



Adaptive Tracing Problematic Area Localization

Masoumeh Nourollahi

Polytechnique Montréal
DORSAL Laboratory

Agenda

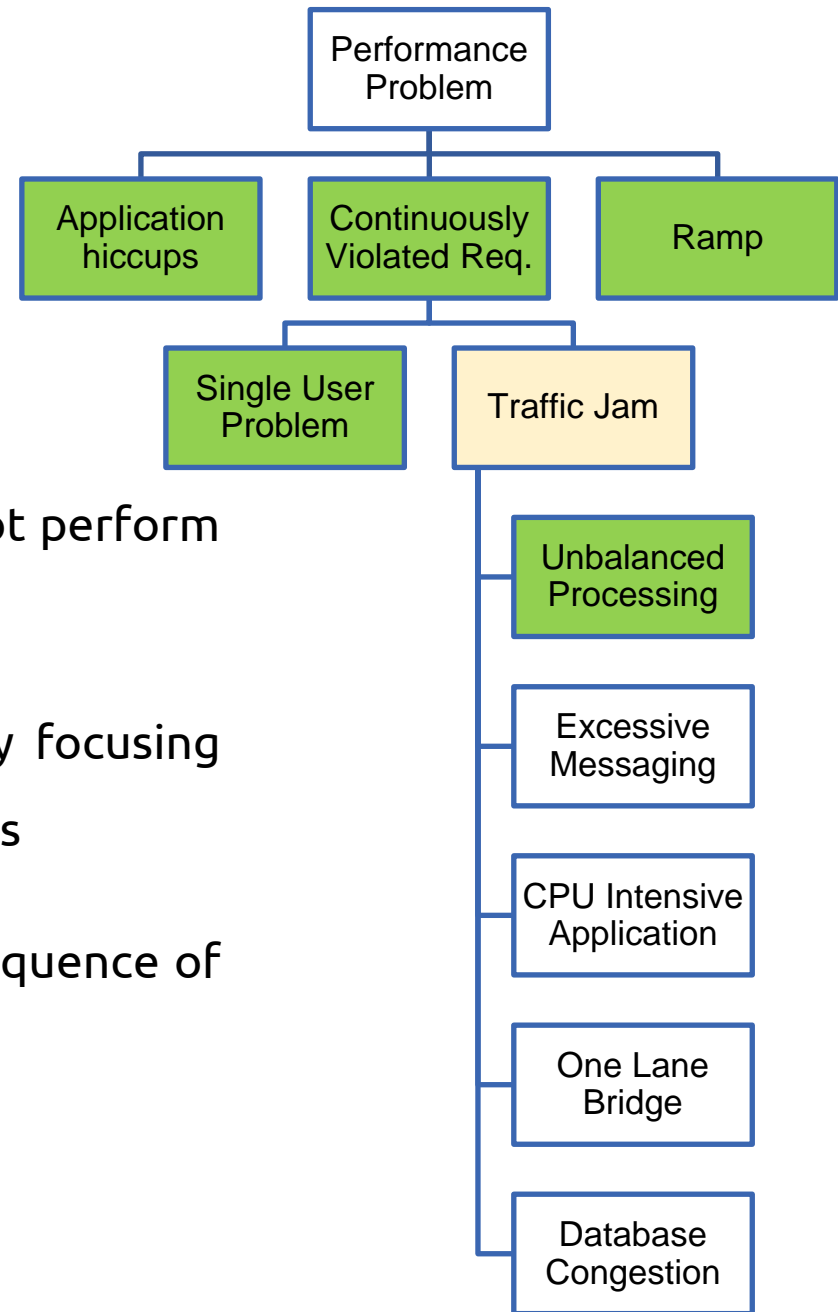
- Research Questions
- Performance problems
- Methodology
- Test setup
- Experiments
 - General performance problem
 - Hiccups
 - Ramp
 - Traffic jam
- Conclusions

