

Auditing Outsourced Services

Cheng Tan

Courant Institute, New York University

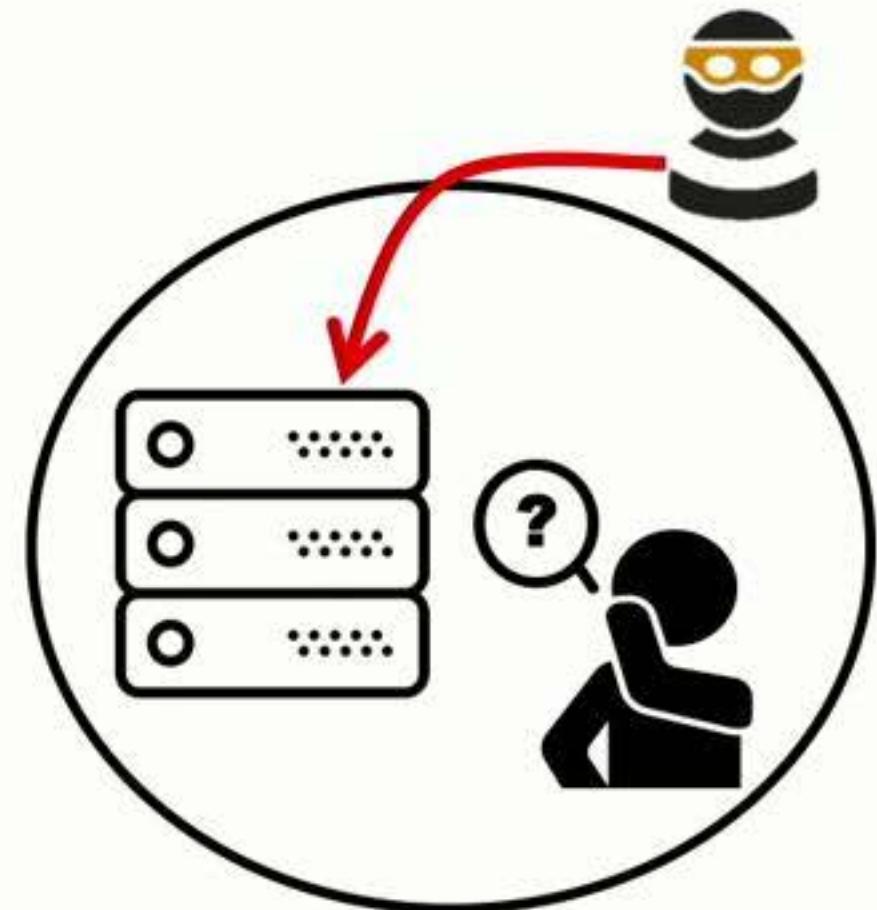
chengtan@cs.nyu.edu



Does my remote service
behave as promised?



Are there unexpected
internal failures?



Are there
external attacks?

Remote servers (clouds) make mistakes

- bugs:
- misconfigs:
- ...

threat **post** **Dangerous Kubernetes Bugs Allow Authentication Bypass, DoS**
 **Security misconfigurations: Code for human error**

Remote servers (clouds) make mistakes

- bugs:
- misconfigs:
- untrusted cloud:

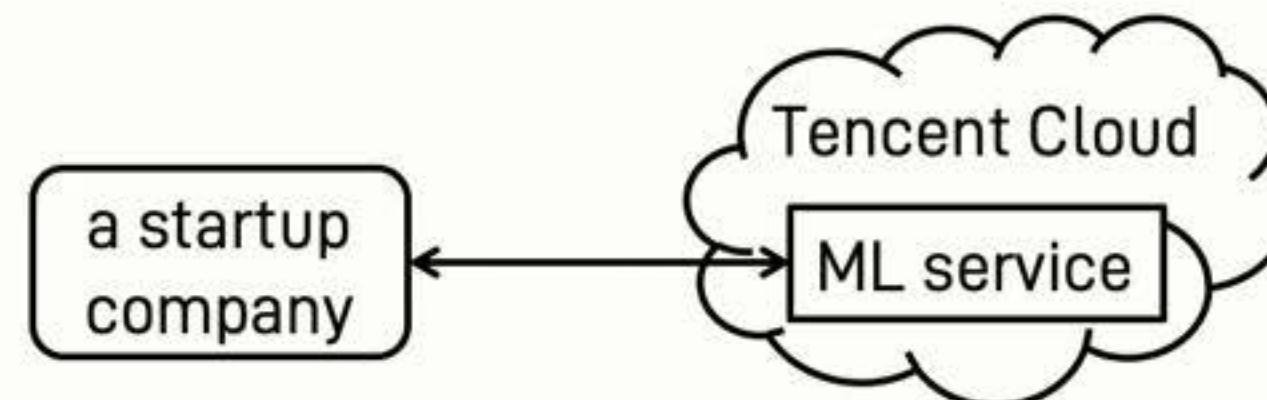
threat **post** Dangerous Kubernetes Bugs Allow Authentication Bypass, DoS



Security misconfigurations:
Code for human error



Tencent Cloud denies technical attack against rival



Remote servers (clouds) make mistakes

- bugs:
- misconfigs:
- untrusted cloud:

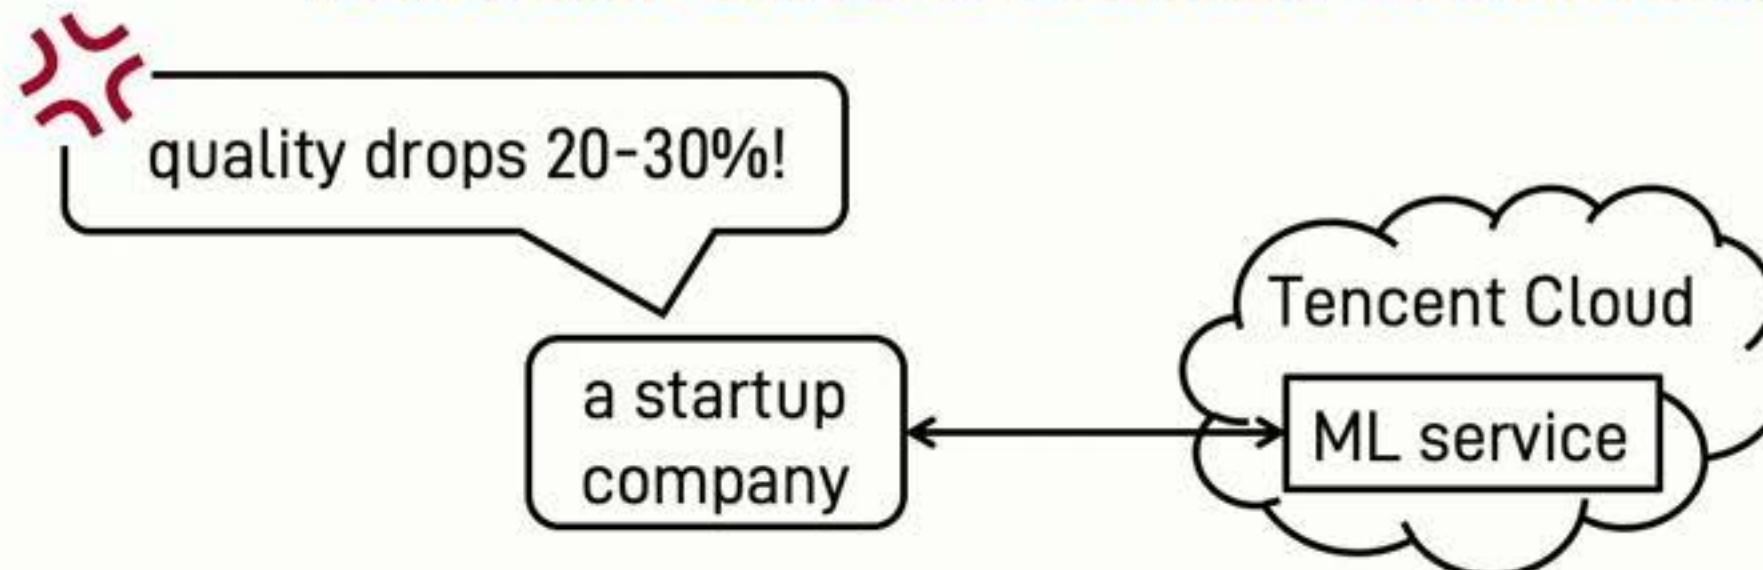
threat **post** Dangerous Kubernetes Bugs Allow Authentication Bypass, DoS



Security misconfigurations:
Code for human error



Tencent Cloud denies technical attack against rival

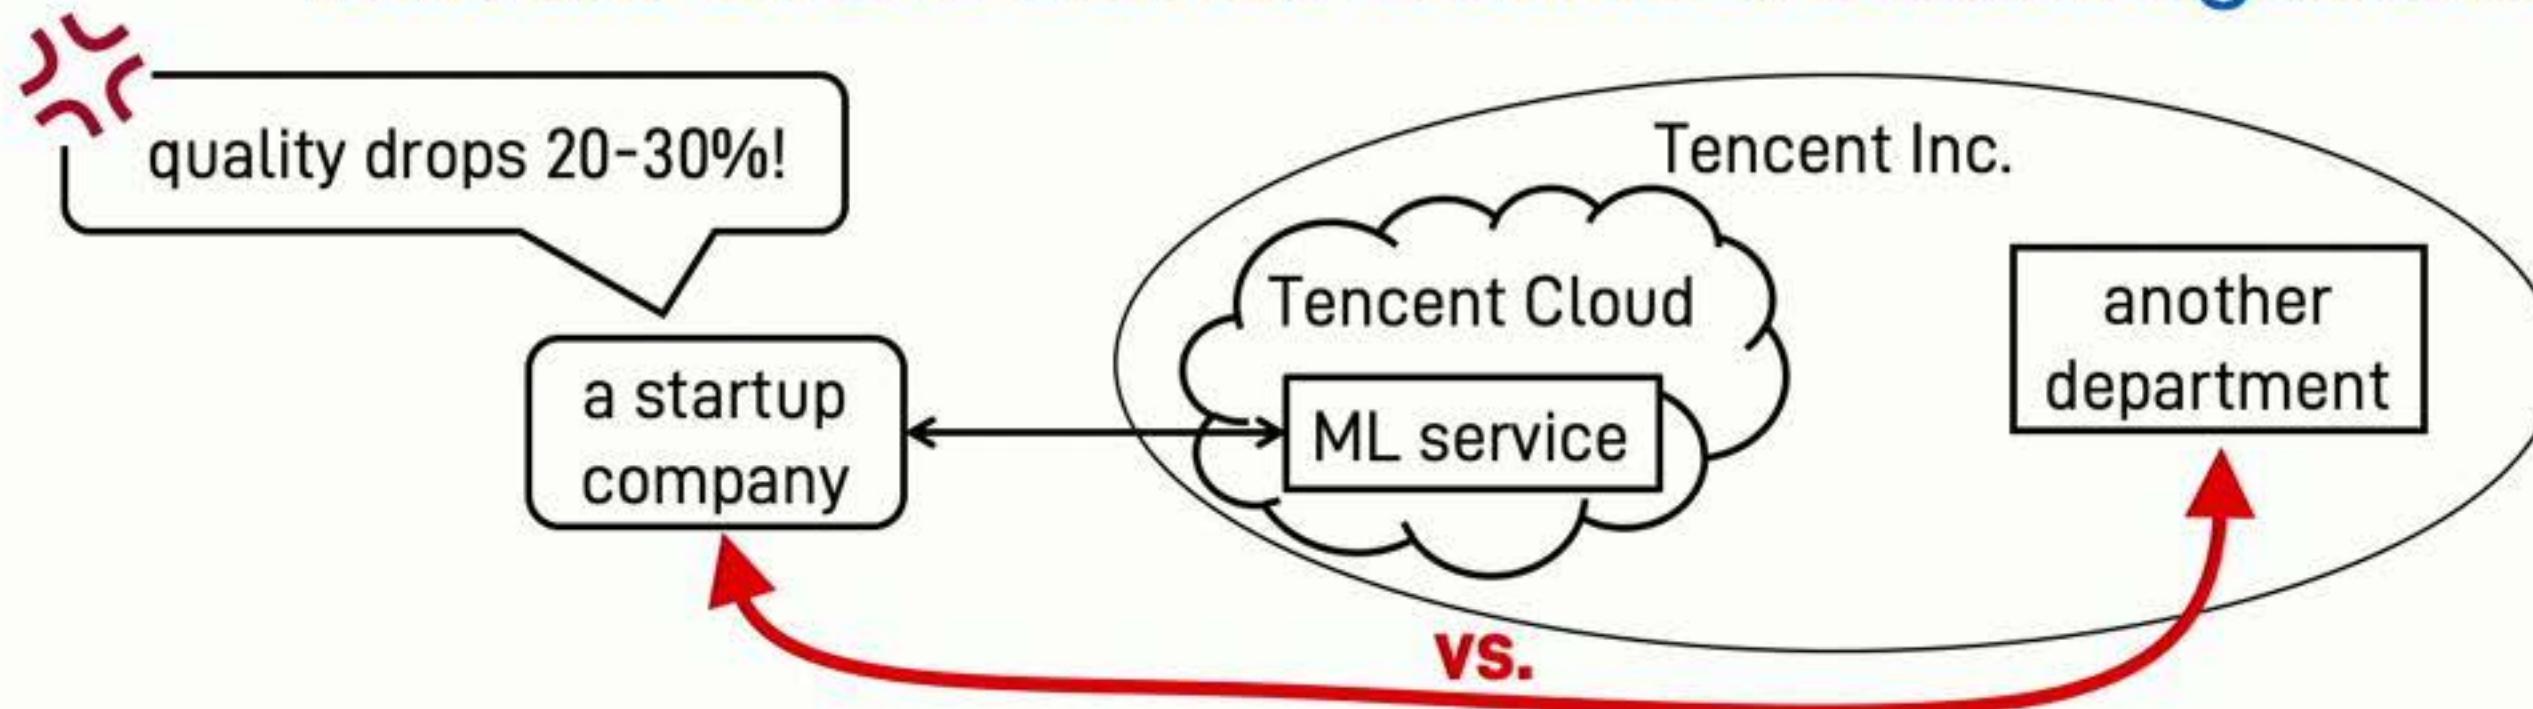


Remote servers (clouds) make mistakes

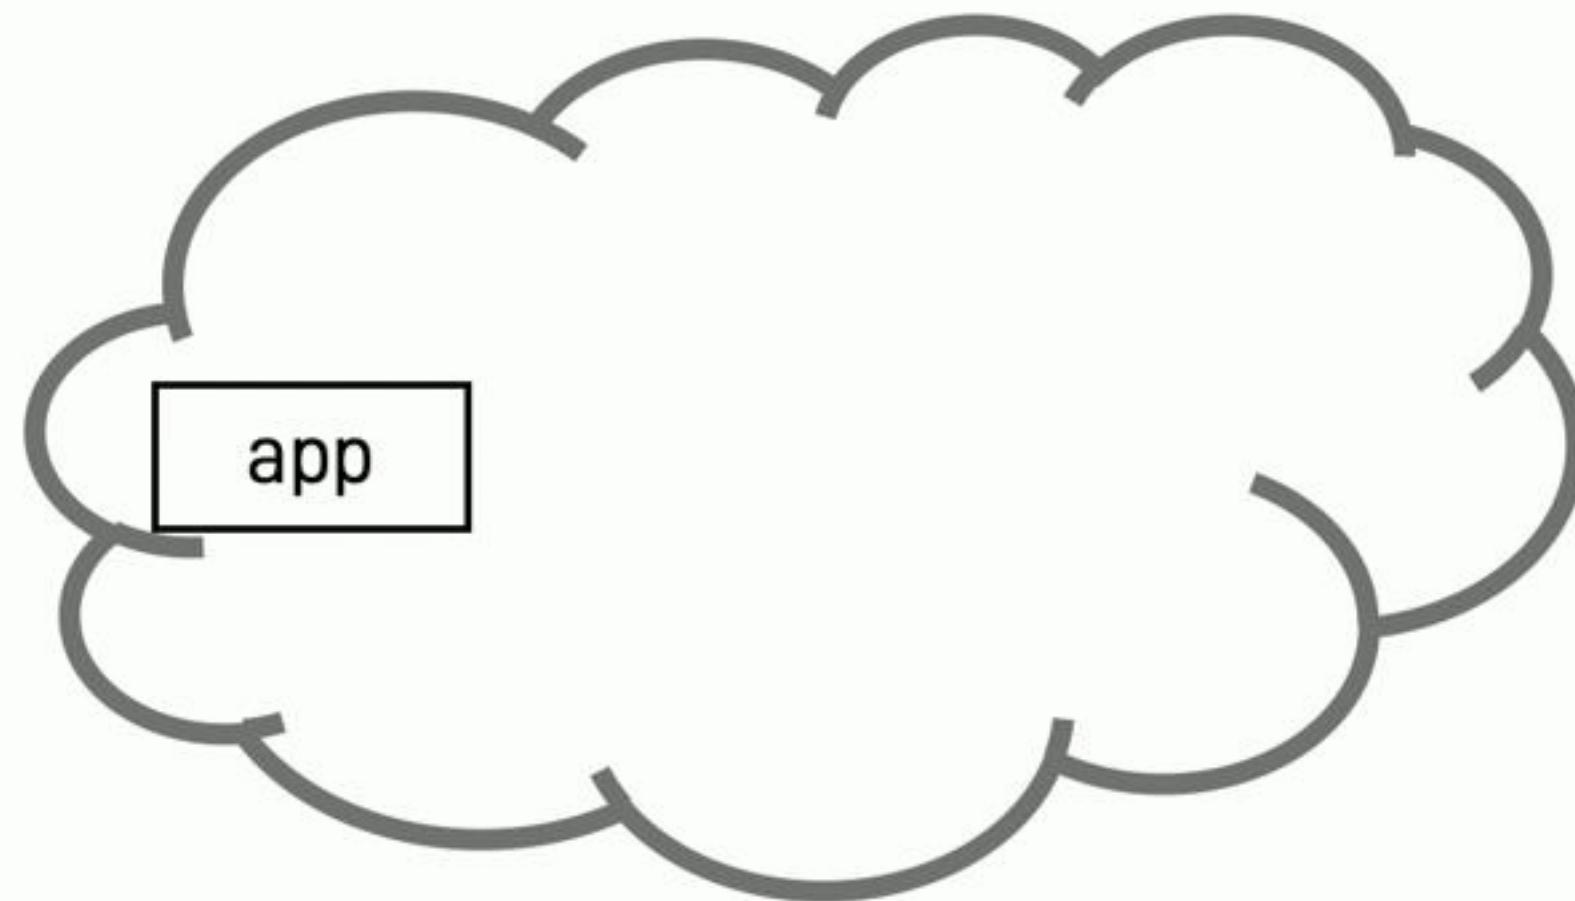
- bugs:
- misconfigs:
- untrusted cloud:



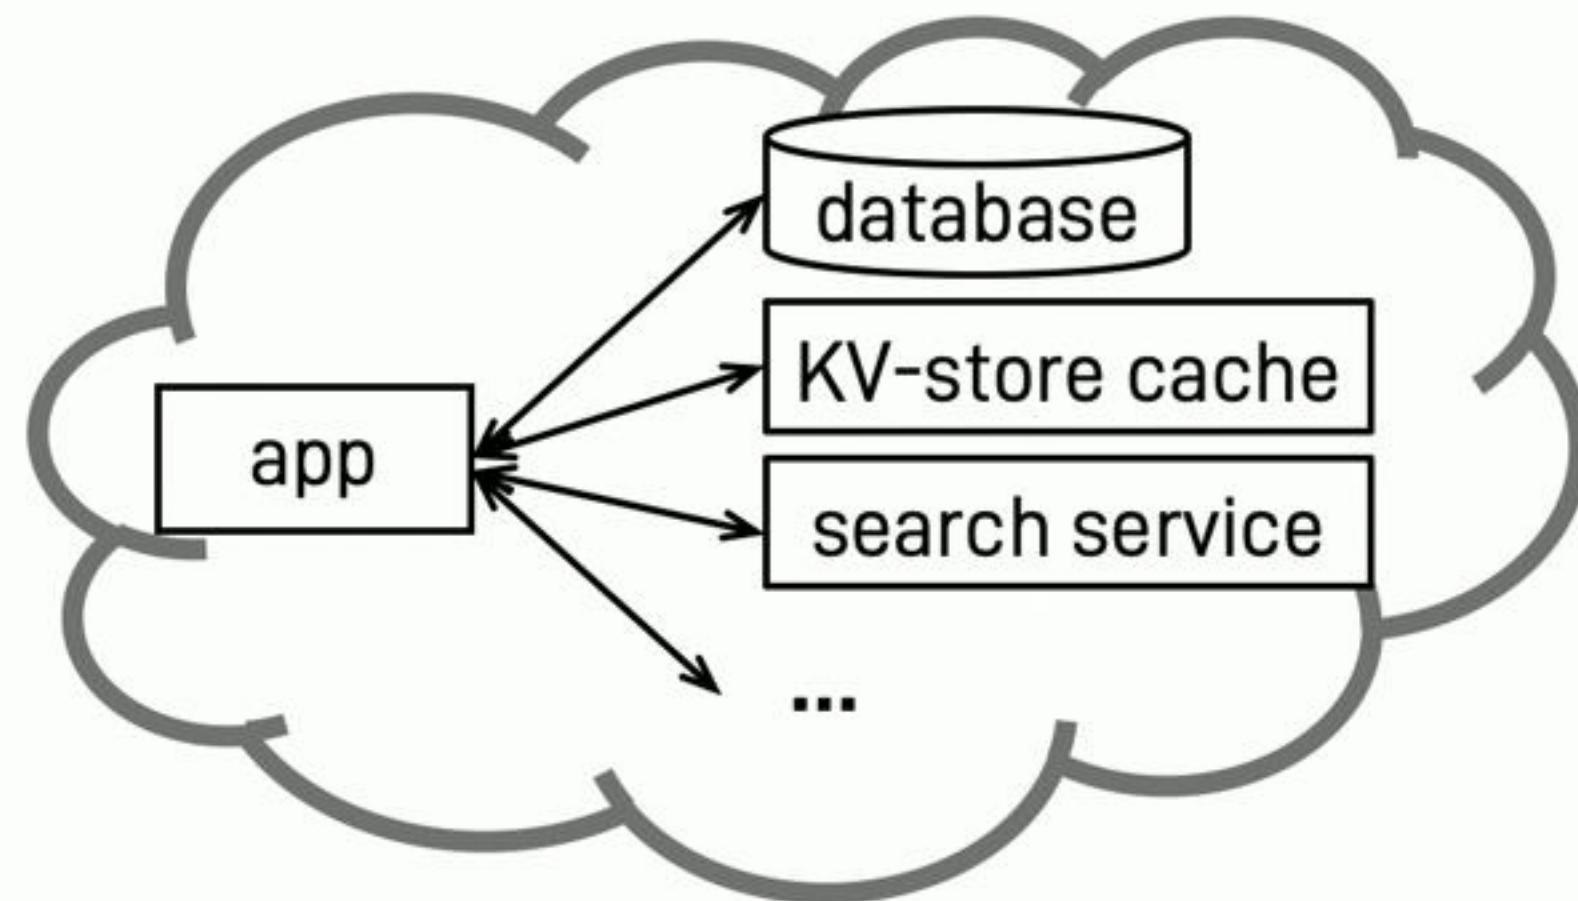
Tencent Cloud denies technical attack against rival



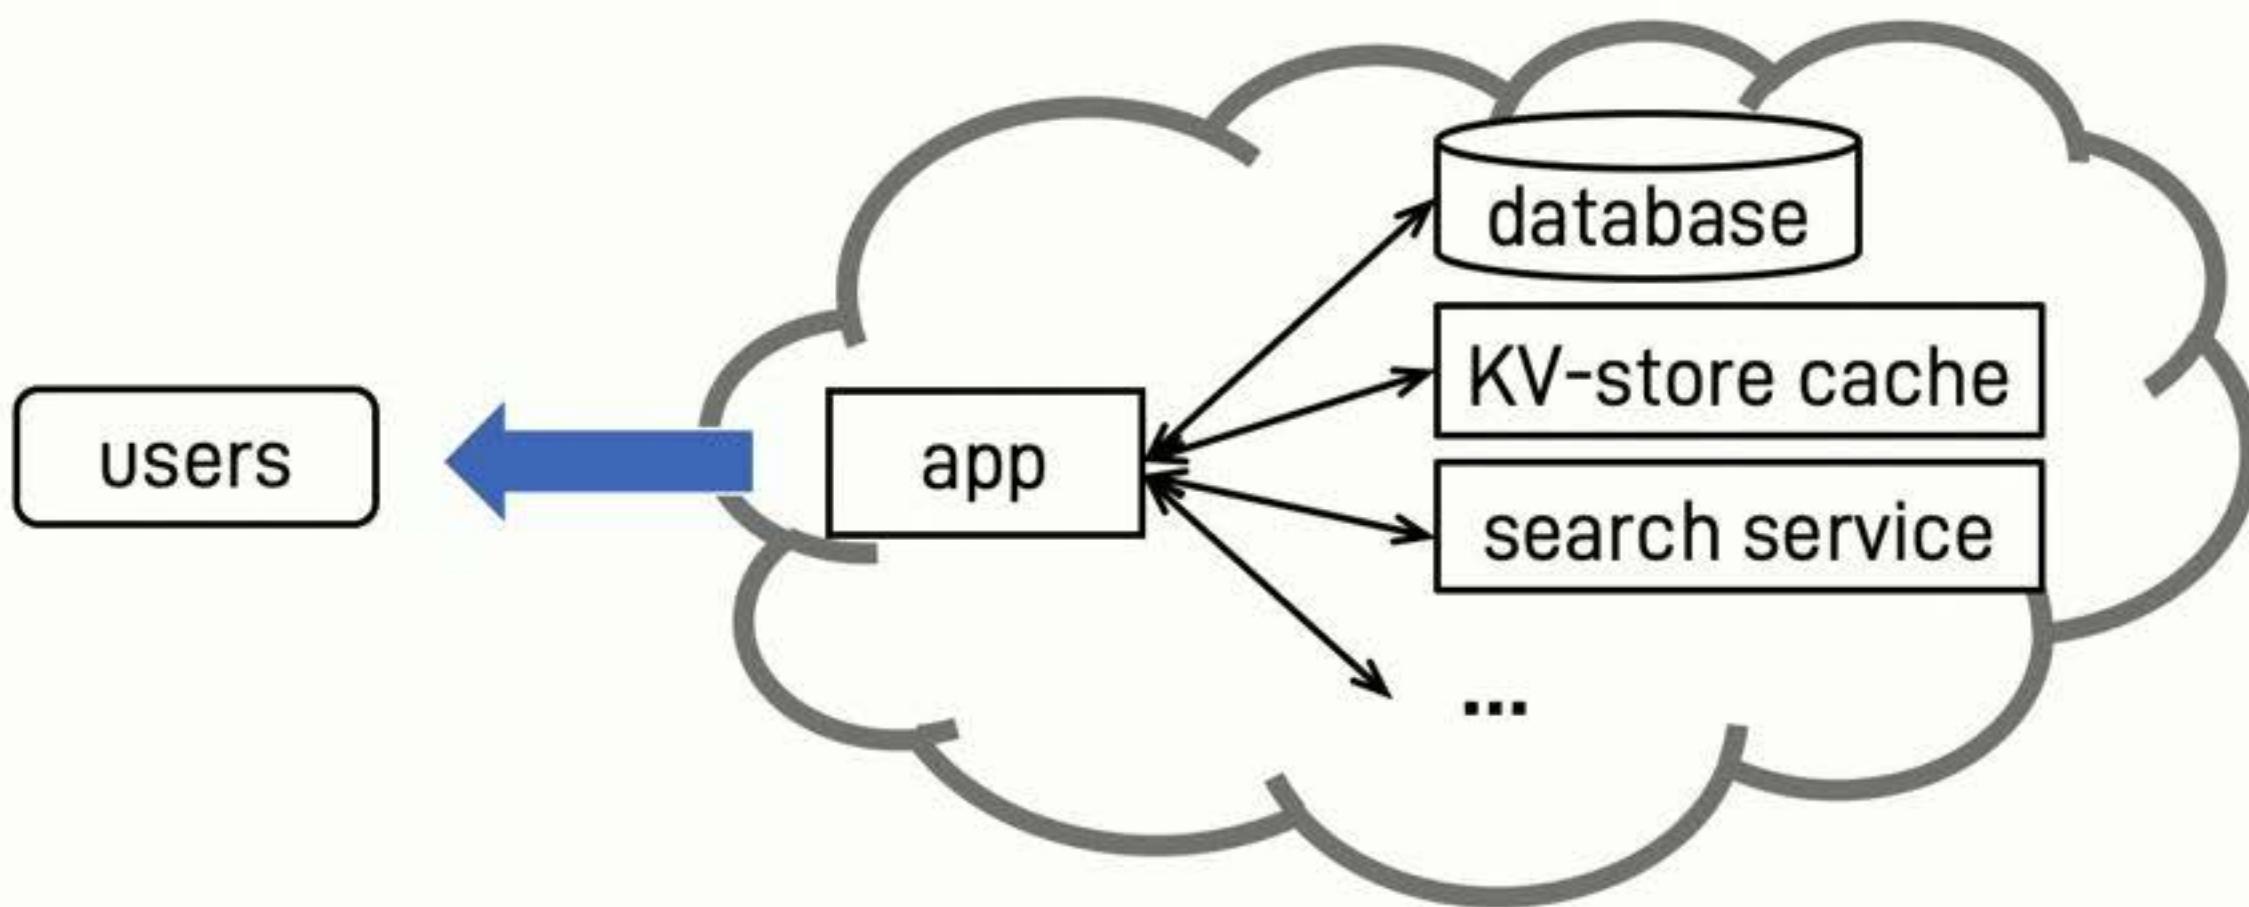
Verify outsourced services with systems



Verify outsourced services with systems



Verify outsourced services with systems

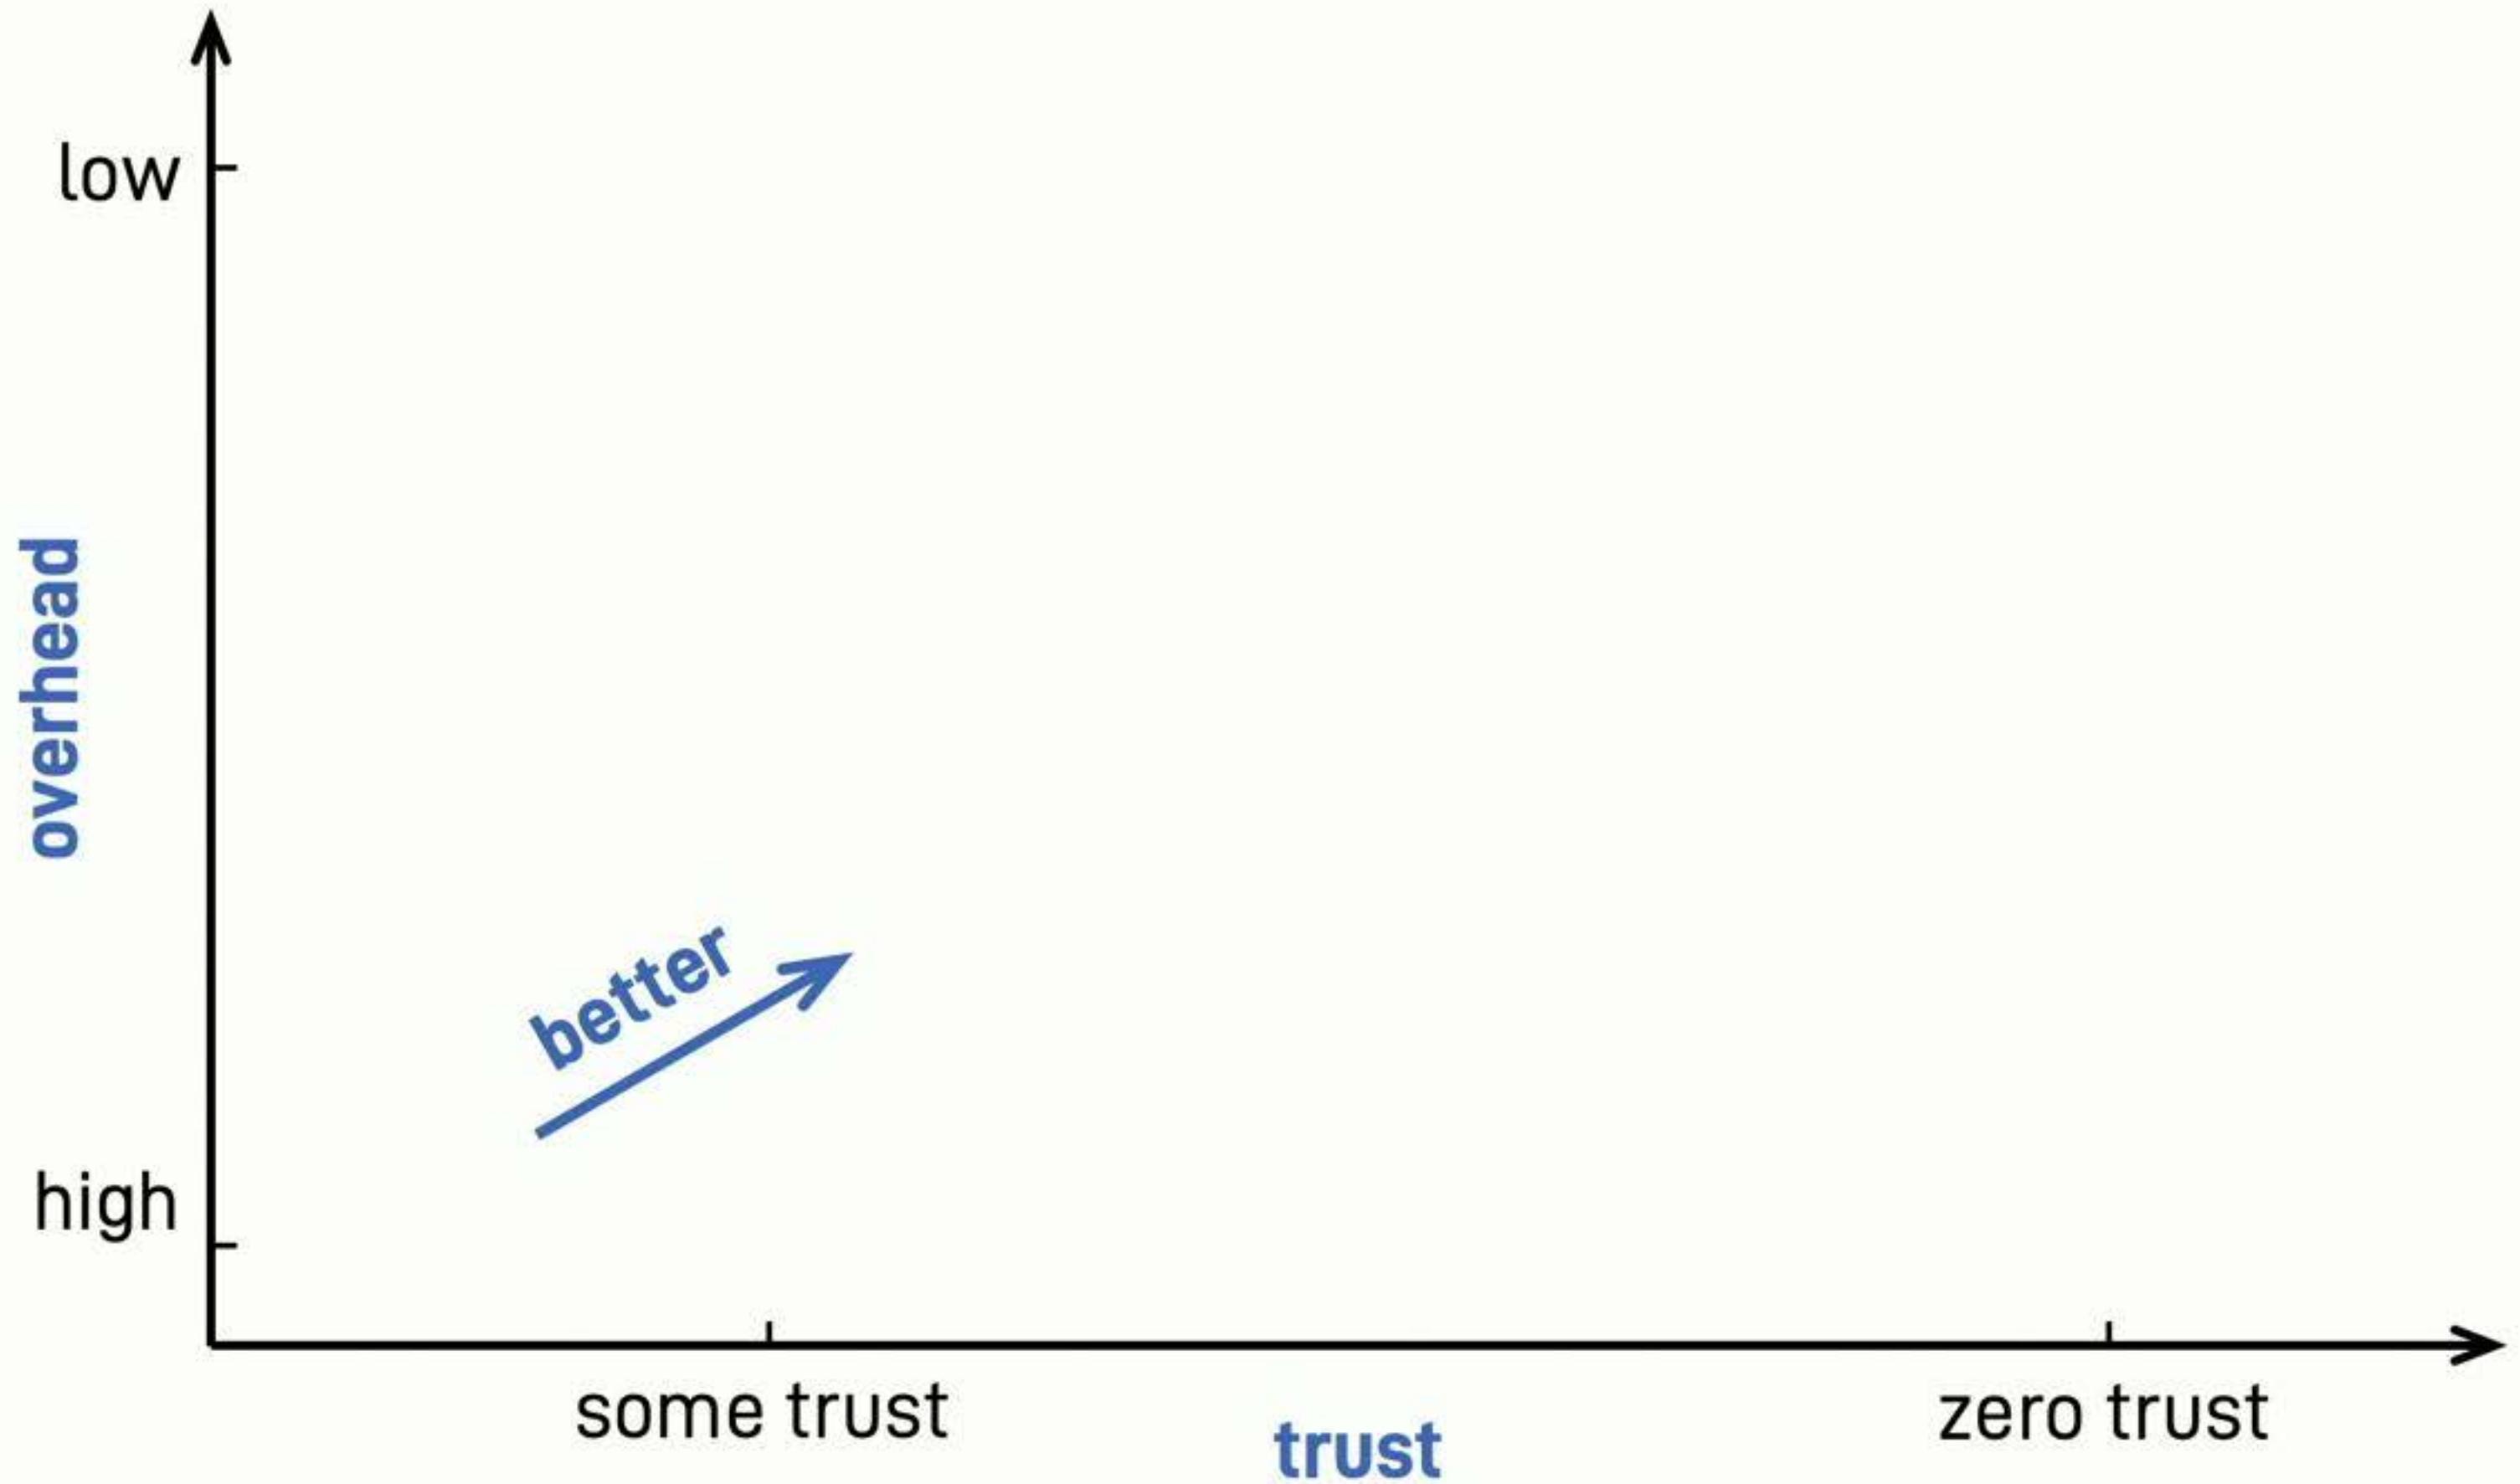


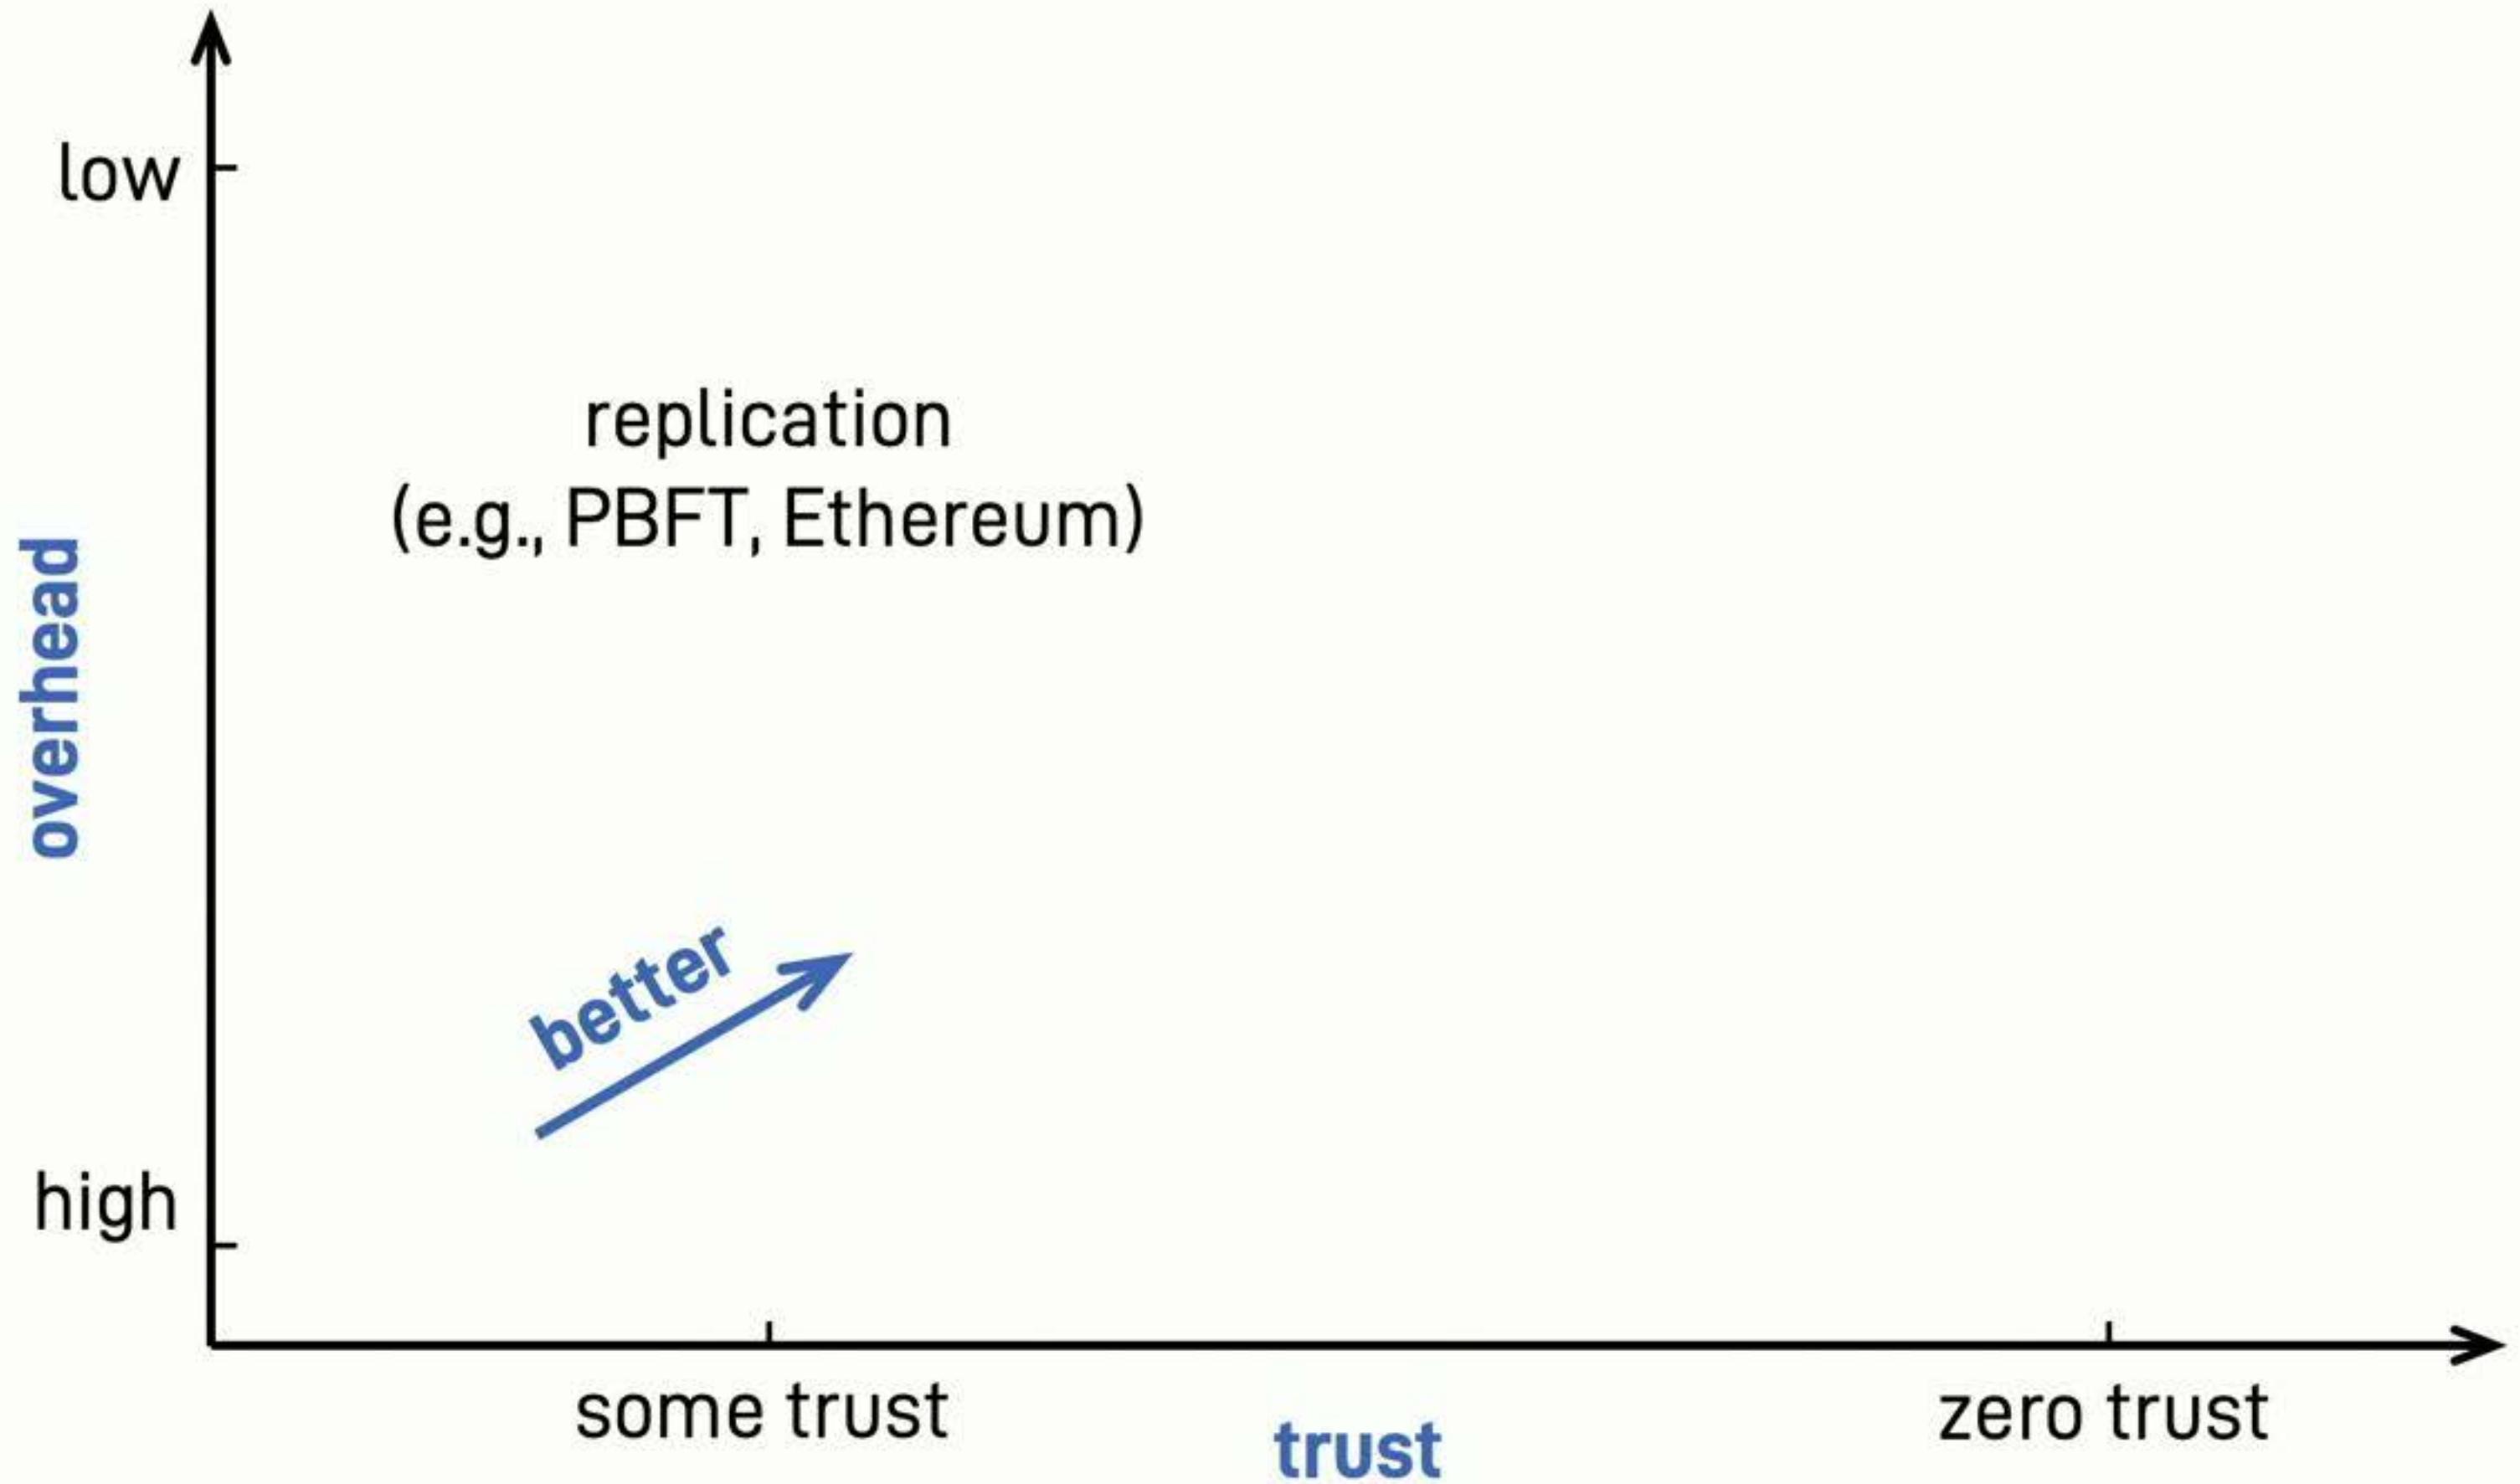
- Honest servers can convince users that app is faithfully executed.
- Buggy or malicious servers cannot pass users' checks.

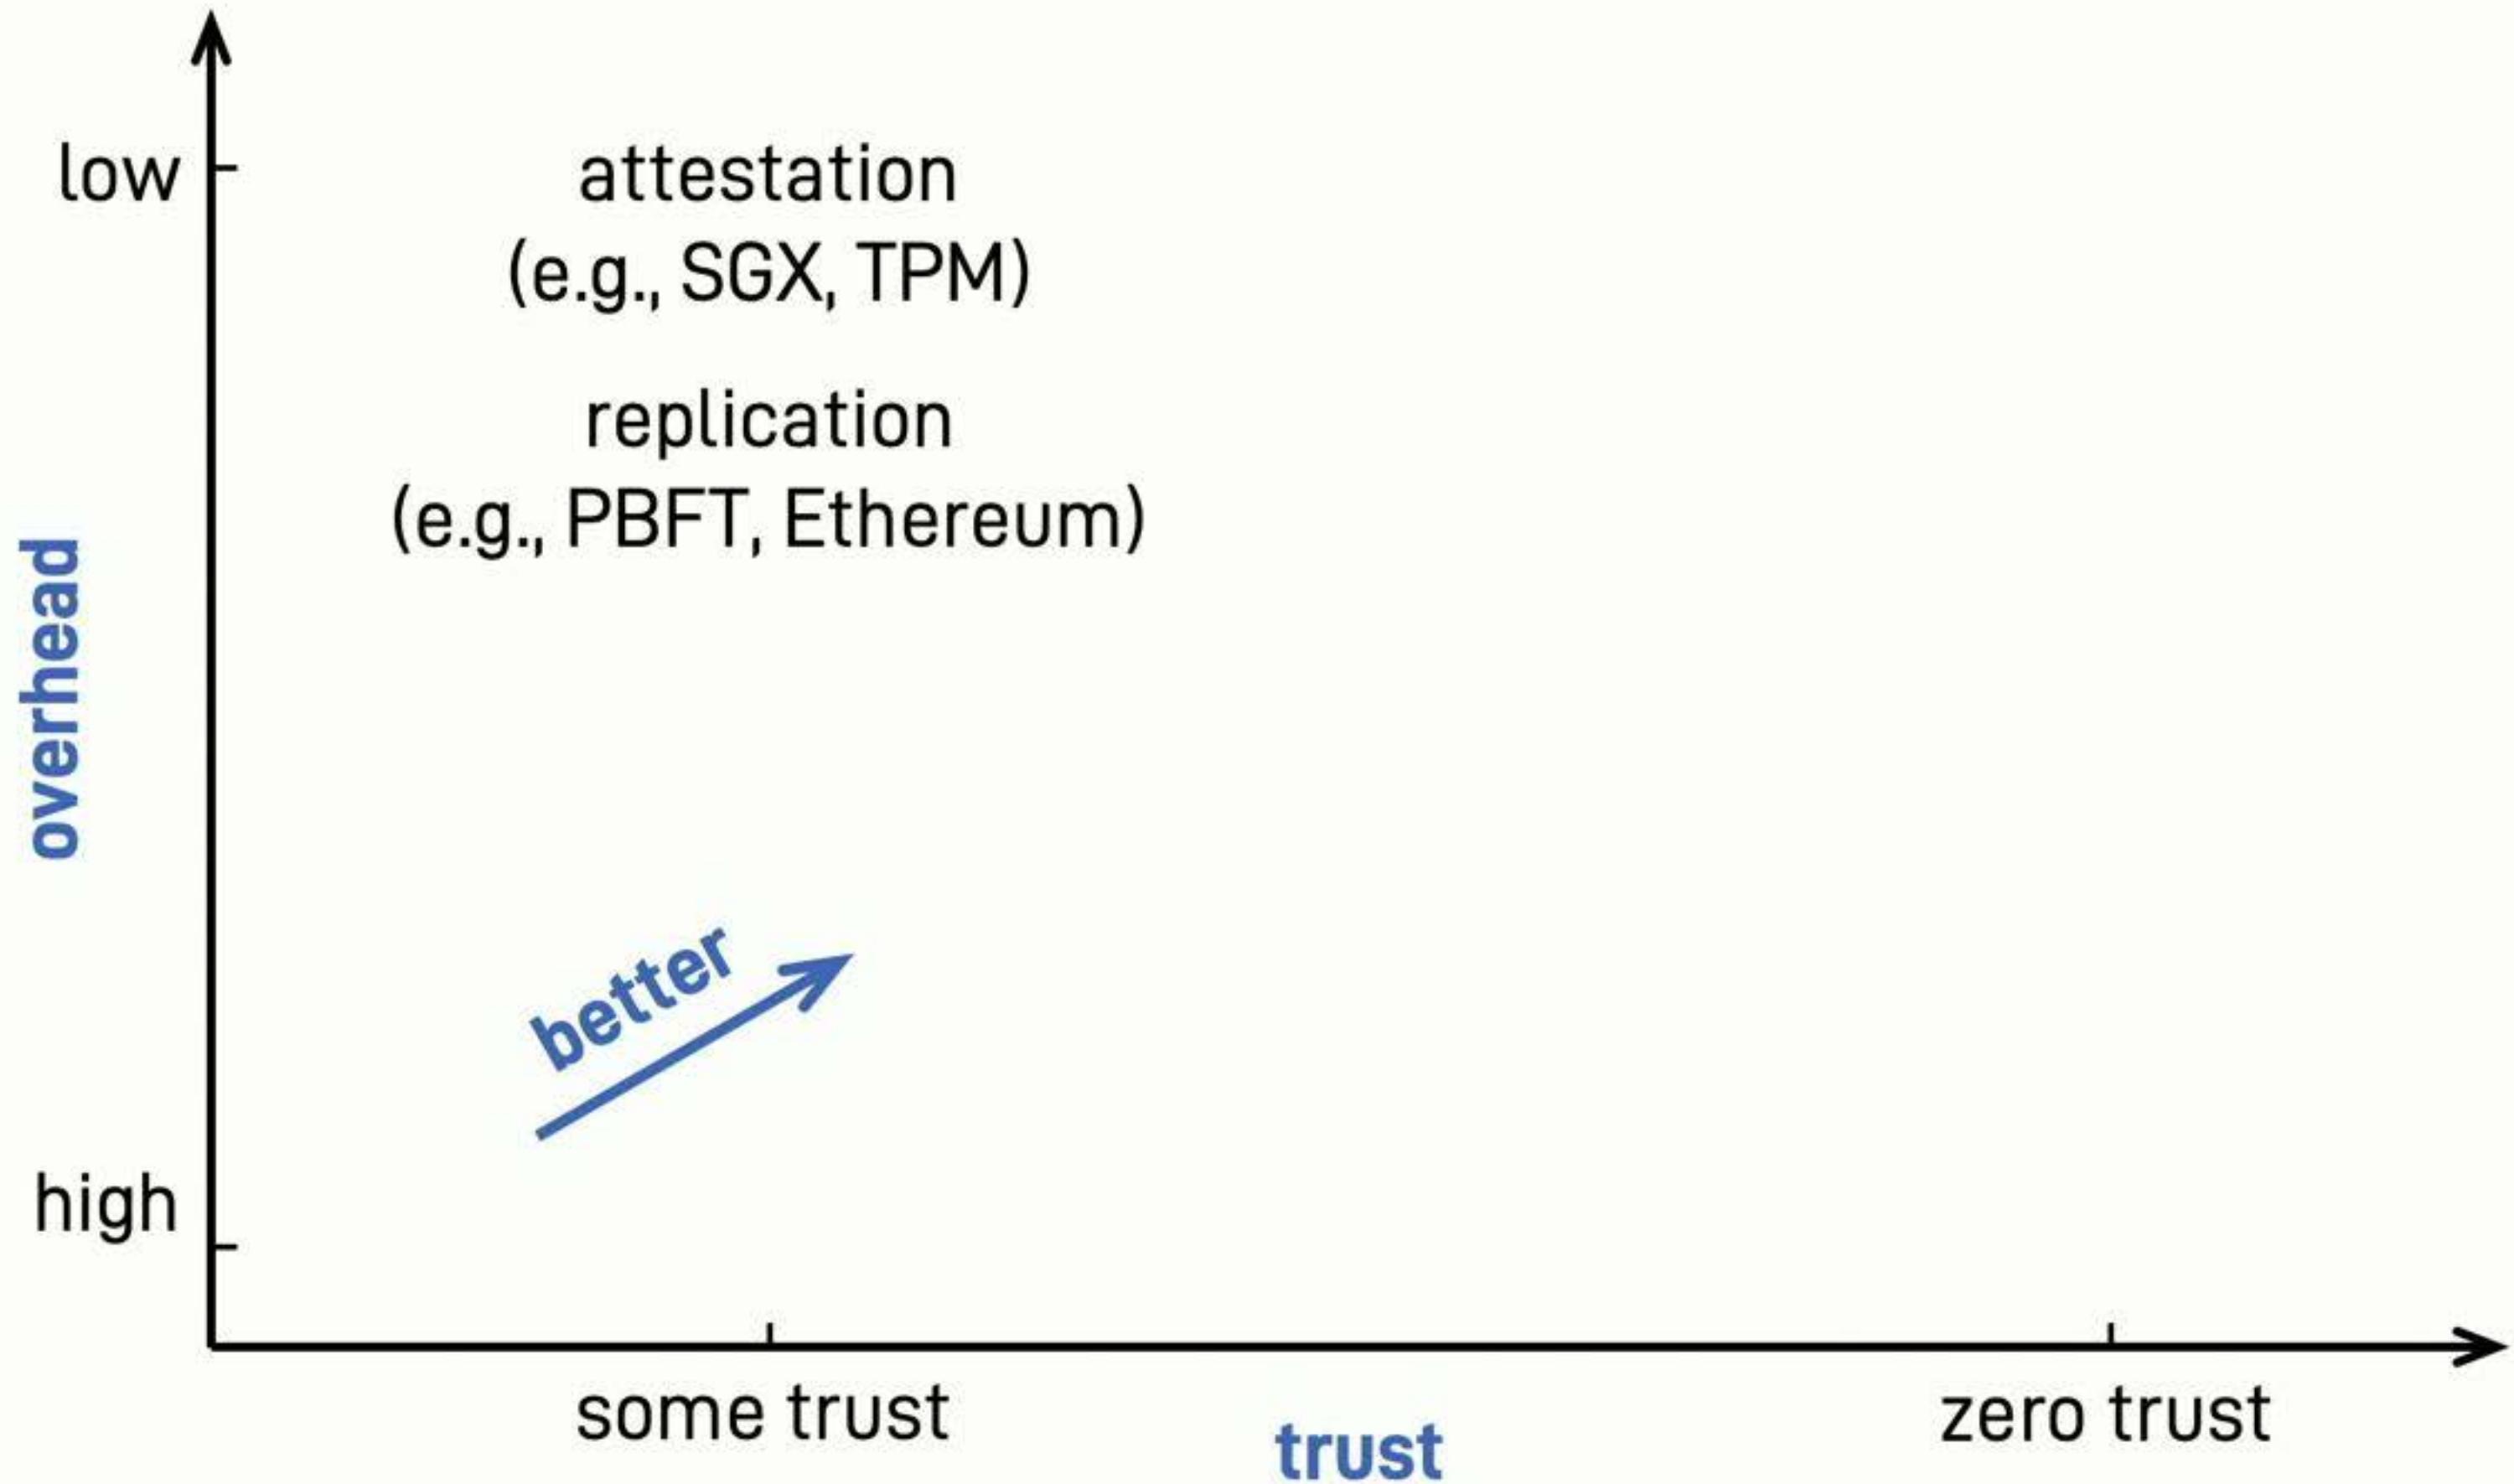
Execution integrity...

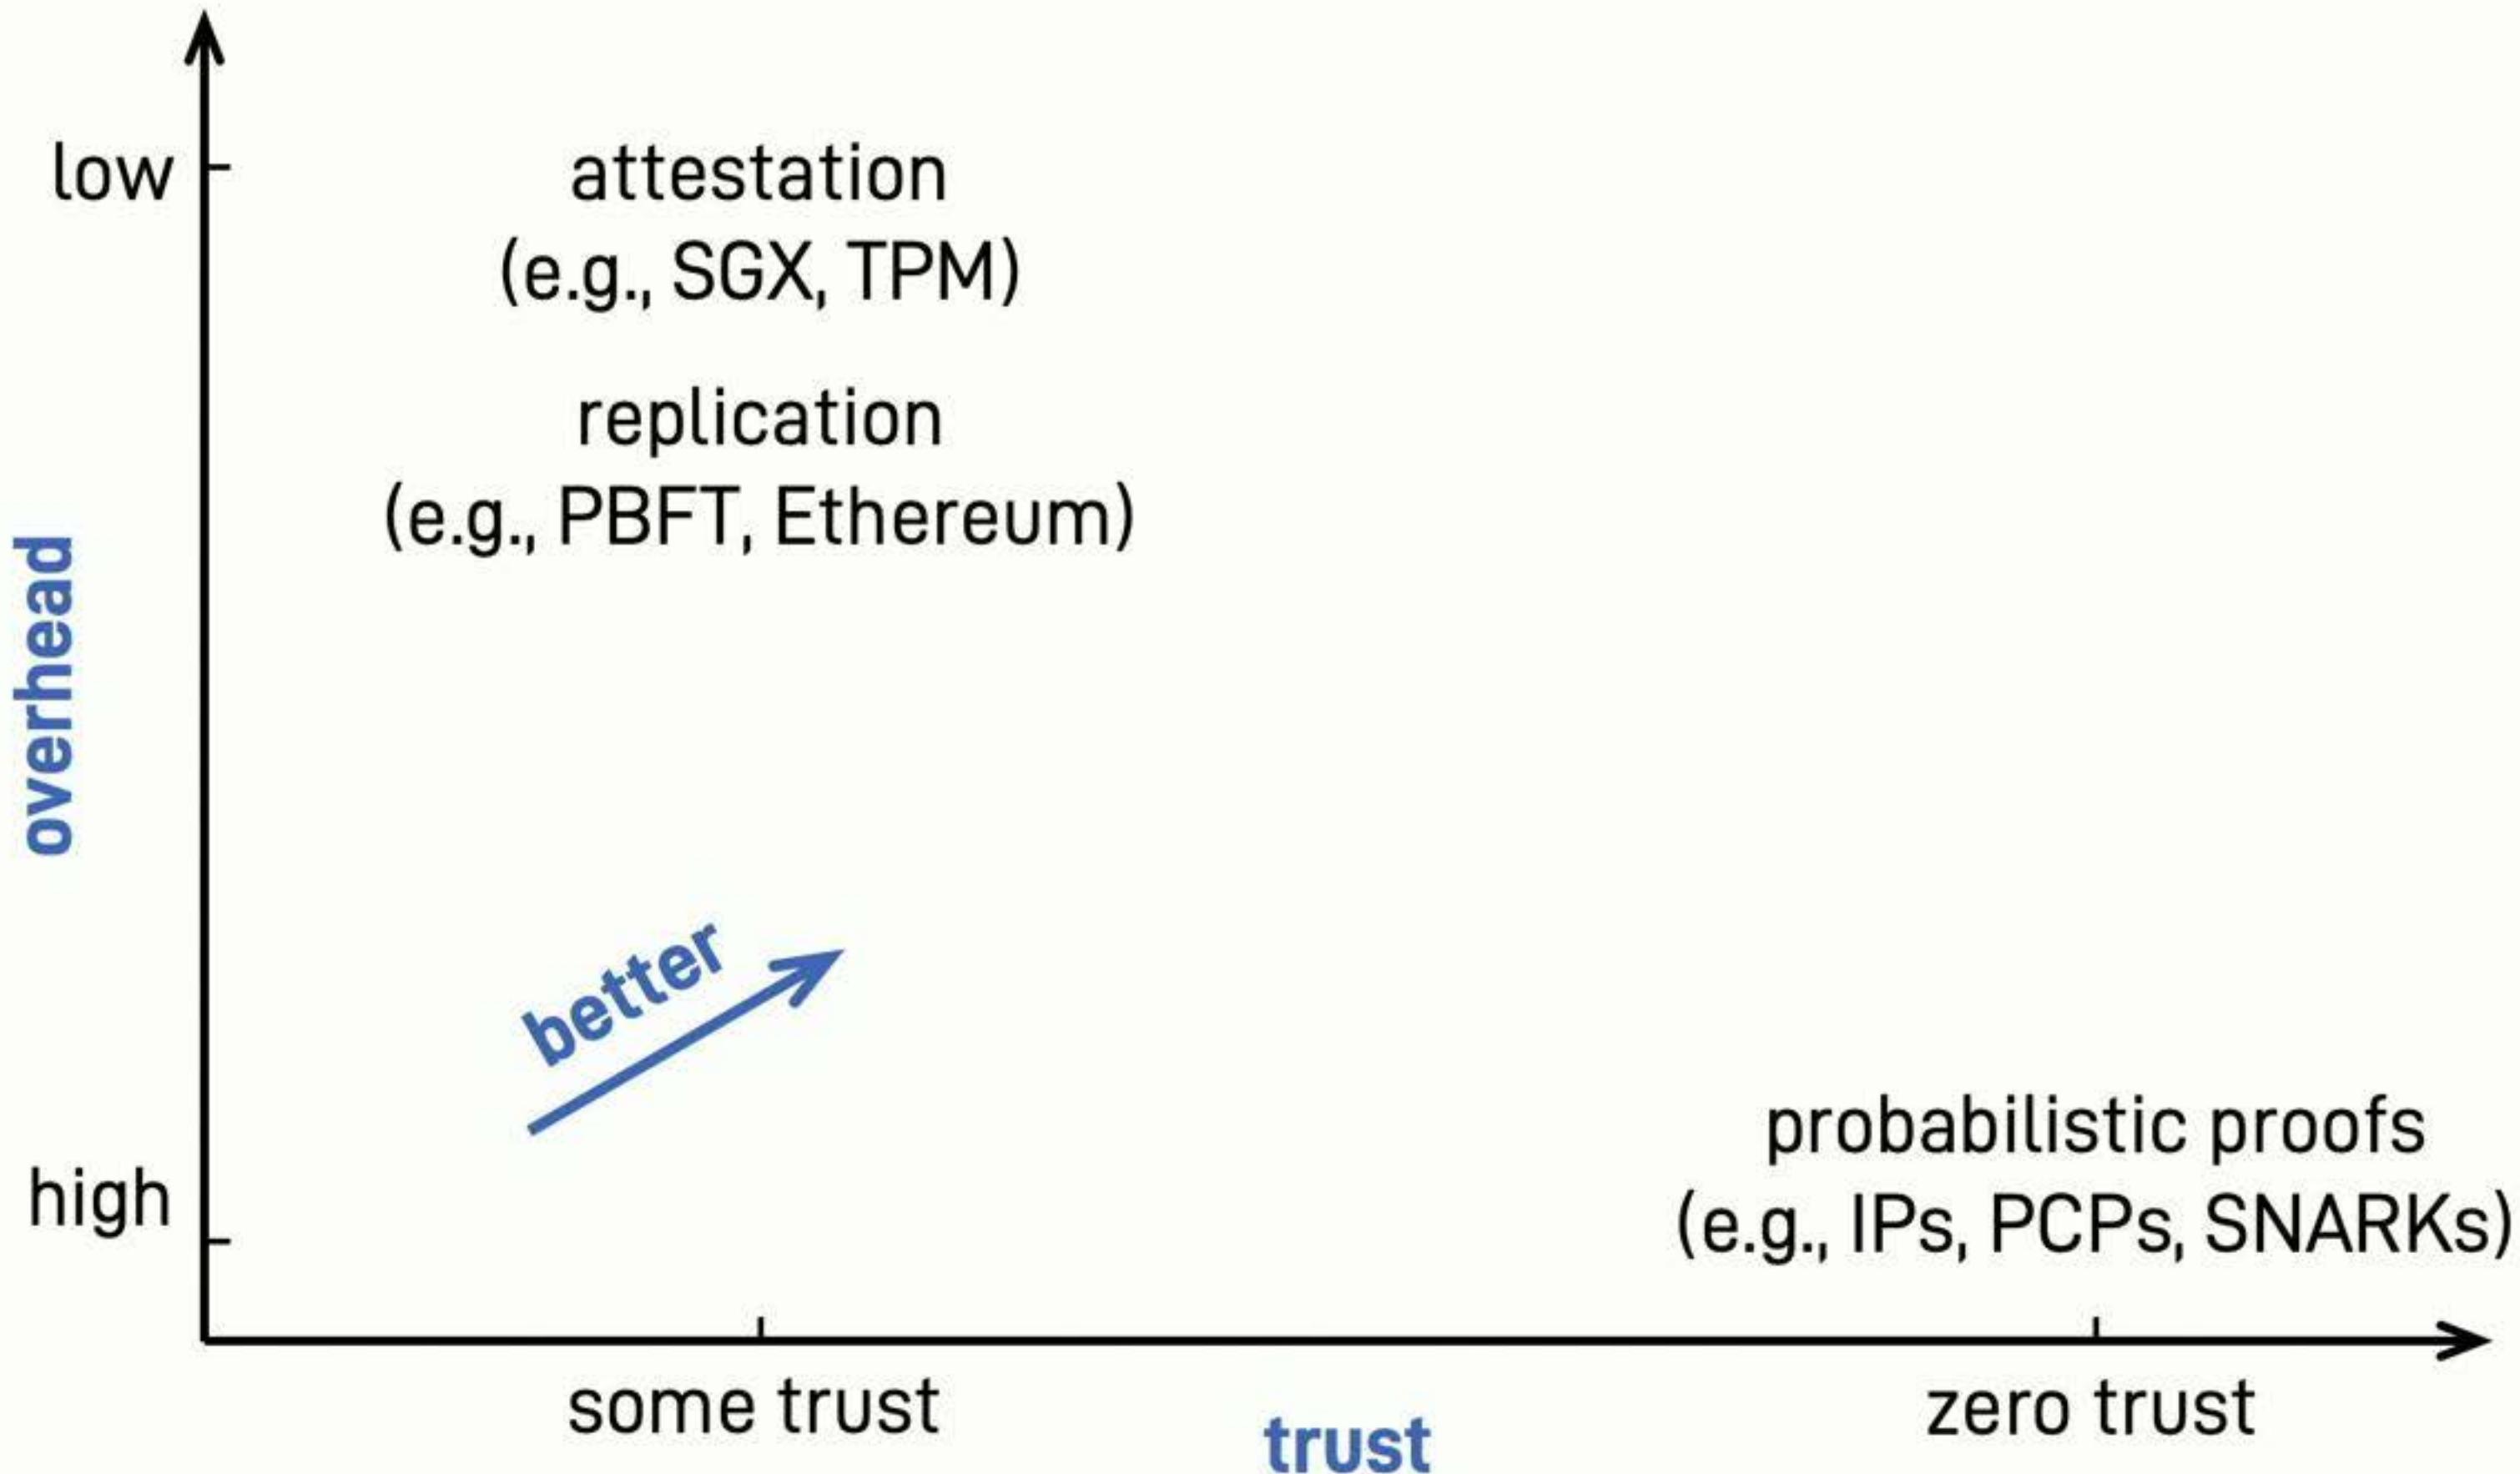
...is about executing code **faithfully**.

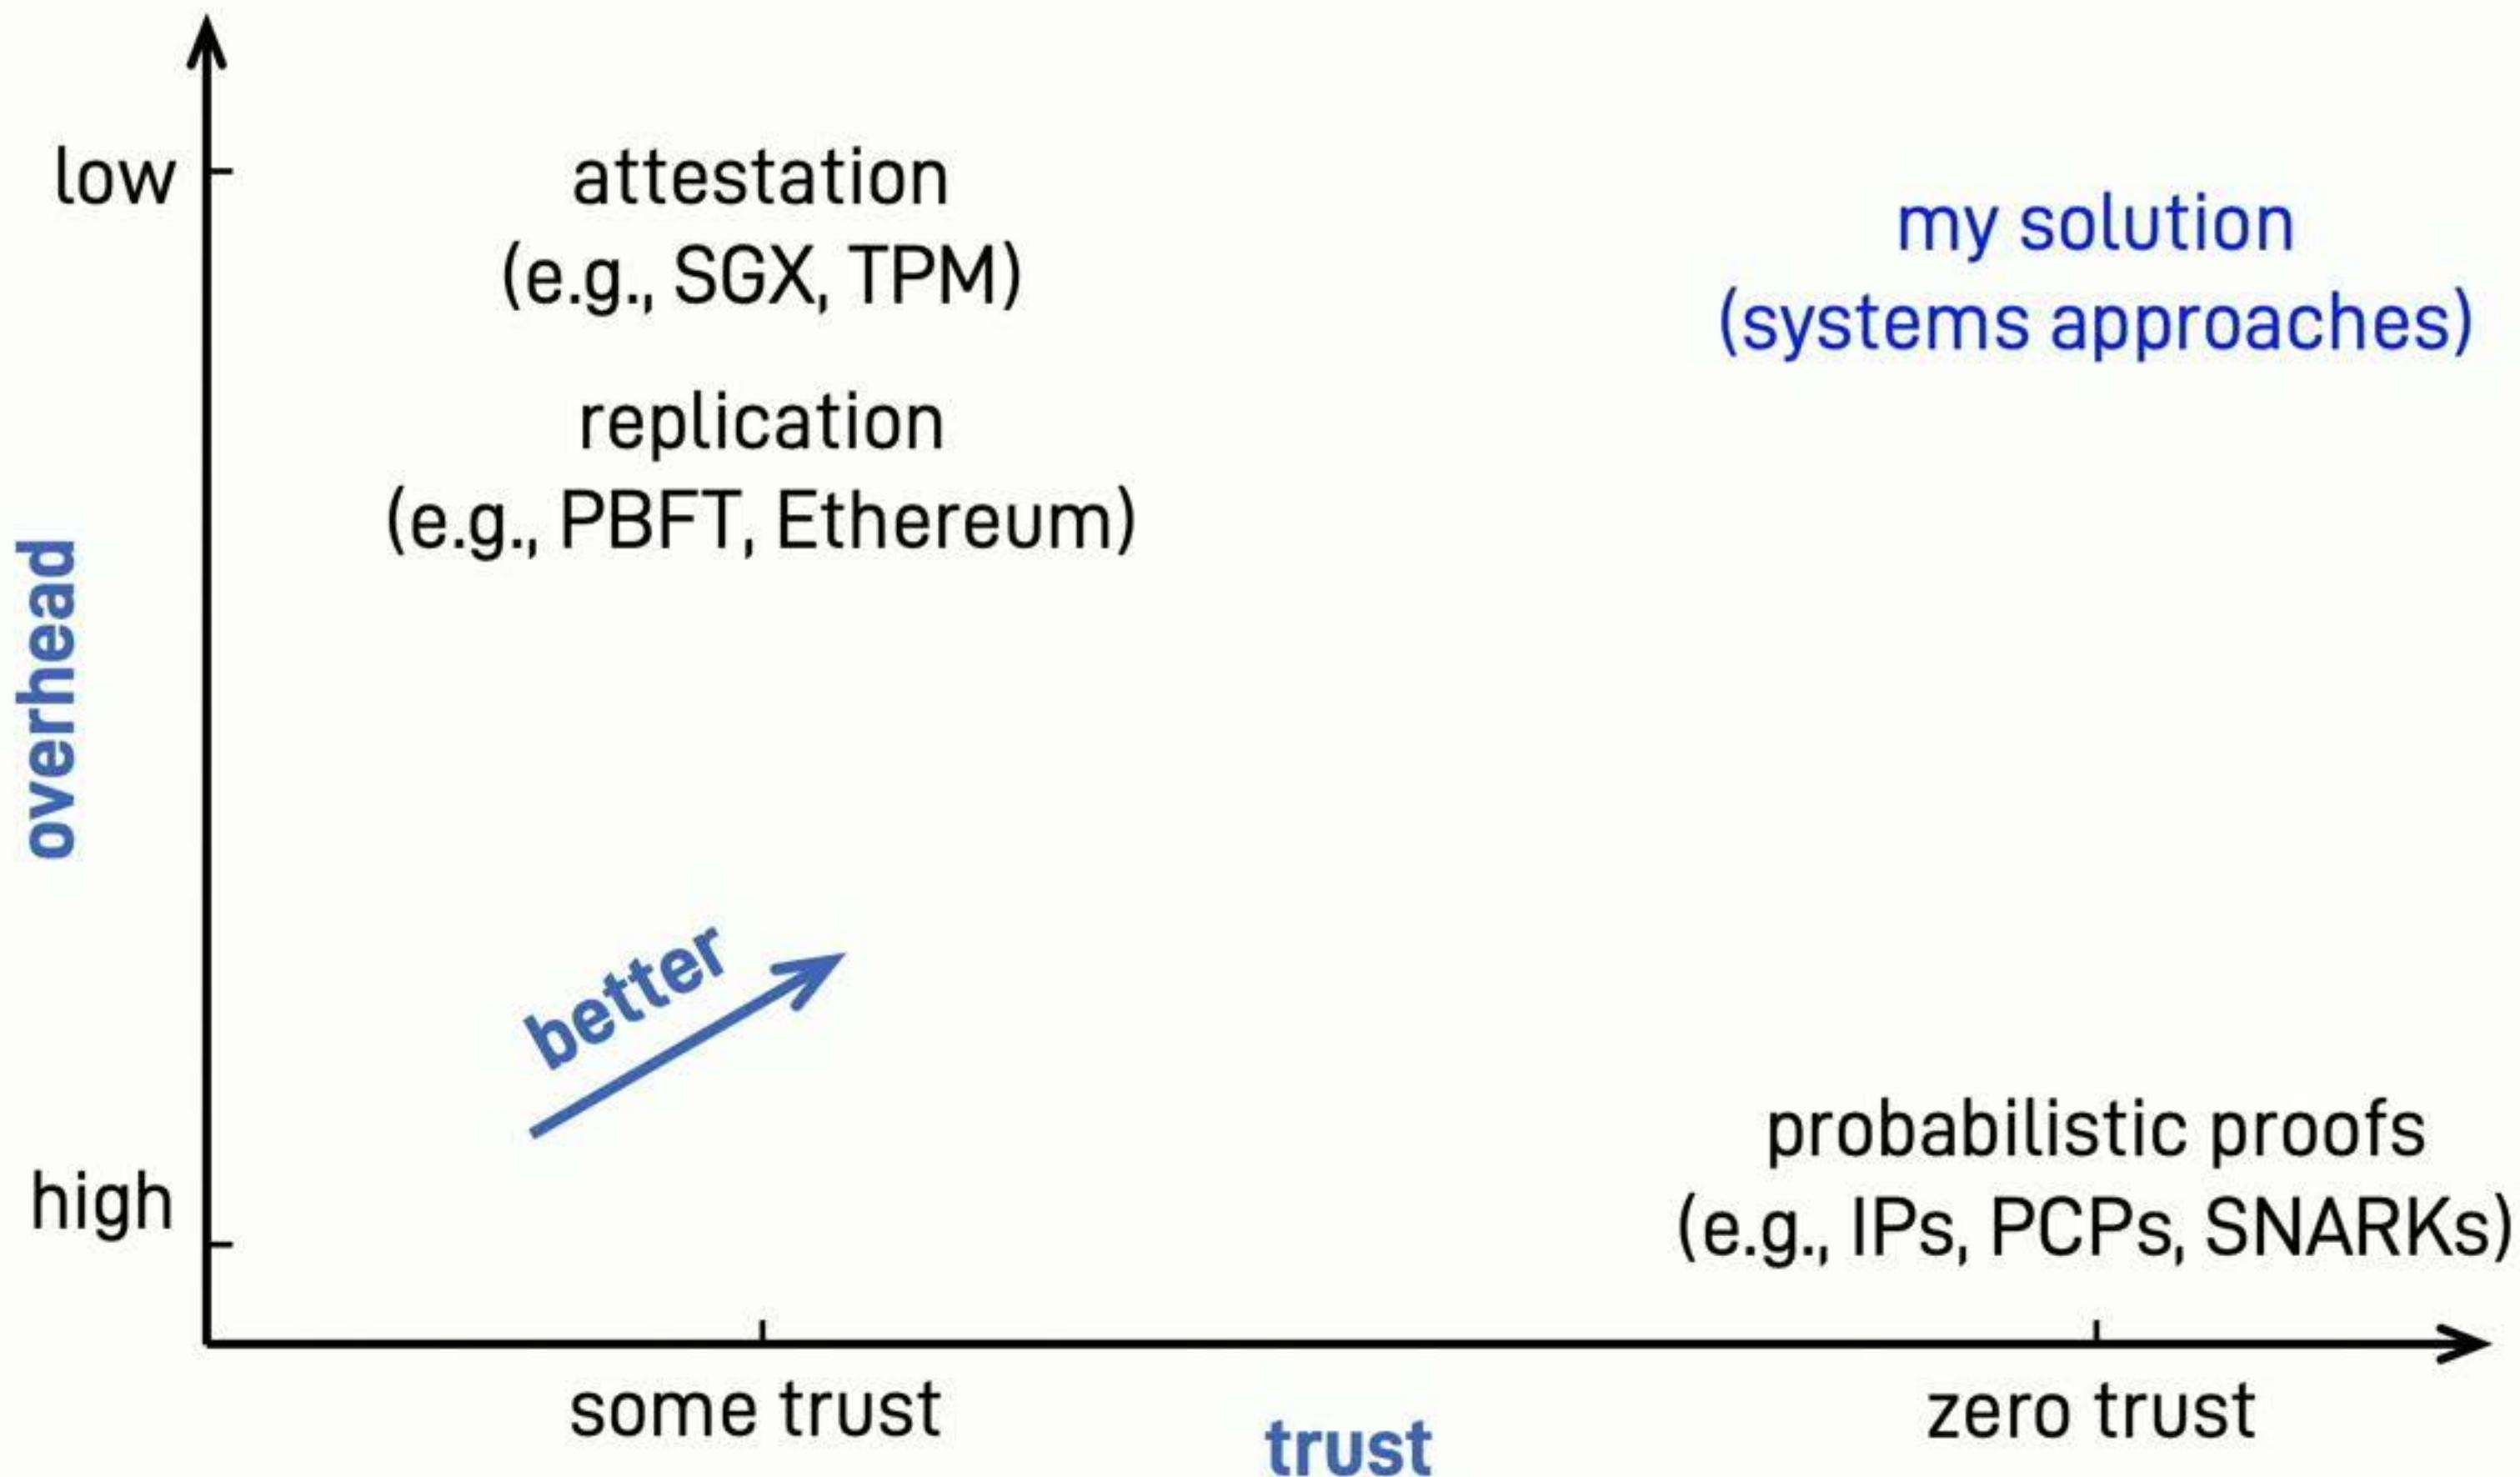
...is separate from program verification.

