



Ekstazi

Practical Regression Test Selection with Dynamic File Dependencies

Milos Gligoric, Lamyaa Eloussi, and Darko Marinov



ISSTA, Baltimore, USA

07/16/2015

CNS-0958199

CCF-1012759

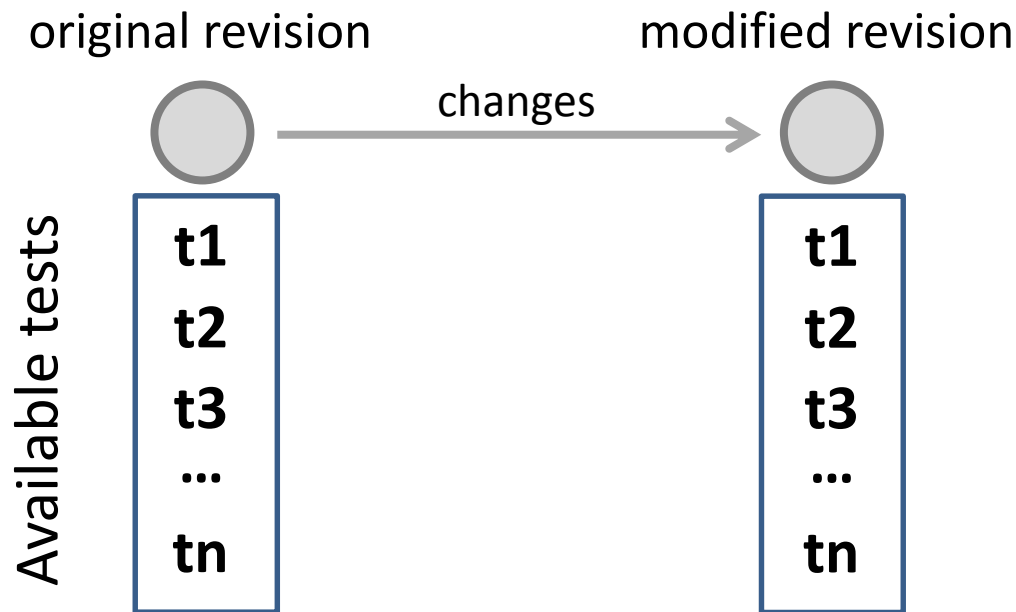
CCF-1421503

CCF-1439957



Regression Testing

- Executes tests for each new code revision
- Checks if changes broke something
- Widely used in industry



Regression Testing is Costly (1)

| test execution time single revision | number of tests single revision |
|--|------------------------------------|
|--|------------------------------------|



~5min

1667



~10min

641534



~45min

1296



~45min

361



~45min

631



~4h

4975



~17h

8663



Run many times each day

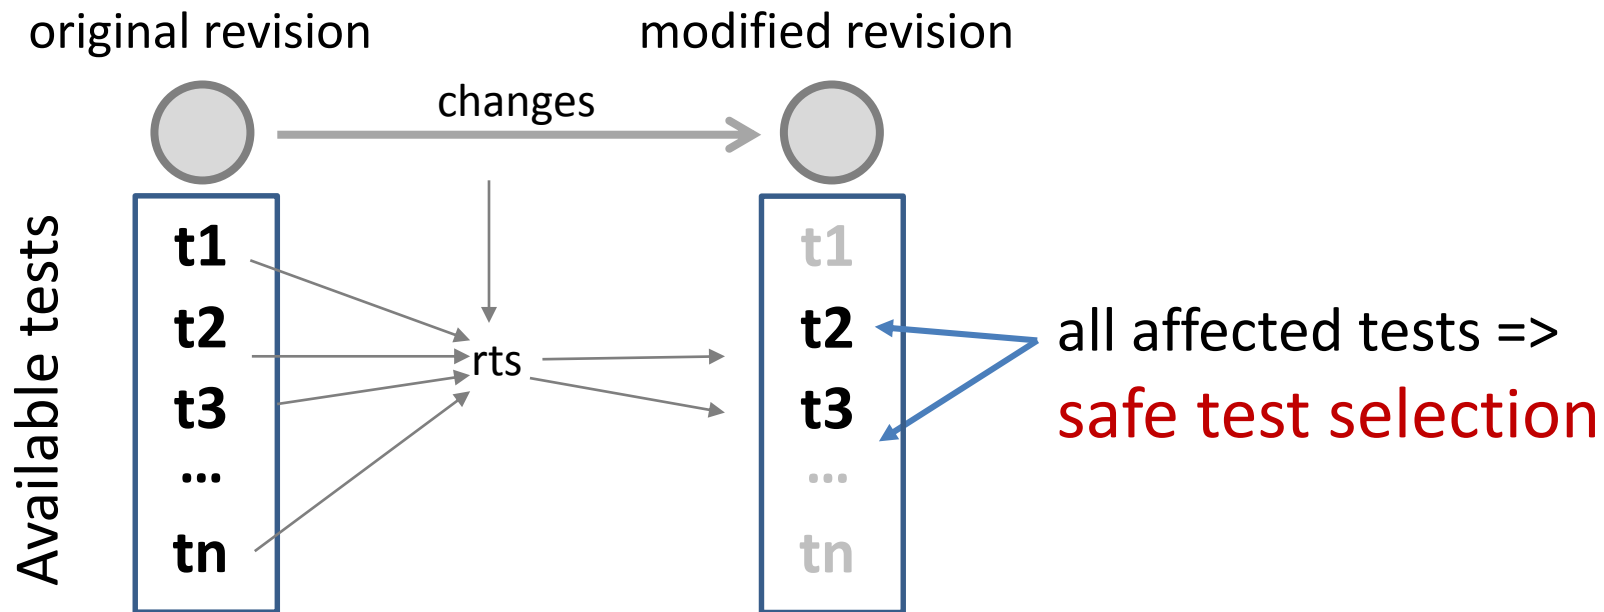
Regression Testing is Costly (2)



