Stateless Model Checking of the Linux Kernel's Hierarchical Read-Copy-Update (Tree RCU)

Michalis Kokologiannakis
National Technical University of Athens, Greece

Konstantinos Sagonas
Uppsala University, Sweden
National Technical University of Athens, Greece

ABSTRACT

Read-Copy-Update (RCU) is a synchronization mechanism used heavily in key components of the Linux kernel, such as the virtual filesystem (VFS), to achieve scalability by exploiting RCU's ability to allow concurrent reads and updates. RCU's design is non-trivial, requires significant effort to fully understand it, let alone become convinced that its implementation is faithful to its specification and provides its claimed properties. The fact that as time goes by Linux kernels are becoming increasingly more complex and are employed in machines with more and more cores and weak memory does not make the situation any easier.

This paper presents an approach to systematically test the code of the main flavor of RCU used in the Linux kernel (Tree RCU) for concurrency errors, both under sequential consistency and weak memory. Our modeling allows Nidhugg, a stateless model checking tool, to reproduce, within seconds, safety and liveness bugs that have been reported for RCU. More importantly, we were able to verify the Grace-Period guarantee, the basic guarantee that RCU offers, on several Linux kernel versions (non-preemptible builds). Our approach is effective, both in dealing with the increased complexity of recent Linux kernels and in terms of time that the process requires. We have good reasons to believe that our effort constitutes a big step towards making tools such as Nidhugg part of the standard testing infrastructure of the Linux kernel.

CCS CONCEPTS

• **Theory of computation → Verification by model checking**; Logic and verification; • **Software and its engineering → Software testing and debugging**; *Formal software verification*;

KEYWORDS

Software model checking, Read-Copy-Update; Linux kernel; Nidhugg

ACM Reference format:
Michalis Kokologiannakis and Konstantinos Sagonas. 2017. Stateless Model Checking of the Linux Kernel's Hierarchical Read-Copy-Update (Tree RCU). In *Proceedings of International SPIN Symposium on Model Checking of Software, Santa Barbara, CA, USA, July 2017 (SPIN'17),* 10 pages.
https://doi.org/10.1145/3092282.3092287

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SPIN'17, July 2017, Santa Barbara, CA, USA
© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.
ACM ISBN 978-1-4503-5077-8/17/07…$15.00
https://doi.org/10.1145/3092282.3092287

1 INTRODUCTION

The Linux kernel is used in a surprisingly large number of devices: from PCs and servers, to routers and smart TVs. For example, in 2015 more than one billion smart phones used a modified version of the Linux kernel [5], and in 2016 almost all modern supercomputers used Linux as well [22]. It is self-evident that the correct and reliable operation of the Linux kernel is of great importance, which renders thorough testing and verification of its components a necessity.

Naturally, this process needs to span all of the kernel's components and subsystems. One particular subsystem with a non-trivial implementation is the Read-Copy-Update (RCU) mechanism [18, 19]. RCU is a synchronization mechanism that provides excellent scalability by enabling concurrent reads and updates.

However, RCU's implementation is quite involved, making its precise modeling arduous. Moreover, the lockless design of its fast-paths, and the fact that it needs to operate in heavily concurrent environments make its modeling and verification process challenging. Extra difficulties for this process stem from the relatively short release cycle of the Linux kernel (there is a new release every approximately two months), the number of changes that are involved in each release, and the increasing complexity of the kernel's code. Still, the fact that concurrency bugs manage to survive (maybe only under particular configurations, architectures and memory models) even after heavy *stress testing* underlines the need for employing software model checking techniques that are able to operate on as big a percentage of the *actual code* of the Linux kernel as possible.

This paper reports on the use of stateless model checking (aka systematic concurrency testing) for testing the core of Tree RCU, the main RCU flavor used in the Linux kernel. In particular, using Nidhugg [2], we were able to verify that its implementation preserves the *Grace-Period (GP) guarantee*, the basic guarantee that RCU offers. Our effort concentrated on a non-preemptible kernel environment, where we also investigated the effects that weak memory models (TSO in particular) may have in RCU's operation. We used the source code from five different versions of the Linux kernel, and verified that the GP guarantee holds in all of them.

In order to strengthen our verification claim, we injected bugs similar to ones that have been reported throughout the development of RCU and, in all cases, our tool was able to come up with scenarios in which they occur. In particular, we were able to demonstrate that a submitted patch, intended to impose a locking design, in reality fixed a much more serious bug, a fact that was previously unknown. We report on this issue and present the exact conditions under which this bug occurs. Finally, as we will show, our technique handles real code employed in today's production systems in a very efficient and scalable way. In fact, Nidhugg copes extremely well with the increasing complexity the newer kernel versions induce.

Overview. After an introduction to RCU and to stateless model checking, we describe details of the implementation of Tree RCU in Section 4. Section 5 presents our modeling of the kernel's environment. In Section 6 we investigate some previously reported failures for an older kernel and present their cause, and in Section 7 we report on the verification of the GP guarantee of RCU. The paper ends with related work, lessons learned, and concluding remarks.

2 READ-COPY-UPDATE (RCU)

Read-Copy-Update is a synchronization mechanism invented by McKenney and Slingwine [18, 19] that is a part of the Linux kernel since 2002. The key feature of RCU is the good scalability it provides by allowing concurrent reads and updates. While this may seem counter-intuitive or impossible at first, RCU allows this in a very simple yet extremely efficient way: by maintaining multiple data versions. RCU is carefully orchestrated in a way that not only ensures that reads are coherent and no data will be deleted until it is certain that no one holds references to them, but also uses efficient and scalable mechanisms which make read paths extremely fast. Most notably, in non-preemptible kernels, RCU imposes zero overhead to readers.

The basic idea behind RCU is to split updates in two phases: the *removal phase* and the *reclamation phase*. During the removal phase, an updater *removes* references to data either by destroying them (i.e., setting them to `NULL`), or by replacing them with references to newer versions of these data. This phase can run concurrently with reads due to the fact that modern microprocessors guarantee that a reader will see either the old or the new reference to an object, and not a weird mash-up of these two or a partially updated reference. During the reclamation phase, the updater frees the items removed in the removal phase, i.e., these items are *reclaimed*. Of course, since RCU allows concurrent reads and updates, the reclamation phase must begin after the removal phase and, more specifically, when it is certain that there are no readers accessing or holding references to the data being reclaimed.

The typical update procedure using RCU looks as follows [18].

(1) Ensure that all readers accessing RCU-protected data structures carry out their references from within an RCU read-side critical section.
(2) Remove pointers to a data structure, so that subsequent readers cannot gain a reference to it (removal phase).
(3) Wait until all pre-existing readers complete their RCU read-side critical section, so that there no one holding a reference to the item being removed.
(4) At this point, there cannot be any readers still holding references to the data structure, which may now be safely freed.

Note that steps 2 and 4 (the reclamation phase) in this procedure are not necessarily performed by the same thread.

Waiting for pre-existing readers can be achieved either by blocking (via `synchronize_rcu()`), or by registering a callback that will be invoked after all pre-existing readers have completed their RCU read-side critical sections (via `call_rcu()`).

In order to formalize some of the aspects presented above, we provide some definitions.

Definition 2.1. Any statement that is not within an RCU read-side critical section is said to be in a *quiescent state*.

```
Initially: int x = 0, y = 0, r_x = 0, r_y = 0;

void reader(void)                void updater(void)
{                                {
  rcu_read_lock();                 WRITE_ONCE(x, 1);
  r_x = READ_ONCE(x);              synchronize_rcu();
  r_y = READ_ONCE(y);              WRITE_ONCE(y, 1);
  rcu_read_unlock();             }
}
```

Figure 1: RCU's Grace-Period guarantee litmus test.

