

Tracing and Monitoring Distributed Multi-Core Systems Project - Progress Meeting -

User Space Trace Abstraction Techniques

Heidar Pirzadeh and Wahab Hamou-Lhadj

Software Behaviour Analysis Lab

{s_pirzad, abdelw}@ece.concordia.ca

Montreal, QC

Dec. 09, 2011

Progress


- Trace abstraction:
 - We continued to develop trace abstraction techniques for user space traces
 - Explored the use of state information in trace abstraction and exploration
 - Developed techniques for automatically extracting important content from a trace
- Anomaly detection:
 - Investigation of different tracing mechanisms
 - Reduction of learning time in building models
 - Reduction of false positives
 - Development of a taxonomy of attacks on the Linux kernel

Our Approach for Trace Abstraction

- Based on the extraction of execution phases from large traces
- What is an execution phase?
 - A segment of program's execution that performs a specific task
- Trace Segmentation: Automatically divide a trace into phases
 - Allow SW engineering to browse traces as a flow of execution phases rather than mere sequence of events


Example

- A trace generated from a compiler will contain the various compiler's phases including parsing, preprocessing, lexical analysis, semantic analysis, etc.
- In most visualization tools, it will look like:

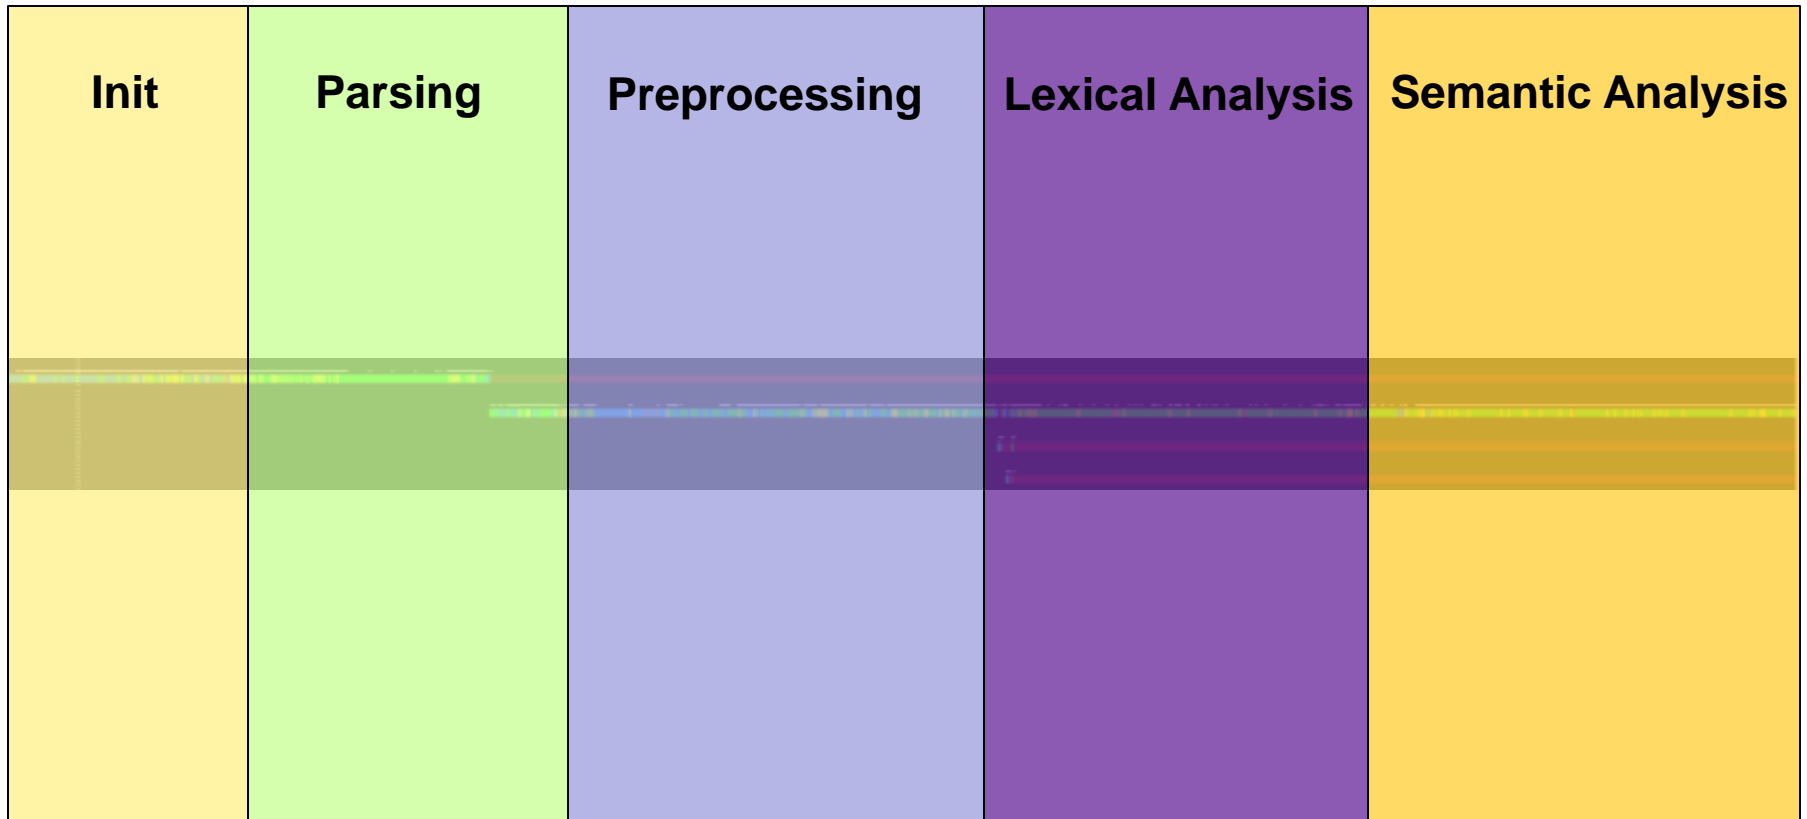


- But how can we tell what happens where?

Visually...




Visually...



A different view...

