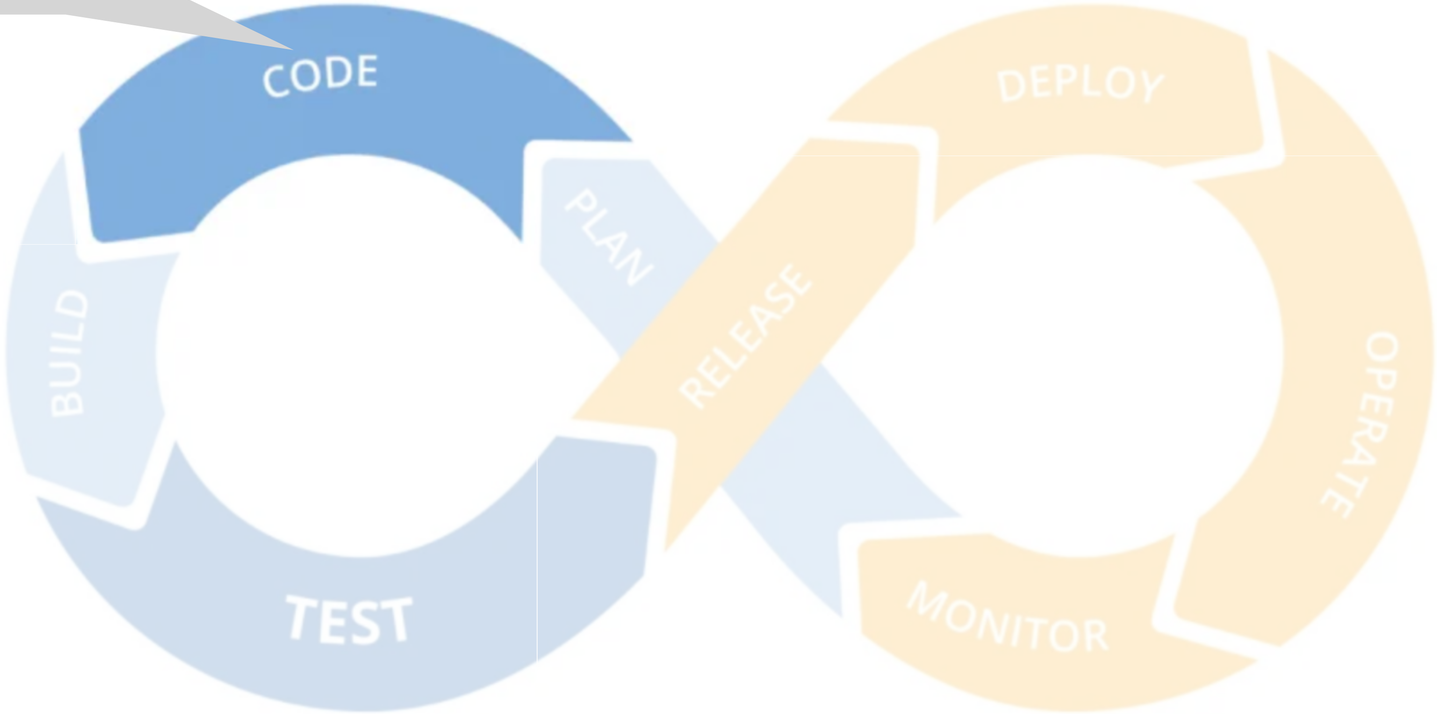

DevOps Phases

Other Phases of Development and Operations

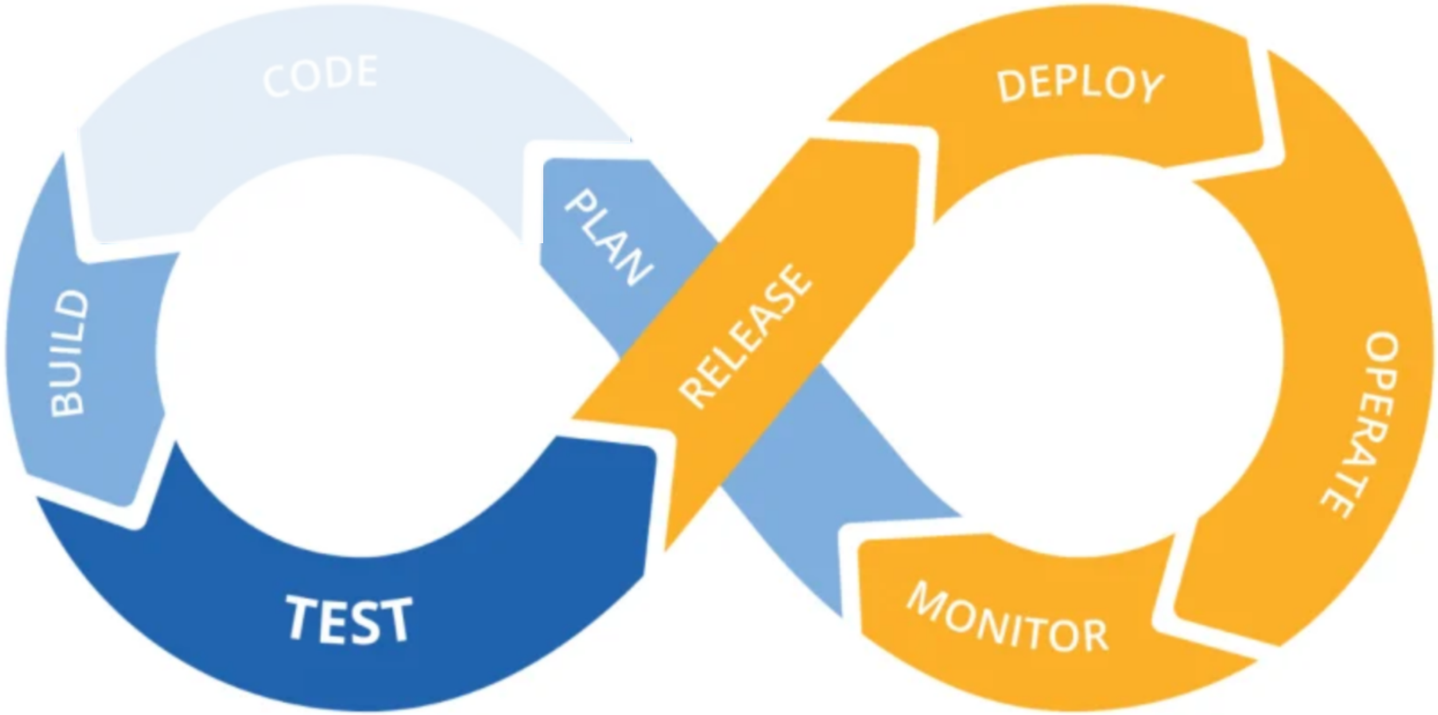

DevOps Phases

High Profile Attacks

solarwinds

- Hacked build system
- Malware in signed code
- “More than 200 victims”

- Feature, not bug
- Disable with configuration setting
- “Most critical vulnerability”

https://www.bloomberg.com/news/articles/2020-12-19/at-least-200-victims-identified-in-suspected-russian-hacking

https://www.theguardian.com/technology/2021/dec/10/software-flaw-most-critical-vulnerability-log-4-shell

UCF

TURING AWARD LECTURE

Reflections on Trusting Trust

To what extent should one trust a statement that a program is free of Trojan horses? Perhaps it is more important to trust the people who wrote the software.

KEN THOMPSON

UCF

Why Bother Breaking In?

Vision

Expand the scope of software analysis beyond the programming language to the entire software ecosystem to further strengthen and secure software.