Statements in quiescent states are not permitted to hold references to RCU-protected data structures. Note that different RCU flavors have different sets of quiescent states.

Definition 2.2. Any time period during which each CPU resides at least once in a quiescent state is called a *grace period*.

Consequently, if an RCU read-side critical section started before the beginning of a specified grace period *GP*, it would have to complete before the end of *GP*. This means that the reclamation phase has to wait for *at least* one grace period to elapse before it begins. Once a grace period has elapsed, there can no longer be any readers holding references to the old version of a newly updated data structure (since each CPU has passed through a quiescent state) and the reclamation phase can safely begin.

RCU Specifications. Let us now present some requirements that every RCU implementation must fulfill. We do not attempt to present a formal or a complete specification for RCU here.[1] Instead, we only present the basic guarantees of RCU.

Grace-Period Guarantee. The fact that in RCU updaters wait for all pre-existing readers to complete their read-side critical sections, constitutes the only interaction between the readers and the updaters. The Grace-Period guarantee is what allows updaters to wait for all pre-existing RCU read-side critical sections to complete. Such critical sections start with the macro `rcu_read_lock()` and end with `rcu_read_unlock()`. What this guarantee means is that the RCU implementation must ensure that any read-side critical sections in progress at the start of a given grace period *GP* will have completely finished (including memory operations, etc.) before that *GP* ends. This very fact allows RCU verification to be focused; every correct implementation has to adhere to the following rule:

> *If any statement in a given RCU read-side critical section CS precedes a grace period GP, then all statements (including memory operations) in CS must complete before GP ends.*

Memory operations are included here in order to prevent the compiler or the CPU from undoing work done by RCU.

In order to see what this guarantee really implies, consider the code fragment in Fig. 1. In this code, since `synchronize_rcu()` has to wait for all pre-existing readers to complete their RCU read-side critical sections, the outcome:

$$\texttt{r_x == 0 \&\& r_y == 1} \tag{1}$$

should be impossible. This is what the Grace-Period guarantee is all about. It is the most important guarantee that RCU provides; in effect, it constitutes the core of RCU.

[1] Requirements for RCU appear at: https://www.kernel.org/doc/Documentation/RCU/.

Publish-Subscribe Guarantee. This guarantee is used in order to coordinate read-side accesses to data structures. The Publish-Subscribe mechanism is used in order for data to be inserted into data structures (e.g., lists), without disrupting concurrent readers. Since updaters run concurrently with readers, this mechanism ensures that readers will not see uninitialized data, and that updaters will have completed all initialization operations before publishing a data structure. For this, RCU offers two primitives: (1) The `rcu_assign_pointer()` primitive, which has similar semantics to C11's `memory_order_release` operation. In effect, it is similar to an assignment but also provides additional ordering guarantees. (2) The `rcu_dereference()` primitive, which can be considered as a subscription to a value of a specified pointer and guarantees that subsequent dereference operations will see any initialization that took place before the `rcu_assign_pointer()` (publish) operation. The `rcu_dereference()` primitive has semantics similar to C11's `memory_order_consume` load, and uses both volatile casts and memory barriers in order to provide the aforementioned guarantee.

3 STATELESS MODEL CHECKING

Stateless model checking [11], also known as systematic concurrency testing, is a technique with low memory requirements that is applicable to programs with executions of finite length. Stateless model checking tools explore the state space of a program without explicitly storing global states. The technique has been successfully implemented in tools such as VeriSoft [12] and CHESS [20].

Stateless model checking tools try to combat the problem of combinatorial explosion in the number of interleavings that need to be examined in order to maintain full coverage of all program behaviors by using *partial order reduction* techniques. Partial order reduction is based on the observation that two interleavings can be considered equivalent if one can be obtained from the other by swapping adjacent, independent execution steps. *Dynamic Partial Order Reduction* (DPOR) techniques capture dependencies between operations of concurrent threads while the program is running [10]. The exploration begins with an arbitrary interleaving whose steps are then used to identify dependent operations and program points where alternative interleavings need to be explored in order to capture all program behaviors.

Stateless model checking and DPOR techniques have been extended to handle memory model non-determinism in addition to scheduling non-determinism. Nidhugg [2], for example, is a stateless model checker for C/C++ programs that use pthreads, which incorporates extensions for finding bugs caused by weak memory models such as TSO, PSO and, partially, POWER and ARM. Nidhugg works on the level of LLVM intermediate representation and employs an effective dynamic partial order algorithm called source-DPOR [1]. In our work we used Nidhugg for all our tests.

In stateless model checking all tests need to be *data-deterministic* in the sense that, in a given state, a given execution step must always lead the system to the same new state. This means that the test case cannot depend on some unknown input or on timing properties (e.g., take some action depending on the value of the clock). In addition, all test cases need to be *finite* in the sense that there must be a bound $n \in \mathbb{N}$ such that all executions of the program terminate within n execution steps.

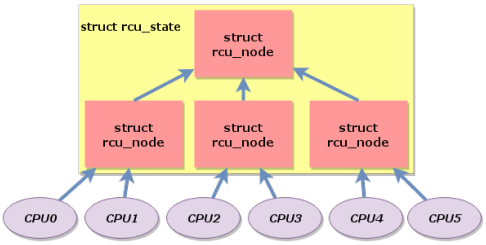

Figure 2: Tree RCU node hierarchy (adapted from [16]).

4 TREE RCU IMPLEMENTATION

The Linux kernel offers many different RCU implementations, each one serving a different purpose. The first Linux-kernel RCU implementation was Classic RCU. A problem with Classic RCU was lock contention due to the presence of one global lock that had to be acquired from each CPU wishing to report a quiescent state to RCU. In addition, Classic RCU had to wake up every CPU (even idle ones) at least once per grace period, thus increasing power consumption.

Tree RCU offers a solution to both these problems since it reduces lock contention and avoids awakening dyntick-idle [21] CPUs. Tree RCU scales to thousands of CPUs easily, while Classic RCU could scale only to several hundred.

Below we present a high-level explanation of Tree RCU along with some implementation details, a brief overview of its data structures, and some use cases that are helpful in understanding how RCU's fundamental mechanisms are actually implemented.

4.1 High-Level Explanation

In Classic RCU each CPU had to clear its bit in a field of a global data structure after it passed through a quiescent state. Since CPUs operated concurrently on this data structure, a spinlock was used to protect the mask, which could suffer from extreme contention.

Tree RCU addresses this issue by creating a heap-like node hierarchy. The key here is that CPUs will not try to acquire the same node's lock when trying to report a quiescent state to RCU; in contrast, CPUs are split into groups and each group will contend for a different node's lock. Each CPU has to clear its bit in the corresponding node's mask once per grace period. The last CPU to check in (i.e., to report a quiescent state to RCU) for each group, will try to acquire the lock of the node's parent, until the root node's mask is cleared. This is when a grace period can end. A simple node hierarchy for a 6-CPU system is presented in Fig. 2. As can be seen, CPUs 0 and 1 will acquire the lower-left node's lock, CPUs 2 and 3 will acquire the lower-middle node's lock, and CPUs 4 and 5 will acquire the lower-right node's lock. The last CPU reporting a quiescent state for each of the lower nodes will try to acquire the root node's lock, and this procedure happens once per grace period.

The node hierarchy created by Tree RCU is tunable, and is controlled, among others, by two `Kconfig` options, namely:

`CONFIG_RCU_FANOUT_LEAF`: Controls the maximum number of CPUs contending for a leaf-node's lock. Default value is 16.

`CONFIG_RCU_FANOUT`: Controls the maximum number of CPUs contending for an inner-node's lock. Default value is 32 for 32-bit systems and 64 for 64-bit systems.

More information can be found at the `init/Kconfig` file.

