

Mining Routing Data from CENIC and Campus Networks

Tina Wong, Packet Design, Inc

Darrell Newcomb, CENIC

Cengiz Alaettinoglu, Packet Design, Inc

Outline


- Internet routing background
- Visualization and analysis algorithms
- Methodology
- CENIC and campus networks data

Outline

- Internet routing background
- Visualization and analysis algorithms
- Methodology
- CENIC and campus networks data

Internet Routing Background

- The Internet comprises of Autonomous Systems (ASes) and relies on hierarchical routing:
 - An AS uses IGP (e.g. OSPF, ISIS, EIGRP, RIP) to route within itself
 - Runs BGP to facilitate inter-AS connectivity



Challenges in BGP Data Mining

- BGP maintains significant amount of state:
 - About 150K prefixes in Internet
 - This would be 300K routes in dual-homed AS
 - We observe 1.5 million routes at a Tier-1 ISP
- BGP sends a lot of messages:
 - Peering session reset means table exchange
 - Major connectivity change produces million messages

Outline

- Internet routing background
- Visualization and analysis algorithms
- Methodology
- CENIC and campus networks data

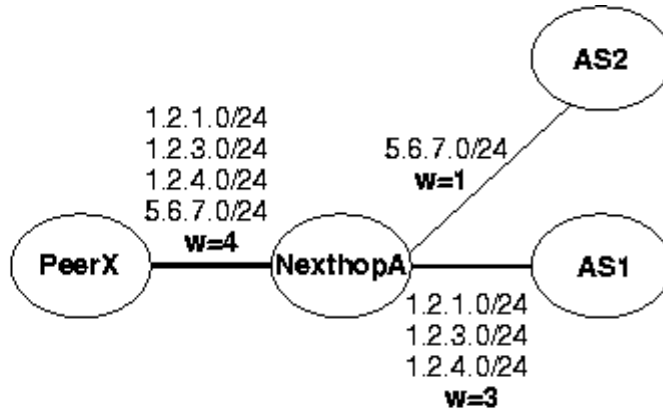
Algorithms

- We have developed two new algorithms to help make sense of BGP
 - A visualization called TAMP (Threshold And Merge Prefixes) that shows the large-scale structure of some set of BGP routes
 - An analysis technique called Stemming to do root-effect analysis of BGP messages
- Both algorithms are driven by BGP routing data alone

TAMP Visualization

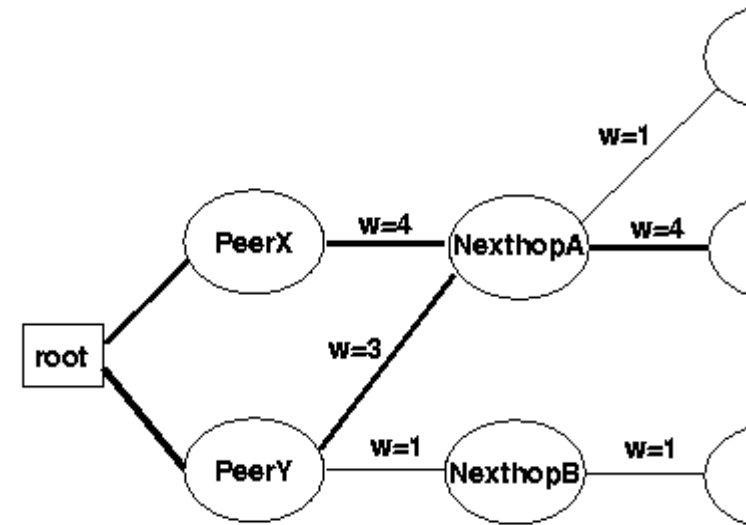
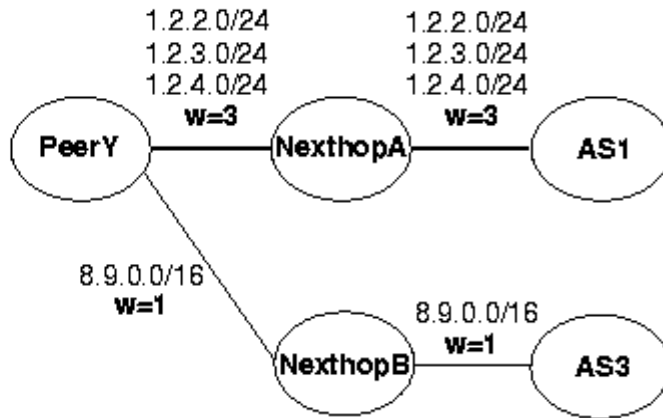
PeerX's RIB

	NH	ASPath
1.2.1.0/24	A	1
1.2.3.0/24	A	1
1.2.4.0/24	A	1
5.6.7.0/24	A	2



PeerY's RIB

	NH	ASPath
1.2.2.0/24	A	1
1.2.3.0/24	A	1
1.2.4.0/24	A	1
8.9.0.0/16	B	3



(a)

(b)

(c)

TAMP Animation

- A router's TAMP tree changes over time as it receives route announcements and withdrawals from its BGP peers
- Track these changes to generate an animation to show routing activities

Stemming

- Detects routing anomalies by finding the most strongly correlated components in a stream of BGP messages
- Each component represents a set of related routing changes
- Works well on both simple tree-like topologies (e.g. enterprises) and complex forest-like topologies (e.g. ISPs)

Outline

- Internet routing background
- Visualization and analysis algorithms
- Methodology
- CENIC and campus networks data

□ Data Collection

- Packet Design Route Explorer
 - Passively peer with iBGP speaking routers
 - Maintain adjacencies with IGP routers in each area
- CENIC-DC network
 - April 04 to present
- CENIC-HPR network
 - Dec 04 to present
- UCB network (ISP routes only)
 - Aug 03 to present

