Service-driven Network Virtualisation Through Multi-Topology Routing

Dr. Ning Wang

University of Surrey
Guildford, UK
http://www.ee.surrey.ac.uk/CCSR/

November 6, 2007, Brussels

- **Background**

  - Separation of value-added services from the physical network infrastructure

  - Physically/logically partition the network resources for supporting heterogeneous services

    - Physical resources: network bandwidth
    - Virtual resources: routing/forwarding tables, different policies for traffic treatment etc.

- **Network Planes (NPs)**

  - Slices of physical/logical network resources that are used for supporting heterogeneous services

  - A multi-dimensional network resource engineering paradigm

Multi-Topology Routing (MTR)

- **Background**

  - Providing separate routing policies or decisions for different *types* of traffic

    - IPv4 vs. IPv6
    - Unicast vs. multicast

  - Examples: Multi-topology OSPF (RFC 4915)

- **MTR for network virtualisation**

  - Providing separate routing policies or decisions for traffic *with different service requirements*

  - Providing resilience support against traffic dynamics (upsurges) and physical link failures

    - Dynamic load balancing across multiple routing topologies
    - Fast re-routing in case of link failures without waiting for IGP re-convergence

- **Overview**

  - Using multiple MT-IGP topologies for supporting distinct Network Planes (service differentiation)

  - Using multiple *equivalent* MT-IGP topologies within one single NP for load balancing and resilience support

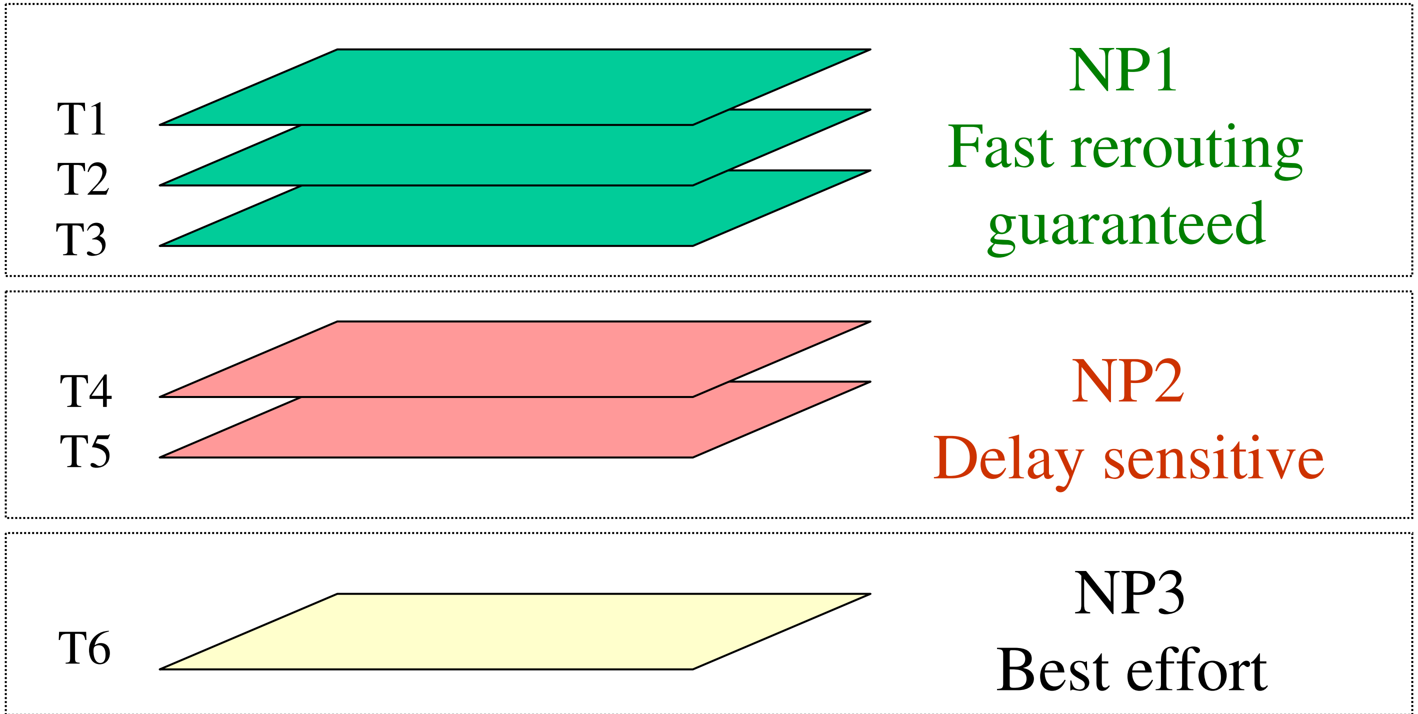

A Framework for MTR-based NP Engineering

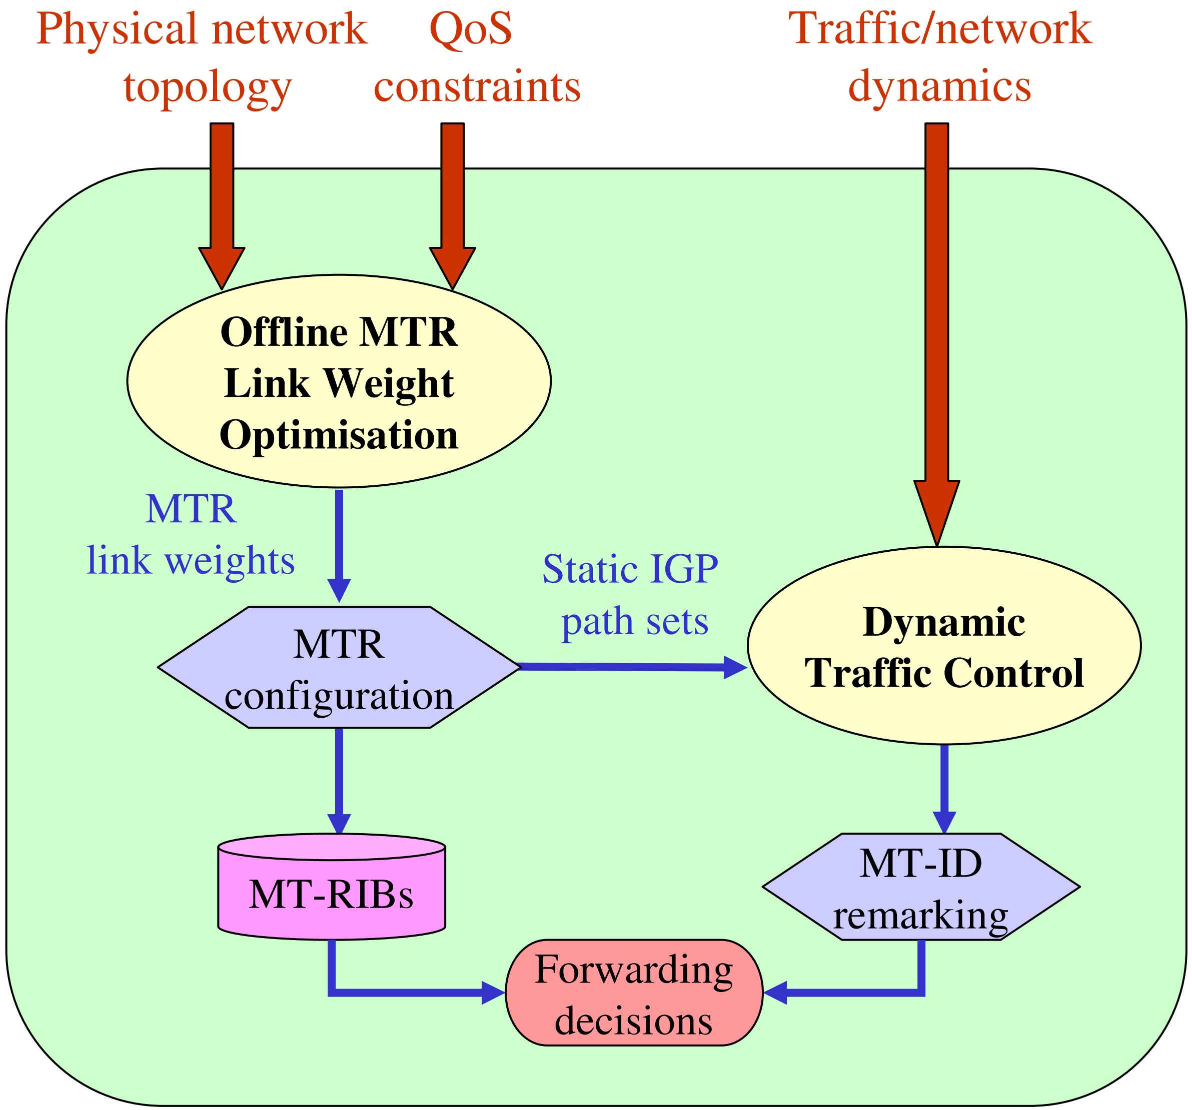

Path Diversity Metric: Degree of Involvement (DoI)

- *Degree of Involvement (DoI)* of a link for an OD PoP pair is the number of times it is included in the shortest IGP paths in different MTR topologies for each OD pair.