4.2 Data Structures

Let us now describe three major data structures of Tree RCU's implementation: `rcu_data`, `rcu_node`, and `rcu_state`. Suppose that a CPU registers a callback that will eventually be invoked. Tree RCU needs to store some information regarding this callback. For this, the implementation maintains some data organized in the per-CPU `rcu_data` structure, which includes, among others: (i) the last completed grace period number this CPU has seen; used for grace-period ending detection (`completed`), (ii) the highest grace period number this CPU is aware of having started (`gpnum`), (iii) a `bool` variable indicating whether this CPU has passed through a quiescent state for this grace period, (iv) a pointer to this CPU's leaf of hierarchy, and (v) the mask that will be applied to the leaf's mask (`grpmask`). Thus, when a CPU registers a callback, it stores it in the respective per-CPU data structure.

Now, when a CPU passes through a quiescent state, it has to report it to RCU by clearing its bit in the respective leaf node. The node hierarchy consists of `rcu_node` structures which include, among others: (i) a lock protecting the respective node, (ii) the current grace period number for this node, (iii) the last completed grace period number for this node, (iv) a bit-mask indicating CPUs or groups that need to check in in order for this grace period to proceed (`qsmask`), (v) a pointer to the node's parent, and (vi) the mask that will be applied to parent node's mask (`grpmask`).

Lastly, the RCU global state, as well as the node hierarchy are included in an `rcu_state` structure. The node hierarchy is represented in heap form in a linear array, which is allocated statically at compile time based on the values of `NR_CPUS` and the `Kconfig` options. Note that small systems have a hierarchy consisting of a single `rcu_node`. This structure contains, among others: (i) the node hierarchy, (ii) a pointer to the per-CPU `rcu_data` variable, (iii) the current grace-period number, and (iv) the number of last completed grace period. There are several values that are propagated through these different structures, e.g., the grace period number. However, this was not always the case, and it was the discovery of bugs that often led to changes in the source code.

Finally, we have already mentioned that Classic RCU had a sub-optimal dynticks interface, and that one of the main reasons for the creation of Tree RCU was to leave sleeping CPUs lie, in order to conserve energy. Tree RCU avoids awakening low-power-state dynticks-idle CPUs using a per-CPU data structure called `rcu_dynticks`. This structure contains, among others: (i) a counter tracking the irq/process nesting level, and (ii) a counter containing an even value for dynticks-idle mode, else containing an odd value. These counters enable Tree RCU to wait only for CPUs that are not sleeping, and to let sleeping CPUs lie. How this is achieved is described below.

4.3 Use Cases

The common usage of RCU involves registering a callback, waiting for all pre-existing readers to complete, and finally, invoking the callback. During all these, special care is taken to accommodate sleeping CPUs, offline CPUs and CPU hotplugs [7], CPUs in user-land, and CPUs that fail to report a quiescent state to RCU within a reasonable amount of time.

Registering a Callback. A CPU registers a callback by invoking `call_rcu()`. This function queues an RCU callback that will be invoked after a specified grace period. The callback is placed in the callback list of the respective CPU's `rcu_data` structure. This list is partitioned in four segments:

(1) The first segment contains entries that are ready to be invoked (`DONE` segment).
(2) The second segment contains entries that are waiting for the current grace period (`WAIT` segment).
(3) The third segment contains entries that are known to have arrived before the current grace period ended (`NEXT_READY` segment).
(4) The fourth segment contains entries that *might* have arrived after the current grace period ended (`NEXT` segment).

When a new callback is added to the list, it is inserted at the end of the fourth segment.

In older kernels (e.g., v2.6.x), `call_rcu()` could start a new grace period directly, but this is no longer the case. In newer kernels, the only way a grace period can start directly by `call_rcu()` is if there are too many callbacks queued and no grace period in progress. Otherwise, a grace period will start from *softirq* context.

Every softirq is associated with a function that will be invoked when this type of softirqs is executed. For Tree RCU, this function is called `rcu_process_callbacks()`. So, when an RCU softirq is raised, this function will eventually be invoked (either at the exit from an interrupt handler or from a `ksoftirq/n` kthread), and will start a grace period if there is need for one (e.g., if there is no grace period in progress and this CPU has newly registered callbacks, or there are callbacks that require an additional grace period). RCU softirqs are raised from `rcu_check_callbacks()` which is invoked from scheduling-clock interrupts. If there is RCU-related work (e.g., if this CPU needs a new grace period), `rcu_check_callbacks()` raises a softirq.

The `synchronize_rcu()` function, which is implemented on top of `call_rcu()` in Tree RCU, registers a callback that will awake the caller after a grace period has elapsed. The caller waits on a completion variable, and is consequently put on a wait queue.

Starting a Grace Period. The `rcu_start_gp()` function is responsible for starting a new grace period; it is normally invoked from softirq context and an `rcu_process_callbacks()` call. However, in newer kernels, `rcu_start_gp()` neither directly starts a new grace period nor initializes the necessary data structures. It rather advances the CPU's callbacks (i.e., properly re-arranges the segments), and then sets a flag at the `rcu_state` structure to indicate that a CPU requires a new grace period. The grace-period kthread is the one that will initialize the node hierarchy and the `rcu_state` structure, and by extension start the new grace period.

The RCU grace-period kthread excludes concurrent CPU-hotplug operations and then sets the quiescent-state-needed bits in all the `rcu_node` structures in the hierarchy corresponding to online CPUs. It also copies the grace period number in all the `rcu_node` structures. Concurrent CPU accesses will check only the leaves of the hierarchy, and other CPUs may or may not see their respective node initialized. But each CPU has to enter the RCU core in order to acknowledge that a grace period has started and initialize its `rcu_data` structure.

This means that each CPU (except for the one on which the grace-period kthread runs) needs to enter softirq context in order to see the new grace period beginning (via `rcu_process_callbacks()`).

The grace-period kthread resolved many races present in older kernels, where when a CPU required a new grace period, it tried to directly initialize the node hierarchy, something that could potentially lead to bugs; see Section 6.

Passing Through a Quiescent State. Quiescent states for Tree RCU (RCU-sched) include: (i) context switch, (ii) idle mode (idle loop or dynticks-idle), and (iii) user-mode execution. When a CPU passes through a quiescent state, it updates its `rcu_data` structure by invoking `rcu_sched_qs()`. This function is invoked from scheduling-related functions, from `rcu_check_callbacks()`, and from the `ksoftirq/n` kthreads. However, the fact that a CPU has passed through a quiescent state does not mean that RCU knows about it. Besides, this fact has been recorded in the respective per-CPU `rcu_data` structure, and not in the node hierarchy. So, a CPU has to report to RCU that it has passed through a quiescent state, and this will happen —again— from softirq context, via the `rcu_process_callbacks()` function; see below.

Reporting a Quiescent State to RCU. After a CPU has passed through a quiescent state, it has to report it to RCU via the function `rcu_process_callbacks()`, whose duties include:

- Awakening the RCU grace-period kthread (by invoking the `rcu_start_gp()` function), in order to initialize and start a new grace period, if there is need for one.
- Acknowledging that a new grace period has started/ended. Every CPU except for the one on which the RCU grace-period kthread runs has to enter the RCU core and see that a new grace period has started/ended. This is done by invoking the function `rcu_check_quiescent_state()`, which in turn invokes `note_gp_changes()`. The latter advances this CPU's callbacks and records to the respective `rcu_data` structure all the necessary information regarding the grace-period beginning/end.
- Reporting that the current CPU has passed through a quiescent state (via `rcu_report_qs_rdp()`, which is invoked from `rcu_check_quiescent_state()`). If the current CPU is the last one to report a quiescent state, the RCU grace-period kthread is awakened once again in order to clean up after the old grace period and propagate the new `->completed` value to the `rcu_node` structures of the hierarchy.
- Invoking any callbacks whose grace period has ended.

