Responsive, Energy-Proportional Computer Networks

Nedeljko Vasić, Dejan Novaković, Satyam Shekhar, Prateek Bhurat, Marco Canini, and Dejan Kostić

Networked Systems Laboratory

EPFL, Switzerland

Research sponsored by FNSNF Swiss National Science Foundation

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Networking Energy Consumption

- Internet's energy consumption is already large
  - US network infrastructure requires ~20 TWh/year
  - Italy's ISP Telecom Italia needs ~2 TWh/year
  - Datacenter's networking is 20% of server energy
- More demands will result in further increases
  - Video streaming, Video-on-Demand, Cloud computing
- CMOS reaching a plateau in power-efficiency
  - Cooling costs of new equipment will increase
    - 1 MW for latest Cisco platform, CRS-1

Network redundancy/variability in traffic

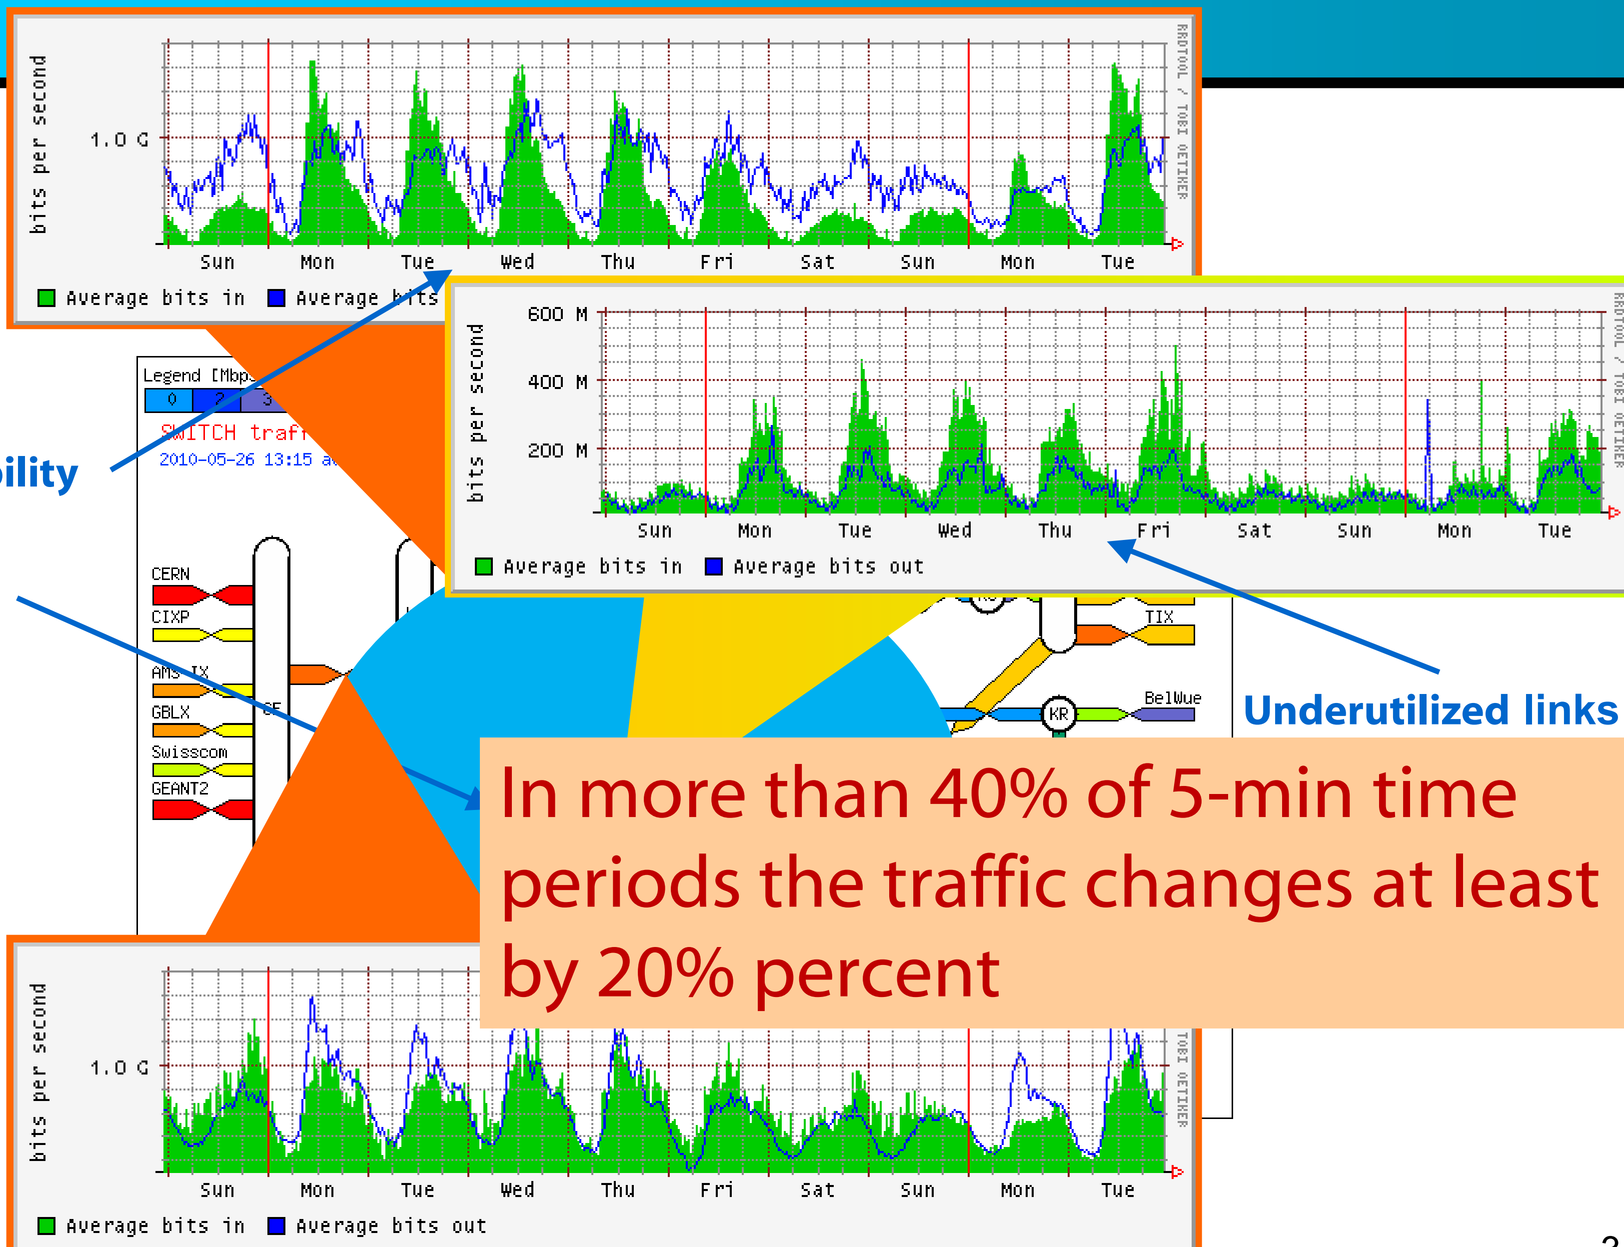

Traffic variability

Redundancy

Underutilized links

In more than 40% of 5-min time periods the traffic changes at least by 20% percent

Datacenter Networks

Large degree of redundancy

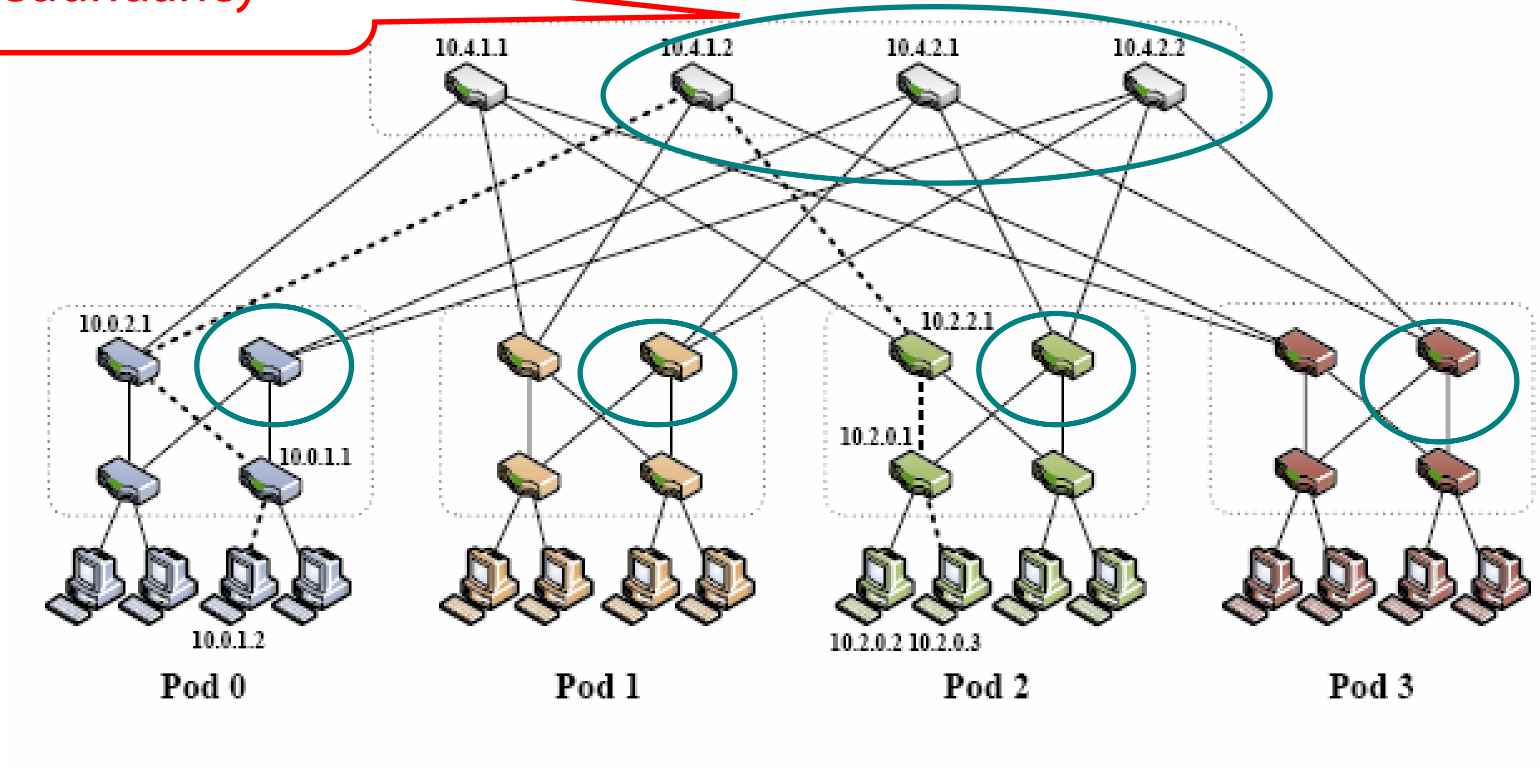

[Al-Fares et al., SIGCOMM '08]