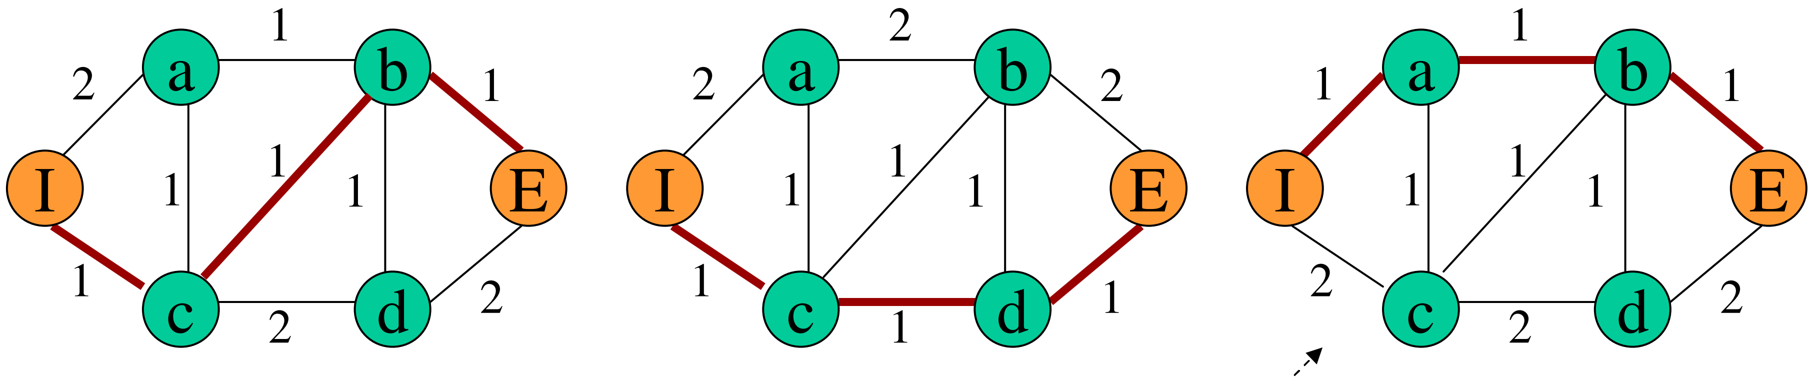

Topology 1 Topology 2 Topology 3

$$FDoI_l^{u,v} = \begin{cases} 1 & \text{if } DoI_l^{u,v} = |R| \\ 0 & \text{Otherwise} \end{cases}$$

$$DoI_{I,c}^{I \to E} = \begin{cases} 2 \\ 3 \end{cases}$$

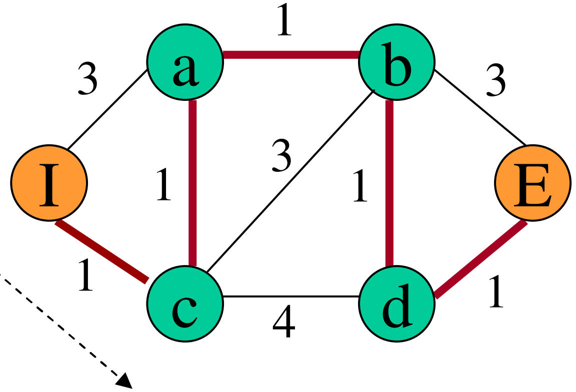

Topology 3'

To minimize: $\sum_{u,v \in V} \sum_{l \in E} FDoI_l^{u,v}$

The GEANT Network Topology (2004)

– 23 PoP nodes
– 74 uni-directional links with bandwidth capacity of 155Mbps, 2.4Gpbs 4.8Gbps and 10Gbps