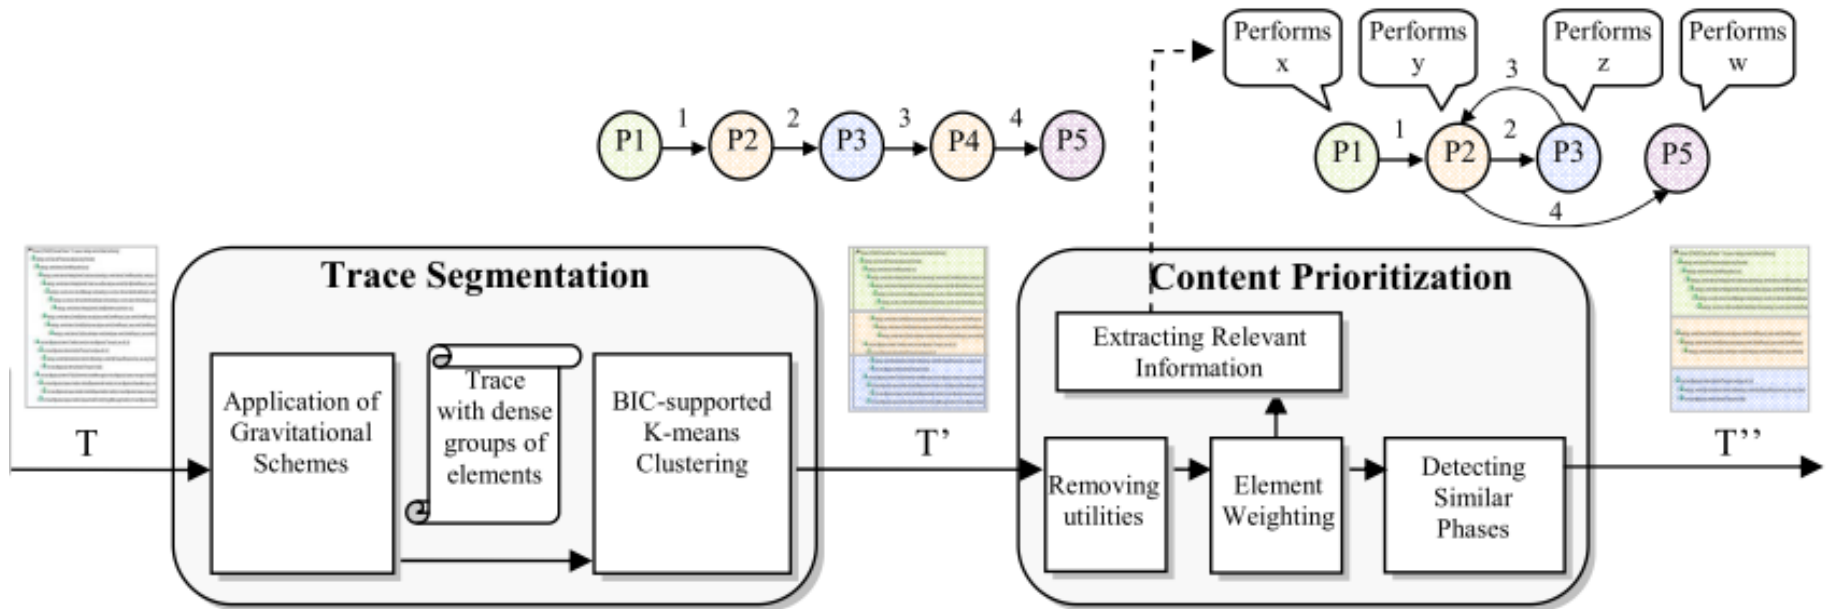
Nested phases can be added



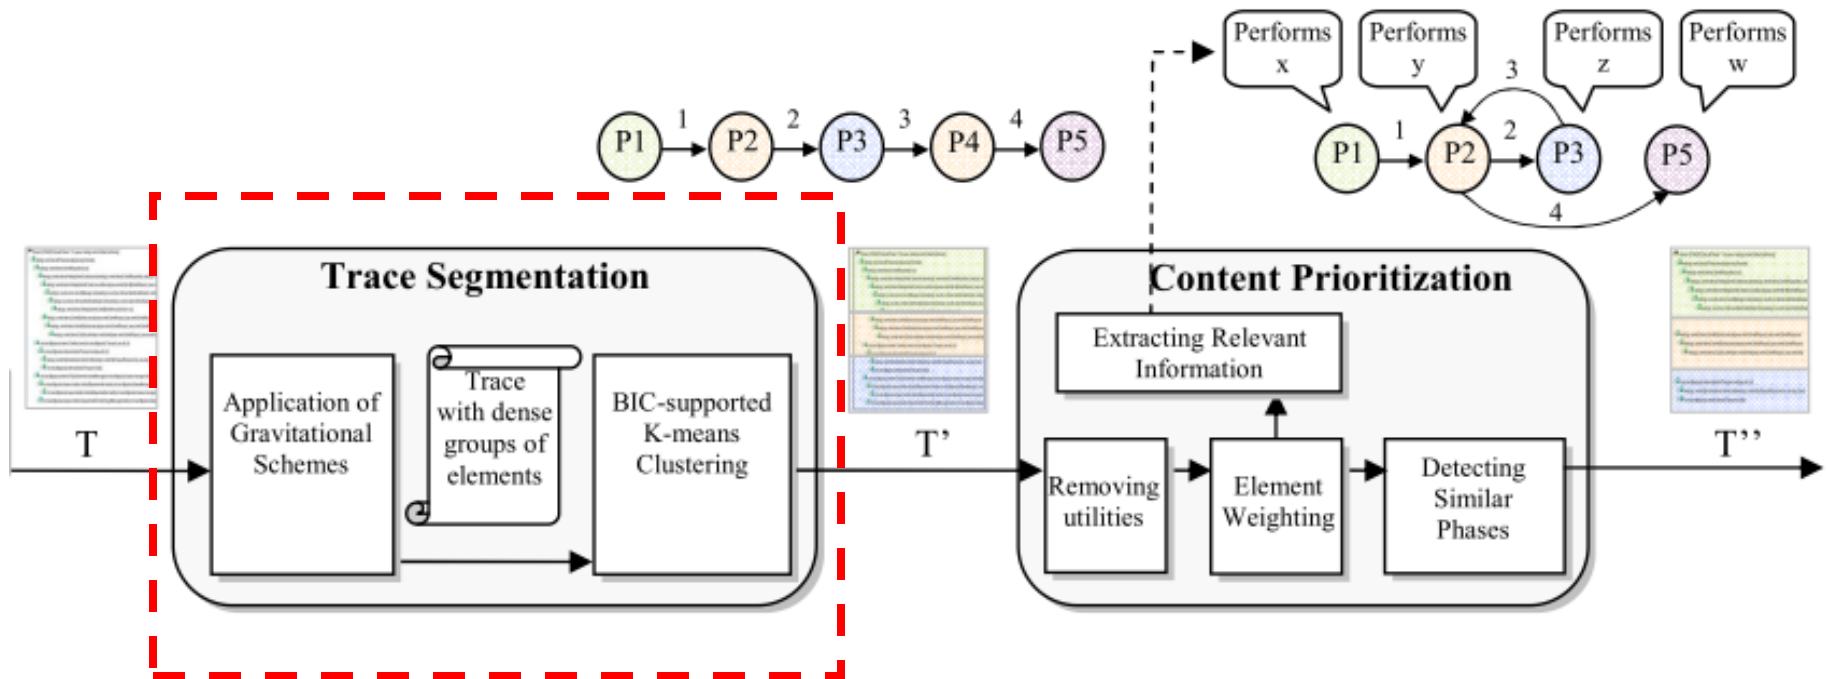
Research Questions?

- How can we automatically extract execution phases from a trace?
- What additional information states can reveal about execution phases?
- How can we extract the main components that implement a specific phase?
- Can we use execution phases to further reduce the size of traces?

Approach: Trace Abstraction Framework




Our Approach: Trace Abstraction Framework




Trace Segmentation Approach

- The scientific foundation comes from the study of the human perception system
 - The ability for humans to group similar items to form objects and shapes
 - Explained using the Gestalt laws of similarity and continuity


Measuring Similarity



Measuring Continuity in Traces with Nesting Levels



Measuring Continuity in Traces with Nesting Levels




Case Study

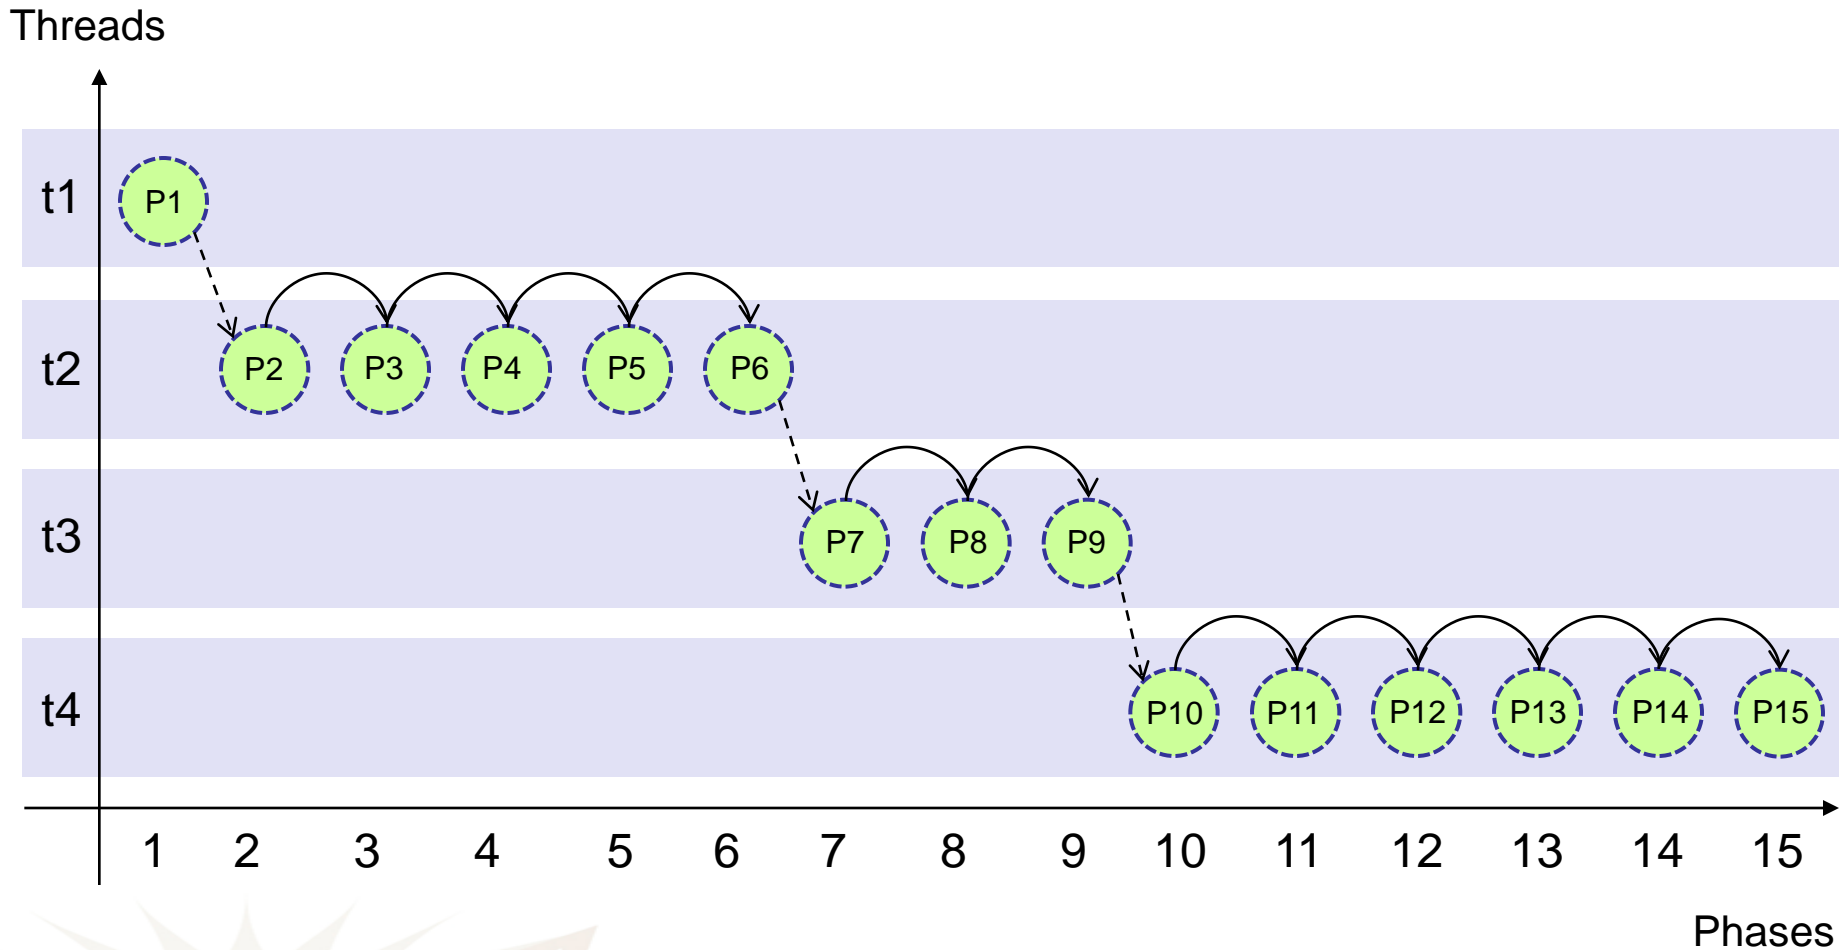
Program: WEKA 3.6.6

Scenario: building a decision tree learning algorithm for classifying data instances.


