CSCI 357: Algorithmic Game Theory

Lecture 21: Incentives in BGP Routing

Shikha Singh


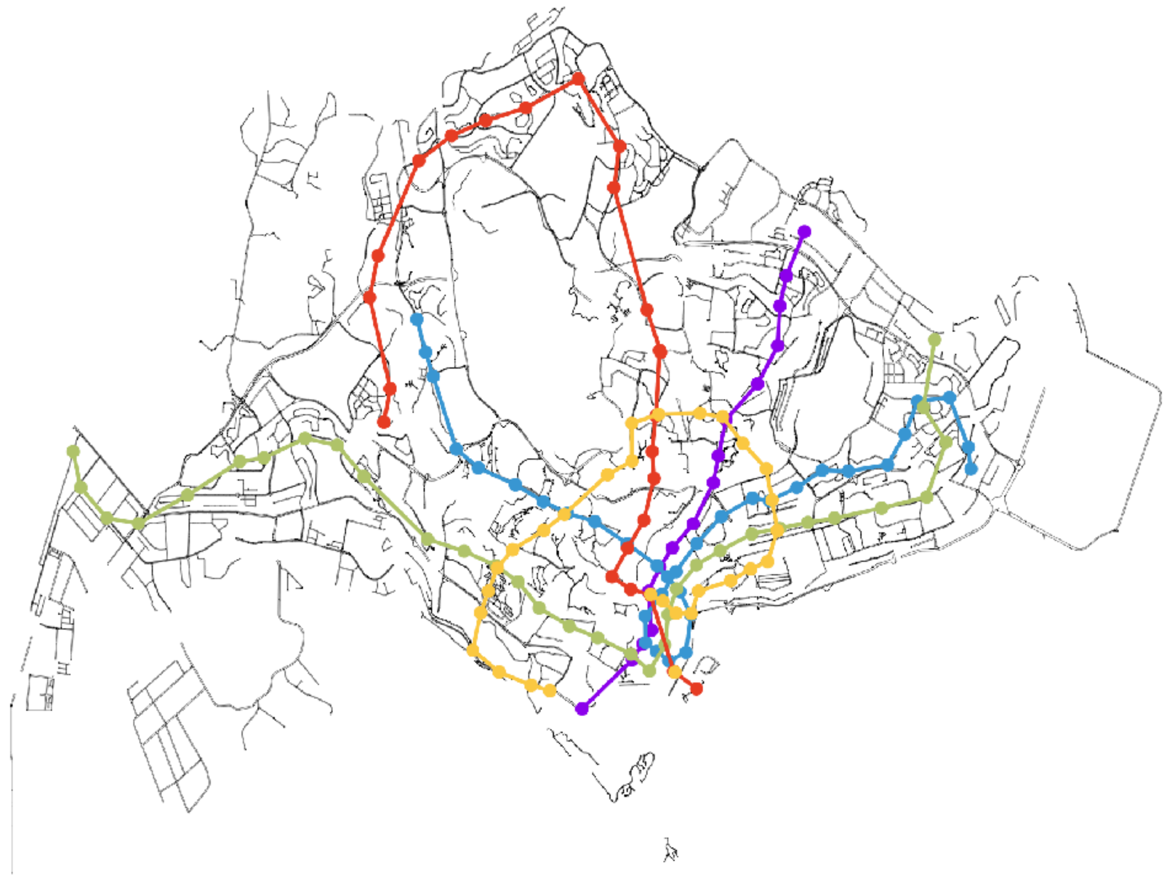

Announcements and Logistics

- **Midterm 2** graded feedback returned:
  - Mean: 90.4 % and Median: 91.5 %
  - Happy with class performance
  - Please let me know if you have questions about the feedback
- Colloquium attendance form emailed out
  - Any thoughts on the talk?

Questions?

Project Logistics

- Will go through 1-page abstracts and reach out if I have comments or concerns
- Project meetings instead of office hours now
  - Sign up ahead of time https://tinyurl.com/357projectmeet
  - Required to meet at least one more before classes end
- Student presentations in class next week
  - Indicate preference (Mon or Thurs) **by Wed 10 pm:** https://tinyurl.com/357present
  - Those who do not fill out will be assigned arbitrarily
- 2-page report (which includes background) due Friday 5 pm
- Final report will be an extension of the 2 page report

Questions?

Plan

- Today: Short lecture on incentives in routing traffic through the internet
  - Tips on project presentation
  - If we have time: stay back and discuss project progress/questions
- Thursday: Short lecture on incentives in cryptocurrencies
  - Course overview wrap up + **SCS forms**
  - Please bring your laptop to class on Thursday!

Routing in the Internet

- So far, we have discussed **"delay-based" routing** which is common in local-area-networks
- Now we discuss a different type of routing:
    - When routing traffic between different local networks
- The **Inter**net is not one network
    - A "network of networks"
- An **autonomous system (AS)** is a centrally controlled collection of routers: a bunch of routers with a common admin
- The internet has around **42,000 ASes**
    - ISPs, universities, businesses, etc.

Autonomous System: controlled by one entity

Comcast network. Source: business.comcast.com

Routing in the Internet

- How do we transmit information from one AS to another?
  - Some ASes are physically connected
  - Others are not and may need to route through one or many intermediate ASes
- **Questions**.
  - Who pays whom?
  - How do we route traffic within an AS
  - How do we coordinate routing across ASes?
  - All of this is relevant to study incentives

Routing Within an AS

- Routing within an AS: usually **shortest path protocols**
- Resemble Dijkstra and Bellman-Ford algorithms from 256
- What are the edge weights?
    - If all edge weights are 1, we only care about "number of hops"
        - Can use a simpler algorithm than Dijkstra/Bellman-Ford
        - Just breadth-first search works
    - Edge weights might depend on "recently delay"
- All routers essentially agree on the best paths
    - Dictated by a central administrator

Comcast network. Source: business.comcast.com

Routing Between ASes

- No central administrator to coordinate
- Completely decentralized network (ASes span the whole globe)
- Traffic is routed based on various agreements:
    - **Paid transit**: An ISP might buy access from another larger ISP, e.g. Williams probably does this
    - **Peering**: two ISPs may agree to exchange traffic for free/reduced rate
- To understand this:
    - ISPs are classified into tiers

Tiers of ISPs

- **Tier 1**: Never buys traffic.
  - Can reach the entire internet through peering or its own infrastructure
- **Examples**: **AT&T**, Sprint, Tata, Telia, Level3 (now part of CenturyLink)
- **Tier 2/3**: Has to pay other providers for access to some parts of the internet
- **Examples**: Williams, **Comcast** (barely), small cable companies
- Different ISPs work out mutual agreements

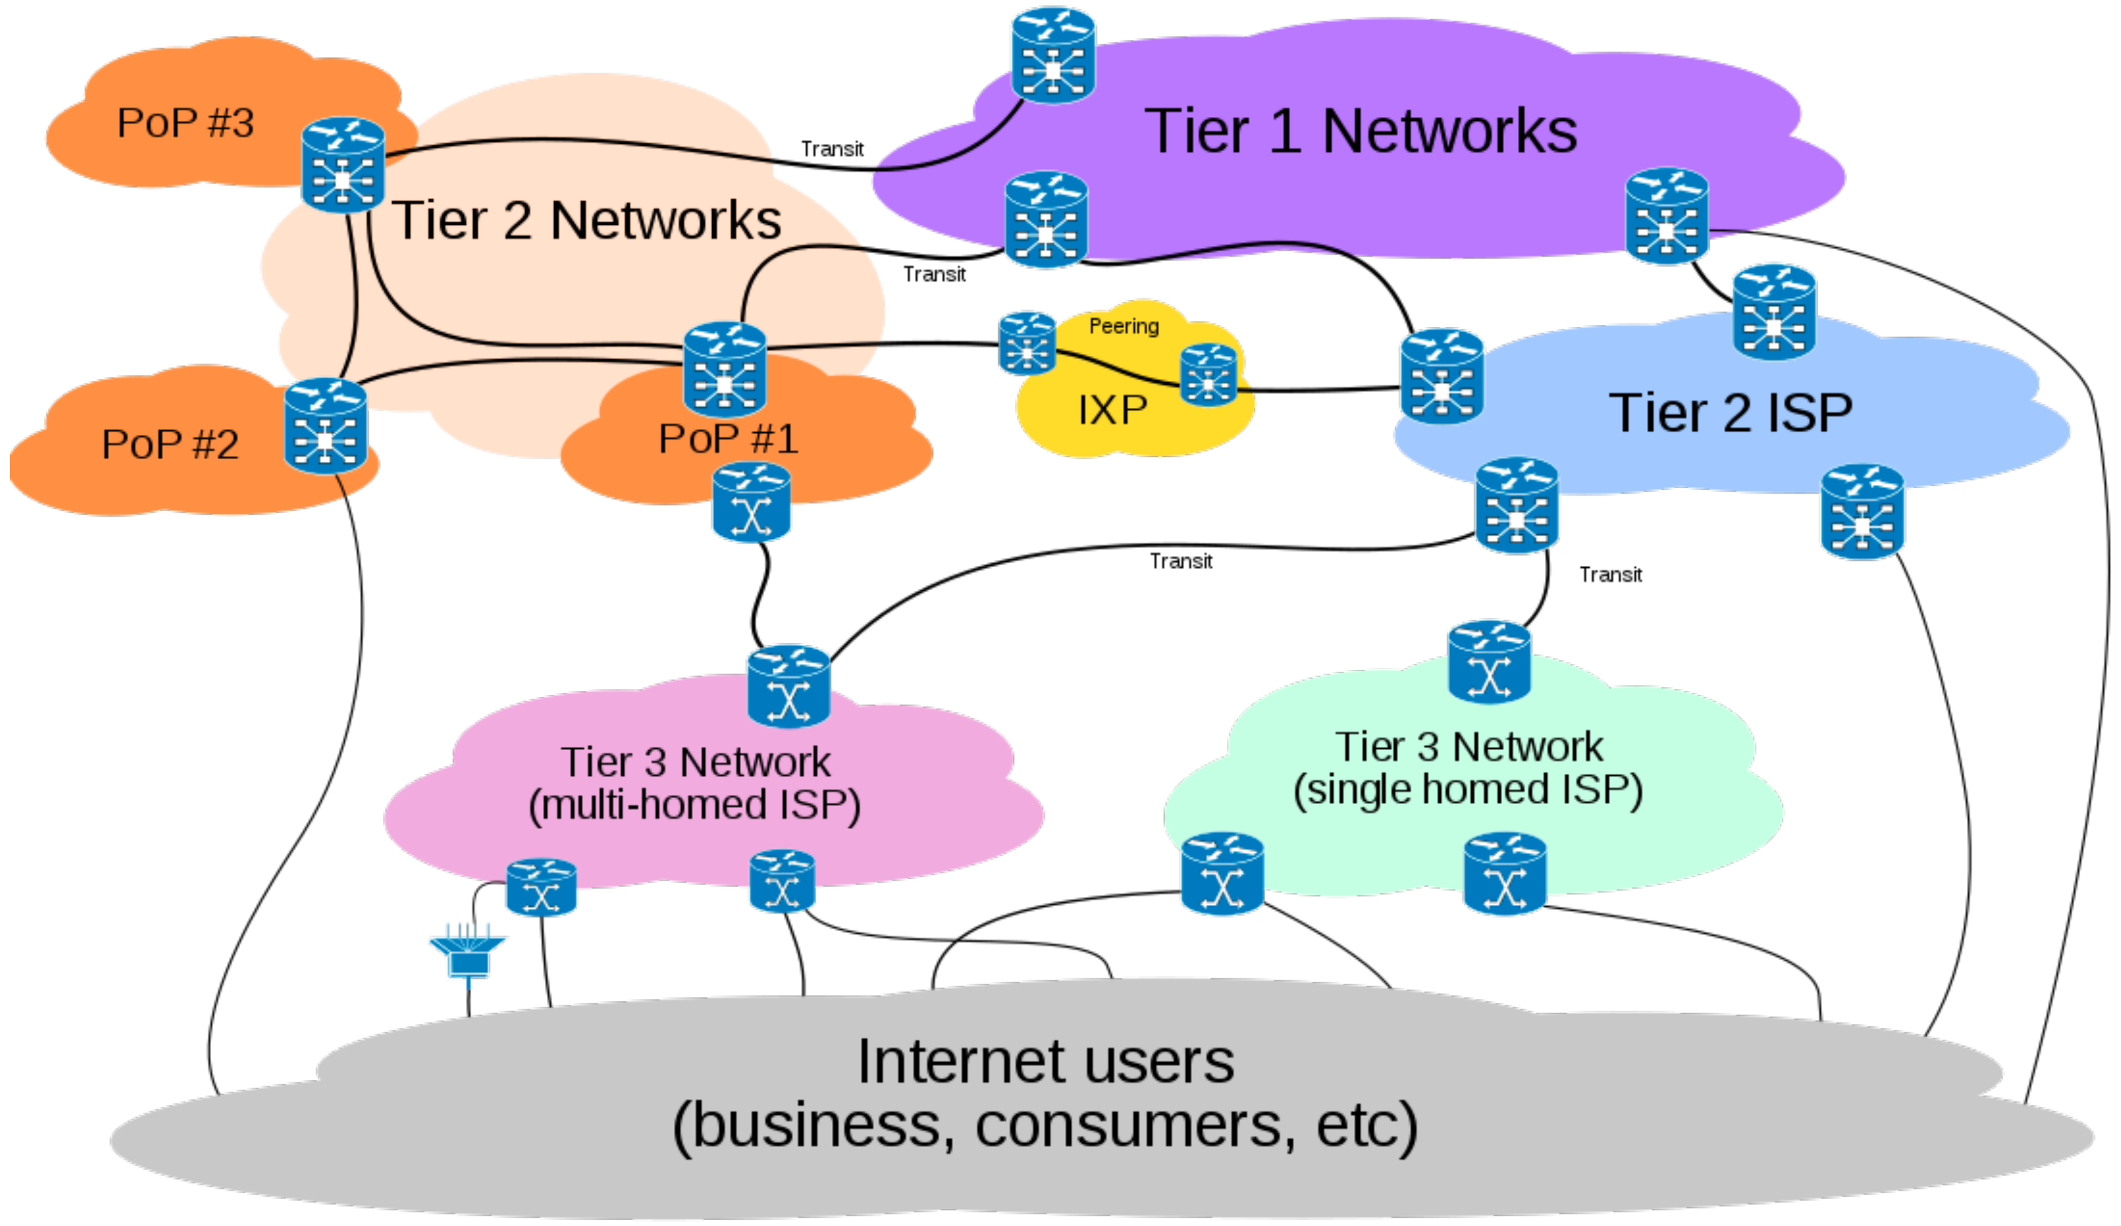

https://en.wikipedia.org/wiki/Tier_1_network#List_of_Tier_1_networks

The AT&T Global Network —

Advanced and Powerful Network Carries More Than 8.52 Petabytes of Data Traffic* on Average Business Day

- Multiprotocol Label Switching (MPLS)-based services** available to 137 countries over 1,600+ nodes.
- 130,000 MPLS customer ports.
- 32 Internet data centers across the globe.
- 535,000 fiber route miles.

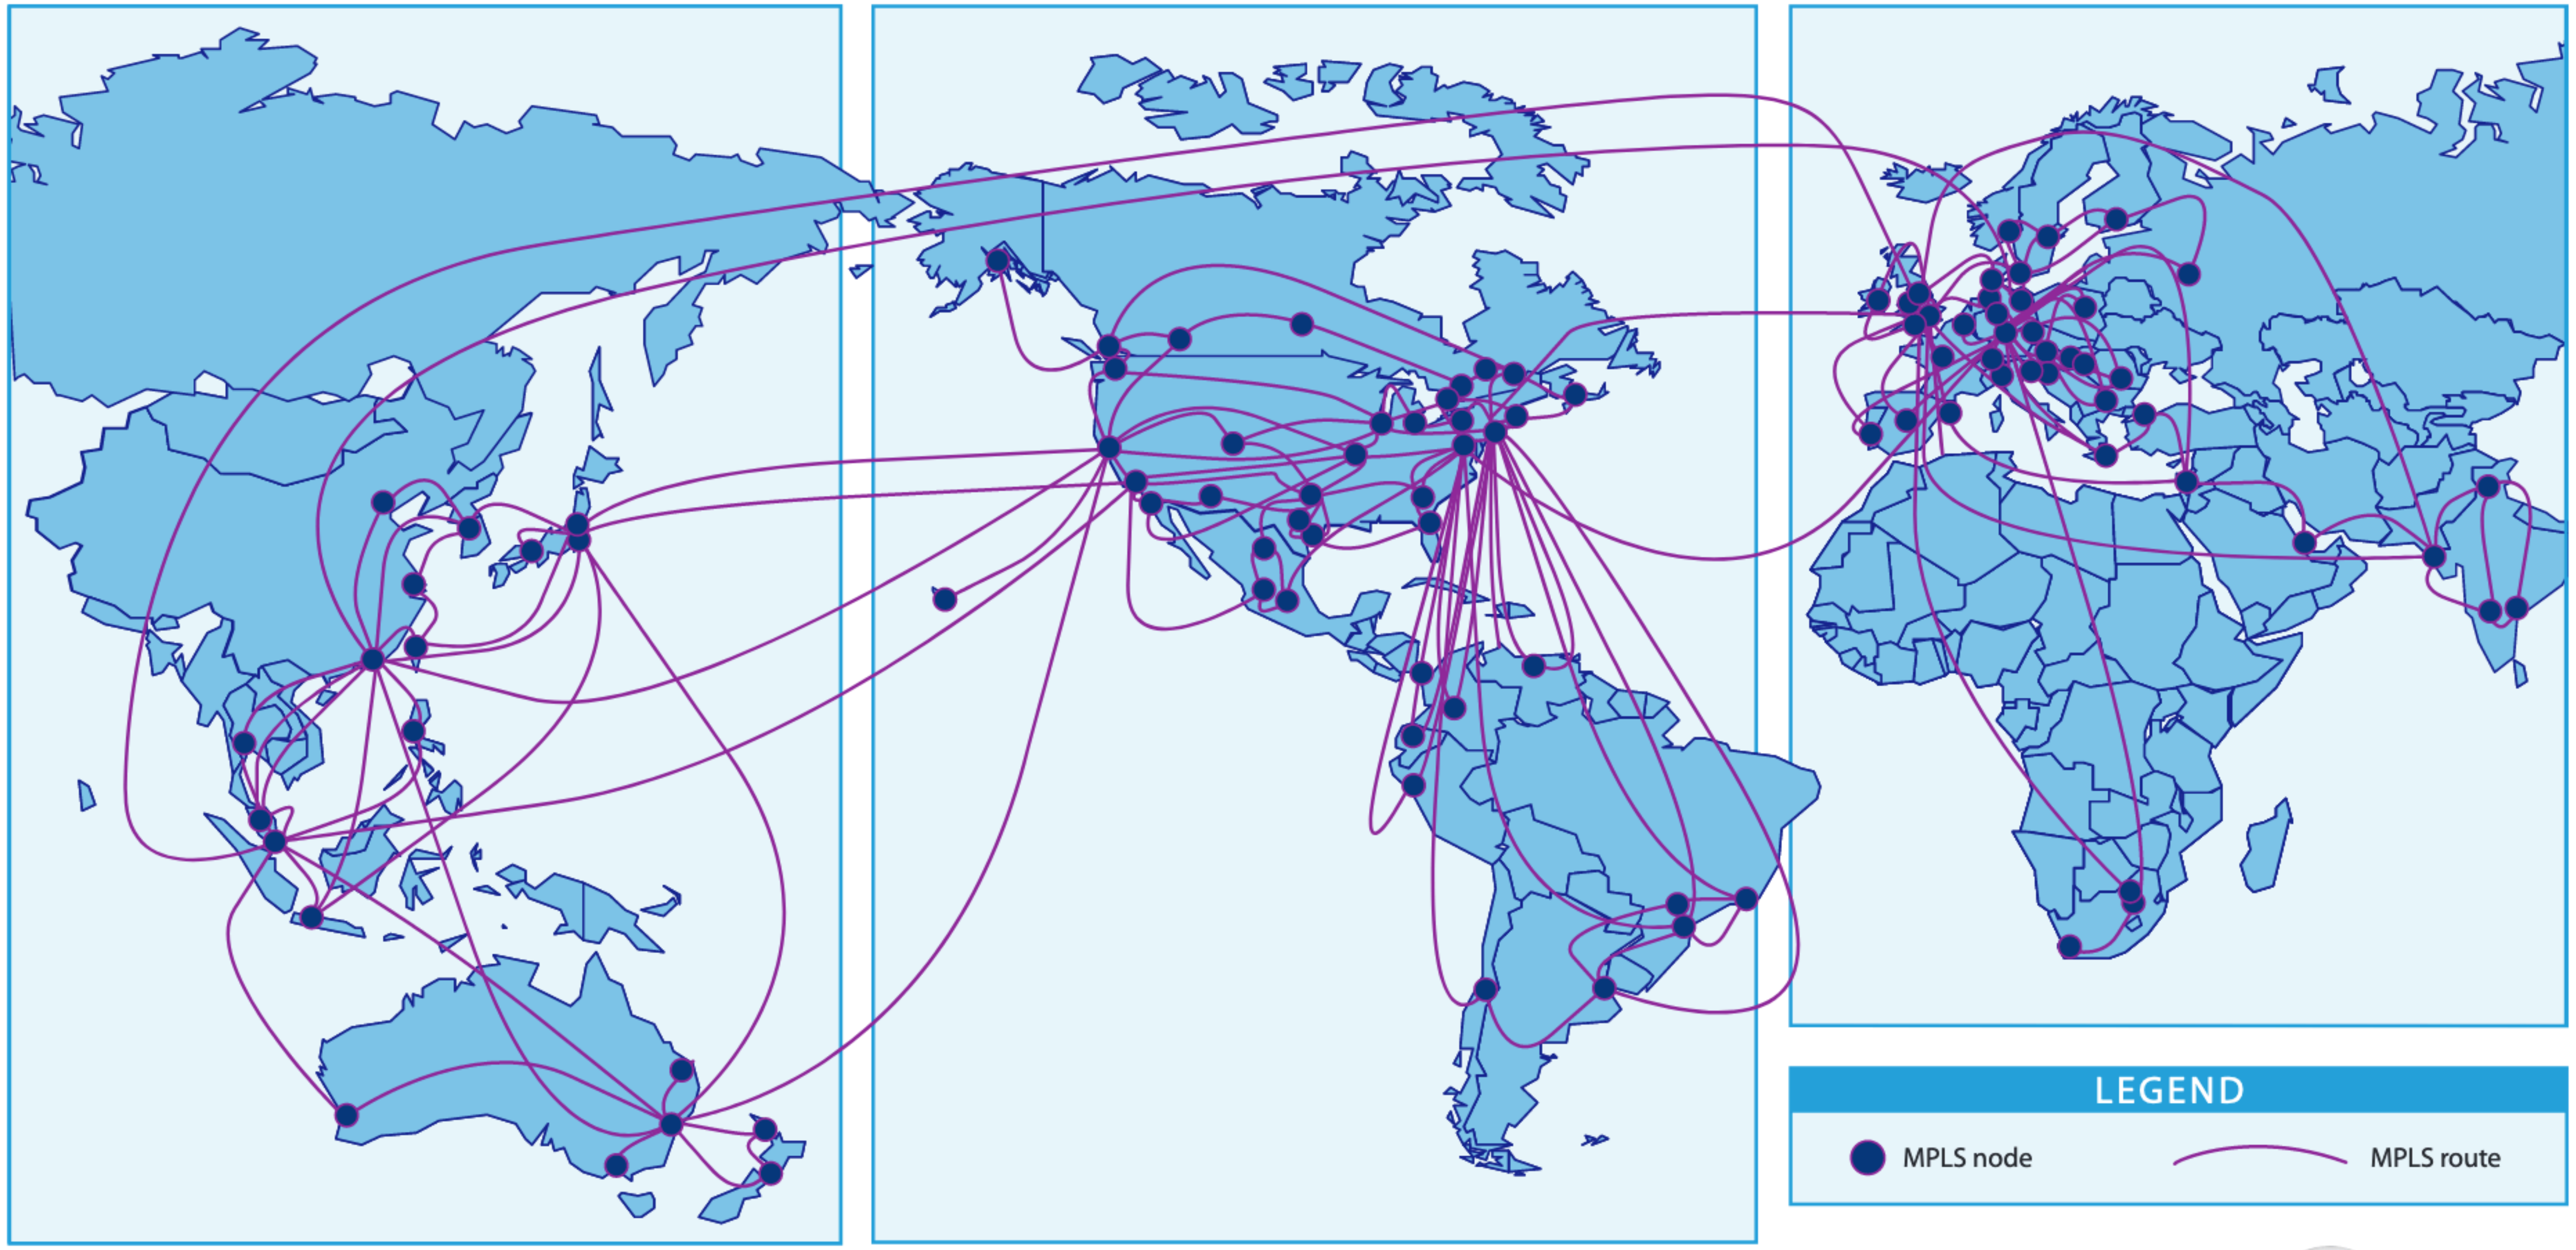

Simplified map: not all nodes/links/routes shown

* Enough data to transmit the digitized contents of the Library of Congress more than 400 times every day.

