Routing

Problem: Given more than one path from source to destination, which one to take?

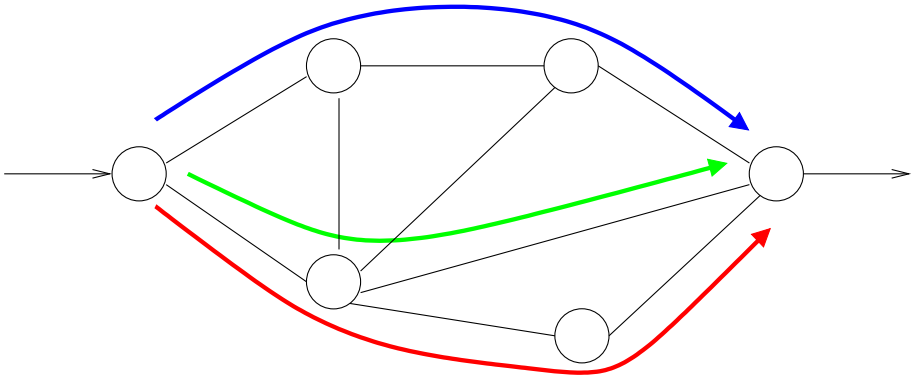

Features:

- Architecture
- Algorithms
- Implementation
- Performance

Architecture

Internet routing: two separate routing subsystems

→ intra-domain: within an organization

→ inter-domain: across organizations

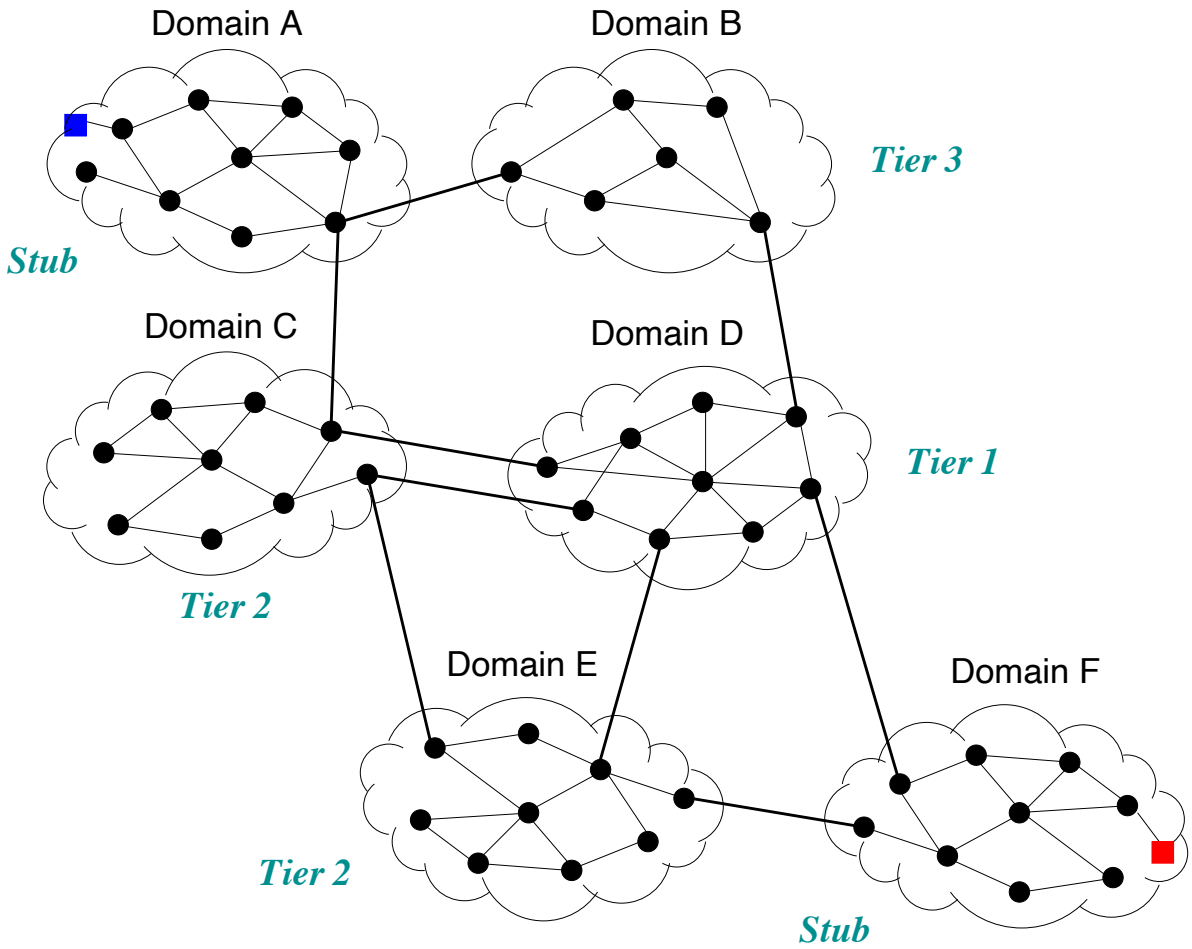

Three levels: LAN, intra-domain, and inter-domain