- * linear increase in the number of revisions per day
- linear increase in the number of tests per revision
 - => **quadratic increase in test execution time**
- 75+ million tests run per day
- 20+ revisions per minute**

Regression Test Selection (RTS)

- Speeds up regression testing
 - Without requiring more computers or energy
- Analyzes changes to a codebase
- Runs only tests whose behavior may be affected



Outline

- Motivation
- **Ekstazi RTS Technique**
- **Tool for Java**
- **Evaluation**
- **Conclusions**




Ekstazi

<http://www.ekstazi.org>

Ekstazi – Example

original revision

modified revision



changes to C2, C3

| | C1 | C2 | C3 | f |
|----|----|----|----|---|
| t1 | ★ | | | |
| t2 | | ★ | | |
| t3 | | | ★ | |
| t4 | ★ | ★ | | ★ |

Ekstazi(original, modified)

| | C1 | C2 | C3 | f |
|----|----|----|----|---|
| t1 | ★ | | | |
| t2 | | ★ | | |
| t3 | | | ★ | |
| t4 | ★ | ★ | ★ | |

```
t1() {  
  C1 obj = new C1();  
  assert(obj.m() == 1);  
}
```


```
class C1 {  
  int m() { return 1; }  
}
```

Analysis

Execution of Selected


Fine- vs. Coarse-grained Dependencies: Methods vs. Files

original revision



changes to p, r

modified revision



| | C1 | C2 | C3 | |
|----|----|----|----|---|
| | m | p | q | r |
| t1 | ★ | | | |
| t2 | | ★ | | |
| t3 | | | ★ | |
| t4 | ★ | | ★ | |

→ rts(original, modified) →

| | C1 | C2 | C3 | |
|----|----|----|----|---|
| | m | p | q | r |
| t1 | ★ | | | |
| t2 | | ★ | | |
| t3 | | | ★ | |
| t4 | ★ | | ★ | |


RTS with method dependencies select: t2, t3

RTS with file dependencies select: t2, t3, t4

Prior RTS Techniques vs. Ekstazi

- 30 years, 300+ papers
- fine-grained dependency
 - methods, basic blocks, statements, ...
- fewer tests
- slow analysis
- unsafe
- coarse-grained dependency
 - dynamically used files (code classes and other files)
- more tests
- fast analysis
- safer

Important: total time = analysis + execution



Intuition for Surprising Results

- Efficient due to these properties
 - Small fraction of files are modified at each revision
 - Small fraction of tests depend on each file
 - Changes are localized
- Safe for (m)any code/file change
 - Unchanged files cannot affect test execution

