

# QUERYING

# TRACES

## ADMIN

◦ Midterm scores, etc ◦ Will send out over weekend/by Monday

◦ Final Project

↳ Haven't really heard back for the cases where I asked for clarification

→ New Task ◦ This week: Post what you plan to work on + rough timeline

START WORKING ON FINAL PROJECT!!

◦ STATE OF THE COURSE

→ Topics of interest?

→ Any interest in student presenters?

Where We ARE

What are traces & how collected

- ↳ Dapper
- ↳ Pivot Tracing
- ↳ Snicker

Potential (but uncommon) uses

- ↳ Replay

Real Uses

- ↳ Critical Path tracing
- ↳ Find Bottlenecks

Today: "Querying" traces.

↳ What request traces do you want to look at?

↳ Depends on what one is trying to do?

↳ Why is tail-latency high?

- Cannot just look at critical path

↳ Must decide what requests to check

↳ Localize what part of a service

...

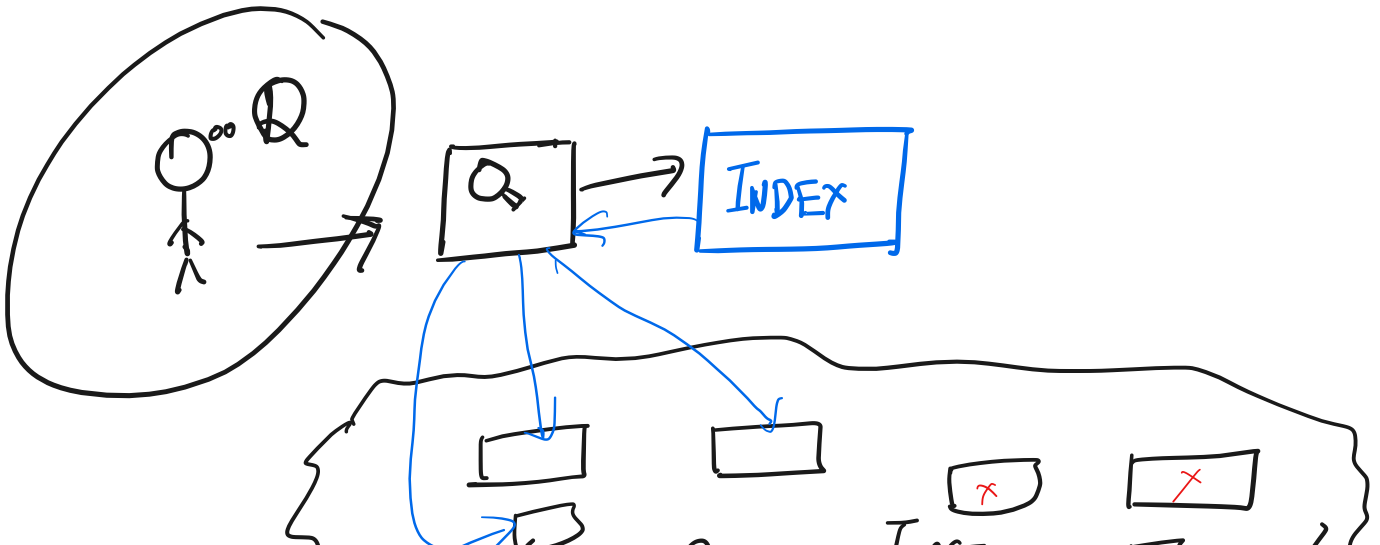
↳ Why are requests slower now?

↳ Pick two or more request traces & compare.

But which ones?

...

What is desirable



Another way to look at the problem: How to build an INDEX?

Q: What type of index is useful?

↳ Depends on common queries.

[tprof] ↳ Find requests whose latency resource utilization is >  $\frac{\text{abs value } n^{\text{th}} \text{ percentile}}{\text{average}}$

[tprof] → Find requests where component has ✓


[GNTA] → Find requests corresponding to some business logic

[GNTA] ↳ Find requests that are affected by a recent code change.

ooo

Of course we also want to allow people to combine these queries.

Want indices with two types of information



(But over what?)

## New Problems

(a) How to update index?

Request  $\xrightarrow{\text{assign}}$  Type

+ Compute aggregate

(b) Index selection: What types to use-

↳ Automatic [tprof] (sort of)

→ Semi-automatic [GMTA]

## Observations for computing aggregates

↳ Need to compute aggregates over groups

Q1: How to group?

- Determines what types of aggregates are meaningful.  
↳ OR WHAT THEY SHOW.

Example: - 99 percentile of B's latency when called from A must be computed only on  $A \rightarrow B$

M:  $\pm 1$  on B when processing

requests of type X

...

