UNIVERSITY OF WATERLOO

CS 856: Programmable Networks

Lecture 11: Software-Defined WAN, Reconfigurable Optical Networks, & Wrap up

Mina Tahmasbi Arashloo

Winter 2023

Research Discovery Days

From the Math Innovation Office (MIO):

"The Math Innovation Office is hosting our inaugural Math & Computing Research Discovery Days, taking place from April 17-18, 2023, which will bring an impressive lineup of keynote speakers, research demonstrations, networking opportunities, and more. This is a great chance for you to learn about leading-edge math research, connect with industry experts and successful entrepreneurs, and network with the math community."

Logistics

- Presentations next week
  - We'll have pastries! 😁
- Final project report due on **Monday, April 10th.**

Final Project Report

- Use the **USENIX paper templates**
- **6 pages**, excluding references.
- Structured as a research paper
  - Abstract
  - Introduction
  - Motivation
  - Related Work
  - Design and implementation
  - Evaluation

Final Project Report

- Depending on your project, your report's organization and section titles could be a little different from the last slide.
- You can look at workshop papers for inspiration:
  - HotNets 2022 (or other years)
  - Short papers at SOSR 2022 (or other years)
  - Workshops at SIGCOMM 2022 (or other years)

Wide Area Networks

- Carry Tbs of data per second
- Internet providers
  - e.g., AT&T
- Private WAN's connecting "business locations"
  - e.g., Microsoft and Google WAN connecting their data centers.

https://en.vcenter.ir/network/wide-area-network-wan/

Wide Area Networks

Very different from data center networks

- Much larger geographical span
- RTTs can be 100s of milliseconds
- Non-symmetric topologies

Finding the "right" paths is crucial to good performance!

