Networking Review & Grand Challenges

Brighten Godfrey

CS 538 January 22 2018

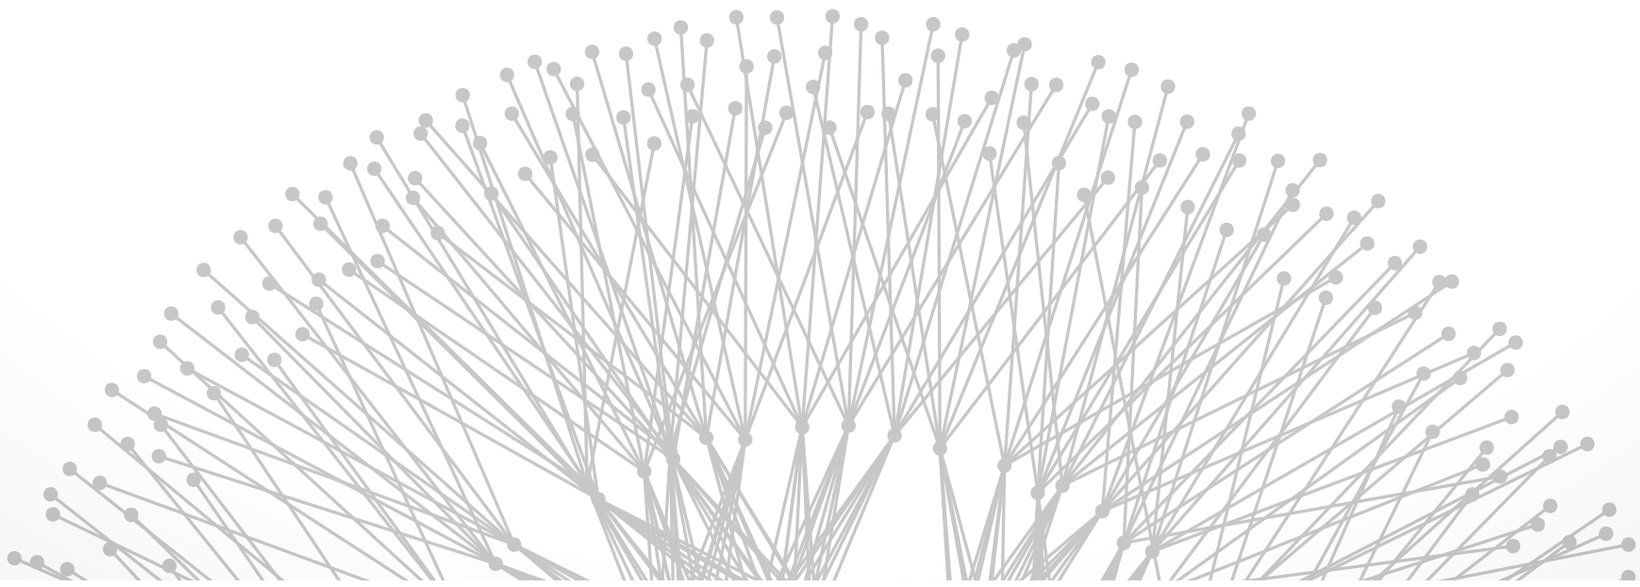

Announcements

Introducing Sangeetha

Key dates posted

- Assignment release, presentations, …
- Of note: Assignment 1 out Wed., due in a week
- Reading topics still subject to adjustment

No laptop use in lecture please

Undergraduate Networking in Three Slides

(including this one)

Layering

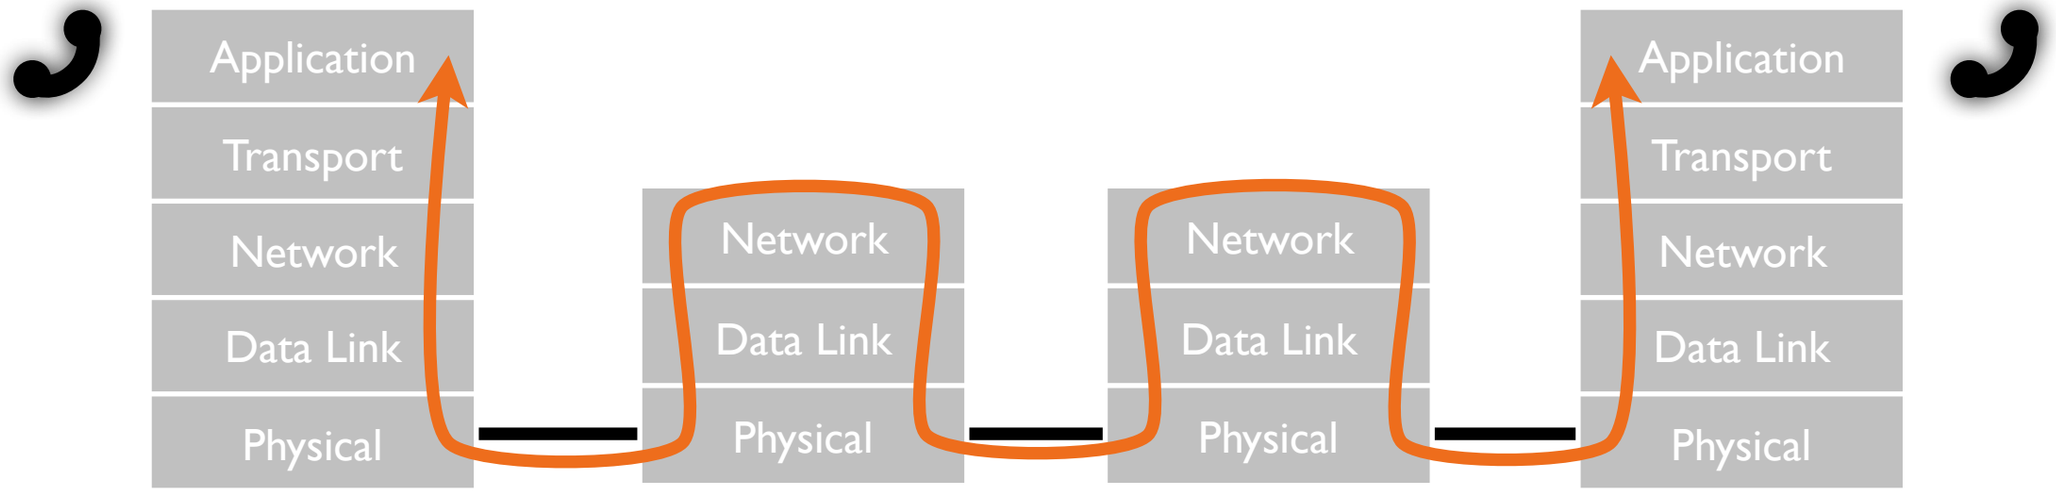

A kind of modularity

Functionality separated into layers

- Layer *n* implements higher-level functionality by interfacing only with layer *n*-1
- Hides complexity of surrounding layers: enables greater diversity and evolution of modules

Layering

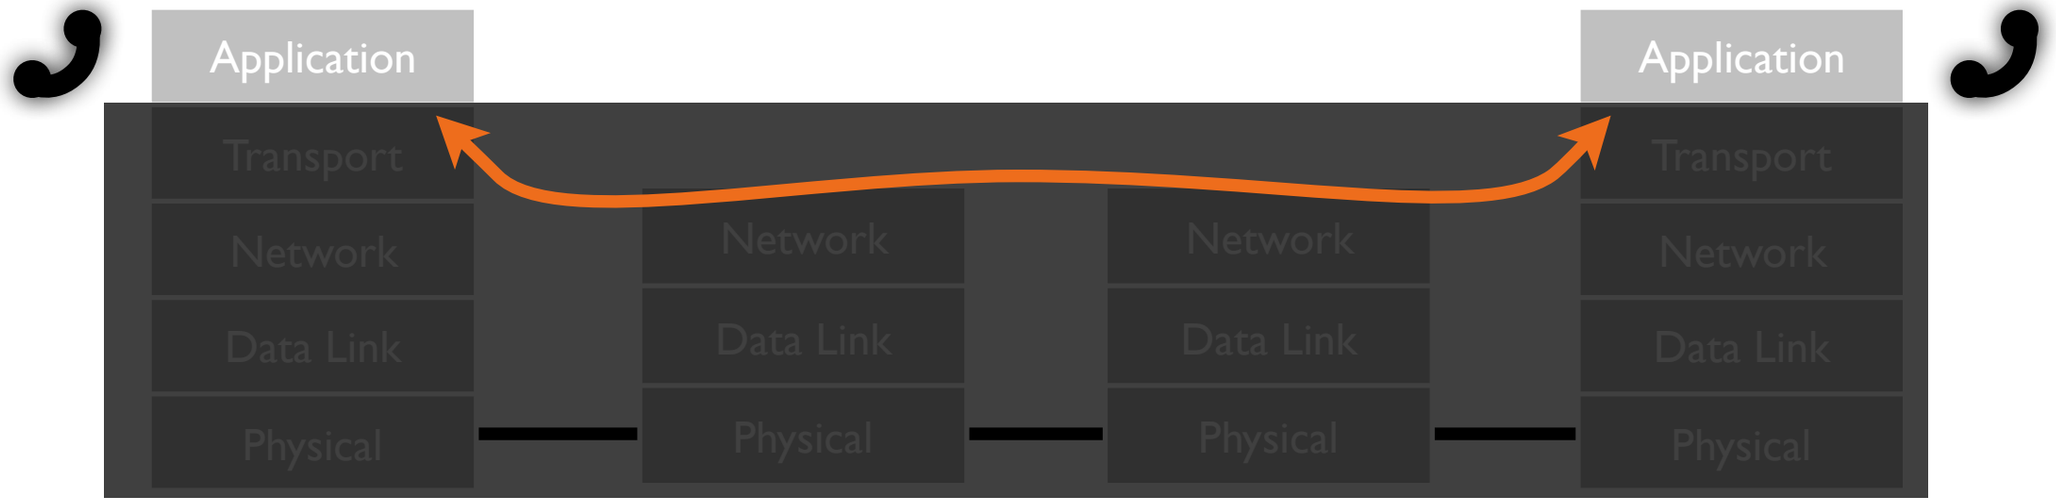

A kind of modularity

Functionality separated into layers

- Layer *n* implements higher-level functionality by interfacing only with layer *n*-1
- Hides complexity of surrounding layers: enables greater diversity and evolution of modules