Starting Small: Analyzing the Build System

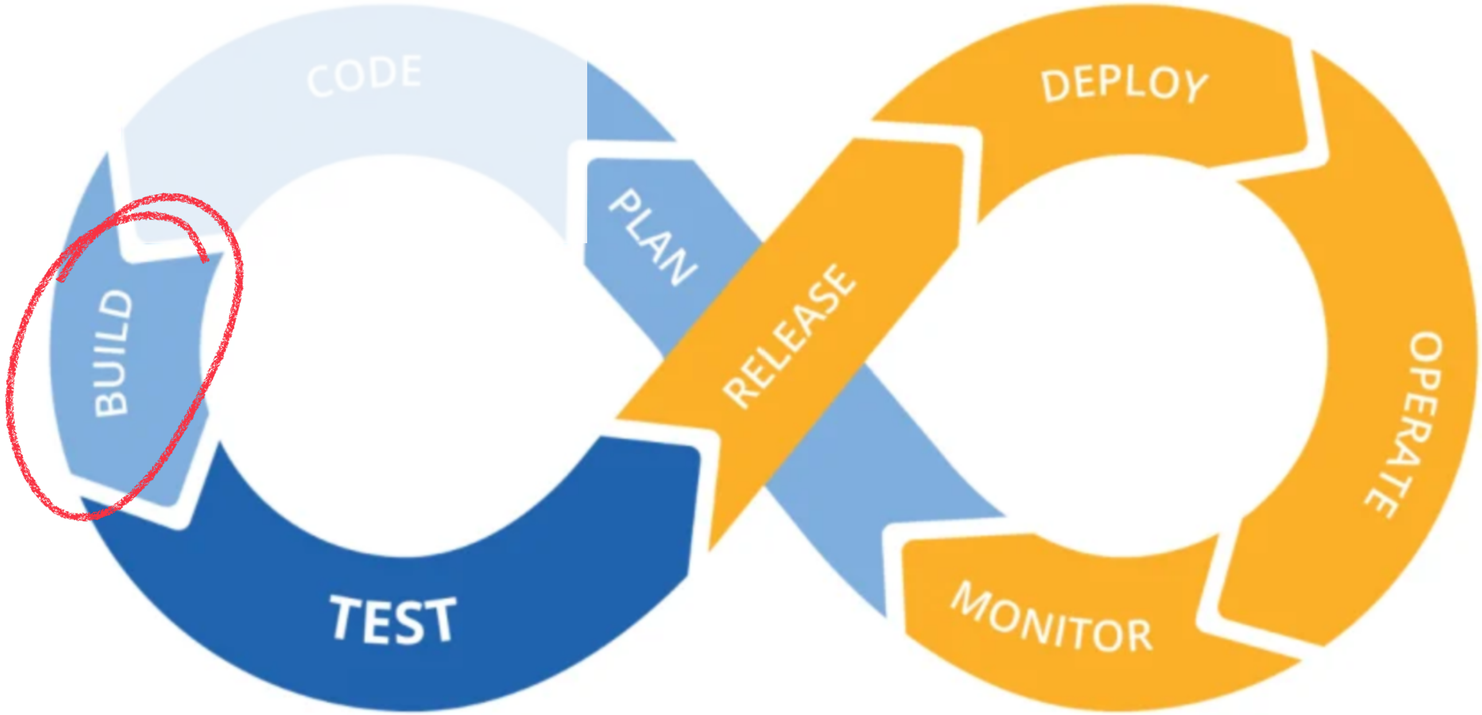

DevOps Phases

The Linux Kernel Build System

Example: Linux Kernel

Linux Kernel

70% of mobile devices

70% of IoT developers

40% of servers

UCF

Can Have Trillions of Program in One Codebase

- Allows system builders to reuse existing software

Configuration options enable/disable features

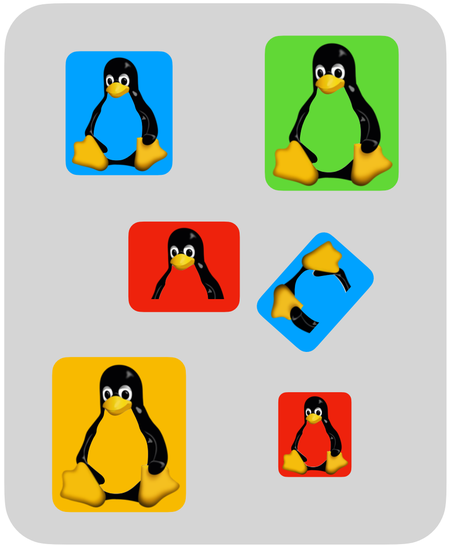

Linux build system generates many variations

Build customized software without reprogramming

Configurability Complicates Maintenance

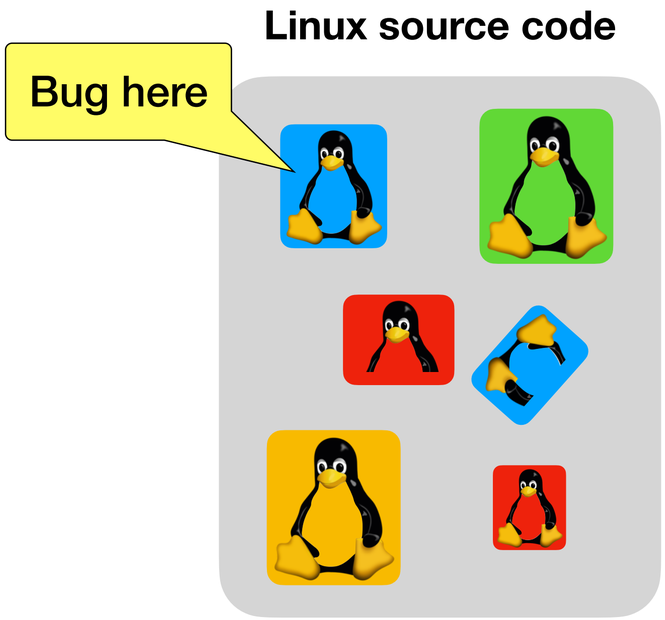

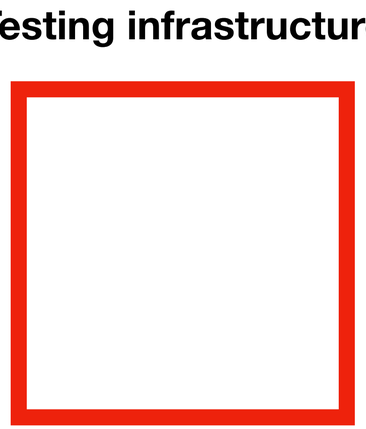

- Even if one variant program is correct, another might be broken
- Have to test all variations that might be used
- Automated testing typically works on one variant at a time

The Linux Kernel has a Very Active Codebase

Linux-next commit history

index : kernel/git/next/linux-next.git

The linux-next integration testing tree — Linux Next Group

about summary refs log tree commit diff stats

Age	Commit message (Expand)	Author	Files	Lines
8 hours	Add linux-next specific files for 20240213 HEAD next-20240213 master	Stephen Rothwell	4	-0/+9642
8 hours	fixup for "drm/amd: Stop evicting resources on APUs in suspend"	Stephen Rothwell	1	-1/+1
9 hours	Merge branch 'for-next/kspp' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	45	-185/+336
9 hours	Merge branch 'bitmap-for-next' of https://github.com/norov/linux.git	Stephen Rothwell	49	-417/+634
9 hours	Merge branch 'for-next/execve' of git://git.kernel.org/pub/scm/linux/kernel/g...	Stephen Rothwell	1	-1/+1
9 hours	Merge branch 'rust-next' of https://github.com/Rust-for-Linux/linux.git	Stephen Rothwell	14	-49/+177
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/mic/linu...	Stephen Rothwell	16	-91/+1183
9 hours	Merge branch 'slab/for-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	7	-126/+115
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/kris...	Stephen Rothwell	3	-20/+42
9 hours	Merge branch 'zstd-next' of https://github.com/terrelln/linux.git	Stephen Rothwell	58	-2594/+4789
9 hours	Merge branch 'mhi-next' of git://git.kernel.org/pub/scm/linux/kernel/git/mani...	Stephen Rothwell	9	-88/+423
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/srin...	Stephen Rothwell	6	-69/+83
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/srin...	Stephen Rothwell	1	-4/+4
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/joel/fsi...	Stephen Rothwell	2	-2/+18
9 hours	Merge branch 'for-next/seccomp' of git://git.kernel.org/pub/scm/linux/kernel/...	Stephen Rothwell	3	-14/+73
9 hours	Merge branch 'ntb-next' of https://github.com/jonmason/ntb.git	Stephen Rothwell	2	-2/+2
9 hours	Merge branch 'libnvdimm-for-next' of git://git.kernel.org/pub/scm/linux/kerne...	Stephen Rothwell	3	-3/+3
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/live...	Stephen Rothwell	0	-0/+0
9 hours	Merge branch 'kunit' of git://git.kernel.org/pub/scm/linux/kernel/git/shuah/l...	Stephen Rothwell	2	-3/+4
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/shuah/li...	Stephen Rothwell	32	-121/+340
9 hours	Merge branch 'pwm/for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/...	Stephen Rothwell	12	-384/+373
9 hours	Merge branch 'renesas-pinctrl' of git://git.kernel.org/pub/scm/linux/kernel/g...	Stephen Rothwell	4	-53/+276
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/linu...	Stephen Rothwell	24	-94/+130
9 hours	Merge branch 'gpio/for-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	46	-626/+2510
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/remo...	Stephen Rothwell	12	-193/+271
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/mkp/...	Stephen Rothwell	24	-883/+727
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/jejb...	Stephen Rothwell	43	-471/+1385
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/tj/c...	Stephen Rothwell	1	-8/+12
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/vkoul/dm...	Stephen Rothwell	15	-242/+615
10 hours	Merge branch 'counter-next' of git://git.kernel.org/pub/scm/linux/kernel/git/...	Stephen Rothwell	2	-2/+1
10 hours	Merge branch 'staging-next' of git://git.kernel.org/pub/scm/linux/kernel/git/...	Stephen Rothwell	38	-4359/+168
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/krzk...	Stephen Rothwell	3	-2/+61
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/vkoul/so...	Stephen Rothwell	2	-4/+2
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/phy/linu...	Stephen Rothwell	40	-1257/+3441
10 hours	Merge branch 'togreg' of git://git.kernel.org/pub/scm/linux/kernel/git/jic23/...	Stephen Rothwell	67	-455/+2724
10 hours	Merge branch 'icc-next' of git://git.kernel.org/pub/scm/linux/kernel/git/djak...	Stephen Rothwell	12	-1116/+1599
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/fpga...	Stephen Rothwell	4	-54/+104
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/coresigh...	Stephen Rothwell	30	-803/+1377
10 hours	Merge branch 'habanalabs-next' of git://git.kernel.org/pub/scm/linux/kernel/g...	Stephen Rothwell	13	-396/+899
10 hours	Merge branch 'char-misc-next' of git://git.kernel.org/pub/scm/linux/kernel/gi...	Stephen Rothwell	6	-10/+35
10 hours	Merge branch 'tty-next' of git://git.kernel.org/pub/scm/linux/kernel/git/greg...	Stephen Rothwell	67	-1615/+2168
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/westeri/...	Stephen Rothwell	12	-50/+278
10 hours	Merge branch 'usb-next' of git://git.kernel.org/pub/scm/linux/kernel/git/greg...	Stephen Rothwell	77	-1038/+5655
10 hours	Merge branch 'driver-core-next' of git://git.kernel.org/pub/scm/linux/kernel/...	Stephen Rothwell	7	-22/+37
10 hours	Merge branch 'for-leds-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	22	-181/+613
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/sre/...	Stephen Rothwell	1	-1/+1
10 hours	Merge branch 'for-firmware-next' of git://git.kernel.org/pub/scm/linux/kernel...	Stephen Rothwell	1	-1/+1
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/pdx8...	Stephen Rothwell	19	-244/+669
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/tj/w...	Stephen Rothwell	9	-302/+1284
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/denn...	Stephen Rothwell	0	-0/+0
10 hours	Merge branch 'next' of https://github.com/kvm-x86/linux.git	Stephen Rothwell	85	-745/+1631
10 hours	Merge branch 'riscv_kvm_next' of https://github.com/kvm-riscv/linux.git	Stephen Rothwell	2	-11/+15
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/kvmarm/k...	Stephen Rothwell	28	-84/+247
10 hours	scsi: core: Make scsi_bus_type const	Ricardo B. Marliere	2	-2/+2
10 hours	Merge branch kvm-arm64/misc into kvmarm/next	Oliver Upton	1	-1/+1
10 hours	Merge branch 'rcu/next' of git://git.kernel.org/pub/scm/linux/kernel/git/paul...	Stephen Rothwell	32	-564/+841
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/trac...	Stephen Rothwell	0	-0/+0
10 hours	Merge branch 'edac-for-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	27	-295/+3949
10 hours	Merge branch 'timers/drivers/next' of git://git.linaro.org/people/daniel.lezc...	Stephen Rothwell	5	-7/+22
10 hours	Merge branch 'master' of git://git.kernel.org/pub/scm/linux/kernel/git/tip/ti...	Stephen Rothwell	204	-1114/+4829
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/broo...	Stephen Rothwell	65	-586/+1037
10 hours	KVM: selftests: Print timer ctl register in ISTATUS assertion	Oliver Upton	1	-1/+1
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/krzk...	Stephen Rothwell	4	-9/+9
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/robh...	Stephen Rothwell	21	-652/+661
10 hours	scsi: core: Really include kunit tests with SCSI_LIB_KUNIT_TEST	Lukas Bulwahn	1	-1/+1
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/pcmoore/...	Stephen Rothwell	2	-4/+2
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/joro/iom...	Stephen Rothwell	9	-483/+457
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/jarkko/l...	Stephen Rothwell	1	-3/+3
10 hours	Merge branch 'next' of git://github.com/cschaufler/smack-next	Stephen Rothwell	2	-41/+86

~30k mailing list messages per month

~6k commits per month, 100s per day

e.g., ~13k commits between v5.12 and v5.13

All These Code Changes Need Testing

Most active 5.12 bug reporters

kernel test robot	184	16.1%
Syzbot	111	9.7%
Abaci Robot	107	9.4%
Dan Carpenter	44	3.9%
Hulk Robot	41	3.6%
Stephen Rothwell	28	2.5%
Randy Dunlap	19	1.7%
Kent Overstreet	12	1.1%
Guenter Roeck	11	1.0%
TOTE Robot	11	1.0%
Colin Ian King	9	0.8%
Andrii Nakryiko	8	0.7%
Juan Vazquez	7	0.6%
Arnd Bergmann	6	0.5%

https://lwn.net/Articles/853039/

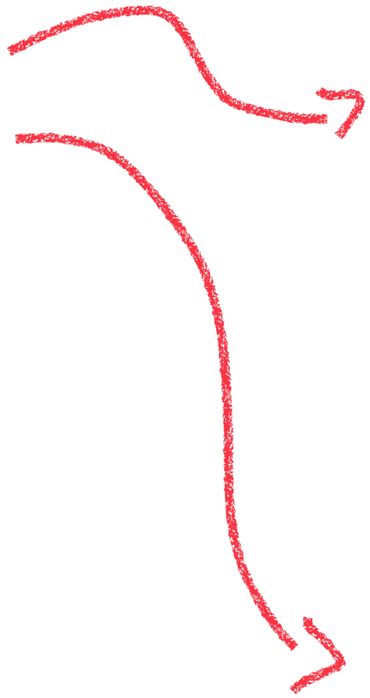

Intel 0-day kernel test robot

- Suite of static and dynamic testing tools
  - compile, boot, performance, etc.
- continuously runs on new commits in linux-next

Google syzbot

- syzkaller system call fuzz tester
- continuously tests the kernel
- runs on linux-next, other versions

