Charting the Expanse of the Cloud

Nitinder Mohan

mohan@in.tum.de

Cloud is expanding …

- More than 2000+ datacenters deployed globally
- Cloud + CDN traffic estimated to be 95% total traffic
- Top 3 cloud providers control 65% of cloud infrastructure

Google Cloud Platform

Azure

aws

2012

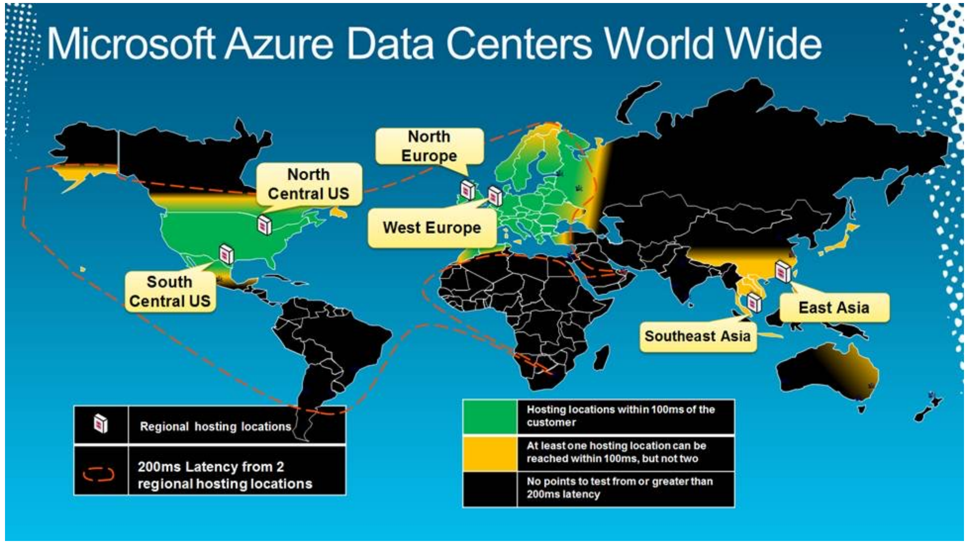

https://joranmarkx.wordpress.com/2012/01/16/microsoft-azure-data-center-locations-worl-wide/

2022

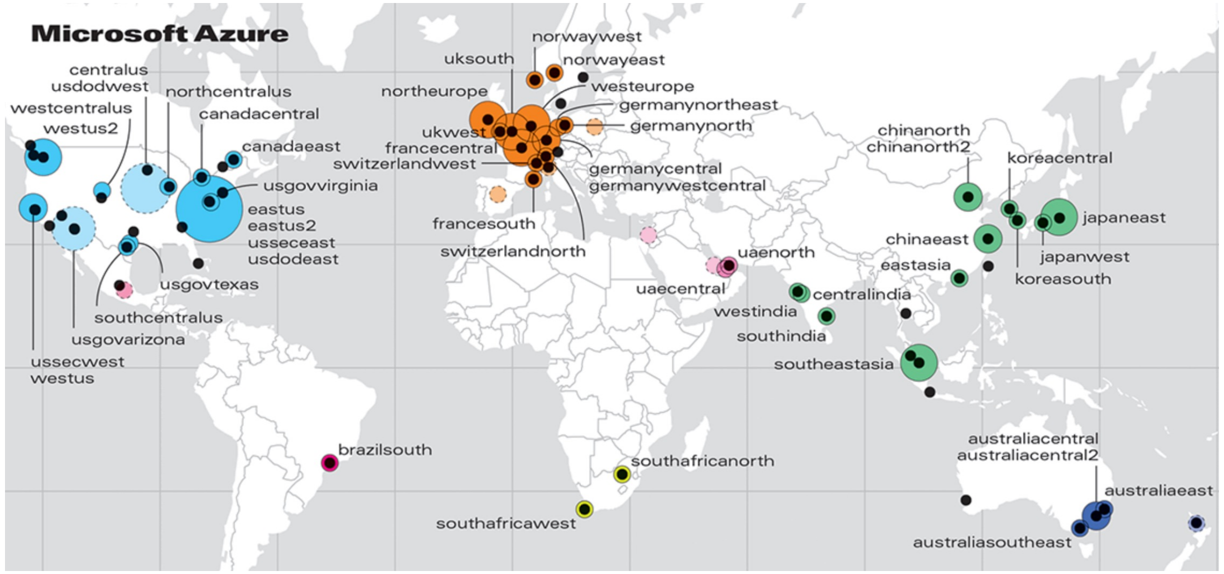

https://infrastructuremap.microsoft.com/explore

… but it might be dispersing

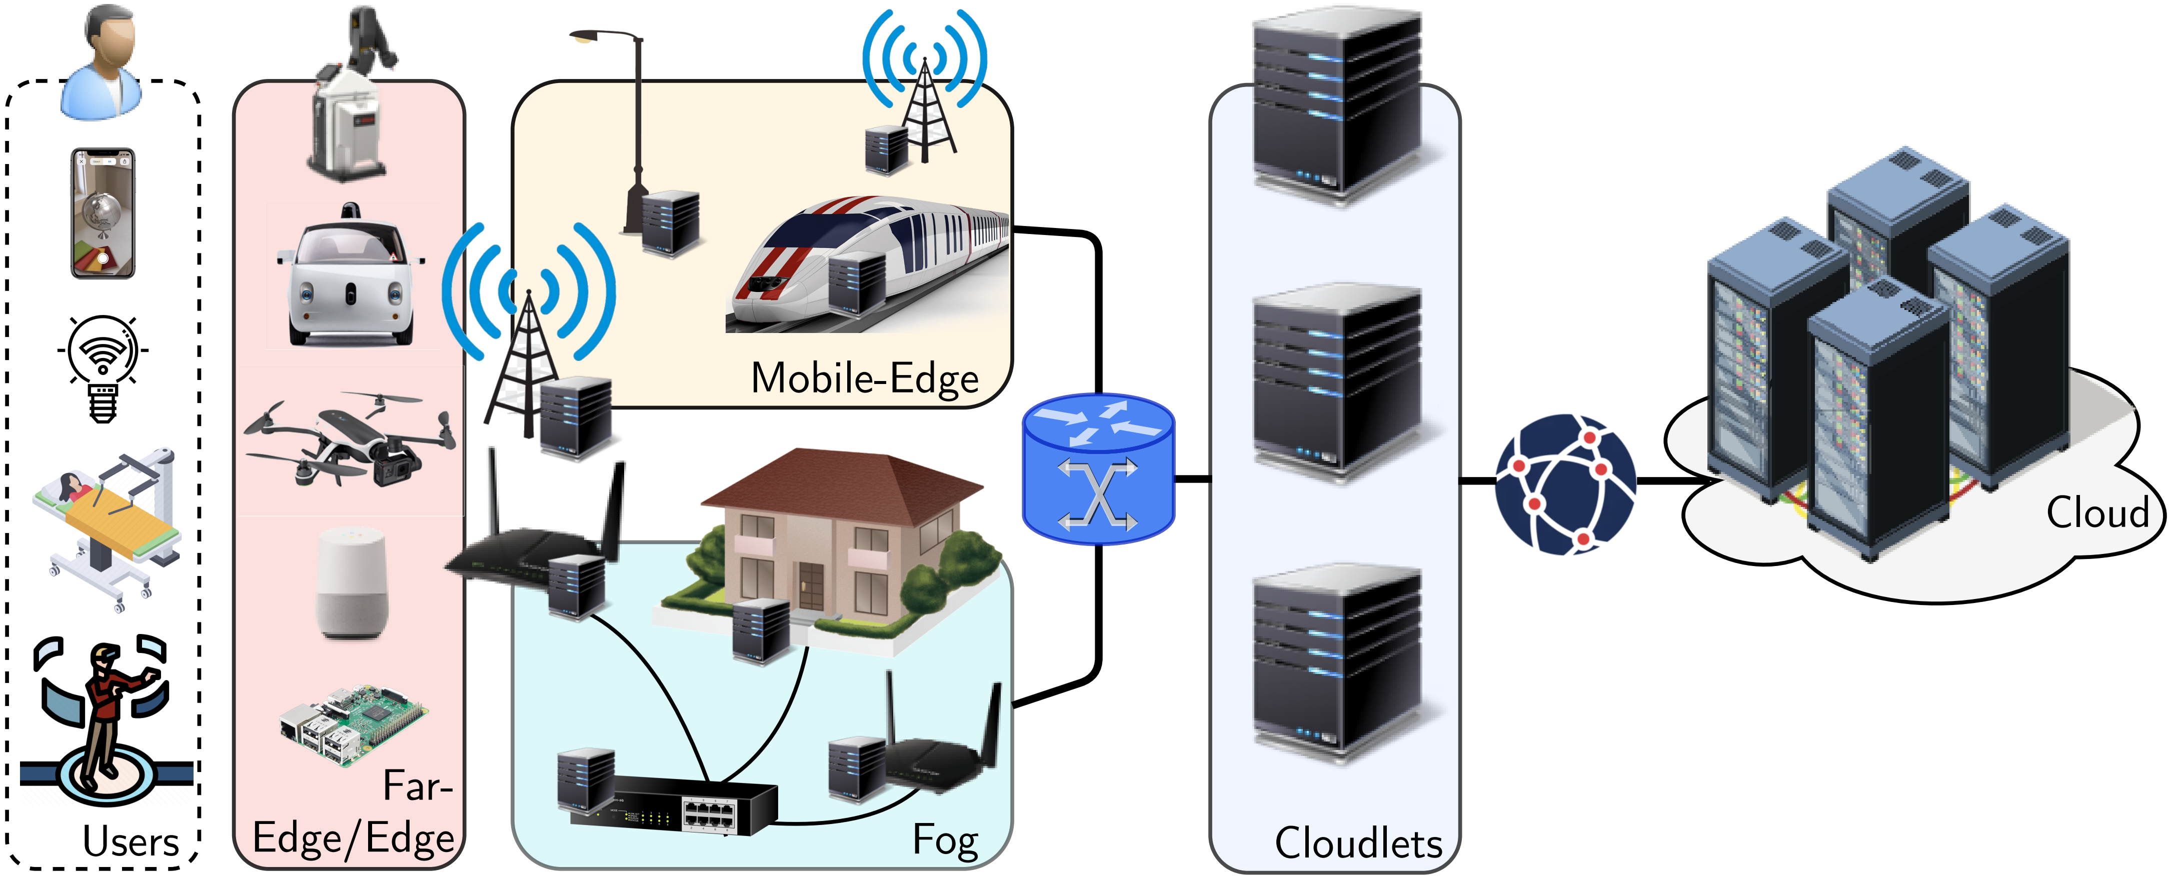

Need for Short Latencies for Applications

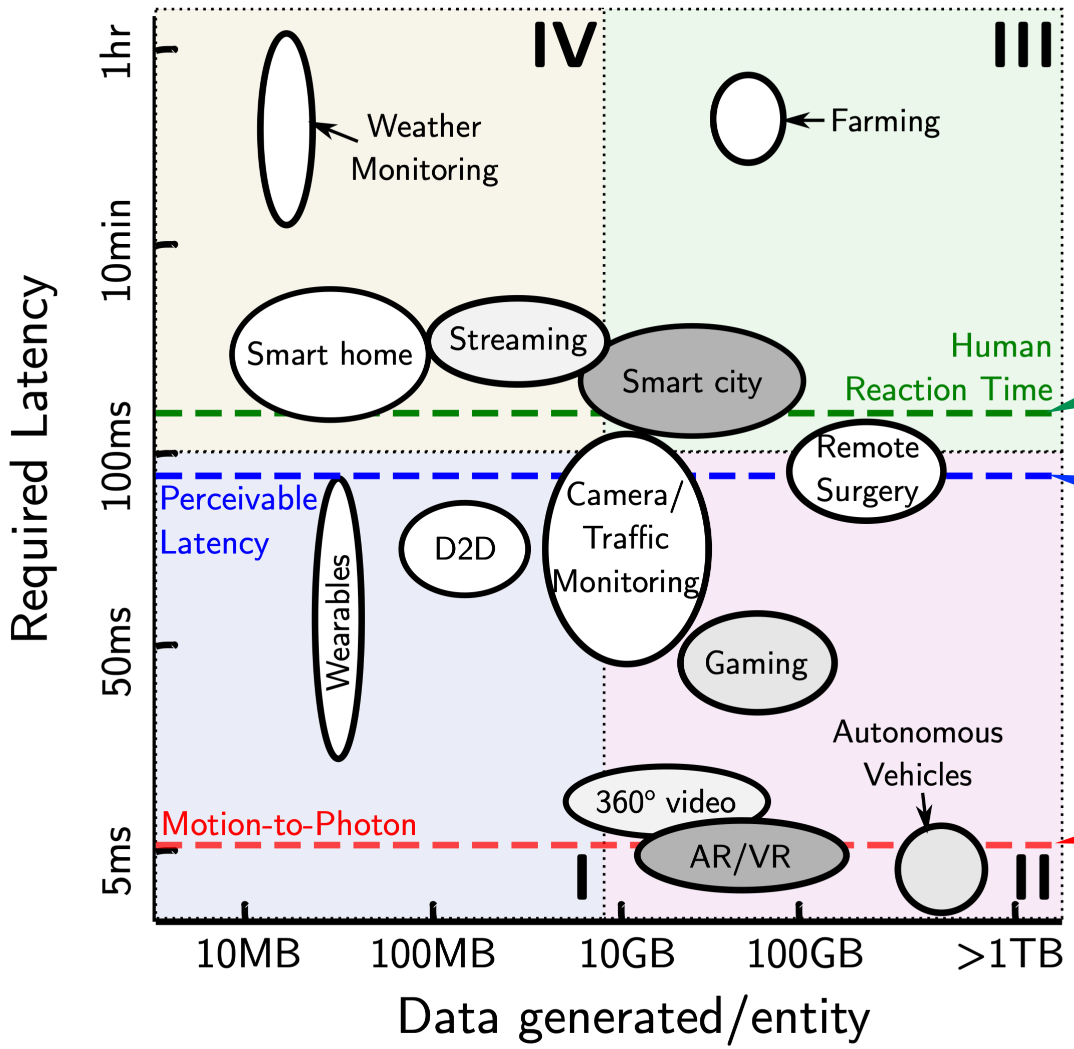

Human Reaction Time (HRT) is delay difference between visual stimulus and associated motor response

- Typically less than 250ms end-to-end
- Required for remote human-in-the-loop applications

Preceivable Latency (PL) is latency threshold when the user starts to experience lags in the application

- Typically within 80-100 ms end-to-end
- Strict requirement for interactive applications

Motion-to-Photon (MTP) is delay between user input and its effect reflecting on the display

- Typically within 10-20 ms end-to-end
- Strict requirement for immersive applications

Mohan, Nitinder, et al. "Pruning edge research with latency shears." *In ACM HotNets* 2020.

