

Mining Performance Regression Inducing Code Changes in Evolving Software

Qi Luo, Denys Poshyvanyk, Mark Grechanik*

College of William and Mary

*University of Illinois at Chicago



MSR 2016
Austin, TX, U.S.

```
public synchronized void regressionFunc ();
static Regressions staticRT = new Regressions();

public void setTeams(Collection<Team> teams) {
    this.teams = teams;
}

Backlog.staticRT.regressionFunc();
```



```
Set<Integer> selectedBacklogIds = this.getSelectedBacklogs();
if(selectedBacklogIds == null || selectedBacklogIds.size() == 0) {
    Collection<Product> products = new ArrayList<Product>();
    productBusiness.storeAllTimeSheets(products);
    for (Product product: products) {
        selectedBacklogIds.add(product.getId());
    }
}
return Action.SUCCESS;
```



```
for(Story child : story.getChildren()) {
    if (child.getId() == story.getId()) {
        continue;
    }
    StoryTreeBranchMetrics childMetrics = this.calculateStoryTreeMetrics(child)
}
return metrics;
```



Automated Detection of Performance Regressions Using Regression Models on Clustered Performance Counters

Weiyi Shang, Ahmed E. Hassan
Software Analysis and Intelligence Lab (SAIL)
Queen's University, Kingston, Ontario, Canada
{swy, ahmed}@cs.queensu.ca

Mohamed Nasser, Parminder Flora
BlackBerry
Waterloo, Ontario, Canada

ABSTRACT

Performance testing is conducted before deploying system updates in order to ensure that the performance of large software systems did not degrade (i.e., no performance regressions). During such testing, thousands of performance counters are collected. However, comparing thousands of performance counters across versions of a software system is very time consuming and error-prone. In an effort to auto-

1. INTRODUCTION

Performance assurance activities aim to identify and eliminate regressions in each newly released version. Performance regressions are response to expected resource utilization and may compromise the

Mining Performance Regression Testing Repositories for Automated Performance Analysis

King Chun Foo¹, Zhen Ming Jiang², Bram Adams², Ahmed E. Hassan², Ying Zou¹, Parminder Flora³
Department of Electrical and Computer Engineering¹
Queen's University
Kingston, ON, Canada
{k.foo, ying.zou}@queensu.ca

School of Computing²
Queen's University
Kingston, ON, Canada
{zmjiang, bram, ahmed}@cs.queensu.ca

Performance Engineering³
Research In Motion
Waterloo, ON, Canada

Abstract— Performance regression testing detects performance regressions in a system under load. Such regressions refer to situations where software performance degrades compared to previous releases, although the new version behaves correctly. In current practice, performance analysts must manually analyze performance regression testing data to uncover performance regressions. This process is both time-consuming and error-prone due to the large volume of metrics collected, the absence of formal performance objectives and the subjectivity of individual

During the course of the test, the running system is re-utilized. After each test, domain-knowledge and prior knowledge of metric values are used to analyze the deviations of metric values from the baseline. A defect report will be generated if the observed deviations are significant. The test concludes that the observed regressions. In a large environment, this process is both time-consuming and error-prone.

Automated Detection of Performance Regressions Using Statistical Process Control Techniques

Thanh H. D. Nguyen, Bram Adams,
Zhen Ming Jiang, Ahmed E. Hassan
Software Analysis and Intelligence Lab (SAIL)
School of Computing, Queen's University
Kingston, Ontario, Canada
{thanhnguyen,bram,zmjiang,ahmed}@cs.queensu.ca

Mohamed Nasser, Parminder Flora
Performance Engineering
Research In Motion (RIM)
Waterloo, Ontario, Canada

ABSTRACT

The goal of performance regression testing is to check for performance regressions in a new version of a software system. Performance regression testing is an important phase in the software development process. Performance regression testing is very time consuming yet there is usually little time assigned for it. A typical test run would outlast

features, code changes might degrade the software's performance. Hence, performance engineers must perform regression tests to make sure that the software still performs as good as previous versions. Performance regression testing is very important to large software systems where a large number of field problems are performance related [19].

Performance regression testing is very time consuming yet

Main Effects Screening: A Distributed Continuous Quality Assurance Process for Monitoring Performance Degradation in Evolving Software Systems

Cemal Yilmaz[†], Arvind S. Krishna[‡], Atif Memon[†], Adam Porter[†], Douglas C. Schmidt[†], Aniruddha Gokhale[‡], Balachandran Natarajan[‡]

[†]Dept. of Computer Science, University of Maryland, College Park, MD 20742

[‡]Dept. of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37235

Developers of highly configurable performance systems often use a type of in-house "regression testing" to ensure their systems have not adversely affected their performance across its large configuration space. Time and resource constraints often limit the scope of this testing.

Although these software parameters promote flexibility and portability, they also require that the software be tested in an enormous number of configurations. This creates serious challenges for developers who must ensure that their decisions, additions, and modifications work across this large (and often changing) configuration space:

- Settings that maximize performance for a particular

V_i

```
public synchronized 1  
    void calculate(); 2  
  
input a, b            3  
A item;  
  
item = new A();        5  
  
item.calculate();      8
```

V_{i+1}

```
public synchronized 1  
    void calculate(); 2  
  
input a, b            3  
A item;  
  
if (a > b)           5  
    item = new A();    6  
else item = A.getItem(); 7  
item.calculate();      8
```



Increased Execution Time

V_i

```
public synchronized 1  
    void calculate(); 2  
  
input a, b            3  
A item;  
  
item = new A();        5  
  
item.calculate();      8
```

V_{i+1}

```
public synchronized 1  
    void calculate(); 2  
  
input a, b            3  
A item;  
  
if (a > b)           5  
    item = new A();    6  
else item = A.getItem(); 7  
item.calculate();      8
```



Increased Execution Time



**Problematic Code
Change**

```
public synchronized void regressionFunc ();
static Regressions staticRT = new Regressions();

public void setTeams(Collection<Team> teams) {
    this.teams = teams;
}

Backlog.staticRT.regressionFunc();
```



```
Set<Integer> selectedBacklogIds = this.getSelectedBacklogs();
if(selectedBacklogIds == null || selectedBacklogIds.size() == 0) {
    Collection<Product> products = new ArrayList<Product>();
    productBusiness.storeAllTimeSheets(products);
    for (Product product: products) {
        selectedBacklogIds.add(product.getId());
    }
}
return Action.SUCCESS;
```



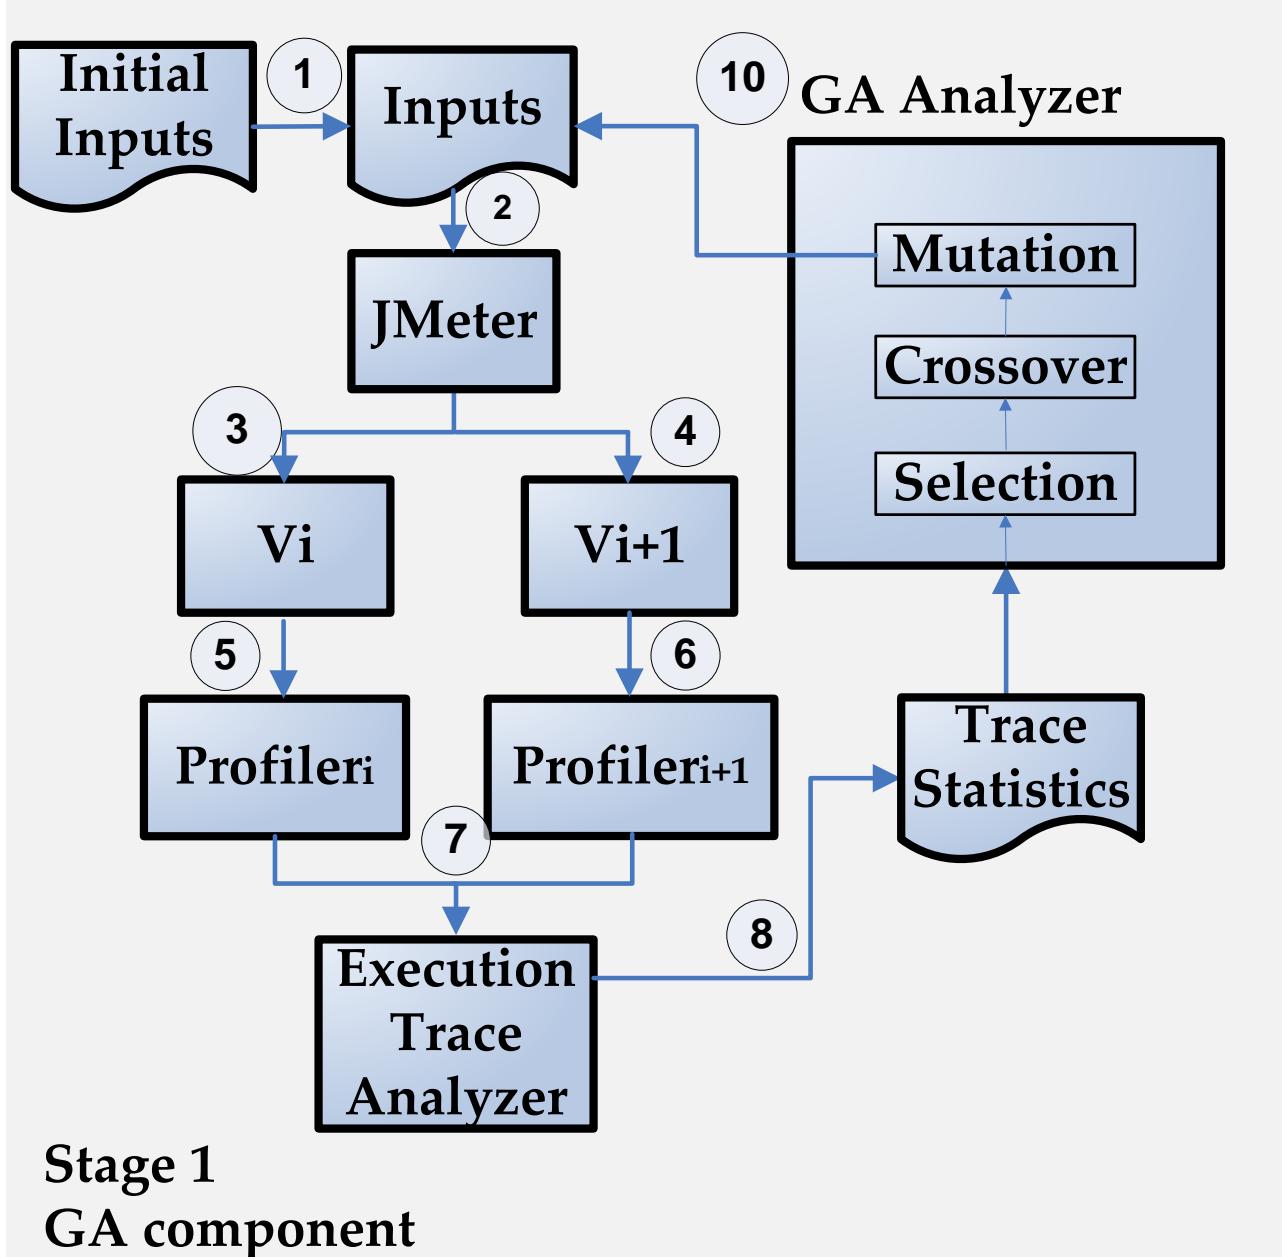
```
for(Story child : story.getChildren()) {
    if (child.getId() == story.getId()) {
        continue;
    }
    StoryTreeBranchMetrics childMetrics = this.calculateStoryTreeMetrics(child)
}
return metrics;
```