** MPLS technology enables high-quality delivery of multiple services over a single IP network infrastructure.

at&t

How Does Traffic Work

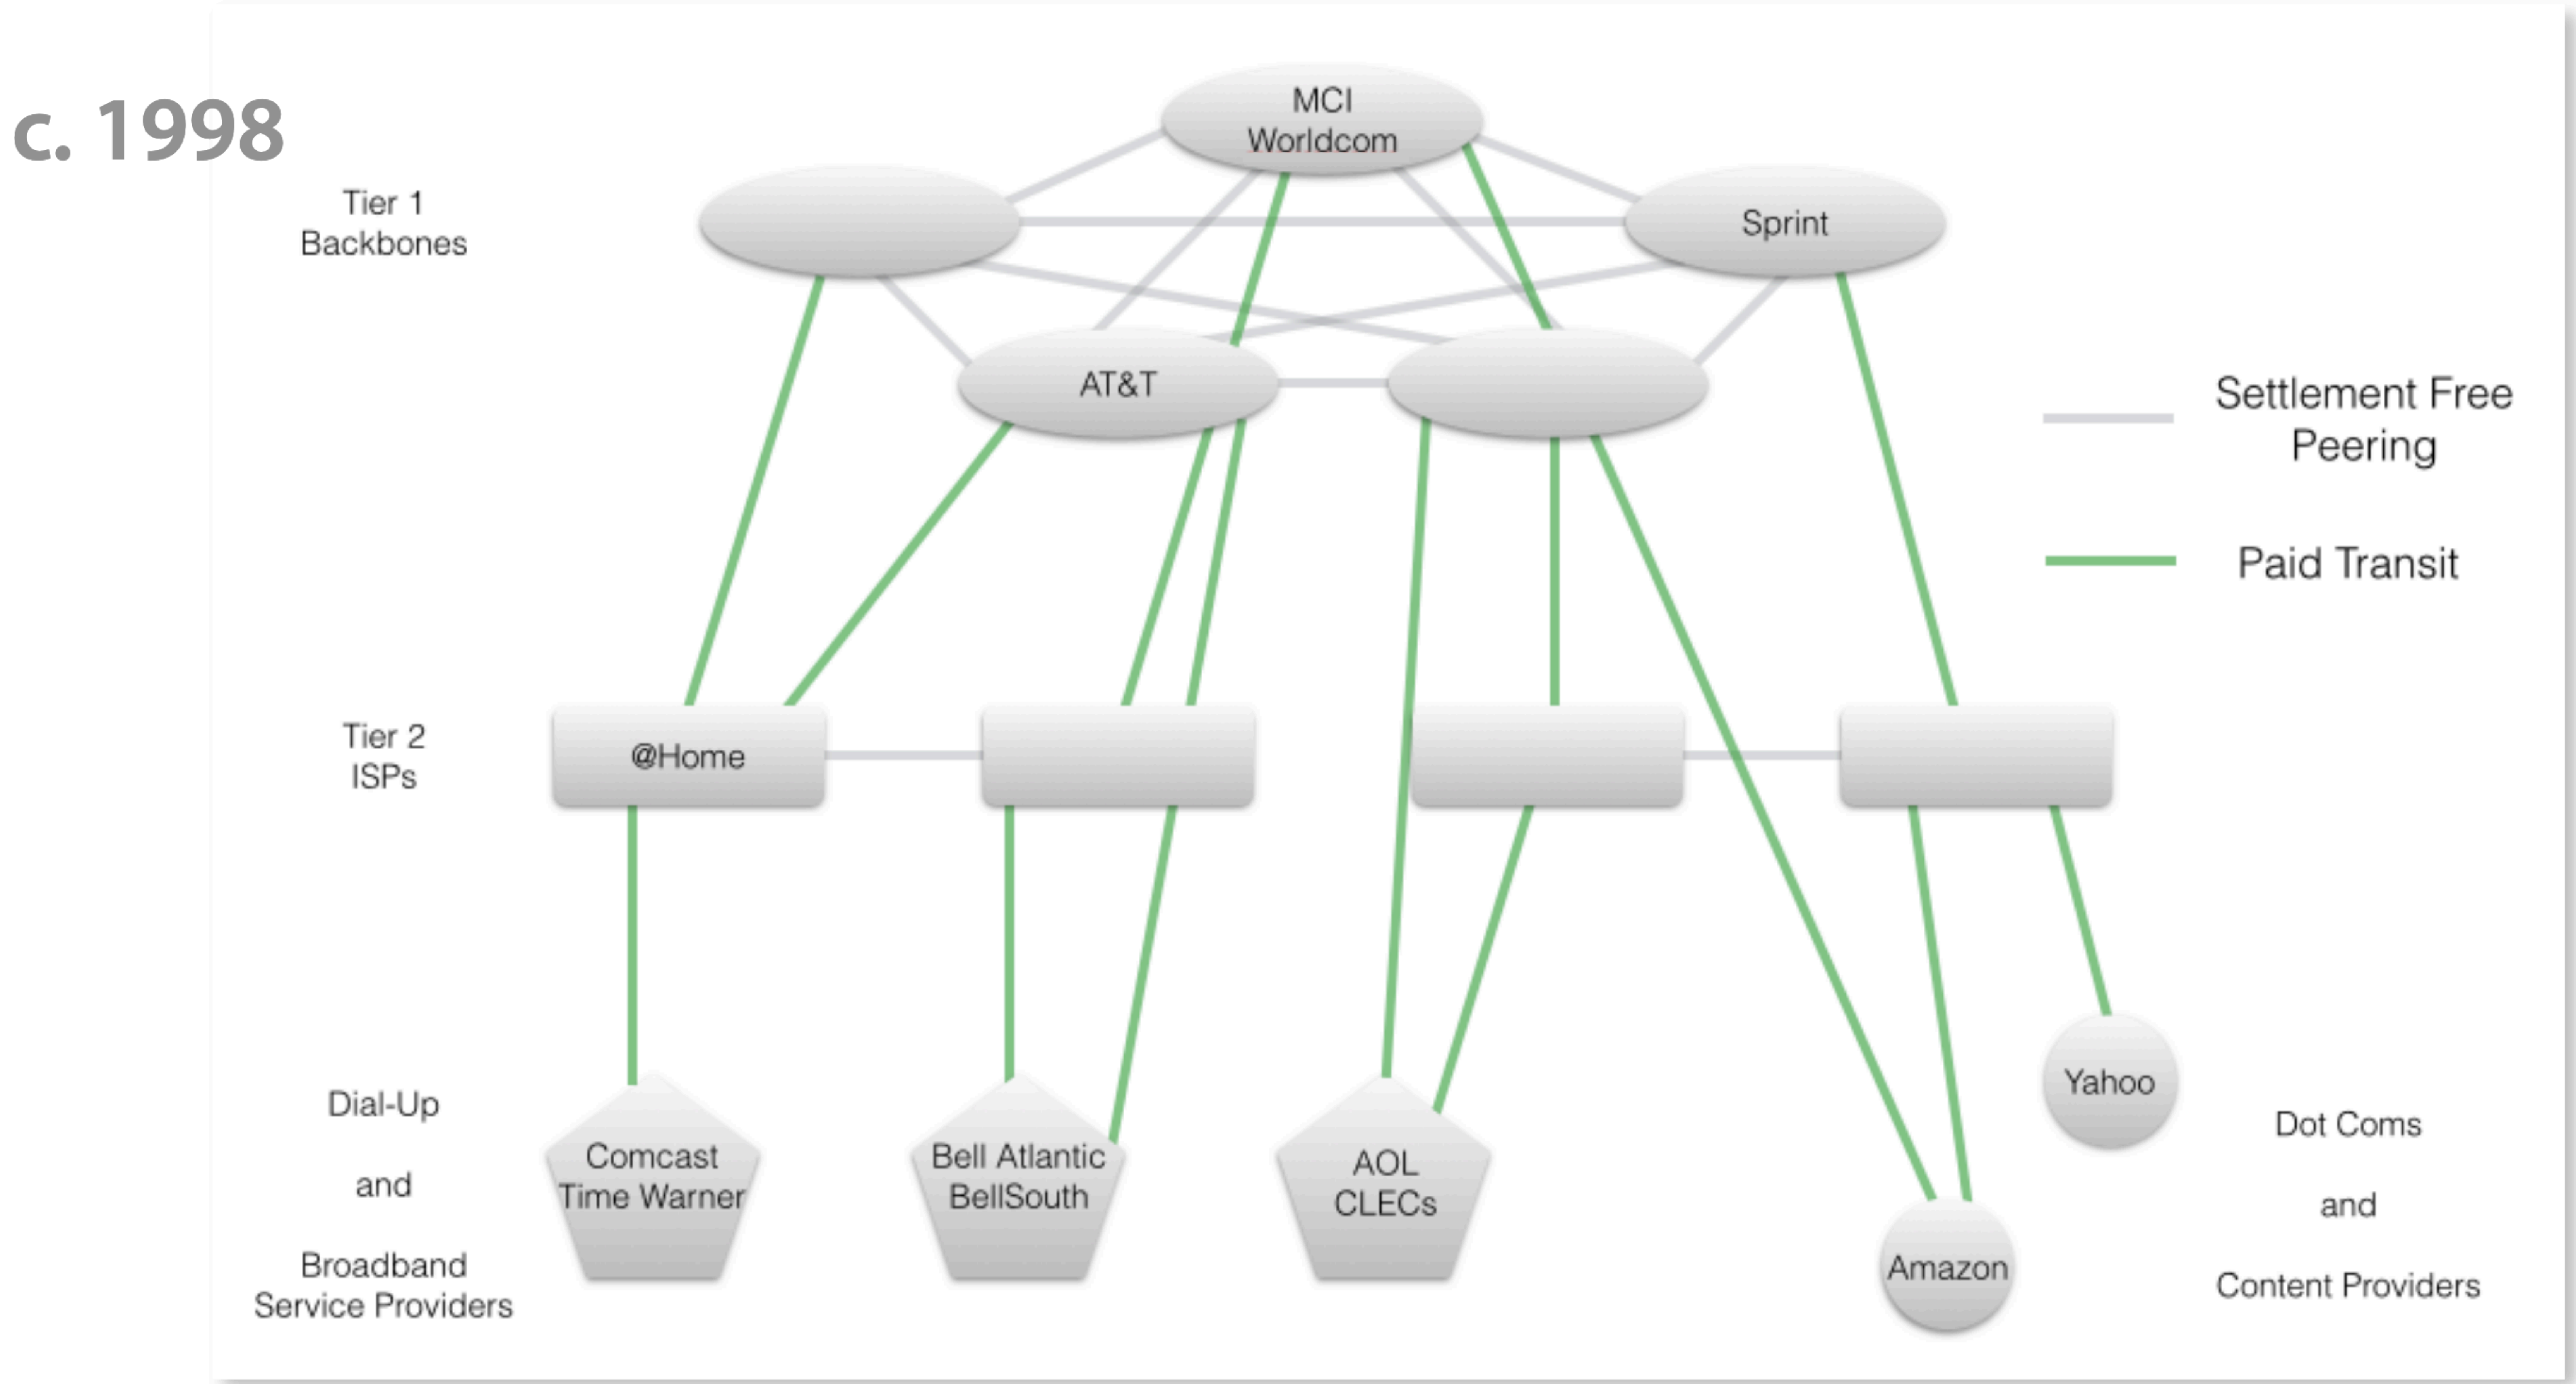

From "How the Net Works: !A Brief History of Internet Interconnection" by Bret Swanson

Peering Agreements

- Tier 2 ISPs can help each other out by directly exchanging traffic for free (avoiding paid network)
- Tier 2s can come to a peering agreement and exchange traffic through an "**internet exchange point"**
    - Many of these: https://www.internetexchangemap.com/
- Side note: who controls these?
    - Many owned by **Packet Clearing House**
    - International nonprofit
    - Also controls DNS

