

# The Role of Software Tracing in Software Maintenance

**Abdelwahab Hamou-Lhadj, PhD.**

Software Behaviour Analysis (SBA) Research Lab  
Department of Electrical and Computer Engineering  
Concordia University, Montreal, QC, Canada

[www.ece.concordia.ca/~abdelw](http://www.ece.concordia.ca/~abdelw)

CIIA'13, Saida, Algeria  
May 5, 2013

# Software Maintenance

- is defined as the modification of a software system after delivery
- accounts for 75% of the time of the software life cycle
- tends to be a human resource intensive process
- incurs very high costs: SW maintenance is estimated to a multi-billion dollar market

# Issues with exiting software

- More than 100 billion lines of code in production in the world
- A large portion of it is unstructured, patched, and badly documented
- Initial design and architecture can no longer be trusted
- High turn-over causes initial developers to move from one company to another
- SW industry tends to be a poorly regulated industry

# As a result

Software engineers must spent a considerable amount of time to understand the system before making any changes to it



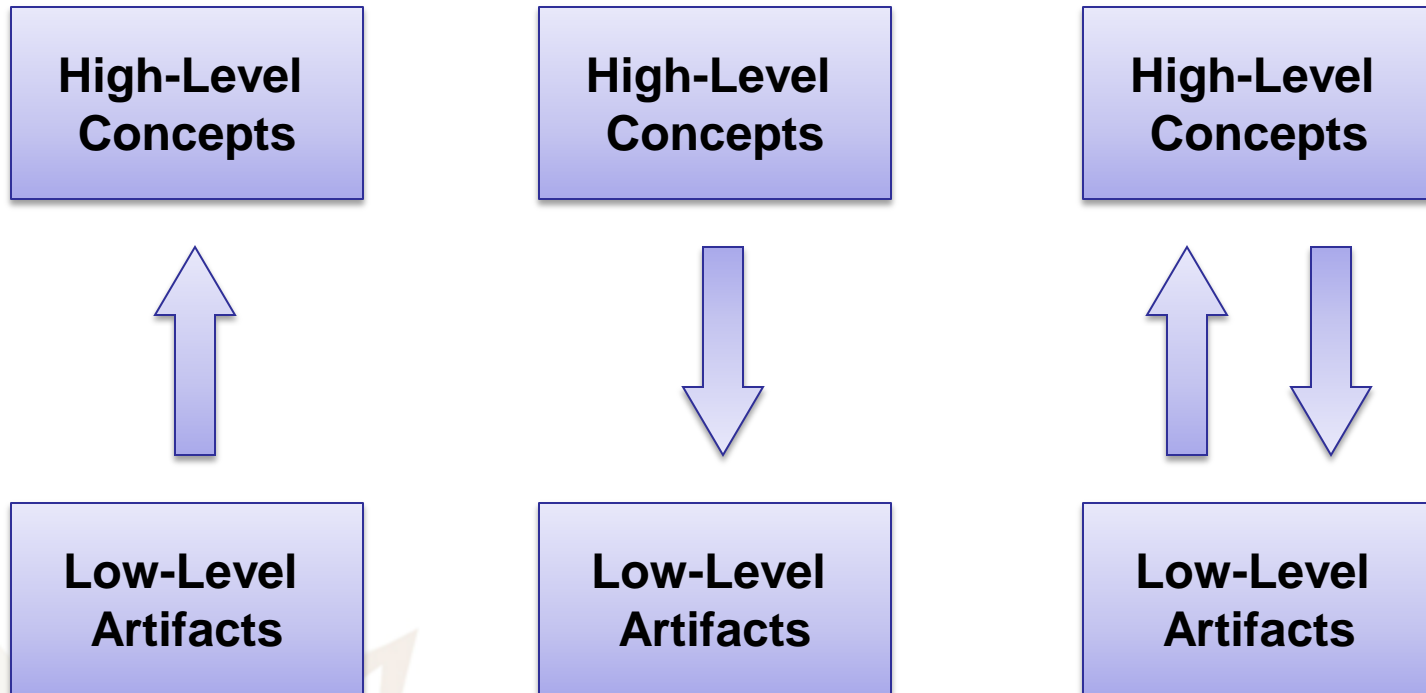
# As a result

Software engineers must spent a considerable amount of time to understand the system before making any changes to it



# How do programmers understand programs?

## Program Comprehension Models





# Examples of execution traces

- Traces of routine (method) calls
- Traces of inter-process communication
- Traces of statement execution
- Traces of communication among subsystems
- Etc.



# Trace Analysis (cont'd)

- Advantages:
  - High focus and resolution
  - Mapping of program inputs to outputs
  - Source code is not needed
- Challenges:
  - Tracing adds overhead to the system
  - Traces are overwhelmingly large
  - Different types of traces may require different processing techniques

# Applications of Trace Analysis: Industrial Projects

Project 1: Tracing and Monitoring Tools for Distributed Multi-Core Systems

Project 2: Diagnostics for Real Time Distributed Multi-Core Architecture in Avionics

Project 3: Finding Faulty Functions from Traces of Field Failures

# Tracing and Monitoring Tools for Distributed Multi-Core Systems

Develop techniques and algorithms to provide a software architecture for low overhead trace generation and analysis tools for complex distributed multi-core systems

# Project Partners



# Trace Generation

- Research thread led by Dr. Michel Dagenais from Polytech de Montreal
- Objectives:
  - Build a tracer with low overhead and no disturbance on the system
  - Offer support for synchronisation in a multi-core environment
  - Offer support for system and user space tracing

# LTTng: Linux Trace Toolkit New Generation


- Instruments the Linux kernel
- Adds 2% overhead to the kernel in the worst case scenario
- Is free and open source
- Is being integrated with the Linux kernel

# Trace Analysis

- Objectives:
  - Simplify the understanding and analysis of very large traces
  - Extract high-level views from raw events
  - Identify the main components that implement the traced scenario
  - Correlate user space and system space traces

# Motivating example


- A trace generated from a compiler:
  - parsing, preprocessing, lexical analysis, semantic analysis
- In most trace visualization tools, it will look like:







- But how can we tell what happens where?




# An execution phase based view



**Init      Parsing      Preprocessing      Lexical Analysis      Semantic Analysis**



# Same problem, different domain: The human perception system



... yet the brain can quickly identify parts of a scene, group them into shapes, and interpret the shapes

## **Segmentation:**

The perceptual system segments local elements against their context and integrates them as objects and regions

## **Global Perception:**

The segmented scene is then quickly scanned with eye movements so as the brain obtains an overall impression of it


## **Preattentive Process:**

The scene is analyzed in more detail by visiting the regions in a certain order. The pop-out effect is an important factor in this process