Network Energy-(un)proportionality

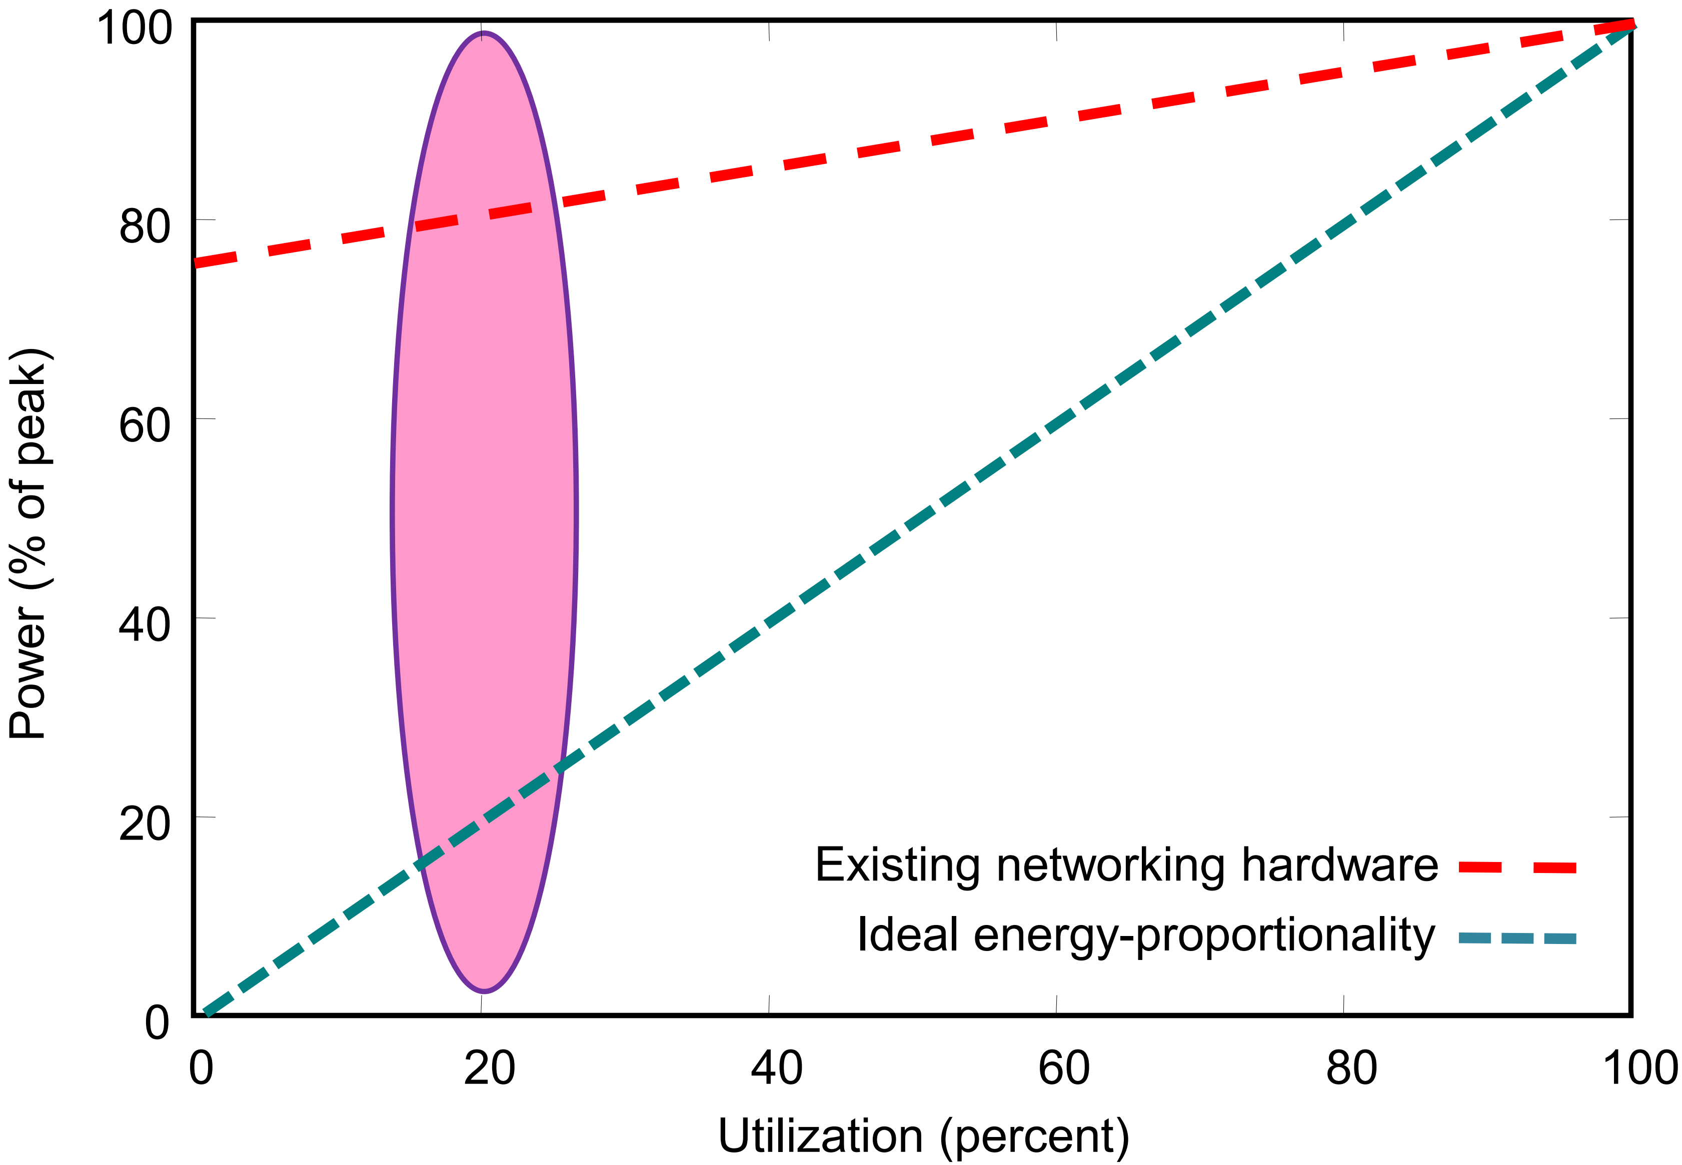

Goal: Energy-Proportional Networks

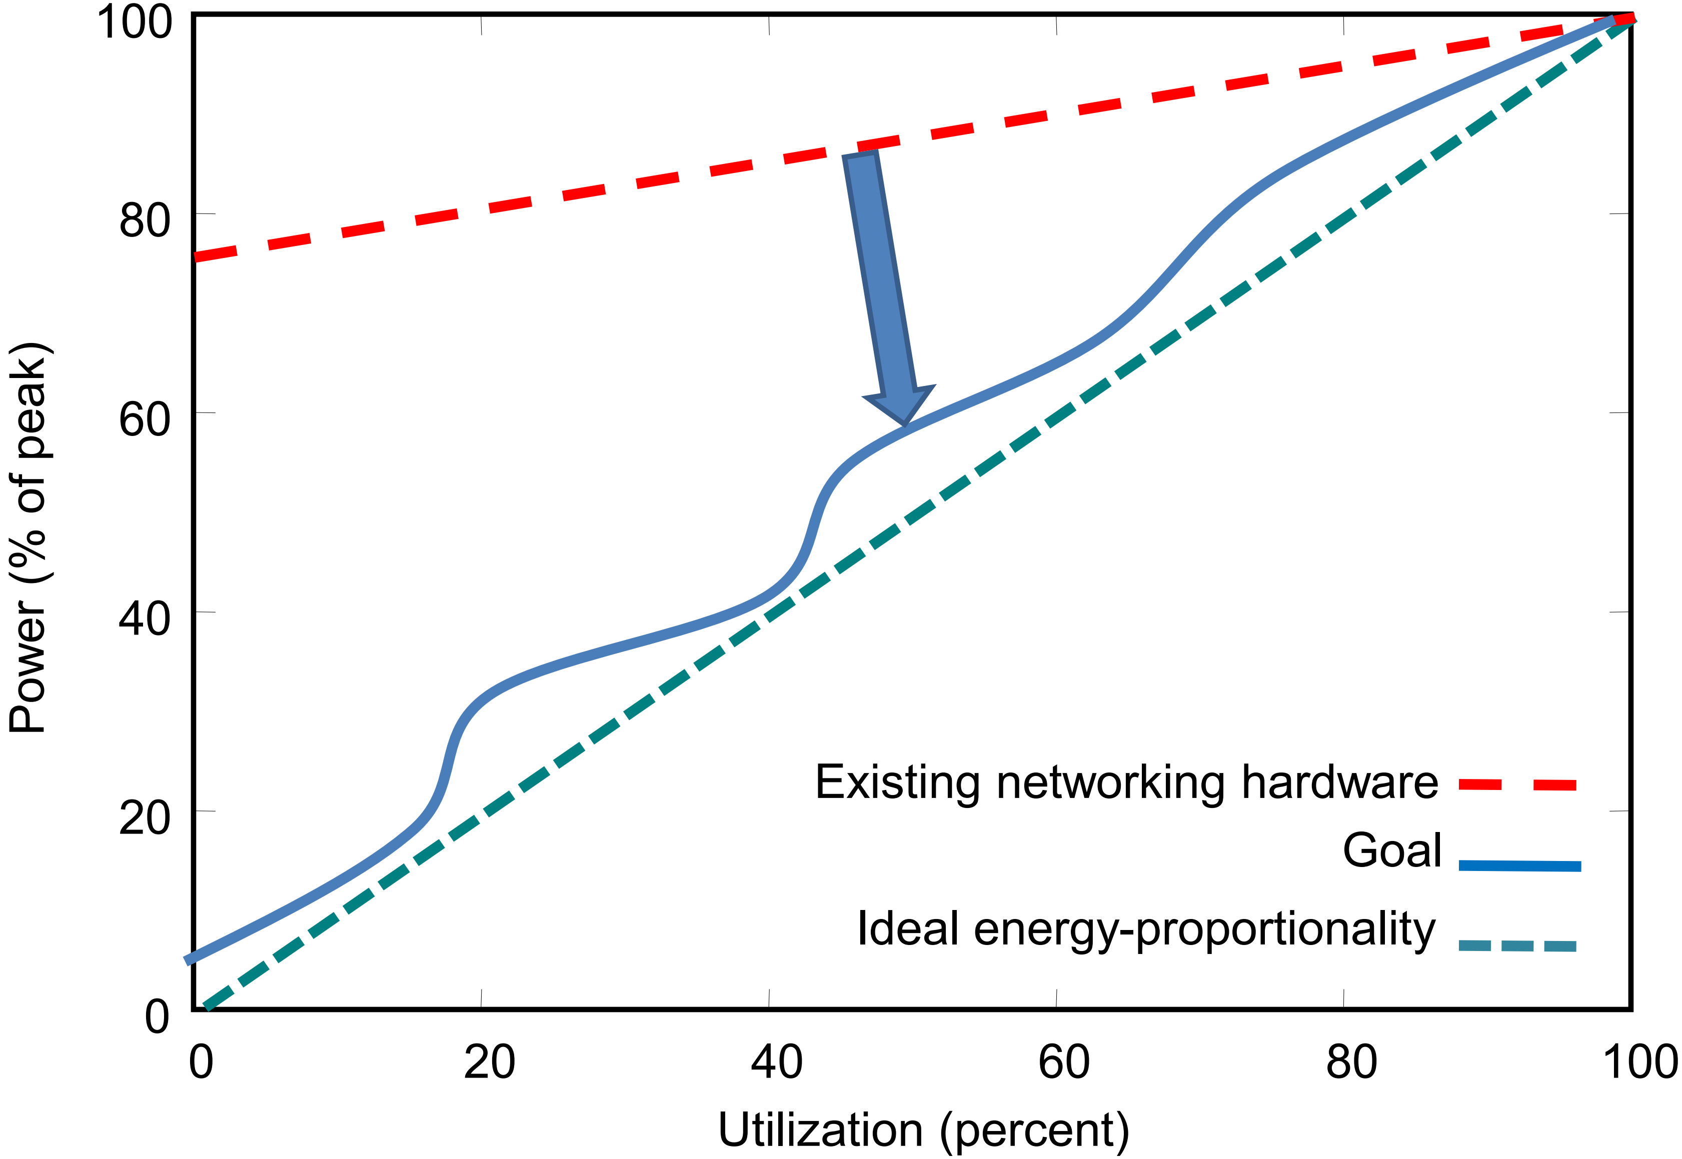

Possible Approach

- Make individual components energy-proportional
  - Implementation and deployment challenges
    - Limits of energy efficiency in CMOS
    - Leakage current
    - Always-on components, …

Our Approach (Ensemble)

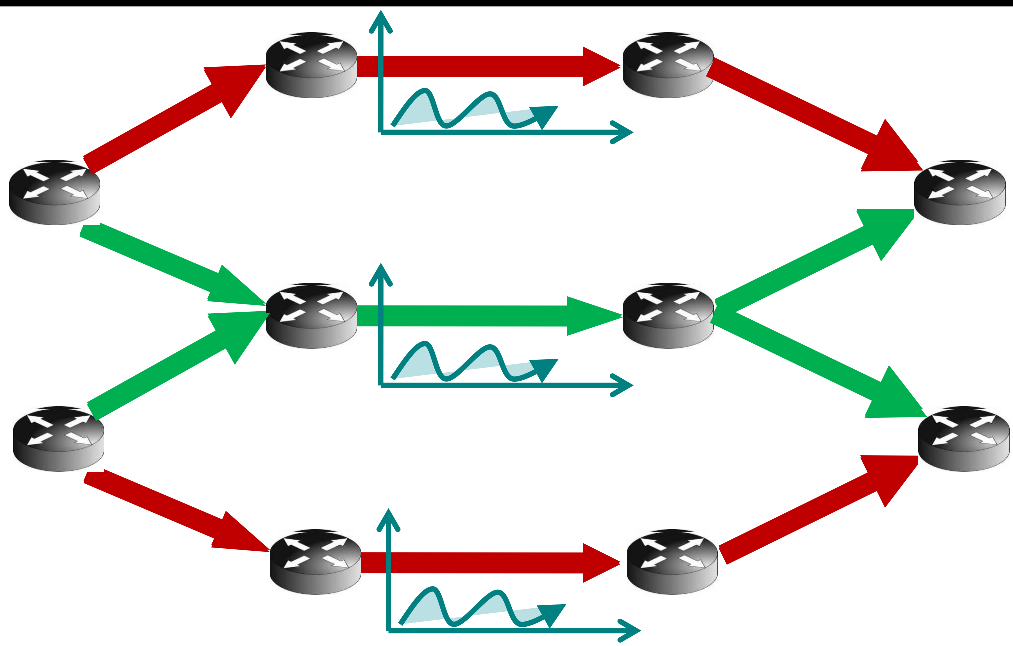

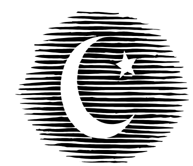

- Dynamically match resources to the load
  → make the **ensemble** energy-proportional

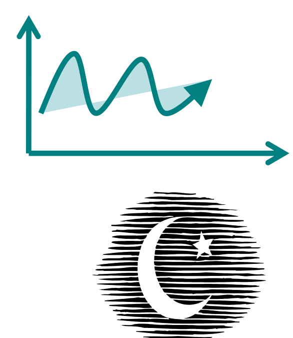

Ensemble Approach Challenges

- Producing significant energy savings
- Meeting the SLOs
- Avoiding oscillations
- Ease of deployment
- Responsiveness to traffic variations

Routing table computation

- Routing that minimizes power consumption
  - Multi-commodity flow problem, but with additional constraints for power objective:

    Links + routers (switches) on/off

  - Problem is NP-complete

When traffic demand changes, optimal routing changes!

