# GreenTE: Power-Aware Traffic Engineering

Beichuan Zhang (University of Arizona)

# Limits being reached

- Routers are becoming more powerful as well as more energy hungry.
  - E.g., a CRS router with max config: 92 Tbps and 1 MegaWatt !
  - Driven by exponential growth of data traffic.

# Limits being reached

- Routers are becoming more powerful as well as more energy hungry.
  - E.g., a CRS router with max config: 92 Tbps and 1 MegaWatt !
  - Driven by exponential growth of data traffic.

- Limits are being reached:
  - Power delivery to the total system as well as to components.
  - Cooling router components
  - Heat removal from the box and the facility

# Limits being reached

- Routers are becoming more powerful as well as more energy hungry.
  - E.g., a CRS router with max config: 92 Tbps and 1 MegaWatt !
  - Driven by exponential growth of data traffic.

- Limits are being reached:
  - Power delivery to the total system as well as to components.
  - Cooling router components
  - Heat removal from the box and the facility

- Router energy consumption is becoming an important issue for ISPs, IXPs, and Data Centers.
  - System performance, financial, environmental implications.

# What can we do?

- Individual routers
  - More energy efficient hardware
  - Better energy management of components given the workload.

# What can we do?

- Individual routers
  - Better energy efficiency in hardware
  - Better energy management of components given the workload.
- Networks
  - Routing affects router's workload, thus its energy consumption.

# What can we do?

- Individual routers
  - Better energy efficiency in hardware
  - Better energy management of components given the workload.

- Networks
  - Routing affects router's workload, thus its energy consumption.
  - ***Traffic engineering for better network-wide energy efficiency?***

# Today's ISP networks

- Built for service availability
  - Overprovision of link capacity
  - Redundant links/paths
  - Load-balancing traffic engineering
  - Lead to low average link utilization.

# Today's ISP networks

- Built for service availability
  - Overprovision of link capacity
  - Redundant links/paths
  - Load-balancing traffic engineering
  - Lead to low average link utilization.
- Not very efficient in using energy
  - Routers/links are up 24x7 at full capacity, regardless of workload.
  - Both opportunities and challenges.

# Opportunities in power-aware TE

- Take advantage of path redundancy and low link utilization.

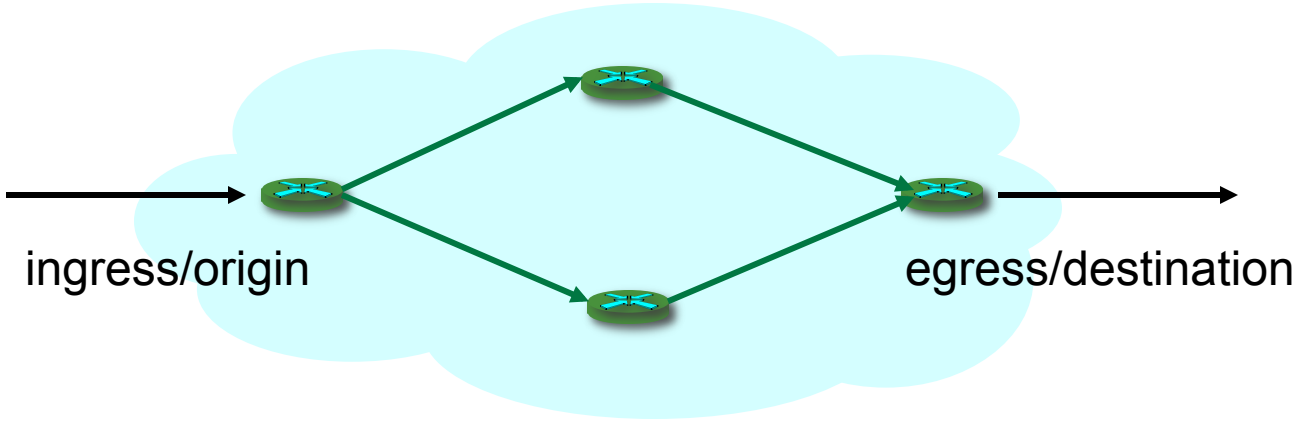

# Opportunities in power-aware TE

- Take advantage of path redundancy and low link utilization.

