Network Neutrality, Service Differentiation, Harms and Benefits

David Clark, CSAIL, MIT

April, 2015

Quick outline

- History of neutrality
  - A bad choice of words…
- How QoS mechanisms work
- Fears and benefits of QoS
- Real-world disputes: Netflix vs. Comcast
- Neutral harms
- Alternative platforms—scope of regulation.
- (if we have time): the pain of the FCC.

History of neutrality

- The Internet has never been “neutral”.
  - We gave priority to interactive apps in 1990.
  - Why? They worked better…
- Mechanisms designed and standardized in 1990’s.
- Total failure to deploy in public Internet.
  - Used today in private networks.
  - Issues are economic and coordination.
    - Who gets paid how much?
    - How debug failures?

Design approaches

- When congestion occurs, reorder pending packets.
  - You heard about this in lecture 12.
  - Obvious point—if system not overloaded, no queue, nothing to manage.
- Find preferable routes.

Queues

- Priority
  - Put traffic into different queues.
  - Service in order of priority.
  - Consequence: shift the variation in delay (jitter) onto less demanding applications.
    - If mis-configured, can starve low priority queues
- Isolation (weighted fair queuing)
  - Put traffic into different queues.
  - Service according to schedule
  - Consequence: flows with “enough” allocation, no queue, so no jitter.
    - If mis-configured, losing queue has bad day. All apps lose.

Alternative routes

- In Internet today (especially among Autonomous Systems) no route diversity based on performance.
  - Too complex to compute
  - Not clear what parameters would drive diversity
  - Not clear whether traffic is too dynamic.
- CDNs do version of this all the time.
  - Pick source that works best.
  - Nothing the ISPs can control.
- We had wrong conception:
  - Single source-single destination is not what is happening.

Why bother?

- Bandwidth is cheap. Just overprovision.
  - True most of the time.
  - But two important exceptions.
- Mobile (cellular) service.
  - Limited spectrum, limited backhaul, lots of congestion.
- Home access link.
  - Sold by peak speed, so queues will form there.
    - Capacity allocation using weighted fair queuing…
  - Peak speed caps may not be best idea, but easy to market.
  - Secondary problem, over-large buffers.

Benefits and fears

- Benefits: when queues *do* form, differentiation can materially improve performance of certain apps.
  - Latency and jitter sensitive apps
  - Skype, games, etc.
- Fears: these tools, if deployed, will be used to impose business-related discrimination against select activities.
  - Hurt Netflix vs. their own product, or Hulu, etc.
  - Ability to monetize QoS will reduce incentive to add capacity. (Intentional under-provisioning)

Real-world issues

- The disputes between Netflix and various ISPs (but Comcast is the most well-known).
- Netflix delivers 1/3 of all off-net traffic coming into access ISPs like Comcast (in the U.S.).
- No matter how that traffic is delivered, the paths have to be especially engineered.
  - Netflix is sort of BYOC (bring your own congestion)
  - Also true for Youtube, big content delivery networks (CDNs).