Layering

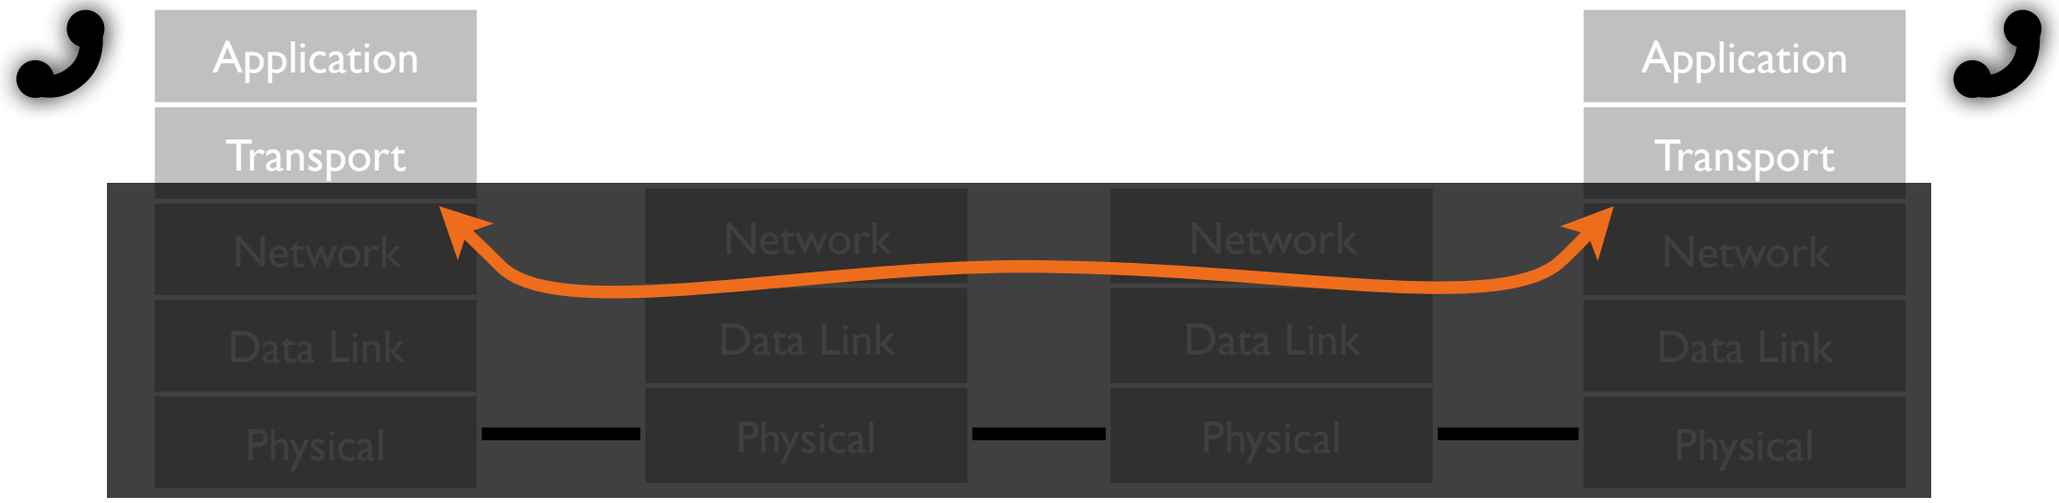

A kind of modularity

Functionality separated into layers

- Layer *n* implements higher-level functionality by interfacing only with layer *n*-1
- Hides complexity of surrounding layers: enables greater diversity and evolution of modules

Layering

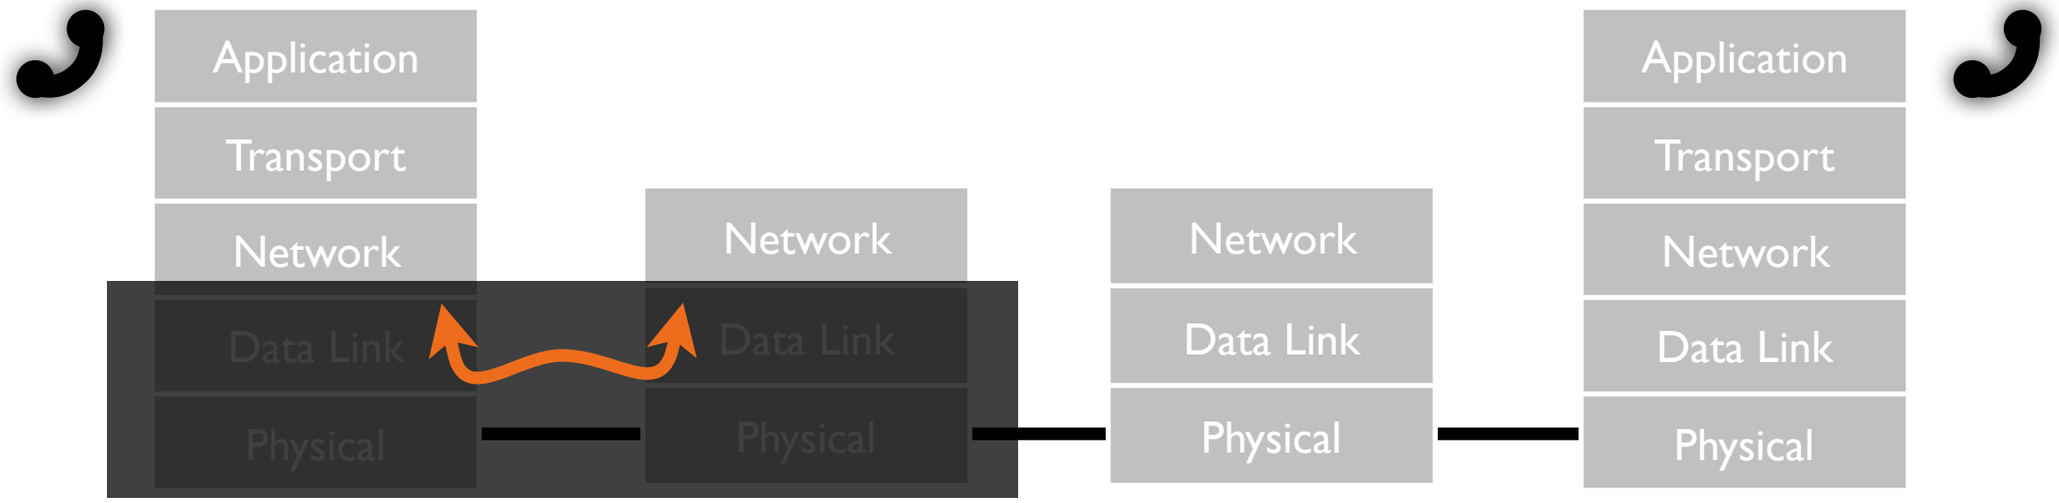

A kind of modularity

Functionality separated into layers

- Layer *n* implements higher-level functionality by interfacing only with layer *n*-1
- Hides complexity of surrounding layers: enables greater diversity and evolution of modules

Layering

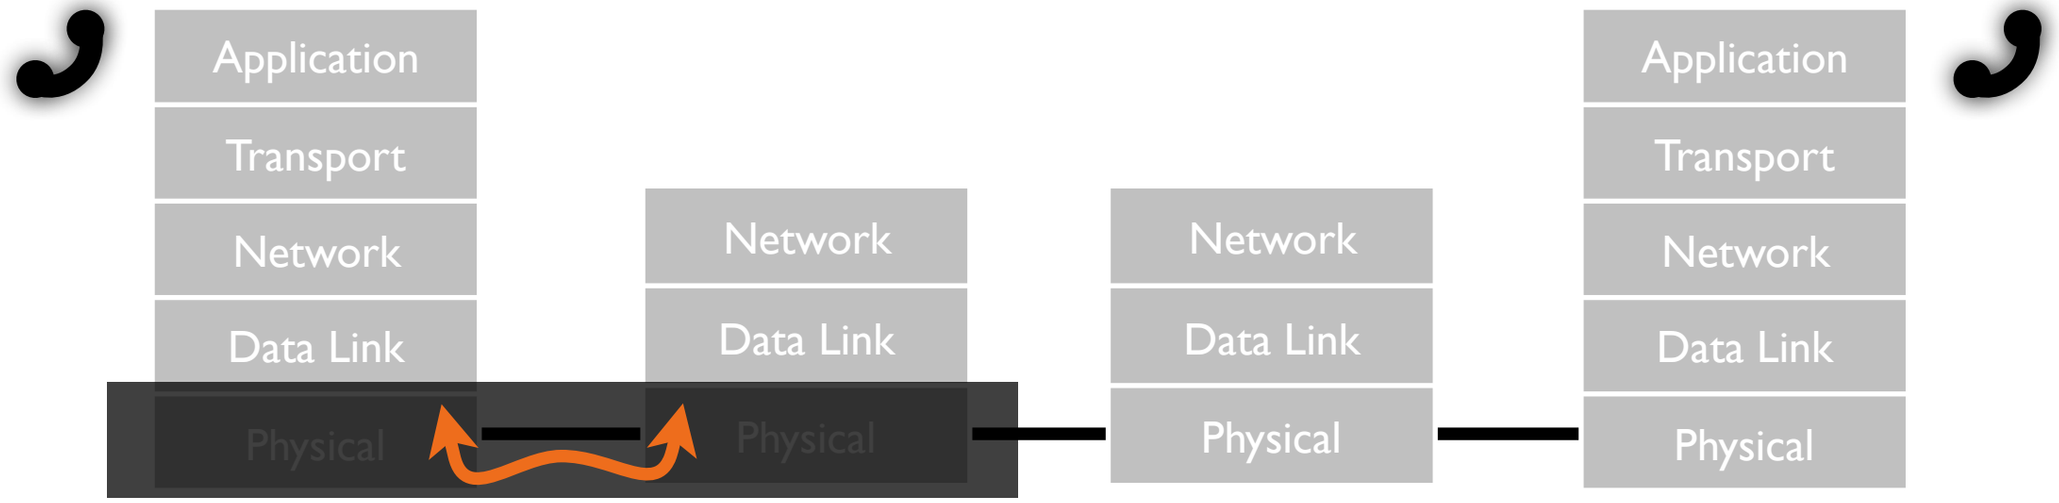

A kind of modularity

Functionality separated into layers

- Layer *n* implements higher-level functionality by interfacing only with layer *n*-1
- Hides complexity of surrounding layers: enables greater diversity and evolution of modules

