Interconnection in the Internet:
the policy challenge

David Clark

MIT CSAIL

November, 2011

Background

- Erosion of old models
  - Revenue neutral peering no longer obvious efficient outcome.
  - Who pays may vary even without recourse to market power argument.
- Dispute between Comcast and L3 attracted attention.
  - CDNs raise specific issues.
  - See our TPRC papers, etc., for more background.
- Policy implications
  - Points up problems with NN antidiscrimination rules.
  - No new regulation now, but that may not remain true.
    - Concern is valid
  - Better data and/or disclosure requirements may help.

Traditional peering

- When two similar ISPs (e.g. tier 1 ISPs) propose to peer:
  - Probably both have similar internal cost structures.
  - Incremental cost is negative for both of them
    - Both save money (cost of transit).
  - No new traffic is associated with the agreement.
  - Complexity of negotiation leads to setting transfer payment W to 0.
    - Economist would say this is often an efficient outcome.

If they are not similar?

- A more common situation today.
  - Negotiation between ISPs of different size.
  - Negotiation between ISPs with different internal cost structures.
  - Negotiation between ISPs with different classes of customers.
    - Obvious current case: ISP1 is broadband access network; ISP2 is CDN.

Different sizes?

- Paper by Odlyzko and Tilly: "A refutation of Metcalf's law…"
  - Metcalf's law: all users value each other equally -> value of net goes up as N^2.
  - Consequence: if two networks connect (e.g. peer), the gain in value to each is **independent of relative size**. (e.g. revenue neutral peering is valid).
  - Their view: users do *not* value each other equally.
  - Consequence: small network get more value from peering, and thus should pay.

