

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

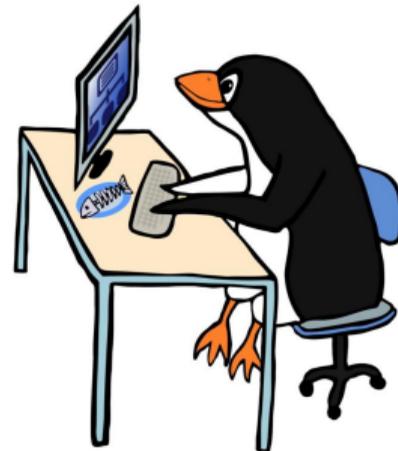
§ Siemens AG, Corporate Research and Technology, Munich

Linux Plumbers Conference, Lisbon

September 9, 2019

Our Overall Goal

Formalising and assessing the Linux
Kernel development process



© Moini
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 *outside* the community

Academic Research

- ▶ Applying statistical (big data, ML, etc.) methods on software development projects
- ▶ Linux kernel serves as a prime example for their evaluation

Commercial Users

- ▶ Measuring progress of certain long-term kernel developments, e.g., mainlining the PREEMPT_RT patch
- ▶ Measuring and monitoring many business-critical open-source projects a scale (CHAOSS)
- ▶ Fulfilling the certification requirement "Development process assessment" in regulated environments (safety- and security-related certified systems)

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 Linux Kernel Maintainer Handbook

Dan Williams

Linux Plumbers 2018

Subsystem Profile

- Describe maintenance of your sub-system as a “bus-factor” document
- Document the policies that each subsystem typically decides differently

Patches or Pull requests
Last day for new feature submissions
Last day to merge features
Non-author Ack / Review Tags Required
Test Suite

Trusted Reviewers
Resubmit Cadence
Time Zone / Office Hours
Checkpatch / Style cleanups
Off-list review

Others?

Subsystem Profile

- Describe maintenance of your sub-system as a “bus-factor” document
- Document the policies that each subsystem typical decides differently

Patches or Pull requests
Last day for new feature submissions
Last day to merge features
Non-author Ack / Review Tags Required
Test Suite

Trusted Reviewers
Resubmit Cadence
Time Zone / Office Hours
Checkpatch / Style cleanups
Off-list review

Others?

Subsystem Profile

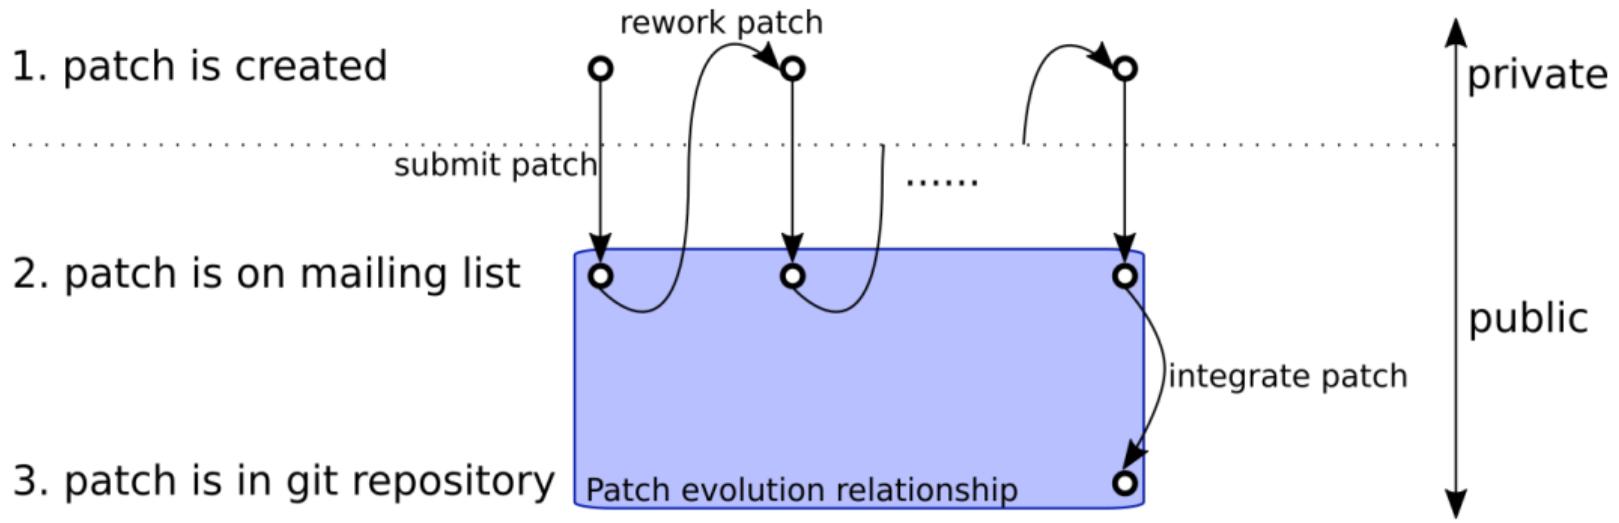
- Describe maintenance of your sub-system as a “bus-factor” document
- Document the policies that each subsystem typically decides differently

Patches or Pull requests
Last day for new feature submissions
Last day to merge features
Non-author Ack / Review Tags Required
Test Suite

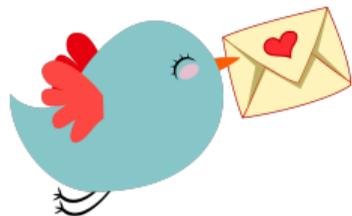
Trusted Reviewers
Resubmit Cadence
Time Zone / Office Hours
Checkpatch / Style cleanups
Off-list review

Others?