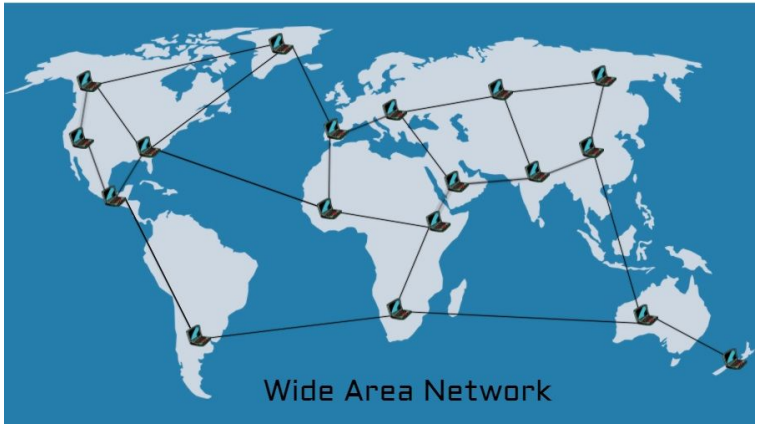

https://en.vcenter.ir/network/wide-area-network-wan/

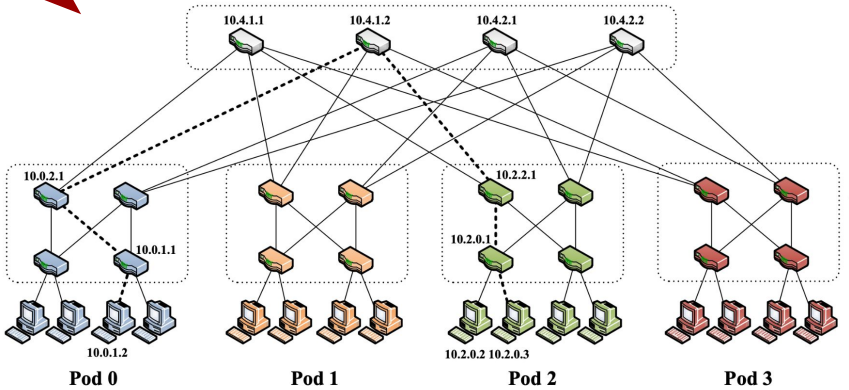

"A Scalable, Commodity Data Center Network Architecture", SIGCOMM'08

Finding the "right" paths in the WAN

Known as the traffic engineering (TE) problem

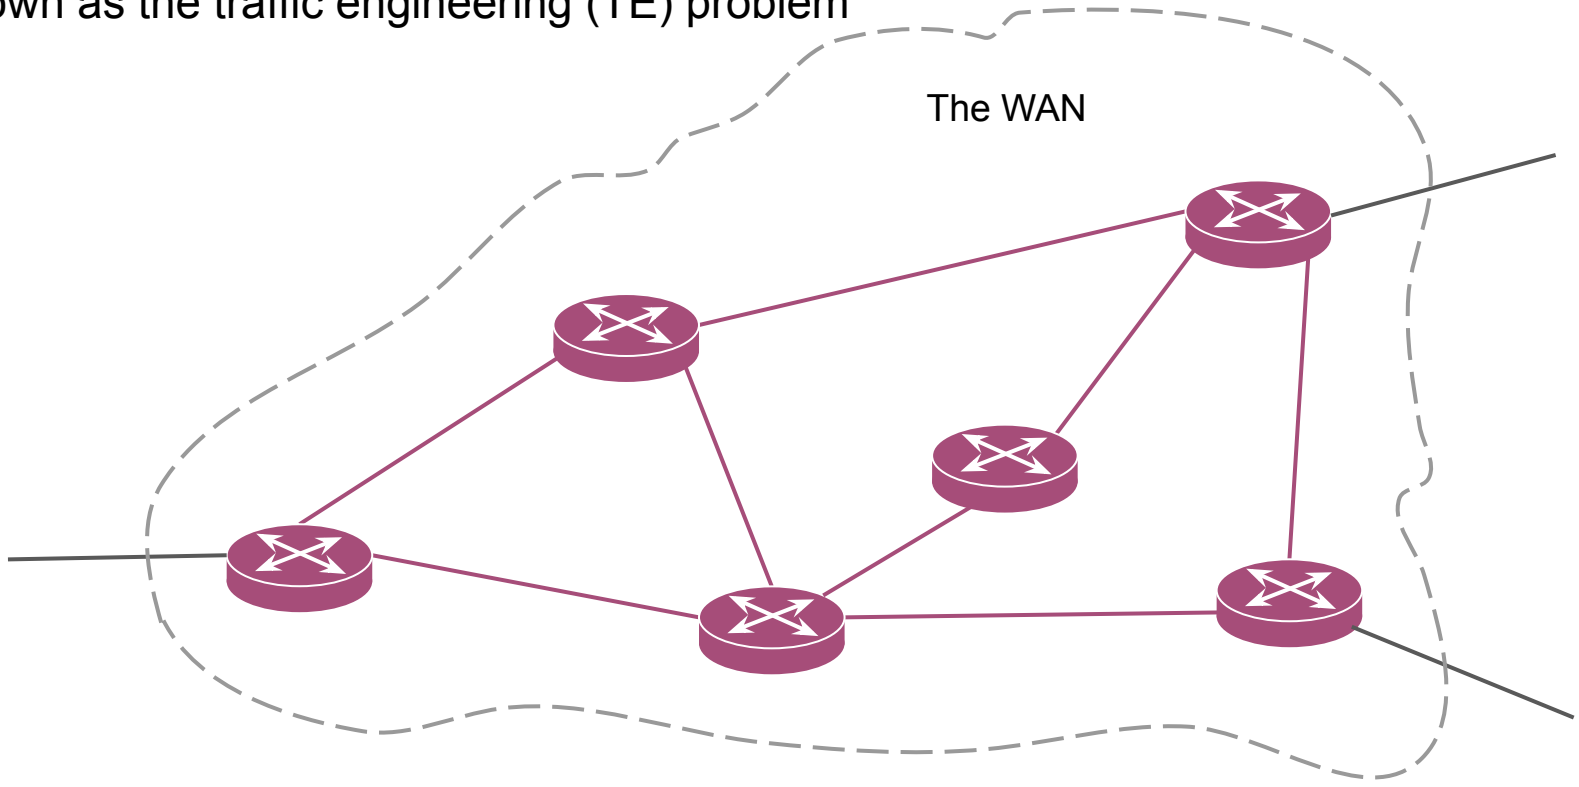

Finding the "right" paths in the WAN

Traffic enters and leaves the network at the "edge" switches/routers

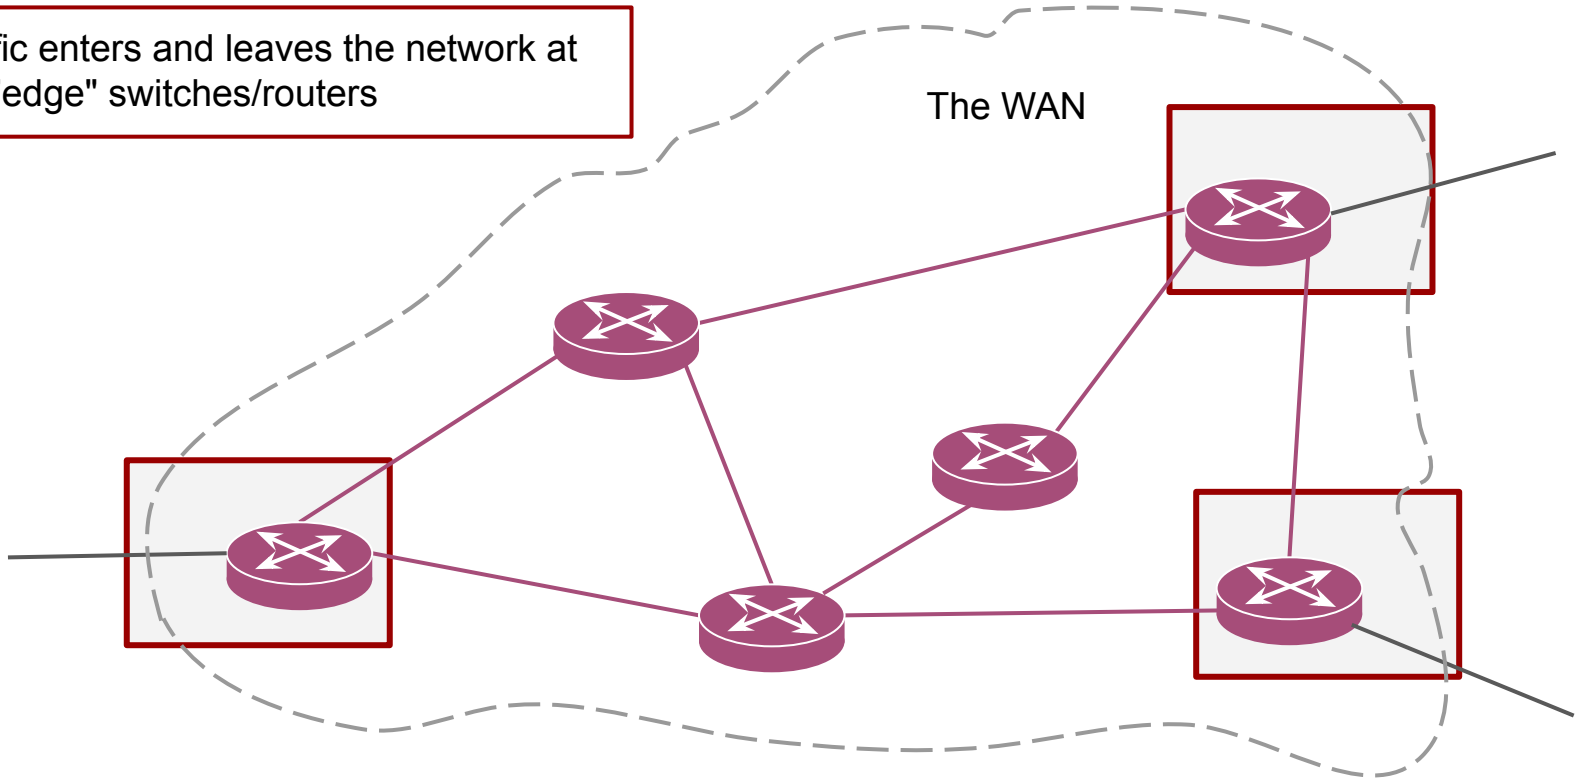

Finding the "right" paths in the WAN

Suppose we want traffic going to 10.0.0.0/24 that enter from ingress switch A to exit from B.