Simple picture

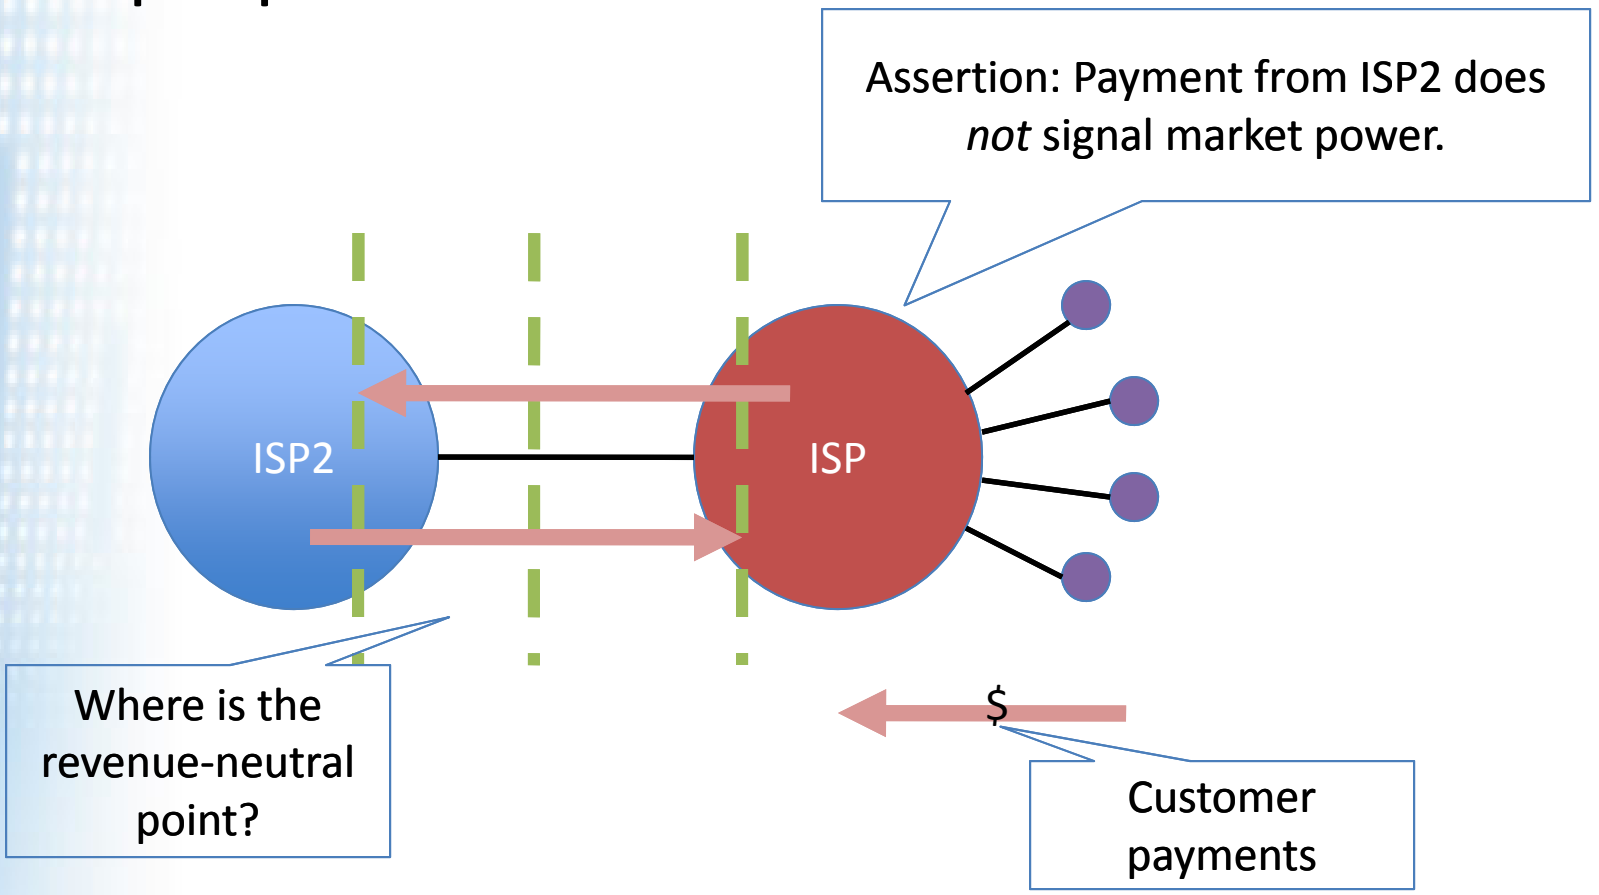

Money flows in from the edges.

- So there must be a point where the flows meet.

Some terminology

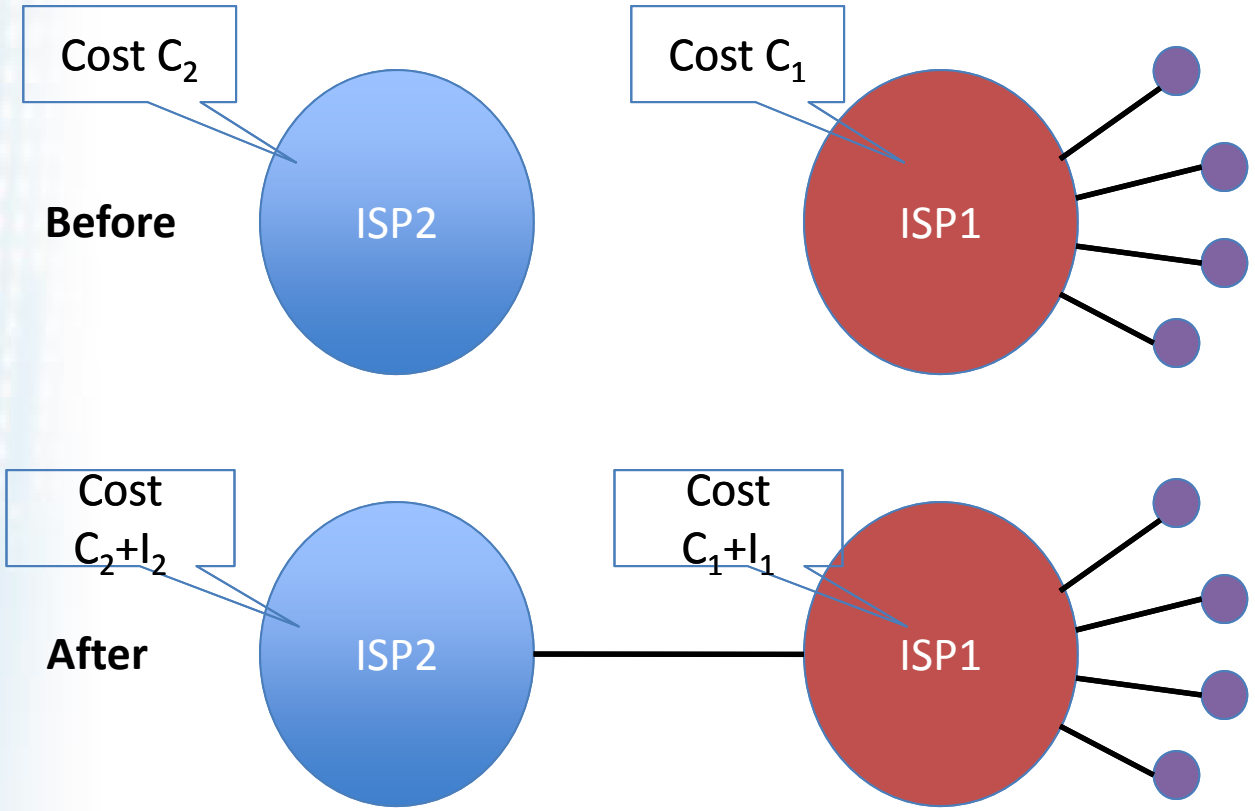

Several cases:

- $I_1<0$; $I_2<0$ (Traditional peering—both sides save money.)
- $I_1>0$; $I_2<0$ (Many CDNs) [POINT OF CONTENTION]
- $I_1<0$; $I_2>0$ (CDN for rural ISP)
- $I_1>0$; $I_2>0$ (No interconnection will occur unless some other party pays)

Cost-based argument

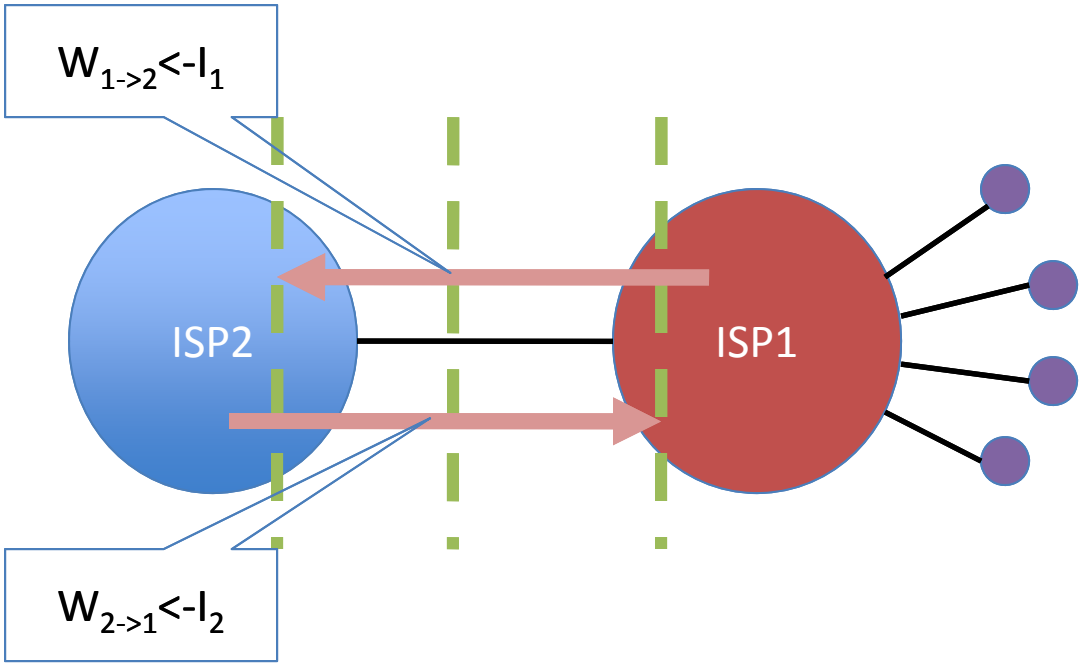

Negotiation to set transfer payment W is bounded by cost savings.
(But actual values for I are hard to determine and not public…)

Refining the questions we should ask

- Assuming that we cannot use actual cost models as an input to any analysis…
- What factors might determine the outcome of bargaining?
- Are there bounds on the outcome that might signal market power and the need for intervention?
  - Could we just use rents (e.g. high profits) as a signal of market power?

Content vs. conduit—the role of “value”.

- With commercial content, there is a payment from consumer to the producer/programmer: the content payment.
  - This is separate from the delivery (conduit) fee.
  - (Content fees take all sorts of forms. Heavily studied by economists.)
- In other circumstances (e.g. telephony) there is no content fee.
  - We just talk to each other—peer production of content.
- If the context of negotiation is access to commercial content, then discussions of value can signal two different things:
  - Attempts to tap into “content payment”.
  - Negotiation over allocation/recovery of the incremental costs.
- Telephone example.
  - “800” numbers.
  - No payment for content but variation in who pays based on perceived value.

History: tapping the content value

- Railroads used to have value-based pricing.
  - Sustained by regulation.
- Trucking undercut high-value pricing with "bytes is bytes" (a ton is a ton) pricing.
  - Contributed to collapse of railroads.
- Lessons:
  - Bad game plan.
  - Presumption: value pricing implies either regulatory intervention or market power.
    - Or that distortion from a cost basis is minimal.

