the economics of network control

Metric-Based Traffic Engineering: A Real World Study.

Arman Maghbouleh
am @ cariden.com

Apricot 2004
Kuala Lumpur, Malaysia

Agenda

- TE Introduction
- Study Outline
  - Networks
  - Routing Models
- Results
- MPLS Notes
- Conclusion

IGP Traffic Engineering

- Manipulate Internal Routing
- Balance Traffic
  - Minimize Maximum Utilization
    - Single-Element Failure Conditions (typical)
- Save Money

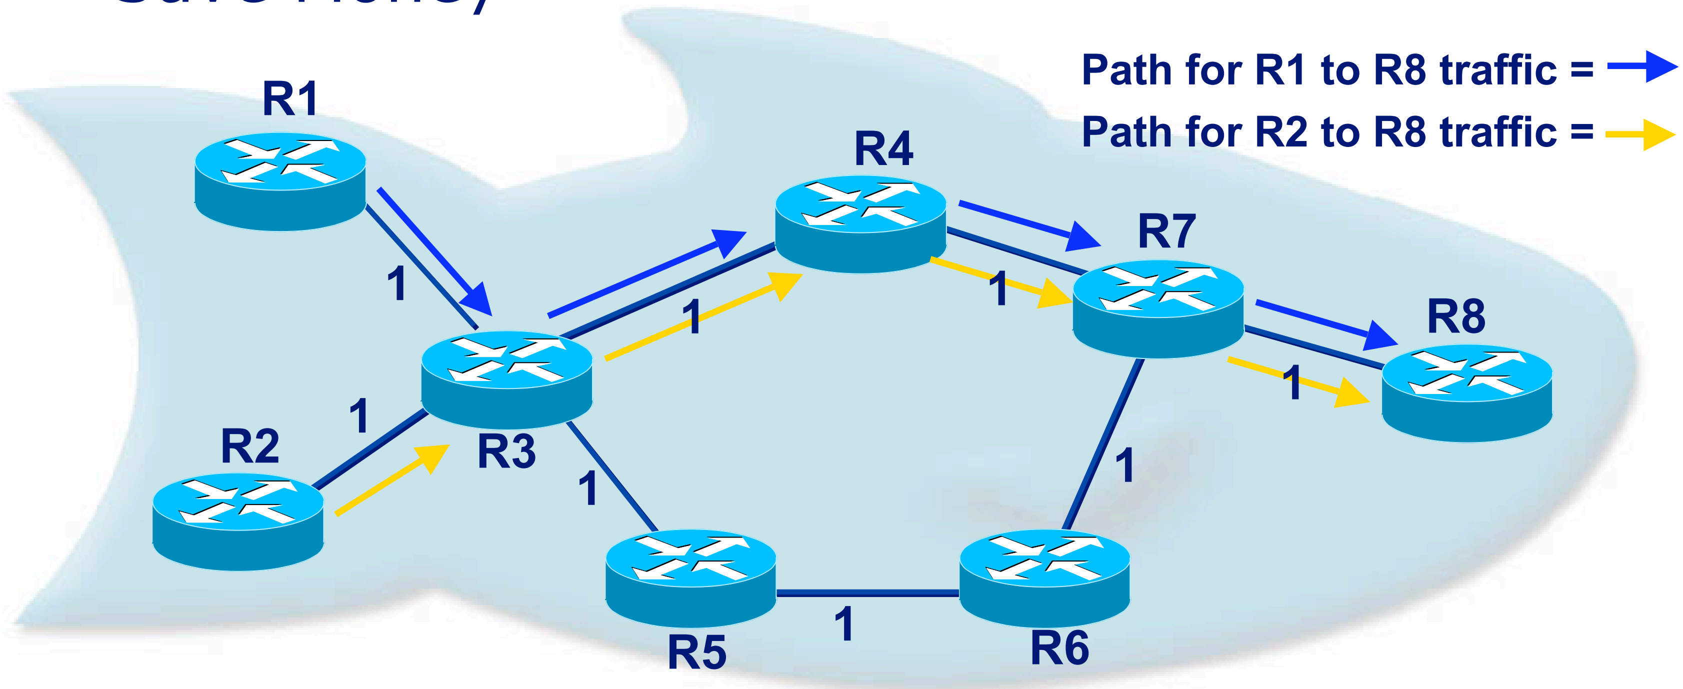

TE Payback

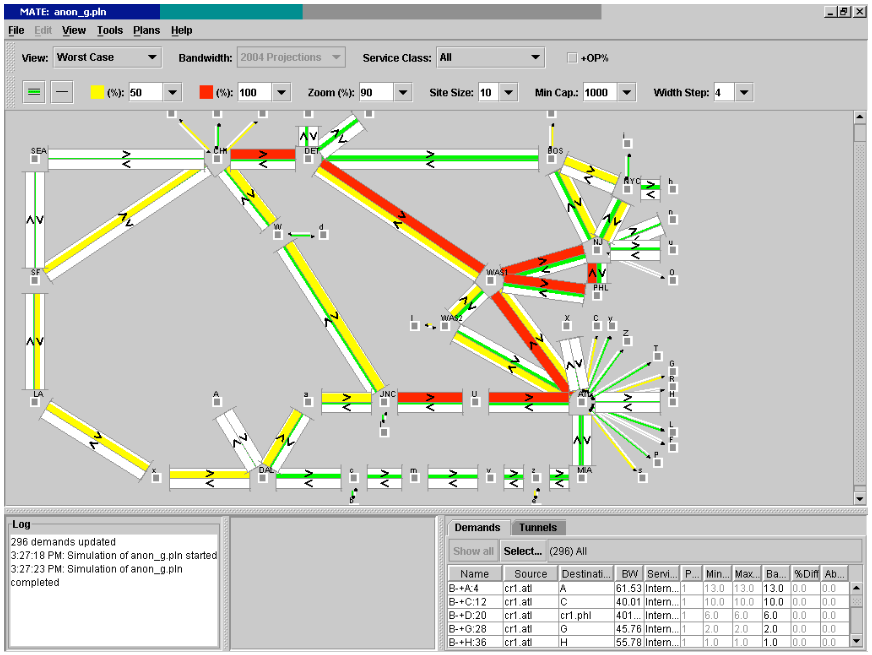

Without TE

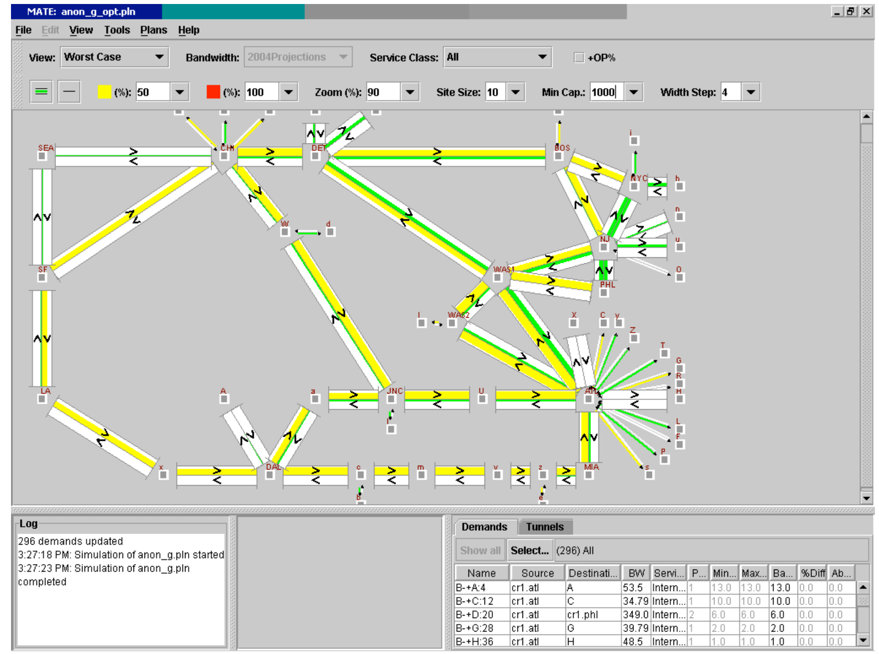

With TE

- Real Example