Layering

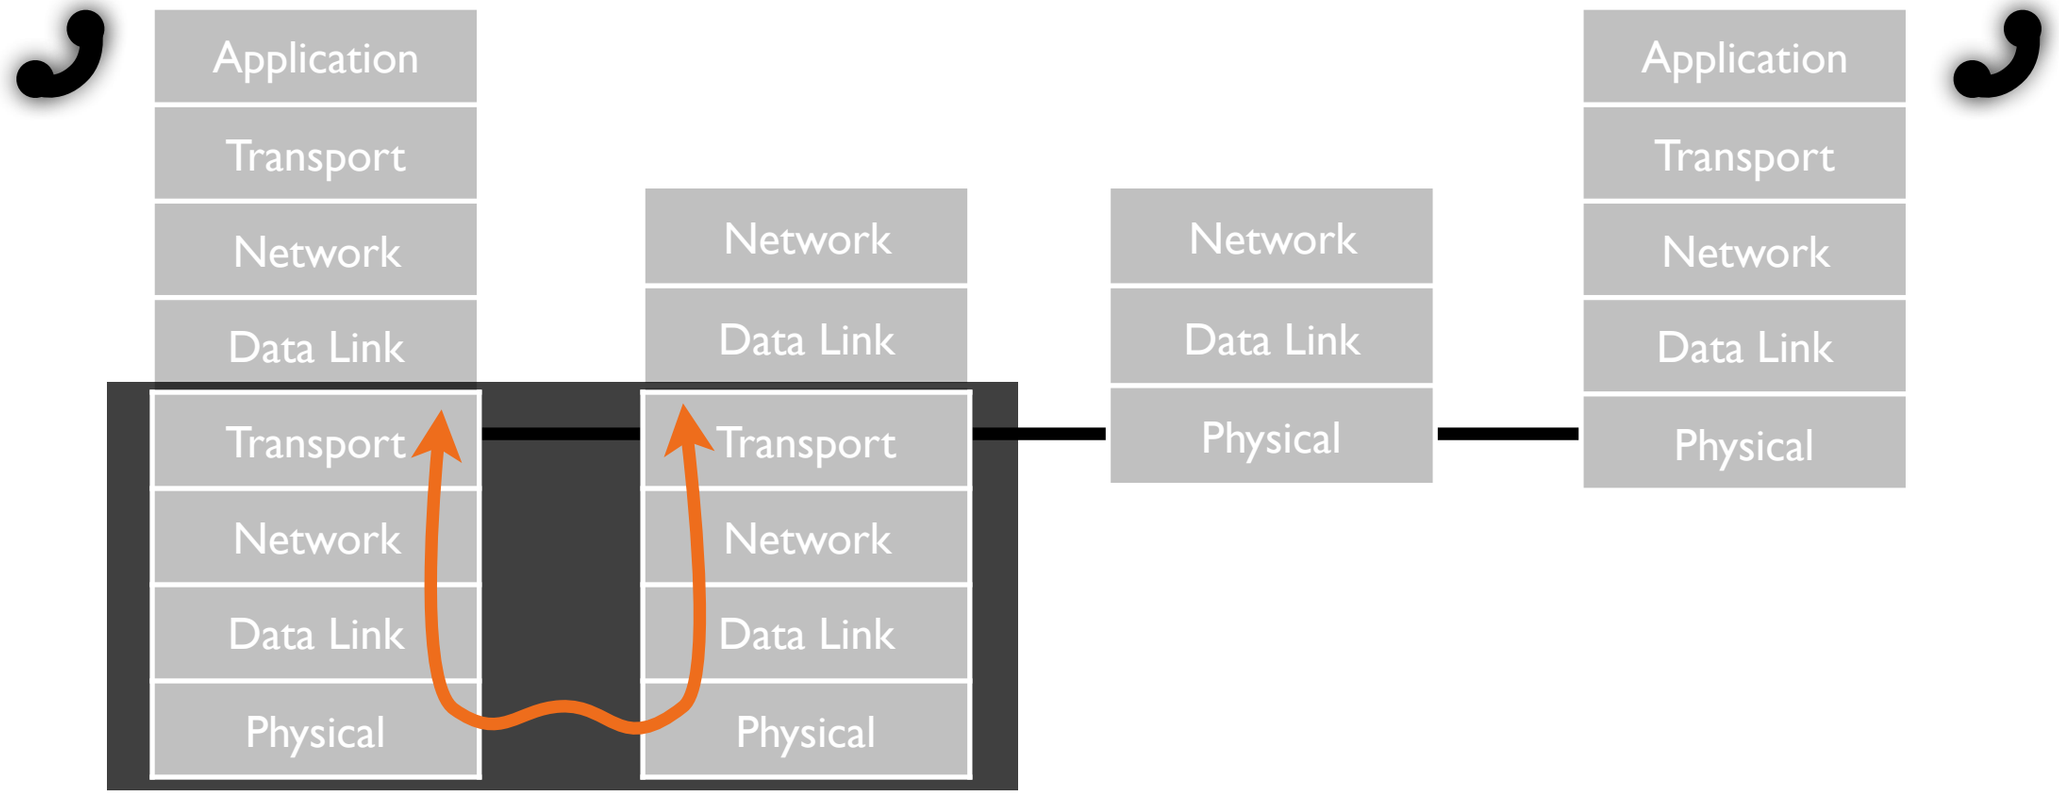

Tunnel

- VPN
- TOR
- VXLAN
- GRE
- MPLS
- …

Common functionality & problems

Application	*Anything you want...*	*Life, the universe, and everything*
Transport	Process-level communication	Reliability, flow control, ordering, congestion, ...
Network	Packets across domains	Independent parties, scale, routing
	Packets across networks	Addressing, heterogeneity, routing
Data Link	Packets on a 'wire'	Framing, errors, addressing
Physical	Encoding of bits	Physics, analog-to-digital

Grand Challenges

Bismuth strontium calcium copper oxide (BSCCO)

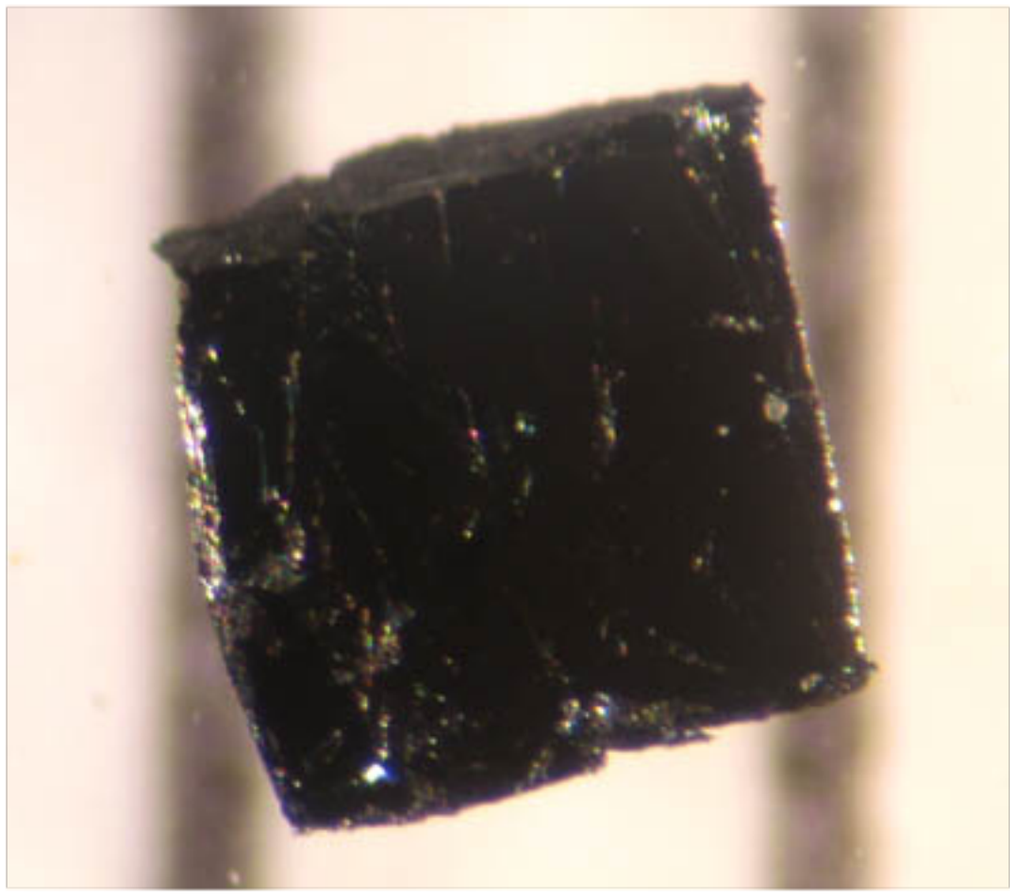

[Photo: James Slezak via Wikimedia]

Superconducts up to about -168°C (-271 °F)

High temperature superconductivity is a "Grand Challenge" for condensed matter physics

Grand Challenges

Widely recognized as among the most important unsolved problems in a field

- P vs. NP
- natural language understanding
- bug-free programs
- moving society to carbon-neutral energy
- preventing cancer
- ...

Grand Challenges in networking?

Getting an A in this class?

GC's in networking

An Informal Survey

1. "What I'm working on!"
2. High level objectives

- Security & privacy
- Reliability
- Usability

- Different than P vs. NP: hard to even define "security"; objectives involve tradeoffs