Can the cloud and centralization cope with growing application requirements?

Methodology → End-Points

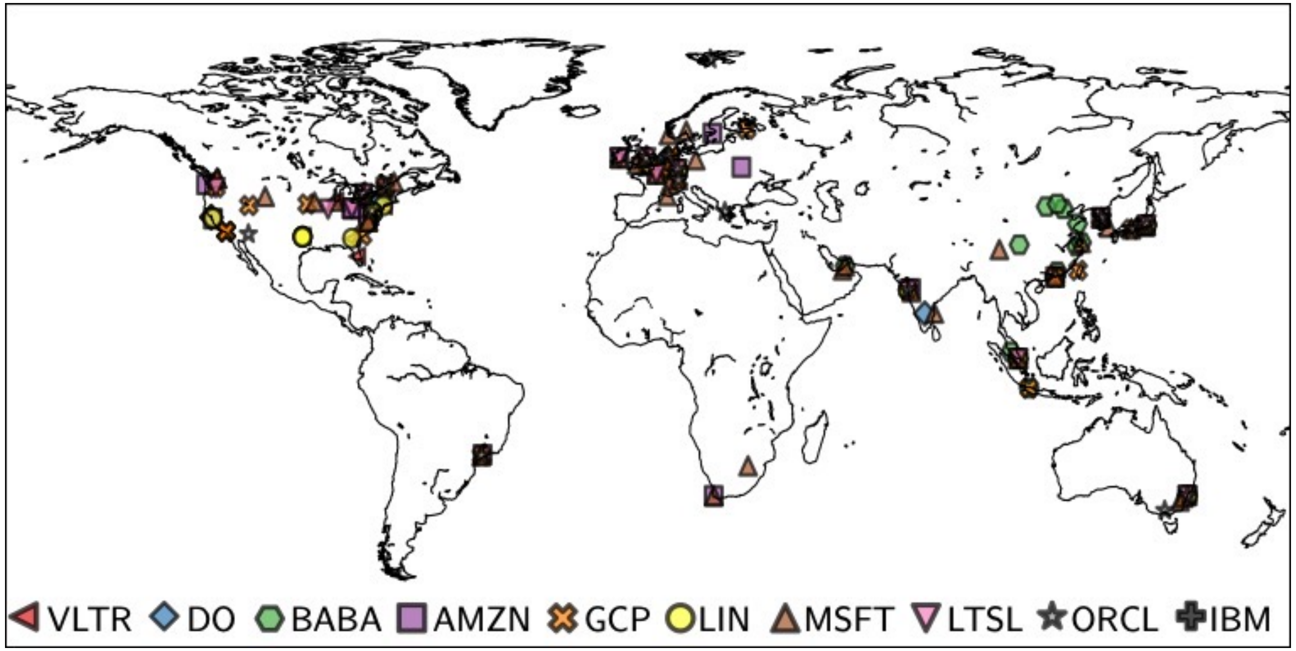

	Datacenters per continent						Backbone
	EU	NA	SA	AS	AF	OC	N/W
Amazon EC2 (AMZN)	6	6	1	6	1	1	Private
Google (GCP)	6	10	1	8	-	1	Private
Microsoft (MSFT)	14	10	1	15	2	4	Private
Digital Ocean (DO)	4	6	-	1	-	-	Semi
Alibaba (BABA)	2	2	-	16	-	1	Semi
Vultr (VLTR)	4	9	-	1	-	1	Public
Linode (LIN)	2	5	-	3	-	1	Public
Amazon Lightsail (LTSL)	4	4	-	4	-	1	Private
Oracle (ORCL)	4	4	1	7	-	2	Private
IBM (IBM)	6	6	-	1	-	-	Semi
Total	**52**	**62**	**4**	**62**	**3**	**12**	

- 195 cloud regions operated by 9 major cloud providers of different sizes
- Public VM in each *compute-services* specific region
- Operators have different WAN deployments to support their networking globally

Methodology → Vantage Points

- More than 8000 RIPE Atlas probes
- Deployed in 184 countries globally
- Installed in heterogeneous environments, e.g. homes, offices
- Fixed network access
- Discarded probes installed in network operator/ cloud infrastructure

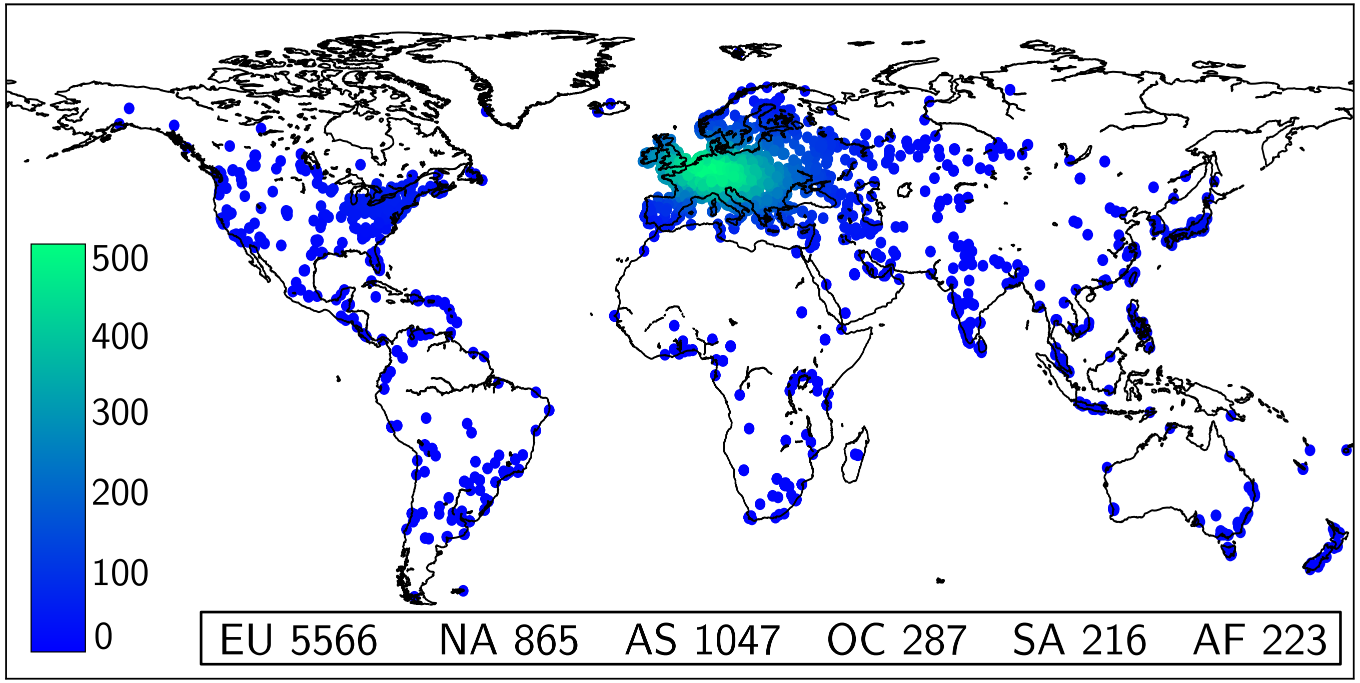

RIPE NCC
RIPE Atlas

Lorenzo Corneo, Maximilian Eder, Nitinder Mohan, Aleksandr Zavodovski, Suzan Bayhan, Walter Wong, Per Gunningberg, Jussi Kangasharju, and Jörg Ott. 2021. Surrounded by the Clouds: A Comprehensive Cloud Reachability Study. In Proceedings of the Web Conference 2021 (WWW '21)

Methodology → Vantage Points

- More than 115,000 Speedchecker (SC) probes
- Deployed in 140 countries globally
- Probes are softwareized applications running on real end user's devices

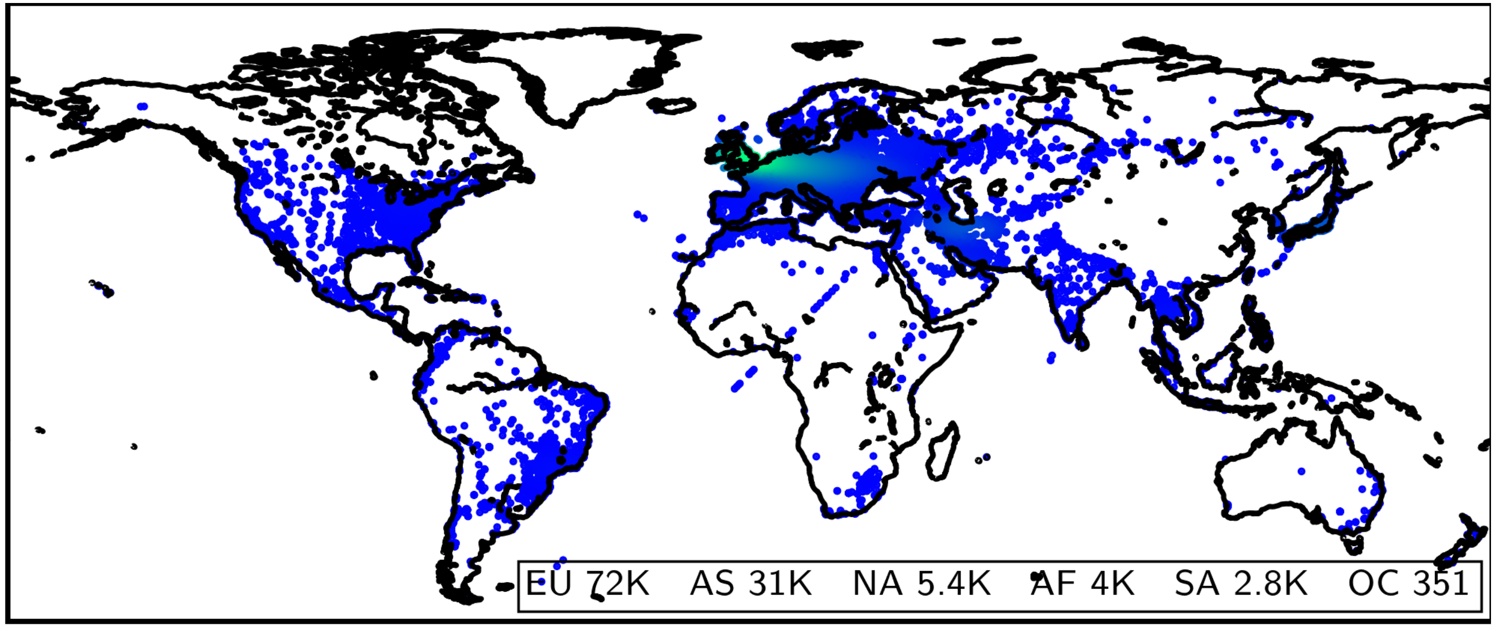

- Covers 95.6% of Internet population (APNIC [5])

SPEEDCHECKER

The Khang Dang, Nitinder Mohan, Lorenzo Corneo, Aleksandr Zavodovski, Jörg Ott, and Jussi Kangasharju. 2021. Cloudy with a chance of short RTTs: analyzing cloud connectivity in the internet. In Proceedings of the 21st ACM Internet Measurement Conference (IMC '21)

Methodology → Vantage Points

- More than 115,000 Speedchecker (SC) probes
- Deployed in 140 countries globally
- Probes are softwareized applications running on real end user's devices

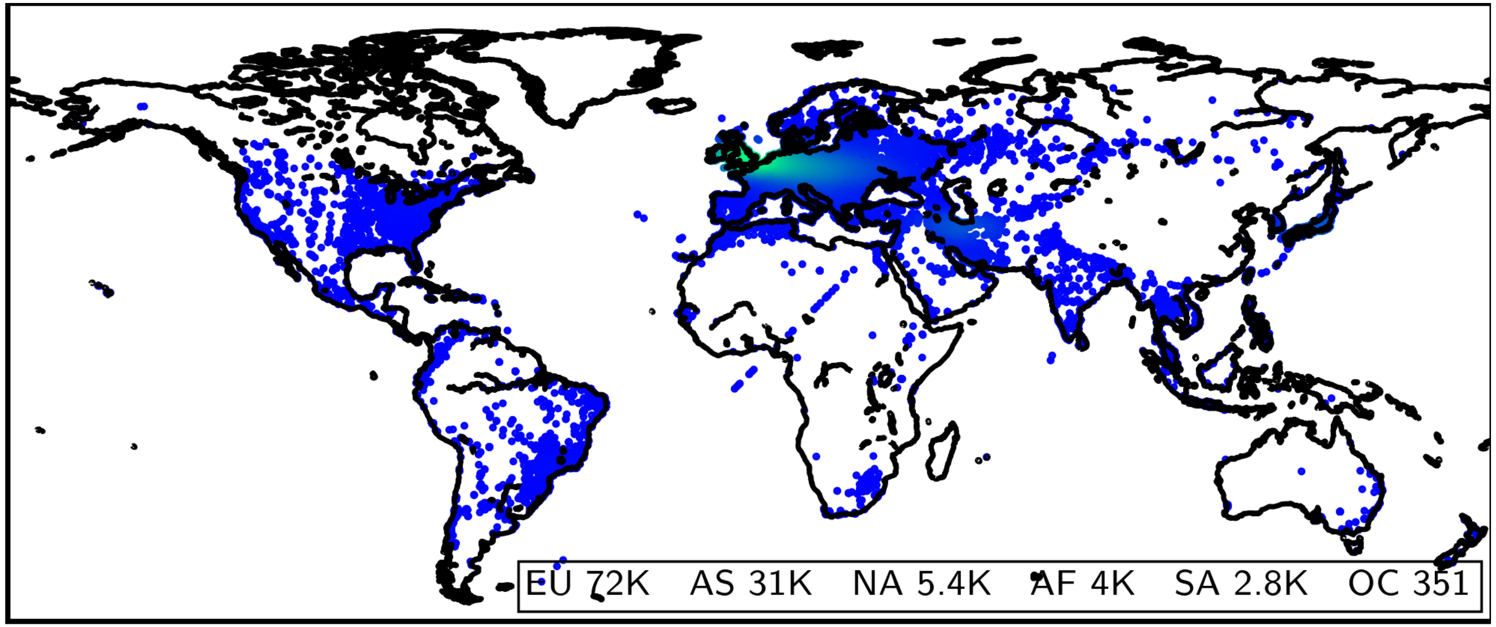

- Covers 95.6% of Internet population (APNIC)