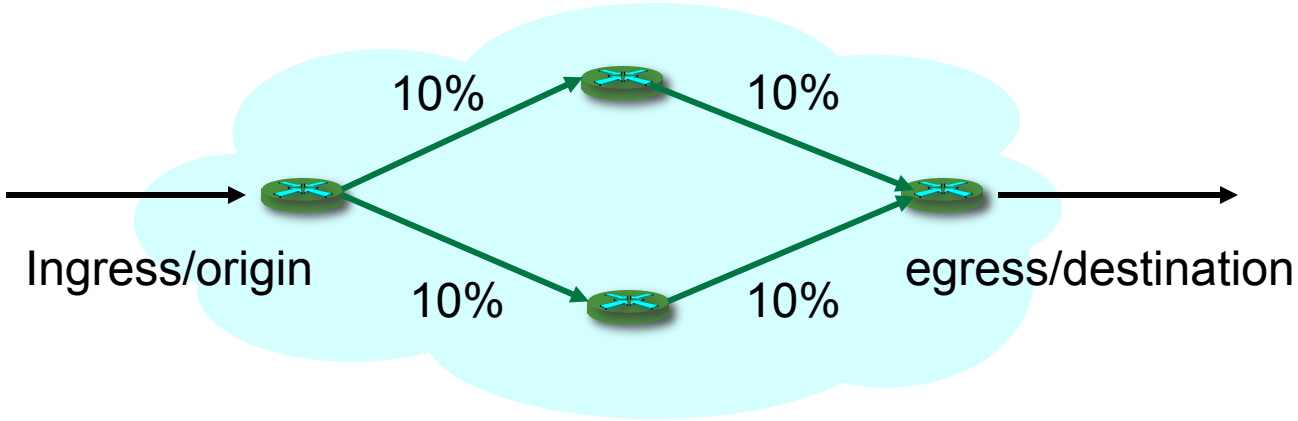

# Opportunities in power-aware TE

- Take advantage of path redundancy and low link utilization.

