




Performance Issues Detection by Comparing Nested Execution Traces in Distributed Systems

Maryam Ekhlas
May 16th, 2022

Polytechnique Montreal
DORSAL Laboratory

Motivation

- Help users understand what is happening when something is different.
 - Performance differences between versions of the application.
 - Performance differences between system configurations.
 - Why some requests are slower than others?



Previous Work

IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, VOL. 47, NO. 2, FEBRUARY 2021

243

Fault Analysis and Debugging of Microservice Systems: Industrial Survey, Benchmark System, and Empirical Study

Xiang Zhou^{ID}, Xin Peng^{ID}, Tao Xie, *Fellow, IEEE*, Jun Sun, Chao Ji, Wenhai Li, and Dan Ding

2020 IEEE 20th International Working Conference on Source Code Analysis and Manipulation (SCAM)

DepGraph: Localizing Performance Bottlenecks in Multi-Core Applications Using Waiting Dependency Graphs and Software Tracing

Naser Ezzati-Jivan	Quentin Fournier	Michel R. Dagenais	Abdelwahab Hamou-Lhadj
Brock University	Polytechnique Montreal	Polytechnique Montreal	Concordia University
nezzatijivan@brocku.ca	quentin.fournier@polymtl.ca	michel.dagenais@polymtl.ca	wahab.hamou-lhadj@concordia.ca