Internet in 1998

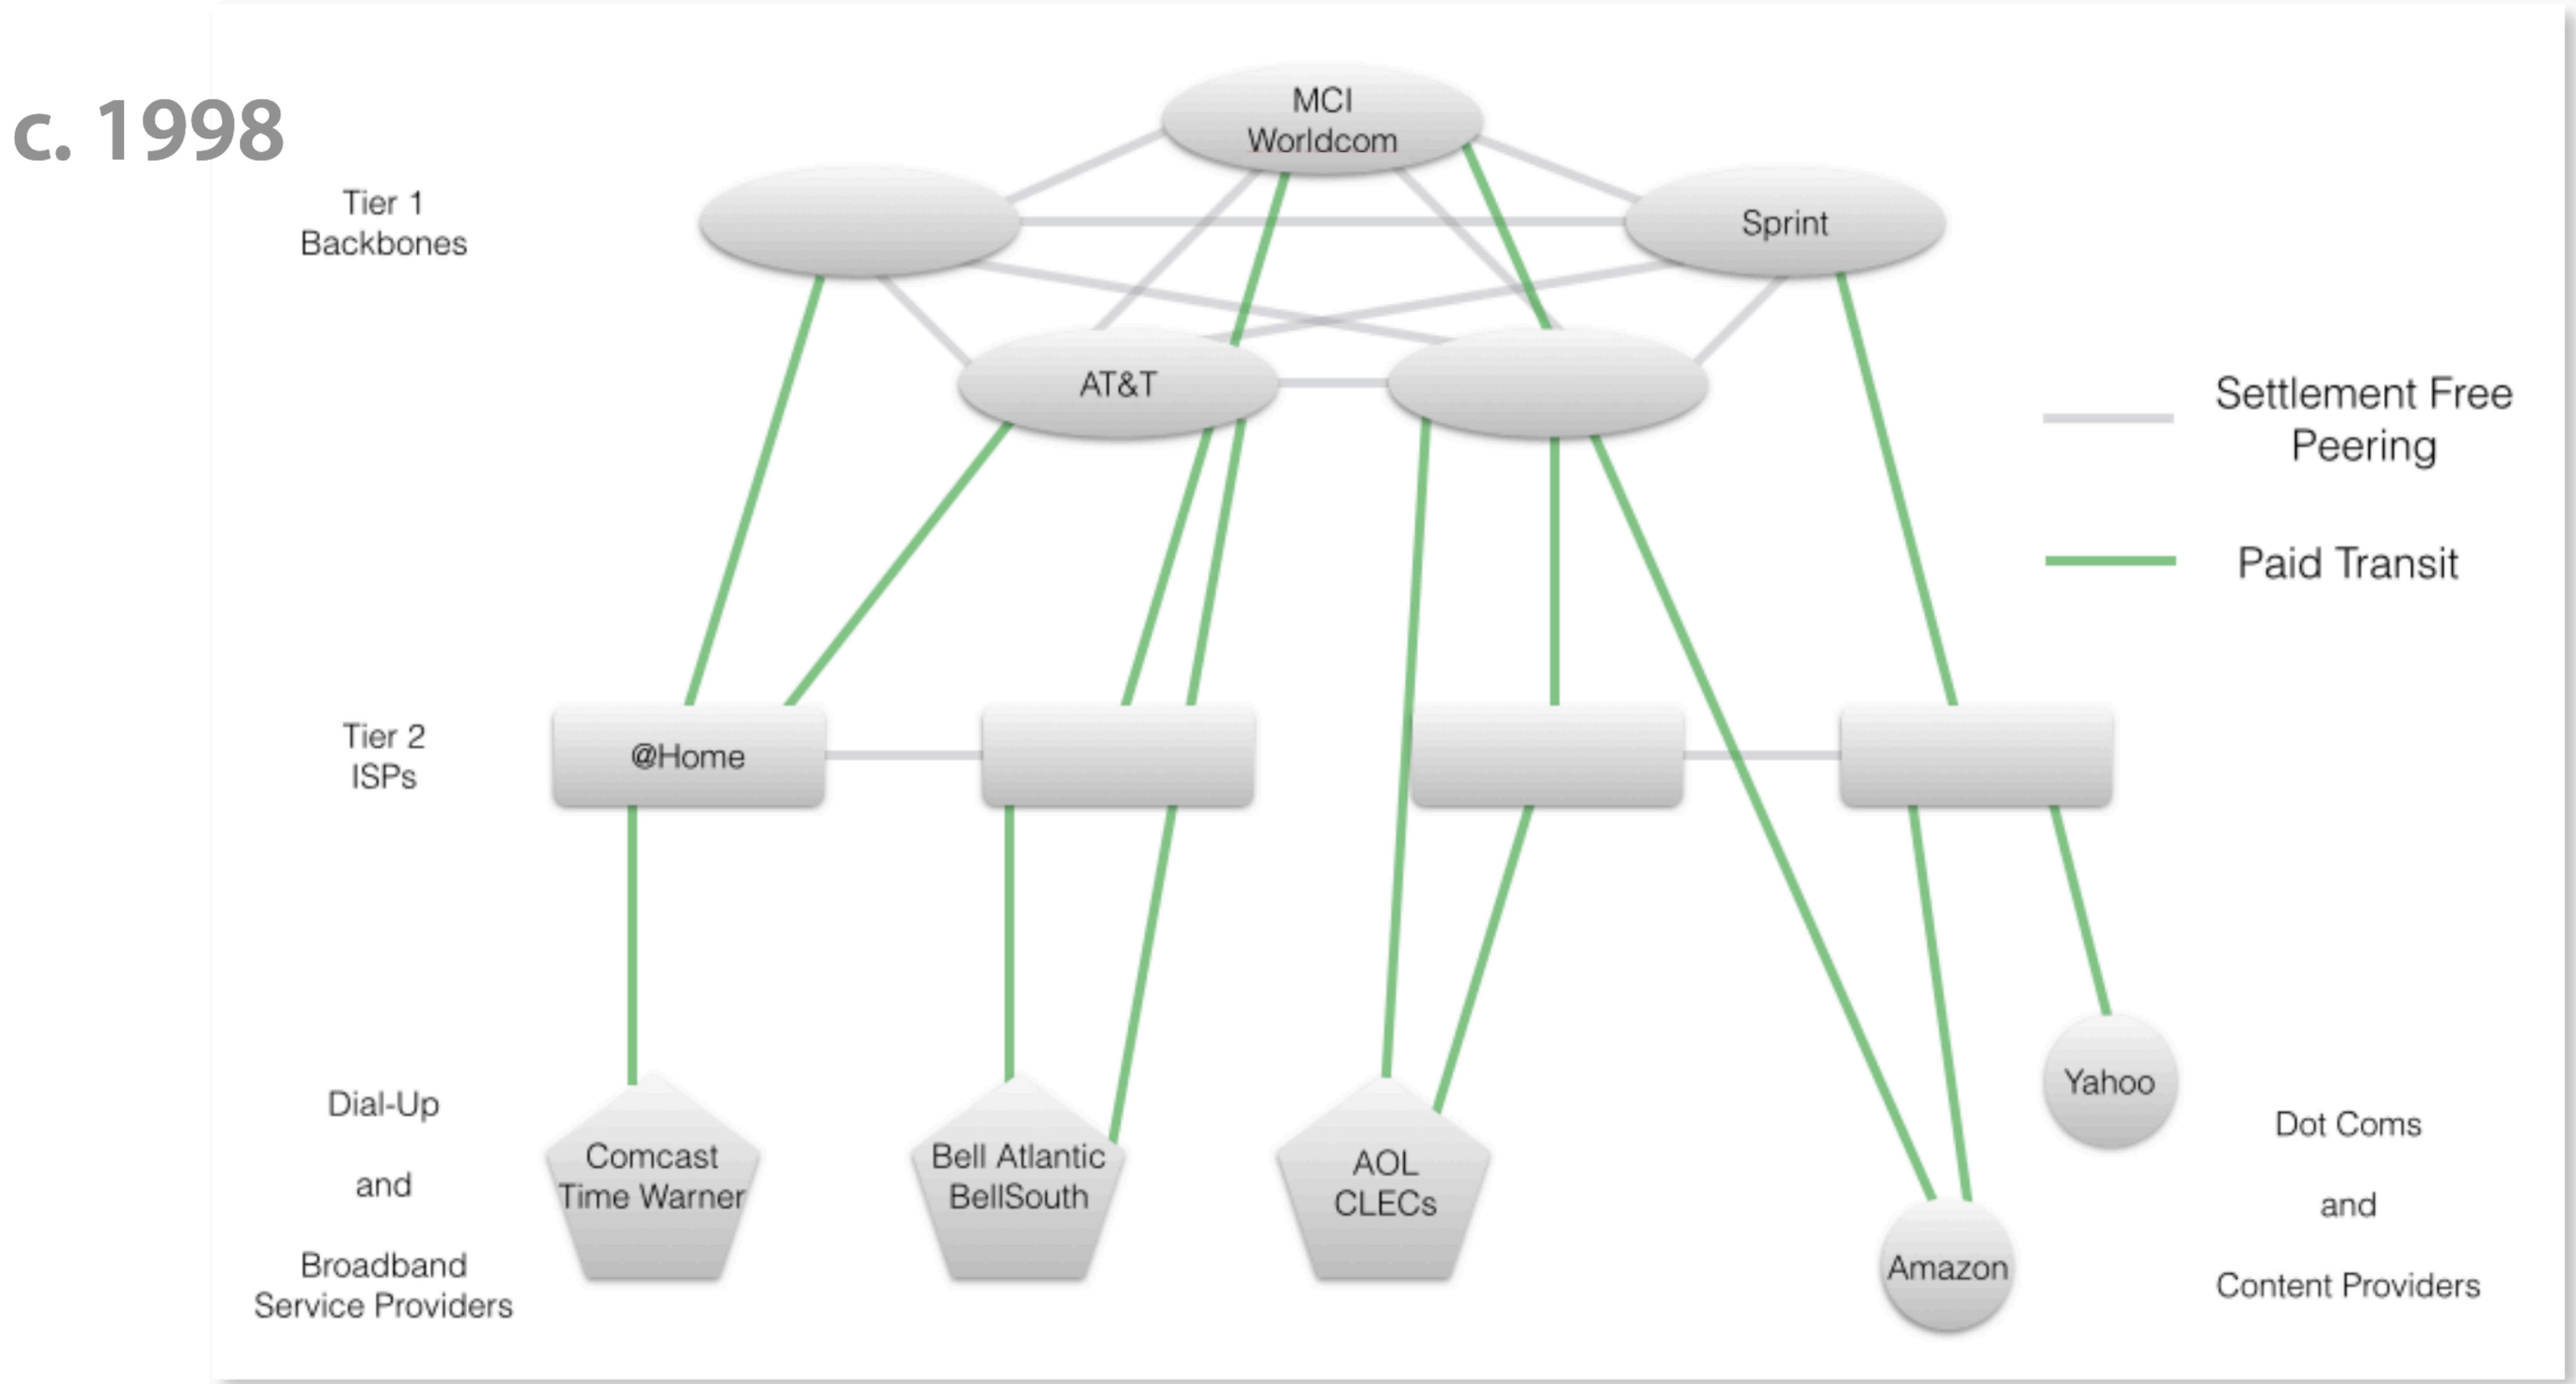

From “How the Net Works: !A Brief History of Internet Interconnection” by Bret Swanson

Internet in 2014

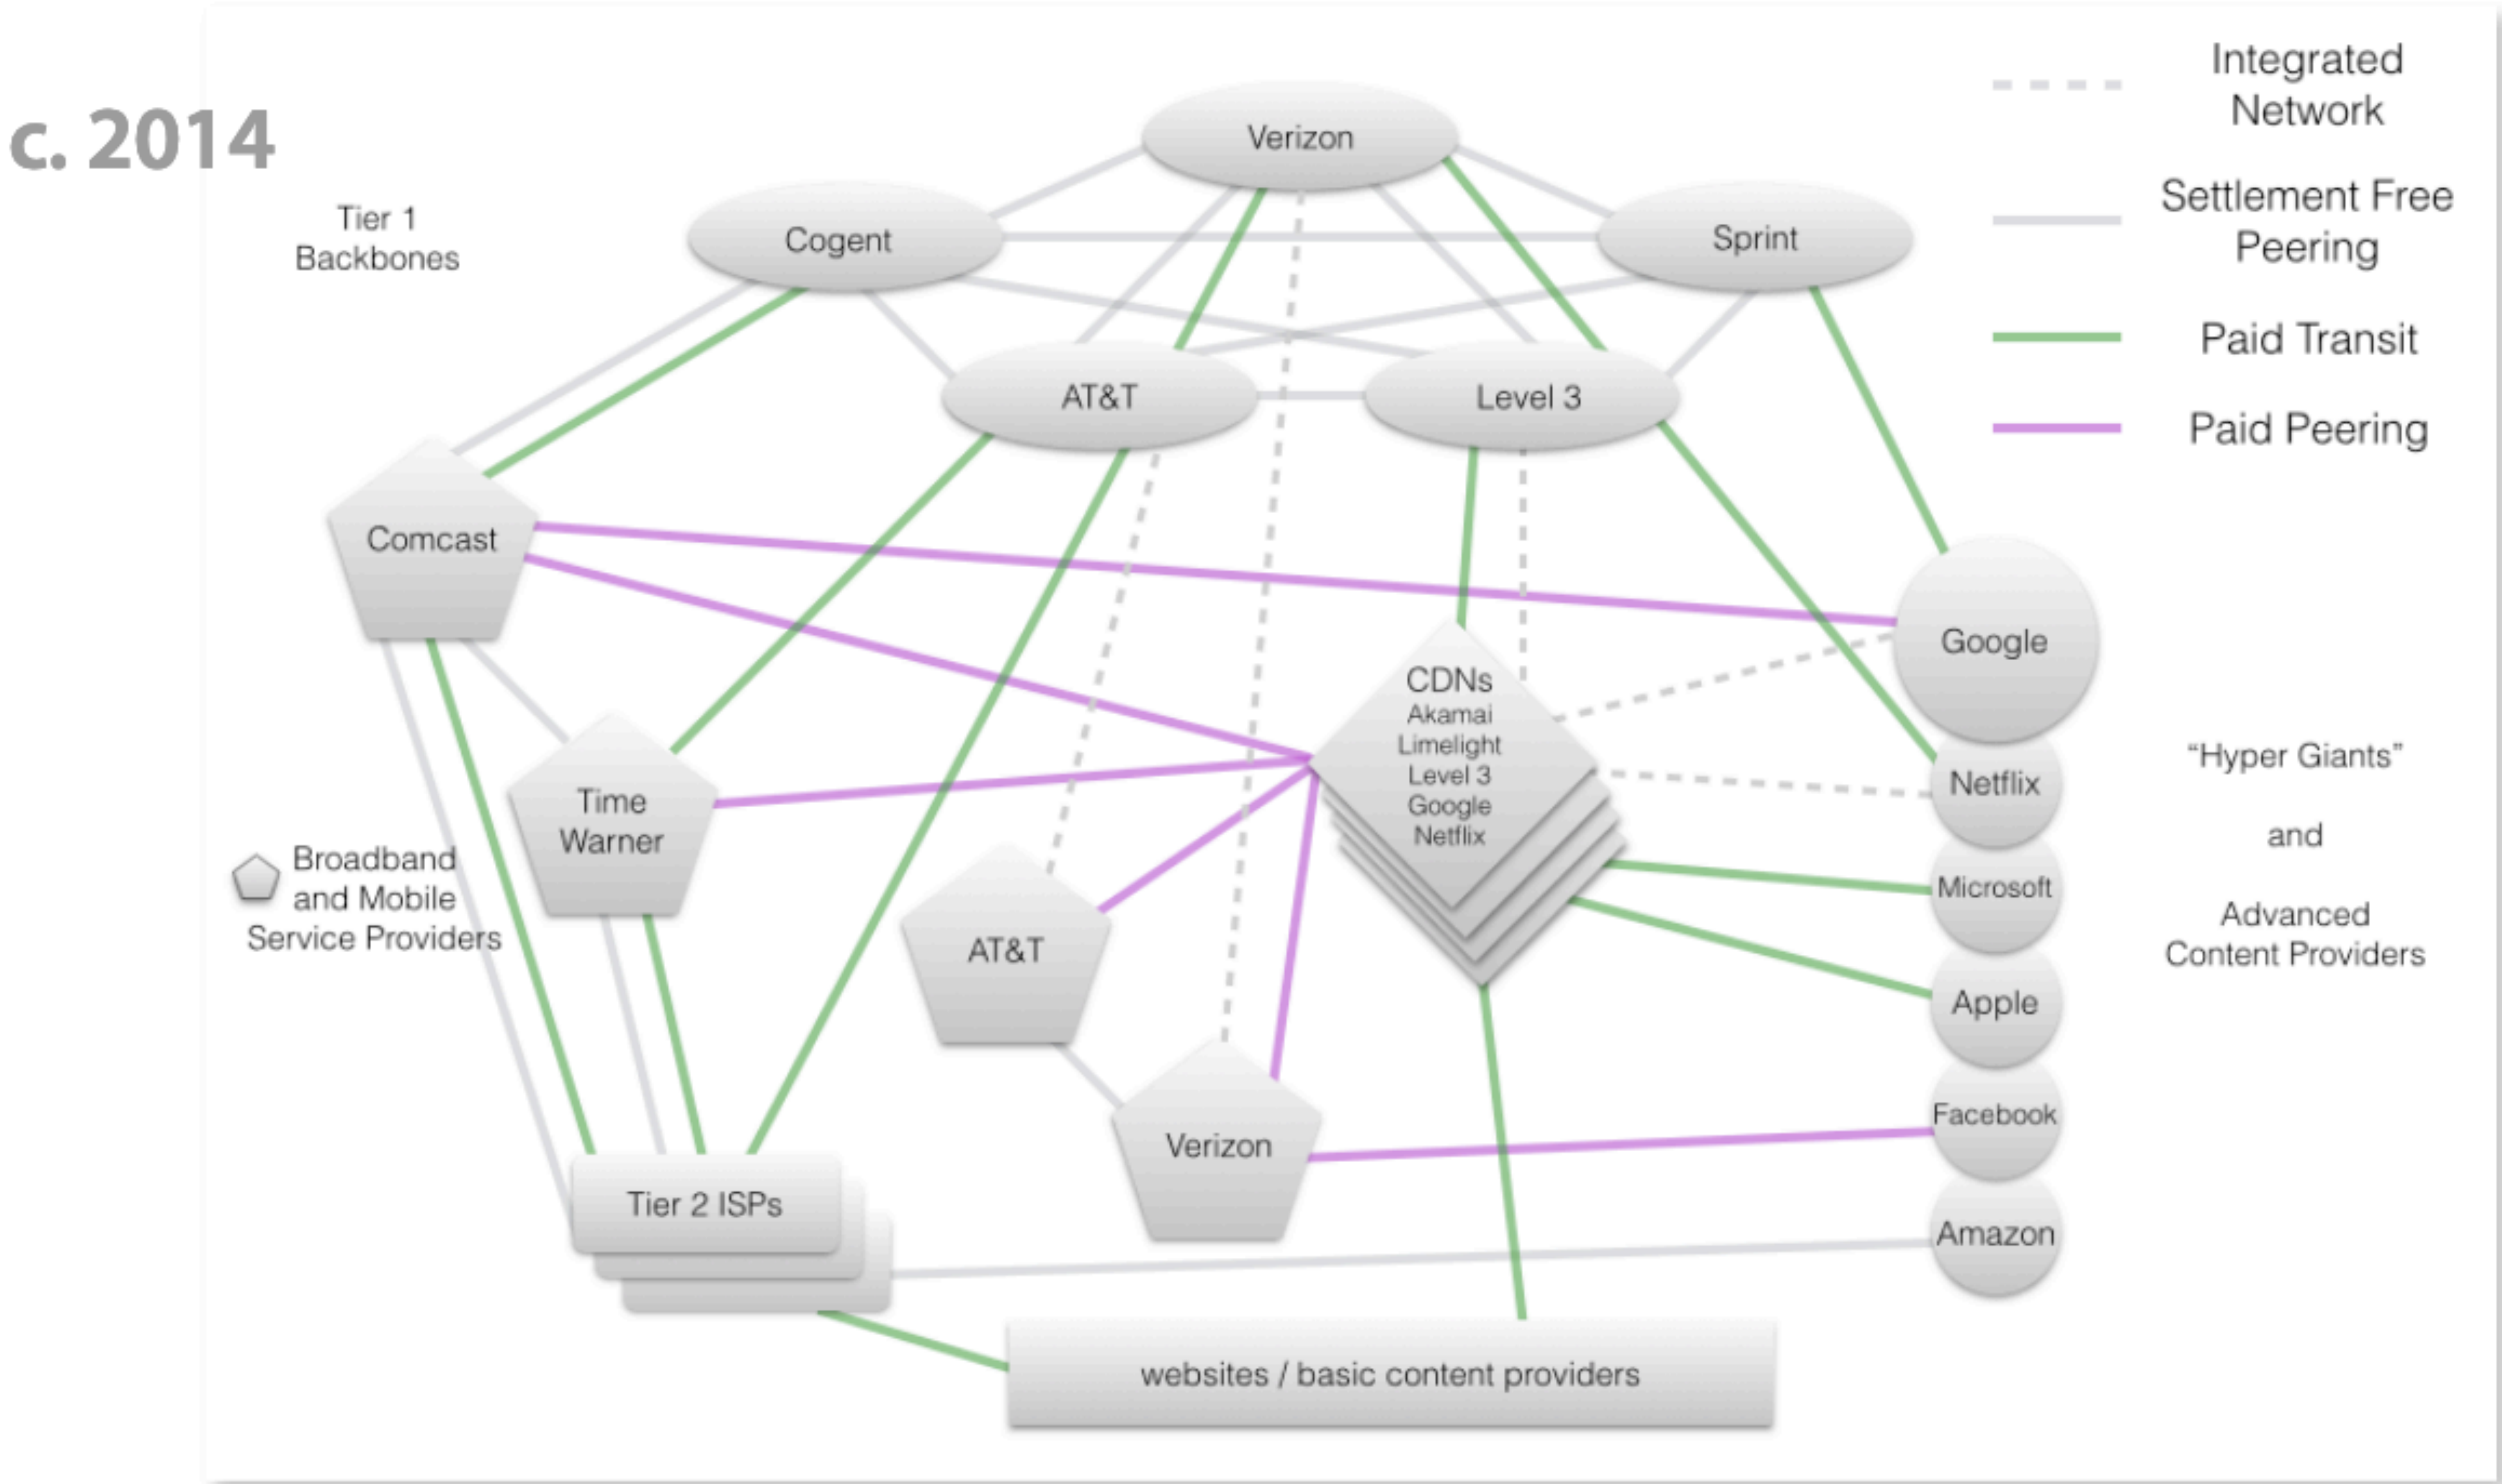

From "How the Net Works: !A Brief History of Internet Interconnection" by Bret Swanson

Routing Between ASes

- How can you find a route from point A to B when they are located in different ASes
    - Potentially requires traversing multiple ASes
- Different ASes may have different preferences:
    - May want to minimize monetary cost of paths
- For example: consider the network given in figure
    - Suppose d is destination for all traffic
    - ASes 1,2 have their preferences labelled
    - Both prefer to route through the other than direct

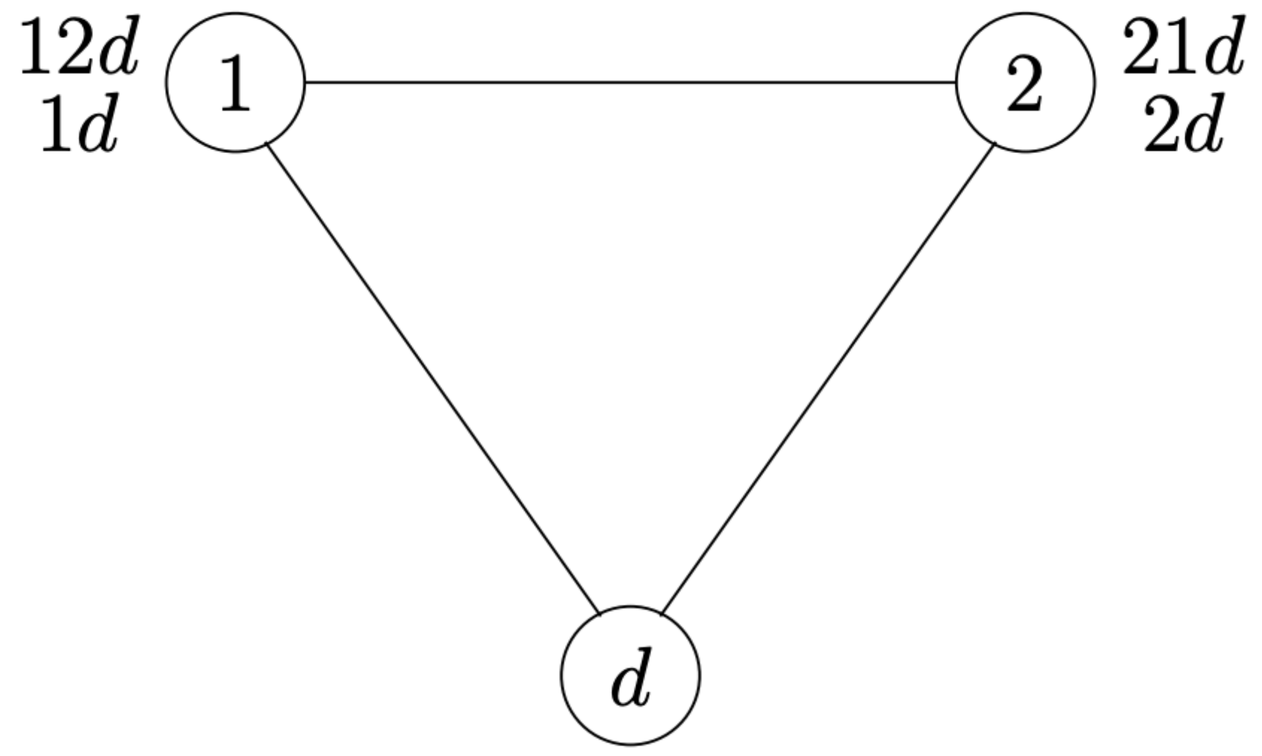

Credit: Roughgarden

Border Gateway Protocol (BGP)

- In the internet, routing between ASes is done using the **Border Gateway Protocol** (BGP)
- We will only discuss a sketch of how it works
    - Ignore many details of the actual protocol
- Fix a destination d (BGP runs in parallel for all choices of d)
- Each destination d broadcasts their presence to neighbor ASes
- Each AS is then supposed to update their own path to d and broadcast to their neighboring ASes
    - All messages are asynchronous

BGP Updates

- At an AS u with neighbors N
    - For each $v \in N$
        - Let P_v be the last path (from v to d) that v announced to u (if no path then is empty)
    - Reset P_u to u's favorite cycle-free path of the form $(u, v) \circ P_v$ (if any, otherwise empty); $\cdot$ is concatenation
    - If P_u changes, announce the new value to all neighbors $v \in N$
- Note that AS u has to avoid cycles: if P_v includes u, cannot route traffic through P_v

This is the "intended behavior": of course ASes may not follow it!