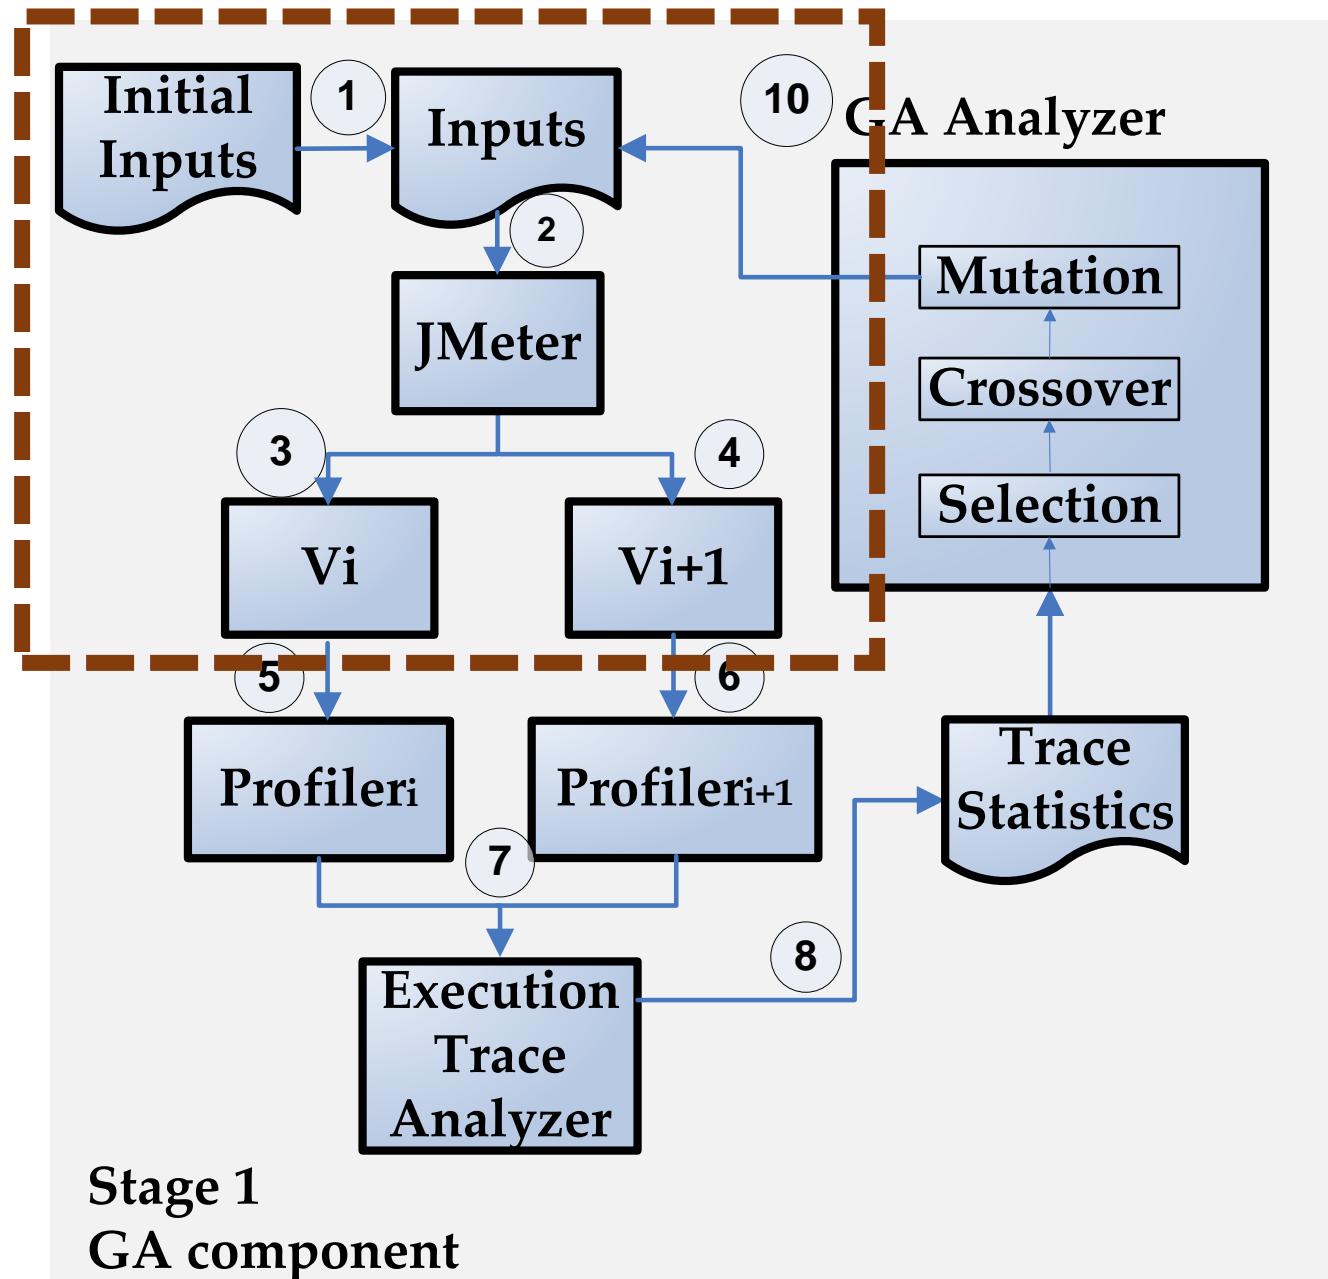


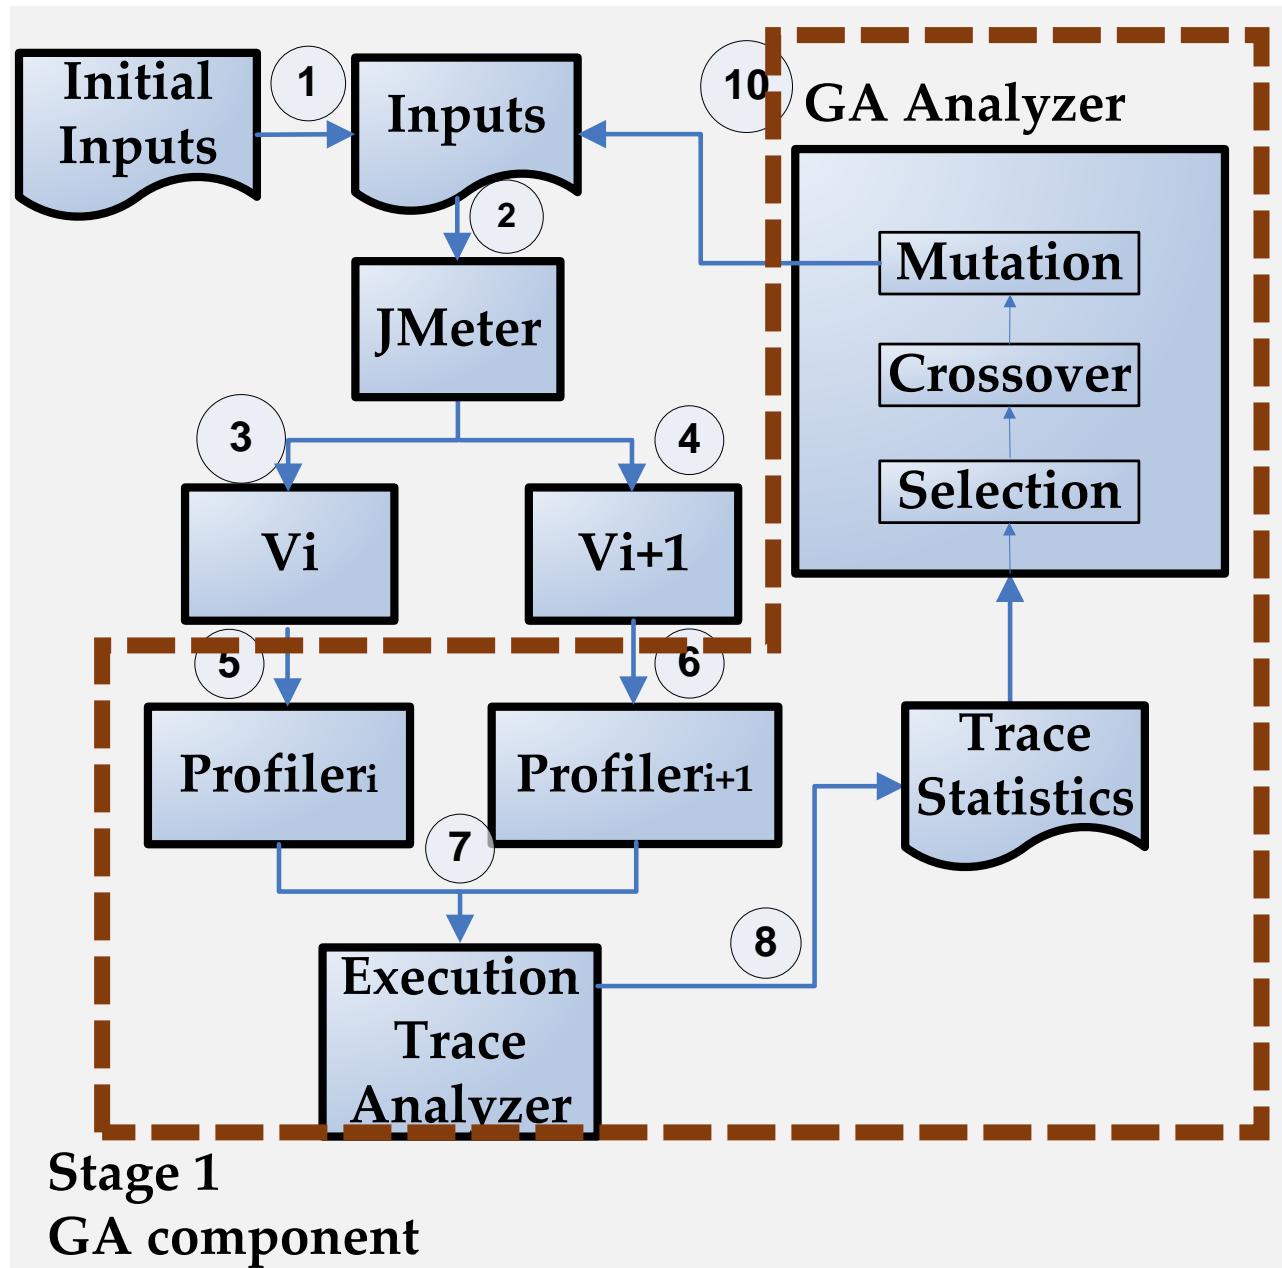
PerfImpact – performance regression testing

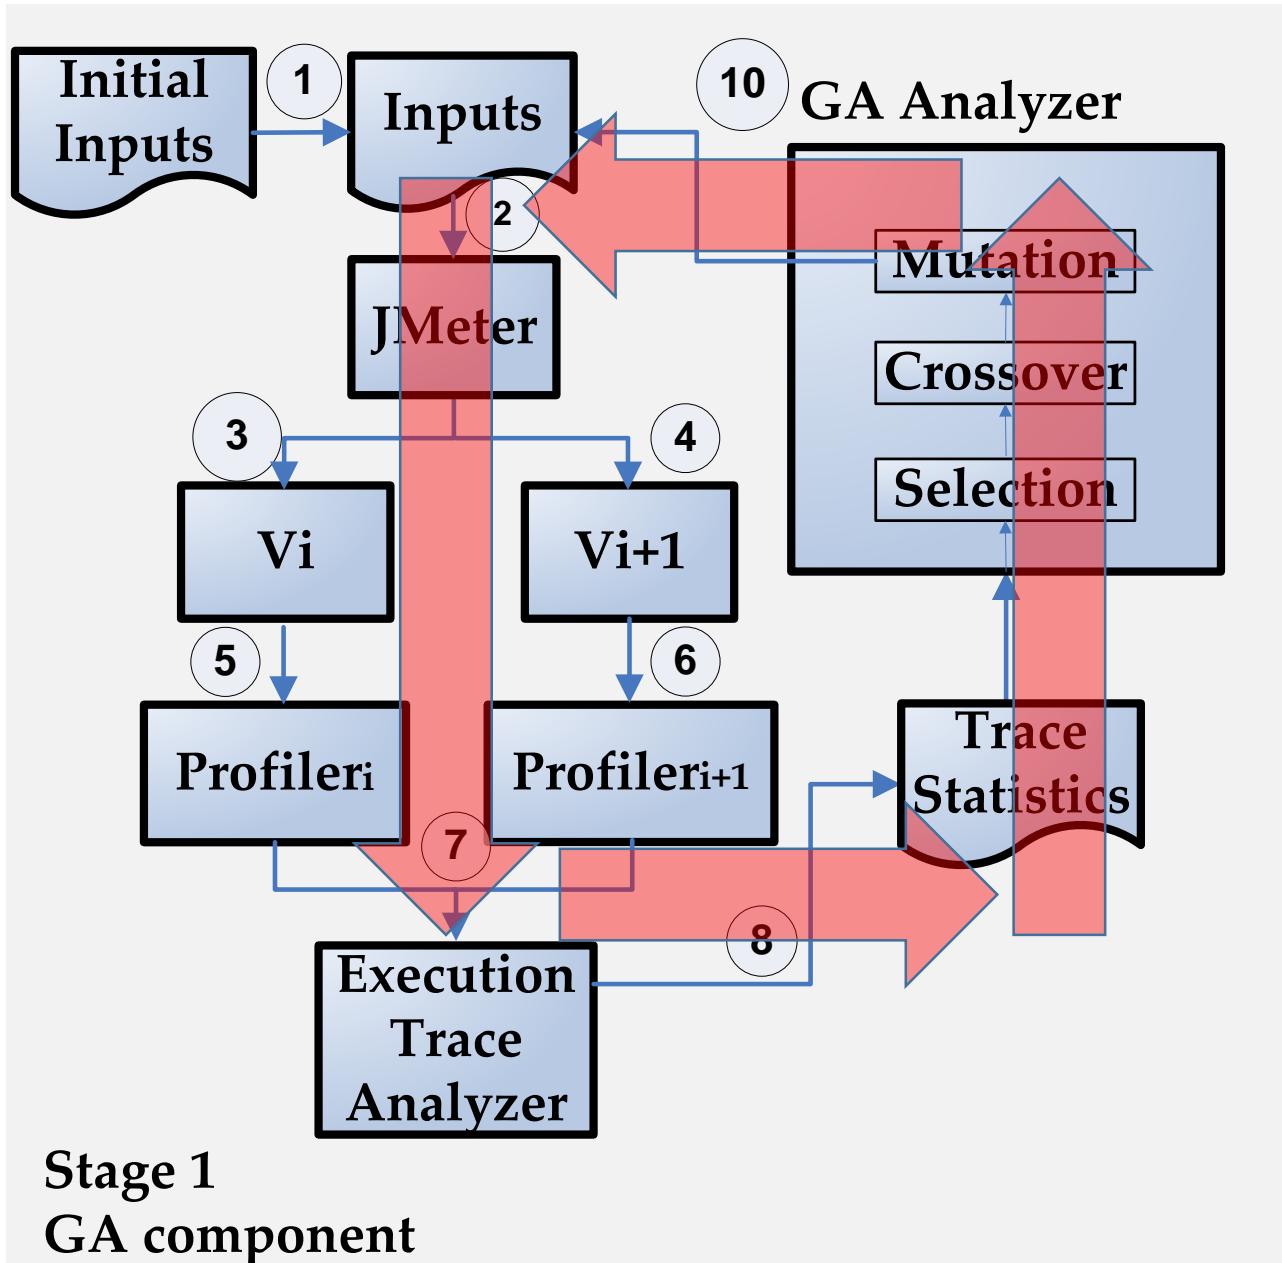
- Search-based input profiling to find specific inputs for exposing performance regressions
- Change impact analysis (CIA) to recommend the problematic code changes

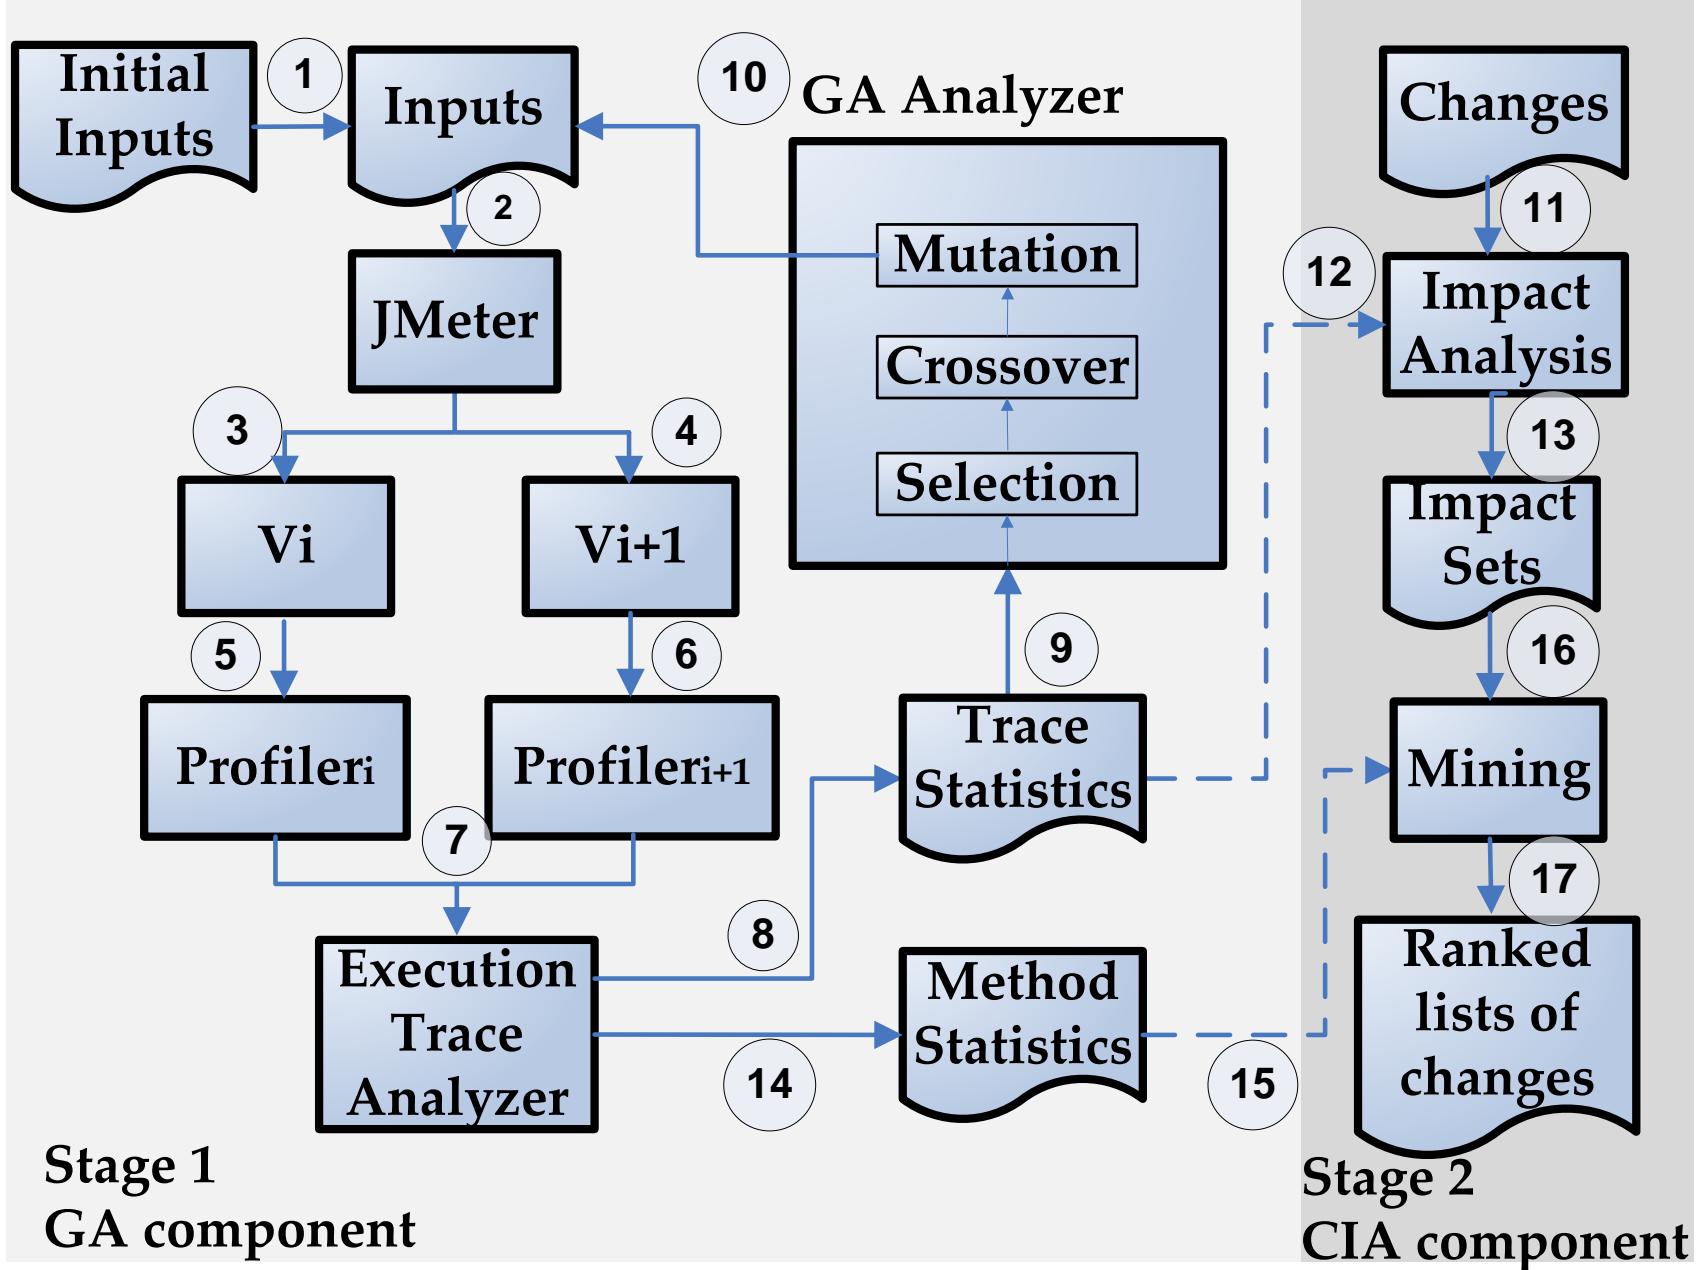


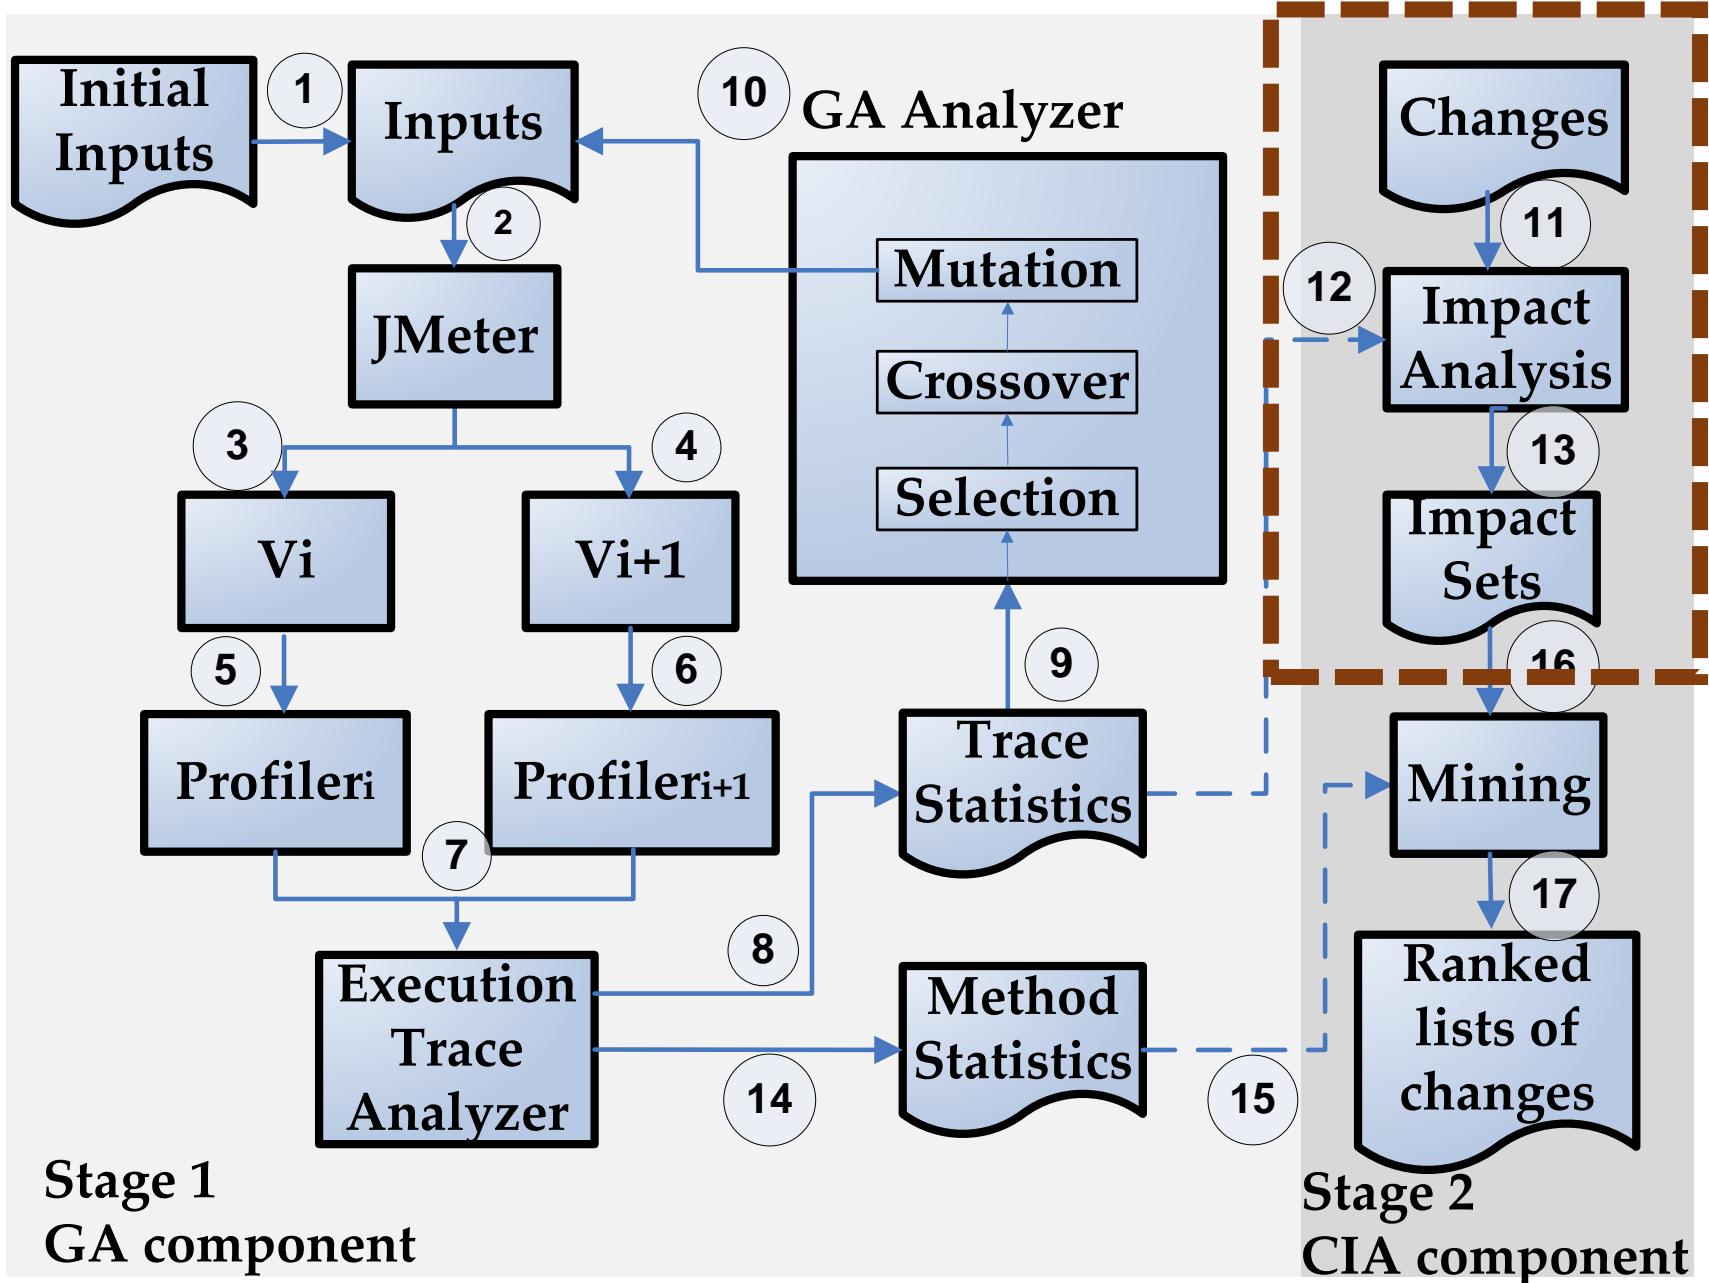


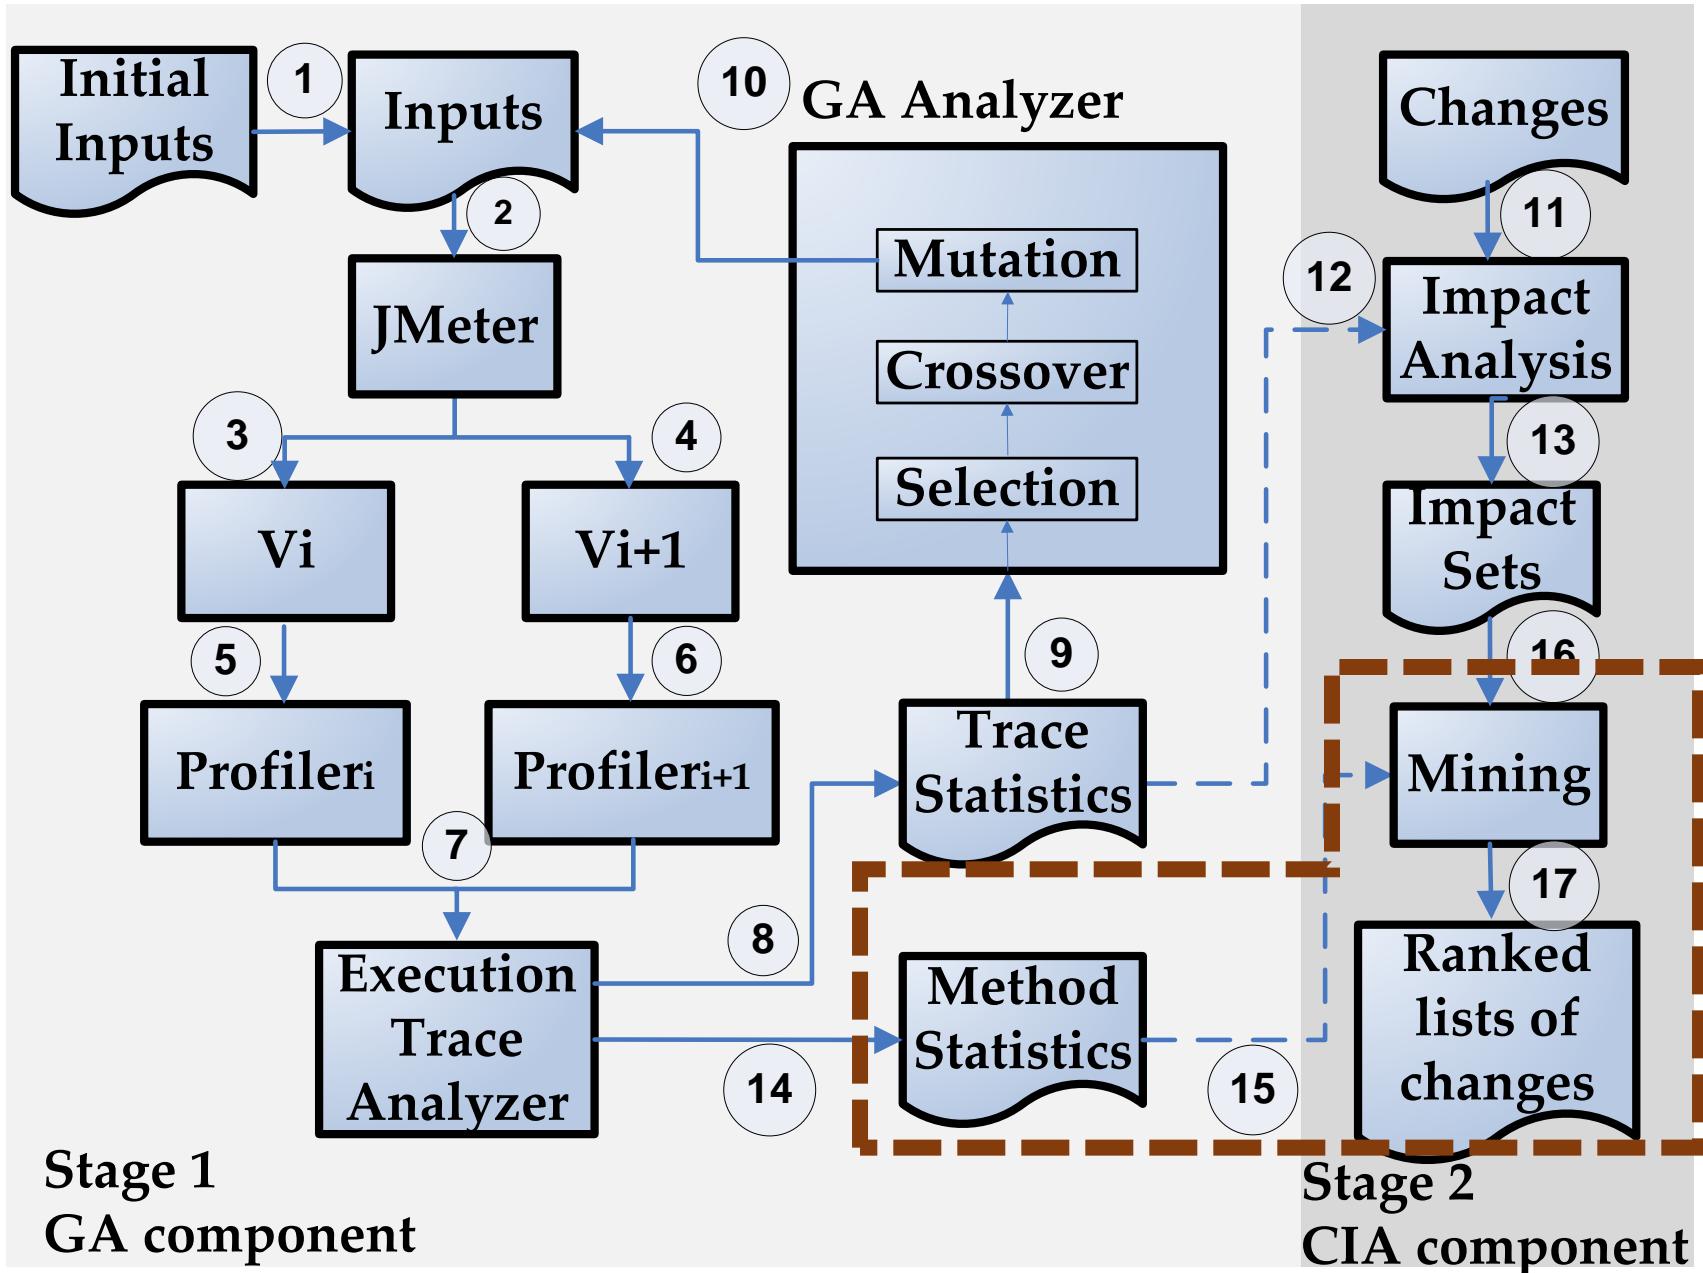


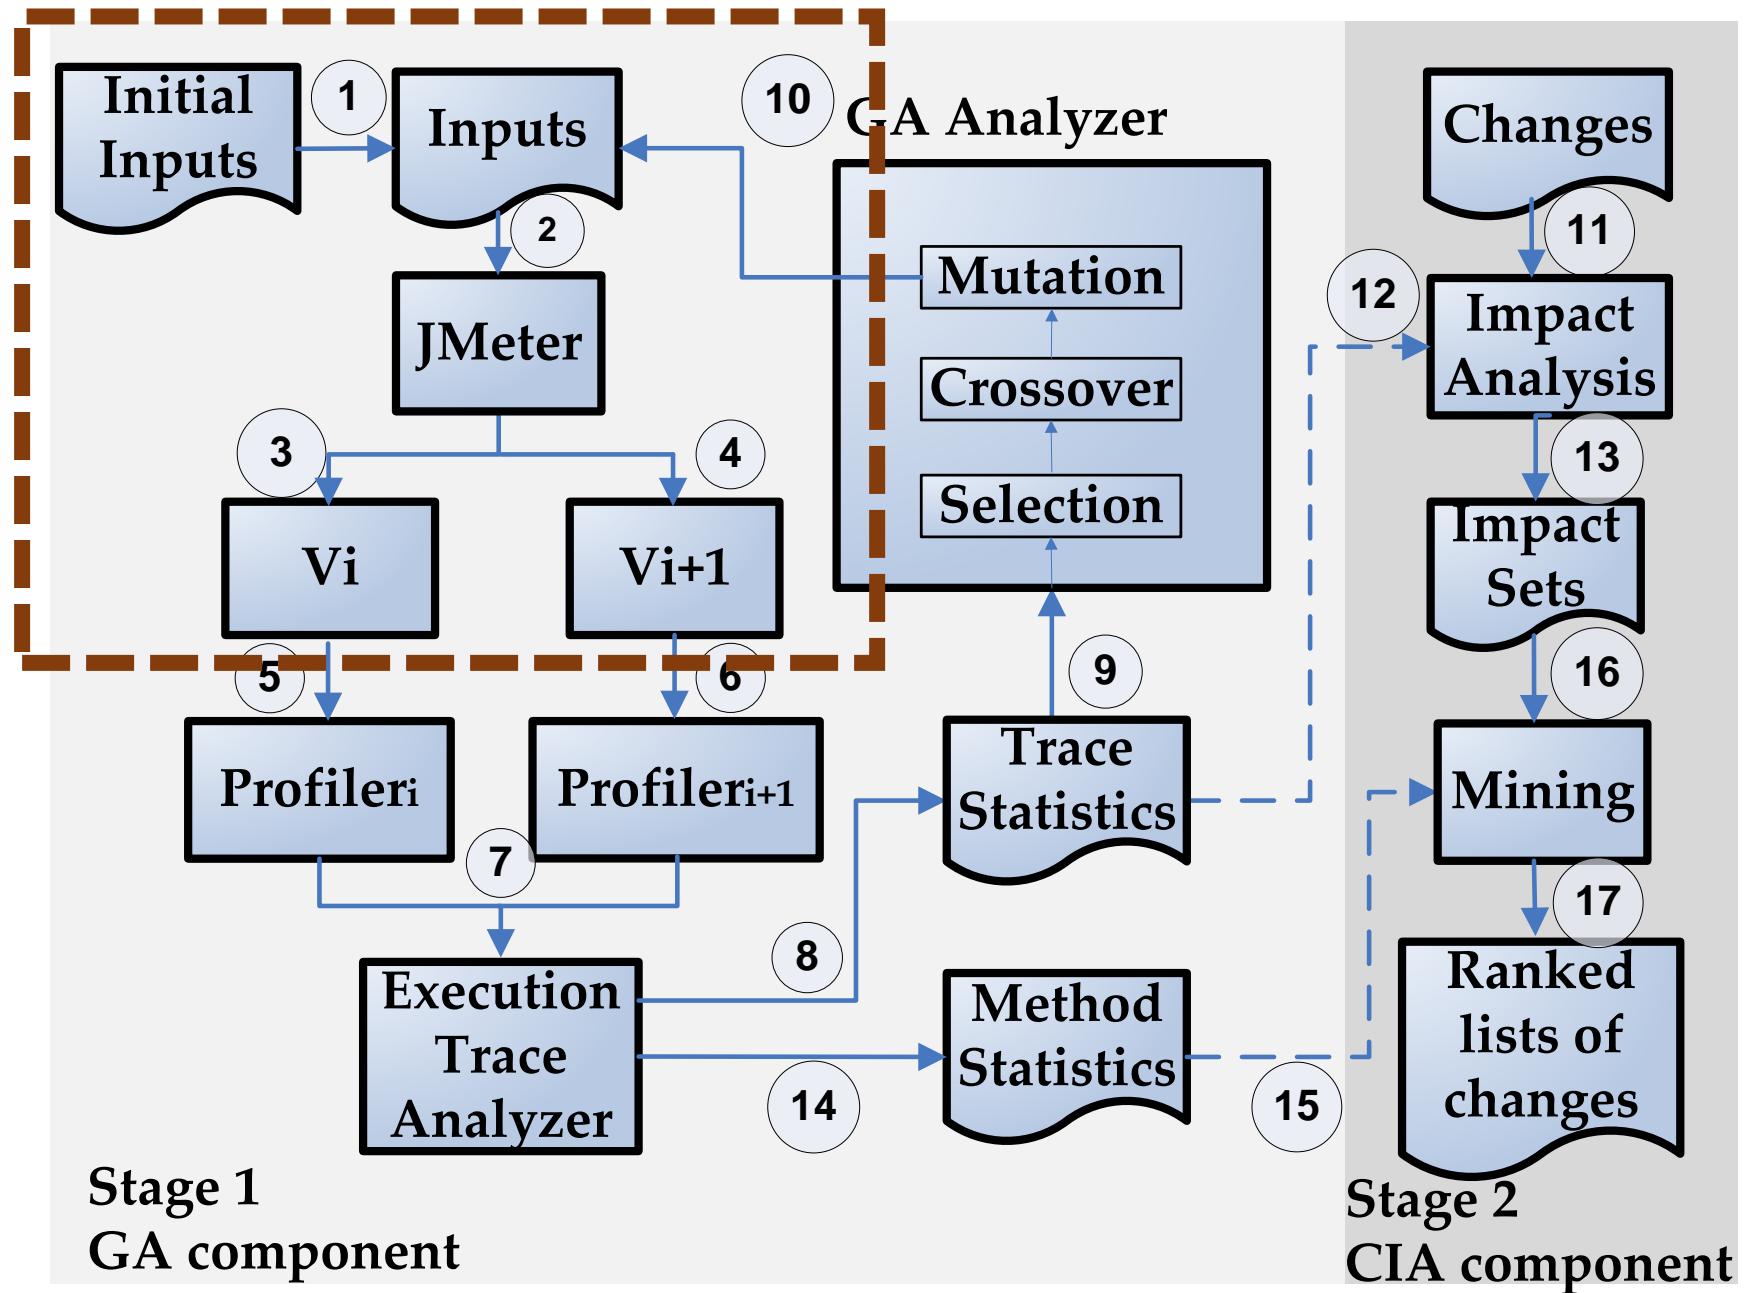


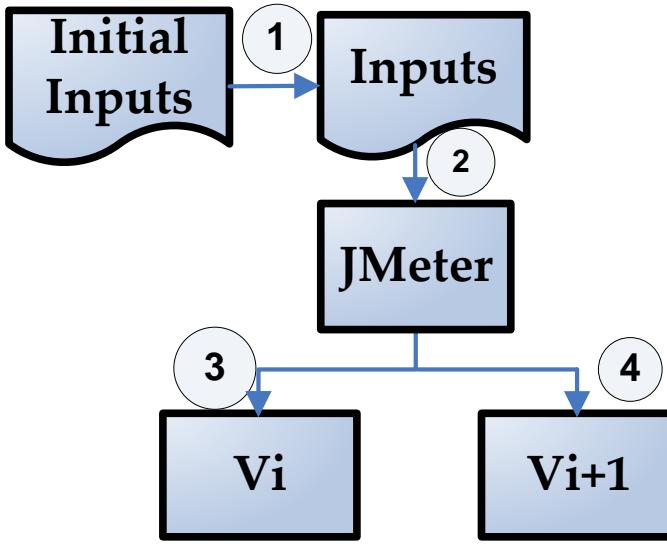












Genes:

Input 1: <http://localhost:8080/Agilefant/editUser.action>

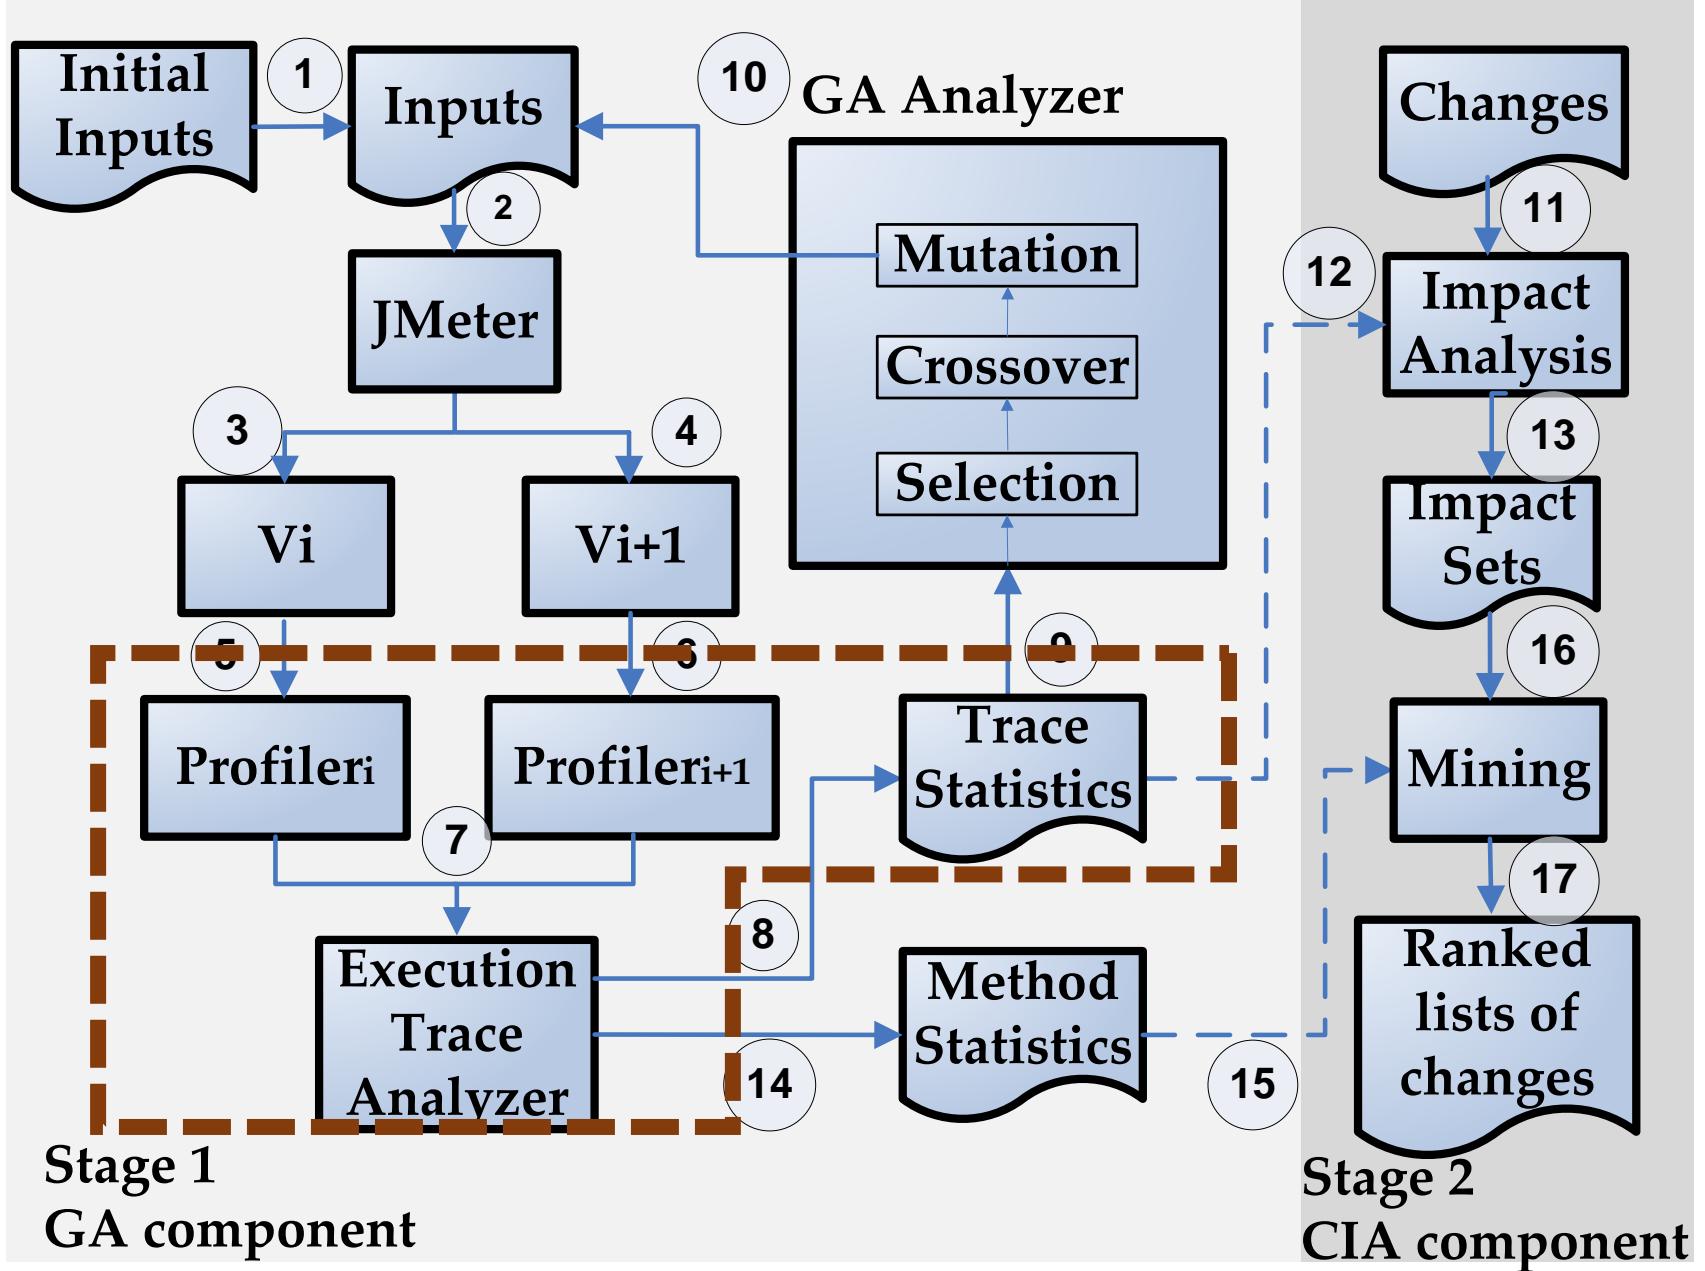
