Trace Compass Scalability Update

Arnaud Fiorini

Polytechnique Montréal
Laboratoire DORSAL

Agenda

1. Background
2. History Tile Backend
3. Experiments
4. Results

Background

- Trace Compass State System Backend:
  - Middle layer storing the analysis results
  - Used in Linux Kernel Analysis, Callstack Analysis and others
- History Tree File:
  - Current implementation in Trace Compass
  - Scales reasonably well
- Partial State System:
  - Stores only part of the intervals
  - Query the trace to dynamically rebuild intervals to answer requests

History Tile Backend

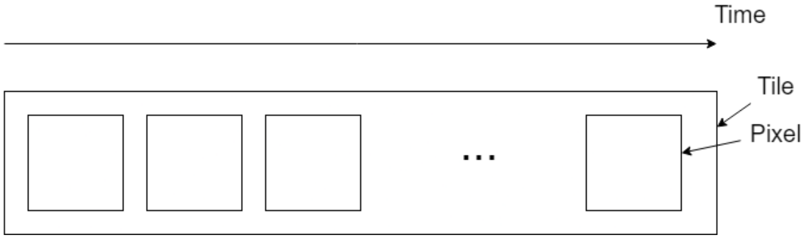

Figure: One *history tile*

History Tile Backend

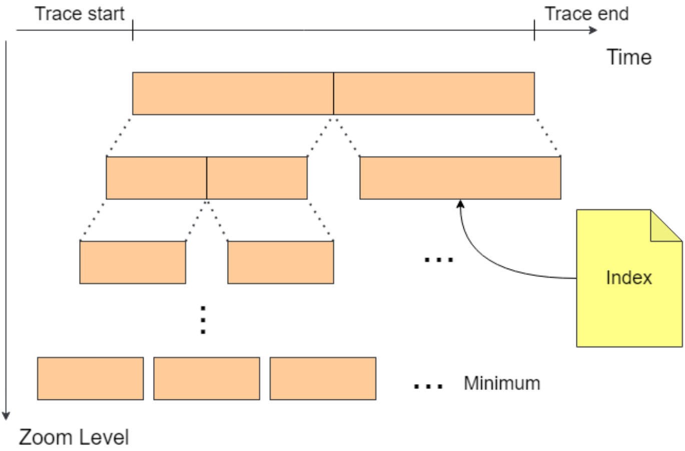

Figure: *History Tile Layers*

History Tile Backend - Answering a query2D

1. Query the starting tile
   - Calculate tile resolution layer
   - If not cached, read the tile on disk
2. Query adjacent tiles
   - Check if the current tile end exceeds the requested end
3. Find out if any intervals were missed:
   - Check the last tile that was queried if all attributes have intervals until the end
   - Return those attributes
4. Run a full query to fetch any missing interval

History Tile Backend - Answering a full query

1. Query the biggest zoom layer tile at time t
2. Continue until all states are found
   - Query the adjacent tile
   - Calculate an *adequate*[1] lower zoom layer
   - Query the tile at that zoom layer

[1] The adequate zoom layer can be calculated according to the minimum size of the missing intervals. Because we checked two complete tiles, the intervals must be at least one tile long. Therefore, it will be contained in any tile where the interval is longer than one pixel.

Experiment

- We will compare three different versions:
  - Current implementation (full)
  - History Tree Tile (tile)
  - Partial State using the History Tree Tile (partial)
- We will run a state system analysis on 8 traces of different sizes (8MB - 28GB) using each back-end
- The result of the analysis will be queried for 4 different zoom levels (showing 95%, 50%, 2.5% and 0.5%)
- The queries were done between two separated nodes in a SSH tunnel.

Results

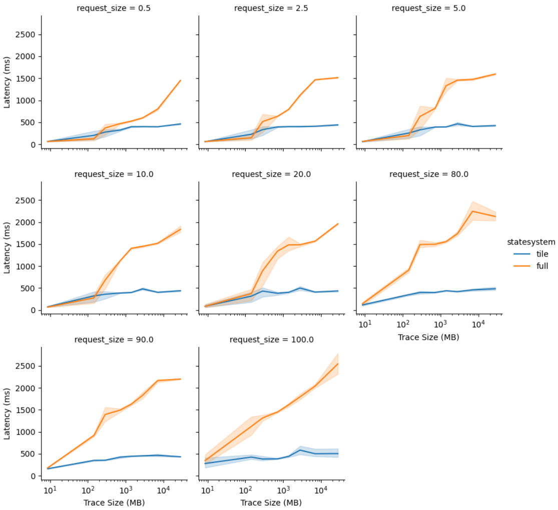
request_size = 0.5
request_size = 2.5
request_size = 5.0
request_size = 10.0
request_size = 20.0
request_size = 80.0
request_size = 90.0
request_size = 100.0
Latency (ms)
0
500
1000
1500
2000
2500
Trace Size (MB)
10^1
10^2
10^3
10^4
statesystem
tile
full

Results

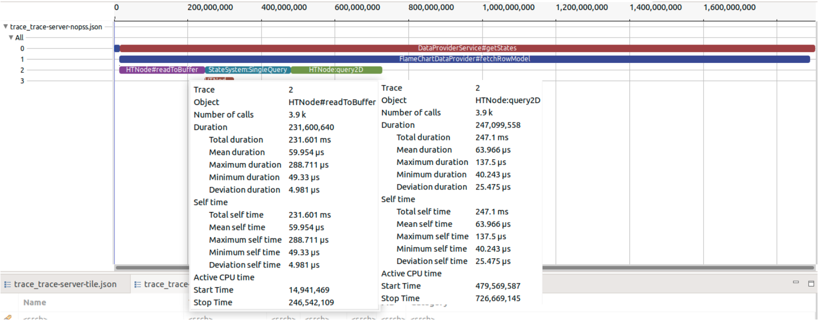

Trace	2
Object	HTNode#readToBuffer
Number of calls	3.9 k
Duration	231,600,640
Total duration	231.601 ms
Mean duration	59.954 μs
Maximum duration	288.711 μs
Minimum duration	49.33 μs
Deviation duration	4.981 μs
Self time	
Total self time	231.601 ms
Mean self time	59.954 μs
Maximum self time	288.711 μs
Minimum self time	49.33 μs
Deviation self time	4.981 μs
Active CPU time	
Start Time	14,941,469
Stop Time	246,542,109