Research Questions

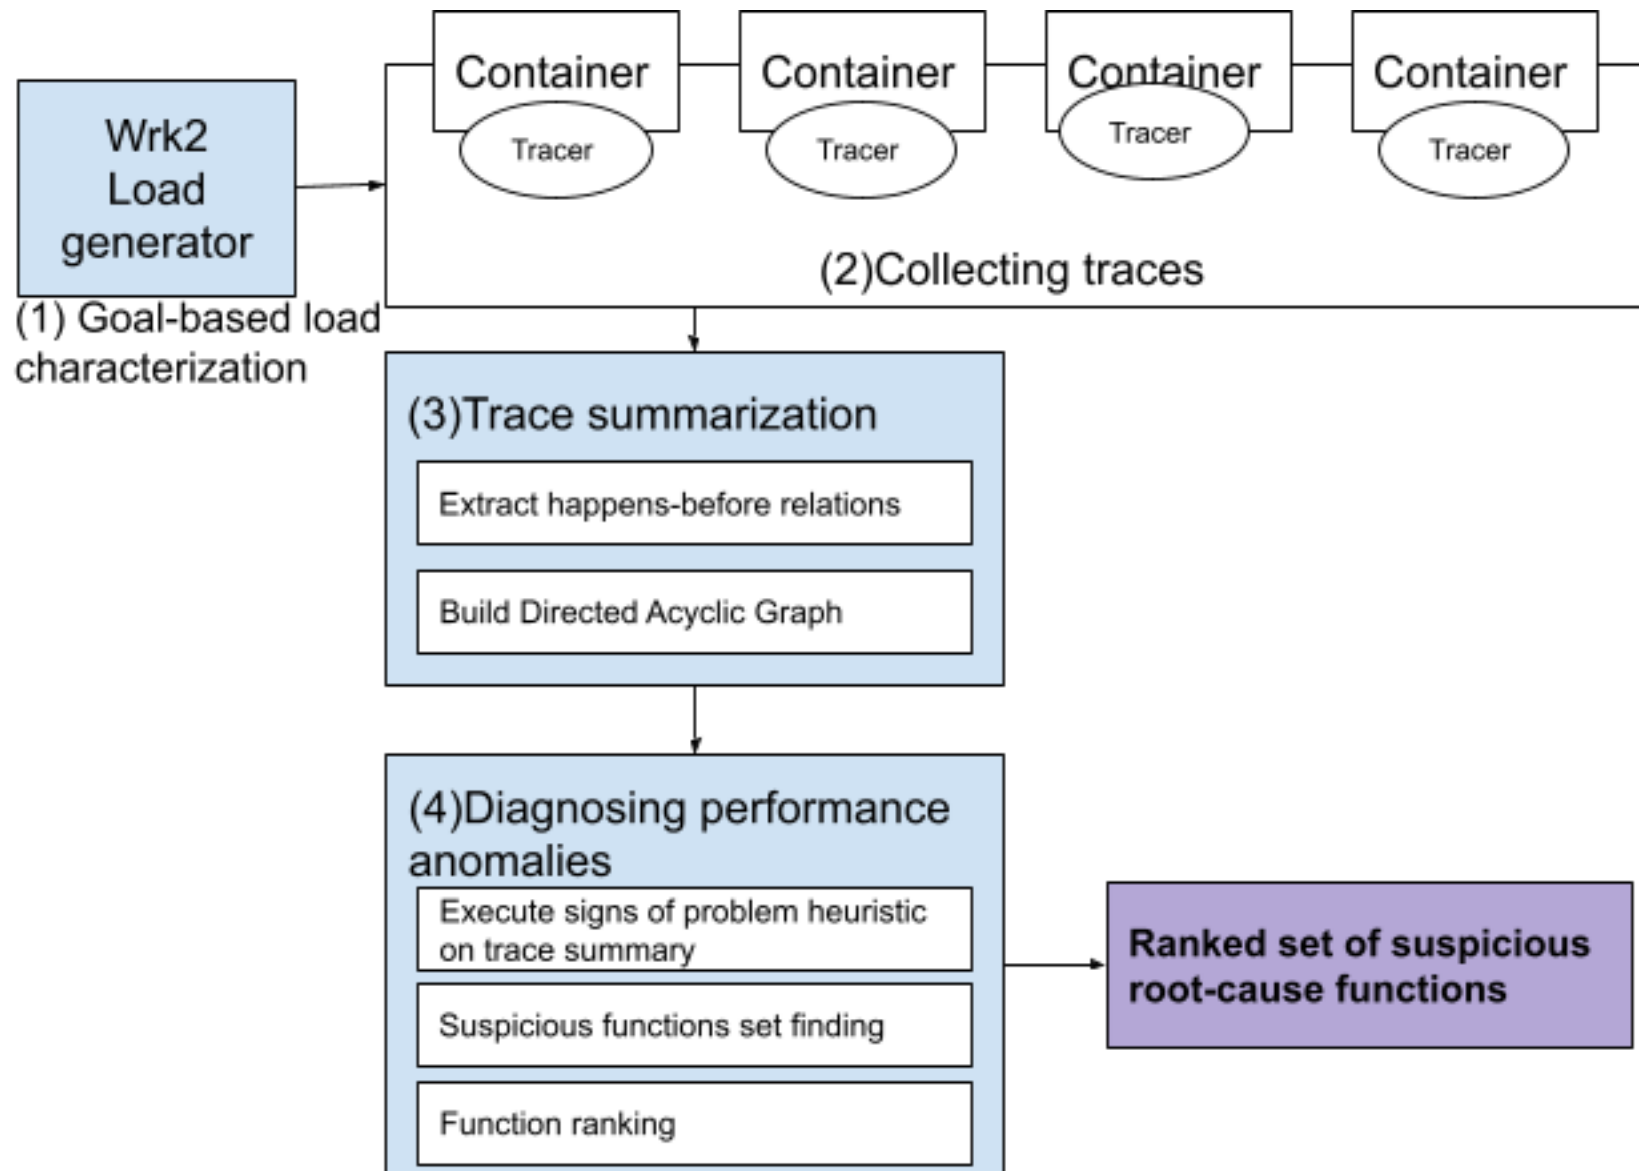
1. Goal based tracing impact on controlling tracing overhead and efficiency
 - By limiting the number of required tracepoints at the start of tracing
2. Determining which tracepoints to activate for goal based tracing
3. Localizing problematic areas in the system, with consecutive tracing adaptation goal in mind
 - By detecting signs of performance anti-patterns