Ekstazi is Safer – Example (1)

revision 0

```
class A {  
  A() {}  
  int m() { return 1; }  
}  
class B extends A {  
  B() {} // calls A()  
}
```

```
test() {  
  B b = new B();  
  assert(b.m() == 1);  
}
```

revision 1

```
class A {  
  A() {}  
  int m() { return 1; }  
}  
class B extends A {  
  B() {}  
  @Override  
  int m() { return 2; }  
}
```

| | A() | B() | A.m() |
|------|-----|-----|-------|
| test | ★ | ★ | ★ |



| | A | B |
|------|---|---|
| test | ★ | ★ |



Ekstazi is Safer – Example (2)

revision 0

```
class A {  
  A() {}  
  public void m() { ... }  
}
```

revision 1

```
class A {  
  A() {}  
  public void m() { ... }  
  public void n() { ... }  
}
```

```
test() {  
  Method[] methods = A.class.getDeclaredMethods();  
  assert(methods.length == 1);  
}
```

| | A() | A.m() |
|------|-----|-------|
| test | | |



| | A |
|------|---|
| test | ★ |




Ekstazi Tool

- Implemented for JVM languages
<http://www.ekstazi.org>
- Technical challenges
 - Efficient monitoring of used classes
 - Handling jar files
 - Supporting parallel execution
 - Fast file comparison
 - Smart hashing (i.e., ignoring debug info in classfiles)

Evaluation – Projects

- 32 projects
- 773,565 tests
- ~5M LOC
- >500 revisions



Project | #TestClasses

| | |
|------------------|-------------|
| Cucumber | 49 |
| JodaTime | 124 |
| Retrofit | 15 |
| CommonsValidator | 61 |
| BVal | 21 |
| CommonsJXPath | 33 |
| GraphHopper | 80 |
| River | 14 |
| Functor | 164 |
| EmpireDB | 23 |
| JFreeChart | 359 |
| CommonsColl4 | 145 |
| CommonsLang3 | 121 |
| CommonsConfig | 141 |
| PdfBox | 94 |
| GSCollections | 1106 |

| | |
|-----------------|-------------|
| CommonsNet | 37 |
| ClosureCompiler | 233 |
| CommonsDBCP | 27 |
| Log4j | 38 |
| JGit | 229 |
| CommonsIO | 84 |
| Ivy | 121 |
| Jenkins | 86 |
| CommonsMath | 461 |
| Ant | 234 |
| Continuum | 68 |
| Guava | 348 |
| Camel | 2015 |
| Jetty | 504 |
| Hadoop | 317 |
| ZooKeeper | 127 |


Project | #TestClasses | Time [sec]

| | | |
|------------------|------|-----------|
| Cucumber | 49 | 8 |
| JodaTime | 124 | 10 |
| Retrofit | 15 | 10 |
| CommonsValidator | 61 | 11 |
| BVal | 21 | 13 |
| CommonsJXPath | 33 | 15 |
| GraphHopper | 80 | 15 |
| River | 14 | 17 |
| Functor | 164 | 21 |
| EmpireDB | 23 | 27 |
| JFreeChart | 359 | 30 |
| CommonsColl4 | 145 | 32 |
| CommonsLang3 | 121 | 36 |
| CommonsConfig | 141 | 39 |
| PdfBox | 94 | 40 |
| GSCollections | 1106 | 51 |

| | | |
|-----------------|------|-------------|
| CommonsNet | 37 | 68 |
| ClosureCompiler | 233 | 71 |
| CommonsDBCP | 27 | 76 |
| Log4j | 38 | 79 |
| JGit | 229 | 83 |
| CommonsIO | 84 | 98 |
| Ivy | 121 | 170 |
| Jenkins | 86 | 171 |
| CommonsMath | 461 | 249 |
| Ant | 234 | 380 |
| Continuum | 68 | 453 |
| Guava | 348 | 469 |
| Camel | 2015 | 1296 |
| Jetty | 504 | 1363 |
| Hadoop | 317 | 1415 |
| ZooKeeper | 127 | 2565 |

Evaluation – Setup

measure execution time




GitHub



git clone PROJECT

git checkout 20 revisions earlier


Evaluation – Apache CXF



- Reduces number of tests: **~15x**
- Reduces test execution time: **~8x**
- Reduces build+test time: **~3x**

Executed on our group's server

Evaluation – Guava



Executed on Amazon EC2

Project | Tests %

| | |
|------------------|----|
| Cucumber | 12 |
| JodaTime | 21 |
| Retrofit | 16 |
| CommonsValidator | 6 |
| BVal | 13 |
| CommonsXPath | 20 |
| GraphHopper | 16 |
| River | 6 |
| Functor | 13 |
| EmpireDB | 18 |
| JFreeChart | 5 |
| CommonsColl4 | 9 |
| CommonsLang3 | 11 |
| CommonsConfig | 20 |
| PdfBox | 12 |
| GSCollections | 29 |

| | |
|-----------------|----|
| CommonsNet | 10 |
| ClosureCompiler | 17 |
| CommonsDBCP | 21 |
| Log4j | 6 |
| JGit | 22 |
| CommonsIO | 12 |
| Ivy | 38 |
| Jenkins | 7 |
| CommonsMath | 6 |
| Ant | 13 |
| Continuum | 10 |
| Guava | 13 |
| Camel | 5 |
| Jetty | 26 |
| Hadoop | 7 |
| ZooKeeper | 20 |