→ LAN: spanning tree, flooding/broadcast

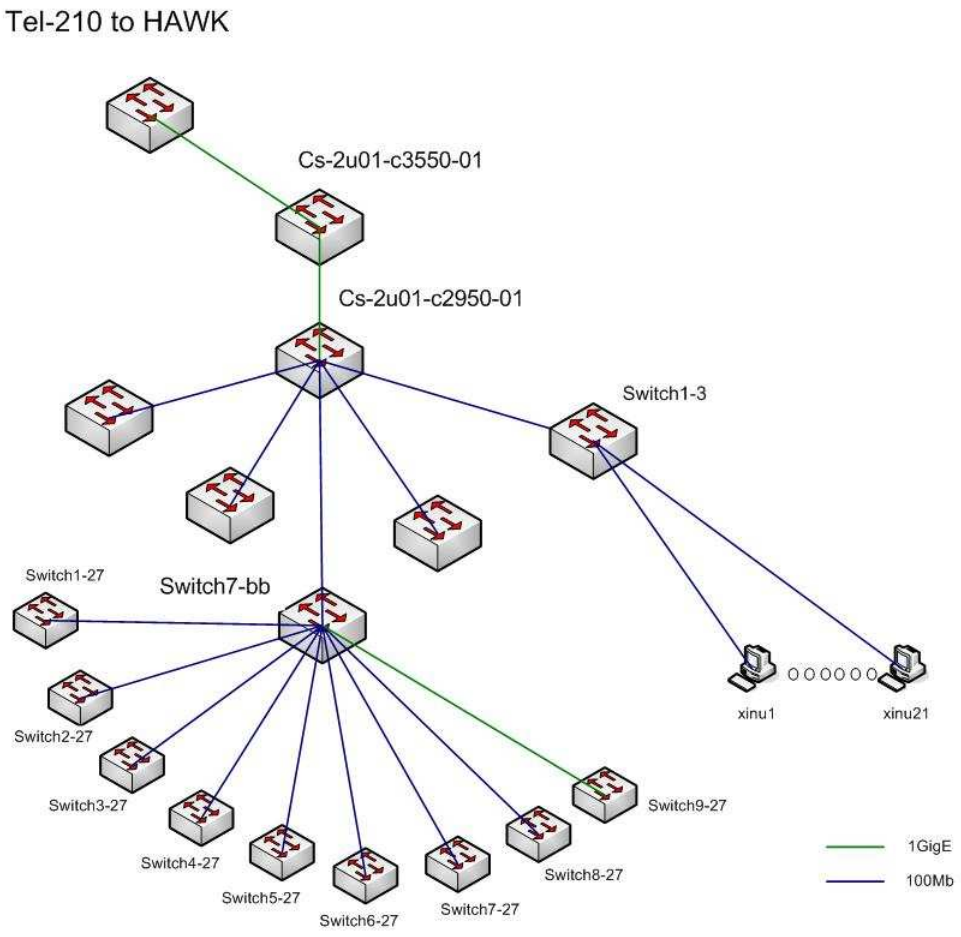

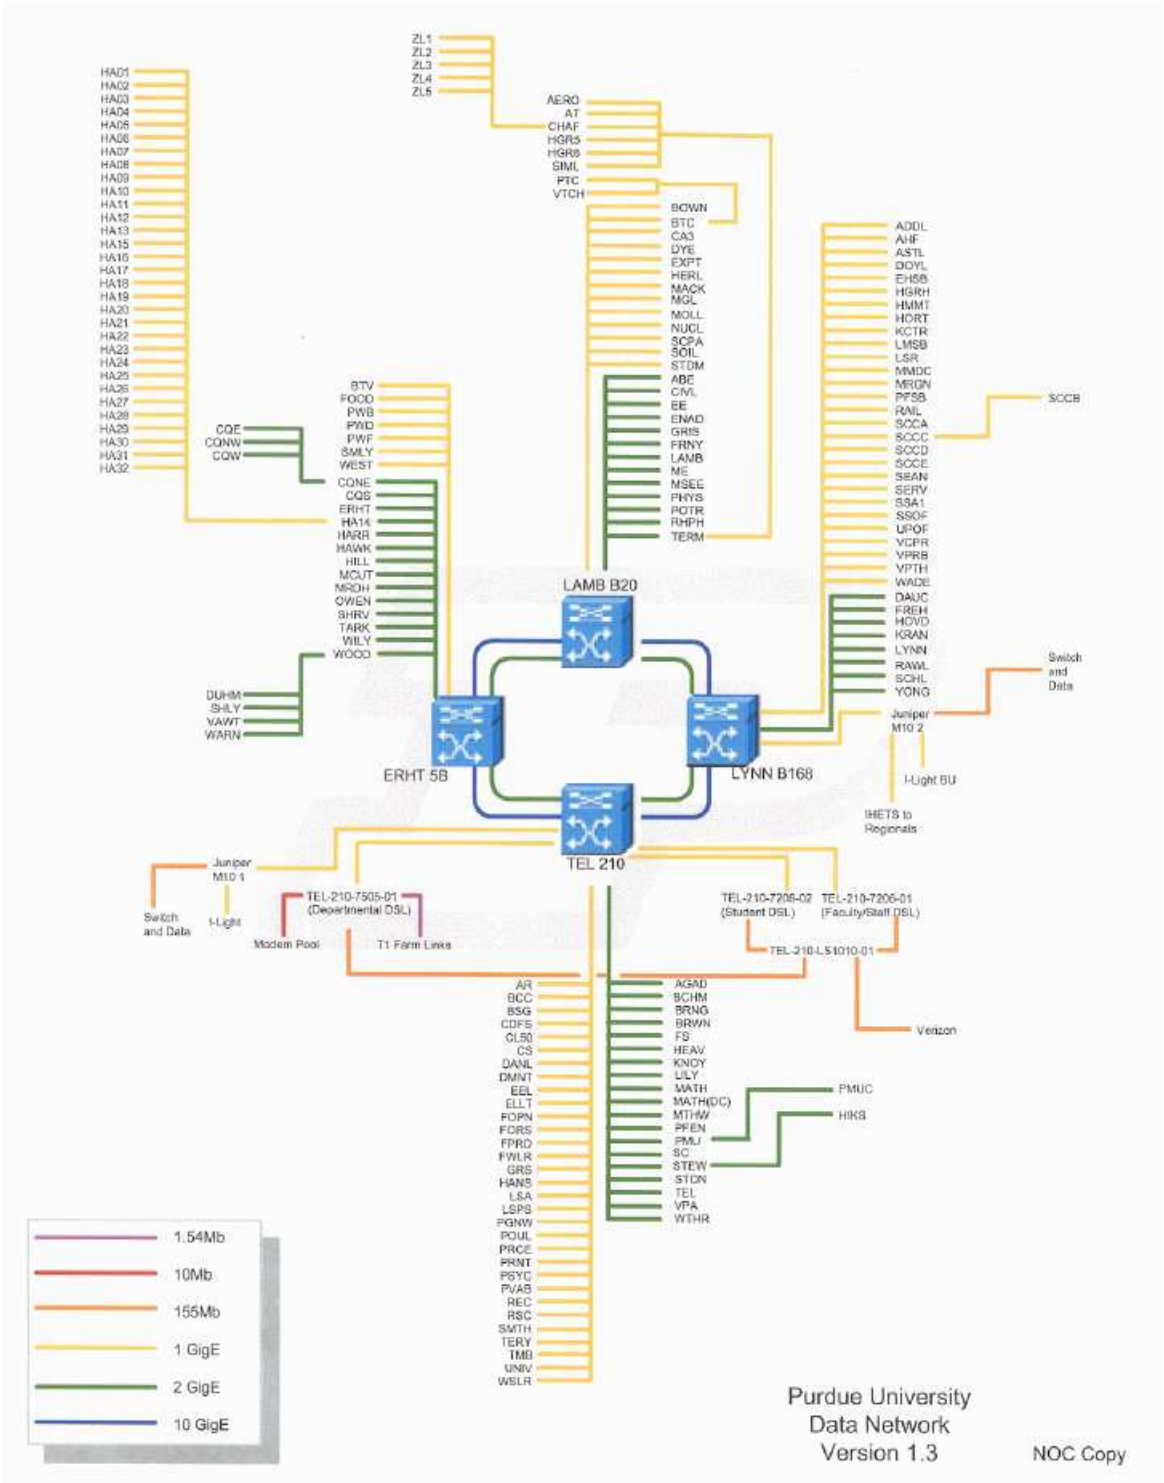
ZL1
ZL2
ZL3
ZL4
ZL5
HA01
HA02
HA03
HA04
HA05
HA06
HA07
HA08
HA09
HA10
HA11
HA12
HA13
HA14
HA15
HA16
HA17
HA18
HA19
HA20
HA21
HA22
HA23
HA24
HA25
HA26
HA27
HA28
HA29
HA30
HA31
HA32
AERO
AT
CHAF
HGR5
HGR6
SIML
PTC
VTCH
BOWN
BTC
CA3
DYE
EXPT
HERL
MACK
MGL
MOLL
NUDL
SCPA
SOIL
STDM
ABE
CIVL
EE
ENAD
GRIS
FRNY
LAMB
ME
MSEE
PHYS
POTR
RHPH
TERM
ADDL
AHF
ASTL
DOYL