Performance Problems Taxonomy

- Determine if the application generally does not perform well
- Identify patterns of performance changes, by focusing on a specific category of performance problems
- Improve root-cause analysis by locating the sequence of function executions containing the pattern



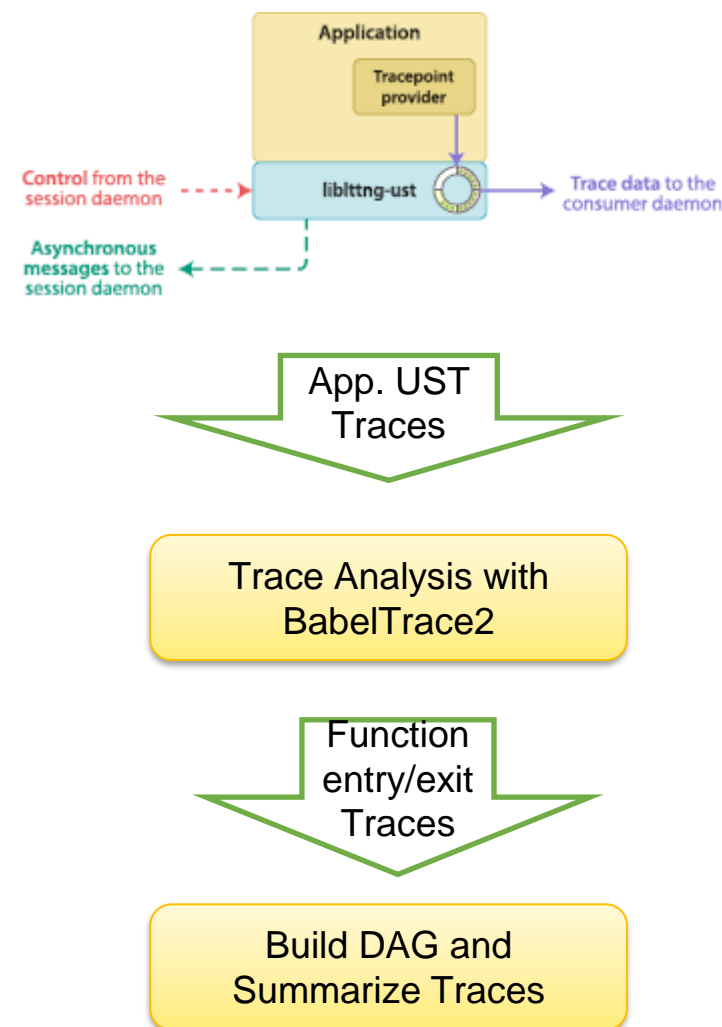
Methodology



Goal-based Tracing

Looking for signs of performance problems

- Looking into signs of performance anti-patterns as the starting point of application tracing
- Provides initial set of candidates for detailed tracing
- Previous work on kernel tracing helps determine main places we should focus on, like candidate services
- Next level is function entry/exit tracing
- The call graph can be constructed from happens-before relations to track root-cause in the chain of events



