Knowledge-Defined Network Orchestration in a Hybrid Optical/Electrical Datacenter Network

Wei Lu (Postdoctoral Researcher)
On behalf of Prof. Zuqing Zhu
University of Science and Technology of China, Hefei, China
zqzhu@ieee.org



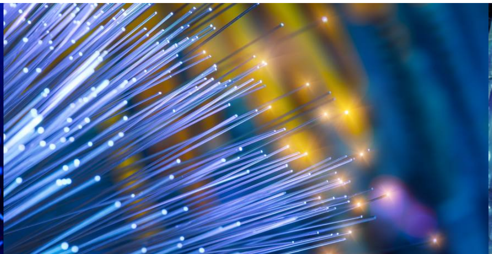


Outline

- Background and Motivations
- Hybrid Optical/Electrical Datacenter Network
- Knowledge-Defined Network Orchestration Mechanism
- Prediction Analytics and Decision Making
- System Implementation and Experimental Demonstration
- Conclusion

Cloud Computing Era

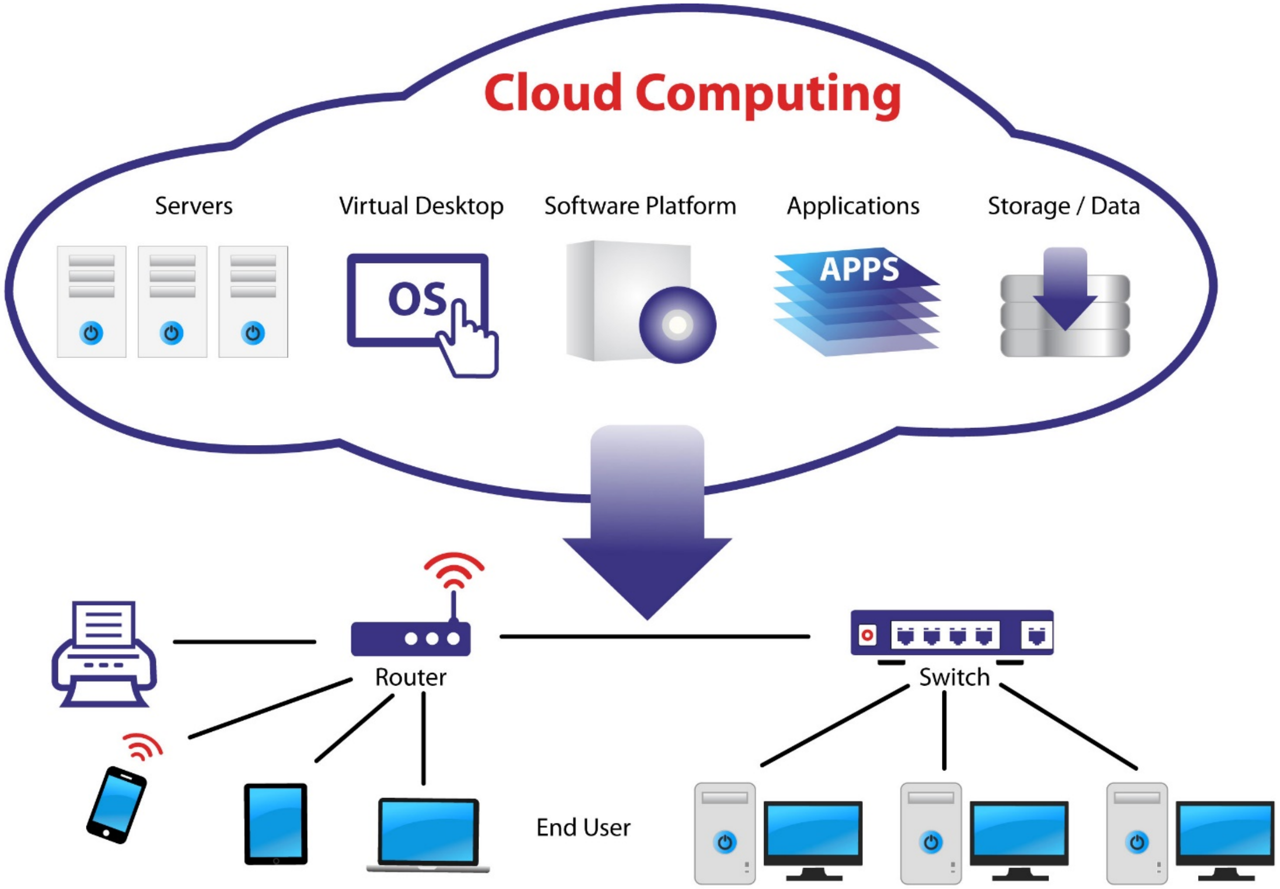

Data Center Networks (DCNs)

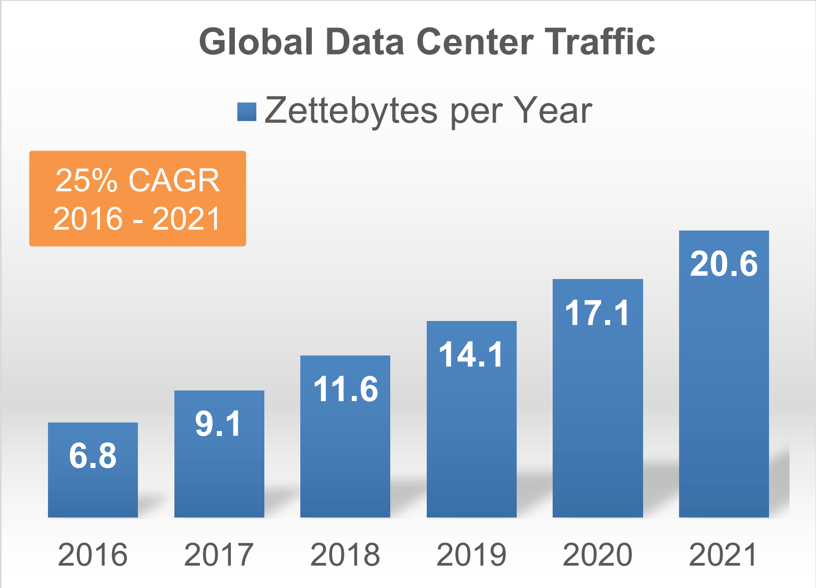

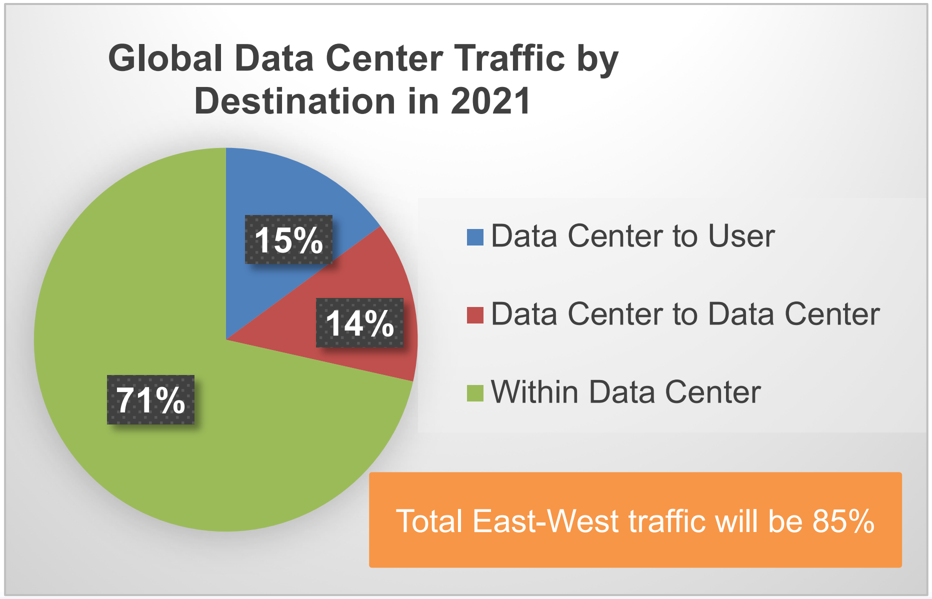

Network Orchestration Mechanism (NO-M)