“Old” concept of interconnection

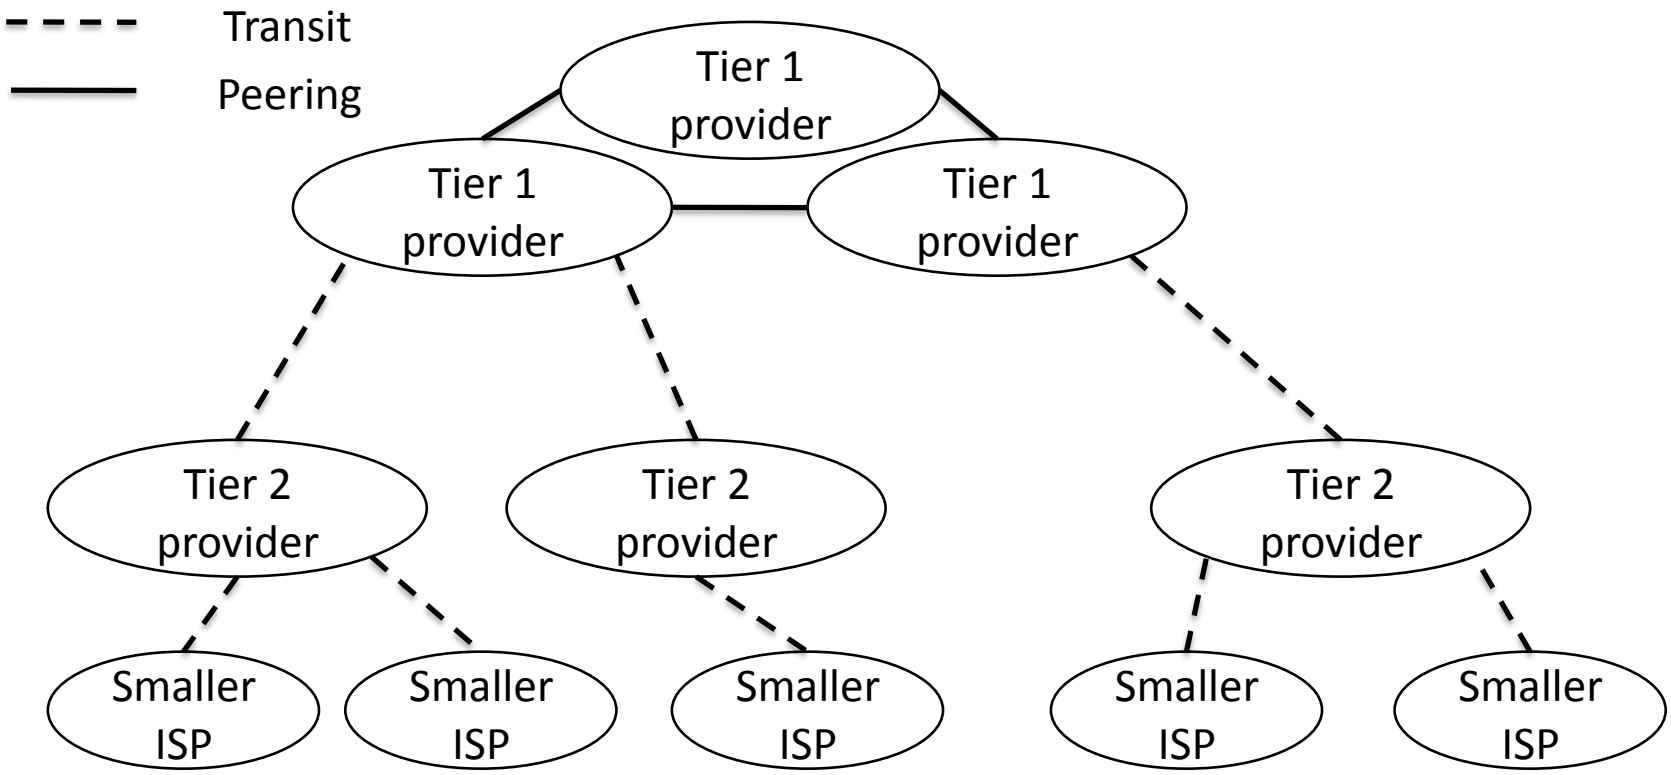

More realistic picture

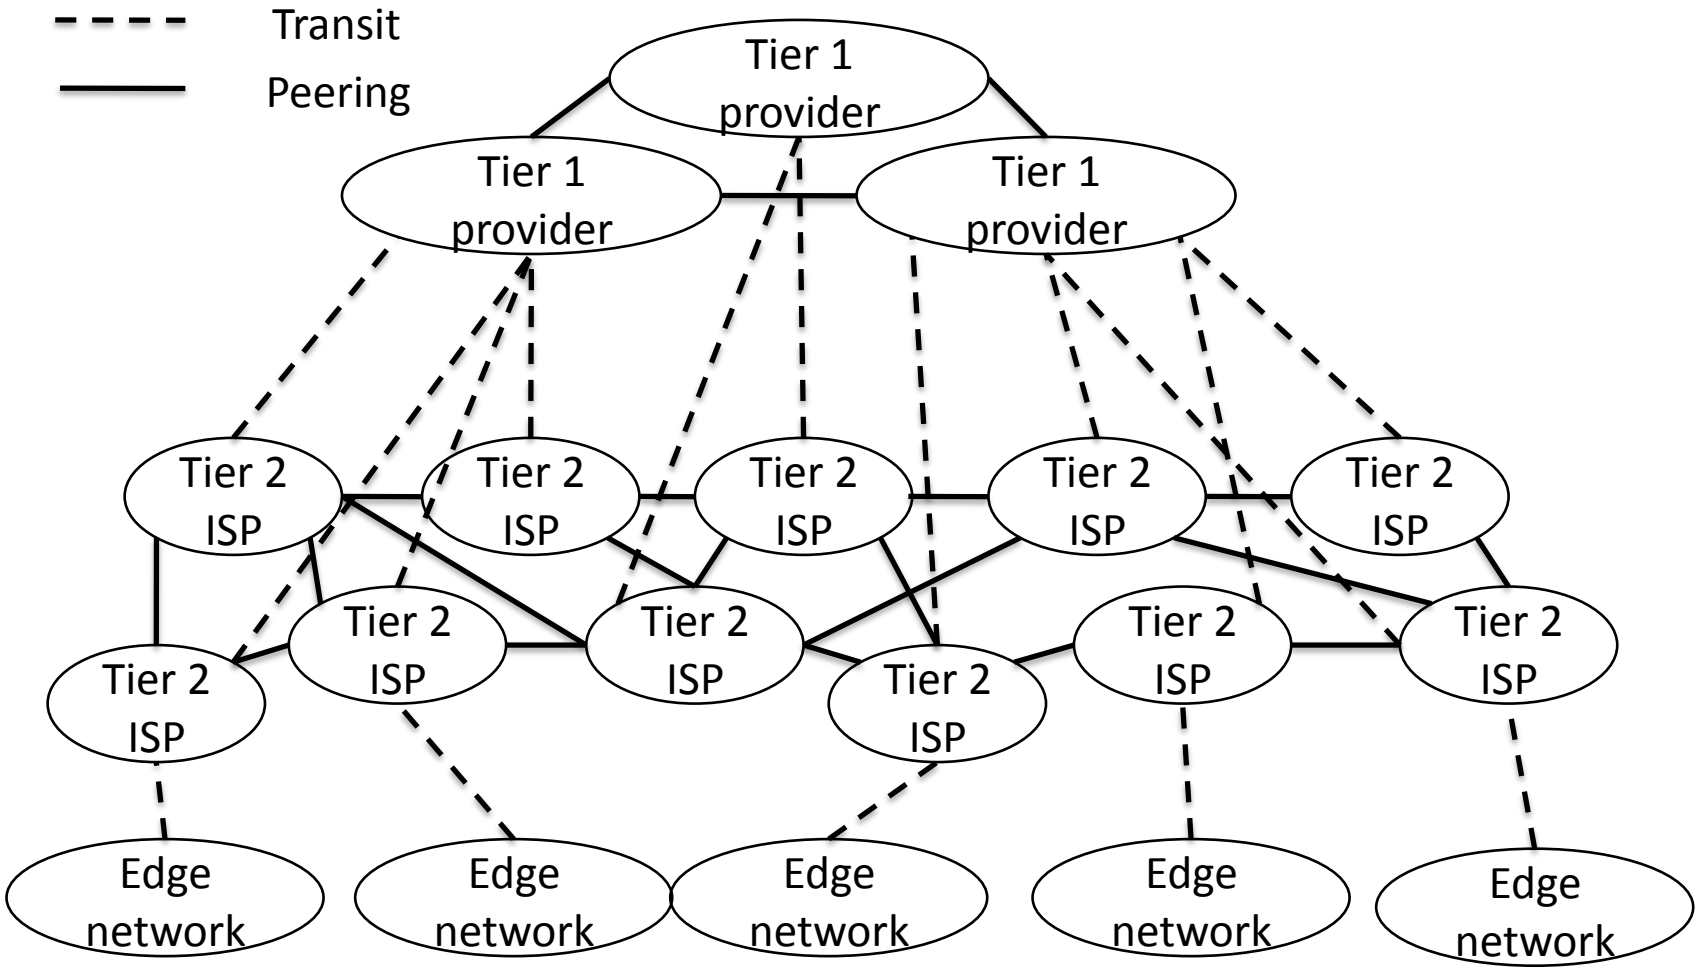

More realistic picture

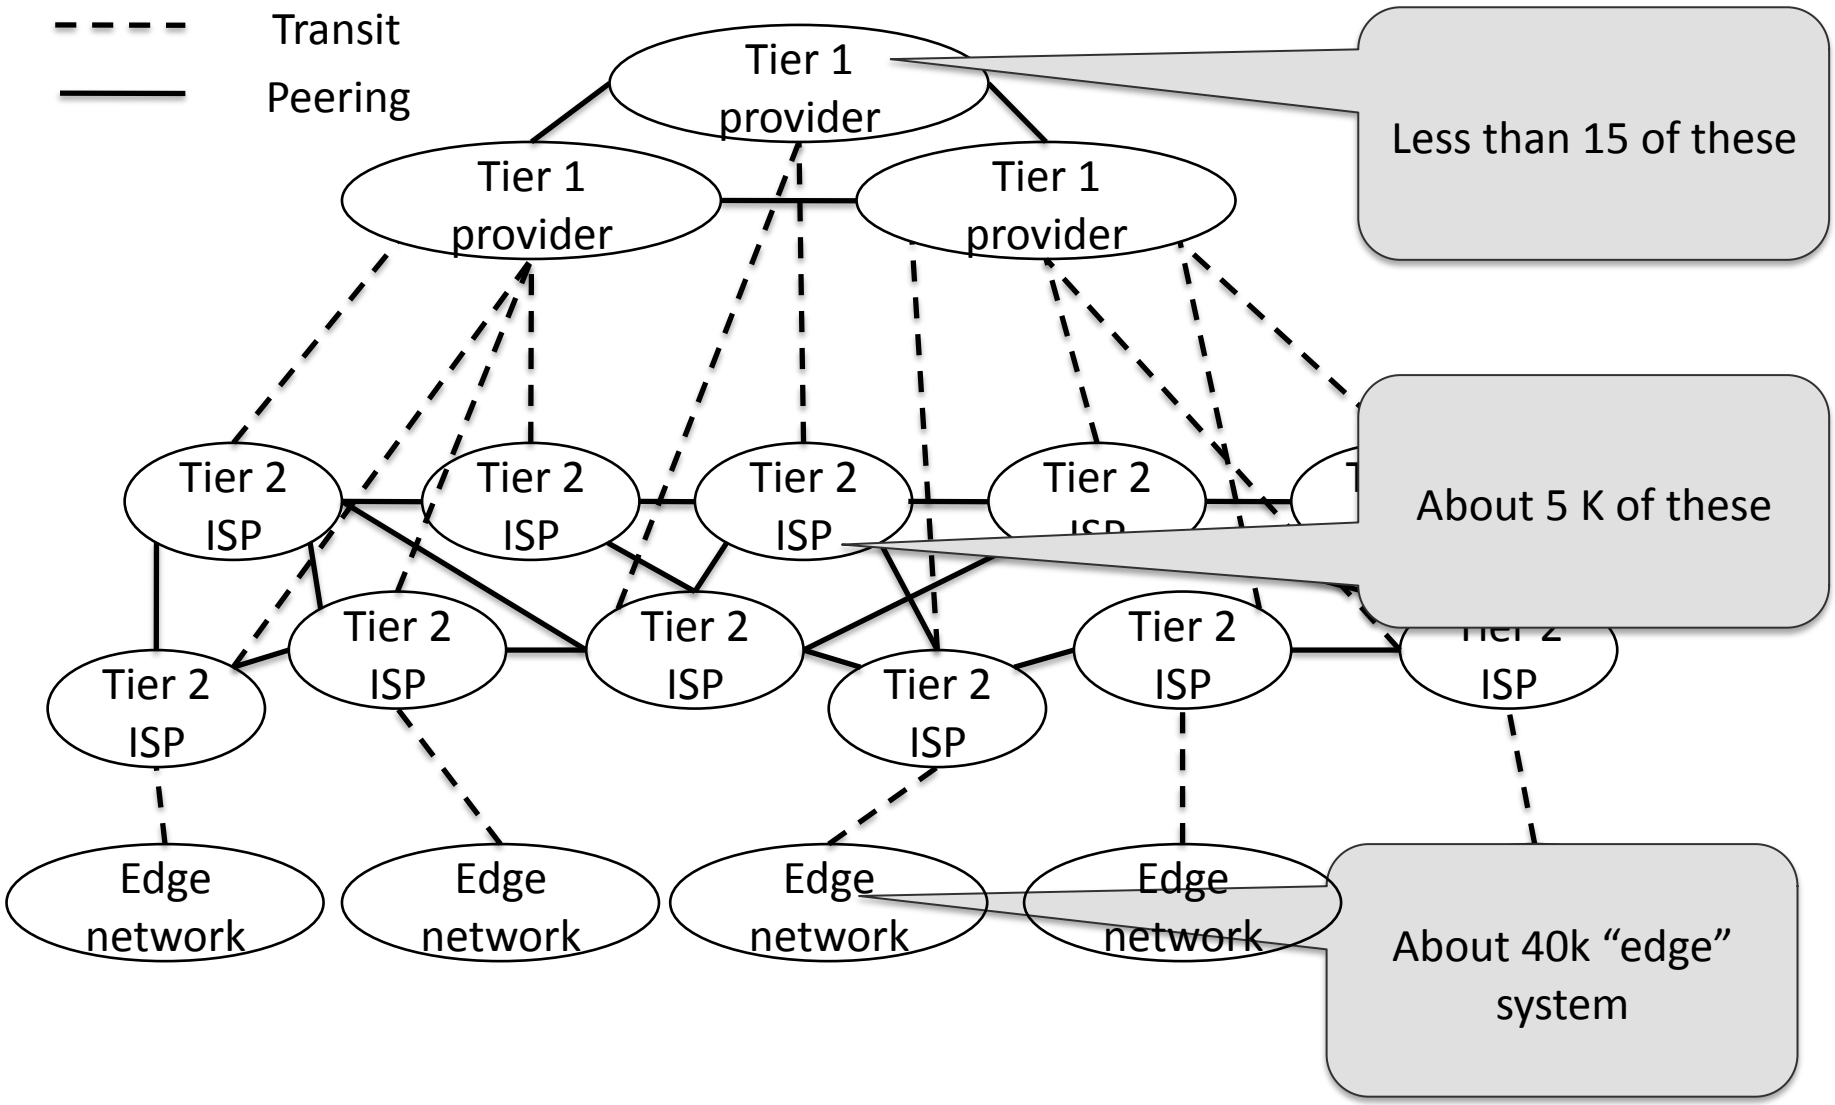

More realistic picture

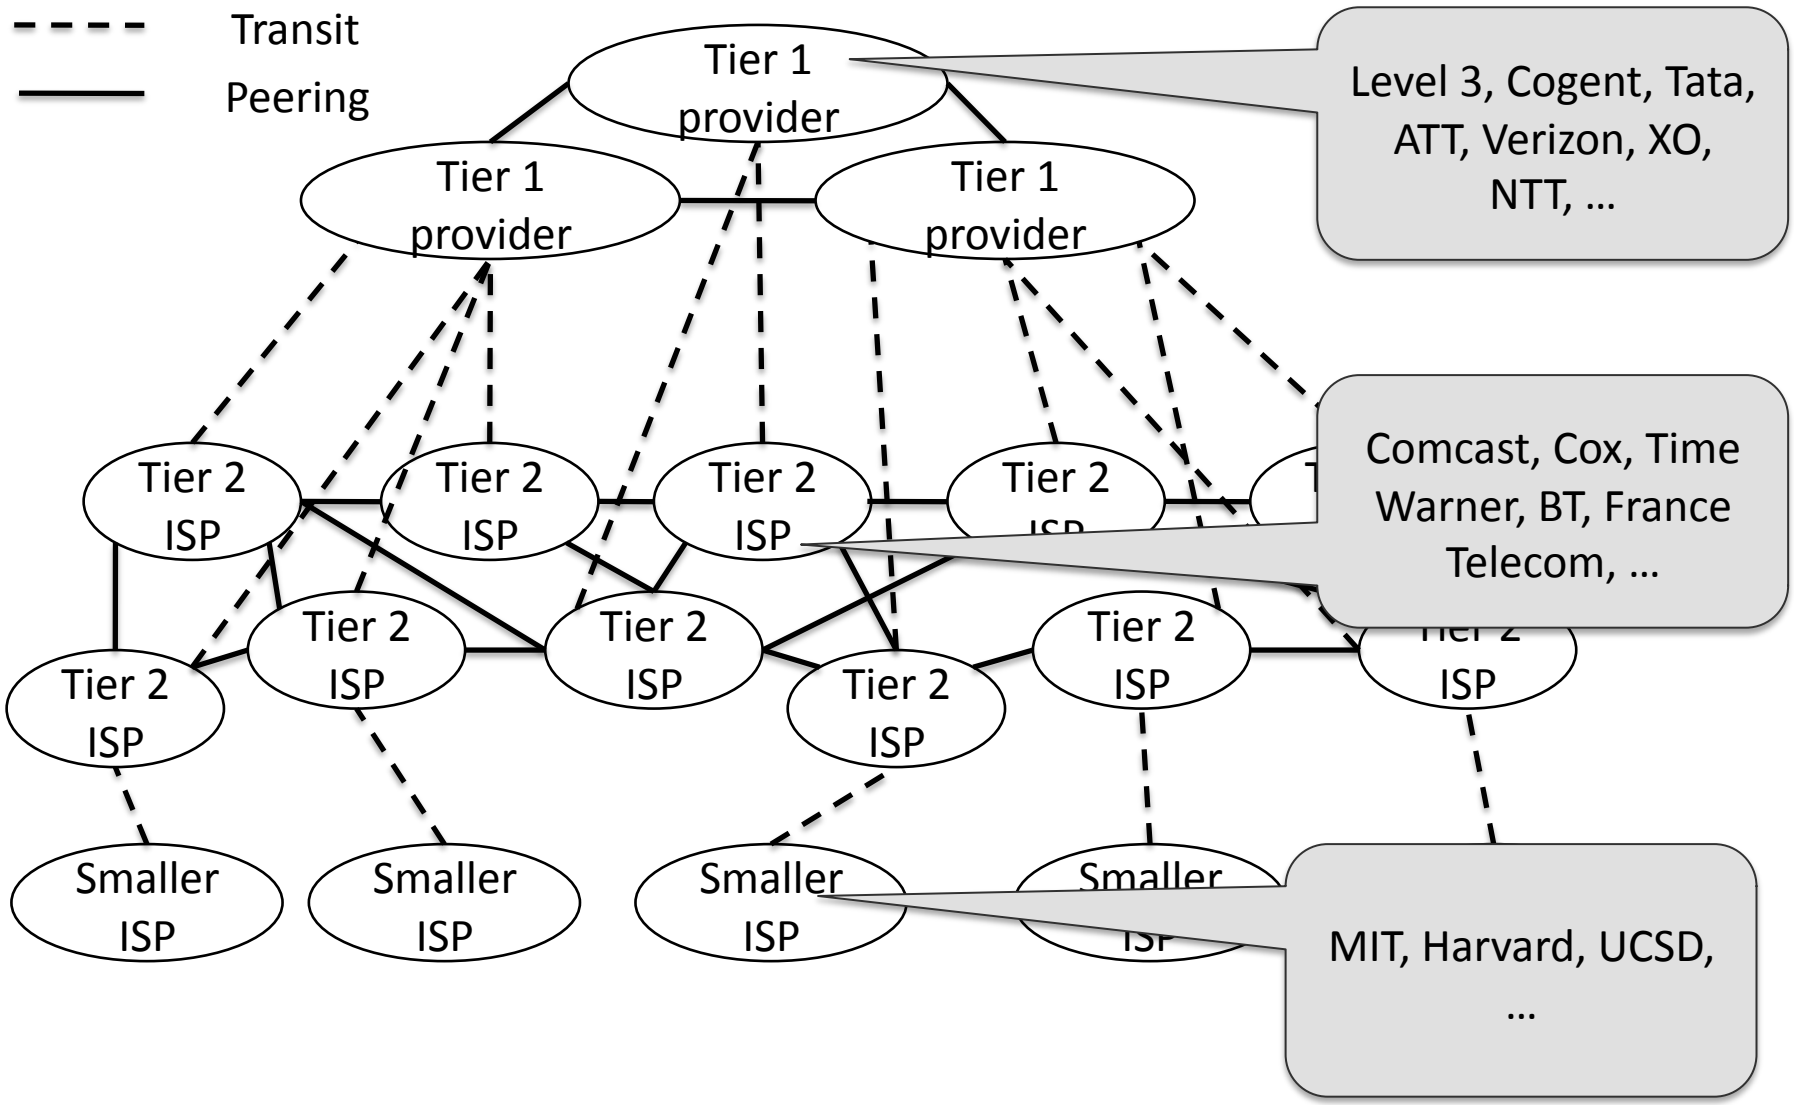

More realistic picture

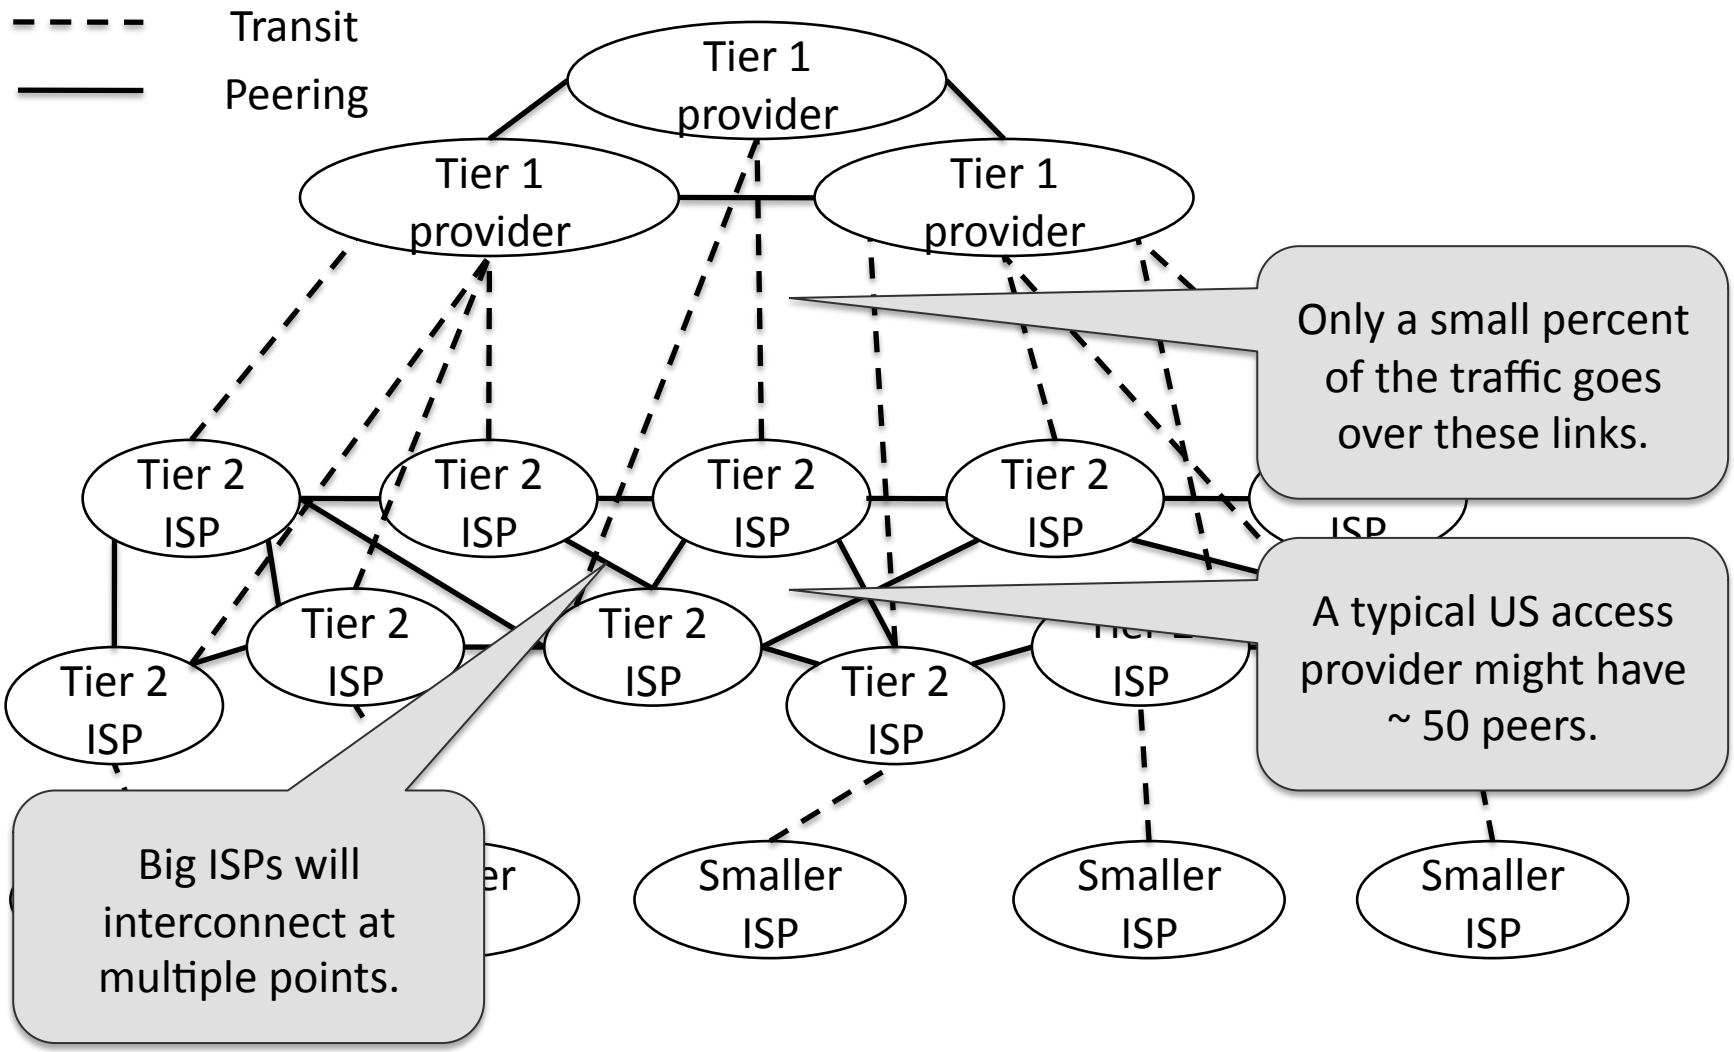

