




Does RCU Really Work?


And if so, how would we know?



Isn't Making Software Work A Solved Problem?

Paul's Installed Base Over The Past Four Decades

Million-Year Bug: Once per million years



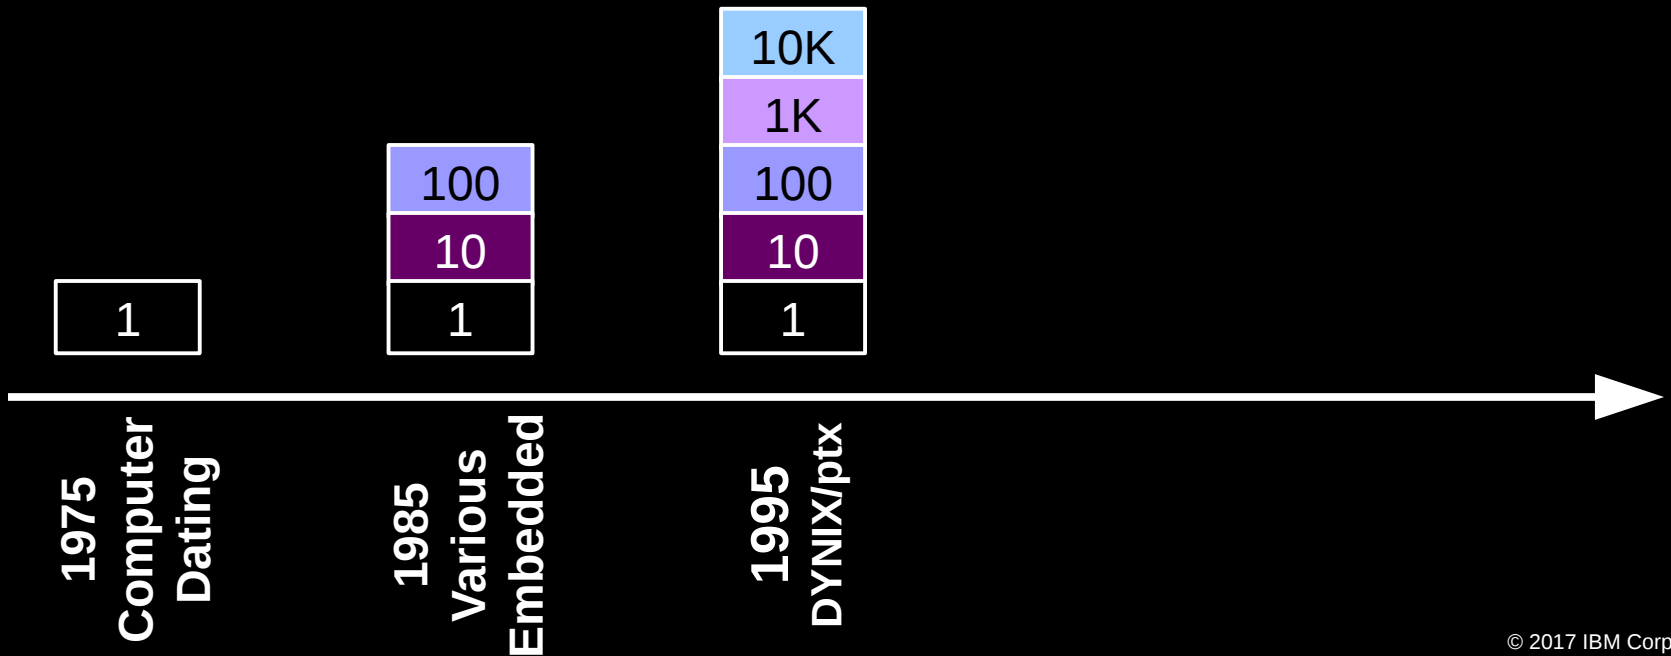
Paul's Installed Base Over The Past Four Decades

