Drongo

Speeding Up CDNs with Subnet Assimilation *from the Client*

CoNEXT '17
Incheon, South Korea
CDN & Caching Session

Authors:
Marc Anthony Warrior
Uri Klarman
Marcel Flores
Aleksandar Kuzmanovic
Northwestern University

Bird's Eye View

- What is Drongo?
- Why we need Drongo
- Performance Analysis
- Thoughts & Conclusions
- Questions

2

What is Drongo?

What is Drongo?

It’s a bird!

What is Drongo?

It’s a bird!

What is Drongo?

It’s a bird!

What is Drongo?

It's a **system that allows *end-users* to enhance the QoS** (quality of service) they get from Content Distribution Networks (**CDN**s)

What is Drongo?

It's a **system that allows *end-users* to enhance the QoS** (quality of service) they get from Content Distribution Networks (**CDN**s)

(in this talk, QoS = ***latency***)

Why Latency?

Why Latency?

- Latency is **time**

Why Latency?

- Latency is **time**
- Latency is **money**
  - Google (Marissa Mayer), Amazon (Greg Linden)
    - Web 2.0 Summet, glinden.blogspot.com

Why Latency?

- Latency is **time**
- Latency is **money**
  - Google (Marissa Mayer), Amazon (Greg Linden)
    - Web 2.0 Summet, glinden.blogspot.com
- Latency is the **bottom line**
  - “What we have found running our applications at Google is that latency is as important, or more important, for our applications than relative bandwidth,” Amin Vahdat (Google)

Drongo helps you
(the end user)
***lower* your own latency!**

Drongo's Effect on Latency

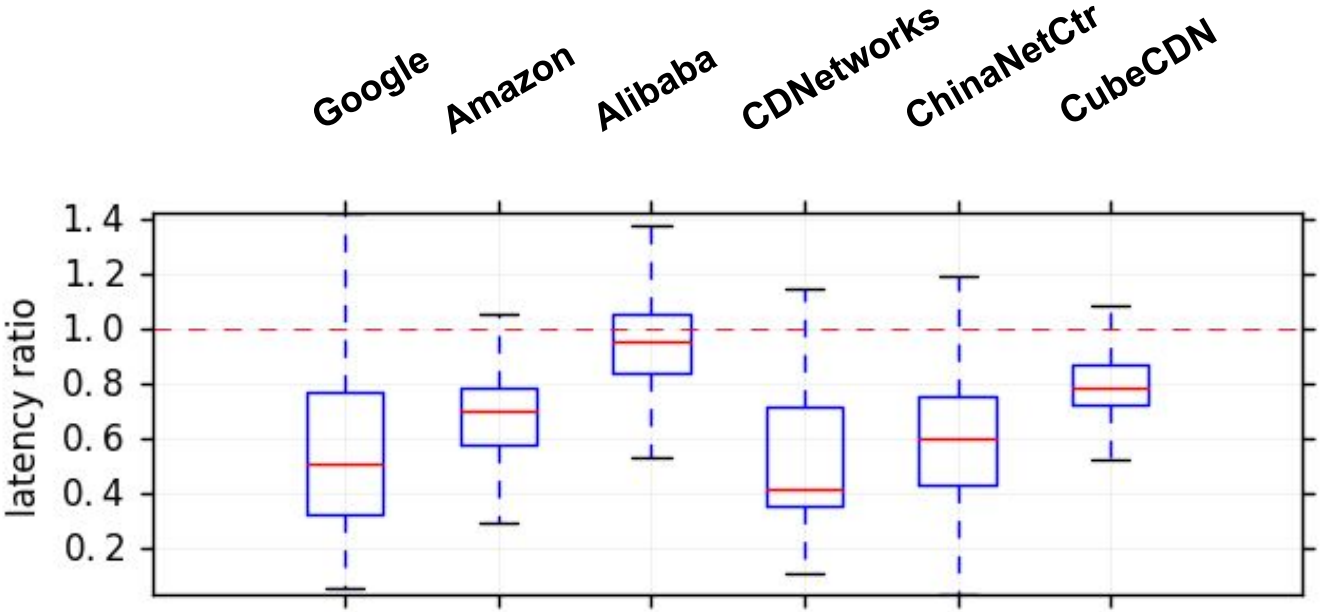

Drongo's Effect on Latency

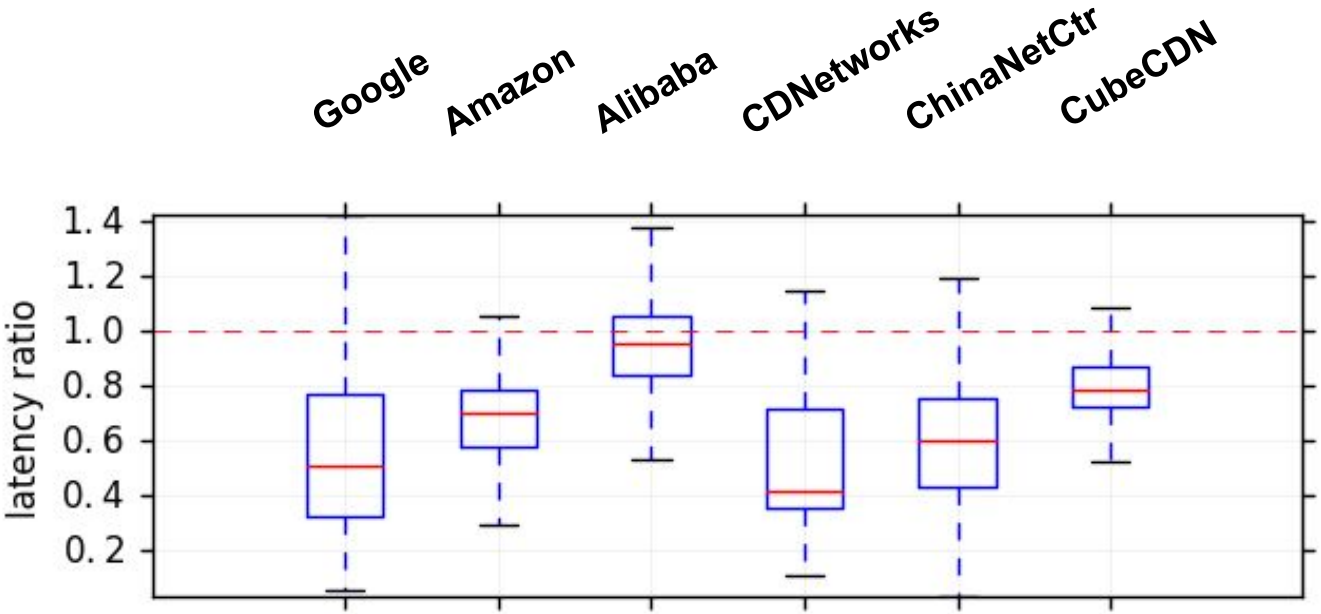

ONLY client-side changes

Example Scenario

Provider wants to serve client

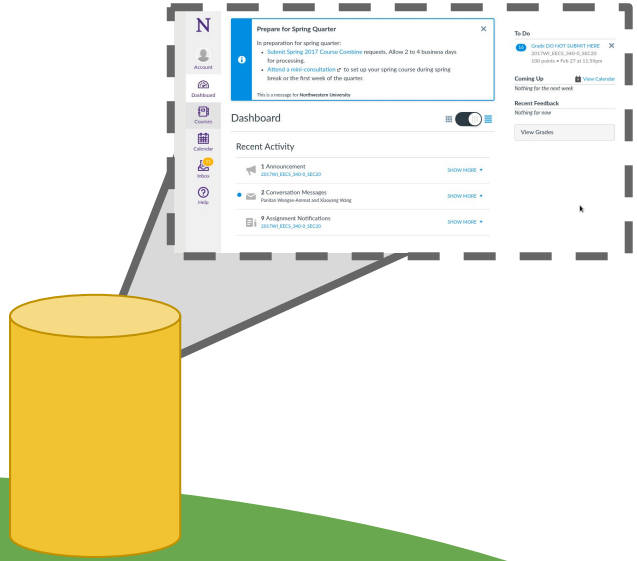

CLIENT

Client is far

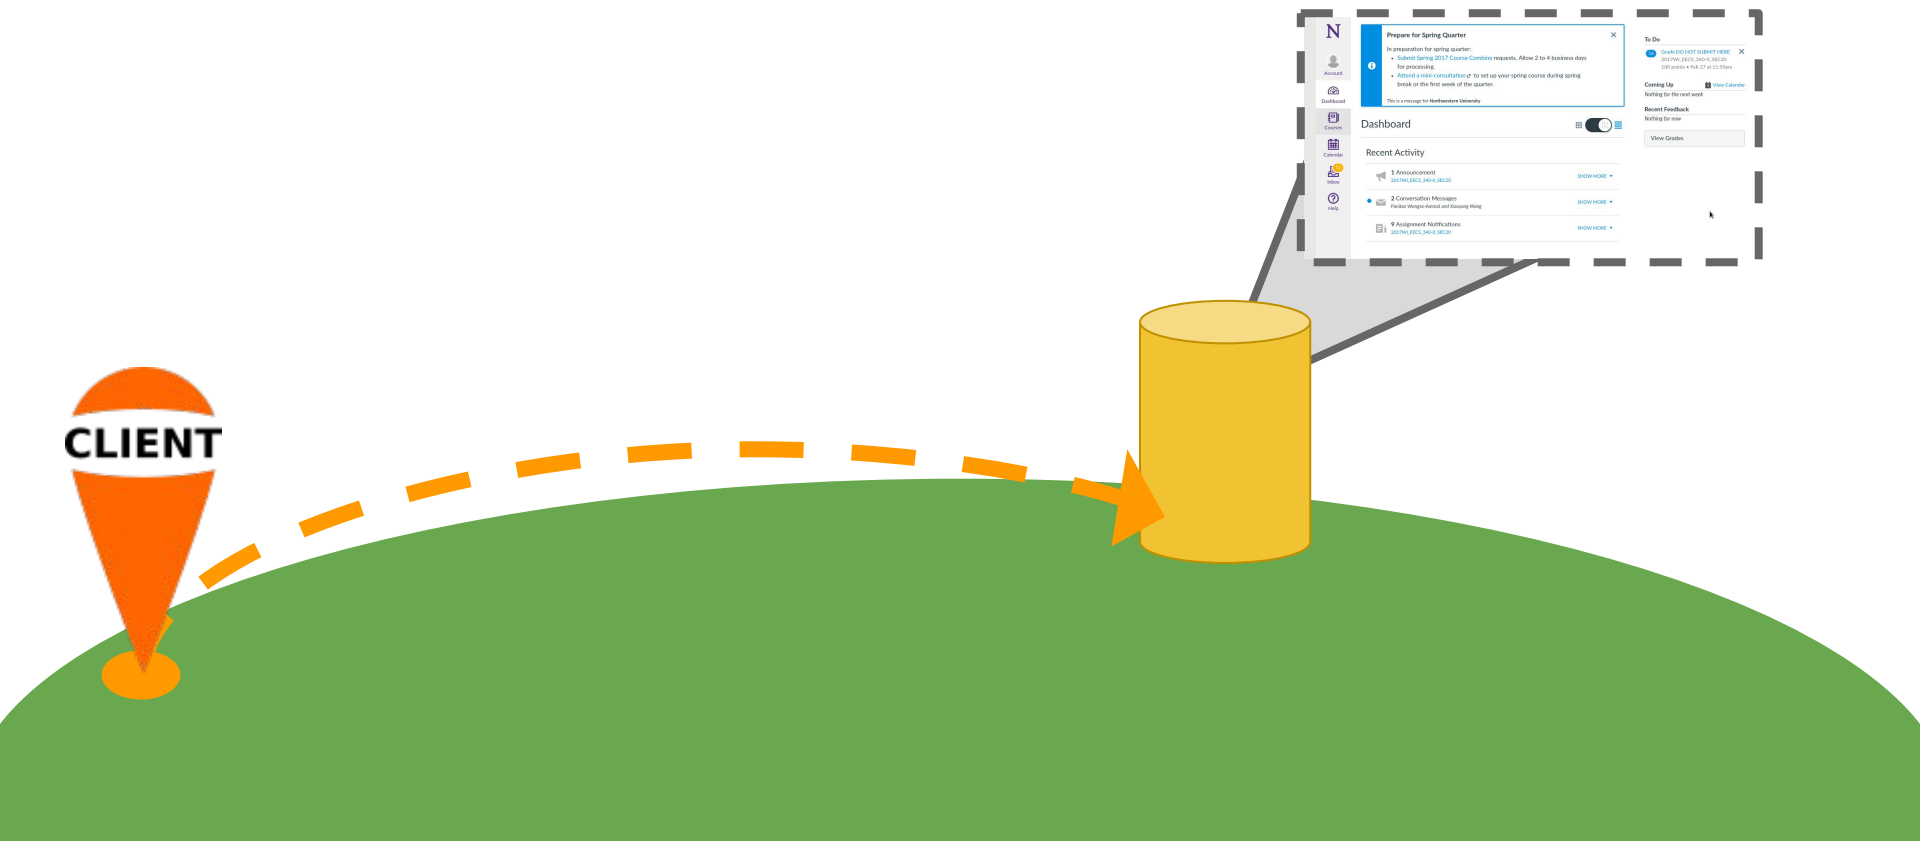

CDN = more replica locations

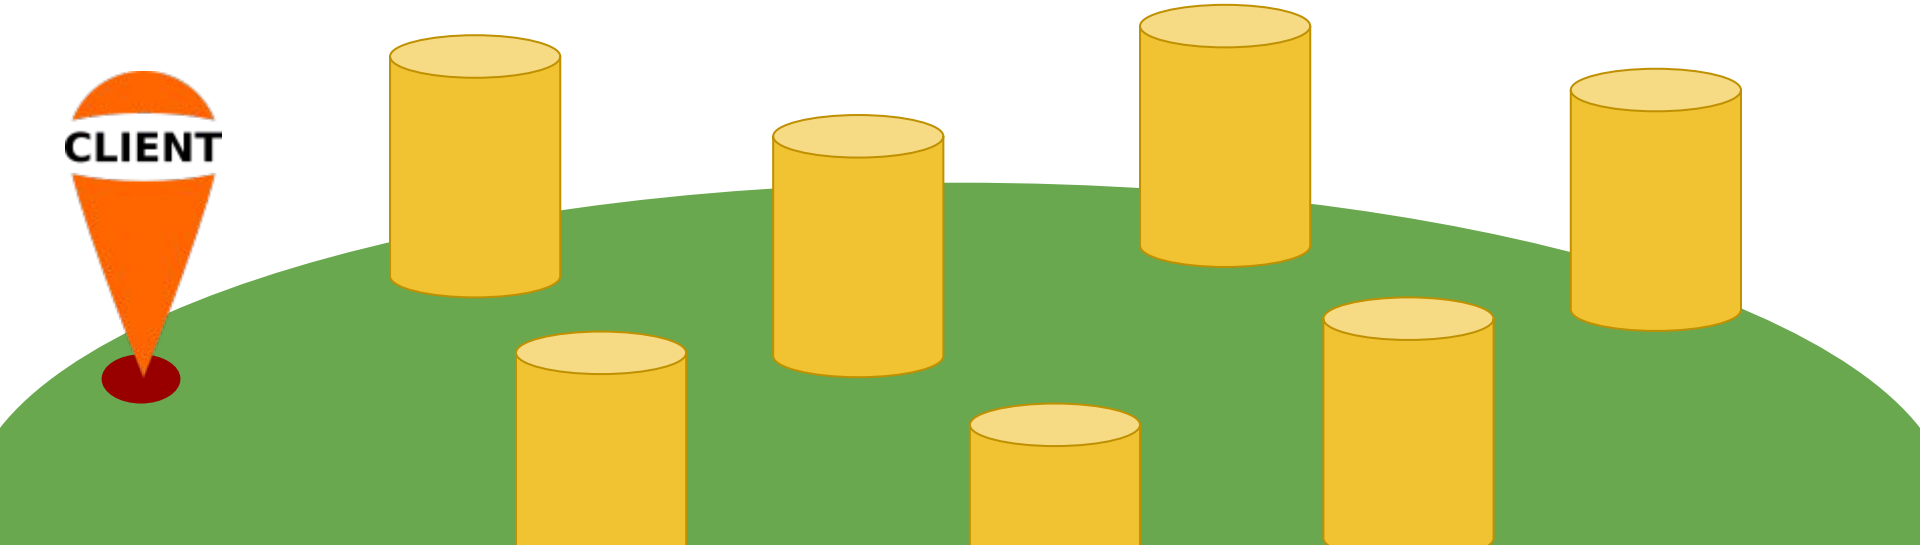

DNS Redirection

Which **replica** serves the client?

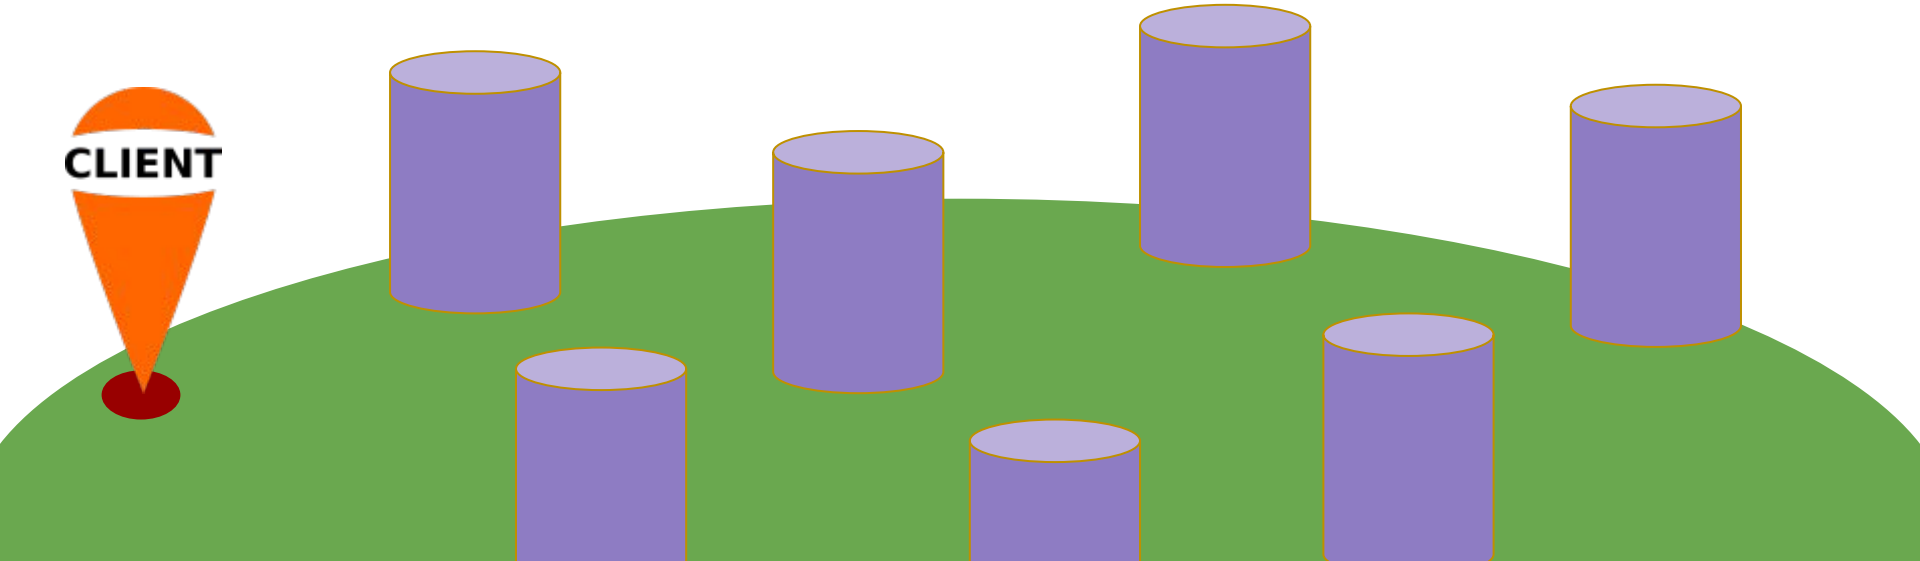

Choose the “closest” server

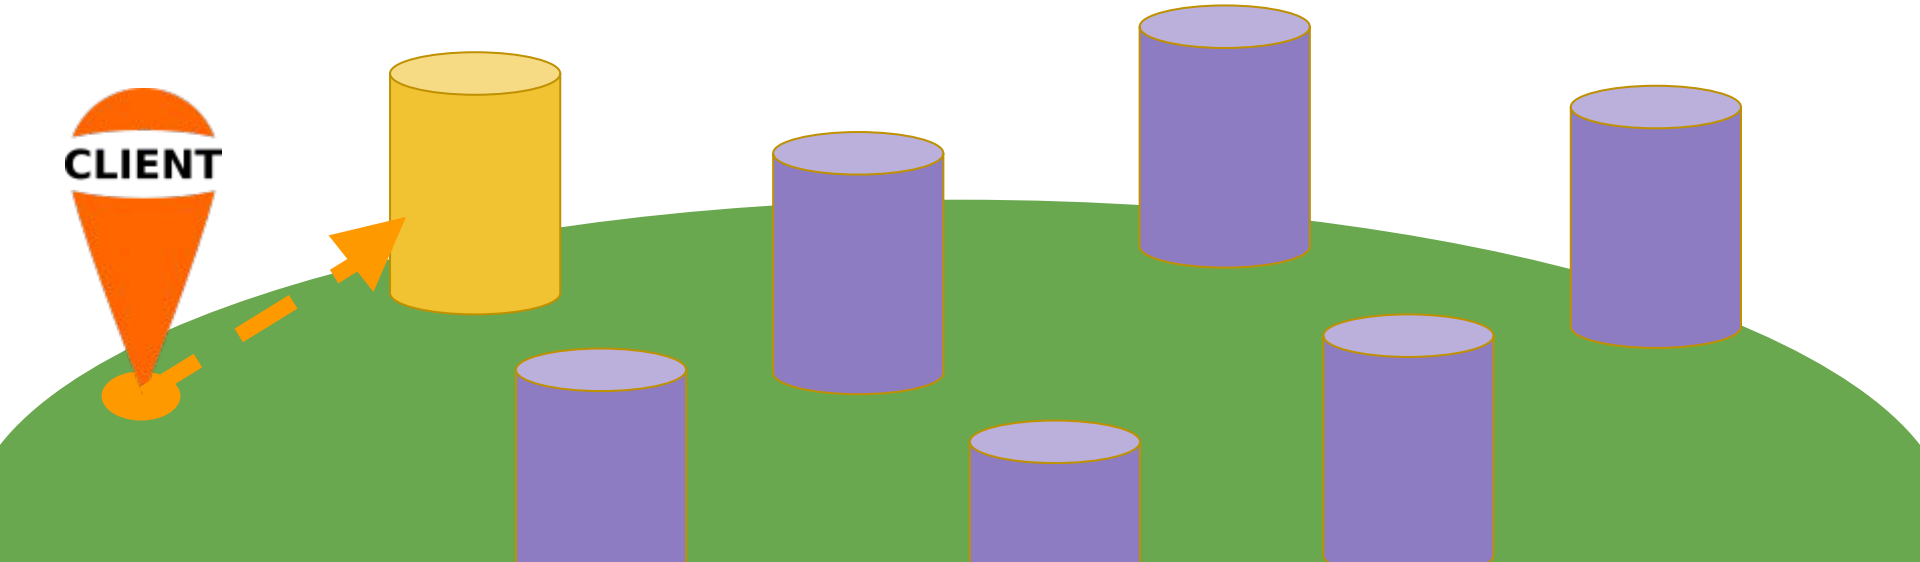

Choose the “closest” server

This choice is **nontrivial**!

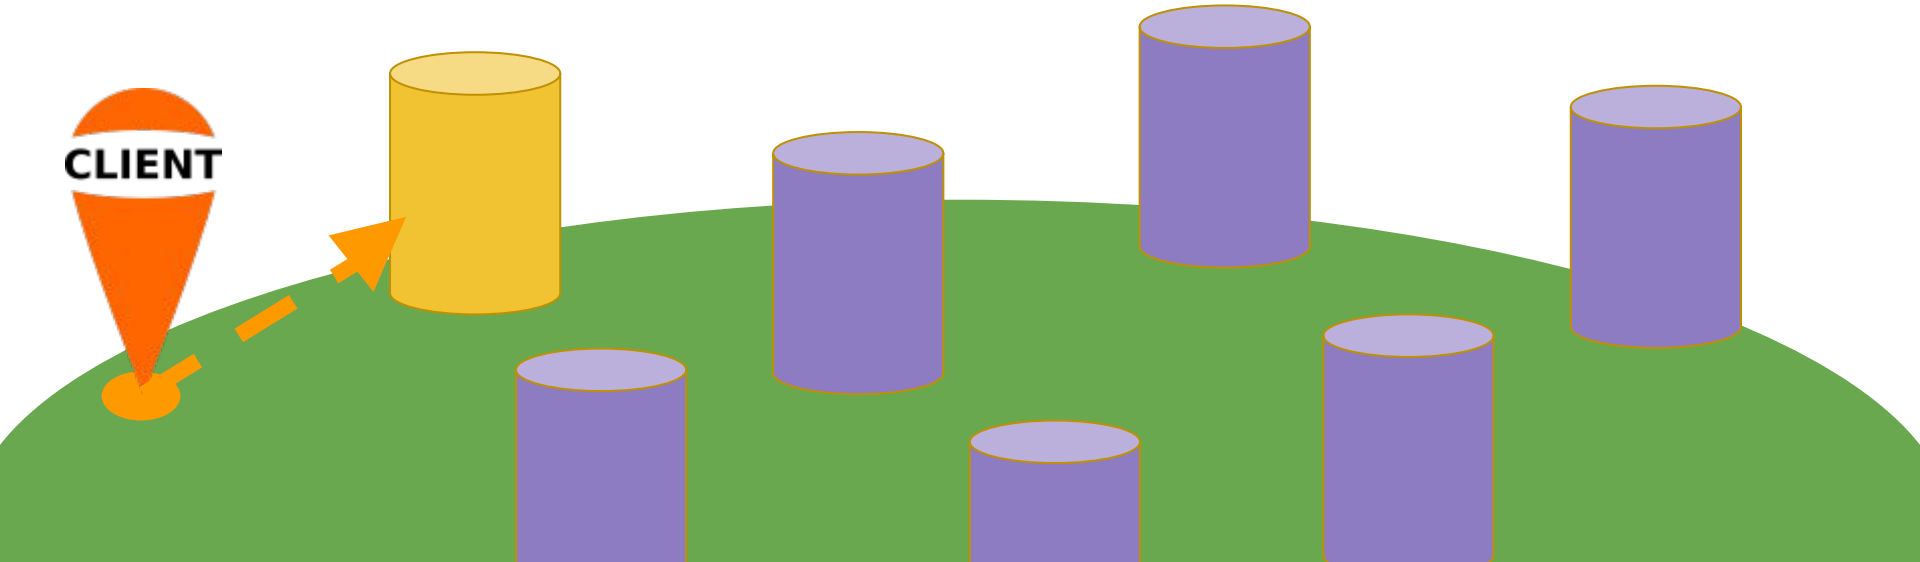

Often Suboptimal Choices!