Input 2: <http://localhost:8080/Agilefant/editProduct.action?productId=5>

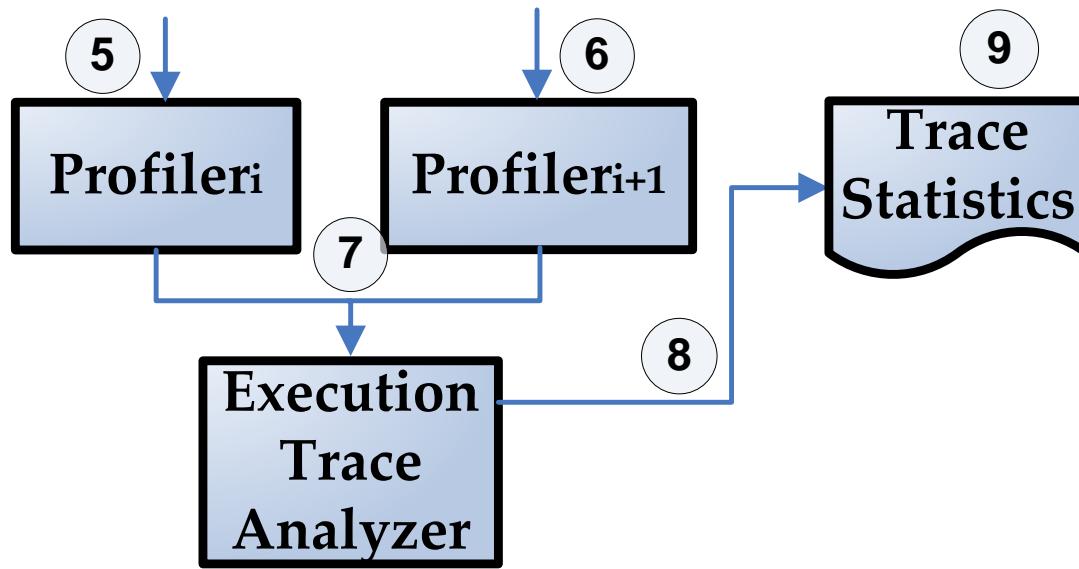
Input 3: <http://localhost:8080/Agilefant/editProduct.action?productId=8>

.....

A chromosome/individual

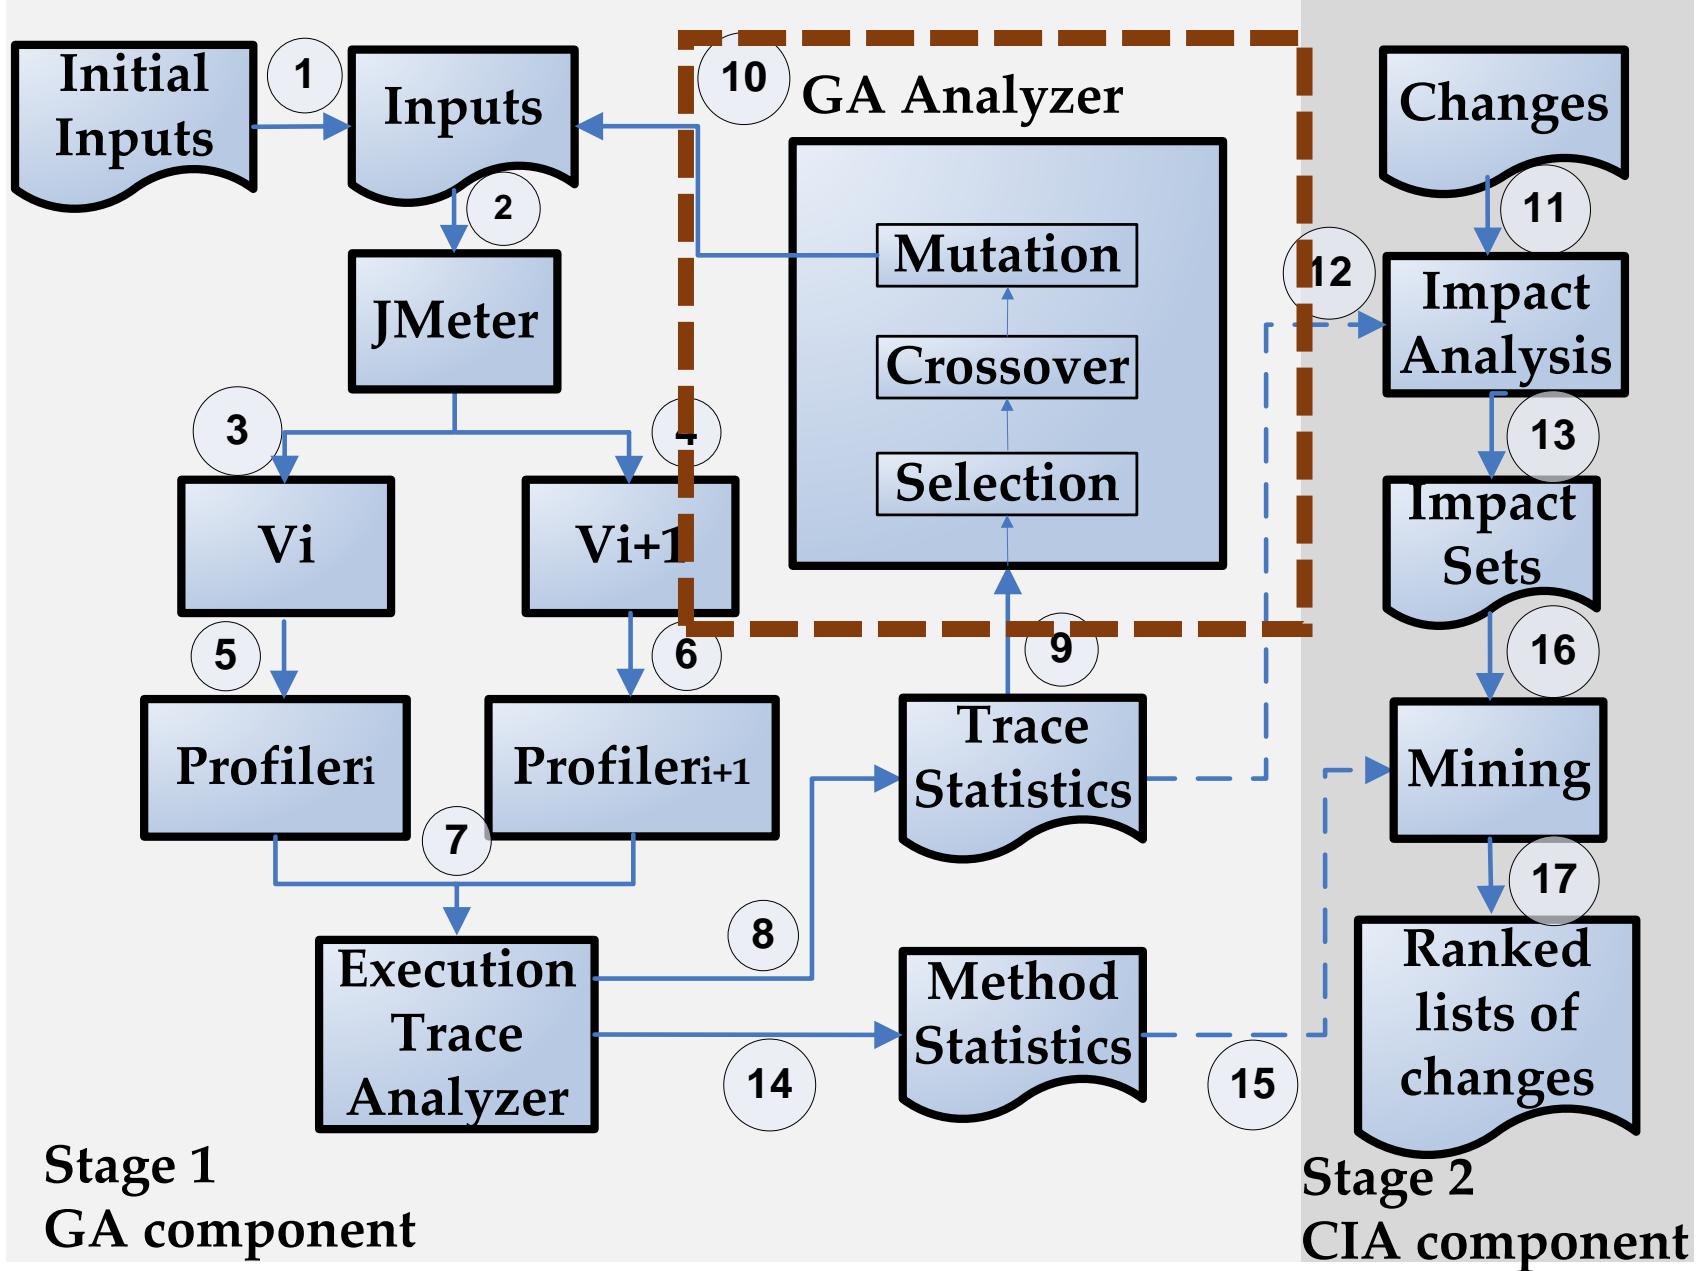
Individual 1: 2, 18, 36, 27, 11, 13, 6, 43, 64, 12, 85, 49, 12, 53, 44, 78, 31, 47



