Tuning Trombones, Smoothing Speedbumps:
Why Internet Speeds in Africa are Slow, and What We Can Do About It

Nick Feamster
Princeton University

Broadband Connectivity in Africa

According to ITU in 2013

- 93 million broadband subscriptions
- 27% growth in past 4 years (**Highest**)
- Broadband associated with economic growth + development

Yet, little is known about performance in Africa and what causes poor performance when it does arise.

How Well Does Broadband Perform?

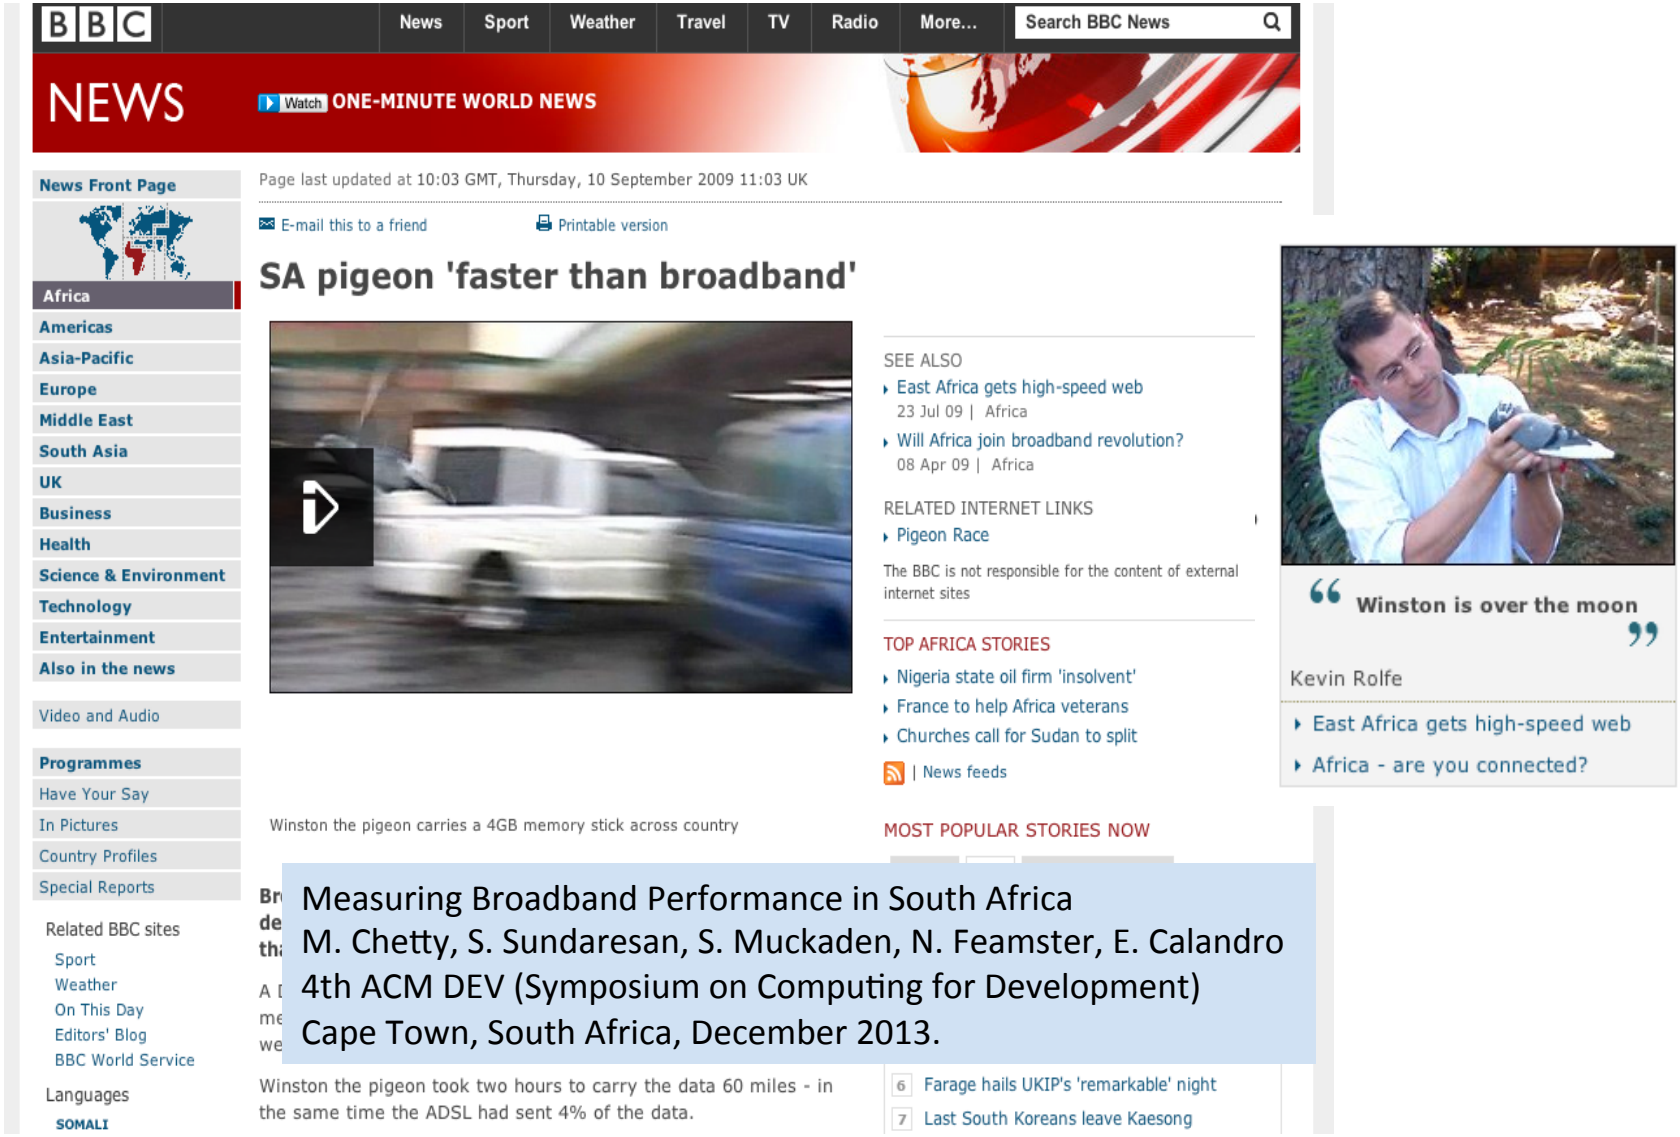

Measuring Broadband Performance in South Africa
M. Chetty, S. Sundaresan, S. Muckaden, N. Feamster, E. Calandro
4th ACM DEV (Symposium on Computing for Development)
Cape Town, South Africa, December 2013.

From Where To Measure?

End Host

Mobile Handset (MySpeedTest)

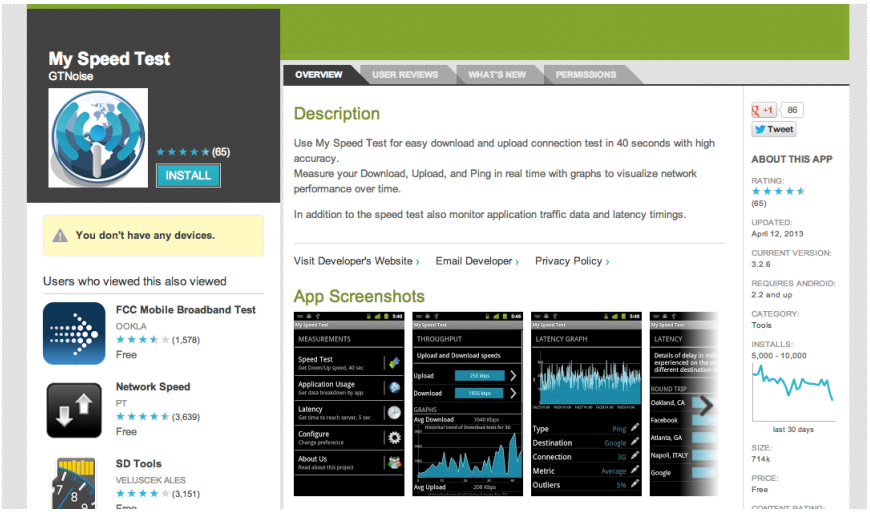

Home Gateway (BISmark)

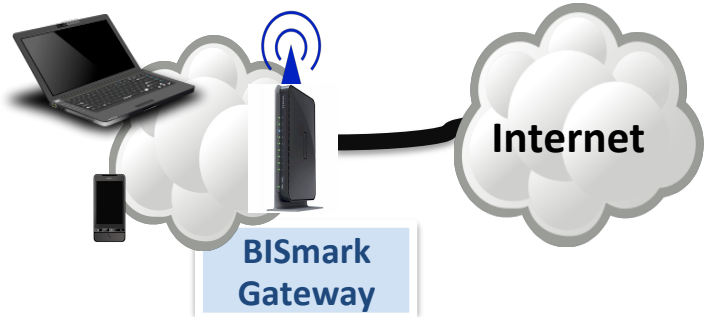

- Fixed-line measurements
- Dongle-based measurements

BISmark: Measurements from Fixed Line Locations

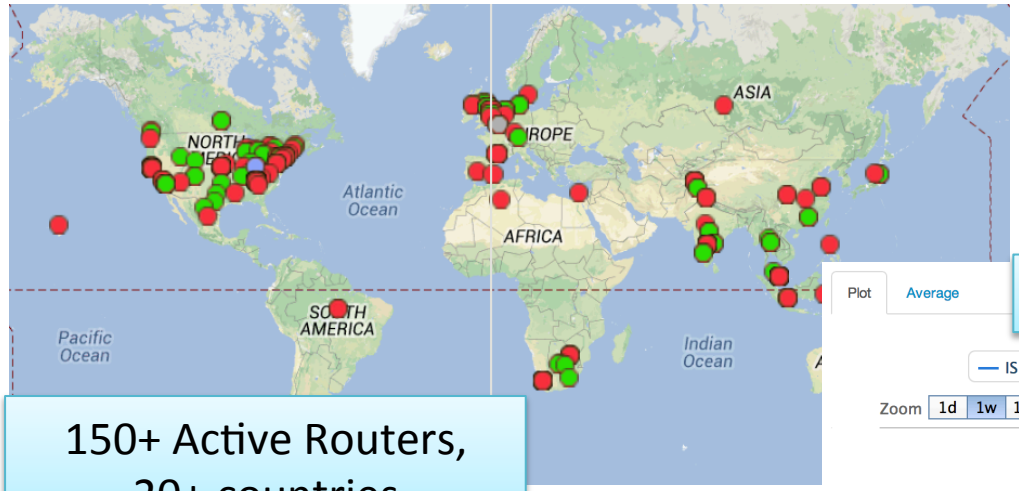

150+ Active Routers, 20+ countries

Example: Cape Town Users

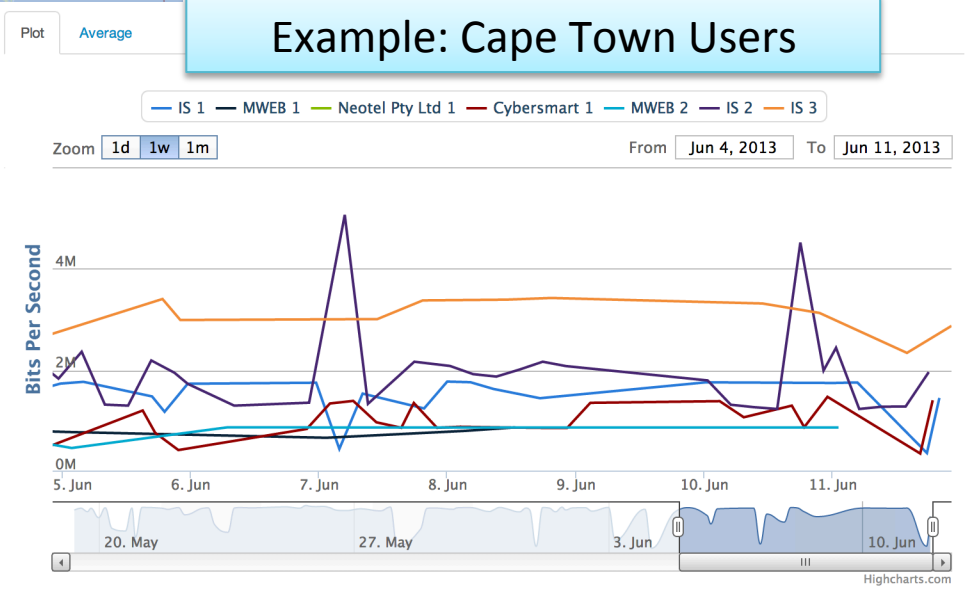

- Users install routers in home networks
- Custom firmware performs periodic measurements
- Can aggregate by country, city, ISP

My Speed Test: Measurements from Mobile Handset

- Periodic latency measurements
- Upload and download throughput tests on demand
- Available in Google Play Store: http://goo.gl/28tx3

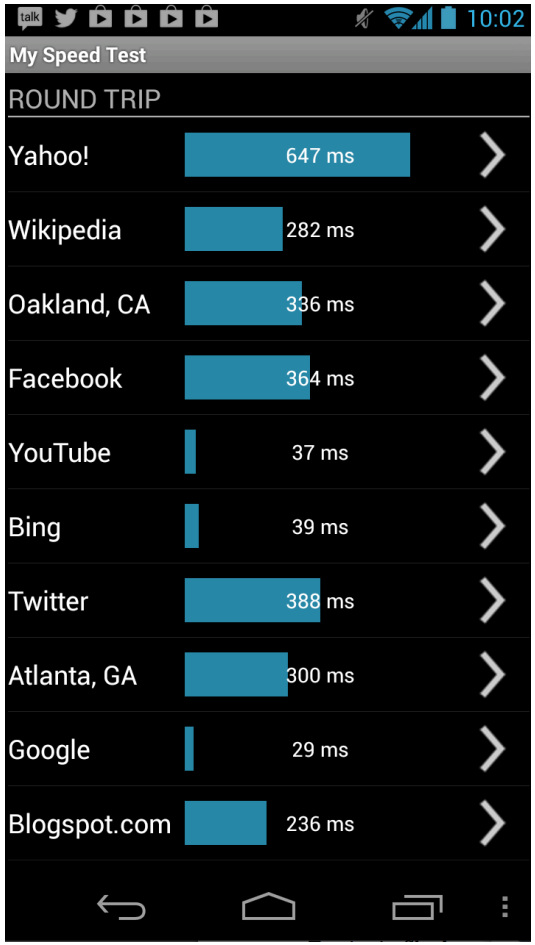

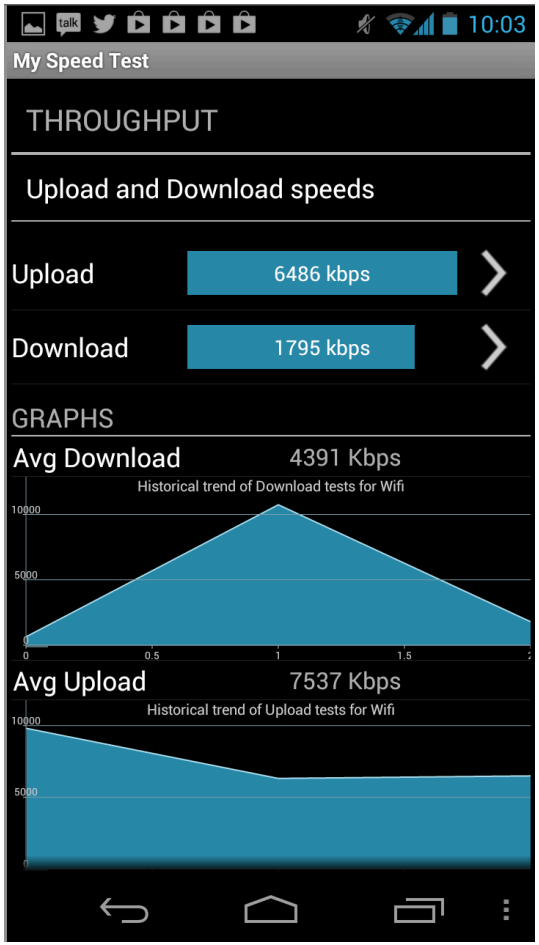

Summary of Results

- Performance consistently underperforms with respect to advertised rates
- Mobile broadband consistently achieves higher throughput than fixed broadband
- (Bad) peering can introduce significant latency, introduce fragility in times of failure (e.g., fiber cut)

Latencies to Nearby Locations are High

Latencies from South Africa to Kenya, Brazil, India are 2x higher than latencies to Europe.

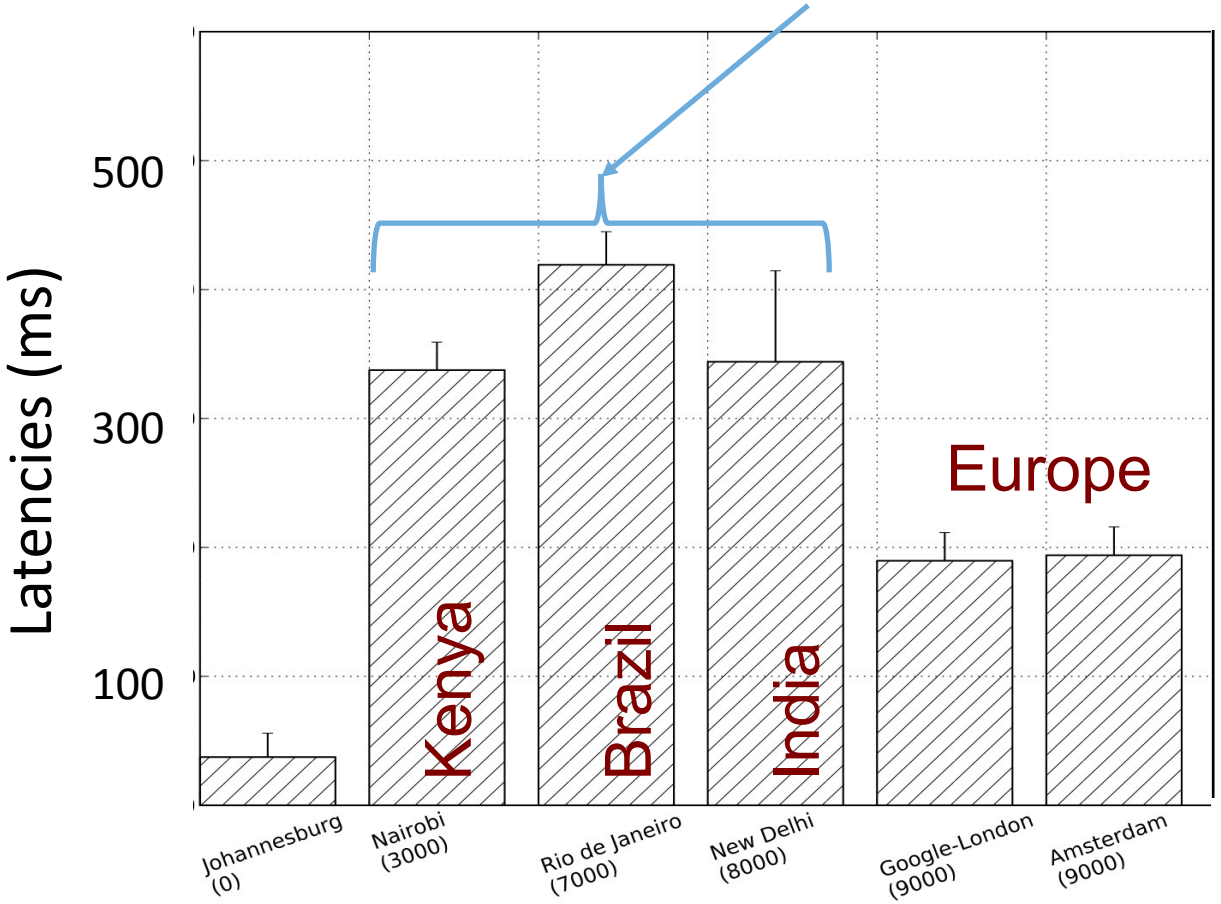

Latencies are Even Higher During Failures

- March 27, 2013 0620 UTC: SWM4 Fiber Cut
- All BISmark hosts could not reach KENet for 3+ hours
- Latencies remain high for another 8+ hours (except for Neotel, in South Africa)

More details: http://connectionmanagement.org

Latencies remain high from Europe and North America. Only one AS in S. Africa, Neotel sees decent latency after 3 hours.

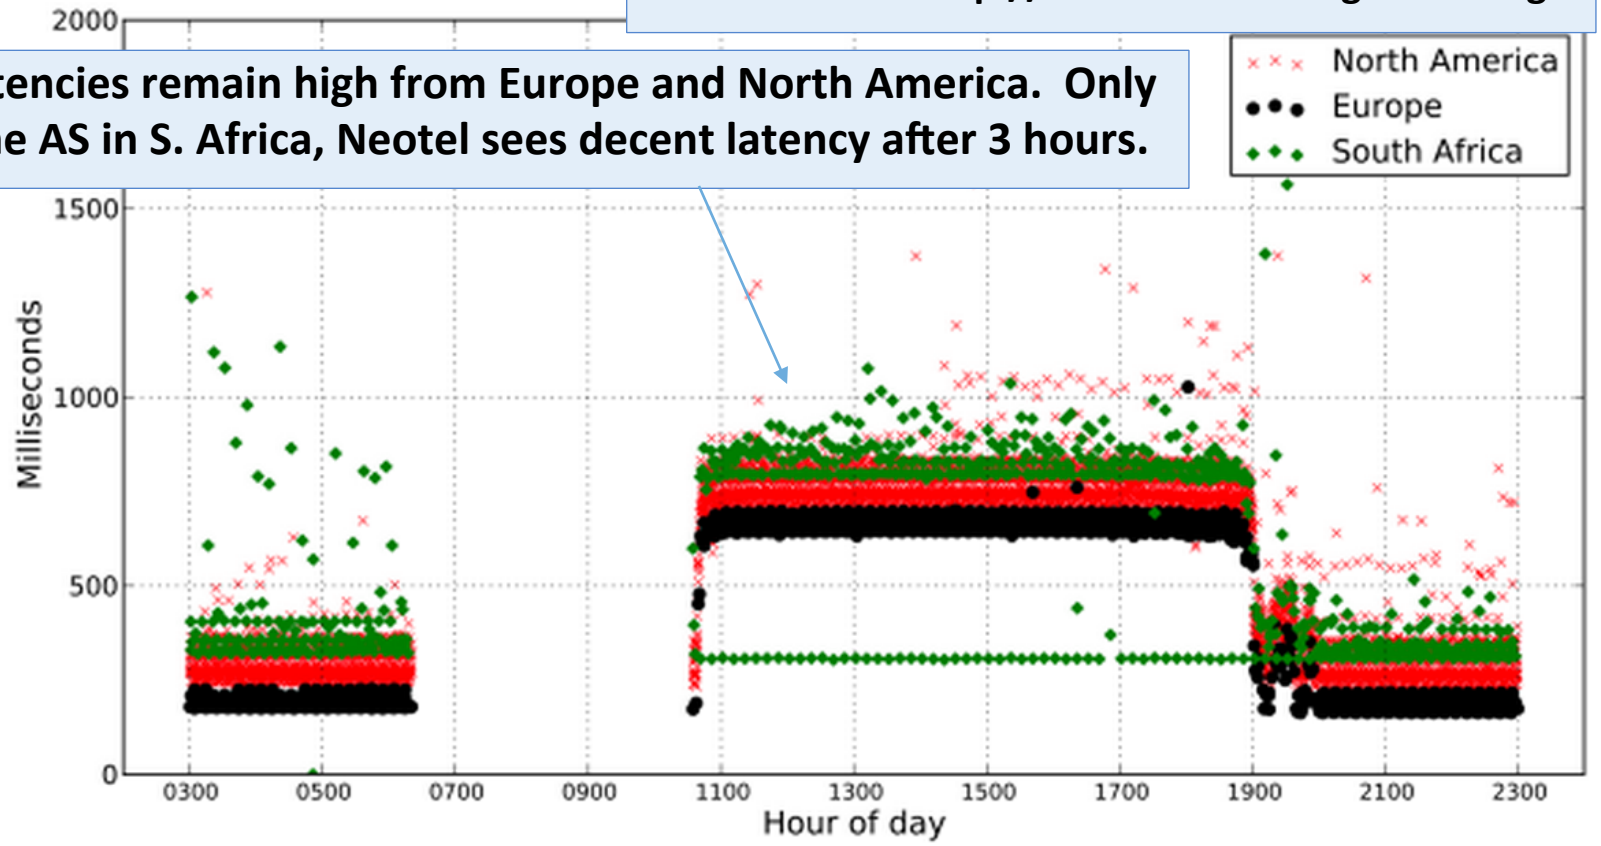

Causes of High Latency: Circuitous Routing Paths

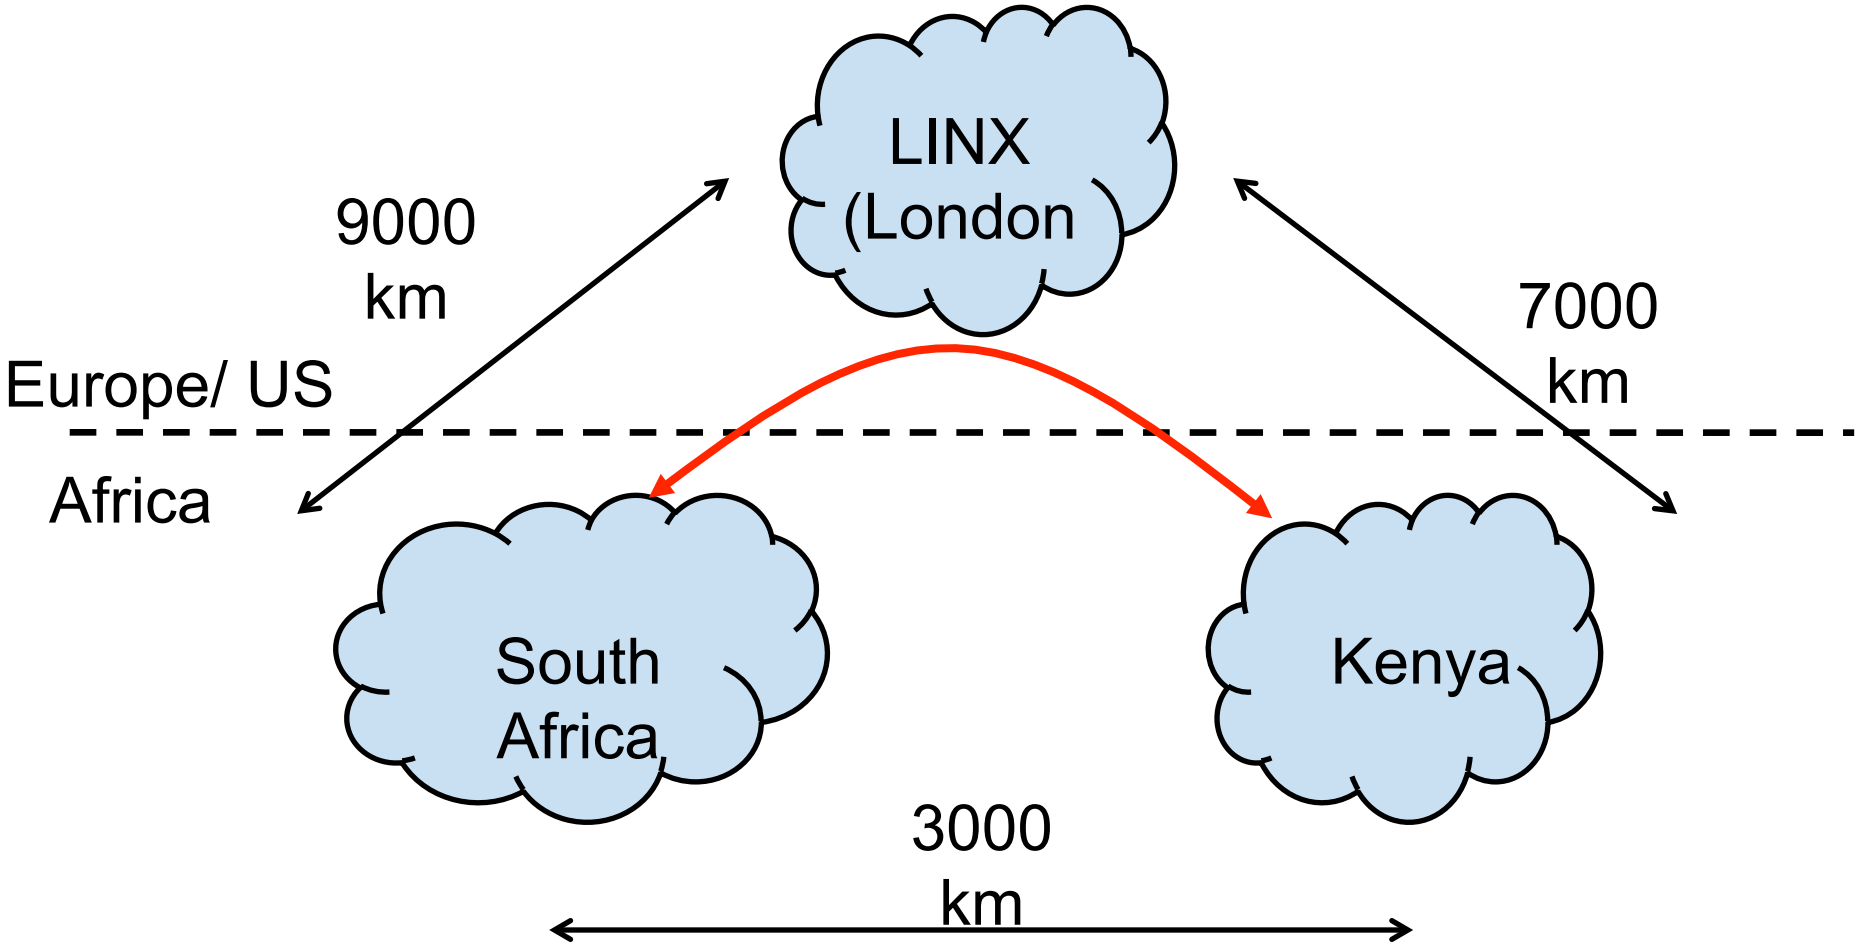

Two Questions

- What is the nature of Internet interconnectivity (between ISPs) in Africa?