SPEEDCHECKER

Comparison with RIPE Atlas:

9K vs 115K total probes

wired vs. wireless

69.2% vs 95.6% Internet users

SC probe density: 12x(EU) 6x(NA) 30-40x(rest)

8K+ vs 12K ASes

Methodology → Experiments

- From VPs to cloud regions in same continent every three hours
  - ICMP and TCP `pings`
  - ICMP and TCP Paris `traceroute`
- For Africa and South America also ran experiments towards nearest continents, i.e. Europe and North America
- Recorded 8M `ping` and 11+M `traceroute` for **approximately two years** (September 2019 – April 2021)

Dataset and scripts publicly available at cloudreachability.github.io

Global Cloud Access Latency

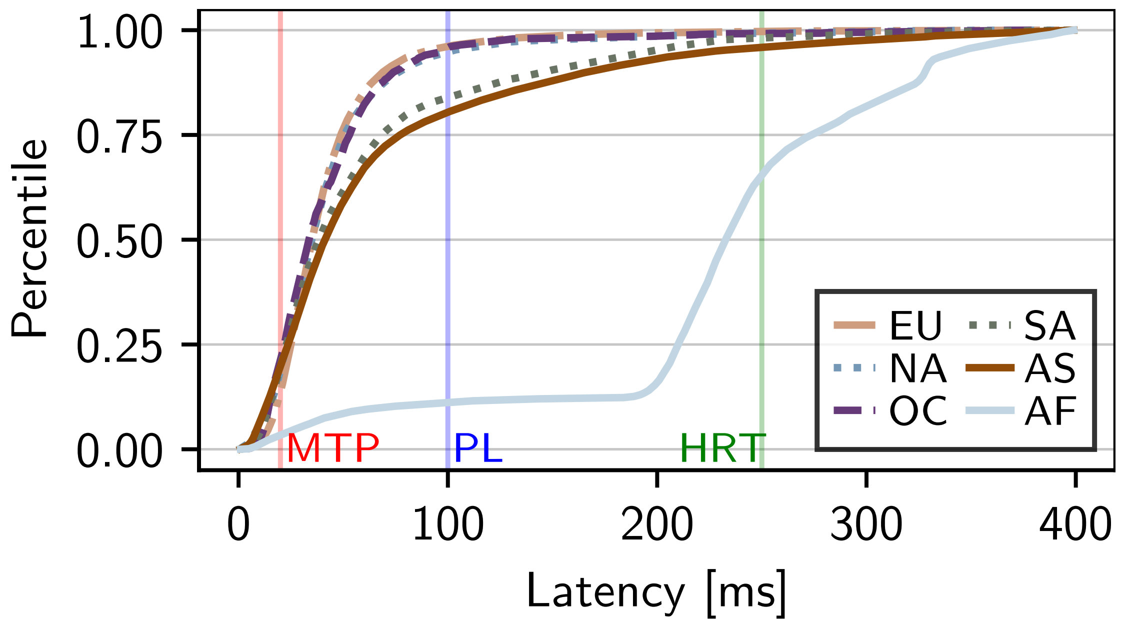

- Achieving MTP thresholds (≤ 20ms) is difficult for real Internet users globally
- Perceivable latency (≤ 100ms) is consistenly achievable in Europe, Oceania and North America

Global Cloud Access Latency

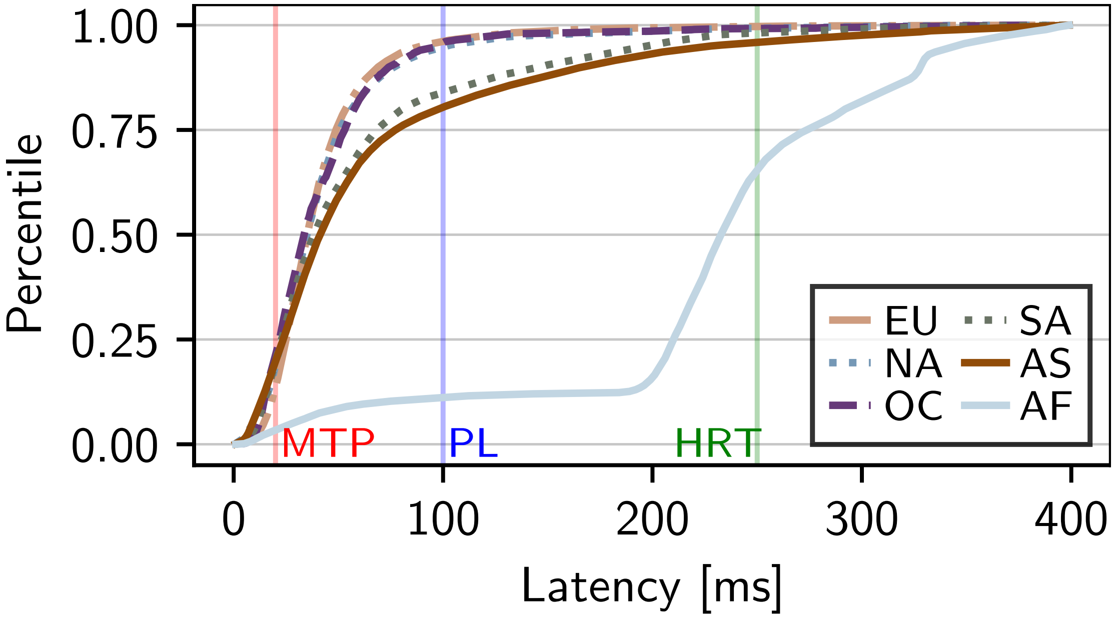

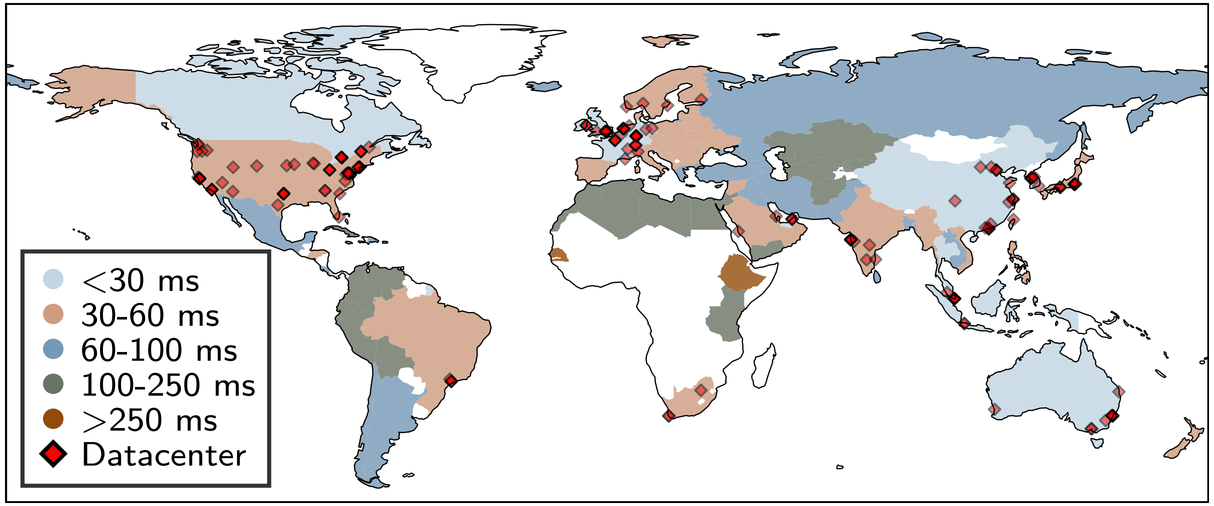

- Achieving MTP thresholds ($\leq$ 20ms) is difficult for real Internet users globally
- Perceivable latency ($\leq$ 100ms) is consistenly achievable in Europe, Oceania and North America
- In South America and Asia, only probes geographically "close" to datacenters enjoy low latencies
- Severe under-provisioning of DCs in Africa results in extremely long latencies
  - Countries in north Africa have shorter latencies towards DCs in Europe than within continent

Where are you measuring from?

- Choice of measurement platform can significantly affect the takeaways.
- RIPE Atlas probes are generally faster than Speedchecker almost globally
- Differences mostly due to density and geographical availability
  - Africa: Atlas is more clustered around south physically closer to DCs
  - South America: SC density in Brazil is much higher than Atlas which drives the results

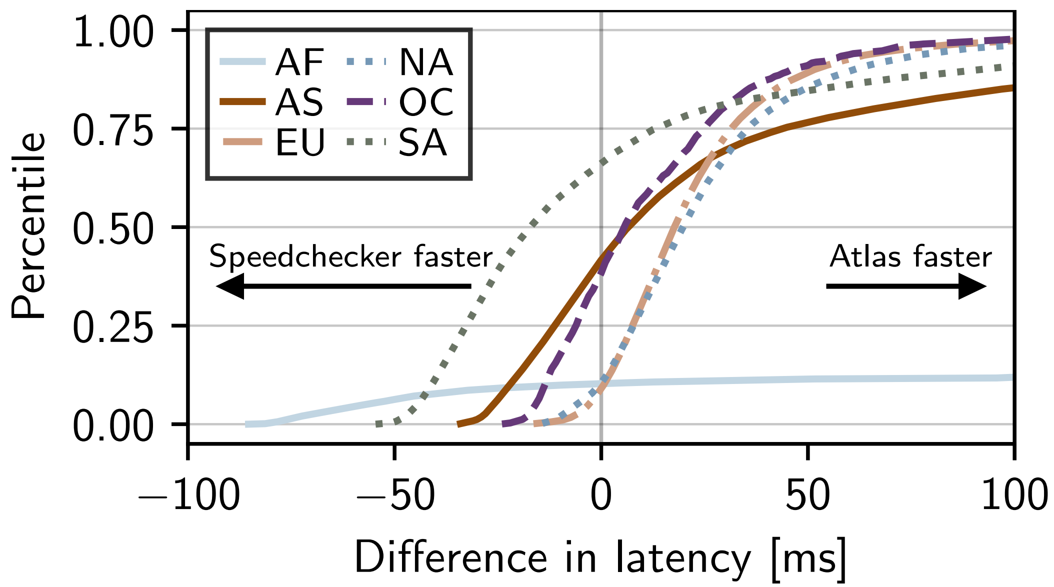

Know thy vantage points!