BGP Example

- Multiple possible fixed points of running BGP
- Which fixed point do we expect to reach?
    - Depends on timing of messages: whichever 1 or 2 finds out first that the other is using a direct path to d can switch and "win"

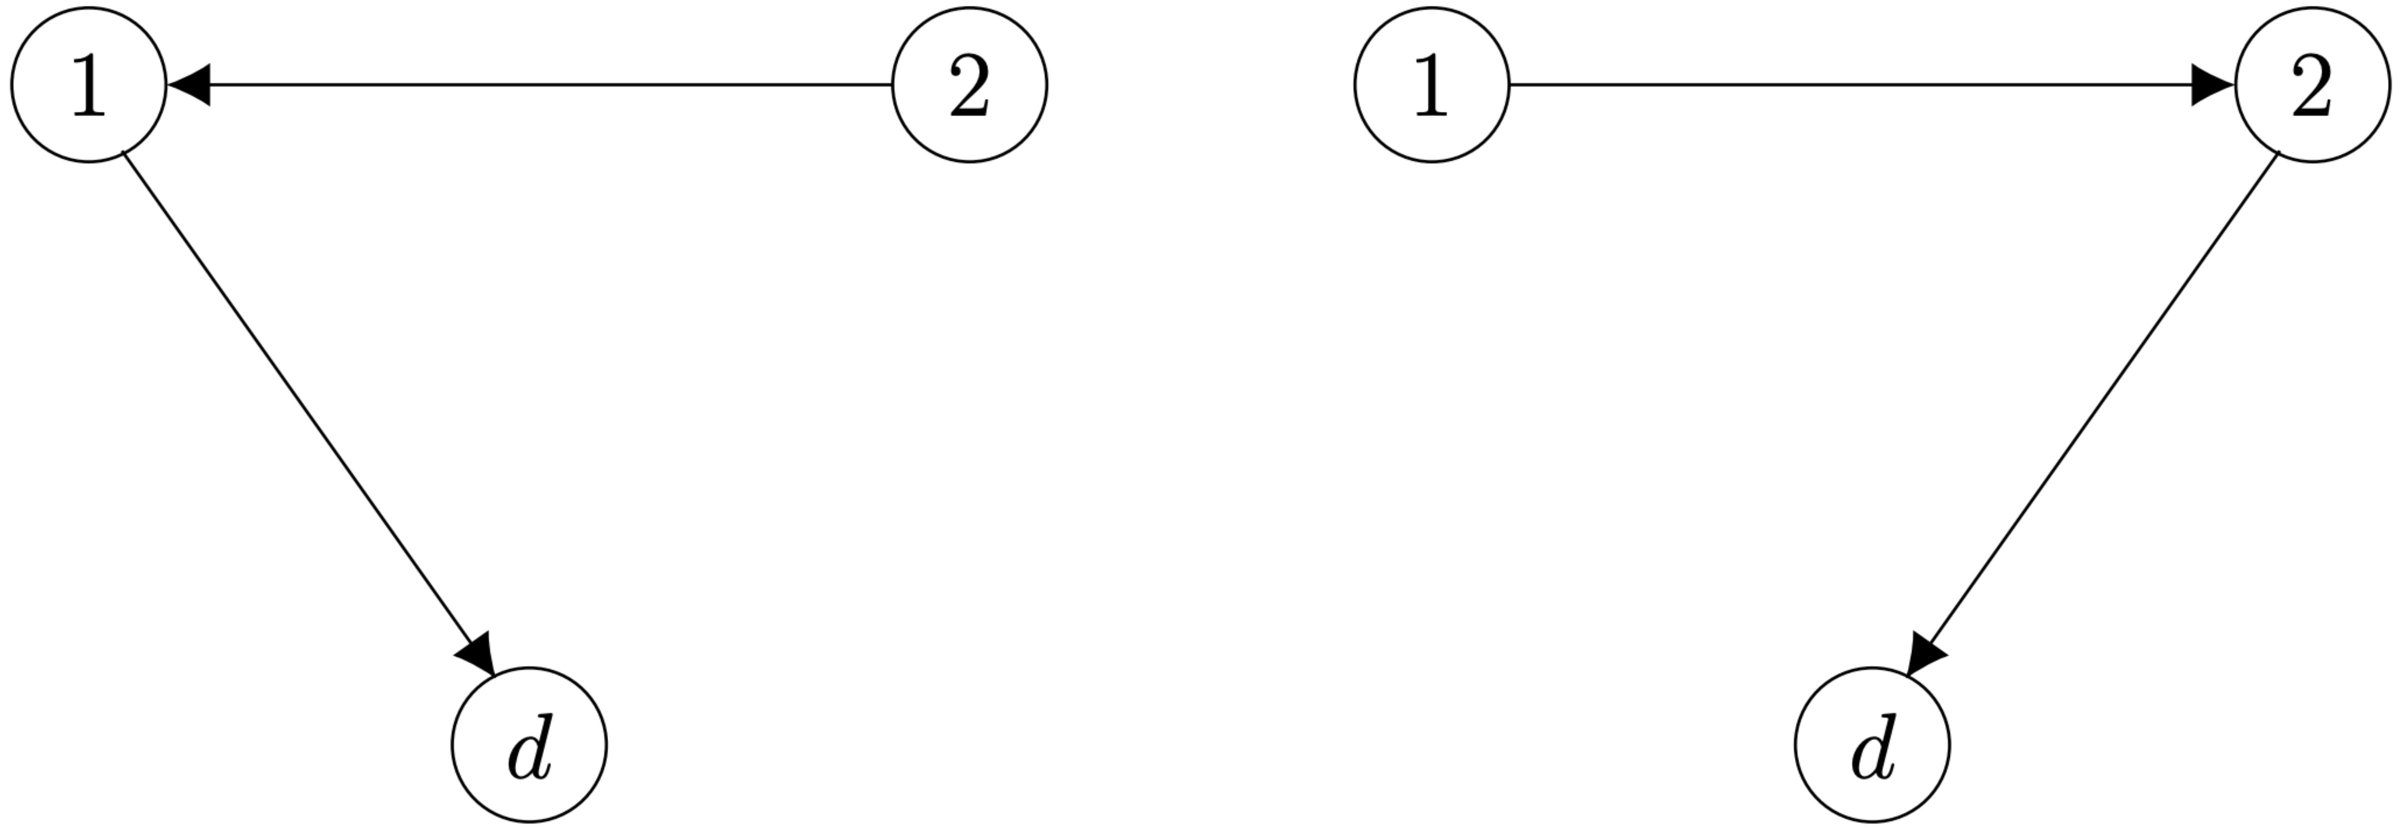

(a) (b)

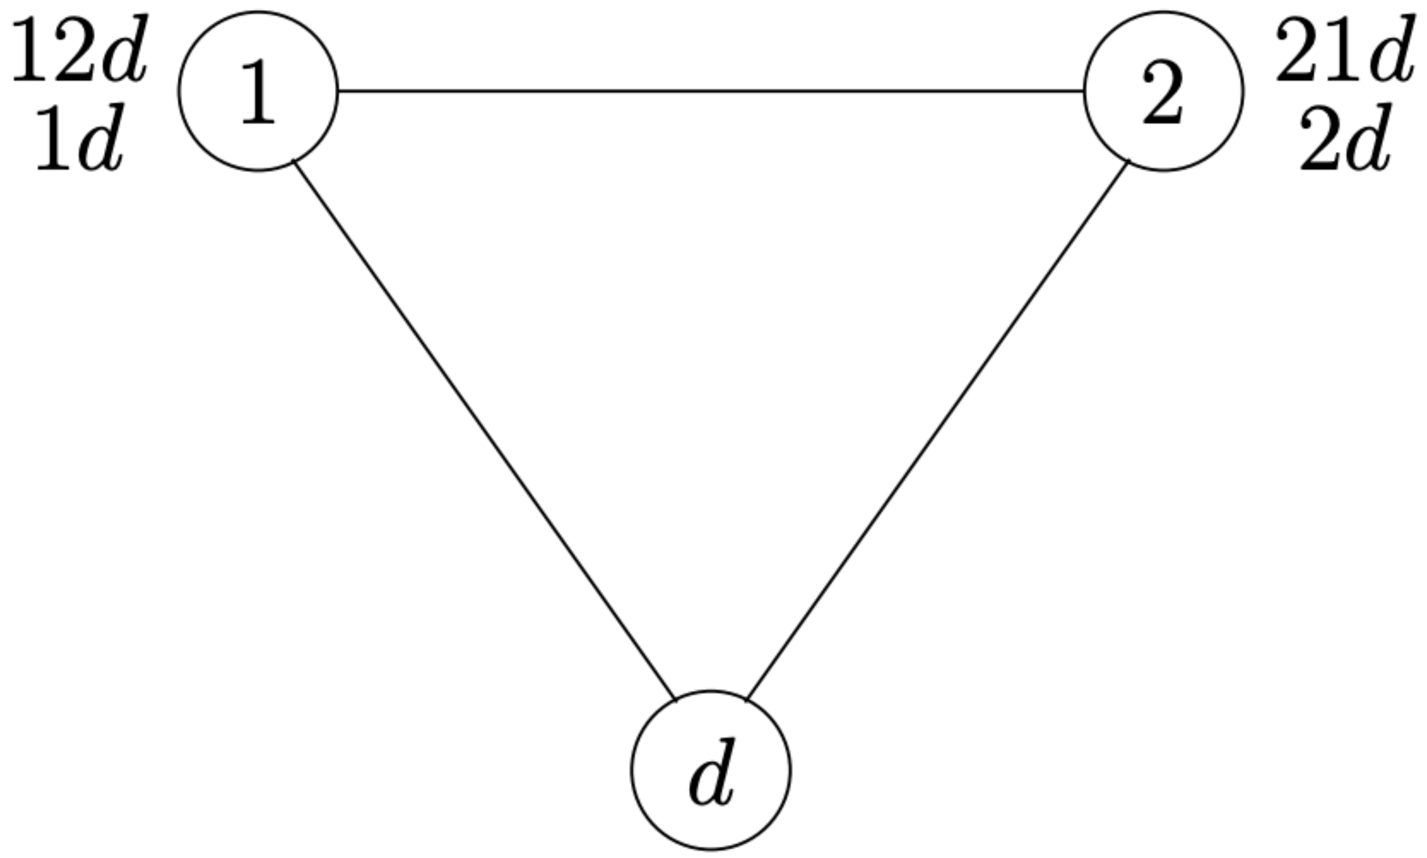

Credit: Roughgarden

Stable Routings: Equilibrium

- In a **stable routing**, no AS wants to change its path to d, given the choices of other ASes and options available to u
- This is a **Nash equilibrium** in a game where
  - **Players**: AS, available **strategies**: neighboring ASes, **payoffs** induced by preferences over paths

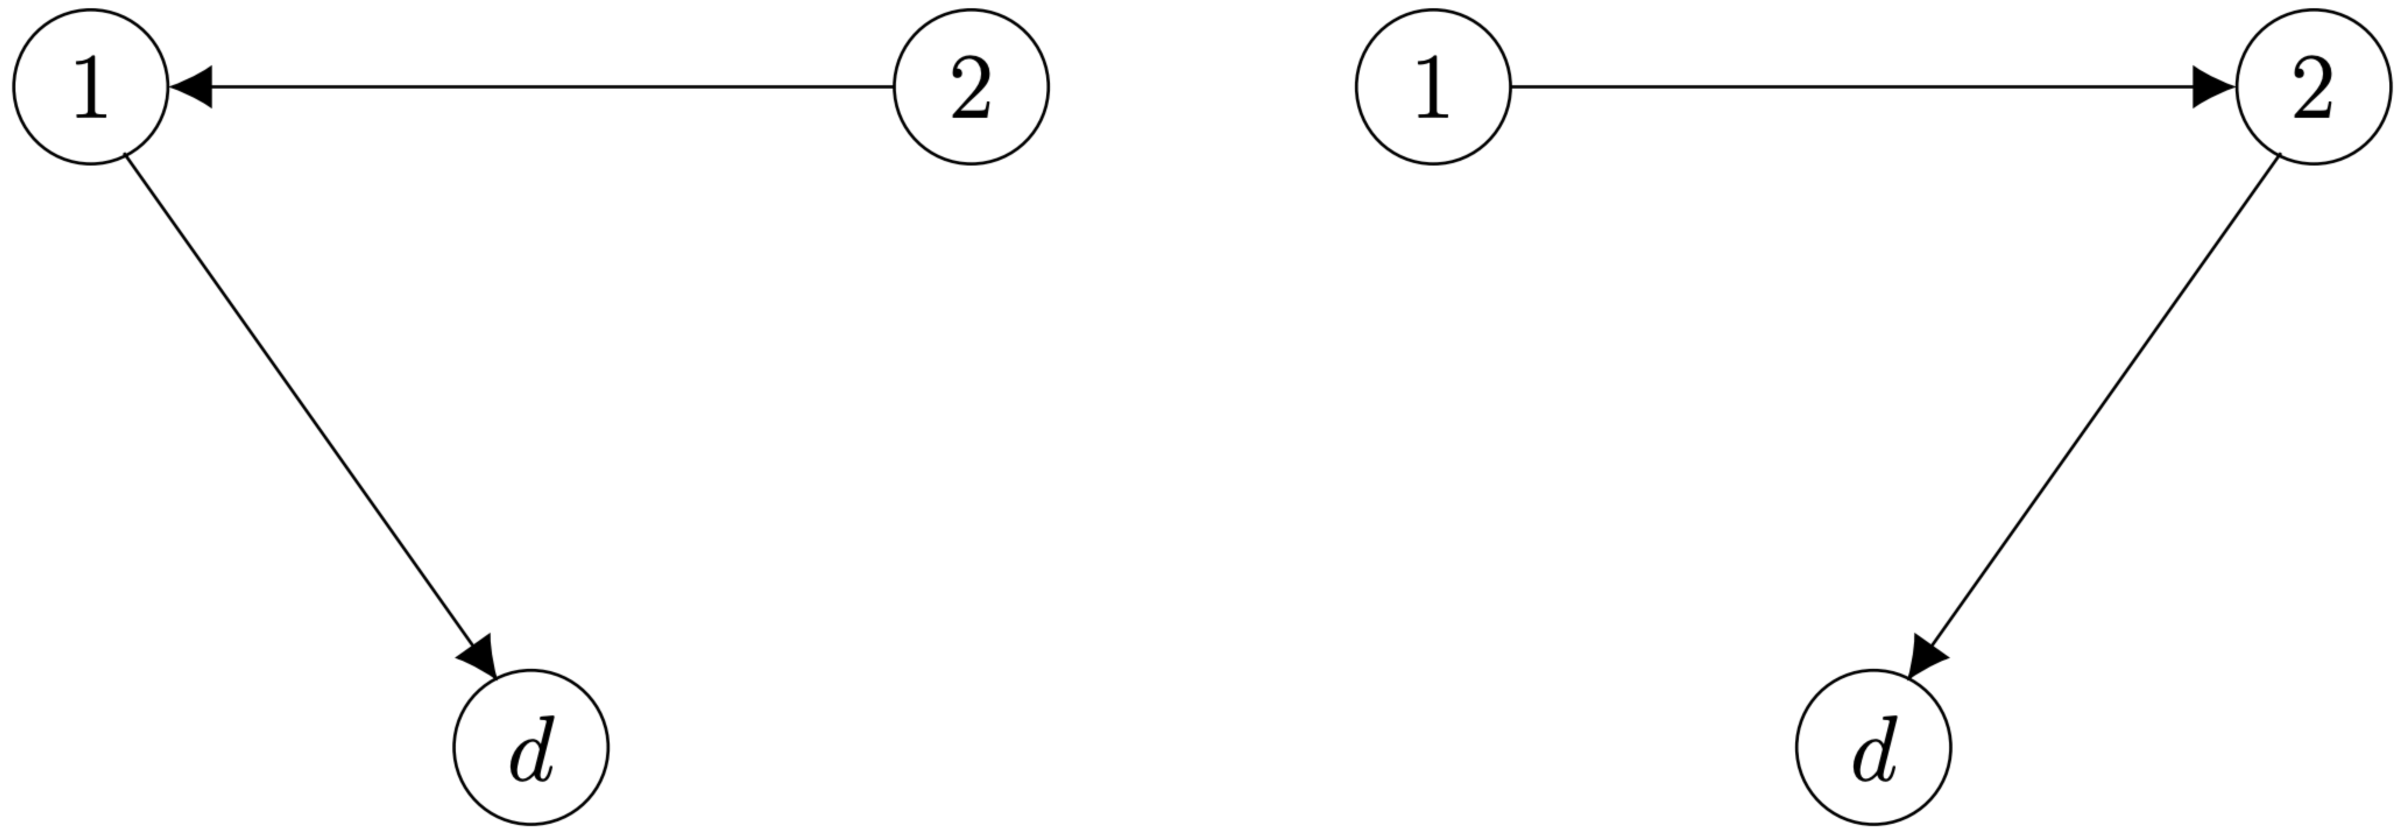

(a) (b)

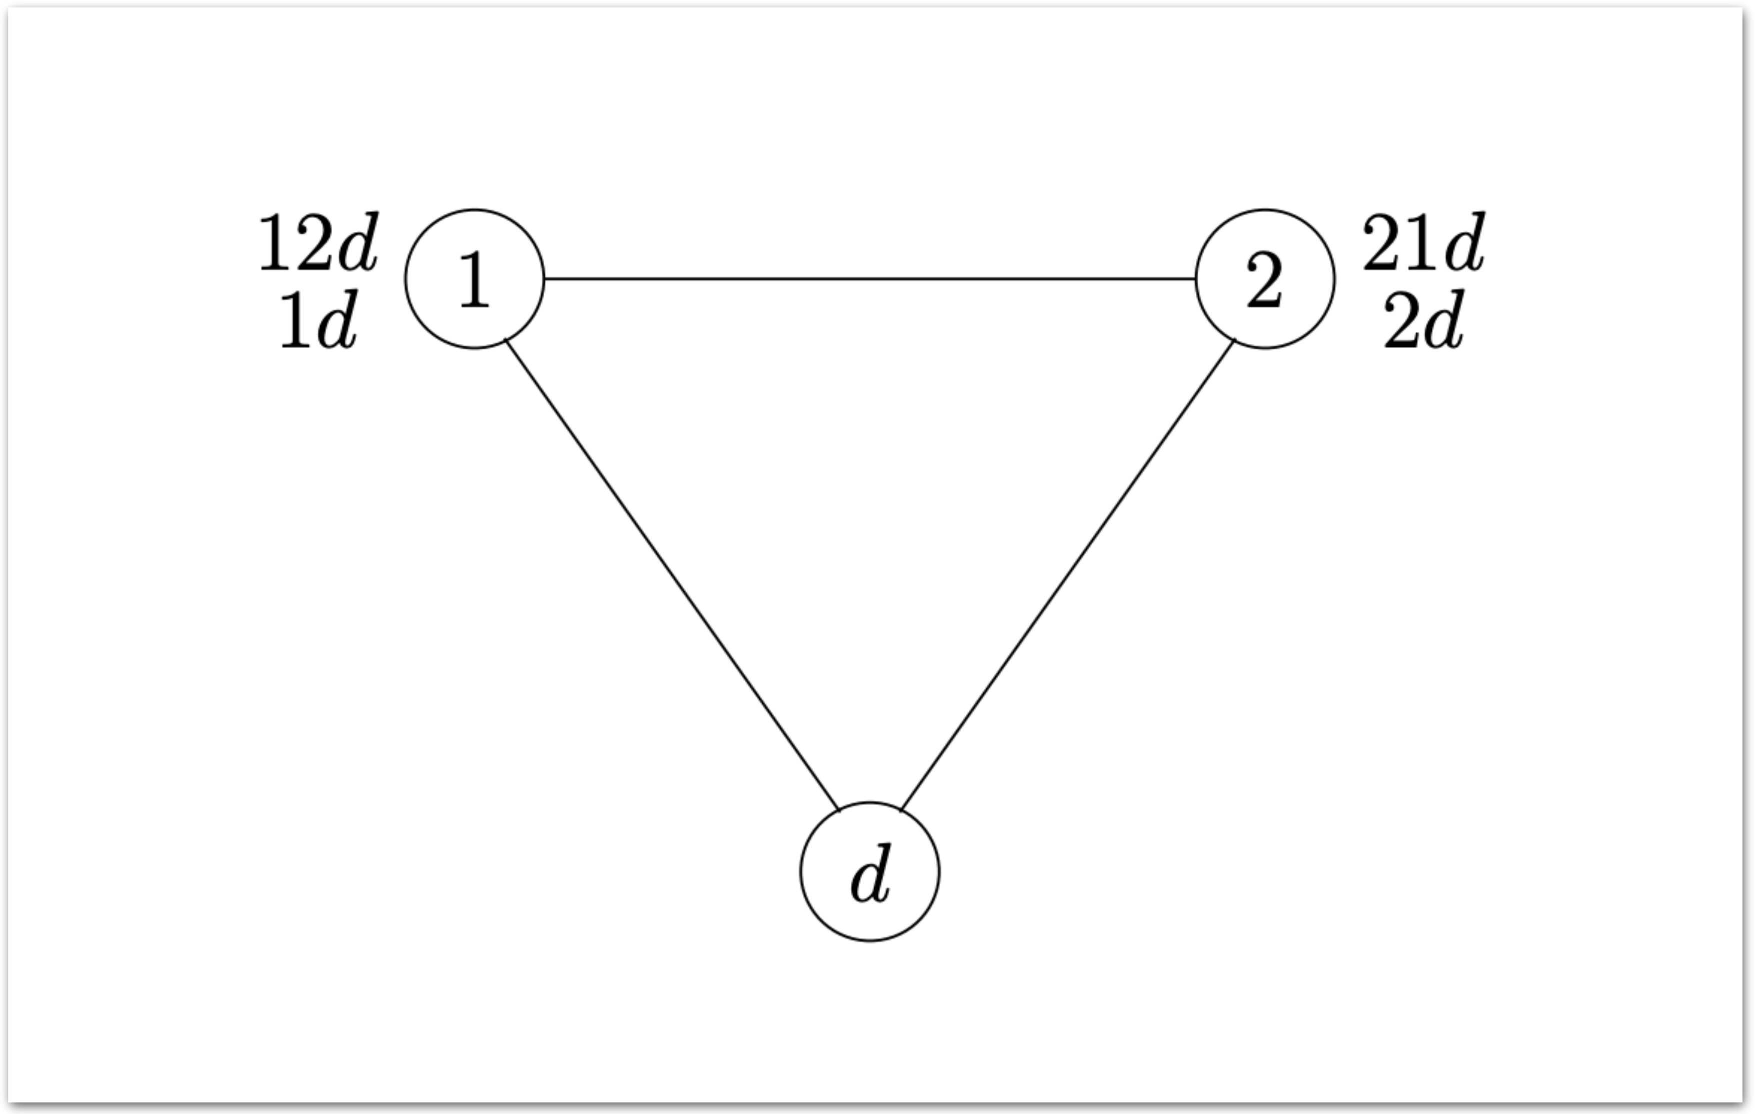

Credit: Roughgarden

Stable Routings: Equilibrium

- Note: a stable routing must a tree, directed into the destination d
- BGP maintains the invariant that out-degree of every AS is at most 1
- BGP also explicitly prevents cycles from forming
- We saw that there can be multiple stable routings
- **Question**. Does a stable routing always exist?