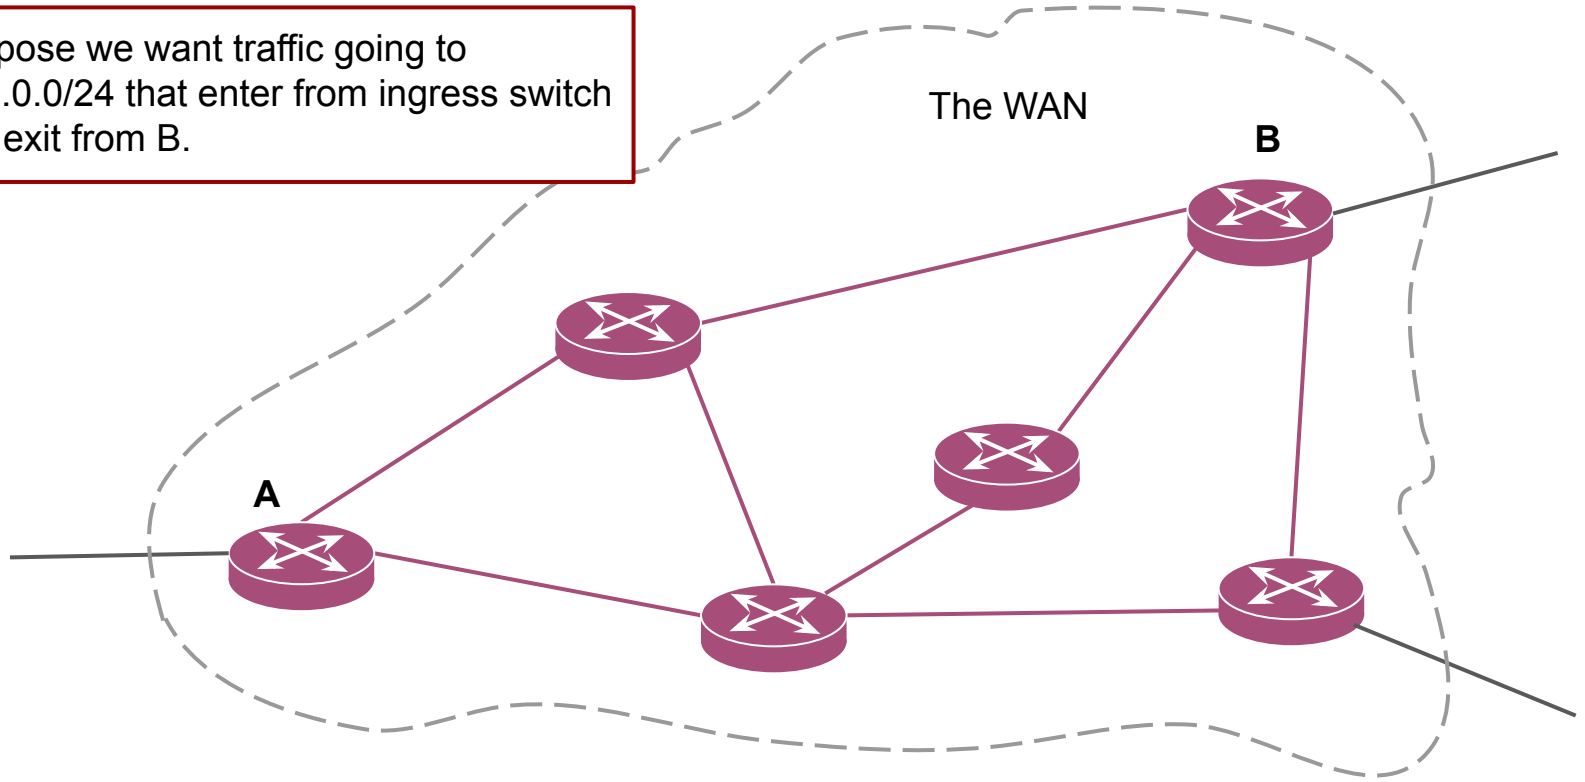

Finding the "right" paths in the WAN

There are many paths between A and B!

A can reserve some of these end-to-end paths.