  - Delay 6 OC-192 Circuits for a year
    (17 circuits under 50% upgrade policy)
  - Capital + Operational Savings ≈ $1M/OC-192/year

Conventional Thinking

- *IP Routing Not Enough Control for TE*
  - *Path computation using a Simple Additive Metric*
  - *Bandwidth availability is not taken into account*
  - *Metric manipulation merely shifts problem*
- *Need Source-Based (ATM/MPLS) for TE*

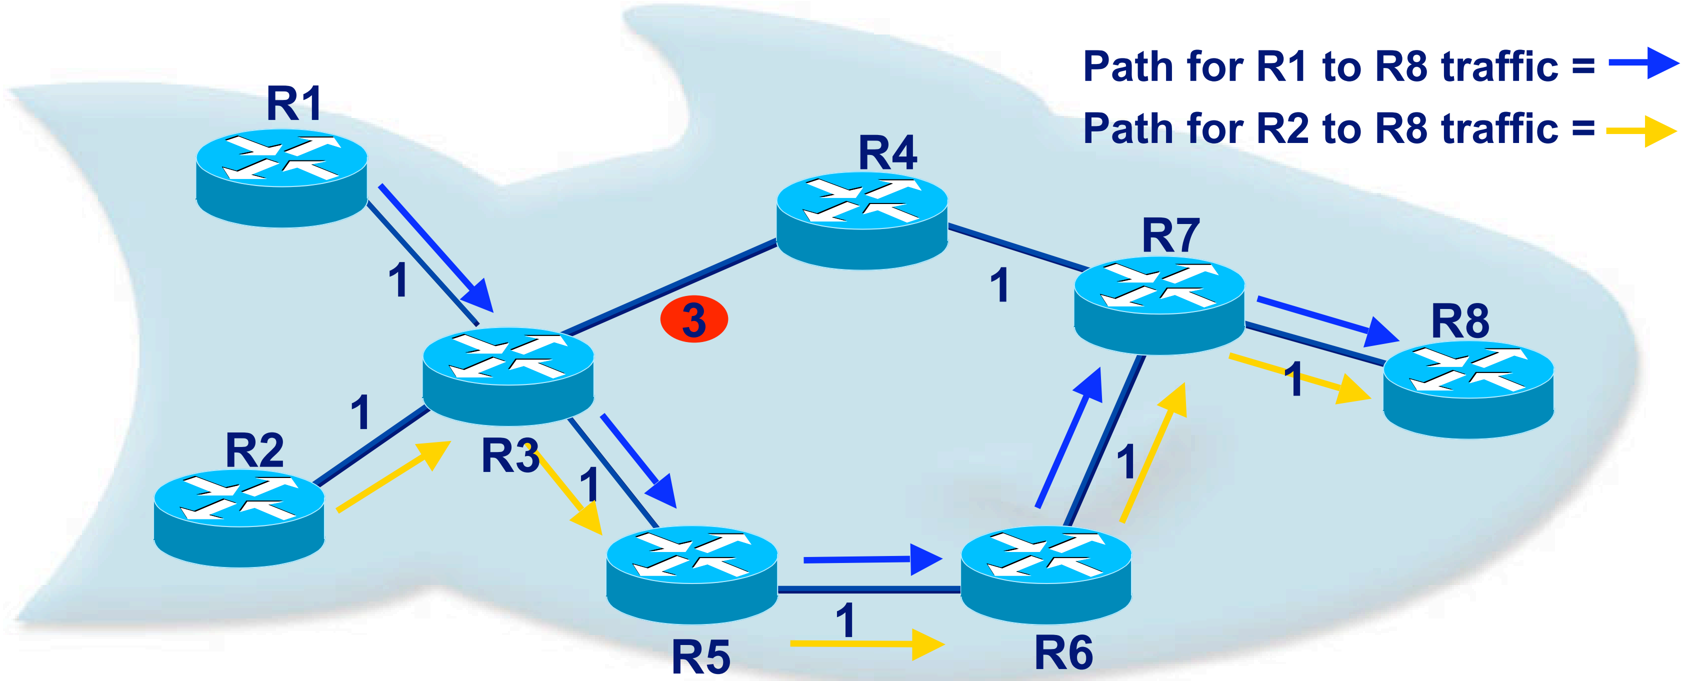

Challenge to Conventional Thinking

- Scientific Advances in Metric Optimization
  - Balance traffic using SPF metrics
  - Use Equal Cost Multipath (ECMP) as necessary