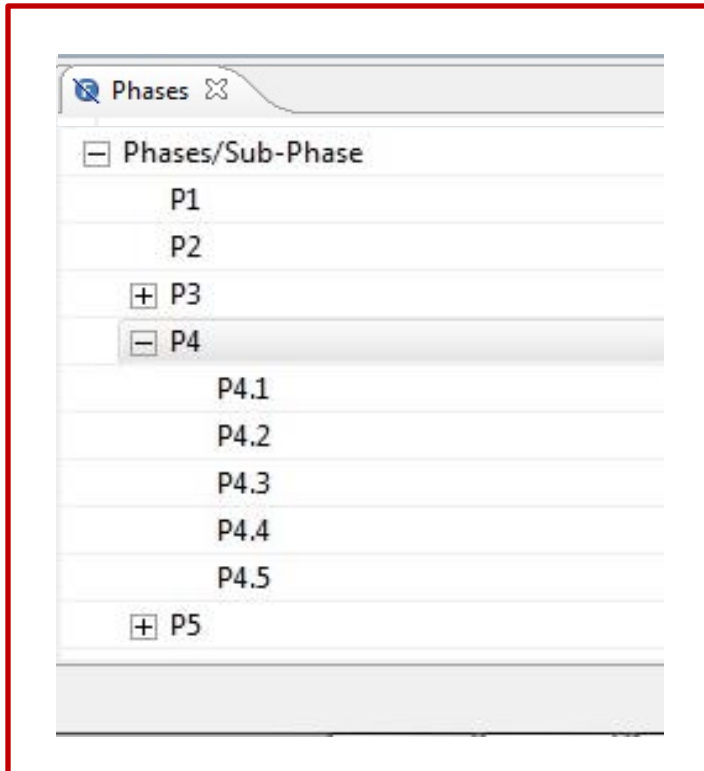
Trace: Multi-threads 1,571,214 events



Phase flow diagram of a Weka trace



Adding phase views to a tool



State Information

- What is a state?
 - The state of the system is the state value of every attribute in the system
 - State has a duration
 - State value, which can really be anything
- Attributes in the kernel-trace state system :
 - CPUs
 - CPUs/0
 - CPUs/0/current_thread
 - Etc.

State Change

Consists of three things:

- timestamp
- attribute
- state value

The state of 'attribute' changed to 'state value'
at time 'timestamp'

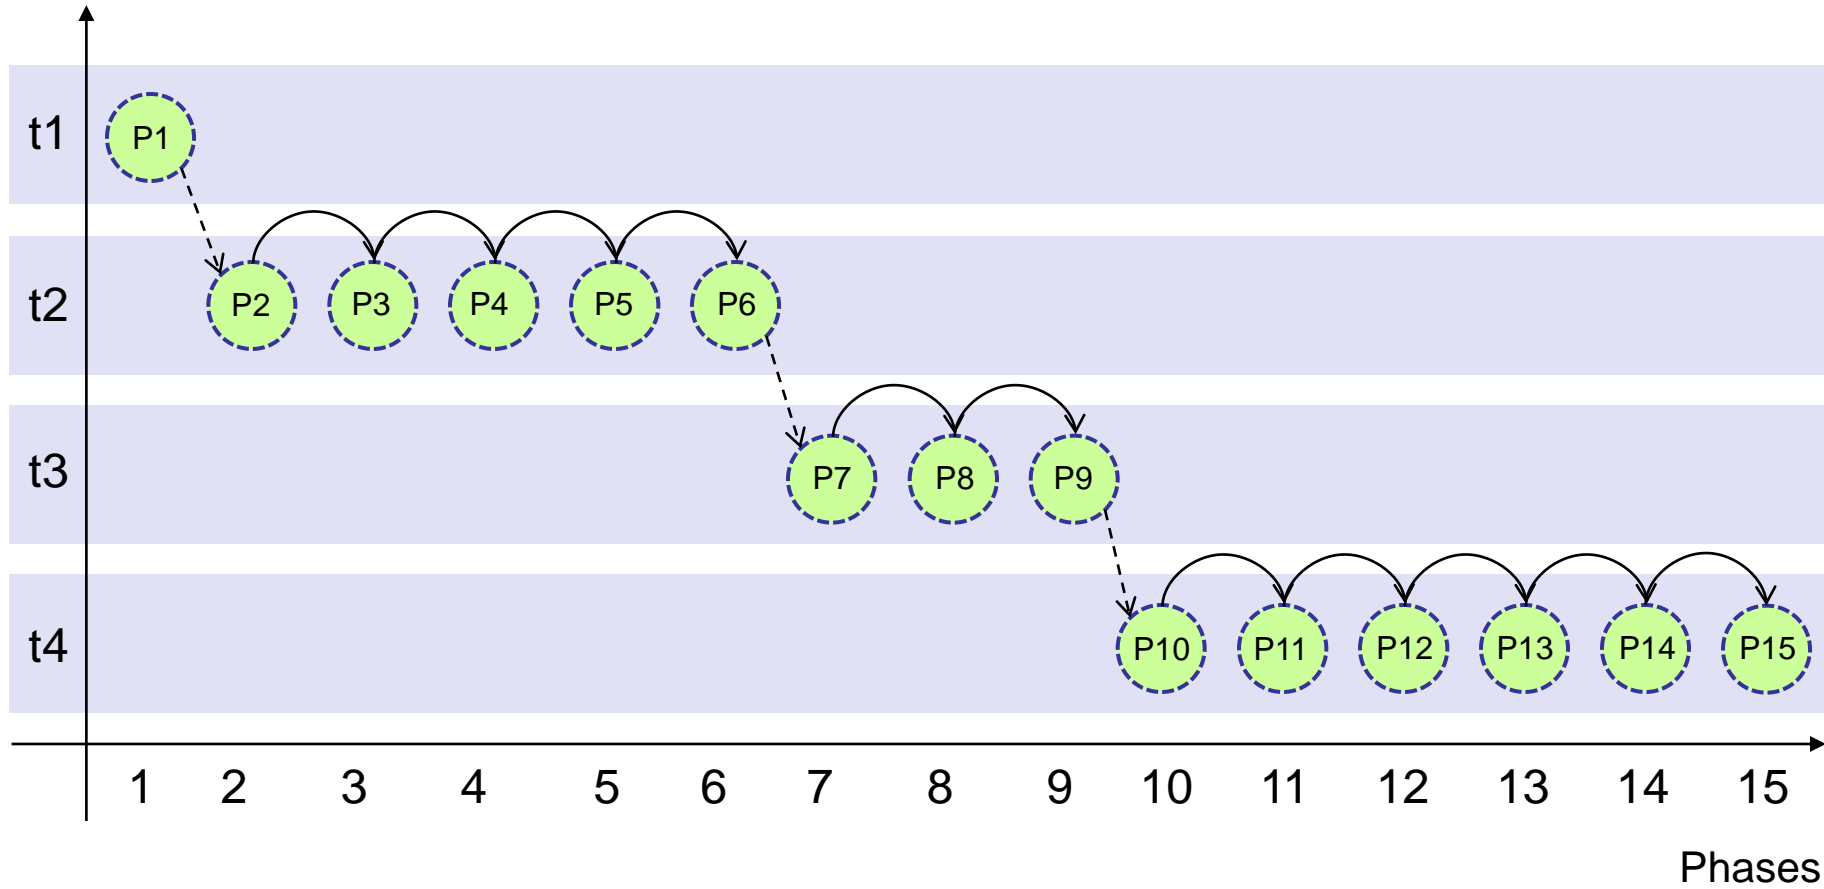
Existing Info

LTTNG Kernel Space Trace:

- Timestamp
- Event (page fault)
- Process ID
- CPU ID
- File Descriptor

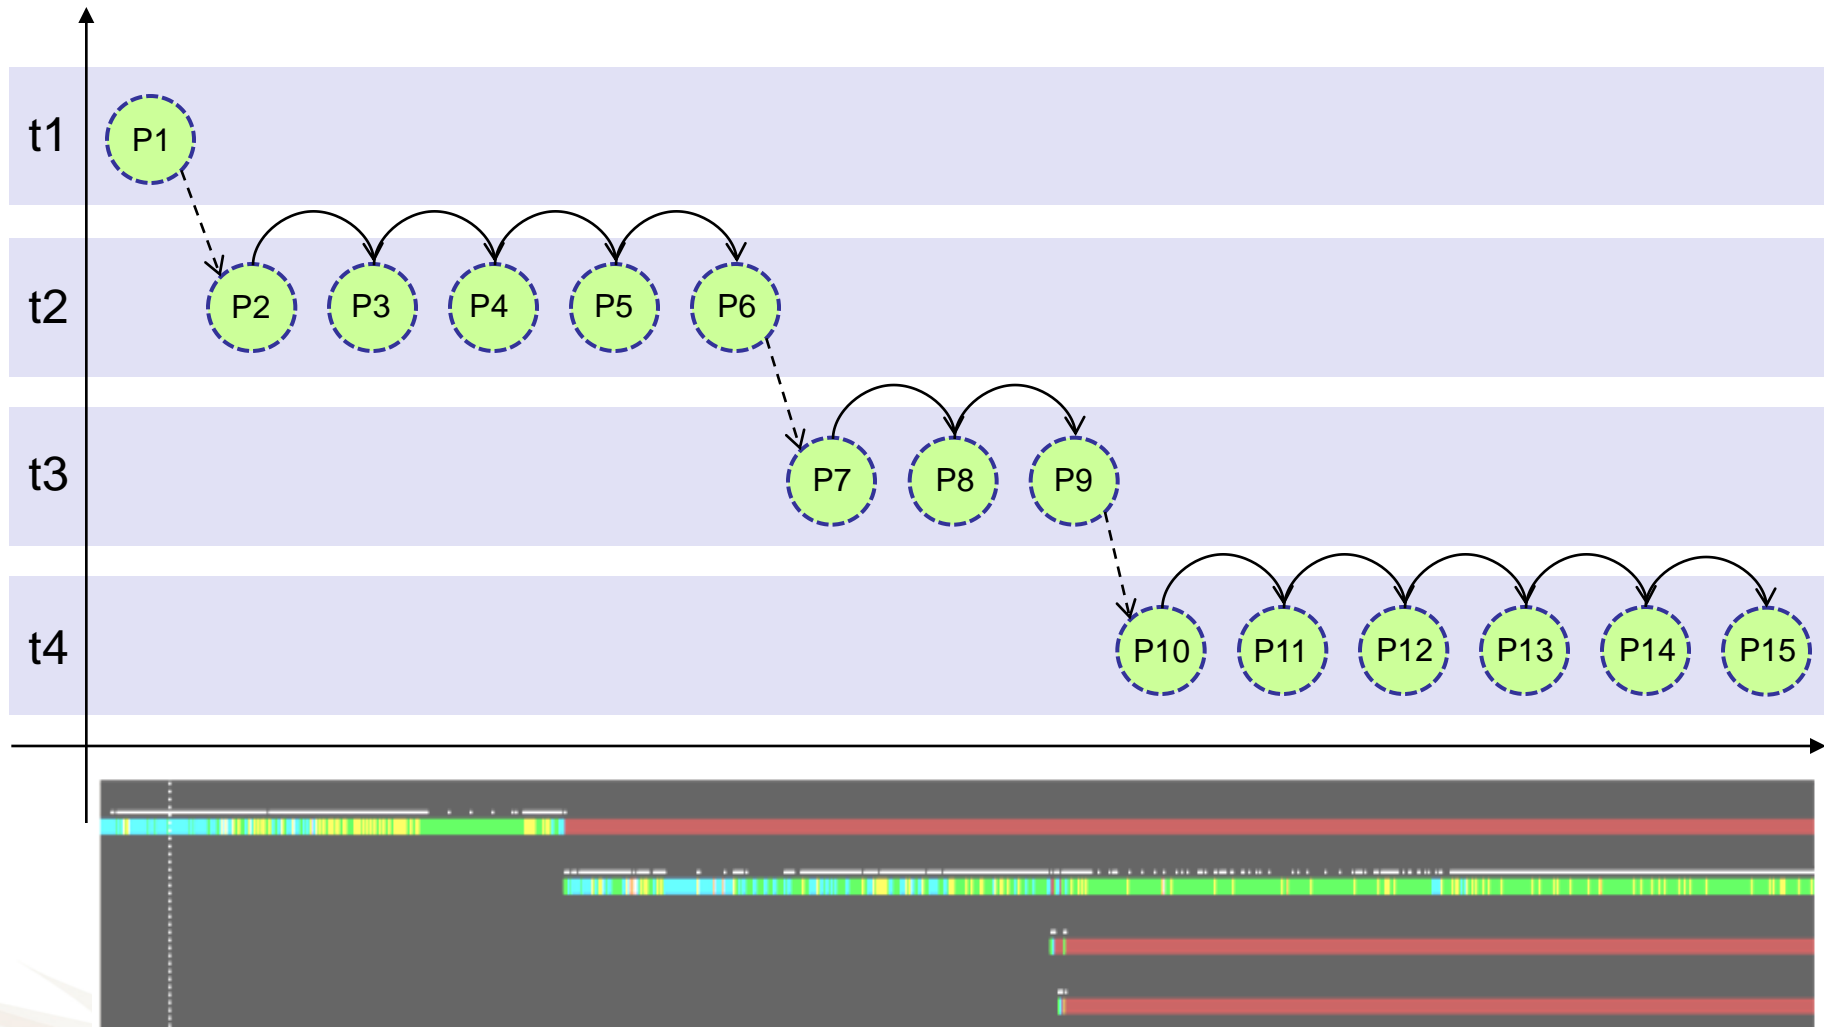
Phase Flow

Threads



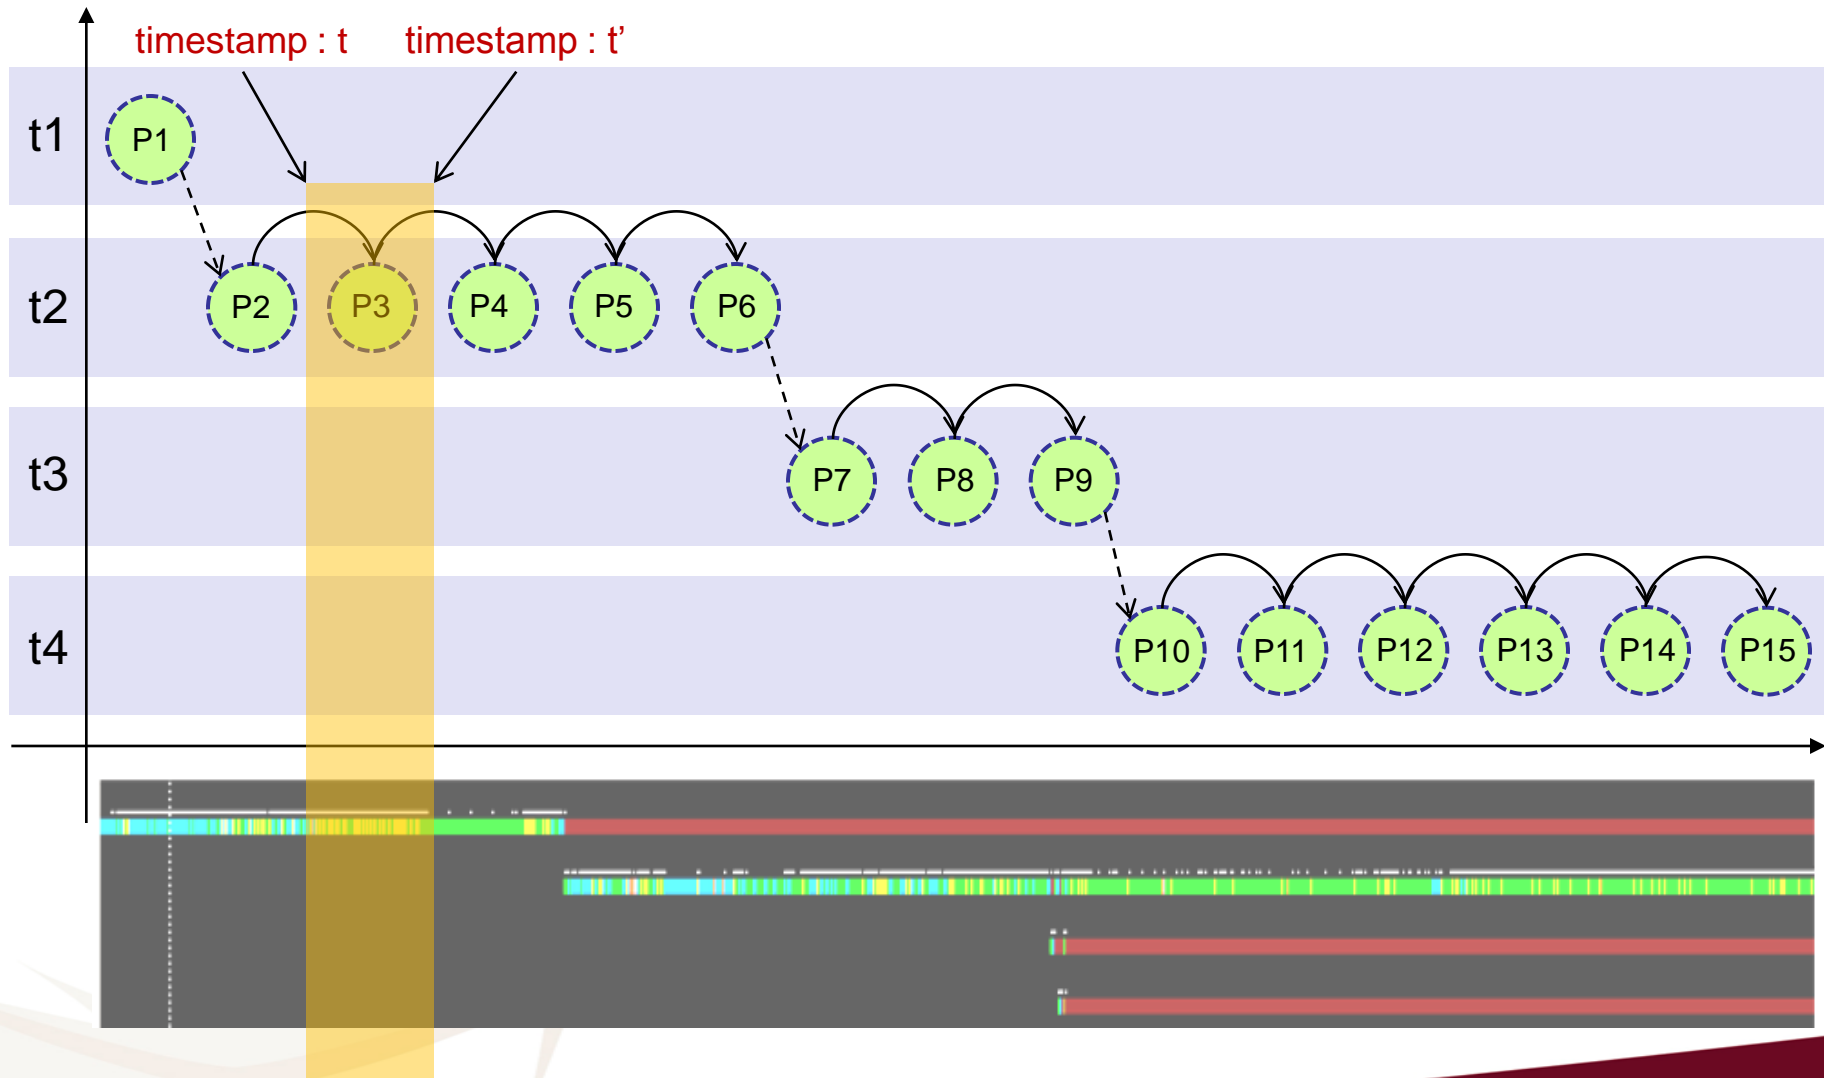
Phases Mapped to Kernel Space Trace