The Build System Causes Blindspots in Testing

Code Hidden by the Build System

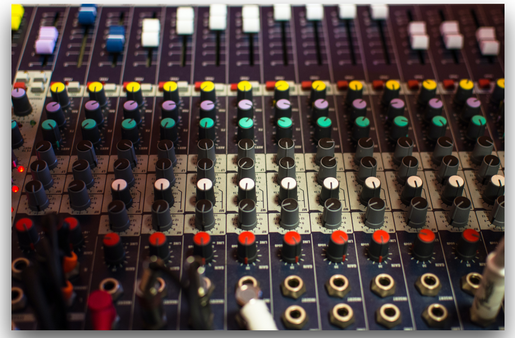

Configuration options
determine what's compiled

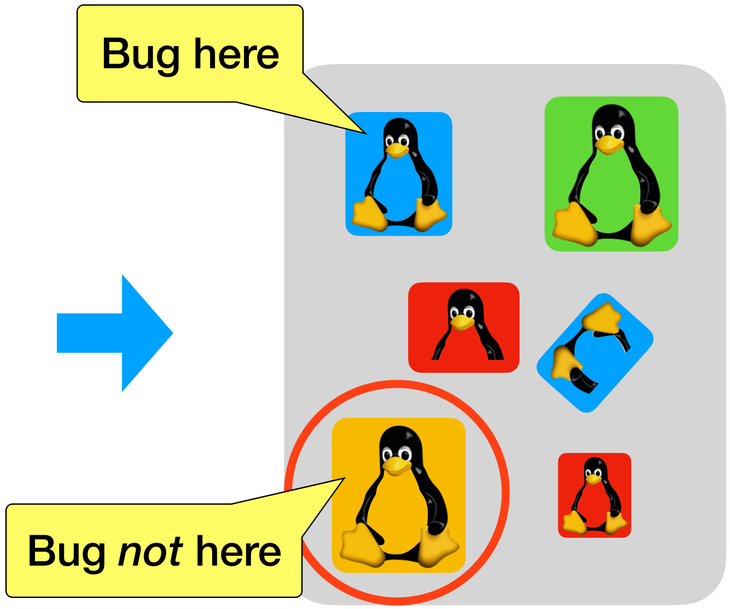

A new bug may only appear
in some configurations

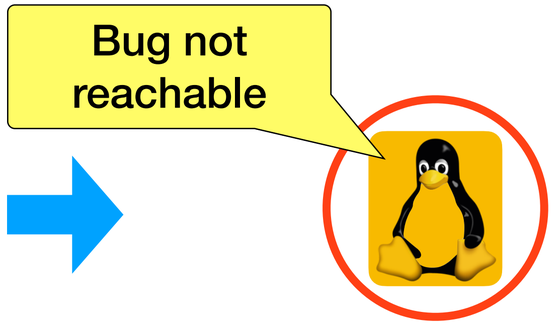

Configuration-dependent
bugs not always reachable

Test Robots Miss Most Code Changes

Configuration	Patch Coverage
Default	22%
Random	30%
Maximal	89%

10x longer build time

No runtime testing

No variation

Test Robots Miss Most Code Changes

Configuration	Patch Coverage
Default	22%
Random	30%
Maximal	89%

Test robots mostly use these configurations

UCF

Maximal Testing is Limited

Most active 5.12 bug reporters

kernel test robot	184	16.1%
Syzbot	111	9.7%
Abaci Robot	107	9.4%
Dan Carpenter	44	3.9%
Hulk Robot	41	3.6%
Stephen Rothwell	28	2.5%
Randy Dunlap	19	1.7%
Kent Overstreet	12	1.1%
Guenter Roeck	11	1.0%
TOTE Robot	11	1.0%
Colin Ian King	9	0.8%
Andrii Nakryiko	8	0.7%
Juan Vazquez	7	0.6%
Arnd Bergmann	6	0.5%

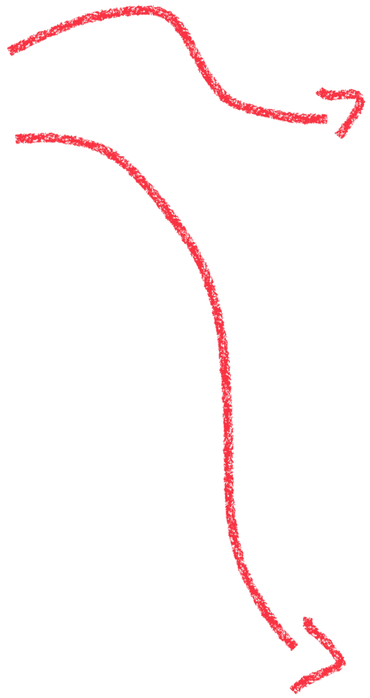

Intel 0-day kernel test robot
- Maximal only for build test

Google syzbot
- Based on default configuration

https://lwn.net/Articles/853039/

Key Problem

Lots of ways to select configuration files for testing
but no guarantee that committed changes get built