Million-Year Bug: Once per ten millenia



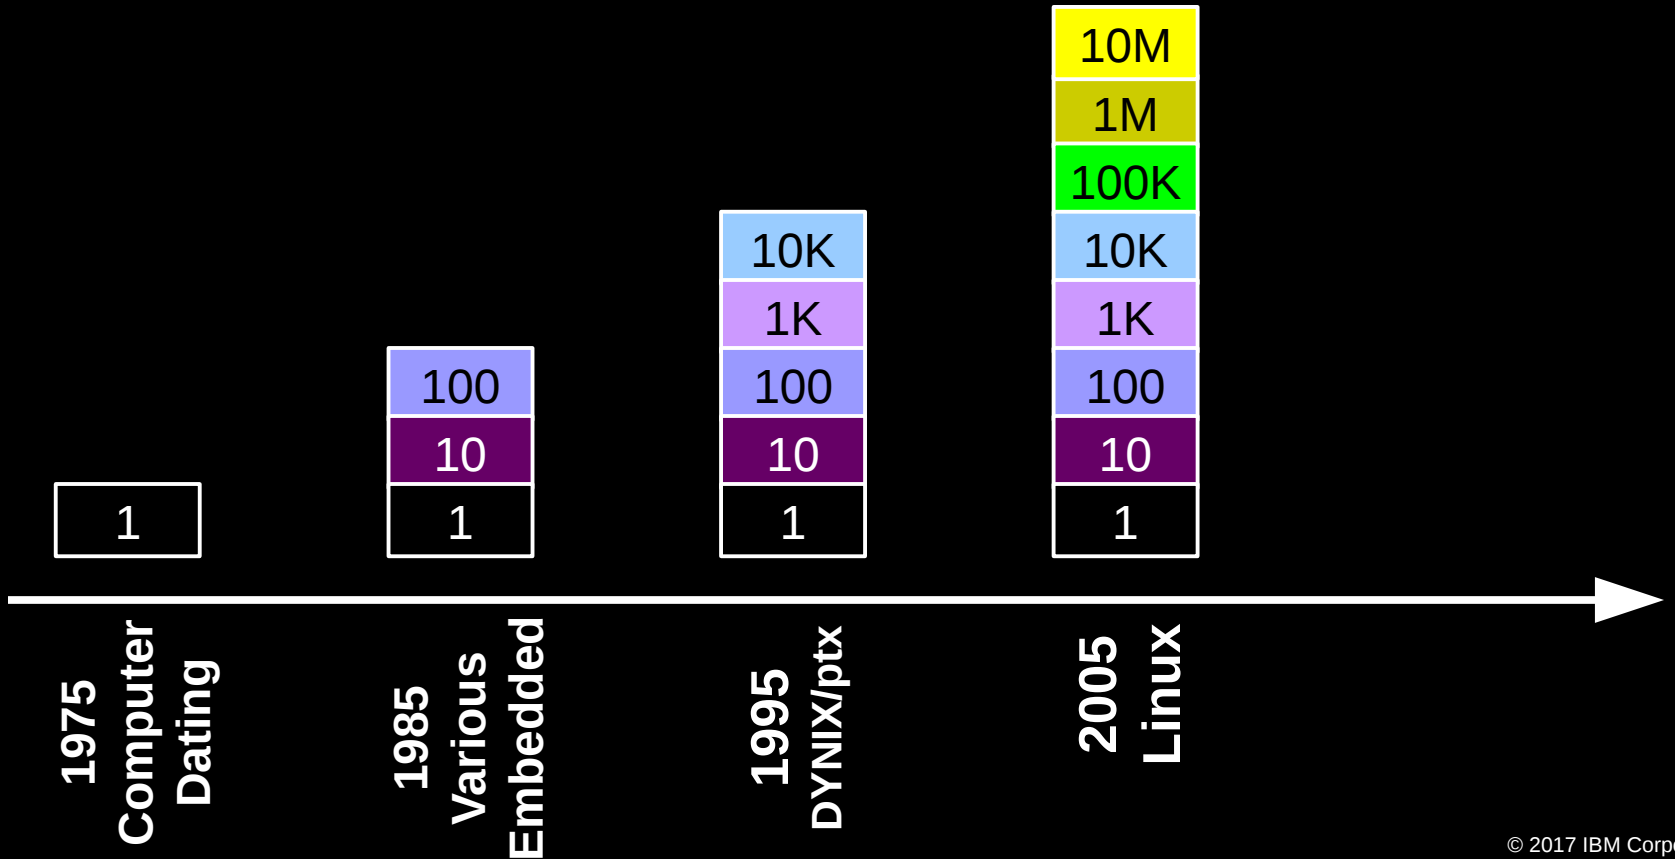
Paul's Installed Base Over The Past Four Decades

Million-Year Bug: Once per century




Paul's Installed Base Over The Past Four Decades

Million-Year Bug: Once a month




Paul's Installed Base Over The Past Four Decades

Million-Year Bug: Several Times per Day



Internet of Things, Anyone???

**Million-Year Bug? You don't want to know...
But Murphy is still alive and kicking!**



Why Stress About Potential Low-Probability Bugs?

- Almost any bug might become a security exploit
 - Internet access means physical presence no longer required
- RCU's low level does not necessarily mean low risk
 - If Row Hammer can hit DRAM, RCU is not invulnerable
- Internet of Things could mean a trillion computers on Earth
 - Even low failure probability translates to huge numbers of failures
 - Some of which might put the general public at risk
 - Linux is already used in some safety-critical applications
 - Murphy transitions from nice guy to real jerk to homicidal maniac
- It is therefore not too early to think about reducing risk
 - And RCU is a good well-contained test case for proofs of concept

Does RCU Really Work? If So, How Would We Know?

Does RCU Really Work? If So, How Would We Know?


- What is RCU (read-copy update) supposed to do?
- What are the odds of RCU “just working”?
- RCU validation

What is RCU Supposed To Do?


What is RCU Supposed To Do? (Brief Overview!)

- Structured deferral: synchronization via procrastination
 - The waiters: *RCU grace periods*
 - `synchronize_rcu()`, `call_rcu()`, ...
 - The waited upon: *RCU read-side critical sections*
 - `rcu_read_lock()` and `rcu_read_unlock`, ...
 - RCU's read-side primitives have exceedingly low overhead, great scalability
- RCU grace periods must wait for pre-existing RCU read-side critical sections
 - How could this possibly be useful? See next slides...
- Other examples of synchronization via procrastination:
 - Reference counting, sequence locking, hazard pointers, garbage collectors
 - Arguably also locking (new acquisition must wait for old acquisition)


What RCU Is Supposed To Do and Not...



RCU Area of Applicability



RCU Applicability to the Linux Kernel



In 1996, I thought I knew everything there was to know about RCU
The Linux kernel community proved me wrong many times!!!

What Are The Odds of RCU “Just Working”?

