Internet Design, Data centers, & Future

Arvind Krishnamurthy

University of Washington

Design Goals

- Primary goal: Multiplexed utilization of existing networks
  - hook up existing L2 protocols => narrow-waist
  - packet switching vs. circuit switching
    - store and forward switches

Secondary Goals

- Survivability
- Support multiple types of services
- Accommodate a variety of networks

- Allow distributed management
- Must be cost effective
- Host attachment with a low level of effort
- Resource accountability

Survivability

- As long as the network is not partitioned, two endpoints should be able to communicate

- Maintain state only at end-points
  - fate-sharing
  - stateless network architecture

- Routing state is held by network
- No failure information is given to ends

Types of Services

- Reliable vs. unreliable
- Realized TCP wasn't always needed
- Separated TCP from IP and introduced UDP

Questions

- What is missing from the list?
- Which goals led to the success of the Internet?

Thought Exercise

- Why are proposals such as S-BGP, IPv6 not taking off?

Doom and Gloom in Networking (circa 2005)

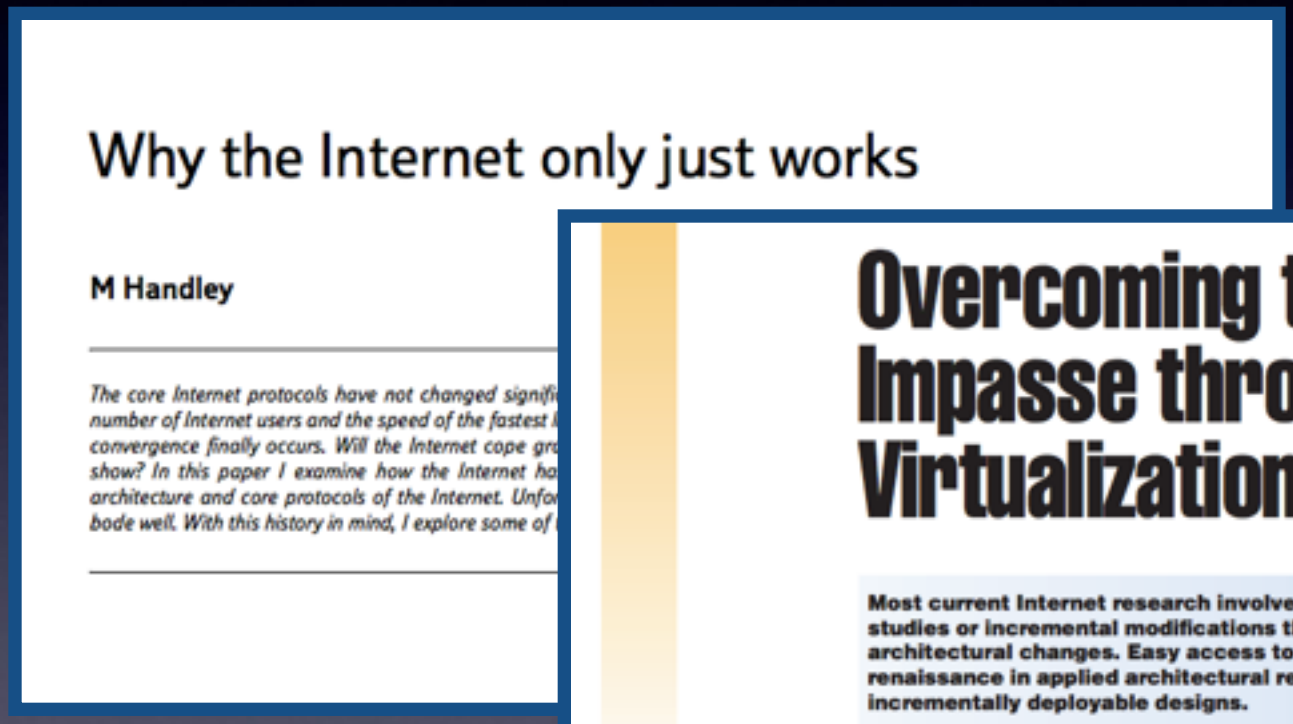

"Unfortunately, the recent history of failed architectural changes does not bode well."

"the prospects for significant change in its existing architecture appear slim.

