

An Exploratory Study Into the Prevalence of Botched Code Integrations

[@wardmuylaert](#)

Ward Muylaert, Wolfgang De Meuter, and Coen De Roover
Software Languages Lab
Vrije Universiteit Brussel
Brussels, Belgium

[@oniroi](#)



Overview

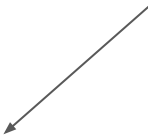
1. Introduction to Conflicts
2. Research Questions
3. Dataset
4. Research Methods
5. Results

Textual Conflicts

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

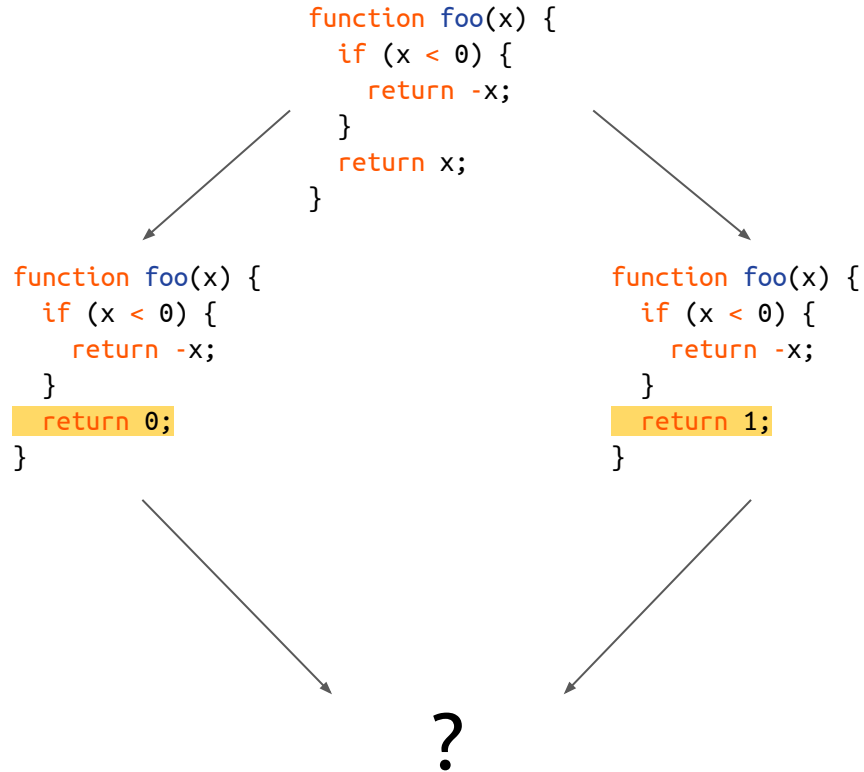
Textual Conflicts

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```



```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return 0;  
}
```