Routing and Peering with Cloud WANs

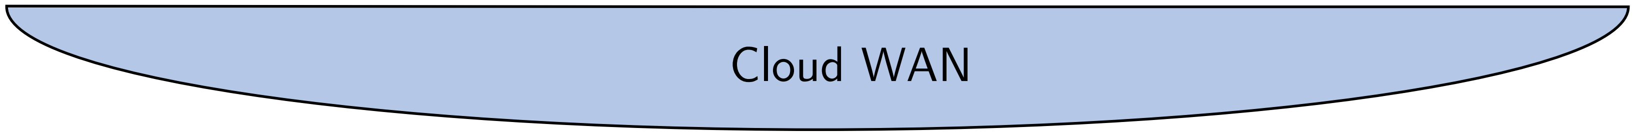

Routing and Peering with Cloud WANs

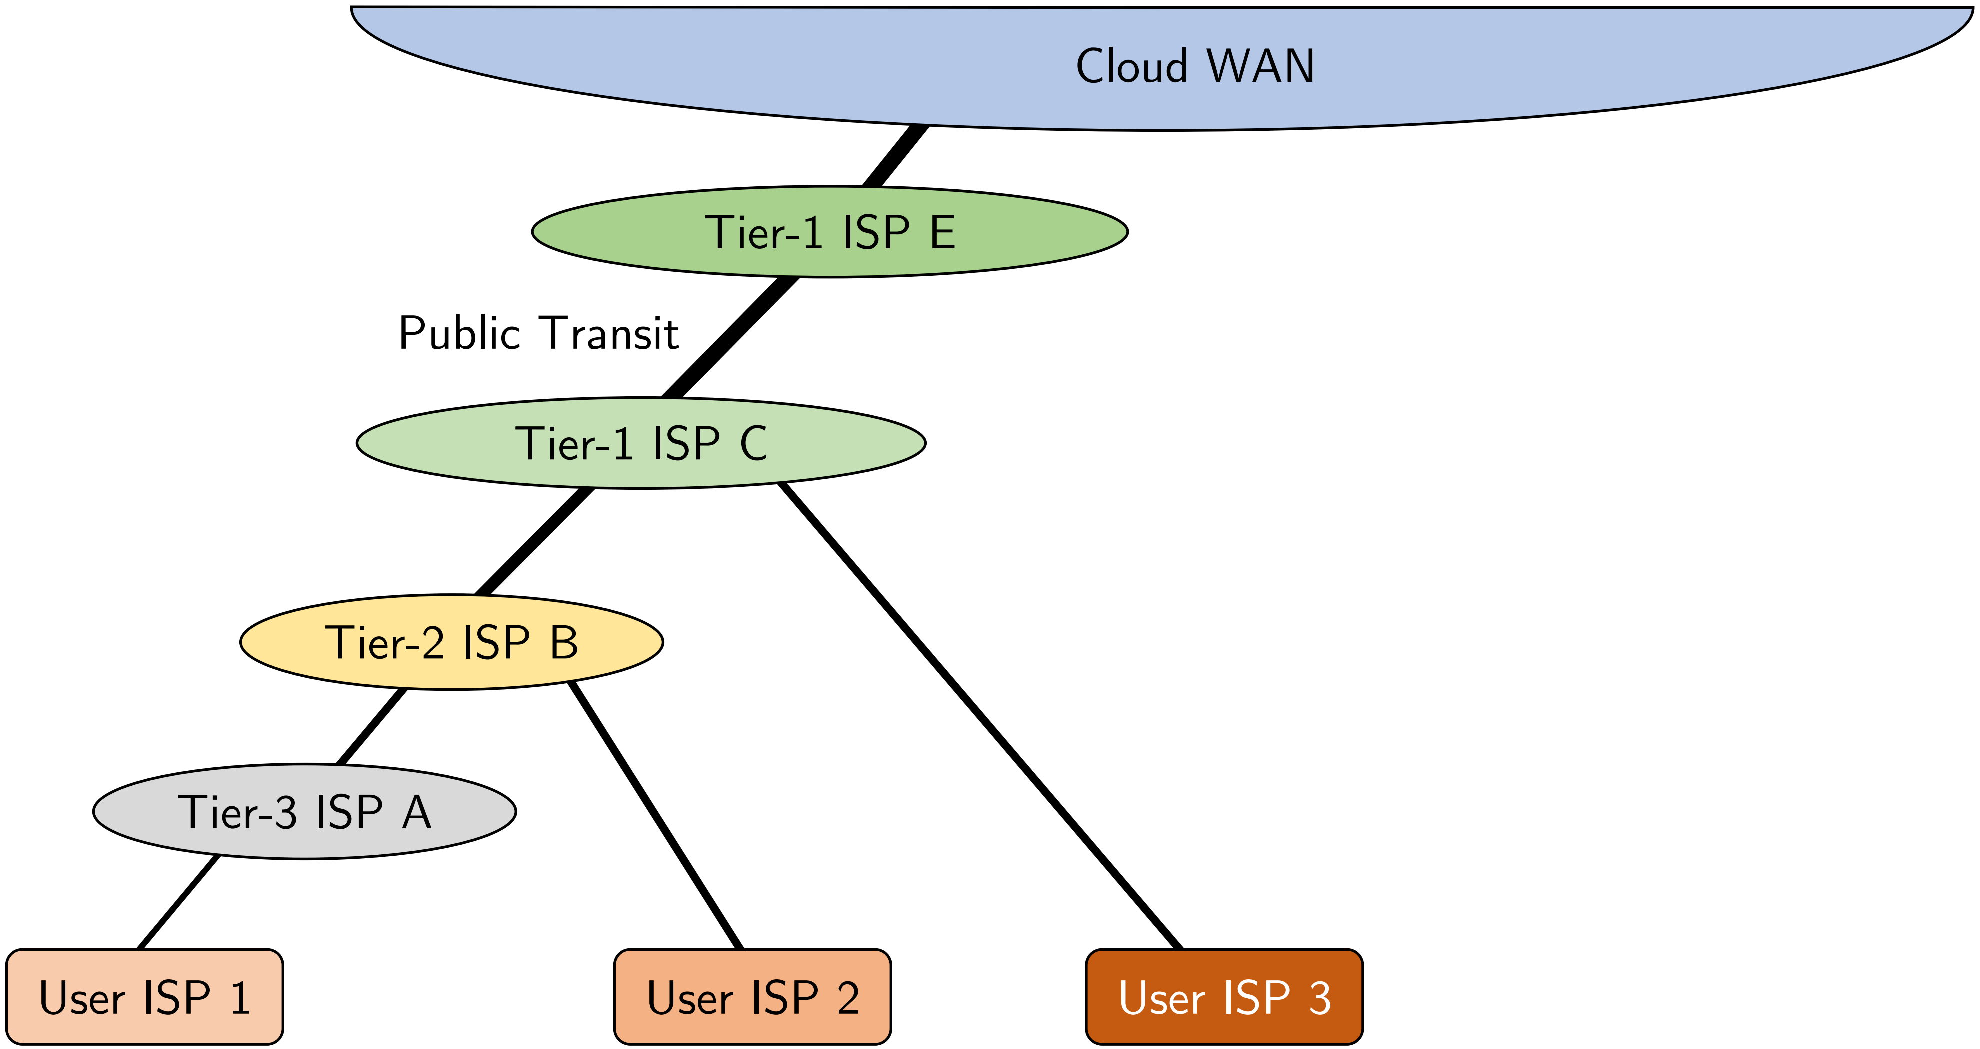

Routing and Peering with Cloud WANs

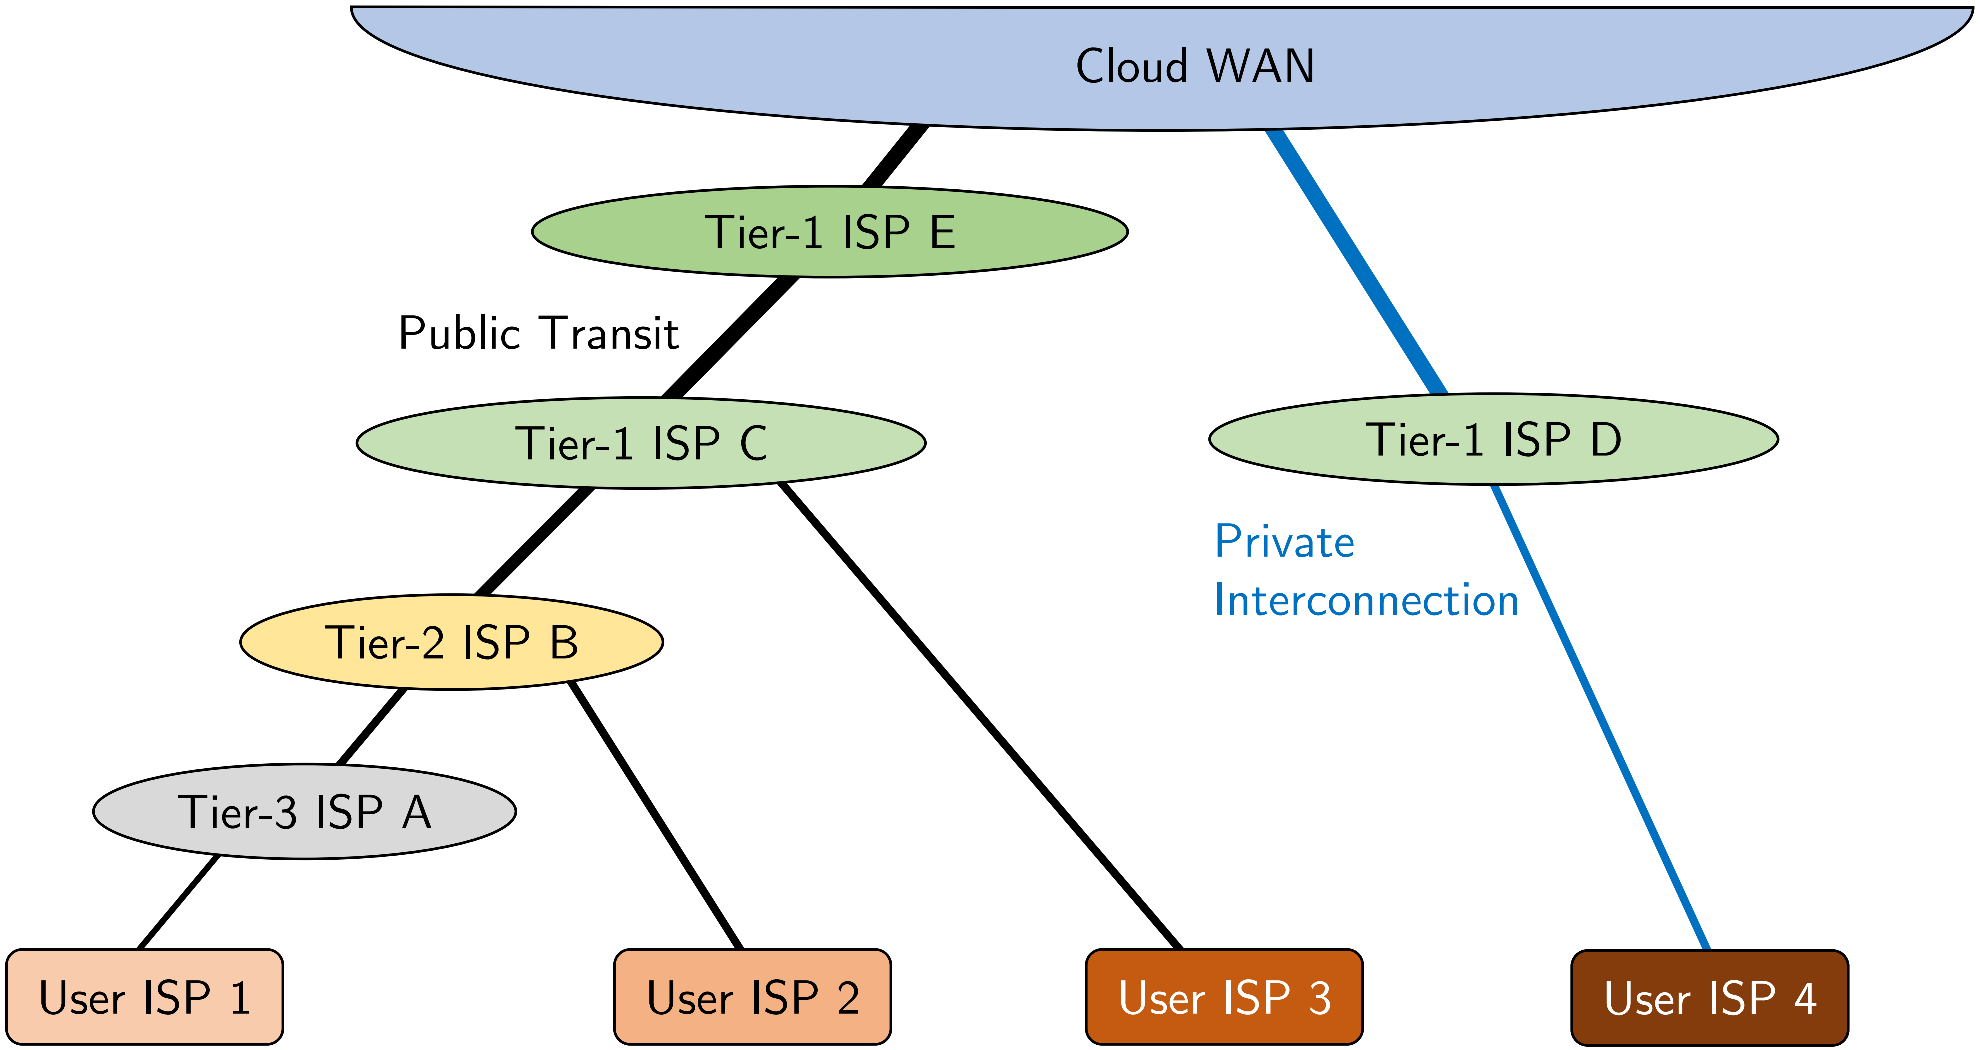

Routing and Peering with Cloud WANs

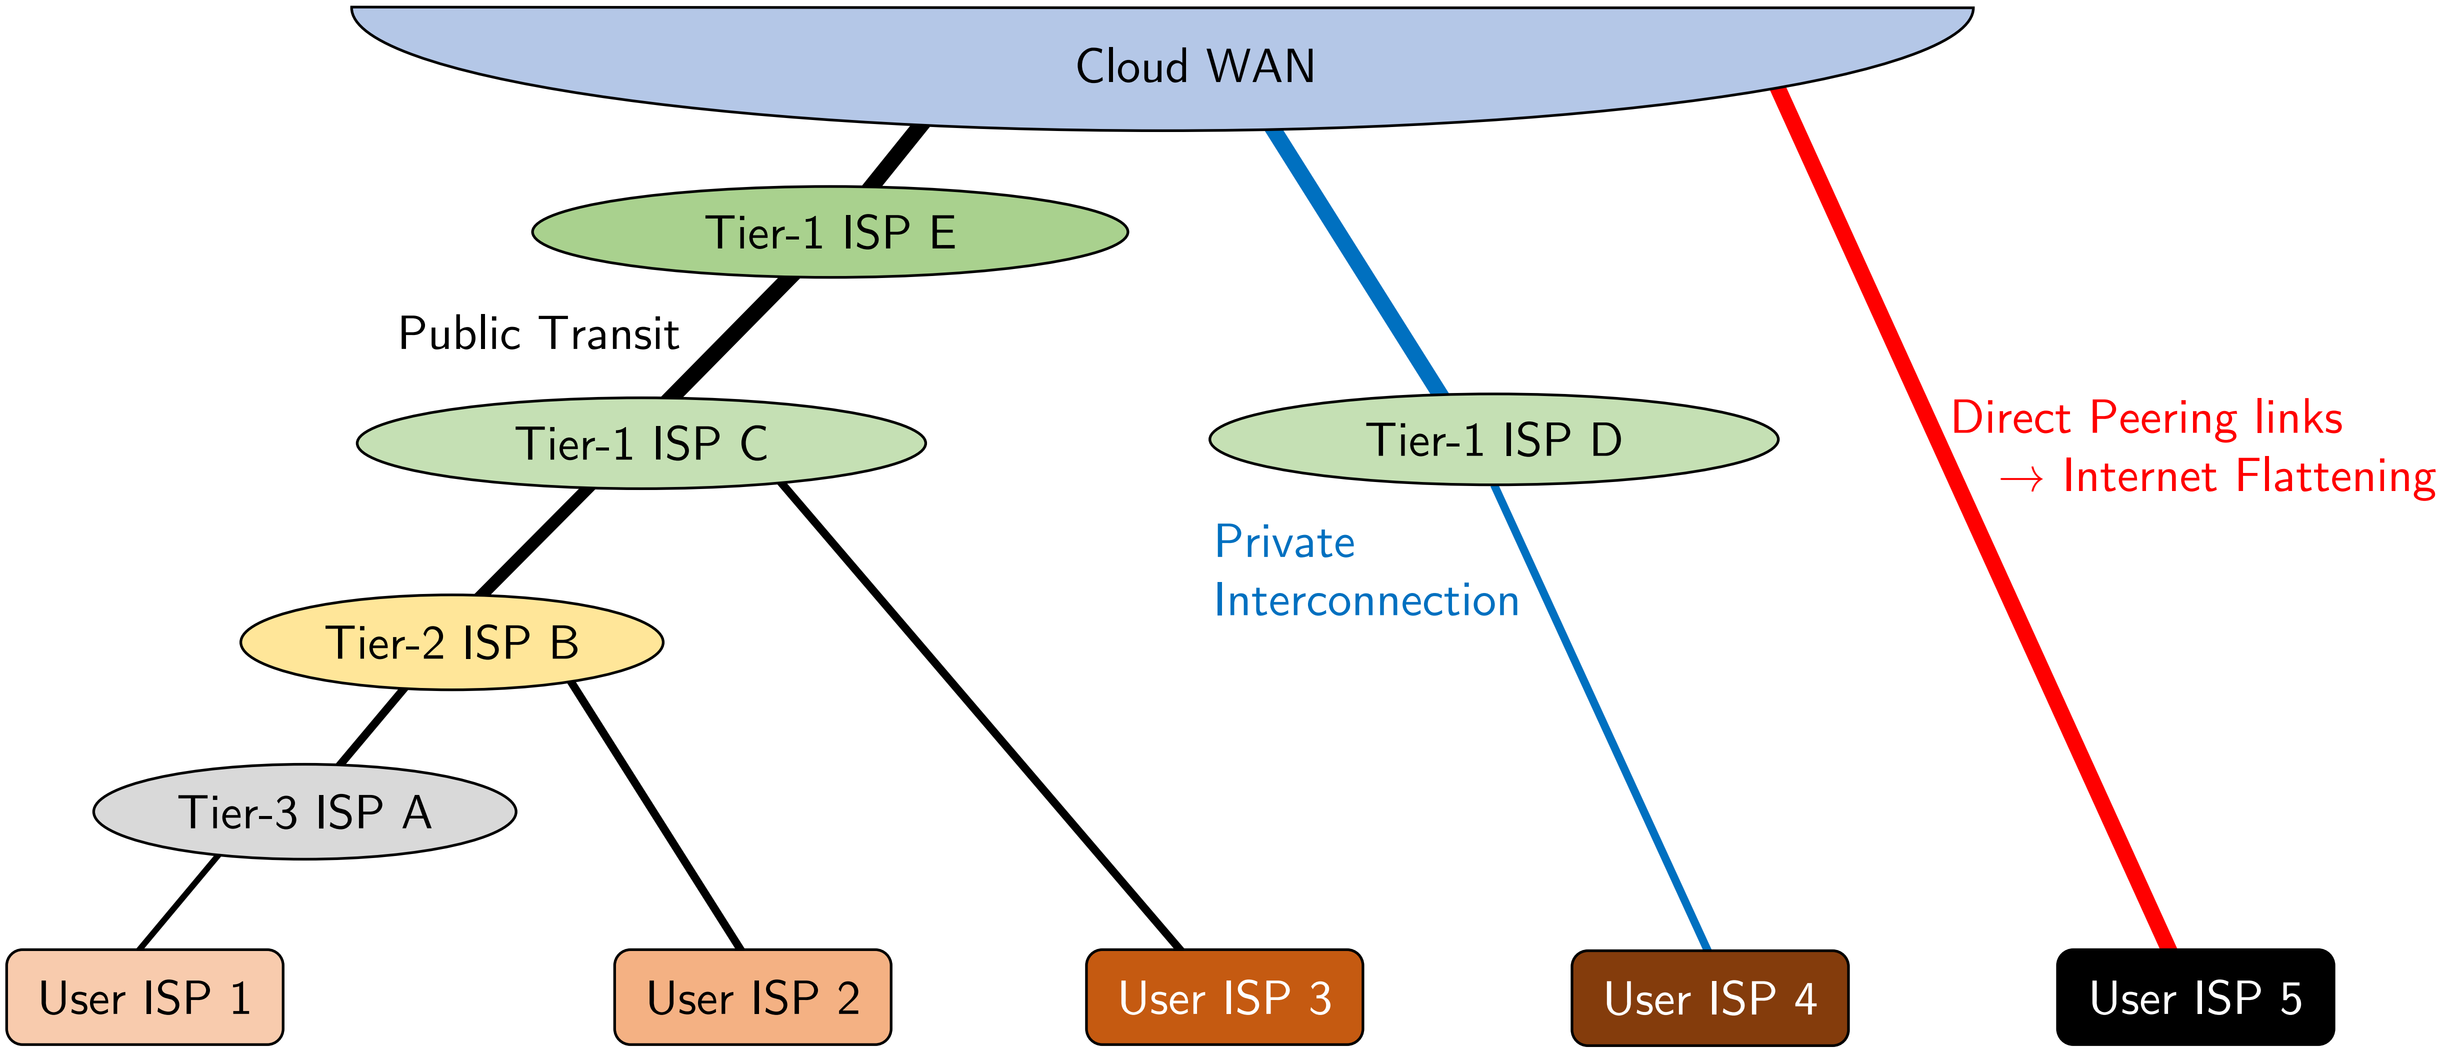

Routing and Peering with Cloud WANs

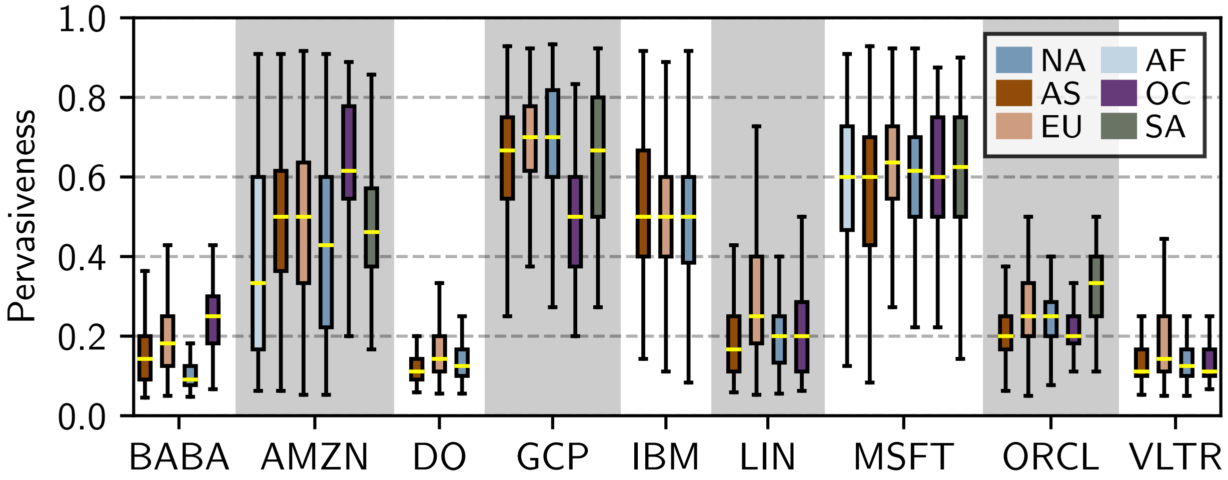

- Cloud providers are actively becoming pervasive – shortening paths between end-users and their private WANs
- Hypergiants (Amazon, Google, Microsoft) have large private WAN that direct peers with most end-user ISPs
