CE 817 - Advanced Network Security
Routing Security I

Lecture 20

Mehdi Kharrazi
Department of Computer Engineering
Sharif University of Technology

Acknowledgments: Some of the slides are fully or partially obtained from other sources. Reference is noted on the bottom of each slide, when the content is fully obtained from another source. Otherwise a full list of references is provided on the last slide.

What is Routing Security?

- Bad guys play games with routing protocols.
- Traffic is diverted.
  - Enemy can see the traffic.
  - Enemy can easily modify the traffic.
  - Enemy can drop the traffic.
- Cryptography can mitigate the effects, but not stop them.

The Enemy's Goal?

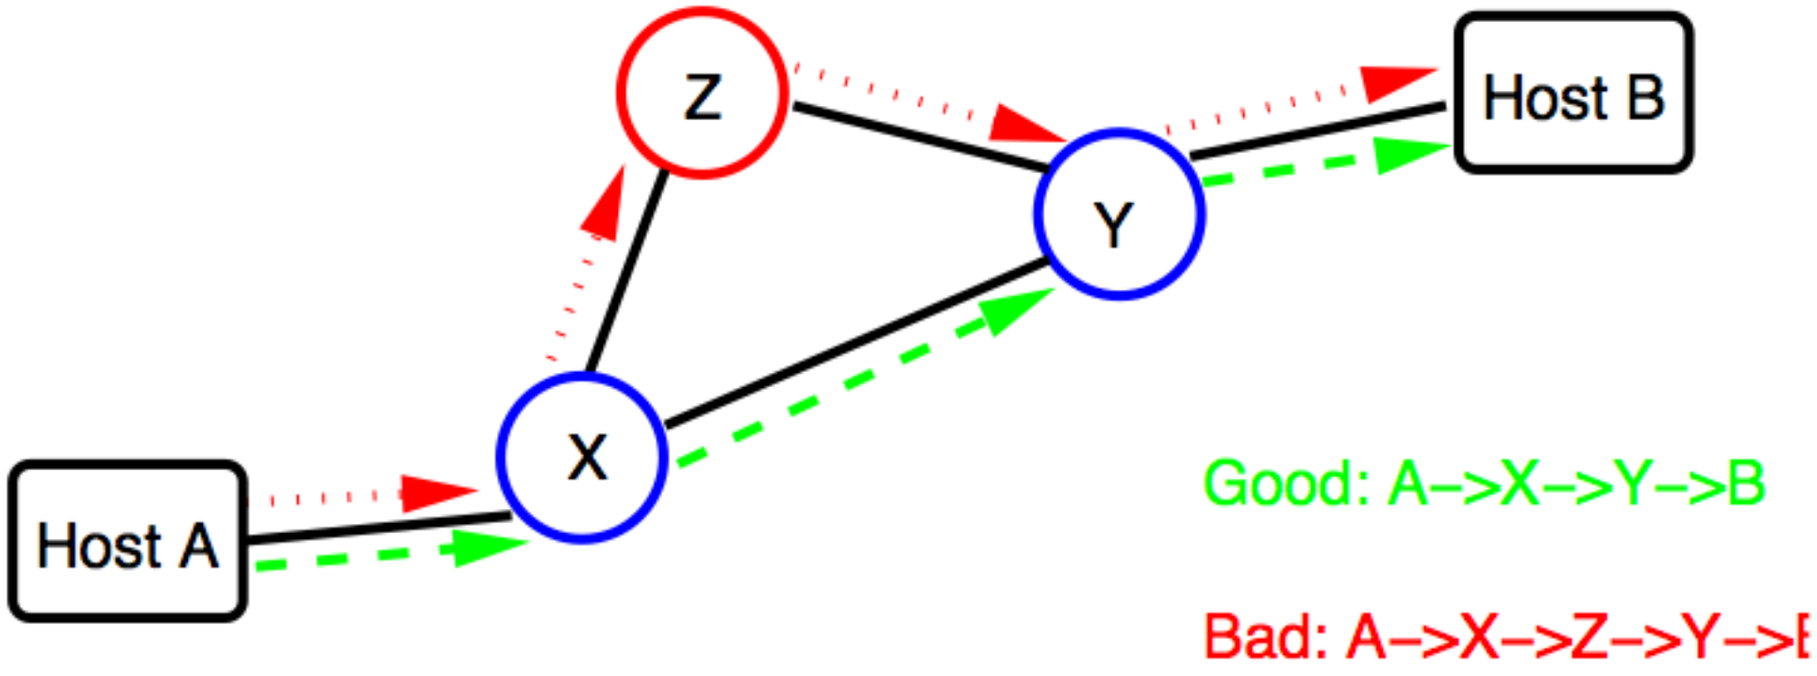

- But how can this happen?

Routing Protocols

- Routers speak to each other.
- They exchange topology information and cost information.
- Each router calculates the shortest path to each destination.
- Routers forward packets along locally shortest path.
- Attacker can lie to other routers.

Normal Behavior

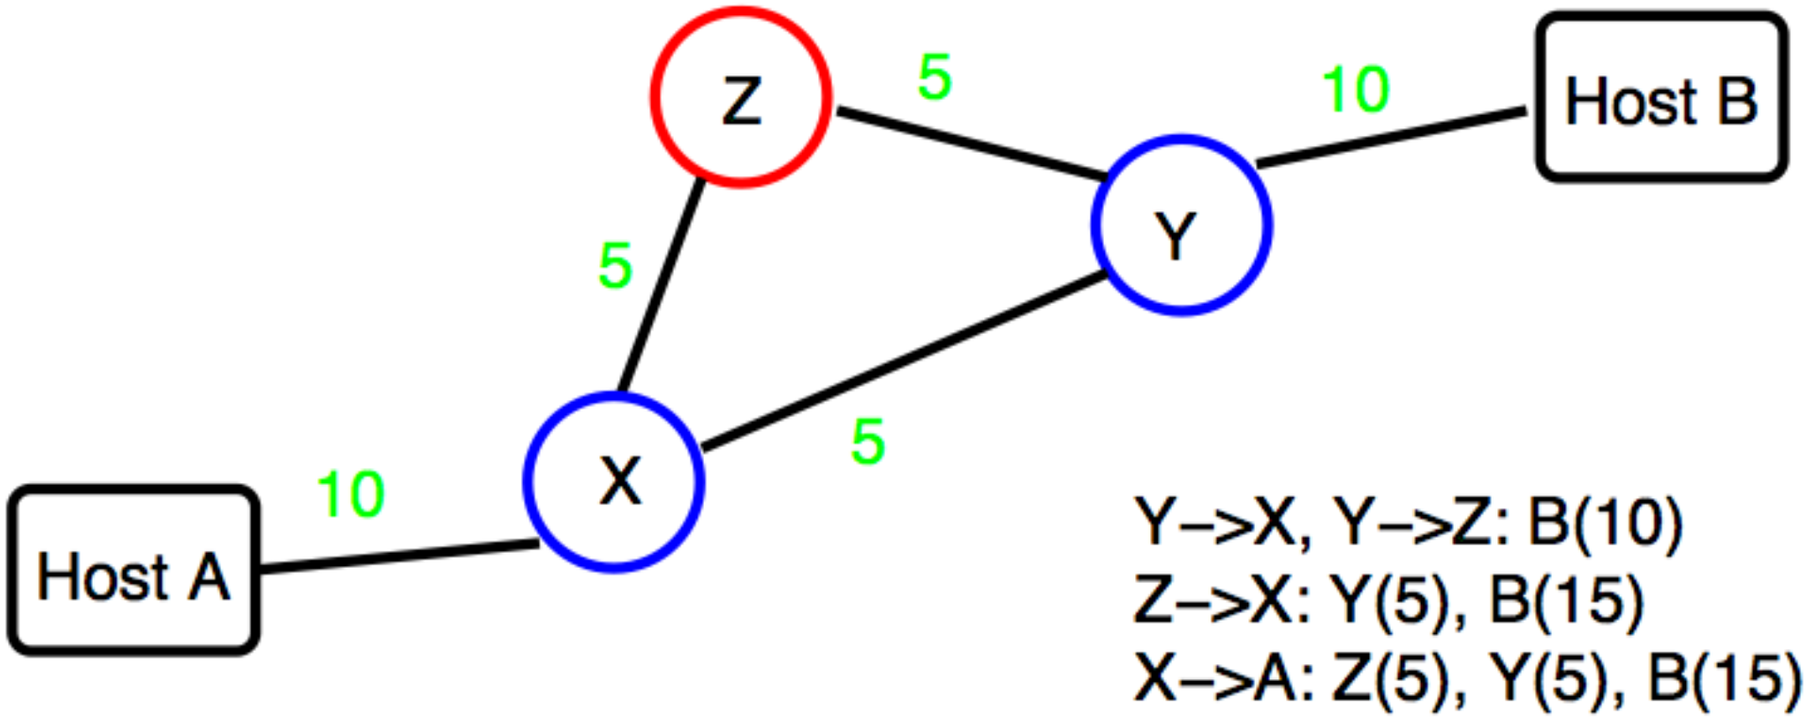

But Z Can Lie

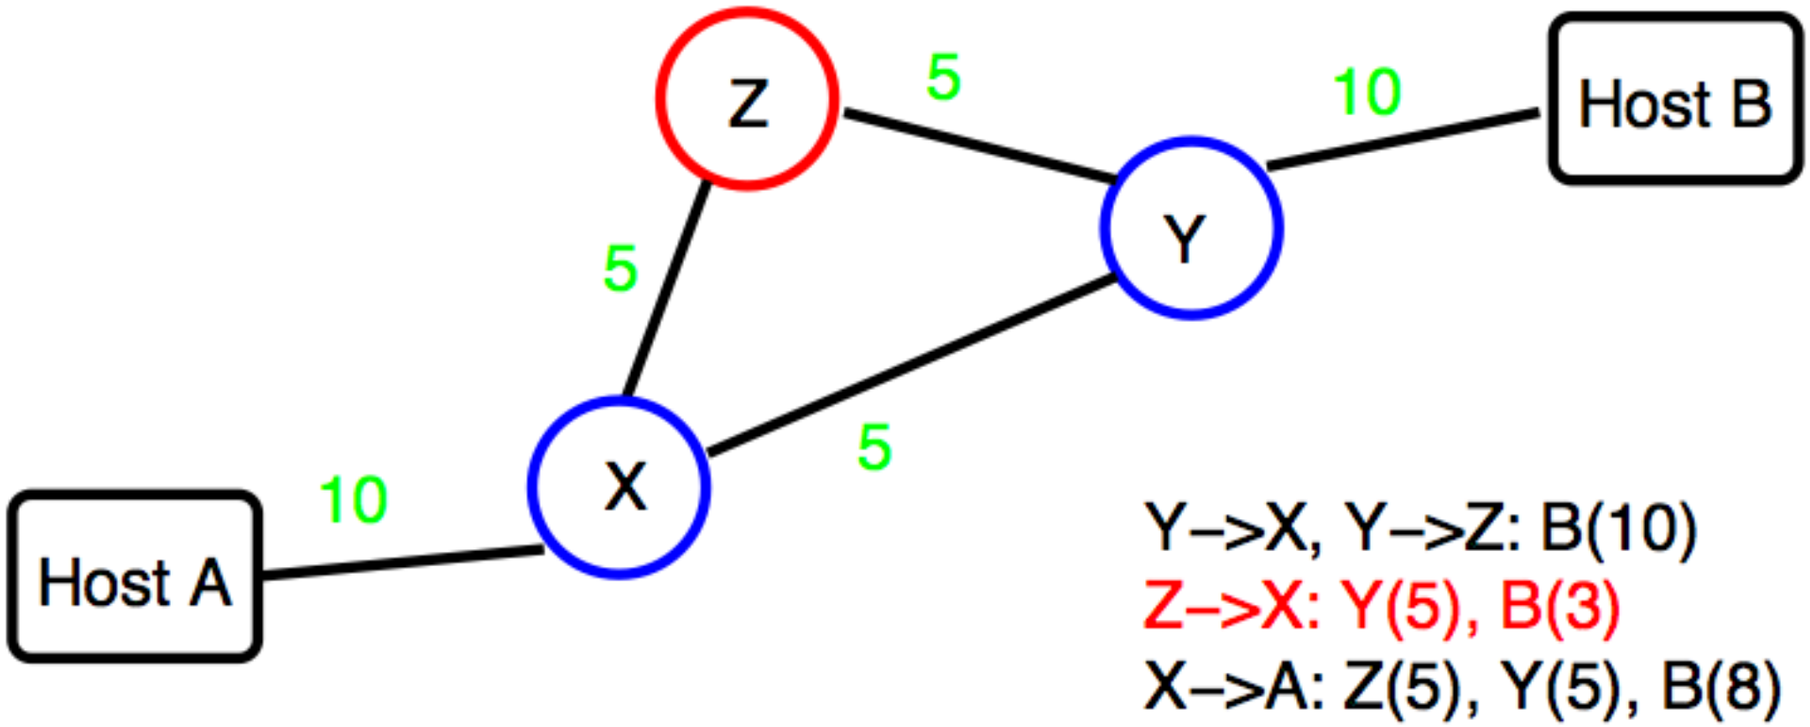

- Note that X is telling the truth as it knows it.