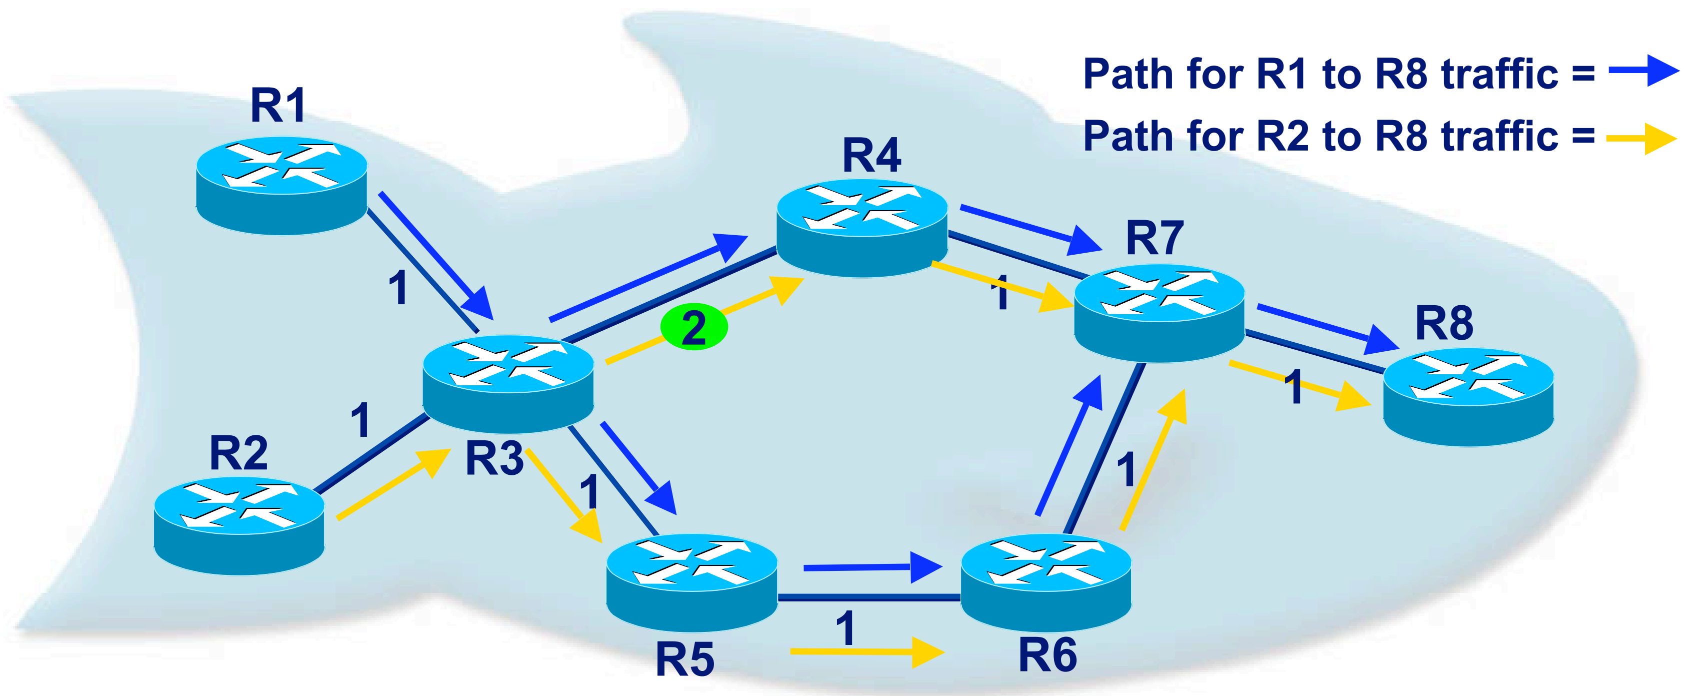

Where is Reality?

Fortz et al.

"… we can find [OSPF] weight settings …[that] get within a few percent of the best possible with general routing,including MPLS."

– (IEEE 2002)

Lorenz et al.

"Source invariant routing can be significantly worse than than per-flow routing."

– (DIMACS 2001)

"… weight setting for OSPF cannot replace MPLS as a traffic engineering tool."

– (IETF-RR list 2001)

Study Outline

- TE Introduction
- Study Outline
  - Networks
  - Routing Models
- Results
- MPLS Notes
- Conclusion

Study

- Six real networks

- Compare Efficiency
  - Optimized Metrics versus

… ———————————◆

Delay Metrics | Explicit Routing | Theoretical Optimal

- Optimizations
  - Objective: Minimize Maximum Utilization
    - UNDER ALL POSSIBLE SINGLE-CIRCUIT FAILURES
  - Inputs: Topology, Link Capacities, Demand Matrix
  - Outputs: Explicit Paths, or Link Metrics

Results (Preview)