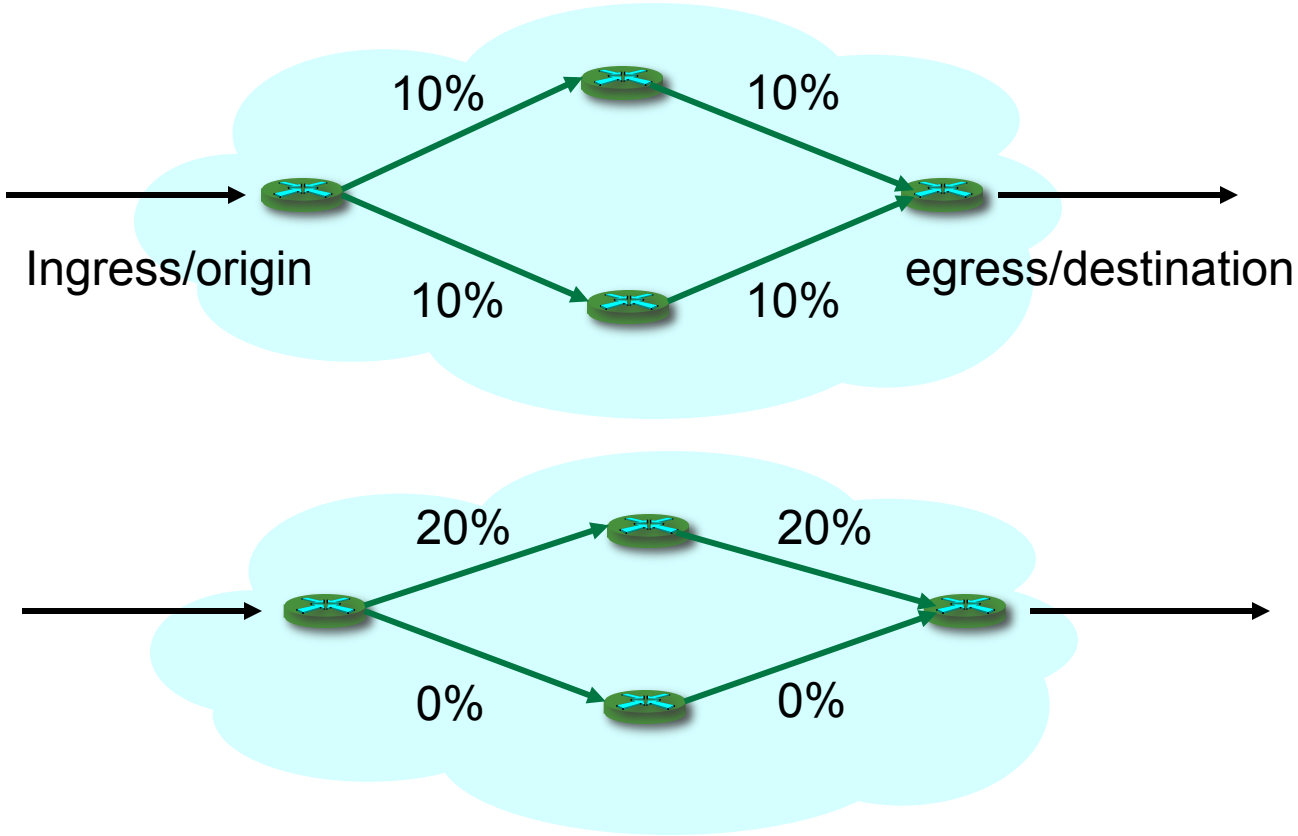

# Research Questions

- Which links to turn on/off and how much traffic each link should carry?
  - An optimization problem.
- How to maintain performance at the desired level?
  - Link utilization and delay
  - Network reliability
- How to realize the traffic distribution?
  - Traffic engineering mechanisms

# The GreenTE Optimization Problem

- Given
  - Topology (links and bandwidth)
  - Power consumption profile of line-cards
  - Traffic matrix

# The GreenTE Optimization Problem

- Given
  - Topology (links and bandwidth)
  - Power consumption profile of line-cards
  - Traffic matrix
- Find a routing solution
  - Paths to be used
  - Traffic split ratio over multiple paths

# The GreenTE Optimization Problem

- Given
  - Topology (links and bandwidth)
  - Power consumption profile of line-cards
  - Traffic matrix
- Find a routing solution
  - Paths to be used
  - Traffic split ratio over multiple paths
- Maximize total energy saving from turning off line-cards

# The GreenTE Optimization Problem

- Given
  - Topology (links and bandwidth)
  - Power consumption profile of line-cards
  - Traffic matrix
- Find a routing solution
  - Paths to be used
  - Traffic split ratio over multiple paths
- Maximize total energy saving from turning off line-cards
- Such that
  - Flow conservation holds
  - Max link utilization (MLU) below a threshold (50%)
  - Delay is bounded (same network diameter or 2x OD delay)

# A heuristic solution

- Search in a set of pre-computed candidate paths instead of all possible paths.
  - k-shortest paths
  - Apply delay constraints in choosing the candidate paths.

- Take another step to balance link load after knowing which links will be on.

# Protocol and operation issues

- A server collects input, solves the problem, and distributes the solution.
  - Need to be done periodically to adapt to traffic changes.
- Piggyback on OSPF-TE
  - TE-LSA for input, TE-Metric for result.
  - Take advantage of LSA flooding
- Forwarding is done via both OSPF and MPLS
  - Make changes by changing MPLS tunnels rather than network-wide OSPF convergence.
- Need to handle periodic control messages such as OSPF Hello.

