Studying the Evolution of Content Providers in the Internet Core

Esteban Carisimo, Carlos Selmo,
J. Ignacio Alvarez-Hamelin, Amogh Dhamdhere

AS ecosystem mutation

NSFNET era (up to 1995)

- Monolithic backbone

Transit era (late 1990s-2000s)

- Densely connected Transit Network
- Eyeballs and CPs on the edge
- Tiered model

Content era (2010s)

- Dominance of multimedia content
- CDNs
- Flat network

CONICET

Irruption of IXPs and CDNs

- **Deployment of CDNs**
  - General-purpose CDNs
  - Ascension of private CDNs
  - Result: CPs densely connected (core)
- **IXPs**
  - Flattening the Internet
  - Let small ASes peer directly with large ASes
  - Co-location with CDNs
- **IXPs+CDNs: Traffic impact**
  - Reduction on Transit traffic
  - Peer-to-peer traffic

Motivations

Google and Netflix Make Land Grab On Edge Of Internet

Many of these deals are secret, but Deepfield Networks knows of about 40 companies that are setting up their own content delivery networks with service providers, according to Craig Labovitz. But he's bound by non-disclosure agreements, and can't name names.

Wired.June 2012

https://www.wired.com/2012/06/cdn/

CONICET

Motivations

Seeking Alpha Portfolio People News Analysis Search by symbol, a

Apple, Microsoft And Facebook Bring More Traffic To In-House CDNs, Impacting Akamai's Media Business

Oct.28.15 | About: Akamai Technologies, (AKAM)

Dan Rayburn
Research analyst, streaming and online video
StreamingMedia

Follow
(819 followers)

Summary

- Akamai reported Q3 earnings and announced that revenue from their media delivery business would be flat or down for Q4, year-over-year.
- Akamai suggested the reason they expect media growth rates to continue to moderate in the "near term" was due to customers having "less traffic growth overall."
- The cause of what Akamai is seeing is a result of Apple, Microsoft and Facebook moving a larger percentage of their traffic to their in-house delivery networks.

October 2015

https://seekingalpha.com/article/3613736-apple-microsoft-facebook-bring-traffic-house-cdns-impacting-akamais-media-business

Motivations

BLOG SEARCH ADVANCED

APNIC

Get IP Manage IP Training Events Research Community

The death of transit?

By Geoff Huston on 28 Oct 2016

Category: Tech matters

Tags: ISPs, data centres, peering

12 Comments

Geoff Huston, October 2016

Is there light at the end of the tunnel for transit providers?

I was struck at a recent NANOG meeting just how few presentations looked at the ISP space and the issues relating to ISP operations, and how many were looking at the data centre environment.

https://blog.apnic.net/2016/10/28/the-death-of-transit/
https://labs.apnic.net/presentations/store/2017-05-25-death-of-transit.pdf

Do CPs belong to the core of the network?

- How to identify if they do? Who are they? Since when?
- Differences on cores by Region?
- How to detect *“up and coming”* CDNs?
- Correlate evolution of connectivity with well-documented business practices

Methodology: k-core decomposition

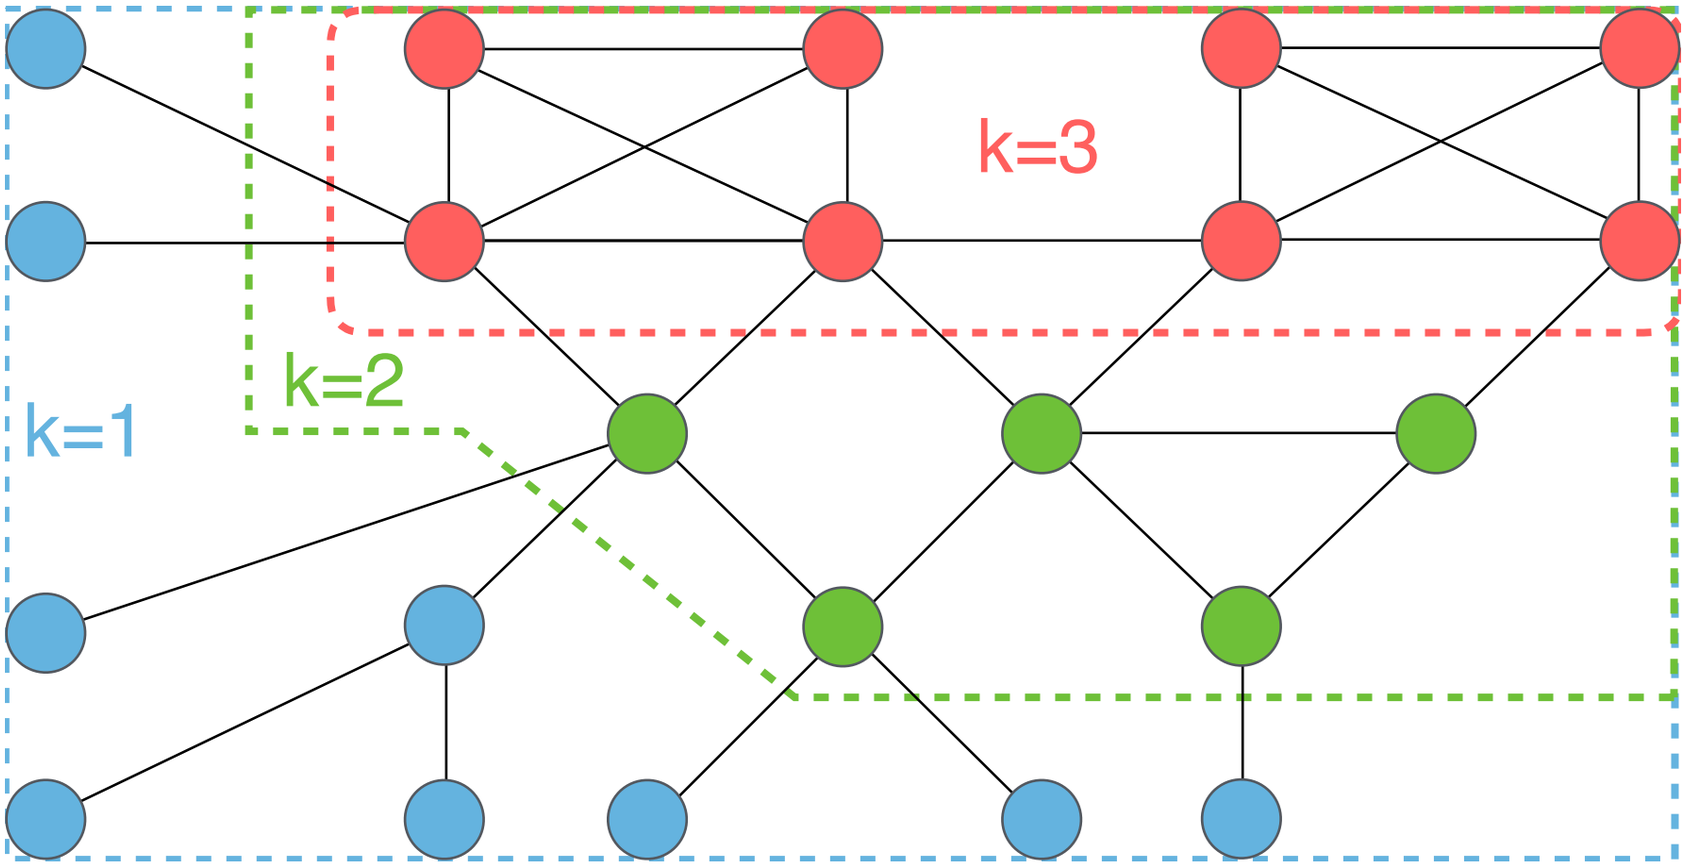

- Core of the network: Densely connected ASes
- How to know who they are: k-cores
- k-cores: *shell-index* given by
  - number **AND** degree of neighbors

