Optimization of Relay Placement for Scalable Virtual Private LAN Services

Mohammad Borhani, Ioannis Avgouleas, Andrei Gurtov

1st ACM SIGCOMM Workshop on Future of Internet Routing & Addressing (FIRA)

LINKÖPING UNIVERSITY

SIGCOMM 2022
Amsterdam

Agenda

Introduction

- VPLS
- HIPLS

Problem Statement

- Full-mesh problem/Relaying
- Relay Placement Problem (RPP)

Proposed Solution

- Approximation Algorithm

Results

Future Works/ Conclusions

Introduction

LINKÖPING UNIVERSITY

Introduction

- Daily, massive amounts of heterogeneous devices are connected to the Internet
- Industry 4.0: Smart sensors/valve controls
- Outbreak of Coronavirus in 2020 [1]
  - Production facilities shut down
  - Remote works, relying on the Internet
- Such devices are often hard to patch/contain exploitable vulnerabilities: Mirai
- Hide such devices from the Internet while allowing remote data administration and updates: VPLS

Introduction- Virtual Private LAN Service (VPLS)

- Over a public network, VPLS provides Layer 2 multipoint Ethernet service
  - Support for legacy protocols [2]
  - Single broadcast domain
  - Cost-effective installation/maintenance

- HIPLS: HIP establishes/maintains/closes IPsec ESP tunnels in VPLS

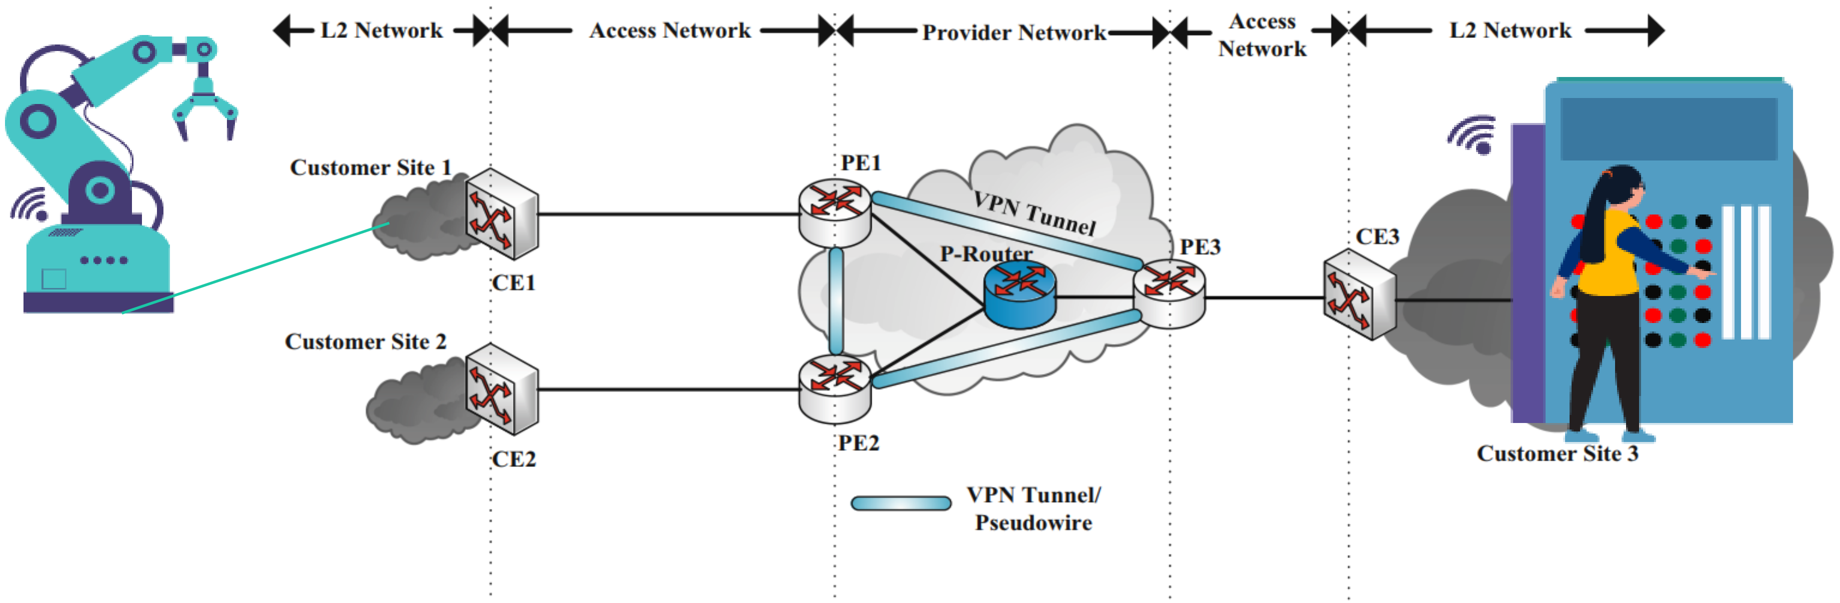

Problem Statement

LINKÖPING
UNIVERSITY

Full-mesh problem

- To interconnect PEs, the provider creates a full-mesh of HIP tunnels
- This reachability model forces routing tables in PEs to grow large
- In a full-mesh setup, each PE should install customers' routes (prefix) to ensure connectivity
- Supporting numerous CEs increases the size of the PE routing table exponentially - TCAM constraints