Related Work

- [Gupta *et al.*, SIGCOMM '03]
  – Greening of the Internet vision
- [Nedevschi *et al.*, NSDI '08]
  – Local actions
- [Chabarek *et al.*, INFOCOM '08]
  – Power-aware network *provisioning*
- [Chiaraviglio *et al.*, GreenCom '09], ElasticTree [Heller *et al.*, NSDI '10], GreenTE [Zhang *et al.*, ICNP '10]
  – Online techniques

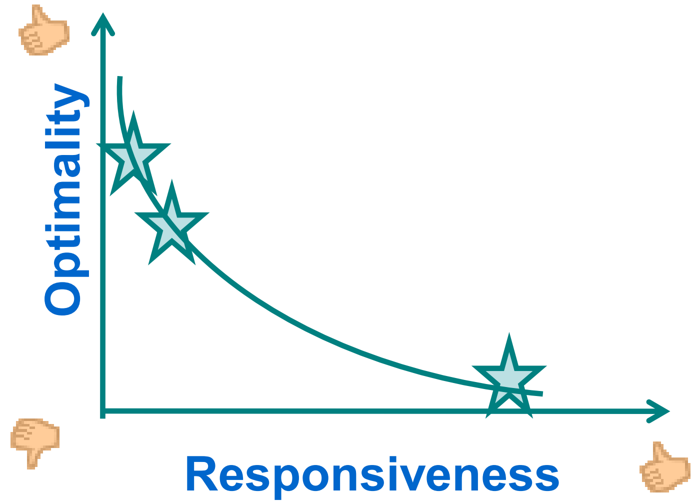

How Often is Recomputation Needed?

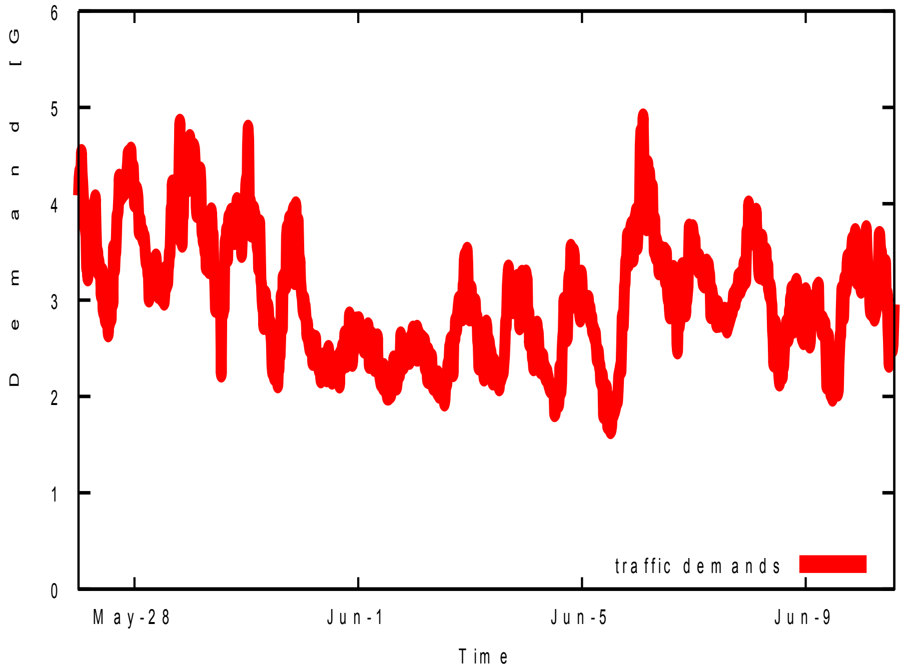

Geant2 - European academic network

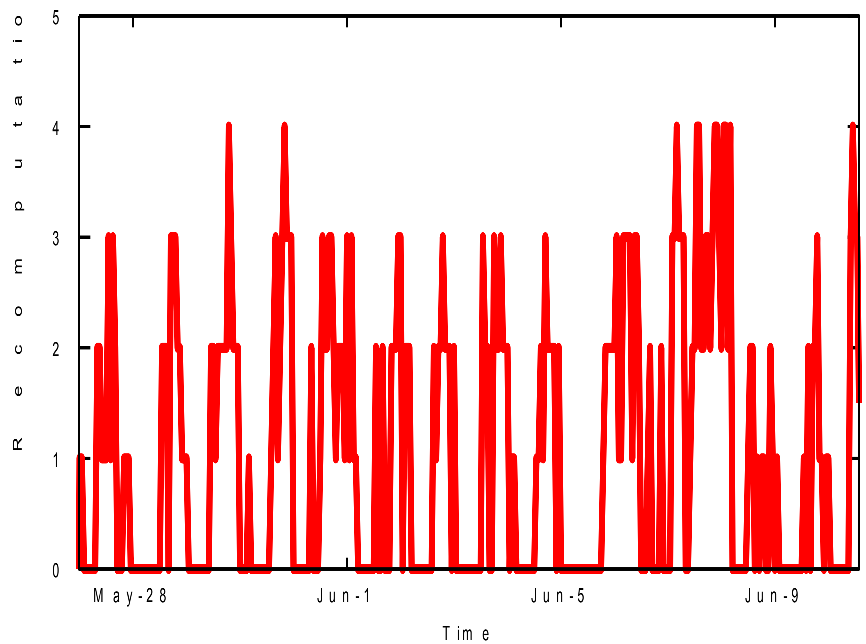

Routing table recomputed
3-4 times per hour!