Textual Conflicts

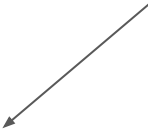


Semantic Conflicts

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

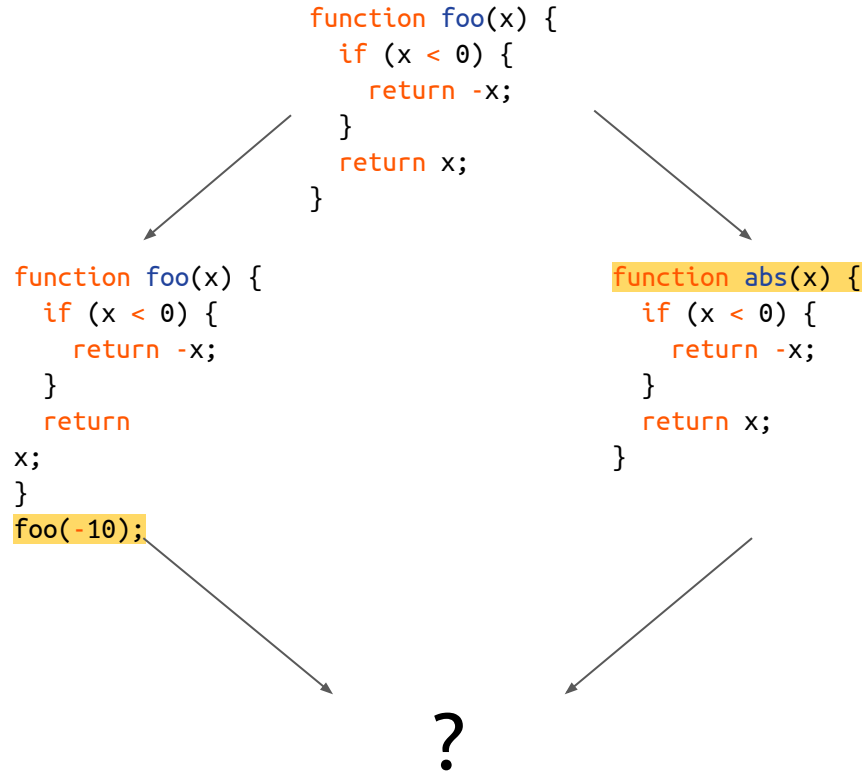
Semantic Conflicts

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```



```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return  
x;  
}  
foo(-10);
```

Semantic Conflicts



Research Questions

RQ1

How often does code integration lead to semantic conflicts?

RQ2


How much effort is needed to fix semantic conflicts after code integration?

RQ3


How long does it take to fix semantic conflicts after code integration?


GitHub


Commits on Jul 7, 2016


 **init/Kconfig: keep Expert users menu together** ...
Randy Dunlap committed with torvalds 13 hours ago

Commits on Jul 6, 2016

 **Merge git://git.kernel.org/pub/scm/linux/kernel/git/davem/net** ...
torvalds committed 19 hours ago


 **Merge tag 'sound-4.7-rc7' of git://git.kernel.org/pub/scm/linux/kerne...** ...
torvalds committed 20 hours ago

 **Merge tag 'chrome-fixes' of git://git.kernel.org/pub/scm/linux/kernel...** ...
torvalds committed 20 hours ago





GitHub


Commits on Jul 7, 2016

 **init/Kconfig: keep Expert users menu together** ...
Randy Dunlap committed with torvalds 13 hours ago

Commits on Jul 6, 2016

 **Merge git://git.kernel.org/pub/scm/linux/kernel/git/davem/net** ...
torvalds committed 19 hours ago

 **Merge tag 'sound-4.7-rc7' of git://git.kernel.org/pub/scm/linux/kerne...** ...
torvalds committed 20 hours ago


 **Merge tag 'chrome-fixes' of git://git.kernel.org/pub/scm/linux/kernel...** ...
torvalds committed 20 hours ago



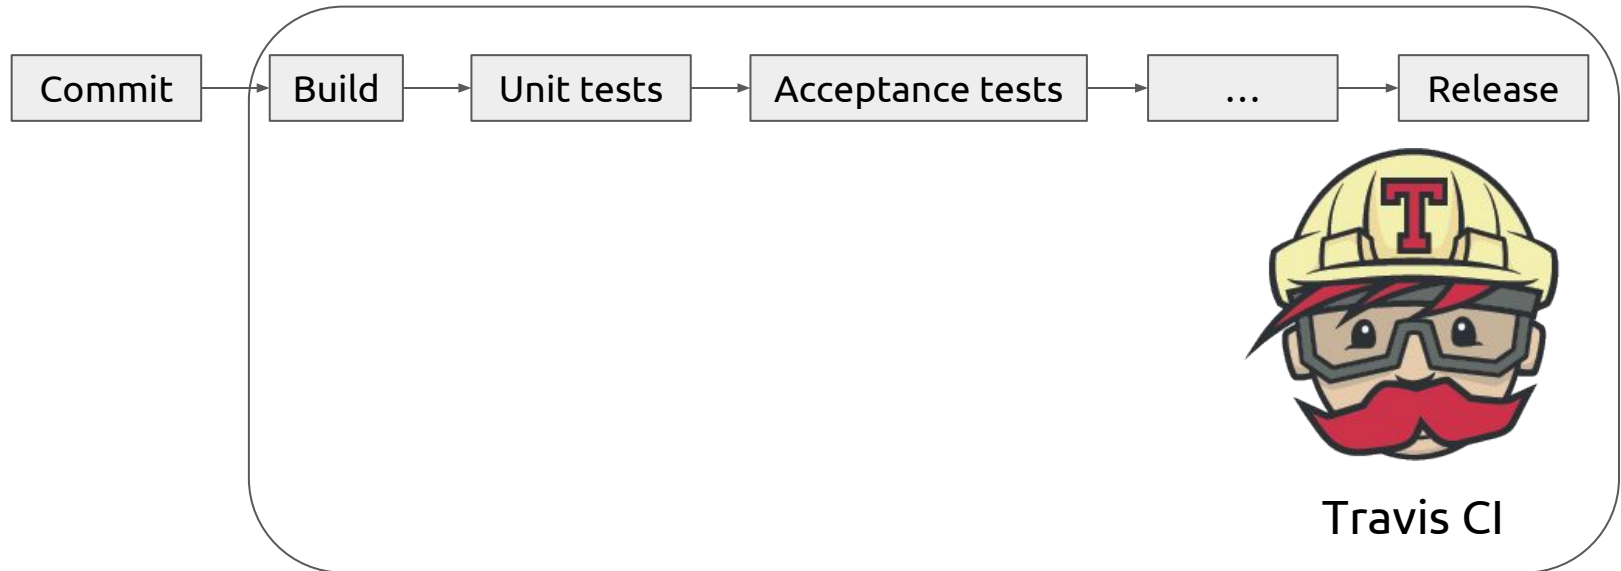
GHTorrent
150GB MySQL dump, 400M commits

