Anycast, Inflation, and Efficiency in the Root DNS

Tom Koch, Ethan Katz-Bassett, Ke Li (Columbia)
Matt Calder (Microsoft, Columbia)
John Heidemann and Calvin Ardi (USC/ISI)

Adapted from Anycast in Context: A Tale of Two Systems. Tom Koch (Columbia University); Ke Li (Columbia University); Calvin Ardi (ISI); Matt Calder (Microsoft, Columbia University); John Heidemann (ISI); Ethan Katz-Bassett (Columbia University). To appear in ACM SIGCOMM 2021.

Overview

- Views of anycast performance differ
  - Some say anycast increases latency [1].
  - Others say anycast performs well [2].

What is going on?

- To understand these differences, we re-evaluate root DNS latency & efficiency.

- By placing *anycast in context*, we find poor anycast performance in the root DNS hardly matters to users.

[1] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet Anycast: Performance, Problems, & Potential. In Proceedings of the 2018 ACM SIGCOMM Conference (Budapest, Hungary).
[2] Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye. 2015. Analyzing the Performance of an Anycast CDN. In Proceedings of the 2015 Internet Measurement Conference (IMC) (Tokyo, Japan).

Motivation

Background and Definitions

Inflation in the Root DNS

Performance in Context

Expansion in the Root DNS

Anycast in Practice

Many systems use anycast

1. DNS (roots, Akamai, Cloudflare, some public DNS)
2. CDNs (Microsoft, Verizon, Cloudflare)
3. Google

Deployments are growing

# of Sites

	2015	Today
Microsoft's CDN	30	110
Root DNS	500	1400

Deployments host latency-sensitive services

Google Cloud VMs can host game engines!

Anycast Gets a Bad Rap

Mixed Reviews

- (Li et. al., SIGCOMM 2018) “While it is not surprising that IP anycast is suboptimal … we find [anycast’s] inefficiencies to be surprisingly excessive.”
    - Deployment investigated: the root DNS.
- (Calder et. al., IMC 2015) “For most clients, anycast performs well despite the lack of centralized control.”
    - Deployment investigated: Microsoft’s anycast CDN.

Questions:

Why do latency-sensitive services use anycast if it hurts performance?

Why are those deployments expanding if there’s “excessive inefficiency”?

Contributions

Analyzing the root DNS using global traces we show that

1. Inefficiency in the root DNS is common, but recursive resolvers preferentially query low-latency root servers helping performance.
2. Latency *hardly matters* for root DNS since caching is so effective.
3. Inflation is not a great indicator of performance, as adding sites is more important for lowering latency.

Motivation

Background and Definitions

Inflation in the Root DNS

Performance in Context

Expansion in the Root DNS

What is IP Anycast?

An approach to routing, where distinct servers (sites) all use the same IP address and serve the same content.

The set of sites is the *deployment*.

Routes to sites are selected via BGP.

Inflation: BGP does not incorporate performance into its decision-making process, so users may take unnecessarily high latency paths.

Potential Benefits of Anycast

1. Simple, scalable content distribution.
2. Seamless handling of certain types of site failure / route withdrawals.
3. DDoS protection.