Two Definitions and a Consequence

Two Definitions and a Consequence

- A ***bug-free software system*** is a trivial software system
- A ***reliable software system*** contains no known bugs

Two Definitions and a Consequence

- A bug-free software system is a trivial software system
- A reliable software system contains no known bugs

- Therefore, any non-trivial reliable software system contains at least one bug that you don't know about
- I assert that Linux-kernel RCU is both non-trivial and reliable, thus containing at least one bug that I don't (yet) know about

Two Definitions and a Consequence

- A bug-free software system is a trivial software system
- A reliable software system contains no known bugs

- Therefore, any non-trivial reliable software system contains at least one bug that you don't know about
- I assert that Linux-kernel RCU is both non-trivial and reliable, thus containing at least one bug that I don't (yet) know about
- But how many bugs?
 - Analyze from a software-engineering viewpoint...

Software-Engineering Analysis

Software-Engineering Analysis

- RCU contains 11,534 lines of code (including comments, etc.)
- 1-3 bugs/KLoC for production-quality code: **11-36 bugs**
 - Best case I have seen: 0.04 bugs/KLoC for safety-critical code
 - Extreme code-style restrictions, single-threaded, formal methods, ...
 - And still way more than zero bugs!!! :-)
- Median age of a line of RCU code is less than four years
 - And young code tends to be buggier than old code!

- We should therefore expect a few tens more bugs in RCU!

RCU Validation

Current RCU Regression Testing

- Stress-test suite: “rcutorture”