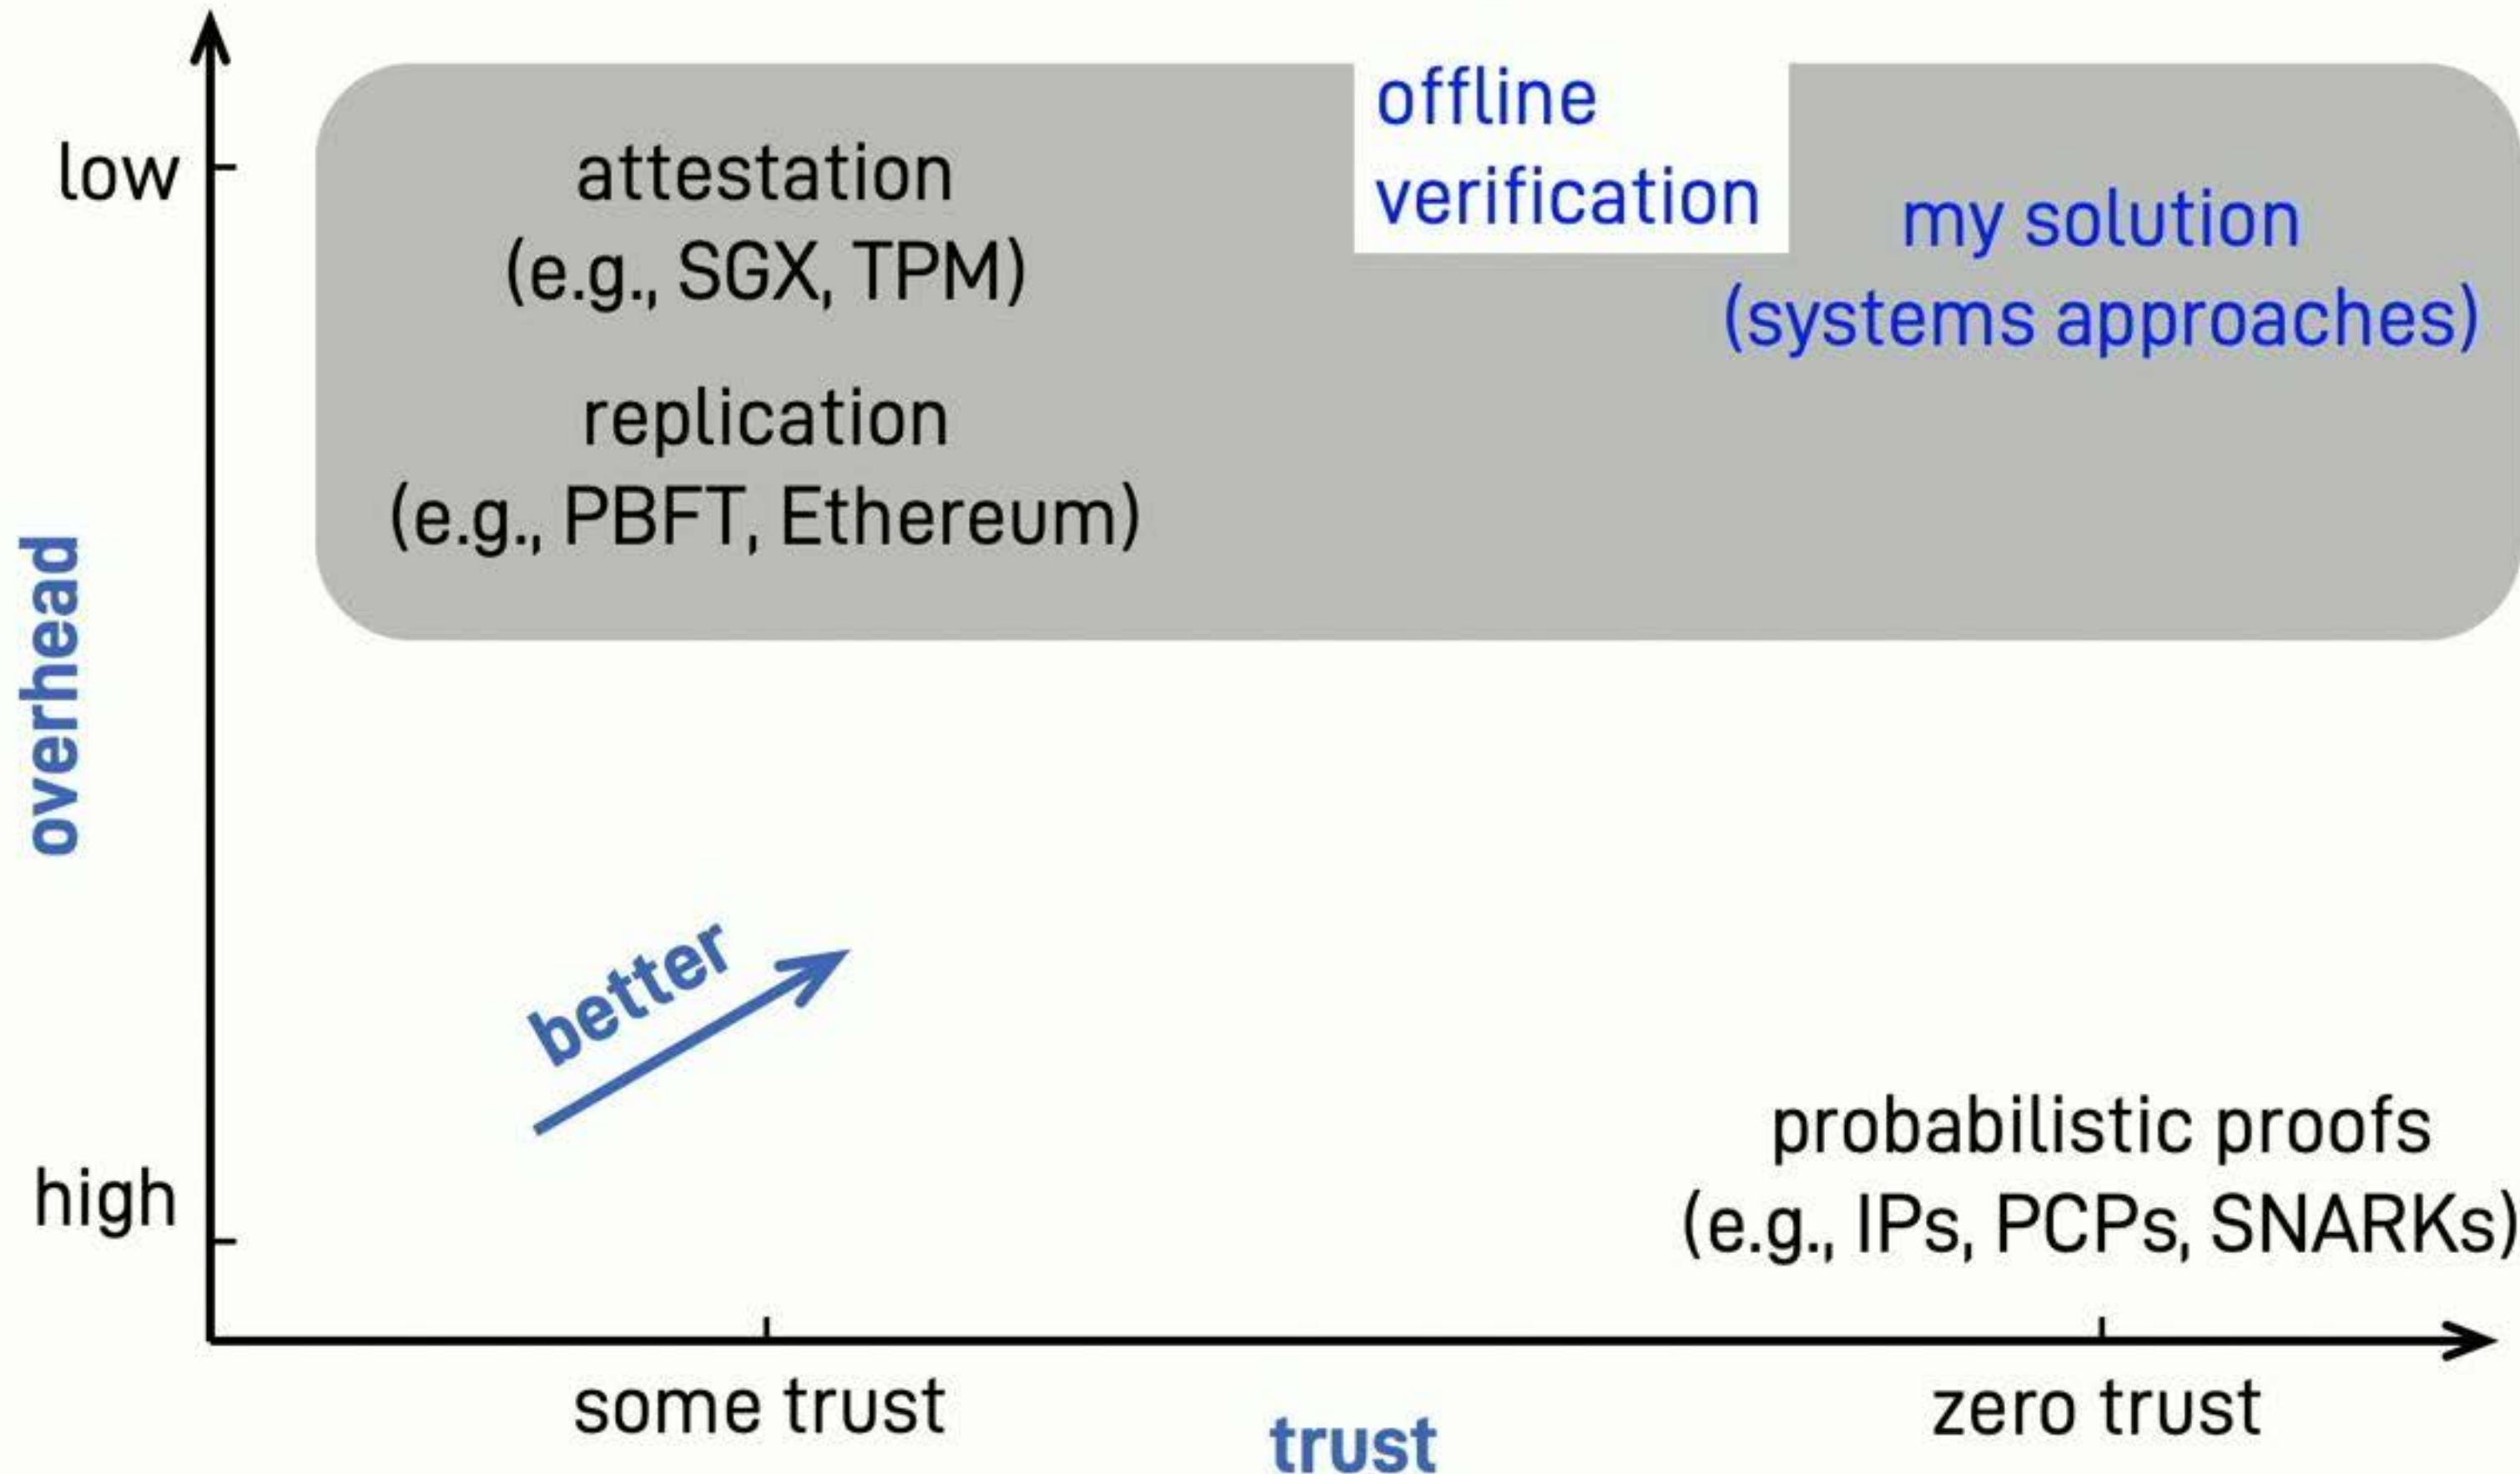


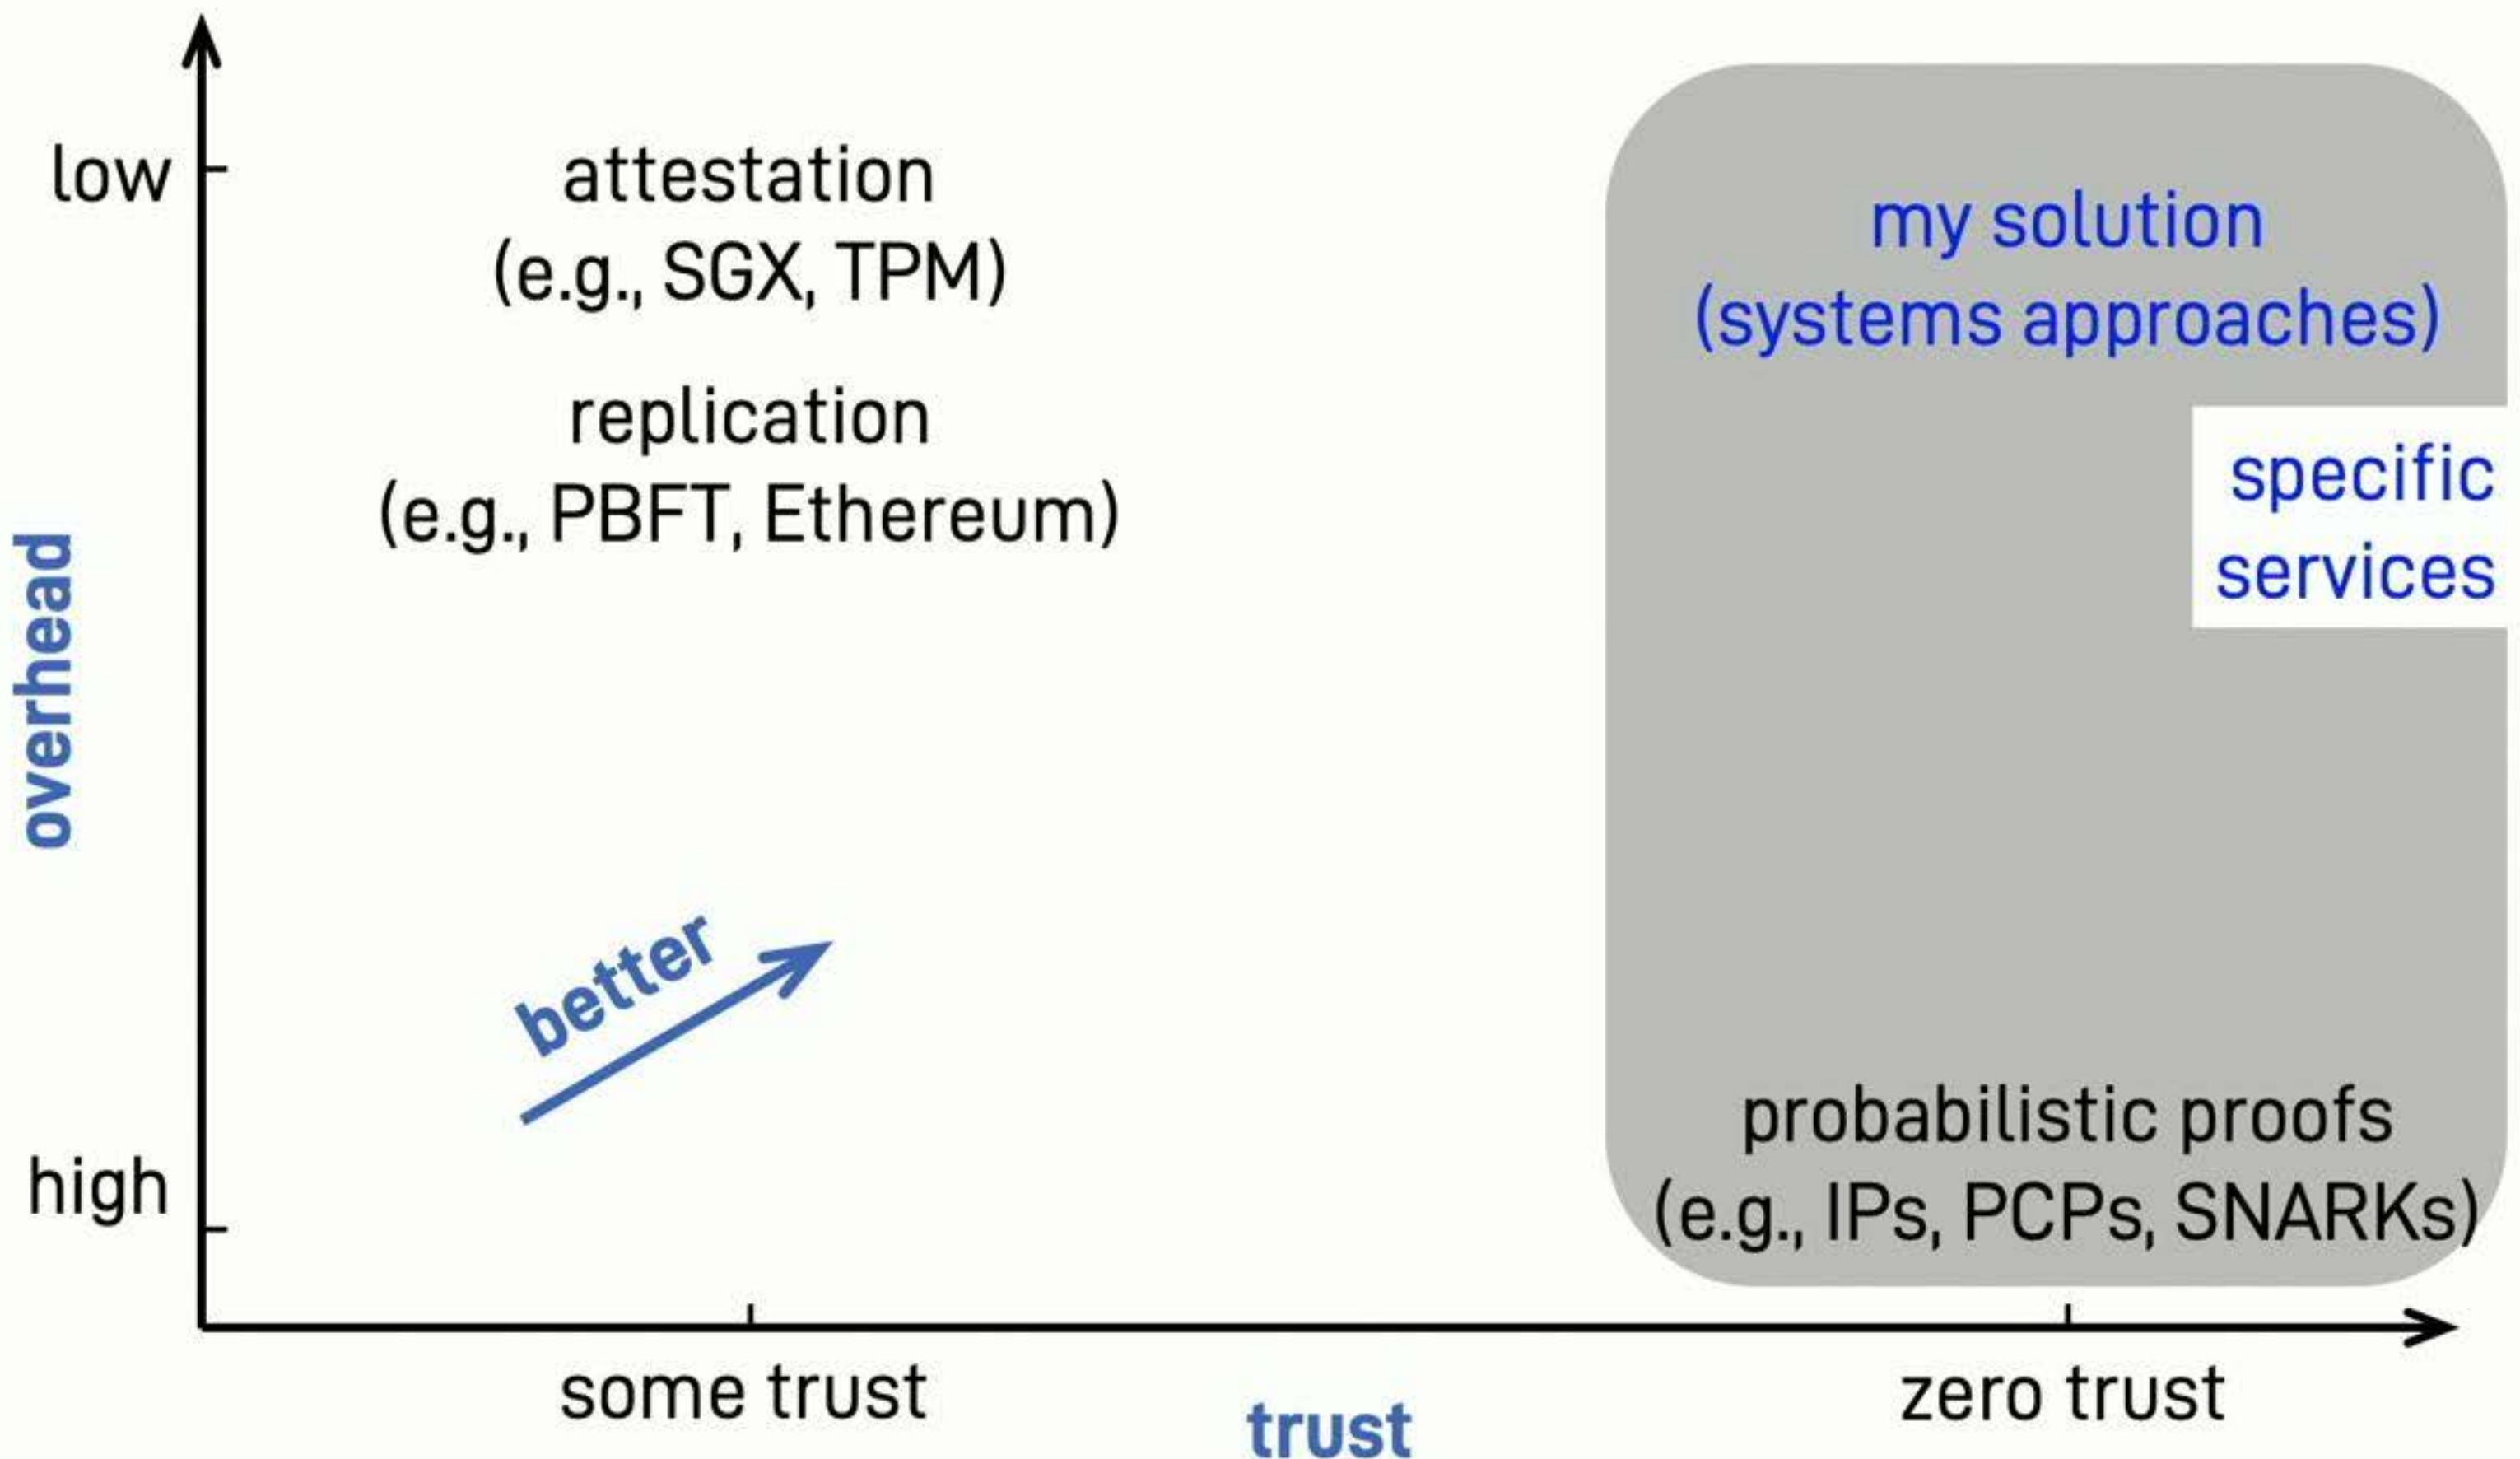


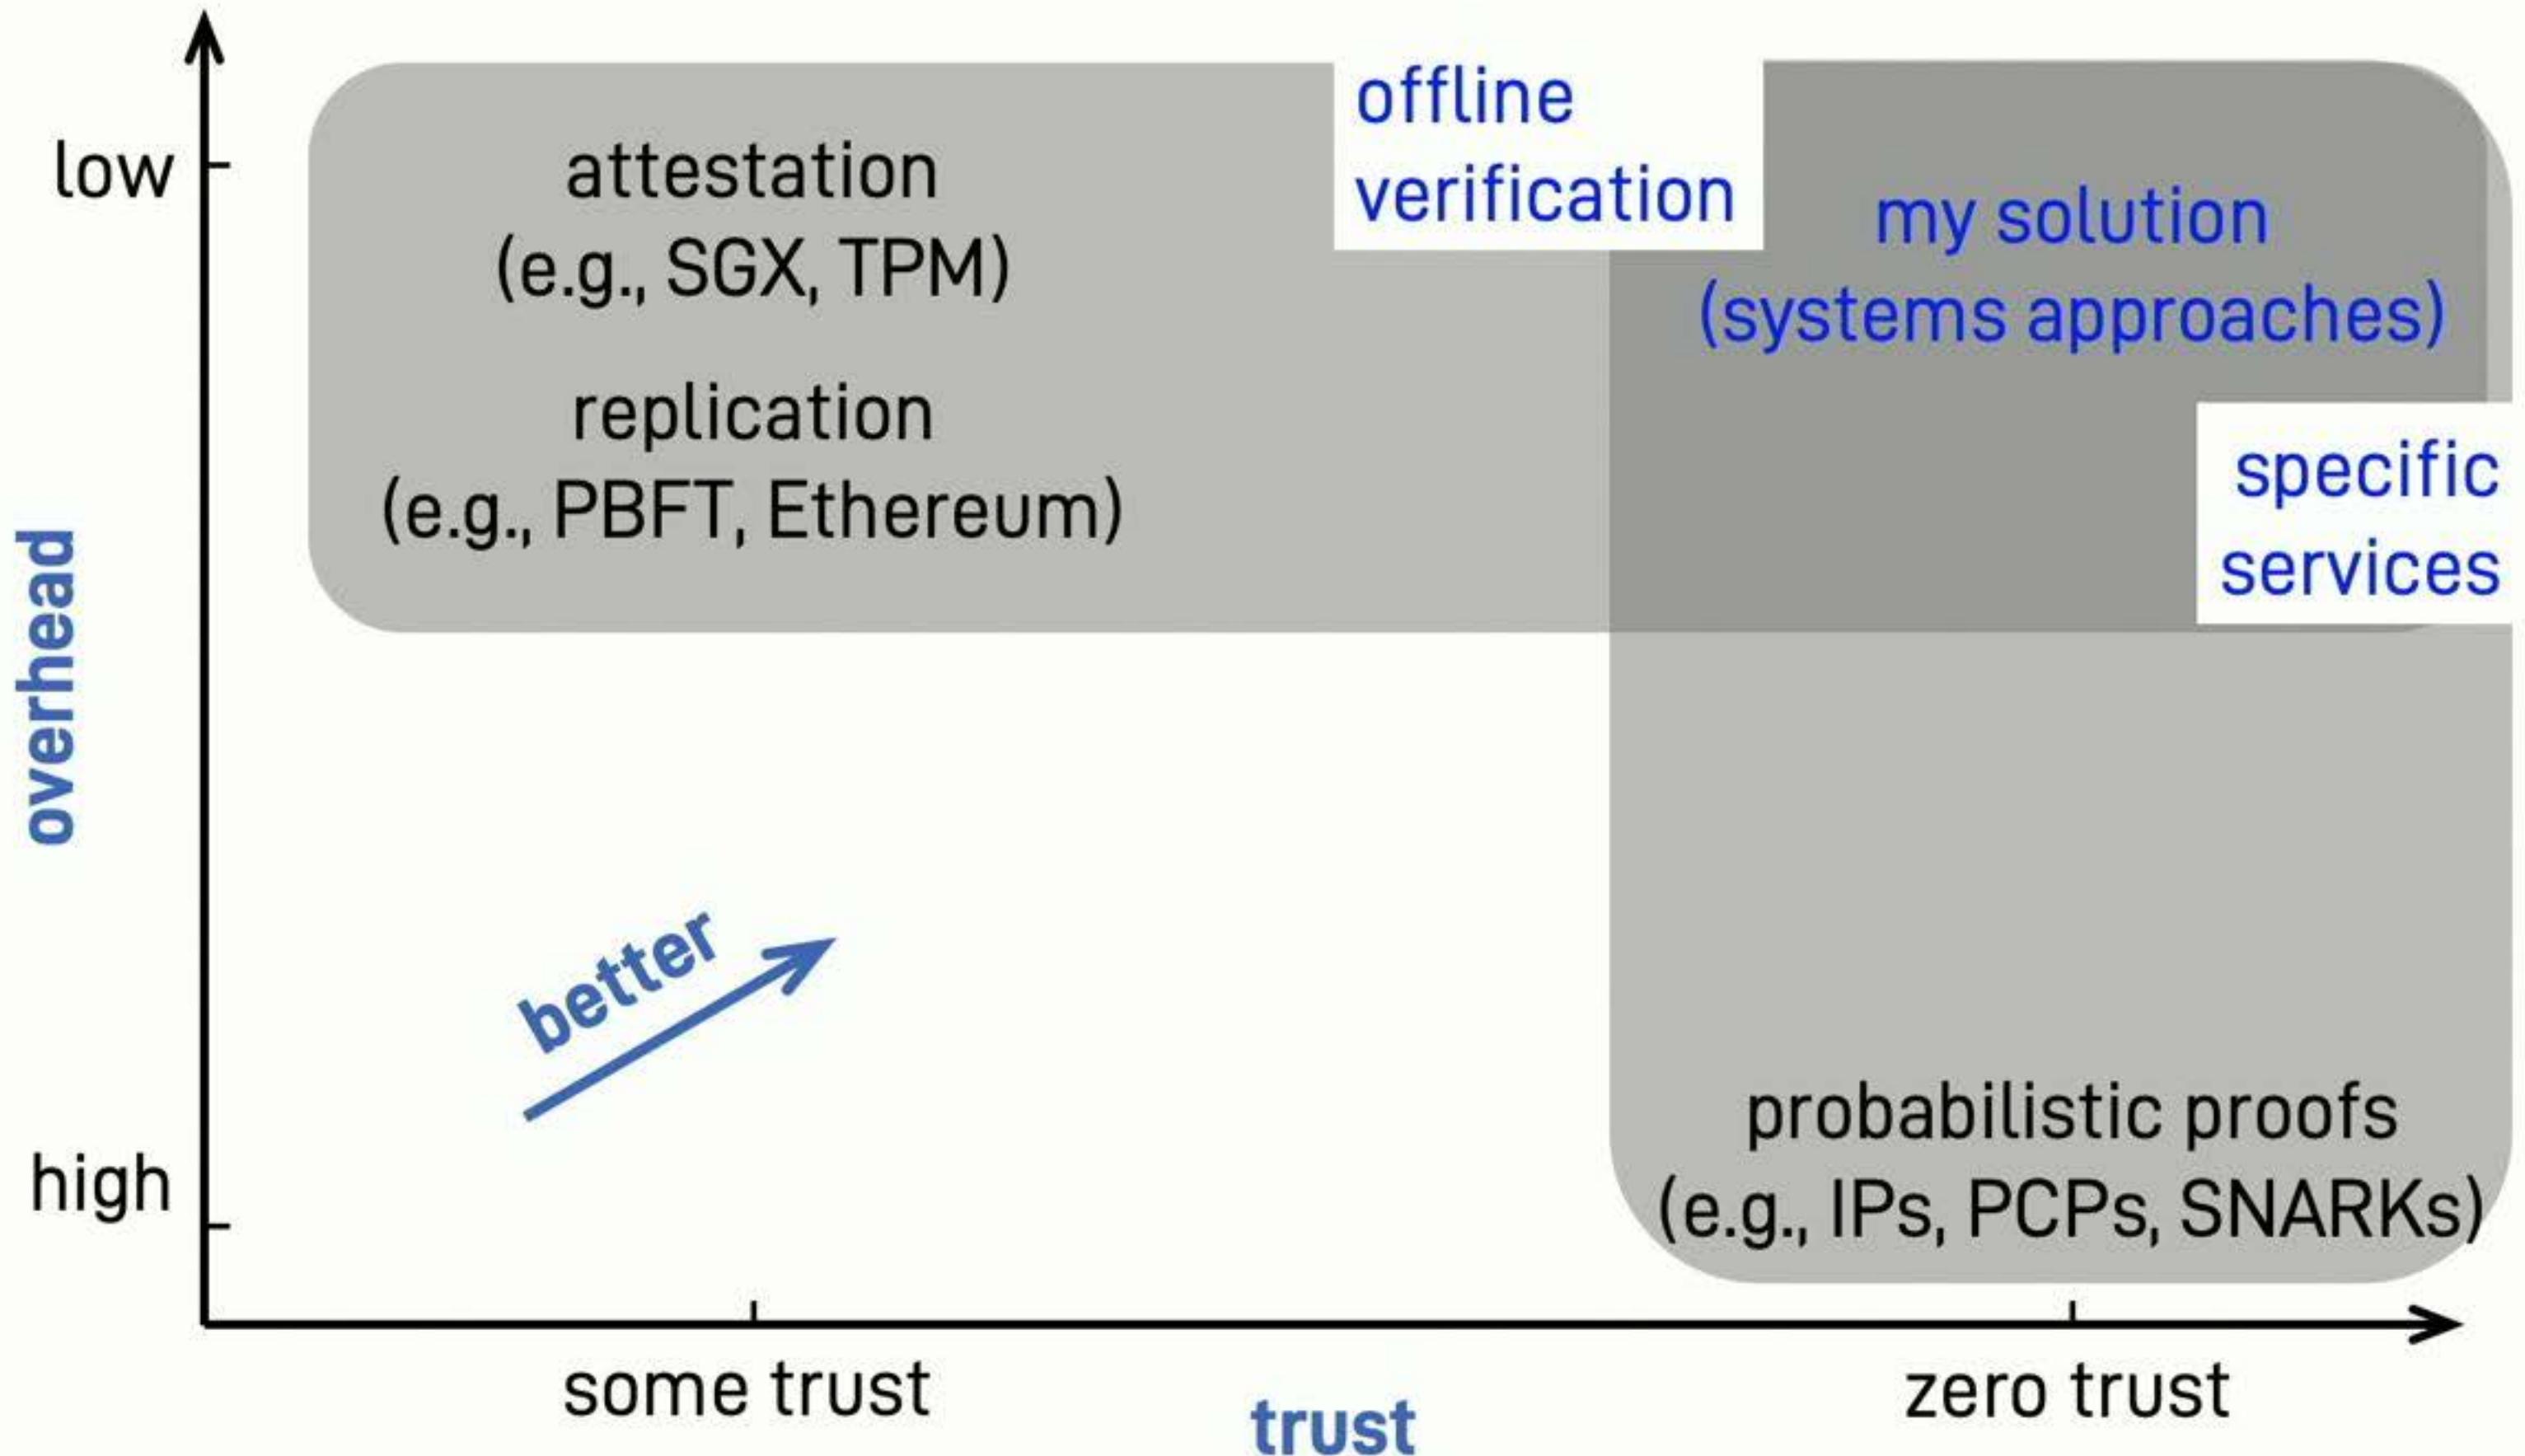




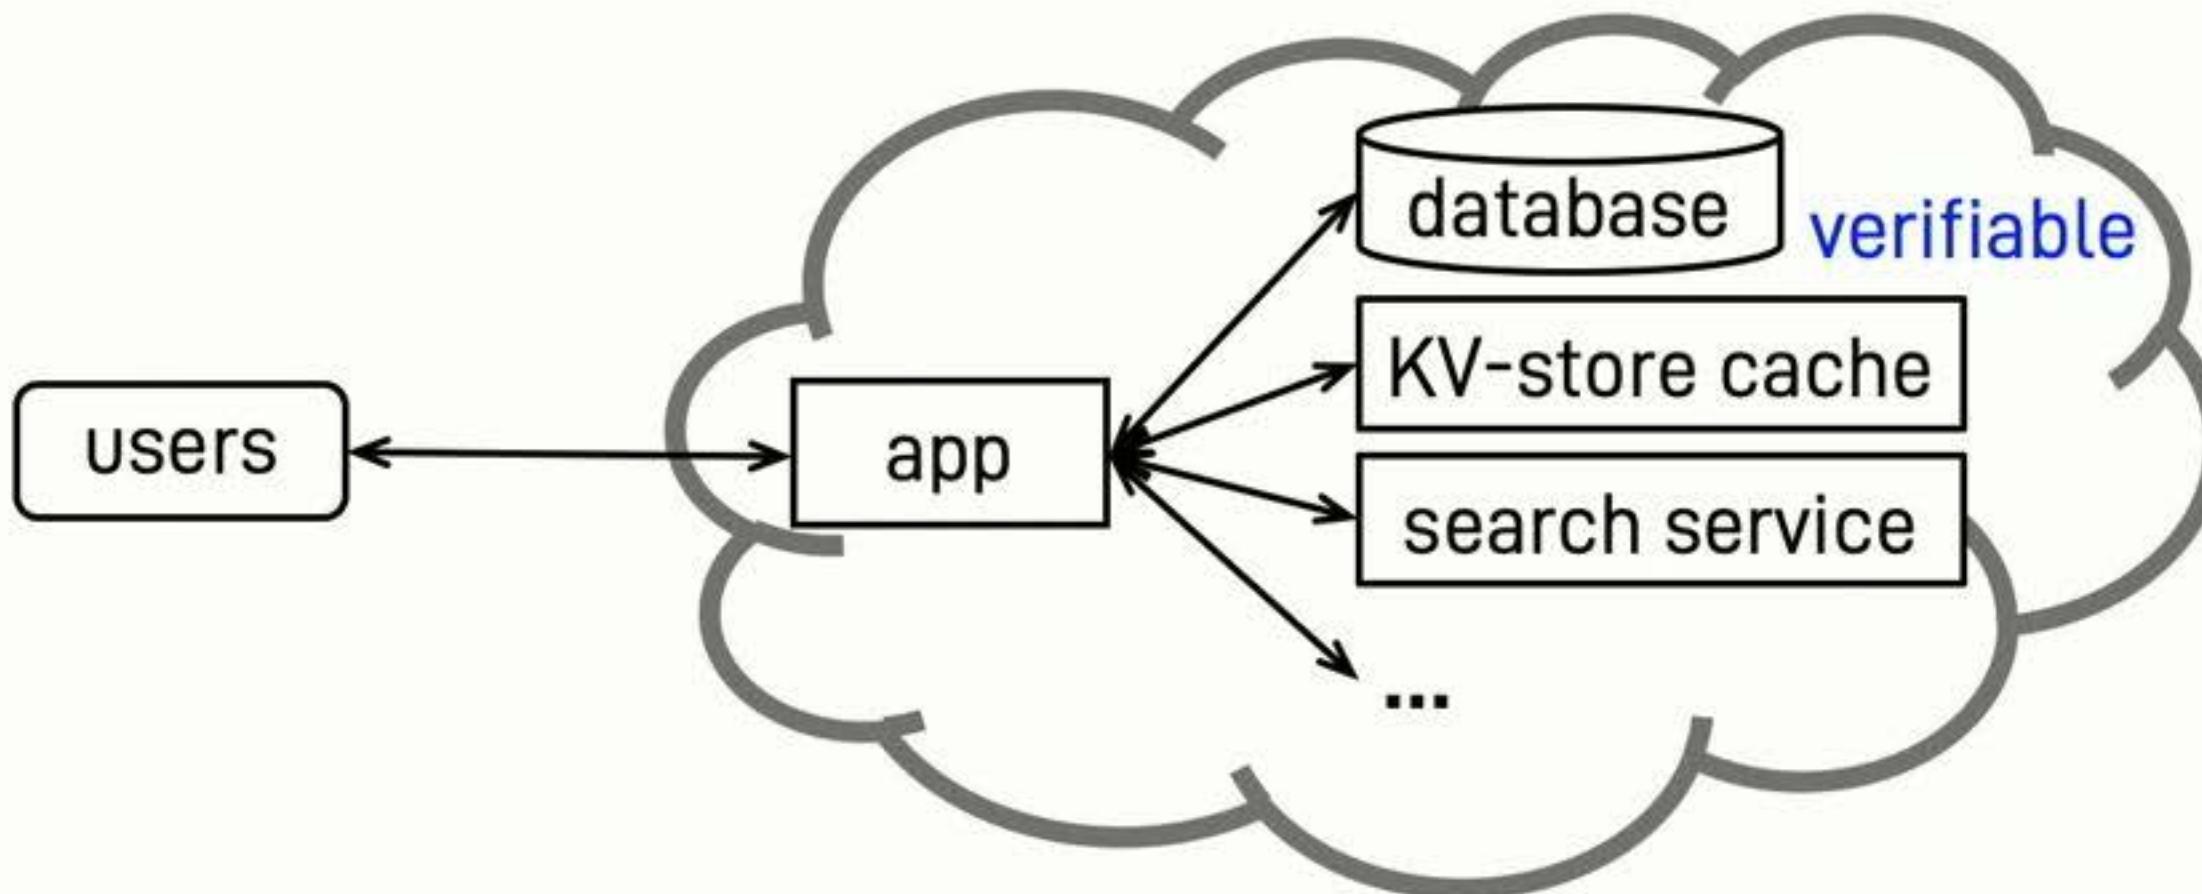




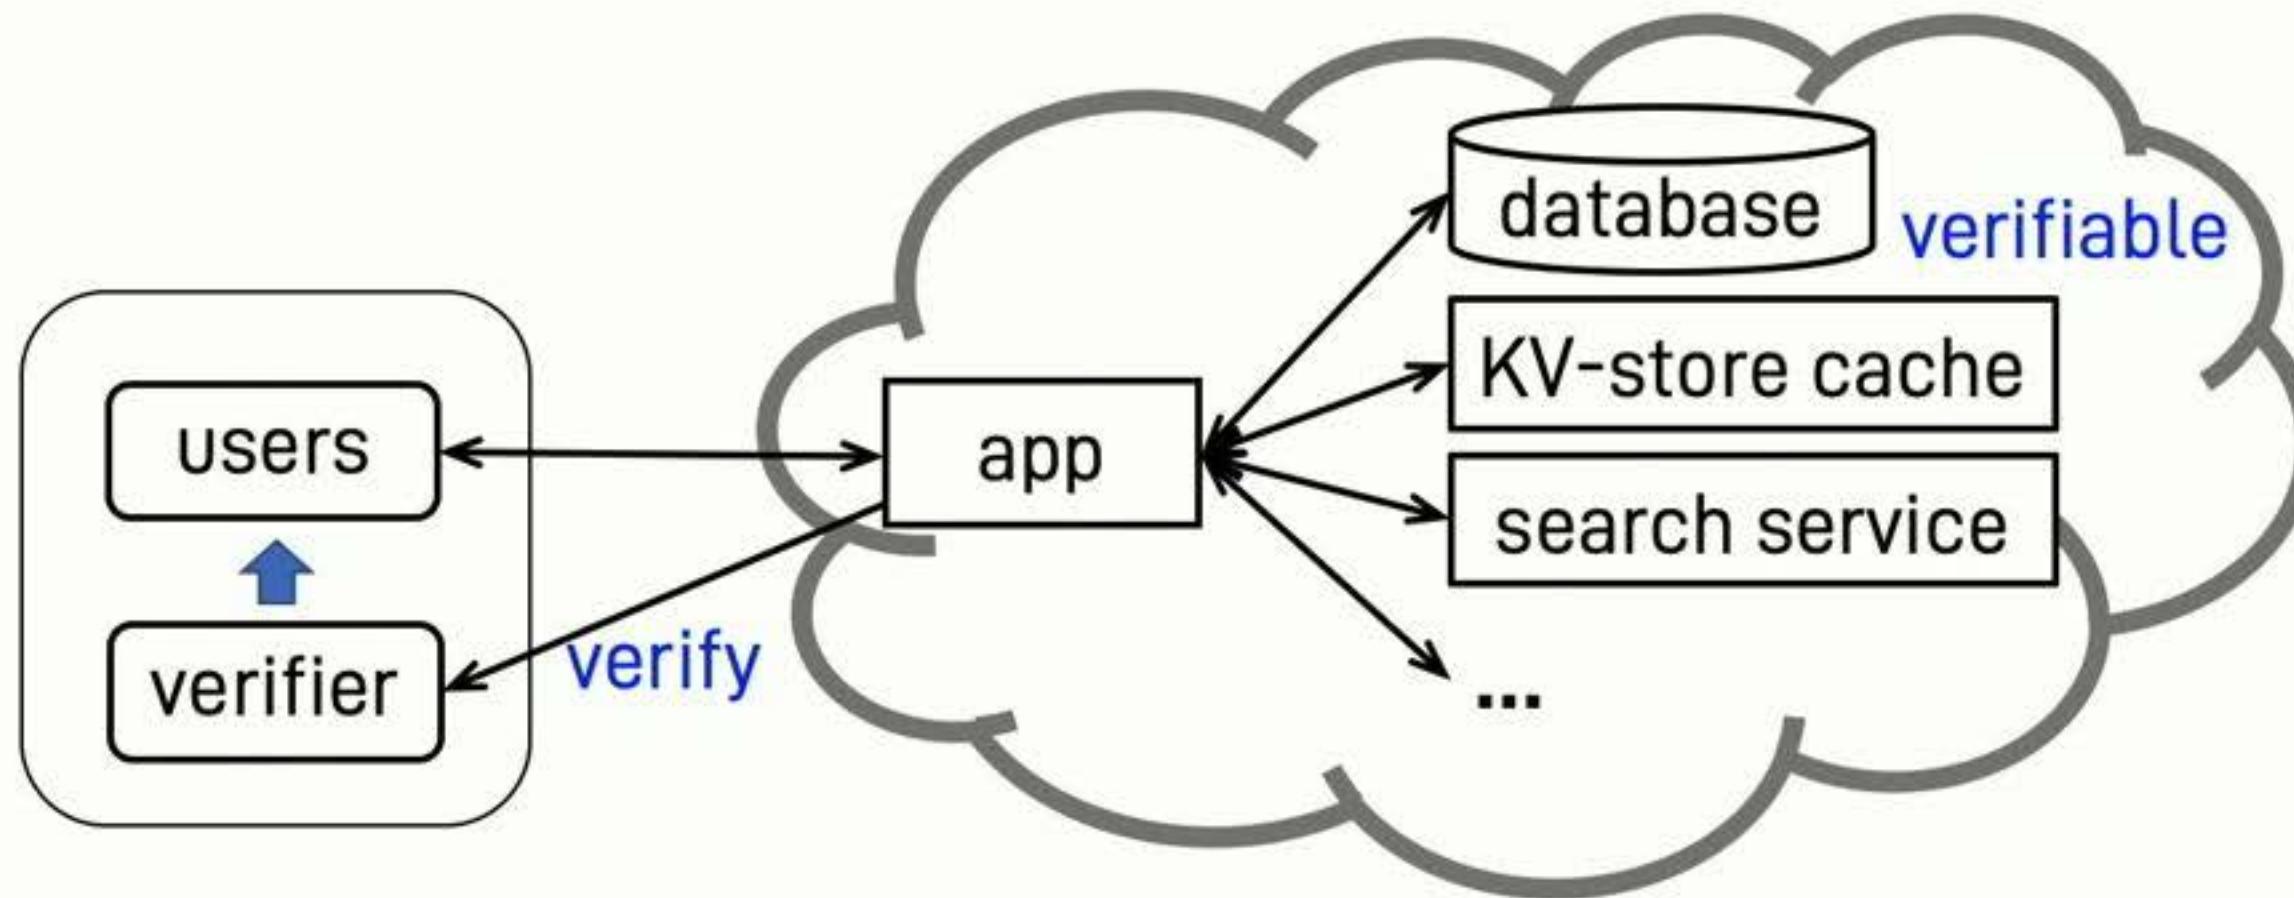




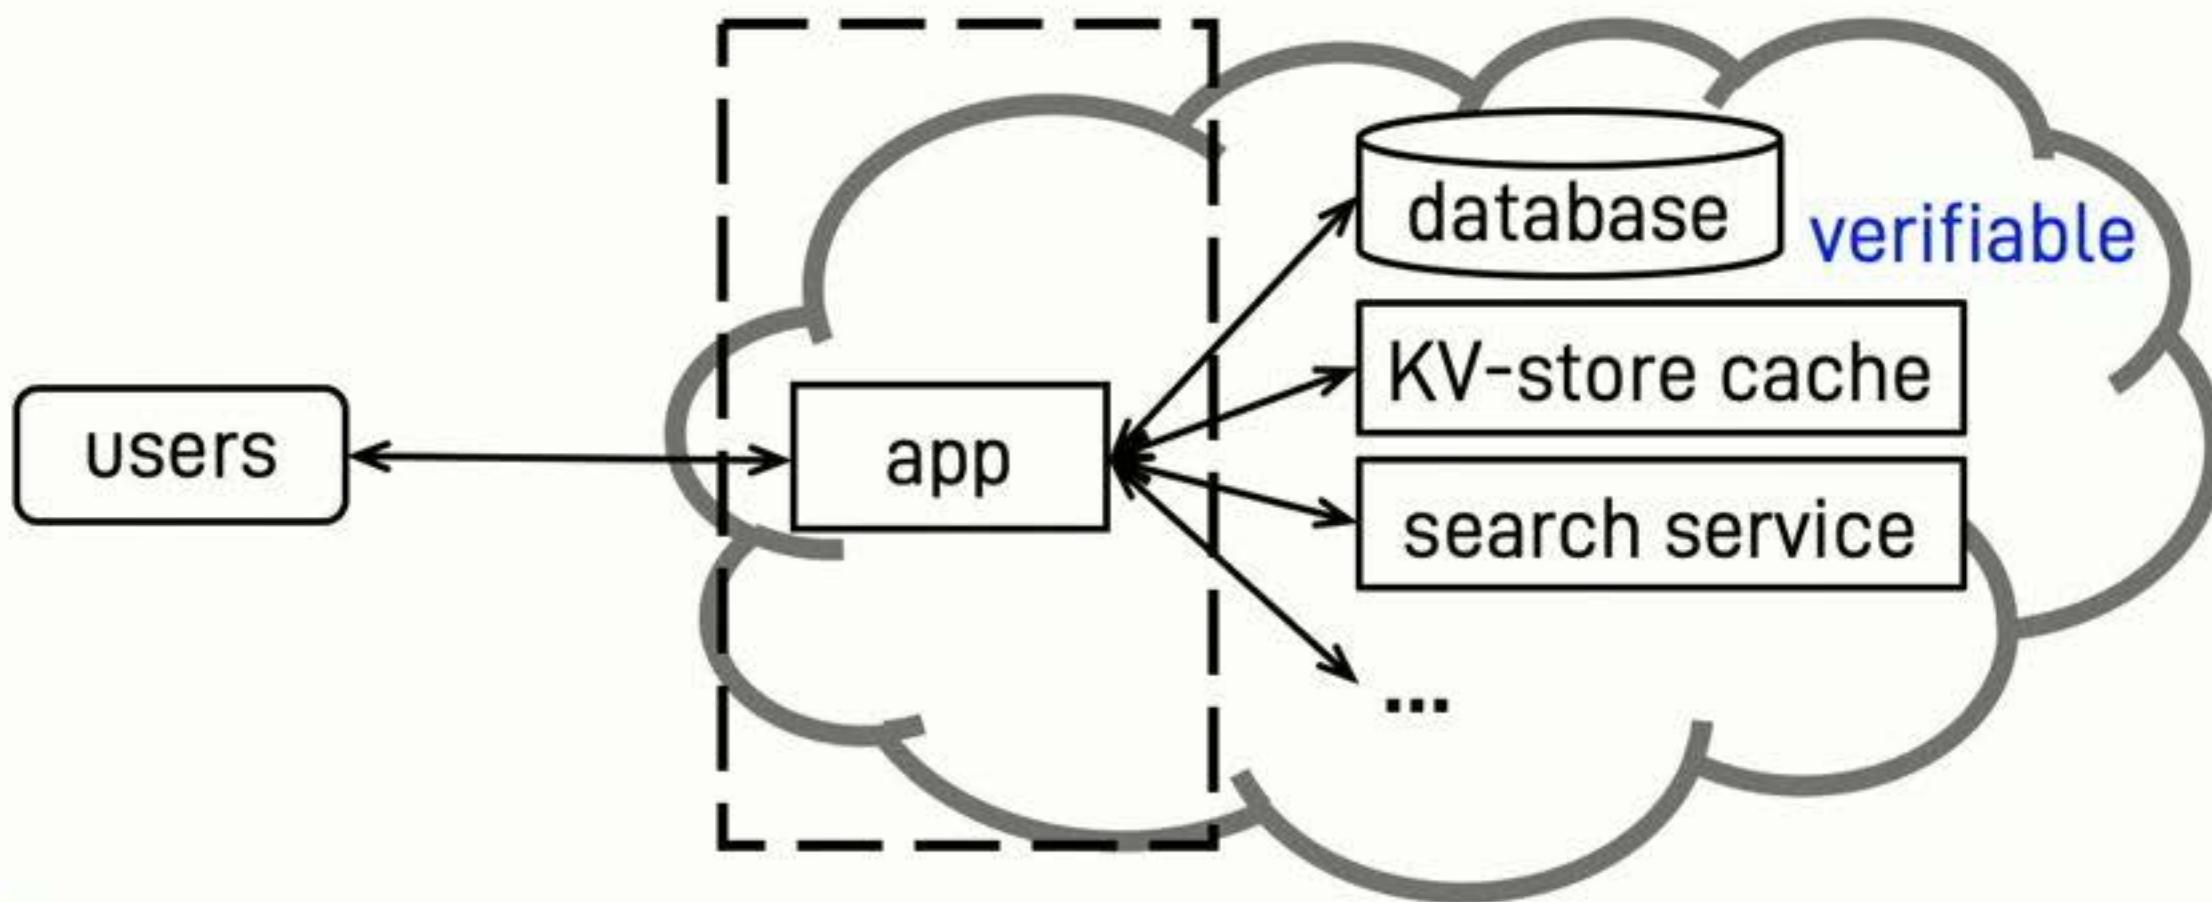
Rest of the talk: verifiable infrastructure



Rest of the talk: verifiable infrastructure



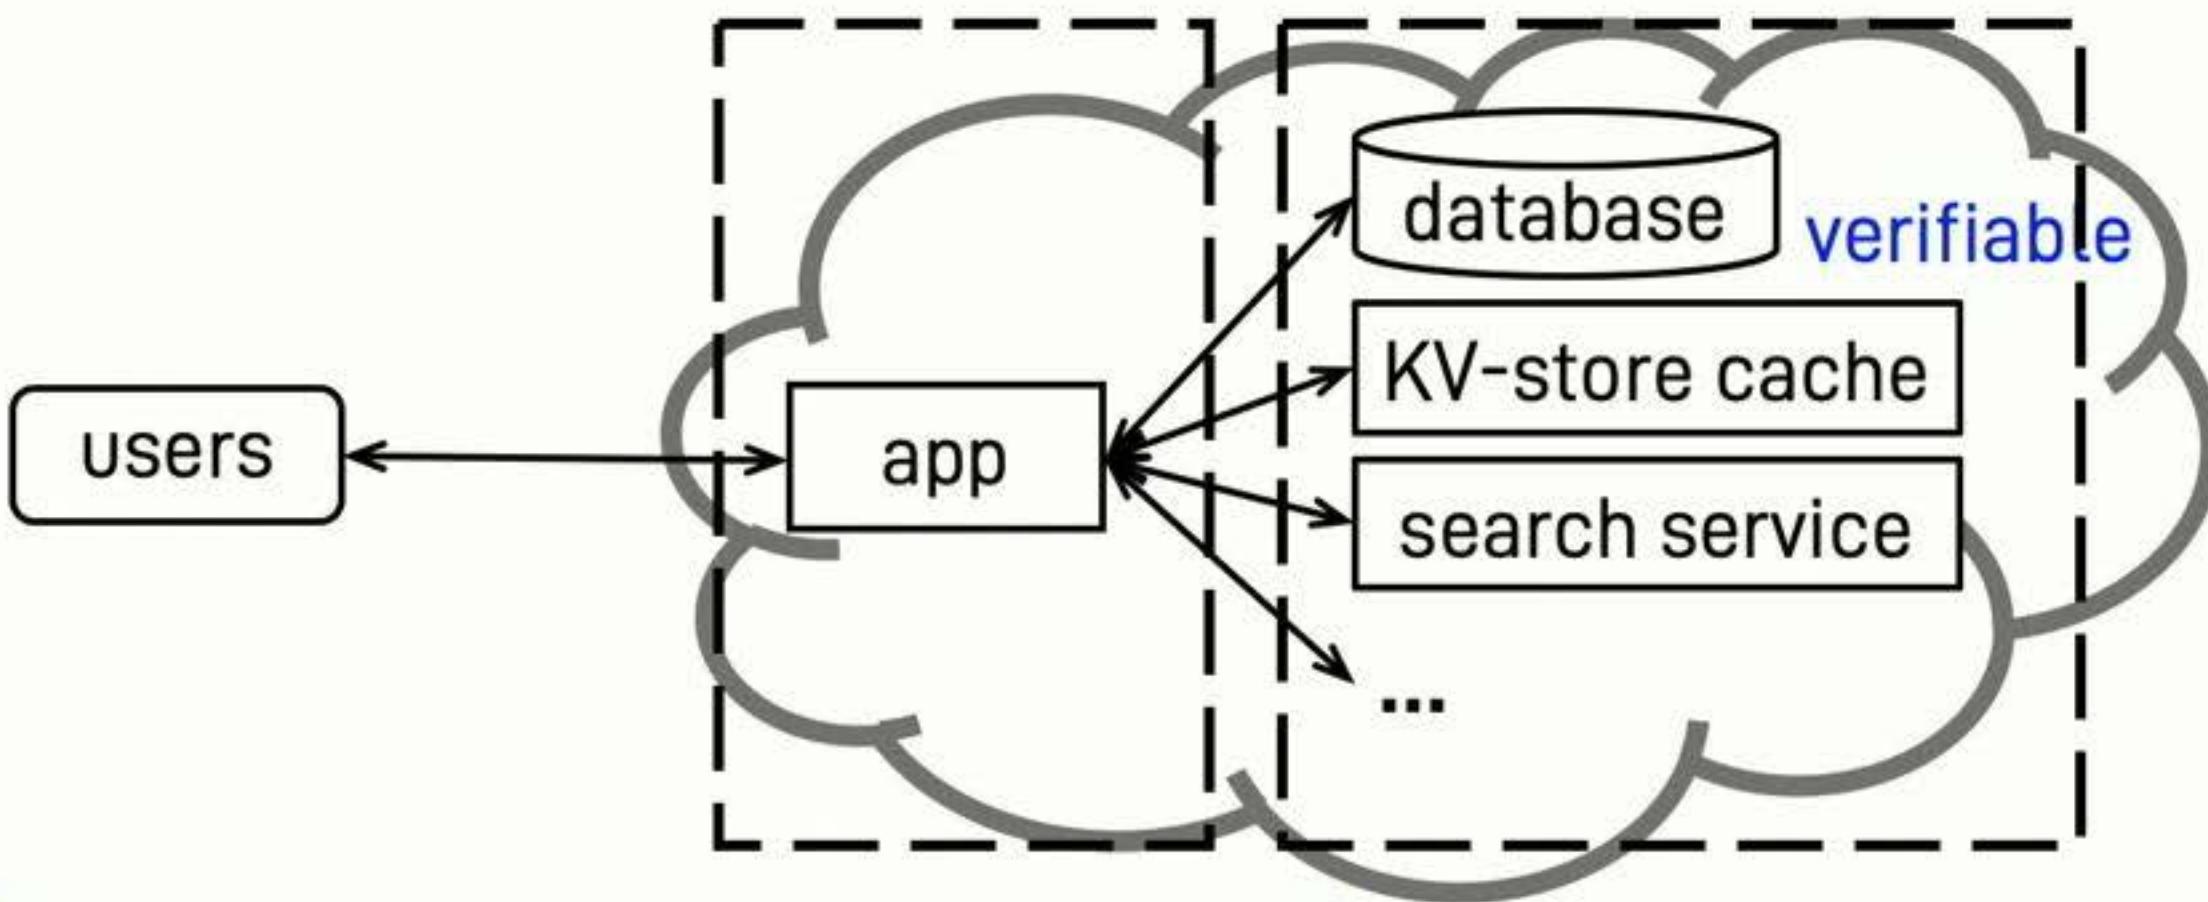
Rest of the talk: verifiable infrastructure



- ① verify white-box services

verifier

Rest of the talk: verifiable infrastructure

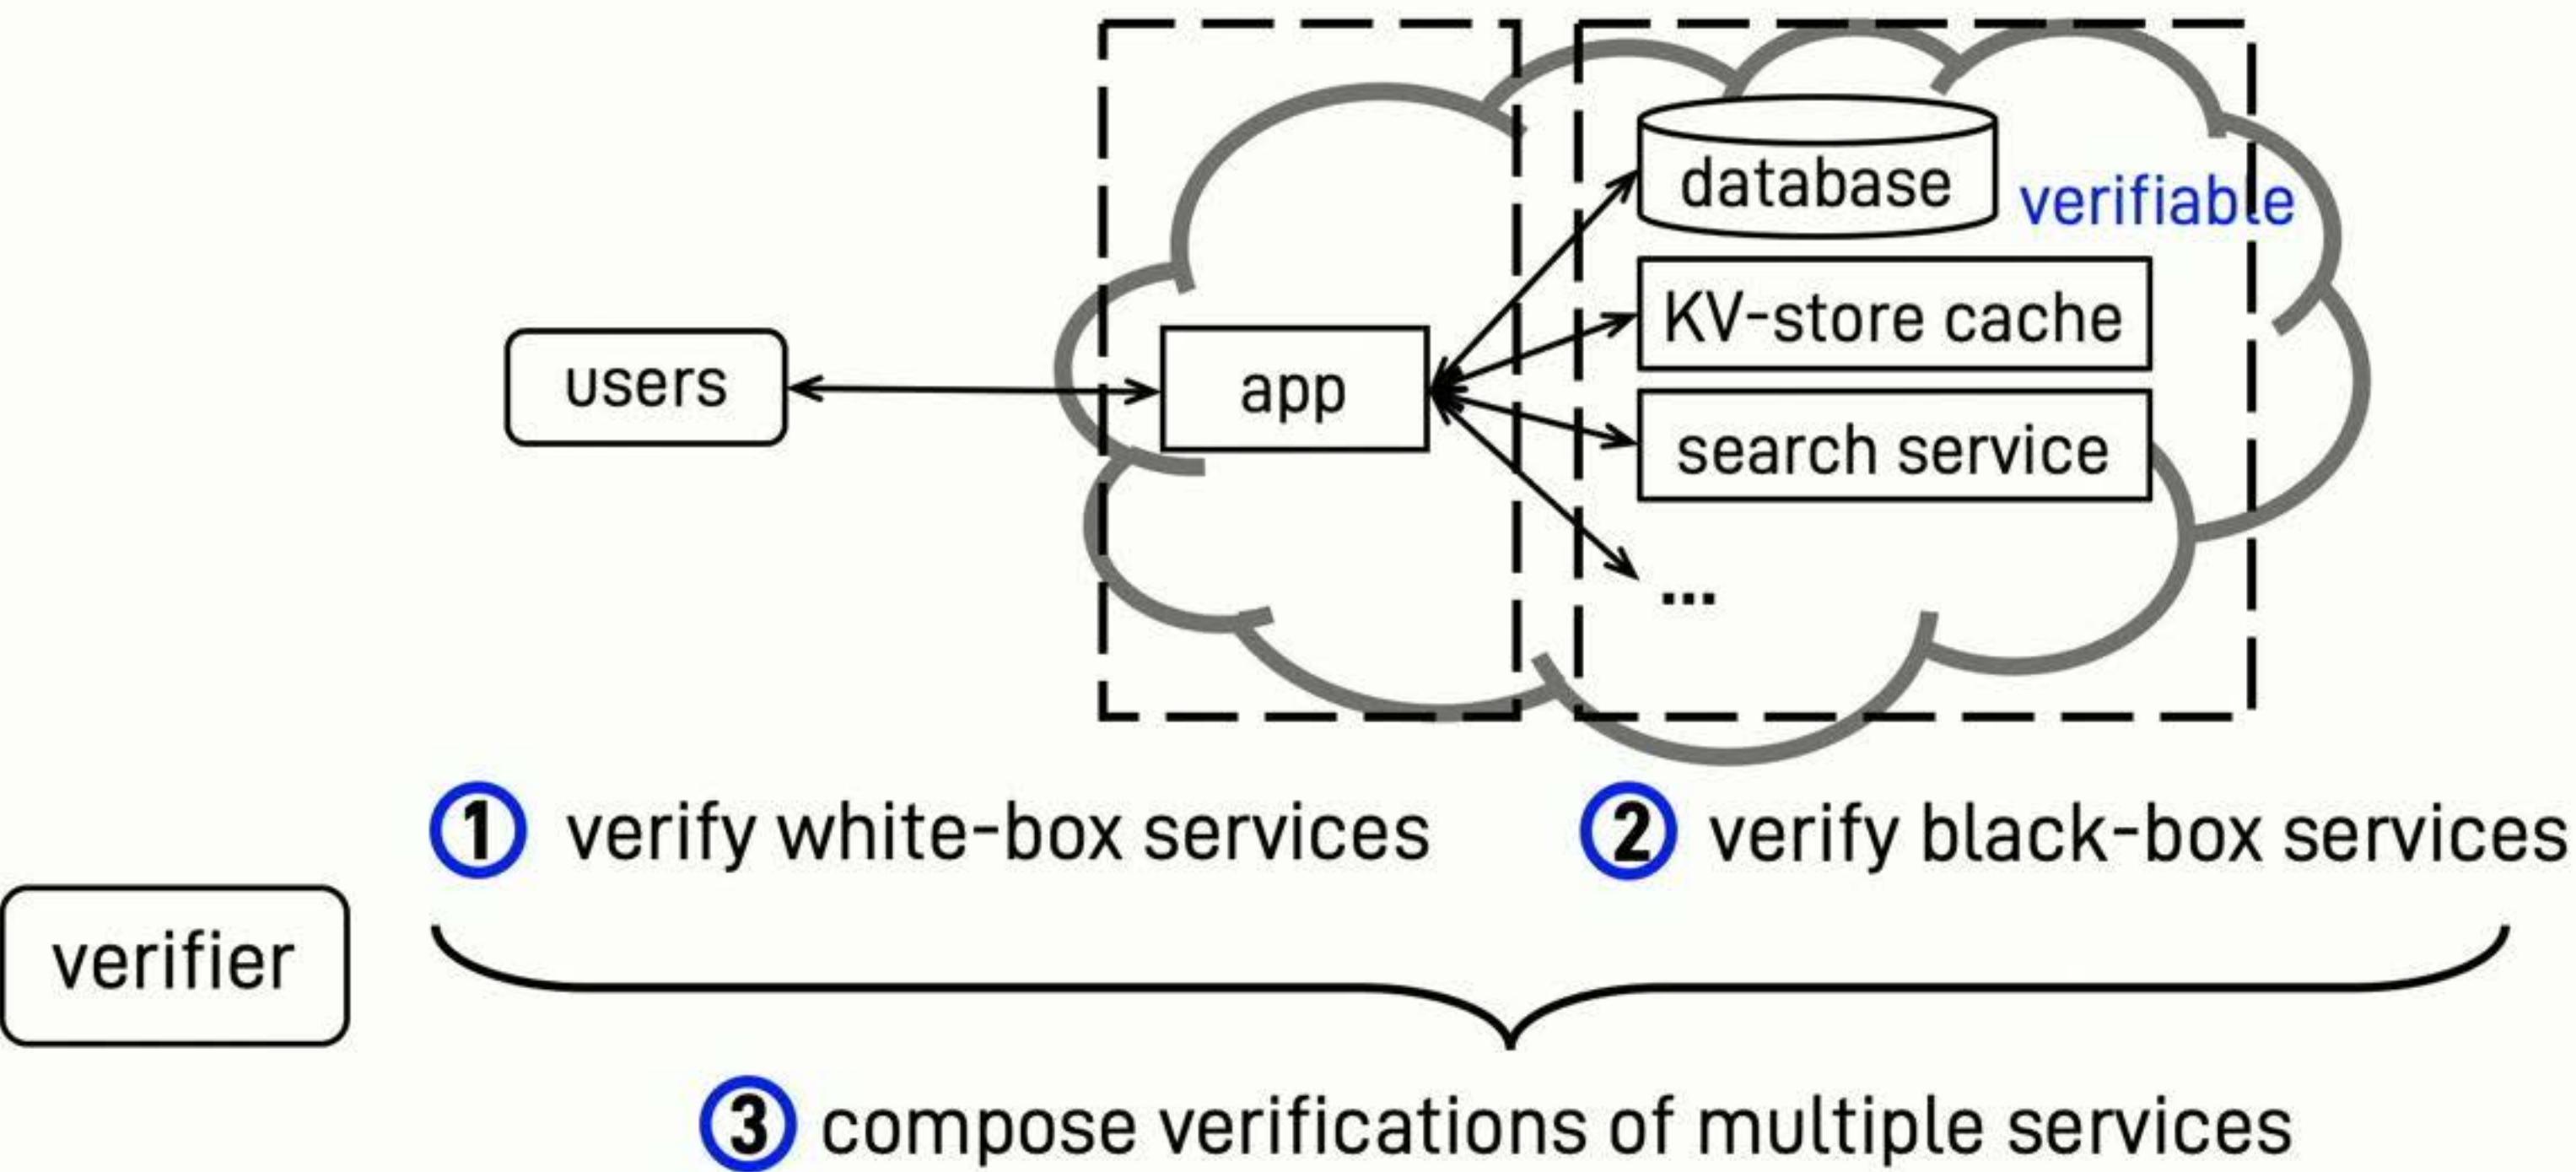


① verify white-box services

② verify black-box services

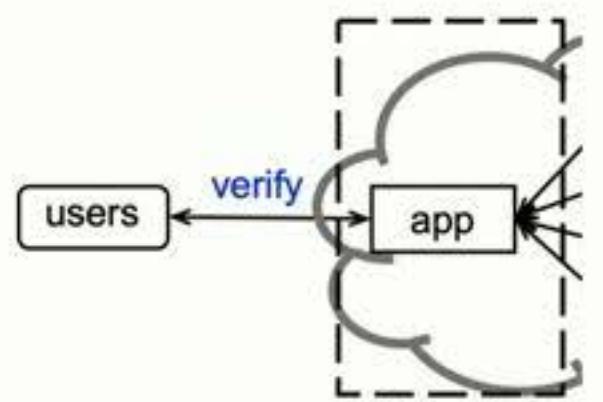
verifier

Rest of the talk: verifiable infrastructure

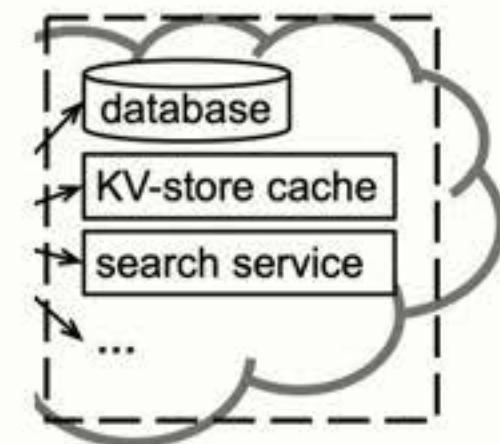


Rest of the talk: verifiable infrastructure

1 Verify white-box services (Orochi [SOSP'17])



2 Verify black-box services (Cobra)



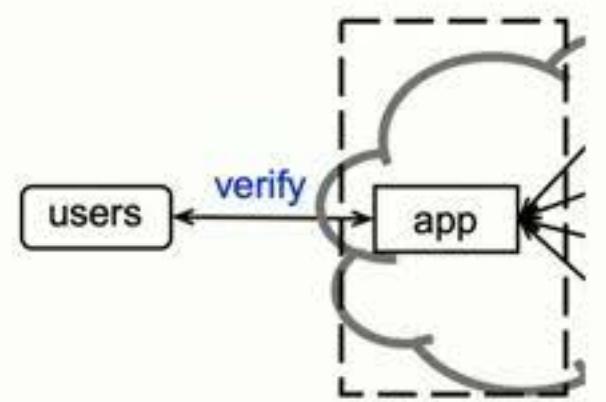
3 Build composable verifiable framework (future work)

4 Other past work

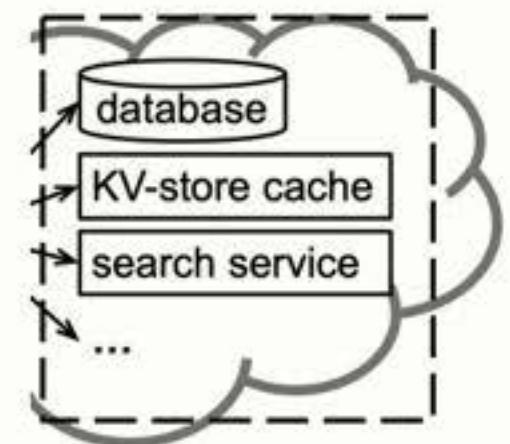
- Troubleshooting data center networks [NSDI'19]
- Protecting secret via security-oriented offloading [EuroSys'15]

Rest of the talk: verifiable infrastructure

1 Verify white-box services (Orochi [SOSP'17])



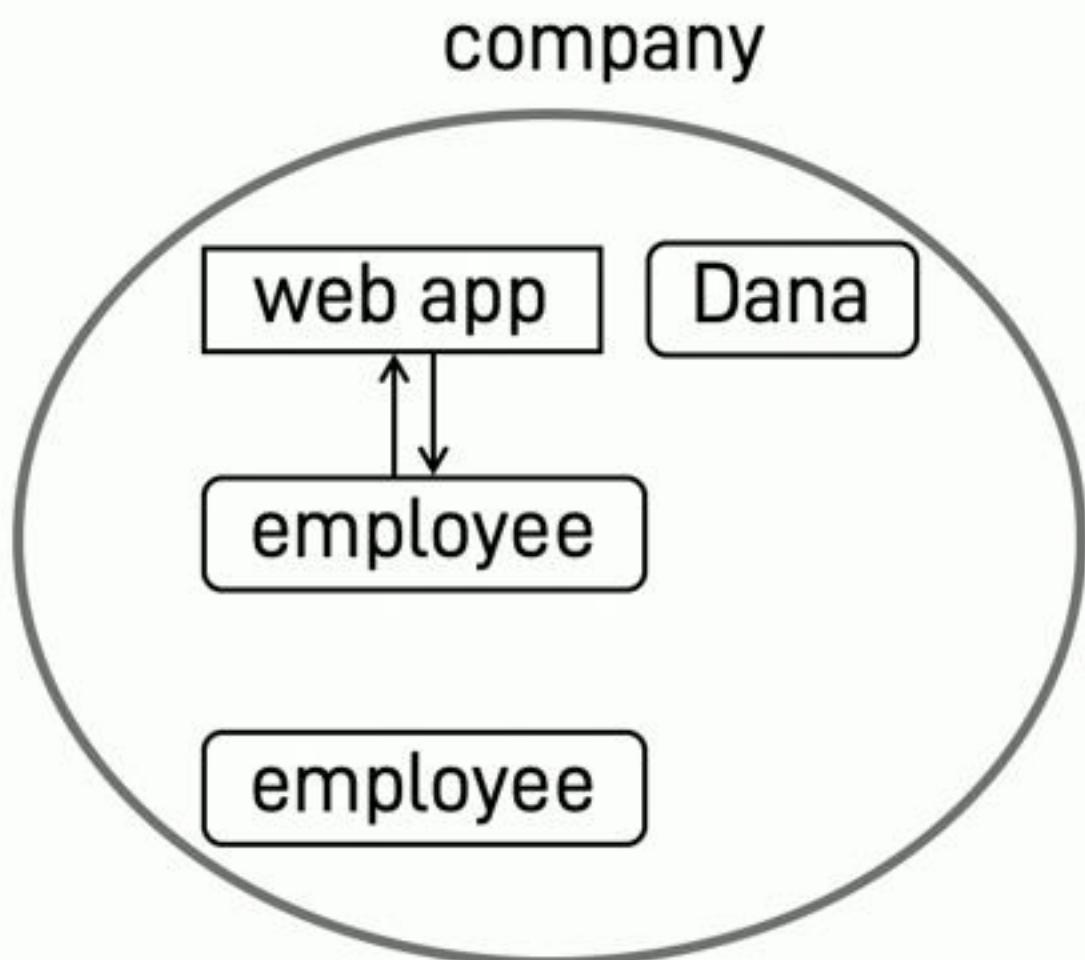
2 Verify black-box services (Cobra)

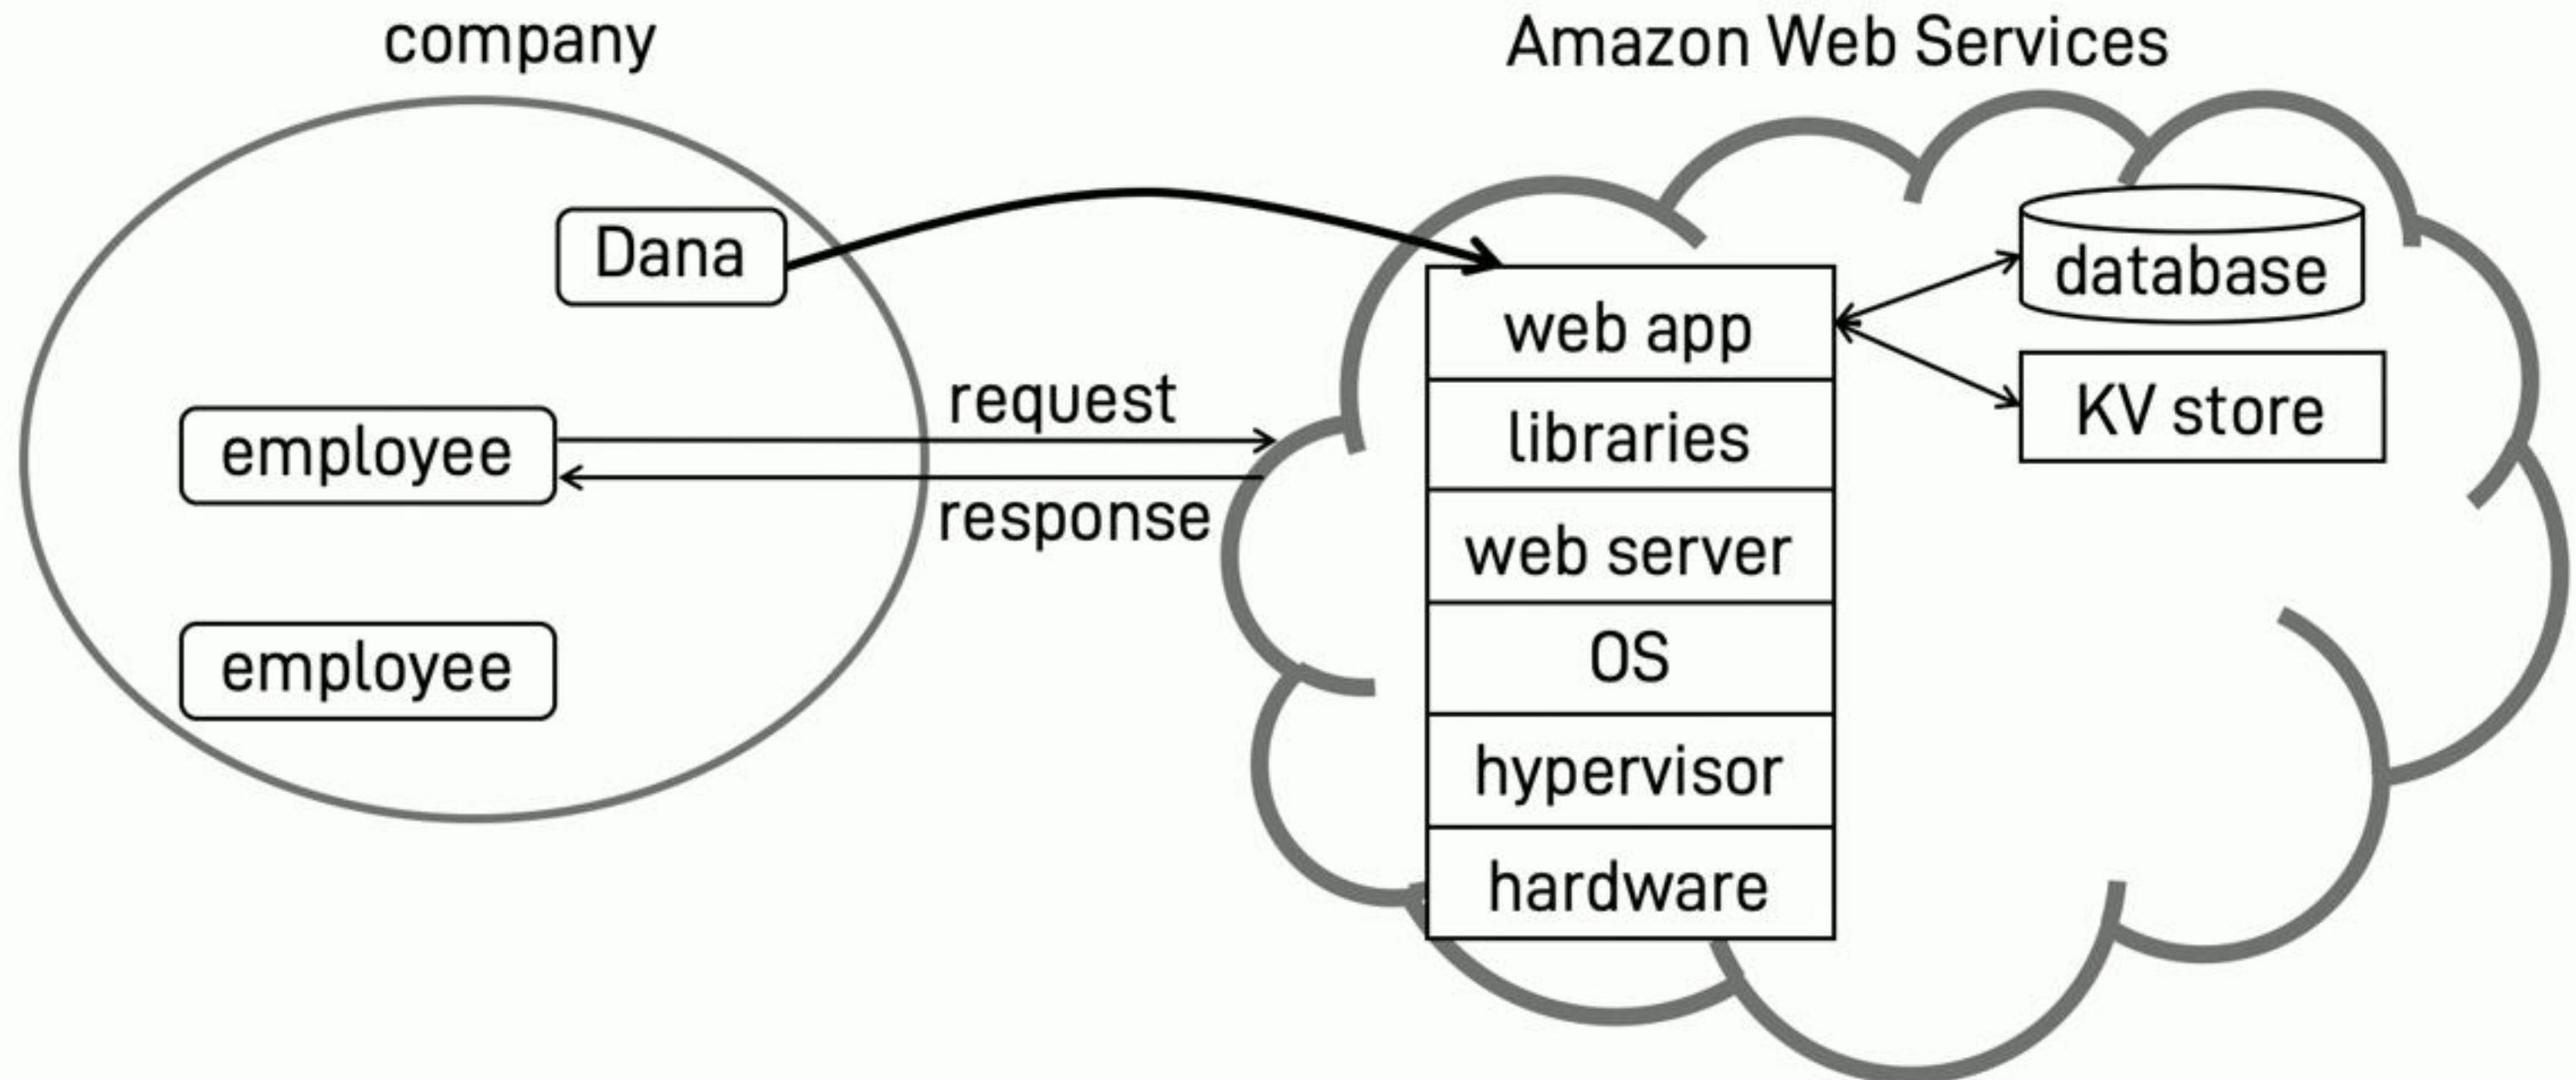


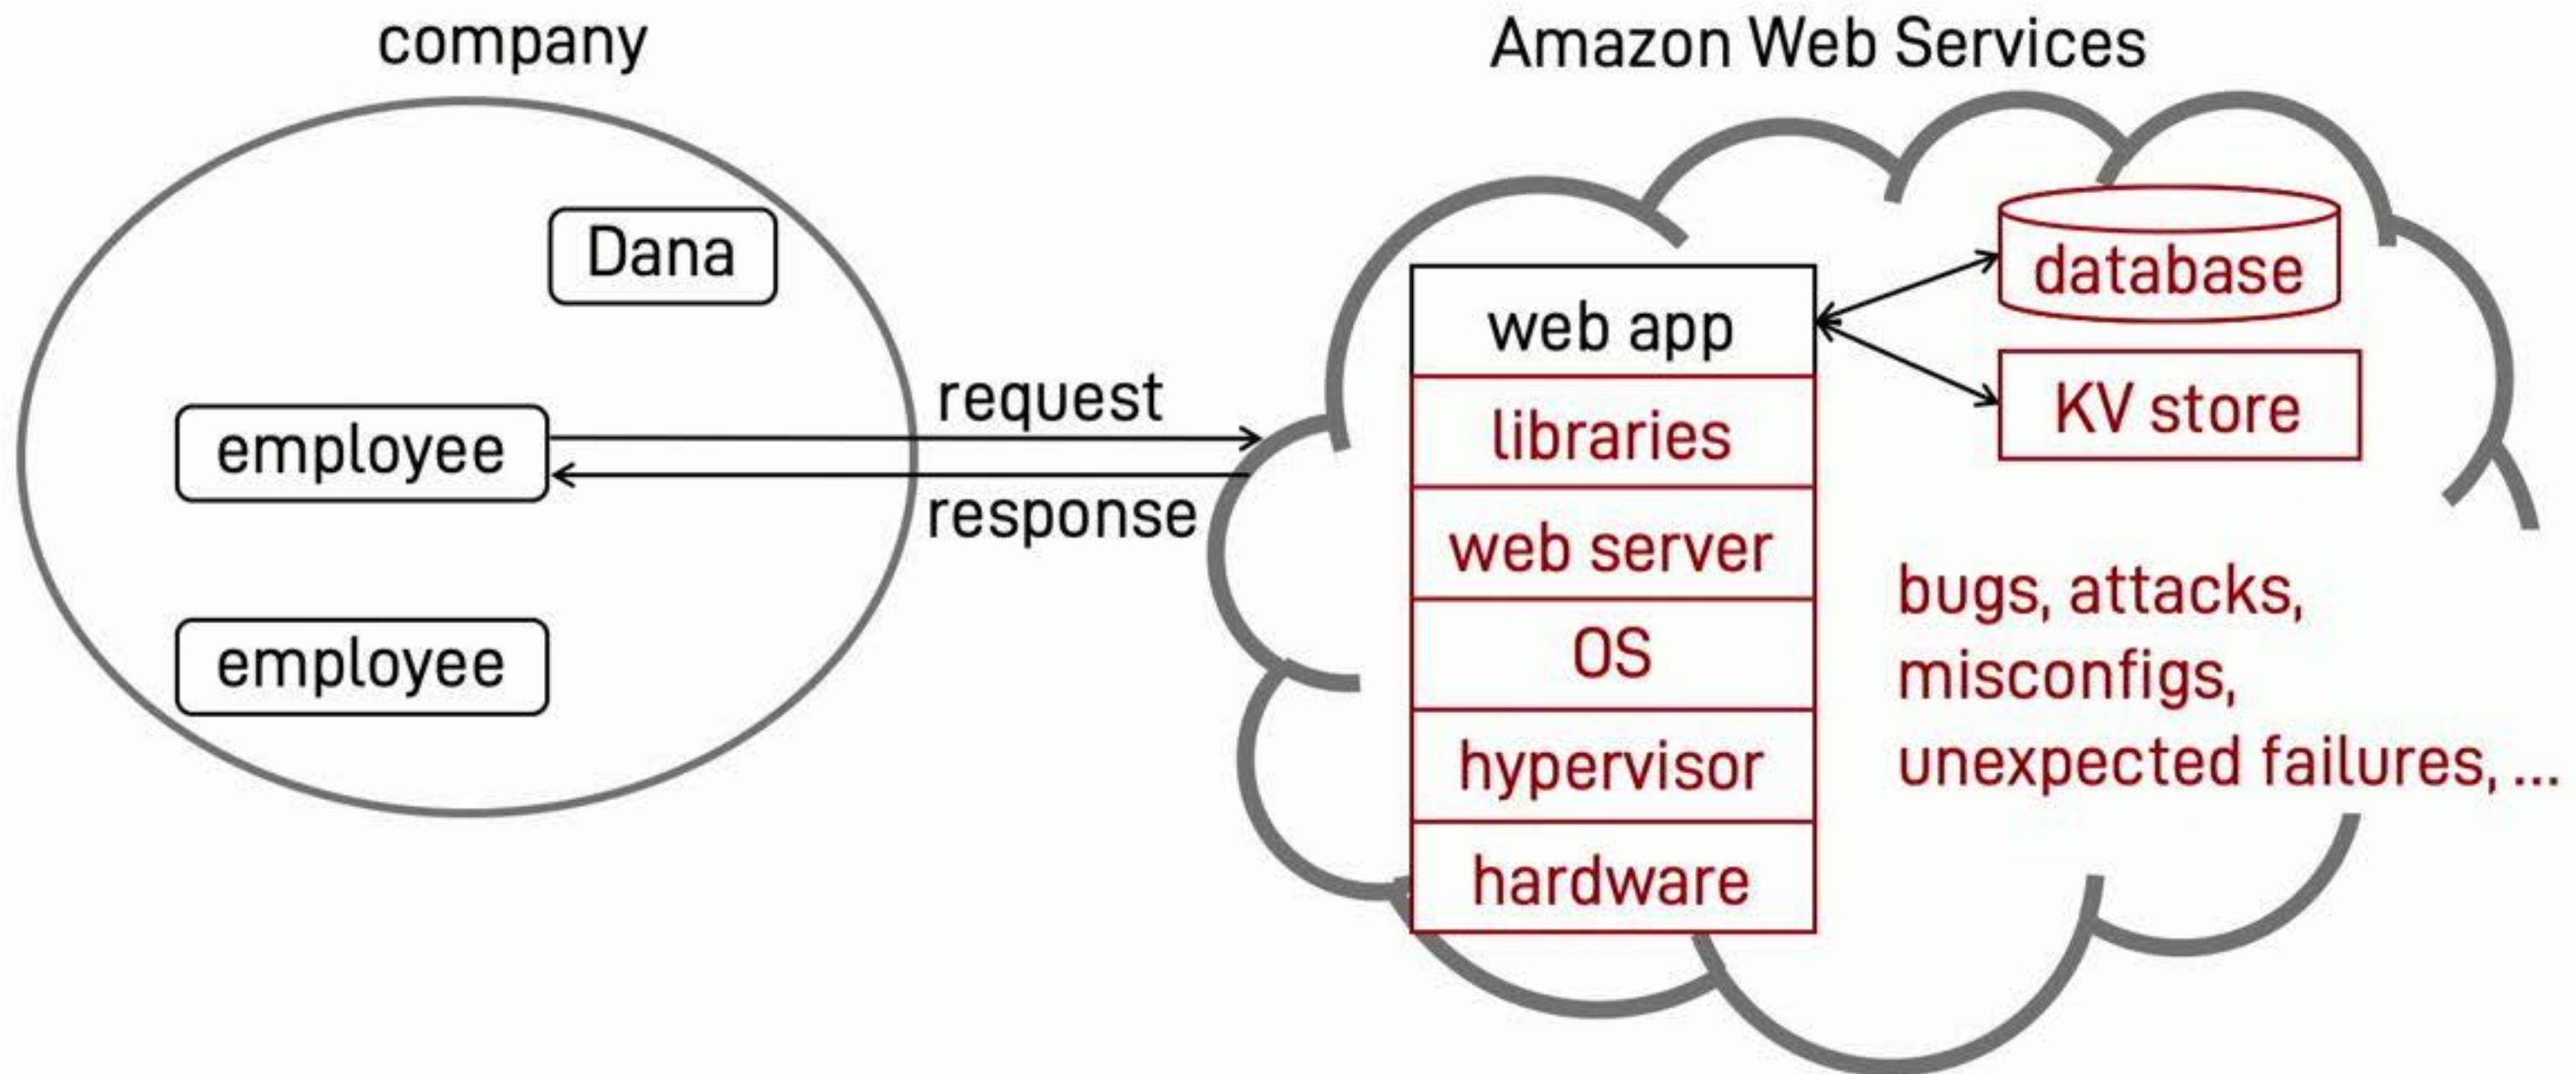
3 Build composable verifiable framework (future work)

4 Other past work

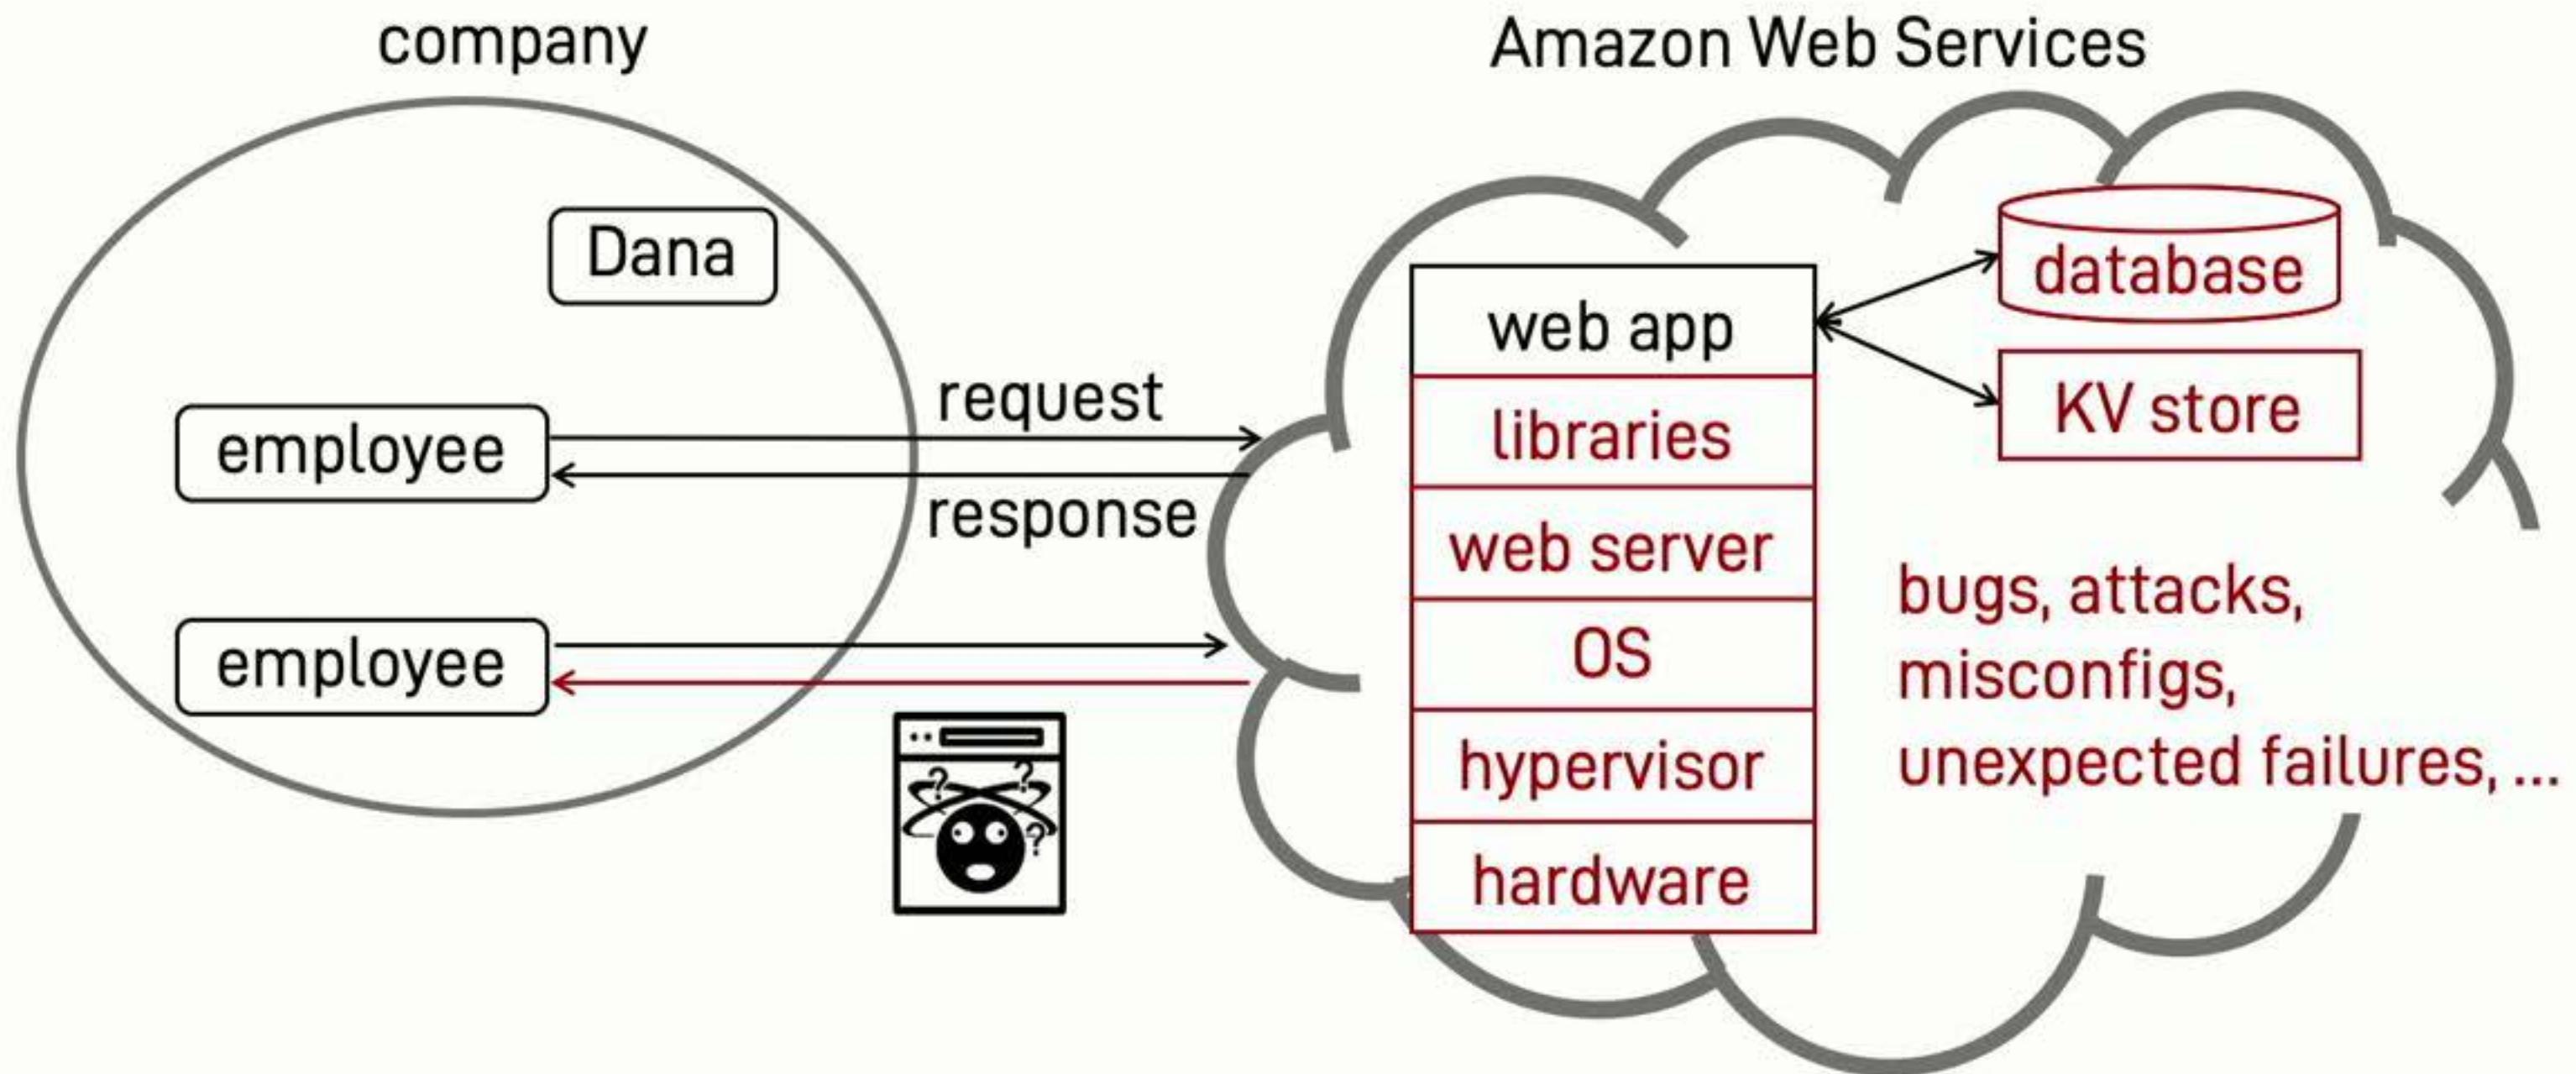
- Troubleshooting data center networks [NSDI'19]
- Protecting secret via security-oriented offloading [EuroSys'15]