Interconnection of CDN

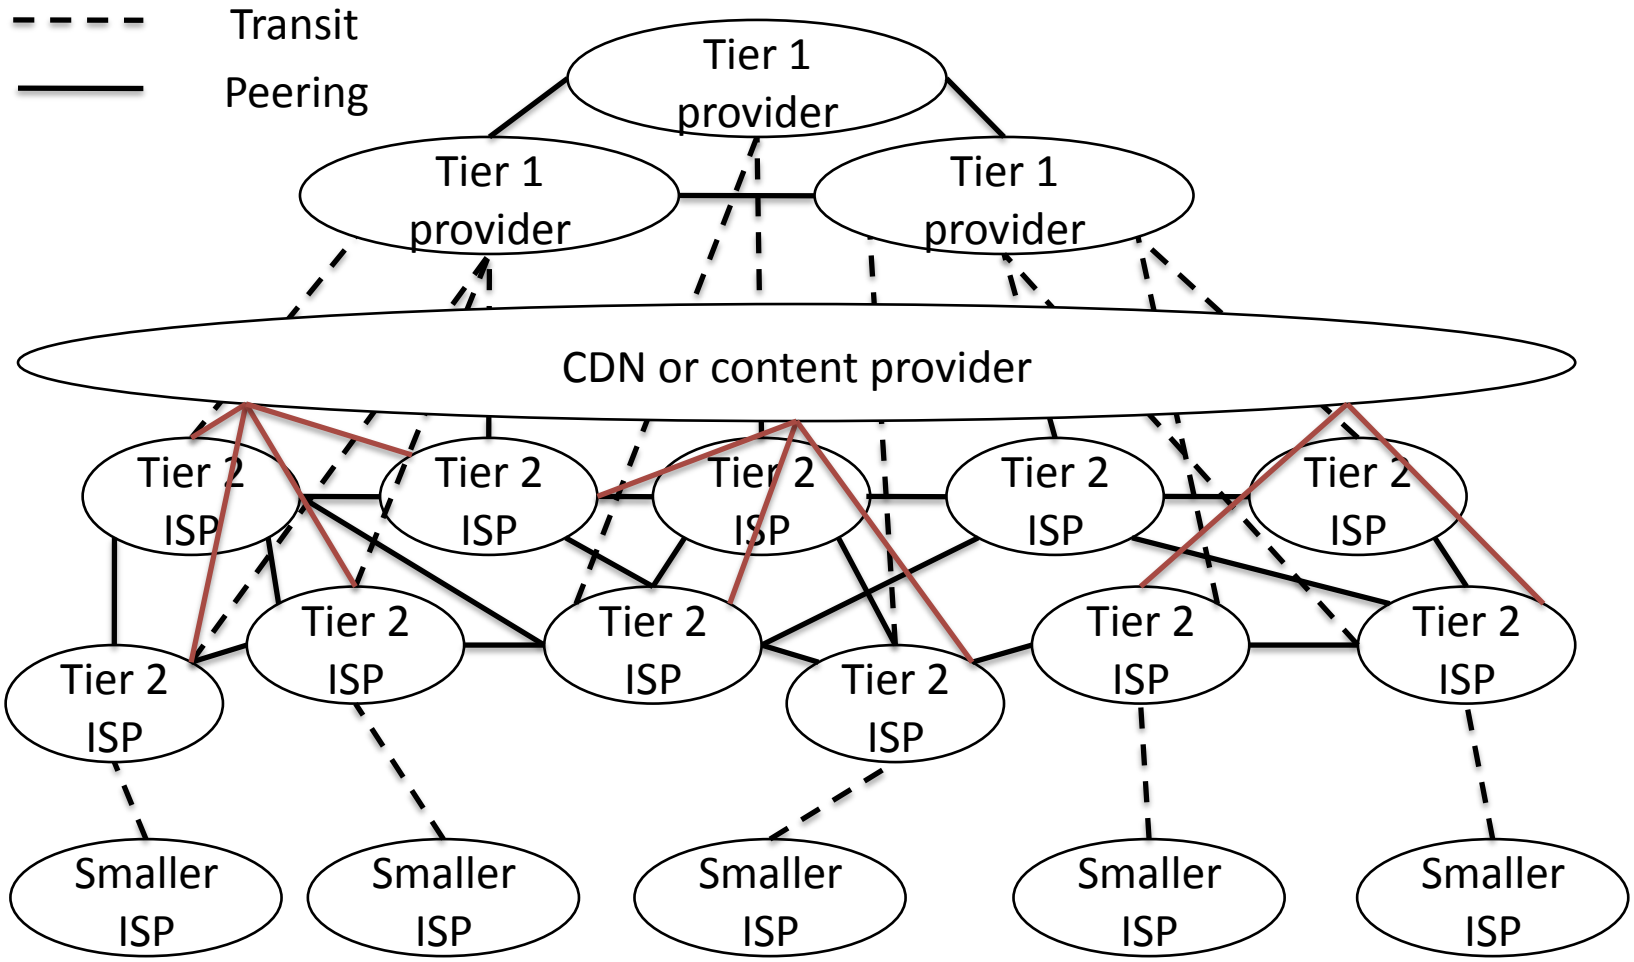

Comcast and Netflix

The current configuration—what Netflix wanted

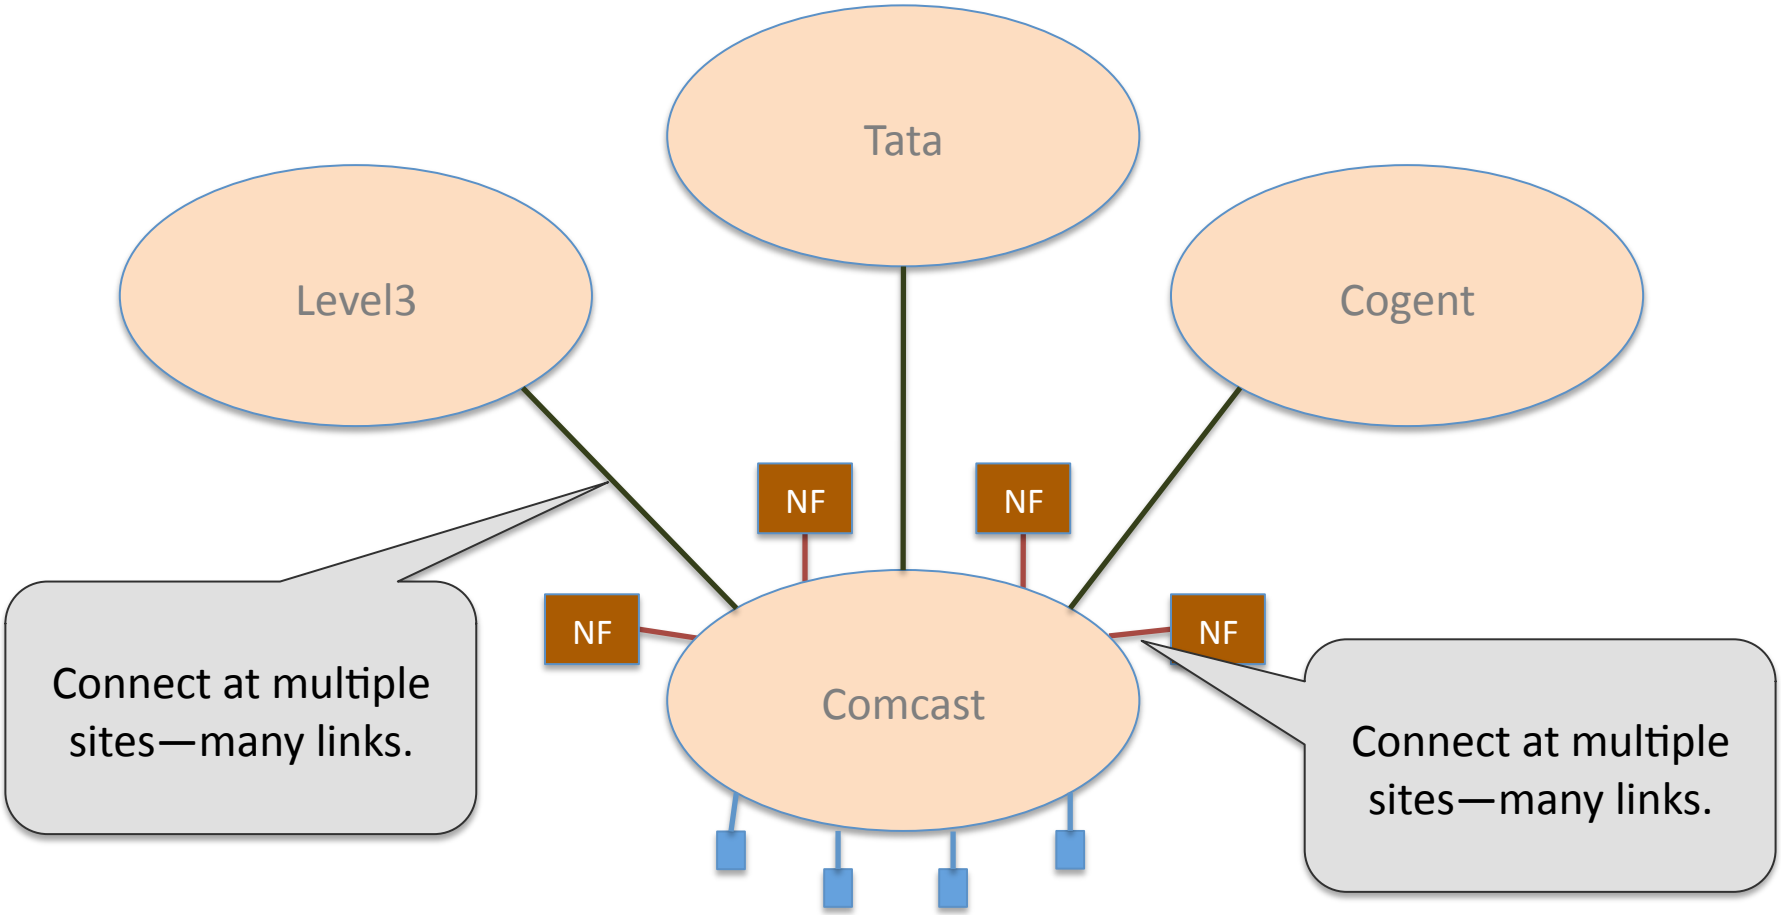

Comcast and Netflix

The prior configuration

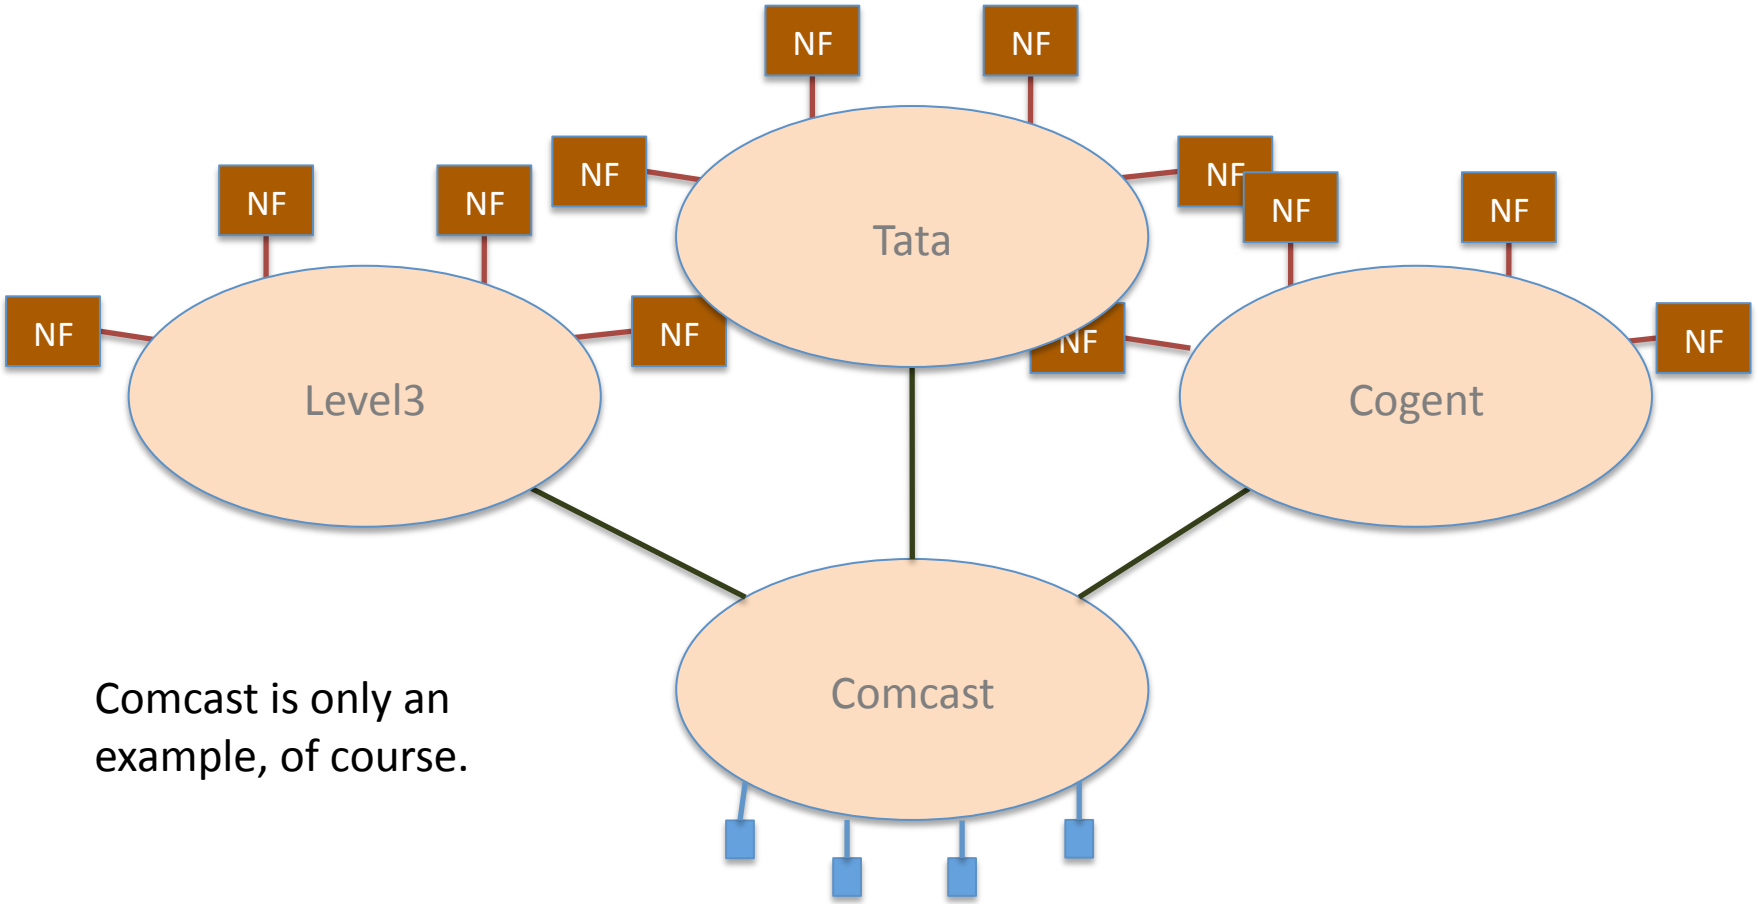

Comcast is only an example, of course.

What was the dispute?

- Did Comcast have an obligation to provide all the ports necessary to receive all that Netflix traffic for free?
- Should Netflix share some of the cost of that interconnection?
- Negotiation was *not* about "premium" access.
  - It was just about who pays to implement *adequate* access.
  - About 3 tb/s of access.

