

Monitoring and Managing QoE in the Wide Area: Programmable Routers to the Rescue?

Nick Feamster

Princeton University

(with apologies to Arpit Gupta, Marco Canini,
Laurent Vanbever, Jen Rexford, and others)

<http://sdx.cs.princeton.edu>

We Have Neither Visibility Nor Control at Interconnection Points


- **No control**
 - Routing **only on destination IP prefixes**
(No customization of routes by application, sender)
 - Can only influence **immediate neighbors**
(No ability to affect path selection remotely)
 - **Indirect** control over data-plane forwarding
(Indirect mechanisms to influence path selection)
- **No visibility**
 - Monitoring is limited to byte counters and NetFlow
 - No precise packet timings along path segments
 - No active probes along path segments

SDX: SDN at Internet Exchanges

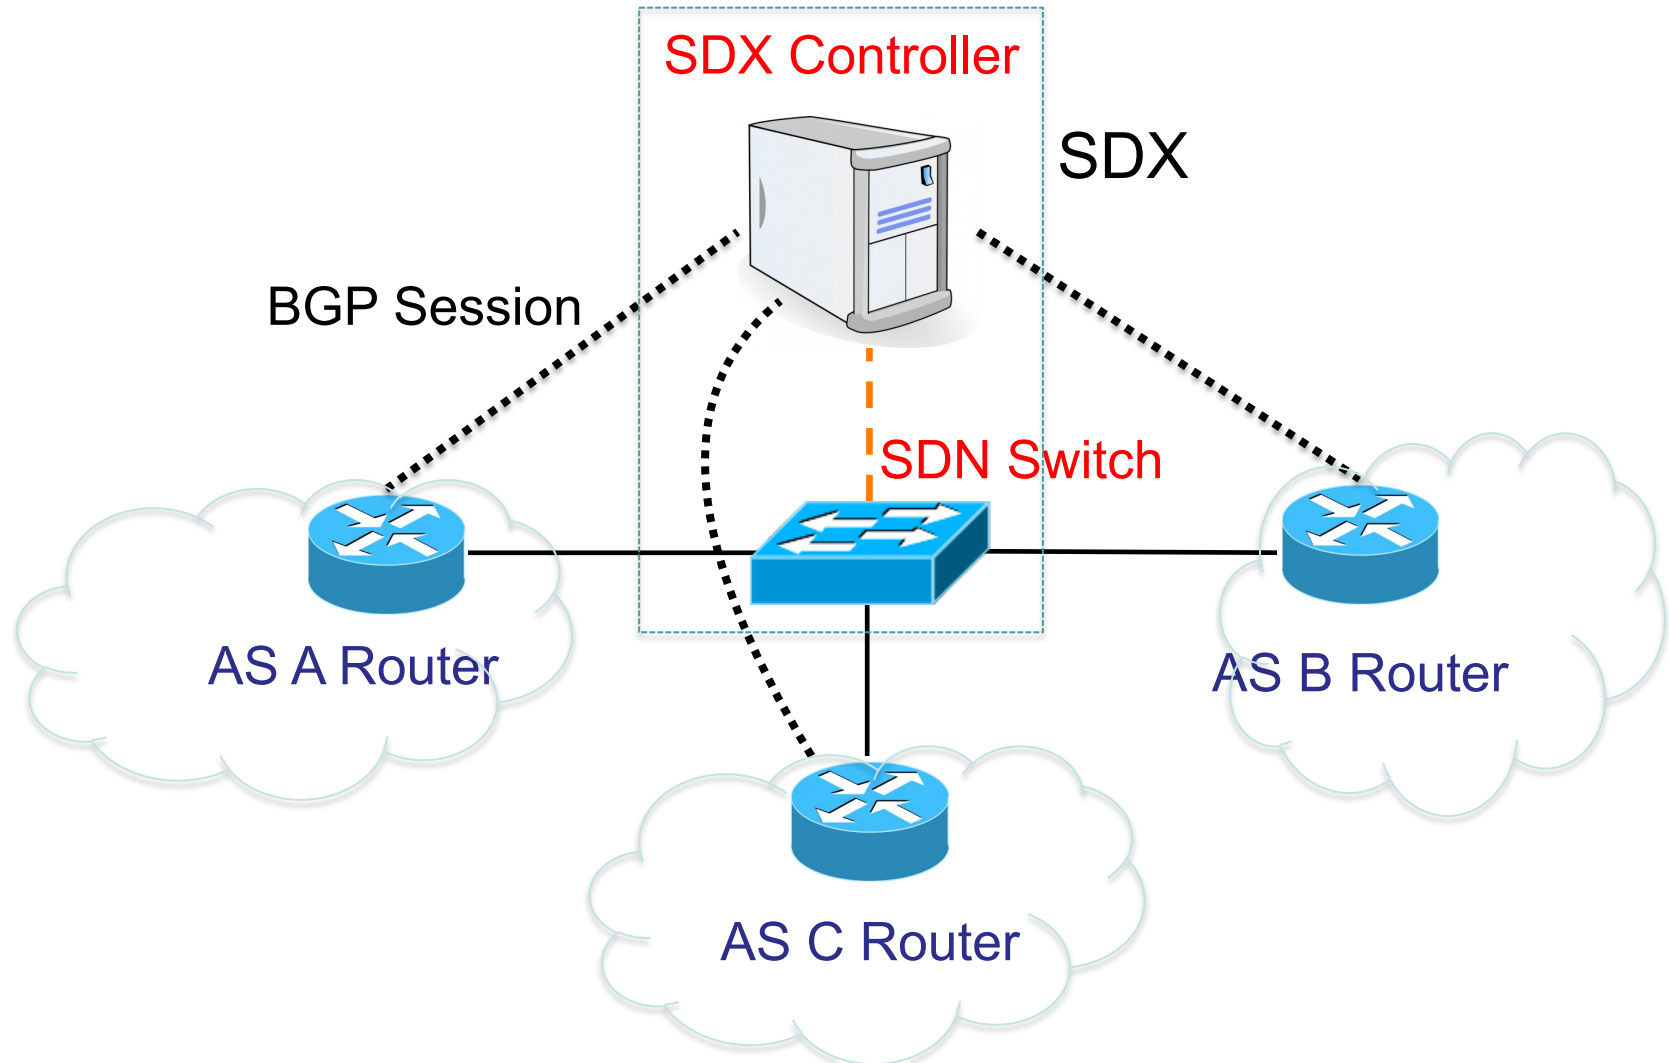
- **Leverage:** SDN deployment even at single IXP can yield benefits for tens to hundreds of ISPs
- **Innovation hotbed:** Incentives to innovate as IXPs on front line of peering disputes
- **Growing in numbers:** ~100 new IXPs established in past three years*