Threads



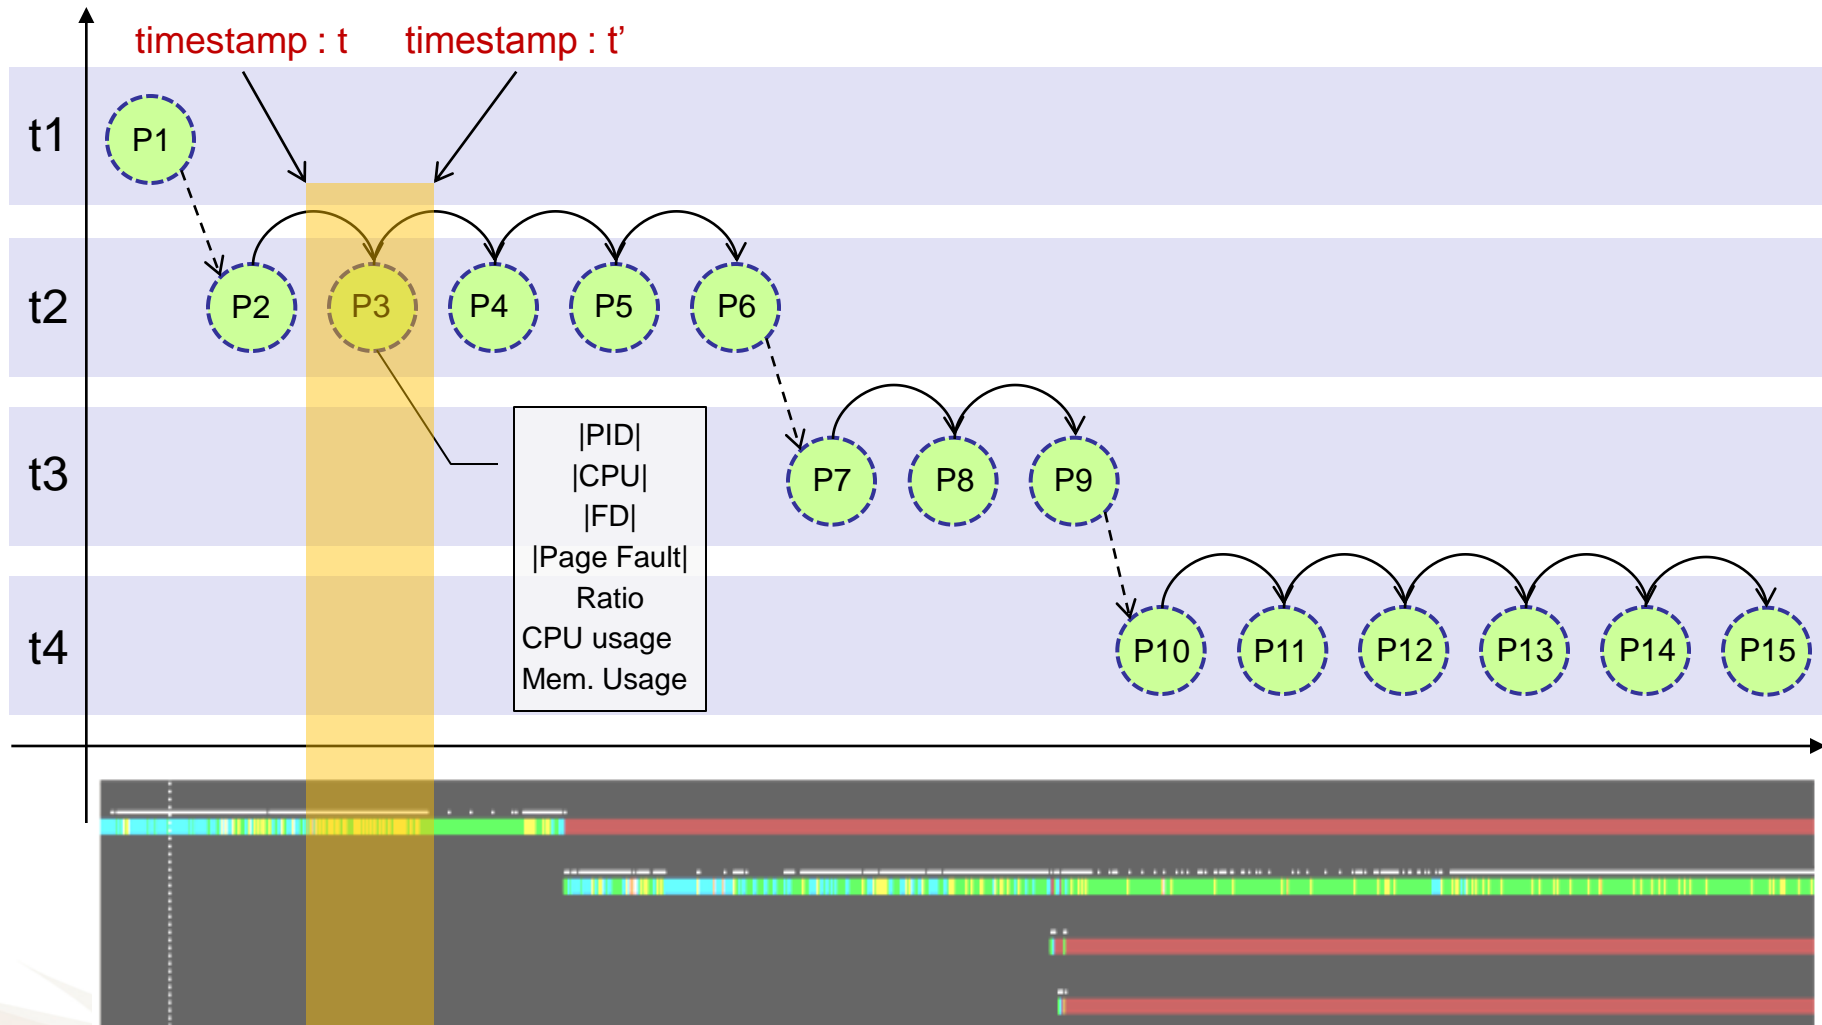
Phases Mapped to Kernel Space Trace

Threads



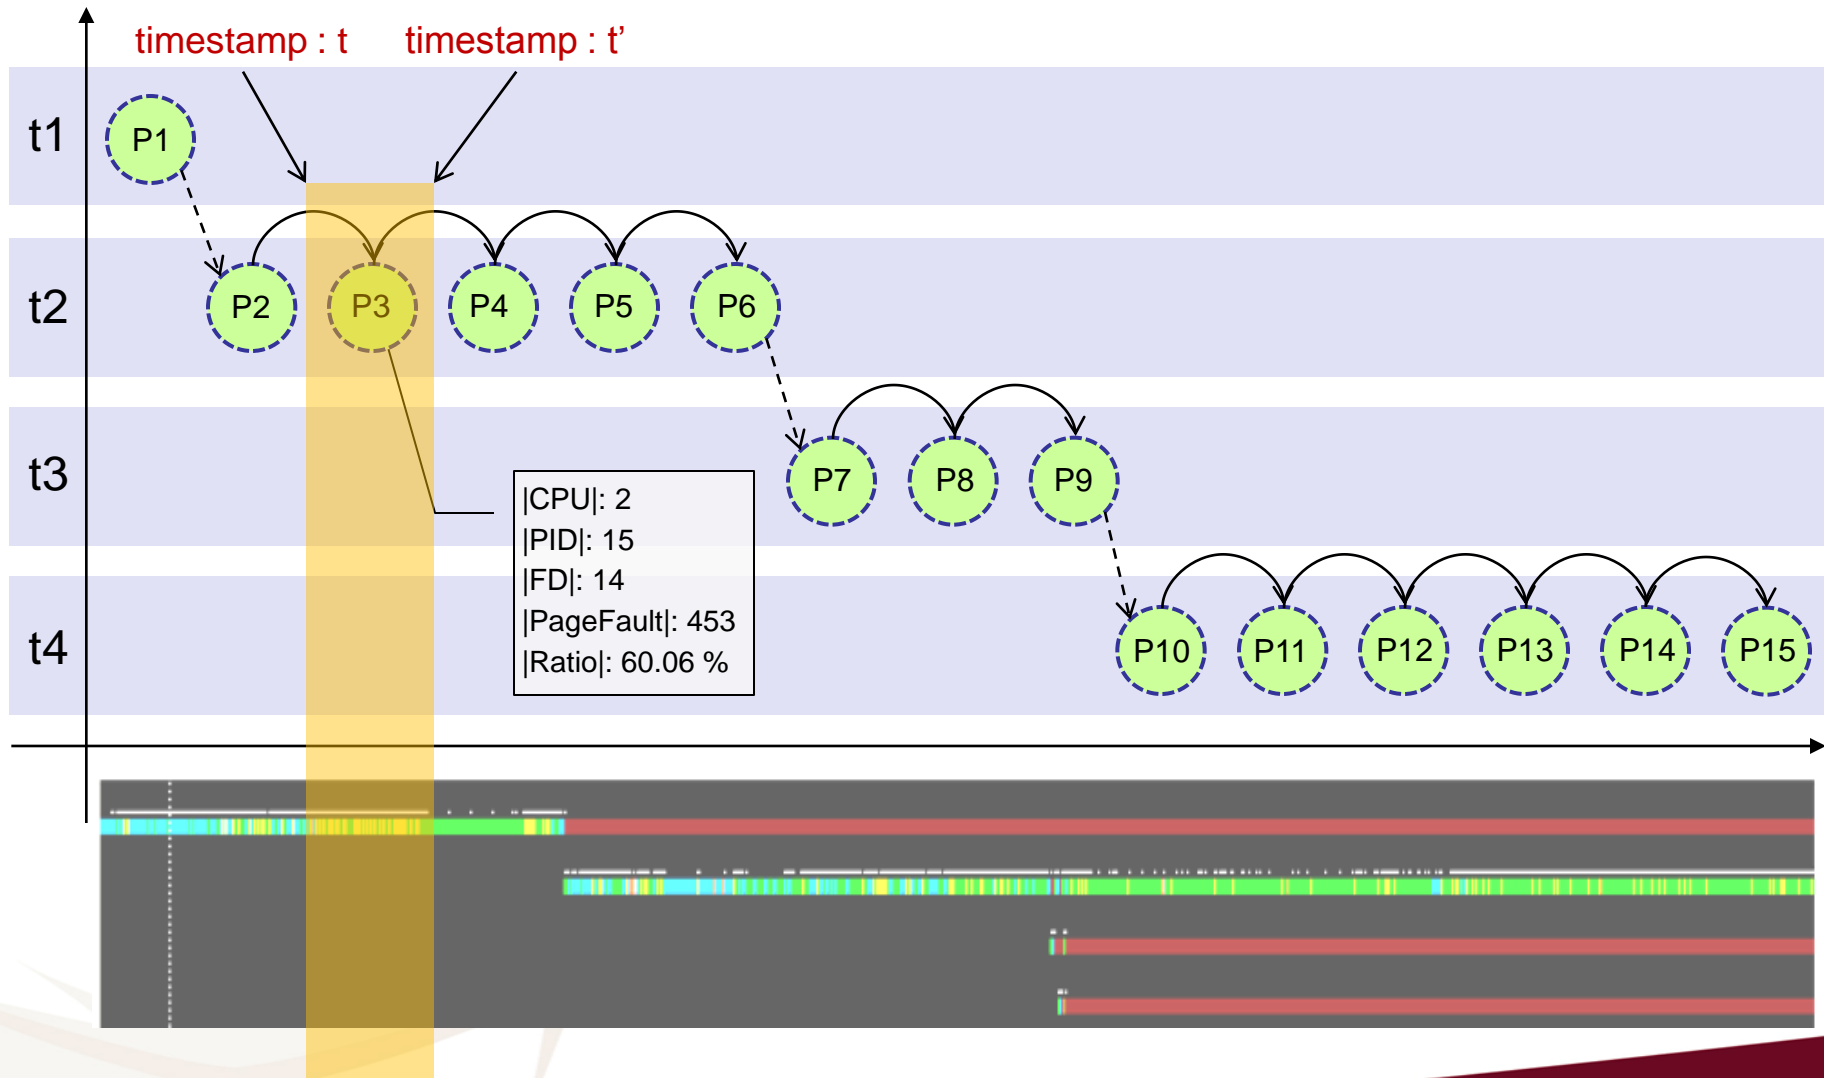
Phases Enriched with State Info

Threads



