Recent Trends in Networking

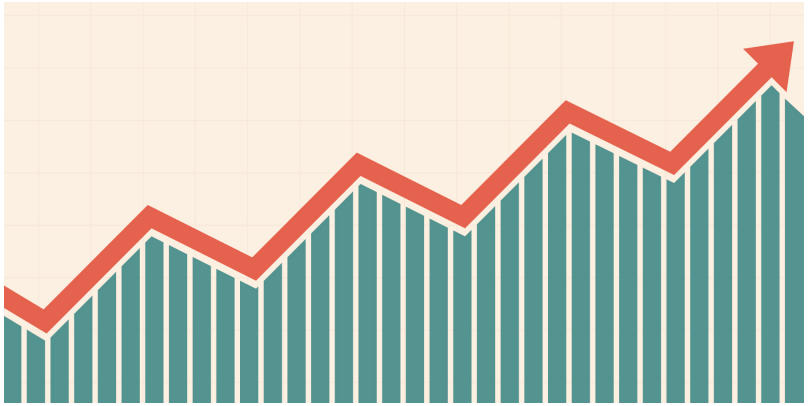

Raj Jain
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@cse.wustl.edu

These slides and audio/video recordings of this class lecture are at:
http://www.cse.wustl.edu/~jain/cse570-19/

Overview

1. Virtualization
2. Edge Computing
3. Disaggregation
4. Black Box to White Box
5. Decomposition
6. Microservices on Bare Metal
7. Intent-Based Policy Management
8. CI/CD
9. Reactive to Proactive
10. Standards to Open Source Software
11. Self-Driving Networks

Trend: Virtualization

- Started with virtual memory ⇒ virtual disk ⇒ Virtual Compute ⇒ Virtual Networks ⇒ Virtual Machines ⇒ Clouds ⇒ Virtual Network Functions ⇒ NFV ⇒ Micro-Services

Computing vs. Networking: Speed

- Moore's Law ⇒ Compute power grew 2x every 18 months

⇒ Memory size grew 2x every 18 months

Amdahl's Law: 1 Bit i/o per instruction

⇒ I/O grew 2x every 18 months

⇒ Networking speeds need to grow 2x every 18 months

⇒ 10 Mbps Ethernet in 1980 to 400 Gbps in 2018

⇒ 40,000x in 38 years ⇒ 2x every 30 months

⇒ Network is still the bottleneck

The Magic Word α

Distance-B/W Principle

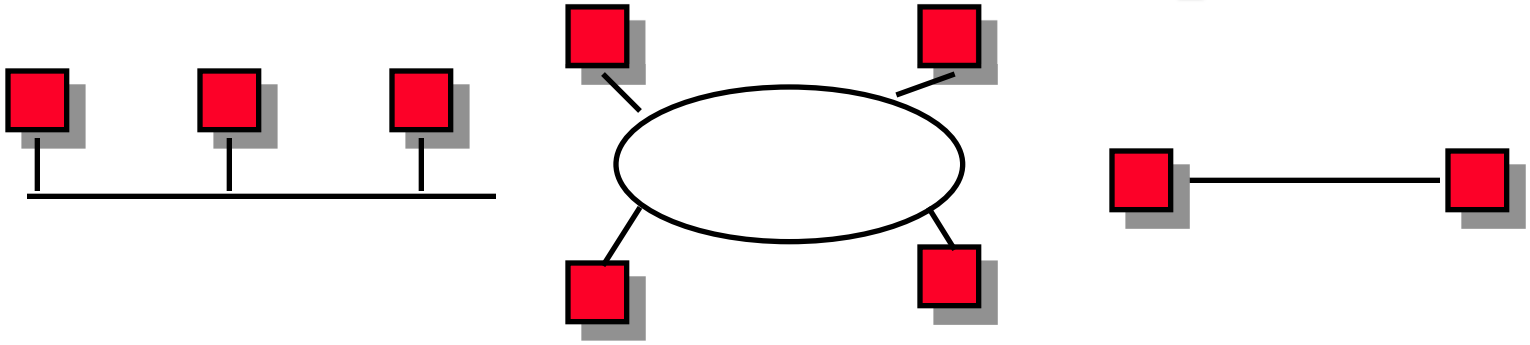

- Efficiency = Max throughput/Media bandwidth
- Efficiency is a decreasing function of α
  α = Propagation delay /Transmission time
  = (Distance/Speed of signal)/{(Transmission size/Bits/sec)}
  = Distance×Bits/sec/{(Speed of signal)(Transmission size)}
- Bit rate-distance-transmission size tradeoff.
- 100 Mb/s $\Rightarrow$ Change distance or frame size