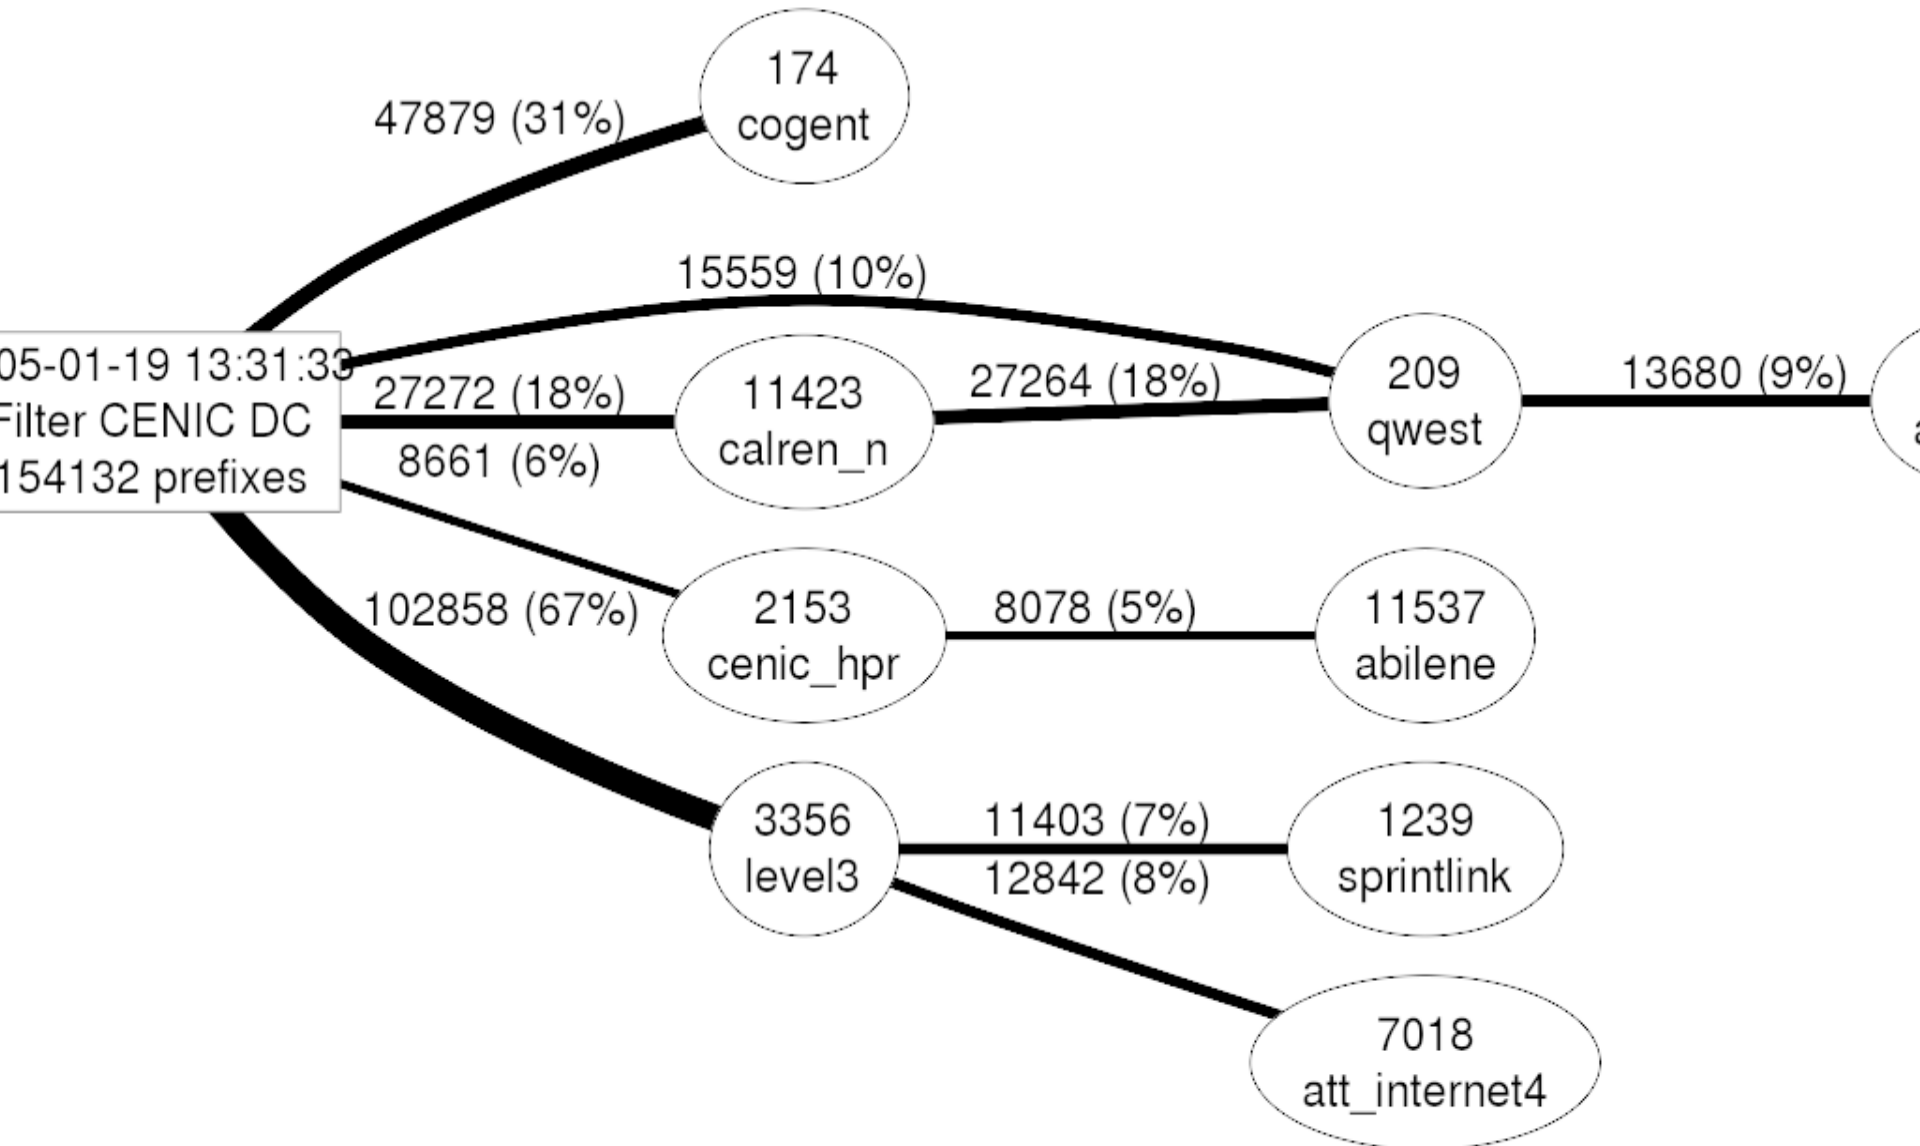
This Presentation

- Archived routing data in database
- Studied it offline with TAMP and Stemming algorithms
 - Although PD REX product is mostly for online work
- We only present some of the results here
 - Focus on inter-domain routing and BGP
 - Focus on most recent data