Towards a formal model of the development process



Linux Kernel development workflow



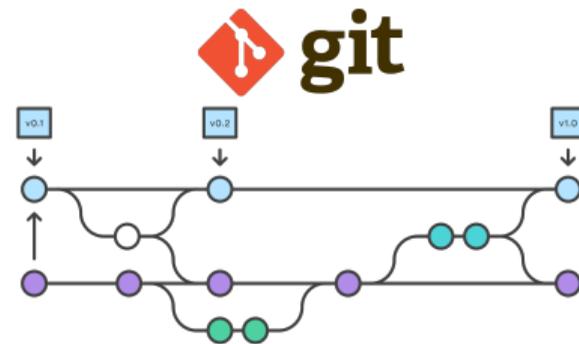
Linux Kernel development workflow



Linux Kernel development workflow



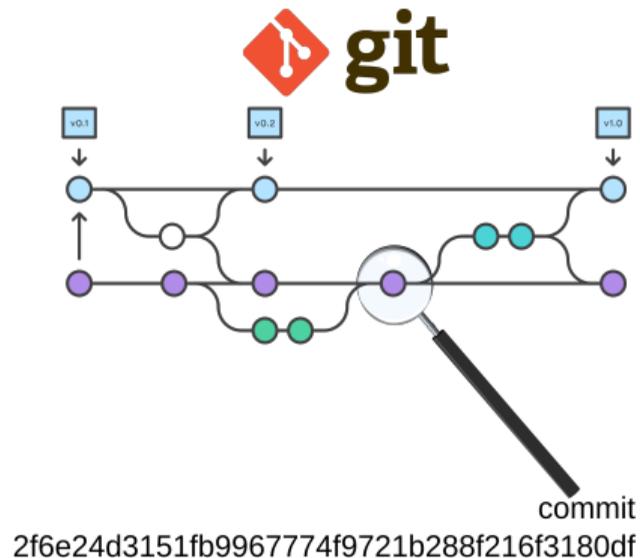
Linux Kernel development workflow



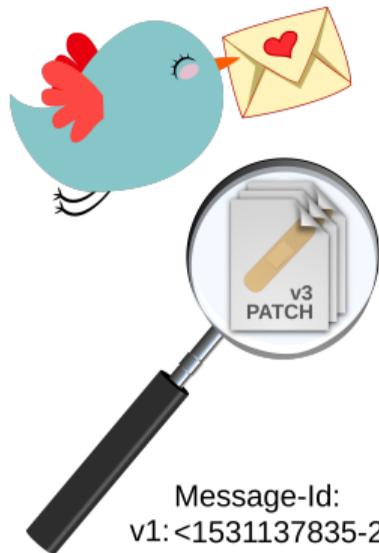
Linux Kernel development workflow



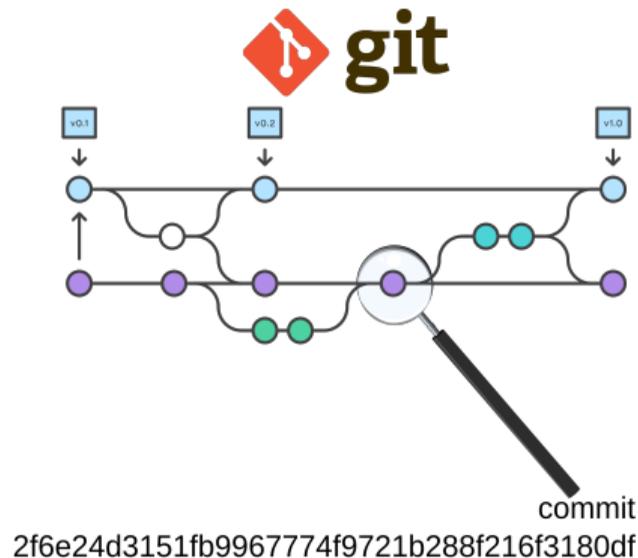
Linux Kernel development workflow



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>

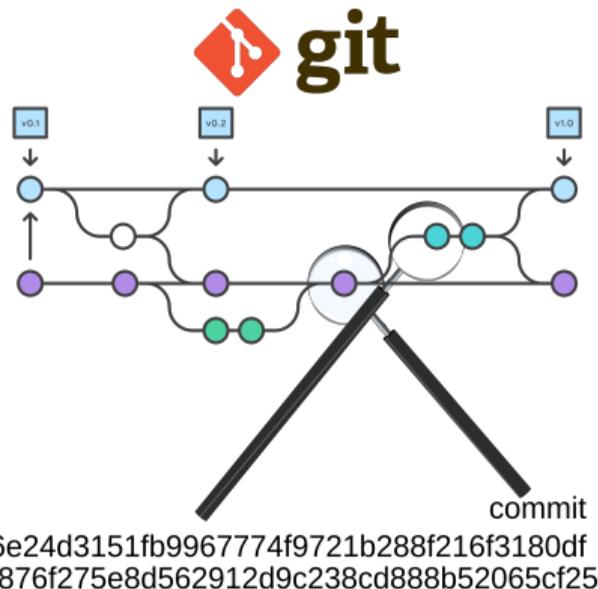


Linux Kernel development workflow

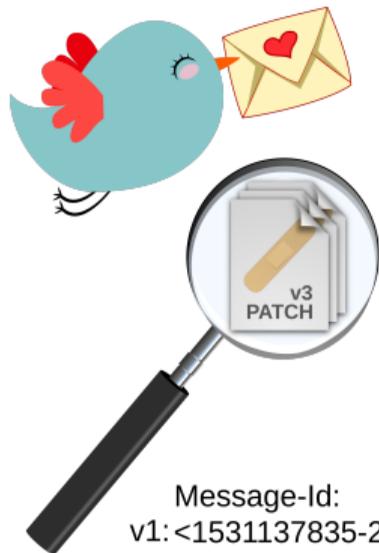


Message-Id:

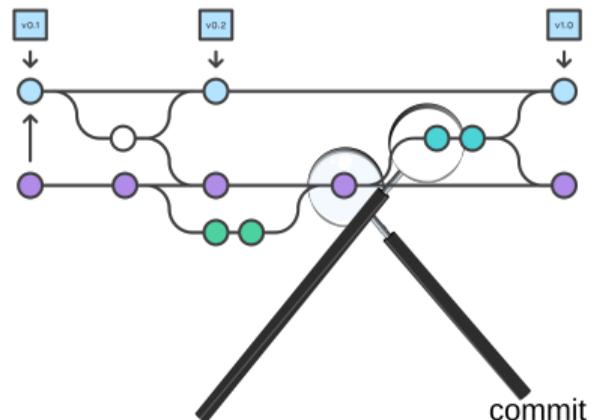
v1: <1531137835-21581-1-git@1wt.eu>
v2: <6739637657-68462-1-git@1wt.eu>
v3: <9717683099-75474-1-git@1wt.eu>



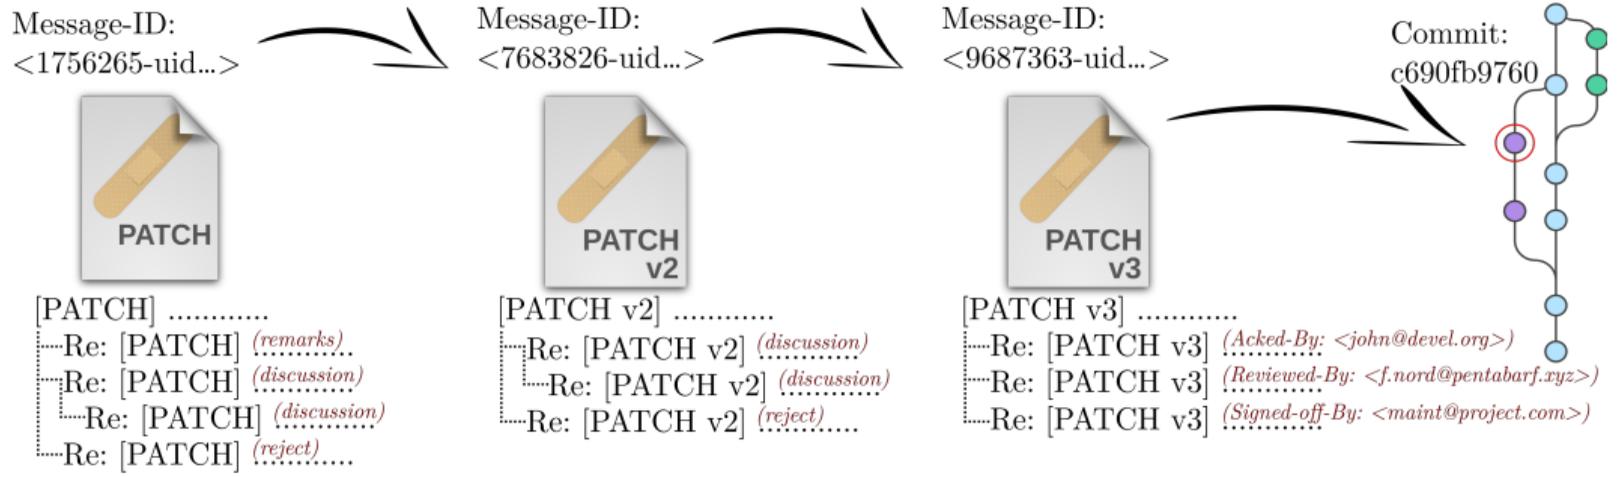
Linux Kernel development workflow



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



commit
 2f6e24d3151fb9967774f9721b288f216f3180df
 89876f275e8d562912d9c238cd888b52065cf25



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

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-3b30p18b2 [...]
```

```
Signed-off-by: Ingo Molnar <mingo@elte.hu>
```

```
diff --git a/kernel/sched.c b/kernel/sched.c
index 1874c74..4cdc91c 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.patrch

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 12 11:15:15 2011 +0200

sched-cure: Fix idle_cpu()

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

diff --git a/kernel/sched.c
index 2054c30..18b218b2 100644
--- a/kernel/sched.c
+++ b/kernel/sched.c
@@ -5037,7 +5037,11 @@
 */
int idle_cpu(int cpu)
{
-   return rq->curr == rq->idle && !rq->nr_running;
+   return rq->curr == rq->idle && !rq->nr_running && !rq->wake_list;
}

/**
```



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

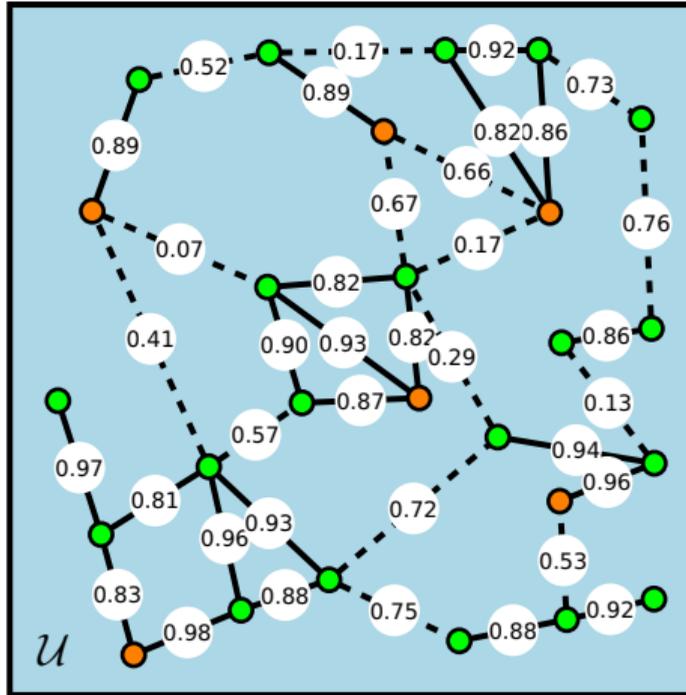
sched: Fix idle_cpu()

On -rt we observed hackbench waking all 400 tasks to a single cpu. This is caused by the interaction of the following patch with the following snippet:

```
[...]
#endif
#ifdef CONFIG_SMP
if (l1ist_empty(&rq->wake_list))
return 0;
else
return 1;
#endif
}

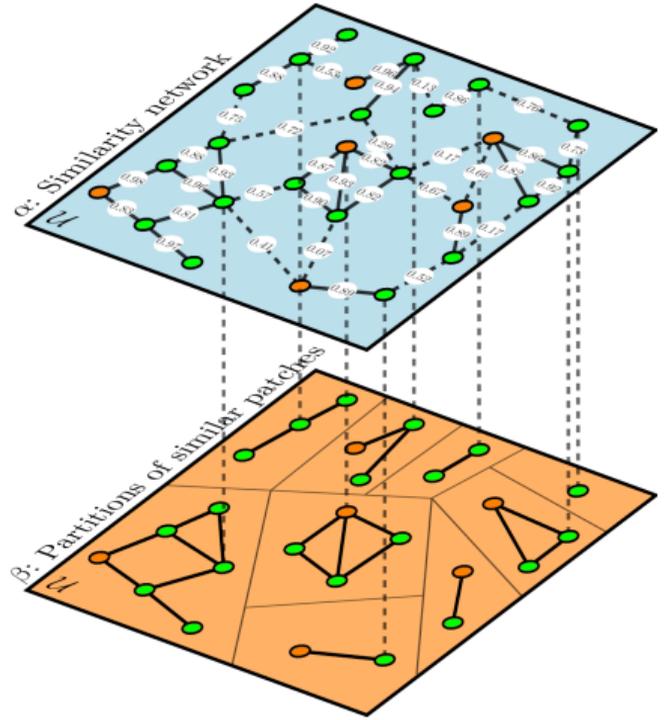
/**
```





Legend

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



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 ListsRalf Ramsauer¹, Daniel Lohmann² and Wolfgang Mauerer³¹Technical University of Applied Sciences Regensburg
University of Bamberg²Siemens AG, Corporate Technology, Munich
r.lohmann@wtr.de, lohmann@ira.uni-bamberg.de, wolfgang.mauerer@othr.de

Abstract—A considerable amount of research on software evolution focuses on making changes in software repositories, but only their pre-integration history.

We present a novel method for tracking this otherwise invisible evolution of software changes on mailing lists by connecting all early instances of changes to their final version repositories. Since artifact modifications on mailing lists are communicated by updates to fragments (i.e., patches) only, identified semantically similar changes in a non-trivial task that user supports solve in a language-independent way. We evaluate our method on high-profile open source software (OSS) projects like the Linux kernel, and validate its high accuracy using an elaborately created ground truth.

Our approach can be used to quantify properties of OSS development processes, which is an essential requirement for using OSS as reliable or safety-critical industrial products, where confidentiality and confidentiality in processes are critical. The high accuracy of our technique allows, to the best of our knowledge, for the first time to quantitatively determine if an open development process effectively aligns with given formal process requirements.

I. INTRODUCTION

Software patches may have come a long way before their final integration into the official branch (known as merge or push) of a project. There are many possible ways of integration. Amongst others, the origin of a patch can be a merge from other developers' repositories (i.e., repositories), or patches from foreign repositories, pull requests on web-based repository managers such as GitHub or GitLab, vendor specific patch stacks, or mailing lists (MLs).