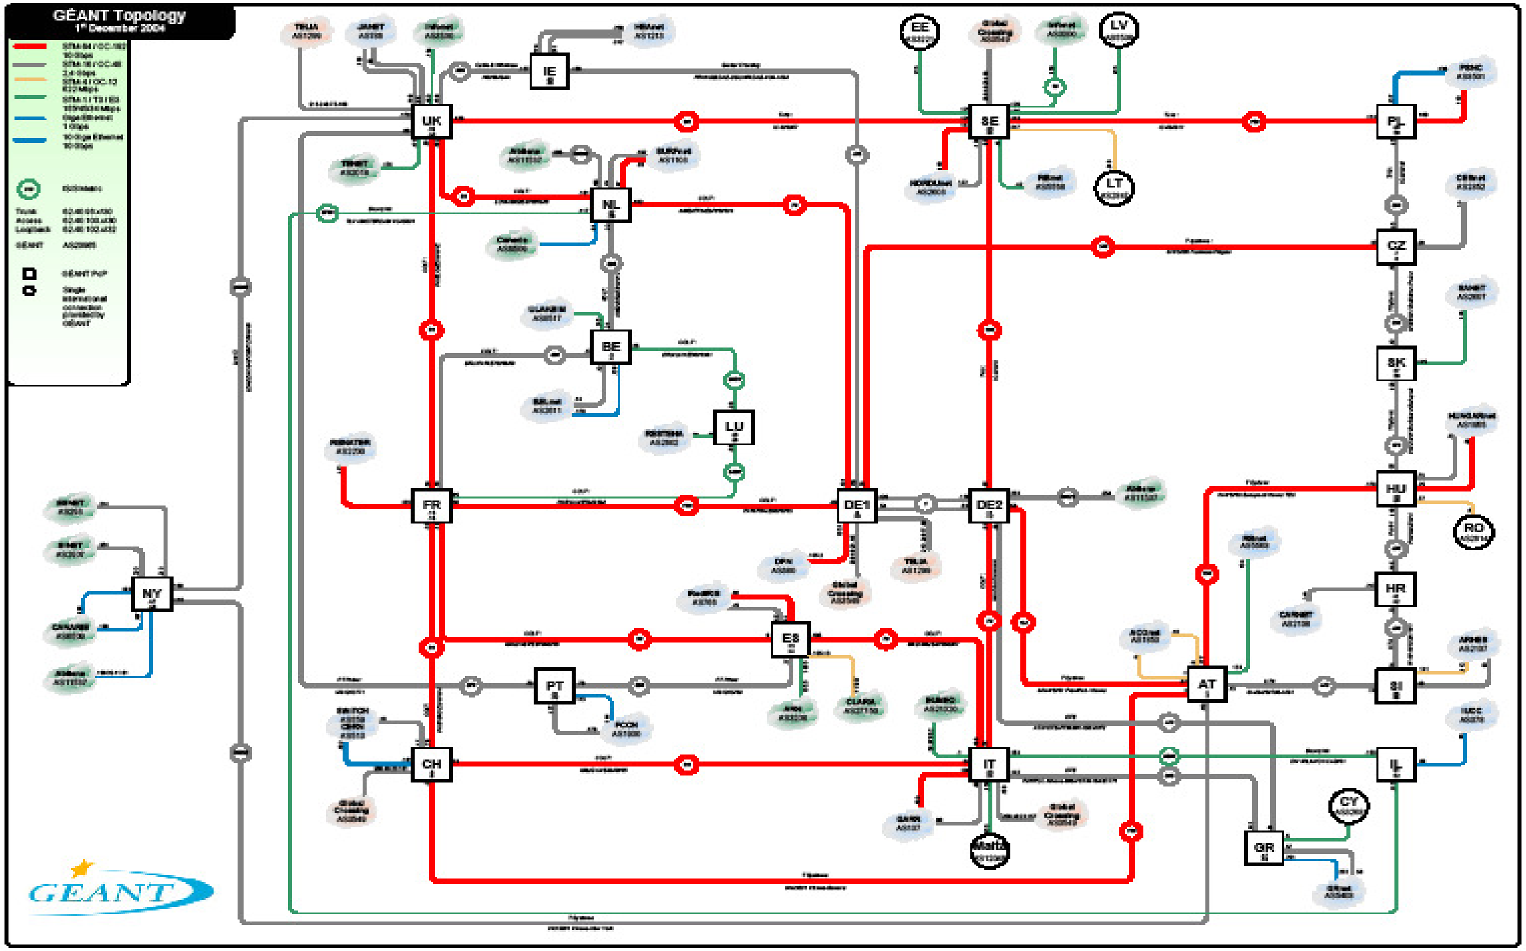

- **Performance metrics**
  - The proportion of OD pairs that can successfully avoid any critical link with *FDoI* (i.e., shared by all routing topologies
  - Path length (as delay constraint for NP2)