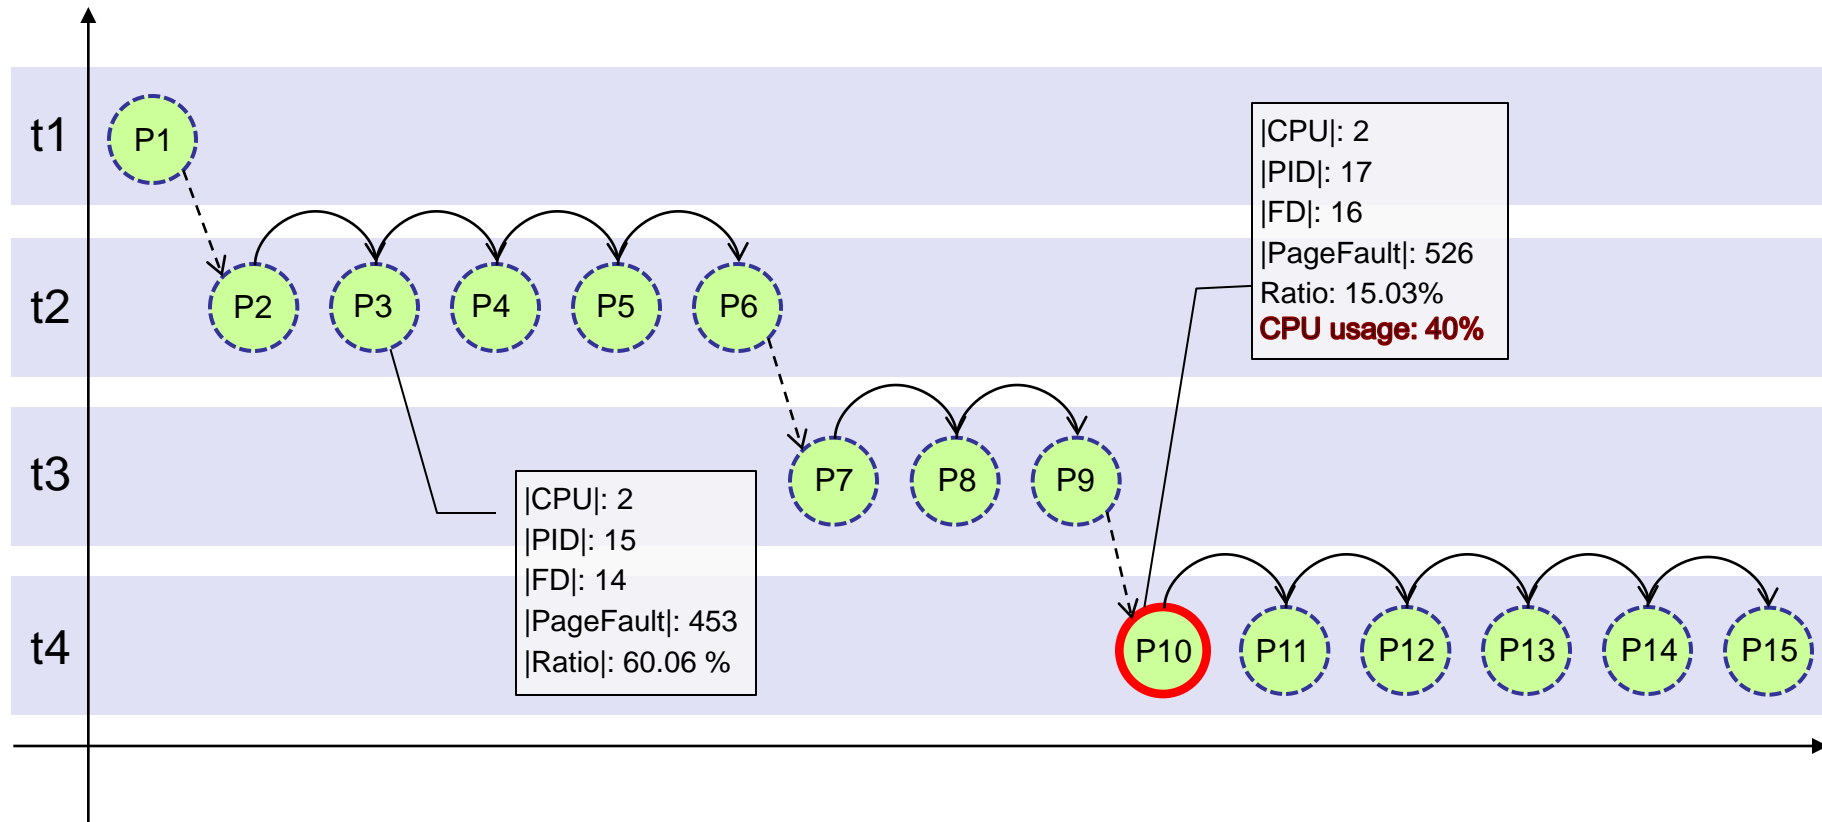
Phases Enriched: Statistics (1)

Threads




Phases Enriched: Statistics (2)

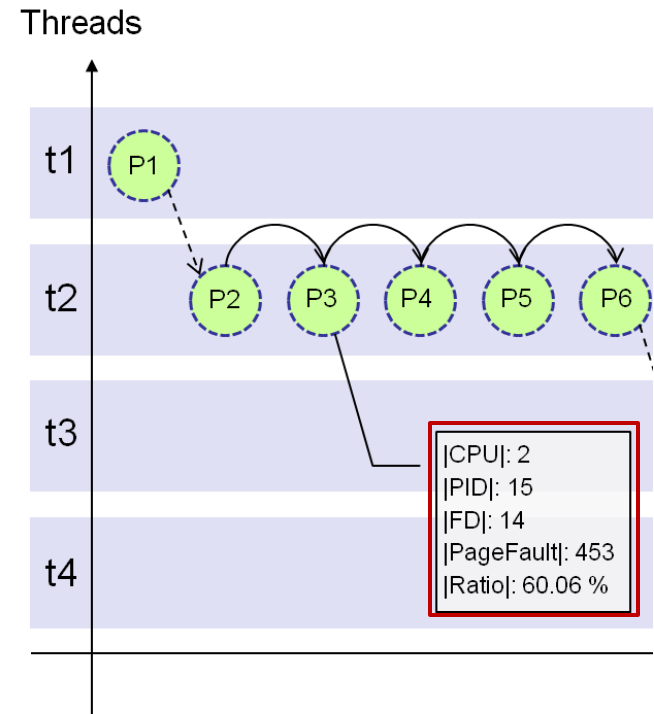
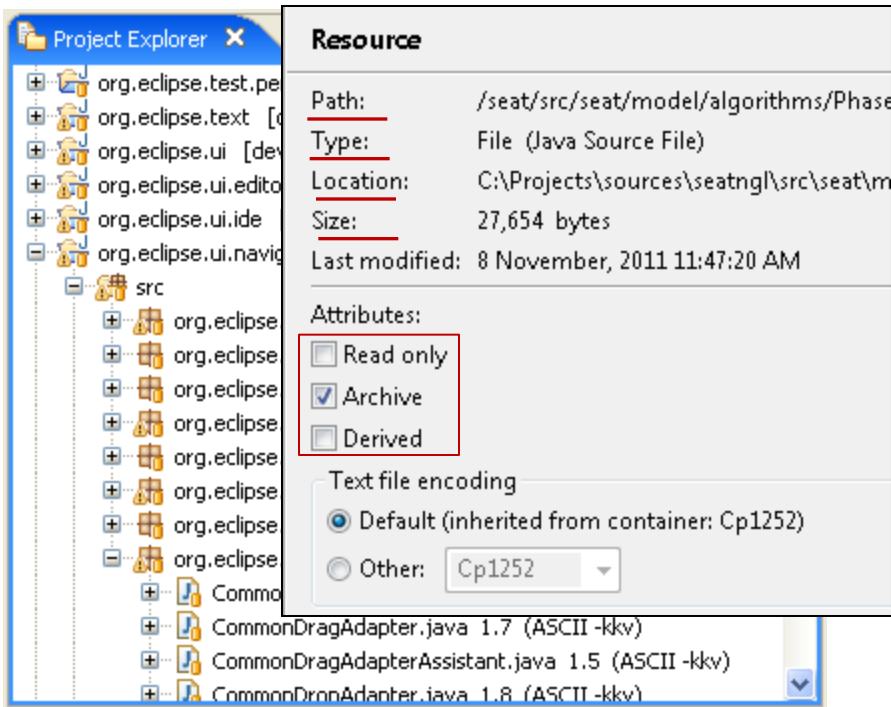
Threads



