Hewlett Packard Labs

Homa: An Efficient Route Reconfiguration Approach in Software Defined WANs.

Diman Zad Tootaghaj, Hewlett Packard Labs, diman.zad-tootaghaj@hpe.com
Faraz Ahmed, Hewlett Packard Labs, faraz.ahmed@hpe.com
Puneet Sharma, Hewlett Packard Labs, puneet.sharma@hpe.com
Mihalis Yannakakis, Columbia University, mihalis@cs.columbia.edu

Hewlett Packard Enterprise

Agenda

- Background
- Network Monitoring
- Network Reconfiguration
    - Failure Recovery
    - Route Updates
- Conclusion and Future works

Software Defined WANs:

Hewlett Packard Enterprise

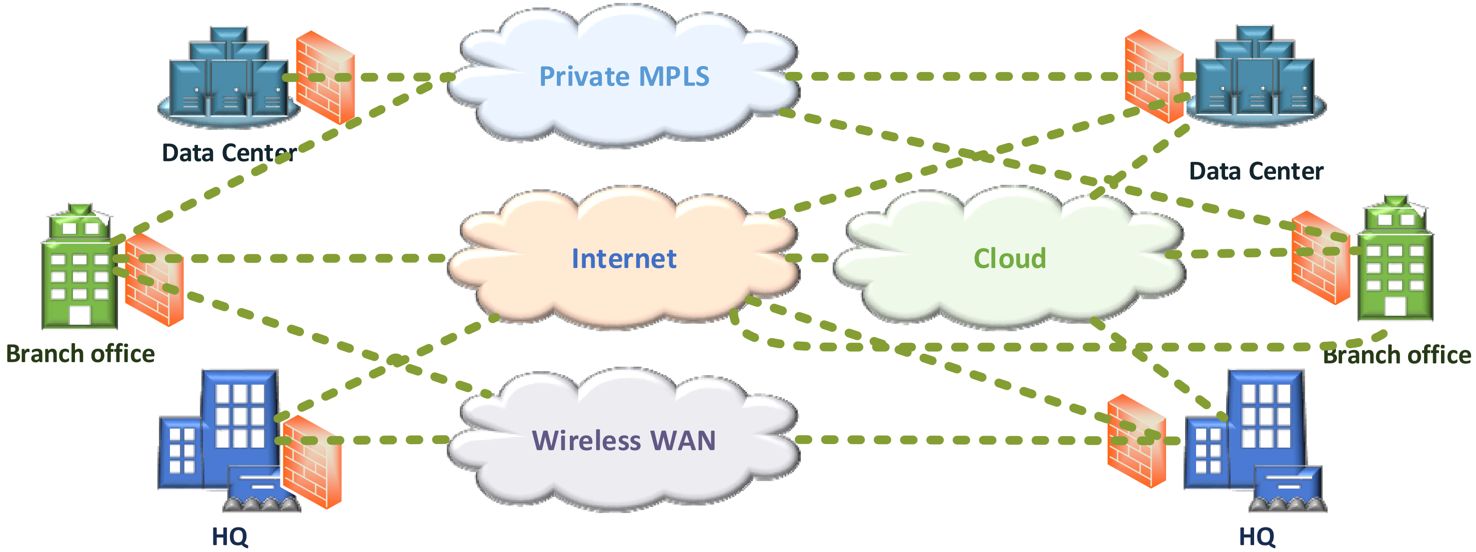

SD-WAN Benefits:

- Lowers management costs
- Boosts network performance
- Reduces admin burden
- Reduces software expenses
- Reduces public cloud data transfer expenses
- Increases endpoint scalability and reduced data center router expenses
- Reduces network OPEX
- Increases network availability and security

SD-WAN types:

- Inter Data Center
- Between Branch offices
- Edge to Core
- Edge to Cloud

Why Reconfigure the Network?

Observations:

Failures are more common in SD-WAN compared to datacenter networks

- SDWAN does not have dedicated control networks
- SDWAN does not have high degree of parallel links
- SDWAN is exposed to physical threats (e.g. cable cuts)
- SDWAN are not deployed in a safe and controlled environment

Old network rules

Network Reconfiguration

New network rules

Background: SD-WAN

1) Lack of Information Visibility

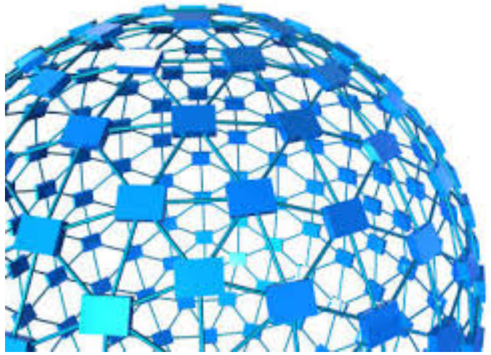

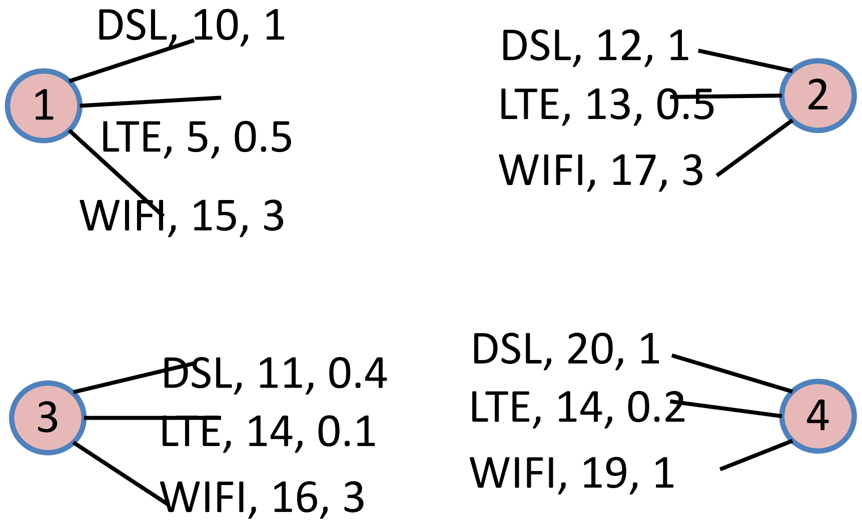

2) Usage-based Cost Model

$$p_{ij}, t = \alpha_{ij,t} + \zeta_{ij,t} . f_{ij,t} \quad \forall (i,j) \in E_o, \forall t \in I$$

