Observations on Routing Policies and Traffic Engineering Practice

Feng Wang and **Lixin Gao**
Department of Electrical and Computer Engineering
University of Massachusetts
{fewang,lgao}@ecs.umass.edu

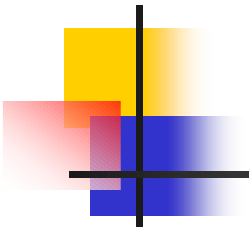

Outline

- Internet Architecture
- Analysis of traffic flow patterns
- Inferring routing policies
- Impact on network performance
- Conclusions

AS Relationships

- Provider to customer: customer pays provider for transit traffic
- Peer to peer: exchange traffic between customers free of change
- Sibling to sibling: transit traffic for each other

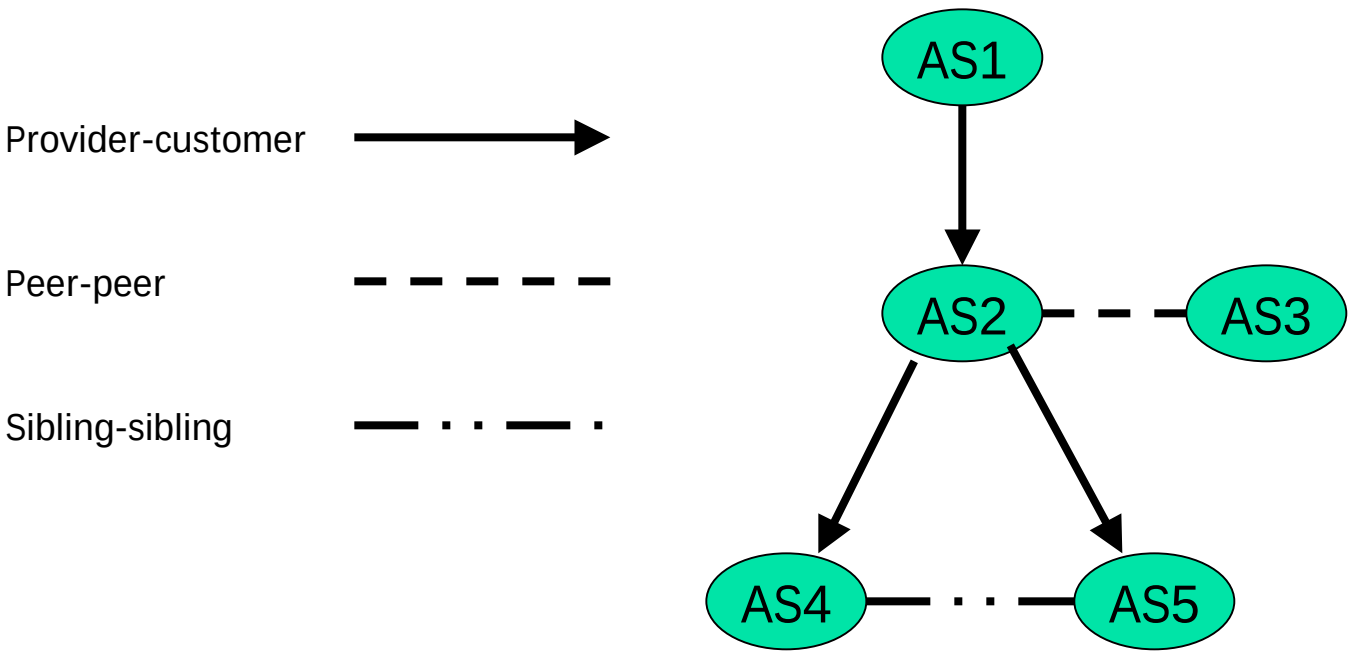

Hierarchical Internet Structure

- Tier-1 AS: providers that access the global Internet and don't buy network capacity from other providers

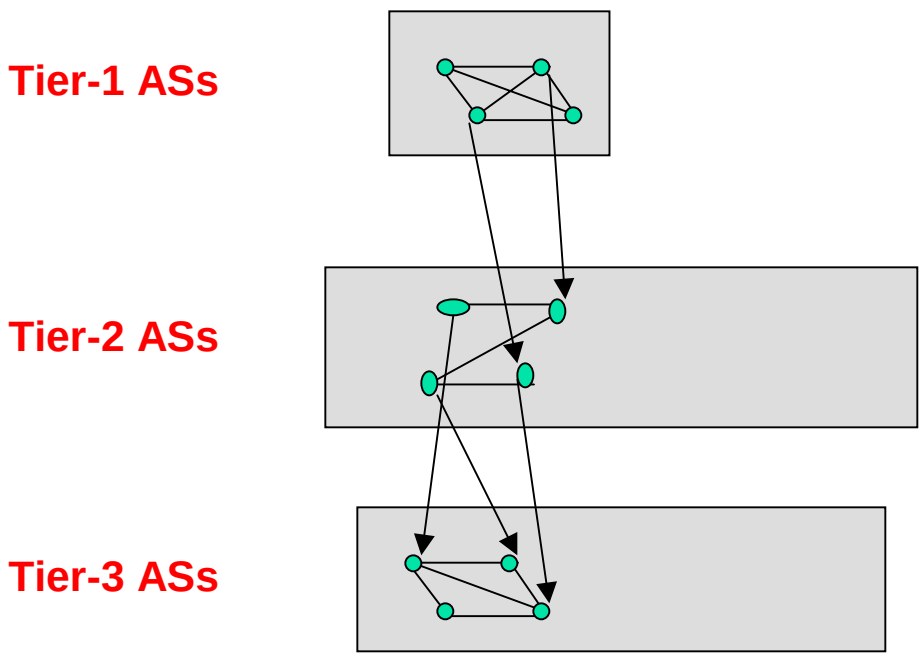

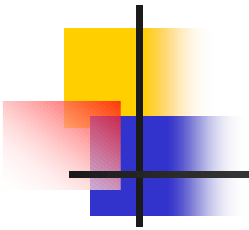

Our Work

- Internet traffic flow patterns
  - Use peer route even if customer route exist?
  - Use provider route even if customer or peer route exist?
- What routing policies cause flow patterns?
- Impact on performance ?

