

Merge-Tree

Visualizing the Integration of Commits into Linux

Evan Wilde Daniel German


University of Victoria

Motivation


- ▶ 900 Root merges per year
- ▶ 50K - 70K Commits per year
- ▶ DAG is very complex




Model



Model




Model



Heuristic

1. Identify the master branch
2. For every commit, recursively calculate shortest path to master

Heuristic



Heuristic Verification

```
Merge: 8cbd84f fd8aa2c
Author: Linus Torvalds <torvalds@linux-foundation.org>
Date: Tue Aug 10 15:38:19 2010 -0700

Merge branch for-linus of git://neil.brown.name/md

* for-linus of git://neil.brown.name/md: (24 commits)
md: clean up do_md_stop
[... edited for the sake of space]
md: split out md_rdev_init
md: be more careful setting MD_CHANGE_CLEAN
md/raid5: ensure we create a unique name for kmem_cache...
```


Heuristic Results

- ▶ Five merges had inaccurate logs causing problems
- ▶ Heuristic identified 79 merges to non-master branches
- ▶ Worked perfectly until 2007
- ▶ Foxtrot breaks heuristic in 2006

Heuristic Results

- ▶ Correctly identified 16860 commits
- ▶ Failed in 1542 commits before Dec. 12 2006
- ▶ 836 appear to be correct


Tool

- ▶ Top-to-bottom
 - ▶ Aggregate
- ▶ Bottom-to-top
 - ▶ See flow into master

An Example




An Example




042dd List Tree

- Merge tag 'regulator-v3.10-rc4' of git://git.kernel.org/pub/scm/linux/kernel/git/broonie/regulator
 - Merge remote-tracking branch 'regulator/fix/palmas' into regulator-linus
 - regulator: palmas: Fix "enable_reg" to point to the correct reg for SMPS10
 - regulator: palmas: Fix incorrect condition
 - Merge remote-tracking branch 'regulator/fix/core' into regulator-linus
 - regulator: core: Correct spelling mistake in comment
 - Merge remote-tracking branch 'regulator/fix/doc' into regulator-linus
 - regulator: Fix kernel-doc generation warnings.
 - Merge remote-tracking branch 'regulator/fix/dbx500' into regulator-linus
 - regulator: dbx500: Make local symbol static

042dd Reingold Tilford Tree



042dd Bubble Tree



042dd Summary

```
Merge tag 'regulator-v3.10-rc4' of git://git.kernel.org/pub/scm/linux/kernel/git/
```

```
Pull regulator fixes from Mark Brown:
```

```
"A few small fixes for v3.10, documentation things in the core and a  
few driver bugs."
```

```
* tag 'regulator-v3.10-rc4' of git://git.kernel.org/pub/scm/linux/kernel/git/bro  
regulator: palmas: Fix "enable_reg" to point to the correct reg for SMPS10  
regulator: palmas: Fix incorrect condition  
regulator: core: Correct spelling mistake in comment  
regulator: dbx500: Make local symbol static  
regulator: Fix kernel-doc generation warnings.
```

042dd Files



Show entries Search:

Filename	Added	Removed
drivers/regulator/core.c	5	2
drivers/regulator/dbx500-prcmu.c	12	12
drivers/regulator/palmas-regulator.c	2	2

Showing 1 to 3 of 3 entries Previous **1** Next

042dd Modules

Show entries Search:

Module 	Count 
regulator	5

Showing 1 to 1 of 1 entries [Previous](#) [1](#) [Next](#)

042dd Authors

Show entries Search:


Author	Added	Removed	Delta	Total
Charles Keepax	1	1	0	2
Kishon Vijay Abraham I	1	1	0	2
Robert P. J. Day	4	1	3	5
Sachin Kamat	13	13	0	26

Showing 1 to 4 of 4 entries Previous **1** Next

Challenges

- ▶ Tree Labels
- ▶ CID's are long and unrecognizable
- ▶ Differences in relative scale

Challenges



Conclusion

Merge trees are a much simpler way to inspect how a commit flows into the master branch.