Recomputation Wastes Energy or Causes Congestion

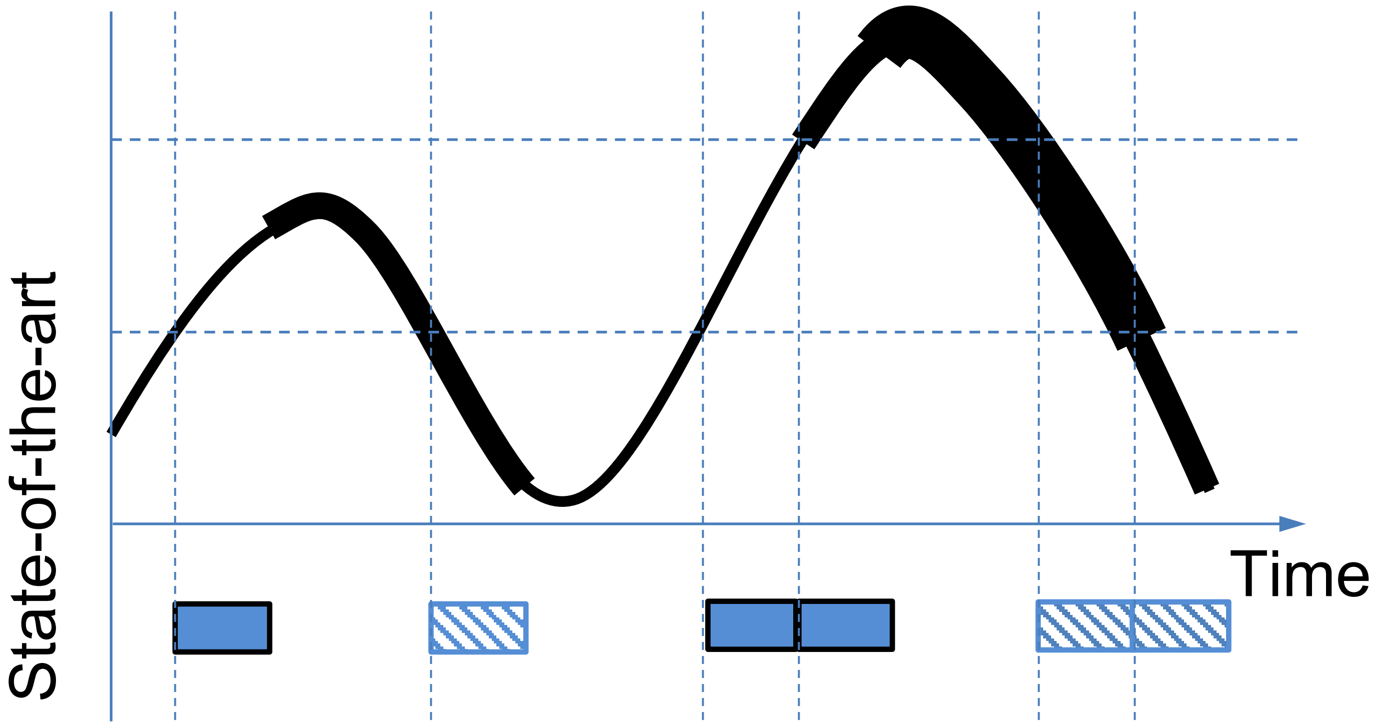

Recomputation causing congestion

Recomputation causing energy waste

REsPoNse

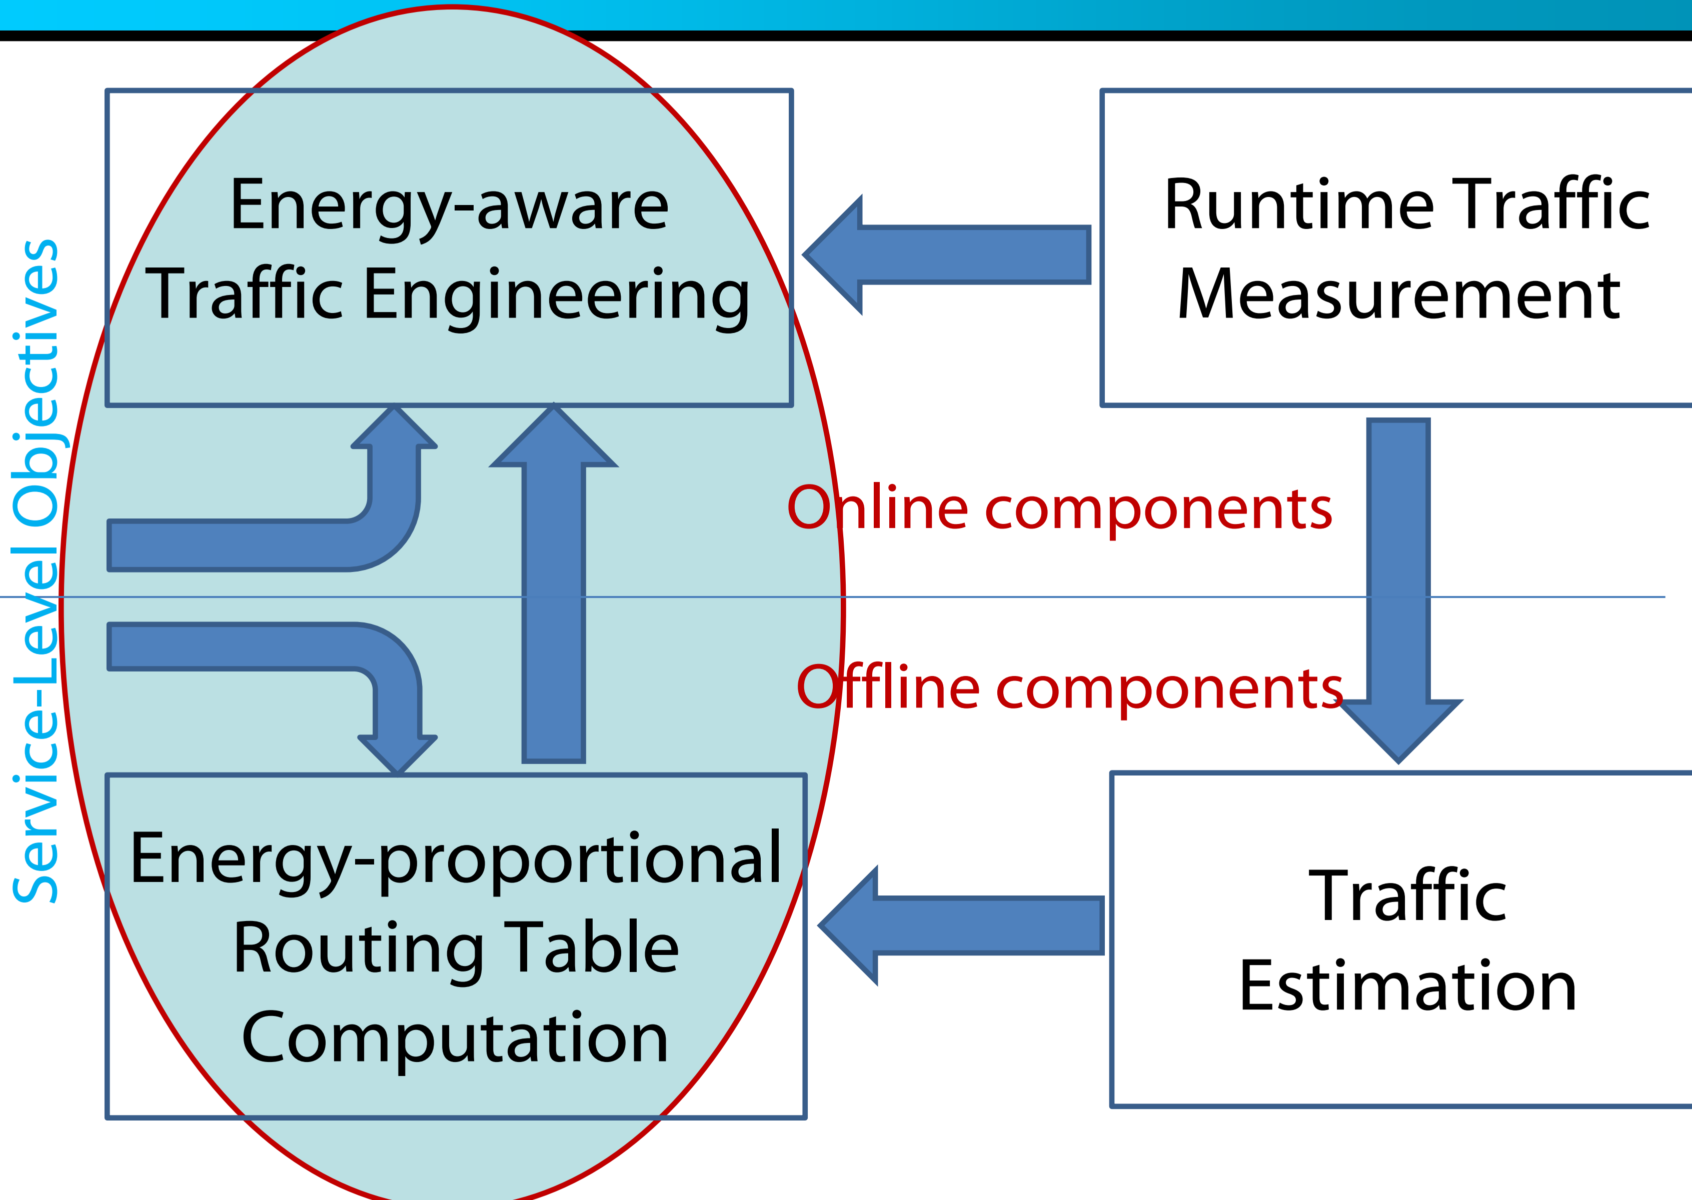
Service-Level Objectives
Energy-aware Traffic Engineering
Runtime Traffic Measurement
Online components
Offline components
Energy-proportional Routing Table Computation
Traffic Estimation

REsPoNse in Action

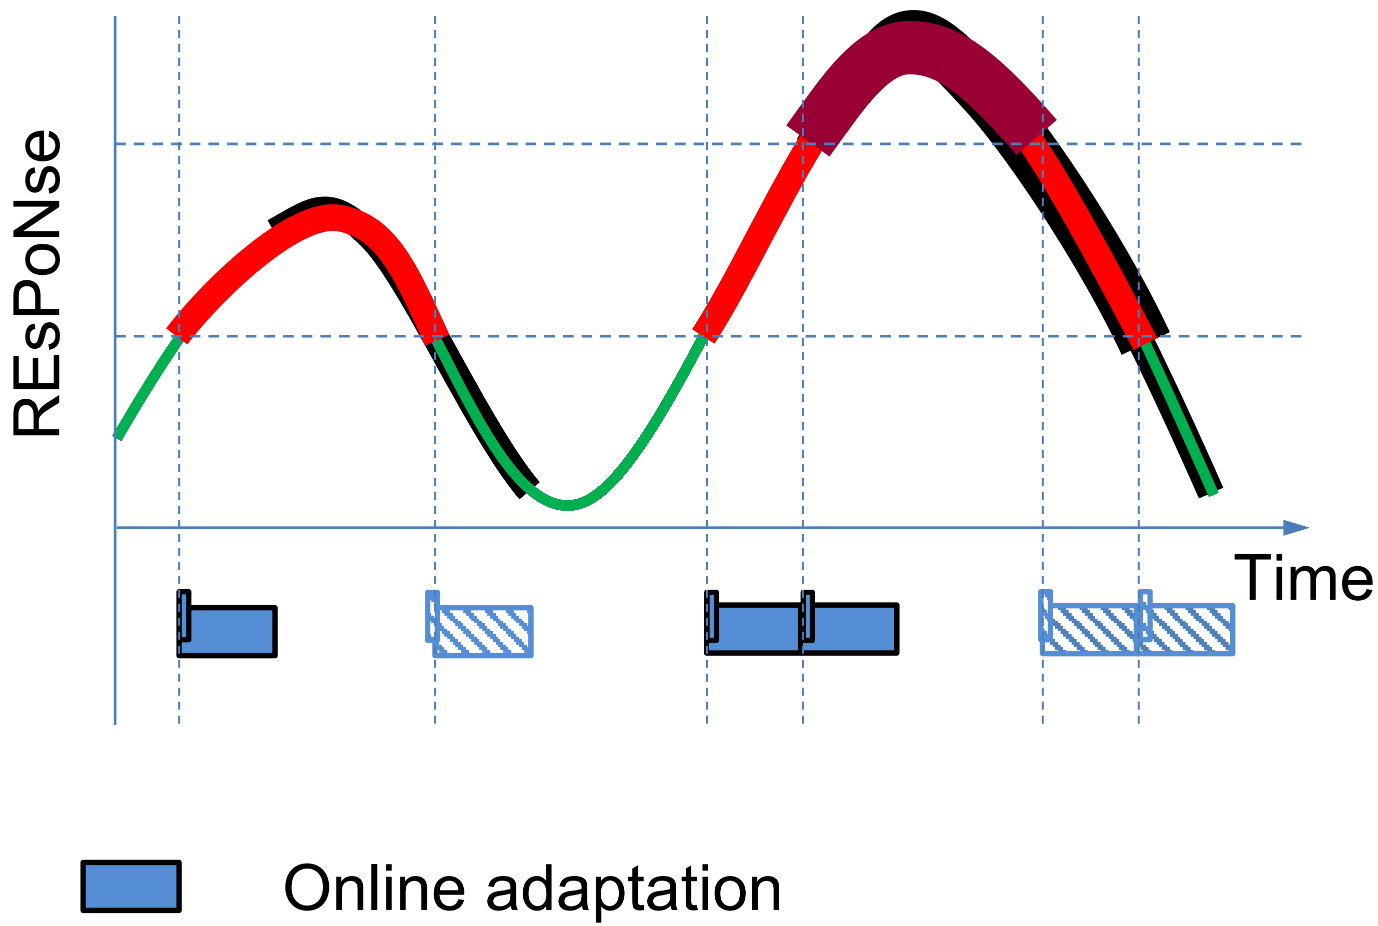

Outline

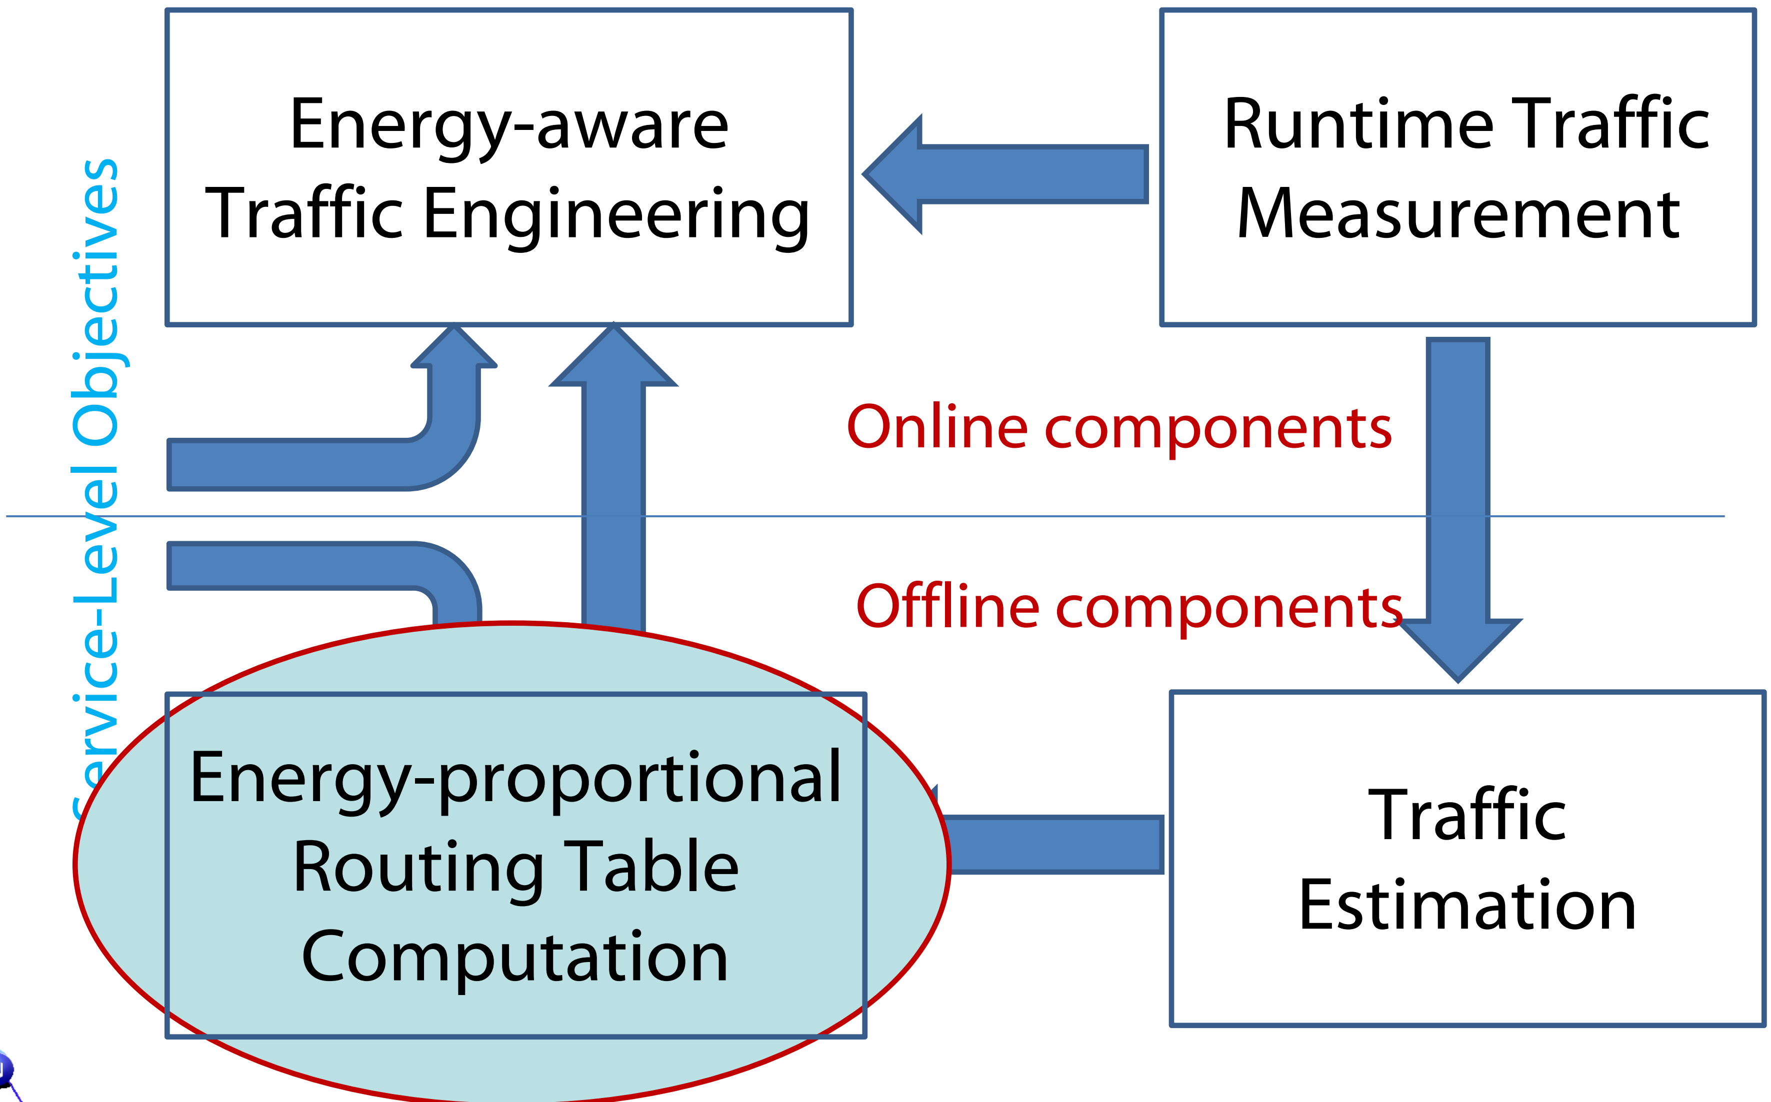
Energy-aware
Traffic Engineering
Runtime Traffic
Measurement
Service-Level Objectives
Online components
Offline components
Energy-proportional
Routing Table
Computation
Traffic
Estimation