# Evaluation

- Using data from real networks
  - Real topology and traffic traces for Abilene and GEANT.
  - Measured topology and generated traffic for Sprint and AT&T

| Network | Usage | Location | Nodes | Links |
|---|---|---|---|---|
| Abilene | Research | US | 12 | 30 |
| GÉANT | Research | Europe | 23 | 74 |
| Sprint | Commercial | US | 52 | 168 |
| AT&T | Commercial | US | 115 | 296 |

# Power saving

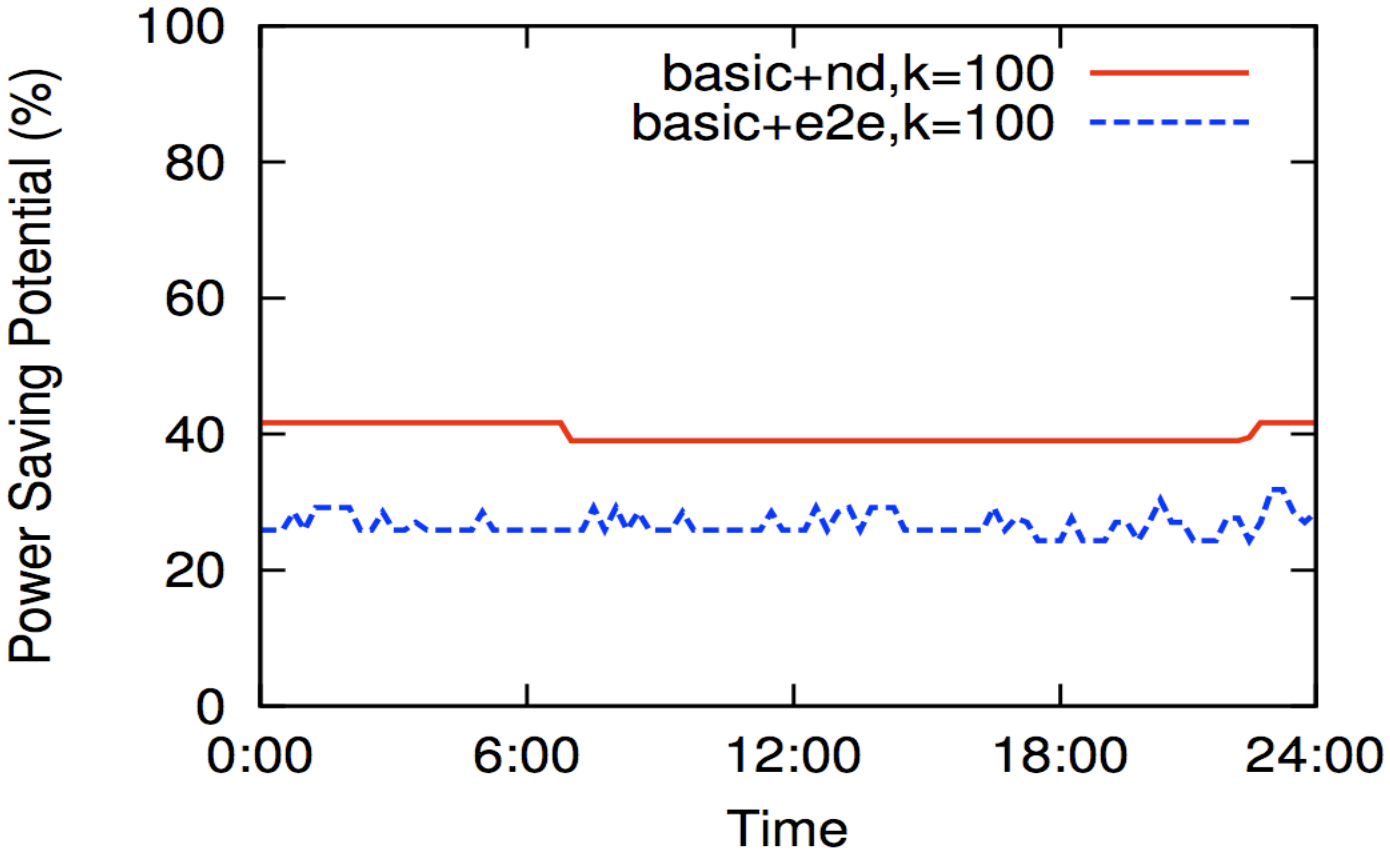


- 27% to 42% of line-card power consumption

# Impact of network load

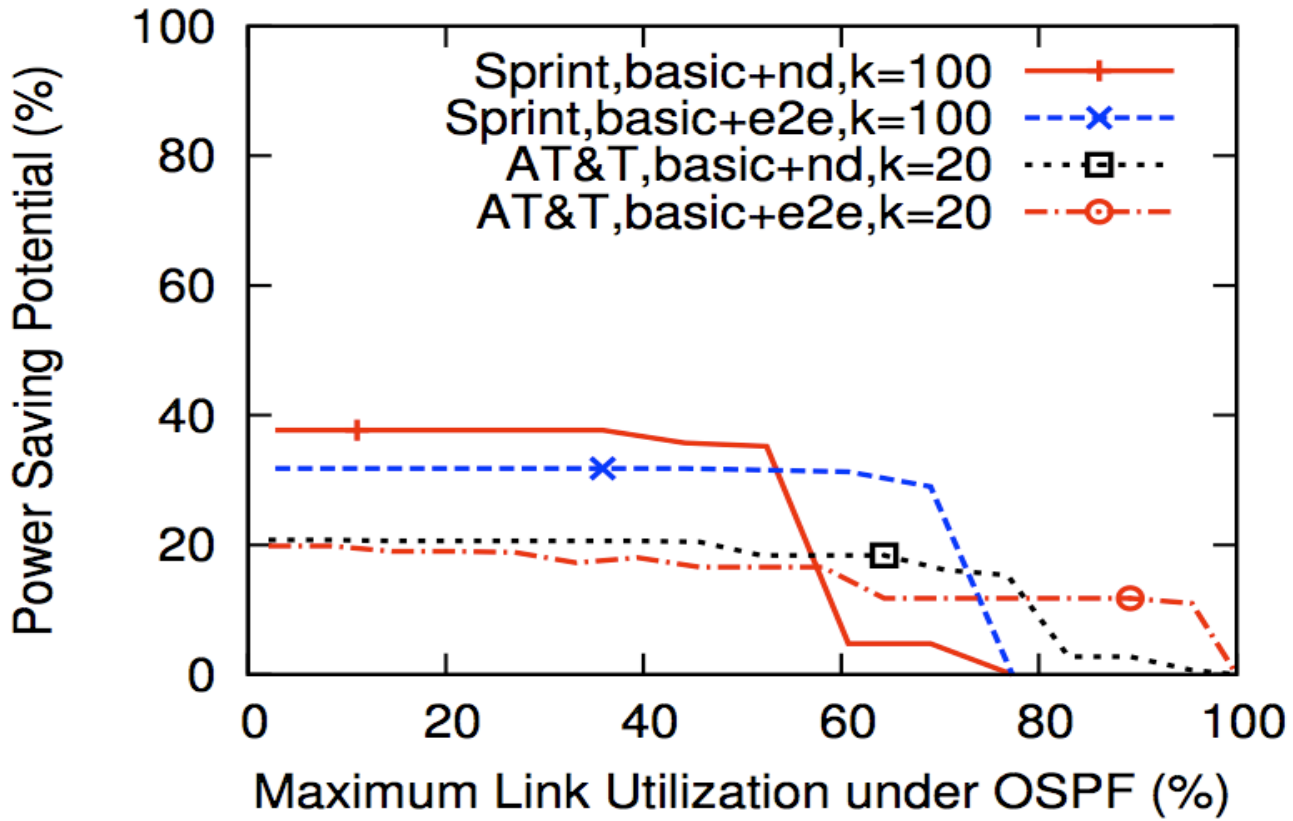


- Links are tuned off under light load, back on under heavy load.