Trend: Edge Computing

- Network protocols have the same performance as long as α is same.
- Multiple short networks are more economical than one long network
- Multiple short networks can transport more bits than one long distance network $\Rightarrow$ Small Cells $\Rightarrow$ Edge computing

Impact of High-Speed Networks

- Distance became inconsequential
  - ⇒ Globalization of information
  - ⇒ Globalization of trade
  - ⇒ Globalization of computing
- No need for local caching of information ⇒ Cloud storage
- No need for local computing ⇒ Cloud computing

Trend: Disaggregation

- Disaggregation of software and hardware
- Software on any hardware ⇒ hardware can change
  Different software on a hardware ⇒ Software can change
- Open source software on commodity hardware

Impact of High-Speed Hardware

- General purpose processors were fast enough for most computation

⇒ No need for special purpose hardware

⇒ All specialization and differentiation via software

⇒ White box networking

Trend: Black Box to White Box

- Black Box: Proprietary
- White Box: Open Source Hardware

Network Operating Systems

- OpenNetworkLinux: Linux distribution for open switches. Part of OCP.
- OpenSwitch: Open Network Operating System for switches under Linux Foundation
- ONIE: Open Network Install Environment. Part of OCP.
- SONiC: Software for Open Networking in Cloud. Open "switch OS" from Microsoft to run on white box switches
- Cumulus Linux: Network OS
- Switch Light OS: Network OS from BigSwitch Network based on ONL
- Big Cloud Fabric: Datacenter switch controller
- Big Monitor Fabric: Network monitoring and Security from BigSwitch networks
- IP Infusion OcNOS: Network OS
- IXIA Vision Edge OS: Monitoring software for open switches
- Pica8 PicOS: Switch operating system based on XORP open router platform. It runs on a Linux kernel and provides network and switching services.
- Pluribus Open Netvisor Linux: Network OS from Pluribus. Partners with Dell open switches.
- SnapRoute FlexSwitch: Set of networking applications, e.g., ARP, DHCP, BGP, VLAN, written in Go by SnapRoute. Accelerated by hardware. Merged in to Open Network Linux at OCP.

Ref: https://www.bigswitch.com/sdn-products/sdn-products/big-monitoring-fabric/overview, https://cumulusnetworks.com/products/cumulus-linux/, https://www.ipinfusion.com/products/ocnos/, https://www.ixiacom.com/products/vision-edge-os, https://en.wikipedia.org/wiki/Pica8, https://www.pluribusnetworks.com/, https://snaproute.com/, http://packetpushers.net/people-get-wrong-snaproutes-flexswitch/, http://onie.org/, https://opennetlinux.org/, https://www.openswitch.net/, https://azure.microsoft.com/en-us/blog/sonic-the-network-innovation-powerhouse-behind-azure/

White Box Switches

- EdgeCore Networks (ACTON)
- Quanta
- HPE
- DNI
- Dell
- Mellanox
- Delta Agema
- Celestica
- Alpha Networks
- Ingrasys
- Inventec
- Netberg

ASICs for White Box Switches

- Broadcom, Marvell, Intel/Fulcrum, Mellanox
- P4 ASICs from Barefoot Networks, and Cavium (by Xplient acquision)
- CPUs: Intel Rangeley, Freescale, ARM A9,

Trend: Decomposition

- Decompose a large complex application in to microservices
- Each service implements a small set of strongly related functions
- Functions that change together should be packaged together
  - Can develop and deploy independent of other services
- Loosely coupled: API binds it to other services. The implementation can be changed without affecting the result.
- Measurable, Testable
- Decomposition by:
  - Subdomains/Departments: Inventory, Order, Delivery, …
  - Capability:
- This allows some functions to be moved to Edge Computing

Ref: http://microservices.io/

Trend: Microservices on Bare Metal

- Clouds provide virtual machines
- There is trend against virtual machines
- Containers on bare metal ⇒ Kata Container

Cloud Native Applications

- Applications that run on multi-cloud
- Run on multiple servers at different locations
- Must be decomposed to allow faults
- Applications assembled as microservices in linux containers
- Cloud Native Computing Foundation (CNCF) – Linux Foundation project

Ref: https://searchitoperations.techtarget.com/definition/native-cloud-application-NCA

Serverless

- Don't need a local server. Just use cloud.
- Developers can work on an application without needing IT organization to provision hardware