Georgios Gousios, The GHTorrent
Dataset and Tool Suite, MSR 2013.

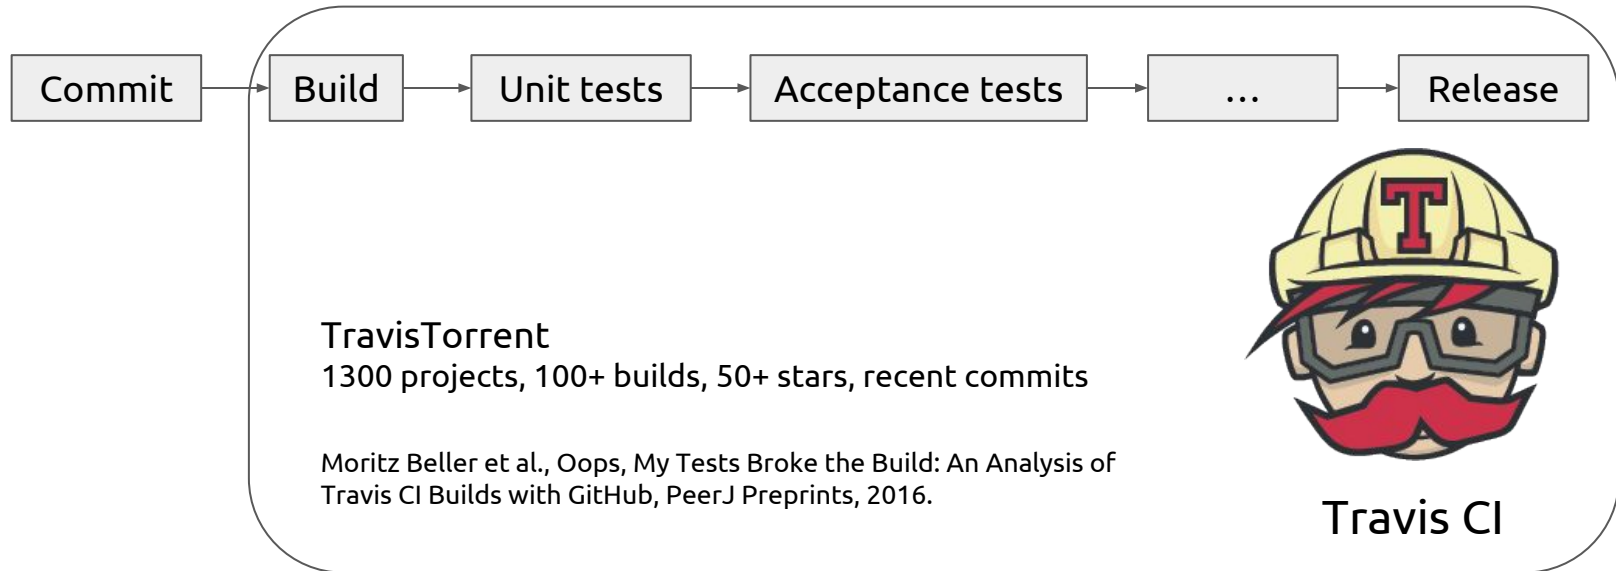
Continuous Delivery



Continuous Delivery



Continuous Delivery



Travis CI API

Builds

state
started_at
finished_at
commit_id
repository_id
...

