

Making Sense of Multi-threaded Application Performance at Scale with NonSequitur

Augustine Wong,Paul Bucci, **Ivan Beschastnikh** Alexandra Fedorova





THE UNIVERSITY OF BRITISH COLUMBIA



The Impact of Performance Bugs

Compromised but operational system Degradation in:



R. Krishna, M. S. Iqbal, M. A. Javidian, B. Ray, and P. Jamshidi. Cadet: Debugging and fixing misconfigurations using counterfactual reasoning.arXiv preprint arXiv:2010.06061, 2020.



The Cost of Fixing Performance Bugs

More difficult to fix than other types of bugs

[W]e found that performance bugs take <u>more time</u> to fix, are <u>fixed by more experienced developers</u> and <u>require more changes</u> to more lines than nonperformance bugs [1].

[1] S. Zaman, B. Adams, and A. E. Hassan. A qualitative study on performance bugs. In 2012 9th IEEE working conference on mining software repositories (MSR), pages 199–208. IEEE, 2012.



A Case Study of Performance Comprehension

We performed an empirical case study to examine the nature of *performance comprehension* [2]



[2] Alexandra Fedorova, Craig Mustard, Ivan Beschastnikh, Julia Rubin, Augustine Wong, Svetozar Miucin, and Louis Ye. 2018. Performance comprehension at WiredTiger. In Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. 83–94.



Analyzing Performance Comprehension

W/ Jira

We analyzed 44 JIRA tickets that tracked performance issues at WiredTiger



Key Findings on Performance Comprehension

WiredTiger developers spent the majority of their performance debugging time investigating latency spikes



Developers were interested in knowing what threads did over time



Developers collected and analyzed large amounts of information



Performance Analysis Tool Requirements

Requirements	Description
R_Latency	Find when latency spikes occur
R_Time	Display thread behavior over time
R_Large	Process large execution traces



Current State of Performance Debugging



Sue Doe is a software developer tasked with investigating latency spikes captured in a large execution trace.



RocksDB trace

Thread	Size			
ID	# Function Calls	# Different Functions		
1	2,466,295	74		
2	2,464,045	75		
3	194,042	247		

What can she do?



Option 1: Use FlameGraphs?



Merges callstack samples together and sorts them alphabetically

Does not show time





Option 1: Use FlameGraphs?







Option 2: Chrome TraceViewer?



Nothing calected. Tan stuff



Option 2: Chrome TraceViewer?





Option 2: Chrome TraceViewer?



Tool Requirements						
R_Latency R_Time R_Large						
\checkmark						

Tries to visualize everything in an execution trace!



Other Tools?

Toolo	Tool Requirements				
10015	R_Latency R_Time		R_Large		
Zinsight [3]	×	\checkmark	\checkmark		
TraceViz [4]	X		\checkmark		
SyncTrace [5]	\checkmark	\checkmark	×		
ExtraVis [6]	×	\checkmark	×		
TraceDiff [7]	×	\checkmark	×		
Compress [8]	×	×	\checkmark		
Phase Finder [9]	X	X	\checkmark		
Sabalan [10]	X	X	\checkmark		

[3] De Pauw et al. (2010)

[4] Dautriche et al. (2016)

[5] Karran et al. (2013)

[6] Cornelissen et al. (2008)

[7] Trumper et al. (2013)

- [8] Hamou-Lhadj et al. (2002)
- [9] Pirzadeh et al. (2010)
- al. (2008) [10] Alimadadi et al. (2018)



Our Solution: Compress + Visualize











1) Merge same callstacks together





1) Merge same callstacks together







2) Arrange groups of merged callstacks

















A box encodes start and end time of the function sequence













Loose causal relationships



Let's Visualize This Trace with NonSequitur



RocksDB

Thread	Size			
ID	# Function Calls	# Different Functions		
1	2,466,295	74		
2	2,464,045	75		
3	194,042	247		





One RegTime Expression Per Thread

















Leave Long Events Alone





Leave Long Events Alone





Use Multiple RegTime Expressions







CallDurationThresh = 1% of a thread's execution time CallGapThresh = 0.1% of a thread's execution time TotalTimeFractionThresh = 3.2% of a thread's execution time





CallDurationThresh = 1% of a thread's execution time CallGapThresh = 0.1% of a thread's execution time TotalTimeFractionThresh = 3.2% of a thread's execution time