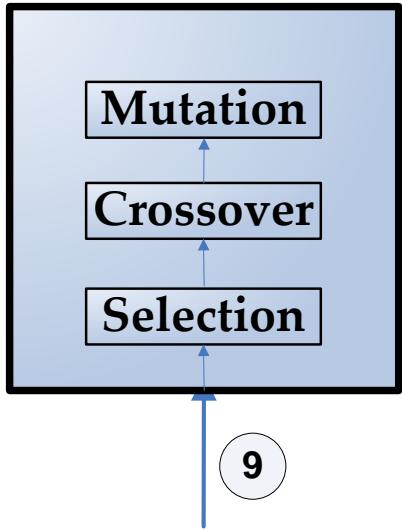


An individual (I_j)

- Execution time in V_i (t_j^i)
- Execution time in V_{i+1} (t_j^{i+1})



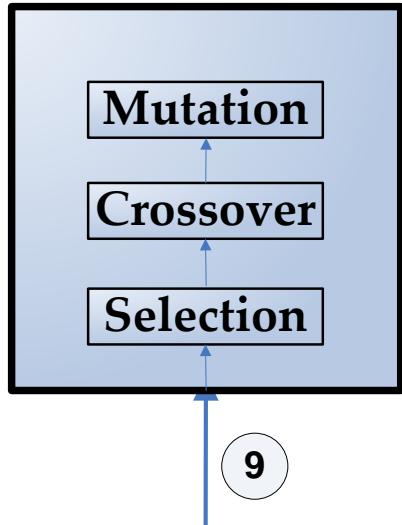
GA Analyzer



- Selection

Fitness value for $I_j = t_j^{i+1} - t_j^i$

GA Analyzer



- Selection

Fitness value for $I_j = t_j^{i+1} - t_j^i$

- Crossover

Parent 1: 2, 18, 36, 27, 11, 13, 6, 43, 64, 12, 85, 49, 12, 53, 44, 91, 79, 23, 3, 19

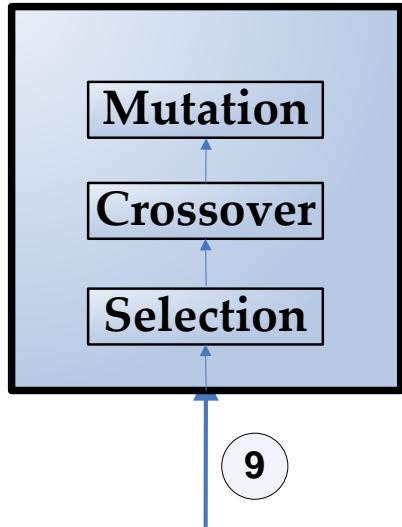
Parent 2: 23, 95, 1, 67, 35, 81, 7, 17, 51, 102, 56, 39, 72, 3, 54, 37, 13, 86, 47, 76



Child 1: 2, 18, 36, 27, 11, 13, 6, 17, 51, 102, 56, 39, 72, 3, 54, 37, 13, 86, 47, 76

Child 2: 23, 95, 1, 67, 35, 81, 7, 43, 64, 12, 85, 49, 12, 53, 44, 91, 79, 23, 3, 19

GA Analyzer



- Selection

Fitness value for $I_j = t_j^{i+1} - t_j^i$

- Crossover

Parent 1: 2, 18, 36, 27, 11, 13, 6, 43, 64, 12, 85, 49, 12, 53, 44, 91, 79, 23, 3, 19

Parent 2: 23, 95, 1, 67, 35, 81, 7, 17, 51, 102, 56, 39, 72, 3, 54, 37, 13, 86, 47, 76



Child 1: 2, 18, 36, 27, 11, 13, 6, 17, 51, 102, 56, 39, 72, 3, 54, 37, 13, 86, 47, 76

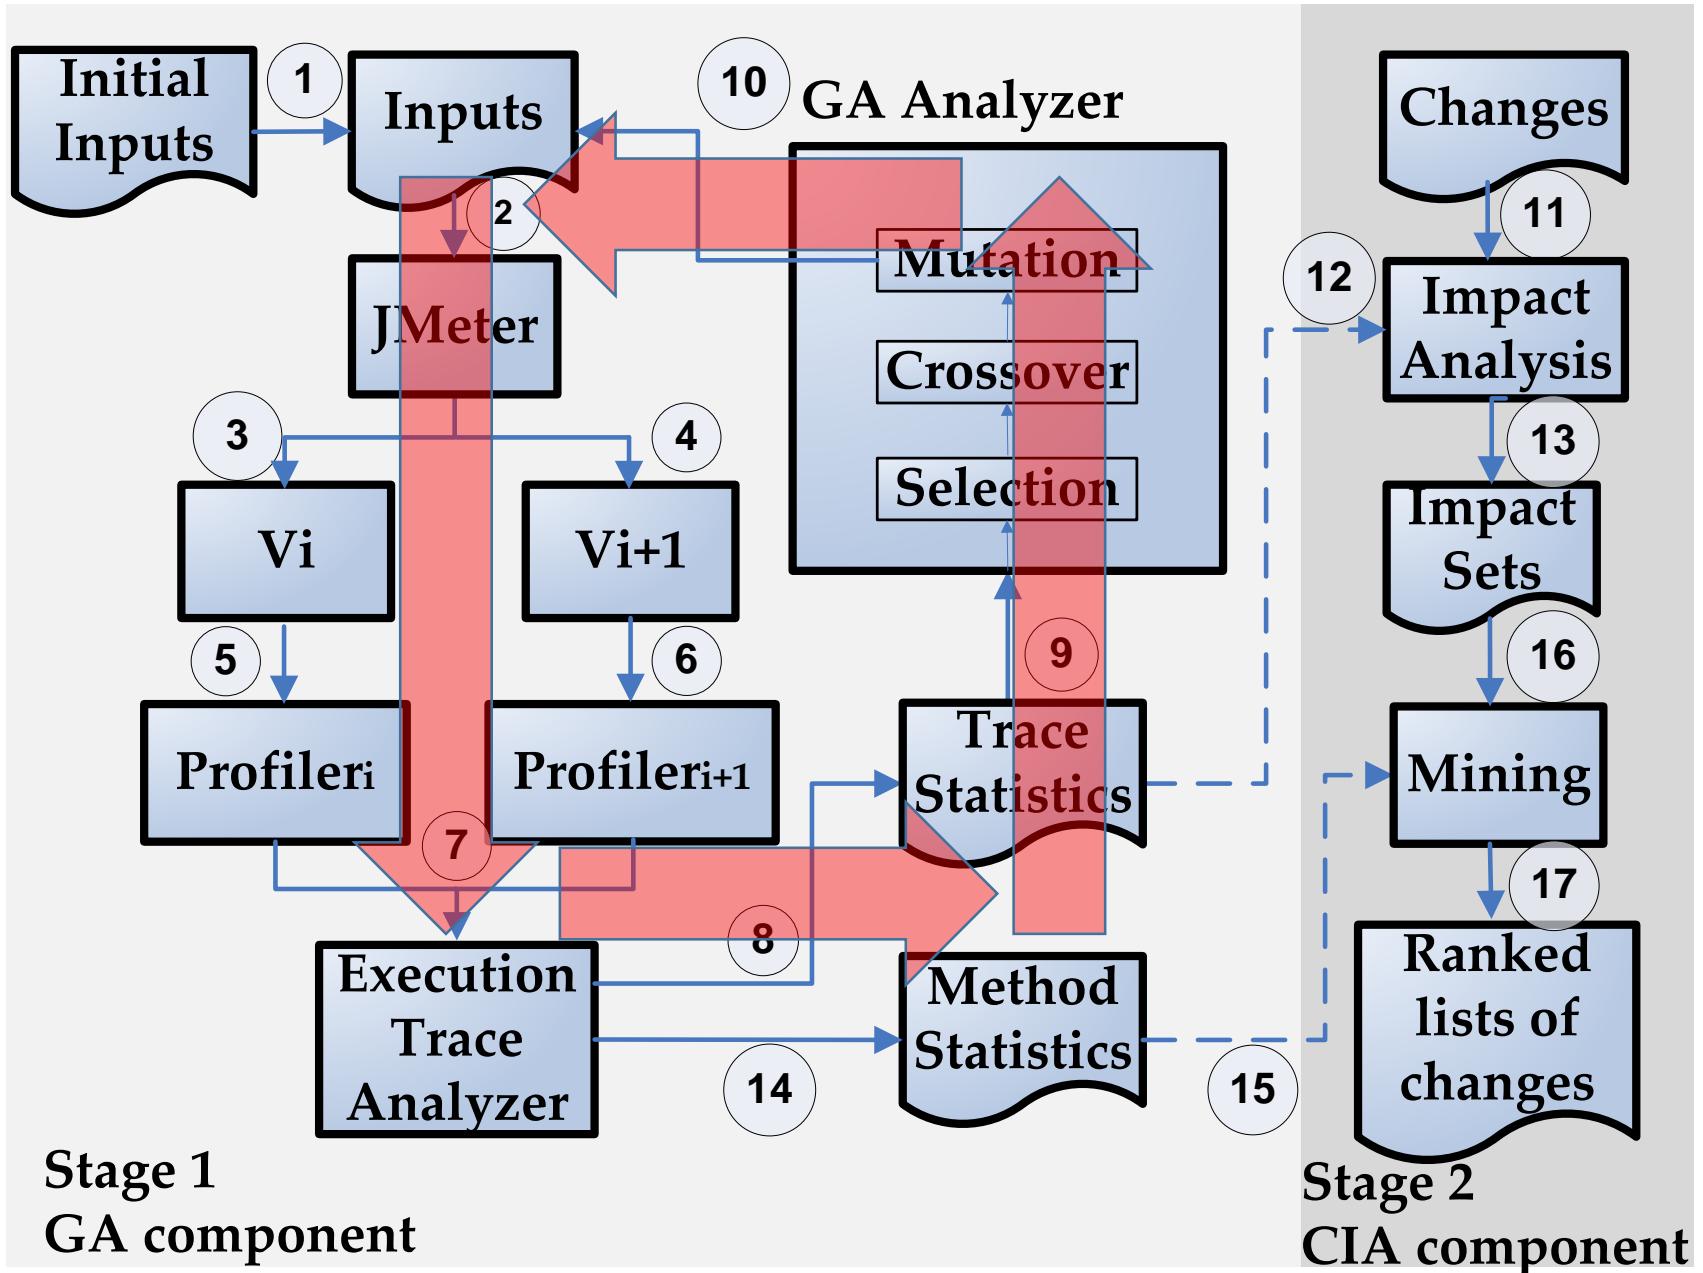
Child 2: 23, 95, 1, 67, 35, 81, 7, 43, 64, 12, 85, 49, 12, 53, 44, 91, 79, 23, 3, 19