SDWAN Network Design Solutions

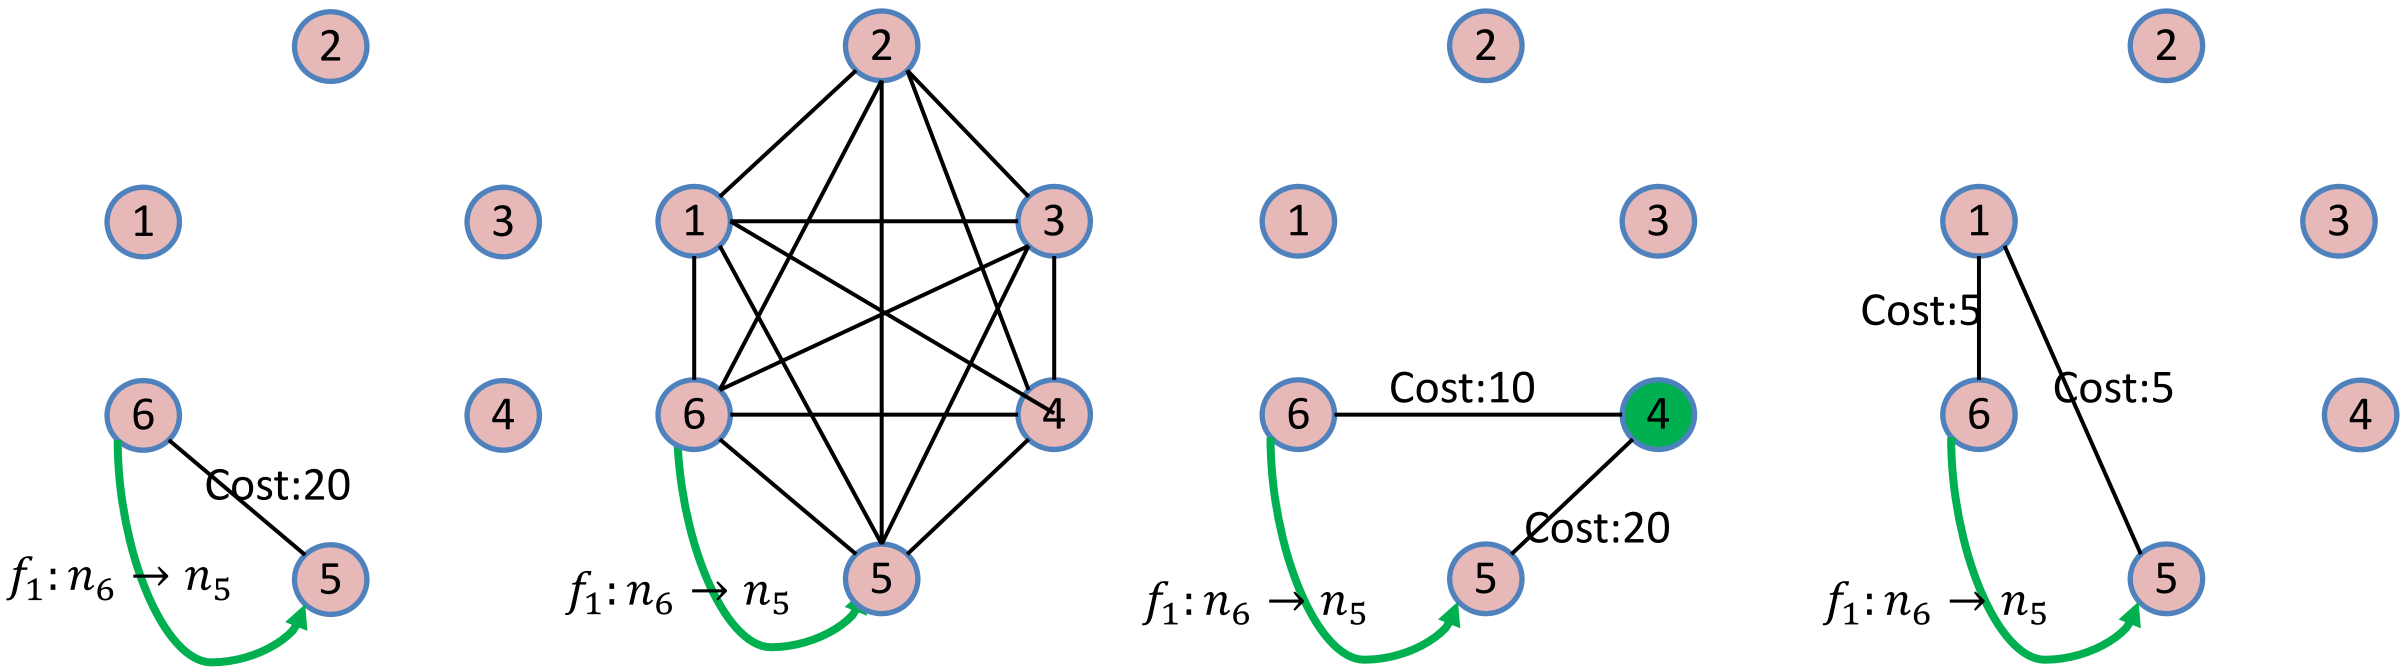

Direct Link | Full Mesh | Hub and spoke | **HOMA**

Motivating Example: Direct Links

Hewlett Packard Enterprise

- Example: Direct Links
- Flows: $f_1 : n_6 \rightarrow n_5, f_2 : n_1 \rightarrow n_5$
- Cost model 1: $10 \times l + f \times l$
- Total cost : 10* 5 +5 and 10*4 +4 = 99

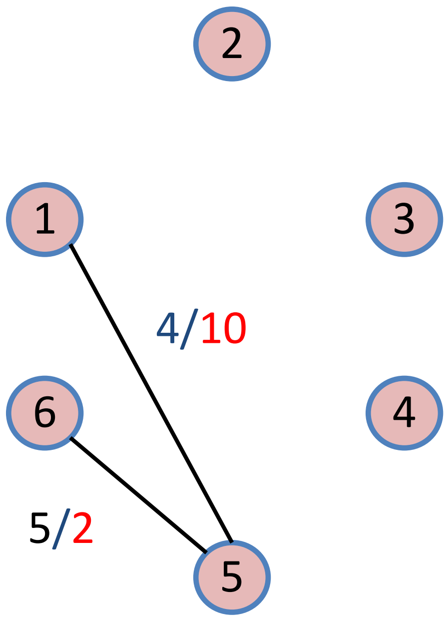

- Cost model 2: $20 \times c + f \times c$
- Total cost : 20* 2 +2 and 20*10 +10 = 252

Motivating Example: Full Mesh

- Example: Full Mesh
- Flows: $f_1: n_6 \rightarrow n_5, f_2: n_1 \rightarrow n_5$
- Cost model 1: $10\times l + f\times l$
- Total cost : 6(20* 5)+5 and 9*(20*4) +4=1329

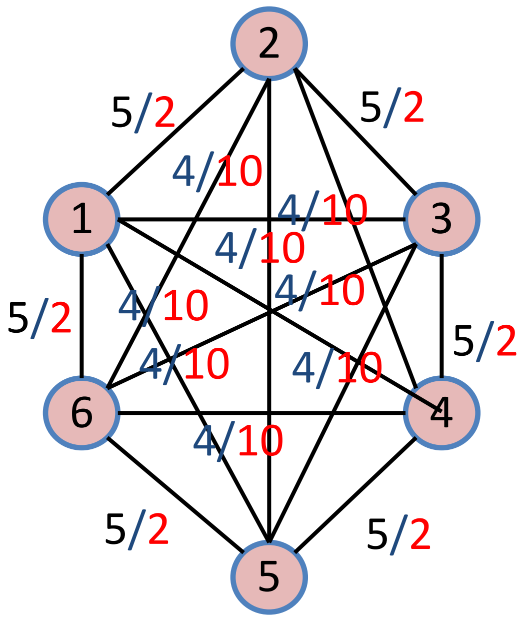

- Cost model 2: $20\times c + f\times c$
- Total cost : 6(20* 2) +2 and 9*(20*10) +10=2052

Motivating Example: H&S

- Example: Hub and Spoke
- Flows: $f_1: n_6 \rightarrow n_5, f_2: n_1 \rightarrow n_5$
- Cost model: $10 \times l + f \times l$
- Total cost 1: (10* 4 +4)*2 and 2*(10*5 +2*5)
- = 208

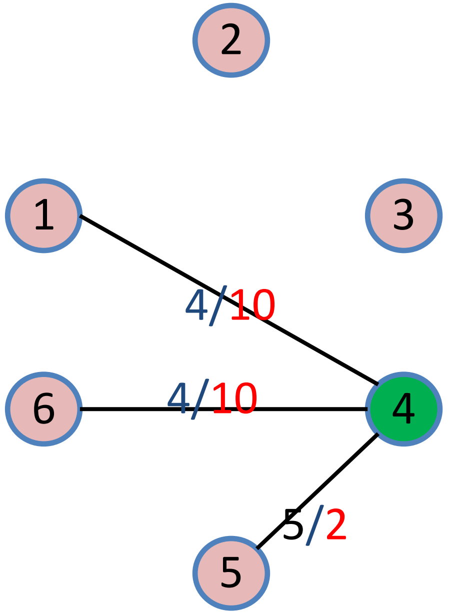

- Cost model 2: $20 \times c + f \times c$
- Total cost : (20* 10 +10)*2 and (20*2+2*2) = 464

