


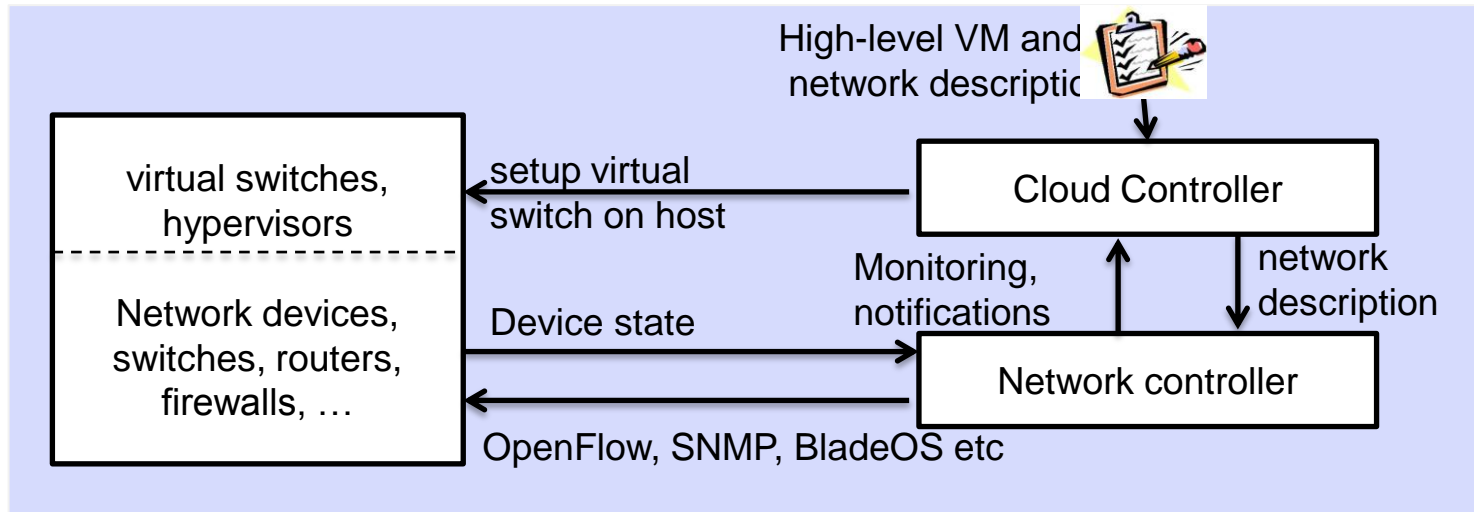
Dynamic Graph Query Primitives for SDN-based Cloud Network Management

Ramya Raghavendra, Jorge Lobo, Kang-Won Lee

IBM T. J. Watson Research Center




Network as a Service Model



- Provide rich set of services
 - isolation, custom addressing, service differentiation etc.
- Interact with variety of network devices
 - support multiple cloud platforms, device management protocols
- Introduces a network controller
 - network-aware VM placement, QoS support, real-time monitoring, diagnostics, management, security etc.

Graph Queries for Network Management

- Support for efficient queries on network graph, e.g., shortest path between two nodes
- Utilized in various network management operations



Network graphs can represent:

- physical network elements such as routers, switches, and servers
- virtual elements such as VMs and virtual switches
- logical elements such as people, processes, web pages, etc.
- and links between them

Algorithm support:

- **shortest path**
- spanning tree
- min flow
- ...