Discussing balance of value is acceptable

- When is it safe(r) to discuss value pricing?
  - One answer: *two-sided markets*
  - Remember the analysis by Odlyzko
- Term in economics:
  - Provider with two customer classes, *which depend on each other.*
- Classic example: singles club.
  - Charge men and women different prices: “ladies’ night”.
- Lots of economic theory on rationale to set prices for classes.
  - The railroad situation was *not* a two-sided market.
    - The different customers were not dependent.
- Important point: relative value is being used as a basis to allocate *cost recovery.*
  - Singles clubs are highly competitive.

Consider the ISP

- Is an ISP a provider in a two-sided market?
  - Are there mutually dependent customer classes?
    - CDN and customer?
- Proposition:
  - Value-driven negotiation over cost recovery is just fine.
  - Extraction of payments associated with the value of content is suspect.
  - So how can we tell the difference?

Specific case: CDNs

- Are CDNs a “class” of customer?
  - Lots of different content payment models.
    - Netflix: consumer pays Netflix pays CDN (pays ISP?).
    - ESPN3: consumer pays ISP pays ESPN (pays MLBAM) pays CDN (pays ISP?).
    - Ad-based: Advertiser pays programmer pays CDN (pays ISP?).
      - Only significant case with additional infusion of money.
  - But all the content seems to be “commercial”.
- To my knowledge, CDNs do not/cannot demand a “value-based component” in their pricing.
  - Highly competitive and commodity.
  - Would seem to suggest that all CDN traffic could validly be put into one “value class” in two-sided analysis.