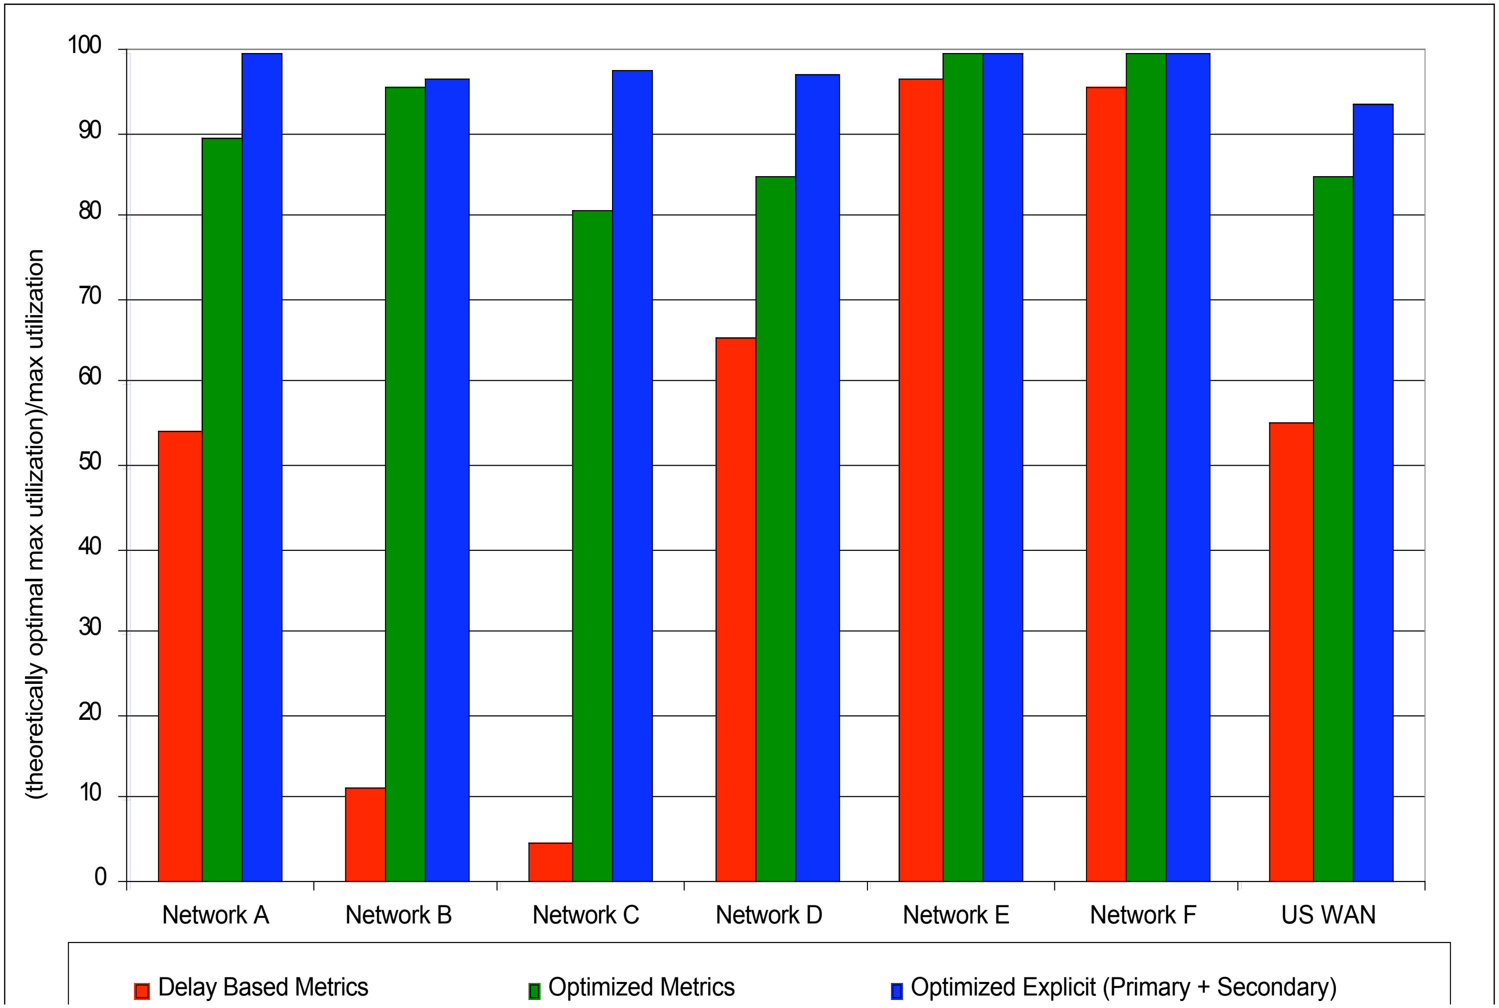

Conclusions (Preview)

- Metric-Based TE Close to Optimal TE
  - Within uncertainty of demand forecasts
- Some topologies limit TE benefits
- SPF Limitations do not affect the bottom-line significantly
- Metric Optimization as alternative or complement to MPLS TE

Networks

- Tier 1, tier 2, content-delivery network
- Global, U.S., Europe
- Some already deployed MPLS
  - Measured versus estimated traffic matrix
- Five operational, one proposed
- Topologies
  - V-O-V
  - Typical U.S. Meshes
  - Global Mesh