The question of last spring

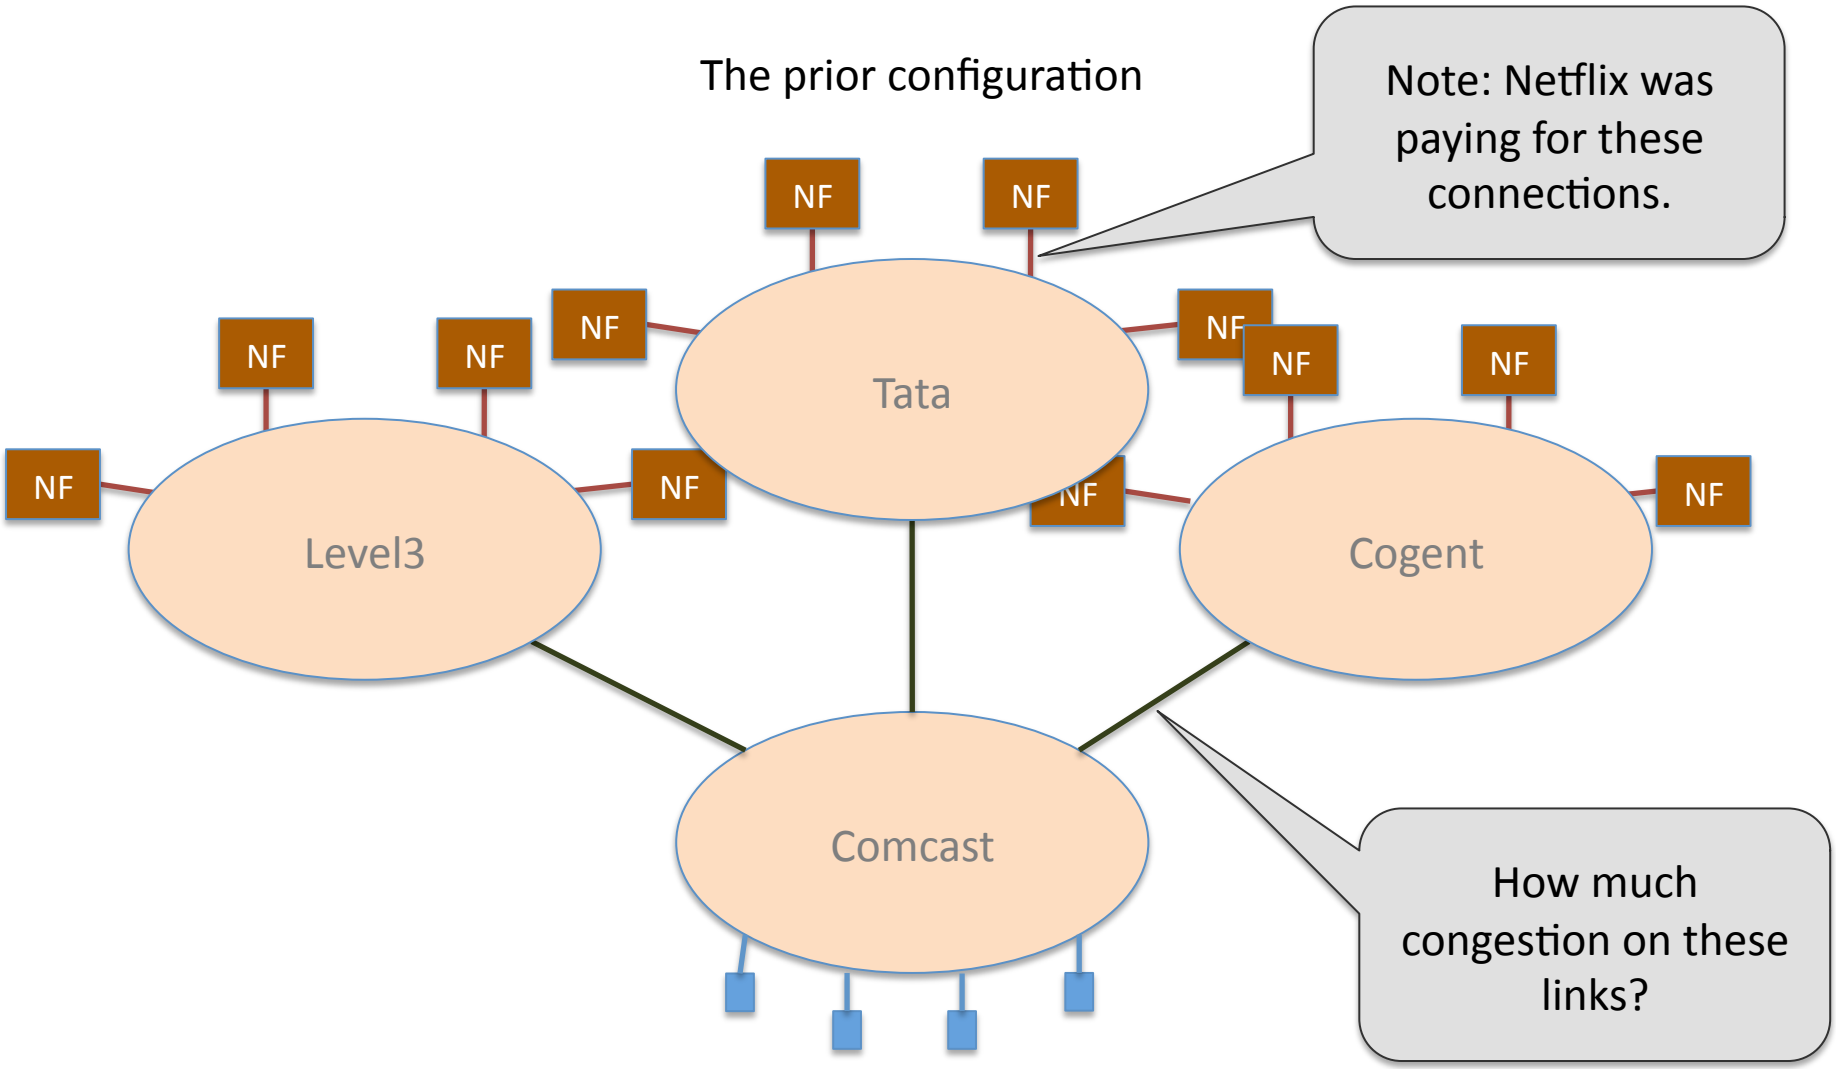

Approach and goal

- Using probe at edge of network, measure level of congestion on specific links in the Internet.
  - No necessary cooperation from ISP.
  - Measure specific link (not path).
- Anticipated outcomes
  - An “atlas” of congestion in the Internet
    - A multi-year effort.
  - Data to inform policy-makers and other non-technical observers.

Method

- Exploit fact that Internet load often shows diurnal pattern.
  - Congestion causes queues to form.
  - Measure Round trip time (RTT) to near and far side of link over time.
  - If link is congested, RTT will go up since measurement packet will sit in the queue.
- Time Series Latency Probes (TSLP)