Cloud DCN Graphs Are Dynamic

Cloud DCNs may experience frequent changes

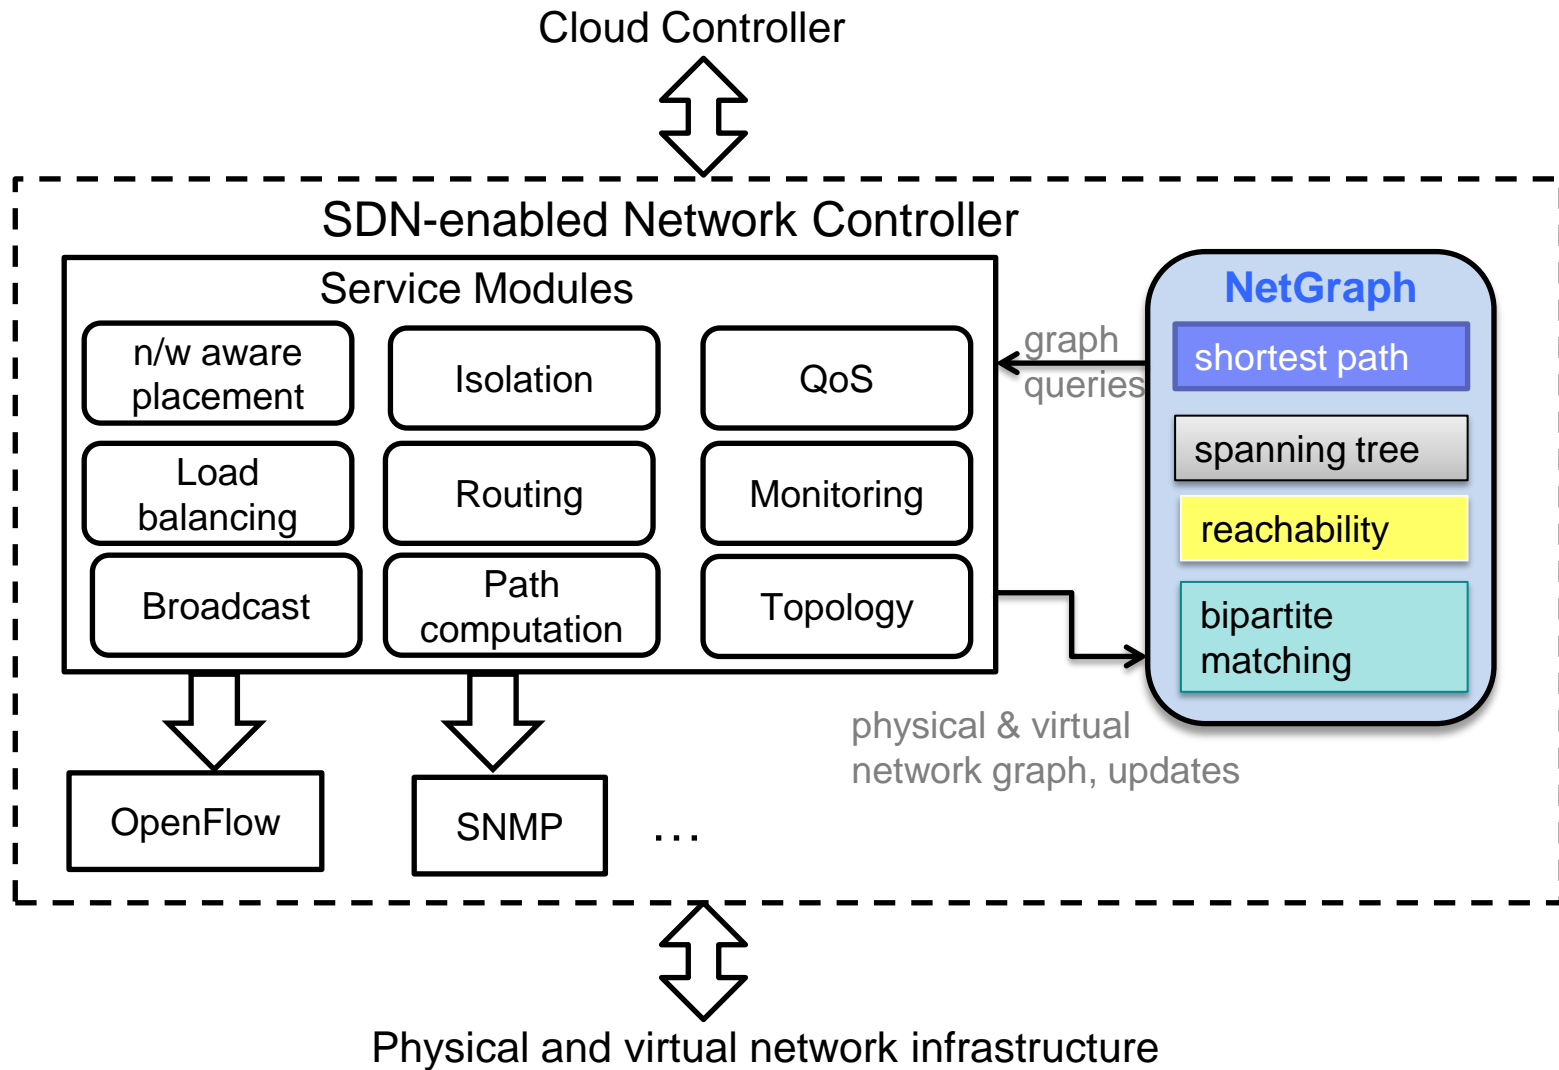
- Deployment of new VMs, removal of old ones
- Migration of VMs
- Layer 2/3 network reconfiguration

Graph updates

- edge/vertex insertion & deletion
- edge weight change (e.g. congestion level)



