

# Service Function Chain Placement in Cloud Data Center Networks: a Cooperative Multi-Agent Reinforcement Learning Approach

Lynn Gao, University of California, Irvine

Yutian Chen, California State University, Long Beach

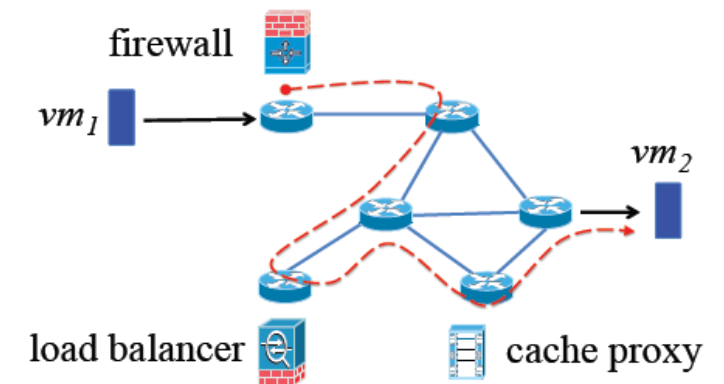
Bin Tang, California State University, Dominguez Hills

# Outline

- Introduction
  - Service Function Chain (SFC)
  - SFC Placement Problem
- Algorithms
  - Multi-agent Reinforcement Learning (MARL)
  - Dynamic Programming (DP)
  - Optimal Exhaustive (Optimal)
- Performance Evaluation
- Conclusions and Future Work

# Service Function Chaining (SFC)

- An SFC, consisting of a sequence of virtual network functions (VNFs) (i.e., firewalls and load balancers), is an effective service provisioning technique in cloud data centers
- Virtual machine (VM) cloud traffic traverses the chain of VNFs in a order to achieve security and performance guarantees
- An example of an SFC of three VNFs:
  - First, the firewall blocks malicious traffic
  - Next, the load balancer avoids network congestion
  - Finally, the cache proxy caches network packets for data access by other cloud

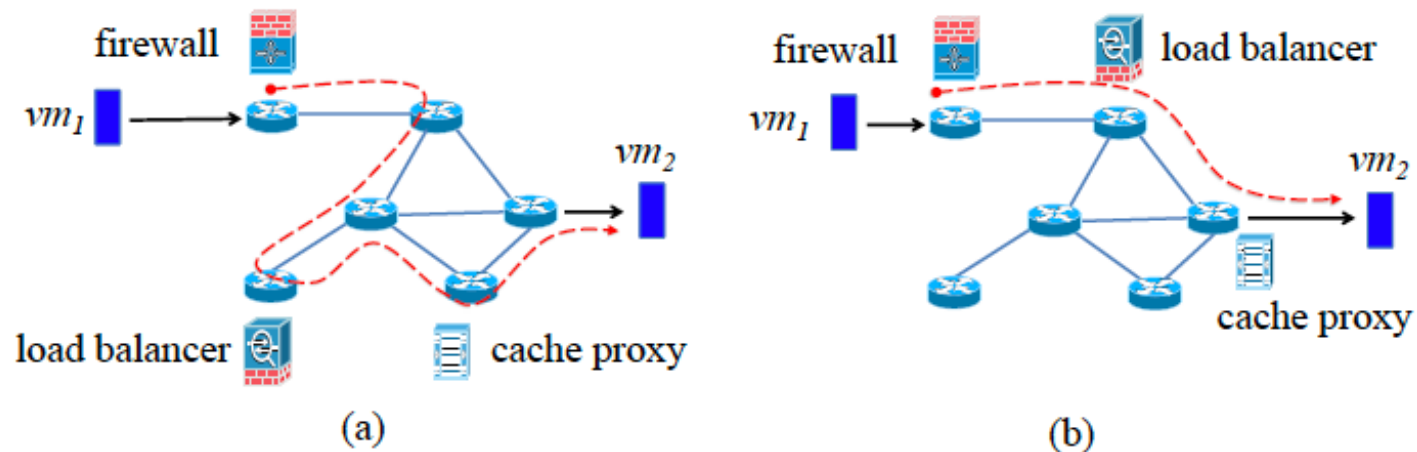


# SFC Placement Problem

- As cloud network resources such as bandwidth and energy are limited in a cloud data center network, one important task for the cloud operator is to install the VNFs at the right locations inside the network to optimize the cloud traffic or user-perceived VM communication delay.
- **SFC Placement Problem:** Given a VM flow in the cloud data center and an SFC, how to place the VNFs of the SFC inside the cloud data center to minimize the communication cost (i.e., network traffic or delay) of the VM flow?