Motivating Example: HOMA

Hewlett Packard Enterprise

- Homa
- Flows: $f_1 : n_6 \rightarrow n_5, f_2 : n_1 \rightarrow n_5$
- Cost model 1: $10 \times l + f \times l$
- Total cost :
  - (10*4+4) and (10*5+5) = 99
- Cost model 2: $20 \times c + f \times c$
- Total cost : (20* 2 +2*2)*2=88

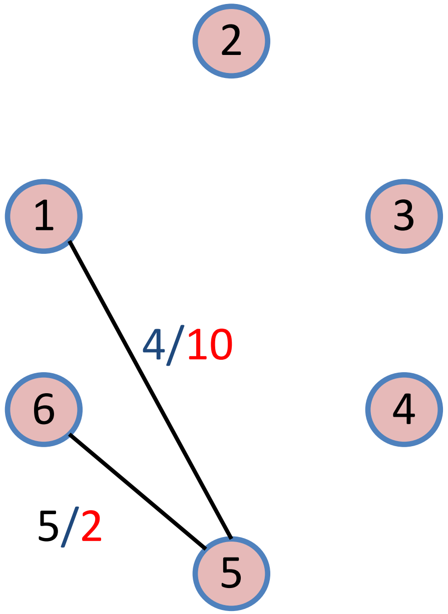

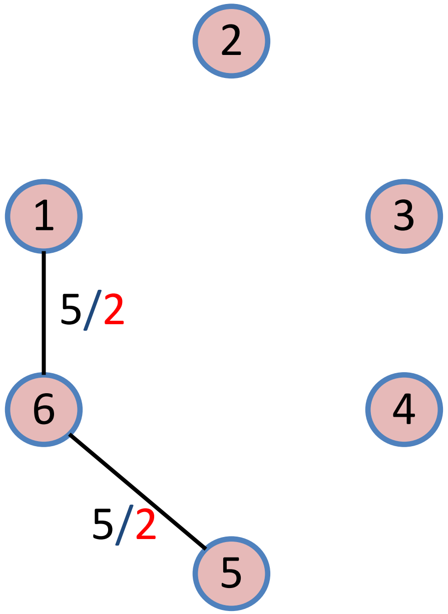

Motivating Example: Comparison

Hewlett Packard Enterprise

- Comparison between Direct Links, H&S, FM, and our approach **Homa**

Approach	Cost Model 1	Cost Model 2
Homa	99	88
Direct Links	99	252
H&S	208	464
Full Mesh	1329	2052

Inferring the underlay topology

Complete information about the underlay topology:

Fraction of overlapping links

$$FOL(e_{ij,t}, e_{nm,t}) = \frac{|l \in e_{ij,t} \cap e_{nm,t} : l \in E_u|}{|E_u|}$$

Inferring the underlay topology:

End-to-end monitors measure the delay

Find the correlation between two virtual links and compare with the autocorrelation to infer the physical network

$$CM(e_{ij,t}, e_{nm,t}) = corr(d_{ij,t}, d_{mn,t})$$

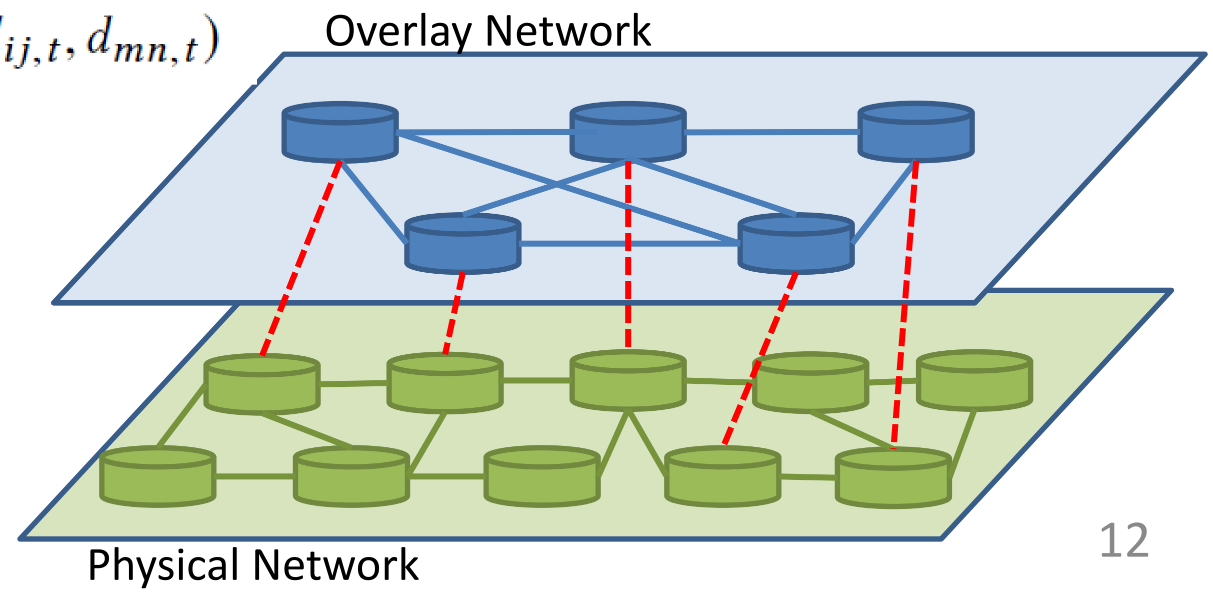

Experimental Results

Impact of direct correlation (shared congested links)