Unreliability: One Example

Internet Routing

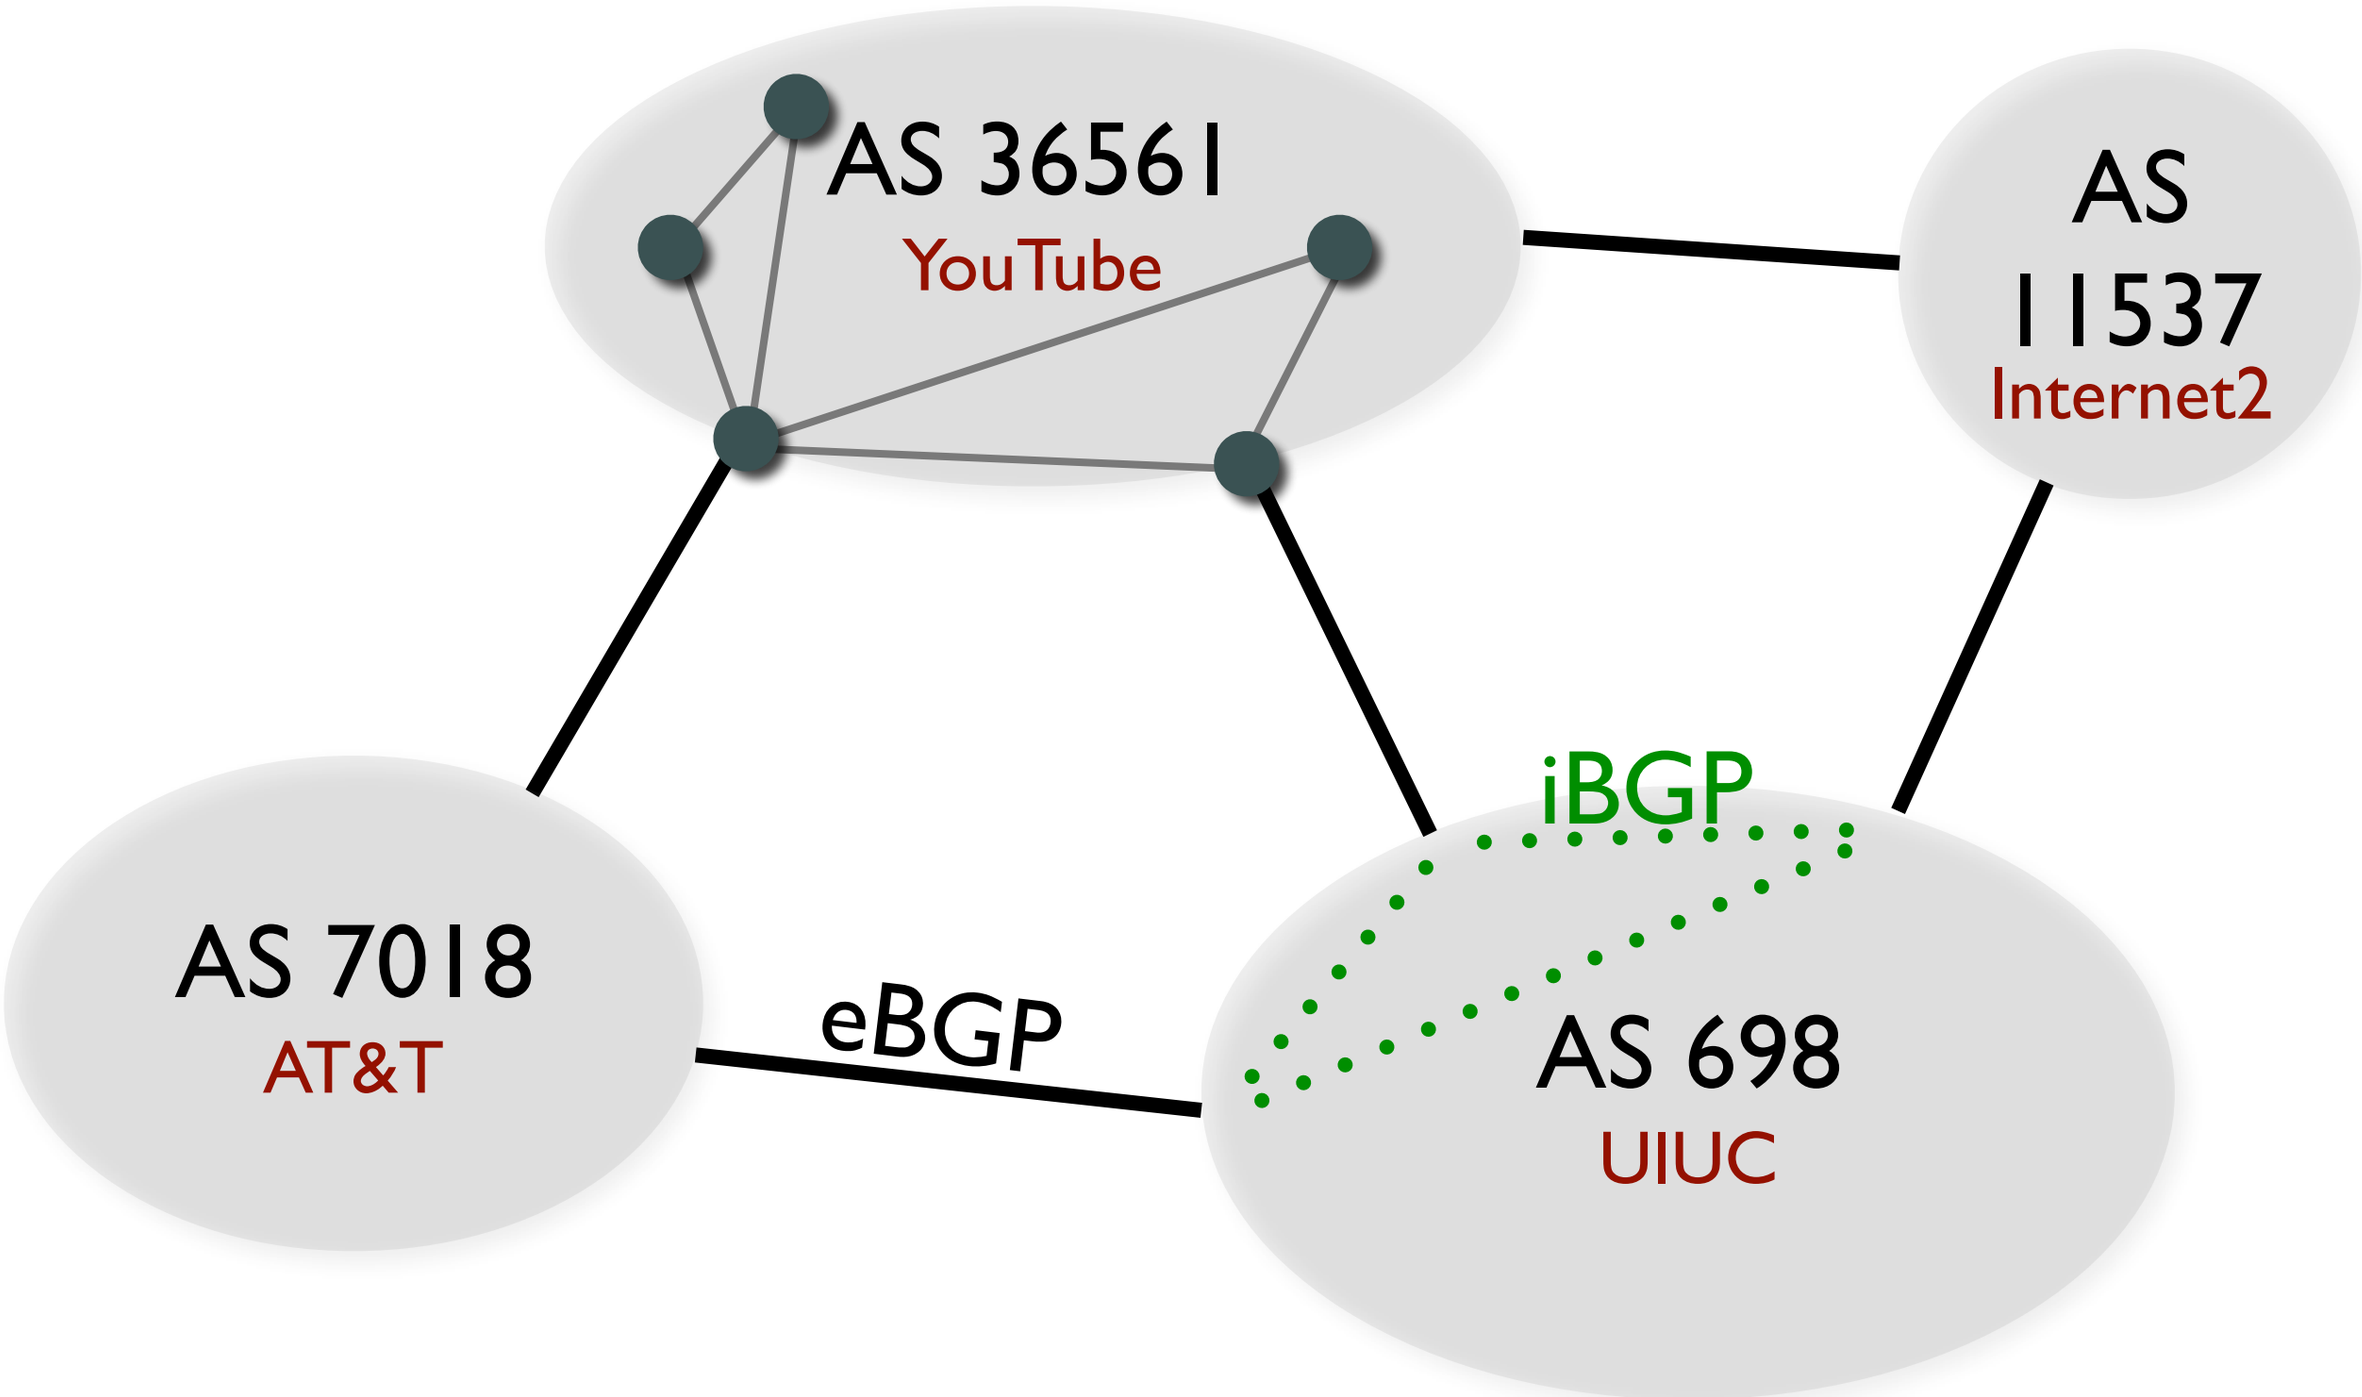

Border Gateway Protocol

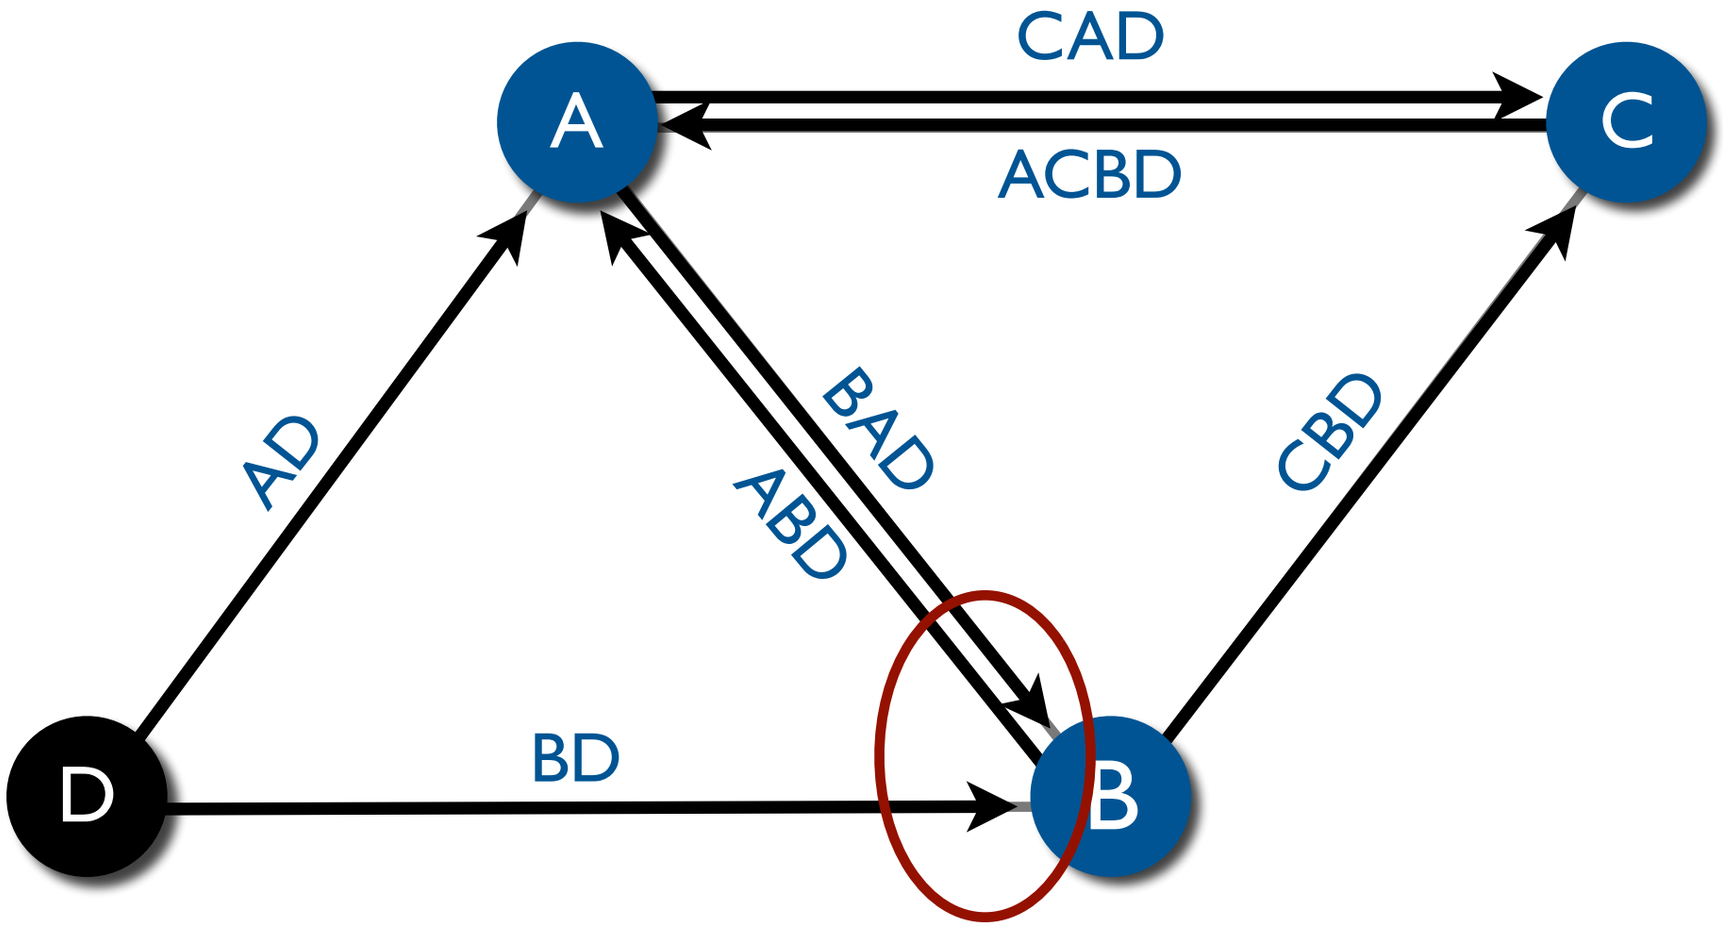

Instability causes outages

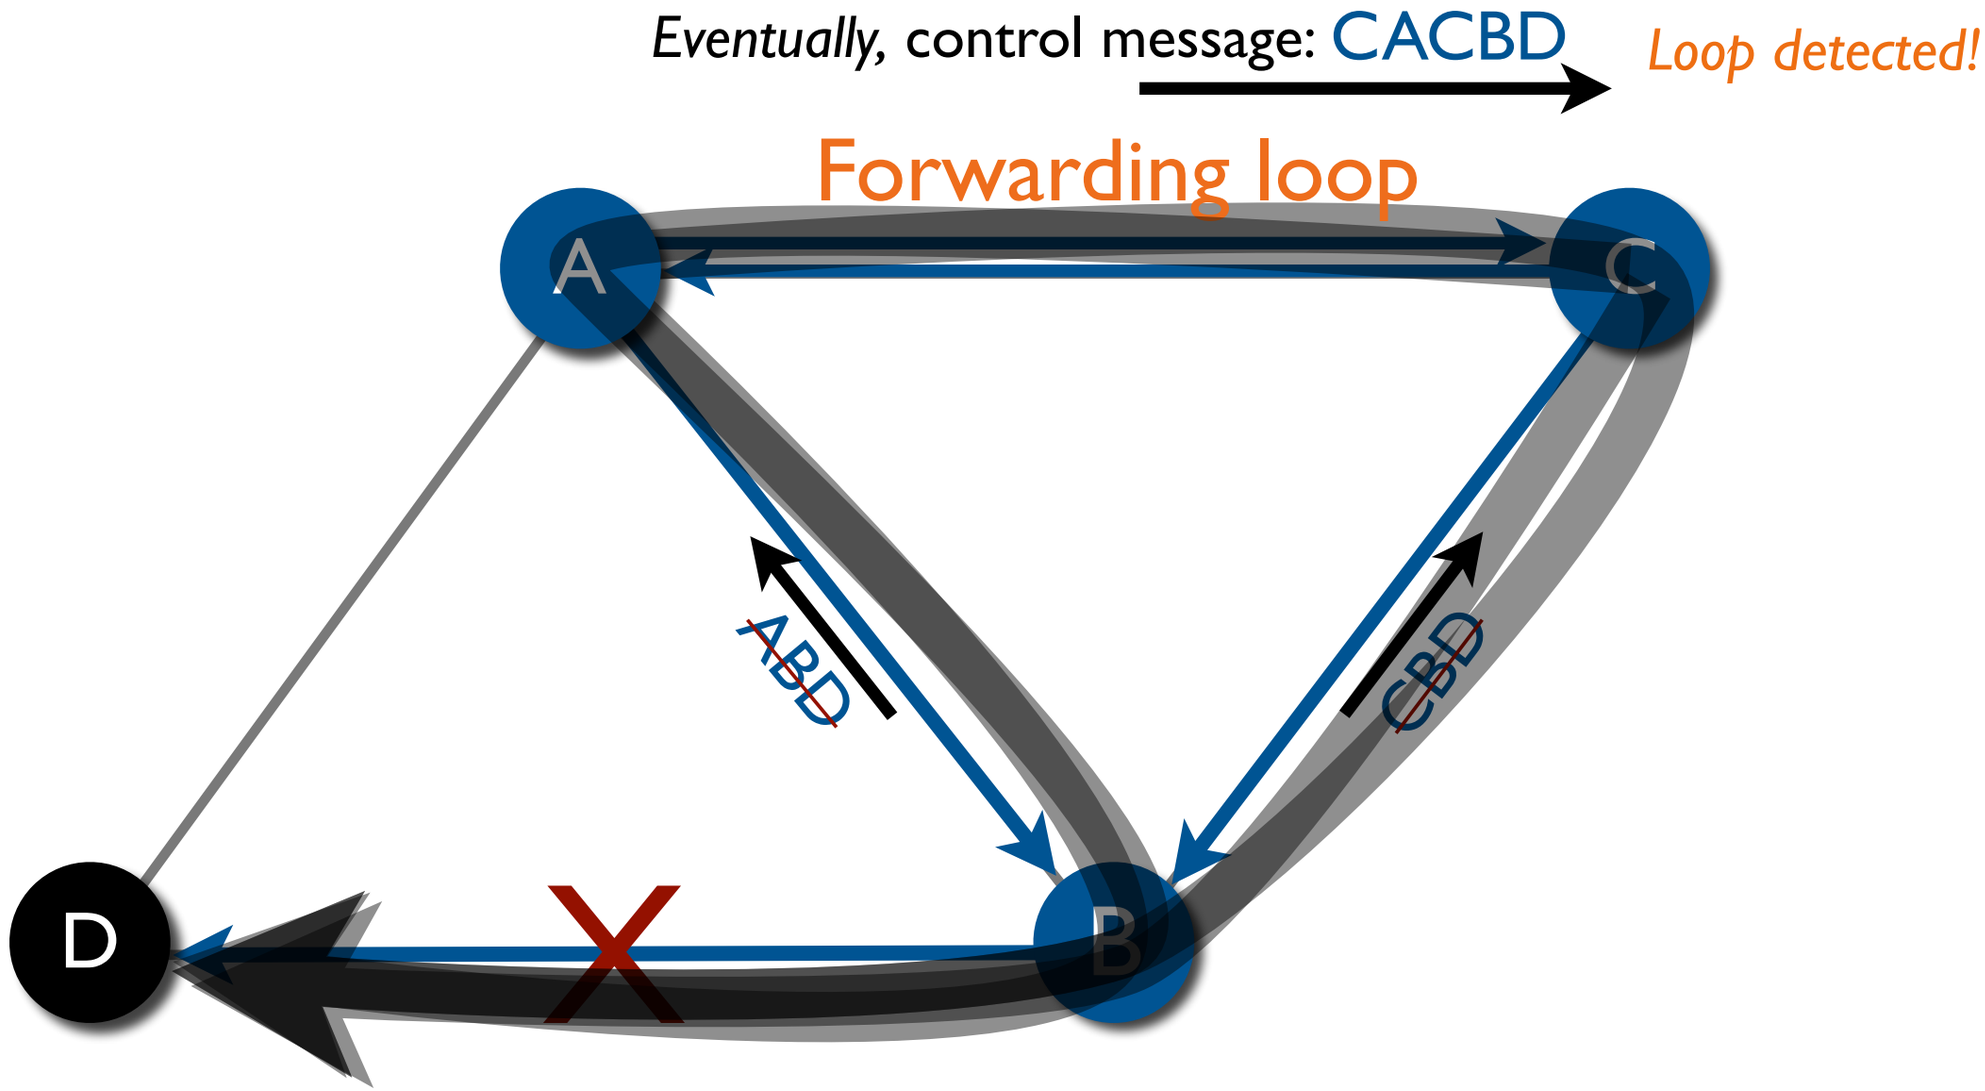

- Link state changes
- Router failures
- Config. changes