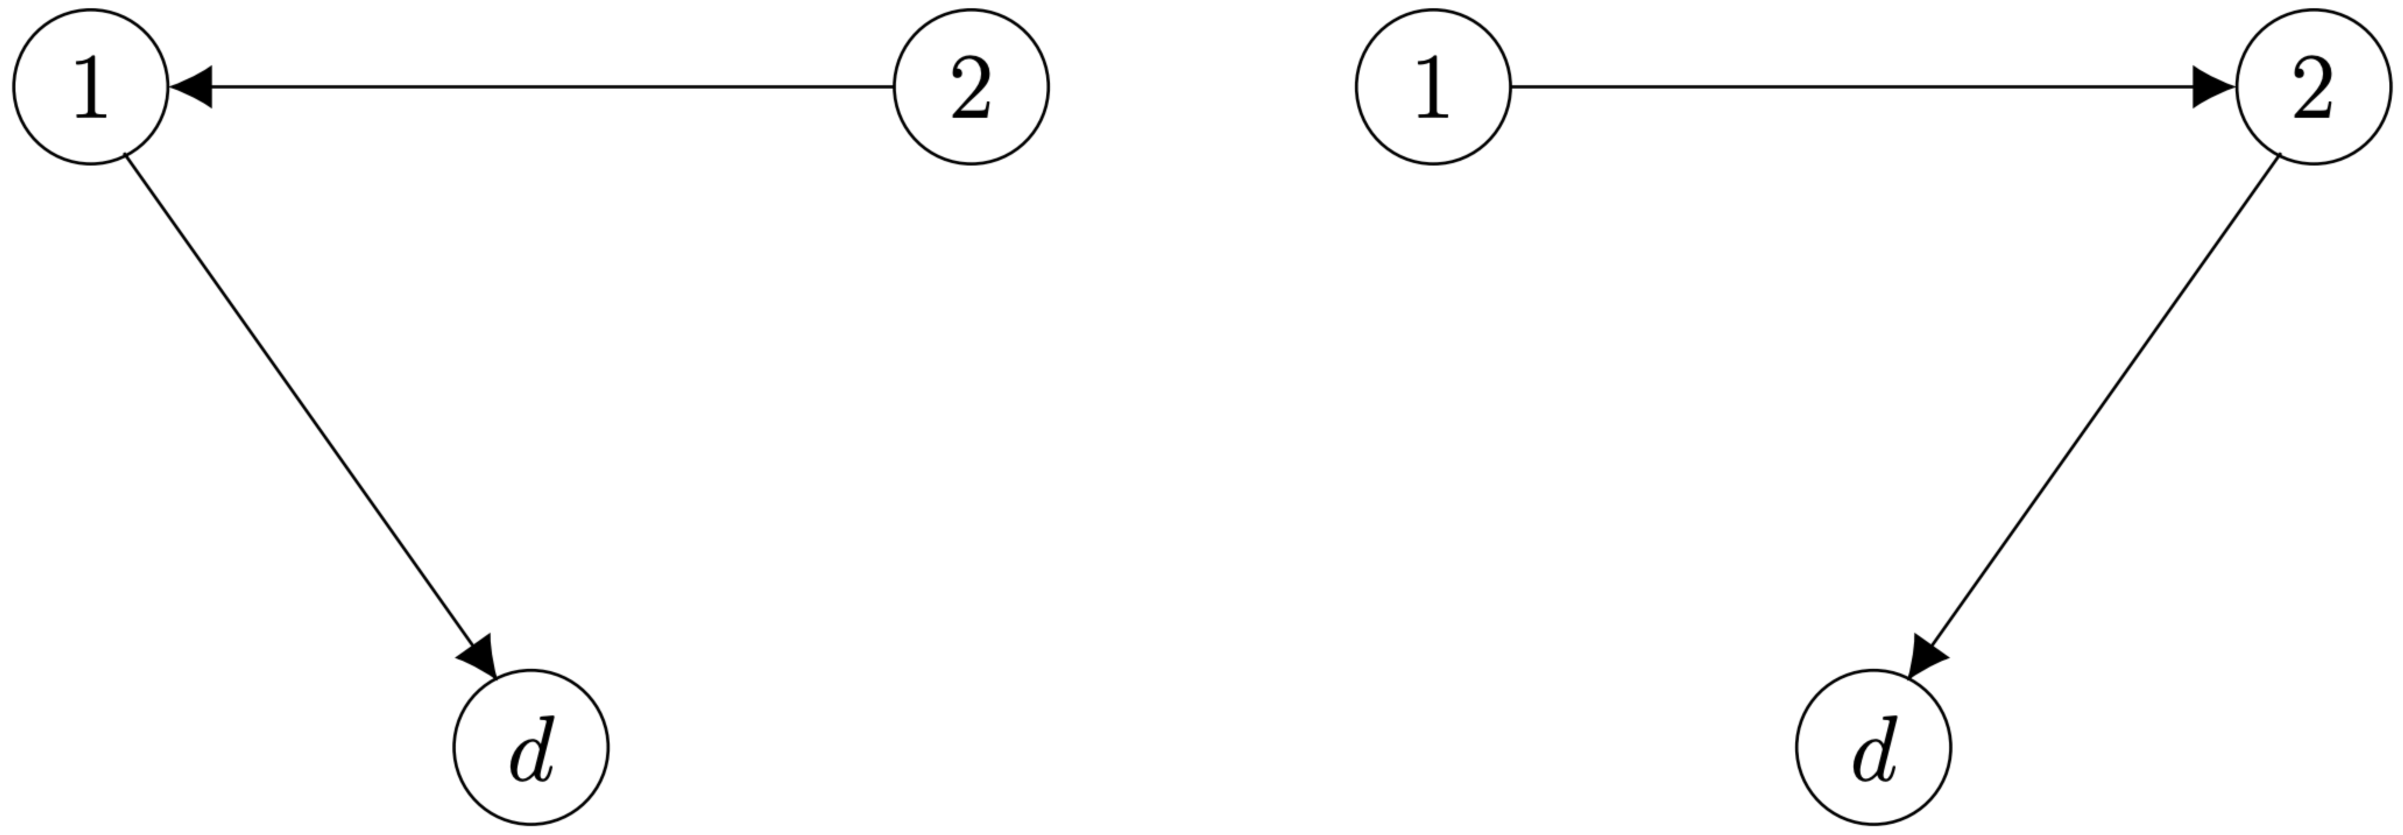

(a) (b)

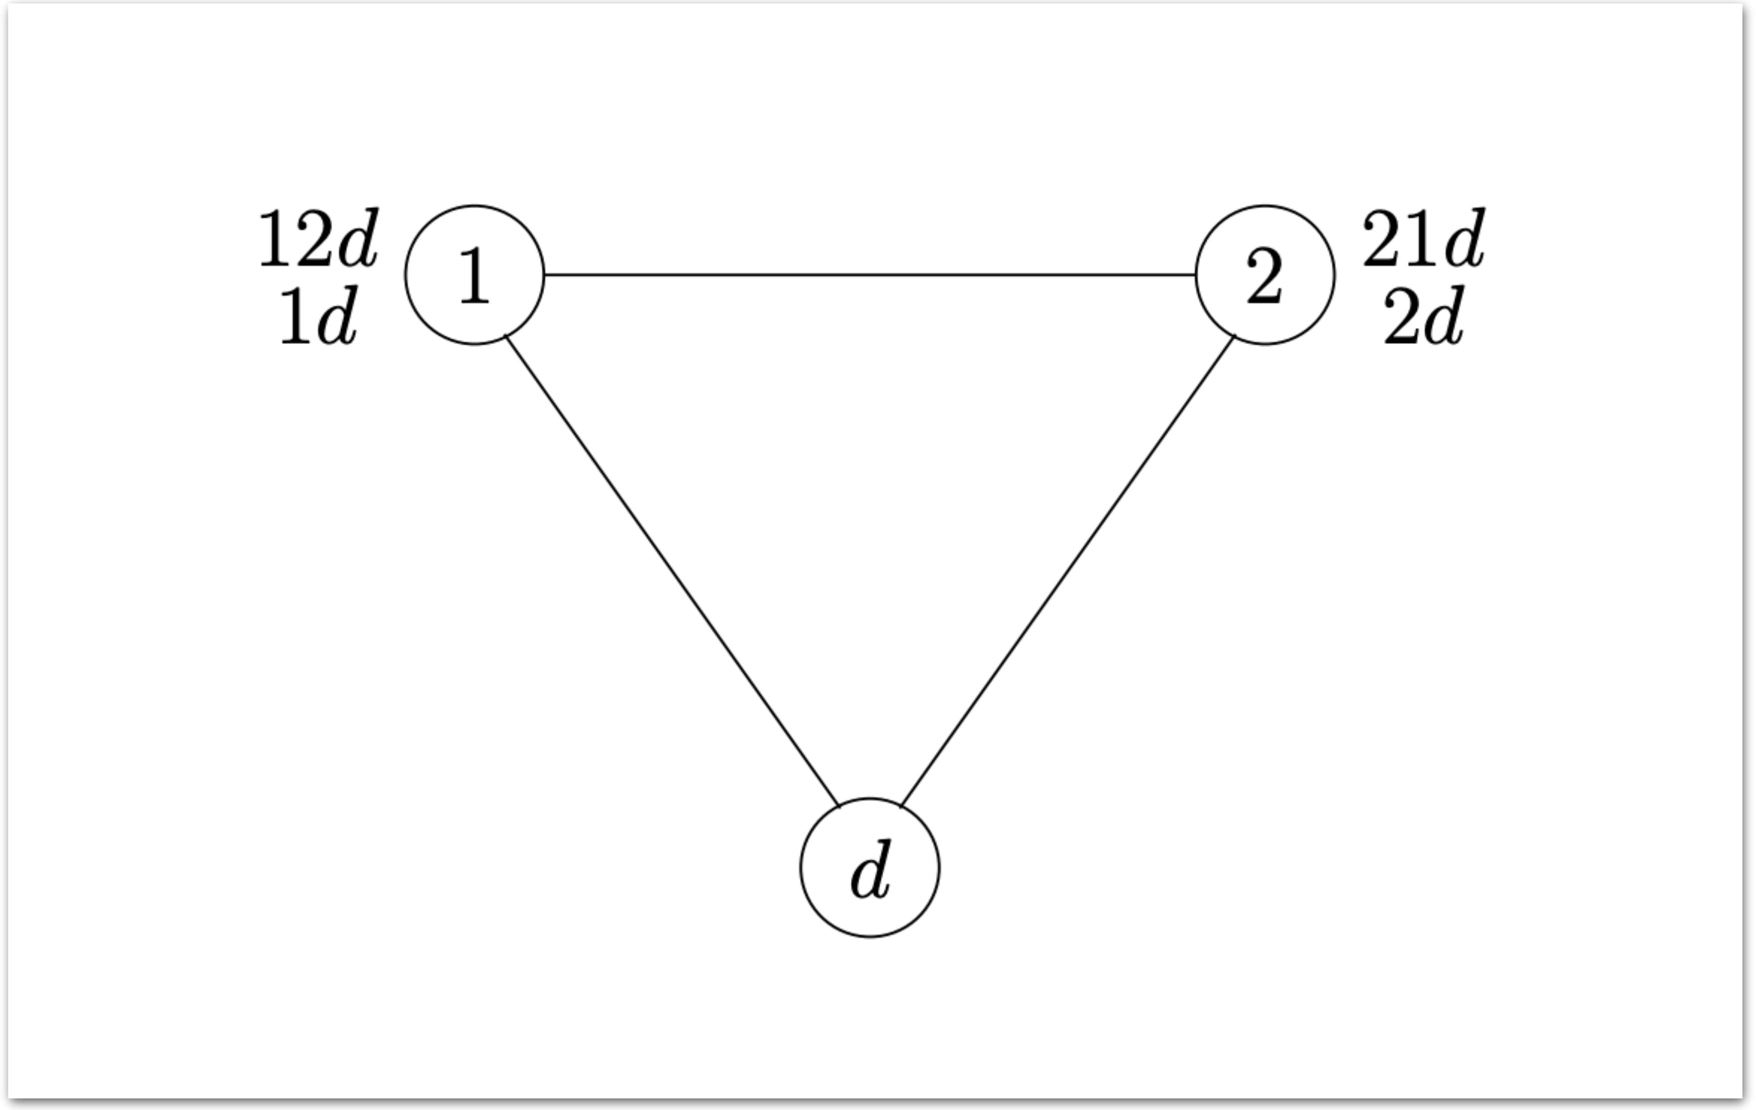

Credit: Roughgarden

Stable Routings: May Not Exist

- Consider the following stable routing network
- Every AS prefers a direct path to d over the empty path
  - Thus, one of the ASes must have a direct path to d

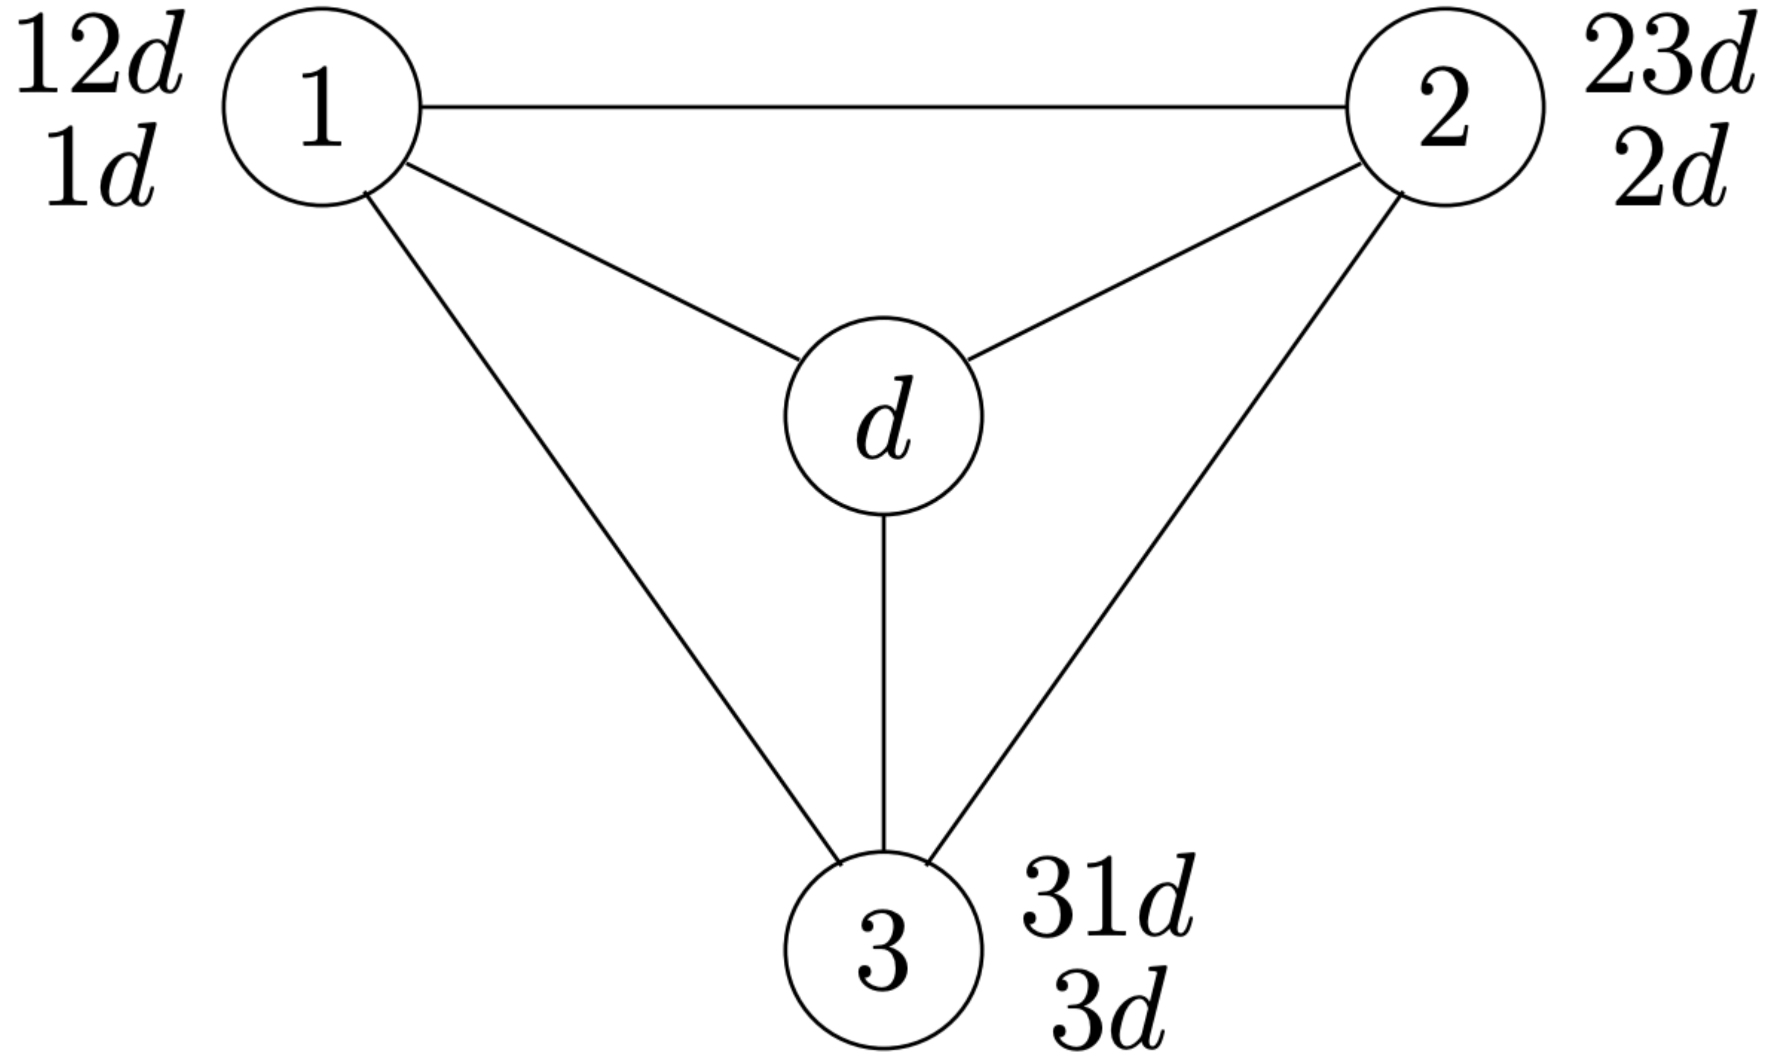

Credit: Roughgarden

Dispute Wheels

- We want to discuss incentives in BGP routing
    - Do ASes have an incentive to follow the intended behavior
- Hard to do if the output of BGP is not well defined in some cases
- Under what cases does the output of BGP make sense?
- **Theorem** [Griffin et al.] If the AS graph has **no dispute wheel**, then:
    - There exists a stable routing
    - This stable routing is unique
    - BGP is guaranteed to converge to the stable routing
- What is a dispute wheel?

232 IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 10, NO. 2, APRIL 2002

The Stable Paths Problem and Interdomain Routing

Timothy G. Griffin, F. Bruce Shepherd, and Gordon Wilfong

Dispute Wheels

- A dispute wheel consists of (for some $k \geq 2$):
  - distinct ASes $u_1, u_2, \ldots, u_k$
  - cycle-free paths $P_1, P_2, \ldots, P_k$ where P_i is a path from u_i to d
  - cycle-free paths $Q_1, Q_2, \ldots, Q_k$, where Q_i is a path from u_i to u_{i+1}

  such that for each $i = 1, \ldots, k$, u_i prefers the indirect path $Q_i \cdot P_{i+1}$ to the direct path P_i (here $u_{k+1} = u_1, P_{k+1} = P_1$)
- Note that P_i's and $Q_i's$ need not be internally disjoint