Flow 1: Ping (ICMP packets) h1->h2
Flow 2: Ping h2 -> h3
Flow 3: UDP connection h3 -> h4

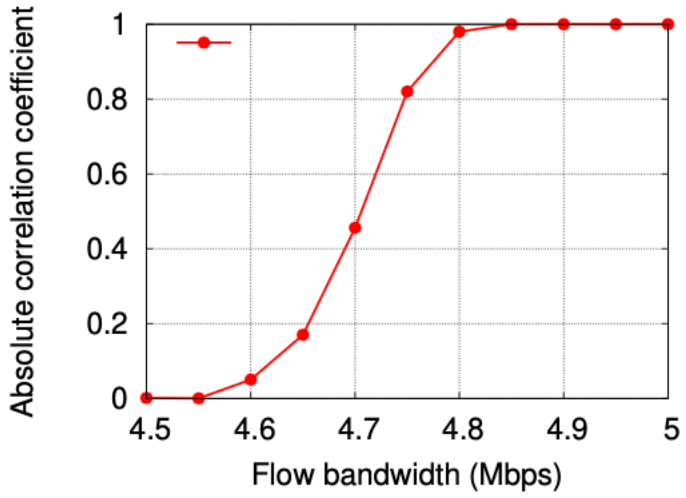

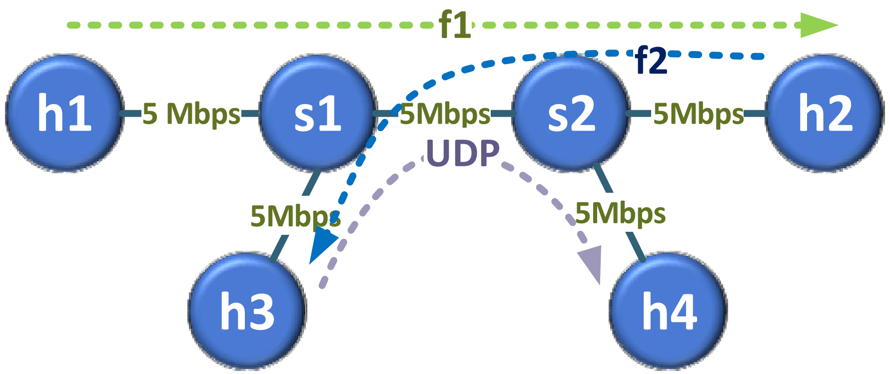

Experimental Results

Impact of indirect correlation

Flow 1: Ping (ICMP packets) h3->h1
Flow 2: Ping h4 -> h2
Flow 3: UDP connection h1 -> h2