V-O-V Networks

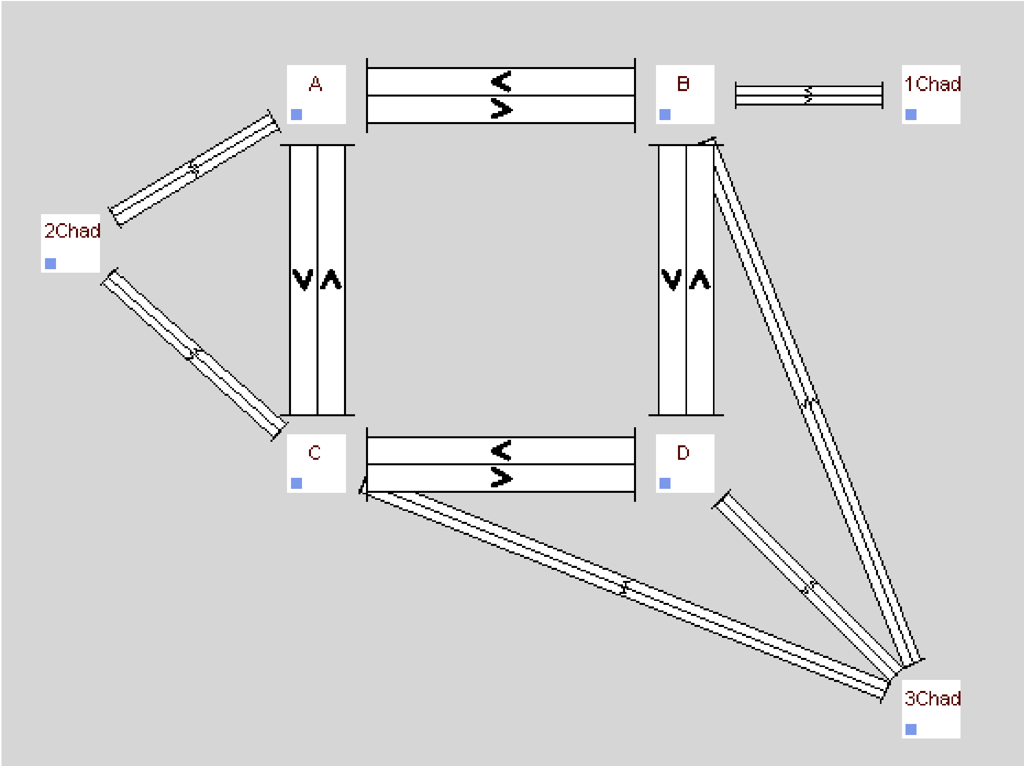

- High capacity simple core
- Peripheral nodes connected
  - Singly, doubly, and infrequently triply

Typical U.S. Backbone

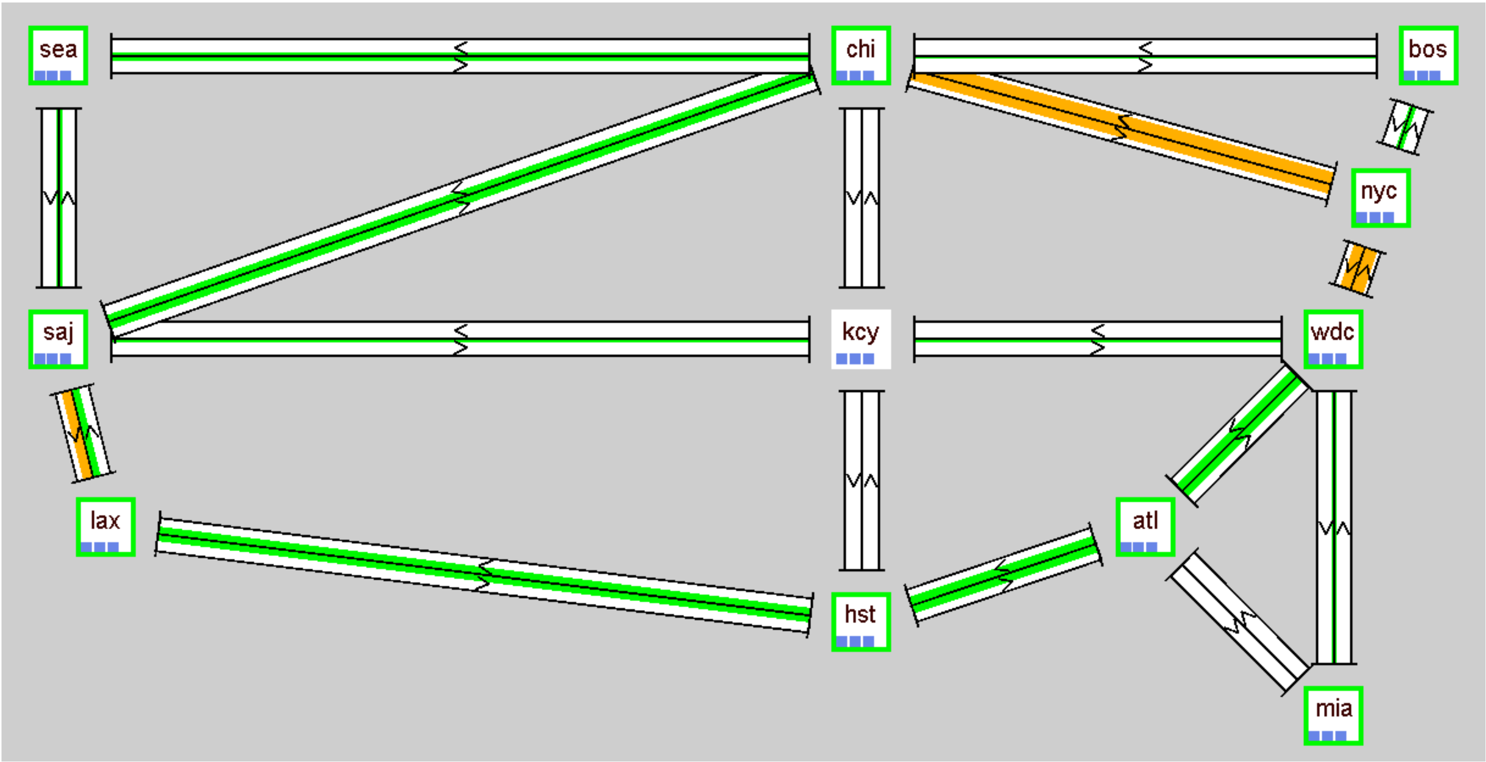

- Three+ paths across country
- Elephants and mice demands

cariden

Plot Legend

- White squares represent sites (PoPs)
  - Small blue squares represent routers
- Lines are physical links
  - Thickness represents capacity
  - Color & fill thickness represents utilization
    - (red >90%, orange >75%)
- Blue arrows represent paths
  - (solid for normal, dashed for failure)
- X represent failure locations

Global Meshes

- No prototype shown for confidentiality
- Combinations of meshes, rings,…
  - Topology bottlenecks across oceans
- Large range of capacities
  - (e.g. OC-3 to OC-192)

Routing Models

- TE Introduction
- Study Outline
  - Networks
  - Routing Models
- Results
- MPLS Notes
- Conclusion

Theoretical Optimal

- Result of multicommodity flow optimizations
  - One per failure scenario
- No shortest-path limitation
  - I.e., possibly source-based routing