# Load balancing

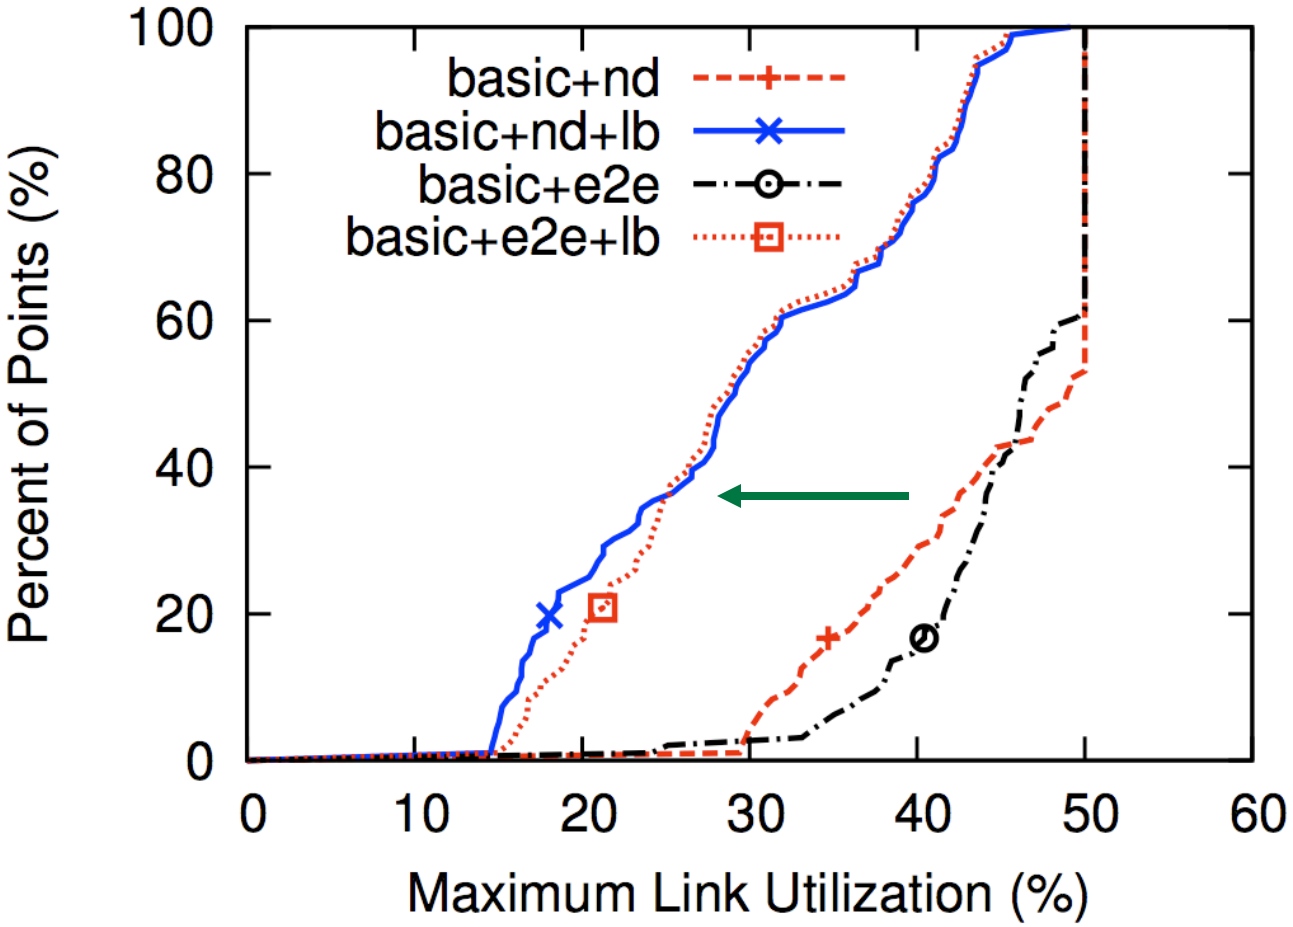


- Significant improvement after load balancing.