Energy-proportional routing paths

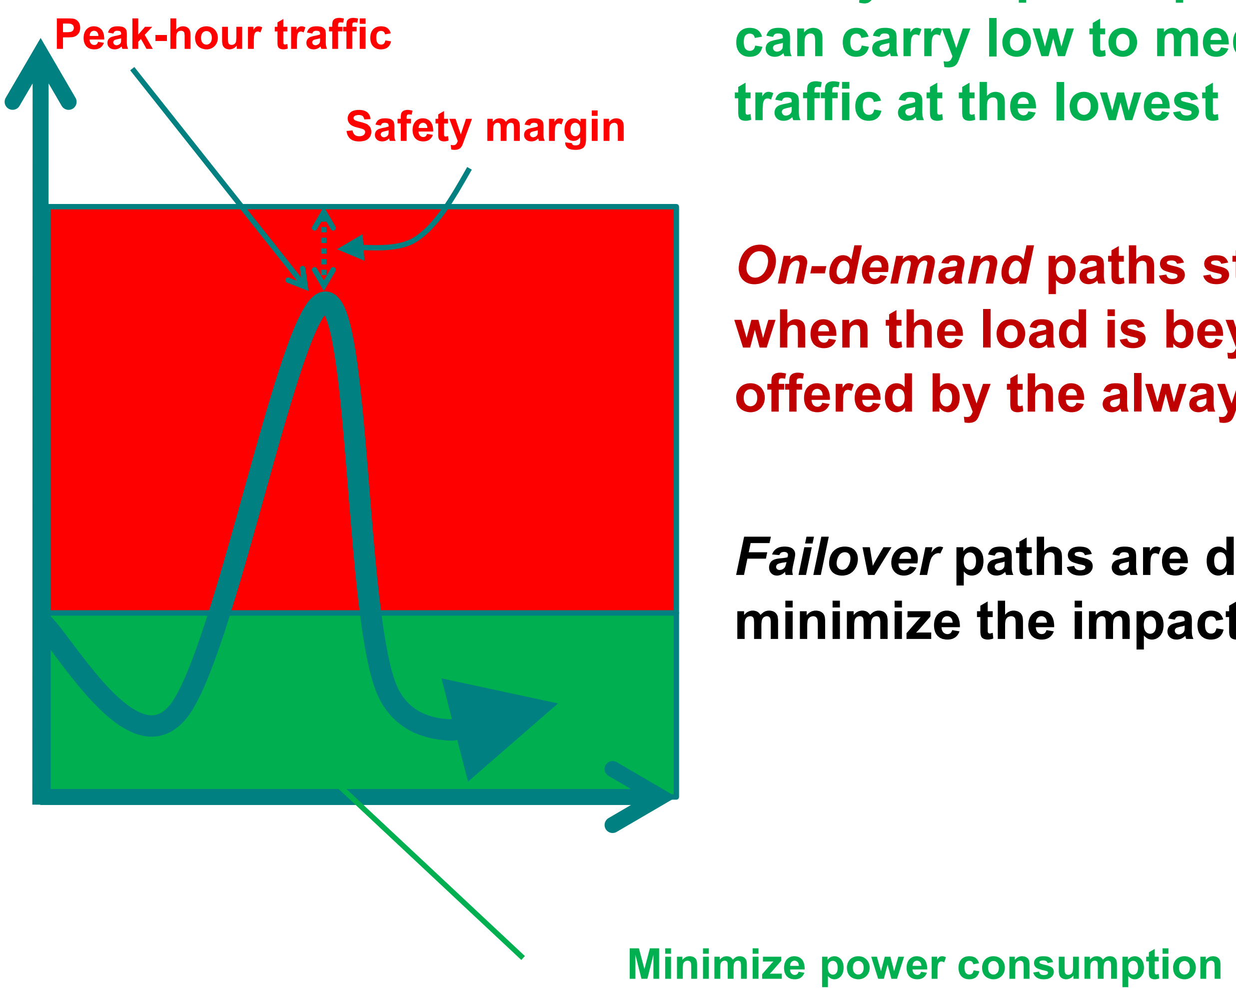

Always-on paths provide a routing that can carry low to medium amounts of traffic at the lowest power consumption

On-demand paths start carrying traffic when the load is beyond the capacity offered by the always-on paths

Failover paths are designed to minimize the impact of single failures

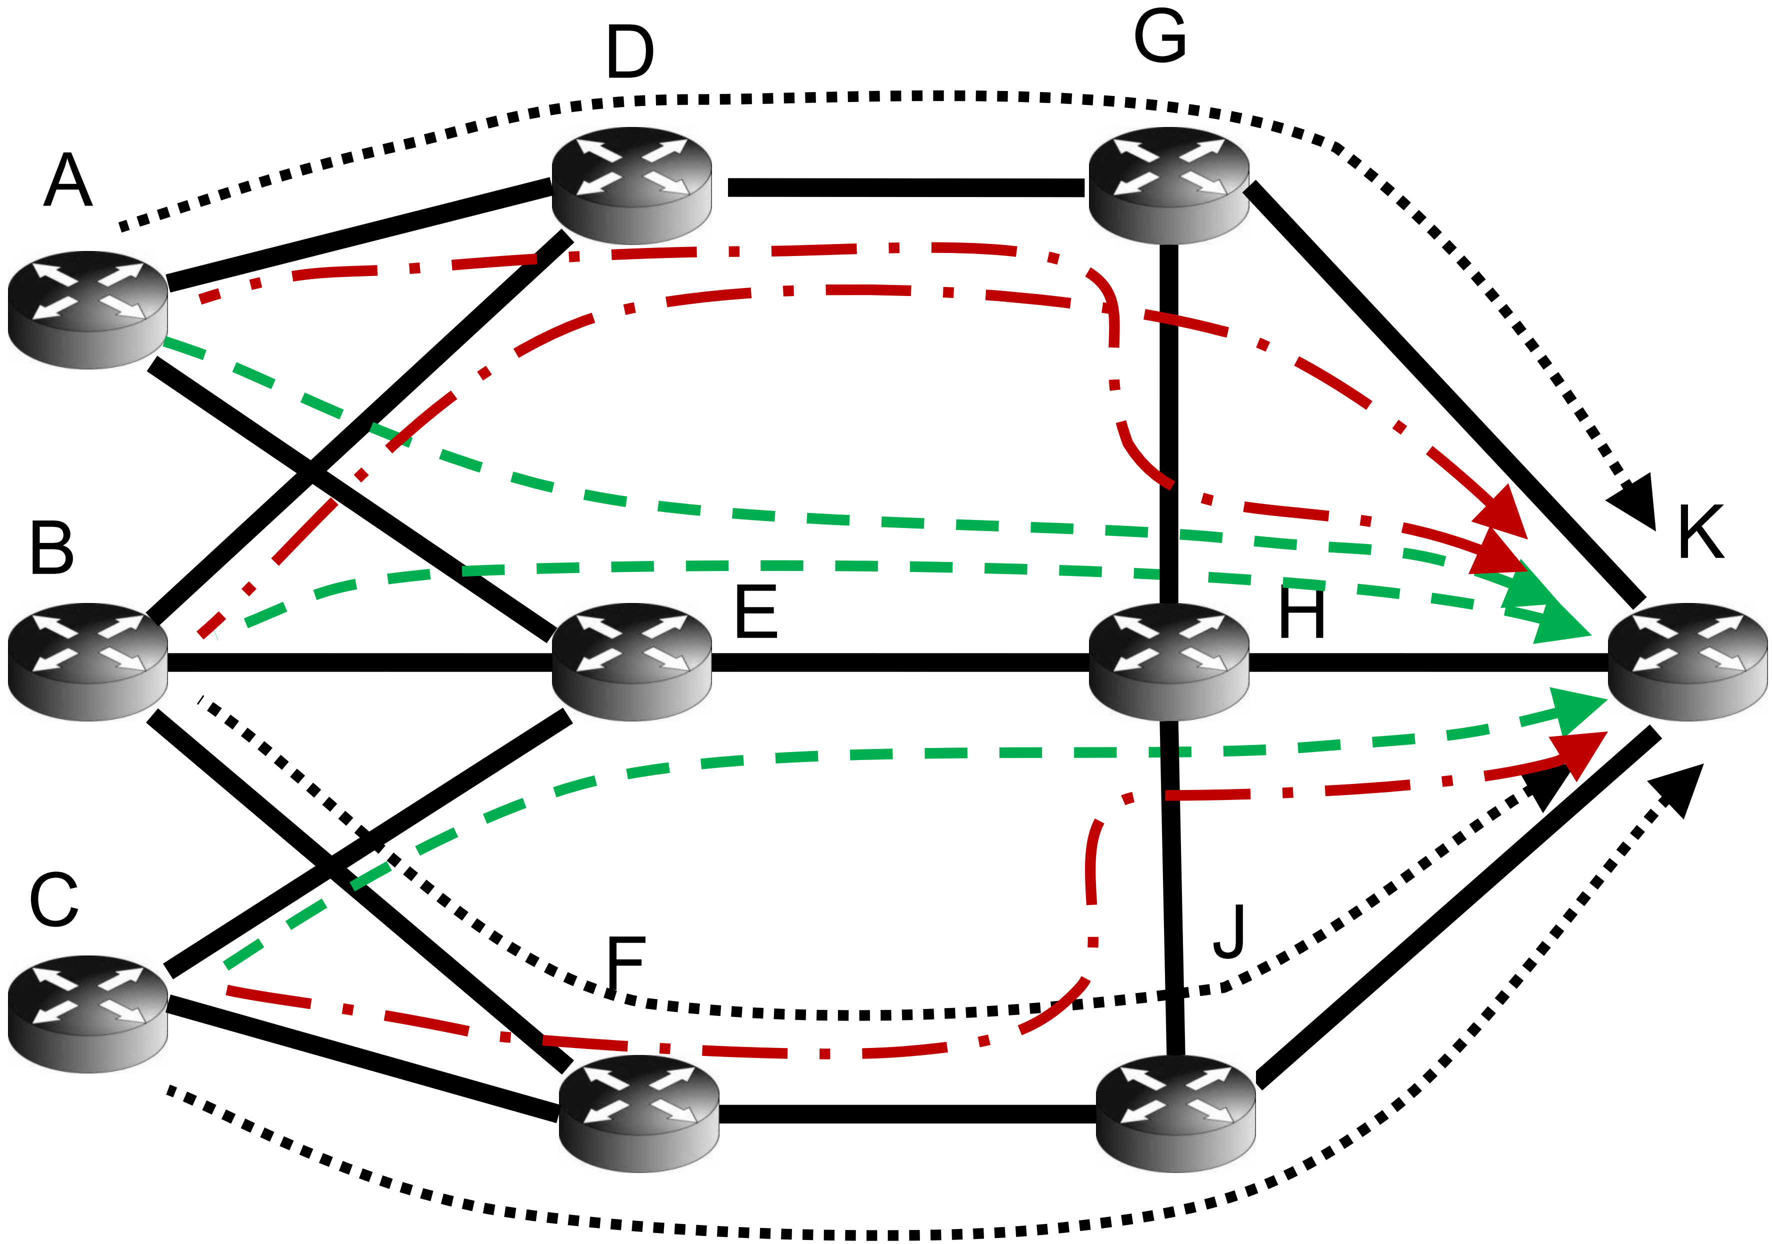

A
B
C
D
E
F
G
H
J
K

Computing REsPoNse paths

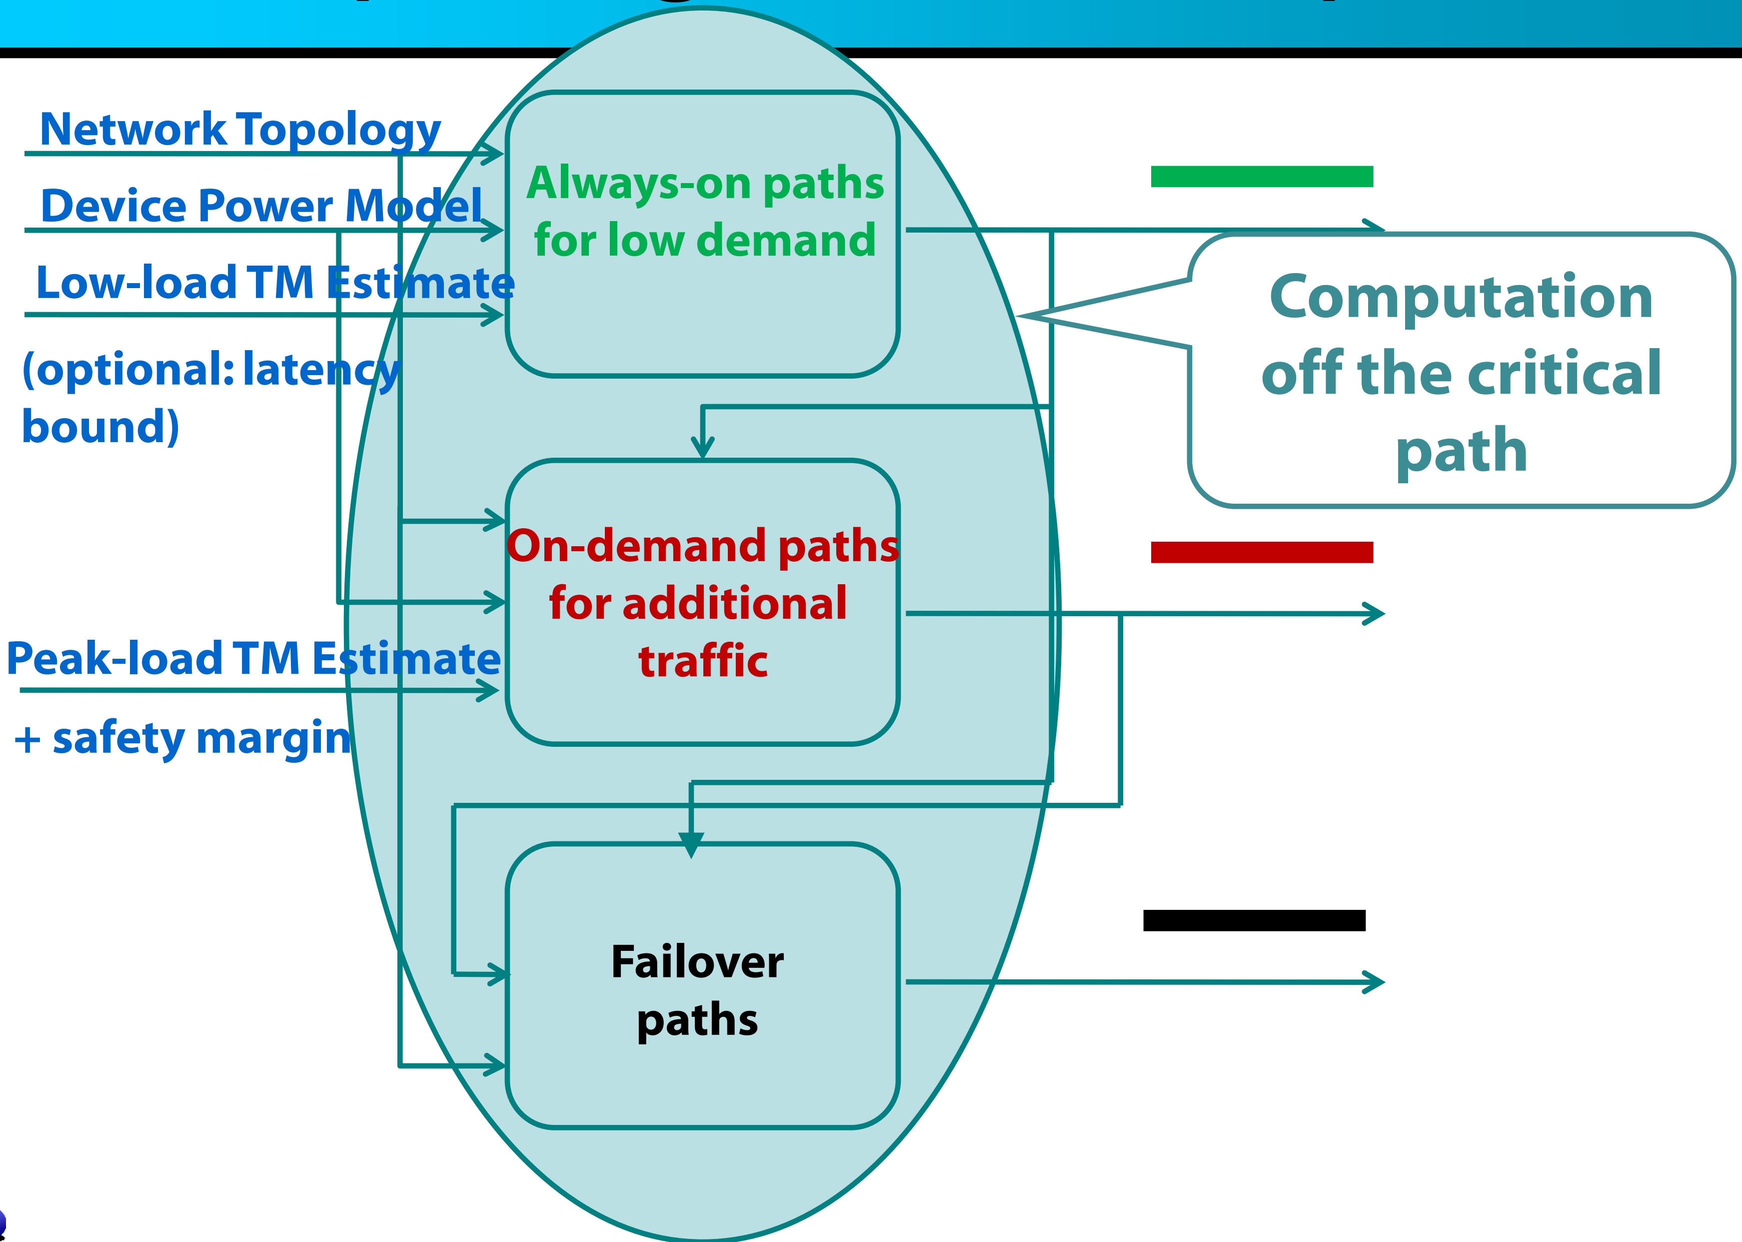

Outline

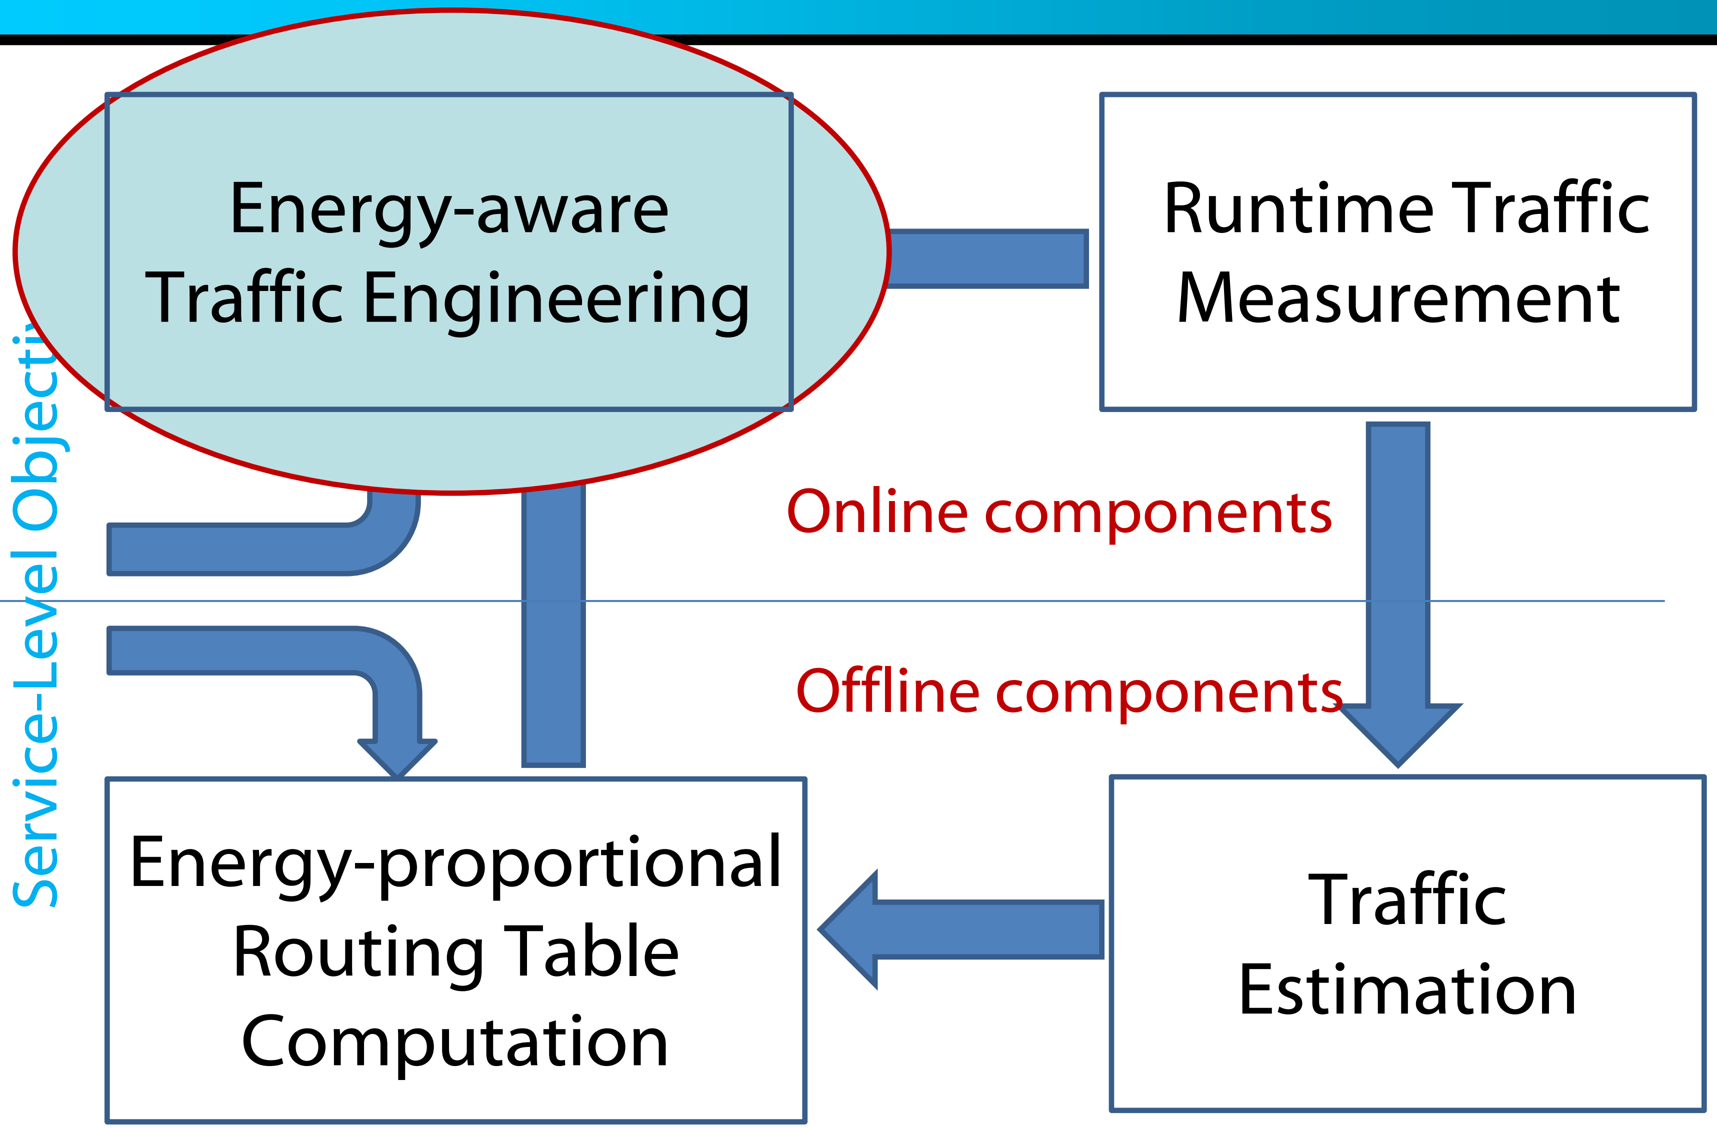
Energy-aware Traffic Engineering
Runtime Traffic Measurement
Online components
Offline components
Service-Level Objecti
Energy-proportional Routing Table Computation
Traffic Estimation

EATe [e-Energy '10]

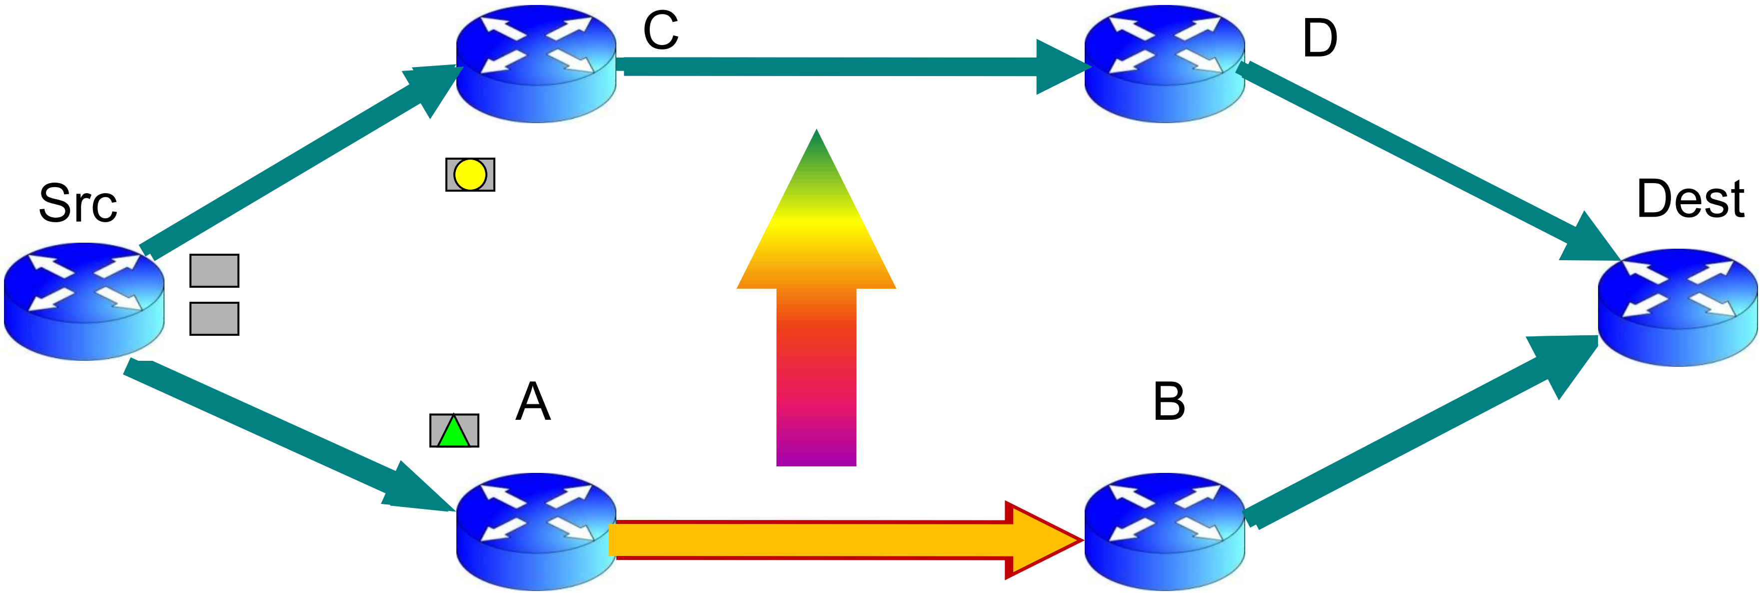

- Online effort to shift traffic to inactivate on-demand paths
  - Intermediate routers mark packets with link load
  - Edge routers collect load info only on alternative paths
    - Scalable

EATe Stability

Stable, converges in a few RTTs

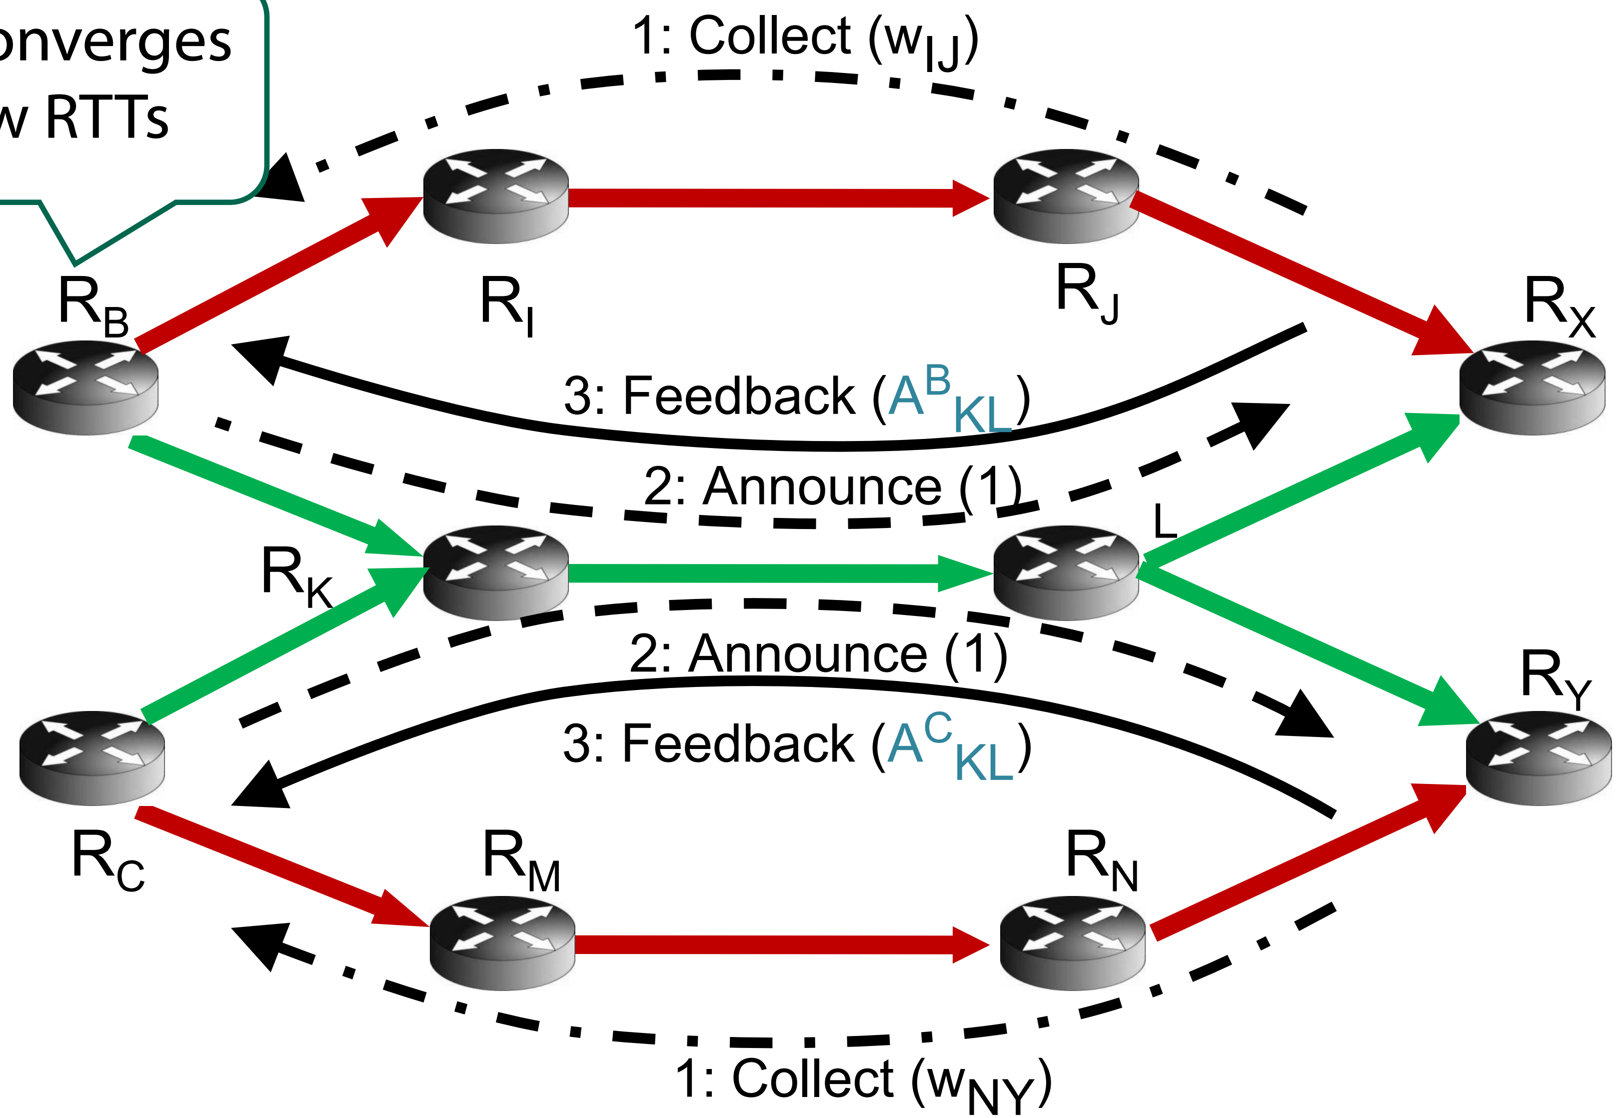

Evaluation Questions

- How energy-proportional is REsPoNse/EATe?
  - ISP topologies
  - Datacenter networks
- How quick is REsPoNse/EATe in shifting traffic?
- What is the impact of traffic aggregation on application performance?

Energy-Proportionality (ISP topology)

Abovenet US ISP topology

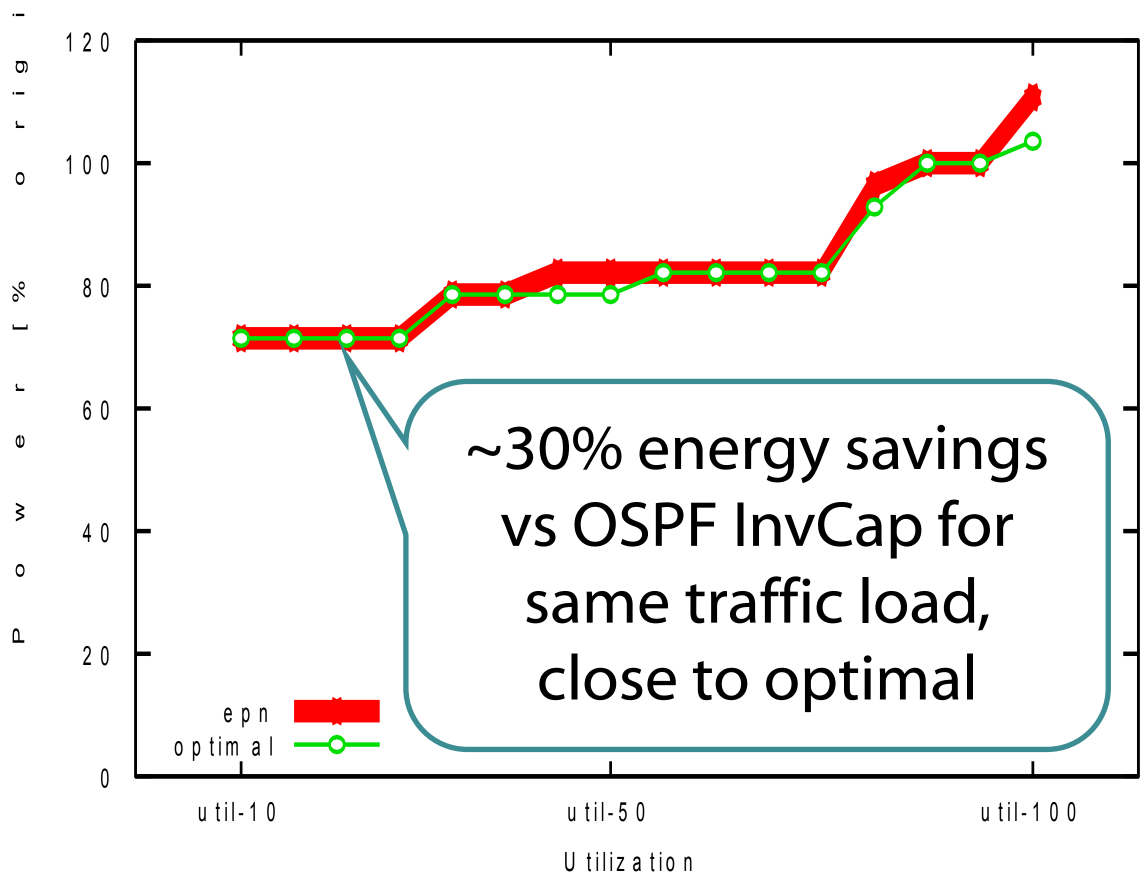

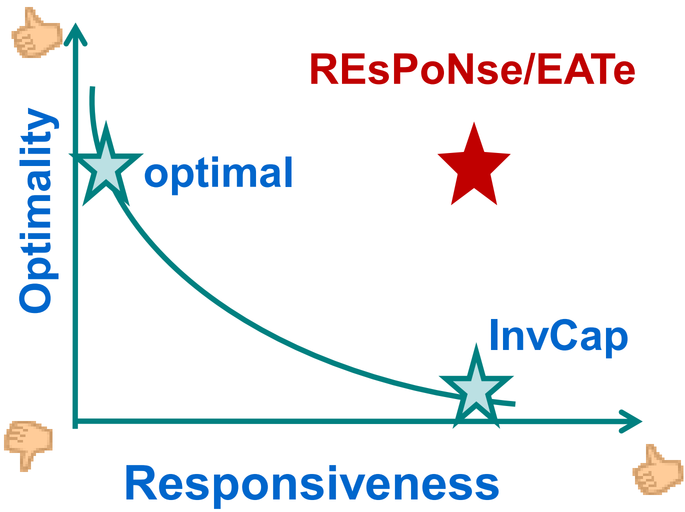

Responsiveness/stability (live)