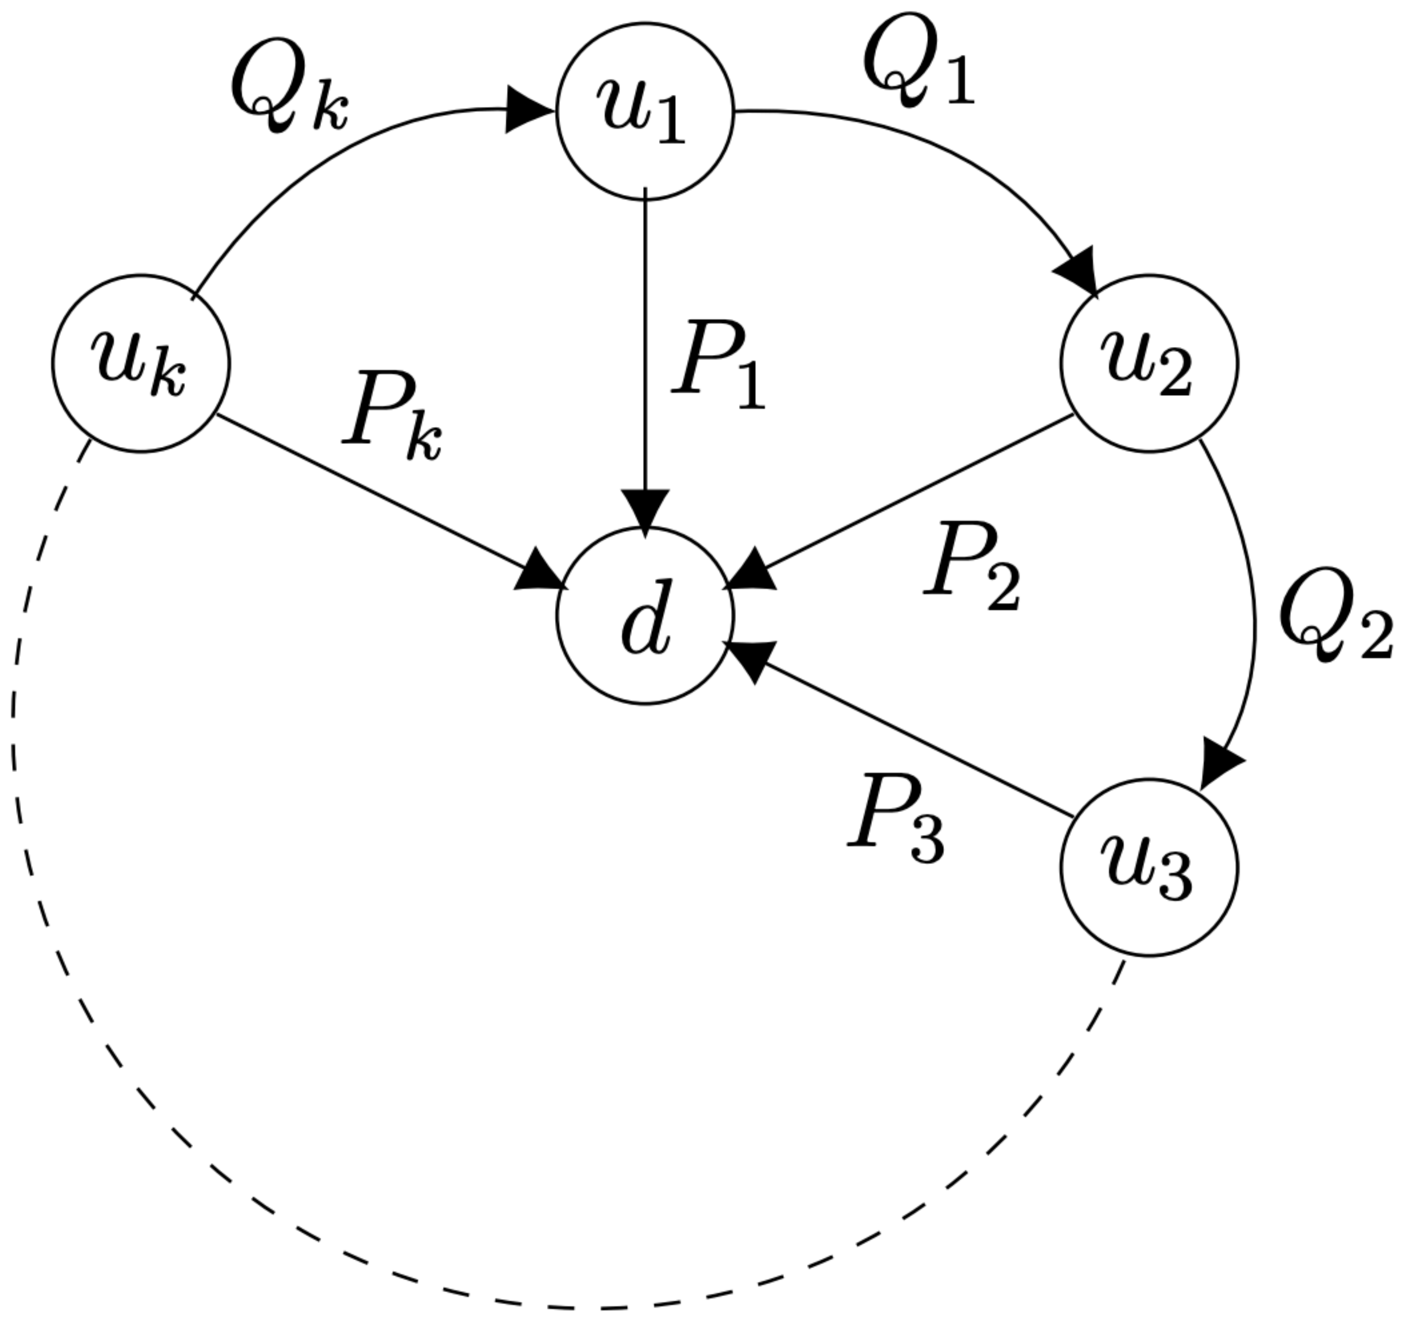

Credit: Roughgarden

Dispute Wheels Example

- $P_1 = 1d$, $P_2 = 2d$
- $Q_1 = 12$, $Q_2 = 21$
- Both ASes prefer $12d$ or $21d$ rather than direct

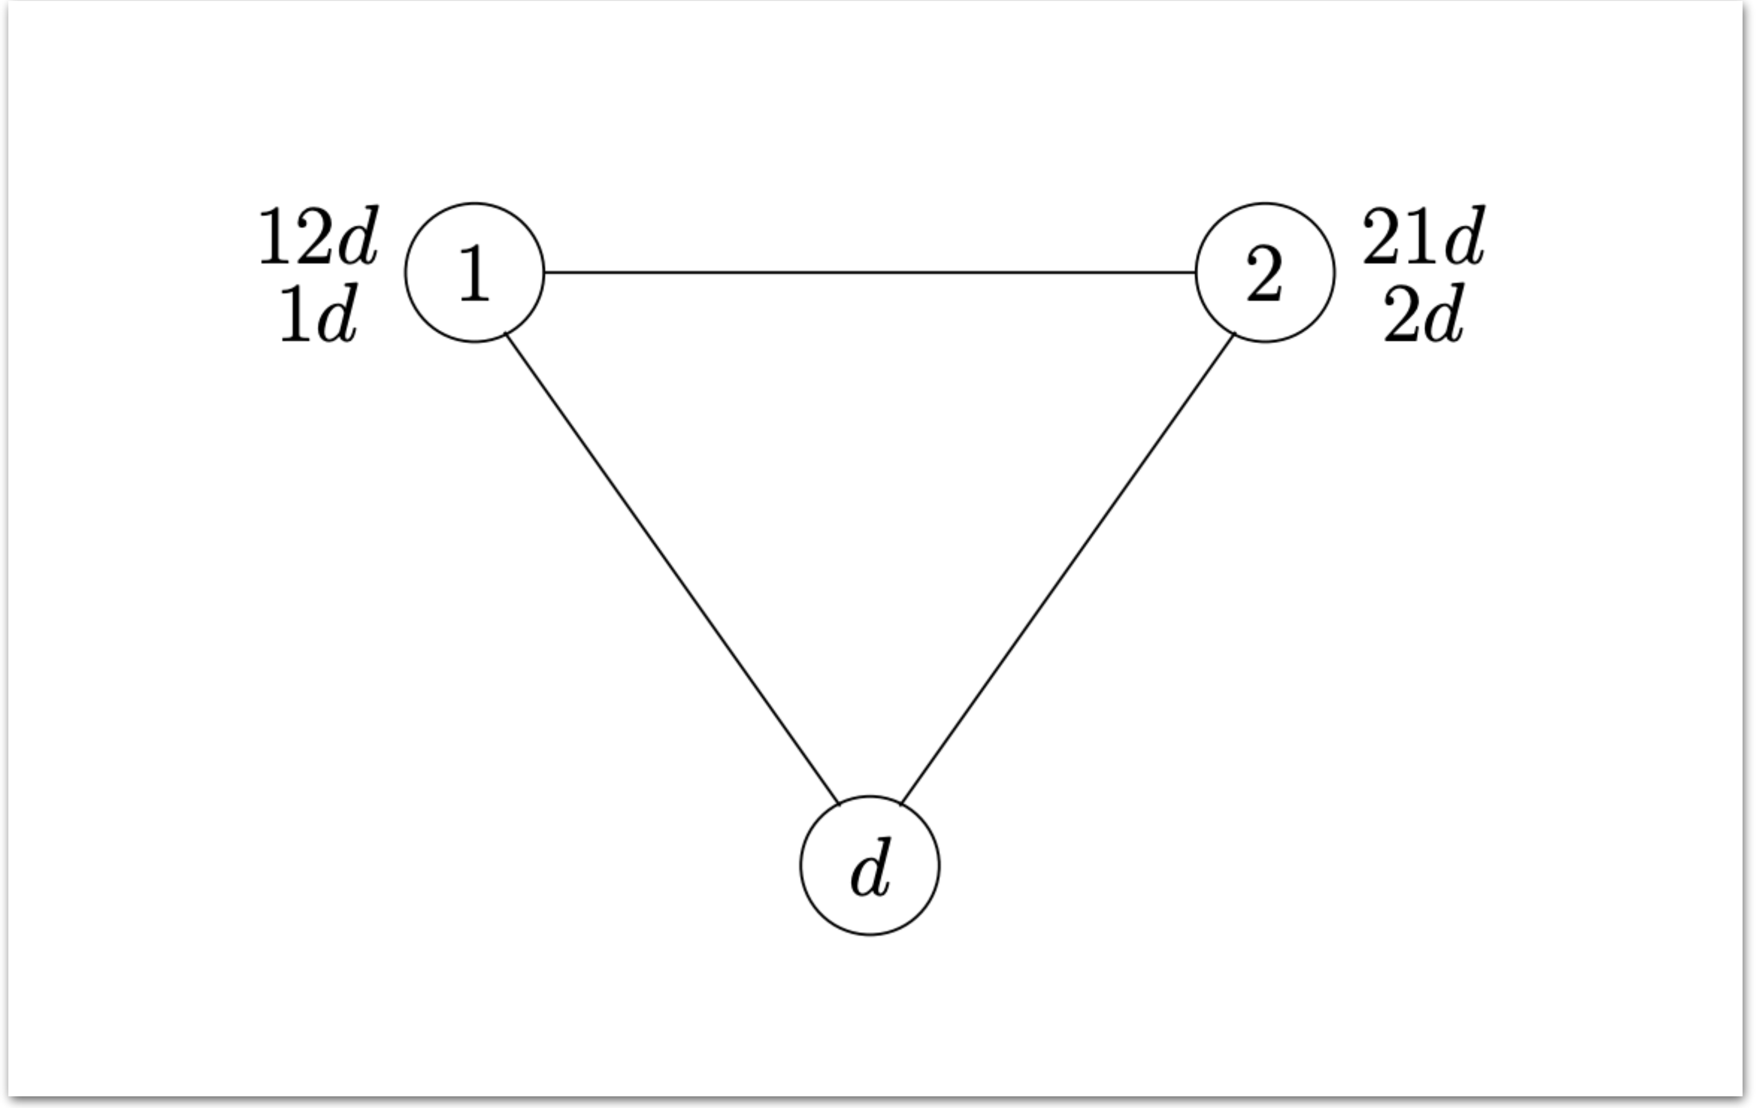

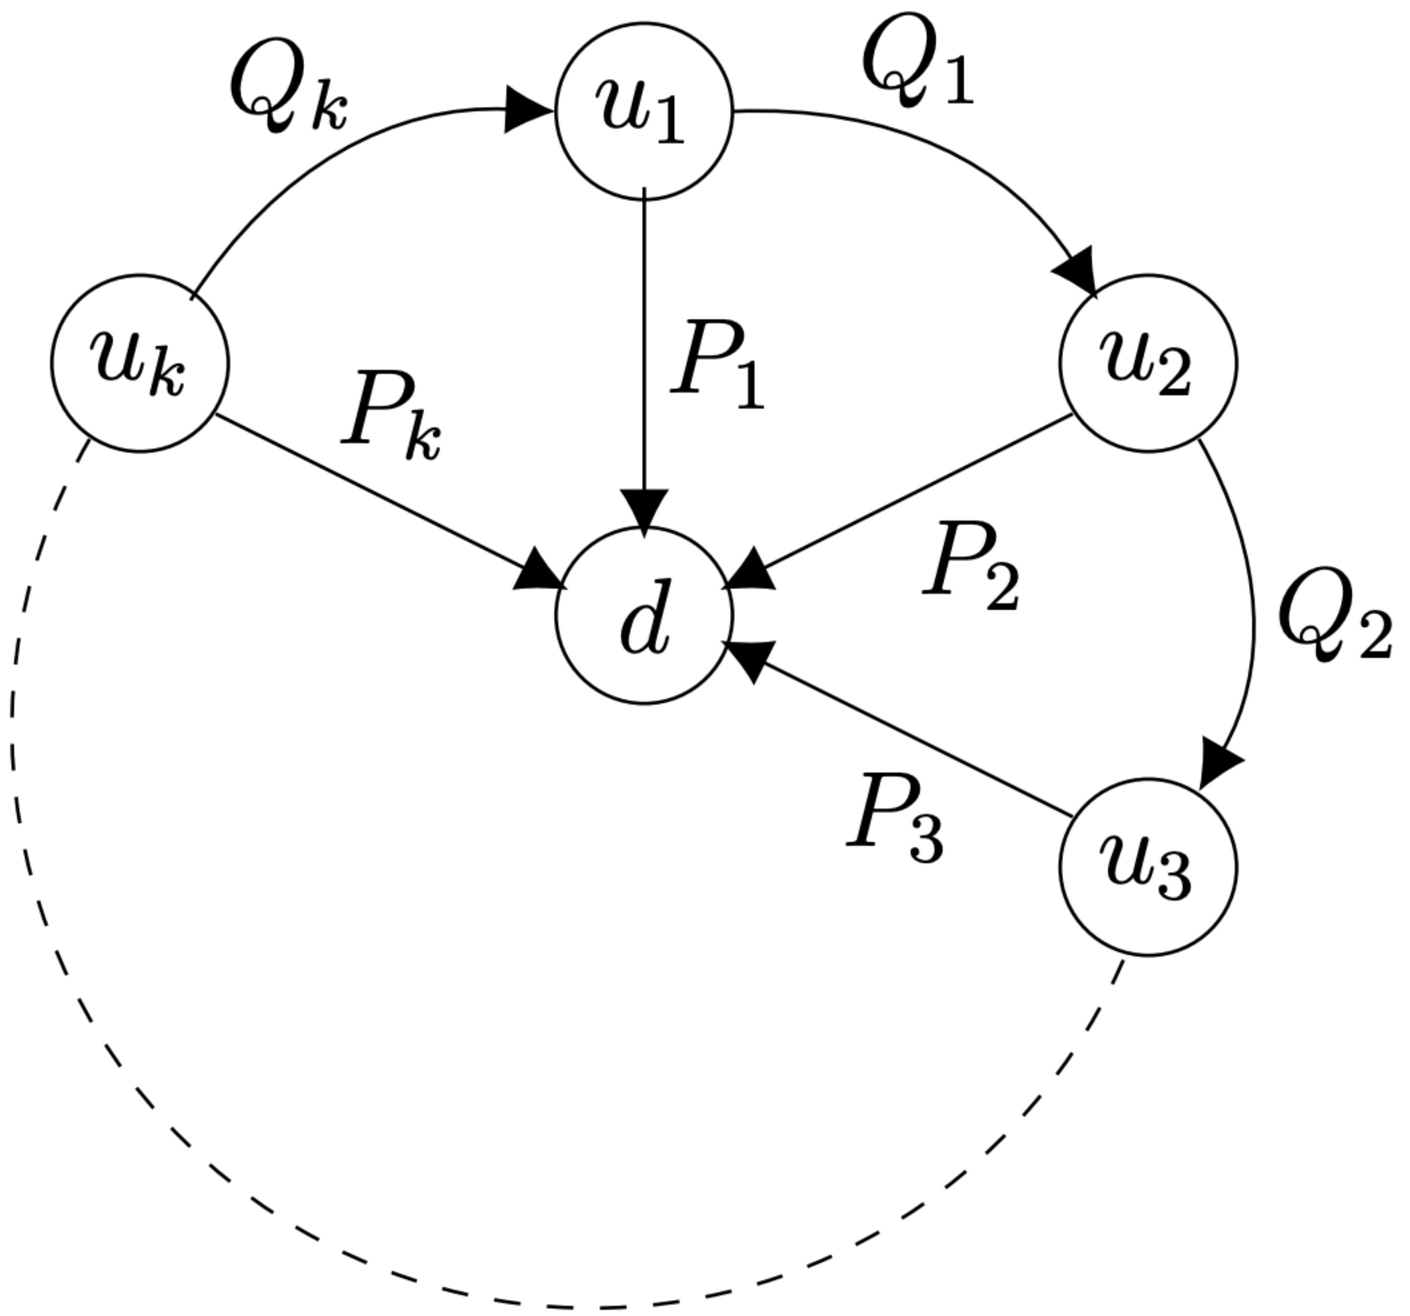

Credit: Roughgarden

Dispute Wheels Example

- $P_i = id$ for $i = 1,2,3$
- $Q_1 = 12,\ Q_2 = 23,\ Q_3 = 31$
- Again each AS prefers indirect over direct