Tracepoint event and probe samples

```
TRACEPOINT_EVENT(
    masoum_tp,
    ComposePostServiceClient_FUNC_send_UploadUserMentions,
    TP_ARGS(
        const char* , event_type,
        const char* , file_name,
        const char* , func_name,
        int , loc),
    TP_FIELDS(
        ctf_string(event_type_field, event_type)
        ctf_string(file_name_field, file_name)
        ctf_string(func_name_field, func_name)
        ctf_integer(int, loc_field, loc)
    )
)
```

```
void ComposePostServiceClient::send_UploadUserMentions(const int64_t req_id, const std::vector<UserMention> & user_mentions, const std::map<std::string, std::string> & carrier)
{
    tracepoint(masoum_tp, ComposePostServiceClient_FUNC_send_UploadUserMentions, "en", __FILE__, __FUNCTION__, __LINE__);
    int32_t cseqid = 0;
    oprot_>writeMessageBegin("UploadUserMentions", ::apache::thrift::protocol::T_CALL, cseqid);

    ComposePostService_UploadUserMentions_pargs args;
    args.req_id = &req_id;
    args.user_mentions = &user_mentions;
    args.carrier = &carrier;
    args.write(oprot_);

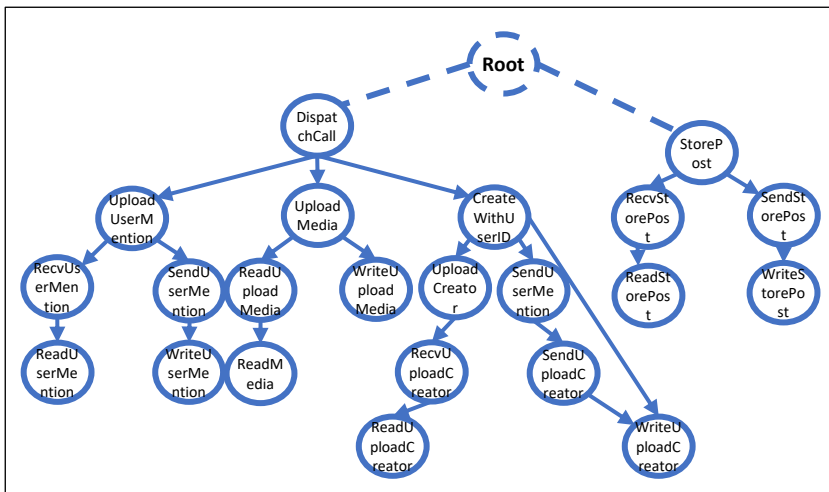
    oprot_>writeMessageEnd();
    oprot_>getTransport()->writeEnd();
    oprot_>getTransport()->flush();
    tracepoint(masoum_tp, ComposePostServiceClient_FUNC_send_UploadUserMentions, "ex", __FILE__, __FUNCTION__, __LINE__);
}
```

Load characterization and Trace Summarization

event_i = (funcId, eventType, timeStamp, ThreadId)

func_i = (funcId, callerId, entryTs, exitTs, duration, threadId)

edge_i = (edgeId, sourceId, targetId, entryTs, exitTs, duration)