10 Click routers in a diamond topology
(16 ms per-hop latency)

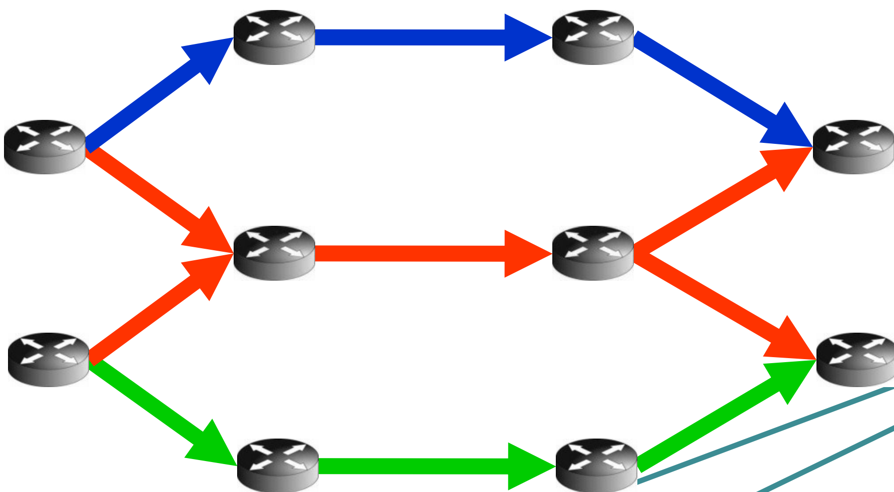

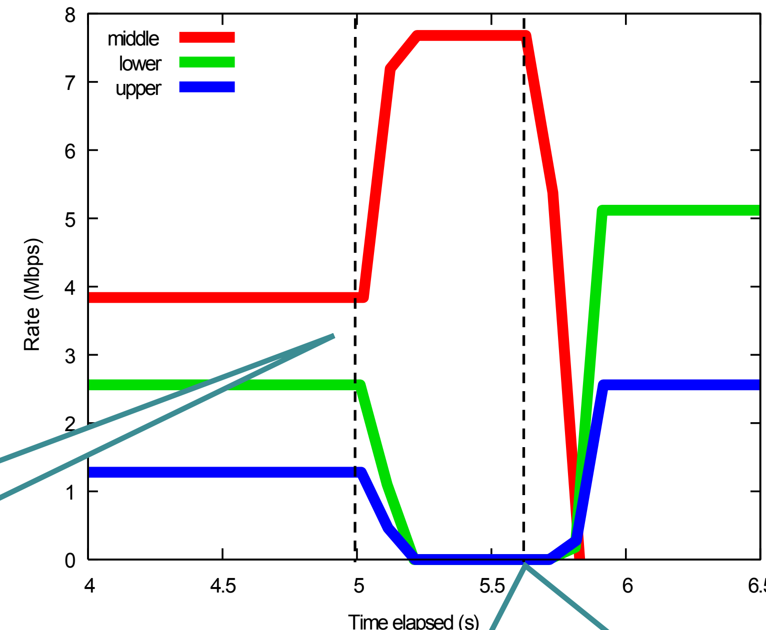

EATe starts running

Link failure

EATe quickly and in a stable manner shifts traffic as needed (either to save energy or to avoid failed links)

Responsiveness/Energy-Proportionality (datacenter)

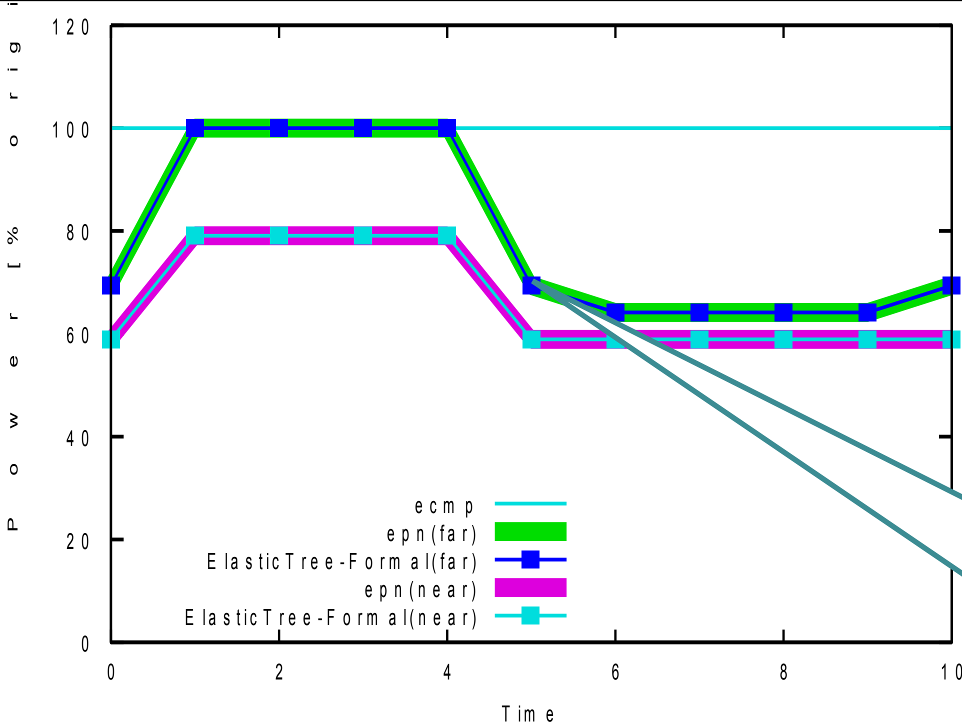

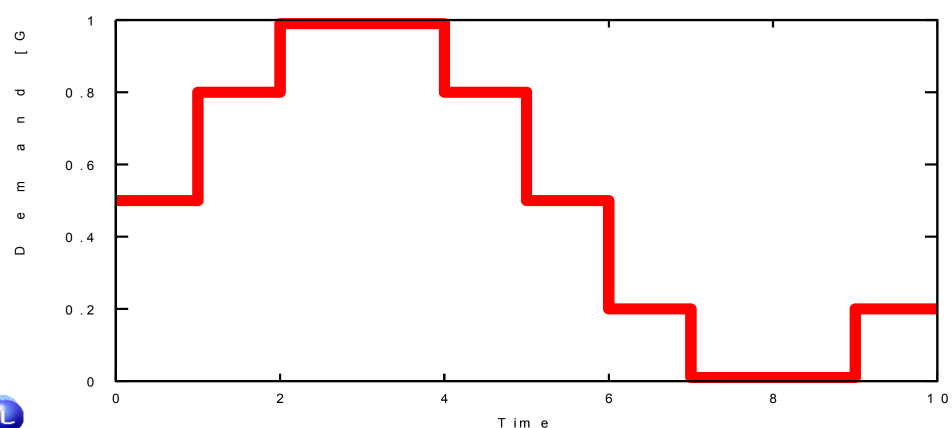

- k=4 fat-tree topology
- Sine-wave demand
  - Each flow [0, 1Gbps]
  - Near (localized)
  - Far (non-localized)

REsPoNse/EATe matches traffic changes, expending the same energy as ElasticTree [Heller *et al.*, NSDI ’10]

Impact on app. performance (live)

Live Modelnet experiment with VoD (BulletMedia [IPTV '07])

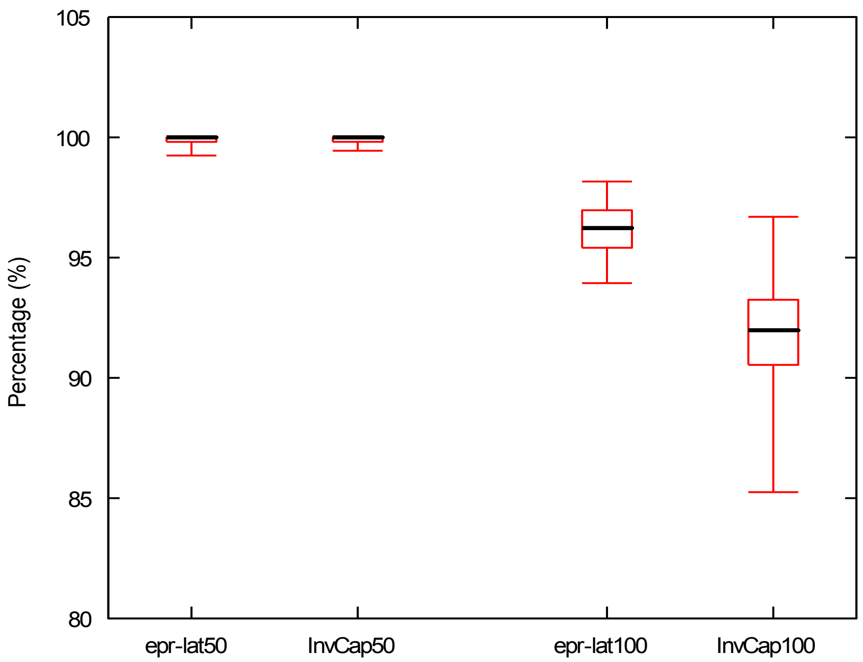

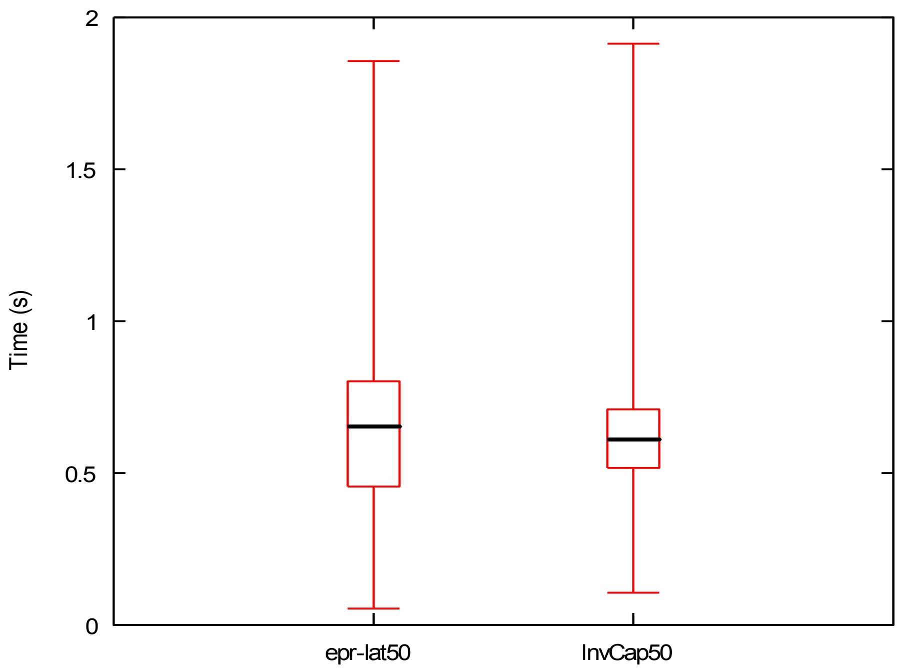

Application performance and end-to-end latency under REsPoNse-LAT is comparable to OSPF-InvCap at both lower and higher utilization levels.

Conclusion

- REsPoNse/EATe
  - Key idea: hybrid offline/online approach
- Properties
  - Stable
  - Incrementally deployable
  - Scalable

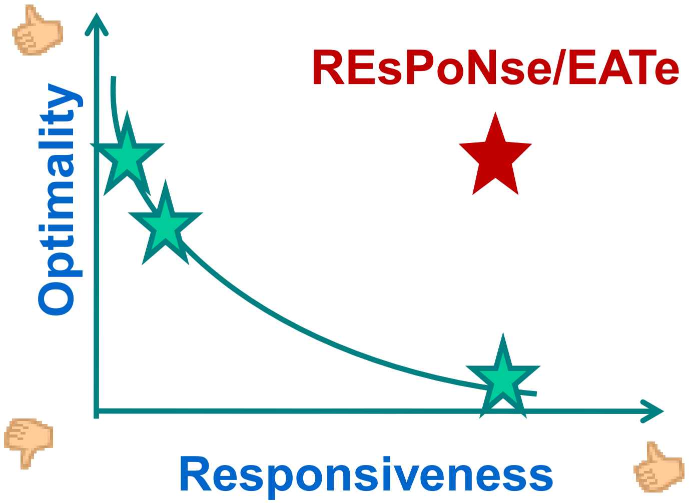

REsPoNse/EATe offers an optimal or a close-to-optimal solution, with good responsiveness