Why is the Problem Hard?

- X has no knowledge of Z's real connectivity.
- Even Y has no such knowledge.
- The problem isn't the link from X to Z; the problem is the information being sent. (Note that Z might be deceived by some other neighbor Q.)

Routing in the Internet

- Two types, internal and external routing.
- Internal (within ISP, company): primarily OSPF.
- External (between ISPs, and some customers): BGP.
- We will focus on External routing

BGP: A Quick Review

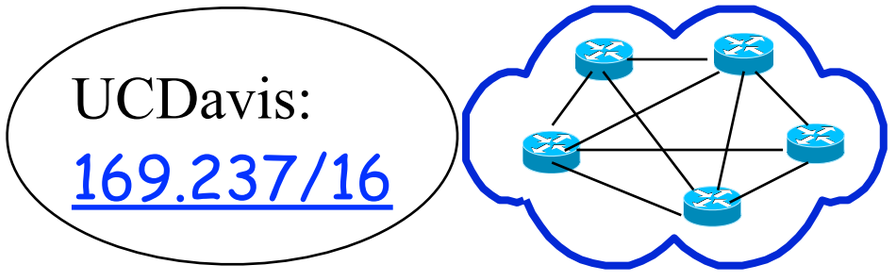

AS6192

"BGP"

- Autonomous System (AS):
  - A set of routers owned by one single system administrative domain
- Address Prefix:

- Example:
  - AS6192 consists of routers in UC Davis
  - UC Davis owns 169.237/16

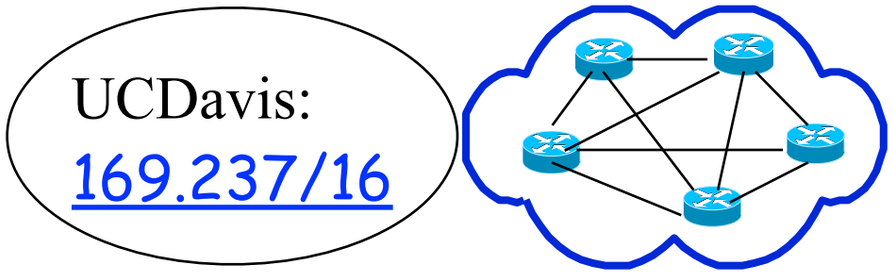

AS6192

"BGP"

- How would I let the whole world know about 169.237/16?
  - I announce that I owned 169.237/16

- More importantly, how would anybody else in the Internet know how to send (or route, forward) a IP packet to 169.237/16?
  - Others would know how to send packets to 169.237/16

Peering ASes

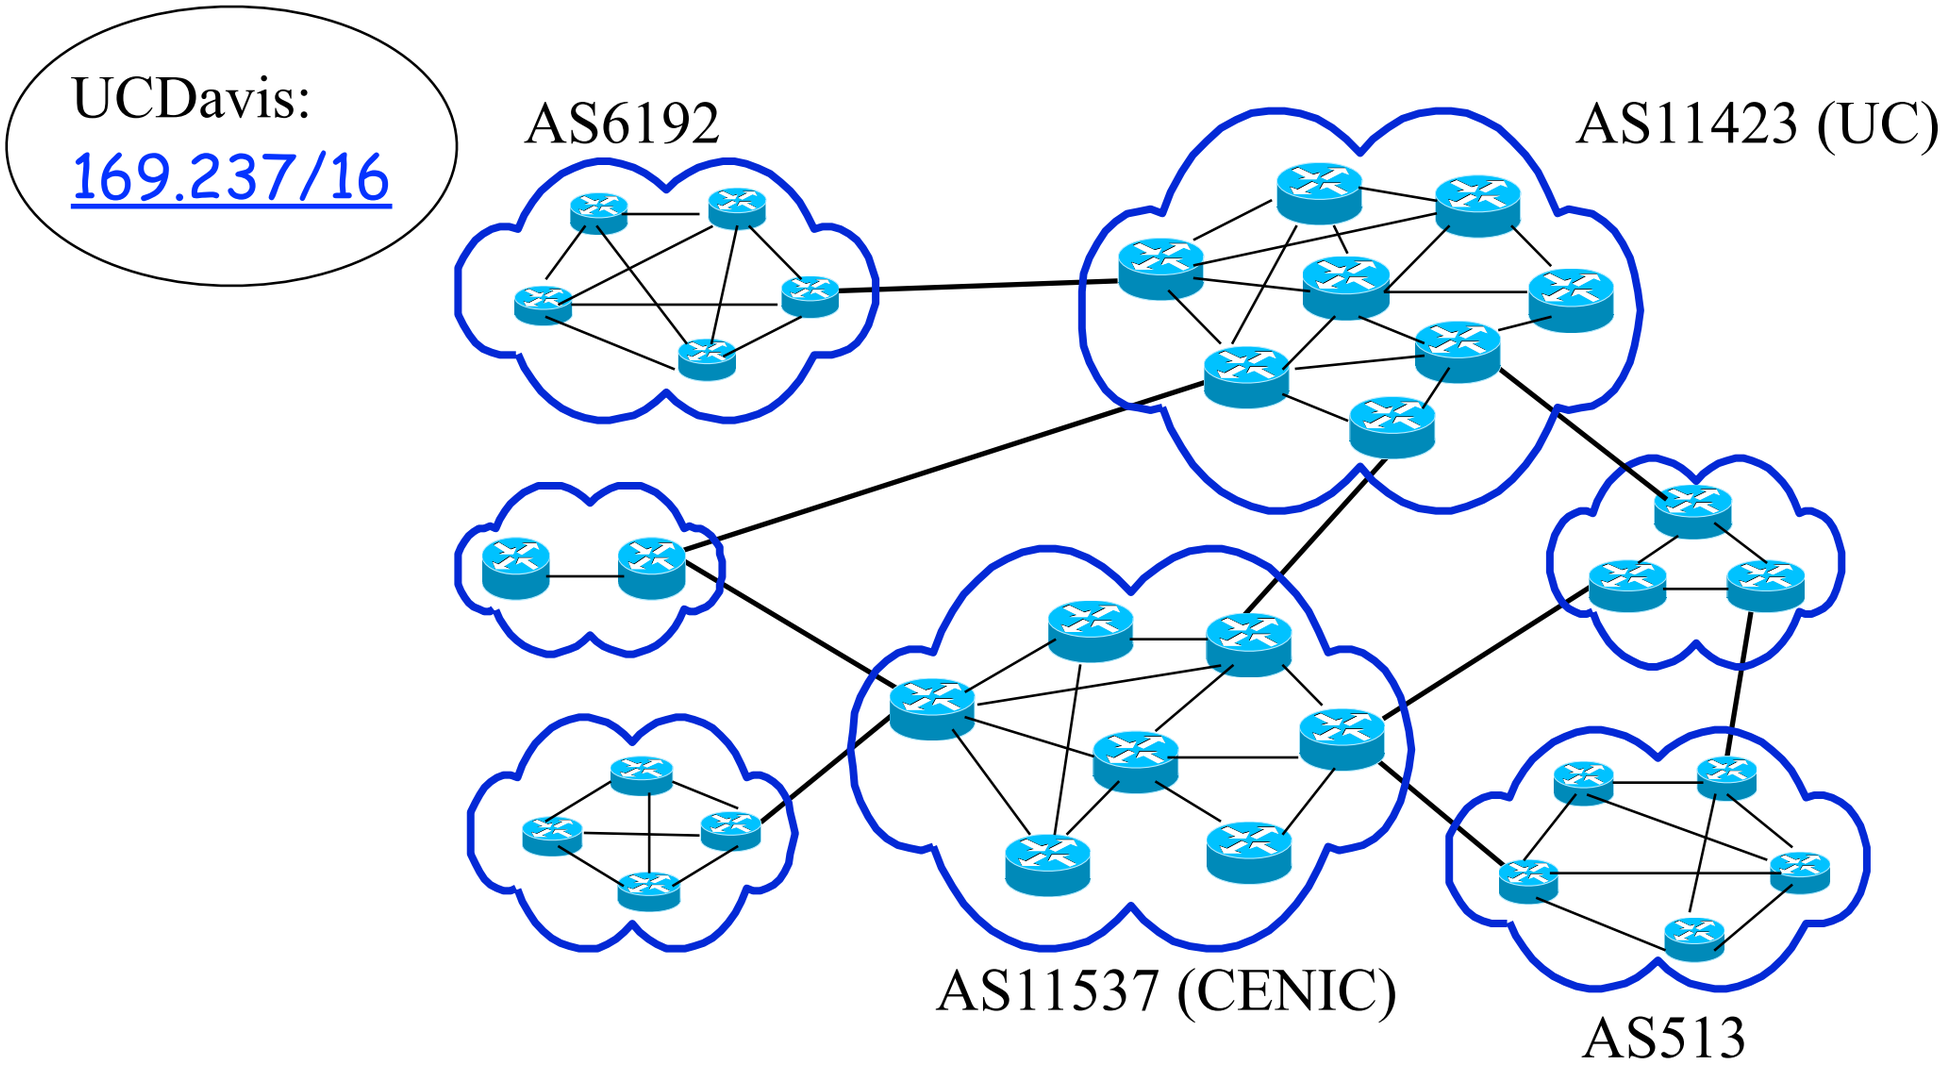

Peering is a local/decentralized trust based on a business contract!