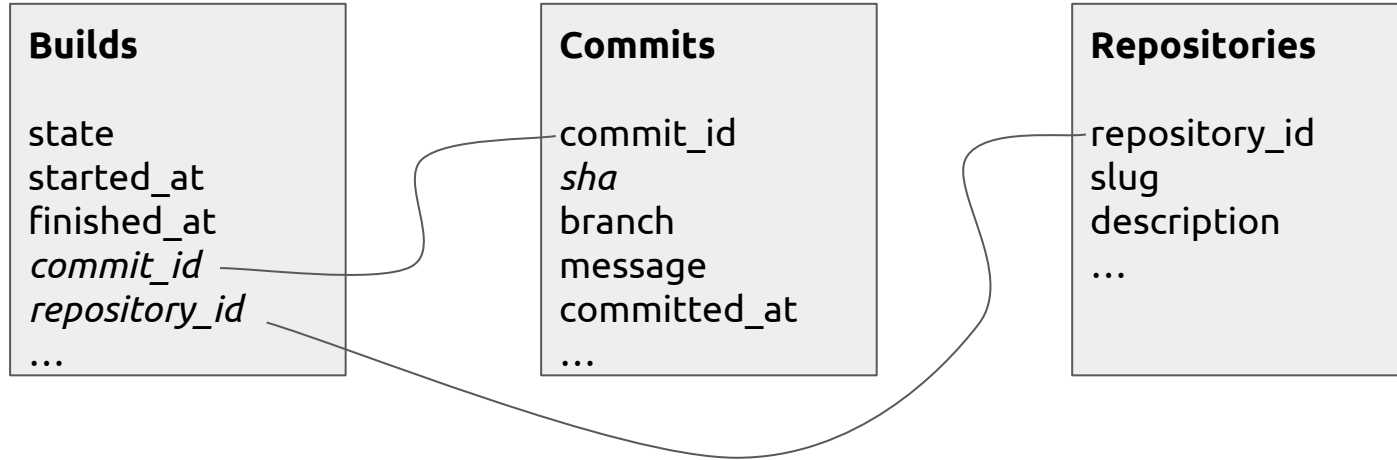
Commits

commit_id
sha
branch
message
committed_at
...

Repositories


repository_id
slug
description
...

Travis CI API



Travis CI API

GitHub



GitForensics
15GB MySQLi dump, 400M commits
Georgios Conikas, The GitForensics Dataset and Tool Suite, MSR 2013.

Builds

state
started_at
finished_at
commit_id
repository_id
...

Commits


commit_id
sha
branch
message
committed_at
...

Repositories

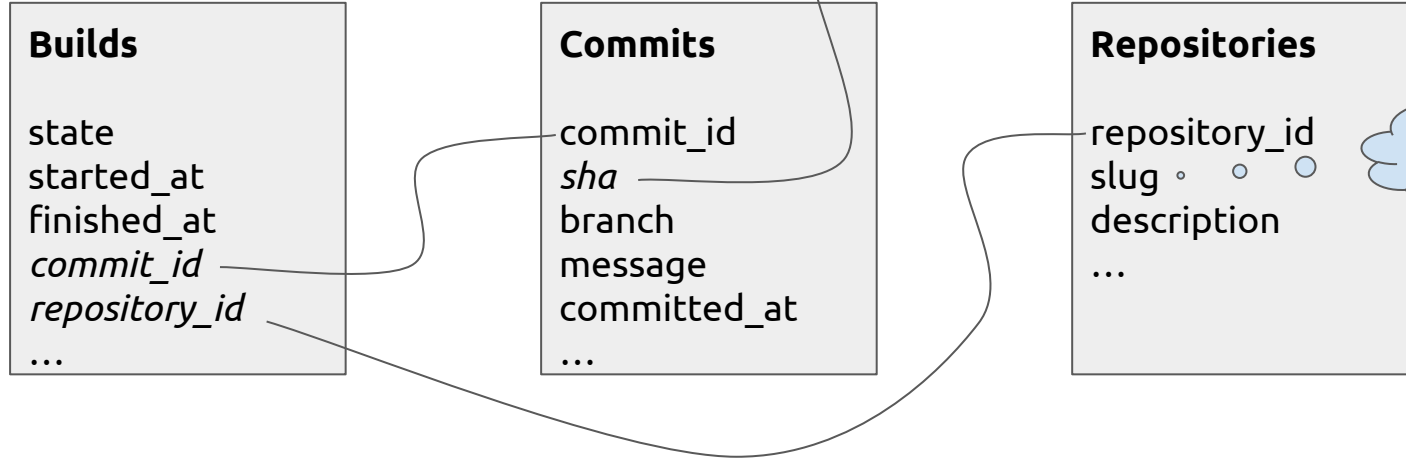
repository_id
slug
description
...