Especially MLs have been in use for software development processes for decades [1]. They have a well-known interface (plain text emails), and come with an elaborate set of text requirements (i.e., a mail user agent). Because of their simplicity, scalability, 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 software systems (e.g., QEMU, Linux, Glibc, etc.), operating

¹This work was supported by Siemens AG, Corporate Research, the German Research Foundation (DFG) under the Special Collaborative Program (SFB 1174/B1). The eleven project has received funding from the ERC under grant agreement 101019067 (EU FP7 grant no. 30344). The 10 authors support their funder's General User's Review 2020 research and innovation programme. It is not funded by the participating members, except from Austria, Germany, Belgium, Italy, Spain and Romania.

systems (e.g., the Linux kernel) or foundations (e.g., Apache, GNU). Mailing lists form the backbone of their development processes [2]. They are not only used to ask questions, file bug reports or discuss general topics, but implement a patch submit-review-improve strategy for software refinement [4] that is typically iterated multiple times before a patch is finally integrated to the repository (cf. 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 outcome of that process, while all intermediate steps leave indirect traces in the repository. Mailing lists 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-technical [12, 22, 40] aspects of software development. This will be possible because MLs contain information on interactions between developers.

Notably, open source components are routinely deployed in industrial fields, and their use is increasingly explored in safety-critical or mission-critical appliances [34], such as medical devices or in automotive products. Especially for core components of a system that implement business-wide differentiating features such as the system software stack or middleware, OSS provides adequate solutions that have already proved to be reliable in other non-critical applications domains. However, non-functional aspects like evidence of quality assurance are also a crucial factor for industry. Deployment of software in safety-critical environments requires conformity with international standards, such as ISO 26262 [29], IEC 61508 [24] or IEC 62304 [23]. This demands certified development processes that implement high standards regarding traceability and reliability of all development decisions, including code writing, reviewing, approval, and maintenance activities like the rationale for which process compliance is to achieve and prove high product quality.