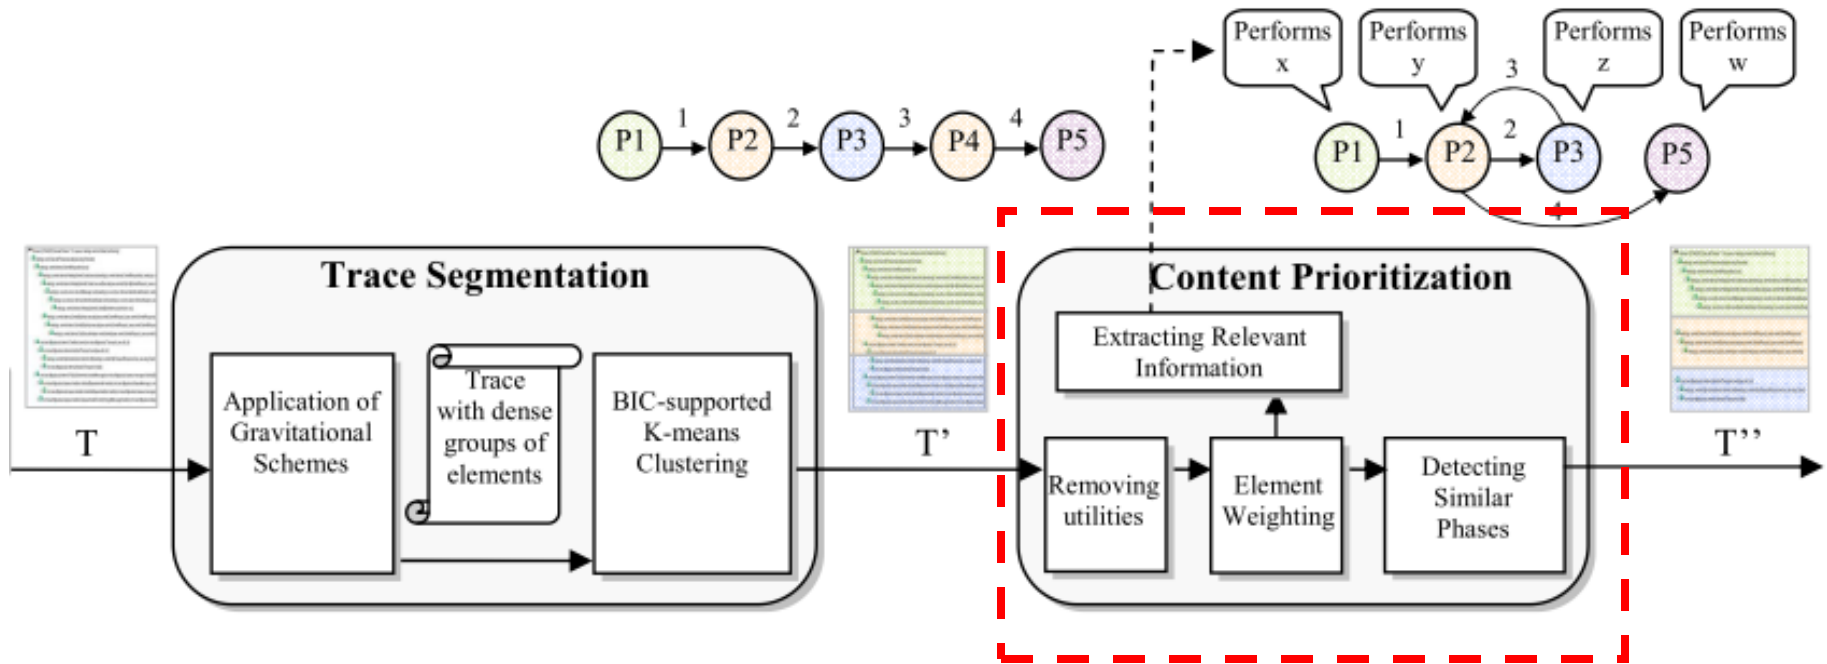
Comparison: Kernel Space vs User Space



Enriched Phase View

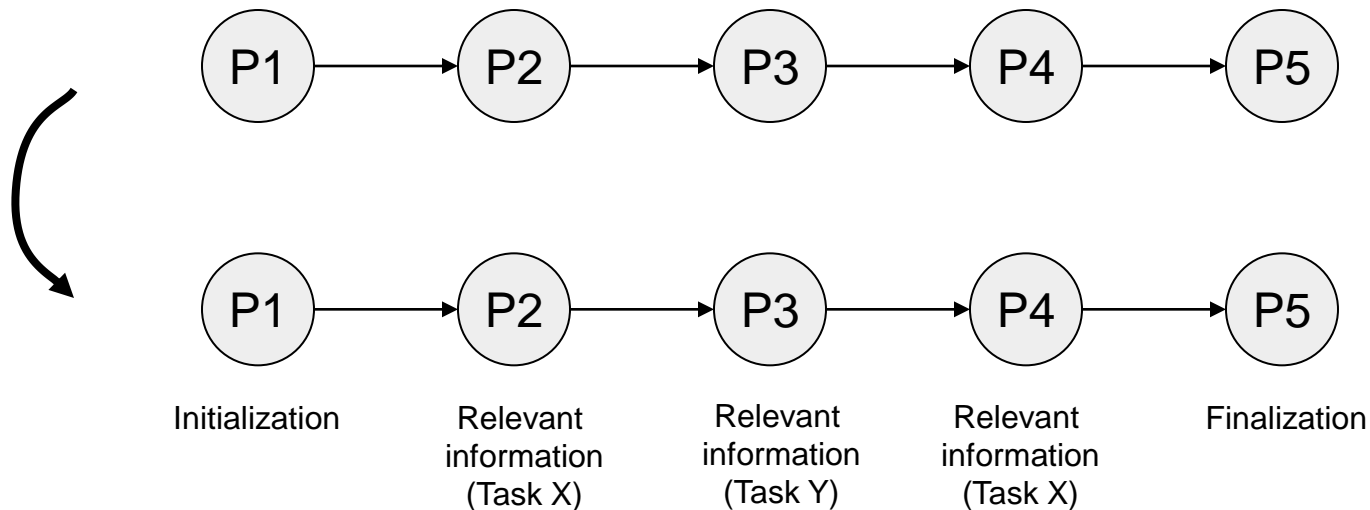


Approach: Trace Abstraction Framework



Content Prioritization

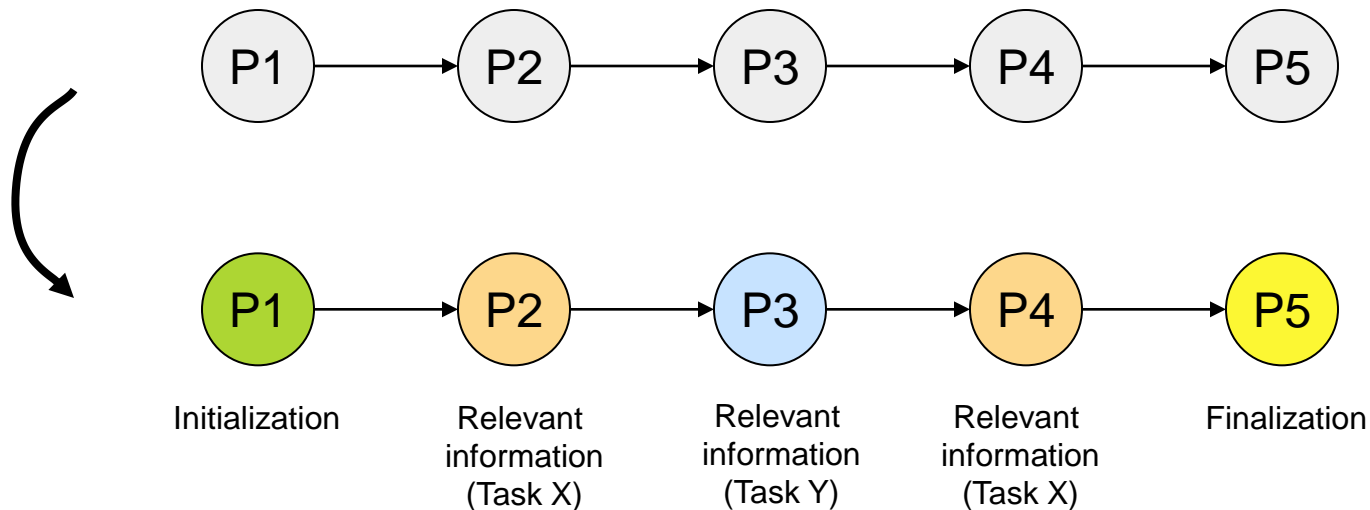
1. Extract representative elements of each phase



- Can give a hint about what is happening in a phase
- Uncover the most relevant elements that implement the traced scenario

Content Prioritization

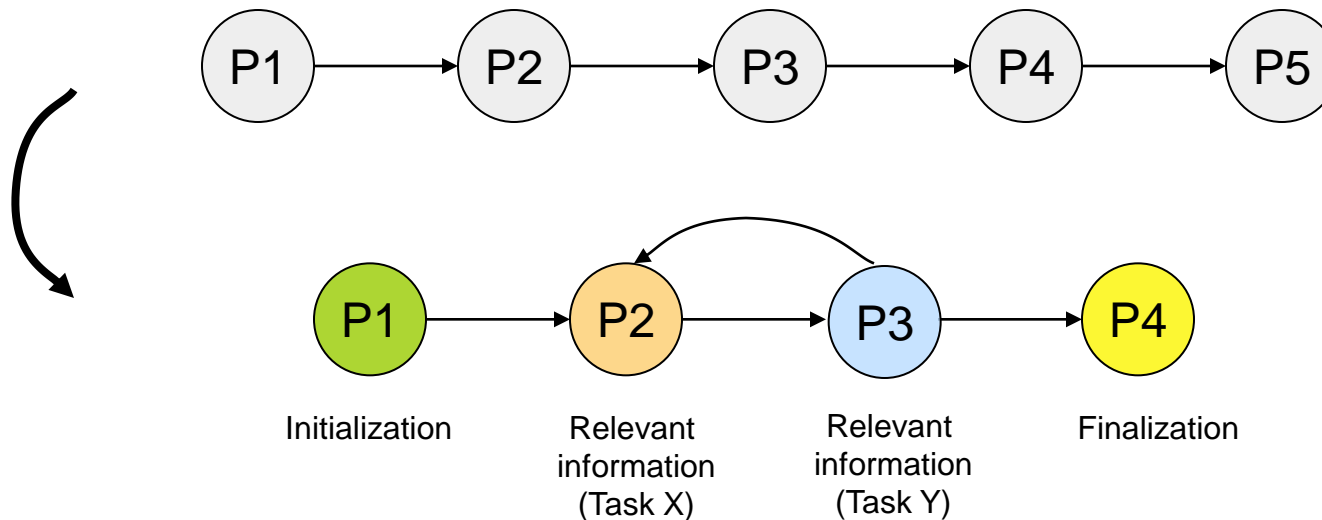
2- Finding similar phases



- Can give a hint about what is happening in a phase
- Uncover the most relevant elements that implement the traced scenario


Content Prioritization

2- Finding similar phases



- Can give a hint about what is happening in a phase
- Uncover the most relevant elements that implement the traced scenario
- Optimized flow of phases

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

Extracting Representative Elements

Trace T	$L_{i,k}$	IG_i	$w_{i,k}$	
Phase 1	a	$L(a)=1$	$G(a)=0.17$	$w(a)=0.45$
	d	$L(d)=1.3$	$G(d)=0.17$	$w(d)=0.58$
	c	$L(c)=1.4$	$G(c)=0.17$	$w(c)=0.66$
	i	$L(i)=1$	$G(i)=0$	$w(i)=0$
	d	$L(e)=1$	$G(e)=0$	$w(e)=0$
	c			
	e			
Phase 2	h	$L(h)=1$	$G(h)=0.47$	$w(h)=0.50$
	i	$L(i)=1$	$G(i)=0$	$w(i)=0$
	m	$L(m)=1.3$	$G(m)=0.47$	$w(m)=0.65$
	n	$L(n)=1$	$G(n)=0.47$	$w(n)=0.50$
	e	$L(e)=1$	$G(e)=0$	$w(e)=0$
	o	$L(o)=1.3$	$G(o)=0.17$	$w(o)=0.24$
	o			
Phase 3	a	$L(a)=1$	$G(a)=0.17$	$w(a)=0.41$
	o	$L(o)=1$	$G(d)=0.17$	$w(d)=0.53$
	d	$L(d)=1.3$	$G(c)=0.17$	$w(c)=0.60$
	c	$L(c)=1.4$	$G(i)=0$	$w(i)=0$
	i	$L(i)=1$	$G(e)=0$	$w(e)=0$
	d	$L(e)=1$	$G(o)=0.17$	$w(o)=0.41$
	c			

$$w_{i,k} = \frac{\overbrace{(\log(ef_{i,k}) + 1) * \log(N / n_i)}^{L_{i,k}} \overbrace{IG_i}^{IG_i}}{\sqrt{\sum_{j=1}^e \underbrace{[(\log(ef_{j,k}) + 1) * \log(N / n_j)]^2}_{\frac{1}{N_k}}}}$$

$$w^{d", Phase 1} = \frac{1.3 * 0.17}{\sqrt{(0.17)^2 + (0.22)^2 + (0.26)^2}} = 0.45$$



Relevant Events Snapshots

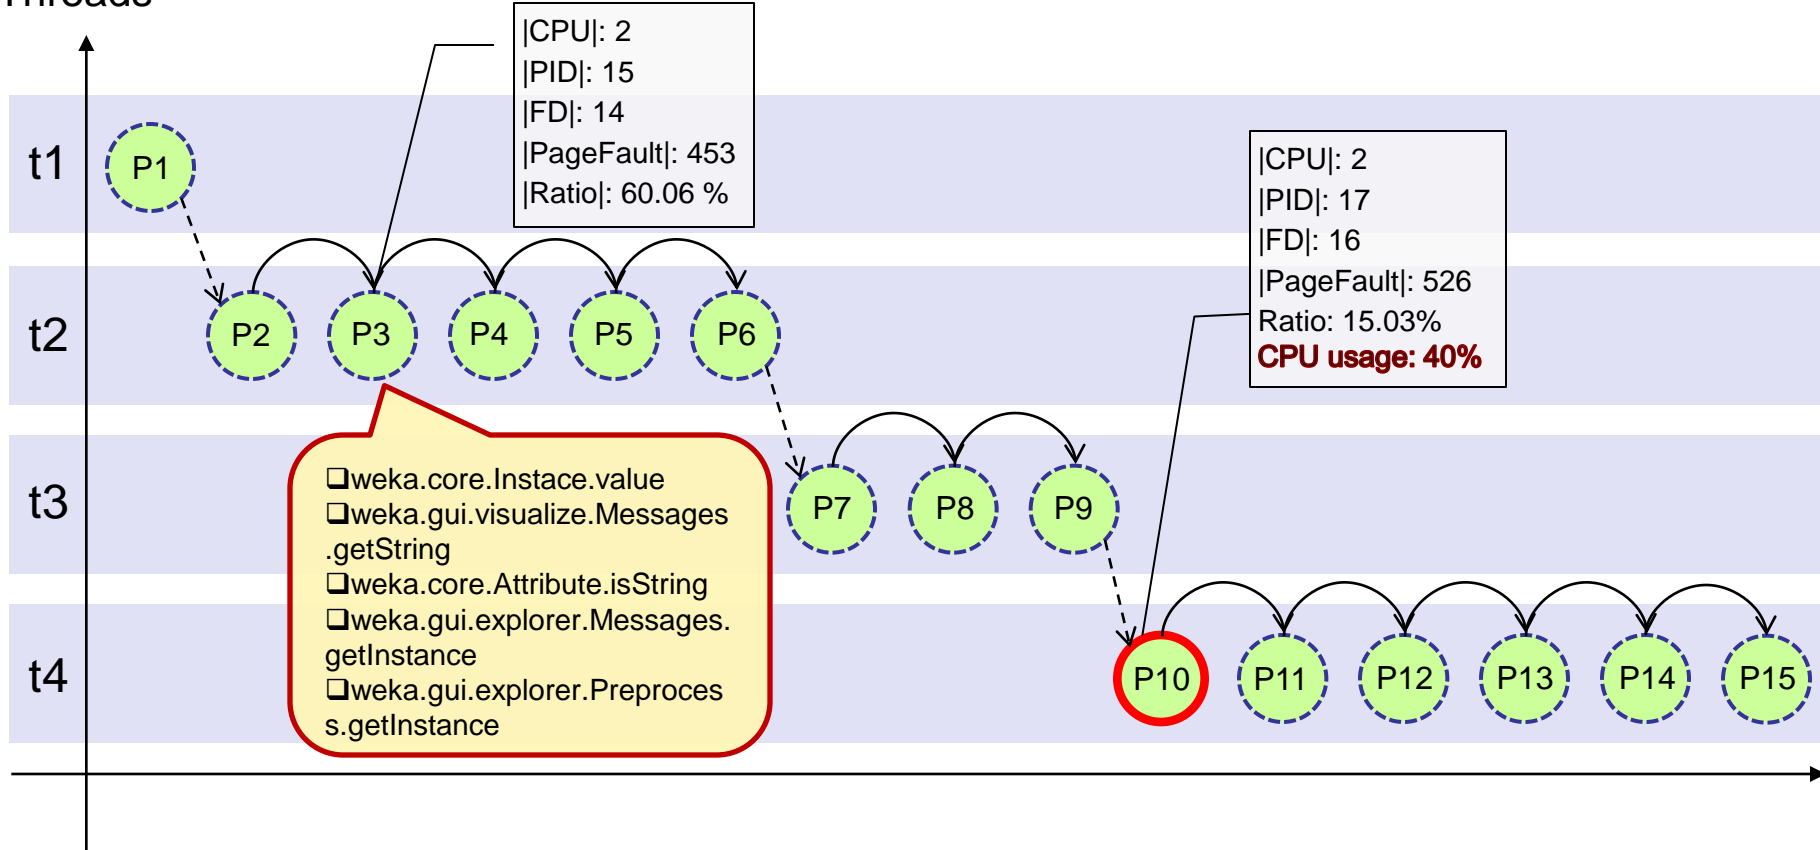
The screenshot displays the Eclipse IDE interface for a Java project named 'Trace/Argouml.trx'. The main workspace shows a phase diagram with five sequential phases labeled P4.1, P4.2, P4.3, P4.4, and P4.5, connected by red arrows. The 'Representatives' view on the right lists the following methods for phase P4:

- org.argouml.uml.cognitive.critics.CrUML.-init-
- org.argouml.cognitive.Critic.-init-
- org.argouml.cognitive.Critic.getHeadline
- org.argouml.cognitive.Critic.setDescription
- org.argouml.cognitive.Designer.getAgency
- org.argouml.uml.cognitive.checklist.Init.newCheckItem
- org.argouml.cognitive.checklist.CheckManager.register
- org.argouml.cognitive.checklist.CheckItem.setCategory
- org.argouml.uml.cognitive.critics.CrUML.getLocalizedString

The 'Phases' view at the bottom shows a tree structure with phases P1, P2, P3, P4, and P5. Phase P4 is highlighted with a red box, indicating its selection.

Case Study: Relevant Events

Threads



Conclusions

- We showed trace abstraction techniques based on execution phases
- We added state information to extracted phases
- We presented techniques for identifying the most relevant components of each phase

Thank you!