

# CSCI599 Internet Measurement

---

Ethan Katz-Bassett

[ethan.kb@usc.edu](mailto:ethan.kb@usc.edu)

Spring 2013

---

# A Quick Survey

Raise your hand if you used the Internet / email:  
...since you got to this room?  
...in the last hour?  
...today?

Raise your hand if you have **MEASURED** the  
Internet before?

---

# The Internet is BIG!

- Tons of users
  - 2.4 billion Internet users
  - More than 5 billion connected devices
- Tons of organizations
  - Hundreds of thousands of networks
  - Tens of thousands of ISPs
- Tons of time
  - Each minute, humankind collectively spends 30 years on Facebook
- Tons of money
  - Google revenue \$29 billion in 2010

---

# Why Measure the Internet?

---

# Why Measure the Internet?

- The Internet is a black box
  - Put traffic in, get traffic out
  - Hard to understand inner workings
  - Even ISPs have little visibility outside their network
- Personal motivations
  - My Internet isn't working!
- Economic motivations
  - Providers need to understand performance and availability to improve
- Infrastructural motivations
  - We depend on the Internet, so it must be resilient
- Scientific motivations
  - One of the largest systems humankind has built and has emergent properties
  - Propose and evaluate new protocols

---

# Aspects of the Internet to Measure

# Aspects of the Internet to Measure

- ❑ Node (router or host) properties
  - *Static*: IP aliases, geography, router role, owner, implementation features
  - *Dynamic*: failures
- ❑ Link properties
  - *Static*: delay, capacity, L2 connection, L2 switches
  - *Dynamic*: loss, reordering, delay variation, failure, utilization,, duplication, corruption
- ❑ Topology properties
  - *Static*: topology (AS, router, IP), routing (AS, IP), location
  - *Dynamic*: workload / traffic matrix
- ❑ Others? Evolution, more about traffic, applications,...

---

*From “Reverse Engineering the Internet”*

---

# Types of Network Measurement

Research varies along a number of dimensions

- Type of measurements
  - Active versus passive
- Type of network
  - Interdomain, intradomain, mobile, wireless, P2P, social network,...
- Type of research
  - New tool, new measurement study, ...
- Duration
  - One-off vs longitudinal



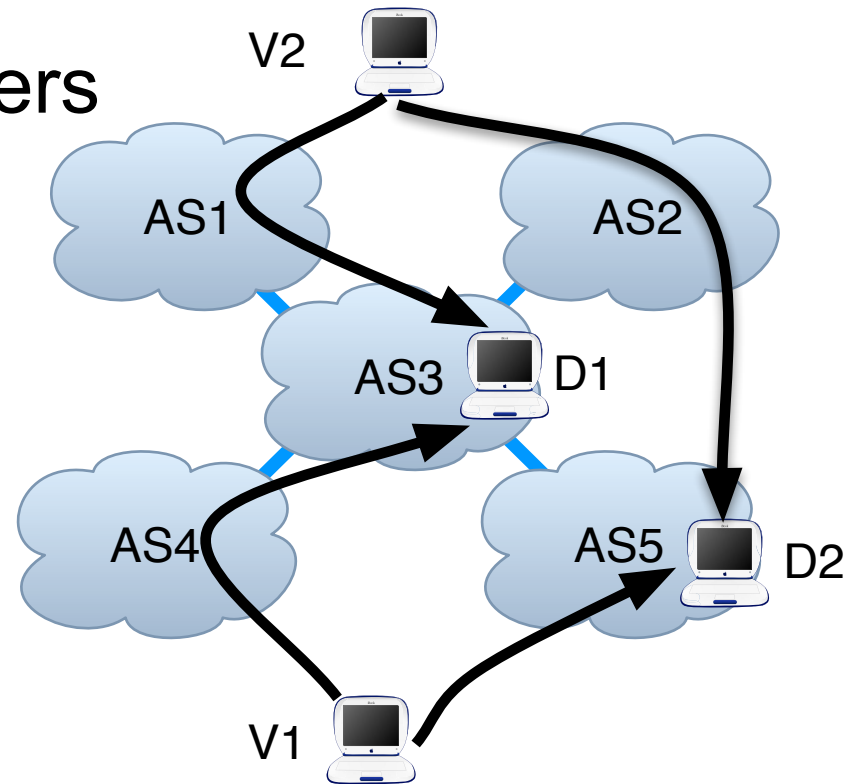
---

# Why is it Challenging?

- Hard to access data or vantage points
- Protocols do not expose information
- Networks do not expose information
  - Must cooperate
  - ...while competing

# Business Relationships Hide Topology

- AS3 peers with all others
- Paths contain at most one peer link
  - Why?
- No way to observe AS3-AS5 or AS3-AS2 from V1 or V2 **no matter what targets we probe**



---

# Questions to Keep in Mind

Do techniques & measurements support claims?

