cariden

the economics of network control

A Practical Perspective on Traffic Engineering

UKNOF 15 – 21 Jan 2010

John Evans
johnevans@cariden.com

Traffic engineering in context

- Network planning and traffic engineering are two faces of the same problem.
  - Network planning:
    - Ensuring there is sufficient capacity to deliver the SLAs required for the transported services
    - e.g. building your network capacity where the traffic is
  - Traffic engineering:
    - Ensuring that the deployed capacity is efficiently used
    - e.g. routing your traffic where the network capacity is

A practical perspective on traffic engineering

- In IP / MPLS networks, traffic engineering is often considered synonymous with RSVP-based MPLS TE
  - It's not the only option …
  - It's also not the only (or even primary) use of MPLS TE in practise
- Traffic engineering is often undertaken without an understanding of the possible benefits … or costs …
- Whilst the concepts are straightforward in theory, there are a number of non-trivial questions to be answered in any deployment

A practical perspective on traffic engineering

- What optimisation objective?
- Strategic or tactical?
- Which approach?
  - IGP TE or MPLS TE
- How often to re-optimise?
- For MPLS TE:
  - Edge mesh, core mesh or by exception?
  - Dynamic or explicit tunnel paths?
  - Tunnel sizing
    - Online or offline sizing
    - How often to resize
- How do you measure the benefit of different approaches?

IP Traffic Engineering: The objective

- What is the primary optimization objective?
  - Either …
    - minimizing maximum utilization in normal working (non-failure) case
  - Or …
    - minimizing maximum utilization under defined failure conditions
- Understanding the objective is important in understanding where different traffic engineering options can help and in which cases more bandwidth is required
  - Other optimization objectives possible: e.g. minimize propagation delay, apply routing policy …
- Ultimate measure of success is cost saving

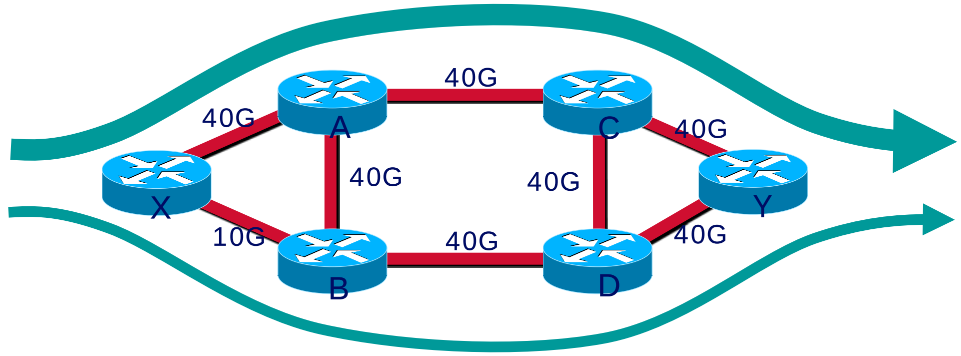

- In this asymmetrical topology, if the demands from X➔Y > 10G, traffic engineering can help to distribute the load when all links are working

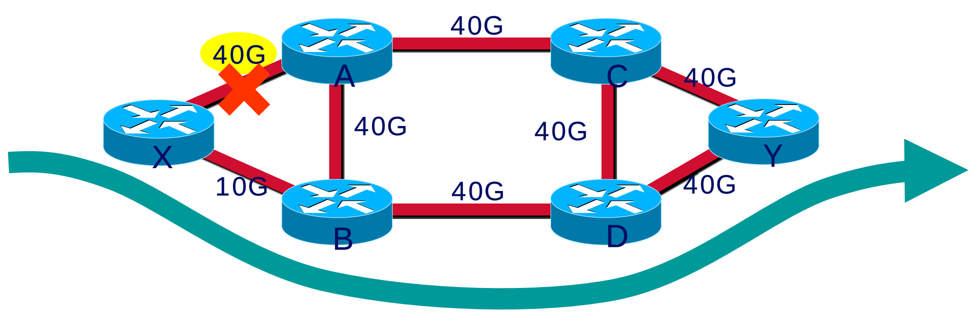

- However, in this topology when optimization goal is to minimize bandwidth for single element failure conditions, if the demands from X➔Y > 10G, TE cannot help - must upgrade link X➔B

Traffic Engineering Approaches

- Technology approaches:
  - MPLS TE
  - IGP Metric based TE (works for IP and MPLS LDP)

- Deployment models:
  - Tactical TE
    - Ad hoc approach aimed at mitigating current congestion
    - Short term operational/engineering process
    - Configured in response to failures, traffic changes
  - Strategic TE
    - Systematic approach aimed at cost savings
    - Medium term engineering/planning process
    - Configure in anticipation of failures, traffic changes
      - Resilient metrics, or
      - Primary and secondary disjoint paths, or
      - Dynamic tunnels, or …