Methodology: k-core decomposition

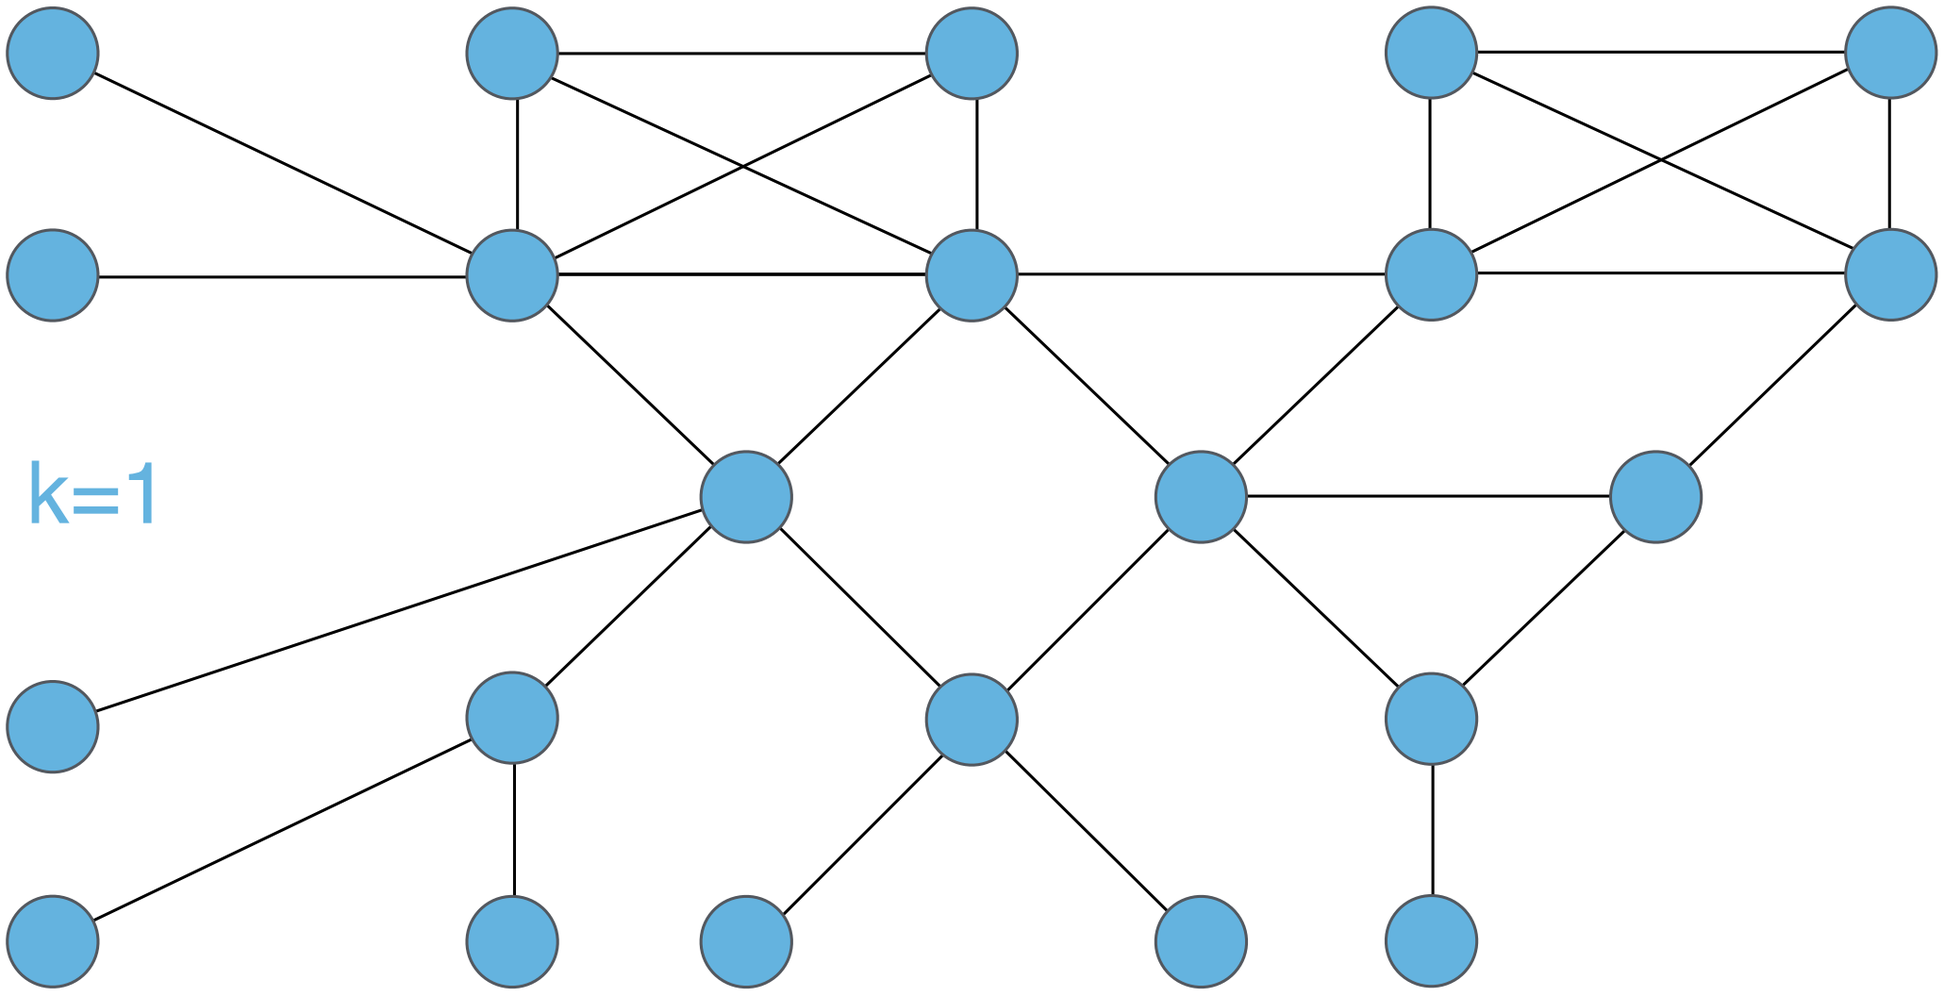

Methodology: k-core decomposition

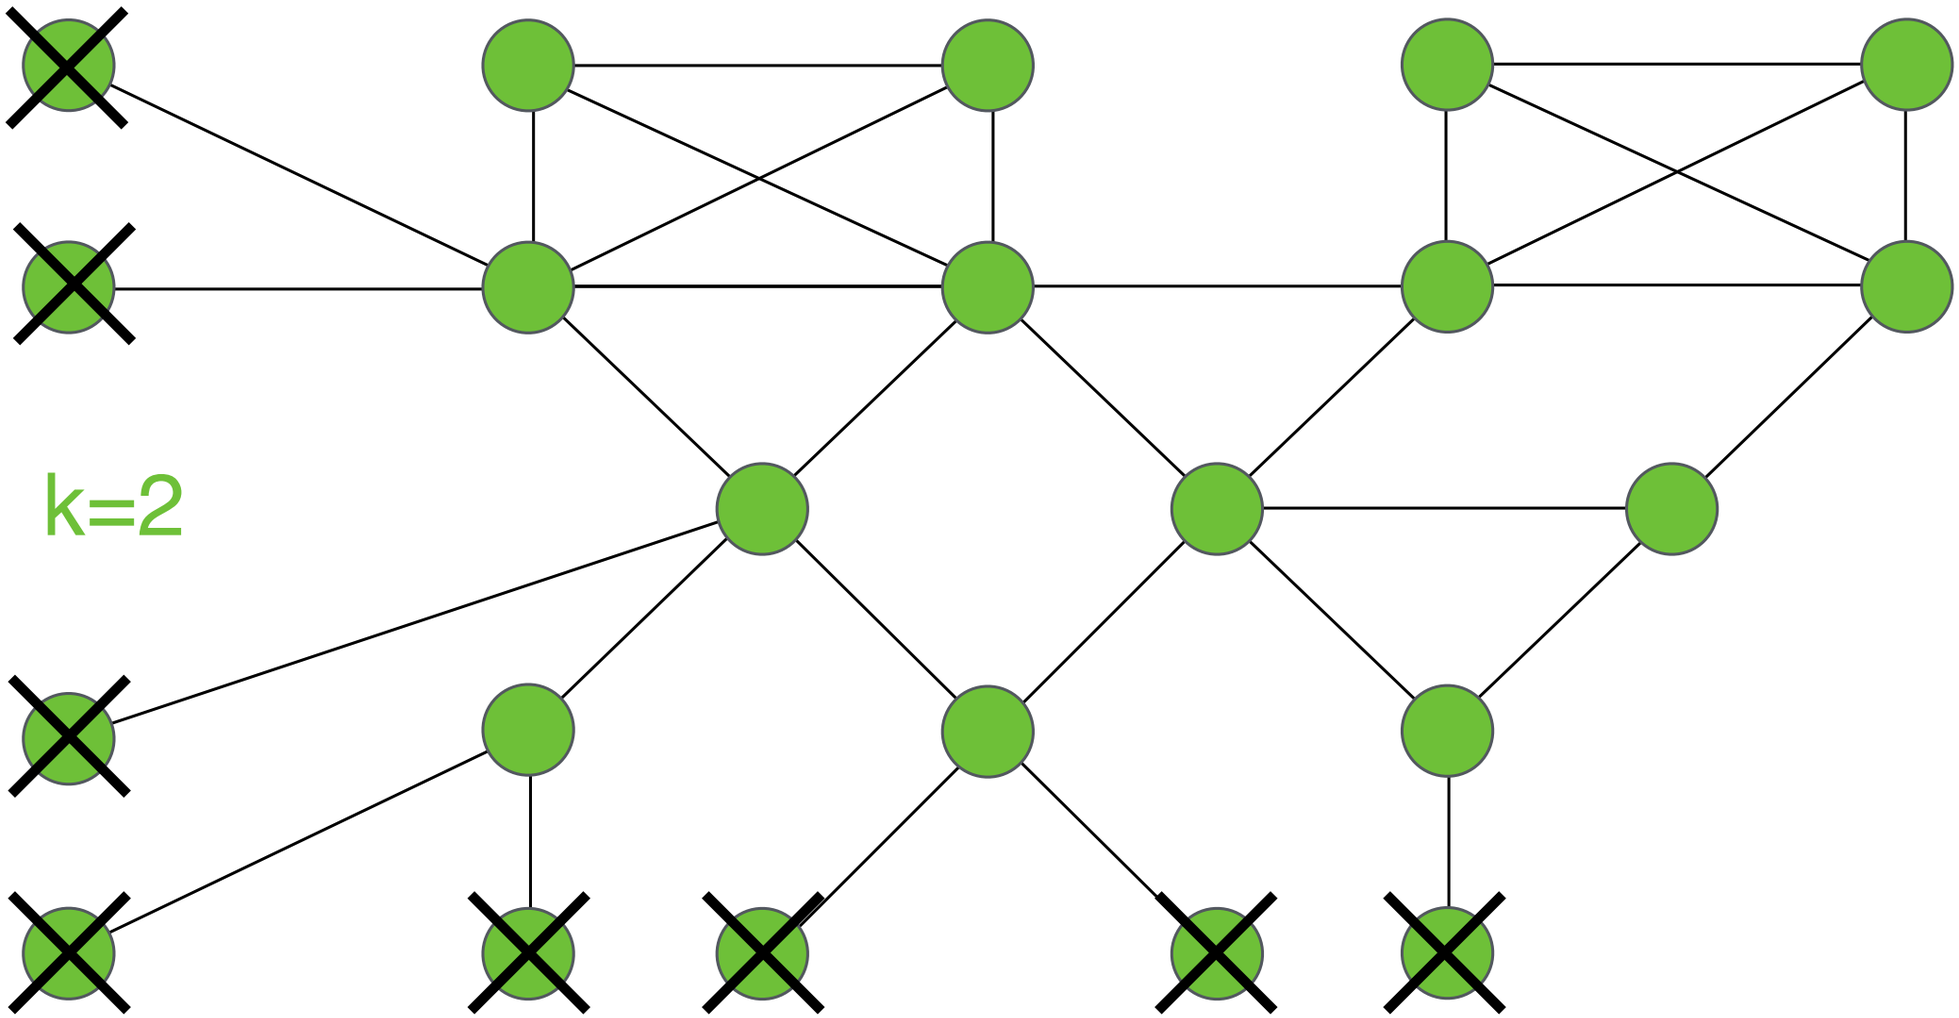

Methodology: k-core decomposition

k=2

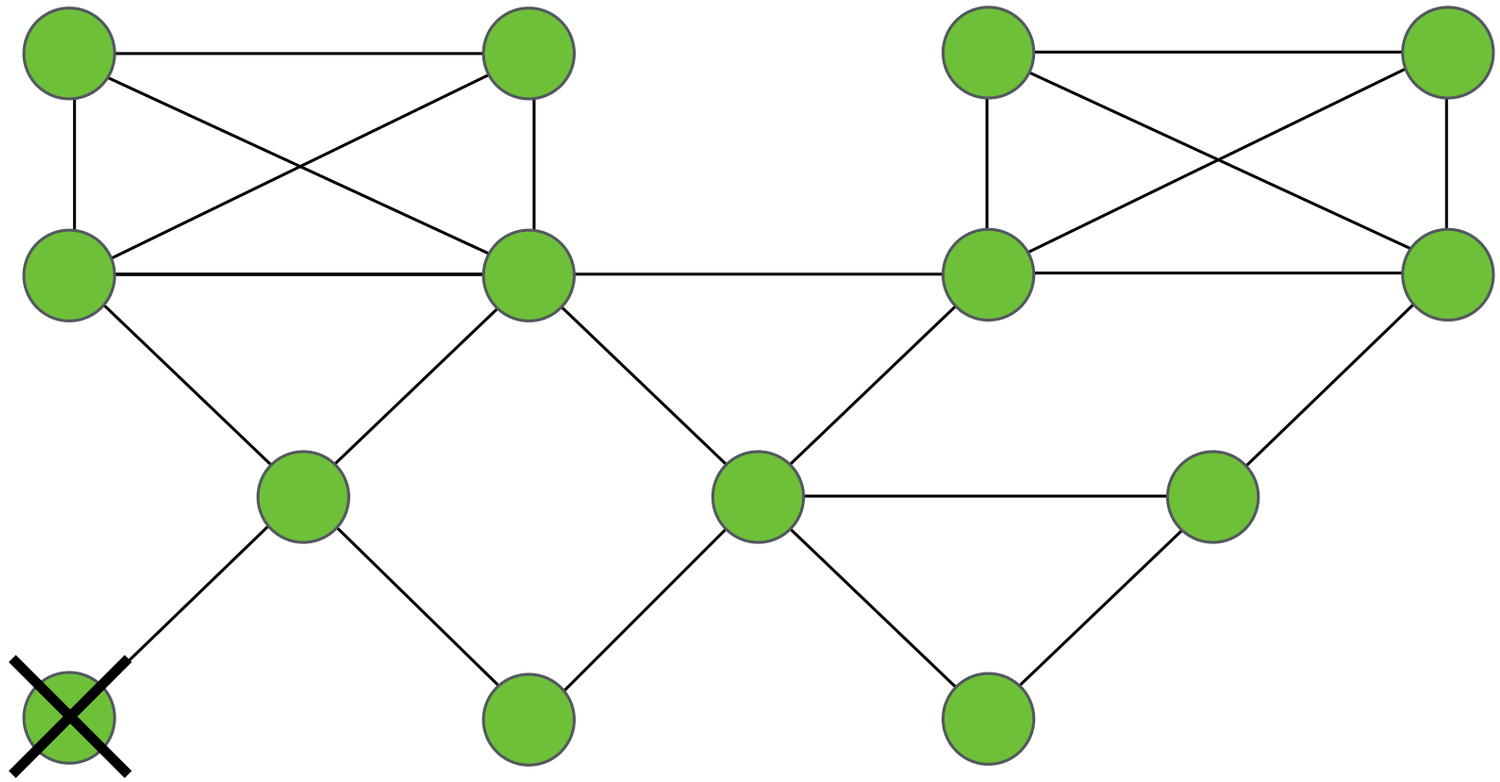

Methodology: k-core decomposition

k=2

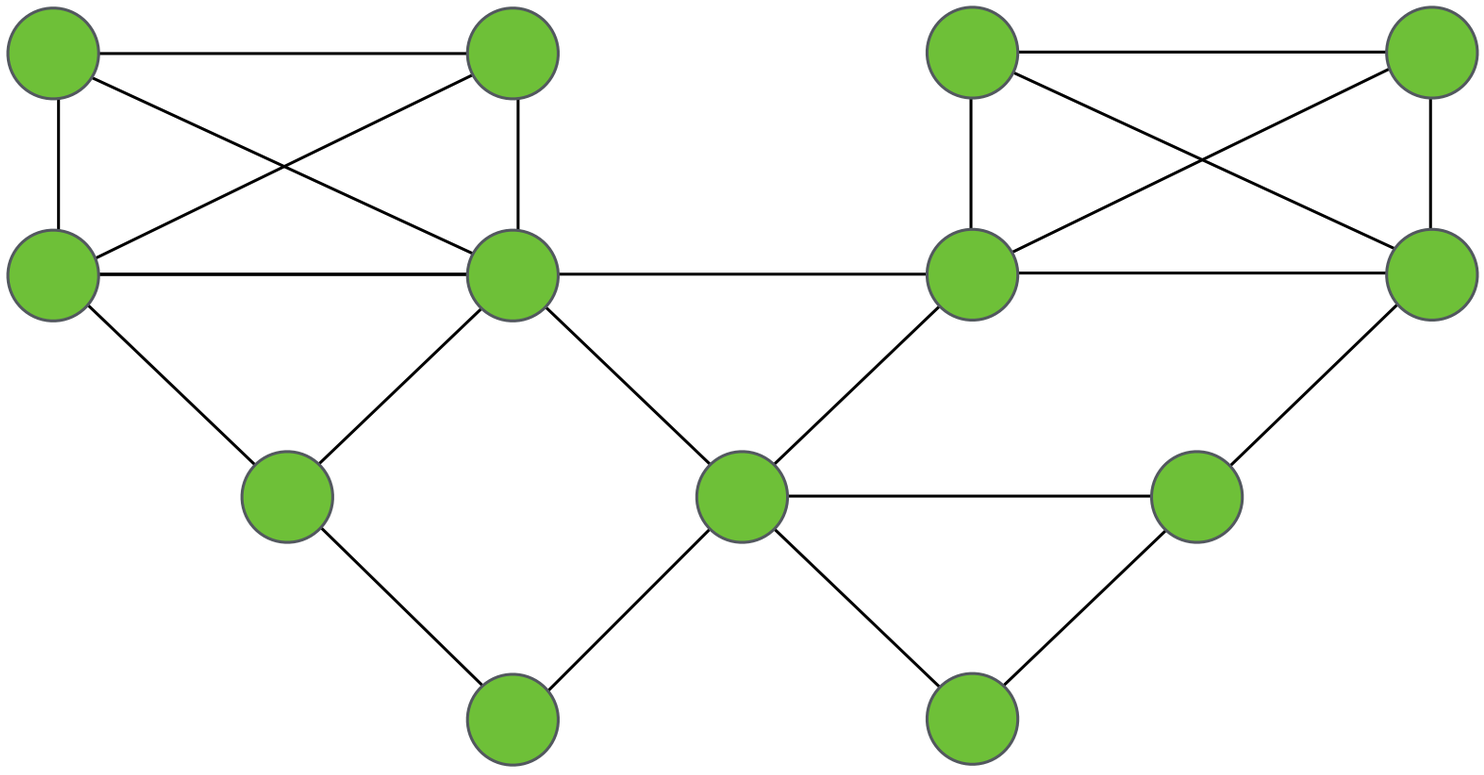

Methodology: k-core decomposition

k=3

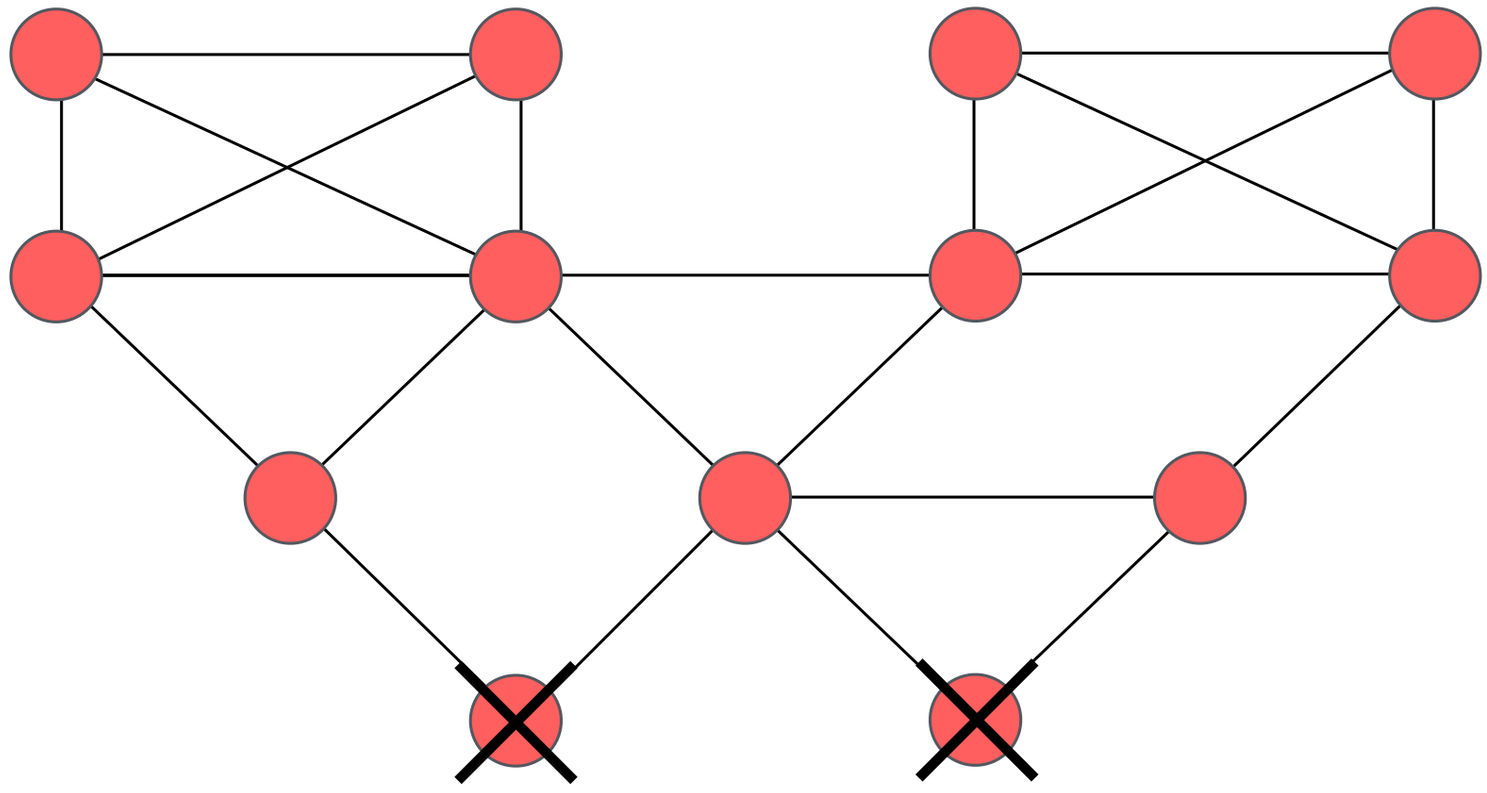

Methodology: k-core decomposition

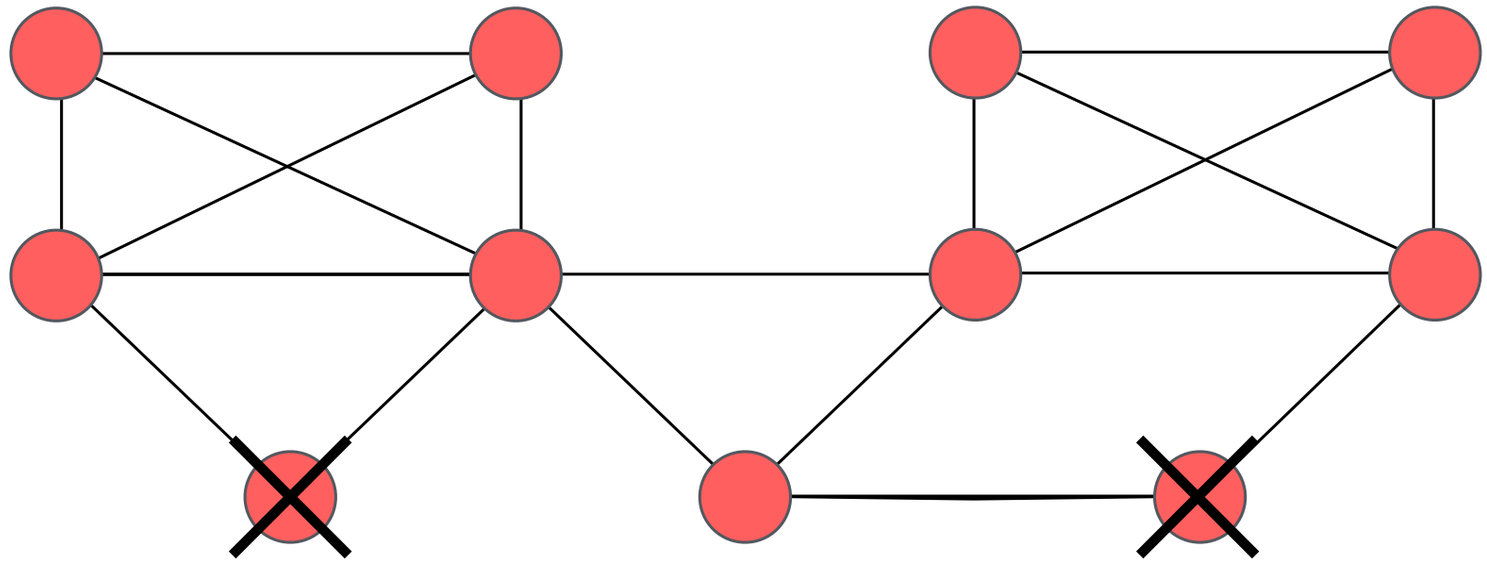

k=3

Methodology: k-core decomposition

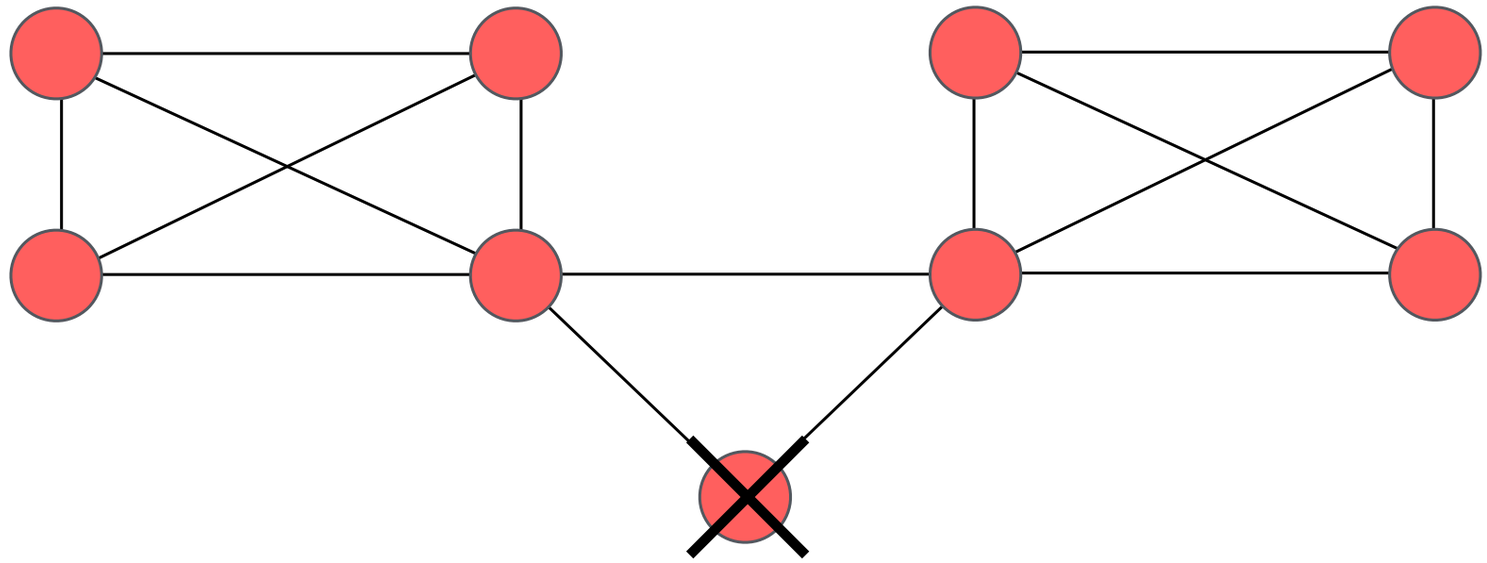

k=3

Methodology: k-core decomposition

k=3

Methodology: k-core decomposition

k=3

Conclusion

If a node belongs to core *k*, it has *k* neighbors node who have at least degree *k*

Methodology: k-core decomposition

k=3

TOPcore: The maximum core (k) in the graph under analysis

Dataset

- **Data sources**
  - CAIDA's publicly available data
  - BGP dumps
    - Routeviews
    - RIPE RIS
  - Ark Traceroute campaigns

- **Details about dataset**
  - Monthly snapshots of AS graph
  - Spans from 1999 to late 2017


Analyses

1. Core evolution of large CPs *(Big Seven)* throughout the years

2. Geographical differences on the core evolution

3. Other members of the core

Analyses

1. **Core evolution of large CPs *(Big Seven)* throughout the years**
2. Geographical differences on the core evolution
3. Other members of the core

Tracking the *Big Seven*

Based on traffic statistics

- Sandvine report
- PeeringDB records

Some caveats

CPs may have several ASNs

- Focus on primaries

In-network caches

- Content is often served from caches on ISPs' address space
- However, peering links are necessary to
  - fill caches, serve dynamic content and serve ISPs that are unwilling to host a cache