CDNs

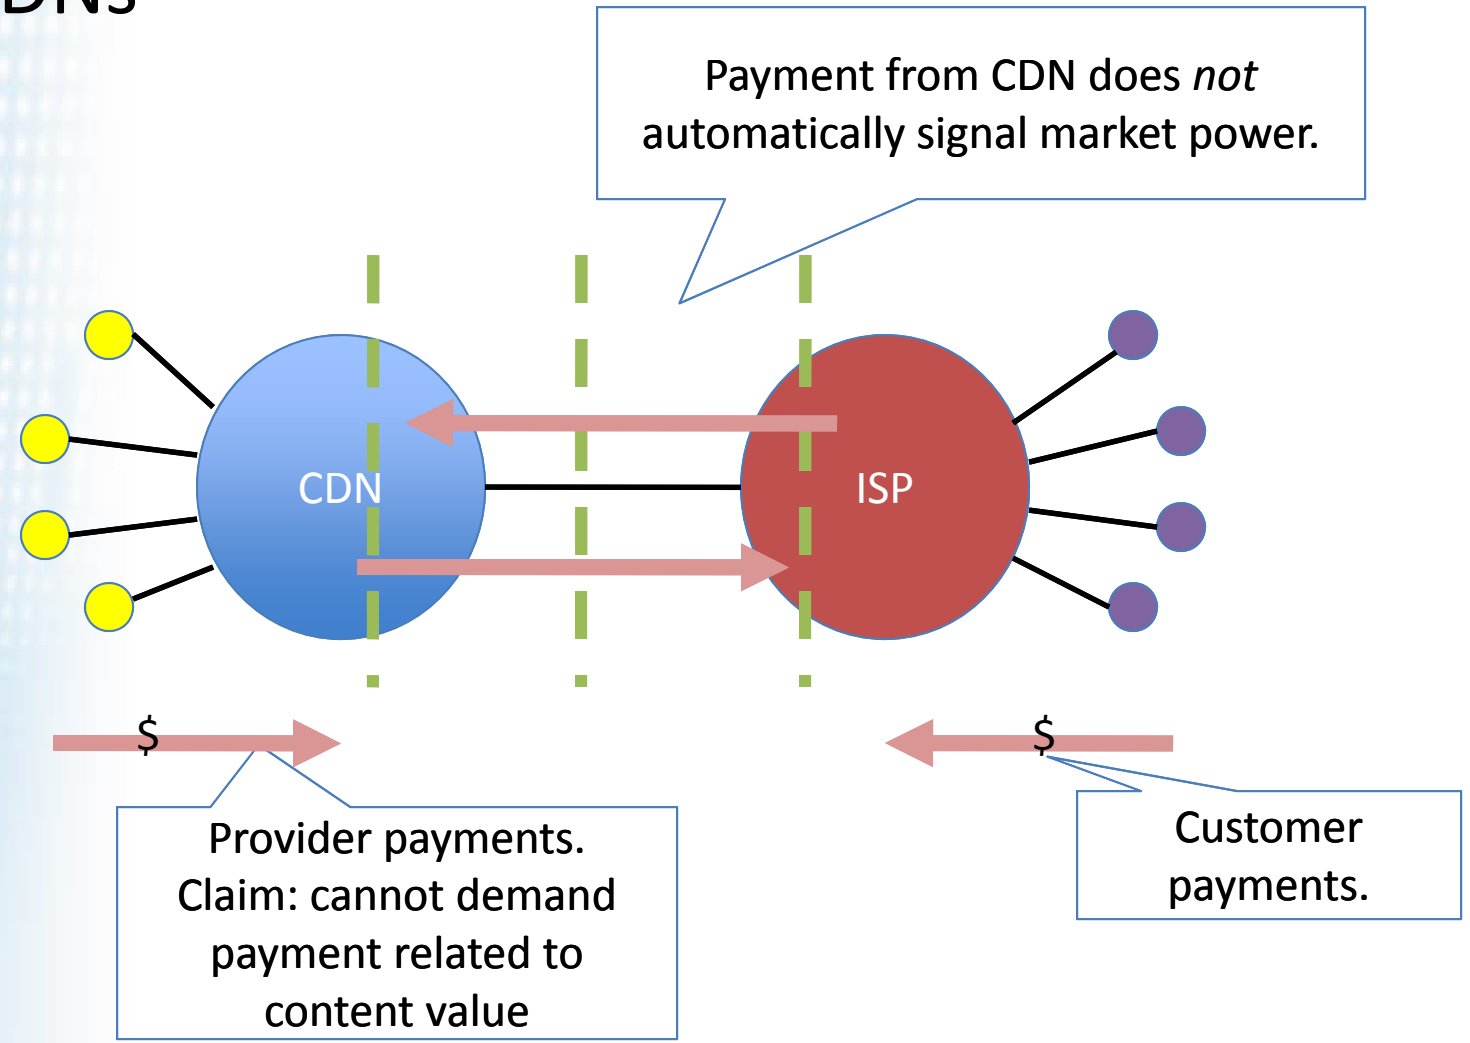

Hypothesis: $I_{CDN} < 0$ (costs go down); $I_{ISP} > 0$ (costs go up).

