



# Comparing distributed execution traces for understanding intermittent latency sources

*Maryam Ekhlas*  
June 1<sup>th</sup>, 2023

Polytechnique Montreal

**DORSAL** Laboratory

## Motivation

---

1. Is it a short-term issue or a recurring one when performance is unsatisfactory?
2. Is this problem due to a change in the code, network problems, or infrastructure and hardware problems?
3. Is buying new hardware or reconfiguring old hardware the most appropriate choice?
4. Which new hardware would provide the most performance improvement for a given budget?

### Microservices performance changes:


1. Source code change.
2. Changes in the type or frequency of requests.
3. Infrastructure change.
4. Resource contention, limitation.
5. Hardware/software configuration change.
6. Combination of the above changes.



## Our Goal


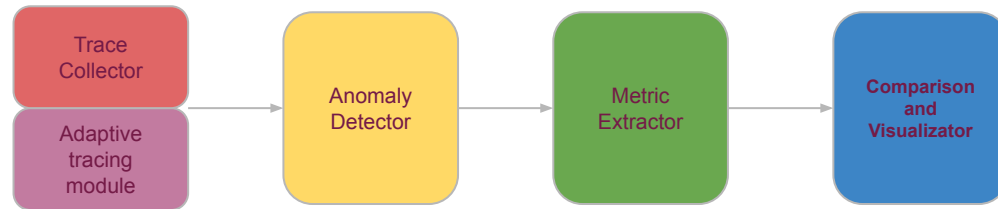
---

1. Finding and locating the performance problem through a distributed system.
2. Comparing two sets of executions to evaluate the differences in terms of performance.
3. Providing sets of views to highlight the differences and speed up problem diagnosis.
4. Identifying the root causes of performance degradation.
5. Minimizing the trace cost.




# Architecture

---





# Metric Extractor



## → Span Delimiters

- ◆ Start
- ◆ End


## → Thread state Delimiters

- ◆ Sched\_switch (new thread starts running on a CPU).
- ◆ Sched\_wakeup (blocked thread becomes ready to run).
- ◆ Softirq\_entry/exit
  - Soft IRQ1 (Timer).
  - Soft IRQ 2,3 (Network).
  - Soft IRQ 4,5 (Block Device).
  - High-Resolution timer (Timer).
- ◆ IRQ 19,23 (USB Device).
- ◆ None (Other Thread).

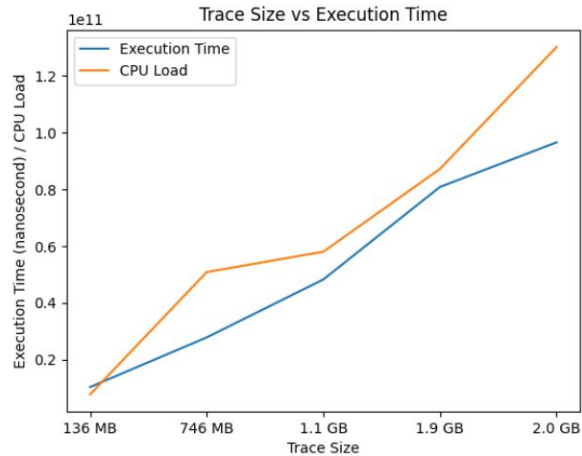





# Evaluation



- Evaluating the precision of our evaluation for each span.
- Evaluating the analysis time based on trace size.