Compared to conventional, orthodox proprietary industrial software, OSS exhibits different dynamics [35], and often requires fundamentally different development processes [15] because of project size and a high number of autonomously governed stakeholders. Because of this nature, OSS projects do not necessarily meet certification criteria [33].

Nevertheless, reviews across different industrial sectors have similar concerns on the use of OSS components [18, 19]

Observing Custom Software Modifications: A Quantitative
Approach of Tracking the Evolution of Patch StacksRalf Ramsauer¹
Technical University of Applied
Sciences Regensburg
ralf.ramsauer@othr.deDaniel Lohmann²
Friedrich-Alexander-University
Erlangen-Nuremberg
lohmann@cs.fu.deWolfgang Mauerer³
Technical University of Applied
Sciences Regensburg
Siemens AG, Munich
wolfgang.mauerer@othr.de

ABSTRACT

Modifications to open-source software (OSS) are often provided in the form of "patch stacks" sets of changes (patches) that modify a given body of source code. Maintaining patch stacks over extended periods of time is problematic since the underlying base project changes frequently. This necessitates a continuous and engineering-intensive adaptation of the stack. Nonetheless, long-term maintenance is an important problem for changes that are not integrated into projects, for instance when they are controversial or only of value to a limited group of users.

We present and implement a methodology to systematically examine the temporal evolution of patch stacks, track non-functional properties like integrability and maintainability, and estimate the overall economic and engineering effort required to successfully develop and maintain patch stacks. Our results provide a basis for quantitative research on patch stacks, including statistical analyses and other methods that lead to actionable advice on the quantitative and long-term maintenance of custom extensions to OSS.

I. INTRODUCTION