Travis CI API

GitHub



GitForensics
150GB MySQLi dump, 400M commits
Georgios Koulias, The GitForensics Dataset and Tool Suite, MSR 2013.



Travis CI API

GitHub



GitHub
15GB MySQL dump, 400M commits
Georgios Conos, The GitHub Dataset and Tool Suite, MSR 2013.

Builds

state
started_at
finished_at
commit_id
repository_id
...

Commits

commit_id
sha
branch
message
committed_at
...

Repositories


repository_id
slug
description
...

i.e., user/repo

1.1M builds

Those same 1300 projects


Identify Merge Commits



Number of parents

```
SELECT c.sha,dcs.state,dcs.repository_id
FROM (
  SELECT DISTINCT sha,state,repository_id
  FROM WardTravisCommits AS wtc
  INNER JOIN WardTravisBuilds AS wtb ON wtb.commit_id = wtc.id
) AS dcs
INNER JOIN commits AS c ON c.sha = dcs.sha
INNER JOIN commit_parents AS cp ON c.id = cp.commit_id
GROUP BY c.id,dcs.state
HAVING COUNT(*) > 1
```

Identify Merge Commits




Number of parents

Get commits
and their builds

```
SELECT c.sha,dcs.state,dcs.repository_id
FROM (
  SELECT DISTINCT sha,state,repository_id
  FROM WardTravisCommits AS wtc
  INNER JOIN WardTravisBuilds AS wtb ON wtb.commit_id = wtc.id
) AS dcs
INNER JOIN commits AS c ON c.sha = dcs.sha
INNER JOIN commit_parents AS cp ON c.id = cp.commit_id
GROUP BY c.id,dcs.state
HAVING COUNT(*) > 1
```

Identify Merge Commits



Number of parents

Get commits
and their builds

```

SELECT c.sha,dcs.state,dcs.repository_id
FROM (
  SELECT DISTINCT sha,state,repository_id
  FROM WardTravisCommits AS wtc
  INNER JOIN WardTravisBuilds AS wtb ON wtb.commit_id = wtc.id
) AS dcs
  
```


As long as they
have two or more
parents

```

INNER JOIN commits AS c ON c.sha = dcs.sha
INNER JOIN commit_parents AS cp ON c.id = cp.commit_id
GROUP BY c.id,dcs.state
HAVING COUNT(*) > 1
  
```


Enough Merge Commits?

584 Projects with 50+ Builds of Merge Commits




Enough Merge Commits?