- What can be done to reduce latency to common Internet services?

Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa
A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, E. Katz-Bassett
Passive and Active Measurement Conference
Los Angeles, CA, March 2014.

Two Questions

- **What is the nature of Internet interconnectivity (between ISPs) in Africa?**

- What can be done to reduce latency to common Internet services?

BISmark: Measurements from Fixed Locations

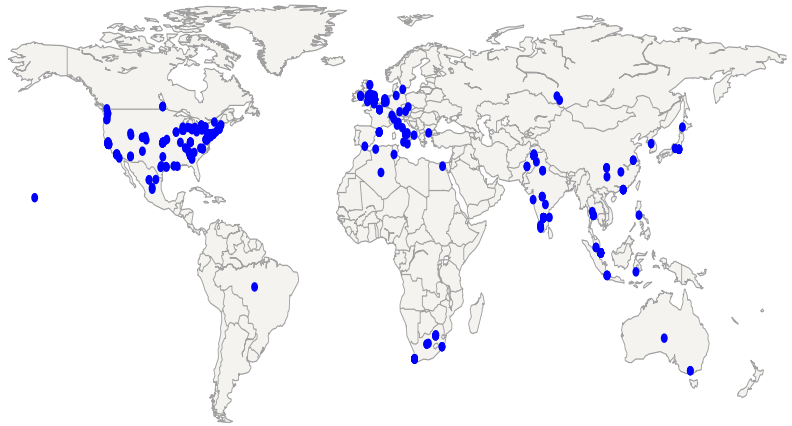

175+ Active Routers, 20+ countries

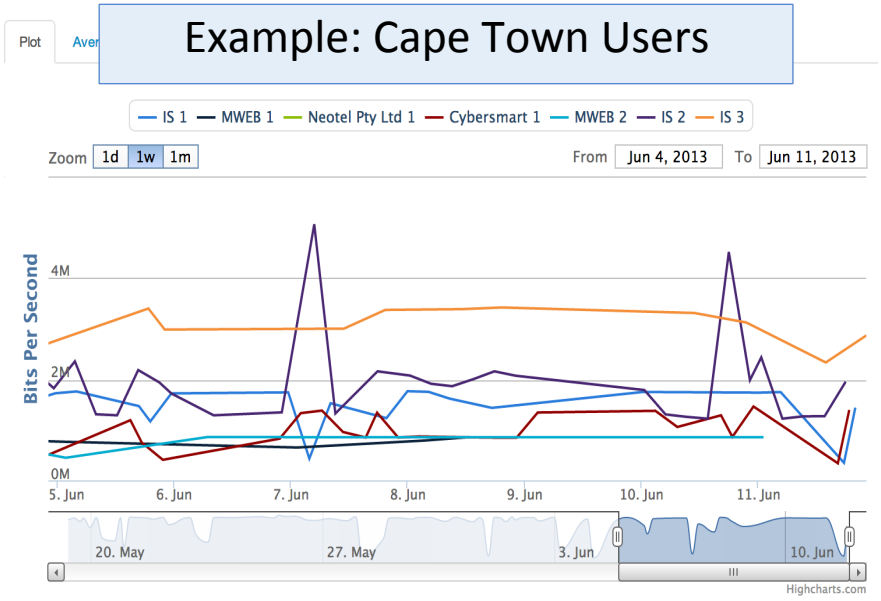

- Users install routers in home networks
- Custom firmware performs periodic measurements
- Can aggregate by country, city, ISP

BISmark Deployment in South Africa

- Periodic latency and throughput measurements
- Traceroutes to explain the cause of path performance
- Router-based deployment
  - 17 home networks, 7 ISPs, all 9 provinces

Destinations for Traceroute Probes

Global M-Lab Servers

Google Caches in Africa

High Latencies to Nearby Locations…

Cape Town (SA) to M-Lab Johannesburg (SA)

…
7, 196.44.0.74, 7.793, South Africa, AS16637
8, 196.223.22.24, 8.338, South Africa, Cape Town IXP
9, 41.164.0.243, 34.679, South Africa, AS36937
…
14, 196.24.45.146, **92.511**, South Africa, AS2018

Cape Town (SA) to M-Lab Nairobi (KE)

…
8, 209.212.111.201, 199.446, South Africa, AS16637
9, 195.66.225.31, 217.301, United Kingdom, London IXP (LINX)
10, 196.32.209.77, 201.569, South Africa, AS36944
…
14, 197.136.0.108, **368.107**, Kenya, AS36914

High Latency

… Circuitous Routing Paths

Cape Town (SA) to M-Lab Johannesburg (SA)

…
7, 196.44.0.74, 7.793, South Africa, AS16637
8, 196.223.22.24, 8.338, South Africa, Cape Town IXP
9, 41.164.0.243, 34.679, South Africa, AS36937
…
14, 196.24.45.146, **92.511**, South Africa, AS2018

Cape Town (SA) to M-Lab Nairobi (KE)

…
8, 209.212.111.201, 199.446, South Africa, AS16637
9, 195.66.225.31, 217.301, United Kingdom, London IXP (LINX)
10, 196.32.209.77, 20[illegible].569, Kenya, AS36944
…
14, 197.136.0.108, **36[illegible].107**, Kenya, AS36914

Traffic leaving Africa

Poor ISP Interconnectivity in Africa

- Reasons
  - Local ISPs not present at regional IXPs
  - IXP participants don’t peer with each other
- Consequences
  - Local traffic does not stay local
  - Paths leave continent

Local ISPs not Present at Regional IXPs

- ISPs prioritize connecting to European IXPs
- Fewer incentives to connect at regional ones