# Gestalt Laws




Law of Similarity



Law of Continuity

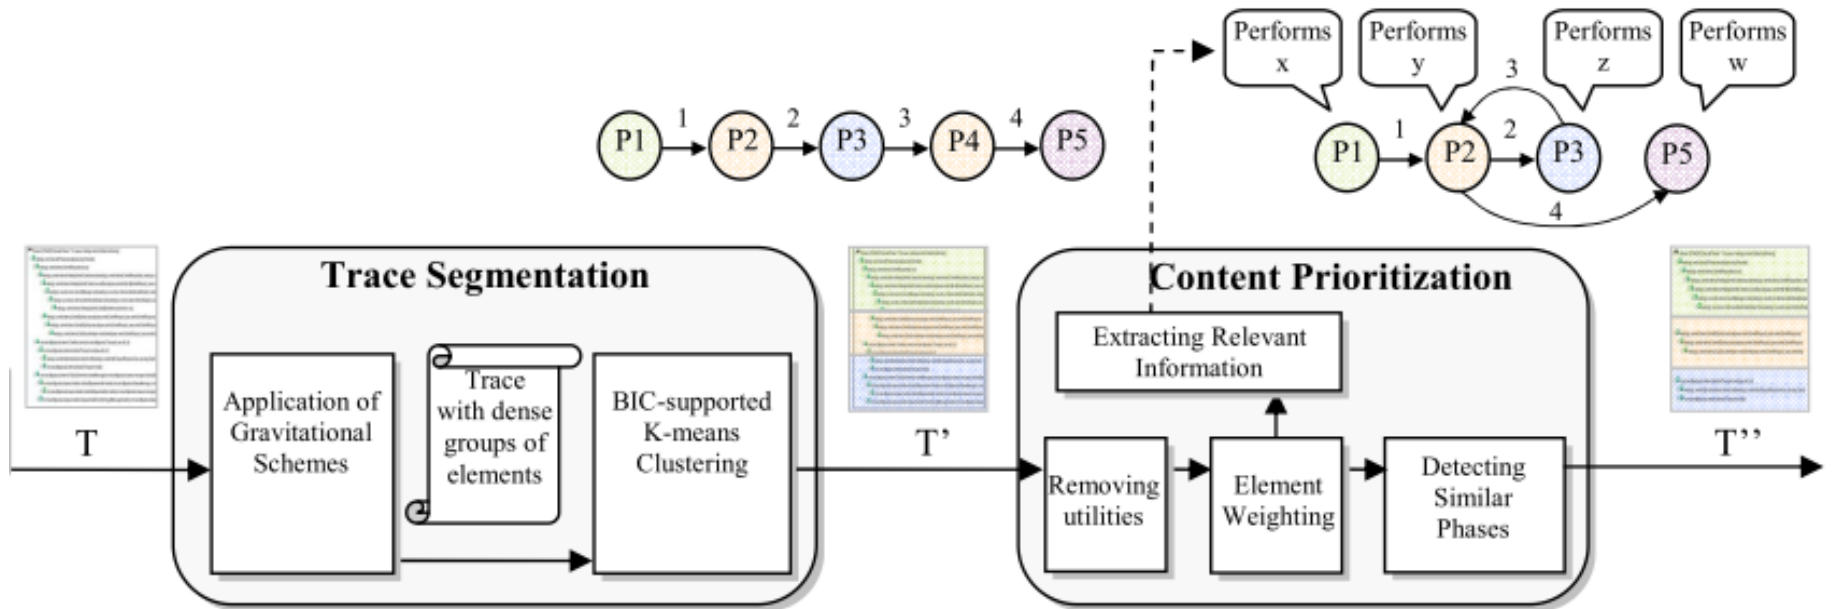


Law of Pragnanz

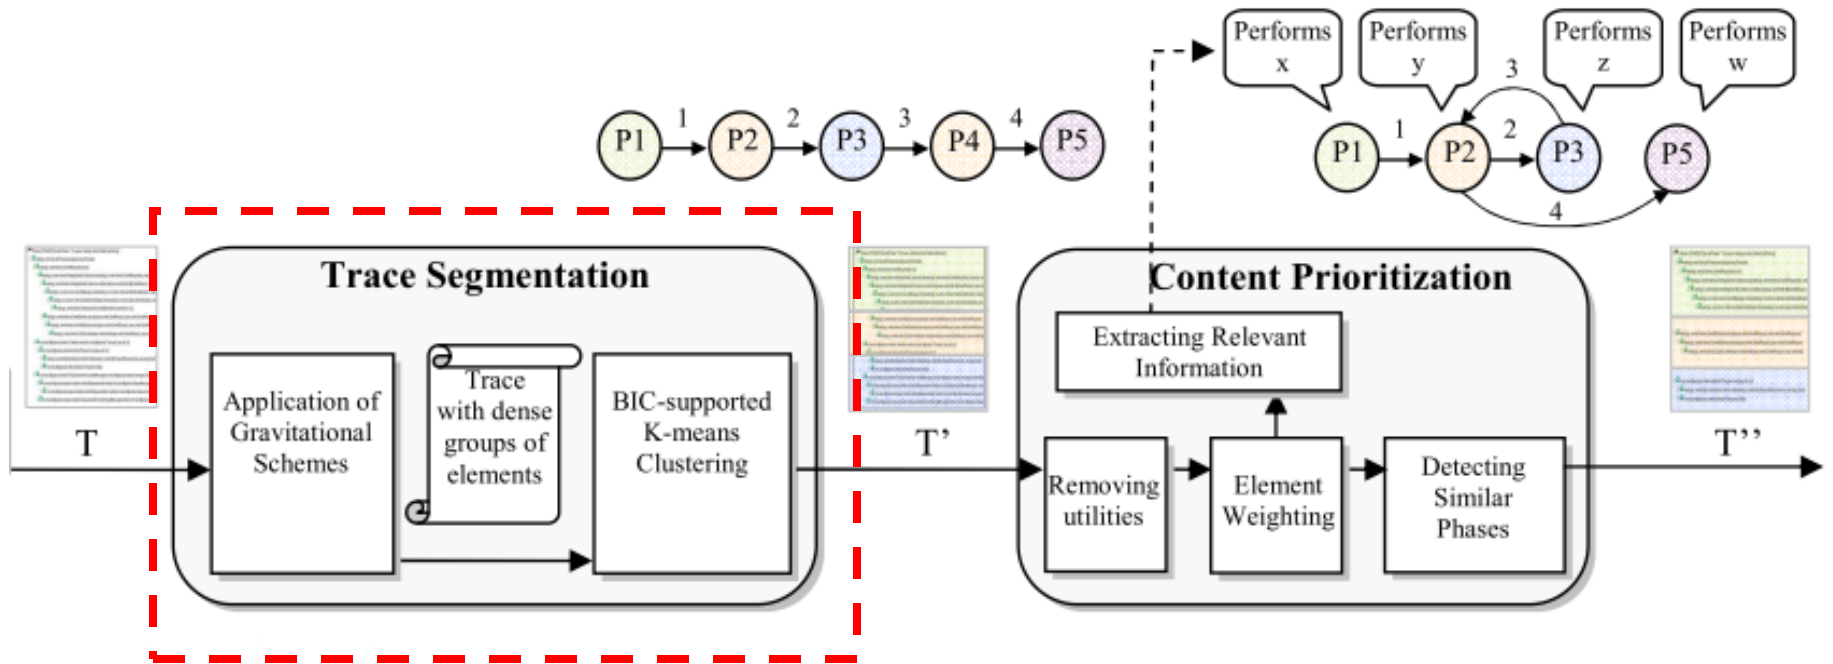


Law of Proximity


# Trace abstraction framework



# Trace abstraction framework




# Repositioning trace events using similarity principle







# Good continuation (cont.)

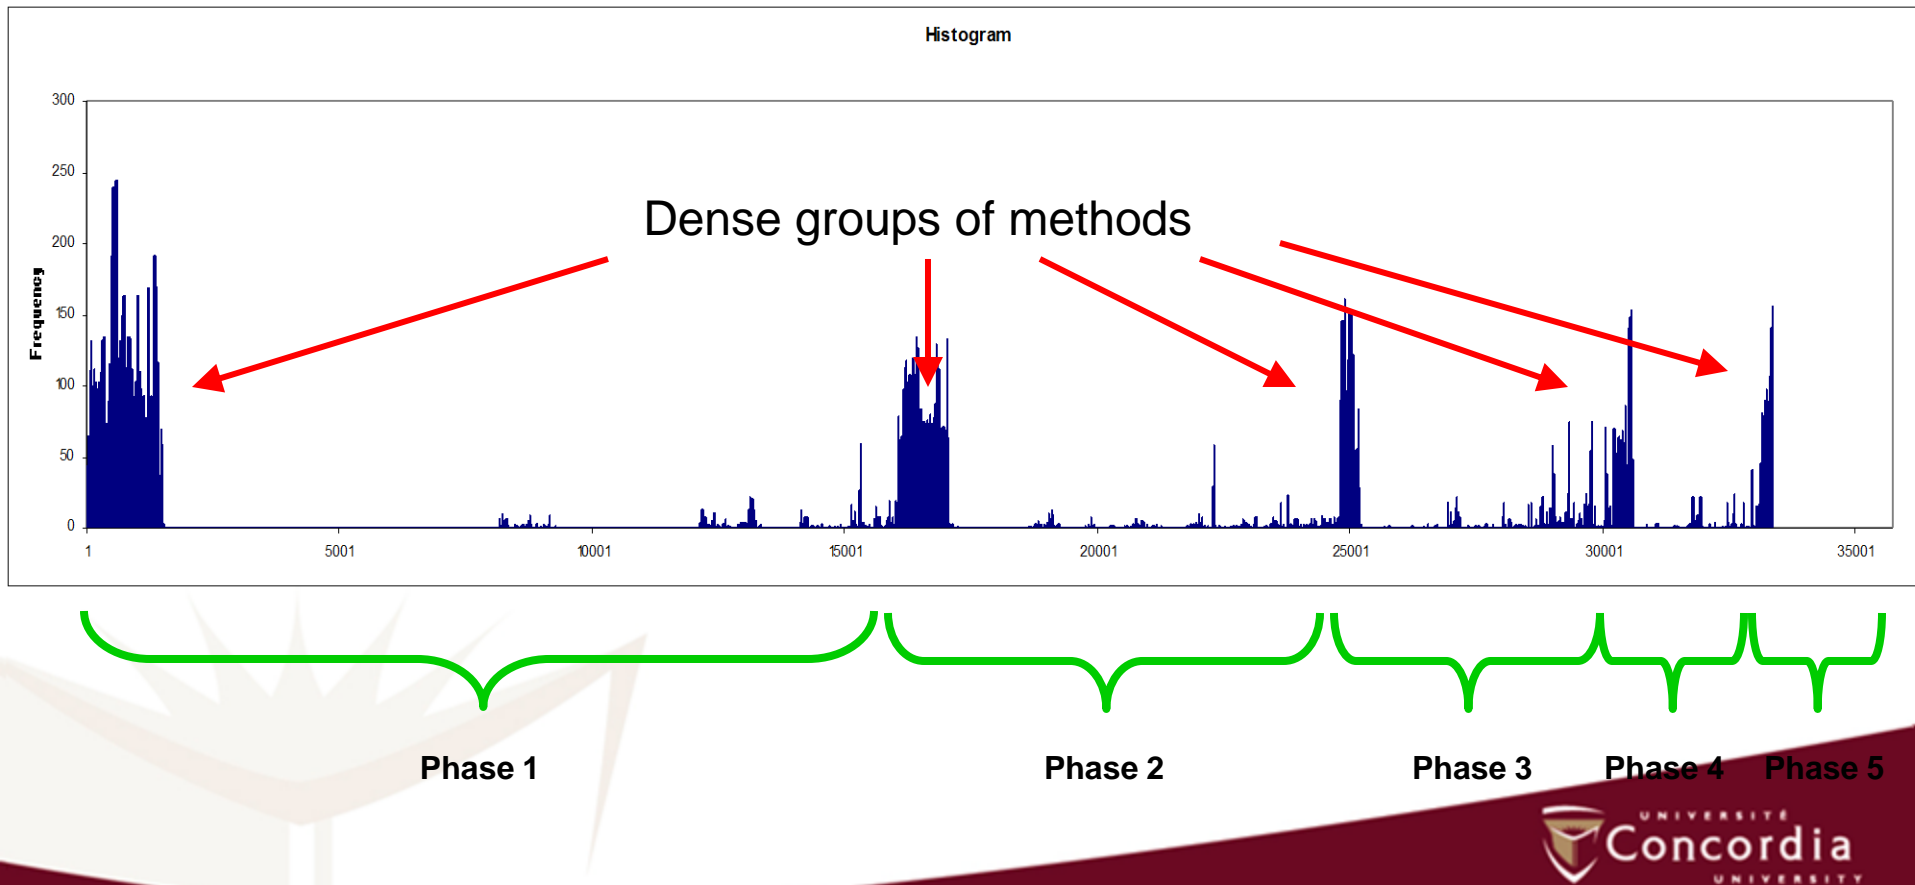


# Evaluation

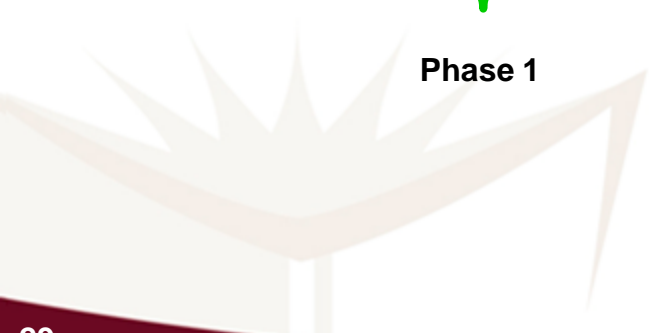
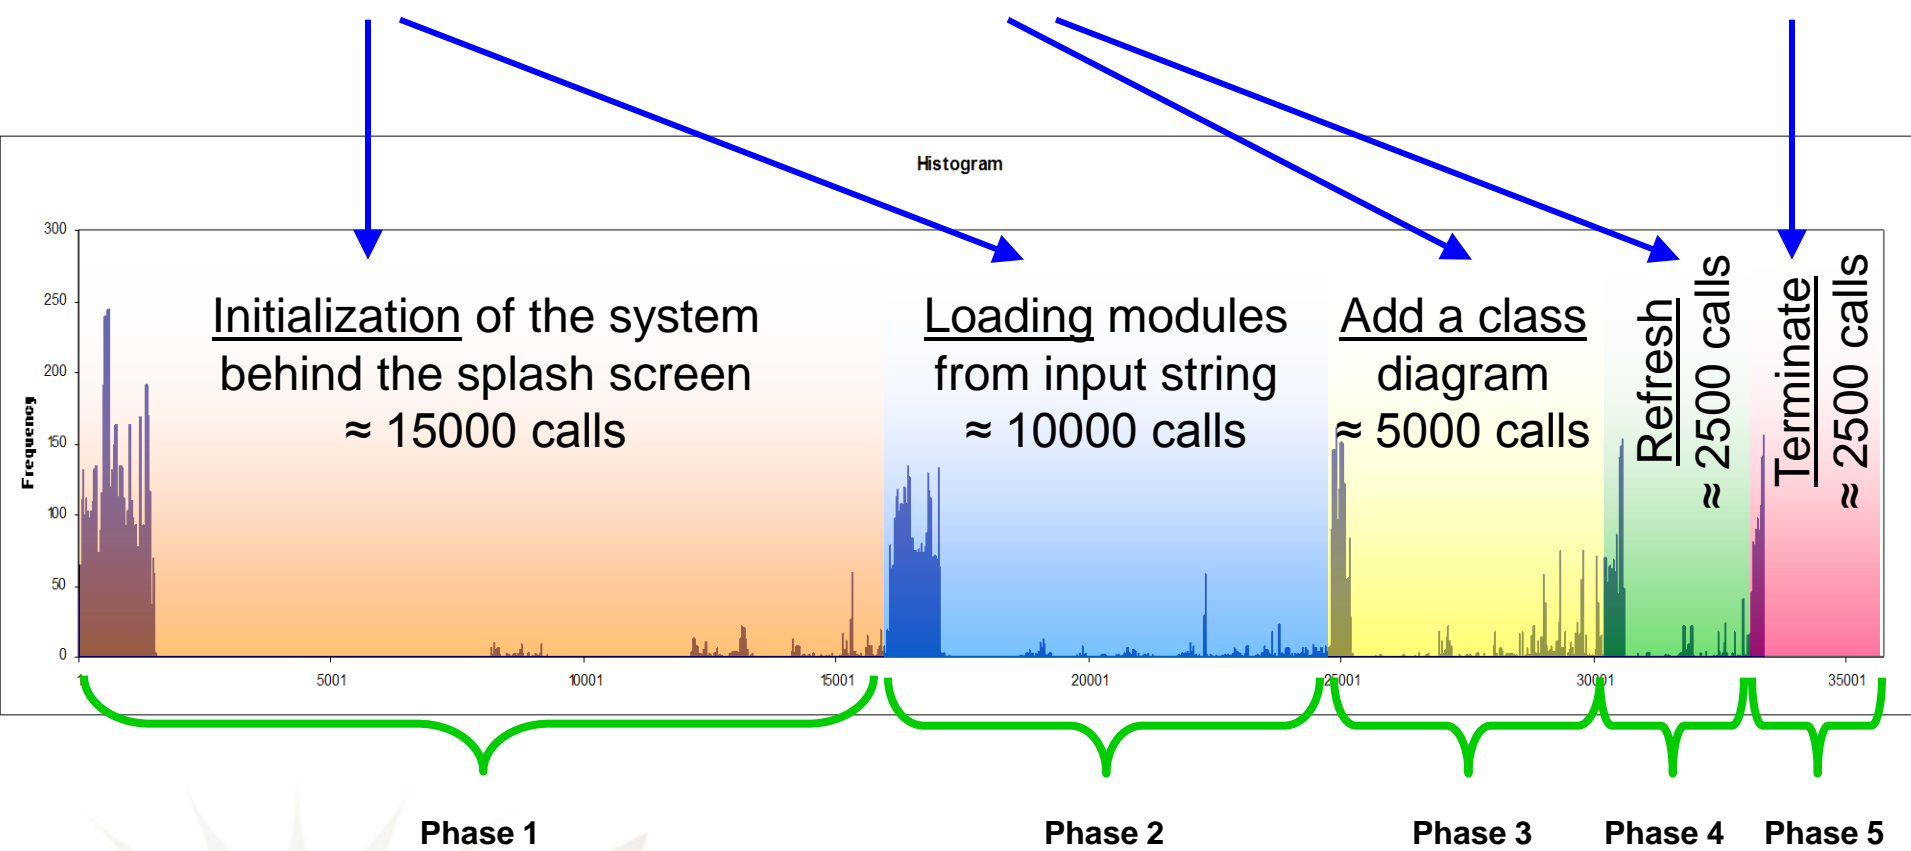
- Target System: ArgoUML
- Scenario: Starting up ArgoUML, drawing a class diagram, quitting ArgoUML
- Trace size: Hundred of thousands of function calls
- Number of distinct routines 2331 = ~33%