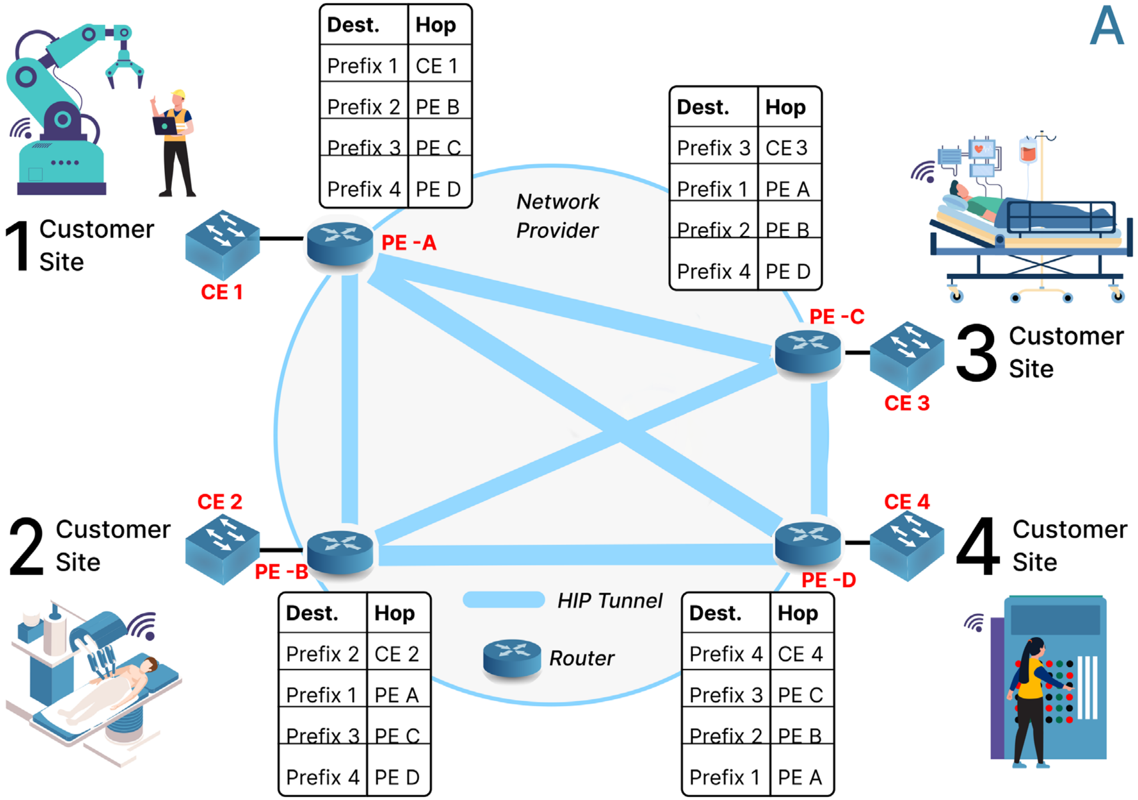

Relaying

- Hub PEs: full-mesh reachability
- Spoke PEs: reach other PEs by relaying through hubs
  - Pros: Relaying can substantially reduce the routing entries on the spoke PEs
  - Cons: more traffic being relayed, which could increase latency for the customer sites
- Optimization problem: minimize the *cost* of activating hubs, hub/spokes connection, and hubs connectivity for relay-based routing