Bounding the bargaining—one example

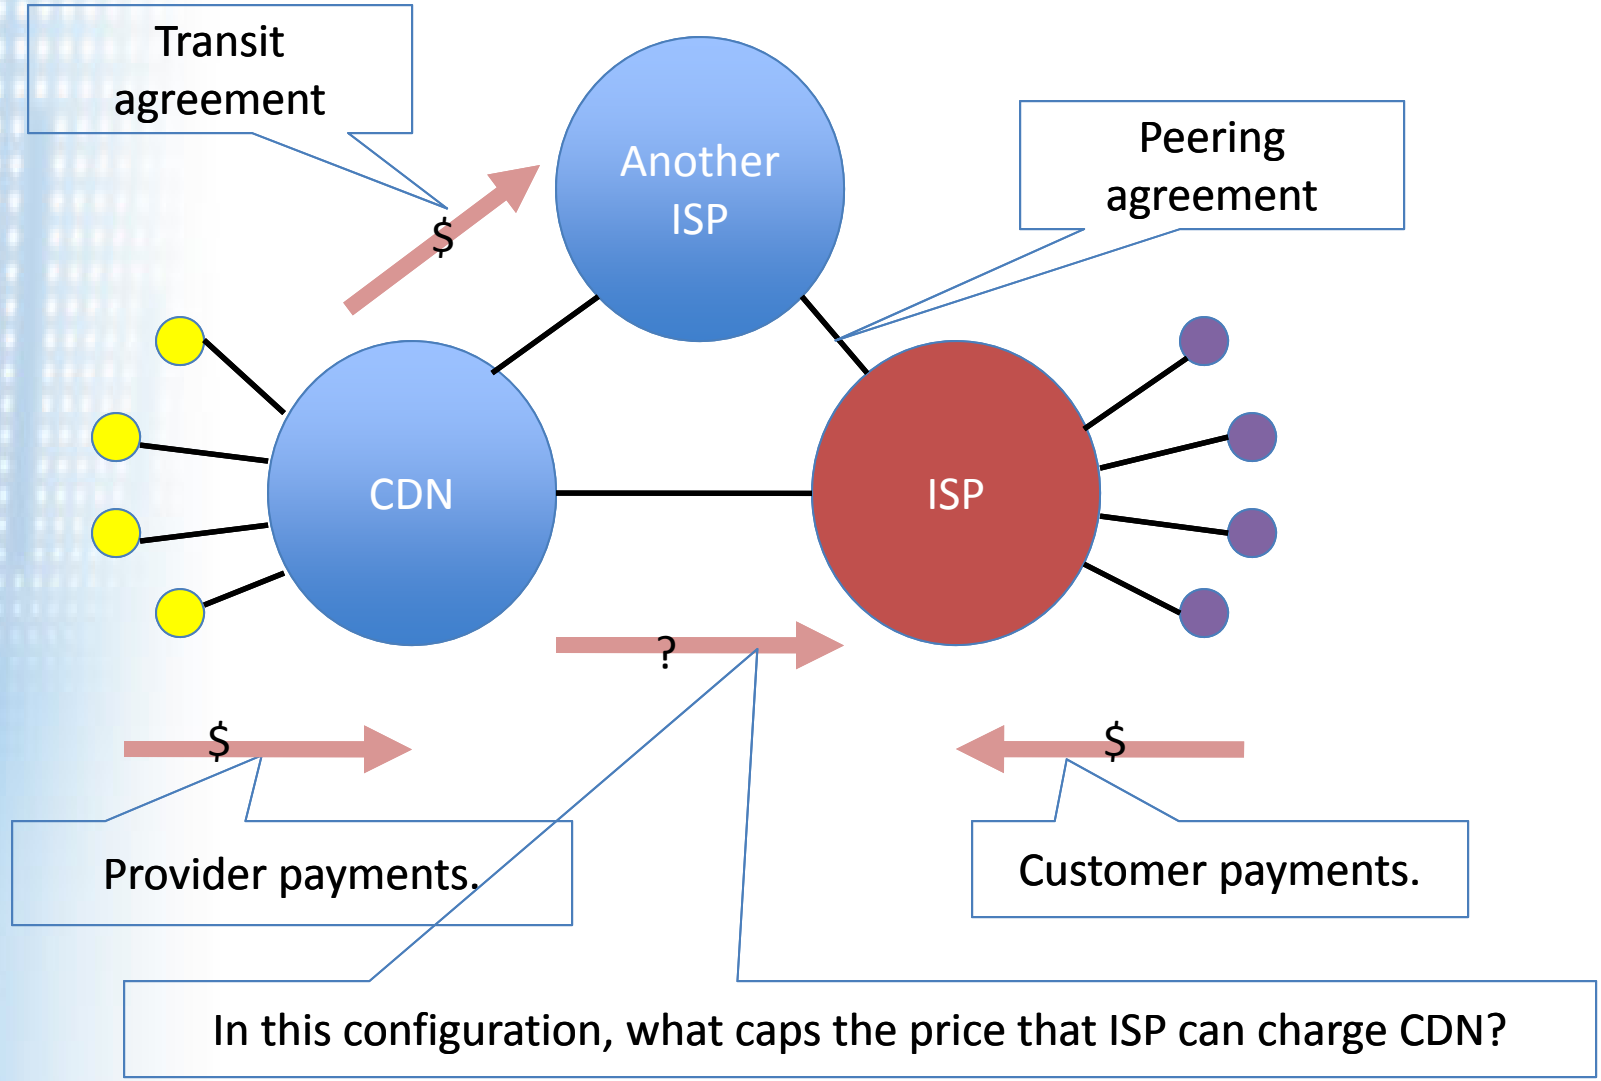

Topological limits

- In the previous picture, why would CDN ever agree to pay ISP more than the cost of transit to CDN if it reaches ISP by its peering partner?
  - Better performance
    - Yes, but probably not sufficient to justify a big distortion.
  - ISP blocks CDN traffic coming from “Another ISP” peer.
    - That would be pretty blatant discrimination.
  - ISP de-peers “Another ISP”, or demands paid peering.
    - Back where we started.
    - But with a tougher job of value discrimination.

Finding the bright line

- Payment from CDN to ISP could be:
  - Extraction of a rent based on the value of the content.
  - Resolution of the “two-sided market” value-based cost allocation.
- Where is the line?
  - A candidate rule: the line is not at “zero payment”, but at a point that is a function of customary transit costs.
    - Perhaps some discount (due to routing restrictions).
    - Perhaps some premium (for enhanced service)?
  - Even though specific agreements and incremental costs are NDA stuff, could a “customary function” emerge?
    - Customary function might be easier to explicate to regulator than cost models for incremental costs I_1 and I_2.
- But nature of negotiation still undefined.
  - Who pays whom how much?

Balance of flows

- One traditional basis for agreeing to revenue-neutral peering has been balance of flows.
  - A long tradition.
  - But actually no obvious basis in cost.
    - Circuits cost the same no matter which way the traffic flows. (Asymmetry leads to under-utilized capacity.)
  - Seems to be based on assumption that balance of flows signals “similarity of character”.
  - Seems to be based on rough rule that value follows the packets.
    - But this could be totally backwards.
- When ISPs are clearly not similar, no clear reason why balance of flows is a good rule of thumb.
  - Level 3 has challenged this idea.