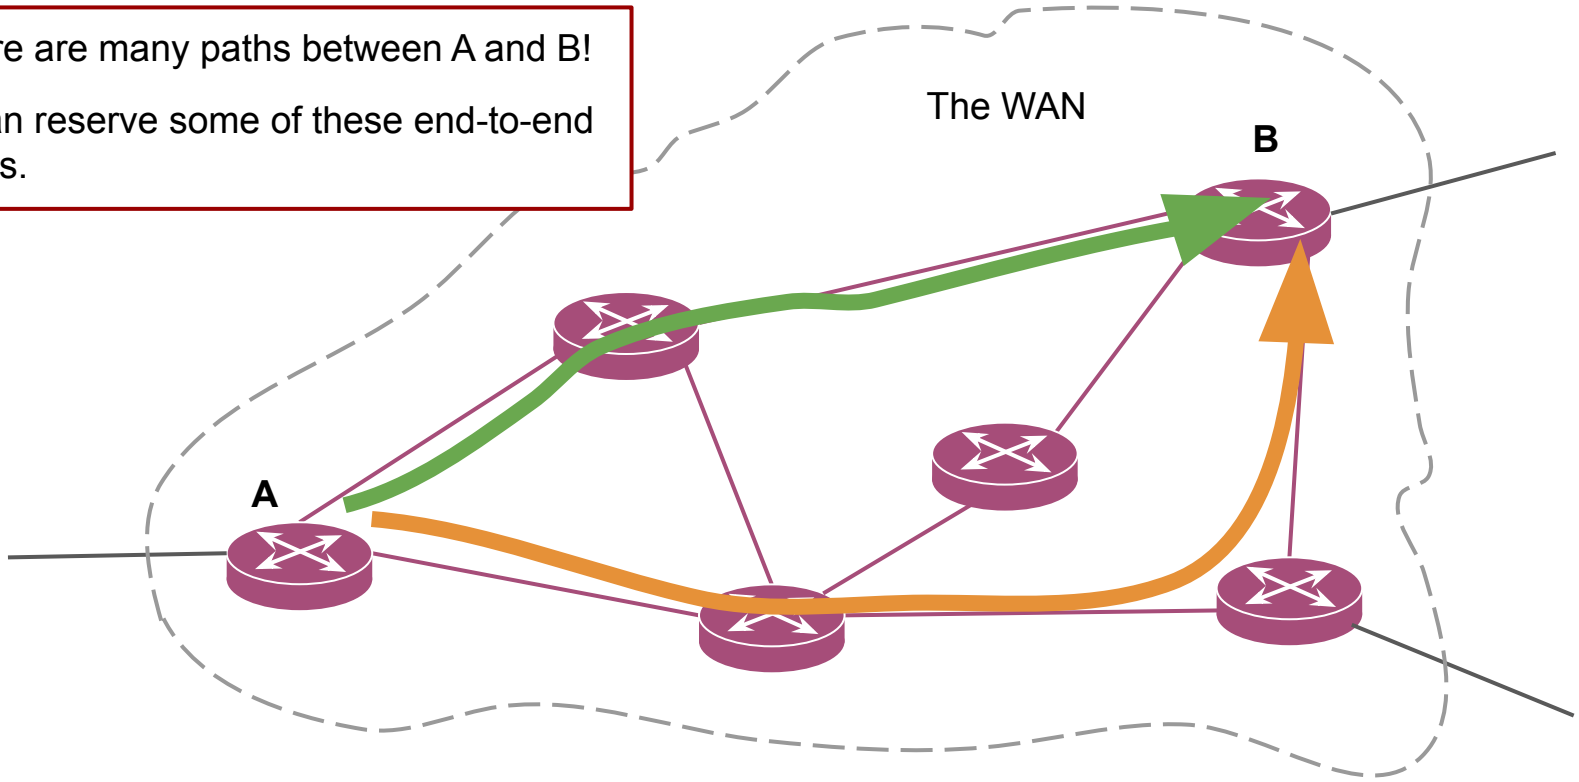

Finding the "right" paths in the WAN

When traffic going to 10.0.0.0/24 enters the network, A can send the traffic along these paths

and be sure the other nodes know how to forward it and have enough resources to do it.

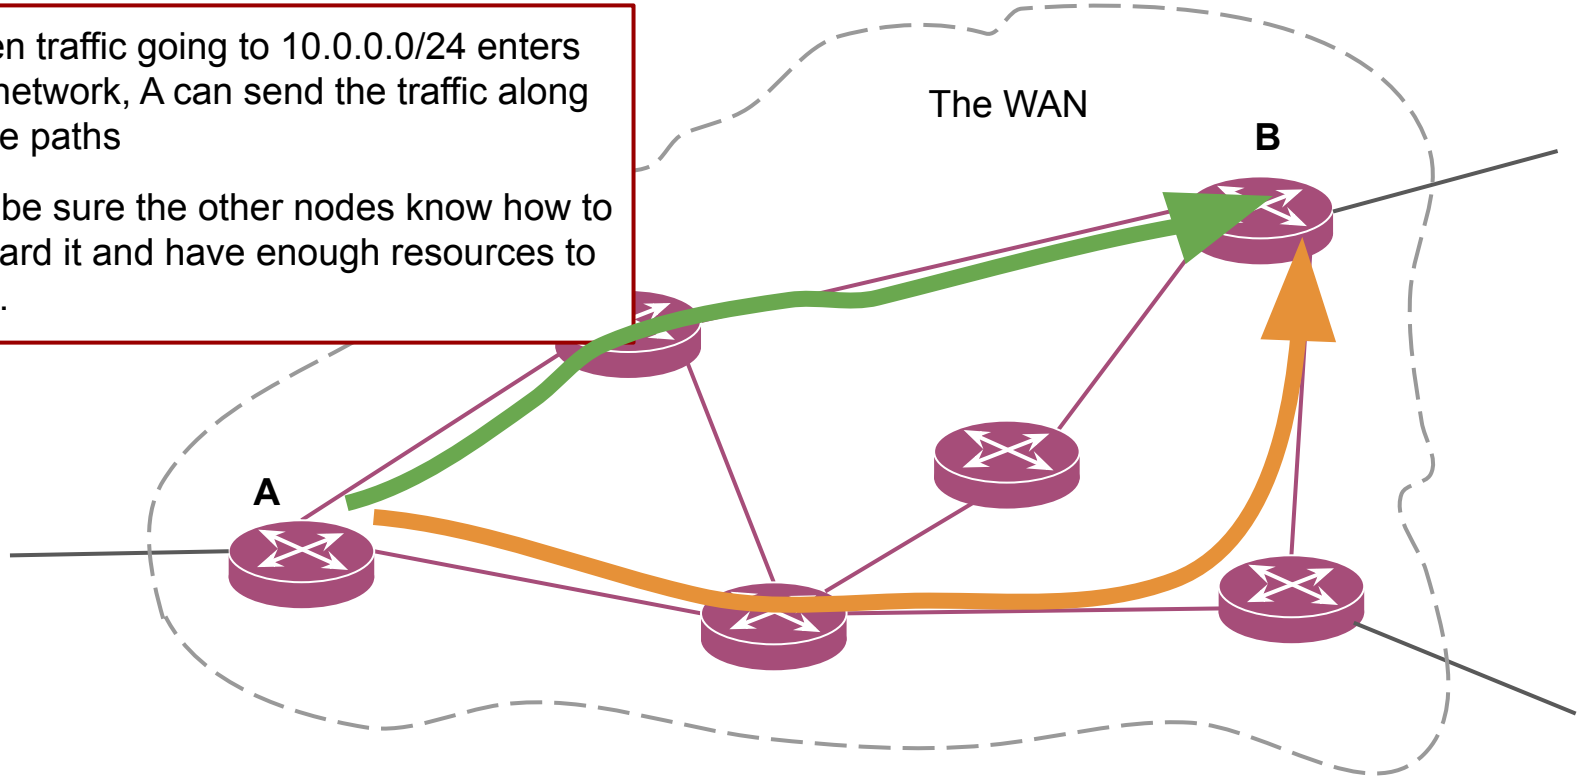

Finding the "right" paths in the WAN

How are paths decided and reserved?

- Routers communicate about the status of their links (e.g., remaining bandwidth, load, etc.)
- Ingress nodes *locally* decides which paths are best for each subset of traffic.
- Ingress node sends a message along the path and asks the nodes to reserve it (signaling).
- If that path is reserved, a message comes back in the opposite direction to setup the "forwarding entries" for the path (label distribution).