AS6192

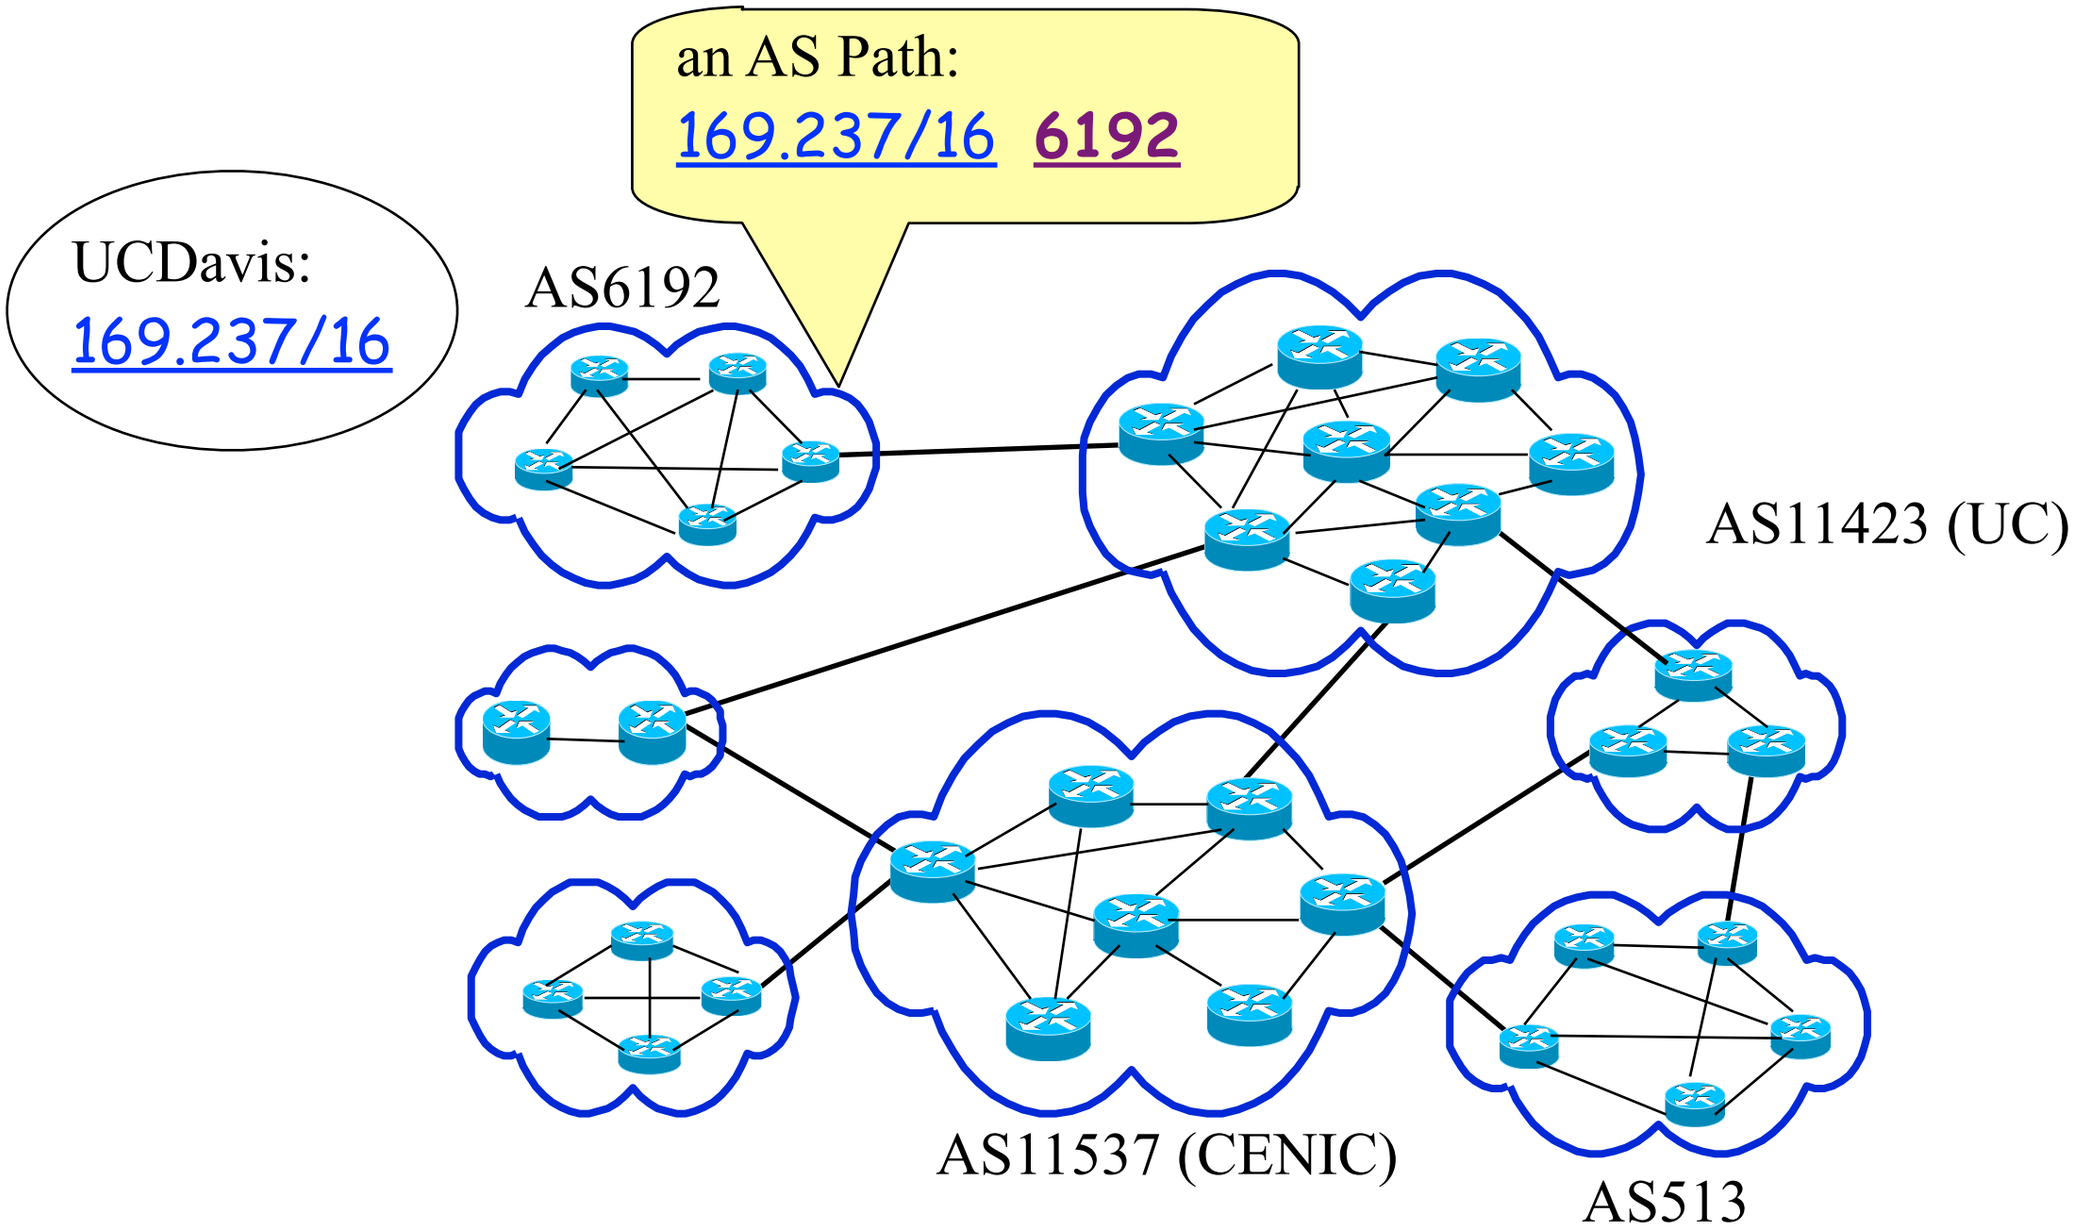
an AS Path:
169.237/16 6192
UCDavis:
169.237/16
AS6192
AS11423 (UC)
AS11537 (CENIC)
AS513