Project | Tests % | Time %

| | | |
|------------------|----|-----|
| Cucumber | 12 | 99 |
| JodaTime | 21 | 107 |
| Retrofit | 16 | 104 |
| CommonsValidator | 6 | 88 |
| BVal | 13 | 138 |
| CommonsJXPath | 20 | 94 |
| GraphHopper | 16 | 85 |
| River | 6 | 35 |
| Functor | 13 | 112 |
| EmpireDB | 18 | 112 |
| JFreeChart | 5 | 80 |
| CommonsColl4 | 9 | 66 |
| CommonsLang3 | 11 | 60 |
| CommonsConfig | 20 | 72 |
| PdfBox | 12 | 80 |
| GSCollections | 29 | 107 |

| | | |
|--------------|----|----|
| CommonsNet | 10 | 21 |
| CommonsMaven | 17 | 62 |
| CommonsDBCP | 21 | 46 |
| JGit | 22 | 65 |
| CommonsIO | 12 | 30 |
| Ivy | 38 | 53 |
| Jenkins | 7 | 74 |
| CommonsMath | 6 | 77 |
| Ant | 13 | 24 |
| Continuum | 10 | 32 |
| Guava | 13 | 45 |
| Camel | 5 | 9 |
| Jetty | 26 | 57 |
| Hadoop | 7 | 38 |
| ZooKeeper | 20 | 43 |

Average All = 68%

Average Long Running = 46%

Ekstazi Without and With Smart Checksum

| Project | WO/ Smart Checksum | | W/ Smart Checksum | |
|------------------|--------------------|-----------|-------------------|-----------|
| | Tests% | Time% | Tests% | Time% |
| CommonsValidator | 7 | 93 | 6 | 88 |
| JFreeChart | 6 | 87 | 5 | 84 |
| CommonsNet | 11 | 28 | 9 | 26 |
| CommonsDBCP | 40 | 60 | 23 | 43 |
| CommonsIO | 21 | 42 | 14 | 30 |
| Ivy | 47 | 63 | 38 | 53 |
| Jenkins | 14 | 72 | 7 | 74 |
| CommonsMath | 6 | 85 | 6 | 75 |
| Camel | 5 | 9 | 5 | 9 |
| Avg. | 17 | 60 | 13 | 54 |

FaultTracer vs. Ekstazi

| Project | FaultTracer | | Ekstazi | |
|------------------|-------------|------------|-----------|-----------|
| | Tests% | Time% | Tests% | Time% |
| JodaTime | 3 | 663 | 21 | 107 |
| CommonsValidator | 1 | 255 | 6 | 88 |
| CommonsXPath | 14 | 294 | 20 | 94 |
| CommonsLang3 | 1 | 183 | 8 | 63 |
| CommonsConfig | 8 | 223 | 19 | 76 |
| CommonsNet | 2 | 57 | 9 | 26 |
| Avg. | 5 | 279 | 14 | 76 |

Ekstazi Users



Zed - actuator services platform

JBoss Operations Network

JBoss Fuse examples

Proprietary banking software

Example from Apache Camel: commit ff94895c

Date: Thu Nov 13 09:17:06 2014 -0600

Including Ekstazi (www.ekstazi.org) profile to optimize execution of the tests

Feedback on Ekstazi



Henryk Konsek @hekonsek · Dec 30

Yay, new version of Ekstazi Maven Plugin is out. Highly recommended! ekstazi.org/maven.html



“Your tool is quite impressive; congratulations!”
an Apache Commons Math developer



Hangout with Google managers and developers

Several requests from various (Apache) developers,
e.g., forcefailing

Conclusions



Ekstazi

- Novel regression test selection technique
 - tracks dynamic file dependencies
 - guarantees safety in many common cases
 - speedups up regression testing for many projects
- Adopted by several popular open-source projects

<http://www.ekstazi.org>

THE UNIVERSITY OF
TEXAS
— AT AUSTIN —



ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Milos Gligoric <milos.gligoric@gmail.com>

