




# **Place Your Locks Well: Understanding and Detecting Lock Misuse Bugs**

**Yuandao Cai, Peisen Yao, Chengfeng Ye, Charles Zhang**  
The Hong Kong University of Science and Technology

# Increasingly Complex Modern Software

- Massive codebases and high concurrency



- Poor software quality costs the US economy **\$2.41 trillion** annually



# Synchronization Primitives against Harmful Races

- Concurrency bugs are hard-to-avoid and extremely harmful!
- Synchronization primitives are in place to synchronize concurrent code
  - preventing various concurrency bugs and vulnerabilities
- Existing works focus on **concurrency bug detection** (**insufficient synchronization**)
  - data races [AI Thokair POPL'23] [chabbi et.al., PLDI'22]...
  - concurrency memory corruption bugs [Yuan et.al., Security'23] [Cai et.al., PLDI'21]...
  - concurrency typestate bugs [ASPLOS'11]...
- Our work focuses on **the misuses of synchronization APIs themselves**
  - currently focusing on locks
  - also causing serious reliability and security issues



# Research Goal and Contributions

1. Understanding the common misuses of locks
  - through a CVE-ID-based empirical study
2. Designing techniques to detect the lock misuses
3. Evaluating and advancing the state-of-the-art bug-finding tools



# An Empirical Study: Setup

- Locks are common synchronization primitives
  - with explicit disciplines for initialization, use, and destruction.
- Study Question 1: What are **the common lock misuses**?
- Study Question 2: What are **the common causes** of those lock misuses?
- **Study Dataset:** 32 CVE IDs assigned between 2010-2021
  - search keywords: e.g., mutex, lock
  - manual validation for CVE ID description



# An Empirical Study: Finding I

## 1. Identifying five general locking discipline violations

- under both sequential and concurrent circumstances
- covering a single thread and multiple threads



## 2. Defining the bug patterns by revealing their characteristics

No.	Misuse Pattern	Bug Description	Concurrency
①	Missing lock releases	A lock is not released after its effective lifetime.	
②	Double locking	A lock is acquired twice.	
③	Using uninitialized locks	A lock is not initialized before using it. A concurrency error occurs when the lock is initialized non-deterministically.	✓
④	Releasing unacquired locks	A lock is released without acquiring it first. A concurrency error occurs when there is another thread holding the lock.	✓
⑤	Cyclic lock acquisitions	Different locks are not acquired in the same order. A concurrency error occurs when each thread in a set waits for the other to release a lock.	✓

# An Empirical Study: Five General Lock Misuses

```
373 static int open_console (UI *ui){
375     if (!CRYPTO_THREAD_write_lock(ui->lock))
376         return 0;
483 }
552 static int close_console (UI *ui){
560     if (status != SS$NORMAL) {
561         ERR_raise_data(..., status);
563         return 0;
564     }
566     CRYPTO_THREAD_unlock(ui->lock);
368     return 1;
369 }
```

```
82 CEN64_THREAD_RETURN_TYPE gdb_thread(...) {
85     pthread_mutex_lock(&gdb->client_mutex);
88     if (gdb->flags & GDB_FLAGS_INITIAL) {
89         pthread_cond_wait(..., &gdb->c_mutex);
90     } else {
91         pthread_mutex_lock(&gdb->client_mutex);
92     }
97     pthread_mutex_unlock(&gdb->client_mutex);
143 }
```

```
330 ret_t cherokee_collector_rrd_new (...){
373     re = pthread_create (..., worker_func, n);
375     ...
379     re = pthread_mutex_init (&n->mutex, NULL);
380     if (re != 0) {
382         return ret_error;
383     }
389 }
```

(1) Missing lock releases (OpenSSL)

(2) Double locking (Cen64)

(3) Using uninitialized locks (Cherokee)

```
459 BIO *OSSL_trace_begin(int category){
465     category = ossl_trace_get_category(category);
466     if (category < 0)
467         return NULL;
473     if (!CRYPTO_THREAD_write_lock(trace_lock))
474         return NULL;
491 }
493 void OSSL_trace_end(int category, BIO * channel){
498     category = ossl_trace_get_category(category);
516     CRYPTO_THREAD_unlock(trace_lock);
519 }
```

```
381 static void *extract_worker_thread_func(...){
385     pthread_mutex_lock(ctxt->mutex);
431     pthread_mutex_lock(&entry->mutex);
433     pthread_mutex_unlock(ctxt->mutex);
442     if (chunk.type == XB_CHUNK_TYPE_EOF) {
443         pthread_mutex_lock(ctxt->mutex);
444         pthread_mutex_unlock(&entry->mutex);
445         my_hash_delete(ctxt->filehash, ...);
446         pthread_mutex_unlock(ctxt->mutex);
470     }
478 }
```

(4) Releasing unacquired locks (OpenSSL)

(5) Cyclic lock acquisitions (MariaDB)

# An Empirical Study: Finding II

- Wreaking severe havoc by triggering lock misuses
  - **denial-of-service with system hang** (concurrent cyclic acquisitions, double locking)
    - CVE-2013-4553, CVE-2014-8131, CVE-2019-14763, CVE-2021-41213,...
  - **memory exhaustion with memory leak** (missing lock releases)
    - CVE-2004-2650, CVE-2018-14660, CVE-2020-12658,...
  - **memory corruption; system crash** (releasing unacquired locks, using uninitialized locks)
    - CVE-2014-1453, CVE-2015-8767, CVE-2017-6353, CVE-2020-10573,...
  - **even privilege escalation and other unidentified issues**
    - CVE-2010-4210, CVE2014-9748, ...
- Relating to other security bugs
  - **atomicity violations** (CVE-2020-10573)
  - **use-after-free** (CVE-2019-14034)
  - **double free** (CVE-2017-6353)

```
98 int search_makelist(search_t *results,...){
145     pthread_mutex_unlock(&conn->lock);
146     int tmp = conn_setup(conn);
147     pthread_mutex_unlock(&conn->lock);
203 }
```