# Delay

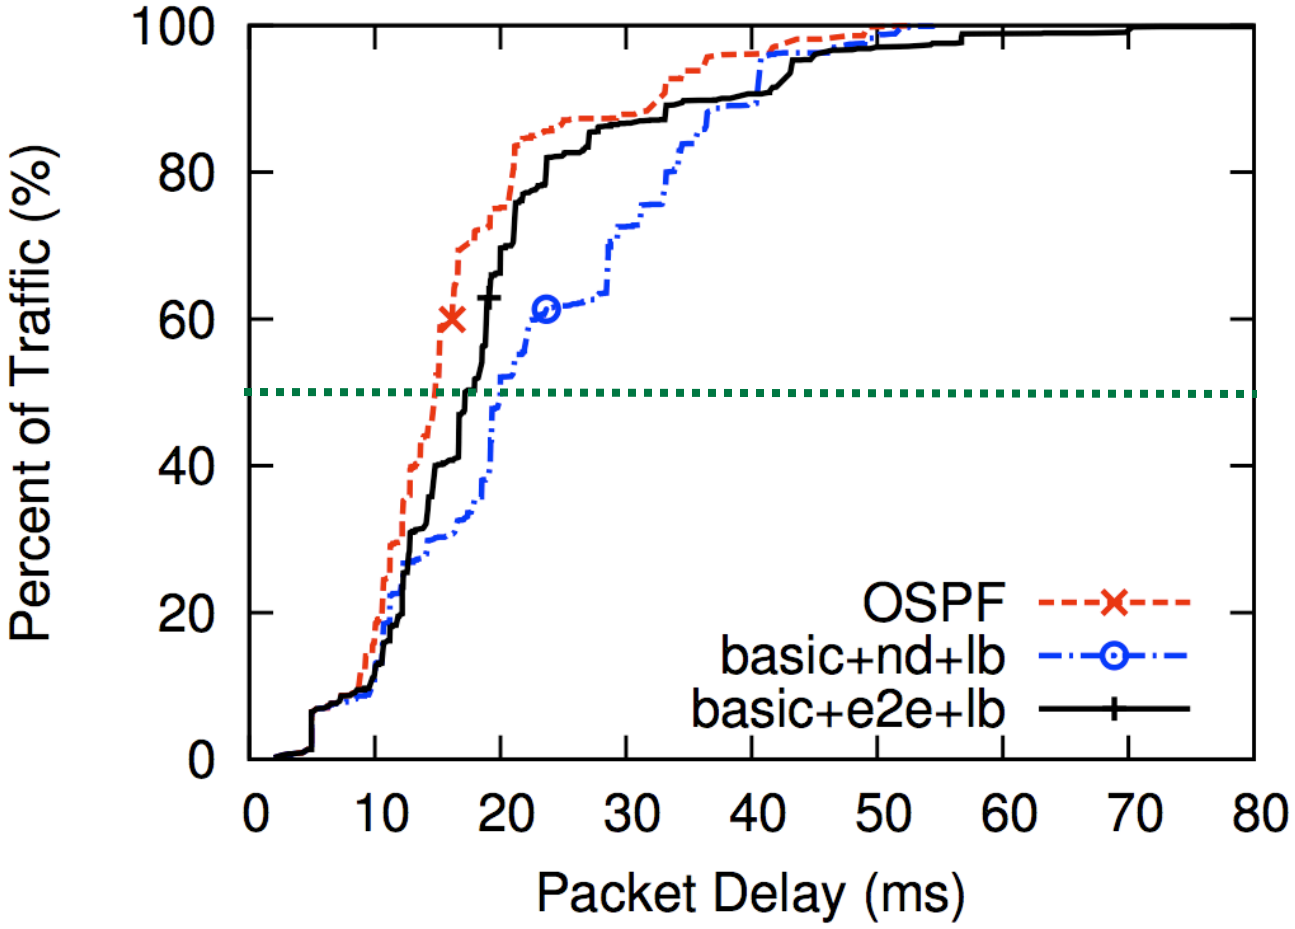


- Medium: 15ms, 17ms, 20ms

# Stability of the solution

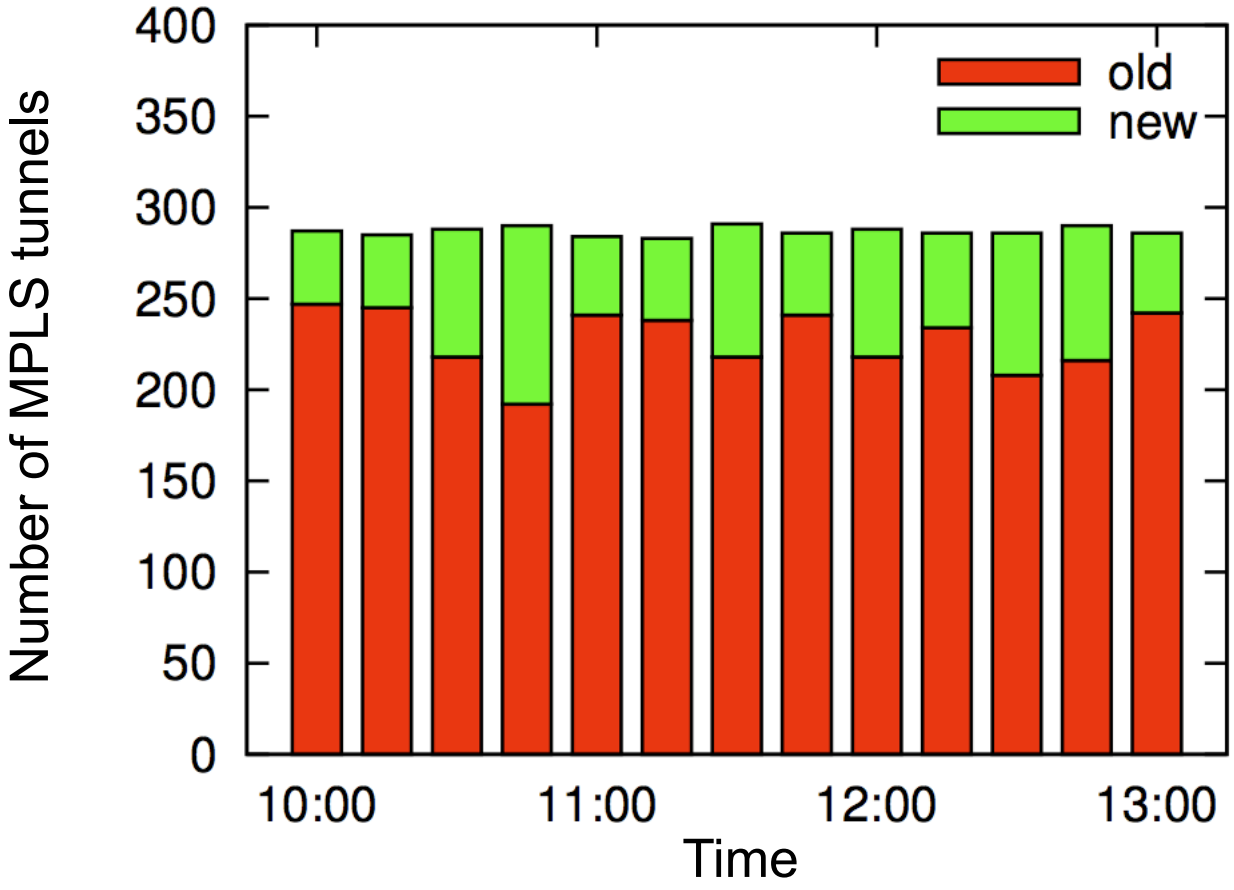


~70% tunnels remain the same from the previous period.

# Conclusion and Future Work

- Power-aware routing and traffic engineering is both feasible and beneficial.
- Need better protocol support and hardware/ system support.
  - E.g., keep-alive messages.
- Need better handling of potential congestion induced by link failures and traffic bursts.
- Distributed solutions?

*Questions and Comments?*