As can be seen, the RCU grace-period kthread is used heavily to coordinate grace-period beginnings and ends. Apart from this, the locks of the nodes in the hierarchy are used to prevent concurrent accesses which might lead to problems.

Entering/Exiting Dynticks-Idle Mode. When a CPU enters dynticks-idle mode `rcu_idle_enter()` is invoked. This function decrements a per-CPU nesting variable (`dynticks_nesting`) and increments a per-CPU counter (`dynticks`), both of which are located in the per-CPU `rcu_dynticks` structure. The `dynticks` counter must have an even value when entering dynticks-idle mode. When a CPU exits dynticks-idle mode `rcu_idle_exit()` is invoked, which increments `dynticks_nesting` and the `dynticks` counter (which must now have an odd value).

However, dynticks-idle mode is a quiescent state for Tree RCU. So, the reason these two variables are needed is the fact that they can be sampled by other CPUs so it can be safely determined if a CPU is (or has been, at some point) in a quiescent state for this grace period. The sampling process is performed when a CPU has not reported a quiescent state for a long time and the grace period needs to end (quiescent state forcing).

Interrupts and Dynticks-Idle Mode. When a CPU enters an interrupt handler, `rcu_irq_enter()` is invoked from `irq_enter()`. This function increments the value of `dynticks_nesting` and if the prior value was zero (i.e., the CPU was in dynticks-idle mode), also increments the `dynticks` counter. When a CPU exits an interrupt handler, `rcu_irq_exit()` decrements `dynticks_nesting` and if the new value is zero (i.e., the CPU is entering dynticks-idle mode), also increments the `dynticks` counter. It is self-evident that entering an interrupt handler from dynticks-idle mode means exiting the dynticks-idle mode. Conversely, exiting an interrupt handler might mean entrance into dynticks-idle mode.

Forcing Quiescent States. If not all CPUs have reported a quiescent state and several `jiffies` have passed, then the grace-period kthread is awakened and will try to force quiescent states on CPUs that have yet to report one. More specifically, the grace-period kthread will invoke `rcu_gp_fqs()`, which works in two phases: in the first phase snapshots of the `dynticks` counters of all CPUs are collected, in order to credit them with implicit quiescent states. In the second phase, CPUs that have yet to report a quiescent state are scanned again, in order to determine if they have passed through a quiescent state from the moment their snapshots were collected. If there are still CPUs that have not checked in, they are forced into the scheduler in order for them to report a quiescent state to RCU.

5 KERNEL ENVIRONMENT MODELING

Let us now present the way we scaffolded a non-preemptible Linux-kernel SMP environment. For this, we had to disable some timing-based warnings, and stub out some primitives used in functions that were not included in our tests (e.g., RCU-expediting related primitives). However, we note that the only changes we made in the source code of Tree RCU involved the replacement of per-CPU variables with arrays; the rest of the source code remains untouched.

5.1 CPU, Interrupts and Scheduling

CPU. Since we emulate an SMP system, we need some kind of mutual exclusion between threads running on the same CPU, for each CPU of the system. Thus, we provide an array of locks (namely `cpu_lock`), with each array entry corresponding to a CPU. When one of these locks is held, the corresponding thread is running on the respective CPU.

We assume that all CPUs are online, that there are no CPU hotplugs, and that `CONFIG_NO_HZ_FULL=n`. All CPUs are initially idle, and when a thread wishes to acquire/relase a CPU, it acquires/releases the CPU's lock and exits/enters idle mode (if necessary).

We also needed to emulate per-CPU variables. In the kernel, these variables are created using special compiler/linker directives, along with some preprocessor directives. However, since these variables require significant runtime support, we used arrays to

emulate them, with each array entry representing the respective CPU's copy of a per-CPU variable.

Since a thread needs to have knowledge regarding the CPU it runs on, we implemented two macros (`set_cpu()` and `get_cpu()`), which manipulate a thread-local variable indicating the CPU on which a thread runs. The CPU on which a thread runs has to be manually set, via `set_cpu()`. The total number of CPUs can be manipulated by setting the `-DCONFIG_NR_CPUS` preprocessor option.

Interrupts and Softirqs. In order to emulate interrupts and softirqs we used an array of locks (`irq_lock`), with each lock corresponding to a CPU. An entry's lock must be held across an interrupt handler by the thread servicing the interrupt on the respective CPU. Of course, the CPU's lock must be already held. In a similar manner, when a thread disables interrupts on a CPU, the same lock has to be acquired. Since we are dealing with non-preemptible kernels, this lock is not contended.

We also needed to model scheduling-clock interrupts (on which RCU relies heavily) and the function `rcu_check_callbacks()`. But, as mentioned, stateless model checking is performed on deterministic programs, meaning that timing-based actions cannot be included in our tests. However, the exact time an interrupt occurs is not so important; what interests us is the implications a timing interrupt might have at a certain point of a program's execution given a concurrency context. Consequently, our version of the interrupt handler invokes `rcu_check_callbacks()` and then, if an RCU softirq is raised, the `rcu_process_callbacks()` function. Of course, we could have just called the `rcu_process_callbacks()` function, but in the Linux kernel this function is not invoked unconditionally, and we wanted our model to be as precise as possible.

Scheduling. The `cond_resched()` function is modeled by having the running thread drop the CPU's lock and then (possibly) reacquire it, but with `rcu_note_context_switch()` being invoked before releasing the lock of the incoming CPU. A better way to model this function would probably have been to drop the current CPU's lock, acquire the lock of a *random* CPU, and then check that no assertion is violated for *every* possible CPU choice. However, doing this requires support for data non-determinism, at least in the form of some suitable built-in (like e.g., `VS_toss(n)`). However, Nidhugg currently does not provide such support. This also explains why so far we have not modeled a preemptible kernel's environment. Our tests aim to be CPU-specific and not thread-specific, in the sense that we care about the actions of each CPU (e.g., entering/exiting a critical section or servicing a softirq) and not about the specific threads that perform these actions.

5.2 Kernel Definitions

Many kernel definitions were copied directly from the Linux kernel. These include data types like `u8`, `u16`, etc., compiler directives like `offsetof()`, macros like `ACCESS_ONCE()`, list data types and functions, memory barriers, as well as various other kernel primitives. On the other hand, many primitives had to be replaced or stubbed; we supplied empty files for `#include` directives, and provided some other definitions based on some specific `Kconfig` options. These include CPU-related definitions (e.g., `NR_CPUS`), RCU-related definitions that are normally configured at compile time (e.g., `CONFIG_RCU_BOOST`), special compiler directives, tracing functions, etc. The `BUG_ON()` macro and its relatives (e.g., `WARN_ON()`) have been replaced by `assert()` statements. Note that we only stubbed primitives irrelevant to our tests (e.g., some primitives related to grace-period expediting), and provided our own definitions for some other primitives in order for them to work with our modeling of the CPUs and interrupts.

All of the definitions we used reside in separate files; these can be copied and reused across multiple kernel versions.

5.3 Synchronization Mechanisms

The emulation of the Linux kernel's synchronization mechanisms used in Tree RCU's implementation is as follows:

Atomic Operations. While we copied the `atomic_t` data type directly from the Linux kernel, this is not the case for atomic operations like `atomic_read()`, `atomic_set()`, etc., since their implementation is architecture dependent. In order to emulate those, we used some `GCC` language extension supported by `clang` [15], the compiler that produces the `LLVM IR` code that Nidhugg analyzes.

Spinlocks and Mutexes. We used `pthread_mutexes` for the emulation of kernel spinlocks and mutexes.