New idea: patterns of usage

- Netflix as example:
- Consider two extremes.
  - All users consume Netflix content equally.
  - Only one user consumes Netflix content.
- If the ISP-CDN connection is revenue neutral:
  - In the former case, users have equal total usage, so the Netflix-specific usage is balanced out.
  - In the latter case, the usage-related costs of this user are being spread across all the users.
  - But not all users have equal usage, as we know.

Distribution of monthly usage.

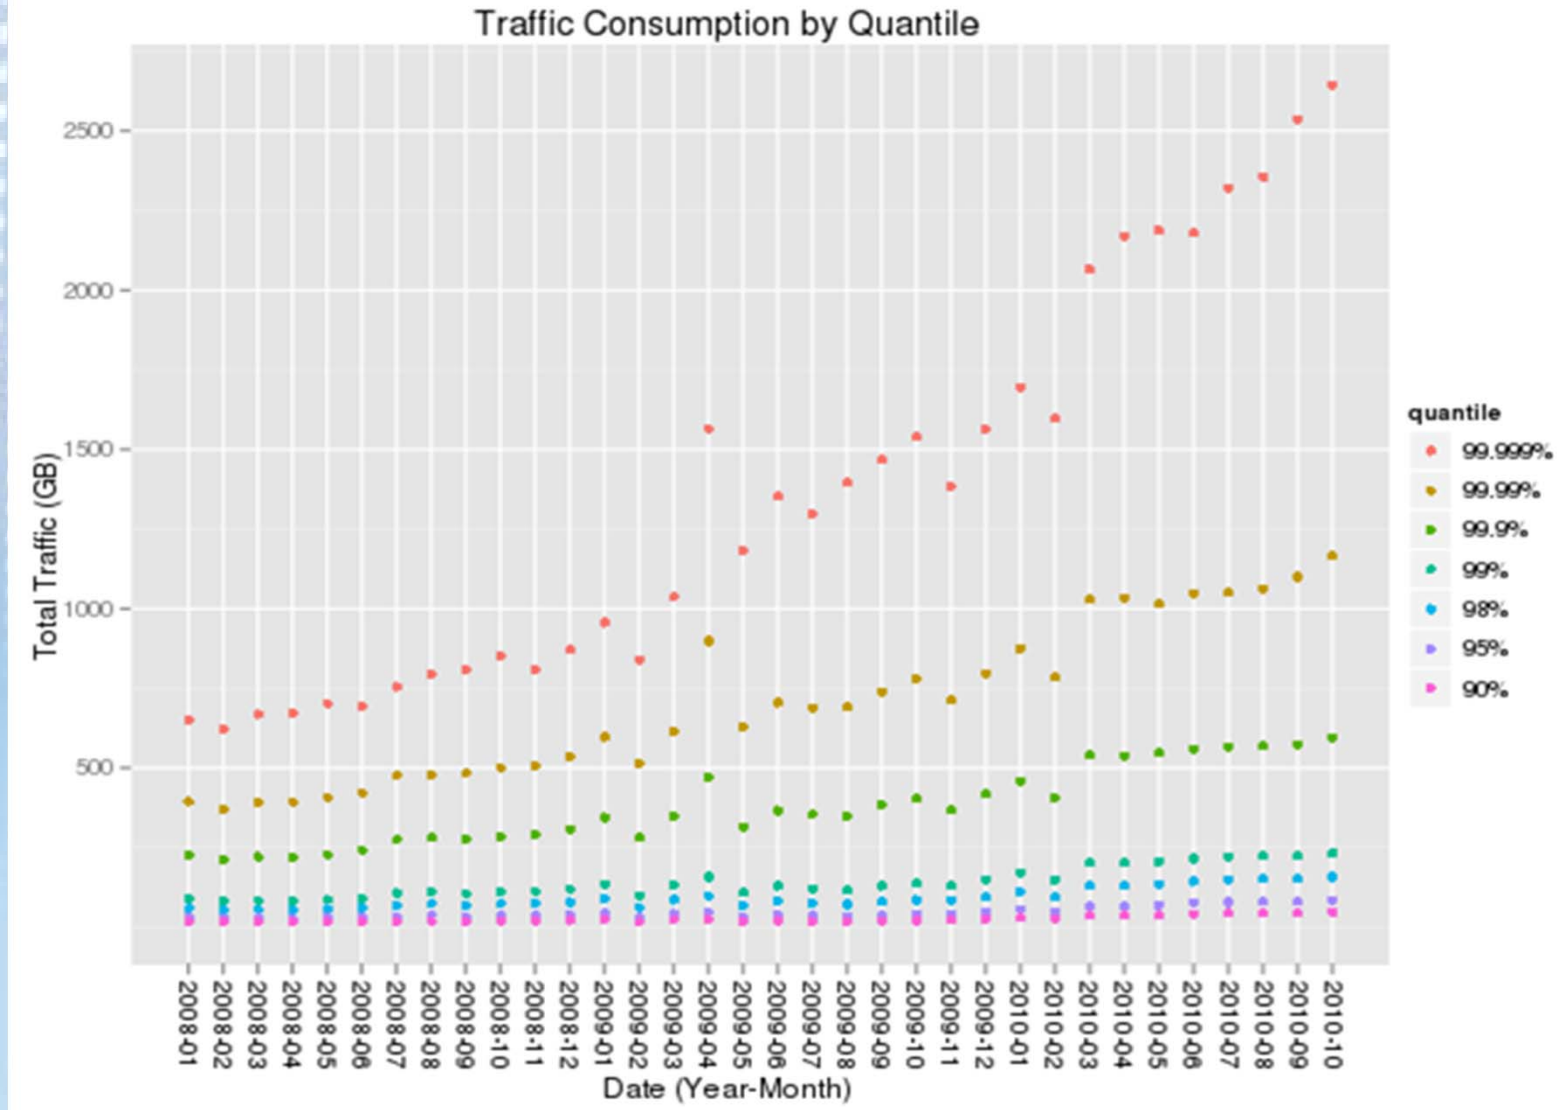

Cross-user subsidy

- If all users pay the same for unequal usage, this is cross-user subsidy.
  - This is *not* a two-sided situation—the users are not mutually dependent.
- Why is this subsidy sustainable?
  - The amount is small and not worth thinking about. (Let's see if this is true.)
  - All participants in the market prefer it.
  - ISPs are not subject to competition.
  - Regulators demand it.

Some real numbers?

- (Based on some data we have from ISPs about user behavior. Report with data out “any day now”.)
- With flat rate pricing, everyone pays the price for the average usage.
  - How to measure usage?
  - What price to assign?

Measuring usage

- Current approach is to measure total monthly usage.
  - An approximation—cost primarily relates to busy-period load.
  - Preliminary data suggests that distribution of monthly usage and busy-period usage are proportional. So a good approximation.
    - But an interesting point. Move from P2P to streaming may have moved the heavy users to the peak period.

Some numbers from our data

- For latest month, mean usage:16.1GB/m.
- Median usage is 4.25GB/m.
- So for median user, usage-related cost is about 3.8x with flat-rate pricing
- The “really big users” are getting a “really big subsidy”, but there are not many of them.
  - The neutral point is the 73%ile user.

Now estimate usage costs.

Cost/GB ($)	Median cost ($)	Flat-rate (mean) cost ($)	Subsidy ($)
.05	.21	.81	.60
.10	.42	1.62	1.20
.20	.84	3.24	2.40
.30	1.26	4.86	3.60
.40	1.68	6.48	4.80

How to recover costs of heavy users

- Ignore the issue and charge flat rate.
- Usage caps and tiers.
  - The users pay directly.