- Precision (and what is kept vs discarded)
- Accuracy
  - How well does tool's abstraction capture phenomena?
- Misconception
  - Does technique capture what we wish to capture?
- Representativeness of measurements

Consider reading “Strategies for Sound Internet Measurement”

---

---

# Calibrating for Sound Measurements

- Examine outliers and spikes
- Self-consistency checks
- Compare multiple measurements
- Synthetic data

---

# This Course

- Organized around aspects of Internet
  - Topology, routes and performance, traffic, applications
- Papers I really like
  - Plenty of great ones that we won't have time for
- How we can try to answer questions that are hard to answer
- Different studies that arrive at different answers for the same question

---

# Course Information

- Instructor  
Ethan Katz-Bassett  
ethan.kb@usc.edu  
SAL 236 by appointment
- Web forum: [www.piazza.com](http://www.piazza.com)
  - Sign up now, with your real name
  - Form project groups, post paper responses, etc
- [www-bcf.usc.edu/~katzbass/teaching/csci599-sp13/](http://www-bcf.usc.edu/~katzbass/teaching/csci599-sp13/)

---

# Structure of Course

- Reading and analyzing research papers
  - Mainly measurement studies and systems
  - Write critiques, comparisons, and observations
- Classroom presentations and discussions
  - Students present the papers
  - We discuss them
- Research project
  - Semester-long project in small groups
  - Novel Internet measurement research

---

# Grading

- Class participation (50%)
  - Written responses to papers on Piazza
  - Presentations
  - Discussions in class and on Piazza
- Project (50%)
  - Project proposal, presentation, final report
  - Research novelty, quality, writing
  - Written review of one other group's report



---

# Expect Hard Work

- Schedule time carefully before committing to this class
- Lots of work
  - Four written paper responses every week
  - Two class presentations on papers
  - Research study with novel ideas and results
  - Submission-quality research paper
- Class is research-oriented
  - Learning and critiquing measurement papers
  - Leads up to your own project

# How to Read

You May Think You Already Know  
How To Read, But...

# You Spend a Lot of Time Reading

- Reading papers for grad classes (like this one!)
- Reviewing papers for conferences/journals
- Giving colleagues feedback on their papers
- Keeping up with work related to your research
- Staying broadly educated about the field
- Transitioning into a new research area
- Learning how to write better papers 😊

**So, it is worthwhile to learn to read *effectively***

# Keshav's Three-Pass Approach: Step 1

- A ten-minute scan to get the general idea
  - Title, abstract, and introduction
  - Section and subsection titles
  - Conclusion
  - Bibliography
- What to learn: the five C's
  - Category: What type of paper is it?
  - Context: What body of work does it relate to?
  - Correctness: Do the assumptions seem valid?
  - Contributions: What are the main research contributions?

# Keshav's Three-Pass Approach: Step 2

- A more careful, one-hour reading
  - Read with greater care, but ignore details like proofs
  - Figures, diagrams, and illustrations
  - Mark relevant references for later reading
- Grasp the content of the paper
  - Be able to summarize the main thrust to others
  - Identify whether you can/should fully understand
- Decide whether to
  - Abandon reading the paper in any greater depth
  - Read background material before proceeding further
  - Persevere and continue on to the third pass

# Keshav's Three-Pass Approach: Step 3

- Several-hour virtual re-implementation of the work
  - Making the same assumptions, recreate the work
  - Identify the paper's innovations and its failings
  - Identify and challenge every assumption
  - Think how you would present the ideas yourself
  - Jot down ideas for future work
- When should you read this carefully?
  - Reviewing for a conference or journal
  - Giving colleagues feedback on a paper
  - Understand a paper closely related to your research

# Other Tips for Reading Papers

- Read at the right level for what you need
  - “Work smarter, not harder”
- Read at the right time of day
  - When you are fresh, not sleepy
- Read in the right place
  - Where you are not distracted, and have enough time
- Read actively
  - With a purpose (what is your goal?)
  - With a pen or computer to take notes
- Read critically
- Organize your notes in one place!