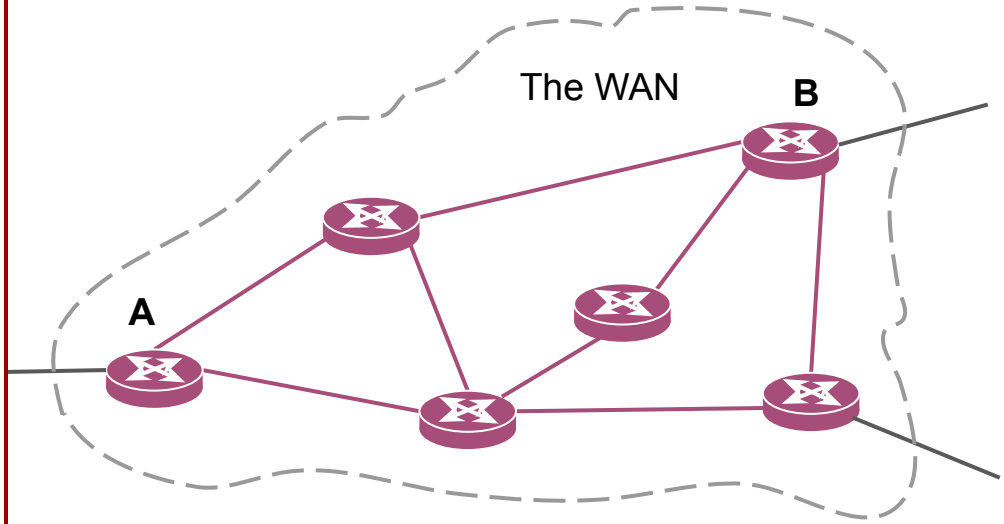

Finding the "right" paths in the WAN

How are paths decided and reserved?

- Routers communicate about the status of their links (e.g., remaining bandwidth, load, etc.)
- Ingress nodes *locally* decides which paths are best for each subset of traffic.
- Ingress node sends a message along the path and asks the nodes to reserve it (signaling).
- If that path is reserved, a message comes back in the opposite direction to setup the "forwarding entries" for the path (label distribution).

This is roughly what happens in RSVP-TE to setup paths for MPLS.

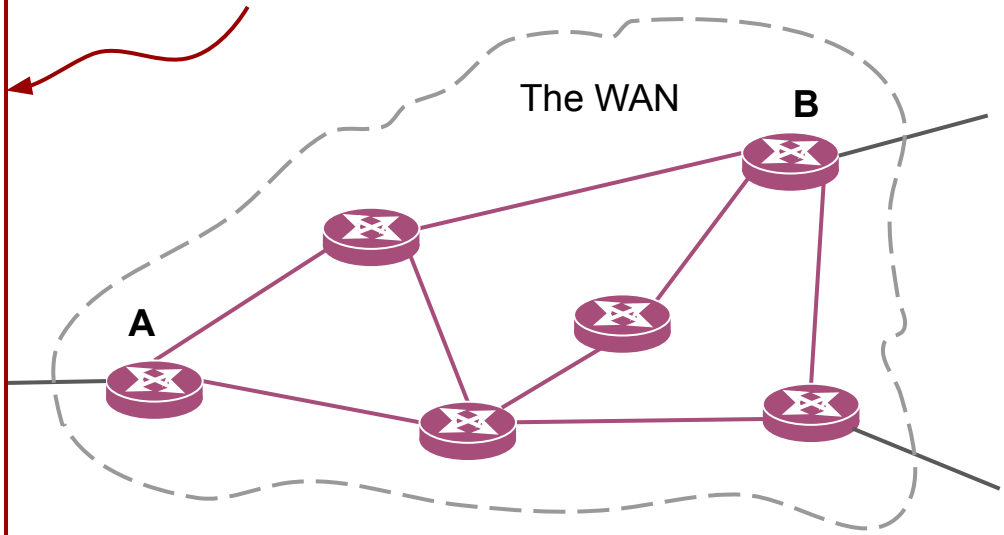

Finding the "right" paths in the WAN

How are paths decided and reserved?

- Routers communicate about the status of their links (e.g., remaining bandwidth, load, etc.)
- Ingress nodes *locally* decides which paths are best for each subset of traffic.
- Ingress node sends a message along the path and asks the nodes to reserve it (signaling).
- If that path is reserved, a message comes back in the opposite direction to setup the "forwarding entries" for the path (label distribution).

Can lead to sub-optimal decisions and ineffective use of resources.

Wouldn't it be great to have a global view and make global decisions? :)

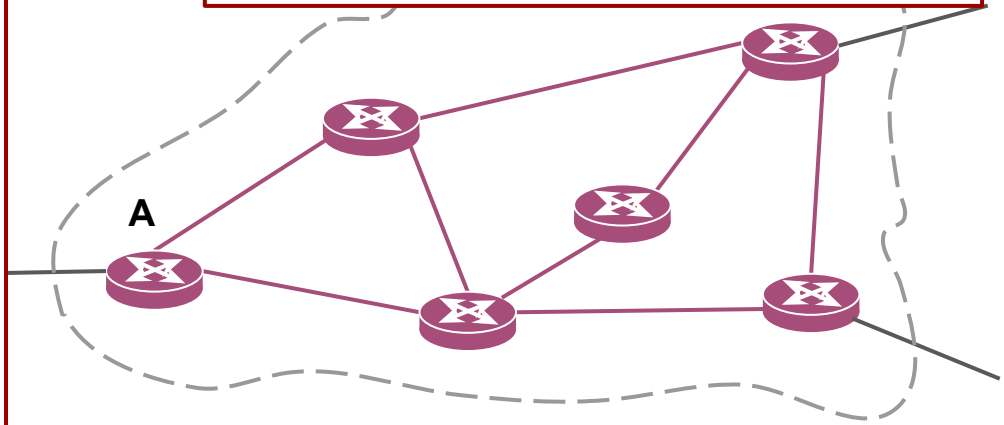

Software-Defined WAN

- Traffic engineering can be formulated as an optimization problem over the entire network
- Using SDN, we can directly solve it and find the optimal solution!
  - Send statistics about the traffic demand from the switches to the controller
  - Have the controller run an optimization problem to find the paths
  - Install the corresponding rules in the switch data plane.