 - <http://lwn.net/Articles/154107/>, <http://lwn.net/Articles/622404/>
- “Intelligent fuzz testing”: “trinity”
 - <http://codemonkey.org.uk/projects/trinity/>
- Test suite including static analysis: “0-day test robot”
 - <https://lwn.net/Articles/514278/>
- Integration testing: “linux-next tree”
 - <https://lwn.net/Articles/571980/>
- Above is old technology – but quite effective
 - 2010: wait for -rc3 or -rc4. 2013: Usually no problems with -rc1
- Formal verification in design, but not in regression testing
 - <http://lwn.net/Articles/243851/>, <https://lwn.net/Articles/470681/>,
<https://lwn.net/Articles/608550/>

January 30, 2017 rcutorture Output

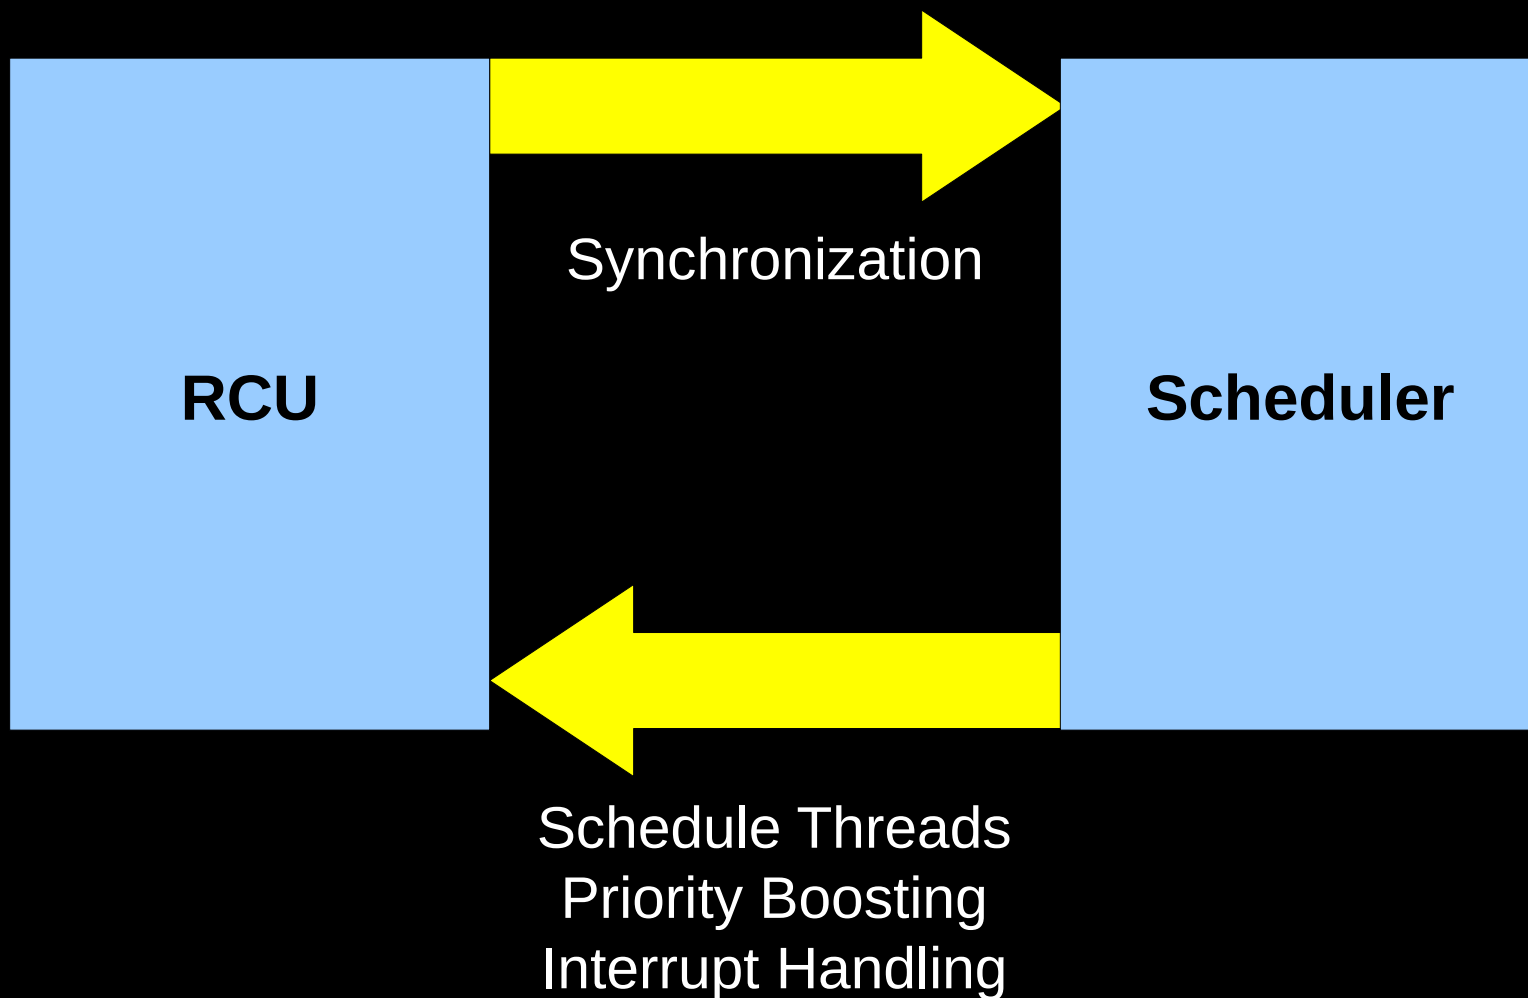
```
tools/testing/selftests/rcutorture/bin/kvm.sh --cpus 50 --duration 1800
SRCU-N ----- 610414 grace periods (5.65198 per second)
SRCU-P ----- 13349 grace periods (0.123602 per second)
TASKS01 ----- 70971 grace periods (0.657139 per second)
TASKS02 ----- 70238 grace periods (0.650352 per second)
TASKS03 ----- 69972 grace periods (0.647889 per second)
TINY01 ----- 8152793 grace periods (75.4888 per second)
TINY02 ----- 17916244 grace periods (165.891 per second)
TREE01 ----- 4376468 grace periods (40.5229 per second)
TREE02 ----- 3034531 grace periods (28.0975 per second)
TREE03 ----- 1048736 grace periods (9.71052 per second)
TREE04 ----- 637788 grace periods (5.90544 per second)
TREE05 ----- 2415024 grace periods (22.3613 per second)
TREE06 ----- 1791390 grace periods (16.5869 per second)
TREE07 ----- 551532 grace periods (5.10678 per second)
TREE08 ----- 1072103 grace periods (9.92688 per second)
TREE09 ----- 7543572 grace periods (69.8479 per second)
```