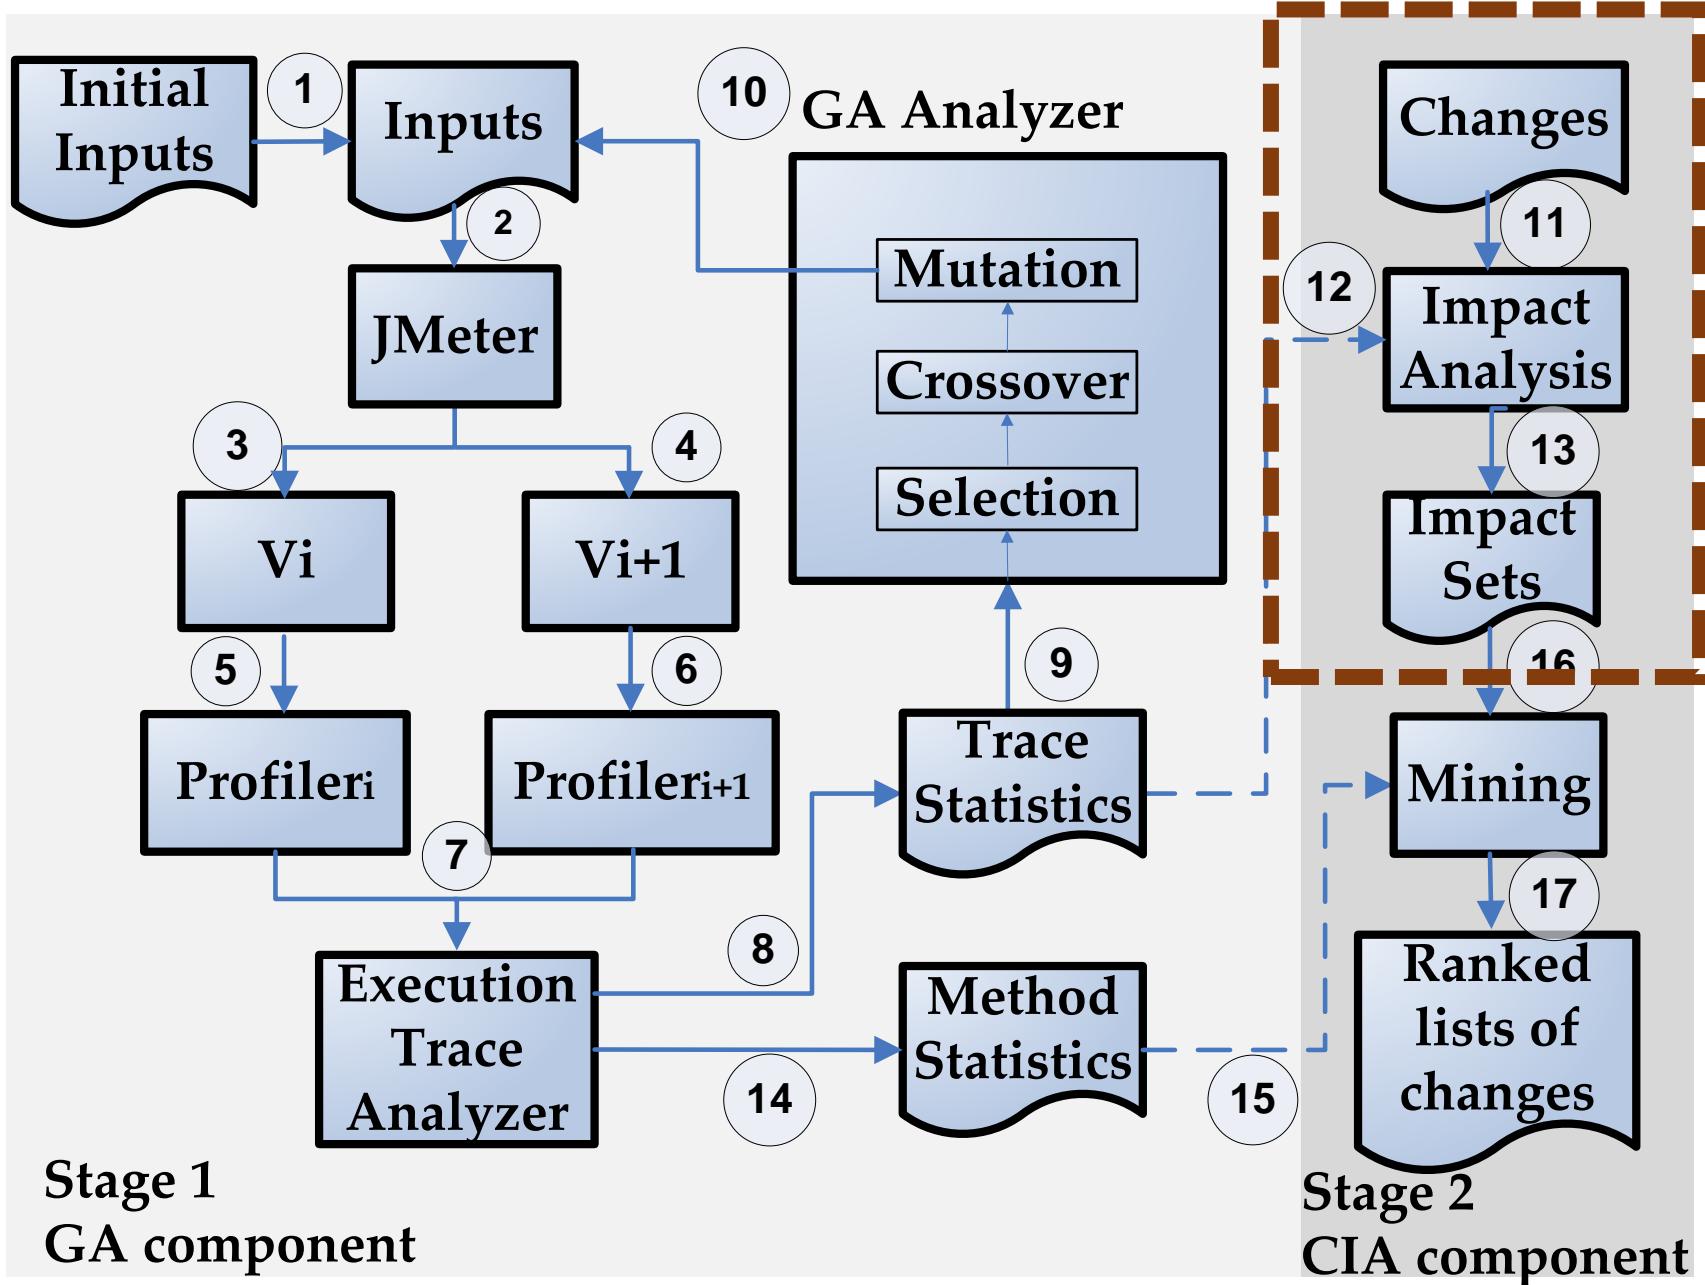
- Mutation

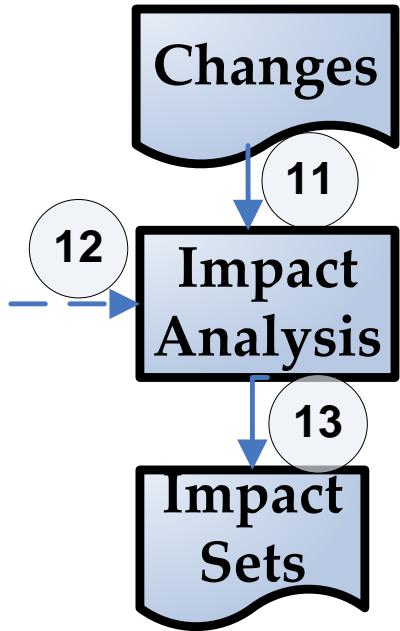
Parent : 2, 18, 36, 27, 11, 13, 6, 43, 64, 12, 85, 49, 12, 53, 44, 91, 79, 23, 3, 19



Child : 2, 18, 36, 27, 11, 13, 6, 43, 64, 73, 85, 49, 12, 53, 44, 91, 79, 23, 3, 19







Change Impact:

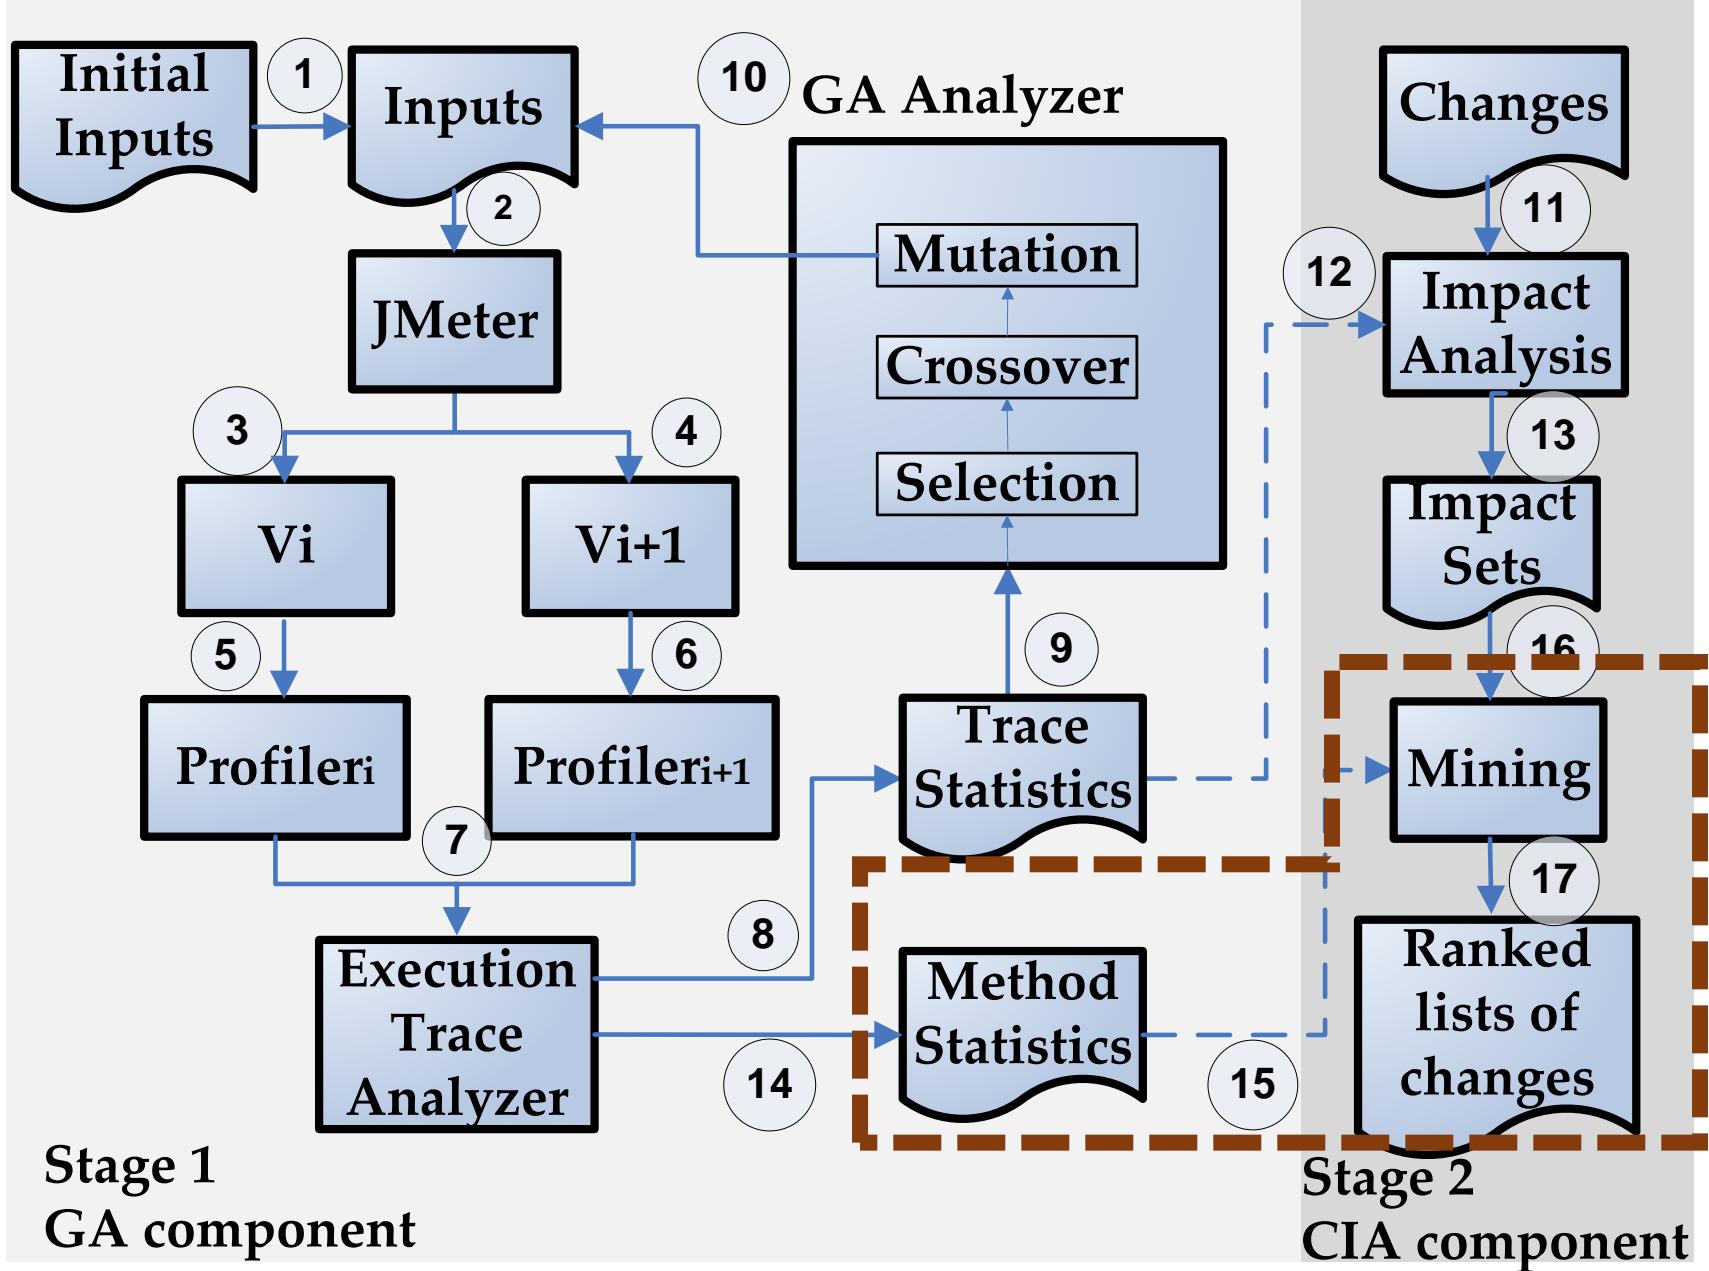
propagates along any (and only)
dynamic paths that pass through the
change (Law et al. ICSE'03)

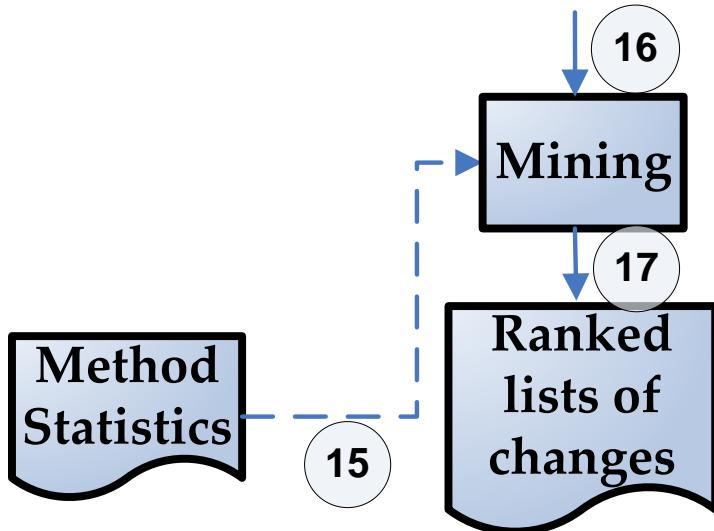
Impact Sets:

$$C_A \Rightarrow M_A, M_D, M_Q, \dots$$

$$C_B \Rightarrow M_B, M_D, M_N, \dots$$

...





Impact Sets:

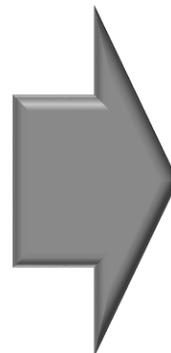
$C_A \Rightarrow M_A, M_D, M_Q, \dots$

12 6 91 ms

$C_B \Rightarrow M_B, M_D, M_N, \dots$

37 6 45 ms

...



Ranked list:

- C_B
- C_A
- ...

Research Questions (RQs)



- RQ₁ - How effective is PerfImpact in finding inputs that likely expose performance regressions
- RQ₂ – Can PerfImpact effectively recommend changes likely responsible for performance regressions

Research Question 1

How effective is PerfImpact in finding inputs that likely expose regressions

→ PerfImpact vs. Random

H_0 : There is no statistically significant difference between PerfImpact and Random