IGP metric-based traffic engineering

- Significant research efforts ...
  - B. Fortz, J. Rexford, and M. Thorup, “Traffic Engineering With Traditional IP Routing Protocols”, IEEE Communications Magazine, October 2002.
  - D. Lorenz, A. Ordi, D. Raz, and Y. Shavitt, “How good can IP routing be?”, DIMACS Technical, Report 2001-17, May 2001.
  - L. S. Buriol, M. G. C. Resende, C. C. Ribeiro, and M. Thorup, “A memetic algorithm for OSPF routing” in Proceedings of the 6th INFORMS Telecom, pp. 187188, 2002.
  - M. Ericsson, M. Resende, and P. Pardalos, “A genetic algorithm for the weight setting problem in OSPF routing” J. Combinatorial Optimization, volume 6, no. 3, pp. 299-333, 2002.
  - W. Ben Ameur, N. Michel, E. Gourdin et B. Liau. Routing strategies for IP networks. Telektronikk, 2/3, pp 145-158, 2001.
  - ...

Comparing TE Approaches

Case Study 1: Performance over Various Networks [Maghbouleh 2002]

- Study on Real Networks
- Single Set of Metrics Achieve 80-95% of Theoretical Best across Failures

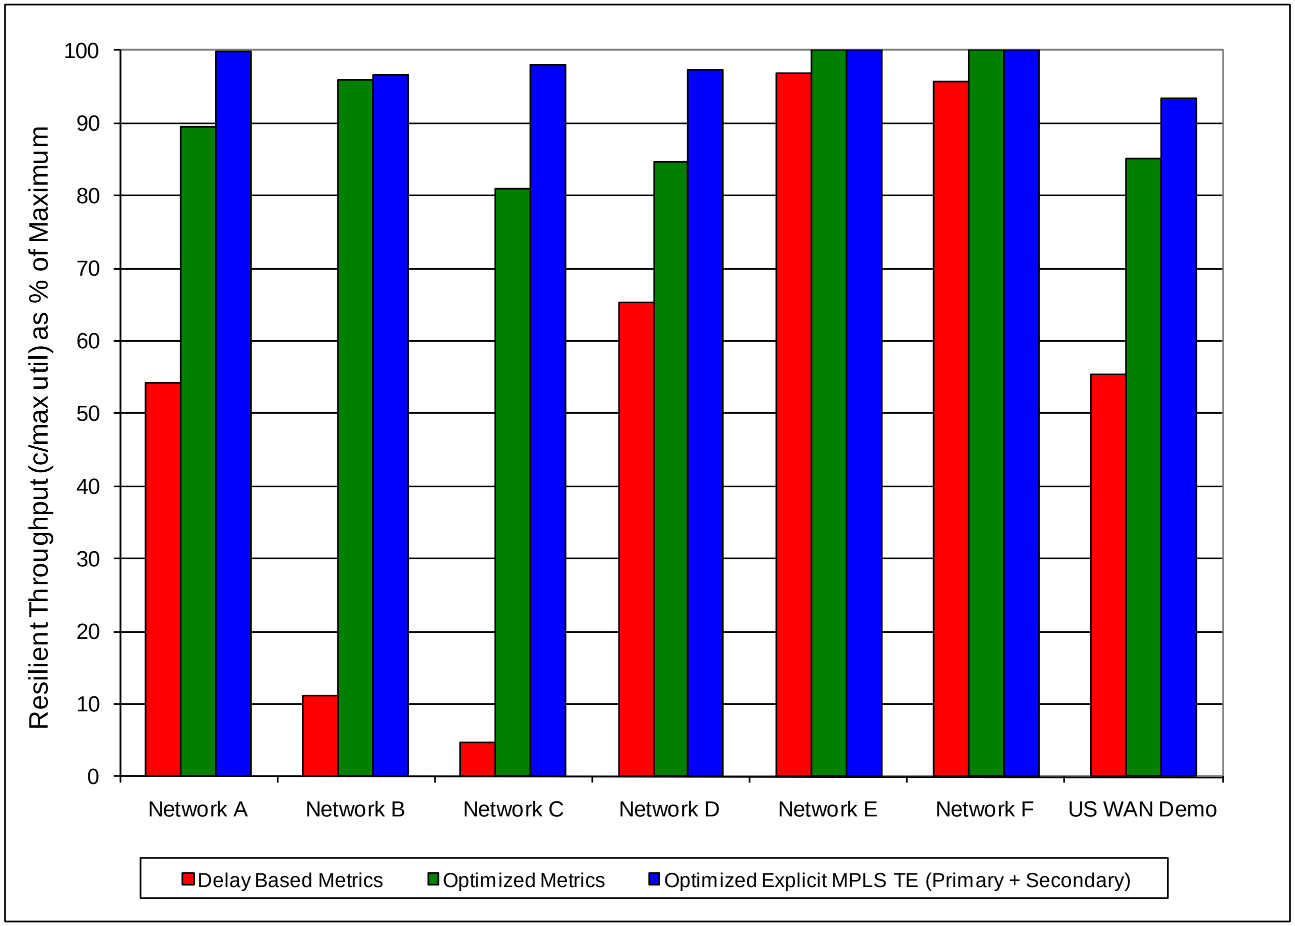

How frequently to reoptimise?