# Write a Response

- 4 papers every week
  - Pick four if there are five papers
  - Everyone should write the responses and post on piazza
  - Due Tuesday 7pm every week
- Response format
  - NOT a conference-style review
  - 1 Paragraph Summary per paper
  - 1+ Paragraphs **NOVEL** Response
    - Novel: Can't be a point that someone else already made
    - => You need to read what others wrote
    - Where appropriate, you can combine responses for multiple papers



# Review Format

- Summary
  - What problem the paper is addressing (1-2 sentences).
  - The core novel ideas or technical contributions of the work
    - What's the 30 second elevator pitch?
    - What should one remember about this paper?
  - A longer description (3-5 sentences) that summarizes the paper's approach, mechanisms, and findings.

# Review Format

- Novel response
  - By Saturday, I'll post some issues you can respond to
  - Or, the following are always good
    - What problems do you see with methodology that paper does not address? (Precision, accuracy, misconception, representativeness)
    - How would the results differ today? Why?
    - What study should we do as followup work?
    - Should we adapt the approach to a new setting?
    - Respond to something another student posted
    - Questions for others to respond to

# How to Give a Presentation

---

# Talk Arrangements

- In general, two students talk each week
- The two students should discuss/coordinate before class
  - I will mark weeks where this is especially important
- Each one covers two papers (about an hour)
- Can add other related papers (talk to me first)
- Read students' responses and incorporate some while preparing presentation
- Please sign up for slots on Piazza

---

# Talk Format

- Assume everyone has read the papers
  - NOT a conference presentation
- Can present papers back-to-back OR synthesize
- Borrow images and slides from authors
- Think critically
  - Is the problem real?
  - Is what we learn useful? Will the work be influential?
- Give your ideas / thoughts AND those of classmates

---

# Talk Format

- Motivation and background
    - What is the problem the paper tries to solve?
    - Why is it important?
  - Key idea
    - What techniques does paper use to solve problem?
    - Think about precision, accuracy, misconception, representativeness
  - Evaluation
    - How does paper demonstrate that technique works?
  - Results
    - What did we learn about Internet? What didn't we?
-

---

# Discussions

- Goal of presentation is to encourage and lead discussion
  - Point out interesting responses from classmates
  - Put questions on slides
- Discussion topics
  - Share ideas on extending the work
  - Comparison across papers
  - Aspects of papers we like/dislike?
- Rest of class expected to participate
  - Help your co-presenter out by participating
  - Everyone should talk in class almost every week

---

# Quiz (30 min)

- Write down your name, MS/PhD, year
- Background: Which networks/systems classes have you taken? Where?
- Questions
  1. Briefly explain how traceroute works.
  2. Draw TCP congestion control window as function of time, marking slow start and steady state.
  3. Briefly explain how BGP works. Why is it “policy routing?”
  4. What is one reason why paths on the Internet can be asymmetric?



# Course Project

---

# Class Project

- Project-oriented class
  - Analyze existing data in a new way
  - Redo old study today
  - Develop new measurement technique
- Papers should help you with project
- Start thinking about a project now!
  
- Today, I'll present some suggestions
  - But, you are free to propose your own
  - Either way, talk to me to refine the idea

---

# Project Timeline

- Form a group (by Jan 30)
    - Two or three students in a group
  - Project proposal (by Feb 13)
    - 2 single-column pages
    - Summarize your idea, related work, project plan
  - Mid-semester status meeting (March 25-29)
  - Project presentation (May 1)
  - Project report (May 13)
  - Review of another group's report (May 15)
-

---

# Finding a Project

- Reuse a project (that is in progress)
    - Your current research
    - A project for another class this semester
    - Must be related to network measurement
  - From a list of potential projects
    - I'll summarize briefly today and am available to provide more details
  - From papers you read
    - Extend ideas, address open questions or limitations
    - Consider glancing ahead in syllabus
    - Ask me for references if you are interested in an area
-