- Medium-sized (IBM, DigitalOcean) have limited WAN size and prefer private interconnections at Tier-1
- Developing CPs (Alibaba, Oracle, Vultr) rely on public Internet
  - Alibaba directly peers within China

Routing and Peering with Cloud WANs

Does direct cloud WAN peering result in latency benefits to users?

In Europe

In Asia

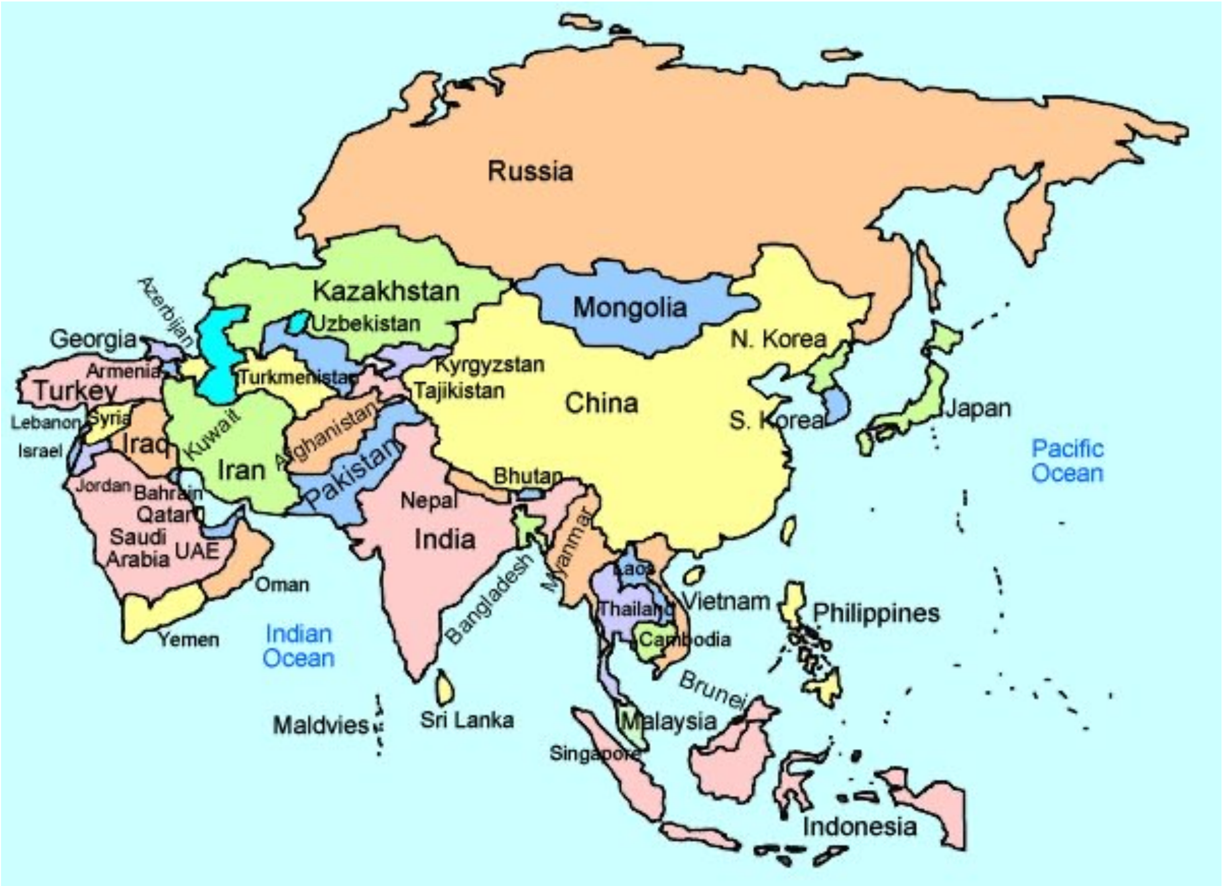

Routing and Peering with Cloud WANs

Does direct cloud WAN peering result in latency benefits to users?

VPs in Germany → DCs in UK

VPs in Japan → DCs in India

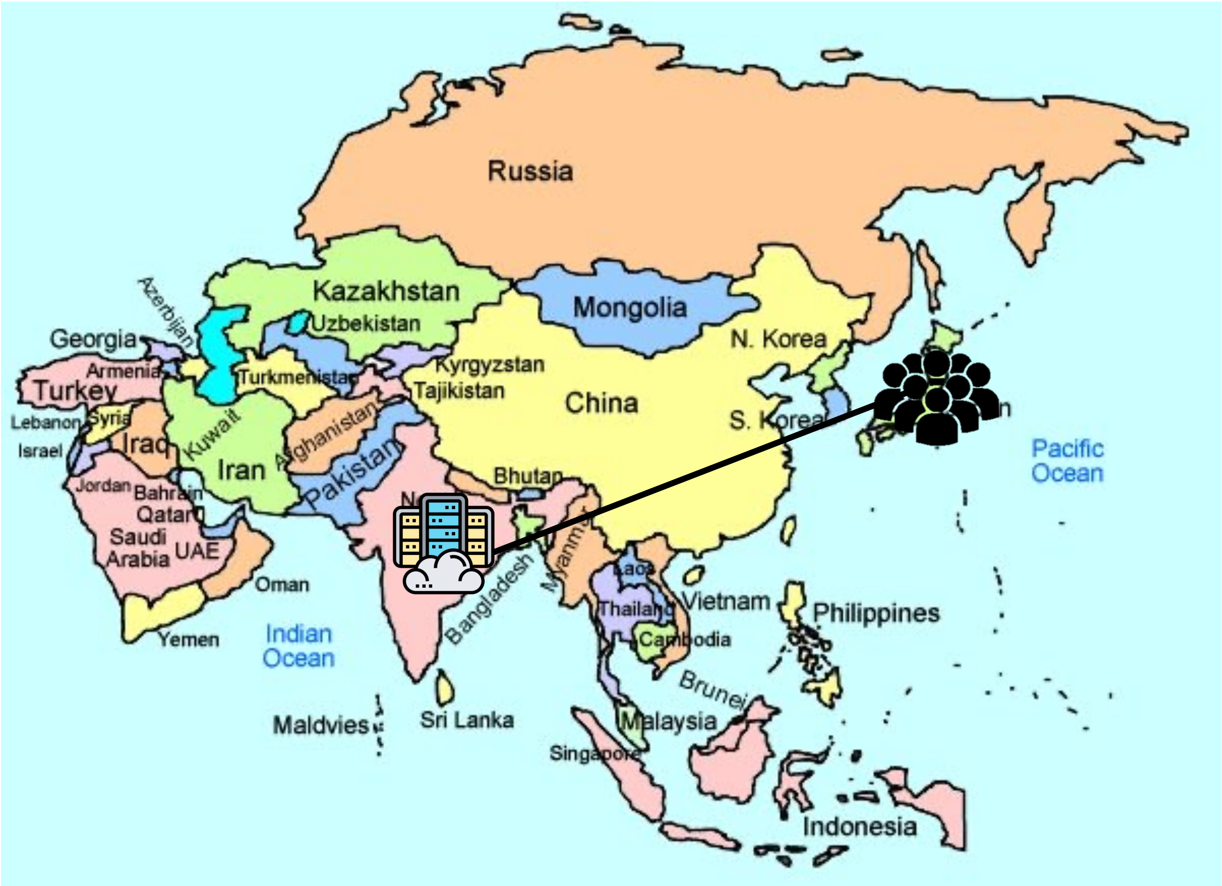

Routing and Peering with Cloud WANs

Does direct cloud WAN peering result in latency benefits to users?

VPs in Germany → DCs in UK

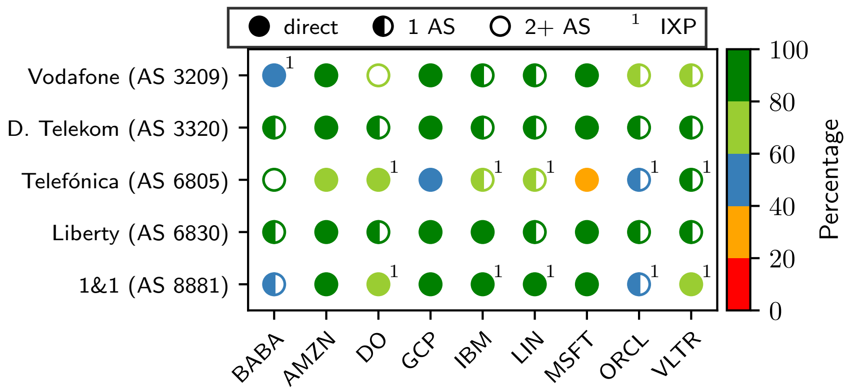

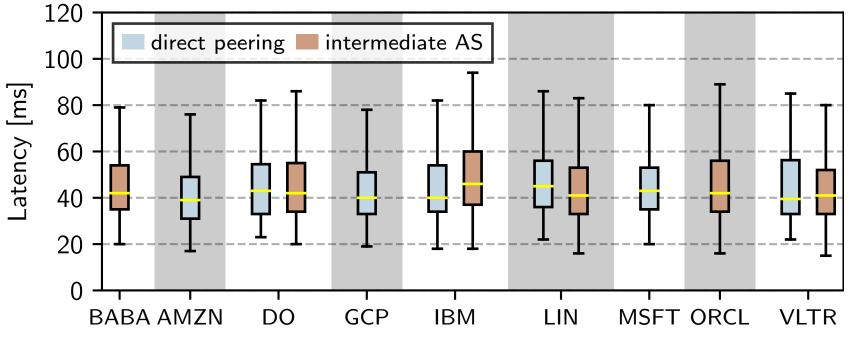

VPs in Japan → DCs in India

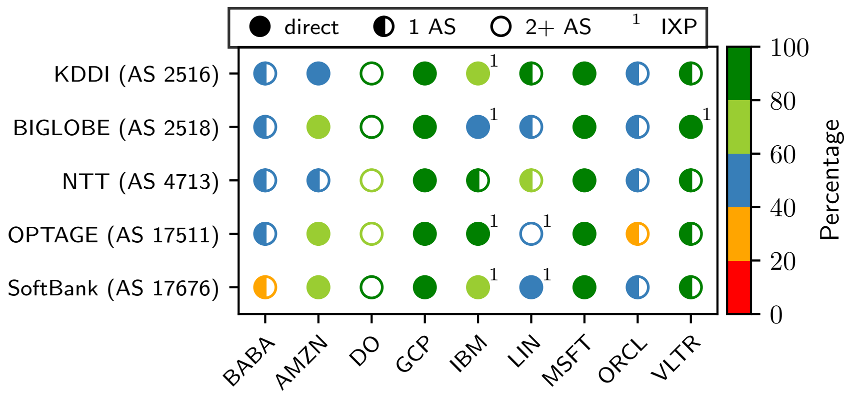

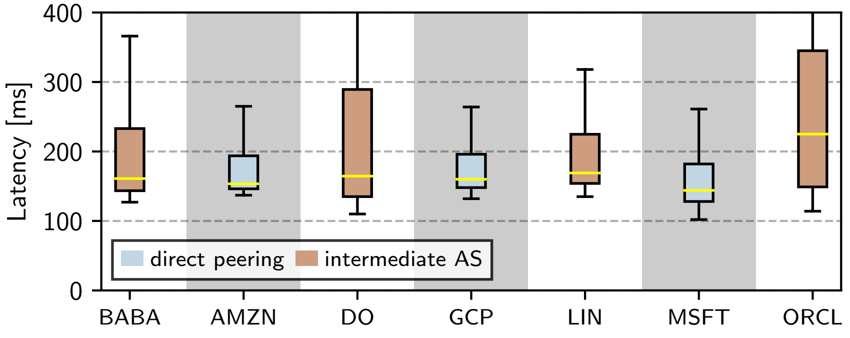

Routing and Peering with Cloud WANs

Does direct cloud WAN peering result in latency benefits to users?

- Direct peering does not bring down cloud access latencies
- In Europe, public Internet is well-provisioned and direct peering has zero effect on latencies
- In Asia, direct peering helps in reducing the long latency tails

We do not capture the impact of peering on network bandwidth and traffic isolation which can be the driving factor for flattening

Influence of the Wireless Last-Mile

- While cloud providers continue to expand their WAN deployment and reduce user path lengths to their network through peering, the last-mile remains out of reach
- Previous studies have shown LTE and WiFi to be the major latency bottleneck
    - Inconsistent latency, packet drops, bufferbloating