Case Study 2: Variance vs. Bandwidth [Telkamp 2003]

- Around 8000 demands between core routers
- Most traffic carried by (relatively) few big demands
  - 97% of traffic is carried by the demands larger than 1 Mbps (20% of the demands!)
- Relative variance decreases with increasing bandwidth
- High-bandwidth demands are well-behaved (predictable) during the course of a day and across days
- Generally little motivation for dynamically changing routing during the course of a day
- Reoptimisation frequency O(days) rather than O(hours)

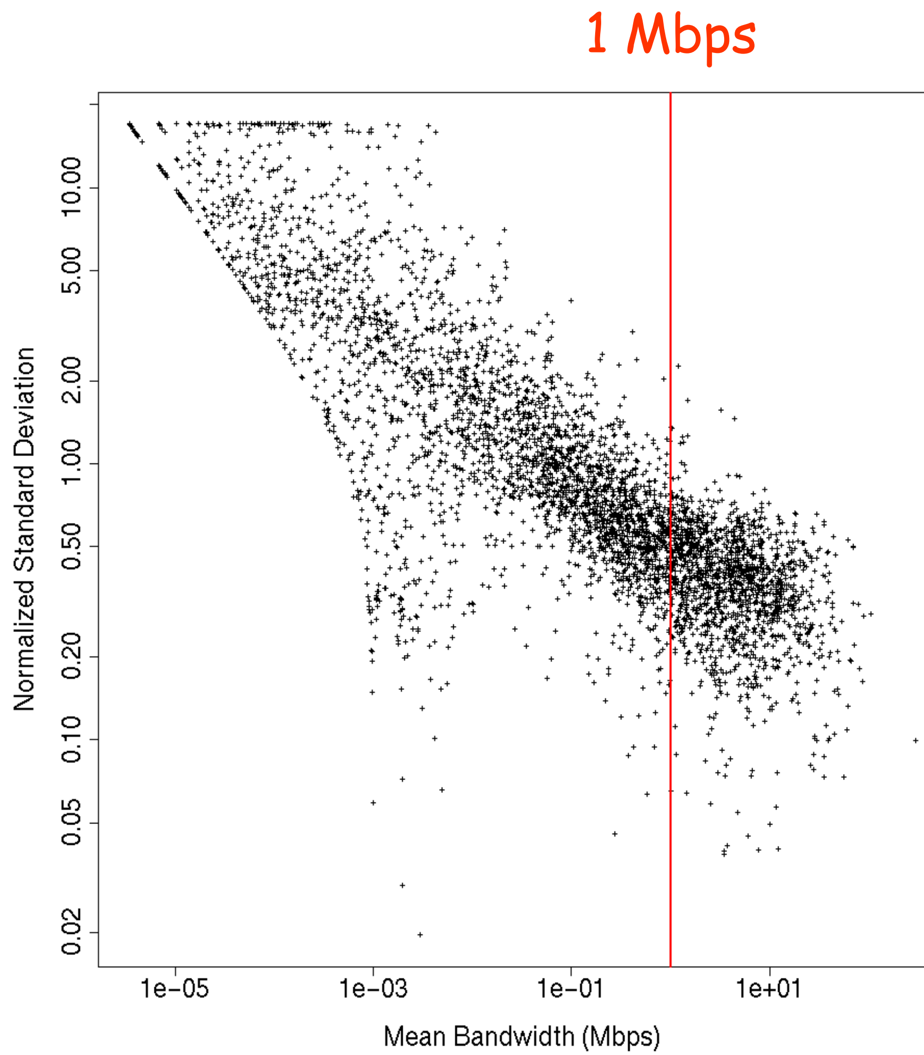

MPLS TE: Edge mesh, core mesh or by exception?

- Edge mesh
  - Requires n*(n-1) tunnels, where n=# of head-ends
  - Significant provisioning and management burden in medium or large networks

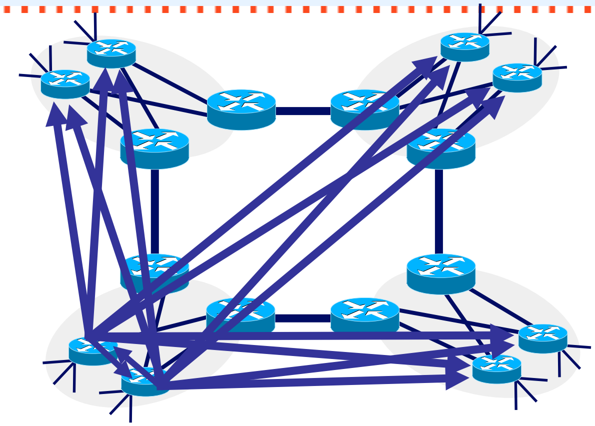

- Core mesh
  - Reduces # tunnels
  - Generally effective for medium to large networks
  - May suffer from “traffic sloshing” – fix with forwarding adjacency