AS6192 -> AS11423

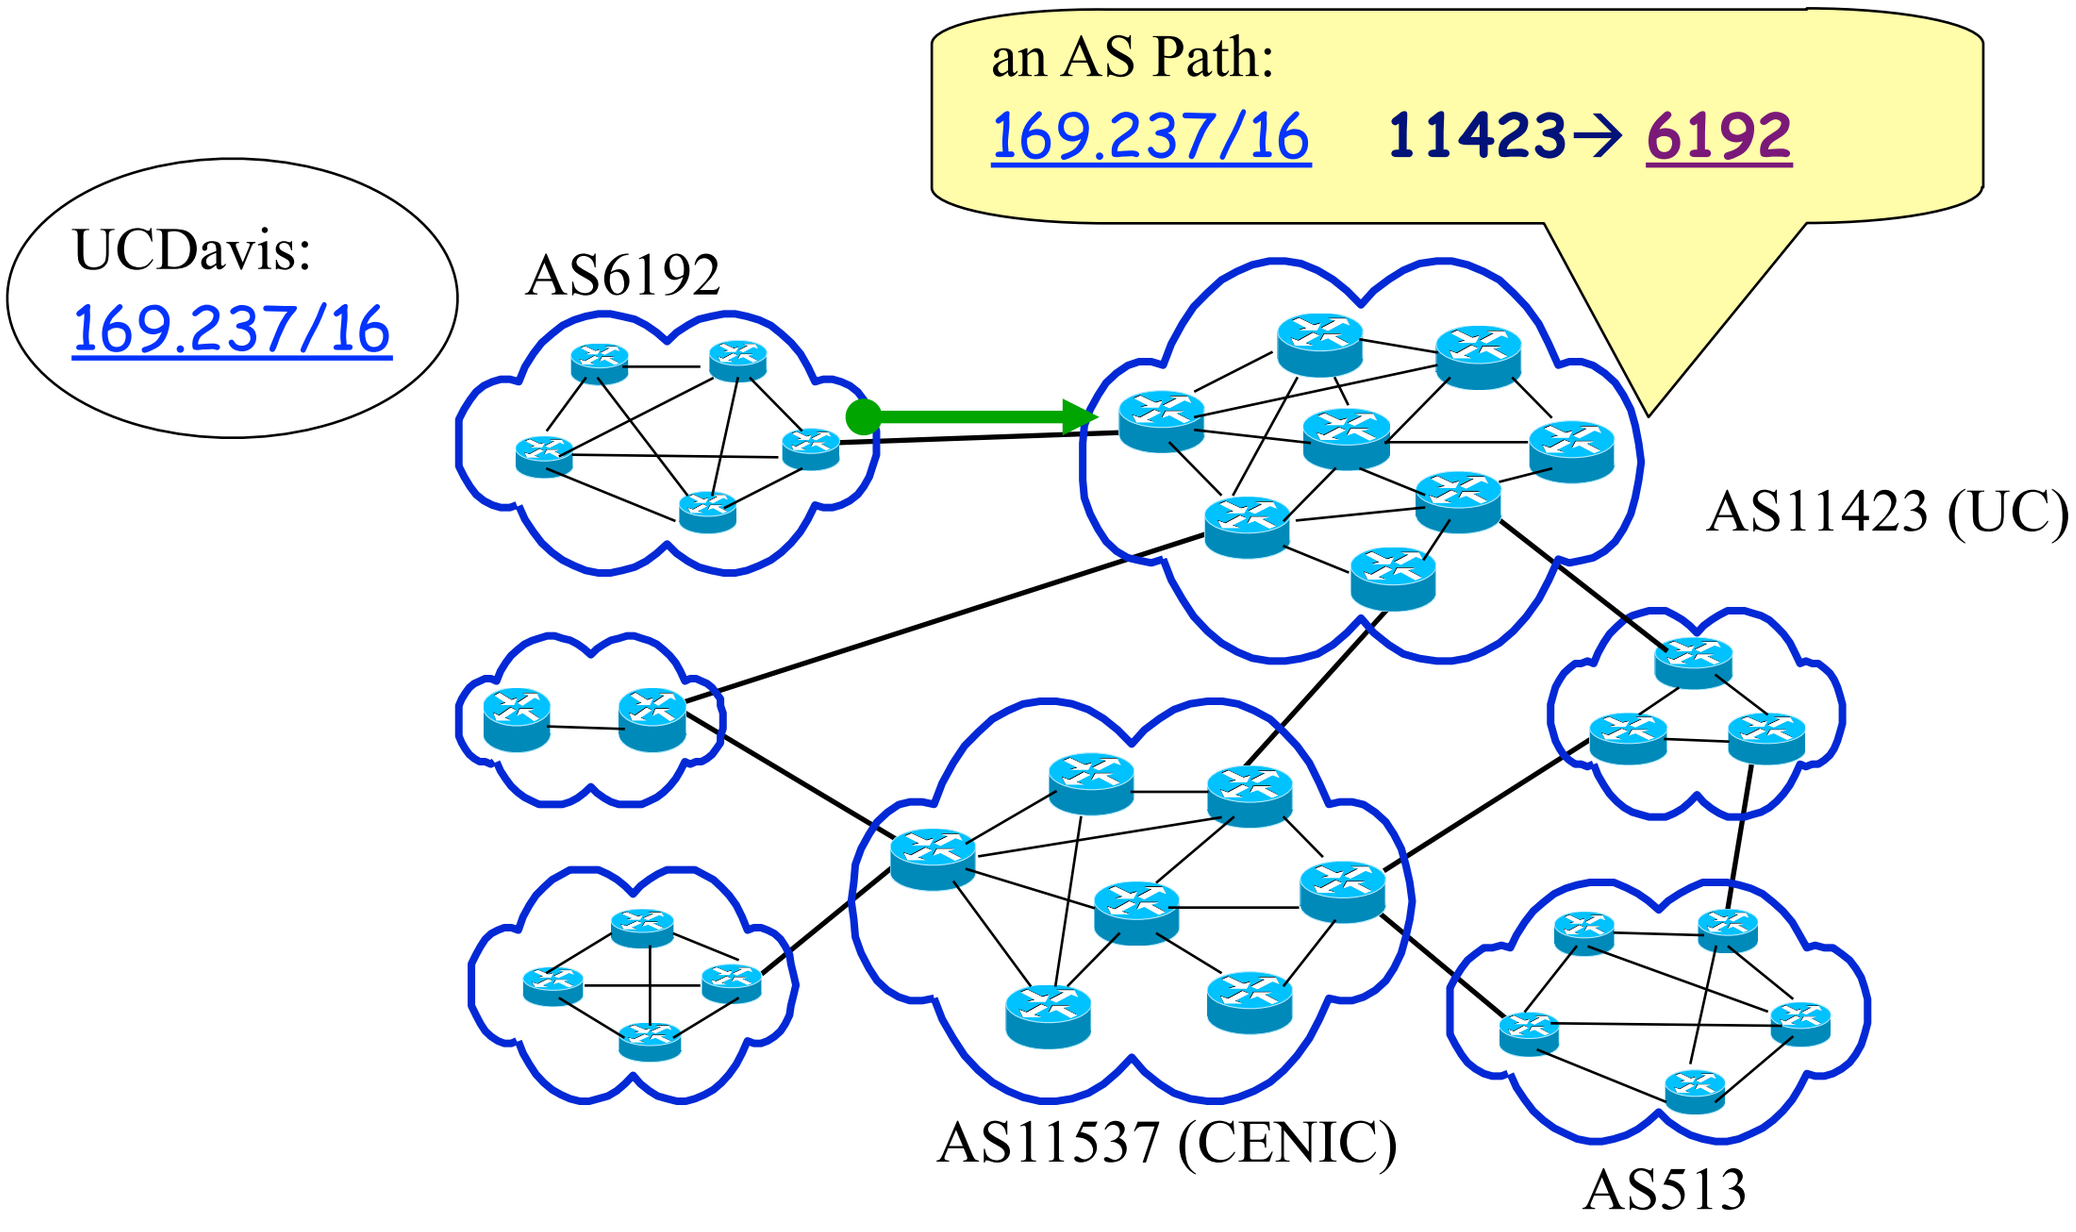

AS11423 -> AS11537

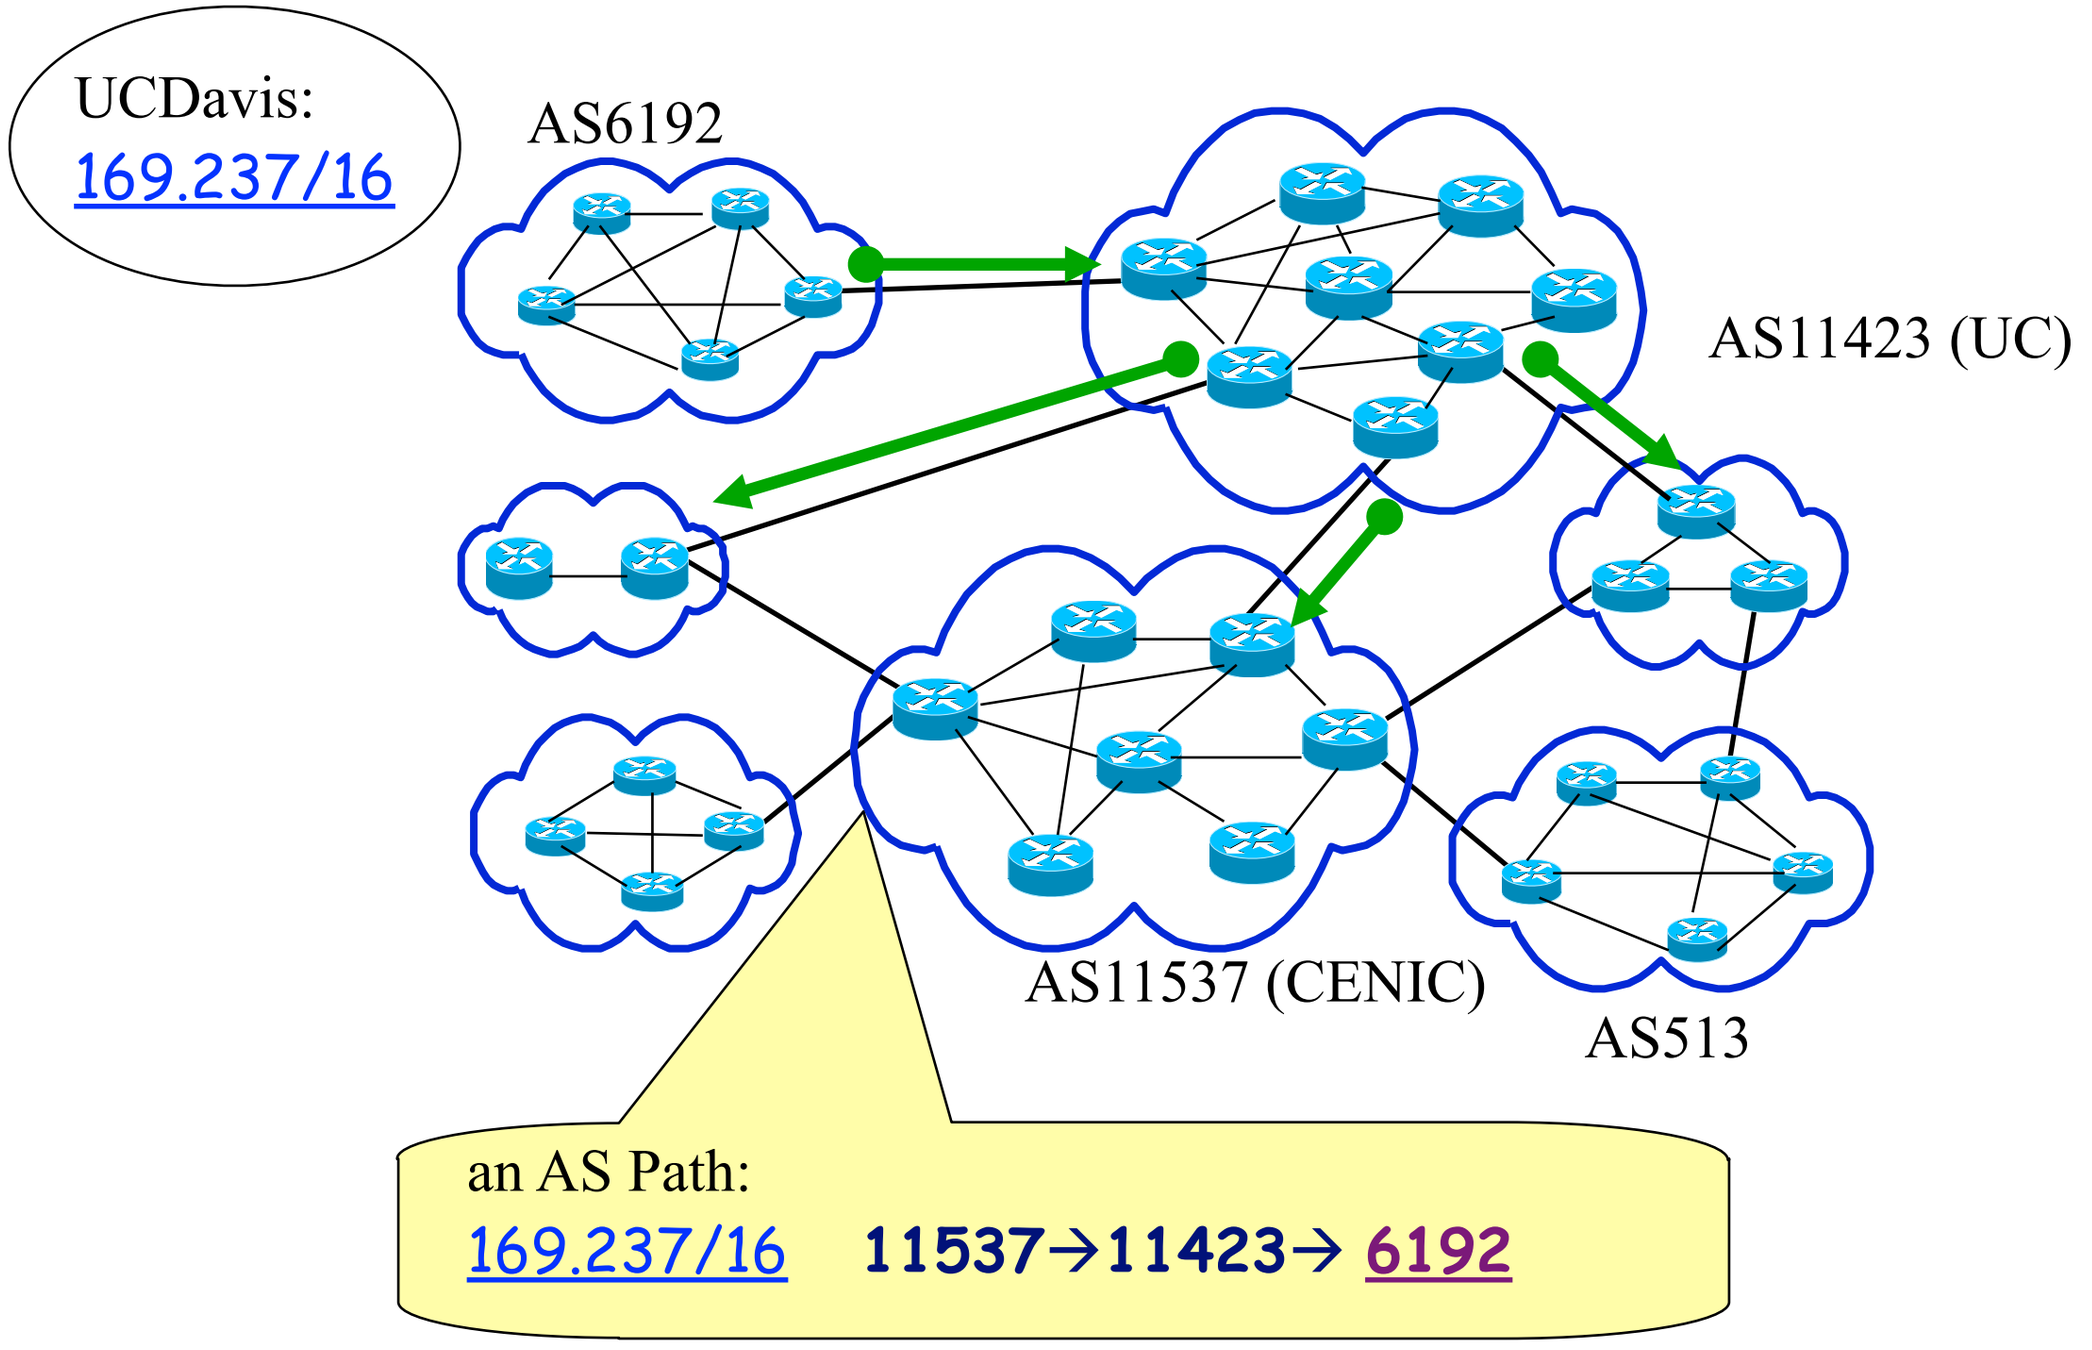

AS11537 -> AS513

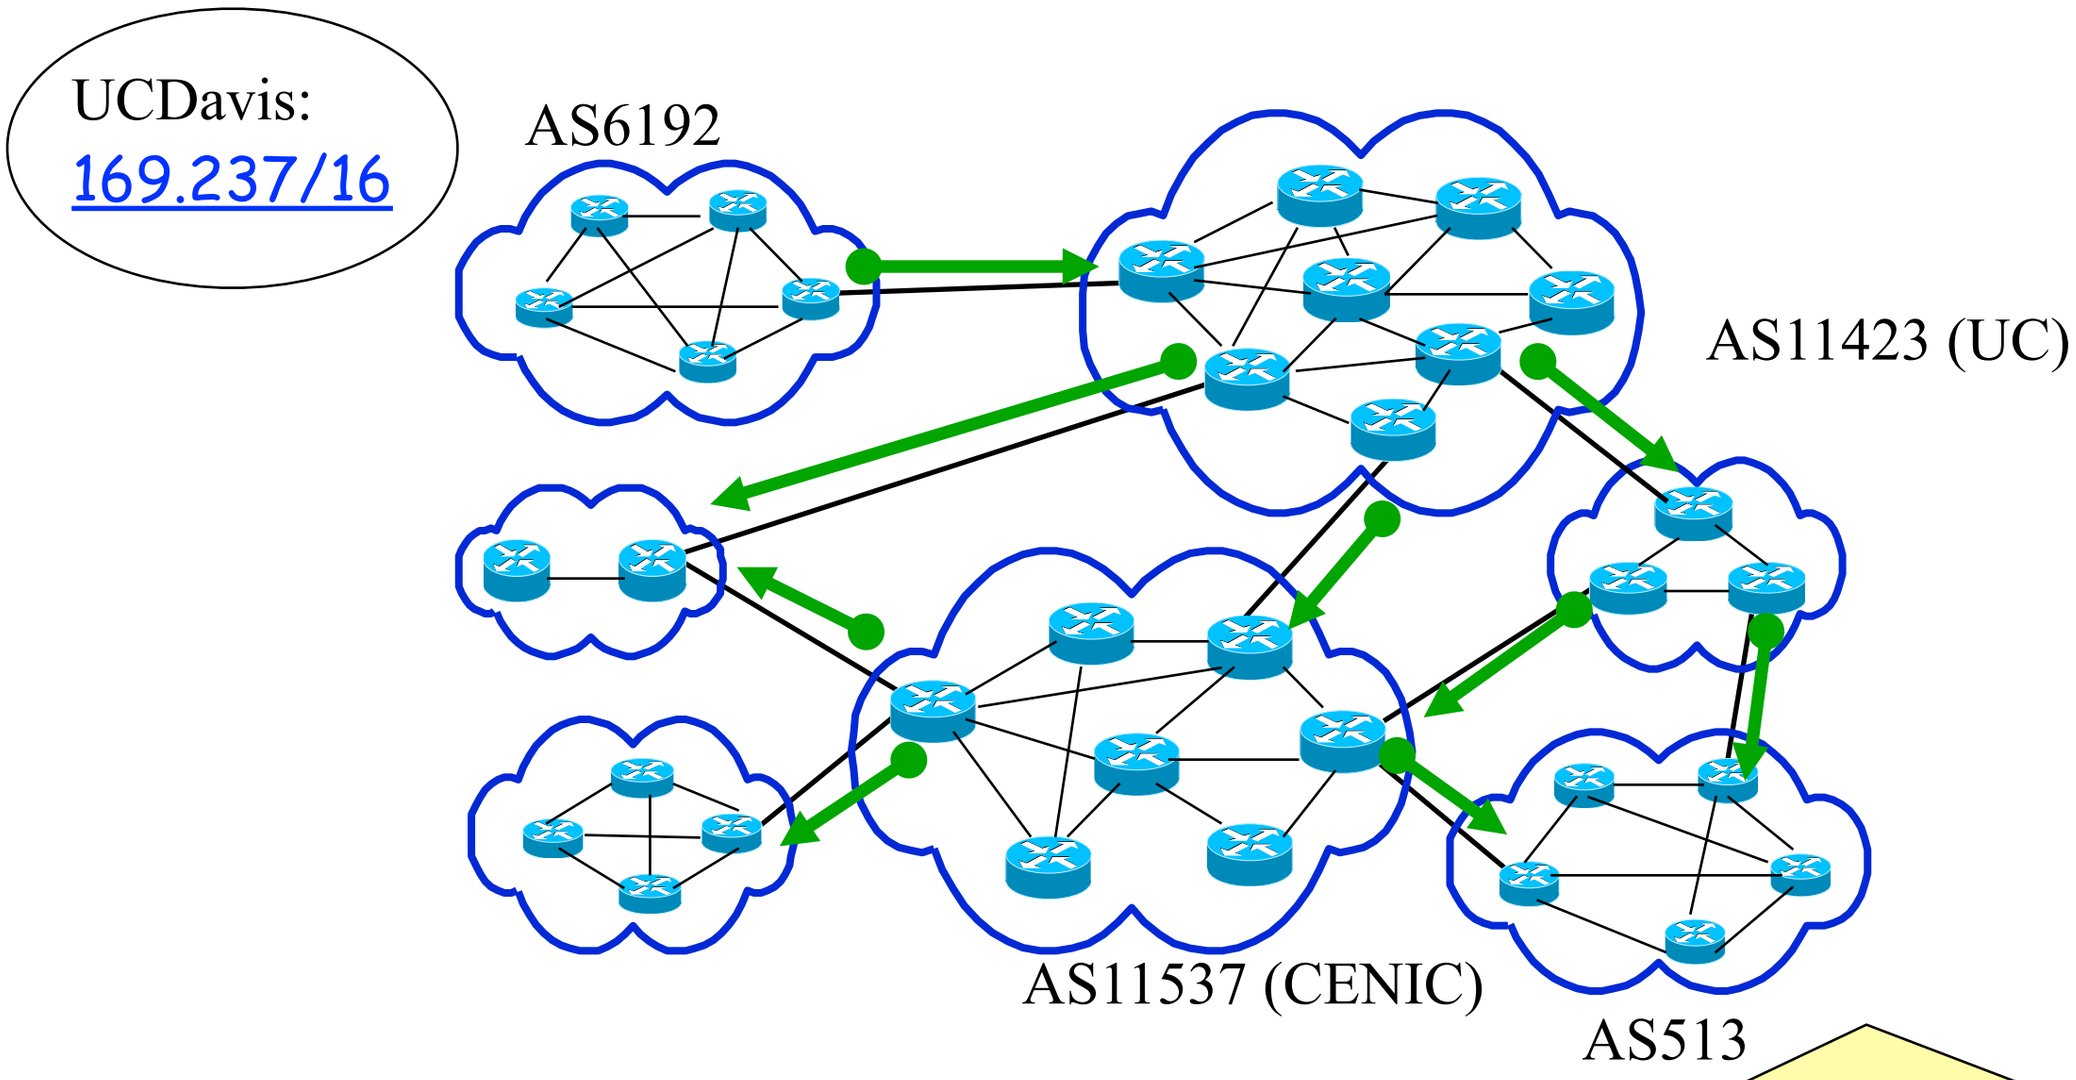

