The Structure and Evolution of the AS-level Internet

Amogh Dhamdhere (CAIDA/UCSD)

Pretty pictures of the Internet

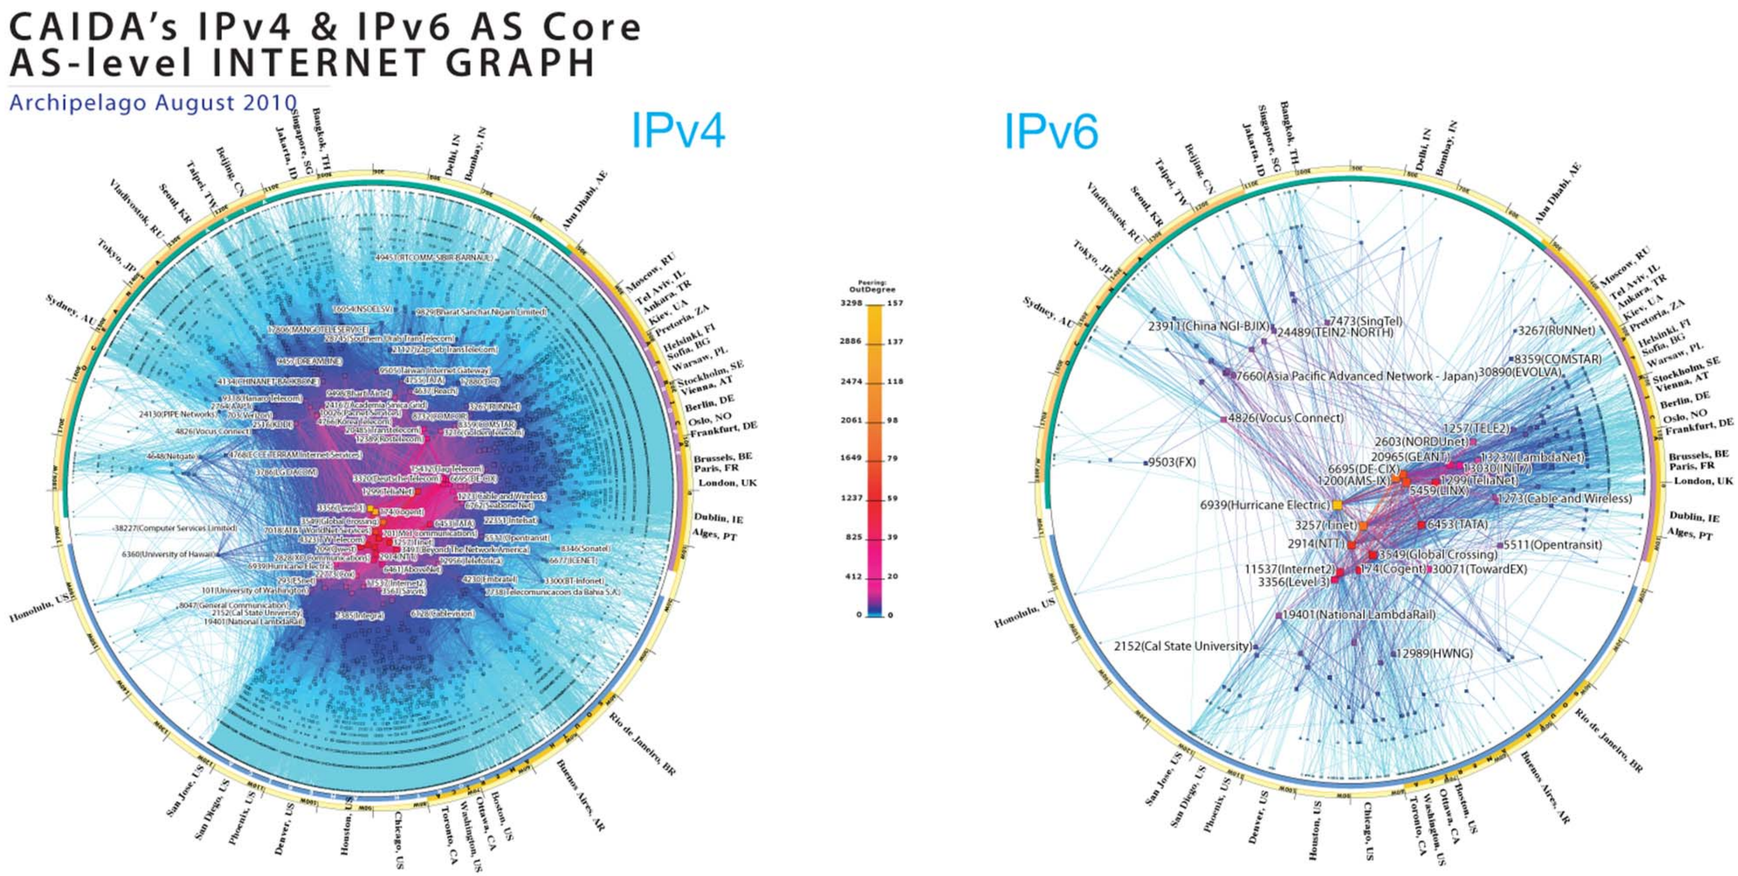

Different Aspects of Internet Topology

- Router-level: How do individual routers connect to each other?
- PoP-level: How are routers organized into "points of presence"?
- AS-level: How to different networks connect to each other?

Different Aspects of Internet Topology

- Router-level: How do individual routers connect to each other?
- PoP-level: How are routers organized into "points of presence"?
- AS-level: How to different networks connect to each other?

AS-level Internet Topology

- The Internet consists of ~40,000 networks
- Each independently operated and managed
  - "Autonomous Systems" (ASes)
- Distributed, decentralized interactions between ASes
- Different AS types based on business function: transit, content, access, enterprise
- Complex structure inside each AS – routers, PoPs, backbone links