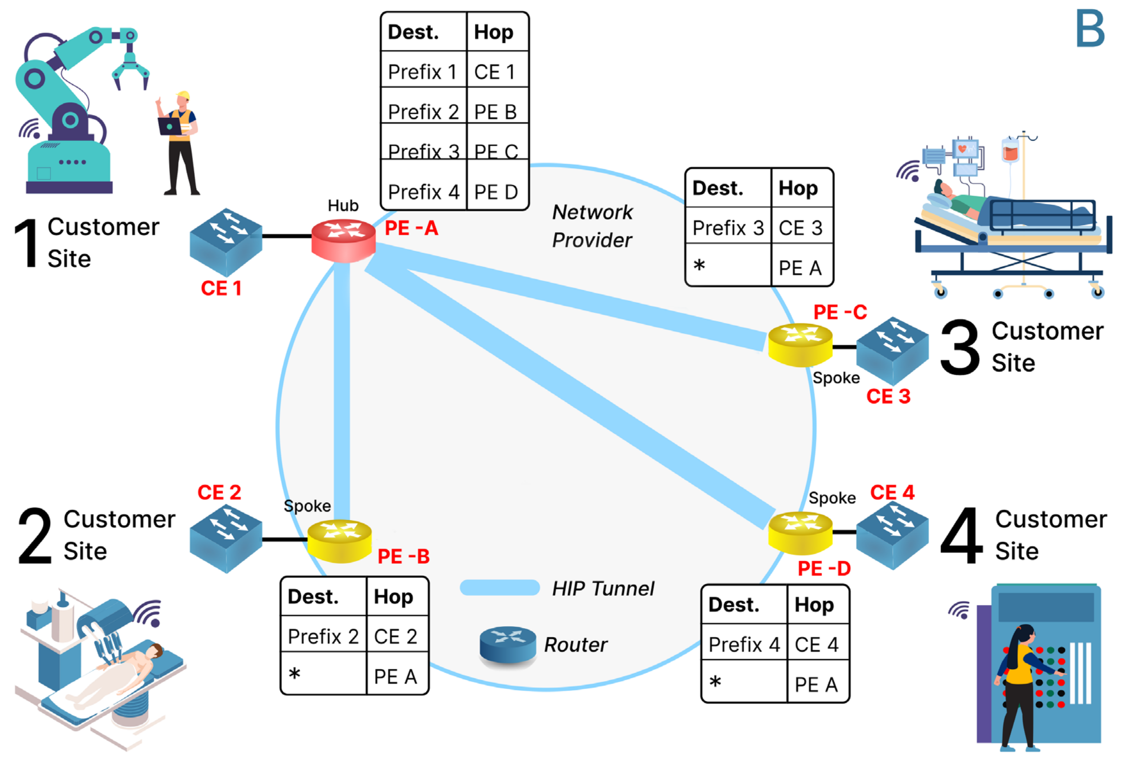

Relay Placement Problem (RPP) – Mathematical Program

- RPP considers:
  - **Hub PE selection**: Selecting a set of active hub: $F \subset \mathcal{F}$
  - **Hub PE assignment**: Assigning spoke PE_j to some hub PE_i
  - **Hub links selection**: selecting links that lower the total cost of connection between hub nodes

