

# The List is Our Process!

## An analysis of the kernel's email-based development process

Ralf Ramsauer\*, Sebastian Duda†, Lukas Bulwahn‡, Wolfgang Mauerer\*§

\* Technical University of Applied Sciences Regensburg

† Friedrich-Alexander University Erlangen-Nürnberg

‡ Hobbyist, active in LF ELISA Project, employed at BMW AG

§ Siemens AG, Corporate Research and Technology, Munich

Embedded Linux Conference Europe, Lyon

October 28, 2019

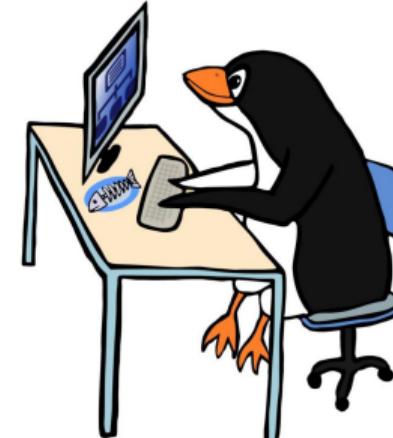


OSTBAYERISCHE  
TECHNISCHE HOCHSCHULE  
REGENSBURG



## Our Overall Goal

Formalising and assessing the Linux  
Kernel development process



© Moini  
[openclipart.org](http://openclipart.org)

## Our Overall Goal

Formalising and assessing the Linux Kernel development process

## Outside / Inside Motivation

- ▶ Safety-Critical Development
- ▶ Development Process Assessment
- ▶ Monitoring (cf. CHAOSS)
- ▶ Fundamentals of Software Engineering



© Moini  
openclipart.org

Motivation from *inside* the community

## Interest of the kernel community itself

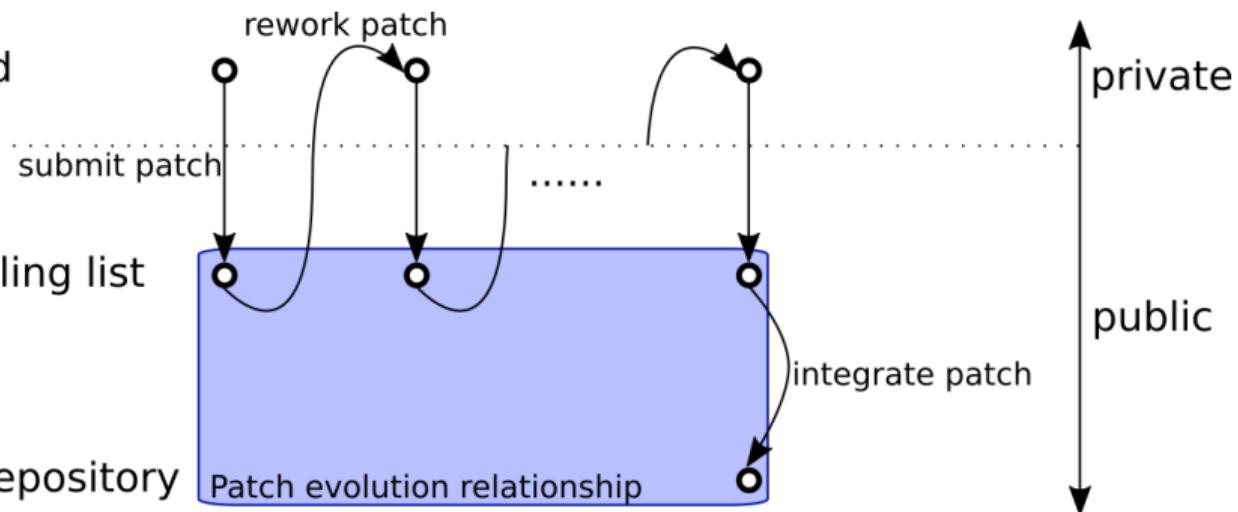
- ▶ D. Williams, Towards a Linux Kernel Maintainer Handbook, LPC 2018
- ▶ J. Corbet, Change IDs for kernel patches, <https://lwn.net/Articles/797613/>
- ▶ [Ksummit-discuss] [MAINTAINERS SUMMIT] Patch version changes in commit logs?
- ▶ [Ksummit-discuss] Allowing something Change-Id (or something like it) in kernel commits

## Towards a formal model of the development process

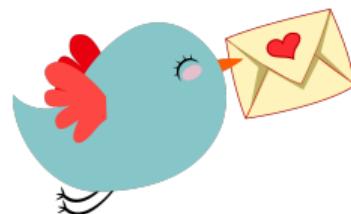
1. patch is created

2. patch is on mailing list

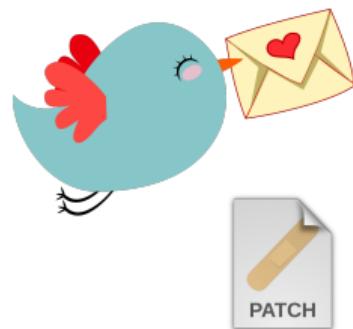
3. patch is in git repository



## Linux Kernel development workflow



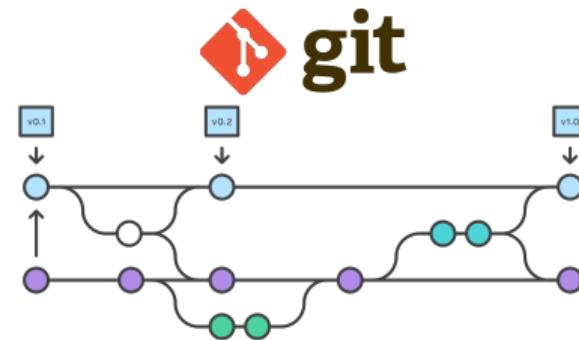
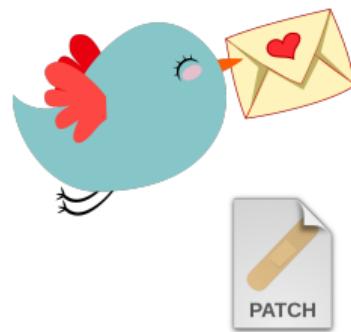
## Linux Kernel development workflow



## Linux Kernel development workflow



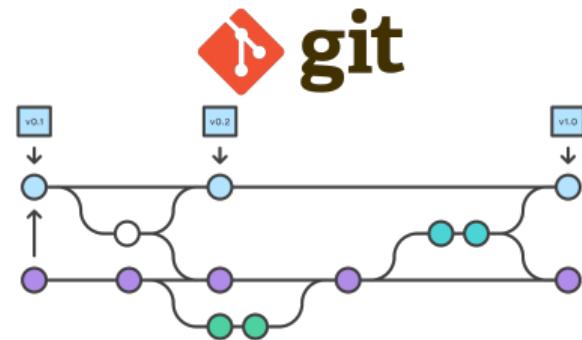
## Linux Kernel development workflow



## Linux Kernel development workflow



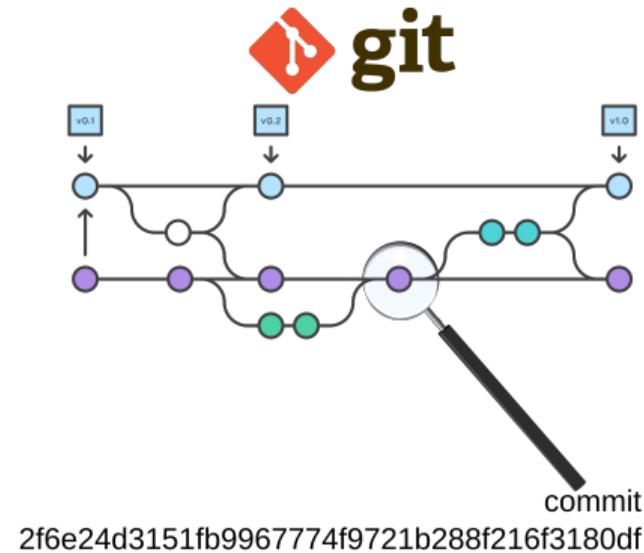
Message-Id:  
<1531137835-21581-1-git@1wt.eu>



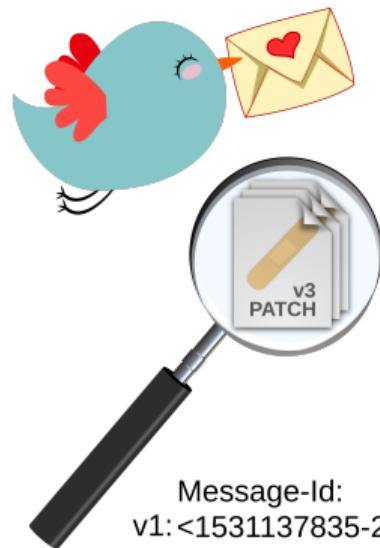
## Linux Kernel development workflow



Message-Id:  
<1531137835-21581-1-git@1wt.eu>

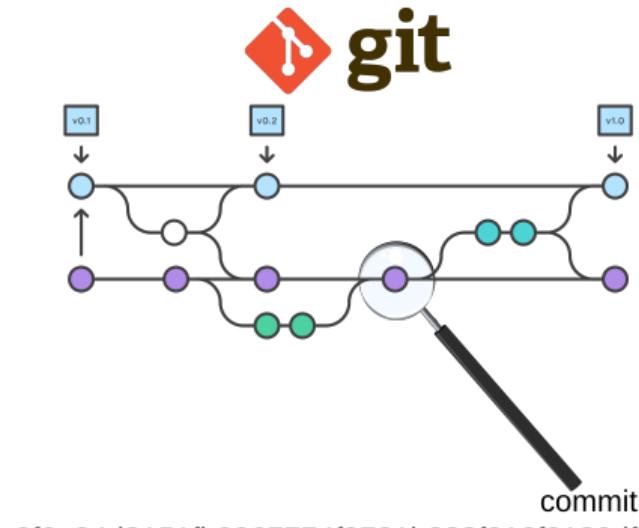


## Linux Kernel development workflow



Message-Id:

v1:<1531137835-21581-1-git@1wt.eu>  
v2:<6739637657-68462-1-git@1wt.eu>  
v3:<9717683099-75474-1-git@1wt.eu>