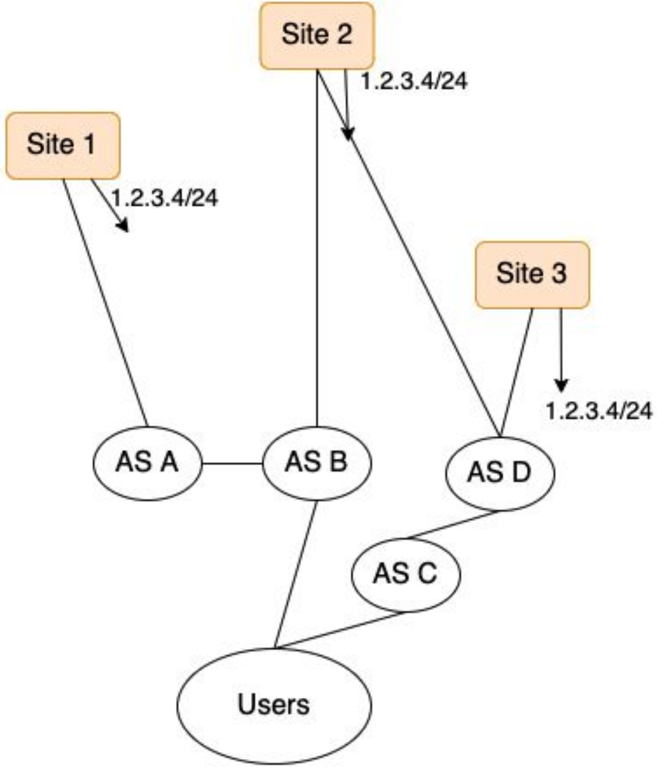

The route from a network of users to 1.2.3.4/24 is selected (via BGP) to be the route through AS B to Site 2.

Users incur 60 ms reaching site 2, even though site 3 offers lower latency.

How Prior Work Calculates Inflation

Prior Work [1-4]: Inflation = Achieved - best alternative = 60 ms - 20 ms = 40 ms

Limitations of this approach:

1. Measuring the “20 ms” path would sacrifice coverage.

2. The “20 ms” path may still be inflated!

[1] Hitesh Ballani and Paul Francis. 2005. Towards a Global IP Anycast Service. In Proceedings of the 2005 ACM SIGCOMM Conference (Philadelphia, PA, USA).
[2] Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye. 2015. Analyzing the Performance of an Anycast CDN. In Proceedings of the 2015 Internet Measurement Conference (IMC) (Tokyo, Japan).
[3] Ricardo de Oliveira Schmidt, John Heidemann, and Jan Harm Kuipers. 2017. Anycast Latency: How Many Sites Are Enough? In International Conference on Passive and Active Network Measurement (PAM) (Sydney, Australia).
[4] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet Anycast: Performance, Problems, & Potential. In Proceedings of the 2018 ACM SIGCOMM Conference (Budapest, Hungary).

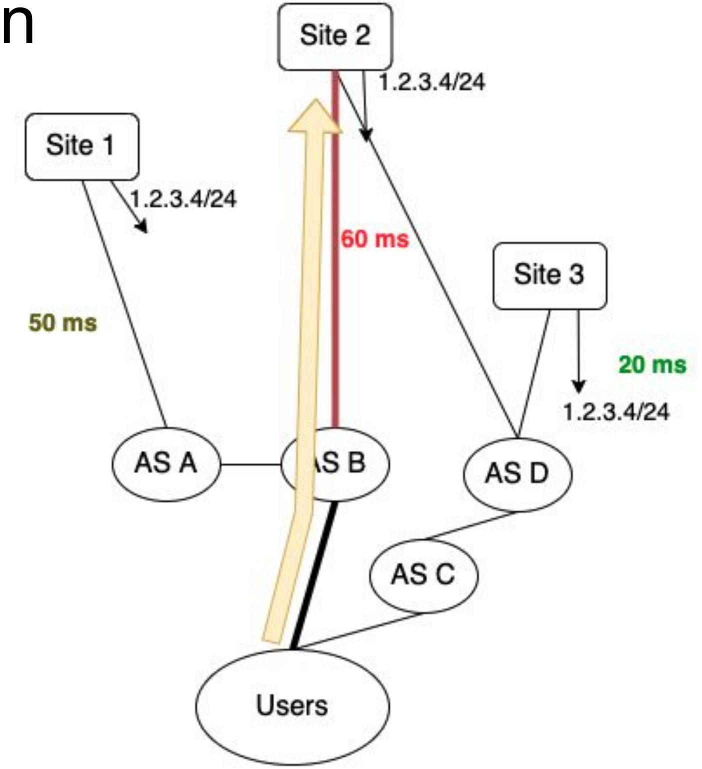

Users incur 40 ms extra compared to the best alternative.