- Resource Utilization
- Energy Consumption
- Quality of Service

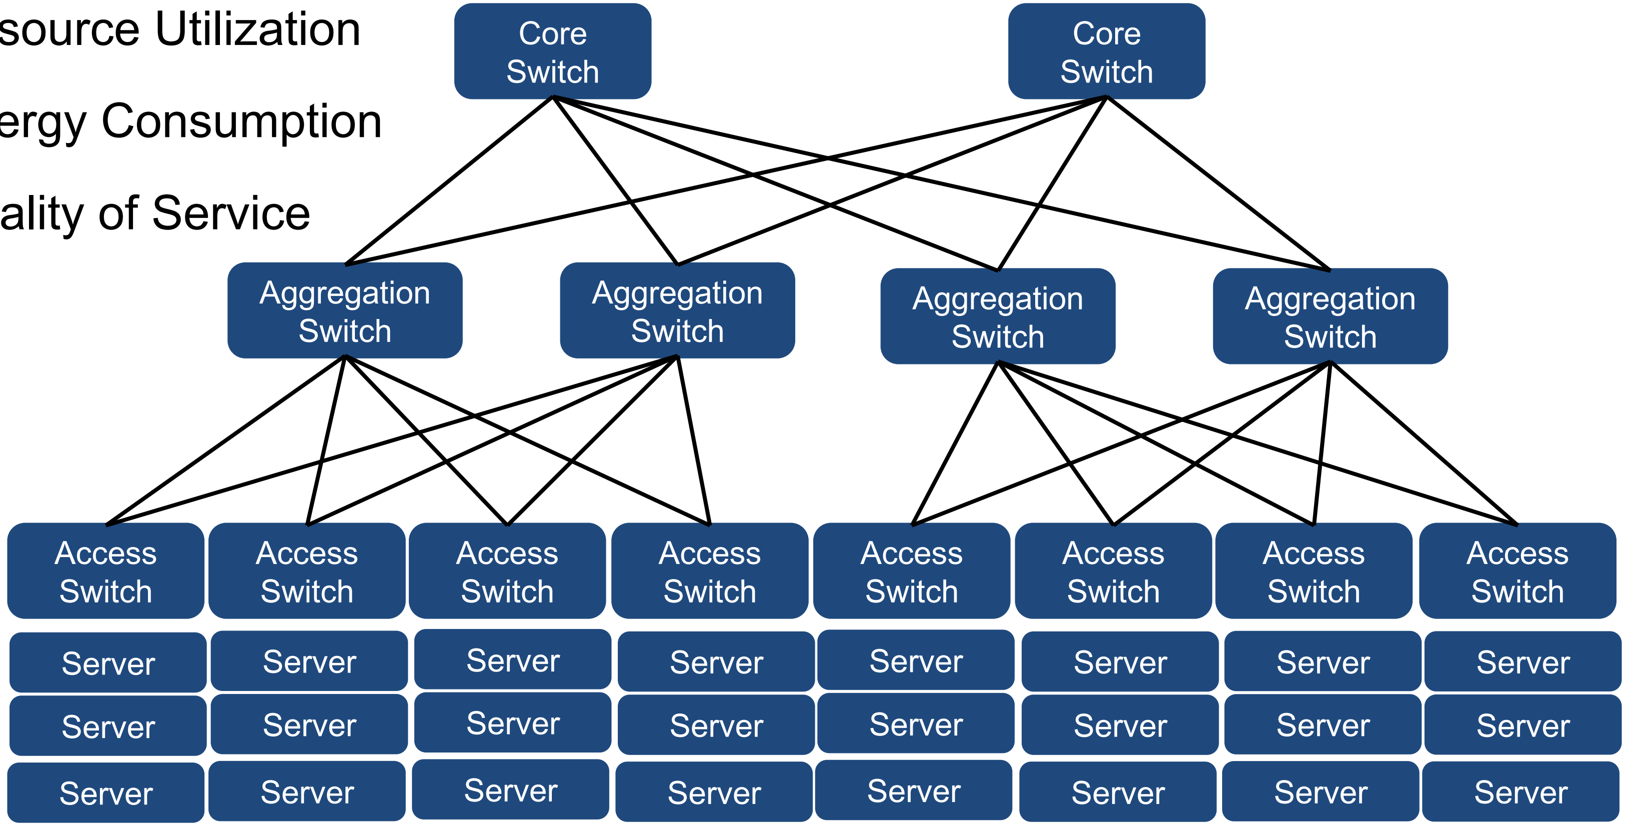

Web Server

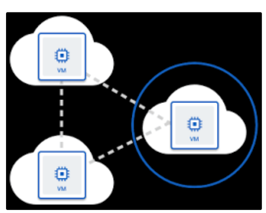

Virtual Network

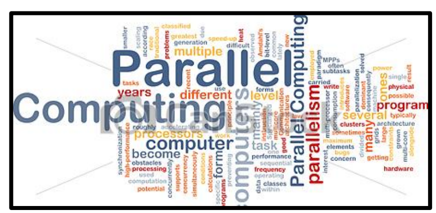

Computing Tasks

Virtual Desktop

Challenges to Realize an Efficient NO-M

- Manage a large number of network elements

Server
Server
Server

Access Switch

Aggregation Switch

Core Switch

……

- Handle computing tasks and traffic in DCNs that are highly dynamic

hadoop Map Reduce

"A 1-millisecond advantage in trading applications can be worth $100 million a year to a major brokerage firm"