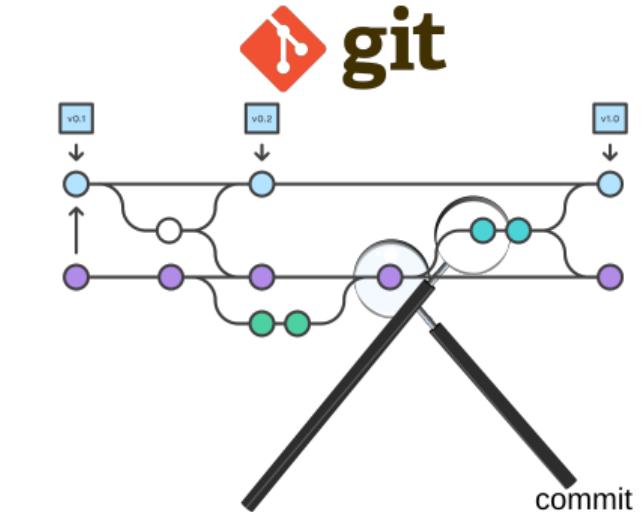


2f6e24d3151fb9967774f9721b288f216f3180df

## Linux Kernel development workflow

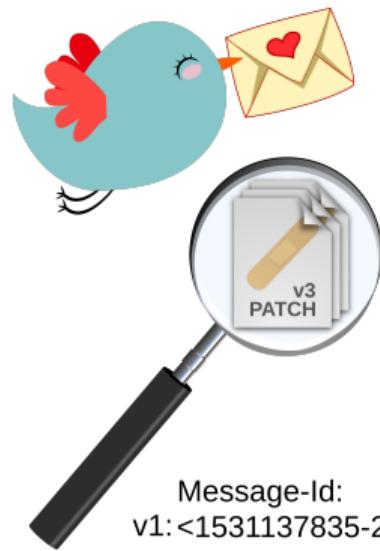


Message-Id:  
v1:<1531137835-21581-1-git@1wt.eu>  
v2:<6739637657-68462-1-git@1wt.eu>  
v3:<9717683099-75474-1-git@1wt.eu>

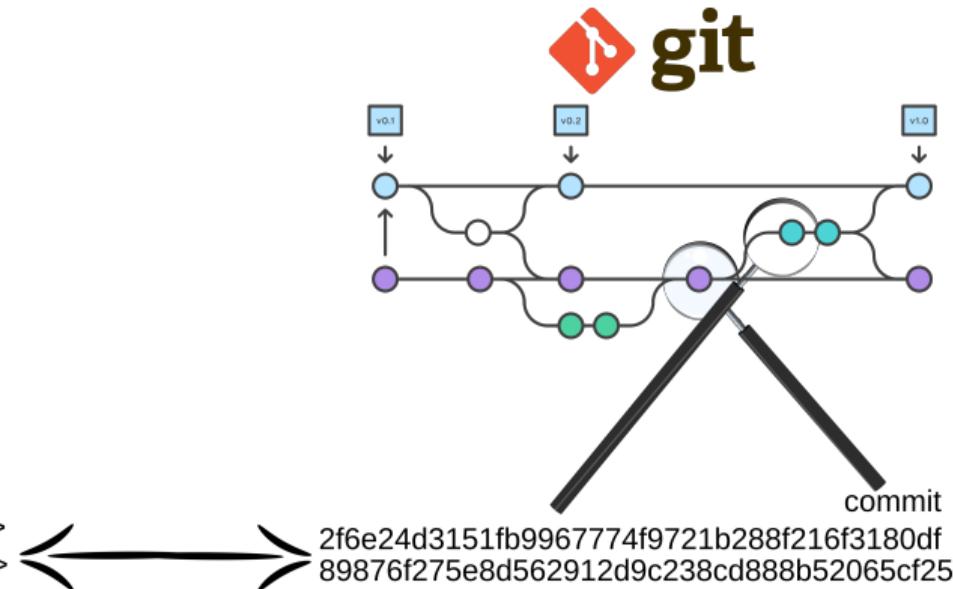


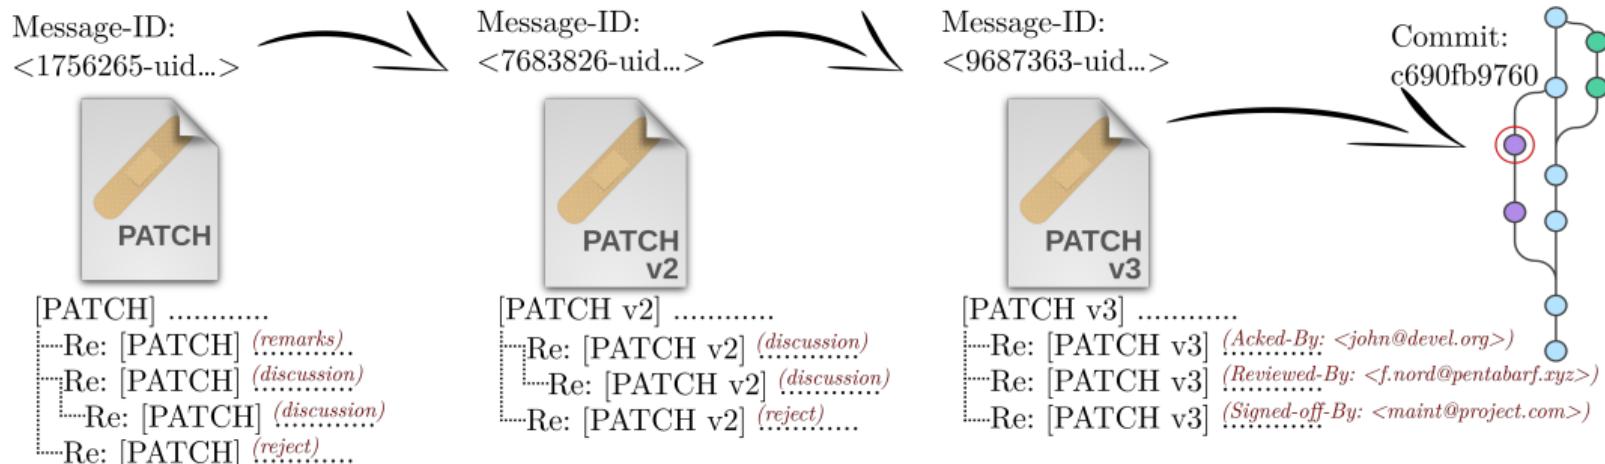
2f6e24d3151fb9967774f9721b288f216f3180df  
89876f275e8d562912d9c238cd888b52065cf25

## Linux Kernel development workflow



Message-Id:  
v1:<1531137835-21581-1-git@1wt.eu>  
v2:<6739637657-68462-1-git@1wt.eu>  
v3:<9717683099-75474-1-git@1wt.eu>





## PaStA - Patch Stack Analysis

- ▶ Detects *similar* patches across different branches
- ▶ Quantify mainlining efforts of off-tree developments (Preempt\_RT, vendor kernels, ...)
- ▶ Works with mailing lists!



Source: [toplock.net.au](http://toplock.net.au)

# Example of similar patches

```
commit 91824d74d6d85f58c63a66b8f2c7993ae246181b
Author: Thomas Gleixner <tglx@linutronix.de>
Date: Mon Sep 12 21:45:49 2011 +0200

sched-cure-utter-idle-accounting-madness.patrch

Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
```

```
diff —git a/kernel/sched.c b/kernel/sched.c
index 205499a..1121a97 100644
--- a/kernel/sched.c
+++ b/kernel/sched.c
@@ -5037,7 +5037,13 @@ EXPORT_SYMBOL(task_nice);
 */
int idle_cpu(int cpu)
{
-    return cpu_curr(cpu) == cpu_rq(cpu)->idle;
+    struct rq *rq = cpu_rq(cpu);
+
+#ifdef CONFIG_SMP
+    return rq->curr == rq->idle && !rq->nr_running && !rq->wake_list;
+#else
+    return rq->curr == rq->idle && !rq->nr_running;
+#endif
}

/**
```



```
commit 908a3283728d92df36e0c7cd63304fd35e93a8a9
Author: Thomas Gleixner <tglx@linutronix.de>
Date: Thu Sep 15 15:32:06 2011 +0200

sched: Fix idle_cpu()

On -rt we observed hackbench waking all 400 tasks to a single
cpu. This is because of select_idle_sibling()'s interaction
with the new ipi based wakeup scheme.
[...snip...]
Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
Signed-off-by: Peter Zijlstra <a.p.zijlstra@chello.nl>
Link: http://lkml.kernel.org/n/tip-3030p18b2[...]
Signed-off-by: Ingo Molnar <mingo@elte.hu>
```

```
diff —git a/kernel/sched.c b/kernel/sched.c
index 1874c74..4cdcc91c 100644
--- a/kernel/sched.c
+++ b/kernel/sched.c
@@ -5138,7 +5138,20 @@ EXPORT_SYMBOL(task_nice);
 */
int idle_cpu(int cpu)
{
-    return cpu_curr(cpu) == cpu_rq(cpu)->idle;
+    struct rq *rq = cpu_rq(cpu);
+
+    if (rq->curr != rq->idle)
+        return 0;
+
+    if (rq->nr_running)
+        return 0;
+
+#ifdef CONFIG_SMP
+    if (!list_empty(&rq->wake_list))
+        return 0;
+#endif
+
+    return 1;
}

/**
```



# Example of similar patches

```
commit 91824d74d6d85f58c63a66b8f2c7993ae246181b
Author: Thomas Gleixner <tglx@linutronix.de>
Date: Mon Sep 12 21:45:49 2011 +0200
```

```
sched-cure-utter-idle-accounting-madness.patch
```

```
Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
```

```
diff --git a/kernel/sched.c b/kernel/sched.c
index 205499a..1121a97 100644
```

```
+     struct rq *rq = cpu_rq(cpu);
+
+ #ifdef CONFIG_SMP
+     return rq->curr == rq->idle && !rq->nr_r
+ #else
+     return rq->curr == rq->idle && !rq->nr_r
+ #endif
```

```
/**
```



```
commit 908a3283728d92df36e0c7cd63304fd35e93a8a9
Author: Thomas Gleixner <tglx@linutronix.de>
Date: Thu Sep 15 15:32:06 2011 +0200
```

```
sched: Fix idle_cpu()
```

On -rt we observed hackbench waking all 400 tasks to a single  
cpu. This is because of select\_idle\_sibling()'s interaction  
with the new ipi based wakeup scheme.