Probing a link

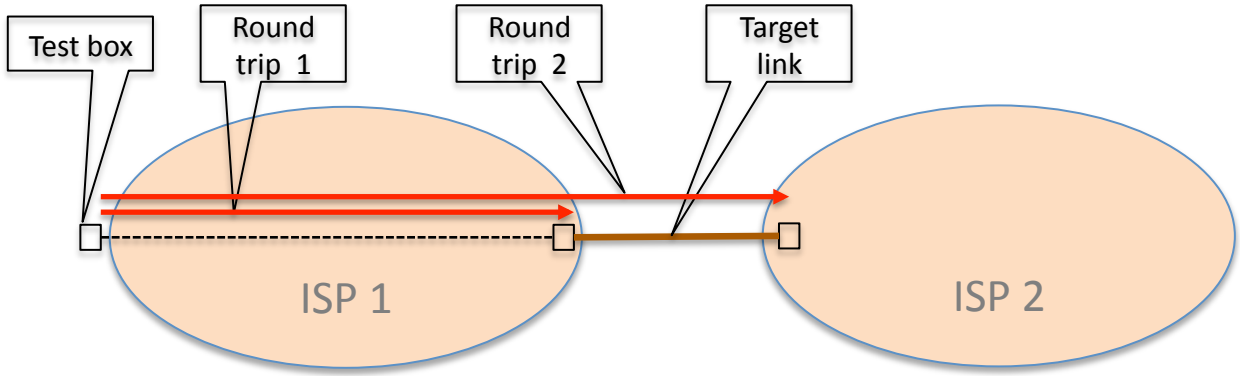

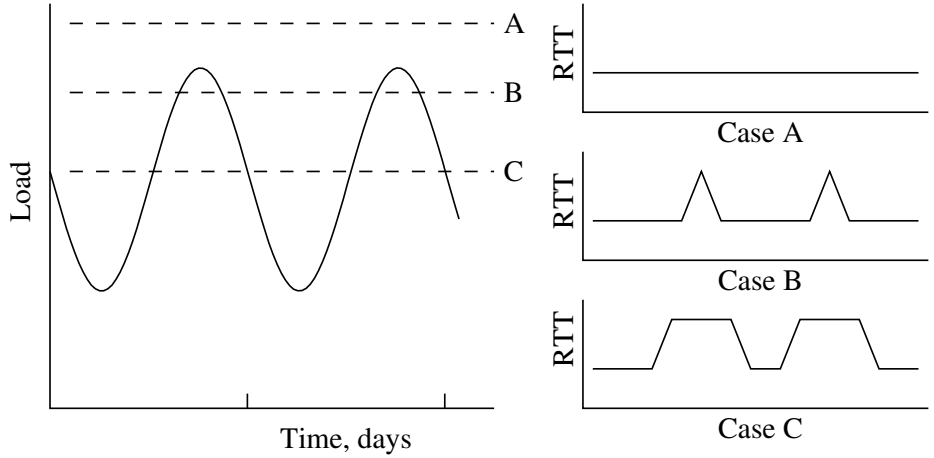

Moderate congestion

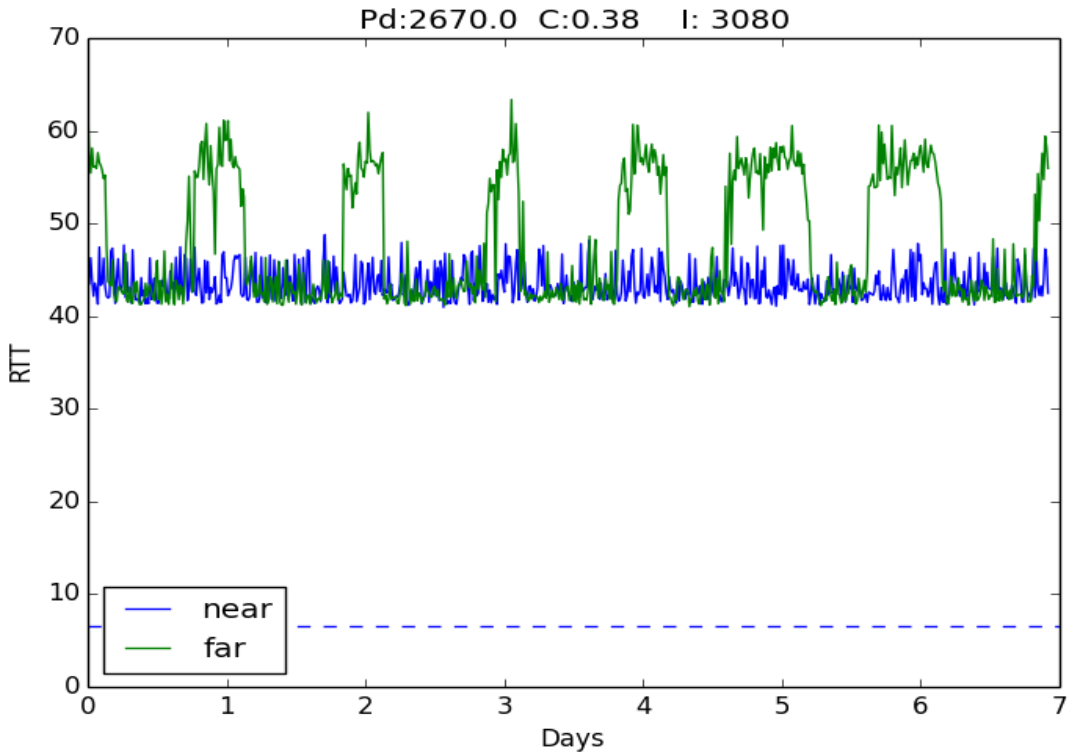

Massive congestion

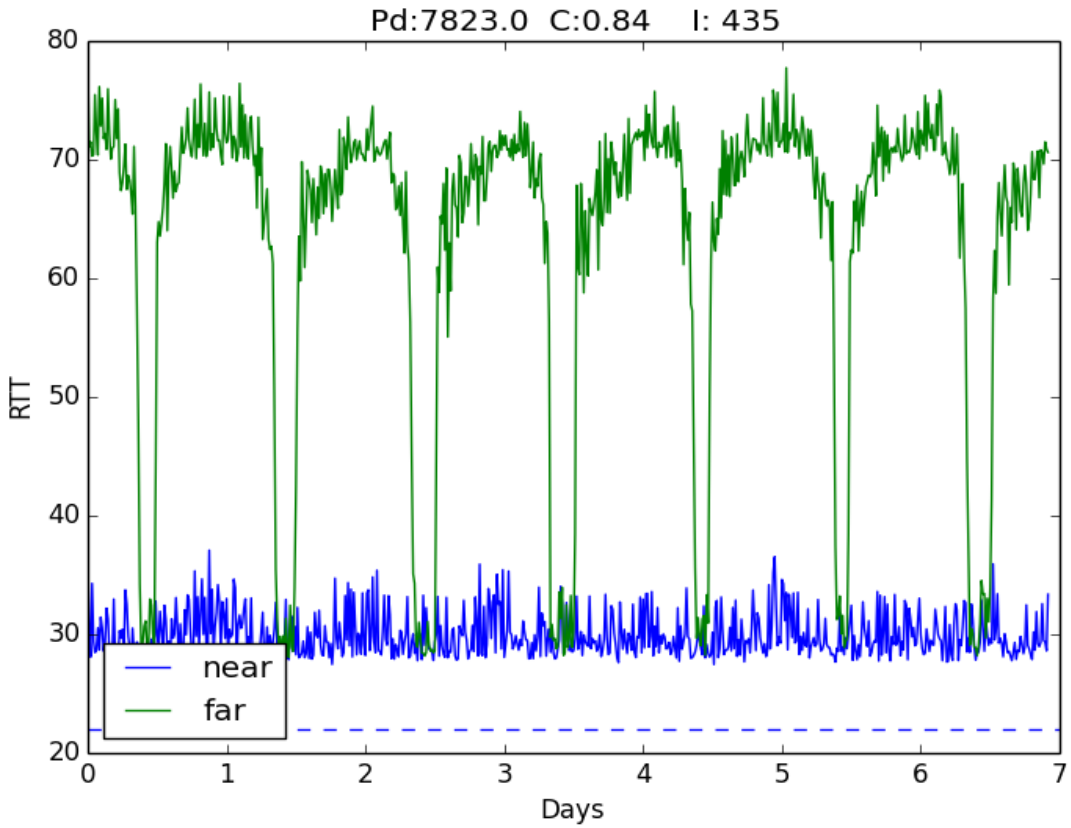

Test infrastructure

- The CAIDA Ark boxes.
  - Widely deployed computers that probe to map Internet topology.
- More than 80 probes around the globe.
  - Currently measuring congestion from about 15 access networks.
  - Physically a Raspberry Pi. (A cute little computer)

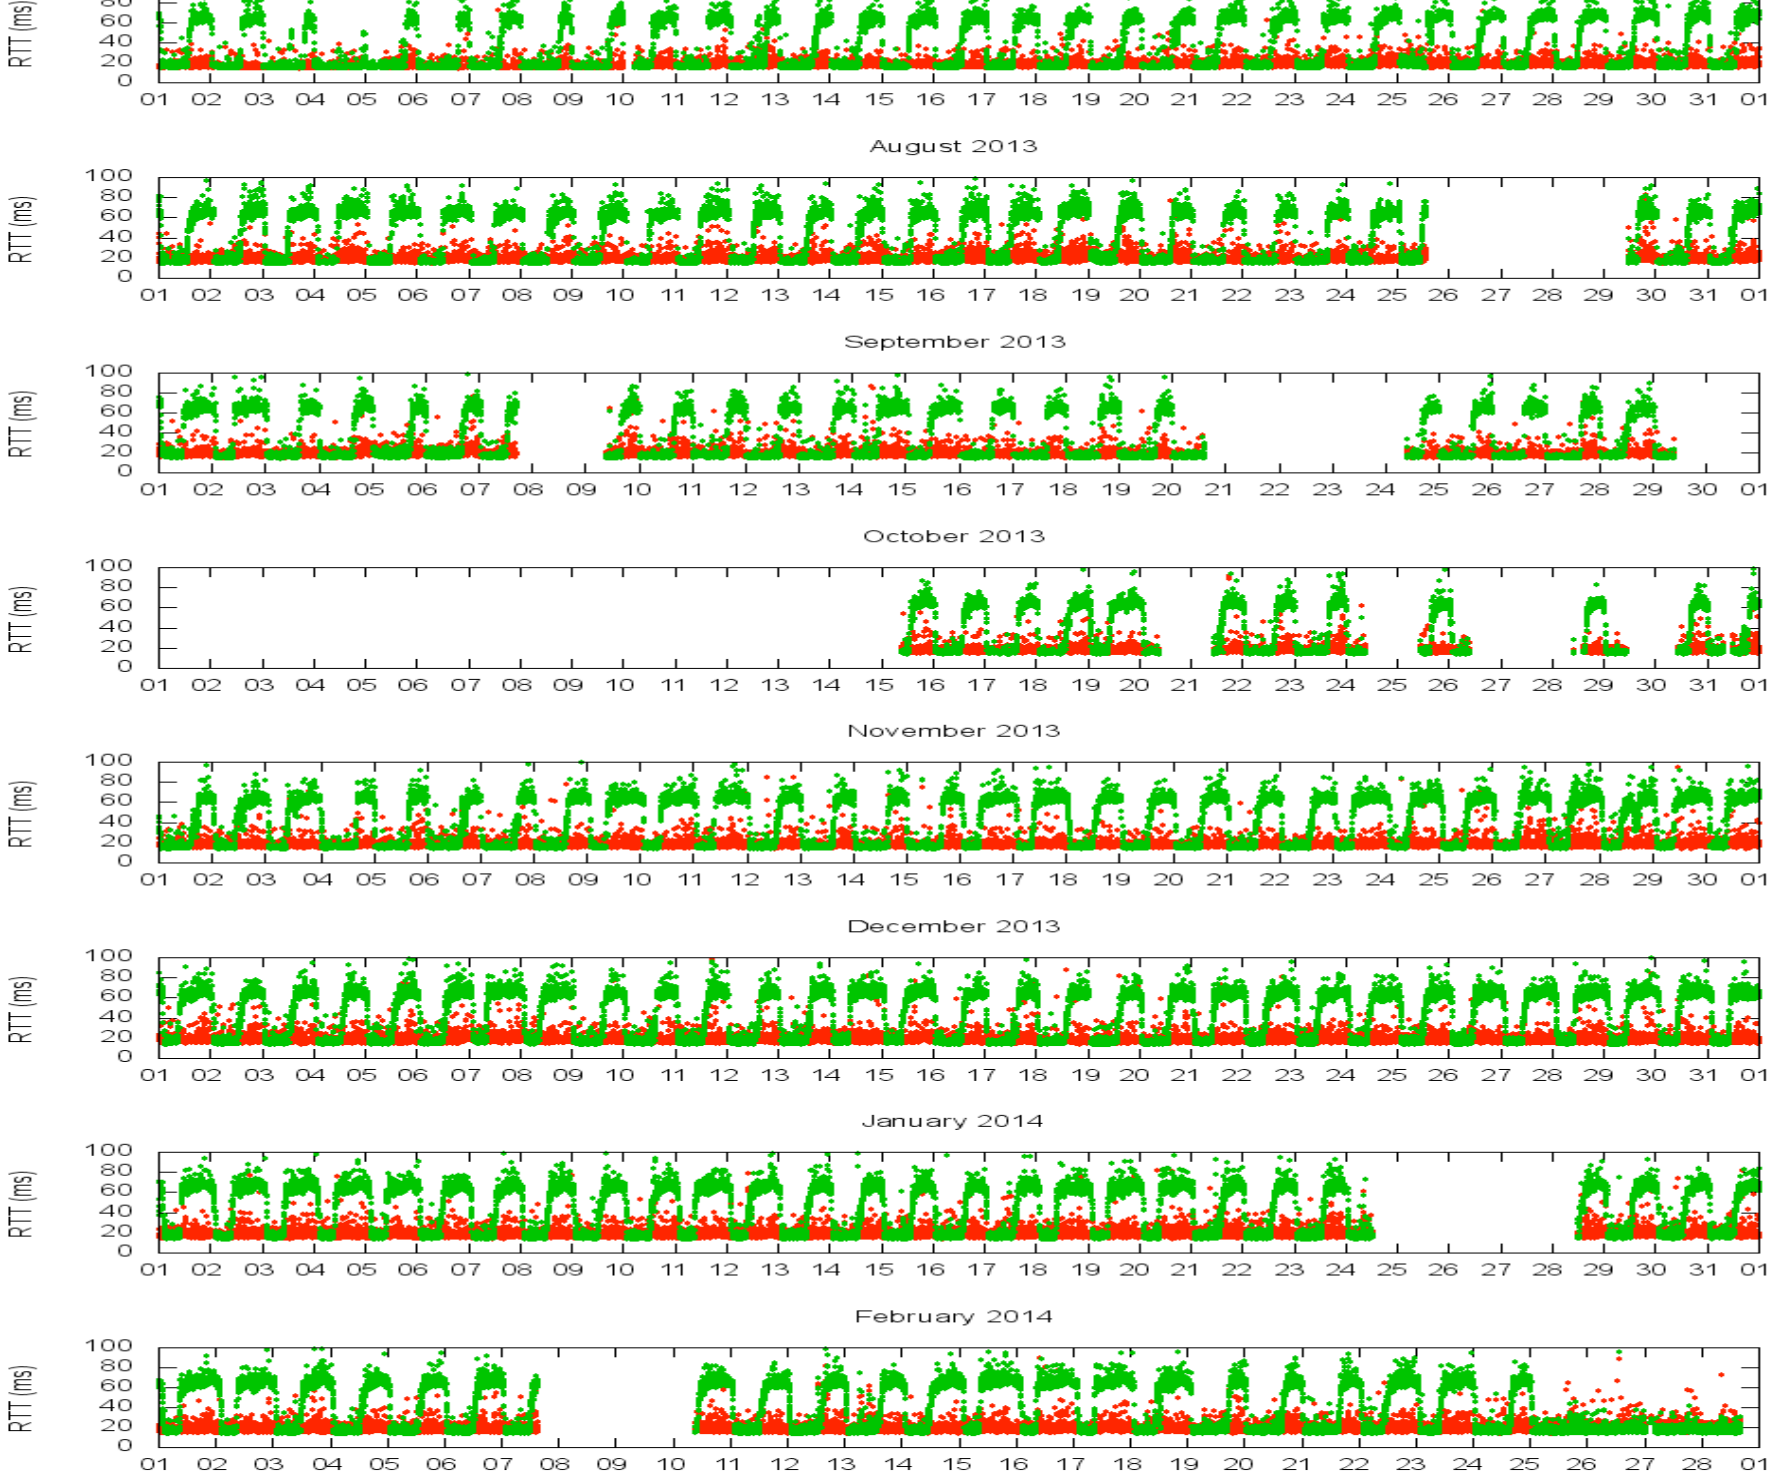
August 2013
September 2013
October 2013
November 2013
December 2013
January 2014
February 2014
RTT (ms)
0
20
40
60
80
100
01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 01