an AS Path:

169.237/16 **513→11537→11423→** 6192

Packet Forwarding

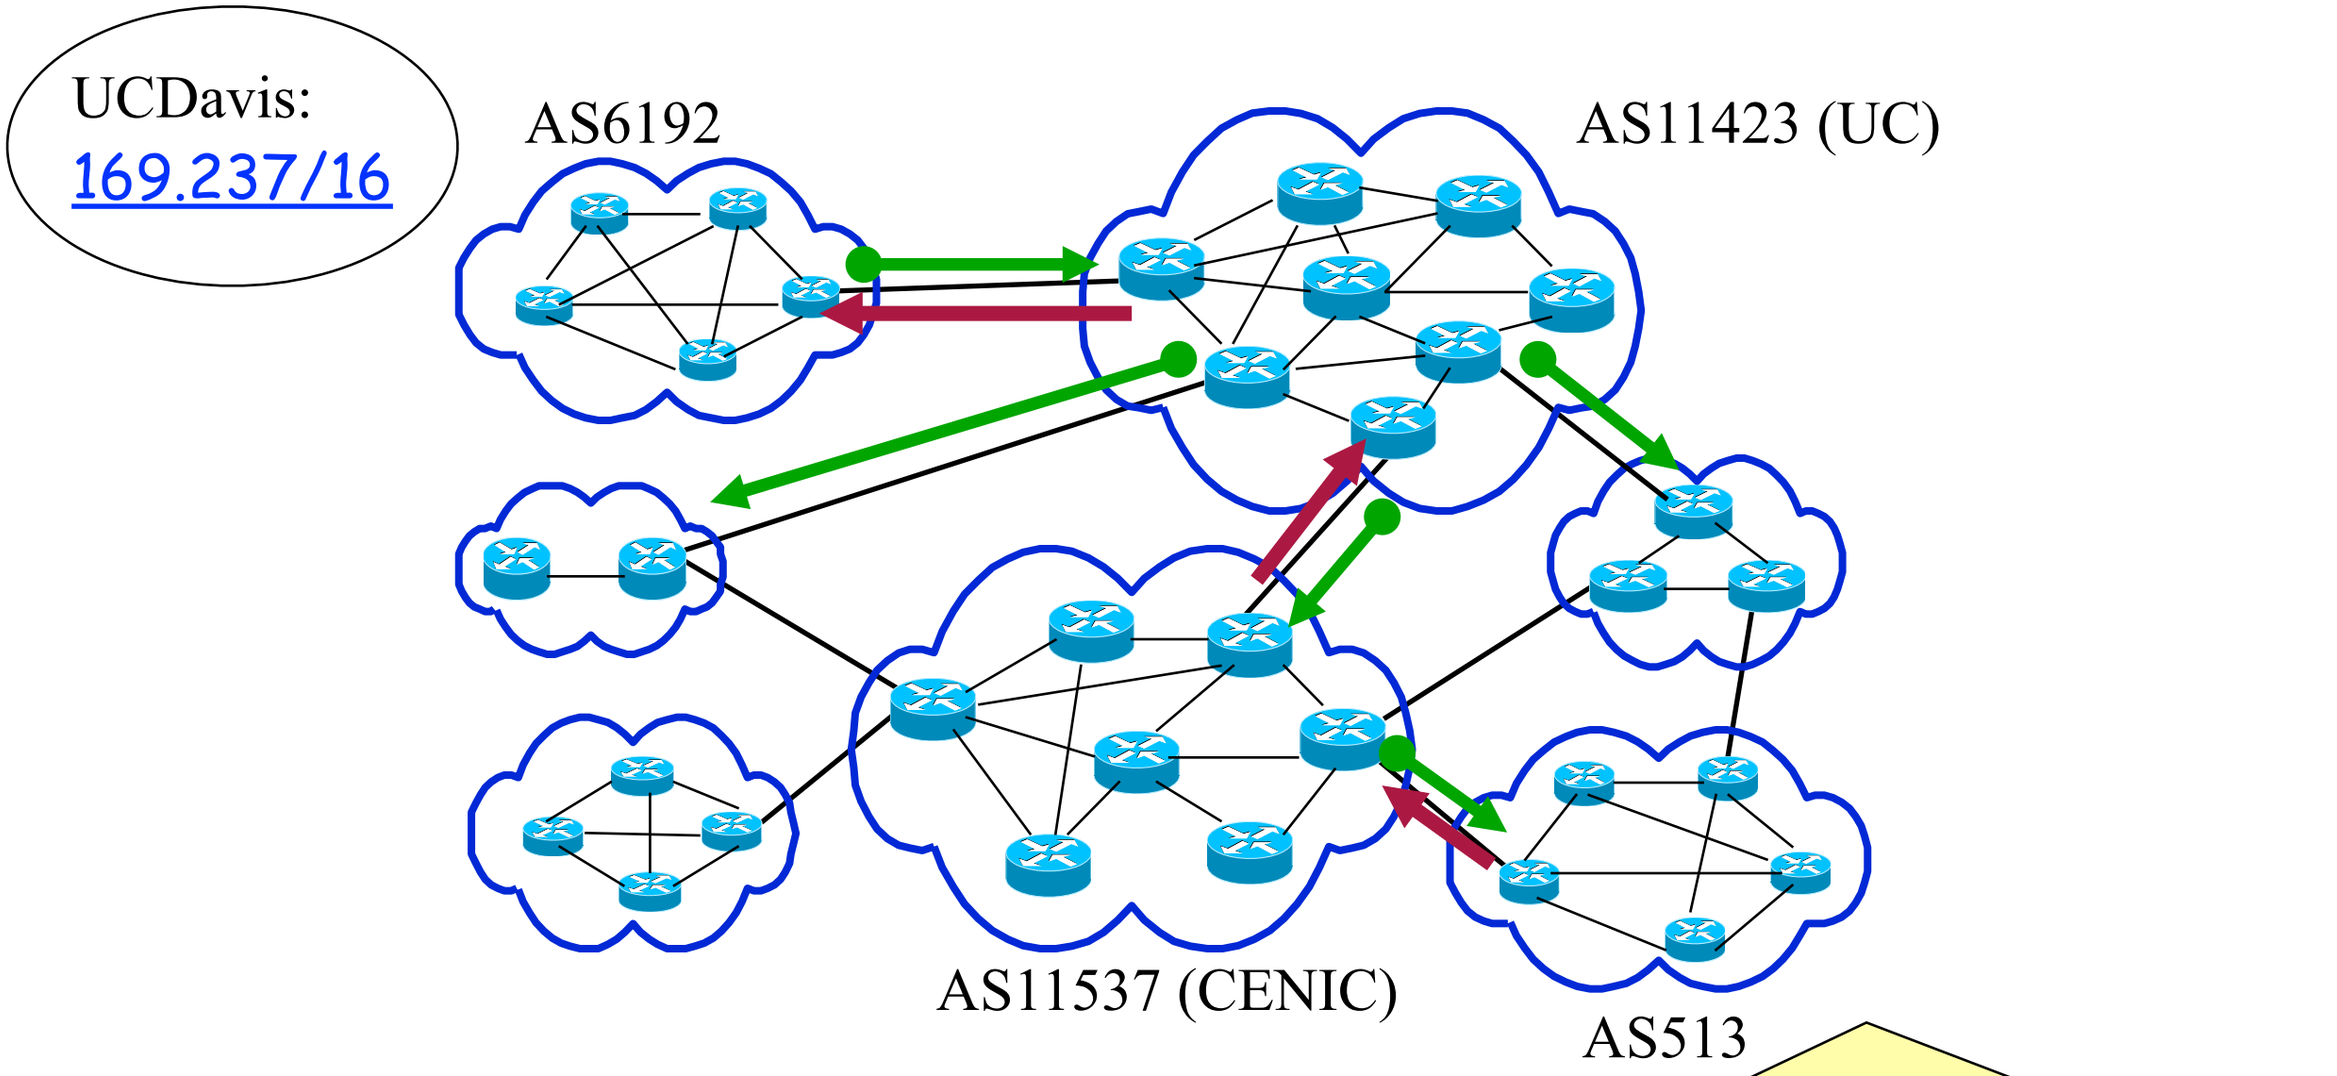

an AS Path:

169.237/16 **513→11537→11423→** 6192

What is Sharif's AS number?