Algorithm 1 Build function happens-before relationship directed acyclic graph

```

1:  $T = func_i \leftarrow function - entry/function - exittraces$ 
2:  $list - edge = []$ 
3:  $list - vertex = []$ 
4: for  $Trace_i$  do
5:   Filter events based on threadId
6:   for  $threadId_i$  do
7:      $stack_i = []$ 
8:     for  $event_k$  do
9:       if  $event_{ktype} \leftarrow func - entry$  then
10:         $Push\ stack_i \leftarrow event_k$ 
11:         $caller - func \leftarrow event_k$ 
12:       end if
13:       if  $event_{ktype} == func - exit$  then
14:        Pop
15:         $callee - func \leftarrow event_k$ 
16:       end if
17:        $entry_{ts} \leftarrow caller - entry_{ts}$ 
18:        $exit_{ts} \leftarrow callee - exit_{ts}$ 
19:        $list - edge_i \leftarrow caller - func, callee - func, entry_{ts}, exit_{ts}$ 
20:        $list - vertex_i \leftarrow caller, callee$ 
21:        $response - time_i \leftarrow (callee - exit_{ts}(i)) - (caller - entry_{ts}(i))$ 
22:     end for
23:   end for
24:   Build directed acyclic graph(list-vertex, list-edge)
25:   for  $edge_k$  do
26:      $vertex_k \leftarrow caller - func$ 
27:      $vertex_{k+1} \leftarrow callee - func$ 
28:      $Graph \leftarrow edge_k, response - time_i, vertex_k, vertex_{k+1}$ 
29:     Calculate weight  $edge_k$ 
30:   end for
31: end for

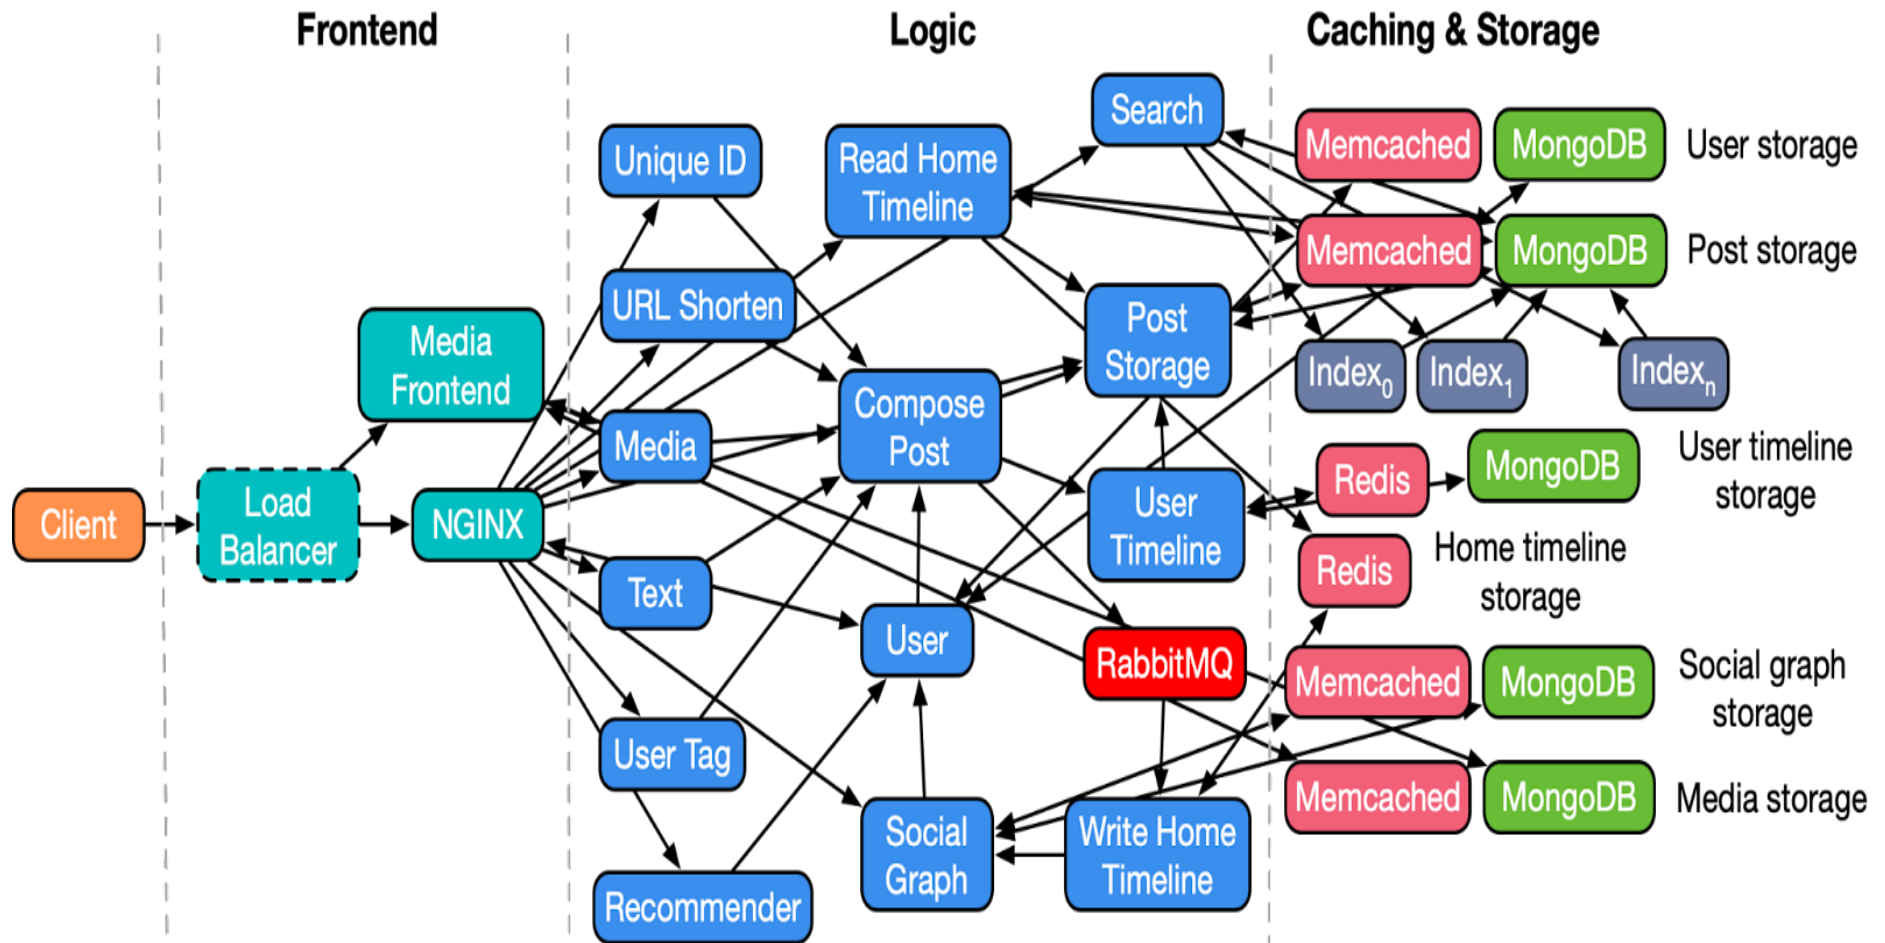
```