$G = (V, E)$	Provider network
$\mathbf{c}$	Cost to use specific link (latency, bandwidth)
$\mathcal{F} \subset V$	Set of potential locations for installing hub PEs
$\mathcal{D} \subset V$	Set of spoke PEs
x_{ij}	Binary variable whether spoke PE_j is connected to hub PE_i
y_i	Whether hub PE_i should be activated
a_i	Cost of activating hub PE_i
d_{ij}	Communication cost between spoke PE_j and PE_i

$$(P) \text{ minimize } \sum_{i \in F} \sum_{j \in \mathcal{D}} d_{ij} x_{ij} + N \sum_{k \in T.edges} c(k) + \sum_{i \in F} a_i y_i$$

Connection cost between spoke PEs and hub PEs

Cost to connect all hub PEs via Steiner Tree

Opening cost for hub PEs

Proposed Solution

Approximation Algorithm for RPP

- RPP is NP-hard problem [3]
- Proposed approach: Sample-Augment Algorithm:
  - Solving the subproblem: SH assignment

$$(SH)\ \text{minimize} \quad \sum_{i \in F} \sum_{j \in \mathcal{D}} d_{ij} x_{ij} + \sum_{i \in F} a_i y_i \tag{2a}$$

$$\text{subject to} \quad \sum_{i \in F} x_{ij} \geq 1,\ j \in \mathcal{D} \tag{2b}$$

$$x_{ij} \leq y_i,\ j \in \mathcal{D} \ \text{ and } \ i \in F \tag{2c}$$

$$x_{ij} \in \{0,1\},\ j \in \mathcal{D} \ \text{ and } \ i \in F \tag{2d}$$

$$y_i \in \{0,1\},\ i \in F \tag{2e}$$

  - Random sample from the problem input
  - Augmenting the result with the solution to original problem

- *Theorem 1: using 3-approximation algorithm for SH problem, 2-approximation for the Steiner Tree problem, Algorithm 1 is an expected 6.6-approximation algorithm for RPP.*

Algorithm 1: Approximation Algorithm for RPP.

```
γ ∈ (0, 1];
F ← ∅;
β ← γ/N;
/* Solving UFLP                                                    */
Execute the 3−approximation algorithm for Spokes-to-Hubs
  (SH) Assignment problem, and obtain the solution as
  H = (F_H, x_ij);
/* Sampling                                                        */
Sample (mark) a spoke PE* at random ;
Sample every other spoke non-marked PE independently
  with probability β;
Let M ={set of marked PEs} ;
/* Augmentation                                                    */
for all i' ∈ F_H if ({j|j ∈ D and x_i'j = 1} ∩M ≠ 0) then
  | F.add(i');
end
Execute the 2−approximated Steiner Tree T on the set M ;
Augment T with adding the shortest paths from each spoke
  PE j ∈ M and its associated hub PE;
Find a tree T'' which spans the F;
Allocate each spoke PE j ∈ D to its closest hub PE in F ;
return {F, T''}
```

Results

Results

- Simulation- Network Topology :
  - AS Network Topology [4]
  - Backbone Network Topology: Internet Topology Zoo

Comparison of Routing Entries for mid-size AS network: full-mesh vs. hub-spoke

#Entries in Routing Tables for All PEs			
#PEs	Full-mesh	Hub-Spoke	#Hubs
100	50100	2092	3
150	109950	3794	5
200	192800	7879	7
250	302250	9889	7
300	440100	12303	9

- The number of installed routing entries is significantly reduced by leveraging the relaying