Networking Woes

- Reasons for this "ossification":
  - Multiple agents ⇒ hard to achieve consensus
  - Incremental deployment ineffective ⇒ no incentives
  - Switch/router designs proprietary and baked in

Networking Resurgence

- Three key enablers:
  - Datacenter networking
  - Software defined networking
  - Bare-metal or open switches

Datacenter Networking

- Single administrative domain
- Rapid & wholesale upgrade
- Performance & cost are core issues

Datacenters Fostering Innovation

- Protocols evolving faster with datacenter
  - E,g., encapsulation protocols have gone through 2-3 generations of evolution in 5 years (VXLAN → NVGRE → Geneve)
- Switch vendors rapidly adding support for protocols
- Switches are becoming more “white boxes"

Typical Switch

- A switch typically consists of two "planes"
- Data plane: process packets based on local forwarding state
  - E.g., lookup destination → output port
- Control plane: compute forwarding state
  - E.g., run distributed protocols to determine routes and forwarding state

Typical Switch (circa 2007)

Control plane: compute the forwarding state

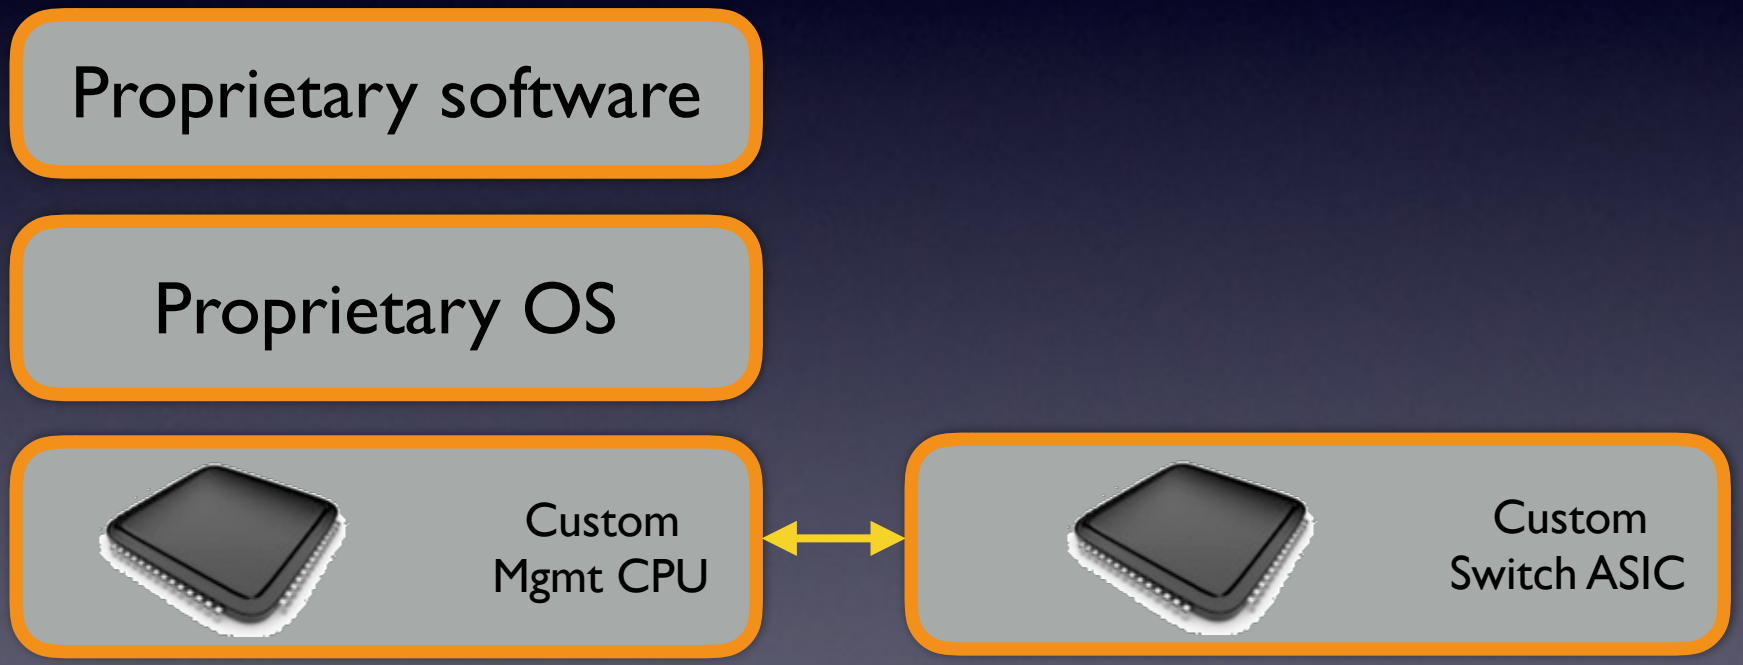

Data plane: process packets with local forwarding state

Software Defined Networking (SDN)

- Clean separation of the control plane and the data plane
- Key enabler: switch API to query topology and install forwarding state