CONICET

Tracking the *Big Seven*

oct-2006

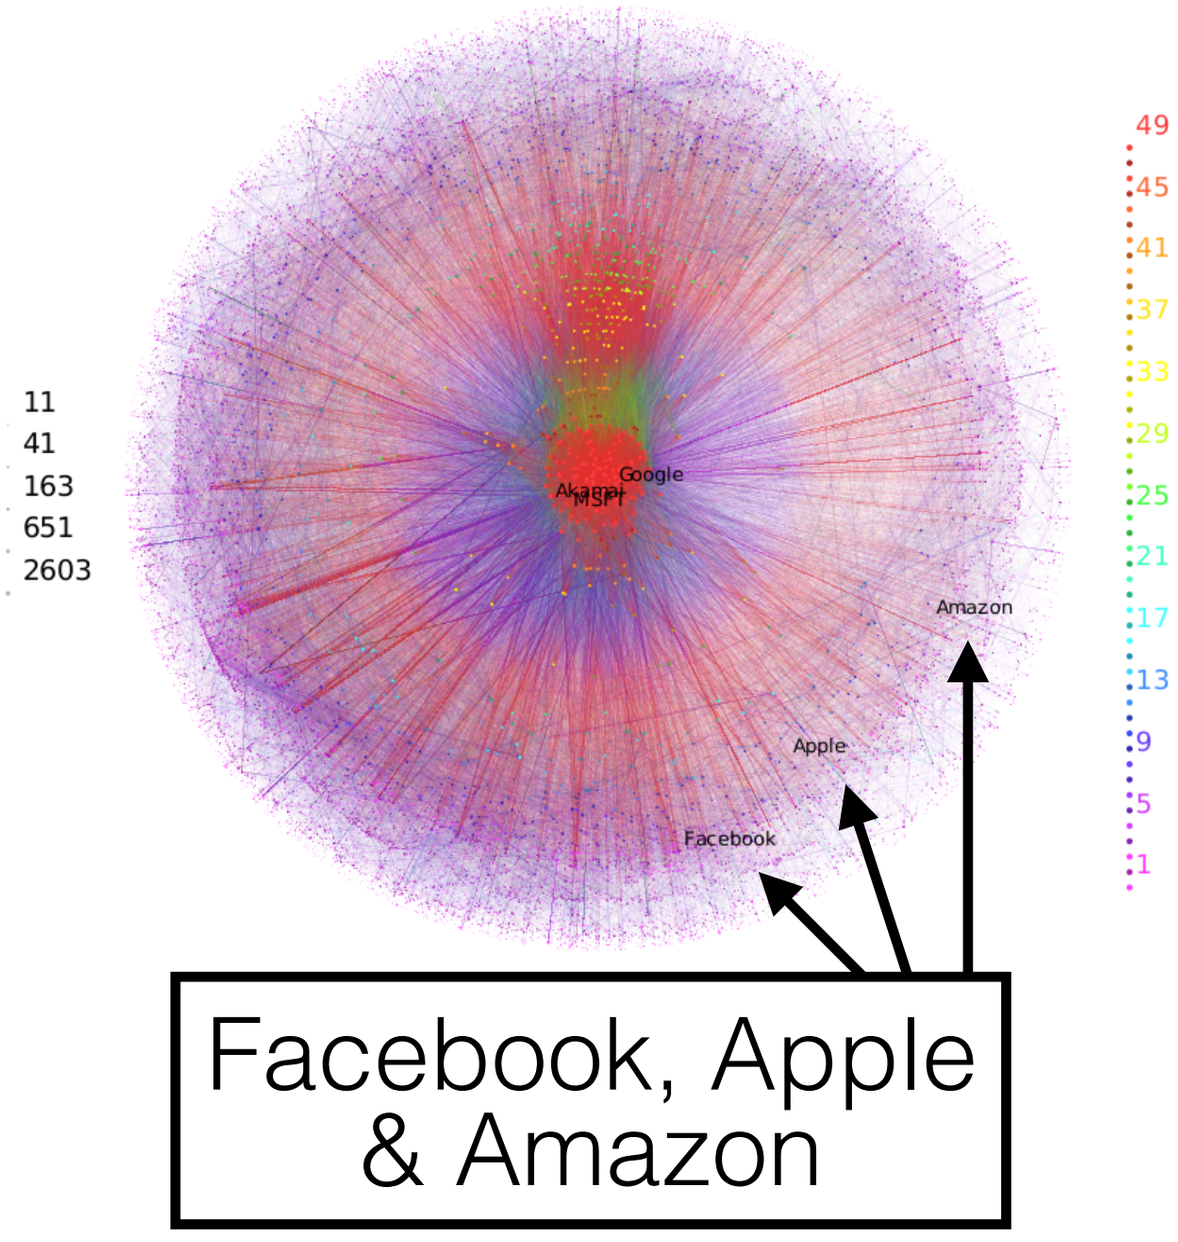

oct-2016

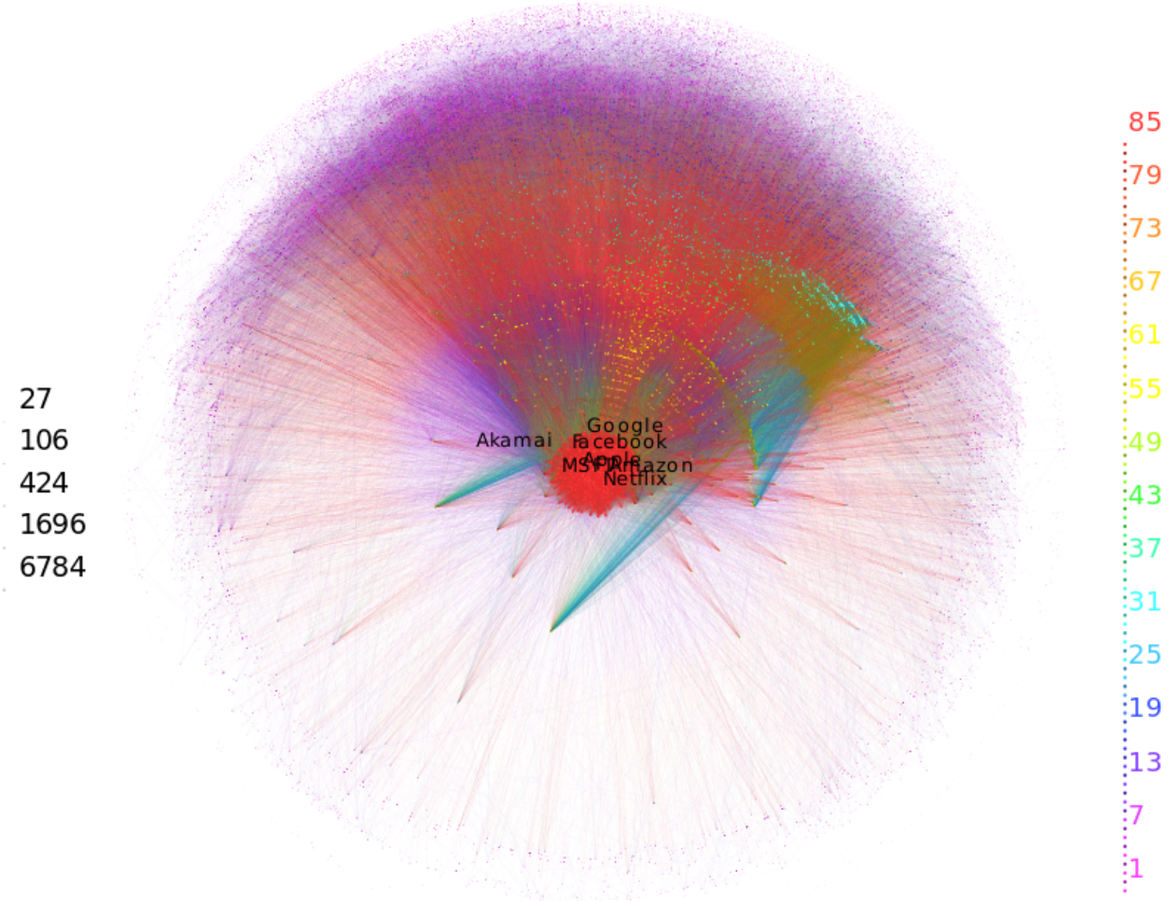

CONICET

Evolution of big seven

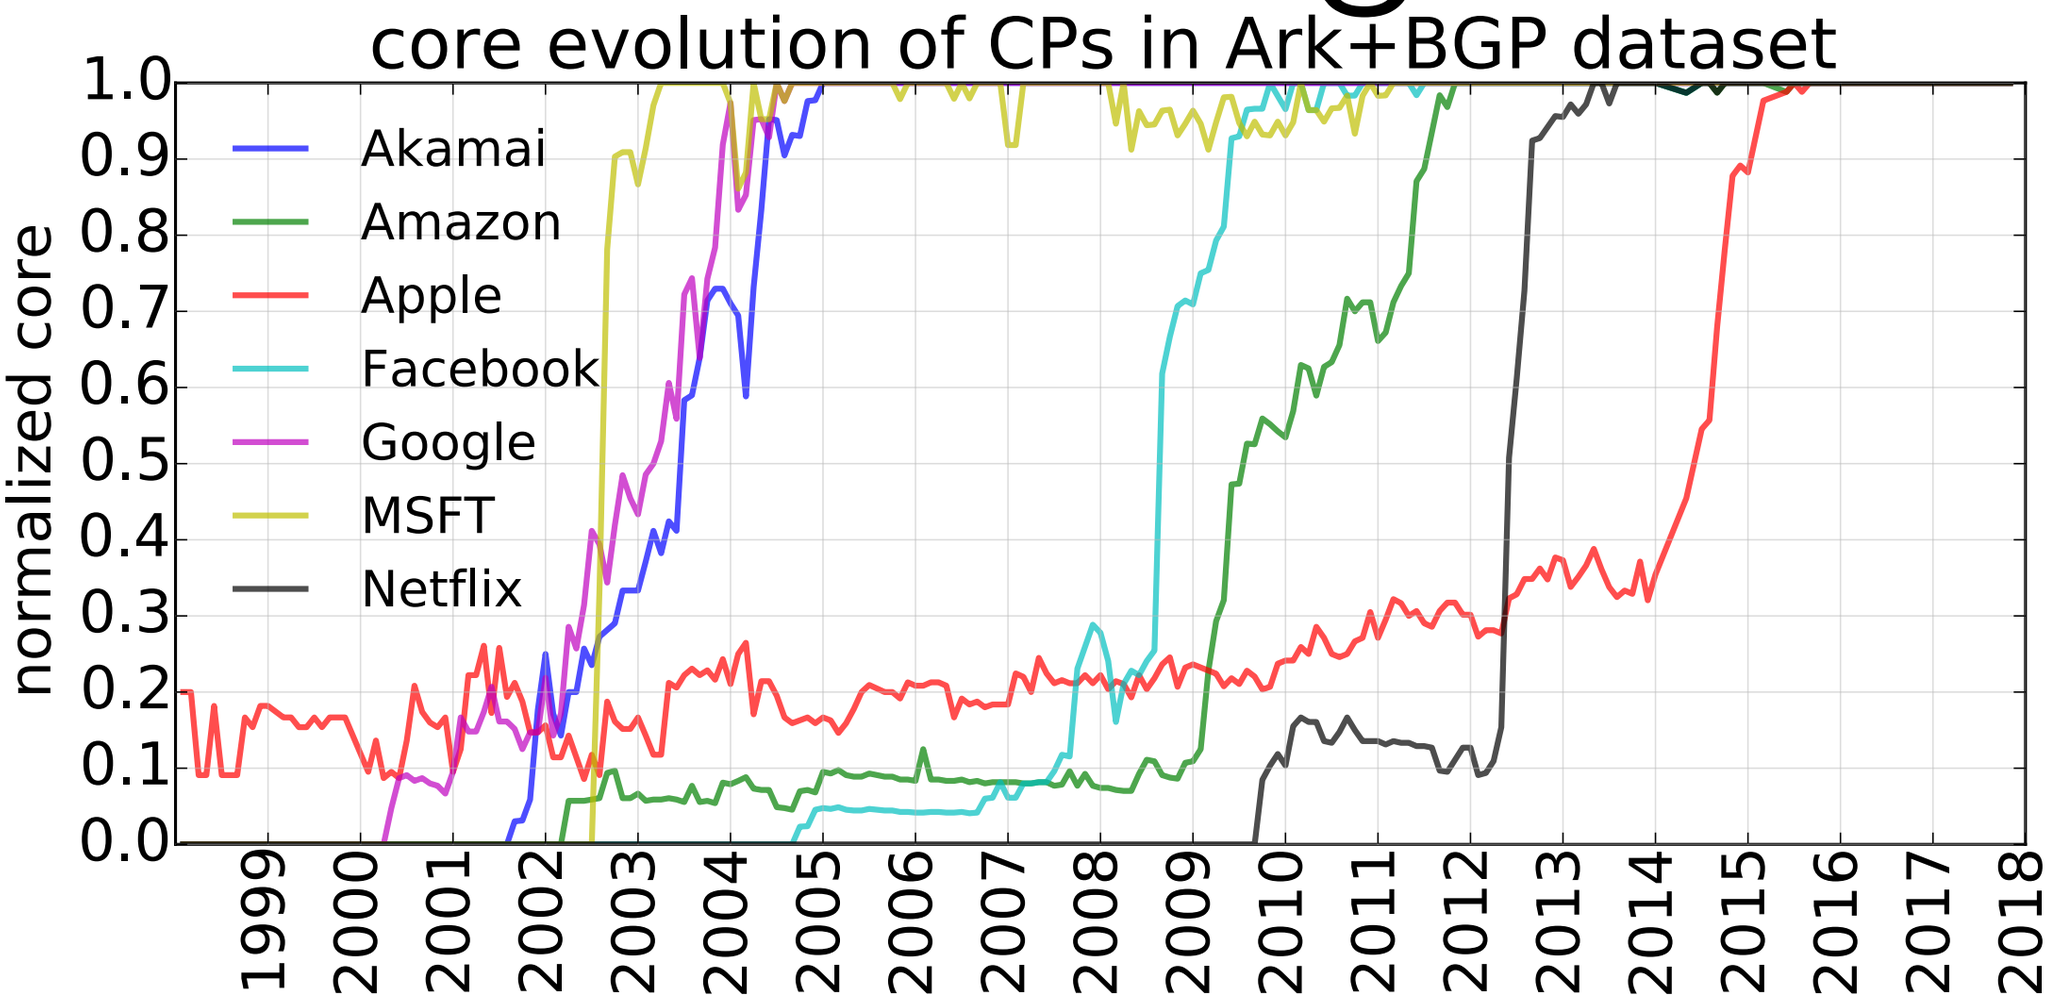

Evolution of big seven

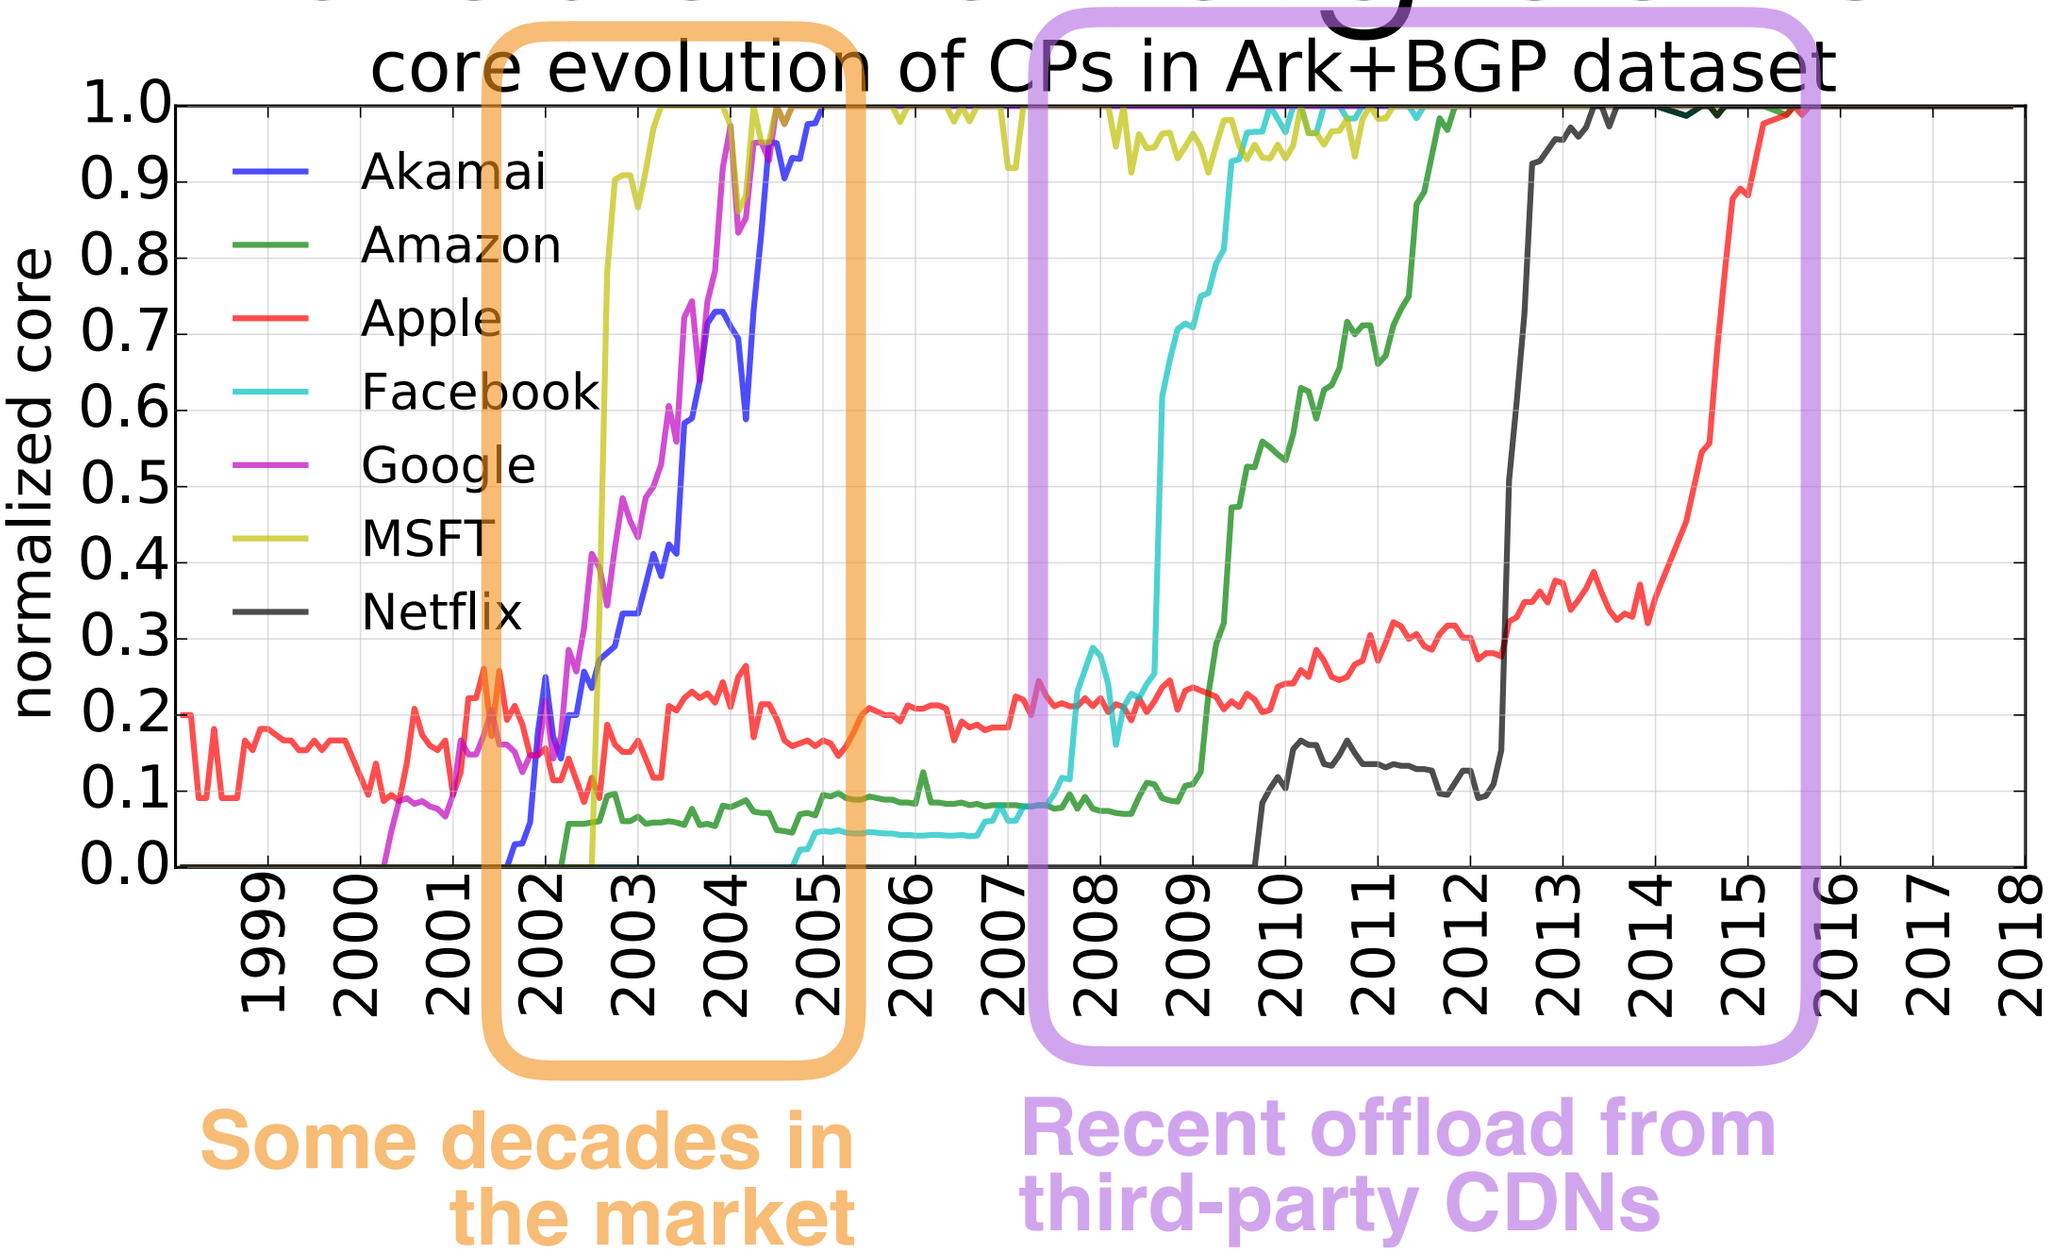

Evolution of Amazon