---

# Some Data Sets You Might Consider

- ISI's IP census and hit lists
  - Which IP addresses respond to ping?
- iPlane
  - 6 years of daily traceroutes: PlanetLab to every prefix
- CAIDA - Ark etc.
  - Daily traceroutes
- Ono
  - Traceroutes between BitTorrent clients
- MLab
  - Hosts a variety of measurement tools and data
- RouteViews / RIPE / UCLA Internet topology
  - Collections of BGP feeds

# Theme: The Mobile Internet

---

# A Local View of Cellular Performance

- Why is mobile performance so bad sometimes?
- Capture traces of common use cases
  - Web search
  - Facebook app
  - Searching for and viewing YouTube video
- Take traces to see what you find
  - TCP dumps
  - Radio state
  - Android network log

---

# Mobile Carrier Topology Evolution

- 2011 paper: Carriers have 4-6 ingresses
  - No WAY that is still accurate
- What do mobile carriers look like now?
- Where does mobile traffic enter the Internet?
  - Important performance implications
- How is this changing over time?
- Use Mobiperf data to see what it looks like now AND set up ongoing analysis for evolution
  
- References: “Cellular Data Network Infrastructure Characterization,” [www.mobiperf.com](http://www.mobiperf.com)



---

# Mobile Performance

- Internet performance on mobile devices is often bad and is poorly understood
  - Environment: location, obstructions, interference
  - Technology: access technology, available spectrum, device features, resource management
- Mobiperf measures from real devices
- Can we disentangle factors?
- Can we produce performance comparisons between carriers and devices?
- Combine measurements from 1000s of users to predict performance for a user

# Theme: Understanding Devices from Afar

---

# Measuring DC and Cloud Networks

- Data centers are now massive networks
- Can you develop techniques to uncover their structure?
  - Topology, routing, switches, applications, etc.
  - Adapt techniques from network tomography
- How does visibility differ by provider?
- Test on Deterlab, where we have ground truth
  
- With professor Minlan Yu

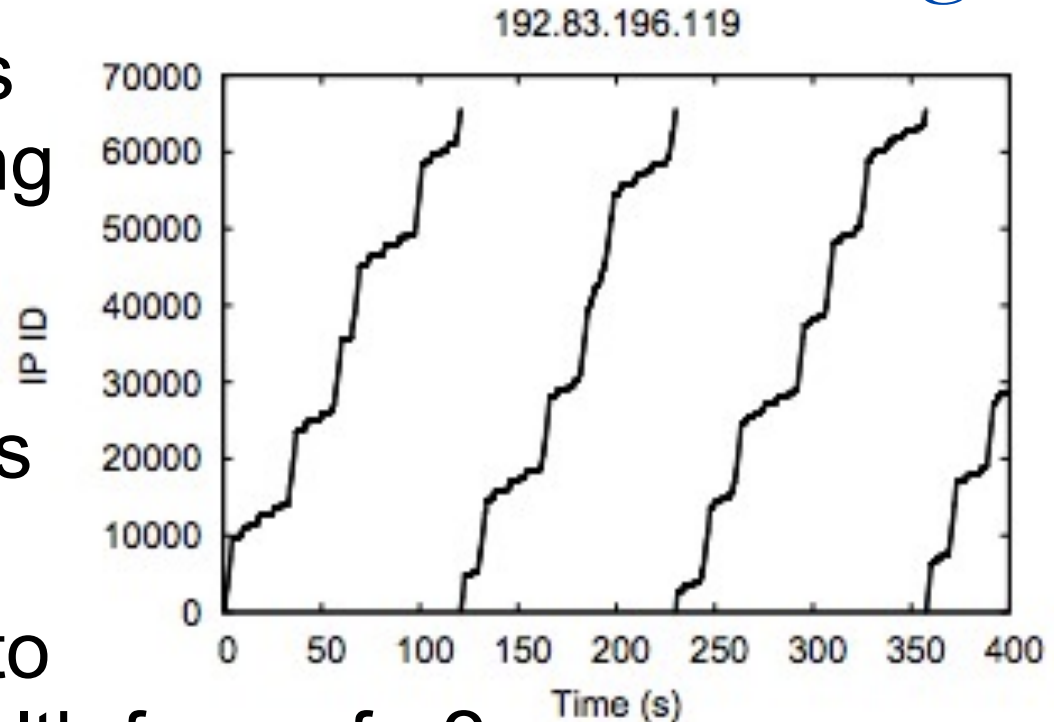
---

# Router Model / Middlebox Traceroute

- Traceroute views all hops the same
- But paths traverse a variety of routers and middleboxes
- Devices have signatures in terms of how they treat certain types of packets and/or in terms of how they are commonly named
- Can you create a traceroute-like tool that captures which devices are along a path?
- With professor Minlan Yu
- References: “Is it Still Possible to Extend TCP?,” “How to Accurately Interpret Traceroute Results,” “DisCarte,” etc