Research Question 2

How effective is PerfImpact in identifying regression-inducing code changes

- Inject nine artificial code changes
- Extract real code changes



Experimental Design

Agilefant



V3.2 vs V3.3

V3.2 vs V 3.5

JPetStore



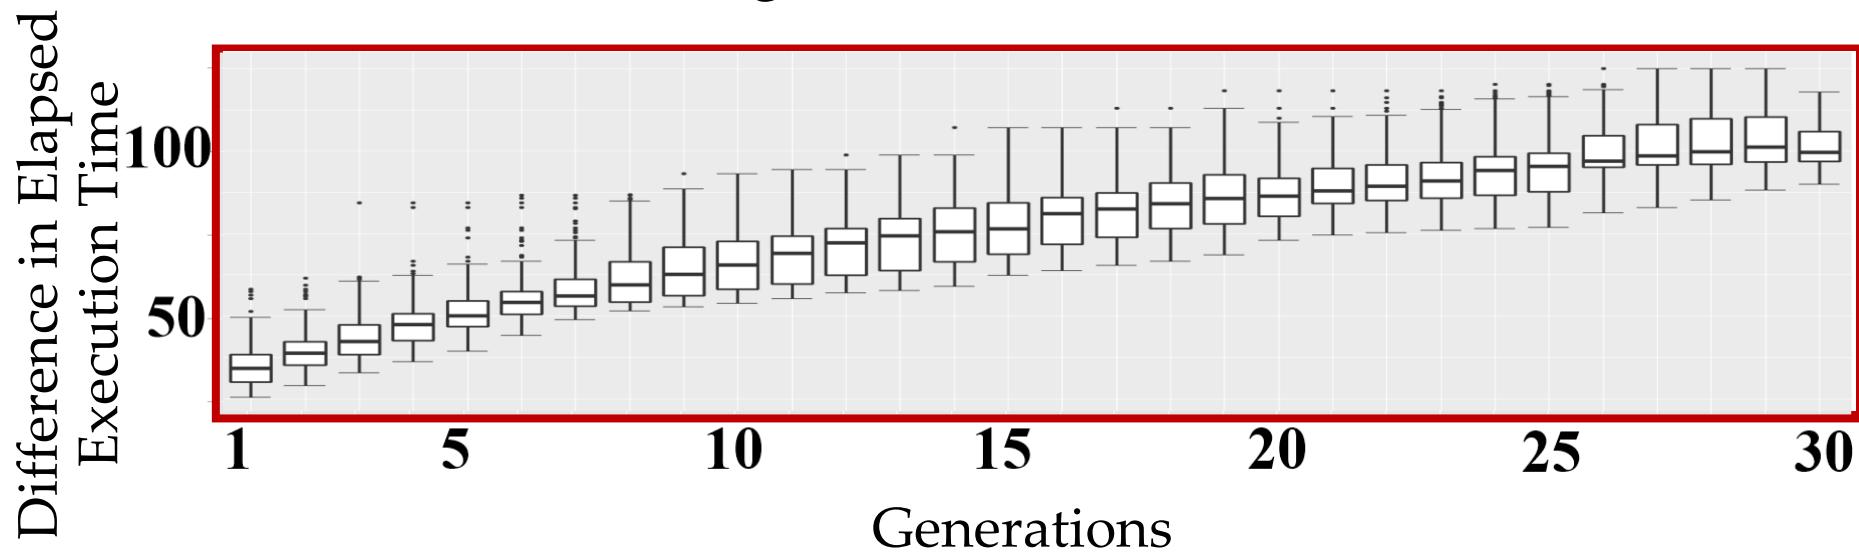
V3.0.0 vs V4.0.5



Total: 187 real changes + 9 artificial changes

RQ₁ – Finding Regression-Specific Inputs

Agilefant V3.2 vs V3.3

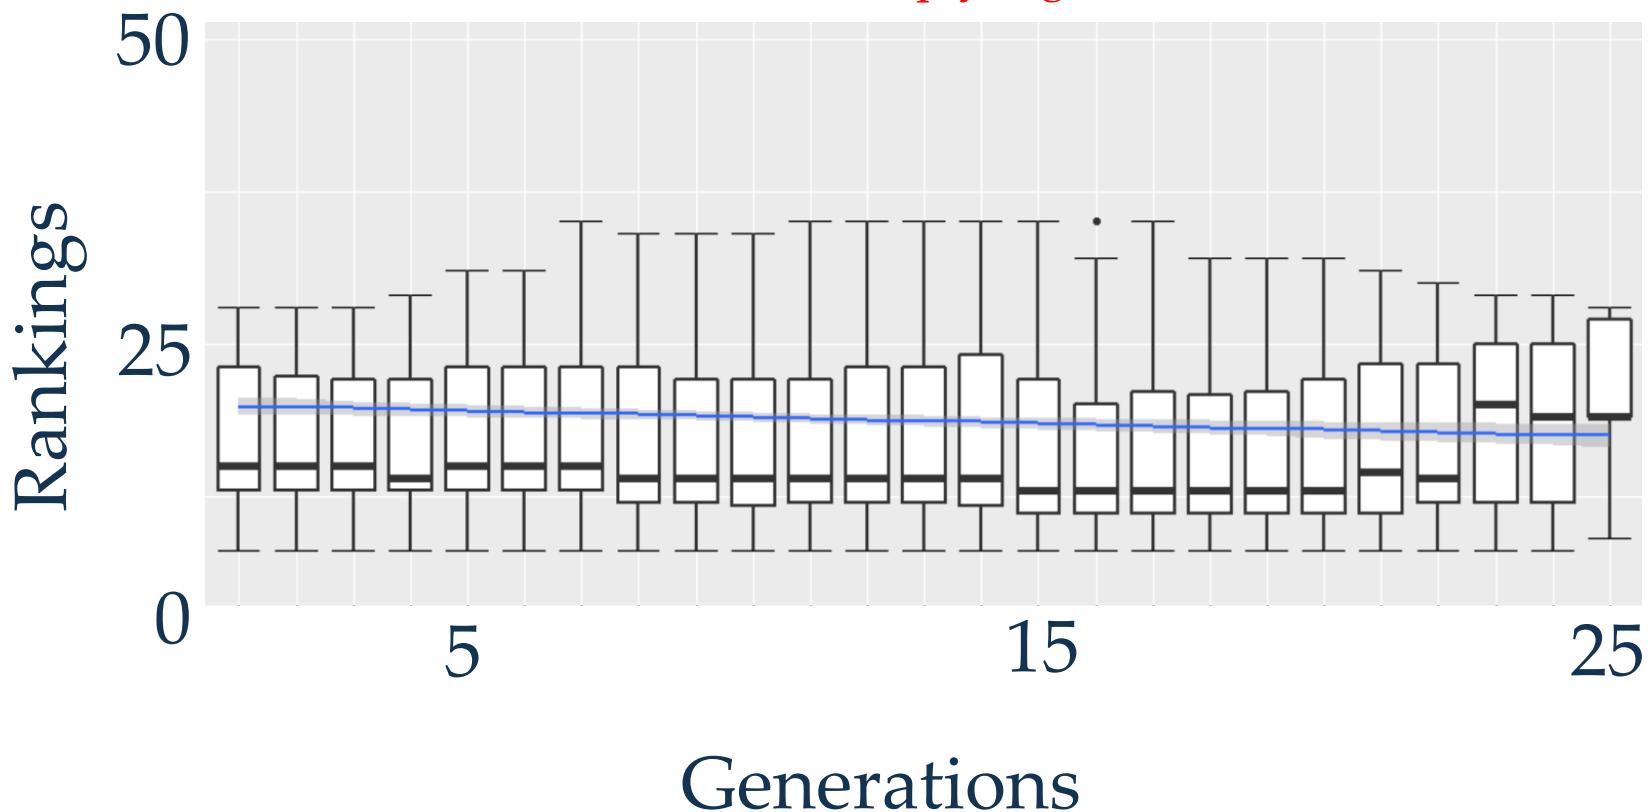


$p < 1.23e - 296 \Rightarrow$ null hypothesis is rejected

RQ₂ -Performance regression in a real code change

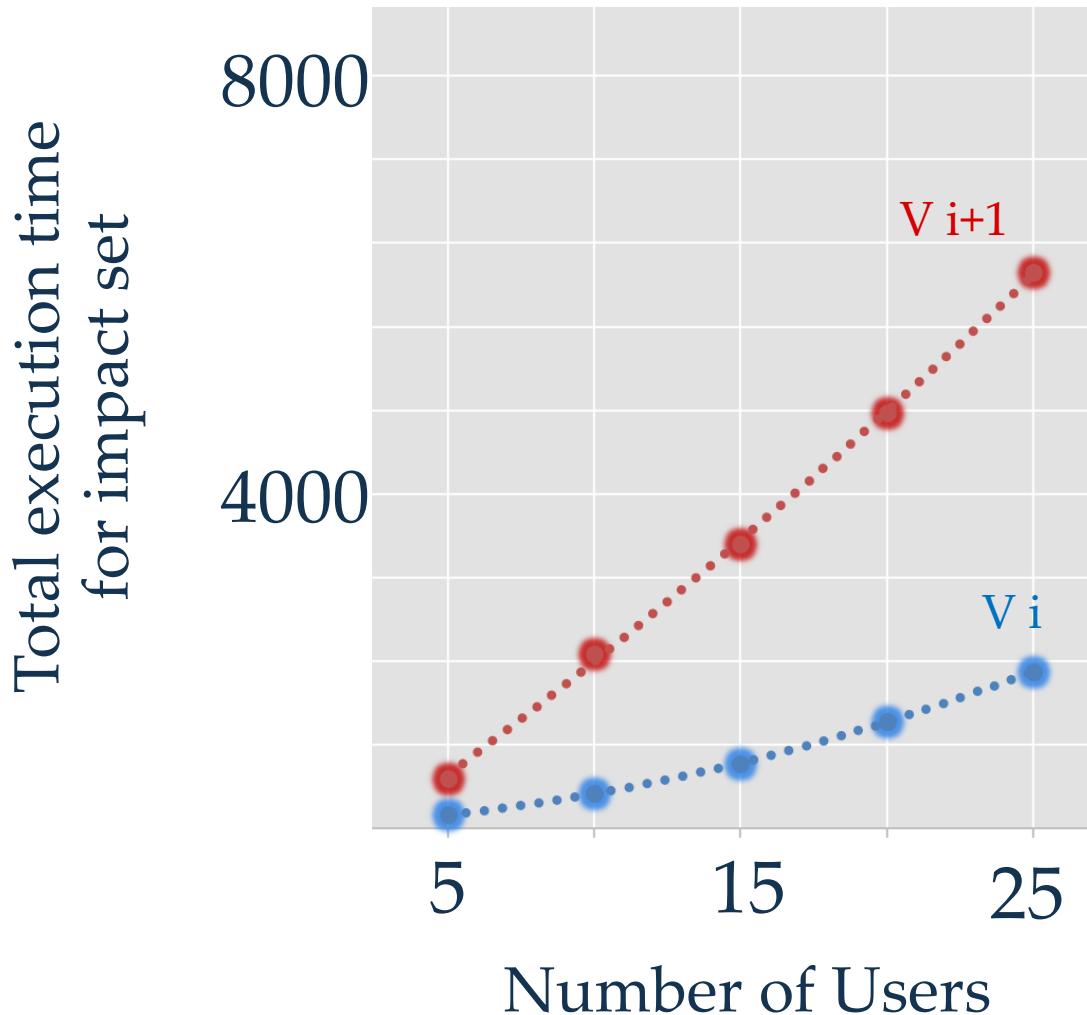
Code change: ProjectBusinessImpl.retrieveLeafStories()

Smaller values imply higher ranks



RQ₂ -Performance regression in a real code change

Code change: ProjectBusinessImpl.retrieveLeafStories()



RQ₂ -Performance regression in a real code change

Code change: ProjectBusinessImpl.retrieveLeafStories()

```
public List<StoryTO> retrieveLeafStories(int projectId, StoryFilters filters) {  
    .....  
    for(Story leafStory : leafStories) {  
        StoryTO tmp = new StoryTO(leafStory); .....  
        Set<Task> tasks = new HashSet<Task>();  
        for (Task task : tmp.getTasks()) {  
            TaskTO taskTO = new TaskTO(task);  
            tasks.add(taskTO);  
        }  
        tmp.setTasks(tasks);  
        .....  
    }  
    return leafStoriesWithRank;  
}
```

Mining Performance Regression Inducing Code Changes in Evolving Software

Qi Luo

College of William and Mary
qluo@cs.wm.edu

Denys Poshyvanyk

College of William and Mary
denys@cs.wm.edu

Mark Grechanik

University of Illinois at Chicago
drmark@uic.edu

ABSTRACT

During software evolution, the source code of a system frequently changes due to bug fixes or new feature requests. Some of these changes may accidentally degrade performance of a newly released software version. A notable problem of regression testing is how to find problematic changes

1. INTRODUCTION

Performance is an important metric of software quality [60, 43], whereas performance testing is a vital activity that developers routinely perform during software development and maintenance to ensure quality [19]. During software evolution, a number of code changes are committed, and

Online appendix:

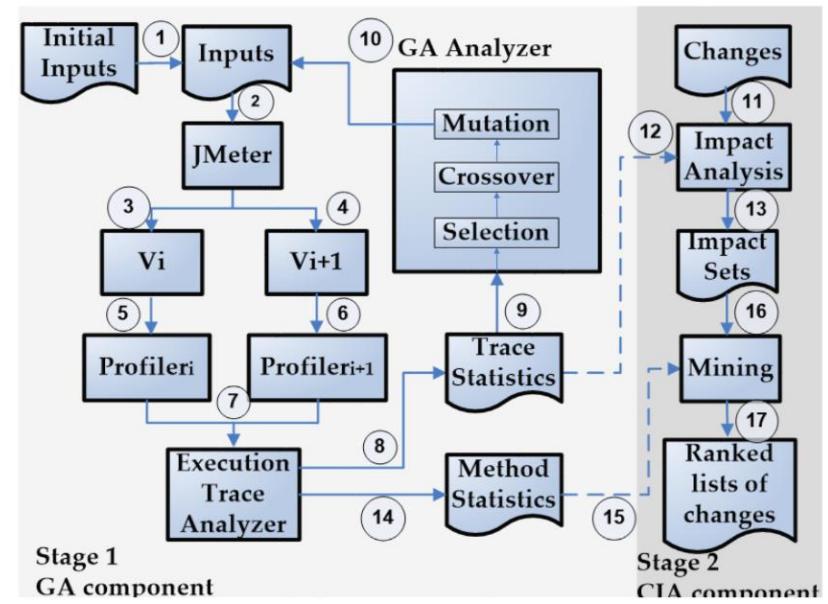
<http://www.cs.wm.edu/semeru/data/MSR16-PerfImpact/>

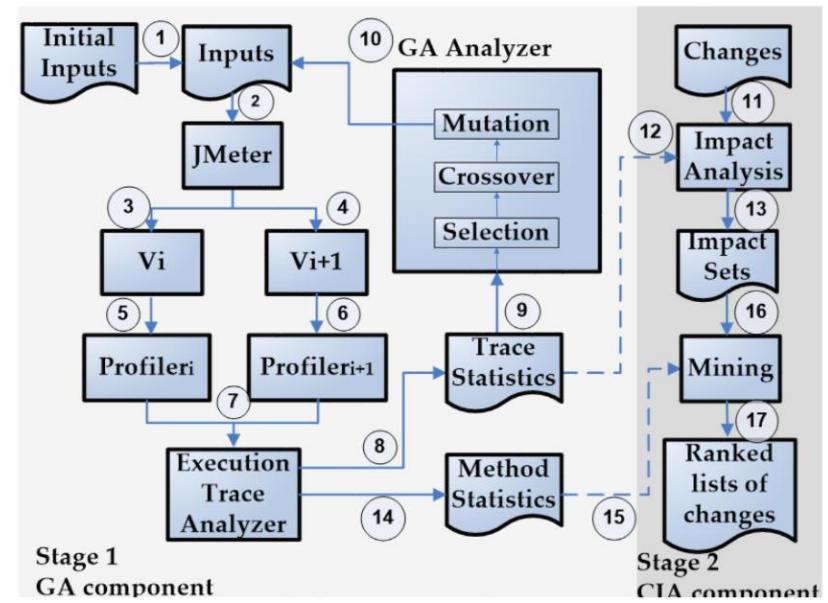
```
public synchronized void regressionFunc();  
static Regressions staticRT = new Regressions();  
public void setTeams(Collection<Team> teams) {  
    this.teams = teams;  
    Backlog.staticRT.regressionFunc();  
}
```

```
Set<Integer> selectedBacklogIds = this.getSelectedBacklogs();  
if(selectedBacklogIds == null || selectedBacklogIds.size() == 0) {  
    Collection<Product> products = new ArrayList<Product>();  
    productBusiness.storeAllProducts(products);  
    for (Product product: products) {  
        selectedBacklogIds.add(product.getId());  
    }  
}  
return Action.SUCCESS;
```

```
for(Story child : story.getChildren()) {  
    if (child.getId() == story.getId()) {  
        continue;  
    }  
    storyTreeBranchMetrics.childMetrics = this.calculateStoryTreeMetrics(child);  
}  
return metrics;
```







Experimental Design



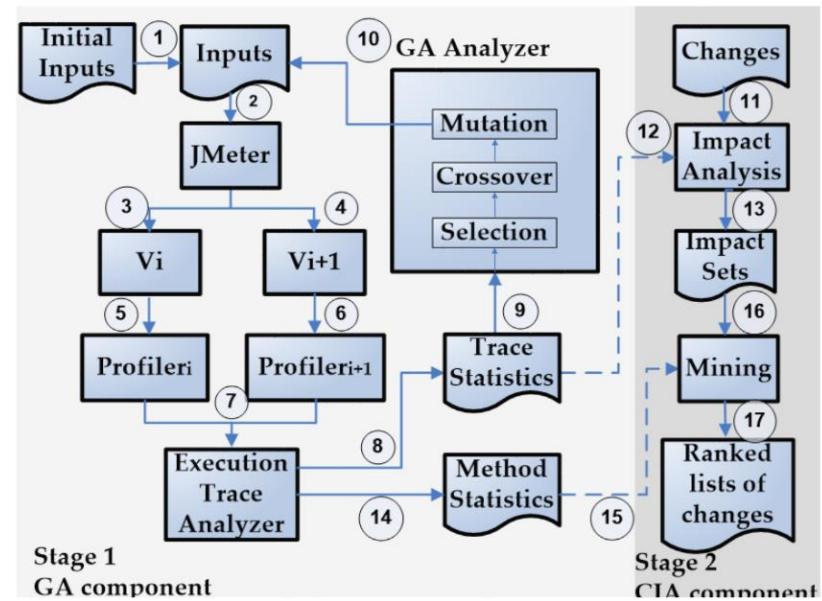
V3.2 vs V3.3
V3.2 vs V 3.5



V3.0.0 vs V4.0.5



Total: 187 real changes + 9 artificial changes



Experimental Design



V3.2 vs V3.3
V3.2 vs V 3.5



V3.0.0 vs V4.0.5



Total: 187 real changes + 9 artificial changes

PerImpact

- PerImpact is effective in finding specific inputs that expose performance regressions
- PerImpact is effective in identifying regression-inducing code changes

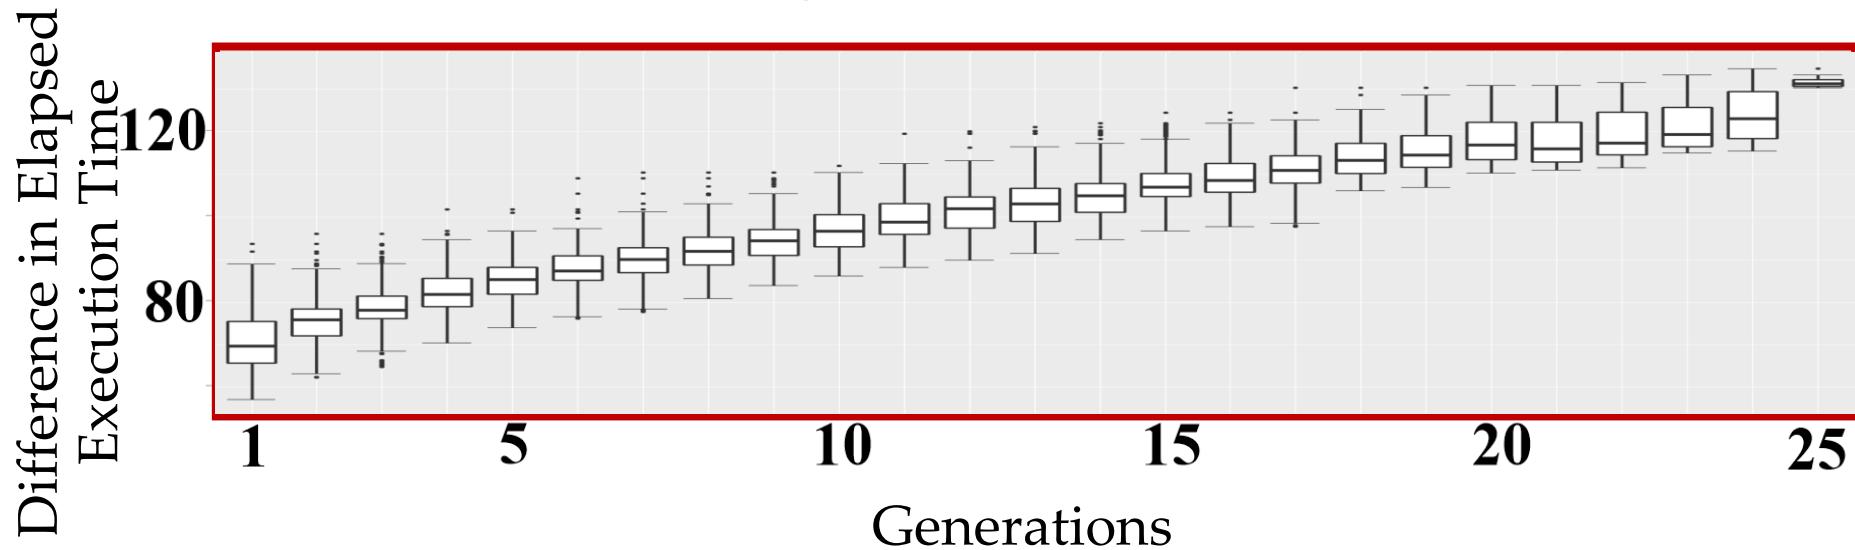
Additional Slides for Questions

GAs – Independent Variables

- Crossover rate – 0.3
- Mutation rate – 0.1
- Number of individuals per generation – 30
- Termination Criterion
 - Maximum limit for the number of generations – 30
 - Average fitness value of every individual in one generation

