On Orchestrating Virtual Network Functions

Raouf BOUTABA

David R. Cheriton School of Computer Science
University of Waterloo

Joint work with Md. F. Bari, S. R. Chowdhury, and R. Ahmed

CNSM 2015, Barcelona, Spain, November 9, 2015

Outline

- Background
- Motivation
- Problem Statement
- Proposed Solutions
- Evaluation
- Conclusion

Background

- Middleboxes have become an integral part of modern networks
- Traditional hardware middleboxes are:
  - Expensive
  - Proprietary
  - Vertically integrated
- Difficult to compose Service Function Chains
  - In a service function chain traffic flows through an ordered sequence of middleboxes
    - Example:
      - Firewall → IDS → Proxy
      - Traffic Analyzer → Firewall → Video Optimizer

Network Function Virtualization

- Packet processing by software middleboxes or Virtualized Network Functions (VNFs)
- VNFs can be deployed on commodity servers
  - E.g., x86 based systems
- VNFs are no longer constrained to fixed network locations
- Service function chains can be composed on the fly
- These features facilitate network optimization

Outline

- Background
- Motivation
- Problem Statement
- Proposed Solutions
- Evaluation
- Conclusion

Motivation

- Hardware middleboxes are:
  - Provisioned for peak traffic
  - Attached at fixed locations in the network
  - Traffic is steered through these attachment points
    - Error-prone configuration process
- In case of VNFs:
  - VNF can be allocated based on demand
  - Network attachment points are not fixed
  - VNFs can be deployed on any compute server
  - Both VNF location and traffic steering can be optimized
    - Autonomic configuration
    - Reduction in CAPEX and OPEX

Outline

- Background
- Motivation
- Problem Statement
- Proposed Solutions
- Evaluation
- Conclusion

Problem Statement

- Given
  - A set of VNF chain requests
  - Physical infrastructure status

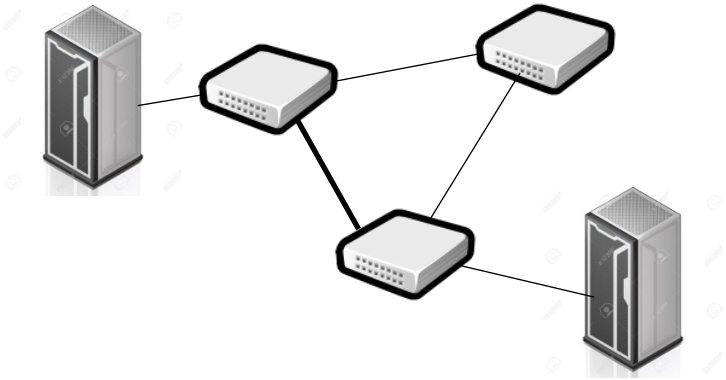

Problem Statement (cont.)

- We need to decide
  - How many VNF instances (VM, container) to deploy?
  - Where to place them?
  - Which VNF (from chain) should be assigned where?
  - How to route traffic between the VNFs?

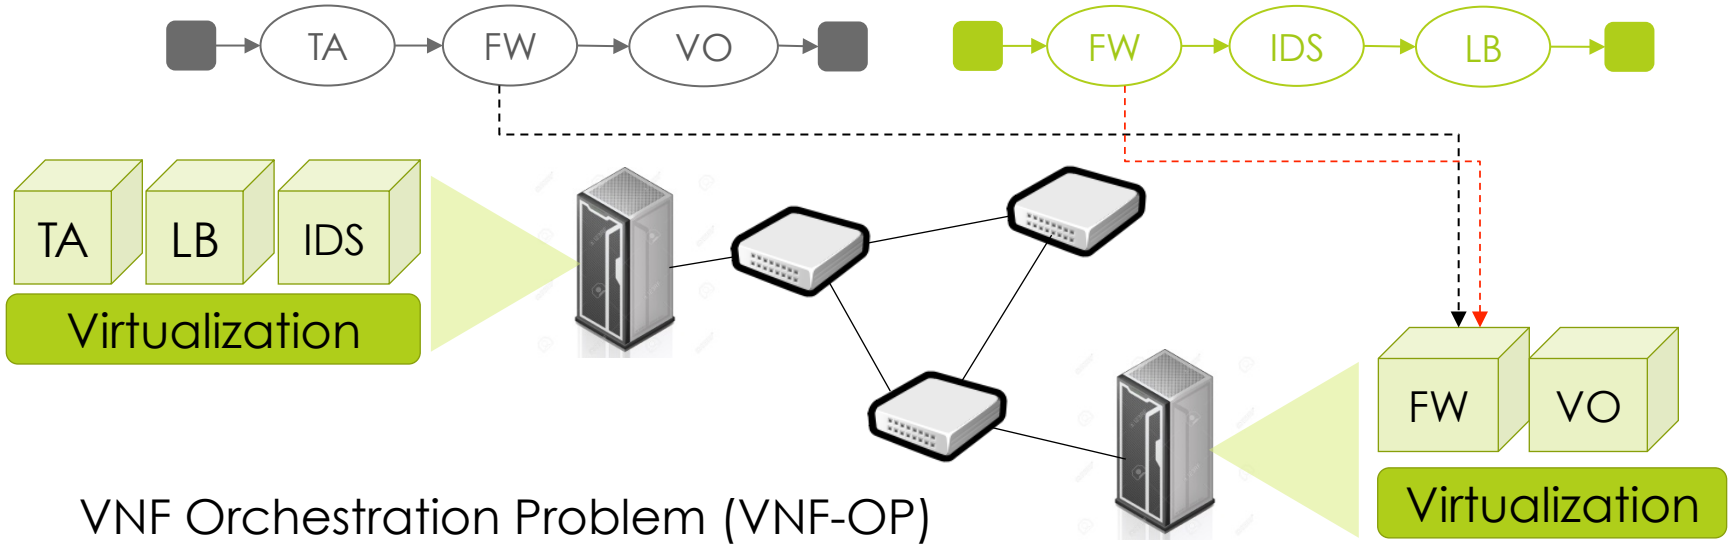

VNF Orchestration Problem (VNF-OP)

Outline

- Background
- Motivation
- Problem Statement
- Proposed Solutions
- Evaluation
- Conclusion

Solution Approach

- VNF-OP is a combination of three problems:
  - Allocating resource for VMs/containers
  - Assigning chain VNFs to these VMs
  - Finding routing paths for the chains
- Mathematical formulation is difficult:
  - Joint optimization results in quadratic constraints
  - Takes a long time to solve even for small problem instances

Optimal Solution

- Proposed approach:
  - Transform physical network

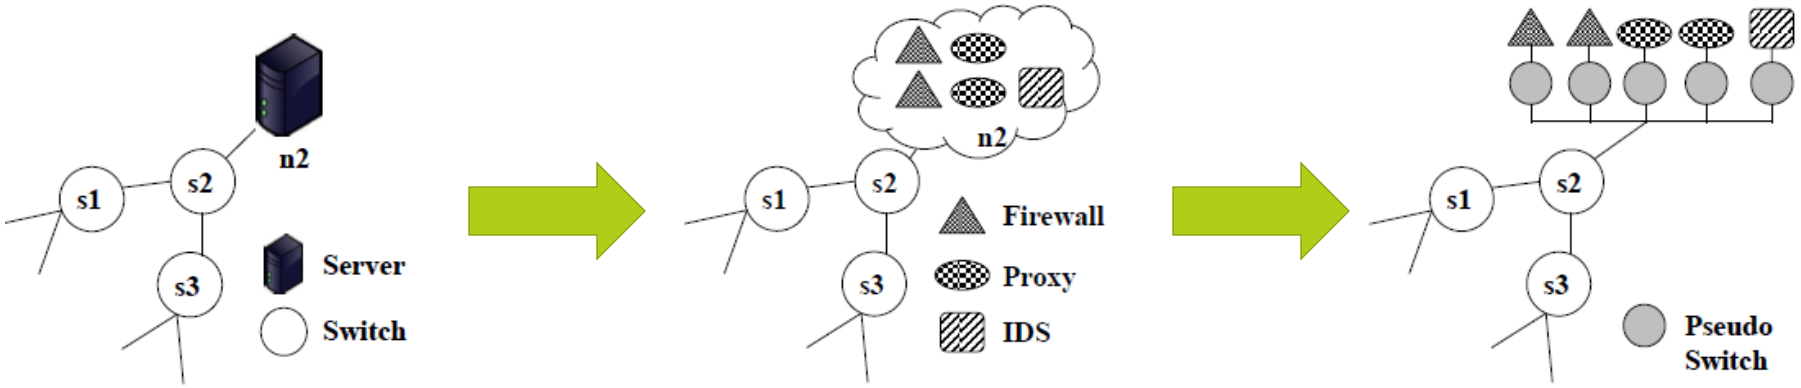

  - Can be formulated as an ILP
  - Much faster than implementation with quadratic constraints

Heuristic

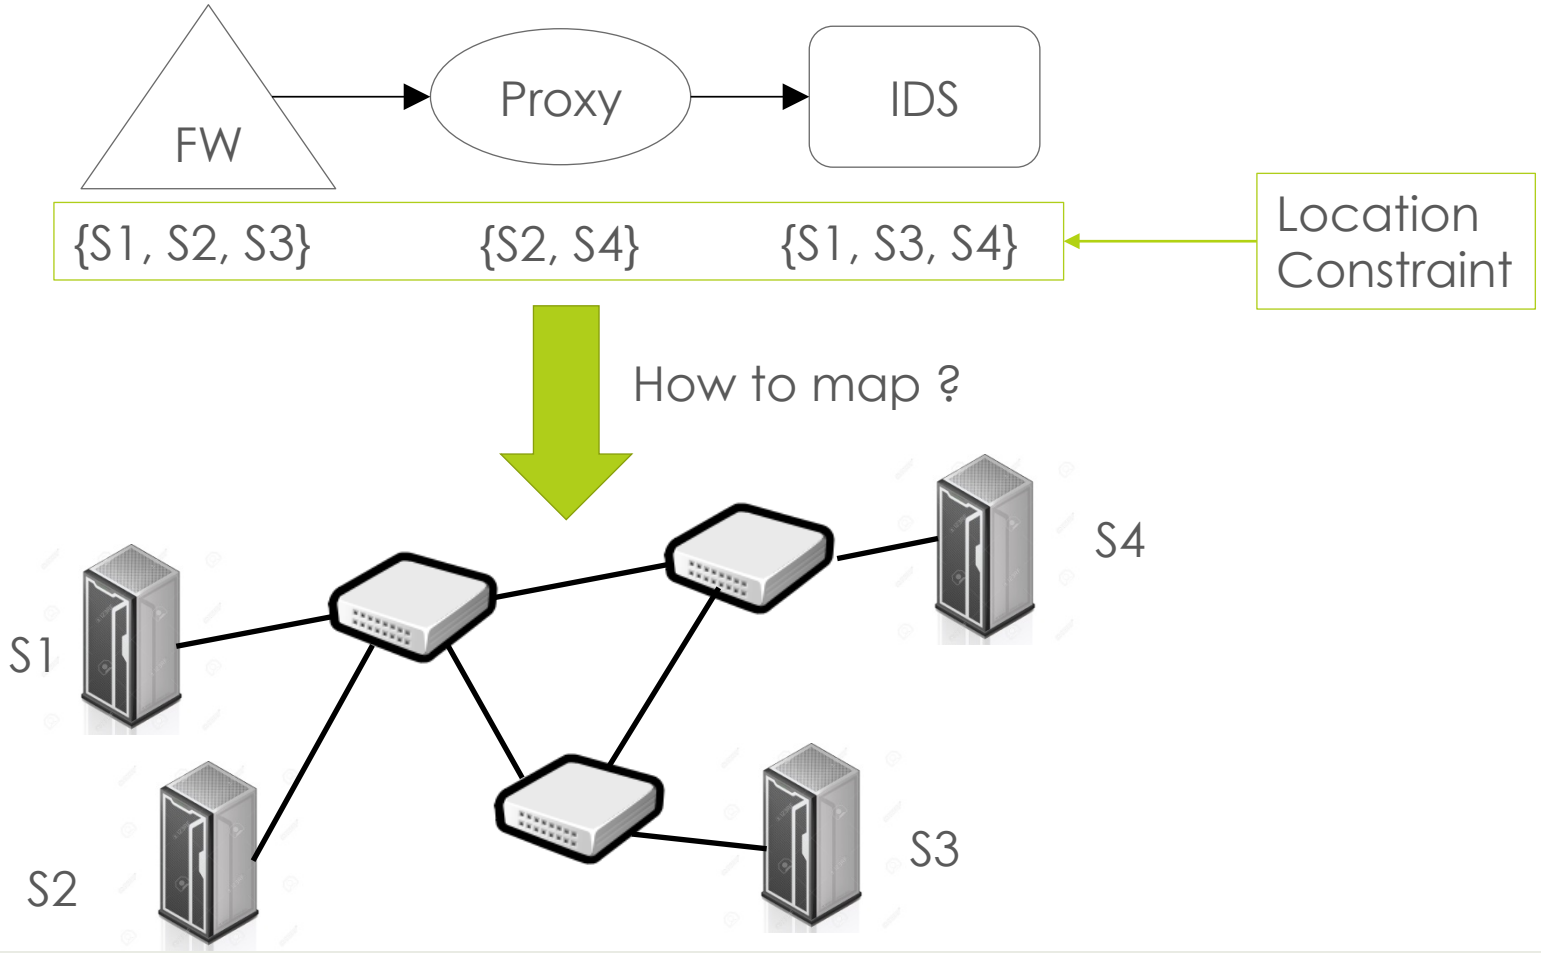
FW
Proxy
IDS
{S1, S2, S3}
{S2, S4}
{S1, S3, S4}
Location Constraint
How to map ?
S1
S2
S3
S4

Heuristic (cont.)

- Create a multi stage graph (one stage for each VNF) as follows:

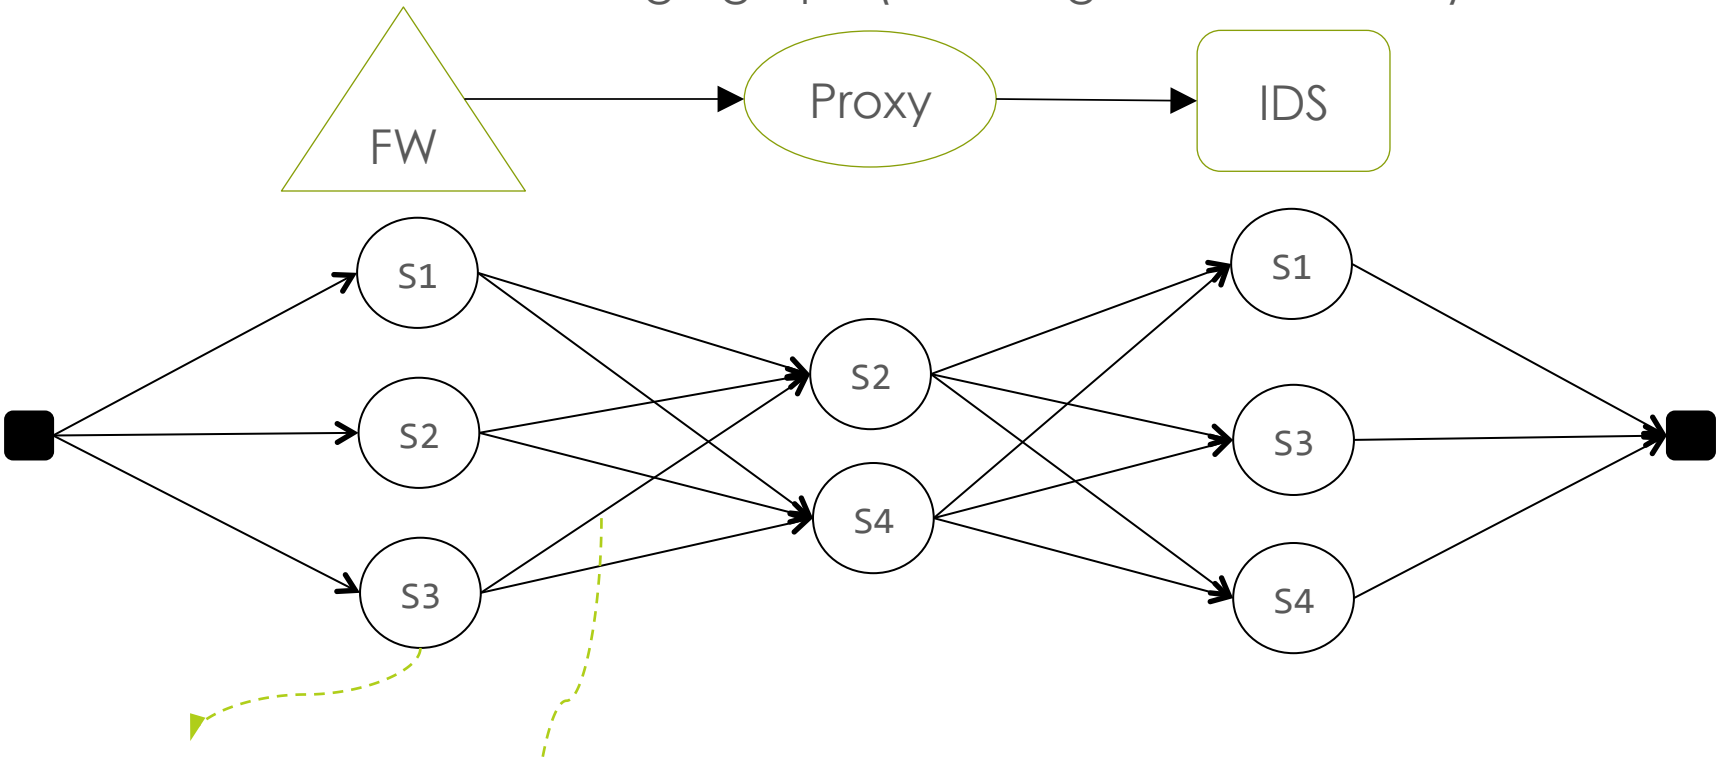

- Objective: Find a path from left most to right most stage that has minimum cost
  - Select exactly one node at each stage

Heuristic (cont.)

- Create a multi stage graph (one stage for each VNF) as follows:

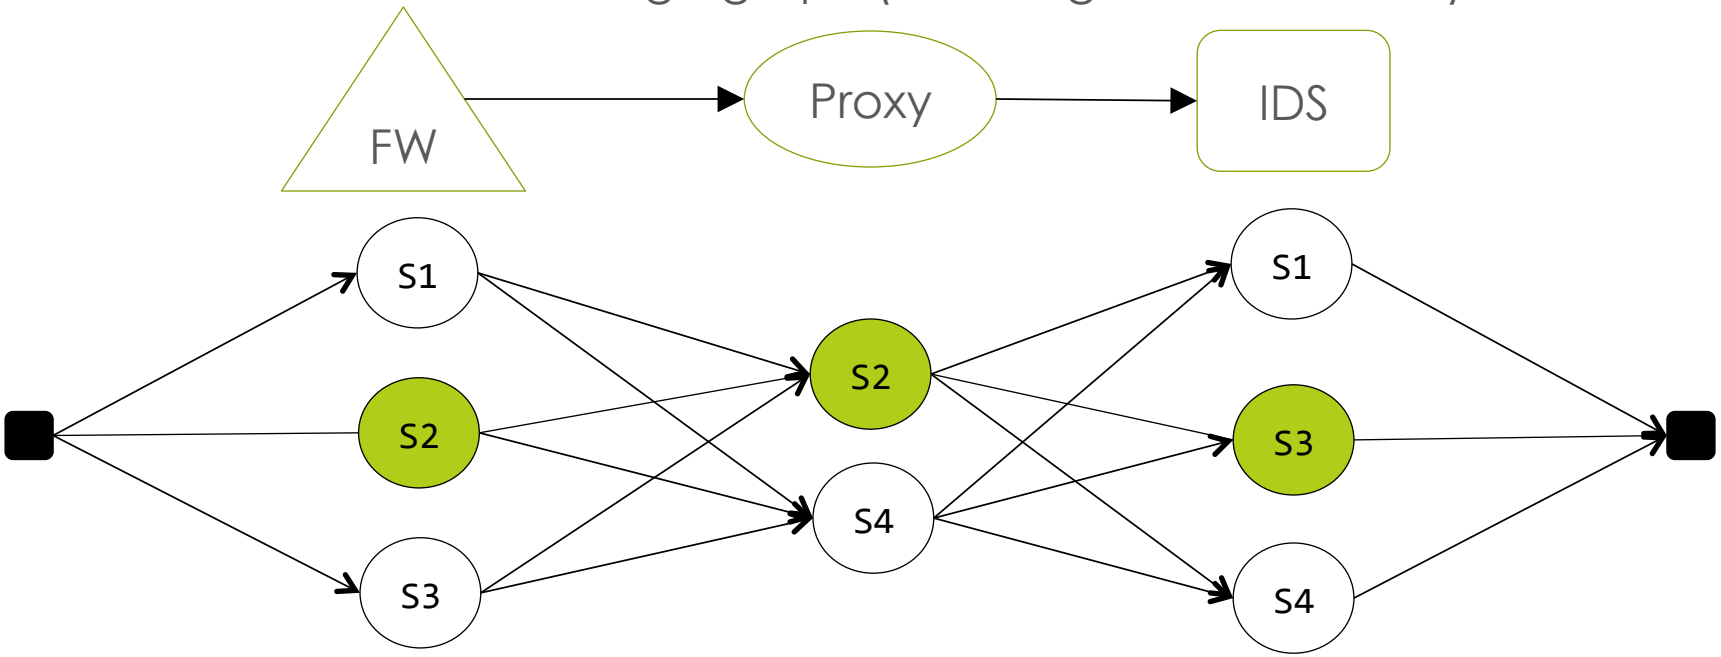

- Objective: Find a path from left most to right most stage that has minimum cost
  - Select exactly one node at each stage
- Similar to assigning tags to an unknown sequence of observations based on known cost function

Solution: Use Viterbi algorithm to find the minimum cost path

- Viterbi is widely used in pattern recognition to assign tags to unknown sequenceS of observations.

Outline

- Background
- Motivation
- Problem Statement
- Proposed Solutions
- Evaluation
- Conclusion

Evaluation Setup

- Two network topologies:
  - Internet2 research network (12 nodes, 15 links)
  - A university data center topology (23 nodes, 42 links)
- Server energy consumption data collected from Intel datasheet
- Hardware middlebox energy consumption data collected from a manufacturer
- Traffic traces
  - Traffic matrix from Internet2 network
  - Data center traffic trace from [1]

[1] T. Benson et al. Network traffic characteristics of data centers in the wild. ACM IMC '10

Evaluation Results

- Hardware Middlebox vs. VNF

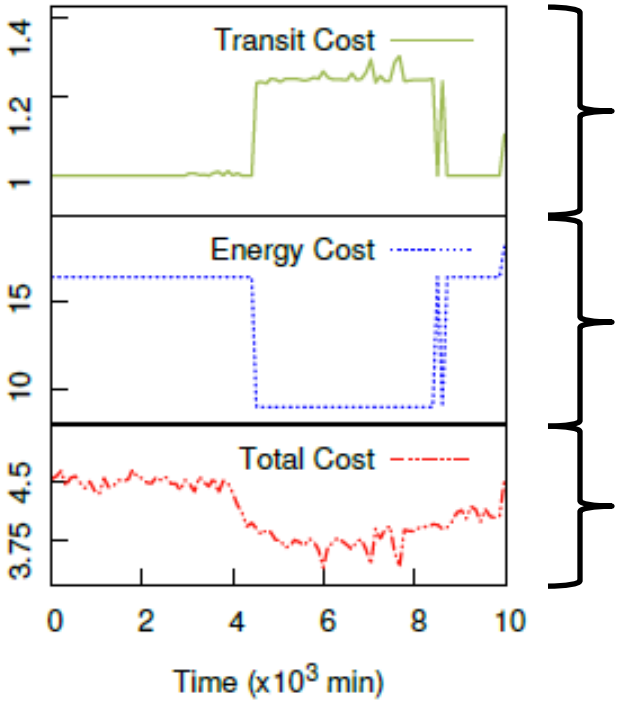

Computed Ratios

$$\frac{Transit\ Cost\ (hardware)}{Transit\ Cost\ (VNF)}$$

$$\frac{Energy\ Cost\ (Hardware)}{Energy\ Cost\ (VNF)}$$

$$\frac{Total\ Cost\ (Hardware)}{Total\ Cost\ (VNF)}$$

$$Total\ cost = Transit\ Cost + Energy\ Cost$$

Evaluation Results (cont.)

- Solution Quality: Heuristic vs. Optimal

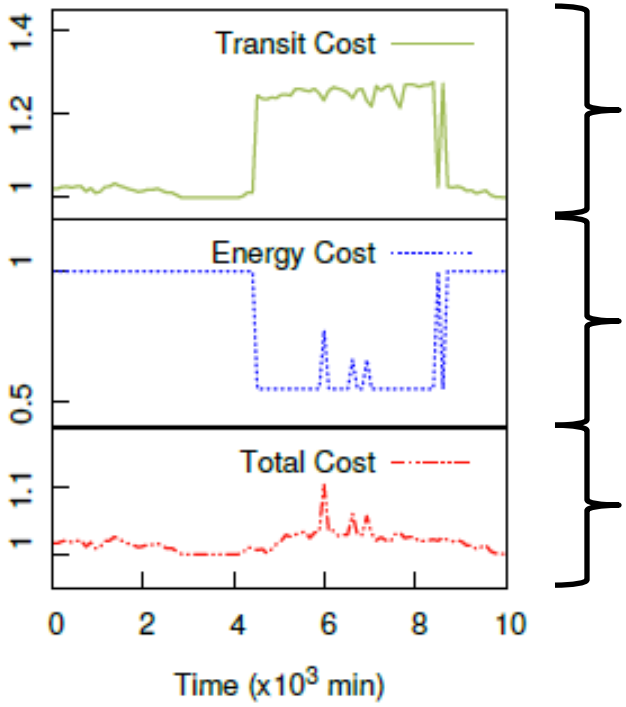

Computed Ratios

$$\frac{Transit\ Cost\ (Heuristic)}{Transit\ Cost\ (Optimal)}$$

$$\frac{Energy\ Cost\ (Heuristic)}{Energy\ Cost\ (Optimal)}$$

$$\frac{Total\ Cost\ (Heuristic)}{Total\ Cost\ (Optimal)}$$

$$Total\ cost = Transit\ Cost + Energy\ Cost$$

Summary of Results

- 4x OPEX reduction by VNFs compared to hardware middleboxes
- Heuristic produces solutions that are within 1.3x the optimal solution
- Heuristic is faster then the optimal
  - 65x for Internet2
  - 3500x for DC network

Outline

- Background
- Motivation
- Problem Statement
- Proposed Solutions
- Evaluation
- Conclusion

Conclusion

- Proposed two solutions to the VNF orchestration problem:
  - ILP formulation that can be easily implemented in CPLEX
  - Heuristic for larger networks
- Experiments suggest that NFV can reduce OPEX by a factor of 4 over hardware middleboxes
- Heuristic produces near-optimal solutions in a fraction of the time required for the CPLEX program

Questions ?