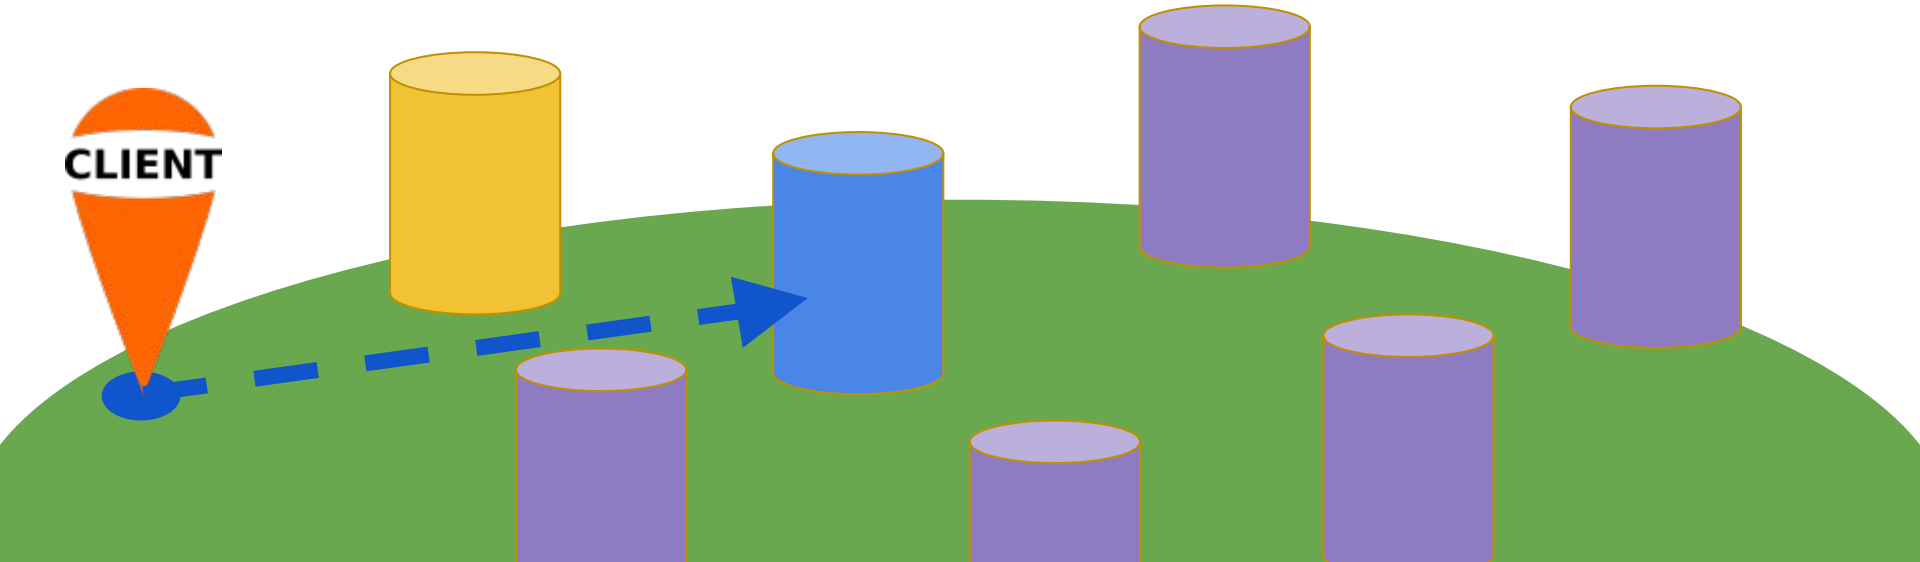

Maybe just a far LDNS…

[Chen - SigComm ’15; Huang - SigComm CCR ‘12; Alzoubi - WWW ‘13; Rula - SigComm ‘14 …]

Ordinary DNS Query

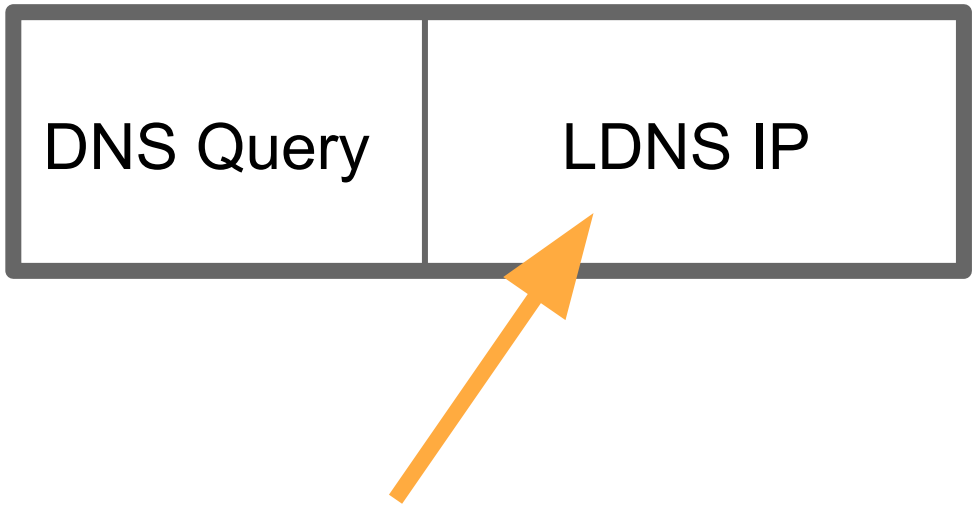

EDNS0 Client-Subnet extension (ECS)

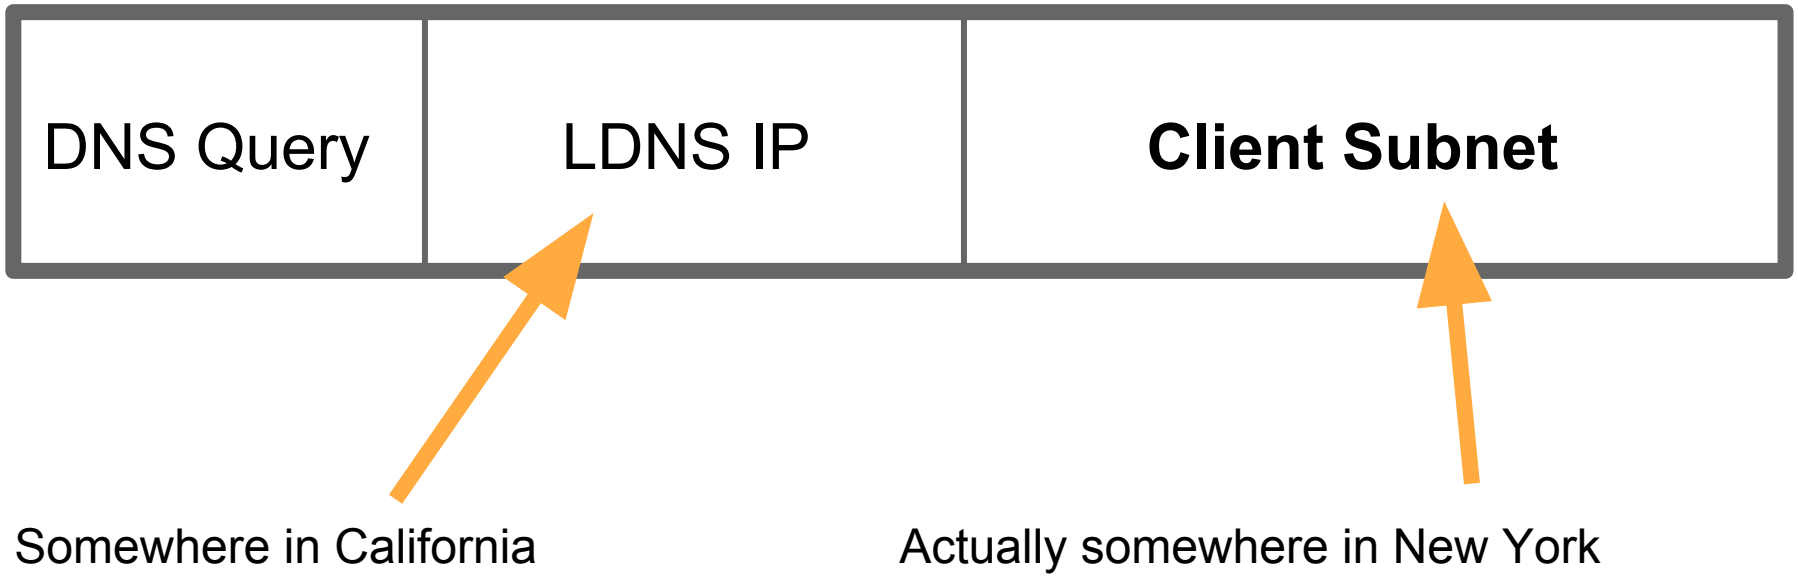

We used ECS:

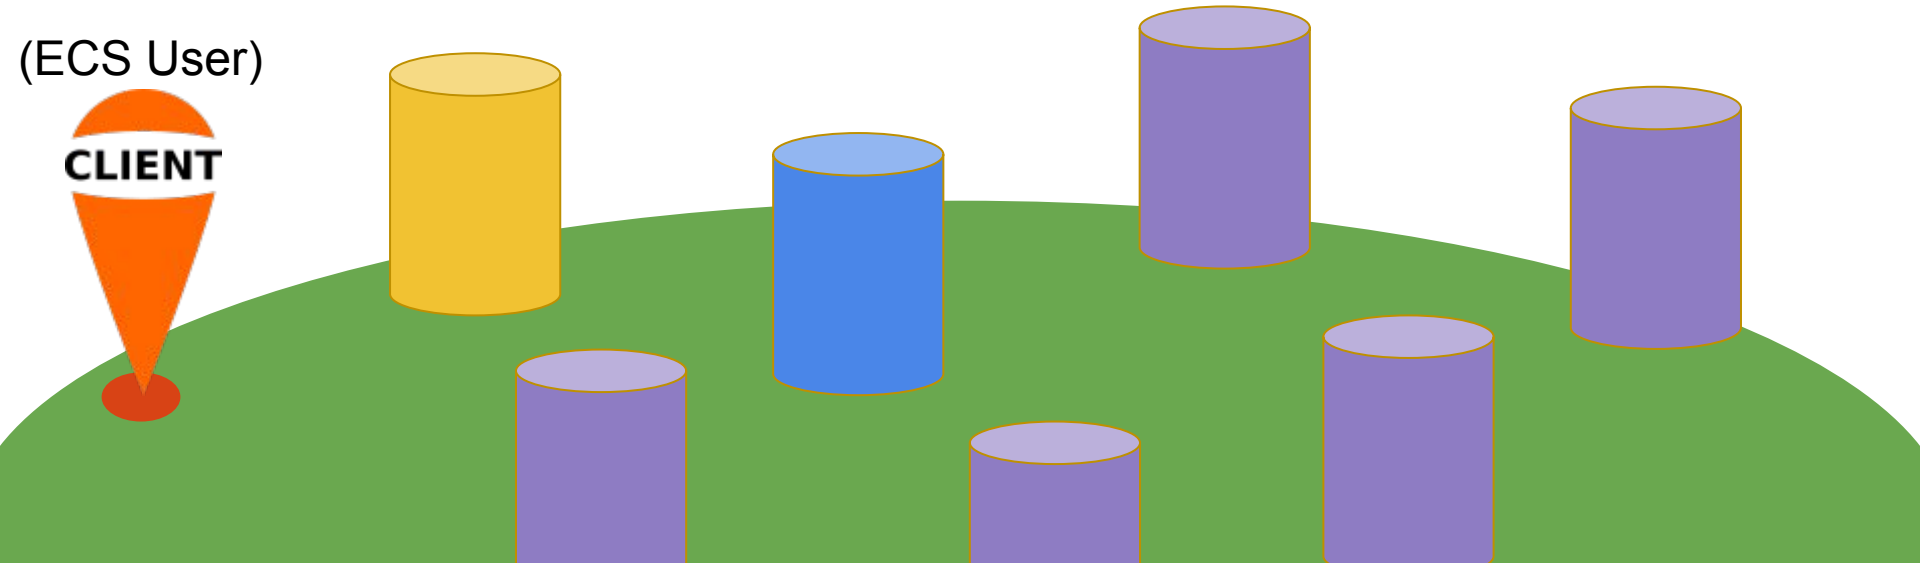

We used ECS:

This still happens

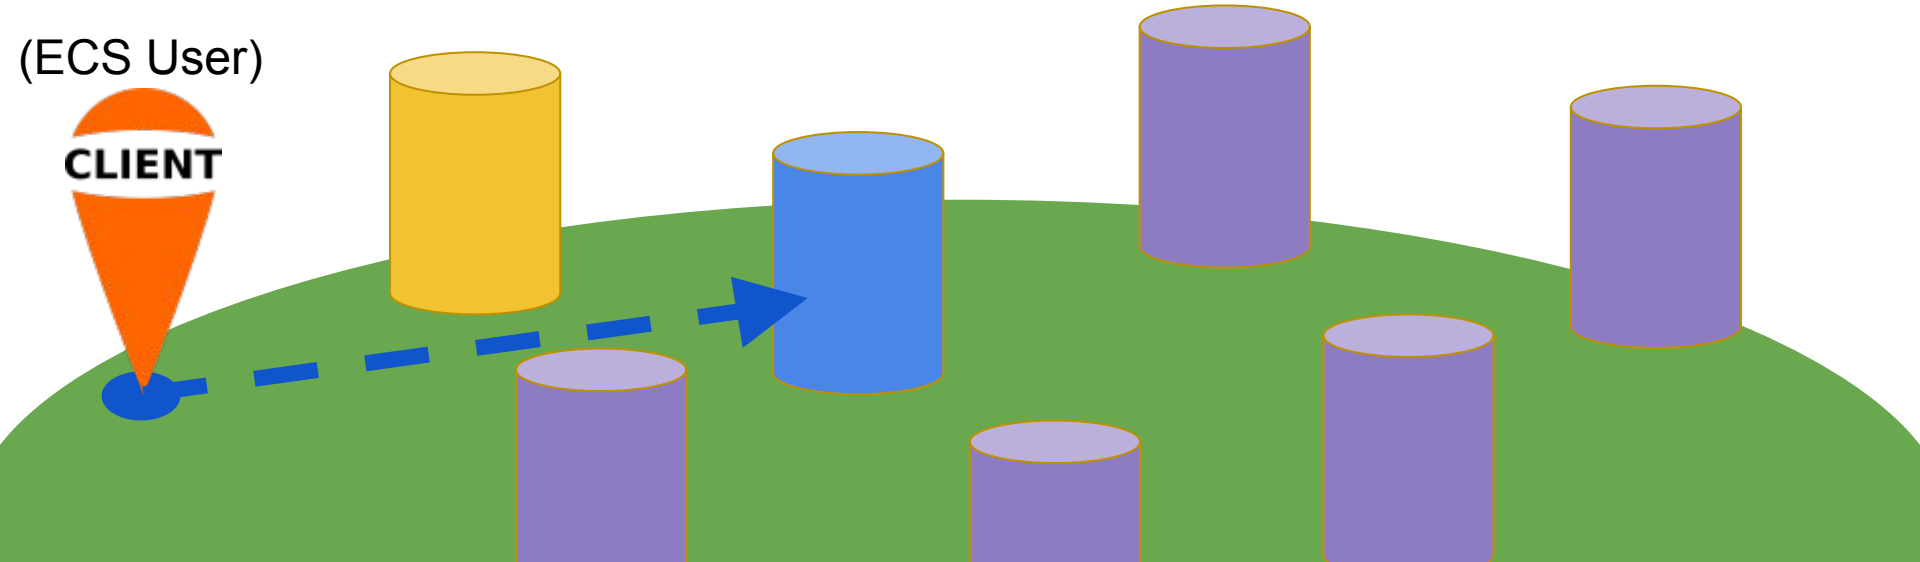

We used ECS:

This still happens

… frequently

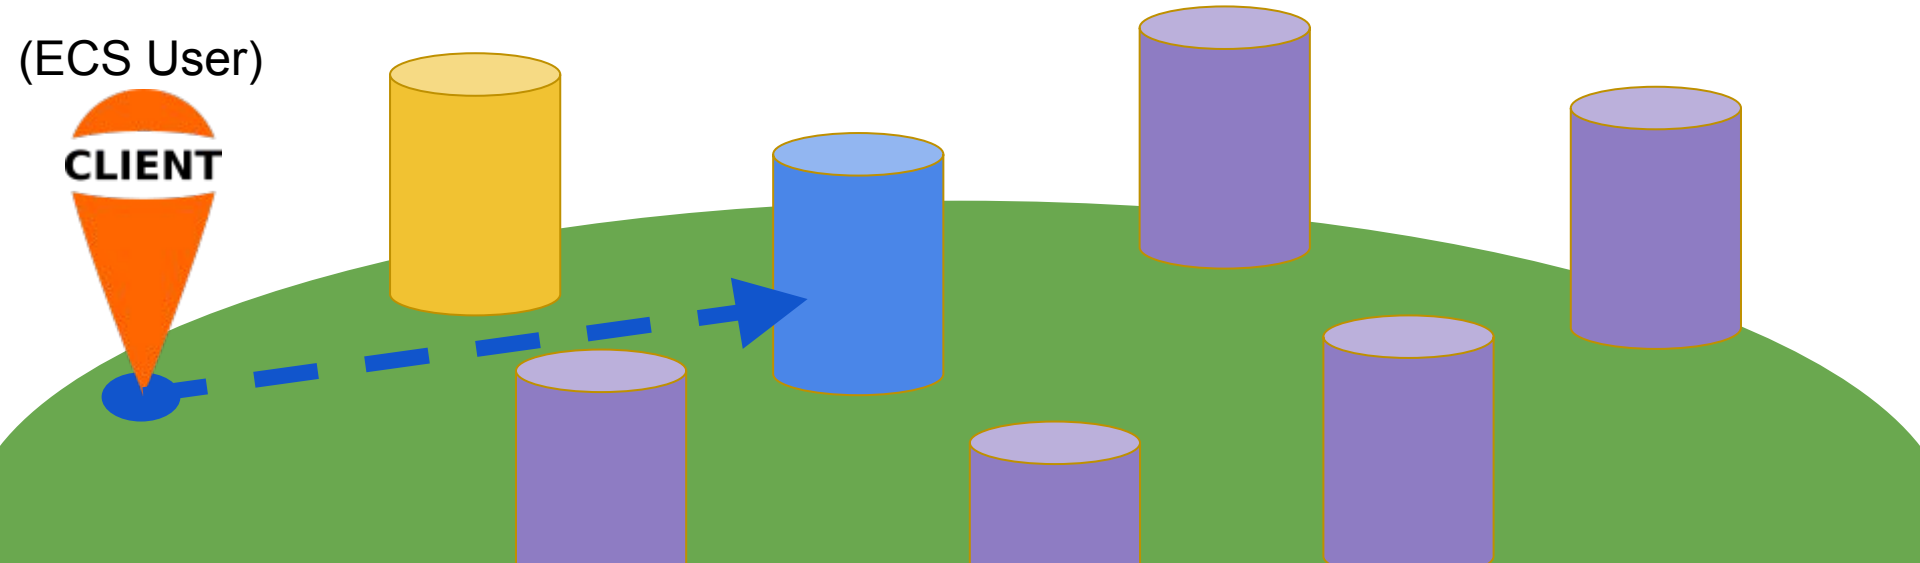

Really? ...

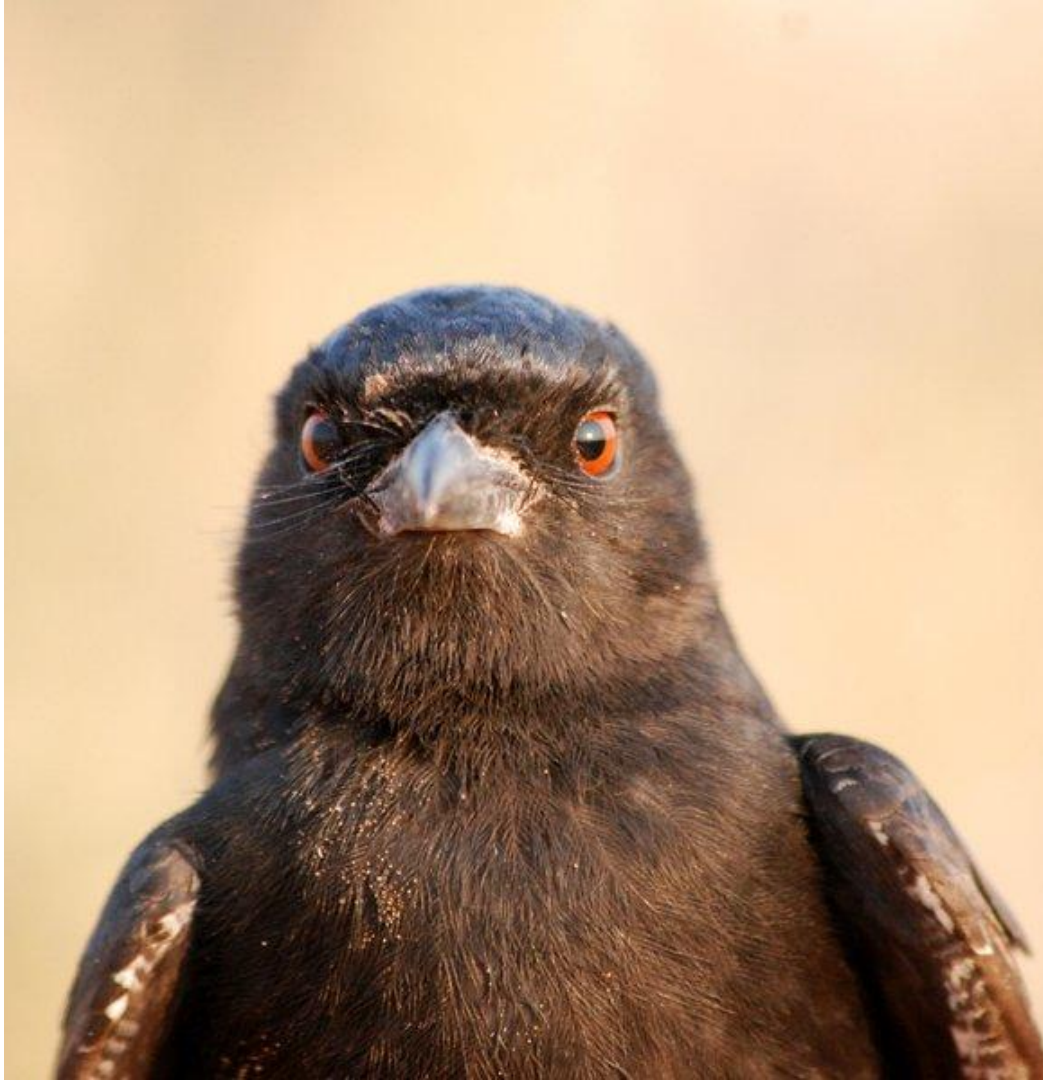

Really? …

YES!
We measured it!

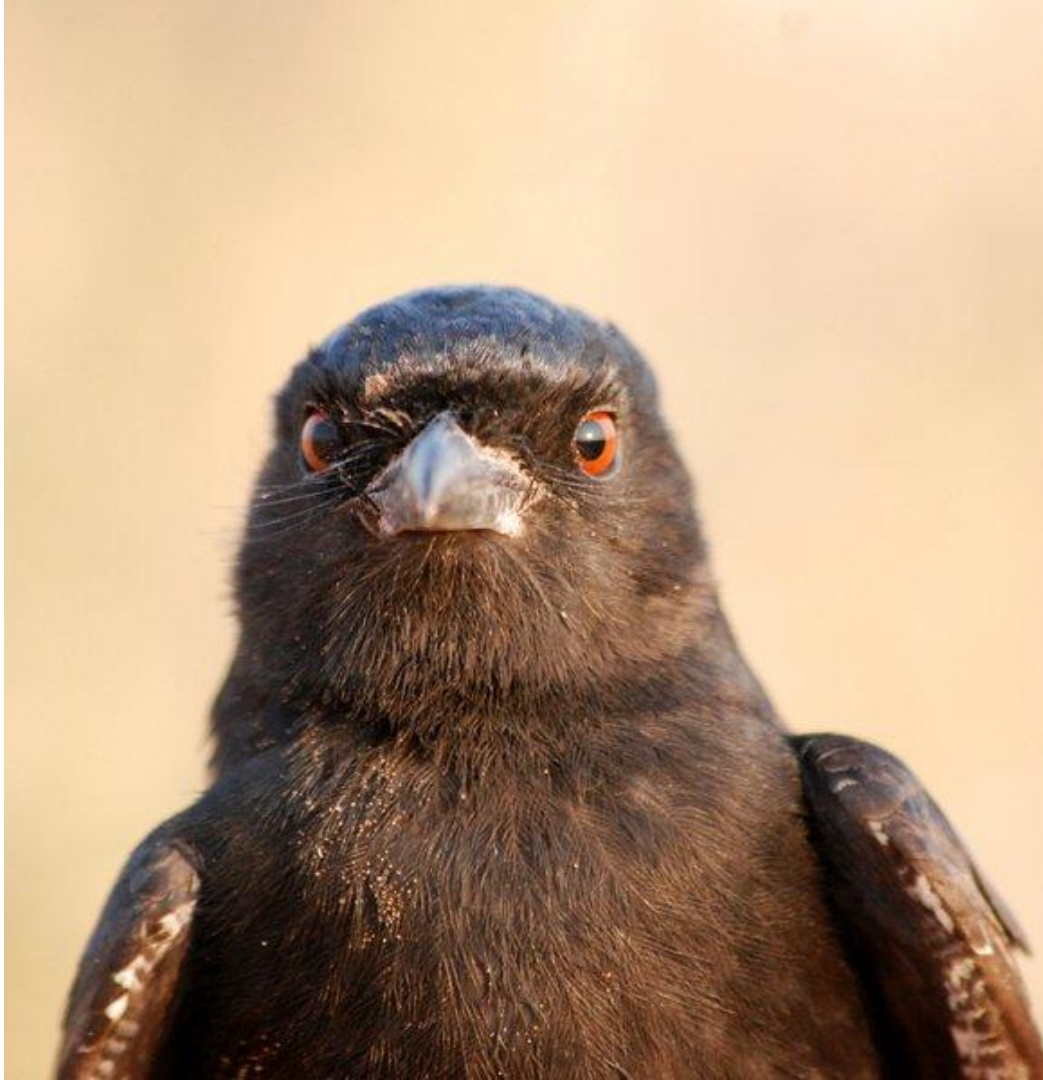

How did we measure it?

How did we measure it?

Find subnets directed to different replicas

Subnet Assimilation

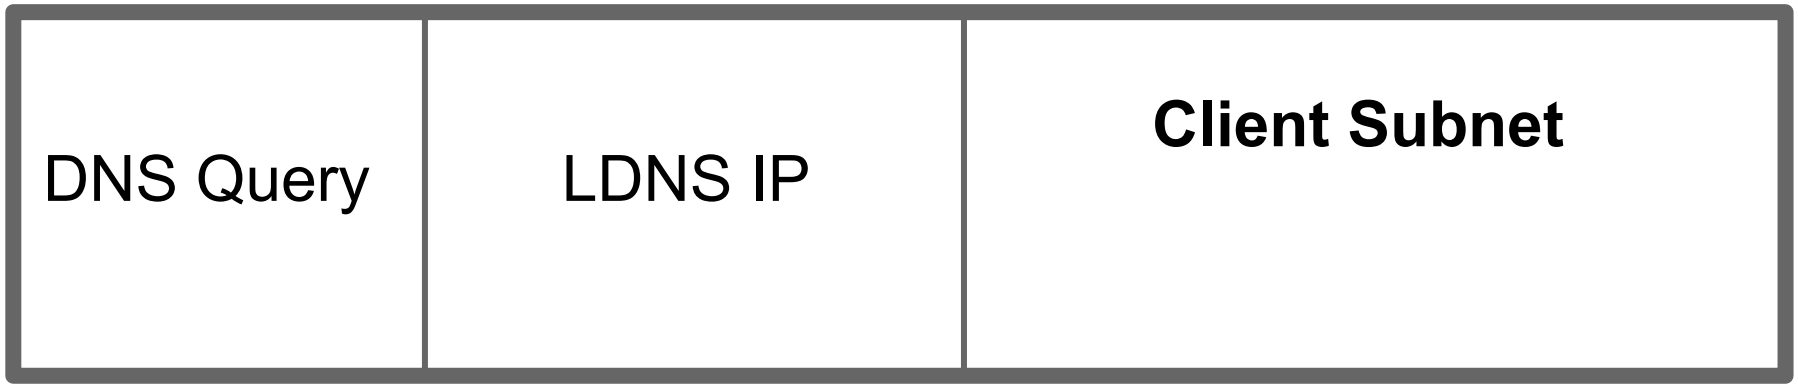

Subnet Assimilation

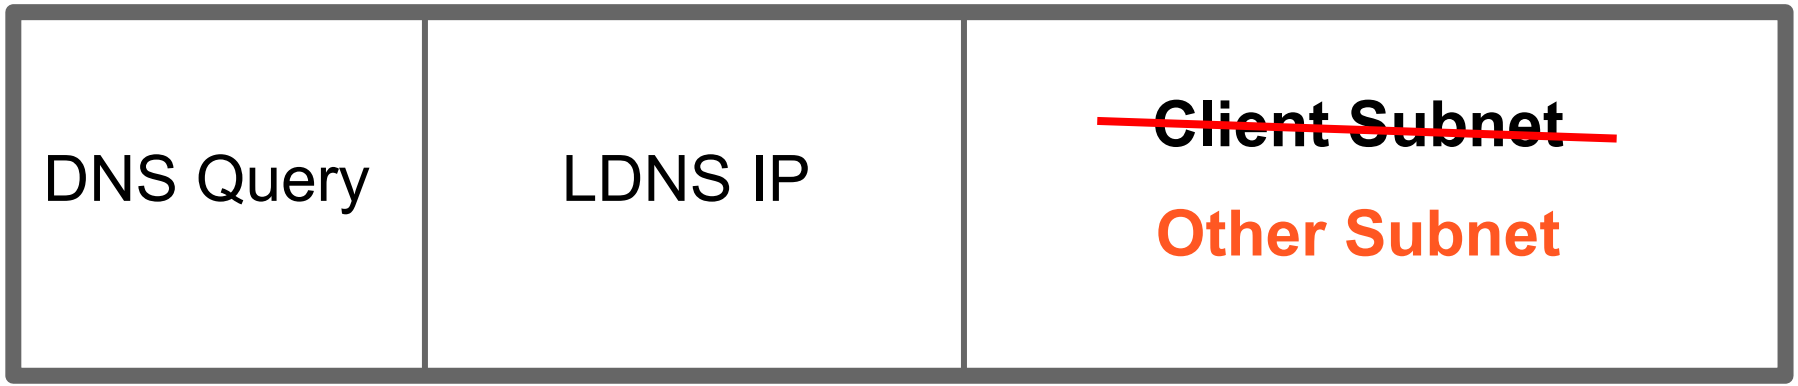

How did we measure it?

Find subnets directed to different replicas

Perform pings and downloads to each replica

How did we measure it?

Find subnets directed to different replicas

Perform pings and downloads to each replica

Identify which subnet resulted in the “best” replica

1. Get "Default" Choice

(use **client's** own subnet for ECS)

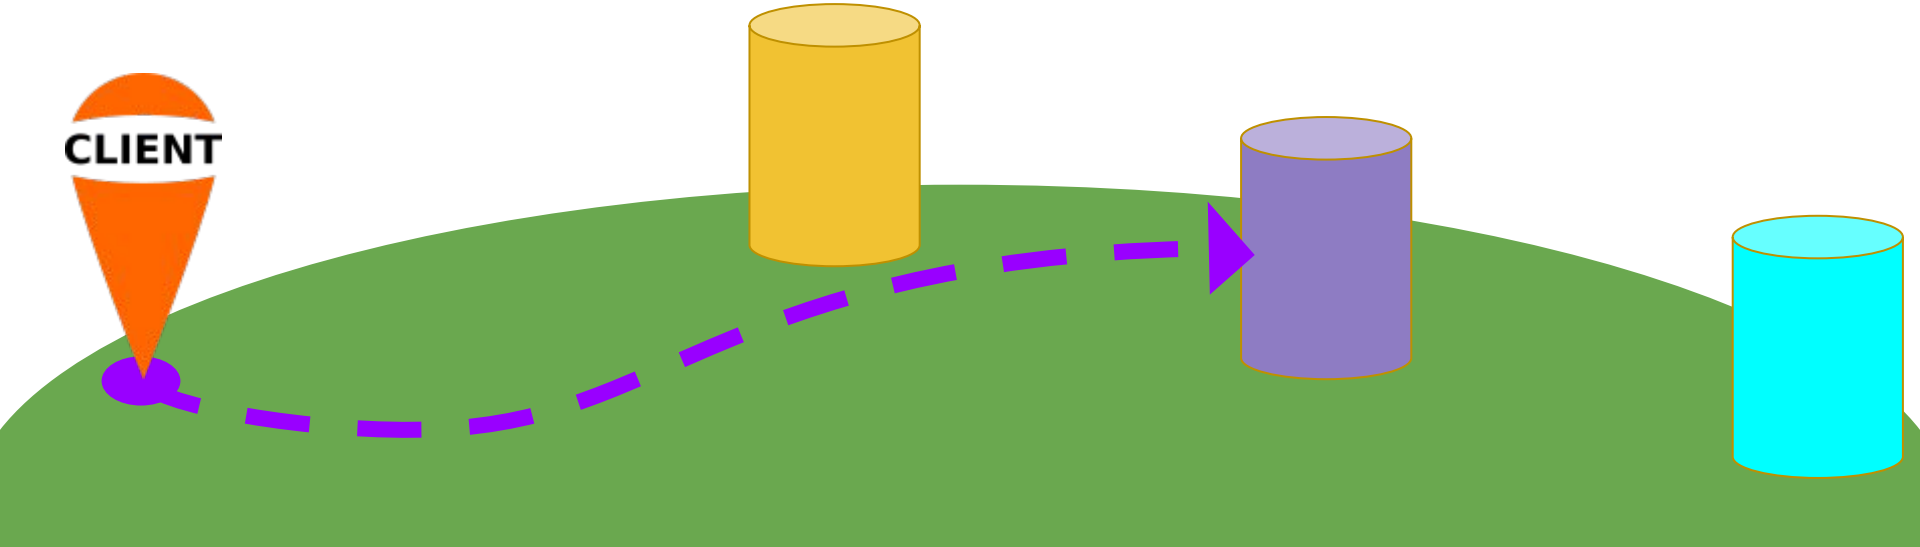

2. Traceroute to default choice

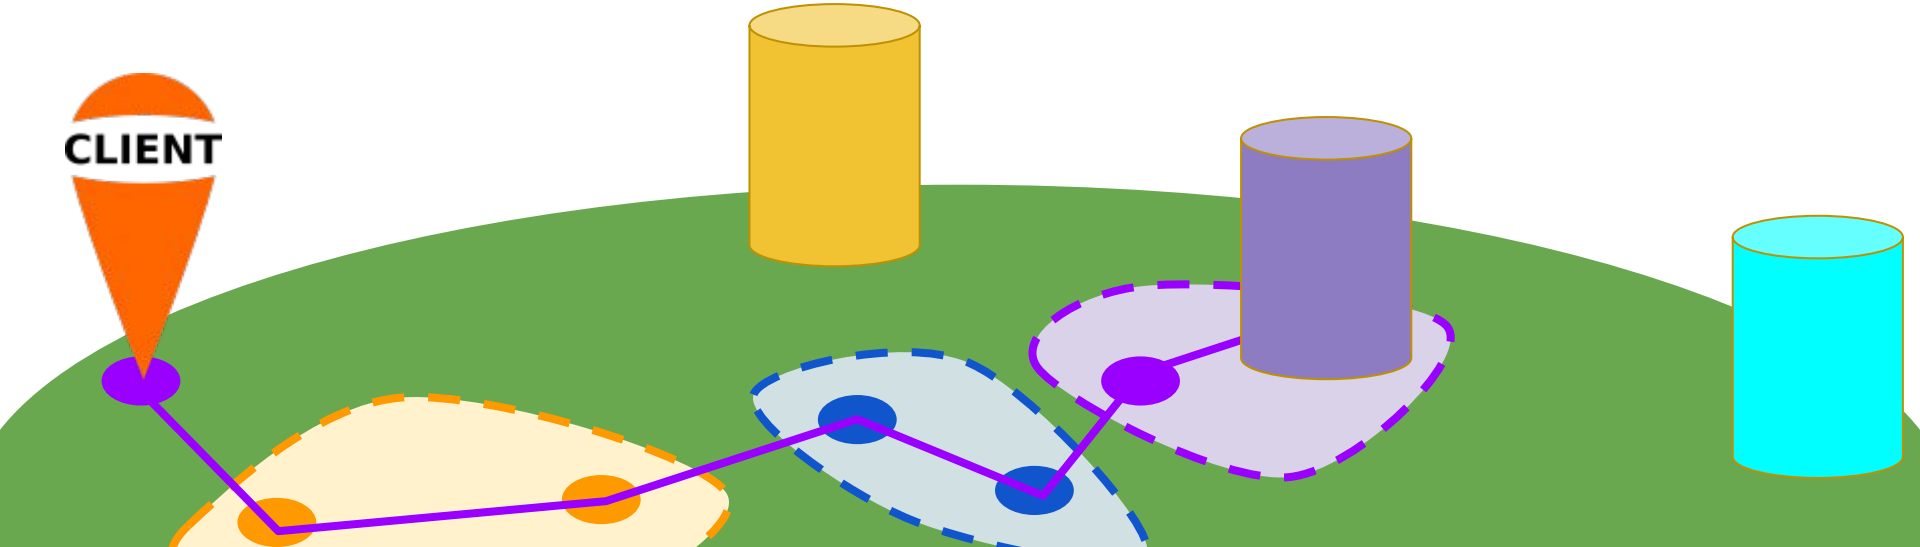

3. Get Hop Subnet Choices

(use **hops'** subnets for ECS)

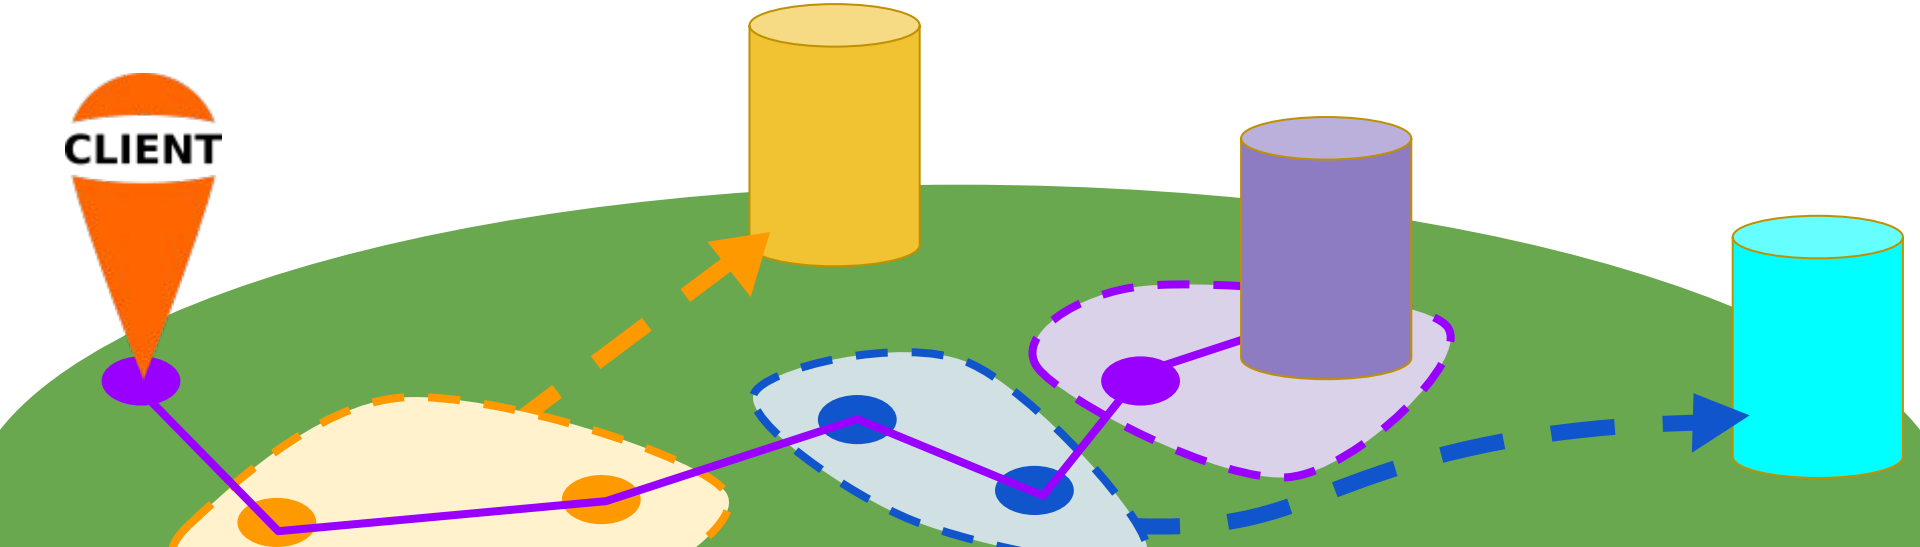

4. Measure Latencies

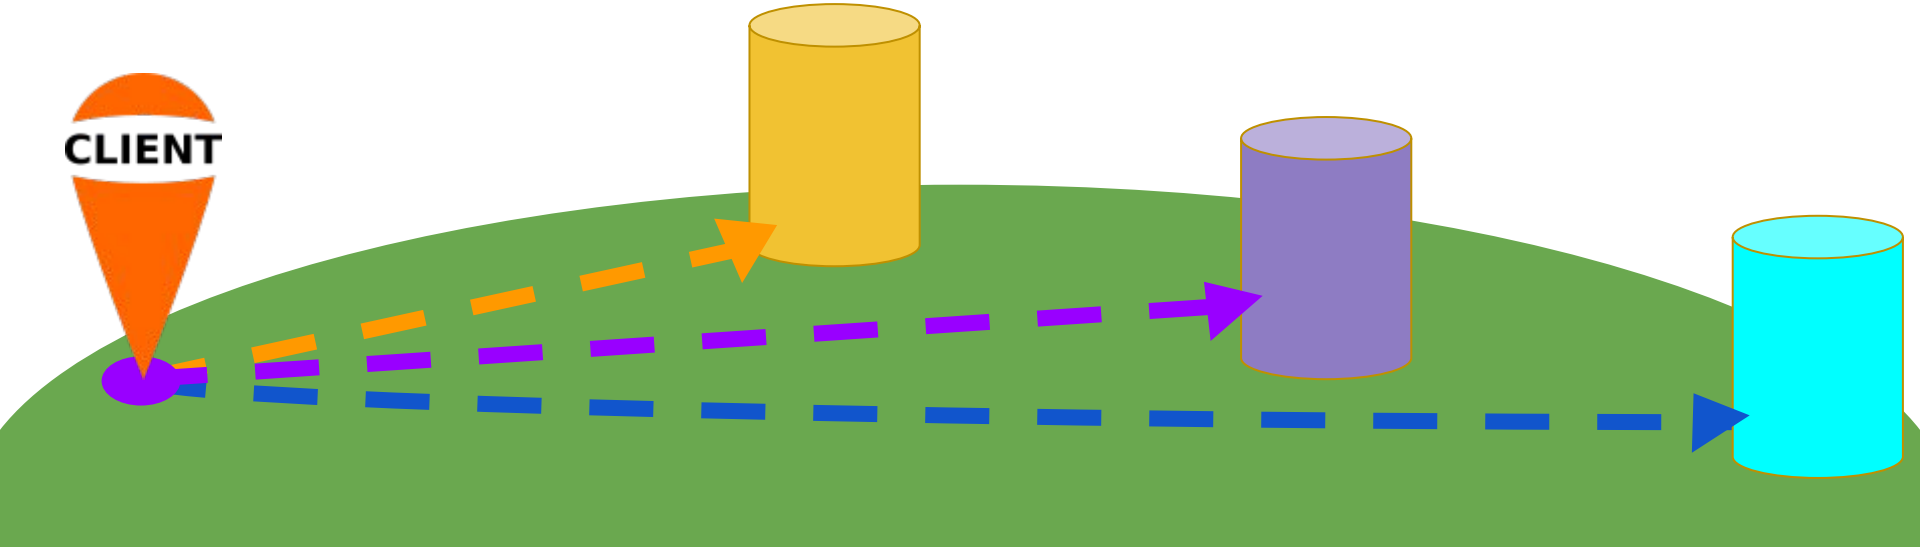

4. Measure Latencies

Steps 1-4: a “trial”

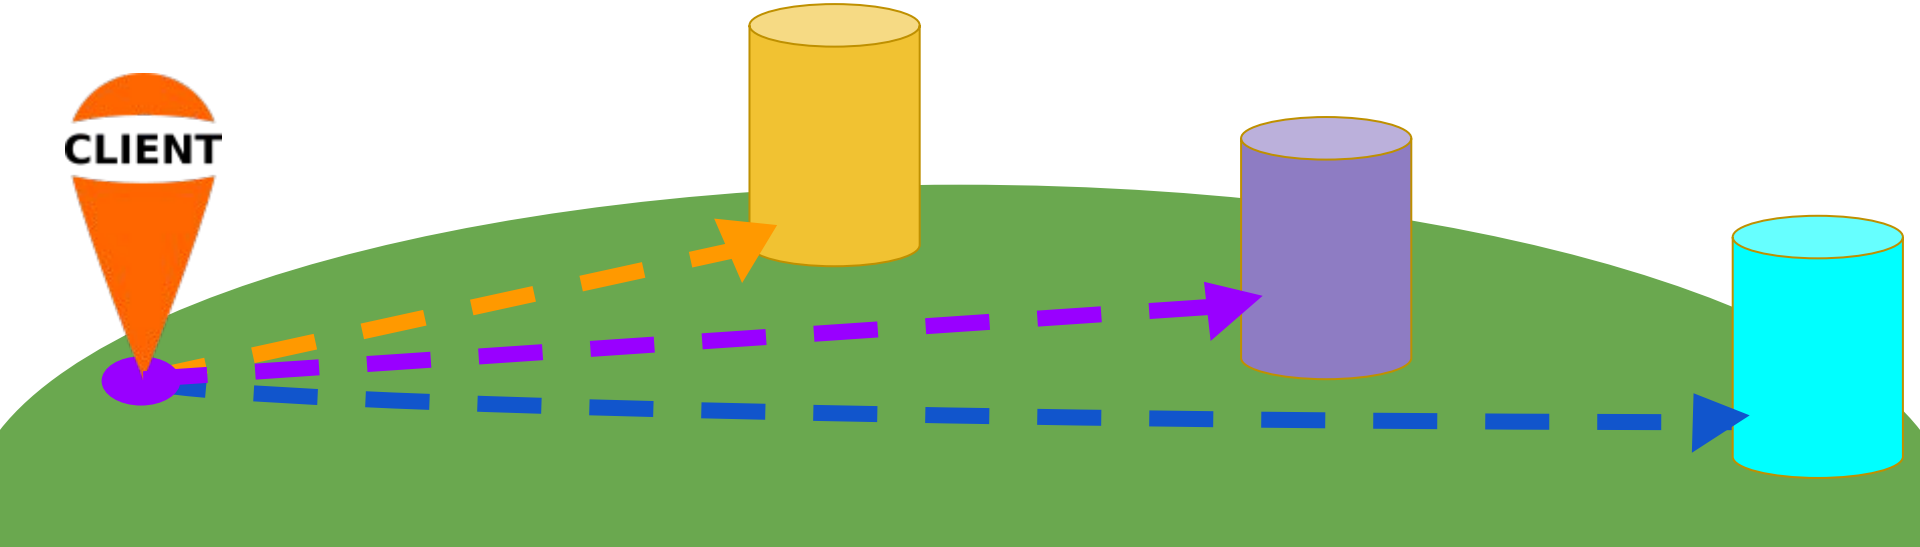

Latency Ratio

Normalize to default choice's RTT

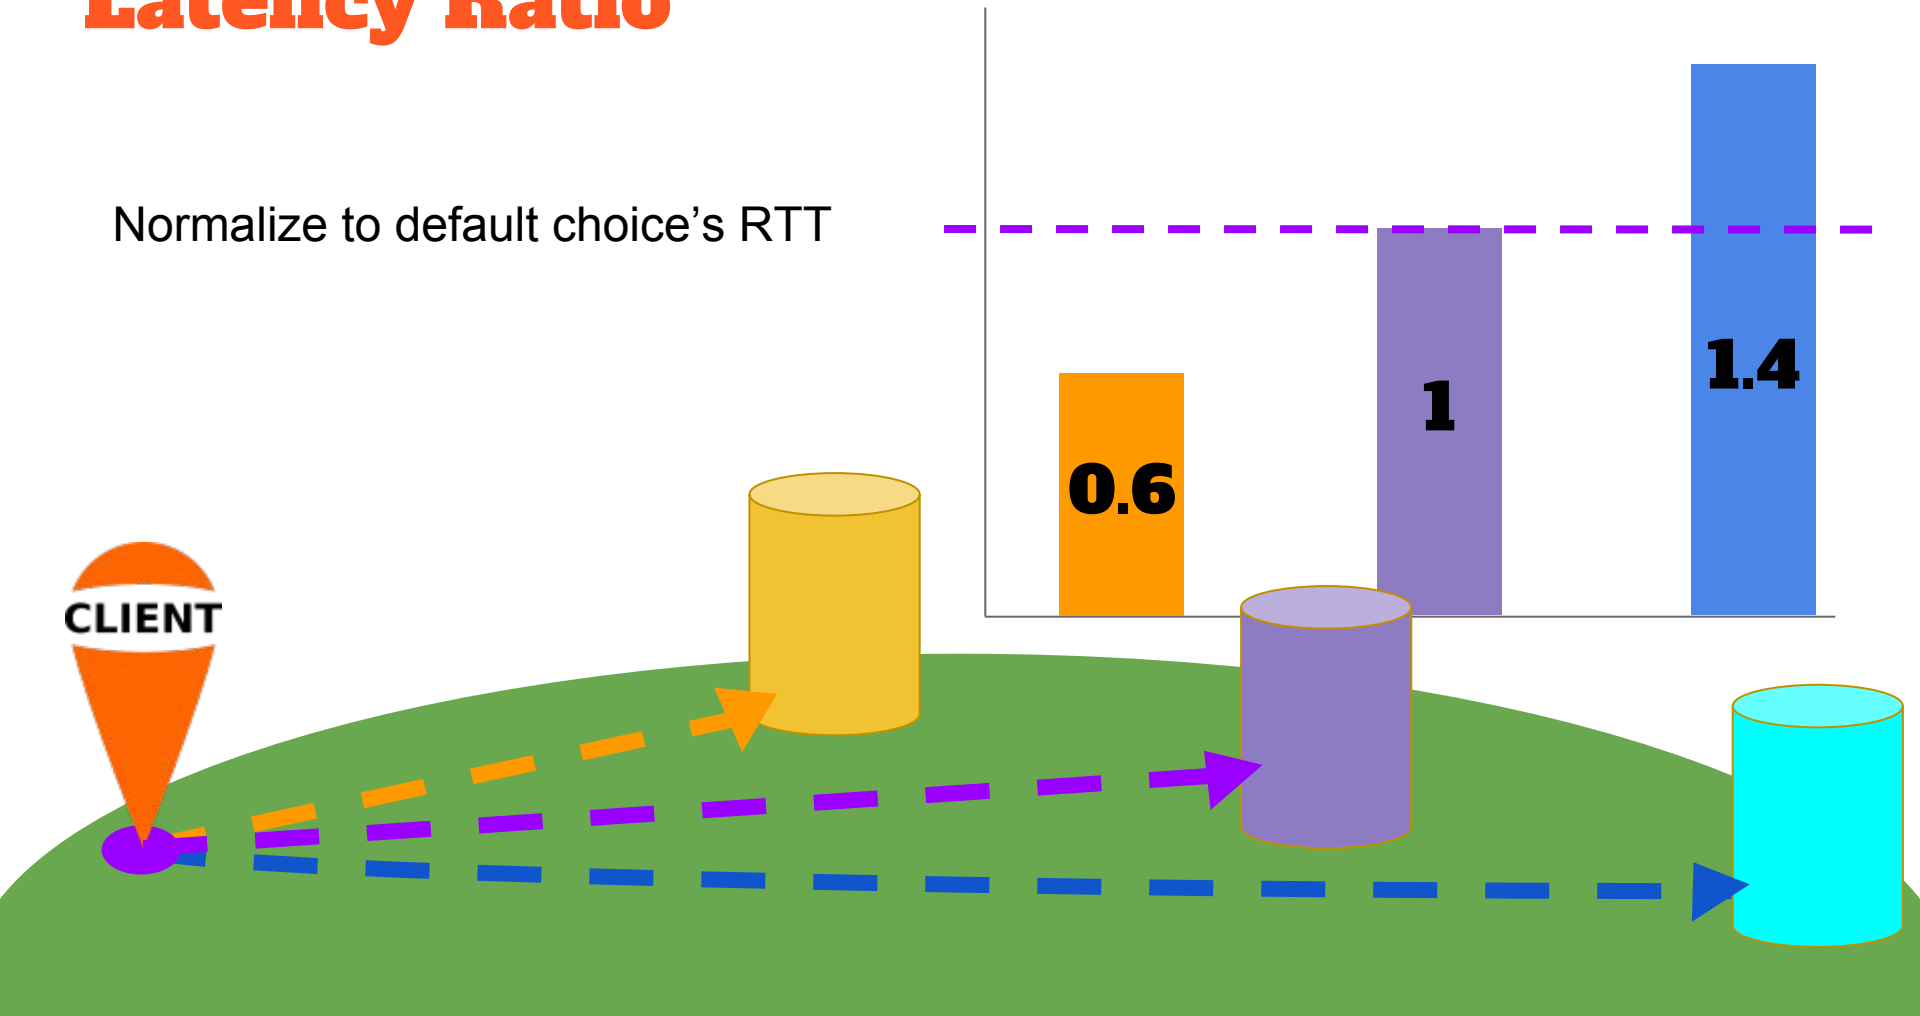

We're looking for this

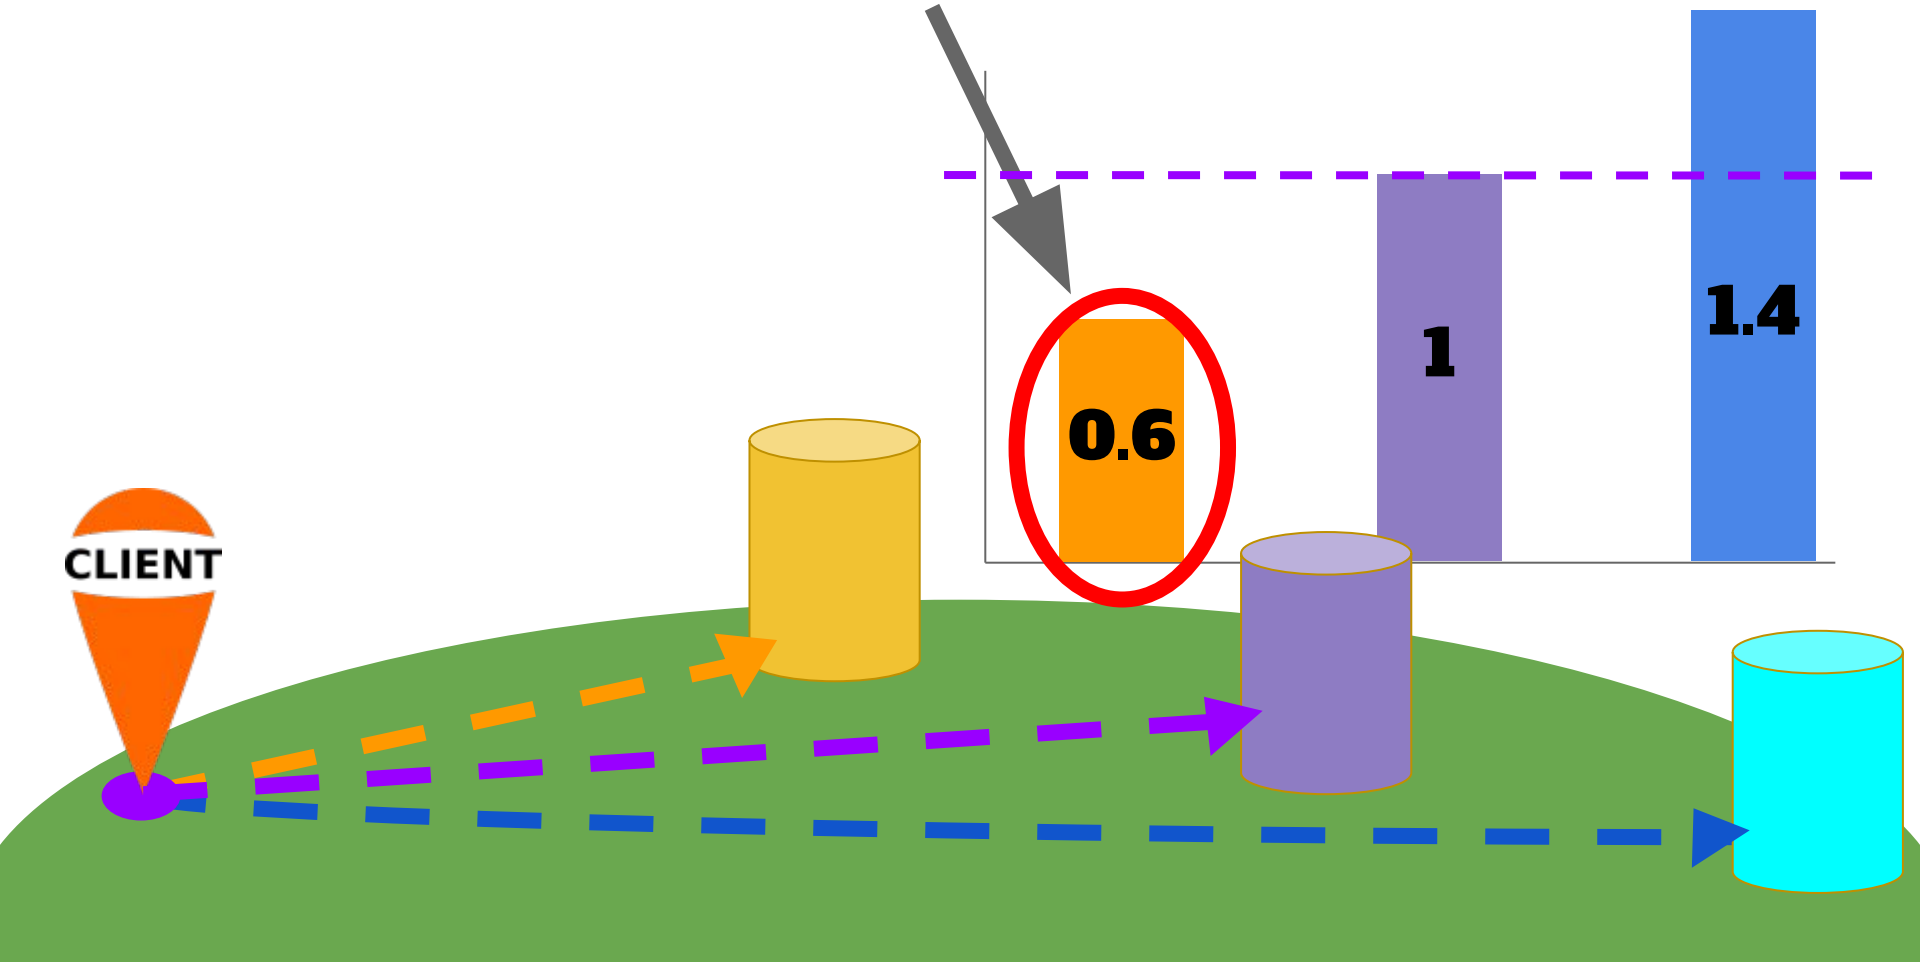

Valley = better choice from hop subnet

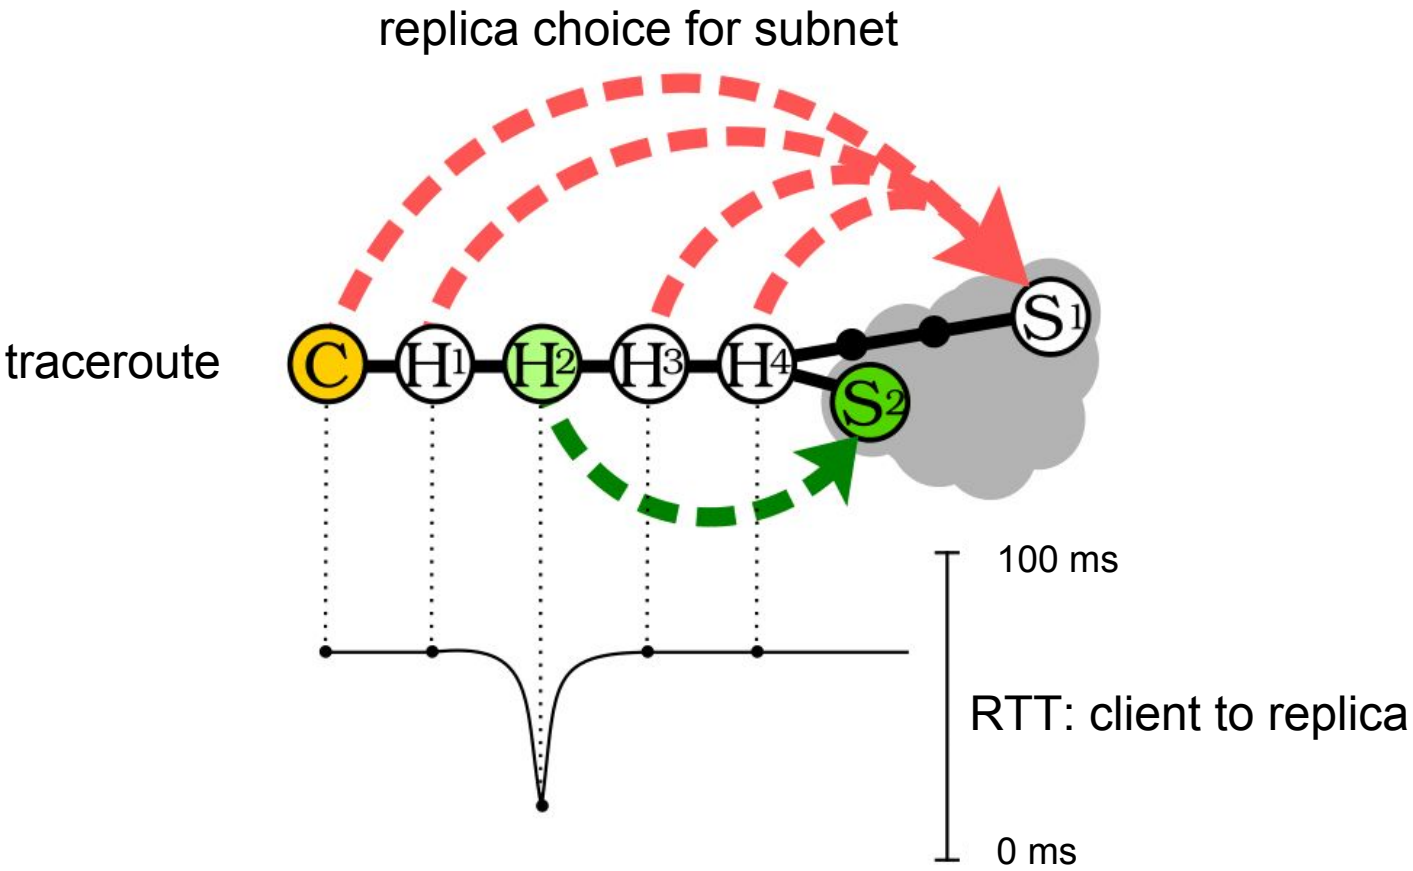

Valley = better choice from hop subnet

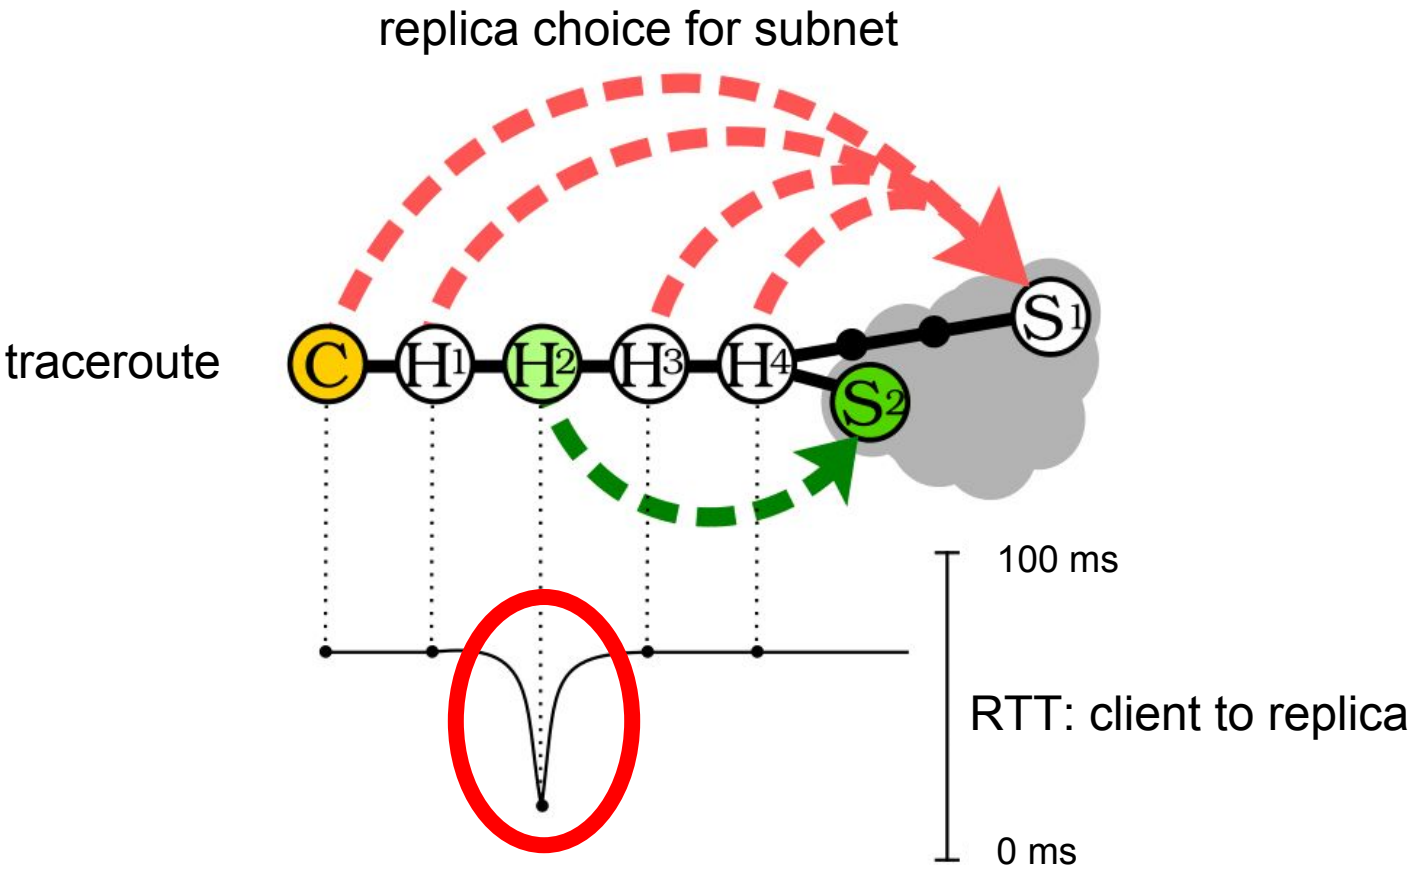

PlanetLab Sees Valleys!