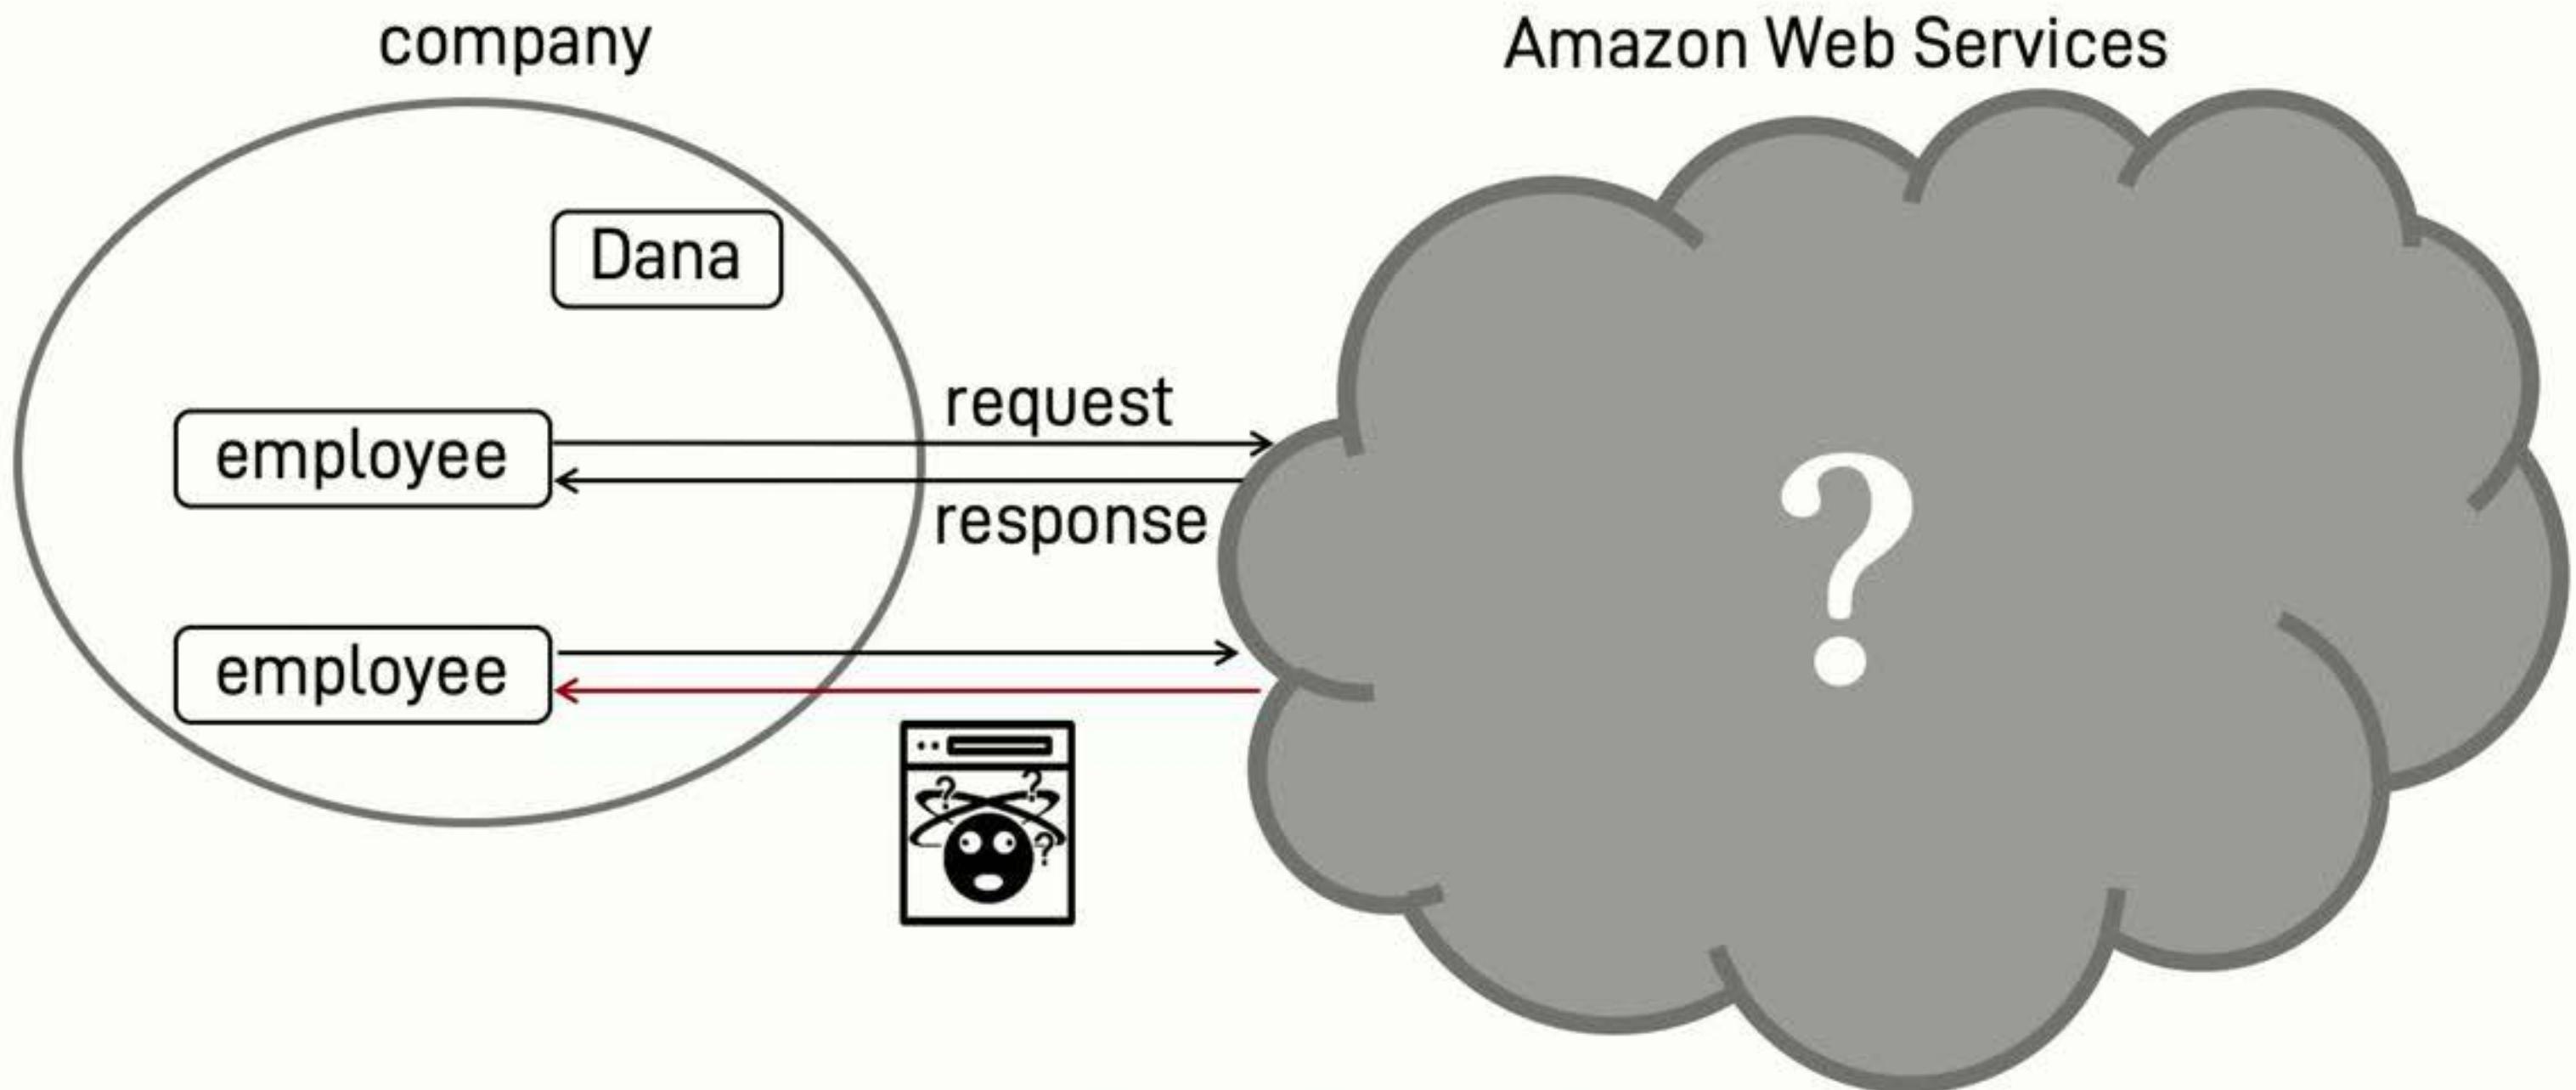




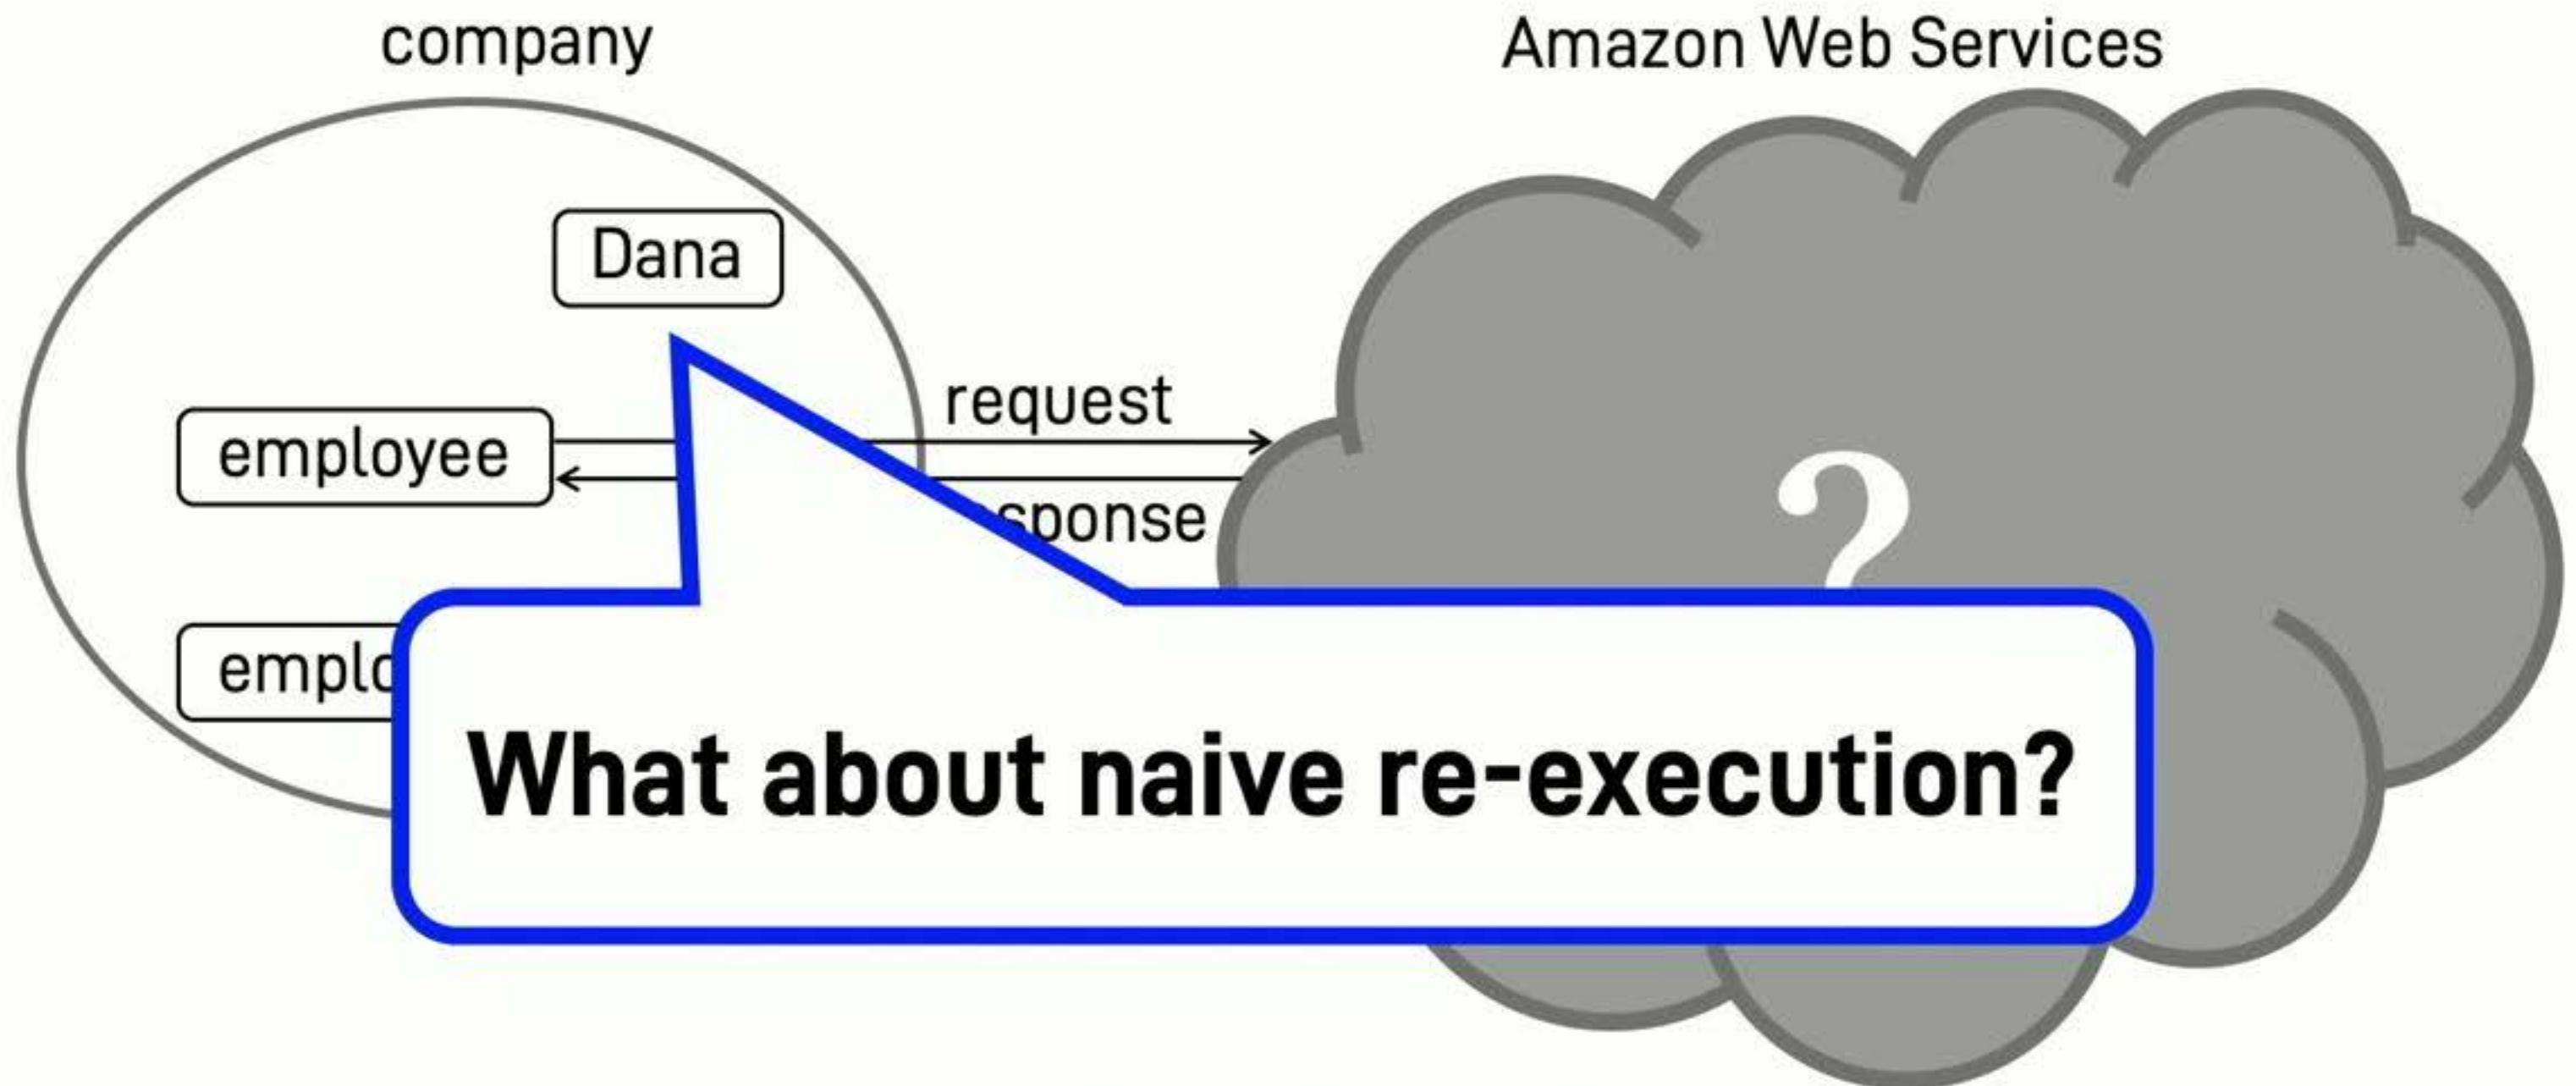
- A lot of things can go wrong...



- A lot of things can go wrong...

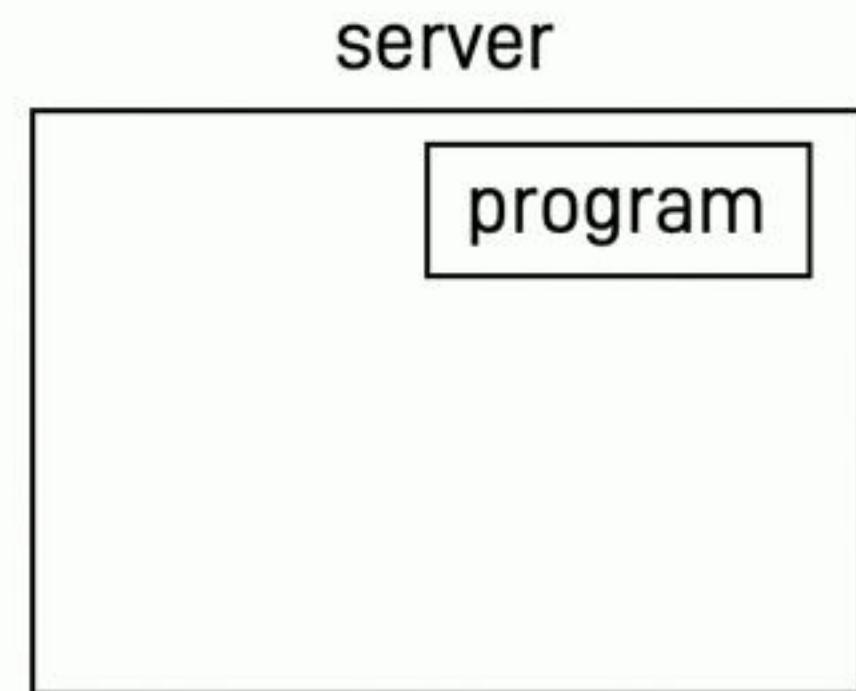


- Dana wants to audit the **end-to-end execution**
responses $\stackrel{?}{=}$ **actual application** + **requests**

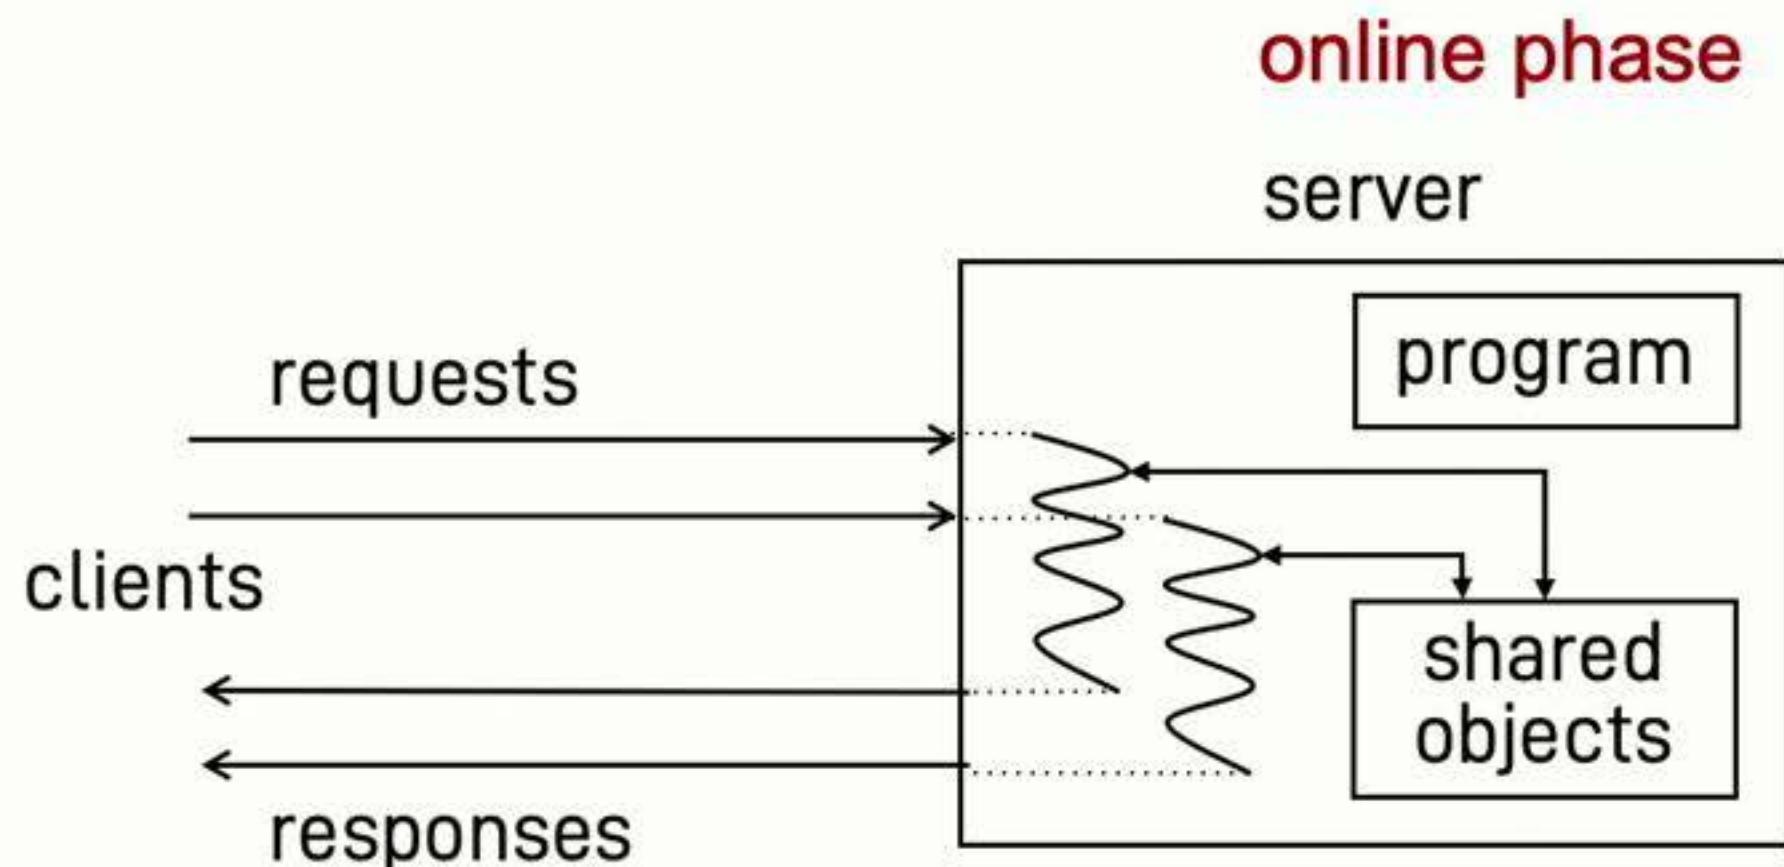


- Dana wants to audit the **end-to-end execution**
responses $\stackrel{?}{=}$ **actual application** + **requests**

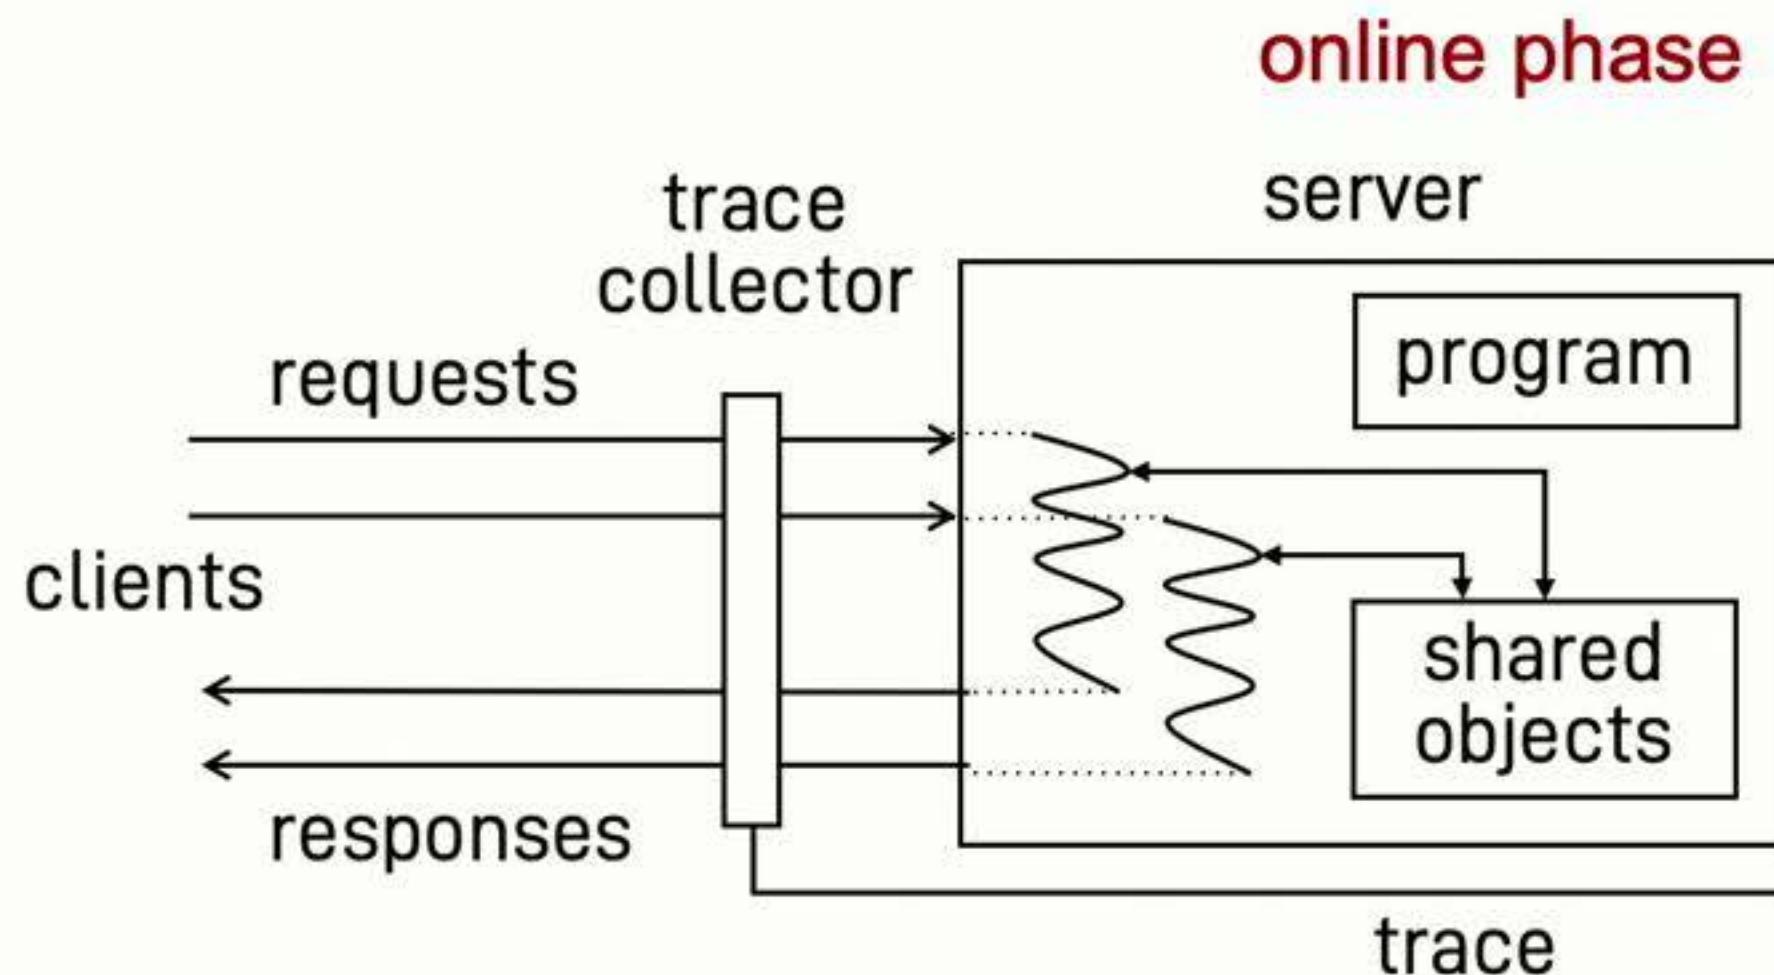
The Efficient Server Audit Problem

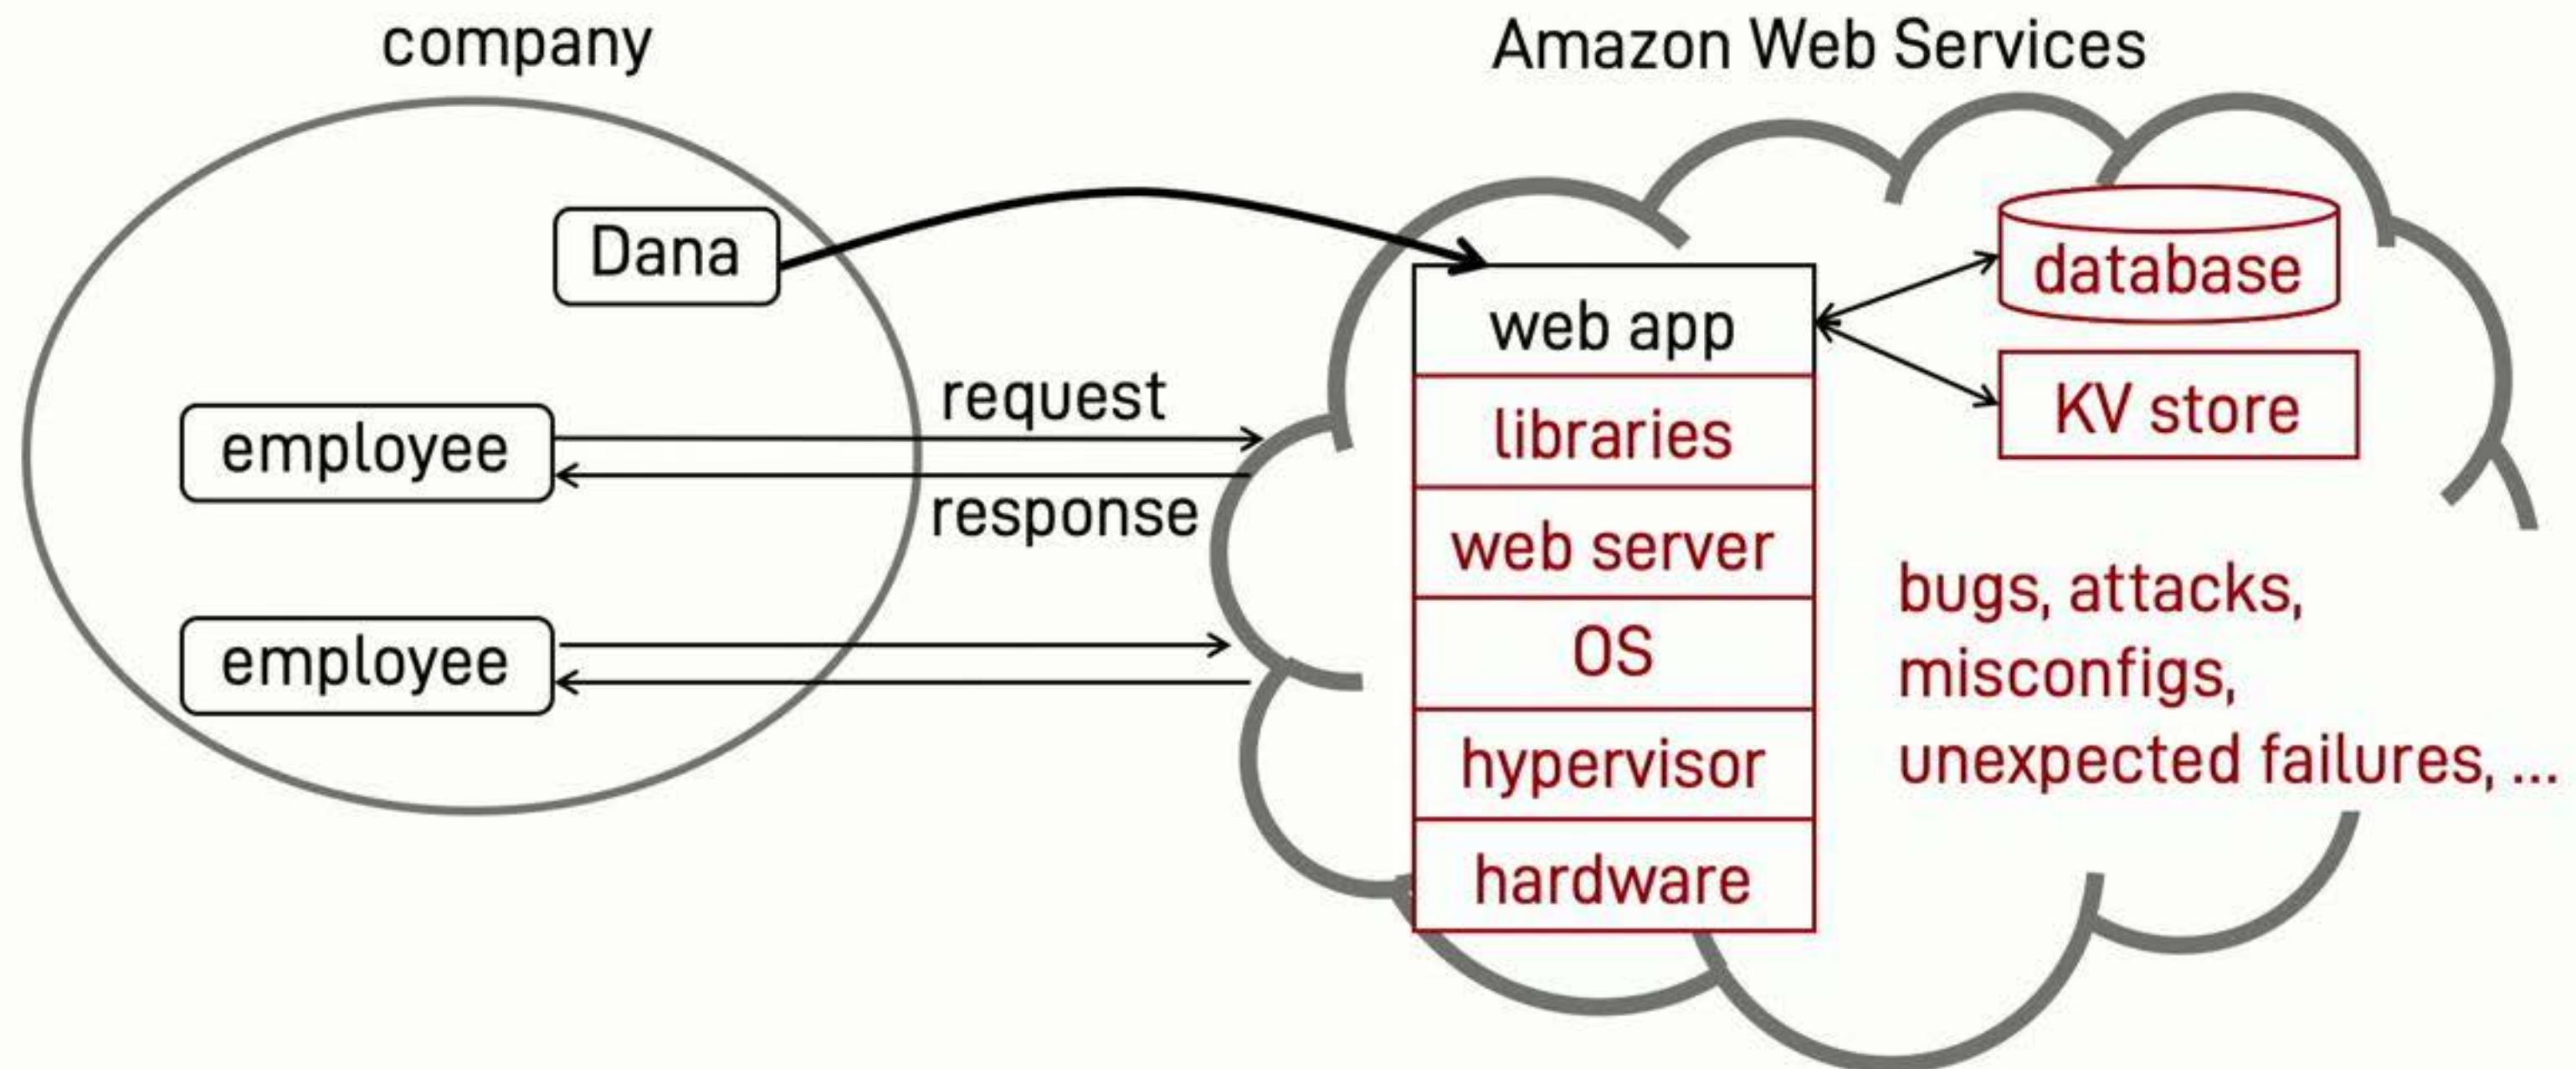


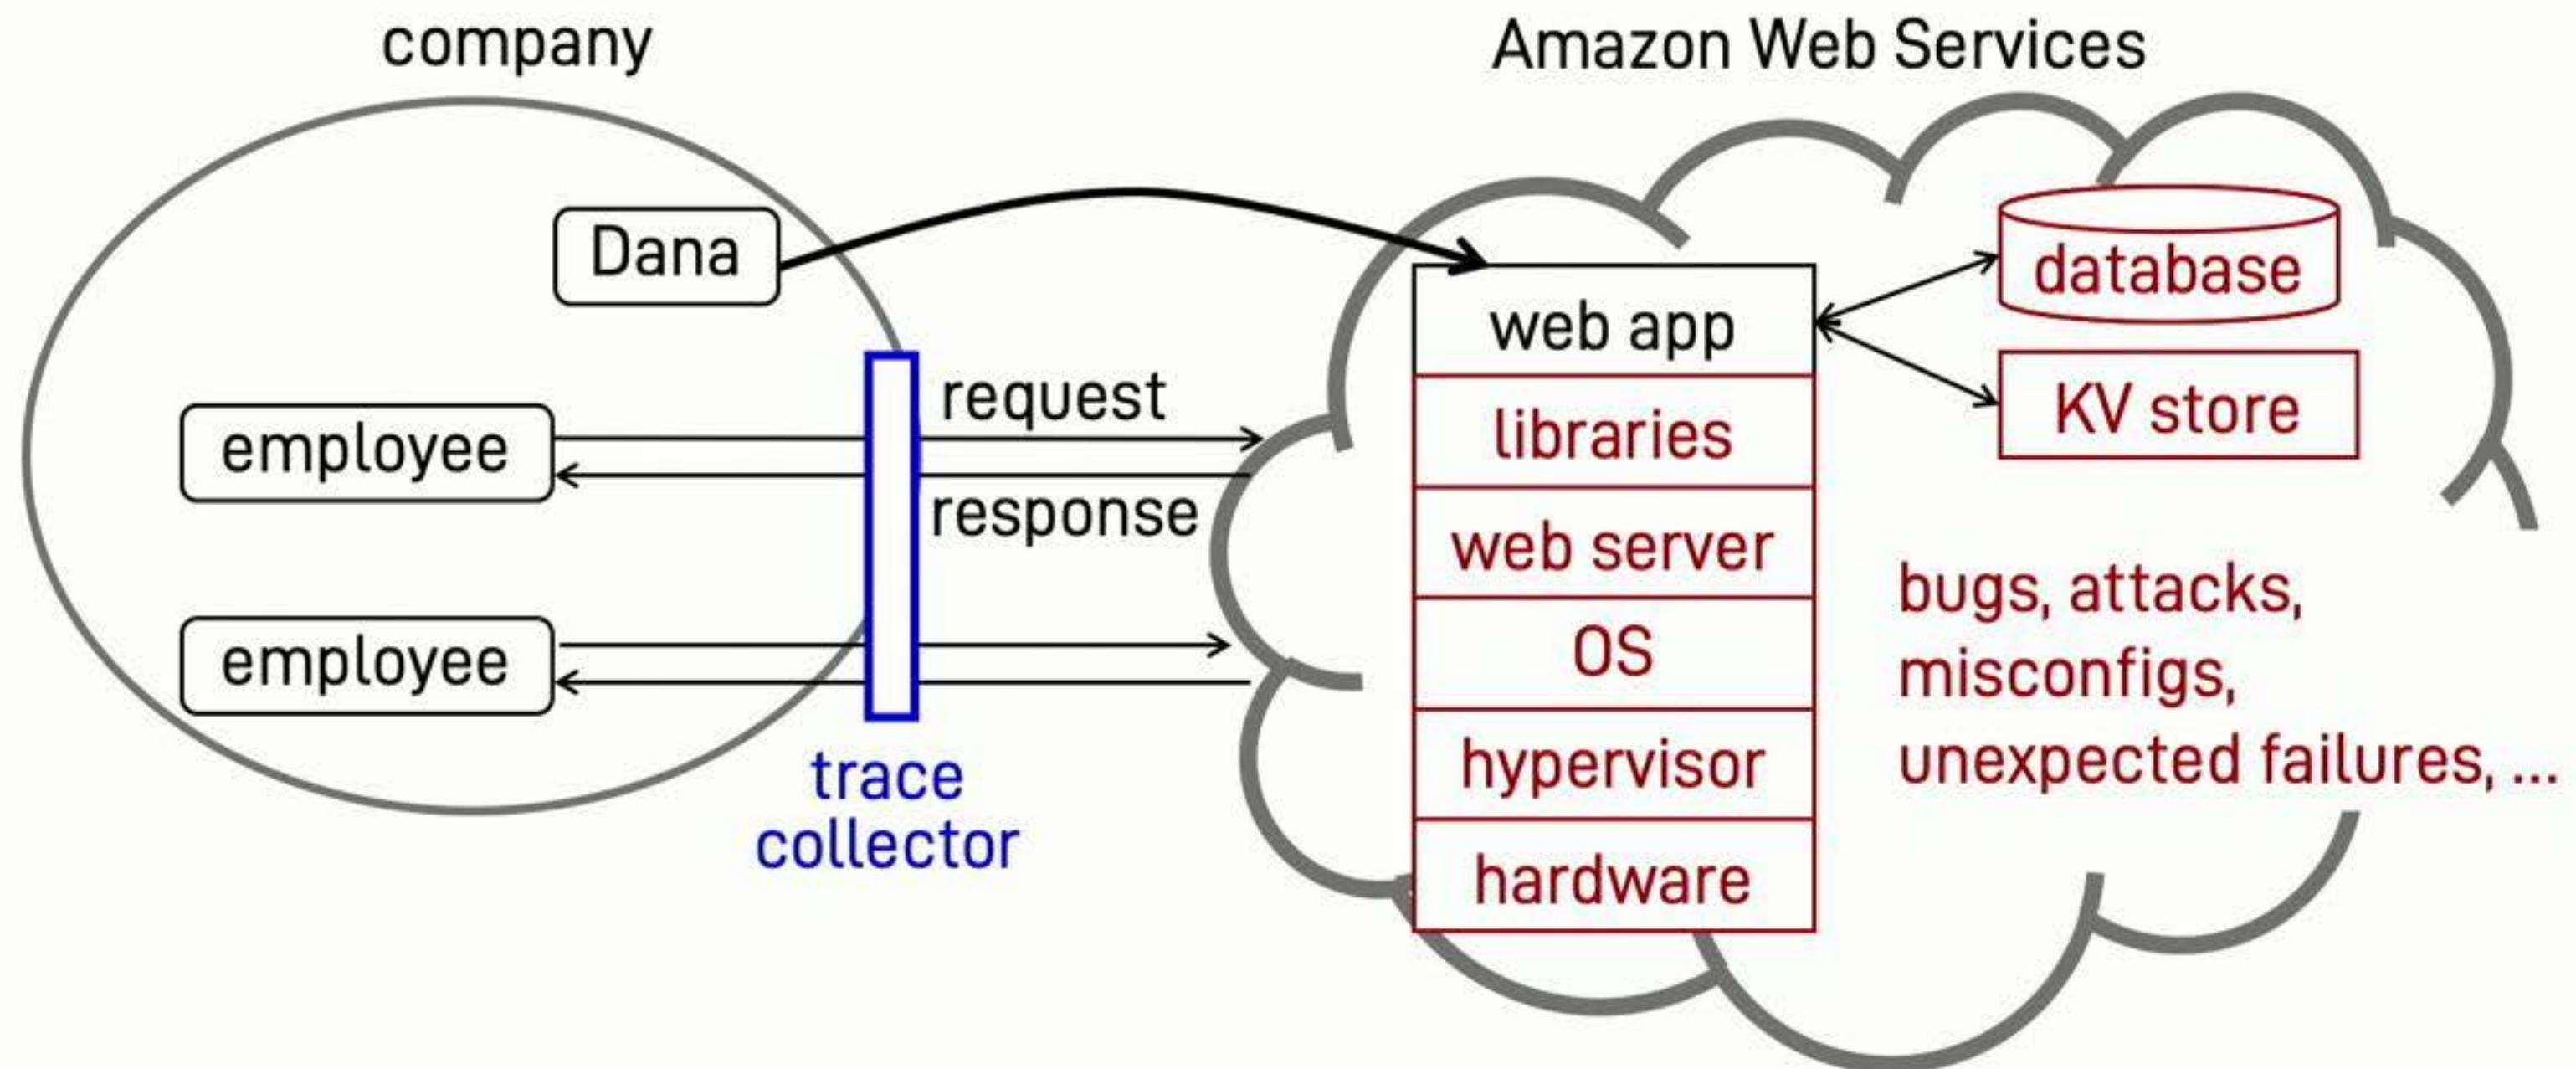
The Efficient Server Audit Problem



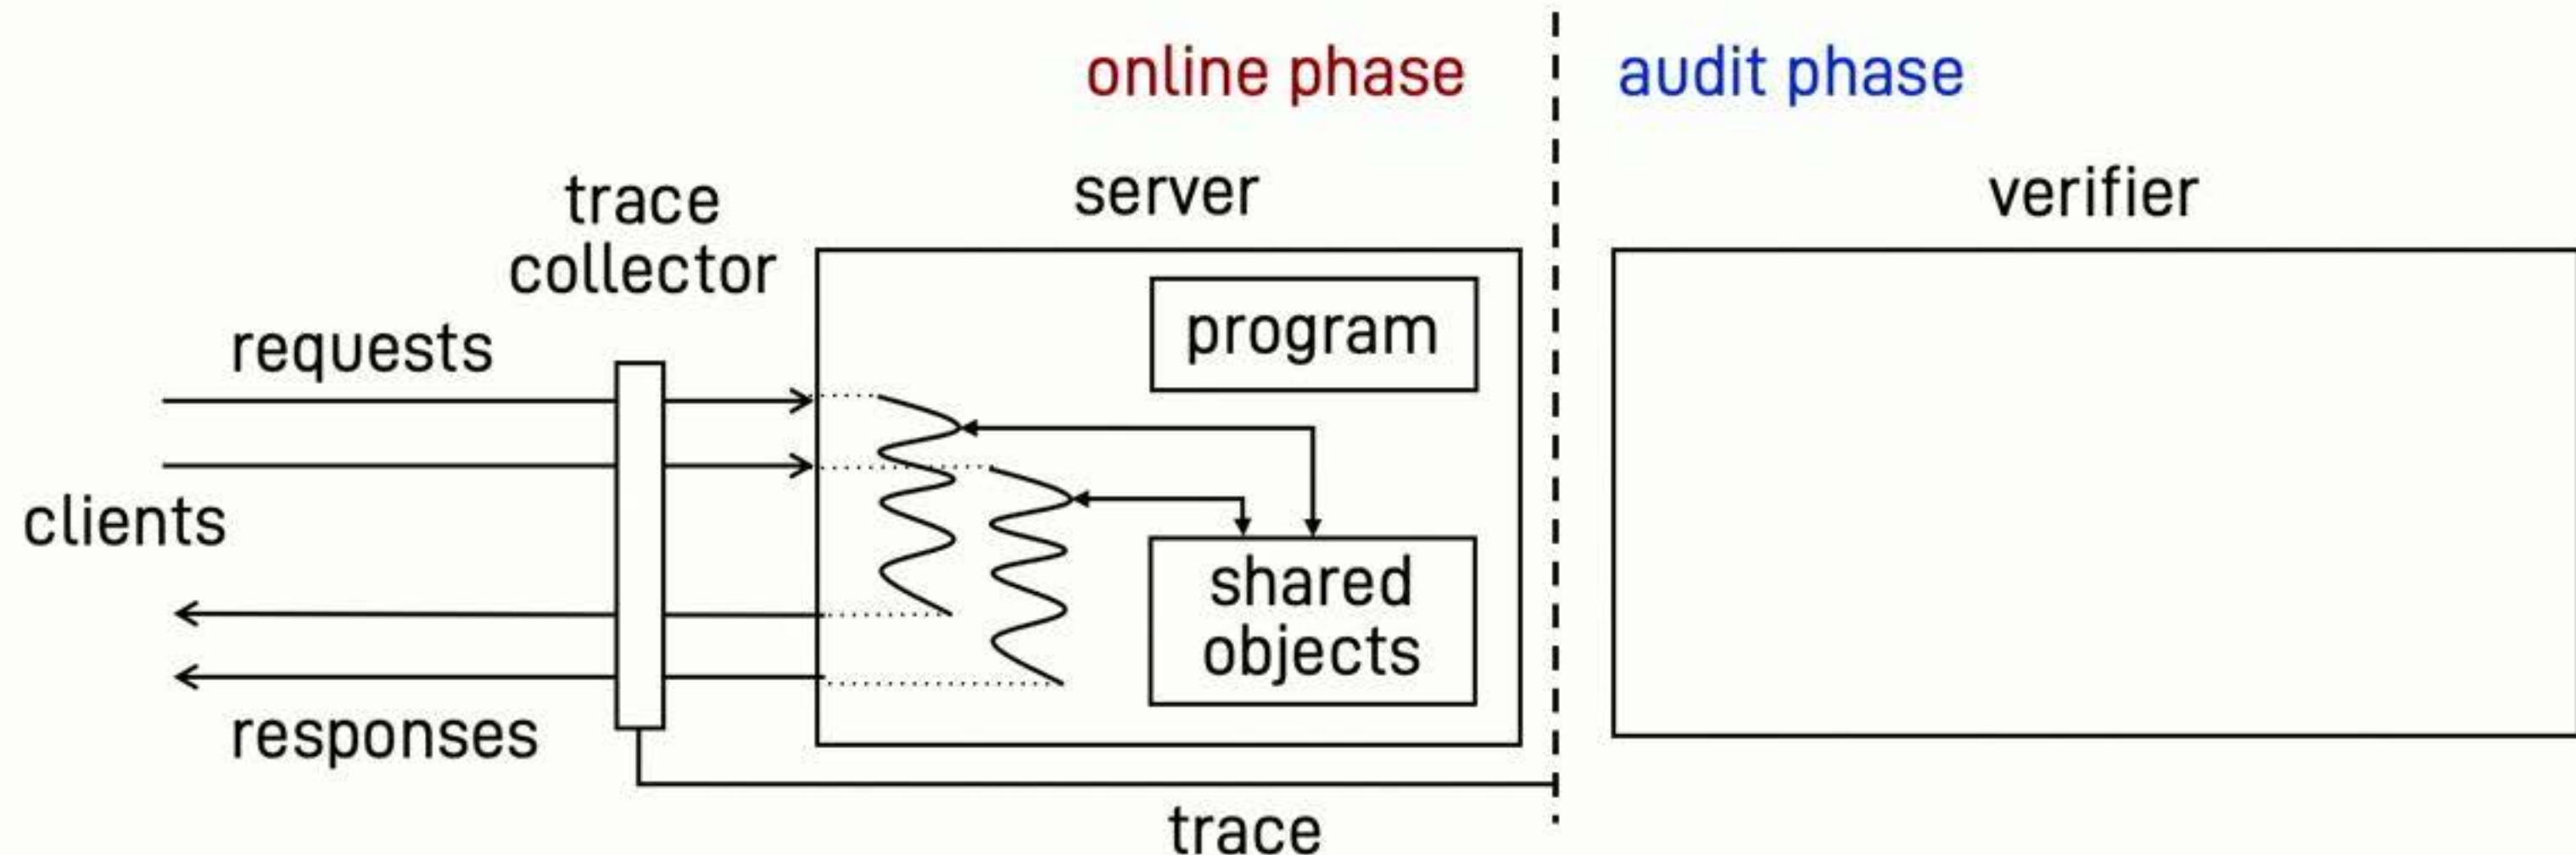
The Efficient Server Audit Problem



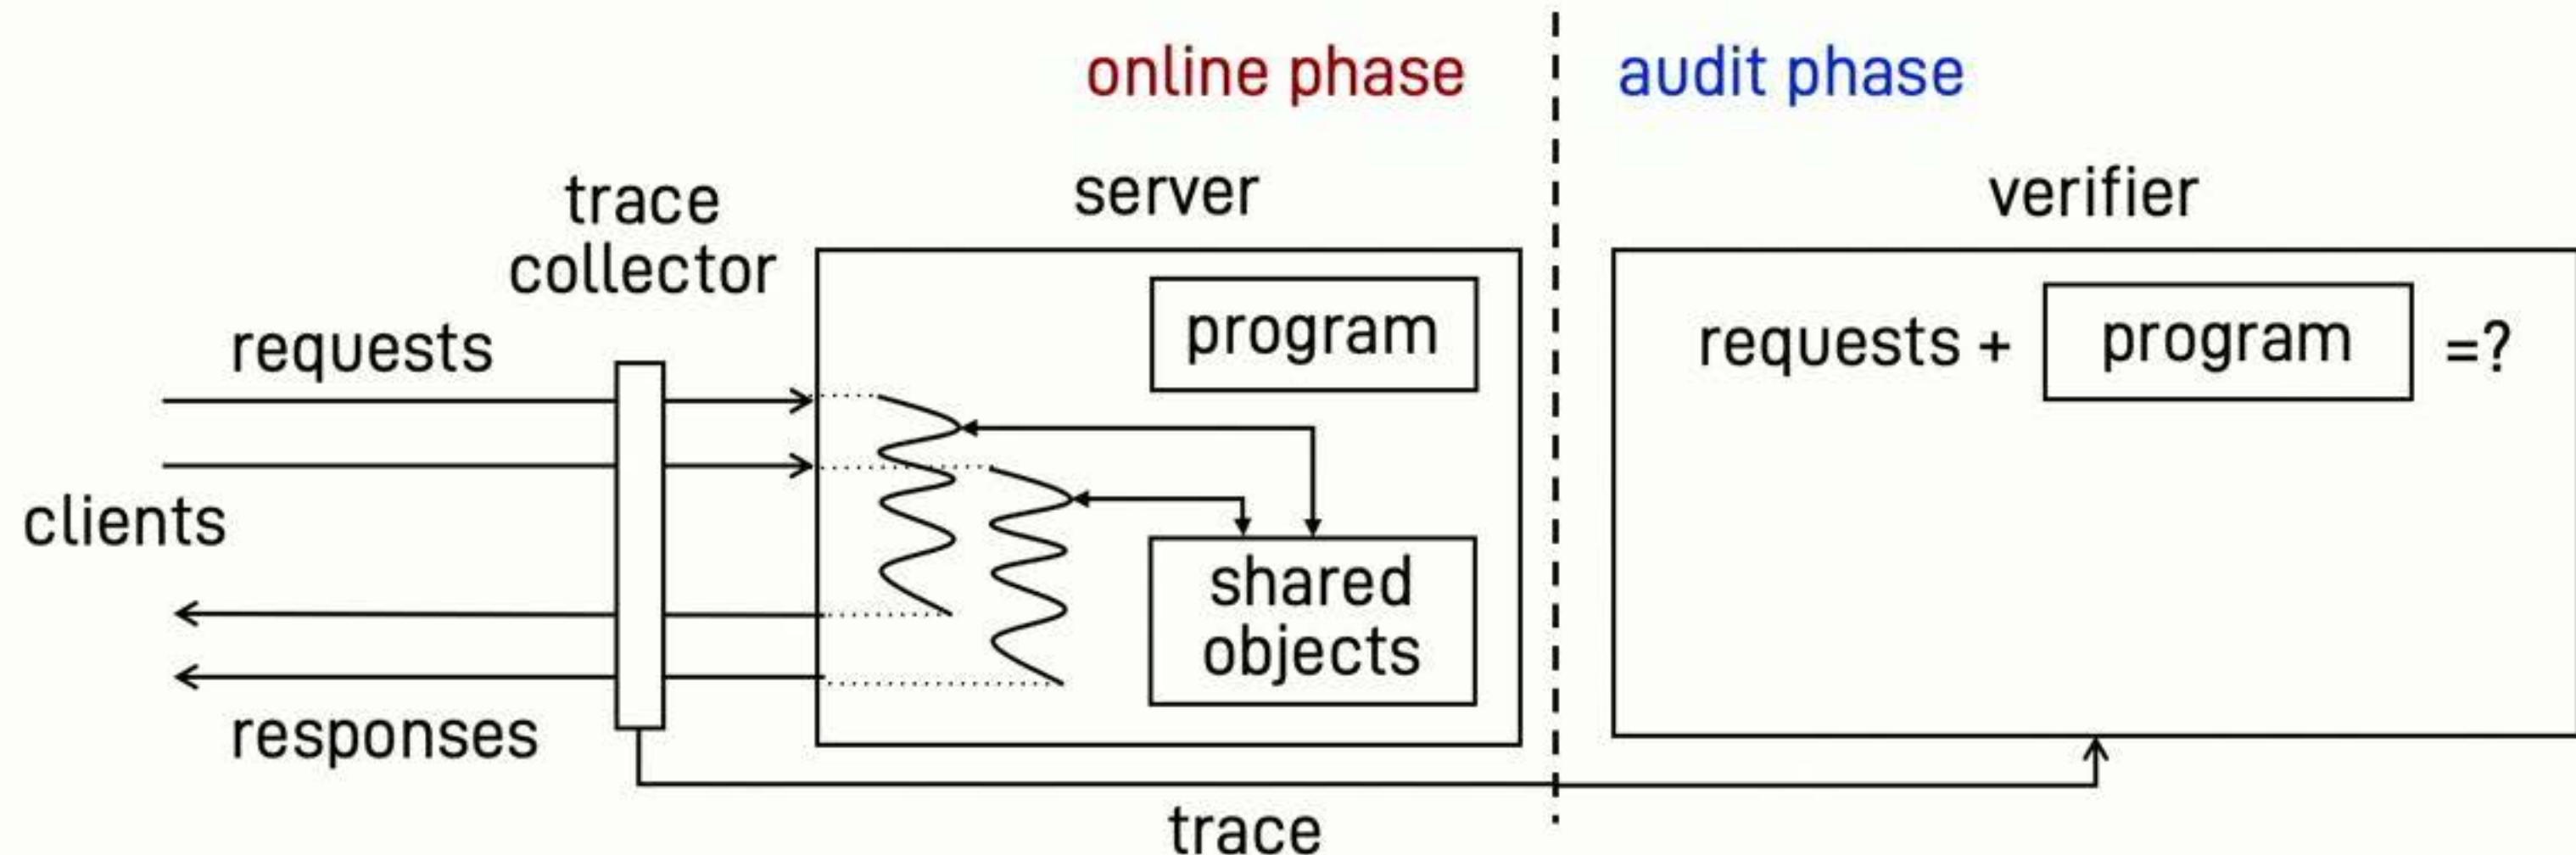




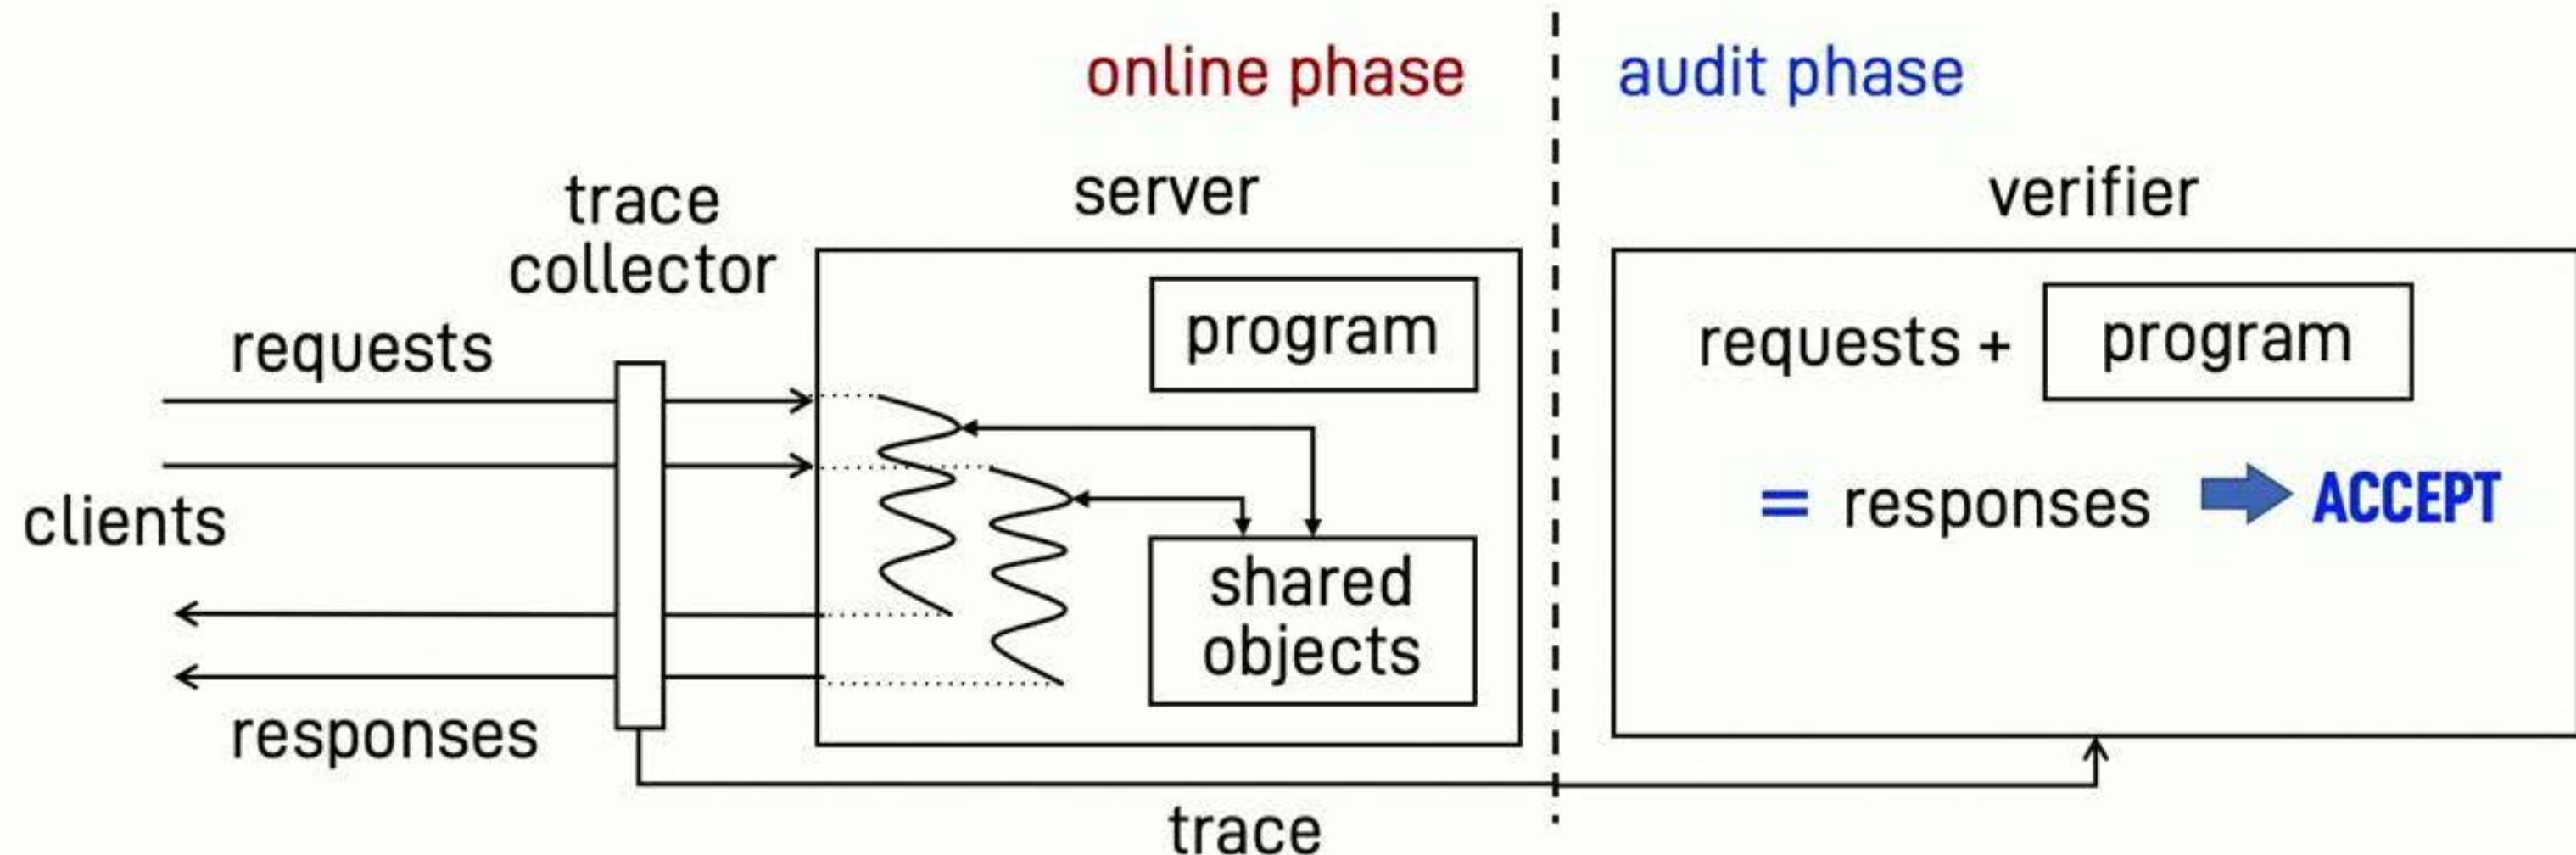
The Efficient Server Audit Problem



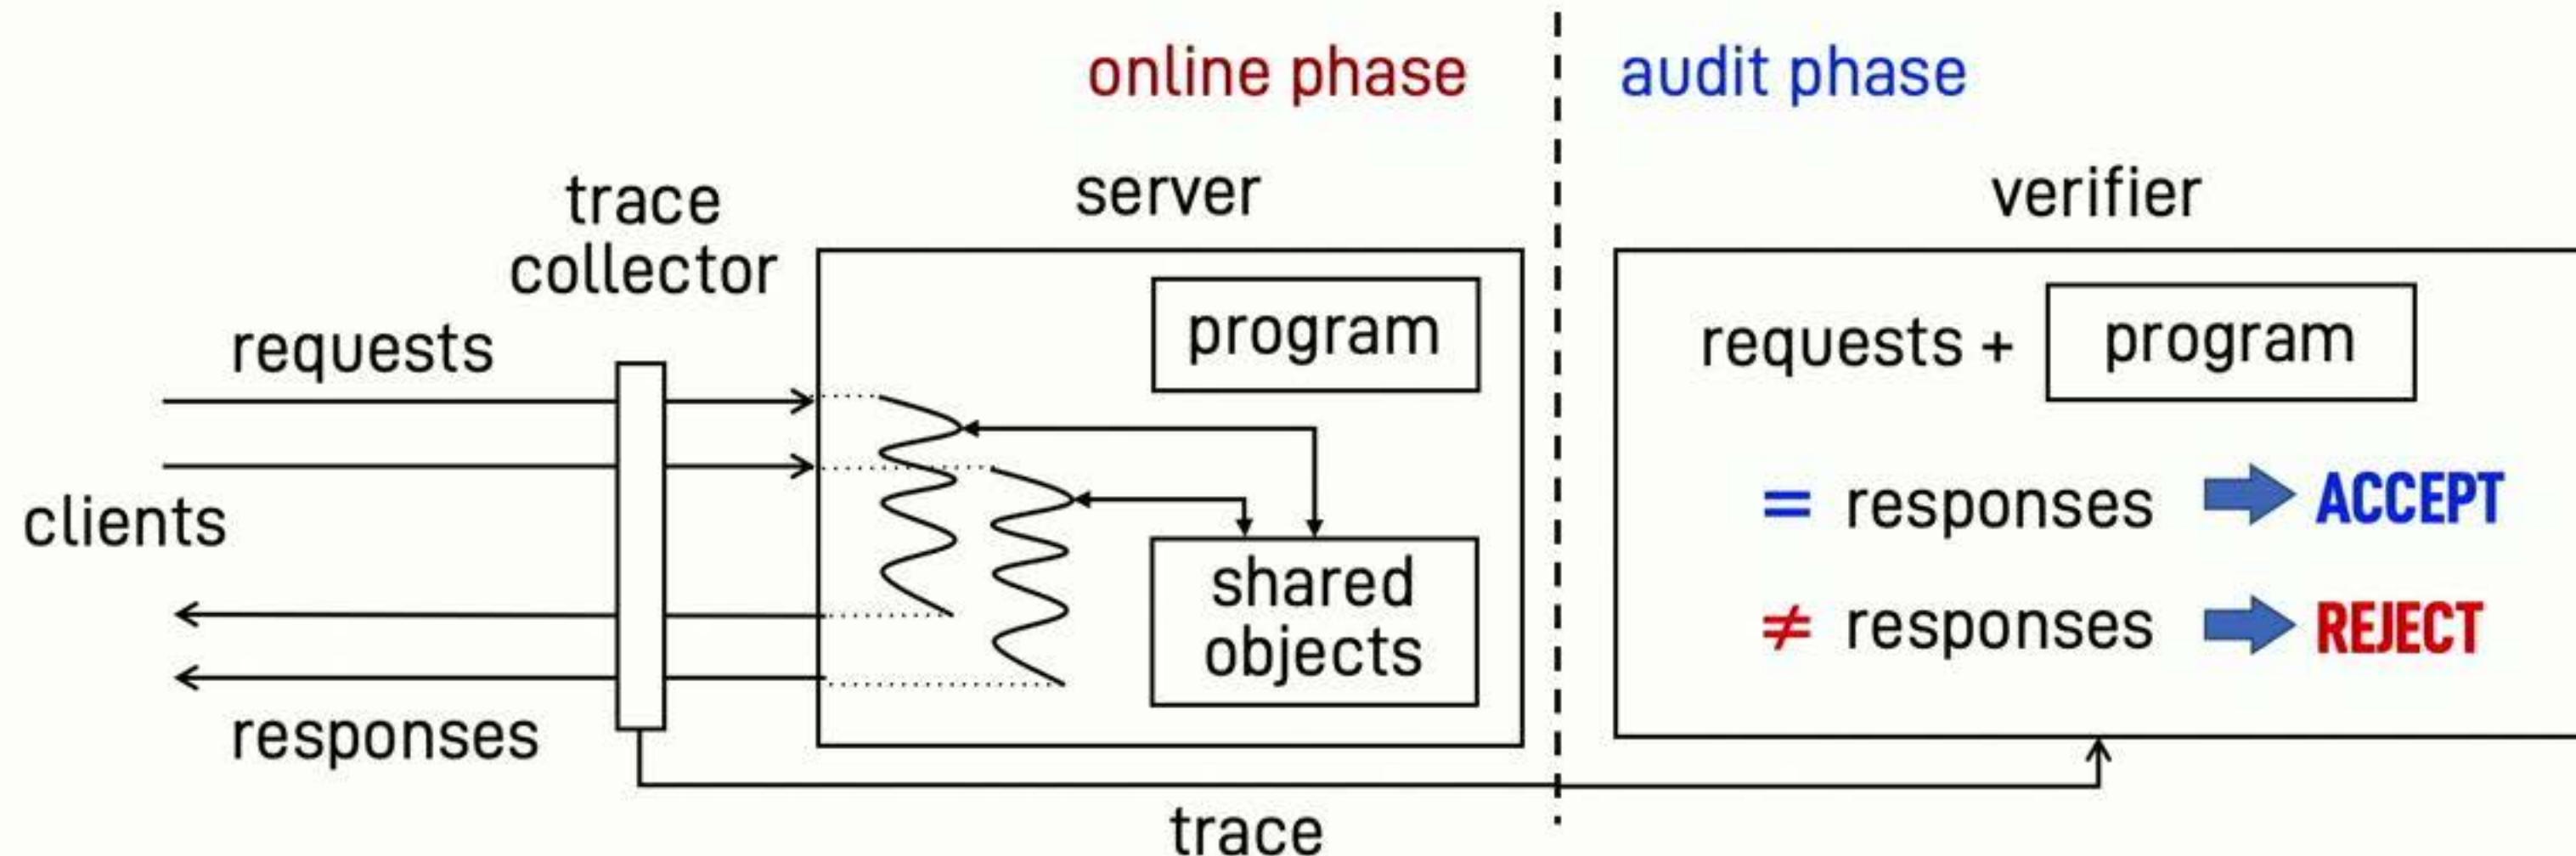
The Efficient Server Audit Problem



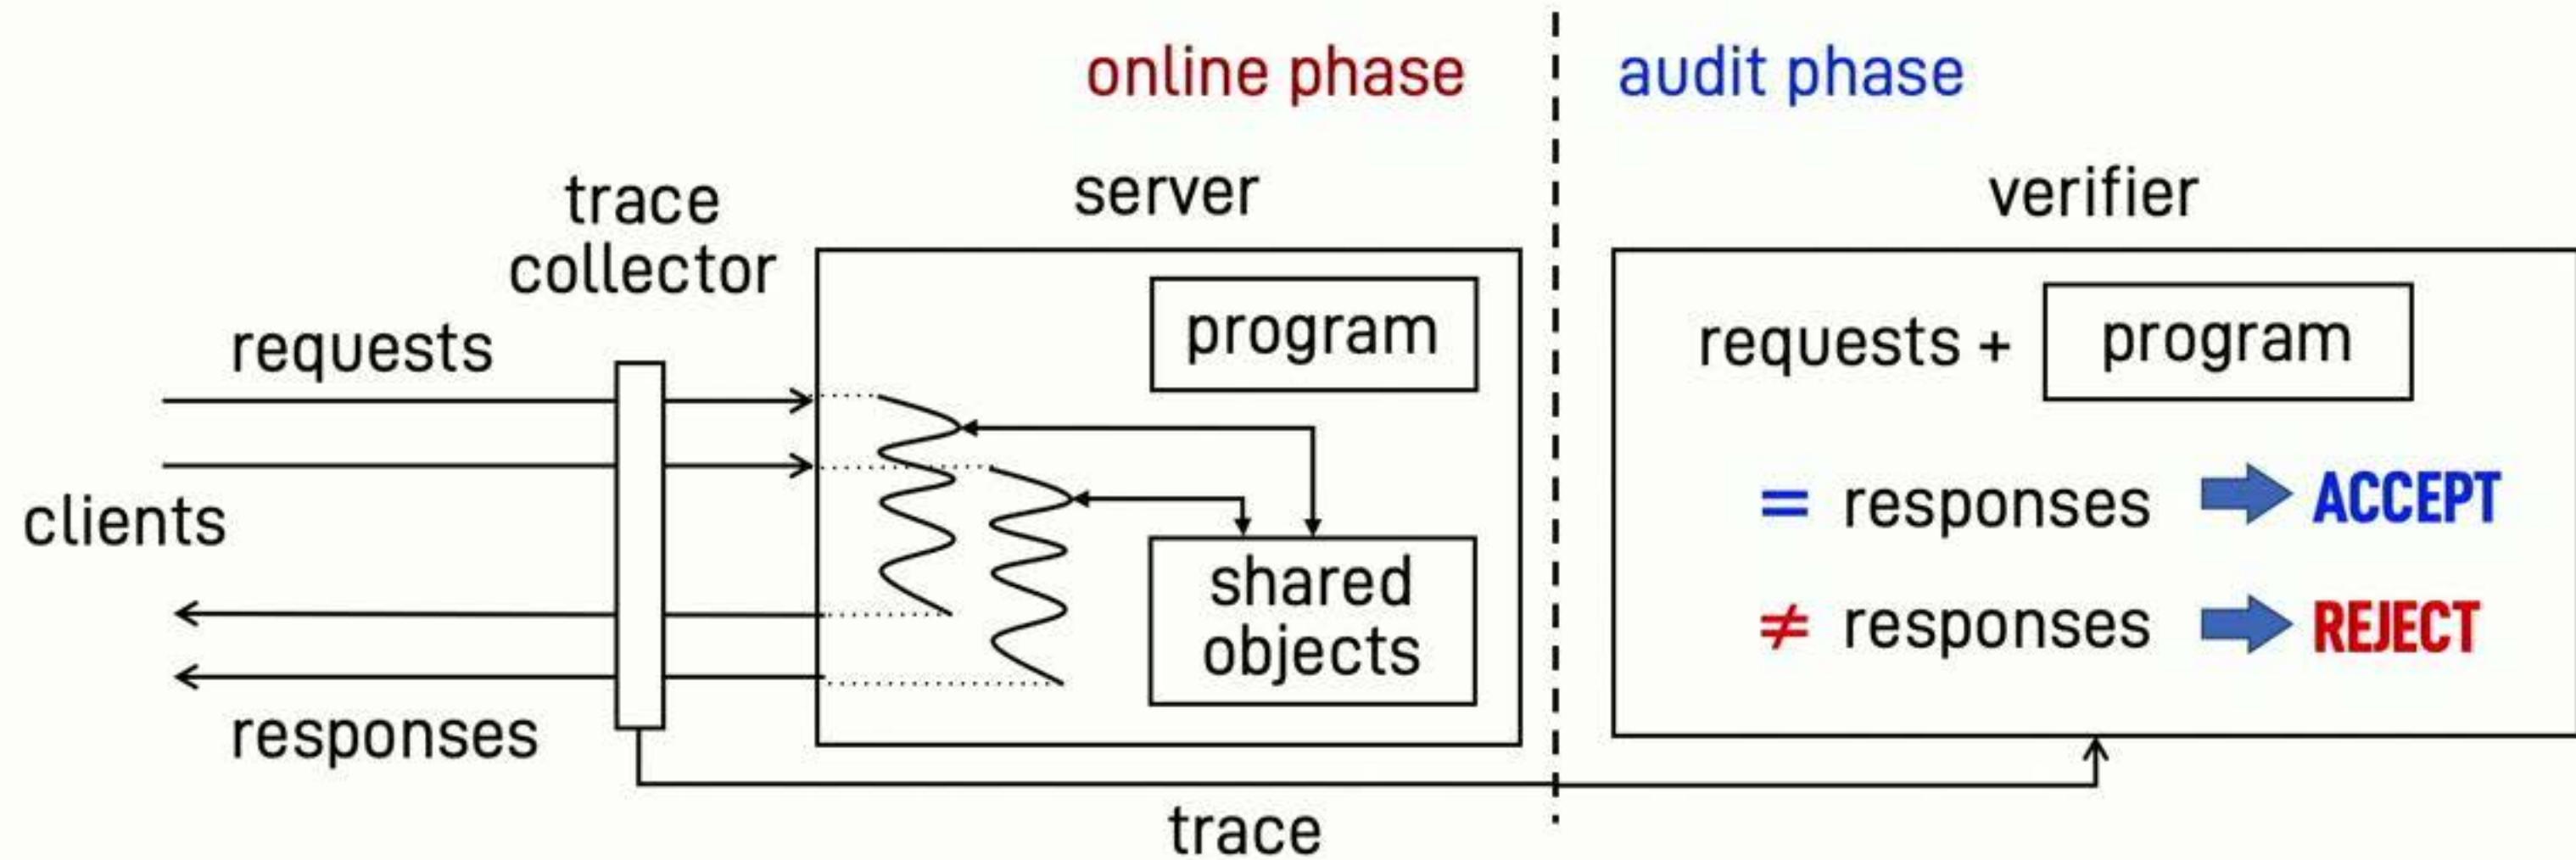
The Efficient Server Audit Problem



The Efficient Server Audit Problem

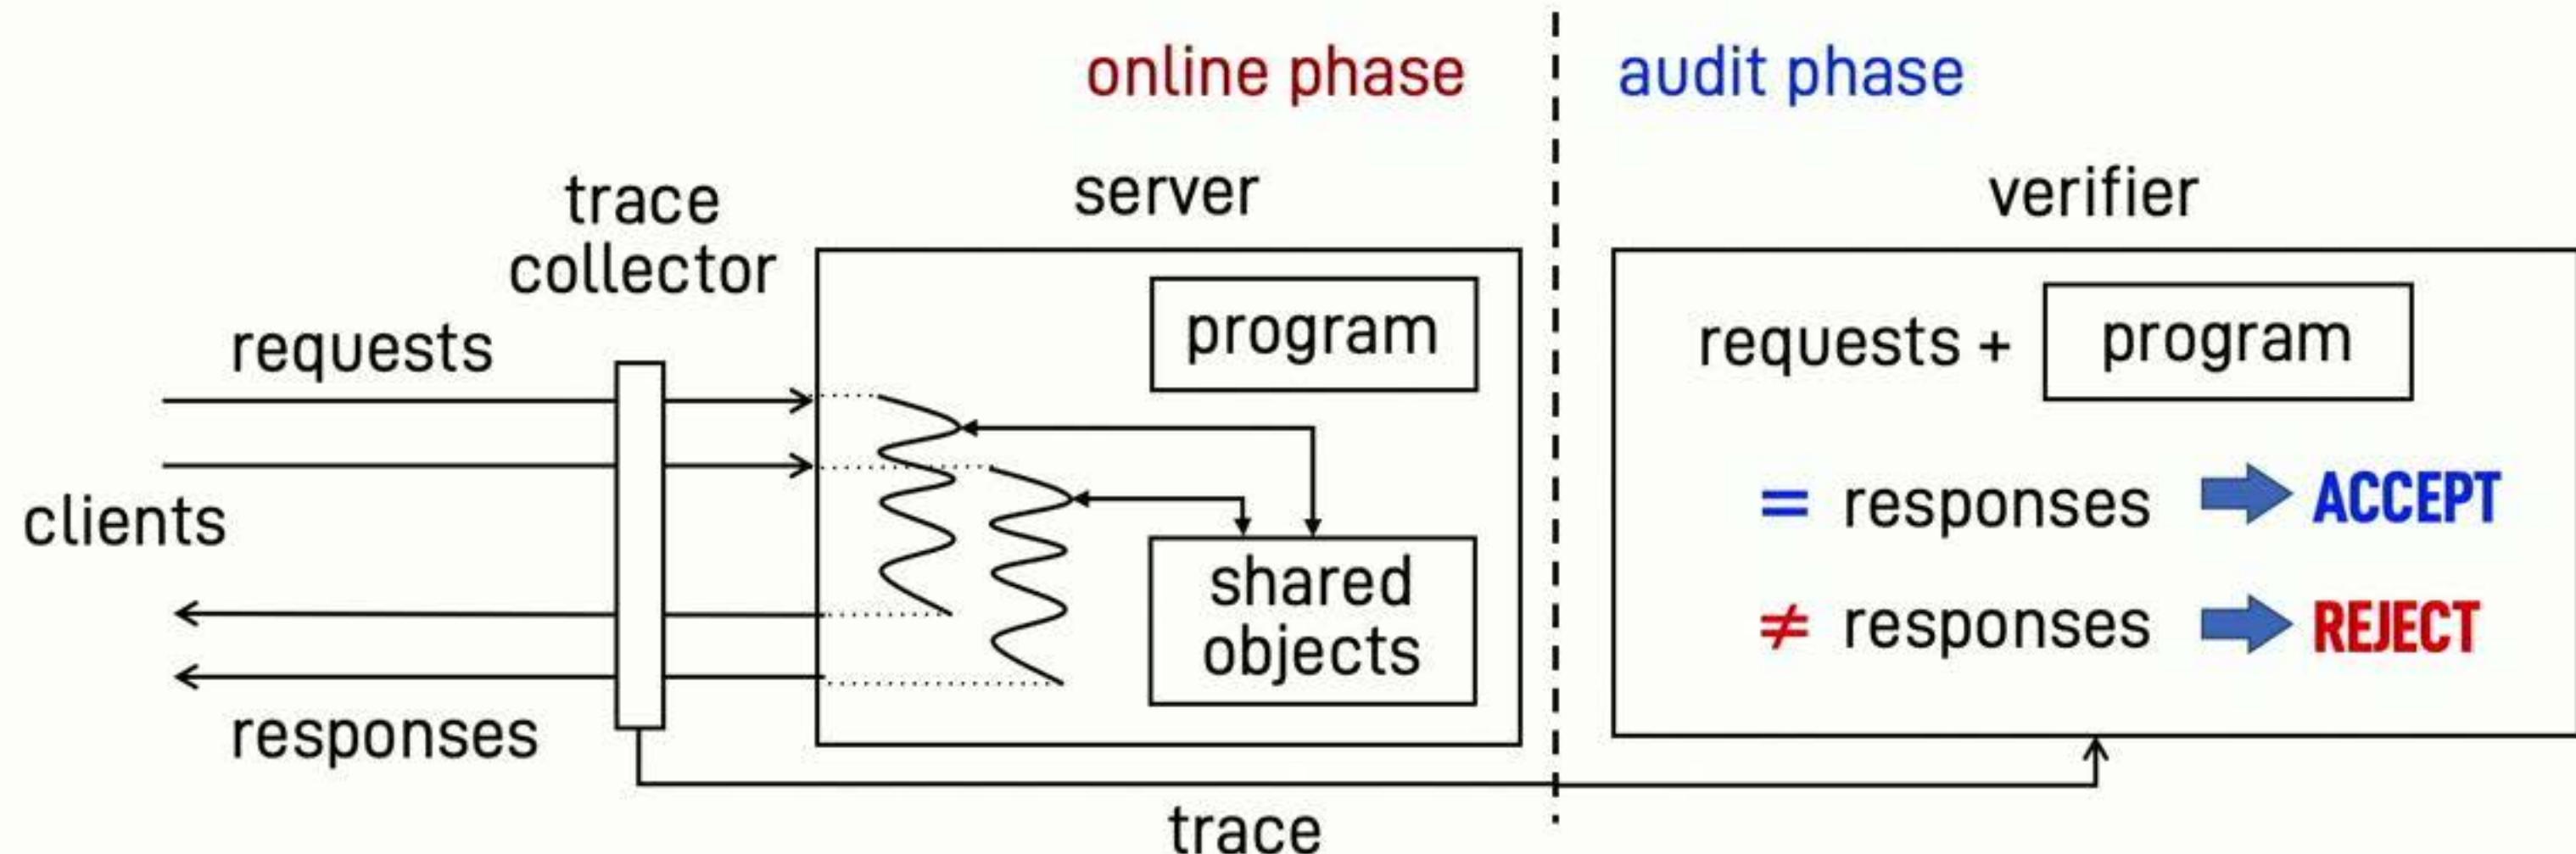


The Efficient Server Audit Problem



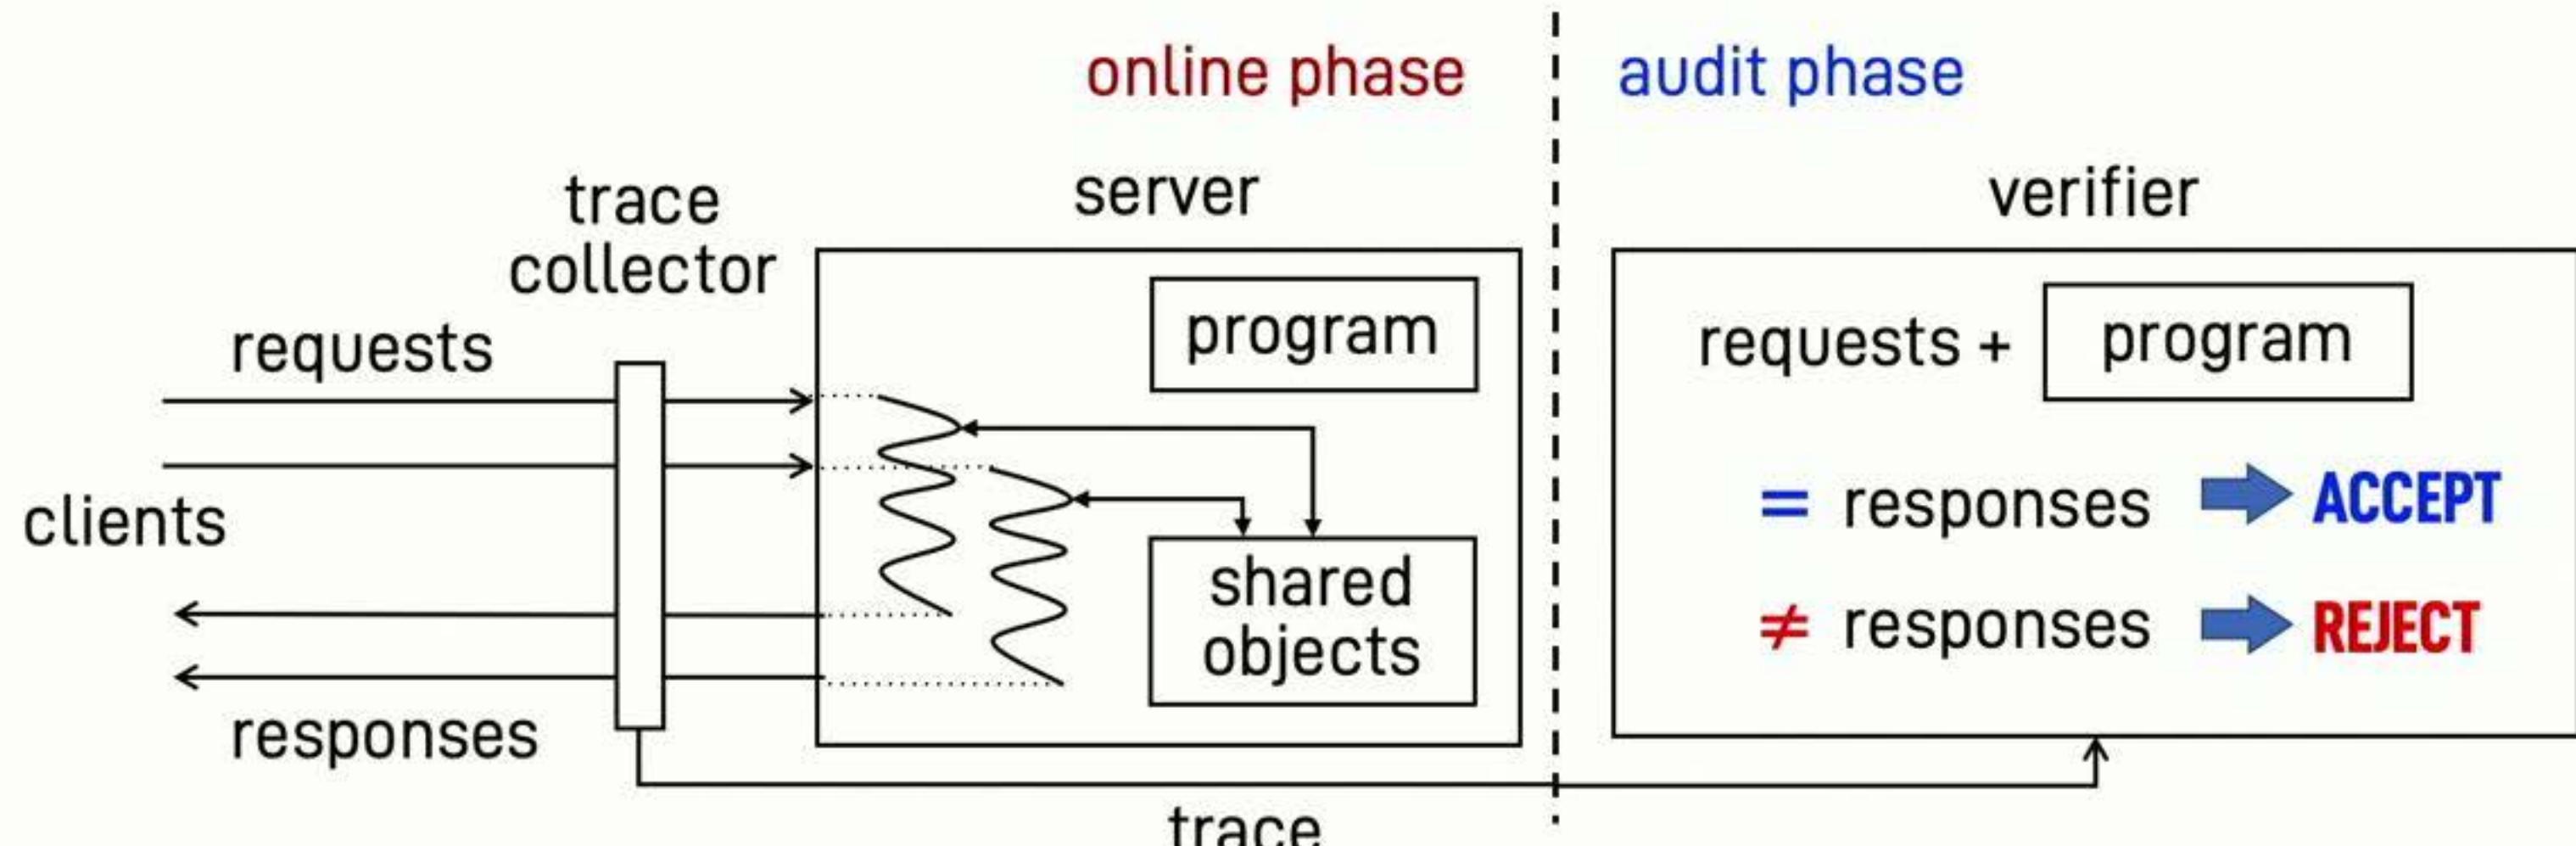
1. server can behave arbitrarily

The Efficient Server Audit Problem



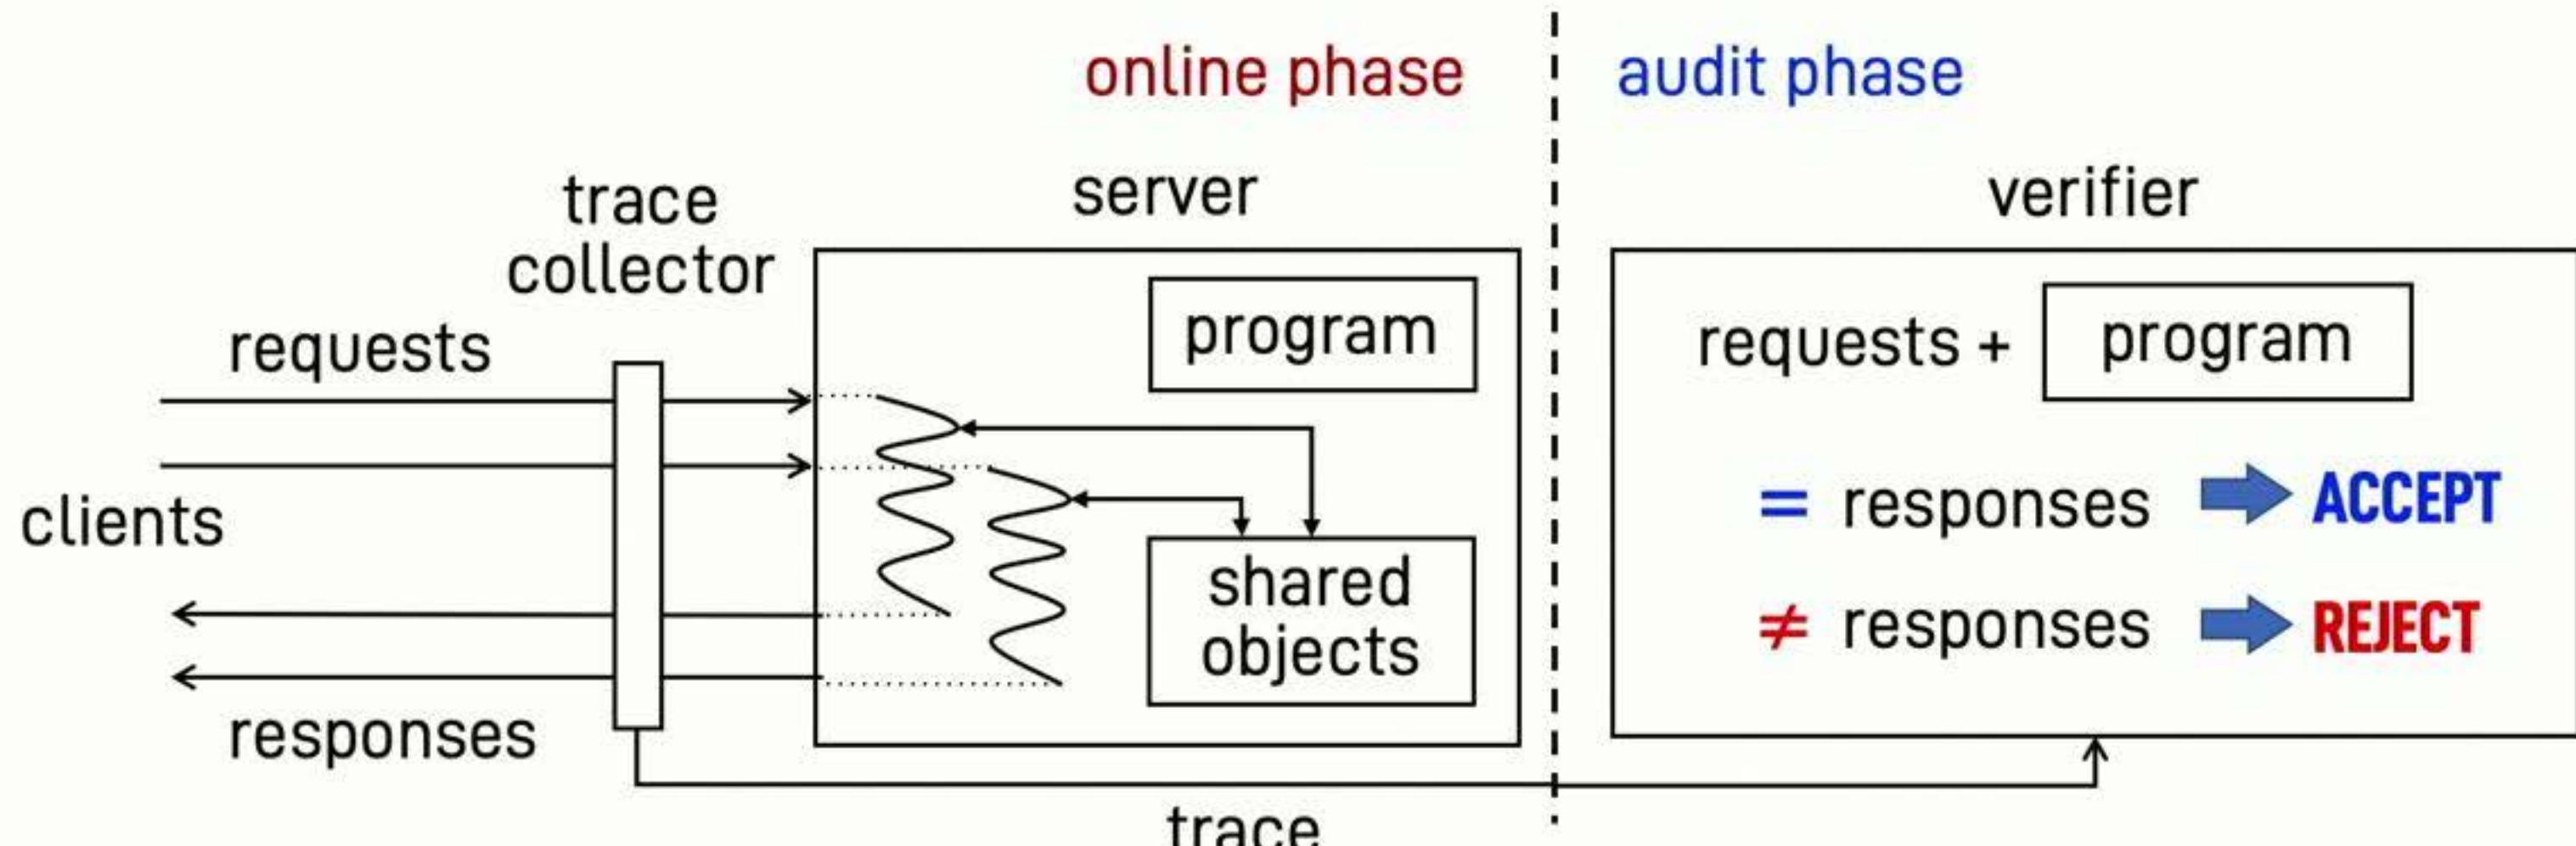
1. server can behave arbitrarily
2. server is concurrent