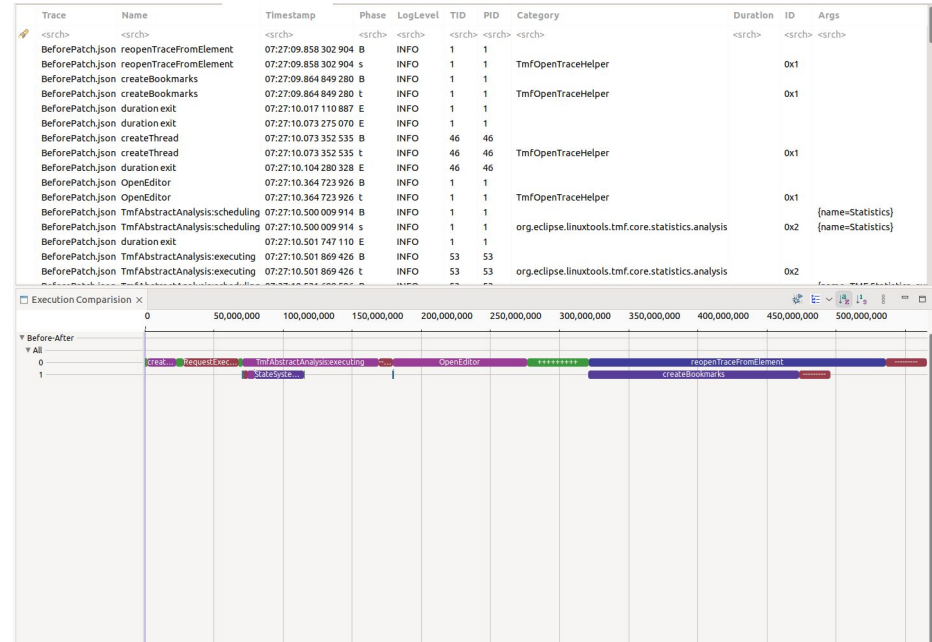
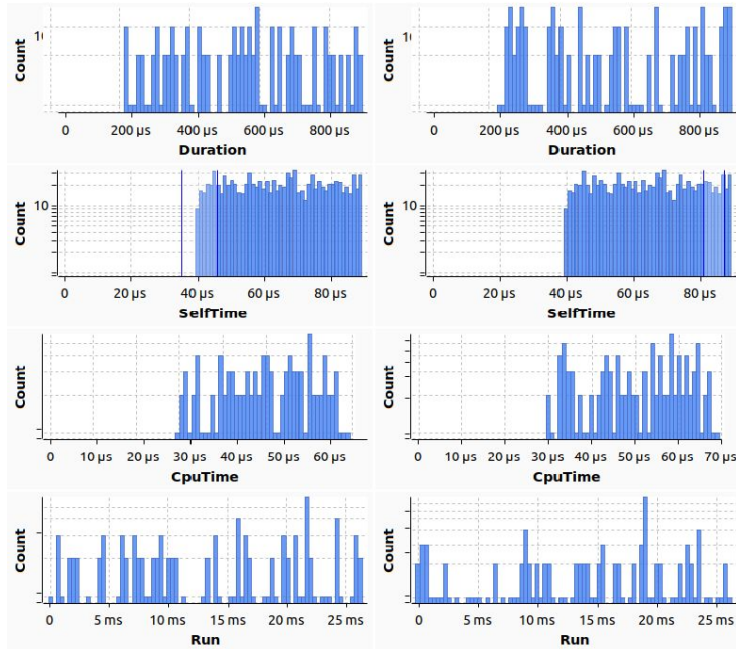



Evaluation Results			
Span Name	Span Id	Jaeger (us)	Our Result (ns)
HTTP GET /customer	4dce9ba044b45859	31,3897	314,206,223
DriverService/FindNearest	118142188d13609c	216,620	216,871,328
HTTP GET /route	7fbc78872f3847df	41,743	42,017,778
HTTP GET /route	068be711cbee624e	69,695	69,698,931
HTTP GET /route	69c577822cda2908	49,599	49,572,453
HTTP GET /route	3ef14dfe688096be	42,478	42,496,332
HTTP GET /route	2c1cdfb1683d2262	31,479	31,481,573
HTTP GET /route	64bceb453d0c6ba	49,609	49,696,770
HTTP GET /route	4d4daf70092f9c01	56,357	56,362,504
HTTP GET /route	5f246b46290f6047	61,724	61,730,655
HTTP GET /route	525a1dc1a5a0a1d5	61,106	61,136,932
HTTP GET /route	602e0b2dad7af025	69,575	69,626,575
HTTP GET: /customer	1c0fd57703da8e0a	315,612	315,548,823
HTTP GET: /route	6d83aaa2ec25eaac	43,402	43,378,498
HTTP GET	6d84d72745b448b3	43,296	43,323,463
HTTP GET	77c0315b649c18e9	315,498	315,536,842
DriverService/FindNearest	43d04e7f1c875ca7	219,060	219,018,811
HTTP GET	1276195a9087245f	50,632	50,653,248
HTTP GET	303f97953df6617a	70,916	70,935,682
HTTP GET: /route	66779dc0155fe817	50,729	70,935,682
HTTP GET: /route	0e2ec848b9b49eb9	71,013	70,979,892







# Comparison and Visualization







# How can we help industrial Developers?




Duration	SpanId	Operation Name	TraceID	ThreadID	UNKNOWN	WAIT BLOCKED	RUN	RUNSYSTEM CALL	INTERRUPTED	WAIT CPU	WAIT UNKNOWN	WAIT FORK
227,507,071	2844007279735121 107	SQL SELECT	818662536074181 6373	964130	0	227630732	11669	7962	0	0	0	0
2,838,960,28	4783316764536372 719	SQL SELECT	368599671838341 9836	964134	0	2810873663	106570 85	3583485	628413	1056556 3	2754456	0



# How can we help industrial Developers?



# How can we help industrial Developers?



- A database query that is executed in the main thread or goroutine causes the thread to block and wait for the operation to complete.
- It prevents other requests from being processed until the computation or query finishes.
- The main thread is blocked until the goroutine signals that the work is completed.