Software-Defined WAN

- Not that easy in reality
  - Networks are large
  - Consistently updating the entire network is hard
  - Solving the optimization problem can take a long time
- Extensive research on how to make it a reality
- Widespread adoption in industry
  - Micorosft SWAN
  - Google B4
  - AT&T SD-WAN

Below the IP layer…

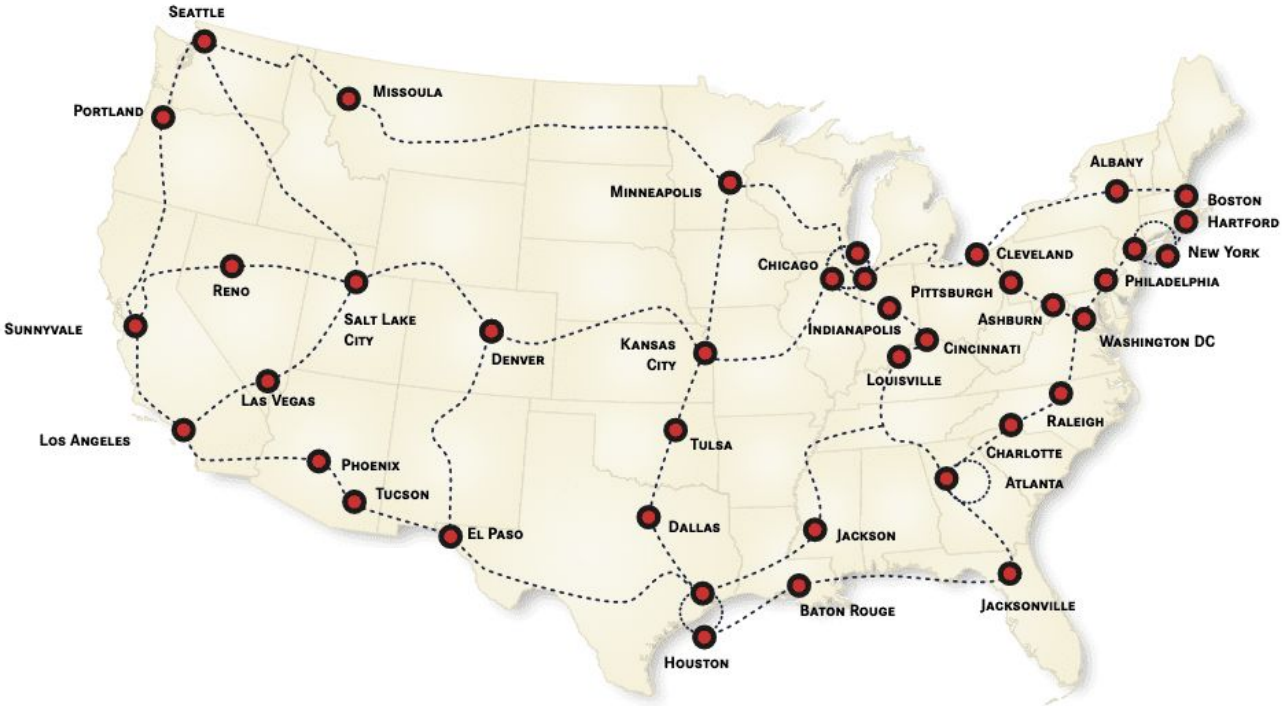

Below the IP layer…

Below the IP layer…

Optical links (Fiber)

Optical cross connects and IP routers co-located

Optical cross connects (no IP router here)

The optical layer is reconfigurable

- Optical cross connects can be configured to switch light between different ports
  - e.g., reconfigurable optical add-drop multiplexer (ROADM)
- By changing which ports are "connected together at the optical layer, we can create different topologies at the IP layer

The optical layer is reconfigurable

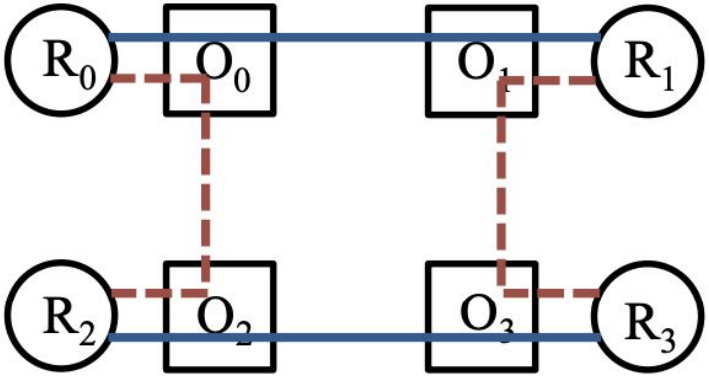

(a) Optical-layer configuration **A**.

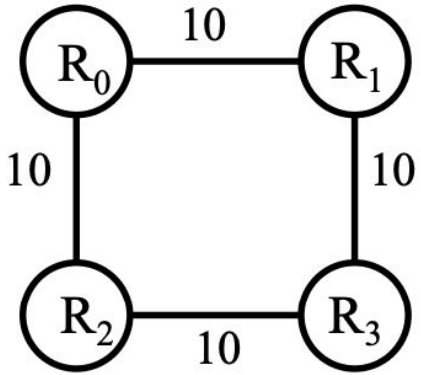

(b) Network-layer topology **A**.

The optical layer is reconfigurable

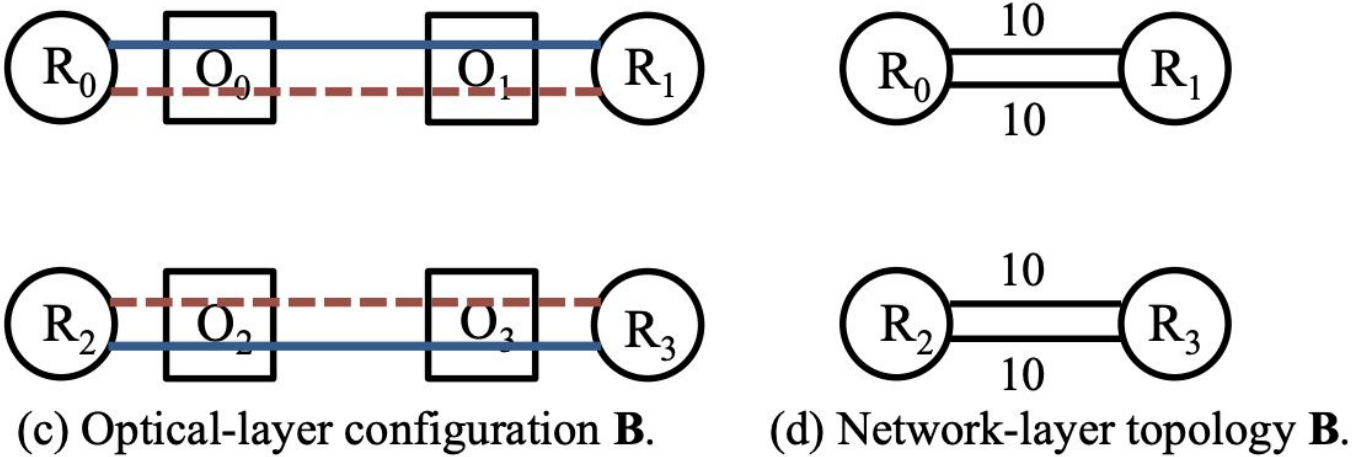

(c) Optical-layer configuration **B**.