# Lightweight Router Health Monitoring

- Router increments IPID when sourcing packet
- IPID rate equals rate router sources packets
- Can we use IPID to monitor router health from afar?
- I have a data set. Can you learn behavior from it?



# Theme: Coverage of Measurement Approaches

# Are there enough traceroute servers?

- Some networks offer public traceroute servers
  - Give operators a different view during problems
  - (and we MIGHT use them in measurement systems)
- Little is known about the set available
- Characterize available servers
  - Census
  - Characterize path diversity and coverage
- How does one choose which to use for a given goal?
- References: “Where the sidewalk ends,” “IXPs: Mapped?,” “Quantifying the Importance of Vantage Point Distribution”

---

# Do Networks Support IP Options?

- Standard part of IP protocol, but rarely used
- Reverse traceroute has been using for years
- Evaluate logs to find:
  - Has support changed? Did reverse traceroute cause change?
  - Where are blind spots?
  - Can we improve the system to adapt?
  - How do, e.g., load-balancers treat options?
- References: “IP Options are not an Option,” “Reverse Traceroute,” “DisCarte: A Disjunctive Internet Cartographer”



# Theme: Measurement in the Large

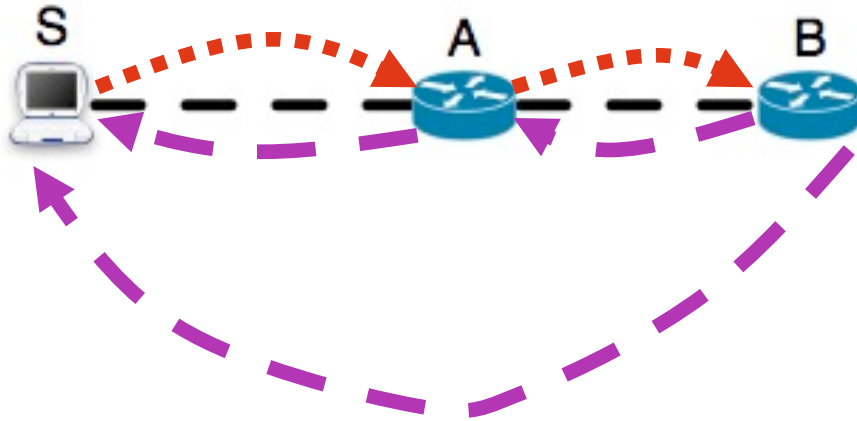
---

# Smart Allocation of Probes

- Systems need to track Internet conditions
- Parts of the Internet change at different rates
- How to allocate probing budget from distributed vantage points to best track changes?
  - Coordinated loss and congestion measurements
  - Predict changes cheaply to focus probes
  - Avoid probing the same change from multiple VPs
- With professor Italo Cunha (UFMG)
- References: “Predicting and Tracking Internet Path Changes,”  
“Deployment of an Algorithm for Large-Scale Topo. Discovery”


# Measuring Link Latency

- Many applications want link latencies
- IP geolocation, ISP performance, performance prediction, ...



- Traditional approach is to assume symmetry:  
$$\text{Delay}(A,B) = ( \text{RTT}(S,B) - \text{RTT}(S,A) ) / 2$$
- Asymmetry skews link latency inferred with traceroute
- We proposed 3 techniques but haven't evaluated widely
  - Can you build a link latency map of the Internet?
- Refs: "Reverse Traceroute," Justine's and Harsha's theses


# Reverse Traceroute Detects Symmetry



Solved  
(S,A)  
(S,C)

- Reverse traceroute identifies symmetric traversal
  - Identify cases when RTT difference is accurate
  - We can determine latency of **(S,A)** and **(S,C)**