Start Time	End Time	Duration	SpanId	OperationName	TraceID	ThreadID	UNKNOWN	WAITBLOCKED	RUN	RUNSYSTEMCALL	INTERRUPTED	WAITCPU	WAITUNKNOWN	WAITFORK
13:40:48.723 909 384	13:40:48.725 449 717	1,540,333	5505750935733821175	HTTP GET: /route	6431679566420986618	216497	0	12363495358	250351 787405	6842	109723	0	0	0
13:40:48.724 442 004	13:40:48.726 096 632	1,654,628	3501446615930315167	HTTP GET: /route	6431679566420986618	216470	0	12363141288	594615 215588	81751	0	14727	0	0
13:40:48.725 146 818	13:40:48.725 997 506	850,688	4087164681460535921	HTTP GET: /route	6431679566420986618	216470	0	12363141288	24295 0	172911 44359	0	0	0	0
13:40:48.725 702 242	13:40:48.726 275 689	573,447	5203168035594131235	HTTP GET: /route	6431679566420986618	216350	0	0	428104 133500	14707	0	0	0	0
13:40:48.726 186 010	13:40:48.726 960 645	774,635	3254686707739110086	HTTP GET: /route	6431679566420986618	216350	0	313045	328829 103943	48028	25384	0	0	0
13:40:48.726 213 732	13:40:48.726 949 865	736,133	3680470341830845931	HTTP GET: /route	6431679566420986618	216499	0	285782	321250 160833	16119	47212	0	0	0
13:41:21.422 380 869	13:41:21.423 696 821	1,315,952	4216478611002514525	HTTP GET: /route	2214312730109965973	217002	0	522091	653124 123904	0	52287	0	0	0
13:41:21.422 411 867	13:41:21.423 318 500	906,633	221128445473134325	HTTP GET: /route	2214312730109965973	217006	0	345405	484217 77955	0	20494	0	0	0
13:41:21.423 020 561	13:41:21.423 483 520	462,959	3752248539699955781	HTTP GET: /route	2214312730109965973	217008	0	112778	162102 146944	27160	16447	0	0	0
13:41:21.423 368 614	13:41:21.423 743 869	375,255	3985339775676182957	HTTP GET: /route	2214312730109965973	217001	0	119501	212316 64709	14655	12311	0	0	0
13:41:21.423 644 192	13:41:21.423 842 996	198,804	5683614158188625367	HTTP GET: /route	2214312730109965973	217005	0	105016	58833 27172	4688	5007	0	0	0
13:41:21.423 749 630	13:41:21.423 986 275	236,645	4101109696620141814	HTTP GET: /route	2214312730109965973	217002	0	0	212238 53167	11921	0	0	0	0
13:41:21.423 802 199	13:41:21.423 986 325	184,126	5768467736216194752	HTTP GET: /route	2214312730109965973	217001	0	159407	59264 20485	4828	5007	0	0	0
13:41:21.423 875 867	13:41:21.424 060 745	184,878	5911336802399416461	HTTP GET: /route	2214312730109965973	217002	0	0	190789 45244	14096	0	0	0	0
13:41:21.424 012 835	13:41:21.424 182 393	169,558	5969567871979466795	HTTP GET: /route	2214312730109965973	217072	0	65562	107707 23856	6081	6051	0	0	0
13:41:21.424 079 560	13:41:21.424 339 137	259,577	5024941312143855339	HTTP GET: /route	2214312730109965973	217002	0	117487	79609 53388	10478	18842	0	0	0
13:40:43.717 127 392	13:40:48.718 272 086	5,001,144	3133712204282027235	SQL SELECT	6431679566420986618	216470	0	4999907384	601355 441102	0	313548	0	0	0
13:41:21.420 958 427	13:41:21.420 969 658	11,231	5881379417971386402	SQL SELECT	2214312730109965973	217002	0	0	132677 1933	0	0	0	0	0



## Future Work

---

- ◆ How do microservice faults or operations performance problems manifest in a distributed system?
- ◆ How can we select groups of executions for our metric extraction and comparison module?
- ◆ When can we start collecting traces both in the kernel and user space?
- ◆ What tracepoints should be active, and which metrics should be collected?



---

# Thank you

**Email:** [maryam.ekhlasi@polymtl.ca](mailto:maryam.ekhlasi@polymtl.ca)