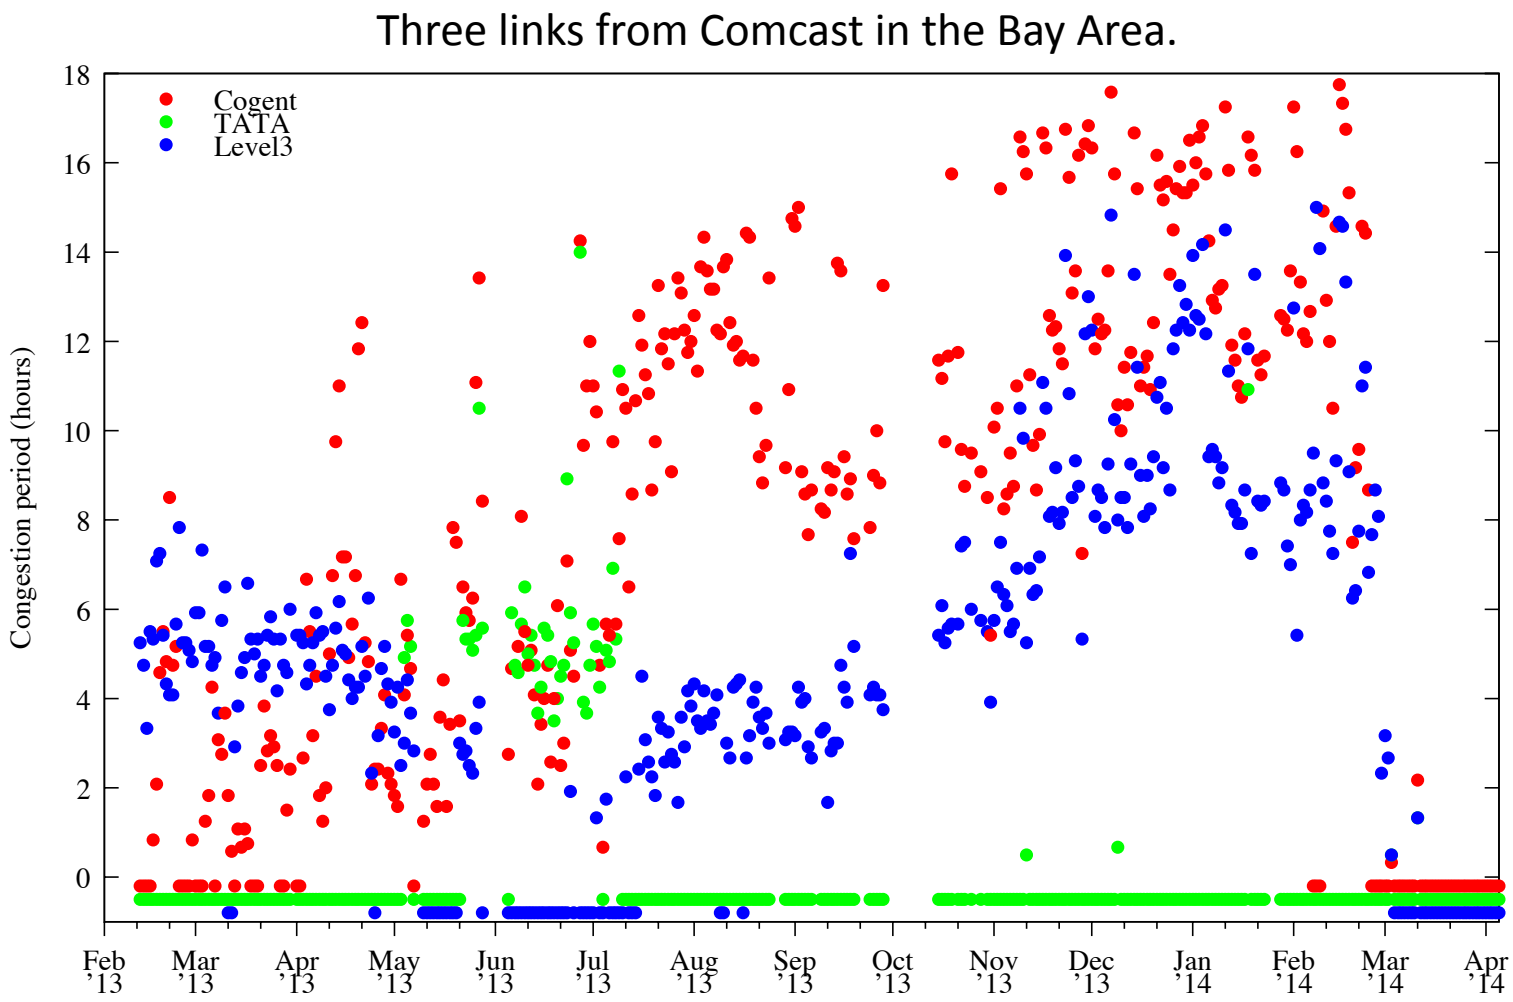

Data analysis and graphing by Matthew Luckie and Amogh Dhamdhere, CAIDA, UCSD

Some observations

- Congestion can come, go, and move quickly.
- Different actors have different options to influence the level of congestion.
  - Sender can pick source. (NB earlier discussion of finding good routes.)
- Adding new capacity is not the only mitigation.
  - Hard to justify investment if congestion may move tomorrow.

The question of last spring

The prior configuration

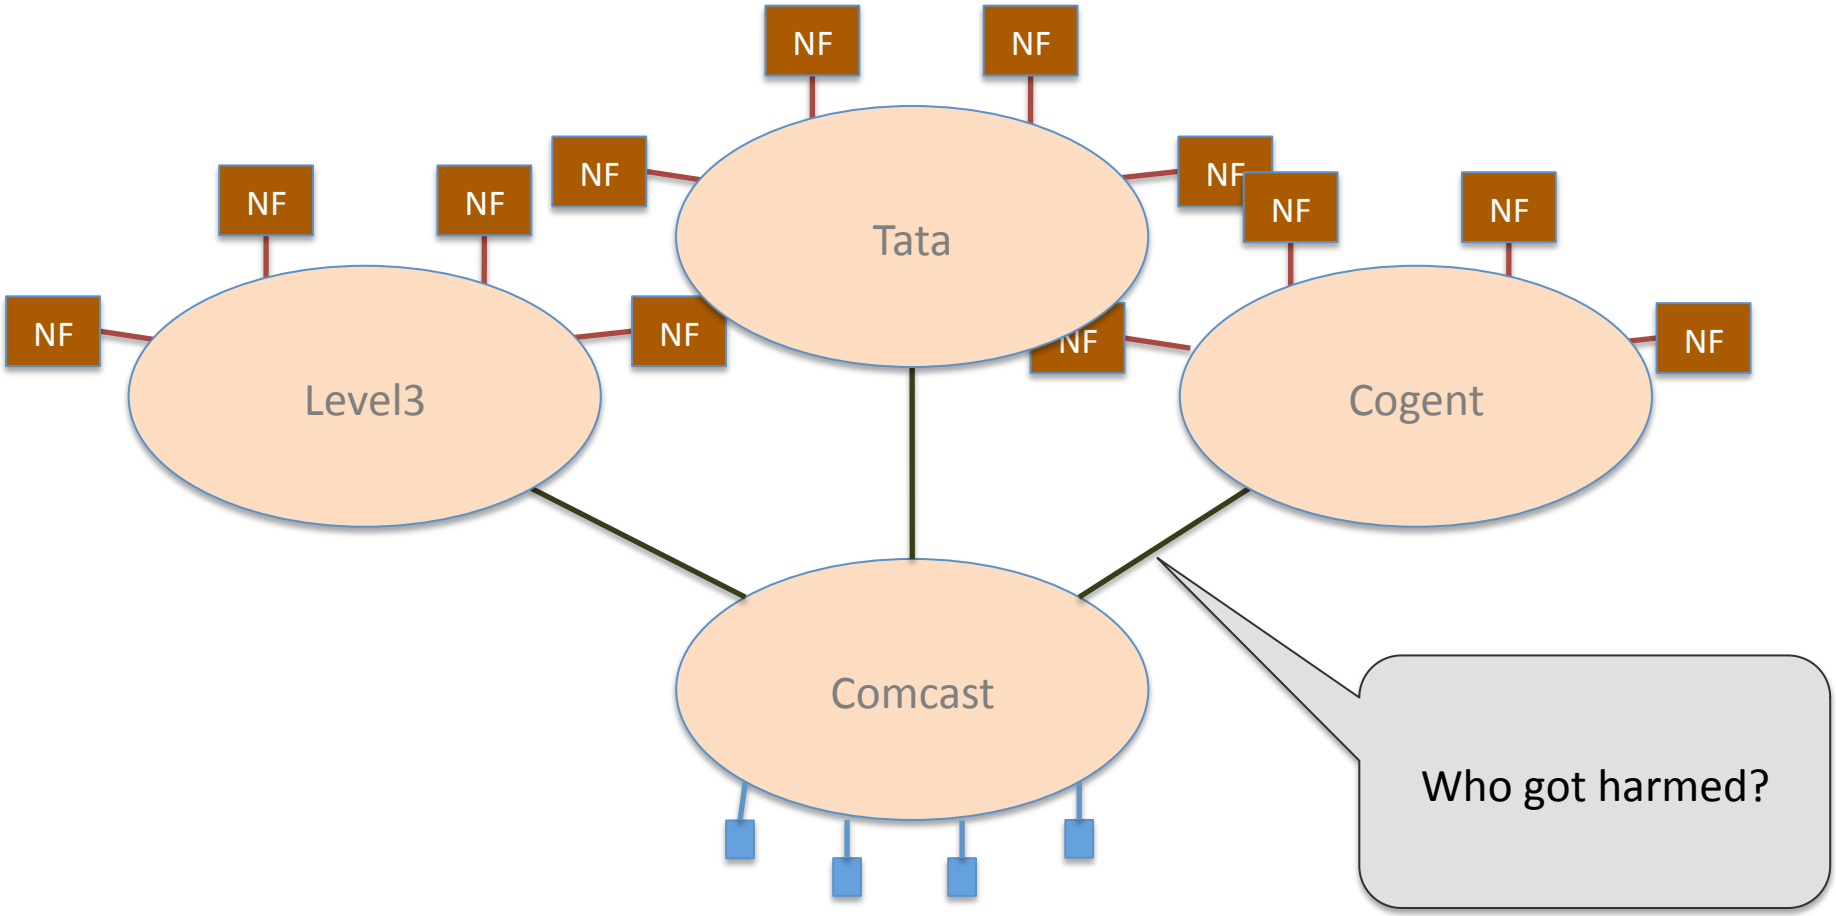

Who got harmed?

- Any traffic flowing across that link.
  - Not just Netflix.
  - Massive collateral damage from dispute.
- What did Level 3 do?
  - Not worth adding capacity or negotiation about who would pay for it.
  - They used fair queuing traffic discrimination to isolate damage to Netflix.
    - OK, FCC—was that good idea or bad?

Neutral harms

- User experience was harmed.
- Customers of Level 3 (and other transit nets) were harmed.
- Was there non-neutrality?
  - Only by Level 3, to protect their other customers.
- Comcast could charge different CDNs different prices.
  - Price discrimination, not traffic discrimination.
  - Should FCC regulate prices for interconnection?
  - Quoth the economist: “All discrimination is price discrimination.”
- ISPs can set usage caps. Hurts high volume video in particular.
  - 300 GB/month is 3 hours a day of HD video.
  - Is this harm acceptable as a way to hinder Netflix?

Alternative platforms

- The global comms infrastructure today is more and more built out of IP technology.
  - But this does *not* mean everything works over the Internet.
  - Lots of private networks built using IP technology.