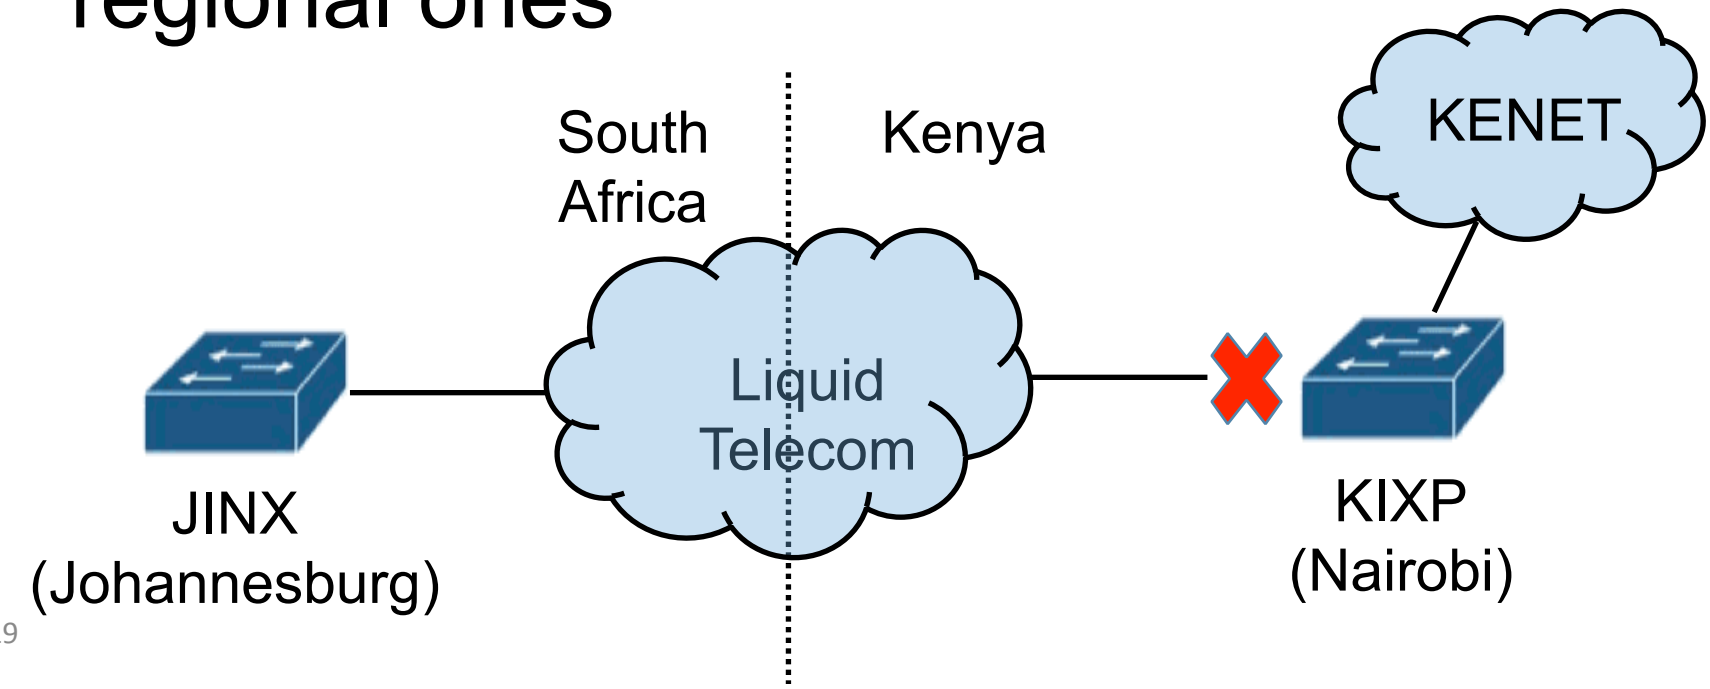

Missing Peering Links at Regional IXPs

- Most content not available locally
- Less incentive to peer with local ISPs

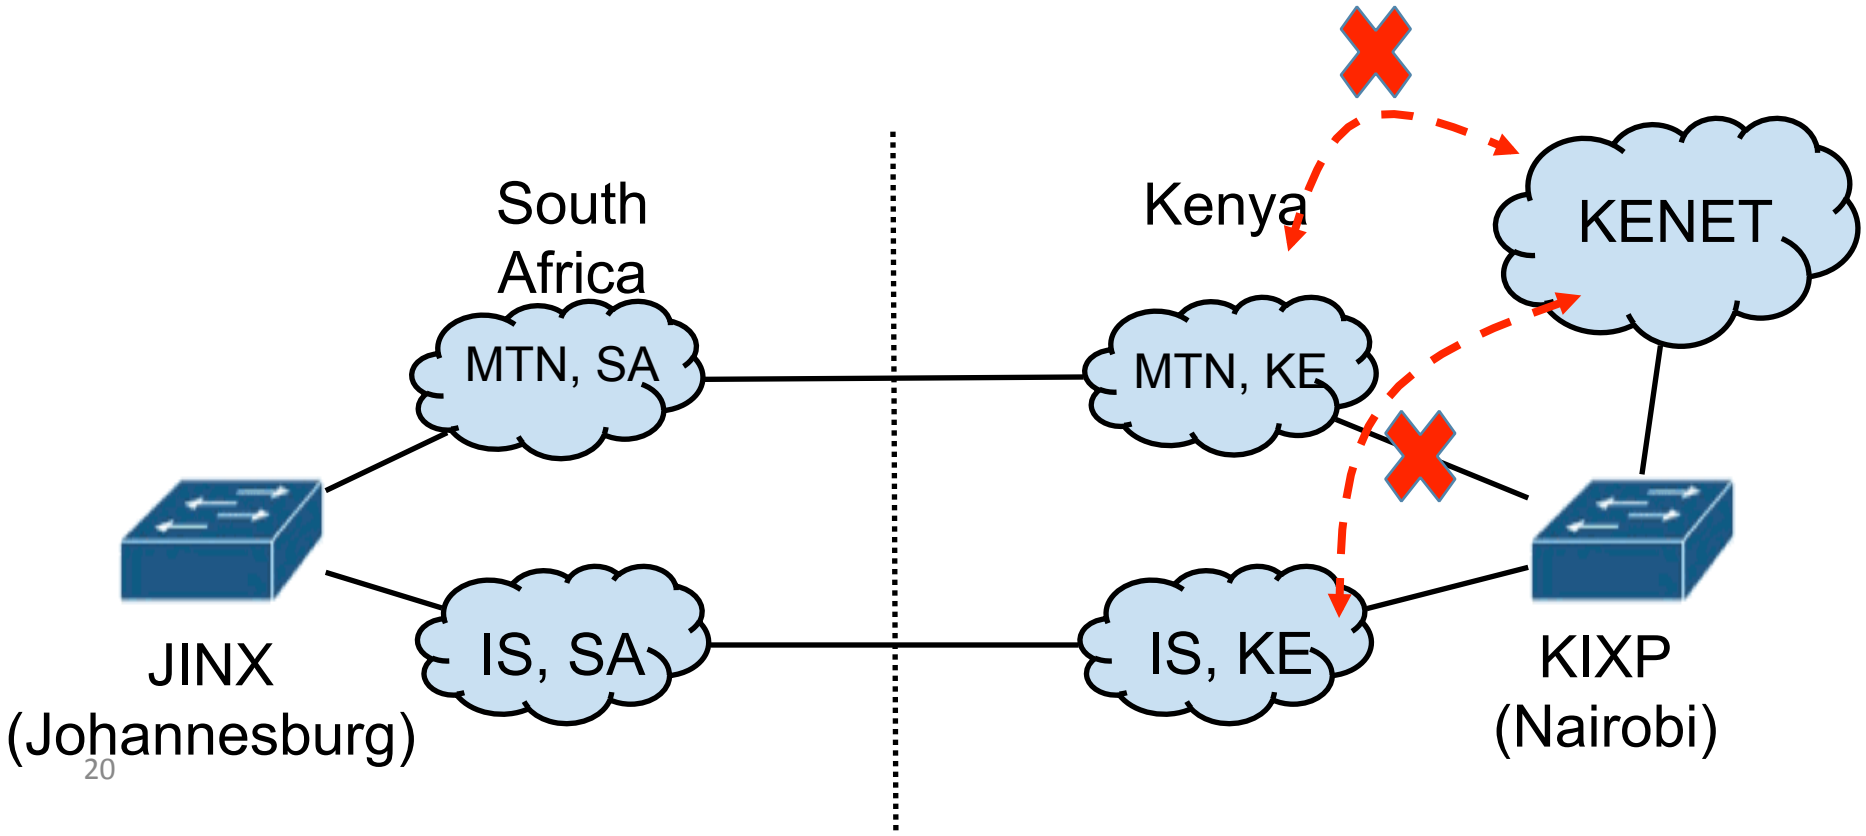

Regional IXPs Only Prevalent on Intra-Country Paths

Within South Africa:
High Fraction of Paths Have at Least one Major Regional IXP

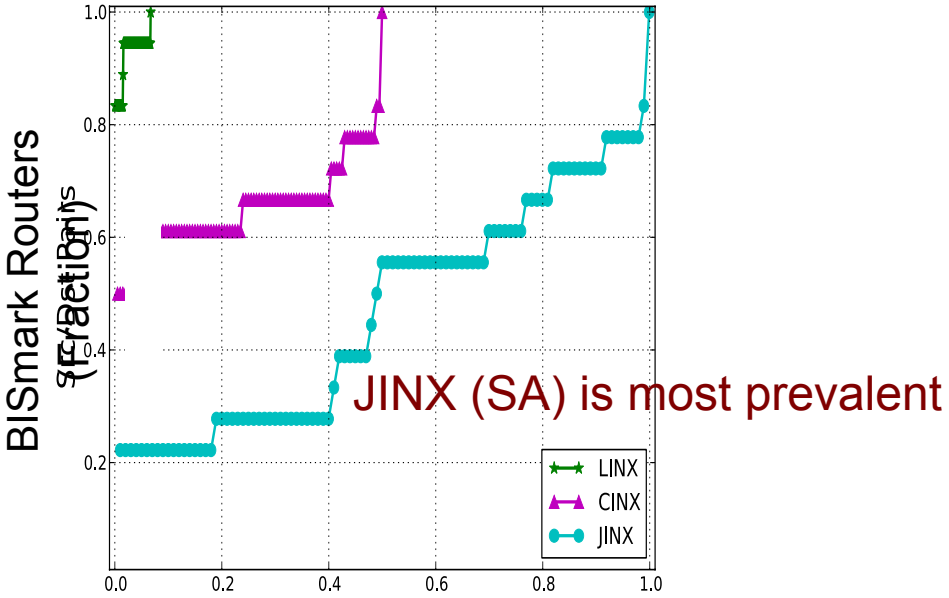

M-Lab Johannesburg

Between South Africa and Kenya: Few Paths have Regional IXPs

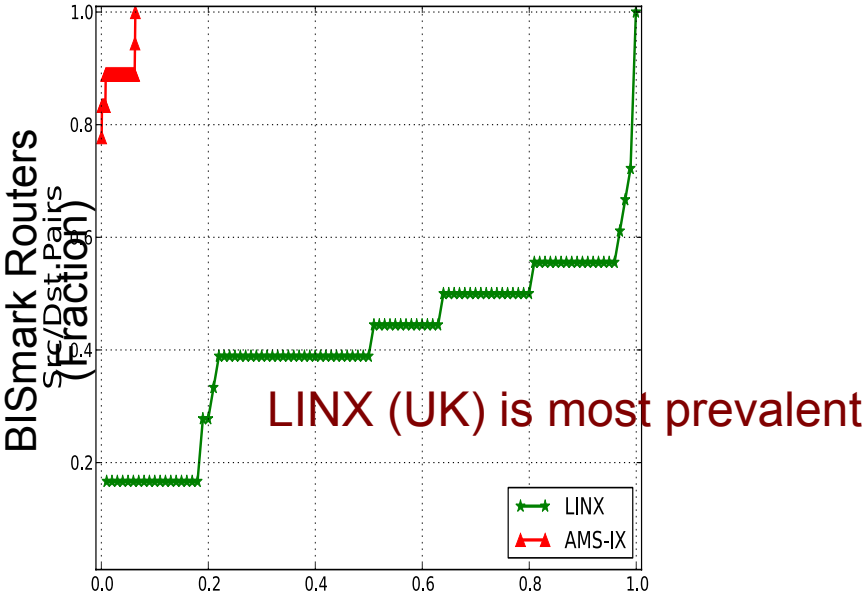

M-Lab Nairobi

Two Questions

- What is the nature of Internet interconnectivity (between ISPs) in Africa?

- **What can be done to reduce latency to common Internet services?**

Solution #1: Add More Caches

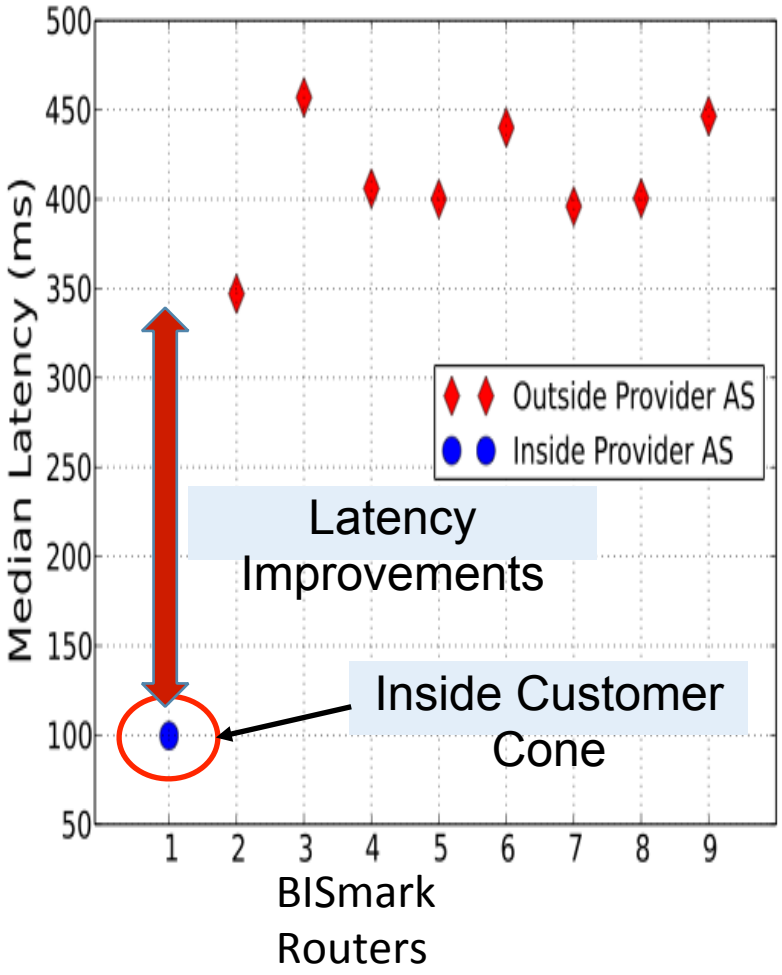

- Traceroute Probes between BISmark routers (eyeball) and Google Cache Node in Uganda (content)
- Google cache hosted by MTN
- Emulates scenario where content is in nearby country

Latency improvements are limited when peering to the cache is not adequate.

Solution #2: Add More Peering Links

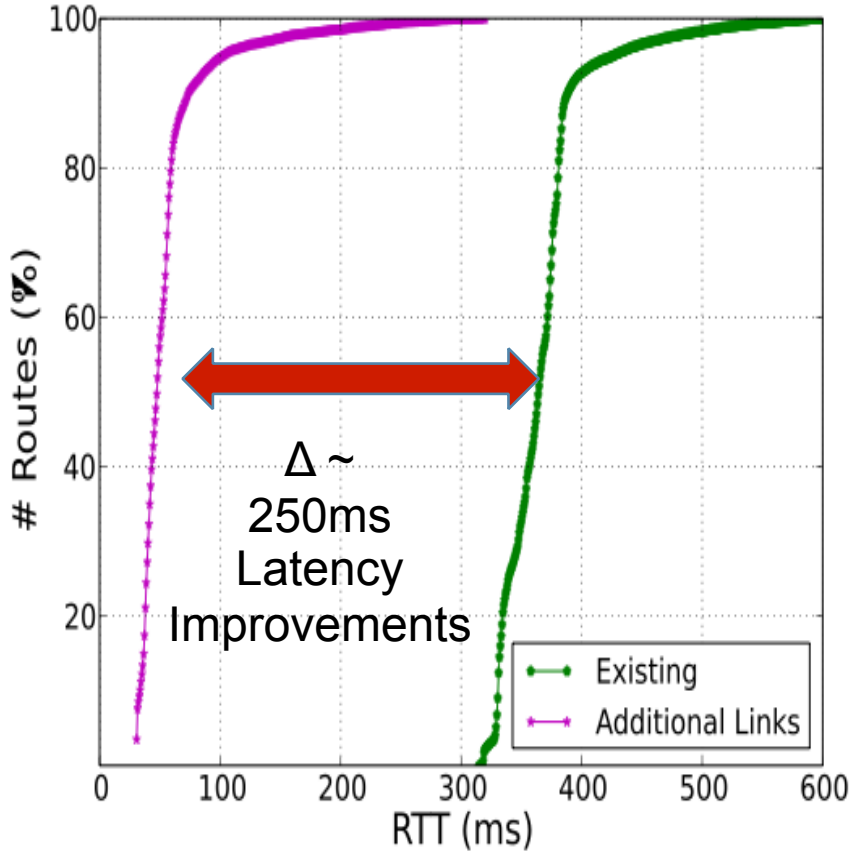

- **Simulation**: Add peering links between all the participants at
  - JINX (Johannesburg)
  - KIXP (Nairobi)
- Emulates scenario where more ISPs connect and peer at regional IXPs

Additional peering links → Significant latency improvements

Better Peering is a Substitute for Additional Caches

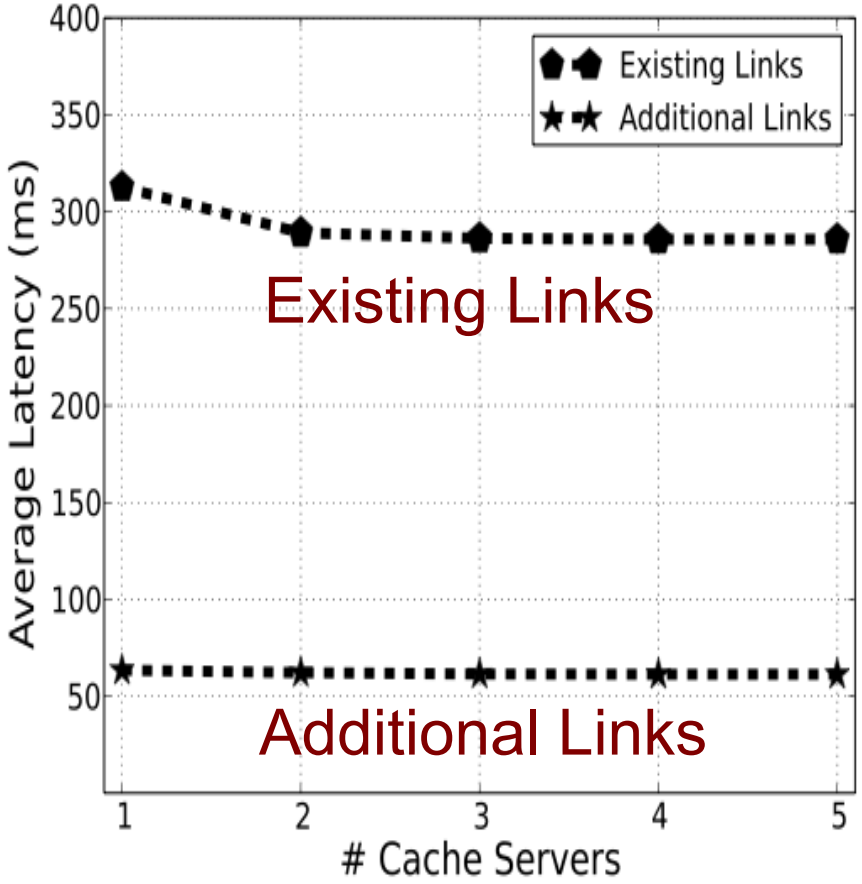

- Experiment:
  - add caches in Kenya
  - traceroute Probe from SA
- Two scenarios
  - Use existing peering links
  - Add more peering links

Additional caches have little effect on average latency (compared to adding more peering links).

Peering in South Africa: Summary

- What is the nature of Internet interconnectivity (between ISPs) in Africa?
  - Many ISPs are not present in regional IXPs
  - Many ISPs do not interconnect at regional IXPs
- What can be done to reduce latency to common Internet services?
  - Peering at regional IXPs can reduce median intra-continent latencies by 250ms
- **Next steps:** Better mechanisms for interconnectivity

Internet Routing is Not Flexible Enough

- Routing only on **destination IP address blocks**
  (No customization of routes by application or sender)
- Can only influence **immediate neighbors**
  (No ability to affect path selection remotely)
- **Indirect** control over packet forwarding
  (Indirect mechanisms to influence path selection)
- Enables only basic packet **forwarding**
  (Difficult to introduce new in-network services)

Valuable Wide-Area Services

- Application-specific peering
  - Route video traffic one way, and non-video another
- Blocking denial-of-service traffic
  - Dropping unwanted traffic further upstream
- Server load balancing
  - Directing client requests to different data centers
- Steering through network functions
  - Transcoders, scrubbers, caches, crypto, …
- Inbound traffic engineering
  - Splitting incoming traffic over multiple peering links

Software Defined Networking (SDN)

- Match packets on **multiple header fields**
  (not just destination IP address)
- Control **entire networks** with a single program
  (not just immediate neighbors)
- **Direct control** over packet handling
  (not indirect control via routing protocol arcana)
- Perform many different **actions** on packets
  (beyond basic packet forwarding)

Deploy SDN at Internet Exchanges

- **Leverage:** SDN deployment even at single IXP can benefit tens to hundreds of providers
  - *Without providers deploying new equipment!*
- **Innovation hotbed:** Incentives to innovate, as IXPs on front line of peering disputes
- **Growing in numbers:**
  - 350-400 IXPs
  - ~100 new IXPs established in past few years

Conventional IXPs

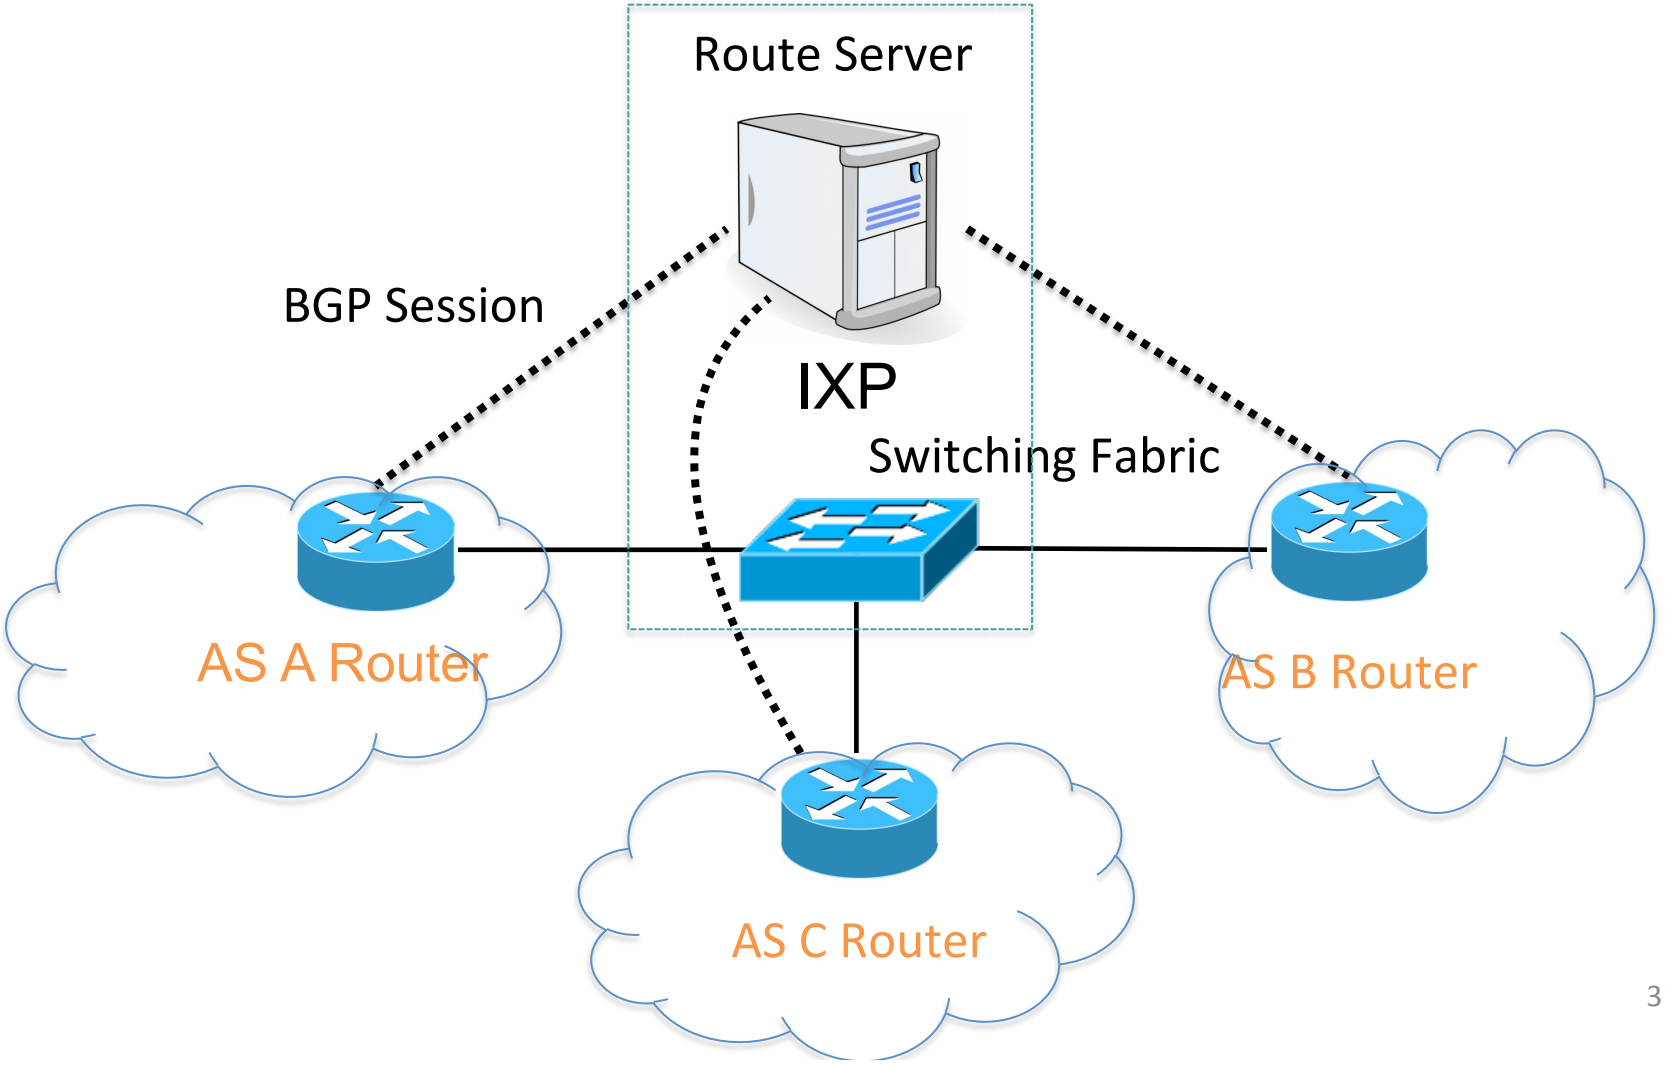

SDX = SDN + IXP

SDX: A Software Defined Internet Exchange
A. Gupta, L. Vanbever, M. Shahbaz, S. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, E. Katz-Bassett
ACM SIGCOMM
Chicago, IL. August 2014.