There are bugs in RCU, and 30 hours of rcutorture failed to find them
This constitutes a critical bug in rcutorture

On the other hand, first time in over a year that I have see this!

How Well Does Linux-Kernel Testing Really Work?

Example 1: RCU-Scheduler Mutual Dependency



So, What Was The Problem?

- Found during testing of Linux kernel v3.0-rc7:
 - RCU read-side critical section is preempted for an extended period
 - RCU priority boosting is brought to bear
 - RCU read-side critical section ends, notes need for special processing
 - Interrupt invokes handler, then starts softirq processing
 - Scheduler invoked to wake ksoftirqd kernel thread:
 - Acquires runqueue lock and enters RCU read-side critical section
 - Leaves RCU read-side critical section, notes need for special processing
 - Because `in_irq()` returns false, special processing attempts deboosting
 - Which causes the scheduler to acquire the runqueue lock
 - Which results in self-deadlock
 - (See <http://lwn.net/Articles/453002/> for more details.)
- Fix: Add separate “exiting read-side critical section” state
 - Also validated my creation of correct patches – without testing!

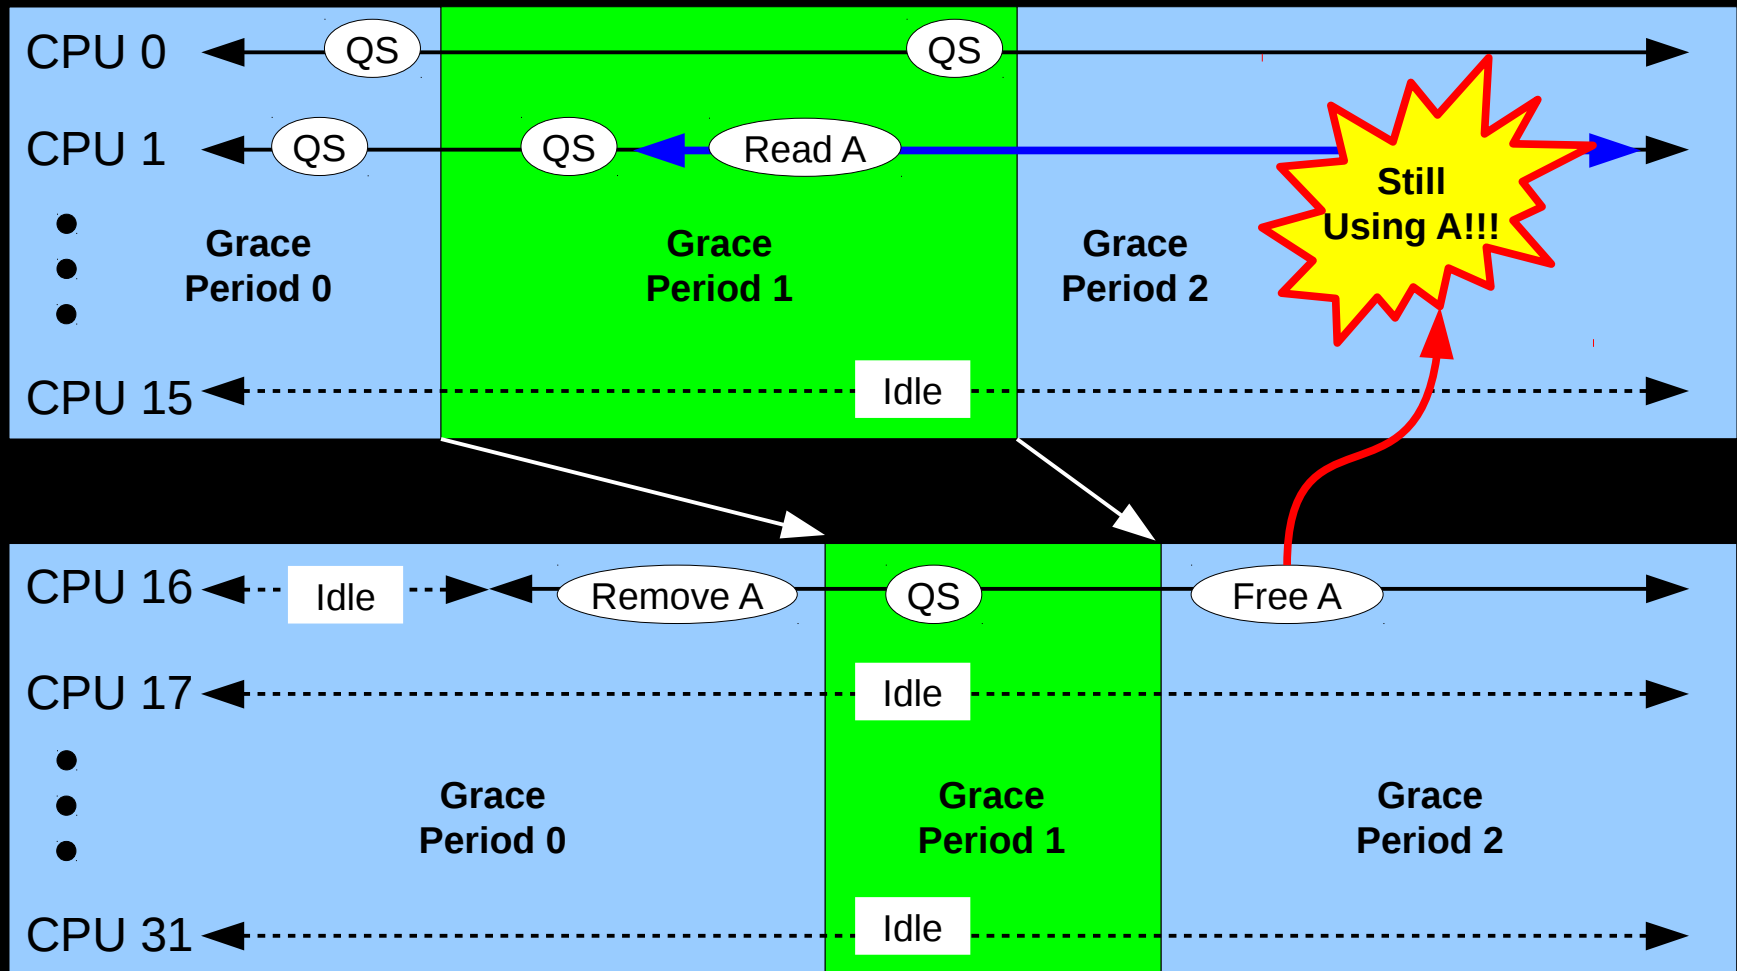
Example 1: Bug Was Located By Normal Testing

Example 2: Grace Period Cleanup/Initialization Bug

1. CPU 0 completes grace period, starts new one, cleaning up and initializing up through first leaf rcu_node structure
2. CPU 1 passes through quiescent state (new grace period!)
3. CPU 1 does rcu_read_lock() and acquires reference to A
4. CPU 16 exits dyntick-idle mode (back on *old* grace period)
5. CPU 16 removes A, passes it to call_rcu()
6. CPU 16 associates callback with next grace period
7. CPU 0 completes cleanup/initialization of rcu_node structures
8. CPU 16 callback associated with now-current grace period
9. All remaining CPUs pass through quiescent states
10. Last CPU performs cleanup on all rcu_node structures
11. CPU 16 notices end of grace period, advances callback to “done” state
12. CPU 16 invokes callback, freeing A (*too bad CPU 1 is still using it*)