- ...

⟹

- Loops
- Detection delay
- Black holes

⟹ FAIL

Instability causes outages

[F. Wang, Z. M. Mao, J. Wang, L. Gao, R. Bush SIGCOMM'06]

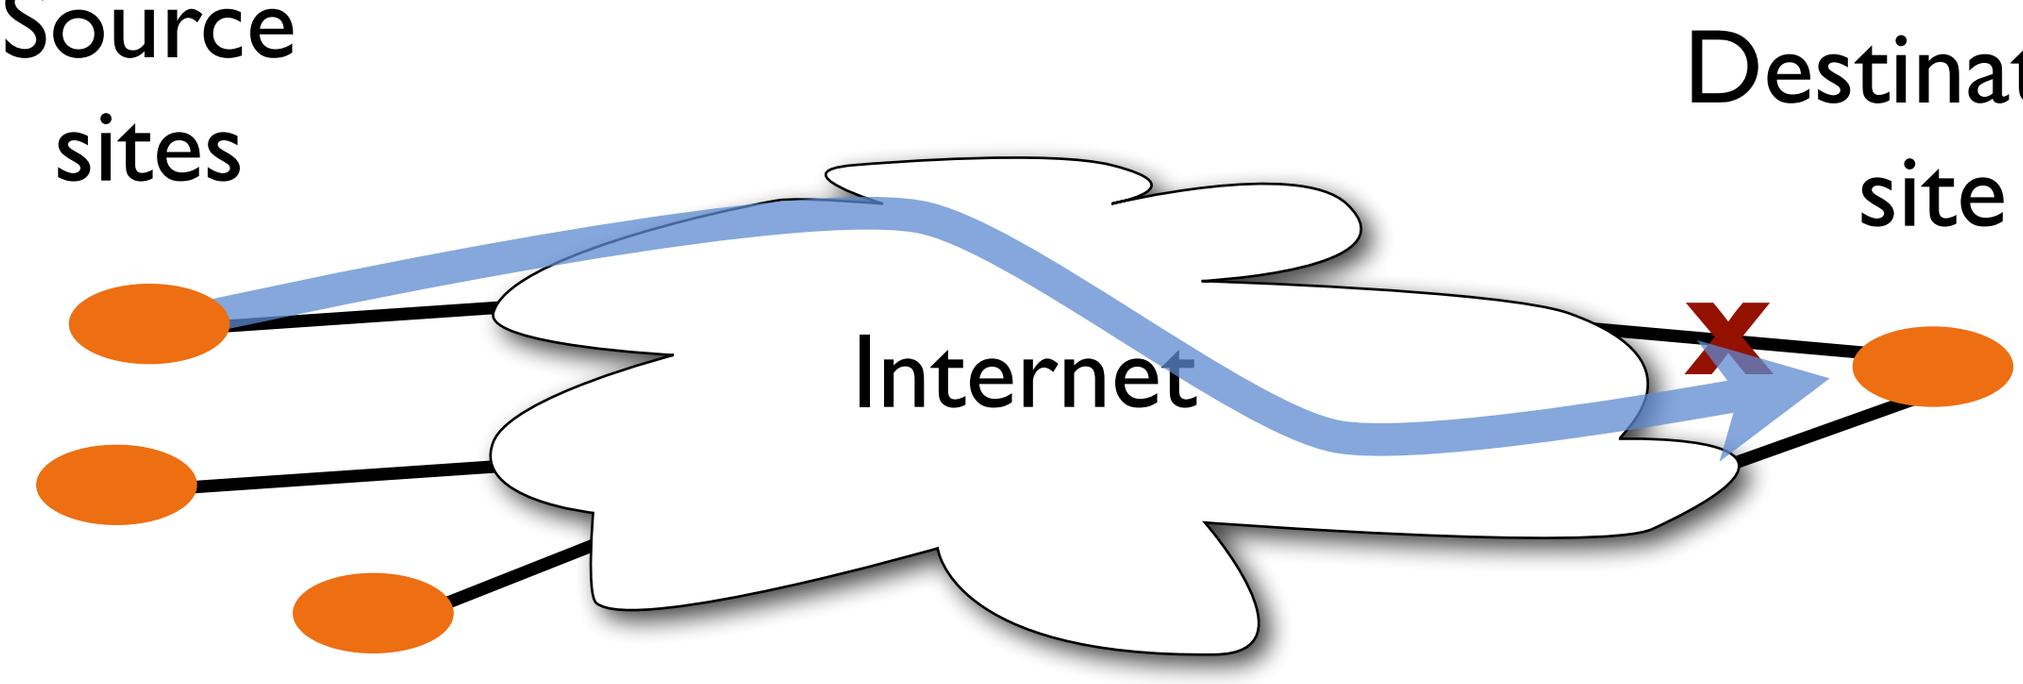

Instability causes outages

[F. Wang, Z. M. Mao, J. Wang, L. Gao, R. Bush SIGCOMM'06]

More outages

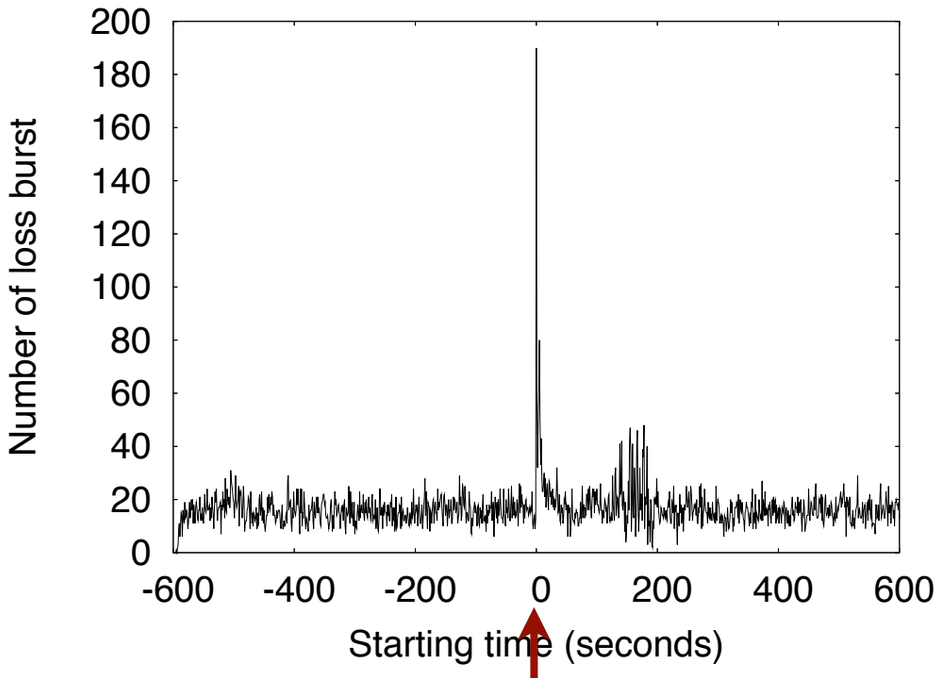

Failure injected

Longer outages

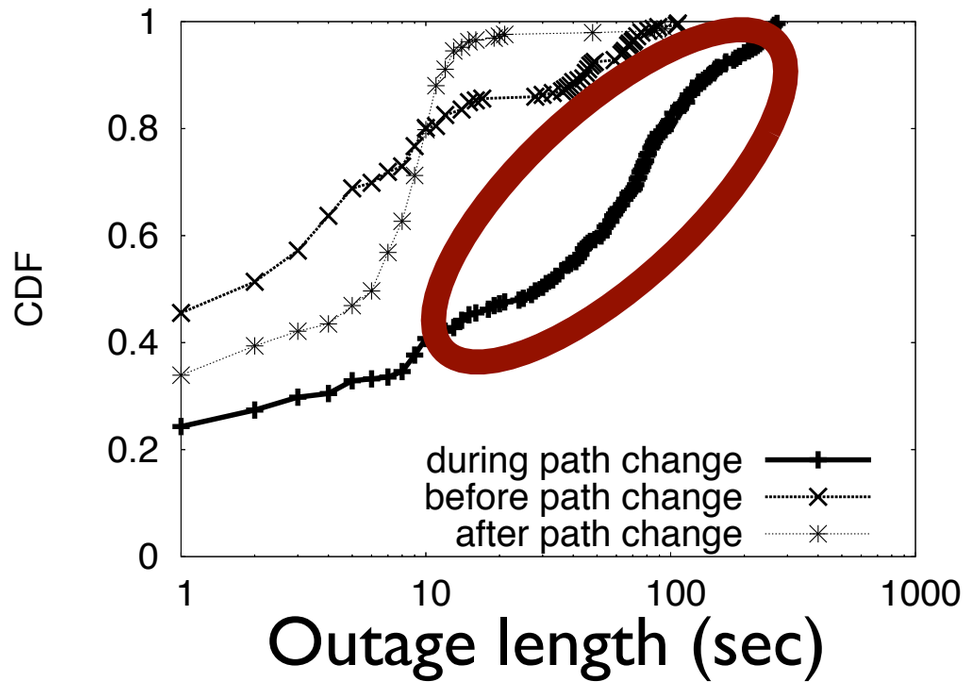

(...and higher latency, packet reordering, router CPU load during instability)

Many sources of unreliability

Congestion

- no end-to-end bandwidth reservations in the Internet

Configuration or software bugs

Failures or delays

- in network, DNS servers, caches, application servers, ...

Insecurity: one example

Prefix hijacking

Anyone can advertise routes for any IP prefix!

How can hijacker get the advertised routes to actually be used by other ASes?

- Announce more specific (longer) prefix than real owner
- Now everyone's traffic is "blackholed"

Can protect against this (Secure BGP), but...

- it's not deployed today
- and even then, can still cleverly (or accidentally) attract traffic and eavesdrop

From hijacking to MITM

Man in the Middle (MITM) attack

- Traffic to a destination redirected (not blackholed) through an attacker
- Attacker can watch everything you do without you noticing

What's the key problem here?

How can attacker forward traffic to destination, if attacker is pretending to *be* the destination?

From hijacking to MITM

Man in the Middle (MITM) attack

- Traffic to a destination redirected (not blackholed) through an attacker
- Attacker can watch everything you do without you noticing

What's the key problem here?

How can attacker forward traffic to destination, if attacker is pretending to *be* the destination?

Let's see how…

[Kapela and Pilosov, DEFCON'08]

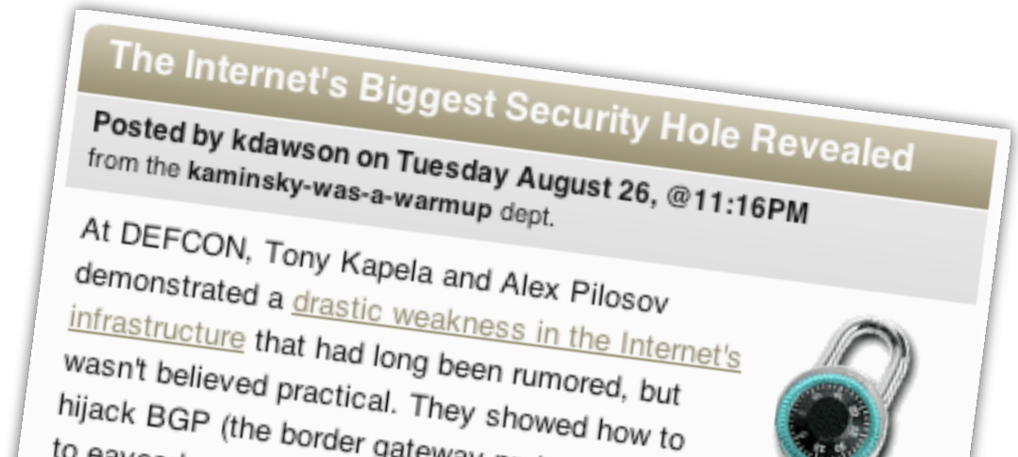

Hijacking + eavesdropping

1. A finds legitimate path ABD for 128.2.0.0/16

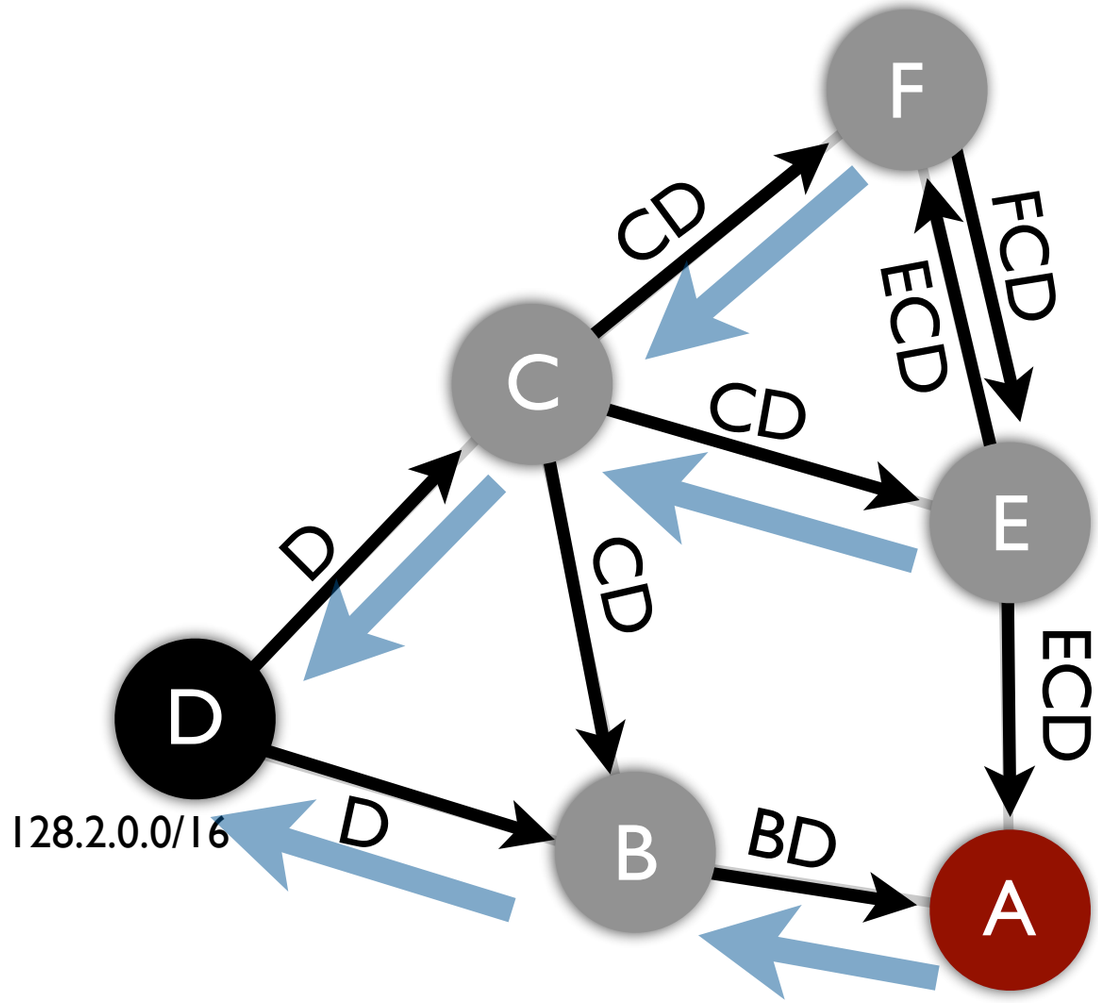

BGP Announcements

Selected routes & flow of data

Hijacking + eavesdropping

1. A finds legitimate path ABD for 128.2.0.0/16
2. A sends semi-bogus announcement of path ABD for 128.2.0.0/17
3. Result:
   - ASes (here B) on real path keep using real path because of loop elimination
   - All other ASes use route through A (/17 beats /16)
4. A forwards traffic to B