We Balance Coverage and Build Size

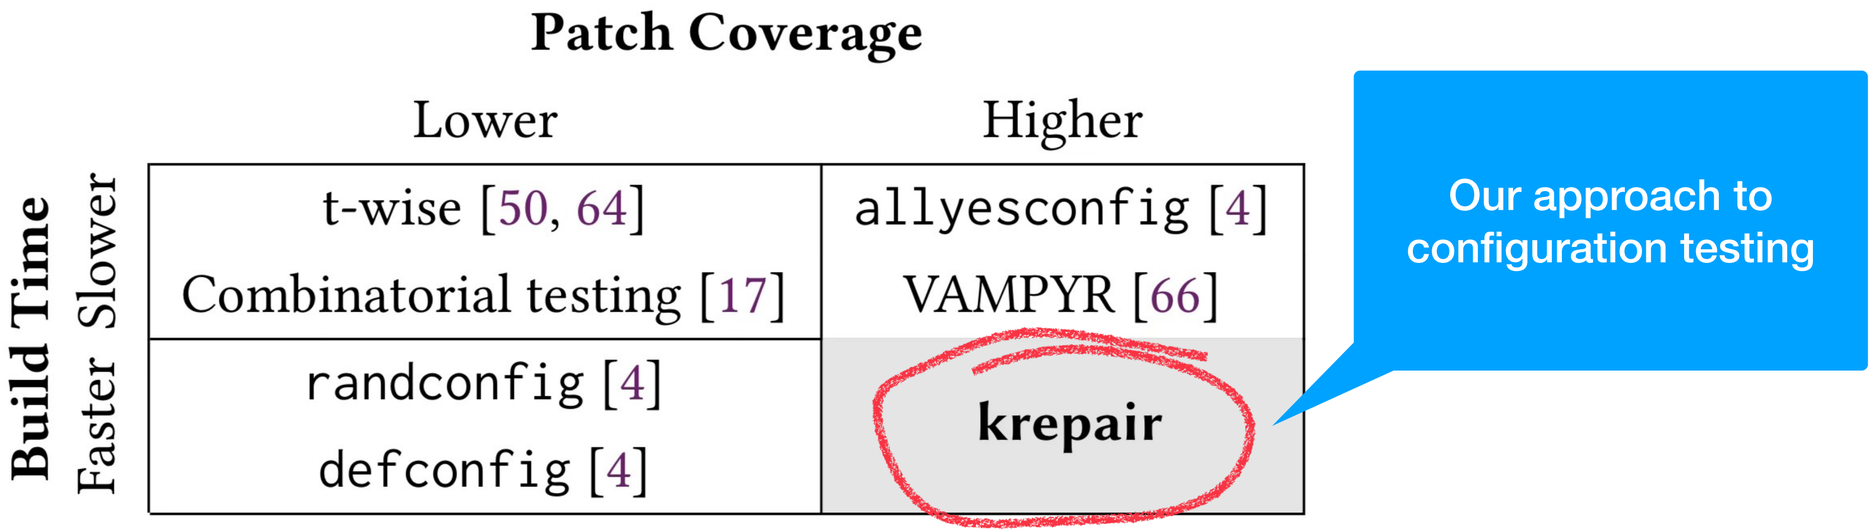

Build System Turns Configurations into Binaries

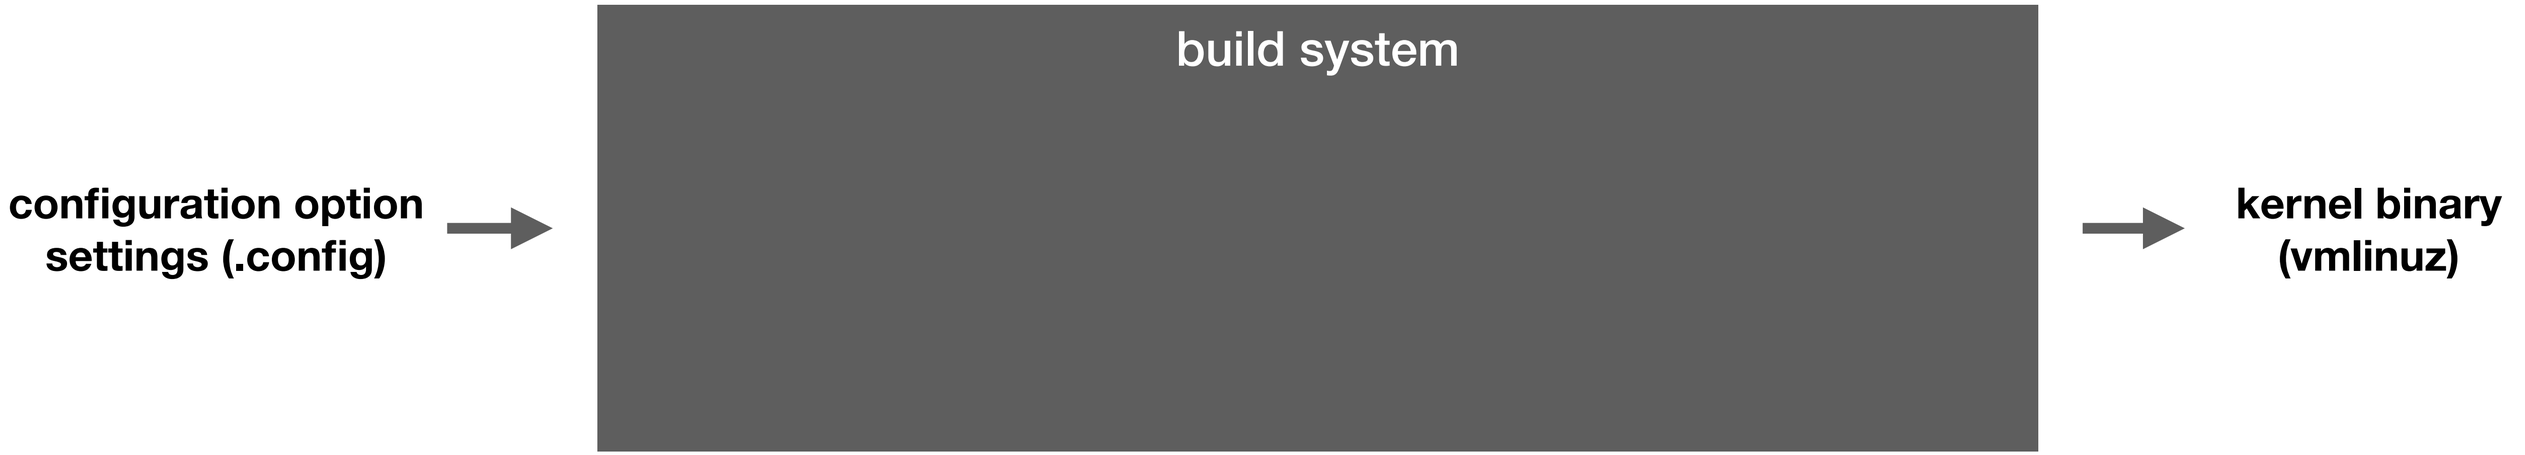

Build System Comprises Several Tools

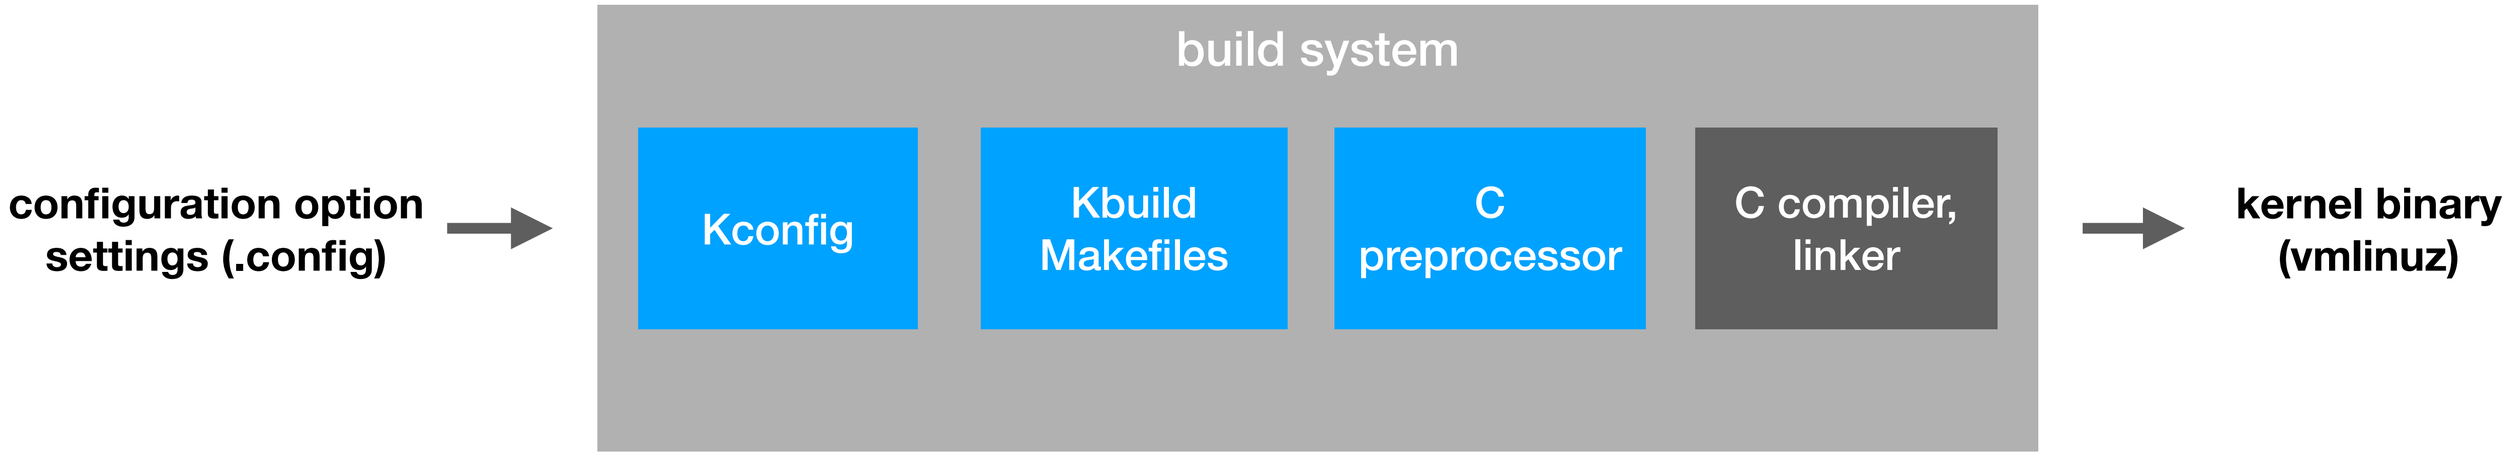

Build System Comprises Several Tools

Our Approach: krepair

Use program analysis on build system to guide configuraiton testing

Our Approach: krepair

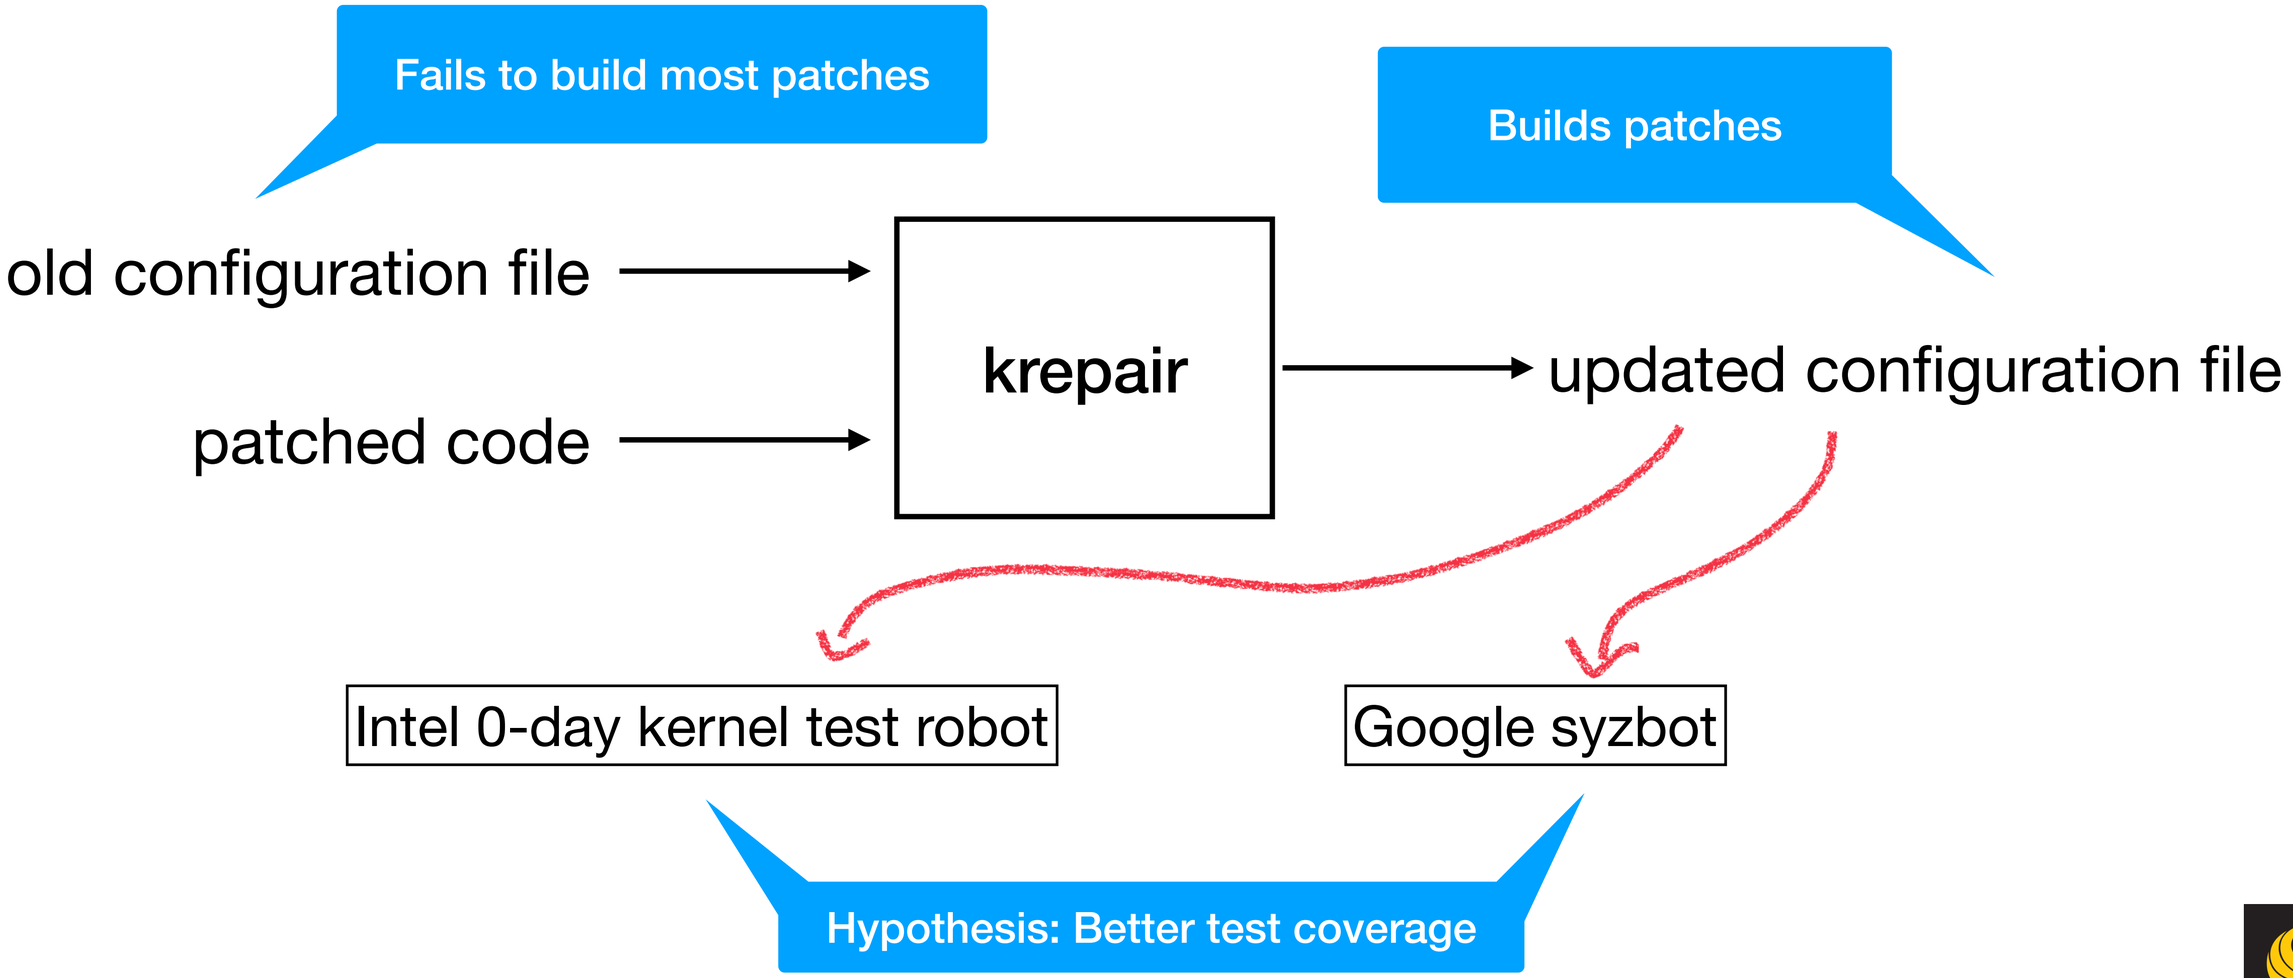

Use Program Analysis on Build Tools

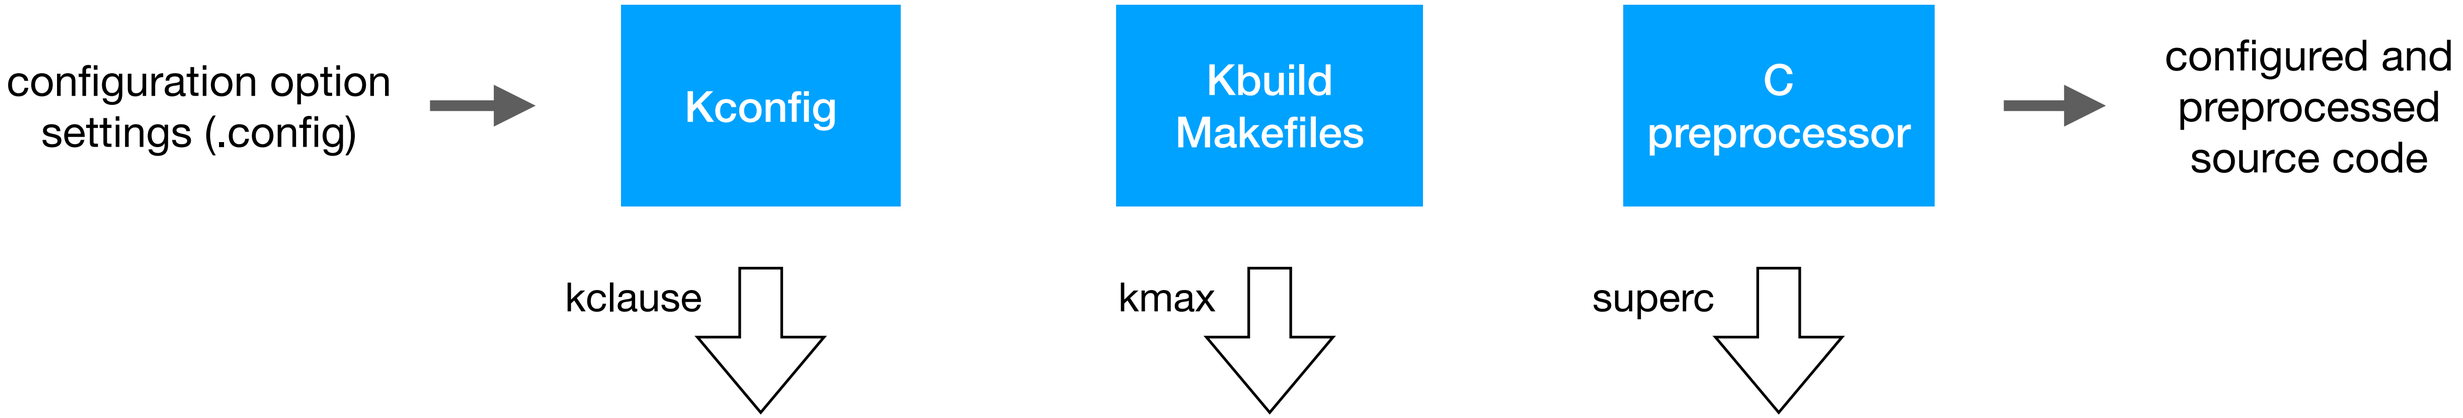

logical constraints on .config files for each file:line

Formally Model Build System Behavior

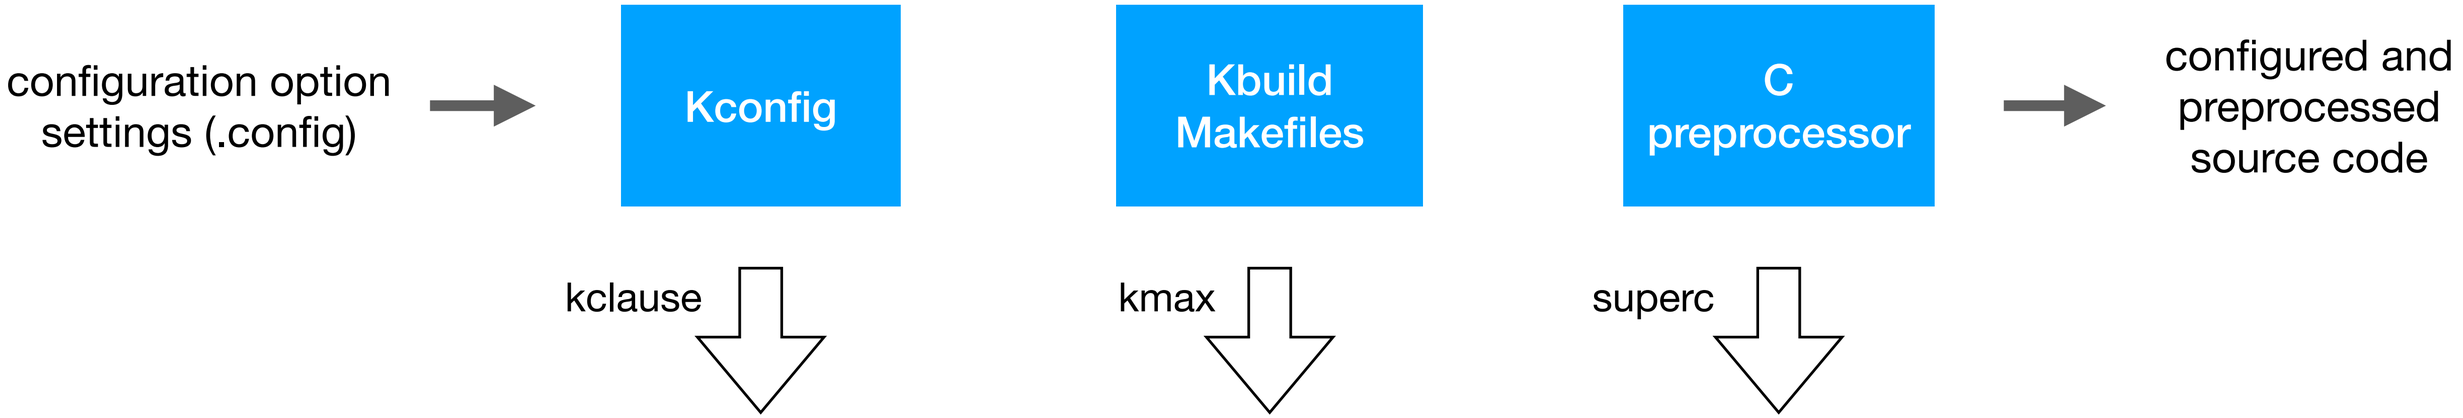