Comparison of Routing Entries for large-scale AS network: full-mesh vs. hub-spoke

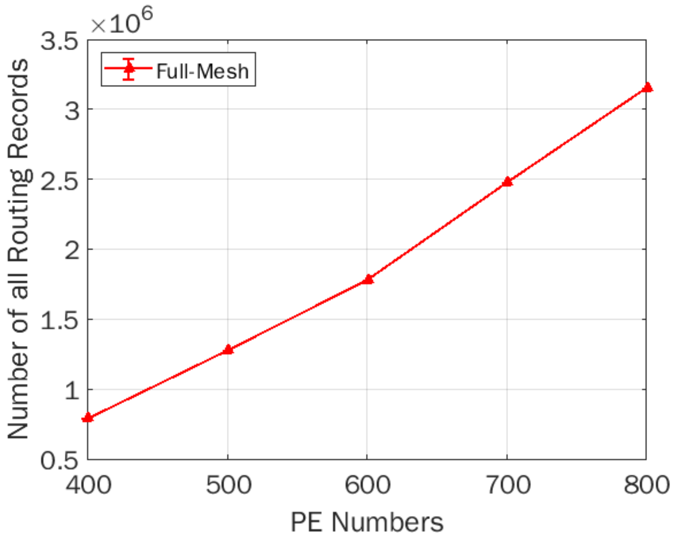

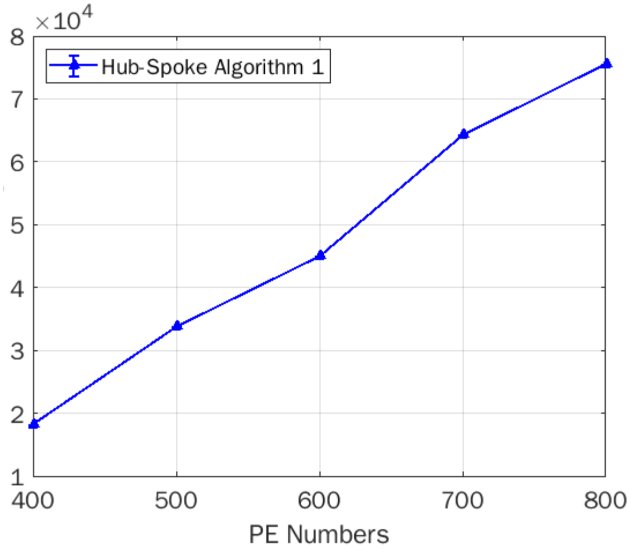

- Increase in number of PEs : the routing entries also increase. However, the increase is significantly greater with full-mesh

Results – Cont.

Solution cost: Random assignment vs. proposed Algorithm

- Proposed Algorithm generates less costly solutions for RPP than random hub placement