DIGITAL ACCESS TO
SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU





HARVARD LIBRARY
Office for Scholarly Communication

Using Performance Variation for Instrumentation Placement in Distributed Systems

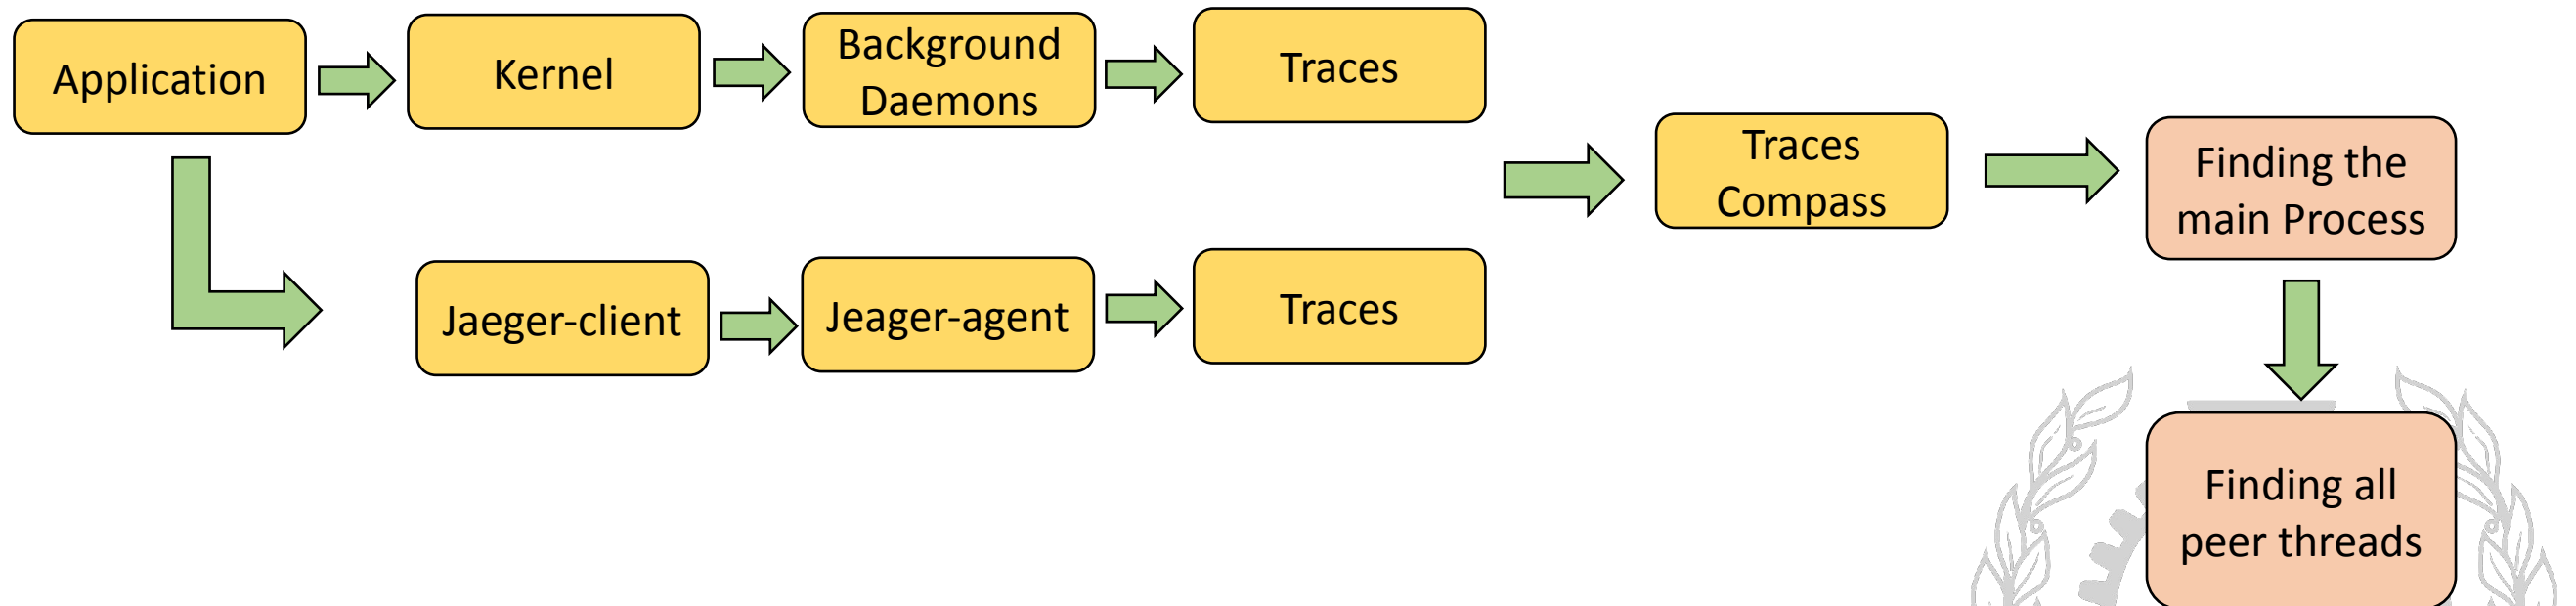


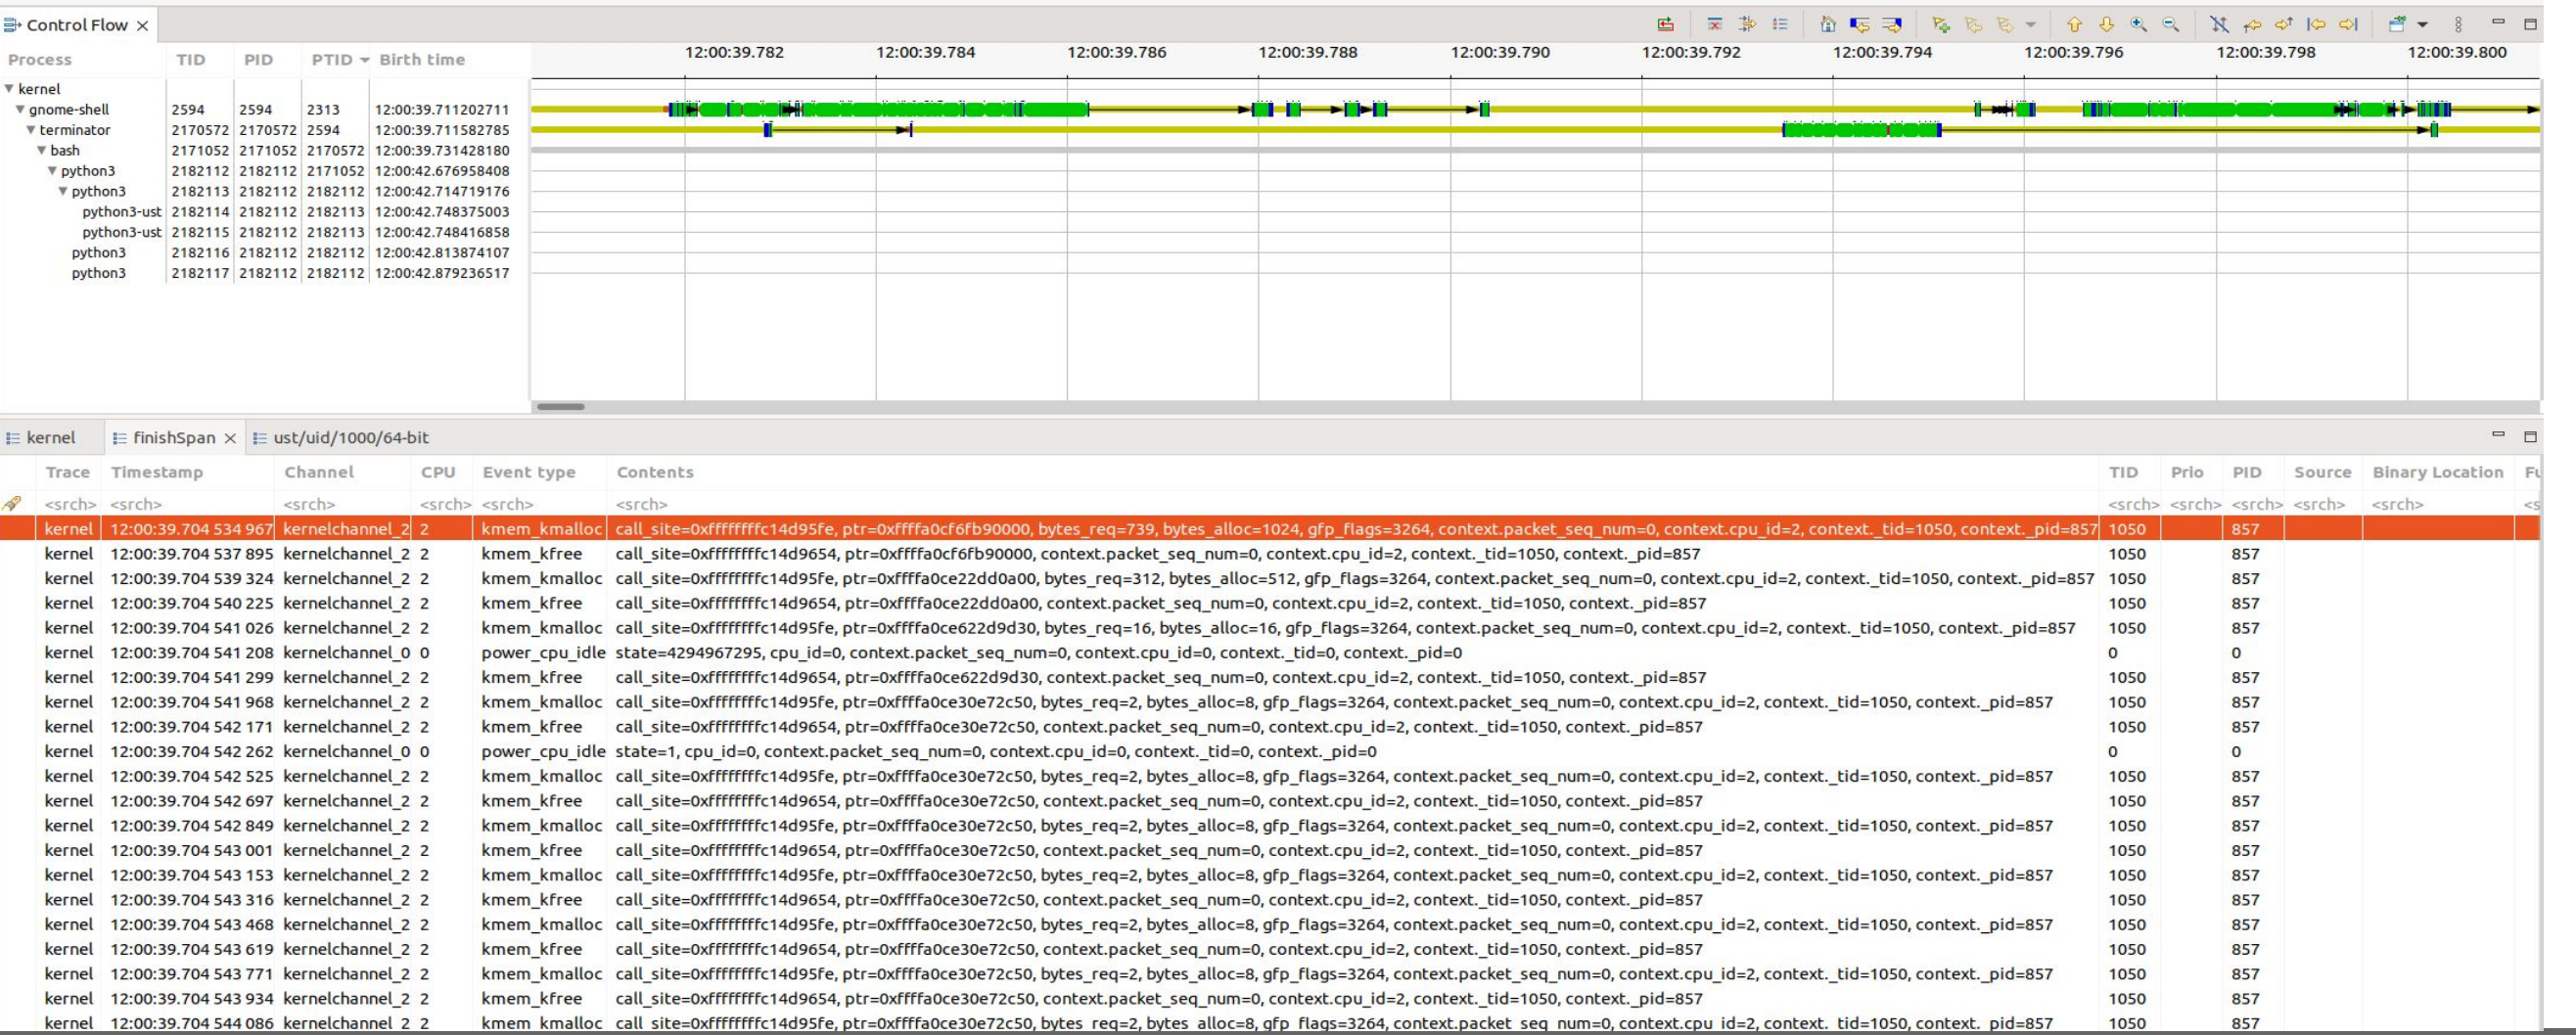
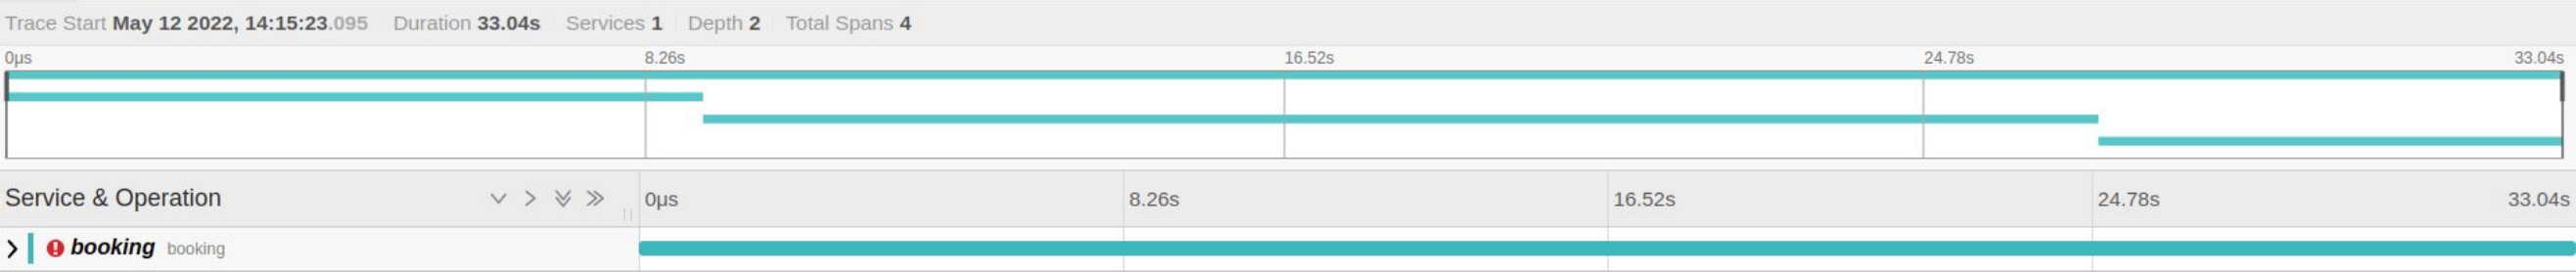
Roadmap