EHSB
HGRH
HMMT
HORT
KCTR
LMSB
LSR
MMDC
MRGN
PFSB
RAIL
SCCA
SCCC
SCCD
SCCE
SEAN
SERV
SSA1
SSOF
UPOF
VCPR
VPRB
VPTH
WADE
DAUC
FREH
HOYD
KRAN
LYNN
RAWL
SCHL
YONG
SCCB
Switch and Data
Juniper M10 2
I-Light BU
IHETS to Regionals
BTV
FOOD
PWB
PWD
PWF
SMLY
WEST
CQE
CQNW
CQW
CQNE
CQS
ERHT
HA14
HARR
HAWK
HILL
MCUT
MROH
OWEN
SHRV
TARK
WILY
WOOD
DUHM
SHLY
VAWT
WARN
LAMB B20
ERHT 5B
LYNN B168
TEL 210
Juniper M10 1
Switch and Data
I-Light
TEL-210-7505-01 (Departmental DSL)
Modem Pool
T1 Form Links
TEL-210-7206-02 (Student DSL)
TEL-210-7206-01 (Faculty/Staff DSL)
TEL-210-LS1010-01
Verizon
AR
BCC
BSG
CDFS
CL50
CS
DANL
DMNT
EEL
ELLT
FOPN
FORS
FPRD
FWLR
GRS
HANS
LSA
LSPS
PGNW
POUL
PRCE
PRNT
PSYC
PVAB
REC
RSC
SMTH
TERY
TMB
UNIV
WSLR
AGAD
BCHM
BRNG
BRWN
FS
HEAV
KNOY
LILY
MATH
MATH(DC)
MTHW
PFEN
PMU
SC
STEW
STON
TEL
VPA
WTHR
PMUC
HIKS
1.54Mb
10Mb
155Mb
1 GigE
2 GigE
10 GigE
Purdue University
Data Network
Version 1.3
NOC Copy

Inter-domain topology:

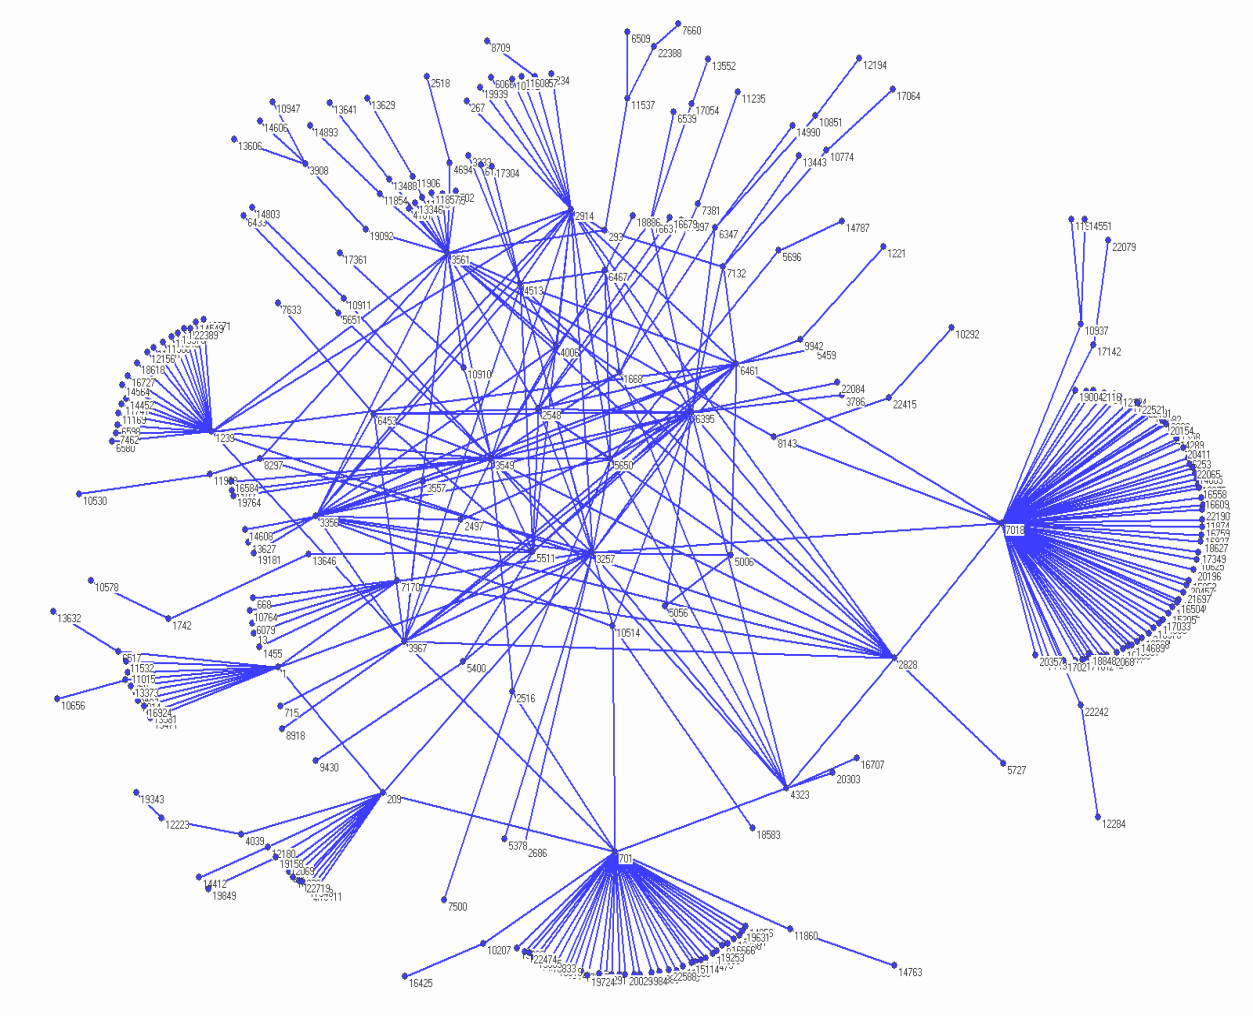

$\longrightarrow$ each dot (or node) is a domain (e.g., Purdue)

$\longrightarrow$ autonomous system (AS): 16- or 32-bit ID

Inter-domain connectivity of Purdue:

- AS 6939 (Hurricane), Internet 2 (AS11164), etc. → INDIANAGIGAPOP (AS 19782) → Purdue (AS 17)
- AS 3356 (Level3), AS 6939 (Hurricane), AS 1299 (Arelion), AS 6461 (Zayo) → WINTEK (AS 11114) → Purdue (AS 17)

→ evolves over time

Indy GigaPoP (19782):

→ part of I-Light (Indiana state-wide project)

→ located at IUPUI

→ provides state-level fiber connectivity including Purdue and IU

Level-3 (AS 3356): Lumen Technologies

→ tier-1

Abilene/Internet2 backbone: www.internet2.edu

Abilene International Network Peers

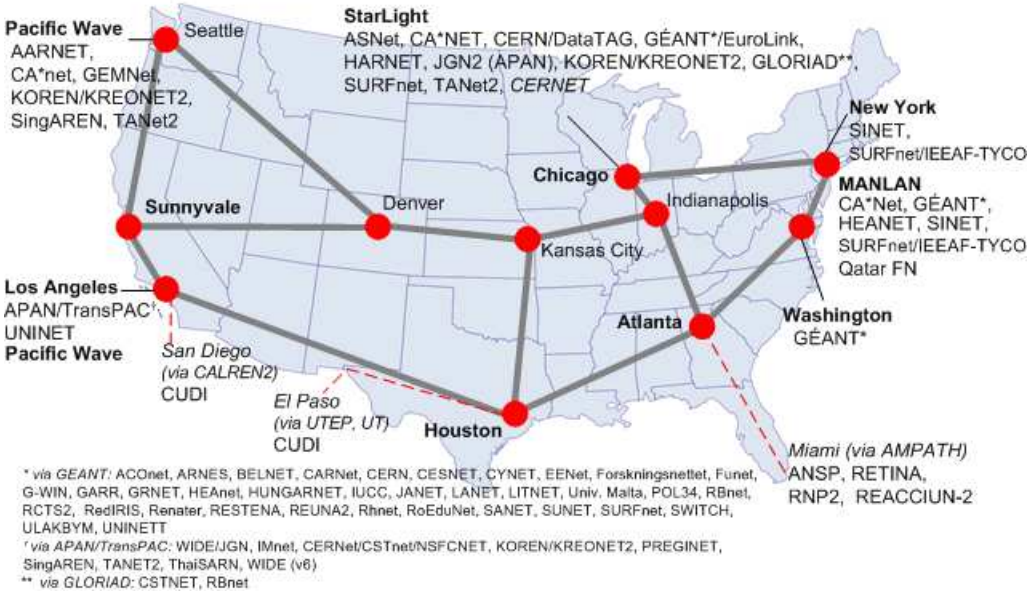

AT&T (AS 7018)'s U.S. PoP topology (inferred):