The Efficient Server Audit Problem



1. server can behave arbitrarily
2. server is concurrent
3. verifier requires less computation power than server

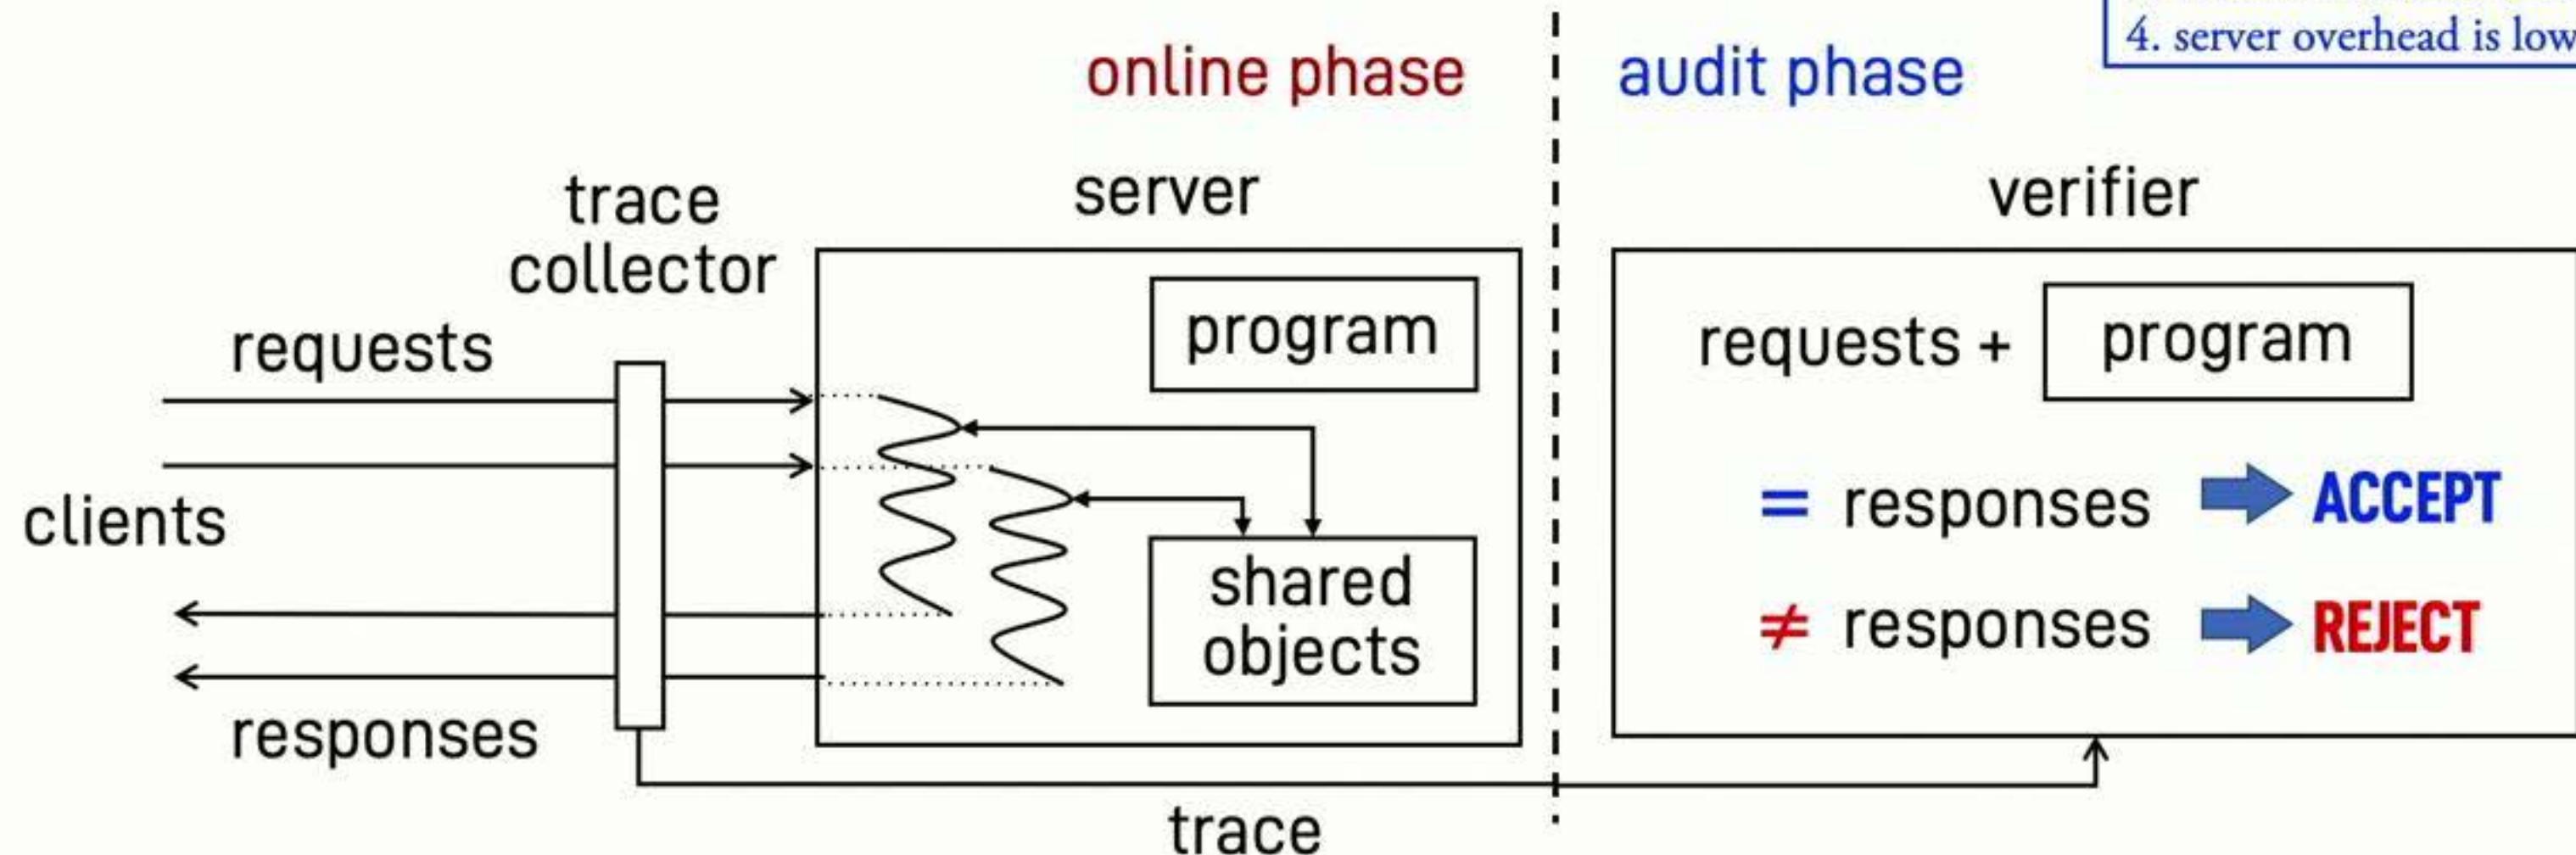
The Efficient Server Audit Problem



1. server can behave arbitrarily
2. server is concurrent
3. verifier requires less computation power than server
4. server overhead is low; legacy applications supported

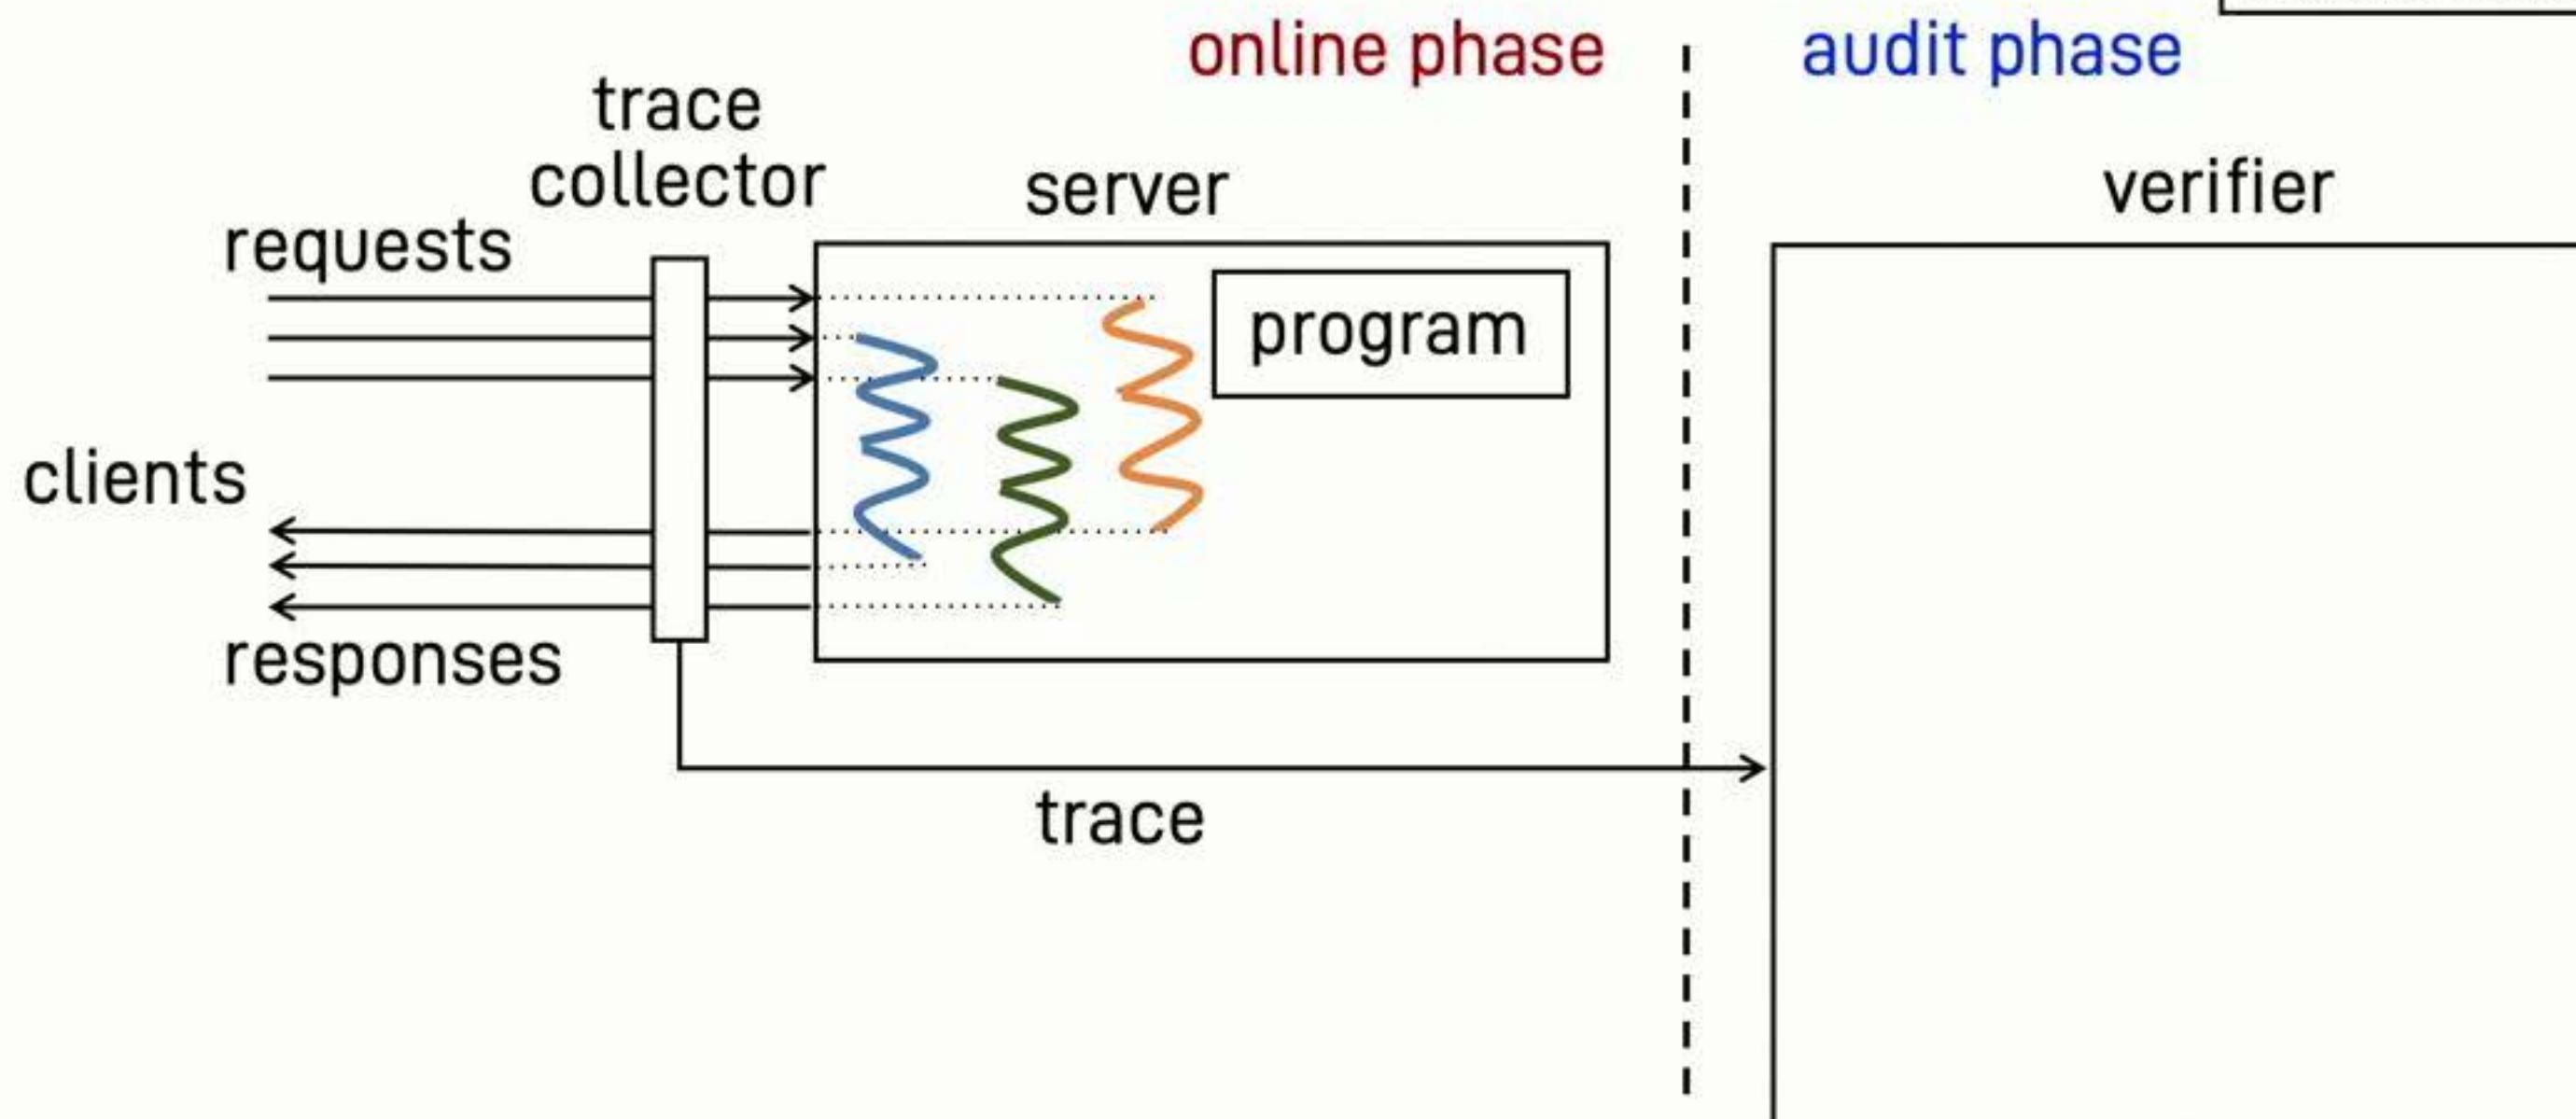
The Efficient Server Audit Problem

1. server is untrusted...
2. server is concurrent
3. verifier is weaker than server
4. server overhead is low...



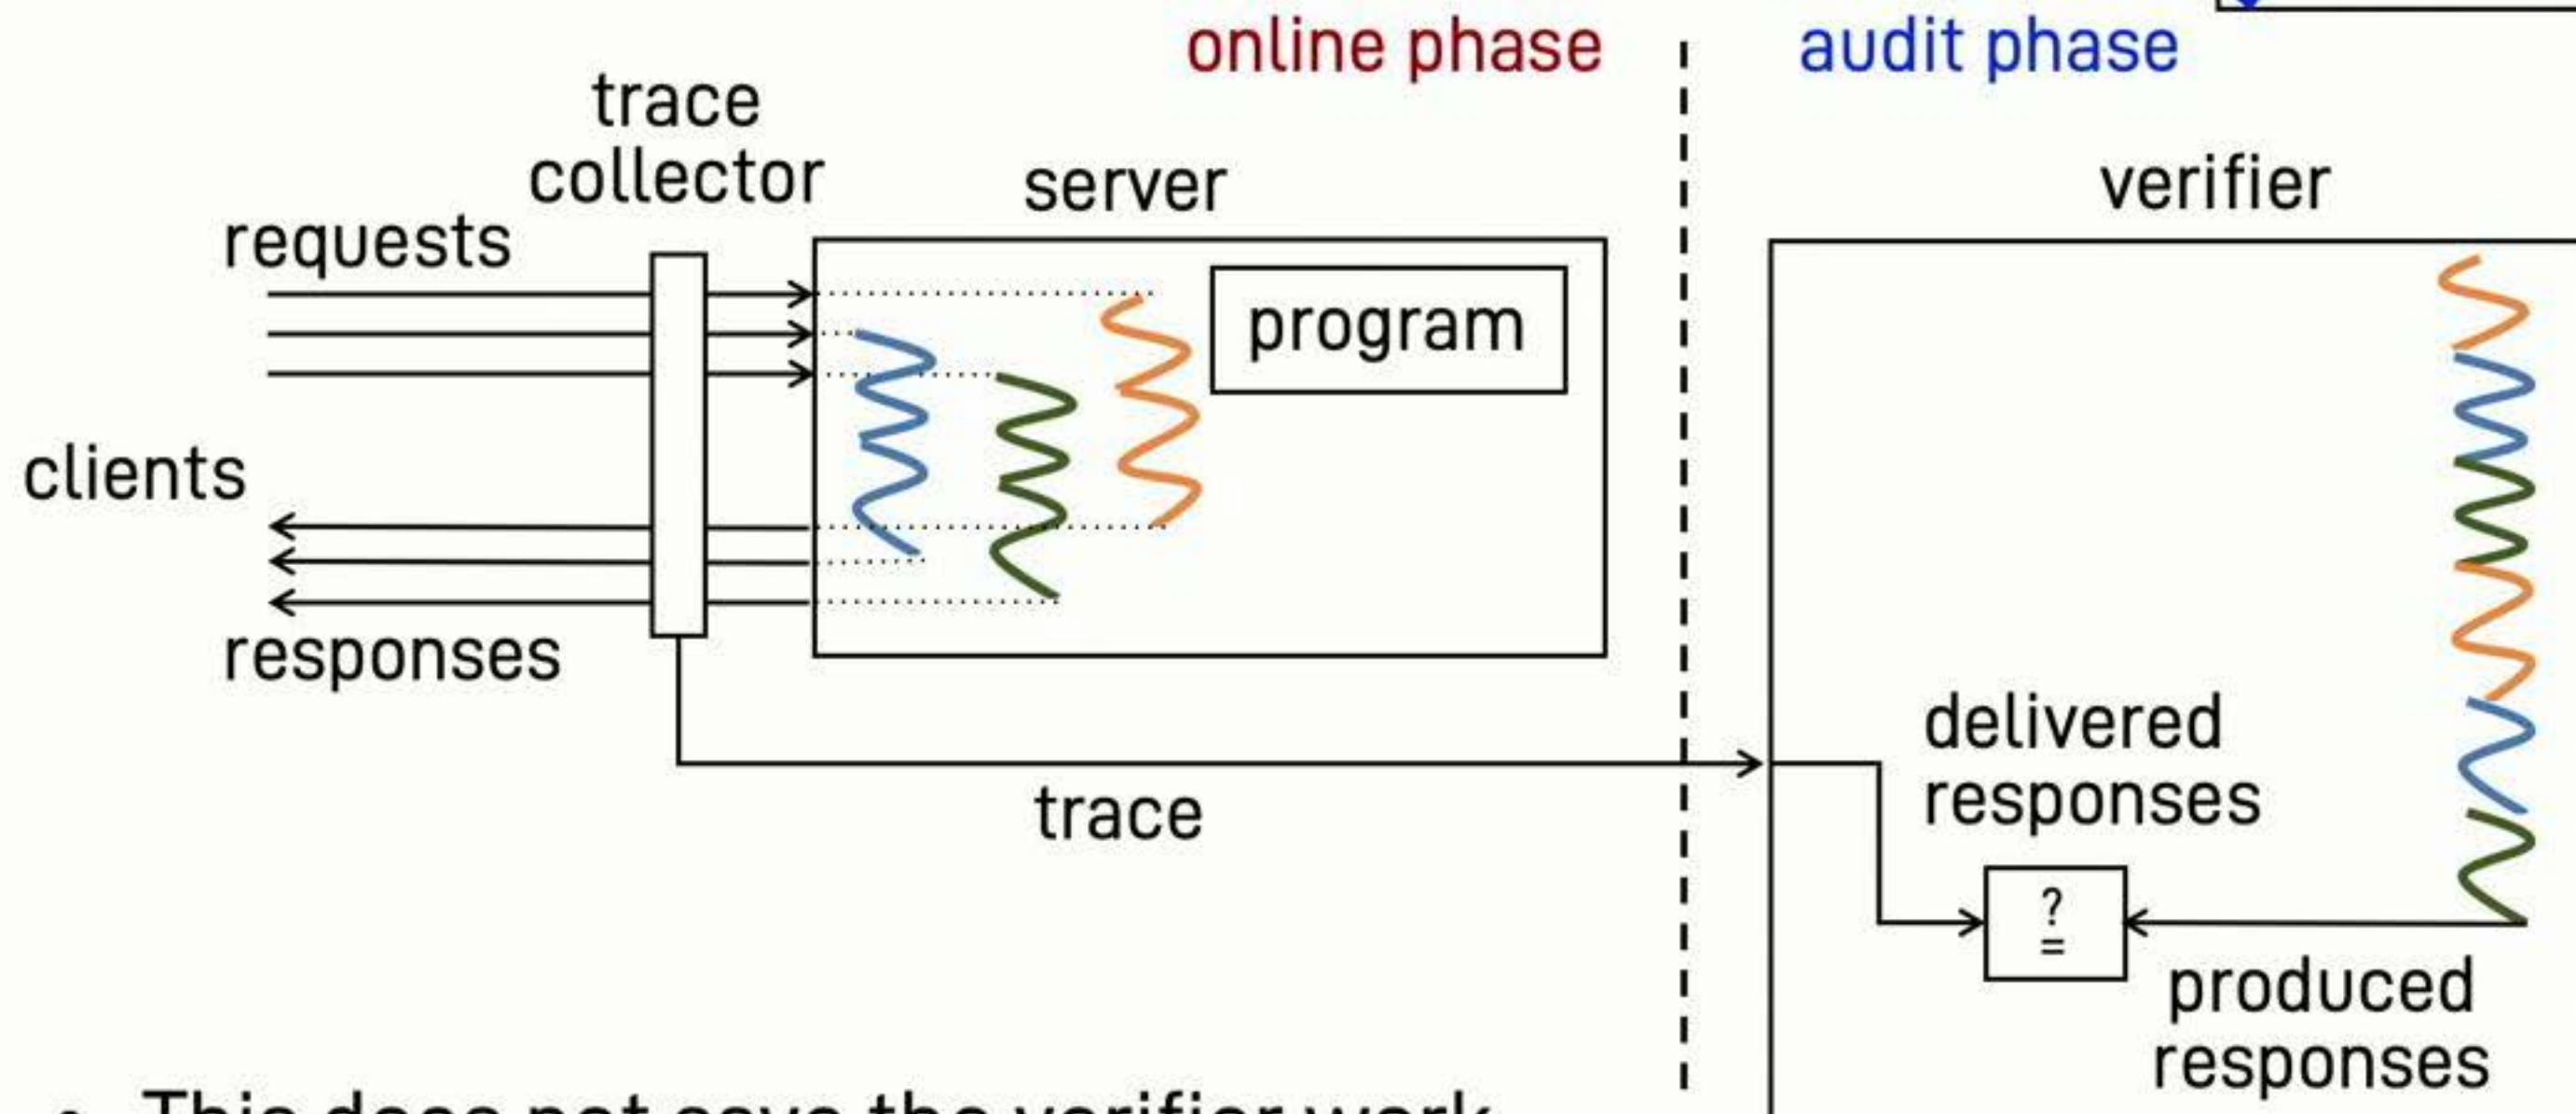
What about naive re-execution?

- 1. server is untrusted...
- 2. server is concurrent
- 3. verifier is weaker than server
- 4. server overhead is low...



What about naive re-execution?

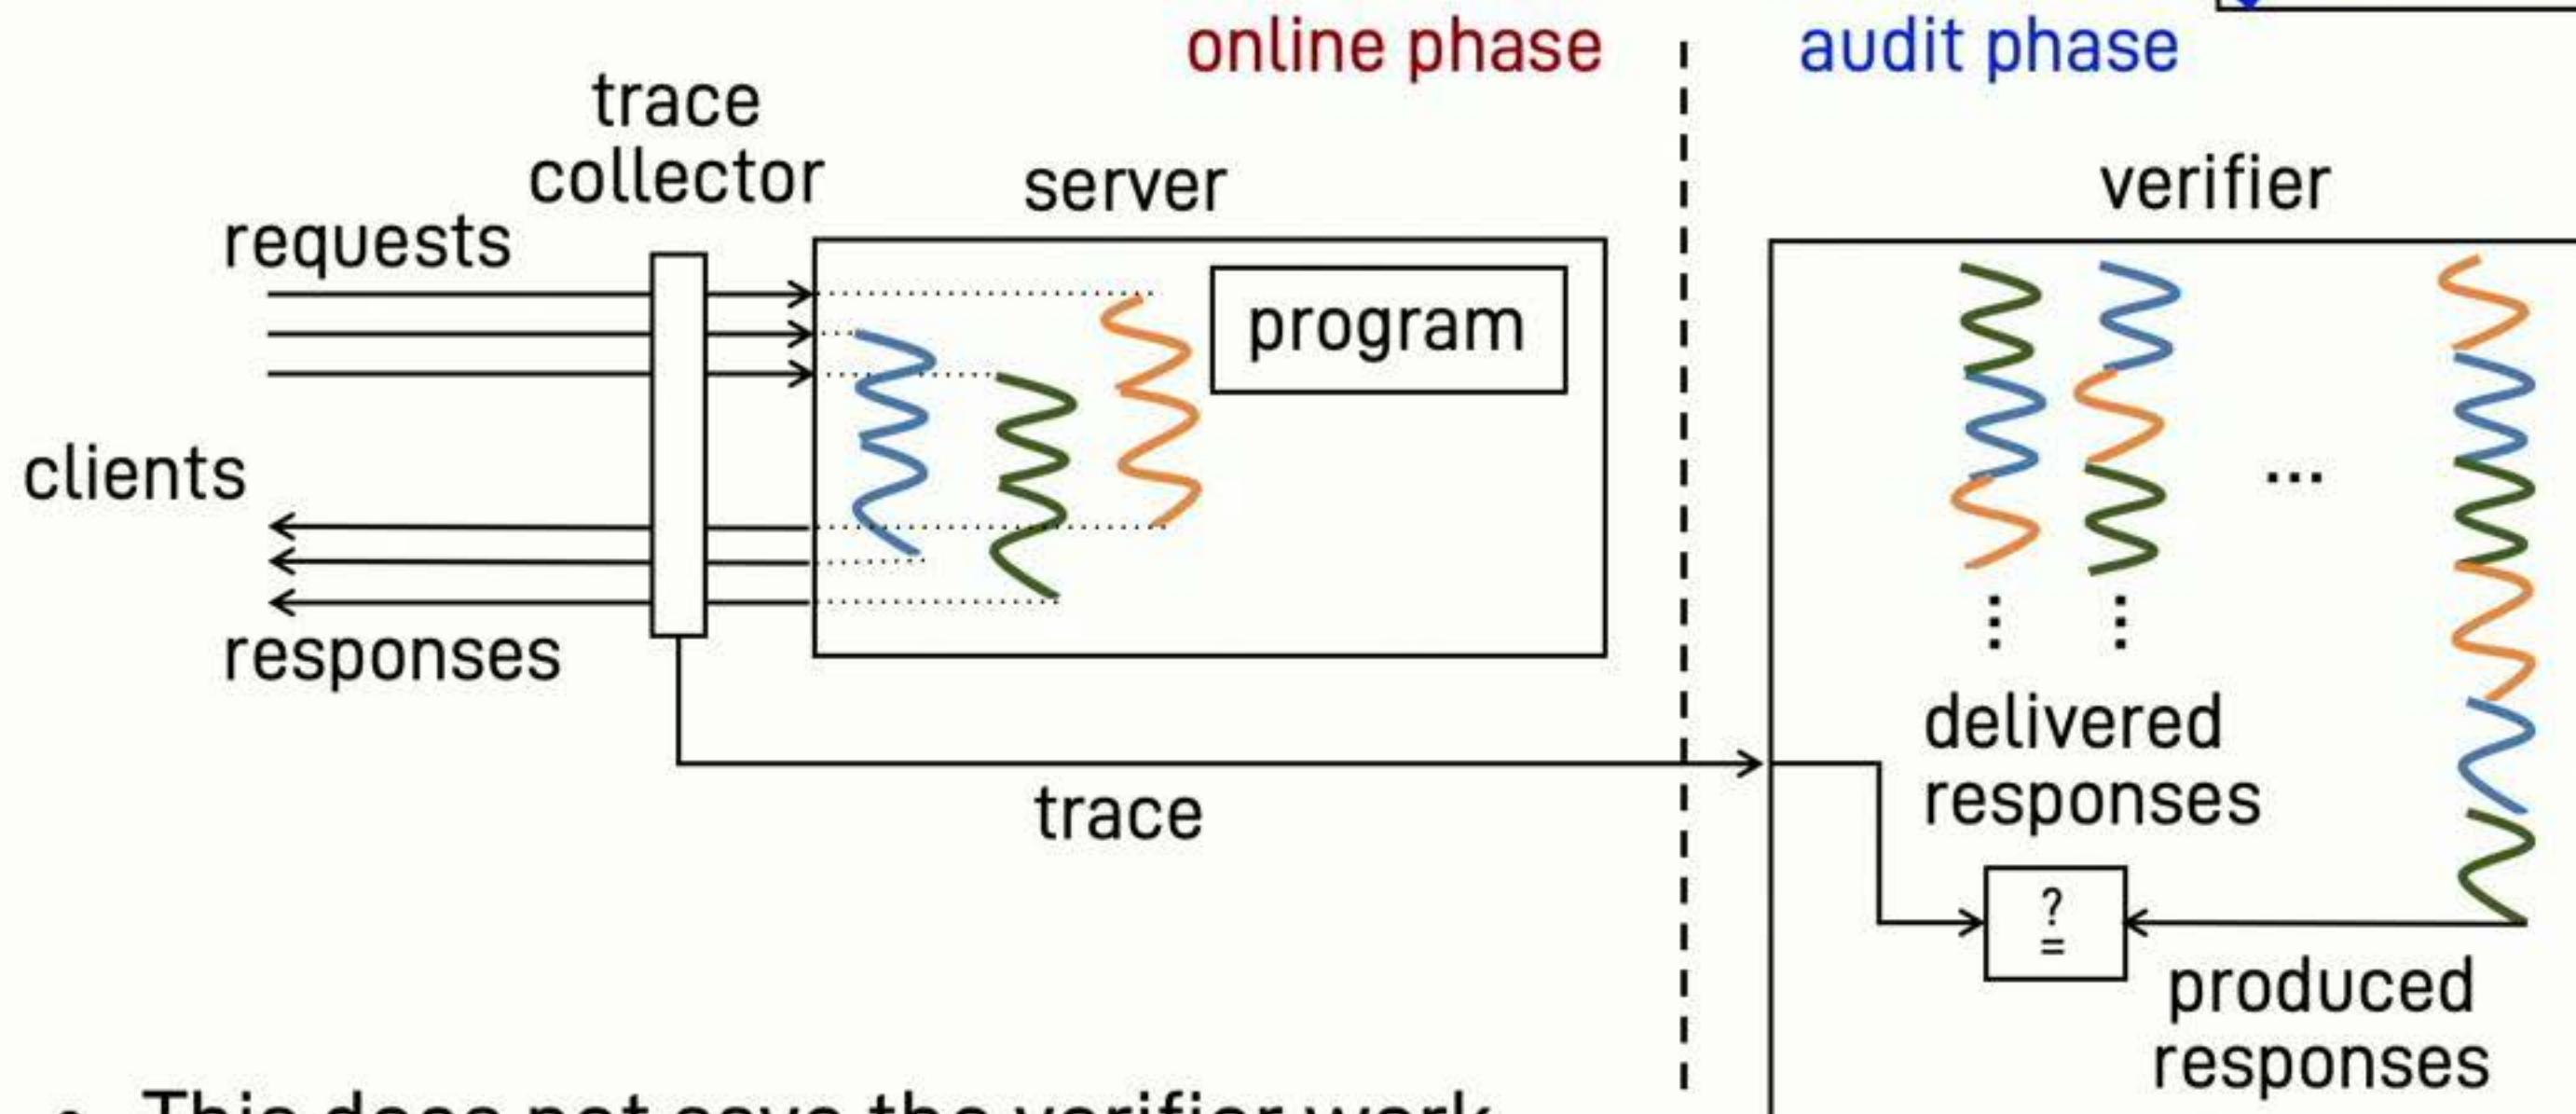
- ✓ server is untrusted...
- ? server is concurrent
- ✗ verifier is weaker than server
- ✓ server overhead is low...



- This does not save the verifier work.

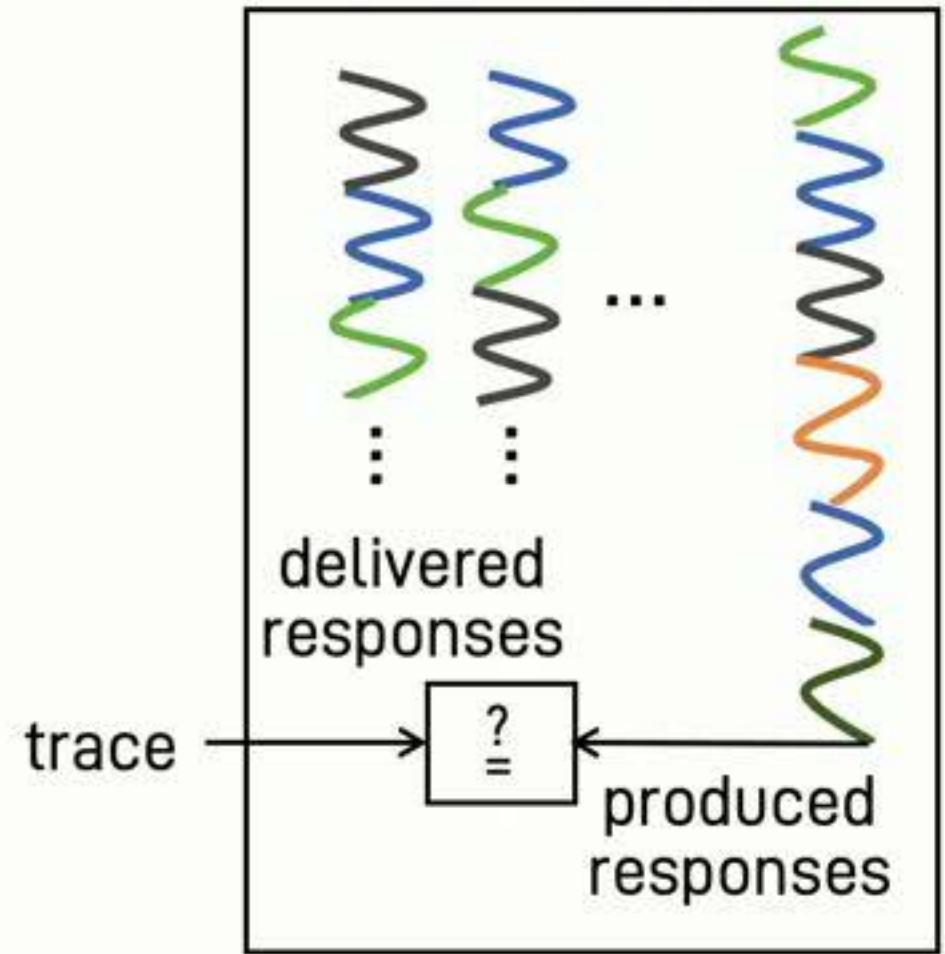
What about naive re-execution?

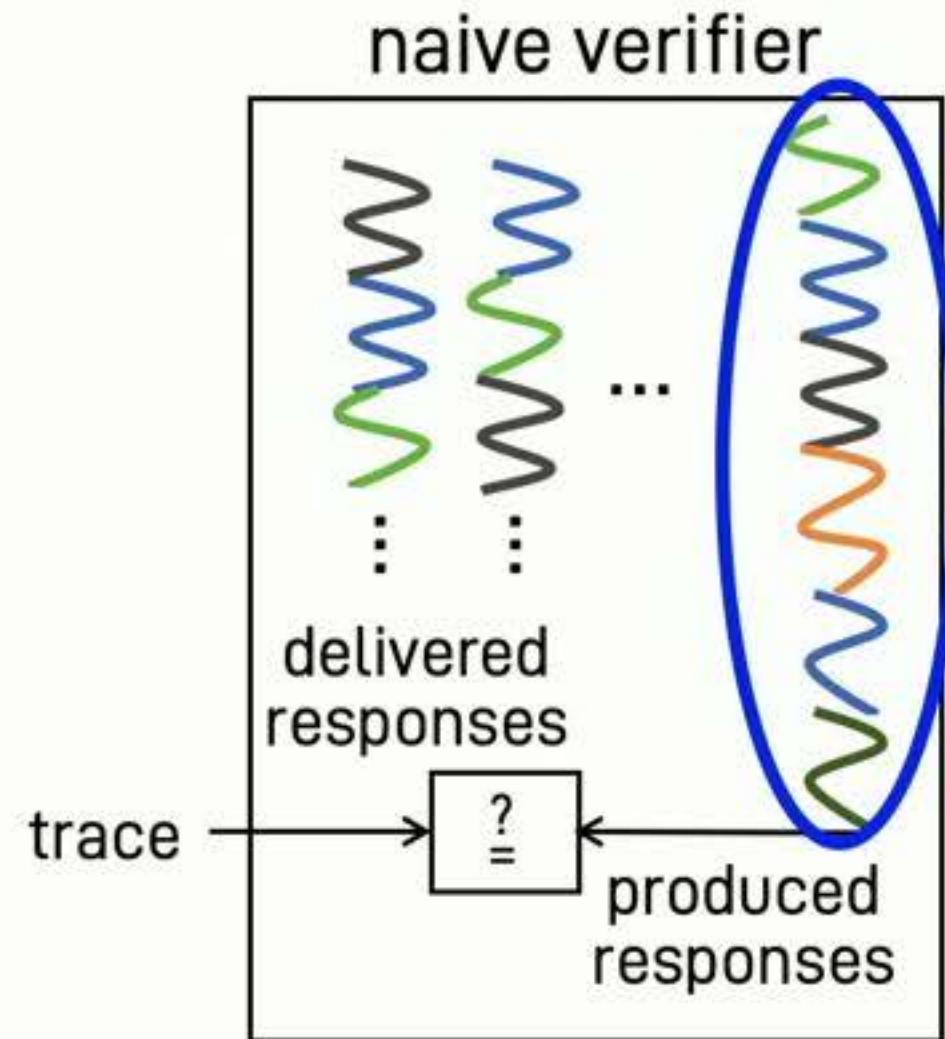
- ✓ server is untrusted...
- ? server is concurrent
- ✗ verifier is weaker than server
- ✓ server overhead is low...



- This does not save the verifier work.
- Due to concurrency, verifier must explore many schedules.

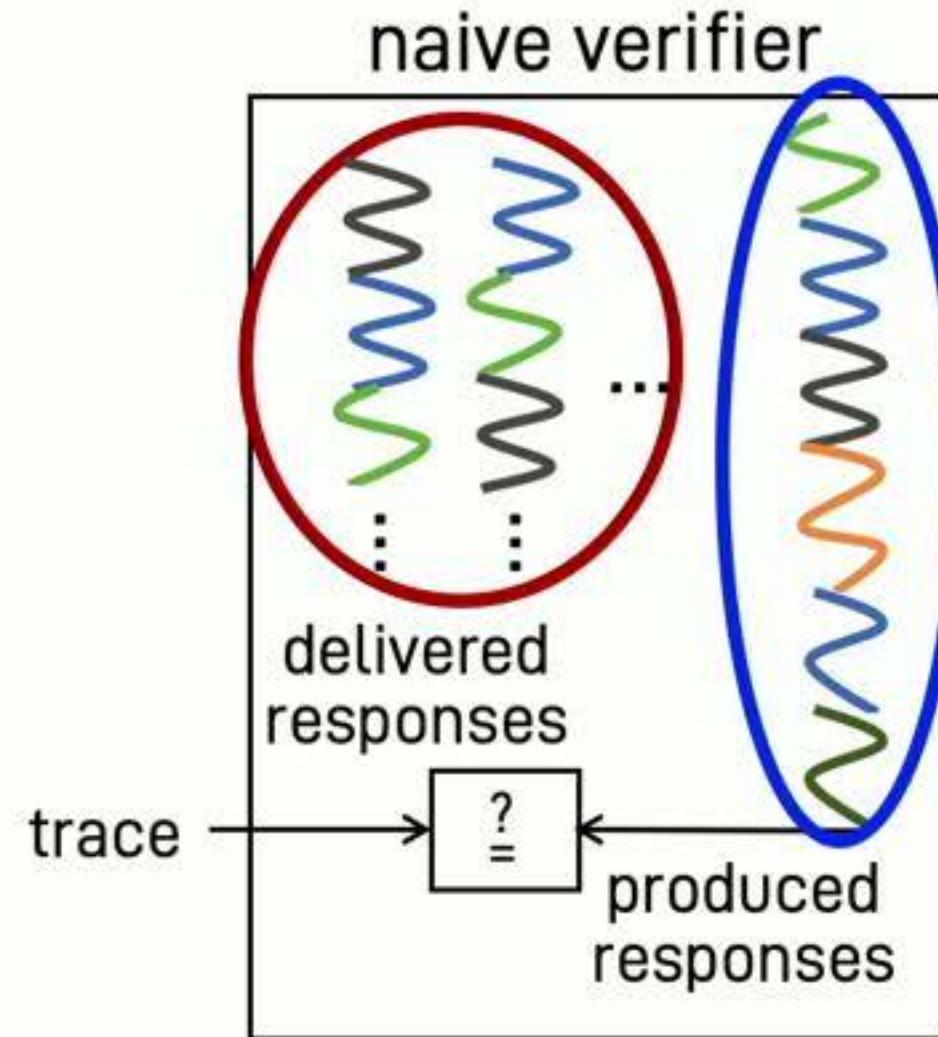
naive verifier





Problem: naive re-execution doesn't save work

Proposal: (somehow) accelerate re-execution

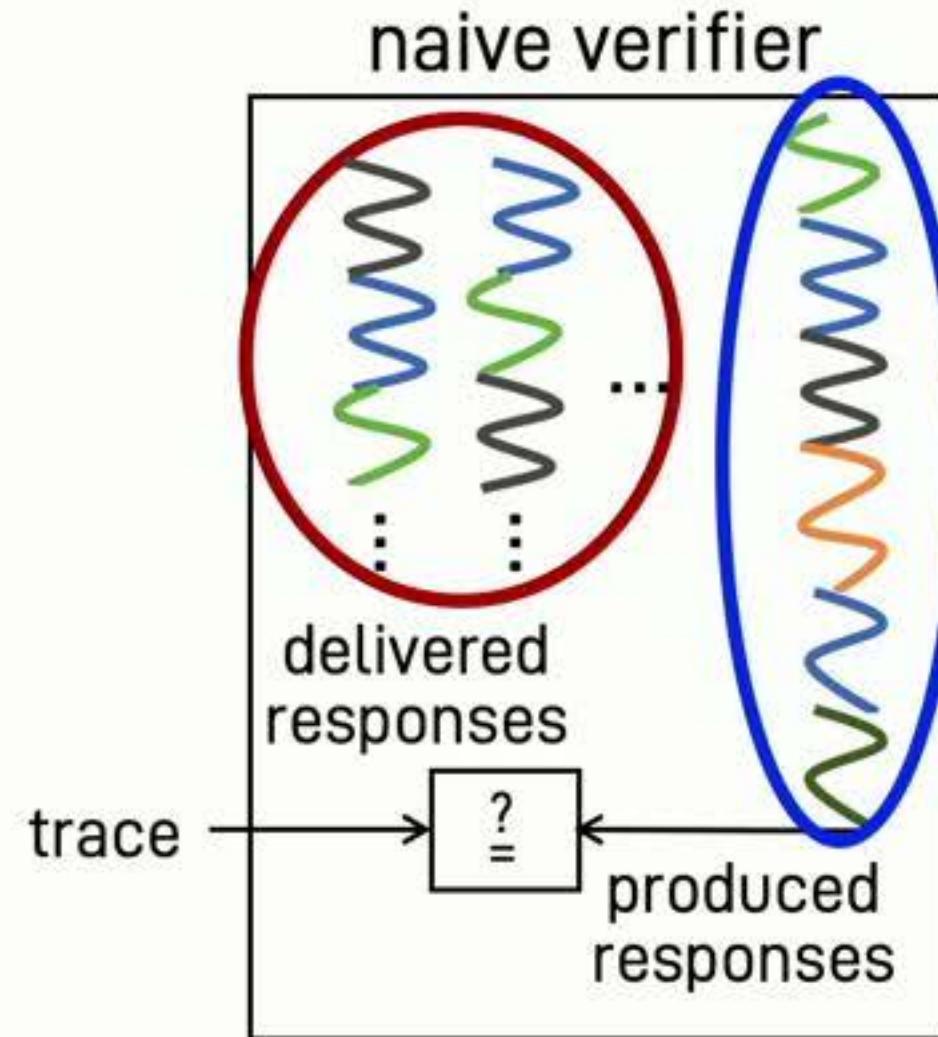


Problem: naive re-execution doesn't save work

Proposal: (somehow) accelerate re-execution

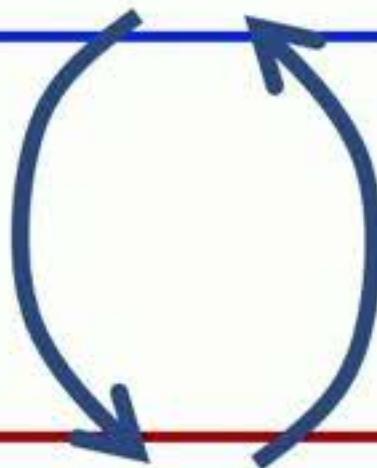
Problem: many schedules to explore

Proposal: ask server for advice, use-and-check it



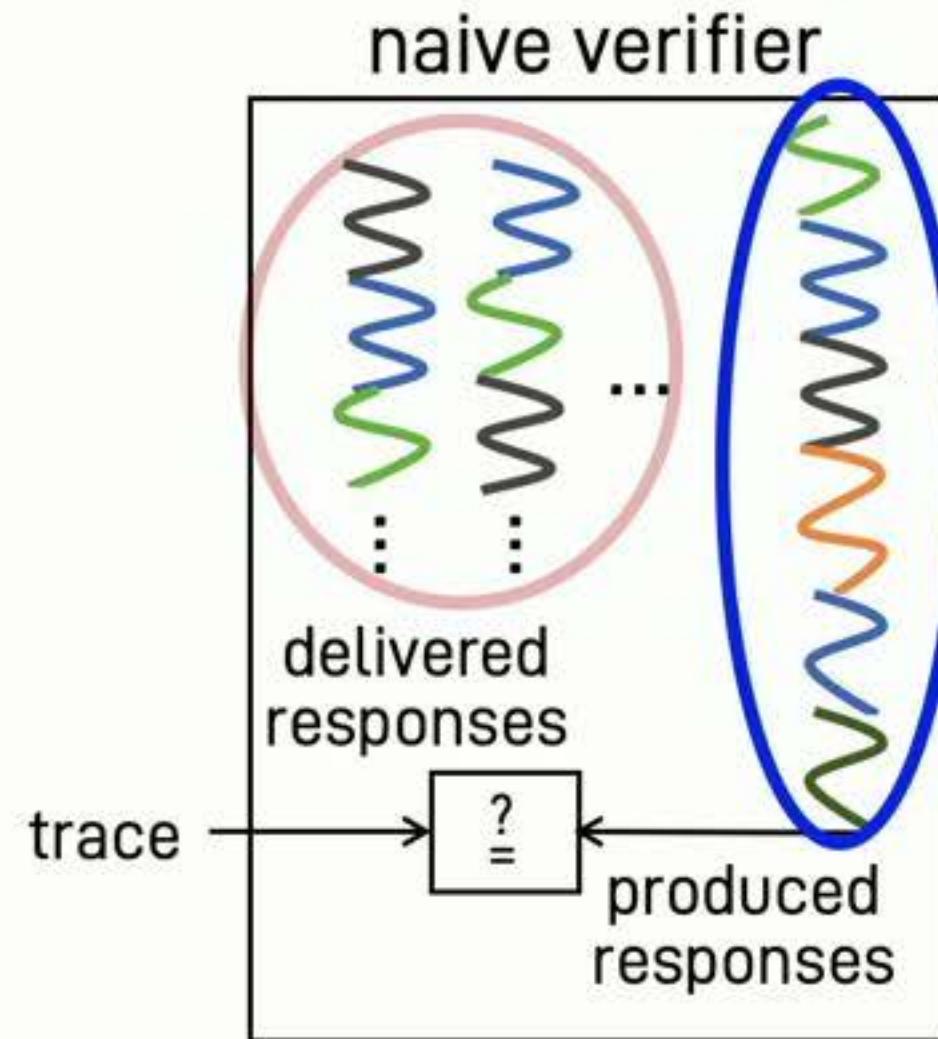
Problem: naive re-execution doesn't save work

Proposal: (somehow) accelerate re-execution



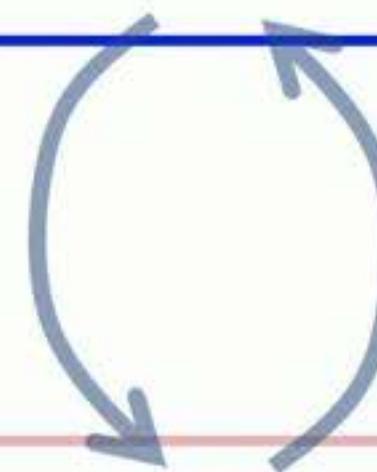
Problem: many schedules to explore

Proposal: ask server for advice, use-and-check it



Problem: naive re-execution doesn't save work

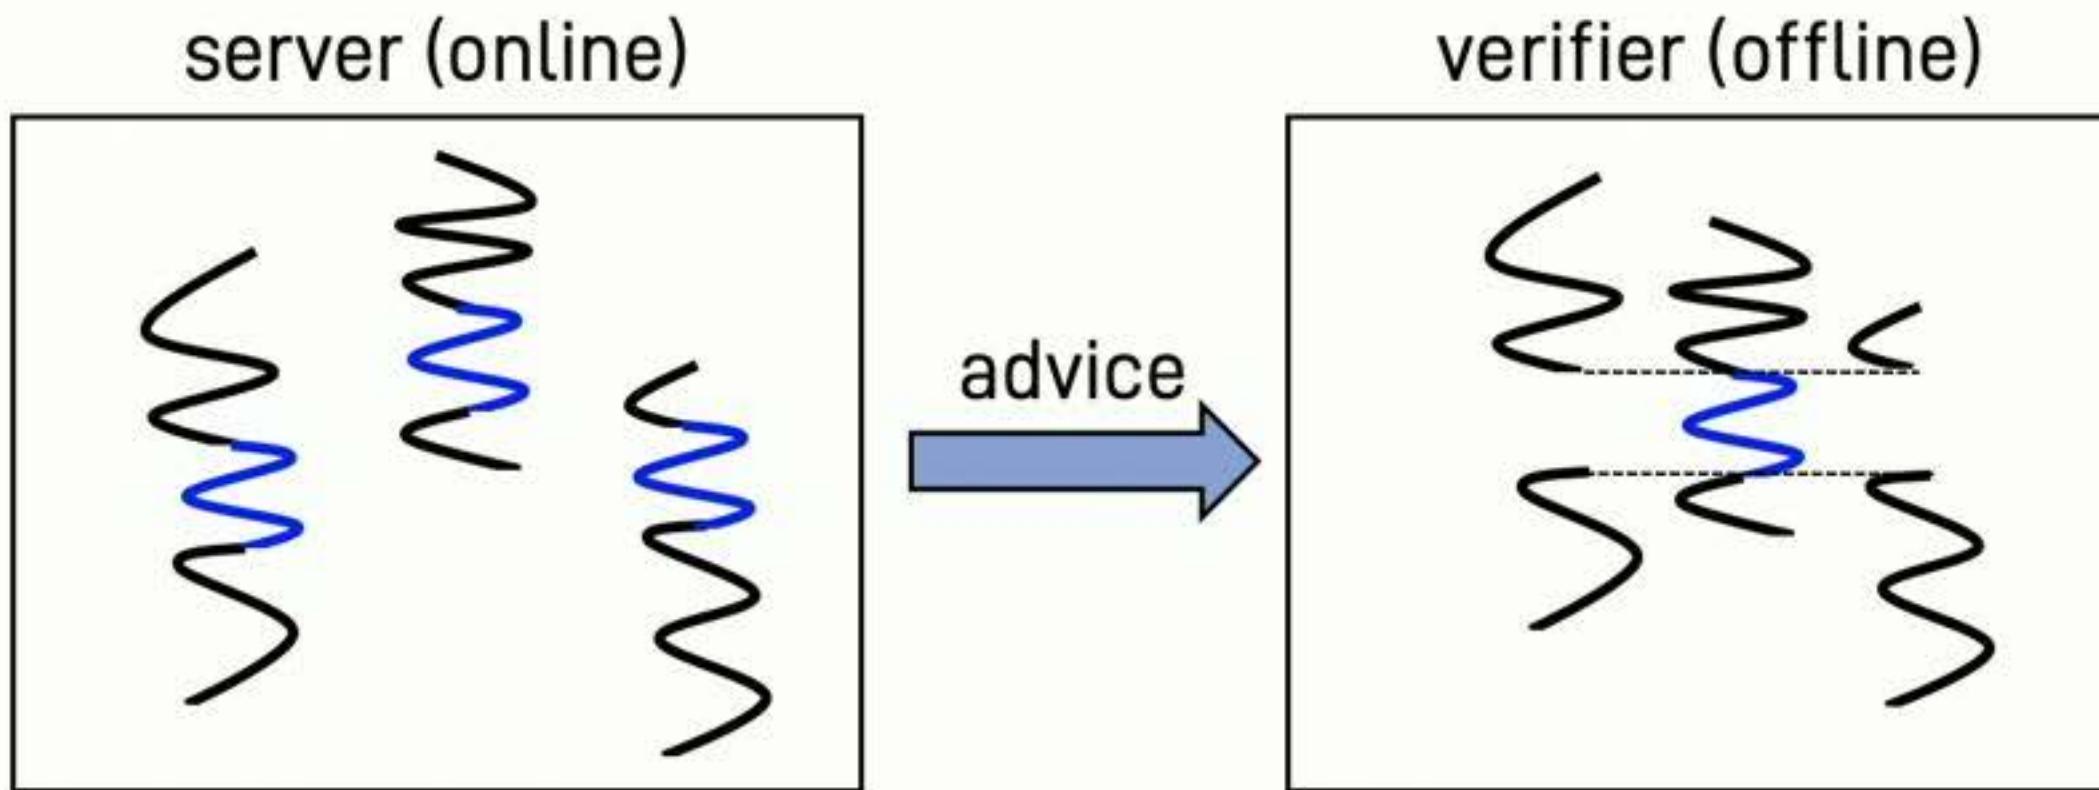
Proposal: (somehow) accelerate re-execution



Problem: many schedules to explore

Proposal: ask server for advice, use-and-check it

Accelerating re-execution: a 30,000-foot view



deduplicate computation across requests

An observation: repeated computation

web app:
HotCRP

page 1

Title
My paper

Submission (PDF, max 100MB)

 400kB 9 Oct 2017 1:59:08pm EDT

Replace: No file chosen

Authors
List the authors one per line, including email address. Reviewers will not be able to see author information. Any authors listed here will be able to edit the submission.

Name
1. Lingfan Yu
2.

page 2

Title
Another Paper

Submission (PDF, max 100MB)

 400kB 9 Oct 2017 12:56:56pm EDT

Replace: No file chosen

Authors
List the authors one per line, including email address. Reviewers will not be able to see author information. Any authors listed here will be able to edit the submission.

Name
1. Cheng Tan
2.

An observation: repeated computation

web app:
HotCRP

page 1

Title
My paper

Submission (PDF, max 100MB)
 400kB 9 Oct 2017 1:59:08pm EDT

Replace: No file chosen

Authors
List the authors one per line, including email address. Reviewers will not be able to see author information. Any authors listed here will be able to edit the submission.

Name
1. Lingfan Yu

page 2

Title
Another Paper

Submission (PDF, max 100MB)
 400kB 9 Oct 2017 12:56:56pm EDT

Replace: No file chosen

Authors
List the authors one per line, including email address. Reviewers will not be able to see author information. Any authors listed here will be able to edit the submission.

Name
1. Cheng Tan

An observation: repeated computation

web app:
HotCRP

page 1

Title
My paper

Submission (PDF, max 100MB)
 400kB ⌂ 9 Oct 2017 1:59:08pm EDT

Replace: No file chosen

Authors
List the authors one per line, including email
not be able to see author information. Any au

Name
1. Lingfan Yu
2.

A black outline of a right-pointing arrow, indicating the direction of the next section.

```
[4228] RetC      O { /home/cheng/orochi
[4229] PopC      O { /home/cheng/orochi
[4230] FPushClsMethodD  O { /home/chi
[4231] FCall      O { /home/cheng/orochi
[4232] String     (Navigation::page) { .
[4233] String     (Navigation::page) { .
[4234] AGetC      (Navigation::page) { /.
[4235] CGetS      (Navigation::page) { /.
[4236] RetC      (Navigation::page) { /h
[4237] UnboxR    O { /home/cheng/orochi
[4238] String     O { /home/cheng/orochi
[4239] NSame      O { /home/cheng/orochi
[4240] JmpZ      O { /home/cheng/orochi
[4241] FPushClsMethodD  O { /home/chi
[4242] FCall      O { /home/cheng/orochi
[4243] String     (Navigation::page) { .
[4244] String     (Navigation::page) { .
```

page 2

Title

Another Paper

Submission (PDF, max 100MB)

 400kB © 9 Oct 2017 12:56:56pm EDT

Replace: Choose File No file chosen

Authors

List the authors one per line, including email address. Reviewers will not be able to see author information to edit the submission.

Name

1. Cheng Tan

2.

A black arrow pointing to the right, indicating the direction of the next step in the process.

```
[4228] RetC      O { /home/cheng/orochi
[4229] PopC      O { /home/cheng/orochi
[4230] FPushClsMethodD  O { /home/ch
[4231] FCall      O { /home/cheng/orochi
[4232] String      (Navigation::page) { .
[4233] String      (Navigation::page) { .
[4234] AGetC      (Navigation::page) { /h
[4235] CGetS      (Navigation::page) { /h
[4236] RetC      (Navigation::page) { /h
[4237] UnboxR      O { /home/cheng/orochi
[4238] String      O { /home/cheng/orochi
[4239] NSame      O { /home/cheng/orochi
[4240] JmpZ      O { /home/cheng/orochi
[4241] FPushClsMethodD  O { /home/ch
[4242] FCall      O { /home/cheng/orochi
[4243] String      (Navigation::page) { .
[4244] String      (Navigation::page) { .
```

have the same control flow

An observation: repeated computation

web app:
HotCRP

page 1

Title

Submission (PDF, max 100MB)
 400kB 9 Oct 2017 1:59:08pm EDT ·

Replace:

Authors
 List the authors one per line, including email
 not be able to see author information. Any au

Name

1.

2.

A small black arrow pointing to the right, indicating the direction of the next step in the sequence.

```
[4228] RetC
[4229] PopC
[4230] FPushC sMethodD
[4231] FCall
[4232] String
[4233] String
[4234] AGetC
[4235] CGetS
[4236] RetC
[4237] UnboxR
[4238] String
[4239] NSame
[4240] JmpZ
[4241] FPushC sMethodD
[4242] FCall
[4243] String
[4244] Salign
```

opcode operands
e.g., ADD 1, 2

have the same control flow

page 2

Title	Another Paper
Submission (PDF, max 100MB)	 400kB <small>© 9 Oct 2017 12:56:56pm EDT</small>
Replace:	<input type="button" value="Choose File"/> No file chosen

A black arrow pointing to the right, indicating the direction of the next step in the sequence.

```
[4228] RetC      O { /home/cheng/orochi.
[4229] PopC      O { /home/cheng/orochi.
[4230] FPushClsMethodD  O { /home/c
[4231] FCall      O { /home/cheng/orochi.
[4232] String     (Navigation::page) { .
[4233] String     (Navigation::page) { .
[4234] AGetC     (Navigation::page) { /.
[4235] CGetS     (Navigation::page) { /.
[4236] RetC      (Navigation::page) { /h
[4237] UnboxR    O { /home/cheng/orochi.
[4238] String     O { /home/cheng/orochi.
[4239] NSame     O { /home/cheng/orochi.
[4240] JmpZ      O { /home/cheng/orochi.
[4241] FPushClsMethodD  O { /home/c
[4242] FCall      O { /home/cheng/orochi.
[4243] String     (Navigation::page) { .
[4244] String     (Navigation::page) { .
```

An observation: repeated computation

web app:
HotCRP

page 1

Title

Submission (PDF, max 100MB)
 400kB © 9 Oct 2017 1:59:08pm EDT

Replace: No file chosen

Authors
 List the authors one per line, including email
 not be able to see author information. Any au
 Name
 1.
 2.

page 2

Title	Another Paper
Submission (PDF, max 100MB)	
 400kB	© 9 Oct 2017 12:56:56pm EDT
Replace:	<input type="button" value="Choose File"/> No file chosen

opcode operands

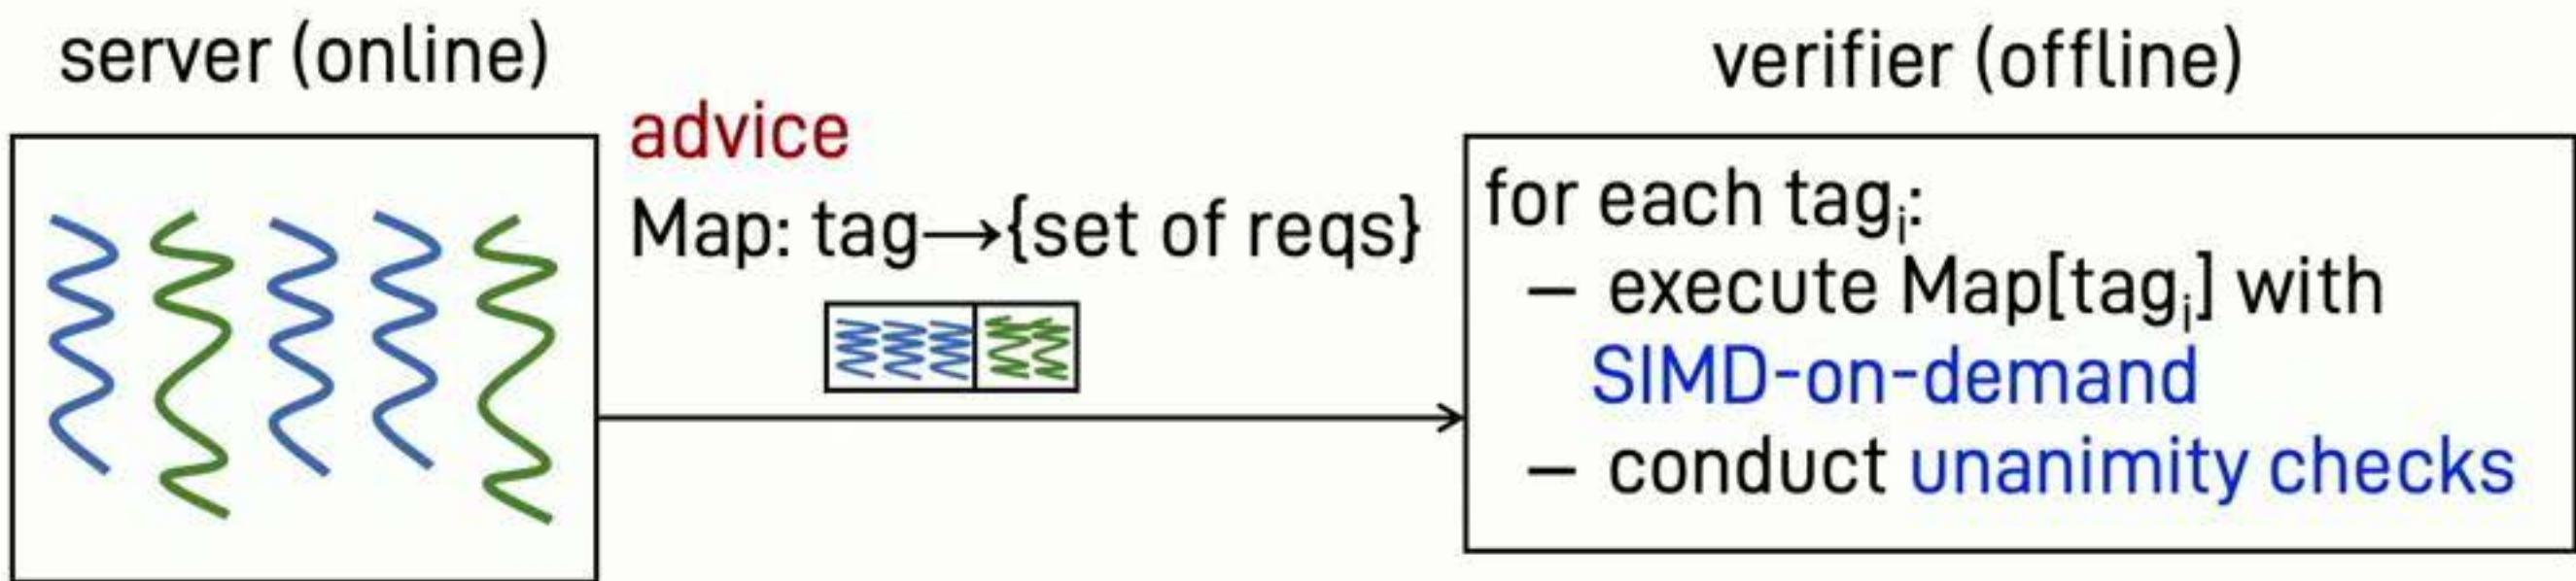
e.g., ADD 1, 2

requires trusting the recorder (server)

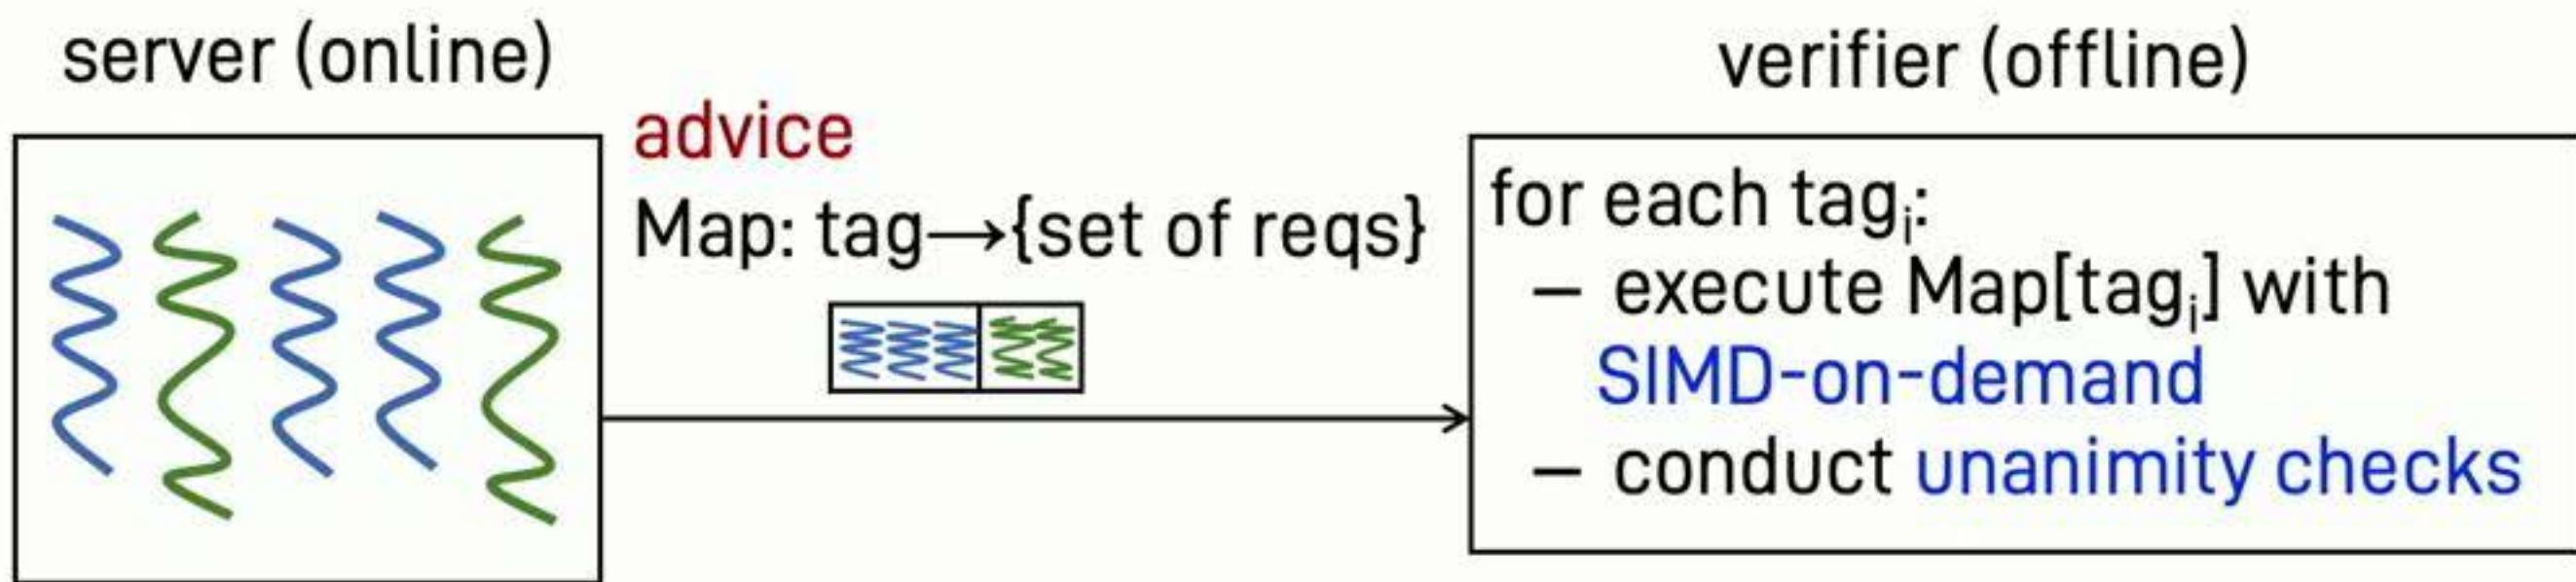
trol flow

```
[4228] RetC      O { /home/cheng/orochi/ }
[4229] PopC      O { /home/cheng/orochi/ }
[4230] FPushC sMethodD      O { /home/cheng/orochi/ }
[4231] FCall      O { /home/cheng/orochi/ }
[4232] String      (Navigation::page) { . }
[4233] String      (Navigation::page) { . }
[4234] AGetC      (Navigation::page) { / }
[4235] CGetS      (Navigation::page) { / }
[4236] RetC      (Navigation::page) { / }
[4237] UnboxR      O { /home/cheng/orochi/ }
[4238] String      O { /home/cheng/orochi/ }
[4239] NSame      O { /home/cheng/orochi/ }
[4240] JmpZ      O { /home/cheng/orochi/ }
[4241] FPushC sMethodD      O { /home/cheng/orochi/ }
[4242] FCall      O { /home/cheng/orochi/ }
[4243] String      (Navigation::page) { . }
[4244] String      (Navigation::page) { . }
```

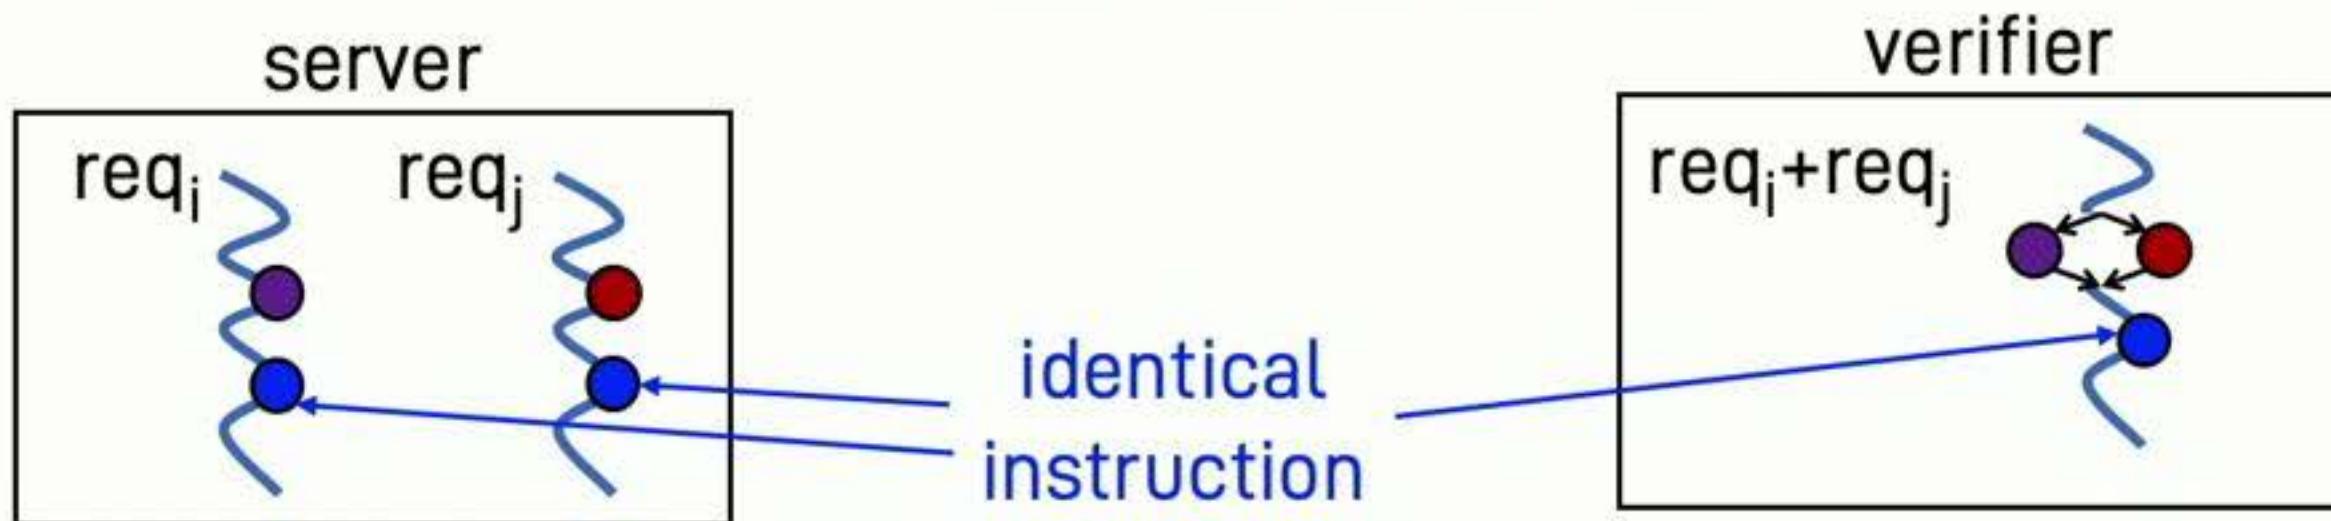
Accelerate re-execution without trusting the server



Accelerate re-execution without trusting the server



- SIMD-on-demand re-executes **identical instructions** once.



SIMD-on-demand eliminates redundant computation

```
main(a, b):  
  c ← a * b  
  c ← c + 1
```

req₁: a=1; b=2

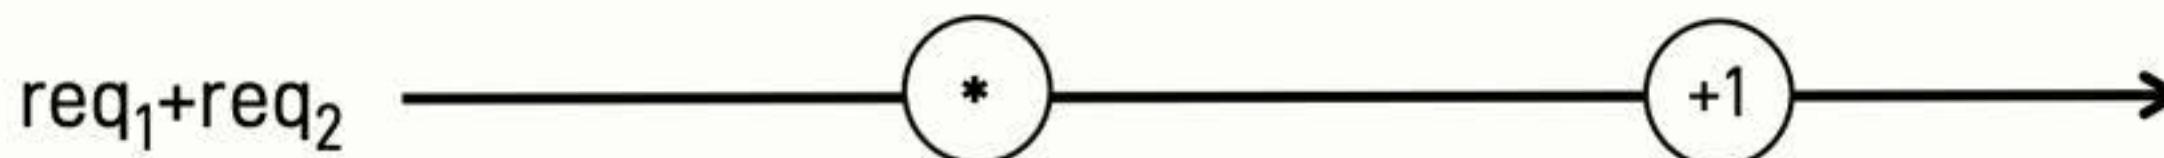
req₂: a=2; b=1

SIMD-on-demand eliminates redundant computation

```
main(a, b):  
  c ← a * b  
  c ← c + 1
```

req₁: a=1; b=2

req₂: a=2; b=1

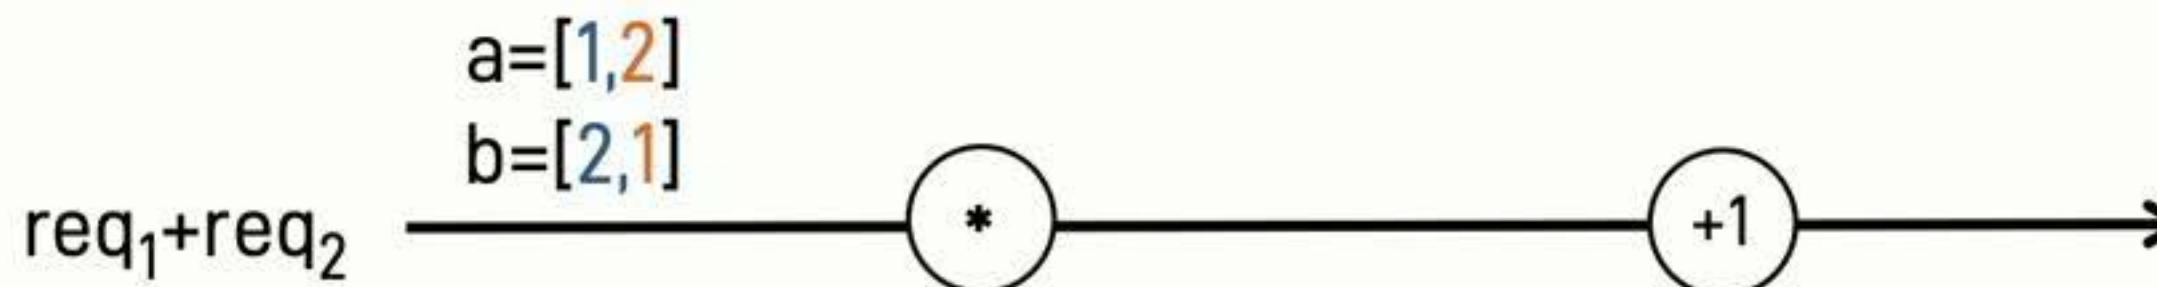


SIMD-on-demand eliminates redundant computation

```
main(a, b):  
  c ← a * b  
  c ← c + 1
```

req₁: a=1; b=2

req₂: a=2; b=1

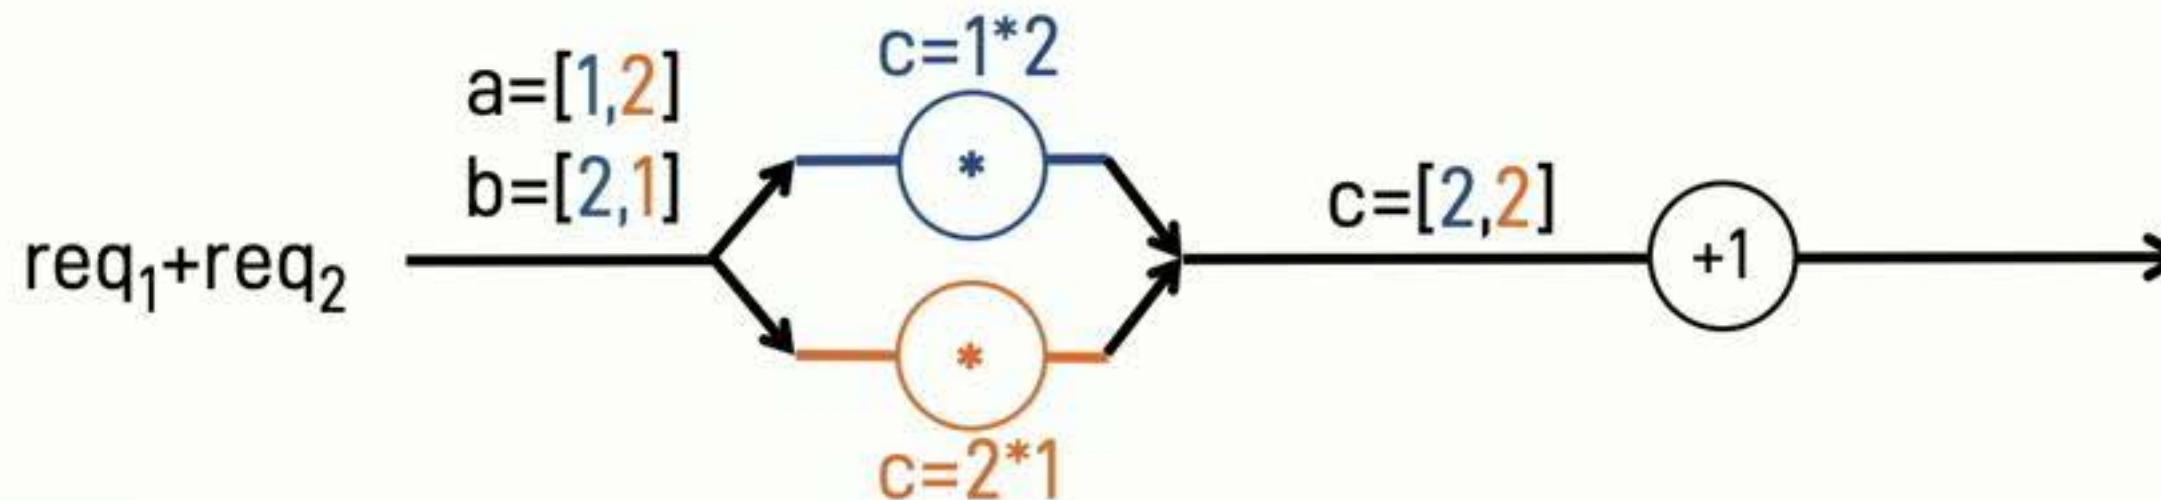


- **Multi-value** represents different values of the same variable.

SIMD-on-demand eliminates redundant computation

```
main(a, b):  
  c ← a * b  
  c ← c + 1
```

req₁: a=1; b=2
req₂: a=2; b=1

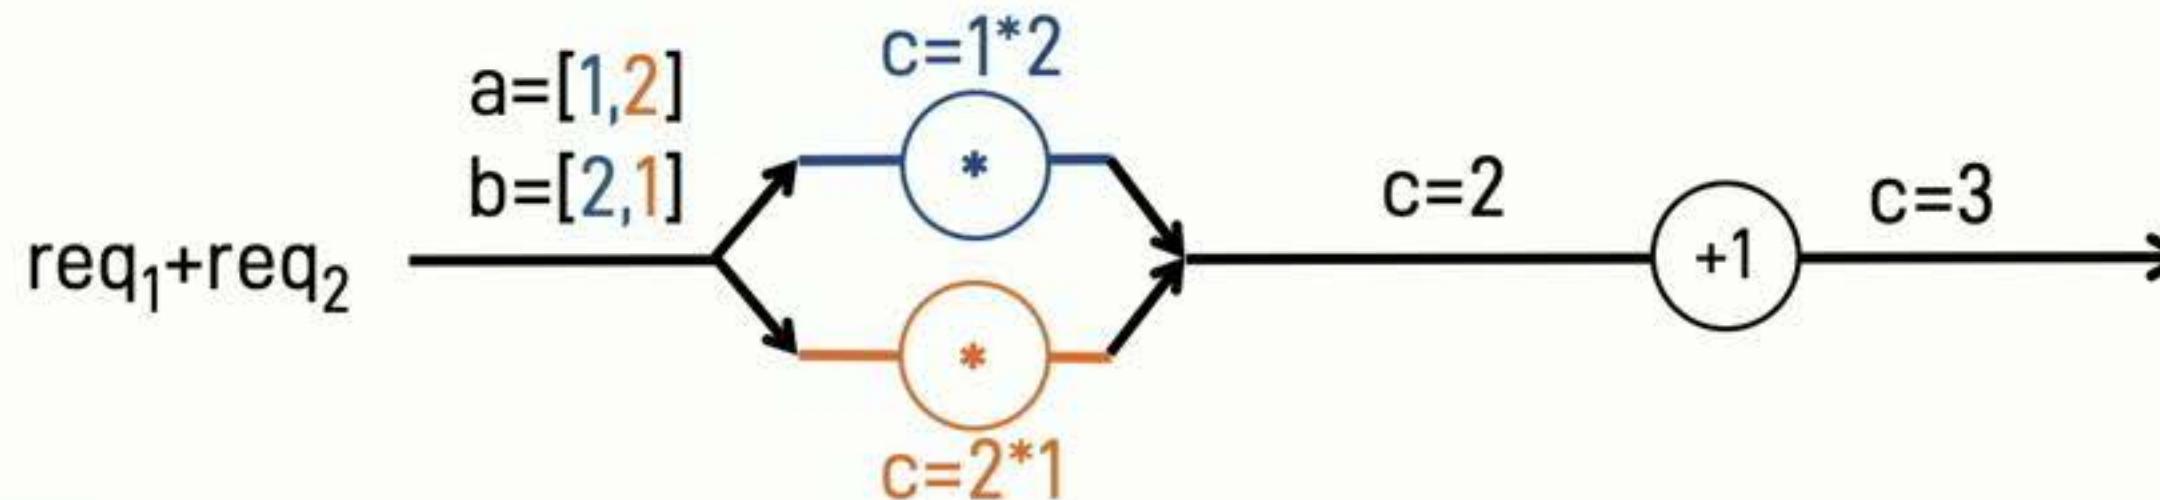


- **Multi-value** represents different values of the same variable.