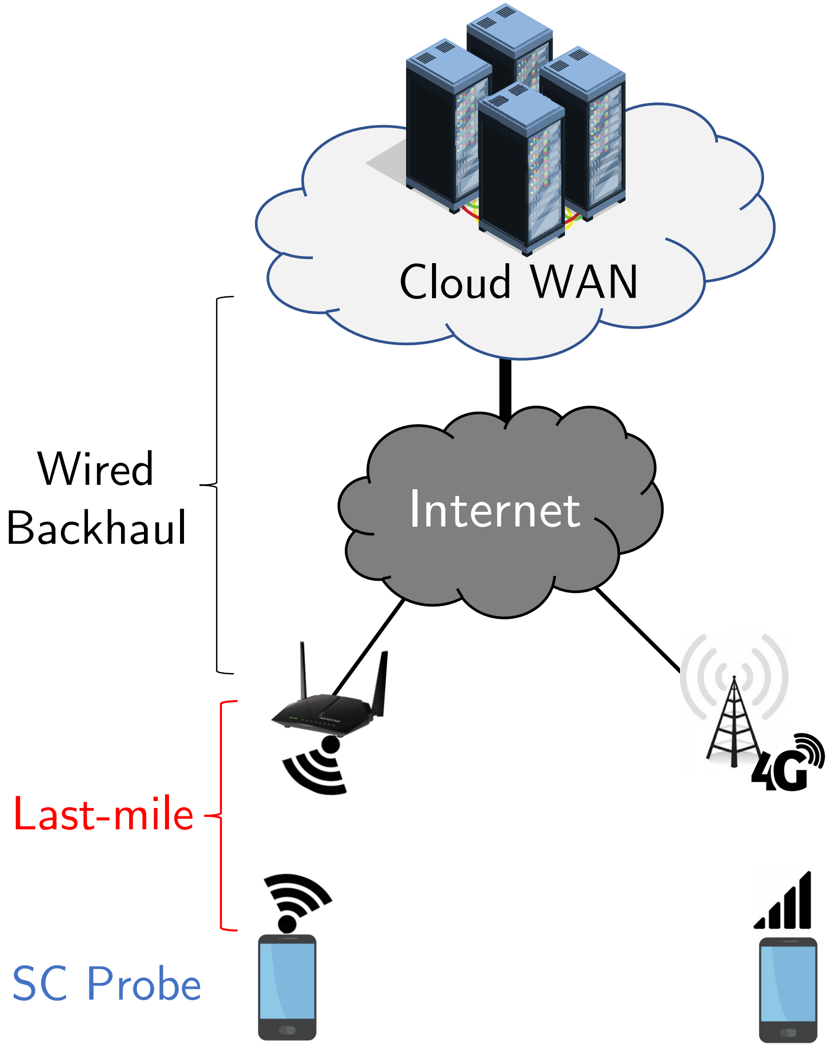

Influence of the Wireless

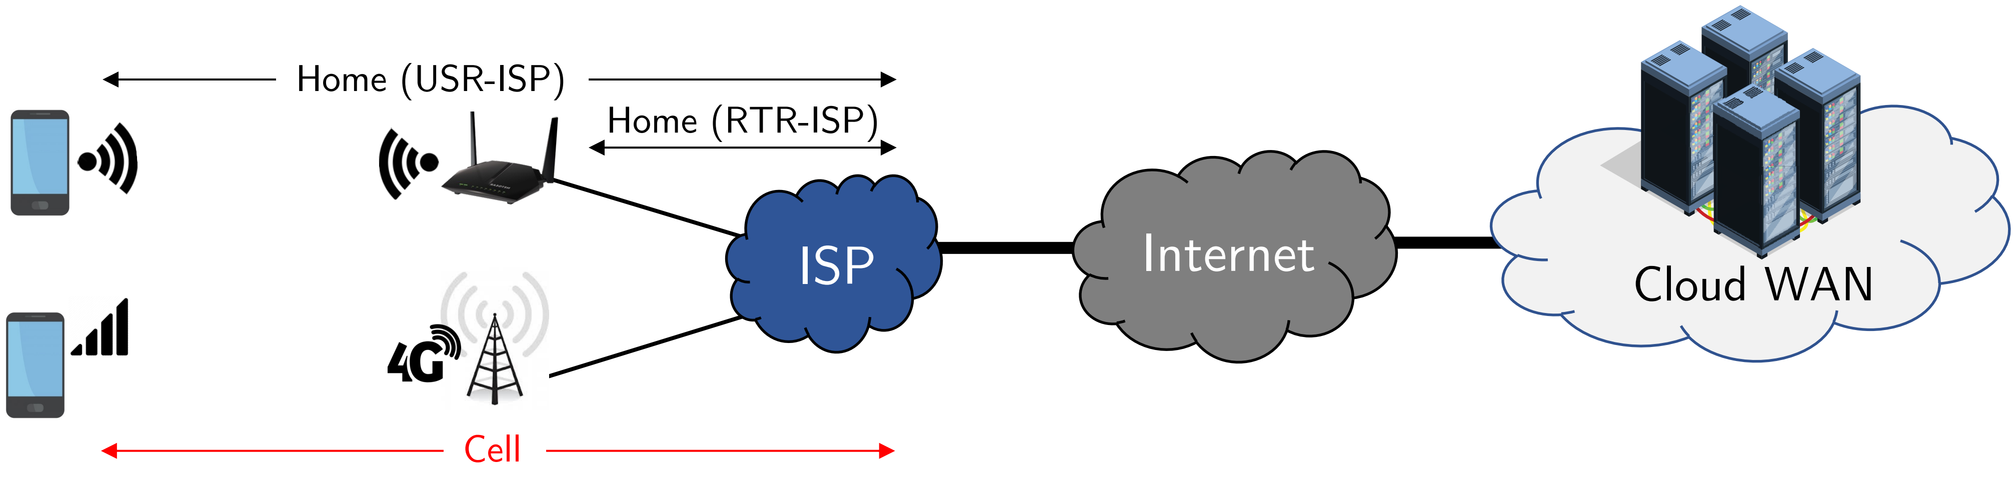

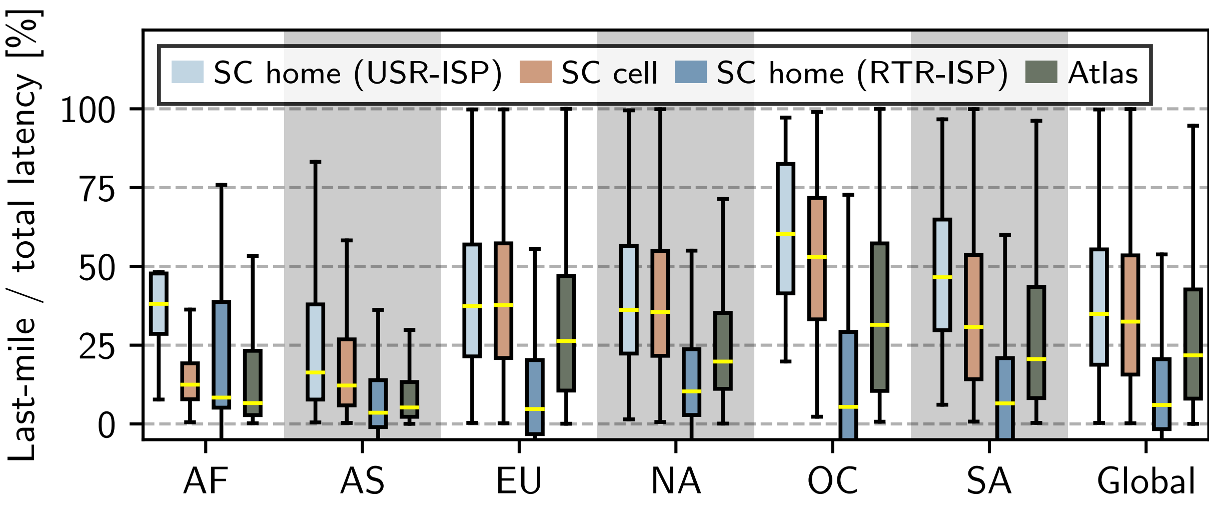

Influence of the Wireless

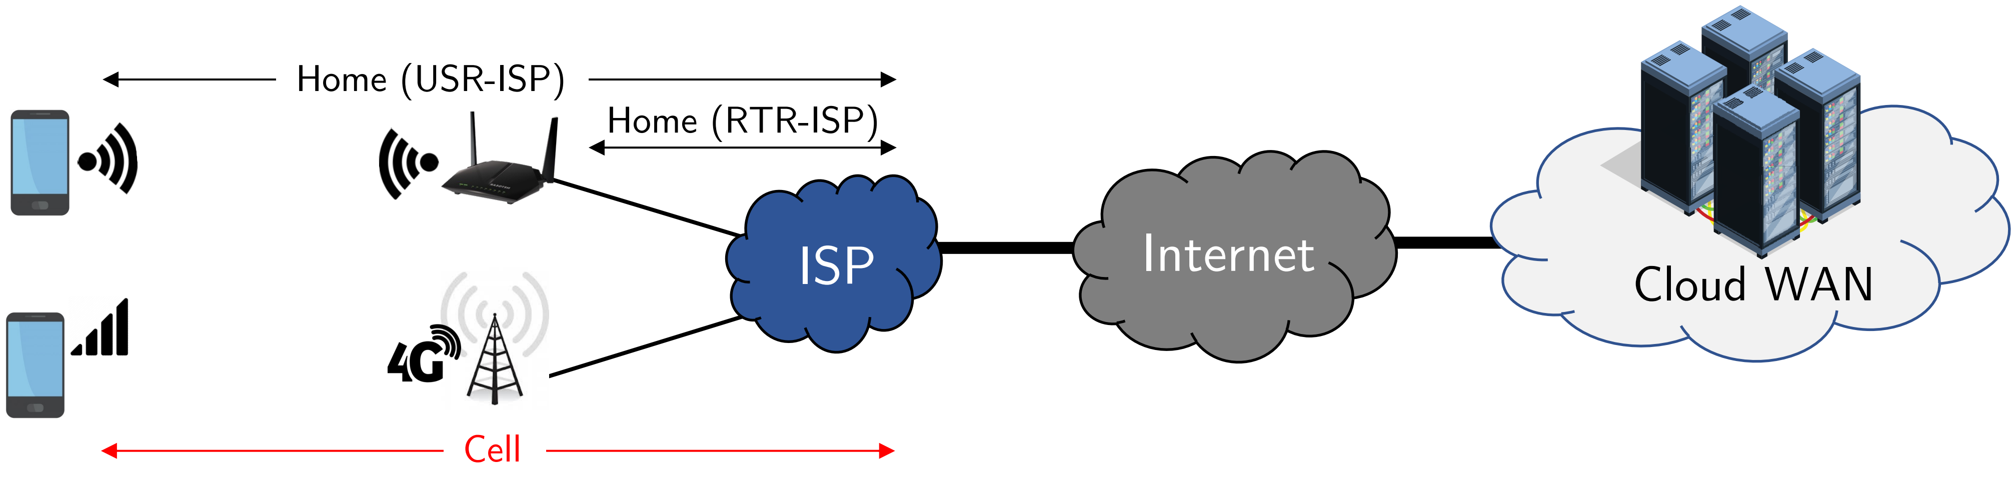

- User-to-ISP latency share is similar for both WiFi and cellular globally

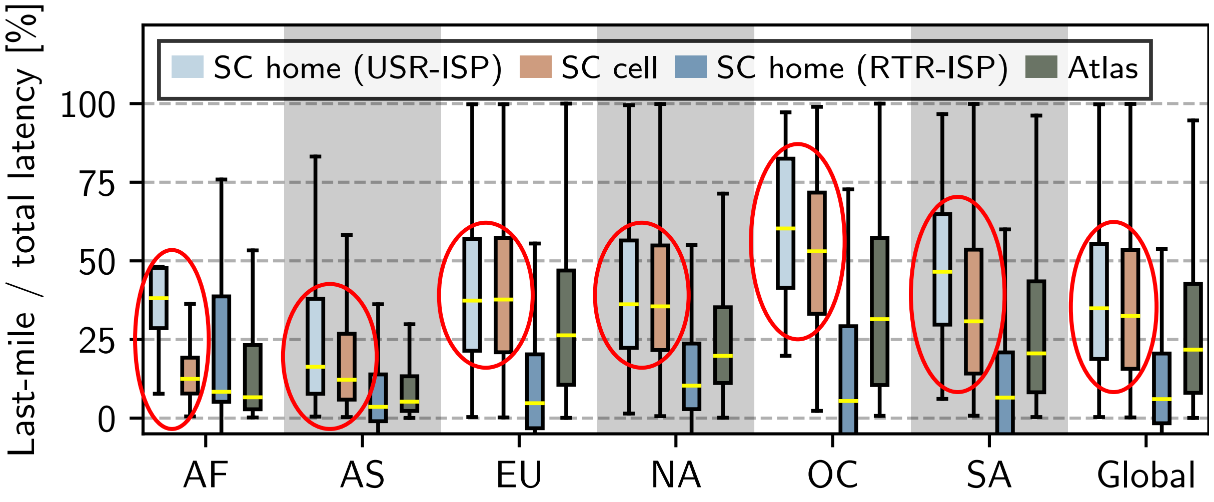

Influence of the Wireless

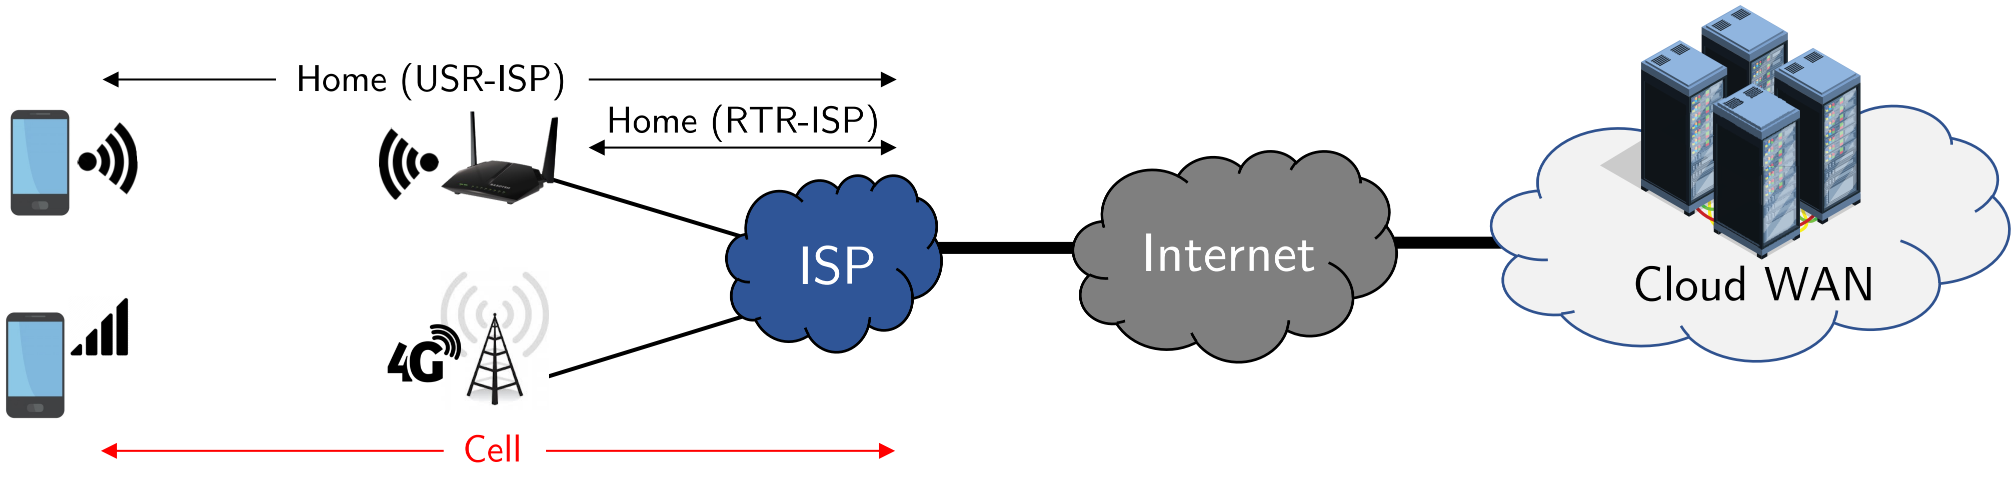

- User-to-ISP latency share is similar for both WiFi and cellular globally
- Latency share of wireless is higher in regions with dense DC deployment (30-50%) since the backhaul is much shorter

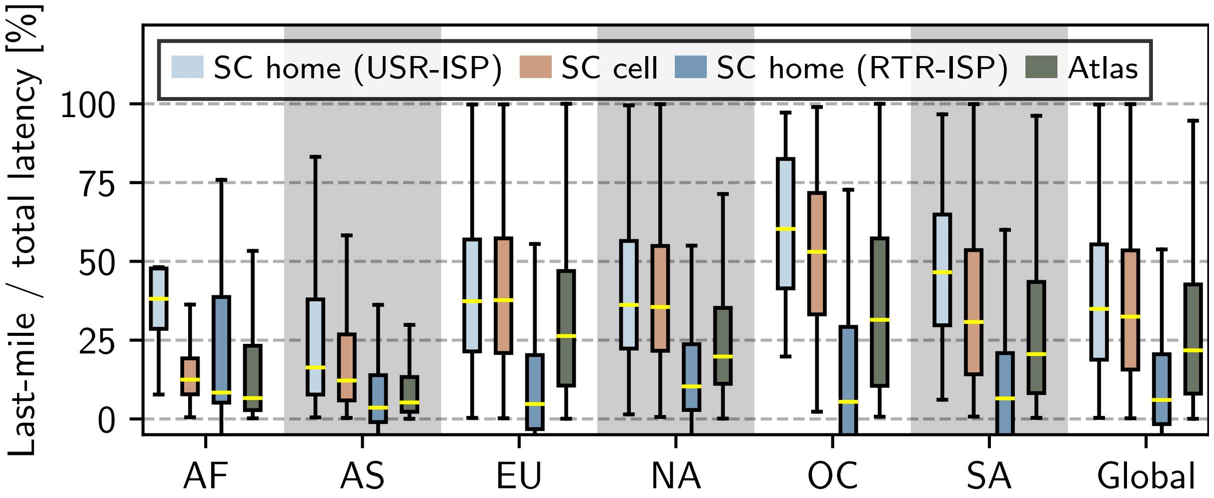

Influence of the Wireless

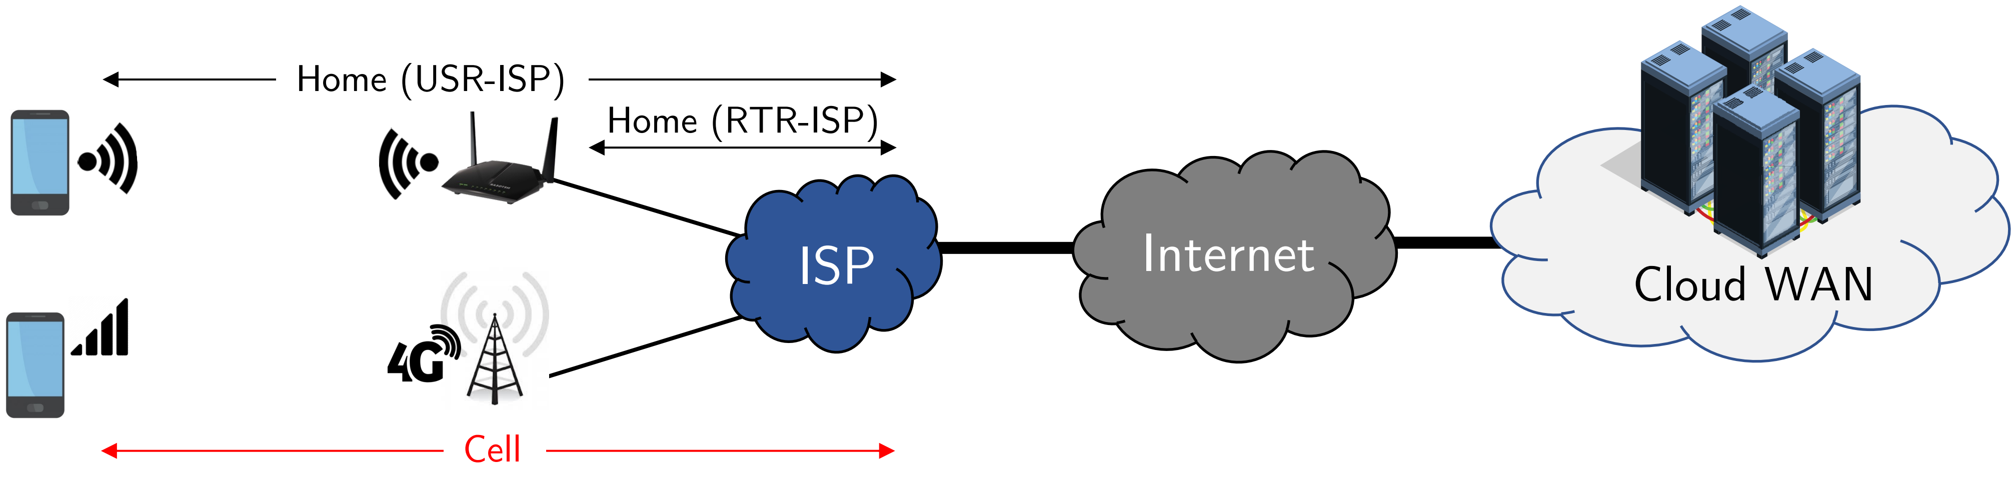

- User-to-ISP latency share is similar for both WiFi and cellular globally
- Latency share of wireless is higher in regions with dense DC deployment (30-50%) since the backhaul is much shorter