100% chance of achieving fast rerouting for any scenario of single link failure or congestion

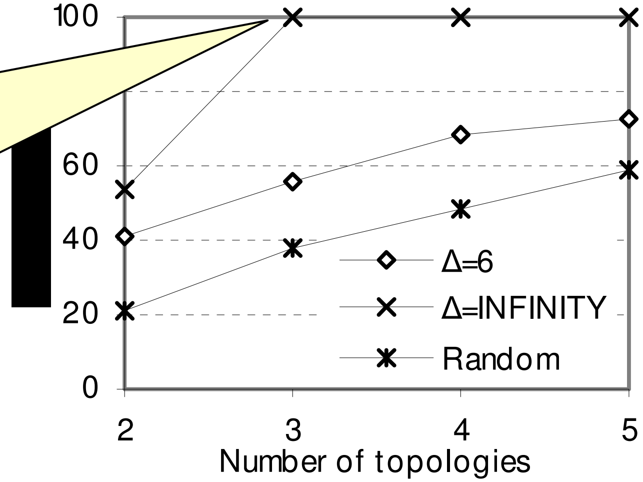

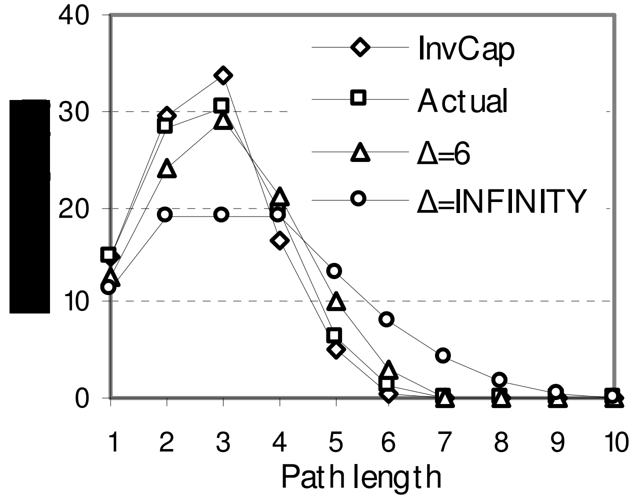

	InvCap	*Actual*	*Δ=6*	*Δ=INF*
Max_len	6	6	6	10
Avg_len	2.69	2.78	3.00	3.57

- **Input**

  - Optimised MTR link weights

  - Traffic/network dynamics

- **Output**

  - Adjustment for splitting ratio for traffic assignment across multiple equivalent routing topologies within each NP

- **Objectives**

  - Load balancing of traffic across topologies at short time scale (e.g., hourly)

  - Perform fast rerouting to alternative routing topologies in case of link failures without waiting for IGP routing re-convergence

- **Assume a centralised TE manager who:**

  - knows the overall network topology

  - gathers and maintains the monitored network performance

  - Periodically calculates the optimised traffic splitting ratio

  - Instructs ingress PoP nodes to enforce the splitting ratio

- *Identify the most utilised link l_{max} in the network*
- For the set of traffic flows (i.e traffic demand between an S-D pair) that are routed through l_{max} in *at least one but not all* the routing topologies, consider each flow f at a time and *incrementally* compute its new traffic splitting ratio among the routing topologies (start from a small proportion demand of f, if succeed exponentially increase the proportion until no further improvements can be made by adjusting the splitting ratio of f
- Go to the next most utilised link and repeat the procedure until no further improvement can be made. The total number of iterations is bound by K