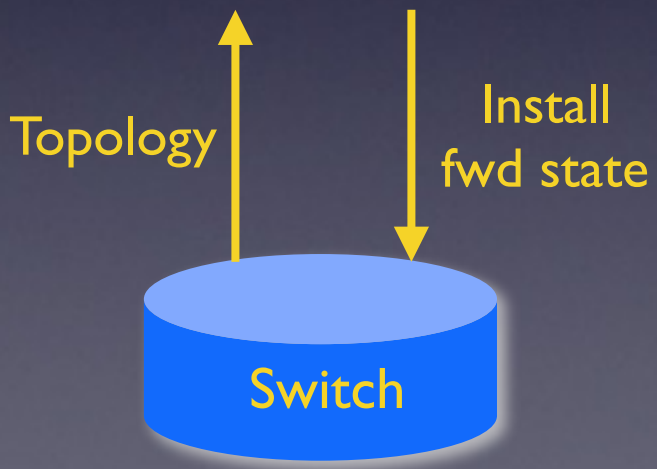

SDN Implications

- No more reliance on custom switch software
- Network *control plane* can be customized and run on traditional servers
  - control plane is "(end-user) software defined"
- Global network view enables simpler control programs

Bare-Metal Switches

- SDN created a market for barebones switches built using "merchant silicon"

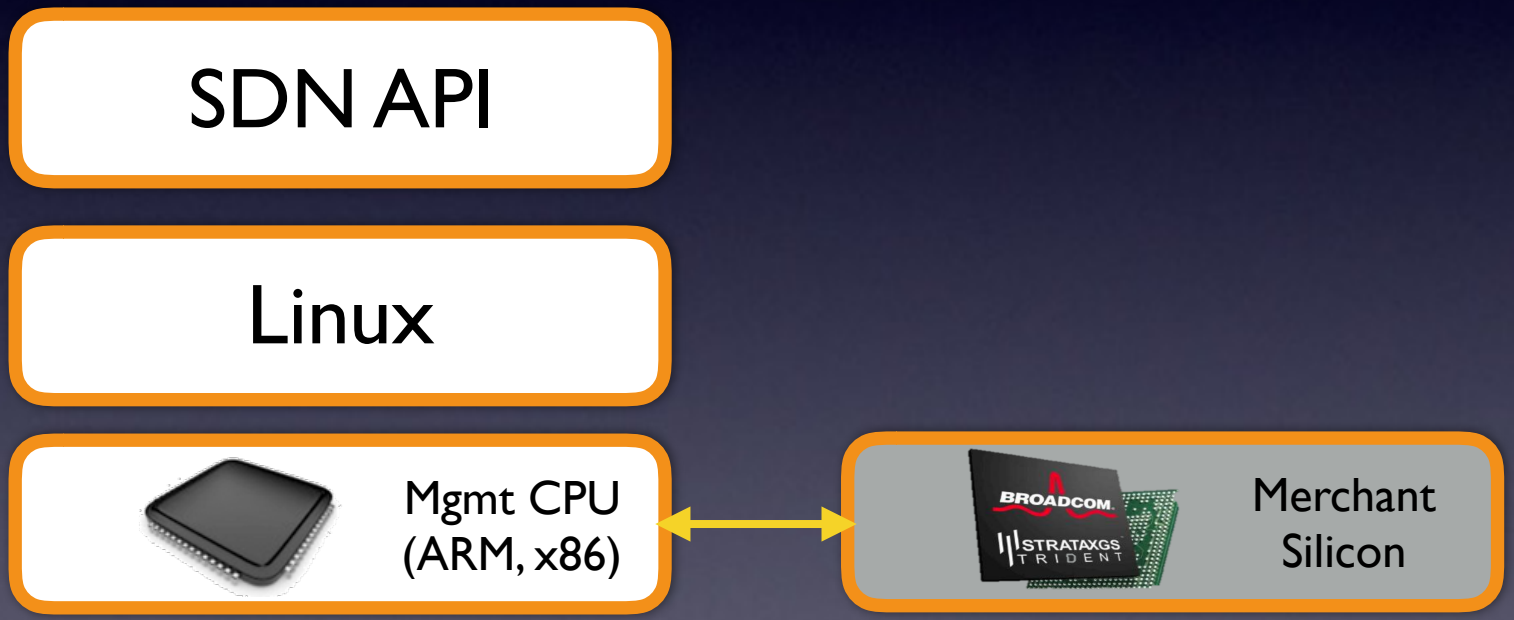

Lowers costs for datacenter operators

More homogeneous switch model

What is next?

- Can we exercise control over switching silicon?
  - Yes! There are reconfigurable switches that allow us to customize even the forwarding data plane.

Course Wrap Up

- Covered a broad set of topics:
  - distributed systems, operating systems, databases, networks
  - classical papers and recent papers
  - examined various issues/tradeoffs in building systems
  - systems building means new algorithms, performance improvements, taste in design, etc.
- Follow-on specialized courses will pick up from where we left off