Software Define Networking

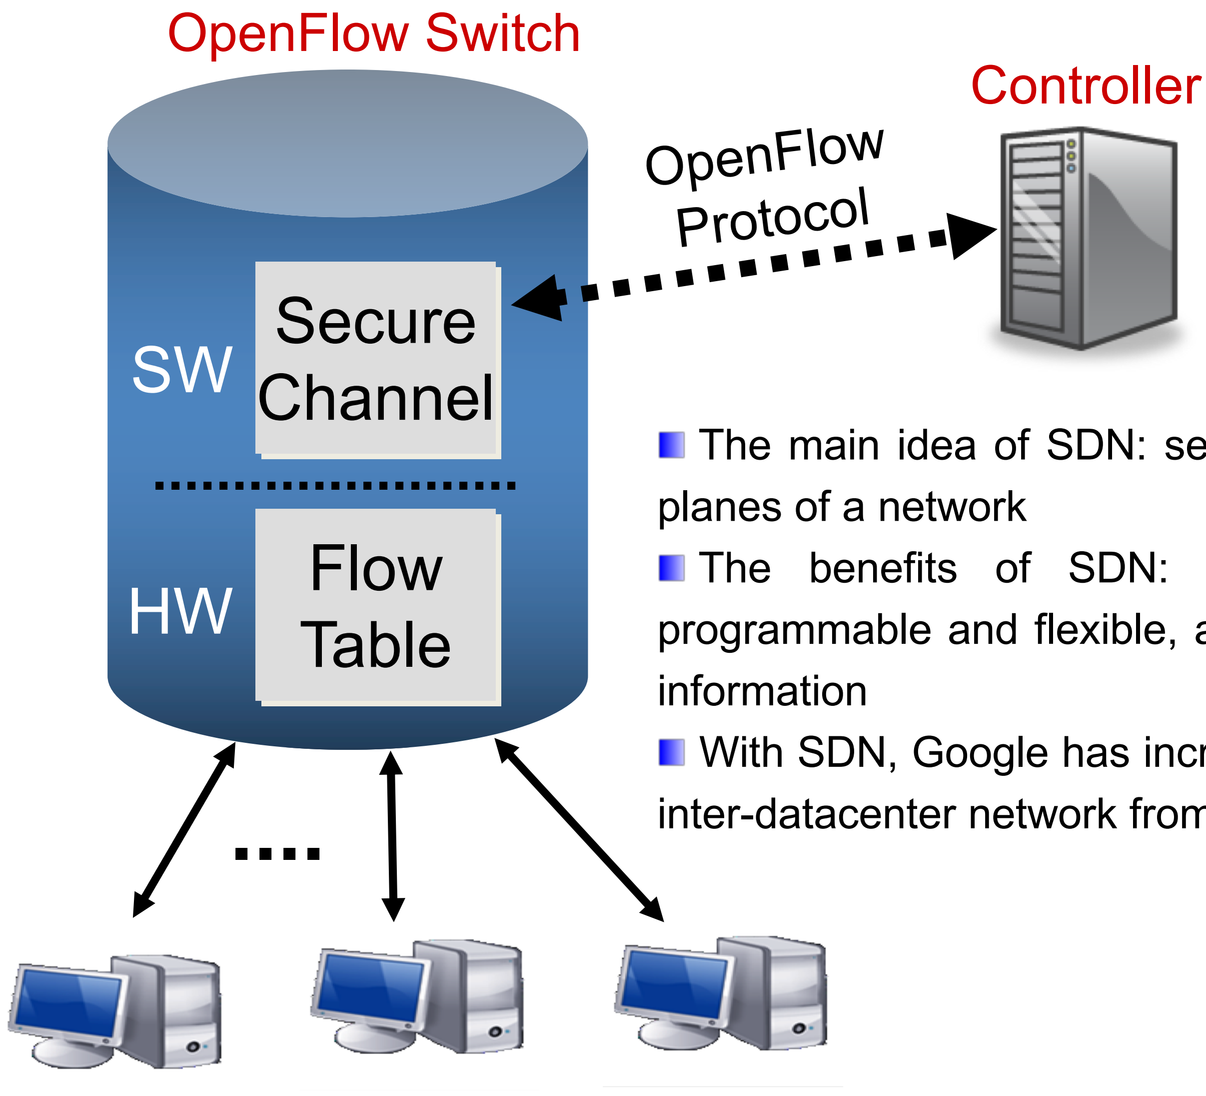

- The main idea of SDN: separating the control and data planes of a network
- The benefits of SDN: making networks be more programmable and flexible, a global review, rich telemetry information
- With SDN, Google has increased the link utilization of its inter-datacenter network from 30% to 95%.

Knowledge Defined Networking

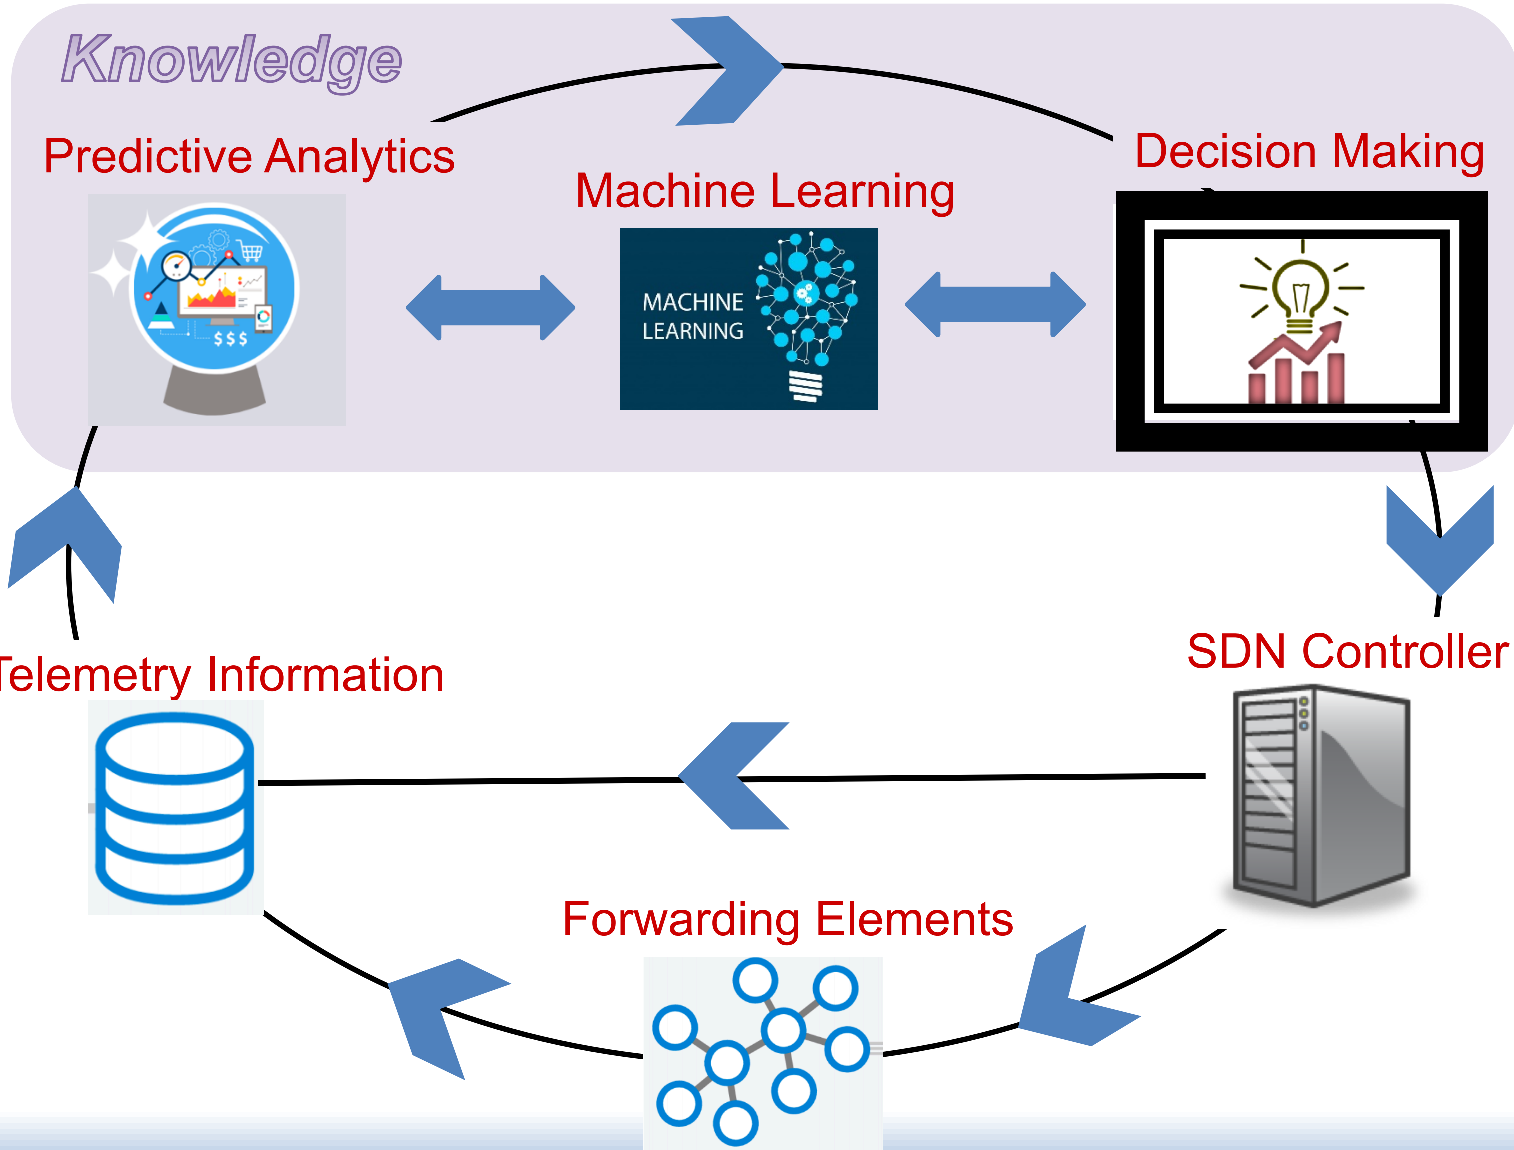

Hybrid Optical/Electrical DCNs (H-O/E DCNs)