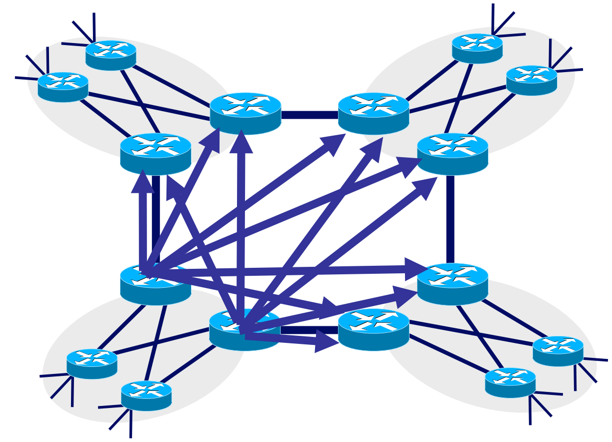

- By exception
  - Useful where the problem is managing a relatively small number or relatively large demands

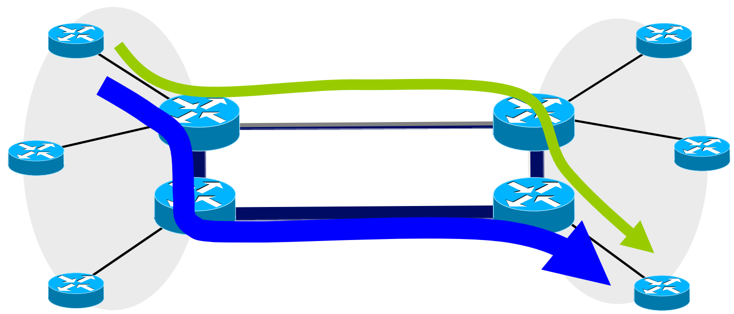

MPLS TE: Edge mesh, core mesh or by exception?

Case Study 3: Example Data from Tier-1 IP Backbone [Telkamp 2007]

- Large network
- Few large nodes contribute to total traffic
  - 20% demands generate 80% of total traffic
- Core mesh or by exception most appropriate solutions

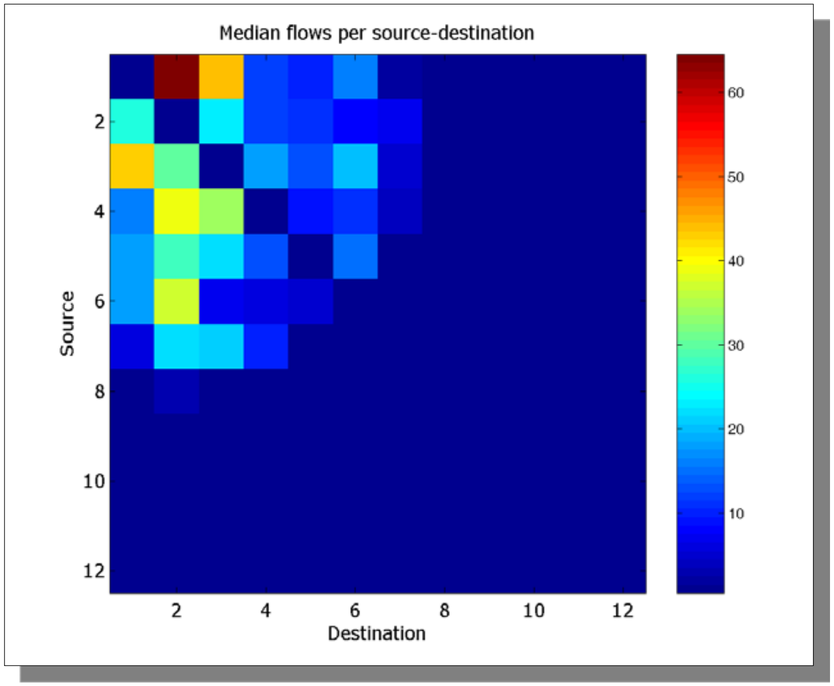

Spatial demand distributions - European subnetwork

MPLS TE deployment considerations

- Dynamic or Explicit tunnel paths?
  - Dynamic path option, i.e. head-end router calculated tunnel path
    - Must specify bandwidths for tunnels
      - Otherwise defaults to IGP shortest path
    - Dynamic tunnels introduce indeterminism and cannot solve "tunnel packing" problem
      - Order of setup can impact tunnel placement
      - Each head-end only has a view of their tunnels
      - Tunnel prioritisation scheme can help – higher priority for larger tunnels
  - Explicit path option, i.e. offline system calculates tunnel path
    - More deterministic, and able to provide better solution to "tunnel packing" problem
      - Offline system has view of all tunnels from all head-ends

Dynamic or Explicit tunnel paths?