(d) Network-layer topology **B**.

SDN for the optical layer?

- The reconfigurability at the optical layer allows us to move capacity from one link to another.
- When deciding what paths the traffic should take, we can also decide the capacities of the links!
- How would that help? Utilizing the network capacity much more effectively

Improving traffic engineering in the WAN

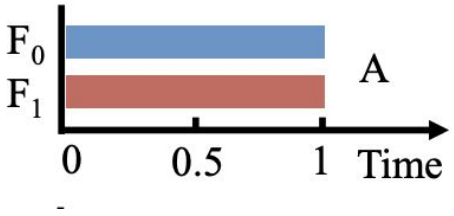

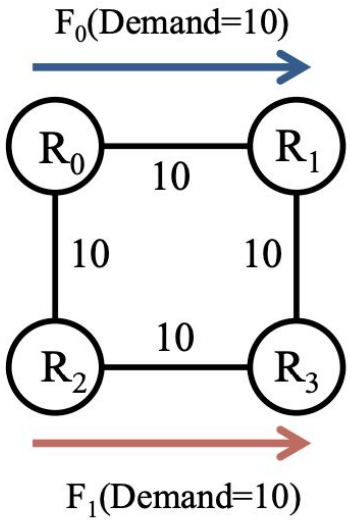

(a) Plan A.

Improving traffic engineering in the WAN

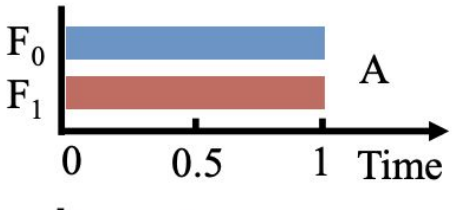

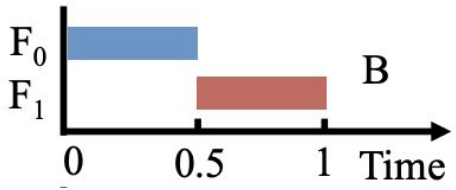

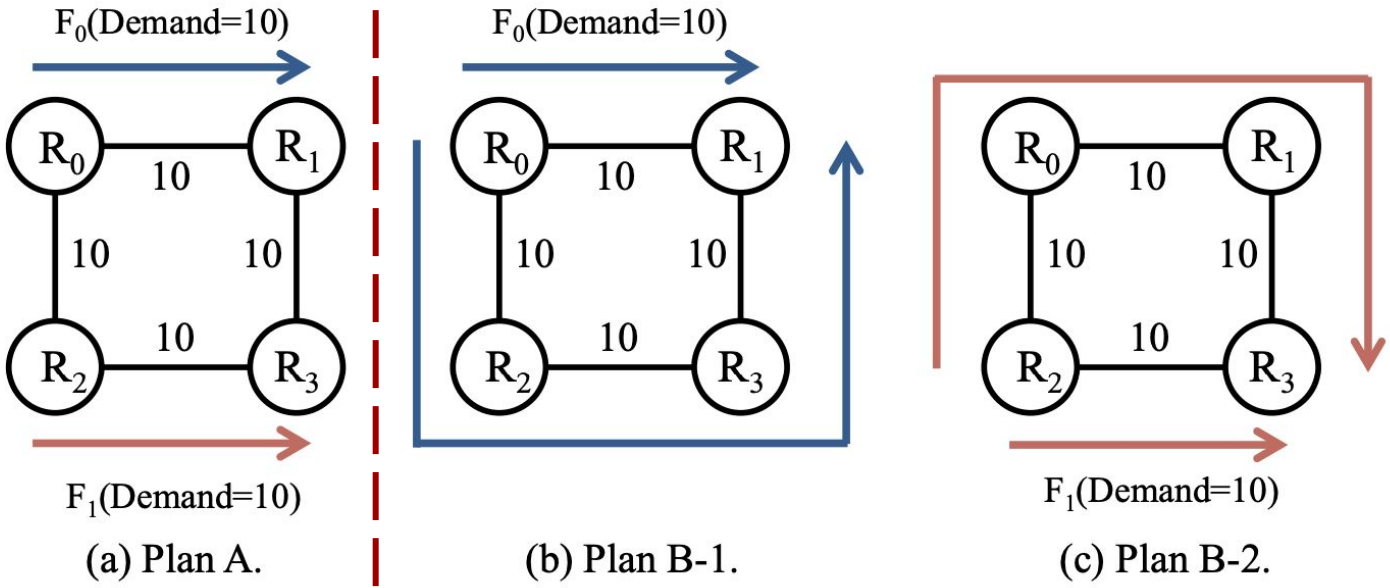

Improving traffic engineering in the WAN

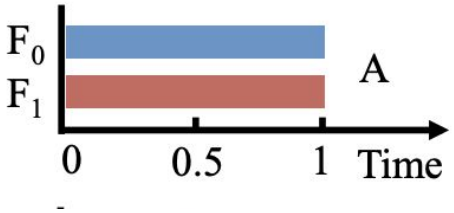

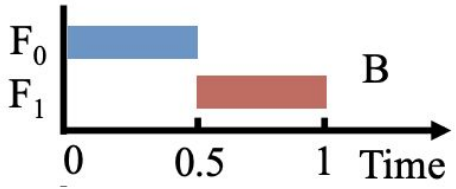

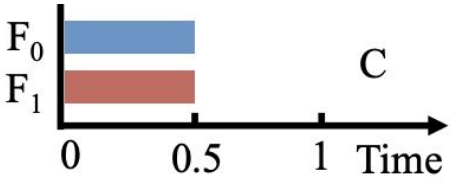

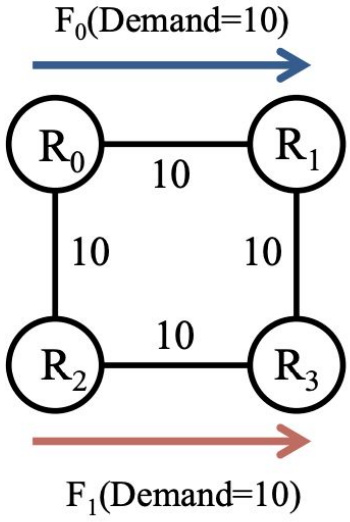

(a) Plan A.