logical constraints on .config files for each file:line

.config file constraints

source code locations

krepair's Algorithm

1. **Analyze**: find patch covering constraints
2. **Reduce**: remove options preventing patch coverage
3. **Repair**: re-add only settings that satisfy patch coverage constraints

1) Figure Out Configuration Constraints for the Patch

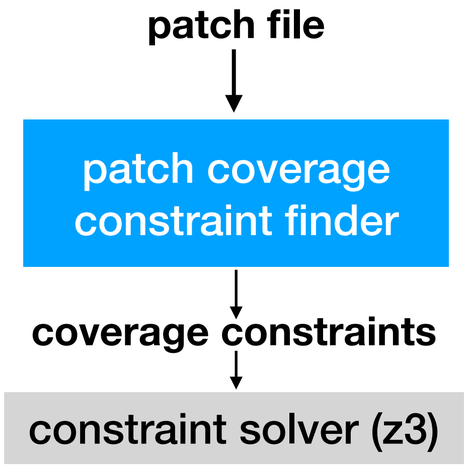

2) Remove Options Preventing Patch from Building

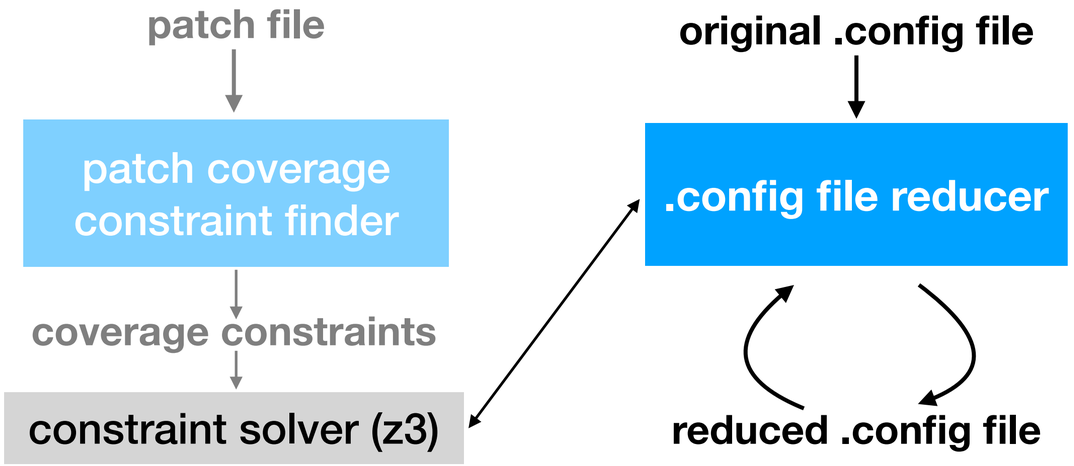

3) Add Back Settings that Satisfy Constraints

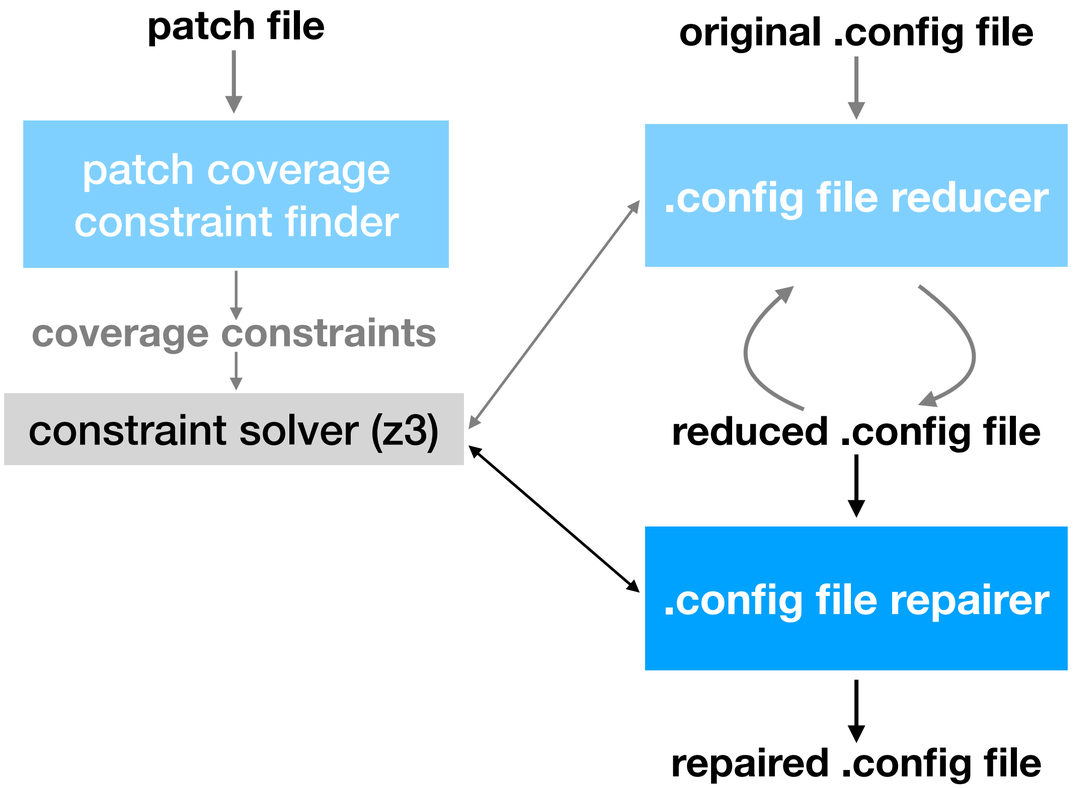

Evaluation

Can we cover patches without exploding build times?