Case Study 4: DT: “IGP Tuning in an MPLS Network”, [Horneffer 2005]

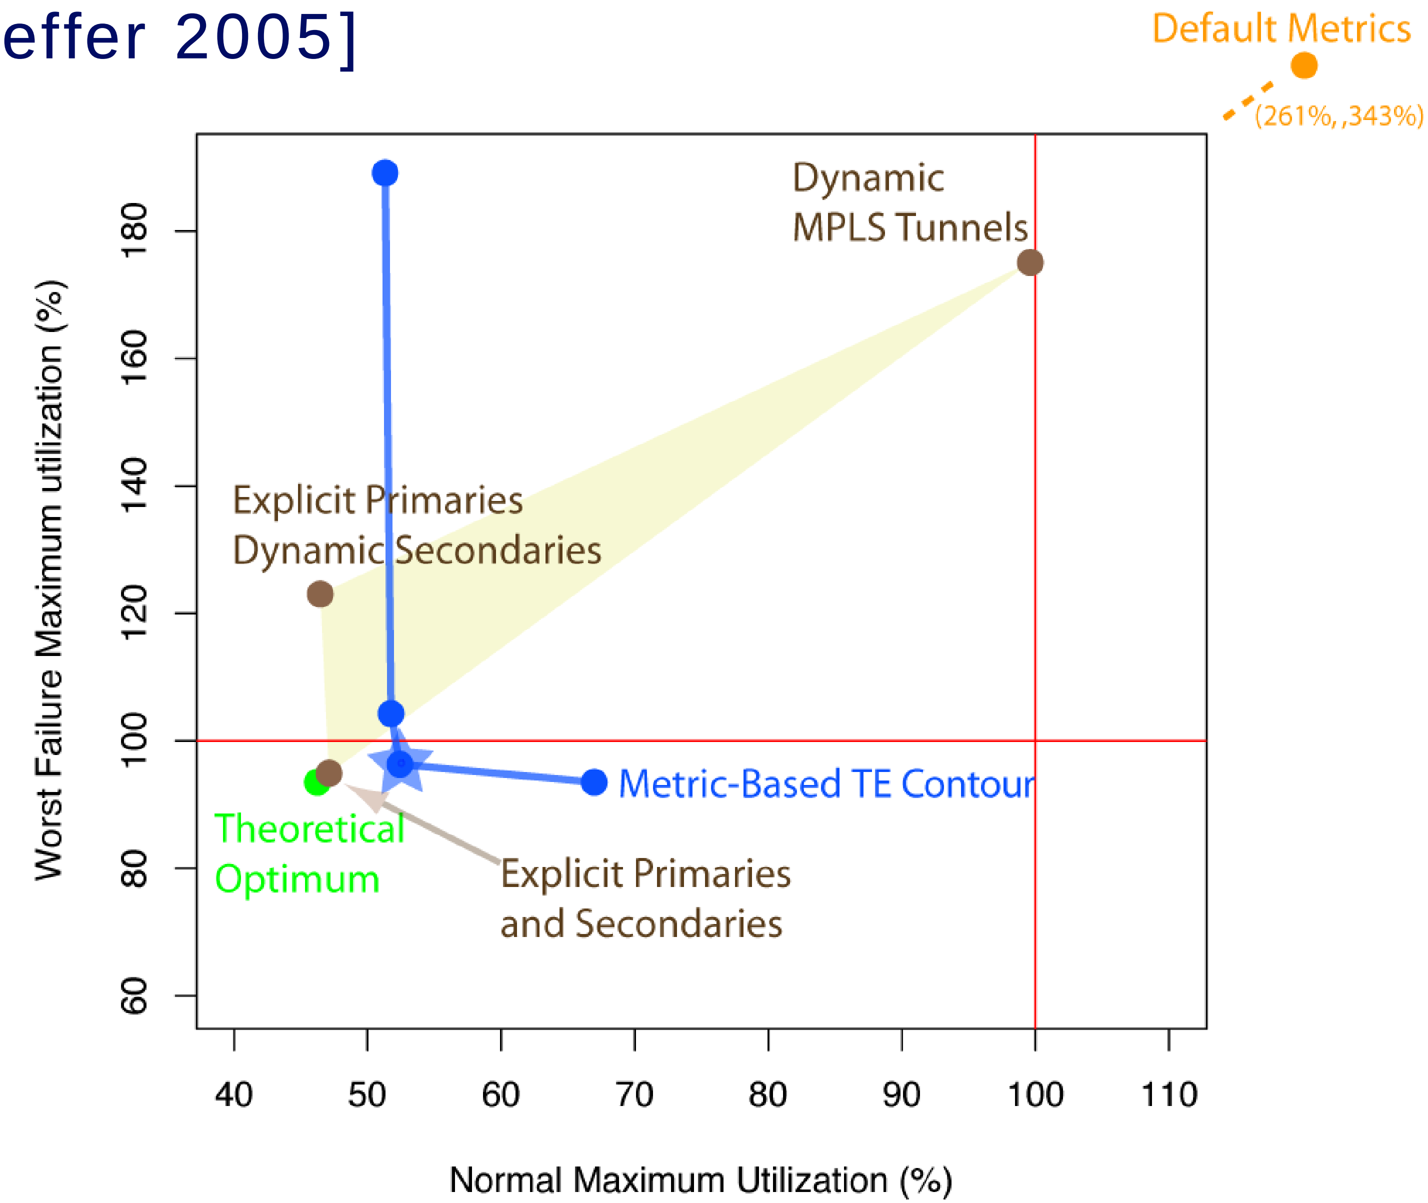

Dynamic or Explicit Tunnel Paths

\+ Default Metrics
x Dynamic MPLS
* Metric-Based TE
o Explicit Pri. + Sec.