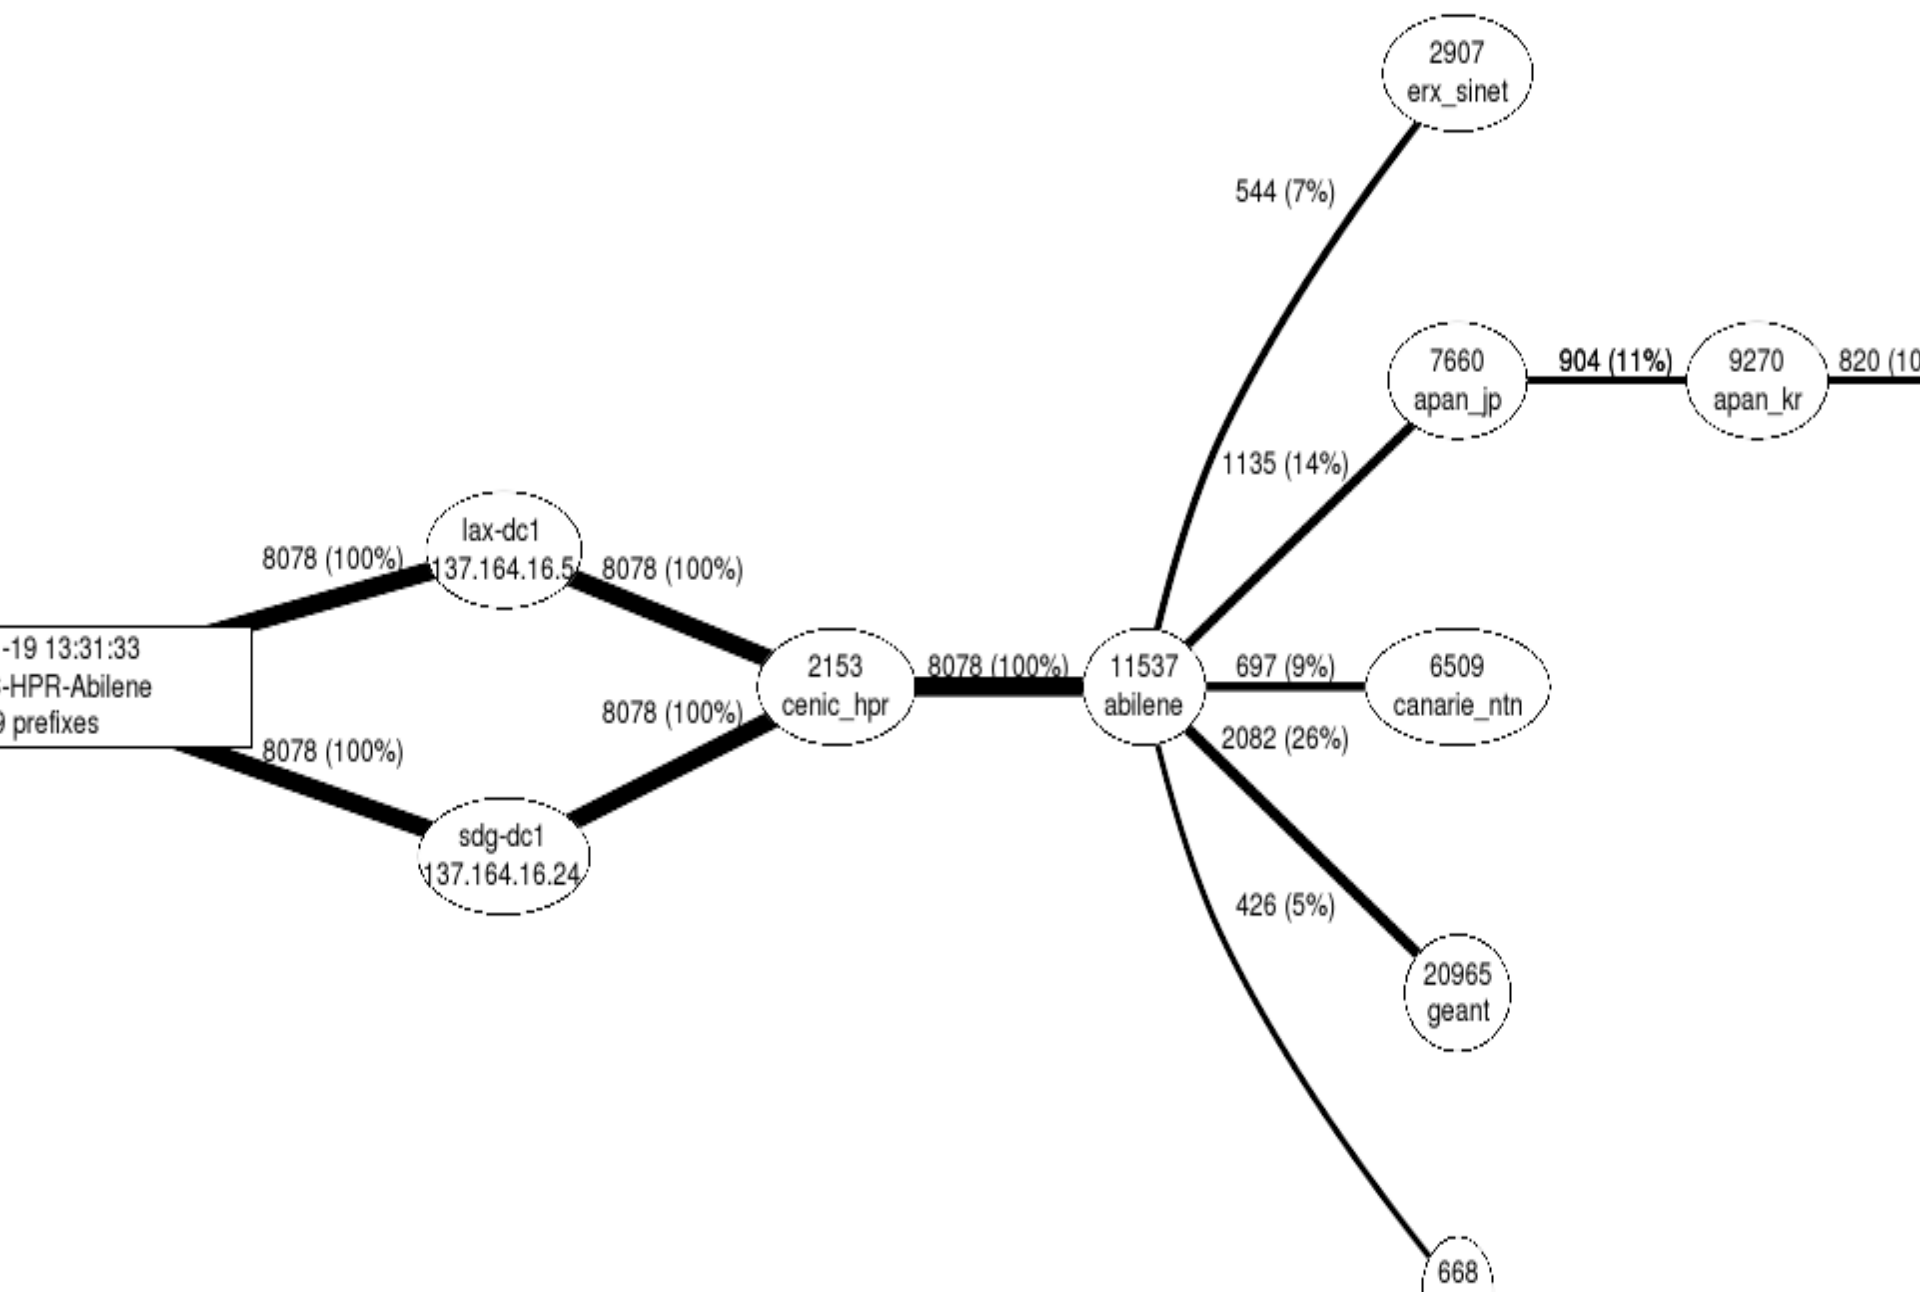
Outline

- Internet routing background
- Visualization and analysis algorithms
- Methodology
- CENIC and campus networks data