How We Calculate Inflation

1. **Geographic Inflation:** Extra Distance

(2,800 km - 1,200 km = 1,600 km)

2. **Latency Inflation:** Extra Latency over Reasonable Lowest Latency

(60 ms - 1,200 km * 1.5 * 1 ms / 200 km = 51 ms)

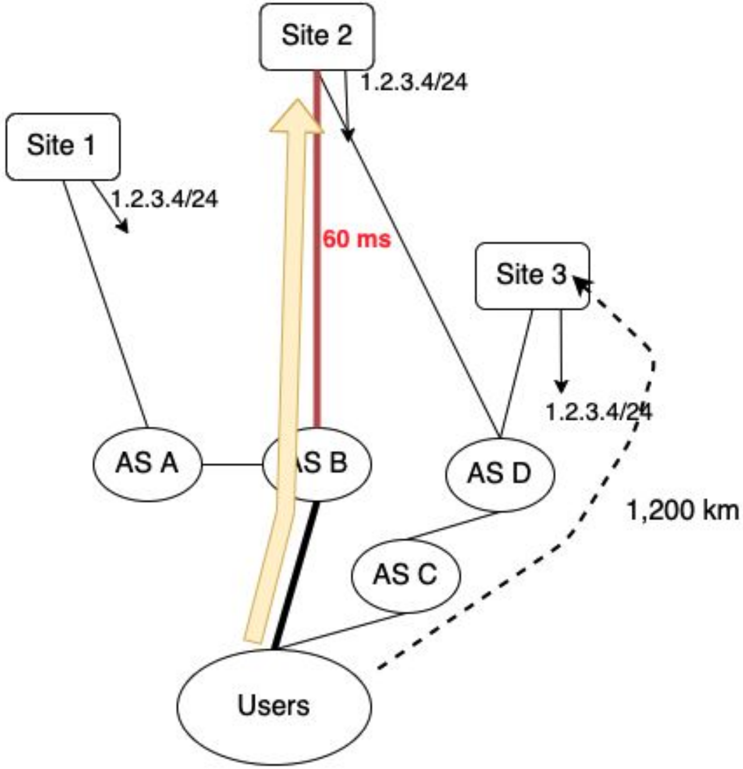

Data Sources

We use the DITL captures from 2018 [1], since they show us global querying behavior and latency to many root letters.

- The 2020 DITL is less complete, and further-anonymized, limiting analysis.

We use user count data from Microsoft.

- Microsoft approximates user counts with # of unique IP addresses

[1] DNS-OARC. 2018. A Day in the Life of the Internet. dns-oarc.net/oarc/data/ditl/2018