RQ₁ – Finding Regression-Specific Inputs

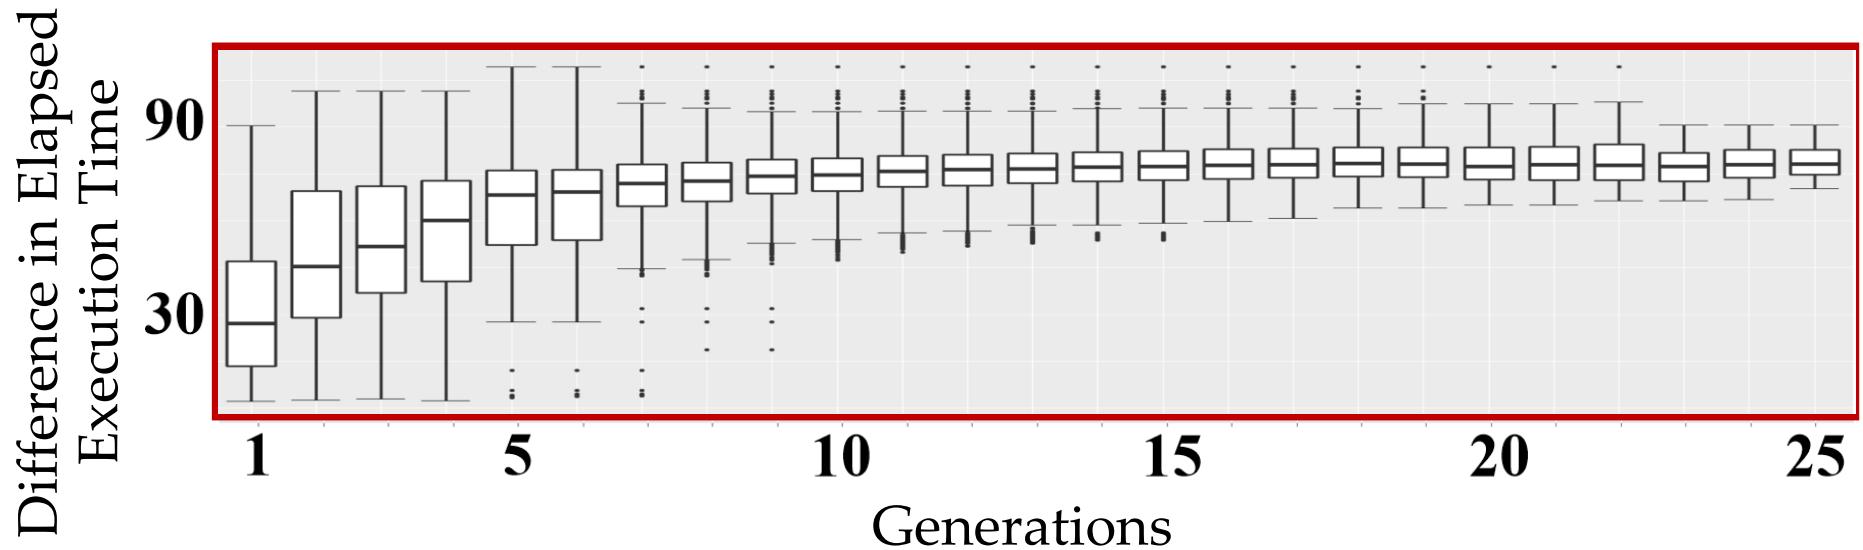
Agilefant V3.2 vs V3.5



$p = 1.37e - 236 \Rightarrow$ null hypothesis is rejected

RQ₁ – Finding Regression-Specific Inputs

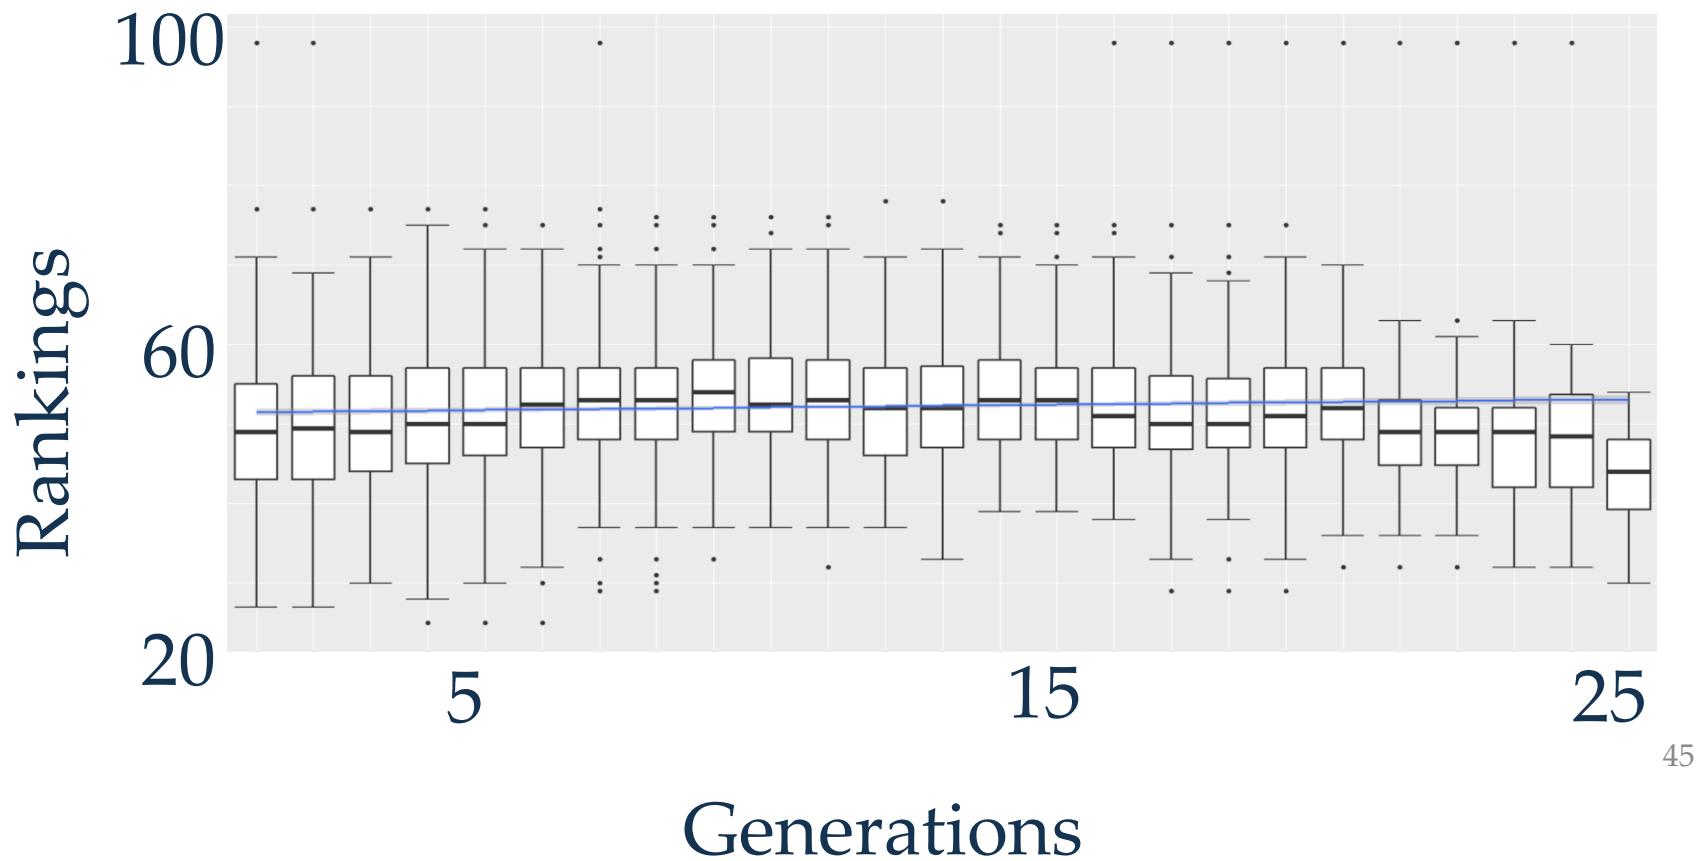
JPetStore V3.0.0 vs V4.0.5



$$p = 2.64e - 198 \Rightarrow \text{null hypothesis is rejected}$$

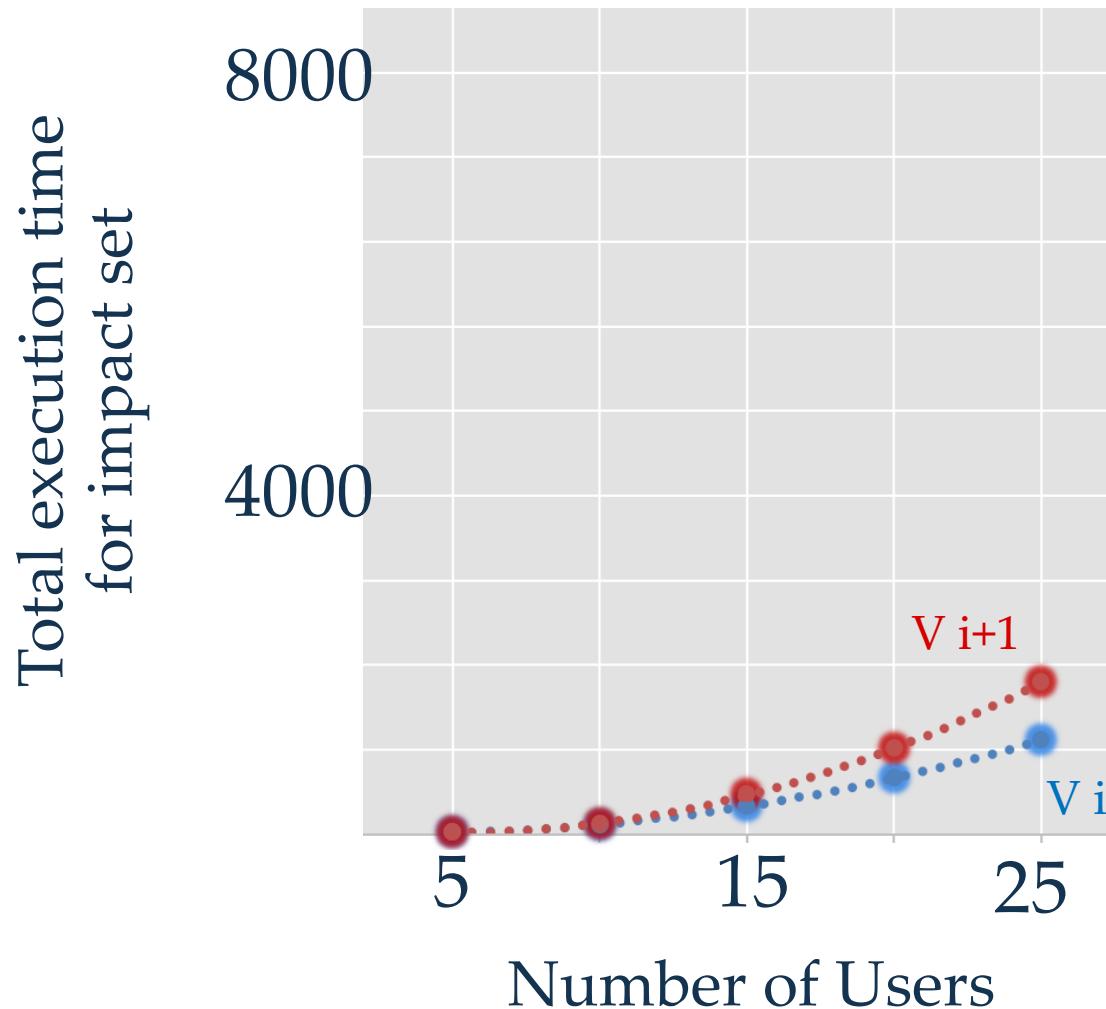
RQ₂ - Performance regression in a real code change

Code change: StoryHierarchyBusinessImpl.calculateStoryTreeMetrics()



RQ₂ -Performance regression in a real code change

Code change: StoryHierarchyBusinessImpl.calculateStoryTreeMetrics()

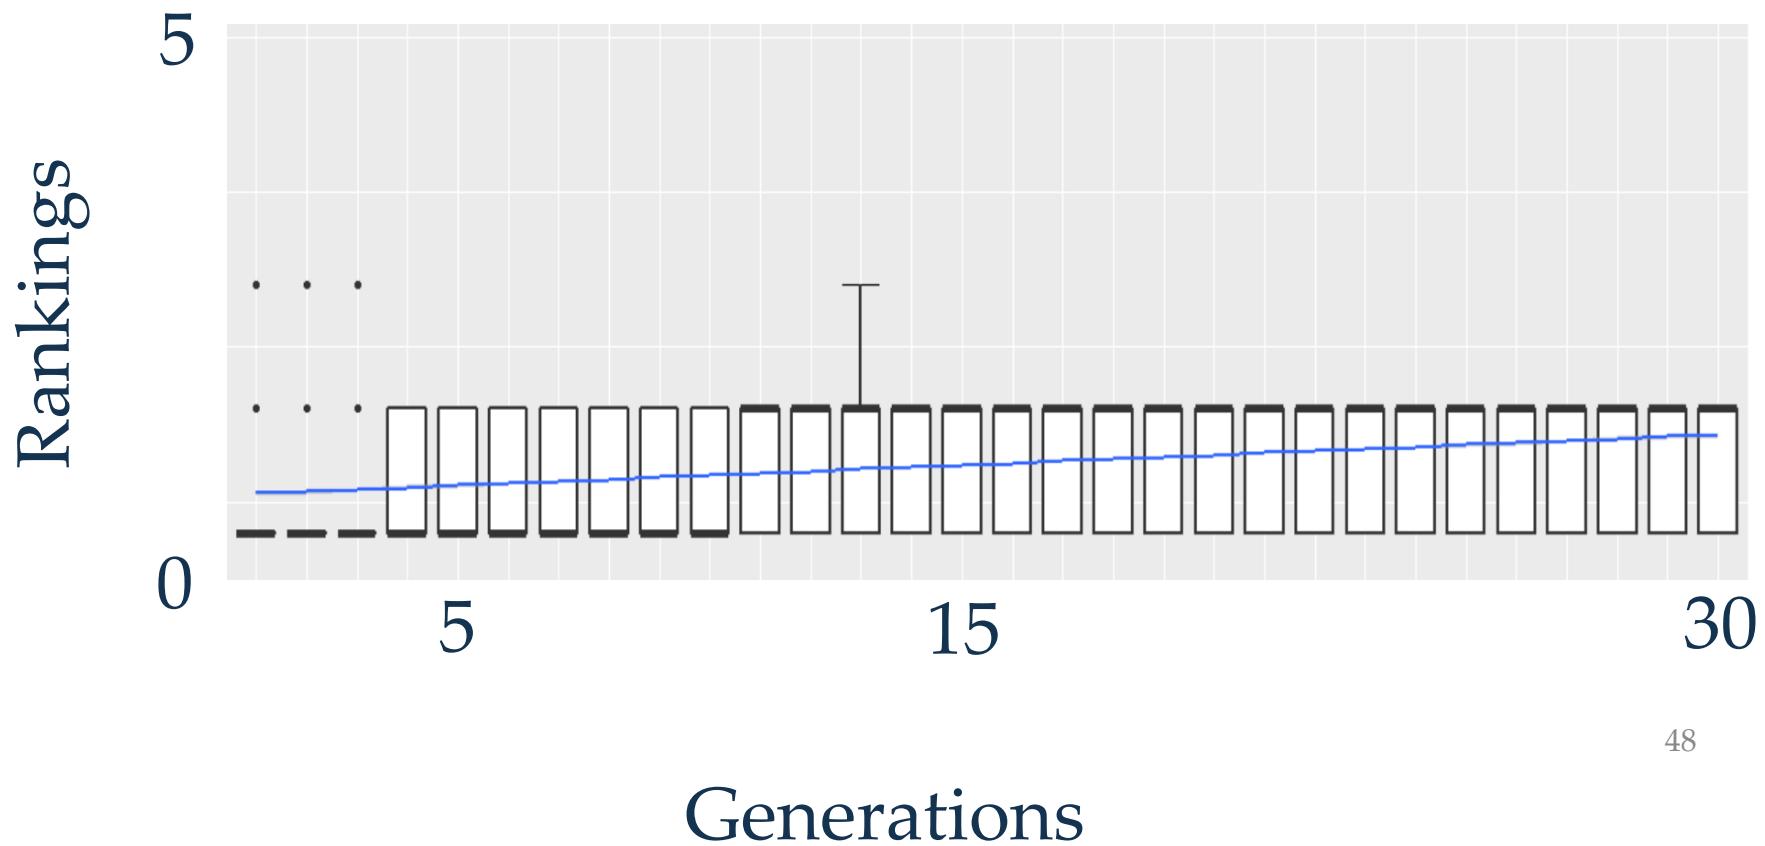


RQ₂ - Performance regression in a real code change

Code change: StoryHierarchyBusinessImpl.calculateStoryTreeMetrics()

```
public StoryTreeBranchMetrics calculateStoryTreeMetrics(Story story) {  
    for(Story child : story.getChildren()) {  
        if (child.getId() == story.getId()) {  
            continue;  
        }  
        StoryTreeBranchMetrics childMetrics = this.calculateStoryTreeMetrics(child);  
    }  
    return metrics;  
}
```

RQ₂ - Performance regression in an artificial code change



RQ₂ - Performance regression in an artificial code change