# Application of trace segmentation

Starting up ArgoUML → Drawing a class diagram → Quitting ArgoUML




Starting up ArgoUML → Drawing a class diagram → Quitting ArgoUML



Starting up ArgoUML → Drawing a class diagram → Quitting ArgoUML


Phase 3: Add a class diagram

Histogram

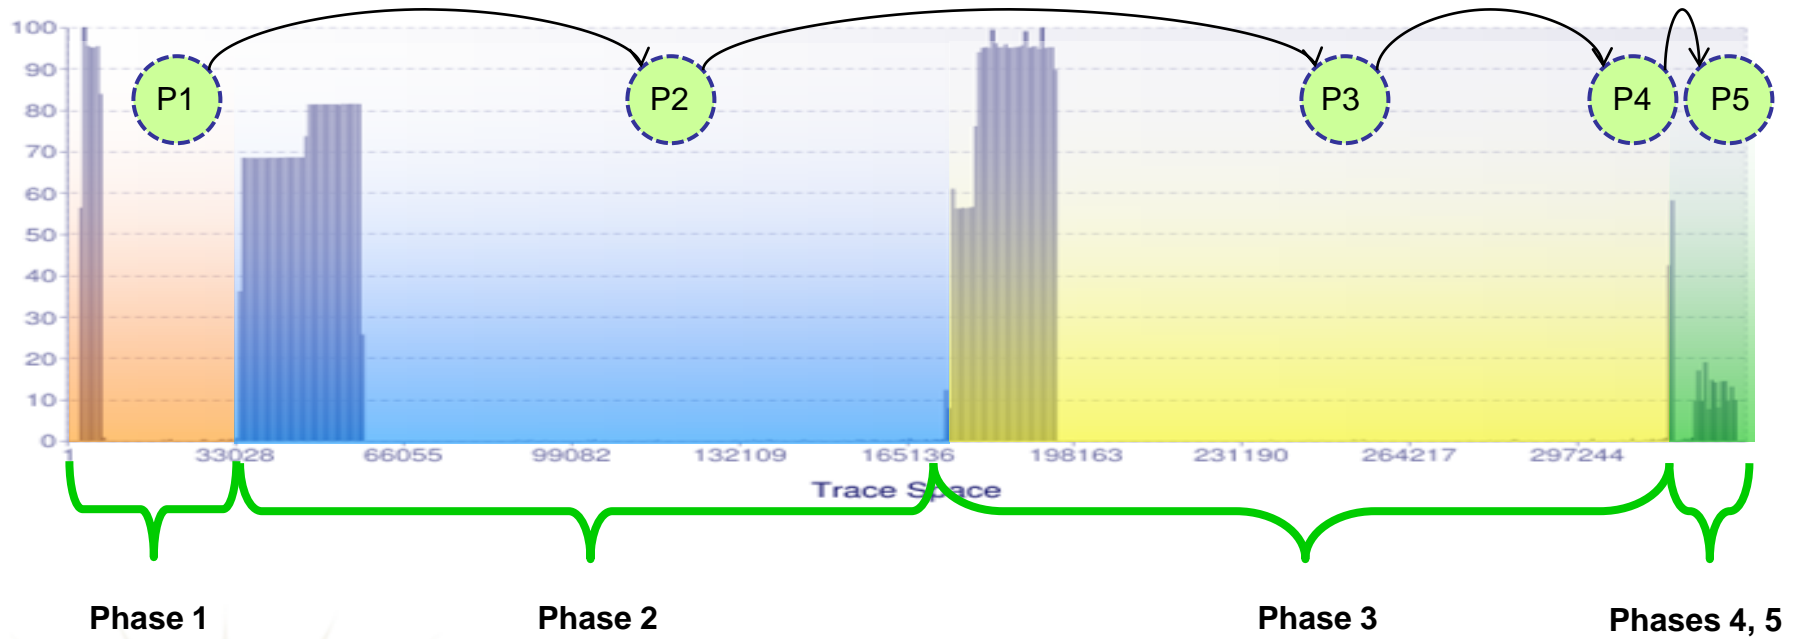


Starting up ArgoUML → Drawing a class diagram → Quitting ArgoUML

Phase 3: Add a class diagram




# Phase flow diagram





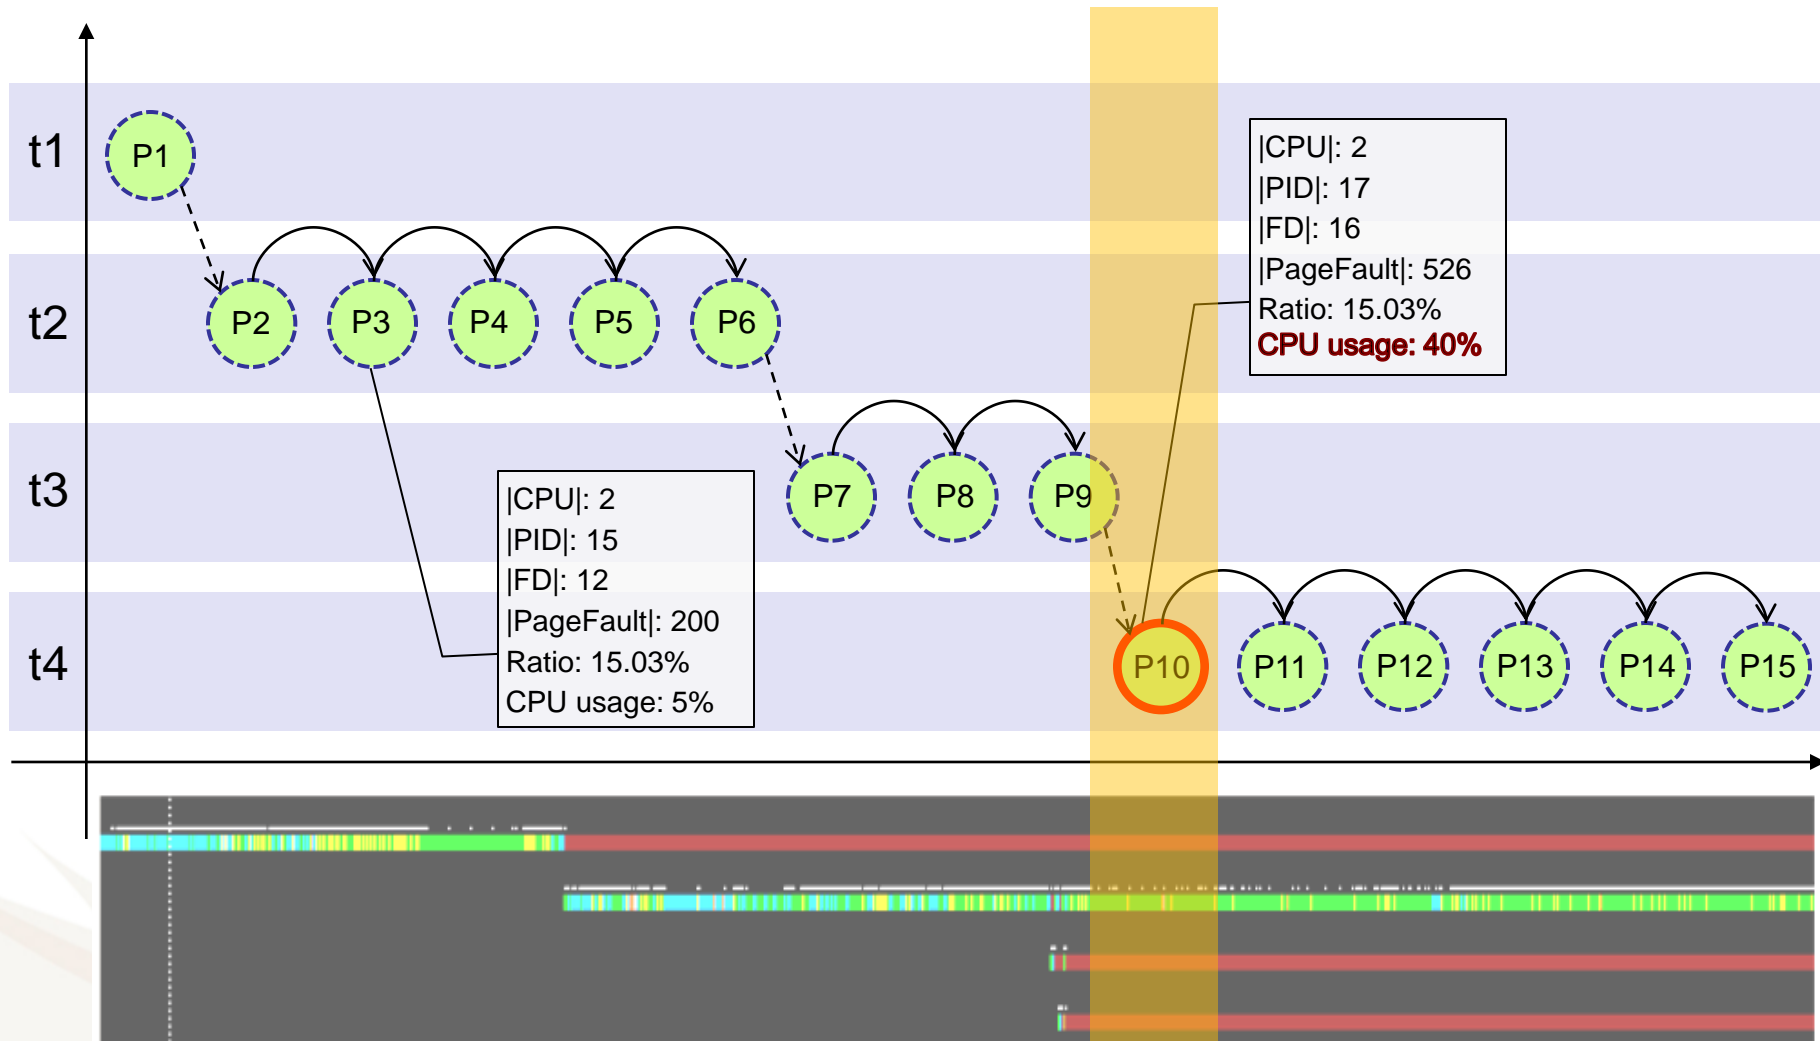
# Combining user and kernel space

Threads

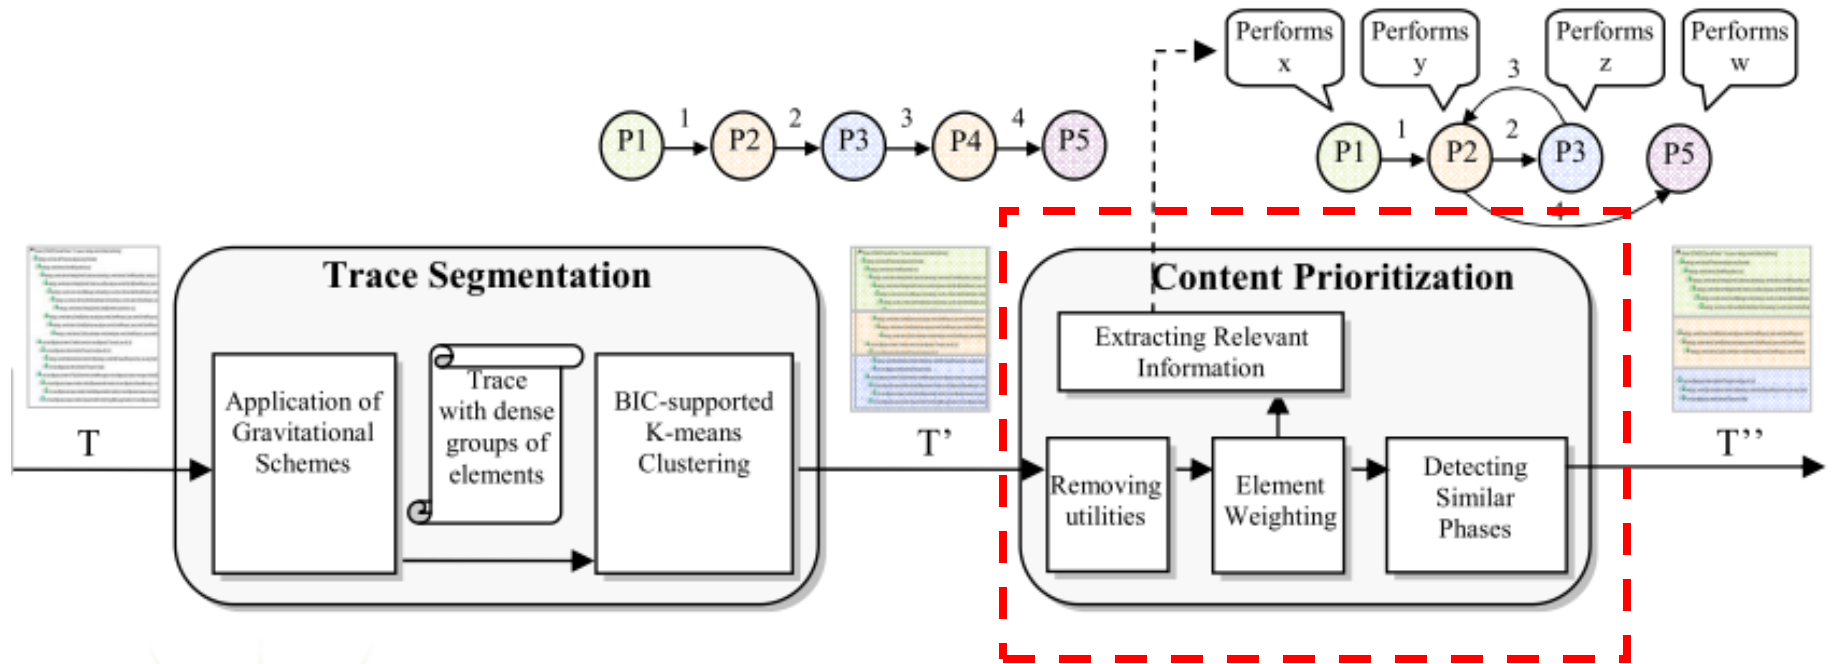


# Adding state information


Threads



# Trace abstraction framework



# Extracting relevant components



- Idea: Elements that are repeated in a phase but are not much shared between phases indicate their relevance to the phase
- This is similar to the concept of term frequency inverse document frequency in the text mining