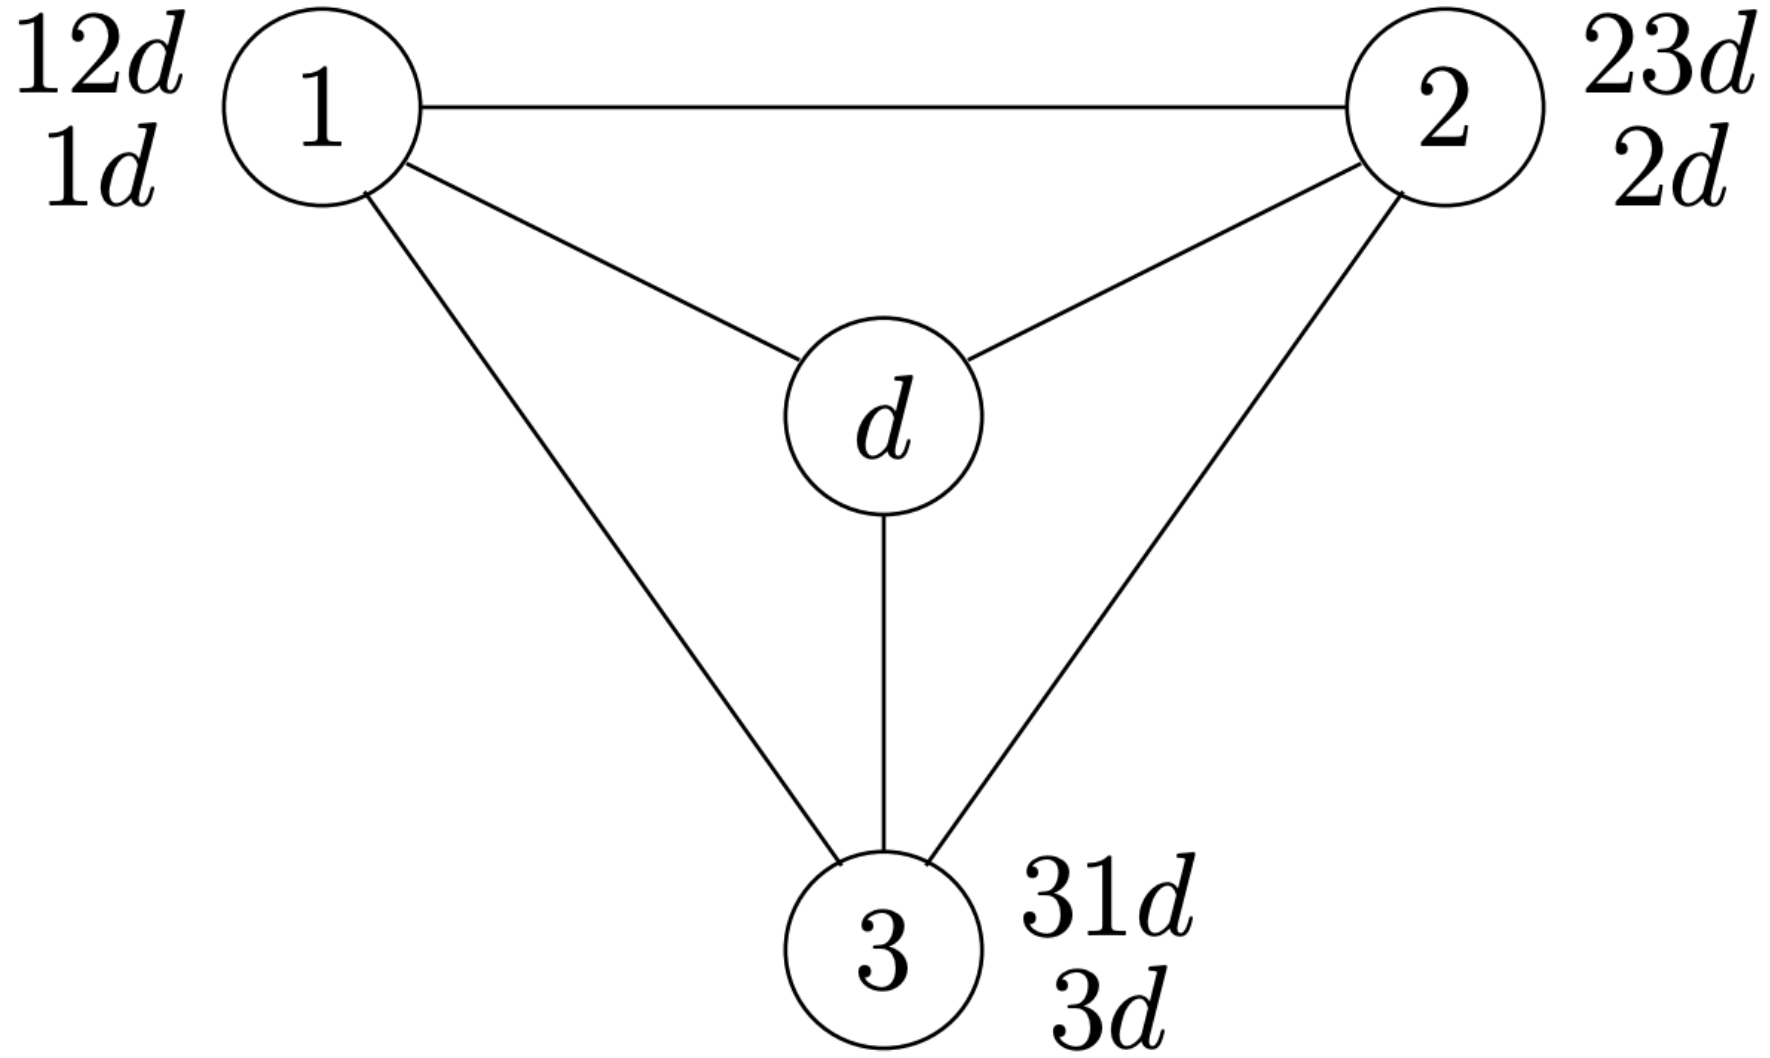

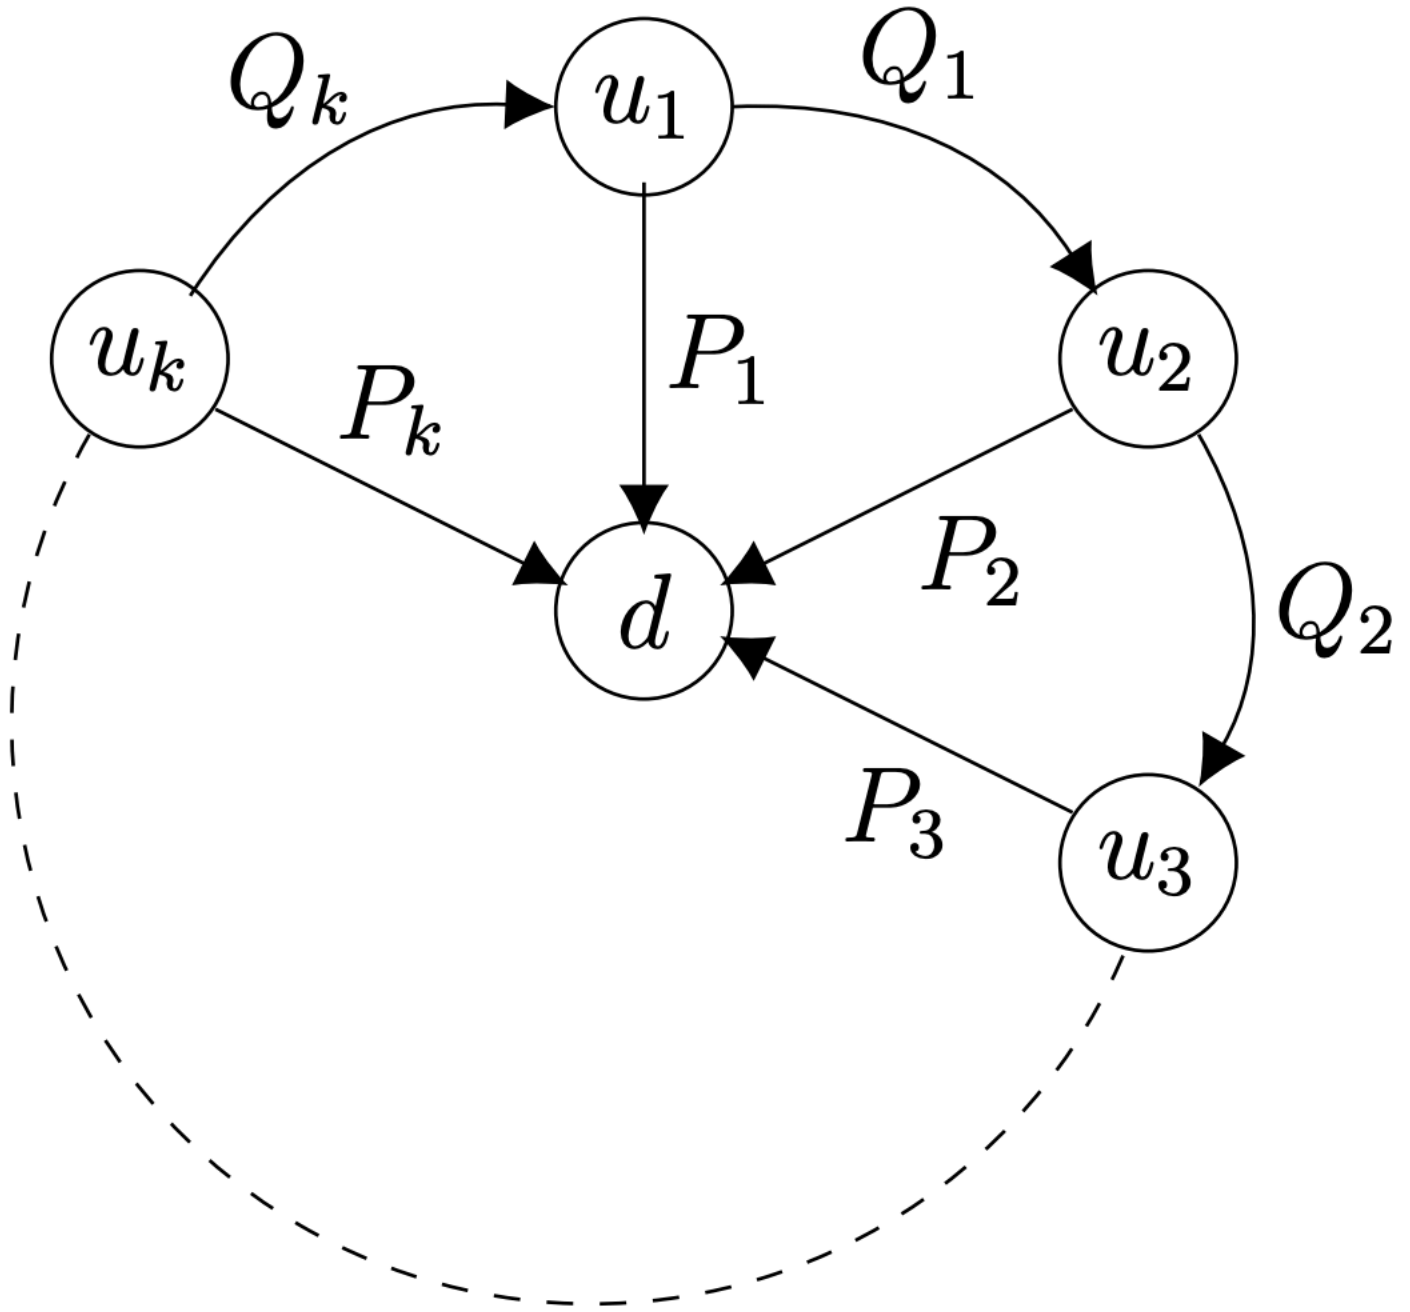

Credit: Roughgarden

Dispute Wheels in Practice

- BGP converges if there are no dispute wheels
- But how strong is the no-dispute-wheel condition?
- **Gao and Rexford** gave justifications about why no dispute wheel condition should generally hold for realistic AS preferences
- **Gao-Rexford conditions** (rule out dispute wheels): every AS prefers
    - to route through a customer over those through a peer
    - to route through a peer over a provider
- Empirical evidence the conditions are approximately true for most ASes

IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 9, NO. 6, DECEMBER 2001 681

Stable Internet Routing Without Global Coordination

Lixin Gao, *Member, IEEE,* and Jennifer Rexford, *Senior Member, IEEE*

Incentive Issues

- We restrict ourselves to AS graphs where BGP converges
- Do ASes have an incentive to follow the protocol?
  - Does any AS have a beneficial unilateral deviation?
- Types of deviations:
  - Choose your path P_u to be something other than the favorite path among the available options
  - **Withhold information** about your path to (some) neighbors
  - **Announce** a path to (some) neighbors that is different from the one you are actually using, possible even a non-existent **fake path**

Fake Path Announcements

- Can and do ASes announce non-existent paths?
- Happens in BGP **all the time**
- In 2008 Pakistan telecom blocked access from Pakistan to Youtube

Pakistan Cuts Access to YouTube Worldwide

Activists in Pakistan on Sunday, protesting the publication of drawings depicting Muhammad. Pakistan blocked YouTube for containing material considered offensive to Muslims.
Rizwan Tabassum/Agence Free-Press — Getty Images

Credit: Roughgarden

Fake Path Announcements

- In 2008 Pakistan telecom blocked access from Pakistan to Youtube
- Pakistan Telecom rerouted Youtube traffic to a local server
- Pakiston Telecom announced this new route to Youtube to some of its neighbors
    - An ISP in HK switched its route to Youtube soon, followed by most of the other ISPs
- **Outcome**: a lot of users could not reach Youtube, Pakistan Telecom inundated with requests

Credit: Roughgarden

Withholding Path Information

- Withholding deviations from neighboring AS can be beneficial
- If AS 1 never announces any paths and AS 2 follows BGP, then the protocol converges to the stable routing preferred by AS 1
- What if there are no dispute wheels (this example does have them)

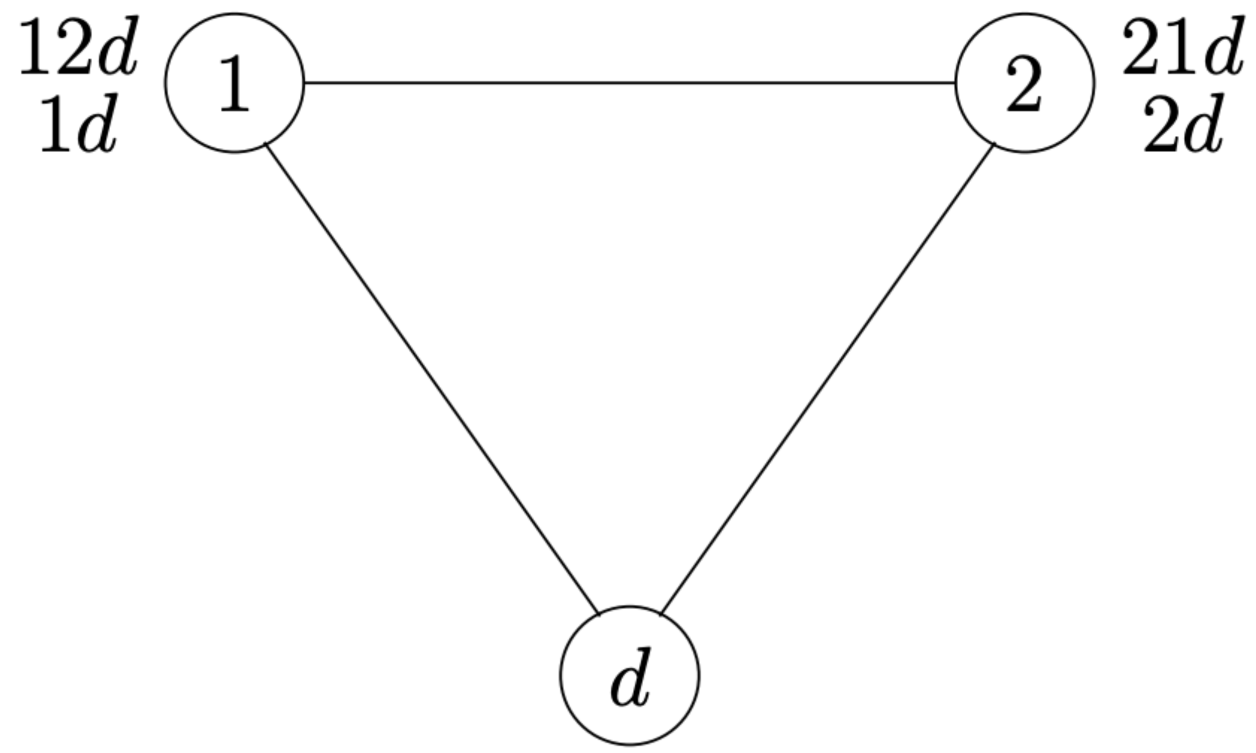

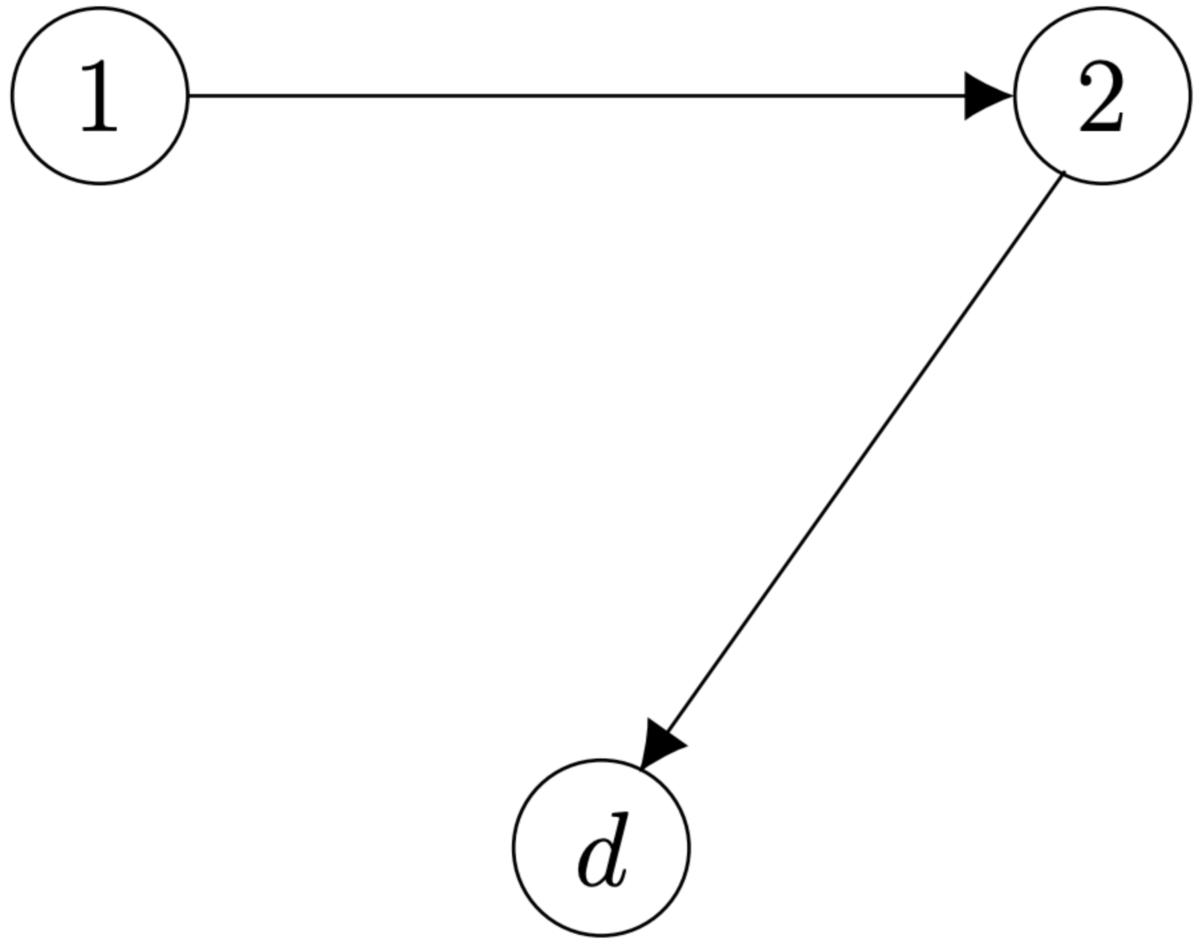

Credit: Roughgarden

Announcing Fake Paths

- If edge $3d$ was in the network and AS 3 preferred this path less than $31d$, the example would have a dispute wheel
- Suppose all ASes are honest, since there is no dispute wheel, BGP converges to a unique stable routing: shown in (a)

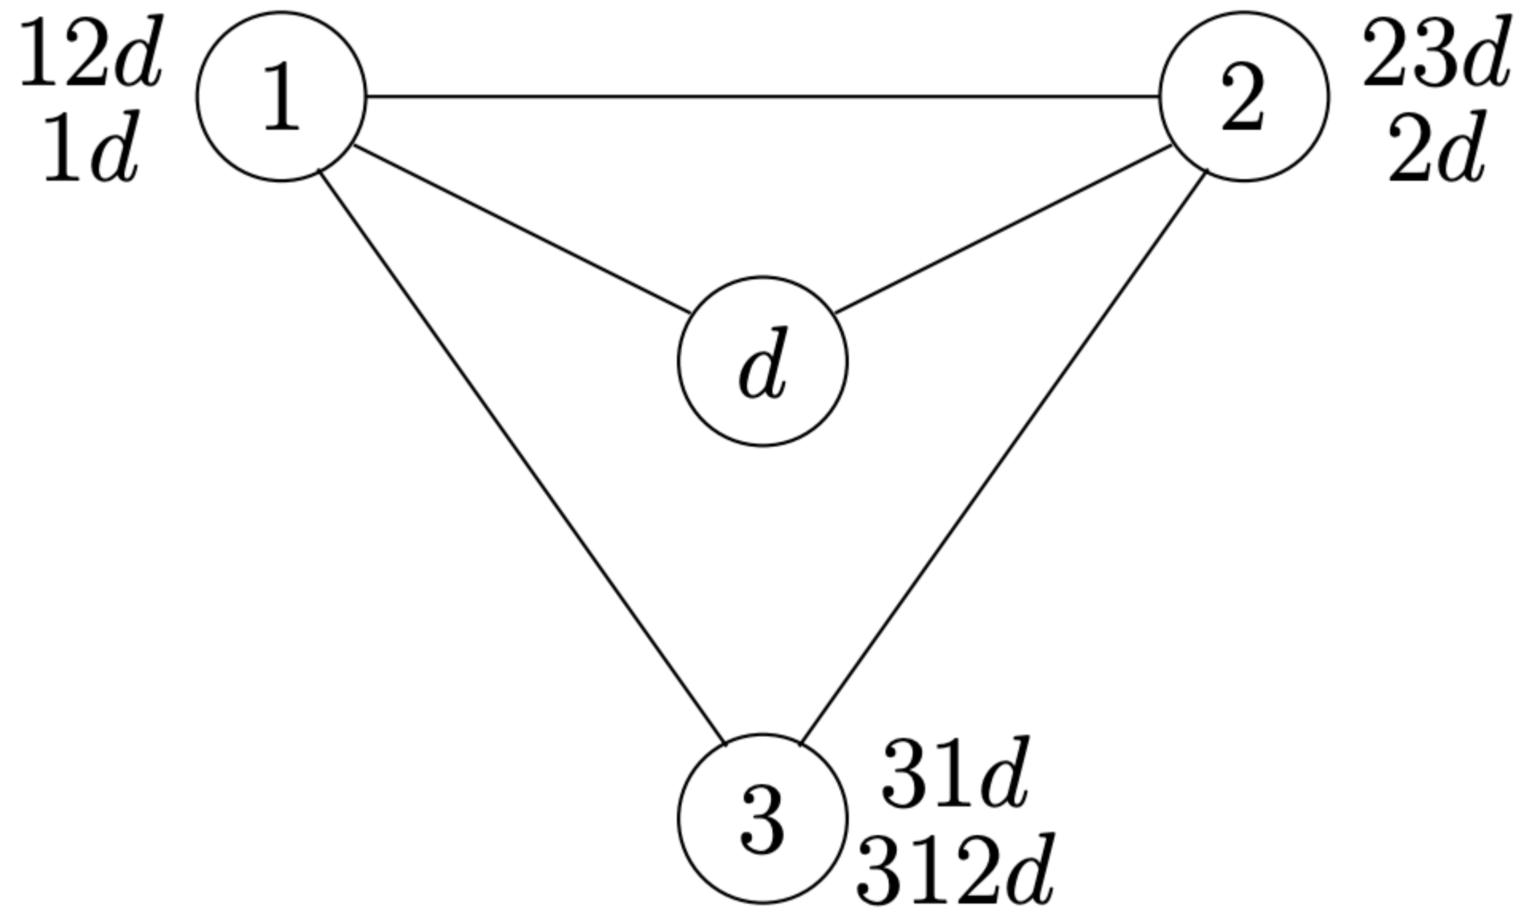

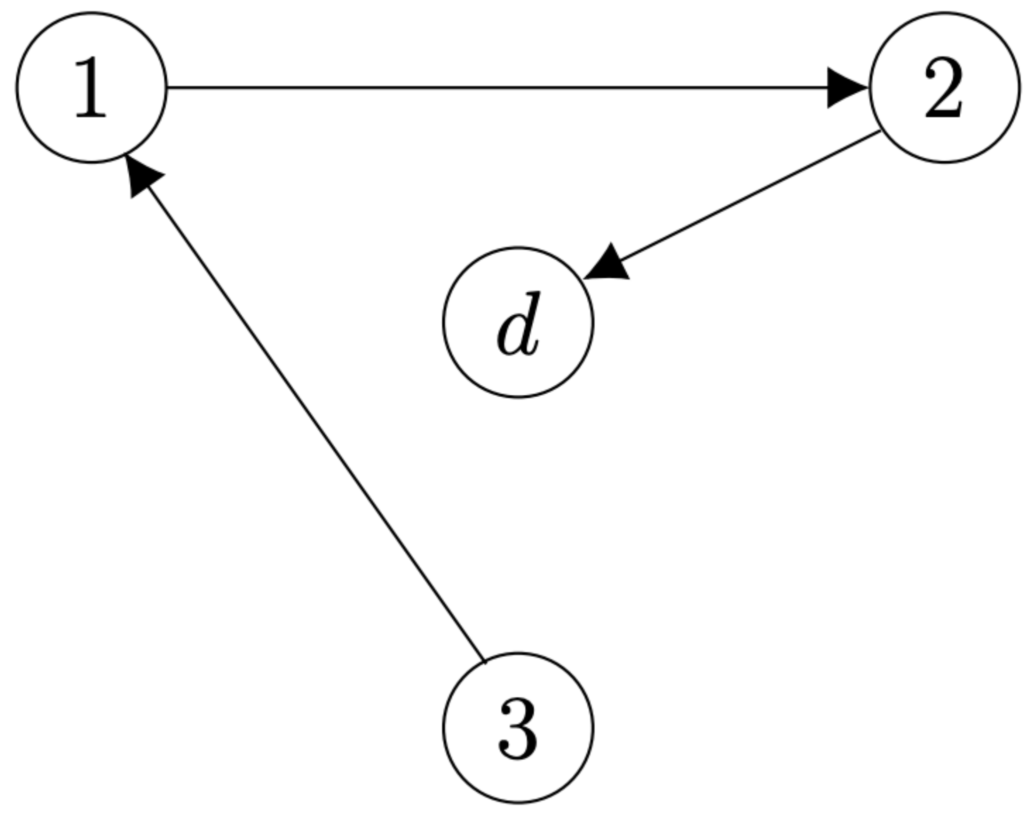

(a)

Credit: Roughgarden

Announcing Fake Paths

- No suppose AS 3 announces the fake path $3d$ to its neighbors
    - Creates a fictitious dispute wheel
- BGP converges to figure (b): AS 2 thinks its traffic is being routed through $23d$, when in fact it is being routed through $31d$