Completions. In order to emulate completion variables, we copied the data type definition directly from the Linux kernel, but we had to model wait queues. Since a thread waiting on a completion is put on a wait queue until some condition is satisfied, we used spin loops in order to emulate wait queues. Nidhugg automatically transforms all spin loops to `__VERIFIER_assume()` statements where, if the condition does not hold, the execution blocks indefinitely. Before waiting on a spin loop, the thread drops the corresponding CPU's lock; it will try to re-acquire it after the condition has been satisfied. Since this is a quiescent state for RCU, the function `rcu_note_context_switch()` (and possibly also the `do_IRQ()` function, in order to report a quiescent state to RCU) could have been invoked before the thread released the CPU's lock. However, if the thread waiting on the completion variable is not the only thread running on the specific CPU, this is unnecessary; these functions can be called from other threads running on the same CPU as well.

6 INVESTIGATING AN OLDER KERNEL BUG

In Section 4.3 we mentioned that the grace-period kthread cleans up after grace-period ends. However, in older kernel versions, the RCU grace-period kthread did not exist; when a CPU entered the RCU core or invoked `call_rcu()`, it checked for grace-period ends by directly comparing the number of the last completed grace period in the `rcu_state` structure with the number of the last completed grace period in the respective `rcu_data` structure. In newer kernels, the `note_gp_changes()` function compares the number of the last completed grace period in the respective `rcu_node` structure with the number of the last completed grace period in the current `rcu_data` structure, *while holding the node's lock*, that way excluding concurrent operations on this node.

In kernel v2.6.32, commit `d09b62dfa336` fixed a synchronization issue exposed by unsynchronized accesses to the `->completed` counter in the `rcu_state` structure [23, 24], which caused the advancement of callbacks whose grace period had not yet expired.

```
completed_snap = ACCESS_ONCE(rsp->completed); /* outside of lock */

/* Did another grace period end? */
if (rdp->completed != completed_snap) {
  /* Advance callbacks.  No harm if list empty. */
  rdp->nxttail[RCU_DONE_TAIL] = rdp->nxttail[RCU_WAIT_TAIL];
  rdp->nxttail[RCU_WAIT_TAIL] = rdp->nxttail[RCU_NEXT_READY_TAIL];
  rdp->nxttail[RCU_NEXT_READY_TAIL] = rdp->nxttail[RCU_NEXT_TAIL];

  /* Remember that we saw this grace-period completion. */
  rdp->completed = completed_snap;
}
```

Figure 3: Snippet of the `rcu_process_gp_end()` function.

Below we will create a test case that shows such a situation, but this test case will also demonstrate that the problem is actually deeper: these unsynchronized accesses also lead to too-short grace periods.

We started by looking at the `rcu_process_gp_end()` function, since the issue was related to it. Figure 3 shows a relevant portion of its code. As can be seen, the access to the `->completed` counter is completely unprotected. So, we injected a `BUG_ON()` statement in the `if`-body to determine if it was possible for a thread to pick up the `->completed` value and then use the `completed_snap` while the `->completed` variable had changed. The answer was affirmative. Our next step was to determine if this could potentially lead to a CPU starting a new grace period without having noticed that the last grace period has ended. Again, an injection of a `BUG_ON()` statement, comparing the current grace period's number with the number of the grace period whose completion was noticed by the CPU, showed that this was possible. With these clues, we constructed a simple test which proved that these unsynchronized accesses can lead to too-short grace periods. The test has a reader seeing changes happening before the beginning of a grace period *and* after the end of the same grace period within a single RCU read-side critical section which, of course, is a violation of the GP guarantee[2].

Let us end this section with some notes regarding this bug:

- The bug does not rely on interactions with the node hierarchy; it existed in both single-node and multi-level hierarchies. (A slightly different test case with the respective `Kconfig` options set appropriately would be required for multi-level hierarchies.)
- Nidhugg reports that this bug is not present in kernel v3.0, which means that it was indeed fixed. In v3.0, `rcu_start_gp()` calls `__rcu_process_gp_end()`, thus guaranteeing that a CPU will see a grace-period ending before a grace-period beginning, something that does not happen in v2.6.32.1. However, the bug was present in previous versions as well, e.g., v2.6.31.1.
- Only two CPUs are required to provoke the bug, and only one of them has to invoke `call_rcu()`.
- Only one grace period is required to provoke the bug, meaning that it does not rely on CPUs being unaware of grace period ends and beginnings (e.g., when a CPU is in dynticks-idle mode). However, this bug *does* require some actions to occur during and after the ending of a grace period, meaning that a simple grace-period guarantee test would not have exposed this bug.
- `force_quiescent_state()` is not required to provoke the bug, although frequent calls to this function would expose it more easily in real-life scenarios.
- This bug is not caused by weak memory ordering; the test fails under sequential consistency.

[2]More information about this test case, the sequence of events that exposes this bug, and all code for Sections 6 and 7 can be found at https://www.github.com/michalis-/rcu.

- Nidhugg produced the violating sequence of events in only 0.56*s* (compilation and Nidhugg transformation time included), and used 30.85MB of memory in total.

7 STATELESS MODEL CHECKING TREE RCU

In this section, we will verify the Grace-Period guarantee of Tree RCU for a non-preemptible Linux kernel environment, using the model we created in Section 5. We have applied this model to five different Linux kernels (v3.0, v3.19, v4.3, v4.7, and v4.9.6), and were able to verify that the actual RCU code satisfies the GP guarantee under both SC and TSO, using a litmus test similar to that of Fig. 1.

7.1 Test Configuration

Let us first briefly discuss our modeling of the Linux kernel. All our experiments focused on the RCU-sched flavor of Tree RCU.

First of all, we model a system with two cores, represented by two mutexes, respectively. We also have three basic threads: the updater, the reader and the RCU grace-period kthread. The RCU-bh grace period kthread is disabled in order to reduce the state space, but it can be re-enabled by setting the `-DENABLE_RCU_BH` preprocessor option. We can assume that the updater and the RCU grace-period kthread run on the same CPU (e.g., CPU0), and that the reader runs on the other CPU (e.g., CPU1). For RCU initialization, the `rcu_init()` function is called. Since there are only two CPUs in our modeling, a single-node hierarchy is created. All CPUs start out idle (`rcu_idle_enter()` is called for each CPU), and `rcu_spawn_gp_kthread()` is called in order to spawn the RCU grace-period kthread.

Of course, interrupt context needs to be emulated as well. In general, even though we do not care about the exact timing of interrupts, it is the occurrence of an interrupt within a specific context that causes a grace period to advance. Thus, we have sprinkled calls to `do_IRQ()` in various points of the test code, which enable the advancement of a grace period. This may not always be the case (i.e., a grace period may not end for some explored executions), but in fact we want to enable both of these scenarios.

7.2 Verifying the Grace-Period Guarantee

All experiments have been run on a 64-bit desktop with an Intel Core i7-3770 processor @ 3.40GHz and 16GB of RAM running Arch Linux 4.10.13-1-ARCH. After running the test with an unroll value in order for the test to be finite, Nidhugg reports that the test is successful for all five kernel versions. Moreover, the process is quite fast. As shown on the first row of Table 1, the verification of the GP guarantee under SC requires less than 18 minutes. Another set of runs, verifying this guarantee under the TSO memory model does not require considerably more time. Nidhugg tells us that there is no possible thread or memory model interleaving that violates the GP guarantee in Tree RCU's implementation.

But, can we really trust these results? After all, there might be a bug in our scaffolding of the Linux-kernel's environment, or there might be a bug in Nidhugg itself. In order to increase our confidence, we injected a number of bugs similar to ones that have occurred in real systems in production over the years. These bugs were added both in the test and the RCU source code. More specifically, we injected two kinds of bugs:

-DASSERT_0: An `assert(0)` statement is inserted after `synchronize_rcu()`. Obviously, this results in a test failure. What this assertion does, however, is that it shows that the grace period *can* end, and that there are *some* explored executions in which it does; i.e., it provides liveness guarantees. We will use this injection in conjunction with some of the next bug injections in order to determine whether the grace period can end or not.