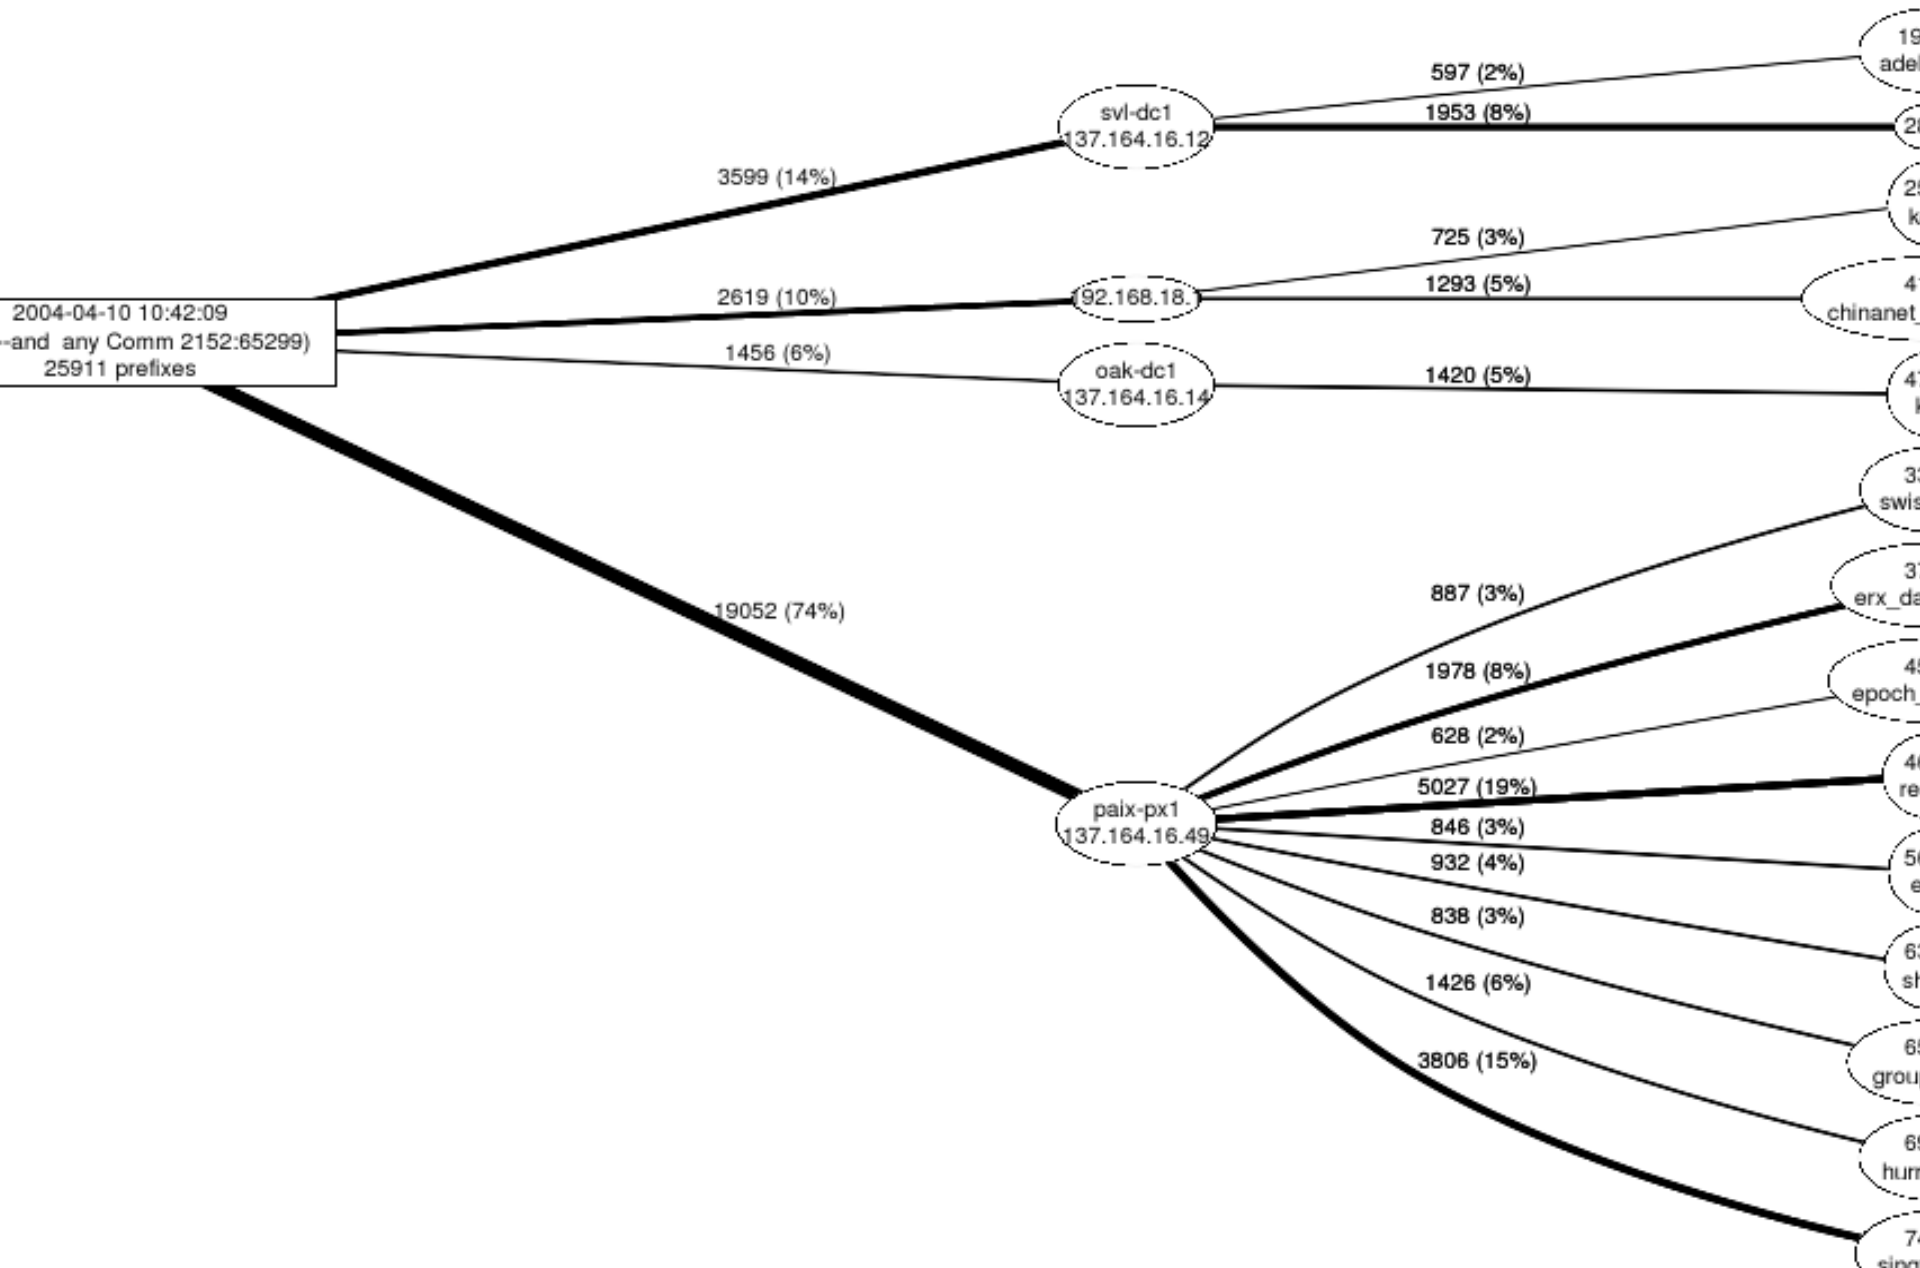
CENIC DC Network



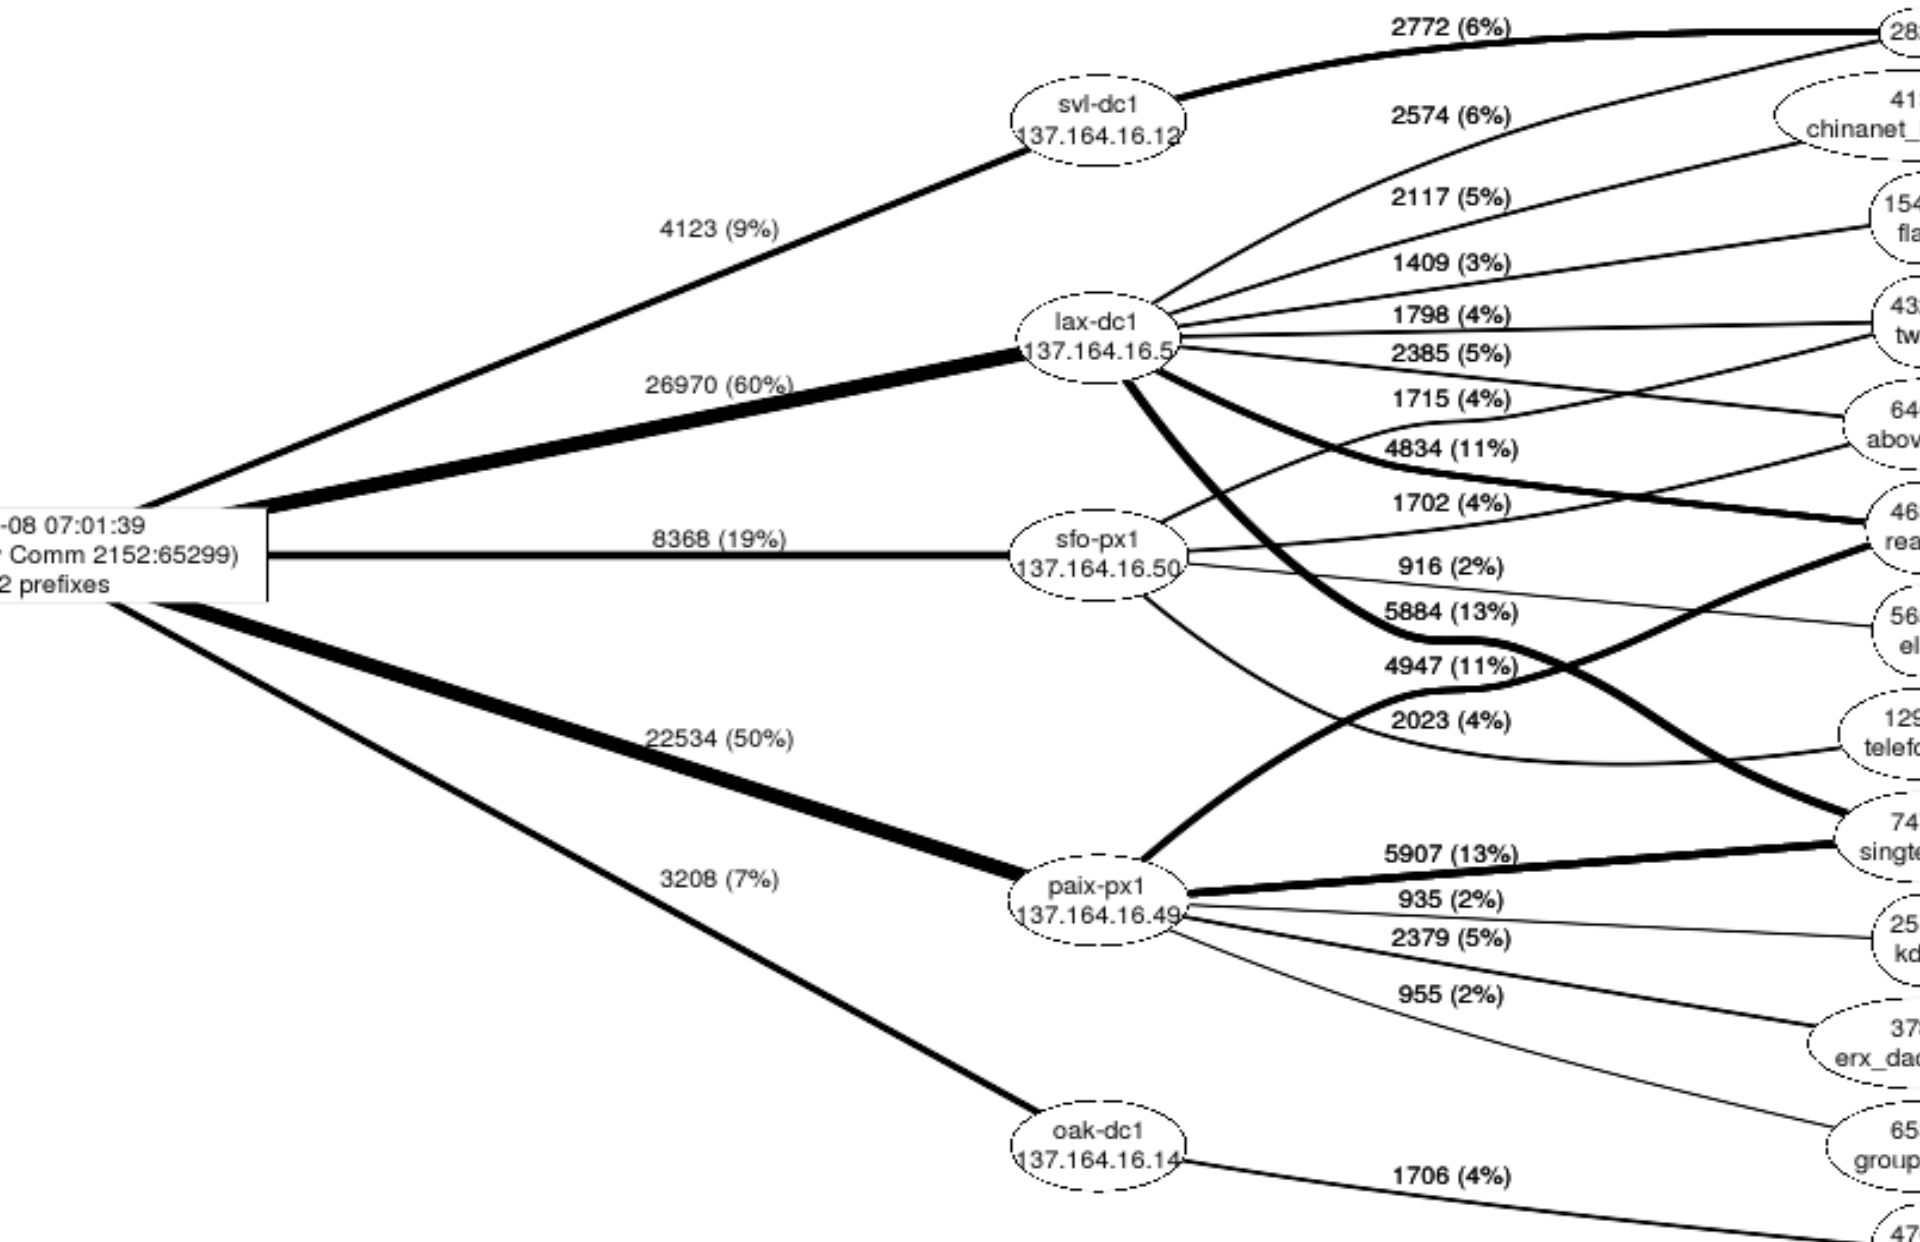
DC – HPR – Abilene