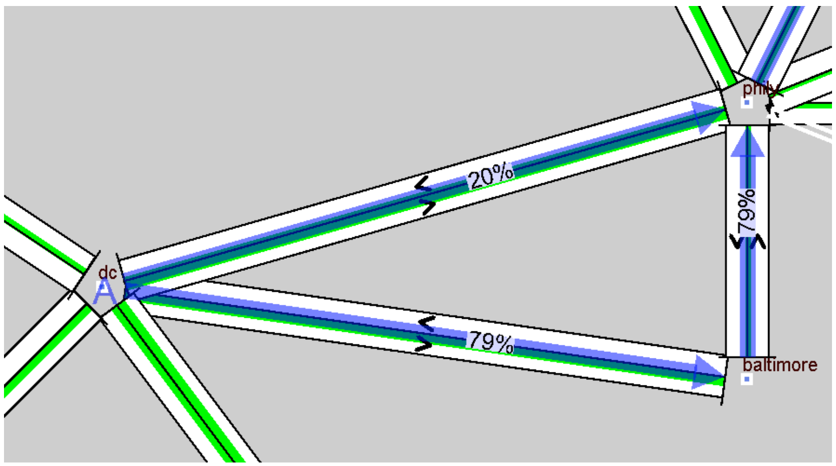

Arbitrary splits of demands

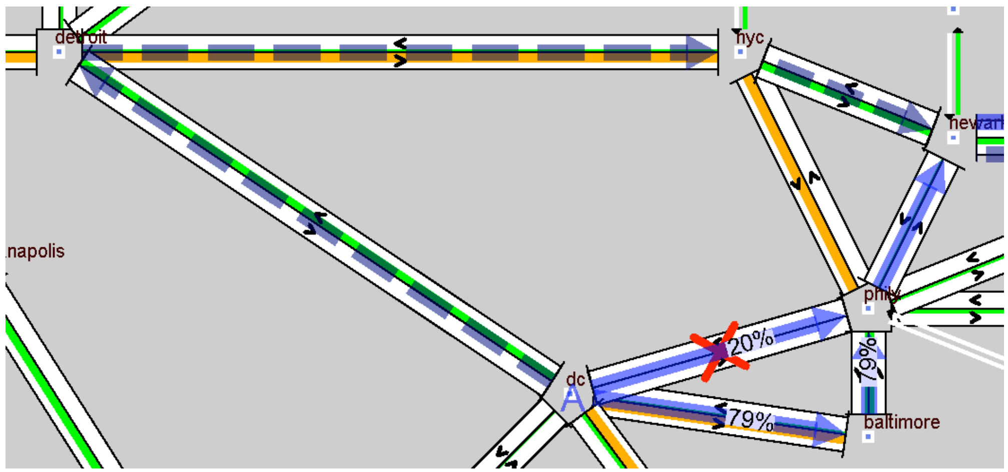

Routing changed on failure

*Real case used with permission.

Shortest-Path Metric Routing

- OSPF, IS-IS
- 1/n Equal-Cost Multipath
- Use Cariden Software to determine metrics
- Single set of metrics designed for resilience

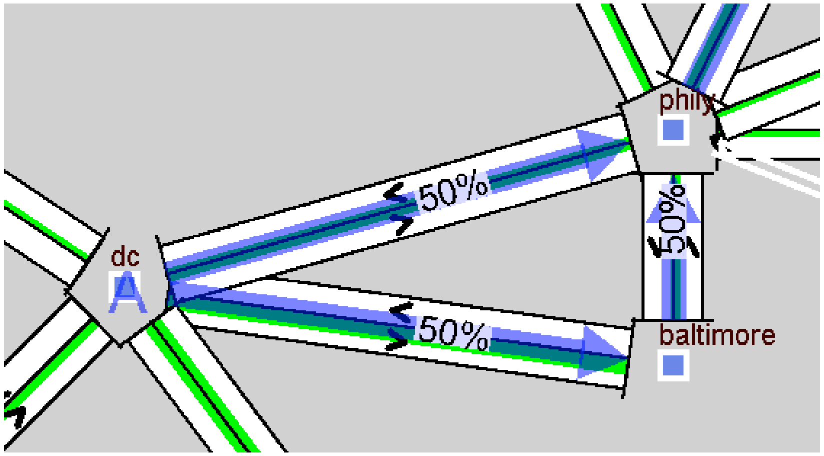

Equal splits on ECMP

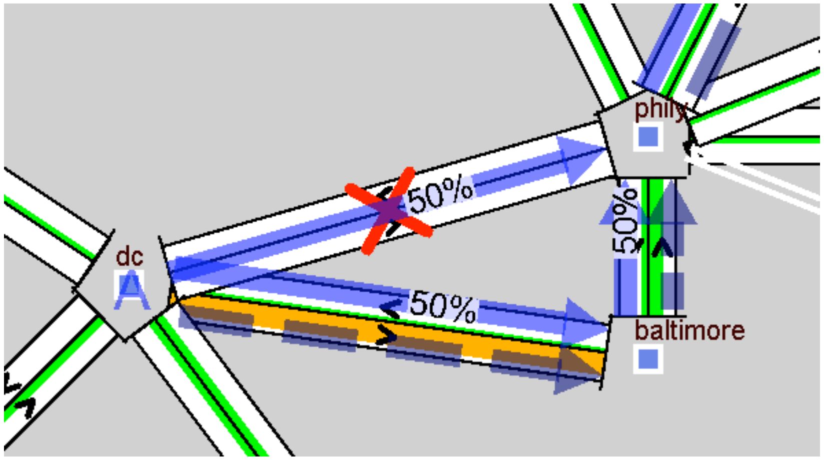

Metrics not change after failure

Explicit Routing

- A primary and secondary path for each source-destination pair
  - Link-diverse secondary
- No reservations
- Used Cariden Software to find optimal paths