-DFORCE_FAILURE_1: This injection forces the reader to pass through and report a quiescent state during its read-side critical section. Of course, this is not permitted and, as expected, results in a failure.

-DFORCE_FAILURE_2: A `return` statement is placed at the beginning of `synchronize_rcu()`. Of course, this results in a test failure since the updater does not wait for pre-existing readers to complete their RCU read-side critical sections, and such critical sections are not permitted to span a grace period.

-DFORCE_FAILURE_3: This injection makes `rcu_gp_init()` clear the node mask (`->qsmask`) variables instead of setting them appropriately. The `rcu_gp_init()` function is invoked from the RCU grace-period kthread at the beginning of each grace period in order to initialize it. Obviously, since the `->qsmask` variables are cleared from the start of the grace period, the grace period can end immediately. In other words, the grace-period kthread does not wait for pre-existing readers to complete. (This can be considered a more complex variant of injection #2.) As expected, this injection results in a test failure.

-DFORCE_FAILURE_4: In this injection the `rcu_gp_fqs()` function is made to clear the `->qsmask` variables instead of waiting for the CPUs to clear their respective bits. Of course, in order for `rcu_gp_fqs()` to clear the `->qsmask` variables, the respective CPUs (in our case, the reader) have to be in dynticks-idle mode (or the CPU must have passed through a quiescent state at some point, since the respective dynticks counters are sampled). Consequently, in our code, CPU0 calls the `rcu_gp_fqs()` function, and CPU1 enters and exits dynticks-idle mode within its RCU read-side critical section, which enables CPU0 to prematurely end the grace period. This can be considered an even more complex variant of injection #2, and results in a test failure, as expected.

-DFORCE_FAILURE_5: This injection makes the function `__note_gp_changes()` clear the bit of the respective node's mask for this CPU (`rnp->qsmask &= ~rdp->grpmask`). This function is called when a CPU enters RCU core in order to record the beginnings and ends of grace periods. However, instead of just recording a grace period beginning, `__note_gp_changes()` is now made to also clear the `->qsmask` bit, which implies that this CPU reported a quiescent state for the new grace period. This results in test failure.

-DFORCE_FAILURE_6: Essentially, what this injection does is delete the `if` statement checking whether a node's mask is zero and calling `rcu_preempt_blocked_readers_cgp()`, in the `rcu_report_qs_rnp()` function. This `if` statement just checks whether the bitmask for this node is cleared in order for a node to acquire its parent's lock. In a real kernel, this should result in too short grace periods, since a signal that will prematurely awake the grace-period kthread is sent, if there are multiple CPUs. In our case, however, it does not lead to too-short grace periods since, in our modeling, `wake_up()` boils down to a no-op – there is no need to wake up someone who is just spinning. However, if we were dealing with a two-level tree, the caller of `rcu_report_qs_rnp()` would move up one level and trigger a `WARN_ON_ONCE()` statement that checks whether the child node's bits are cleared. Hence, this test automatically sets the number of CPUs to `CONFIG_RCU_FANOUT_LEAF + 1` (i.e., to 17, since the default value of `CONFIG_RCU_FANOUT_LEAF` is 16 in these kernels). Also, this test requires the use of a higher unroll value because there are some loops that need to be unrolled at least as many times as the number of CPUs used plus one. So, we used an unroll value of 19 for this case.

-DLIVENESS_CHECK_1: This eliminates the need for a CPU to pass through a quiescent state by setting `rdp->qs_pending` to zero in `__note_gp_changes()`. This function updates the per-CPU `rcu_data` structure and, since `rdp->qs_pending` is set to zero, there is no need for a CPU to report a quiescent state to RCU, which prevents grace periods from completing. When the injection is used in conjunction with `-DASSERT_0`, no execution triggers the `assert(0)` statement after `synchronize_rcu()`.

-DLIVENESS_CHECK_2: A `return` statement is placed at the beginning of the `rcu_sched_qs()` function. In effect, this means that CPUs cannot record their passing through a quiescent state in the respective `rcu_data` structures, something that also prevents grace periods from completing. Used in conjunction with `-DASSERT_0` this bug injection also results in no executions triggering the assertion, thus signifying a liveness violation.

-DLIVENESS_CHECK_3: A `return` statement is placed at the beginning of `rcu_report_qs_rnp()`. This means that CPUs cannot report their passing through a quiescent state to RCU, which in turn means that grace periods cannot complete. This injection also needs to be used together with `-DASSERT_0` to discover the liveness violation.

Figure 4: Description of the bug injections we used, identified by the preprocessor option that enables them.

Table 1: Results for Tree RCU litmus test on five Linux kernel versions (time in seconds).

	v3.0			v3.19			v4.3			v4.7			v4.9.6		
	Time		*Traces*	*Time*		*Traces*	*Time*		*Traces*	*Time*		*Traces*	*Time*		*Traces*
Preprocessor Options	SC	TSO	*Explored*	SC	TSO	Explored	SC	TSO	*Explored*	SC	TSO	*Explored*	SC	TSO	*Explored*
–	**227.37**	**260.38**	**19398**	**594.33**	**651.32**	**24760**	**1041.85**	**1147.20**	**28996**	**411.25**	**416.03**	**11076**	**1033.12**	**1125.30**	**28996**
-DASSERT_0	2.09	2.30	145	1.47	1.66	37	1.70	1.77	29	1.93	2.09	29	2.06	2.16	29
-DFORCE_FAILURE_1	2.11	2.27	146	1.51	1.67	41	1.99	1.97	33	2.09	2.26	33	2.13	2.32	33
-DFORCE_FAILURE_2	0.43	0.38	4	0.63	0.66	3	0.80	1.03	3	1.18	1.22	3	1.14	1.27	3
-DFORCE_FAILURE_3	27.60	31.02	2372	399.89	433.38	13264	334.13	316.14	8114	290.50	329.10	8114	307.72	338.26	8114
-DFORCE_FAILURE_4	1.35	1.46	84	3.01	3.13	79	1.87	2.08	24	3.09	3.17	43	2.99	3.26	43
-DFORCE_FAILURE_5	58.15	64.80	4888	1.18	0.98	9	1.22	1.39	9	1.57	1.66	9	1.60	1.65	9
-DFORCE_FAILURE_6	1.14	1.41	1	3.34	3.33	2	5.28	5.13	2	10.40	10.34	2	10.80	10.93	2
-DLIVENESS_CHECK_1 -DASSERT_0	24.91	26.42	2024	10.73	11.32	608	11.27	10.88	488	10.15	11.64	488	10.77	11.85	488
-DLIVENESS_CHECK_2 -DASSERT_0	50.28	52.37	3888	10.46	11.13	608	12.30	11.95	516	11.55	12.49	516	11.80	13.18	516
-DLIVENESS_CHECK_3 -DASSERT_0	26.38	28.98	2184	12.48	13.63	688	11.37	11.23	488	11.66	12.75	532	11.89	13.27	532

(1) Bugs that make the grace period too short, thus permitting an RCU read-side critical section to span the grace period.
(2) Bugs that prevent the grace period from ending.

Both kinds of bug injections represent RCU failures. Injections of the first kind result in a test failure, since the GP guarantee is violated. Injections of the second kind have to be used with an `assert(0)` statement after `synchronize_rcu()`. If this assertion does not trigger for *any* execution of the litmus test, then the grace period does not end for any execution, which in turn signifies that a successful —as opposed to a failed— completion of the test is a liveness violation.

Figure 4 contains information about the bug injections keyed by the define macro that enables each test. For the `FORCE_FAILURE_6` test, an unroll value of 19 has been used and `CONFIG_NR_CPUS` has been set appropriately, while for all other tests an unroll value of 5 has been used. All tests had the desired outcome, something that increases our confidence in our modeling and the verification result for the GP guarantee of Tree RCU's implementation that we report.

7.3 Results and Discussion

In Table 1, the "Time" columns represent the total wall-clock time in seconds (compilation and Nidhugg transformation time are included). As can be seen, there is very little overhead when going from SC to TSO, which shows the power of stateless model checking with source DPOR [1] and chronological traces [2] in this setting. In fact, we show only one "Traces Explored" column for each kernel, since in all tests the total number of explored executions is the same for both SC and TSO. The reason for that is that there are a lot of memory fences in the code of Tree RCU, which prevent store buffer reorderings from happening. But even if reorderings were possible, all bug injections here do not rely on the employed memory model but instead violate the assertions algorithmically. As expected, since the model checking is stateless, the memory requirements are very low (~35MB in most cases except for `FORCE_FAILURE_6` where the memory is increased to ~105MB due to the higher unroll value).

The most interesting row here is the first one, shown in bold. Here Nidhugg needs to explore the complete set of traces in order

to verify that the GP guarantee does indeed hold for Tree RCU's implementation. In rows with failure injections, exploration stops as soon as the failure is detected. How fast this happens depends on the order in which traces are explored. In some cases failures are detected immediately (in the first traces and in less than two seconds) and in other cases only after many traces have been explored.

It can be observed that the number of explored traces varies between different kernel versions. There are fewer traces explored in v3.0 than in v3.19 and v4.3, due to the absence of the grace-period kthread in the first; this thread contains infinite loops which generate many races that Nidhugg tries to reverse. Note that in v3.0 the `-DFORCE_FAILURE_3` and `-DFORCE_FAILURE_5` injections are liveness checks due to the absence of the grace-period kthread. In v4.7 the explored traces decrease dramatically due to the replacement of `rcu_gp_kthread_wake()` with `swake_up()` in `rcu_report_qs_rsp()`. The former performed a check which read a variable that was written by the grace-period kthread (among others), and generated far too many races. In v4.9.6, however, this change was reverted and the explored traces are the same as those of kernel v4.3. Overall, it is obvious that irrespective of the kernel's growth in size, Nidhugg provides an efficient and scalable way to test such a big codebase since it only depends on races on shared variables, and not on the general complexity of the source code.

8 SOME LESSONS LEARNED

The testing process was very educational. We learned lessons both regarding the construction of the model and the model itself, and how to deal with the combinatorial explosion in the number of interleavings that a stateless model checking tool needs to explore.

Arguably, the most valuable lesson learned was the way a Linux kernel model can be constructed. Initially, the way an SMP system should be emulated was not obvious, and the construction of the model had to be precise. Both of these posed two non-trivial challenges; with the kernel occupying more than 15MLOC, the isolation and testing of only the ingredients we cared about was of extreme importance. Still, we managed to use the source code from the kernel directly, and the constructed model is reusable, which means that it can be used again for further RCU testing.

Of course, confining the state space was not in any way an easy task as well. First of all, as far as the model is concerned, the most important design decision we had to make was the way the interrupts are modeled. Initially, we tried to emulate interrupts with per-CPU threads invoking the interrupt handler repeatedly, but unfortunately this approach rendered the state space extremely large. Apart from this, plenty of other design choices were made and most of them are described in Section 5. As fas as the verification of Tree RCU is concerned, multiple different configurations were tried and did not affect the outcome. We chose the one mentioned at Section 7.1 because the state space was considerably smaller. The reason for that, although not obvious from the beginning, is that the updater and the grace-period kthread are mutually exclusive and take advantage of each other's context switches. In addition, we could have ignored the RCU grace-period kthread and invoked `rcu_gp_init()` and `rcu_gp_cleanup()` appropriately, in order to further reduce the state space. However, we wanted our model to be as precise as possible, so we did not resort to such approximations.

9 RELATED WORK

Previous work on RCU verification includes the expression of RCU's formal semantics in terms of separation logic [13] and the verification of user-space RCU in a logic for weak memory [25]. A virtual architecture to model out-of-order memory accesses and instruction scheduling has been proposed [8], and a verification of user-space RCU has been done using the SPIN model checker [9]. Alglave *et al.* verified that RCU's *actual* kernel code preserves data consistency of the object it is protecting [4] using CBMC [6]. Subsequently, McKenney [17] verified the Grace-Period guarantee for Tiny RCU (a flavor of RCU for uniprocessor systems). Finally, mutation testing strategies have been applied to RCU's code [3] as well.

Concurrently with our work, Liang *et al.* used CBMC to verify the Grace-Period guarantee for Tree RCU [14]. However, compared to the work presented here, their approach has some limitations. First of all, due to CBMC's limited support for lists, their modeling does not include callback handling. This has some implications for verification. The most basic one is that bugs in the callback handling mechanism (e.g., a bug similar to the one we reproduced in Section 6) can not be exposed. Considering the fact that RCU's update side primitives are based on callback handling, this limitation is serious. For example, primitives like `call_rcu()` were not included in the tests, and `synchronize_rcu()`'s implementation (which, in reality, is based on `call_rcu()`) had to be emulated. This in turn means that only the underlying grace period mechanism was modeled, and not the callback mechanism that mediates between that mechanism and `synchronize_rcu()`. A second limitation is that the grace-period kthread was not included in the tests. Although in older kernel versions the grace-period kthread did not exist, for newer Linux kernels excluding the kthread from the tests implies alteration of the kernel's operation. In addition, this thread's exclusion means that the way a grace period started and ended also needs to be changed, since the grace-period kthread plays a crucial role in these operations. Finally, the approach of Liang *et al.* does not include the emulation of dynticks-idle mode. In our approach, the dynticks-idle mode is indeed modeled, and our results show that the basic properties of the `dyntick` counters do hold.

Despite the simpler modeling and these limitations, CBMC needs more than 11 hours and 34GB of memory in order to claim successful verification for Tree RCU in kernel v4.3 under TSO [14], whereas Nidhugg only needs 19 minutes and 102MB of memory. More generally, our results are orders of magnitude better, which we attribute to the different algorithms that the two tools employ.

On the other hand, CBMC's underlying algorithm in principle also handles data non-determinism, something that stateless model checking tools in general, and Nidhugg in particular, do not consider. Still, we do not see how data non-determinism plays any role in the verification of the Grace-Period guarantee of Tree RCU for non-preemptible builds. Some supporting evidence for this claim offers the fact that the bug injections we listed in Section 7 are a proper superset of those identified by CBMC.[3] Furthermore, because our approach does include callback handling, we were able to reproduce an older, real kernel bug that was caused by premature callback advancements, which could potentially lead to too short grace

[3] Injections `-DFORCE_FAILURE_1` and `-DFORCE_FAILURE_4` are not considered by Liang *et al.* [14]; the latter due to not modeling the dynticks-idle mode.

periods that violate the GP guarantee. As explained, this bug can not be reproduced with CBMC, due to its limited support for lists.

10 CONCLUDING REMARKS

We described a way to construct a test suite for the systematic concurrency testing of Linux kernel's RCU mechanism. For this, we emulated a non-preemptible Linux-kernel SMP environment and, using the stateless model checking tool Nidhugg, we managed to verify the most basic guarantee that RCU provides for the main flavor used in the Linux kernel, namely, Tree RCU.

More specifically, we verified the Grace-Period guarantee for five different kernel versions, under both a sequentially consistent and a TSO memory model. For all our tests, we used the source code from the Linux kernel directly, with only a handful of changes, which can be and have been scripted.