Special-purpose software, like embedded control, medical analysis, or other domain-specific applications, is often composed of contributions from general-purpose projects that provide basic building blocks. Because of their widespread use on top of them, they fulfill certain additional requirements, while the development of new ones, the primary benefit of the base project, proceeds independently.

Especially for software with high dependability requirements, it is crucial to keep up to date with maintainers' latest fixes that can be applied and any general features have to be introduced, as diverging software branches are hard to maintain and hard to subsume into systems [6]. Possible drawbacks often surface in the form of patch stacks: lower-granular modifications include the Forward IT Linux mobile extension, exhibited by Android software projects, maintaining patch

Permitted to make digital or hard copies of all or part of this work for personal or classroom use provided the original author and source are credited. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

© 2018 Copyright held by the owner(s). Publication rights licensed to ACM.
10.1145/3211111.3207701.201811

stacks can become a significant issue in terms of effort and costs.

Our toolset PatchStack (Patch Stack Analysis) quantitatively analyzes the evolution of patch stacks by mining Git [3] repositories and produces data that can serve as input for statistical analysis. It recognizes different releases of stacks and groups similar patches (patches that lead to similar modifications) into equivalence classes. This allows us to compare those classes against the base project to measure integrability and influence of the patch stack on the base project. Patches that remain on the external stack across releases are classified as persistent and are highlighted to reflect the maintenance cost of the whole stack. A fine-grained classification of different patch types that depends on the actual modification code function as a measure for the successiveness of the stack.

In summary, we claim 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 PongMail [36], a real-time extension to the Linux kernel that enjoys widespread use in industrial applications for more than a decade, yet has not been integrated into standard Linux. We analyze its influence on maintainers and visualize the development phases of the stack.

2. APPROACH