Trace	2
Object	HTNode:query2D
Number of calls	3.9 k
Duration	247,099,558
Total duration	247.1 ms
Mean duration	63.966 μs
Maximum duration	137.5 μs
Minimum duration	40.243 μs
Deviation duration	25.475 μs
Self time	
Total self time	247.1 ms
Mean self time	63.966 μs
Maximum self time	137.5 μs
Minimum self time	40.243 μs
Deviation self time	25.475 μs
Active CPU time	
Start Time	479,569,587
Stop Time	726,669,145

Results

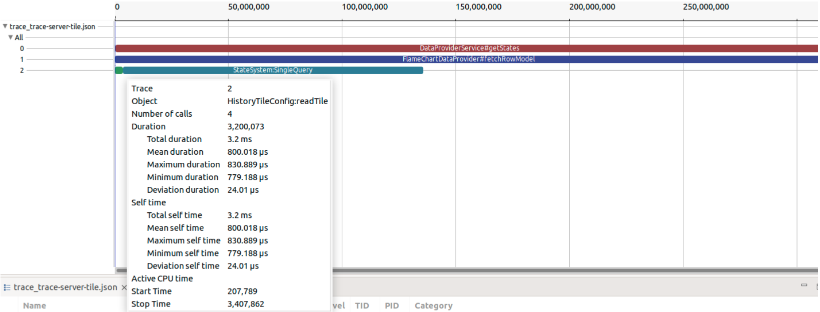
0
50,000,000
100,000,000
150,000,000
200,000,000
250,000,000
trace_trace-server-tile.json
All
0
1
2
DataProviderService#getStates
FlameChartDataProvider#fetchRowModel
StateSystem:SingleQuery
trace_trace-server-tile.json
Name
vel
TID
PID
Category

Trace	2
Object	HistoryTileConfig:readTile
Number of calls	4
Duration	3,200,073
Total duration	3.2 ms
Mean duration	800.018 µs
Maximum duration	830.889 µs
Minimum duration	779.188 µs
Deviation duration	24.01 µs
Self time	
Total self time	3.2 ms
Mean self time	800.018 µs
Maximum self time	830.889 µs
Minimum self time	779.188 µs
Deviation self time	24.01 µs
Active CPU time	
Start Time	207,789
Stop Time	3,407,862

Results

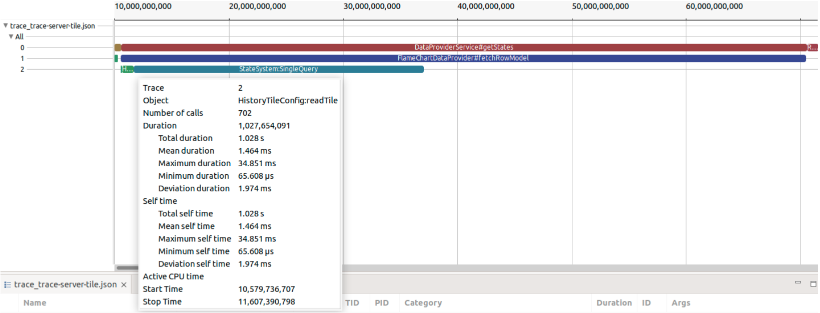

Trace	2
Object	HistoryTileConfig:readTile
Number of calls	702
Duration	1,027,654,091
Total duration	1.028 s
Mean duration	1.464 ms
Maximum duration	34.851 ms
Minimum duration	65.608 µs
Deviation duration	1.974 ms
Self time	
Total self time	1.028 s
Mean self time	1.464 ms
Maximum self time	34.851 ms
Minimum self time	65.608 µs
Deviation self time	1.974 ms
Active CPU time	
Start Time	10,579,736,707
Stop Time	11,607,390,798

Results

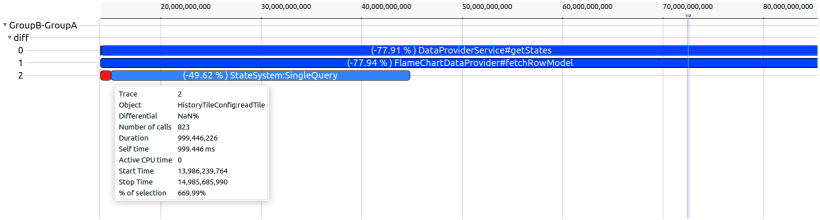
20,000,000,000
30,000,000,000
40,000,000,000
50,000,000,000
60,000,000,000
70,000,000,000
80,000,000,000
GroupB-GroupA
diff
0
1
2
(-77.91 %) DataProviderService#getStates
(-77.94 %) FlameChartDataProvider#fetchRowModel
(-49.62 %) StateSystem:SingleQuery

Trace	2
Object	HistoryTileConfig:readTile
Differential	NaN%
Number of calls	823
Duration	999,446,226
Self time	999.446 ms
Active CPU time	0
Start Time	13,986,239,764
Stop Time	14,985,685,990
% of selection	669.99%