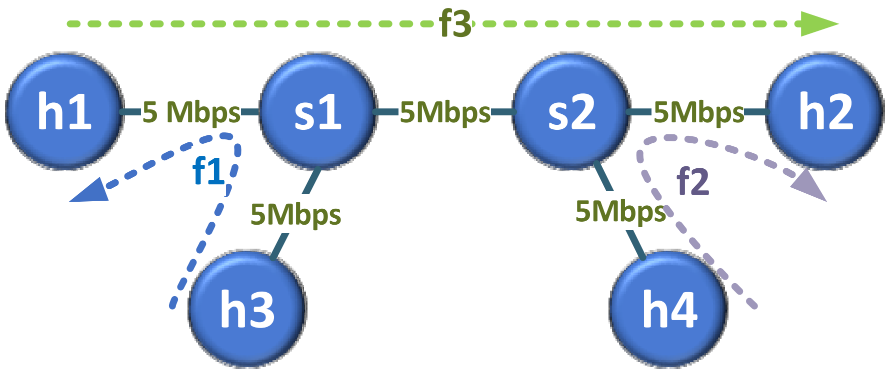

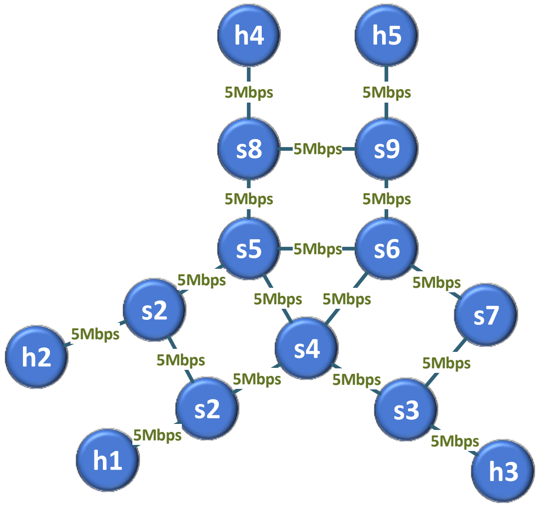

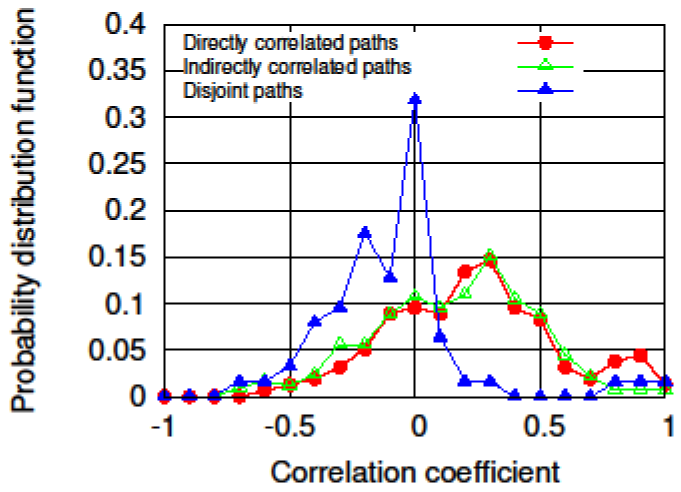

HOMA: Minimum Cost overlay

Overlay network design

Find an overlay network topology that satisfies the demand with minimum cost

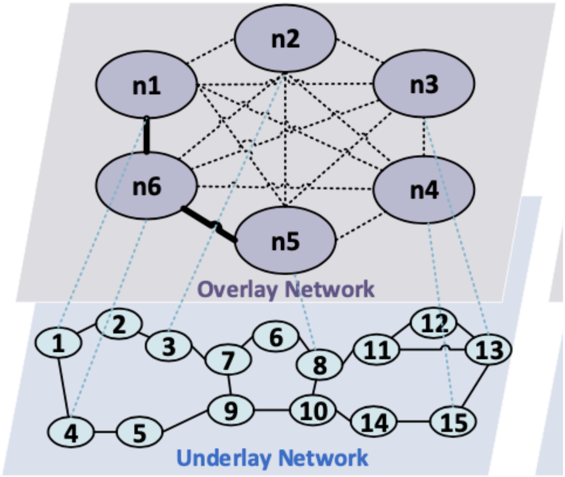

Cost Model:

Fixed cost: Greedy solves it optimally: Use Kruskal's algorithm

Variable cost: Homa

$$p_{ij}, t = \alpha_{ij,t} + \zeta_{ij,t} . f_{ij,t} \quad \forall (i,j) \in E_o, \forall t \in I$$

[1] R. T. Wong. Worst-case analysis of network design problem heuristics. SIAM Journal on Algebraic Discrete Methods, 1980.

$Homa^1$: Minimum cost network update

Hewlett Packard Enterprise

Update the network with minimum cost

$$\textbf{minimize} \sum_{t\in I}\sum_{(i,j)\in E_o} \alpha_{ij,t}\lambda_{ij,t} + \sum_{t\in I}\sum_{(i,j)\in E_o} \zeta_{ij,t} \sum_{h\in H} (x_{ij,t}^h + x_{ji,t}^h).\, d_h$$ Turn on the links with minimum cost

$$\textbf{subject to} \sum_{h\in H} (x_{ij,t}^h + x_{ji,t}^h).\, d_h \le c_{ij,t}, \qquad \forall (i,j) \in E_o, \forall t \in I \quad \text{(1a)}$$

$$\sum_{t\in I}\sum_{j\in V_o} x_{ij,t}^h = \sum_{t\in I}\sum_{k\in V_o} x_{ki,t}^h + sign(b_i^h), \qquad \forall i \in V_o, \forall h \in H \quad \text{(1b)}$$

$$\sum_{h\in H}\sum_{j\in V_o} x_{ij,t}^h.\, d_h \le c_{i,t}, \qquad \forall i \in V_o, \forall t \in I \quad \text{(1c)}$$

$$\sum_{h\in H}\sum_{j\in V_o} x_{ji,t}^h.\, d_h \le c_{i,t}, \qquad \forall i \in V_o, \forall t \in I \quad \text{(1d)}$$

$$\lambda_{ij,t} \ge \frac{\sum_{h\in H}[x_{ij,t}^h + x_{ji,t}^h]}{|H|}, \qquad \forall t \in I, \qquad \forall (ij) \in E_o \quad \text{(1g)}$$

$$\sum_{t\in I}\sum_{h'\in H, h'\ne h}\sum_{(m,n)\in E_o}\sum_{(i,j)\in E_o} x_{ij,t}^h.\, x_{mn,t}^{h'}.\, corr(x_{ij,t}^h, x_{mn,t}^{h'}) \le \Omega_h, \qquad \forall h \in H \quad \text{(1h)}$$ QoS constraint

$$x_{ij}^h, \lambda_{ij,t} \in I \quad \forall (i,j) \in E_o, \forall h \in H \quad \text{(1i)}$$

Min-Cost is NP hard:

Baselines: 1) using direct links, 2) Hub and Spoke, 3) Full Mesh

Homa: Our approach

1. Homa is a mythical bird, that is said to be phoenix-like, consuming itself in fire every few hundred years, only to rise anew from the ashes.

$Homa^1$: Minimum cost network update

Hewlett Packard Enterprise

Update the network with minimum cost

Homa:

Unit demands: randomized greedy algorithm: provides an approximation guarantee of $O(e^{\sqrt{\ln|H| ln ln|H|}})$

General demands: randomized greedy algorithm: provides an approximation guarantee of $O\left(e^{\sqrt{\ln|H| ln ln|H|}}\right).logD$

When we have $|H|$ source-destination pairs.