In general, a patch stack (also known as patch set) is defined as a set of patches (commits) that are developed and maintained independently of the base project. While many modifications include the Forward IT Linux mobile extension, the Linux L3SM (Long Term Support Initiative) kernel, and vendor-specific Android stacks needed to port the system to a particular hardware. In many cases, patch stacks are applied on top of individual releases of an upstream version, but they do not necessarily have to be developed in a linear way [1]. The contents of the patch set derive of a base project are identified by the set of commit hashes that do not occur in the maintainers' project.

Our analysis is based on the following assumptions:

¹https://github.com/OTH/PSA

Data Acquisition

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



Let there be chaos

- ▶ Broken encoding
- ▶ BÅse64
- ▶
- ▶ 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 (Westeuropäische
Sommerzeit)

X-Mailer: Microsoft Outlook Express 6.00.2900.3028

We analyse...

- ▶ v2.6.39..linus/master
- ▶ \approx 610K commits
- ▶ Mails: 2011-05-01 - 2018-12-31
- ▶ \approx 3M mails
- ▶ Lists: All Public Inboxes from lore.kernel.org

We analyse...

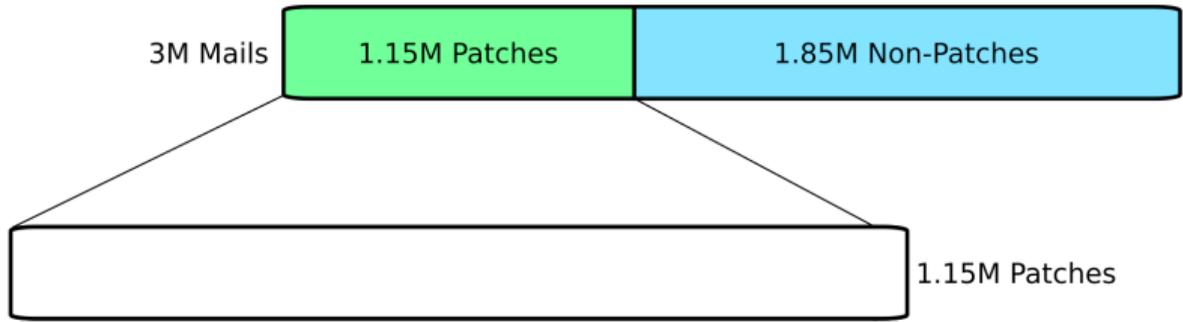
- ▶ v2.6.39..linus/master
- ▶ ≈610K commits
- ▶ Mails: 2011-05-01 - 2018-12-31
- ▶ ≈3M mails
- ▶ Lists: All Public Inboxes from lore.kernel.org
- ▶ linux-amlogic, **linux-arm-kernel**, linux-i3c, linux-mtd, linux-riscv, linuxppc-dev, cocci, linux-block, linux-bluetooth, linux-btrfs, linux-cifs, linux-clk, 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

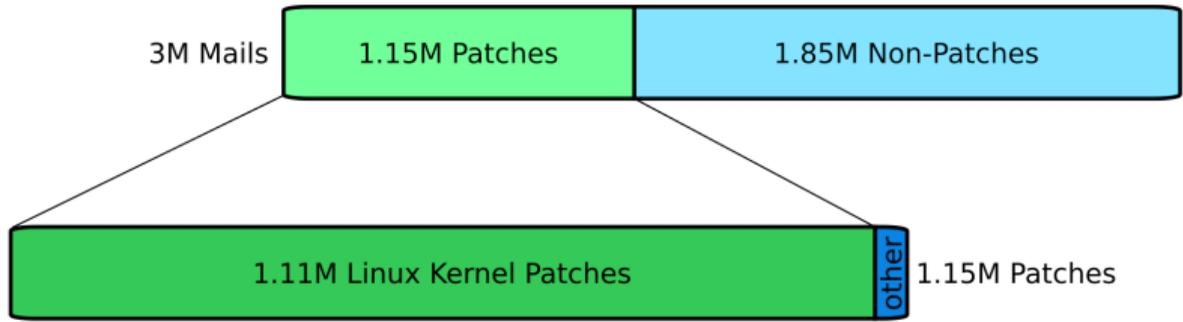
3M Mails



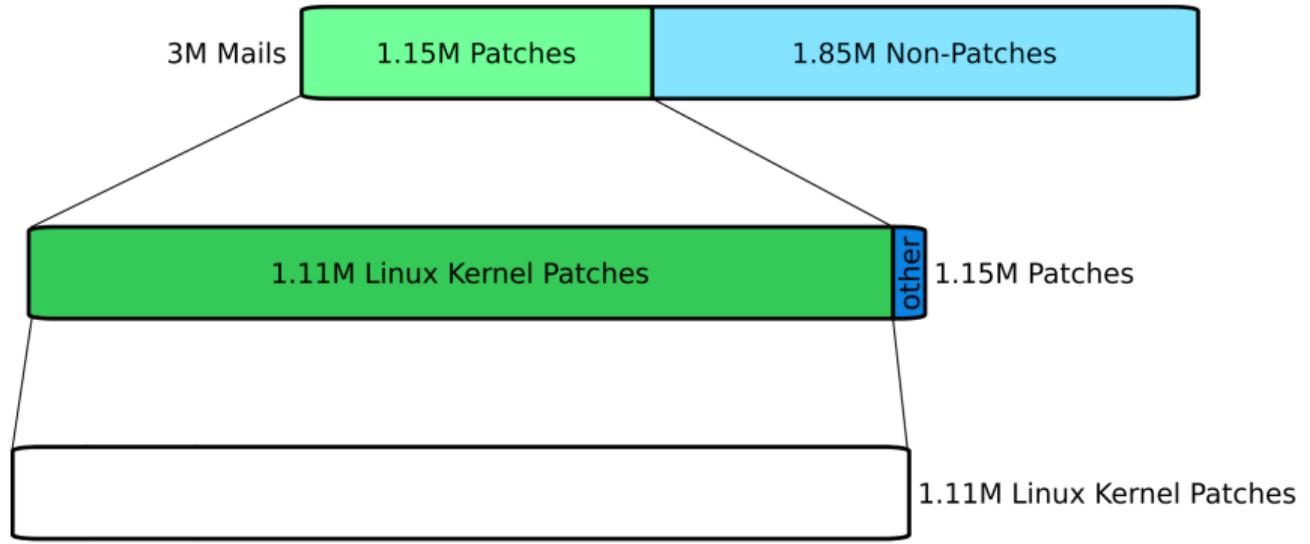
2011-05-2018-12



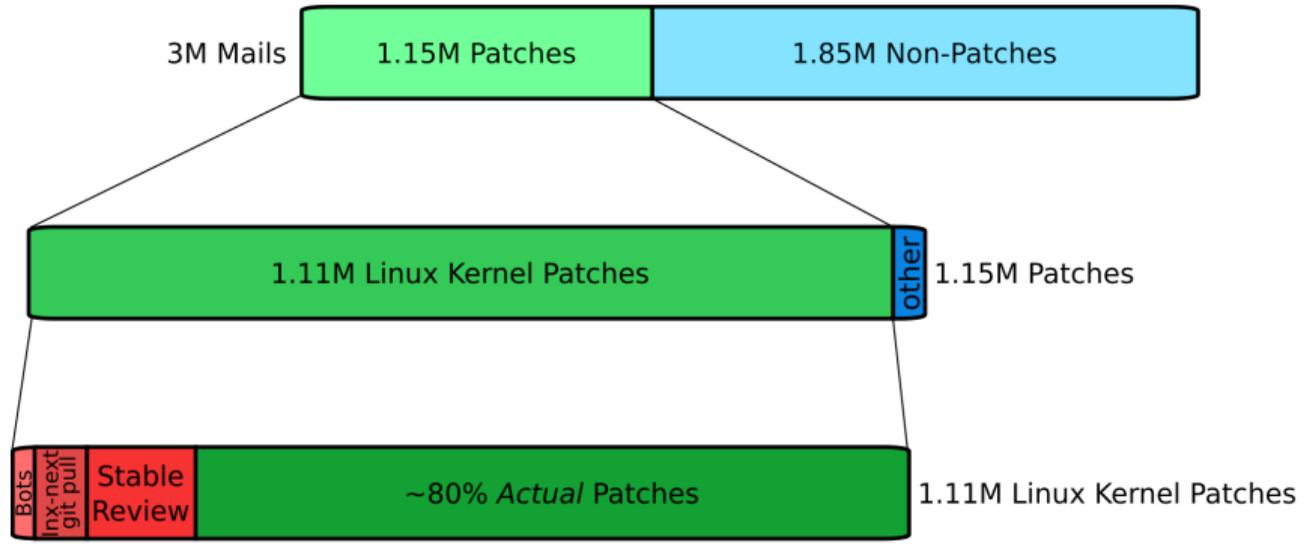
2011-05-2018-12



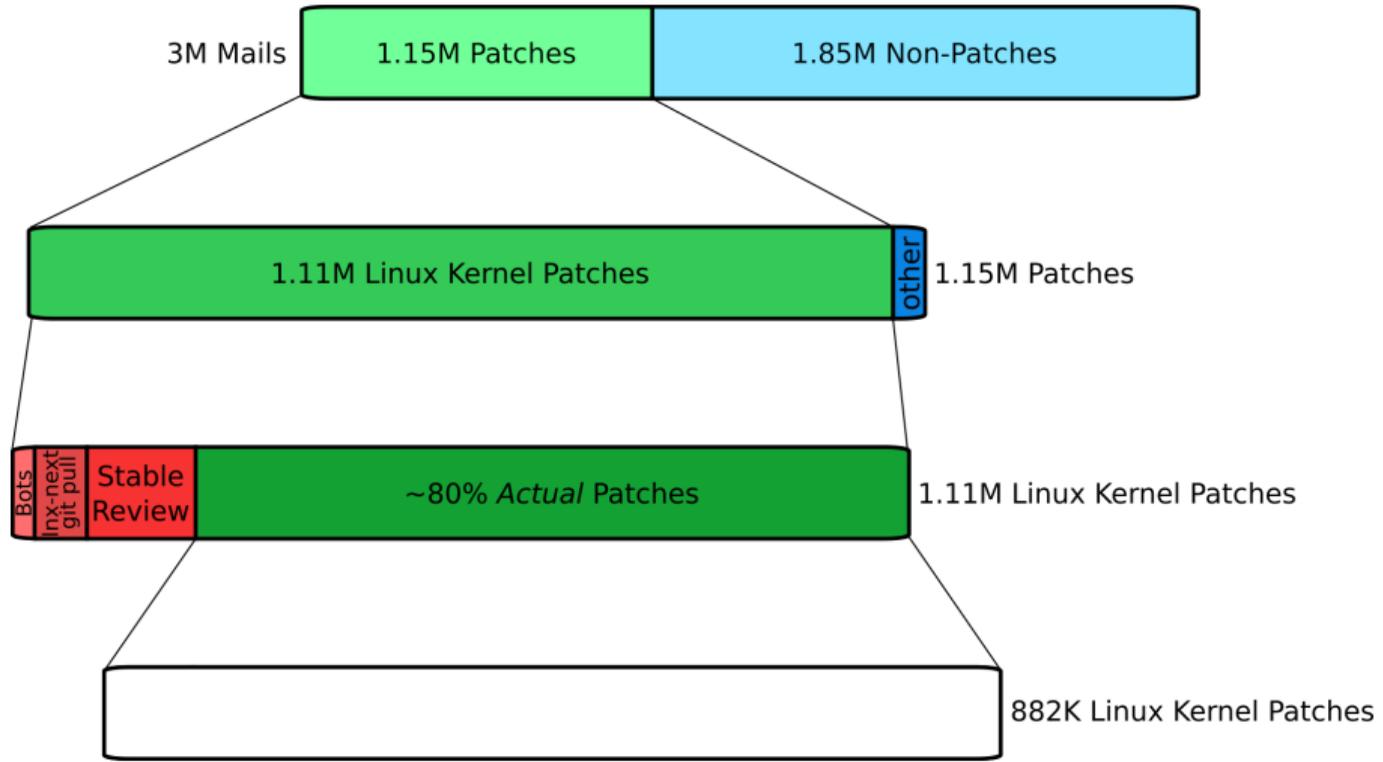
2011-05-2018-12



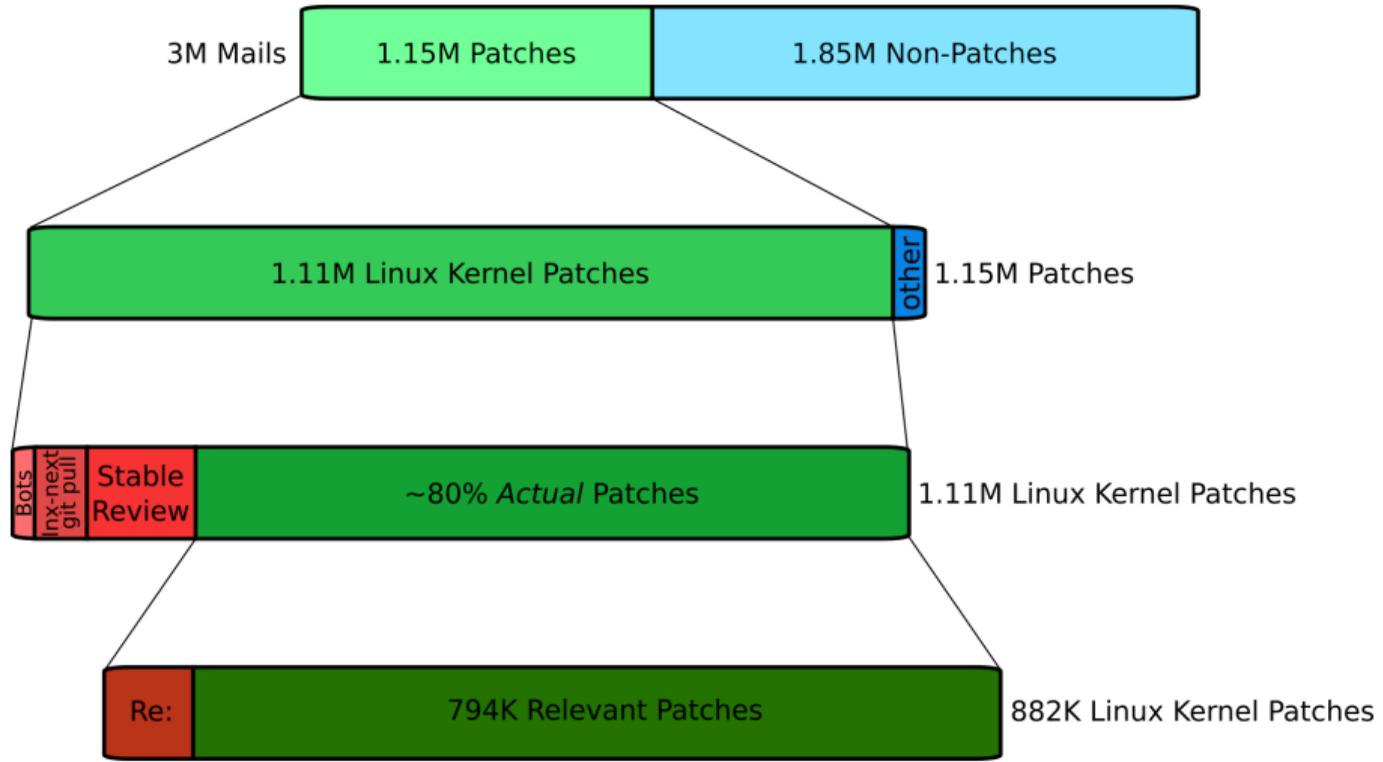
2011-05-2018-12



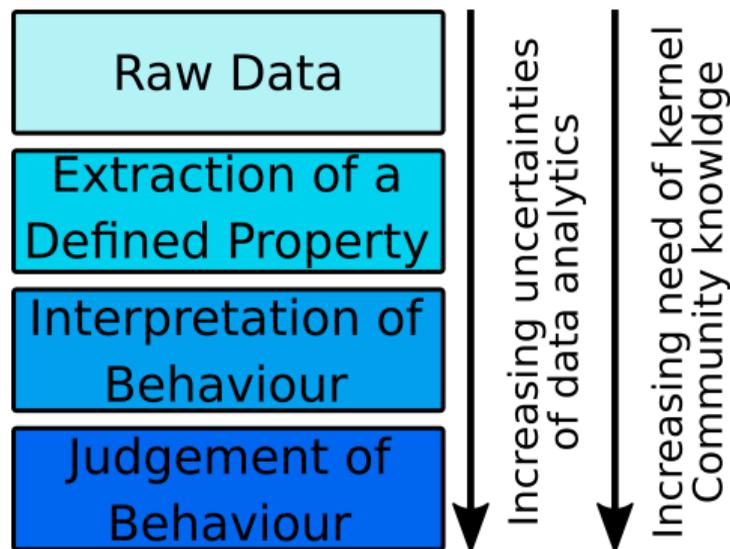
2011-05-2018-12



2011-05-2018-12



Increasing Levels of Data Aggregation



Our main work

- ▶ Focus on the first two levels and *selectively* interpret observations to *describe* behaviour
- ▶ We do **NOT** try and **NOT** intend to judge behaviour

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

Research Question

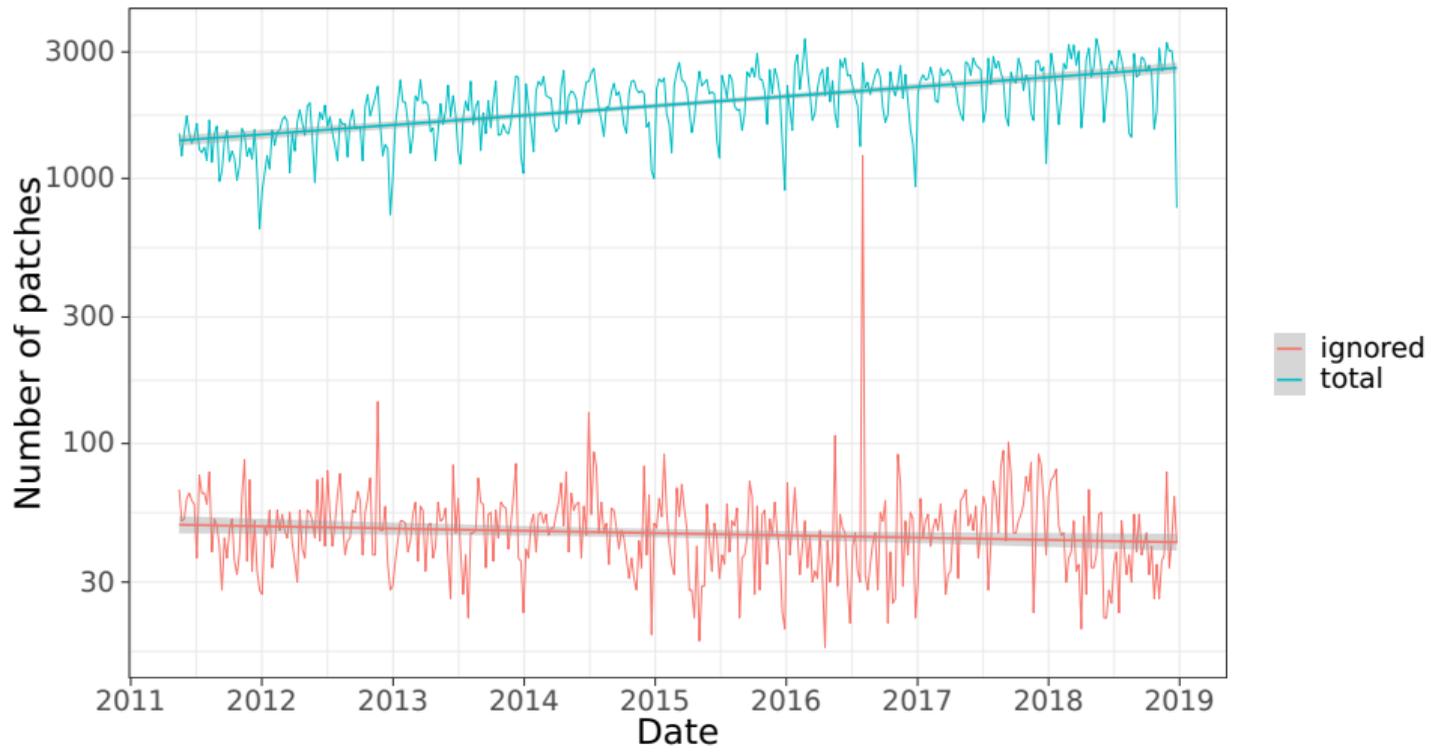
Are there specific characteristics for ignored patches?

Ignored Patches

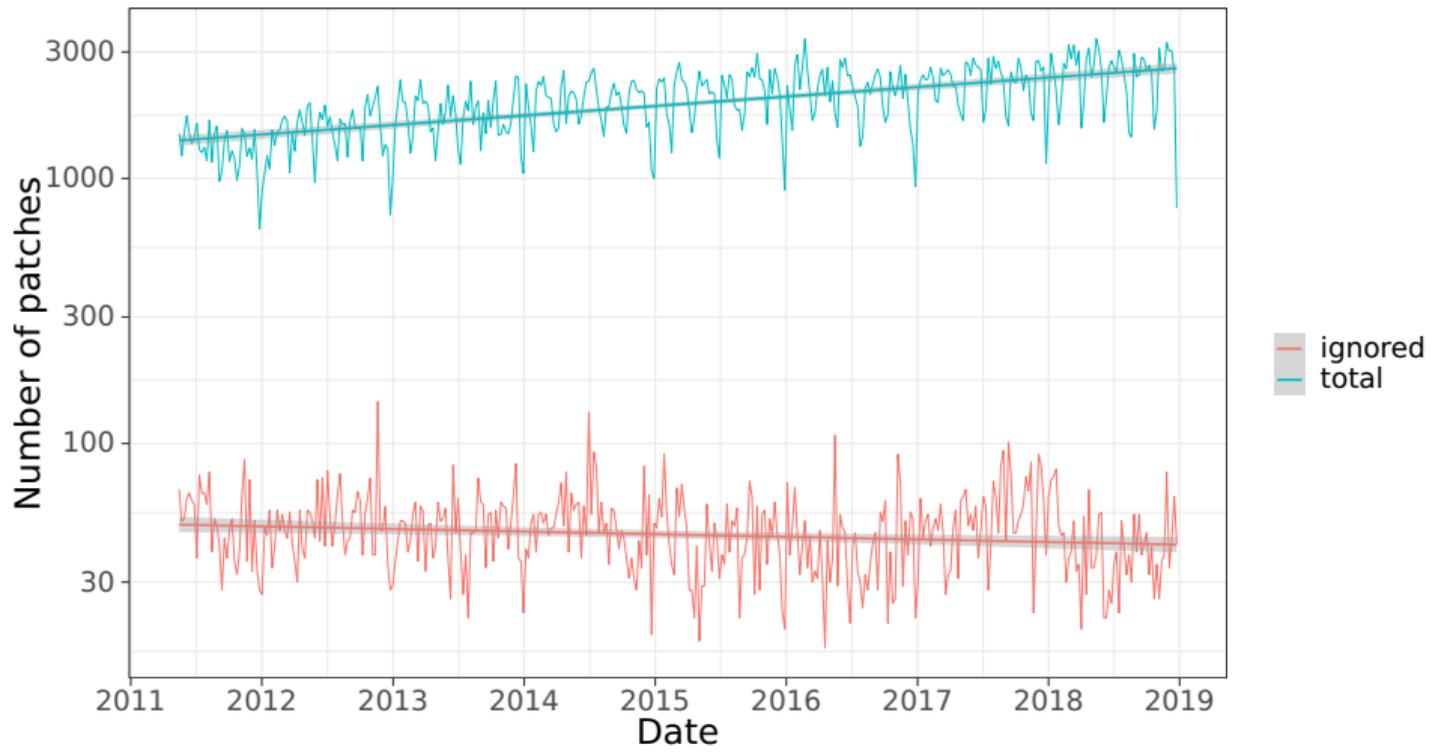