# An Example of SFC Placement

- A small cloud data center of six switches
- One VM communication flow ( $vm_1$ ,  $vm_2$ )
- SFC placement (a): traversing six switches, with six network hops
- SFC placement (b): traversing three switches, with two network hops
- (b) is more network-efficient than (a)



# SFC Placement is Equivalent to $k$ -Stroll Problem


- **$k$ -Stroll Problem**

- **Given**

- a weighted undirected graph
- source  $s$  and destination  $t$ ,
- an integer  $k$

- **Goal:** find an  $s$ - $t$  path or walk (i.e., a stroll) of minimum length that visits at least  $k$  distinct nodes excluding  $s$  and  $t$ .

- $k$ -stroll is NP-hard

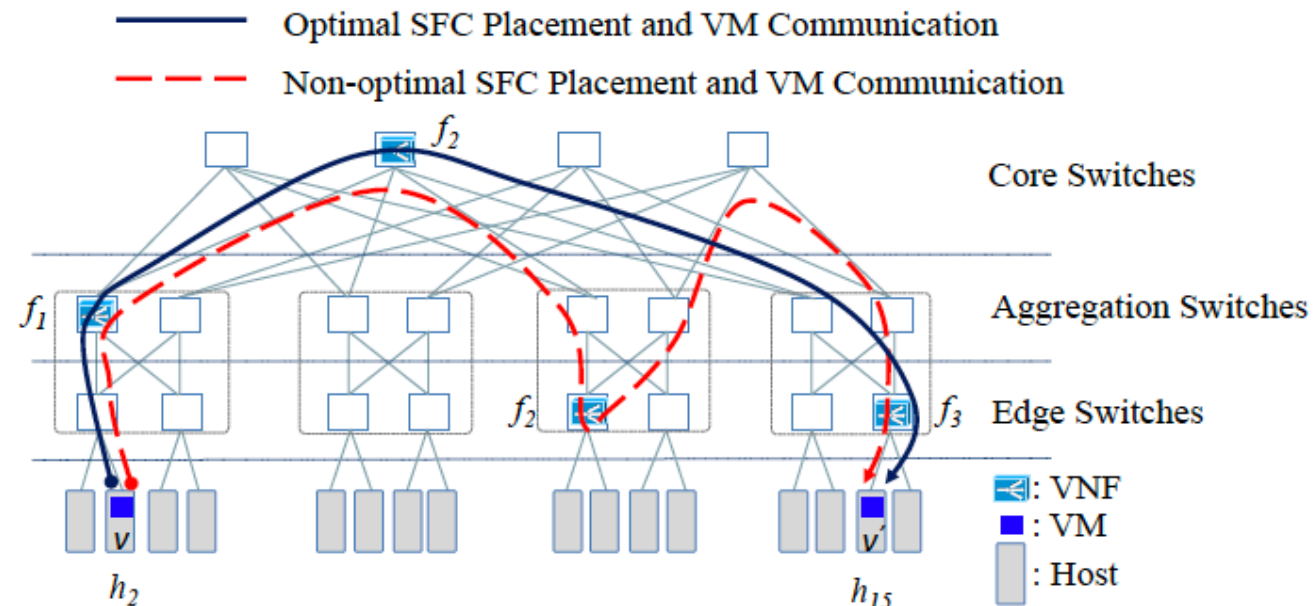


Optimal 2-stroll between  $s$  and  $t$ :  
 $s, D, t, C, t$  with cost 6

# SFC Placement in Fat-tree Data Centers

- One pair of communicating VMs ( $v, v'$ ),  $v$  is at host  $s(v)$  and  $v'$  at  $s(v')$
- SFC with  $n$  VNFs:  $f_1, f_2, \dots, \text{and } f_n$
- *Place the  $n$  VNFs to minimize communication cost:*