Algorithm 1: Greedy algorithm for *Min-Cost* problem

Pick a random permutation of demands (without loss of generality let $d_1, d_2, ..., d_{|H|}$ be the random permutation).
Initialize the residual capacity of all links:
$c_{\tilde{ij},t} = c_{ij,t}\ \forall(ij) \in E_o, t \in I$
for $i = 1$ *to* $|H|$ **do**
  Set $d'_i = \frac{|H|}{i}$.
  Find the shortest constraint path using CSP algorithm (Algorithm 2) to find the smallest cost of routing d'_h units of demand, using the network constructed for the previous $i-1$ pairs.
  if *the correlation constraint for all existing flows are satisfied by adding the new flow* **then**
    Route a single unit of demand between s_i and t_i along the path selected at the i-th iteration.
    Update the residual capacity $c_{\tilde{ij},t}$ of all links along the selected path: $c_{\tilde{ij},t} = c_{\tilde{ij},t} - 1$
  else
    Drop d_i.

Inflated demands

Failure Resiliency

Hewlett Packard Enterprise

Average number of disrupted flows (a), and cost of network reconfiguration (b), in Homa that uses indirect links and traditional SD-WAN that uses direct link architecture, as we increase the number of failures.

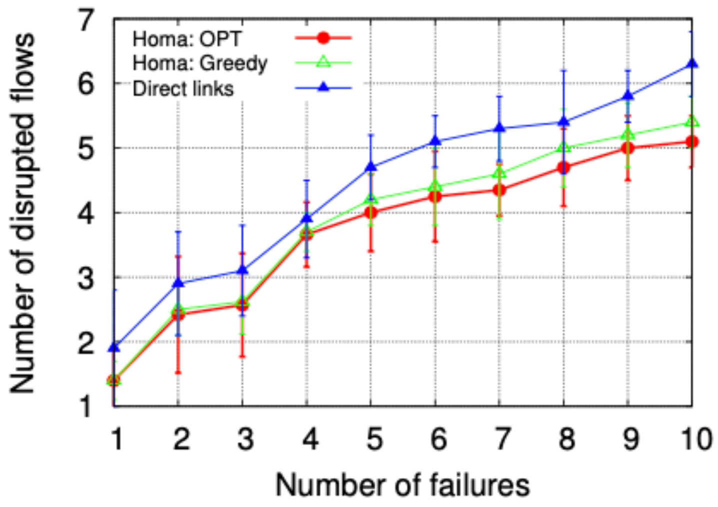

(a) Disrupted Flows.

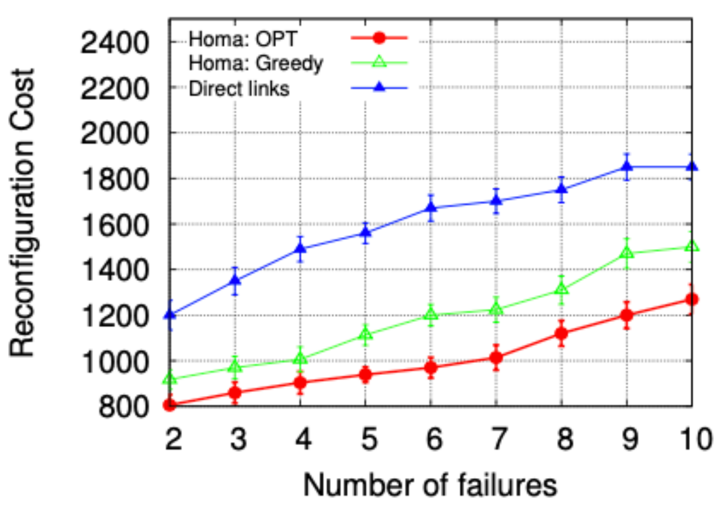

(b) Cost.

Demand Loss

Average percentage of satisfied demands, as we increase the number of failures and demand pairs.