Intent-Based Policy Management

- Intent: Tell what you want done
  Not how you want it done
  Example: Tell phone where you want to go. Fastest time, Shortest path, highway, …
- Invariance: Intent doesn't change if the network changes, devices fail, …
- Portability: Independent of infrastructure, equipment vendors, service providers, protocols used, media used, …
- Compose-ability: Can use any combination of infrastructure, …
- Scalable: From one to billions. Single controllers aren't scalable.
- Action requires context: Actions need to adopt to sudden changes in infrastructure,… ⇒ Rule-based, table-based approaches may not work
- Opendaylight has a new project on Network Intent Composition (NIC)

Ref: https://www.sdxcentral.com/articles/contributed/network-intent-summit-perspective-david-lenrow/2015/02/
https://wiki.opendaylight.org/view/Project_Proposals:Network_Intent_Composition

Trend: CI/CD

- Continuous Integration/Continuous Delivery
- Developer to Operations (DevOps)

Trend: Reactive to Proactive

- Past: Find failed components. Find why it failed?
- Now: Find components that will fail soon
- Anticipate what the future will be and act before it actually happens
- Causal Analysis: Why it failed

Black Box AI to Transparent AI

- Machine learning is currently a blackbox
- ML algorithms are developed/used without domain expertise
- Data cleanliness, labeling, feature extractions, all require domain knowledge, e.g.,
  What is the distance between Port 80, Port 81, and Port 8080?
- Data Imbalance (1 in a million packet is an attack packet).
- Use Synthetic data is used ⇒ Garbage-In, Garbage-Out
- Results are stated without model validation.
- Explainability issue ⇒ No idea of why the results are what they are

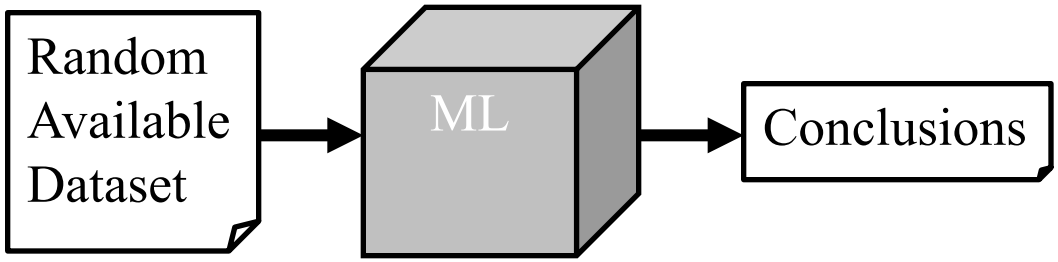

Standards are Slow

- Initially, Standards ⇒ Interoperability
  If all companies implemented the same standard in the same way, they could interoperate
- Standards=Compromises ⇒ We agree to disagree
  All differing opinions are part of the standard as option
  The standard is approved but different companies choose different options ⇒ No Interoperability
- Need Interoperability organizations
  - WiFi ⇒ Approves the subset of standard that is mandatory
- All this introduces delay ⇒ The standard is out of date when it is ready for implementation
- Example: IEEE 802.22-2011 Cognitive Radio (TV White spaces) Regional Area Networks. Started 2004. First standard in 2011. Revised in 2014.
- IEEE 802.11ah-2016 Long-Range WiFi for IoT. Started 2010. Taken over by competition: ZigBee, LoraWAN, …
- Standards are static. Can't change quickly.

Ref: http://www.ieee802.org/11/Reports/802.11_Timelines.htm

Standards are not Open

- Open ⇒ Anyone can implement it without fee
- IETF allows "non-discriminatory and reasonable licensing fee"
  ⇒ Not really open
- Open Source Initiative (OSI) Criteria:
  - No intentional secrets
  - Free and publicly available
  - All patents must be royalty-free for unrestricted use
  - No license agreements, NDA, or paperwork to implement
  - Not dependent on non-open standards

Ref: https://en.wikipedia.org/wiki/Open_standard, https://opensource.org/osr

Trend: Open Source Software

- Standardization to Rough Consensus and Running Code
- IETF has ~100 working groups
  Open Linux Foundation has >100 open source networking projects. Their website can't be kept uptodate.
- 4 Opens:
  - Open Source
  - Open Design
  - Open Development
  - Open Community