Diagnosing performance anomalies

Anti-pattern	Test service	Fault Injection method	Number of traces
Generic performance Problem	Compose Post	Delayed code by adding a long time function	100,000
Hiccup	User Timeline	Delay in 20% of cases on "WriteUserTimeline" function	100,000
Traffic Jam	Compose Post	Changed code to not release lock mutex	100,000
Ramp	Compose Post	No changes	100,000

Test setup

DeathStarBench (SocialNetwork)- test setup.



DeathstarBench- Social Network Service

- Supported actions in DeathStarBench
 - Create text post (optional media: image, video, shortened URL, user tag)
 - Read post
 - Read entire user timeline
 - Receive recommendations on which users to follow
 - Search database for user or post
 - Register/Login using user credentials
 - Follow/Unfollow user

Service	Total LoCs	Communication Protocol	Unique Microservices	Per-language LoC breakdown
Social Network	68061	RPC	36	34% C, 23% C++, 18% Java, 7% node.js, 6% Python, 5% Scala, 3% PHP, 2% Javascript, 2% Go

Code Composition

Experiments

Experiment 1: Generic performance problem

Trace Characterization

Execute signs of
problem heuristic on
trace summary

Suspicious functions set finding

Ranking suspicious functions

Algorithm 2 Generic Performance Problem detection

```

1: Run heavy load on the intended service
2:  $Traces \leftarrow Collect_{function - entry / function - exit} traces$ 
    $caller - callee - relation \leftarrow Algorithm1(Traces)$ 
4: for  $caller - callee - relation_i$  do
   Calculate  $mean - duration - edge_i \leftarrow mean(response - time_i)$ 
6:   for  $edge_i$  do
      $Threshold - mid_i \leftarrow mean - duration - edge_i$ 
8:      $Threshold - factors \leftarrow 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2, 10^3, 10^4$ 
     for  $Threshold - factor_k$  do
10:        $Threshold_{ik} \leftarrow (Threshold - mid_i) * (Threshold - factor_k)$ 
     end for
12:   end for
   end for
14: for  $edge_i$  do
   for  $Threshold_{ik}$  do
16:     if  $response - time_i < Threshold_{ik}$  then
        $has - problem - flag_{ik} \leftarrow False$ 
18:     else
        $has - problem - flag_{ik} \leftarrow True$ 
20:     end if
   end for
22:    $Cumulative_i \leftarrow size(has - problem - flag_i == True)$ 
   end for
24:  $Cumulative - Ratio_i \leftarrow (Cumulative_i) / (size_i)$ 
    $Cumulative - Ratio - set \leftarrow Cumulative - Ratio_i$ 
26:  $Performance - problem - ranking \leftarrow Sort(Cumulative - Ratio - set)$ 