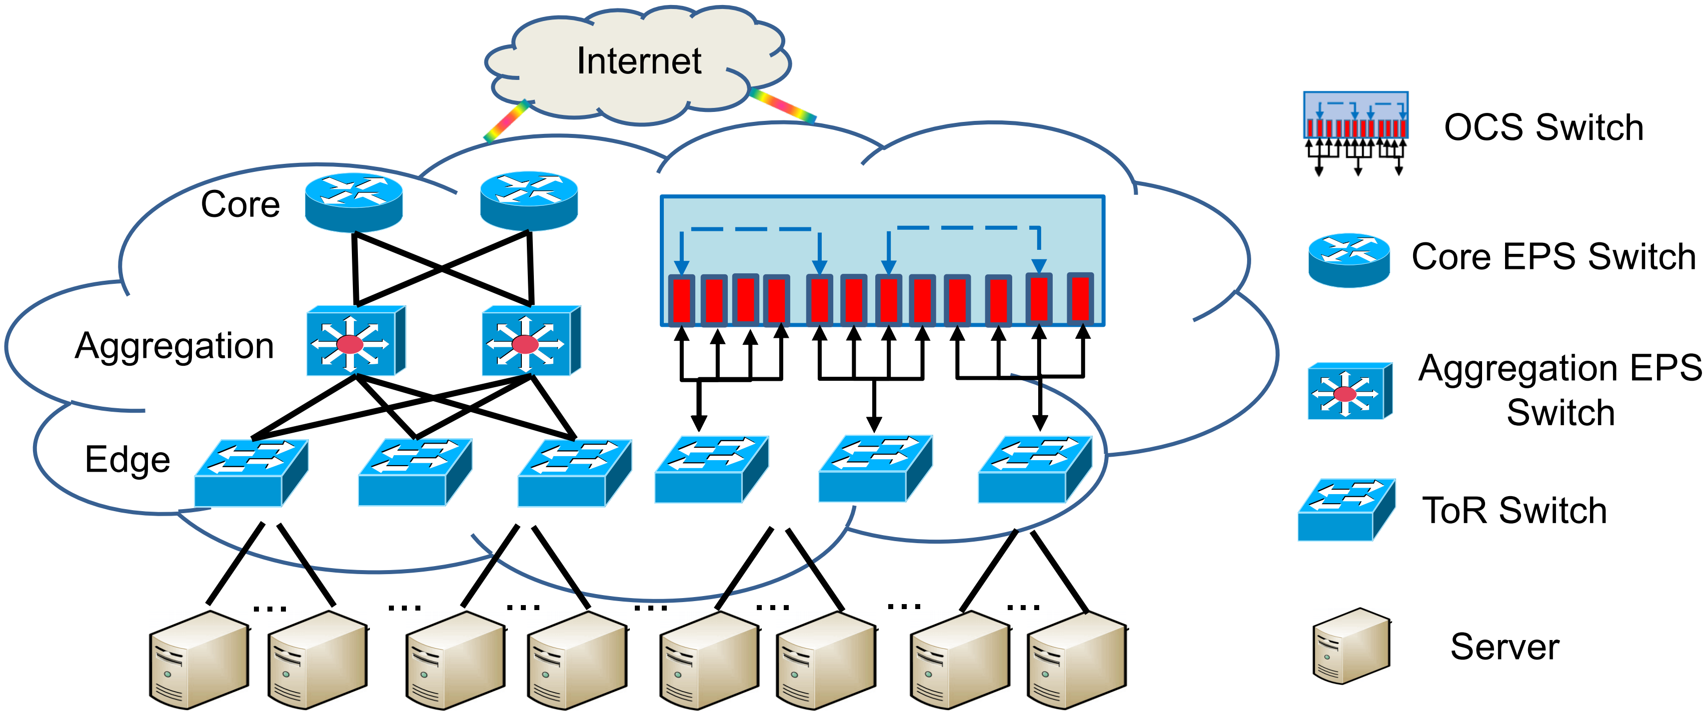

- Electrical Inter-Rack Interconnection
  - Packet-level switching granularity
  - Quick response upon request
  - Delay-sensitive and/or highly dynamic inter-rack traffic
  - Congestion caused by oversubscription

- Optical Inter-Rack Interconnection
  - Wavelength-level switching granularity
  - Pump high-throughput traffic through
  - Bandwidth-sensitive and/or long-lasting inter-rack traffic
  - Relatively long path setup latency