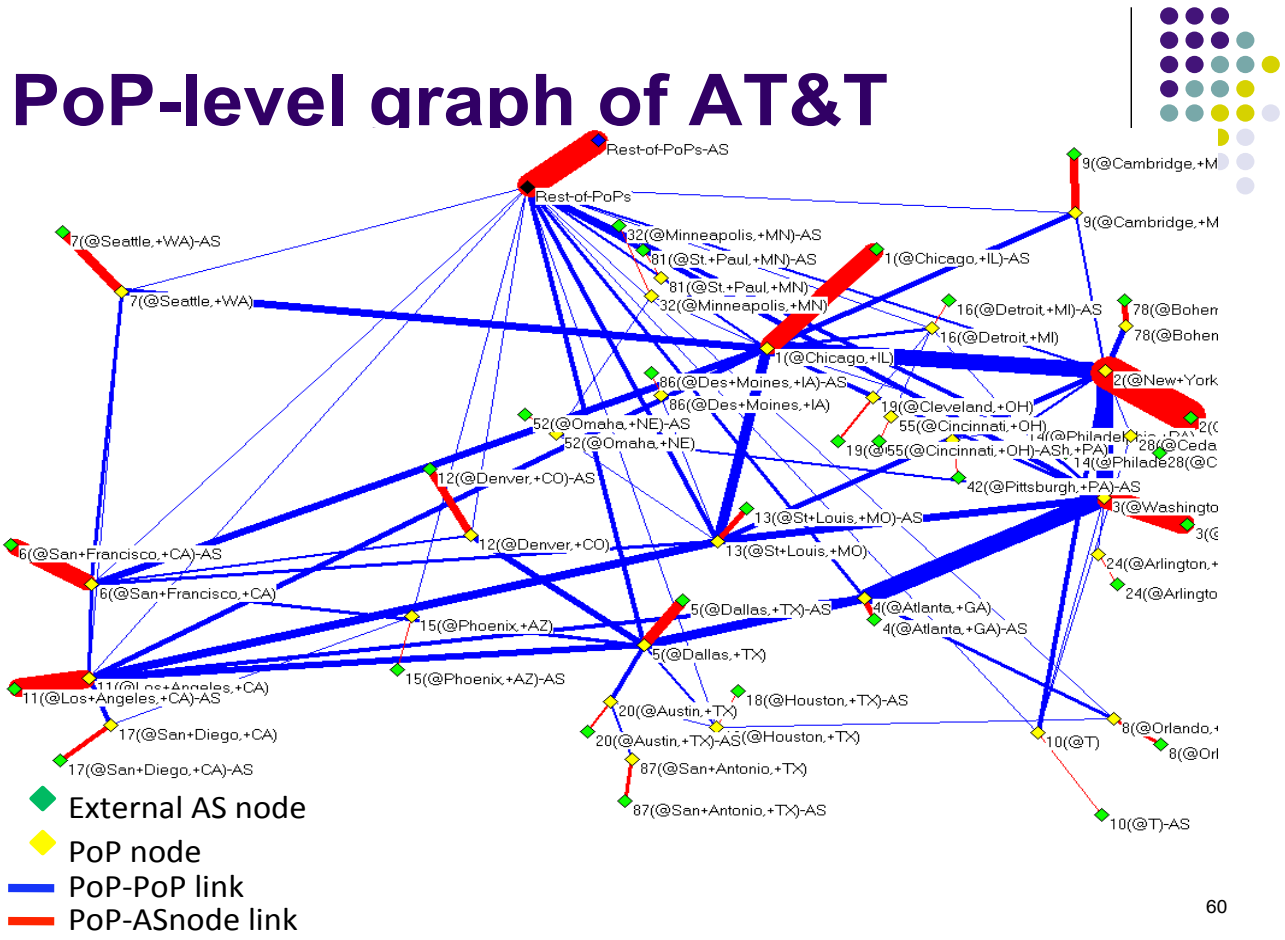

AT&T's Chicago PoP connectivity (inferred):

Chicago PoP of AT&T

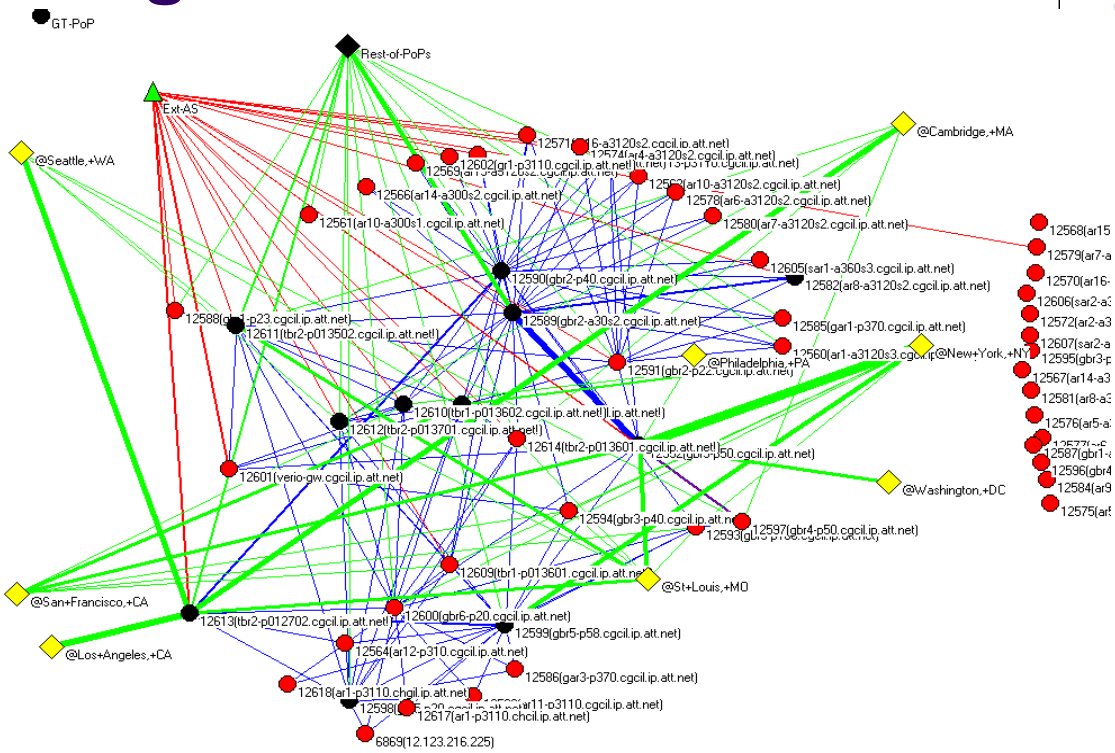

Granularity of routing network:

- router

  → IP packet routing

  → also: label switching within large ISPs)

  → note: LAN routing is invisible

- domain: autonomous system

  → 16- or 32-bit identifier ASN

  → extended to 32-bit in 2007

  → assigned by IANA along with IP prefix block (CIDR)

Network topology

→ i.e., connectivity

- router graph

  → node: router/switch

  → edge: physical link between two routers

- AS graph

  → node: AS

  → edge: physical link between 2 or more border routers

  → sometimes at exchange point/network

Router type:

- access router

  → collects traffic from devices of a domain/network

  → distributes traffic to devices of a domain/network

- border router

  → interface between two or more domains

  → packet crosses administrative boundary

- backbone router

  → routers that form intradomain network

  → e.g., Purdue's backbone routers (ring)

Traditional intra-domain/enterprise router topology emphasizes vertical traffic

→ north-south

→ vertical aggregation/deaggregation

→ e.g., fat tree

Recent data center trend: horizontal or east-west traffic.

→ driven by workload demand

→ e.g., AI

→ Clos/leaf-spine network

AS type:

- stub AS: customer AS
  - → no forwarding
  - → may be multi-homed: more than one provider
  - → e.g., Purdue
- transit AS
  - → provide connectivity to stub AS's and smaller transit AS's
  - → tier-1: global reachability and no provider above
  - → tier-2 or tier-3: regional providers as well as customers of tier-1 AS's

AS graph:

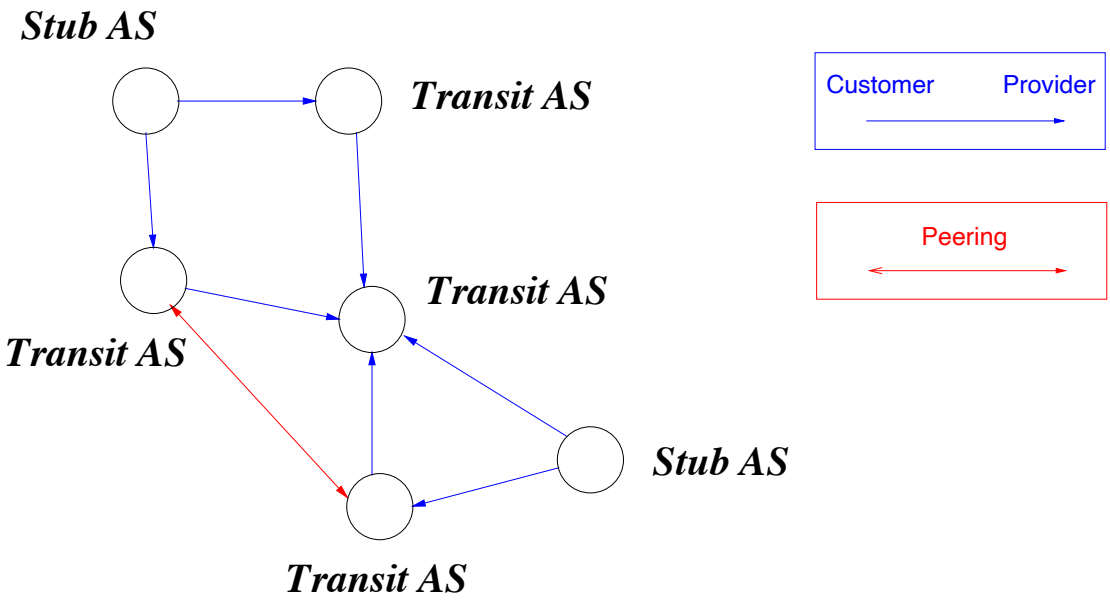

Inter-AS relationship: bilateral

- customer-provider: customer subscribes bandwidth from provider

  → customer can reach provider's reachable IP space

- peering:

  → only the peer's IP address and below

  → the peer's provider's address space: invisible

Common peering cases:

- among tier-1 providers
  - → ensures global reachability
  - → exclusive club
- among tier-2 providers
  - → regional providers
  - → economic factors
- among stubs
  - → economic factors
  - → e.g., content provider and access ("eyeball") provider
  - → e.g., Time Warner and AOL