*<https://prefix.pch.net/applications/ixpdir/summary/growth/>

Conventional IXPs



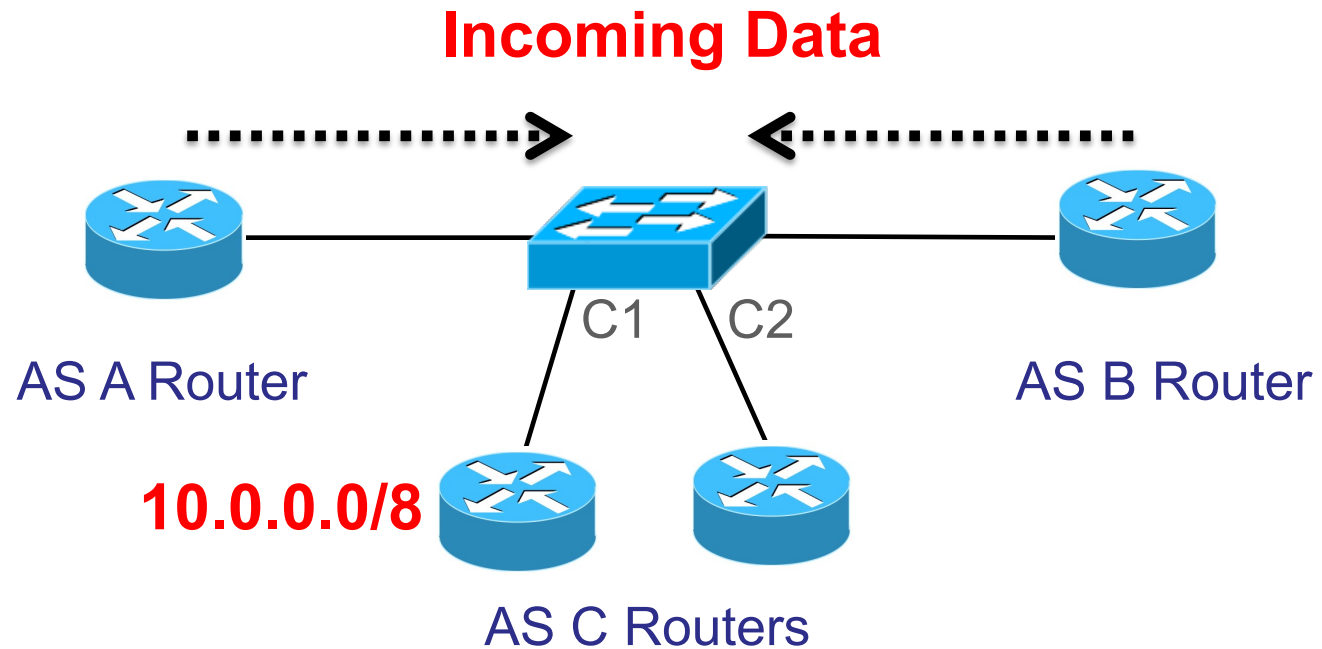
SDX = SDN + IXP



Improving Control


- Forwarding on **multiple header fields**
(not just destination IP prefixes)
- Ability to **control entire network** with a single software program (not just immediate neighbors)
- **Direct control** over data-plane forwarding
(not indirect control via control-plane arcana)