System Architecture Design

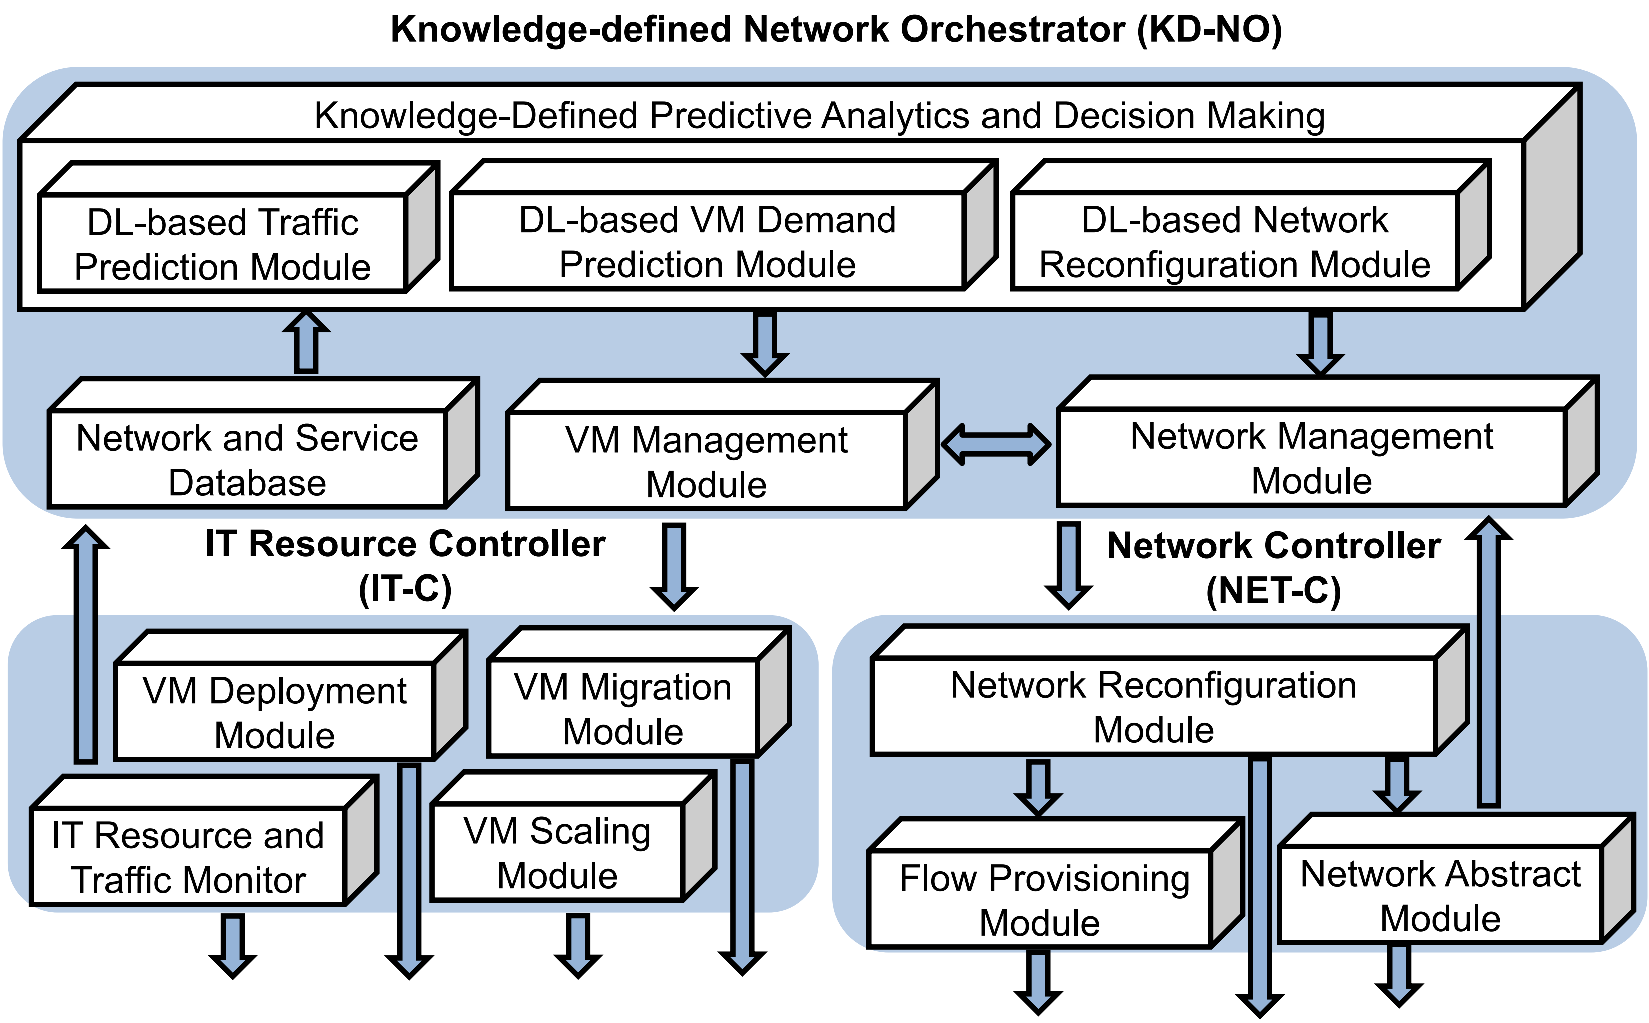

Deep Neural Network (DNN)

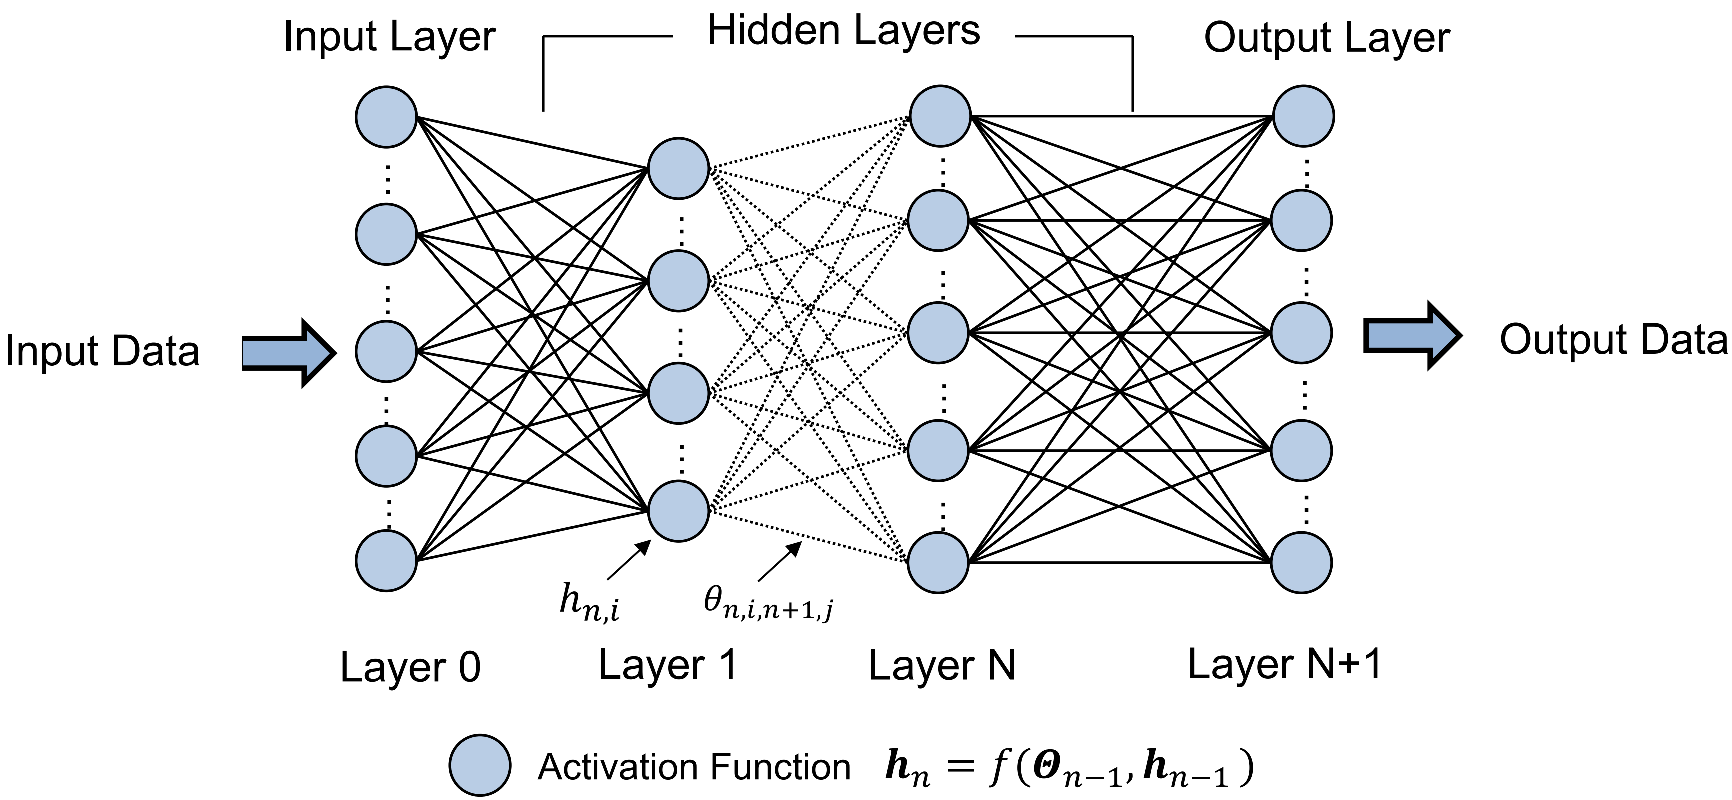

Case Study: VM Migration

- Objective: improve energy-/resource-efficiency in the H-O/E DCN
- Predictive Analytics
  - Traffic matrix between VMs
  - Inter-rack traffic matrix
- Decision Making
  - Which VMs to be migrated?
  - Where to migrate?