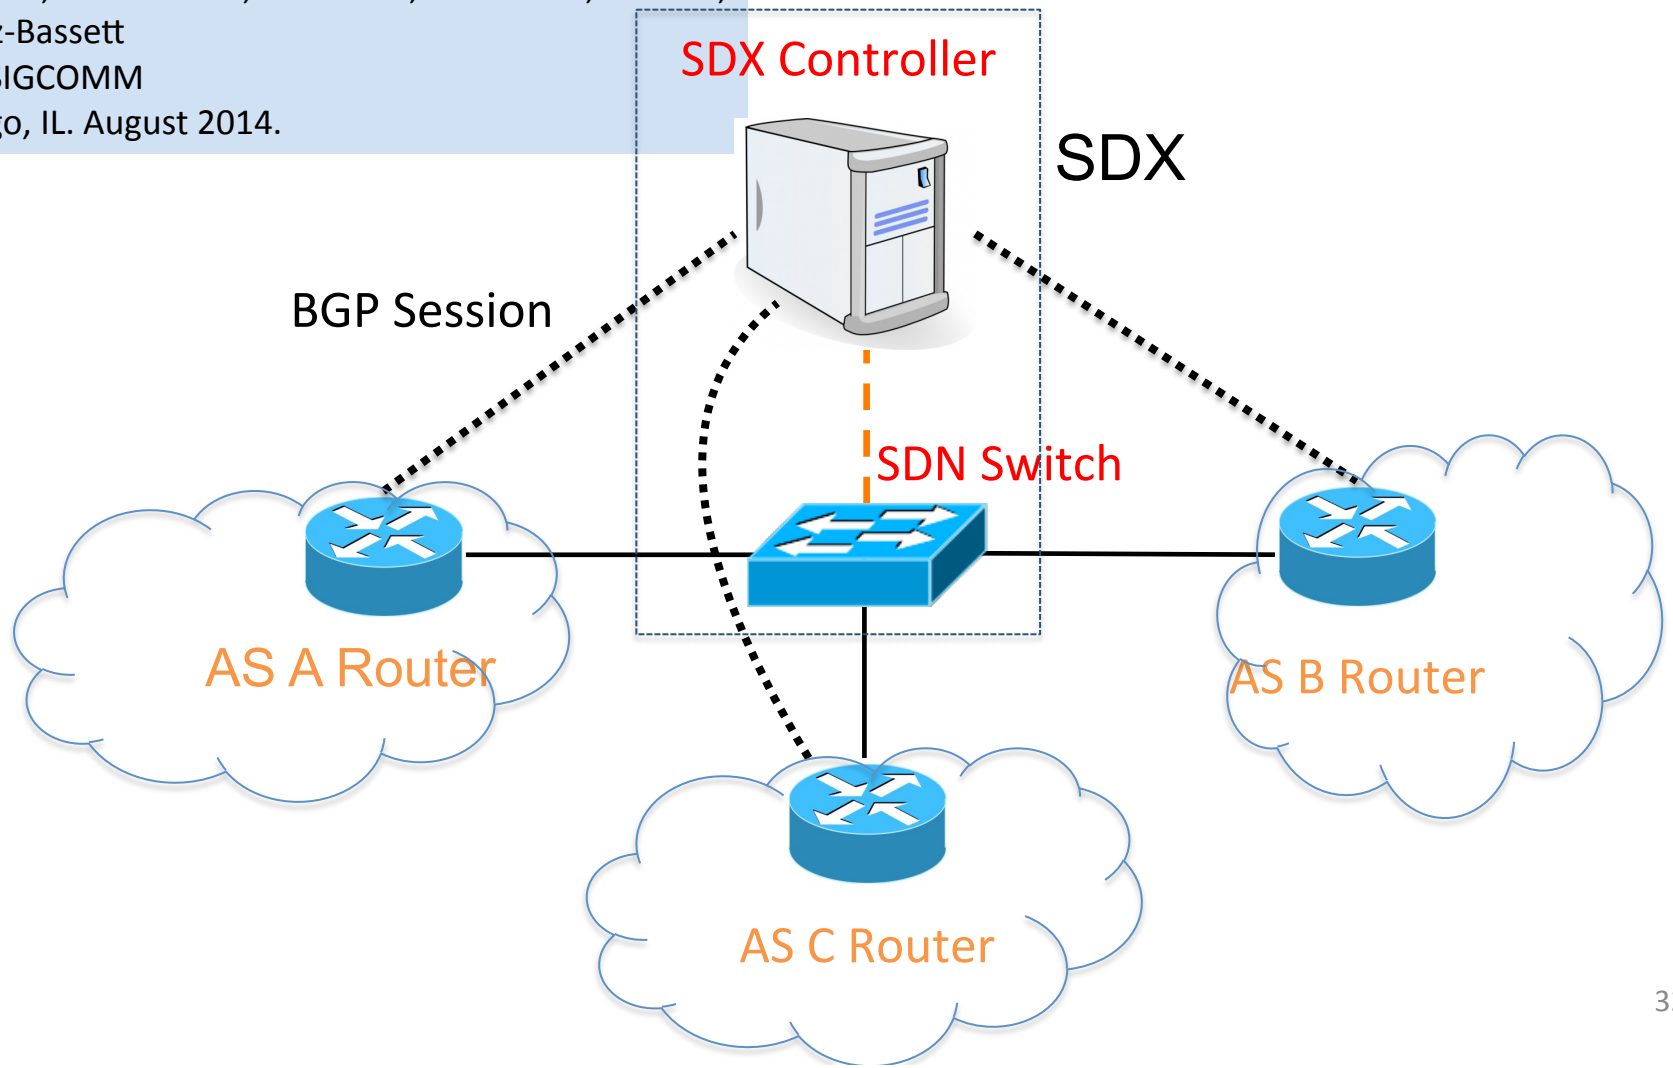

Prevent DDoS Attacks

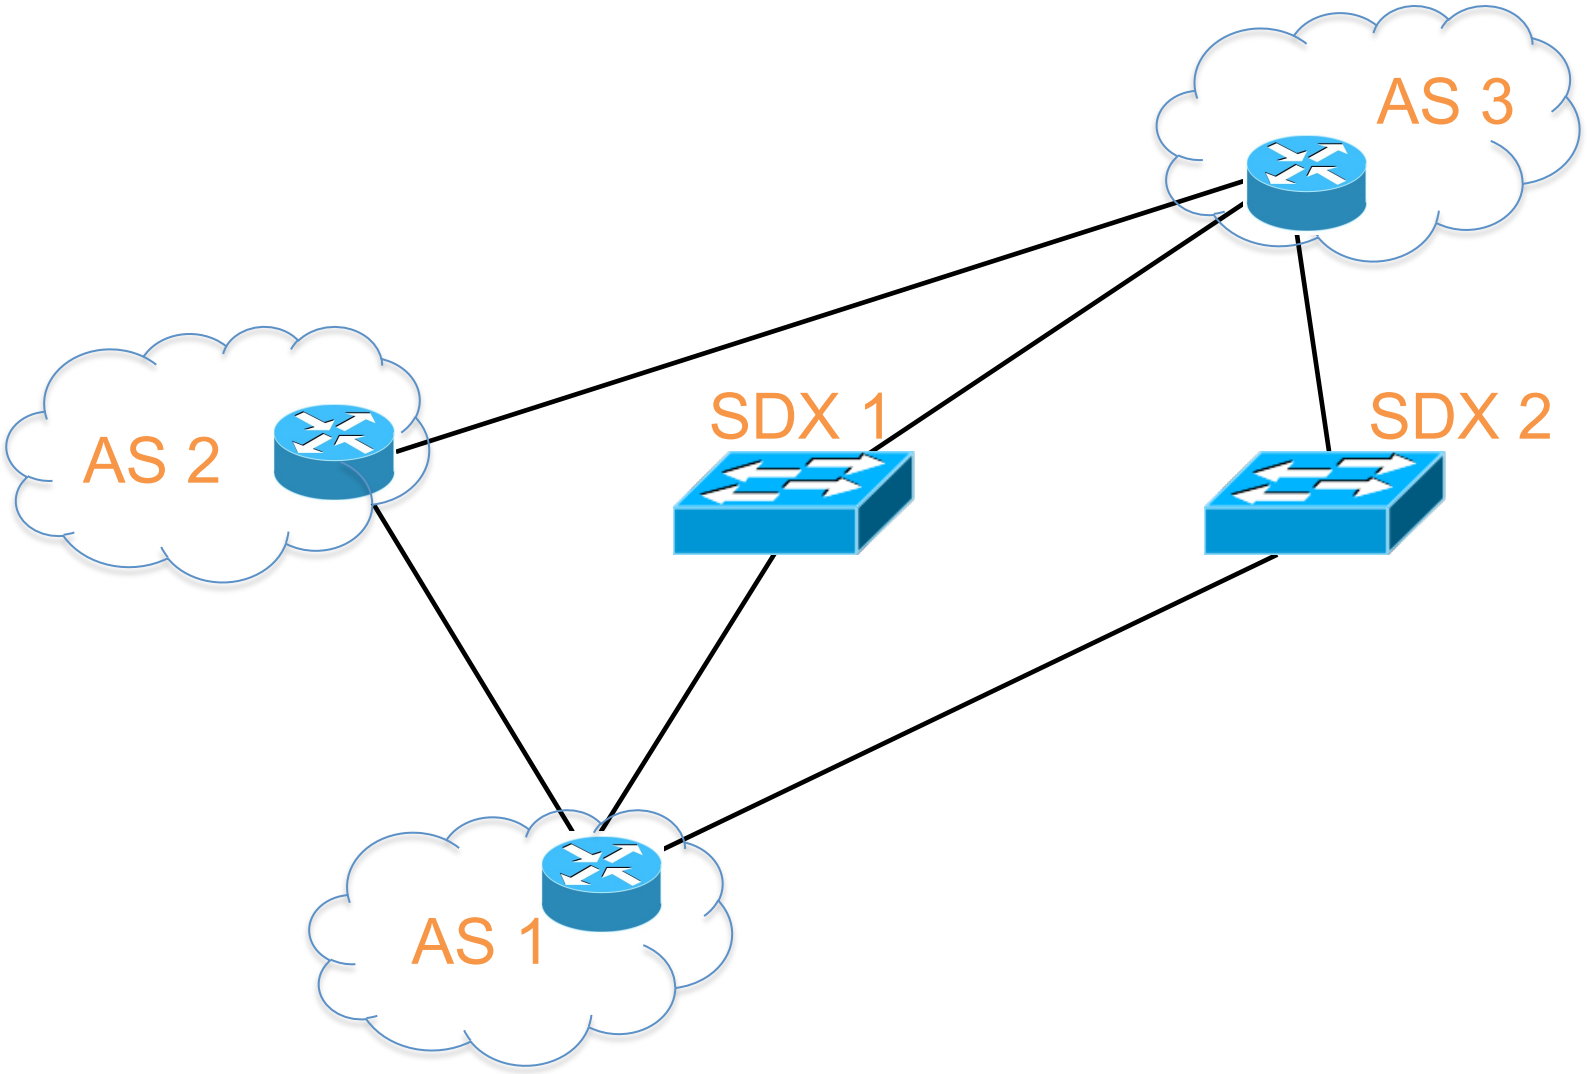

Prevent DDoS Attacks

AS1 under attack originating from AS3

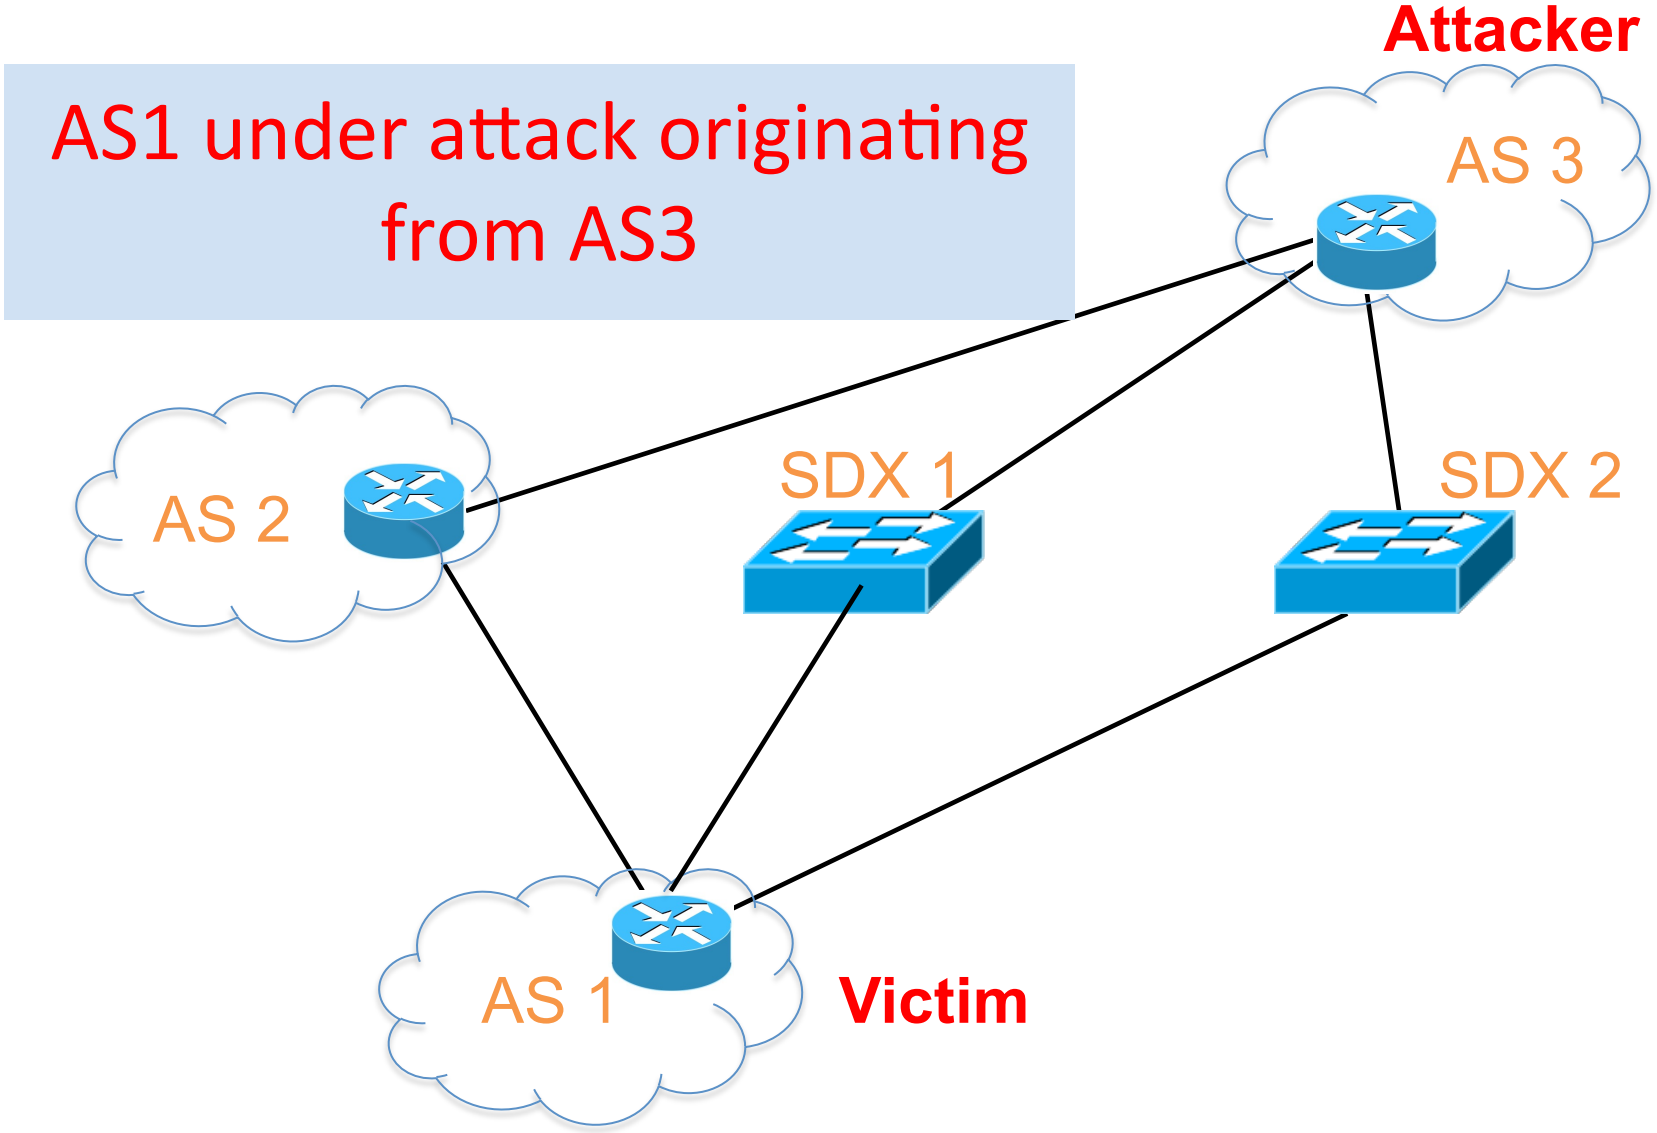

Use Case: Prevent DDoS Attacks

AS1 can remotely block attack traffic at SDX(es)

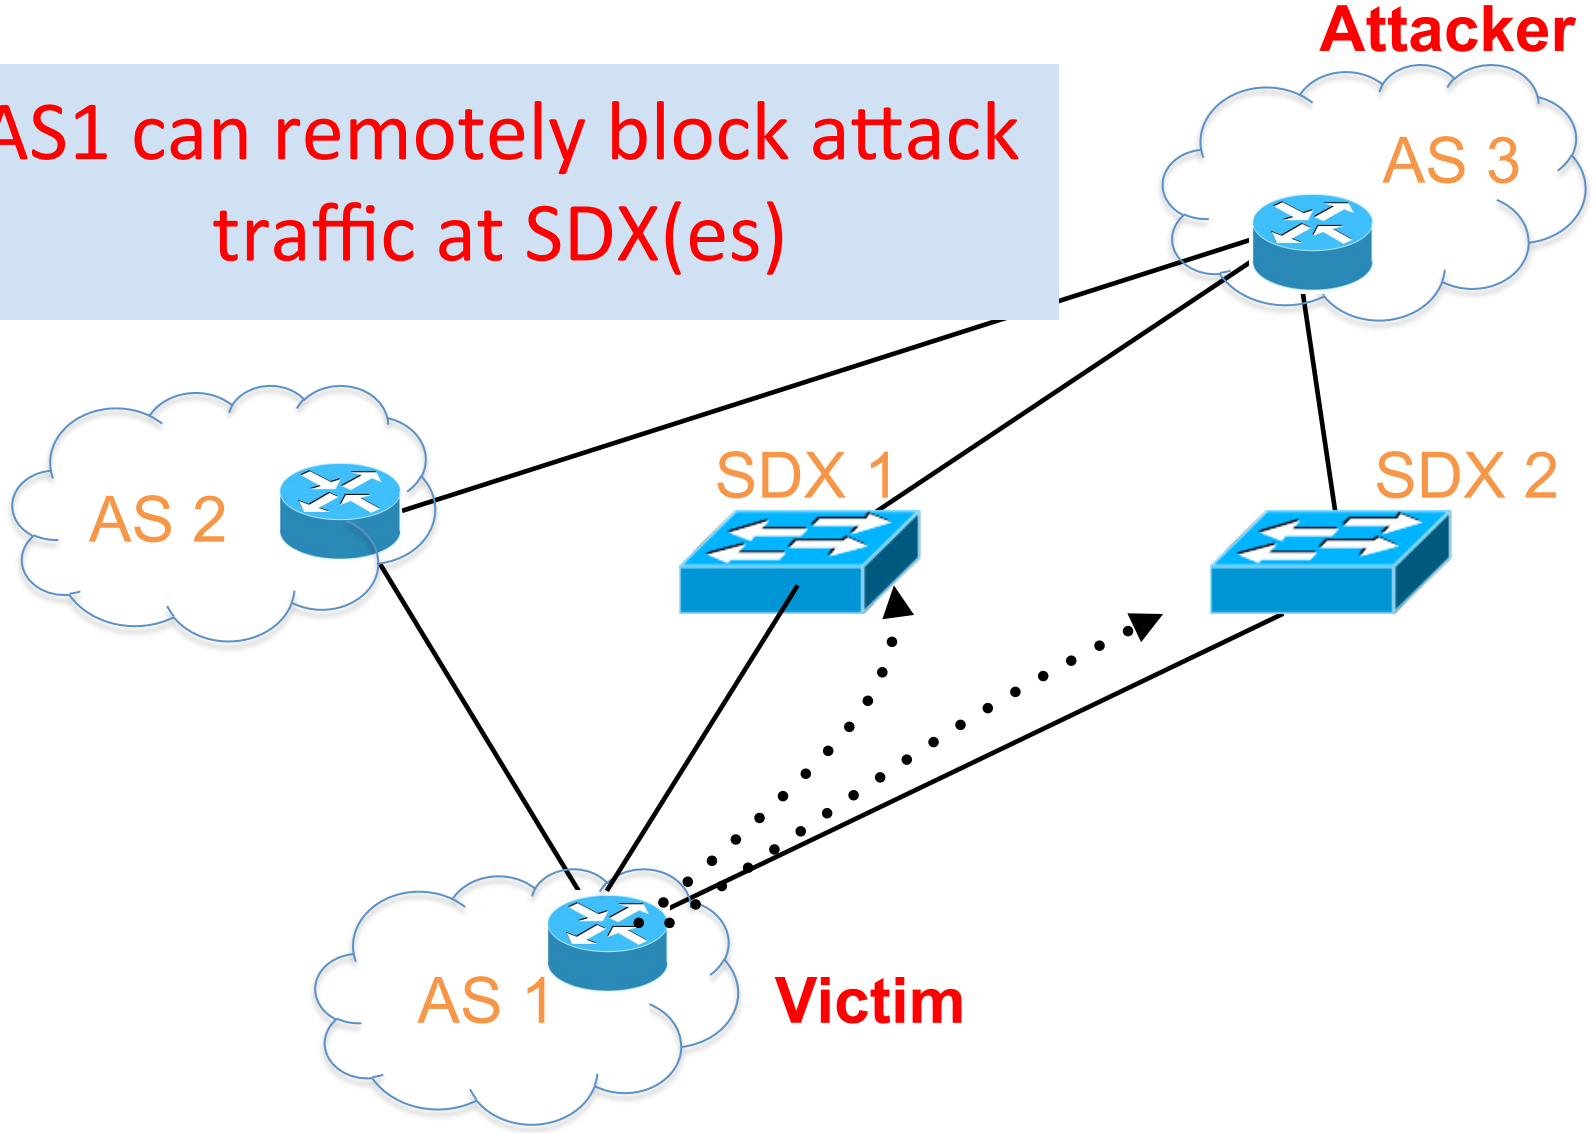

SDX-based DDoS protection vs. Traditional Defenses/Blackholing

- **Remote influence**

  Physical connectivity to SDX not required

- **More specific**

  Drop rules based on multiple header fields, source address, destination address, port number …