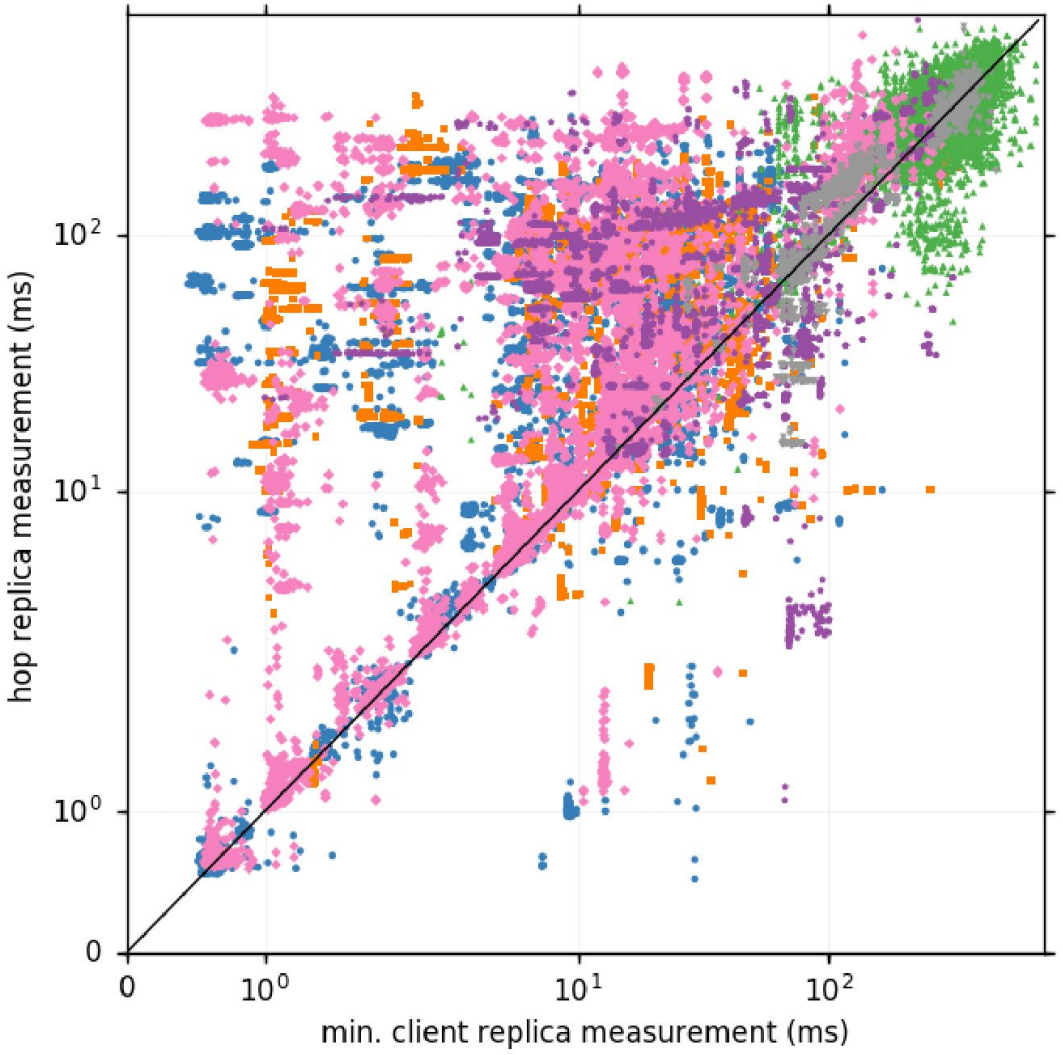

PlanetLab Sees Valleys!

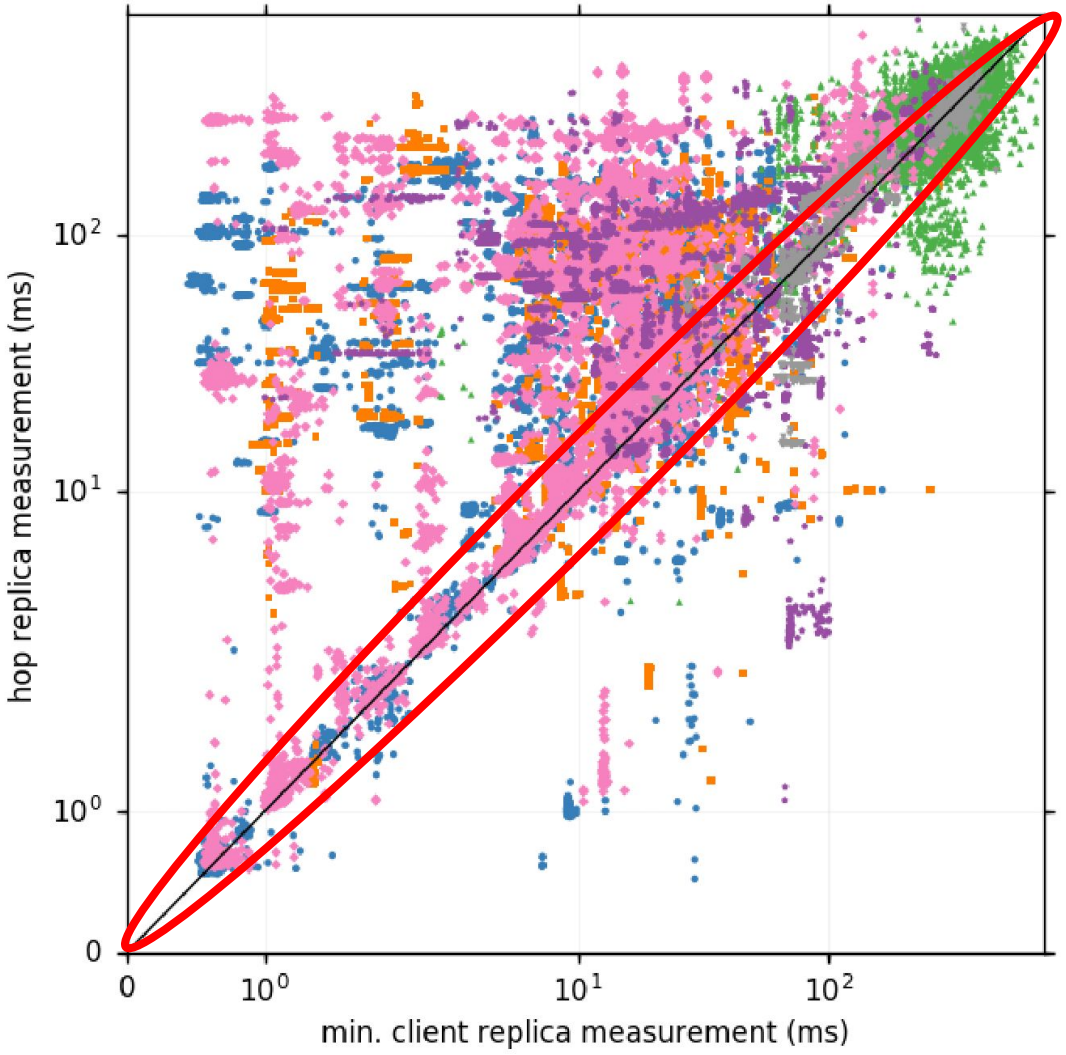

PlanetLab Sees Valleys!

- **Google**: 20.24%
- **Amazon**: 14.02%
- **Alibaba**: 33.68%
- **CDNetworks**: 15.61%
- **ChinaNetCenter**: 27.42%
- **CubeCDN**: 38.58%

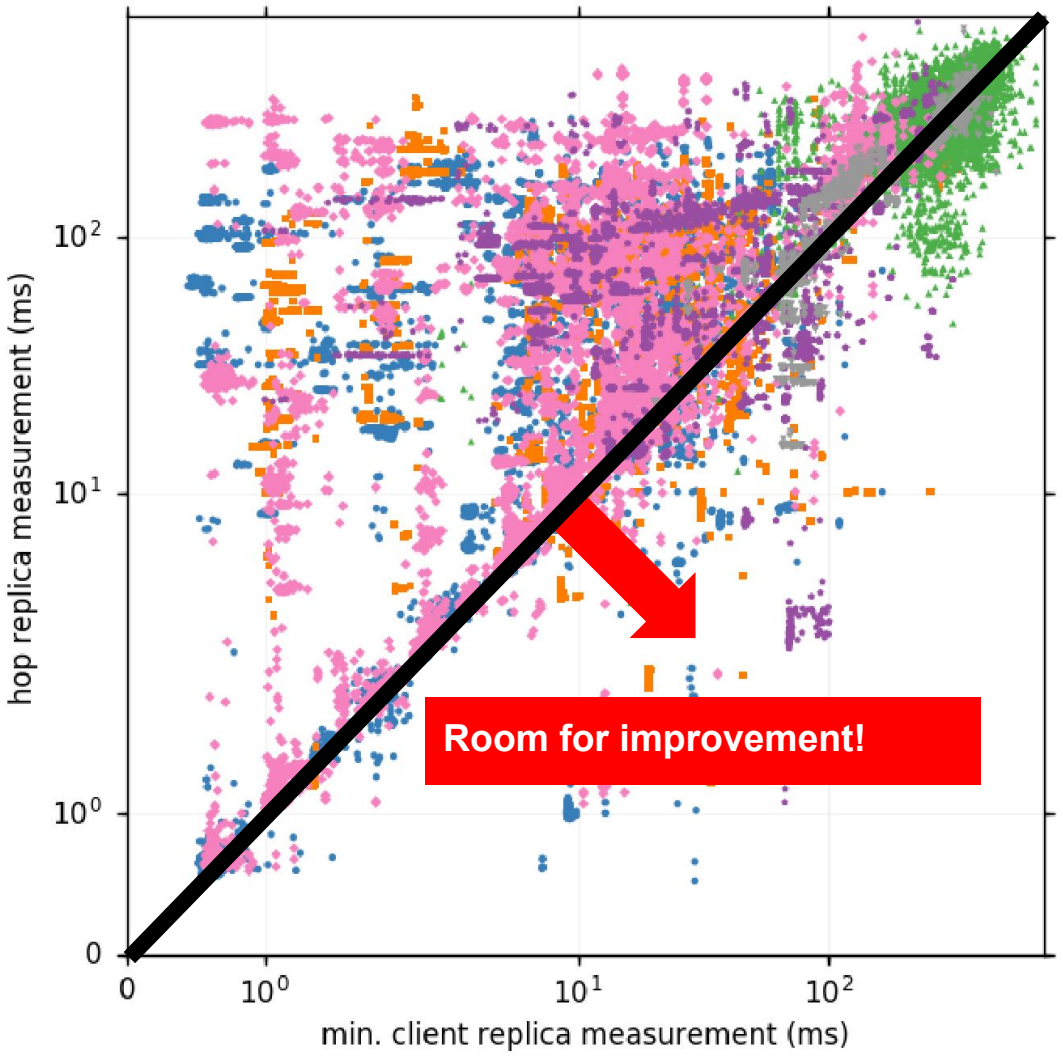

5.

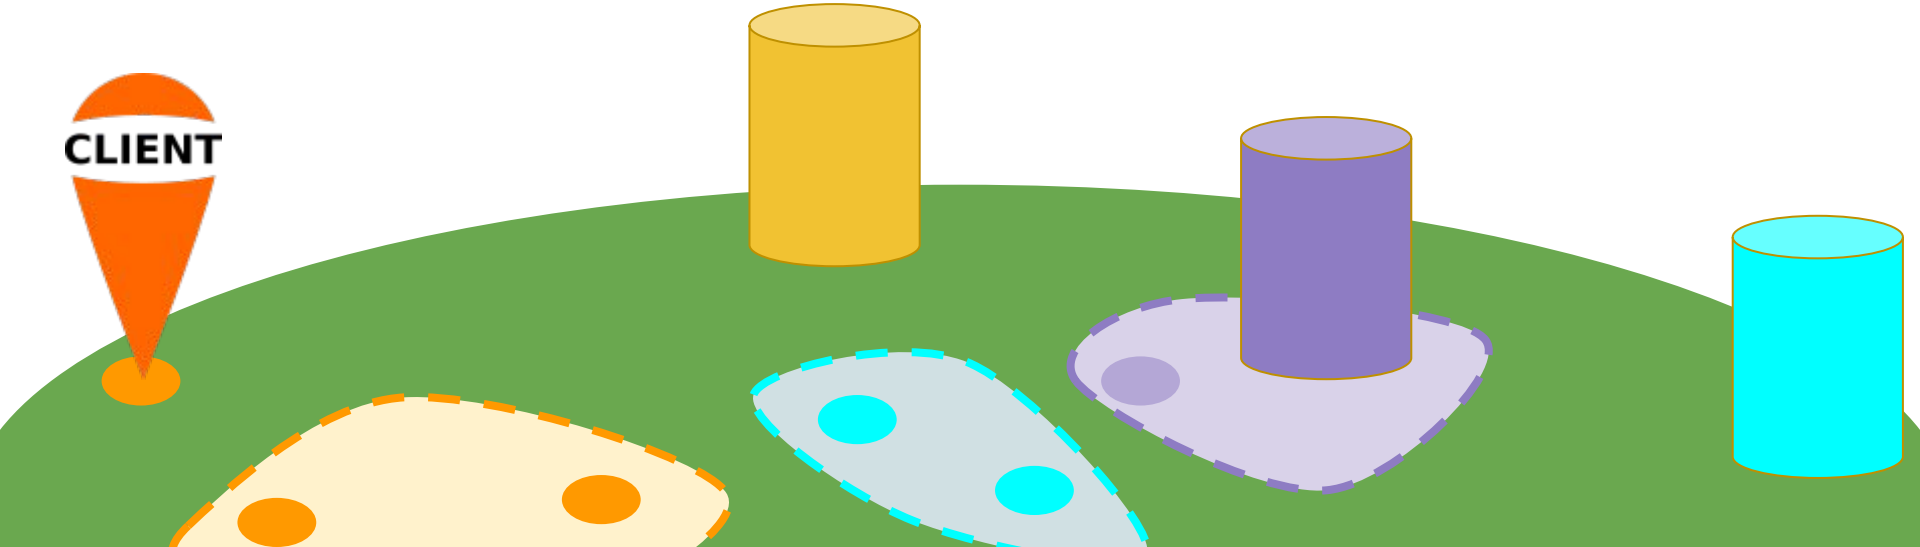

5. Use best subnet for ECS

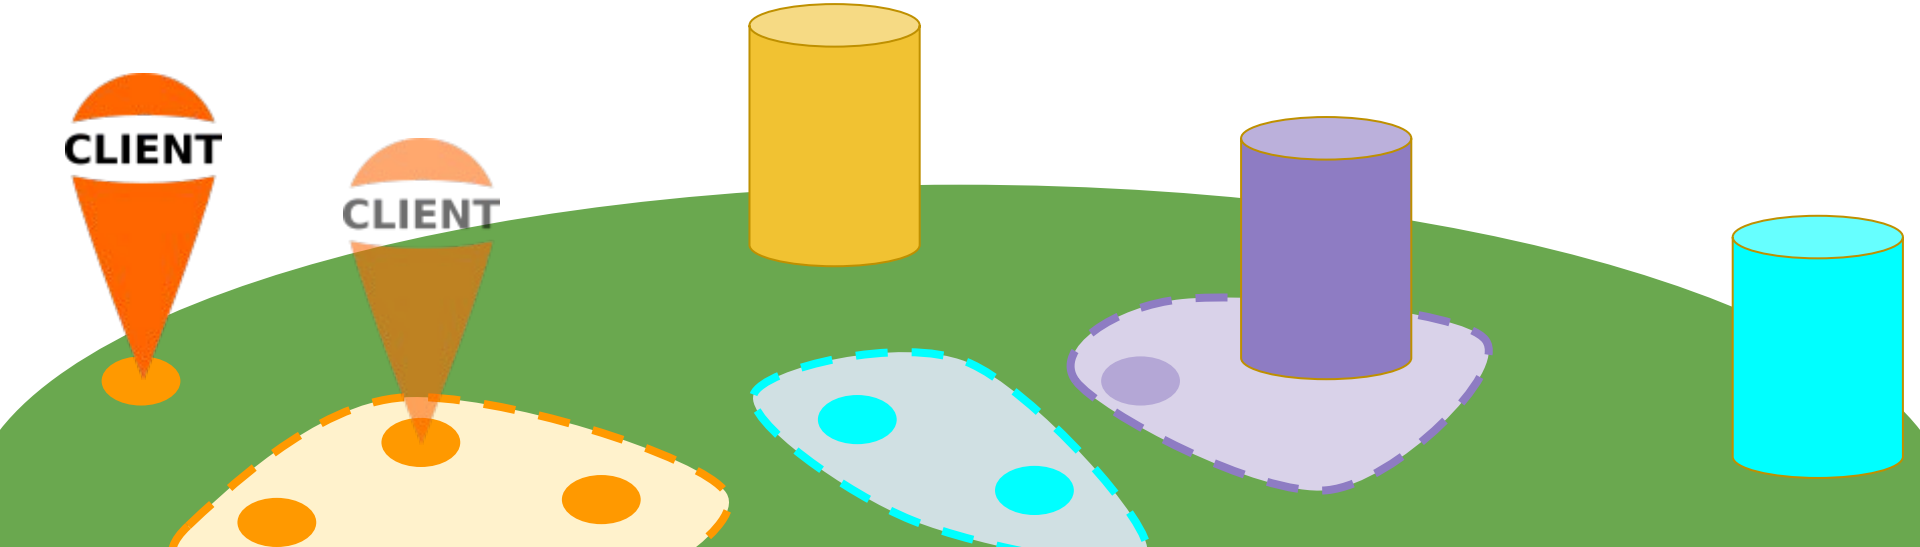

5. Use best subnet for ECS

Get best mapping!

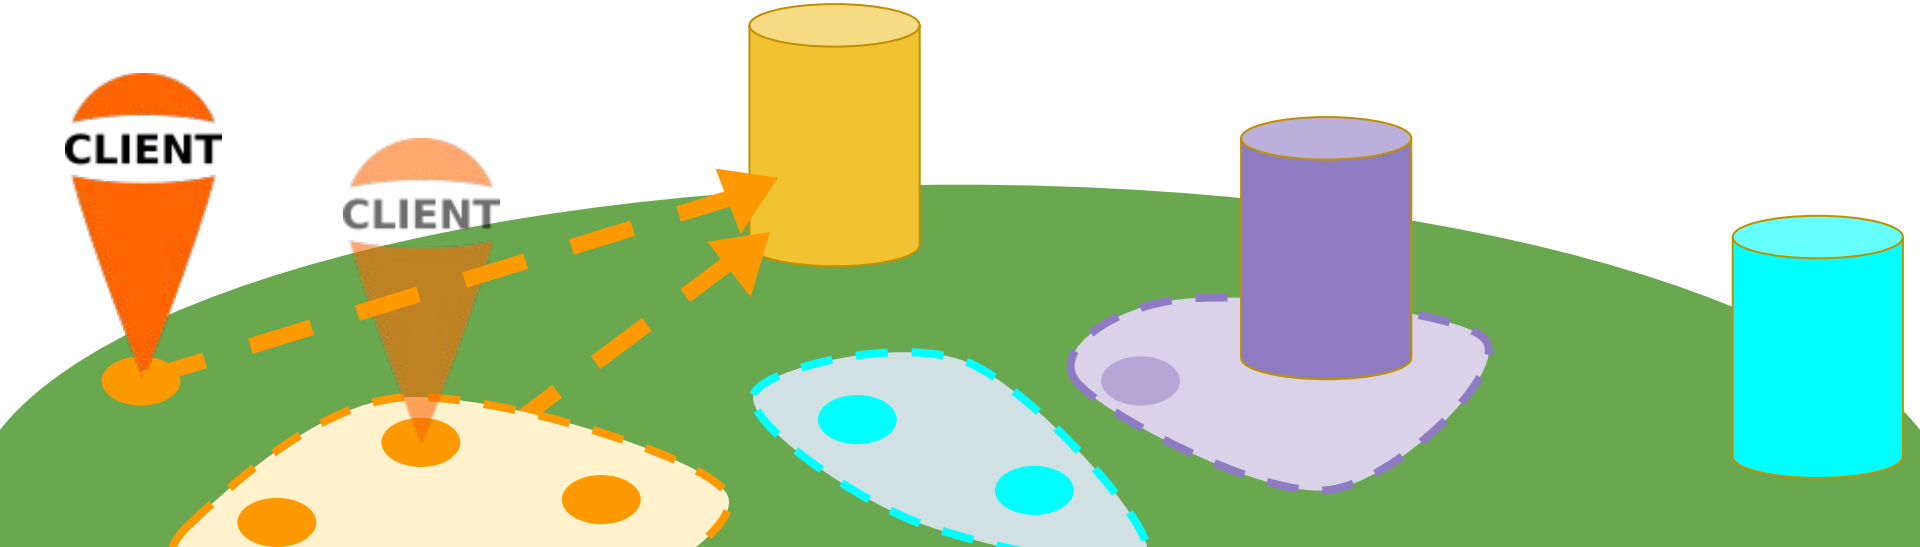

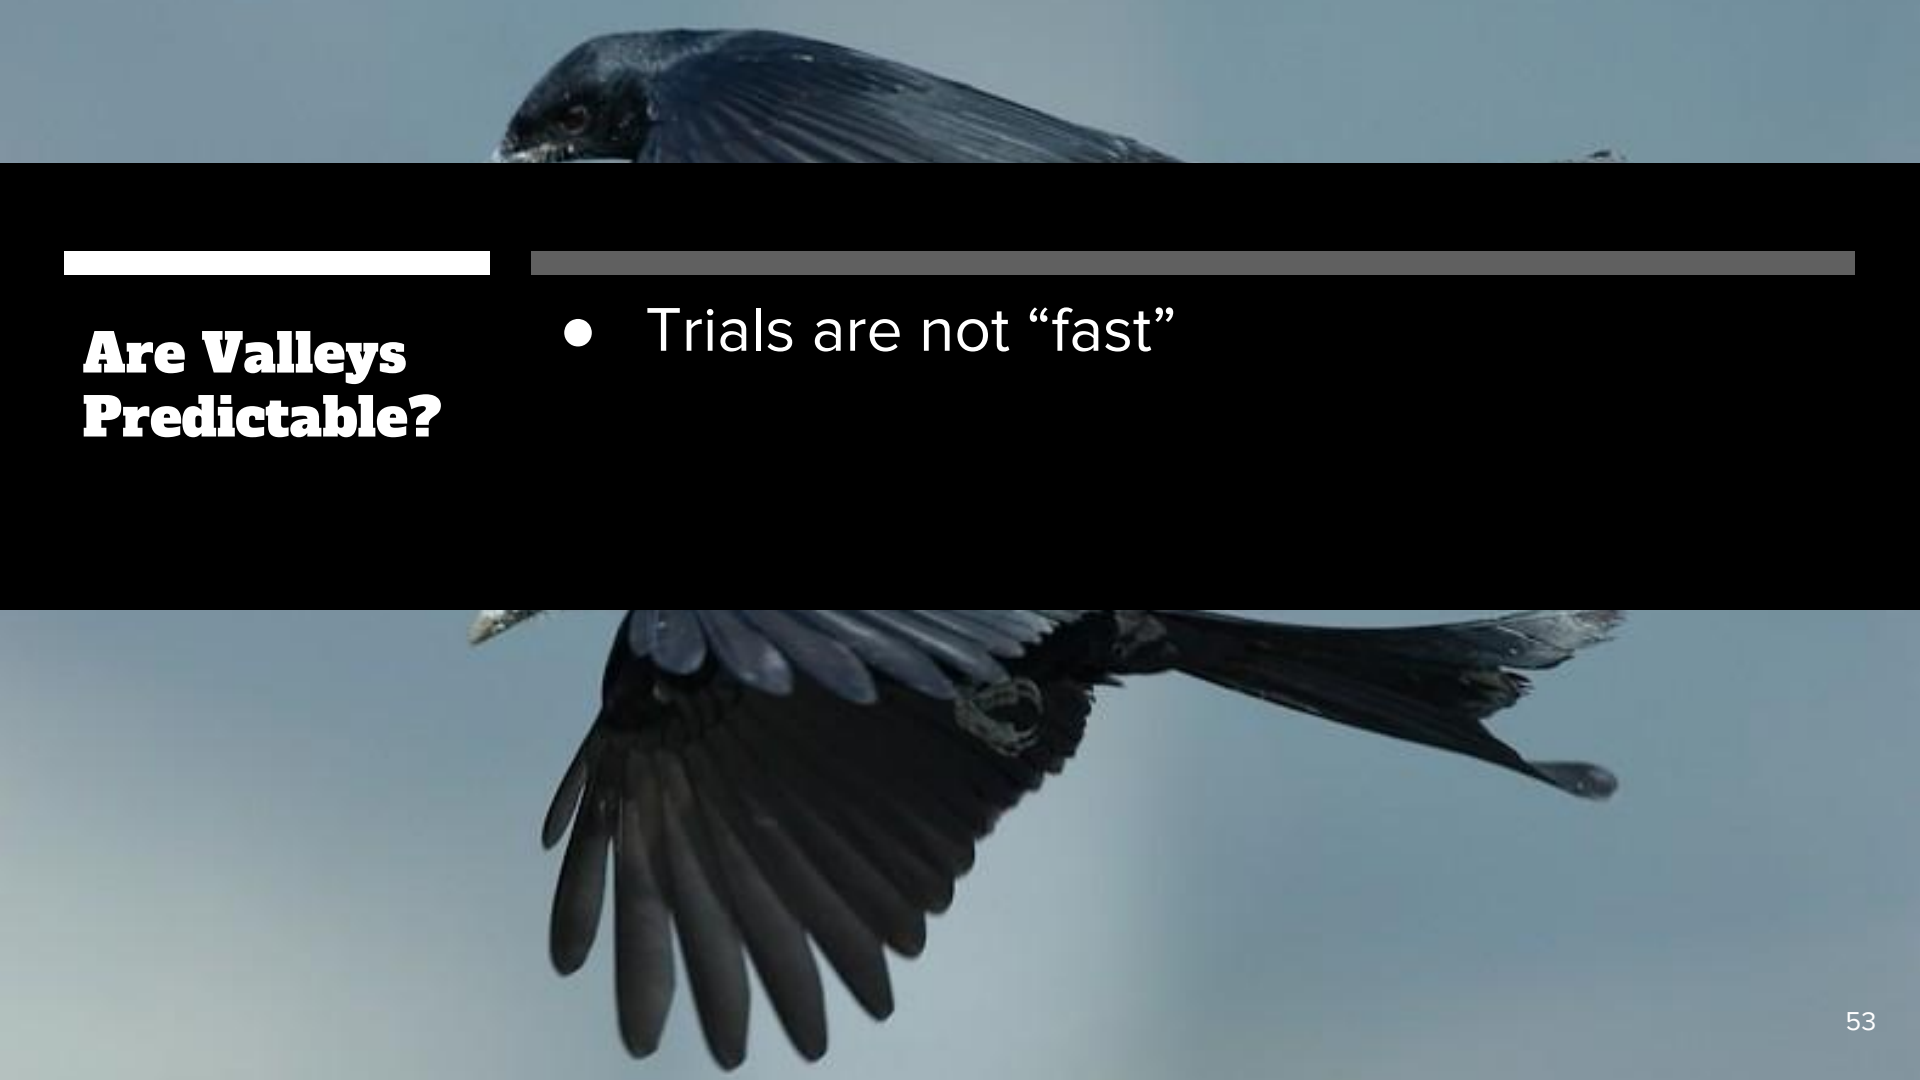

Are Valleys Predictable?

- Trials are not “fast”

Are Valleys Predictable?

- Trials are not “fast”
- We want valleys “on the fly”

Are Valleys Predictable?