Comparing to Default and Maximal

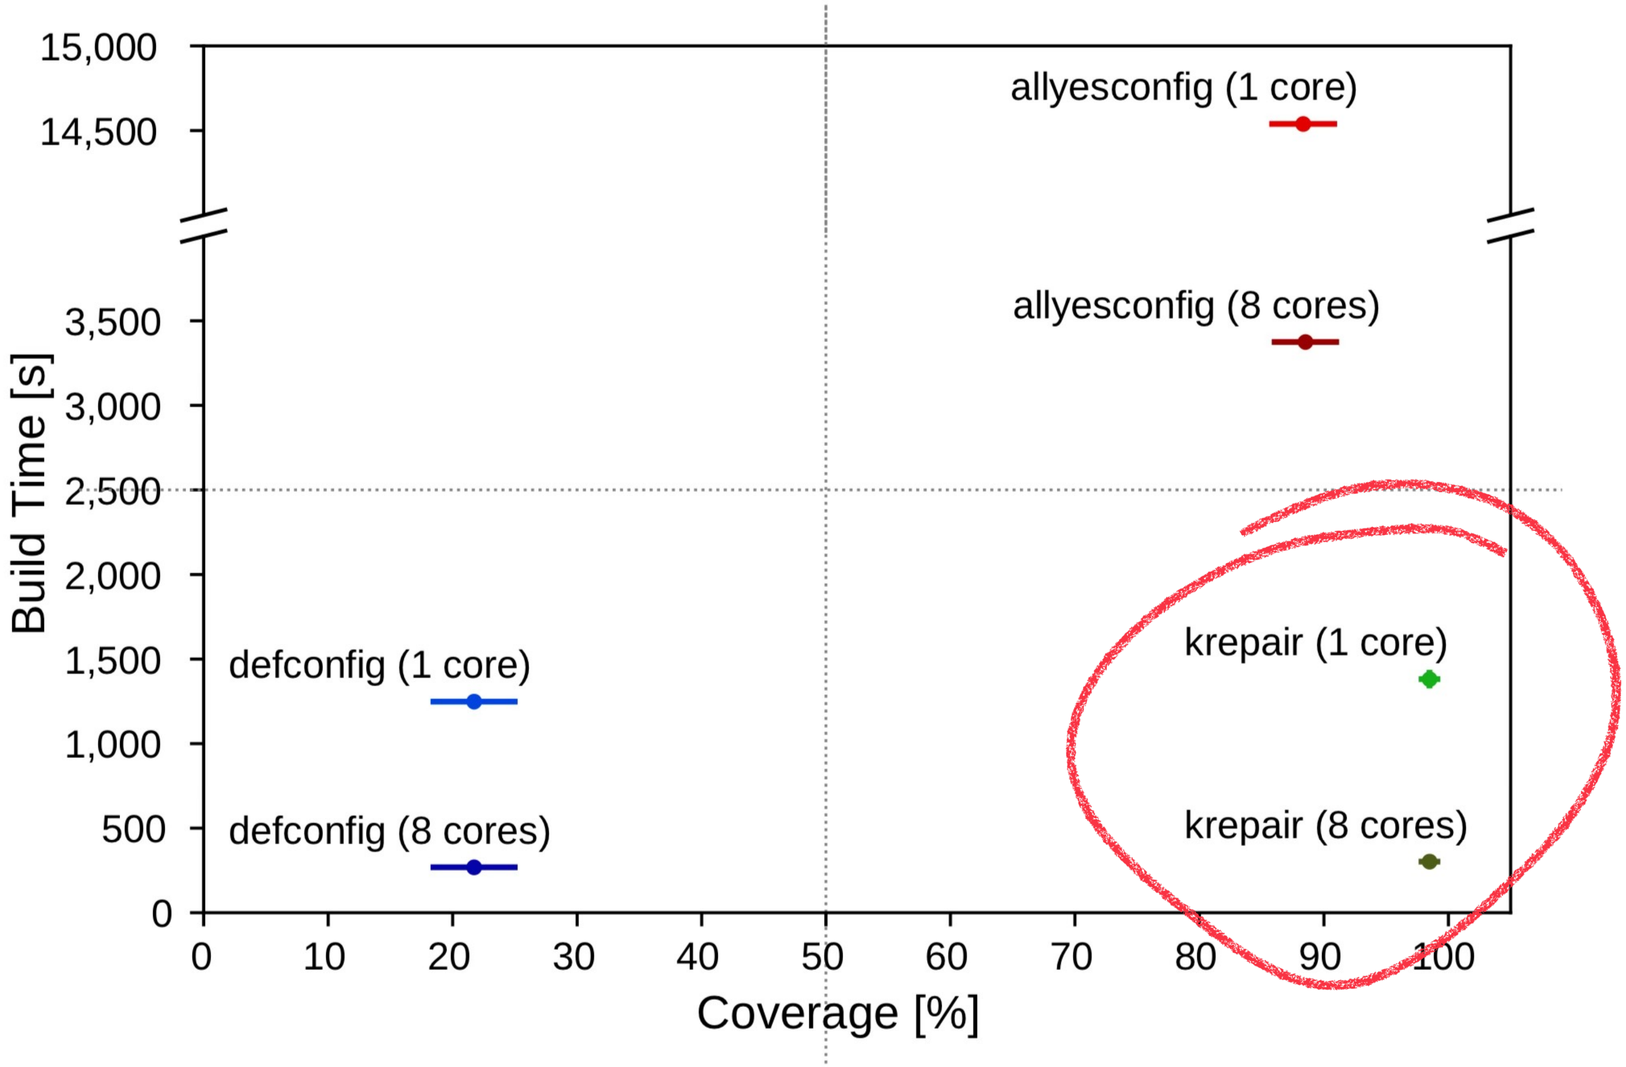

Check Out the Paper for More

Maximizing Patch Coverage for Testing of Highly-Configurable Software without Exploding Build Times

NECIP FAZIL YILDIRAN, University of Central Florida, USA
JEHO OH, University of Texas, USA
JULIA LAWALL, Inria, France
PAUL GAZZILLO, University of Central Florida, USA

https://paulgazzillo.com/papers/fse24.pdf

UCF

Next Steps: Improving Fuzzer Coverage

CVE-2023-3161

Public on January 24, 2023
Last Modified: February 14, 2024 at 9:50:15 PM UTC

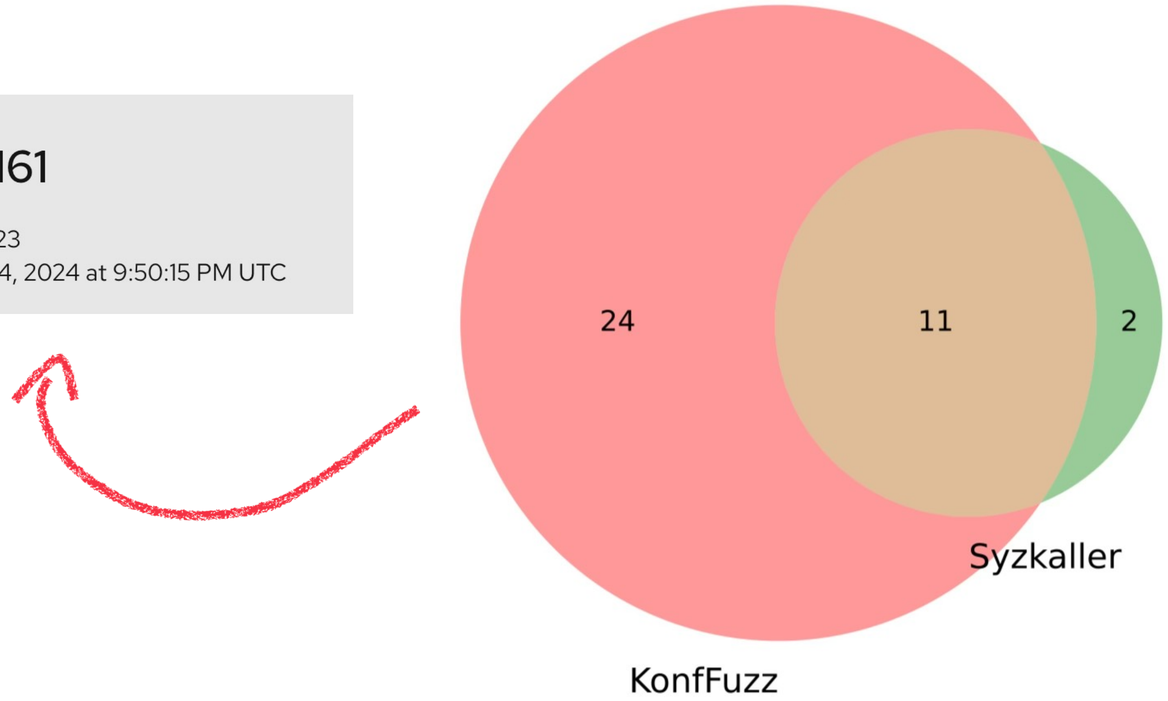

Figure 4: 12-hour *new* bugs found by KONFFUZZ and Syzkaller.