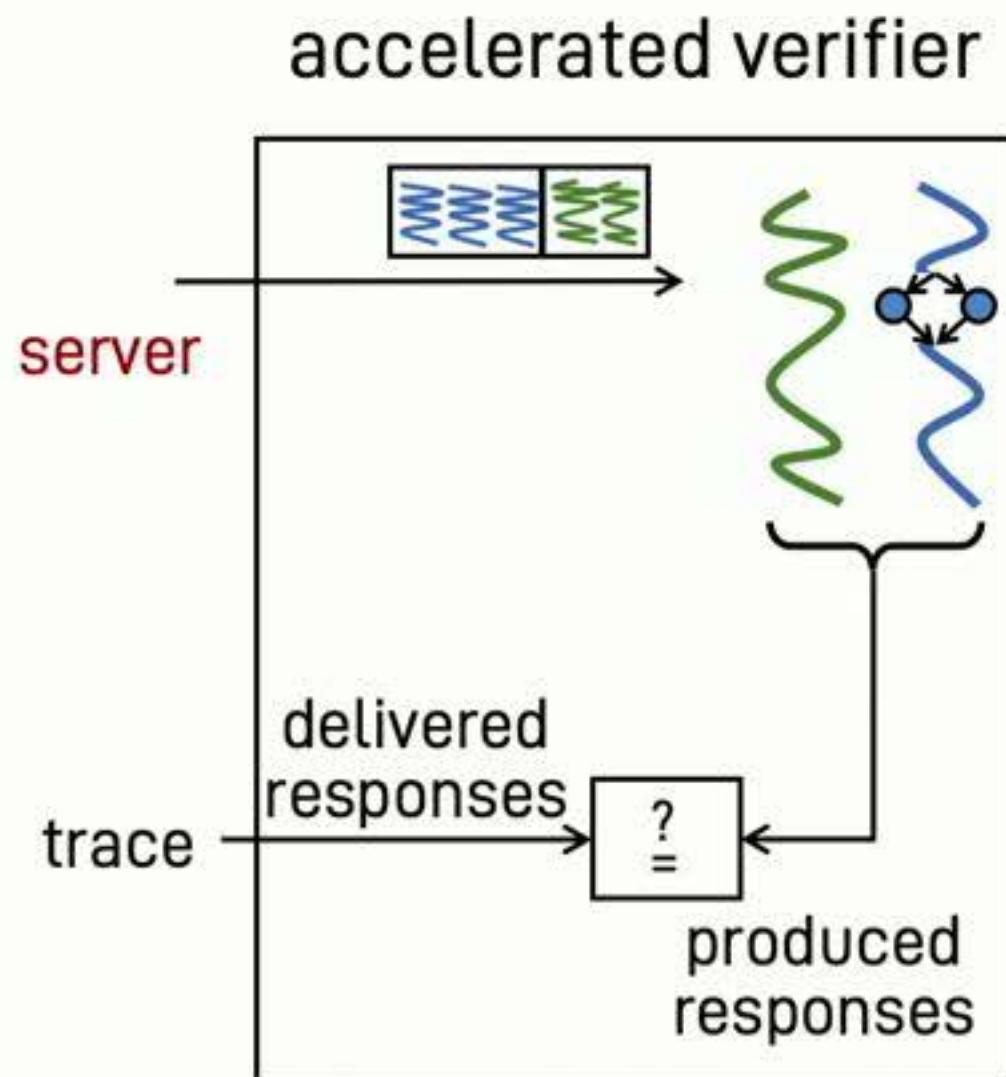
SIMD-on-demand eliminates redundant computation

```
main(a, b):  
  c ← a * b  
  c ← c + 1
```

req₁: a=1; b=2
req₂: a=2; b=1

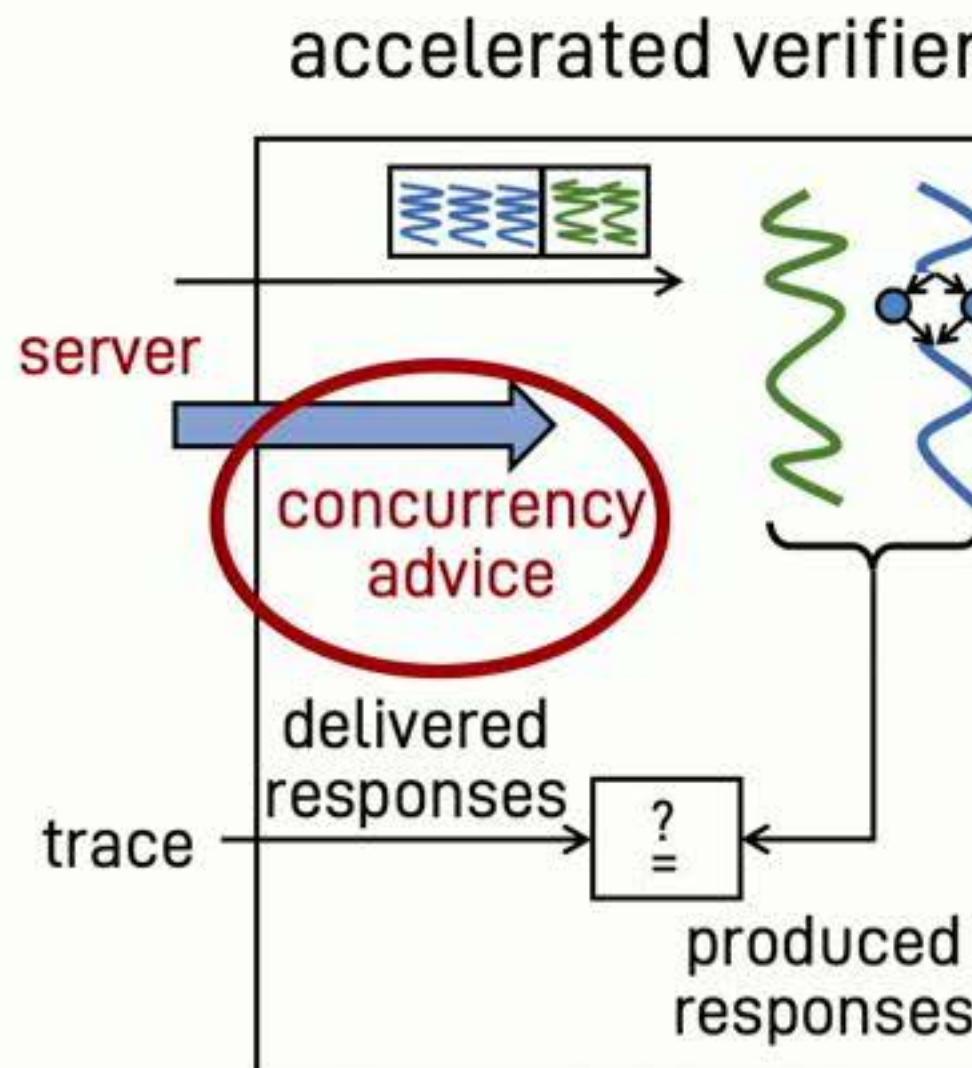


- **Multi-value** represents different values of the same variable.
- Verifier **collapses** multi-value to scalar if possible.



Problem: naive re-execution doesn't save work

Solution: deduplicated re-execution



Problem: naive re-execution doesn't save work

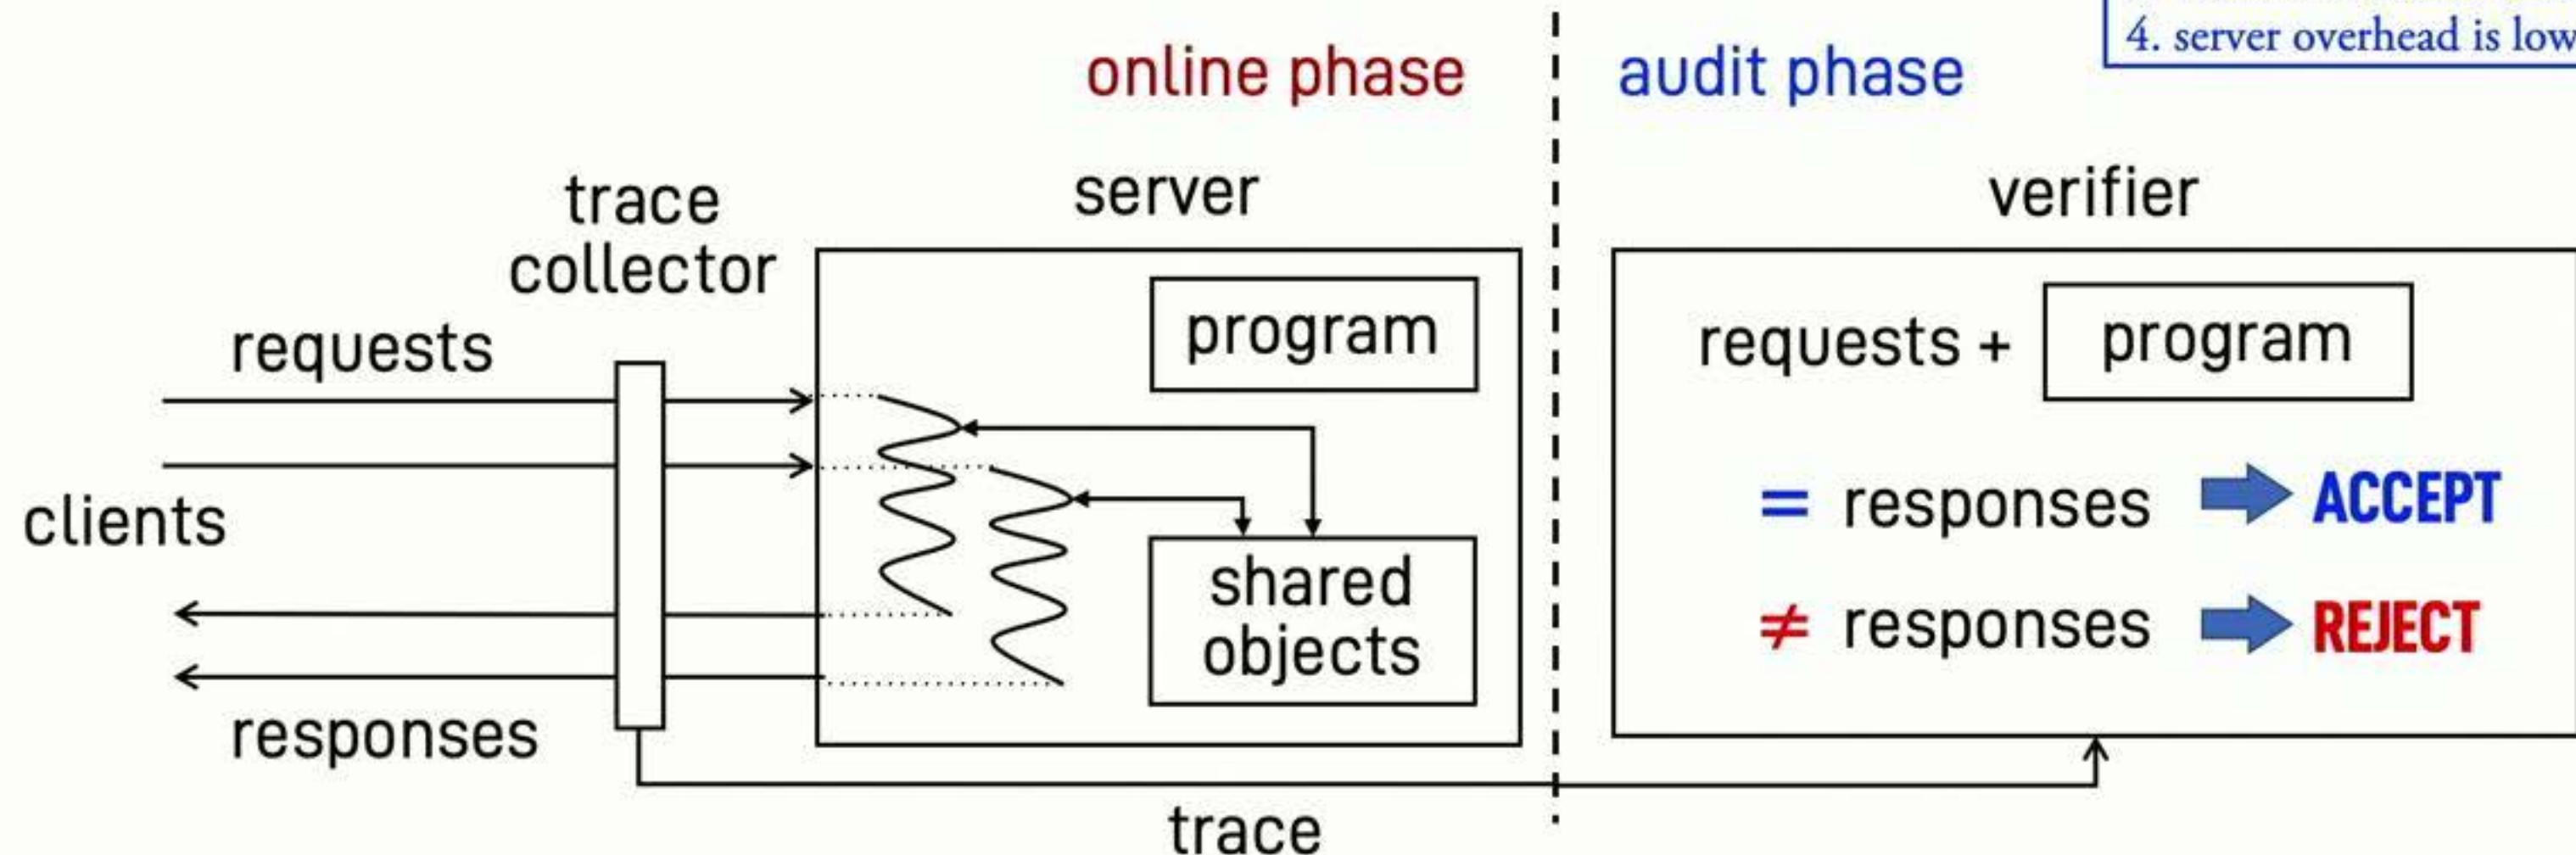
Solution: deduplicated re-execution

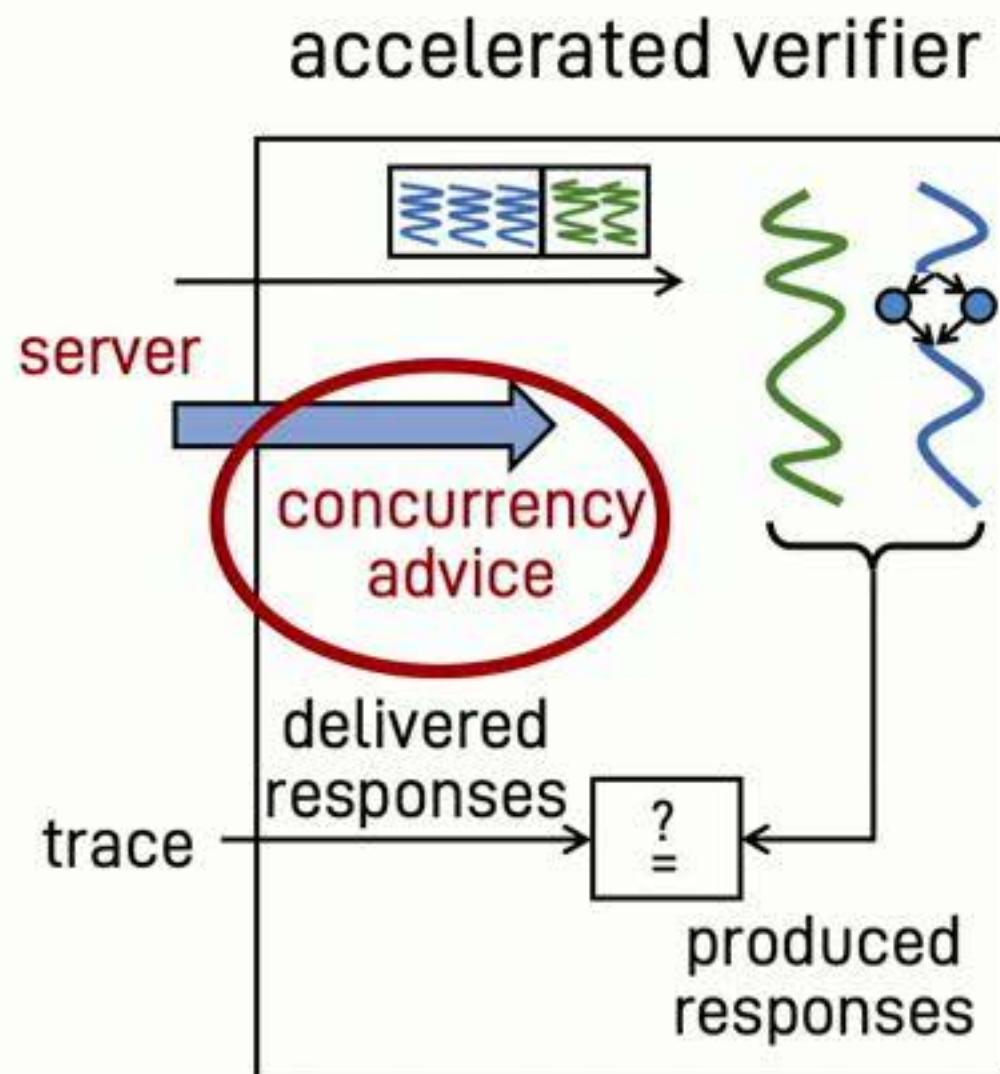
Problem: many schedules to explore

Proposal: (somehow) use but don't trust advice

The Efficient Server Audit Problem

1. server is untrusted...
2. server is concurrent
3. verifier is weaker than server
4. server overhead is low...





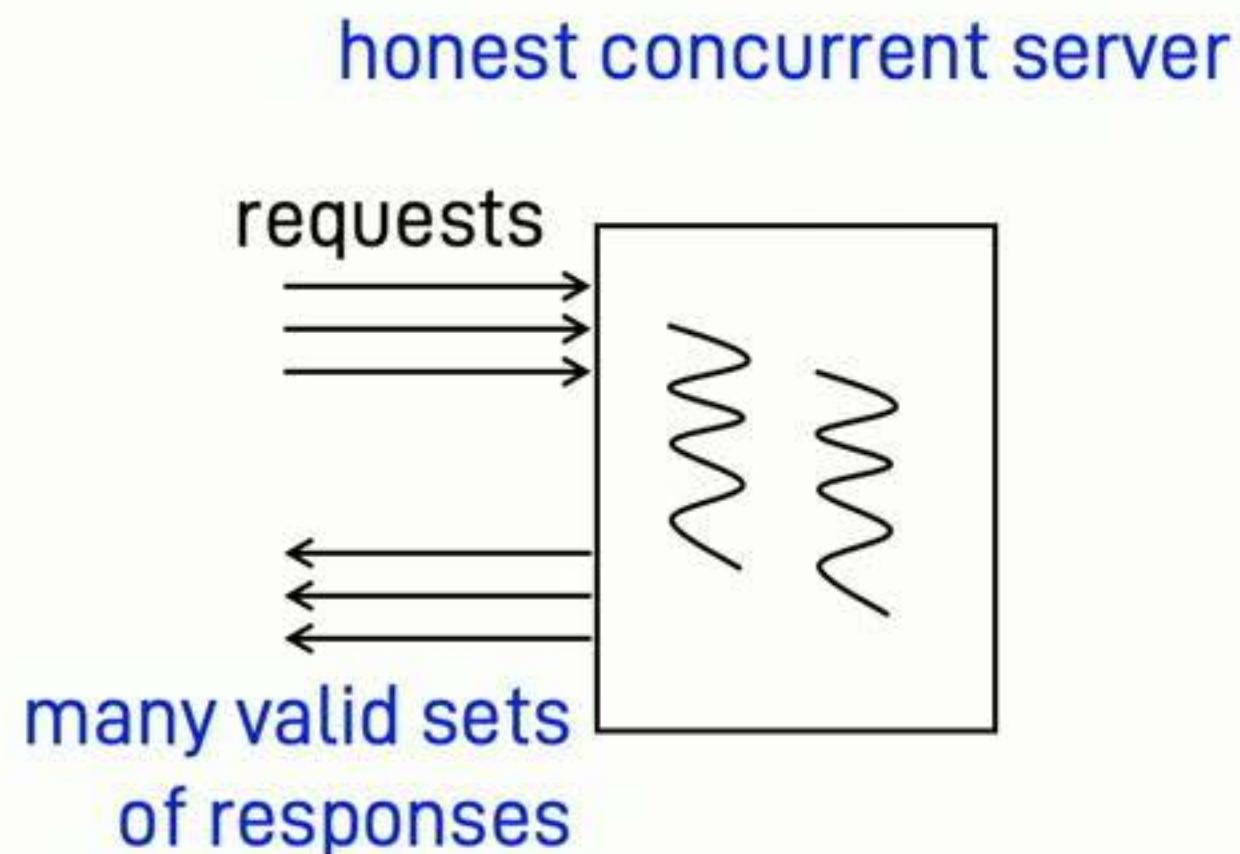
Problem: naive re-execution doesn't save work

Solution: deduplicated re-execution

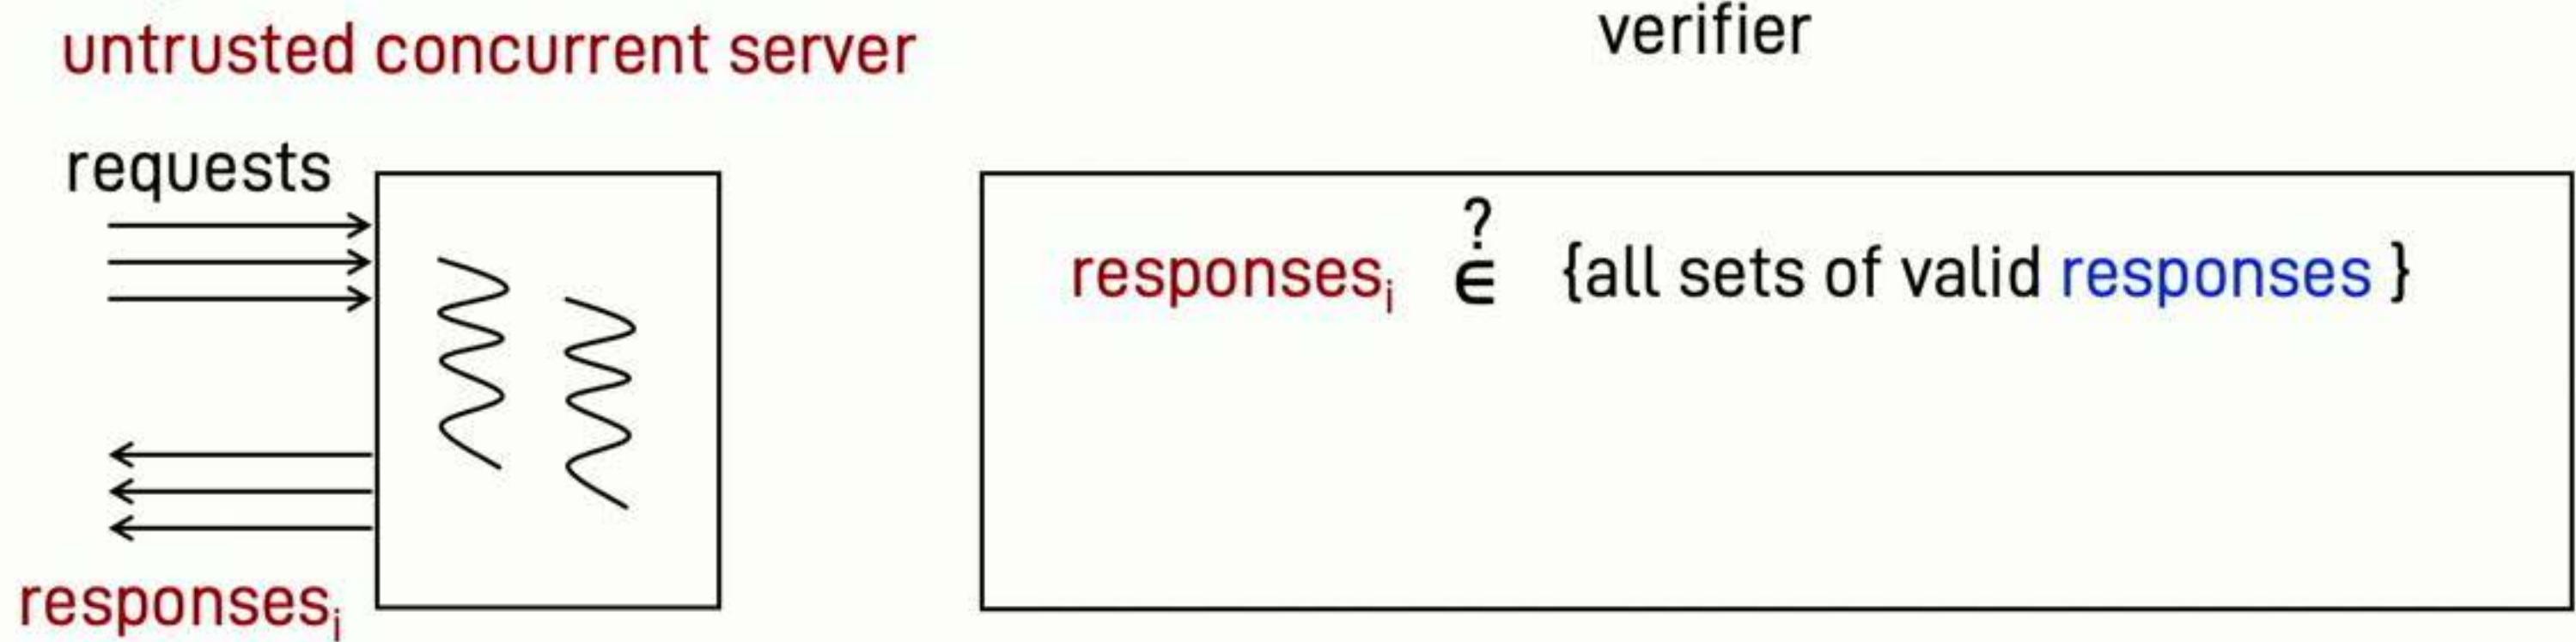
Problem: many schedules to explore

Proposal: (somehow) use but don't trust advice

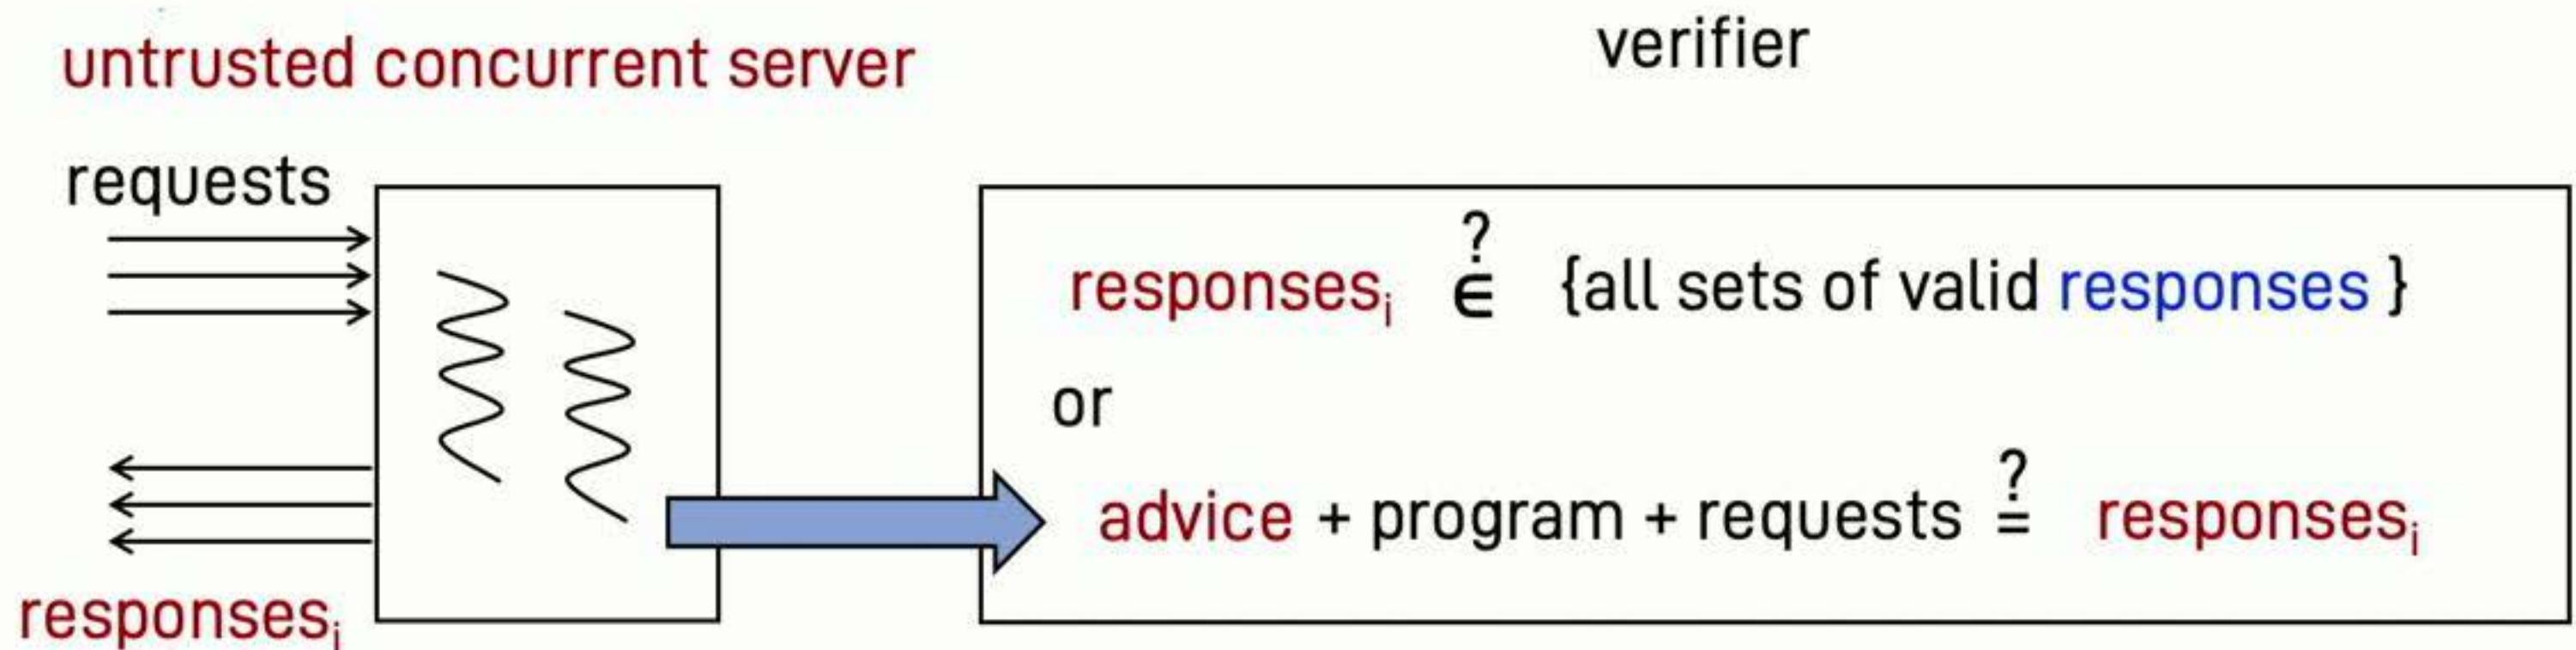
Concurrency and untrusted server are in tension



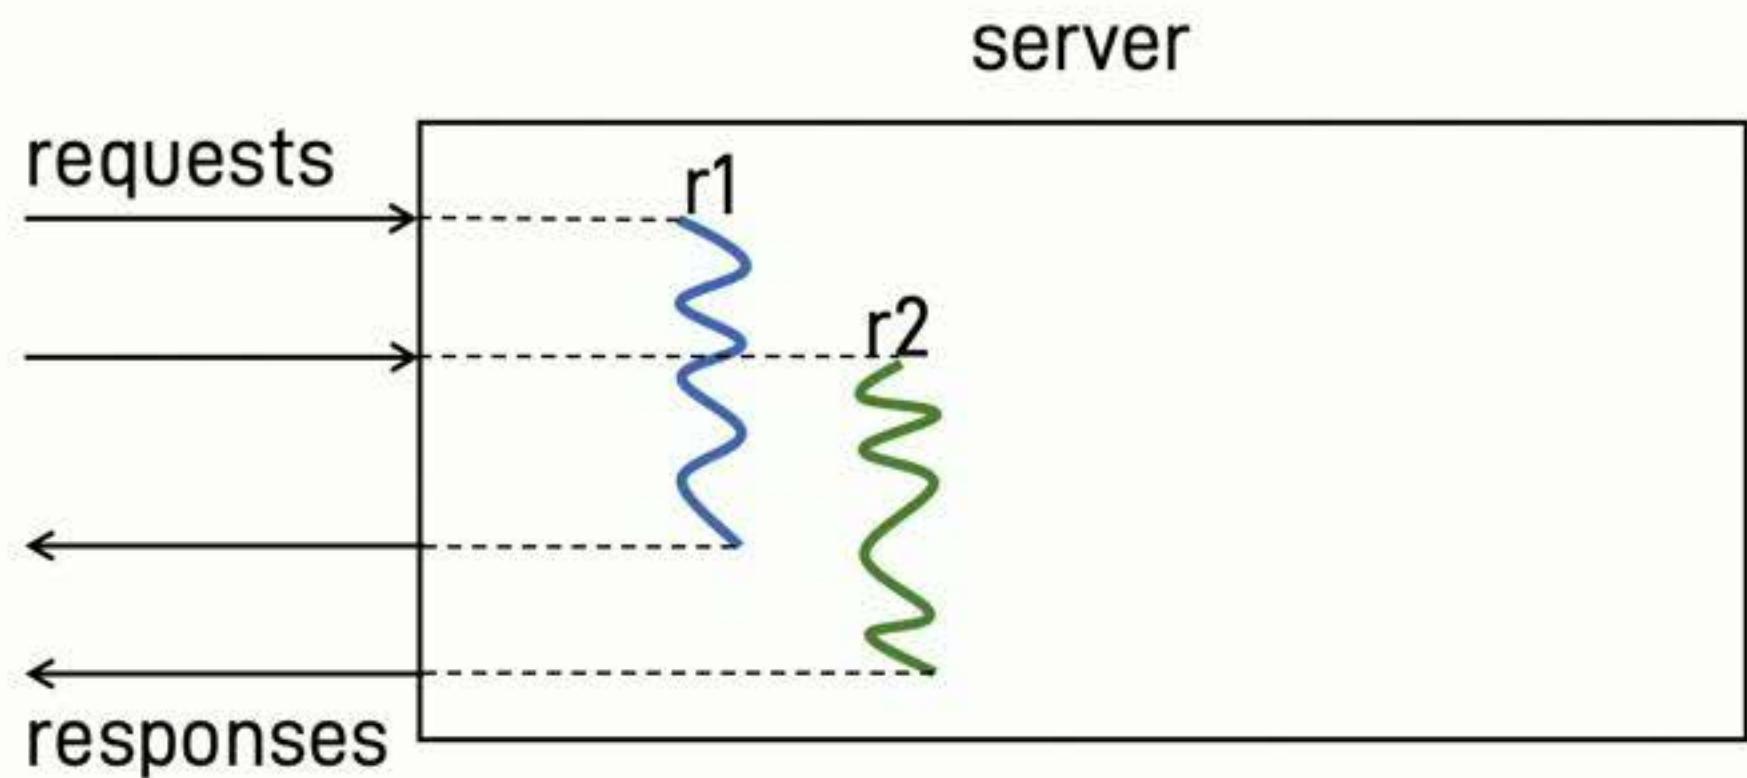
Concurrency and untrusted server are in tension



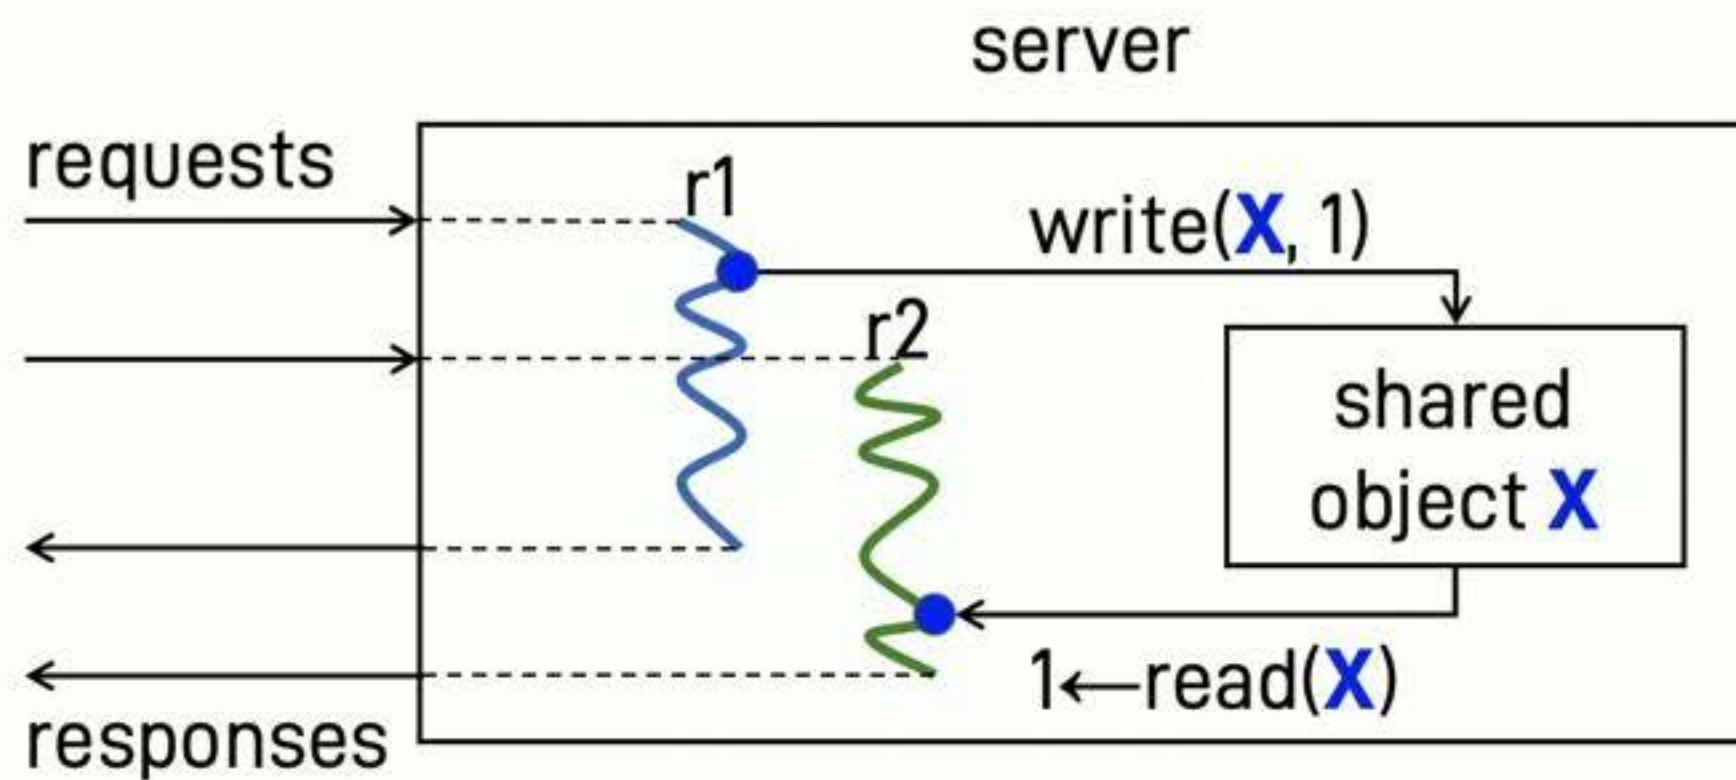
Concurrency and untrusted server are in tension



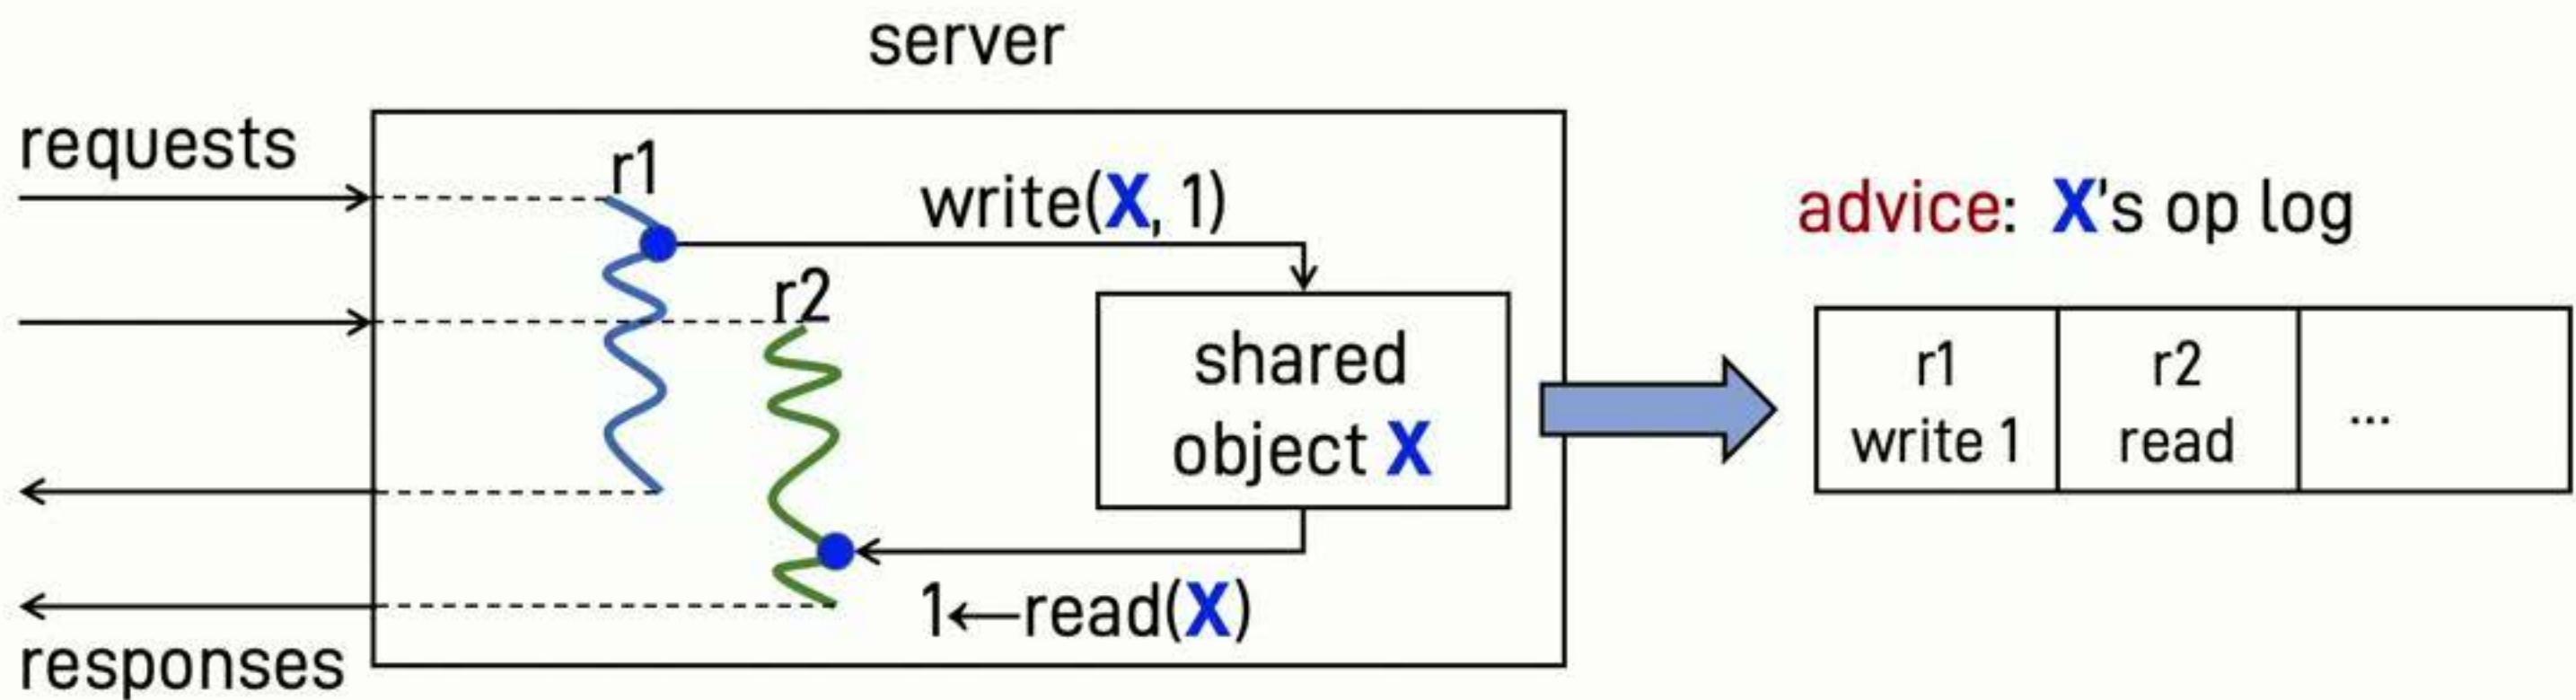
Concurrency model



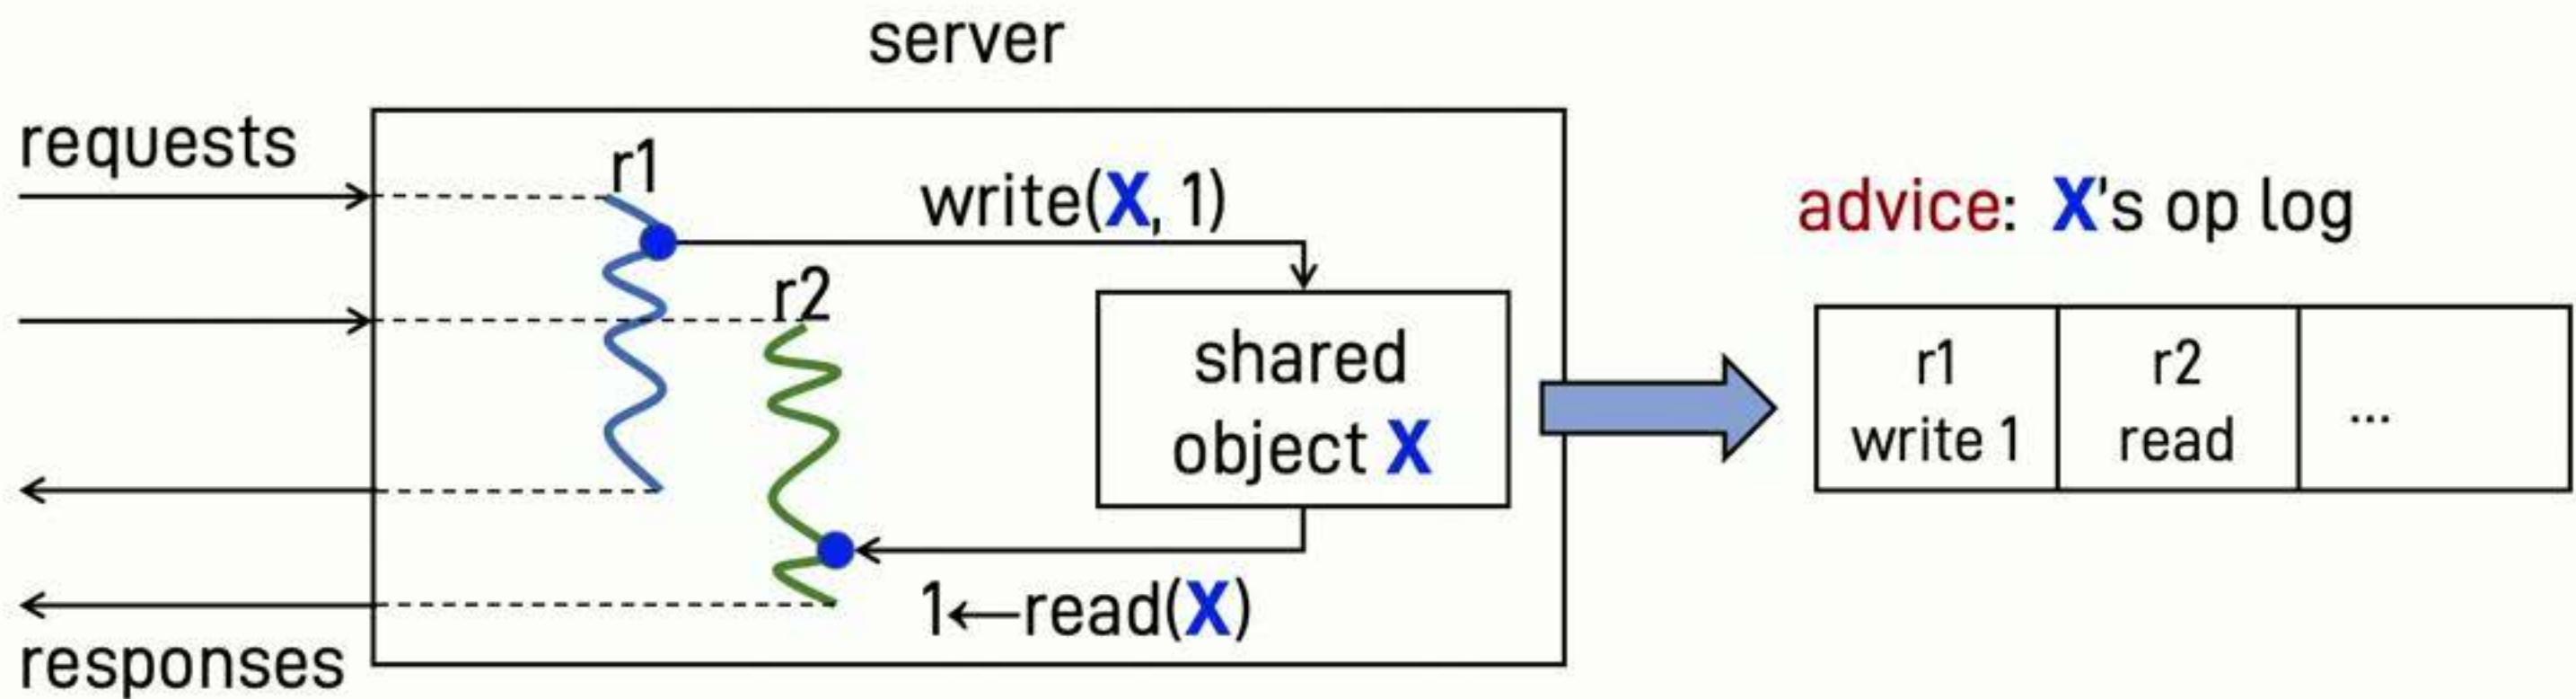
Concurrency model



Concurrency model



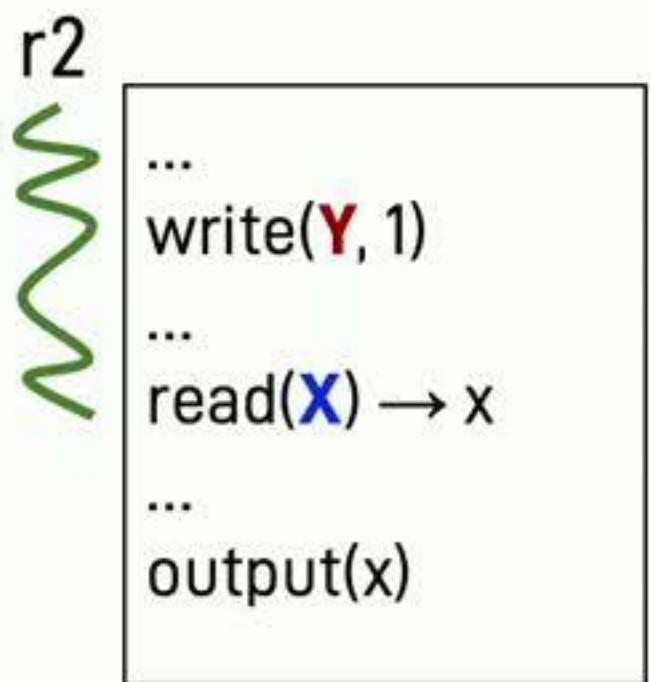
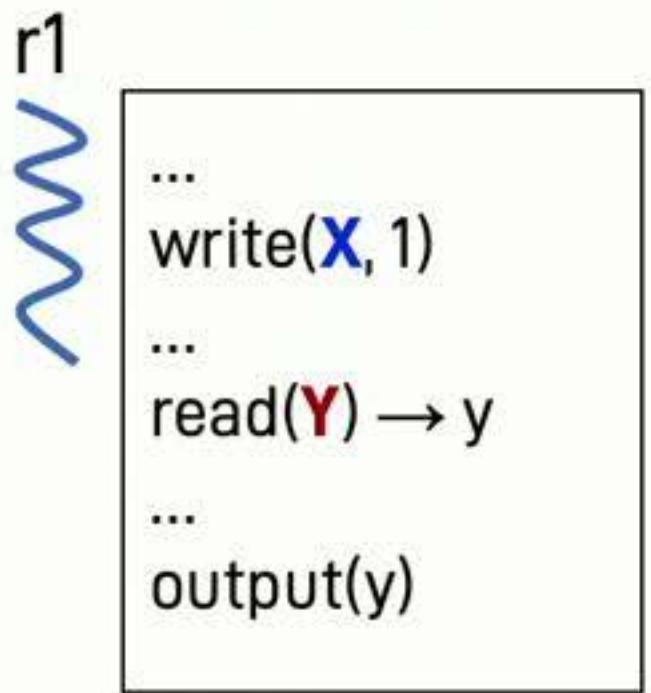
Concurrency model



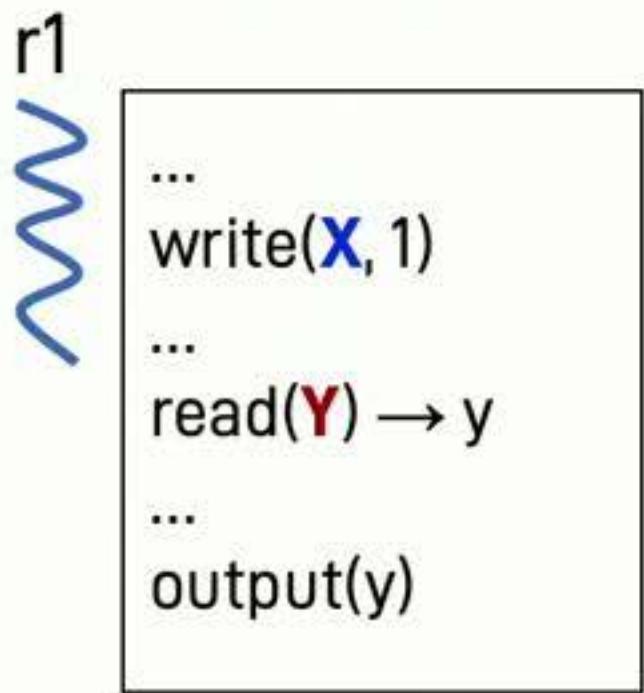
- What about simple re-execution according to op logs?

X=0, Y=0

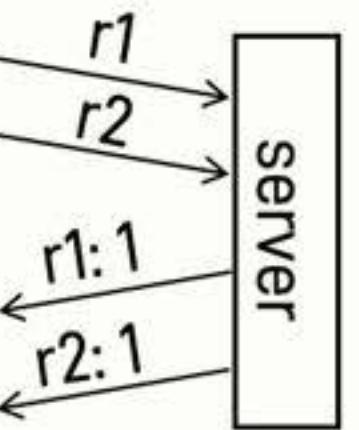
trace



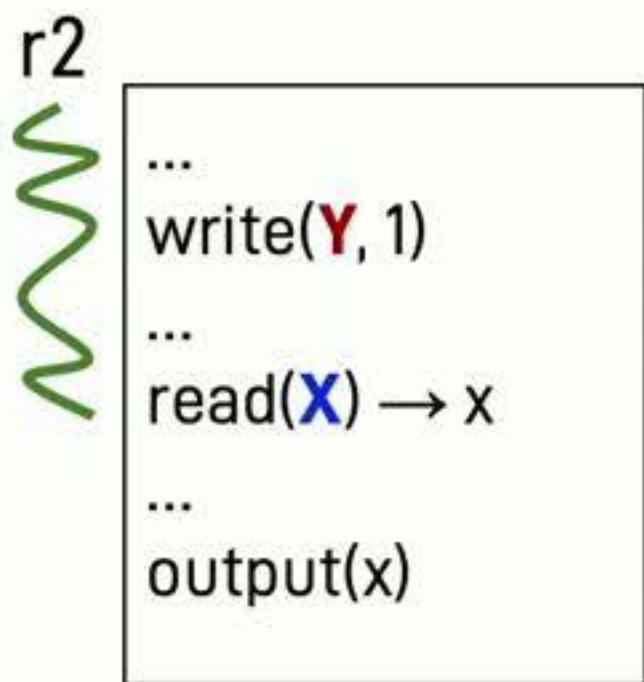
X=0, Y=0



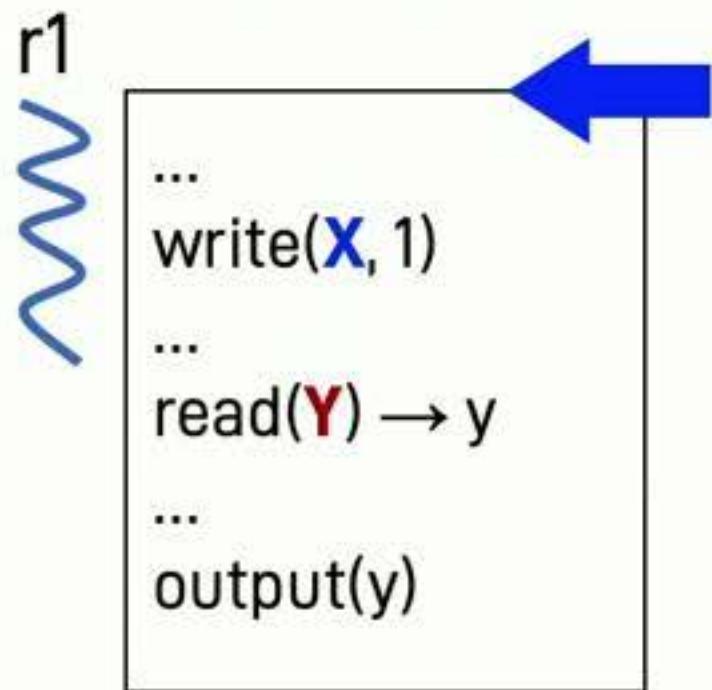
trace



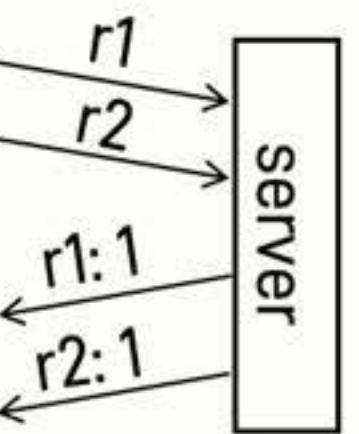
?



X=0, Y=0

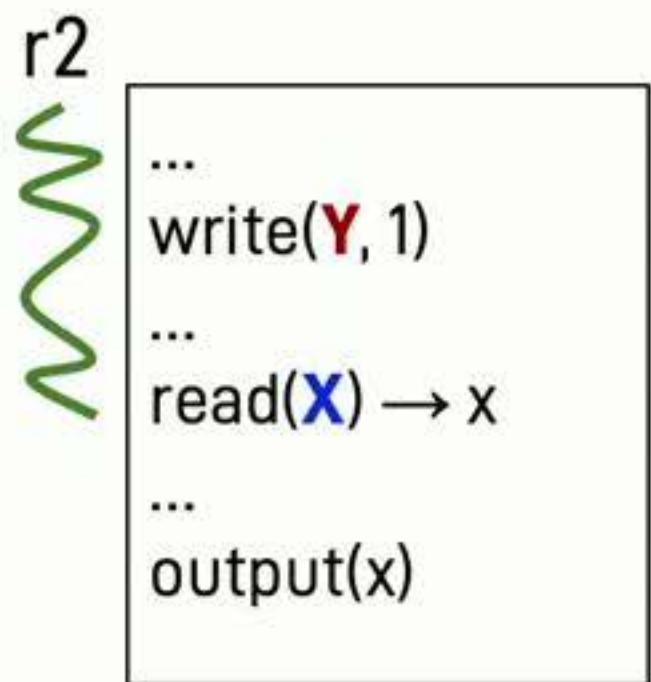


trace

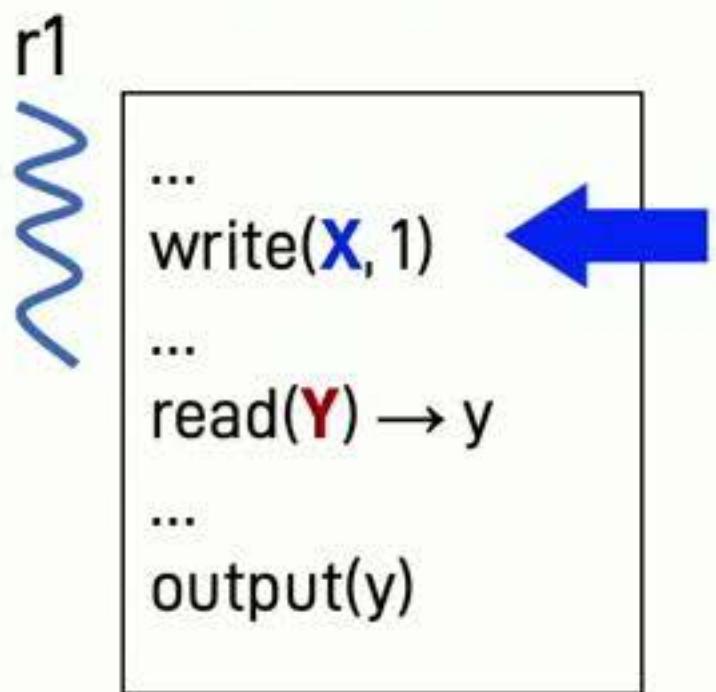


X = 0, Y = 0

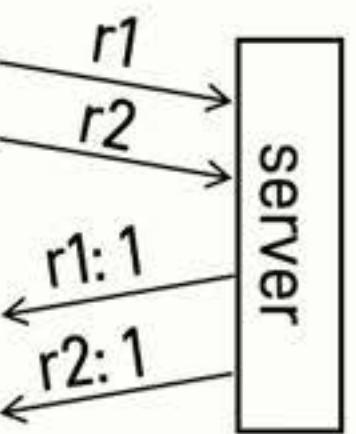
Accept



X=0, Y=0

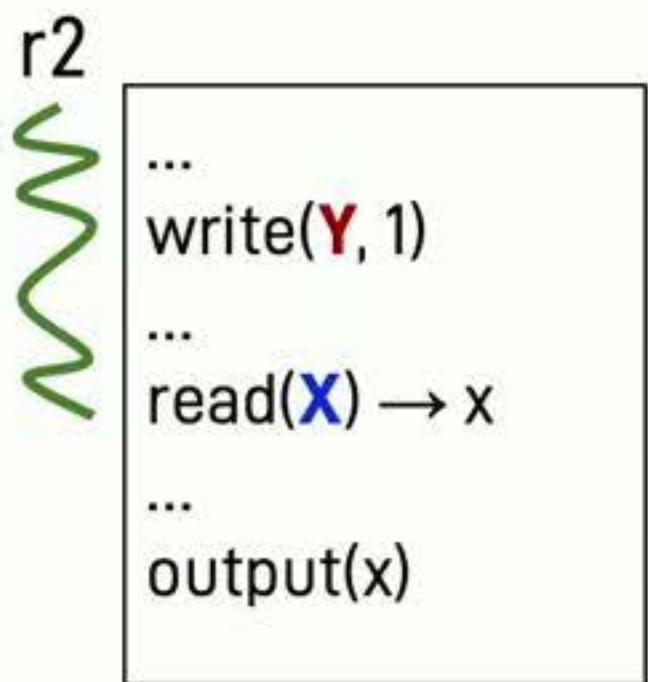


trace

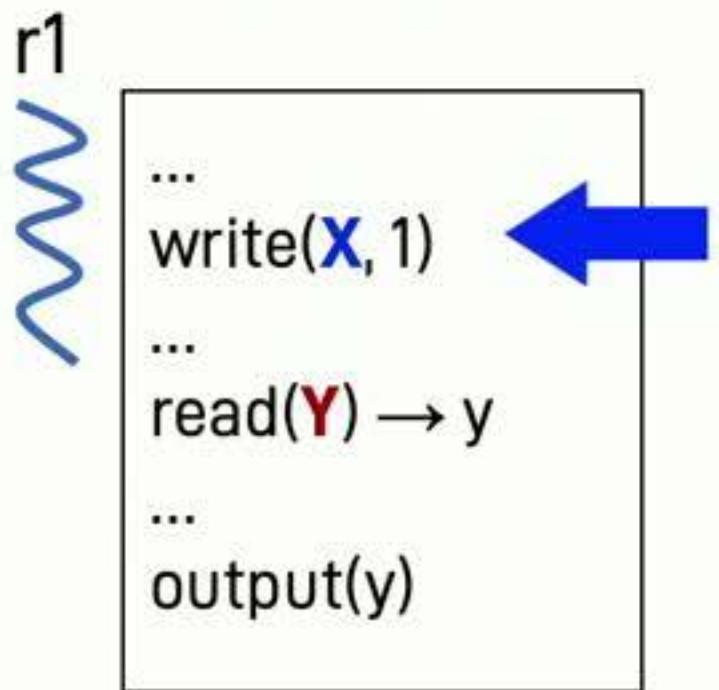


X = 1, Y = 0

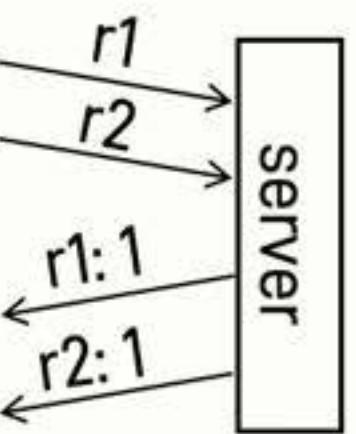
Accept



X=0, Y=0

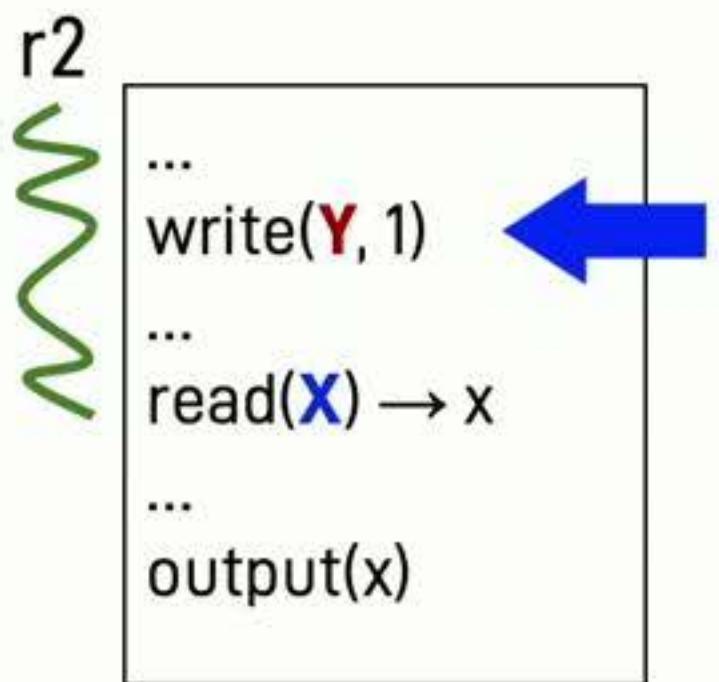


trace



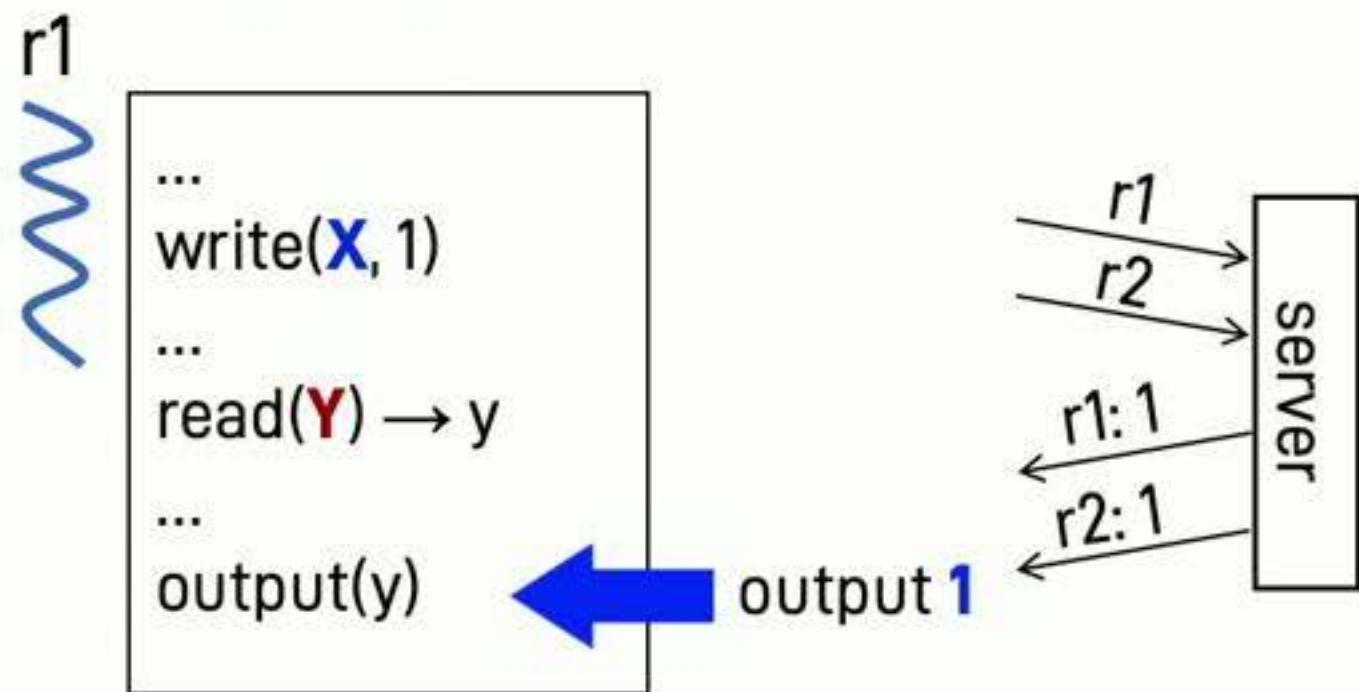
X = 1, Y = 1

Accept



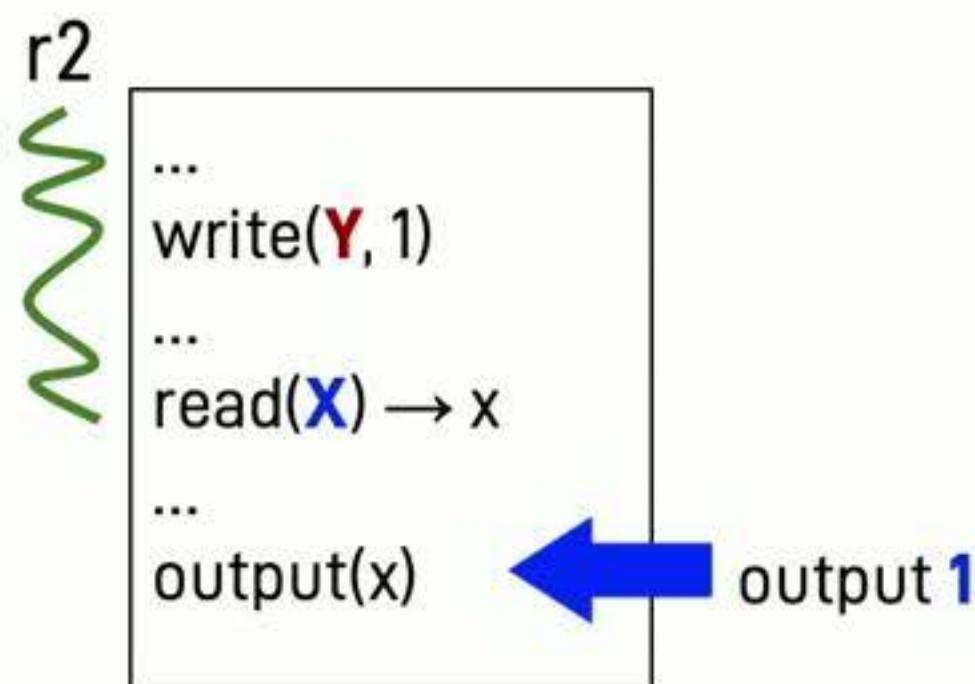
X=0, Y=0

trace

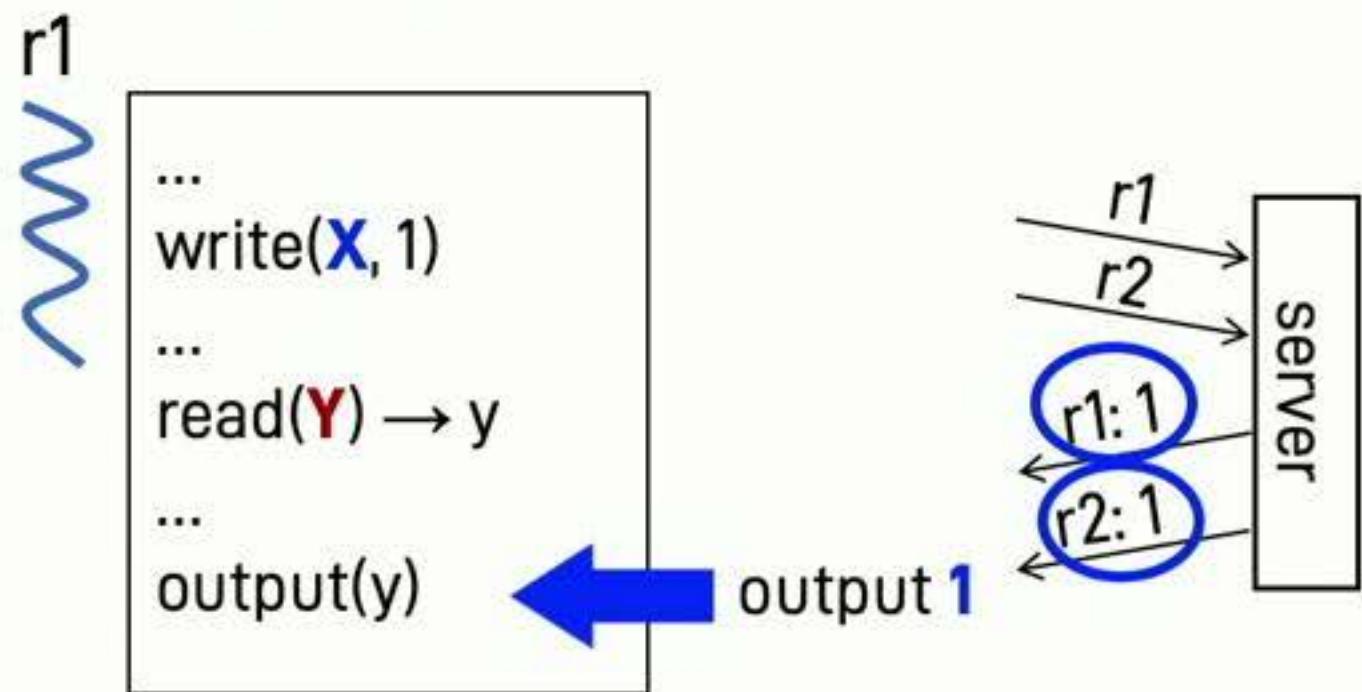


X = 1, Y = 1

Accept



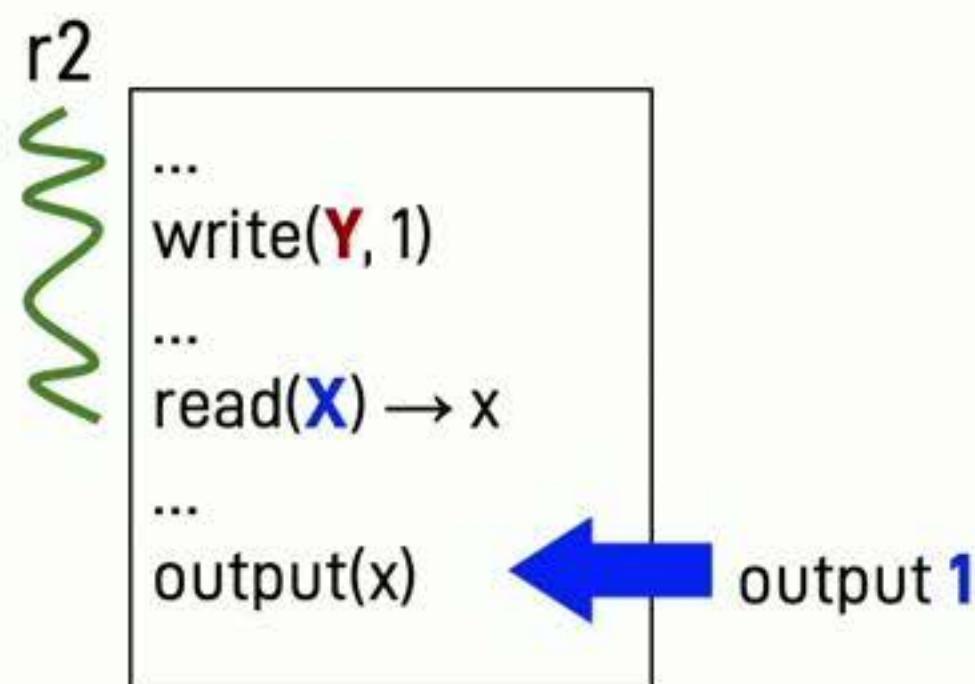
X=0, Y=0



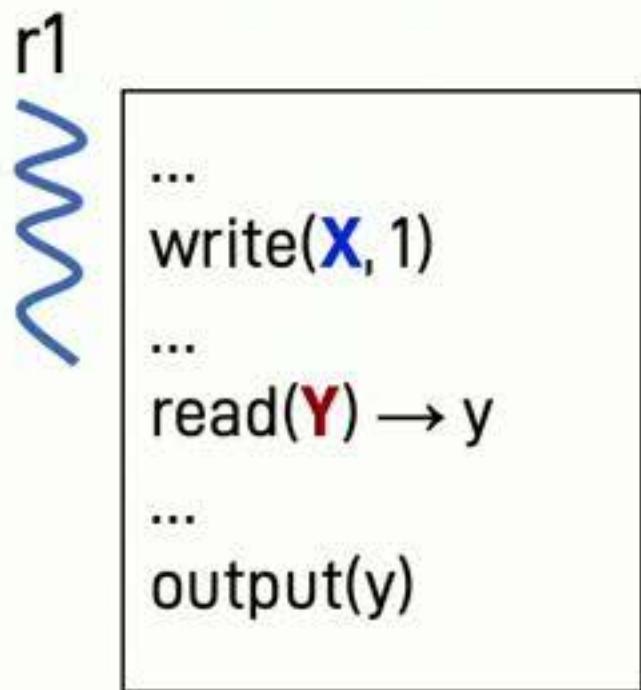
X = 1, Y = 1

Accept

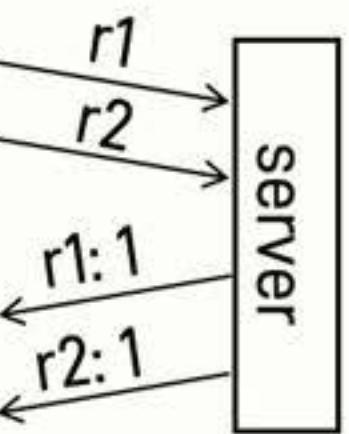
(a)



X=0, Y=0



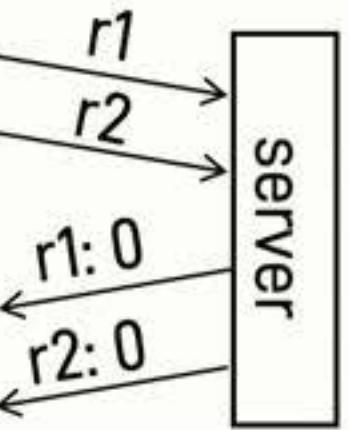
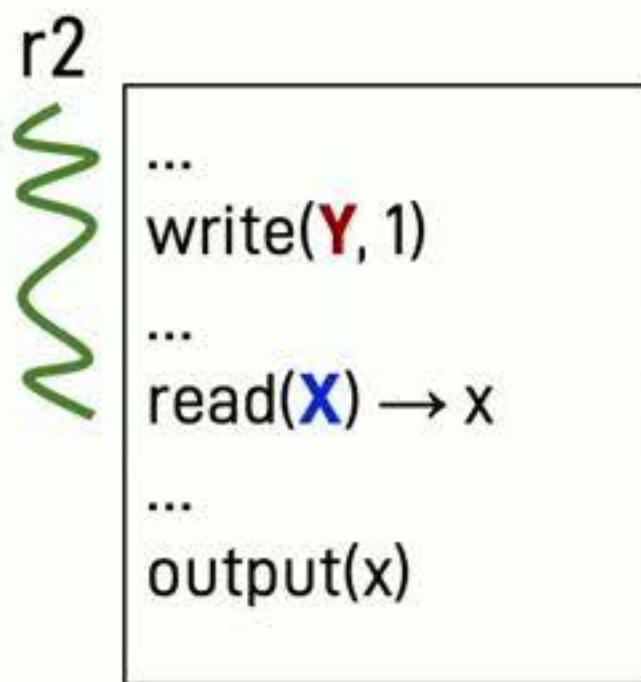
trace



X = 1, Y = 1

Accept

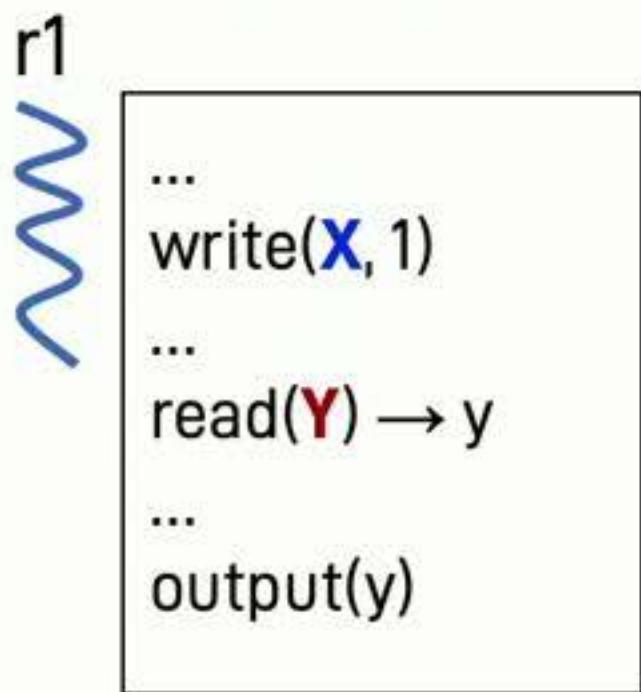
(a)



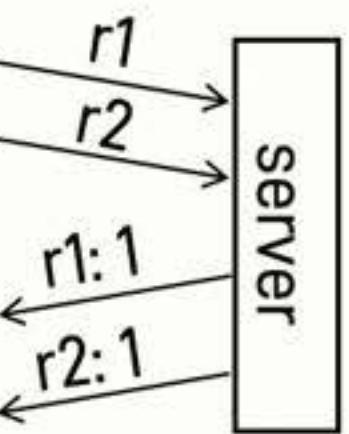
?

(b)

X=0, Y=0



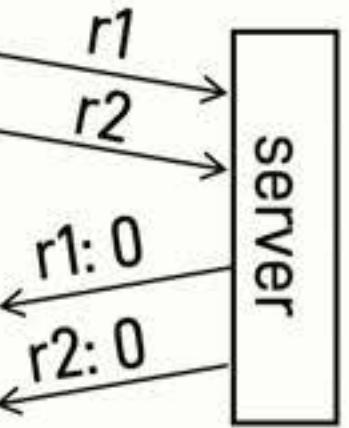
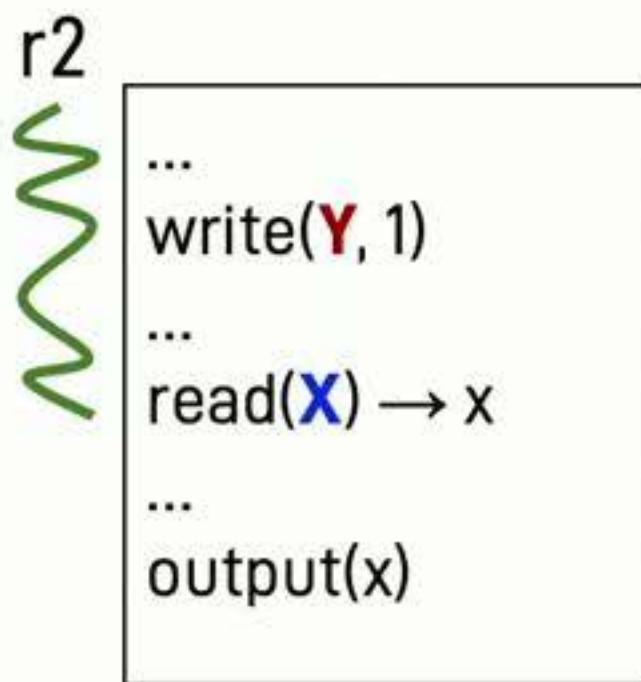
trace



X = 1, Y = 1

Accept

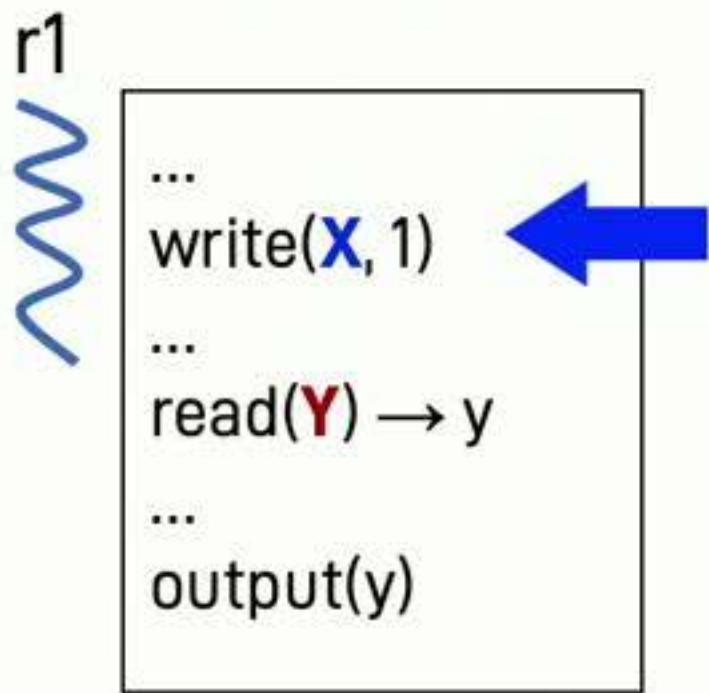
(a)



Reject

(b)

X=0, Y=0

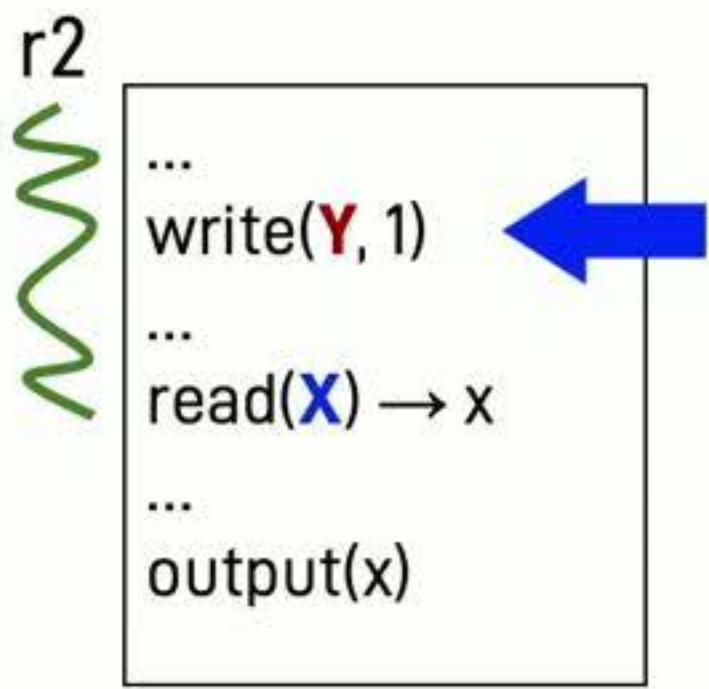


trace

X = 1, Y = 1

Accept

(a)

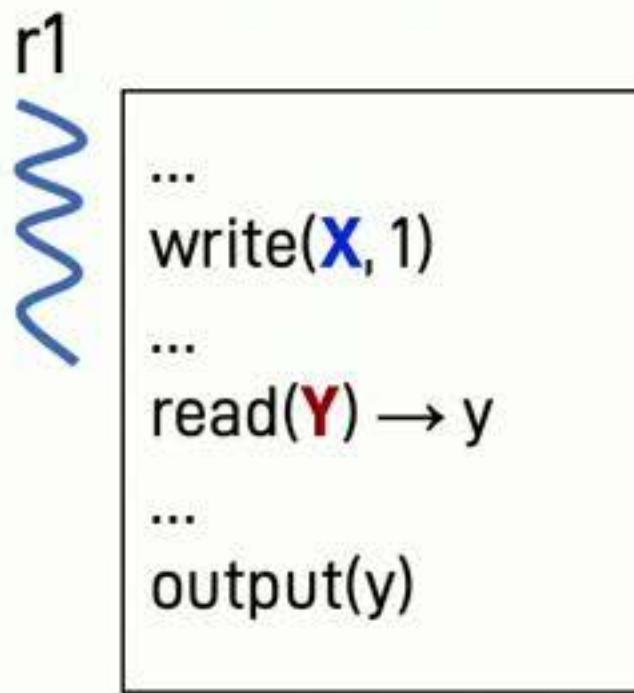


X = 1 or Y = 1

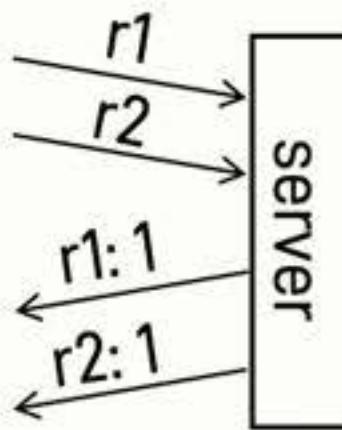
Reject

(b)

X=0, Y=0



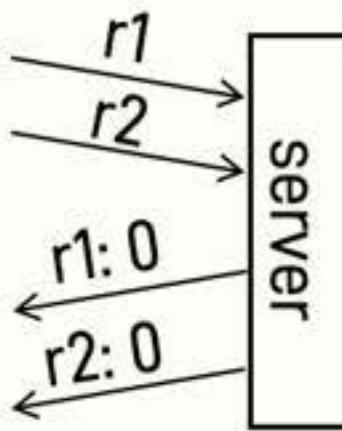
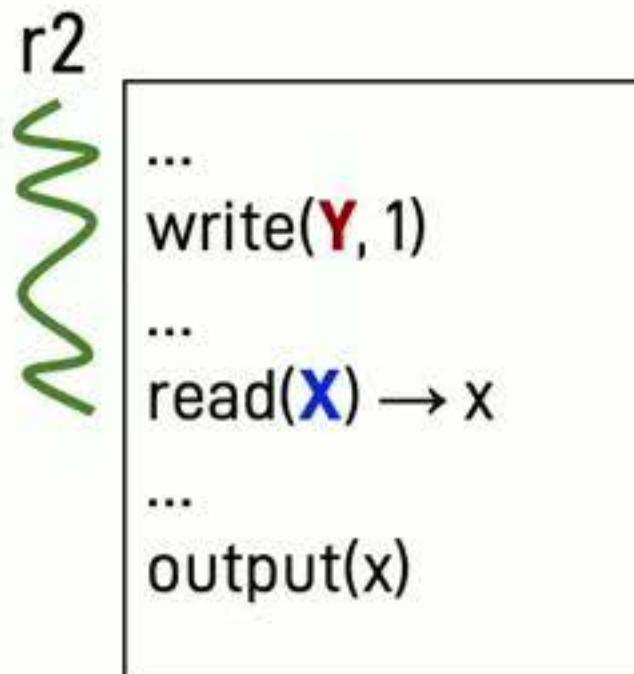
trace



op logs

X's log	r1	r2	...
	write 1	read	...
Y's log	r2	r1	...
	write 1	read	...

(a)



X's log

X's log	r2	r1	...
	read	write 1	...
Y's log	r1	r2	...
	read	write 1	...