The AS-level Internet

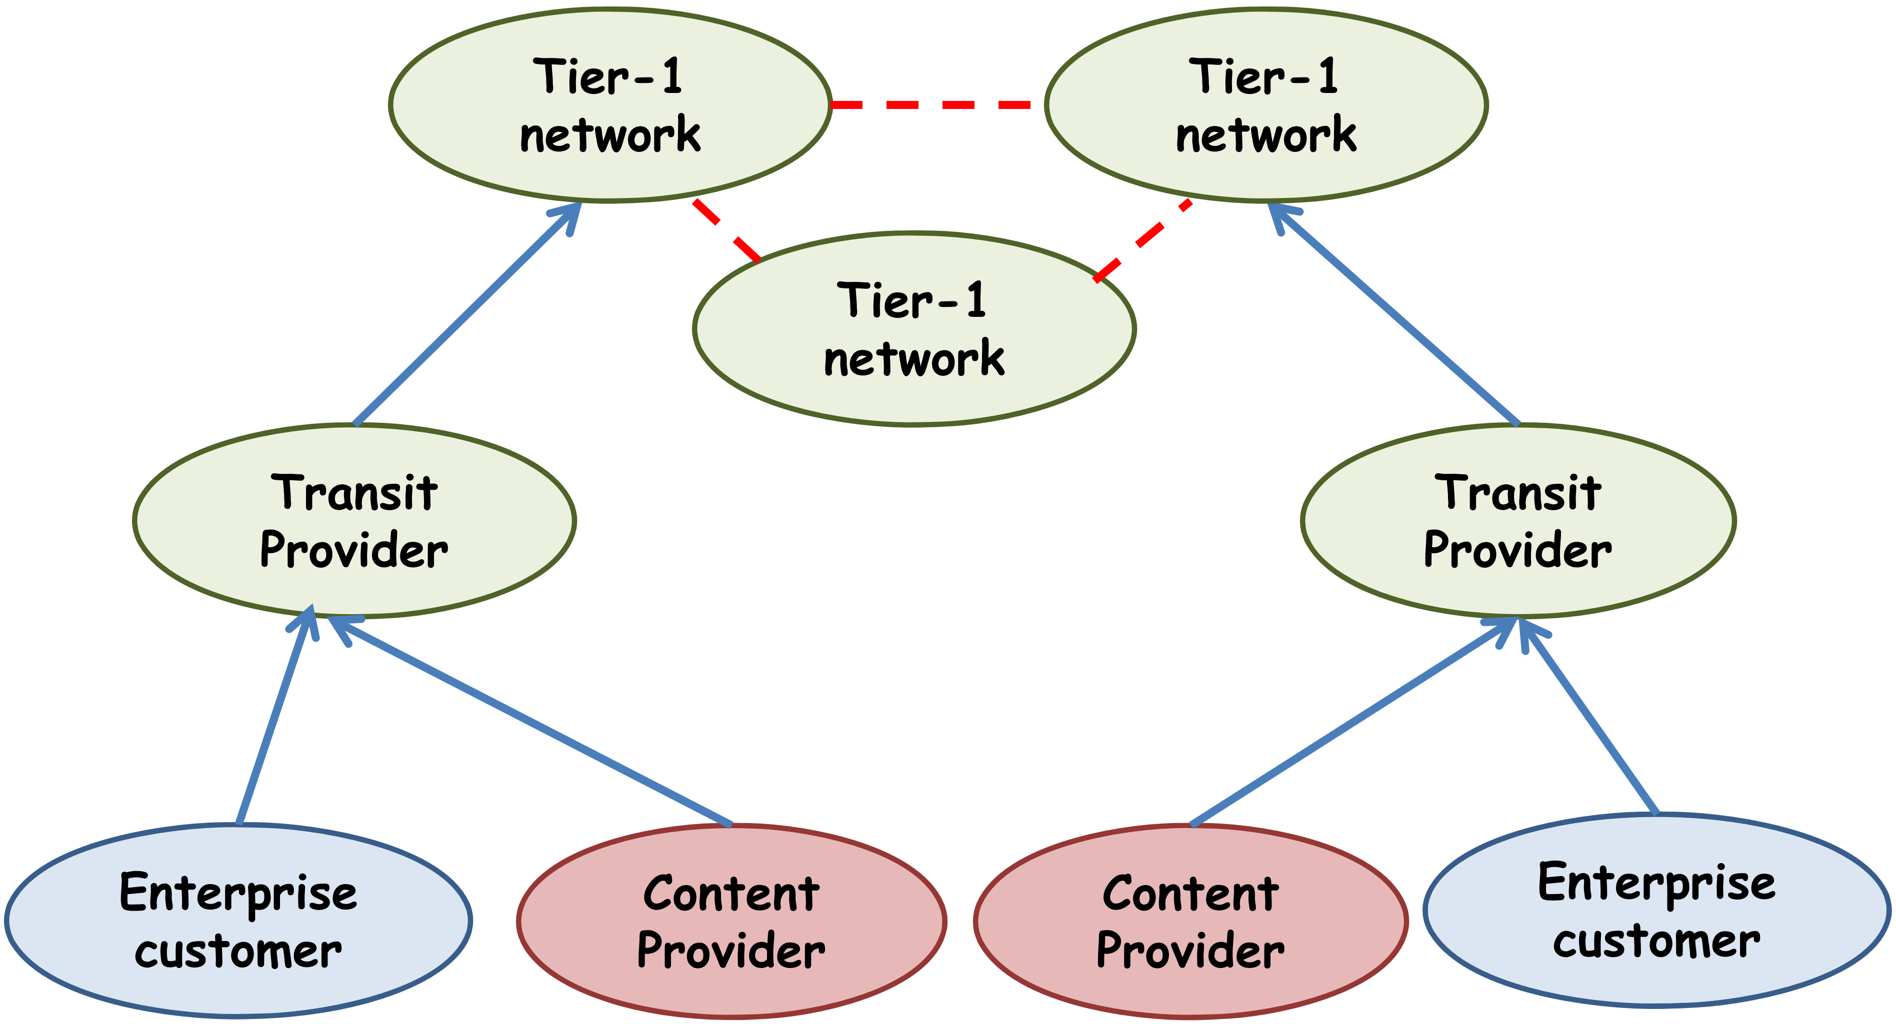

The AS-level Internet

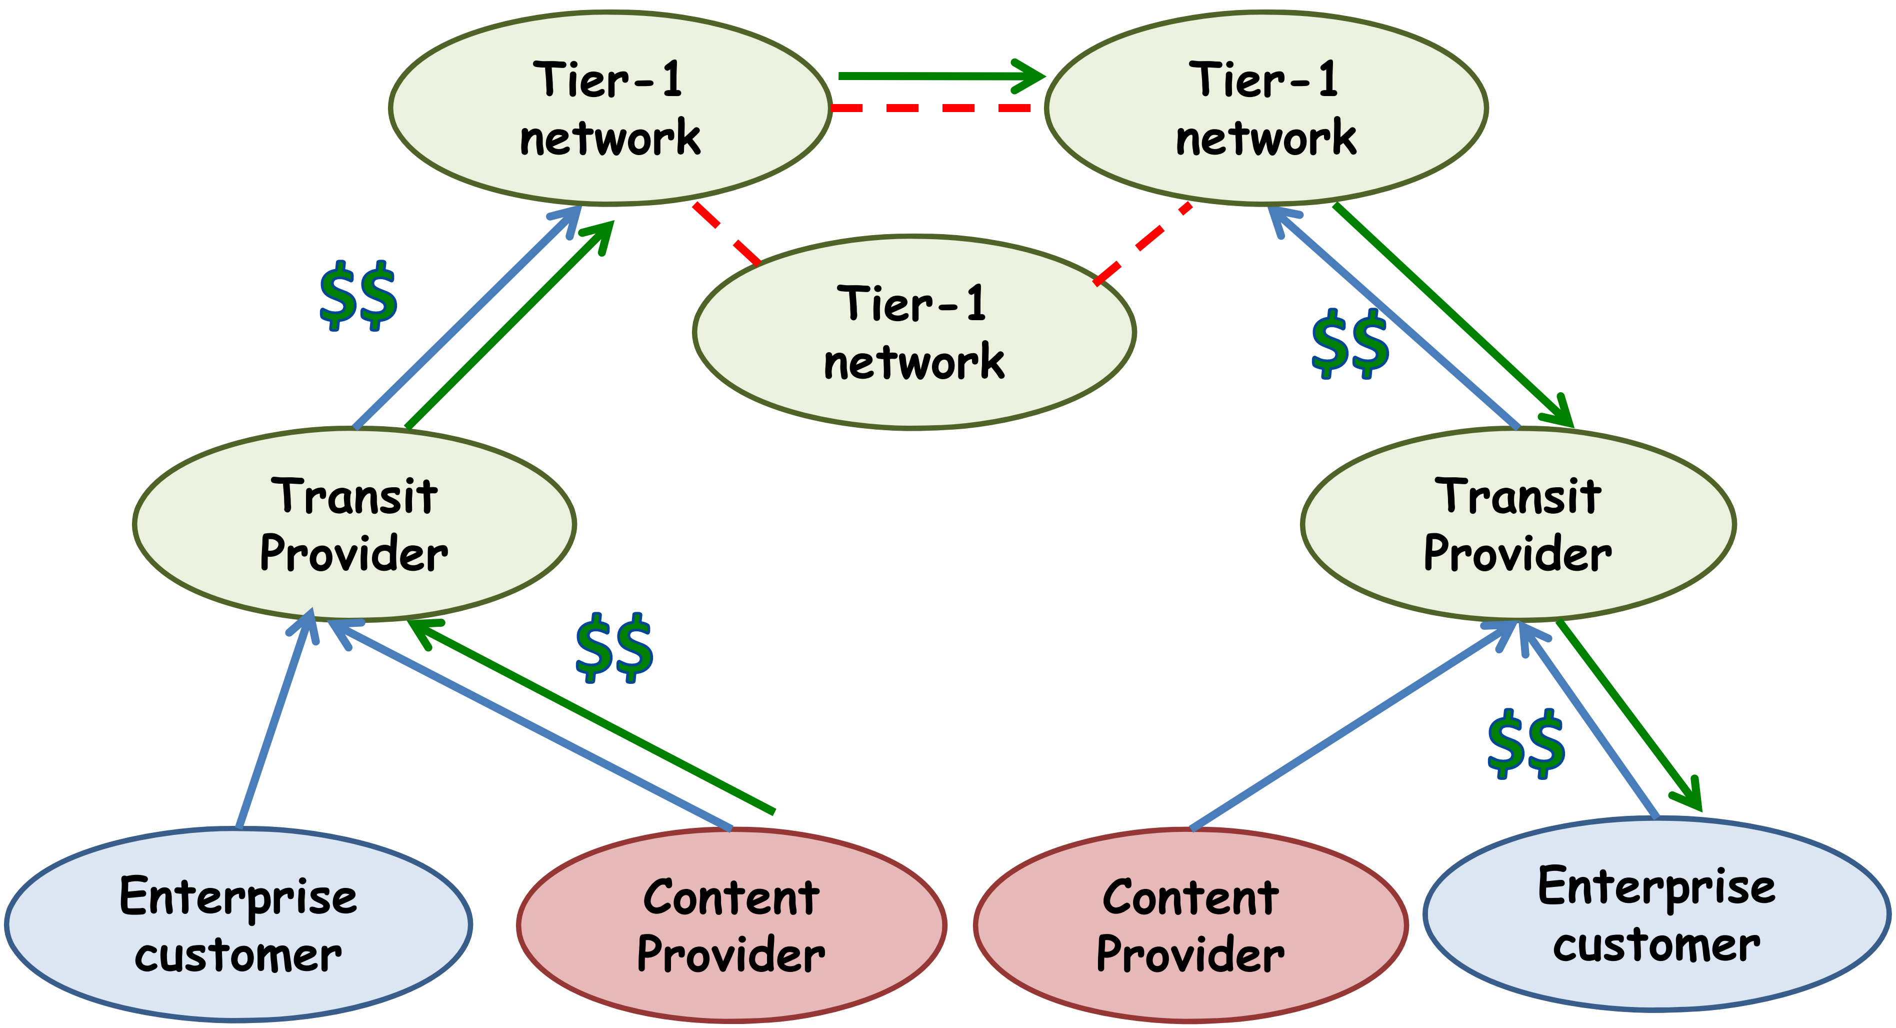

The AS-level Internet

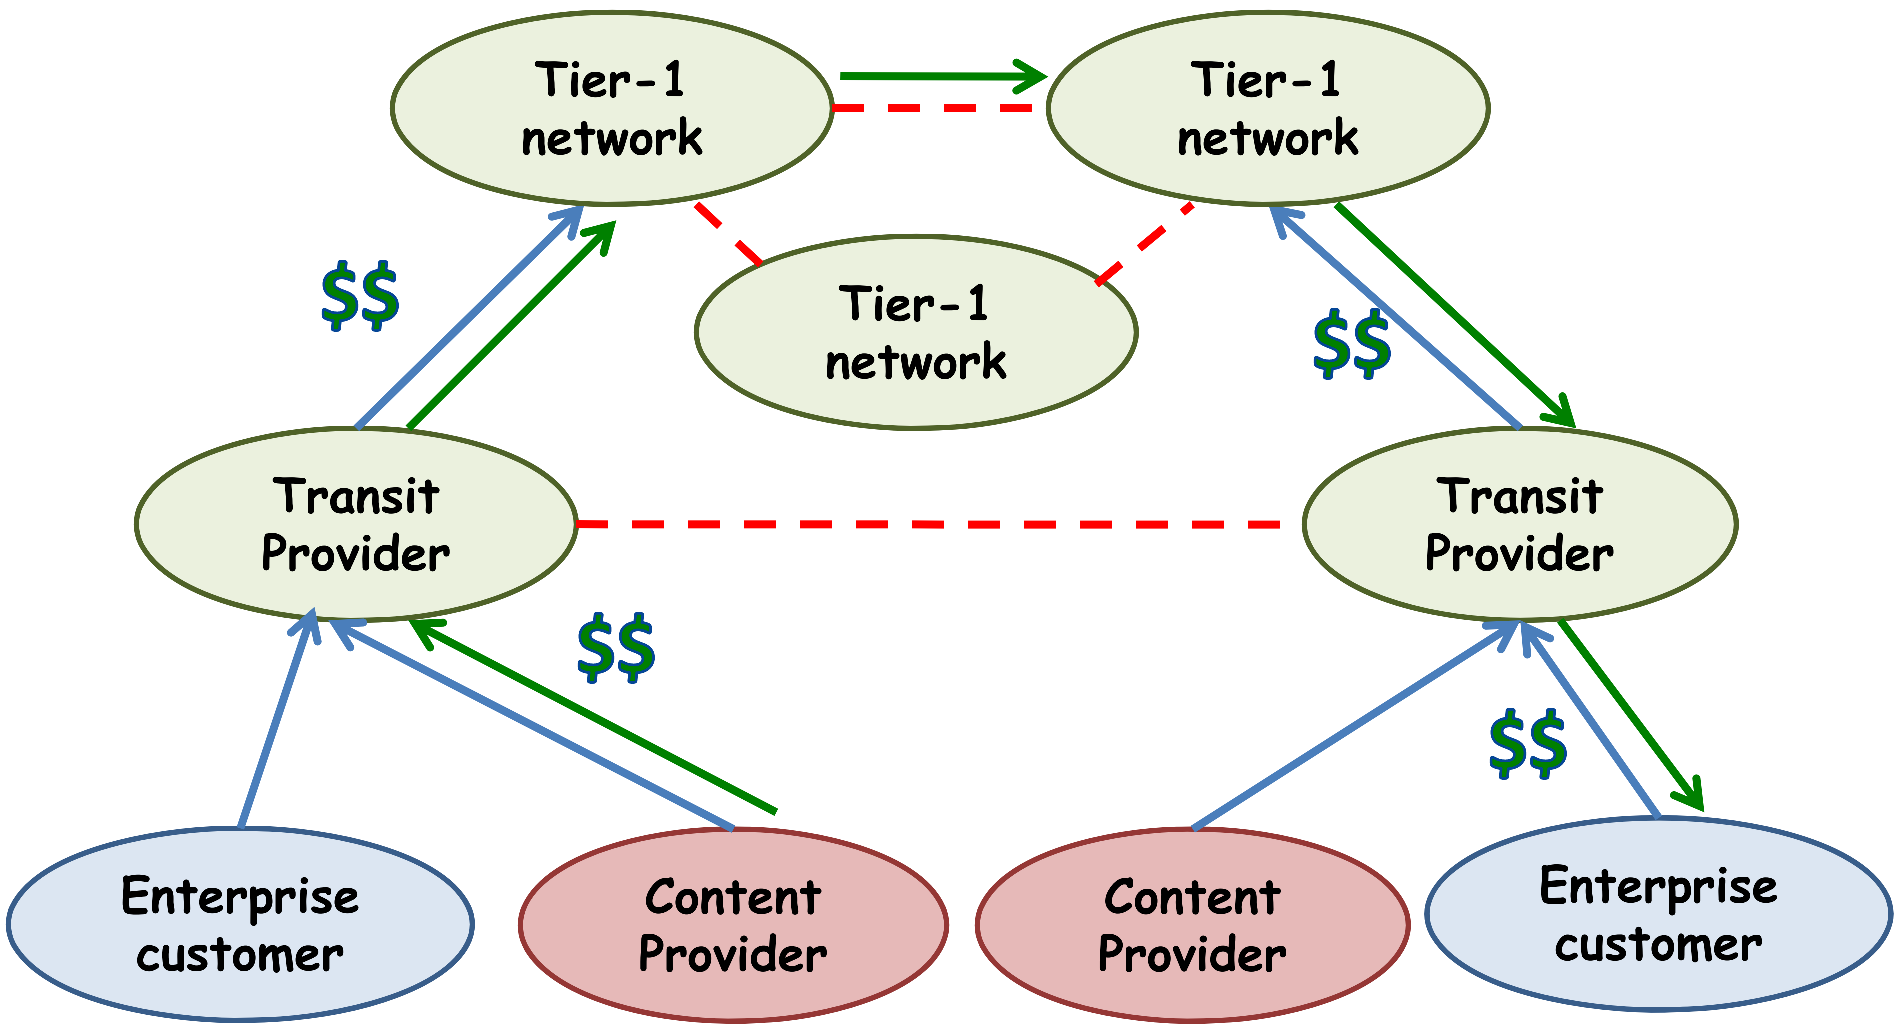

The AS-level Internet

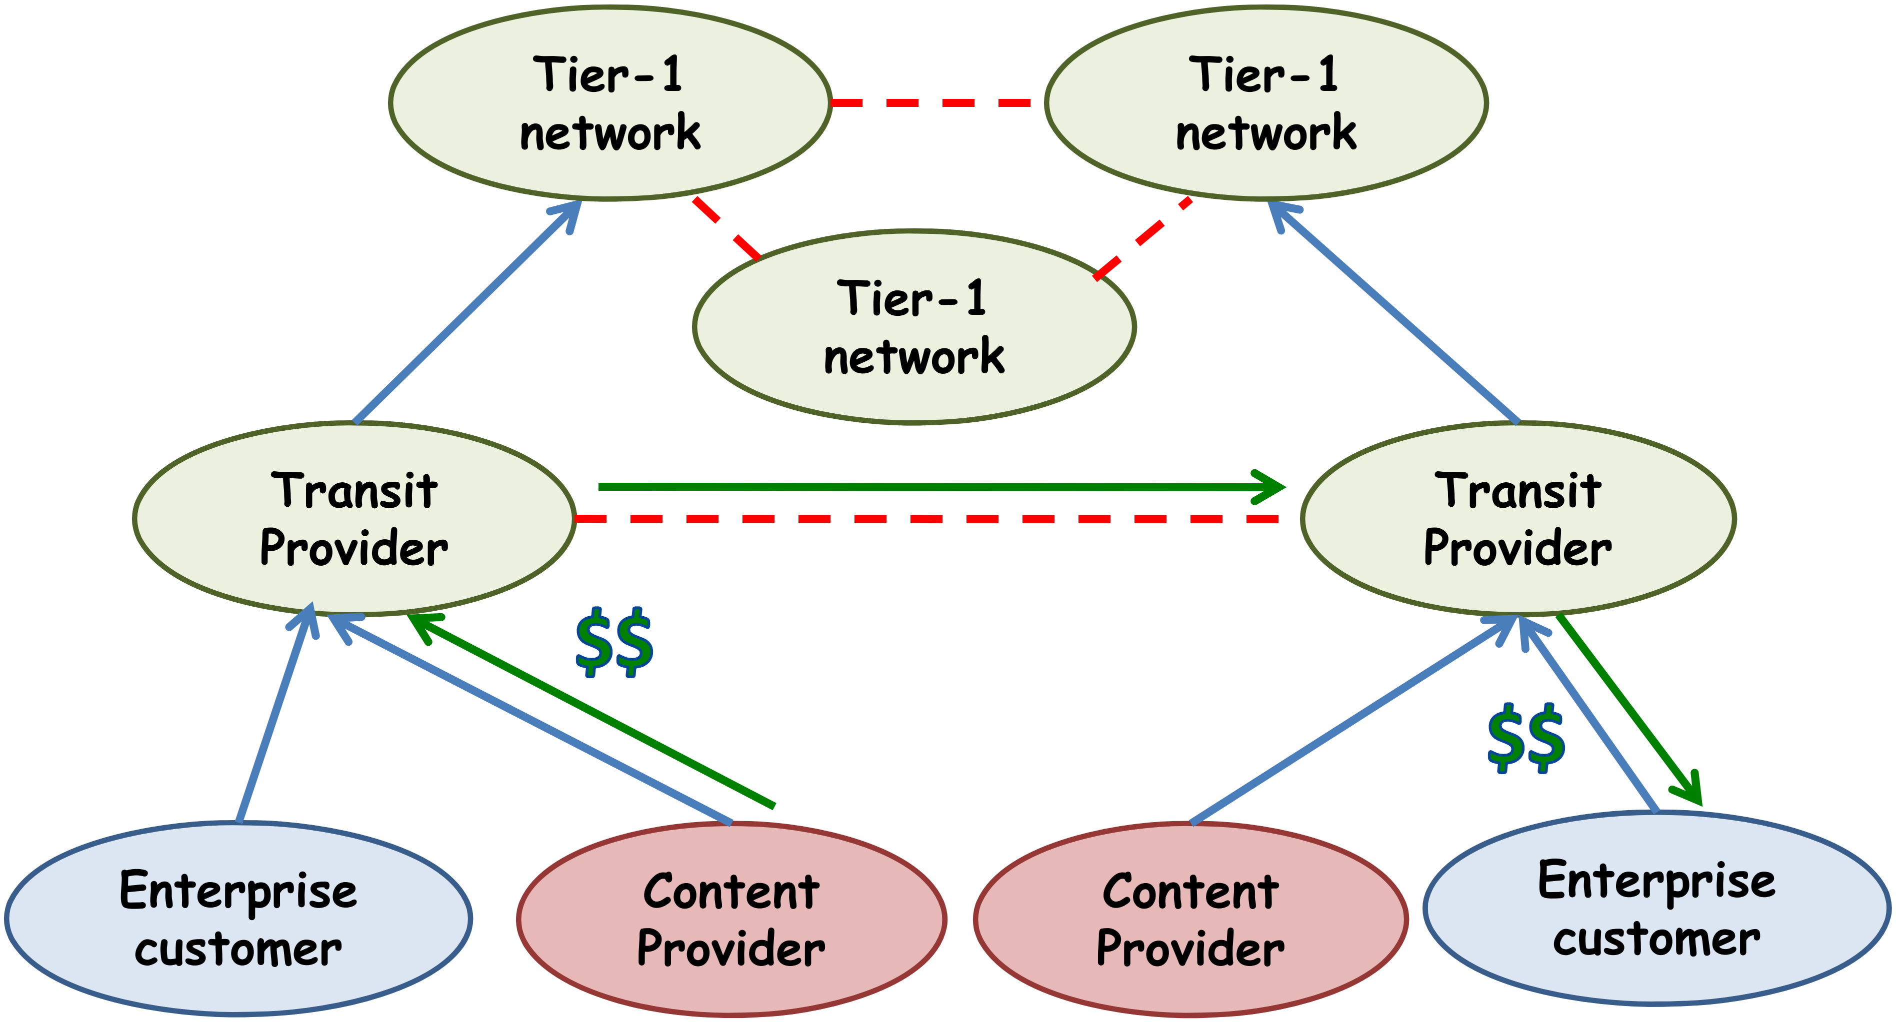

The AS-level Internet

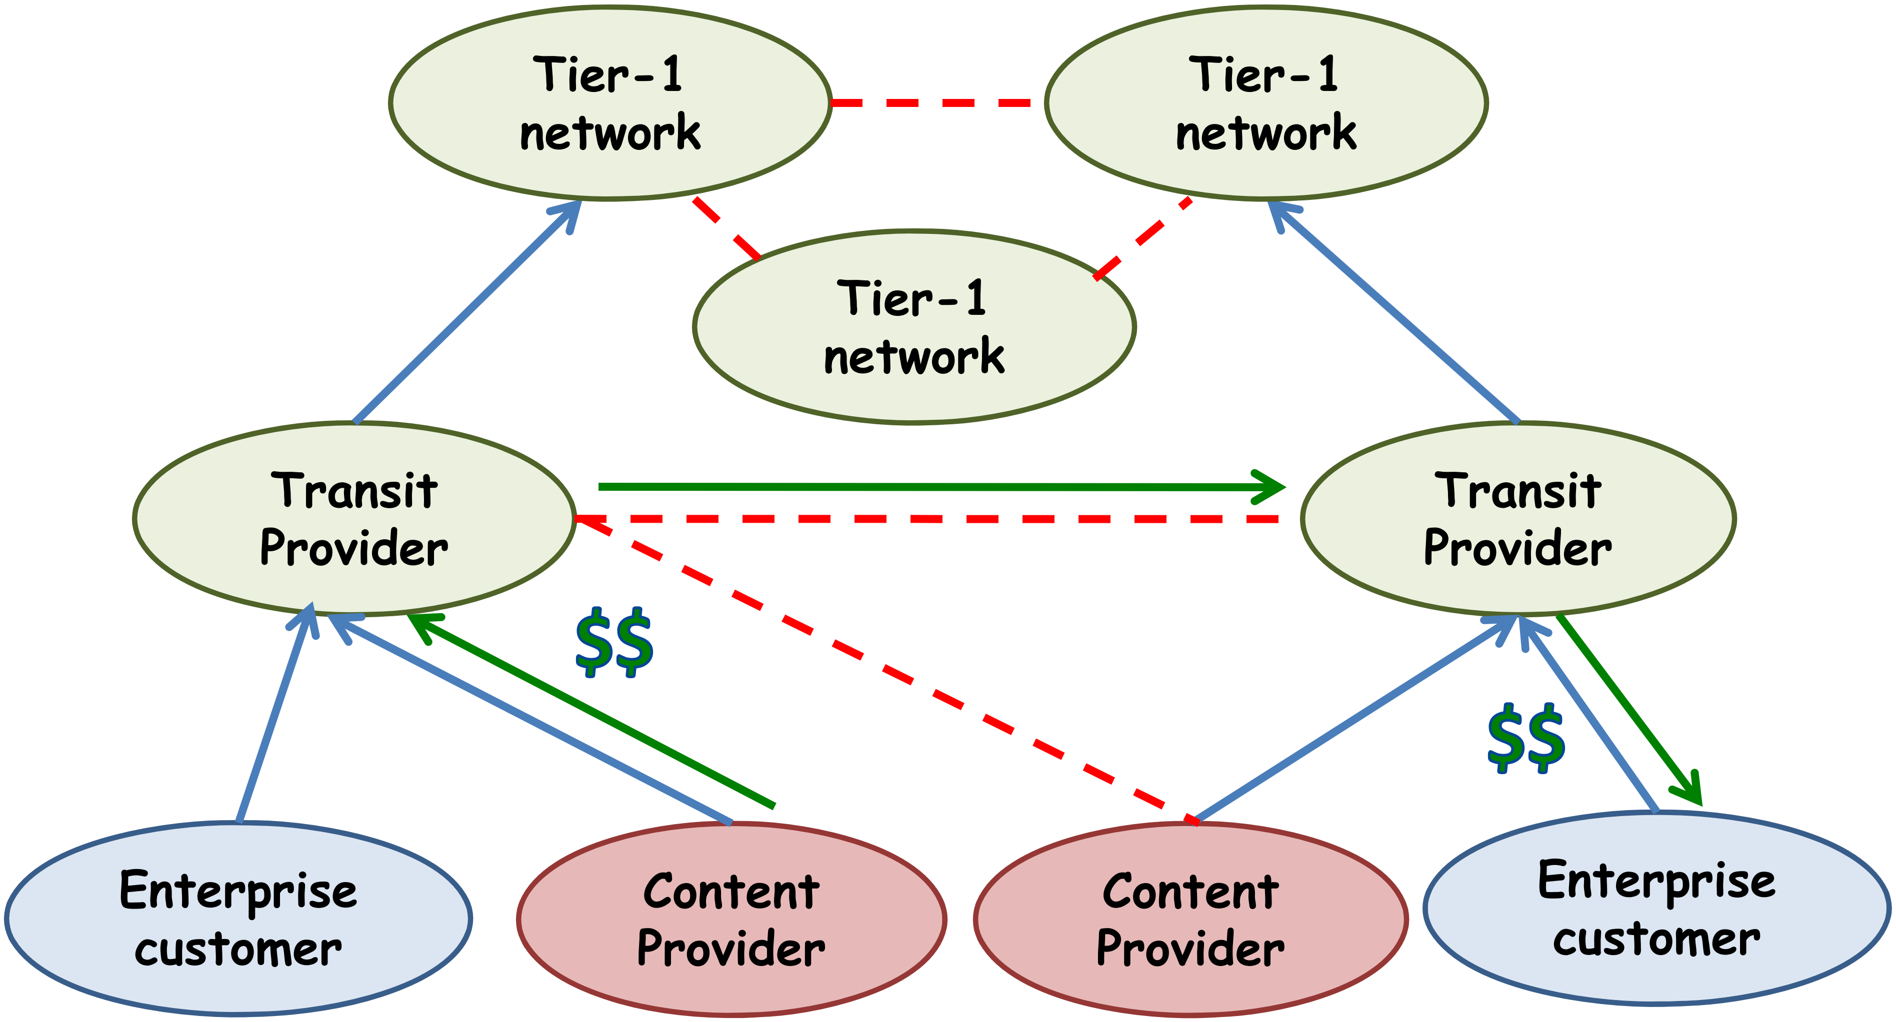

The AS-level Internet

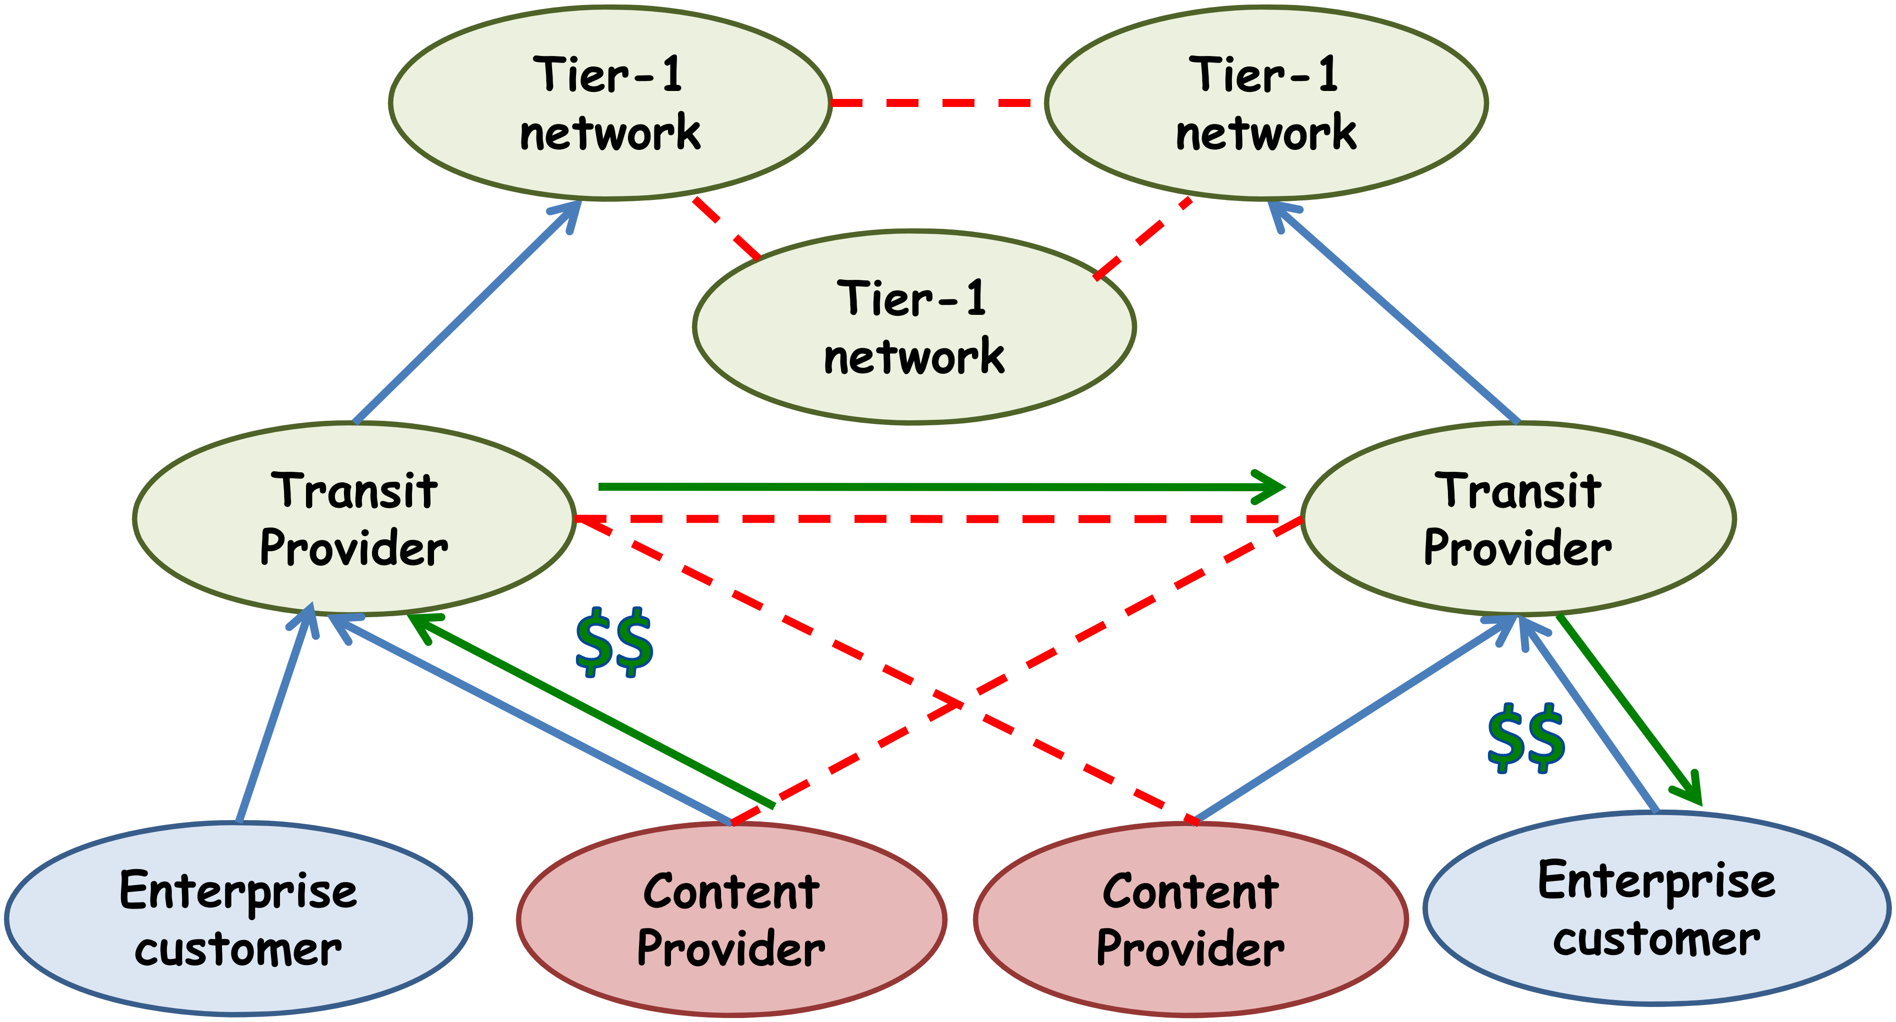

The AS-level Internet

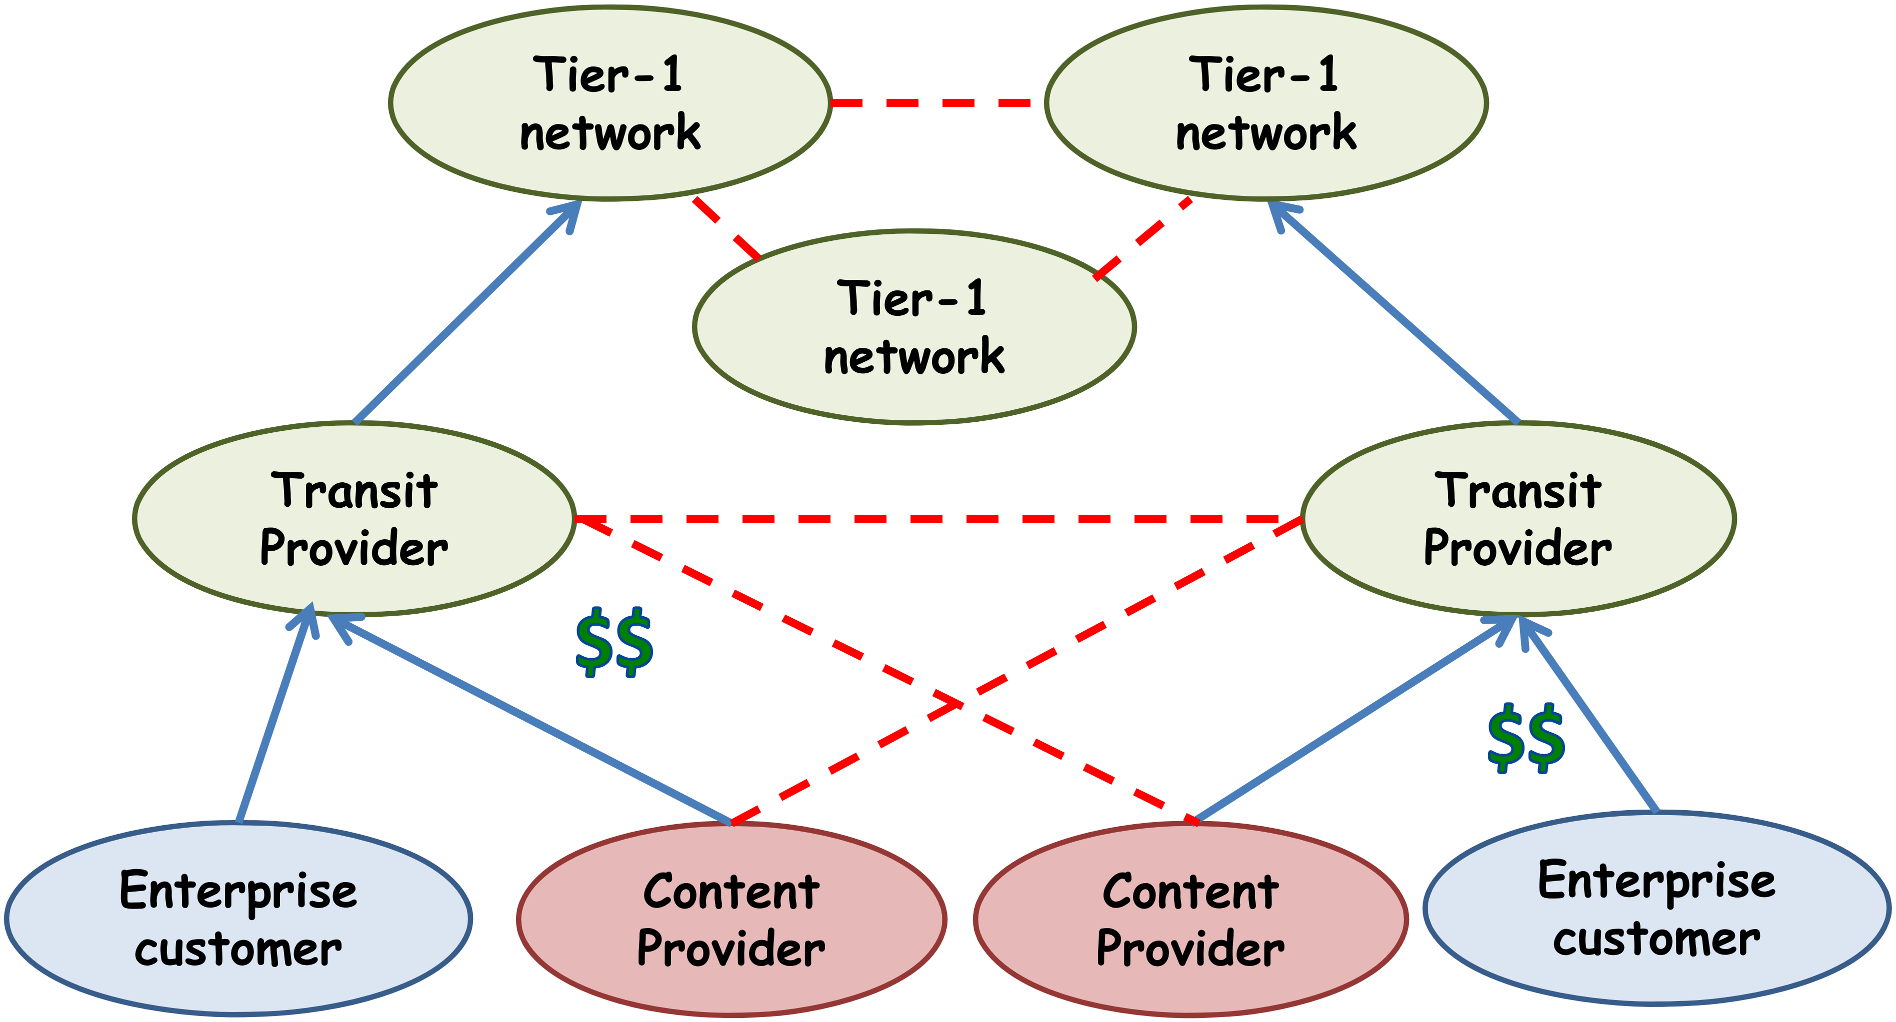

The AS-level Internet

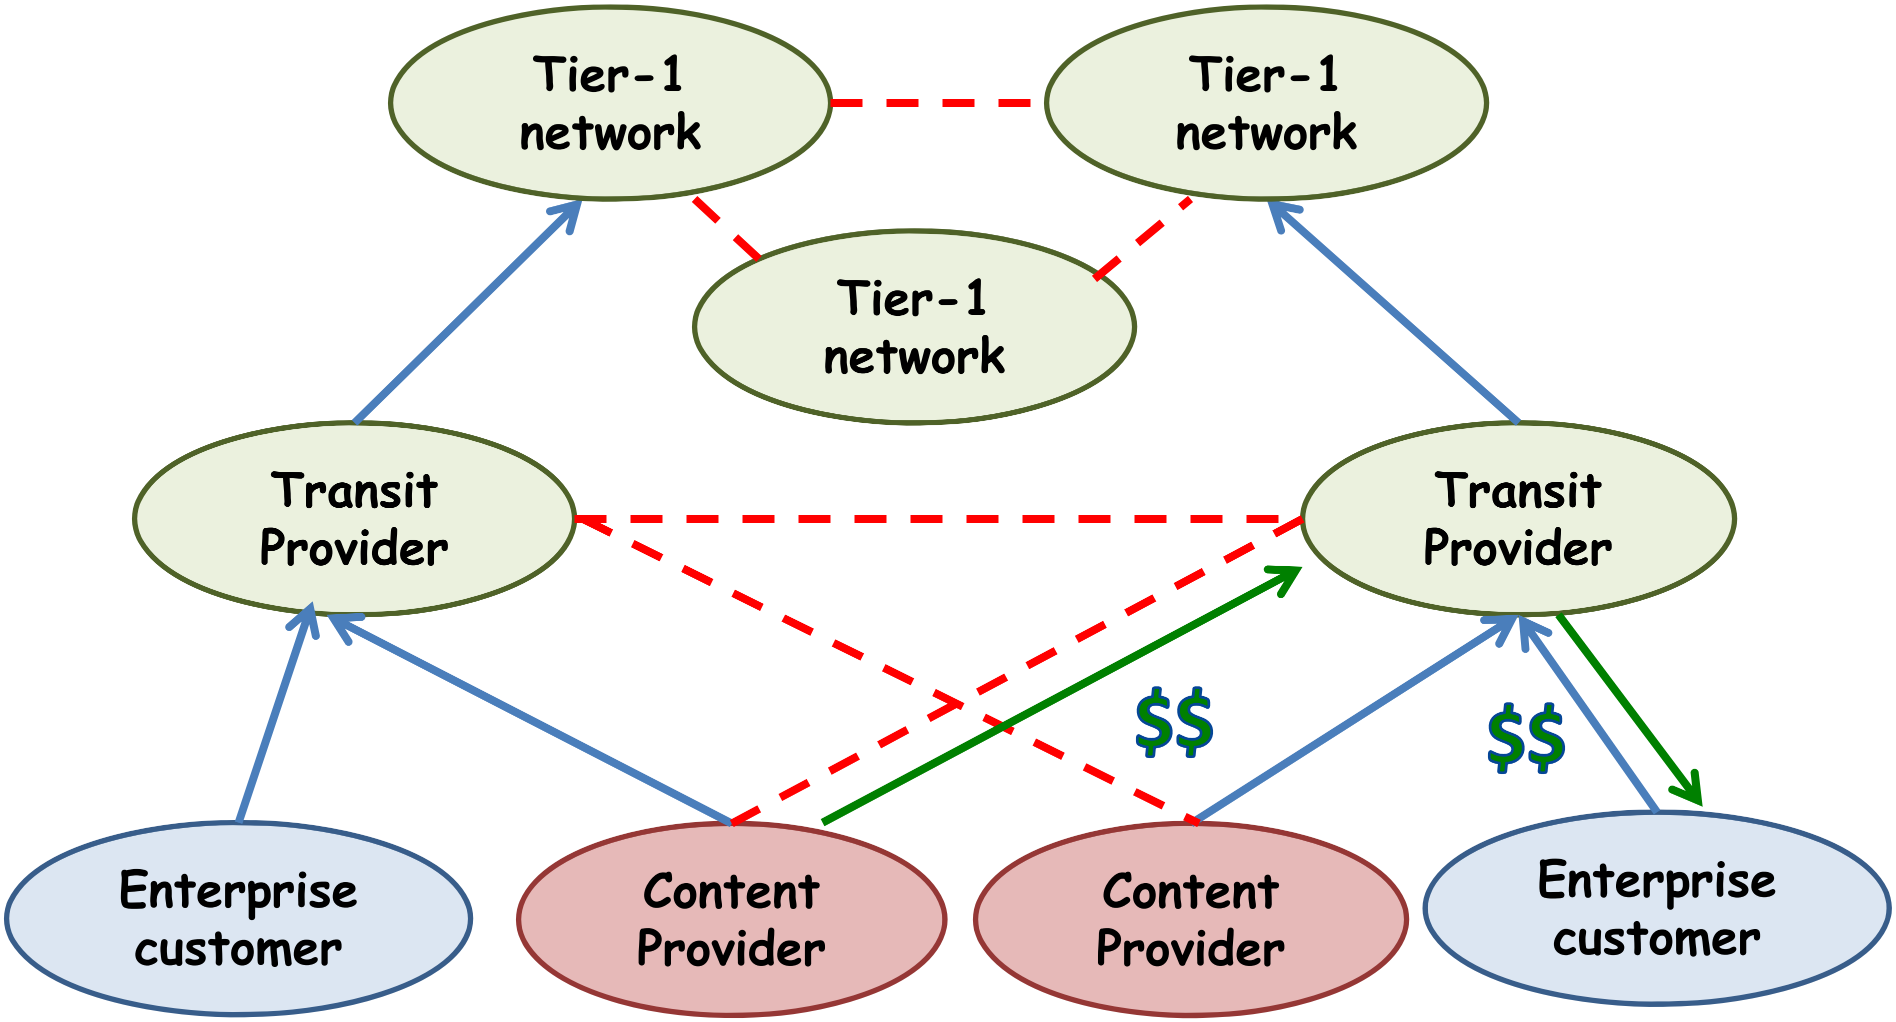

Topology, traffic, money

- An interdomain link represents a business relationship
- Complex interaction between topology, traffic flow and the flow of money
- Topology and business relationships determine traffic flow; traffic flow determines flow of money
- Topology and business relationships both evolve!