- Trials are not “fast”
- We want valleys “on the fly”
- We **need** to find valley-prone subnets

Testing Persistence

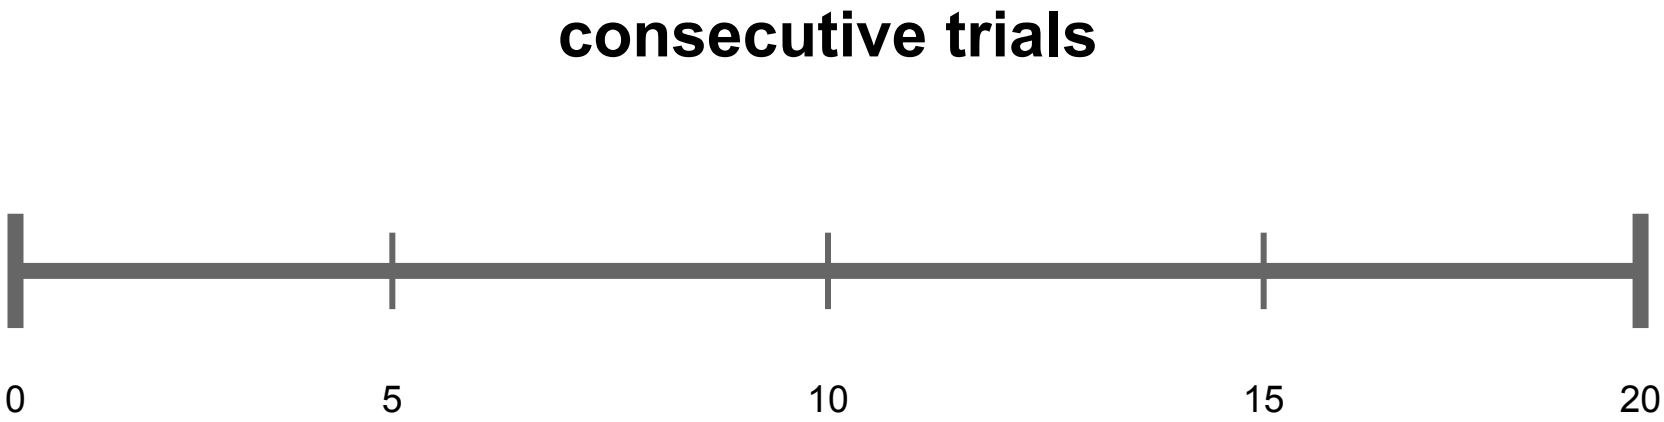

Testing Persistence

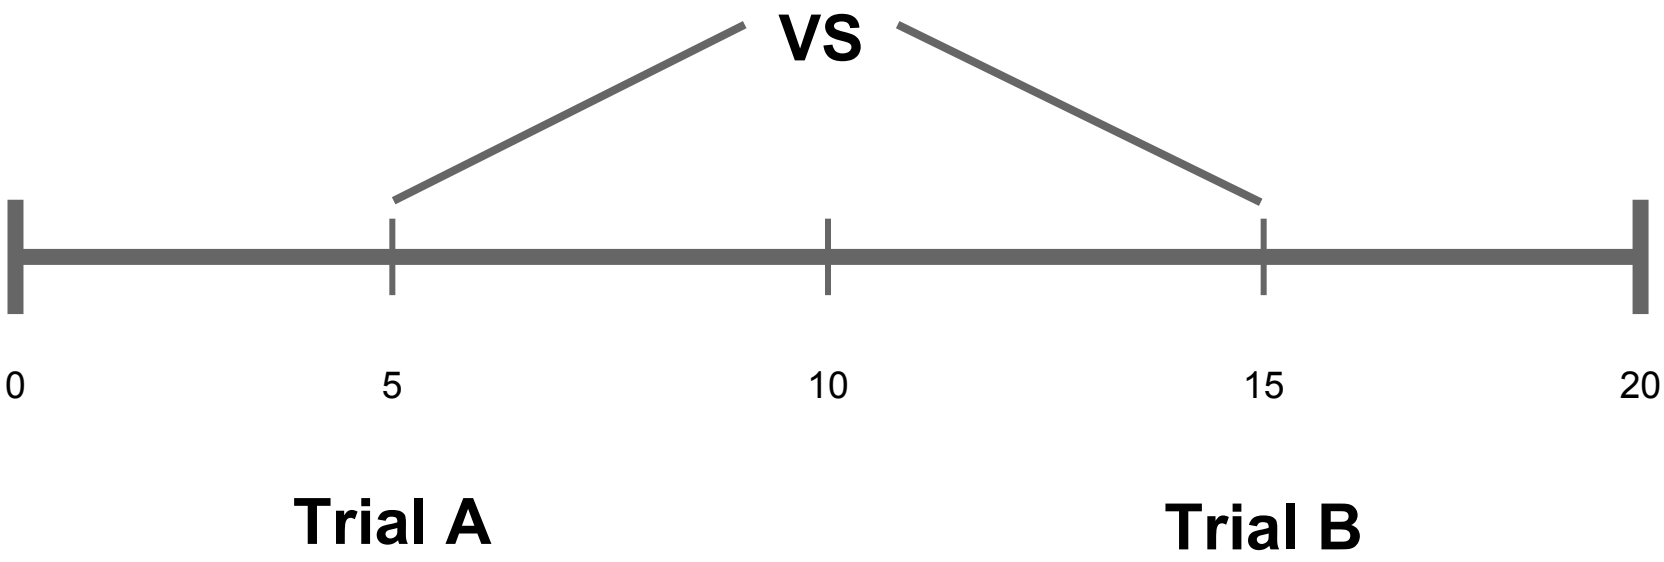

Latency Ratio Difference Over Time

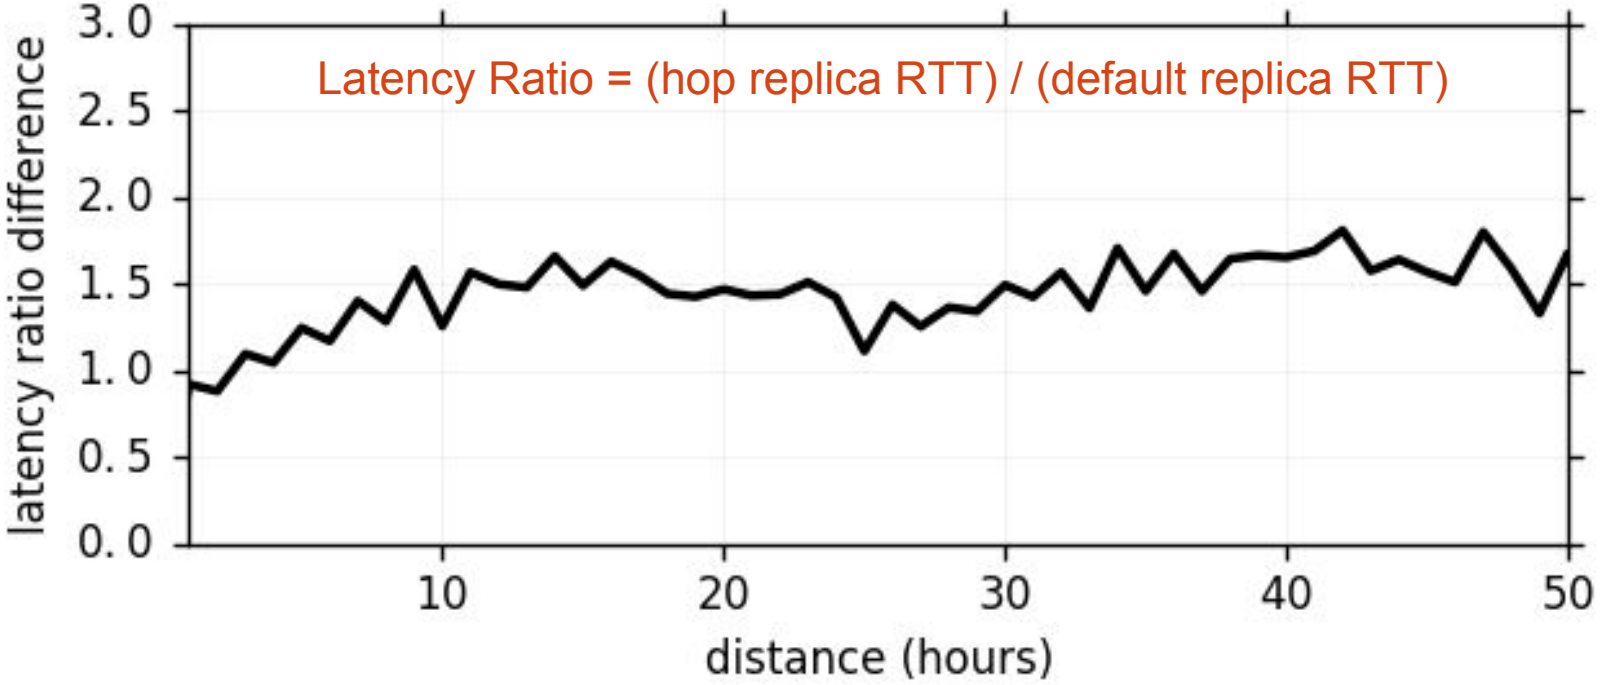

Testing Persistence

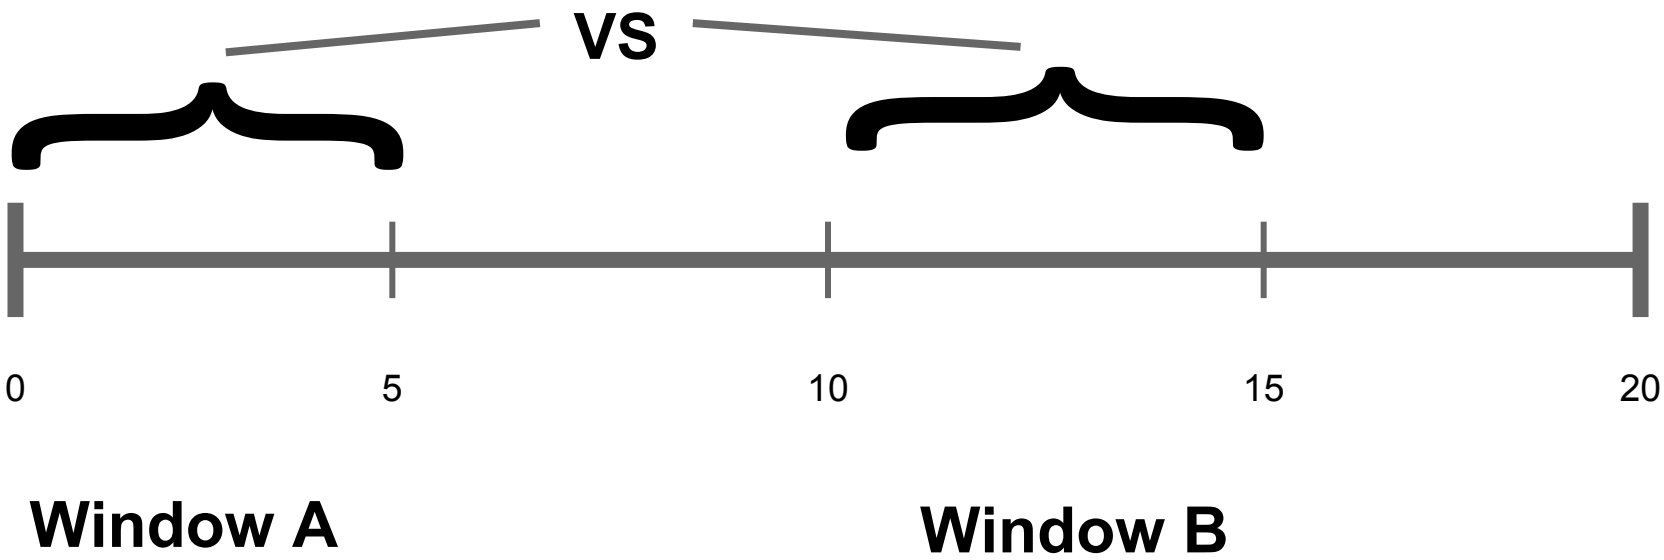

Testing Persistence

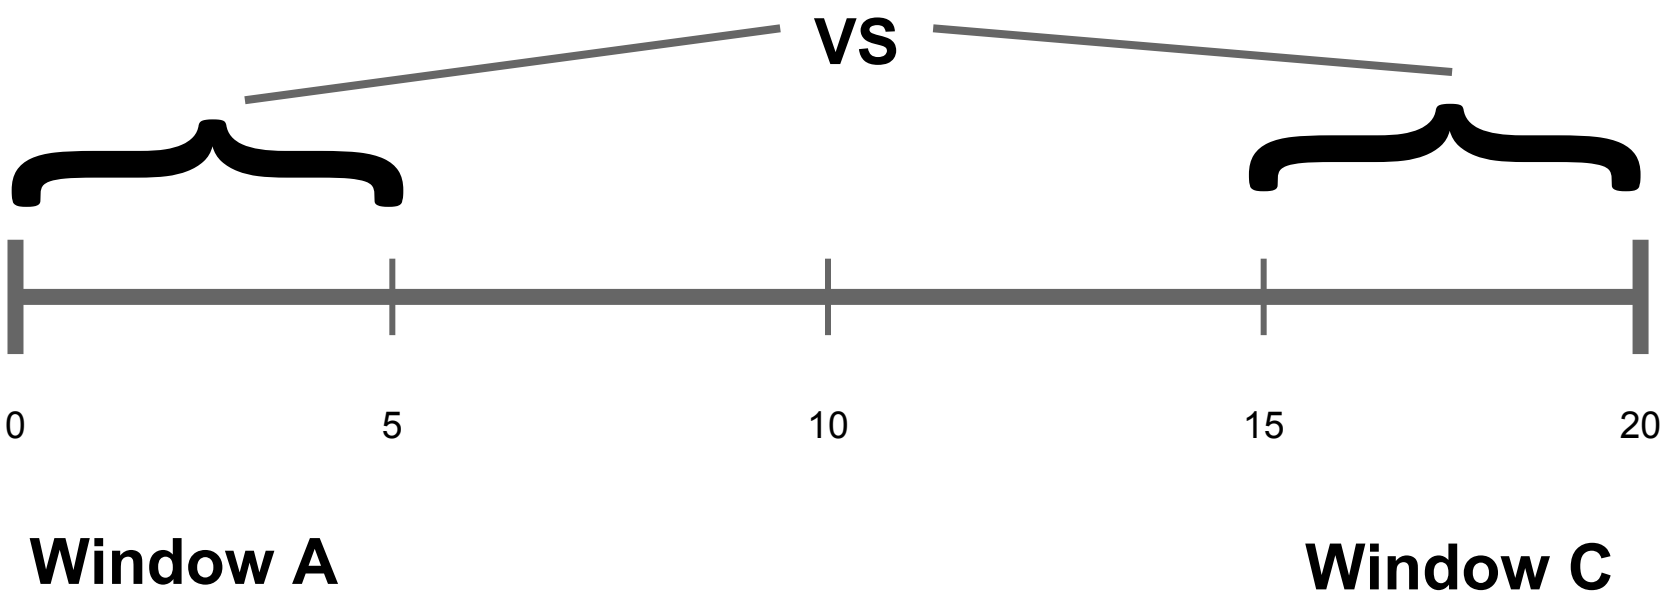

Testing Persistence

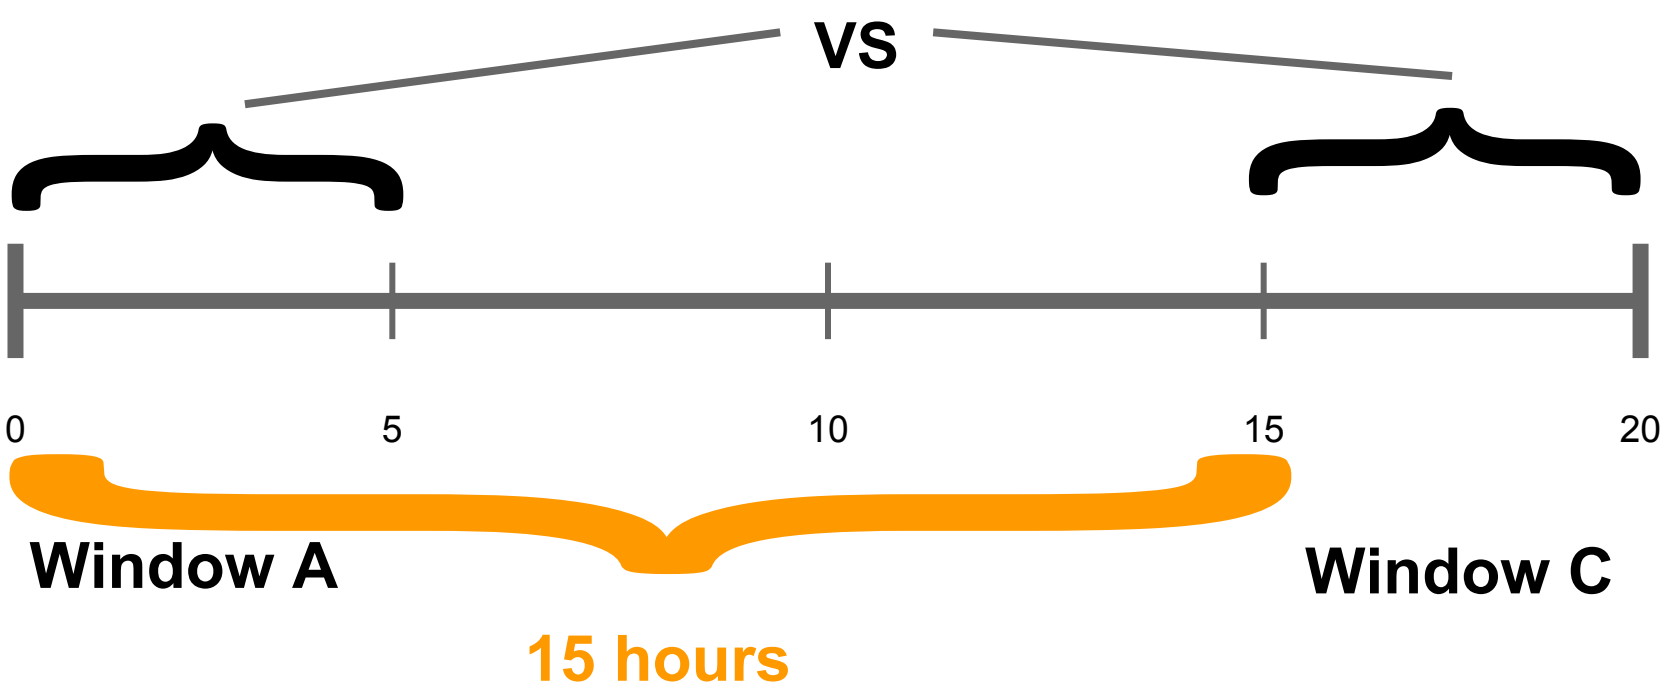

Latency Ratio Difference Over Time

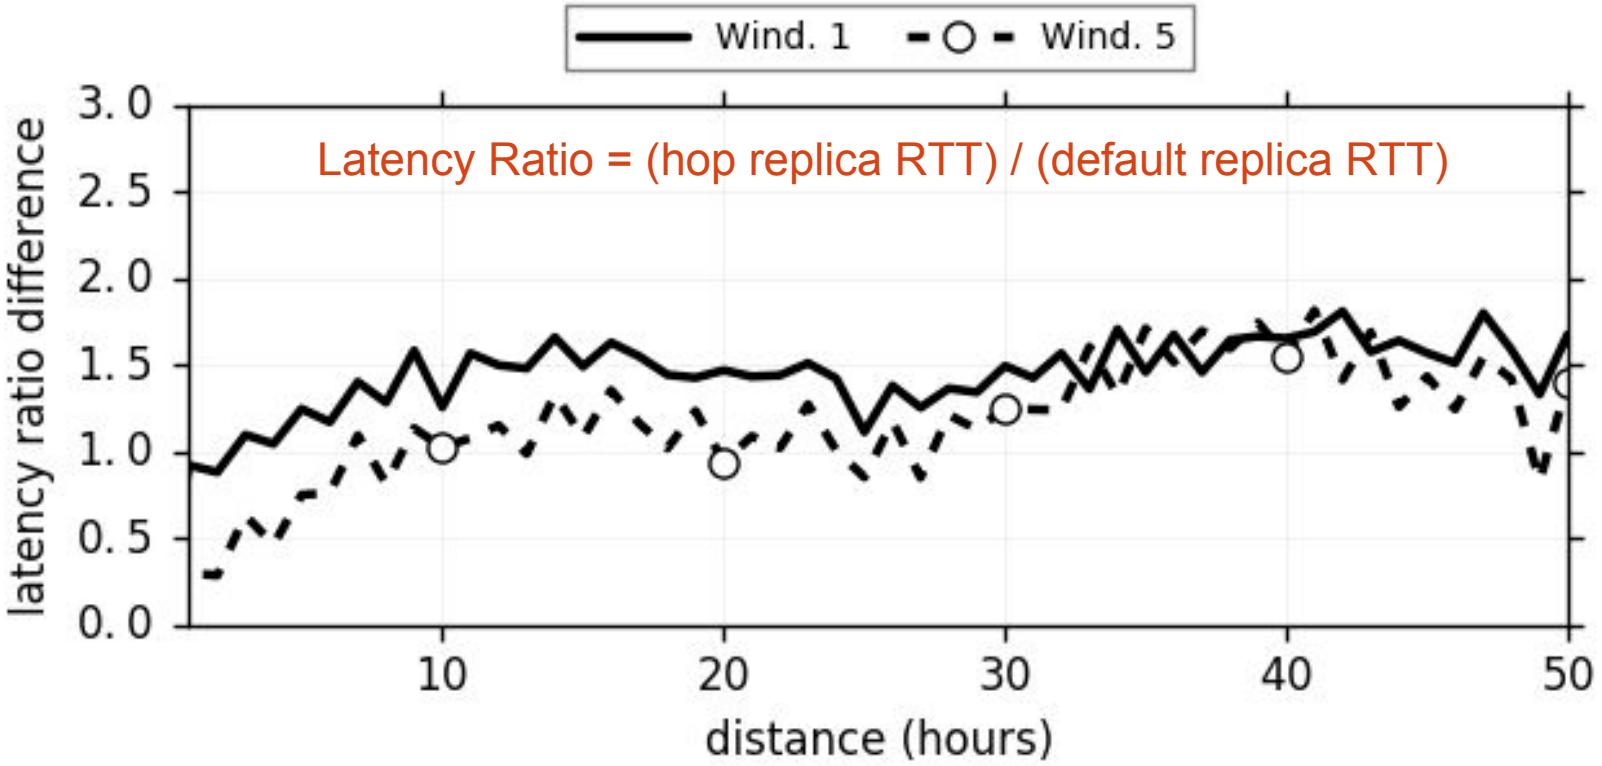

Latency Ratio Difference Over Time

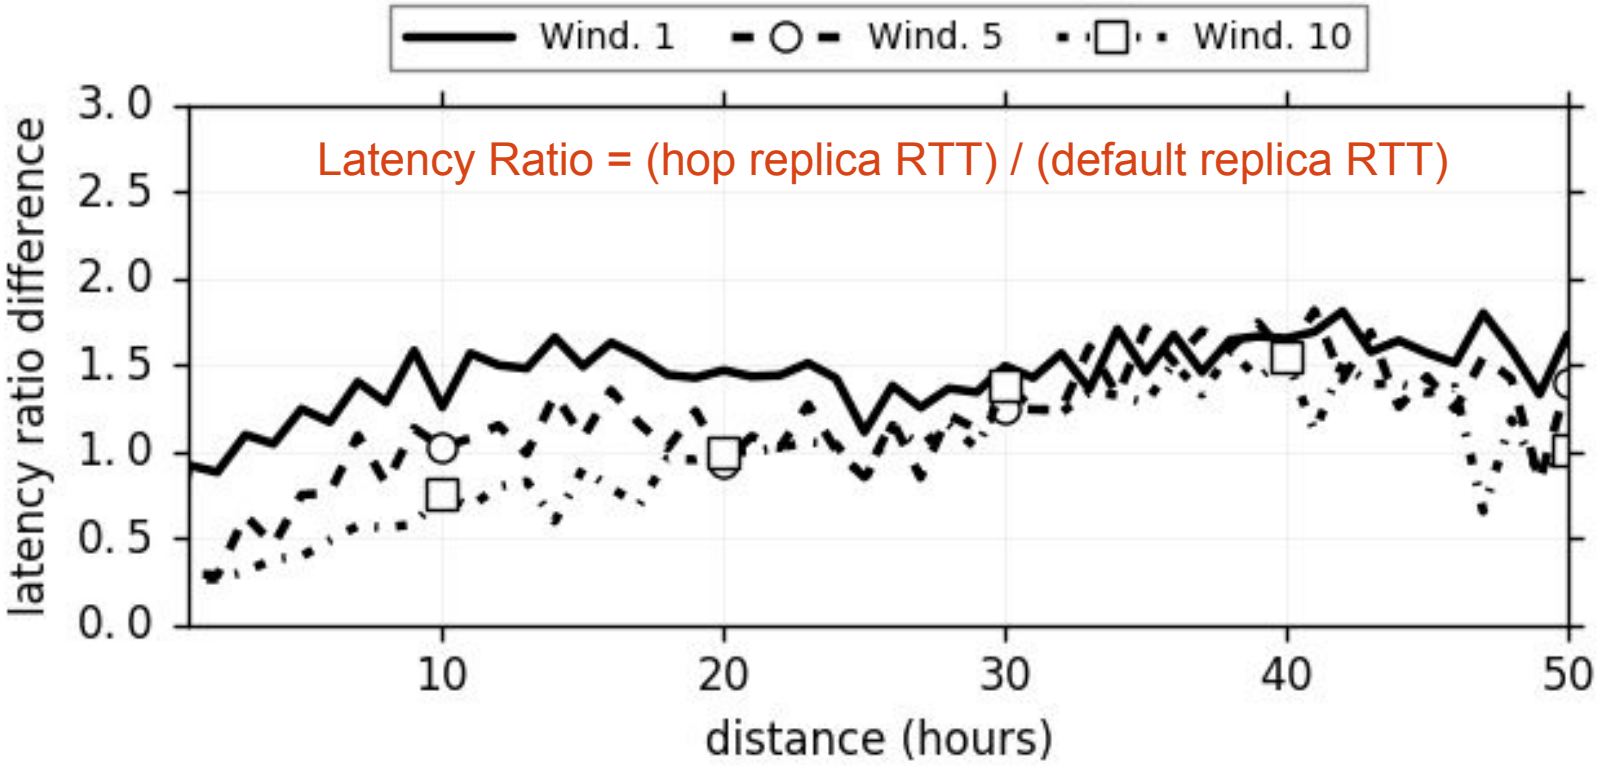

Latency Ratio Difference Over Time

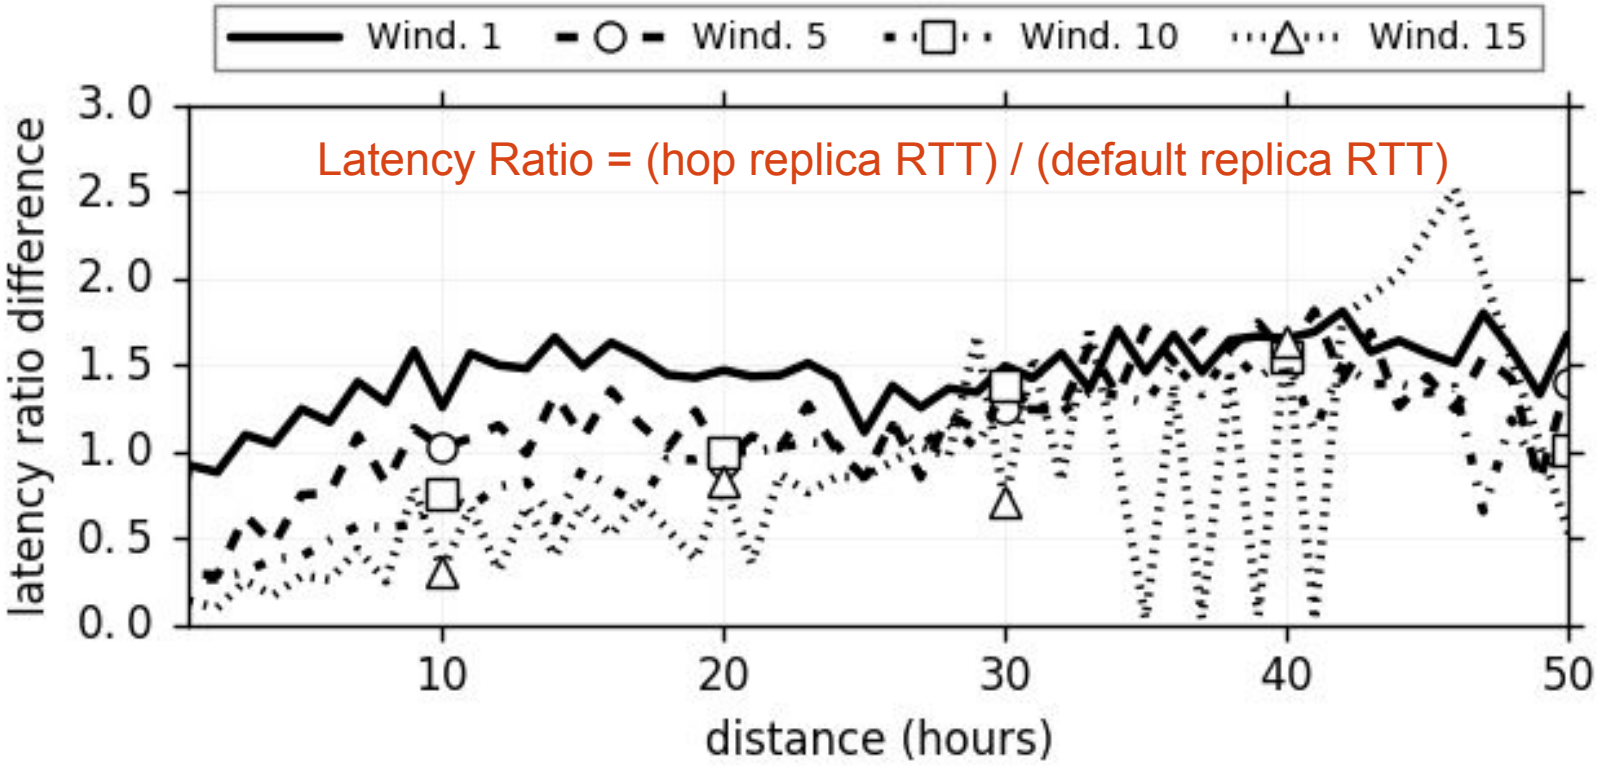

Latency Ratio Difference Over Time

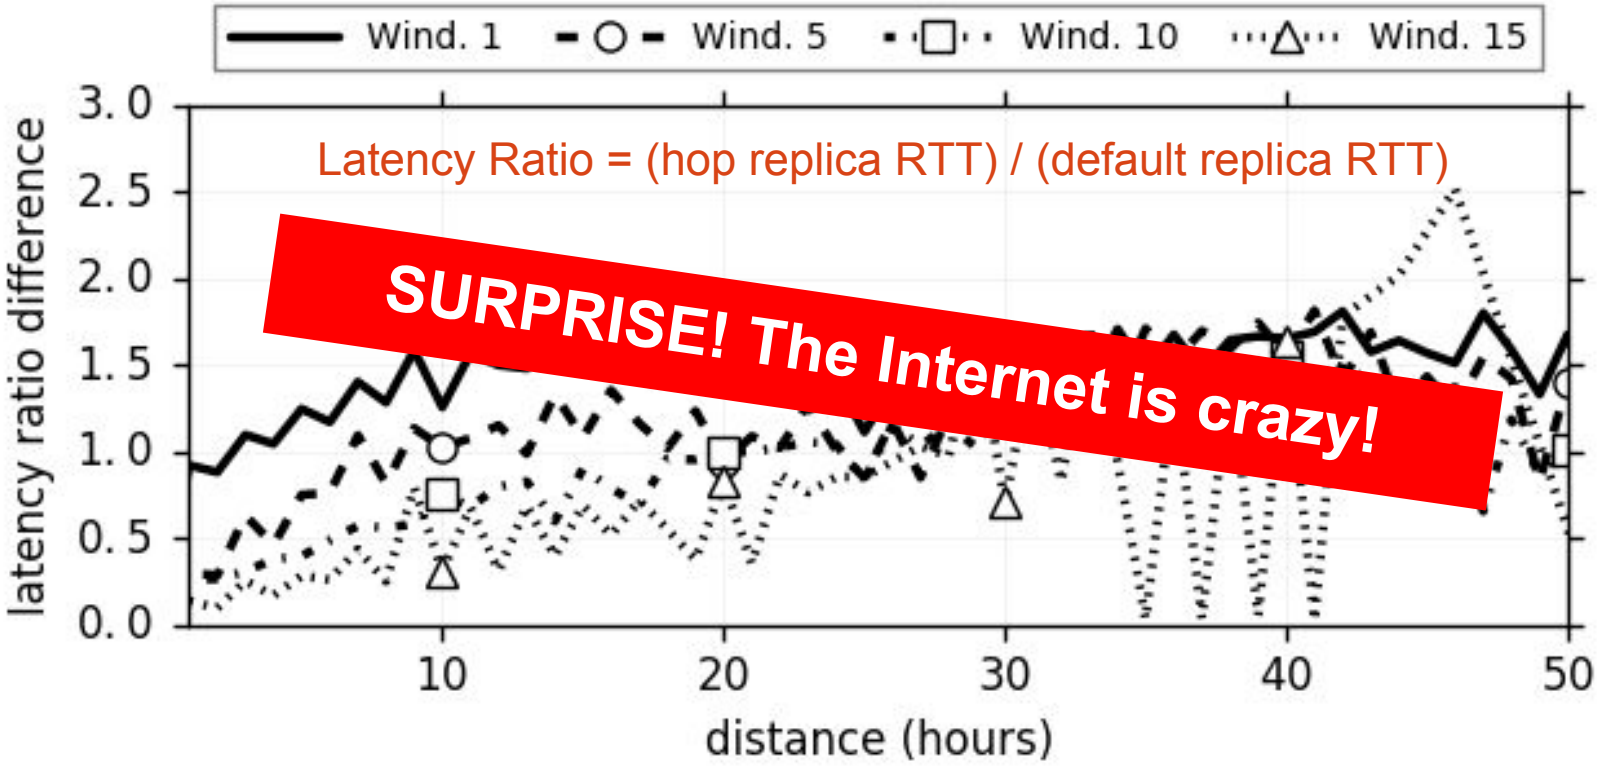

Filter: at *least* one valley

Subnet A	{0,0,0,0,0,V,0,0,0,0,0,0,V}
Subnet B	{0,0,0,0,0,0,0,0,0,0,0,0,0}
Subnet C	{V,V,V,V,0,0,0,0,V,V,V,0,V}