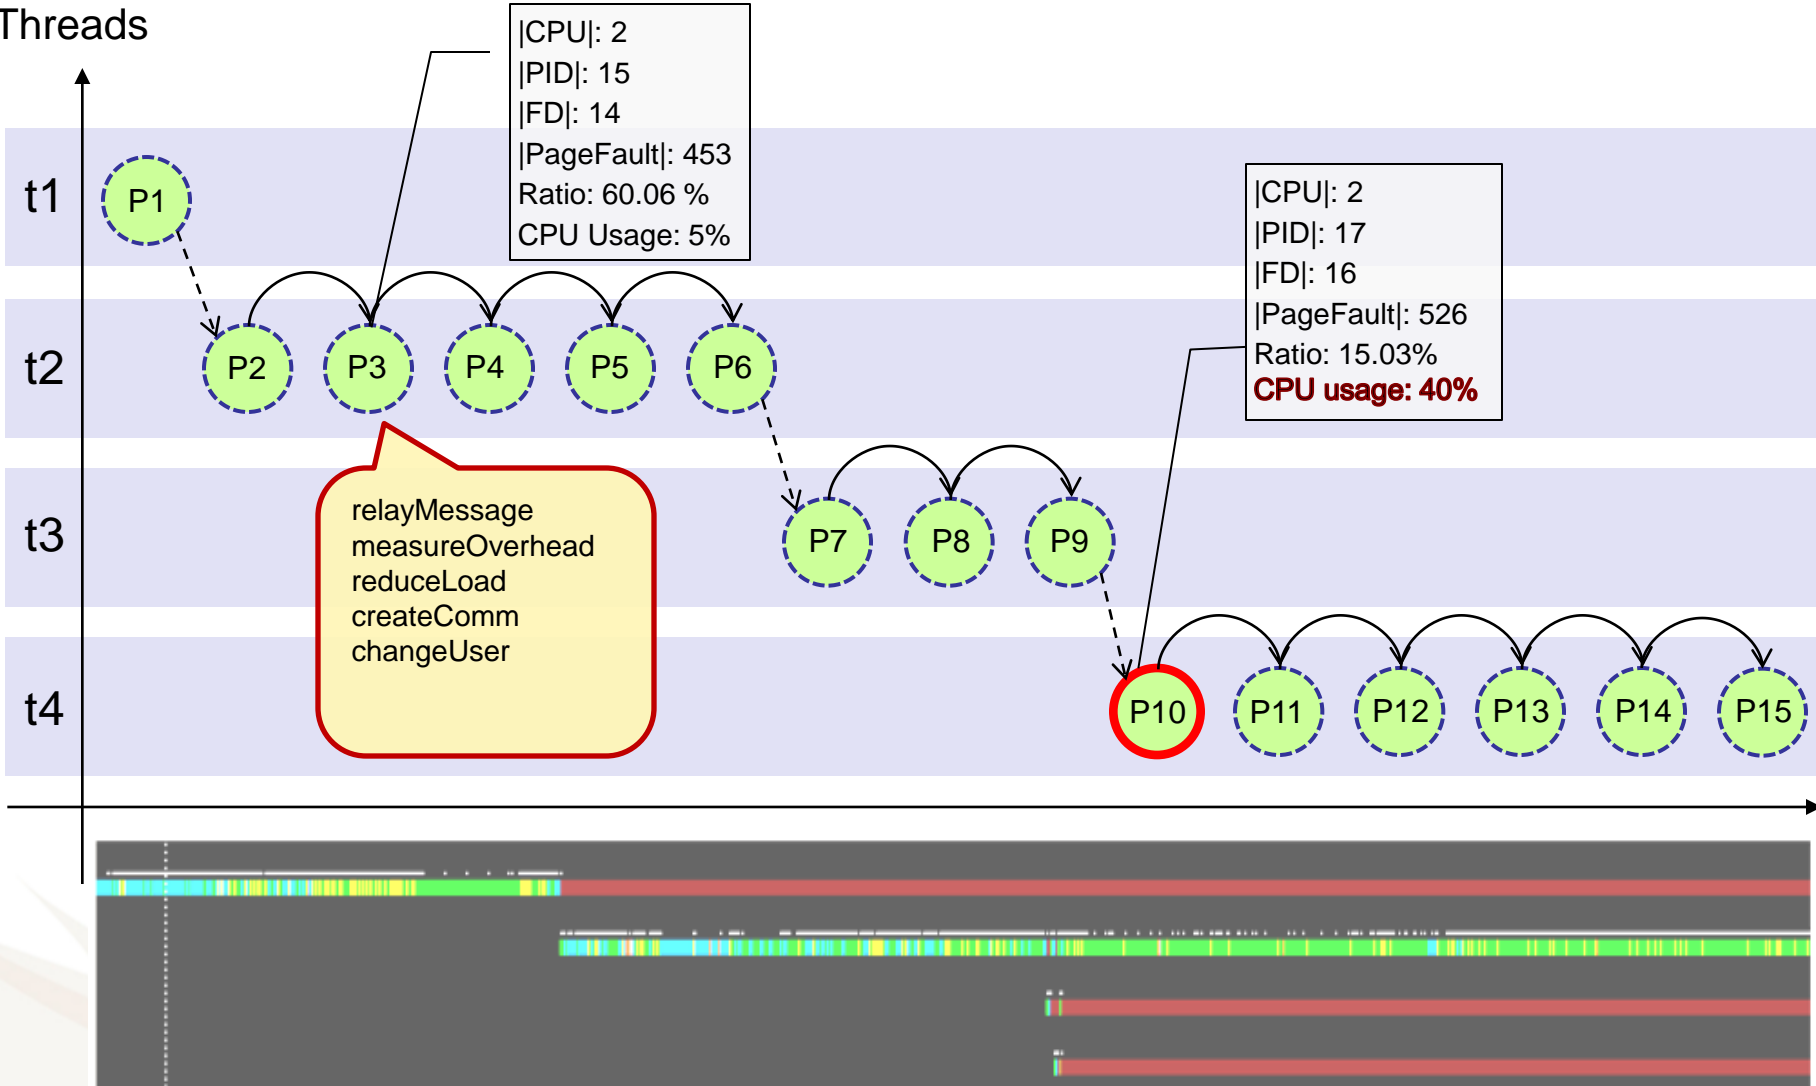
Document 1: Shipment of gold damaged in a fire

Document 2: Delivery of silver arrived in a silver truck

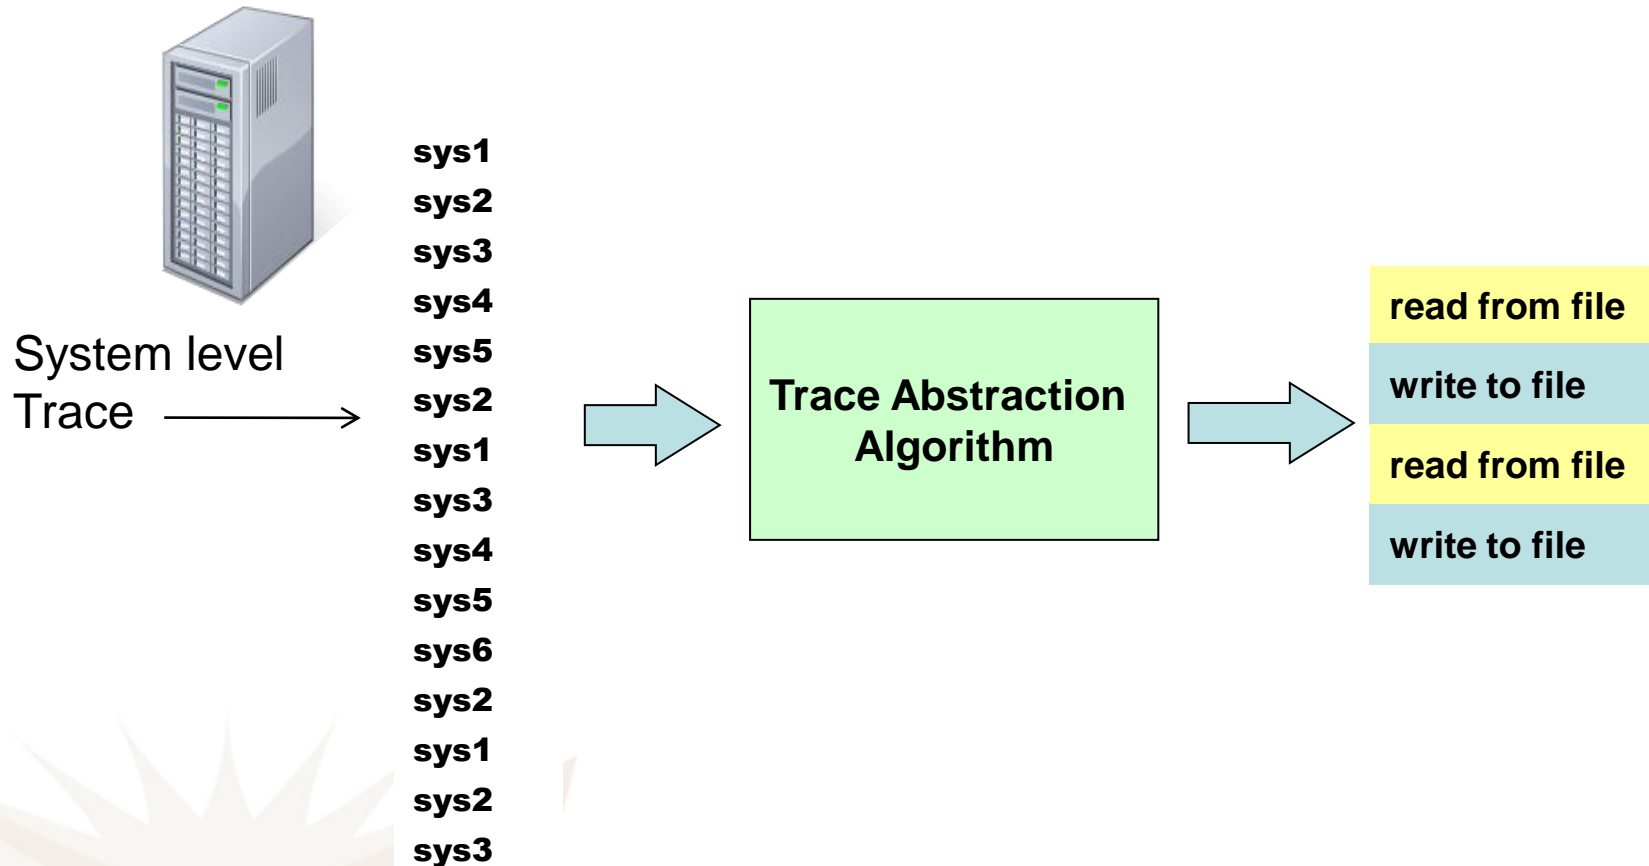
Document 3: Shipment of gold arrived in a truck