(b)

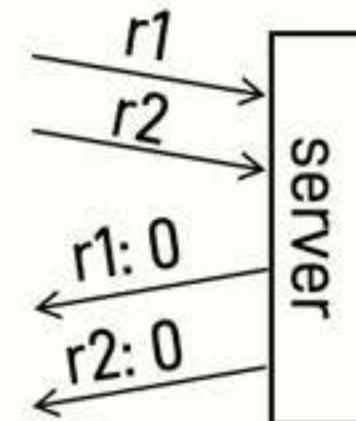
X=0, Y=0

r1

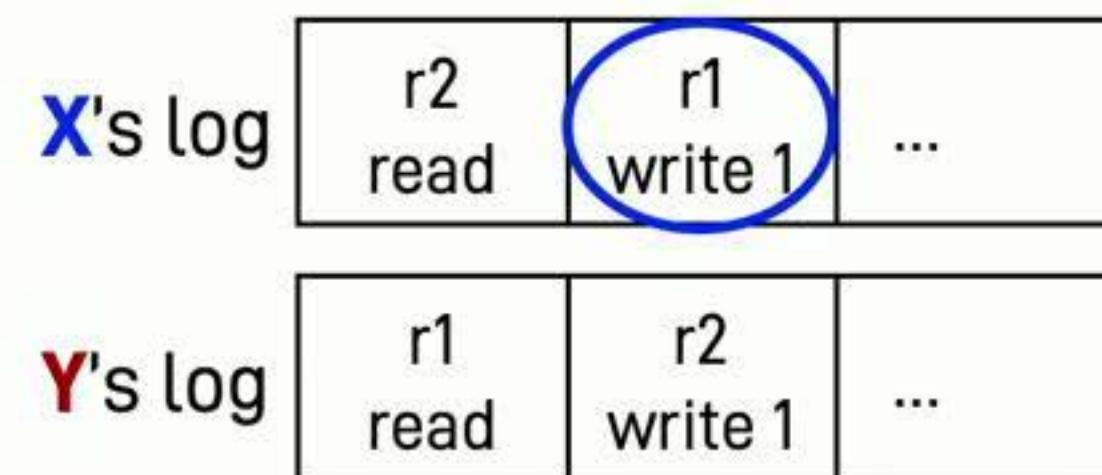
```
...  
write(X, 1) ←  
...  
read(Y) → y  
...  
output(y)
```

r2

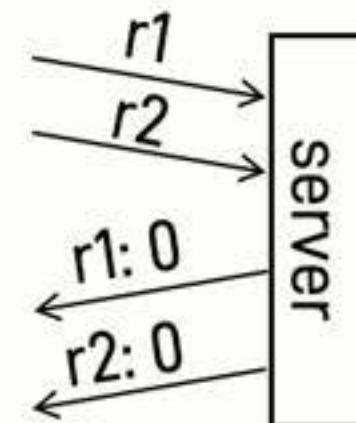
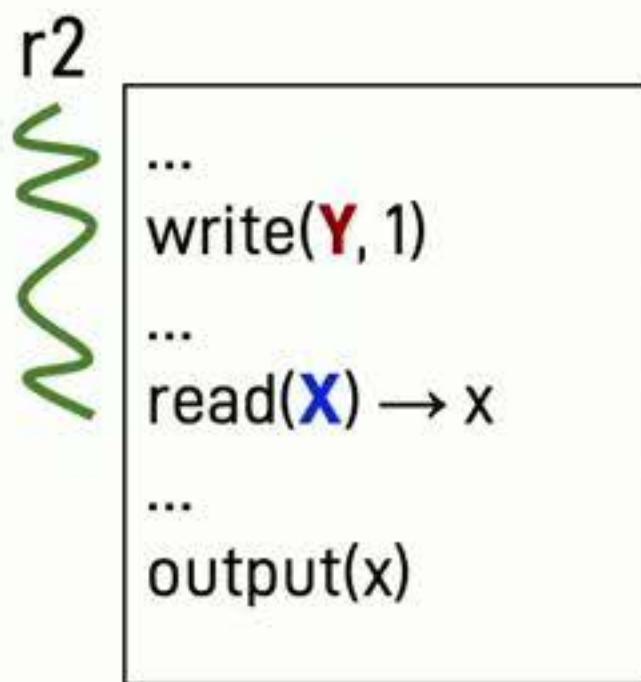
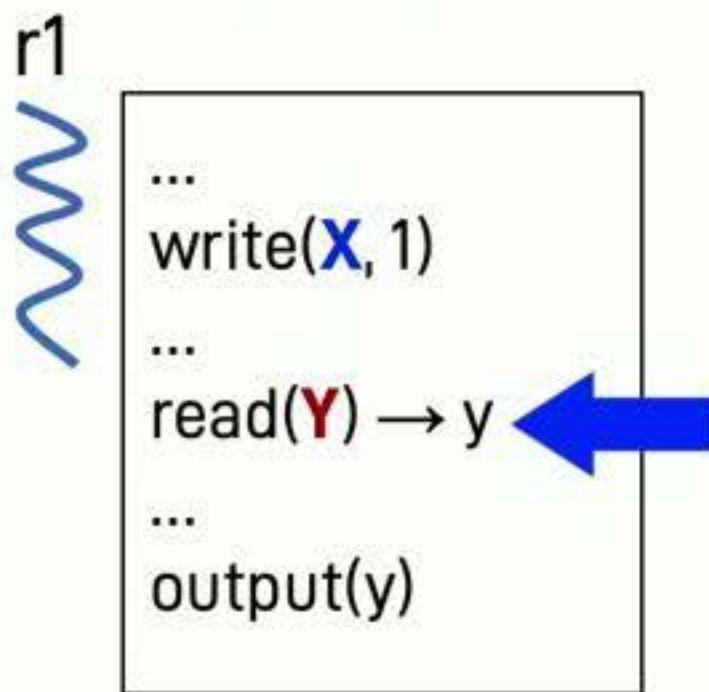
```
...  
write(Y, 1)  
...  
read(X) → x  
...  
output(x)
```



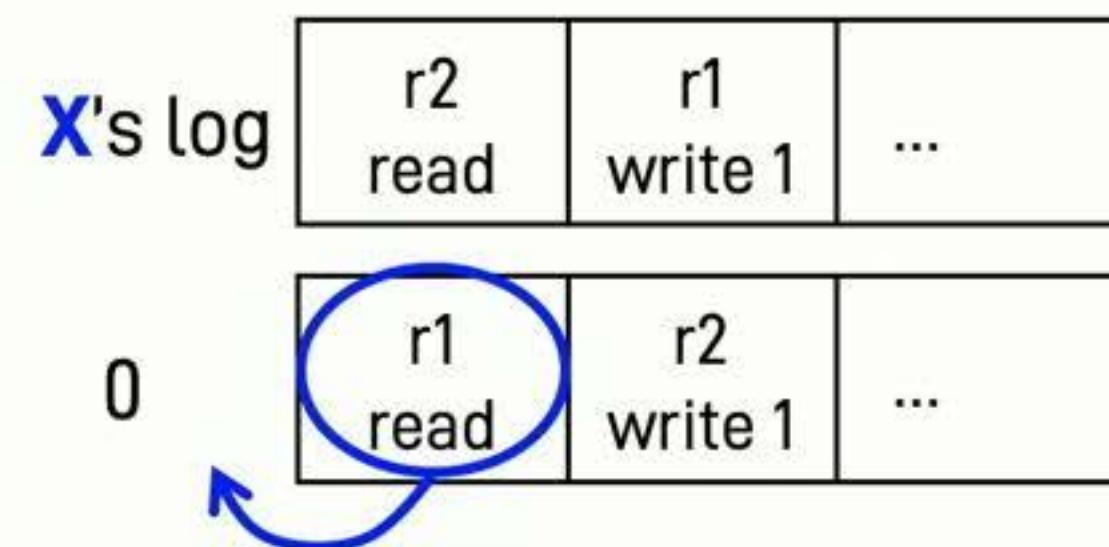
Re-execution according to op logs
wrongly accepts (b)



X=0, Y=0

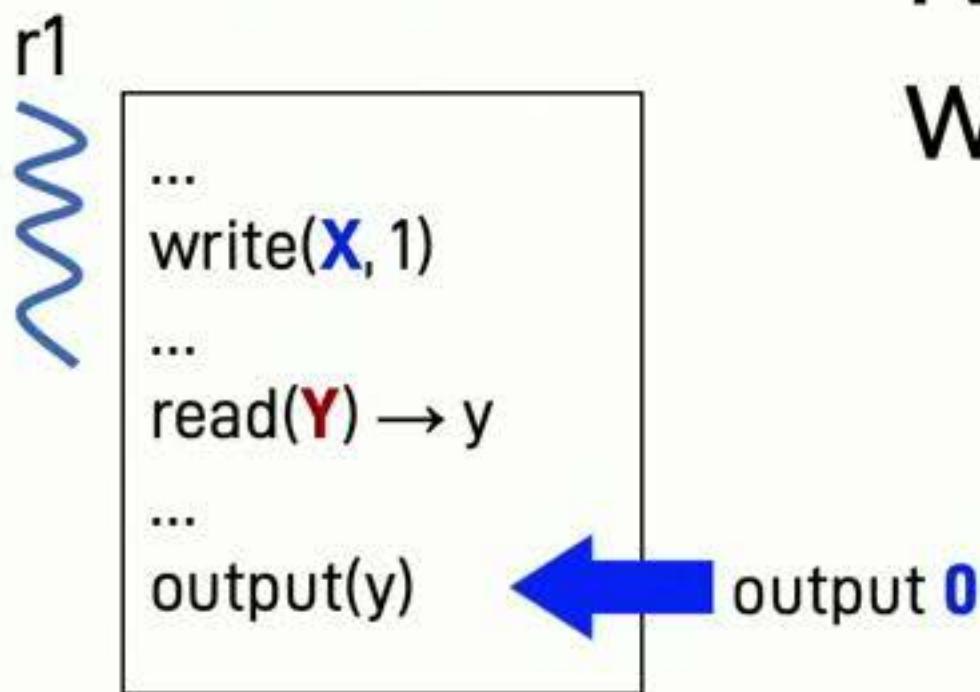


Re-execution according to op logs
wrongly accepts (b)

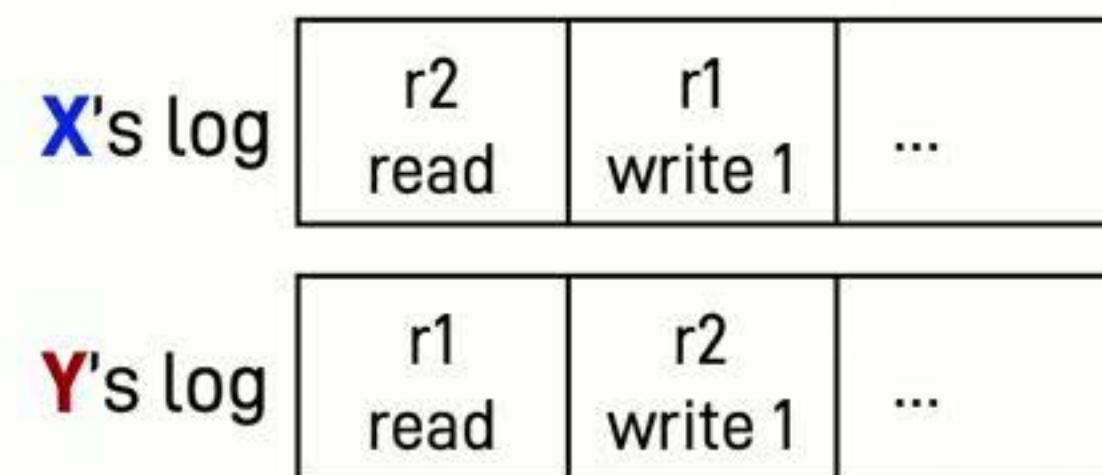
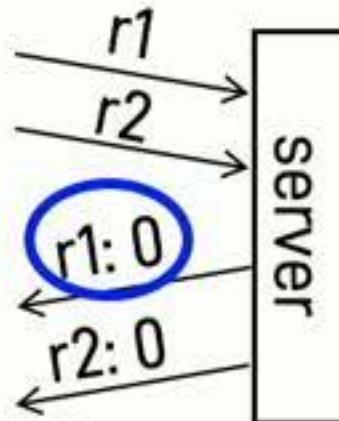
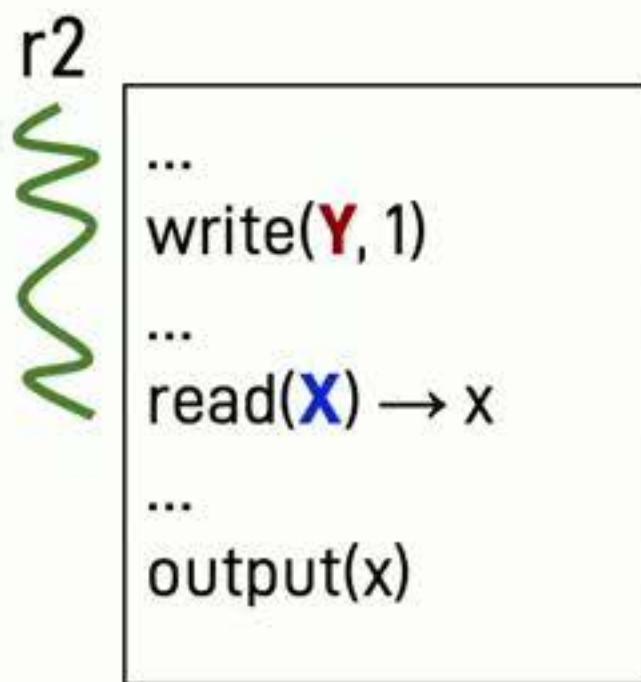


(b)

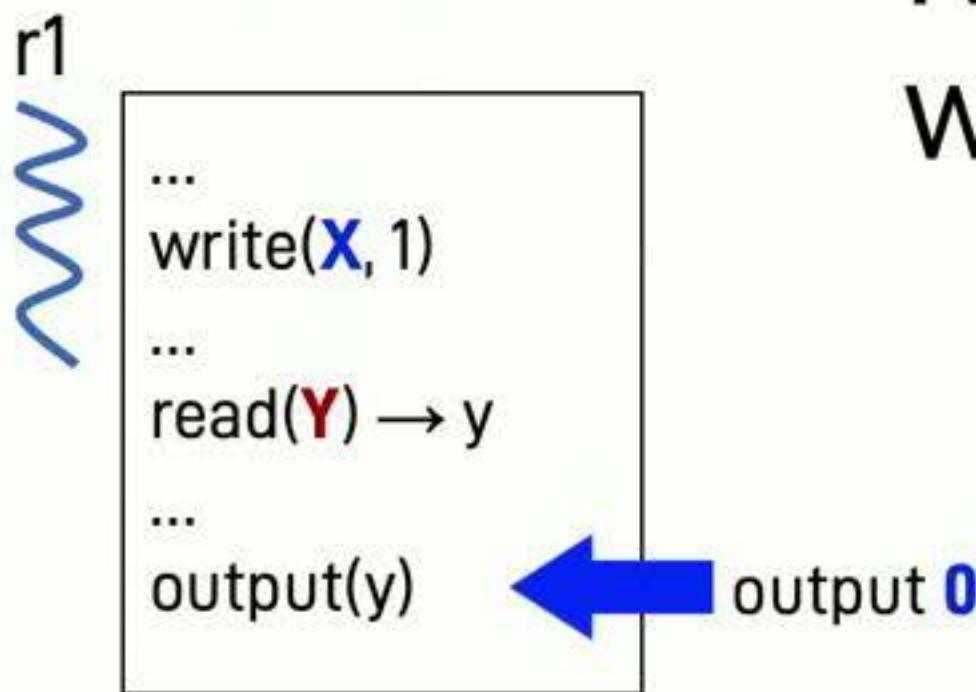
X=0, Y=0



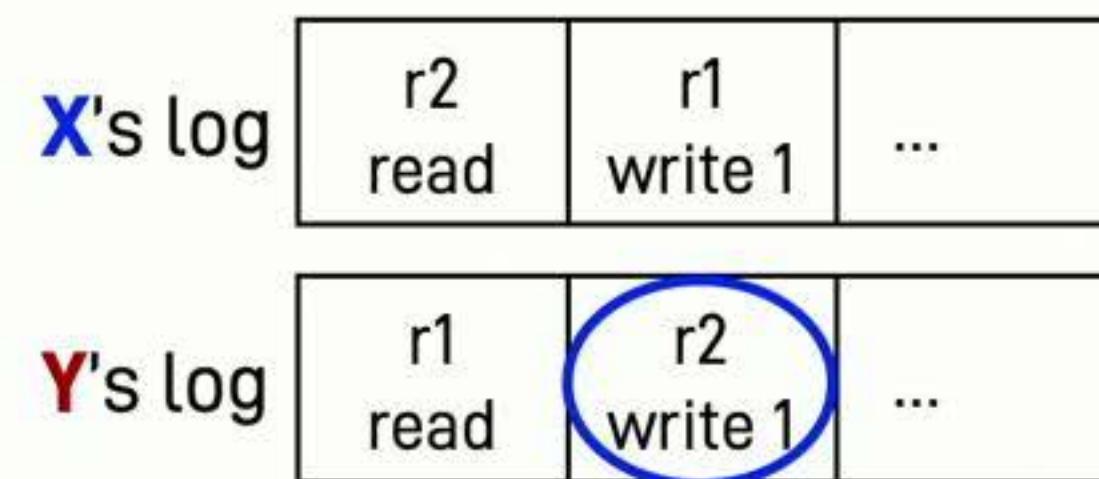
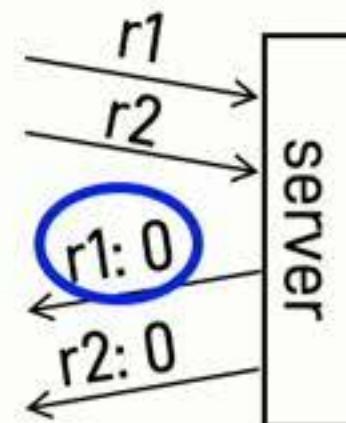
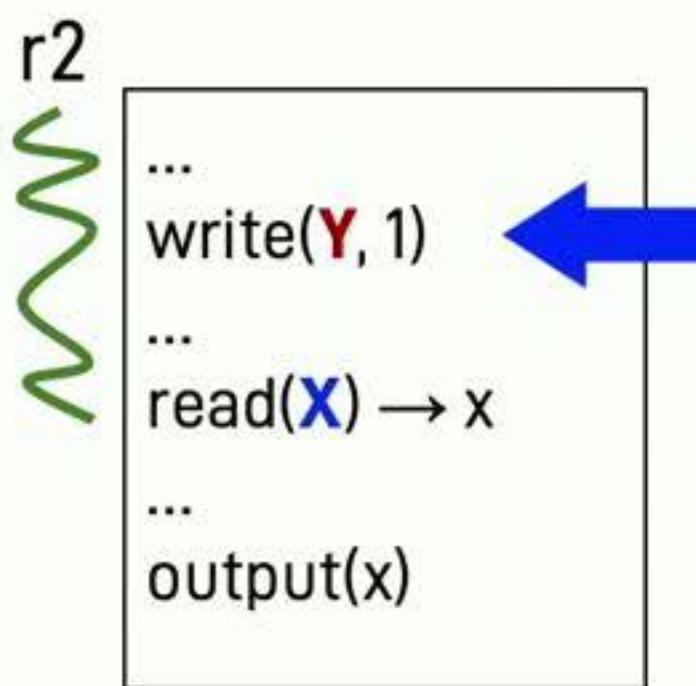
Re-execution according to op logs
wrongly accepts (b)



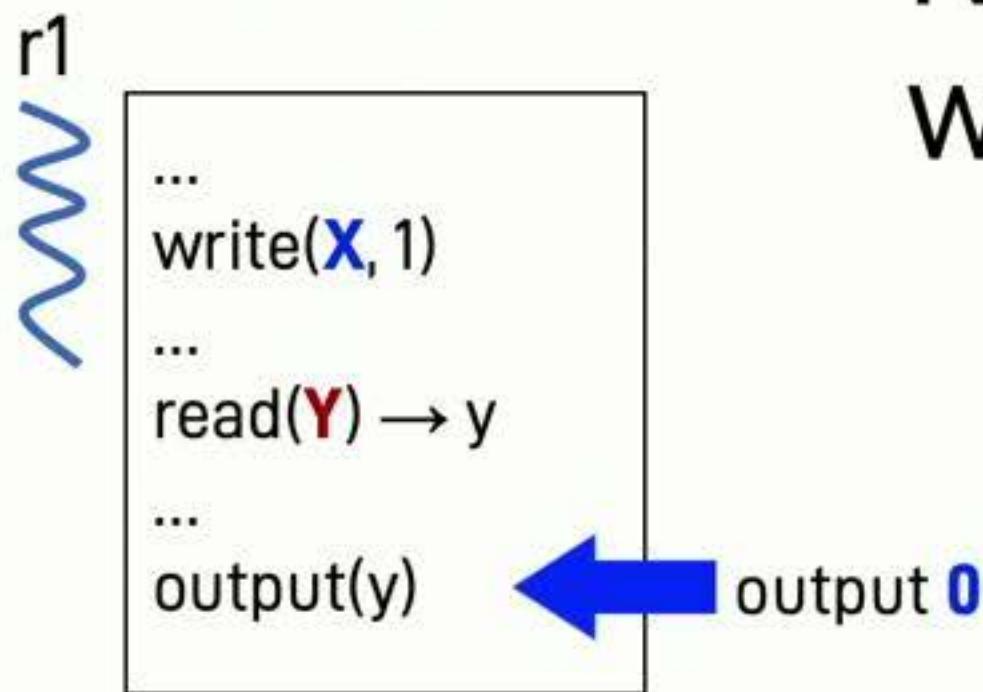
X=0, Y=0



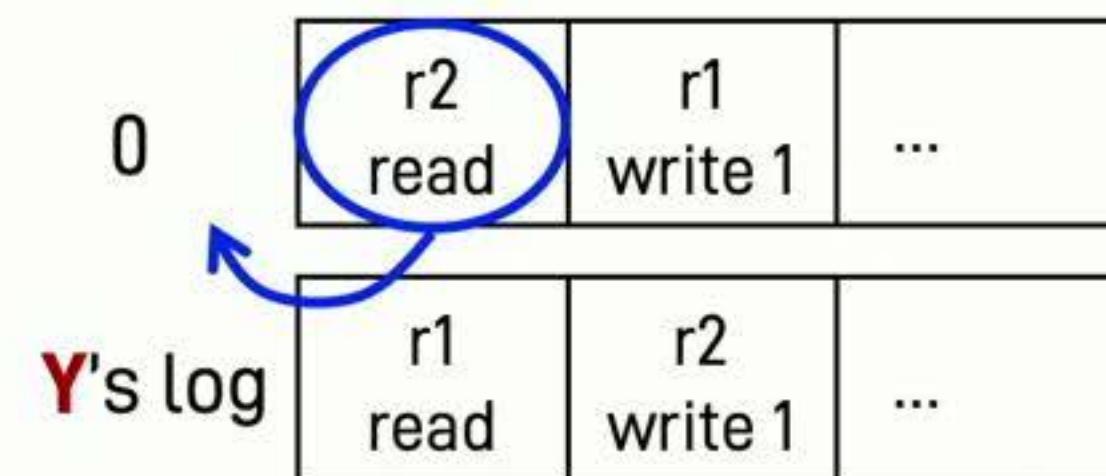
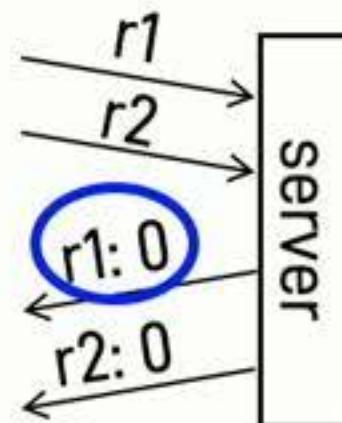
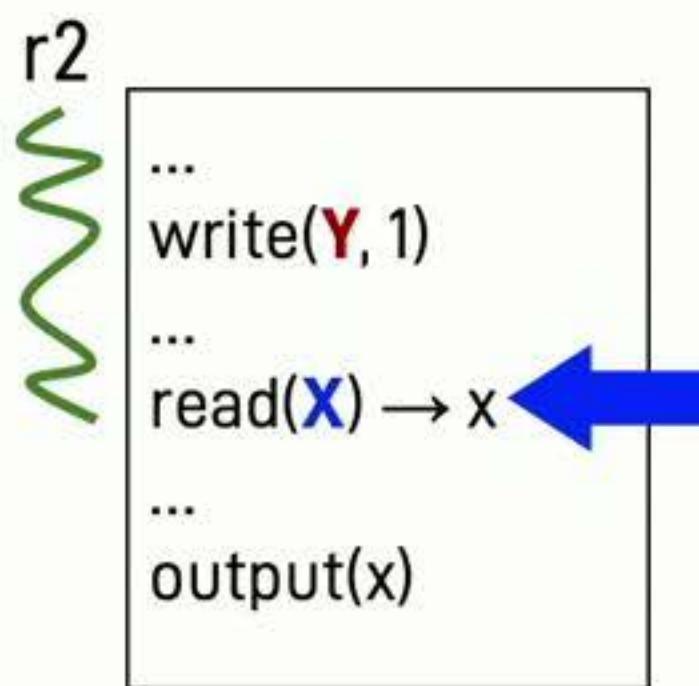
Re-execution according to op logs
wrongly accepts (b)



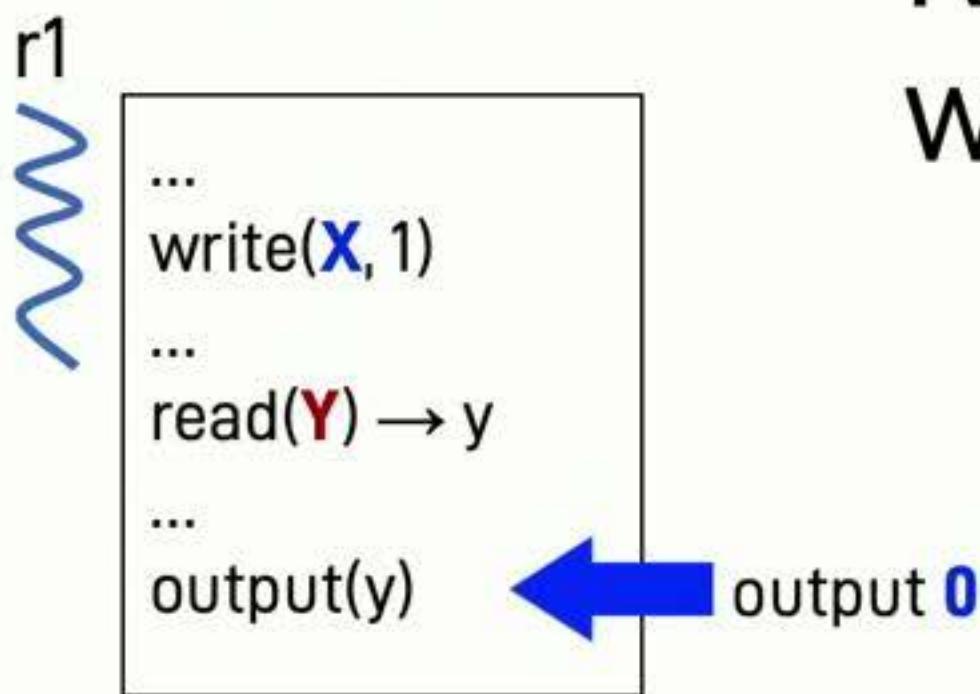
X=0, Y=0



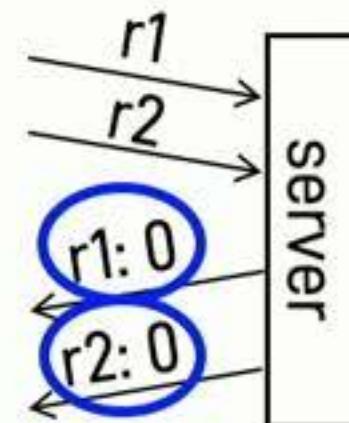
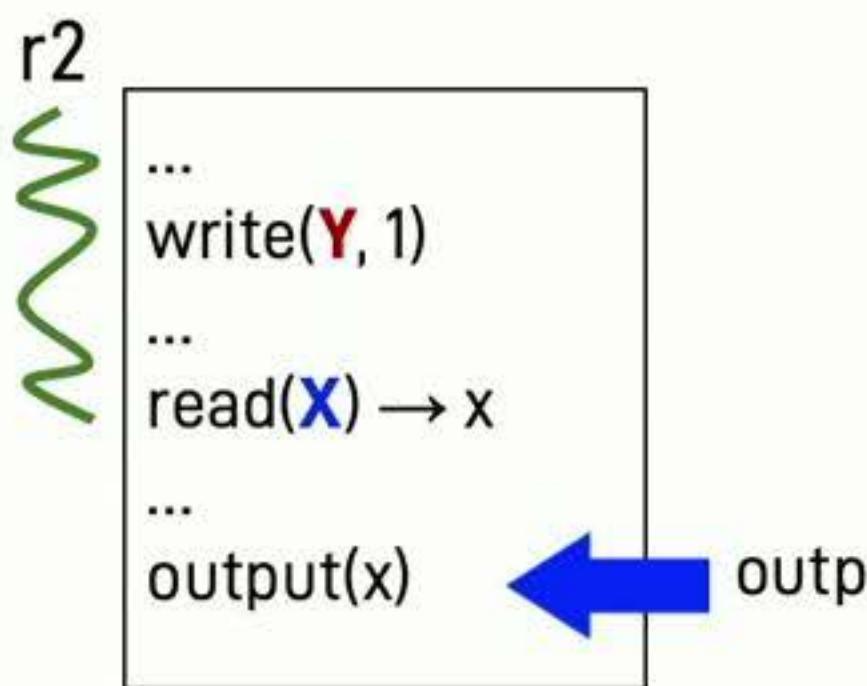
Re-execution according to op logs
wrongly accepts (b)



X=0, Y=0



Re-execution according to op logs
wrongly accepts (b)



X's log

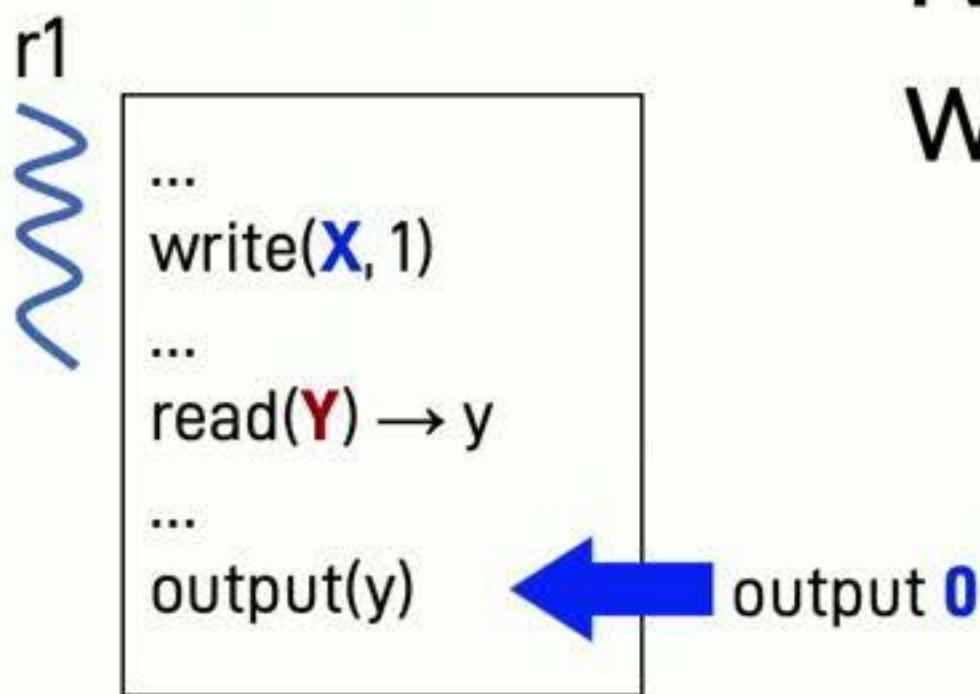
r2	read	r1	write 1	...
----	------	----	---------	-----

Y's log

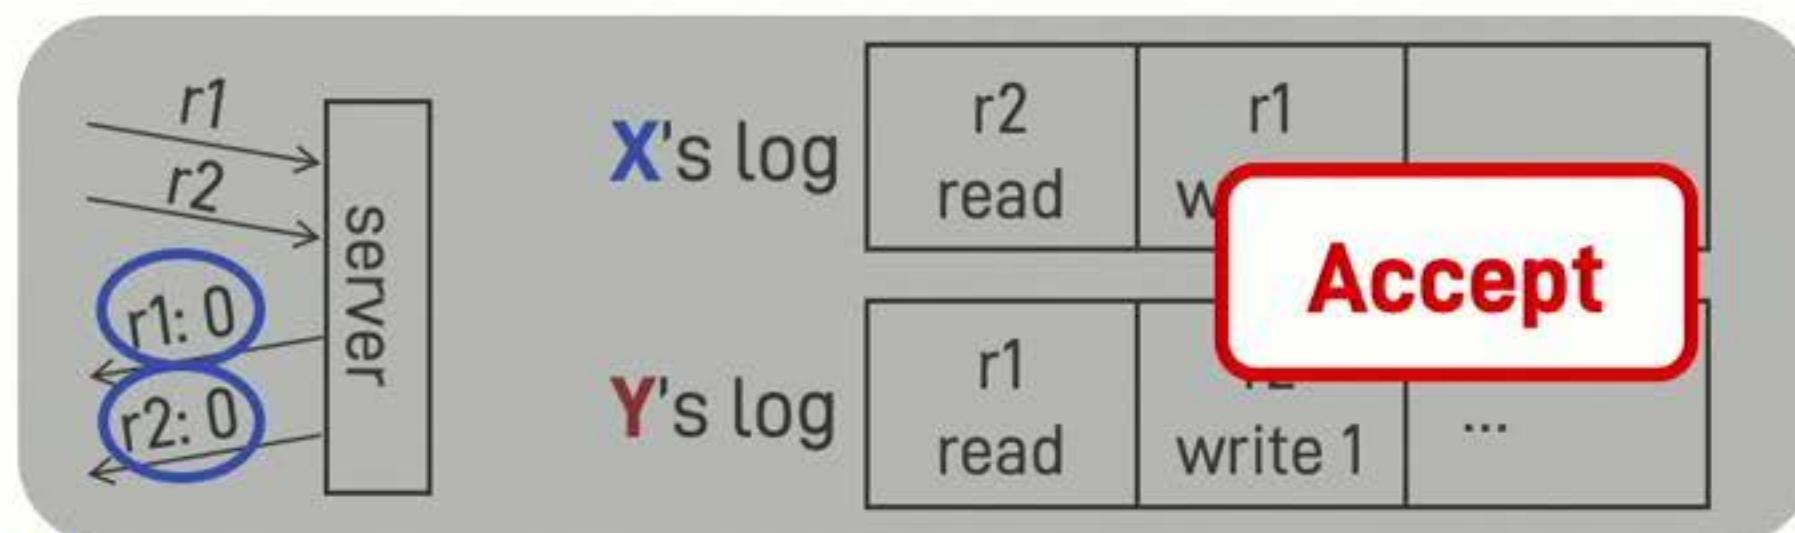
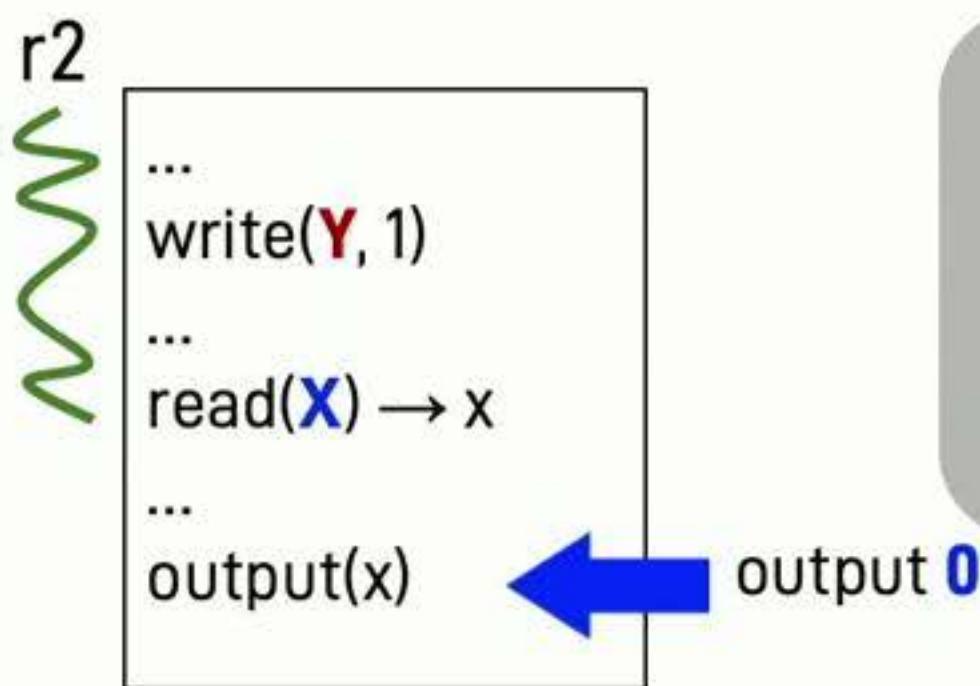
r1	read	r2	write 1	...
----	------	----	---------	-----

(b)

$X=0, Y=0$

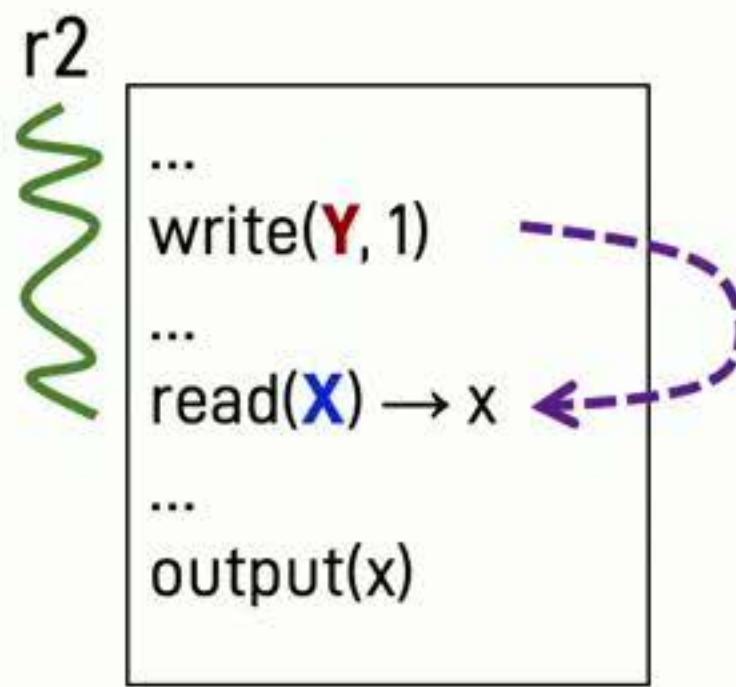
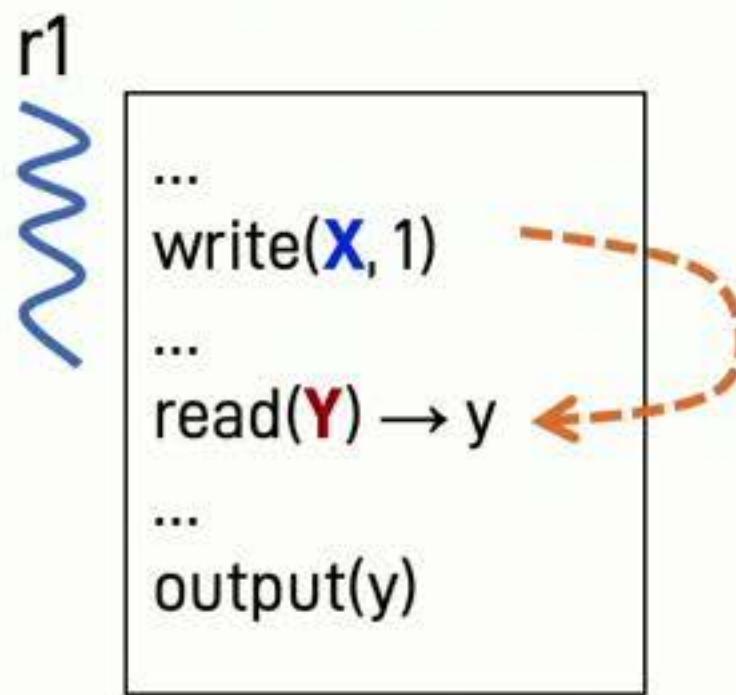


Re-execution according to op logs
wrongly accepts (b)

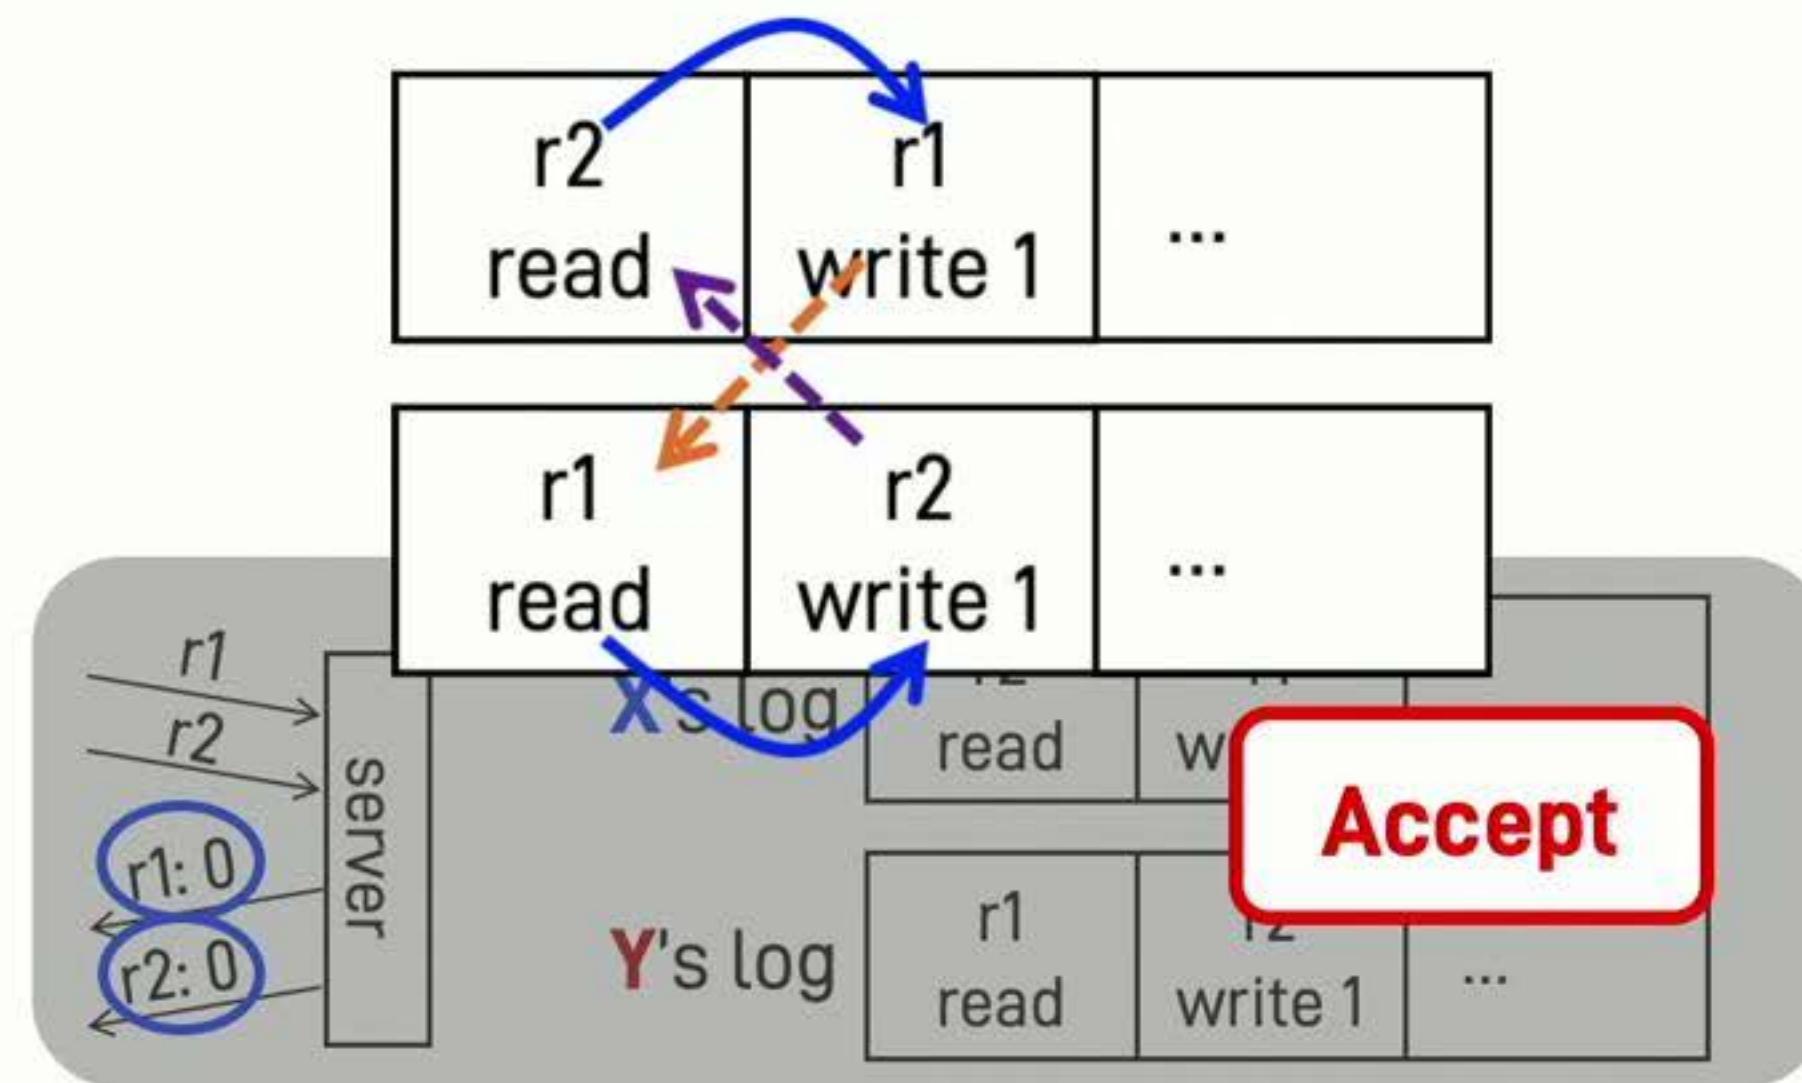


(b) X

X=0, Y=0



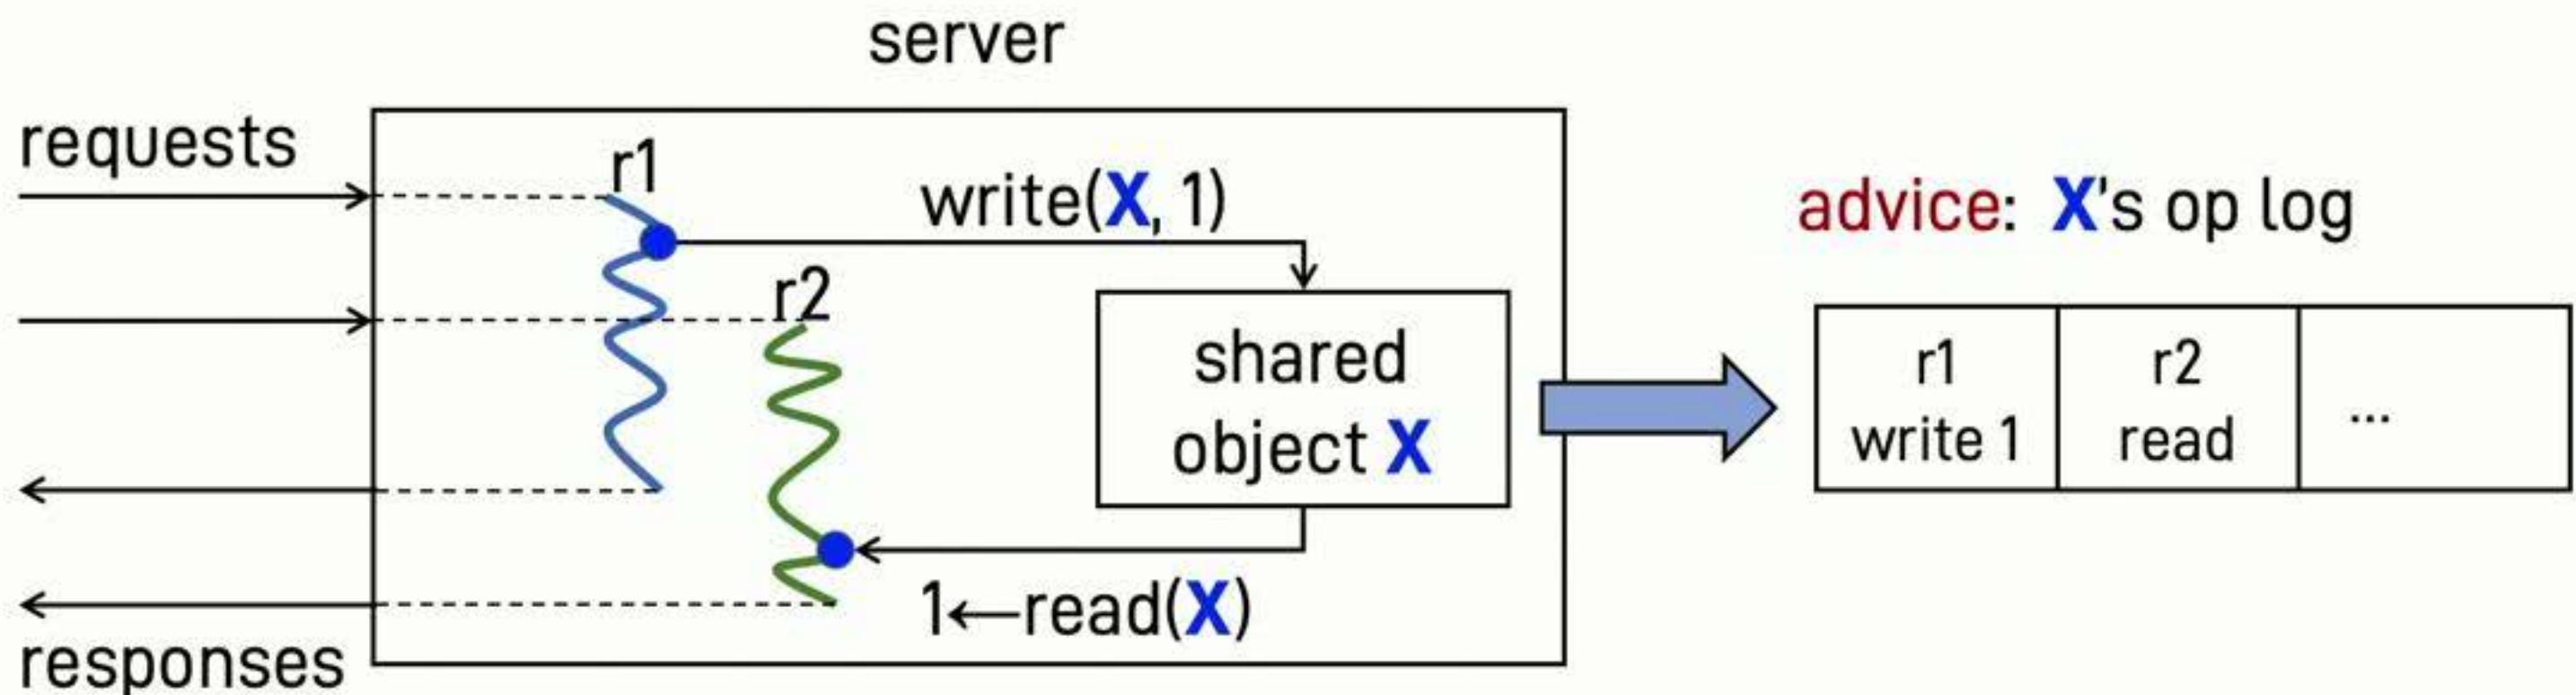
Re-execution according to op logs
wrongly accepts (b)



Solution must consider trace and program

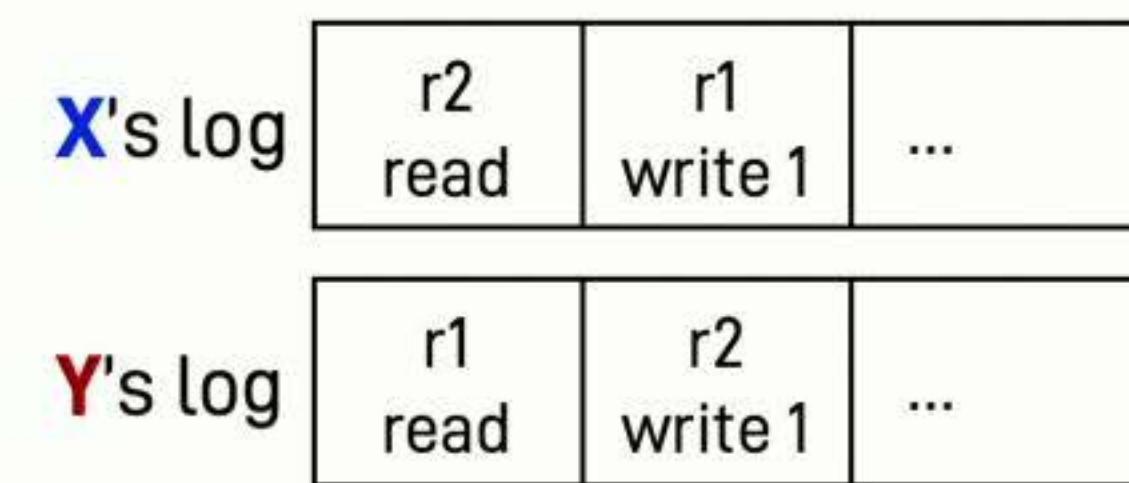
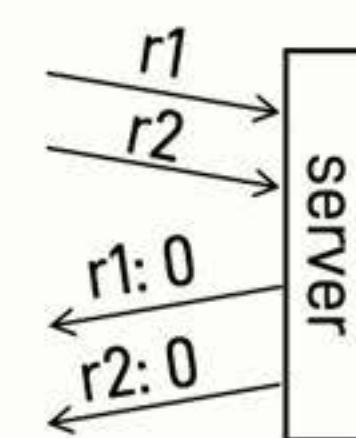
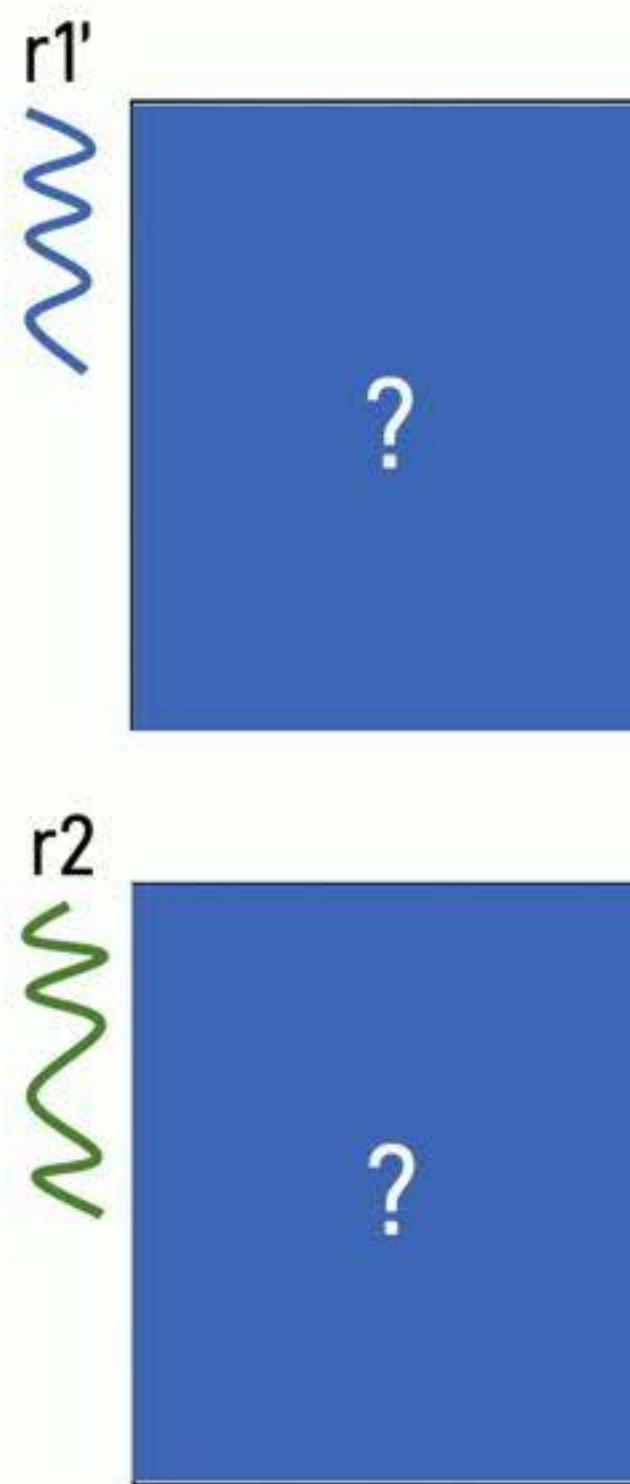
- Validating the **alleged op logs** requires
 - request order (from trace)
 - program order (from program)
 - operation order (from **op logs**)
- **Consistent ordering verification** builds a graph...
...that includes all info above and check acyclicity

Concurrency model



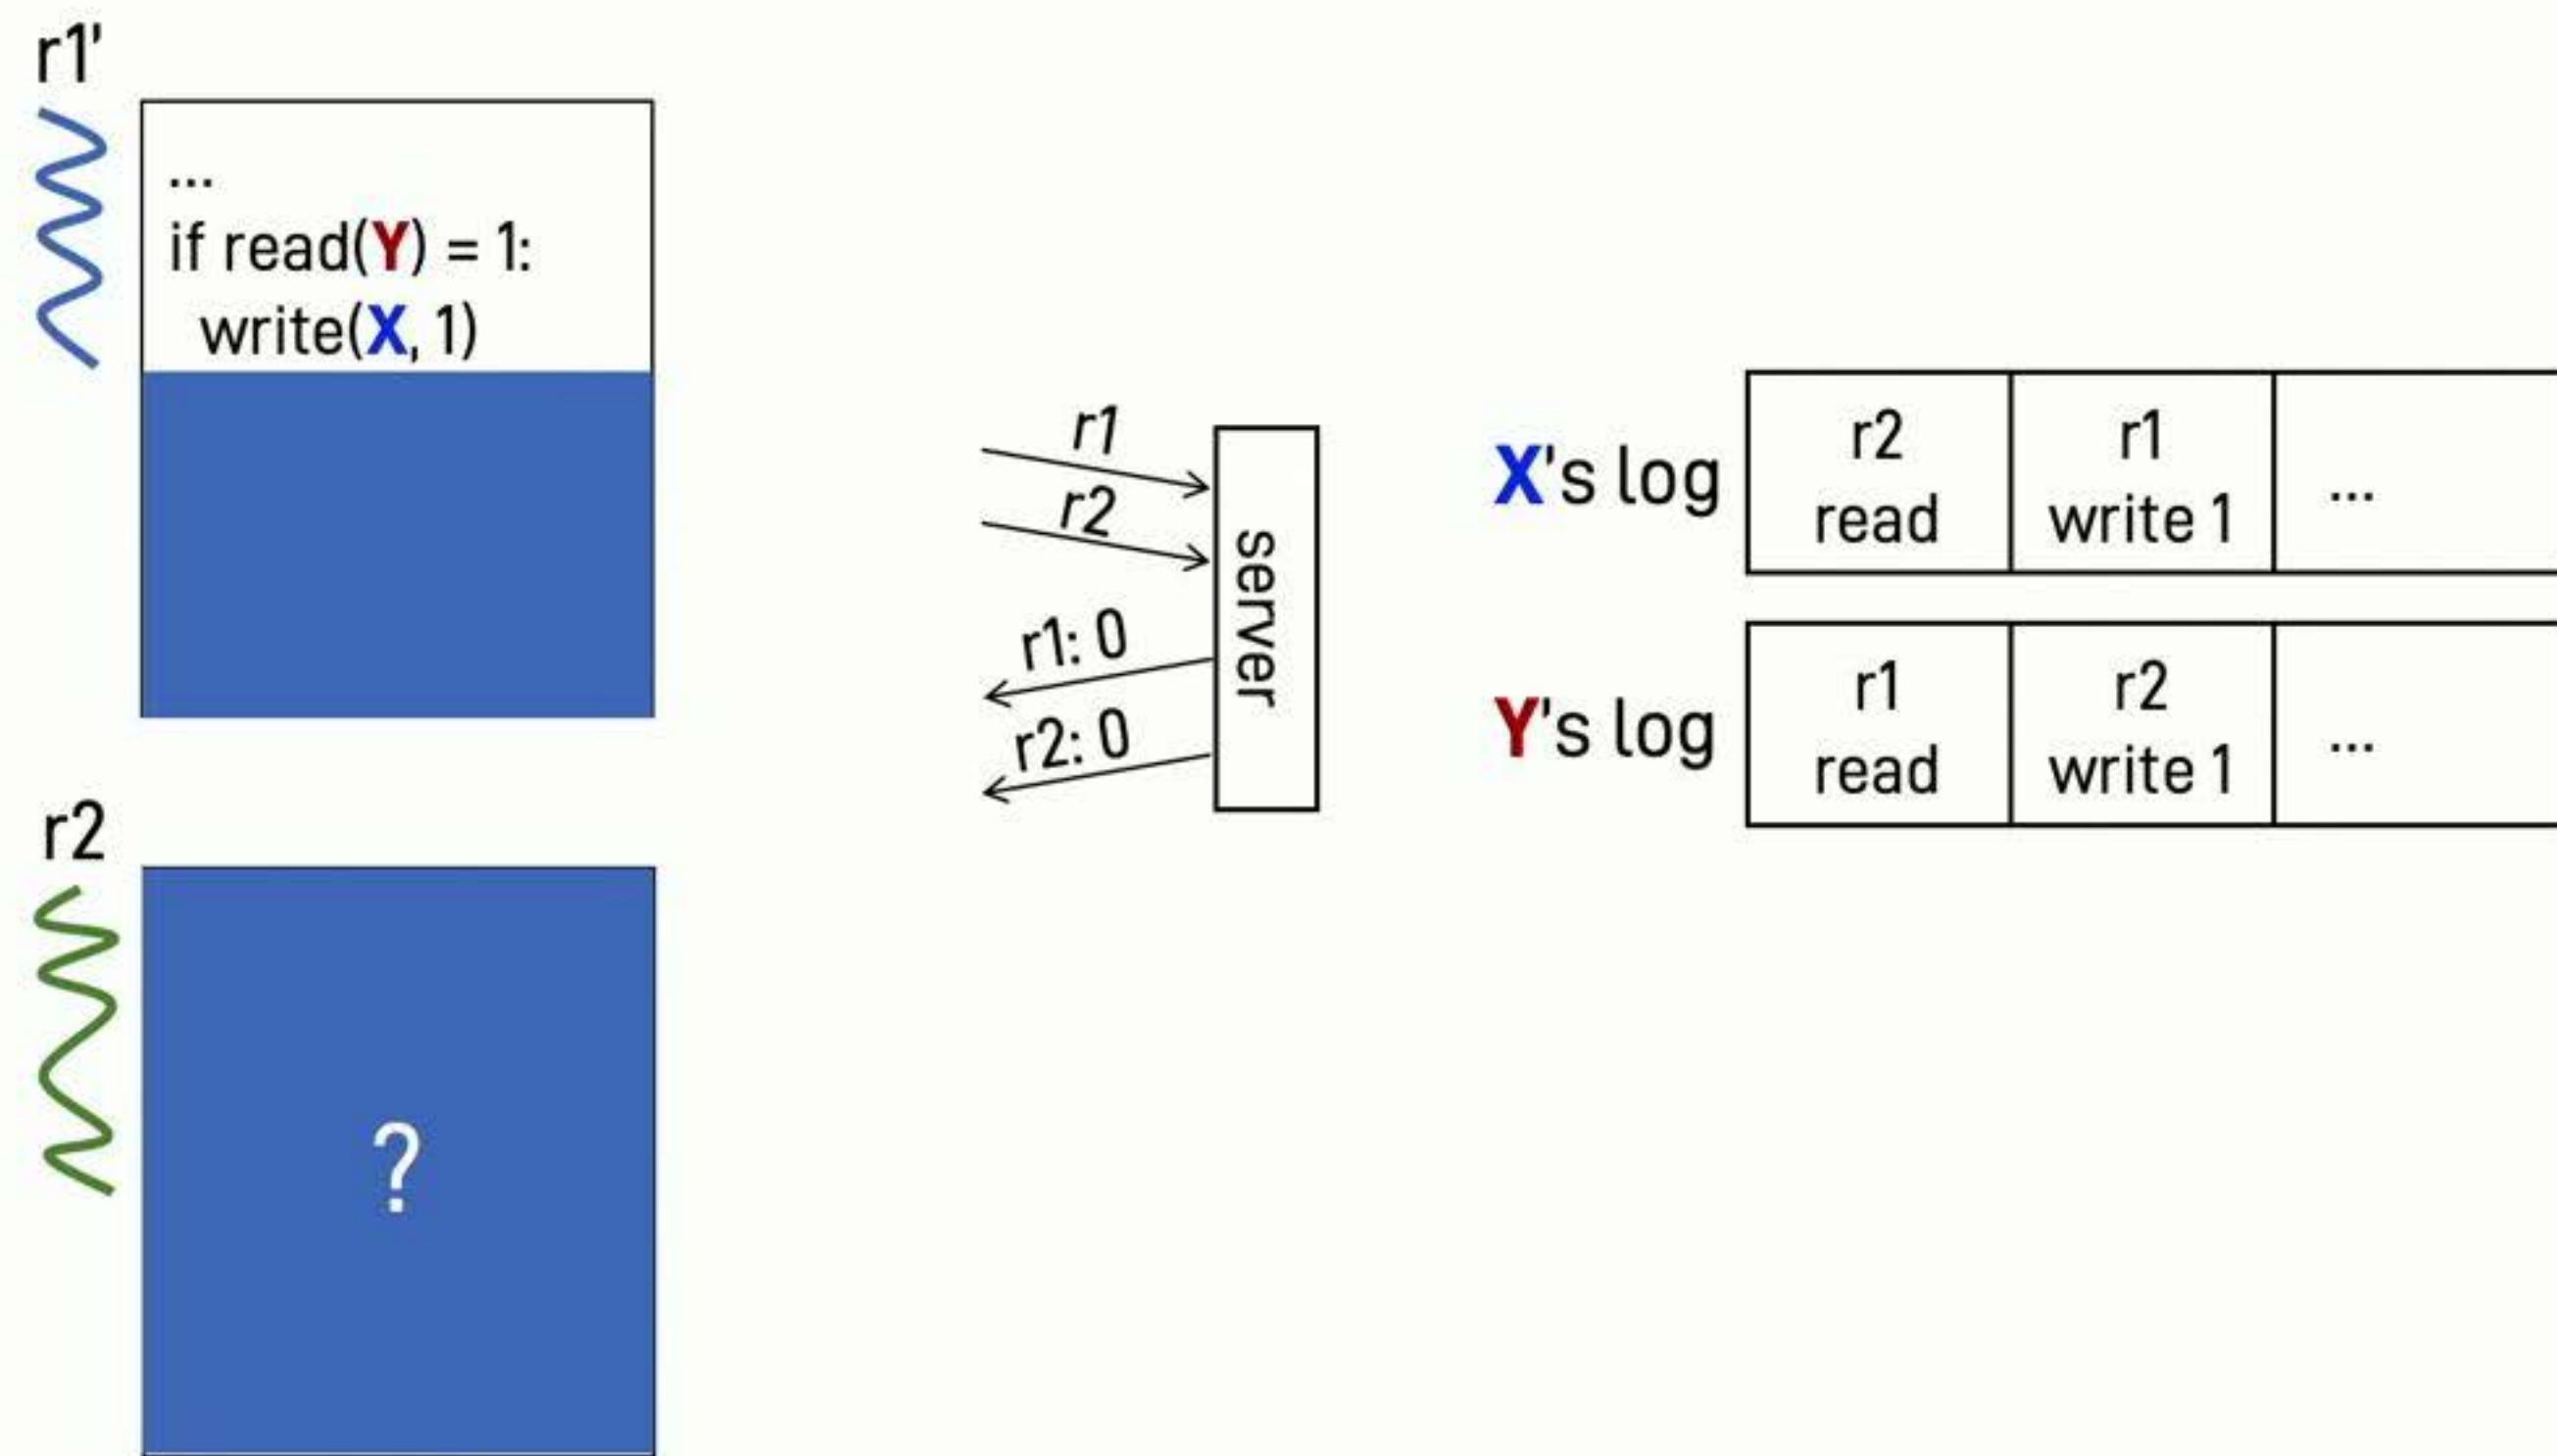
- What about simple re-execution according to op logs?

However, the actual problem is harder

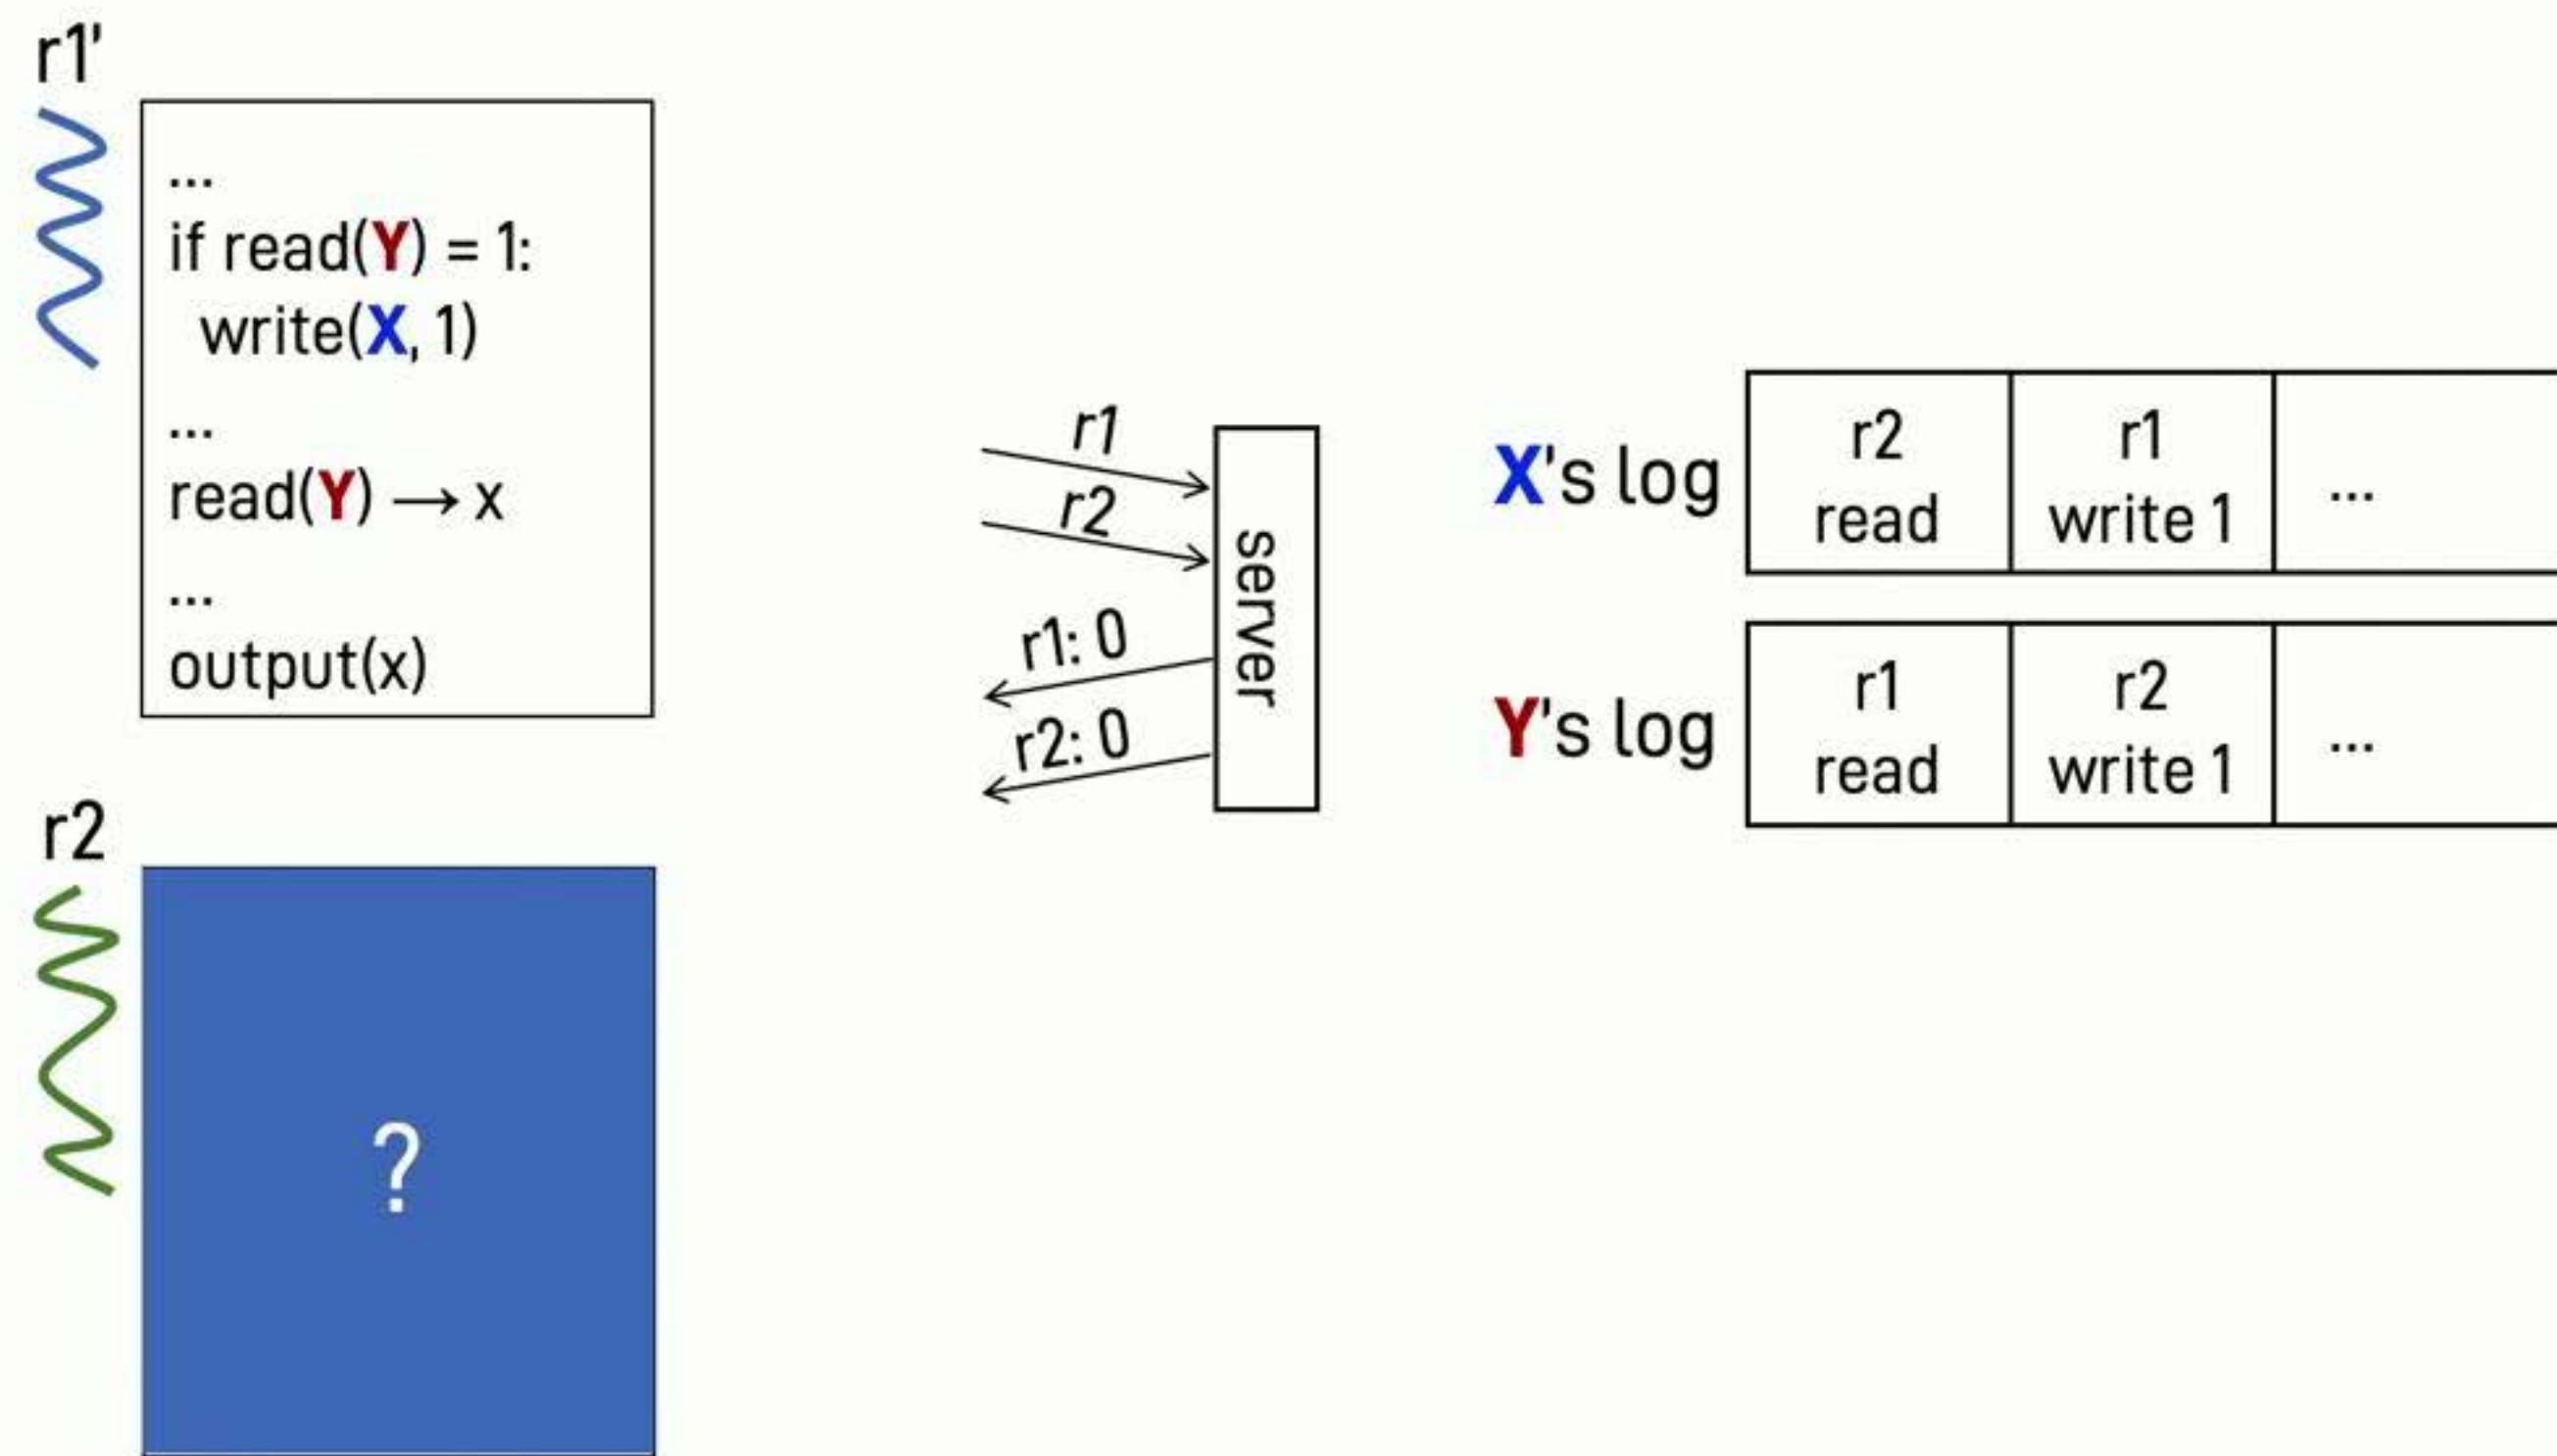


(b)