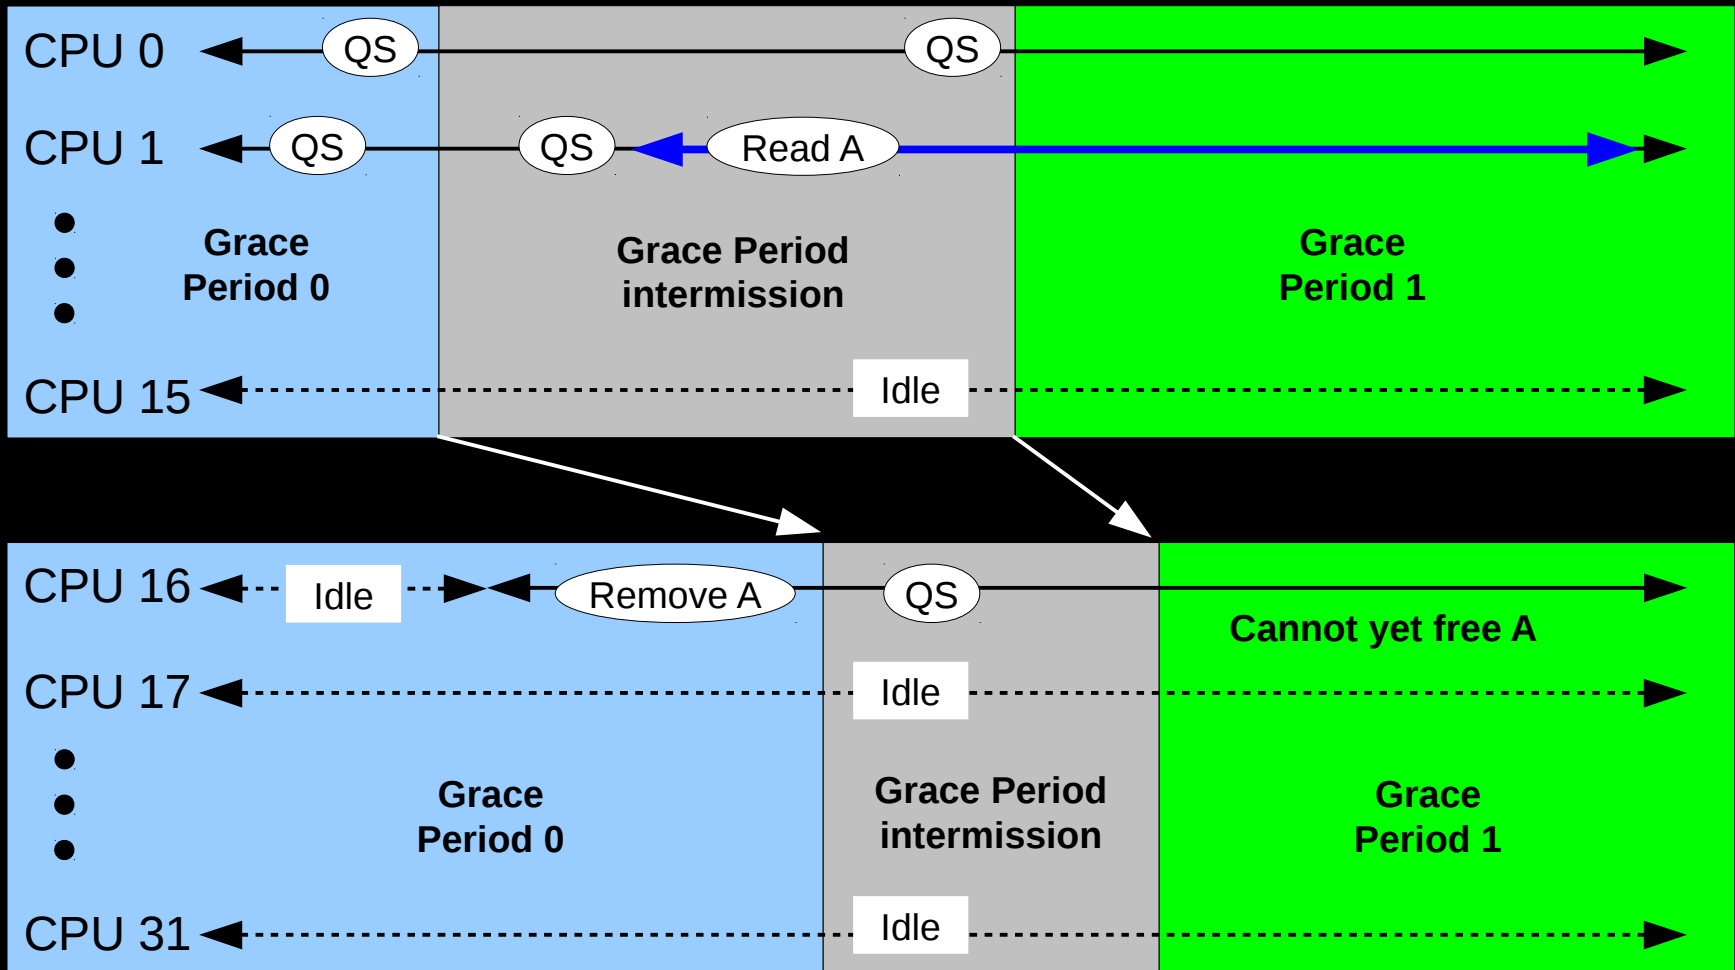
Not found via Linux-kernel validation: In production for 5 years!

Example 2: Grace Period Cleanup/Initialization Bug



Note: Remains a bug even under SC

Example 2: Grace Period Cleanup/Initialization Fix



Example 1 & Example 2 Results

- Example 1: Bug was located by normal Linux test procedures
- Example 2: Bug was missed by normal Linux test procedures
 - Not found via Linux-kernel validation: In production for 5 years!
 - On systems with up to 4096 CPUs...
- Both are bugs even under sequential consistency
- Normal testing is not bad, but improvement is needed
- Can Linux-kernel RCU validation do better?
- But first, what is the validation problem that must be solved?