Observe: Can always compute an aggregate on demand, but not efficient.

Desirable: Maintain stats ahead of time by choosing groupings that are likely to be DESIRABLE.

What Groupings are desirable

- Types indexed
  - Request type
  - user
  - Business flow

- Finer grained structures
  - Service

- Request-level call graph
  - subspans?
- Will discuss later.

...

...

Q2: What aggregates to maintain!

Ex. Request type



Note: Related to but not the same as groupings.  
Why?

So really, we are back to the question of what granularity of stats/aggregation are people likely to find useful.


The obvious ones:

- Request type

- By service.

Looking Beyond (sort of)





Similar control/data flow.

Why?

Within a service

...  
fl(...)

lock(x) } Time spent here ?  
 : }  
 unlock(x) } ↳ Or elsewhere.



f2(...)

...

Why interesting?

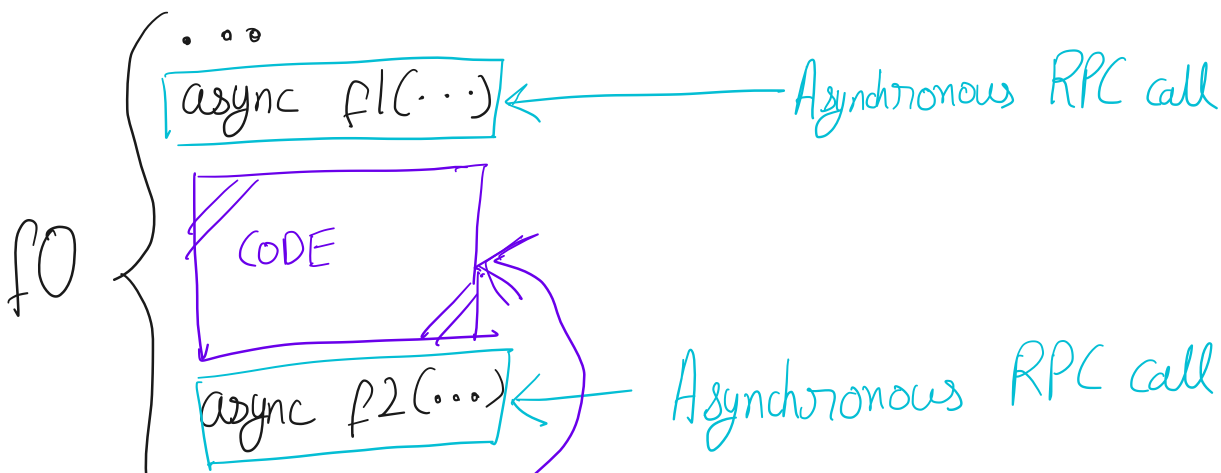
Problem

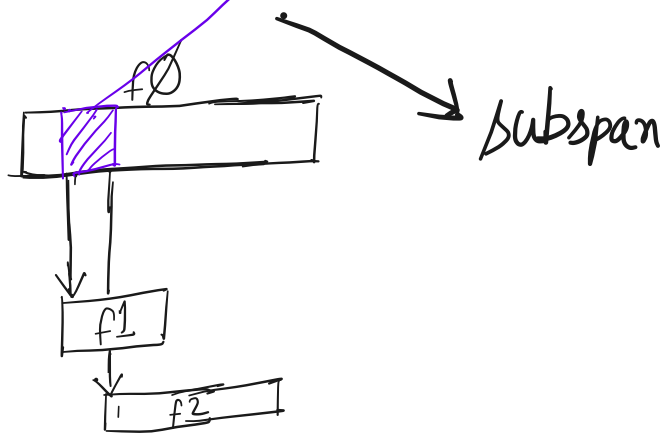
↳ Generally do not have this information\*

↳ Could add it but...

↳ Can we infer it?


Answer: Not really, but can try to approximate in some cases









More generally

-  $f_0$  cannot be waiting for an RPC to complete (?)



Problem: Not very accurate.





Can we use previous tricks to do better?

- Aggregate by dependency/service

- Aggregate by "Business Process"

## Where We Are

- Have a way to assign type to request

↳ Reusing some ideas from before + information in traces

\* Request type

\* Dependency information → request type  
← service granularity  
by mining previous traces

\* Call-graph → spans

\* ...

- Have a way to aggregate metrics for each type

↳ Using info contained directly in traces

+ some inferred information

Q: How to use any of this?

- Great that one can query and get requests, but how do people use it in general?

- CPA: Points out bottleneck - allows one to decide where to focus optimization efforts.

- But in general often want to determine what is different

↳ How do slow & fast request differ?

→ How do old & new system differ for requests?

→ ..