Trend 11: Managed to Self-Driven Networks

- **Self-Discover**: Find its components
- **Self-Organize and Self-configure**: Trending. Predict.
- **Auto-Manage** = Auto-BSS (bill)/Auto-OSS (provision)
- **Self-Monitor**: Counters and Probes. Telemetry
- **Self-Diagnose and Self-Heal**: Self-Report to human operator

Network Manager

Ref: Kireerti Kompella, https://datatracker.ietf.org/meeting/98/materials/slides-98-nmrg-self-driving-networks

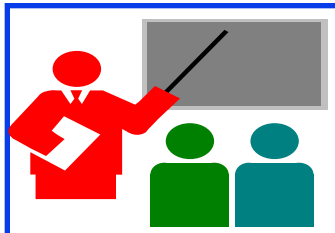

Summary

#	Past	Present/Future
1	SDN	Disaggregation
	Proprietary Black Boxes	Standardized White Boxes
2	Control	Orchestration
3	Clouds	Micro-Clouds
4	Core	Edge
5	Services	Micro-services
6	Orchestration of Switches	Orchestration of Multi-Cloud
7	Virtualization	Containerization
8	Standards	Open-Source SW
9	Centralized	Distributed
10	Smart	Intelligent
11	Managed	Self-Driven

References

- Kireerti Kompella, https://datatracker.ietf.org/meeting/98/materials/slides-98-nmrg-self-driving-networks
- http://microservices.io/
- https://www.sdxcentral.com/articles/contributed/network-intent-summit-perspective-david-lenrow/2015/02/
- http://www.ieee802.org/11/Reports/802.11_Timelines.htm
- https://opensource.org/osr
- https://github.com/cncf/landscape
- http://www.spdk.io/
- https://dirtycow.ninja/
- https://github.com/kata-containers/community

Wikipedia Links

- https://en.wikipedia.org/wiki/Native_cloud_application
- https://en.wikipedia.org/wiki/CNCF
- https://en.wikipedia.org/wiki/Linux_Foundation
- https://en.wikipedia.org/wiki/Kubernetes
- https://en.wikipedia.org/wiki/Microservices
- https://en.wikipedia.org/wiki/Proactive_maintenance
- https://en.wikipedia.org/wiki/CI/CD
- https://en.wikipedia.org/wiki/Continuous_delivery
- https://en.wikipedia.org/wiki/Continuous_integration
- https://en.wikipedia.org/wiki/CICD
- https://en.wikipedia.org/wiki/Continuous_deployment
- https://en.wikipedia.org/wiki/Docker_(software)
- https://en.wikipedia.org/wiki/White_box_(computer_hardware)

Wikipedia Links (Cont)

- https://en.wikipedia.org/wiki/Edge_computing
- https://en.wikipedia.org/wiki/Fog_computing
- https://en.wikipedia.org/wiki/Edge_device
- https://en.wikipedia.org/wiki/Micro_data_center
- https://en.wikipedia.org/wiki/Cloudlet
- https://en.wikipedia.org/wiki/Network_virtualization
- https://en.wikipedia.org/wiki/Mobile_virtualization

Acronyms

- AI Artificial Intelligence
- API Application Programming Interface
- ARM Advanced RISC Machines
- ASICs Application Specific Integrated Circuits
- BGP Border Gateway Protocol
- BSS Business Support Systems
- CI/CD Continuous Integration/Continuous Delivery
- CNCF Cloud Native Computing Foundation
- CRUD Create Read Update and Delete
- DevOps Developer to Operations
- DHCP Dynamic Host Control Protocol
- IEEE Institution of Electrical and Electronic Engineers
- IETF Internet Engineering Task Force
- IoT Internet of Things
- IT Information Technology
- NCA Native Cloud Application

Acronyms (Cont)

- NDA Non-Disclosure Agreement
- NIC Network Intent Composition
- OCP Open Compute Project
- ONIE Open Network Install Environment
- ONL Open Network Linux
- OS Operating System
- SONiC Software for Open Networking in the Cloud
- TV Television
- VLAN Virtual Local Area Network
- WiFi Wireless Fidelity
- XORP eXensible Open Router Platform

Scan This to Download These Slides

Raj Jain
http://rajjain.com

Related Modules

CSE567M: Computer Systems Analysis (Spring 2013),
https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),

https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcgy5e_10TiDw

Wireless and Mobile Networking (Spring 2016),
https://www.youtube.com/playlist?list=PLjGG94etKypKeb0nzyN9tSs_HCd5c4wXF

CSE571S: Network Security (Fall 2011),

https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u

Video Podcasts of Prof. Raj Jain's Lectures,
https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw