



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



Towards Learning From Linux Kernel Configurations' Failures with Clang

(for improving Clang support with the kernel)

Prof. Mathieu Acher

@acherm

(University of Rennes 1, Inria/IRISA)





Research: Learning the Linux Configuration space

- 15K+ options; 10^{6000} configurations?
- R1: Finding/fixing/preventing configuration bugs
- R2: Predicting kernel size out of a configuration
 - yes, without compiling it and whatever kernel version (and perhaps compiler version)
 - hint: vmlinux size from 7Mb to 1,8Gb
- R3: Predicting compilation time out of a configuration
- TUXML to build lots of configurations
<https://github.com/TuxML/tuxml>
 - Could well be KernelCI, ongoing “bridge” coming





Sometimes configurations fail

version	nbuild	nfailures	percentage_failure
4.13.3	95854	3164	3.3
4.15	44154	500	1.13
4.20	23830	330	1.38
5.0	20398	446	2.19
5.4	26281	434	1.65
5.7	22343	822	3.68
5.8	25151	806	3.2

make randconfig with
CONFIG_X86=y and
CONFIG_X86_64=y

compiler: gcc 6.3

Disclaimer 1: not all failures are due to an actual bug in the kernel...

Disclaimer 2: some failures are due to the same bug

Disclaimer 3: haven't done yet neither an in-depth analysis nor bug reports or data sharing (if you're interested in, ping me!) for all versions



LINUX PLUMBERS CONFERENCE

August 24-28, 2020



Sometimes configurations fail



Classification problem: predict the class (BUILD/FAILURE) out of options values

```
# Processor type and features
#
# CONFIG_ZONE_DMA is not set
# CONFIG_SMP is not set
# CONFIG_X86_FEATURE_NAMES is not set
# CONFIG_X86_FAST_FEATURE_TESTS is not set
CONFIG_X86_X2APIC=y
CONFIG_X86_MPPARSE=y
CONFIG_GOLDFISHER
# CONFIG_INTEL_RDT_A is not set
```

build **passing**

```
CONFIG_PM_WAKELOCKS=y
CONFIG_PM_WAKELOCKS_LIMIT=100
CONFIG_PM_WAKELOCKS_GC=y
CONFIG_PM=y
# CONFIG_PM_DEBUG is not set
CONFIG_PM_CLK=y
CONFIG_PM_GENERIC_DOMAINS=y
CONFIG_WQ_POWER_EFFICIENT_DEFAULT=y
CONFIG_PM_GENERIC_DOMAINS_SLEEP=y
```

build **passing**

```
CONFIG_VM_EVENT_COUNTERS=y
CONFIG_SLUB_DEBUG=y
# CONFIG_SLUB_MEMCG_SYSFS_ON is not set
# CONFIG_COMPAT_BRK is not set
# CONFIG_SLAB is not set
CONFIG_SLUB=y
# CONFIG_SLOB is not set
GE_DEFAULT is not set
ELIST_RANDOM is not set
ELIST_HARDENED is not set
GE_ALLOCATOR=y
ARTIAL=y
A_VERIFICATION=y
```

```
CONFIG_PM_WAKELOCKS=y
CONFIG_PM_WAKELOCKS_LIMIT=100
CONFIG_PM_WAKELOCKS_GC=y
CONFIG_PM=y
# CONFIG_PM_DEBUG is not set
CONFIG_PM_CLK=y
CONFIG_PM_GENERIC_DOMAINS=y
CONFIG_WQ_POWER_EFFICIENT_DEFAULT=y
CONFIG_PM_GENERIC_DOMAINS_SLEEP=y
```

build **passing**

X86_64	TRACING	...	DRM_VBOXVIDEO	GENERIC_ALLOCATOR	DRM_TTM	BUILD?	SIZE (Mb)
1	0	...	1	1	0	BUILD	8.3
1	1	...	1	0	1	FAILURE	78.9
1	0	...	1	1	0	BUILD	97.3
1	1	...	1	1	1	BUILD	22.4
1	0	...	1	0	2	FAILURE	34.8
1	1	...	1	1	1	BUILD	176.8
1	1	...	1	1	0	BUILD	54.3

Do you recognize a pattern?
(matrix is 95K rows and 12K columns)



**LINUX
PLUMBERS
CONFERENCE**

August 24-28, 2020

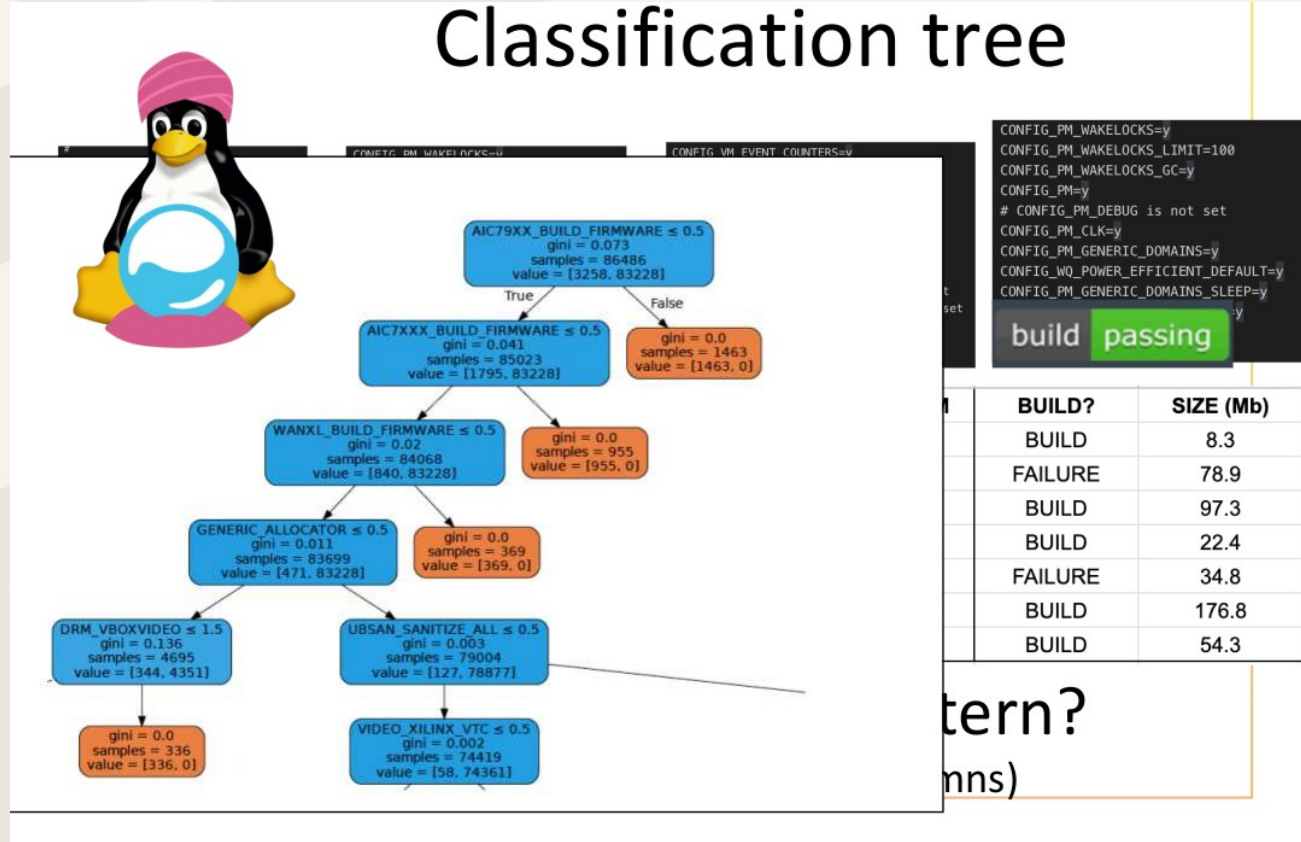


Configuration bug location/
understanding/prevention:
**TUXML can help to
pinpoint
faulty (combinations of
options**