- **Coordinated**

  Drop rules can be coordinated across multiple IXPs

Inbound Traffic Control

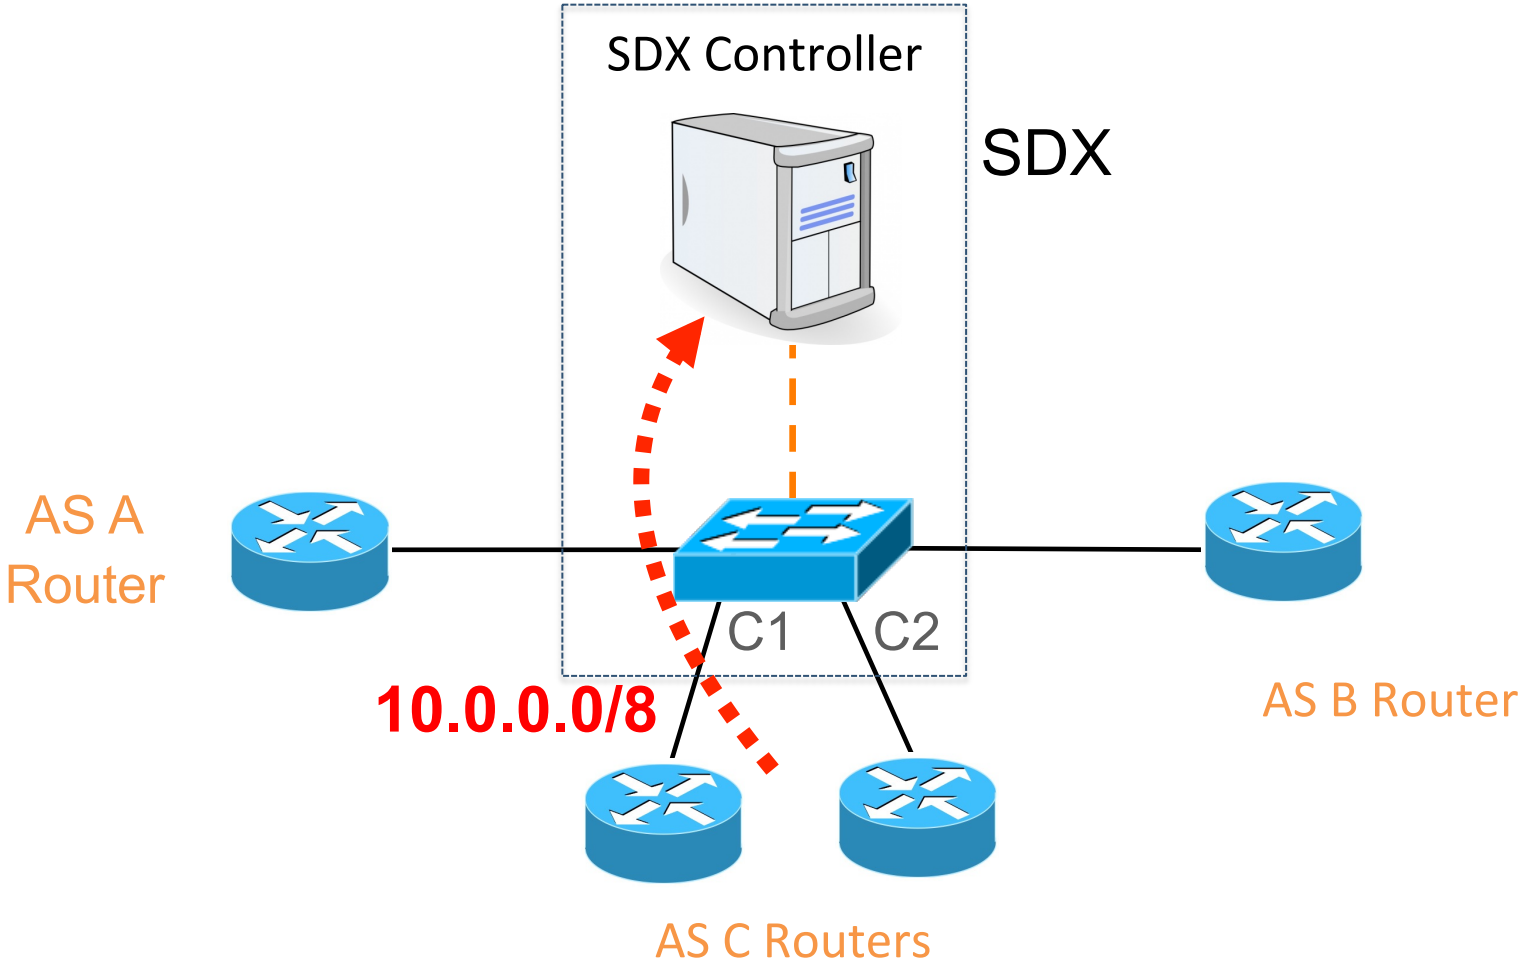

Inbound Traffic Control

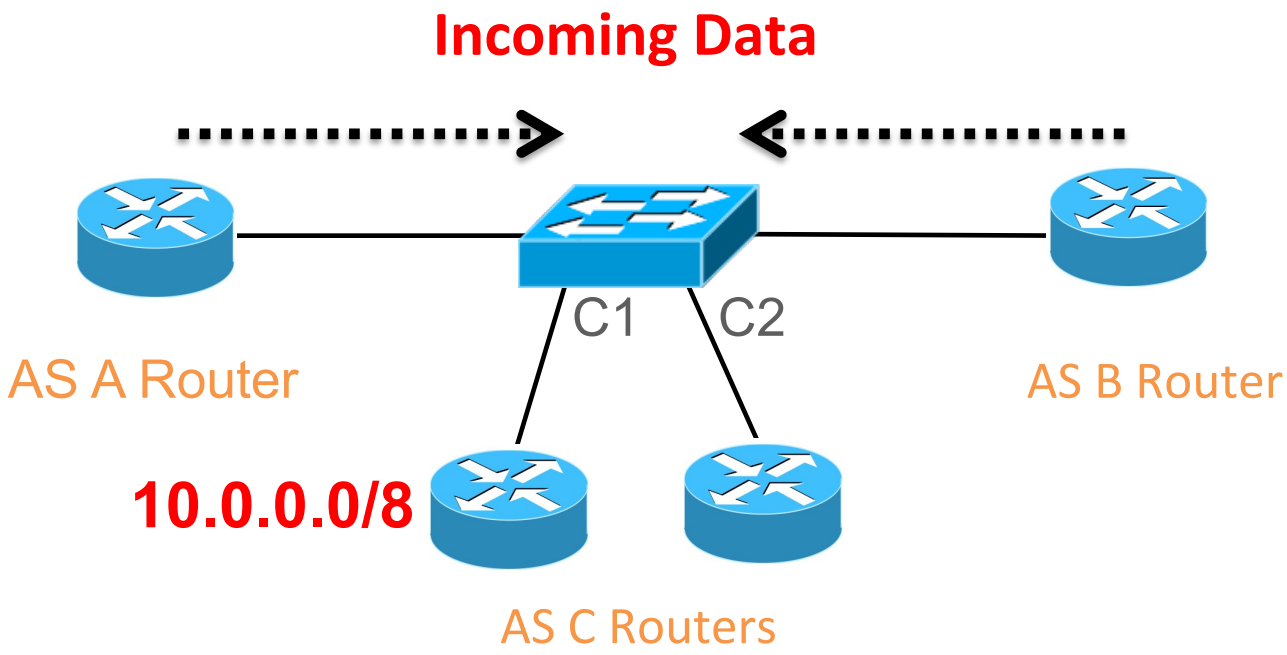

Incoming Traffic	Out Port	Using BGP	Using SDX
dstport = 80	C1		

Inbound Traffic Control

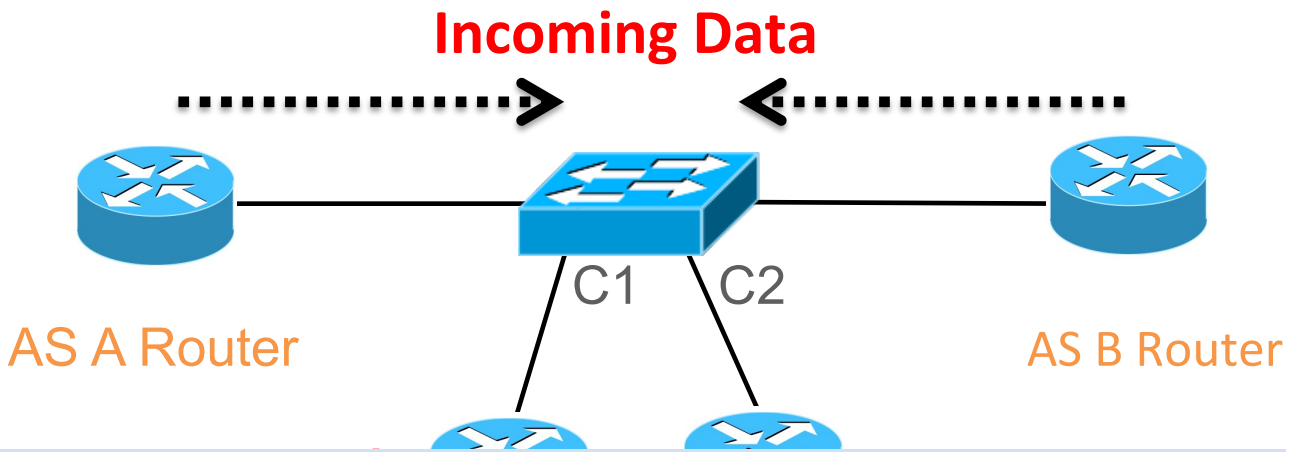

Fine grained policies not possible with BGP

Incoming Traffic	Out Port	Using BGP	Using SDX
dstport = 80	C1	?	

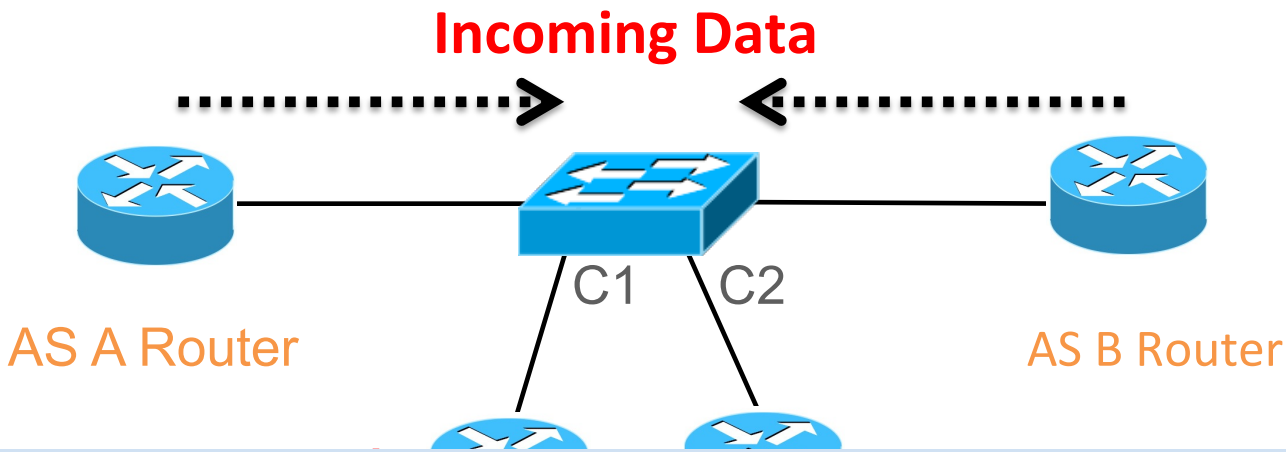

Enables fine-grained traffic engineering policies

Incoming Traffic	Out Port	Using BGP	Using SDX
dstport = 80	C1	?	match(dstport =80)→ fwd(C1)

Building SDX is Challenging

- Programming **abstractions**

  How networks define SDX policies and how are they combined together?

- **Interoperation** with BGP

  How to provide flexibility w/o breaking global routing?

- **Scalability**

  How to handle policies for hundreds of peers, half million address blocks, and matches on multiple header fields?

Building SDX is Challenging

- Programming **abstractions**

  How networks define SDX policies and how are they combined together?

- **Interoperation** with BGP

  How to provide flexibility w/o breaking global routing?

- **Scalability**

  How to handle policies for hundreds of peers, half million prefixes and matches on multiple header fields?