Trace Collection

- Instrumenting Jaeger source code with LTTng.
 - Defining the start and endpoint of each span.
- Converting JSON to DAG.
- Illustrating the communication between the main thread and its peer on the kernel side.





Strategy

- Local problem
 - Bad code
 - ✓ Infinite loop.
- System problem
 - System misconfiguration.
 - Workload on the CPU/ Memory/ Network.
 - ✓ Using benchmark tools for creating HW load.

```
WARNING: --num-threads is deprecated, use --threads instead
sysbench 1.0.18 (using system LuaJIT 2.1.0-beta3)
```

```
Running the test with following options:
Number of threads: 16
Initializing random number generator from current time
```

```
Prime numbers limit: 10000
```

```
Initializing worker threads...
```

```
Threads started!
```

```
CPU speed:
  events per second: 45867.53
```

```
General statistics:
  total time:          10.0003s
  total number of events: 458710
```

```
Latency (ms):
  min:                0.19
  avg:                 0.35
  max:                24.35
  95th percentile:    0.35
  sum:                159949.73
```

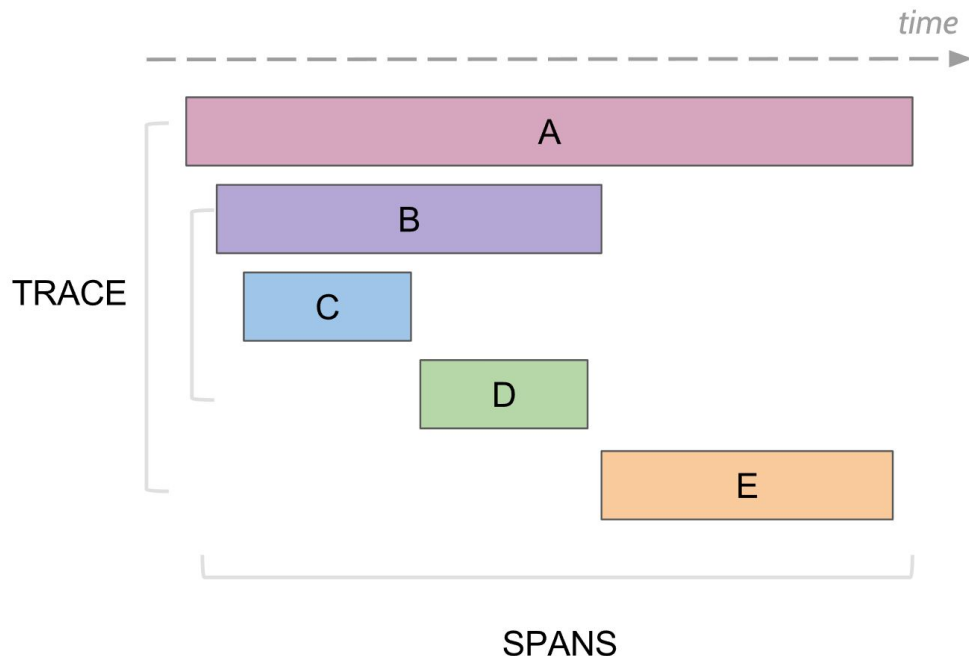
```
Threads fairness:
```

```
1 [|||||] 100.0% 5 [|||||] 100.0% 9 [|||||] 100.0% 13 [|||||] 100.0%
2 [|||||] 100.0% 6 [|||||] 100.0% 10 [|||||] 100.0% 14 [|||||] 100.0%
3 [|||||] 100.0% 7 [|||||] 100.0% 11 [|||||] 100.0% 15 [|||||] 100.0%
4 [|||||] 100.0% 8 [|||||] 100.0% 12 [|||||] 100.0% 16 [|||||] 100.0%
Mem[|||||] 15.4G/30.7G Tasks: 295, 2050 thr; 16 running
Swp[|||||] 0K/2.00G Load average: 3.80 1.27 0.73
Uptime: 4 days, 05:58:48
```