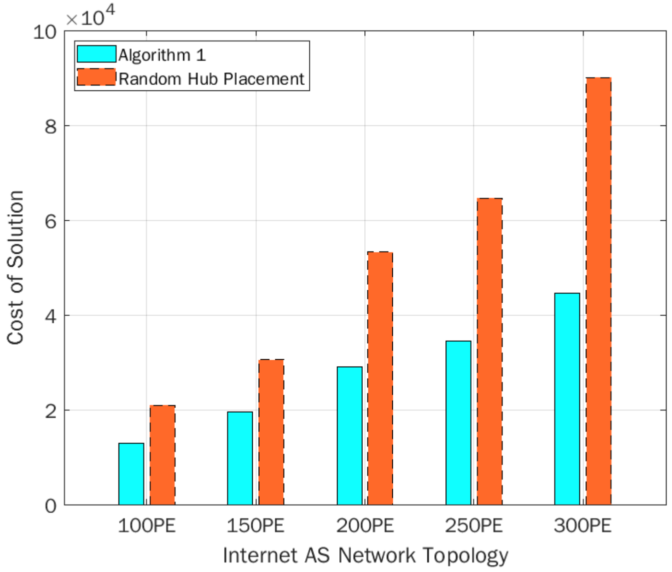

Mininet latency experiment for full-mesh and hub-spoke

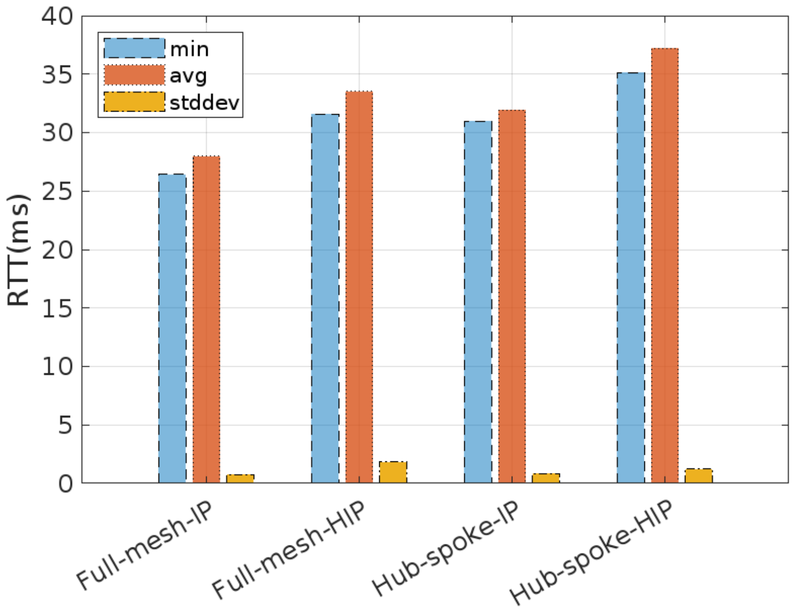

Comparison of path traversed by full-mesh vs hub-spoke

- The ratio of average increase in path length in traversed distance caused by hub-spoke design is less than 1.6

Extra Path Traversed by Hub-Spoke

Network	Location	#Nodes	Ratio	Margin of Error
Backbone, Transit	US	51	1.387	1.3873 ±0.115 (±8.26%)
Backbone, Customer	NL	50	1.536	1.5361 ±0.0774 (±5.04%)

HIPLS needs to endure an extra ~ 3.6 ms delay in hub-spoke compared to full-mesh on average RTT

Future Works/Conclusions

LINKÖPING UNIVERSITY

Conclusions

- RPP: Reducing the total routing tables size in secure VPLS that minimizes:
  - The number of hubs (hubs deployment cost)
  - The connection cost between spokes and hubs
  - The cost of connecting hubs
- We proposed algorithm with provable guarantees: approximate but efficient solution
- Simulation results: 90% reduction in routing table entries with a slight increase in tunnel path length
- We currently lack accurate traffic pattern/topologies data for real-world VPLS deployments
- Future works:
  - Add/remove PEs in an online manner
  - Resiliency support in problem formulation

References

[1] Shen, D. Guo, F. Long, L. A. Mateos, H. Ding, Z. Xiu, R. B. Hellman,A. King, S. Chen, C. Zhang, and H. Tan, “Robots under covid-19 pandemic: A comprehensive survey,”IEEE Access, vol. 9, pp. 1590–1615, 2021.

[2] K. Gaur, A. Kalla, J. Grover, M. Borhani, A. Gurtov, and M. Liyanage, “A survey of Virtual Private Lan Services (VPLS): Past, present and future,” Computer Networks, vol. 196, p. 108245, 2021

[3] Michael R. Garey and David S. Johnson. 1979. Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., USA.

[4] Ahmed Elmokashfi, Amund Kvalbein, and Constantine Dovrolis. 2010. On the Scalability of BGP: The Role of Topology Growth. IEEE Journal on Selected Areas in Communications 28 (2010), 1250–1261.

Thank you!

mohammad.borhani@liu.se