Commodity Peering: Apr 2004



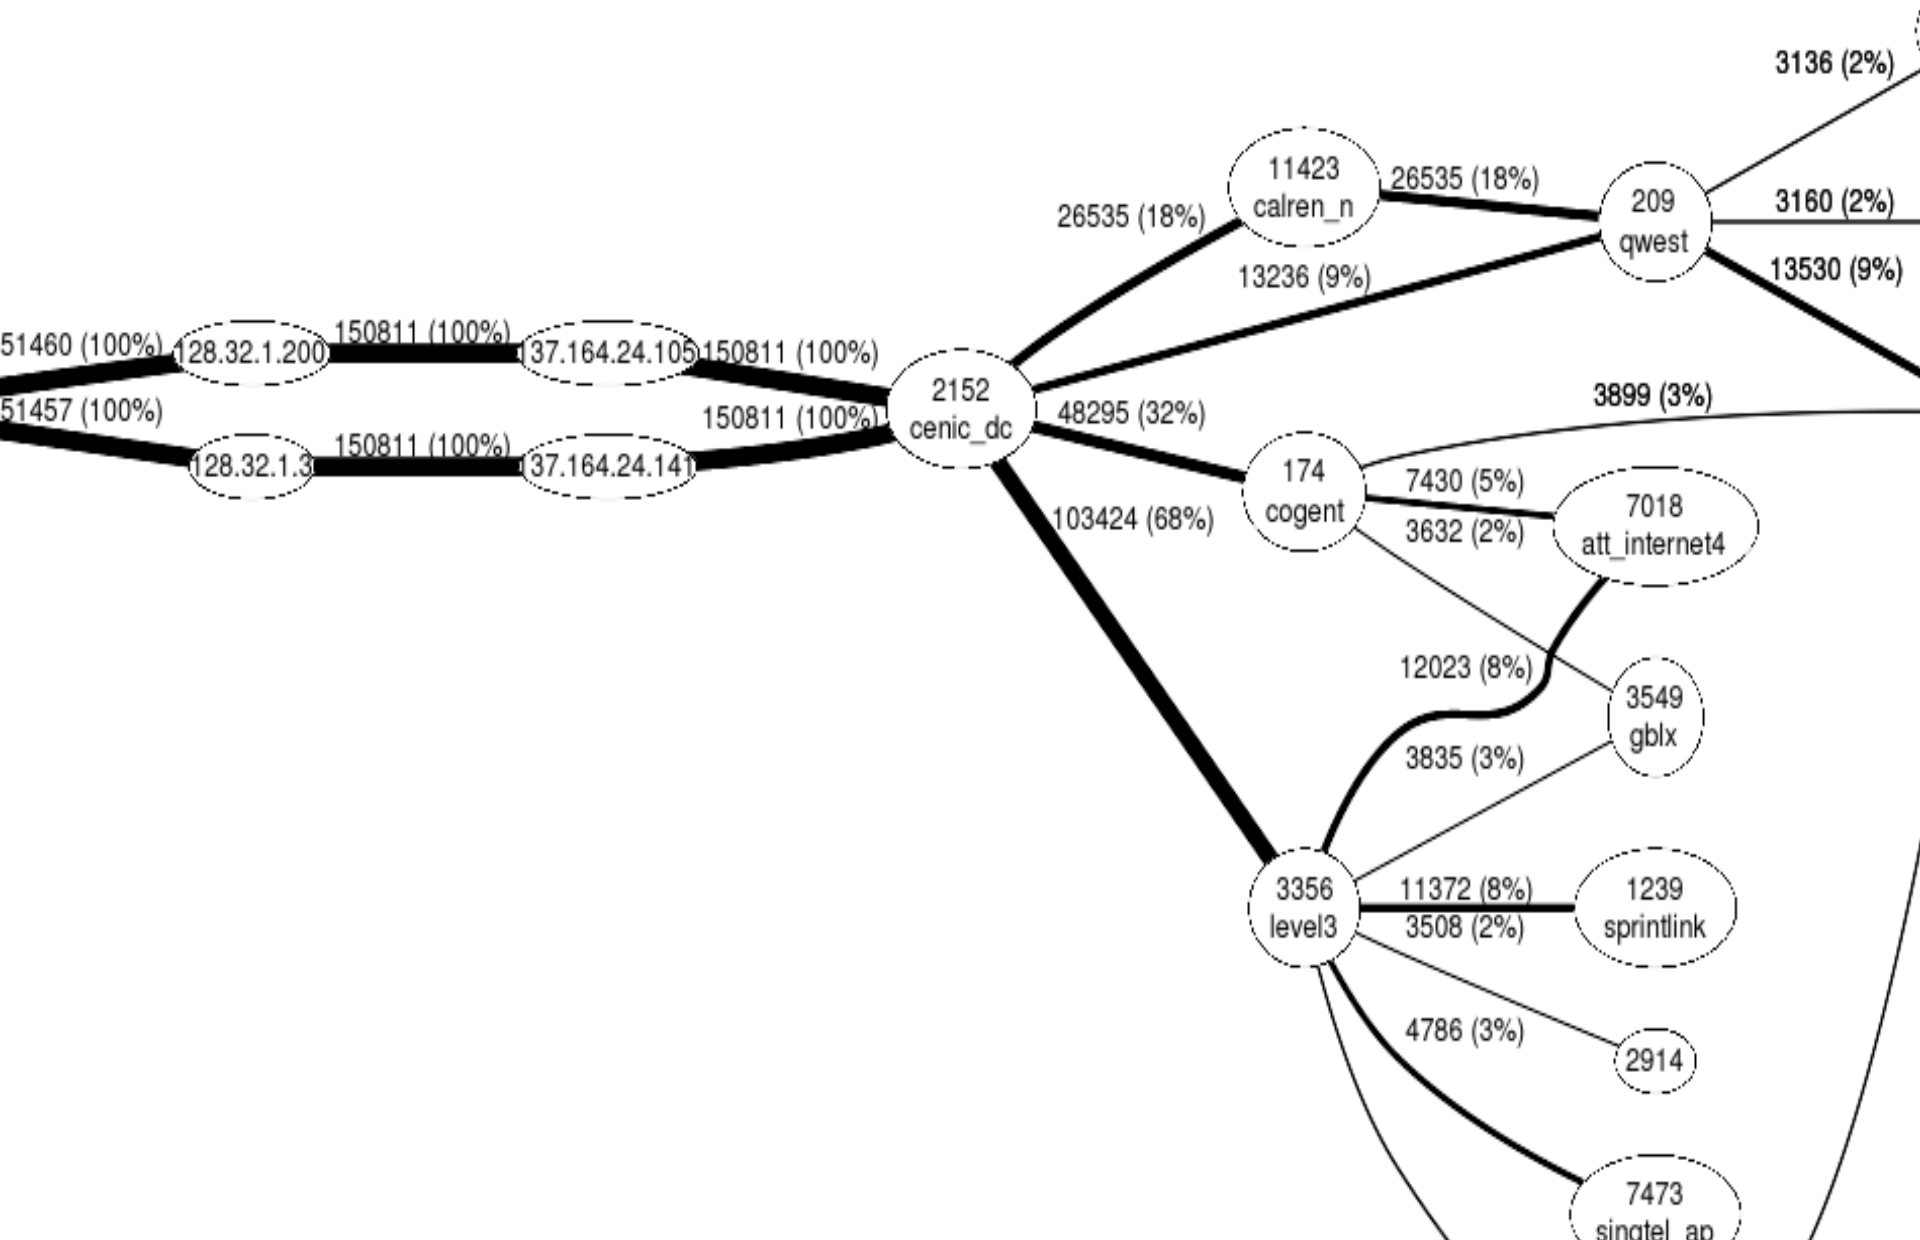
Commodity Peering: Mar 2005




Berkeley's POV: Aug 2003



Berkeley's POV: Jan 2005



Berkeley-DC-QWest-Giga-I2




HPR BGP Community for TE

- Add community ASN:2153 to prefixes for transit over HPR
- LocalPref routes with ASN:2153 to prefer over DC versions

ASN	Name	#Prefixes
195	SDSC	37
25	UCB	16
7377	UCSD	7
131	UCSB	4
31	Caltech	4
6192	UCD	3