12657

212.72.64.0/19

213.131.192.0/19

12660

81.31.160.0/19

213.233.160.0/19

12692

140.204.0.0/16

161.71.0.0/16

12711

212.48.224.0/19

212.48.228.0/24

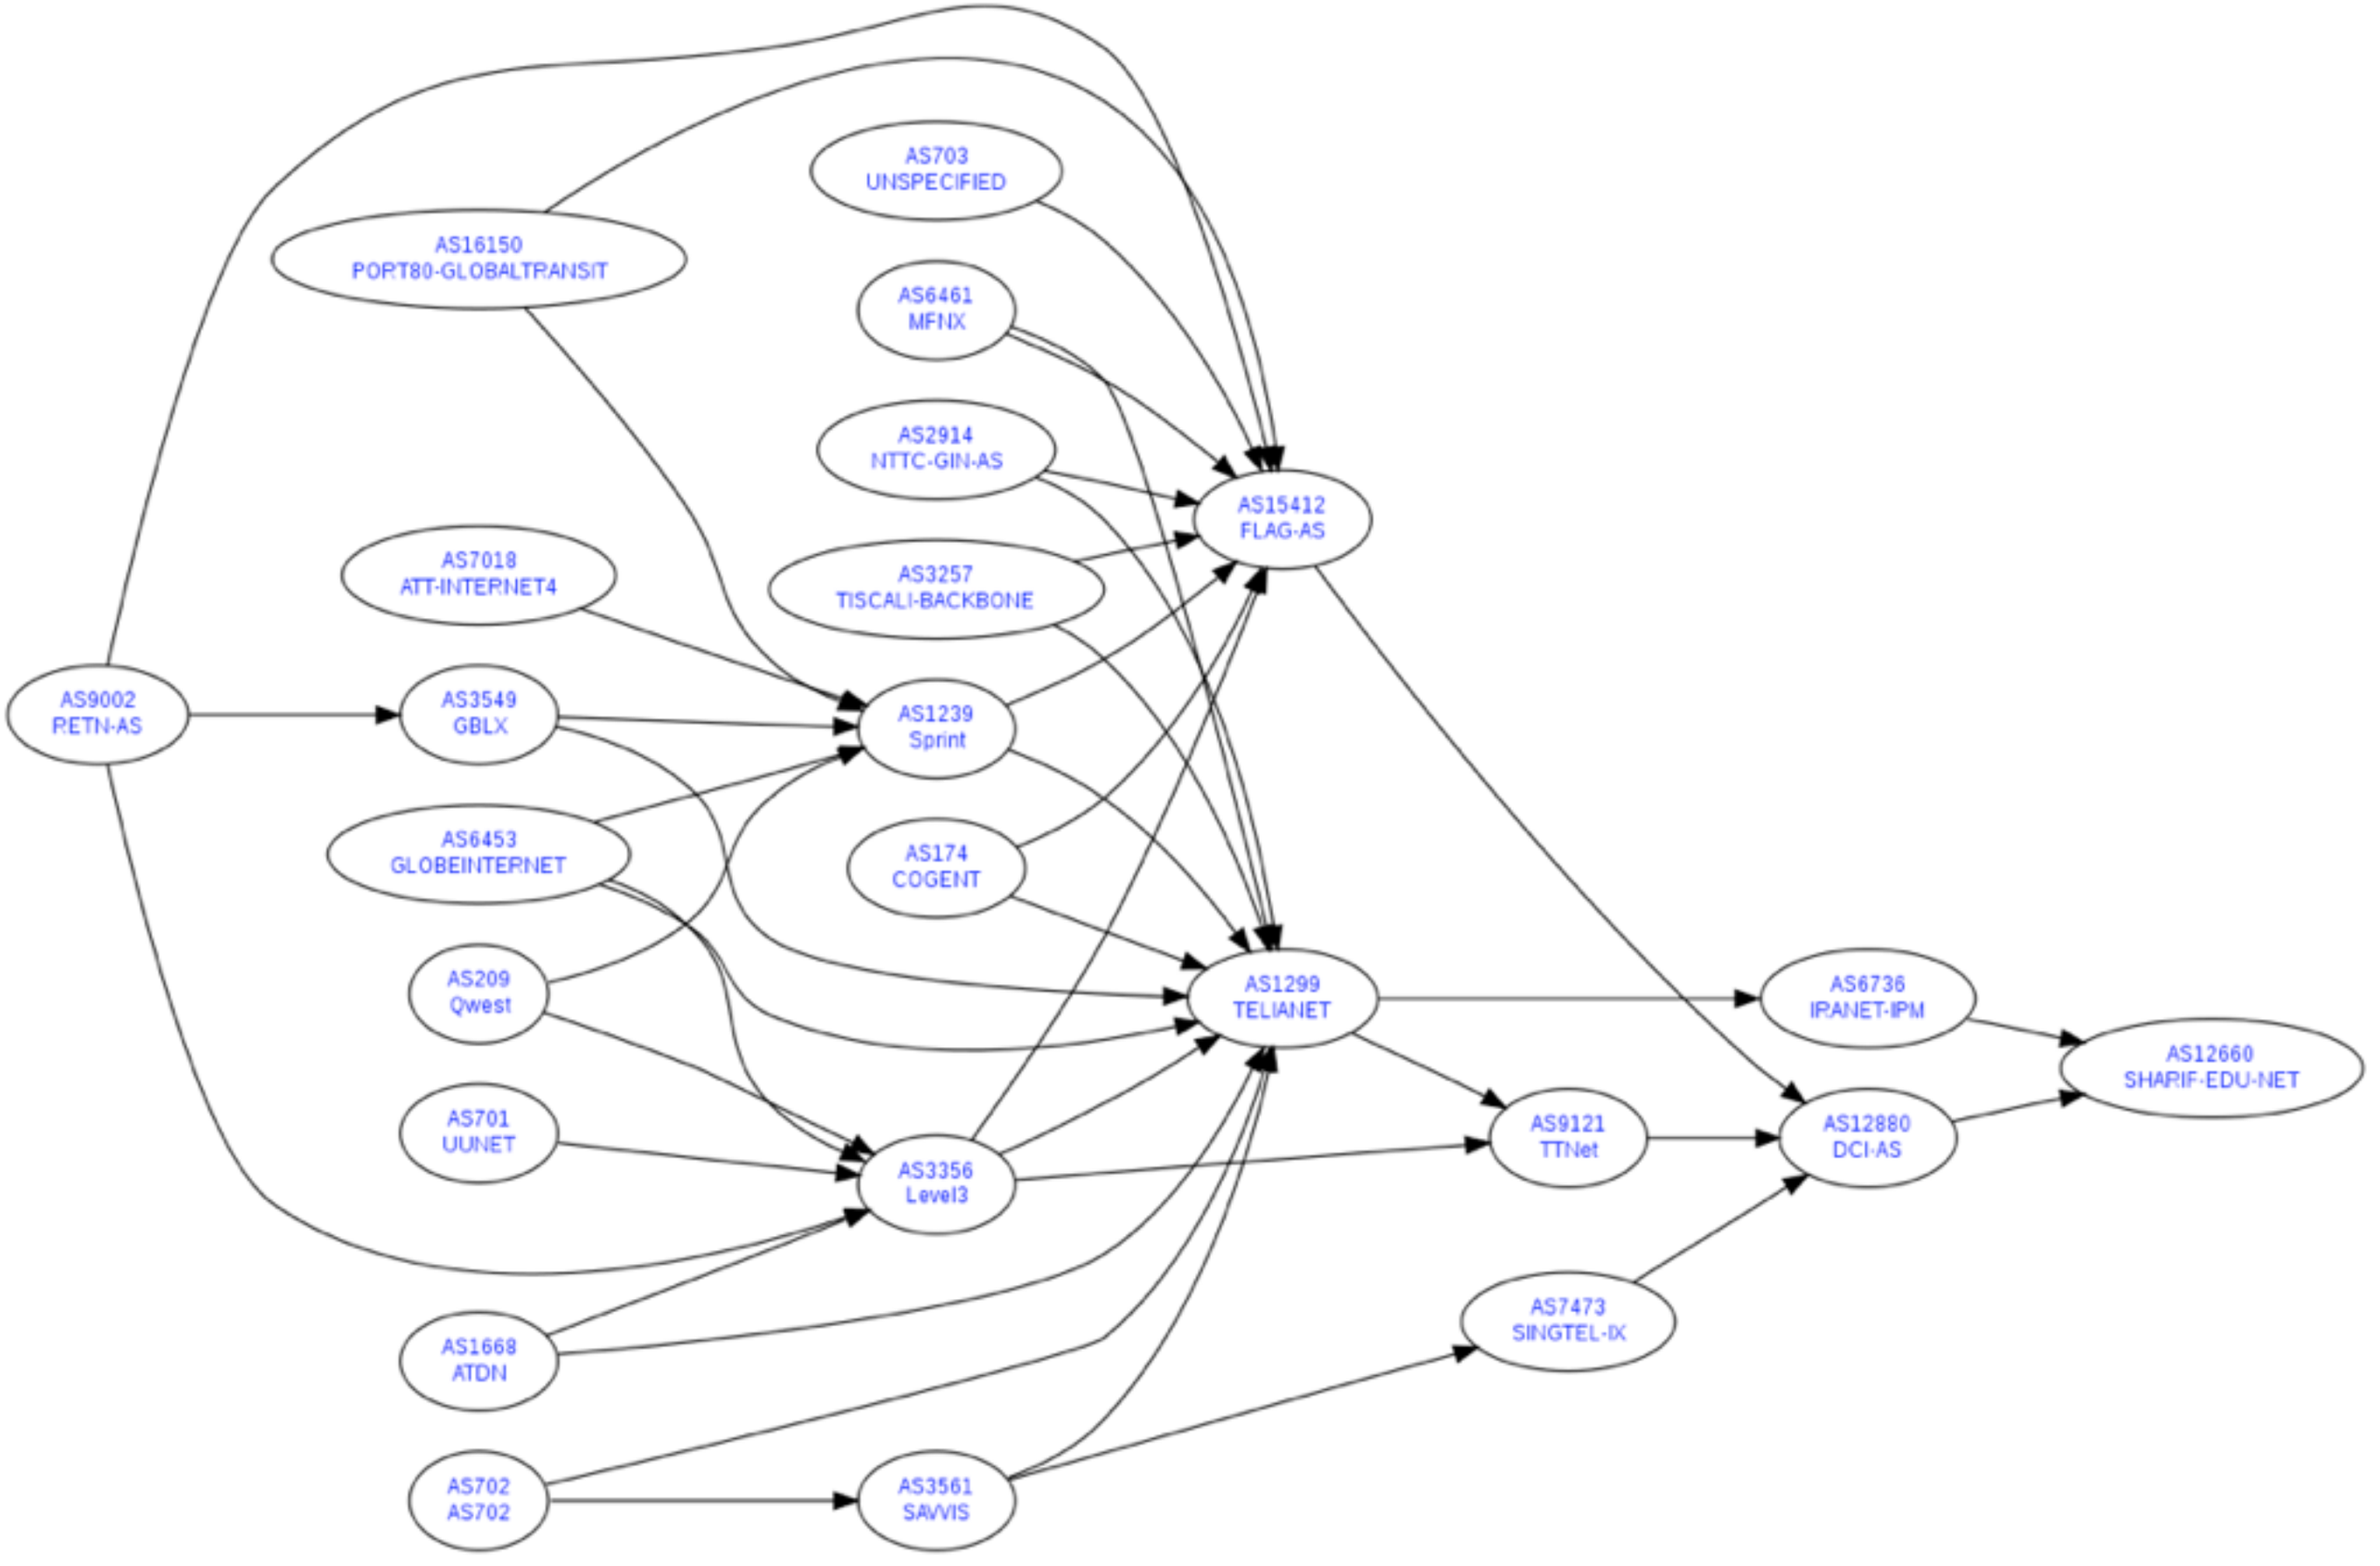
AS703
UNSPECIFIED
AS16150
PORT80-GLOBALTRANSIT
AS6461
MFNX
AS2914
NTTC-GIN-AS
AS15412
FLAG-AS
AS7018
ATT-INTERNET4
AS3257
TISCALI-BACKBONE
AS9002
RETN-AS
AS3549
GBLX
AS1239
Sprint
AS6453
GLOBEINTERNET
AS174
COGENT
AS209
Qwest
AS1299
TELIANET
AS6736
IRANET-IPM
AS12660
SHARIF-EDU-NET
AS701
UUNET
AS9121
TTNet
AS12880
DCI-AS
AS3356
Level3
AS7473
SINGTEL-IX
AS1668
ATDN
AS702
AS702
AS3561
SAVVIS

The Scale of the “Internet”

- 20464 Autonomous Systems
- 167138 IP Address Prefixes announced

- Every single AS must maintain the routing table such that it knows how to route the traffic toward any one of the 167138 prefixes to the right destination.

- BGP is the protocol to support the exchange of routing information for ALL prefixes in ALL ASes.

BGP Session Security

Security Goals for BGP

- Secure message exchange between neighbors
  - Confidential BGP message exchange
  - No denial of service
- Validity of the routing information
  - Origin authentication
    - Is the prefix owned by the AS announcing it?
  - AS path authentication
    - Is AS path the sequence of ASes the BGP update traversed?
  - AS path policy
    - Does the AS path adhere to the routing policies of each AS?
- Correspondence to the data path
  - Does the traffic follow the advertised AS path?

TCP Connection Underlying BGP Session

- BGP session runs over TCP
  - TCP connection between neighboring routers
  - BGP messages sent over TCP connection
  - Makes BGP vulnerable to attacks on TCP
- Main kinds of attacks
  - Against confidentiality: eavesdropping
  - Against integrity: tampering
  - Against performance: denial-of-service
- Main defenses
  - Message authentication or encryption
  - Limiting access to physical path between routers
  - Defensive filtering to block unexpected packets

Attacks Against Confidentiality

- Eavesdropping
  - Monitoring the messages on the BGP session
  - … by tapping the link(s) between the neighbors
- Reveals sensitive information
  - Inference of business relationships
  - Analysis of network stability
- Reasons why it may be hard
  - Challenging to tap the link
    - Often, eBGP session traverses just one link
    - … and may be hard to get access to tap it
  - Encryption may obscure message contents
    - BGP neighbors may run BGP over IPSec

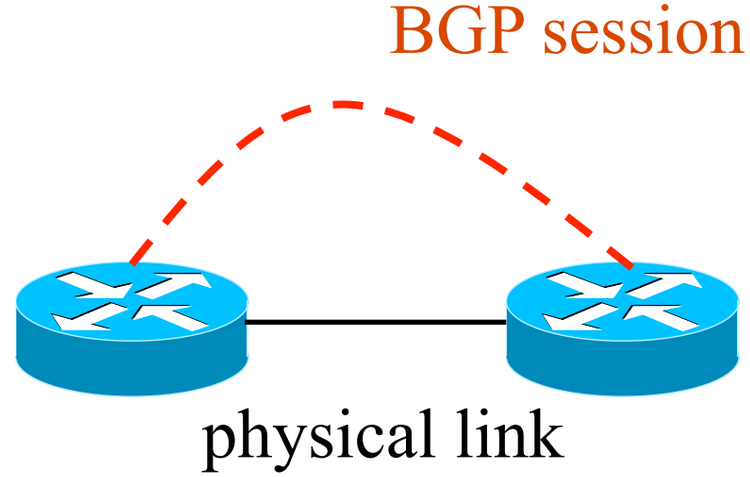

Attacking Message Integrity

- Tampering
  - Man-in-the-middle tampers with the messages
  - Insert, delete, modify, or replay messages
- Leads to incorrect BGP behavior
  - Delete: neighbor doesn't learn the new route
  - Insert/modify: neighbor learns bogus route
- Reasons why it may be hard
  - Getting in-between the two routers is hard
  - Use of authentication (signatures) or encryption
  - Spoofing TCP packets the right way is hard