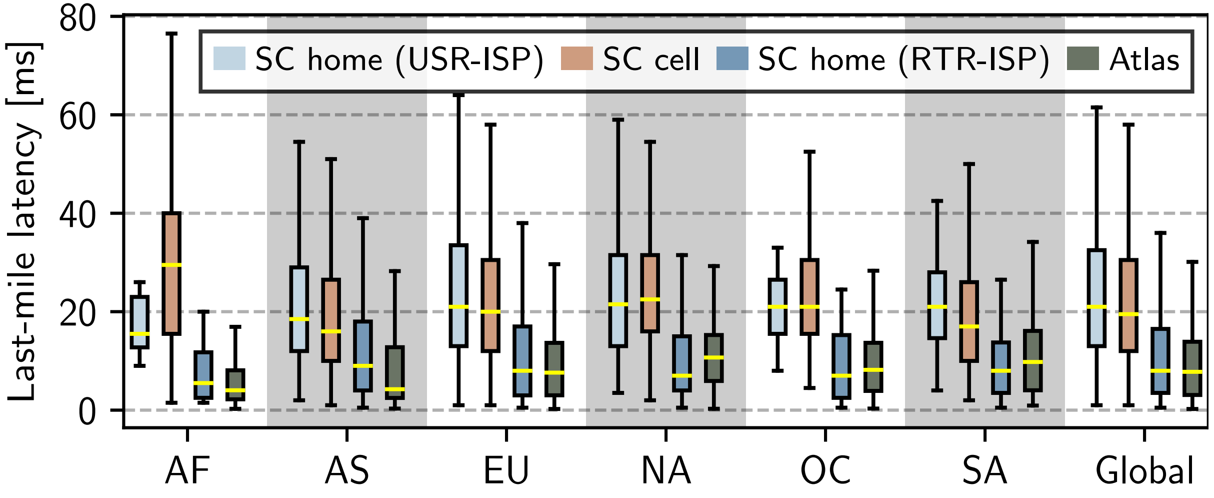

Influence of the Wireless

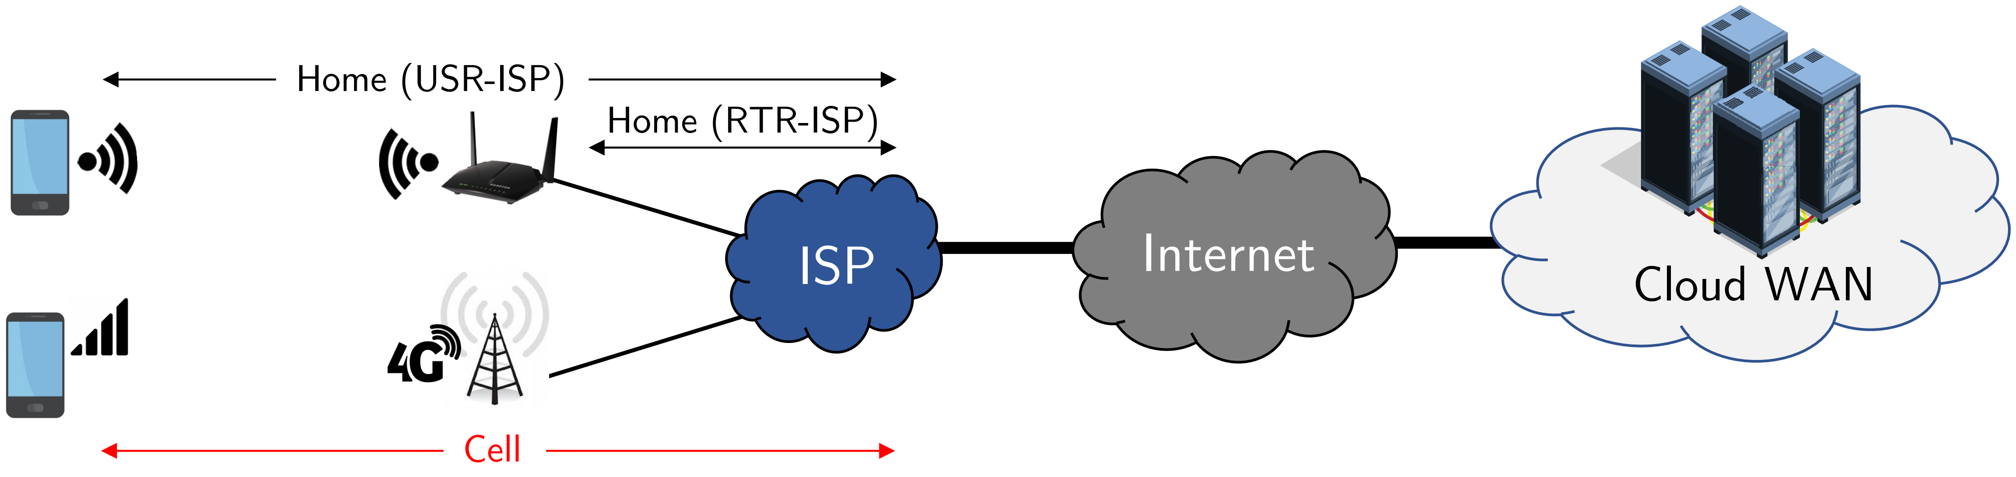

- User-to-ISP latency share is similar for both WiFi and cellular globally
- Latency share of wireless is higher in regions with dense DC deployment (30-50%) since the backhaul is much shorter
- Absolute latency due to last-mile is similar for both *home* (WiFi+Router) and cellular (only wireless), i.e. $\gtrsim$ 20ms

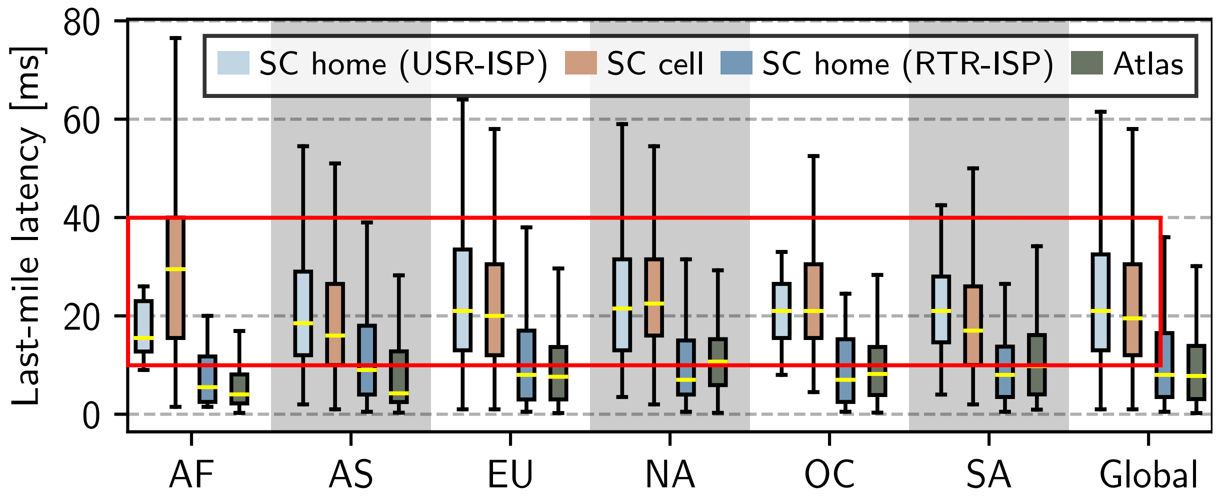

Influence of the Wireless

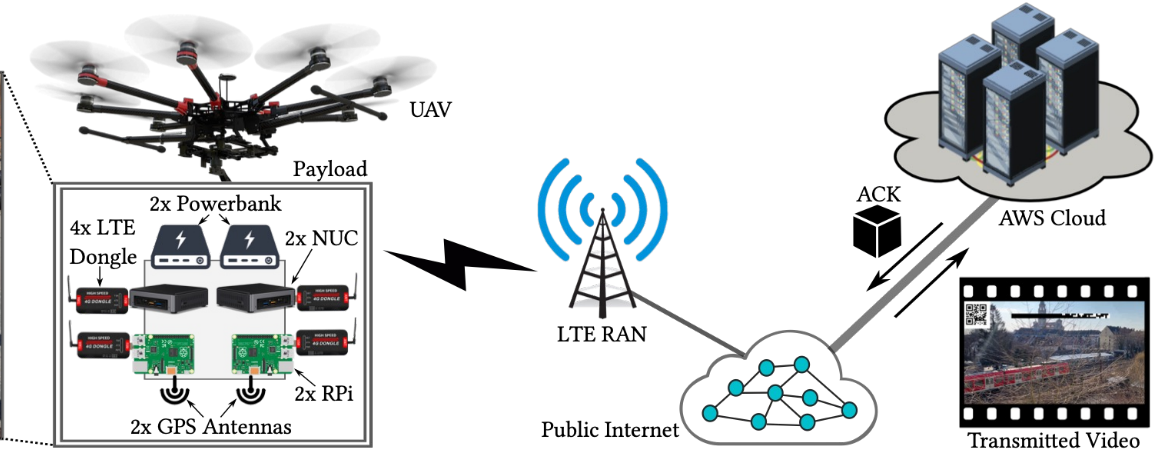

Last-mile latencies can be much worse for next-gen apps

- In the air
- Under load

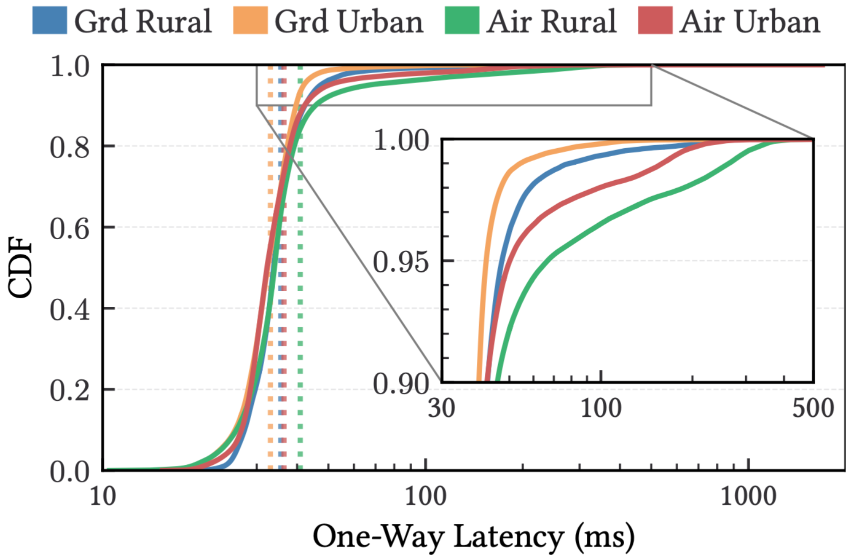

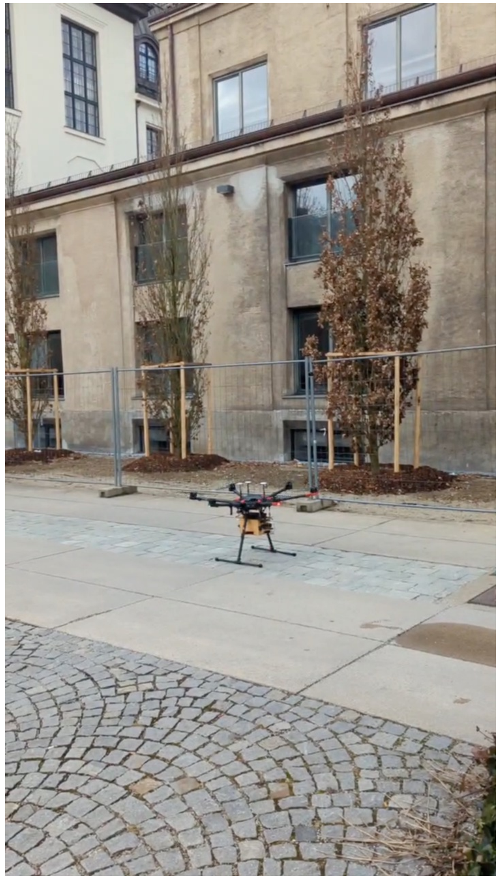

Aygün Baltaci, Hendrik Cech, Nitinder Mohan, Fabien Geyer, Vaibhav Bajpai, Jörg Ott, Dominik Schupke: Analyzing Real-time Video Delivery over Cellular Networks for Remote Piloting Aerial Vehicles. Accepted for publication at the ACM Internet Measurement Conference 2022 (IMC 2022)

Centralization vs Decentralization – Points to Ponder

TECH

Google wins cloud deal from Elon Musk's SpaceX for Starlink internet connectivity

PUBLISHED THU, MAY 13 2021·8:30 AM EDT | UPDATED THU, MAY 13 2021·3:55 PM EDT

AT&T to run its mobility network on Microsoft's Azure for Operators cloud, delivering cost-efficient 5G services at scale

June 30, 2021 | Microsoft News Center

- Cloud provider's reach will continue to increase
- Barrier-to-entry (in terms of investment) for new competition will keep becoming higher

Centralization vs Decentralization – Points to Ponder

TECH

Google wins cloud deal from Elon Musk's SpaceX for Starlink internet connectivity

PUBLISHED THU, MAY 13 2021·8:30 AM EDT | UPDATED THU, MAY 13 2021·3:55 PM EDT

- Cloud provider's reach will continue to increase
- Barrier-to-entry (in terms of investment) for new competition will keep becoming higher

AT&T to run its mobility network on Microsoft's Azure for Operators cloud, delivering cost-efficient 5G services at scale

June 30, 2021 | Microsoft News Center

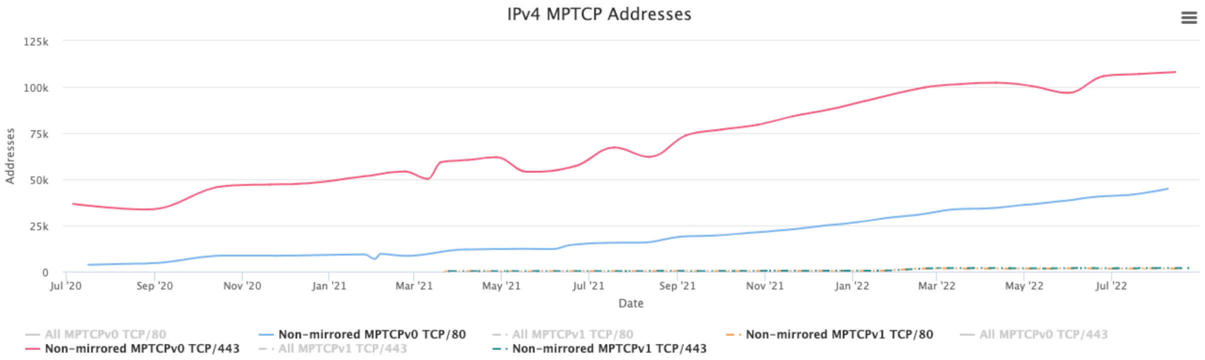

Large-scale protocol adoption

Stress control

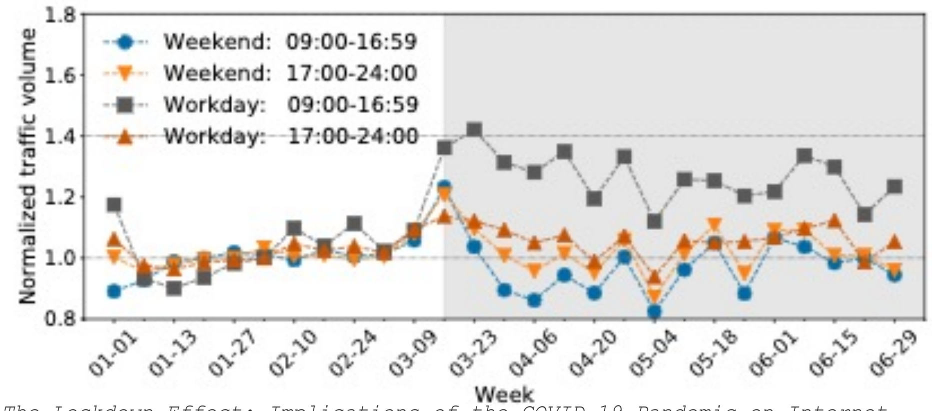

F. Aschenbrenner, T. Shreedhar, O. Gasser, N. Mohan and J. Ott, "From Single Lane to Highways: Analyzing the Adoption of Multipath TCP in the Internet," *2021 IFIP Networking Conference*

Feldmann et al. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. ACM Internet Measurement Conference (IMC '20).