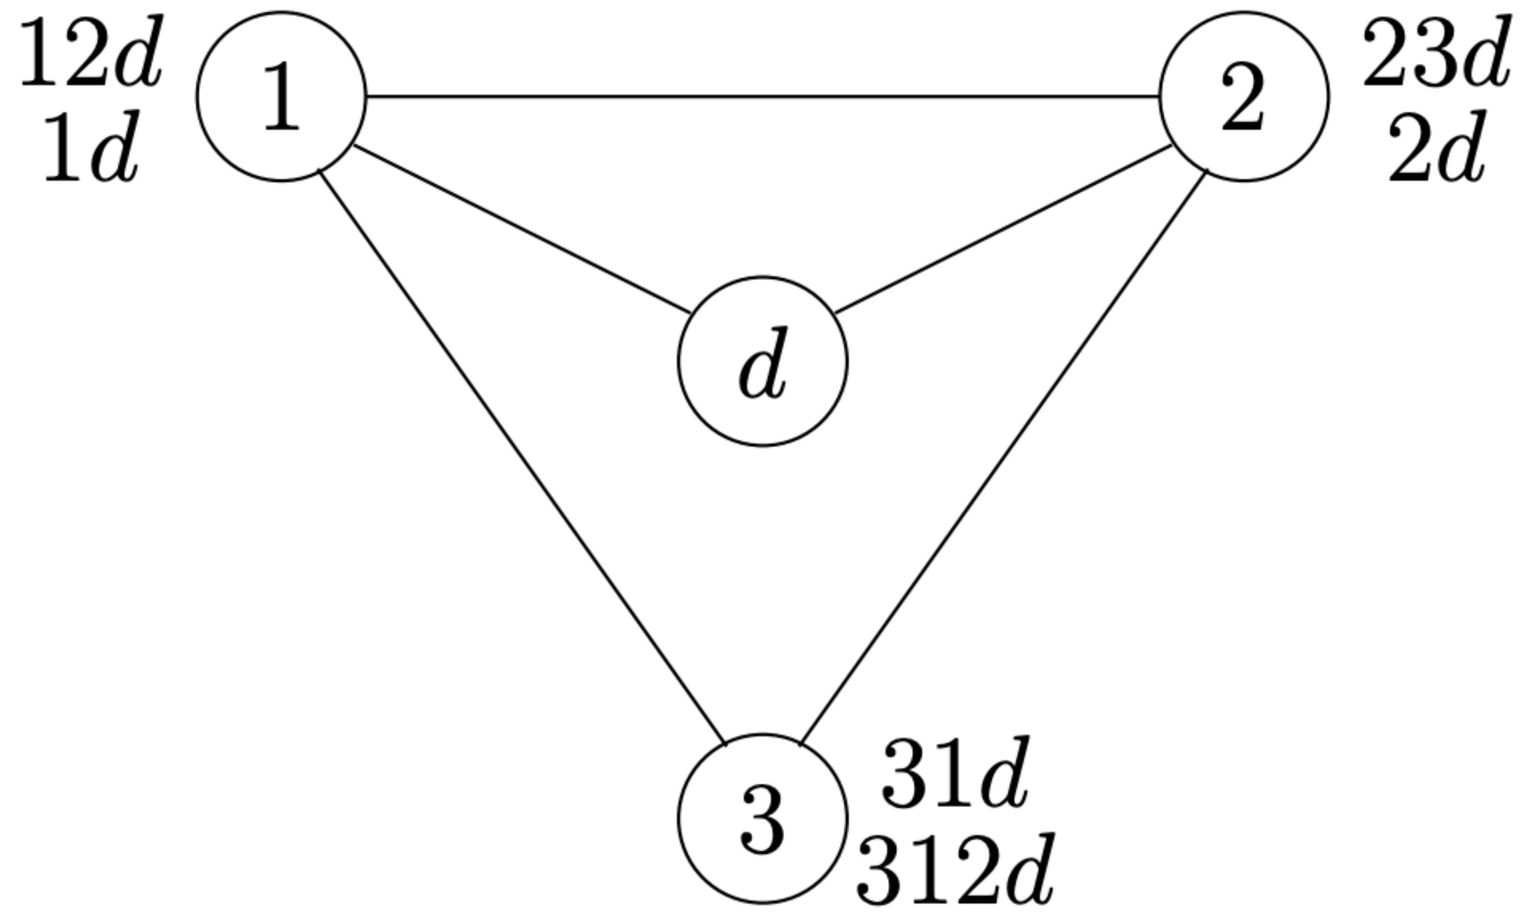

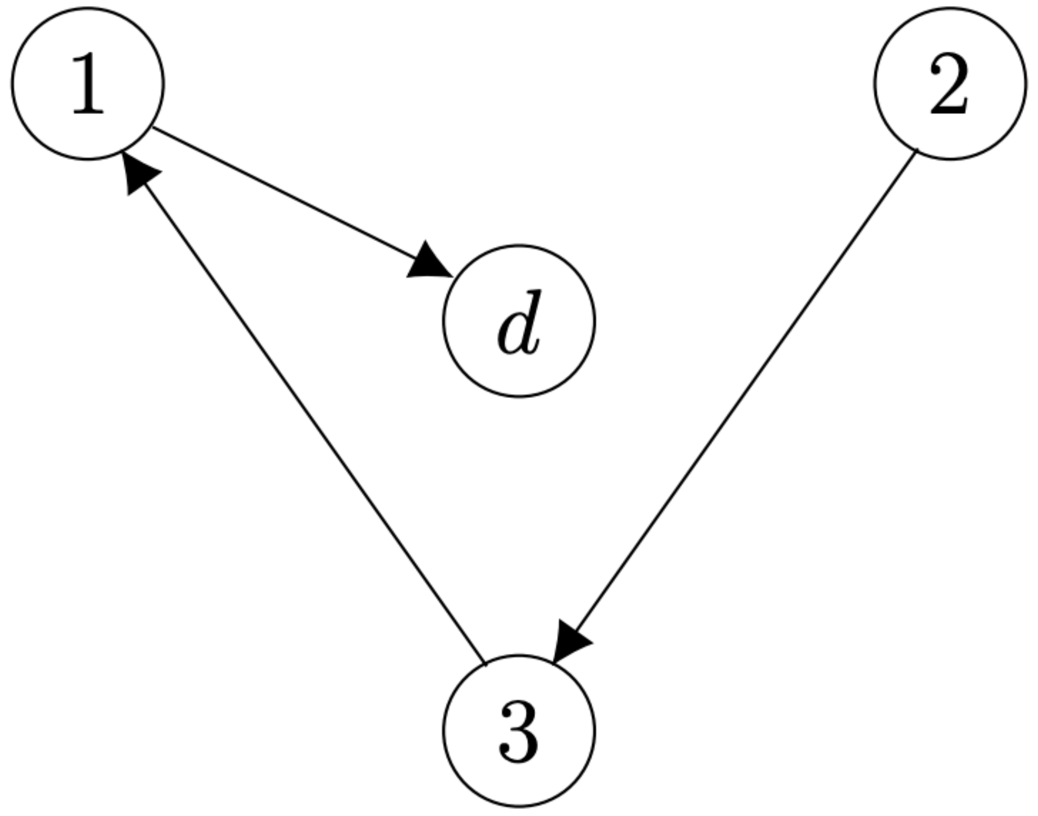

(b)

Credit: Roughgarden

Announcing Fake Paths

- If AS 2 knew the path wasn't being routed through $23d$: would not be stable but without that, it is
- AS 3 gets its favorite path when announcing a fake route, compared to being truthful

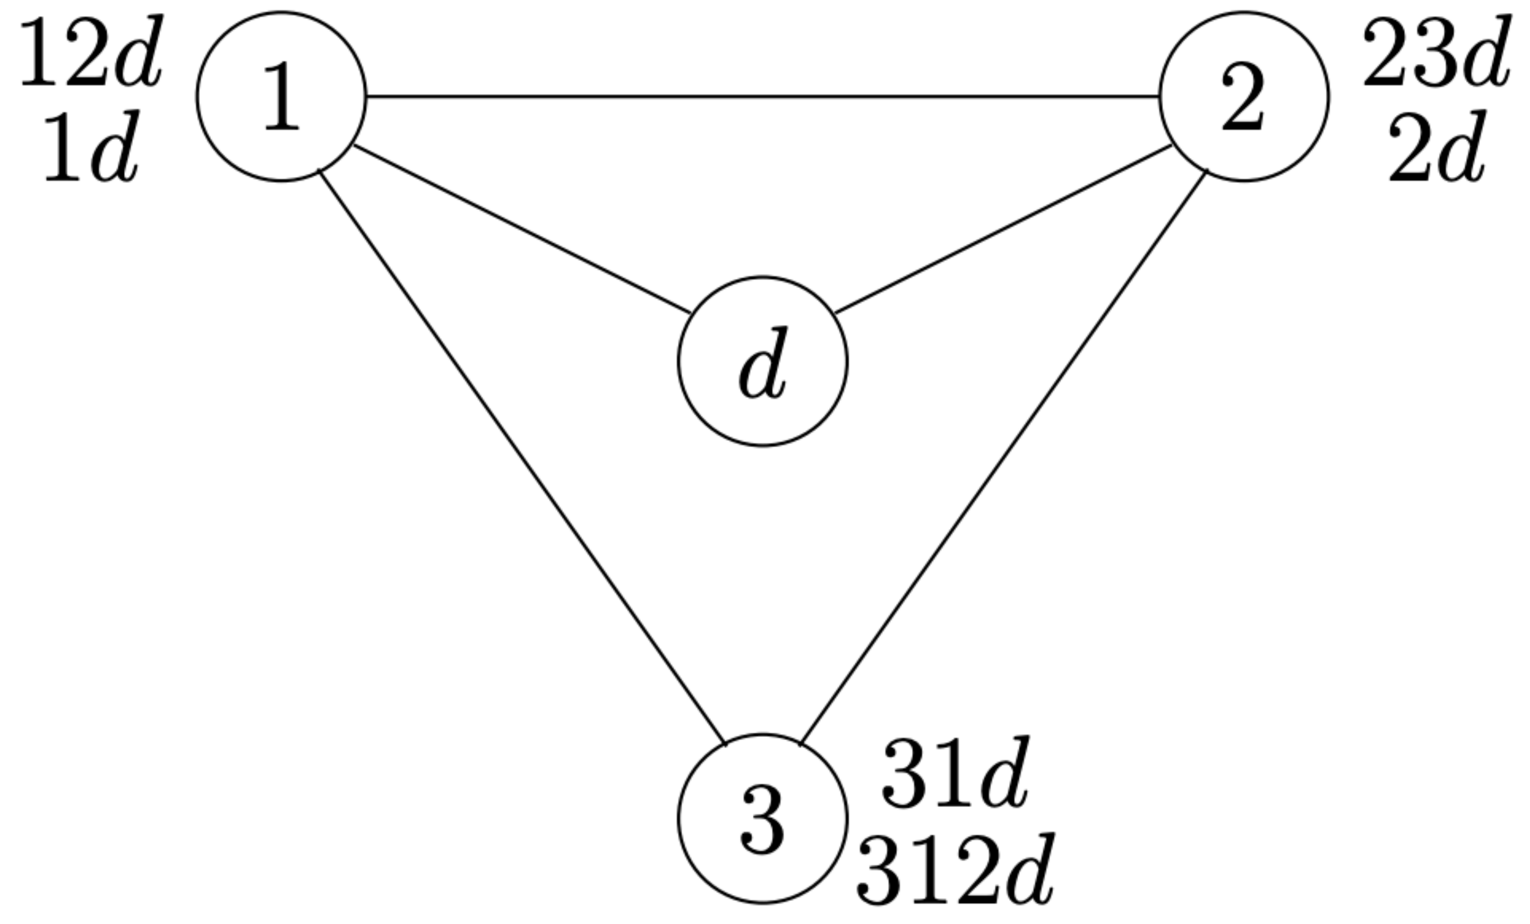

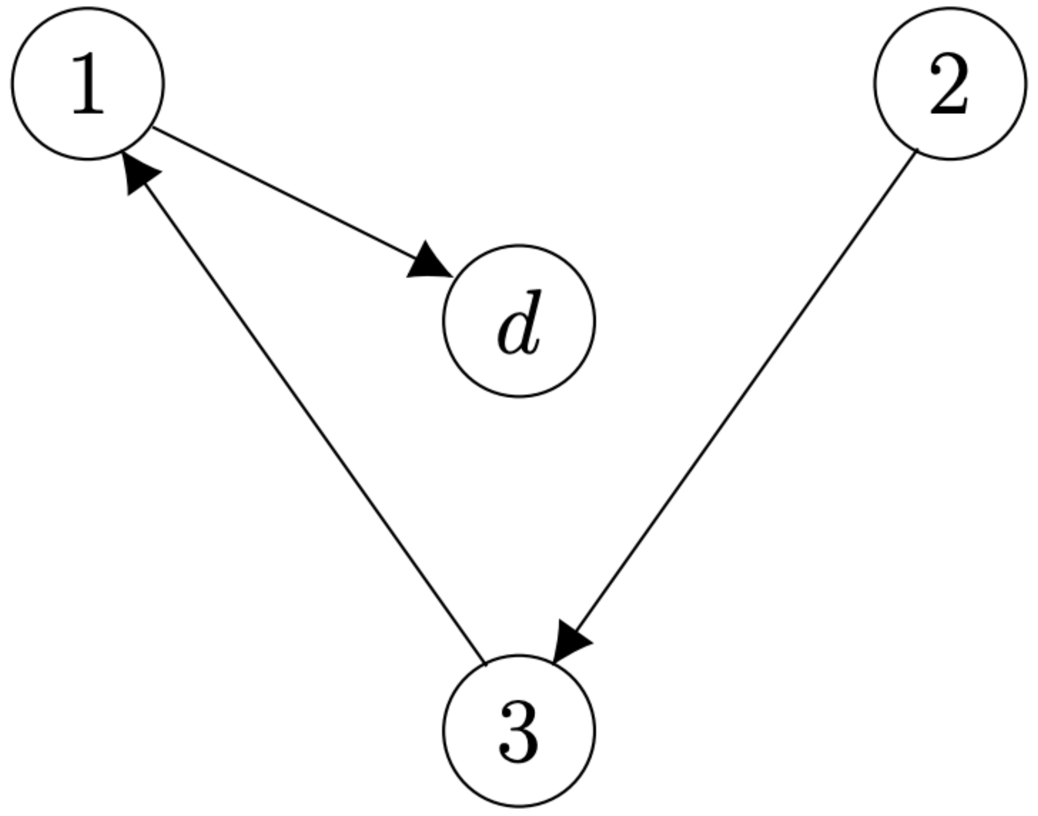

(b)

Credit: Roughgarden

BGP with Path Verification

- While fake path announcements are possible and frequent in current BGP protocol, there is work being done to eliminate it
    - BGPsec protocol uses cryptographic signatures to verify the existence of announced paths
- Suppose we are able to detect and disallow fake path announcements
    - **Question**. Is BGP then incentive compatible?

Credit: Roughgarden

BGP with Path Verification

- **Theorem**. (Levin et al.). Assuming that no AS can announce a path that doesn't exist and the AS graph does not have a dispute wheel
  - No AS has an incentive to unilaterally deviate from the intended behavior in BGP (assuming other ASes follow BGP)

SIAM J. COMPUT.
Vol. 40, No. 6, pp. 1892–1912

INTERDOMAIN ROUTING AND GAMES*

HAGAY LEVIN†, MICHAEL SCHAPIRA‡, AND AVIV ZOHAR‡

Abstract. We present a game-theoretic model that captures many of the intricacies of *interdomain routing* in today's Internet. In this model, the strategic agents are source nodes located on a network, who aim to send traffic to a unique destination node. The interaction between the agents is dynamic and complex—asynchronous, sequential, and based on partial information. Best-reply dynamics in this model capture crucial aspects of the de facto standard interdomain routing protocol, namely, the Border Gateway Protocol (BGP). We study complexity and incentive-related issues in this model. Our main results show that in realistic and well-studied settings, BGP is incentive-compatible. That is, not only does myopic behavior of all players *converge* to a "stable" routing outcome, but no player has motivation to unilaterally deviate from BGP. Moreover, we show that even *coalitions* of players of *any* size cannot improve their routing outcomes by collaborating. Unlike the vast majority of works in mechanism design, our results do not require any monetary transfers (to or by the agents).

Credit: Roughgarden

Rationality and Traffic Attraction: Incentives for Honest Path Announcements in BGP

Sharon Goldberg*
Princeton University

Shai Halevi
IBM Research

Aaron D. Jaggard†
Rutgers University

Vijay Ramachandran‡
Colgate University

Rebecca N. Wright§
Rutgers University

ABSTRACT

We study situations in which autonomous systems (ASes) may have *incentives* to send BGP announcements differing from the AS-level paths that packets traverse in the data plane. Prior work on this issue assumed that ASes seek only to obtain the best possible *outgoing path* for their traffic. In reality, other factors can influence a *rational* AS's behavior. Here we consider a more natural model, in which an AS is also interested in *attracting incoming traffic* (*e.g.,* because other ASes pay it to carry their traffic). We ask what combinations of BGP enhancements and restrictions on routing policies can ensure that ASes have no incentive to lie about their data-plane paths. We find that protocols like S-BGP alone are insufficient, but that S-BGP does suffice if coupled with additional (quite unrealistic) restrictions on routing policies. Our game-theoretic analysis illustrates the high cost of ensuring that the ASes honestly announce data-plane paths in their BGP path announcements.

http://ccr.sigcomm.org/online/files/p267-goldberg.pdf

Incentive-Compatible Interdomain Routing*

(Extended Abstract)

Joan Feigenbaum†
Yale University
New Haven, CT, USA
jf@cs.yale.edu

Vijay Ramachandran‡
ICSI
Berkeley, CA, USA
vijayr@icsi.berkeley.edu

Michael Schapira§
The Hebrew University
Jerusalem, Israel
mikesch@cs.huji.ac.il

ABSTRACT

The routing of traffic between Internet domains, or *Autonomous Systems* (ASes), a task known as *interdomain routing*, is currently handled by the Border Gateway Protocol (BGP) [17]. Using BGP, autonomous systems can apply semantically rich routing policies to choose interdomain routes in a distributed fashion. This expressiveness in routing-policy choice supports domains' autonomy in network operations and in business decisions, but it comes at a price: The interaction of locally defined routing policies can lead to unexpected global anomalies, including route oscillations or overall protocol divergence (see, *e.g.*, [20]). Networking researchers have addressed this problem by devising constraints on policies that guarantee BGP convergence without unduly limiting expressiveness and autonomy (see, *e.g.*, [7, 8]).

In addition to taking this engineering or "protocol-design" approach, researchers have approached interdomain routing from an economic or "mechanism-design" point of view. It is known that lowest-cost-path (LCP) routing can be implemented in a truthful, BGP-compatible manner [3] but that several other natural classes of routing policies cannot [2, 5]. In this paper, we present a natural class of interdomain-routing policies that is more realistic than LCP routing and admits incentive-compatible, BGP-compatible implementation. We also present several positive steps toward a general theory of incentive-compatible interdomain routing.

http://cs.brown.edu/courses/csci2950-g/papers/ic-routing.pdf

Project Report & Presentation

Project Report

- 2-page version which includes Background due this Friday
- Final report (at least 6 pages) due May 19 at 11 pm
  - See examples from Fall 2020 on GLOW
- A LaTeX template is now available at https://www.overleaf.com/read/tdqthvfvmrxt
- Submit PDF via Github repository; if your project consists of code
  - Make sure to include code in the repository along with a README file on how to run and test it

Grading Rubric

- **Scholarship / Background**: Project involves reading, understanding, and contextualizing the most relevant research related to the topic. Write up adequately describes the background on the problem studied and how the approach is related to the existing literature. Resources used are appropriately cited.
- **Contributions**: Consistent with the timeline, and makes adequate contributions.
- **Correctness:** The work presented does not contain technical mistakes or logical gaps. Any programs written must be with clear documentation and a README on how they should be run to reproduce the results. Theoretical results should include correct and complete proofs.
- **Creativity:** Creative component: e.g., goes beyond what has been suggested by the instructor, this can include interesting empirical analysis, theoretical observations and insights, extensions or conjectures.
- **Presentation**: Approach and results are presented in a coherent and appealing way that makes it easy to follow. For example, figures and examples are used to explain concepts, and data visualization or plotting is used to convey empirical results.

Project Presentation Goals

- To get a good **grade**, but is about much more
- **Sell the topic/work**. Pitch your work to the class
  - Why is it cool?
  - What were the interesting things you learnt/found?
  - Share your enthusiasm!
- **Challenges and techniques**. Share the challenges you faced and how you overcame them
- **Connect**. Tie to the topics we learnt in the class!

Project Presentation Logistics

- We have a total of 14 projects: 10 pairs, 4 solo
- We will split this into two parts: 7 talks on each
    - Monday (May 9) and Thursday (May 12)
    - Based on preference + diversity of topics
- Monday talks will be a bit behind of Thursday in terms of contributions till date: this is to be expected
- Each presentation must be no more than 9 minutes
    - 1 minute for question & answer
- Rule of thumb: a minute per slide in the worst case
    - Try to keep it to 9 slides

Presentation Outline

- **Problem definition** and motivation
  - What problem (or topic) did you focus on?
  - Why is it important?
- What were the **goals** of your project?
- What **related literature** did you look at?
- What are your **contributions**?
- **Results**. What were your findings/ what did you learn?
- **Conclusion** and progress

Theory vs Empirical

- **Theory project:** synthesize the material and present by connecting to topics covered in class
  - What tools/techniques did you learn/use?
  - Goal: to teach to the class!
  - What are the interesting future directions/takeaways?
- **Empirical/Implementation project**:
  - What tools/techniques did you learn/use?
  - Teach the algorithm you implemented to the class
  - What were your results/findings?

Presentation Tips

Tell a story

- What is the main takeaway message?
- Create a narrative around it
- Elements of a good story
    - *Characters, settings, the plot, the conflict, and the resolution*

Presentation Tips

Highlight your contributions & efforts

- **You** and **your work** is central to the theme
- Share **your** challenges
- How **you** overcame them (effort)
- What did **you** learn from this project?
  - Reflect on **your** work

Presentation Tips

Pictures and Commentary

- Use pictures: figures, graphs, clip arts!
- No one likes staring at a wall of text
- Point to draw attention!
- Text on slides is for the audience, not for you
    - Do not read slides word-for-word!

Presentation Tips

Less is More

- Don't try to cover too much!
- Be realistic about content that fits in 7 mins
    - Presentation vs report: each serve a different purpose
- A minute a slide is a good rule of thumb

Acknowledgments

- These lecture slides are partly based off the notes by Tim Roughgarden:
    - http://timroughgarden.org/f16/l/l8.pdf

Credit: Roughgarden