Filter: at *least* one valley

Subnet A	{0,0,0,0,0,V,0,0,0,0,0,0,V}
~~Subnet B~~	~~{0,0,0,0,0,0,0,0,0,0,0,0,0}~~
Subnet C	{V,V,V,V,0,0,0,0,V,V,V,0,V}

Filter: at *least* one valley

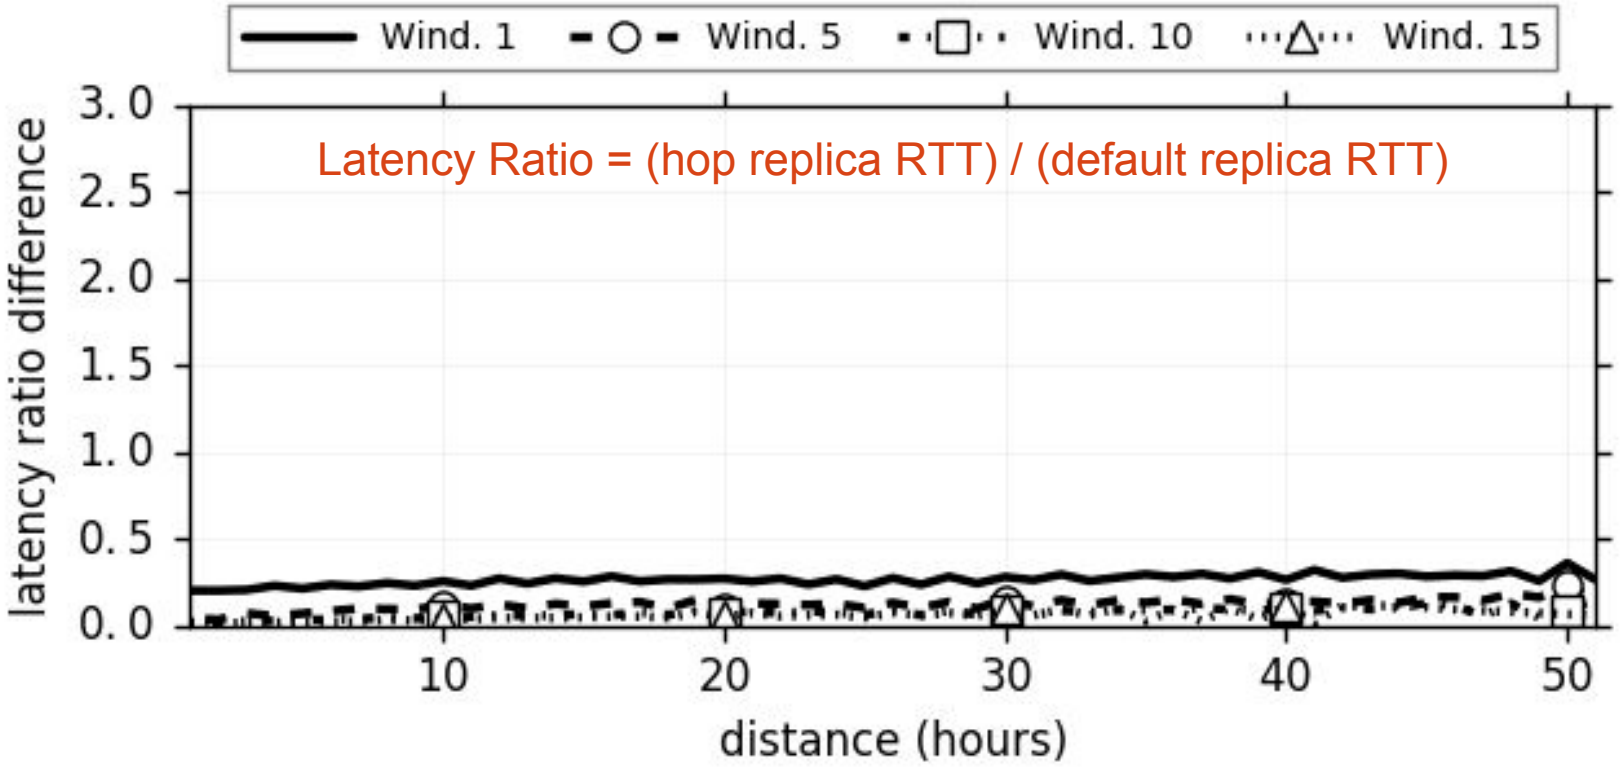

Filter: at *least* one valley

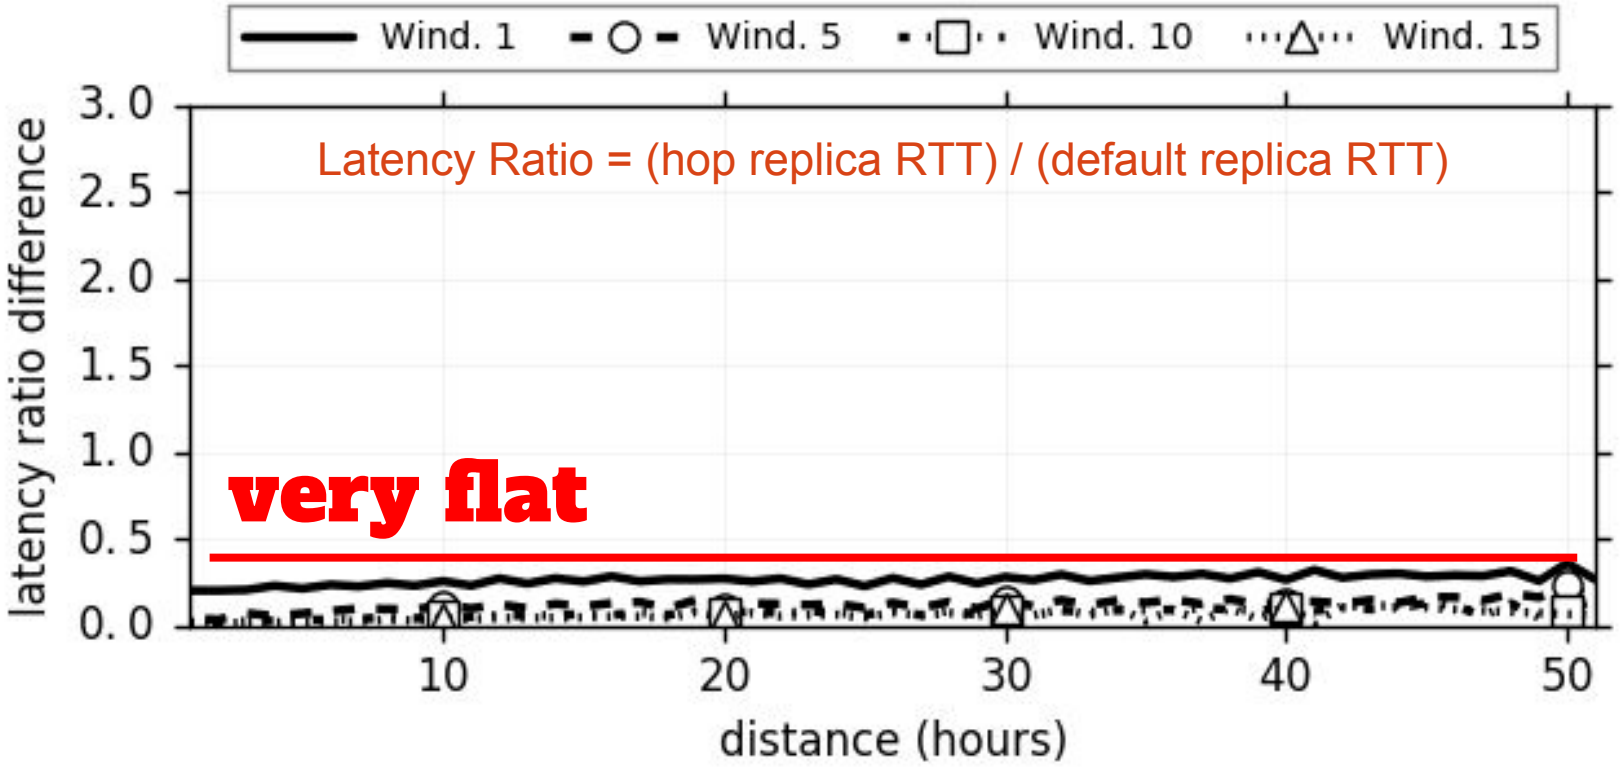

Filter: at *least* one valley

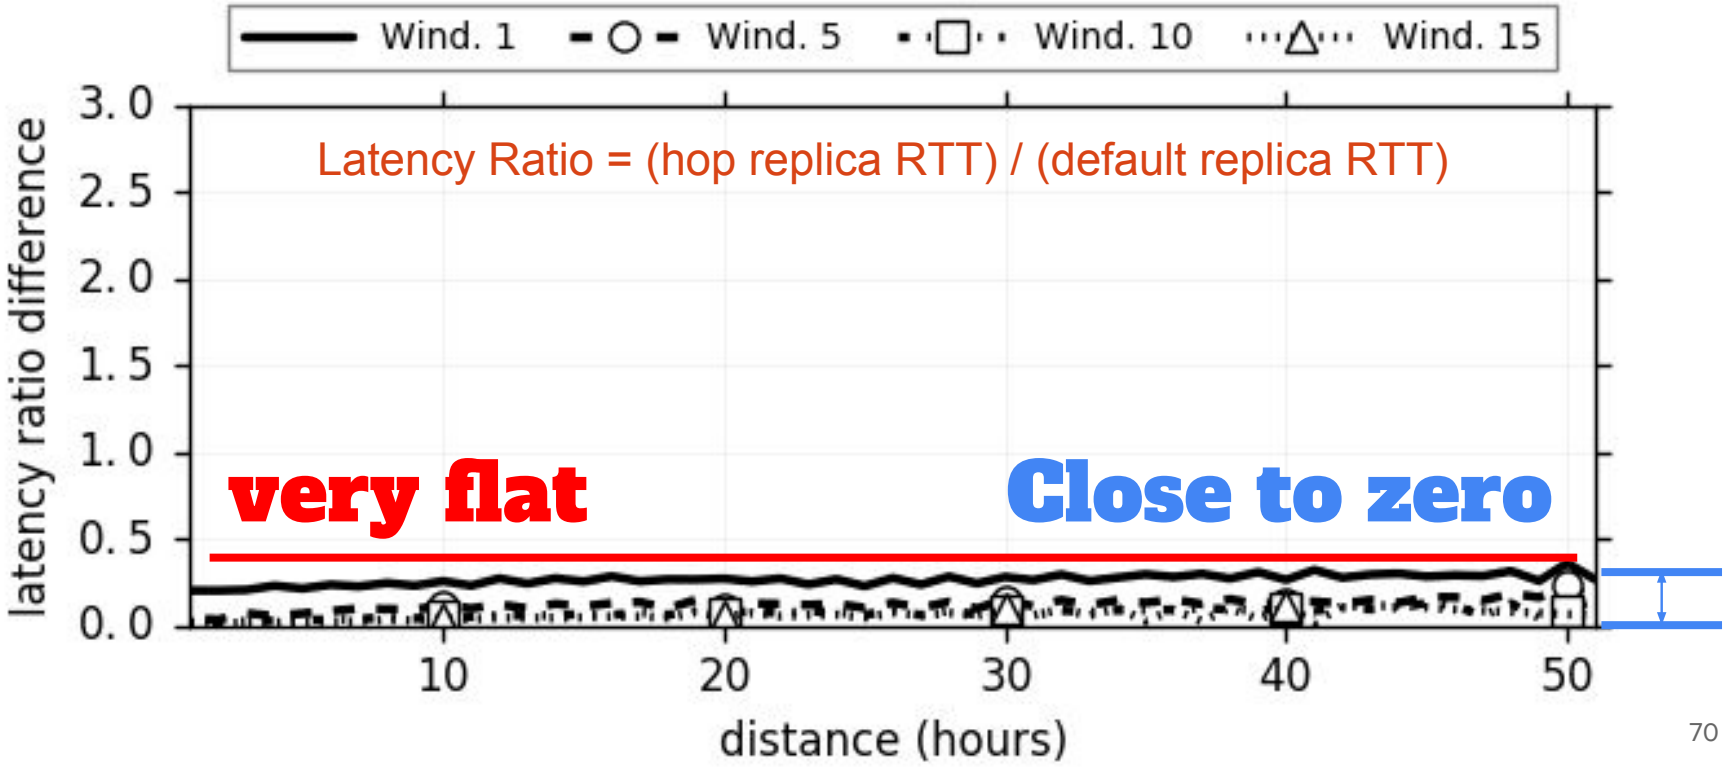

Parameter Exploration

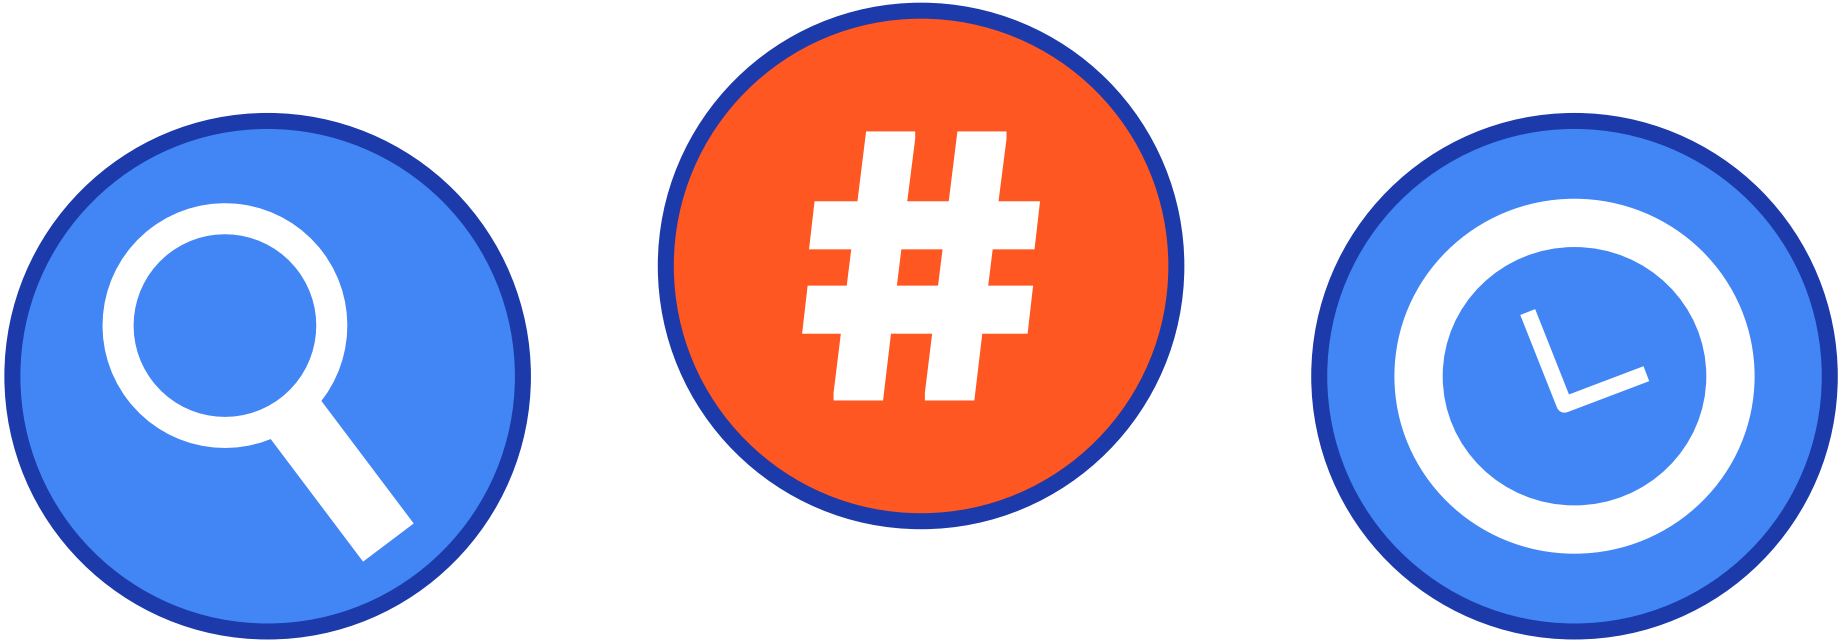

V_{thresh} = How *deep* are the valleys from *useful* subnets?

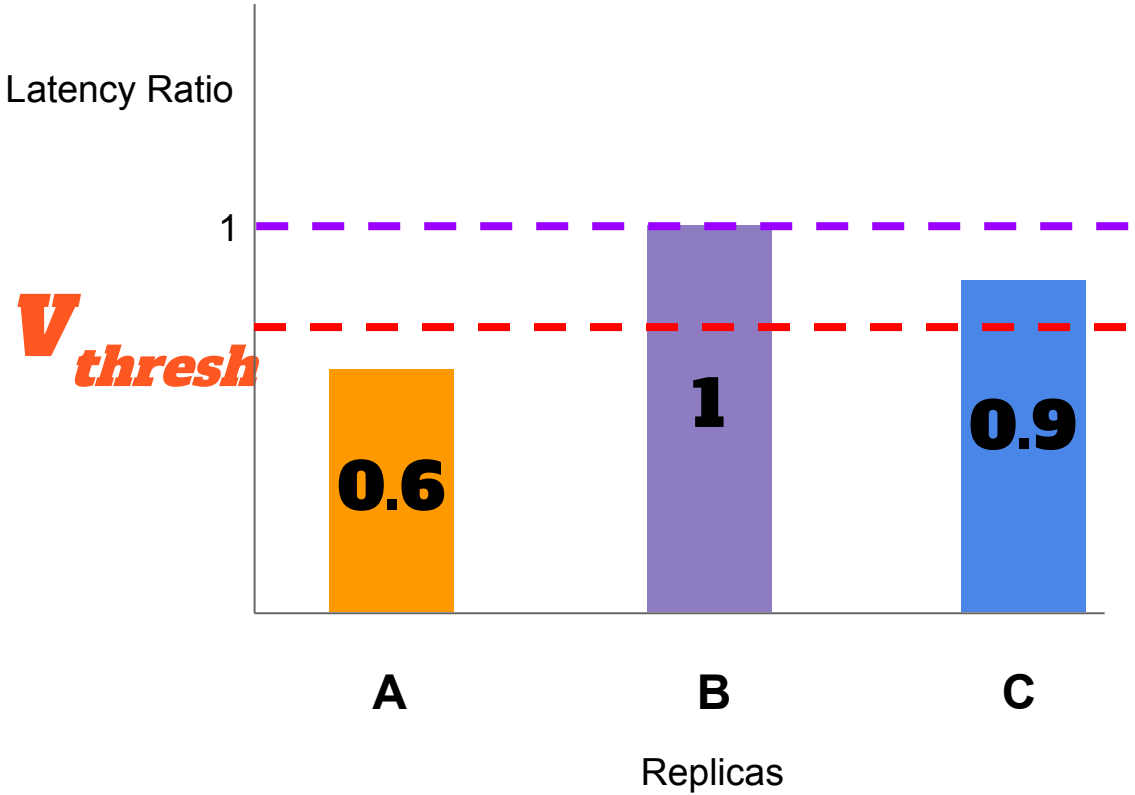
Latency Ratio
1
V_{thresh}
0.6
1
0.9
A
B
C
Replicas

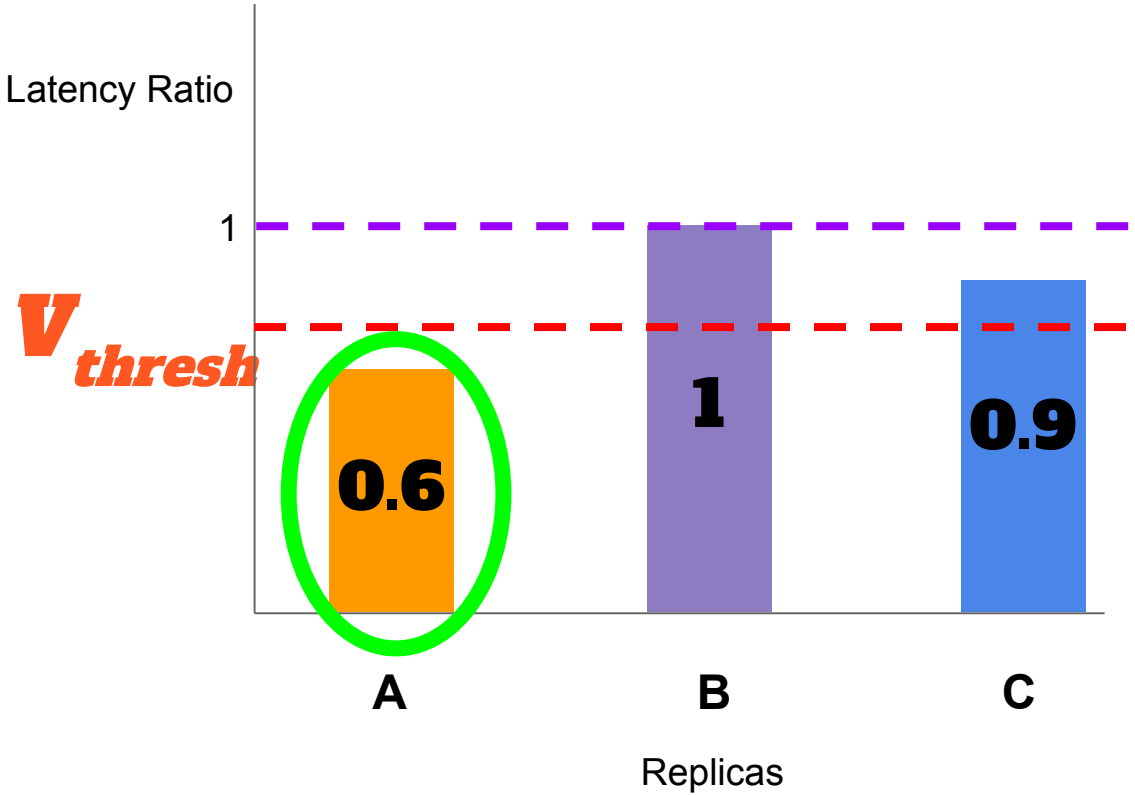
Latency Ratio
1
V_{thresh}
0.6
1
0.9
A
B
C
Replicas

V_{freq} = How *often* do valleys occur in *useful* subnets?