# Identifying main content

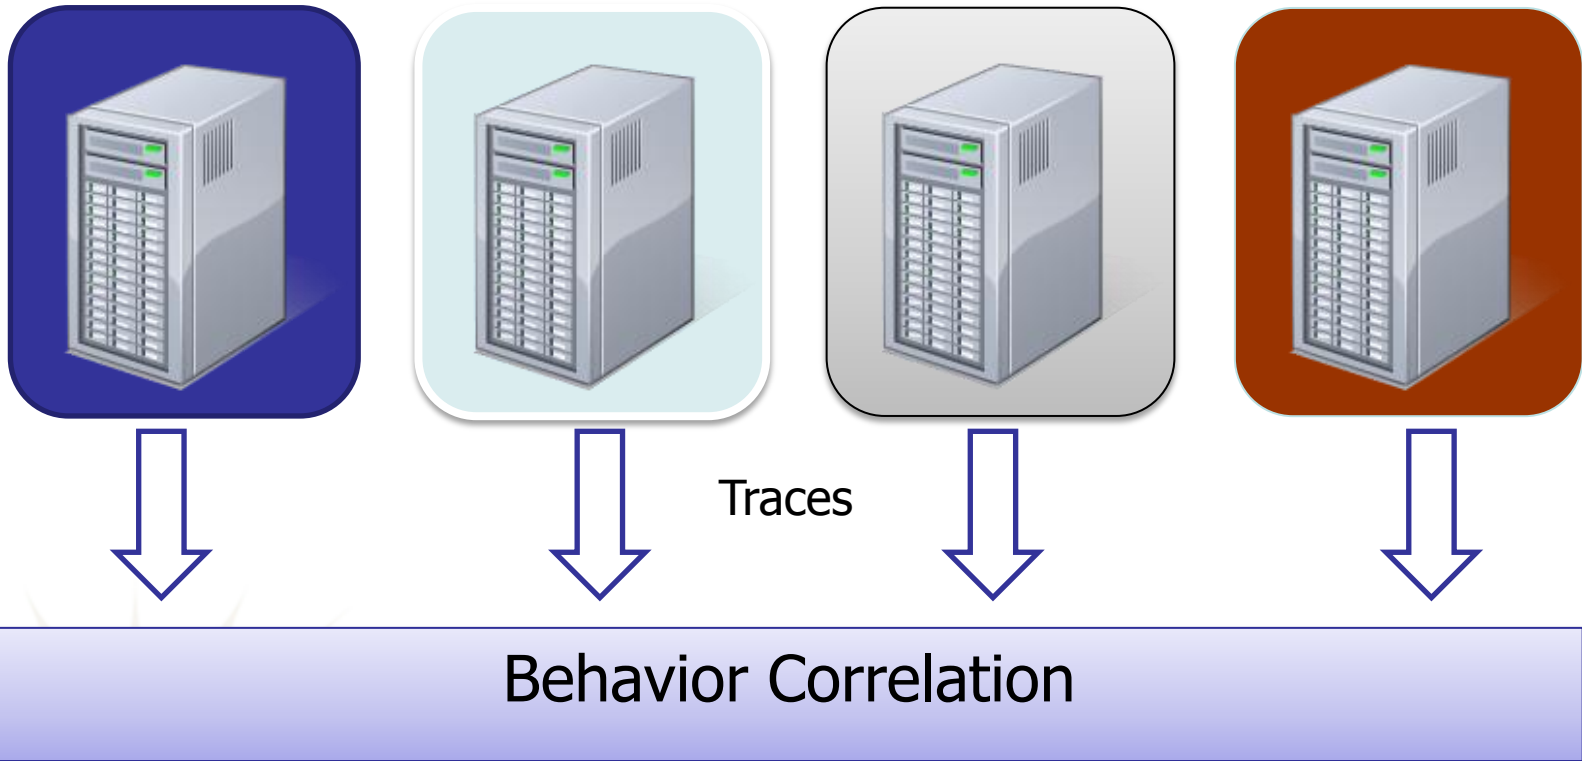
Threads



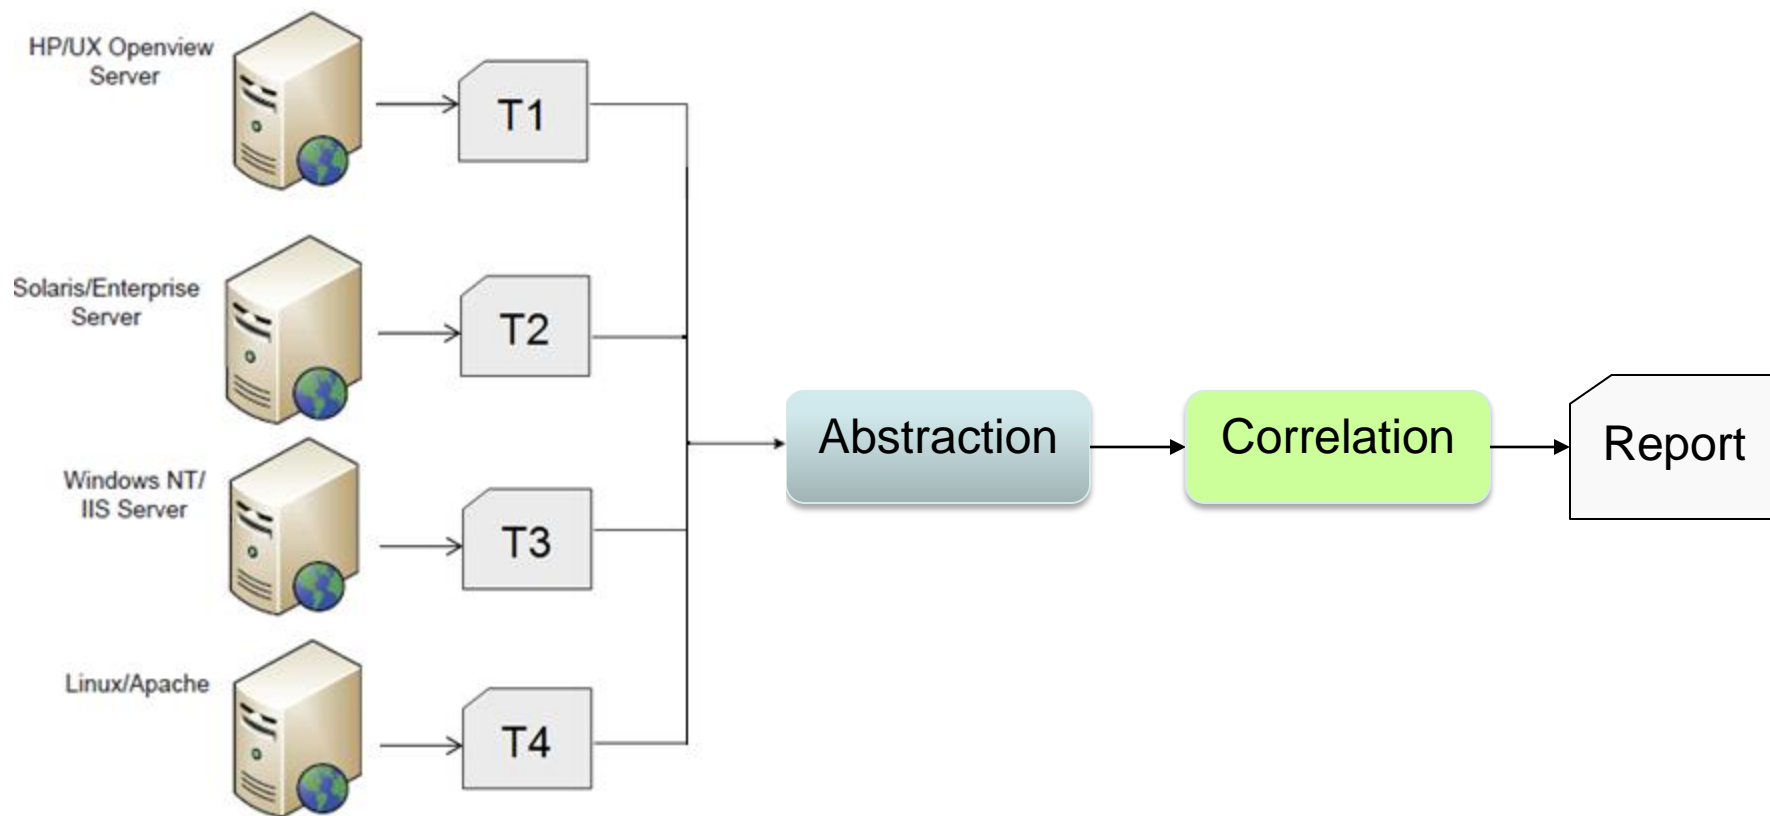
# Abstraction of System Call Traces



# Fault Tolerance: Redundancy and Diversity

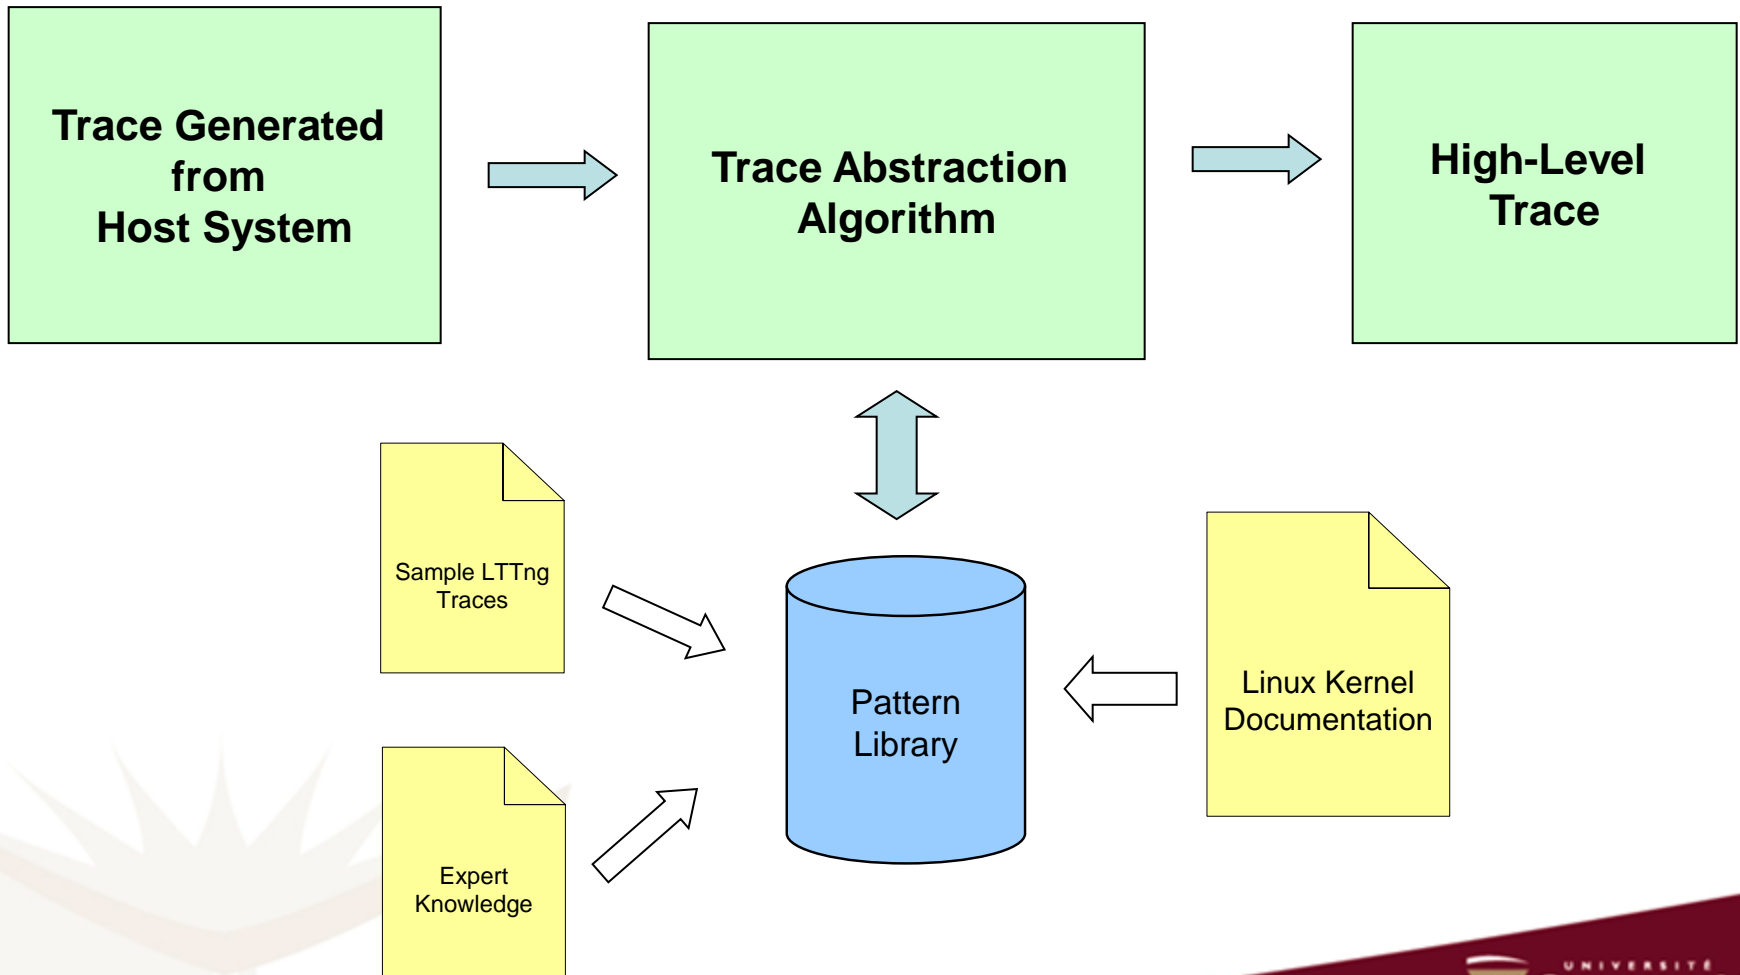


# OS Diversity






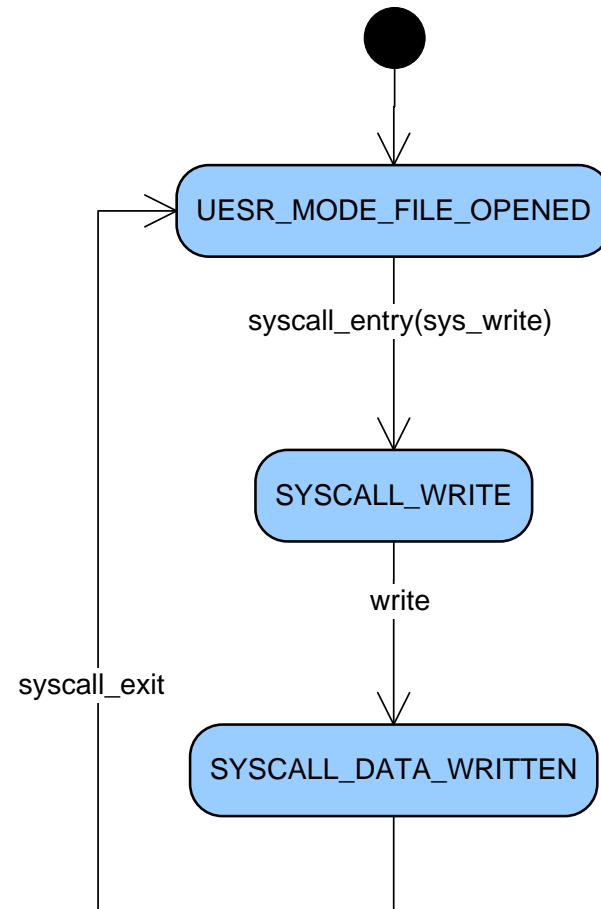
# Kernel-Level Trace Abstraction



# File Read & Write Patterns



Write to File



Read from File

# Evaluation

- Two nodes: Linux and BSD
- Failures are simulated on BSD
- We are able to detect and recover from most failures
- Abstraction is a crucial step for behavior correlation to be effective
- Similarity based on pattern detection provides accurate measures