584 Projects with 50+ Builds of Merge Commits

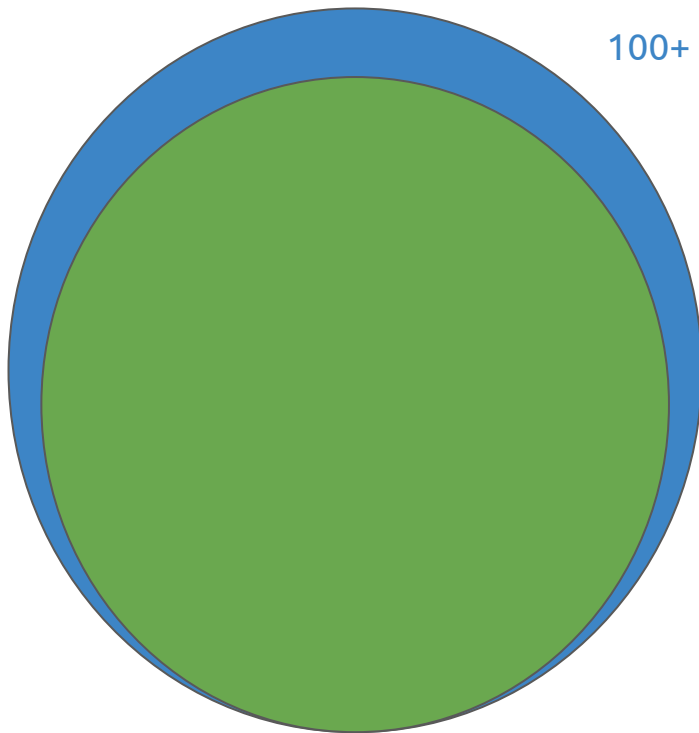


Dataset



100+ builds, 50+ stars, recent commits: 1300 projects


Dataset



100+ builds, 50+ stars, recent commits: 1300 projects

Still on Travis CI: 1248 projects

Dataset



100+ builds, 50+ stars, recent commits: 1300 projects

Still on Travis CI: 1248 projects

50+ builds of merge commits: 584 projects

What Is Failure?

✗ master Brandon Fish	[Truffle] Fix Range#{inspect,to_s} taint specs	🔴 #13717 failed 1f6820a	🕒 5 hrs 42 min 2 sec 📅 about 18 hours ago
✓ truffle-array-inspect Brandon Fish	[Truffle] Refactoring Array#inspect updates for PR	🟢 #13715 passed a7ba831	🕒 5 hrs 40 min 10 sec 📅 about 20 hours ago



Travis CI

What Is Failure?

✗ master Brandon Fish	[Truffle] Fix Range#{inspect,to_s} taint specs	#13717 failed 1f6820a	🕒 5 hrs 42 min 2 sec 📅 about 18 hours ago
✓ truffle-array-inspect Brandon Fish	[Truffle] Refactoring Array#inspect updates for PR	#13715 passed a7ba831	🕒 5 hrs 40 min 10 sec 📅 about 20 hours ago




Travis CI




- ✓ passed
- ✗ failed
- ✗ errored

cancelled
started

Failed Merge Commit Builds per Project




Failed Merge Commit Builds per Project



1/2 projects: 1/6 merges fail
1/4 projects: 1/3 merges fail

Dataset revisited



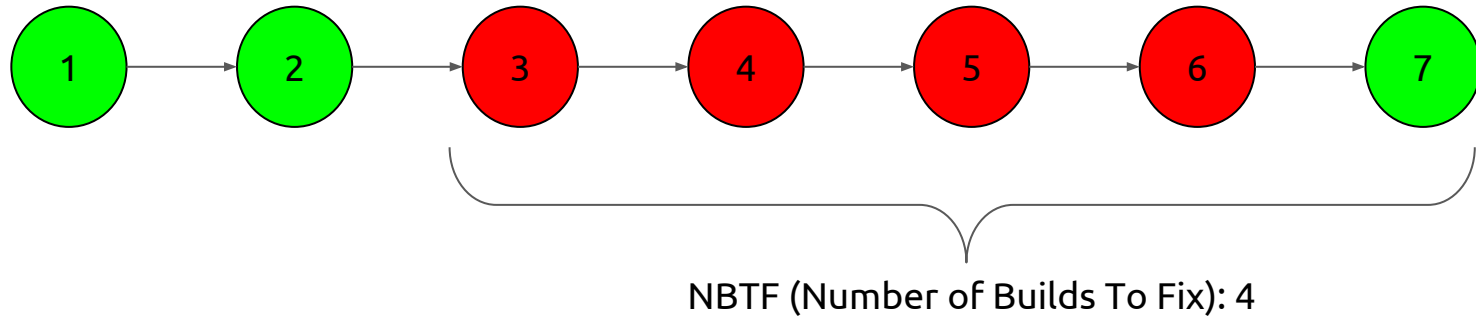
100+ builds, 50+ stars, recent commits: 1300 projects

Still on Travis CI: 1248 projects


50+ builds of merge commits: 584 projects

Without outliers: 559 projects


How Much Effort in Fixing the Build?



How Much Effort in Fixing the Build?




How Much Effort in Fixing the Build?



How Much Effort in Fixing the Build?