Inbound Traffic Control



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

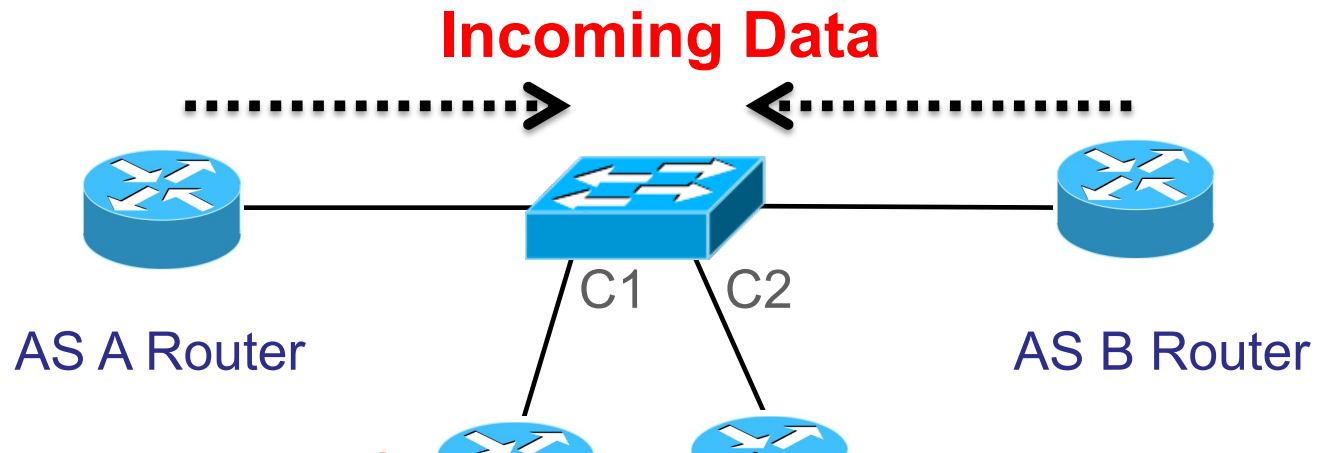
Inbound Traffic Control



Fine grained policies not possible with BGP

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

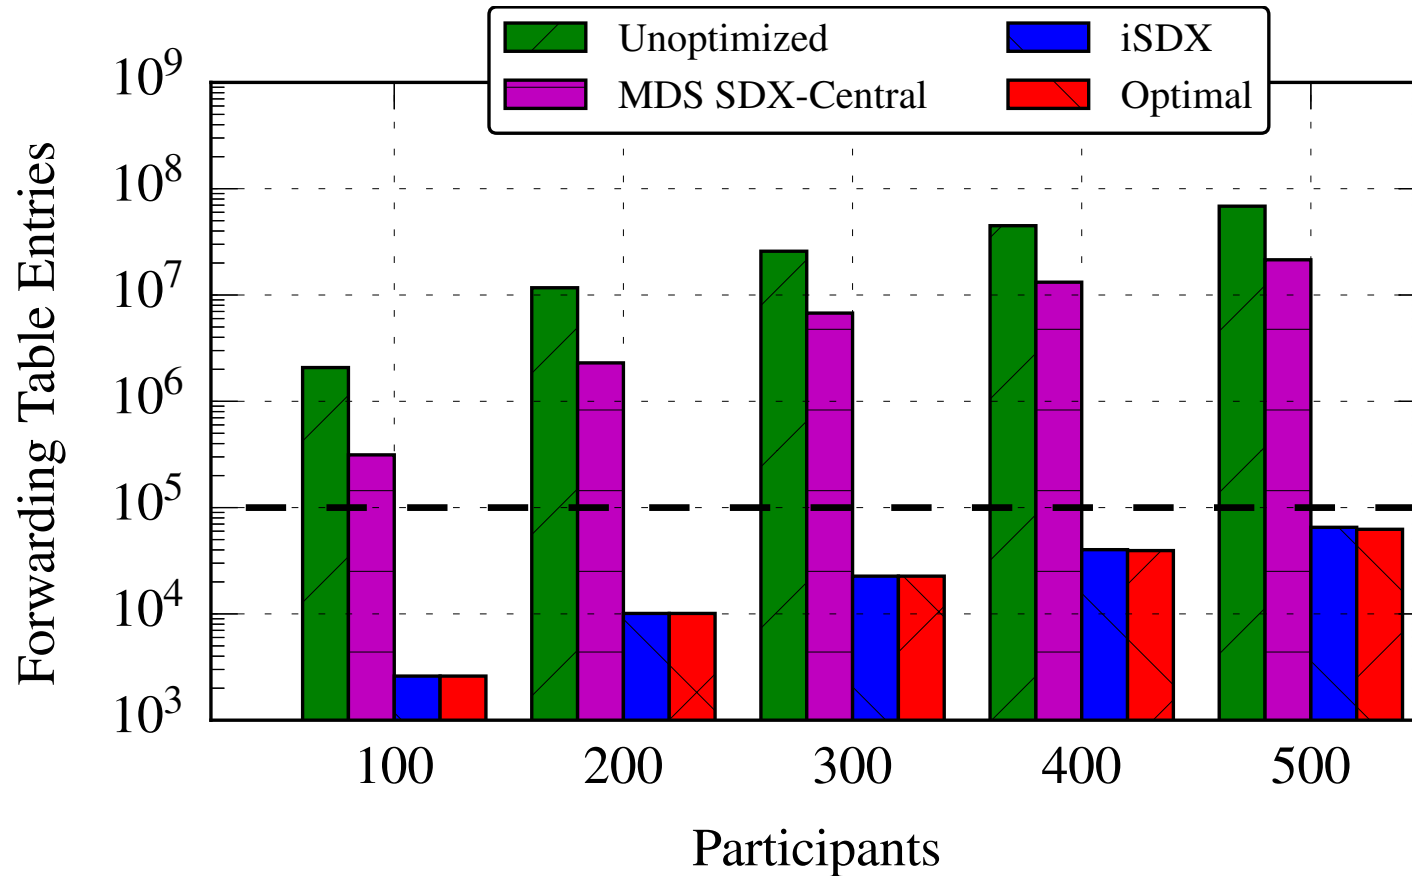
Inbound Traffic Control



Enables fine-grained traffic engineering policies

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

We Can Do This at IXP-Scale



**BGP Routes and Updates for DE-CIX
in a Commodity Hardware Switch.**

Better Visibility: In-Band Network Telemetry

- Network elements collect, report, modify state **in-real time** as data packets go through switch.
- **Writing state into packets**
 - (switch, in, out) tuples
 - Latency
 - Link Utilization
- **Dynamic counters** based on different hash buckets
- **Dynamic actions** based on (say) counter thresholds
 - Dynamic rule creation
 - Reactive probing

What is Possible

- In-band traceroute/topology discovery
- Per-hop latency/loss/utilization recording
- Active probing based on counter thresholds
- Dynamic redirection of traffic flows
- ...

SDX Software Release

- Running code with full BGP integration
- Used in Coursera course
(10s of thousands of students)
- Ongoing Deployment Efforts:
 - Inter-agency exchange (NSA)
 - Large IXP (DE-CIX)

<http://sdx.cs.princeton.edu>

Compared to SNMP

- Faster
- Less CPU Overhead
- In-band (exporting data, not polling)
- Can pre-filter
- Packet-level granularity
- Publish/subscribe