(even more effective when
combined with failures
clustering, see details in
research papers)

Sometimes configurations fail

Classification tree



tern?
nn)



Sometimes configurations fail

- A failure raises two questions:
 - Is it a bug of the Linux kernel?
 - Avoid false positives or duplication of bug reports
 - Pinpoint responsible options (usually it's not one individual option)
 - Preventive fix at Kconfig level to avoid build of configurations that will fail anyway
 - Synthesize a fix or wait for a fix (humans to the rescue!)
 - Is it a bug of the compilation environment?
 - No it's the Linux doc ;)
 - CI service bug e.g., TUXML, not KernelCI ;)
 - **Compiler "bug"** e.g., gcc, not clang ;)



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



Claim: Configuration failures can... help!

- To pinpoint Clang-specific issues
 - compare with gcc: same build status?
- Identify problematic combinations of options for Clang
 - Avoid duplication of bug reports
 - Preventive fix to jump to other problematic cases
 - Help understanding/fixing Clang-related bugs
 - You can know where the problem potentially comes from
 - Help improving Clang



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



Design Study (brainstorm)

- **Large-scale compilation with Clang of random configs.**
- Which clang version? (and other specificities: clang options?)
- Which architecture? (eg x86_64?)
- Which Linux version? (eg 5.4 LTS?)
- Which configurations?
 - randconfig
 - which options should be (de-)activated anyway?
- Which observations? (warnings? build status? size? compilation time?)
- Hypothesis (poll)
 - % build failures with clang?
 - (much) more than gcc?



Construct Validity (brainstorm)

- Too costly?
 - ~260 seconds per config (with gcc/x86_64) on 16-core, recent machines
- Too many failures? (let's be provocative)
 - Overwhelming: “we” have already tons of bugs to resolve and fix...
 - Irritating: fixing “esoteric” configurations may decrease motivation or complicate the task
- We don't know after all... so let's try!
- **Seeking volunteers to help “designing” and technically realizing this experiment**
 - Can well be done with KernelCI/kcidb
 - Need Clang-specific knowledge



**LINUX
PLUMBERS
CONFERENCE**

August 24-28, 2020



Supplementary material



LINUX
PLUMBERS
CONFERENCE

August 24-28, 2020



Research papers

<https://hal.inria.fr/hal-02147012/>

Learning From Thousands of Build Failures of Linux Kernel Configurations

- Mathieu Acher, Hugo Martin, Juliana Alves Pereira, Arnaud Blouin, Djamel Eddine Khelladi, Jean-Marc Jézéquel
- <https://hal.inria.fr/hal-02147012>



Learning Very Large Configuration Spaces: What Matters for Linux Kernel Sizes

- Mathieu Acher, Hugo Martin, Juliana Pereira, Arnaud Blouin, Jean-Marc Jézéquel, Djamel Eddine Khelladi, Luc Lesoil, Olivier Barais
- <https://hal.inria.fr/hal-02314830>



<https://hal.inria.fr/hal-02314830>

Talk at Embedded Linux Conference Europe (2019, Lyon)

<https://www.youtube.com/watch?v=UBghs-cwQX4&feature=youtu.be>



Team



- DiverSE research team <http://diverse-team.fr>
 - Hugo Martin, Juliana Alves Pereira, Arnaud Blouin, Jean-Marc Jézéquel, Djamel Eddine Khelladi, Luc Lesoil, Olivier Barais to name a few
- TUXML team at ISTIC / University of Rennes 1
 - Paul Saffray, Alexis Le Masle, Michaël Picard, Corentin Chédotal, Gwendal Didot, Dorian Dumanget, Antonin Garret, Erwan Le Flem, Pierre Le Luron, Mickaël Lebreton, Fahim Merzouk, Valentin Petit, Julien Royon Chalendar, Cyril Hamon, Luis Thomas, Alexis Bonnet
- Georges Aaron Randrianaina (ENS Rennes)
- IGRIDA <http://igrida.gforge.inria.fr>
 - Morillon Pascal, Deltel Charles