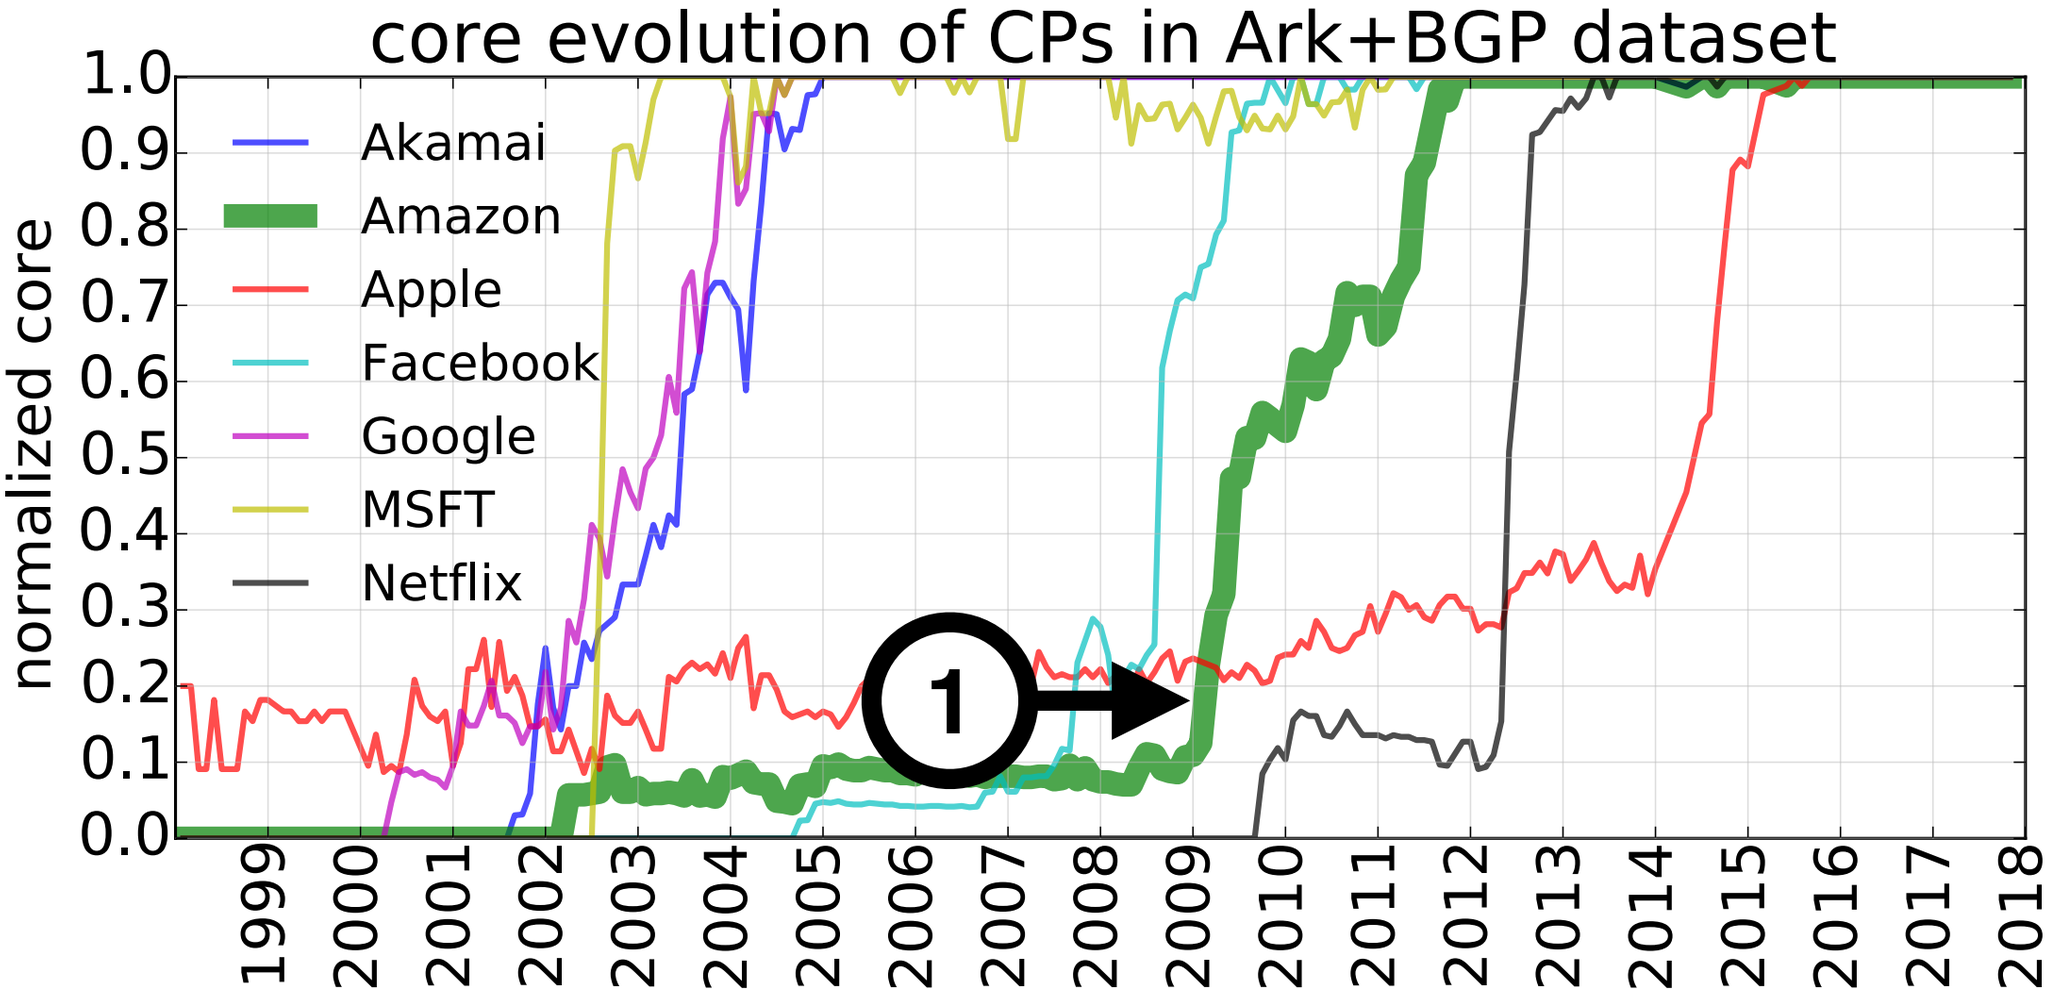

1. Release of DC in Northern California

Evolution of Amazon

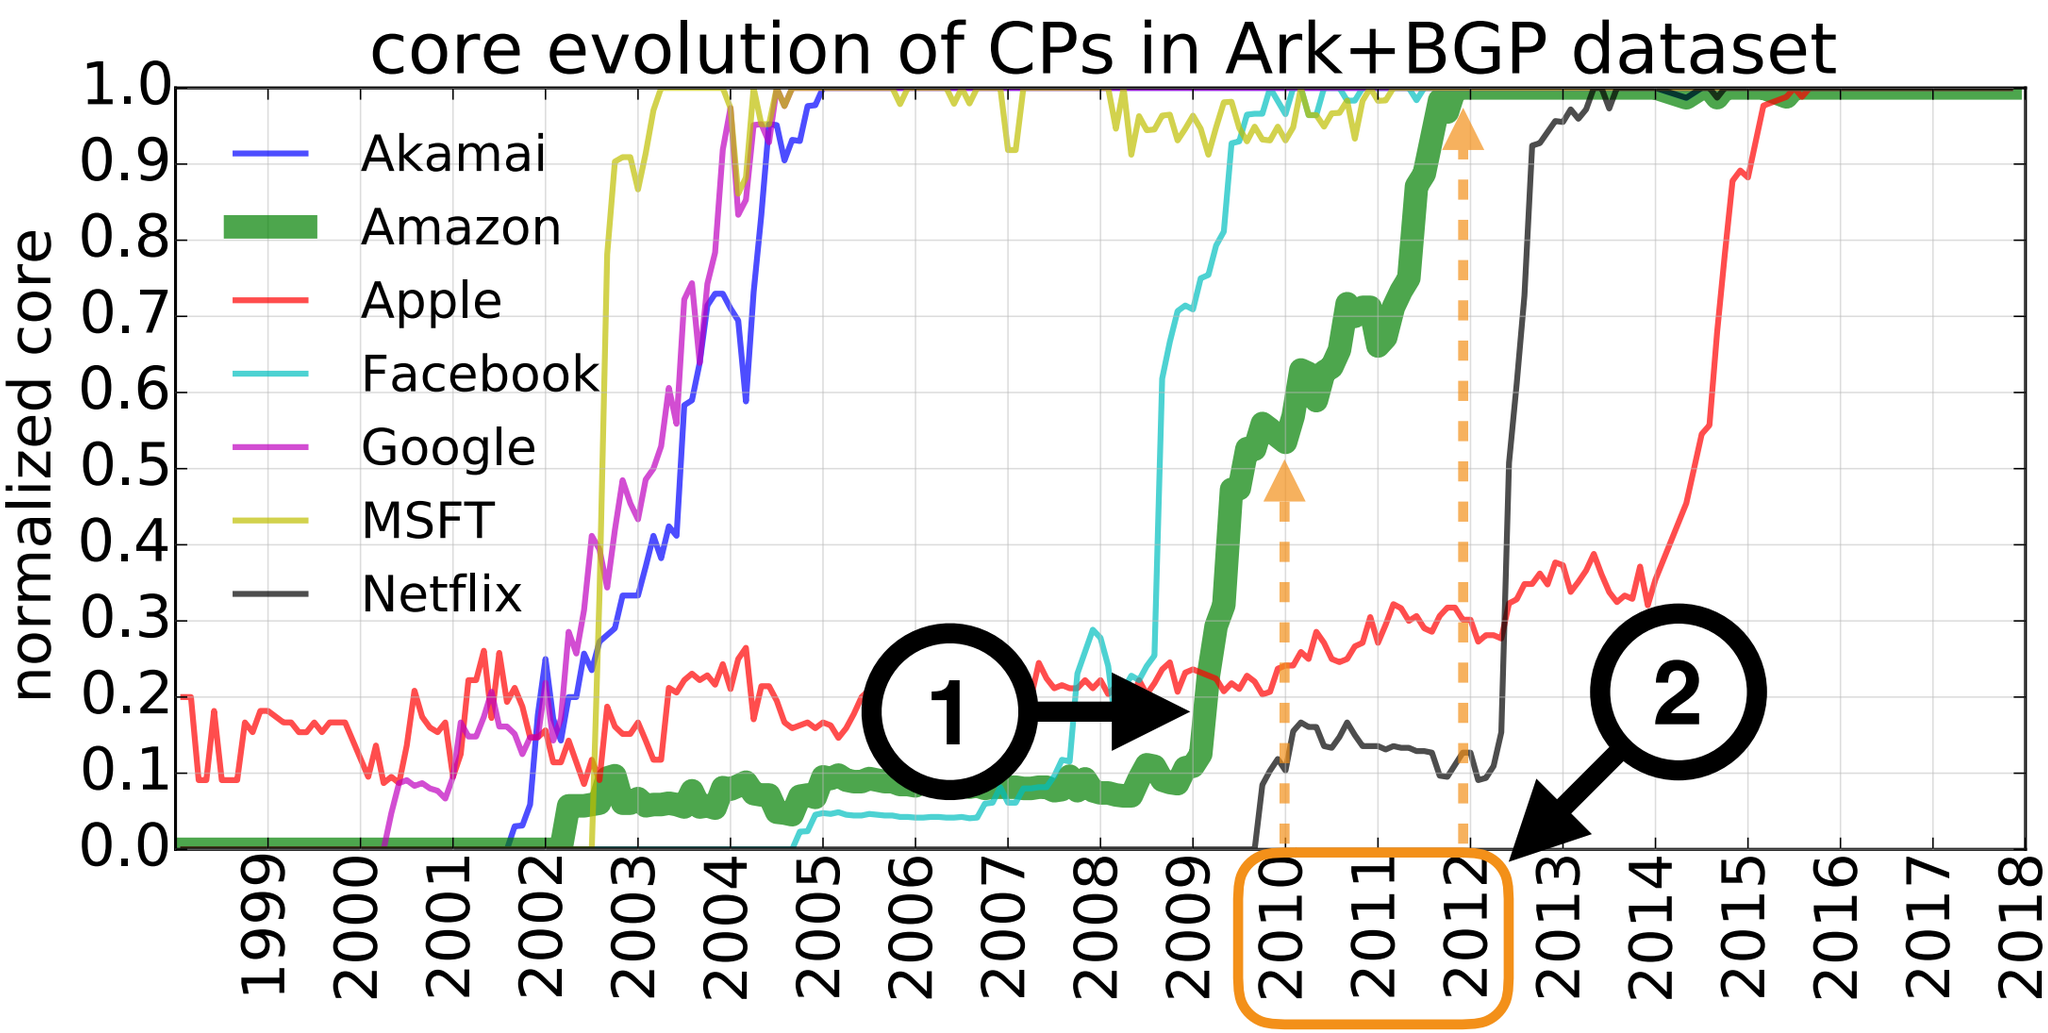

1. Release of DC in Northern California
2. Overseas expansion (Brazil, Europe & Asia)

Evolution of Amazon

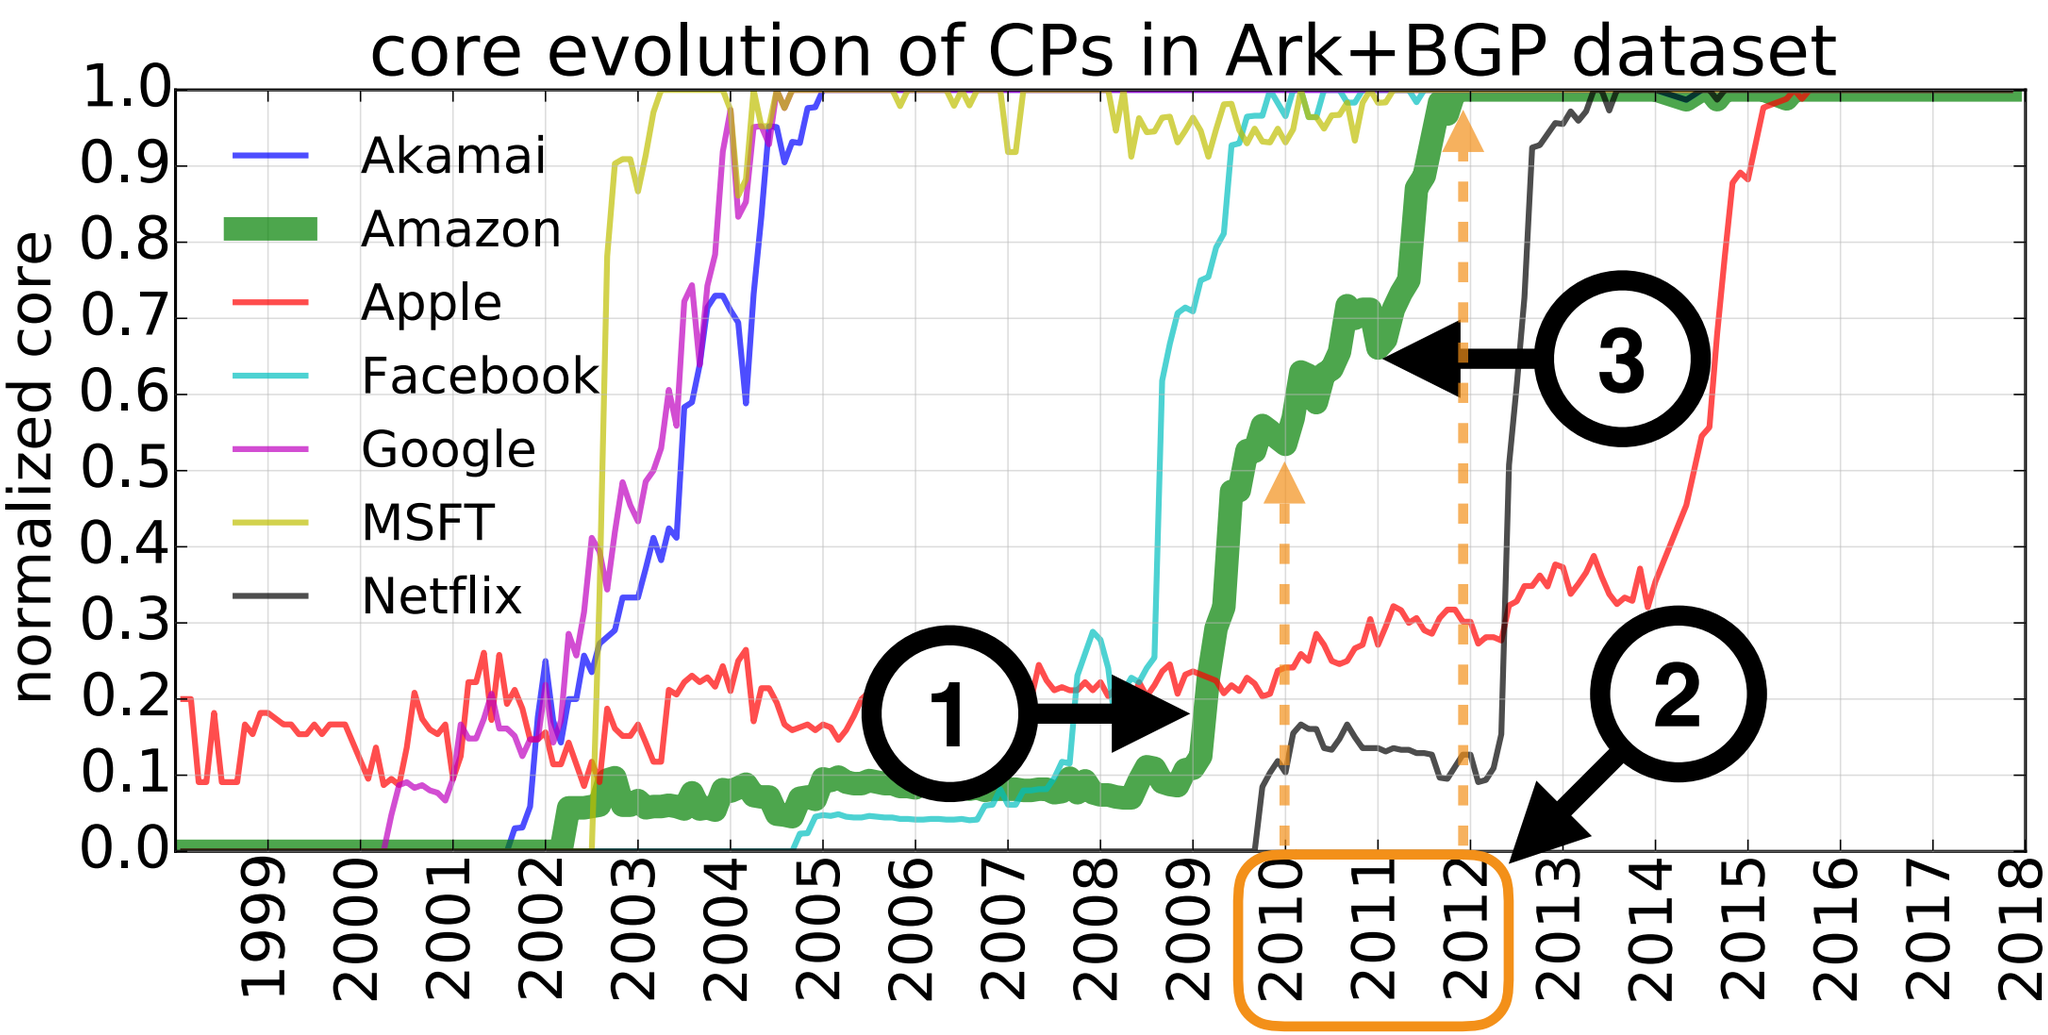