Lamyaa Eloussi <eloussi2@Illinois.edu>

Darko Marinov <marinov@Illinois.edu>



Ekstazi – Maven Integration

```
<plugin>  
<groupId>org.ekstazi</groupId>  
<artifactId>ekstazi-maven-plugin</artifactId>  
<version>4.6.0</version>  
<executions>  
  <execution>  
    <id>ekstazi</id> <goals> <goal>select</goal> </goals>  
  </execution>  
</executions>  
</plugin>
```

Method and Class Granularities

| Project | Method | | | Class | | |
|-----------------|-----------|------------|-----------|-----------|------------|-----------|
| | Entities% | Time (AEC) | Time (AE) | Entities% | Time (AEC) | Time (AE) |
| Bval | 16 | 138 | 94 | 13 | 138 | 97 |
| ClosureCompiler | 20 | 96 | 53 | 17 | 62 | 50 |
| CommonsColl4 | 7 | 81 | 60 | 9 | 66 | 55 |
| CommonsConfig | 19 | 76 | 57 | 20 | 72 | 58 |
| CommonsDBCP | 23 | 43 | 37 | 21 | 46 | 39 |
| CommonsIO | 14 | 30 | 24 | 12 | 30 | 24 |
| CommonsLang3 | 8 | 63 | 51 | 11 | 60 | 53 |
| CommonsMath | 6 | 75 | 17 | 6 | 77 | 16 |
| CommonsNet | 9 | 26 | 22 | 10 | 21 | 21 |
| Cucumber | 13 | 105 | 78 | 12 | 99 | 76 |
| EmpireDB | 13 | 117 | 100 | 18 | 112 | 99 |
| Avg. | 13 | 87 | 61 | 14 | 77 | 59 |

Name Meaning

www.vocabulary.com

If you've ever been so happy that the rest of the world seemed to disappear, you've felt *ecstasy* — a feeling or state of intensely beautiful bliss.

Ekstazi – Results

| Project | Tests% | Time% |
|-------------------|-----------|-----------|
| Camel | 5 | 9 |
| CommonsDBCP | 40 | 60 |
| CommonsIO | 21 | 42 |
| CommonsMath | 6 | 85 |
| CommonsNet | 11 | 28 |
| CommonsValidator | 7 | 93 |
| Ivy | 47 | 63 |
| Jenkins | 14 | 72 |
| JFreeChart | 6 | 87 |
| Avg. (all) | 17 | 60 |

Ekstazi Without and With Smart Checksum

| Project | WO/ Smart Checksum | | W/ Smart Checksum | |
|-------------------|--------------------|-----------|-------------------|-----------|
| | Tests% | Time% | Tests% | Time% |
| Camel | 5 | 9 | 5 | 9 |
| CommonsDBCP | 40 | 60 | 23 | 43 |
| CommonsIO | 21 | 42 | 14 | 30 |
| CommonsMath | 6 | 85 | 6 | 75 |
| CommonsNet | 11 | 28 | 9 | 26 |
| CommonsValidator | 7 | 93 | 6 | 88 |
| Ivy | 47 | 63 | 38 | 53 |
| Jenkins | 14 | 72 | 7 | 74 |
| JFreeChart | 6 | 87 | 5 | 84 |
| Avg. (all) | 17 | 60 | 13 | 54 |

FaultTracer vs. Ekstazi

| Project | FaultTracer | | Ekstazi | |
|------------------|-------------|------------|-----------|-----------|
| | Tests% | Time% | Tests% | Time% |
| CommonsConfig | 8 | 223 | 19 | 76 |
| CommonsXPath | 14 | 294 | 20 | 94 |
| CommonsLang3 | 1 | 183 | 8 | 63 |
| CommonsNet | 2 | 57 | 9 | 26 |
| CommonsValidator | 1 | 255 | 6 | 88 |
| JodaTime | 3 | 663 | 21 | 107 |
| Avg. | 5 | 279 | 14 | 76 |