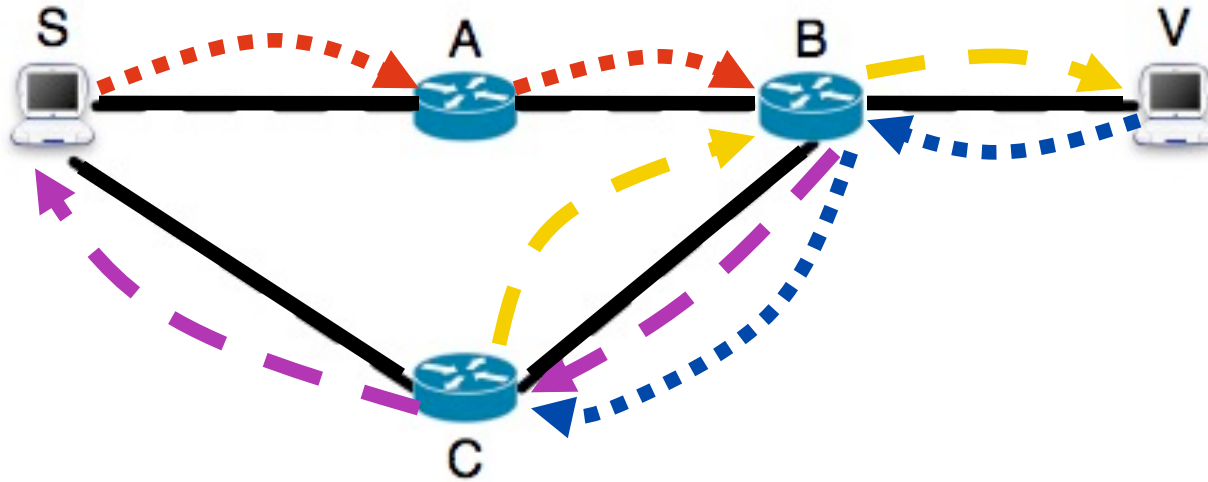
# Reverse TR Constrains Link Latencies



Solved  
(S,A)  
(S,C)

- Build up system of constraints on link latencies of all intermediate hops
  - Traceroute and reverse traceroute to all hops
  - $RTT = \text{Forward links} + \text{Reverse links}$

# Reverse TR Constrains Link Latencies




Solved

(S,A)  
(S,C)  
(V,B)  
(B,C)  
(A,B)

- Build up system of constraints on link latencies of all intermediate hops
  - Traceroute and reverse traceroute to all hops
  - $RTT = \text{Forward links} + \text{Reverse links}$

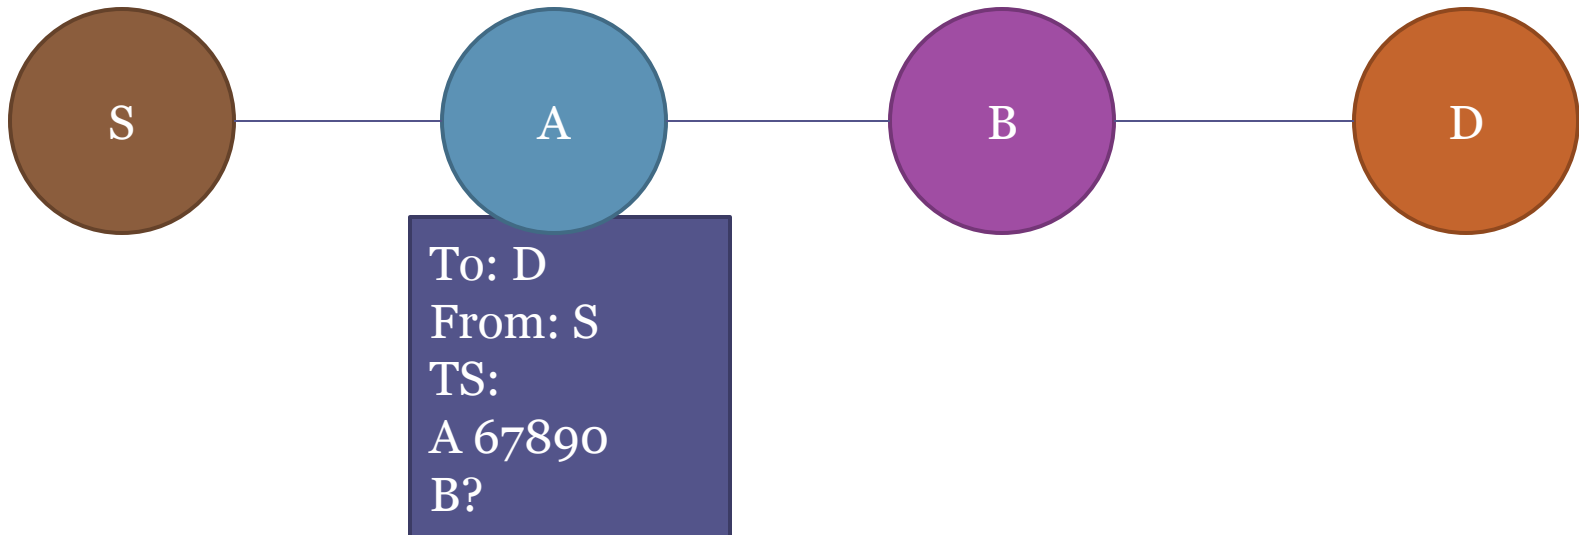
# A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps



# A Timestamp Measurement

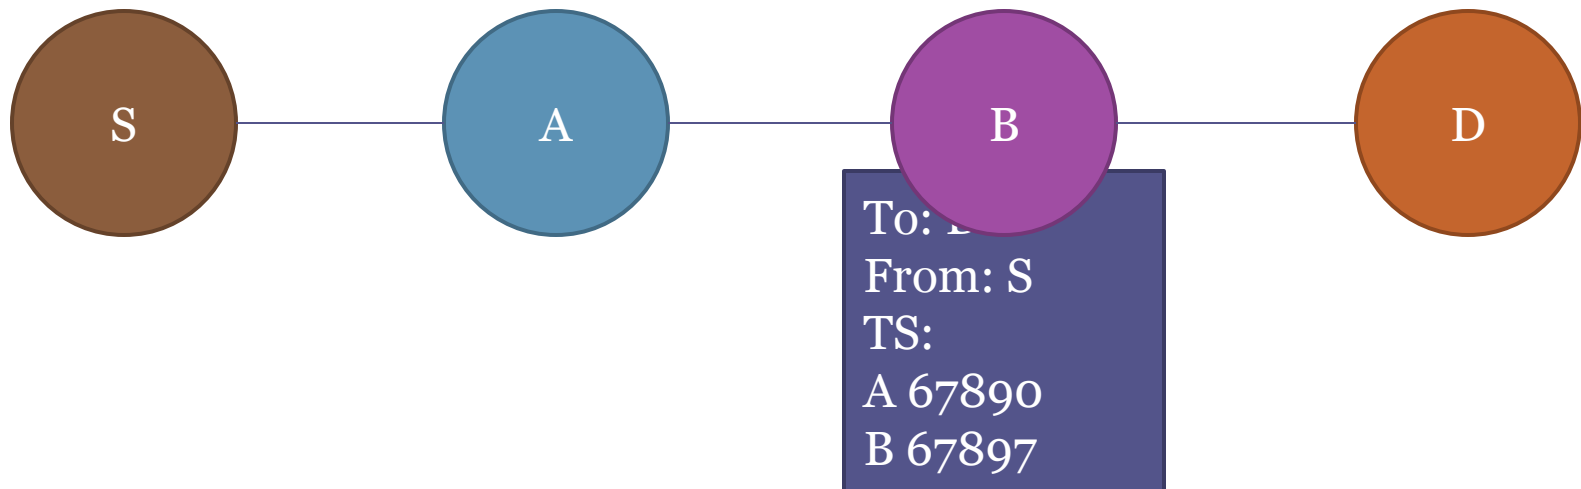
- Send a probe that traverses an A-B link, and ask A and B each for timestamps