By the numbers...

- ▶ lore.kernel.org lists 2011-2018: \emptyset 2.5%
 - ▶ 2011: \emptyset 3.9%
 - ▶ 2015: \emptyset 2.1%
 - ▶ 2018: \emptyset 1.6%
- ▶ vs. ALL mailing lists: (data 2019-05 - today)
 - ▶ \emptyset 3.3%

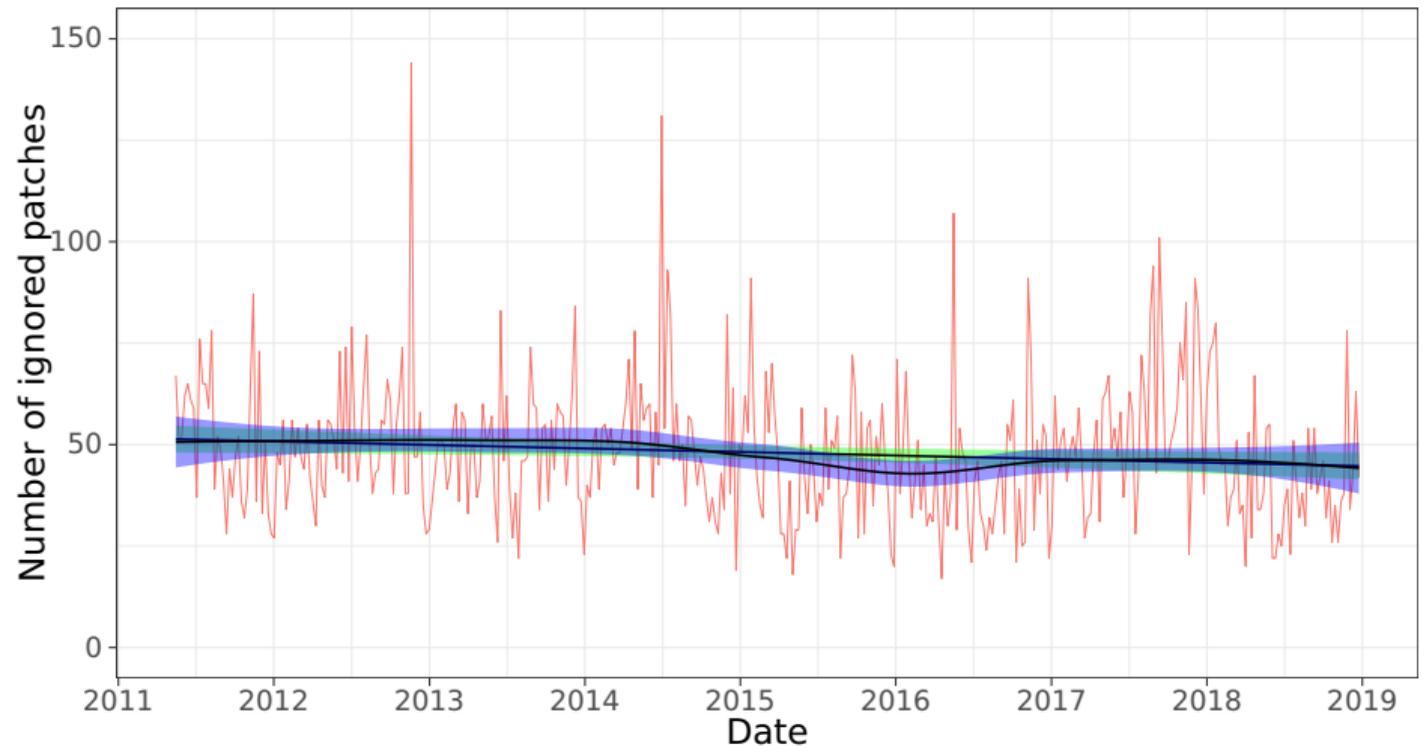
Evolution of ignored patches



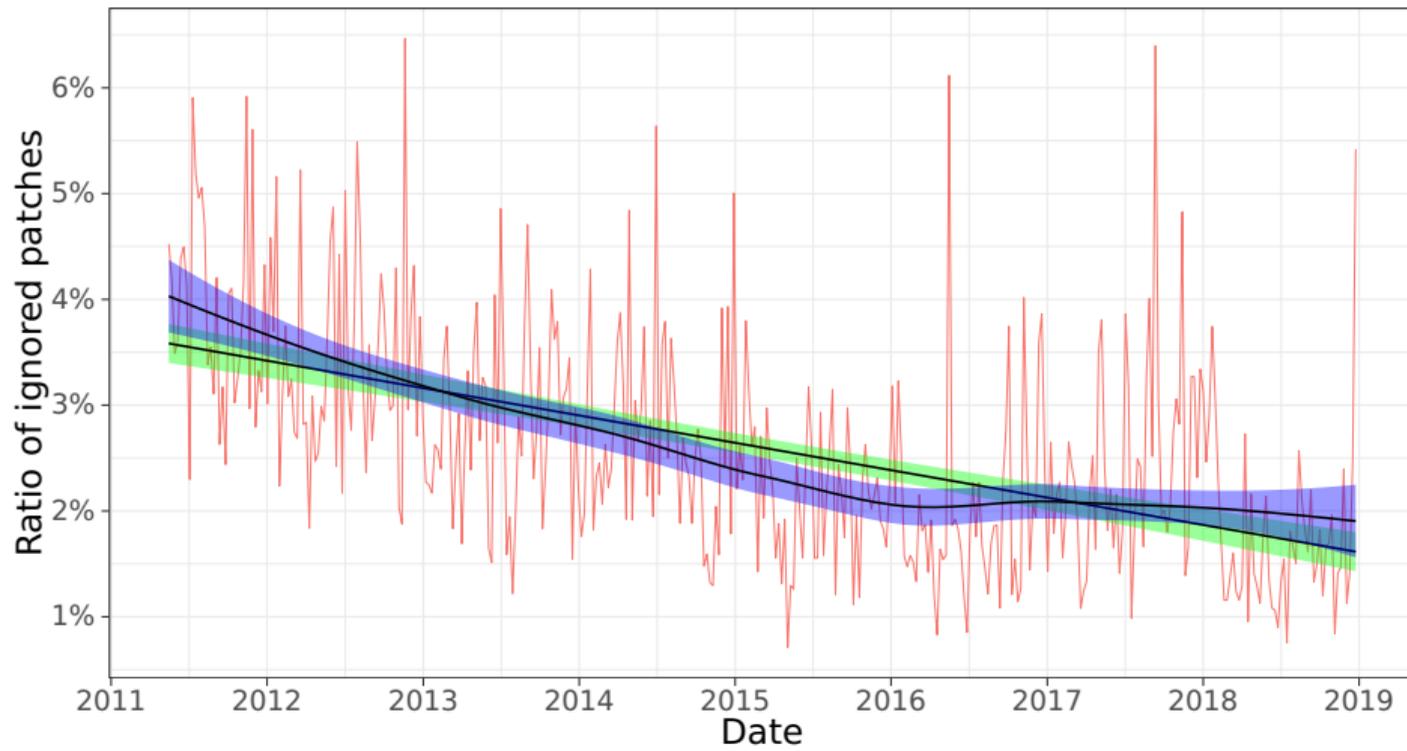
Evolution of ignored patches



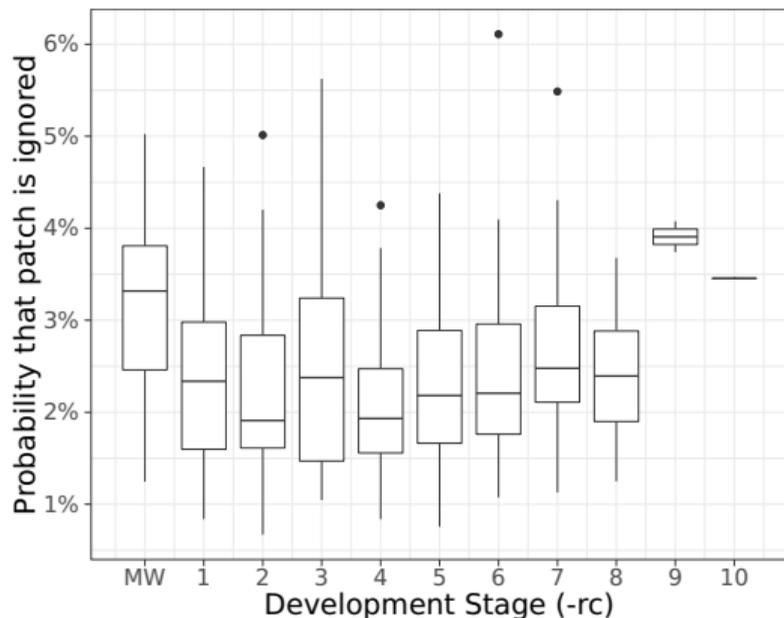
Evolution of ignored patches



Evolution of ignored patches



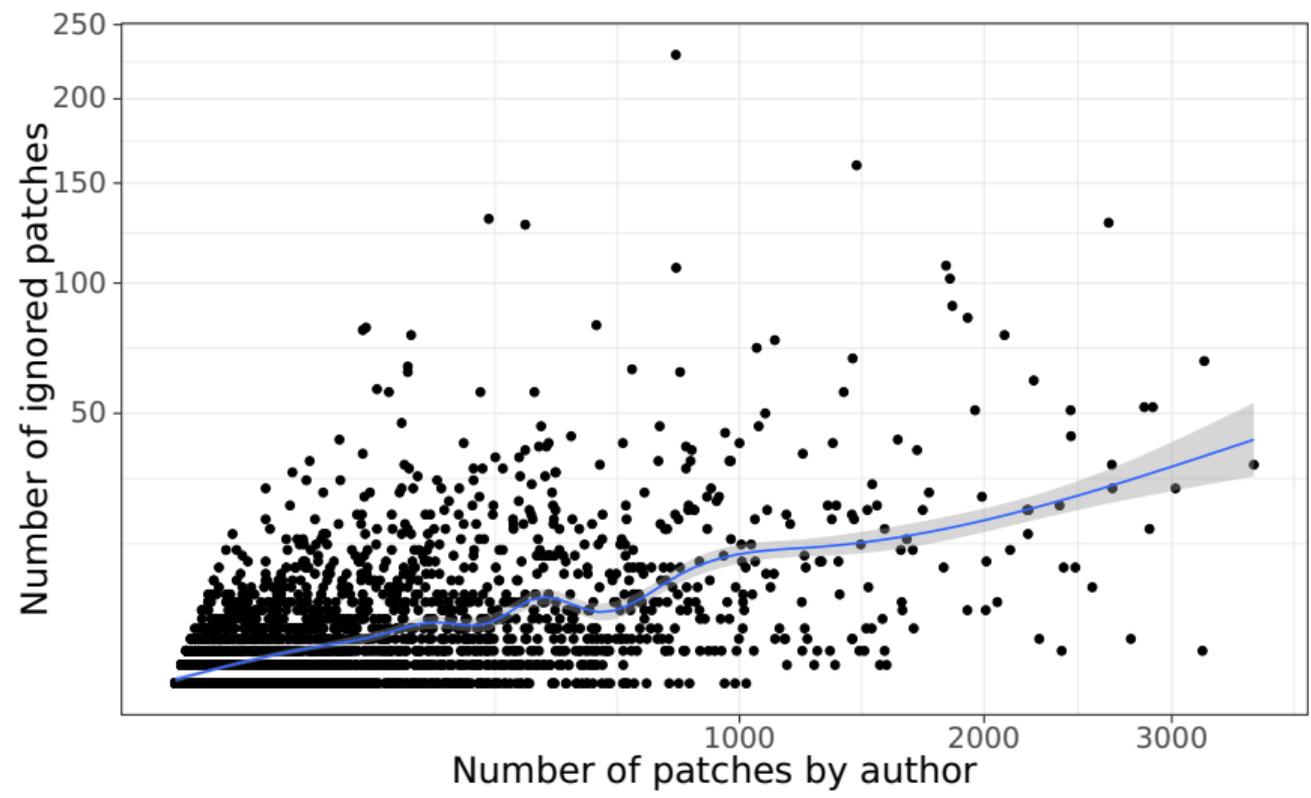
Does it matter *when* a patch is sent?



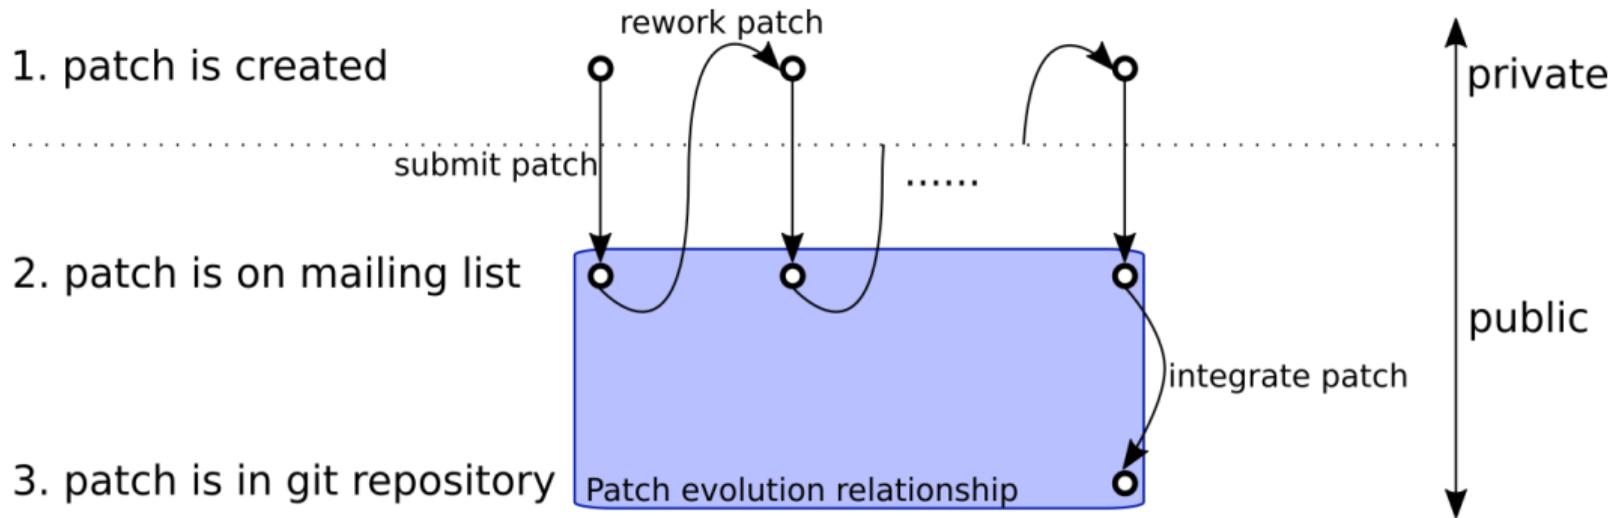
Distribution of ratio ignored/total patches grouped by linux kernel development stage

Insights

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



Towards a formal model of the development process



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 (≈ 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 <jannah@google.com>

Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>

OTH OSTBAYERISCHE TECHNISCHE HOCHSCHULE REGENSBURG 2. Methodology

Towards a formal model of the development process

1. patch is created
2. patch is on mailing list
3. patch is in git repository

Ramsauer, Bulwahn, Duda, Mauerer The List is Our Process! September 9, 2019 6 / 30

OTH OSTBAYERISCHE TECHNISCHE HOCHSCHULE REGENSBURG 2. Methodology

Linux Kernel development workflow

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

commit: 296e24d3151ff9967774f9721b2892163180df
89876275e8d562912d9c238cd888b52065c125

Ramsauer, Bulwahn, Duda, Mauerer The List is Our Process! September 9, 2019 7 / 30

OTH OSTBAYERISCHE TECHNISCHE HOCHSCHULE REGENSBURG 2. Methodology

2011-05-2018-12

Ramsauer, Bulwahn, Duda, Mauerer The List is Our Process! September 9, 2019 17 / 30

OTH OSTBAYERISCHE TECHNISCHE HOCHSCHULE REGENSBURG 3. Evaluation

Evolution of ignored patches

Ramsauer, Bulwahn, Duda, Mauerer The List is Our Process! September 9, 2019 21 / 30

What is useful to measure?
How to improve the process?

BACKUP

What fraction of patches correctly addresses MAINTAINERS

Definition

A patch correctly addresses MAINTAINERS if...

- ▶ it addresses **one** ML as listed by `get_maintainers.pl`
- ▶ **OR:** it addresses **one** maintainers as listed by `get_maintainers.pl`

What fraction of patches correctly addresses MAINTAINERS