  - (Nobody seems to like actual per-byte pricing.)
- The “other side” pays—the provider or the CDN.
  - The bargaining over cost recovery I have been referring to.

An enhanced negotiation

- When two ISPs negotiate:
  - Specifically an access ISP and a CDN
- First try the "balance of flows" rule and see if both sides are satisfied.
- If not, consider the degree of uniformity of destinations from the CDN across the ISP
  - If uniform, ISP *might* agree that it is satisfactory to have the users pay cost directly.
    - If all the money comes from the users, what difference does it make?
  - If highly non-uniform, ISP *might* ask payment so that these customers are not being subsidized by all the other users.
    - Alternative is to go to usage tiers and bill users directly.

Using caps to get the CDN to pay

- Example: Australia
  - Low monthly caps.
  - Providers and CDNs pay for “premium service” so that their bytes do not count against the cap.
    - Called “un-metering”.

Usage caps change the negotiation

- What price should an access ISP charge a CDN to “un-meter” content?
  - It is not capped by customary transit.
  - CDN does not have the flexibility to change routes.
  - So seems to increase greatly the power of the access ISPs.
  - Which implies that usage caps and price tiers may attract regulatory “attention”.

Review: assumptions and ideas

- CDN market is competitive
  - CDNs do not partake of the content payment.
    - (Except perhaps as collection agent.)
- Interconnection can result in significant incremental costs, positive or negative.
  - May be many reasons.
- Transit costs may provide a customary basis to cap reasonable outcomes of negotiation.

Some summary thoughts about policy

- Rational discussion of interconnection cannot be separated from discussion of the carriage of commercial content.
- In the context of commercial content, consider:
  - *To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet*, consumers are entitled to access the lawful Internet content of their choice. (Early FCC principle.)
- Retail pricing policy (e.g. whether to have low usage caps) will end up being tangled with bulk interconnection negotiation.