Caltech's TE Case

- Come through LosNettos

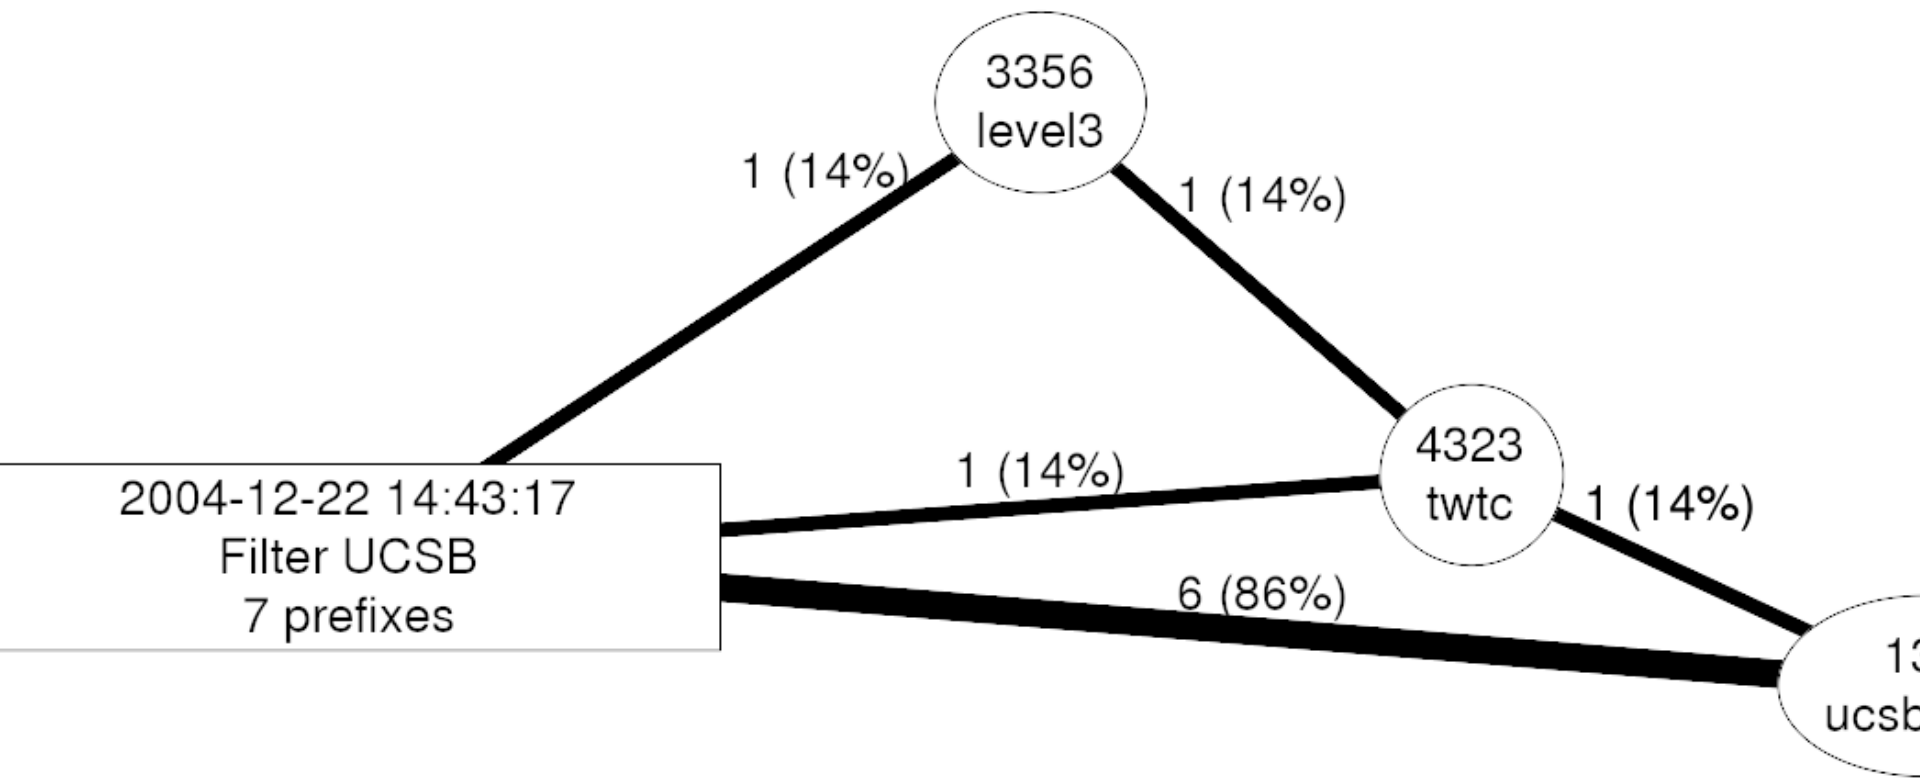


UCSB Multihoming

CENIC's POV

- Multihomed to CENIC and Time-Warner
- Advertise more specific 192.150.217.0/24 to Time-Warner, with 4x prepends
 - ASPath = 3356 4323 131 131 131 131
 - ASPath = 4323 131 131 131 131
- Advertise less specific 192.150.216.0/23 to CENIC, without prepends
 - ASPath = 131

UCSB Multihoming CENIC's POV



Without Prepend

Filtered events: cenicDC/BGP [(-and Prefix 128.111.0.0/16 any)]

Sort by: Any

Time	Router	Operation	Neighbor/ Prefix	Attributes
2004-12-21 07:25:21.475133	137.164.16.12	Withdraw	128.111.0.0/16	AS Path: 131 (IGP) Local-Pref: 100 MED: 100 Communities: 131:65498 2152:131 2152:652 Next Hop: 137.164.23.9
2004-12-21 07:26:53.010171	137.164.16.12	Announce	128.111.0.0/16	AS Path: 131 (IGP) Local-Pref: 100 MED: 100 Communities: 131:65498 2152:131 2152:652 Next Hop: 137.164.23.9

2004-12-21 06:00:00 - 2004-12-21 10:29:38



entries

Route Convergence Times with Path Exploration

(this is an SVG animation)

Turkish Telecom Leaked Routes

CENIC's POV

- Dec 24 2004, from midnight to late morning
- From a few prefixes to a few thousand prefixes
- Most leaks were relatively short duration
- Prefix hijack victims include
 - GE, Army Research Lab, ATT, Level3
 - etc

Turkish Telecom Leaked Routes CENIC's POV

(this is an SVG animation)

Turkish Telecom Leaked Routes Berkeley's POV

(this is an SVG animation)