Case Study 5: Anonymous network …

- TE Options:
  - Default metrics (no TE)
  - Metric based TE (MATE)
  - Dynamic MPLS TE
    - Mesh of CSPF tunnels in the core network
    - "Sloshing" causes congestion under failure scenarios
  - Explicit MPLS TE (MATE)
  - Failures cases considered
    - Single-circuit, circuit+SRLG, circuit+SRLG+Node
    - Plot is for single-circuit failures
- Cariden MATE software for simulations and optimizations

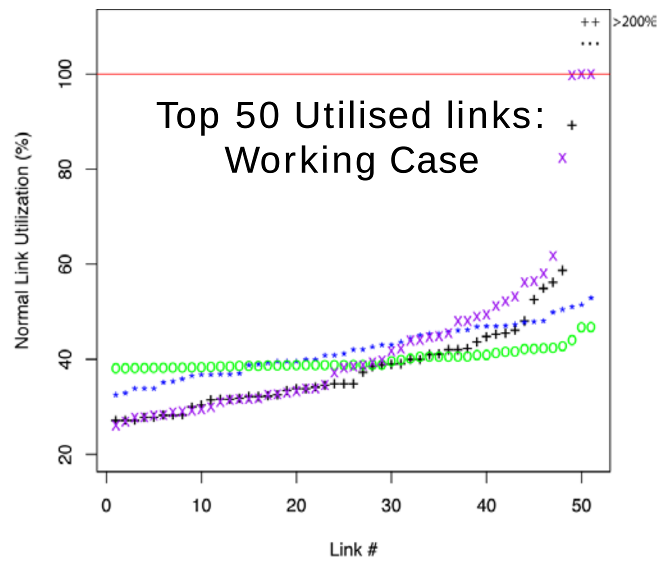

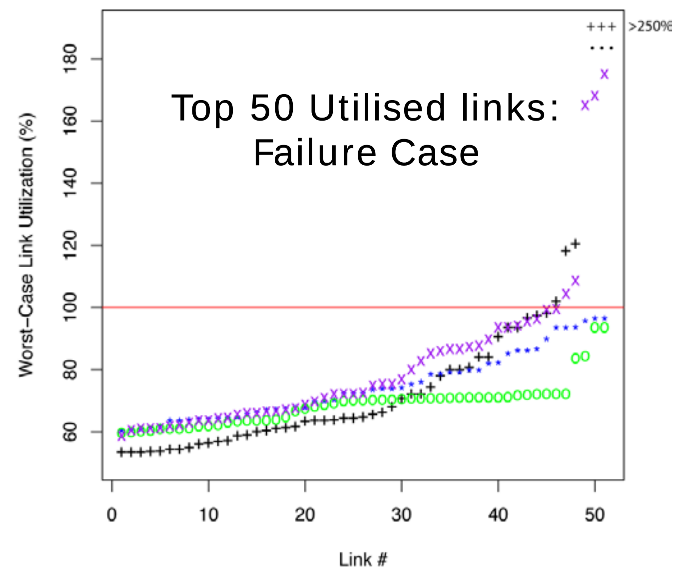

Tunnel Sizing: online vs. offline

- MPLS TE tunnel bandwidth is a one dimensional parameter – no concept of rate/burst
- Tunnel sizing matters ...
  - Needless congestion if actual load >> reserved bandwidth
  - Needless tunnel rejection if reservation >> actual load
- Online vs. offline sizing:
  - Online sizing, i.e. by head-end router: autobandwidth + dynamic path option
    - Router automatically adjusts reservation (up or down) based on traffic observed in previous time interval
    - Time interval is important
    - Tunnel bandwidth is not persistent (lost on reload)
  - Offline sizing, i.e. is specified to head-end router by external system
    - If using explicit path options …
      - … it doesn't really matter (as long as not so high that tunnels are rejected)
    - If using dynamic path options …
      - Use same tunnel sizing heuristic as is used for capacity planning
      - set bw to percentile (e.g. P95) of projected max load over time between optimisations

Tunnel Sizing: How frequently to resize?