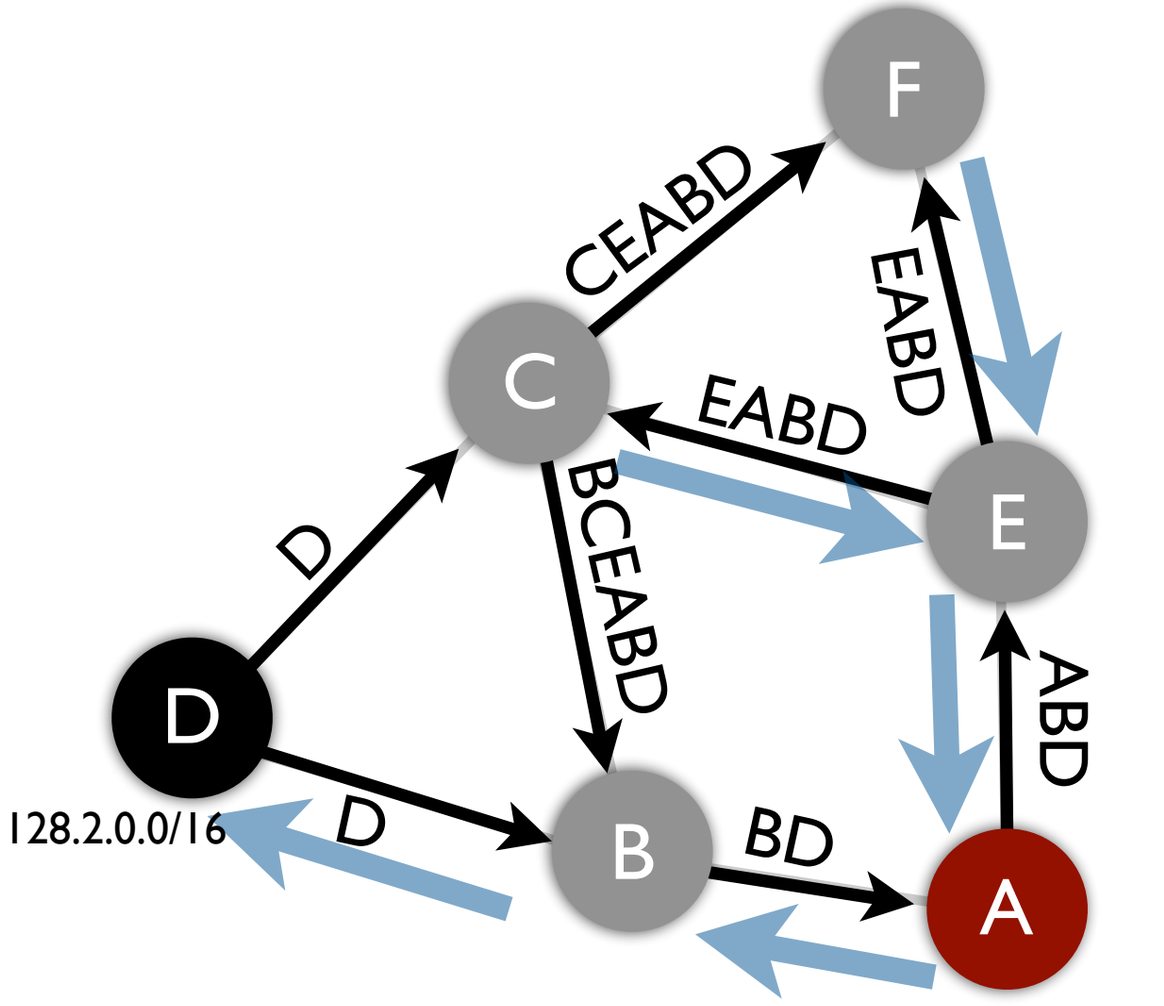

Hijacking + eavesdropping

Kapela & Pilosov also described how to spoof traceroute information to be even more undetectable.