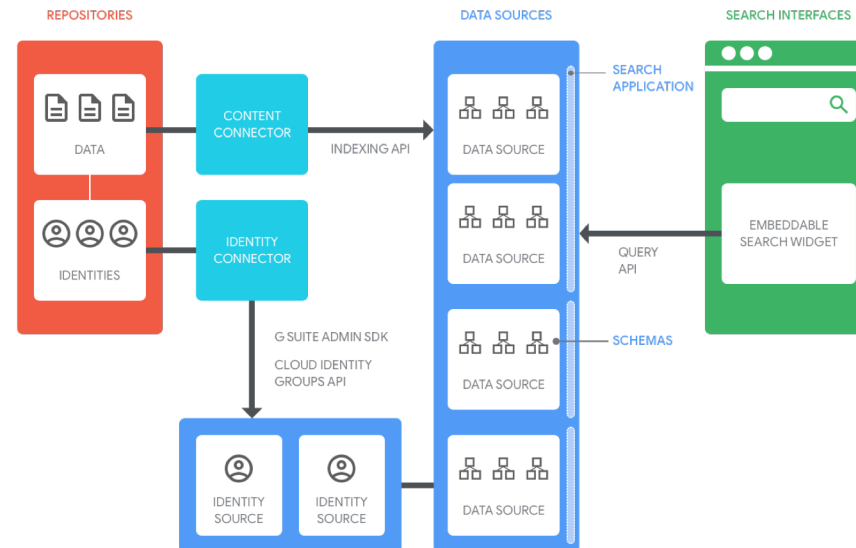
$$C_c(p) = \sum_{j=1}^{n-1} c(p(j), p(j+1)) + (c(s(v), p(1)) + c(p(n), s(v'))).$$



# SFC Placement in Fat-tree Data Centers

- One pair of communicating VMs ( $v, v'$ ),  $v$  is at host  $s(v)$  and  $v'$  at  $s(v')$
- SFC with  $n$  VNFs:  $f_1, f_2, \dots$ , and  $f_n$
- Place the  $n$  VNFs to minimize communication cost:

$$C_c(p) = \sum_{j=1}^{n-1} c(p(j), p(j+1)) + (c(s(v_i), p(1)) + c(p(n), s(v'_i))).$$

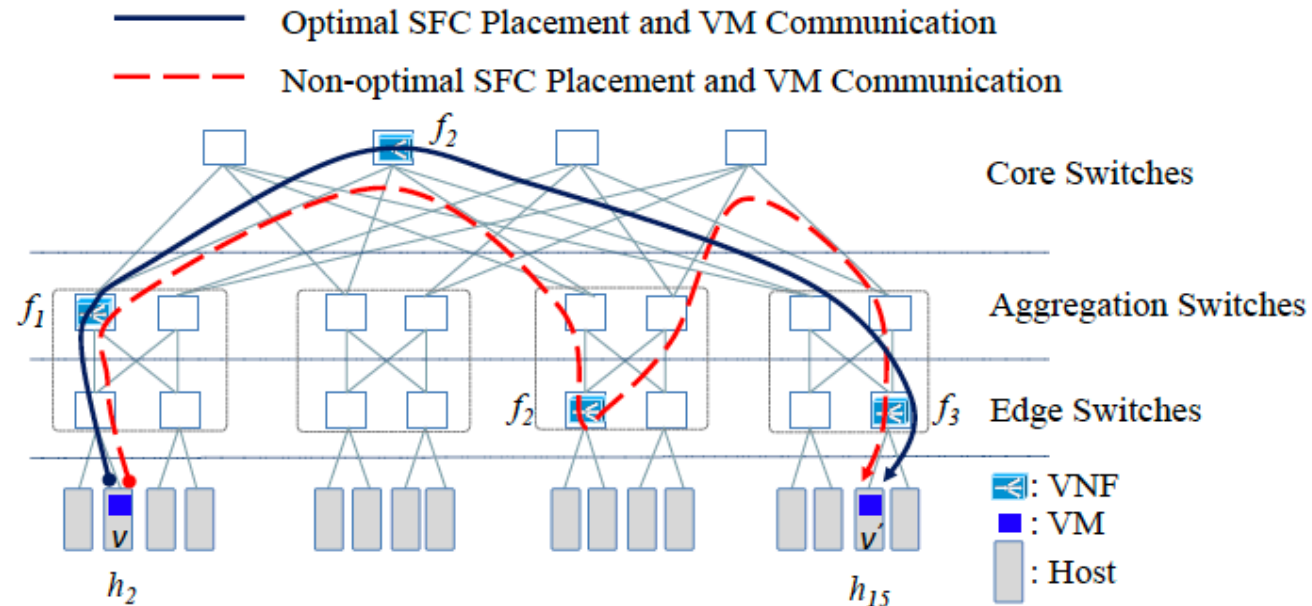




# SFC Placement in Fat-tree Data Centers

- One pair of communicating VMs ( $v, v'$ ),  $v$  is at host  $s(v)$  and  $v'$  at  $s(v')$
- SFC with  $n$  VNFs:  $f_1, f_2, \dots$ , and  $f_n$
- *Place the  $n$  VNFs to minimize communication cost:*

$$C_c(p) = \sum_{j=1}^{n-1} c(p(j), p(j+1)) + (c(s(v), p(1)) + c(p(n), s(v'))).$$



# Algorithms for SFC Placement

- MARL algorithm
- DP algorithm
- Optimal algorithm

# MARL Algorithm

- Uses reinforcement learning and Q-learning
  - Finite set of states  $S$
  - Finite set of actions  $A$
  - State transition function  $t$
  - Reward function  $r$
  - Value based methods
  - Agents explore the world by taking actions based on the state of the world

$$Q(s, a) \leftarrow (1 - \alpha) \cdot Q(s, a) + \alpha \cdot [r(s, a) + \gamma \cdot \max_b Q(t, b)]$$

# MARL Algorithm - Action Selection Rule

- Exploitation: move to node  $t$  where

$$t = \operatorname{argmax}_{u \in U} \left\{ \frac{[Q(s,u)]^\delta}{[w(s,u)]^\beta} \right\}$$

- Exploration: move to node  $t$  following distribution

$$p(s, t) = \frac{[Q(s,t)]^\delta / [w(s,t)]^\beta}{\sum_{u \in U} [Q(s,u)]^\delta / [w(s,u)]^\beta}$$

# MARL Algorithm

- All the  $m$  agents start at host  $s(v)$
- Takes place in episodes (iterations)
- Each episode has two stages:
  - Stage 1: each agent finds its own  $k$ -stroll following action rules
  - Stage 2: finds the smallest cost  $k$ -stroll and updates reward and  $Q$ -value for edges in this  $k$ -stroll
- Terminate after specified number of iterations or within proximity to compared DP or Optimal algorithms
- Time complexity:  $O(Nmk)$ 
  - $N$ : number of episodes
  - $M$ : number of agents
  - $k$ : number of VNFs

**Algorithm 1** MARL Algorithm for SFC Placement.

**Input:** A data center graph  $G(V = V_s \cup V_h, E)$ ,  $s(v_1)$ ,  $s(v'_1)$ , and an SFC  $(f_1, f_2, \dots, f_n)$ .

**Output:** A  $k$ -stroll from  $s(v_1)$  to  $s(v'_1)$ ; that is, a switch  $p(j) \in V_s$  to place each of the  $k$  VNFs  $f_j \in \mathcal{F}$  and the cost  $C_c(p)$  of the  $k$ -stroll.

**Notations:**  $i$ : index for switches;  $j$ : index for agents;

$U_j$ : the set of nodes unvisited by agent  $j$ , initially  $U_j = V_s$ , the set of switches;

$L_j$ : the path taken by agent  $j$ , initially empty;

$l_j$ : the length of  $L_j$ , initially zero;

$r_j$ : the node where agent  $j$  is located currently;

$Q(u, v)$ :  $Q$ -value of edge  $(u, v)$ , initially  $\frac{|E|}{|V| \cdot \sum_{(u,v) \in E} w(u,v)}$ ;

$p$ : an array storing the distinct switches on  $s(v_1)$ - $s(v'_1)$  stroll;

$\alpha$ : learning rate,  $\alpha = 0.1$ ;

$\gamma$ : discount factor,  $\gamma = 0.3$ ;

$W$ : a constant value of 10 following [17];

```

1. Initially all the  $m$  agents are at host  $s(v)$ , i.e.,  $r_j = s(v), 1 \leq j \leq m$ ;
2. while (termination condition is not met)
3.   for ( $i = 1; i <= k; i++$ ) // Finding the  $k$  switches to place VNFs
4.     for ( $j = 1; j \leq m; j++$ ) // Agent  $j$ 
5.       Agent  $j$  decides the next node  $s_j$  to move to following action rule
         in Definition 1;
6.        $L_j = L_j \cup \{s_j\}$ ;
7.        $l_j = l_j + w(r_j, s_j)$ ;
8.        $Q(r_j, s_j) = (1 - \alpha) \cdot Q(r_j, s_j) + \alpha \cdot \gamma \cdot \max_{z \in U_j} Q(s_j, z)$ ; // Q-value
9.        $r_j = s_j$ ; // Agent  $j$  moves to switch  $s_j$ ;
10.       $U_j = U_j - \{s_j\}$ ; // Switches not yet visited by agent  $j$ 
11.    end for;
12.  end for;
13.  for ( $j = 1; j \leq m; j++$ ) // Agent  $j$  ends at destination host  $s(v')$ 
14.     $L_j = L_j \cup \{s(v')\}$ ;
15.     $l_j = l_j + w(r_j, s(v'))$ ;
16.     $Q(r_j, s_j) = (1 - \alpha) \cdot Q(r_j, s_j) + \alpha \cdot \gamma \cdot \max_{z \in U_k} Q(s_j, z)$ ; // Q-value
17.     $r_j = s(v')$ ;
18.  end for;
19.  Let  $j^* = \operatorname{argmin}_{1 \leq j \leq m} l_j$  be the agent with a  $k$ -stroll of smallest length;
20.  for (each edge  $(u, v) \in L_{j^*}$ )
21.     $r(u, v) = \frac{W}{l_{j^*}}$ ; // Update reward value  $r(u, v)$ ;
22.     $Q(u, v) \leftarrow (1 - \alpha) \cdot Q(u, v) + \alpha \cdot [r(u, v) + \gamma \cdot \max_b Q(v, b)]$ ; // Q-value
23.  end for;
24. end while;
25. RETURN The switch  $p(j) \in V_s$  to place VNF  $f_j \in \mathcal{F}$  and the cost  $C_c(p)$ .

```

Stage 1

Stage 2

# Dynamic Programming Algorithm

- Based on the observation that although finding the shortest stroll visiting  $k$  distinct nodes is NP-hard, finding the shortest stroll of  $k$  edges can be solved optimally and efficiently using DP
- Time complexity:  $O(k|V|^4)$

**Algorithm 2** A DP Algorithm for SFC Placement Problem.

**Input:** A complete graph  $G'(V', E')$ ,  $s(v)$ ,  $s(v')$ , and an SFC  $(f_1, f_2, \dots, f_k)$ .

**Output:** cost of an  $s(v)$ - $s(v')$  stroll in  $G'$  visiting at least  $k$  distinct switches.

**Notations:**  $e$ : index for edges;  $i$ : index for switches;

$c(u, s(v'), e)$ : cost of a  $u$ - $s(v')$  stroll with  $e$  edges, initially  $+\infty$ ;

$successor(u, s(v'), e)$ :  $u$ 's successor in a  $u$ - $s(v')$  stroll with  $e$  edges, initially  $-1$ ;

$r$ : number of edges needed on  $s(v)$ - $s(v')$  stroll, initially  $k+1$ ;

$p$ : an array storing distinct switches on  $s(v)$ - $s(v')$  stroll;

$num$ : the number of distinct switches in  $p$ ;

$found$ : true if it has found a  $s(v)$ - $s(v')$  stroll with at least  $k$  distinct switches, initially false;

1.  $V' = \{u_1, \dots, u_{|V'|}\}$ , let  $u_a = s(v)$  and  $u_{|V'|} = s(v')$ ;

2.  $\forall u_i, u_j \in V'$  with  $i \neq j$ ,  $c(u_i, u_j, 1) = c_{u_i, u_j}$ ,  $successor(u_i, u_j, 1) = u_j$ ,

$successor(u_j, u_i, 1) = u_i$ ;  $\forall u_i \in V'$ ,  $c(u_i, u_i, 1) = +\infty$ ,  $successor(u_i, u_i, 1) = -1$ ;

3. **while** ( $\neg found$ )

4.   **for** ( $e = 2; e \leq r; e++$ )   // edges in  $u_i$ - $s(v')$  stroll

5.       **for** ( $i = 1; i \leq |V'| - 1; i++$ )                                 // node  $u_i$

6.           **for** (each  $u$ ,  $u \neq u_i \wedge u \neq s(v') \wedge u_i \neq successor(u, s(v'), e - 1)$ )

7.               **if** ( $c(u_i, s(v'), e) > c_{u_i, u} + c(u, s(v'), e - 1)$ )

8.                    $c(u_i, s(v'), e) = c_{u_i, u} + c(u, s(v'), e - 1)$ ;

9.                    $successor(u_i, s(v'), e) = u$ ;

10.               **end if**;

11.    $num = 0$ ;  $p = \phi$  (empty set),  $e--$ ;

12.    $b = successor(s(v), s(v'), e)$ ;

13.   **while** ( $e > 1$ )

14.       **if** ( $b \neq s(v) \wedge b \neq s(v') \wedge b \notin p$ )

15.            $p[num] = b$ ;  $num++$ ;

16.       **end if**;

17.        $e--$ ;

18.        $b = successor(b, s(v'), e)$ ;

19.   **end while**;

20.   **if** ( $num < k$ )  $r++$ ; // less than  $k$  distinct switches

21.   **else**  $found = true$ ;

22. **end while**;

23. Place  $f_1, \dots, f_k$  on the first  $k$  switches stored in  $p$ ;

24. **RETURN**  $c(s(v), s(v'), r)$ .

# Optimal Algorithm

**Algorithm 3** Exhaustive and Optimal SFC Placement.



**Input:** A data center graph  $G(V, E)$ ,  $s(v)$  and  $s(v')$ , and an SFC  $(f_1, f_2, \dots, f_k)$ .

**Output:** A VNF placement  $p$  and the total cost  $C_c(p)$ .

1.  $C_c(p) = +\infty$ ;
2. Among all  $|V_s| \cdot (|V_s| - 1) \cdot \dots \cdot (|V_s| - k + 1)$  SFC placements, find  $p$  that gives the minimum cost  $C_c(p)$ ;
3. **RETURN**  $p$  and  $C_c(p)$ .


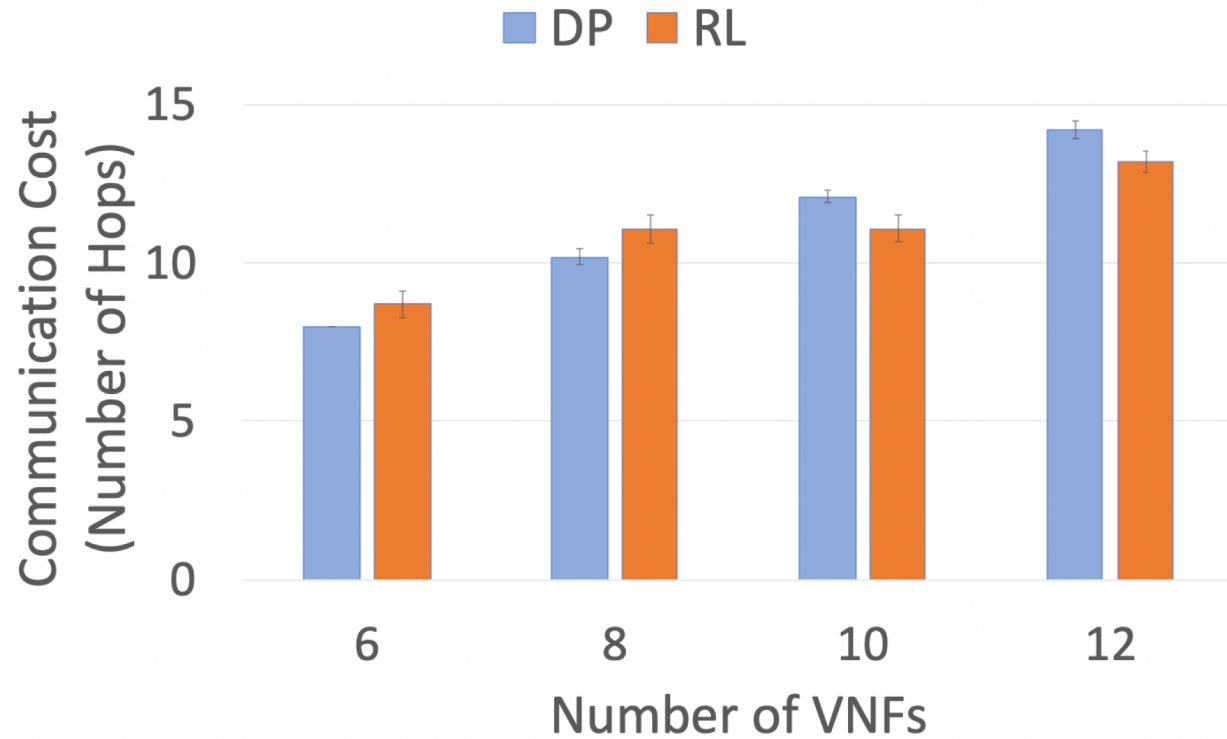
- Enumerates all the SFC placements and finds the one with minimum cost
- Time complexity:  $O(|V|^k)$

# Performance Evaluation – K=4 Fat Trees

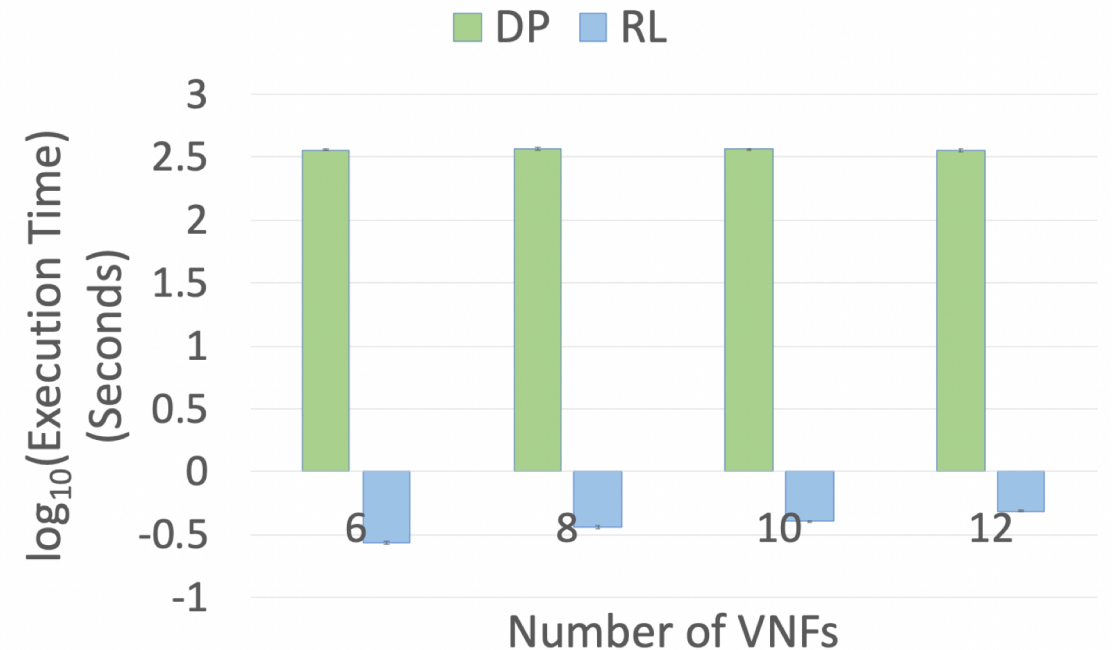





# Performance Evaluation – K=8 Fat Trees



# Performance Evaluation – K=8 Fat Trees



$$|V| = \frac{5}{4} * K^2$$

Algorithm	Time Complexity
MARL Algorithm	$O(N * K^3 * k)$
Dynamic Programming	$O(k * K^8)$

# Performance Evaluation - # of Agents in MARL

Table 2: Varying number of agents  $m$  in RL. Here,  $k = 10$  and  $K = 6$ .

Number of Agents $m$	1	5	10	15	20
Communication Cost (number of hops)	23.3	17.6	14.7	13.7	13.4
Execution Time (seconds)	65.54	27.72	9.07	2.70	0.22
Number of Iterations	169.6	75.05	22.25	10.7	3.1

# Conclusions and Future Work

- Novelties
  - The core of the the SFC placement problem is the k-stroll problem.
  - Solved k-stroll with RL with good performance
  - Under the k-stroll modeling of SFC placement, deep reinforcement learning is not necessary.
- Future
  - Convergence and convergence speed of our RL algorithm.
  - Balance cooperation with competition to find a near-optimal SFC placement efficiently in a rational and game theoretical environment.