https://github.com/paulgazz/kmax

Bugs in the Build System Itself

Bugs here too

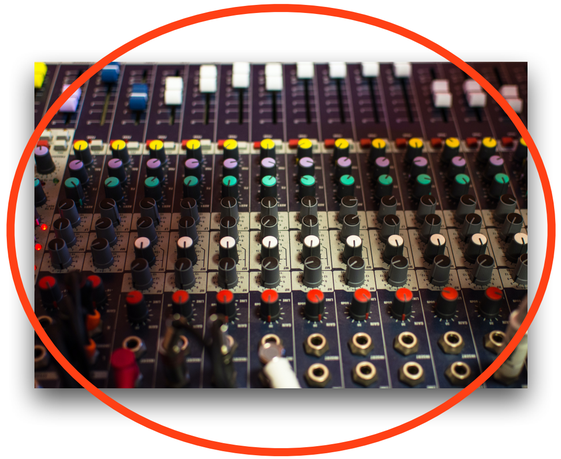

Configuration options determine what's compiled

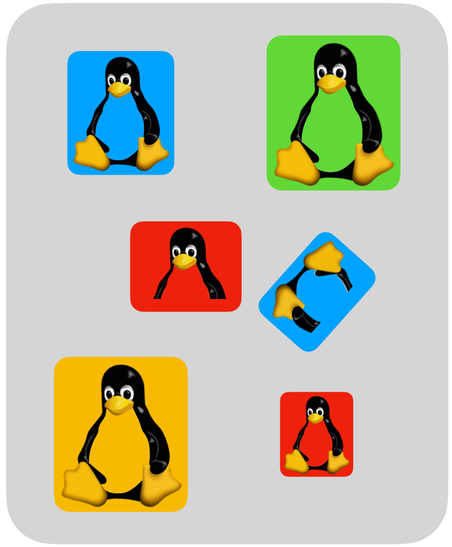

A new bug may only appear in some configurations

Configuration-dependent bugs not always reachable

our goal: automatically analyze the build system

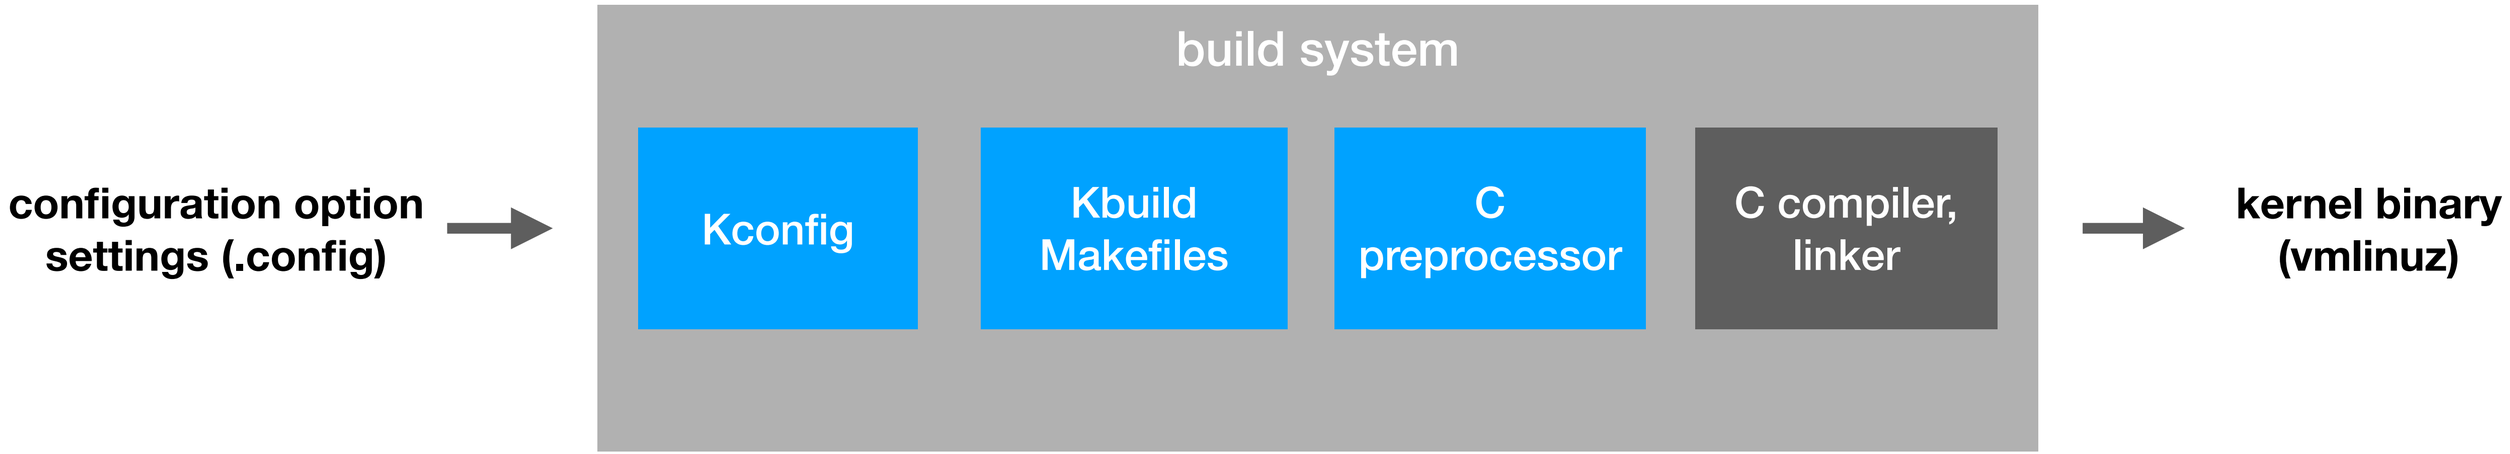

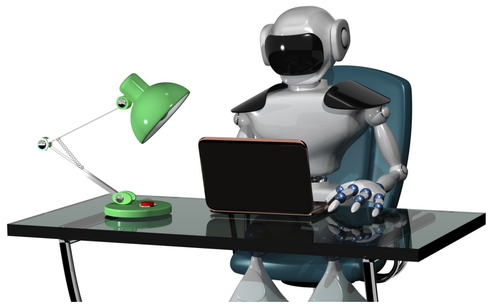

today's focus: Kconfig's unmet dependency bugs

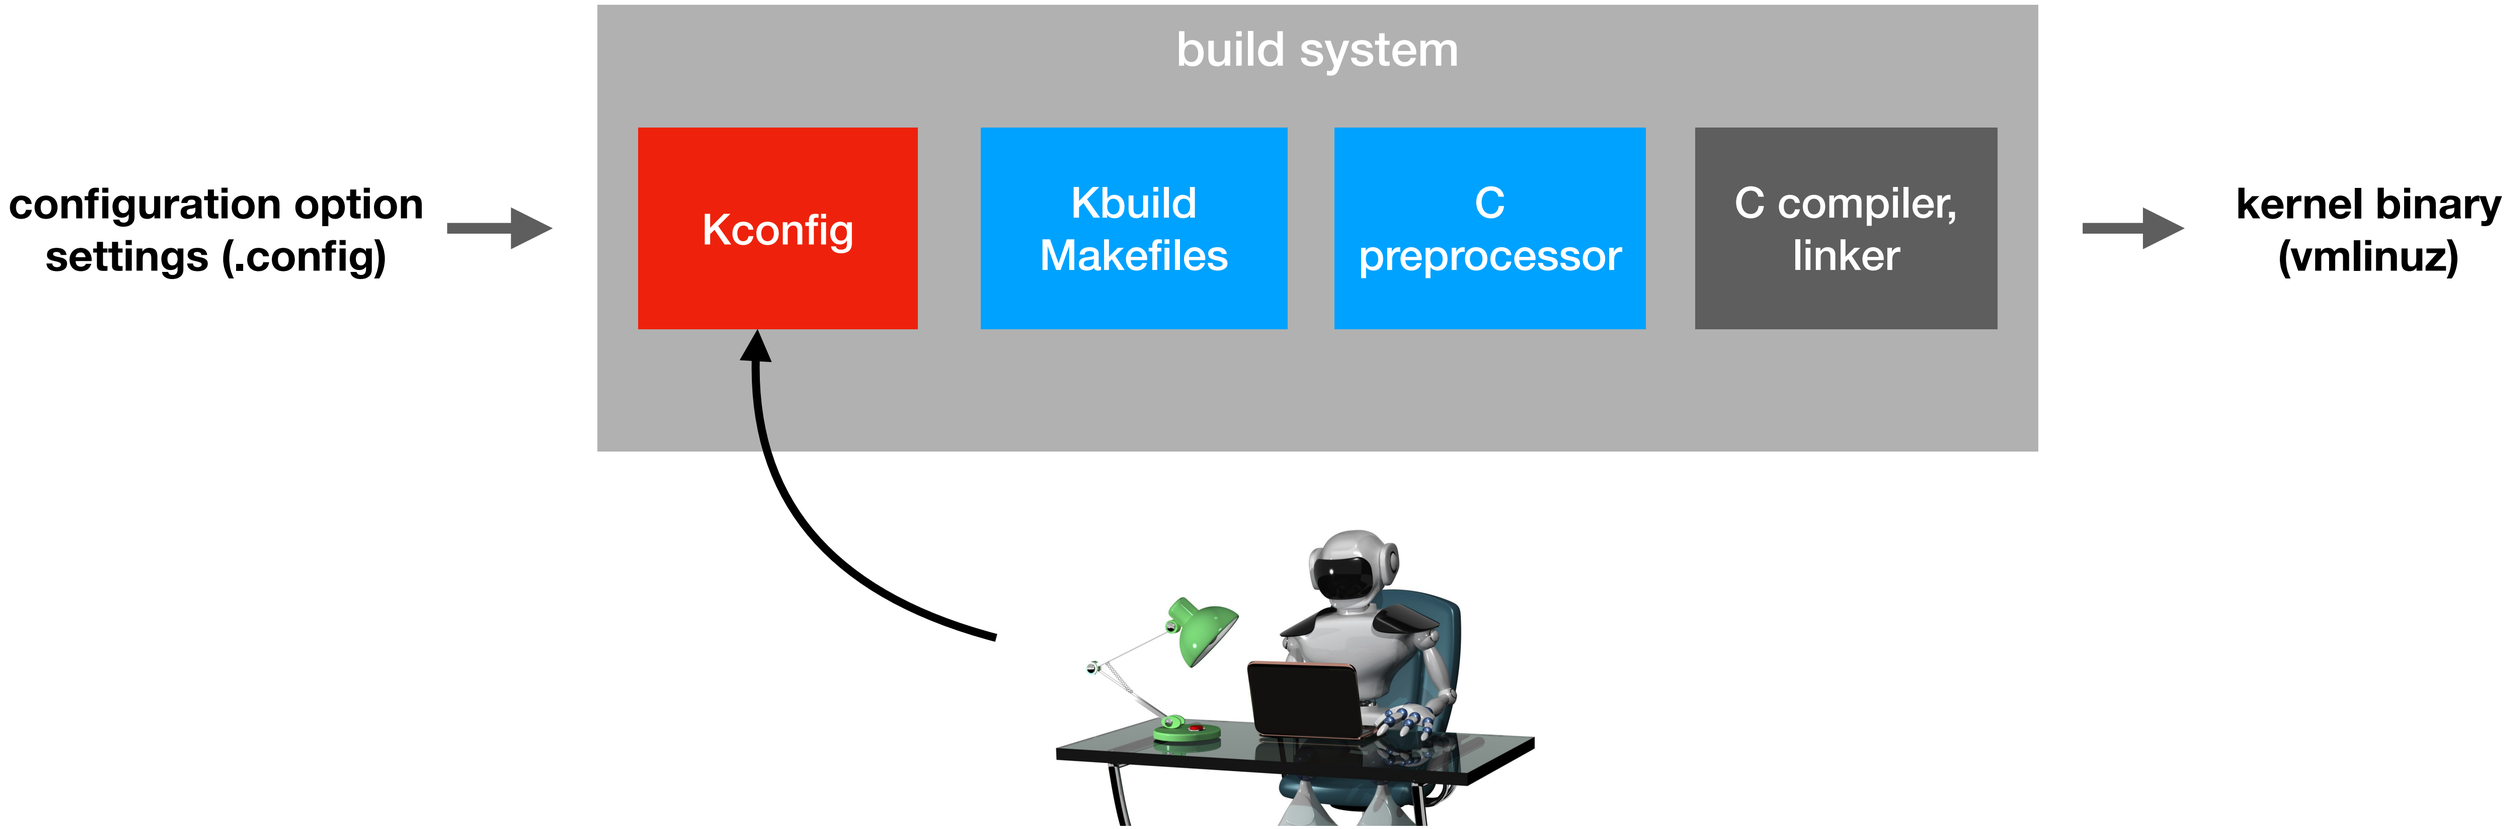

kconfig language example

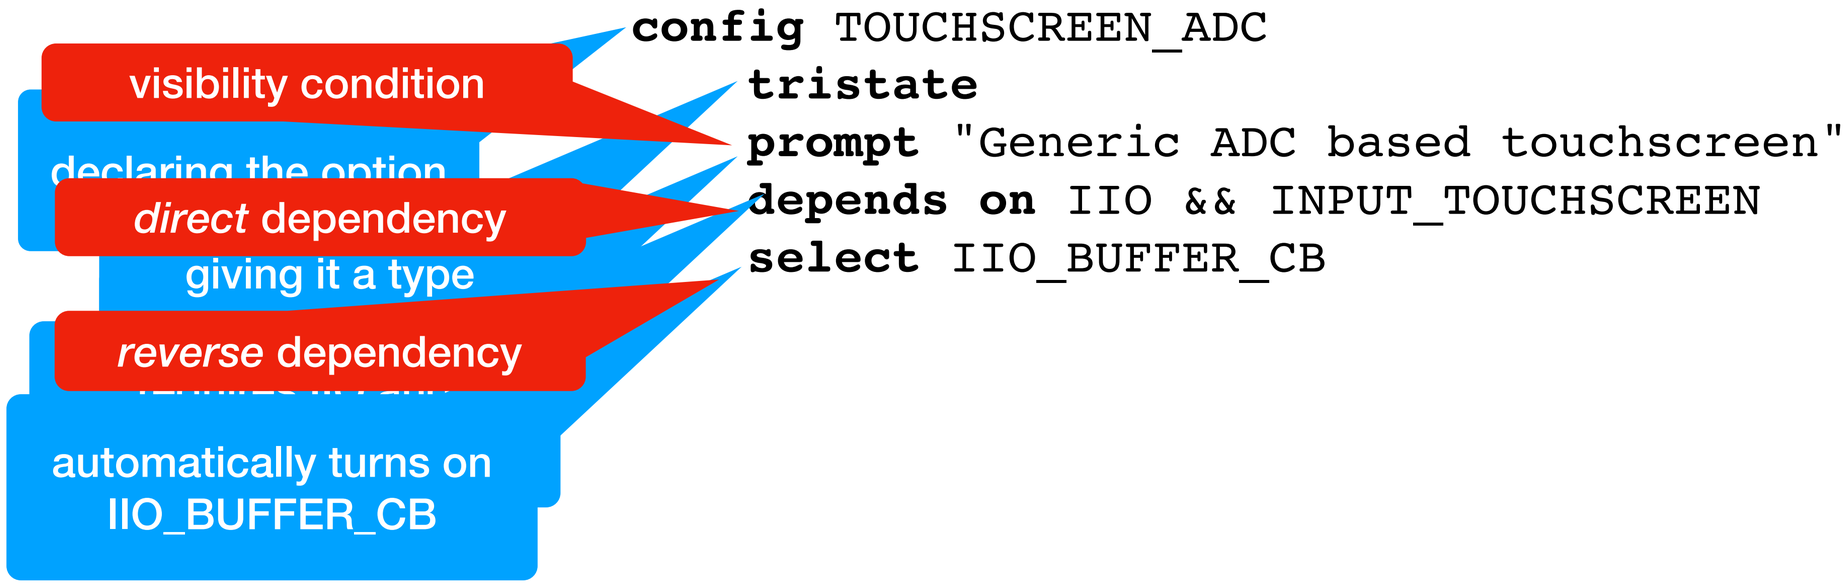

the unmet dependency bug

> “**select should be used with care. select will force a symbol to a value without visiting the dependencies.** By abusing select you are able to select a symbol FOO even if FOO depends on BAR that is not set. In general use select only for non-visible symbols (no prompts anywhere) and for symbols with no dependencies. That will limit the usefulness but on the other hand avoid the illegal configurations all over”

https://www.kernel.org/doc/html/latest/kbuild/kconfig-language.html

an unmet dependency bug in the wild

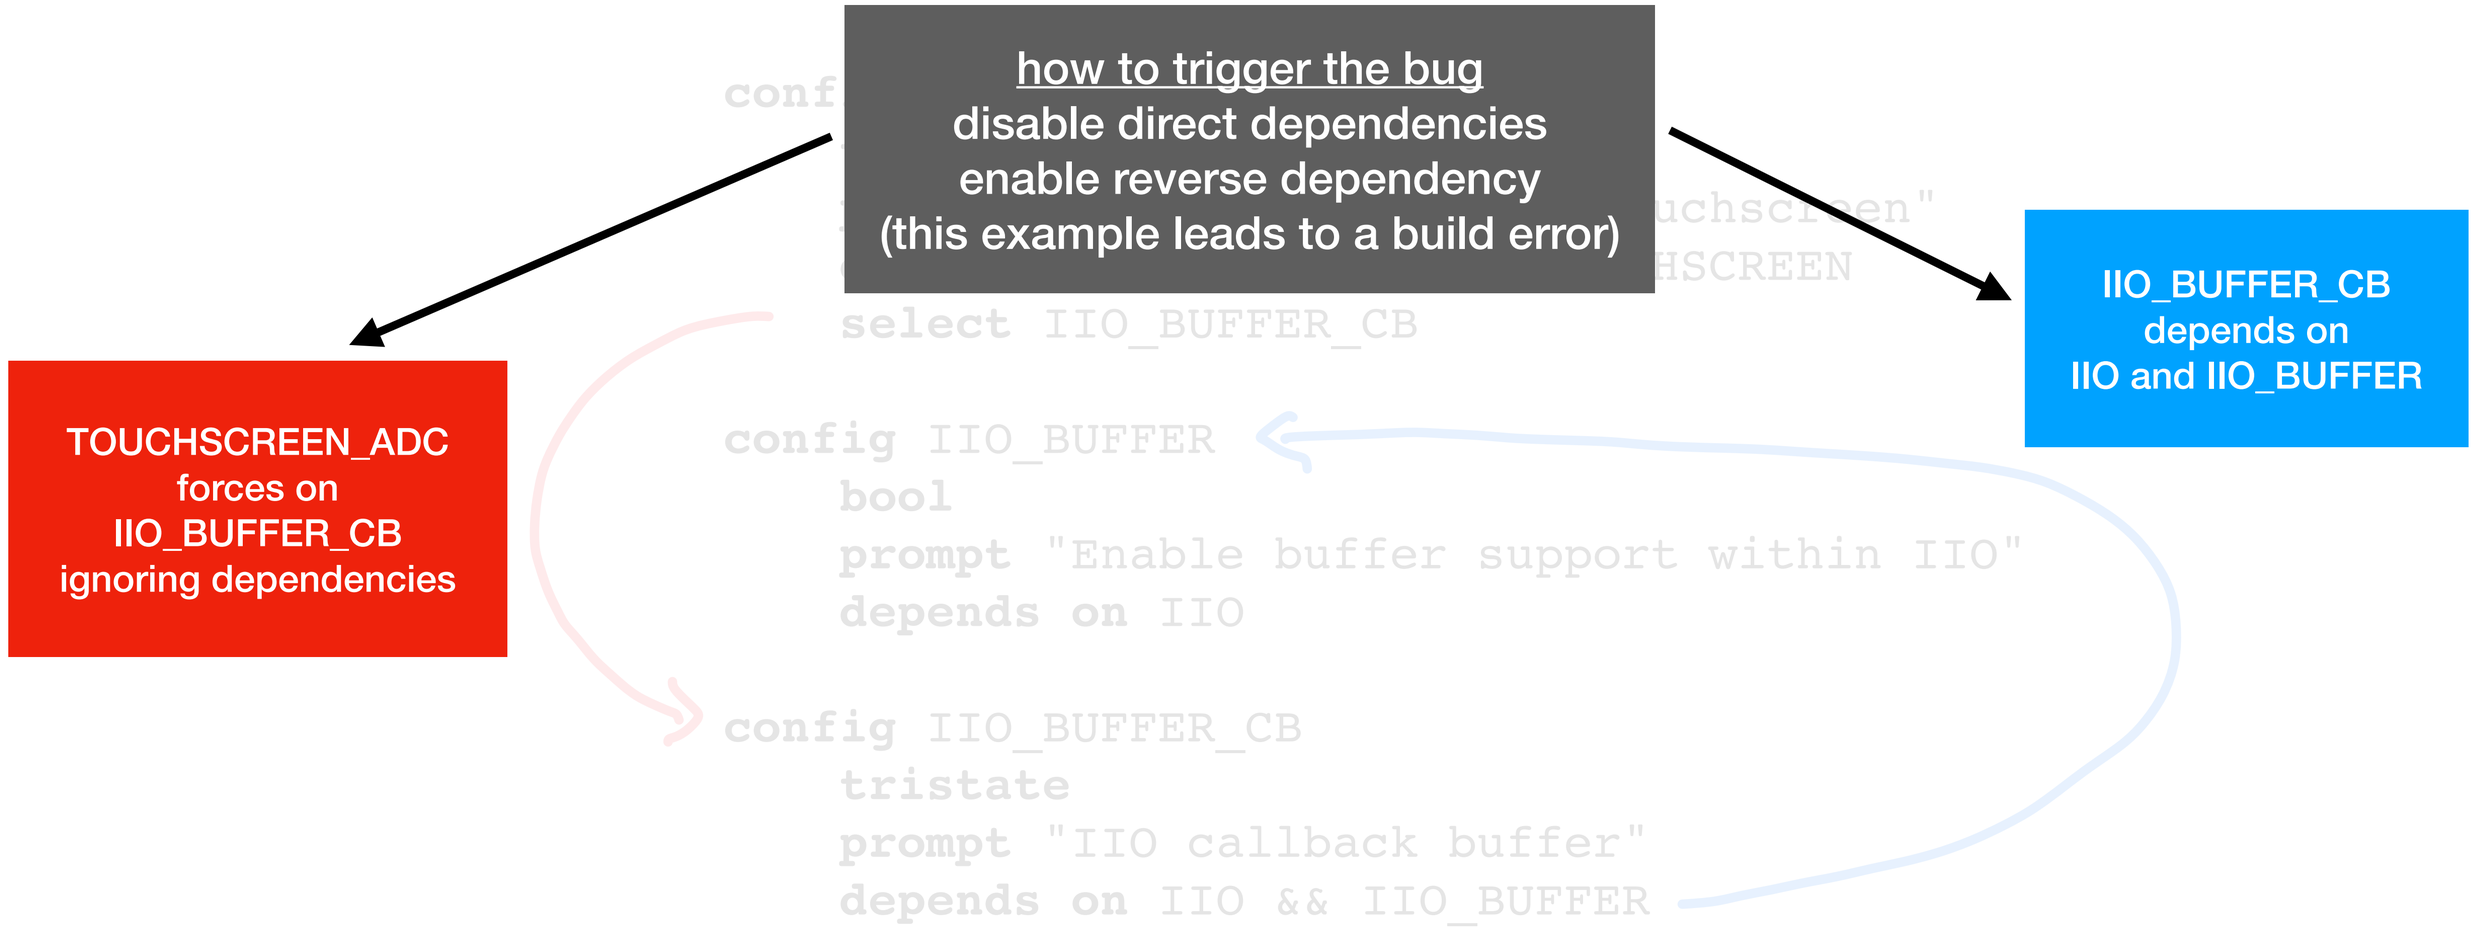

turn “are there unmet dependency bugs?”
into a Boolean satisfiability problem.

use an off-the-shelf solver to get an answer

we first model Kconfig in symbol logic

TOUCHSCREEN_ADC implies IIO and INPUT_TOUCHSCREEN

IIO_BUFFER implies IIO

IIO_BUFFER_CB implies (IIO && IIO_BUFFER or TOUCHSCREEN_ADC)

we check every select for an unmet dependency

if this formula is ever true, i.e., SAT,
then an unmet dependency is possible

TOUCHSCREEN_ADC and (IIO and INPUT_TOUCHSCREEN)
and IIO_BUFFER_CB and not (IIO and IIO_BUFFER)

Integrated into Intel Kernel Test Robot

- **kismet**: unmet depencnecy bug finder
- 100s of unmet dependency bugs found
- https://lore.kernel.org/all/?q=kismet

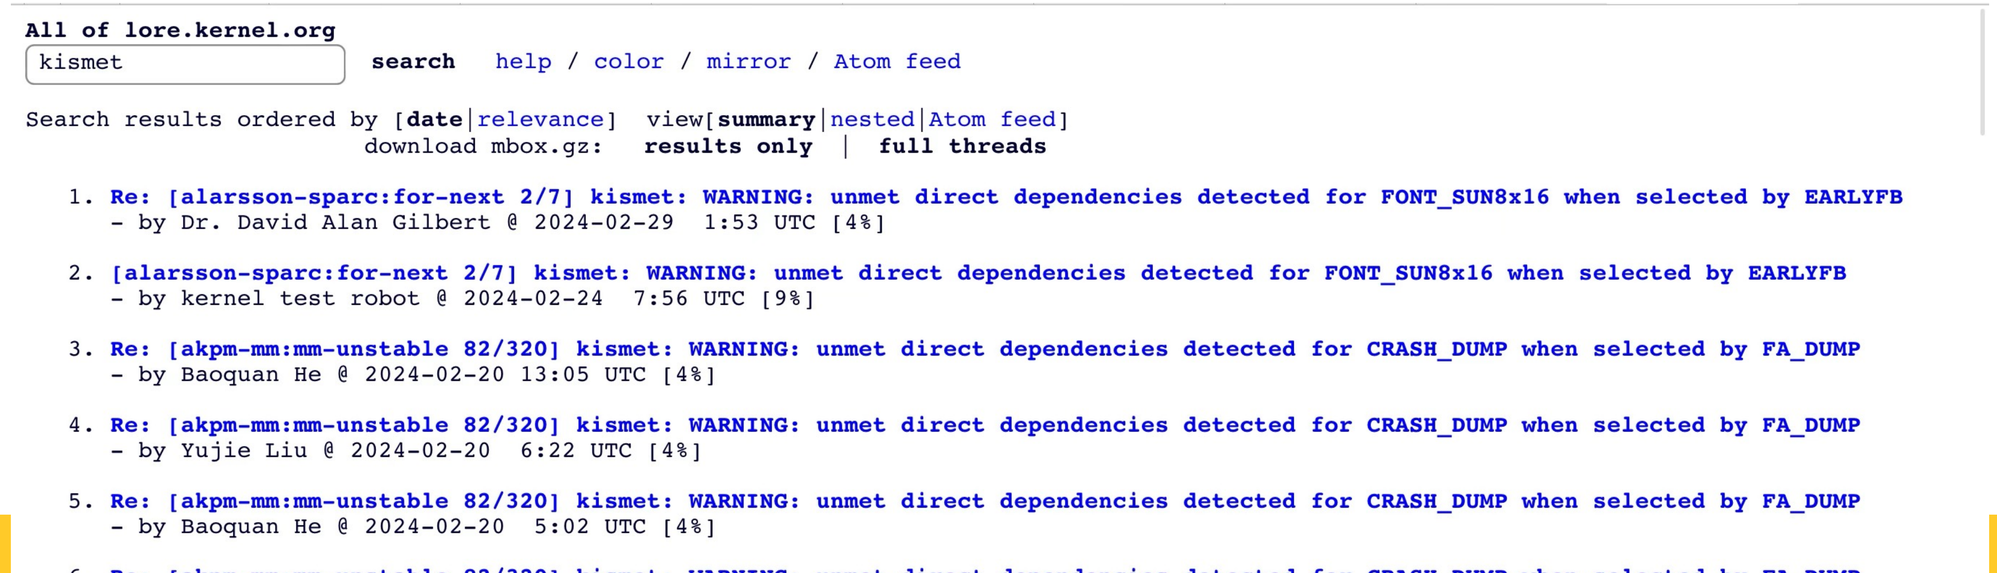

```
> kismet warnings: (new ones prefixed by >>)
> >> kismet: WARNING: unmet direct dependencies detected for FONT_SUN8x16 when
selected by EARLYFB
>
>    WARNING: unmet direct dependencies detected for FONT_SUN8x16
>      Depends on [n]: FONT_SUPPORT [=y] && (FRAMEBUFFER_CONSOLE [=n] && (FONTS [=n]
|| SPARC [=y]) || BOOTX_TEXT)
>      Selected by [y]:
>      - EARLYFB [=y] && SPARC64 [=y]

Hmm just looking at that; I'd seen and nailed those on the PPC side
but hadnt hit them on SPARC before.
(I'm really tempted to simplify the heck out of that line, I'm not sure
why it's anywhere near so complicated)
```

Other Directions in Our Lab

- Static analysis of all configuration simulatenously (SugarC, Varalyzer)
- Static analysis of build system code (SuperC, kmax, kclause, kismet)
- Improving fuzz testing (krepair, KonfFuzz)
- Translating away from C (Maki, Macroni)

PhD Students (Past and Present)

Necip Yildiran
PhD 2022, Google

Jeho Oh (UT Austin)
PhD 2022, Apple
(Co-advised w/Batory)

Brent Pappas
3rd Year PhD

Sanan Hasanov
2nd Year PhD

Find Out More Here

pgazz.com