

## **Trace Coordinator**

Ahmad Faour

Polytechnique Montréal DORSAL Laboratory

## Agenda

- 1. Proposed Solution
- 2. Results
- 3. Discussion

#### Reminder - Trace Coordinator



#### Introduction - Use cases

#### • Target use cases

- High Performance Computing with MPI Cluster
- Microservice with Kubernetes Cluster

#### • Other interesting use cases

- LTTng Log Rotation
- Similarity of queries
- Client-Server

#### Results - Time Graph API

- Process Tree (Id)
- States
- Arrows



#### Results - XY API

- Process Tree (Id)
- Series of points



#### Results - Parameters

- 2 sets of 20 traces for a total of 4GB and 40 GB
- Different screen resolutions (1080p, 2K et 4K)
- Execute on many sets of Trace Compass Server (1, 5, 10 and 20)

## Results - Trace and Experiment API

Trace File Size	Number of Workers	Open Trace (s)	Create Experiment (s)
4GB	$     \begin{array}{c}       1 \\       5 \\       10 \\       20     \end{array} $	$\begin{array}{c c} 2.83 \\ 1.17 \\ 0.95 \\ 0.89 \end{array}$	$277.15 \\ 67.81 \\ 35.39 \\ 17.60$
40GB	151020	$  \begin{array}{c} 1.78 \\ 1.40 \\ 1.06 \\ 1.03 \end{array}  $	$2850.97 \\ 845.37 \\ 444.04 \\ 152.89$

## Results - TimeGraph API

		Ge	Get TimeGraph		
Screen Resolution	Trace Number	er	States	Arrows	
	File Size of Worke	rs (s)	(s)	(s)	
1080p	$\begin{array}{c}1\\4\text{GB}&10\\20\end{array}$	$\begin{array}{r} 362.62 \\ 43.56 \\ 31.05 \\ 24.65 \end{array}$	$4.46 \\ 2.42 \\ 1.97 \\ 1.77$	$7.30 \\ 1.76 \\ 0.91 \\ 0.53$	
	Size of the Respon	se 3.1MB	21.3MB	523.2KB	
	40GB $\begin{array}{c} 1\\ 5\\ 10\\ 20\end{array}$	$\begin{array}{r} 3430.79 \\ 447.29 \\ 256.73 \\ 219.49 \end{array}$	$28.40 \\ 18.77 \\ 2.51 \\ 1.90$	$36.16 \\ 8.78 \\ 8.12 \\ 5.46$	
	Size of the Respon	se 139.5MB	24.0MB	32.8KB	
2K	$\begin{array}{c}1\\4\mathrm{GB}&10\\20\end{array}$	$315.31 \\ 43.72 \\ 30.98 \\ 23.31$	$4.50 \\ 3.14 \\ 2.79 \\ 2.54$	5.47 1.30 0.93 0.58	
	Size of the Respon	se 3.1MB	35.4MB	989.3KB	
	$\begin{array}{c} 1 \\ 40 \text{GB} & \begin{array}{c} 1 \\ 5 \\ 10 \\ 20 \end{array}$	$\begin{array}{c} 3532.46 \\ 445.16 \\ 261.71 \\ 220.66 \end{array}$	51.66 20.14 9.46 5.79	$86.04 \\ 11.98 \\ 8.68 \\ 2.47$	
	Size of the Respon	se 139.5MB	39.7MB	250.6 KB	
4K	$\begin{array}{c} 1\\ 4\text{GB} \\ 20 \end{array}$	$\begin{array}{r} 380.32 \\ 43.16 \\ 31.22 \\ 24.75 \end{array}$	$7.28 \\ 5.54 \\ 4.39 \\ 4.36$	$7.66 \\ 1.80 \\ 1.22 \\ 0.98$	
	Size of the Respon	se 3.1MB	63.1MB	2.1MB	
	$\begin{array}{c} 1 \\ 40 \text{GB} \\ 20 \end{array}$	$\begin{array}{c c} 2928.45\\ 448.91\\ 269.43\\ 214.35 \end{array}$	$44.55 \\ 21.77 \\ 10.85 \\ 6.91$		
	Size of the Respon	se 139.5MB	71.2MB	1.0MB	

## Results - XY API

	Í	Get XY	
Screen	Trace result Number	Tree	Models
Resolution	File Size of Workers	(s)	(s)
1080p	15	145.57 36.99	10.90 3 47
	4GB 10 20	33.71 30.24	2.12
	Size of the Response	1.5MB	7.1MB
	15	3282.95	704.70
	40GB 10 20	339.27 309.84	66.00 37.05
	Size of the Response	130.8MB	7.2MB
	1	141.95	13.75
	4GB 10	35.56 32.18	$\frac{4.31}{2.68}$
	Size of the Response	1.5MB	11.8MB
2K	1	4992.27	893.76
	40GB 10 20	340.31 303.97	93.08 50.41
	Size of the Response	130.8MB	12.0MB
4K	1	142.48	21.30
	4GB 10 20	32.40	3.76
	Size of the Response	1.5MB	21.1MB
	1	3304.87	1435.14 267.29
	40GB 10 20	349.94 299.12	$137.52 \\ 77.09$
	Size of the Response	130.8MB	21.6MB

## Critical Path Analysis



Trace Coordinator – Ahmad Faour

## Analysis on dependencies between traces









Noticed:

- The index is a universal identifier (UUID)
- The UUID is the key to matching different vertices across traces

## Analysis on dependencies between traces\*



\*Unfortunately, we were not able to implement the analysis.

#### Discussion - Return on the result

- Greater performance gains as trace size increases
- Depending on the query, the computational load may be concentrated on certain analysis nodes
- Serialization, a bottleneck at the coordinator level

#### Conclusion\_

- Improving trace analysis and visualization tools in the context of high performance computing systems
- Abstraction layer to expose visualization models for each analysis
- Queries based on desired rows and resolution to avoid overloading



# Thank you for listening!