Motivation

Background and Definitions

Inflation in the Root DNS

Performance in Context

Expansion in the Root DNS

Inflation in the Root DNS and Microsoft's CDN

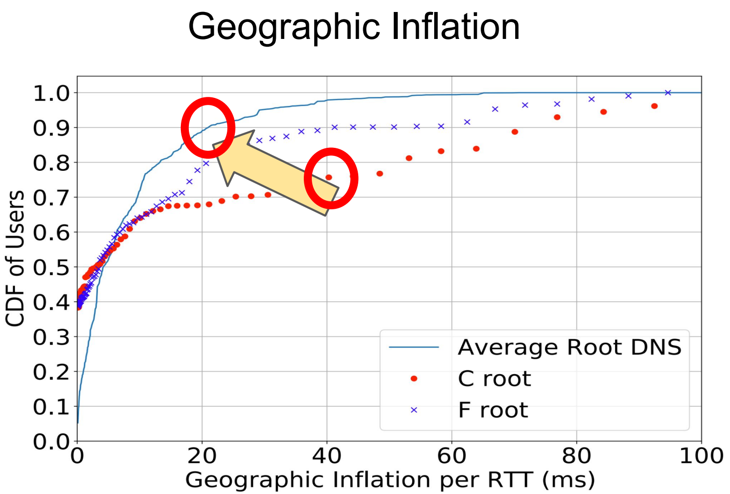

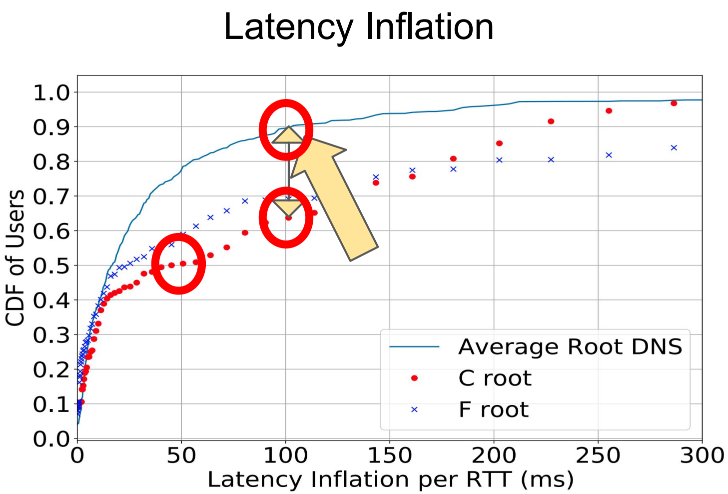

1. Individual root letters have lots of inflation (e.g., C root).
2. Average root DNS inflation is not as bad as in individual letters.
3. Studying **both types** of inflation is valuable.

Measuring Inflation in Context

How does inflation impact users?

- For root DNS, one thinks “not a lot” due to caching.
- But there is **a lot** of attention placed on root DNS latency, and deployments are growing.

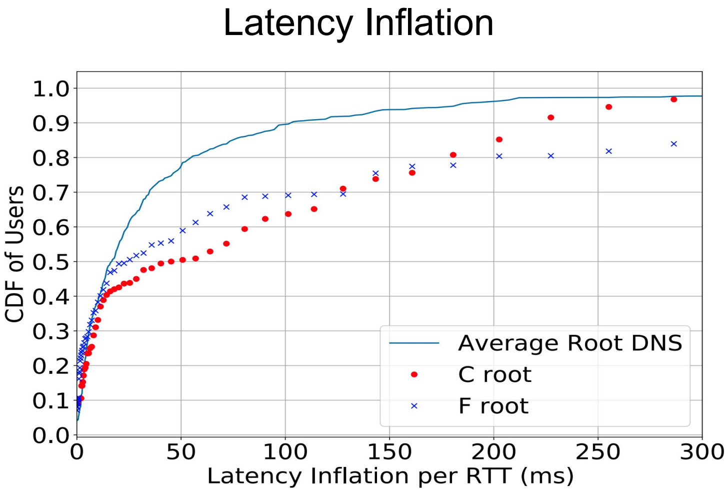

Motivation

Background and Definitions

Inflation in the Root DNS

Performance in Context

Expansion in the Root DNS

Root DNS Latency and Inflation Hardly Matter

We amortize root DNS queries over users to get queries per day.

$$Q_{day}(RR) = \frac{N_Q(RR)}{N_U(RR)}$$

# of queries per RR in the root traces

# of Microsoft users using recursive

Takeaway: Most users rarely interact with the root DNS, but interact far more than one might expect.