Releasing unacquired locks that leads to **atomicity violations** (Axel)



# Detecting the Five Lock Misuses with Lockpick

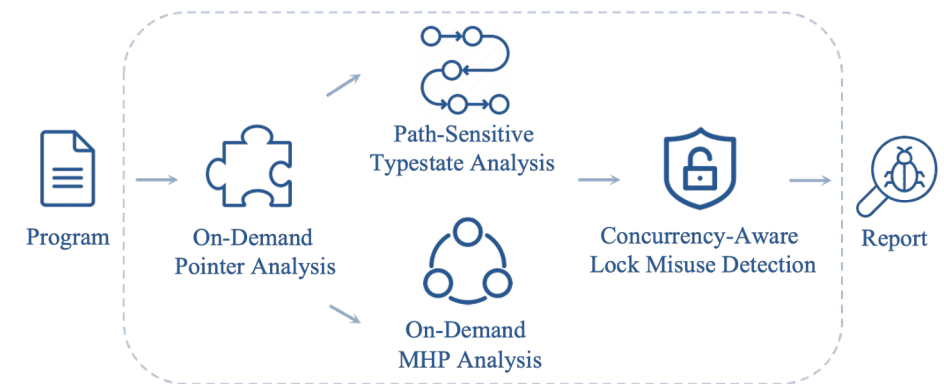
- **Lock misuse formulation:** characterizing lock misuses with a **finite-state machine (FSM)**

1. Model the states of lock objects using tpestates
2. Capture the state transitions of lock objects with a new FSM
3. Capture the lock misuses by tracking the state transitions




- **Lock misuse detection:** detecting lock misuses with **several customized techniques**

1. Path-sensitively track the tpestates of locks
2. Reason about the MHP relations of statements
3. Flag the lock misuses based on tpestate violations




# Implementation and Experiment Setup

- Lockpick is built upon the LLVM infrastructure and the Z3 SMT solver
  - a **soundy** implementation to reach both high efficiency and precision
    - unrolling loops twice, ignoring inline assembly, pointer arithmetic
  - a value-flow-based pointer analysis
    - on-demand flow-, context-sensitive pointer analysis
  - path conditions are encoded as first-order logic formulae over bit-vectors
- Question 1: How **effective and practical** is Lockpick at uncovering lock misuses in mature open-source software systems?
- Question 2: How does Lockpick perform **compared to the state-of-the-art tools**?




# (1) Highlights: Effectiveness on Bug Finding

- Finding **203 developer-confirmed bugs** in over 80 well-checked software programs
  - **184** of them have been fixed (at the time of publication)
  - finding various kinds of bugs
  - hiding for an average of **7.4** years



The distributions of bug type




The distributions of hidden time (Year)



- **16 CVE IDs** have been assigned for multiple bugs with high security impacts
  - CVE-2021-41141, CVE-2021-43429, CVE-2022-31621, CVE-2022-31624, CVE-2022-31623, CVE-2022-31622, CVE-2022-30027, CVE-2022-37869, CVE-2022-37868, CVE-2022-38791, CVE-2022-37874, CVE-2022-37875, CVE-2022-37876, CVE-2022-37871, CVE-2022-37872...

## (2) Highlights: Advancement over Previous Tools

- **Baselines:** SVF, L2D2 (built on Infer), Clang Static Analyzer
- **Benchmarks:** ten popular software programs with **35.8 MLoC**
- **Efficiency:** being able to analyze big programs like Linux kernel in about five hours
- **Precision:** embracing better precision than other tools
- **Recall:** being able to discover **26 past CVE IDs** in C/C++ programs (2010-2021)
  - other tools cannot reach



Project	KLoC	SVF		L2D2		CSA		LOCKPICK	
		#FP	#R	#FP	#R	#FP	#R	#FP	#R
Cherokee	55	3	6	99% <sup>†</sup>	483	22	26	1	5
Curl	135	3	4	0	0	0	0	0	2
PJSIP	434	32	43	98% <sup>†</sup>	505	0	2	2	15
OpenSIPS	477	25	55	0	0	0	0	5	40
OpenSSL	490	66	68*	0	0	0	0	2	6
WolfSSL	944	16	20	3	3	0	1	3	11
MySQL	4,152	0	0*	100% <sup>†</sup>	1,157	95	99*	3	10
MariaDB	4,697	96% <sup>†</sup>	141*	100% <sup>†</sup>	4,993	100% <sup>†</sup>	229*	9	27
FreeBSD	8,457	66	81	NA	NA	0	0	12	31
Linux	15,987	88% <sup>†</sup>	328*	NA	NA	0	0	19	57
<b>FPR</b>	—	85.1%		99.2%		93.7%		27.5%	

**Thank you for your listening!**

# Questions & Answers

More details can be found in our paper:

[https://www.usenix.org/system/files/sec23fall-prepub-298\\_cai-yuandao.pdf](https://www.usenix.org/system/files/sec23fall-prepub-298_cai-yuandao.pdf)

Bug and CVE ID lists can be found:

<https://drive.google.com/file/d/1HY7PydeDga-850ZOn3YPACnX7hRws8DG/view>