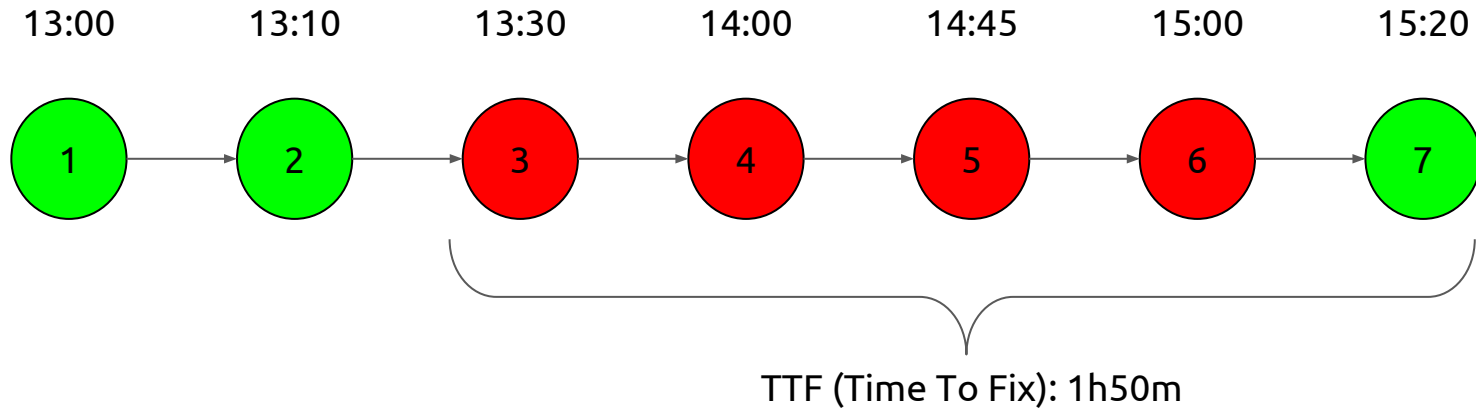


NBTF only proxy for effort



Most builds are fixed easily.
1/4 are not or are ignored.

How Long to Fix the Build?

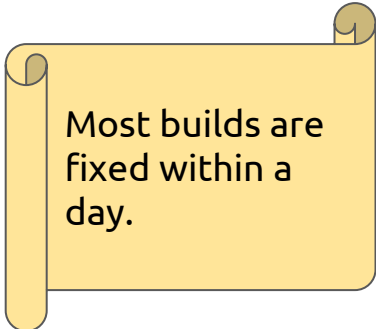


How Long to Fix the Build?

Time range (disjunct)	Amount	Relative	Cumulative
Less than 1 hour	8031	28.2%	28.2%
Less than 12 hours	6062	21.3%	49.4%
Less than 1 day	2392	8.4%	57.8%
Less than 7 days	5556	19.5%	77.3%
Less than 1 month (30 days)	3321	11.6%	88.9%
Less than 1 year (365 days)	2967	10.4%	99.3%
More than 1 year	198	0.7%	100%

How Long to Fix the Build?

Time range (disjunct)	Amount	Relative	Cumulative
Less than 1 hour	8031	28.2%	28.2%
Less than 12 hours	6062	21.3%	49.4%
Less than 1 day	2392	8.4%	57.8%
Less than 7 days	5556	19.5%	77.3%
Less than 1 month (30 days)	3321	11.6%	88.9%
Less than 1 year (365 days)	2967	10.4%	99.3%
More than 1 year	198	0.7%	100%



Most builds are fixed within a day.

How Long to Fix the Build?


Time range (disjunct)	Amount	Relative	Cumulative
Less than 1 hour	8031	28.2%	28.2%
Less than 12 hours	6062	21.3%	49.4%
Less than 1 day	2392	8.4%	57.8%
Less than 7 days	5556	19.5%	77.3%
Less than 1 month (30 days)	3321	11.6%	88.9%
Less than 1 year (365 days)	2967	10.4%	99.3%
More than 1 year	198	0.7%	100%




Time taken
 \neq
 time worked
 on

Most builds are
 fixed within a
 day.

Number of Builds To Fix vs Time To Fix



Number of Builds To Fix vs Time To Fix



Answers to Research Questions

RQ1

Conflicts after code integration happen often.

For half of the projects, the build did not succeed for 1/6 merge commits.

RQ2

Multiple builds are needed to fix the majority of botched code integrations.