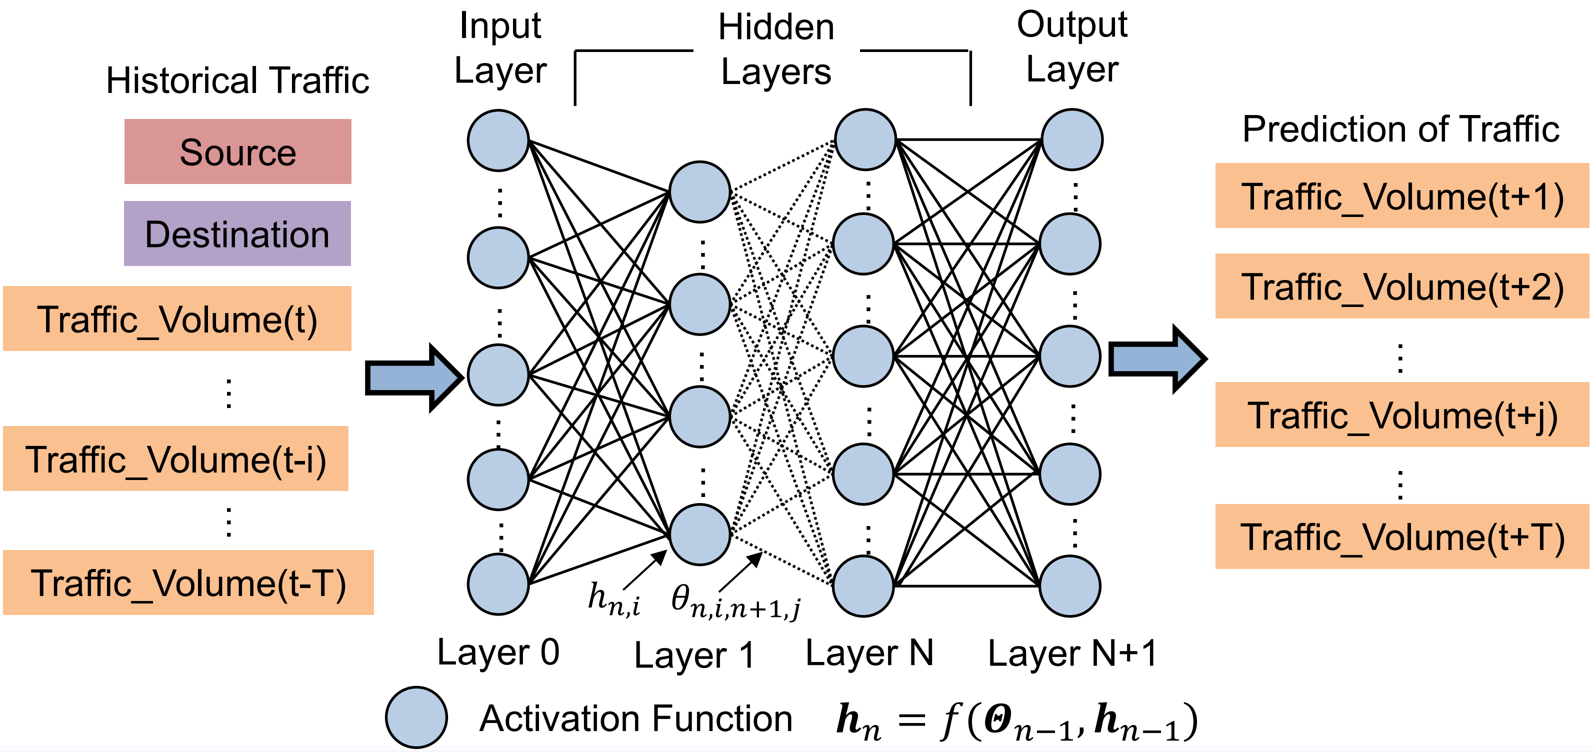

Case Study: VM Scaling

- Objective: ensure good performance of VMs
- Predictive Analytics
  - IT resources demanded by VMs
- Decision Making
  - Which VMs to be scaled and how?

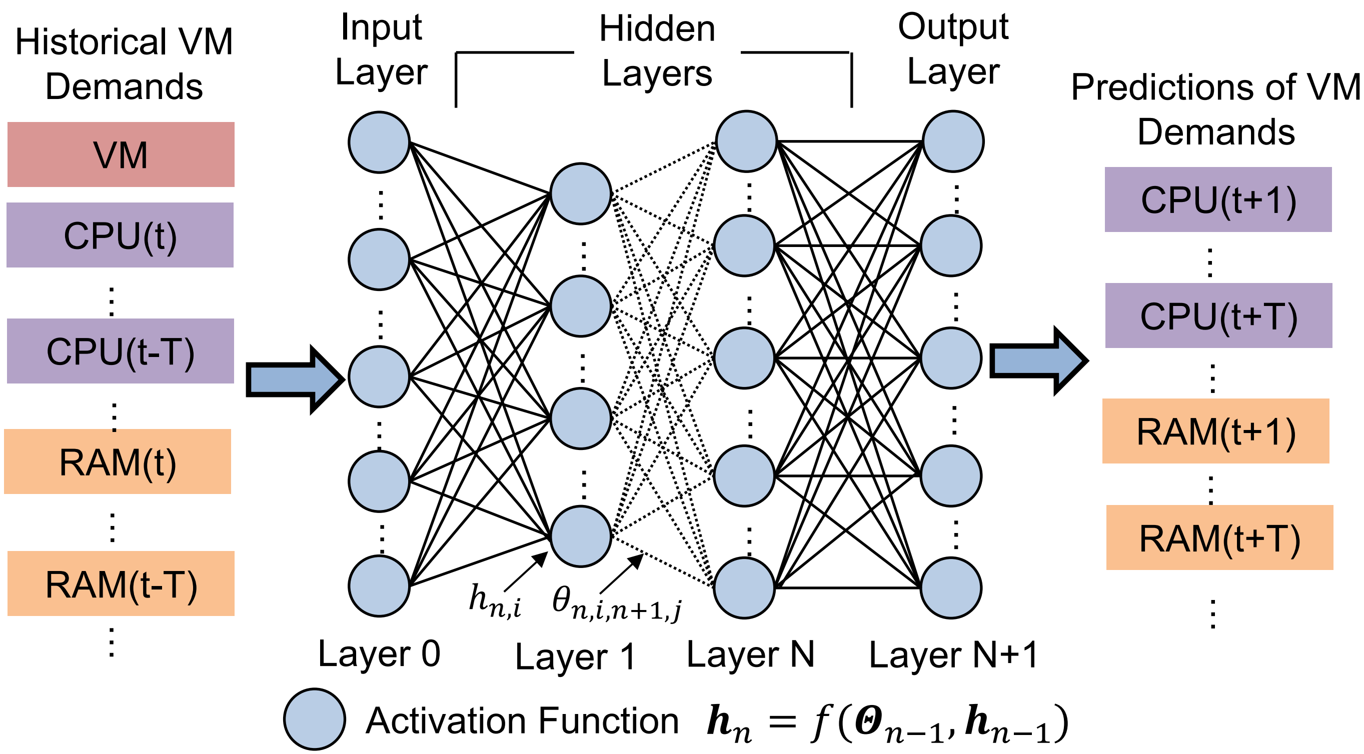

Case Study: H-O/E DCN Reconfiguration

- Objective: reduce average data-transfer delay
- Predictive Analytics
  - Input: inter-rack traffic and DCN configuration
  - Output: average data-transfer delay
- Decision Making
  - Optimal DCN configuration