NonSequitur Design Drawbacks



X-axis not linear

Each callstack has a minimum horizontal pixel width



Navigating NonSequitur: Highlight Functions

Thread ID	Enter	a function name:					
1 ▲							
3	#	Function	Color				
	38	rocksdb::WriteThread::CompleteParallelMemTableWriter(rocksdb::WriteThread::Write*)					•
	39	rocksdb::WriteThread::ExitAsBatchGroupLeader(rocksdb::WriteThread::WriteGroup&, and					
*	40	rocksdb::WriteThread::JoinBatchGroup(rocksdb::WriteThread::Writer*)					
	41	rocksdb::WriteThread::LinkOne(rocksdb::WriteThread::Writr					
	42	rocksdb::log::Writer::AddRecord(rocksdb::Slice					•
Thread 1							
			i i i i i	- 10-10-10-10-10-10-10-10-10-10-10-10-10-1	* *		
Thread 2							
			15 IC IC	-100	(72 V	-3 - 0 - 0 - 0	
Thread 3							



Navigating NonSequitur: Linked Highlighting



Linked highlighting to help users find correlations in activities *across* threads.









NonSequitur Eval: The Research Questions

- **RQ1:** Does NonSequitur help analyze large execution traces to understand what threads are doing over time?
- **RQ2:** Does NonSequitur help analyze large execution traces to find when unusually long function latencies occur?
- **RQ3:** Does NonSequitur help analyze large execution traces to learn what one thread was doing during the time when another thread was blocked or delayed?
- **RQ4:** For small traces that can be visualized with other tools, does NonSequitur lead to false conclusions in questions regarding thread activity over time as compared to other tools that do not drop information from the trace?



NonSequitur Eval: The Research Questions

- **RQ1:** Does it help explain what threads are doing over time?
- **RQ2:** Does it help find unusually long function latencies?
- **RQ3:** Does it help with correlating thread activities?
- **RQ4:** Does it mislead on small traces?



NonSequitur Eval: The User Study





NonSequitur Eval: Trace Sizes

Normal WiredTiger Trace 0.8M function calls across 24 threads

Large WiredTiger Trace 50M function calls across 28 threads



Large RocksDB Trace 39M function calls across 32 threads



NonSequitur Eval: The Alternative Tools

Chrome TraceViewer for the normal WiredTiger Trace





NonSequitur Eval: The Alternative Tools

OpTrack for the large traces





NonSequitur Eval: The Tasks

Tasks	Description
T1	Understand what threads are doing over time
T2	Find when outlier events occur
Т3	Understand what one thread was doing during the time when another thread was blocked or delayed.



NonSequitur Eval: Difference Scores

Difference Score > 0











NonSequitur Eval: RQ1 Finding

Does it help explain what threads are doing over time?

Large WiredTiger Trace

50M function calls across 28 threads

Large RocksDB Trace

39M function calls across 32 threads

NonSequitur vs OpTrack





NonSequitur Eval: RQ2 Finding

Does it help find unusually long function latencies?





NonSequitur Eval: RQ3 Finding

Does it help with correlating thread activities?





NonSequitur Eval: RQ4 Finding

Does it mislead on small traces?

•• **Normal WiredTiger Trace** 0.8M function calls across 24 threads •• NonSequitur VS ... Tasks Chrome TraceViewer Τ2



T3

T3



Eval: Three case studies

Study 1: Reproduce slow LSM Performance in WiredTiger Study 2: Slow Performance with the mmap Option in WiredTiger Study 3: RocksDB memtable Concurency



Eval: Three case studies

Study 1: Reproduce slow LSM Performance in WiredTiger

Study 2: Slow Performance with the mmap Option in WiredTiger Study 3: RocksDB memtable Concurency





Eval: compression ratio

High compression ratios with RegTime!

System	Thread	# unique	n _{before}	n _{after}	Compression
		events	2	-	Ratio
RocksDB	0	74	2,466,295	323	7,635
RocksDB	1	316	2,762,774	2,485	1,111
WiredTiger	0	115	7,784,936	546	14,258
WiredTiger	1	81	6,935,791	209	33,185
Chrome	0	34	274,361	96	2,857
Chrome	1	8	46,464	56	829



Contributions

Making Sense of Multi-threaded Application Performance with **NonSequitur**



The lossy compression performed by RegTime does not negatively impact a developer's ability to understand what threads are doing over time.



NonSequitur can help developers perform performance analysis tasks on large execution traces more accurately.



https://github.com/auggywonger/nonsequitur_vis/