Case Study 6: Anonymous network …

- Resizing can be online (i.e. by head-end router) or offline (i.e. by external system)
- Possible inefficiencies or congestion if periodicity too low
- Online resizing too often can result in “bandwidth lag”
- Periodically readjust – O(days) rather than O(hours)

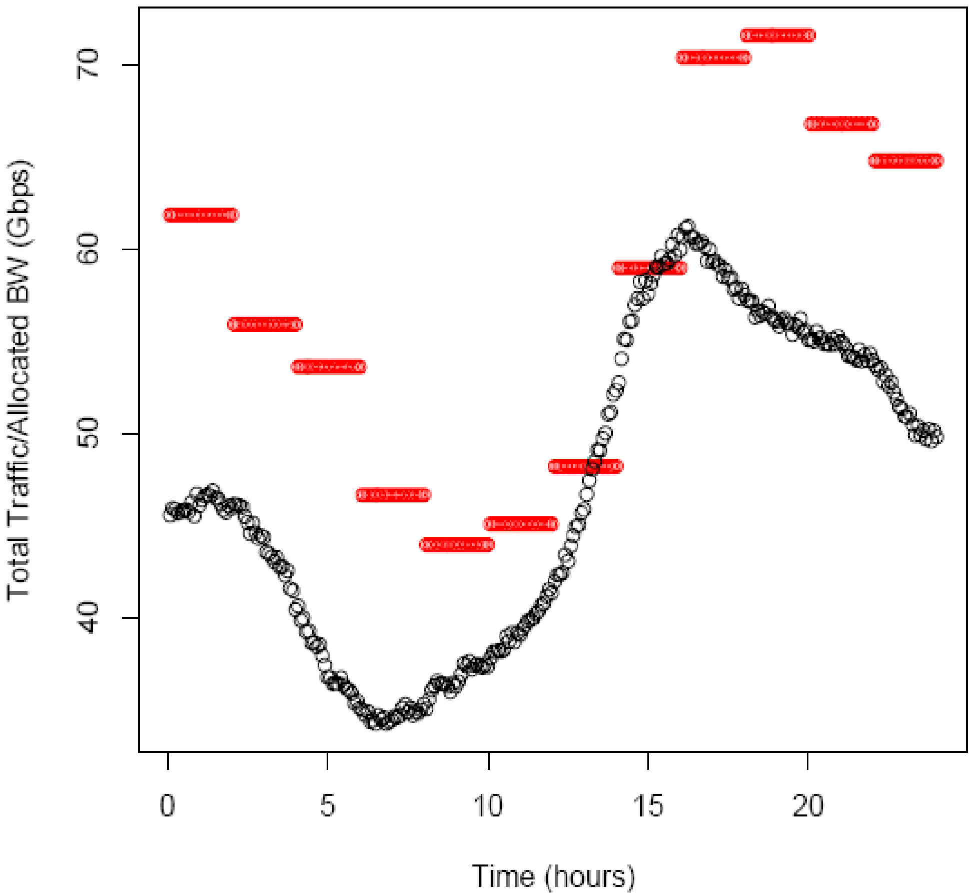

“online sizing: bandwidth lag”

A practical perspective on traffic engineering

- Need to define whether optimising for working case or failure case
- Deployment choices
  - Tactical vs. strategic
  - IGP metric based TE vs. RSVP TE
  - RSVP-TE
    - Choice of core, edge mesh or by exception
    - Explicit path options can be more deterministic/optimal, but requires offline tool
    - Offline tunnel sizing allows most control – use same tunnel sizing heuristic as is used for capacity planning
    - Re-optimisation and resizing O(days) is generally sufficient
- How do you measure the benefit of different approaches?

Network Planning Methodology

- Need to know traffic matrix to be able to simulate and compare potential approaches
- Ultimate measure of success is cost saving