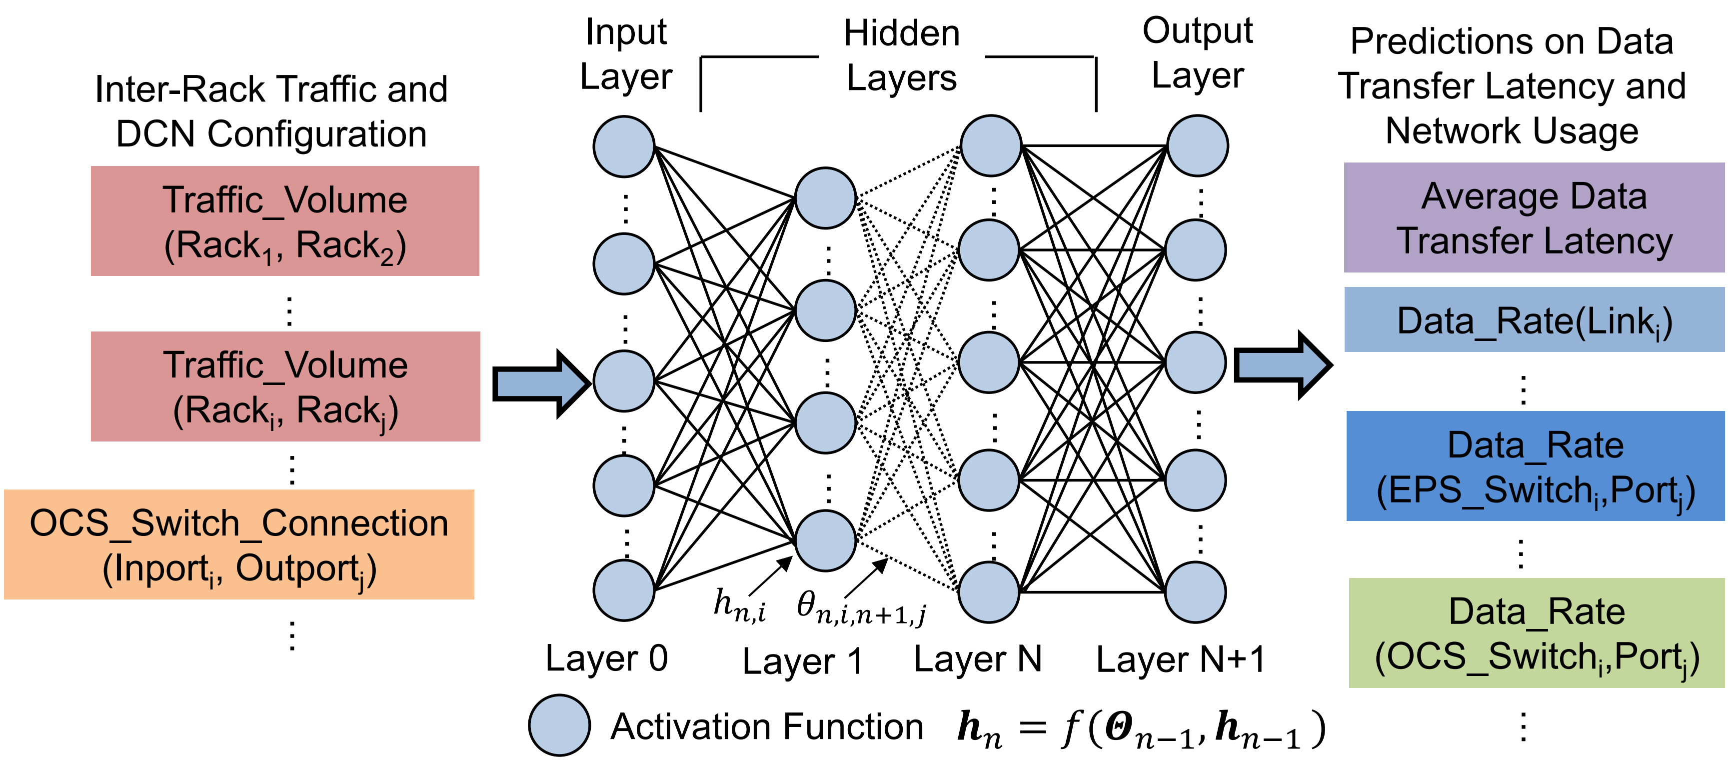

System Implementation: Data Plane

- Electrical Inter-Rack Interconnections

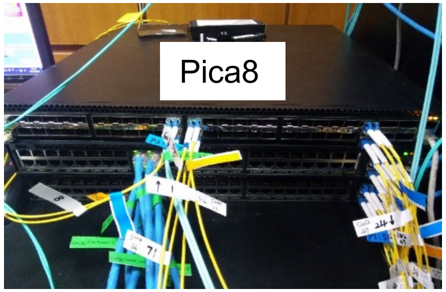

- Optical Inter-Rack Interconnections

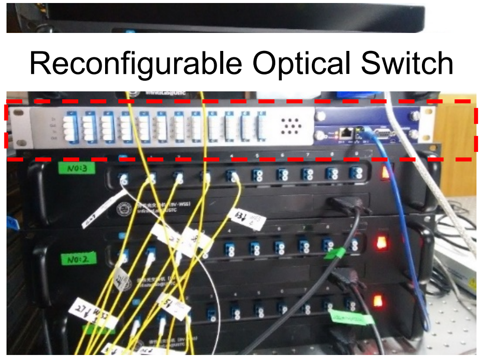

- Servers in Racks

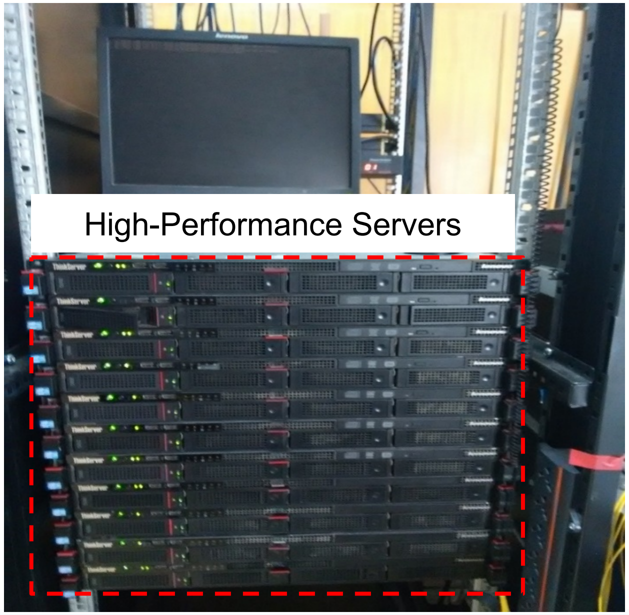

System Implementation: Control Plane

- IT-C: manage IT resource in a DC
  - OpenStack:
    - VM deployment and scaling
    - IT resource and traffic monitoring
- NET-C: control intra-rack and inter-rack connections
  - ONOS:
    - Flow provisioning
    - Network abstract and configuration
- KD-NO: coordinate IT-C and NET-C
  - VM and network management
  - Predictive analytics and decision making

Results in VM Migration

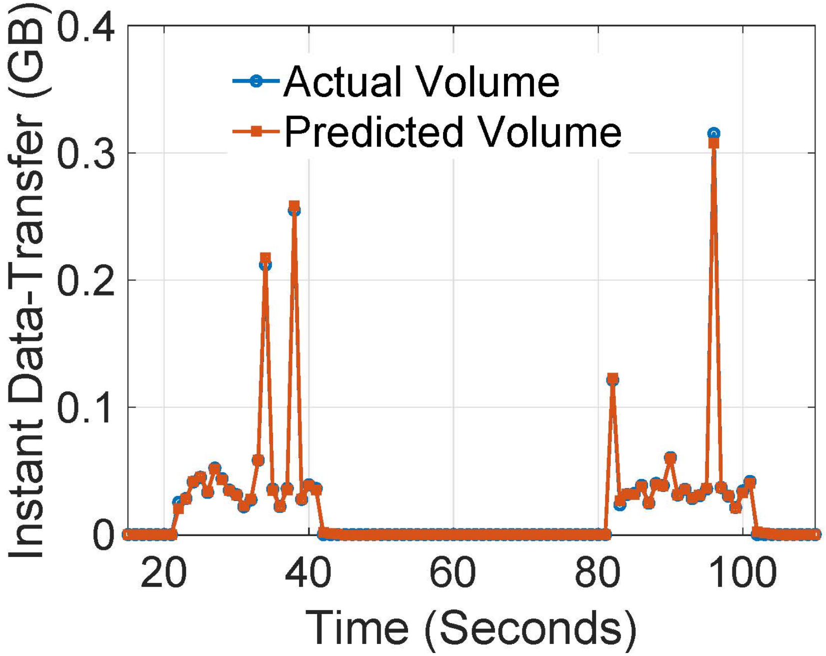

(a) Predicted and actual traffic volumes

	Rack 1	Rack 2	Rack 3
Rack 1	0	1.8489	0.0060
Rack 2	0.0044	0.0218	0.8262
Rack 3	0.0042	0.9130	0.0506

(1) without VM Migration (GB)

	Rack 1	Rack 2	Rack 3
Rack 1	0	0.7935	1.0614
Rack 2	0.0042	0.0002	0.0090
Rack 3	0.0044	0.9130	1.7802

(2) with VM Migration (GB)