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
460523	maryam	20	0	34456	9460	8044	S	1587	0.0	1:31.73	sysbench cpu run --num-threads=16
460537	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.76	sysbench cpu run --num-threads=16
460532	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.76	sysbench cpu run --num-threads=16
460535	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.76	sysbench cpu run --num-threads=16
460525	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.75	sysbench cpu run --num-threads=16
460533	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.75	sysbench cpu run --num-threads=16
460536	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.75	sysbench cpu run --num-threads=16
460527	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.74	sysbench cpu run --num-threads=16
460534	maryam	20	0	34456	9460	8044	R	100.	0.0	0:05.74	sysbench cpu run --num-threads=16
460526	maryam	20	0	34456	9460	8044	R	99.4	0.0	0:05.74	sysbench cpu run --num-threads=16
460528	maryam	20	0	34456	9460	8044	R	99.4	0.0	0:05.74	sysbench cpu run --num-threads=16
460539	maryam	20	0	34456	9460	8044	R	99.4	0.0	0:05.73	sysbench cpu run --num-threads=16
460538	maryam	20	0	34456	9460	8044	R	99.4	0.0	0:05.68	sysbench cpu run --num-threads=16
460531	maryam	20	0	34456	9460	8044	R	98.7	0.0	0:05.69	sysbench cpu run --num-threads=16
460530	maryam	20	0	34456	9460	8044	R	98.1	0.0	0:05.70	sysbench cpu run --num-threads=16
460524	maryam	20	0	34456	9460	8044	R	96.8	0.0	0:05.71	sysbench cpu run --num-threads=16
460529	maryam	20	0	34456	9460	8044	R	96.2	0.0	0:05.65	sysbench cpu run --num-threads=16
17277	maryam	20	0	6835M	376M	123M	S	7.1	1.2	23:35.48	/usr/bin/gnome-shell
387848	maryam	20	0	13804	7204	3644	R	1.9	0.0	32:01.25	htop
17088	maryam	20	0	1945M	220M	142M	S	1.3	0.7	25:52.19	/usr/lib/xorg/Xorg vt2 -displayfd 3 -auth /run/user/1000/gdm/Xauthority


Strategy

- Span or swim lane showing call graph relationship
- DAG also shows call graph and the task's relationship
 - Concurrent behavior
 - Synchronization points
 - Asynchronous Behavior



Reference:

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37365088>



Strategy



- Extracting critical path

- Concurrency

Finding the path between the start and endpoint of the execution while we are not waiting for something else.


- No concurrency



The workflow of the application.