Attacking Message Integrity

- Tampering
  - Man-in-the-middle tampers with the messages
  - Insert, delete, modify, or replay messages
- Leads to incorrect BGP behavior
  - Delete: neighbor doesn't learn the new route
  - Insert/modify: neighbor learns bogus route
- Reasons why it may be hard
  - Getting in-between the two routers is hard
  - Use of authentication (signatures) or encryption
  - Spoofing TCP packets the right way is hard
    - Getting past source-address packet filters
    - Generating the right TCP sequence number

Denial-of-Service Attacks, Part 1

- Overload the link between the routers
  - To cause packet loss and delay
  - … disrupting the performance of the BGP session
- Relatively easy to do
  - Can send traffic between end hosts
  - As long as the packets traverse the link
  - (which you can figure out from traceroute)
- Easy to defend
  - Give higher priority to BGP packets
  - E.g., by putting packets in separate queue

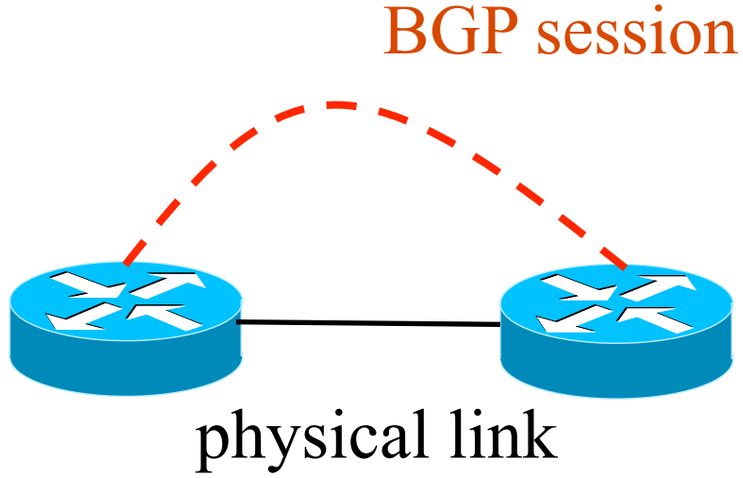

Denial-of-Service Attacks, Part 2

- Third party sends bogus TCP packets
  - FIN/RST to close the session
  - SYN flooding to overload the router
- Leads to disruptions in BGP
  - Session reset, causing transient routing changes
  - Route-flapping
- Reasons why it may be hard
  - Spoofing TCP packets the right way is hard
    - Difficult to send FIN/RST with the right TCP header
  - Packet filters may block the SYN flooding
    - Filter packets to BGP port from unexpected source
    - … or destined to router from unexpected source

Exploiting the IP TTL Field as a Defense

- BGP speakers are usually one hop apart
  - To thwart an attacker, can check that the packets carrying the BGP message have not traveled far
- IP Time-to-Live (TTL) field
  - Decremented once per hop
    - Avoids packets staying in network forever
- Generalized TTL Security Mechanism (RFC 3682)
  - Send BGP packets with initial TTL of 255
  - Receiving BGP speaker checks that TTL is 254
  - … and flags and/or discards the packet others
- Hard for third-party to inject packets remotely

Validity of the routing information:

Origin authentication

IP Address Ownership and Hijacking

- IP address block assignment
  - Regional Internet Registries (ARIN, RIPE, APNIC)
  - Internet Service Providers
- Proper origination of a prefix into BGP
  - By the AS who owns the prefix
  - … or, by its upstream provider(s) in its behalf
- However, what's to stop someone else?
  - Prefix hijacking: another AS originates the prefix
  - BGP does not verify that the AS is authorized
  - Registries of prefix ownership are inaccurate

Prefix Hijacking

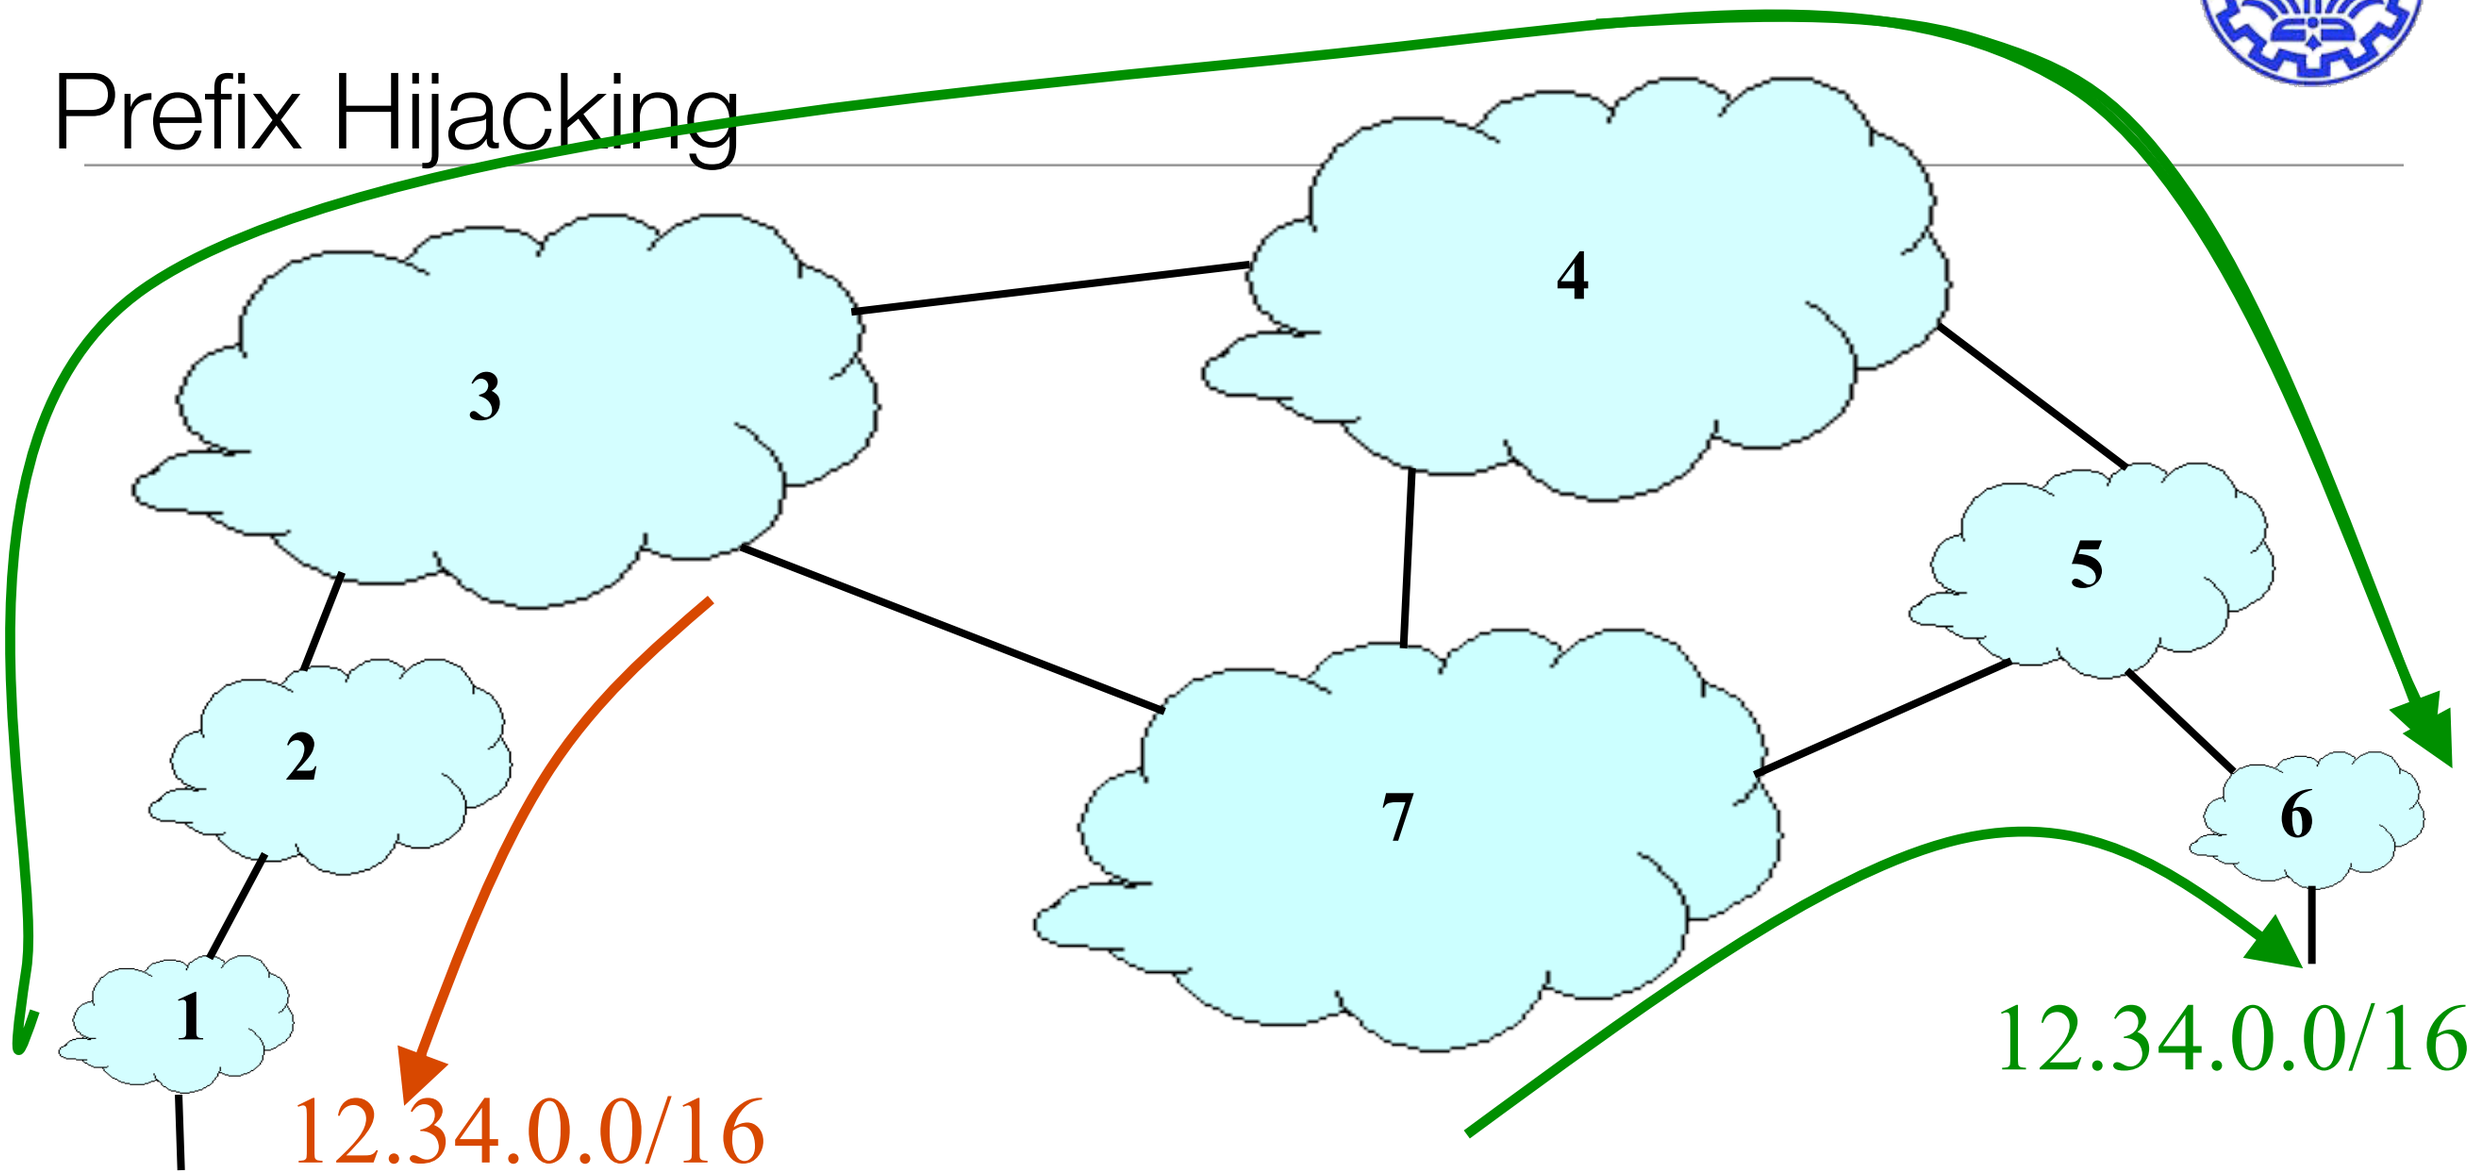

- Consequences for the affected ASes
  - Blackhole: data traffic is discarded
  - Snooping: data traffic is inspected, and then redirected
  - Impersonation: data traffic is sent to bogus destinations