# A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps



# A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps



# Components of Timestamp Values

- We can subtract A's timestamp from B
- The difference is 7 milliseconds
- But what does this difference comprise?

```
To: Steve  
From: D  
TS:  
A 67890  
B 67897
```

$$TS(A) - TS(B) = \text{latency} + \text{skew}(A,B) + \text{queue}$$

# Components of Timestamp Values

- We can subtract A's timestamp from B
- The difference is 7 milliseconds
- But what does this difference comprise?

```
To: Steve  
From: D  
TS:  
A 67890  
B 67897
```

$$TS(A) - TS(B) = \text{latency} + \text{skew}(A,B) + \text{queue}$$



Can ignore by  
taking the min  
across several  
measurements

# Components of Timestamp Values

- We can subtract A's timestamp from B
- The difference is 7 milliseconds
- But what does this difference comprise?

```
To: Steve
From: D
TS:
A 67890
B 67897
```


$$TS(A) - TS(B) = \text{latency} + \text{skew}(A,B) + \text{queue}$$



Still need to get rid  
of this!

# Canceling out Skew

- What if we could measure the B-A Link in the opposite direction?
- With many PlanetLab nodes, we can find a path that crosses the link in the opposite direction



To: D2  
From: S2  
TS:  
B?  
A?

# Canceling out the Skew

$$\Delta_1 = \text{TS}(A) - \text{TS}(B) = \text{latency} + \text{skew}(A,B) + \text{queue}$$

To: Steve  
From: D  
TS:  
A 67890  
B 67897

$$\Delta_2 = \text{TS}(A) - \text{TS}(B) = \text{latency} - \text{skew}(A,B) + \text{queue}$$

To: Ethan  
From: D2  
TS:  
B 67900  
A 67912

$$\text{So... latency} = \Delta_1 + \Delta_2 / 2$$

---

# IPv4 vs IPv6 Performance

- Do IPv4 and IPv6 see different performance along routes?
  - Throughput, latency
  - Do different address families see different support or prioritization?
  - Do they traverse different middleboxes?
- MLab is interested in this and has data
- References: MLab, “Measuring the Deployment of IPv6: Topology, Routing, and Performance”



# Evolution of Internet Routes

- iPlane has 6 years of daily traceroutes from 100s of vantage points to 100Ks of prefixes
- What long-term trends in routing can be observed in the data?
  - Have routes grown shorter?
  - Have some prefixes benefited more?
  - How has topology (AS, PoP, router) evolved?
- TBs of data -> analyze on a cluster or cloud using something like Hadoop
- With professor Harsha Madhyastha at UCR  
[http://iplane.cs.ucr.edu/iplane\\_logs](http://iplane.cs.ucr.edu/iplane_logs)
- References: “Has Internet Delay Gotten Better or Worse?”

---

# Understanding Latency at Scale

- Each day, each iPlane vantage point probes the prefixes in a random order
- Over time, we have many samples of latency
- Have latencies reduced over time?
- Is it possible to model path latency as a function of time of day, even though a single path may traverse several time zones?
- With professor Harsha Madhyastha at UCR  
[http://iplane.cs.ucr.edu/iplane\\_logs](http://iplane.cs.ucr.edu/iplane_logs)
- References: “Has Internet Delay Gotten Better or Worse?”

# What if iPlane predicts multiple paths?

- iPlane predicts path performance between arbitrary hosts by first predicting the path
- Predicts path by:
  - Using traceroutes to build a model of routing
  - Stitching traceroutes to build new paths
  - Predicting which is most likely, throwing away rest
- What if we keep the rest?
  - Use them to build distribution of performance?
  - Predict path changes?
- References: “iPlane”, “iPlane Nano”

# Theme: CDN and Cloud Performance

---

# Understanding CDNs' Client Mapping

- Content delivery networks use DNS to direct client to a replica
- What do CDN deployments look like?
- Is user served by same data center over time?
  - Short term switches: misconfigurations, outages
  - Long term switches: Reflect AND drive network evolution
- With Xun Fan and Matt Calder

---

# What Eyeballs and Content See

- Not sure this will pan out, but...
  - Group at Georgia Tech has 3 unique testbeds
    - Hosts collocated with Level 3 CDN PoPs
    - Project BISmark
      - Home WiFi routers instrumented for measurement
    - BGPMux
      - ASN and prefix available for experiments
      - Five universities as providers
  - How do we combine these views to assess routing and CDN performance, identify problems, and develop schemes to optimize?
  - With professor Nick Feamster (Georgia Tech)
-

---

# That's It!

- Sign up for Piazza
- Start thinking about projects and groups
- Check for my prompts for the first week's papers
- Read papers and post responses