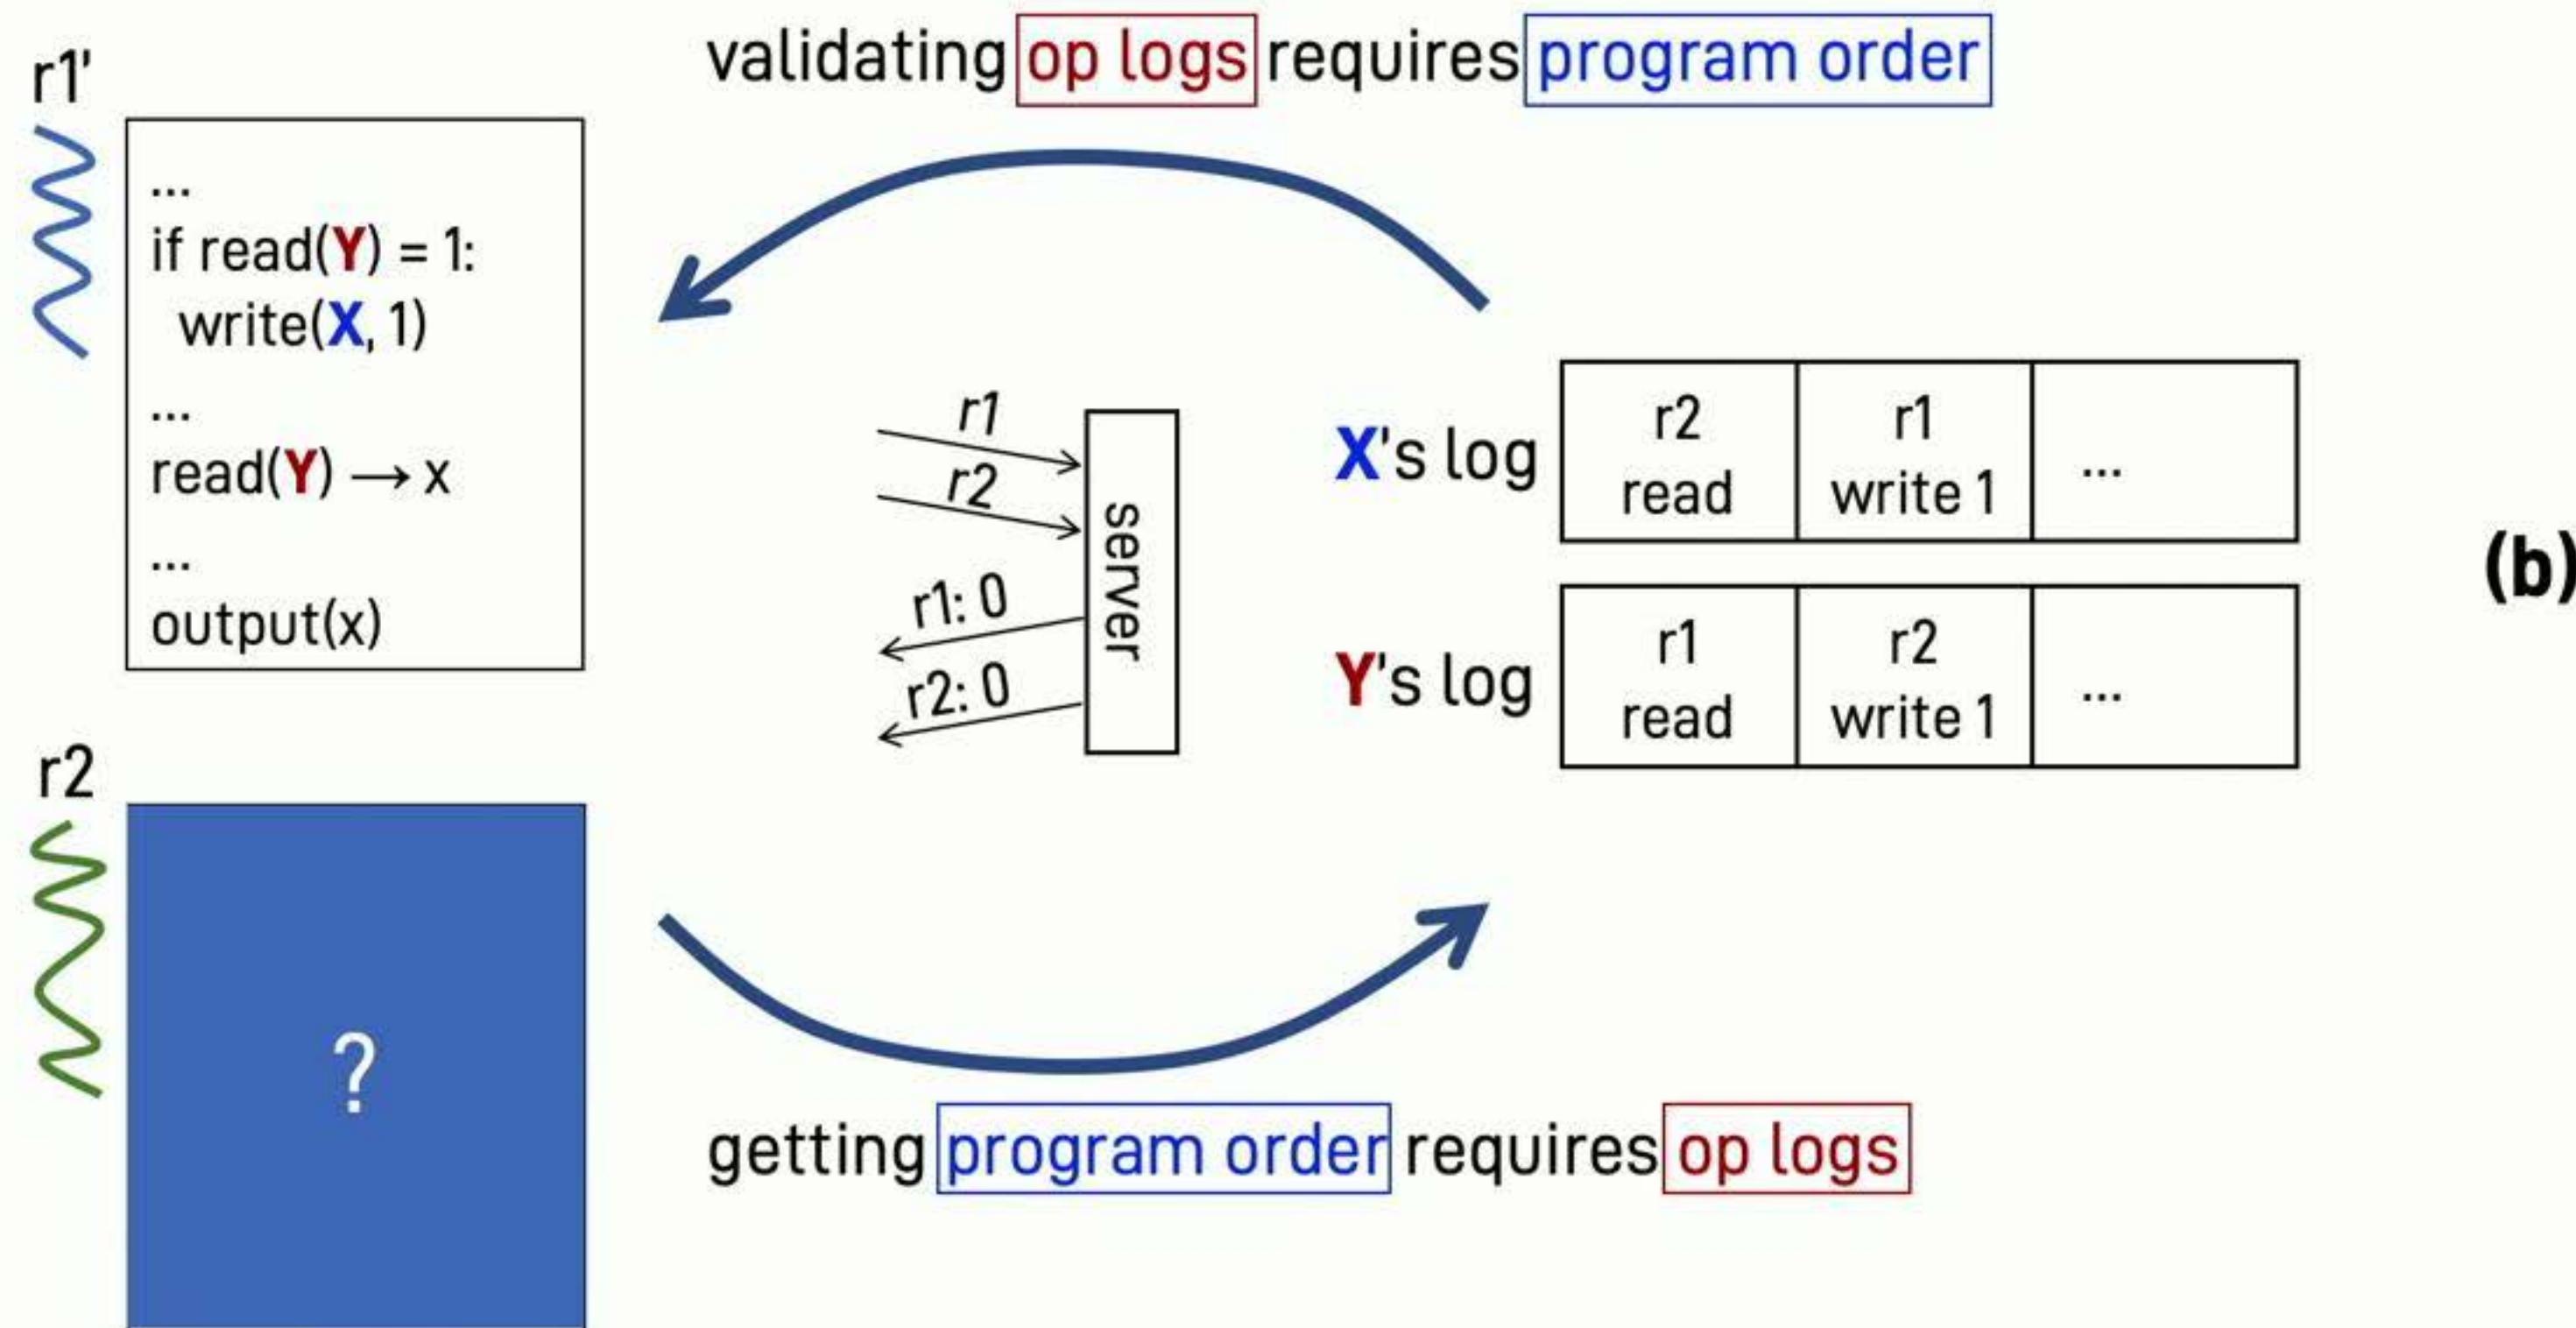
However, the actual problem is harder



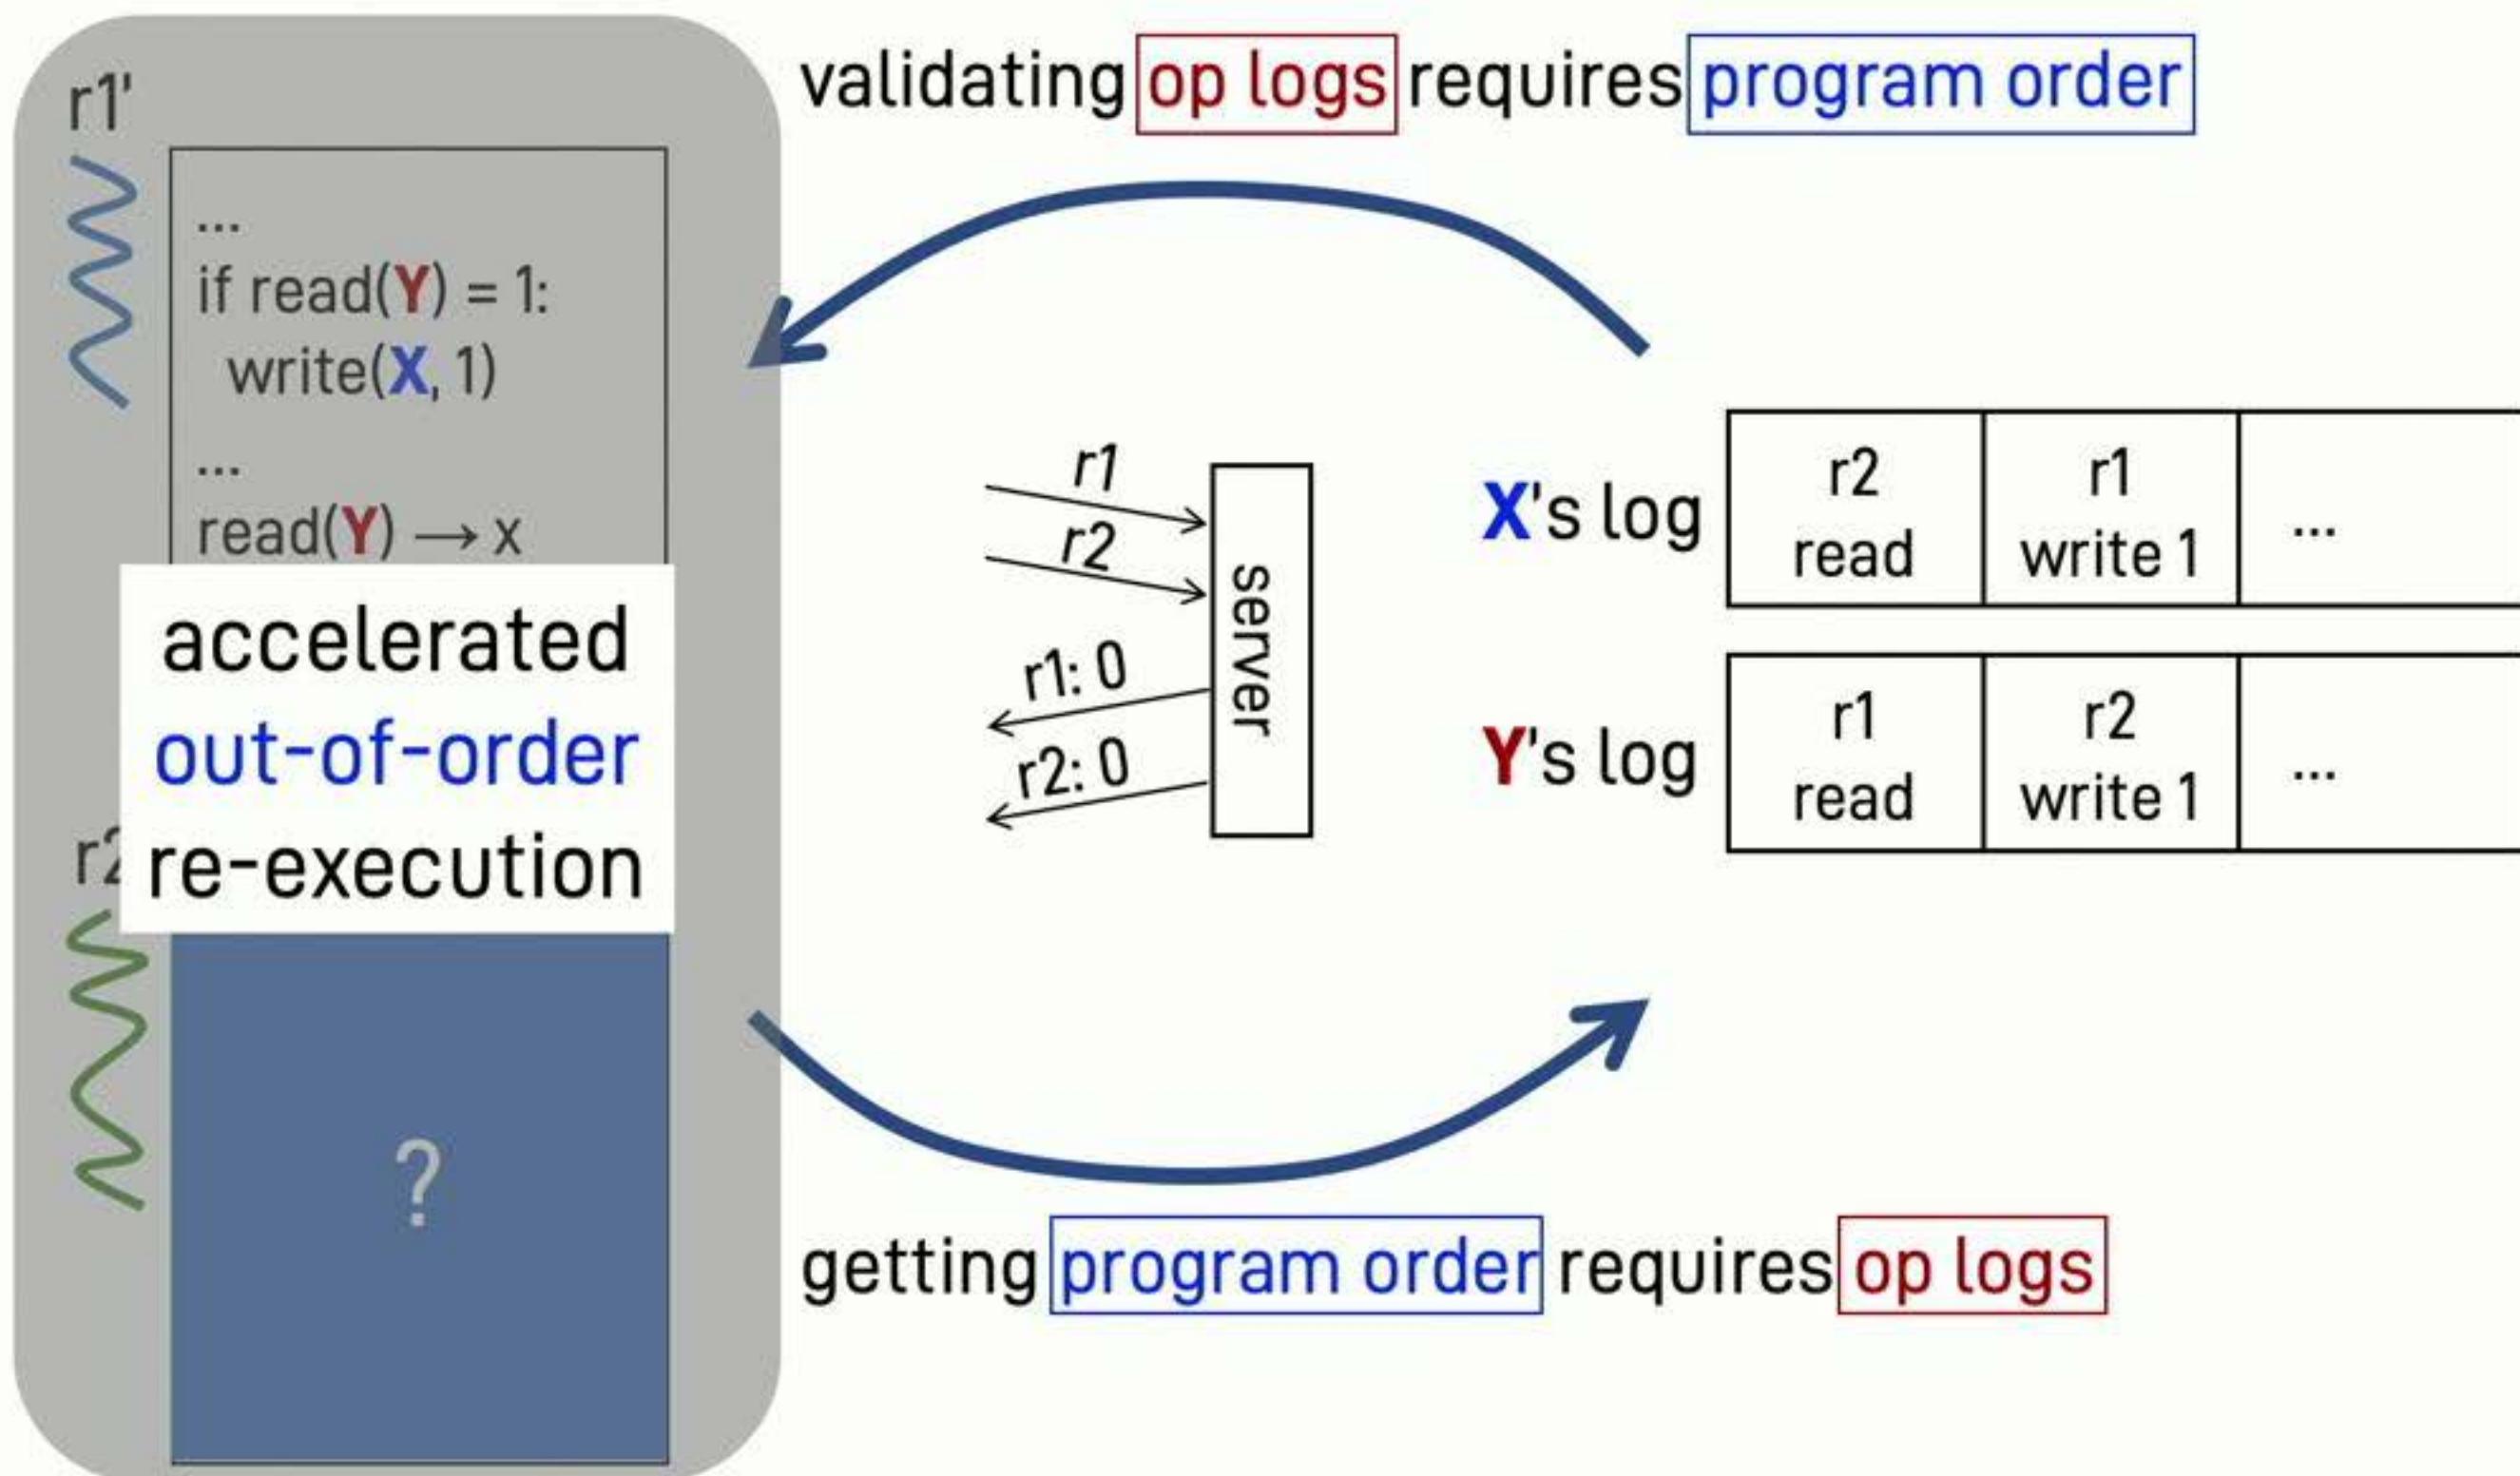
However, the actual problem is harder



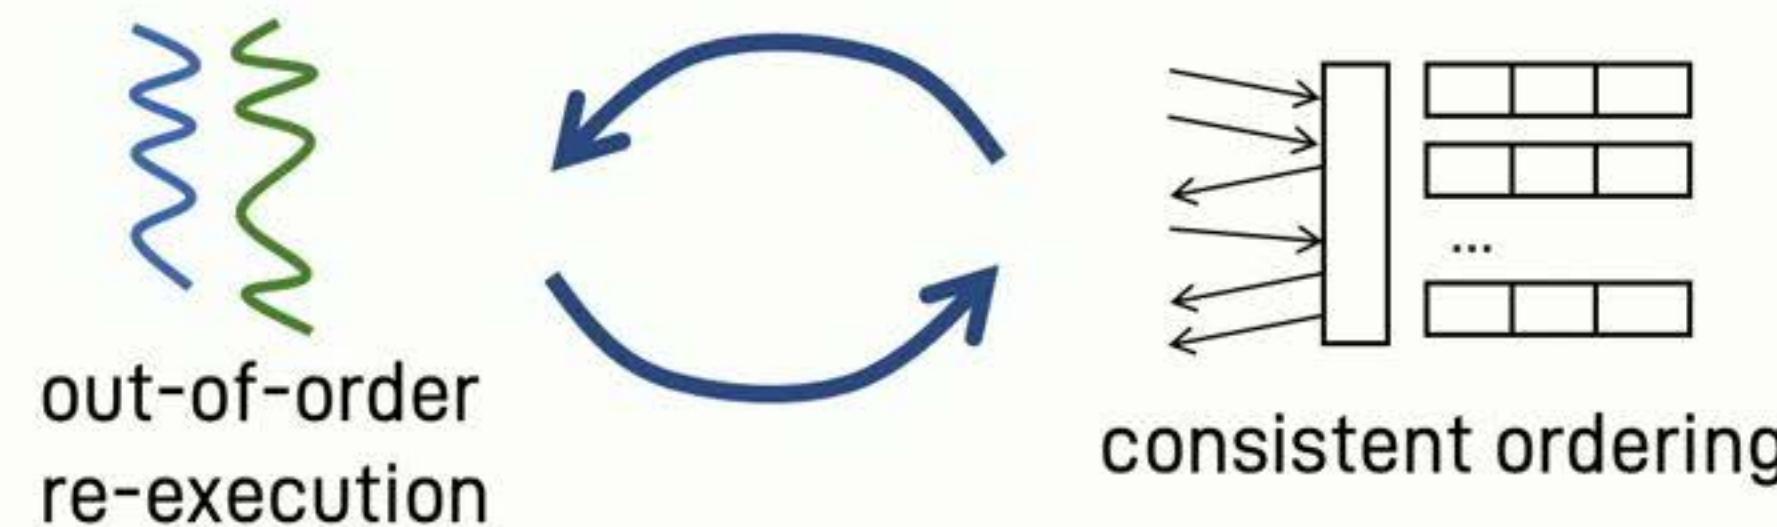
However, the actual problem is harder



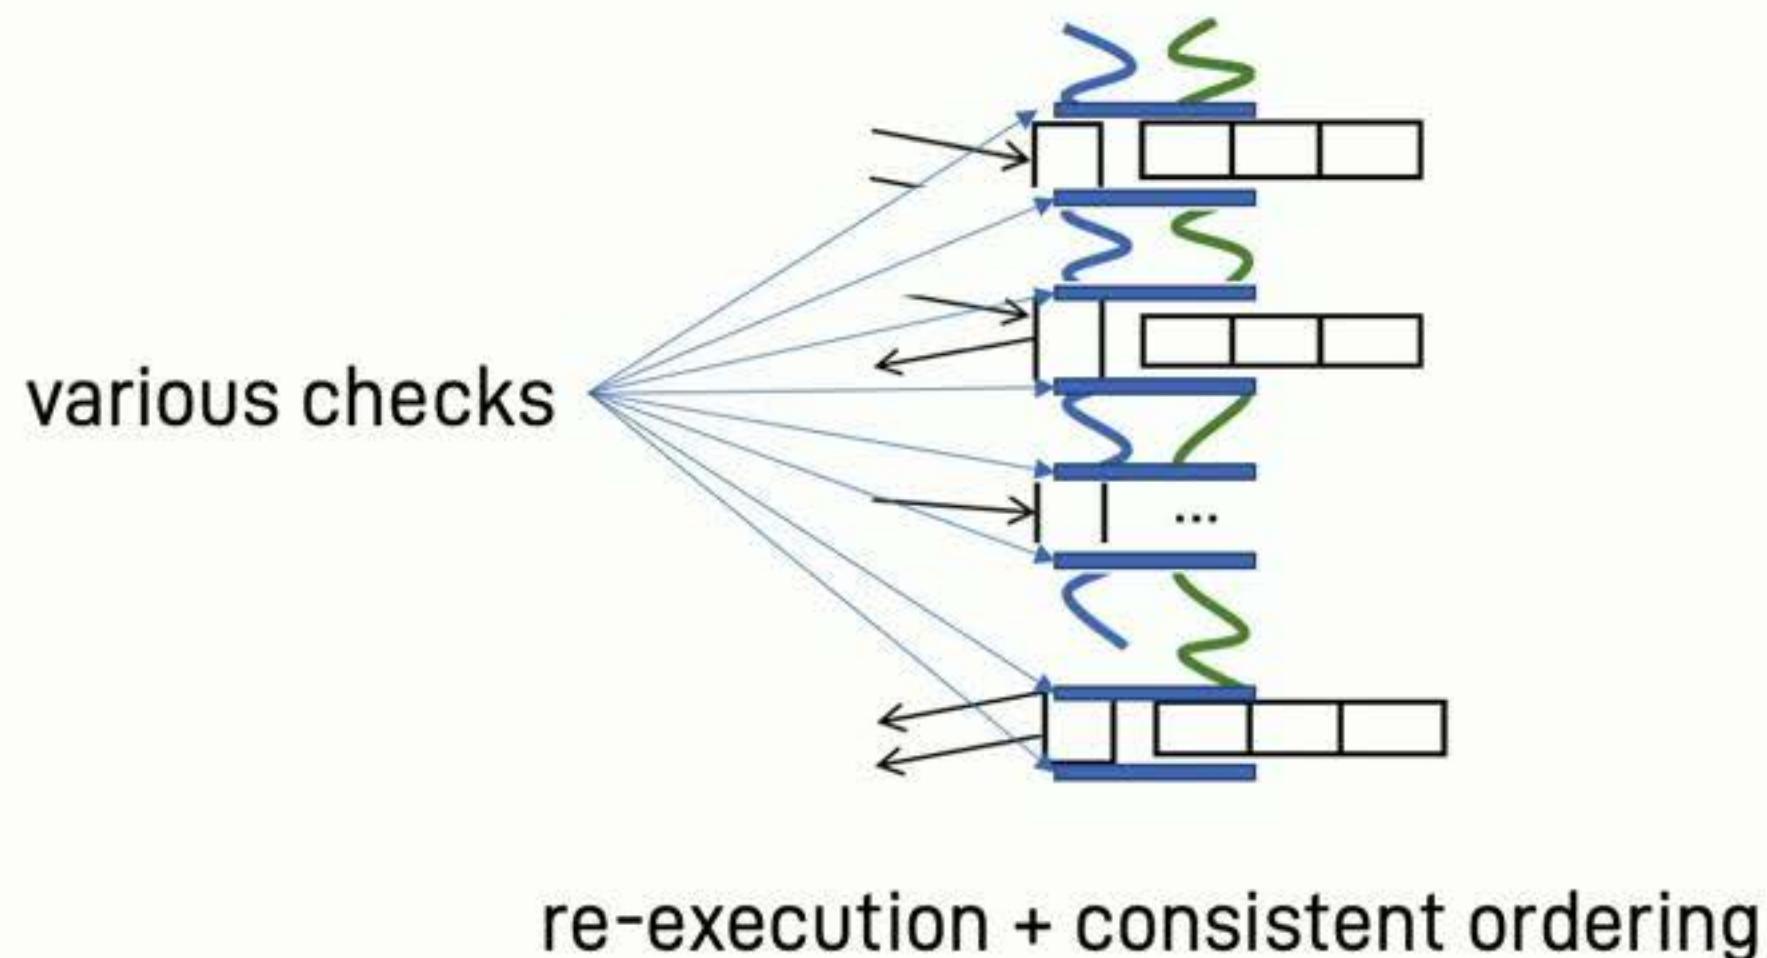
However, the actual problem is harder



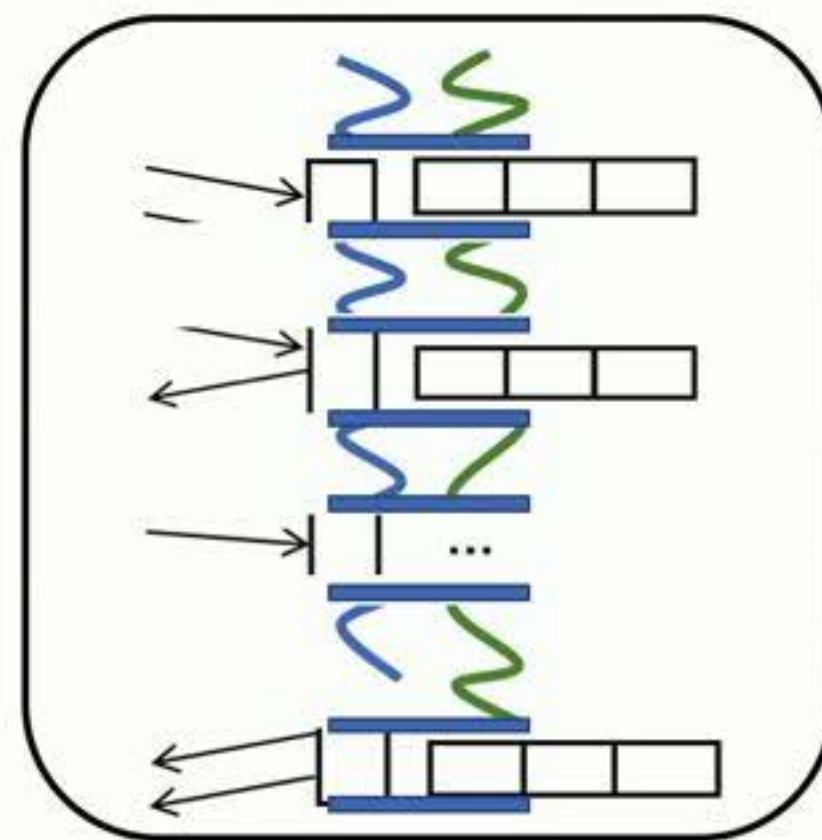
Orochi has a co-designed verification protocol



Orochi has a co-designed verification protocol



Orochi has a co-designed verification protocol



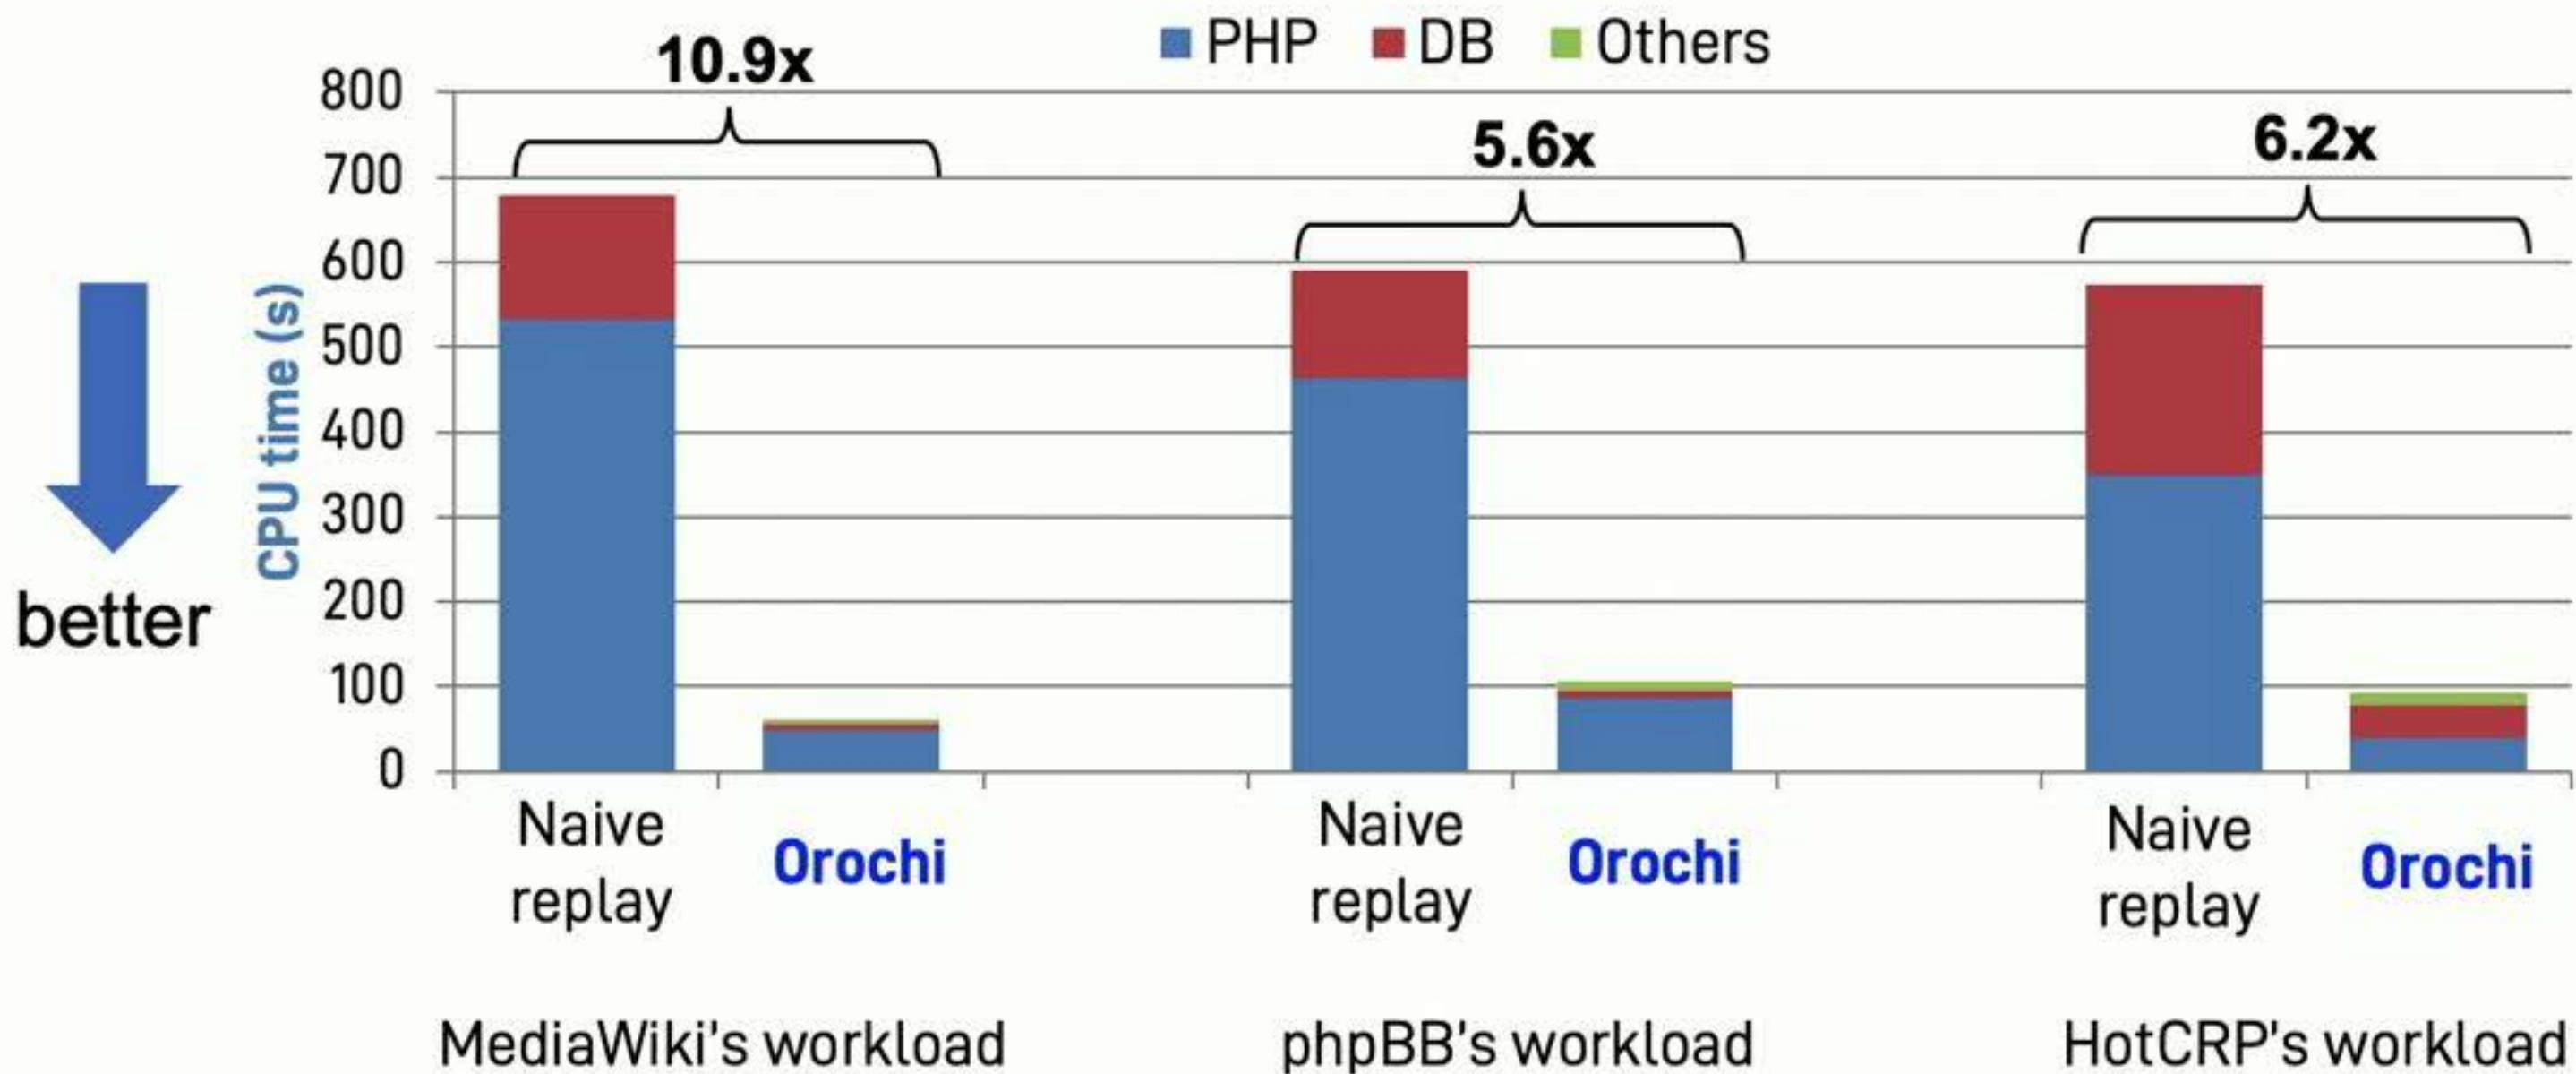
- Verification protocol is proved to be **correct** ...
... meaning { Completeness: honest server \Rightarrow verifier accepts
Soundness: verifier accepts \Rightarrow honest server

Evaluation setup

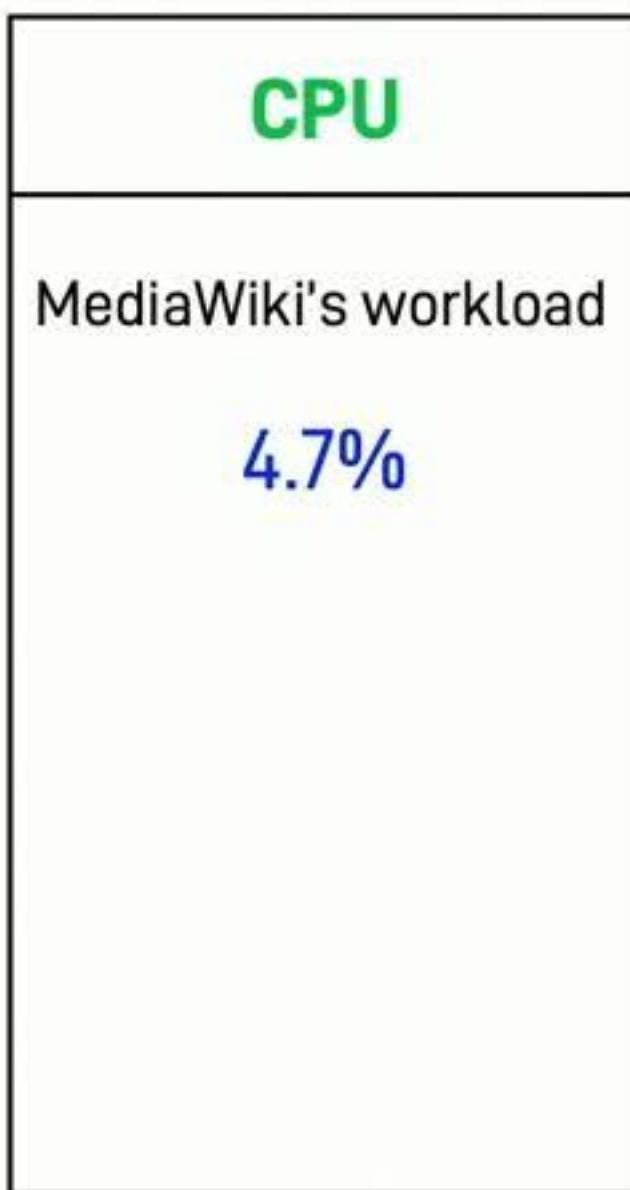
- Applications:
 - MediaWiki, phpBB and HotCRP
- Workloads:
 - MediaWiki: Wikipedia 2007 trace
 - phpBB: 7-day's posts from CentOS forum
 - HotCRP: Simulation of SIGCOMM'09

Is the verifier efficient?

Orochi's verifier achieves speedups compared to naive replay



What are the (server's) CPU/network/storage costs?



What are the (server's) CPU/network/storage costs?

CPU	Network		
MediaWiki's workload	trace (per req)	advice (per req)	Orochi's overhead
4.7%	MediaWiki's workload	7.1KB	1.7KB 11.4%

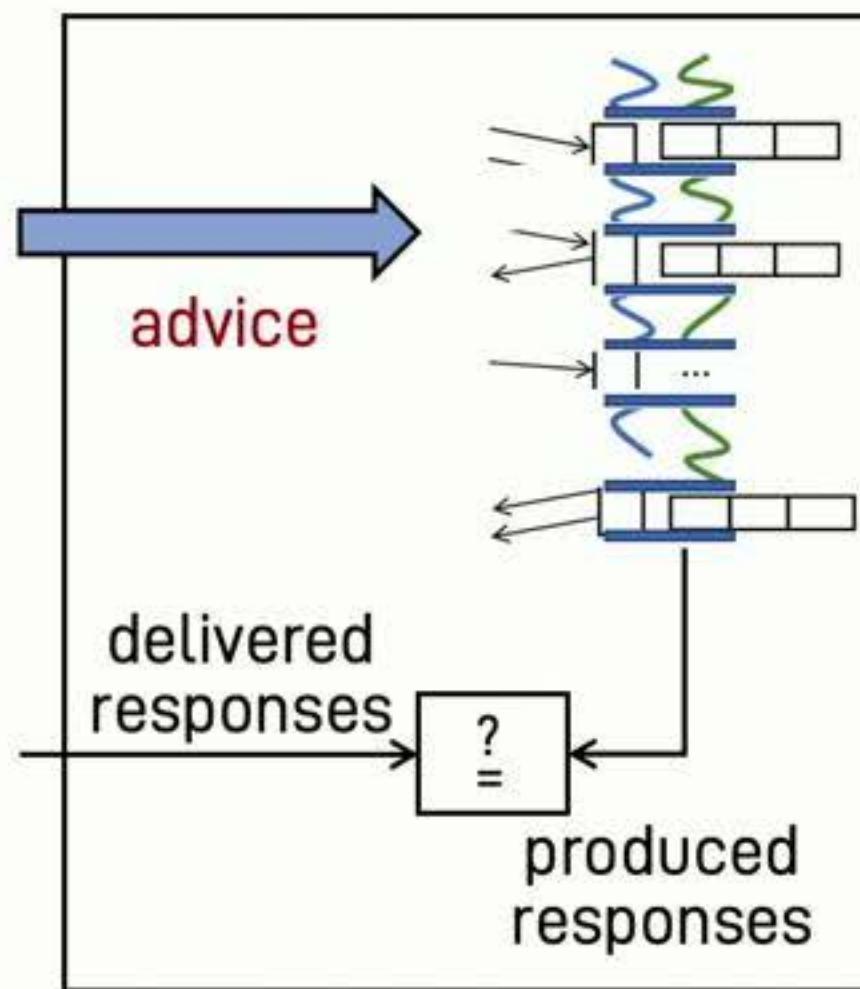
What are the (server's) CPU/network/storage costs?

CPU	Network			Storage
MediaWiki's workload	trace (per req)	advice (per req)	Orochi's overhead	MediaWiki's workload
4.7%	MediaWiki's workload	7.1KB	1.7KB	11.4%

What are the (server's) CPU/network/storage costs?

CPU	Network			Storage
MediaWiki's workload	trace (per req)	advice (per req)	Orochi's overhead	MediaWiki's workload
4.7%	MediaWiki's workload	7.1KB	1.7KB	1.0x
phpBB's workload	phpBB's workload	5.7KB	0.3KB	11.4%
8.6%	HotCRP's workload	3.2KB	0.4KB	2.7%
HotCRP's workload	HotCRP's workload			10.9%
5.9%				1.5x

Orochi's verifier



Problem: naive re-execution doesn't save work

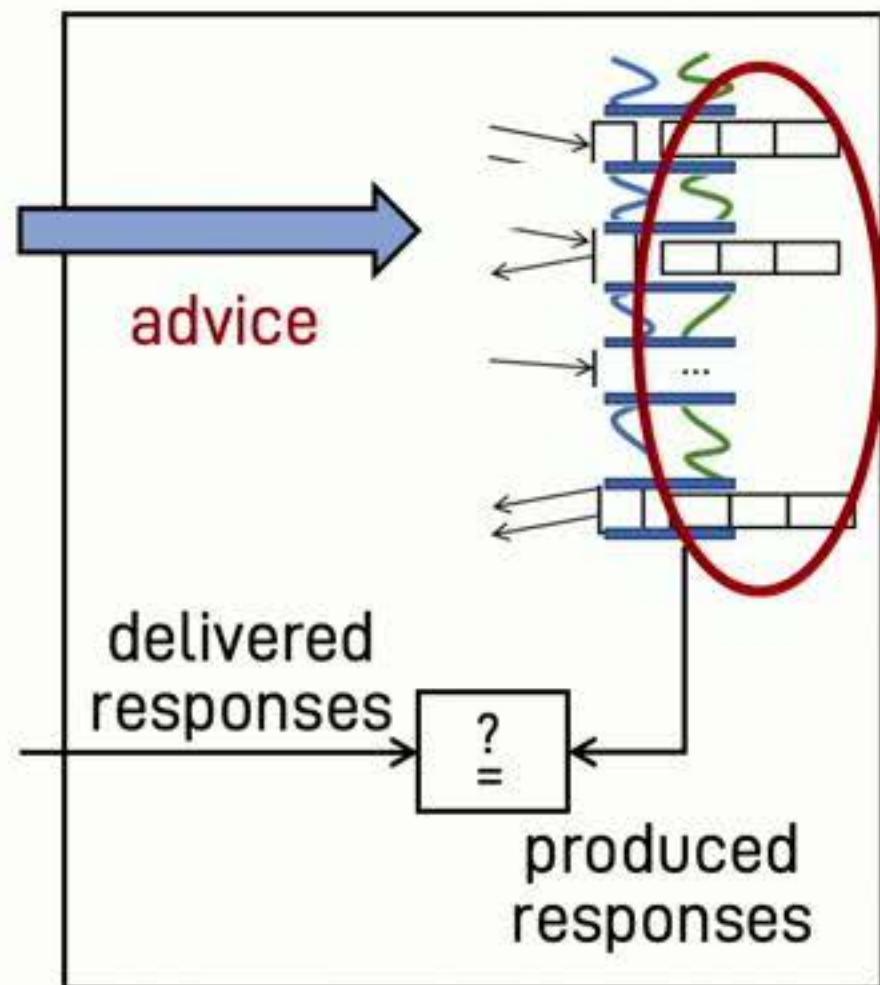
Solution: deduplicated re-execution



Problem: concurrent and untrusted server

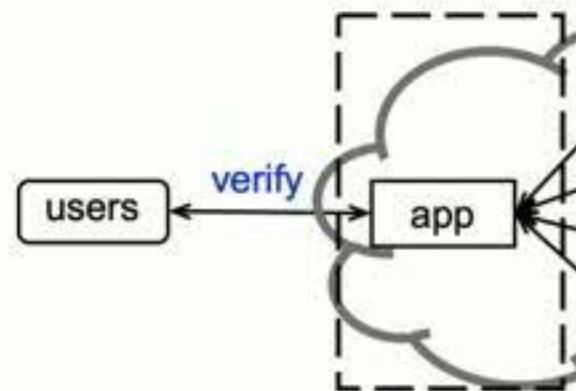
Solution: co-design verification protocol

Orochi's verifier

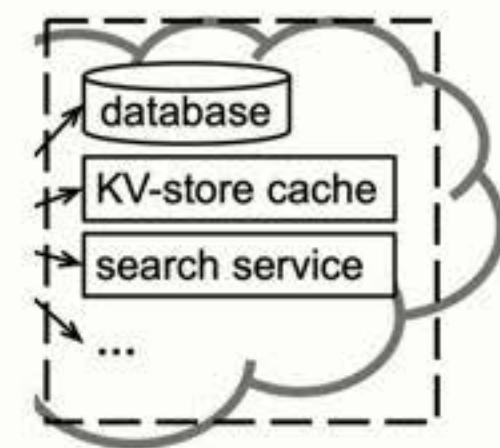


op logs may be **unavailable** (e.g., black-box databases)

1 Verify white-box services (Orochi [SOSP'17])



2 Verify black-box services (Cobra: verify Serializability)



3 Build composable verifiable framework (future work)

4 Other past work

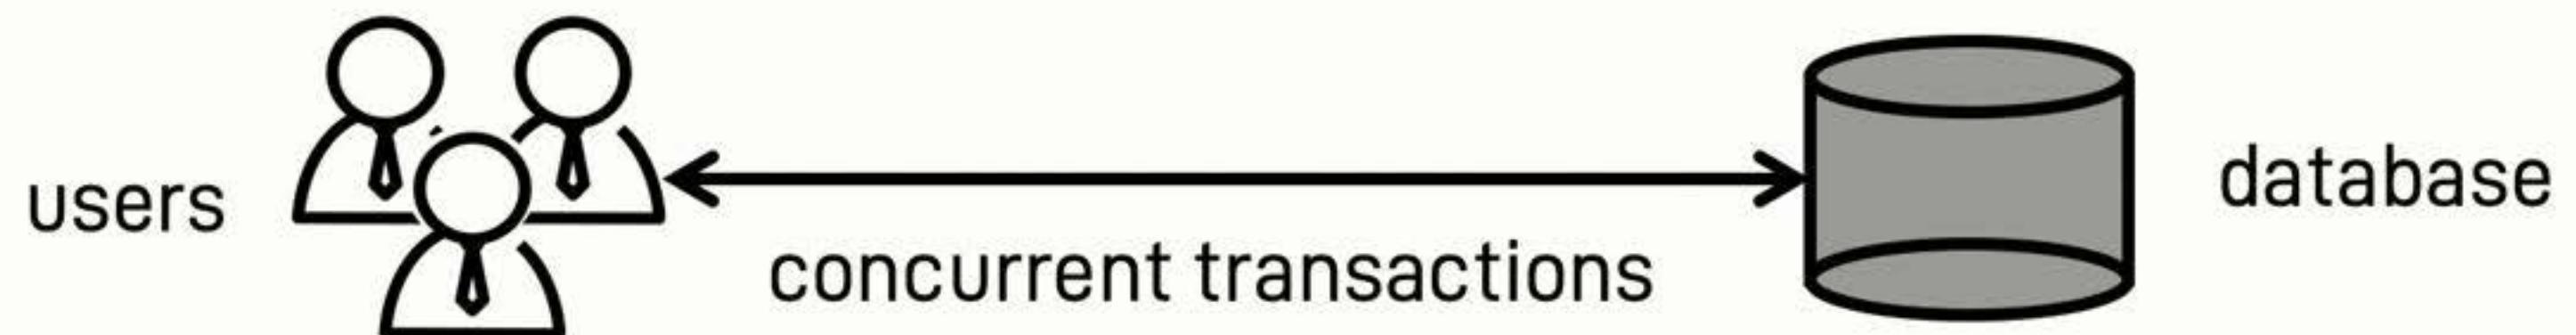
- Troubleshooting data center networks [NSDI'19]
- Protecting secret via security-oriented offloading [EuroSys'15]

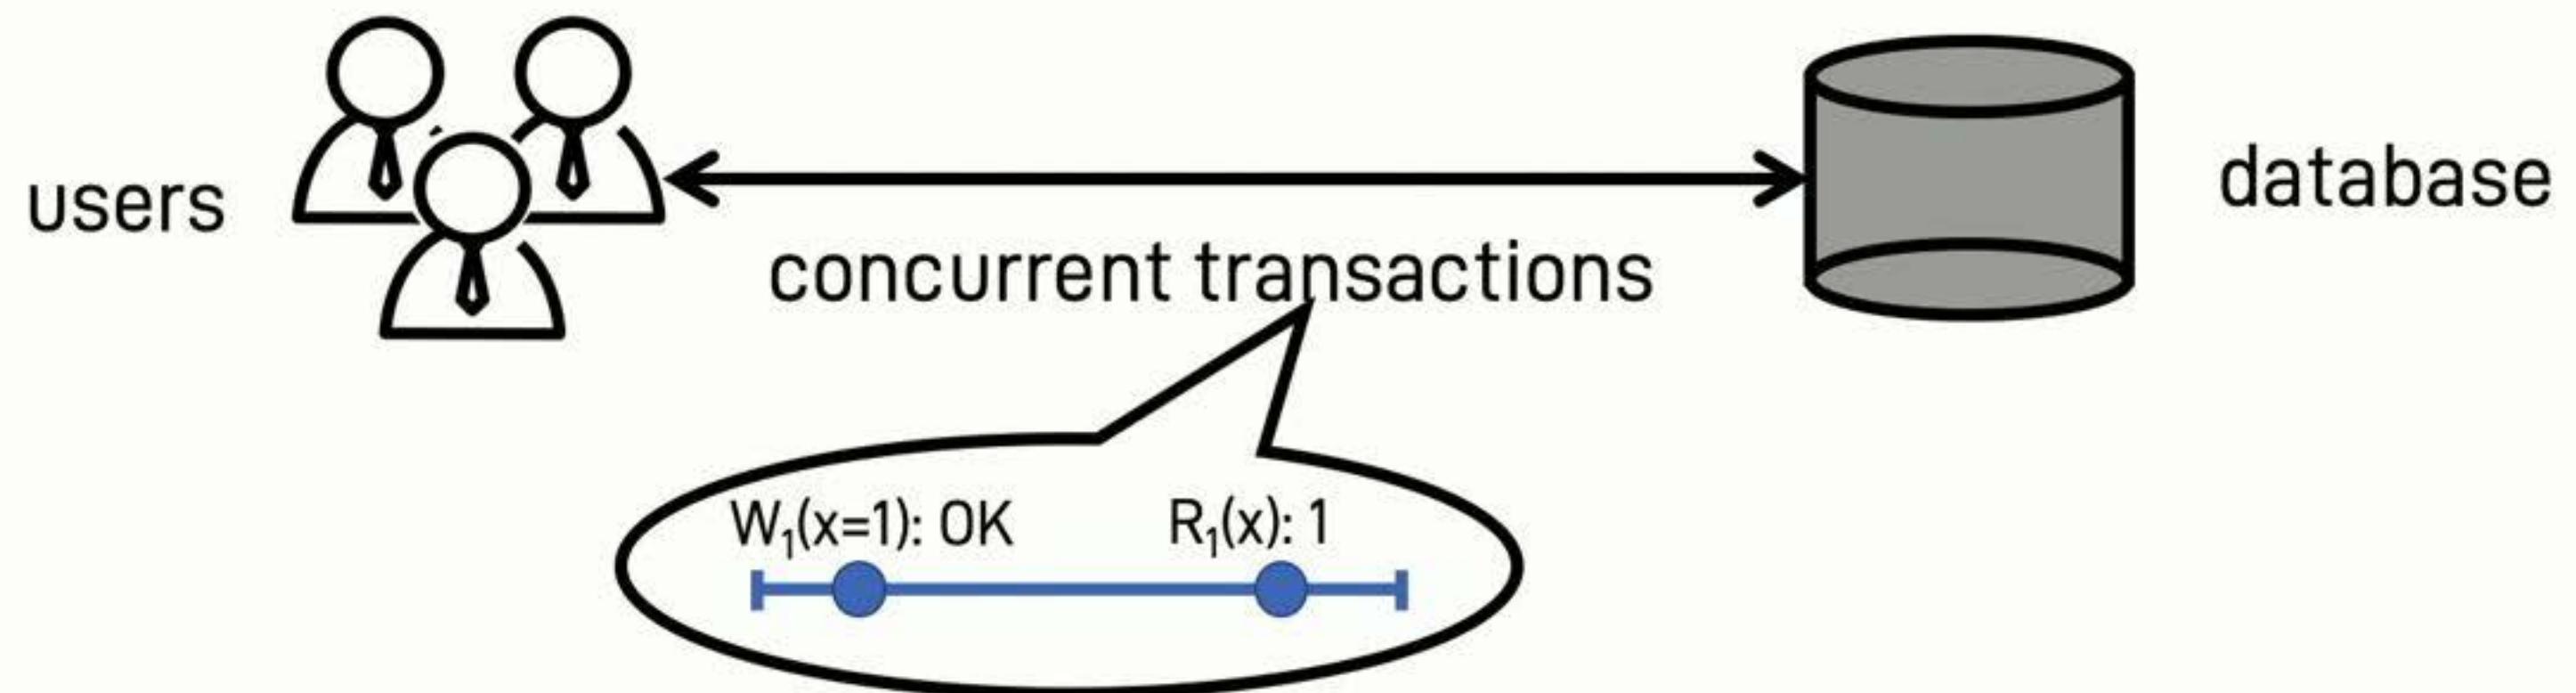
Why Serializability (SER)?

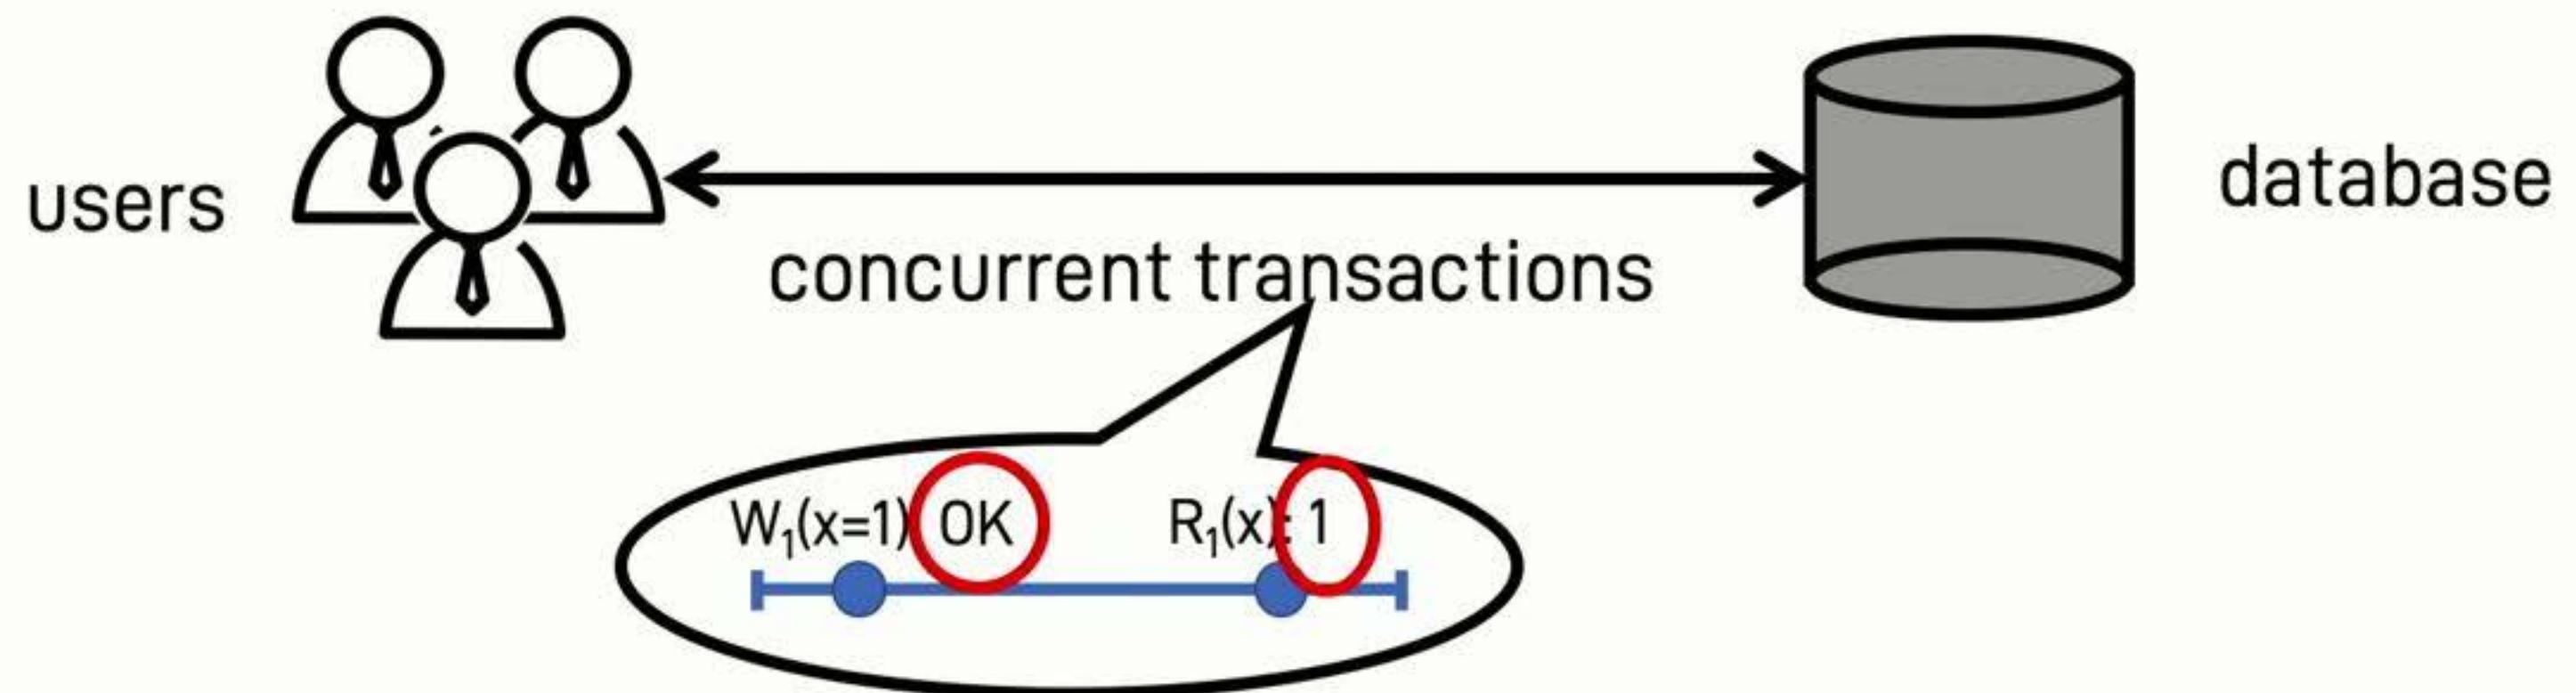
gold standard
isolation level

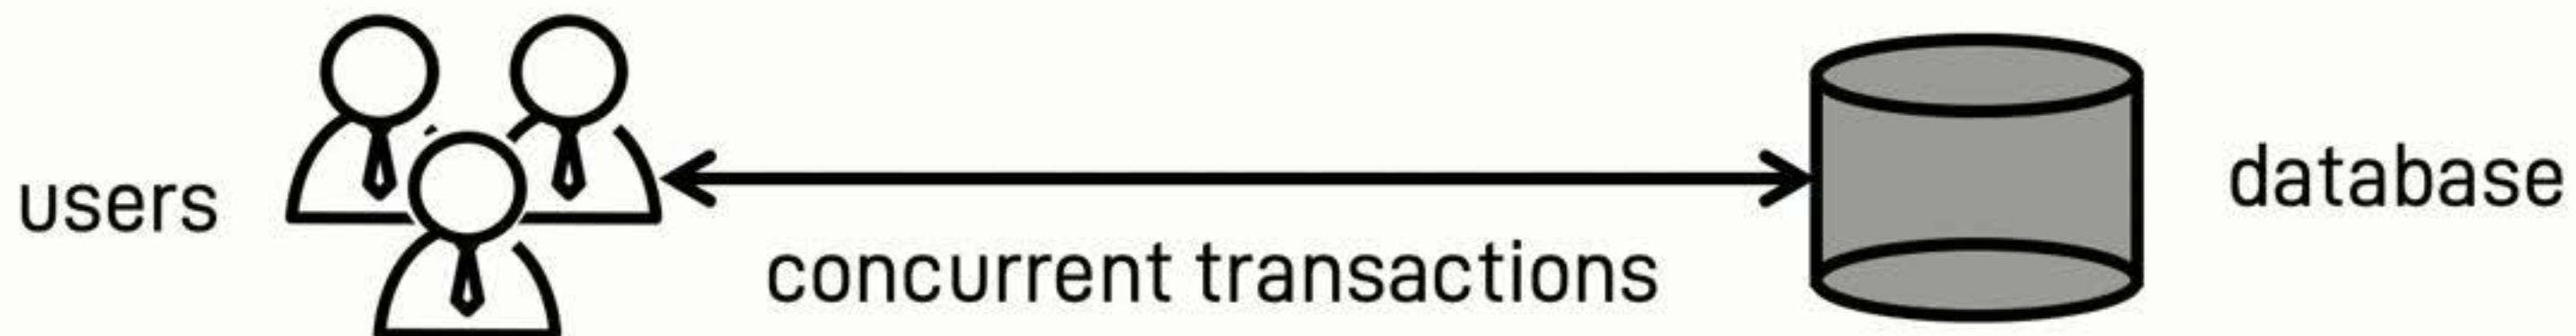


fundamental
and
challenging









verifier Are concurrent transactions SER?

concurrent transactions $\stackrel{?}{=}$ sequential execution
of these transactions

concurrent transactions

$\stackrel{?}{=}$

sequential execution
of these transactions

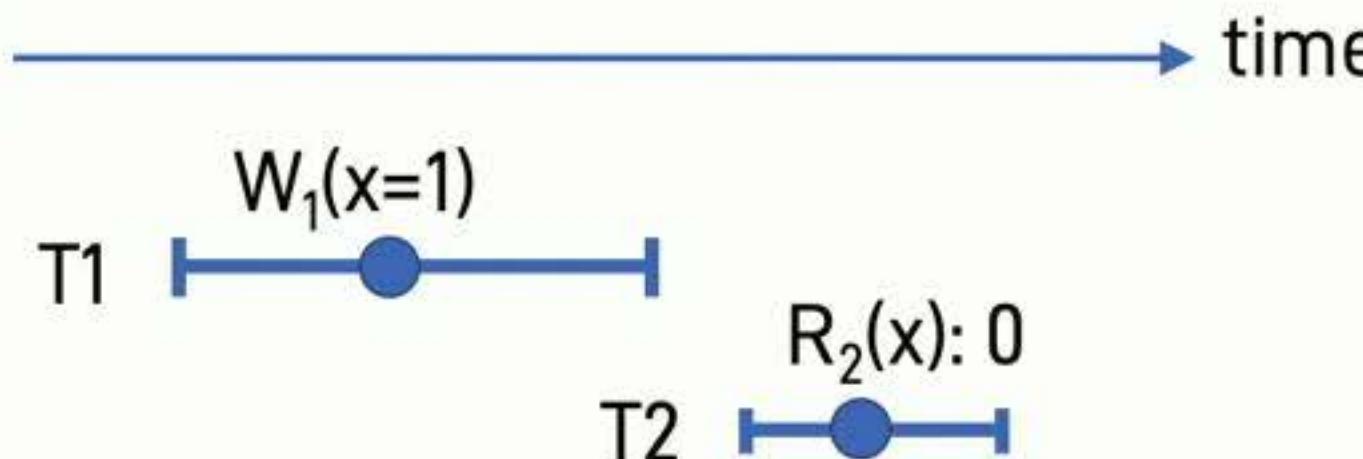
- Checking **view-SER** is an NP-complete problem [Papadimitriou 79]

concurrent transactions

=
?

sequential execution
of these transactions

- Checking **view**-SER is an NP-complete problem [Papadimitriou 79]
- Challenge: SER doesn't respect real-time order

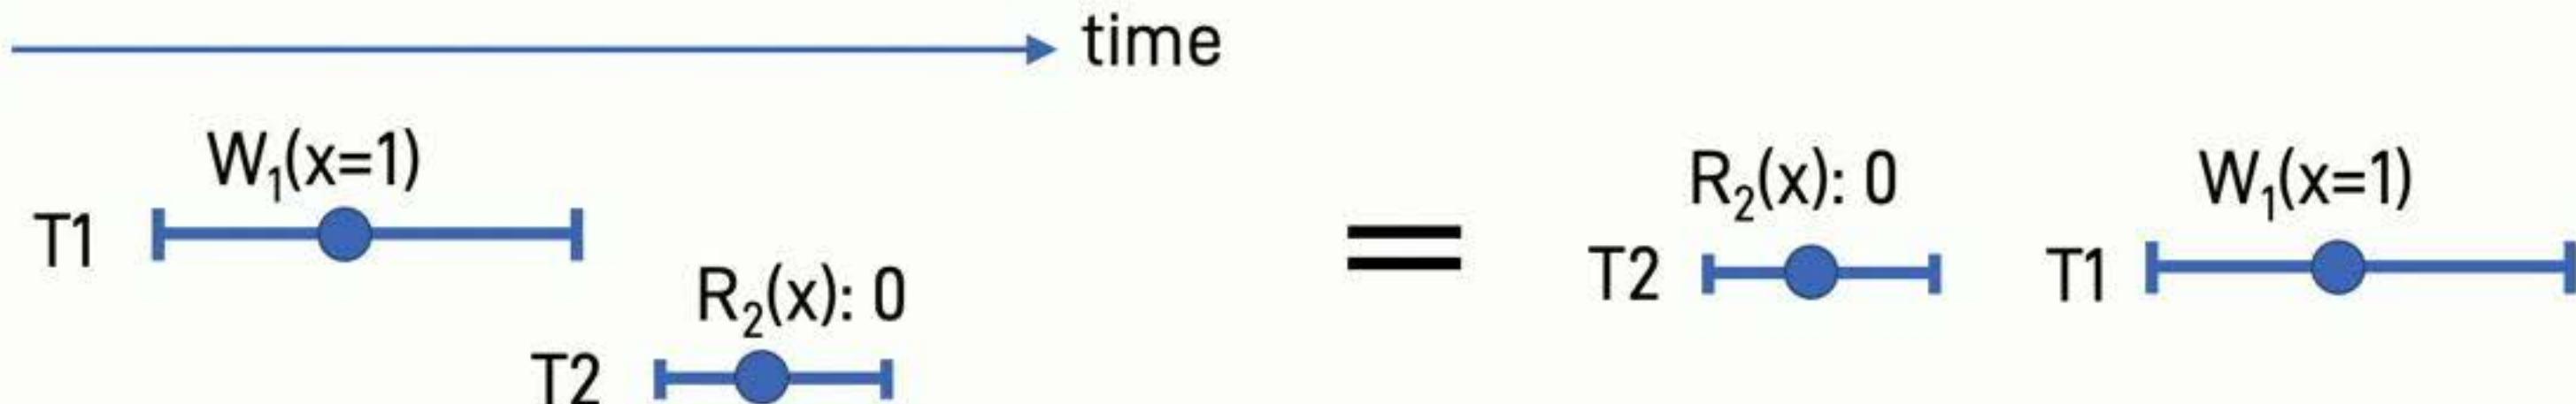


concurrent transactions

= ?

sequential execution
of these transactions

- Checking **view**-SER is an NP-complete problem [Papadimitriou 79]
- Challenge: SER doesn't respect real-time order



Cobra aims at real-world workloads

- Intuition: advances of SAT/SMT solvers

Cobra aims at real-world workloads

- Intuition: advances of SAT/SMT solvers
- Starting point (Papadimitriou's construction):

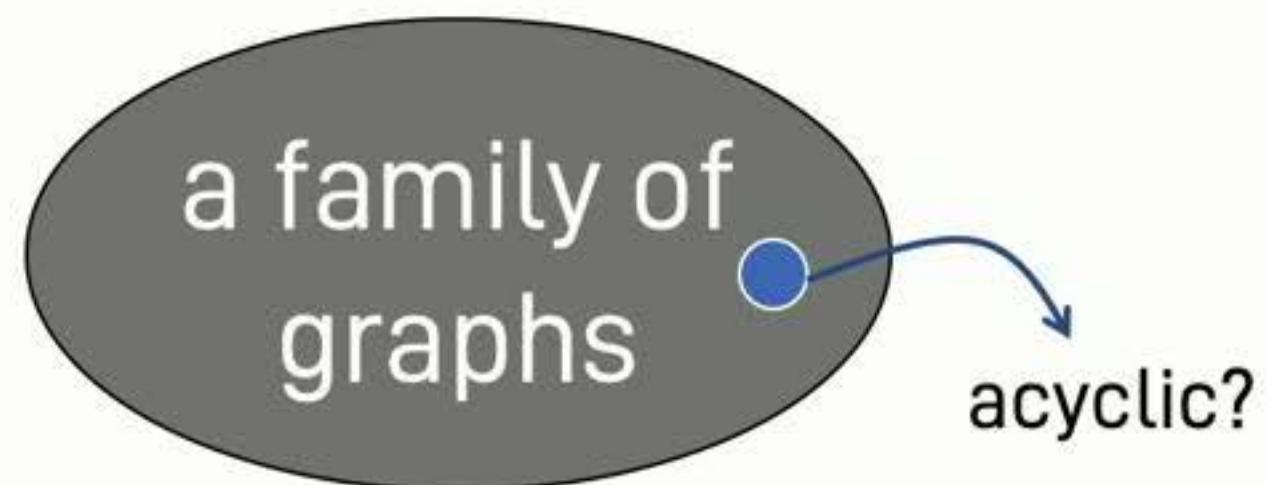
Cobra aims at real-world workloads

- Intuition: advances of SAT/SMT solvers
- Starting point (Papadimitriou's construction):

concurrent transactions

$\stackrel{?}{=}$

sequential execution
of these transactions



Cobra aims at real-world workloads

- Intuition: advances of SAT/SMT solvers
- Starting point (Papadimitriou's construction):

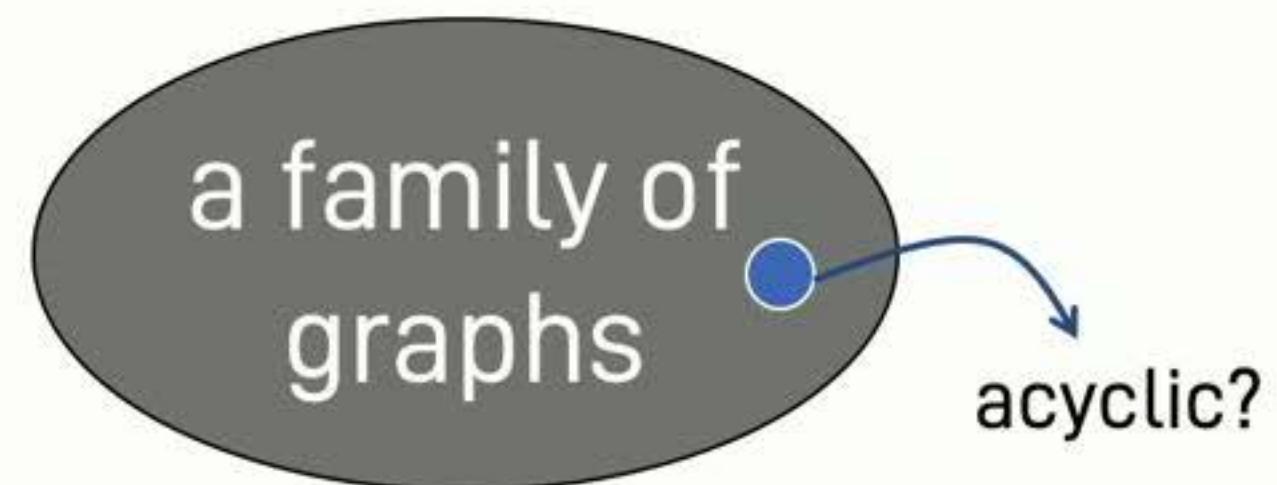
concurrent transactions

$\stackrel{?}{=}$

sequential execution
of these transactions



equivalent



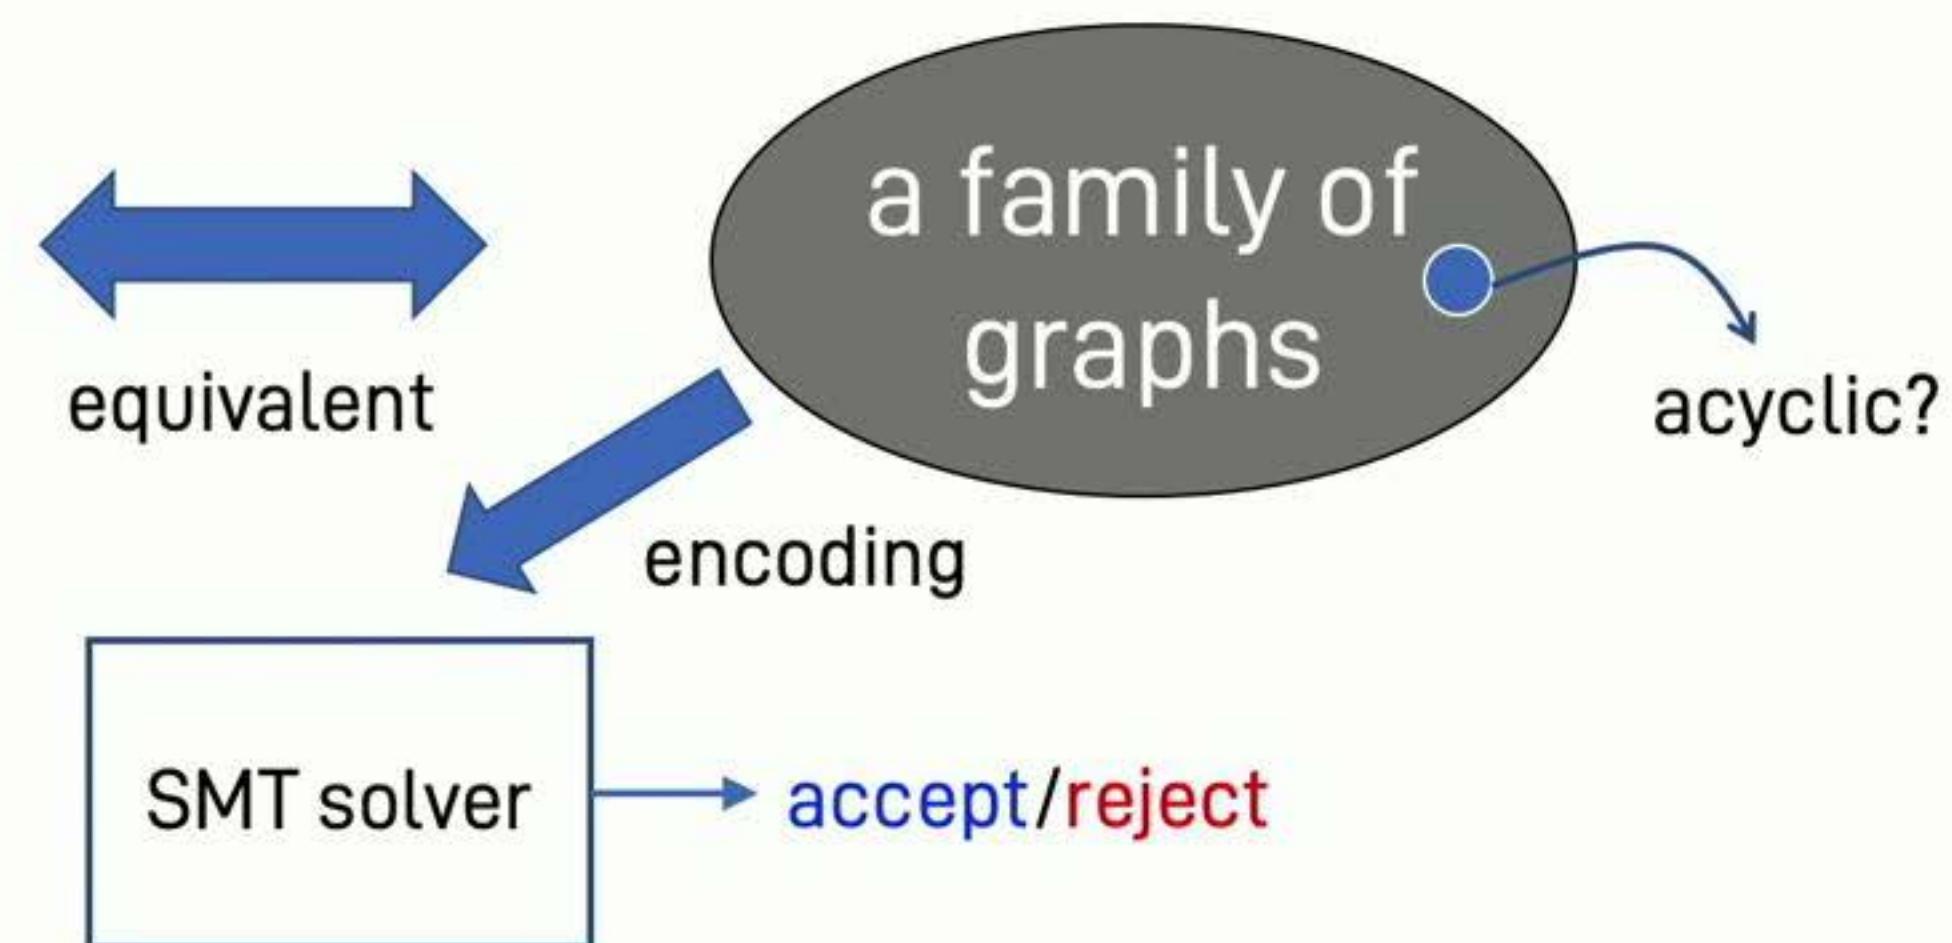
Cobra aims at real-world workloads

- Intuition: advances of SAT/SMT solvers
- Starting point (Papadimitriou's construction):

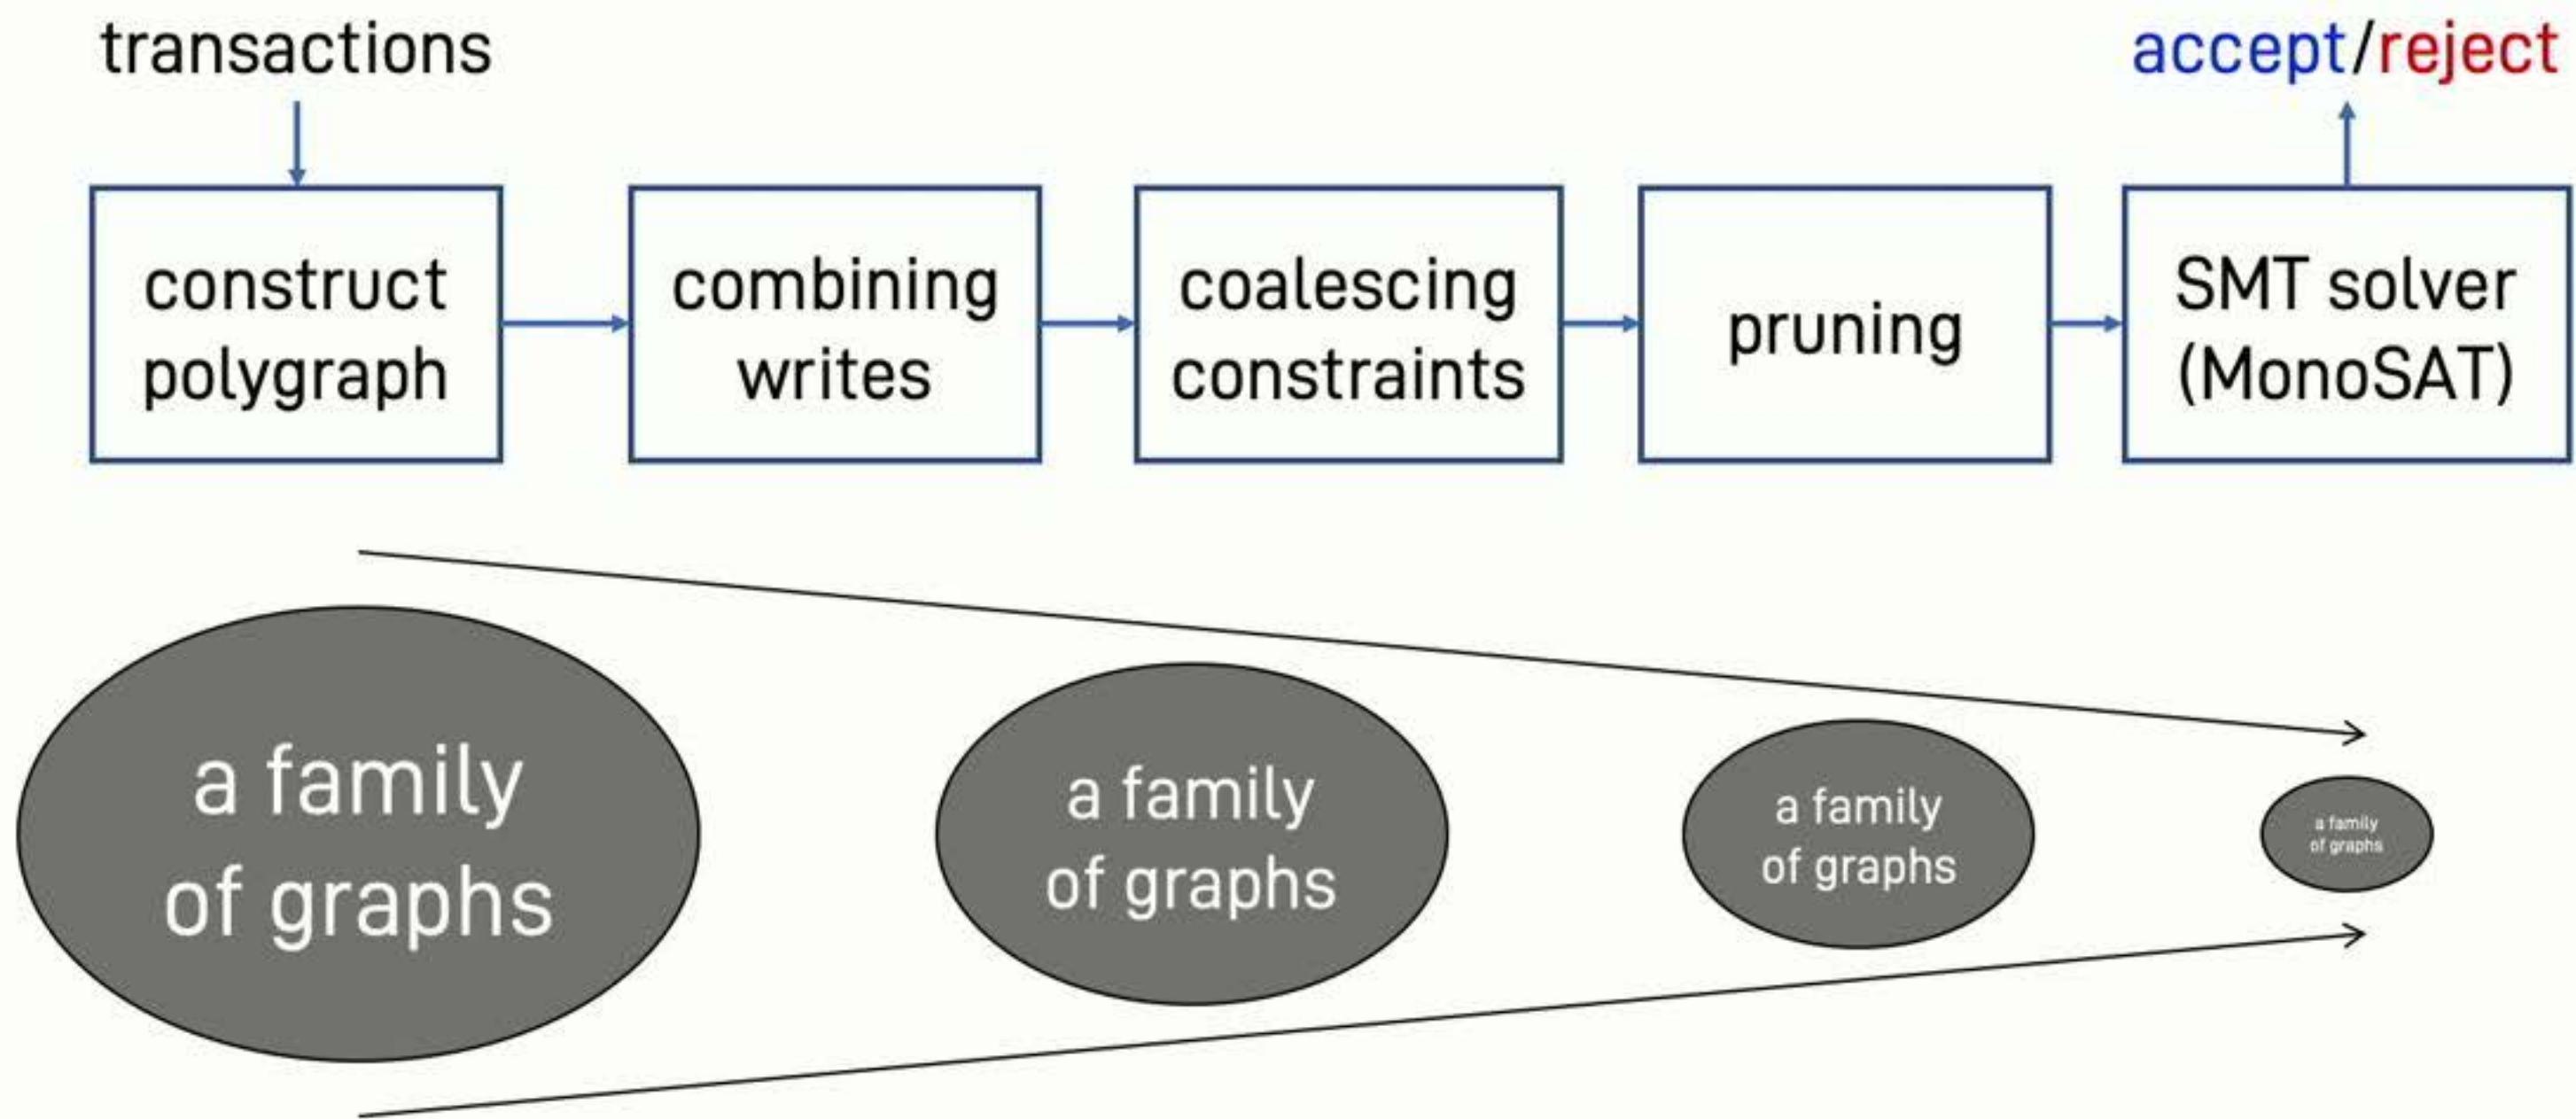
concurrent transactions

$\stackrel{?}{=}$

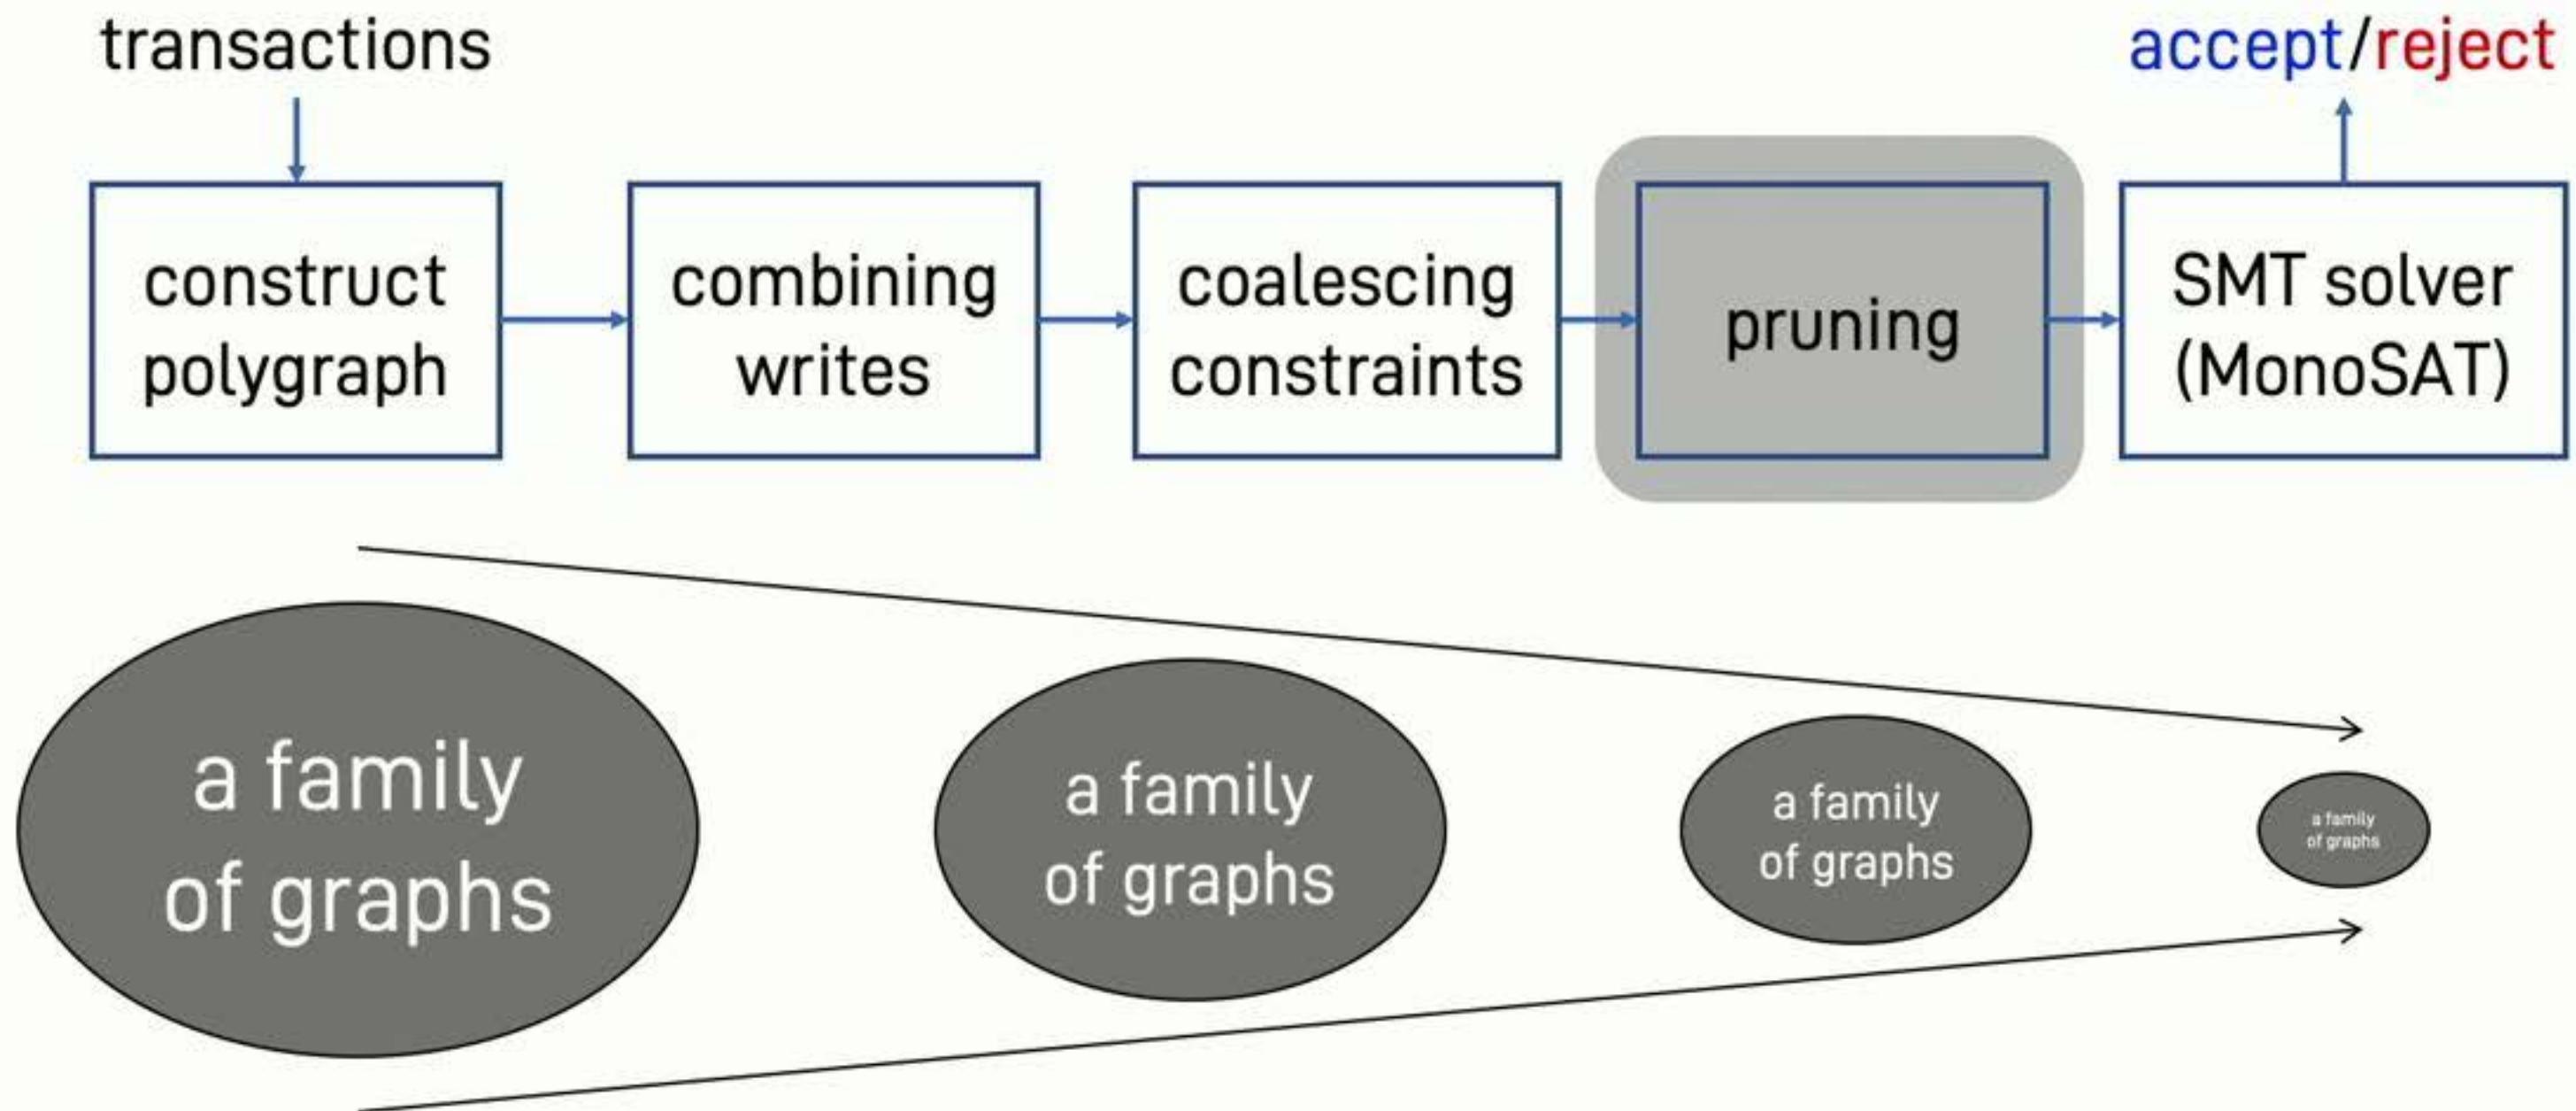
sequential execution
of these transactions



Cobra: narrowing the search space



Cobra: narrowing the search space



Pruning via graph paths (reachability)

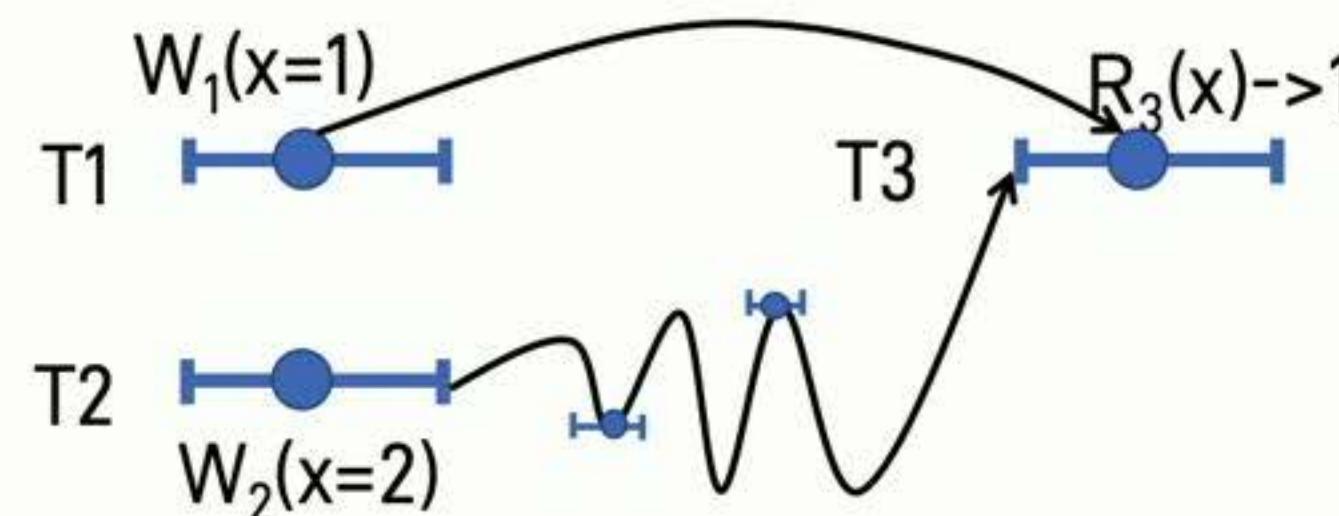
- idea: reduce #graphs by adding constraints to the family

Pruning via graph paths (reachability)

- idea: reduce #graphs by adding constraints to the family
 - 1) what constraints can be inferred from reachability?
 - 2) how to get reachability efficiently?

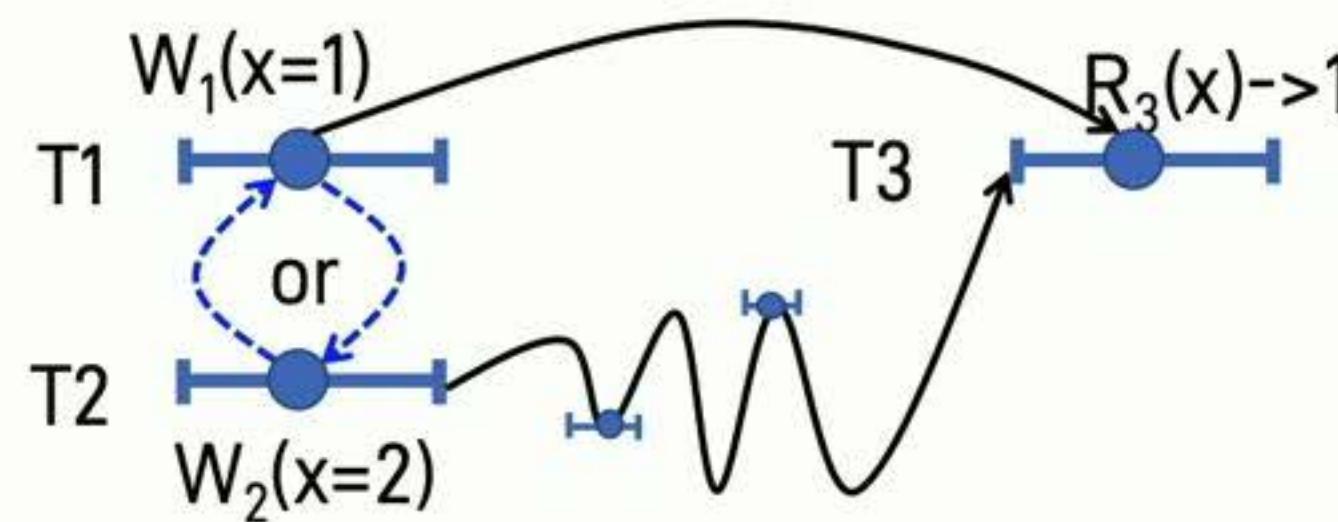
Pruning via graph paths (reachability)

- idea: reduce #graphs by adding constraints to the family
 - 1) what constraints can be inferred from reachability?
 - 2) how to get reachability efficiently?



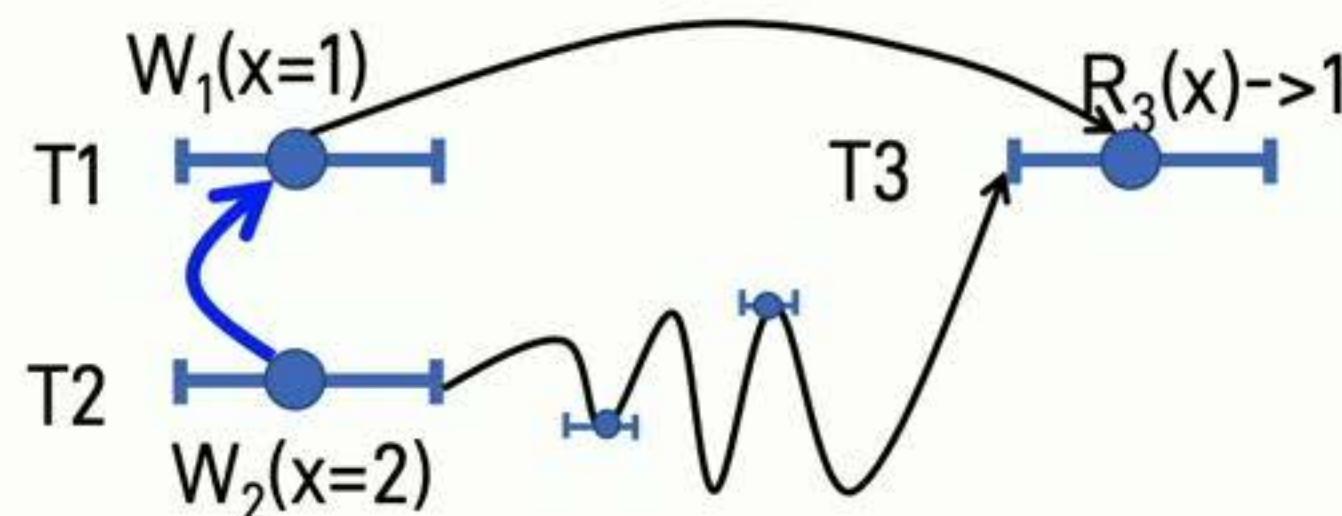
Pruning via graph paths (reachability)

- idea: reduce #graphs by adding constraints to the family
 - 1) what constraints can be inferred from reachability?
 - 2) how to get reachability efficiently?