# Tracing and Monitoring Framework

The screenshot displays the LTTng Kernel - Eclipse SDK interface. The main window is divided into several panes:

- Project Explorer:** Shows a project structure with folders like 'Experiments [1]' and 'Traces [5]', containing sub-traces such as 'firefox3-kernel', 'firefox3-ust', 'highusage', 'kernel', and 'mutex'.
- Control Flow:** A table listing processes being traced. The 'firefox' process is highlighted, showing its TID (13937), PTID (13920), and birth time (09:40:25.073133923).
- Trace:** A timeline view showing the execution of the 'firefox' process, with a vertical line indicating the current time (09:40:25.087).
- Events - firefox3:** A table of system events. The 'sys\_ioctl' event is highlighted, showing its timestamp (09:40:25.0874229), source (5), and content (arg=140736015176848, cmd=3224389159, fd=6).
- Histogram:** A bar chart showing the frequency of events over time. The current event is at 1339422025.087422965, and the window span is 0.002814750 seconds.

Users: Ericsson, Google, IBM, and many more

# Diagnostics for Real Time Distributed Multi-Core Architecture in Avionics

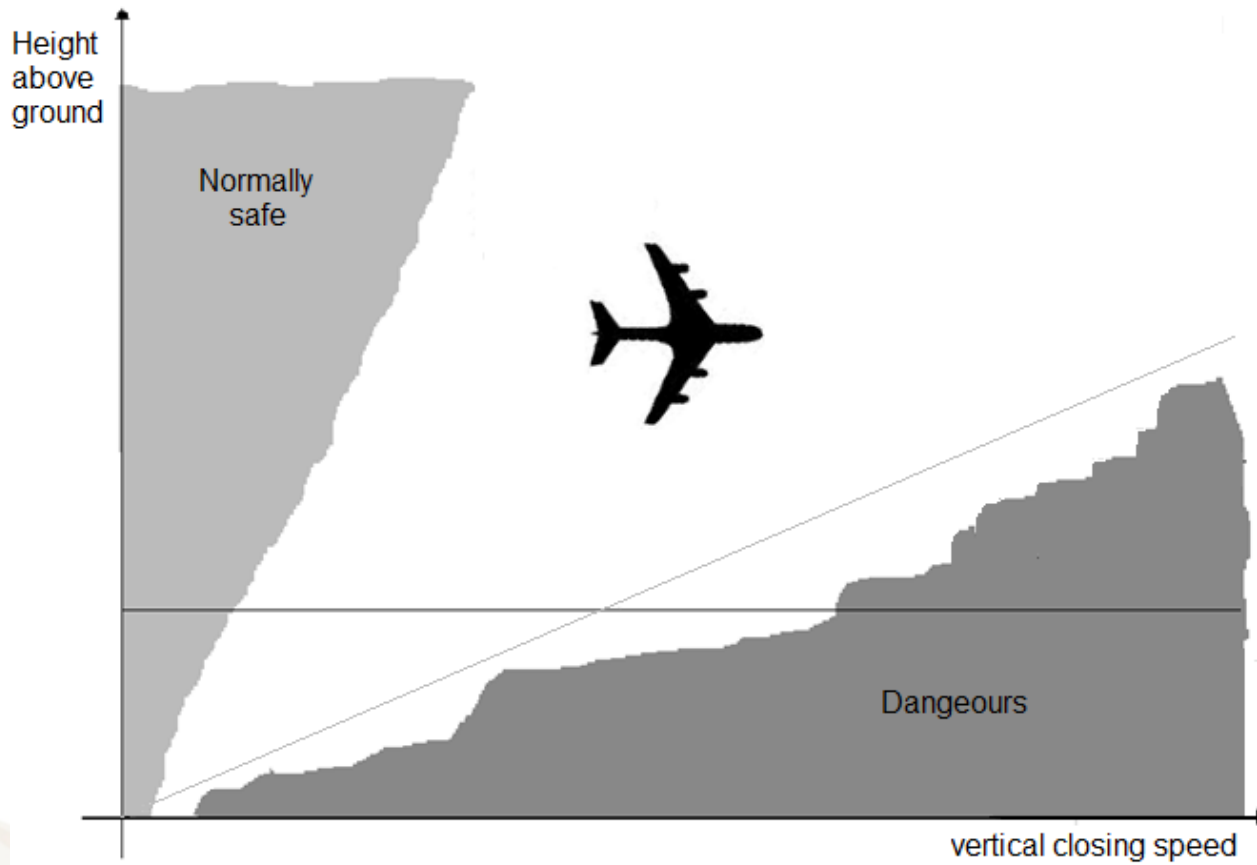
Build efficient algorithms for low overhead, low disturbance tracing of real-time embedded multi-core systems and simulators

Develop special purpose performance analysis debugging, and feature location modules for avionic systems