Topology, traffic, money

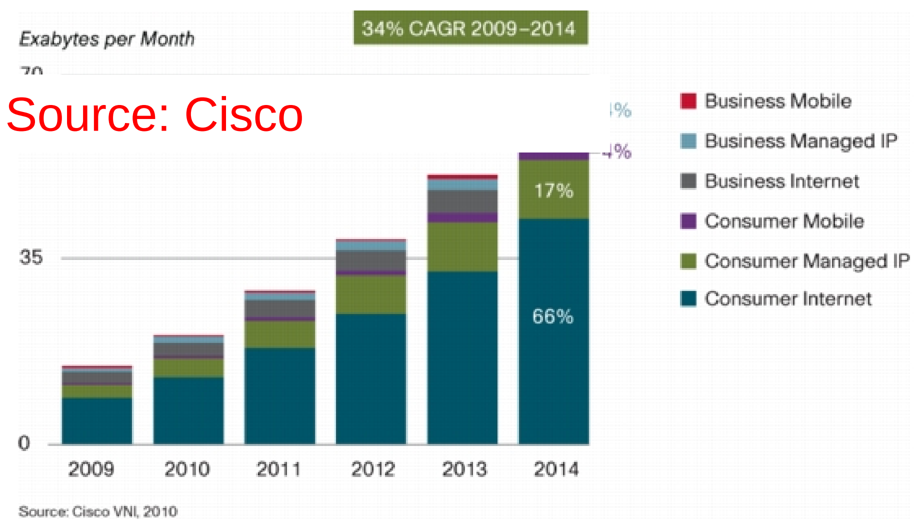

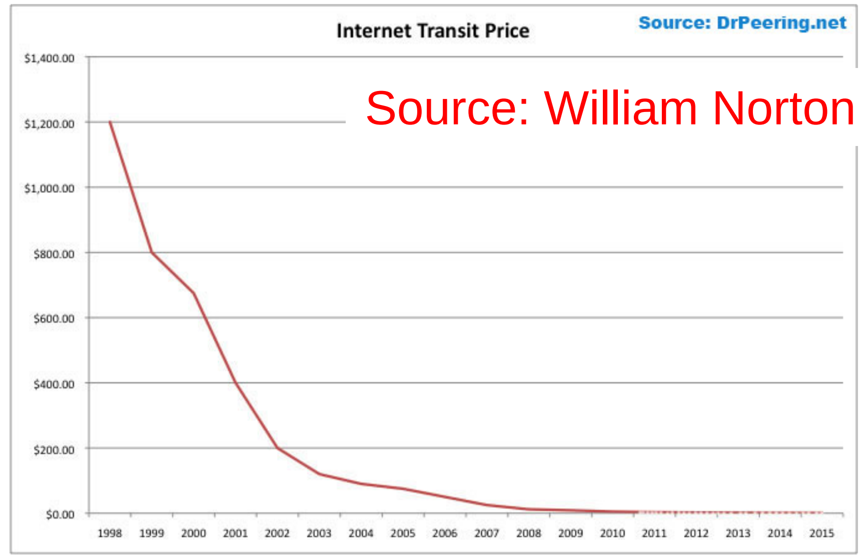

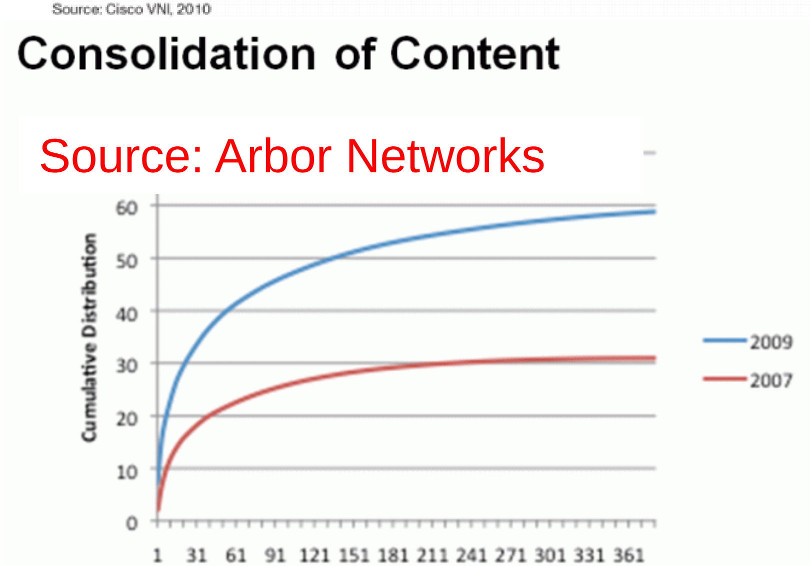

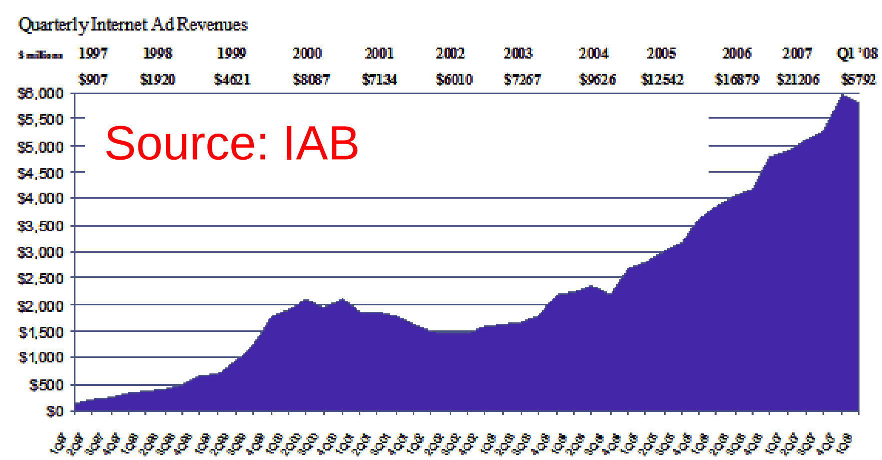

Outline

- AS topology as a network of business relationships
- Measuring the AS-level Internet
- Structural properties of the AS-level Internet
- Evolution of the AS-level Internet
- Modeling the structure and evolution of the AS-level Internet

AS Business Relationships

- A link between two ASes represents a business relationship
- Broad spectrum of business relationships
- Research literature has mostly considered the two extremes
- Customer-provider: customer pays provider for transit to the rest of the Internet
- Peering: Networks provide access to their respective customers (usually for free)

Business Relationships and Routing

- Business relationships influence AS routing decisions
- “Valley-free, prefer-customer, prefer-peer” routing policy

Business Relationships and Routing

- Business relationships influence AS routing decisions
- "Valley-free, prefer-customer, prefer-peer" routing policy

Do not advertise routes from a provider/peer to another provider/peer

Business Relationships and Routing

- Business relationships influence AS routing decisions
- “Valley-free, prefer-customer, prefer-peer” routing policy

Prefer a customer route (revenue generating) over a peer (free) or provider (paid) route

Business Relationships and Routing

- Business relationships influence AS routing decisions
- "Valley-free, prefer-customer, prefer-peer" routing policy

Prefer a peer route (free) over a provider route (paid)