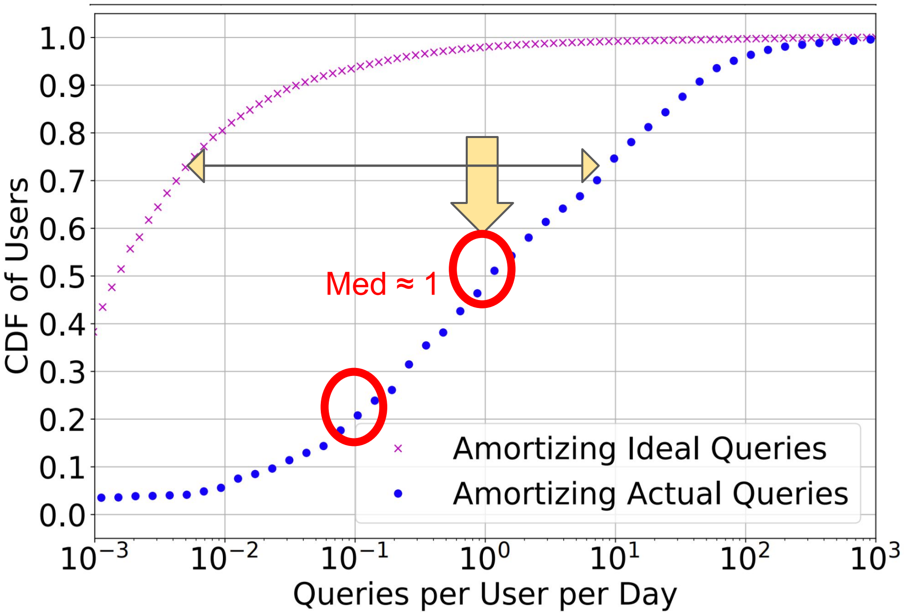

Root latency hardly matters to users, but 'Ideal' shows that there are many unnecessary queries to the root DNS.

Motivation

Background and Definitions

Inflation in the Root DNS

Performance in Context

Expansion in the Root DNS

Larger Deployments Lower Latency *and* Efficiency

Fewer users visit their closest site in larger deployments, but the site users hit is pretty close.

Having low latency options is more important than hitting the closest site.

Efficiency is a misleading metric!

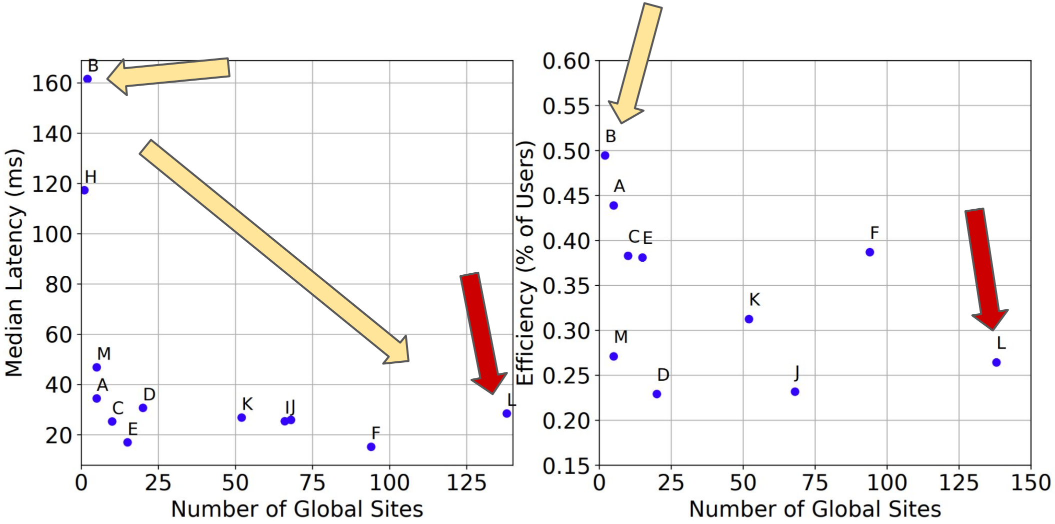

Efficiency is the percent of users that visit their closest site.
Larger deployments offer lower latency, but are less efficient.

What Caused All the Growth?

We asked root DNS operators why they expanded, and whether the expansion will continue.

Surprisingly, **many listed decreasing latency** as a reason for expanding.

Expansion will likely slow.

Past		Future	
Reason for Growth	**Number of Orgs**	**Future Growth Trend**	**Number of Orgs**
Latency	8	Acceleration of Growth	1
DDoS Resilience	9	Decceleration of Growth	4
ISP Resilience	5	Maintain Growth Rate	4
Other	3	Cannot Share	1

DITL -- Please Keep it Going!

DITL was an invaluable source of data since it offered us a global view of recursive querying behavior.

Thanks to DNS-OARC and all the root letters who participated!

The broad, global view of the root DNS *system* allowed us to piece together a coherent story.

Anonymization limits utility -- perhaps there are other ways to preserve privacy?

Conclusions

One should assess anycast **in context of the system**.

- Comparing root DNS latency and anycast inflation is unproductive.
- Root DNS operators should not deploy new sites to lower latency.

Inflation is a misleading metric -- absolute latency is more important.

Component performance is not representative of system performance.

Questions?