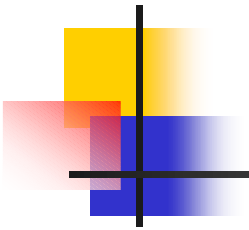

Related Work

- Savage et al (SIGCOMM'99)
  - there are better alternate paths
- Padmanabhan et al (SIGCOMM'01)
  - there are circuitious routes
- Tangmunarunkit et al (INFOCOM'00)
  - Longer paths due to routing policies

Traffic Flow Patterns

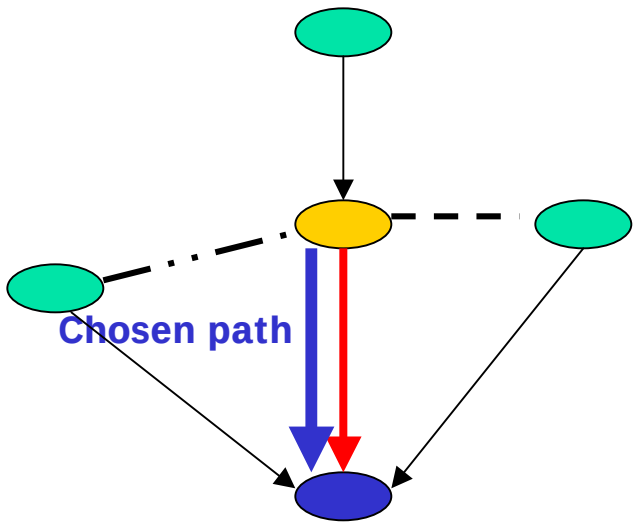

Cu/other

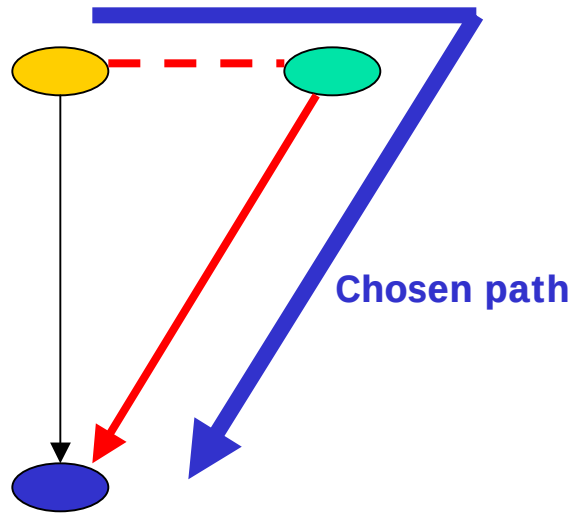

Pe/Cu

More Traffic Flow Patterns

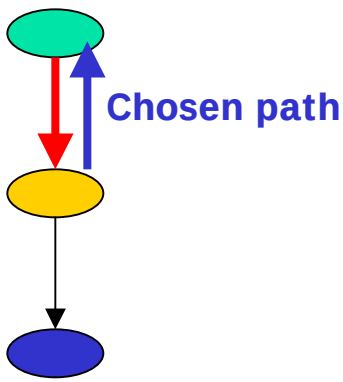

Pr / Cu

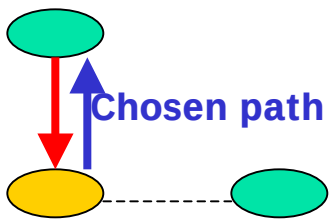

Pr / Pe

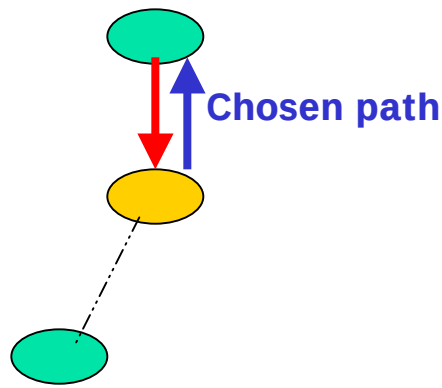

Pr / Si

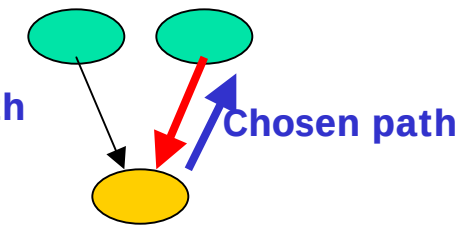

Pr / (otherPr)

More Traffic Flow Patterns

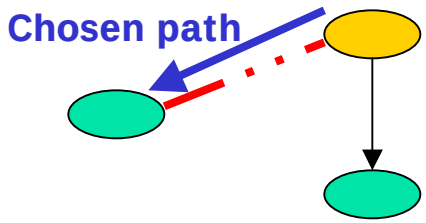

Si/Cu

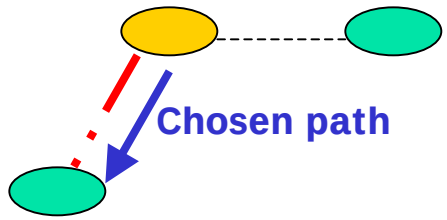

Si/Pe

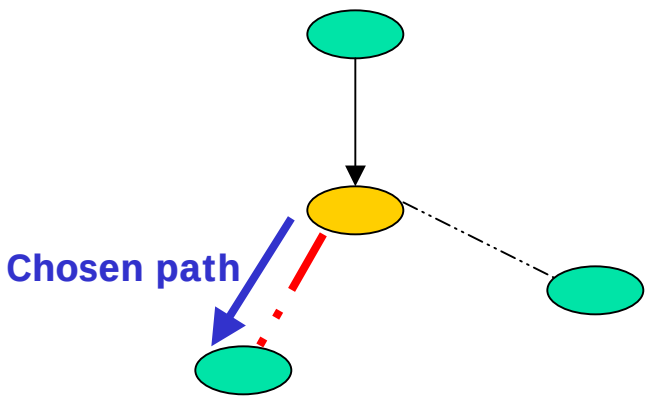

Si/(Si, Pr)

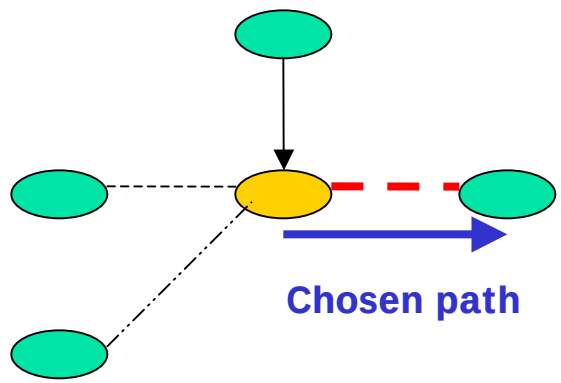

Pe/(Pe,Si,Pr)

Best Available Path

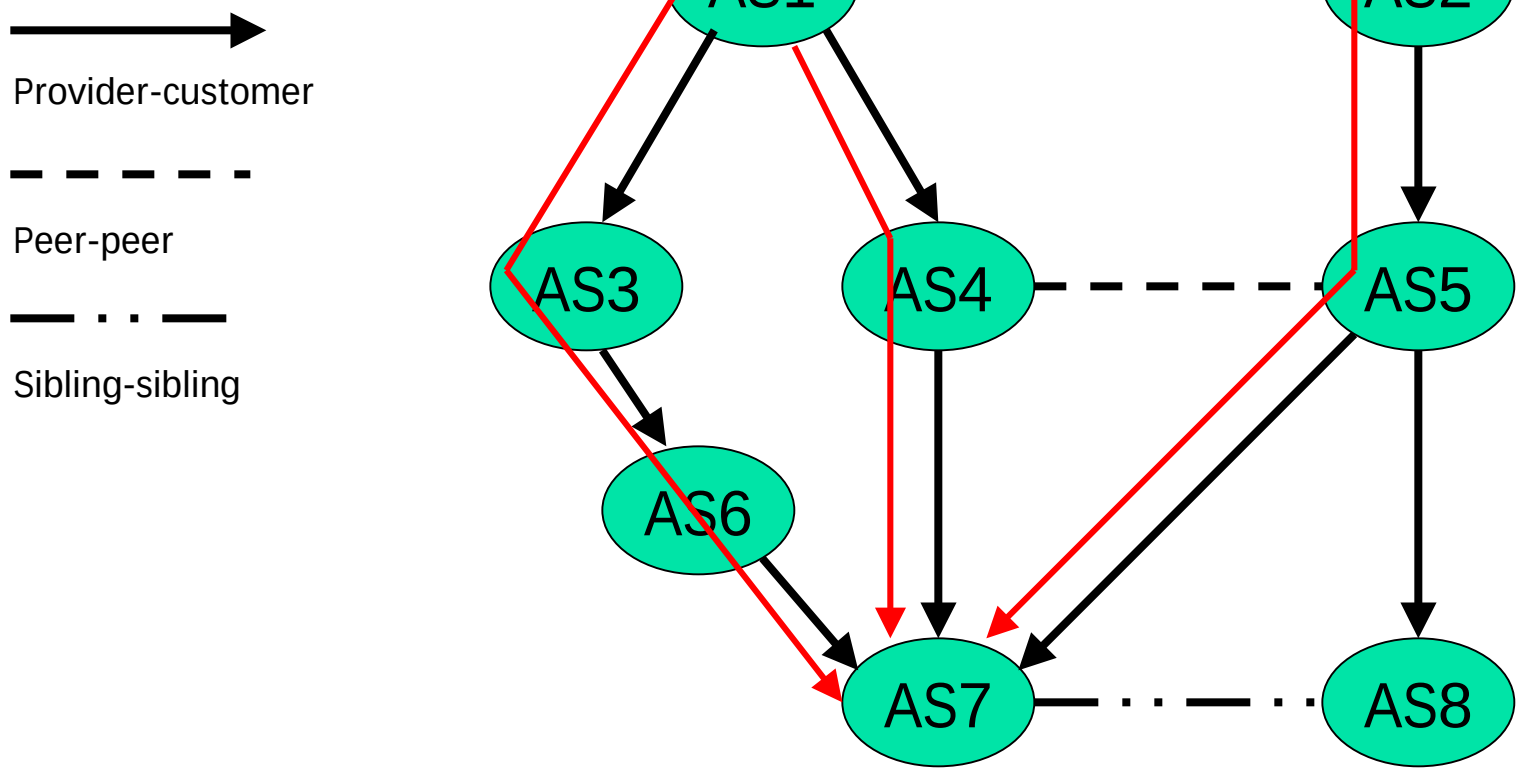

AS path {1 4 7} is the best available path

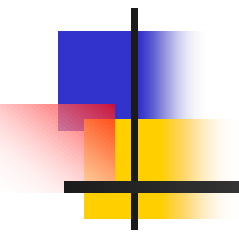

Analysis of traffic flow patterns for Tier-1 ASs

Use AS relationships inferred from Gao'00

Traffic Flow Patterns

- For AS1, 16% of prefixes belong to the Pe/Cu category
- For AS3549, 5% of prefixes belong to the Pe/Cu category

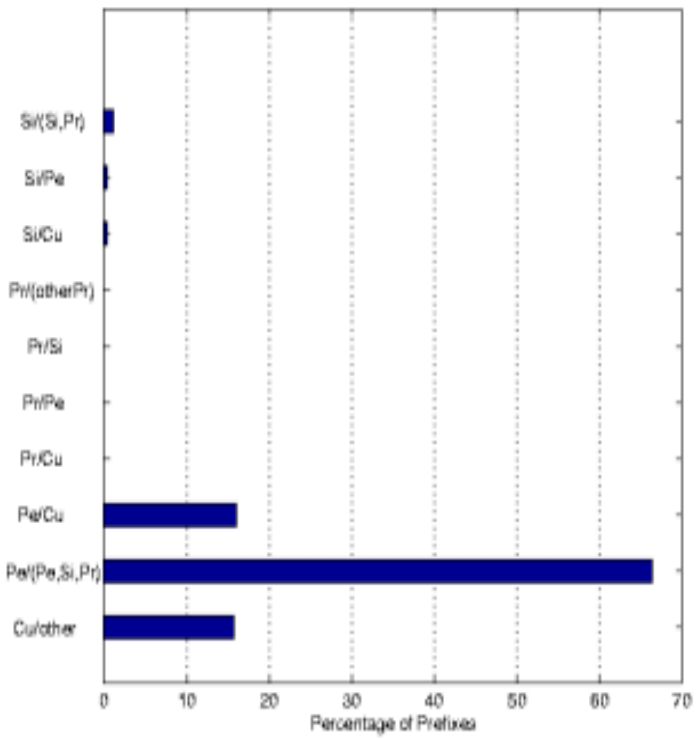

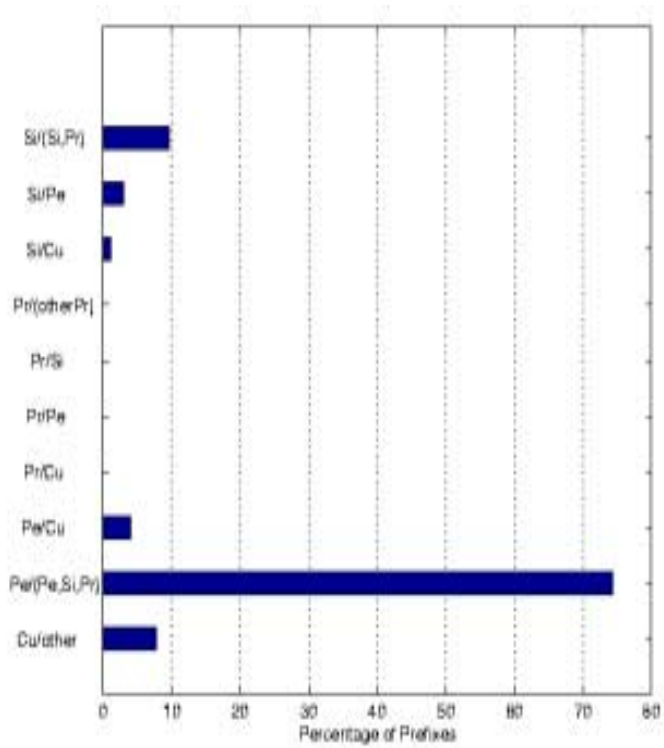

AS1 Jan 4, 2002 AS3549

Minimum Number of Prefixes Belonging to Pe/Cu

Next hop AS is one of ten Tier-1 ASs

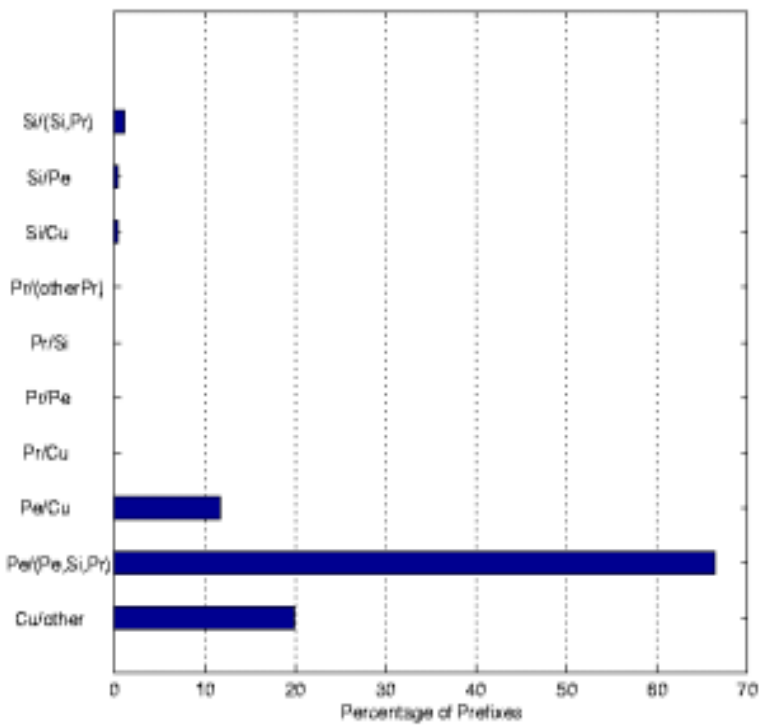

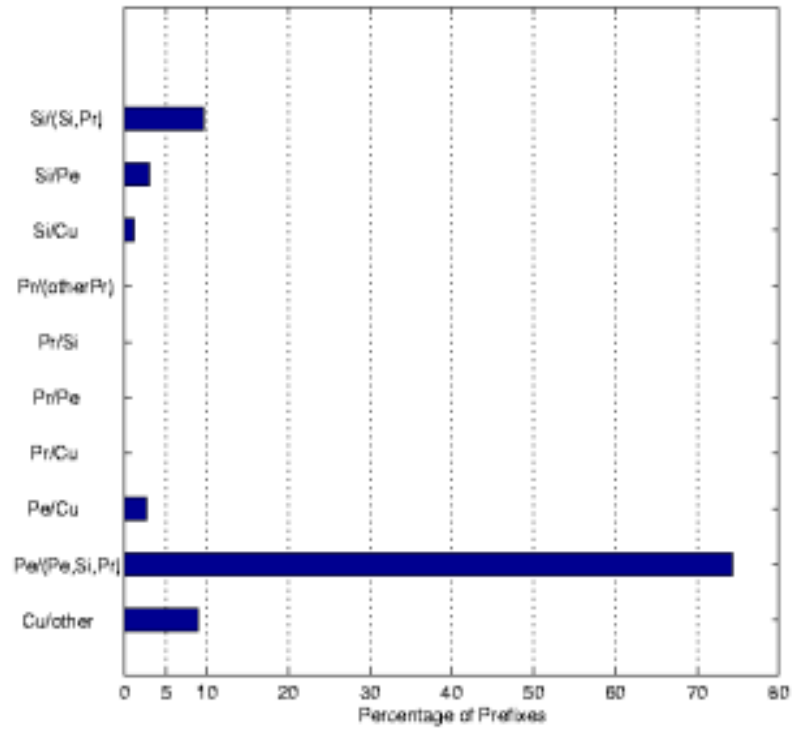

AS1 Jan 04, 2002 AS3549

Prevalence on 7 Tier-1 ASs

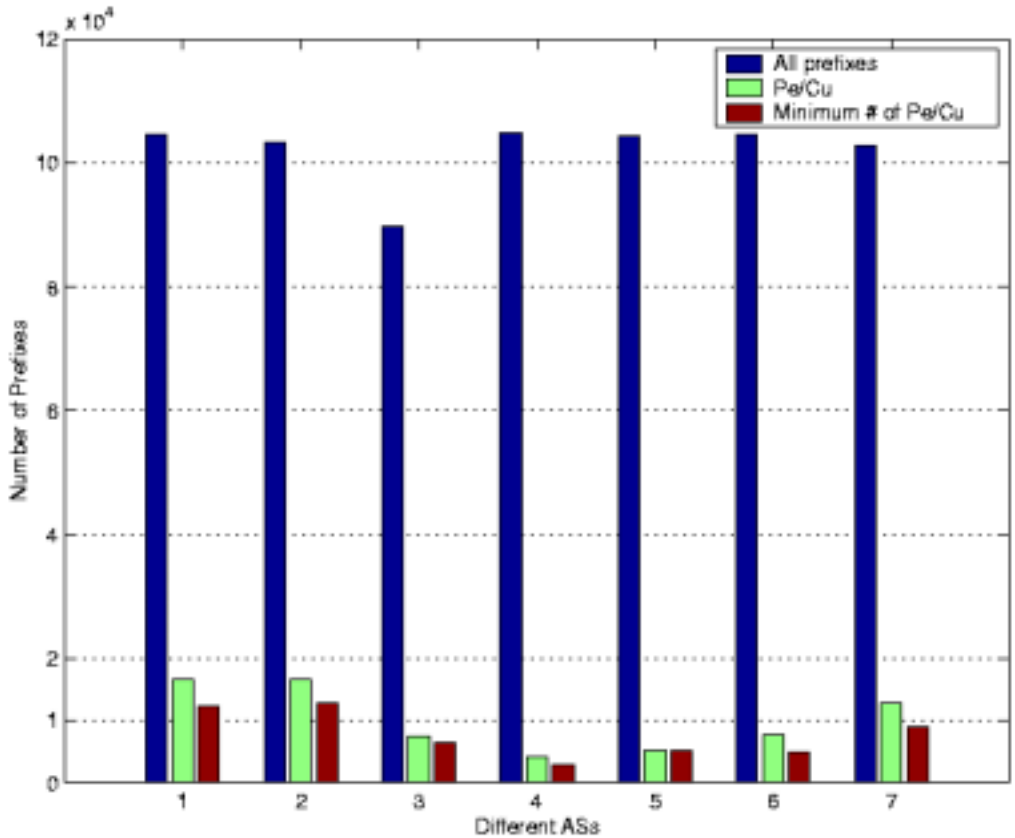

Jan 04, 2002

Comparing the Length of Chosen and the Best Available Paths

	AS1	AS3549
# of prefixes	104554	104779
# of Pe/Cu prefixes	16698	4220
Minimum # of Pe/Cu prefixes	12255	2872
# of Pe/Cu prefixes that take longer than best available path	4887 (30%)	1782 (42%)
Minimum # of Pe/Cu prefixes that take longer than best available path	3802	1379

Persistence of Traffic Flow Patterns

For AS 1

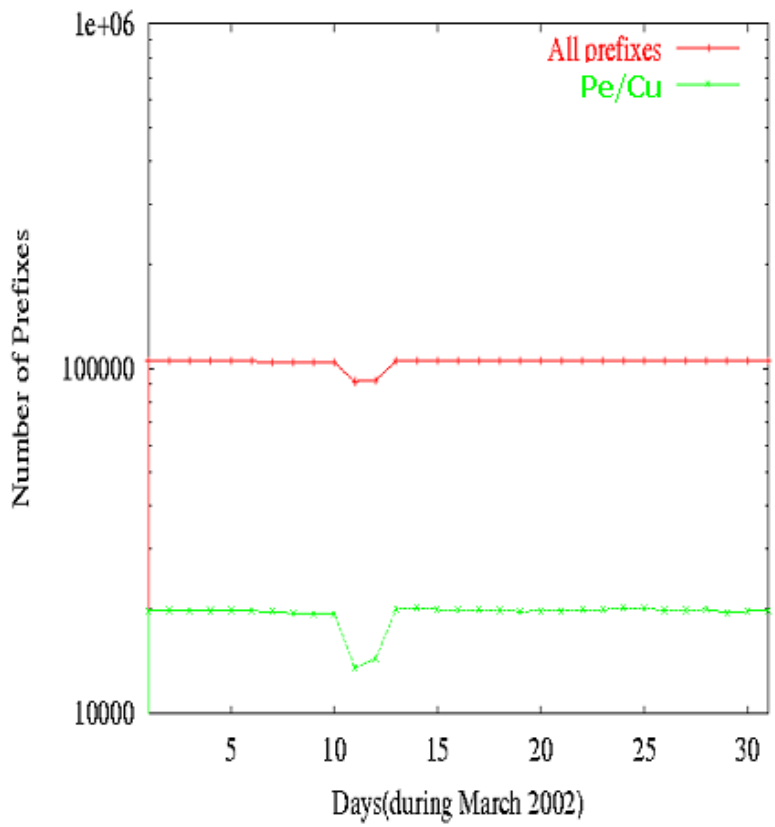

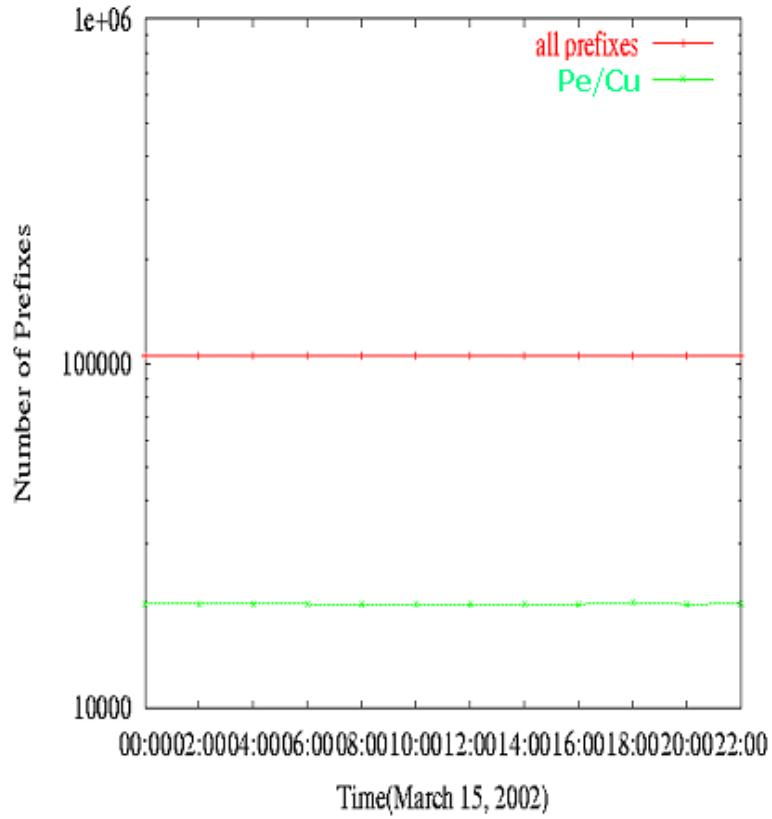

Three possible causes of traffic flow patterns

- Import policies:
  - Routing preference Anomaly: peer route has higher local_pref than customer route
  - Equal Local_Pref: peer route has the same local_pref as customer route. Other attributes (AS path length, MED etc) is the cause.
- Export policies:
  - No routes: the best available route is not received.

Local_Pref in AS1

AS relationships	Local_Pref					
	200	110	100	90	50	30
Customer(506)	1	1	476	26	2	0
Peer(46)	0	0	2	1	42	1

Consistency of next-hop AS with Local_Pref

AS number	Percentage of Ass(%)		
	customer	peer	Mapped provider
AS1	99.2	93.5	0
AS3549	98.9	97.5	0

Distribution of three possible causes

	Percentage(%)		
	No routes	Routing preference anomaly	Equal Local_Pref
AS1	97	1	2
AS3549	98	1	1

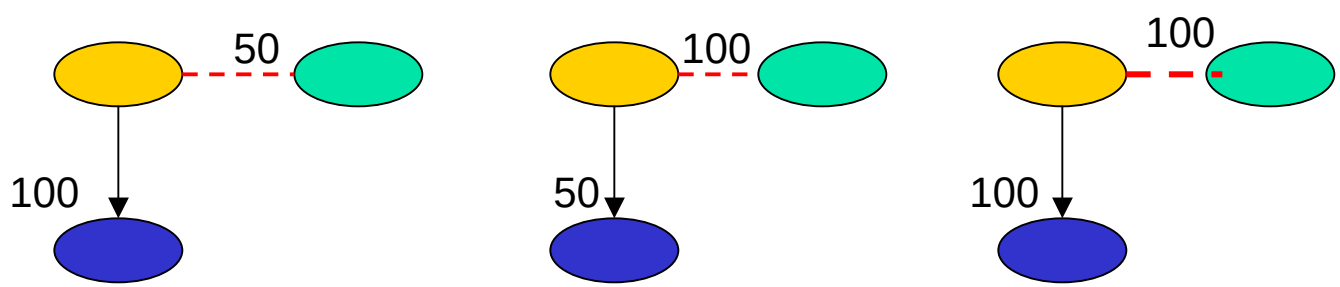

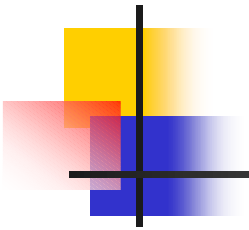

Inferring Export Policies

- Direct providers receive routes: direct providers in the best available path receive announcements from originating AS.
  - Origin AS might announce prefix to direct provider with no export
- Direct providers do not receive routes: direct providers in the best available path do not receive announcements from originating AS
  - Origin AS does not announce prefix to direct provider

Methodology of inferring export policies

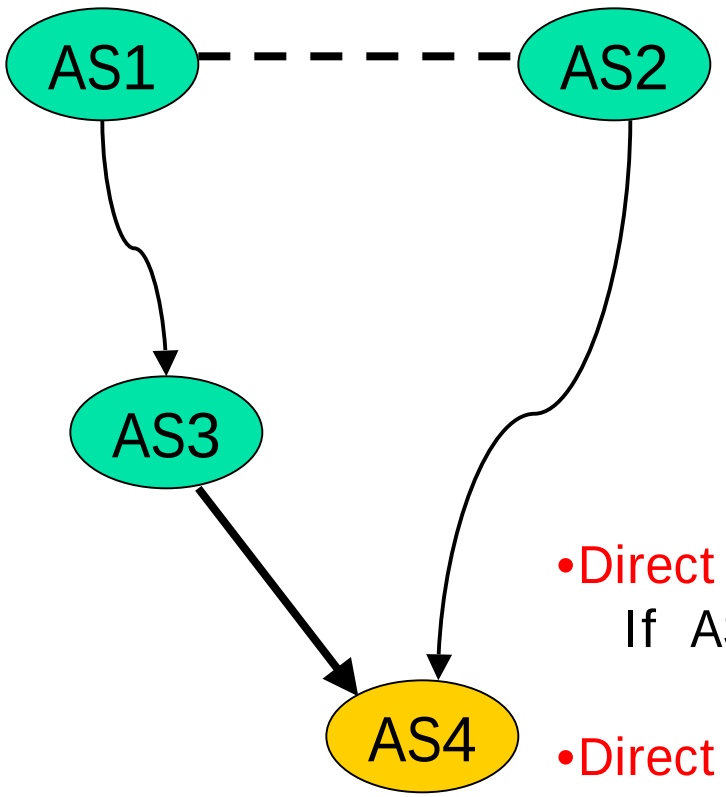

- Direct provider AS3 receives routes:
  If AS path {3 4} appears in BGP table
- Direct provider AS3 does not receive routes:
  If AS path {3 1 2 4} appears in BGP table

Percentage of two causes

class	Number	Percentage
Direct providers receive routes	10042	88%
Direct providers do not receive routes	1416	12%

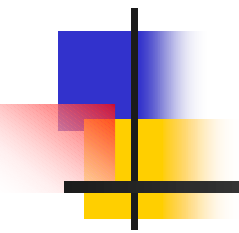

Impact on network performance

Traceroute from AS7018

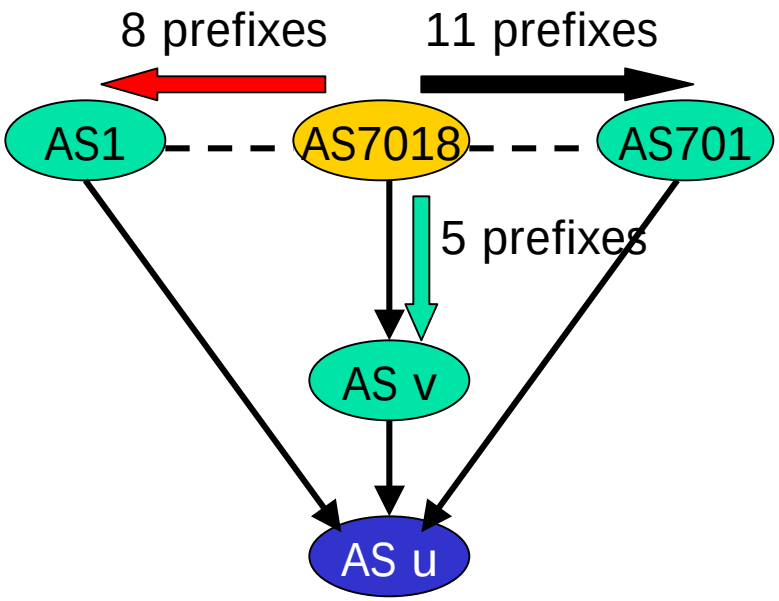

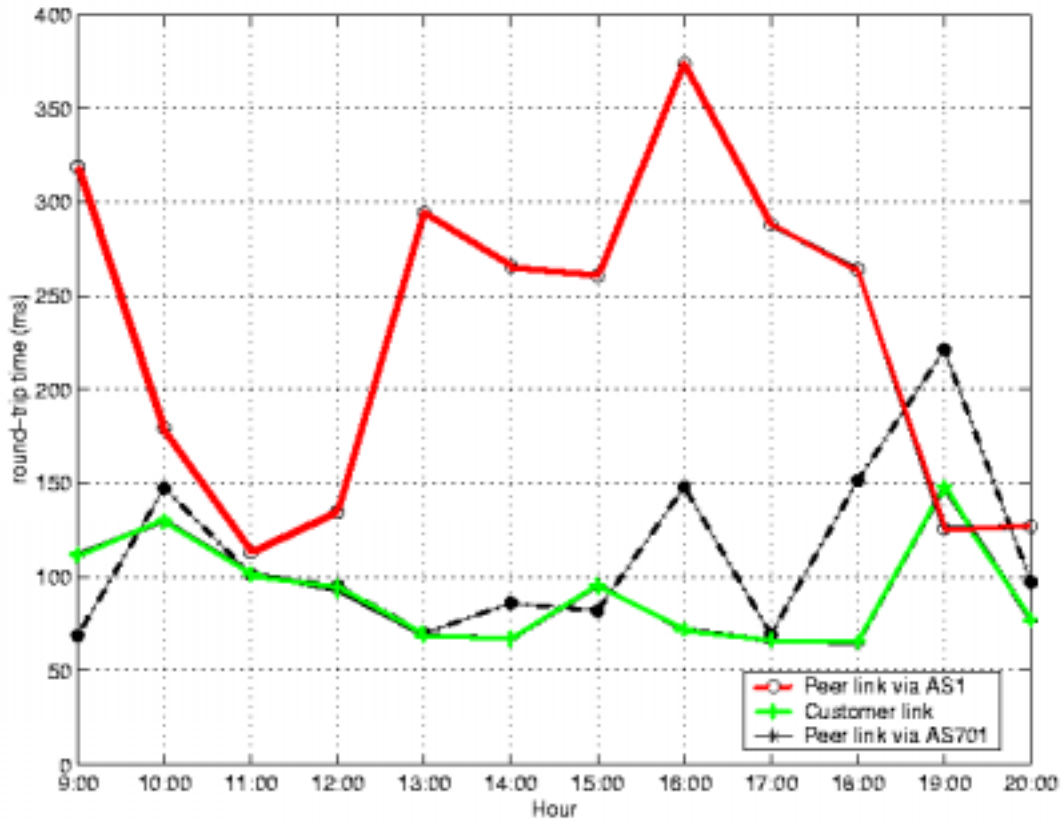

Traceroute from AS1

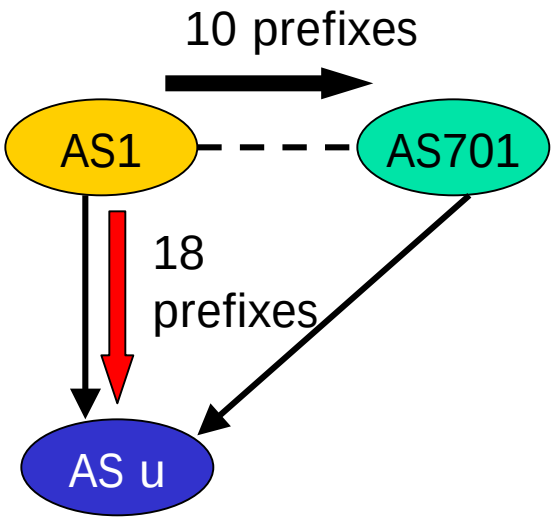

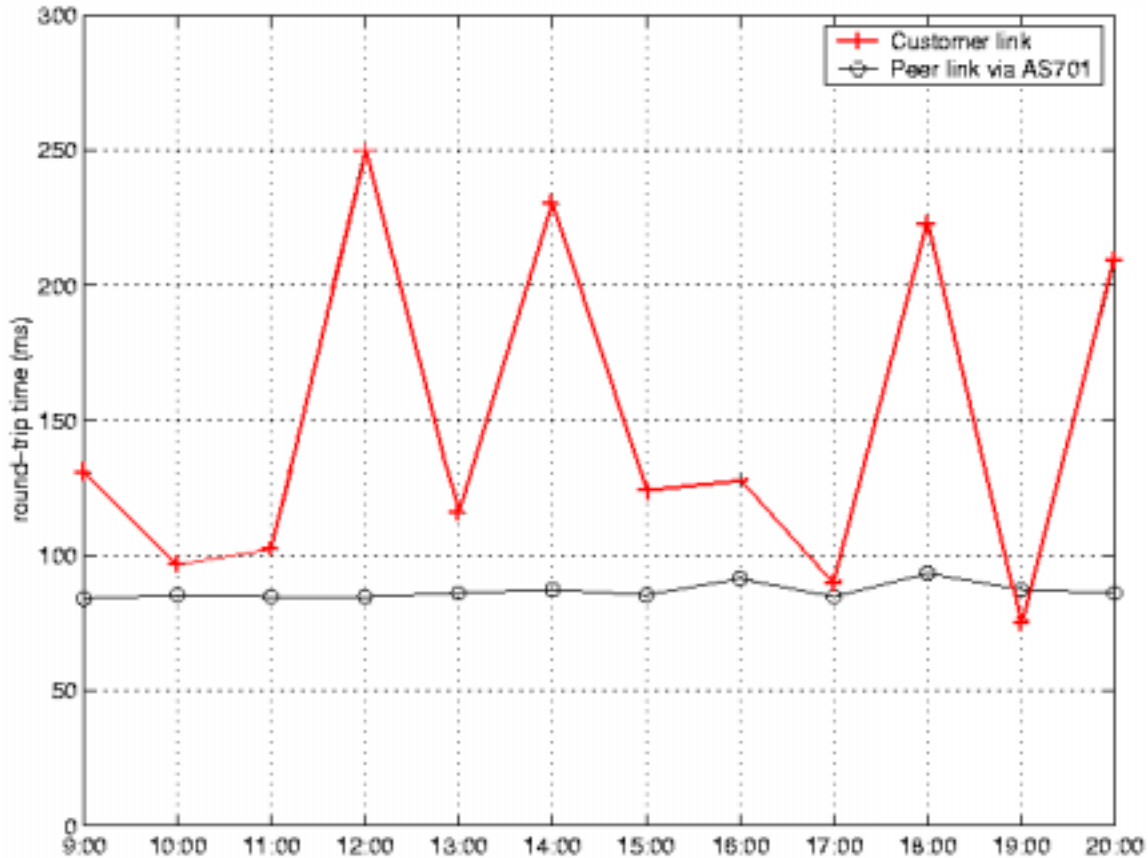

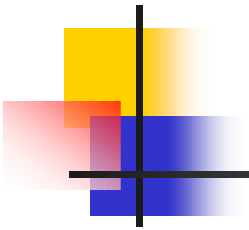

Conclusions

- Observations on some unexpected traffic flow patterns
- Observations on traffic engineering practices
  - active traffic engineering of origin ASs
- Impact on Performance