Strategy

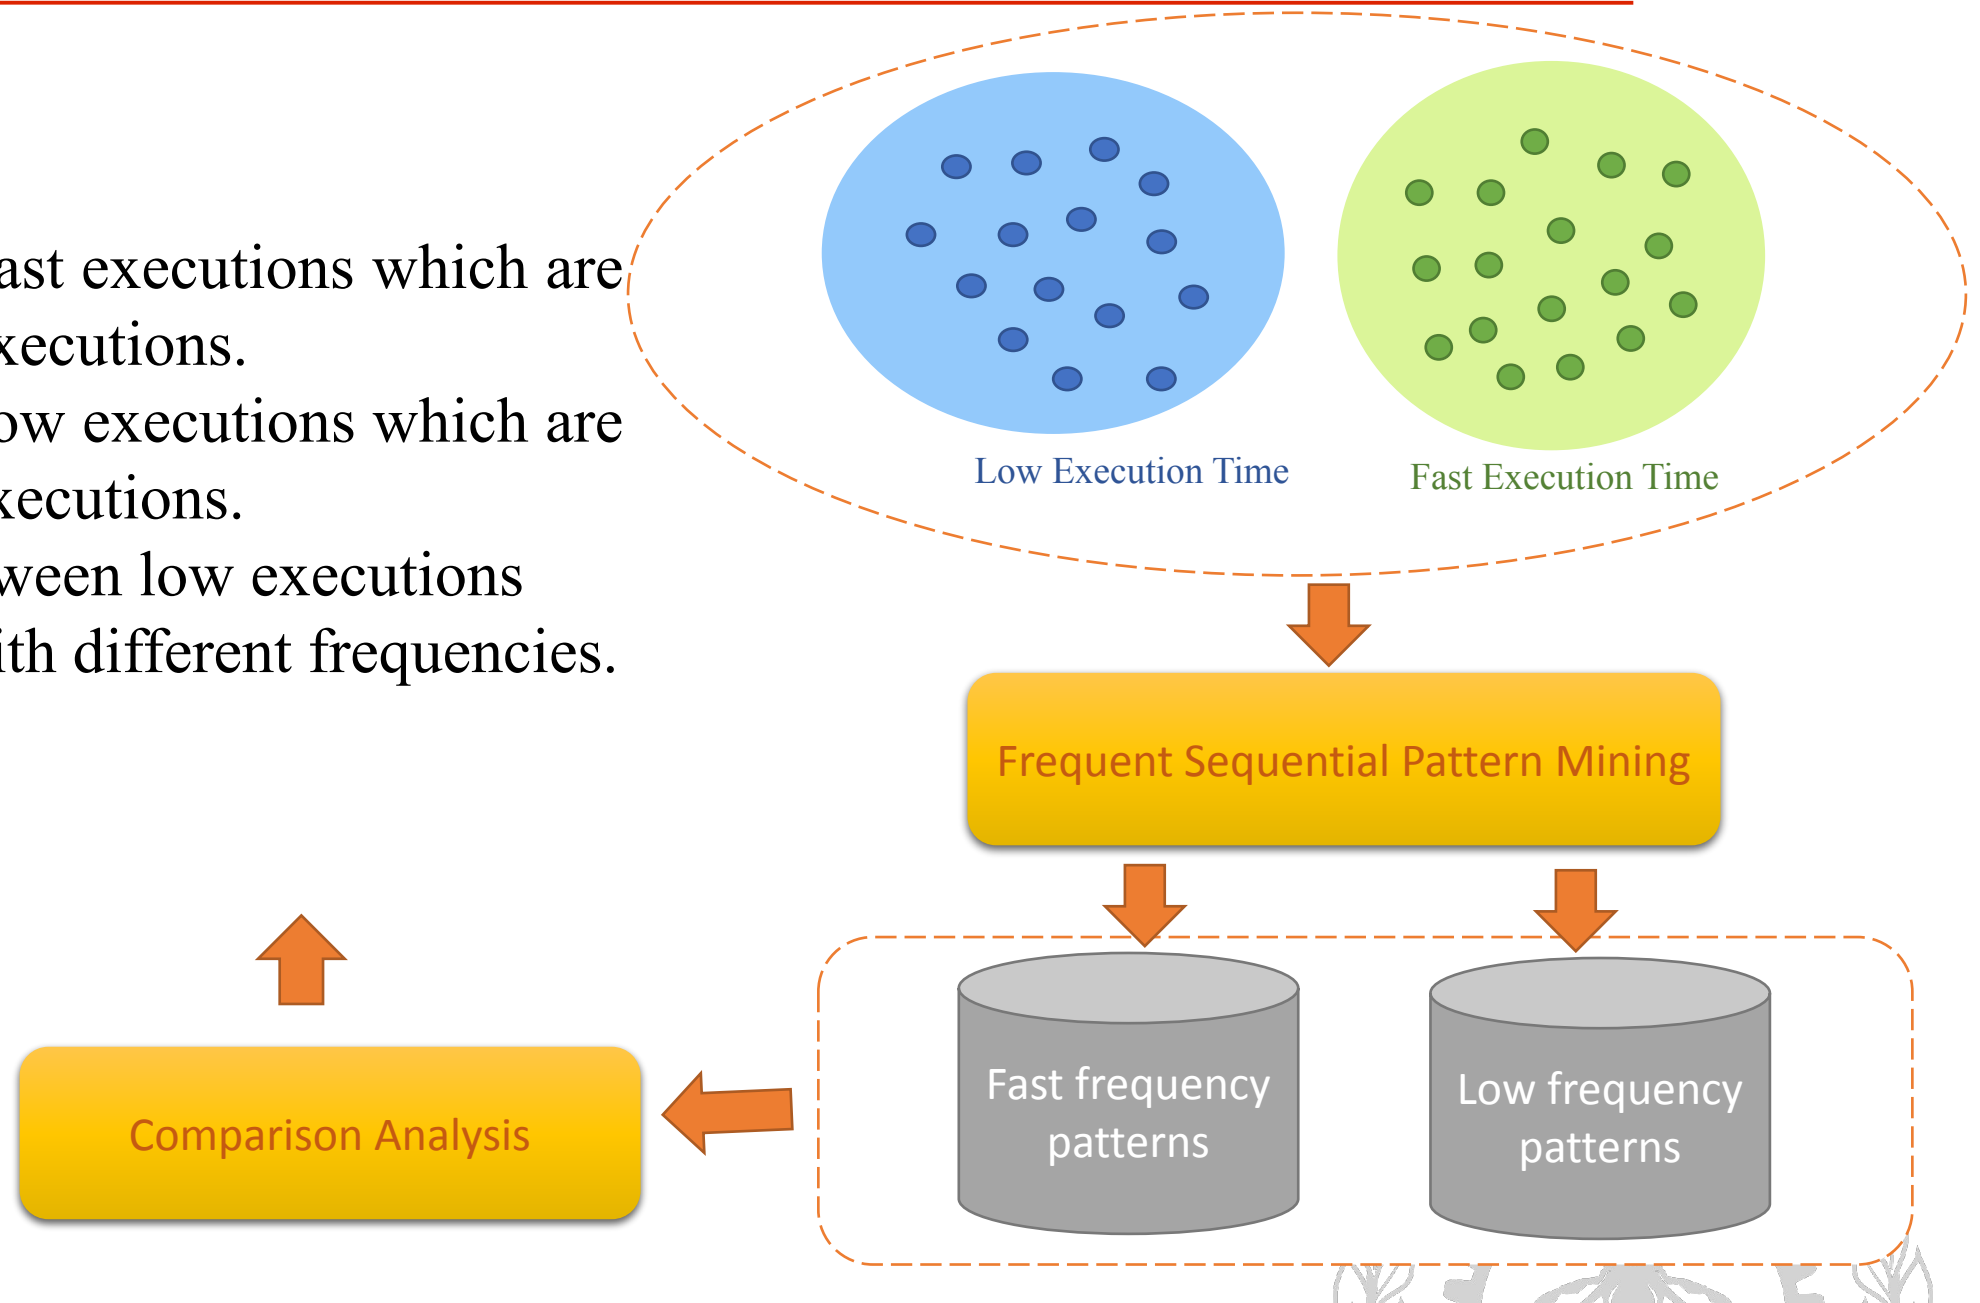
- Categorization

- **Thresholding**
 - Number of page fault
 - Memory usage
 - Execution time
 - ...
- **Clustering**  Combination of performance metrics



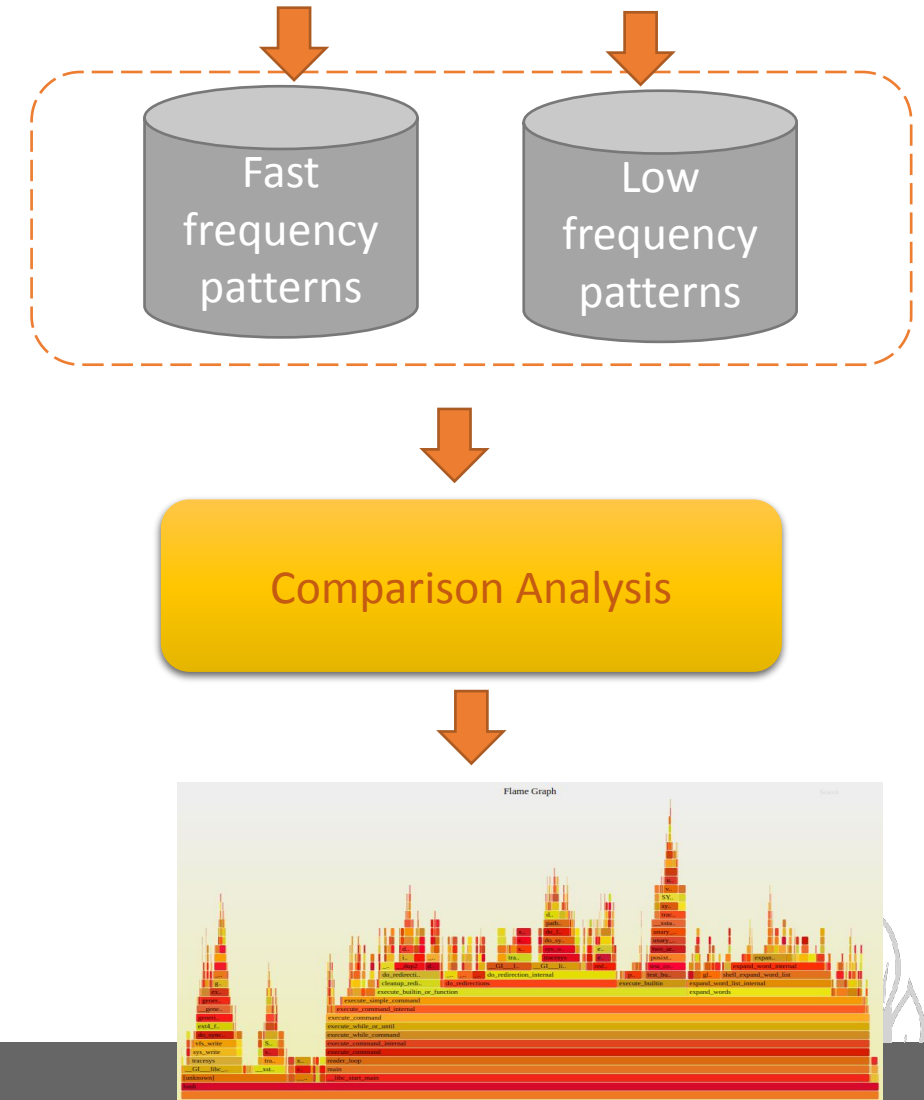
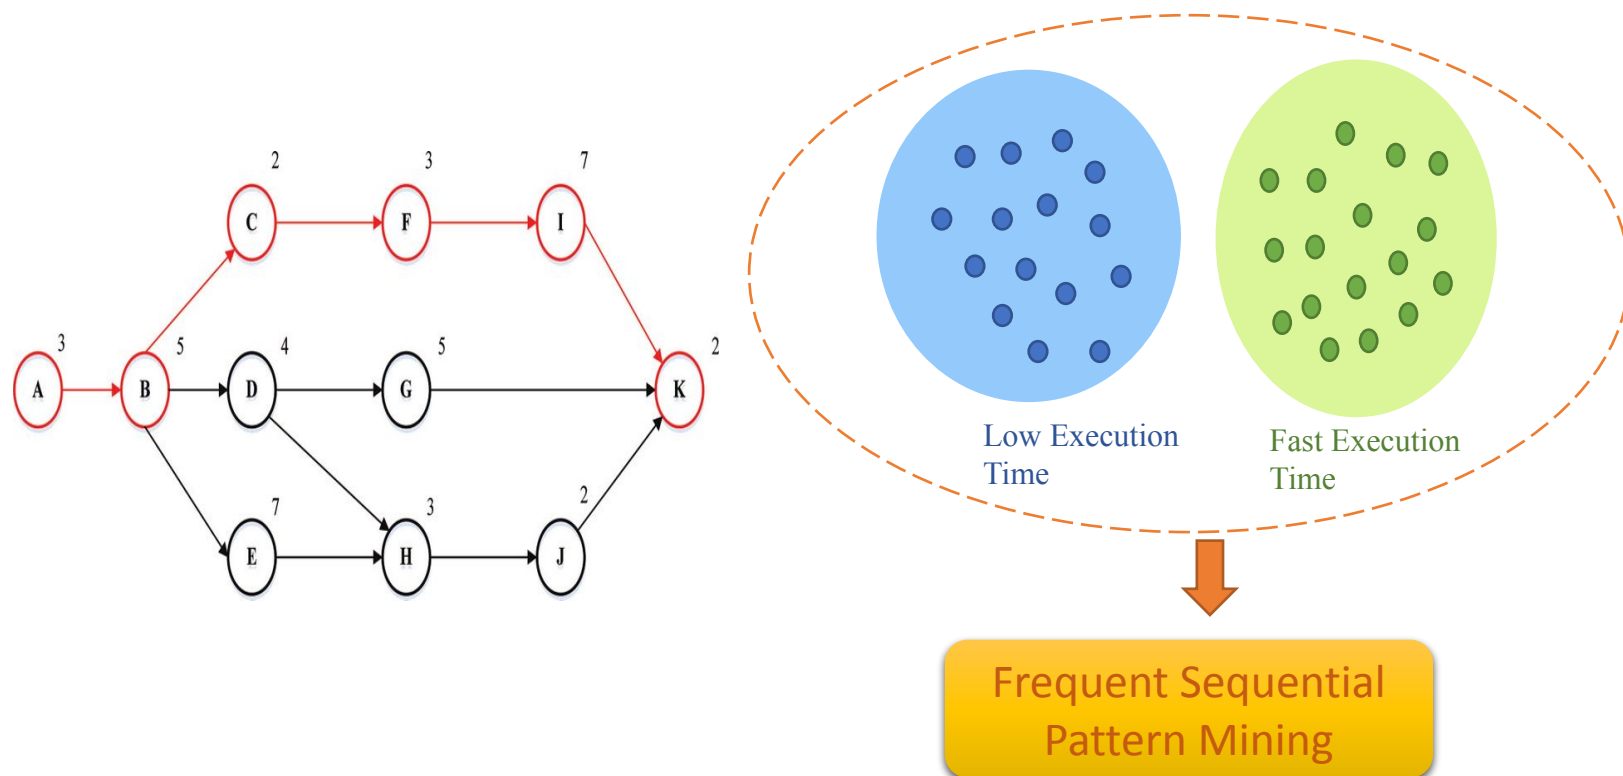
Strategy

- Comparison:
 - Frequent patterns of fast executions which are not available in low executions.
 - Frequent patterns of low executions which are not available in fast executions.
 - Common patterns between low executions and fast executions with different frequencies.



Future Goal

- Extracting critical path based on the DAG.
- Categorizing executions.
- Localizing performance problems.



Thank you

Email: maryam.ekhlasi@polymtl.ca

Source Code: <https://github.com/maryamekhlasi/jaeger-client-python.git>