**More Than 1.5 Billion Linux Instances Running!!!
Woo-Hoo!!! Linux Has Won!!!**

**More Than 1.5 Billion Linux Instances Running!!!
Woo-Hoo!!! Linux Has Won!!!**

But How The #@\$&! Do I Validate RCU For This???

How The #@\$&! Do I Validate RCU For This???

- A race condition that occurs once in a million years happens ***several times per day*** across the installed base
 - I am very proud of rcutorture, but it simply cannot detect million-year races when running on a reasonable test setup
 - Even given expected improvements in rcutorture
 - Even with help from mutation testing
 - Groce et al., “How Verified is My Code? Falsification-Driven Verification”
<https://www.cs.cmu.edu/~agroce/ase15.pdf>

RCU Validation Options?

- Other failures mask RCU's, including hardware failures
 - I know of no human artifact with a million-year MTBF
 - But I do know of Linux uses that put the public's safety at risk...
- Increasing CPUs on test system increases race probability
- Rare critical operations forced to happen more frequently
- Knowledge of possible race conditions allows targeted tests
 - Plus other dirty tricks from 25 years of testing concurrent software
 - Provide harsh environment to force software to evolve quickly
- Formal verification used for some aspects of RCU design

RCU Validation Options?

- Other failures mask RCU's, including hardware failures
 - I know of no human artifact with a million-year MTBF
 - But I do know of Linux uses that put the public's safety at risk...
- Increasing CPUs on test system increases race probability
- Rare critical operations forced to happen more frequently
- Knowledge of possible race conditions allows targeted tests
 - Plus other dirty tricks from 25 years of testing concurrent software
 - Provide harsh environment to force software to evolve quickly
- Formal verification used for some aspects of RCU design

- Should I use formal verification in RCU's regression testing?

Formal Verification and Regression Testing: Requirements

Formal Verification and Regression Testing: Requirements

- (1) Either automatic translation or no translation required
 - Automatic discarding of irrelevant portions of the code
 - Manual translation provides opportunity for human error

Formal Verification and Regression Testing: Requirements

- (1) Either automatic translation or no translation required
- (2) Correctly handle environment, including memory model
- (3) Reasonable memory and CPU overhead
- (4) Map to source code line(s) containing the bug
- (5) Modest input outside of source code under test
 - Preferably glean much of the specification from the source code itself (empirical spec!)
 - Specifications are software and can have their own bugs

Formal Verification and Regression Testing: Requirements

- (1) Either automatic translation or no translation required
 - Automatic discarding of irrelevant portions of the code
 - Manual translation provides opportunity for human error
- (2) Correctly handle environment, including memory model
 - The QRCU validation benchmark is an excellent cautionary tale
- (3) Reasonable memory and CPU overhead
 - Bugs must be located in practice as well as in theory
 - Linux-kernel RCU is 15KLoC and release cycles are short
- (4) Map to source code line(s) containing the bug
 - “Something is wrong somewhere” is not a helpful diagnostic: I **know** bugs exist
- (5) Modest input outside of source code under test
 - Preferably glean much of the specification from the source code itself (empirical spec!)
 - Specifications are software and can have their own bugs
- (6) Find relevant bugs
 - Low false-positive rate, weight towards likelihood of occurrence (fixes create bugs!)

Formal Validation Tools Used and Regression Testing

▪ Promela and Spin

- Holzmann: “The Spin Model Checker”
- I have used Promela/Spin in design for more than 20 years, but:
 - Limited problem size, long run times, large memory consumption
 - Does not implement memory models (assumes sequential consistency)
 - Special language, difficult to translate from C

▪ ARMMEM and PPCMEM (2)

- Alglave, Maranget, Pawan, Sarkar, Sewell, Williams, Nardelli: “PPCMEM/ARMMEM: A Tool for Exploring the POWER and ARM Memory Models”
 - Very limited problem size, long run times, large memory consumption
 - Restricted pseudo-assembly language, manual translation required

▪ Herd (2, 3)

- Alglave, Maranget, and Tautschnig: “Herding Cats: Modelling, Simulation, Testing, and Data-mining for Weak Memory”
 - Very limited problem size (but much improved run times and memory consumption)
 - Restricted pseudo-assembly language, manual translation required

C Bounded Model Checker (CBMC): Usage

- C Bounded Model Checker (CBMC) applies long-standing hardware verification techniques to software
- Easy to use: Given recent Debian-derived distributions:


```
sudo apt-get install cbmc
```



```
cbmc filename.c
```
- If no combination of inputs can trigger an assertion or cause an array-out-of-bounds error, it prints:

```
VERIFICATION SUCCESSFUL
```
- And since 2015, CBMC handles concurrency!!!

How Does CBMC Work?



Scorecard For Linux-Kernel C Code (Incomplete)

	Promela	PPCMEM	Herd	CBMC
(1) Automated				
(2) Handle environment	(MM)		(MM)	(MM)
(3) Low overhead				SAT?
(4) Map to source code				
(5) Modest input				
(6) Relevant bugs	???	???	???	???
Paul McKenney's first use	1993	2011	2014	2015

Promela MM: Only SC: Weak memory must be implemented in model

Herd MM: Some PowerPC and ARM corner-case issues

CBMC MM: Only SC and TSO

Note: All four handle concurrency! (Promela has done so for 25 years!!!)