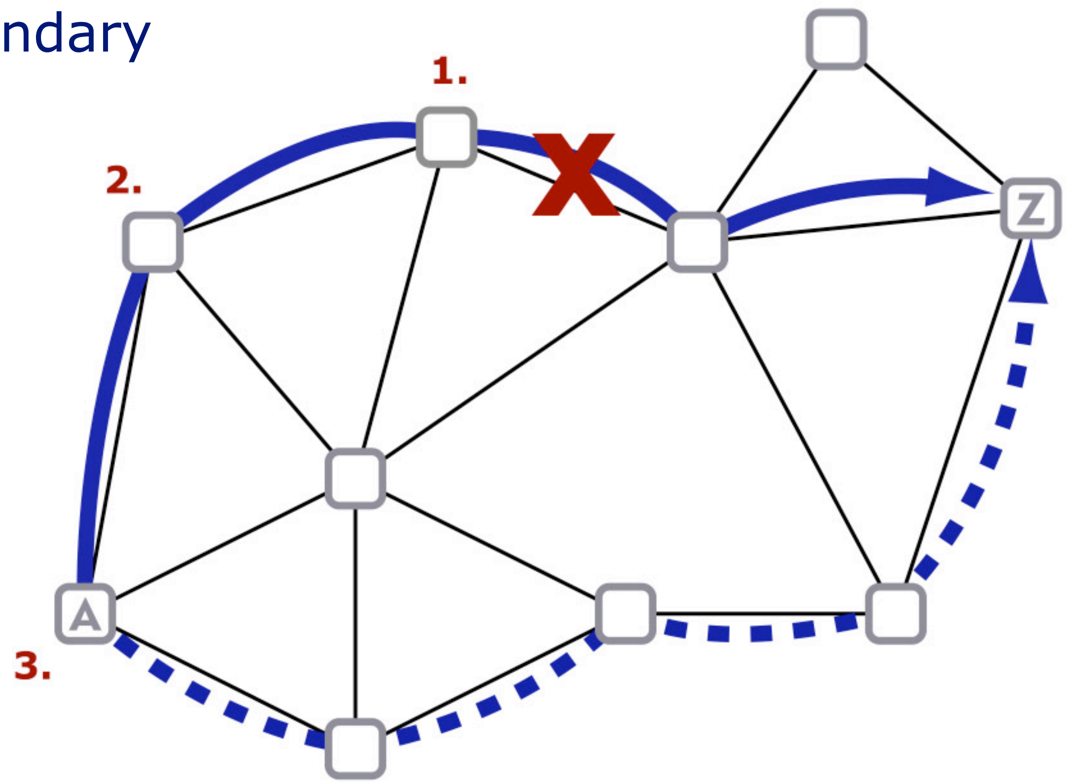

Results

- TE Introduction
- Study Outline
  - Networks
  - Routing Models
- Results
- MPLS Notes
- Conclusion