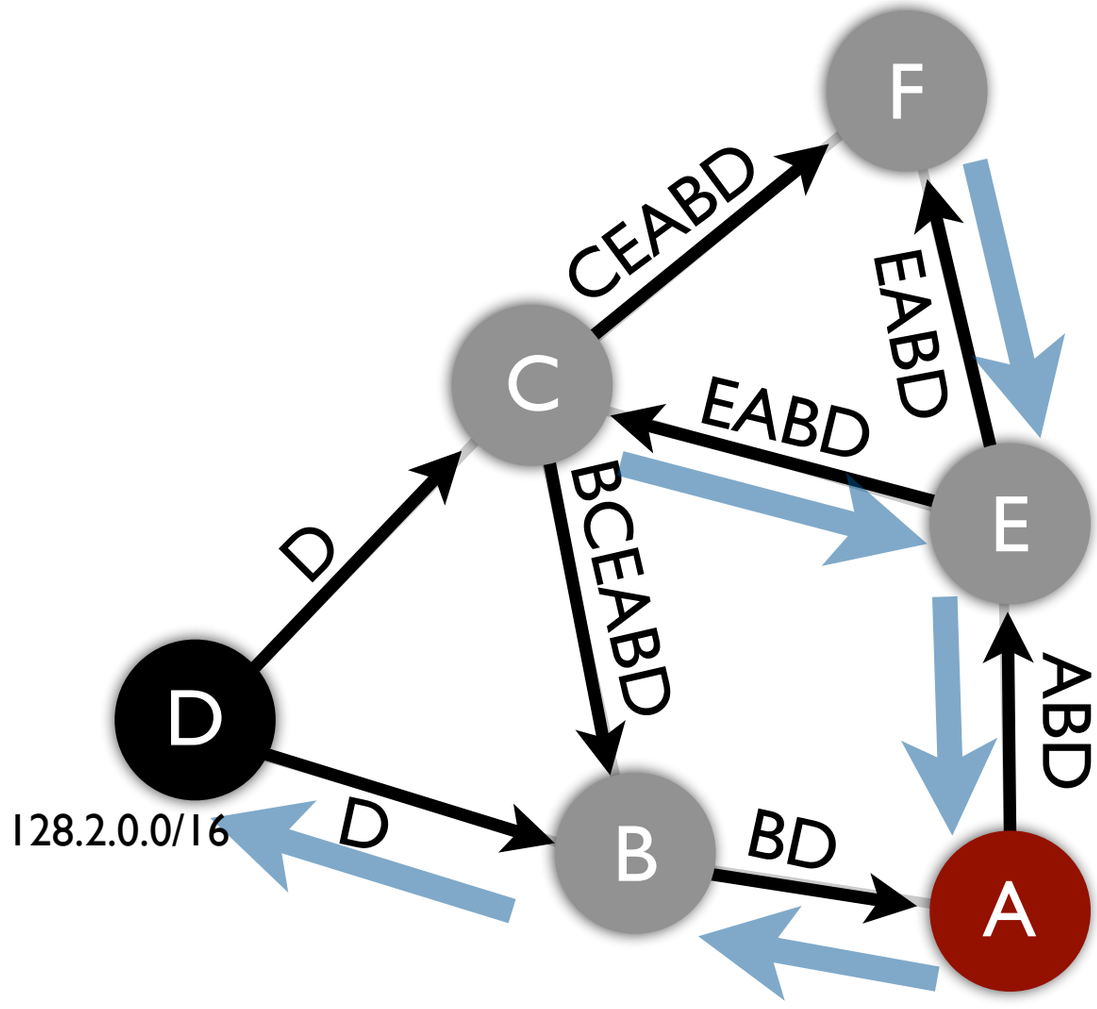

BGP Announcements

Selected routes & flow of data

January 5, 2017 incident

Routes to several pornographic sites (and later Apple iTunes) change

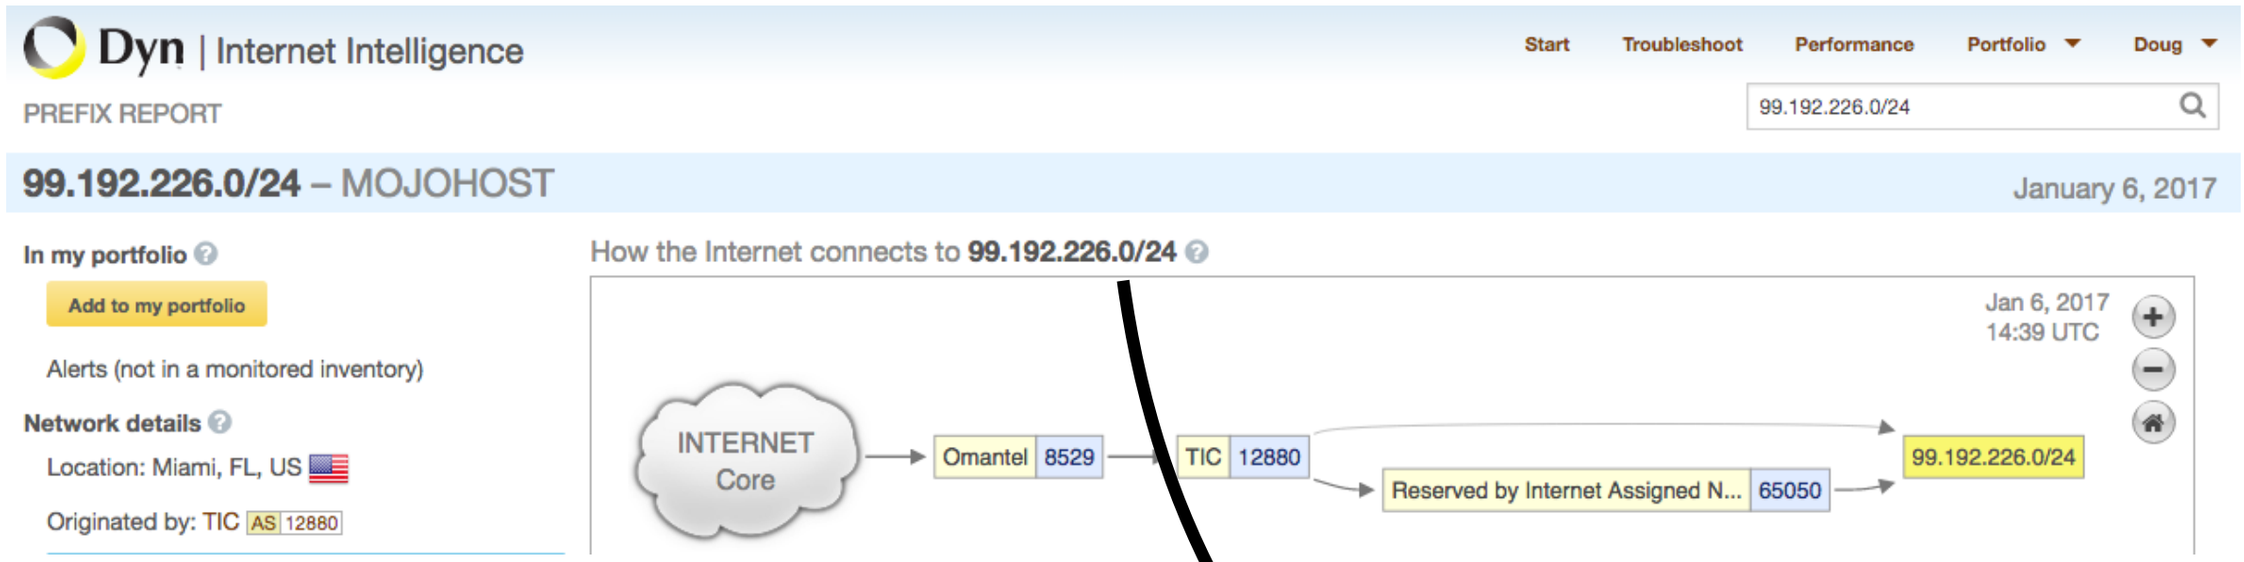

Iranian state ISP

January 5, 2017 incident

Routes to several pornographic sites (and later Apple iTunes) change

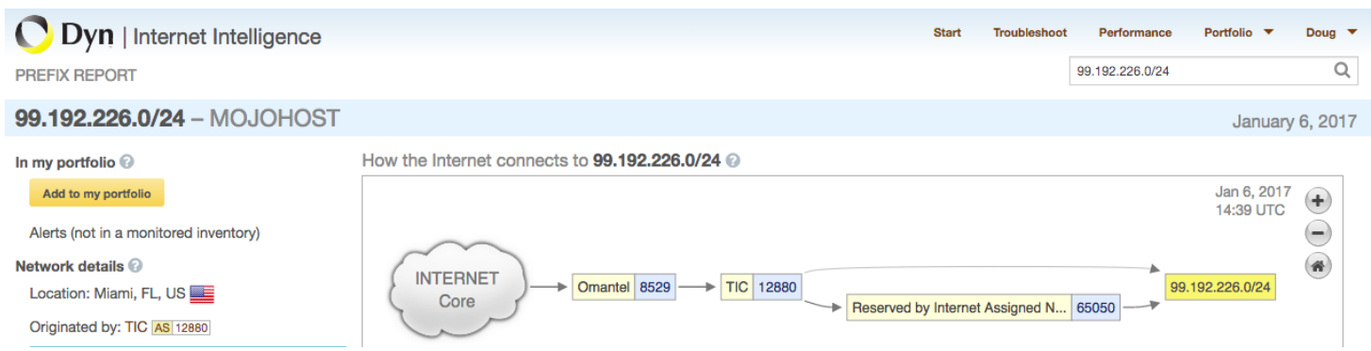

Recovery after owner finds out and takes action

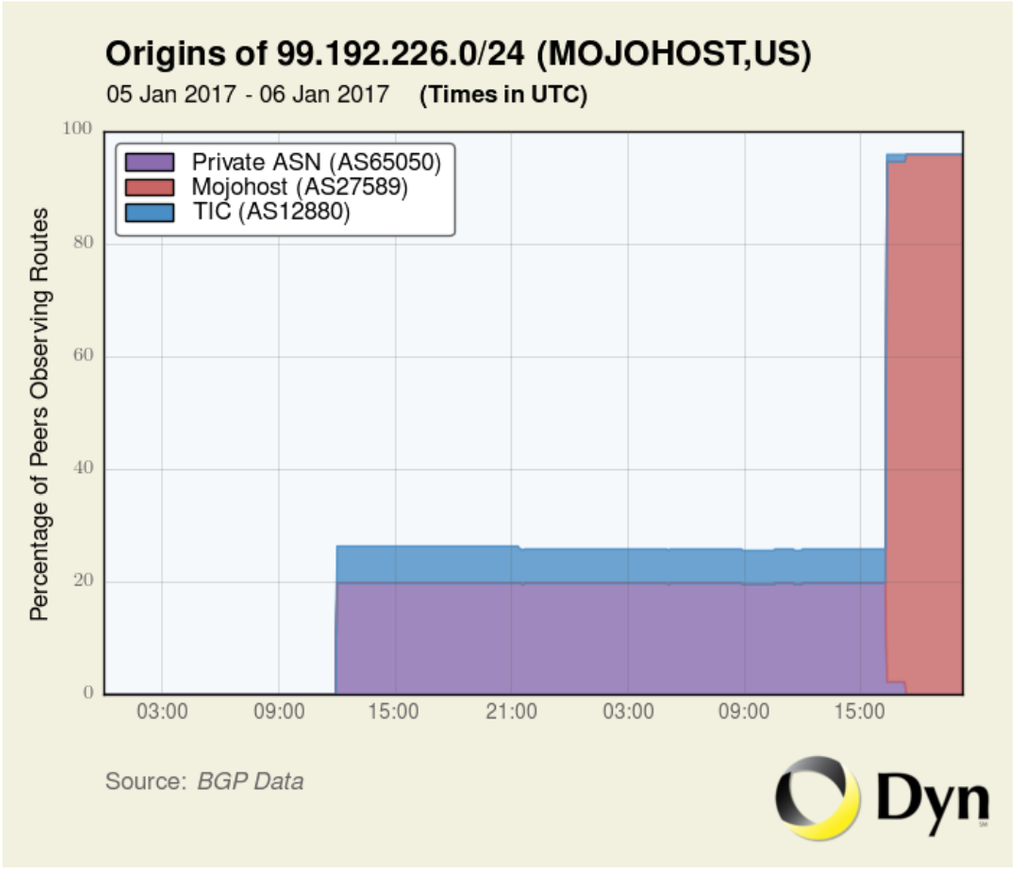

Source:
http://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/

December 12, 2017 incident

AS 39523 (DV-LINK-AS, Russia) hijacks routes

- 80 IP prefixes from Russian networks and also Google, Apple, Facebook, Microsoft, NTT, Riot Games & more
- Two intervals of about 3 minutes
- Announcement propagated largely via Hurricane Electric

Sources:
https://bgpmon.net/popular-destinations-rerouted-to-russia/

https://dyn.com/blog/recent-russian-routing-leak-was-largely-preventable/

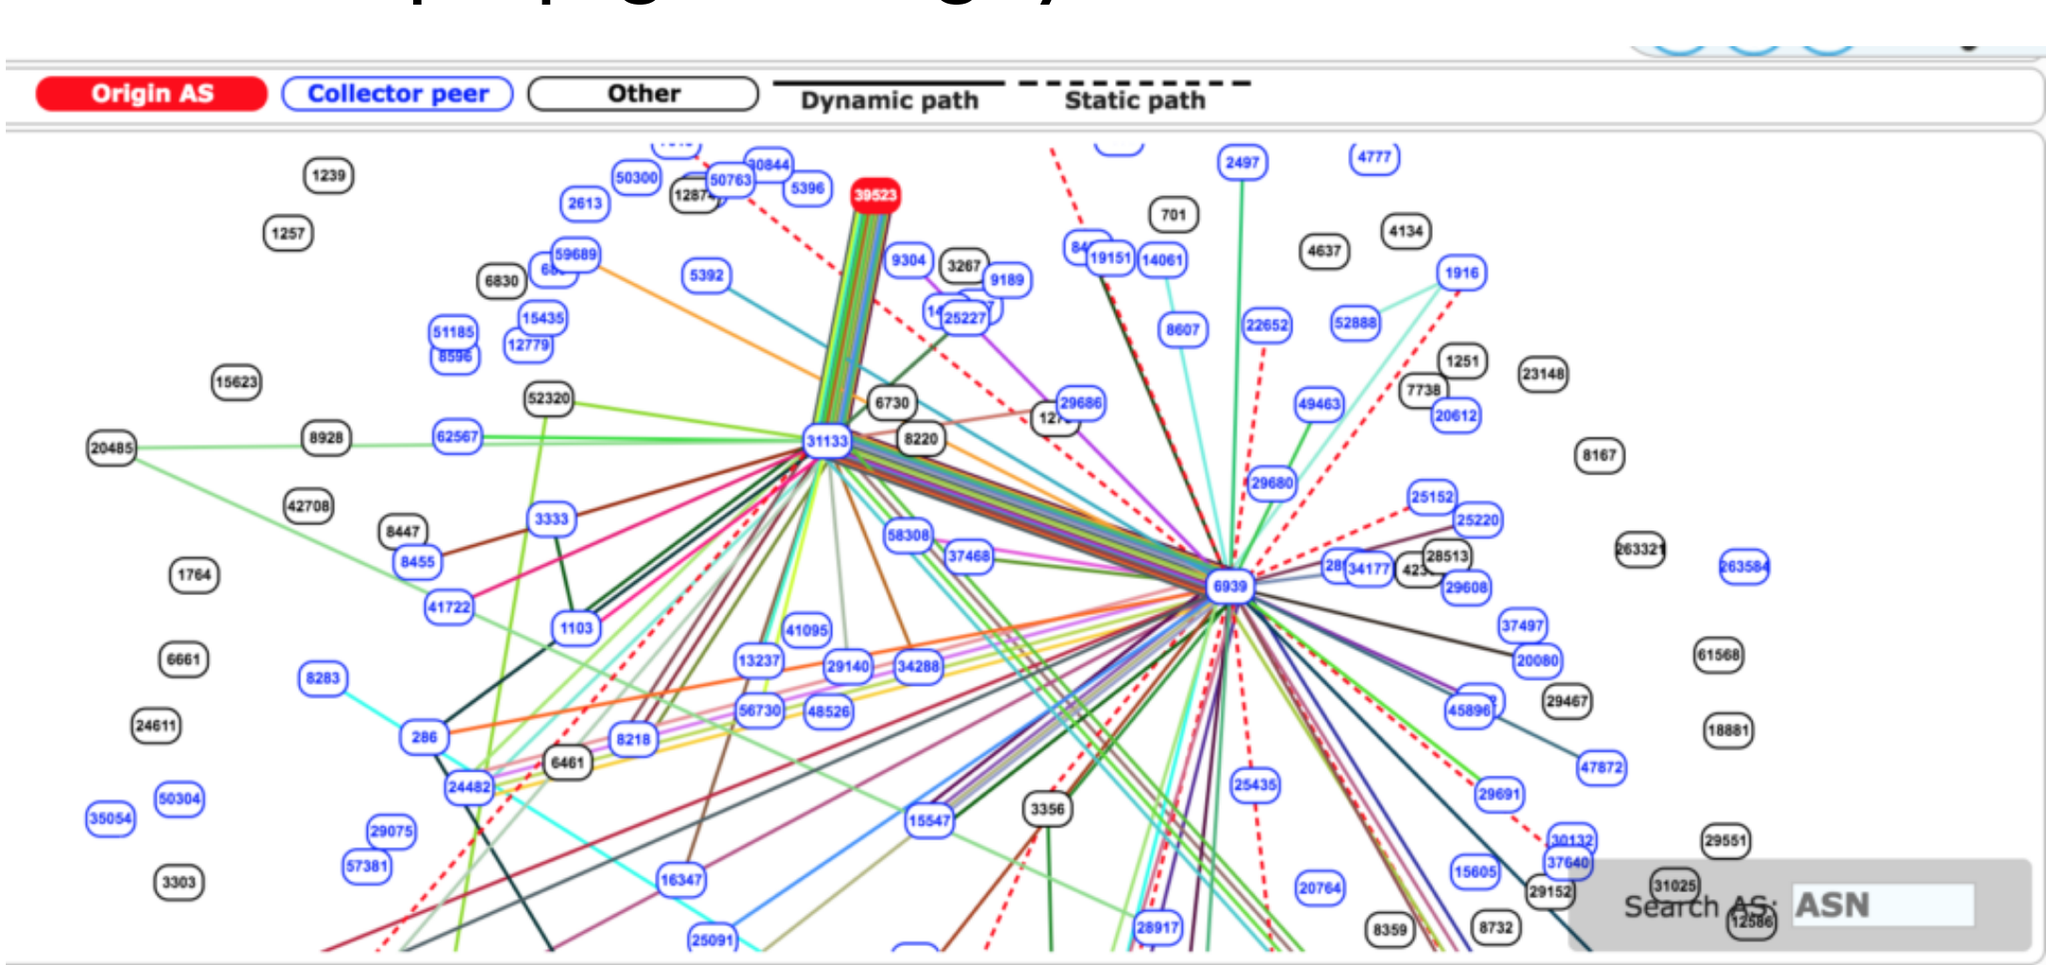

Grand Challenges in networking

An Informal Survey

1. "What I'm working on!"
2. Nebulous high level objectives
   - Security & privacy • Complexity
   - Reliability
   - Usability
3. Why does networking lack a crisp Grand Challenge?
   - Infrastructure needs to support highly diverse and dynamic goals, applications, and environments

Grand Challenges in networking

Meta-challenge:

How do we make the Internet evolvable?

Announcements

Reviews due by 11:59 pm Tuesday:

- A protocol for packet network intercommunication (Cerf and Kahn, 1974)
- The Design Philosophy of the DARPA Internet Protocols (Clark, 1988)

Micro-presentations

For those of you looking for project teams, tell us

- Your technical background
- Areas you're interested in studying, if you have ideas