[..snip..]

Signed-off-by: Thomas Gleixner <tglx@linutronix.de>

```
+     struct rq *rq = cpu_rq(cpu);
+
+     if (rq->curr != rq->idle)
+         return 0;
+
+     if (rq->nr_running)
+         return 0;
+
+ #ifdef CONFIG_SMP
+     if (!list_empty(&rq->wake_list))
+         return 0;
+ #endif
+
+     return 1;
+
+     if (!list_empty(&rq->wake_list))
+         return 0;
+ #endif
+
+     return 1;
}
```



# Example of similar patches

commit 91824d74d6d85f58c63a66b8f2c7993ae246181b

```
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Mon Sep 13 15:32:06 2011 +0200

 sched-cure [ ... ]
 sched-cure #endif
 sched-cure #ifdef CONFIG_SMP
 sched-cure     cpu_rq(cpu)
 sched-cure return
 sched-cure rq
 sched-cure rq->curr
 sched-cure rq->idle
 sched-cure -    re
 sched-cure +    st
 sched-cure +    rq->nr_running
 sched-cure +    rq->wake_list
 sched-cure +    struct
 sched-cure +    return rq->curr == rq->idle && !rq->nr_running;
 sched-cure +    #endif
 sched-cure +}
 sched-cure */
 sched-cure
```



commit 908a3283728d92df36e0c7cd63304fd35e93a8a9

Author: Thomas Gleixner <tglx@linutronix.de>

Date: Thu Sep 15 15:32:06 2011 +0200

sched: Fix idle\_cpu()

On -rt we observed hackbench waking all 400 tasks to a single  
cpu. This was due to the interaction between the idle() and the

```
[ ... ]
#endif
#endif
if
Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
Signed-off-by: Christian Rauch <ra@chello.nl>
Link: http://lkml.org/lkml/2011/9/15/18b2 [...]
Signed-off-by: Christian Rauch <ra@chello.nl>
```

diff --git

index 2054

--- a/kern

+++ b/kern

@@ -5037,7

\*/
int idle\_c

{

- re

+ st

rq->nr\_running

rq->wake\_list

struct

+ return rq->curr == rq->idle && !rq->nr\_running;

#endif

}

/\*
\*\*

On -rt we observed hackbench waking all 400 tasks to a single

cpu. This was due to the interaction between the idle() and the

```
[ ... ]
#endif
#endif
if
Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
Signed-off-by: Christian Rauch <ra@chello.nl>
Link: http://lkml.org/lkml/2011/9/15/18b2 [...]
Signed-off-by: Christian Rauch <ra@chello.nl>
```

diff --git

index 18b2

--- a/kern

+++ b/kern

@@ -513

\*/
int idle

{

- re

+ st

rq->curr

rq->idle

rq->nr\_running

rq->wake\_list

struct

+ return 0;

#endif

+ if (!llist\_empty(&rq->wake\_list))

+ return 0;

#endif

+ return 1;

}

/\*
\*\*



# Example of similar patches

commit 91824d74d6d85f58c63a66b8f2c7993ae246181b

Author: Thomas Gleixner  
Date: Mon Jul 11 13:50:18 2011 +0200

sched-cure  
Signed-off-by: Thomas Gleixner <tglx@linutronix.de>

diff —git  
index 2054  
--- a/kern  
+++ b/kern

@@ -5037,7  
\* /  
int idle\_rq(struct rq \*rq)

{  
- requeue\_task(rq);  
+ struct st;  
+ #ifdef CONFIG\_SMP  
+ requeue\_all\_tasks\_on\_rq(rq);  
+ #else

+ return rq->curr == rq->idle && !rq->nr\_running;  
+ #endif  
}

/\*\*



commit 908a3283728d92df36e0c7cd63304fd35e93a8a9

Author: Thomas Gleixner <tglx@linutronix.de>

Date: Thu Sep 15 15:32:06 2011 +0200

sched: Fix idle\_cpu()

On -rt we observed hackbench waking all 400 tasks to a single  
cpu. This was due to the interaction between sched\_wake() and  
idle\_cpu().

[...]  
#endif  
#ifdef CONFIG\_SMP  
#if defined(CONFIG\_SMP) & !defined(CONFIG\_SCHED\_DEADLINE)  
#endif

diff —git  
index 18b218b2  
--- a/kern  
+++ b/kern  
@@ -513,7  
\* /  
int idle\_rq(cpu\_struct \*rq)

{  
- requeue\_all\_tasks\_on\_rq(rq);  
+ struct st;  
+ #ifdef CONFIG\_SMP  
+ requeue\_all\_tasks\_on\_rq(rq);  
+ #else

+ return rq->curr == rq->idle && !rq->nr\_running;  
+ #endif  
}

/\*\*

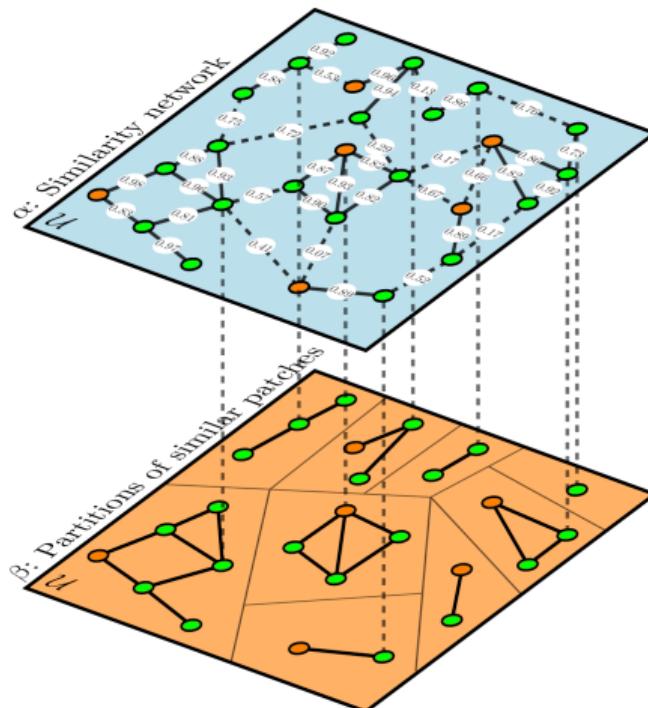
return 0;  
+ if (!rq->idle & !rq->nr\_running & !list\_empty(&rq->wake\_list))  
+ return 0;

#endif  
+ return 1;  
}

/\*\*



**Diff similarity: 0.875**



### Legend

- ▶ green nodes: patches on MLs
- ▶ orange nodes: commits in repository
- ▶ edges: similarity of patches/commits
  - ▶ dashed: similarity below thres
  - ▶ solid: similarity above thres

## Interested in the techniques? More details in:

### The List is the Process: Reliable Pre-Integration Tracking of Commits on Mailing Lists

Ralf Ramsauer<sup>\*</sup>, Daniel Lohmann<sup>†</sup> and Wolfgang Mauerer<sup>‡</sup>

<sup>\*</sup>Technical University of Applied Sciences Regensburg

<sup>†</sup>University of Passau

<sup>‡</sup>Siemens AG, Corporate Technology, Munich

ralf.ramsauer@fh-reg.de, lohmann@ca.tu-passau.de, wolfgang.mauerer@ethz.ch

Abstract—A considerable corpus of research on software engineering focuses on creating effective repositories, but often omits their pre-integration history.

We present a novel method for tracking the software engineering cycle by observing the evolution of changes in their first version in repositories. Since artifact modifications on mailing lists are communicated by update messages, we propose a patch stack analysis approach to study changes in a non-intrusive way. That is, our approach does not require any modifications to the mailing list or to the source code. Instead, it analyzes the evolution of patches in mailing lists of open-source software (OSS) projects like the Linux kernel, and studies its high accuracy using an elaborate control experiment.

Our approach is used to quantify properties of OSS development processes, which is an essential requirement for our approach to be useful. We show that the quality of the process and conformance to processes are notable. The high accuracy of our technique allows, at the best of our knowledge, to automatically and quantitatively determine if open development processes effectively align with formal process requirements.

#### 1. INTRODUCTION

Software patches may have come a long way before their final integration into the official branch (known as master or trunk) of an OSS project. In many cases, patches are developed in separate branches, either in the same repository as the main project, or in a different repository. Another reason is the origin of a patch can be unknown from other developers' repositories (i.e., integration of branches or patches from foreign repositories, pull requests on web-based repository managers such as GitHub, vendor patches, patch stacks, and mailing lists (MLs)).