1. Release of DC in Northern California
2. Overseas expansion (Brazil, Europe & Asia)
3. Creation of WHOIS record for Amazon NS zone

CONICET

Evolution of Google

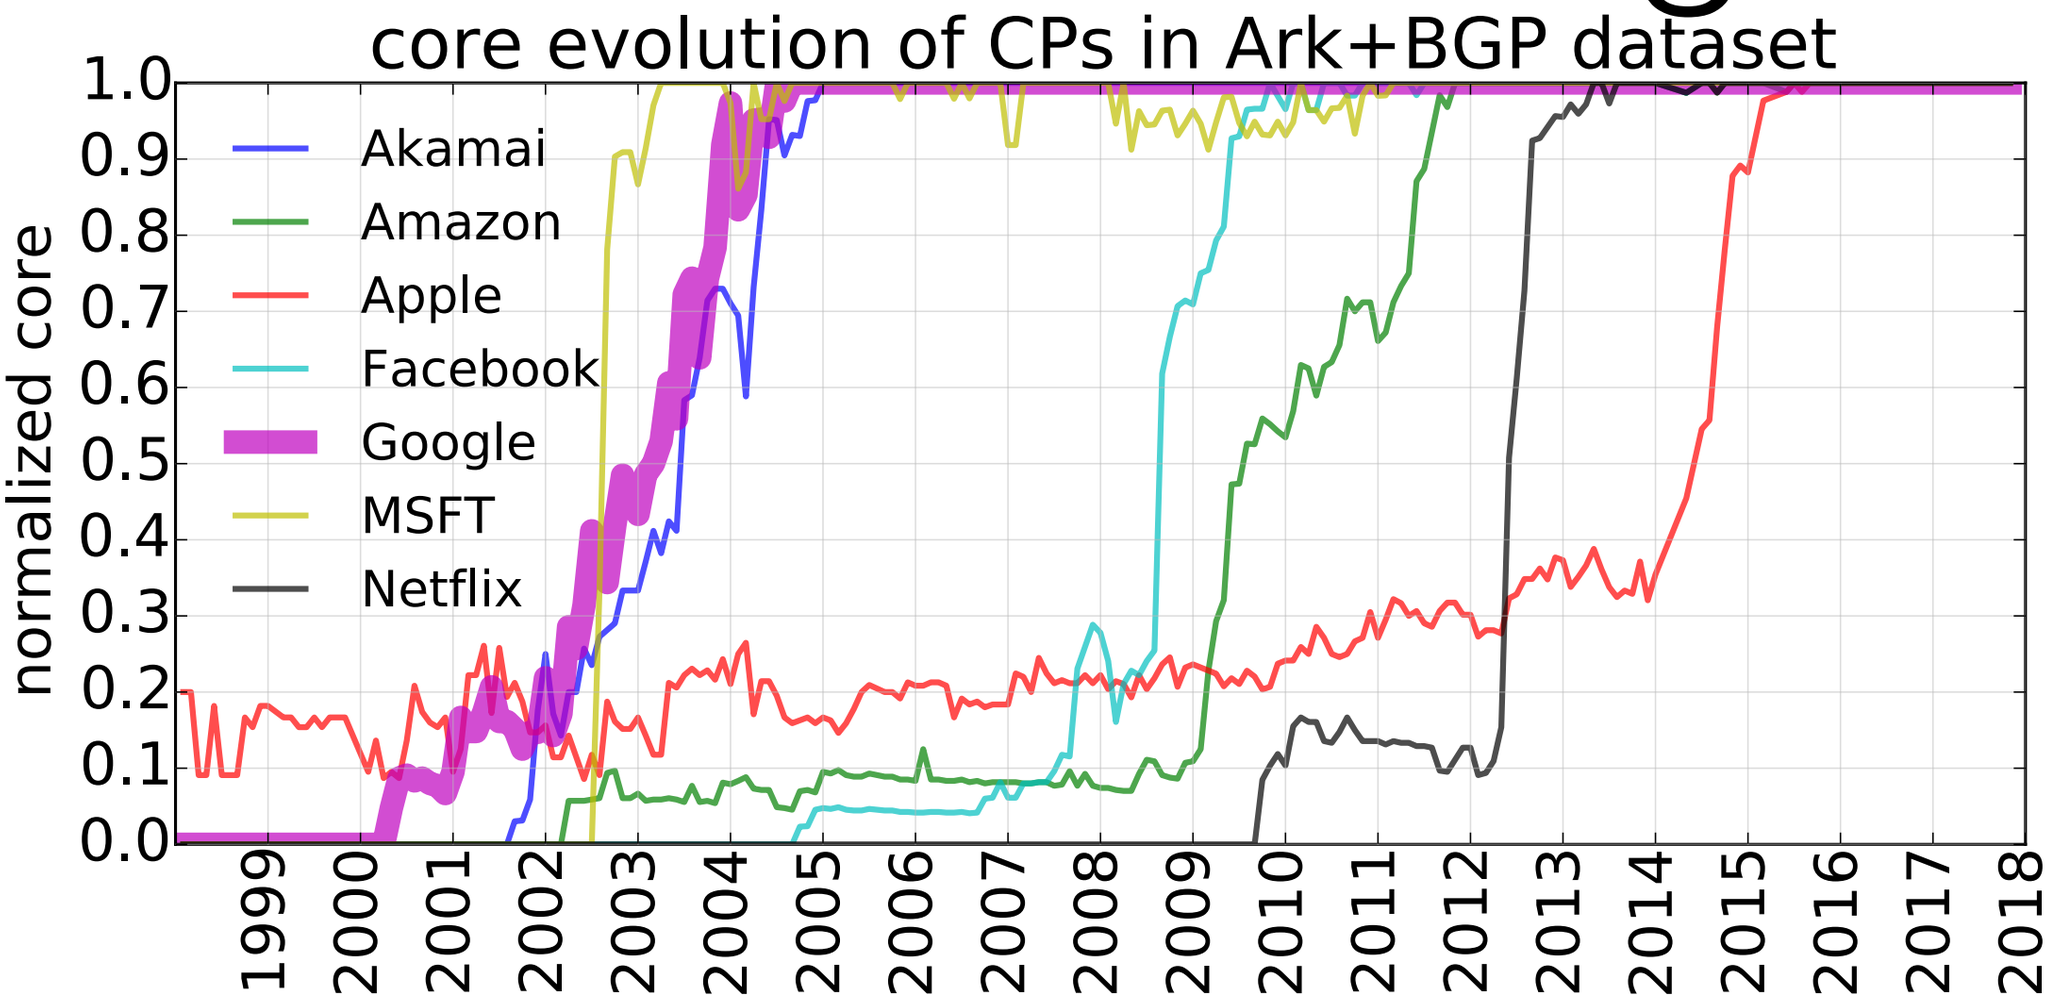

- Reached TOPcore before YouTube acquisition in 2006
- No CDN back then, but peer with TIER-1s

Evolution of Netflix

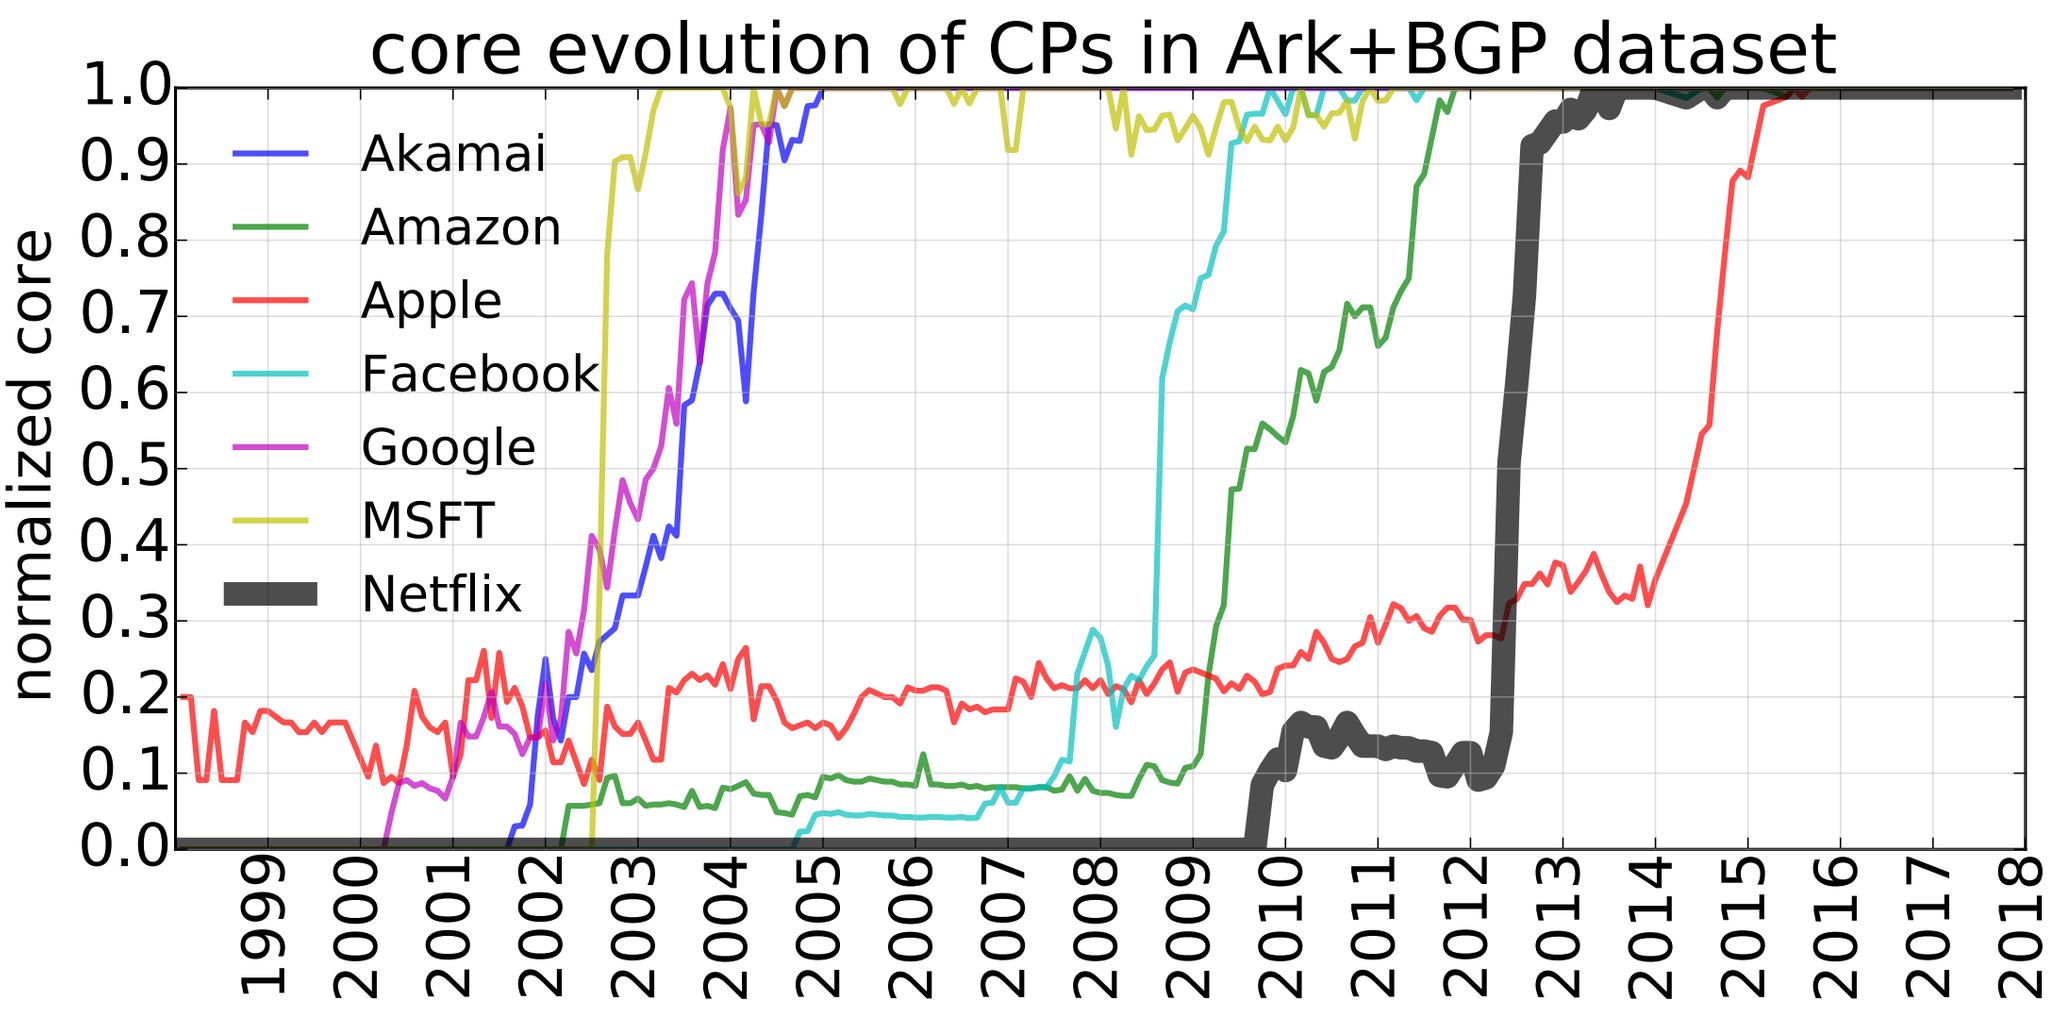

- Major source of traffic in the US since 2011
- In 2012 Netflix moved from Akamai to its own CDN
- OCA rollout led to some legal disputes

Analyses

1. Core evolution of large CPs *(Big Seven)* throughout the years
2. **Geographical differences on the core evolution**
3. Other members of the core

Analyses

1. Core evolution of large CPs *(Big Seven)* throughout the years
2. Geographical differences on the core evolution

3. Other members of the core

Other CPs in the core

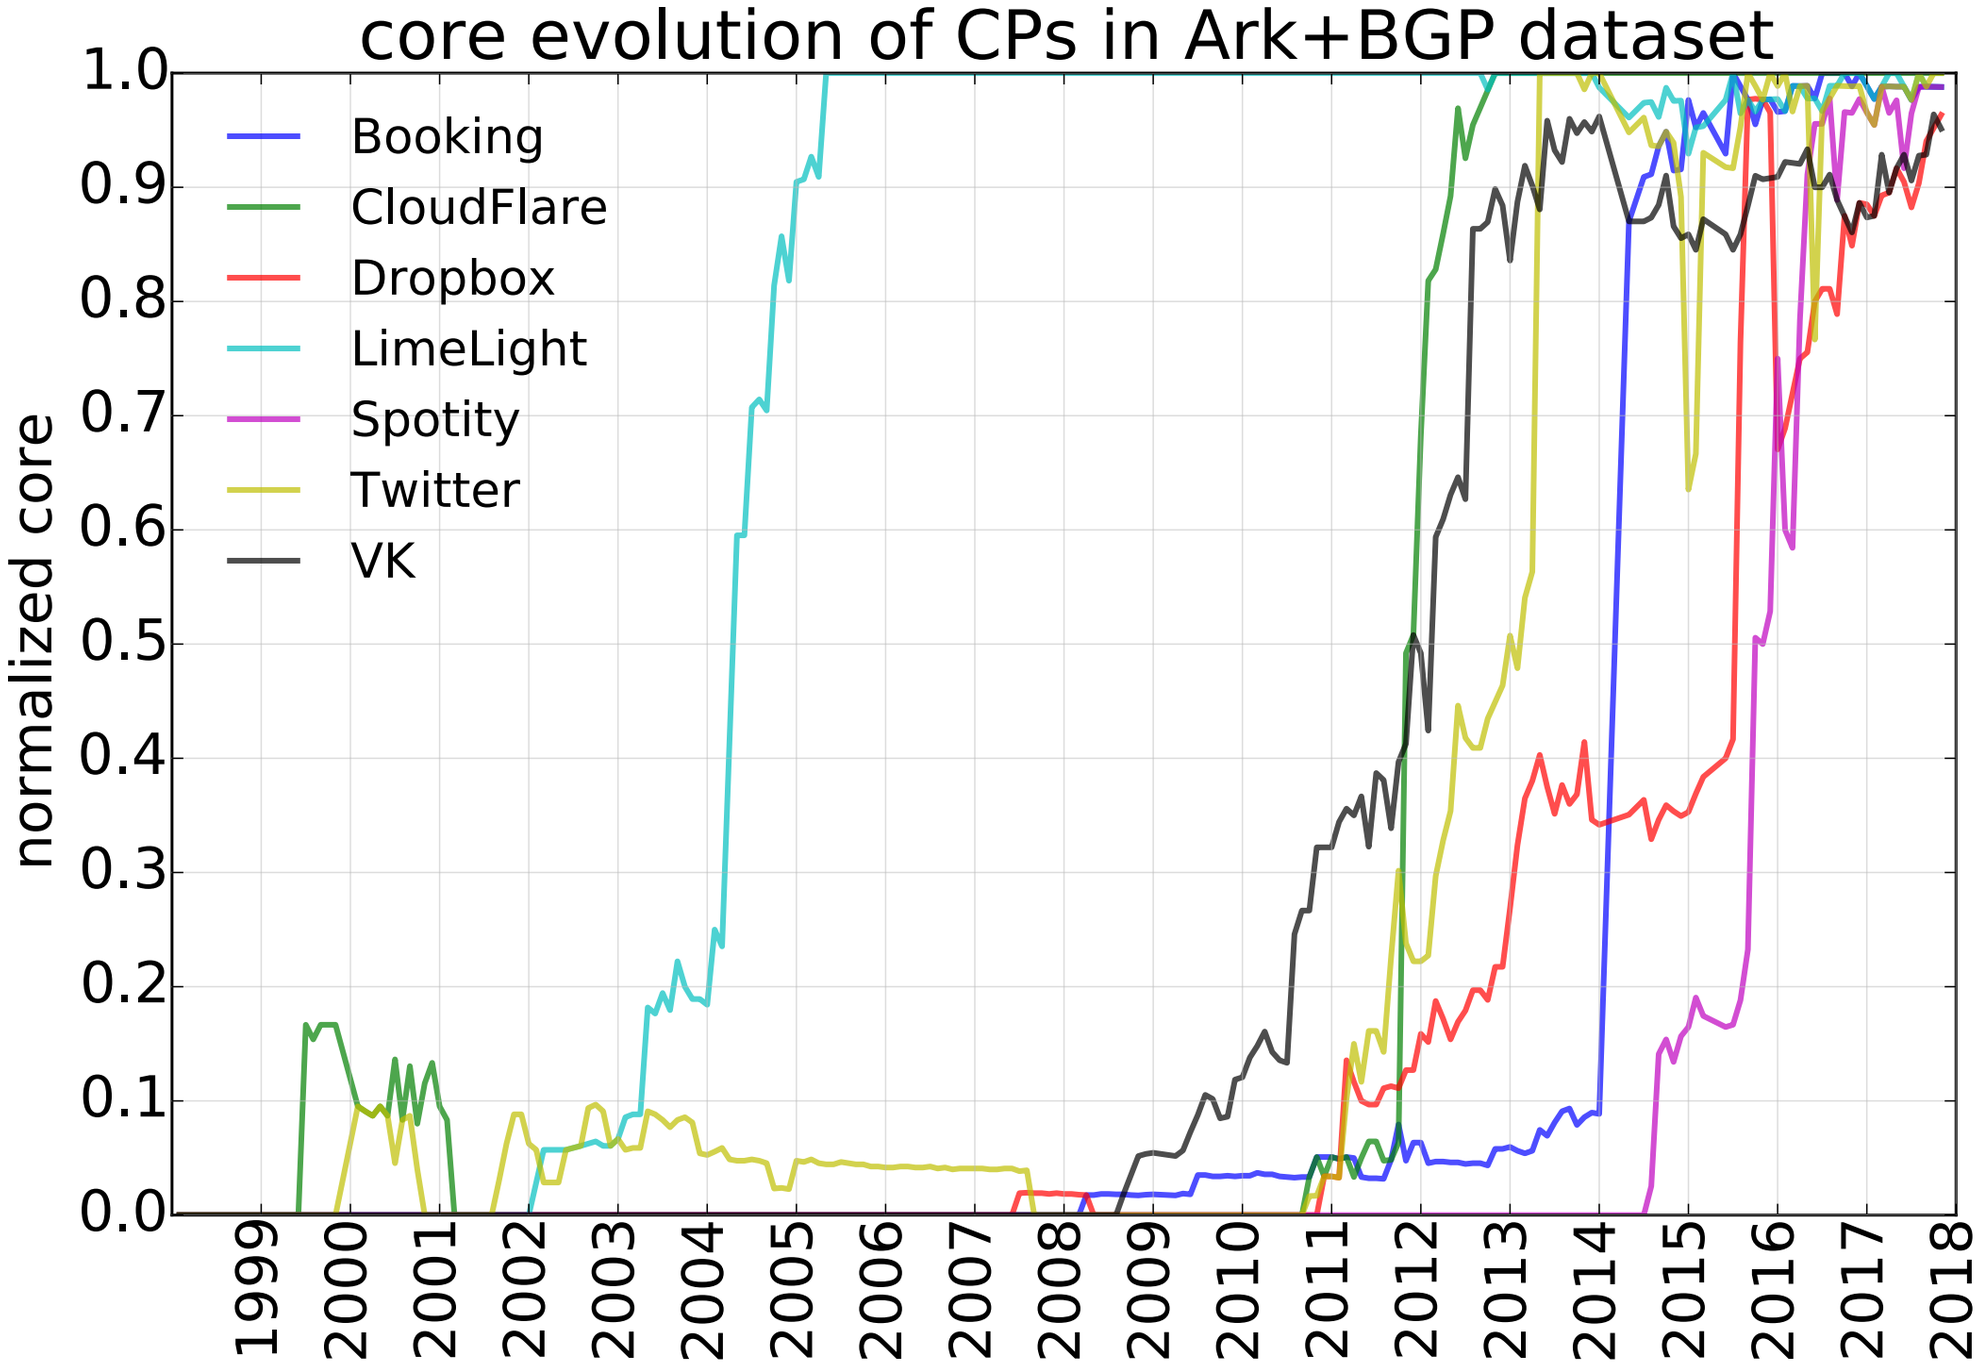

Arrival at the core

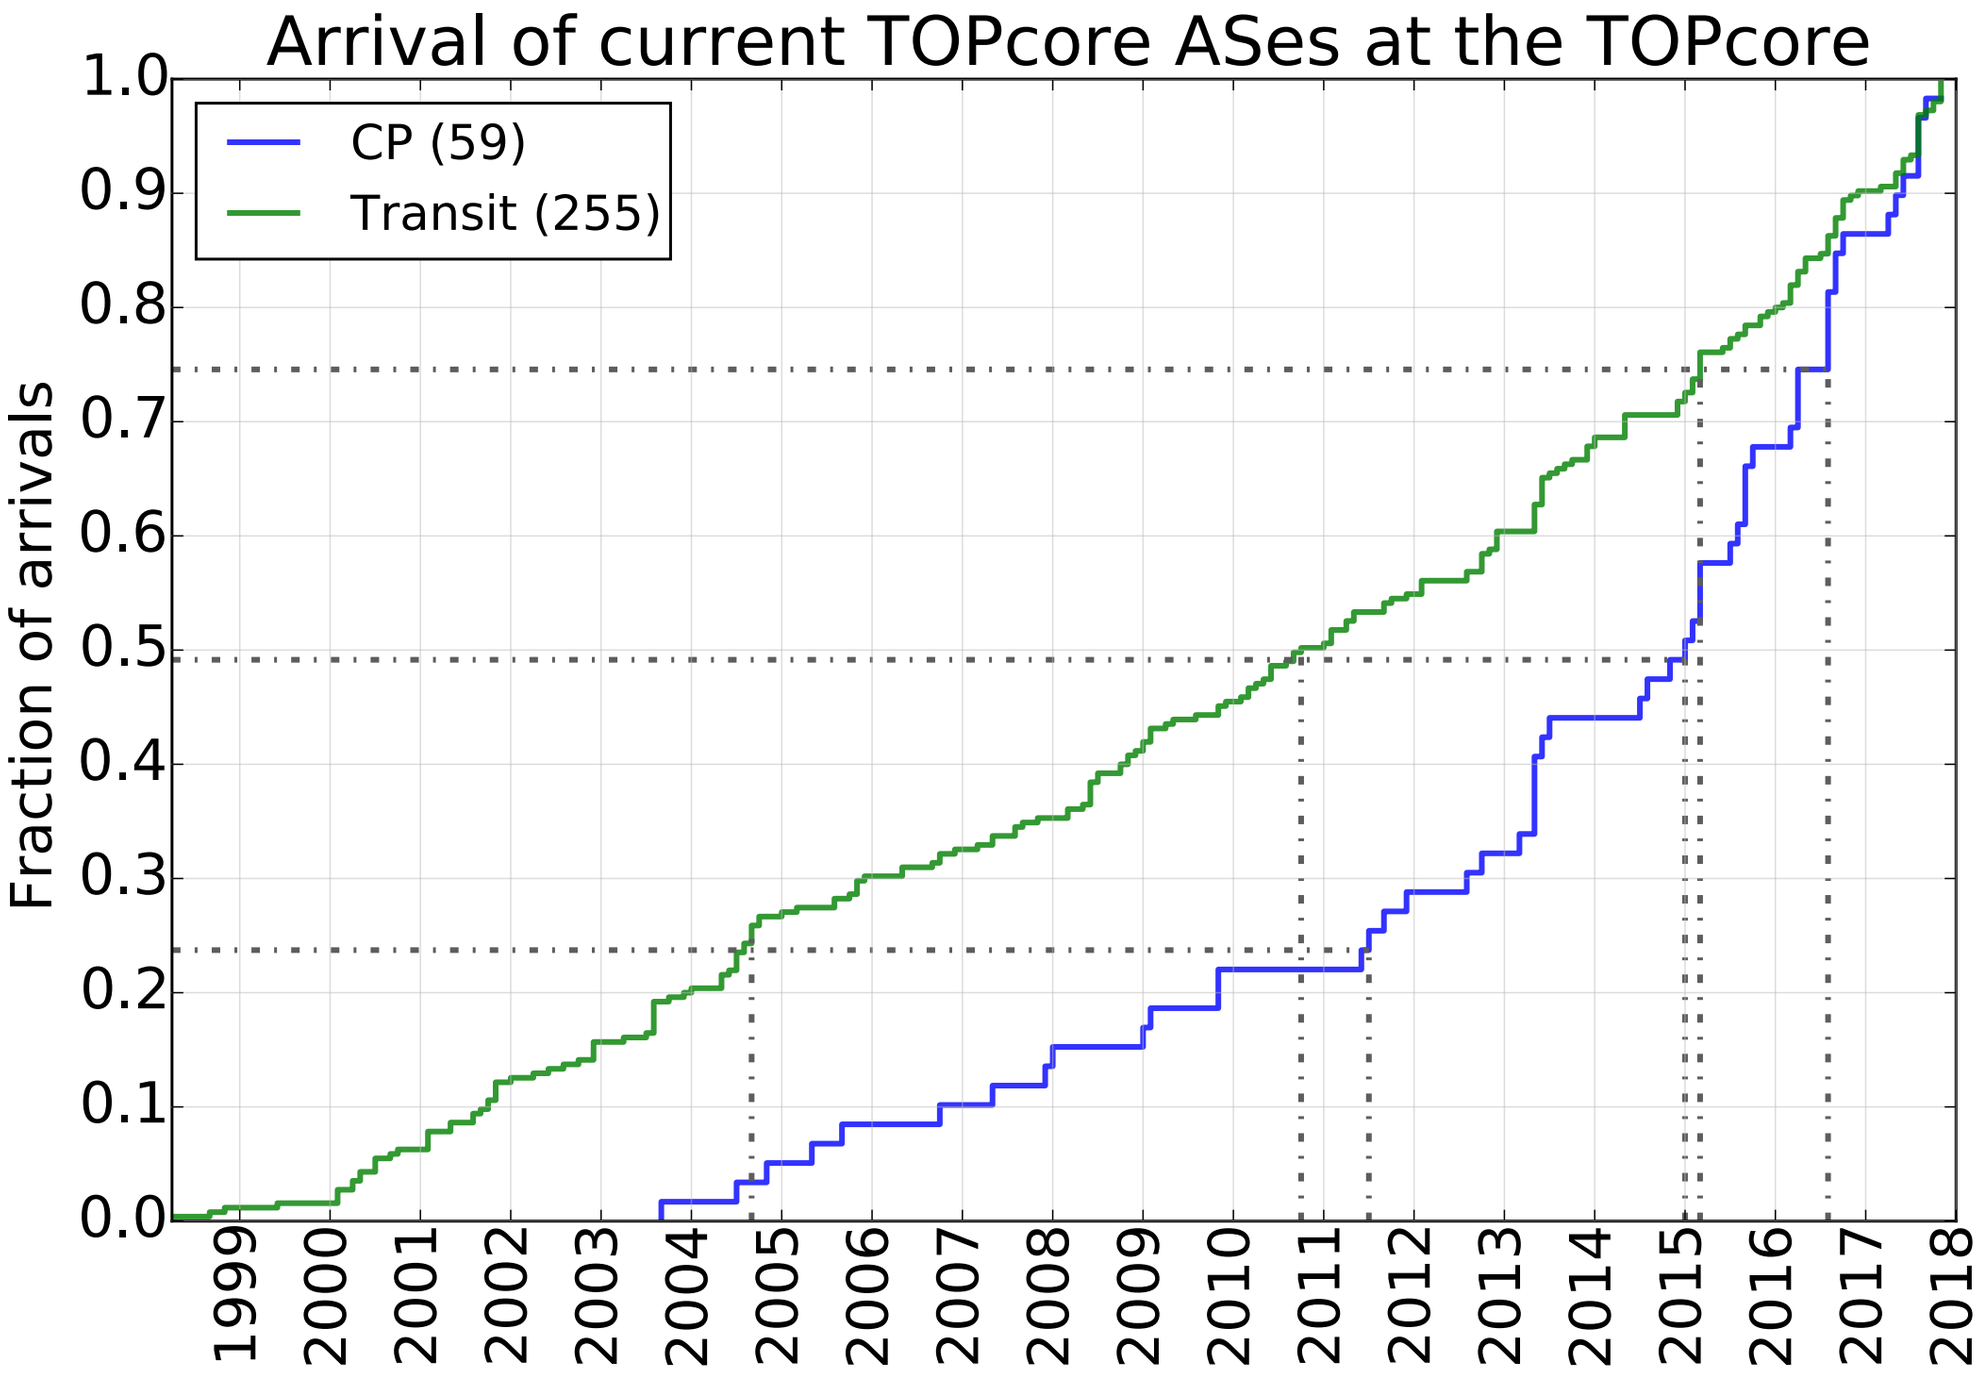

Conclusions

- Large CPs deploy and run densely connected ASes
- Content is moving to private CDNs
- Match expansion of private CDNs with business strategies
- Create a website to allow users to look for other patterns

Website

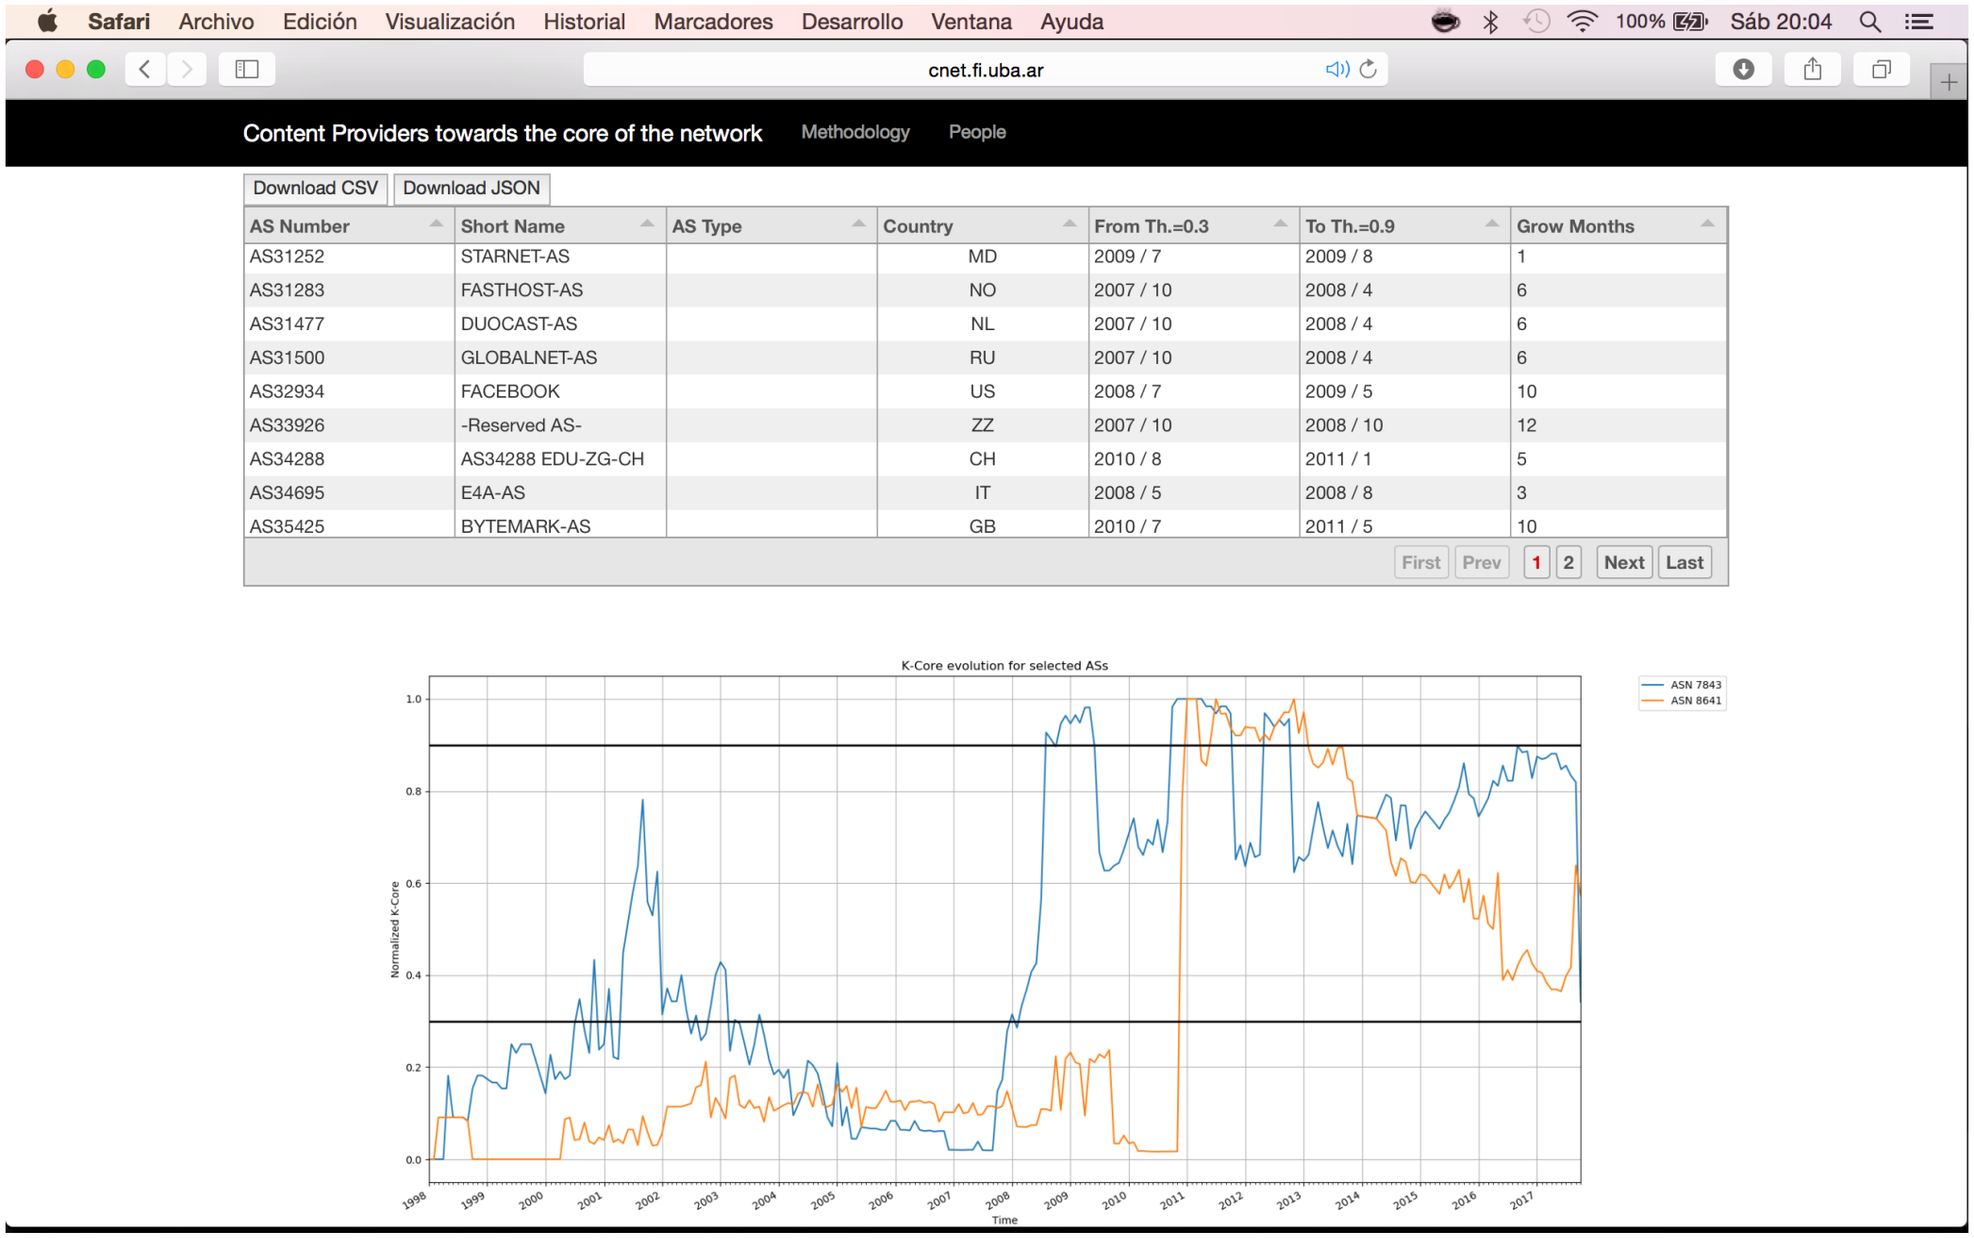

Visit: http://cnet.fi.uba.ar/TMA2018

Thanks!

Questions?

CONICET