RQ3

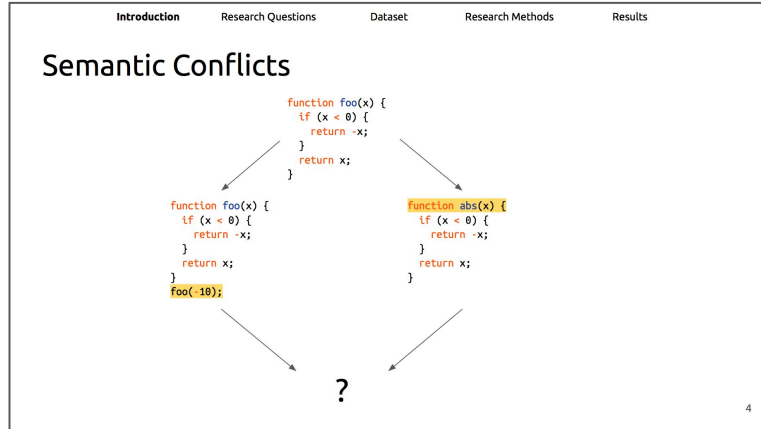
The majority of botched code integrations are fixed within a day.

Future

- Refine effort estimation using commit information
- Refine failure categorization via analysis of Travis CI logs
- Help developers with merging

Summary

Summary



Summary

Introduction Research Questions Dataset Research Methods Results

Semantic Conflicts

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}  
foo(-10);
```

```
function abs(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

?

4


Introduction Research Questions **Dataset** Research Methods Results

Continuous Delivery

```
Commit → Build → Unit tests → Acceptance tests → ... → Release
```

TravisTorrent
1300 projects, 100+ builds, 50+ stars, recent commits

Moritz Beller et al., Oops, My Tests Broke the Build: An Analysis of Travis CI Builds with GitHub, PeerJ Preprints, 2016.



Travis CI

7

Summary

Introduction Research Questions Dataset Research Methods Results

Semantic Conflicts

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}  
foo(-10);
```

```
function abs(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

?

4


Introduction Research Questions Dataset Research Methods Results

Continuous Delivery

```
graph LR  
  Commit --> Build  
  Build --> UnitTests[Unit tests]  
  UnitTests --> AcceptanceTests[Acceptance tests]  
  AcceptanceTests --> Dots[...]  
  Dots --> Release
```

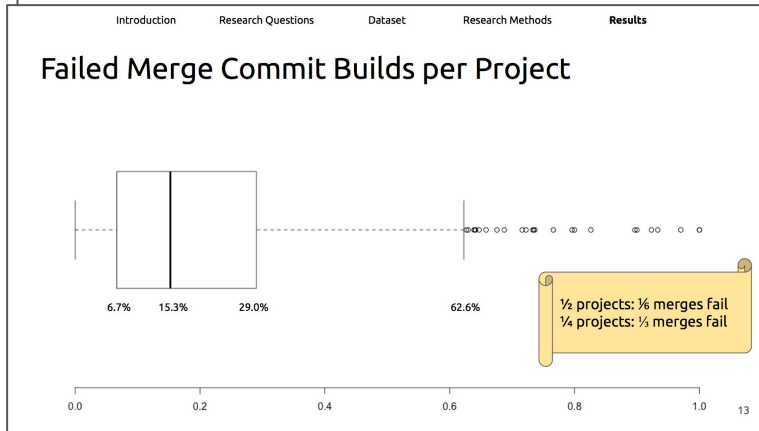
TravisTorrent
1300 projects, 100+ builds, 50+ stars, recent commits

Moritz Beller et al., Oops, My Tests Broke the Build: An Analysis of Travis CI Builds with GitHub, PeerJ Preprints, 2016.



Travis CI

7



Summary

Introduction Research Questions Dataset Research Methods Results

Semantic Conflicts


```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

```
function foo(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}  
foo(-10);
```

```
function abs(x) {  
  if (x < 0) {  
    return -x;  
  }  
  return x;  
}
```

?

4



Introduction Research Questions Dataset Research Methods Results

Continuous Delivery

```
graph LR  
  Commit --> Build  
  Build --> Unit_tests[Unit tests]  
  Unit_tests --> Acceptance_tests[Acceptance tests]  
  Acceptance_tests --> Ellipsis[...]  
  Ellipsis --> Release
```

TravisTorrent
1300 projects, 100+ builds, 50+ stars, recent commits

Moritz Beller et al., Oops, My Tests Broke the Build: An Analysis of Travis CI Builds with GitHub, PeerJ Preprints, 2016.

Travis CI

7

Future

- Refine effort information & aim to publish
 - Number of files
 - Number of lines of code
- Help developers with merging

21