Directly Program the SDX Switch

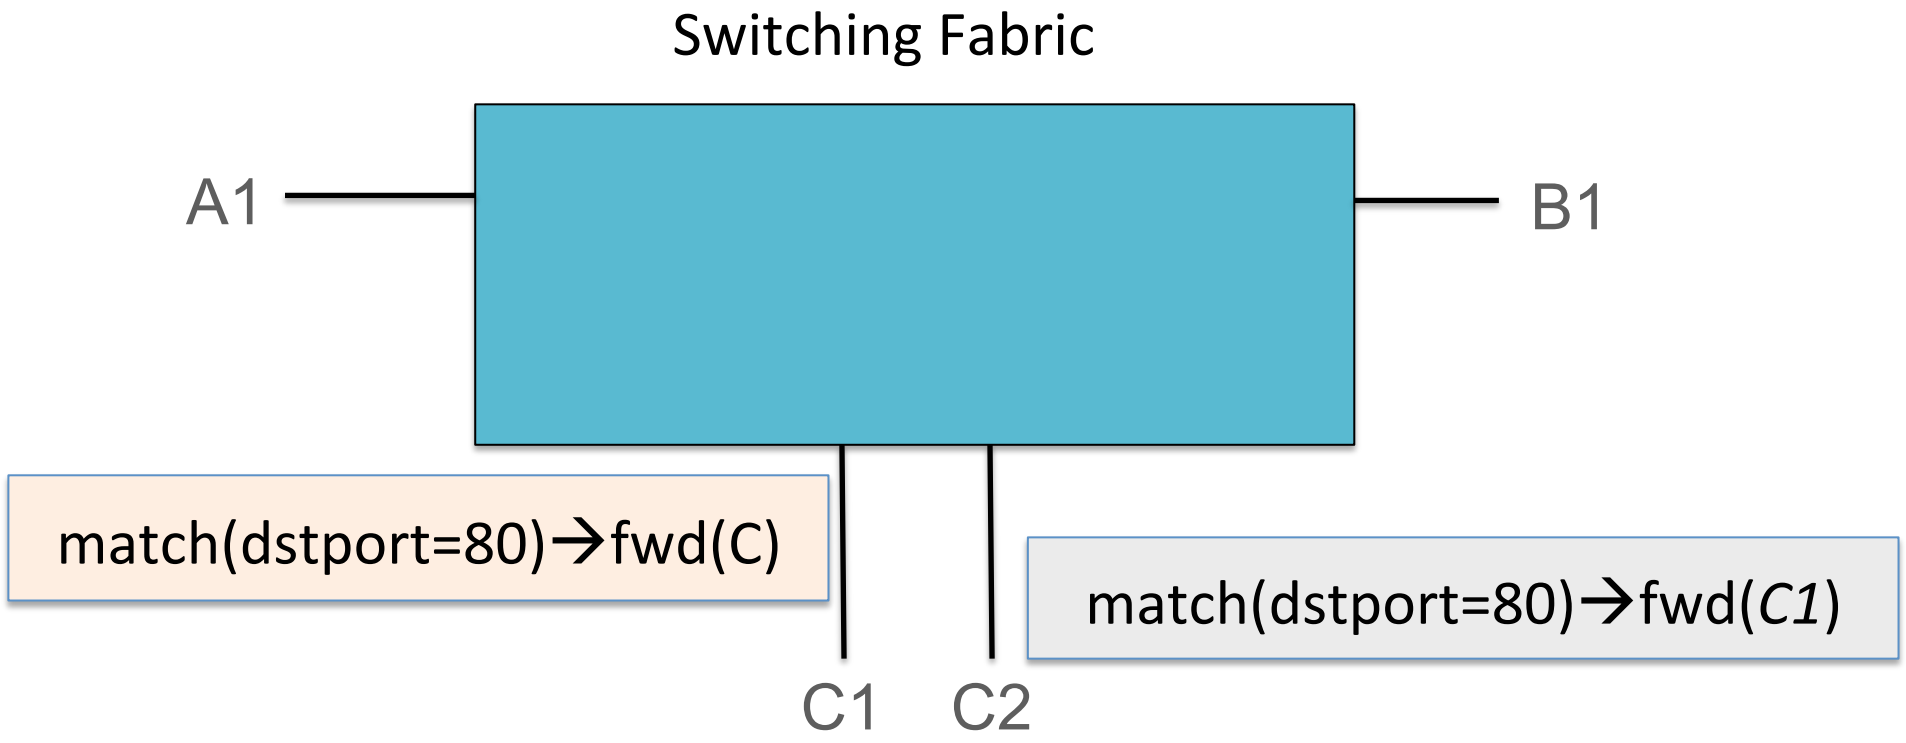

AS A & C directly program the SDX Switch

Virtual Switch Abstraction

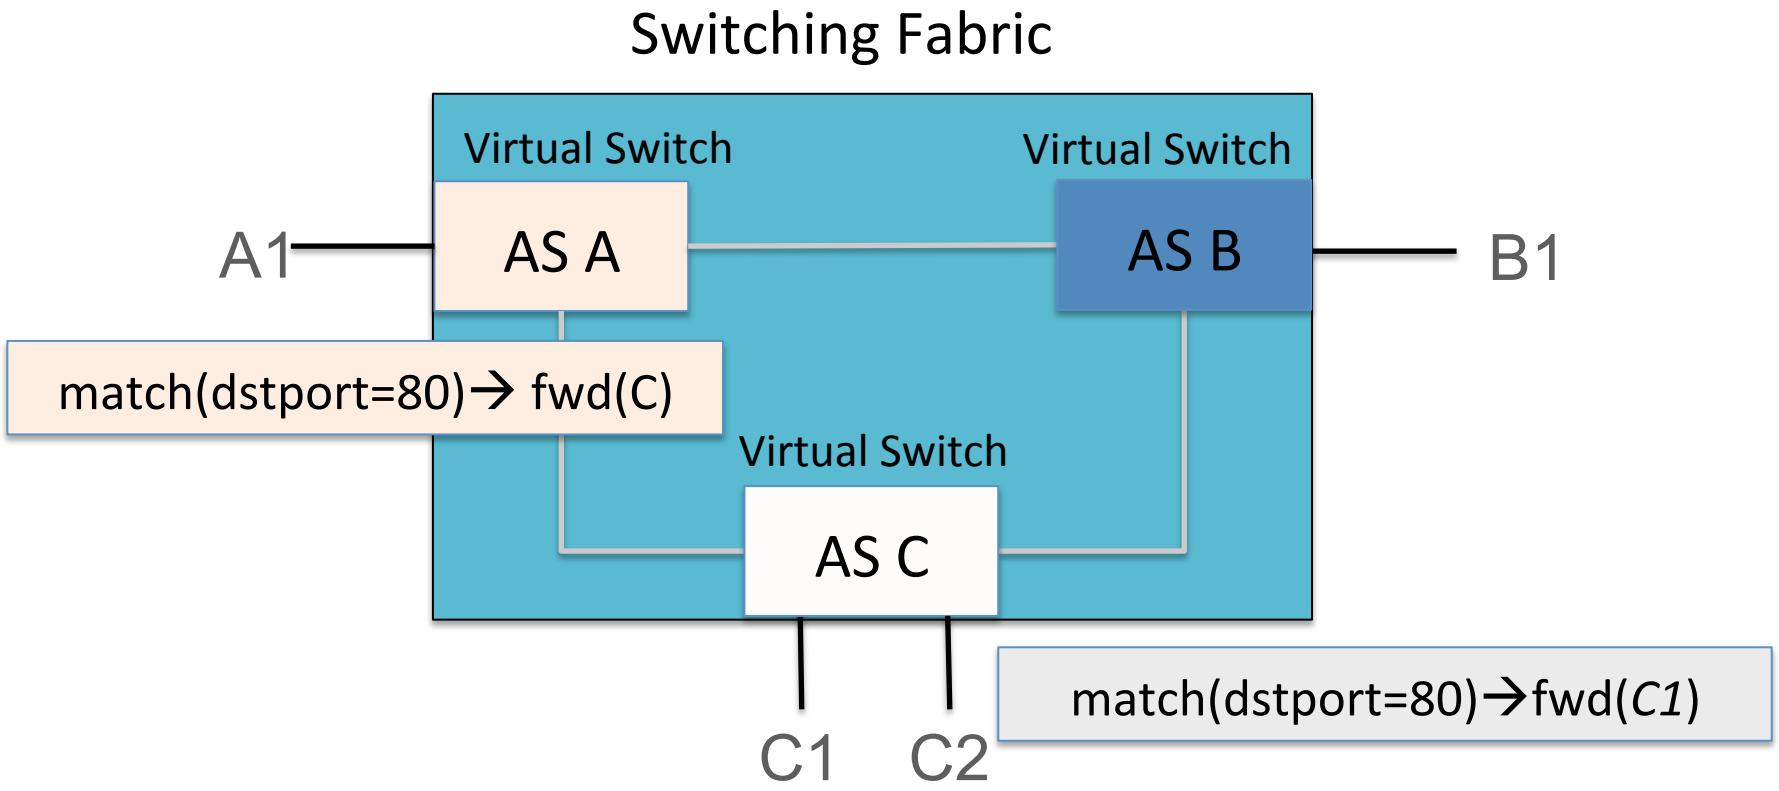

Each AS writes policies for its own virtual switch

Combining Participant's Policies

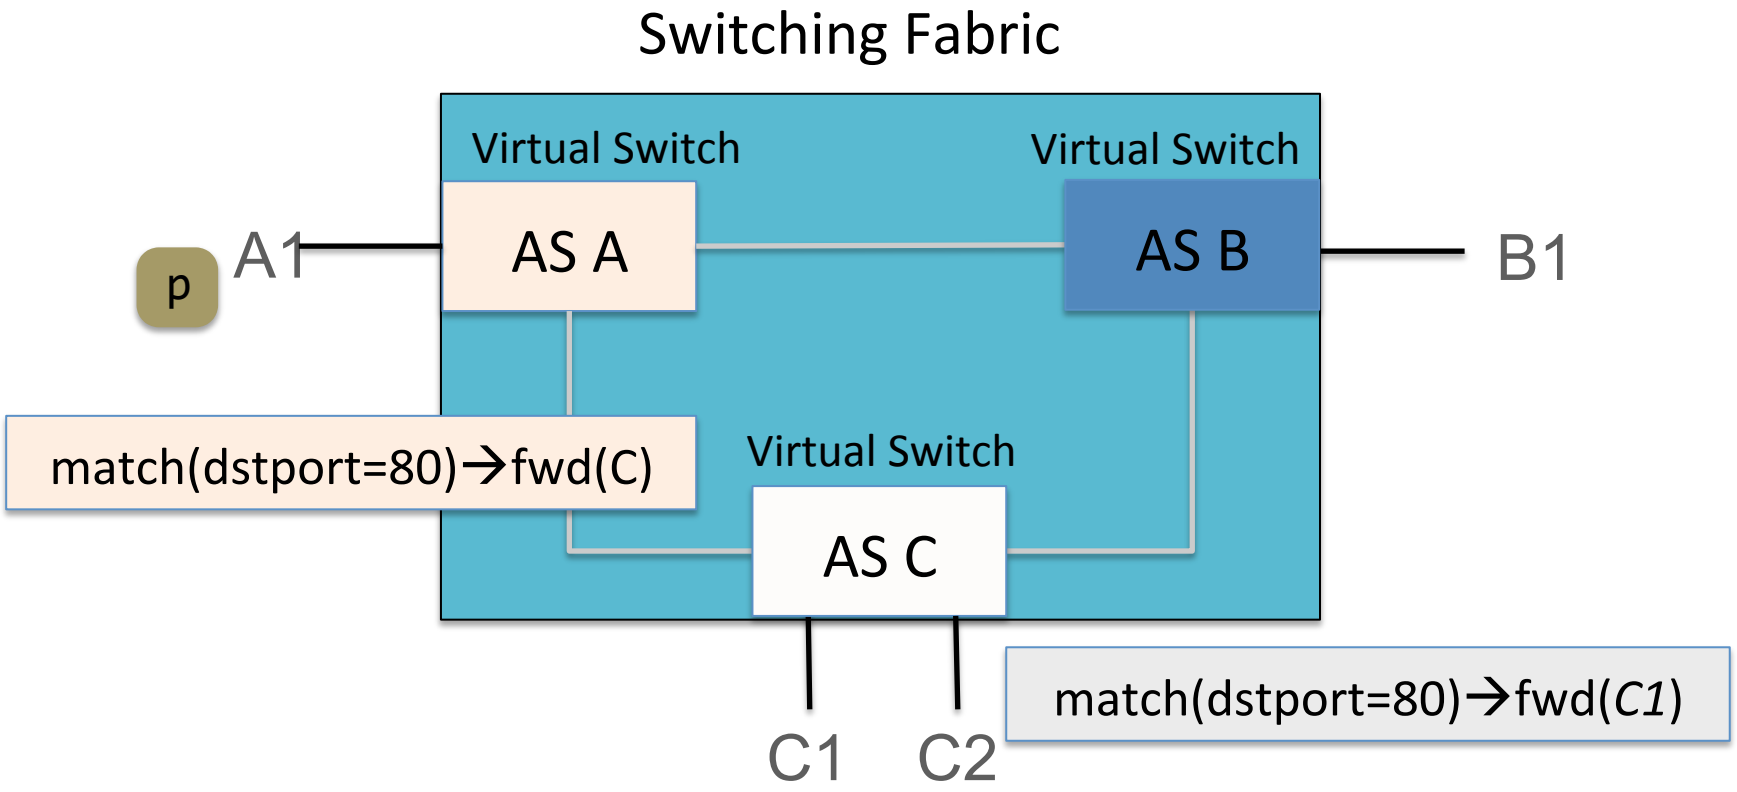

Synthesize: match(inport=A1 & dstport=80) → fwd(C1)

Building SDX is Challenging

- Programming **abstractions**

  How networks define SDX policies and how are they combined together?

- **Interoperation** with BGP

  How to provide flexibility w/o breaking global routing?

- **Scalability**

  How to handle policies for hundreds of peers, half million address blocks, and matches on multiple header fields?

SDX Platform

- Running code with full BGP integration
  - Github available from http://sdx.cs.princeton.edu
- SDX testbeds:
  - Transit Portal for "in the wild" experiments
  - Mininet for controller experiments
- Exploring deployment opportunities
  - Princeton, DOD/IC, GENI, SOX, Internet2, ESnet
  - Regional IXPs in US, Europe, and Africa

SDX: Summary

- The Internet is changing
  - New challenges for content delivery
  - Increasing importance of IXPs
- SDN can let providers innovate
  - New capabilities and abstractions
- Next steps
  - Operational deployments
  - Additional SDX applications
  - Distributed exchange points

Getting Involved

- UCT Honors Project (To Be Announced)
  - Mobile Measurement
  - Performance, Throttling, Censorship
- Collaborations with UCT Networking Lab
- Summer Internship at Princeton

Come for an Internship!

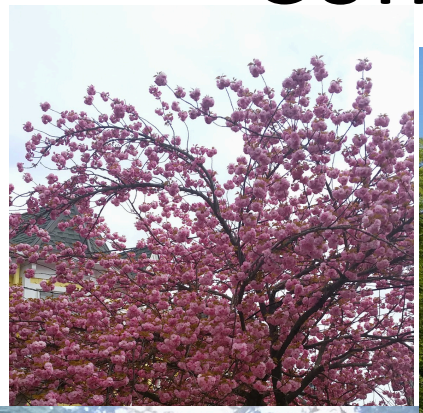

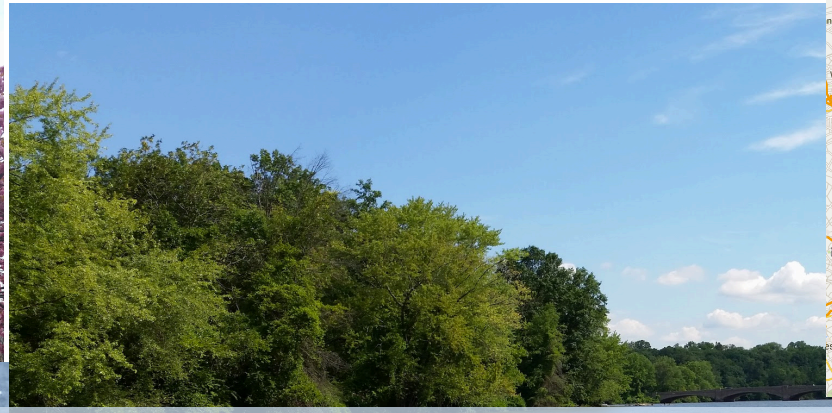

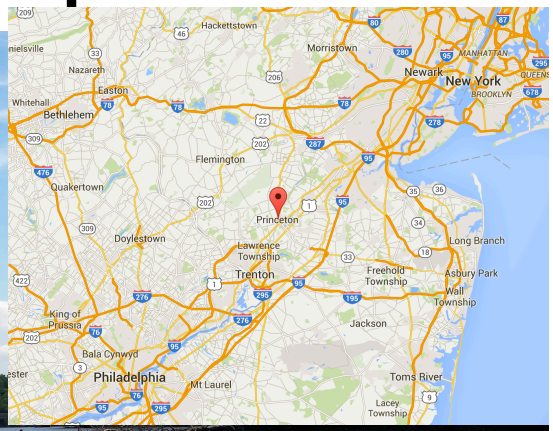

- **Princeton International Student Internship Program (ISIP) starts Summer 2016!**

- The program will occur only during the summer months, when on-campus housing is available.
- Internship applicants must be currently enrolled undergraduates at institutions outside of the United States.
- Get in touch: feamster@cs.princeton.edu