Project | #TestClasses

| | |
|------------------|-----------|
| Cucumber | 8 |
| JodaTime | 10 |
| Retrofit | 10 |
| CommonsValidator | 11 |
| BVal | 13 |
| CommonsXPath | 15 |
| GraphHopper | 15 |
| River | 17 |
| Functor | 21 |
| EmpireDB | 27 |
| JFreeChart | 30 |
| CommonsColl4 | 32 |
| CommonsLang3 | 36 |
| CommonsConfig | 39 |
| PdfBox | 40 |
| GSCollections | 51 |

| | |
|-----------------|-------------|
| CommonsNet | 68 |
| ClosureCompiler | 71 |
| CommonsDBCP | 76 |
| Log4j | 79 |
| JGit | 83 |
| CommonsIO | 98 |
| Ivy | 170 |
| Jenkins | 171 |
| CommonsMath | 249 |
| Ant | 380 |
| Continuum | 453 |
| Guava | 469 |
| Camel | 1296 |
| Jetty | 1363 |
| Hadoop | 1415 |
| ZooKeeper | 2565 |

Project | Time [sec] | #TestClasses | Tests%

| | | | |
|------------------|----|------|-----------|
| Cucumber | 8 | 49 | 12 |
| JodaTime | 10 | 124 | 21 |
| Retrofit | 10 | 15 | 16 |
| CommonsValidator | 11 | 61 | 6 |
| BVal | 13 | 21 | 13 |
| CommonsJXPath | 15 | 33 | 20 |
| GraphHopper | 15 | 80 | 16 |
| River | 17 | 14 | 6 |
| Functor | 21 | 164 | 13 |
| EmpireDB | 27 | 23 | 18 |
| JFreeChart | 30 | 359 | 5 |
| CommonsColl4 | 32 | 145 | 9 |
| CommonsLang3 | 36 | 121 | 11 |
| CommonsConfig | 39 | 141 | 20 |
| PdfBox | 40 | 94 | 12 |
| GSCollections | 51 | 1106 | 29 |

| | | | |
|-----------------|------|------|-----------|
| CommonsNet | 68 | 37 | 10 |
| ClosureCompiler | 71 | 233 | 17 |
| CommonsDBCP | 76 | 27 | 21 |
| Log4j | 79 | 38 | 6 |
| JGit | 83 | 229 | 22 |
| CommonsIO | 98 | 84 | 12 |
| Ivy | 170 | 121 | 38 |
| Jenkins | 171 | 86 | 7 |
| CommonsMath | 249 | 461 | 6 |
| Ant | 380 | 234 | 13 |
| Continuum | 453 | 68 | 10 |
| Guava | 469 | 348 | 13 |
| Camel | 1296 | 2015 | 5 |
| Jetty | 1363 | 504 | 26 |
| Hadoop | 1415 | 317 | 7 |
| ZooKeeper | 2565 | 127 | 20 |

Project | Time [sec] | #TestClasses | Tests % | Time %

| | | | | |
|------------------|----|------|----|------------|
| Cucumber | 8 | 49 | 12 | 99 |
| JodaTime | 10 | 124 | 21 | 107 |
| Retrofit | 10 | 15 | 16 | 104 |
| CommonsValidator | 11 | 61 | 6 | 88 |
| BVal | 13 | 21 | 13 | 138 |
| CommonsJXPath | 15 | 33 | 20 | 94 |
| GraphHopper | 15 | 80 | 16 | 85 |
| River | 17 | 14 | 6 | 35 |
| Functor | 21 | 164 | 13 | 112 |
| EmpireDB | 27 | 23 | 18 | 112 |
| JFreeChart | 30 | 359 | 5 | 80 |
| CommonsColl4 | 32 | 145 | 9 | 66 |
| CommonsLang3 | 36 | 121 | 11 | 60 |
| CommonsConfig | 39 | 141 | 20 | 72 |
| PdfBox | 40 | 94 | 12 | 80 |
| GSCollections | 51 | 1106 | 29 | 107 |

| | | | | |
|-----------------|------|------|----|-----------|
| CommonsNet | 68 | 37 | 10 | 21 |
| ClosureCompiler | 71 | 233 | 17 | 62 |
| CommonsDBCP | 76 | 27 | 21 | 46 |
| Log4j | 79 | 38 | 6 | 62 |
| JGit | 83 | 229 | 22 | 65 |
| CommonsIO | 98 | 84 | 12 | 30 |
| Ivy | 170 | 121 | 38 | 53 |
| Jenkins | 171 | 86 | 7 | 74 |
| CommonsMath | 249 | 461 | 6 | 77 |
| Ant | 380 | 234 | 13 | 24 |
| Continuum | 453 | 68 | 10 | 32 |
| Guava | 469 | 348 | 13 | 45 |
| Camel | 1296 | 2015 | 5 | 9 |
| Jetty | 1363 | 504 | 26 | 57 |
| Hadoop | 1415 | 317 | 7 | 38 |
| ZooKeeper | 2565 | 127 | 20 | 43 |