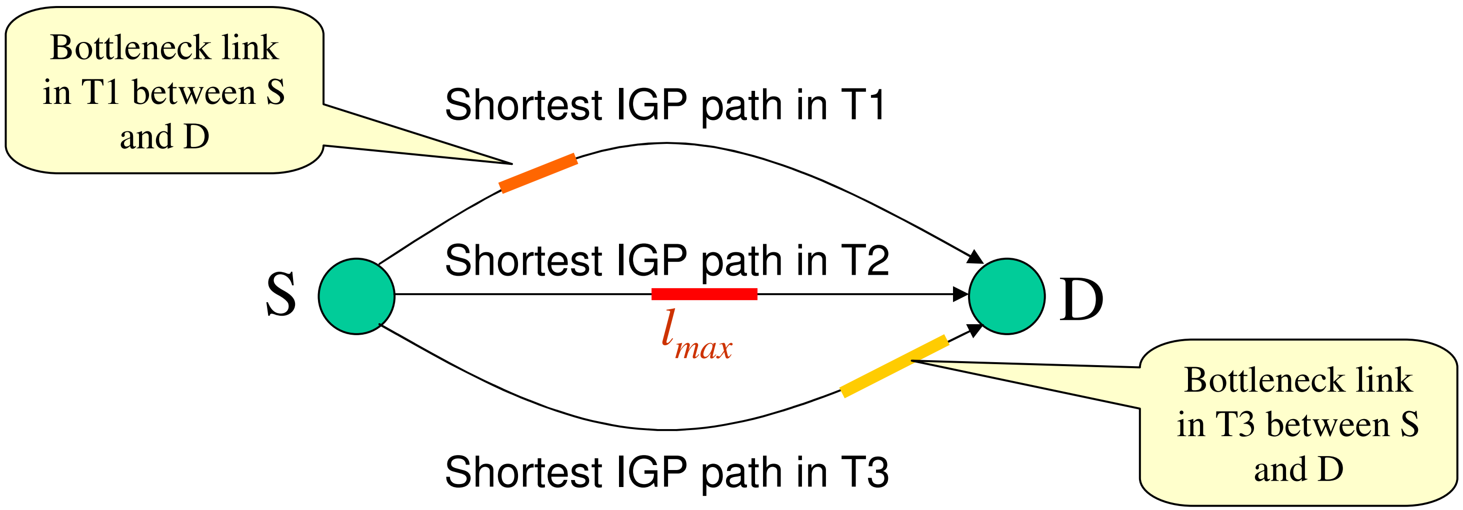

- **Traffic Matrix Series**

  - The GEANT network

    - TM obtained every 15 minutes (Apr. 2005)
    - Test with traffic traces for 7 consecutive days (more than 600 distinct traffic matrices)
    - Dataset obtained from the TOTEM project

      http://totem.info.ucl.ac.be/

- **Performance metrics:**

  - Maximum intra-domain link utilization (*MLU*)

  - Network cost (piece-wise linear function)

- **Algorithm comparison:**

  - Link weight setting inverse to capacity (InvCap)

  - Actual link weight setting by the operators (Actual)

  - Robust link weight setting considering multiple TMs (Multi-TM)

  - Optimal

GEANT Performance (Max. Link Utilisation)

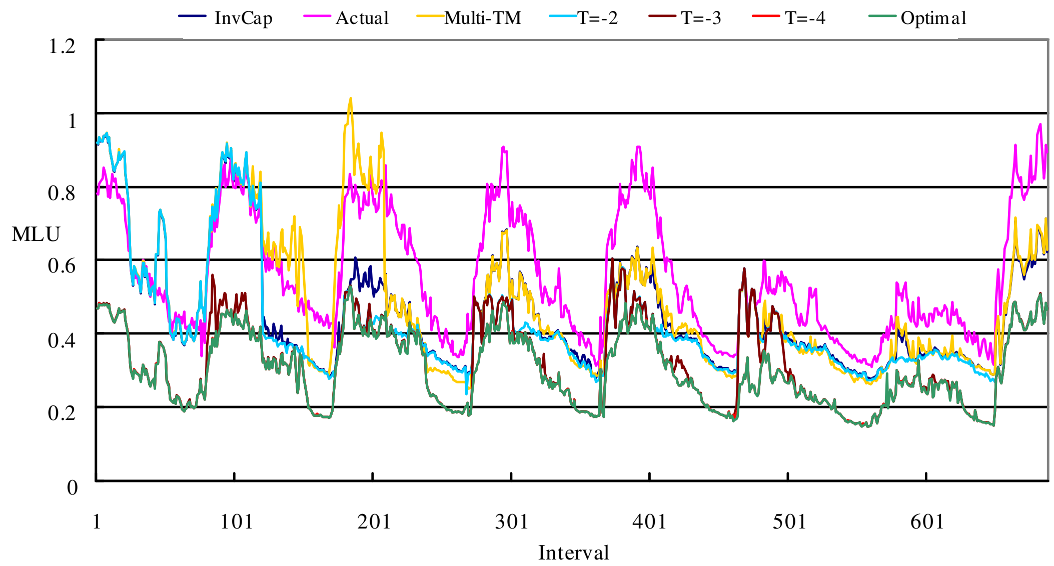

NP1
No delay constraint

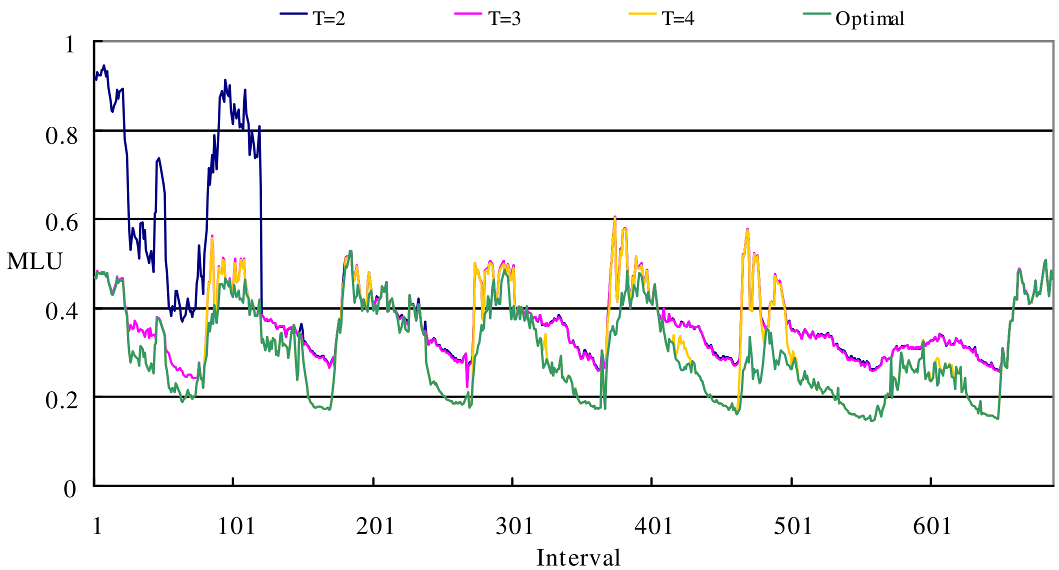

NP2
With delay constraint
(Δ=6)

Statistics on *MLU* Performance

	AMU	HMU	PNO
Optimal	30.05%	52.82%	-
InvCap	45.72%	94.41%	1.6%
Actual	55.47%	96.91%	0%
Multi-TM	48.56%	100.04%	0.44%
NP1 (2T, Δ = INF)	42.9%	92.61%	13.08%
NP1 (3T, Δ = INF)	31.95%	60.36%	78.34%
NP1 (4T, Δ = INF)	30.08%	52.88%	99.56%
NP2 (2T, Δ = 6)	41.85%	94.41%	17%
NP2 (3T, Δ = 6)	36.63%	60.61%	26.02%
NP2 (4T, Δ = 6)	31.86%	60.22%	78.78%

***Average maximum link utilization* (*AMU*):** the average value of the *MLU* across all the traffic matrices during the seven-day period;

***Highest maximum link utilization* (*HMU*):** the highest value of the *MLU* across all the traffic matrices during the period.

***Proportion to near-optimal performance* (*PNO*):** the percentage over all the traffic matrices in which our scheme can achieve near-optimal performance. We define here the meaning of near-optimal to be the *MLU* that is within 3% of gap to the optimality.

- **Motivation**
  - Network virtualisation through multi-topology IGP routing
  - To achieve service differentiation across multiple Network Planes
  - To enable dynamic traffic control against traffic burst and single link failures

- **Algorithm comparison:**
  - A two phase NP engineering scheme is designed and implemented
    - Offline optimisation of MTR link weight for maximising path diversity within a domain
    - Short-time scale traffic control according to the monitored performance against traffic dynamics

- **Outcome**
  - Simulation experiments based on operational network topology show that near-optimal TE performance can be obtained with only a few multi-topology routing topologies