```

General performance problem outputs

- Set of suspicious function calls
- Automatically finding the performance requirements for each function pair (edge), that is required in other problems analysis

Experiment Results

Test case	Load intensity	Services tested	Number of test functions	Functions ranked as problematic
Generic performance problem- Base Load	1 Thread 1 Connection 10 Seconds 1 Requests/Second	read-user-timeline compose-post read-home-timeline	195	No problems observed
Generic performance problem- Normal Load	5 Thread 5 Connections 10 Seconds 5 Requests/Second	read-user-timeline compose-post read-home-timeline	195	Next slide
Generic performance problem- Stress Load	500 Thread 500 Connections 10 Seconds 2000 Requests/Second	read-user-timeline compose-post read-home-timeline	228	Next slide

Experiment Results- Ranking suspicious functions

#	(Service1, Service2)	Edge (Func1, Func2)	Probability
1	(composepost_processor, composepost_UploadUserMentions_args)	(process_UploadUserMentions, read)	0.5
2	(composepost_processor, composepost_UploadUrls_args)	(UploadUrl, read)	0.5
3	(composepost_processor, composepost_UploadUrls_result)	(process_UploadUrl, write)	0.5
4	(composepost_processor, composepost_UploadUserMentions_result)	(process_UploadUserMentions, write)	0.5
5	(composepost_client, composepost_processor)	(dispatchCall, process_UploadUrls)	0.5
6	(composepost_client, composepost_processor)	(dispatchCall, process_UploadUserMentions)	0.5
7	(composepost_UploadUrls_args, socialnetworktypes_Url)	(read, read)	0.5
8	(composepost_UploadUserMentions_args, socialnetworktypes_UserMention)	(read, read)	0.5
9	(composepost_processor, composepost_UploadUniqueld_args)	(process_UploadUniqueld, read)	0.24
10	(composepost_UploadCreator_args, socialnetworktypes_Creator)	(read, read)	0.09

Experiment 2: Hiccup performance problem

Inputs

- Response time series of each thread functions
- Threshold T that shows max response time acceptable for that function
- Maximum allowed proportion the cumulative hiccup time may cover, of the whole experiment duration

Output

- Set of functions that are hiccup candidates

Algorithm 3 Hiccup Performance Problem detection

```

Run normal load on the intended service
2:  $Traces \leftarrow Collect\ function - entry / function - exit\ traces$ 
    $caller - callee - relation \leftarrow Algorithm1(Traces)$ 
4: Calculate performance requirement threshold
   for  $caller - callee - relation_i$  do
6:   Calculate  $mean - duration - edge_i \leftarrow mean(response - time_i)$ 
     for  $edge_i$  do
8:      $Threshold - mid_i \leftarrow mean - duration - edge_i$ 
      $Threshold - factors \leftarrow 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2, 10^3, 10^4$ 
10:    for  $Threshold - factor_k$  do
         $Threshold_{i,k} \leftarrow (Threshold - mid_i) * (Threshold - factor_k)$ 
12:    end for
     end for
14: end for
   Calculate moving percentile size
16: for  $edge_i$  do
     for  $edge_{i,k}$  do
18:    $occurrence - diff_k = entry - timestamp(edge_{i,k+1}) - entry -$ 
      $timestamp(edge_{i,k})$ 
      $occurrence - diff_i \leftarrow occurrence - diff_k$ 
20:    $occurrence - mean_i \leftarrow mean(occurrence - diff_i)$ 
      $moving - percentile_i \leftarrow (occurrence - mean_i) * (number - of -$ 
      $occurrences)$ 
22:   end for
     for  $Threshold_k$  do
24:     for  $moving - percentile_i$  do .....
     end for
26:   end for
      $Cumulative_i \leftarrow size(has - problem - flag_i == True$ 
28: end for
      $Cumulative - Ratio_i \leftarrow (Cumulative_i) / (size_i)$ 
30:  $Cumulative - Ratio - set \leftarrow Cumulative - Ratio_i$ 
      $Performance - problem - ranking \leftarrow Sort(Cumulative - Ratio - set)$ 

```

Experiment Results

Test case	Load intensity	Services tested	# of test functions	Functions ranked as problematic
Hiccup performance problem- Normal Load	5 Thread 5 Connections 10 Seconds 5 Requests/Second	read-user-timeline compose-post read-home-timeline	195	Next slide
Hiccup performance problem- Heavy Load	500 Thread 500 Connections 10 Seconds 2000 Requests/Second	read-user-timeline compose-post read-home-timeline	228	Next slide

Normal and Stress load shows the same results for Hiccup problem identification!