Xfinity Xbox IPTV

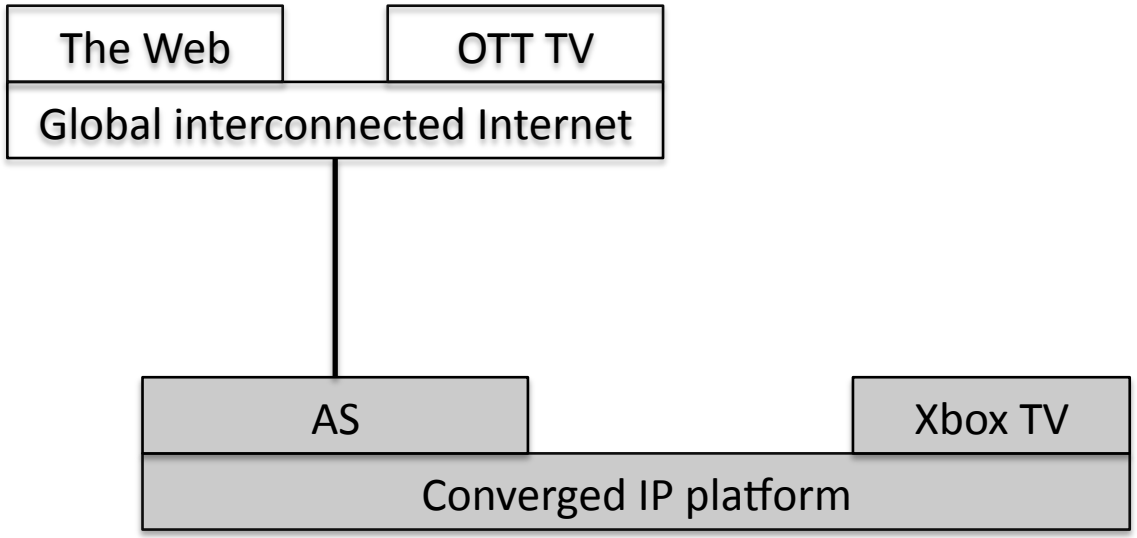

Comcast

Evolution of cable access

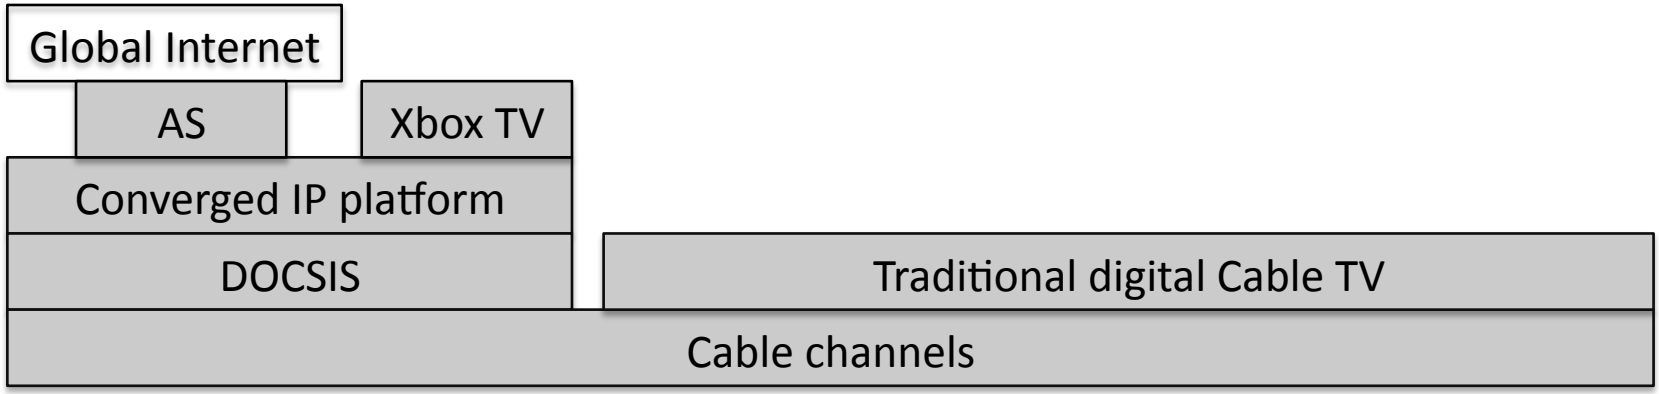

Today

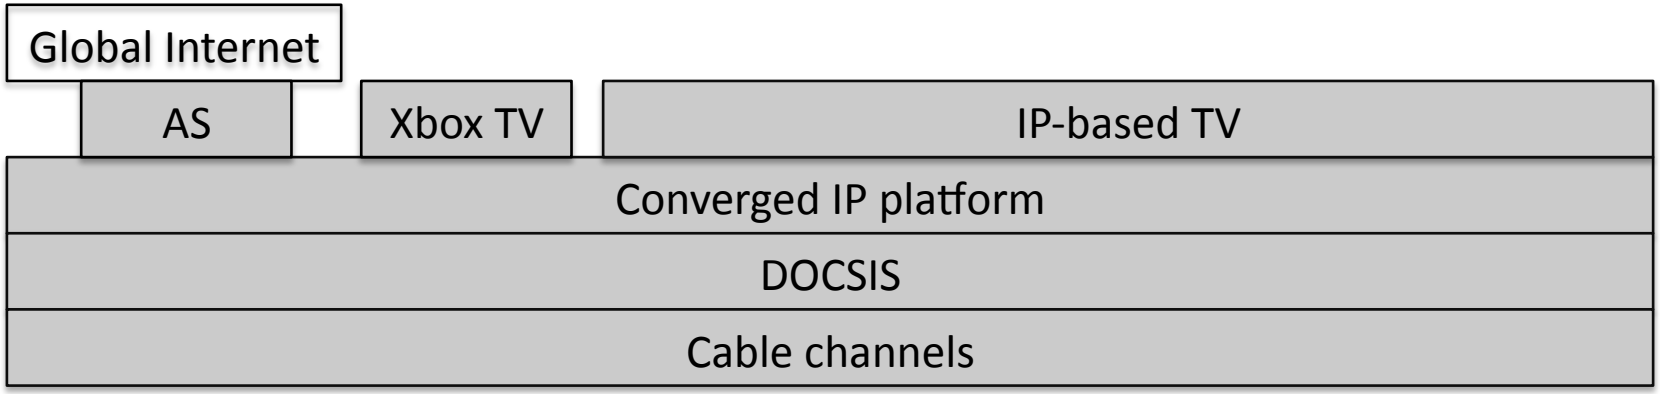

Future

Third-party on both nets

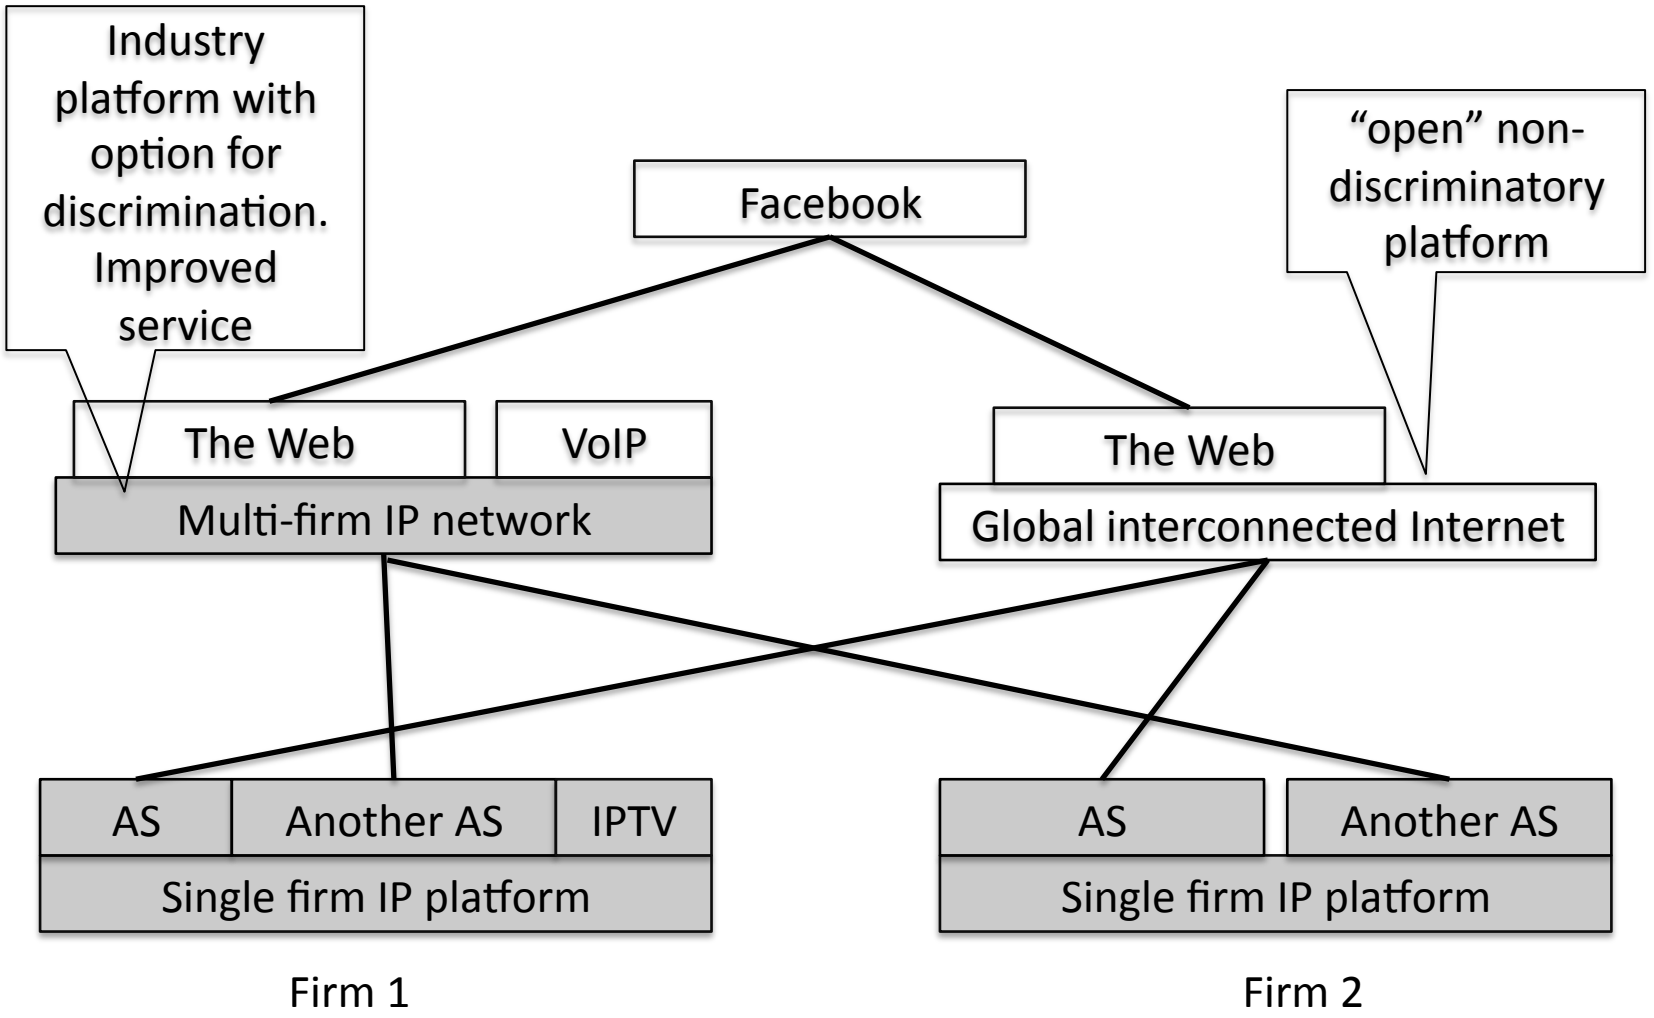

Regulatory issues

- Carrier and third-party services on top of the single-firm and multi-firm IP platform can compete with third-party services running over the Internet.
- Why would they do this?
  - Better quality of user experience (QoE).
    - Protected by fair queuing schemes.
  - Alternative models of interconnection.
- Is the ability of the carrier to exploit this single-firm platform potentially an abuse of market power, or a natural consequence of investing in the facilities?
  - So why is the Internet not a competitive platform?
  - Perverse incentive not to invest in public Internet?
    - Quoth the European ISP: “The ‘dirt road’ future..I like that concept.”

Protecting the Internet

- Option 1: Allow the Internet platform to offer equal quality as the single-firm platform.
  - Implies that it will not be neutral.
  - Ignore what happens on the single-firm platform if the Internet service is "good enough".
    - Ugh---how define that.
    - Answer: measures of QoE.
- Option 2: limit the ability of the owner to exploit the single firm platform.
  - What right does the regulator have to mandate that?
- Option 3: structural separation or neutrality requirements on the single-firm platform.
  - And again, what right…