Business Relationships Affect Traffic Flow

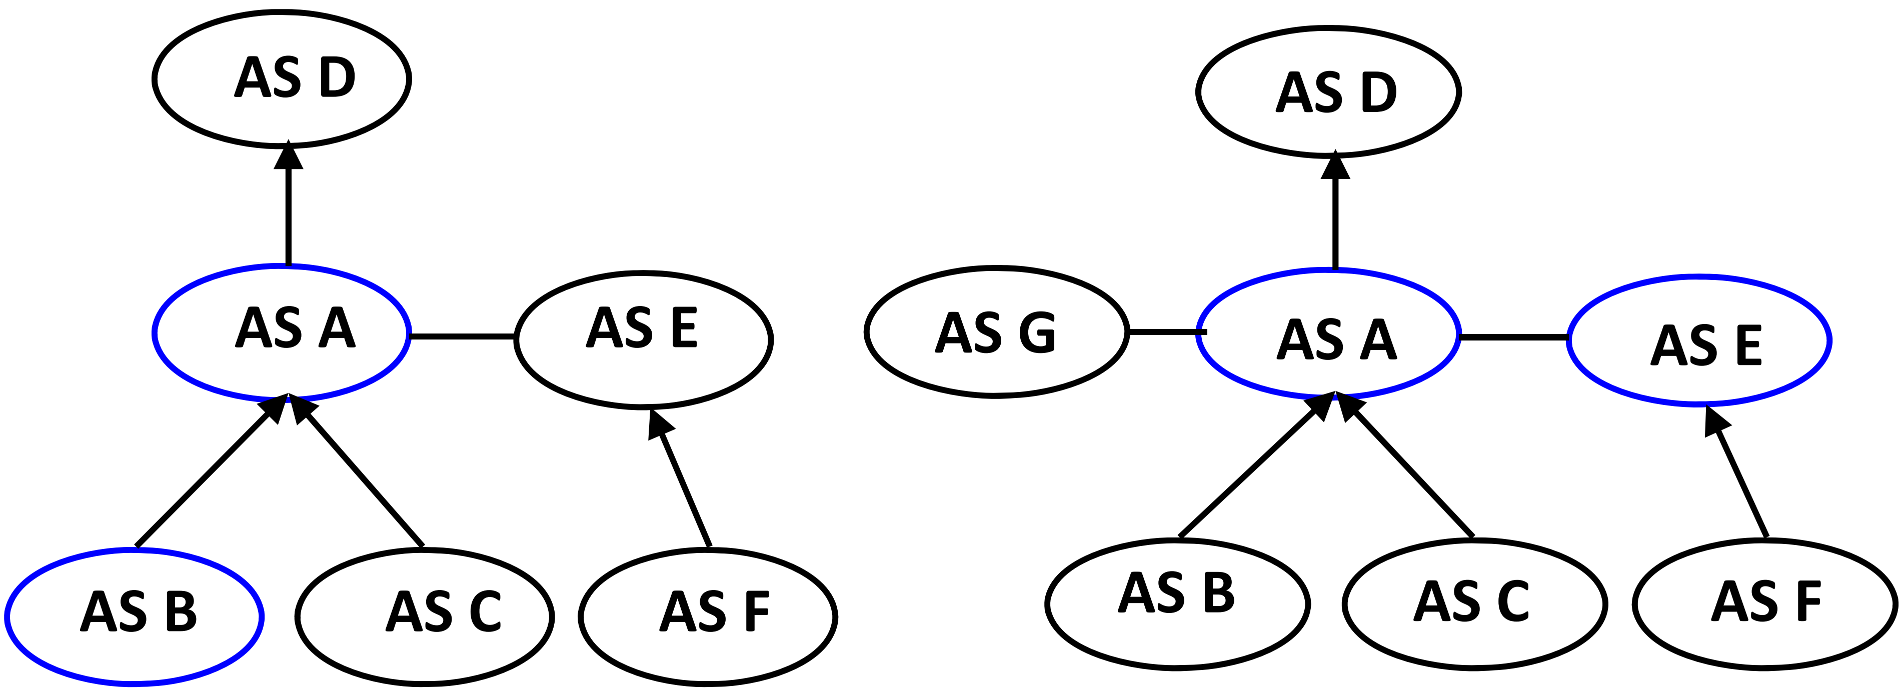

Business Relationships Affect Traffic Flow

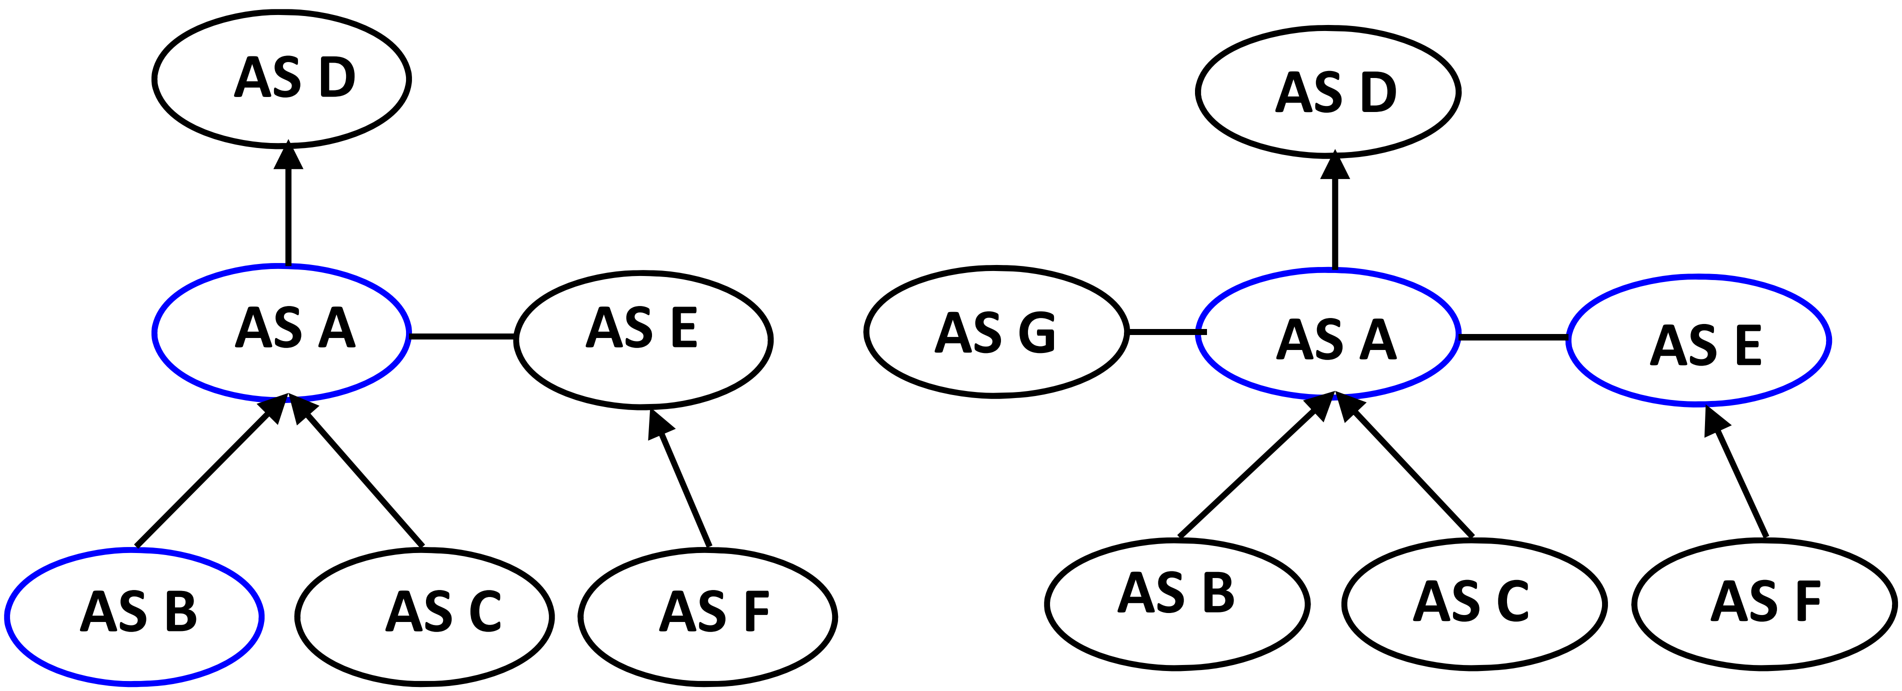

AS B customer of AS A

Business Relationships Affect Traffic Flow

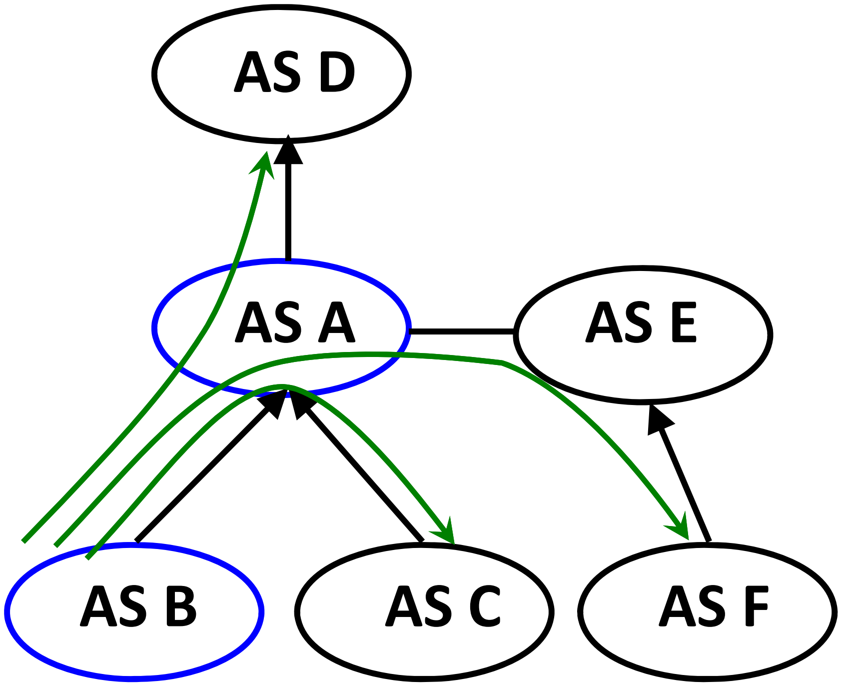

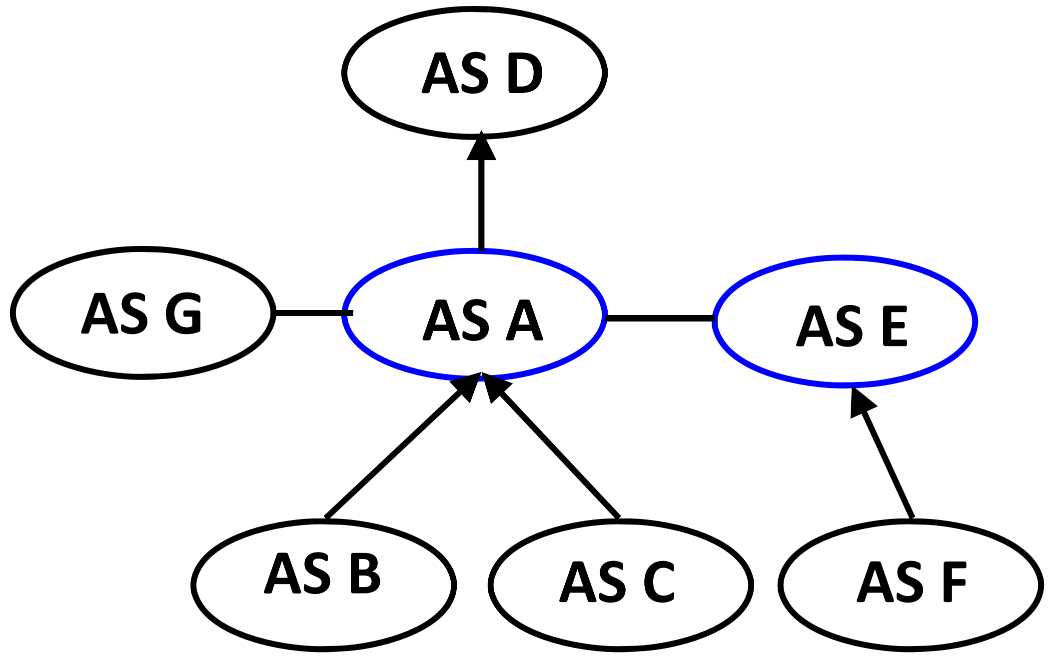

AS B customer of AS A

Business Relationships Affect Traffic Flow

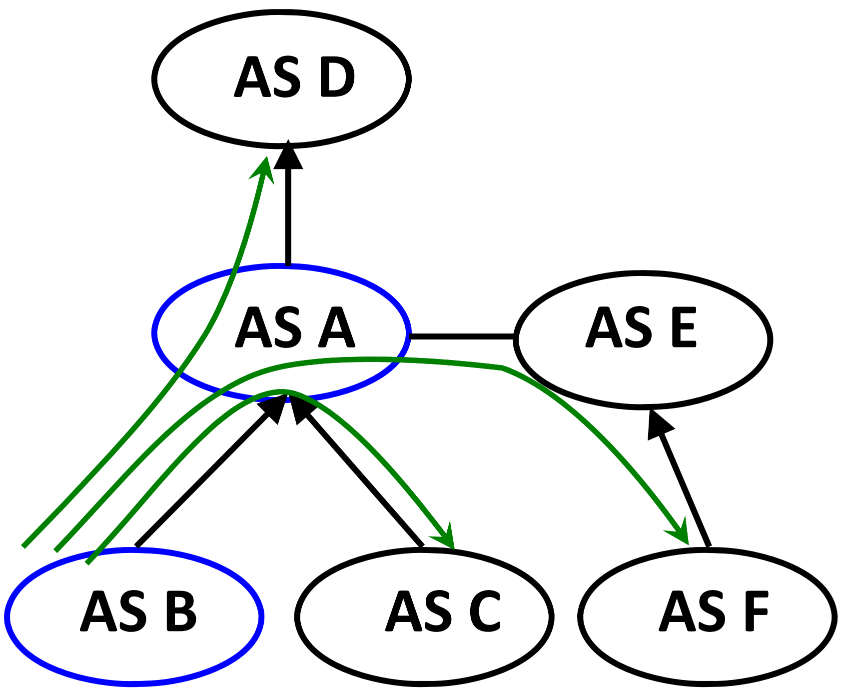

AS B customer of AS A

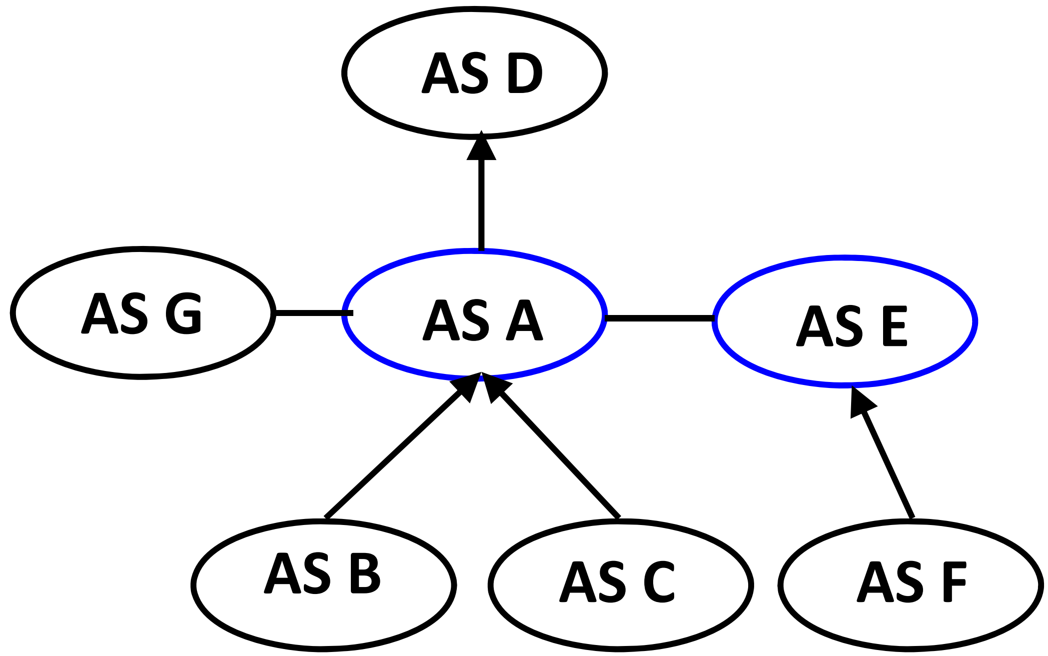

AS A and AS E are peers

Business Relationships Affect Traffic Flow

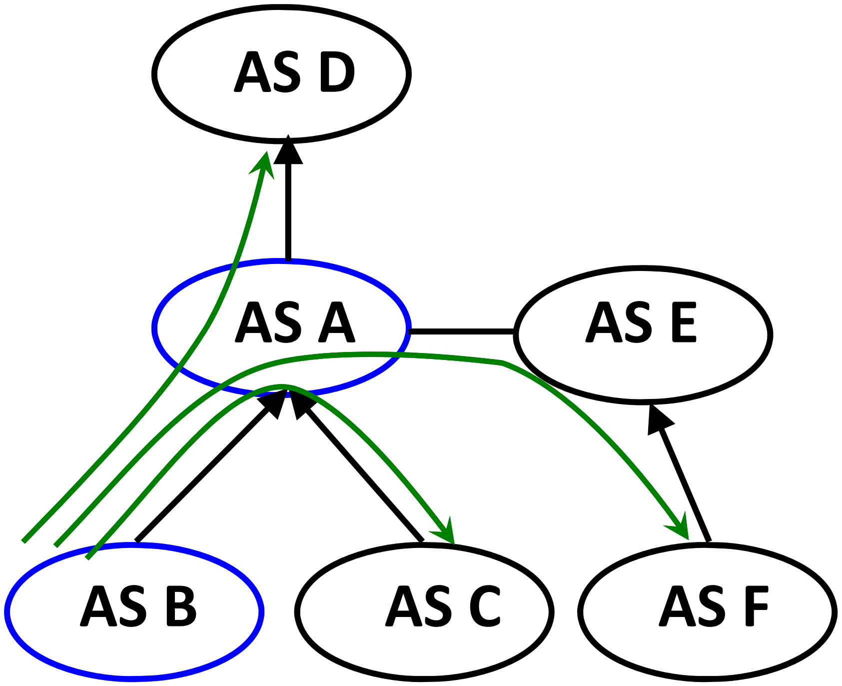

AS B customer of AS A

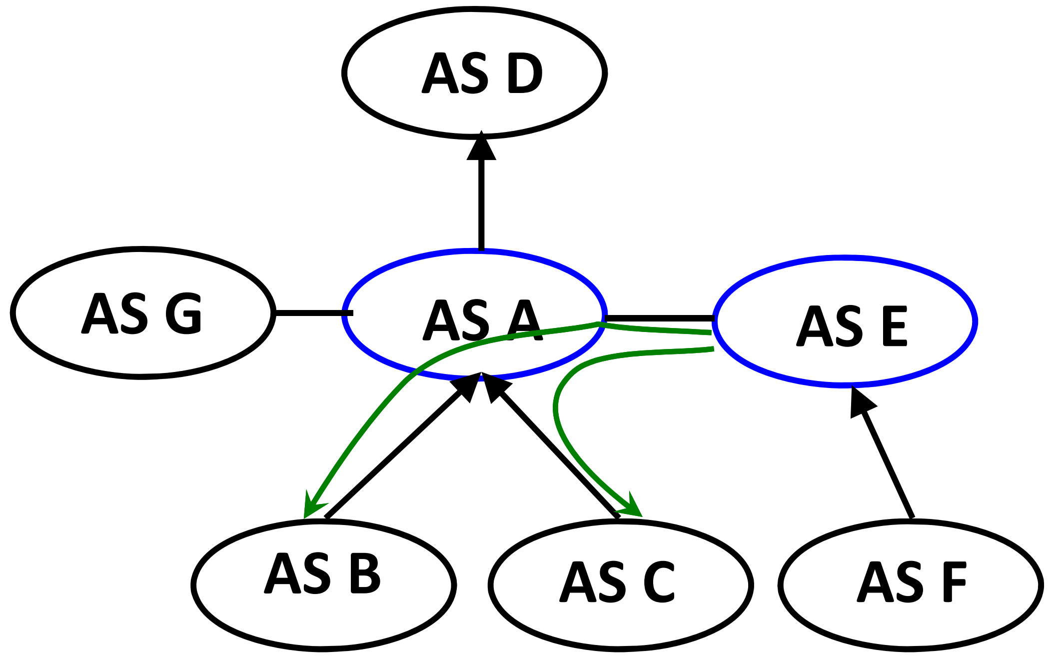

AS A and AS E are peers

Business Relationships Affect Traffic Flow

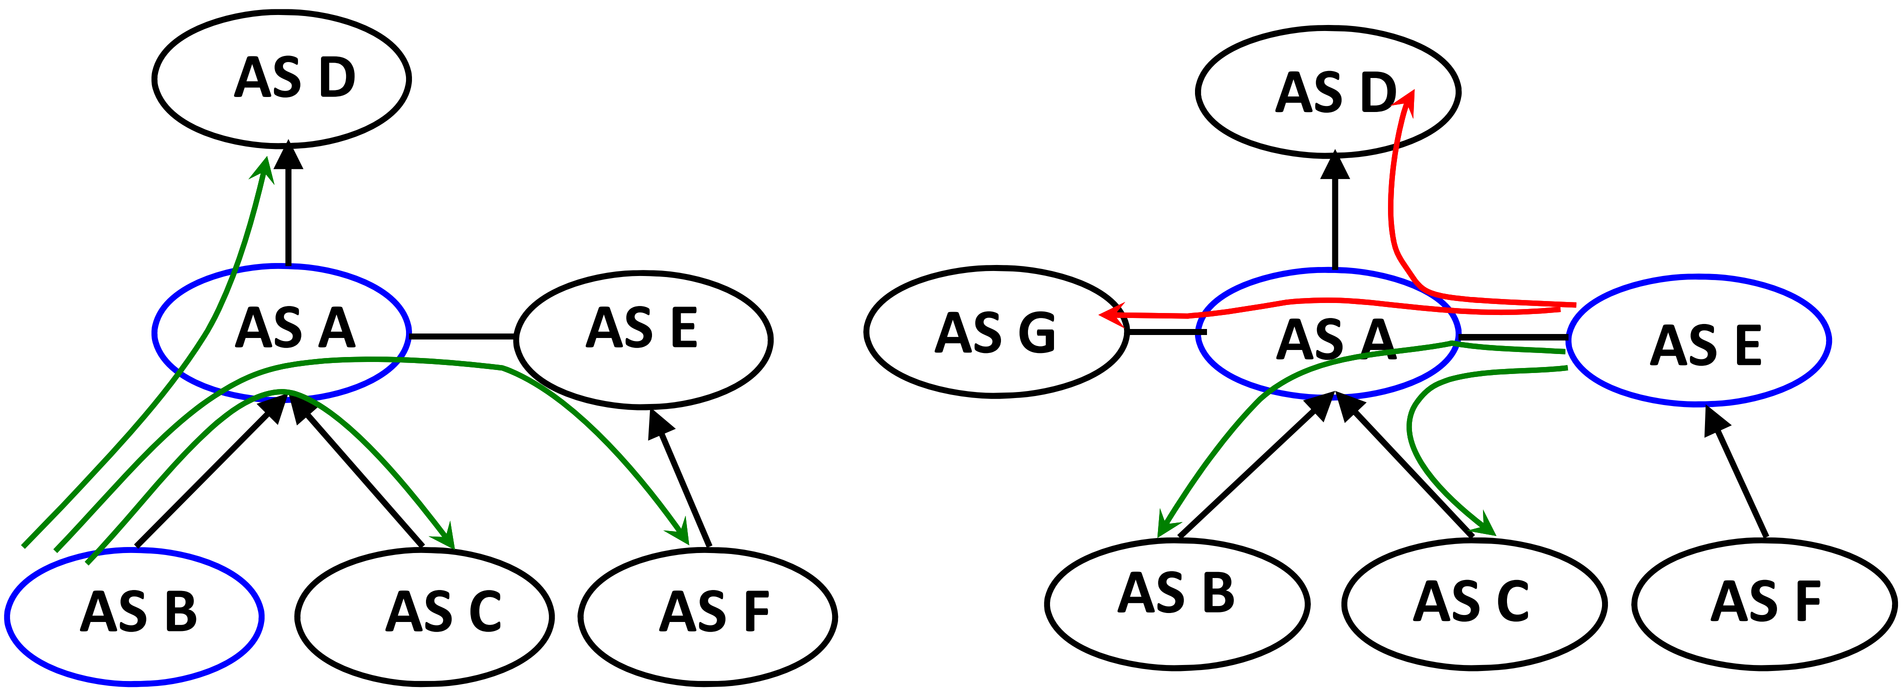

AS B customer of AS A

AS A and AS E are peers

Outline

- AS topology as a network of business relationships
- Measuring the AS-level Internet
- Structural properties of the AS-level Internet
- Evolution of the AS-level Internet
- Modeling the structure and evolution of the AS-level Internet

Measuring the AS-level Internet

- Large-scale traceroute projects (Ark, DIMES, etc.)
  - Issue traceroute from a set of vantage points
  - Convert IP-level paths into AS-level paths
  - Combine AS paths to construct AS topology
- Several issues with converting traceroute paths to AS-level paths
  - Third-party addresses
  - IXPs
- Sampling biases*: vantage points must be distributed, and probe the entire routed Internet

Lakhina, Byers, Crovella, Xie, "Sampling Biases in IP Topology Measurements", IEEE Infocom 2003.

Measuring the AS-level Internet

- BGP route monitors are ASes that volunteer to provide BGP feeds
  - Collect AS paths from each BGP monitor towards each routed prefix
  - Construct AS topology by combining AS paths from multiple vantage points
- Routeviews/RIPE RIS are two projects that have been collecting BGP feeds from volunteer ASes for many years
  - Currently about 400 volunteer ASes
- “Cleaner” to construct AS topology from BGP snapshots

Monitor Placement Matters!

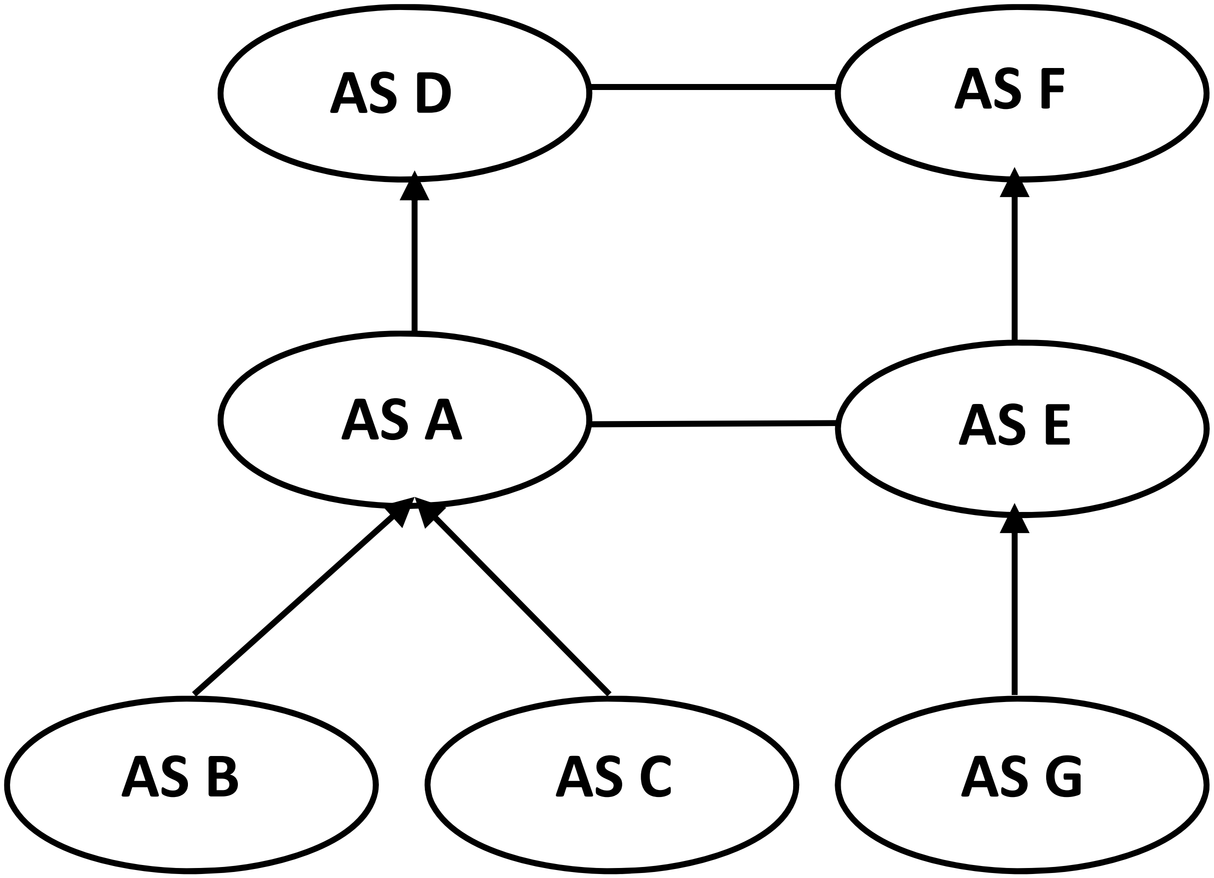

Monitor Placement Matters!

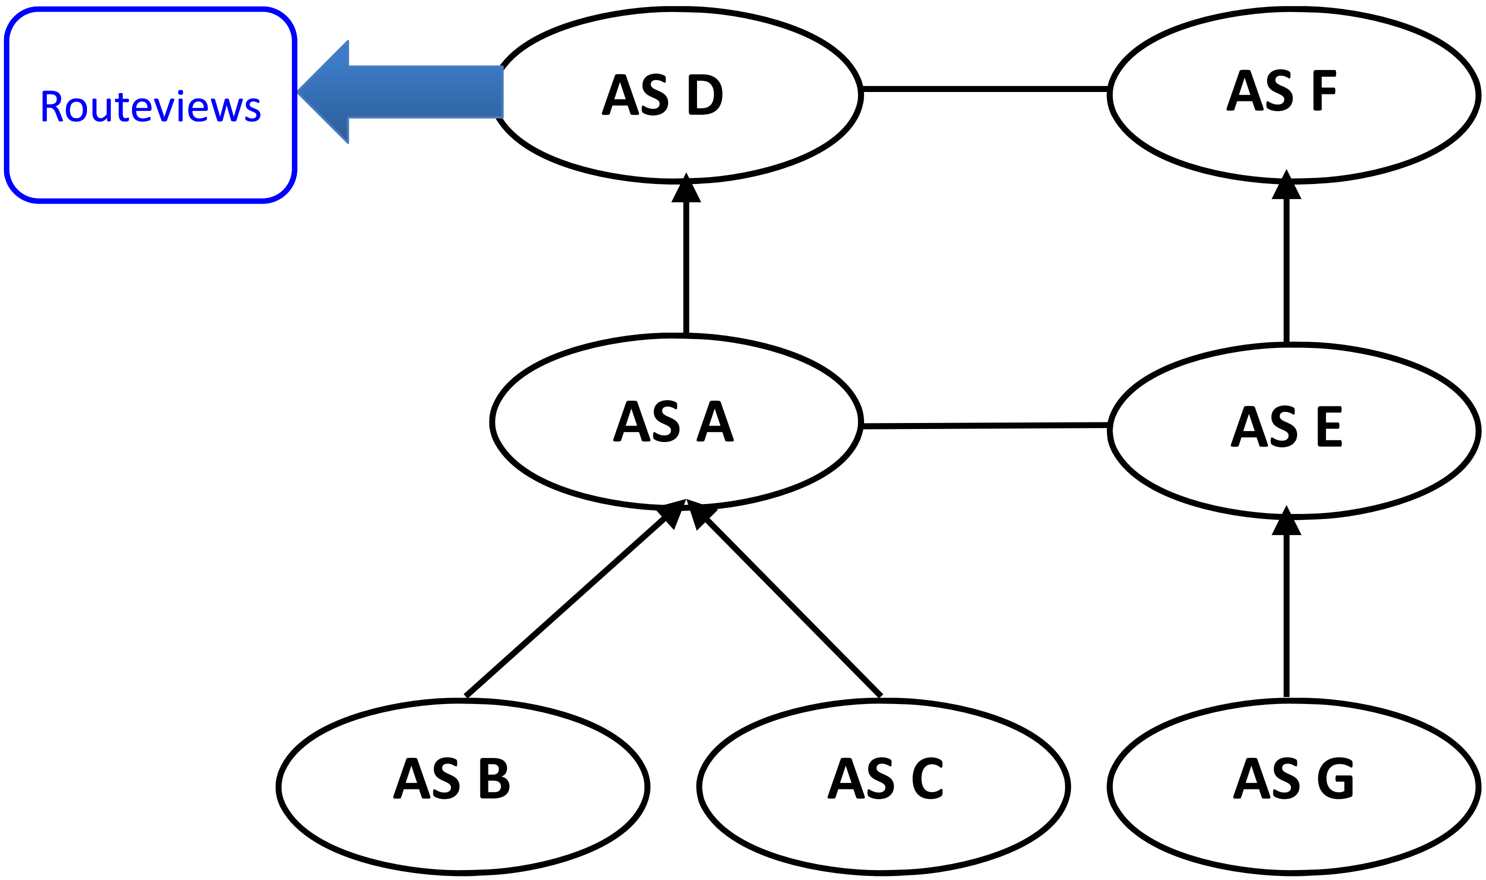

Monitor Placement Matters!

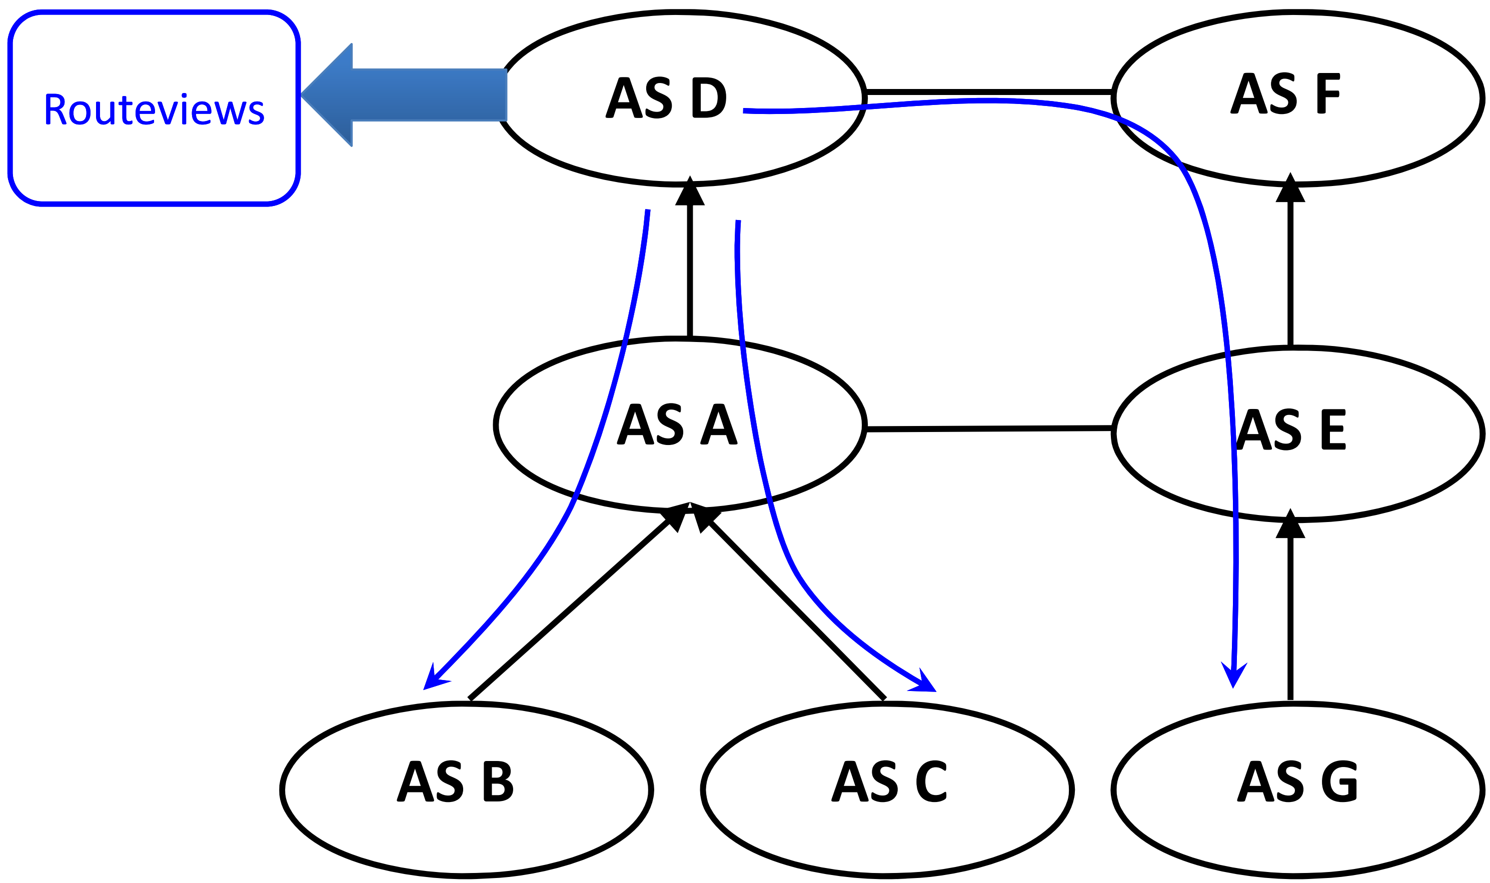

Monitor Placement Matters!

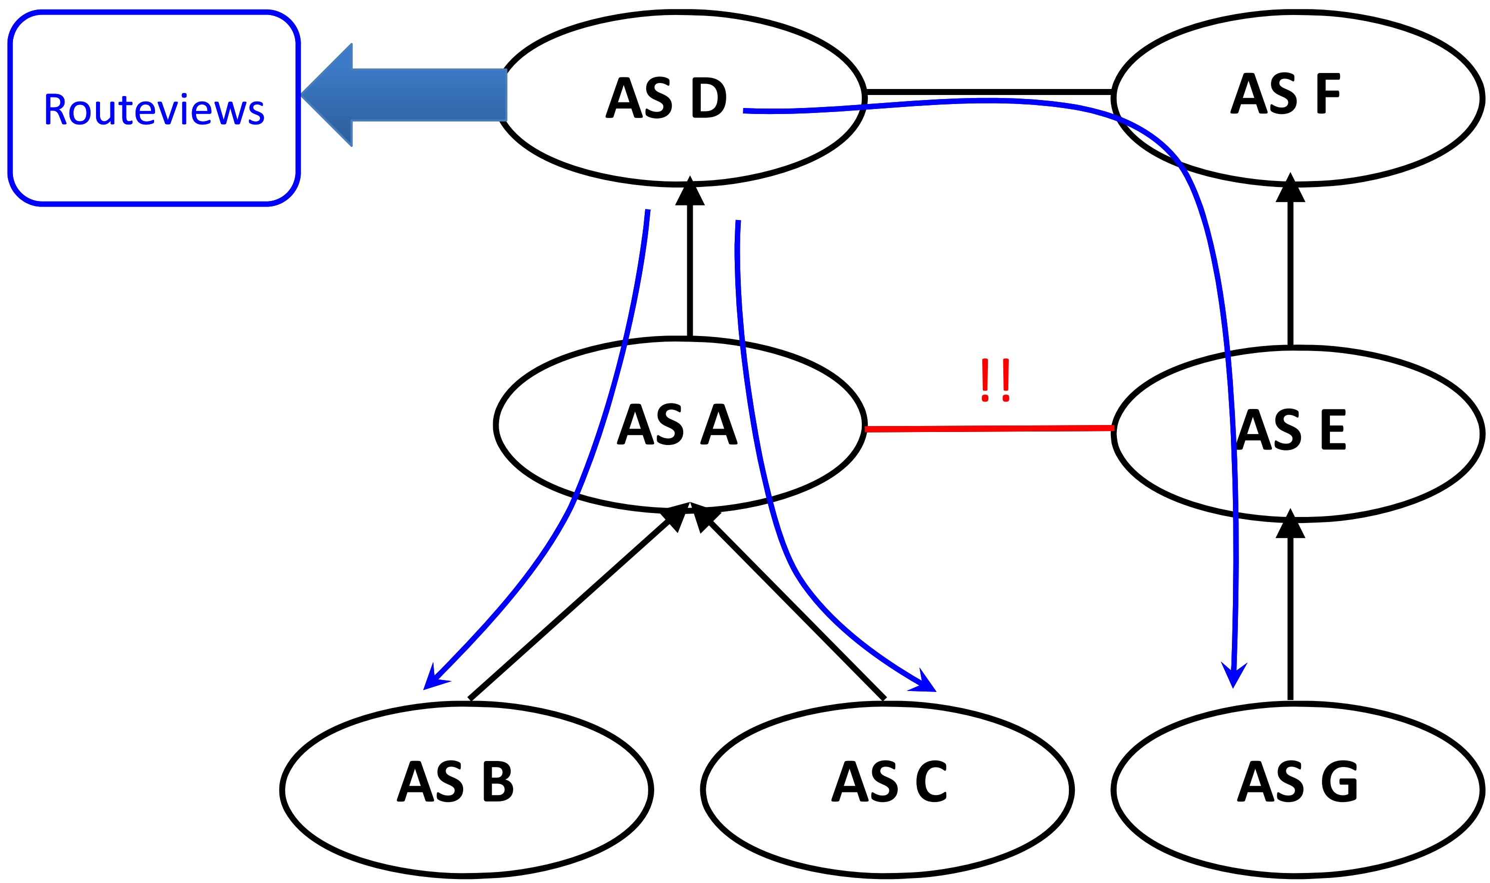

Monitor Placement Matters!

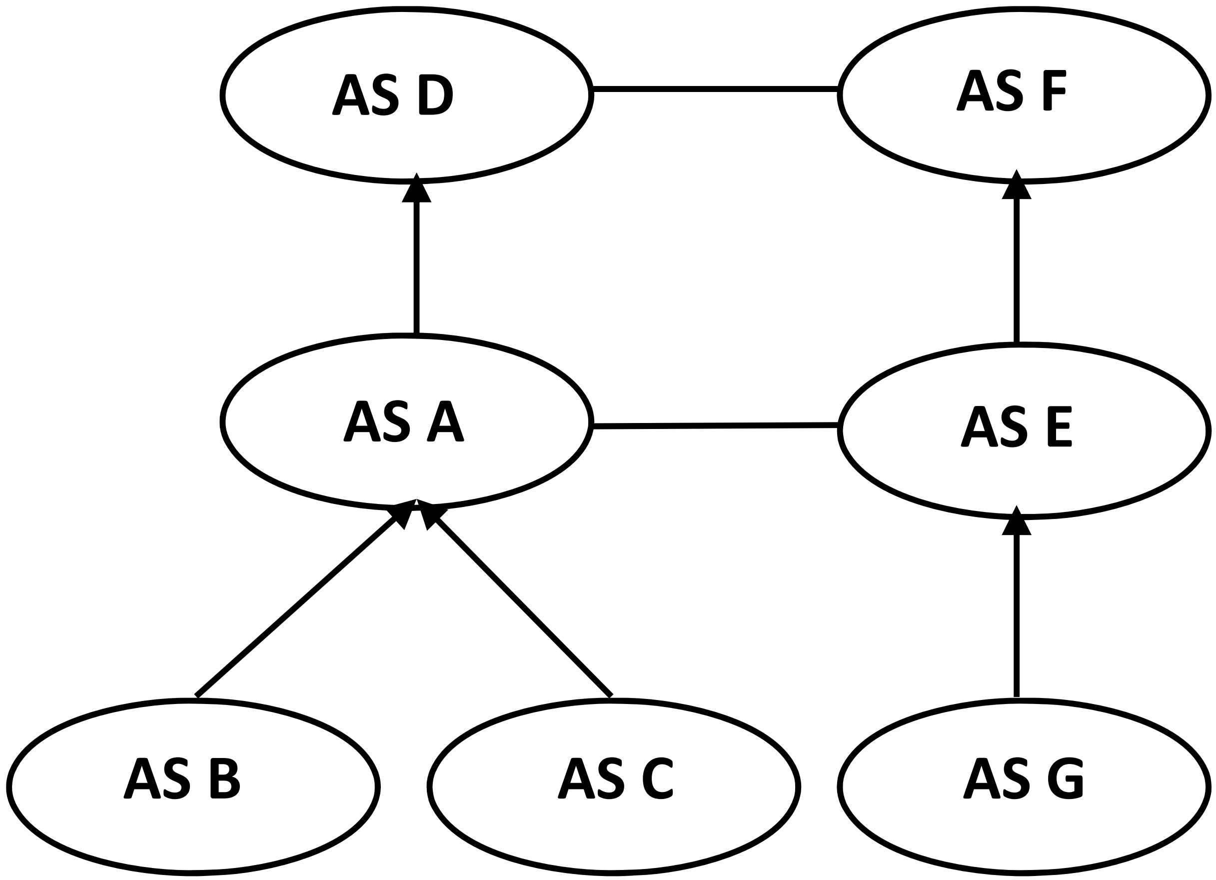

Monitor Placement Matters!

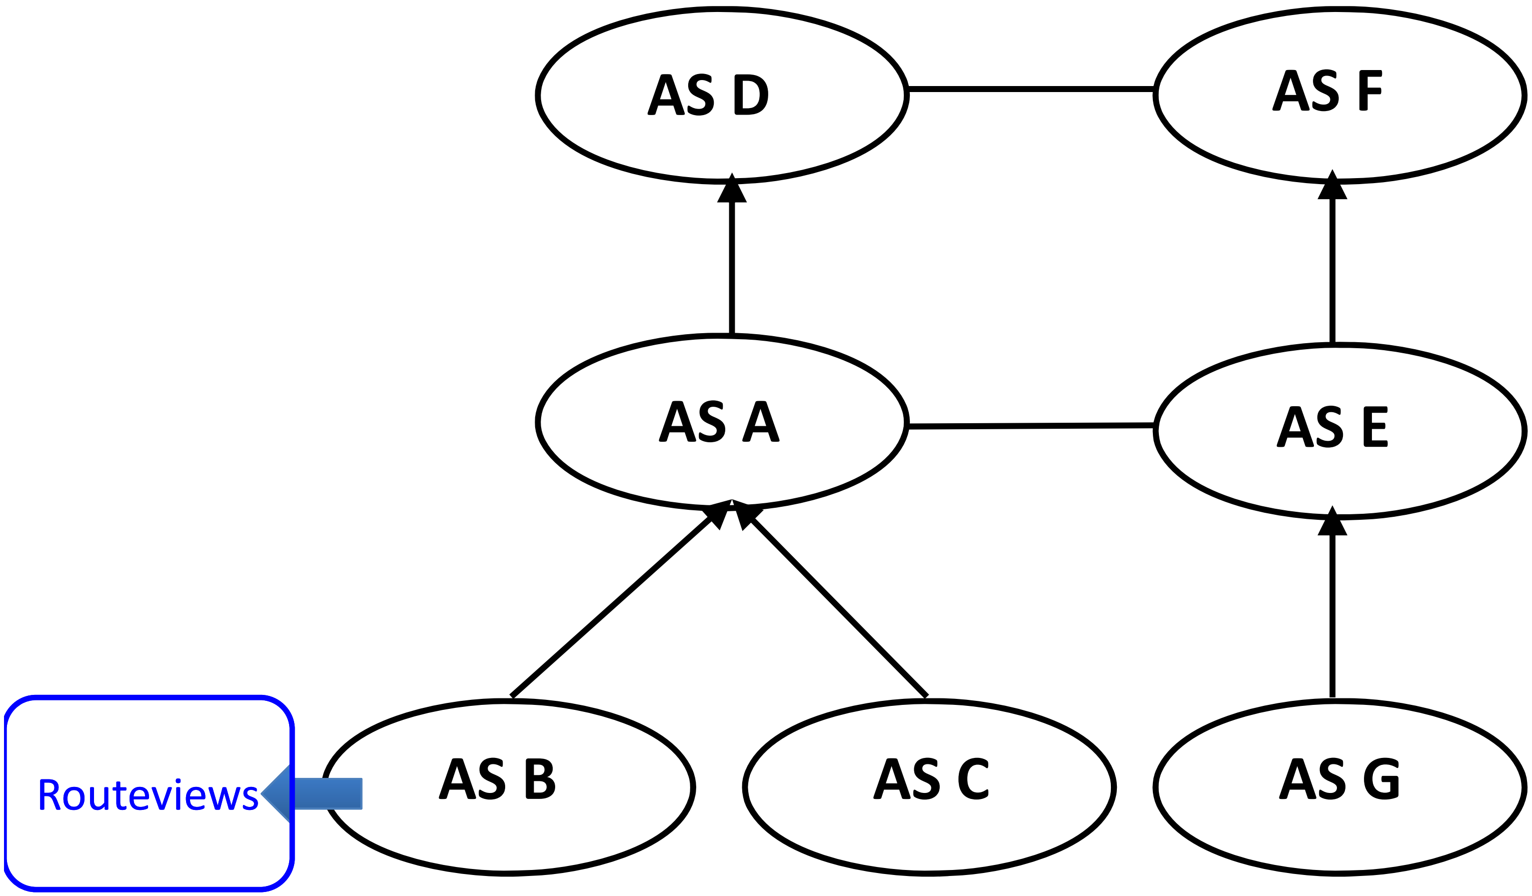

Monitor Placement Matters!

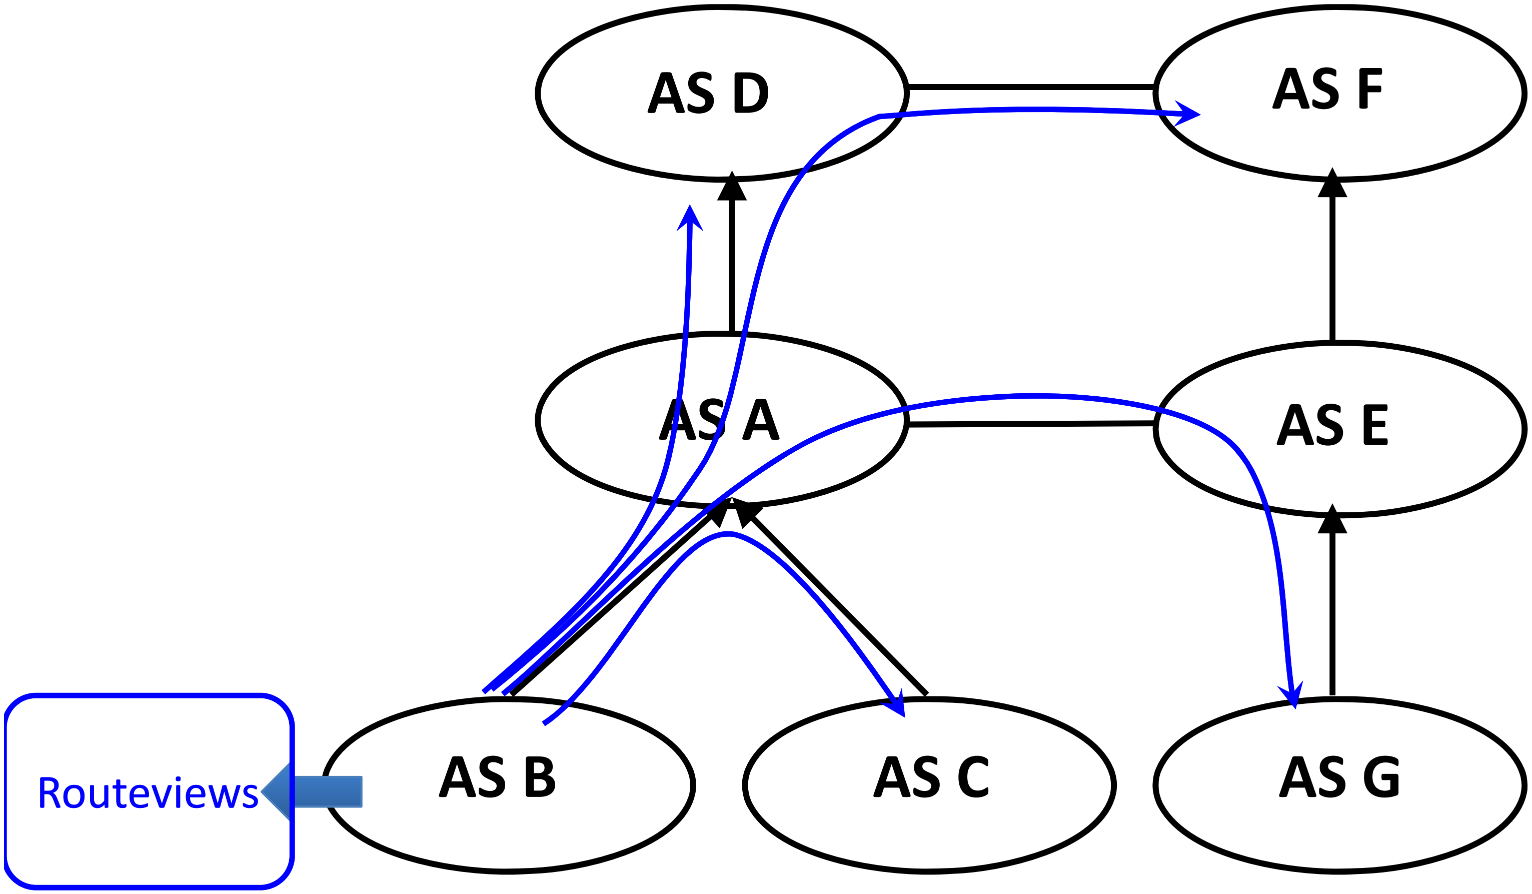

Monitor Placement Matters!

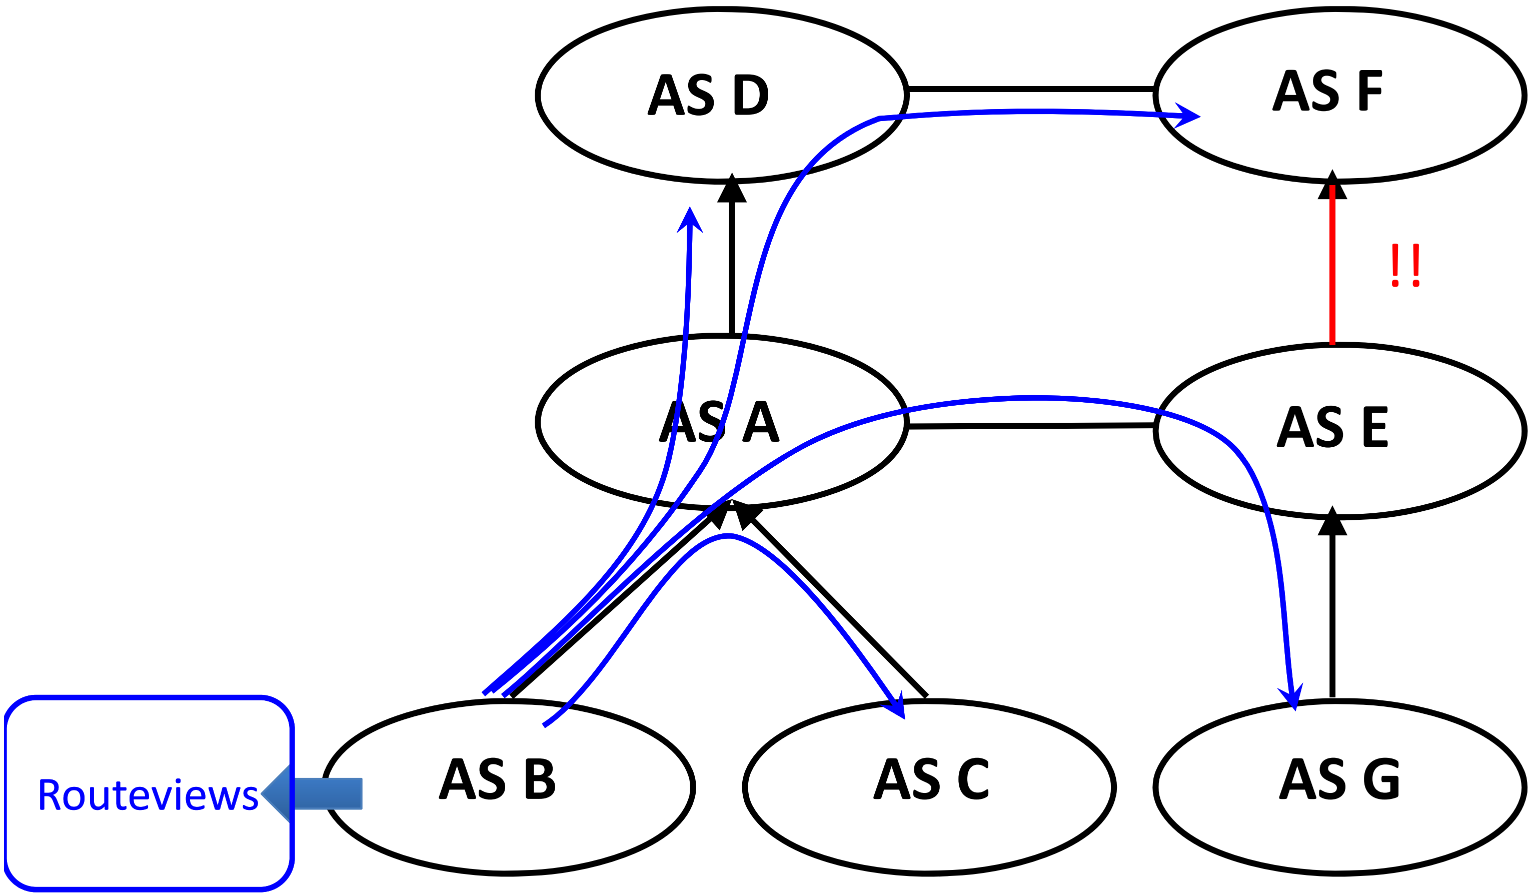

(In)visibility of AS topology

- How much of the topology do we miss by using a limited set of BGP vantage points?
- ASes: We see almost all ASes
- Customer-provider links: We see almost all customer-provider links
- Peering links: We likely miss a significant fraction of peering links in the Internet

How many peering links do we miss?

- To observe a peering link A-B, we need vantage points at A, or B, or at an AS hierarchically lower than A and B
- But we only have ~400 monitors, and many of them do not provide a full BGP feed
- Various estimates of the missing number of peering links: up to 60% missing*, up to 90% missing for tier-2 networks and Content Providers*

Chang, Willinger, “Difficulties Measuring the Internet’s AS-level Ecosystem”, ISS 2006

Oliveira, Pei, Willinger, Zhang, Zhang, “In Search of the Elusive Ground Truth: The Internet’s AS-level Connectivity Structure”, Sigmetrics 2008

IXPs

- Internet Exchange Points (IXPs) are a mostly ignored part of the interdomain connectivity ecosystem
- IXPs provide a shared fabric for “public” peering
  - A network can potentially connect to every other network at the IXP
  - Often no route filters: each network could potentially exchange traffic with every other network
- Currently >400 IXPs around the world, and their number and popularity is increasing

Anatomy of a large European IXP

- Ager et al.* measured connectivity and traffic at a large European IXP with ~400 members
- ~67% of all possible interdomain links at the IXP were found to exist!
- More peering links at this one IXP than were estimated to exist in the entire Internet
- Takeway: The public view is missing a large part of the interdomain connectivity picture!

Ager, Chatzis, Feldmann, Sarrar, Uhlig, Willinger, "Anatomy of a Large European IXP", Sigcomm 2012

Back to AS relationships

- We would really like to know the business relationship associated with an interdomain link
- Unfortunately, these are proprietary – networks are reluctant to give these away
- Recall that ASes are known to use the "valley-free, prefer-customer, prefer-peer" routing policy – policies manifest themselves in routes
- Leverage this assumption to infer business relationships based on observed BGP paths

AS relationship inference algorithms

- Gao*proposed the first (and most widely used) AS relationship inference algorithm
- Many refinements in subsequent years: Subramaniam et al., Zhang et al., Di Battista et al., Dimtropoulos et al., Gregori et al.
- Unfortunately limited validation of these algorithms; ground truth hard to obtain

Gao, "On Inferring Autonomous System Relationships in the Internet", IEEE/ACM Transactions on Networking, 2001

CAIDA's AS-rank

AS Ranking | Org Ranking | Information for a single AS | Information for a single Org | Background | Data Sources | Help   Org Ranking Help

The top Organizations ranked by customer cone size are displayed below.
For information about a specific Org, enter its name:

Look up an Org by name [Search]

Dataset: 2012.10.01 [Change dataset]

Table **shows** 10 of 43174 Orgs, sorted by number of ASes in customer cone [update view]

Org rank	Org name	Num. ASes	customer cone						AS degree	Org degree
			Number of			Percentages of all				
			ASes	IPv4 prefixes	IPv4 addresses	ASes	IPv4 Prefixes	IPv4 Addresses		
1	Level 3 Communications	18	31,937	336,953	1,835,879,372	75%	80%	71%	4,842	4,255
2	TeliaNet Global Network	5	16,753	173,505	708,992,583	39%	41%	27%	850	755
3	Cogent/PSI	3	15,924	169,875	694,172,644	37%	40%	27%	3,747	3,432
4	Tinet SpA	2	15,147	166,783	639,815,449	35%	39%	25%	956	856
5	NTT America, Inc.	7	13,249	156,544	674,237,239	31%	37%	26%	884	770
6	TATA Communications	6	9,028	140,057	561,900,698	21%	33%	21%	943	856
7	TELECOM ITALIA SPARKLE S.p.A.	3	9,001	122,246	460,017,217	21%	29%	18%	343	309
8	MCI Communications Services, Inc. d/b/a Verizon Business	44	8,857	139,353	872,872,874	20%	33%	34%	2,219	1,940
9	Sprint	14	7,101	119,608	831,151,193	16%	28%	32%	1,000	884
10	Qwest Communications Company, LLC	10	6,660	112,307	651,247,290	15%	26%	25%	1,844	1,659

data sources

geolocation	database	2012.06.25	netacuity	
organization	whois	0000.00.00	JPNIC, KRNIC, LACNIC	
		2012.06.29	AFRINIC, APNIC, ARIN, LACNIC, RIPE	
topology	BGP	2012.10.01, 2012.10.02, 2012.10.03, 2012.10.04, 2012.10.05	ripe	rrc00, rrc01, rrc03, rrc04, rrc05, rrc06, rrc07, rrc10, rrc11, rrc12, rrc13, rrc14, rrc15
			routeviews	eqix, isc, jinx, kixp, linx, routeviews2, routeviews6, saoppaulo, sydney, telxatl, wide

Support for this work is provided by the U.S. Department of Homeland Security's Science and Technology Directorate (Project N66001-08-C-2029), the National Science Foundation Internet Laboratory for Empirical Network Science (Project CNS-0958547), and Cisco's University Research Program.

CAIDA's AS-rank

- Luckie et al.* developed a new AS-relationship inference algorithm
- Solicited ground truth via a "corrections" interface
- Assembled largest collection of AS-relationship ground-truth to date
- Current algorithm is ~99% accurate for both customer-provider and peering links

Luckie, Huffaker, Dhamdhere, Claffy, "Inferring AS Relationships and Customer Cones", in preparation

Outline

- AS topology as a network of business relationships
- Measuring the AS-level Internet
- Structural properties of the AS-level Internet
- Evolution of the AS-level Internet
- Modeling the structure and evolution of the AS-level Internet

Structural Properties

- Rich area of research over the last decade+ starting with the discovery of power laws by Faloutsos et al.*
- Also a lot of controversy: is it a power-law or not?
- Eventual agreement: degree distribution is highly skewed ☺
- How do IXPs change the degree distribution?

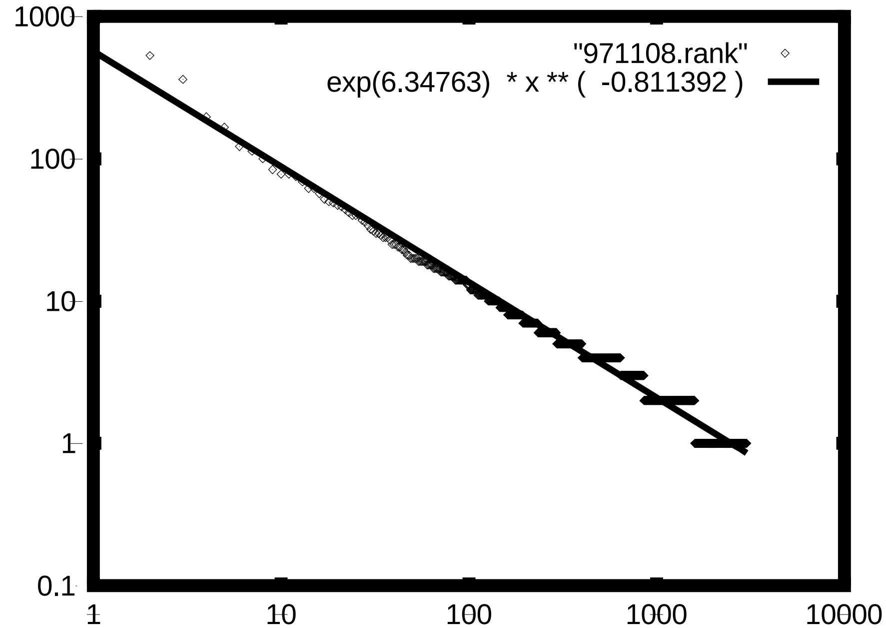

(a) Int-11-97

Faloutsos, Faloutsos, Faloutsos, "On Power Law Relationships in the Internet Topology", ACM Sigcomm 1999

Structural Properties

- Small-world properties: Measured AS graphs show strong clustering and almost constant average path lengths
- Basic topological properties such as degree distribution and clustering have been invariant over time

The Traditional Hierarchy

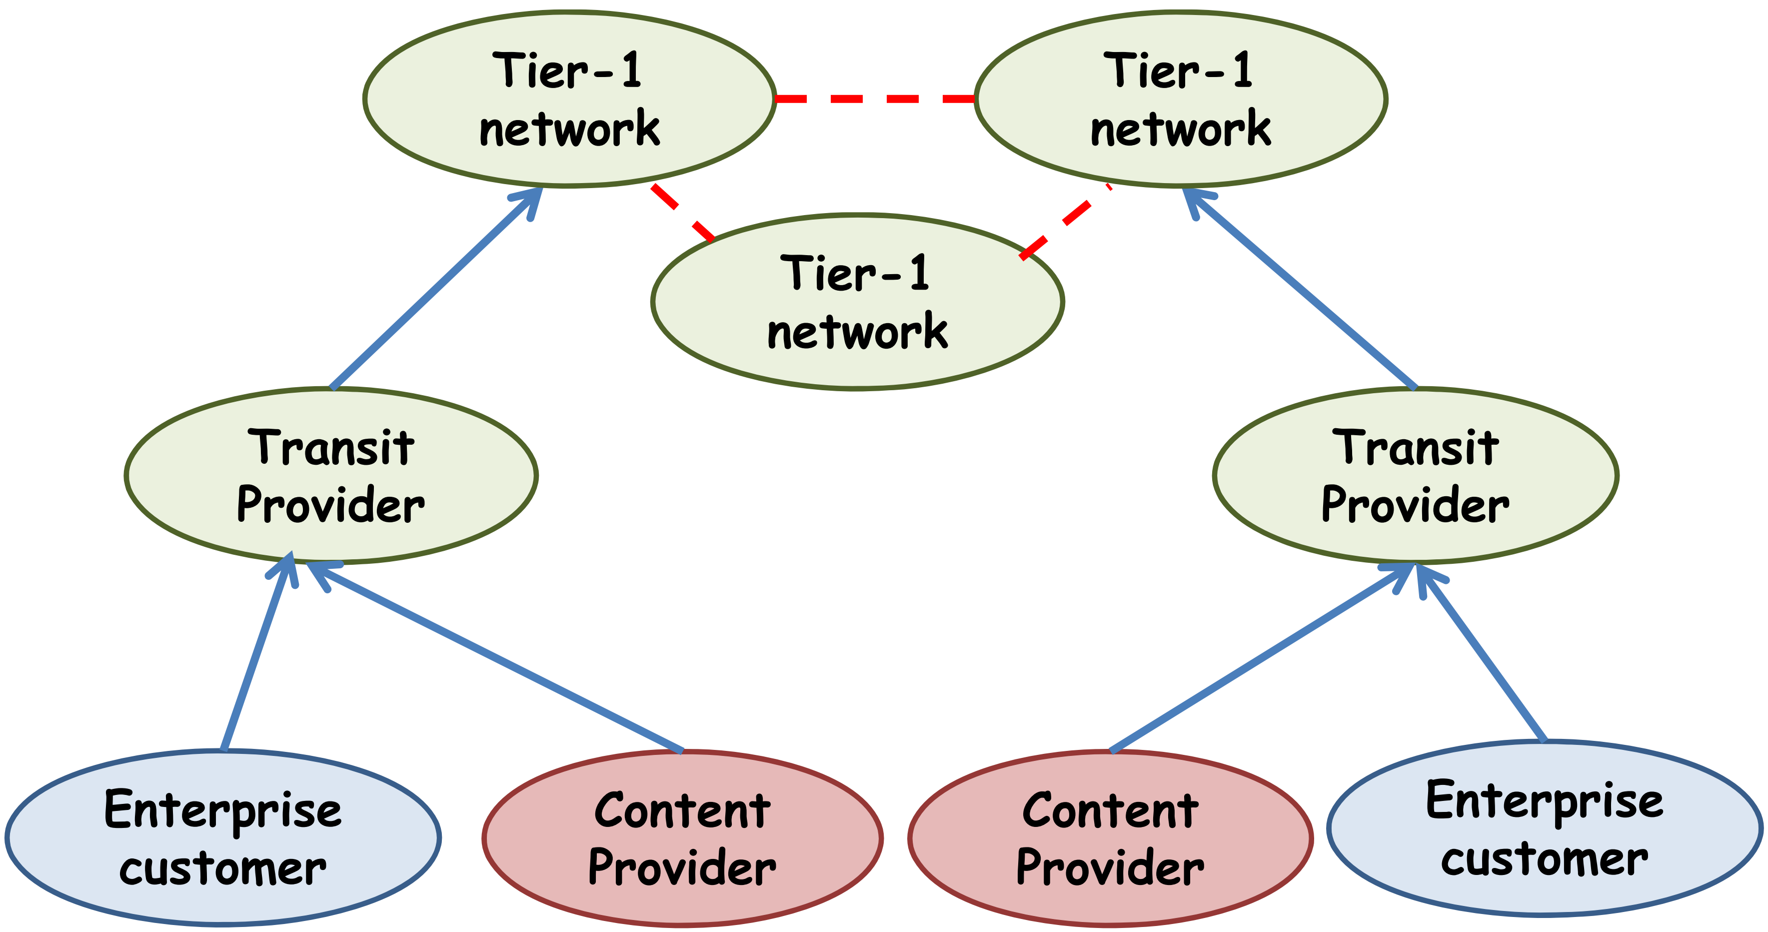

The Traditional Hierarchy

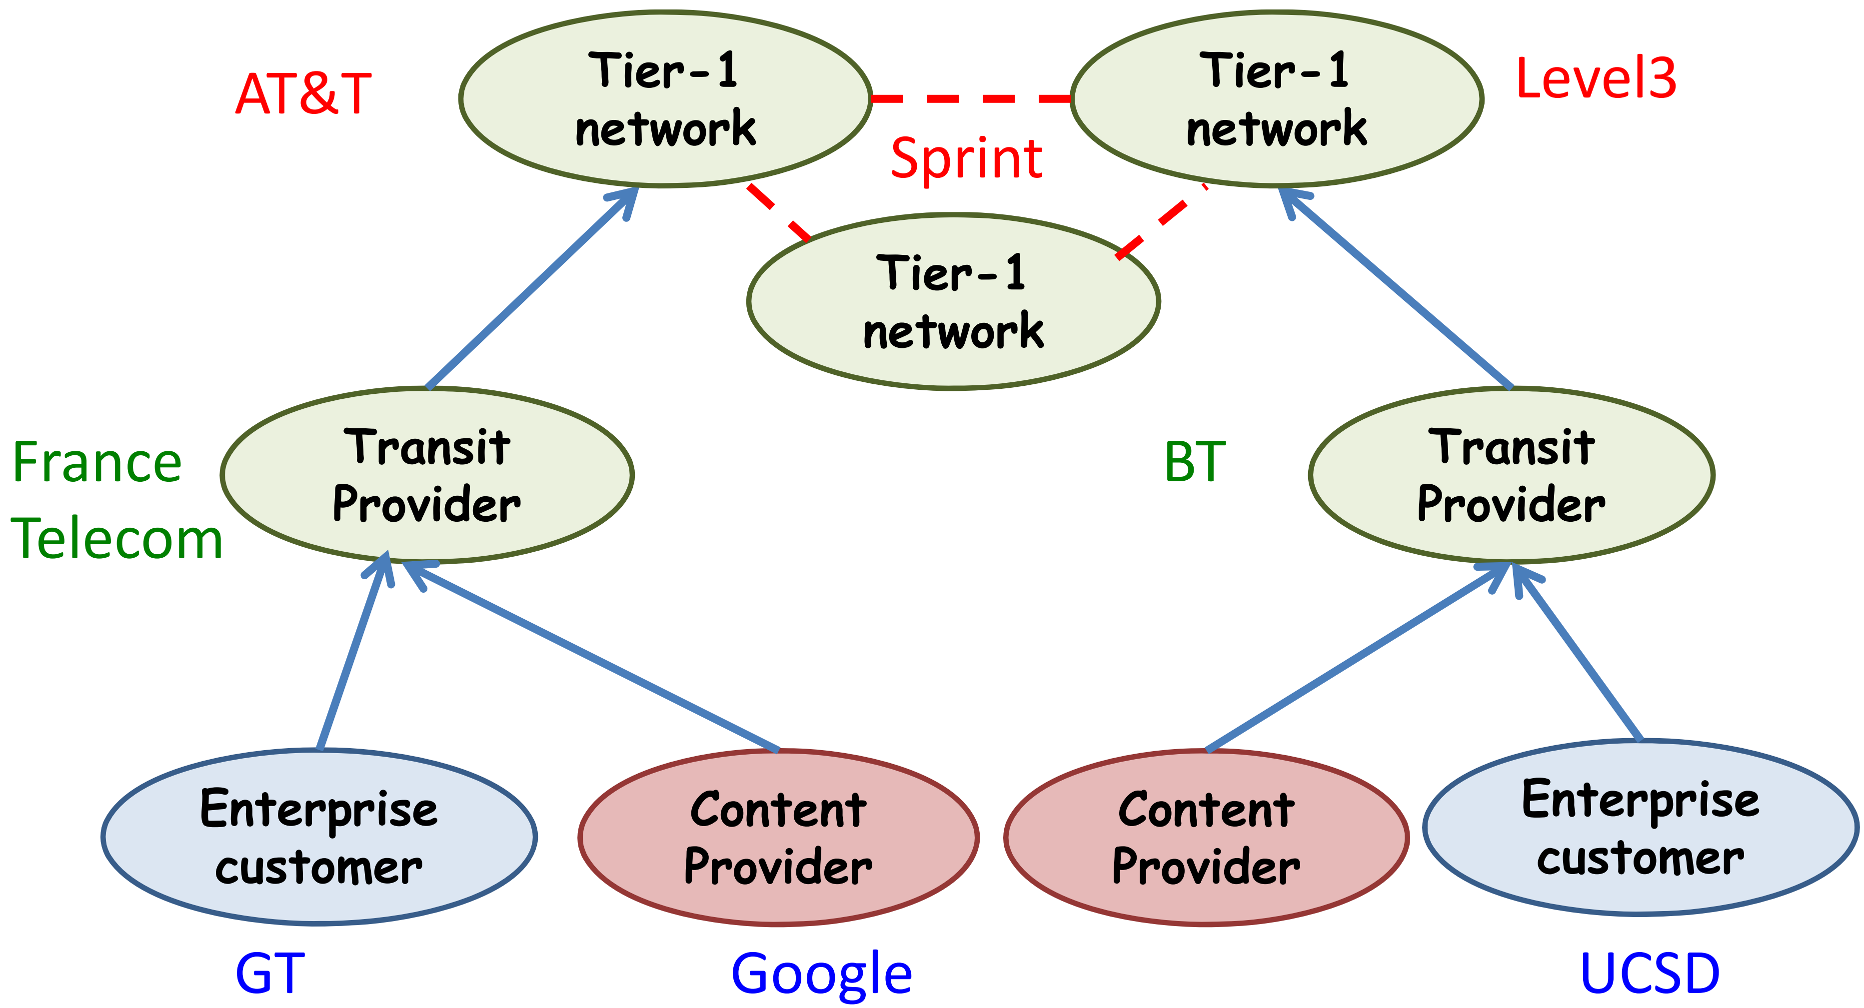

The Traditional Hierarchy

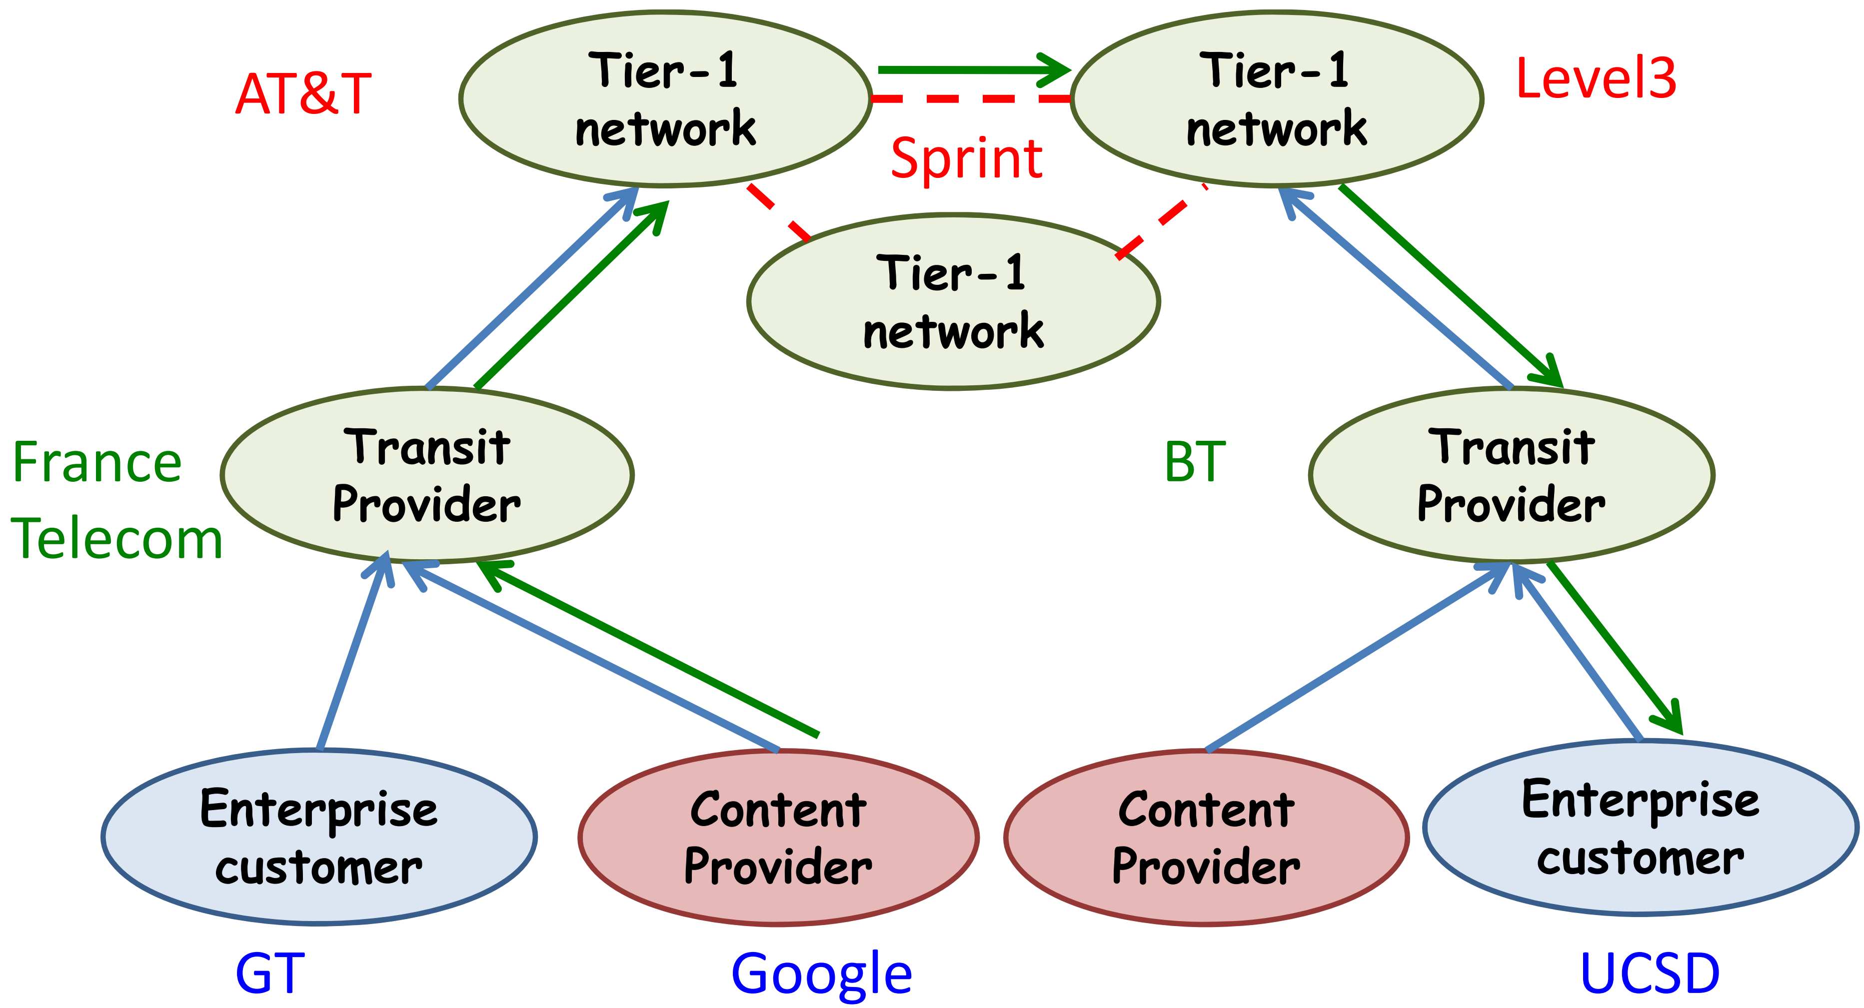

The Traditional Hierarchy

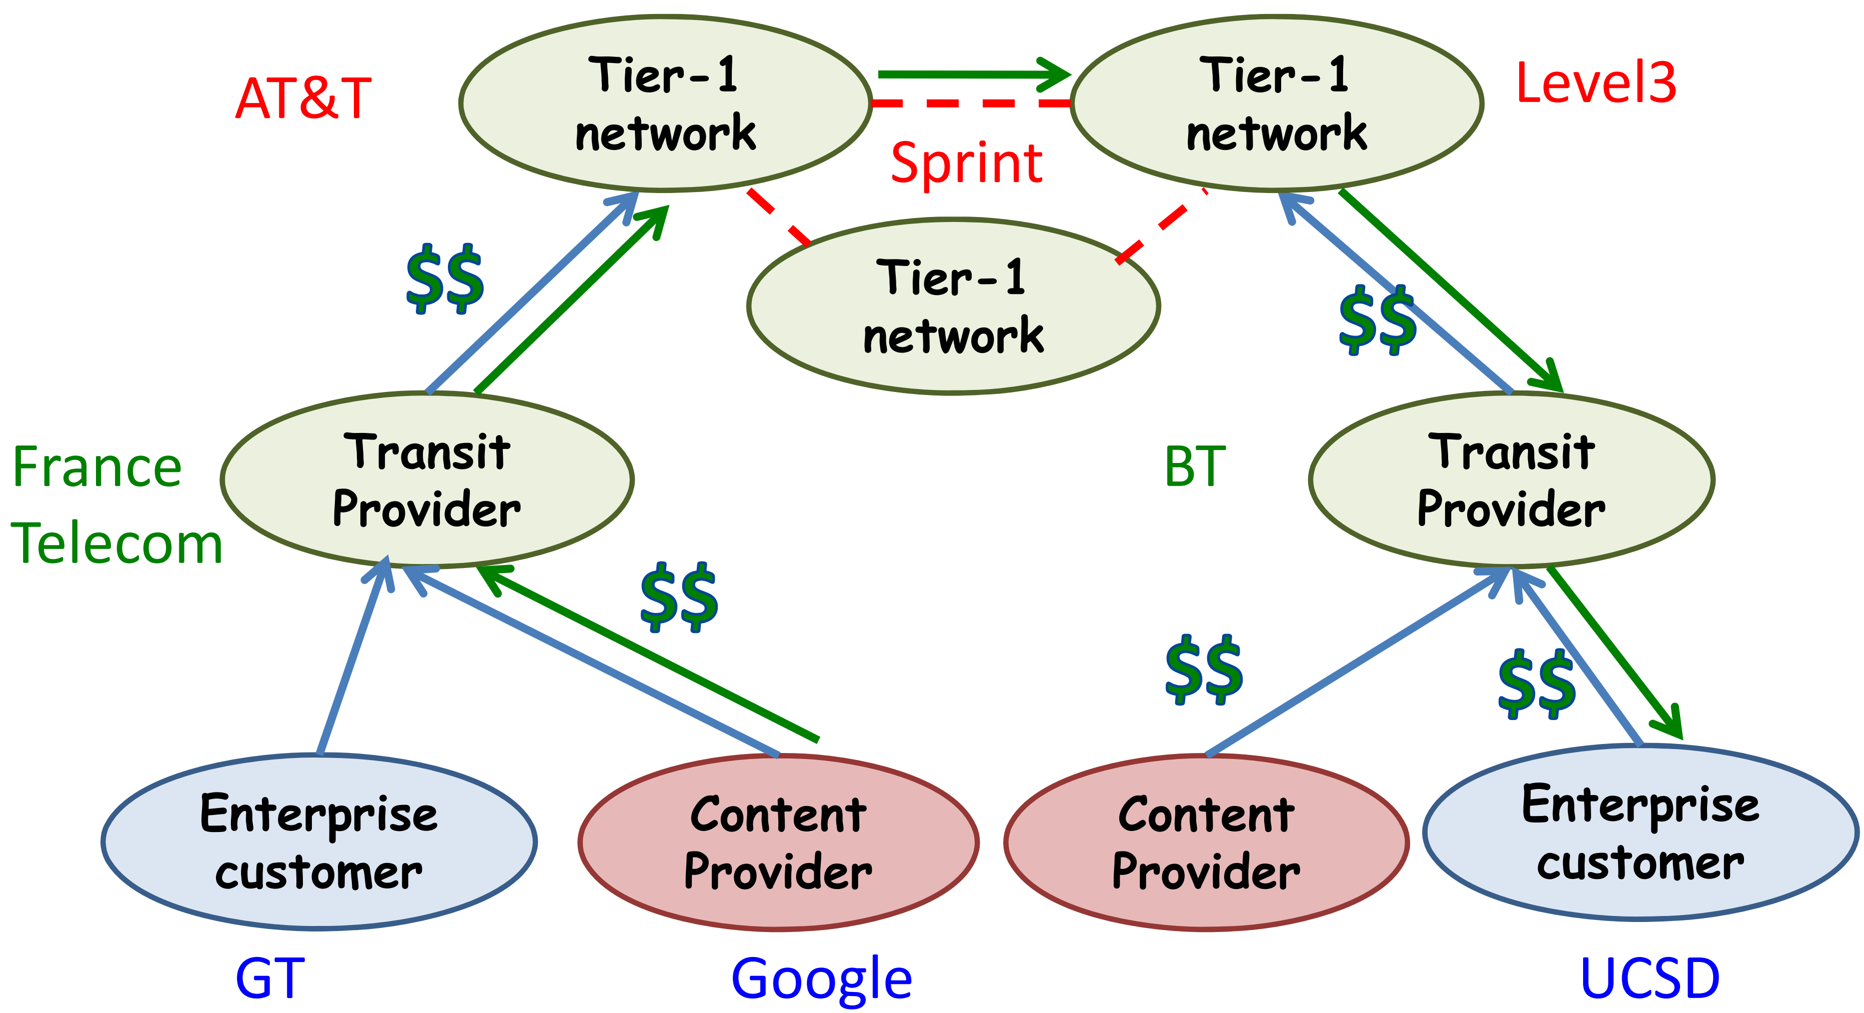

Tier-1 clique over time [Luckie et al.]*

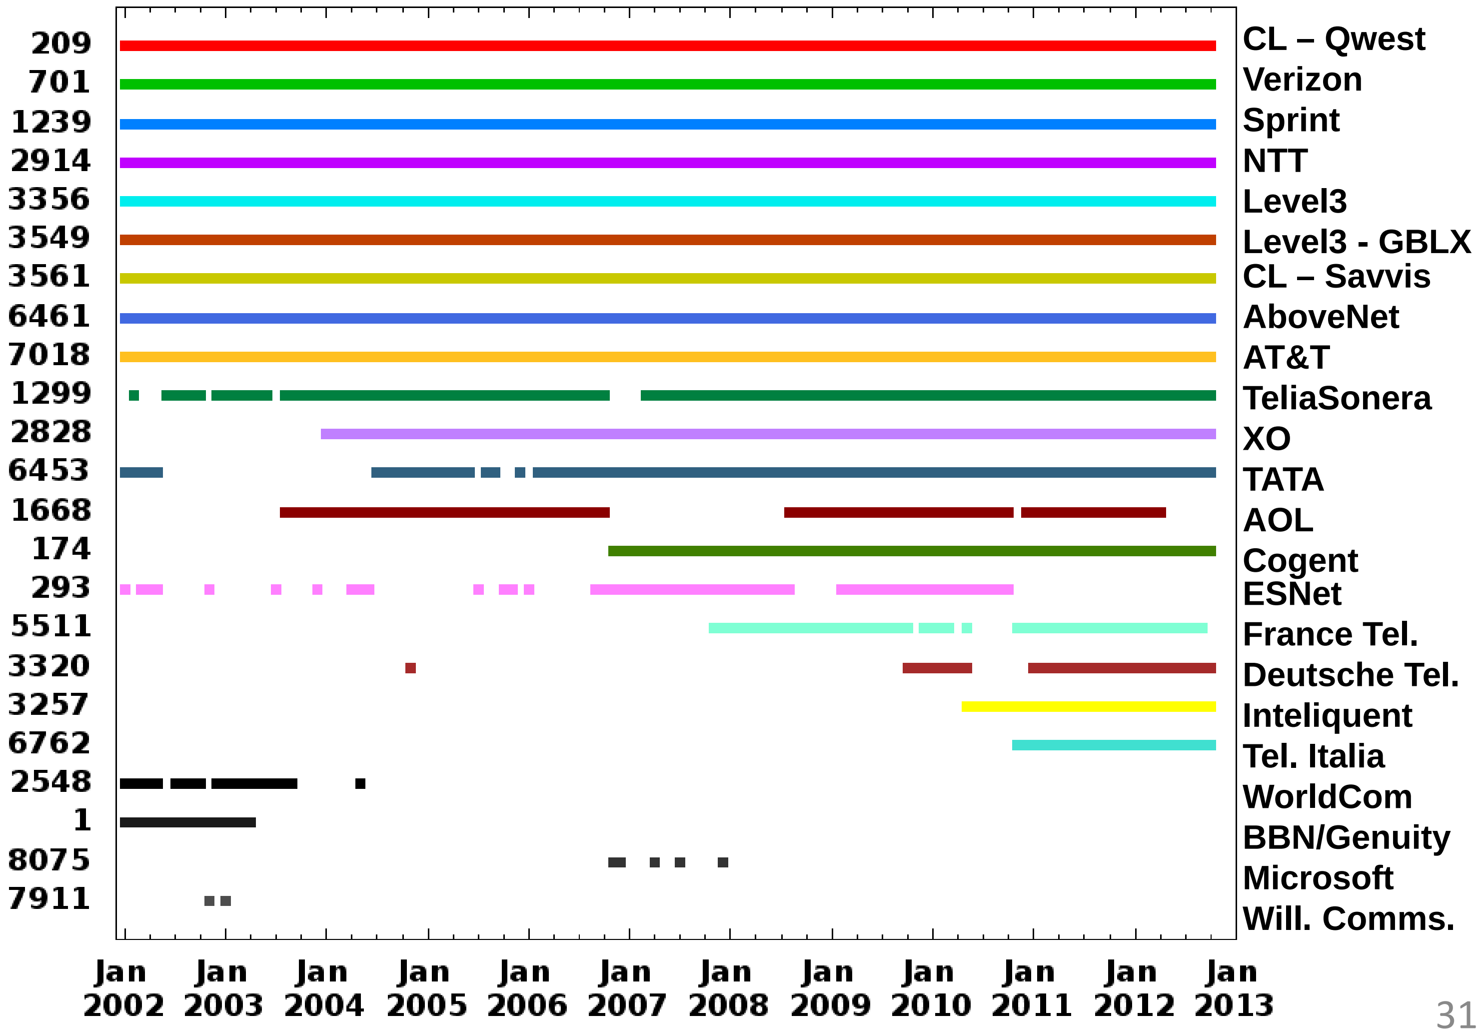

Luckie, Huffaker, Dhamdhere, Claffy, “Inferring AS Relationships and Customer Cones”, in preparation

Outline

- AS topology as a network of business relationships
- Measuring the AS-level Internet
- Structural properties of the AS-level Internet
- Evolution of the AS-level Internet
- Modeling the structure and evolution of the AS-level Internet

Growth Trends

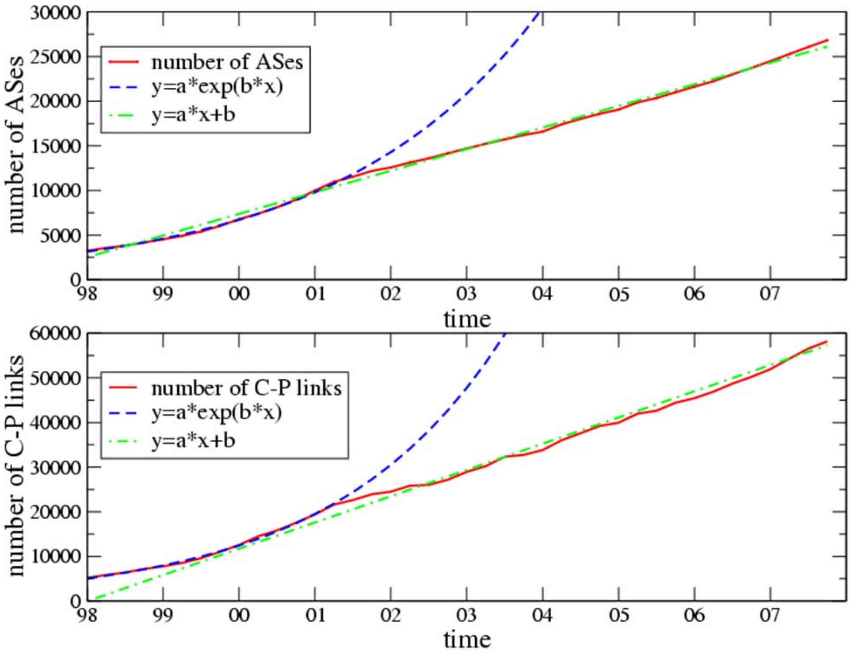

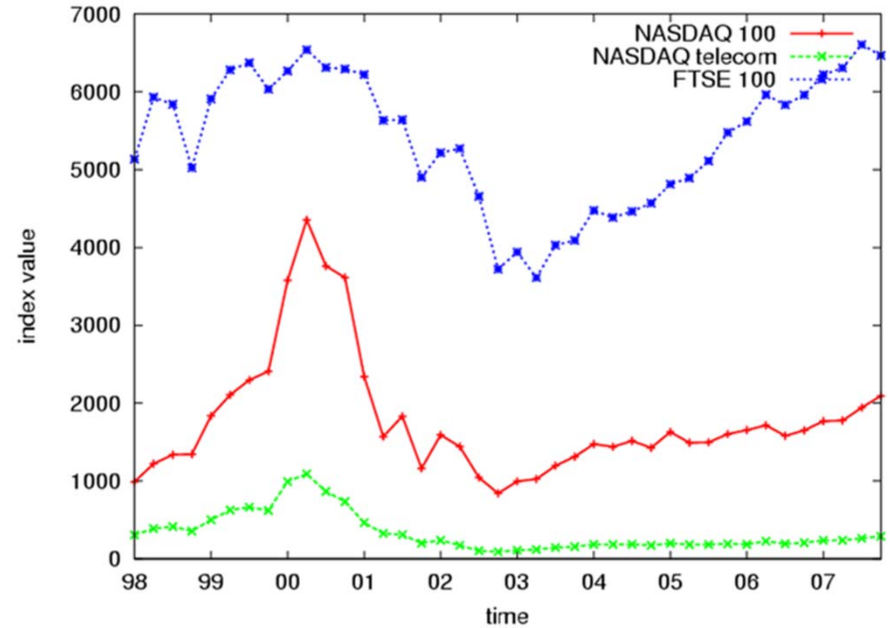

- Number of CP links and ASes showed initial exponential growth until mid-2001 followed by linear growth until today
- Change in trajectory coincided with stock market crash in North America in mid-2001

Dhamdhere, Dovrolis, “Twelve Years in the Evolution of the Internet Ecosystem”, IEEE/ACM Transactions on Networking, 2011.

Classification of ASes based on business function

- Four AS types:
  - Enterprise customers (EC)
  - Small Transit Providers (STP)
  - Large Transit Providers (LTP)
  - Content, Access and Hosting Providers (CAHP)
- Based on customer and peer degrees
- Classification based on decision-trees
  - 80-85% accurate

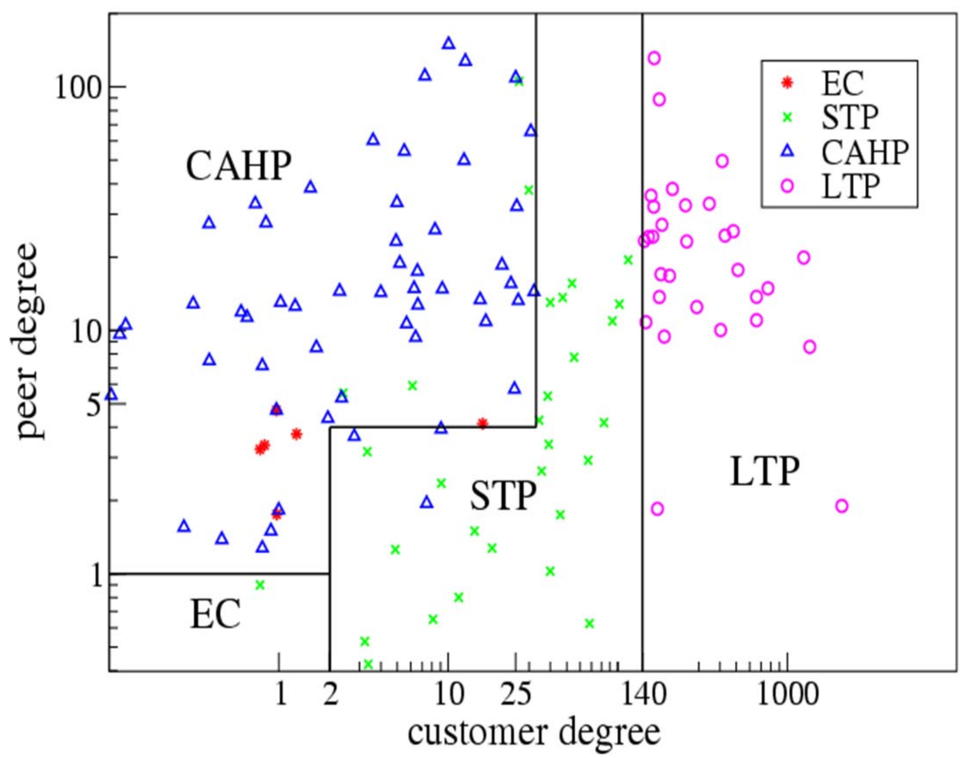

Dhamdhere, Dovrolis, “Twelve Years in the Evolution of the Internet Ecosystem”, IEEE/ACM Transactions on Networking, 2011.

Evolution of AS types

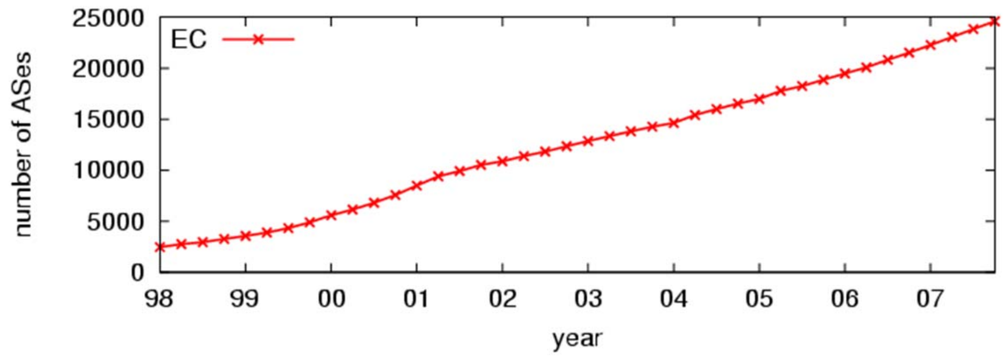

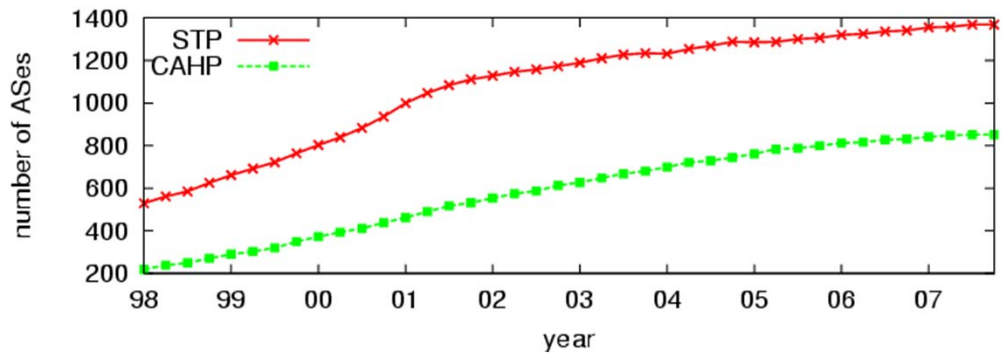

- Slow growth of STPs (30% increase since 2001)
- EC population produces most growth (150% increase since 2001)

Dhamdhere, Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem", IEEE/ACM Transactions on Networking, 2011.

Path lengths stay constant

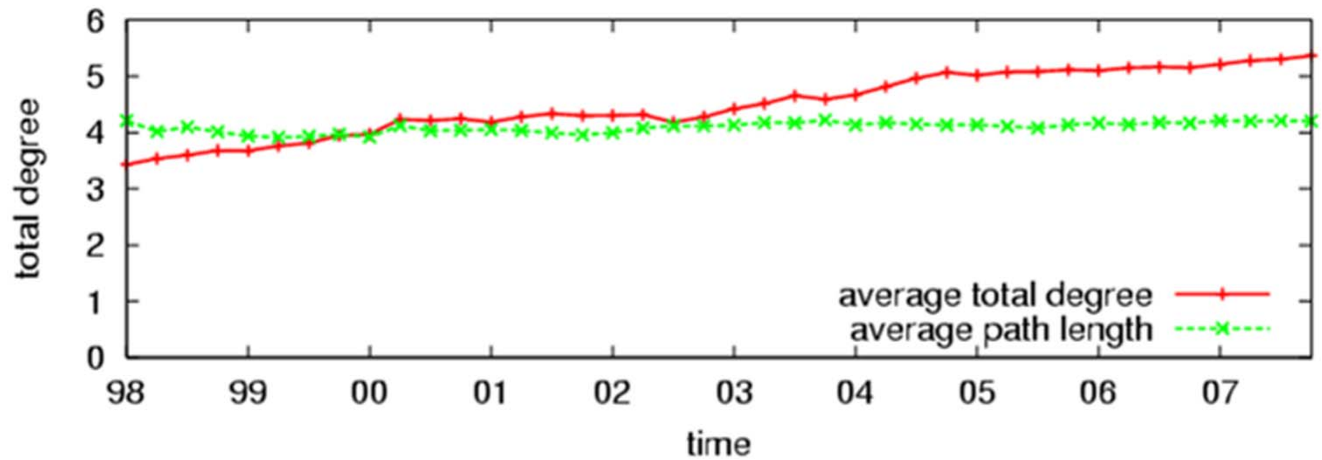

- Number of ASes has grown from 5000 in 1998 to 42000 in 2012
- Average path length constant at ~4 AS hops
- Densification?

Dhamdhere, Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem", IEEE/ACM Transactions on Networking, 2011.

Where does densification happen?

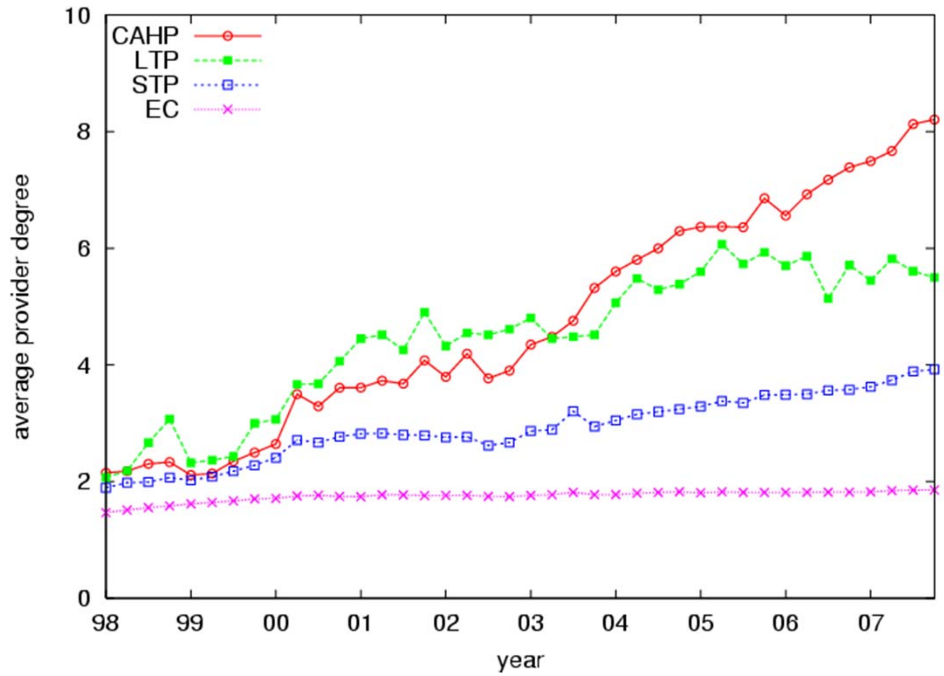

- CAHPs have increased their multihoming degree significantly (avg. 8 providers for CAHPs today)
- Multihoming degree of ECs almost constant (avg. < 2)
- Densification of the Internet occurs at the core

Dhamdhere, Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem", IEEE/ACM Transactions on Networking, 2011.

Flattening (topology)

- Gill et al.* measured geographical expansion by content providers
- Major CPs are increasingly building out their own networks
- Routes increasingly bypass tier-1 networks

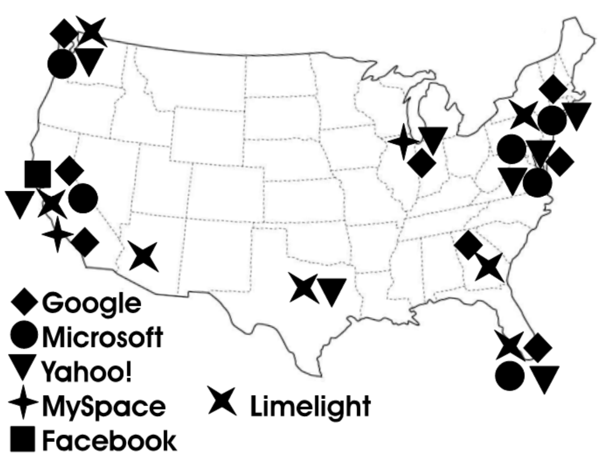

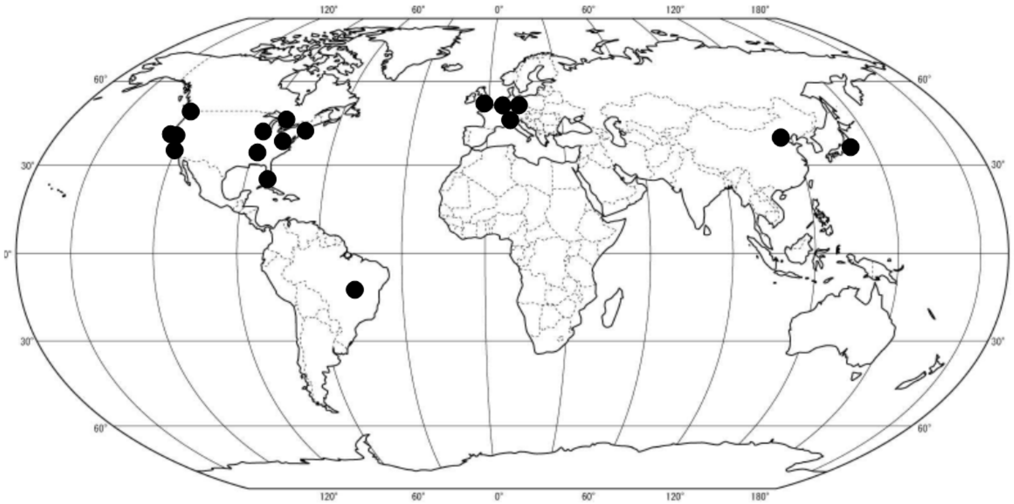

Gill, Arlitt, Li, Mahanti, "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles, or Contrived Collapse", PAM, 2008.

Flattening (traffic)

- Arbor networks measurements of interdomain traffic*
- Traffic consolidates: a few large “supergiants”
- Traffic bypasses tier-1 networks; flows directly on peering links

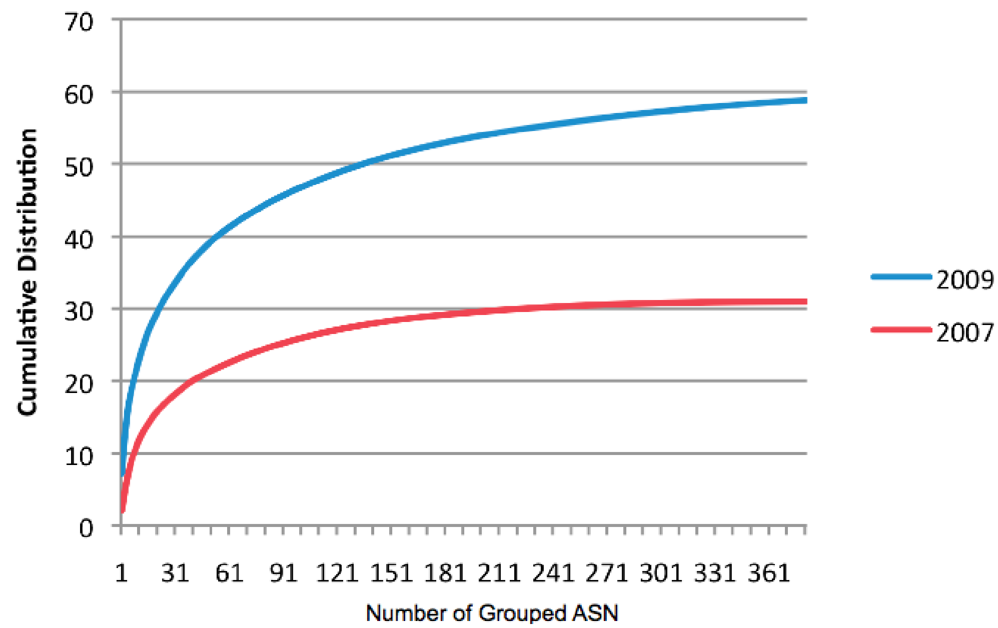

Labovitz, Iekel-Johnson, McPherson, Oberheide, Jahanian, “Internet Inter-domain Traffic”, ACM Sigcomm, 2010.

“Open” Peering

- peeringDB: An online database where networks volunteer information about peering
- Lodhi et al.* measured peering policies advertised in peeringDB
- A majority of networks advertise an “open” peering policy --- willing to peer with anyone!
- Contributes further to flattening of the topology?

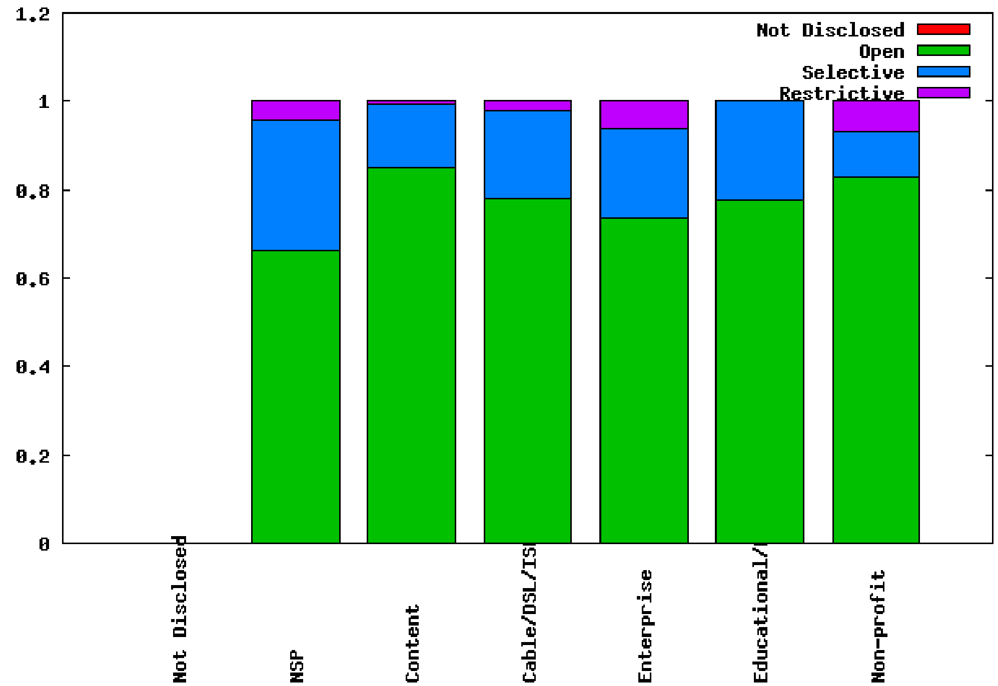

*Lodhi, Dhamdhere, Dovrolis, “Analysis of Peering Strategy Adoption by Transit Providers in the Internet”, NetEcon 2012.

Outline

- AS topology as a network of business relationships
- Measuring the AS-level Internet
- Structural properties of the AS-level Internet
- Evolution of the AS-level Internet
- Modeling the structure and evolution of the AS-level Internet

Why care for topology models?

- Simulation: For many applications, we cannot simulate the Internet at-scale
  - Need to scale down (or scale up) topologies

- Evolution: We'd like to know how the topology evolves, and what it might be heading towards

- Prediction: We'd like to predict traffic flows and (more interestingly) economic flows: who makes money? Who doesn't?

Top-down models

- Basic idea: Start will well-known properties of Internet topology, produce a model that reproduces those properties
- Example properties: degree distribution, clustering, diameter, betweenness, hierarchy – mostly graph-theoretic metrics
- Pros: perfect for producing synthetic topologies that match certain statistics of the measured topology, small number of parameters

Preferential Attachment

- Barabasi and Albert*: Simple "rich get richer" model that produces power-law degree distributions
- Several follow-up models: Better match degree distribution, as well as other properties, e.g., clustering
- Cons: The data that these models use as input can be incomplete and messy, these models are not necessarily predictive

Barabasi, Albert, "Emergence of Scaling in Random Networks", Science, 1999.

Bottom-up models

- Fundamentally different approach to modeling topology structure and evolution
- Model the incentives and actions of individual actors, let global properties "emerge"
  - E.g., network design incentives*, economic incentives
- Pros: can be designed to capture operational realities, can be used to study dynamics
- Cons: Difficult to parameterize, computationally expensive to simulate

Fabrikant, Koutsoupias, Papadimitriou, "Heuristically Optimized Trade-offs: A New Paradigm for Power Laws in the Internet", ICALP, 2002.

ITER*

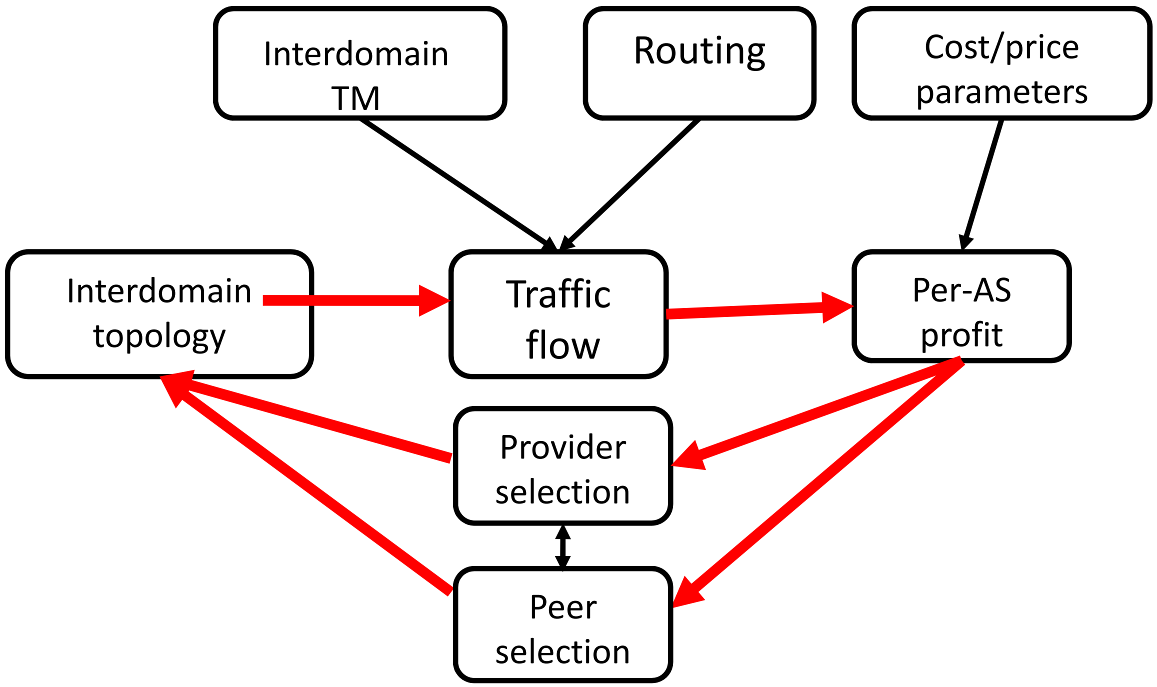

- Agent based computational model
- Model the complex feedback loops between topology, traffic, and economics
- Compute an equilibrium: no network has an incentive to change connectivity

Dhamdhere, Dovrolis, "The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh", ACM CoNEXT, 2010.

Using ITER to model flattening

- The Old Internet (late 90s – 2007)
- Content providers generated small fraction of total traffic
- Content providers were mostly local
- Peering was restrictive

- The New Internet (2007 onwards)
- Content providers generate large fraction of total traffic
- Content providers are present everywhere
- Peering is more open

ITER: Traffic Transiting Transit Providers

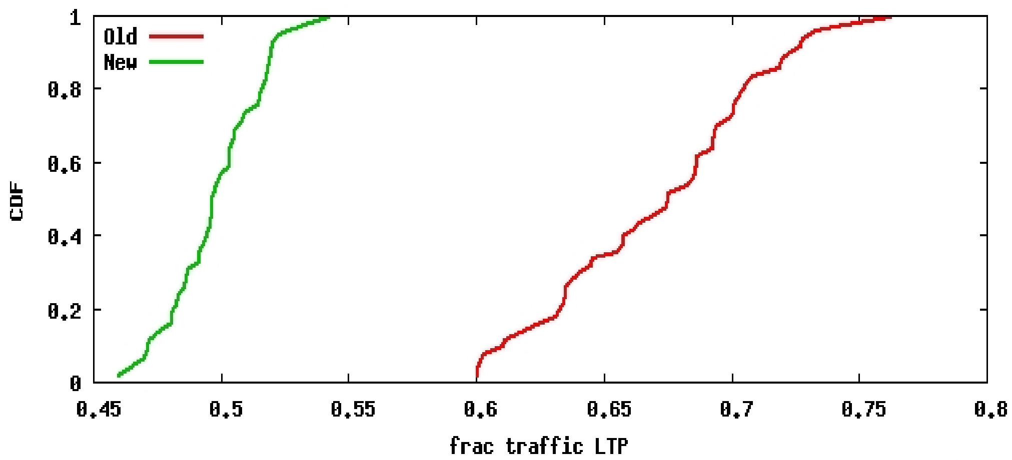

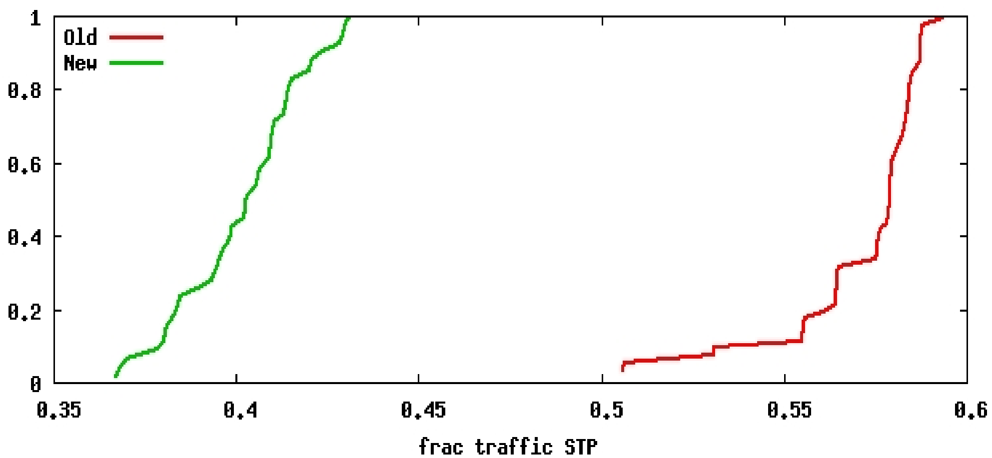

- Traffic bypasses transit providers
- More traffic flows directly on peering links
- Implication: Transit providers lose money!
- Content providers get richer

Back to the Real World

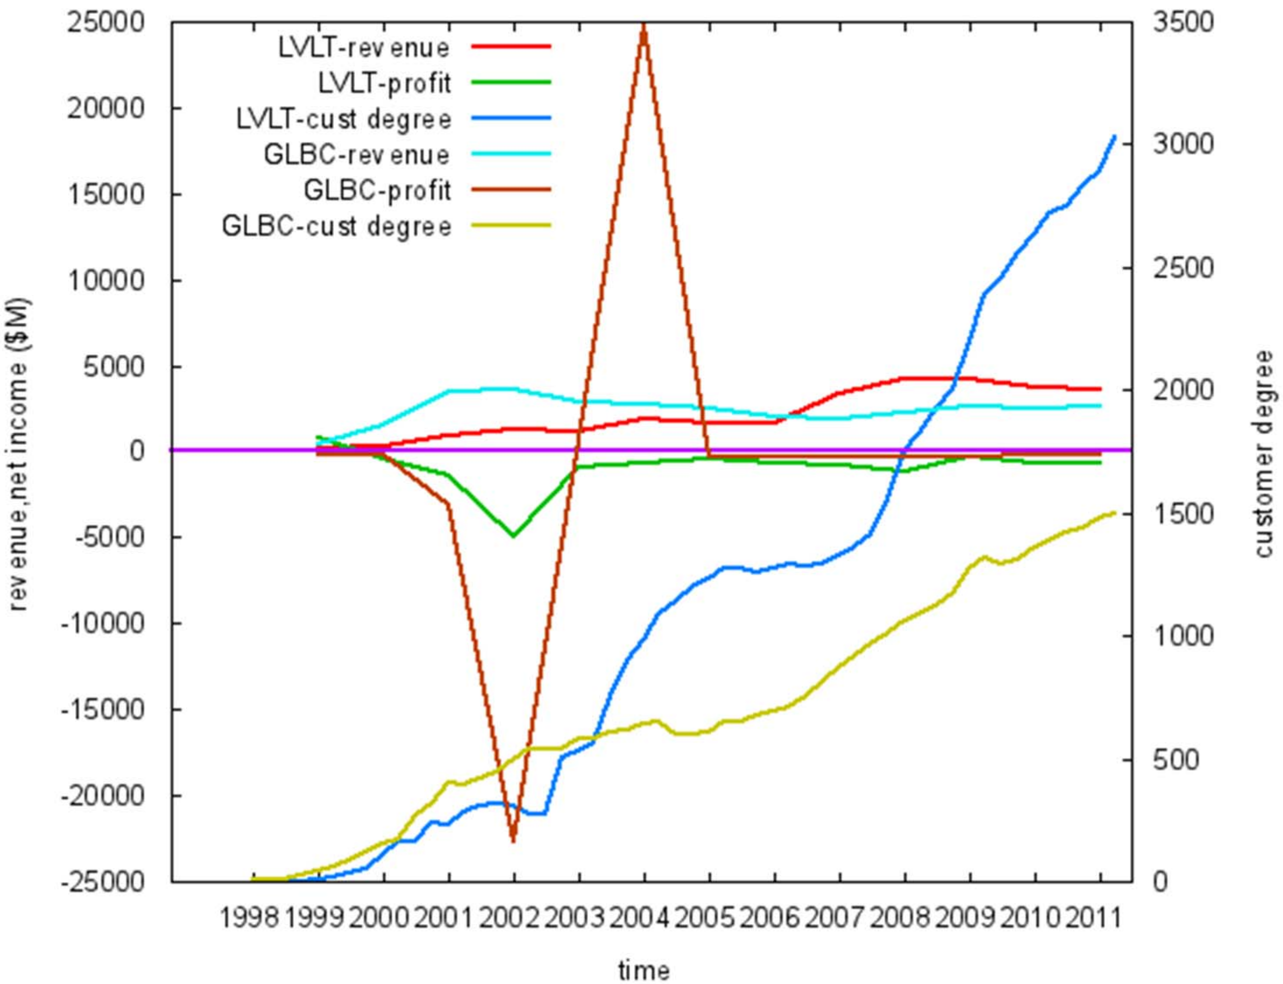

Back to the Real World

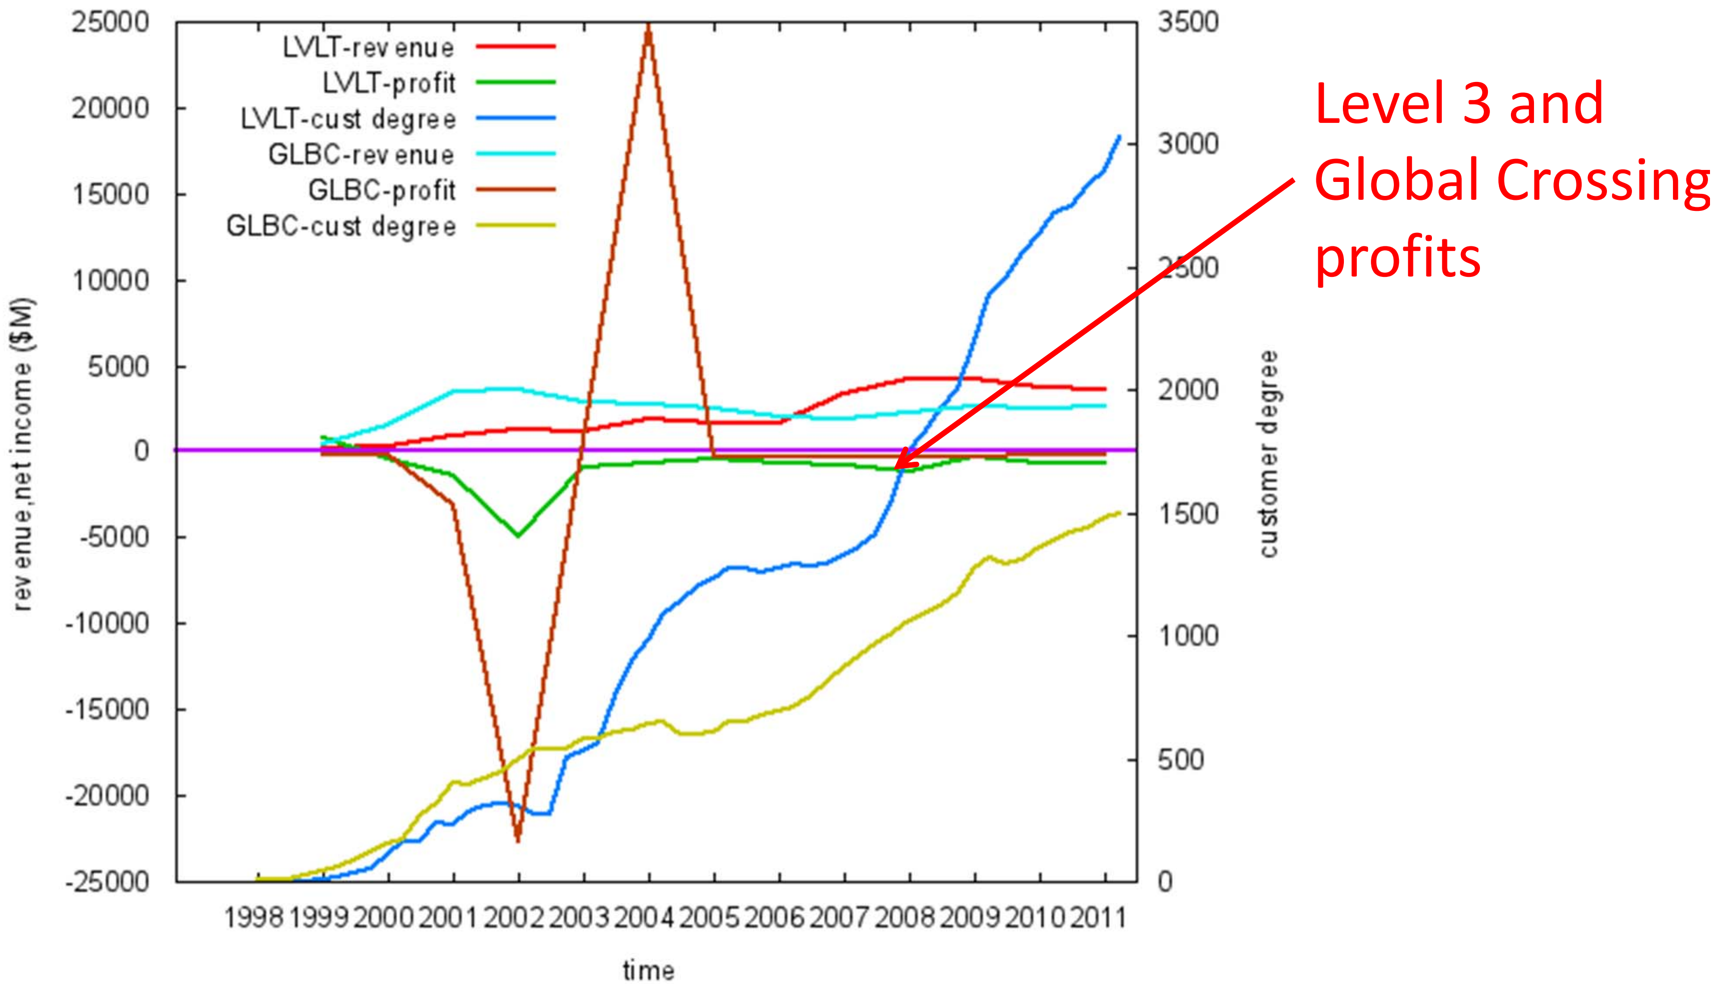

Back to the Real World

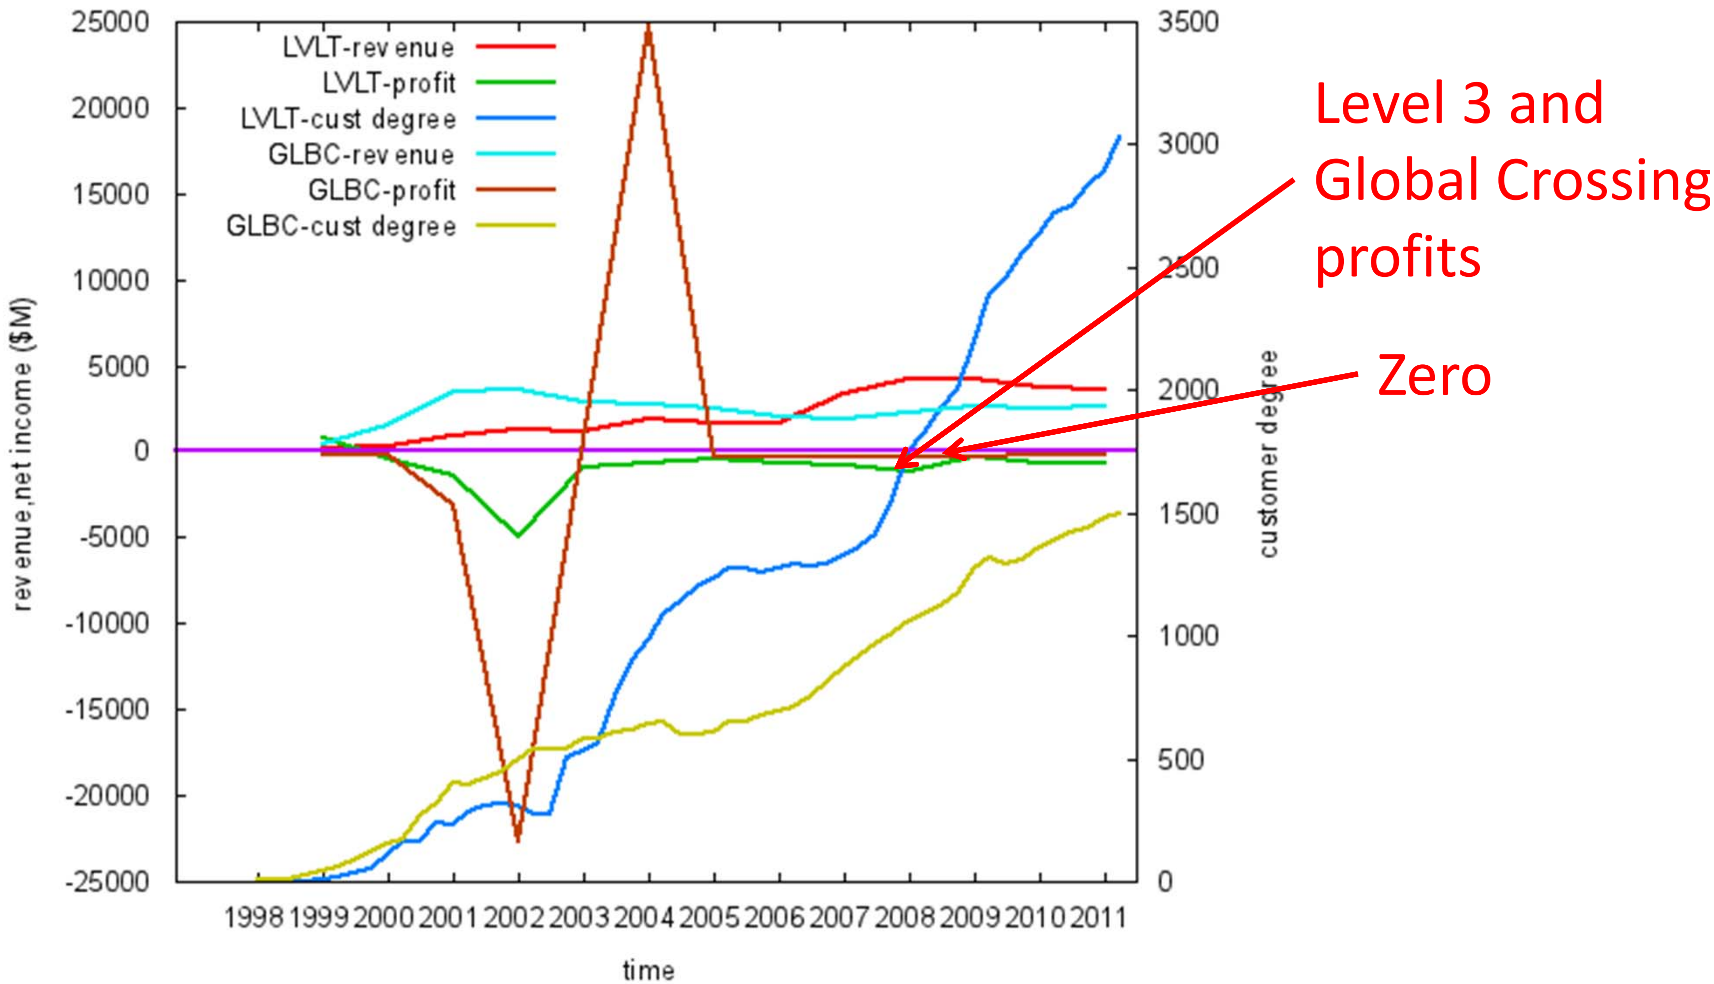

GENESIS*

- Agent based interdomain network formation model
- Incorporates many real-world constraints in provider/peer selection
- Focuses on strategy selection by ASes
- Objective of a network: Maximize economic fitness
- Choose the peering strategy that maximizes fitness

Lodhi, Dhamdhere, Dovrolis, “GENESIS: An Agent-based Model of Interdomain Netowork Formation, Traffic Flow, and Economics”, IEEE Infocom, 2012.

Using GENESIS to study strategy adoption*

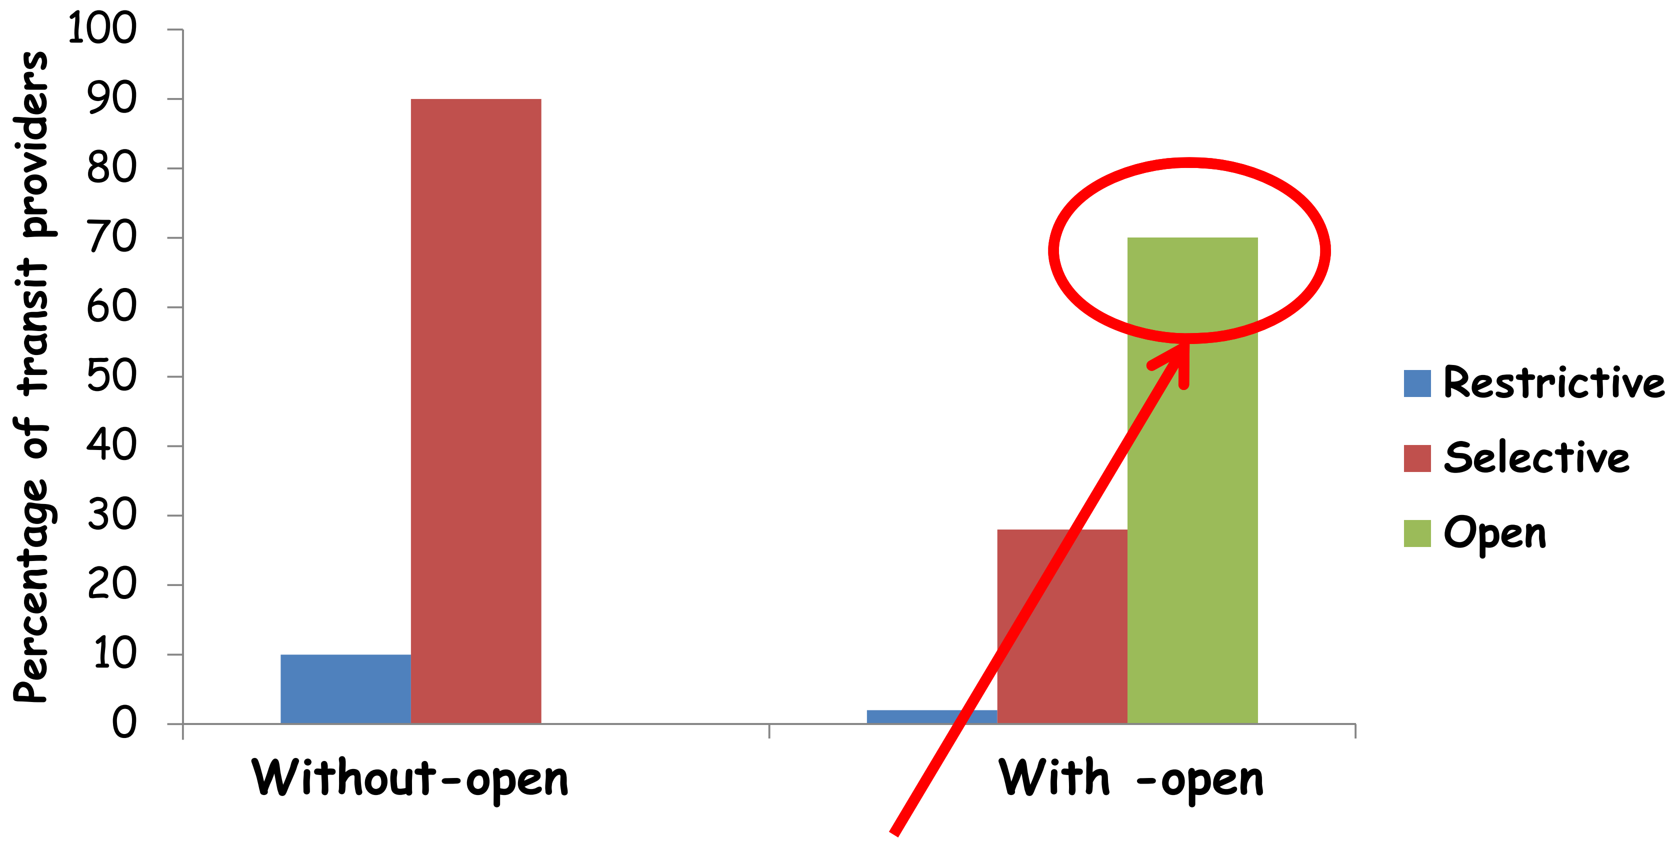

Matches very well with data from peeringDB

Lodhi, Dhamdhere, Dovrolis, "Analysis of Peering Strategy Adoption by Transit Providers in the Internet", NetEcon 2012.

Thanks! Questions?

amogh@caida.org

www.caida.org/~amogh