Experiment Results- Ranking suspicious functions

#	(Service1, Service2)	Edge (Func1, Func2)	Probability
1	(usertimeline_client, usertimeline_client)	(WriteUserTimeline, recv_WriteUserTimeline)	0.27
2	(poststorageservice_client, poststorageservice_client)	(StorePost, recv_StorePost)	0.1265
3	(usertimeline_client, usertimeline_result)	(recv_WriteUserTimeline, read)	0.01711
4	(composepost_client, composepost_processor)	(dispatchCall, process_UploadUniqueld)	0.0147
5	(composepost_processor, composepost_UploadCreator_args)	(process_UploadCreator, read)	0.0075
6	(composepost_processor, composepost_UploadUniqueld_args)	(process_UploadUniqueld, read)	0.0040
7	(composepost_processor, composepost_UploadText_args)	(process_UploadText, read)	0.0030
8	(composepost_client, composepost_client)	(dispatchCall, dispatchCall)	0.0024
9	(composepost_client, composepost_processor)	(dispatchCall, process_UploadCreator)	0.0024
10	(composepost_client, composepost_processor)	(dispatchCall, process_UploadText)	0.0009

Experiment 3: Ramp performance problem- Time Window Strategy

There is an occurrence of ramp, if a significant difference between response time of the beginning and end of the load test is observed

1. For each path/func a pairwise comparison of the response time series of neighbouring experiments is conducted
2. For this comparison response time sets for each experiment is gathered and bootstrapped to be normally distributed. Then a t-test is applied to compare
3. If comparison of starting and ending sets results in significant growth of response time, then ramp anti-pattern is observed



Experiment Results


Test case	Load intensity	Services tested	# of test functions	(Service, Function) ranked as problematic
Ramp performance problem- Constant Heavy Load	<ul style="list-style-type: none"> 5 intervals of putting heavy load and then gathering traces in base load Base load configuration: 1 Thread, 1 Connections, 10 Seconds, 5 Requests/Second <ul style="list-style-type: none"> Heavy load configuration: 200 Thread, 200 Connections, 120 Seconds, 400 Requests/Second 	read-user-timeline compose-post read-home-timeline	128	1.(HomeTimelineServiceProcessor, dispatchCall) 2.(MediaServiceProcessor, dispatchCall)
Ramp performance problem- Constant Stress Load	<ul style="list-style-type: none"> 5 intervals of increasingly putting heavy load and then gathering traces in base load Base load configuration: 1 Thread, 1 Connections, 10 Seconds, 5 Requests/Second <ul style="list-style-type: none"> Base to max Heavy load configuration: 200 Thread, 200 Connections, 120 Seconds, 400 Requests/Second 	read-user-timeline compose-post read-home-timeline	128	1.(HomeTimelineServiceProcessor, dispatchCall) 2.(MediaServiceProcessor, dispatchCall)

Constant heavy load and Increasingly Heavier load had the same results in our experiments!

Experiment 4: Traffic Jam performance problem- Time Window Strategy

There is an occurrence of Traffic Jam, if a significant difference between response time of all load tests is observed

1. For each path/func a pairwise comparison of the response time series of neighbouring experiments is conducted
2. For this comparison response time sets for each experiment is gathered and bootstrapped to be normally distributed. Then a t-test is applied to compare
3. If comparison of all sets results in significant growth of response time, then Traffic Jam anti-pattern is observed



Experiment Results

Test case	Load intensity	Services tested	# of test functions	(Service, Function) ranked as problematic
Traffic Jam performance problem	Put heavy Load 1: 20 Thread, 20 Connections 120 Seconds, 20 Requests/Second	read-user-timeline compose-post read-home-timeline	239	1.(HomeTimelineServiceProcessor, dispatchCall)
	Put heavy Load 2: 100 Thread, 100 Connections 120 Seconds, 100 Requests/Second			2.(MediaServiceProcessor, dispatchCall)
	Put heavy Load 3: 200 Thread, 200 Connections 120 Seconds, 400 Requests/Second			3.(SocialGraphServiceProcessor, dispatchCall)
	Put heavy Load 4: 500 Thread, 500 Connections 120 Seconds, 2000 Requests/Second			4.(GetFollowers_result, write)
				5.(GetFollowers_result, read)

Conclusions

- Ranking candidate functions to trace in more detail, based on the provided tracing goals
- Increased efficiency of tracing by focusing on specific tracing goals