Interaction of NetGraph with Network Controller



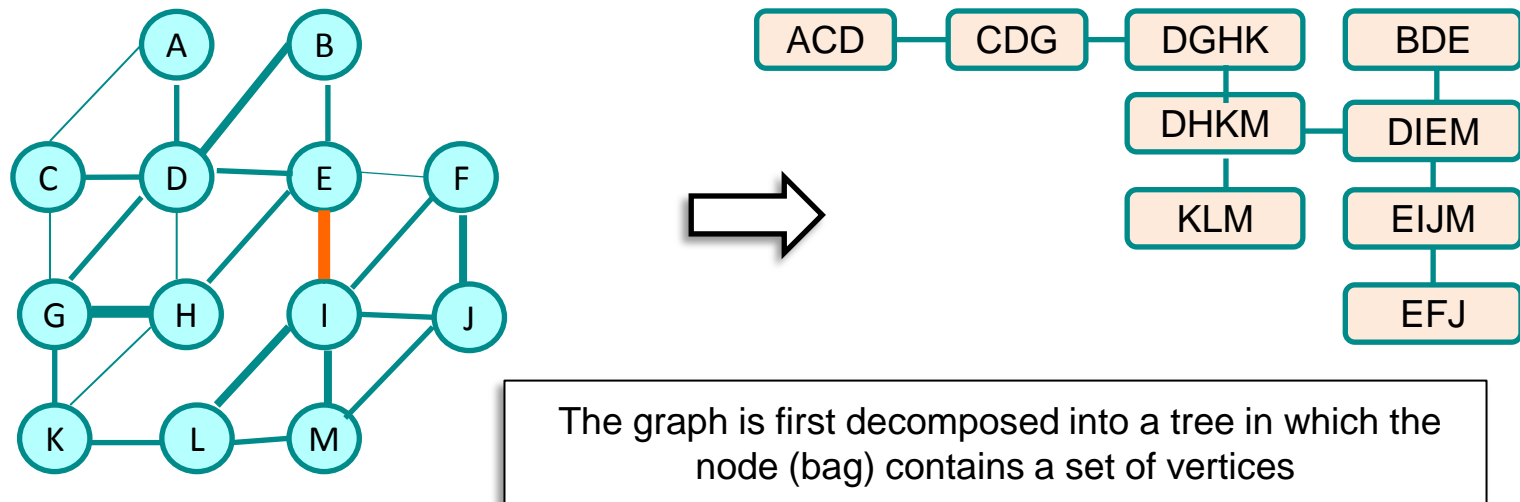
All Pair Shortest Path as Graph Query Example

Algorithm	Technique	Running time	Memory
Dijkstra's algorithm	Static algorithm. Shortest paths are recomputed on update	HIGH	HIGH
Dynamic shortest path algorithms*	Maintain subgraphs and recompute queries on affected subgraph	LOW	HIGH
Index-based shortest path algorithm (our approach [§])	Maintain tree-decomposition based index Update tree nodes that are affected by update	LOW	LOW

* G. Ramalingam and T. Reps. An incremental algorithm for a generalization of the shortest-path problem

§ Extensions to F. Wei. TEDI: efficient shortest path query answering on graphs published in SIGMOD 2010

Index-based shortest path computation (our approach)



Utilize tree decomposition-based shortest paths method:

- Shortest path search can be executed in a bottom-up manner
- Query time is decided by the height and the bag cardinality of the tree, instead of the size of the graph
- Updates are restricted to the bags containing the edge and bags affected by edge update

Interfaces

Maintain network graph

- **Network topology: nodes and links**
- Graph **class** with Vertices **and** Edges
- **Topology updates**
(node1, node2, edgeweight)



Compute graph query

- **Node and edge computations**
CountDegree, CountNeighbors, etc.
- **Algorithmic functionalities**
ComputeMST(S), IsSubPath(P1, P2)
- **Shortest paths**
 - ComputeAPSP
 - ComputeSSSP(S)
 - ...

Experimental Evaluation


- Query update time
- Memory requirement

Data Center Network





Synthetic DCN topologies with 1024 nodes, 16 first level, 48 second level switch ports

Different topologies considered





Gnutella P2P Network



- Graphs represent connections in Gnutella P2P network during August of 2002
- The graphs vary from 8,111 to 22,686 nodes and 20,781 to 65,373 edges.

Internet AS Graph




Snapshots of the connections between Autonomous Systems (AS) taken from the University of Oregon Route Views Archive Project

The resulting graphs (AS 500, . . . ,AS 3000) have 500 to 3,000 nodes and 2,406 to 13,734 edges, with edge weights as high as 20,000

Discussion

- Networking tasks rely on **graph abstractions**
- Cloud DCNs experience **network dynamics**
- **NetGraph**: shared graph library for underlying graph operations
 - Receive topology updates and recompute graph queries
 - Implemented shortest path queries as a specific example
 - Scalable due to tree-decomposition based index maintenance
 - Architecture extensible to other graph queries



Questions?

Ramya Raghavendra
rraghav@us.ibm.com