Especially MLs have been in use for software development processes for decades [17]. They have a long history (plain text emails), and come with an absolute minimum of tool support. As a result, they are often considered as legacy systems. Despite their simplicity, stability, reliability and interface robustness, they are still widely used in many open-source software (OSS) projects. In particular, mailing lists are a core infrastructure component of long-living OSS projects such as low-level systems software (e.g., QEMU, U-Boot, GRUB, etc.).

This work was supported by Siemens AG, Corporate Research, the German Research Foundation (DFG) under grant SFB 1053, and the Bavarian Ministry of Economic Affairs, Energy, and Technology (BMEV). The DFG project has received funding from the ECSEL Joint Undertaking (EJSU), a public–private partnership between the European Union's Horizon 2020 research and innovation program, EJSU, its associated partners, and the European Commission. It is co-financed by the consortium members, grants from Austria, Germany, Italy, Spain and Switzerland.

systems (e.g., the Linux kernel) or foundations (e.g., Apache, GNU). Mailing lists form the backbone of the software development process [23]. They are not only used to ask questions, file bug reports or discuss general topics, but implement a patch submit-review-approve strategy for stepwise refinement [41] that is typically started immediately before a patch is finally integrated into the mainline [17, Fig. 1].

Therefore, MLs contain a huge amount of information on the pre-integration history of patches. A commit in a repository may be the earliest of the process, when the intermediate state of the patch is in the system. Mailing lists also allow us to analyze development history and code evolution, but also enable us to inspect reviewing and maintenance processes. They further allow inferring organizational [30] and socio-cultural [11, 12, 13, 40] aspects of software development. This is possible because MLs contain information on interactions between developers.

Nevertheless, open source components are mostly deployed in industrial, safety-critical, and mission-critical applications such as medical devices or in automotive products. Especially for core components of a system that implement business-wise non-functional features such as the system-wide stock of medical devices, the quality of the patch set is of great importance. Another aspect, the origin of a patch can be unknown from other developers' repositories (i.e., integration of branches or patches from foreign repositories, pull requests on web-based repository managers such as GitHub, vendor patches, patch stacks, and mailing lists (MLs)).

Especially MLs have been in use for software development processes for decades [17]. They have a long history (plain text emails), and come with an absolute minimum of tool support. As a result, they are often considered as legacy systems. Despite their simplicity, stability, reliability and interface robustness, they are still widely used in many open-source software (OSS) projects such as low-level systems software (e.g., QEMU, U-Boot, GRUB, etc.).

Compared to conventional, often more proprietorial industrial software development environments [35], and often requires fundamentally different development processes [15] because of project size and a high number of massively distributed stakeholders. Because of this nature of OSS, projects do not necessarily need to conform to standard practices [13].

Open source, nevertheless, across different industrial sectors share similar concerns on the use of OSS components [18, 19]:

### Observing Custom Software Modifications: A Quantitative Approach of Tracking the Evolution of Patch Stacks

Ralf Ramsauer  
Technical University of Applied Sciences Regensburg  
ralf.ramsauer@fh-reg.de

Daniel Lohmann  
Friedrich-Alexander University Erlangen-Nürnberg  
lohmann@cs.fau.de

Wolfgang Mauerer  
Technical University of Applied Sciences Regensburg  
wolfgang.mauerer@ethz.ch

#### ABSTRACT

Motivation to open-source software (OSS) are often provided via the term of "patch stacks"—sets of changes (patches) that modify a given body of code. Maintaining patch stacks is a common task in OSS development, which is usually the underlying base project changes frequently. This necessitates a continuous and engineering-intensive adaptation of the patch stack to the base project (so-called "stack-walking") into stepwise evolution. This follows several important problems for changes that are not integrated into projects, for instance when they are controversial or only of minor importance.

We present and explore a methodology to systematically capture the temporal evolution of patch stacks, track their evolution, and measure their quality, stability, and consistency, and estimate the eventual research and engineering effort required to successfully develop and maintain patch stacks. Our methodology consists of three main steps: analysis on patches (including statistical analysis and other methods that lead to actionable advice on the construction and long-term maintenance of custom extensions to OSS).

#### 1. INTRODUCTION

Special purpose software, like industrial control, medical analysis, or other domain-specific applications, often can provide basic building blocks. Custom modifications implemented on top of these building blocks, however, are classified as controversial and are hypothesized to reflect the maintenance status of the whole stack. Thus, the gradual introduction of different patch types that display the actual modifications could function as a measure for the inconsistency of the stack.

In this paper, we make the following contributions:

- We provide an approach and tool for observing the evolution of patch stacks.
- We propose a language-independent semi-automatic algorithm based on string distances that is suitable for detecting similar patches on patch stacks.
- We provide a case study on Preempt-RT [30], a real-life extension of the Linux kernel that enforces widespread real-time guarantees on a wide range of embedded devices, yet has not been integrated into standard Linux. We measure its influence on maintain and visualize the development dynamics of the stack.

#### 2. APPROACH

In general, a patch stack (also known as patch set) is defined as a set of patches (changes) that are developed and later integrated into the base project. Well-known examples include the Preempt-RT Linux runtime extension, the Linux Long Term Support Initiative kernel, and the Yocto Project [1]. Patch stacks are often applied as a particular backbone. In many cases, patch stacks are applied as top of individual releases of an upstream version, or as a top of a specific feature branch (so-called "feature branch"). The contents of the patched version of a base project are identified as the set of commit hashes that do not appear in the original base project.

Our analysis is based on the following assumptions:

<https://github.com/BL/PatchIt>

## Data Acquisition

- ▶ Dumps from gmane.org etc.
- ▶ kernel.org public inboxes
  - ▶ some lists, prehistoric data
  - ▶ <https://lore.kernel.org/lists.html>
  - ▶ Some lists are imports from gmane.org :-)
- ▶ Our own collection



## Data Acquisition

- ▶ Dumps from gmane.org etc.
- ▶ kernel.org public inboxes
  - ▶ some lists, prehistoric data
  - ▶ <https://lore.kernel.org/lists.html>
  - ▶ Some lists are imports from gmane.org :-)
- ▶ Our own collection
  - ▶ 200 lists, since May '19
  - ▶ <https://github.com/linux-mailinglist-archives>



© Penguins of Madagascar  
20th Century Fox

## Let there be chaos

- ▶ Br?ken encoding
- ▶ BÃse64
- ▶ ▶
- ▶ MUAs
- ▶ Bots
- ▶ HTML
- ▶ Automated mails
- ▶ non-Linux patches
- ▶ Stable reviews
- ▶ Malformed recipients
- ▶ ...



Message-Id: <74851t0\$h3103wn0\$Delldi Fri, 9 Mar 71685 18:45:56  
+0000

Date: Mon, 08 Aug 05 04:01:15 ?x?\_???????

Date: Tue, 27 Mar 22001 13:42:39 +0200 (Westeurop?ische Sommerzeit)

X-Mailer: Microsoft Outlook Express 6.00.2900.3028

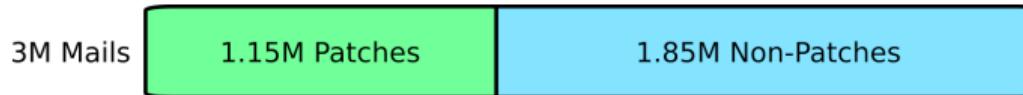
## We analyse...

- ▶ v2.6.39..linus/master
- ▶ ≈610K commits
- ▶ Mails: 2011-05-01 – 2018-12-31
- ▶ ≈3M mails
- ▶ Lists: All Public Inboxes from [lore.kernel.org](http://lore.kernel.org)

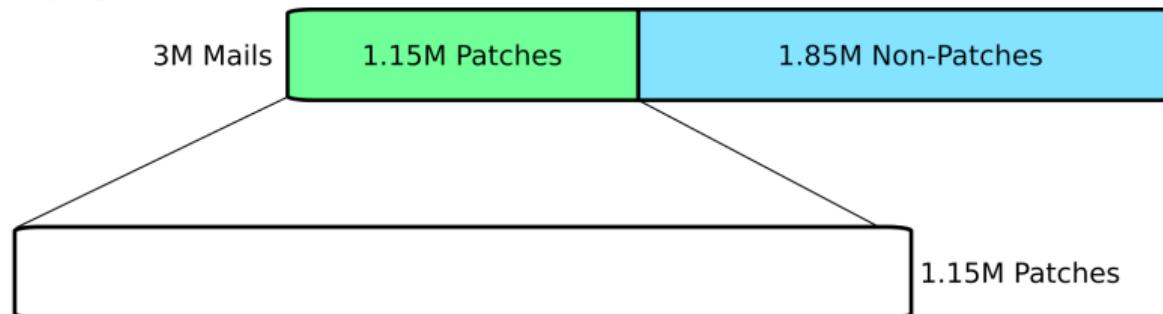
## We analyse...

- ▶ v2.6.39..linus/master
- ▶ ≈610K commits
- ▶ Mails: 2011-05-01 – 2018-12-31
- ▶ ≈3M mails
- ▶ Lists: All Public Inboxes from [lore.kernel.org](http://lore.kernel.org)
- ▶ [linux-amlogic](#), **linux-arm-kernel**, [linux-i3c](#), [linux-mtd](#), [linux-riscv](#), [linuxppc-dev](#), [coccinelle](#), [linux-block](#), [linux-bluetooth](#), [linux-btrfs](#), [linux-cifs](#), [linux-clk](#), [linux-crypto](#), [linux-ext4](#), [linux-fsdevel](#), [linux-hwmon](#), [linux-iio](#), [linux-integrity](#), **linux-kernel**, [linux-media](#), [linux-mips](#), [linux-modules](#), [linux-next](#), **netdev**, [linux-nfs](#), [linux-parisc](#), [linux-pci](#), [linux-renesas-soc](#), [linux-rtc](#), [linux-security-module](#), [linux-sgx](#), [linux-trace-devel](#), [linux-watchdog](#), **linux-wireless**

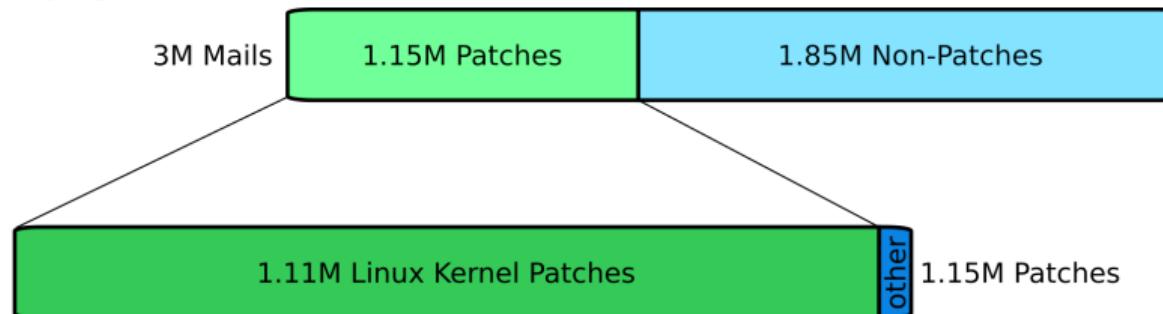
2011-05-2018-12



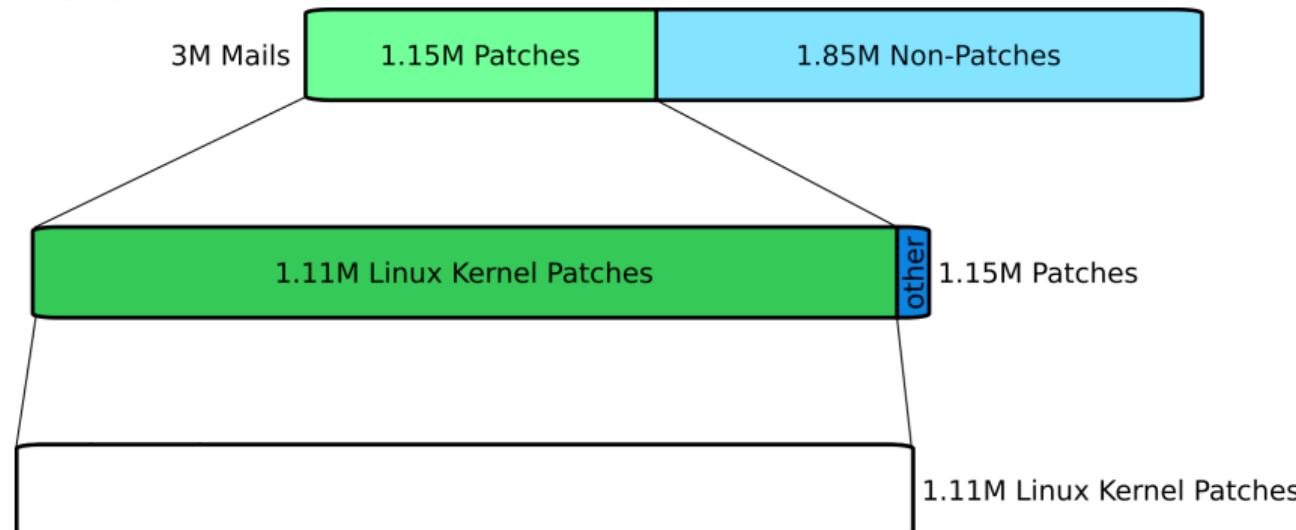
2011-05-2018-12



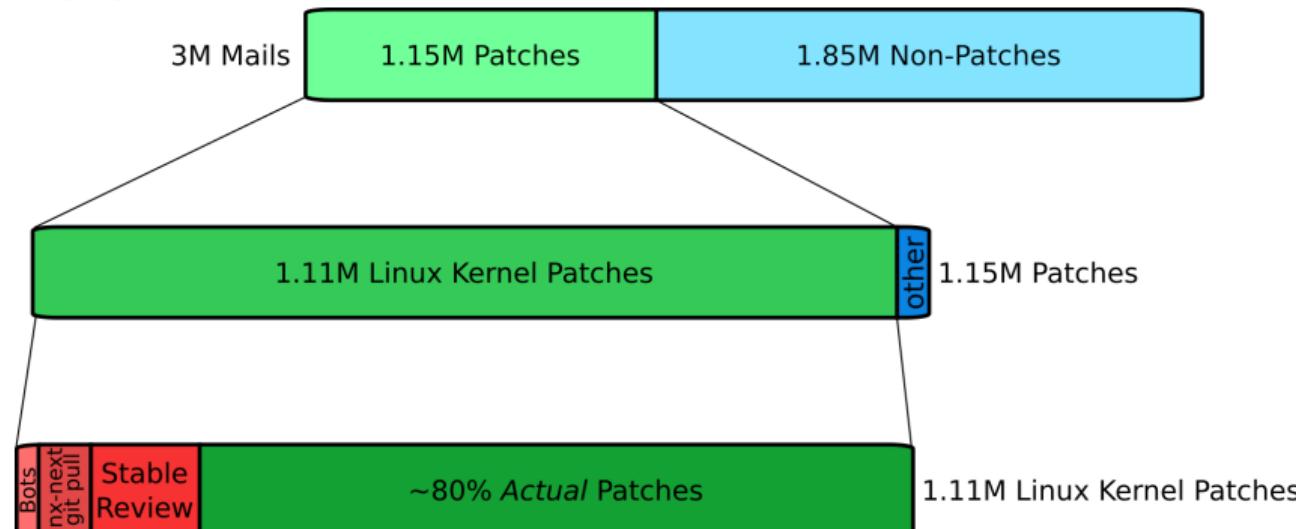
2011-05-2018-12



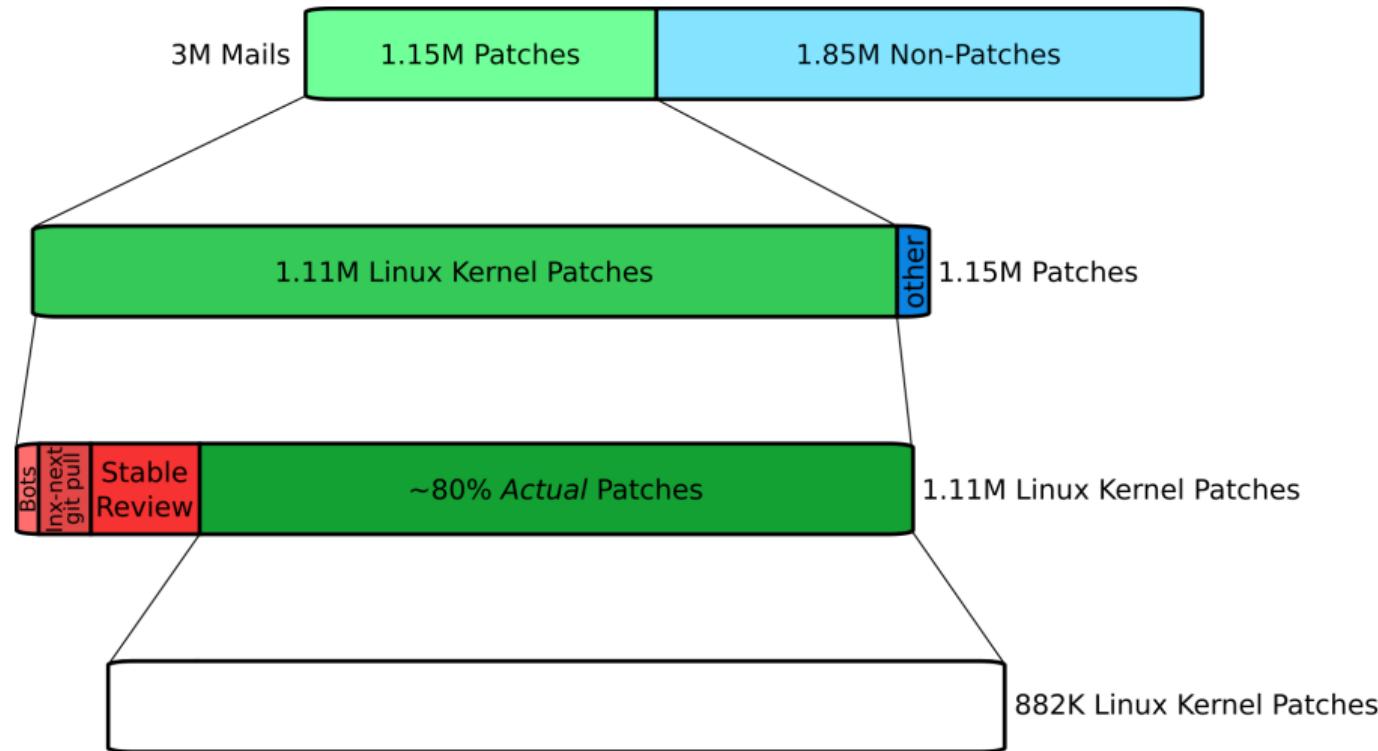
2011-05-2018-12



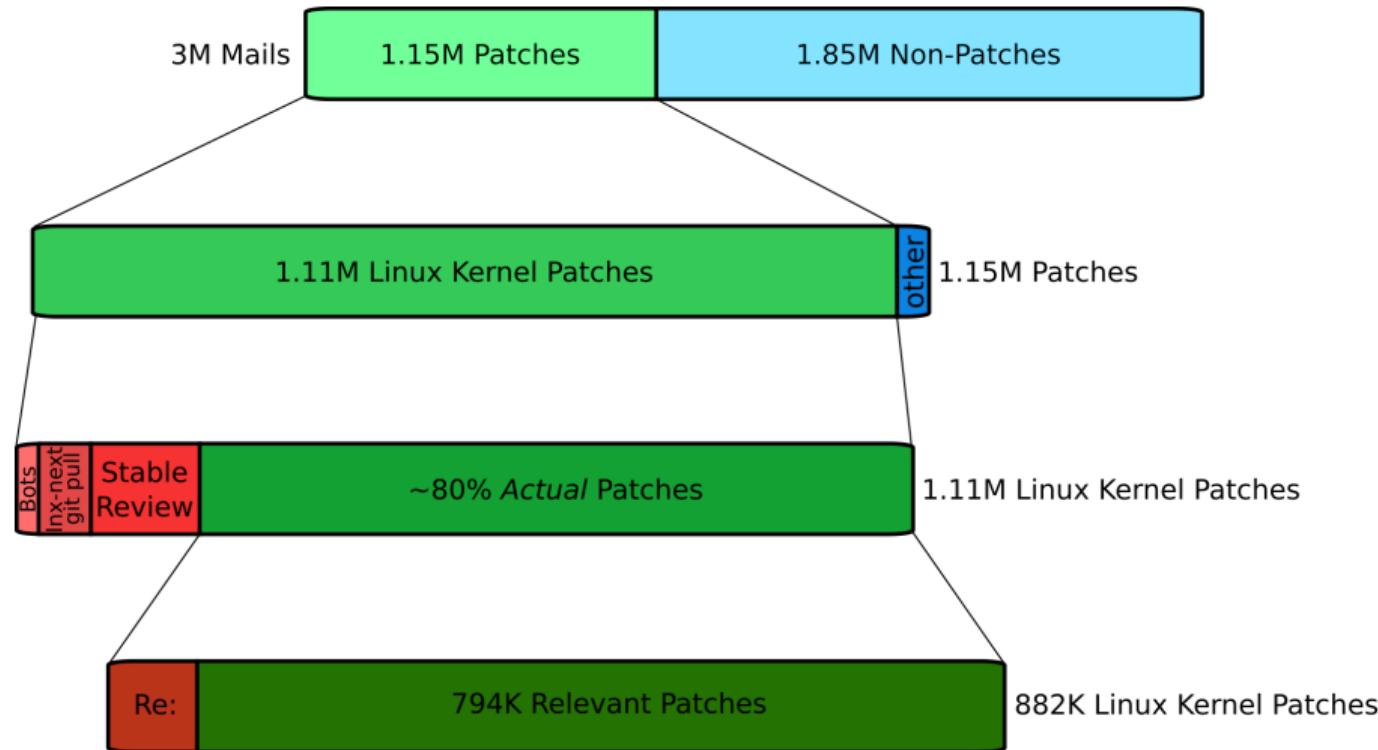
2011-05-2018-12



2011-05-2018-12



2011-05-2018-12



## Ignored Patches

### Research Question

Are there specific characteristics for ignored patches?

### Definition

A patch on a ML is *ignored* if...

- ▶ ... the thread of the patch has no responses from persons other than the author
- ▶ ... the patch was not accepted upstream
- ▶ ... all related patches (e.g., revisions in other series) were ignored

## Ignored Patches

## By the Numbers...

- ▶ **lore.kernel.org lists 2011-2018: ø2.5% ignored patches**
  - ▶ 2011: ø3.9%
  - ▶ 2015: ø2.1%
  - ▶ 2018: ø1.6%

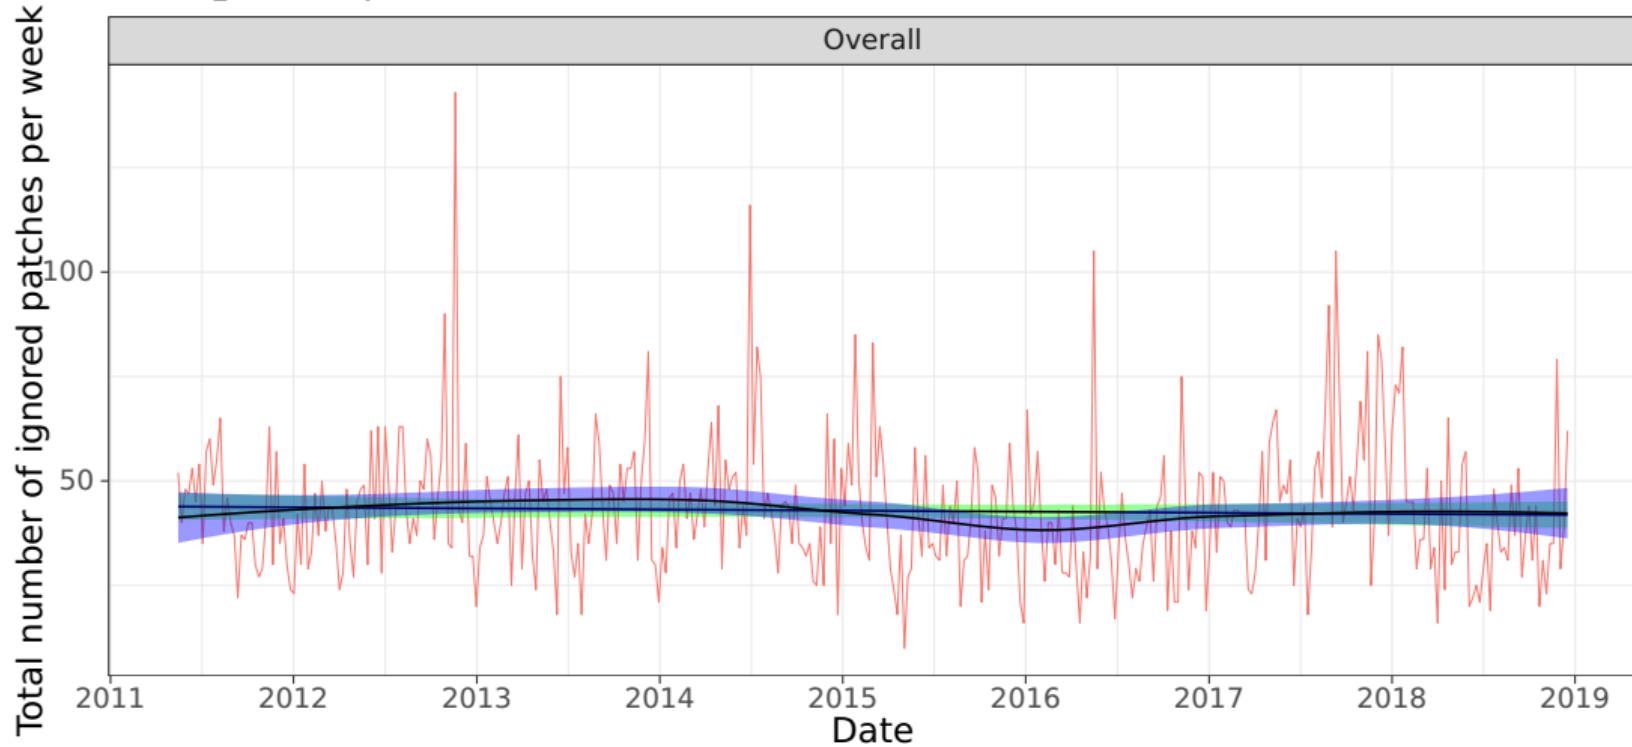
## Evolution of ignored patches



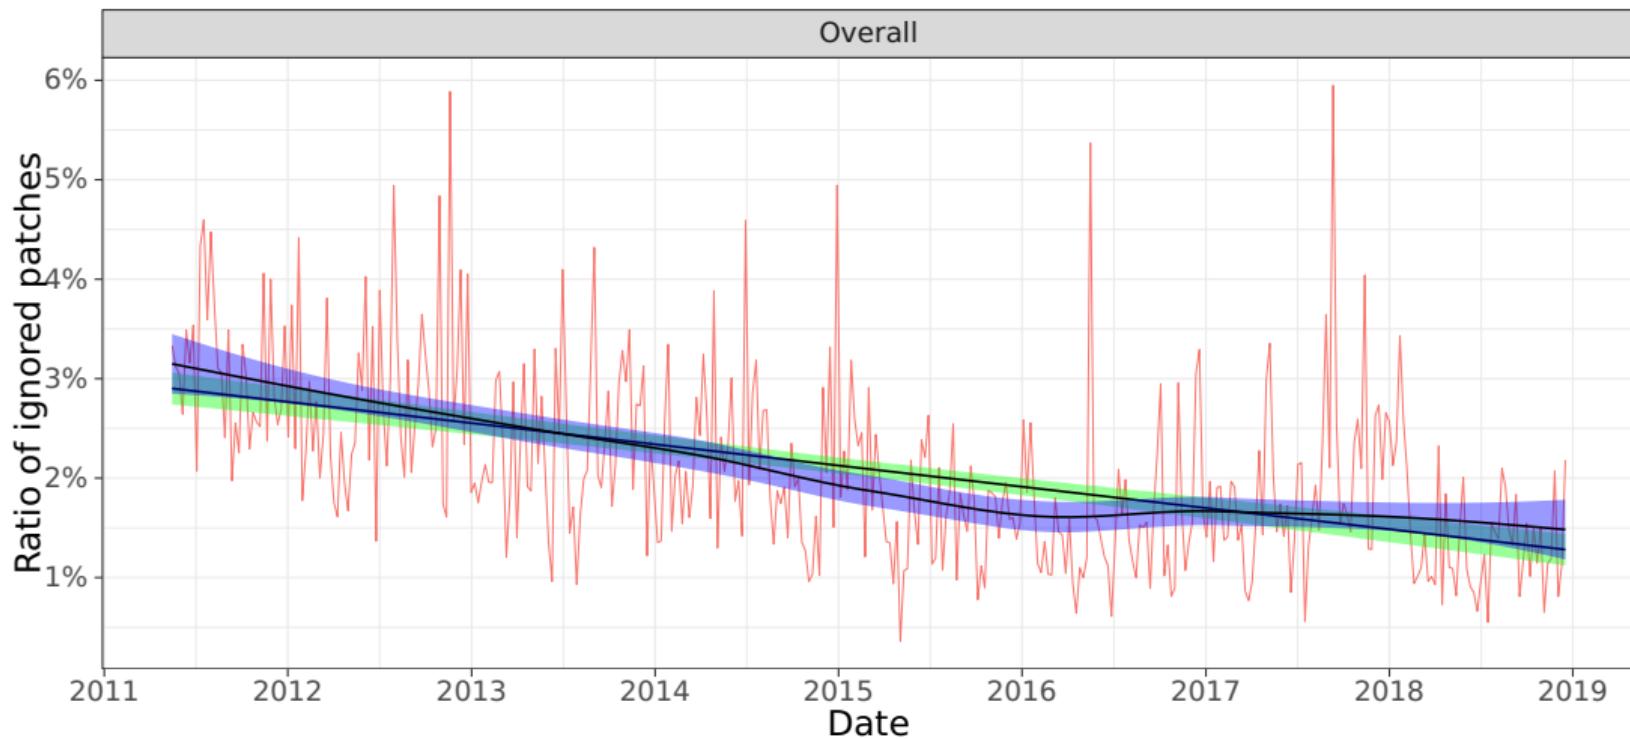
## Evolution of ignored patches



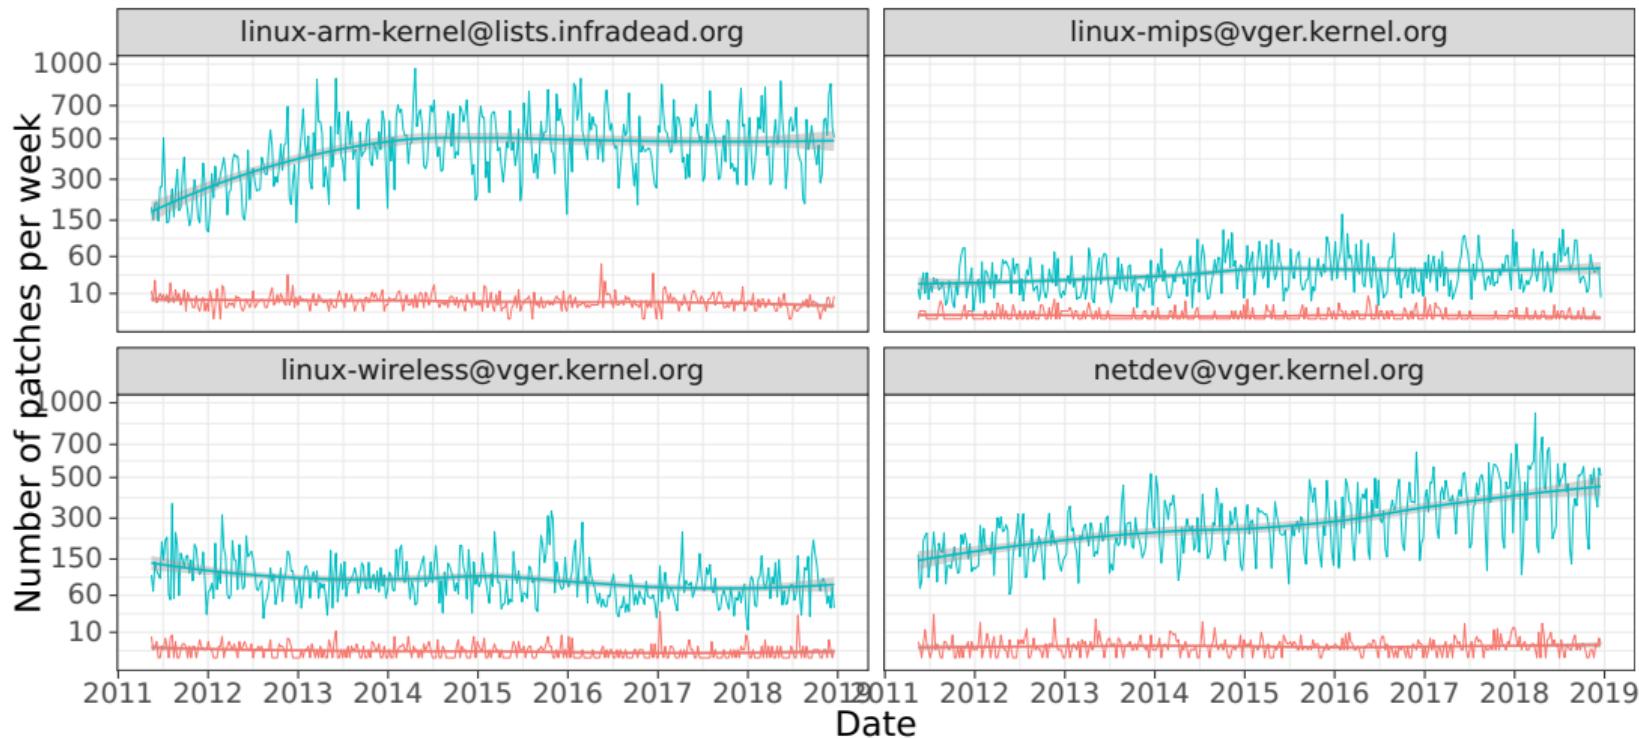
## Evolution of ignored patches



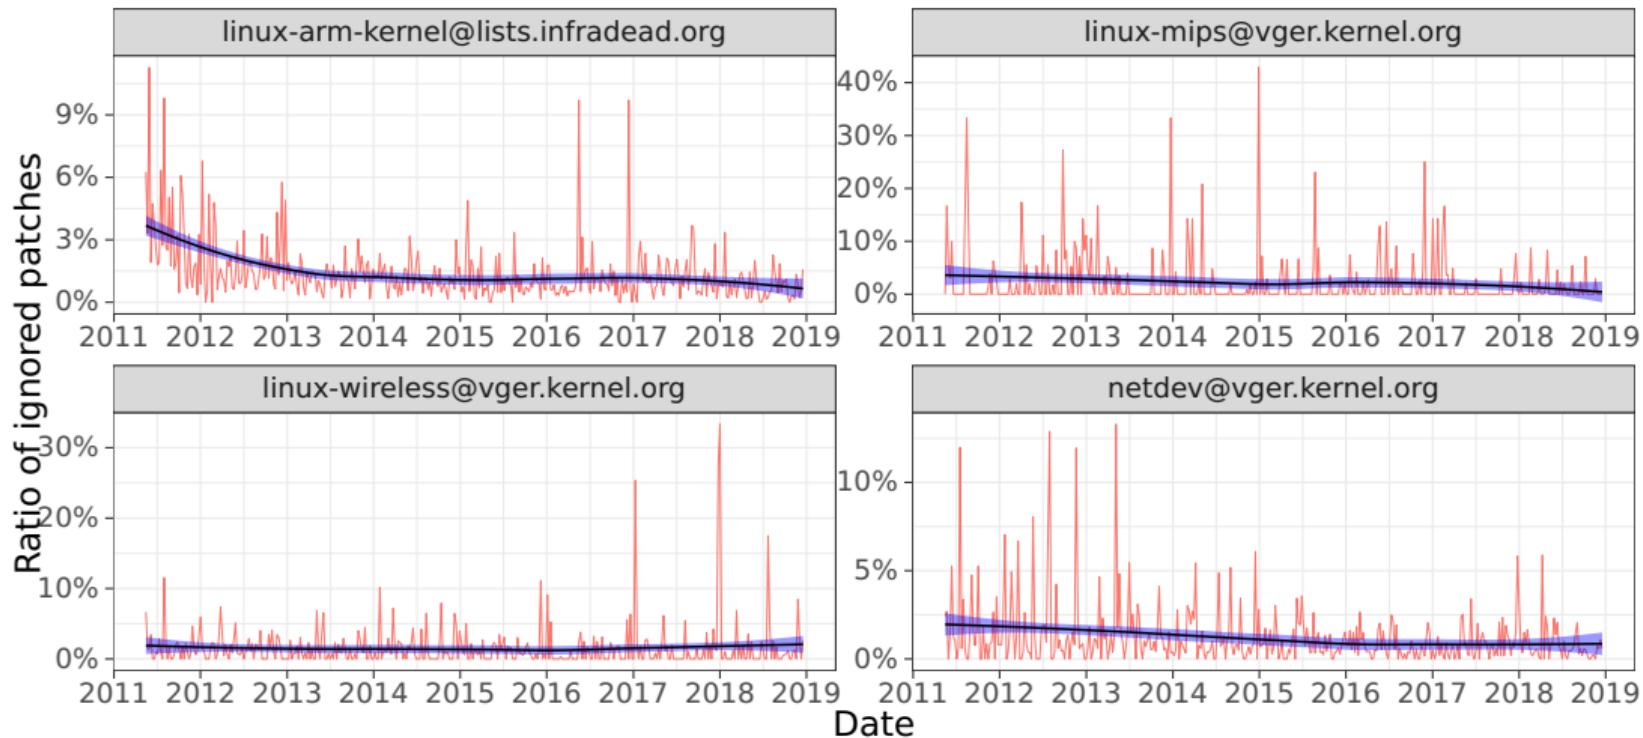
## Evolution of ignored patches



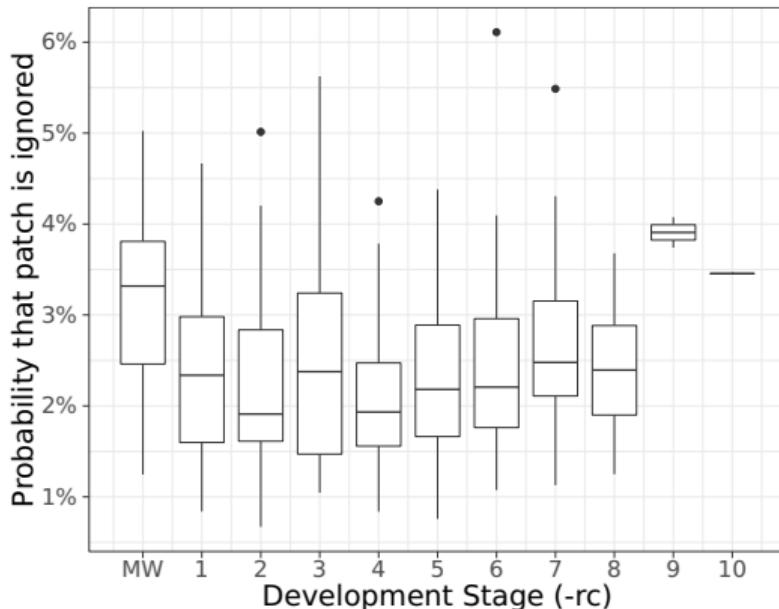
## Evolution of ignored patches



## Evolution of ignored patches



Does it matter *when* a patch is sent?



Distribution of ratio ignored/total patches grouped by linux kernel development stage

### Insights

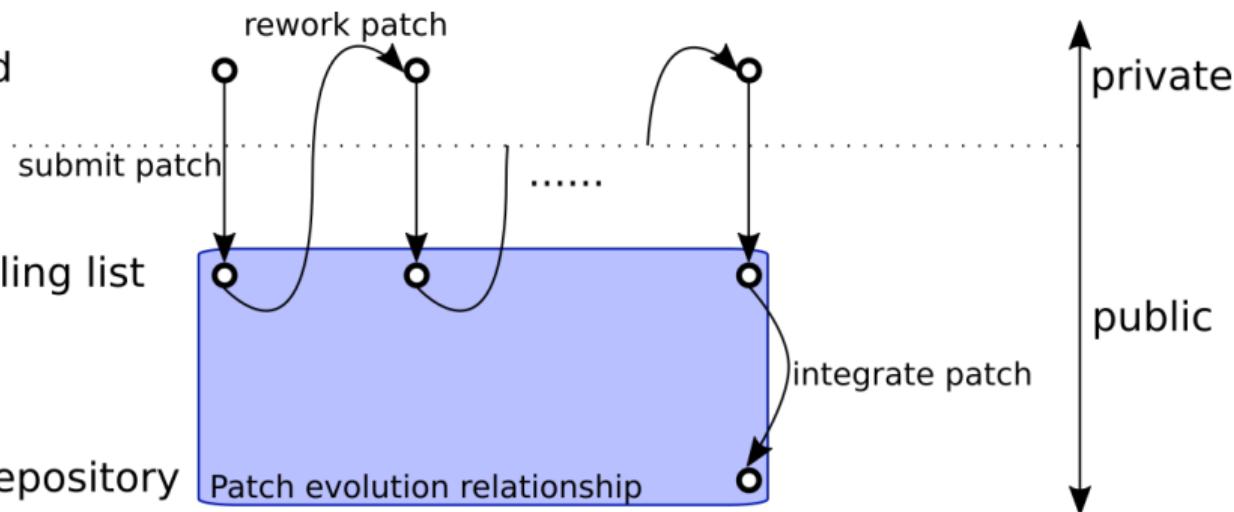
- ▶ Largely independent of the development stage
- ▶ Slightly higher chance of ignorance during merge window

## Towards a formal model of the development process

1. patch is created

2. patch is on mailing list

3. patch is in git repository



## Off-list Patches

### Definition

An *off-list patch* is a patch that...

- ▶ ...has been included in Linus' git repository
- ▶ ...has never been sent to any public mailing list

## Off-list Patches

### Definition

An *off-list patch* is a patch that...

- ▶ ...has been included in Linus' git repository
- ▶ ...has never been sent to any public mailing list

### Results

- ▶ Identified 80 commits with PaStA heuristics from v5.1-rc1..v5.1 ( $\approx 1800$  commits)
- ▶ Manually assessed 60 commits and identified 24 off-list patch commits

## Off-list Patches

### The obvious

- ▶ Reverting patches is discussed on mailing list, the reverting patch is not sent.
- ▶ Very few patches from maintainers are actually off-list patches

### The less obvious

- ▶ Some off-list patches are clearly some security-related issues
- ▶ Patches from some subsystem maintainers are often off-list

```
commit c7084edc3f6d67750f50d4183134c4fb5712a5c8
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Fri Apr 5 15:39:26 2019 +0200

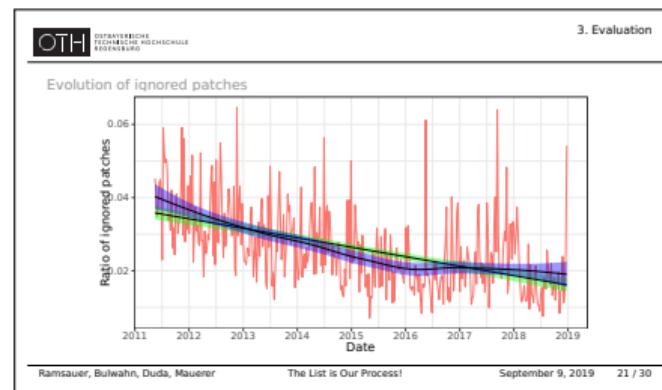
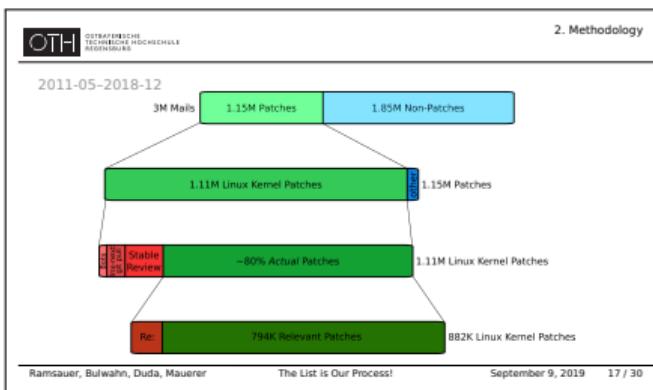
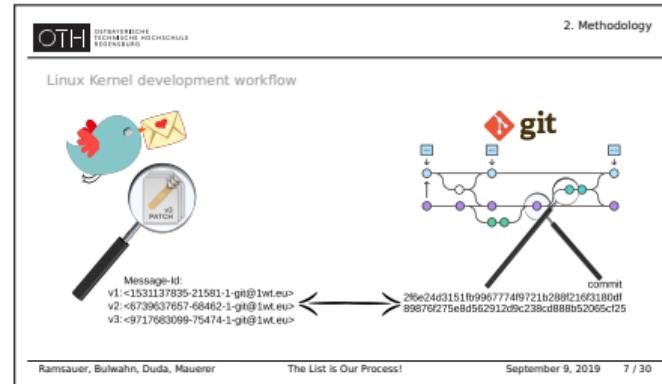
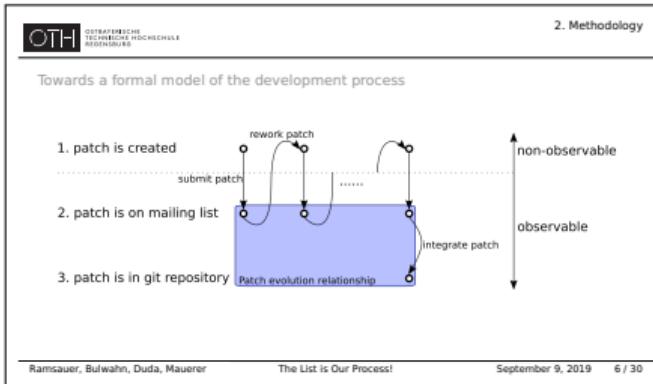
tty: mark Siemens R3964 line discipline as BROKEN
```

The n\_r3964 line discipline driver was written in a different time, when SMP machines were rare, and users were trusted to do the right thing. Since then, the world has moved on but not this code, it has stayed rooted in the past with its lovely hand-crafted list structures and loads of "interesting" race conditions all over the place.

After attempting to clean up most of the issues, I just gave up and am now marking the driver as BROKEN so that hopefully someone who has this hardware will show up out of the woodwork (I know you are out there!) and will help with debugging a raft of changes that I had laying around for the code, but was too afraid to commit as odds are they would break things.

Many thanks to Jann and Linus for pointing out the initial problems in this codebase, as well as many reviews of my attempts to fix the issues. It was a case of whack-a-mole, and as you can see, the mole won.

Reported-by: Jann Horn <jannh@google.com>  
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>  
Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>



# Thank you!

{ralf.ramsauer,wolfgang.mauerer}@othr.de

sebastian.duda@fau.de

lukas.bulwahn@gmail.com