# Project Partners



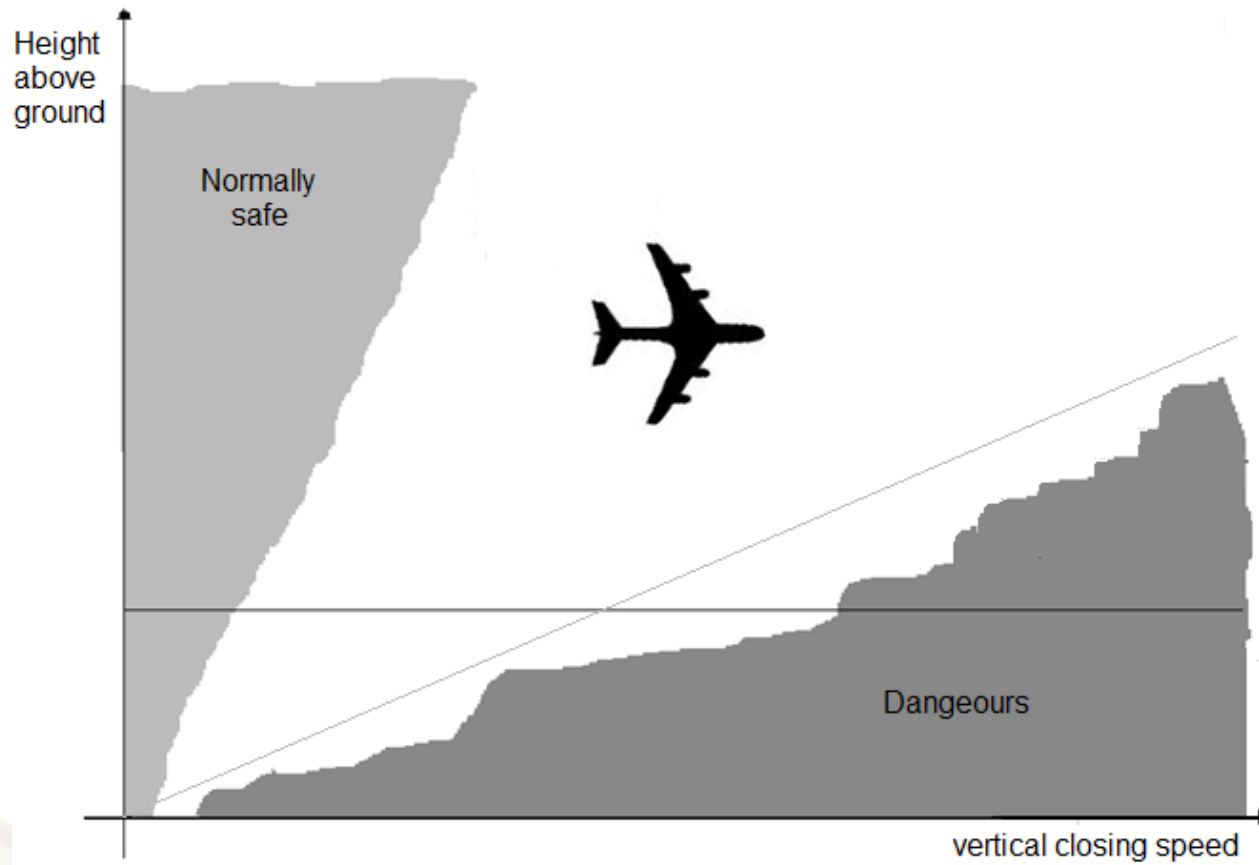
# Motivating scenario



IOS

FLOA:	
Warning Obstacle	FALSE
FLOA:	
Avoid Obstacle	FALSE

# Motivating scenario

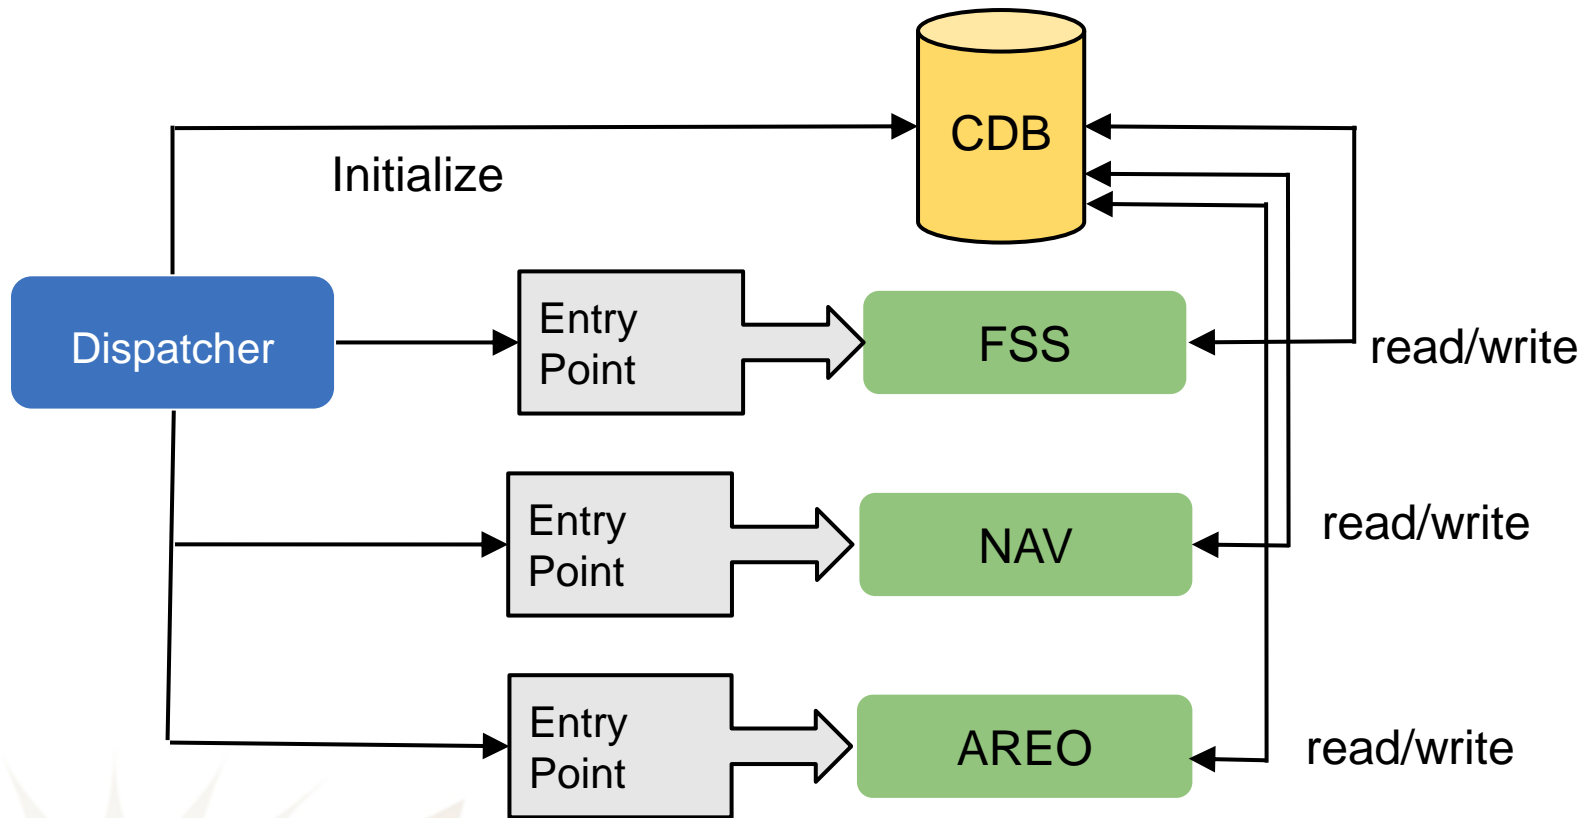


IOS

FLOA:	
Warning Obstacle	TRUE
FLOA:	
Avoid Obstacle	TRUE

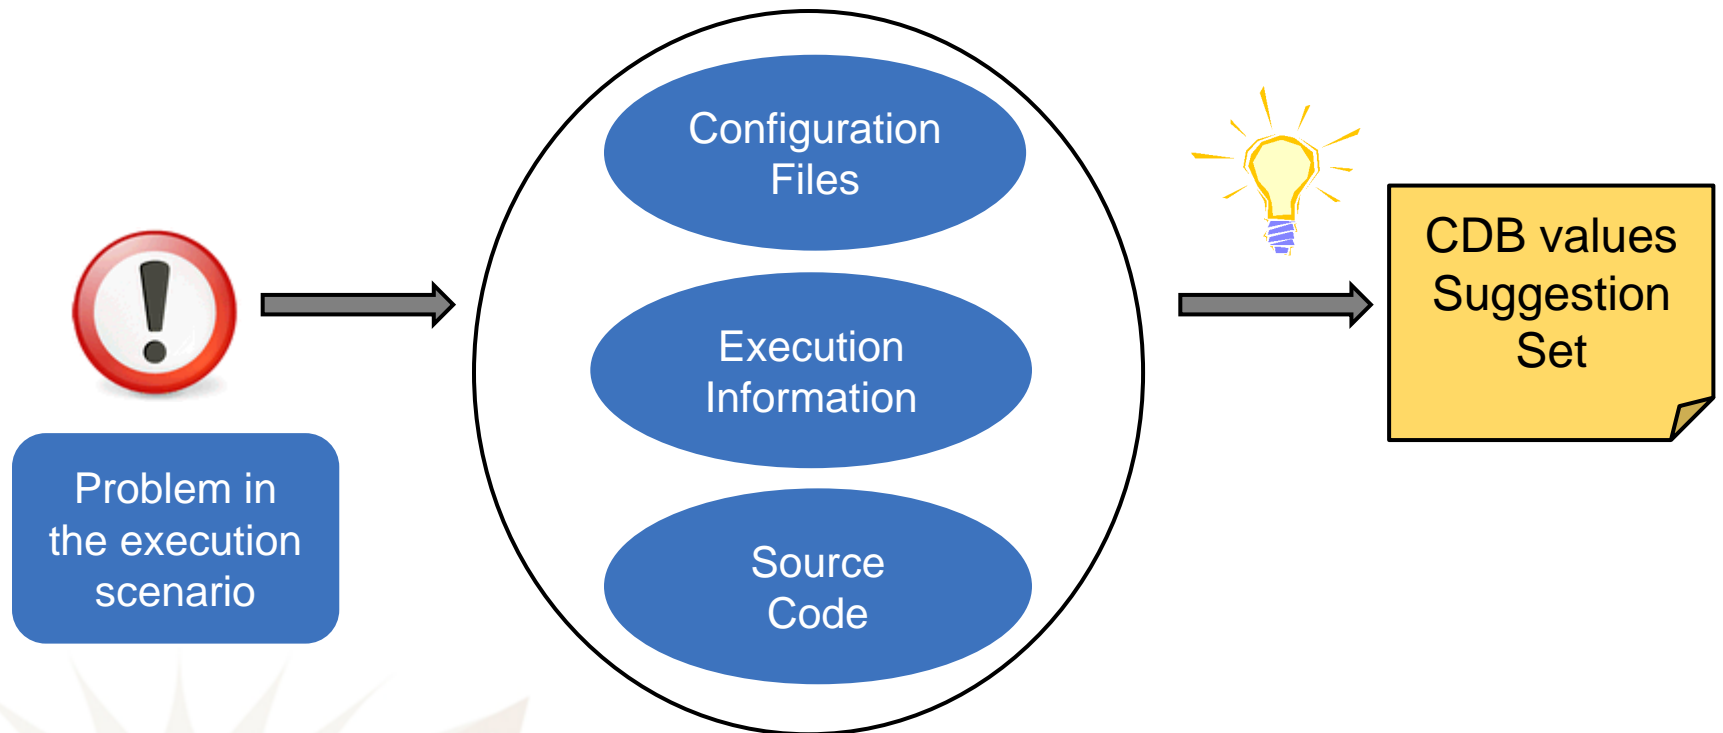


# CAE - Architecture

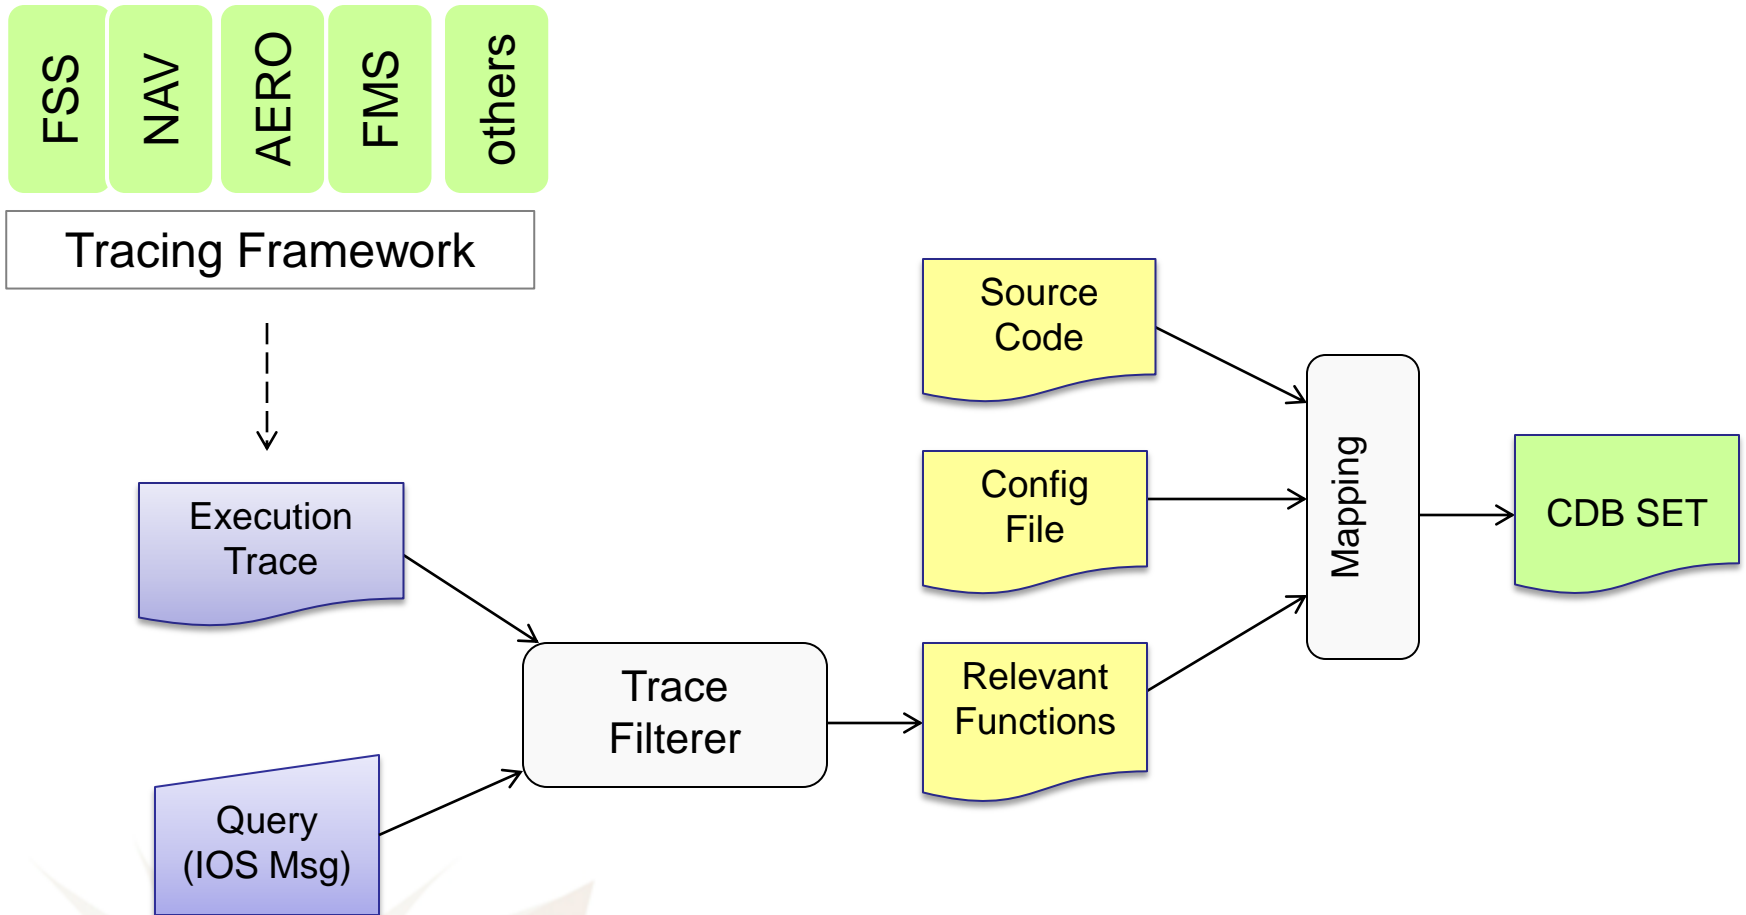




# Need: Automatic identification of CDB values for a specific failure



# Proposed Solution




# Evaluation

Scenario	Aircraft Condition	Trace size (millions)	# CDB vars in config.	Relevant CDBs	Retrieved CDBs	Precision	Recall
TAWS Mode1	Altitude: 900 feet Vertical speed: -3000 feet/min	20	1720	4	1	25%	50%
TAWS Mode4B	Altitude: 300 feet Airspeed: 50 knots Gears Position: down Flaps Position: in flight	8	1620	4	19	21%	100%
TAWS Mode4A	Altitude: 400 feet Airspeed: 50 knots Gears Position: up Flaps Position: landing	4	1499	5	28	19%	100%

# Finding Faulty Functions from the Traces of Field Failures

Improve the troubleshooting process to increase the productivity of software engineers by reducing the number of field reports to be analysed