Scorecard For Linux-Kernel C Code

	Promela	PPCMEM	Herd	CBMC	Test
(1) Automated					
(2) Handle environment	(MM)		(MM)	(MM)	
(3) Low overhead				SAT?	
(4) Map to source code					
(5) Modest input					
(6) Relevant bugs	???	???	???	???	
Paul McKenney's first use	1993	2011	2014	2015	1973

So why do anything other than testing?

- Low-probability bugs can require expensive testing regimen
- Large installed base will encounter low-probability bugs
- Safety-critical applications are sensitive to low-probability bugs

Other Possible Approaches

- By-hand formalizations and proofs
 - Stern: Semi-formal proof of URCU (2012 IEEE TPDS)
 - Gotsman: Separation-logic RCU semantics (2013 ESOP)
 - Tasserotti et al.: Formal proof of URCU linked list: (2015 PLDI)
 - Excellent work, but not useful for regression testing
- seL4 tooling: Lacks support for concurrency and RCU idioms
 - Might be applicable to Tiny RCU callback handling
 - Impressive work nevertheless!!!
- Apply Peter O'Hearn's Infer to the Linux kernel
- Nidhugg: Work by Michalis Kokologiannakis and Kostis Sagonas
 - <https://github.com/michalis-/rcu/blob/master/rcupaper.pdf>
 - Appears to be more scalable than CBMC, but some restrictions
 - Nevertheless, Nidhugg finds all my injected bugs

Summary and Challenges

Summary

- RCU's specification is empirical
- RCU's implementation is unlikely to be bug-free, reliable though it might be
- Currently relying on stress testing augmented by mutation analysis, adding formal verification
 - Formal verification currently weak on forward-progress guarantees
 - And has not yet found any RCU bugs that I didn't already know about
 - But RCU validation is difficult, so I am throwing everything I can at it!!!

Challenges

- Find bug in `rcu_preempt_offline_tasks()`
 - Note: No practical impact because this function has been removed
 - <http://paulmck.livejournal.com/37782.html>
- Find bug in `RCU_NO_HZ_FULL_SYSDLE`
 - <http://paulmck.livejournal.com/38016.html>
- Find bug in RCU linked-list use cases
 - <http://paulmck.livejournal.com/39793.html>
- Find lost wakeup bug in the Linux kernel (or maybe qemu)
 - Heavy rcutorture testing with CPU hotplug on two-socket system
 - Detailed repeat-by: <https://lkml.org/lkml/2016/3/28/214>
 - Can you find this before we do? (Sorry, too late!!!)
- Find any other bug in popular open-source software
 - A verification researcher has provoked a SEGV in Linux-kernel RCU

More Challenges (AKA Current Limitations)

- Incorporate Linux-kernel memory model into analysis
 - And/or the ARM and PowerPC memory models
- Detect race conditions leading to deadlocks and hangs
 - CBMC and Nidhugg can detect unconditional deadlocks and hangs
- Analyze bugs involving networking and mass storage
- Use induction techniques to fully analyze indefinite recursion and unbounded looping
 - Spinloops should be easy: Yes, there are halting-problem limitations
- Analyze larger programs: RCU is not exactly huge!!!
 - Automatically decompose large programs and combine results?

To Probe Deeper (RCU)

- <https://queue.acm.org/detail.cfm?id=2488549>
 - “Structured Deferral: Synchronization via Procrastination” (also in July 2013 CACM)
- <http://doi.ieeecomputersociety.org/10.1109/TPDS.2011.159> and <http://www.computer.org/cms/Computer.org/dl/trans/td/2012/02/extras/ttd2012020375s.pdf>
 - “User-Level Implementations of Read-Copy Update”
- <git://ltnng.org/userspace-rcu.git> (User-space RCU git tree)
- <http://people.csail.mit.edu/nickolai/papers/clements-bonsai.pdf>
 - Applying RCU and weighted-balance tree to Linux mmap_sem.
- http://www.usenix.org/event/atc11/tech/final_files/Triplett.pdf
 - RCU-protected resizable hash tables, both in kernel and user space
- http://www.usenix.org/event/hotpar11/tech/final_files/Howard.pdf
 - Combining RCU and software transactional memory
- <http://wiki.cs.pdx.edu/rp/>: Relativistic programming, a generalization of RCU
- <http://lwn.net/Articles/262464/>, <http://lwn.net/Articles/263130/>, <http://lwn.net/Articles/264090/>
 - “What is RCU?” Series
- <http://www.rdrop.com/users/paulmck/RCU/RCUdissertation.2004.07.14e1.pdf>
 - RCU motivation, implementations, usage patterns, performance (micro+sys)
- http://www.livejournal.com/users/james_morris/2153.html
 - System-level performance for SELinux workload: >500x improvement
- http://www.rdrop.com/users/paulmck/RCU/hart_ipdps06.pdf
 - Comparison of RCU and NBS (later appeared in JPDC)
- <http://doi.acm.org/10.1145/1400097.1400099>
 - History of RCU in Linux (Linux changed RCU more than vice versa)
- <http://read.seas.harvard.edu/cs261/2011/rcu.html>
 - Harvard University class notes on RCU (Courtesy of Eddie Koher)
- <http://www.rdrop.com/users/paulmck/RCU/> (More RCU information)

Legal Statement

- This work represents the view of the author and does not necessarily represent the view of IBM.
- IBM and IBM (logo) are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries.
- Linux is a registered trademark of Linus Torvalds.
- Other company, product, and service names may be trademarks or service marks of others.

Questions?