To show that our emulation of the kernel's environment is sound and to further strengthen our results, we injected RCU failures in our tests, inspired from real bugs that occurred throughout RCU's deployment in production, and Nidhugg was able to identify them all. Moreover, we demonstrated that a patch that applied a well-defined locking design to a variable in an older kernel [24] resolved a much more complex issue that was in effect a bug. We identified and reproduced this bug, providing the exact circumstances under which it occurred. In addition, we tested whether the bug exists in later kernel versions and the answer was negative.

Our work demonstrates that stateless model checking tools like Nidhugg can be used to test *real* code from today's production systems with large code bases. The small time and memory consumption of our tests, especially considering the size and the dynamic nature of the code base tested, underlines the strength of our approach. All the above, along with the fact that our model of the kernel's environment was reused across different kernel versions show that stateless model checking tools can be integrated in Linux kernel's regression testing, and that they can produce useful results. We are currently working together with Linux developers towards making this integration a reality.

Still, we are not yet at a point where we can claim with certainty that the complete implementation of Tree RCU is bug-free; there may be bugs in components of Tree RCU that are not included in our modeling and our tests. In addition, although the GP guarantee is the most significant correctness property of RCU, there are are many other requirements that RCU must meet. Thus, our work could be extended to include more aspects of RCU, and test them under different memory models (e.g., POWER). For example, we could construct tests that include quiescent-state forcing, grace-period expediting and CPU hotplugs. The same applies for the full-dynticks mode which was fully merged in the kernel only relatively recently. Last but not least, the scalability of our results renders the construction of test cases and techniques aiming at the thorough testing of the preemptible Tree RCU extremely interesting as well.

ACKNOWLEDGMENTS

We are very much obliged to Paul E. McKenney for all his help, advice, and suggestions throughout this effort. Also, this paper would not have been possible without Nidhugg; we thank its main developer, Carl Leonardsson, for the hard work he has put into it.

REFERENCES

[1] Parosh Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. 2014. Optimal Dynamic Partial Order Reduction. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 373–384. https://doi.org/10.1145/2535838.2535845

[2] Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas. 2015. Stateless Model Checking for TSO and PSO. In *Tools and Algorithms for the Construction and Analysis of Systems (LNCS)*, Vol. 9035. Springer, Berlin, Heidelberg, 353–367. https://doi.org/10.1007/978-3-662-46681-0_28

[3] Iftekhar Ahmed, Alex Groce, Carlos Jensen, and Paul E. McKenney. 2015. How Verified is My Code? Falsification-Driven Verification. In *Proceedings of the 30th IEEE/ACM International Conference on Automated Software Engineering*. IEEE Computer Society, Washington, DC, USA, 737–748.

[4] Jade Alglave, Daniel Kroening, and Michael Tautschnig. 2013. Partial Orders for Efficient Bounded Model Checking of Concurrent Software. In *Computer Aided Verification (LNCS)*, Vol. 8044. Springer, Berlin, Heidelberg, 141–157. https://doi.org/10.1007/978-3-642-39799-8_9

[5] John Callaham. 2015. Google says there are now 1.4 billion active Android devices worldwide. (29 Sept. 2015). http://www.androidcentral.com/google-says-there-are-now-14-billion-active-android-devices-worldwide

[6] Edmund M. Clarke, Daniel Kroening, and Flavio Lerda. 2004. A Tool for Checking ANSI-C Programs. In *Tools and Algorithms for the Construction and Analysis of Systems (LNCS)*, Kurt Jensen and Andreas Podelski (Eds.), Vol. 2988. Springer, Berlin, Heidelberg, 168–176. https://doi.org/10.1007/978-3-540-24730-2_15

[7] CPUHotplugDoc 2016. CPU Hotplug in the Kernel. (Dec. 2016). https://www.kernel.org/doc/Documentation/core-api/cpu_hotplug.rst

[8] Mathieu Desnoyers, Paul E. McKenney, and Michel R. Dagenais. 2013. Multi-core Systems Modeling for Formal Verification of Parallel Algorithms. *SIGOPS Oper. Syst. Rev.* 47, 2 (July 2013), 51–65. https://doi.org/10.1145/2506164.2506174

[9] Mathieu Desnoyers, Paul E. McKenney, Alan S. Stern, Michel R. Dagenais, and Jonathan Walpole. 2012. User-Level Implementations of Read-Copy Update. *IEEE Trans. Parallel Distrib. Syst.* 23, 2 (Feb. 2012), 375–382. https://doi.org/10.1109/TPDS.2011.159

[10] Cormac Flanagan and Patrice Godefroid. 2005. Dynamic Partial-order Reduction for Model Checking Software. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 110–121. https://doi.org/10.1145/1040305.1040315

[11] Patrice Godefroid. 1997. Model checking for programming languages using VeriSoft. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 147–186.

[12] Patrice Godefroid. 2005. Software Model Checking: The VeriSoft Approach. *Formal Methods in System Design* 26, 2 (2005), 77–101.

[13] Alexey Gotsman, Noam Rinetzky, and Hongseok Yang. 2013. Verifying Concurrent Memory Reclamation Algorithms with Grace. In *Programming Languages and Systems (LNCS)*, Vol. 7792. Springer, Berlin, Heidelberg, 249–269. https://doi.org/10.1007/978-3-642-37036-6_15

[14] Lihao Liang, Paul E. McKenney, Daniel Kroening, and Tom Melham. 2016. Verification of the Tree-Based Hierarchical Read-Copy Update in the Linux Kernel. (Oct. 2016). http://arxiv.org/abs/1610.03052

[15] LLVM-Atomics 2017. LLVM Atomic Instructions and Concurrency Guide. (2017). http://llvm.org/docs/Atomics.html#libcalls-atomic

[16] Paul E. McKenney. 2008. Hierarchical RCU. (2008). http://lwn.net/Articles/305782/

[17] Paul E. McKenney. 2015. Verification Challenge 4: Tiny RCU. (9 March 2015). http://paulmck.livejournal.com/39343.html

[18] Paul E. McKenney and John D. Slingwine. 1998. Read-Copy Update: Using Execution History to Solve Concurrency Problems. In *Parallel and Distributed Computing and Systems*. 509–518.

[19] Paul E. McKenney and Jonathan Walpole. 2007. What is RCU, Fundamentally? (17 Dec. 2007). http://lwn.net/Articles/262464/

[20] Mandanlal Musuvathi, Shaz Qadeer, Thomas Ball, Gerald Basler, Piramanayagam Arumuga Nainar, and Iulian Neamtiu. 2008. Finding and Reproducing Heisenbugs in Concurrent Programs. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation*. USENIX Association, Berkeley, CA, USA, 267–280.

[21] NOHZDoc 2017. NO_HZ: Reducing Scheduling-Clock Ticks. (April 2017). https://www.kernel.org/doc/Documentation/timers/NO_HZ.txt

[22] Abhishek Prakash. 2016. Linux now runs on 99.6% of top 500 supercomputers. (15 Nov. 2016). https://itsfoss.com/linux-99-percent-top-500-supercomputers/

[23] RCU-Bug-AdditionalPatches 2009. RCU: Clean up locking for `->completed` and `->gpnum` fields. (Oct. 2009). https://lkml.org/lkml/2009/10/30/212

[24] RCU-Bug-Synch 2009. RCU: Fix synchronization for `rcu_process_gp_end()` uses of `->completed` counter. (Nov. 2009). https://lkml.org/lkml/2009/11/4/69

[25] Joseph Tassarotti, Derek Dreyer, and Viktor Vafeiadis. 2015. Verifying Read-Copy-Update in a Logic for Weak Memory. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, New York, NY, USA, 110–120. https://doi.org/10.1145/2737924.2737992