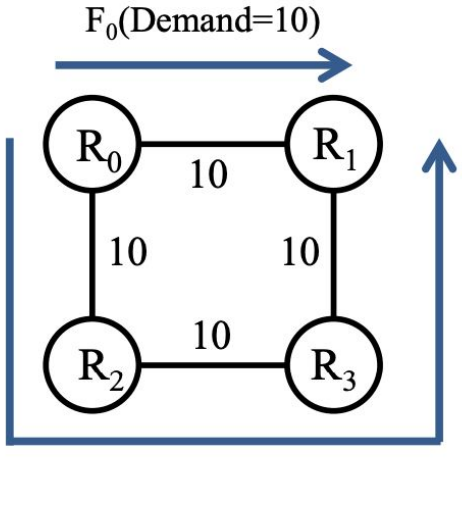

(b) Plan B-1.

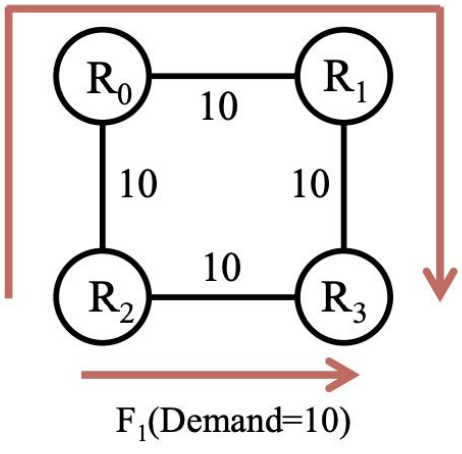

(c) Plan B-2.

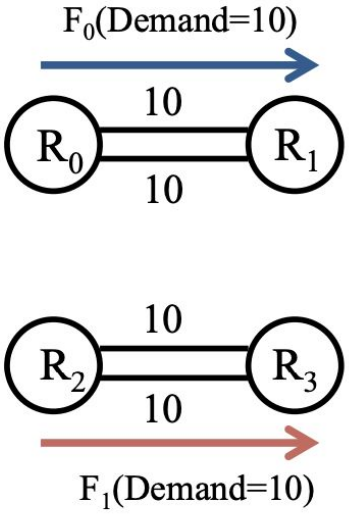

(d) Plan C.

Reconfigurable optics can be beneficial in data centers too

- Configure paths between ToRs based on demand (similar to WANs)
- Reduces the need for over-provisioning (similar to WANs)
- Reduce cabling costs :)

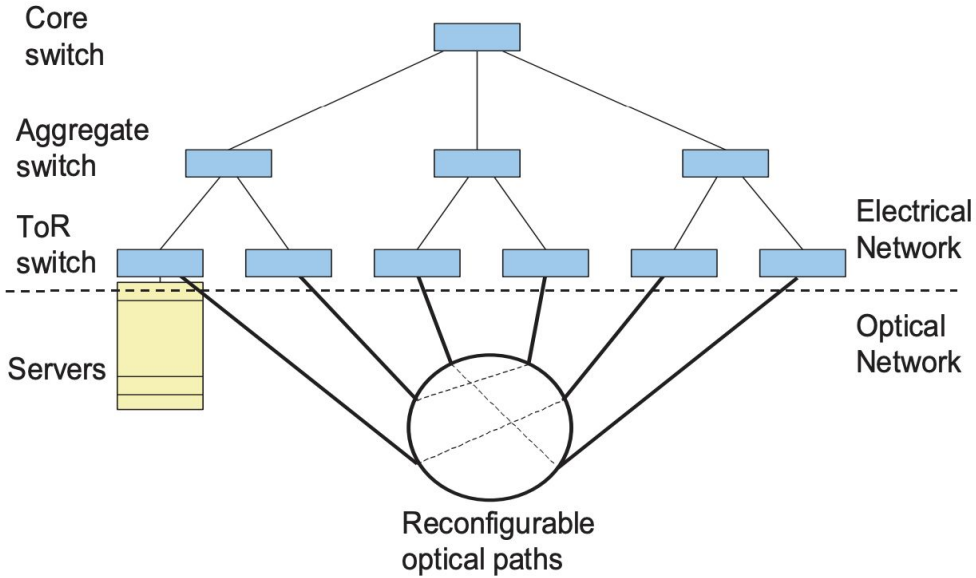

"c-Through: Part-time Optics in Data Centers", SIGCOMM'10

Wrap Up - Programmable Networks

- Week 1: SDN and OpenFlow
  - Distributed indirect control is challenging and error-prone!
  - A top-down approach: make control decisions (e.g., paths) centrally and directly install rules in the switches to implement it.
- Week 2: Programming the Data Plane with P4
  - Why stop at just installing rules?
  - Let's make the data plane pipeline programmable!
  - Programmable parsing, specifying match-action tables and layouts, etc.

Wrap Up - Programmable Networks

- Week 3: Programmable Switch Architectures
  - Network hardware should be really fast!
  - How do we design flexible but fast network hardware?
- Week 4: "Smart" Network Interface Cards
  - The end-points do network processing too, and sometimes that becomes the overhead
  - Can we program network interface cards to accelerate all of a part of it?
  - What is the right hardware architecture and programming model?

Wrap Up - Programmable Networks

- Week 5: Programming Software Network Stacks
  - It was always *possible* to change software-based network processing
  - But it's not easy, specially for low-level high-performance code!
  - What kind of programming abstractions can help? How about automation?
- Week 6: Applications to Traditional Networks
  - Even without direct programmatic control, abstraction and automation can help improve network management significantly!

Wrap Up - Programmable Networks

- Week 7: Network Verification
  - Can we formally verify that the network is doing what we want it to do?
  - How does network programmability help?
- Week 8: Flexible & Fine-Grained Network Monitoring
  - How do we use network programmability to find out what is happening in the network in real-time?

Wrap Up - Programmable Networks

- Week 9: Applications to Transport and Network QoS
  - A little help from the network can go a long way in improving performance and providing differentiated service.
  - How does network programmability help with that?
- Week 10: In-Network Computing
  - Can we accelerate distributed applications in the network? Should we?

Wrap Up - Programmable Networks

- Week 11: Software-Defined WAN, Reconfigurable Optical Networks
- Week 12: Project presentations
    - Looking forward to it!

That's it for CS 856!

Modern networks are large & complex systems.

Be it a clean-slate fully programmable network

or a traditional one

abstraction and automation is the way to go!