TRAINING WINDOW

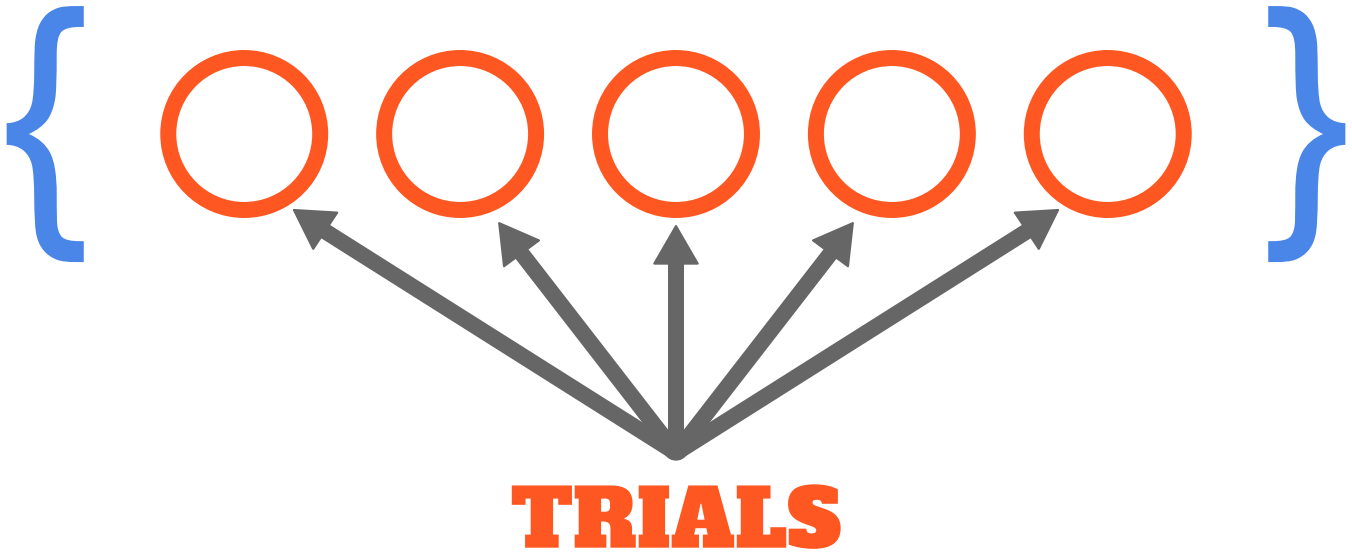
TRIALS

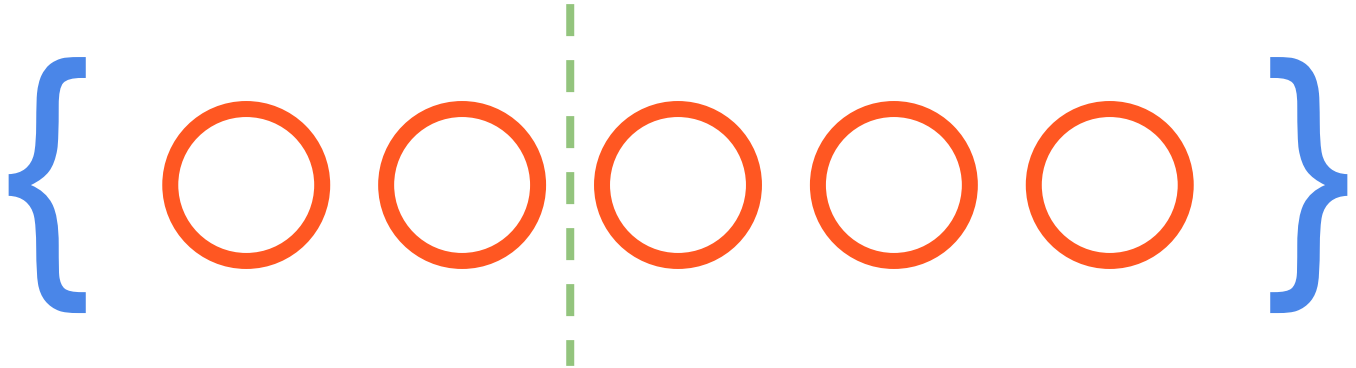

$$v_{freq} = \frac{2}{5}$$

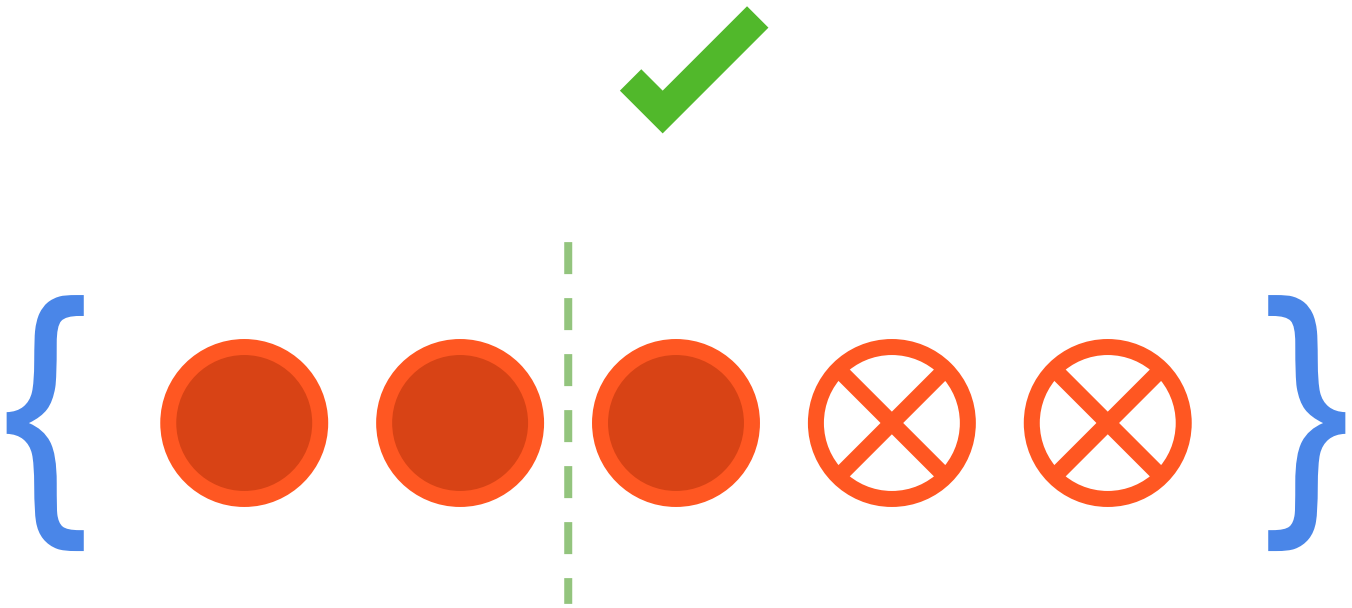

$v_{freq} = 2/5$

Valley-Prone Subnet

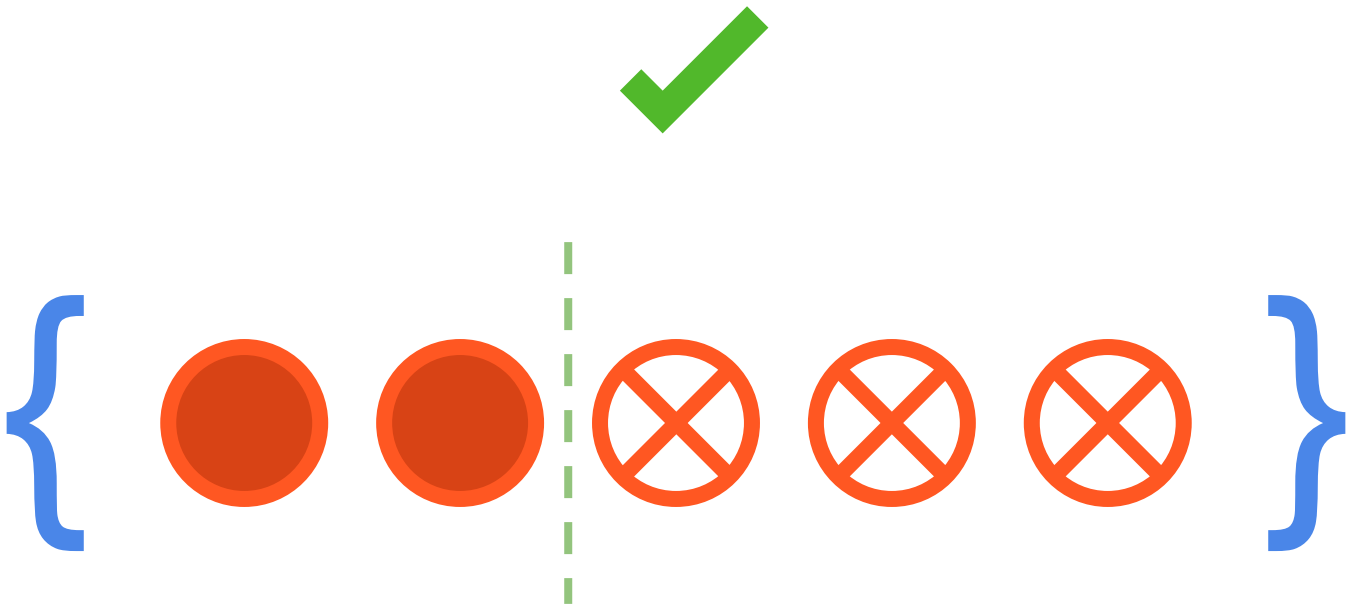

$v_{freq} = \frac{2}{5}$

Valley-Prone Subnet

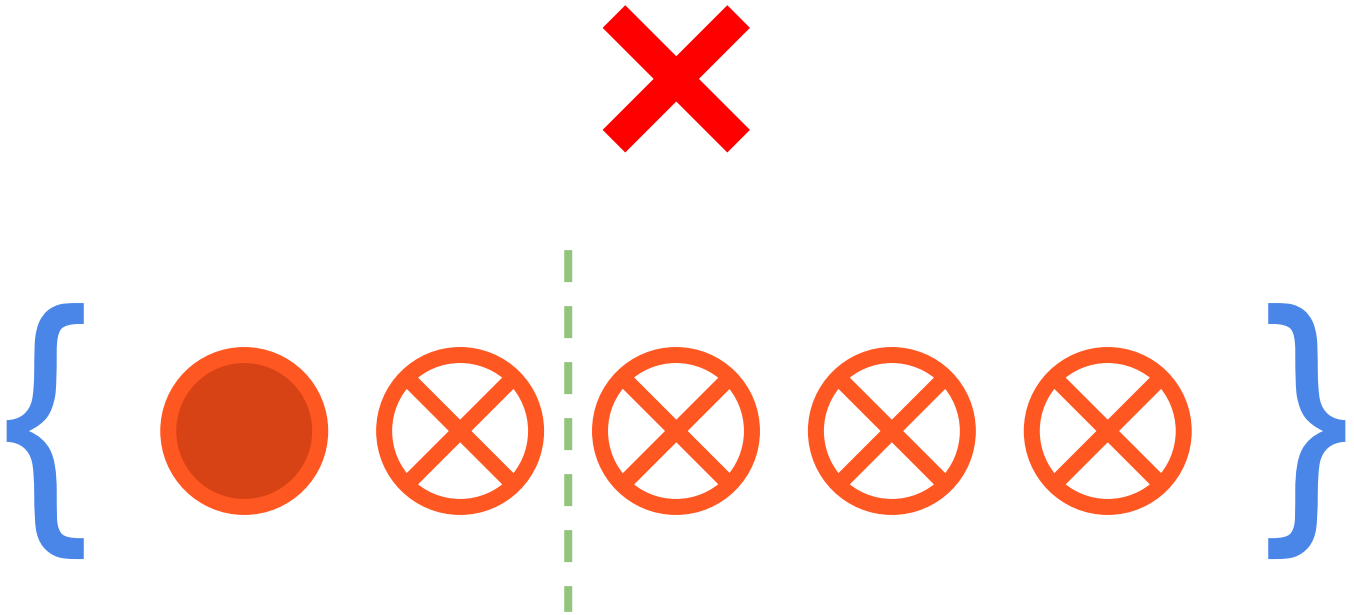

$v_{freq} = \frac{2}{5}$

NOT Valley-Prone Subnet

Overview of Drongo:

1. Collect training window

Overview of Drongo:

1. Collect training window

2. Count the # of sufficiently deep valleys

Overview of Drongo:

1. Collect training window

2. Count the # of sufficiently deep valleys

3. Apply subnet assimilation
    a. Training window is *already* complete
    b. Both parameters met

System Wide Performance

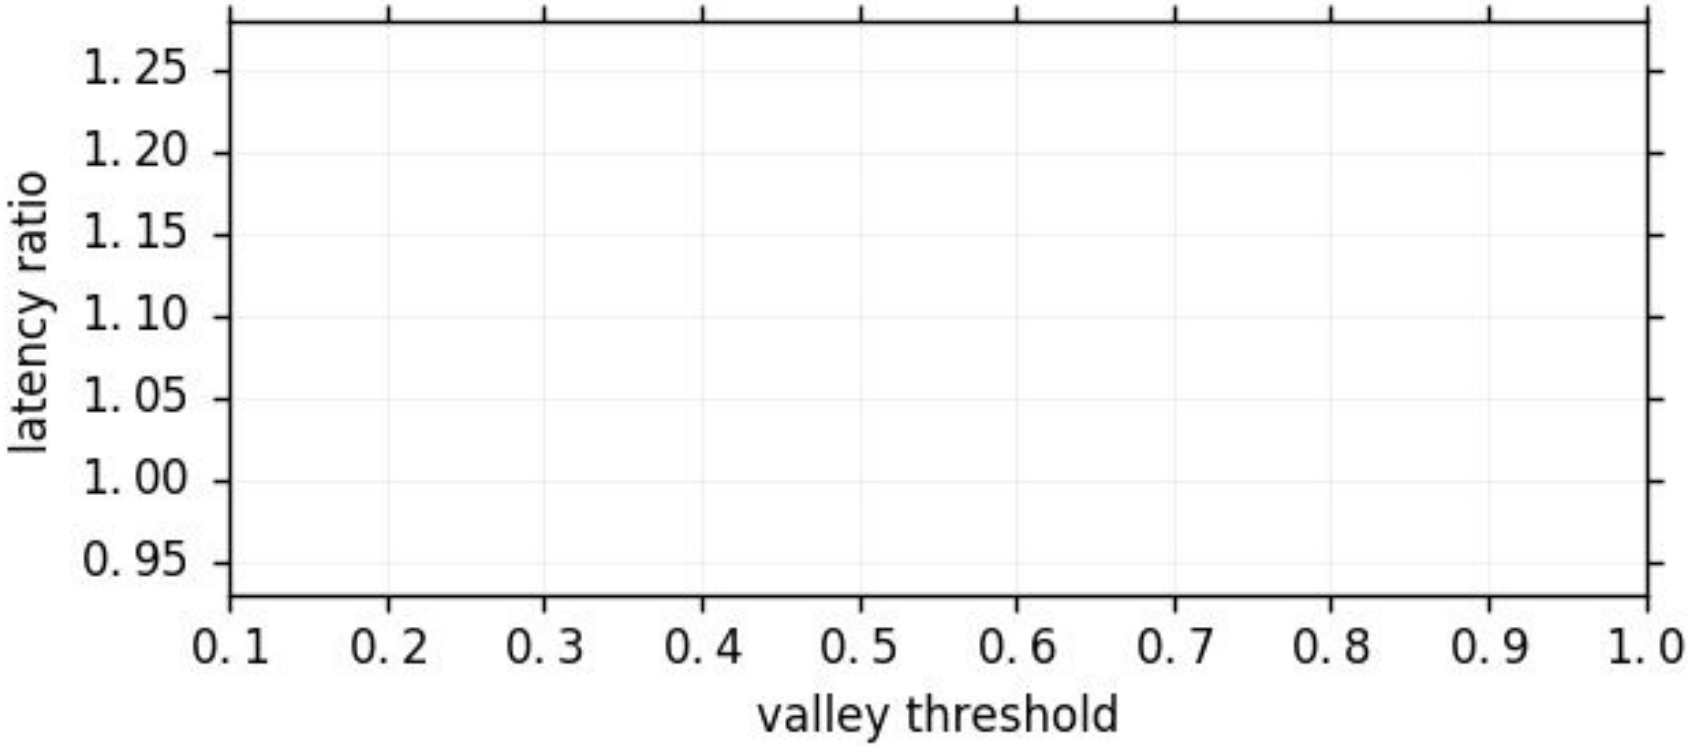

System Wide Performance

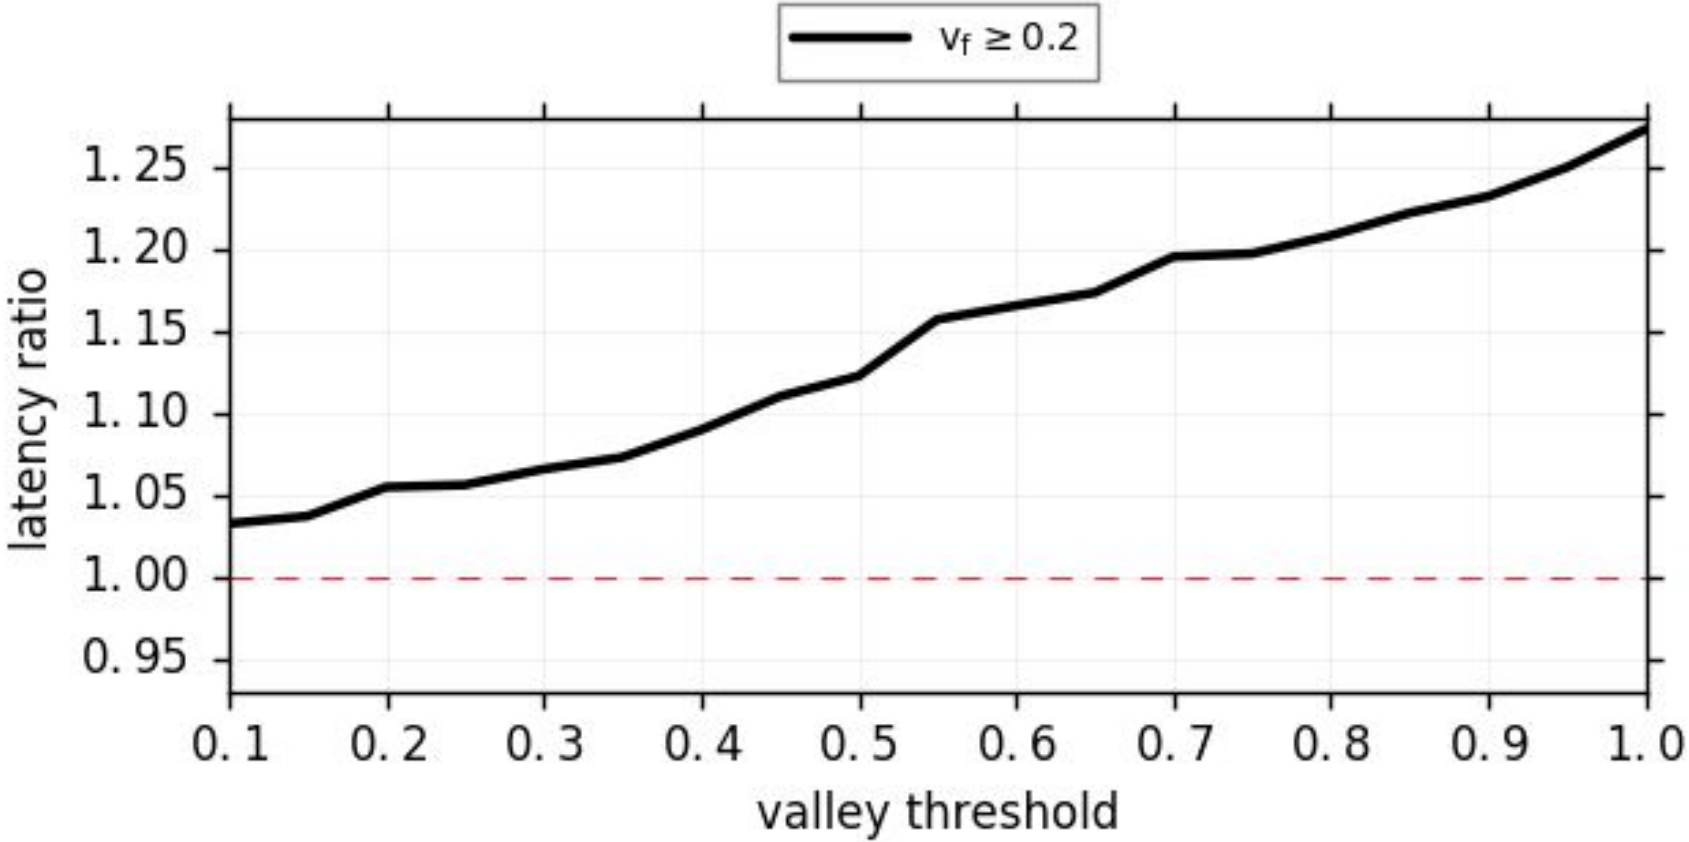

System Wide Performance

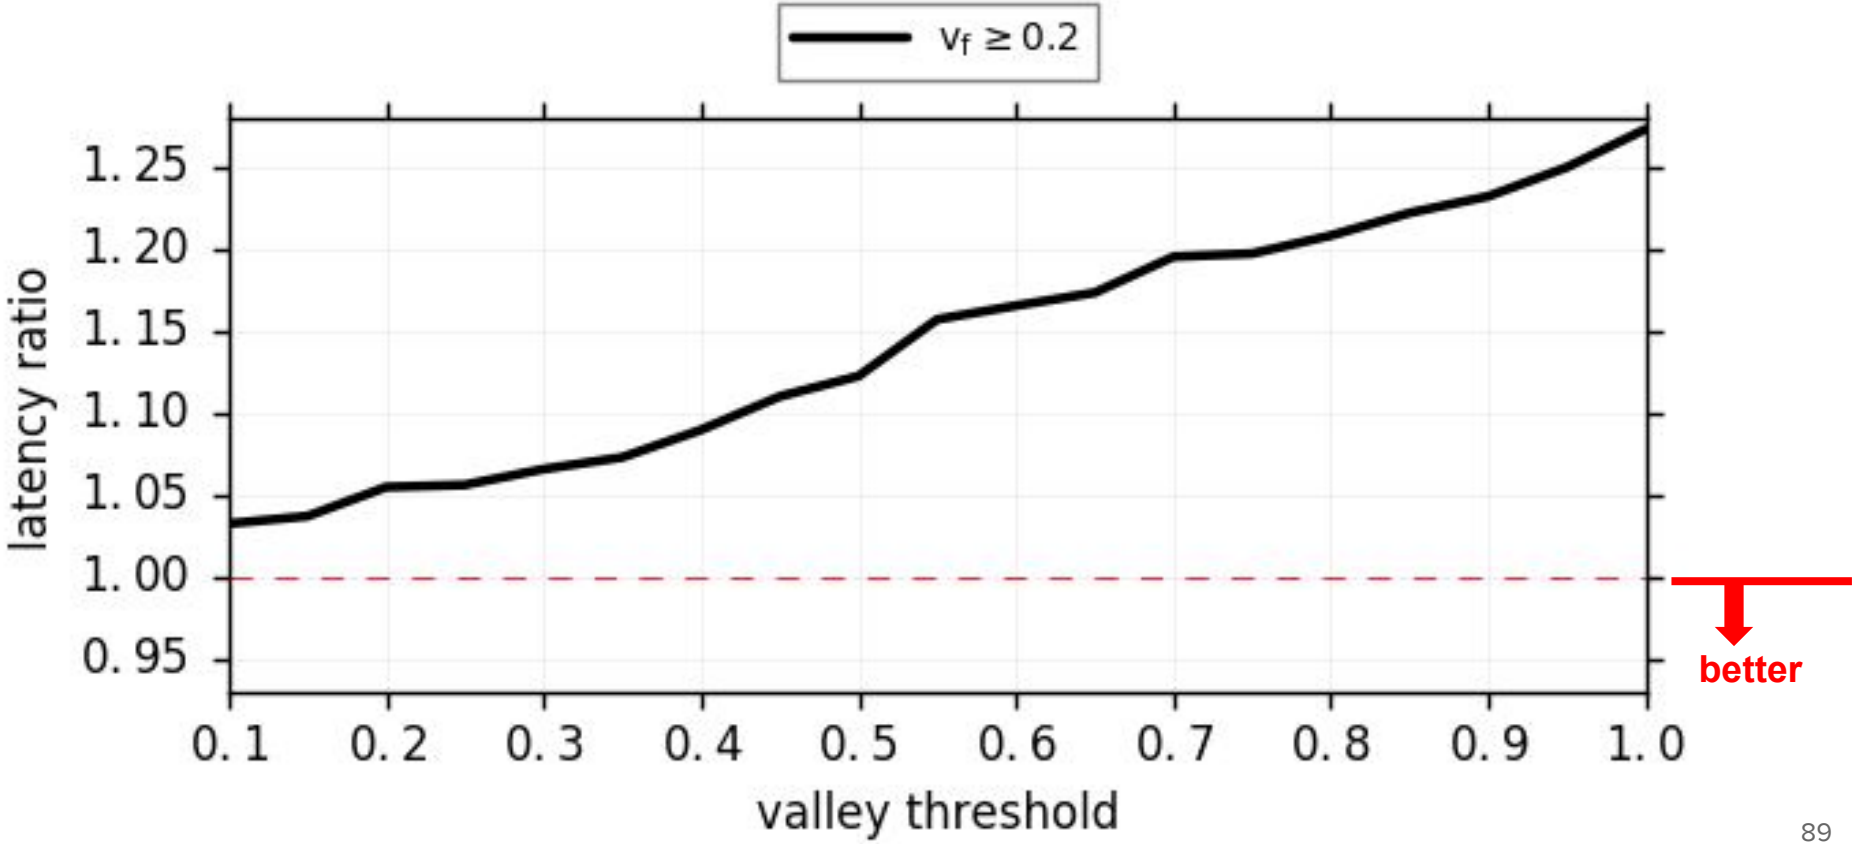

System Wide Performance

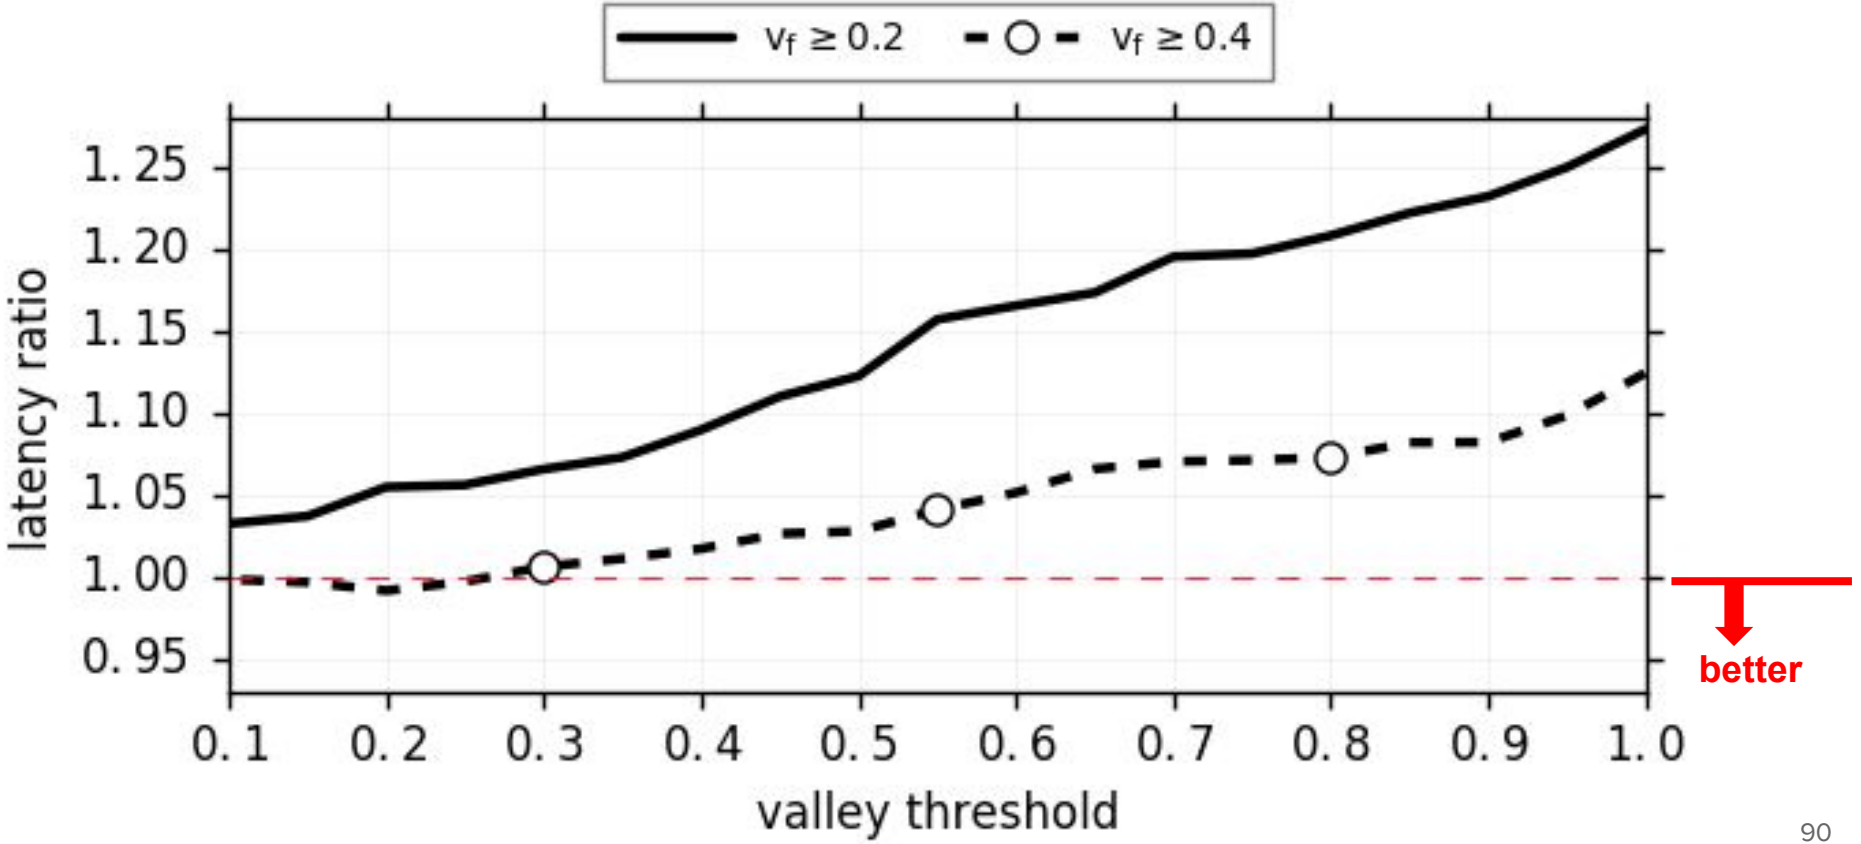

System Wide Performance

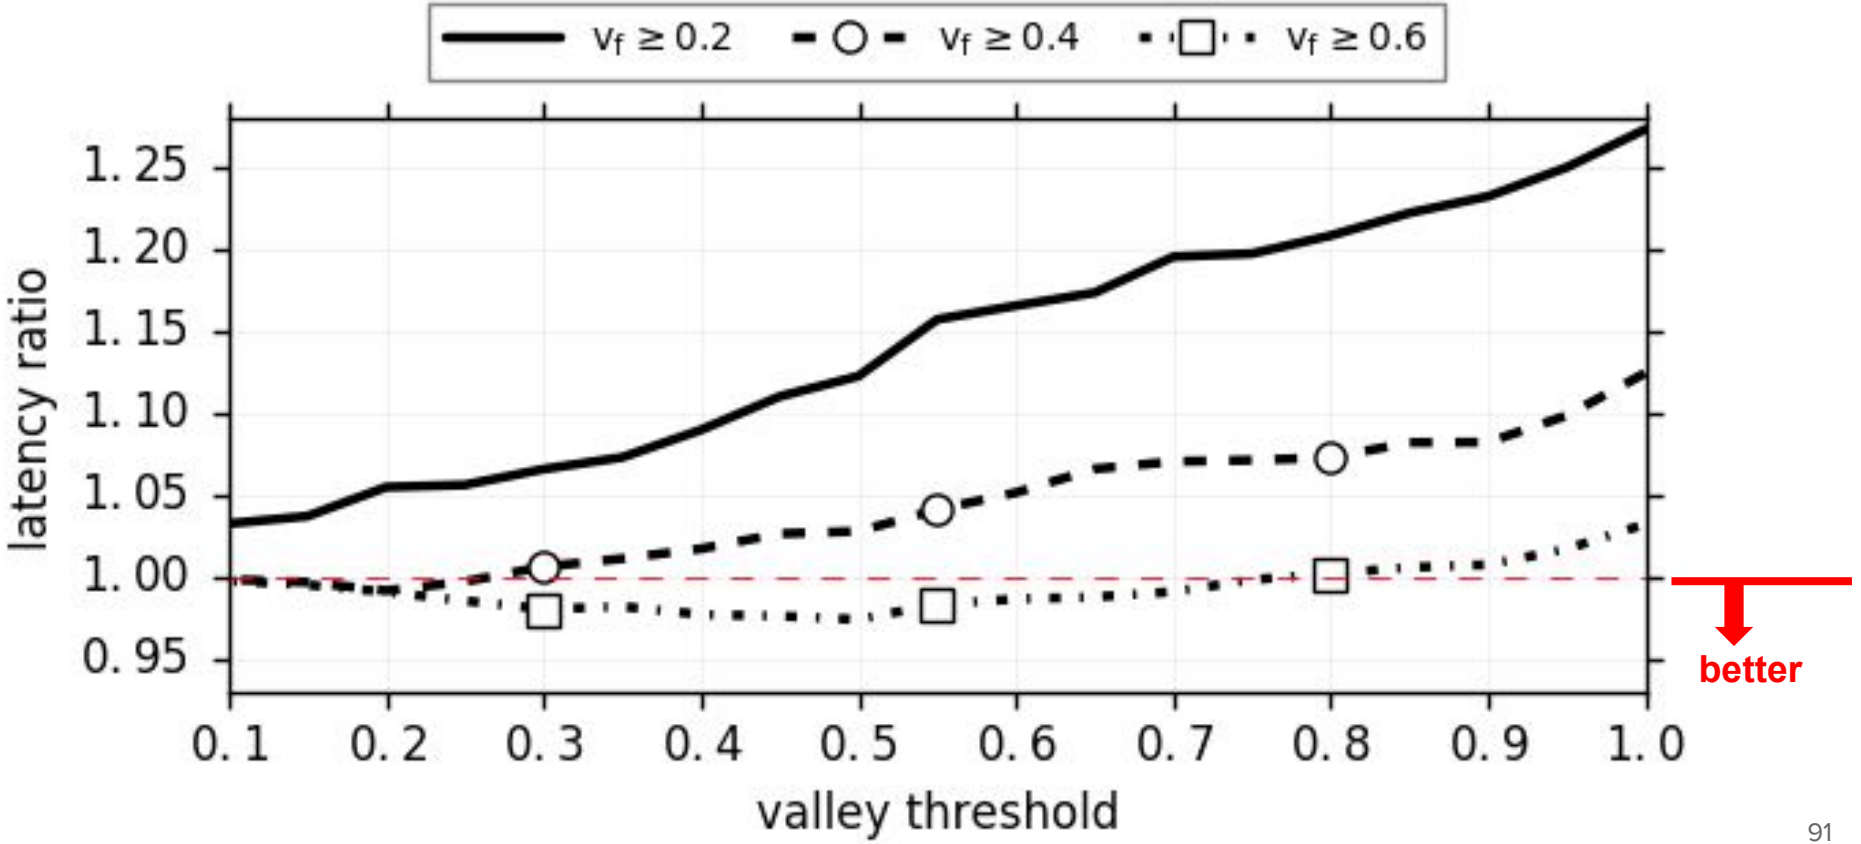

System Wide Performance

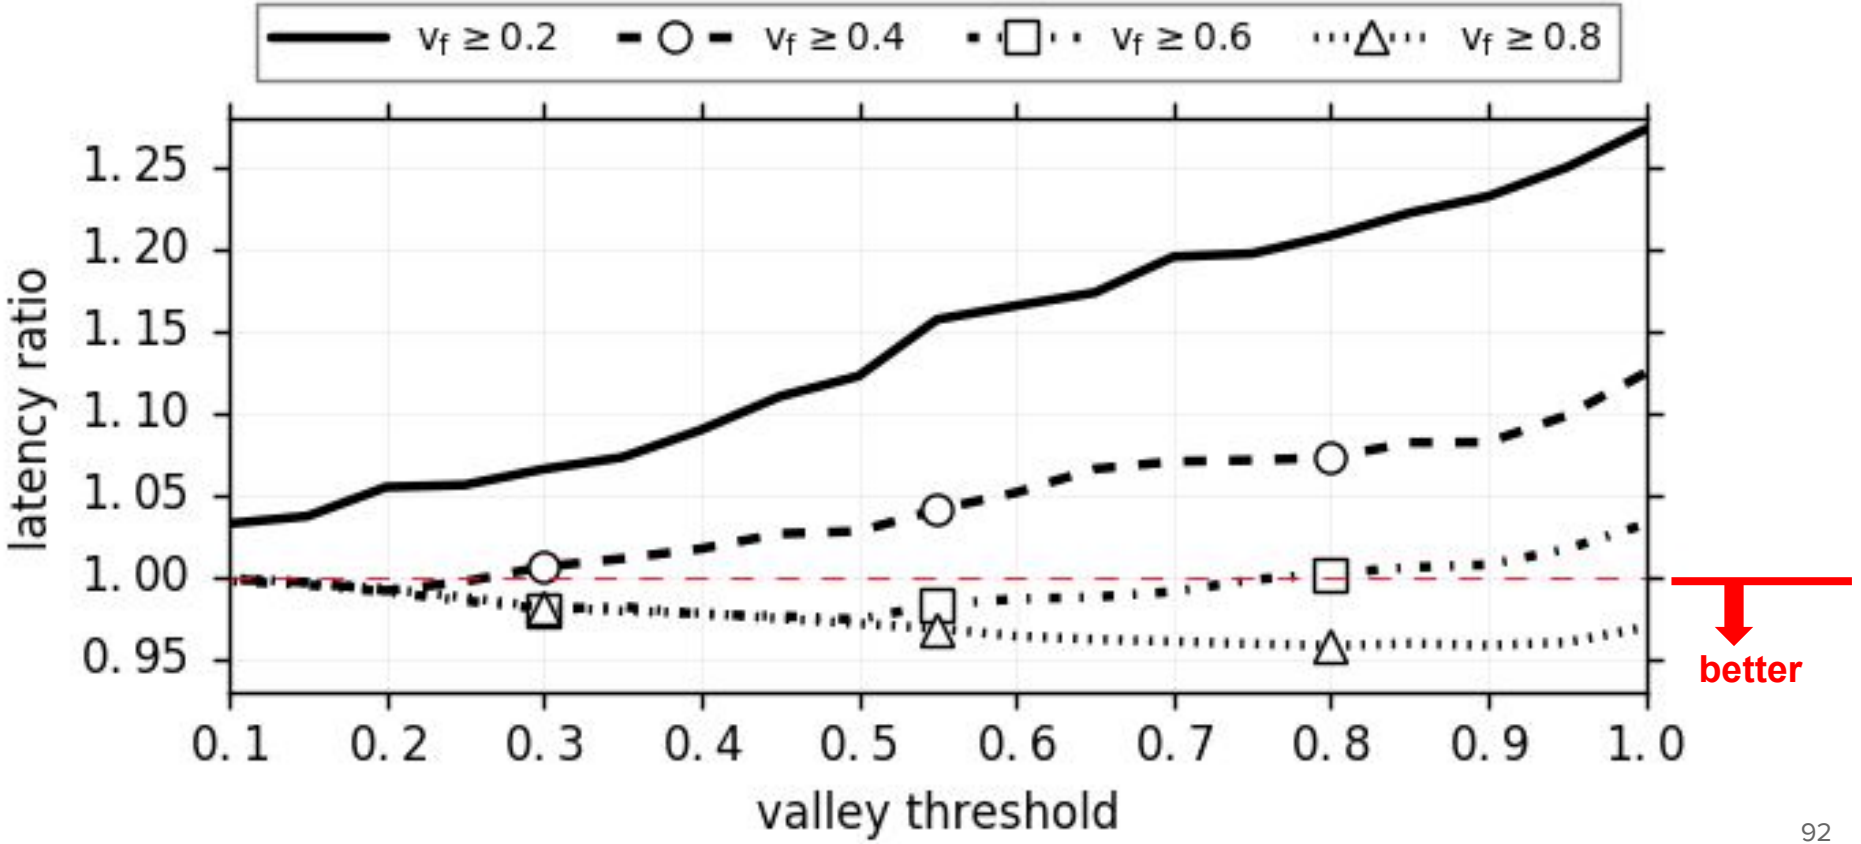

System Wide Performance

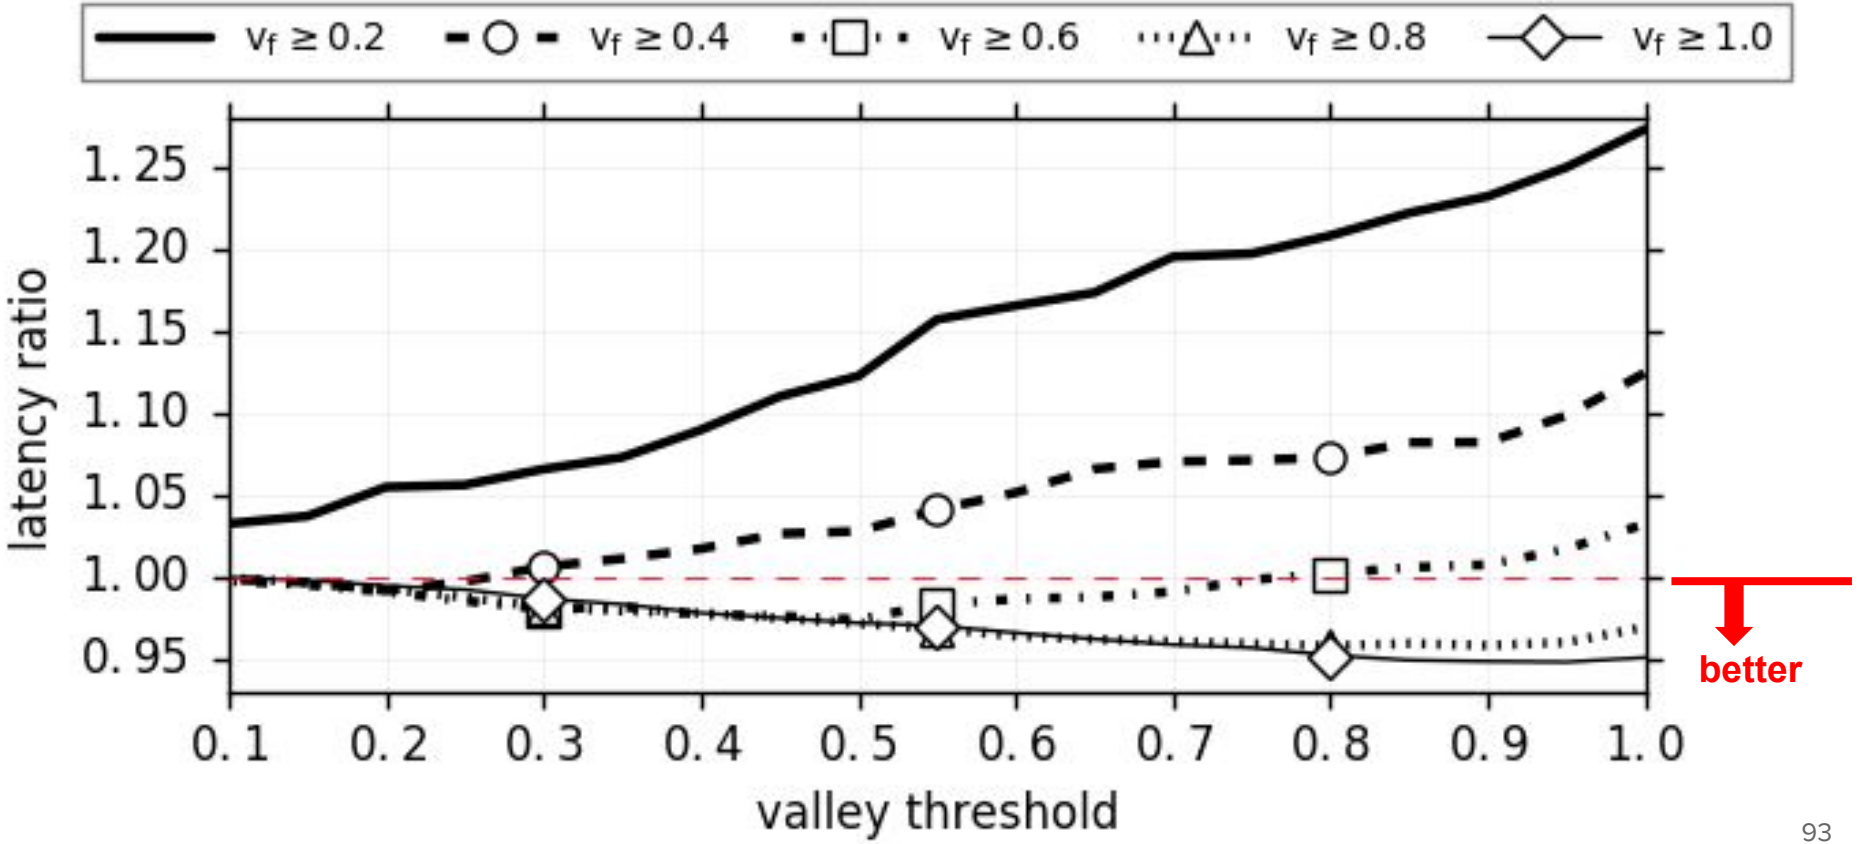

System Wide Performance

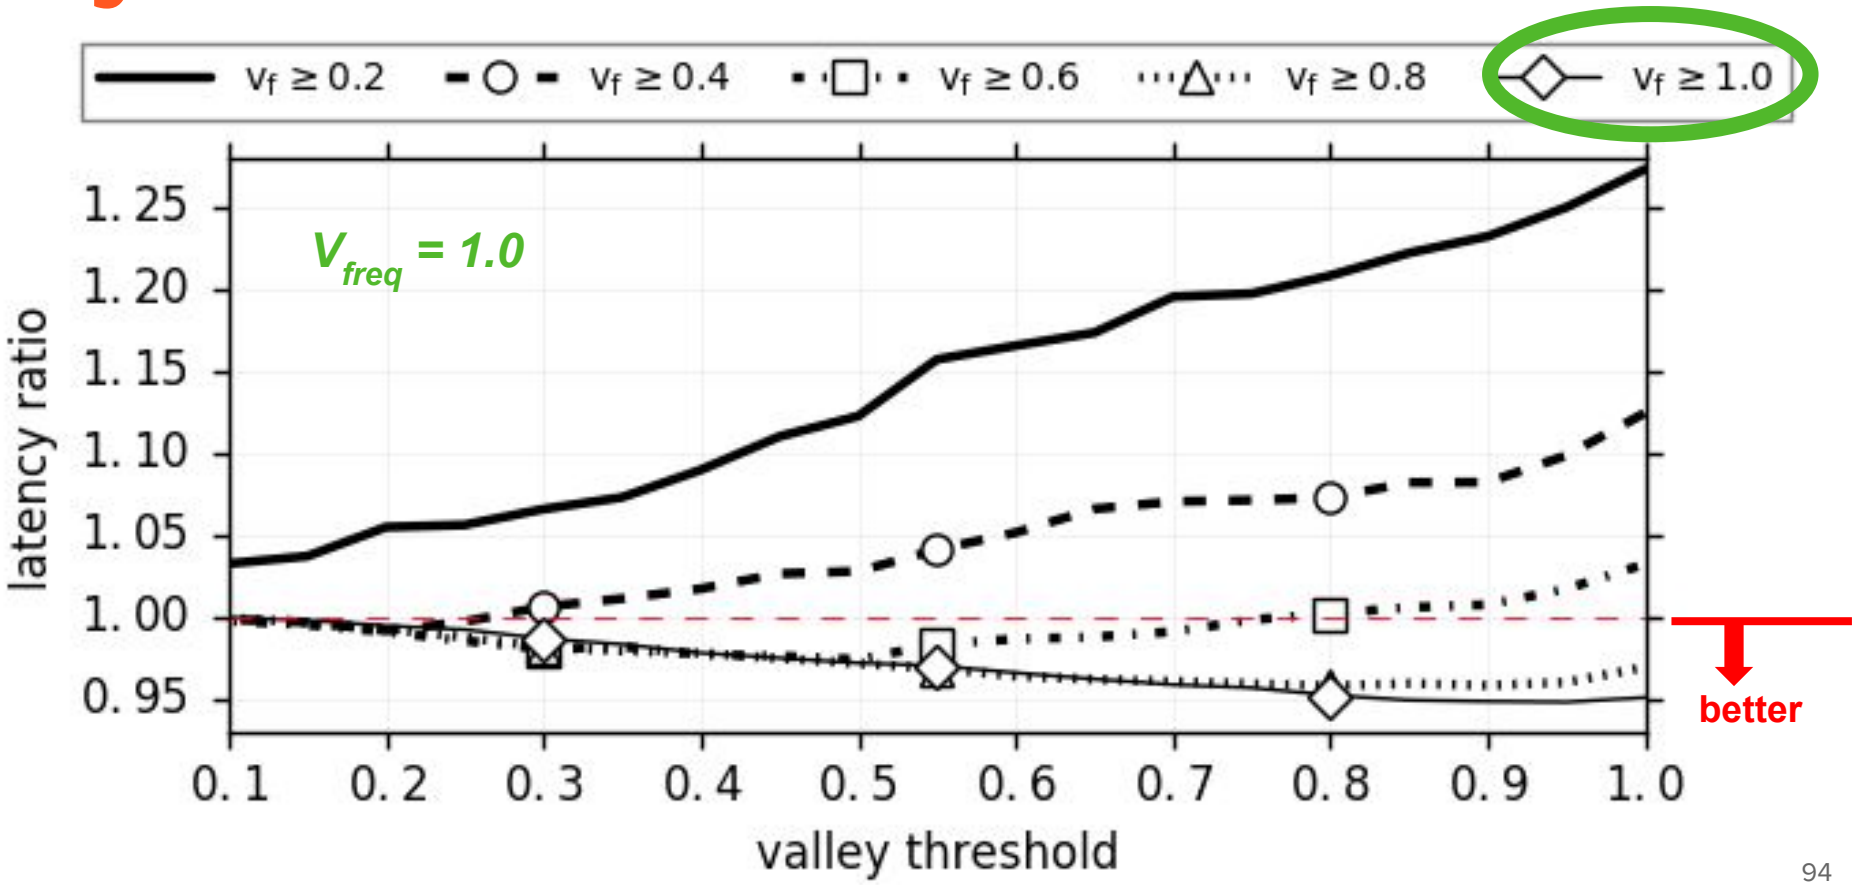

System Wide Performance

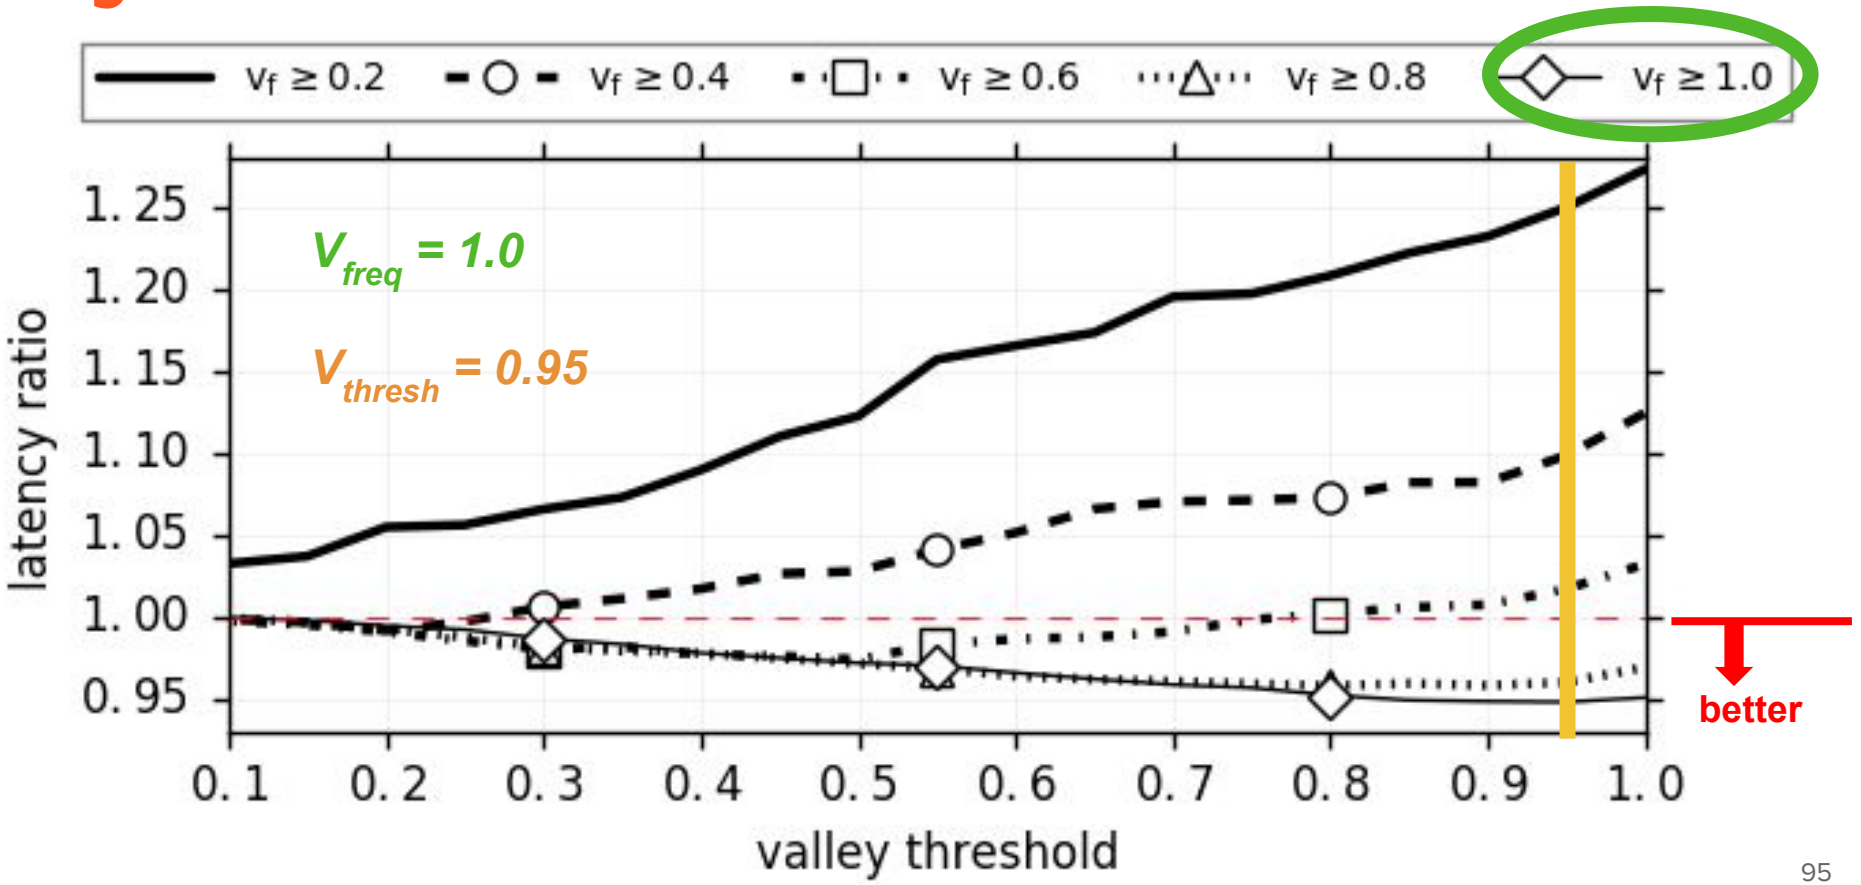

Switch Quality

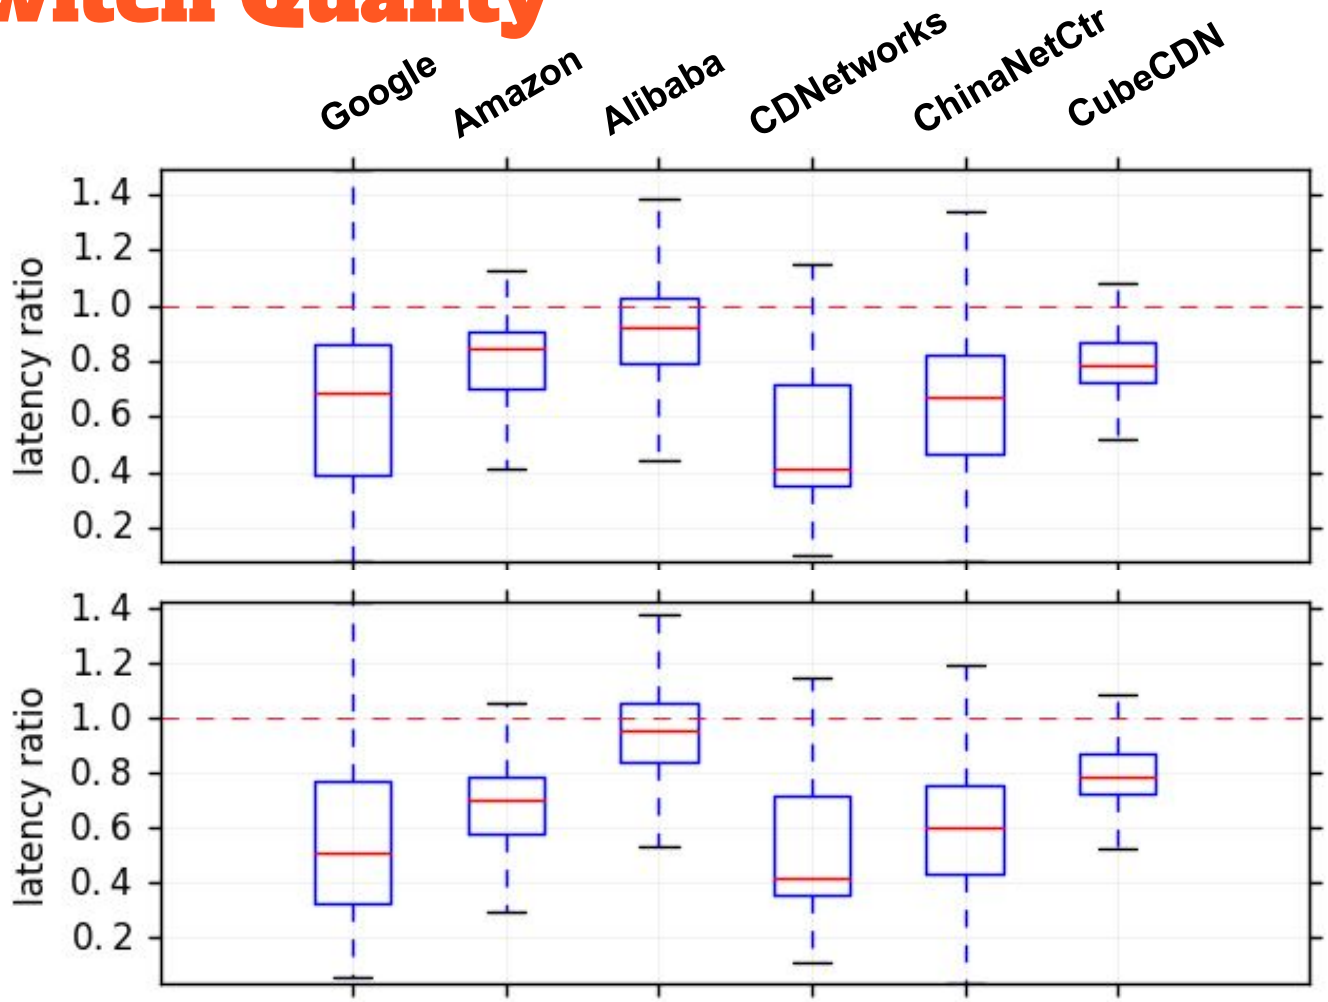

Global Params

Per Prov. Params

Conclusion & Insights

- CDNs have a lot of room for improvement

Conclusion & Insights

- CDNs have a lot of room for improvement
- Clients can help

Conclusion & Insights

- CDNs have a lot of room for improvement
- Clients can help
- Low requirements

Conclusion & Insights

- CDNs have a lot of room for improvement
- Clients can help
- Low requirements
- Can provide 50% improvement

Questions?

Clients Affected

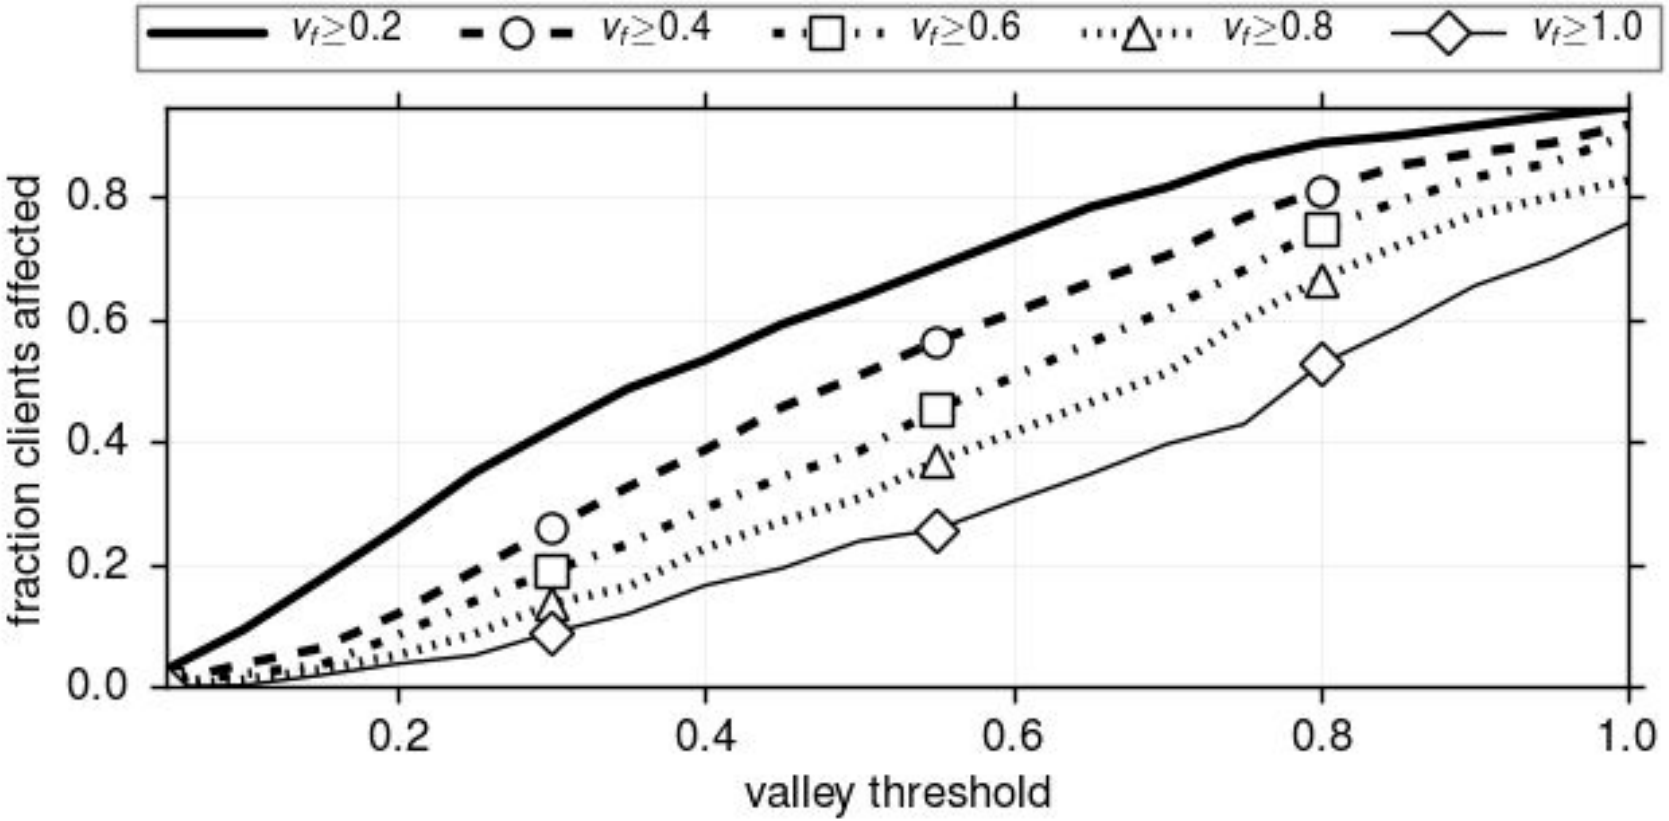

Per Provider Overall Performance

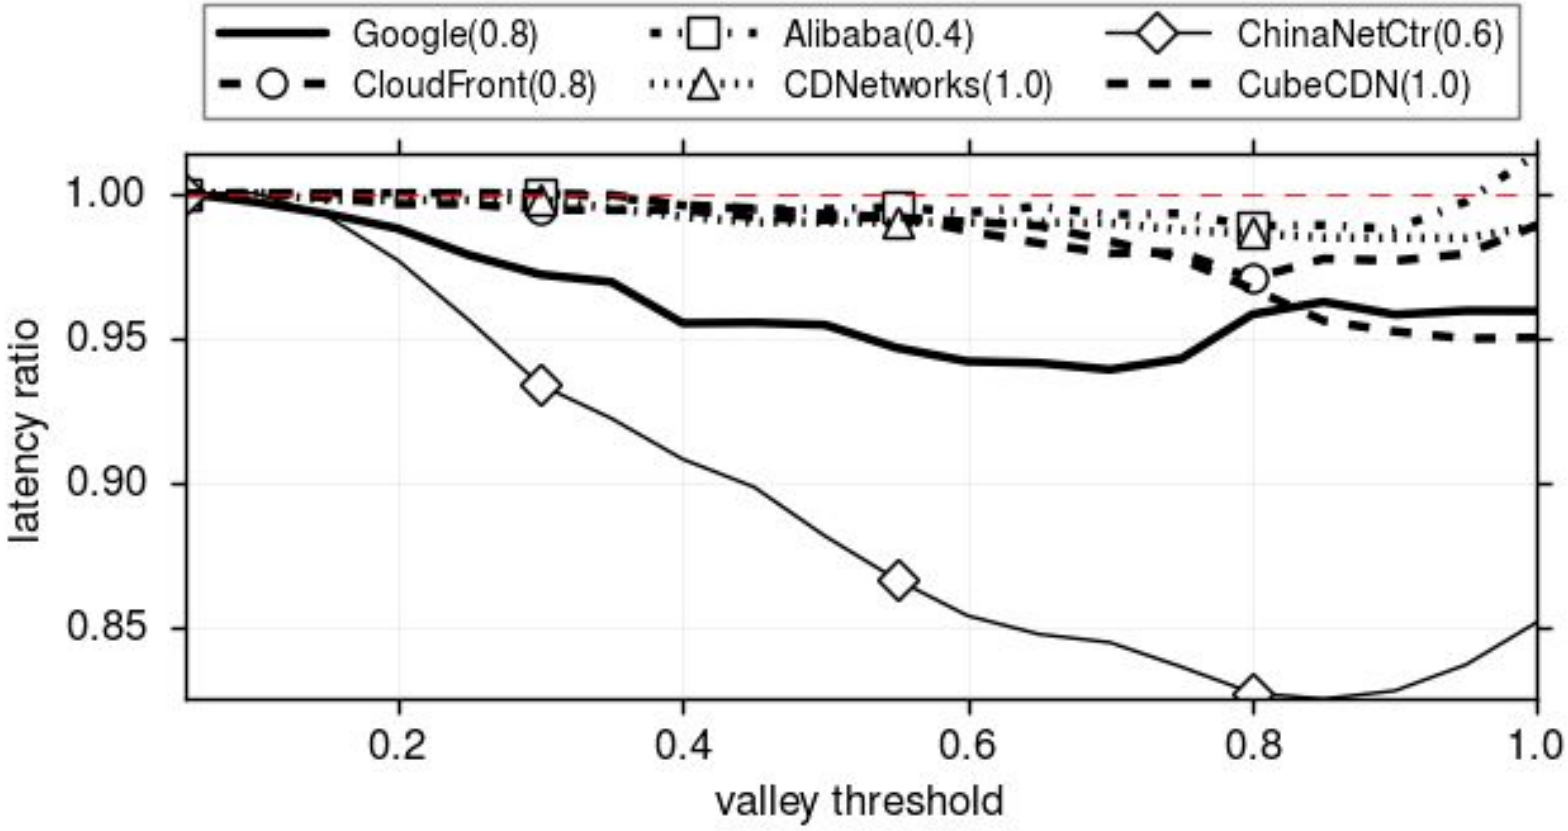

Performance of Drongo's choices