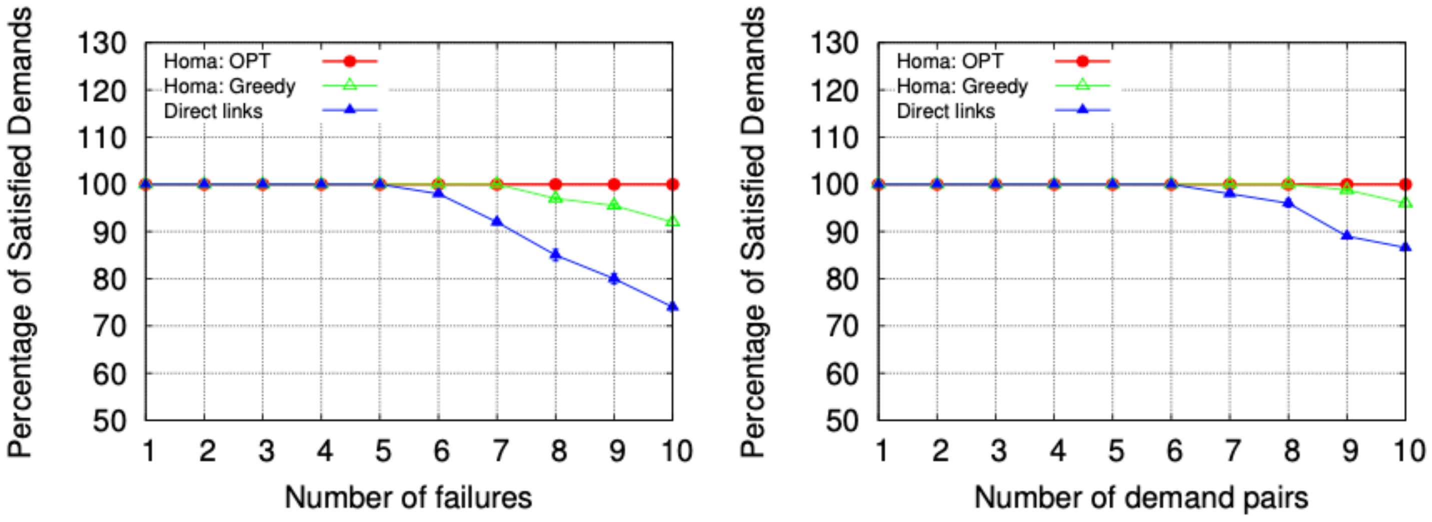

(a) Increasing Failures. (b) Increasing Demands.

Inferring the underlay topology

Comparison between (a) network reconfiguration cost, (b) number of disrupted flows, when we have full information about the underlay network topology and when using the correlation matrix from the underlay network topology.

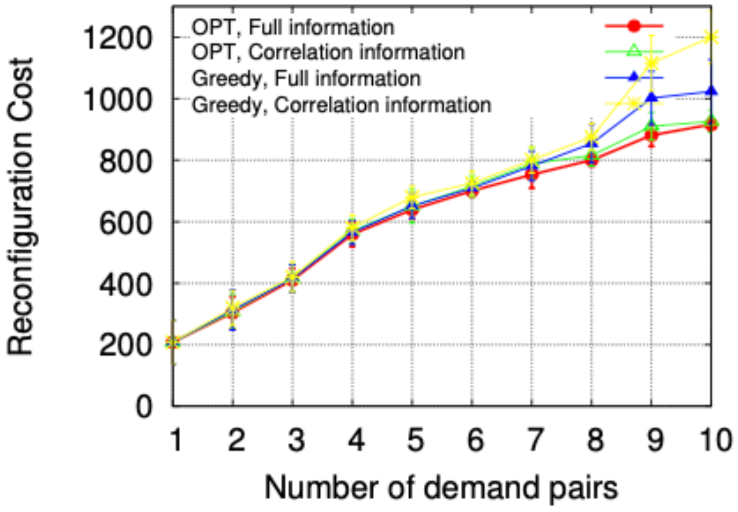

(a) Reconfiguration Cost.

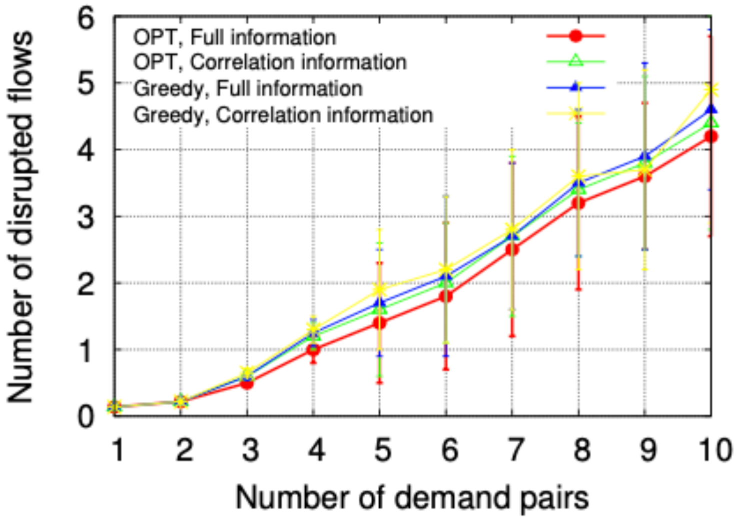

(b) Disrupted Flows.

Conclusion

- **Key ideas**:
  – Overlay Network with minimum cost.
    - Failure, service/traffic changes, topology changes
    - Monitoring the performance of overlay links
    - Satisfying QoS constraint
- **Results:**

  Lower disruption cost w.r.t baseline, lower number of shared links, low demand loss

  Trade-off between:
    - Disruption
    - Demand loss
    - Congestion
    - Execution time
    - Reconfiguration cost

Hewlett Packard Labs

Thank you

Diman Zad Tootaghaj, Hewlett Packard Labs, diman.zad-tootaghaj@hpe.com
Faraz Ahmed, Hewlett Packard Labs, faraz.ahmed@hpe.com
Puneet Sharma, Hewlett Packard Labs, puneet.sharma@hpe.com
Mihalis Yannakakis, Columbia University, mihalis@cs.columbia.edu

Hewlett Packard Enterprise

Hewlett Packard Enterprise

References

[1] Michel, O., & Keller, E. (2017, May). SDN in wide-area networks: A survey. In *Software Defined Systems (SDS), 2017 Fourth International Conference on* (pp. 37-42). IEEE.

[2] Jain, S., Kumar, A., Mandal, S., Ong, J., Poutievski, L., Singh, A., ... & Zolla, J. (2013, August). B4: Experience with a globally-deployed software defined WAN. In *ACM SIGCOMM Computer Communication Review* (Vol. 43, No. 4, pp. 3-14). ACM.

[3] Lei, T., Lu, Z., Wen, X., Zhao, X., & Wang, L. (2014, May). SWAN: An SDN based campus WLAN framework. In *Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems (VITAE), 2014 4th International Conference on* (pp. 1-5). IEEE.