Results

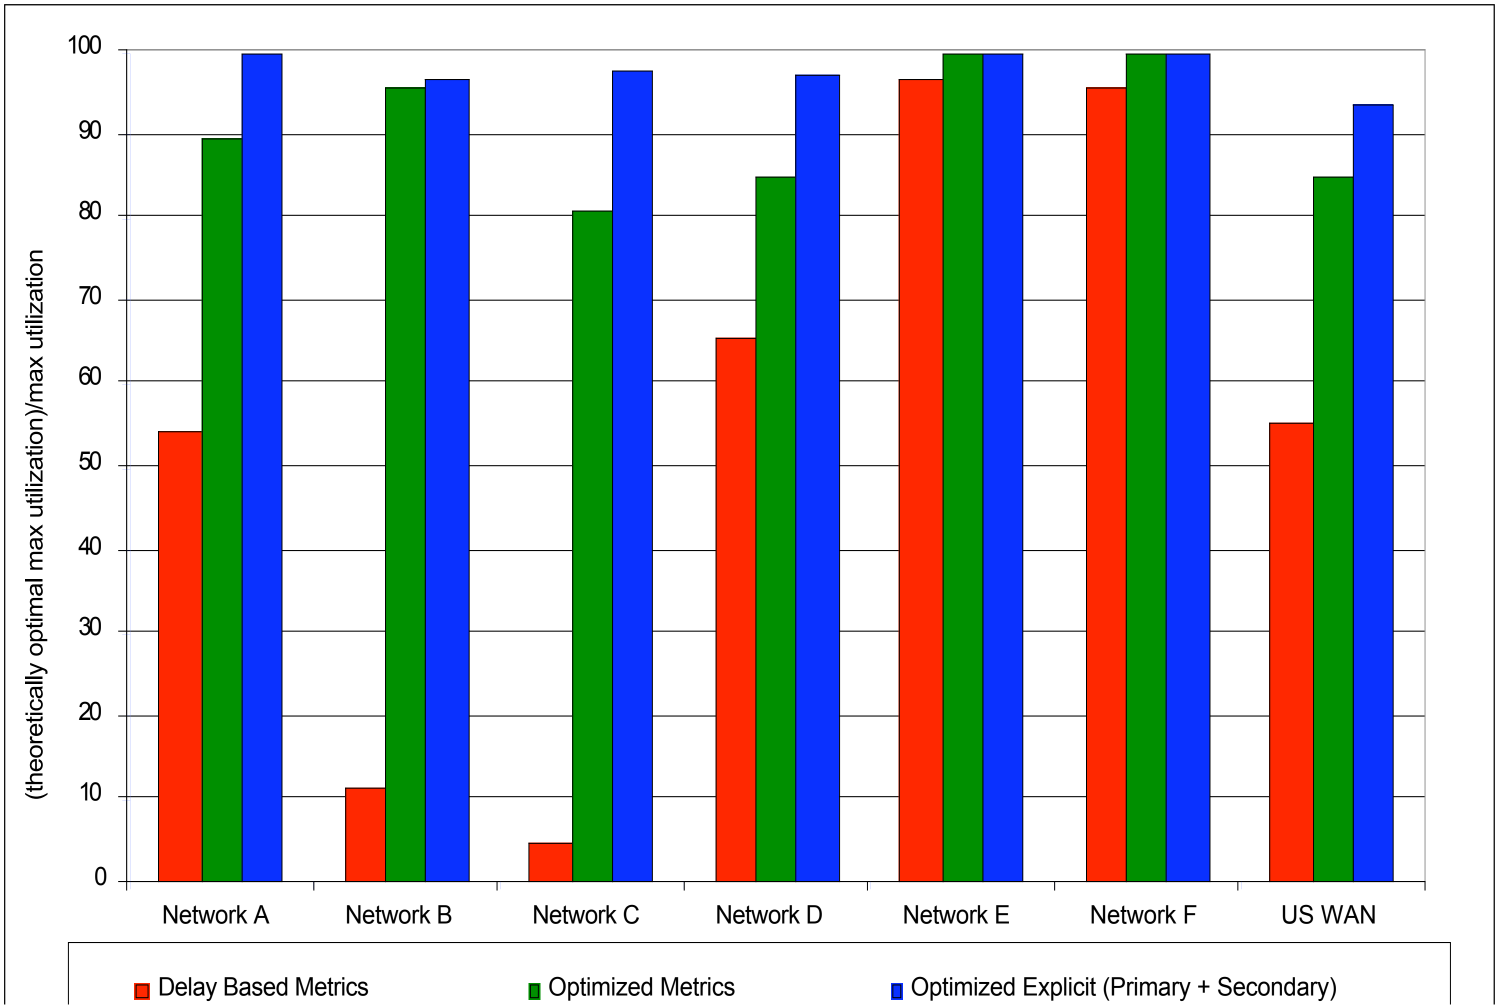

(theoretically optimal max utilization)/max utilization
100
90
80
70
60
50
40
30
20
10
0
Network A
Network B
Network C
Network D
Network E
Network F
US WAN
Delay Based Metrics
Optimized Metrics
Optimized Explicit (Primary + Secondary)

cariden

Results (in text)

- Can optimize SPF metrics within 80%-95% of maximum theoretical efficiency
  - … trivially at 100% for V-O-V topologies

- Explicit routing around 90-95% of theoretical limit

Outline

- Introduction
- Study Outline
  - Networks
  - Routing Models
- Results
- MPLS Notes
- Conclusion

Hybrid: SPF-Explicit

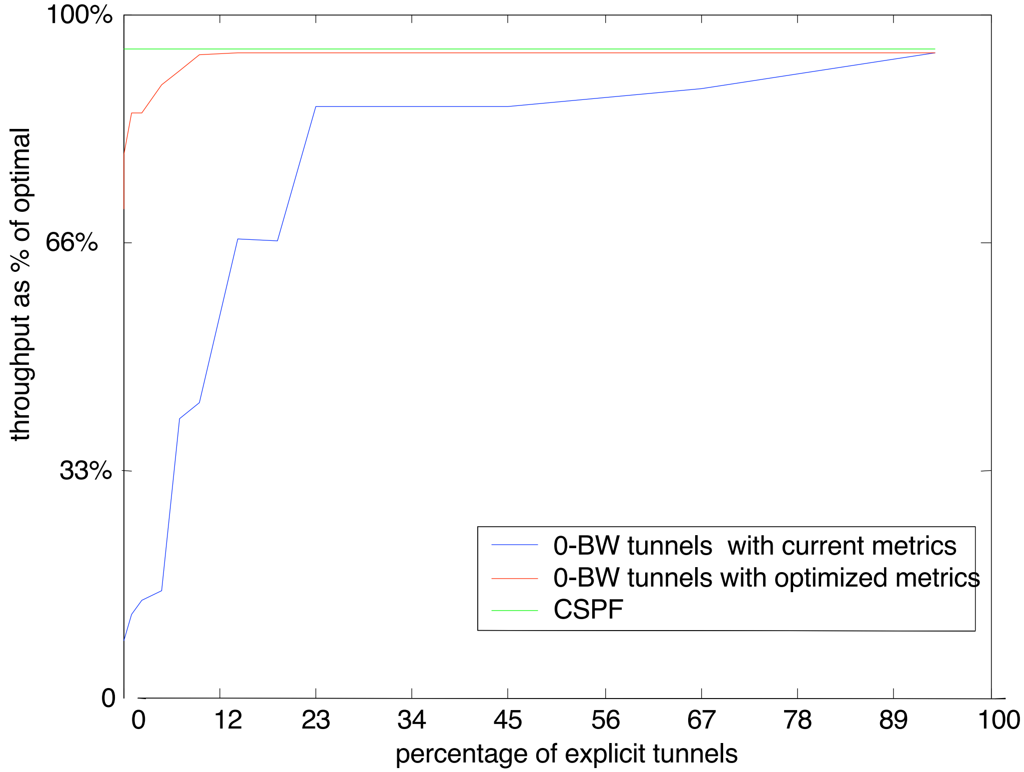

X-axis tunnels were explicitly routed from largest to smallest.

- Metric optimization + explicit routes as needed

  "We expect this is not an unreasonable approach."
  -Randy Bush (NANOG 26)

  - **Also: Ben-Ameur et al. France Telecom, draft-wang-te-hybrid-approach-00.txt**

- Few tunnels explicit if start with good metrics

***Real case used with permission.**

Conclusions

- Metric-Based TE Close to Optimal TE
  - Within uncertainty of demand forecasts
- Some topologies limit TE benefits
- SPF Limitations do not affect the bottom-line significantly
- Metric Optimization as alternative or complement to MPLS TE

References

- B. Fortz, J. Rexford, and M. Thorup, "Traffic Engineering With Traditional IP Routing Protocols" in 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.
- Cariden "IGP Traffic Engineering Case Study", Cariden Technologies, Inc., October 2002.
- B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights" in Proceedings of IEEE INFOCOM, March 2000.
- B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world" IEEE Journal on Selected Areas in Communications, volume 20, pp. 756-767, May 2002.
- 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.

Acknowledgments

- Andrew Lange (Cable & Wireless/Exodus)
- Thomas Telkamp (Global Crossing)
- Clarence Filsfils (Cisco)
- Glen Larwill (Comcast)
- Jennifer Rexford (AT&T Labs)