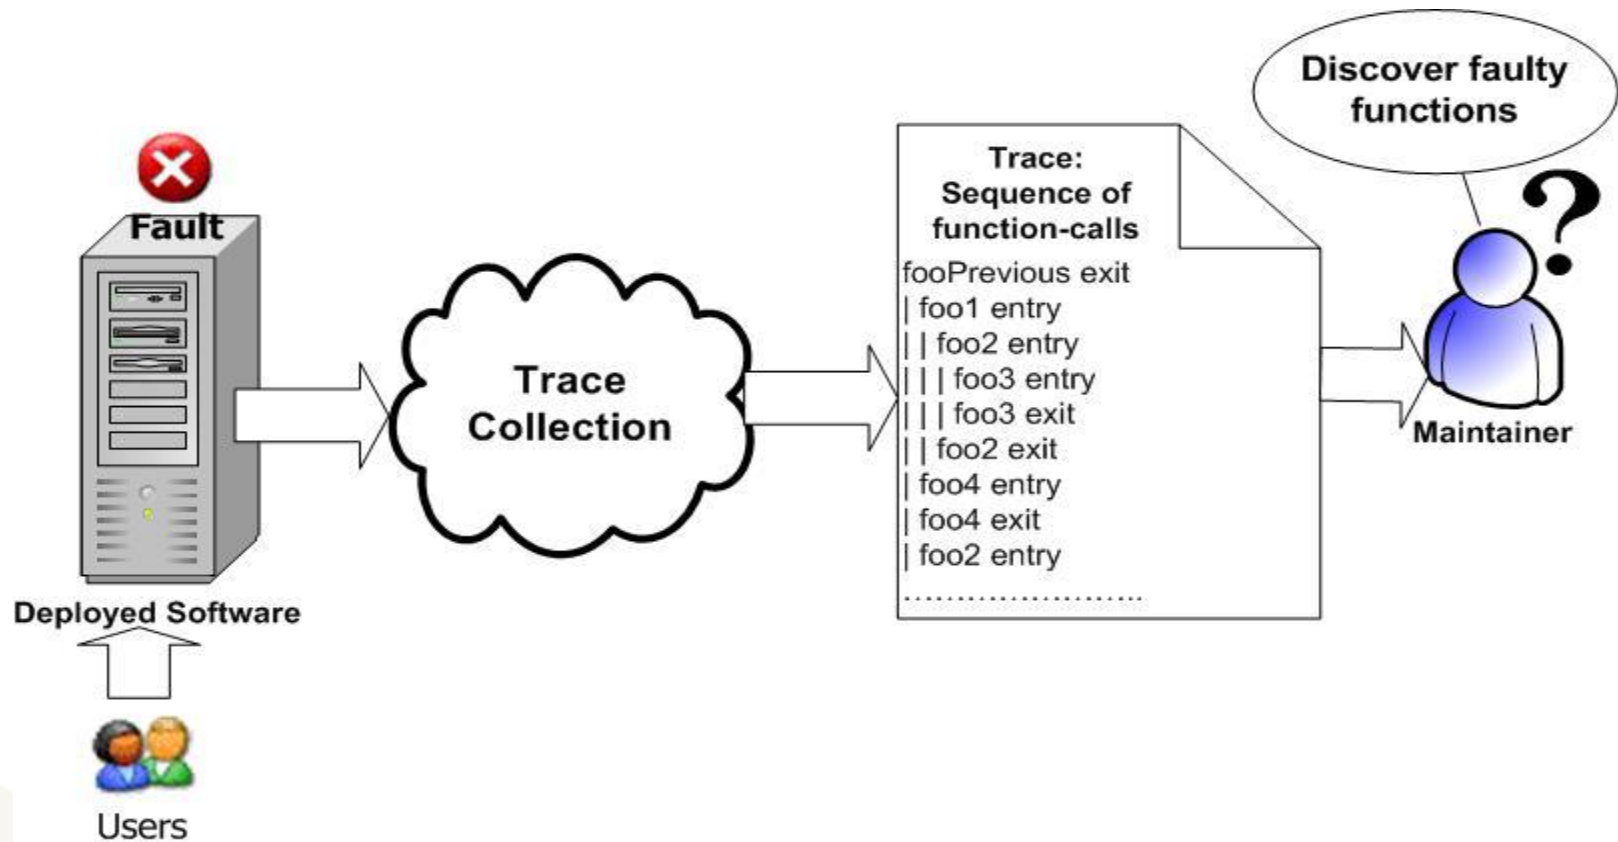
**NSERC  
CRSNG**



**Western**  
UNIVERSITY • CANADA



# Finding Faulty Functions from the Traces of Field Failures



# Approach

Train models using one-against-all approach on trace patterns to predict faulty functions in new failed traces

Ranking		
	Function	Probability
Rank 1	foo5	0.708
Rank 2	foo2	0.27
Rank 3	foo1	0.08
Rank 4	foo4	0.02

New field failed trace


```
.....
1 foo1 exit
2 foo2 entry
3 | foo10 entry
4 || foo1 entry
```

Patterns

```
foo23
foo4
foo1
foo4 → foo1
foo23 → foo4
foo23 : foo1 → foo4
```

.....

```
1 foo23 exit
2 foo4 entry
3 | foo1 entry
4 || foo4 entry
```





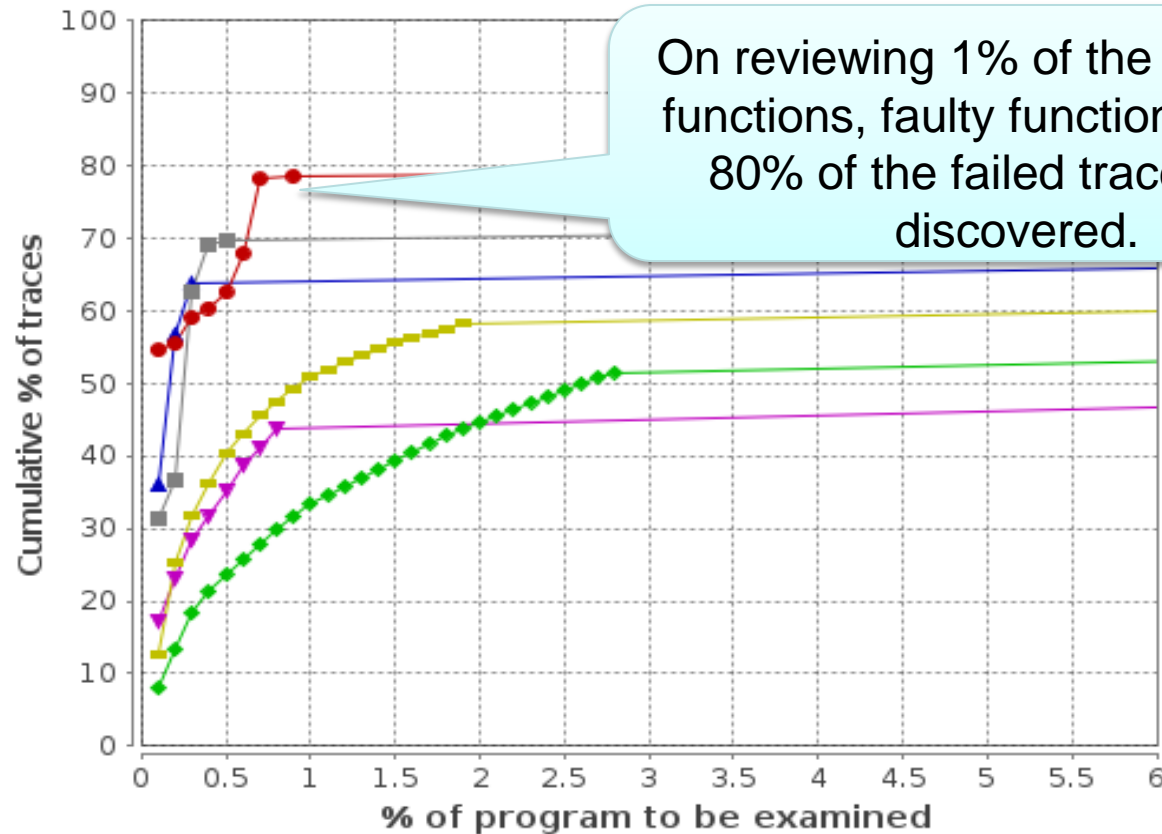
# Application to the IBM system

20+ million LOC, 300+ components, approx. 200 K+ functions, traces of size up to 4GB (44 million function-calls), and 82% rediscoveries of field faults.

	# Failed Traces	# Faulty Comp.	# Faulty Func.
Release 1	269	52	65
Release 2	337	35	47
Release 3	99	30	31
Total Distinct Faults (Union)		65	103

# Results on the IBM system

Classification on individual releases



On reviewing 1% of the program's functions, faulty functions in up to 80% of the failed traces were discovered.

- F007 on release 1
- F007 on release 2
- ▲ F007 on release 3
- ◆ Straw-man classification on release 1
- Straw-man classification on release 2
- ▼ Straw-man classification on release 3

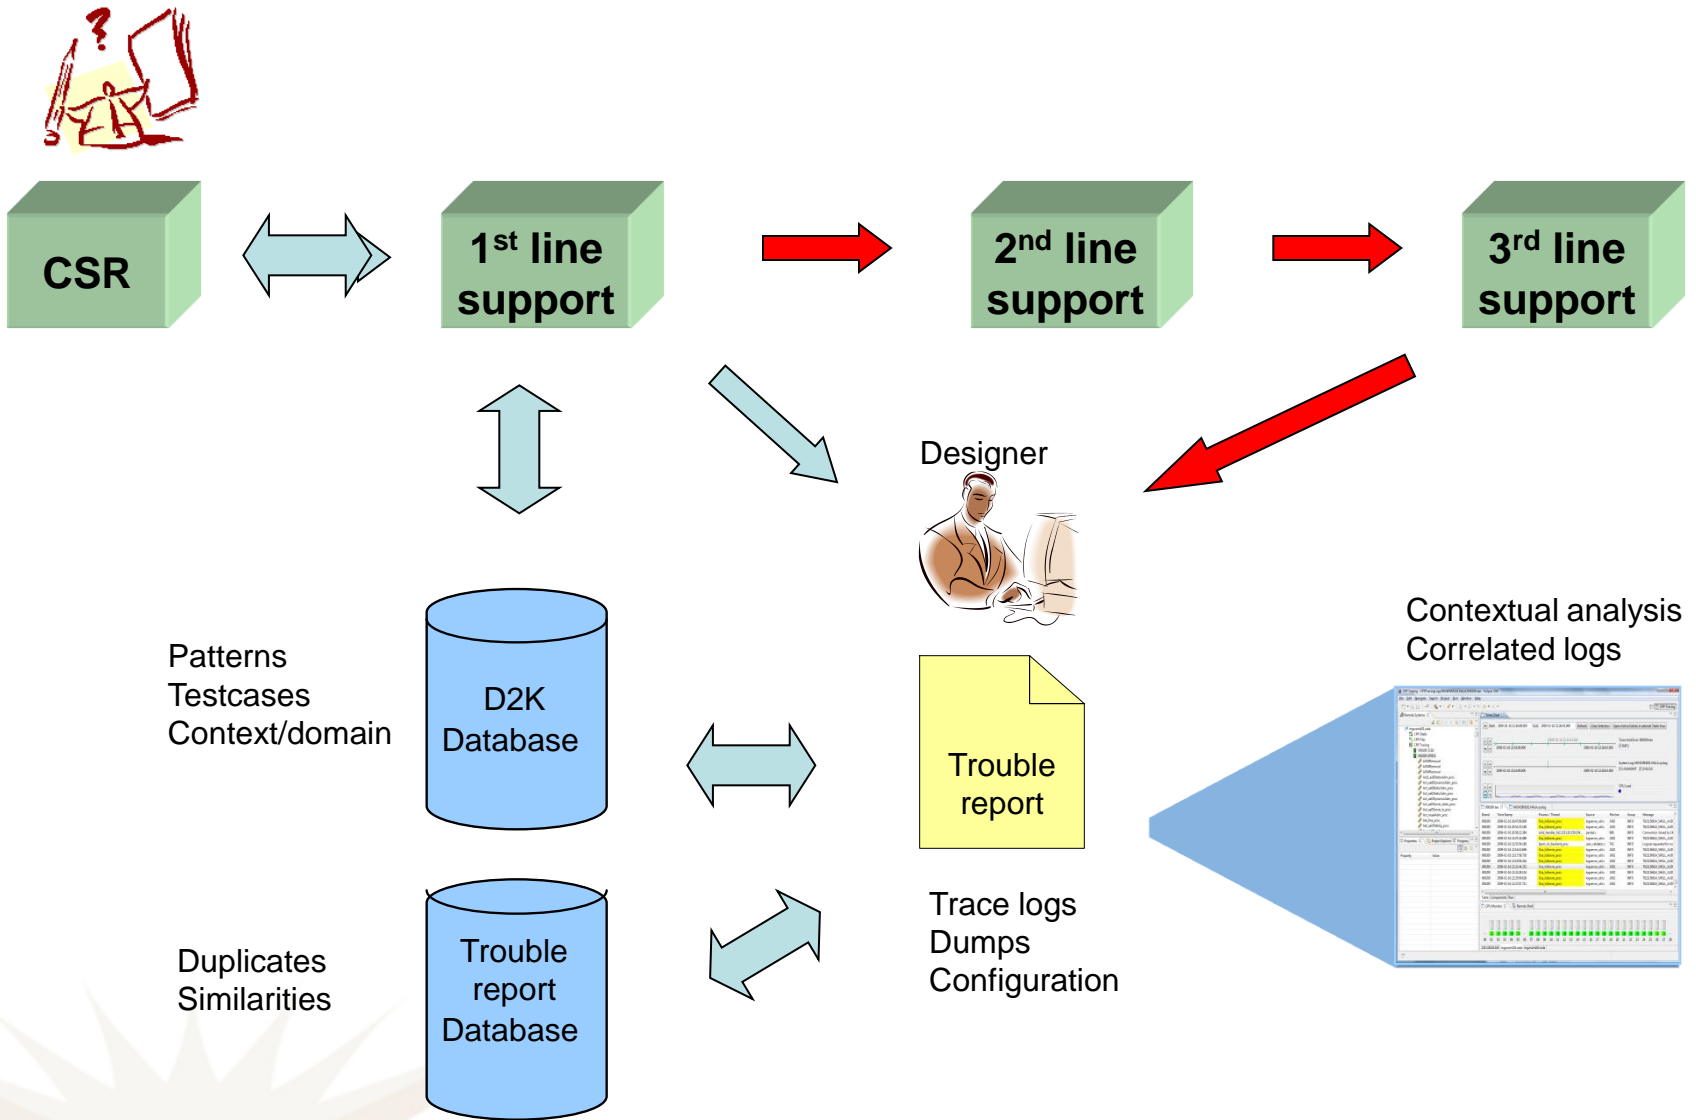
# From Data To Knowledge for Better System Maintenance - D2K Project

Enable and implement efficient use of analytical techniques to achieve revenue targets within risk limits by continuously improving the end-to-end software maintenance process



# D2K Objectives

1. Identify changes to improve current software maintenance process and information flow
2. Investigate automated solutions for fault discovery, diagnosis, and prediction
3. Provide better analysis capabilities to software engineers
4. Help software engineers focus on the real problem rather than spending time on irrelevant information



# Conclusion

Trace analysis is useful for many software engineering applications including software maintenance and evolution, performance analysis, software resilience, and cyber security

## Future

Invest in an end-to-end Enterprise Tracing Platform (ETP) for trace generation, modeling, abstraction, and analytics to support forward and background engineering tasks

**Merci!**