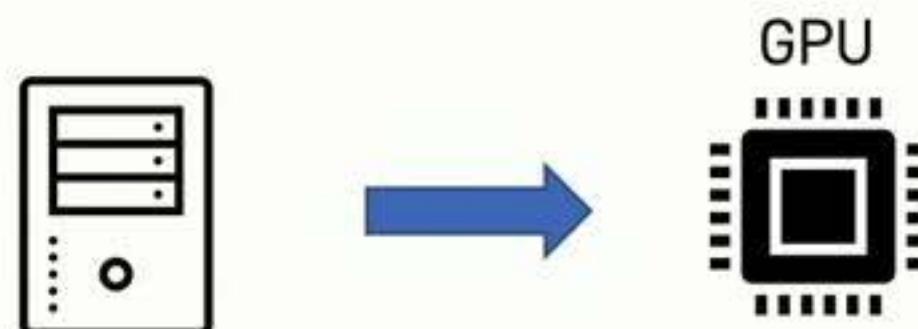
Pruning via graph paths (reachability)

- idea: reduce #graphs by adding constraints to the family
 - 1) what constraints can be inferred from reachability?
 - 2) how to get reachability efficiently?



Pruning via graph paths (reachability)

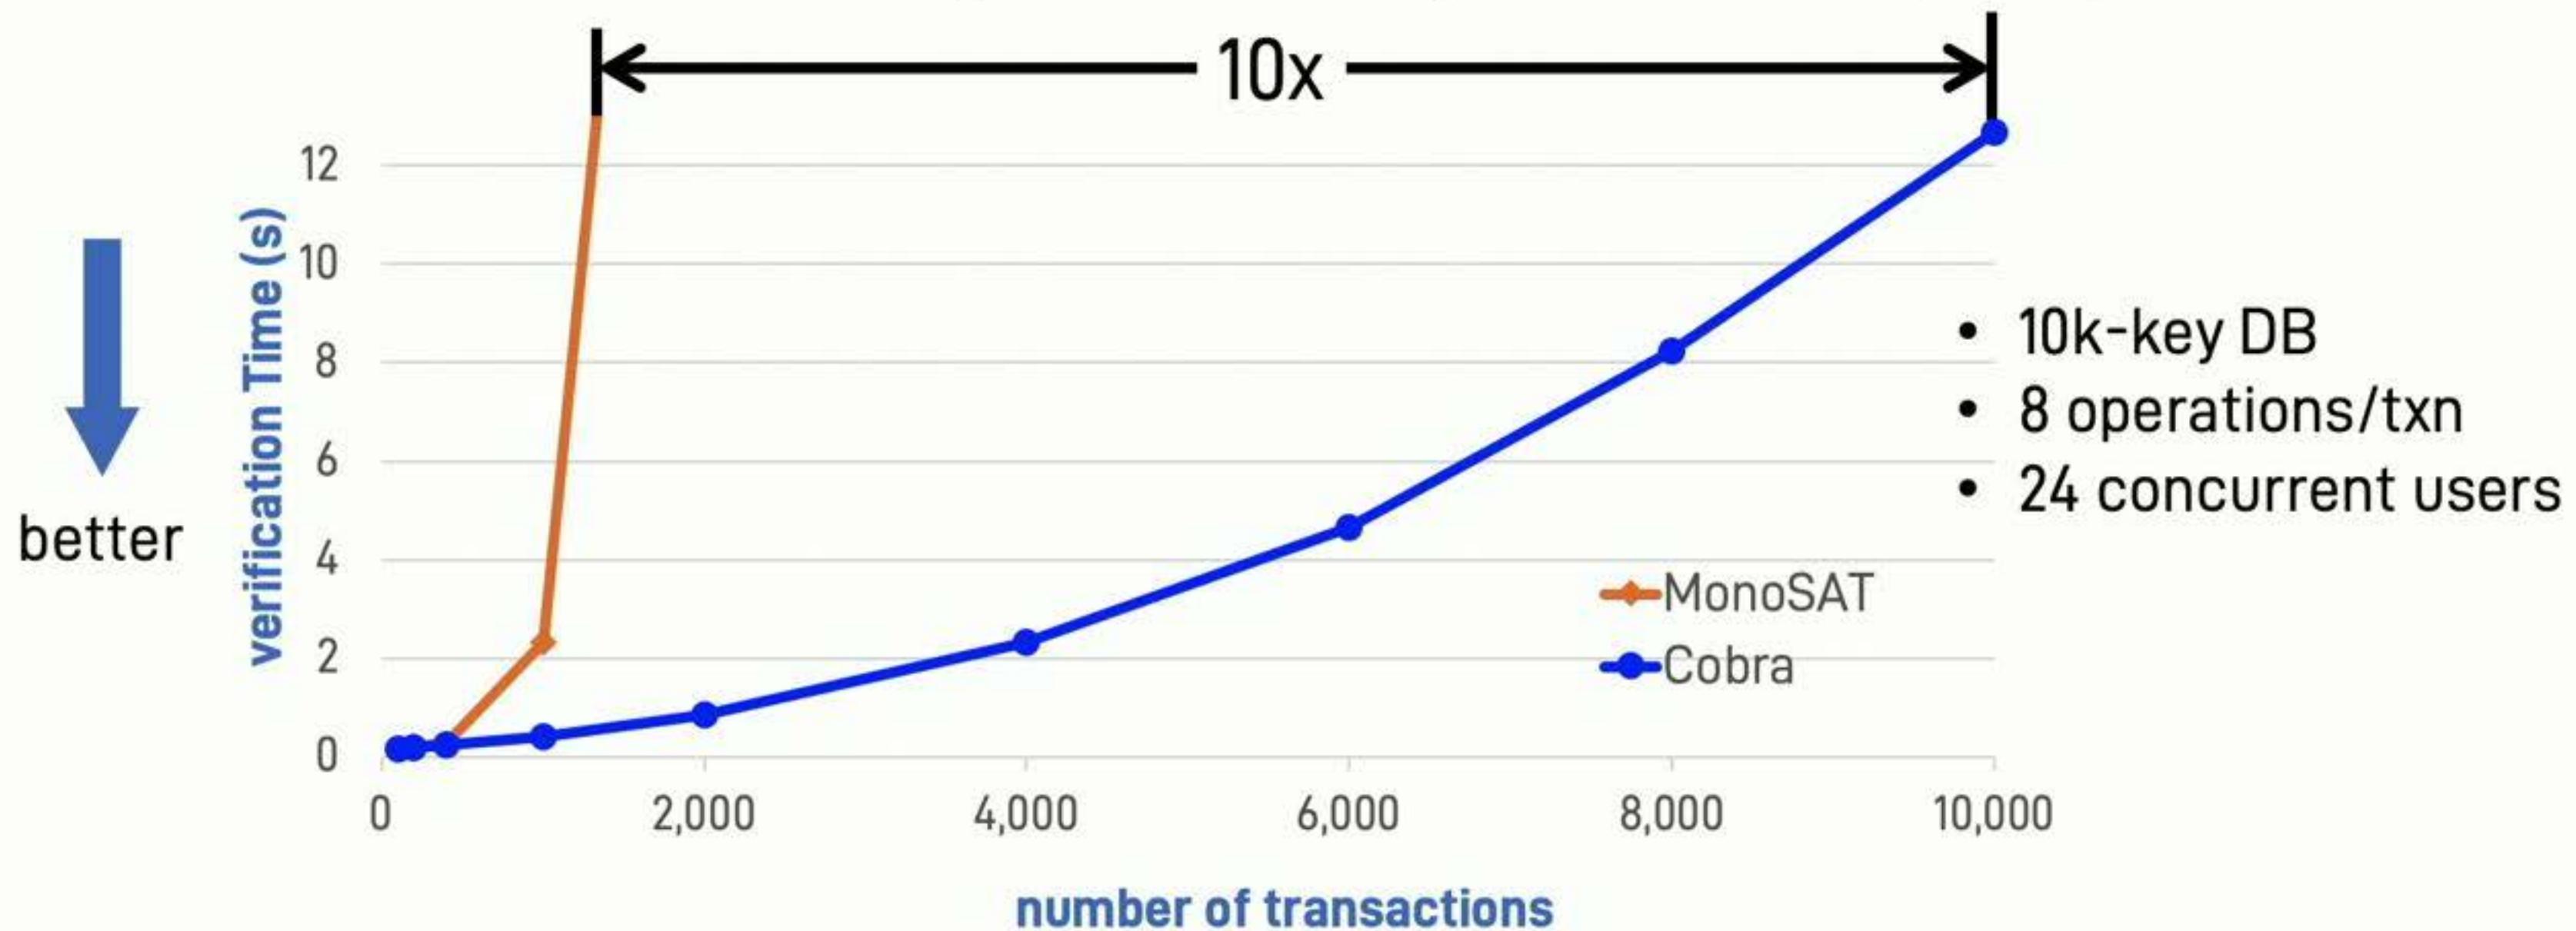
- idea: reduce #graphs by adding constraints to the family
 - 1) what constraints can be inferred from reachability?
 - 2) how to get reachability efficiently?



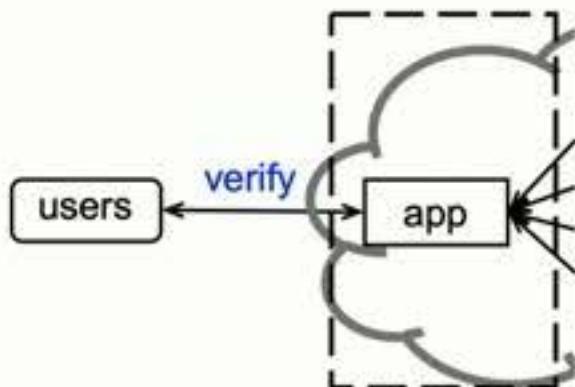
calculating reachability
using **Matrix Multiplication**

Cobra can handle 10x larger workloads

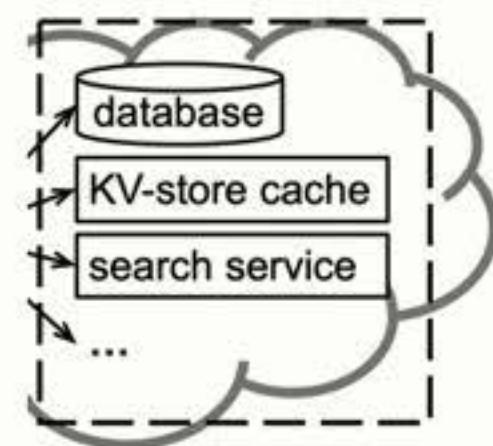
RW benchmark: read-only and write-only transactions (50:50)



1 Verify white-box services (Orochi [SOSP'17])



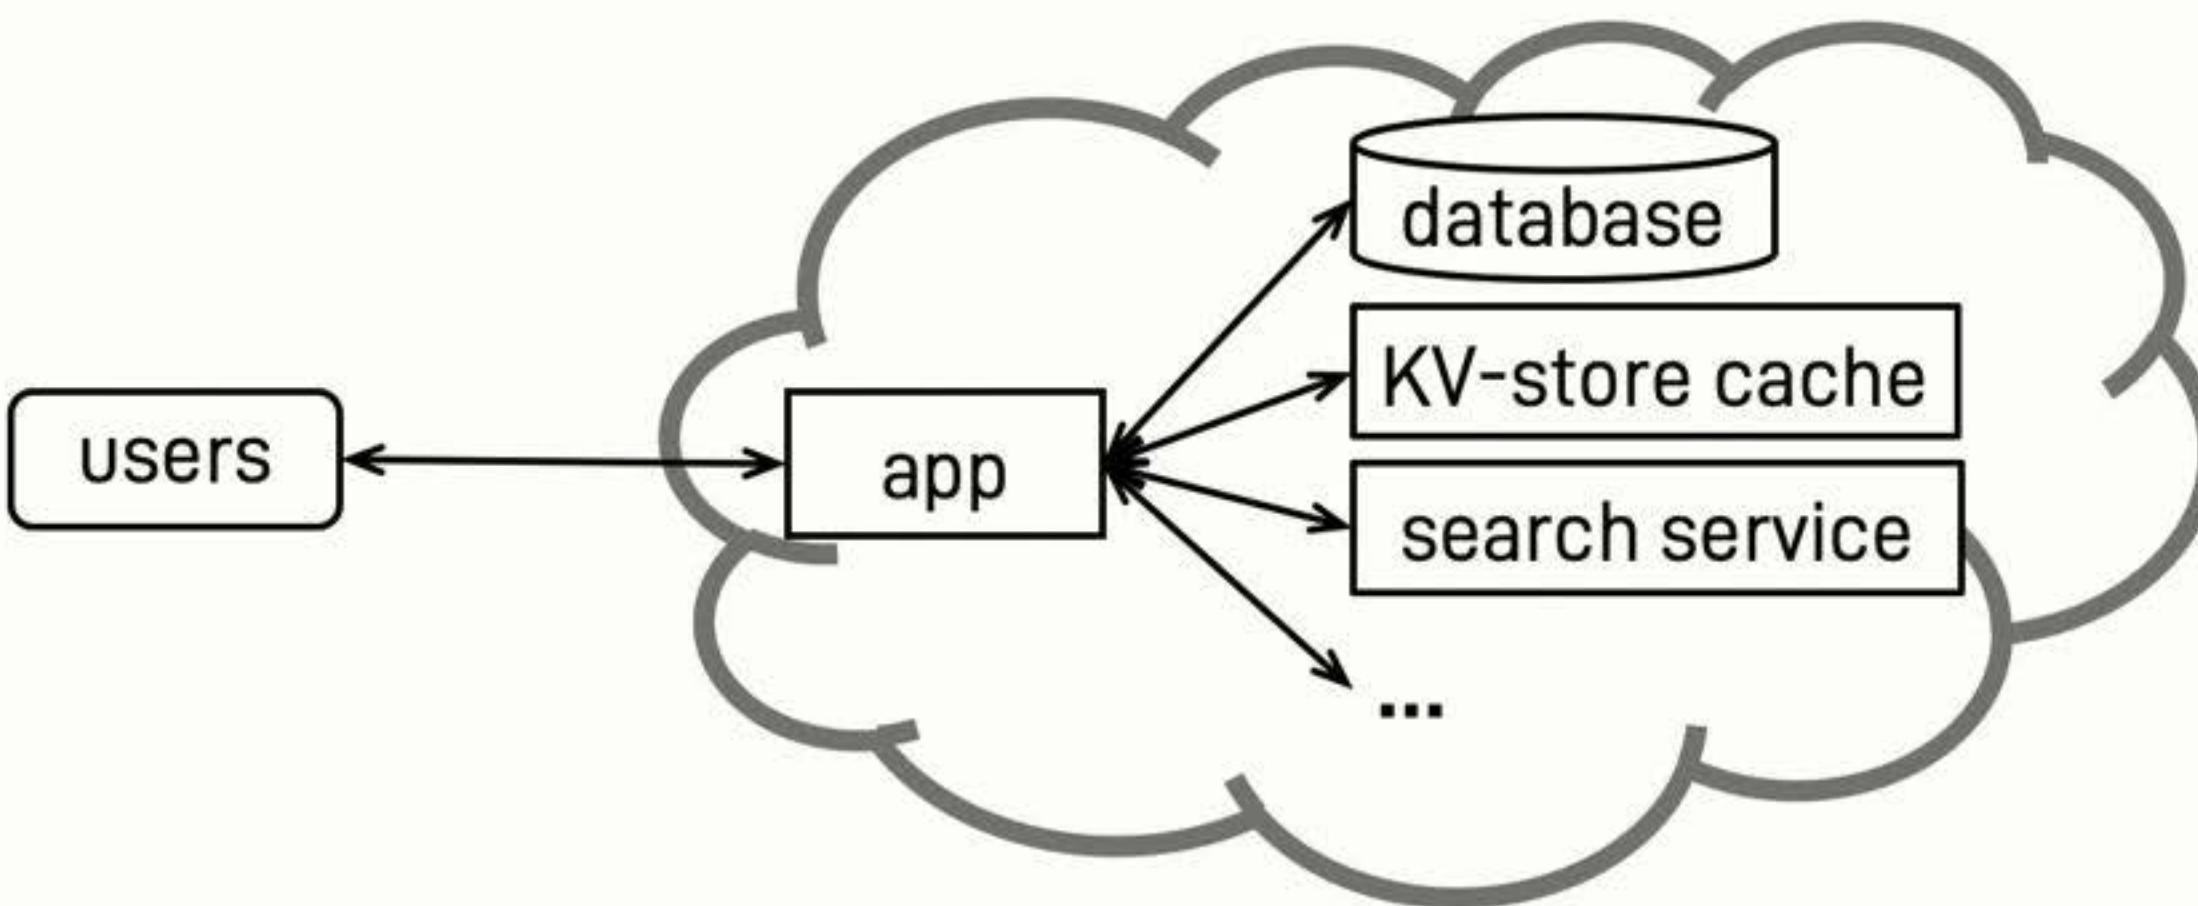
2 Verify black-box services (Cobra: verify Serializability)

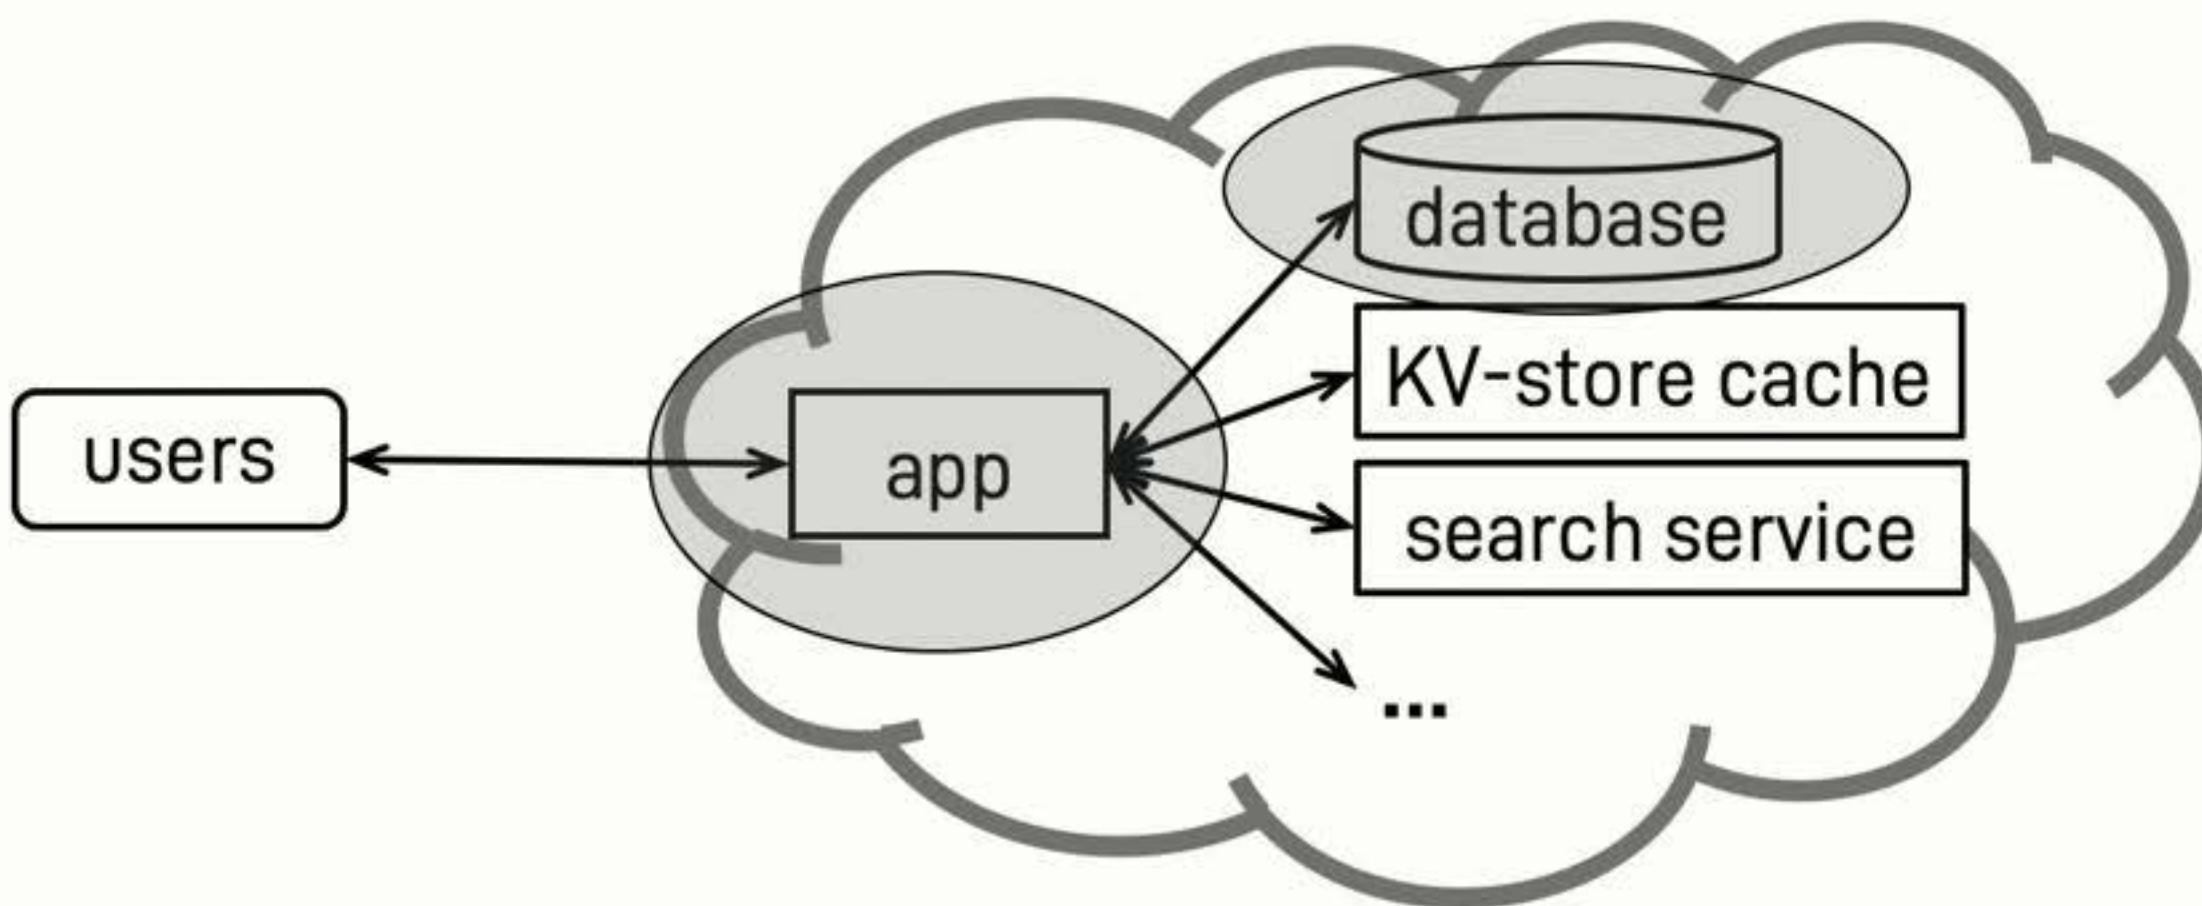


3 Build composable verifiable framework (future work)

4 Other past work

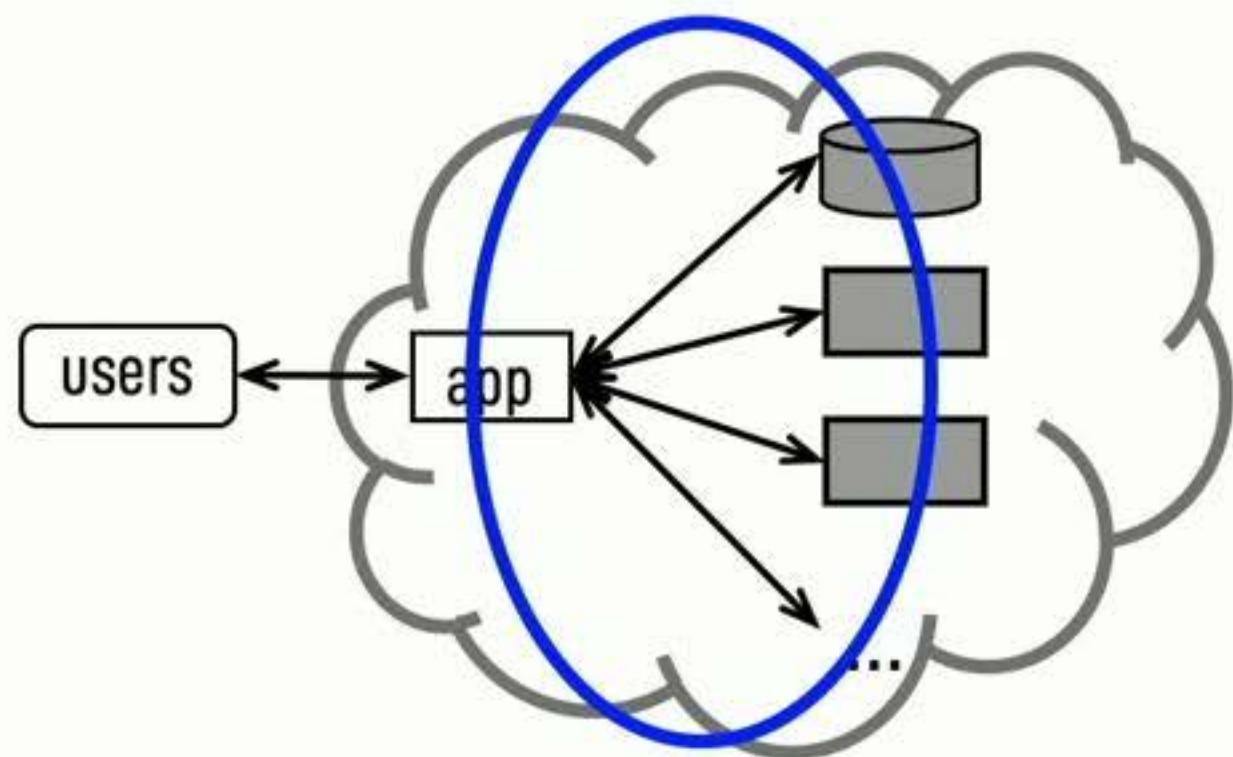
- Troubleshooting data center networks [NSDI'19]
- Protecting secret via security-oriented offloading [EuroSys'15]





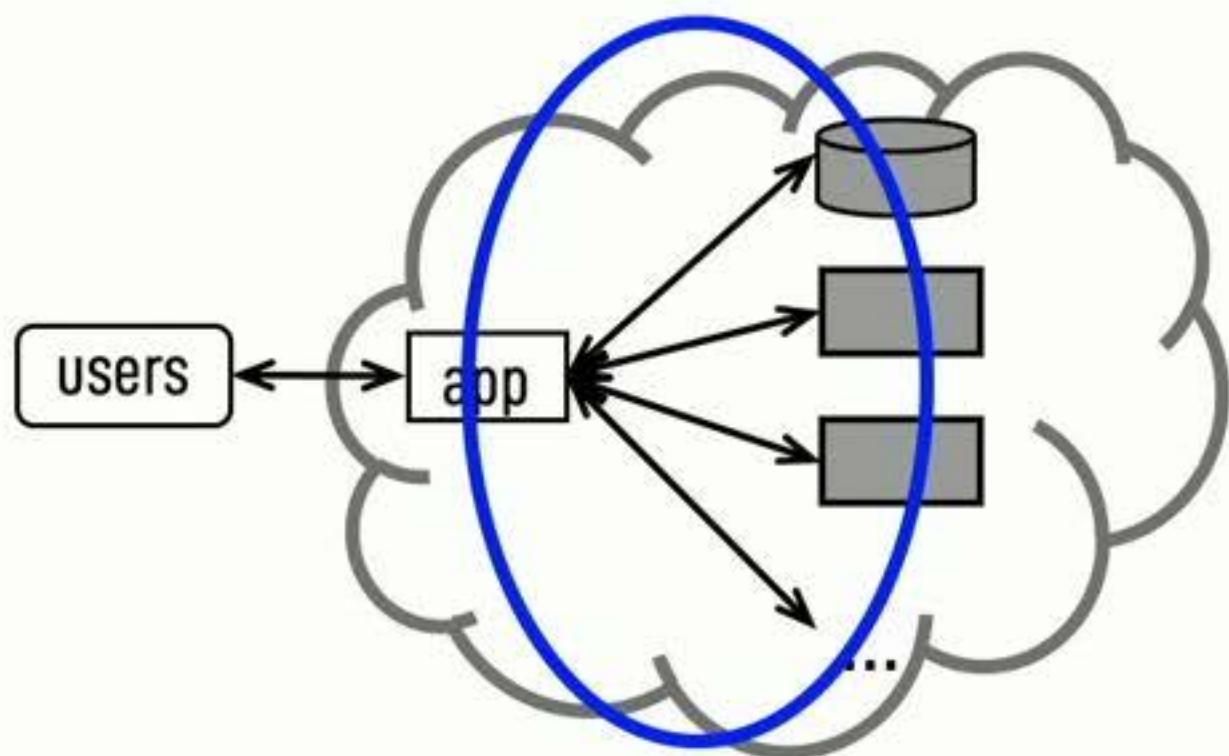
Question #1:

How to compose verifications of multiple services?

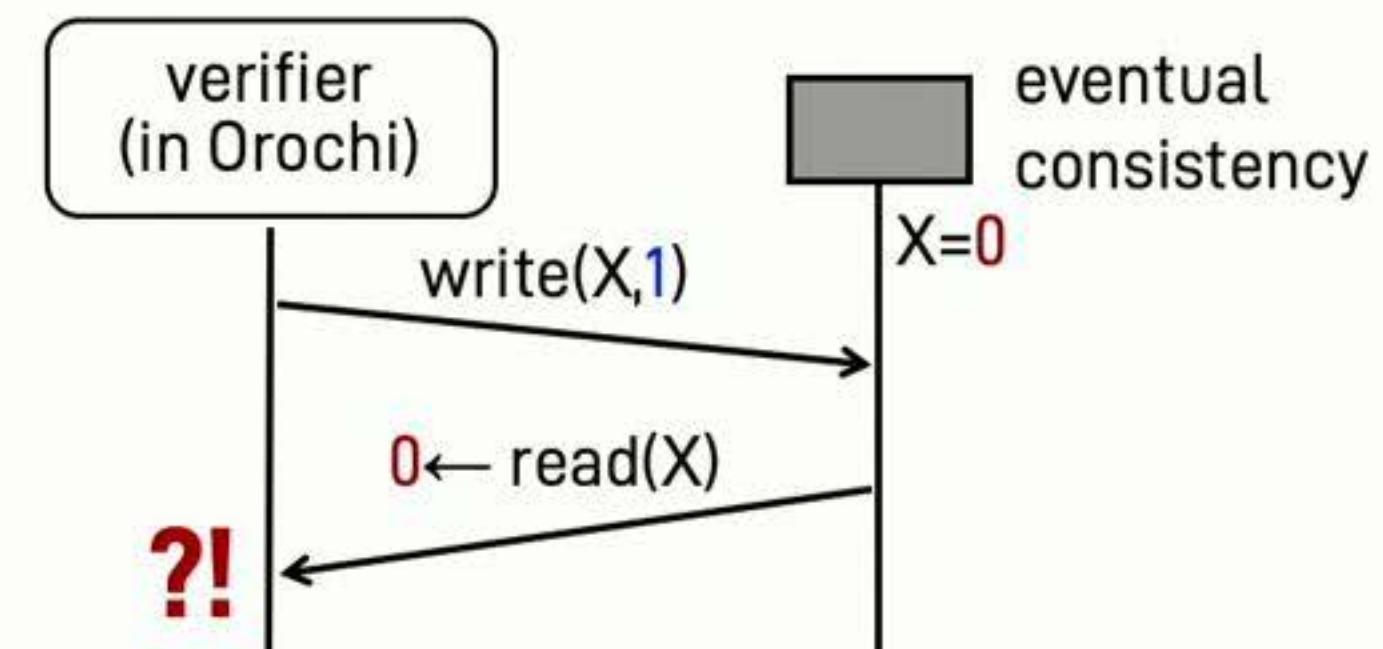


Question #1:

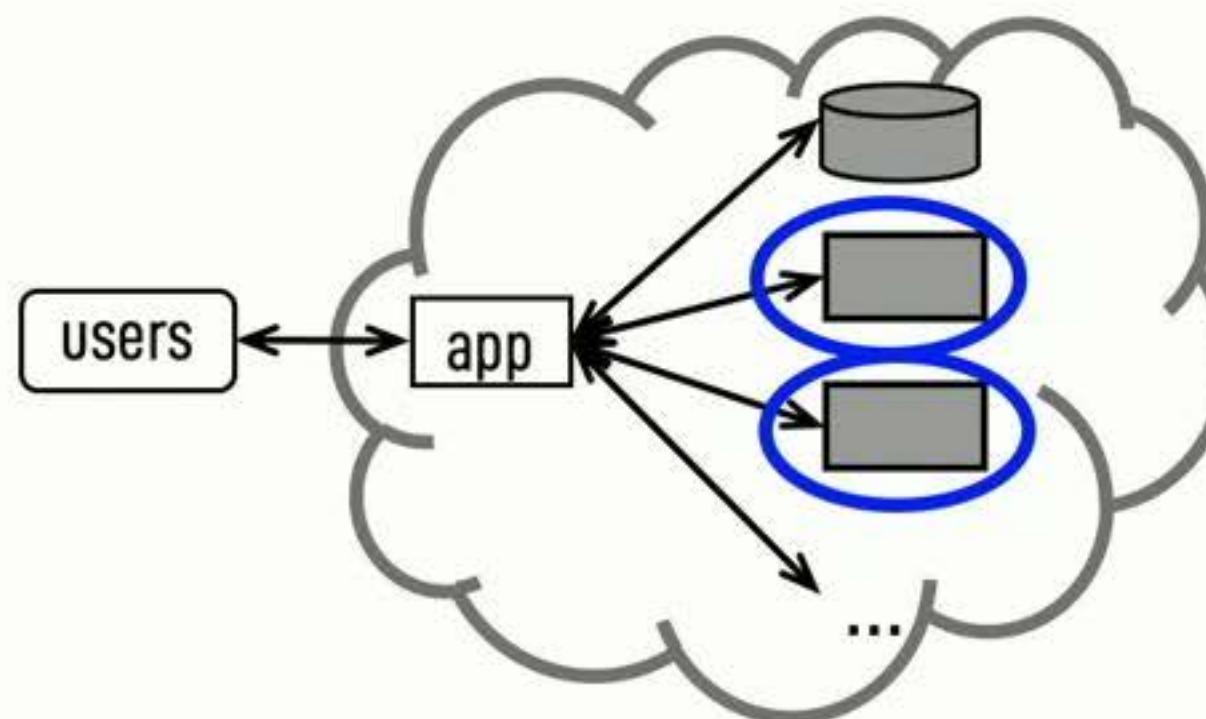
How to compose verifications of multiple services?



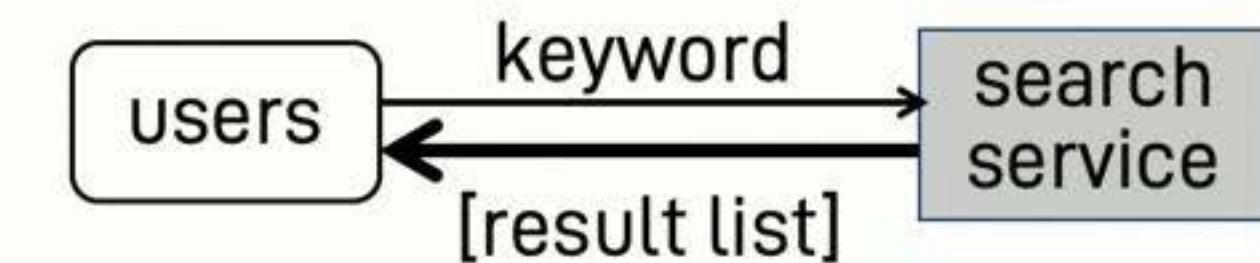
- **Challenge:** some services do not respect real-time order



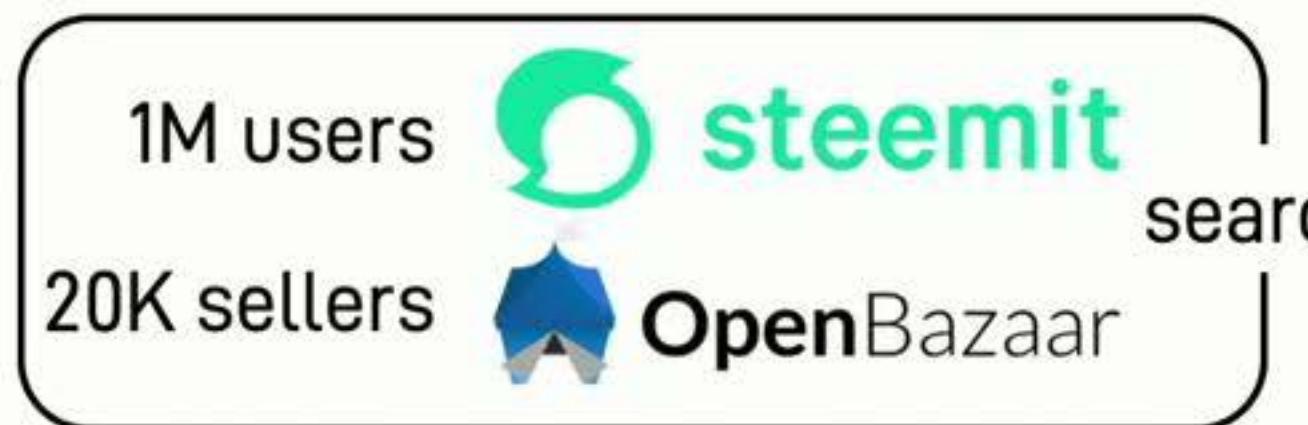
Question #2: How to verify various black-box services?



for example, a search service



Verifiable search engine

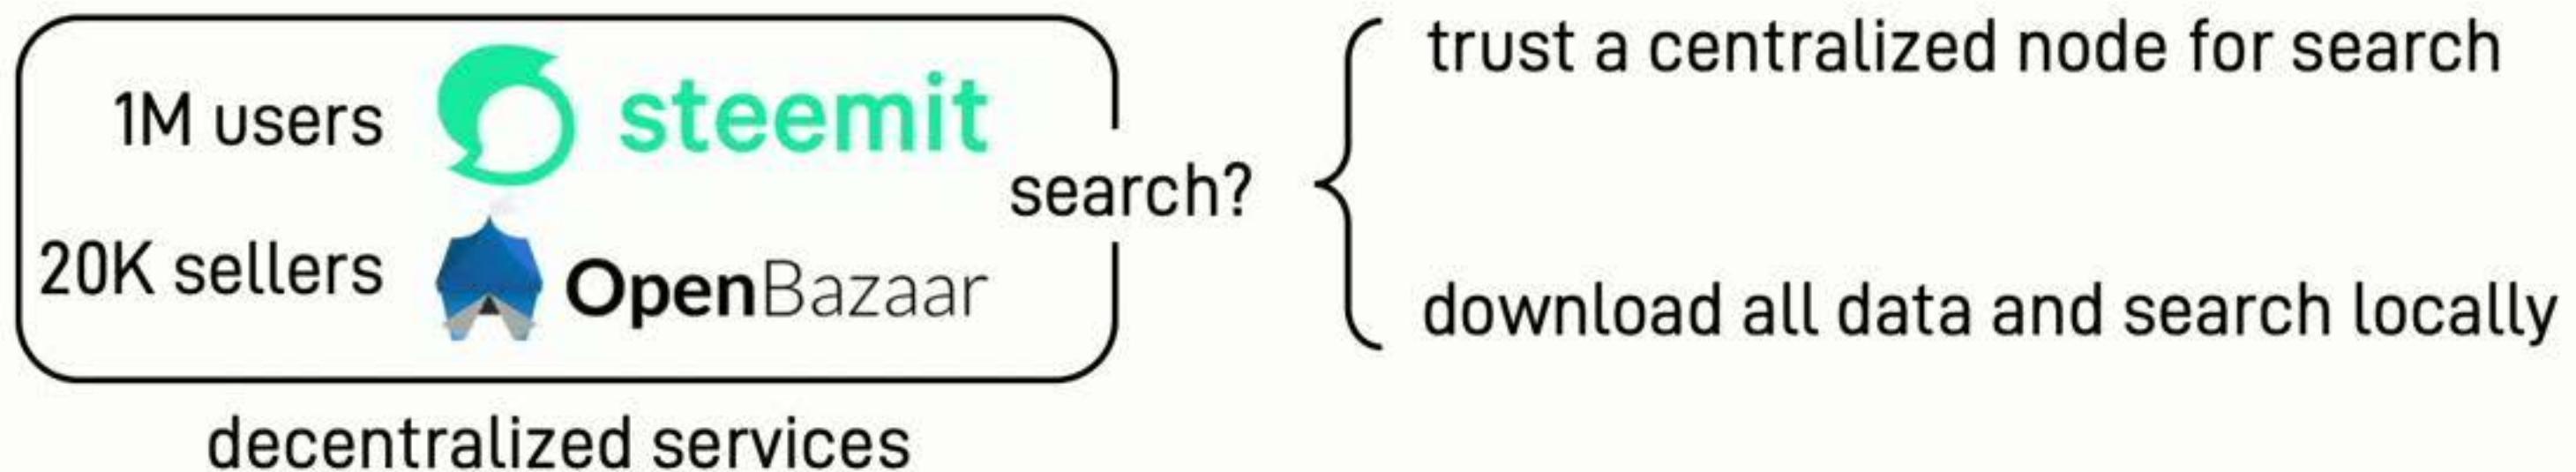


decentralized services

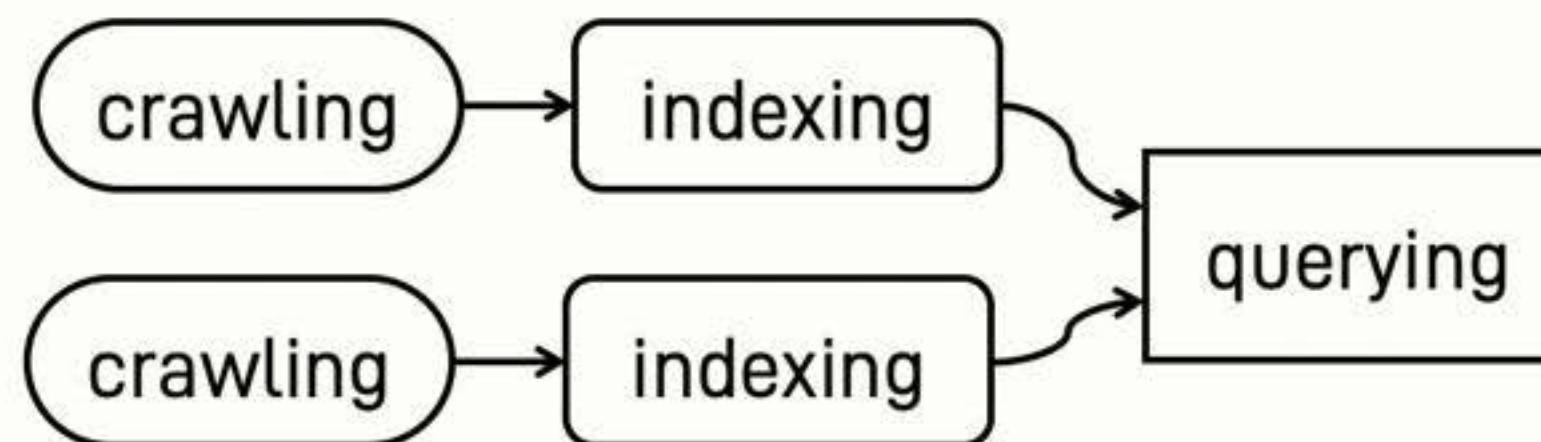
trust a centralized node for search

download all data and search locally

Verifiable search engine

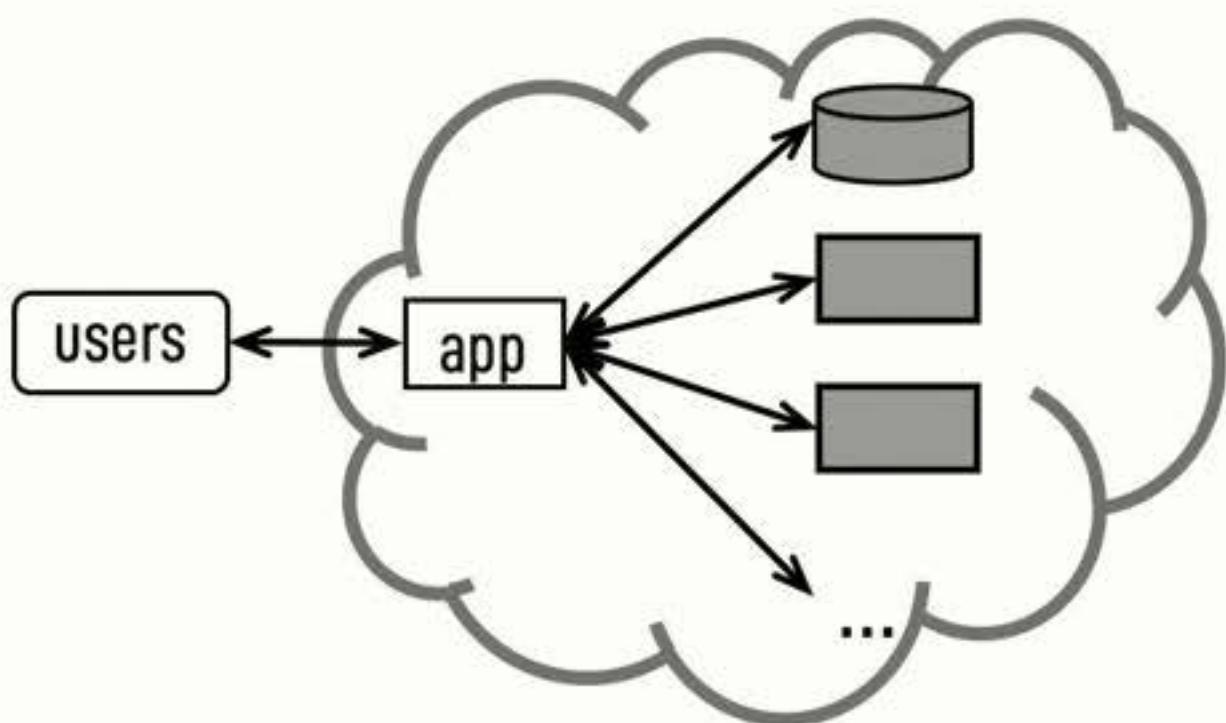


- Requirement (challenge): whole search pipeline needs to be verifiable



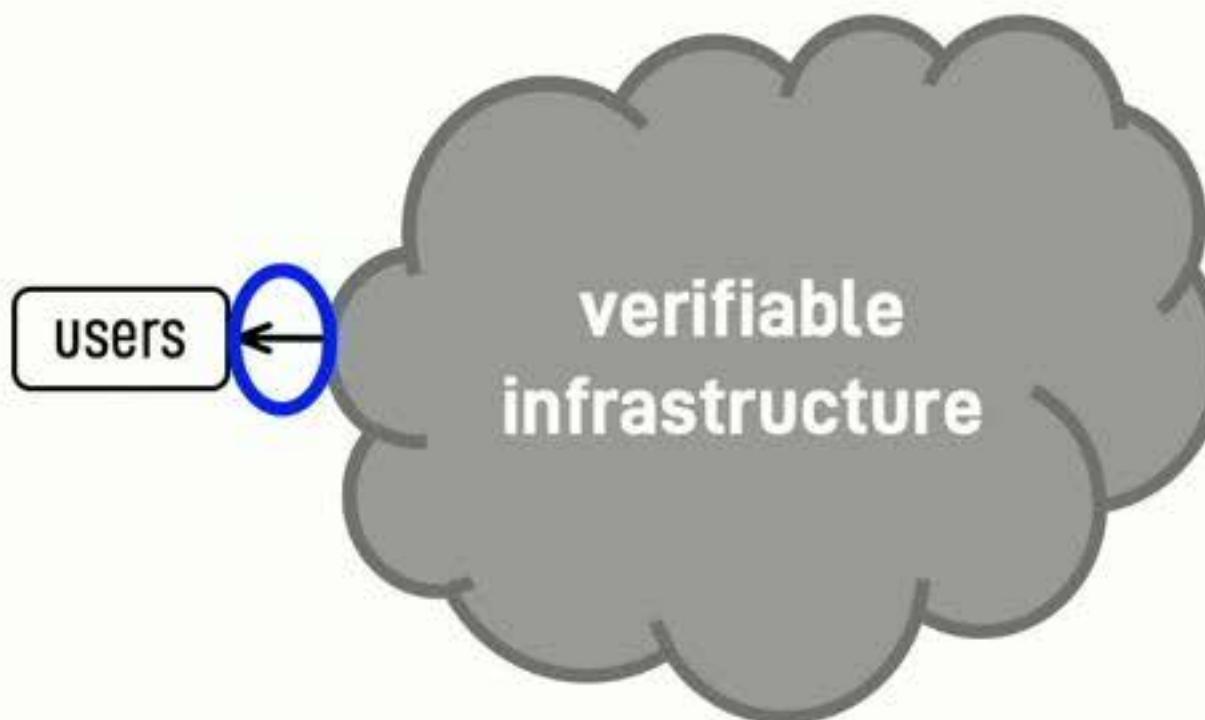
Question #3:

How to verify properties other than execution integrity?



Question #3:

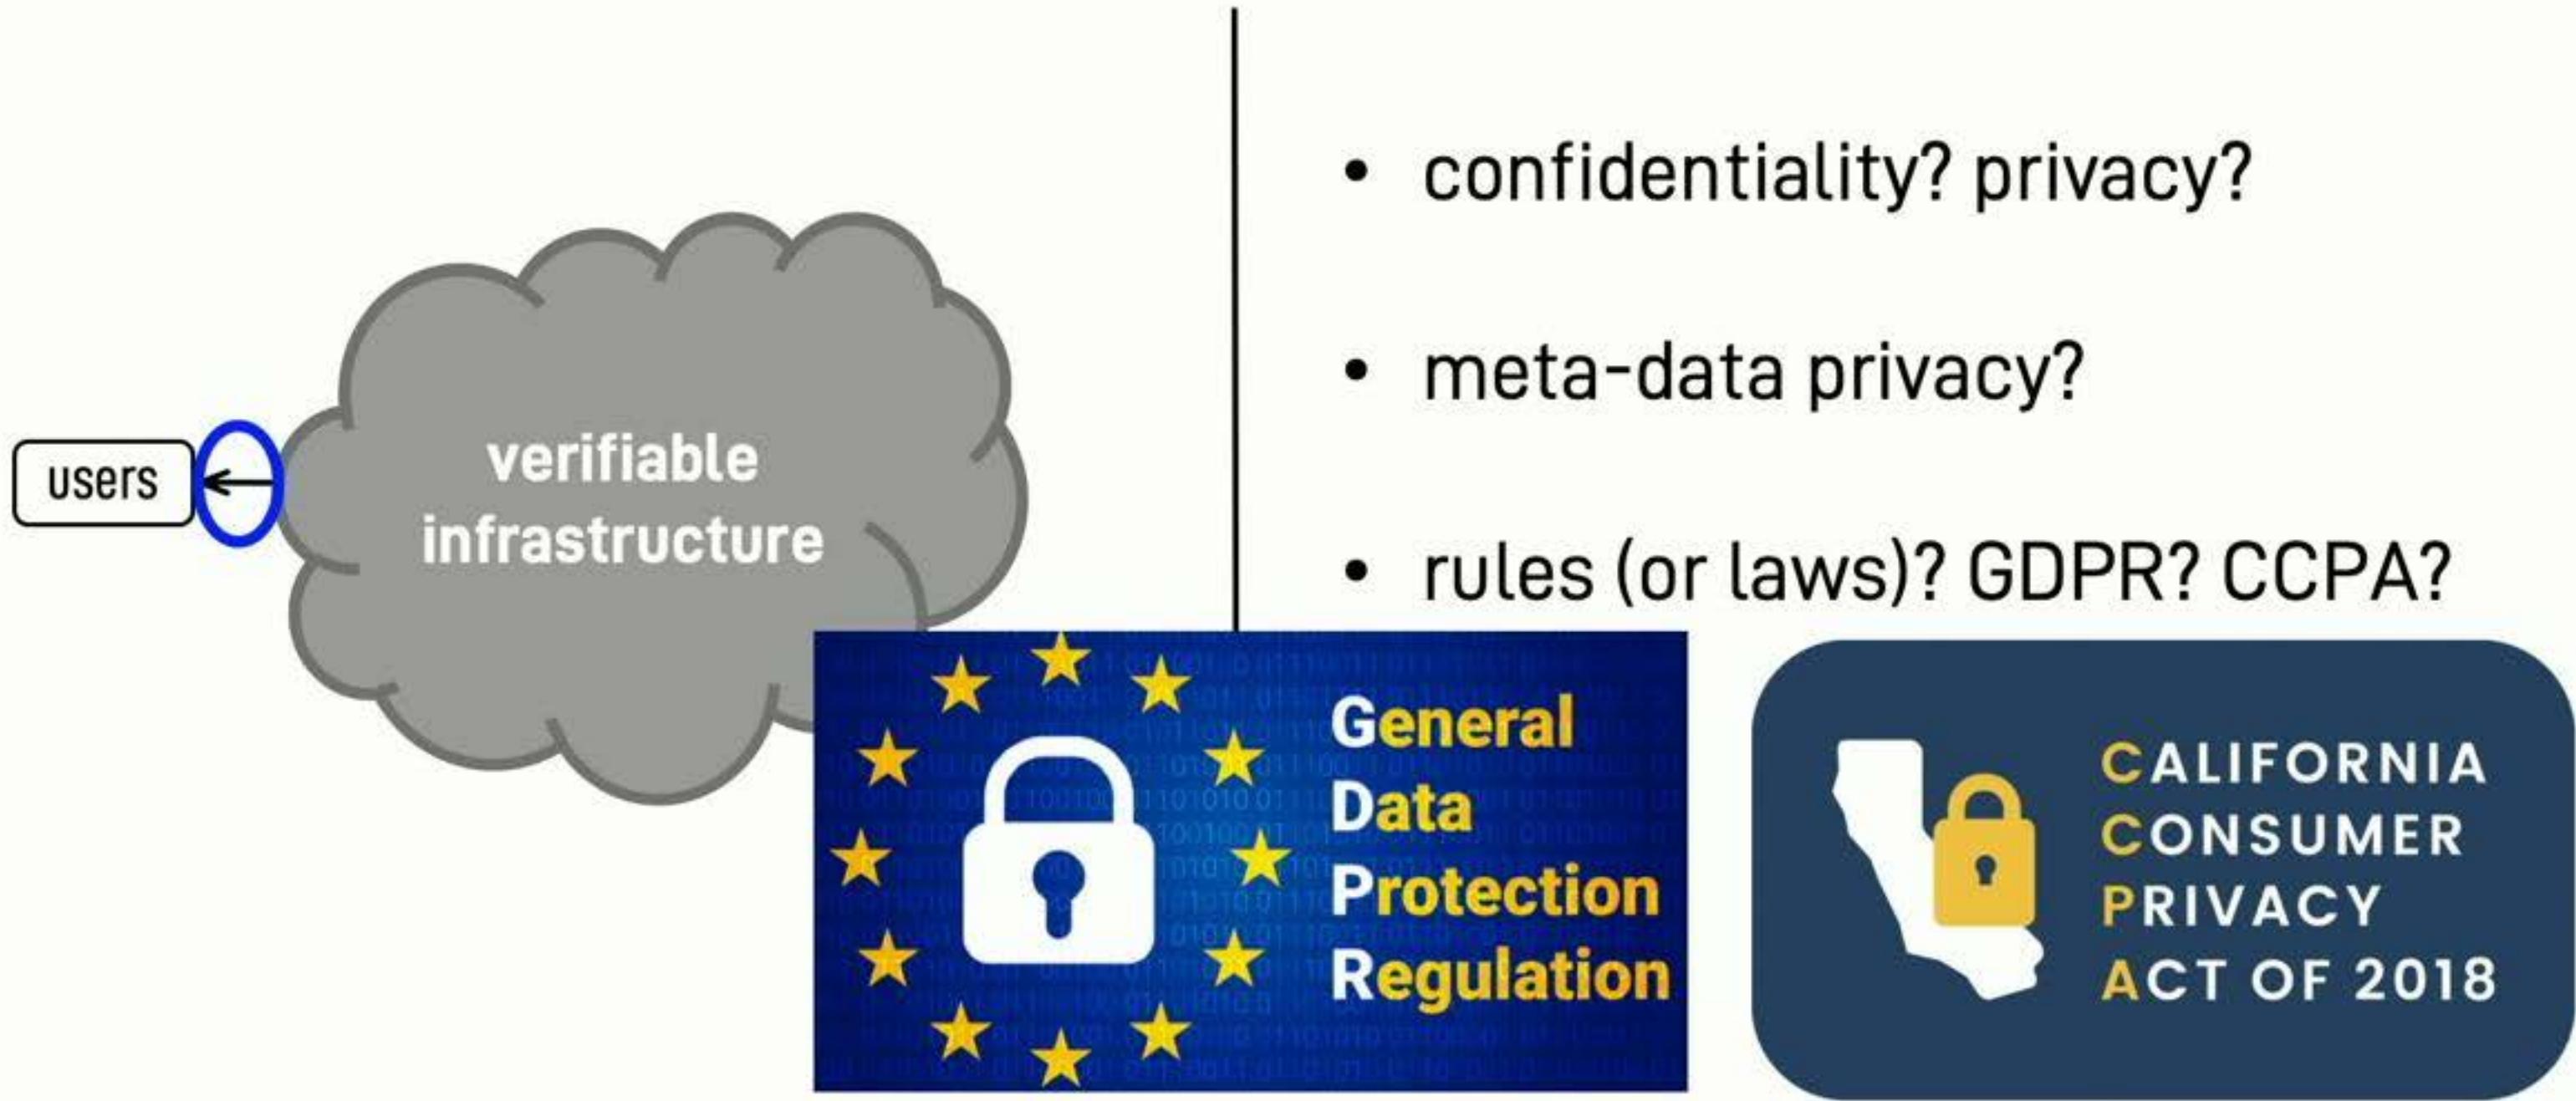
How to verify properties other than execution integrity?



- confidentiality? privacy?
- meta-data privacy?

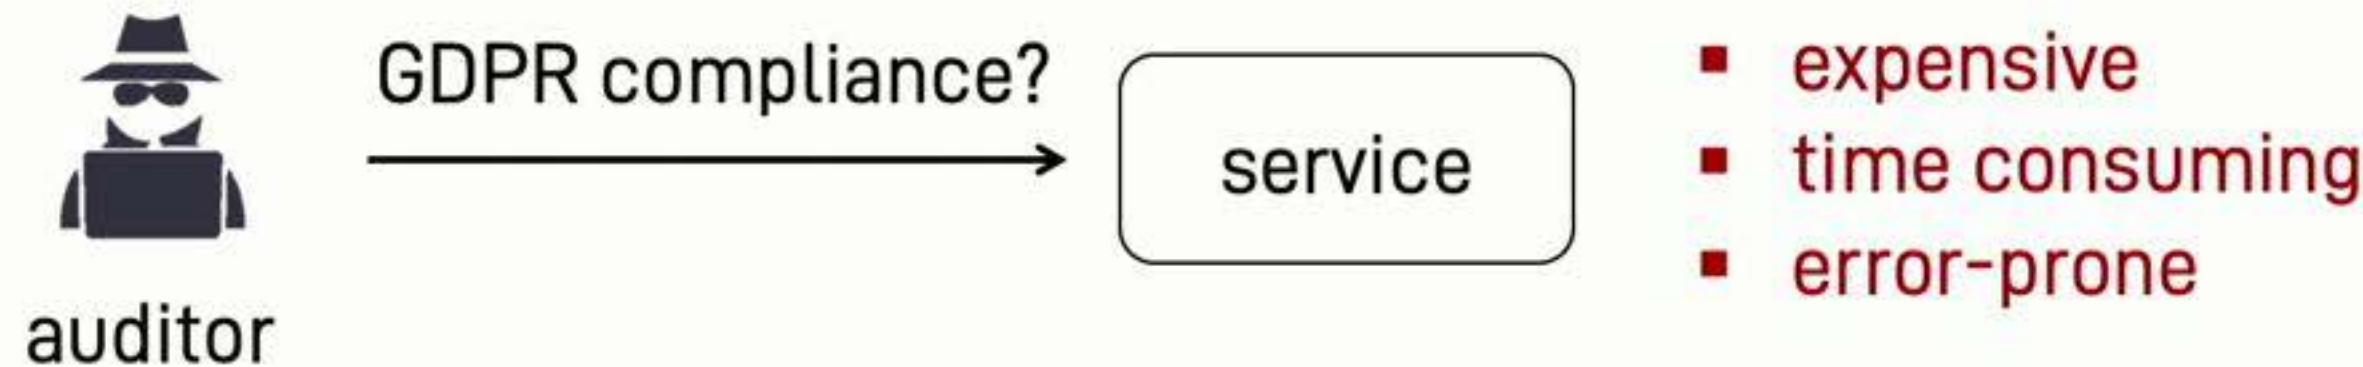
Question #3:

How to verify properties other than execution integrity?



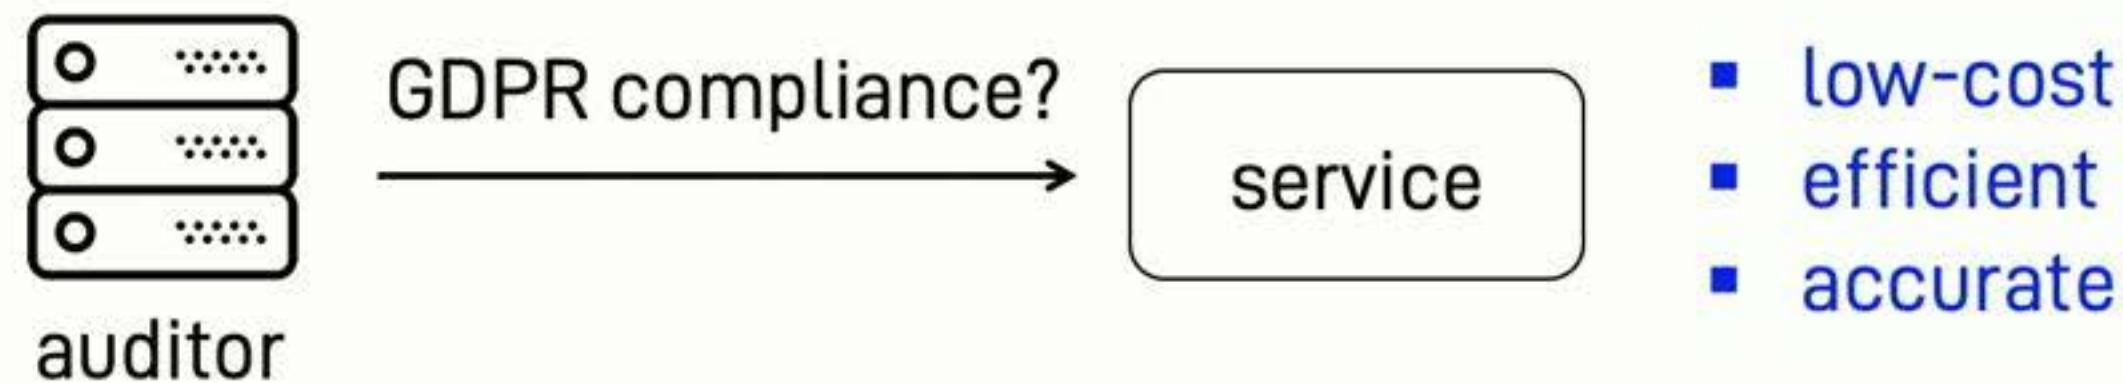
A provable GDPR framework for web apps

- Motivation



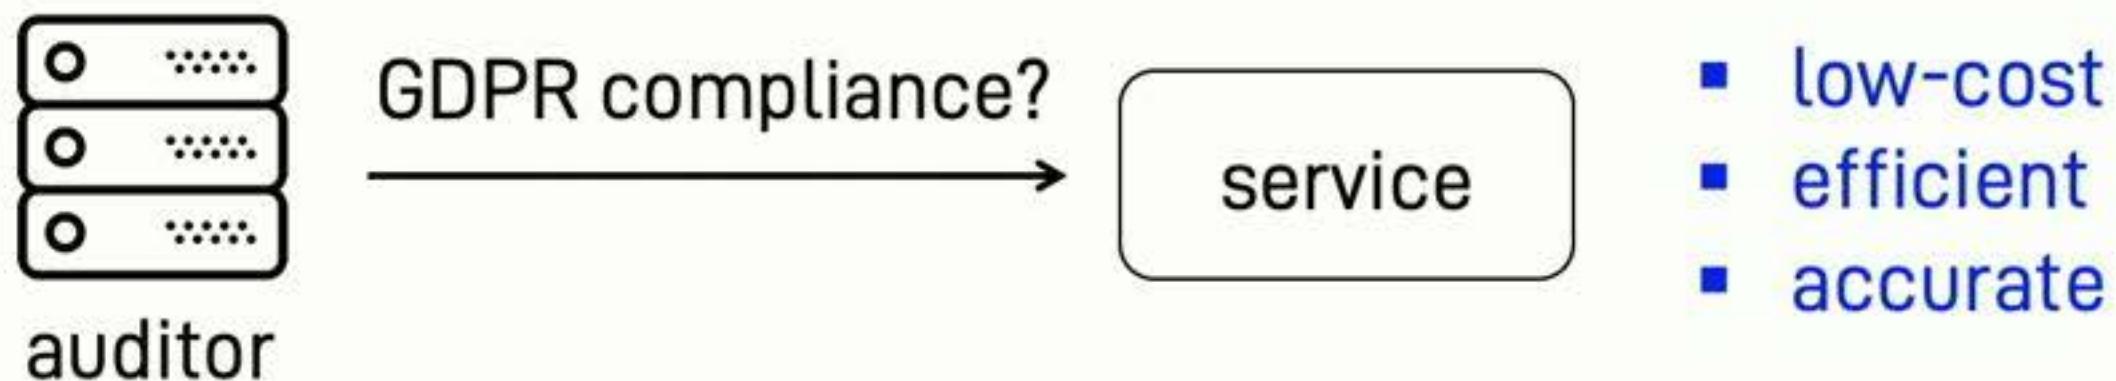
A provable GDPR framework for web apps

- Motivation



A provable GDPR framework for web apps

- Motivation

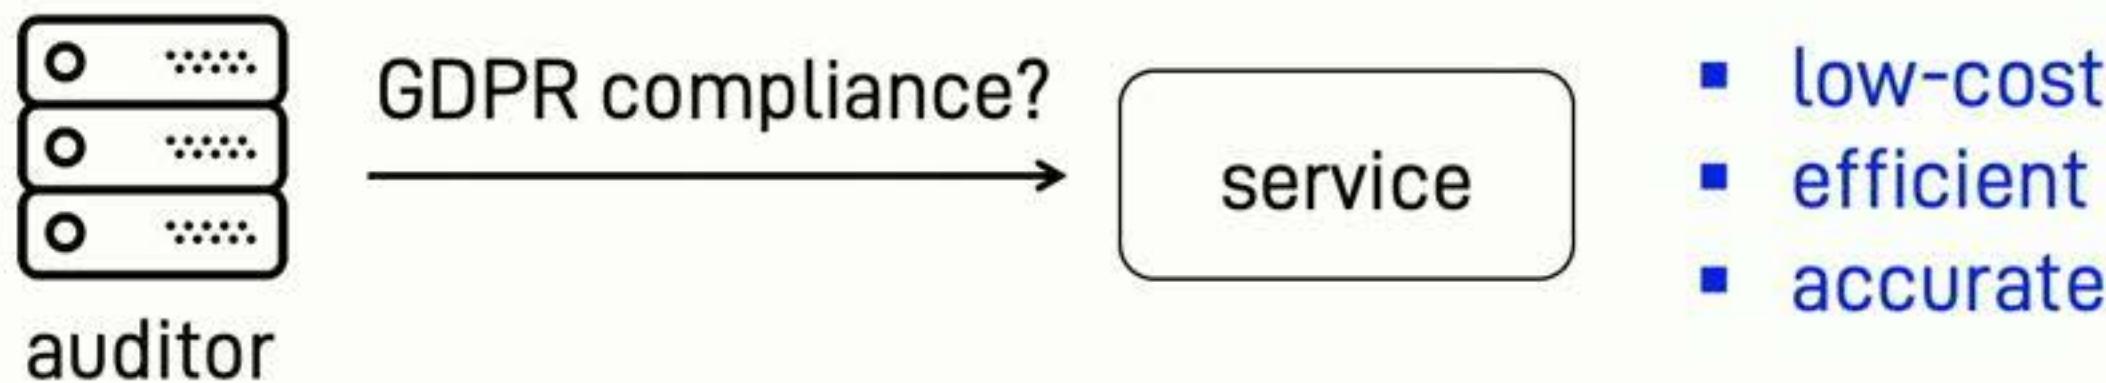


- Requirements (challenges):

- a machine-checkable GDPR definition
- audit puts zero trust on the service
- audit must be efficient

A provable GDPR framework for web apps

- Motivation

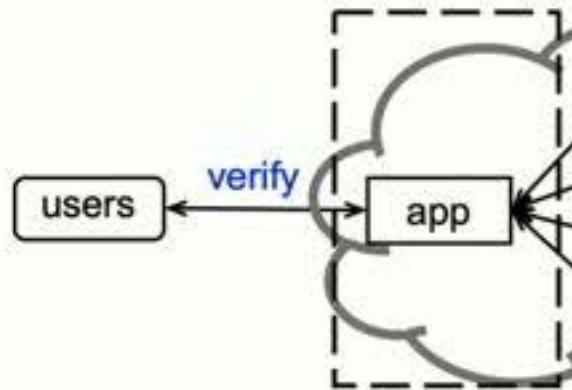


- Requirements (challenges):

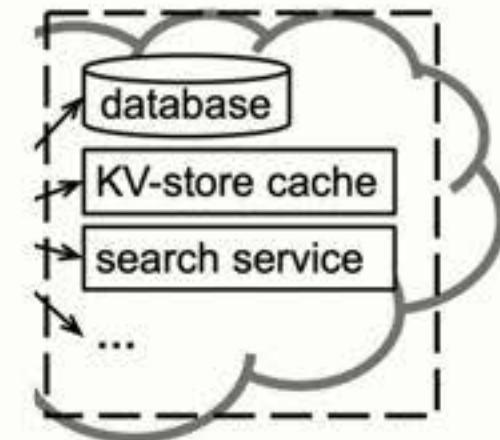
- a machine-checkable GDPR definition
- audit puts zero trust on the service
- audit must be efficient

} Orochi and Cobra can help.

1 Verify white-box services (Orochi [SOSP'17])



2 Verify black-box services (Cobra: verify Serializability)



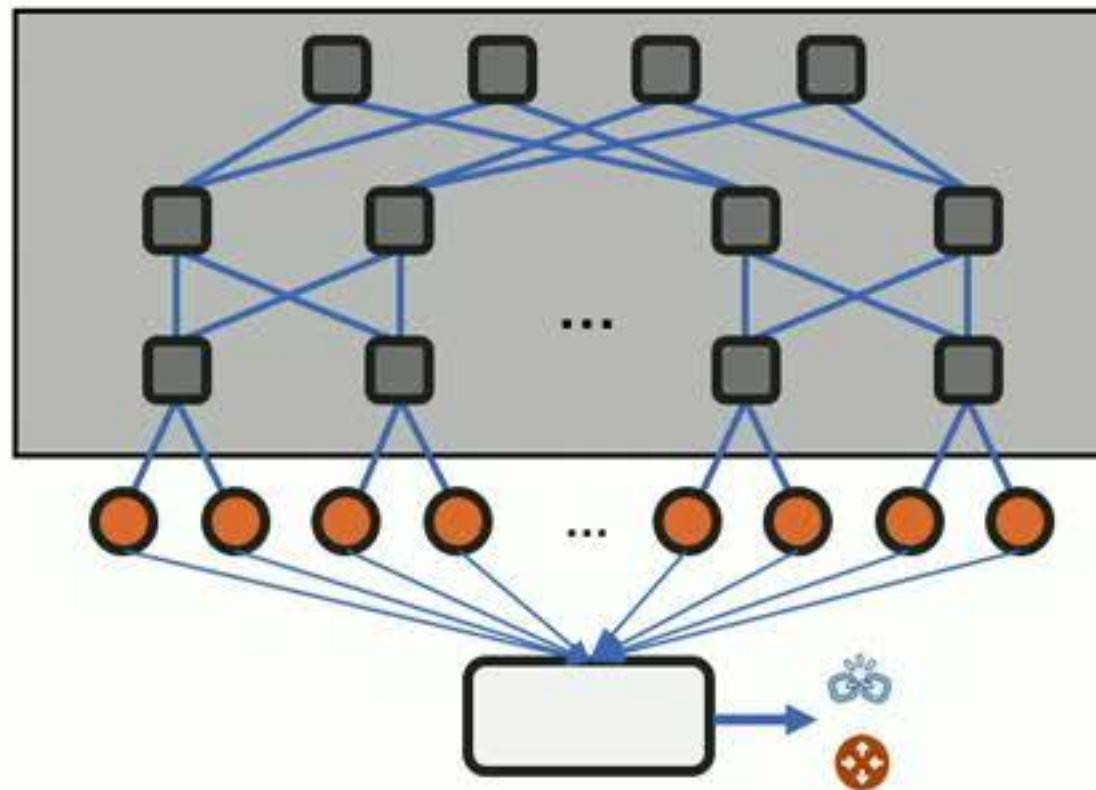
3 Build composable verifiable framework (future work)

4 Other past work

- Troubleshooting data center networks [NSDI'19]
- Protecting secret via security-oriented offloading [EuroSys'15]

NetBouncer [TJGZWDBX, NSDI'19] localizing failures in data center networks

without trusting any
failure information from network

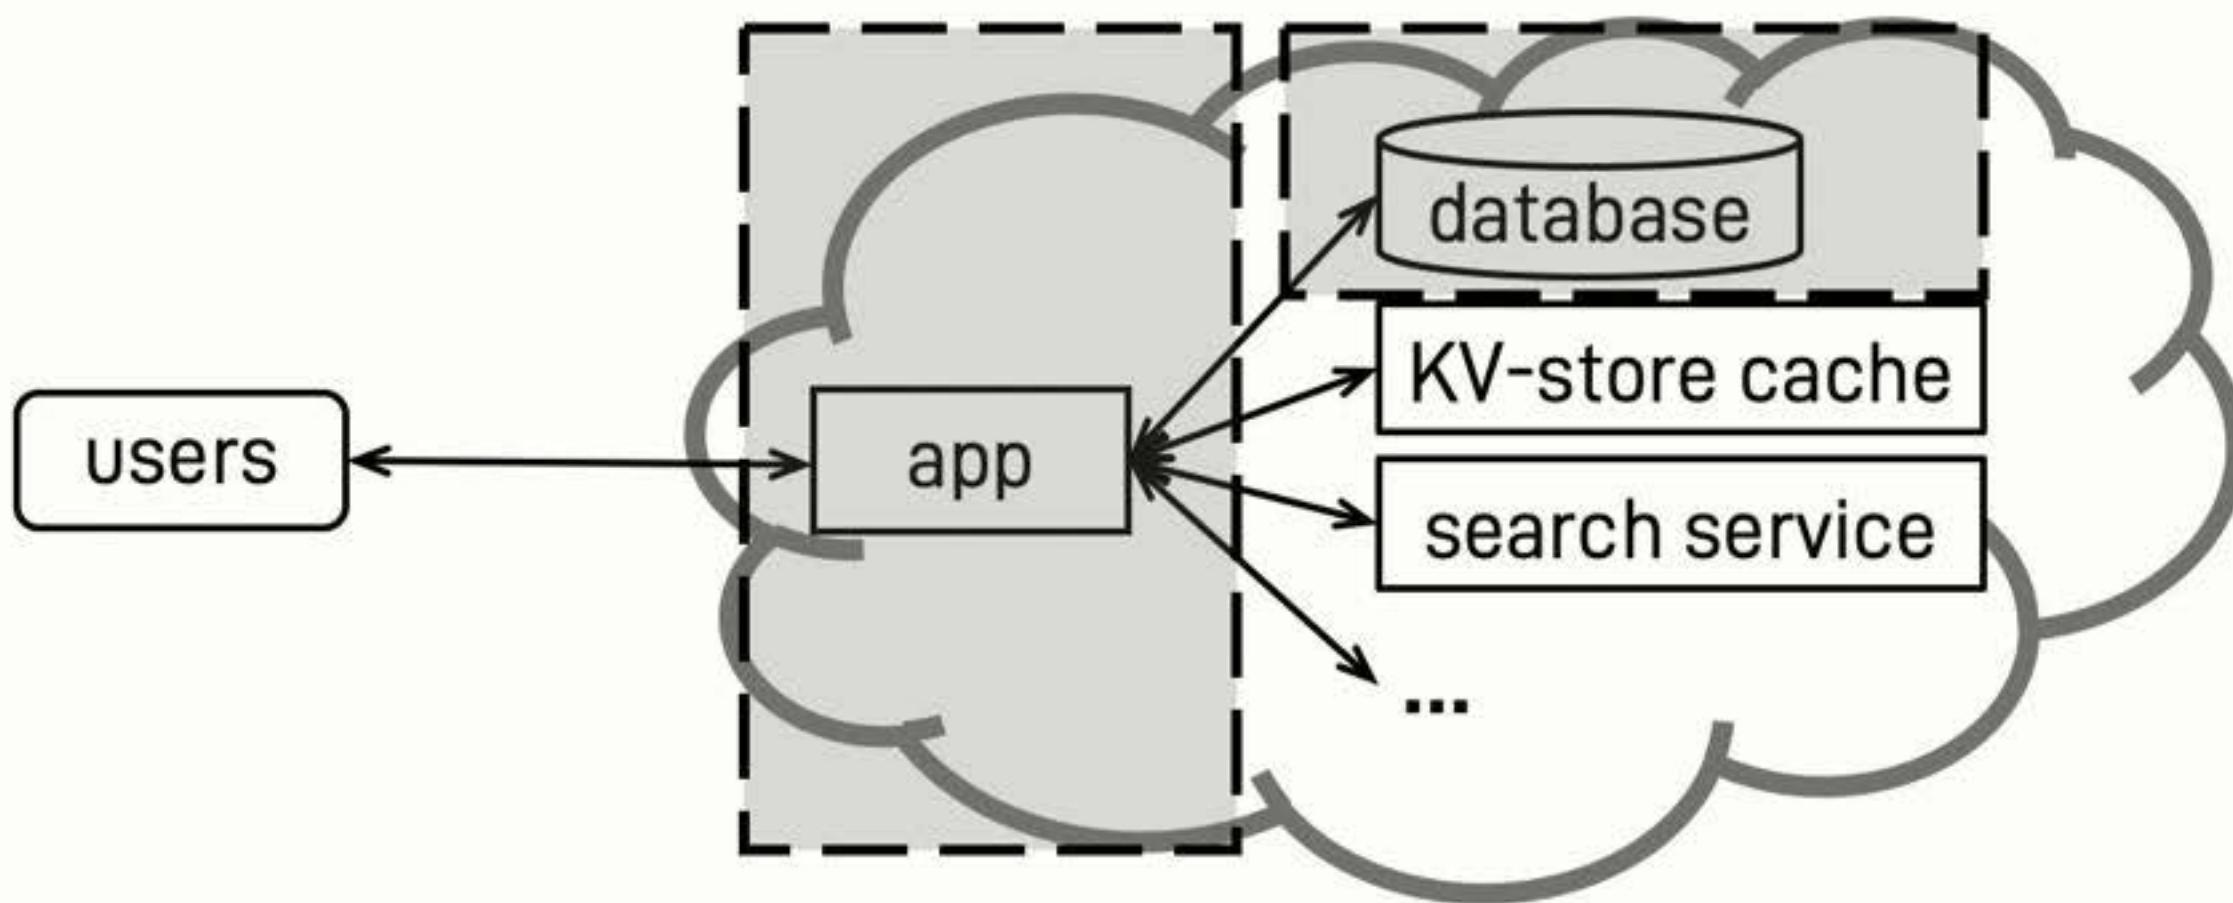


infer failures with end-to-end
observation of network packet loss

- deployed in Microsoft Azure
- detected overlooked failures

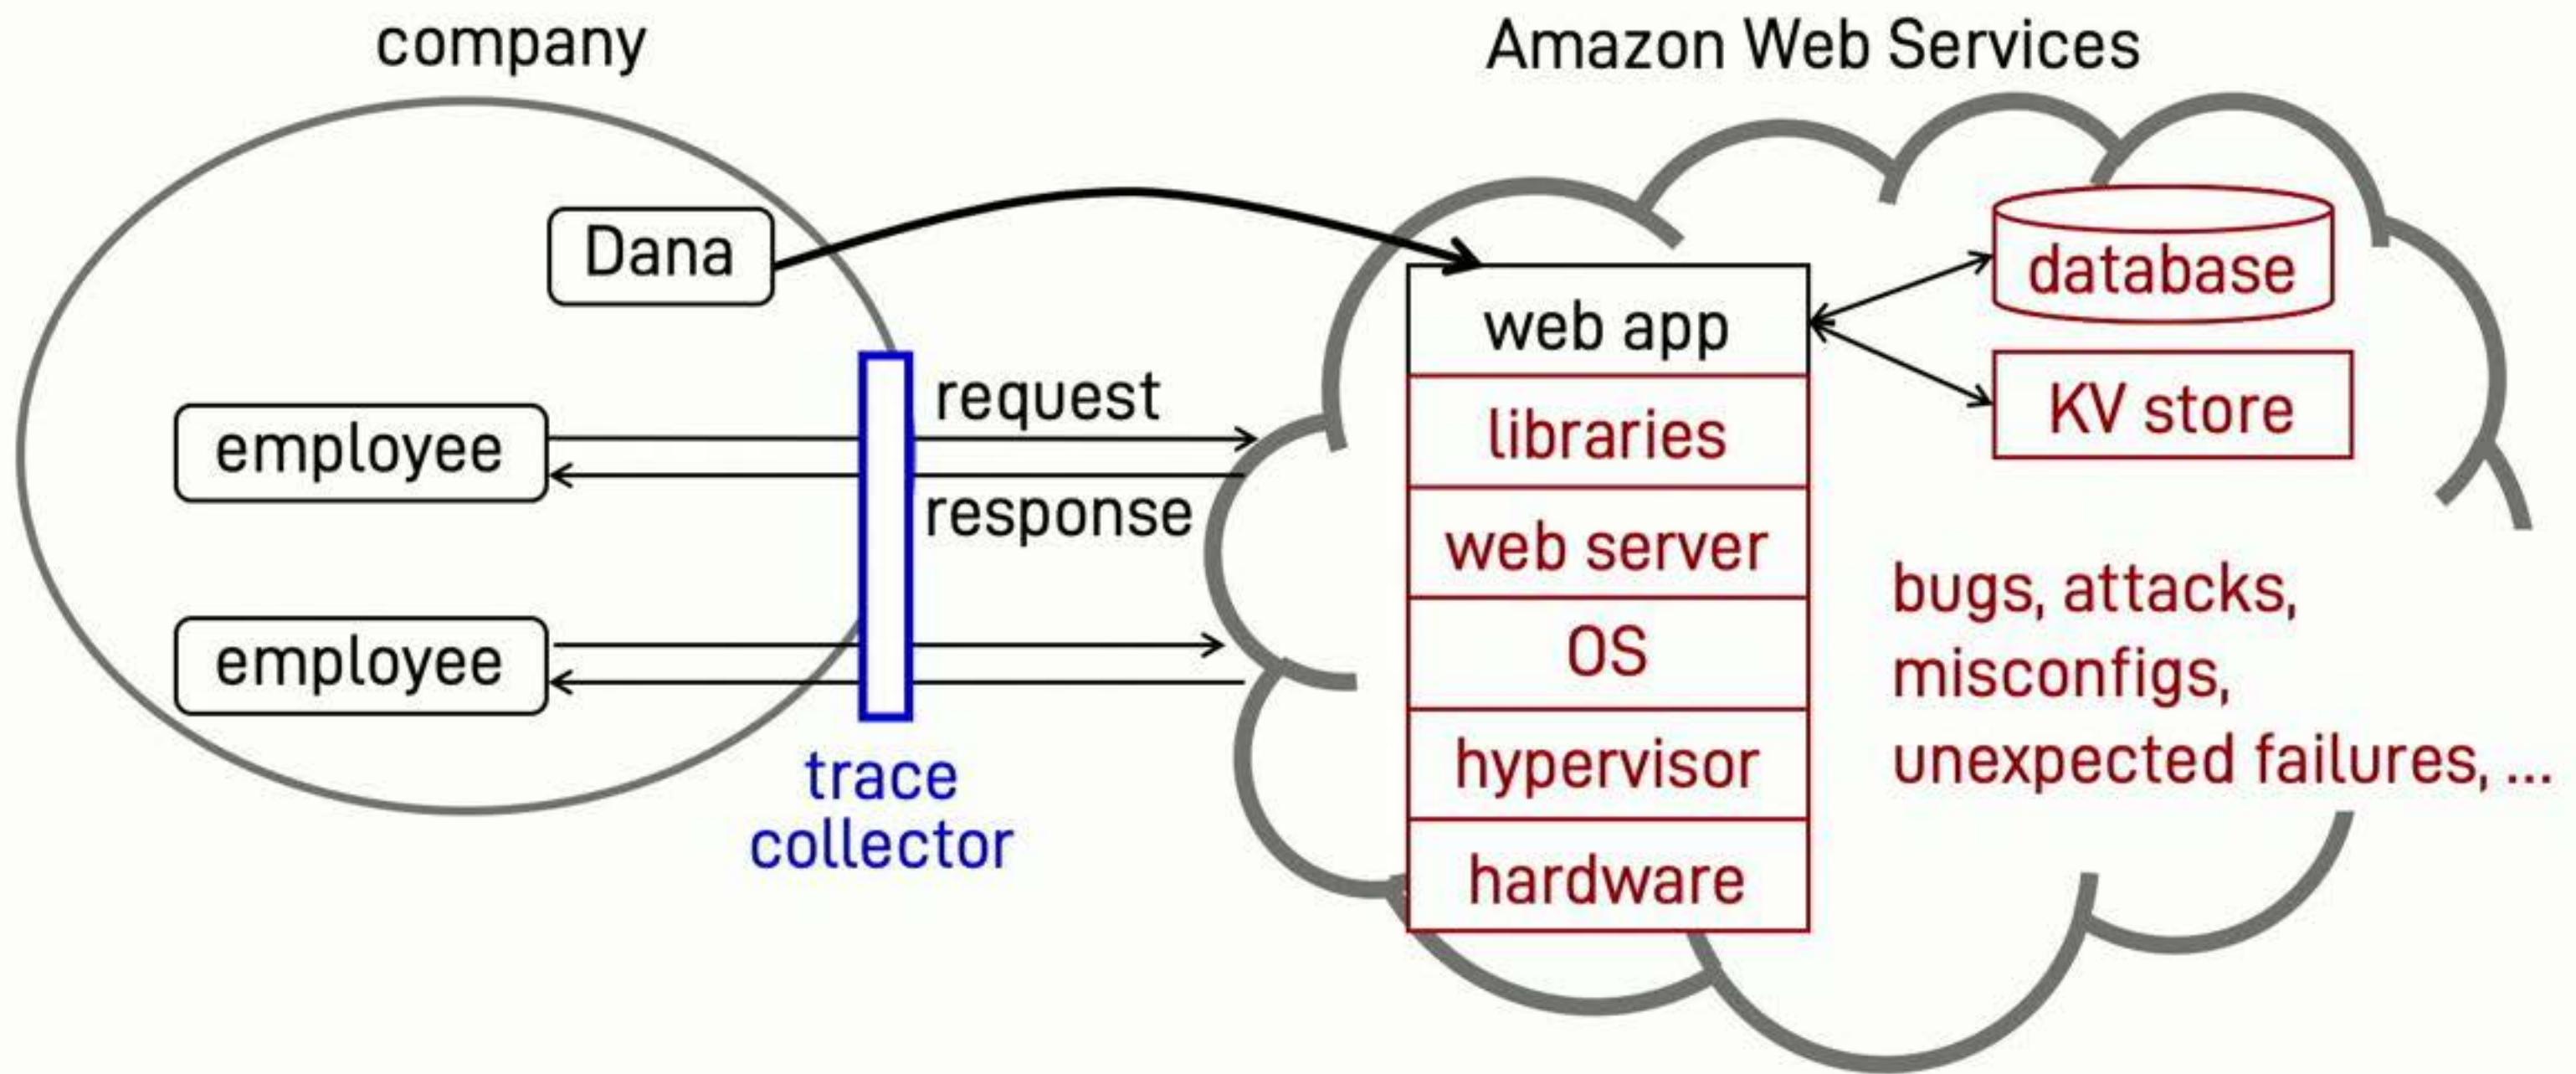
Verifiable infrastructure

... enables users to verify outsourced services.



Orochi:
verifying white-box services

Cobra:
verifying black-box databases



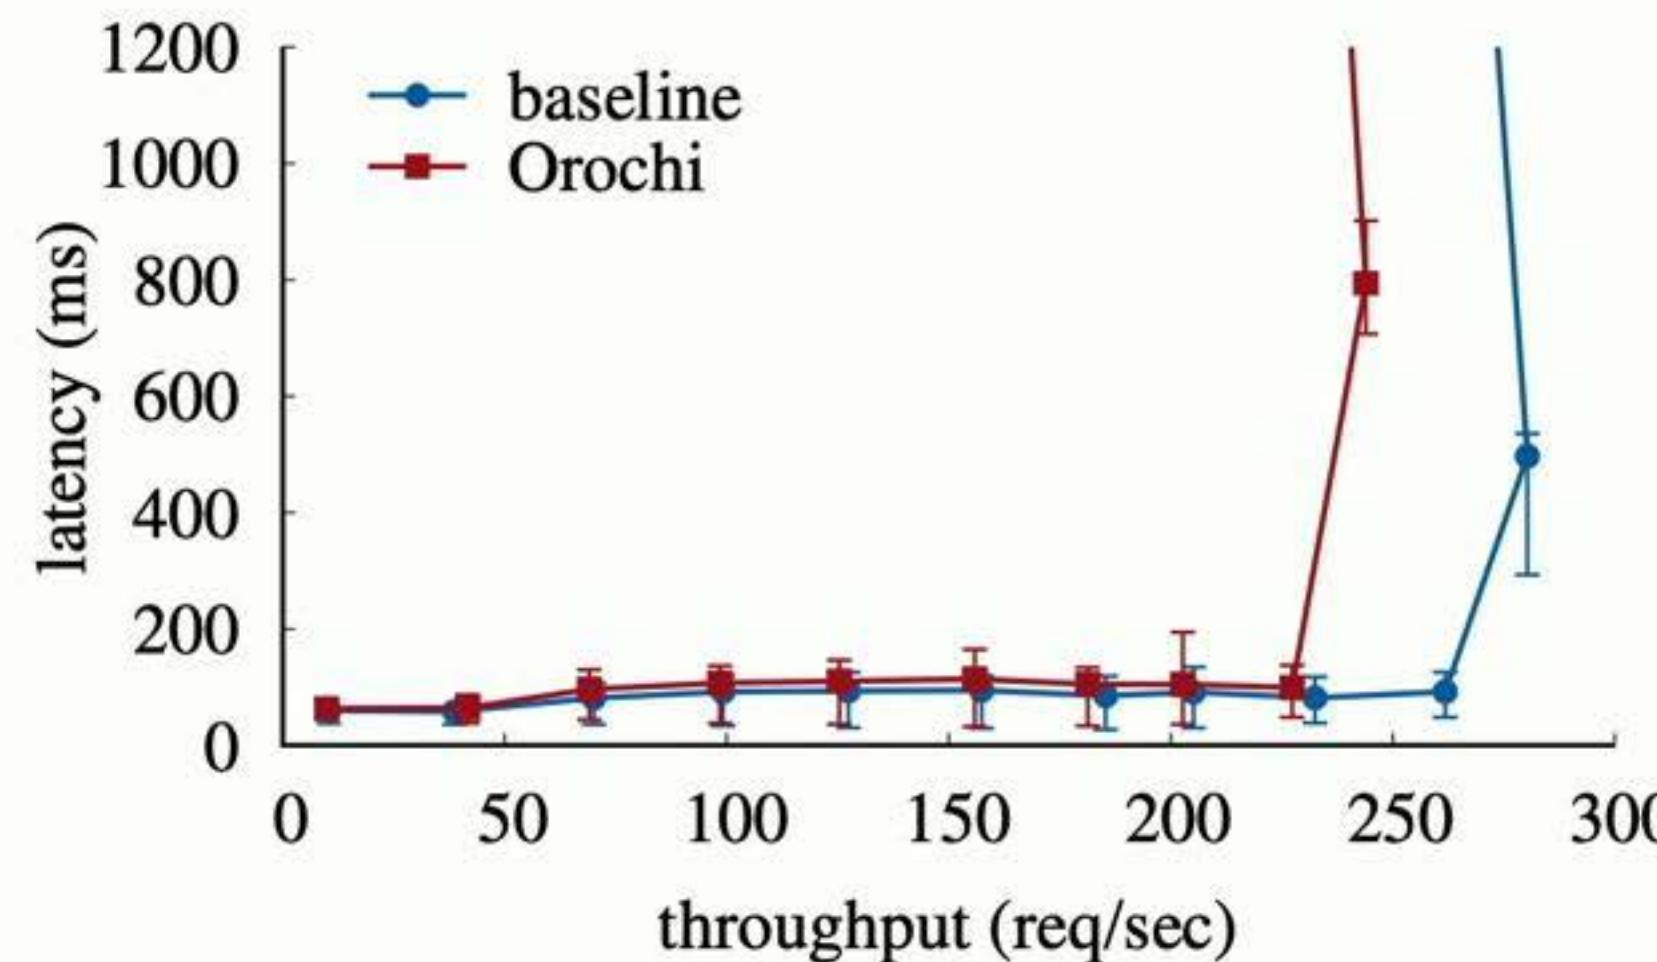
What are the (server's) CPU/network/storage costs?

CPU	Network			Storage
MediaWiki's workload	trace (per req)	advice (per req)	Orochi's overhead	MediaWiki's workload
4.7%	MediaWiki's workload	7.1KB	1.7KB	11.4%

What are the (server's) CPU/network/storage costs?

CPU	Network			Storage
MediaWiki's workload	trace (per req)	advice (per req)	Orochi's overhead	MediaWiki's workload
4.7%	MediaWiki's workload	7.1KB	1.7KB	1.0x
phpBB's workload	phpBB's workload	5.7KB	0.3KB	11.4%
8.6%	HotCRP's workload	3.2KB	0.4KB	2.7%
HotCRP's workload	HotCRP's workload			10.9%
5.9%				1.5x

Orochi imposes small overheads on throughput and latency



phpBB's workload

Why Serializability (SER)?

gold standard
isolation level



fundamental
and
challenging