(b) Data-transfer volumes for with and without VM migration

Results in VM Scaling

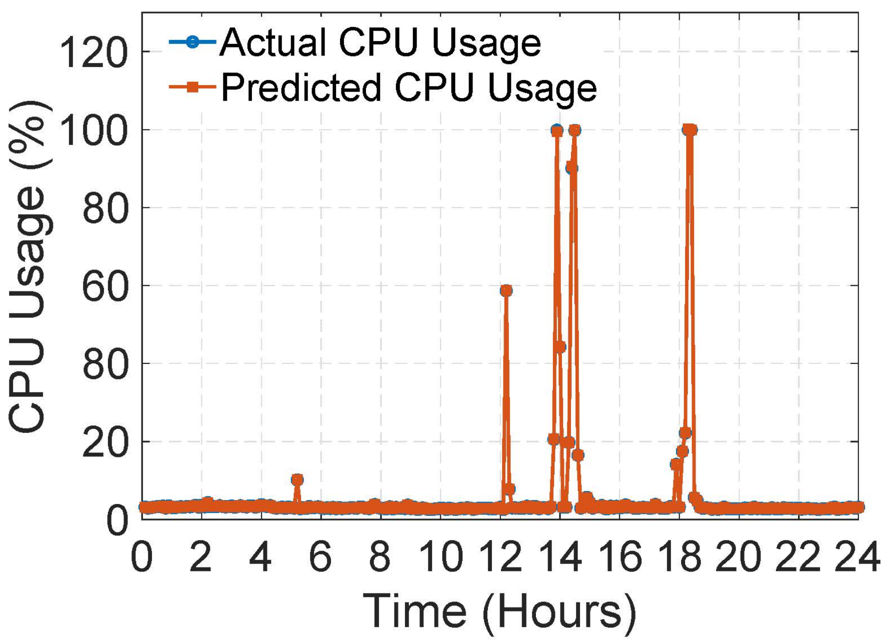

(a) Predicted and actual CPU usages of a VM

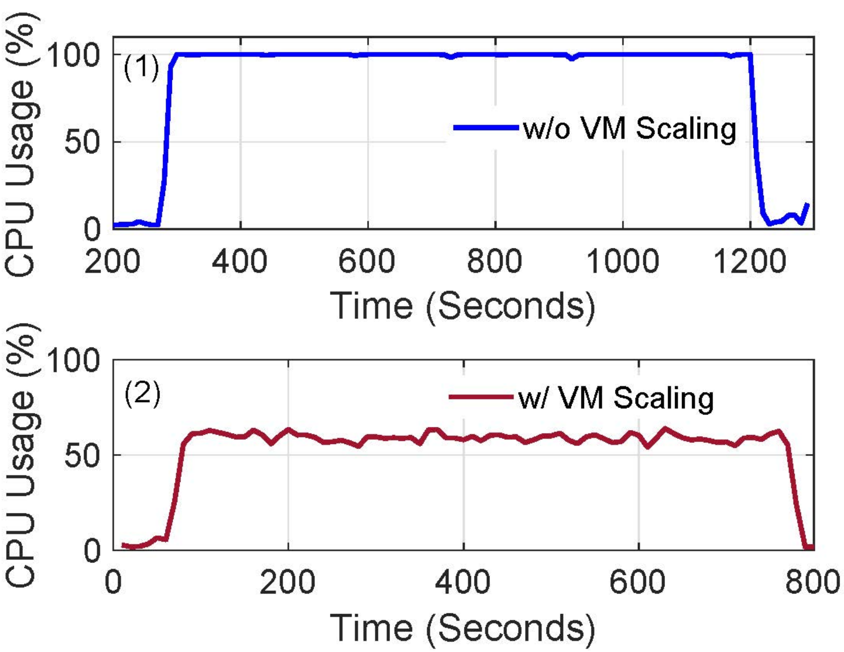

(b) CPU usage with and without VM scaling

Results in H-O/E DCN Configuration

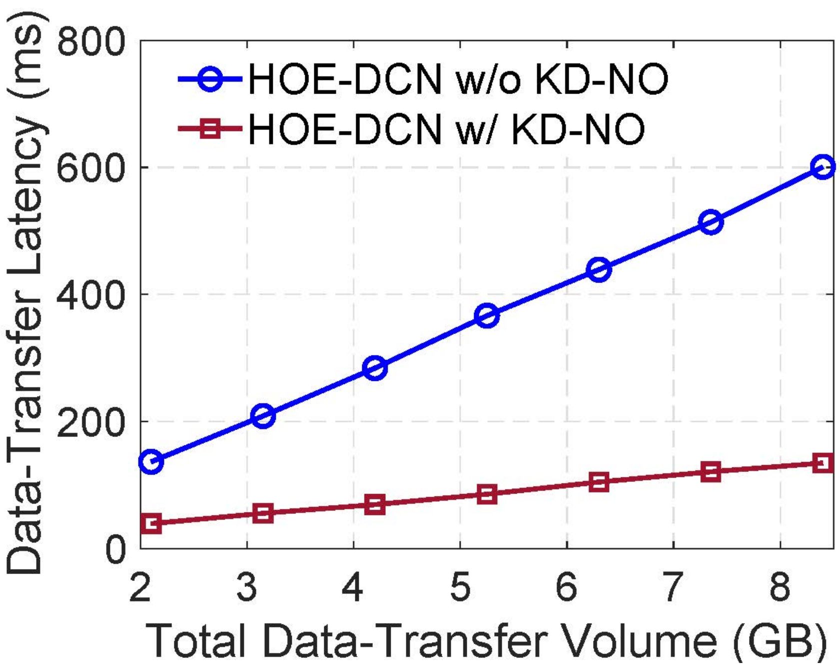

(a) Average data-transfer latency

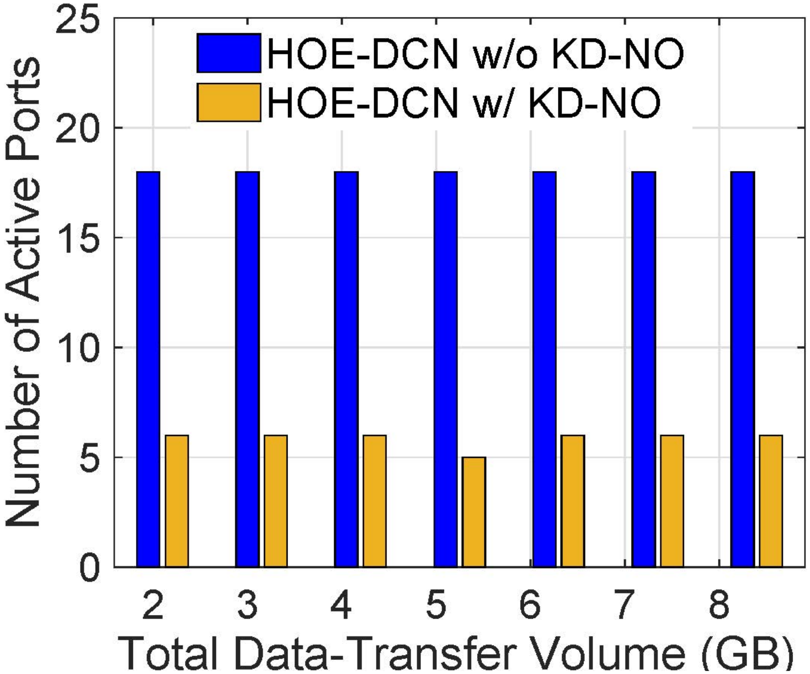

(b) Number of active switch ports

Conclusion

Knowledge-Defined Network Orchestration Mechanism

- Why? Challenges?
- Design a knowledge-defined NO-M in an H-O/E DCN system

Predictive Analytics and Decision Making

- Design three DL-based AI modules for VM migration, scaling and DCN configuration, respectively

System Implementation and Experiment Demonstration

- Implement a real network testbed to prototype the proposed design
- Achieve intelligent decision making and automatic management
- Improve the performance of service provisioning

Future Directions

- Incorporate more intelligent in the designed knowledge-defined NO-M
- Make the designed knowledge-defined NO-M application-driven