Hijacking is Hard to Debug

- Real origin AS doesn’t see the problem
  - Picks its own route
  - Might not even learn the bogus route
- May not cause loss of connectivity
  - E.g., if the bogus AS snoops and redirects
    - … may only cause performance degradation
- Or, loss of connectivity is isolated
  - E.g., only for sources in parts of the Internet
- Diagnosing prefix hijacking
  - Analyzing updates from many vantage points
  - Launching traceroute from many vantage points

Sub-Prefix Hijacking

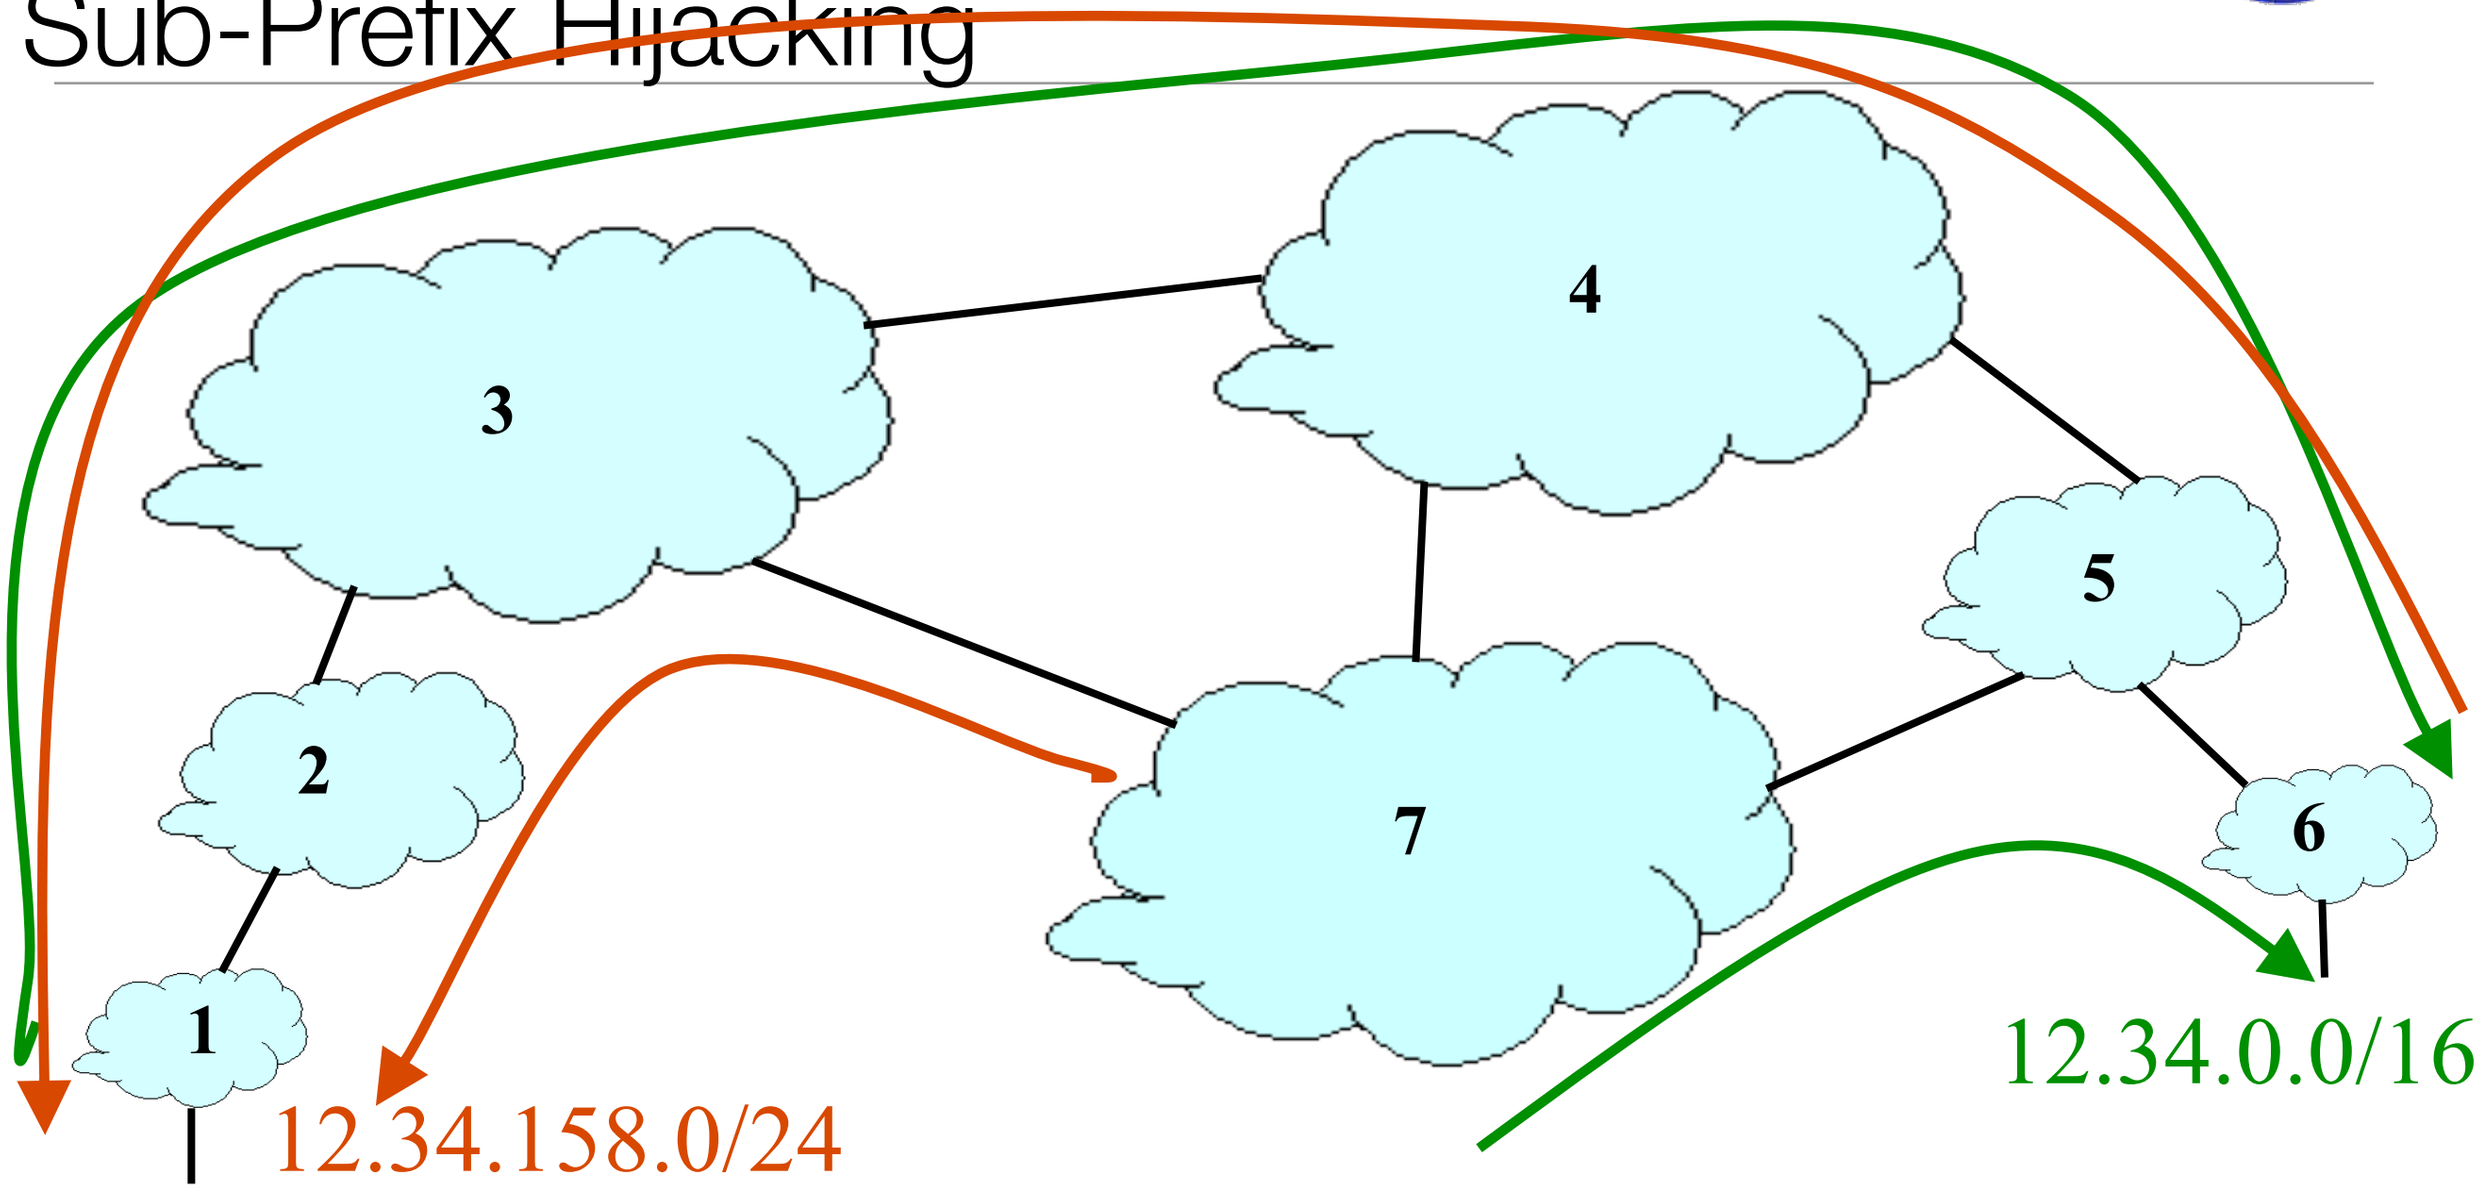

- Originating a more-specific prefix
  - Every AS picks the bogus route for that prefix
  - Traffic follows the longest matching prefix

How to Hijack a Prefix

- The hijacking AS has
  - Router with eBGP session(s)
  - Configured to originate the prefix
- Getting access to the router
  - Network operator makes configuration mistake
  - Disgruntled operator launches an attack
  - Outsider breaks in to the router and reconfigures
- Getting other ASes to believe bogus route
  - Neighbor ASes not filtering the routes
  - … e.g., by allowing only expected prefixes
  - But, specifying filters on peering links is hard

The February 24 YouTube Outage

- YouTube (AS 36561)
  - Web site www.youtube.com
  - Address block 208.65.152.0/22
- Pakistan Telecom (AS 17557)
  - Receives government order to block access to YouTube
  - Starts announcing 208.65.153.0/24 to PCCW (AS 3491)
  - All packets directed to YouTube get dropped on the floor
- Mistakes were made
  - AS 17557: announcing to everyone, not just customers
  - AS 3491: not filtering routes announced by AS 17557
- Lasted 100 minutes for some, 2 hours for others

Timeline (UTC Time)

- 18:47:45
  - First evidence of hijacked /24 route propagating in Asia
- 18:48:00
  - Several big trans-Pacific providers carrying the route
- 18:49:30
  - Bogus route fully propagated
- 20:07:25
  - YouTube starts advertising the /24 to attract traffic back
- 20:08:30
  - Many (but not all) providers are using the valid route

http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml

Timeline (UTC Time)

- 20:18:43
  - YouTube starts announcing two more-specific /25 routes
- 20:19:37
  - Some more providers start using the /25 routes
- 20:50:59
  - AS 17557 starts prepending (“3491 17557 17557”)
- 20:59:39
  - AS 3491 disconnects AS 17557
- 21:00:00
  - All is well, videos of cats, monkeys, etc. doing foolish things can be watched again.

http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml

Another Example: Spammers

- Spammers sending spam
  - Form a (bidirectional) TCP connection to a mail server
  - Send a bunch of spam e-mail
  - Disconnect and laugh all the way to the bank
- But, best not to use your real IP address
  - Relatively easy to trace back to you
- Could hijack someone's address space
  - But you might not receive all the (TCP) return traffic
  - And the legitimate owner of the address might notice
- How to evade detection
  - Hijack unused (i.e., unallocated) address block in BGP
  - Temporarily use the IP addresses to send your spam

Security Goals for BGP

- Secure message exchange between neighbors
  - Confidential BGP message exchange
  - No denial of service
- Validity of the routing information
  - Origin authentication
    - Is the prefix owned by the AS announcing it?
  - AS path authentication
    - Is AS path the sequence of ASes the BGP update traversed?
  - AS path policy
    - Does the AS path adhere to the routing policies of each AS?
- Correspondence to the data path
  - Does the traffic follow the advertised AS path?

Acknowledgments/References

- [Bellovin06] COMS W4180 — Network Security Class Columbia University, Steven Bellovin, 2006.
- [Wu07] ecs 236, Intrusion Detection, S. Felix Wu, UC Davis, Winter 2007.
- [Rex05] COS 461, Computer Networks, Jennifer Rex, Princeton University, Spring 2005.