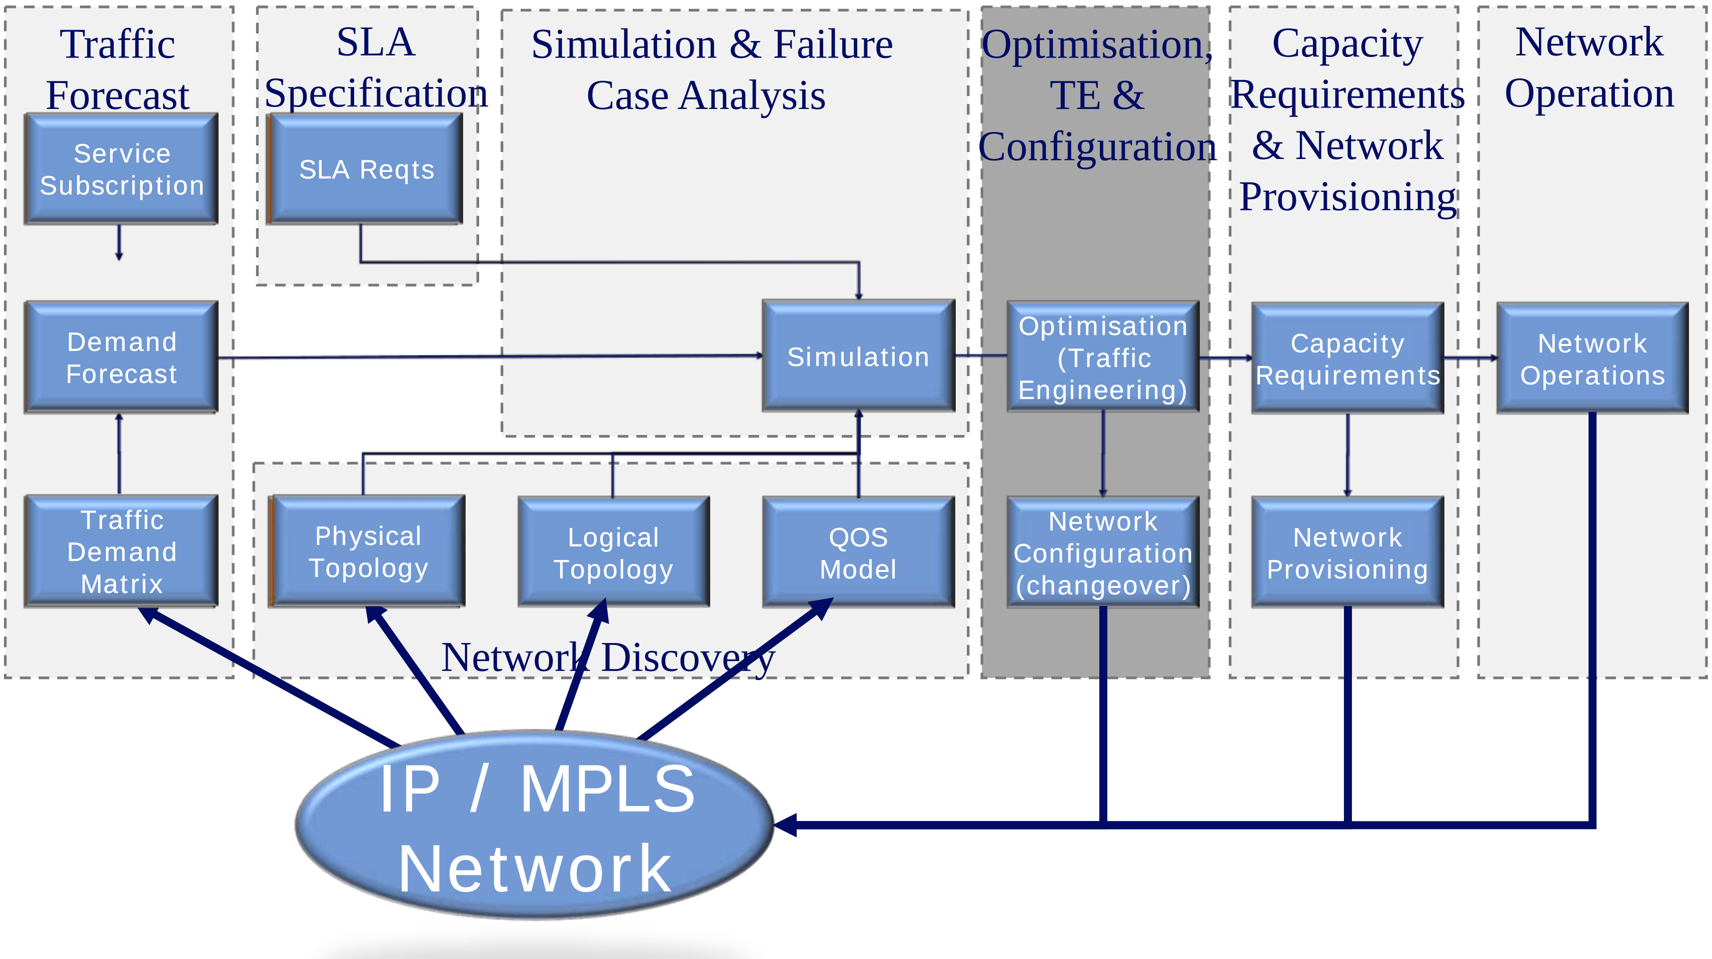

References

- Horneffer 2005
  - Martin Horneffer, “IGP Tuning in an MPLS Network”, NANOG 33, February 2005, Las Vegas
- Maghbouleh 2002
  - Arman Maghbouleh, “Metric-Based Traffic Engineering: Panacea or Snake Oil? A Real-World Study”, NANOG 26, October 2002, Phoenix
- Telkamp 2003
  - Thomas Telkamp, “Backbone Traffic Management”, Asia Pacific IP Experts Conference (Cisco), November 4th, 2003, Shanghai, P.R. China
- Telkamp 2007
  - Thomas Telkamp, Best Practices for Determining the Traffic Matrix in IP Networks V 3.0, NANOG 39, February 2007, Toronto

- All available from:
  - http://www.cariden.com/technologies/papers.html

cariden

the economics of network control

- Web: http://www.cariden.com
- Phone: +44 117 230 2405
- Address: One Redcliff Street
  Bristol, BS1 6TP
  United Kingdom