Robust guarantees for networks with flexible routing

*Sanjay Rao**

Joint work with Yiyang Chang and Mohit Tawarmalani^*

School of Electrical and Computer Engineering and*

*Krannert School of Management^*

Purdue University

State of network verification

- Network management: ad-hoc process in practice
- Contrast to software/hardware: design/verification tools a 10B$ industry [Mckeown, Sigcomm Keynote 2012]
- Significant progress in recent years
  - Correctness of data-plane (Anteater, HSA, Veriflow, ….)
  - Programming language and SMT-based approaches (Frenetic, Batfish, NoD ,….)
- Much of the focus on verifying properties such as:
  - No routing blackholes, honoring reachability policies etc.

Our work

- Go beyond verification of data-plane correctness
- An early step at formally reasoning about quantitative network properties
- Focus on a class of problems that seek to:
  - Guarantee a network can adequately cope with a range of traffic demands and failure scenarios
  - Guarantee acceptable link utilizations across traffic demands and failures

Key contributions

- Optimization framework for provable bounds on link utilizations across traffic demands and failures for a given network design
- Key challenge:
  - Routing flexibility leads to intractable non-convex, (possibly non-linear) problems
- Approach:
  - Draw on relaxations of non-linear problems (LP hierarchies)
  - Stronger bounds than can be obtained with oblivious strategies

Rest of the talk.. Two concrete case studies

Can a network cope with failures?

- Given upto ***f*** links may simultaneously fail, what is the worst case utilization of any link across all failure scenarios?
- Routing may be chosen in flexible fashion to adapt to any given failure.

Formulating utilization verification as an optimization problem

- Given a network design *t*, find the worst case utilization across all links *e*, across all failure scenarios *z* of interest, assuming optimal routing *y* for each scenario

$$\max_{z \in Z} \ \min_{y} \ \max_{e \in E} \ U_e(t, z, y)$$

All failure scenarios

Best routing for given scenario

Highest utilization across edges for given routing and failure scenario

Formulating utilization verification as an optimization problem

$$\max_{z \in Z} \boxed{\min_{y} \max_{e \in E} U_e(t, z, y)}$$

LP: Dualize for a a maximization Problem

z_{ij} = 1 if link <I,j> fails

$$\min_{y} \; S$$

$$Sc_{ij}(1 - z_{ij}) \geq \sum_{s,t} y_{ijst} \;\; \forall i, j$$

$$\sum_{j,k} y_{ijst} - \sum_{j,k} y_{jist} = \begin{cases} 0 & \forall s, t, i \neq s, i \neq t \\ d_{st} & \forall s, t, i = s \\ -d_{st} & \forall s, t, i = t \end{cases}$$

$$y_{ijst} \geq 0 \;\; \forall i, j, s, t$$

Formulating utilization verification as an optimization problem

$$\max_{u,\lambda,z} \quad \sum_{i,t} d_{it} u_{it}$$

$$\text{s.t.} \quad u_{it} - u_{jt} \leq \lambda_{ij} \quad \langle i,j \rangle \in E$$

$$\sum_{\langle i,j \rangle \in E} \lambda_{ij} c_{ij} (1 - z_{ij}) = 1$$

$$\sum_{\langle i,j \rangle \in E} z_{ij} = f$$

$$z_{ij} \in \{0,1\}; \quad u_{ij} \geq 0; \quad v_{tt} = 0$$

Intractability of problem

$$\begin{aligned}
\max_{u,\lambda,z} \quad & \sum_{i,t} d_{it} u_{it} \\
\text{s.t.} \quad & u_{it} - u_{jt} \leq \lambda_{ij} \quad \langle i,j \rangle \in E \\
& \sum_{\langle i,j \rangle \in E} \lambda_{ij} c_{ij} (1 - z_{ij}) = 1 \\
& \sum_{\langle i,j \rangle \in E} z_{ij} = f \\
& z_{ij} \in \{0,1\}; \quad u_{ij} \geq 0; \quad v_{tt} = 0
\end{aligned}$$

Appears non-linear. But we can prove bounds on the dual variables if graph connected after ***f*** failures. Can be linearized. Resulting problem still an MILP

Obtaining tractable relaxations

- RLT relaxations: general approach to relax non-convex problems into tractable LP
- Family of relaxations
- Higher levels of hierarchy
  - Converge to optimal value of the non-convex problem
  - Incur higher complexity

RLT relaxation: example

Min $xy - x + y$

$2 \le x \le 3$

$3 \le y \le 4$

Relaxation steps:

1. Multiply constraints with each other

   Example: $(x-2)(y-3) \ge 0 \Rightarrow xy - 2y - 3x + 6 \ge 0$

2. Replace products of variables xy, x^2, y^2 by new variables

3. Higher levels of RLT relaxation => multiply multiple constraints with each other

Our LP for utilization verification under failures

- First level RLT relaxation
- Minor change to original primal formulation to add slack, which constraints dual more and achieves tighter relaxations

Comparison with R3 (Sigcomm 2010)

- R3: Determines whether utilization < 1 or not under ***f*** failures
- Approach:
  - Convert failures into virtual demands
  - Use oblivious routing like strategies to get a tractable LP
- Main advantages of our approach:
  - Tighter relaxations
  - Can provide actual utilizations (not just whether above 1).
    - Useful to detect which failure scenarios are bad, which link’s capacity gets exceeded and by how much
  - Approach generalizes to other problems

Results: Abilene

# of edge failure	R3 (Sigcomm 10)	Our LP relaxation	MIP
0	0.122	0.122	0.122
1	0.372	0.163	0.163
2	0.622	0.244	0.244
3	0.872	0.488	0.488

- Each cell: utilization of most congested link
- Each edge: 2 parallel edges
- Real traffic matrix

Our RLT-based LP relaxation matches optimal, with tighter bounds than oblivious relaxation (R3)

Results: GEANT

# of edge failure	R3	Our LP relaxation	MIP
0	0.096	0.096	0.096
1	0.196	0.107	0.107
2	0.329	0.120	0.120
3	0.489	0.137	0.137
4	0.649	0.160	0.160
5	0.809	0.192	Insufficient resources to finish
6	0.849	0.240	
7	0.889	0.320	
8	0.929	0.480	
9	0.996	0.959	

- Each edge: 5 parallel edges.
- Traffic matrix: gravity model
- Runtime
  - Our LP relaxation: tens to hundreds of seconds
  - MIP: hours to tens of hours

Case Study II: MPLS tunnel selection

- Tunnels between ingress and egress to ensure a BGP free core
- With demand shifts: switch traffic across k pre-selected tunnels
- Desirable to change tunnels less frequently
  - Require changes to flow tables of internal switches
- For a given choice of tunnels, are utilizations of all links across all traffic demands of interest within acceptable limits?

Formulating utilization verification with tunneling

$$\max_{d \in D} \min_{y} \max_{e \in E} U_e(t, d, y)$$

t:Given choice of tunnels
D:Set of traffic demands
y:Split across tunnels for a given demand

Formulating utilization verification with tunneling

$$
\begin{aligned}
\max_{u,\lambda,d} \quad & \sum_{j \in J} \sum_{p \in P} c_{jp} d_p u_j \\
\text{s.t.} \quad & \sum_{j \in J} a_{ij} u_j - \sum_{e \in E} \lambda_e b_{ie} \leq 0 \quad i \in I \\
& \sum_{e \in E} \lambda_e = 1 \\
& \sum_{p \in P} h_{pl} d_p \leq q_l \quad l \in L \\
& (u_j)_{j \in J} \geq 0; \quad (\lambda_e)_{e \in E} \geq 0
\end{aligned}
$$

Formulating utilization verification with tunneling

$$
\begin{aligned}
\max_{u,\lambda,d} \quad & \sum_{j\in J}\sum_{p\in P} c_{jp} d_p u_j \\
\text{s.t.} \quad & \sum_{j\in J} a_{ij} u_j - \sum_{e\in E} \lambda_e b_{ie} \le 0 \quad i \in I \\
& \sum_{e\in E} \lambda_e = 1 \\
& \sum_{p\in P} h_{pl} d_p \le q_l \quad l \in L \\
& (u_j)_{j\in J} \ge 0; \quad (\lambda_e)_{e\in E} \ge 0
\end{aligned}
$$

Bi-linear objective

Relaxations considered

1. RLT relaxations

2. Oblivious relaxations

$$\max_{d \in D} \min_{y} \max_{e \in E} U_e(t, d, y)$$

is upper-bounded by

$$\min_{t \in T, y \in Y} \max_{d \in D} \max_{e \in E} U_e(t, d, y).$$

Theoretical results

- **Theorem:** The RLT relaxation is tighter than the oblivious relaxation
- **Proposition:** For predicted demands expressed as a convex combination of historical traffic matrices, it is sufficient to consider the corner points. The verification problem is an LP
- Side result: General set of conditions that explain why oblivious formulation is tractable, and the verification problem is not

Evaluation of tunnel selection strategies

- Tunnel selection strategies
  - K-Shortest (e.g., SWAN, Sigcomm 13)
  - Shortest-Disjoint (e.g., SOL, NSDI 16)
  - Robust tunnel selection
    - Oblivious routing + tunnel decomposition

Bounds on utilization (Abilene)

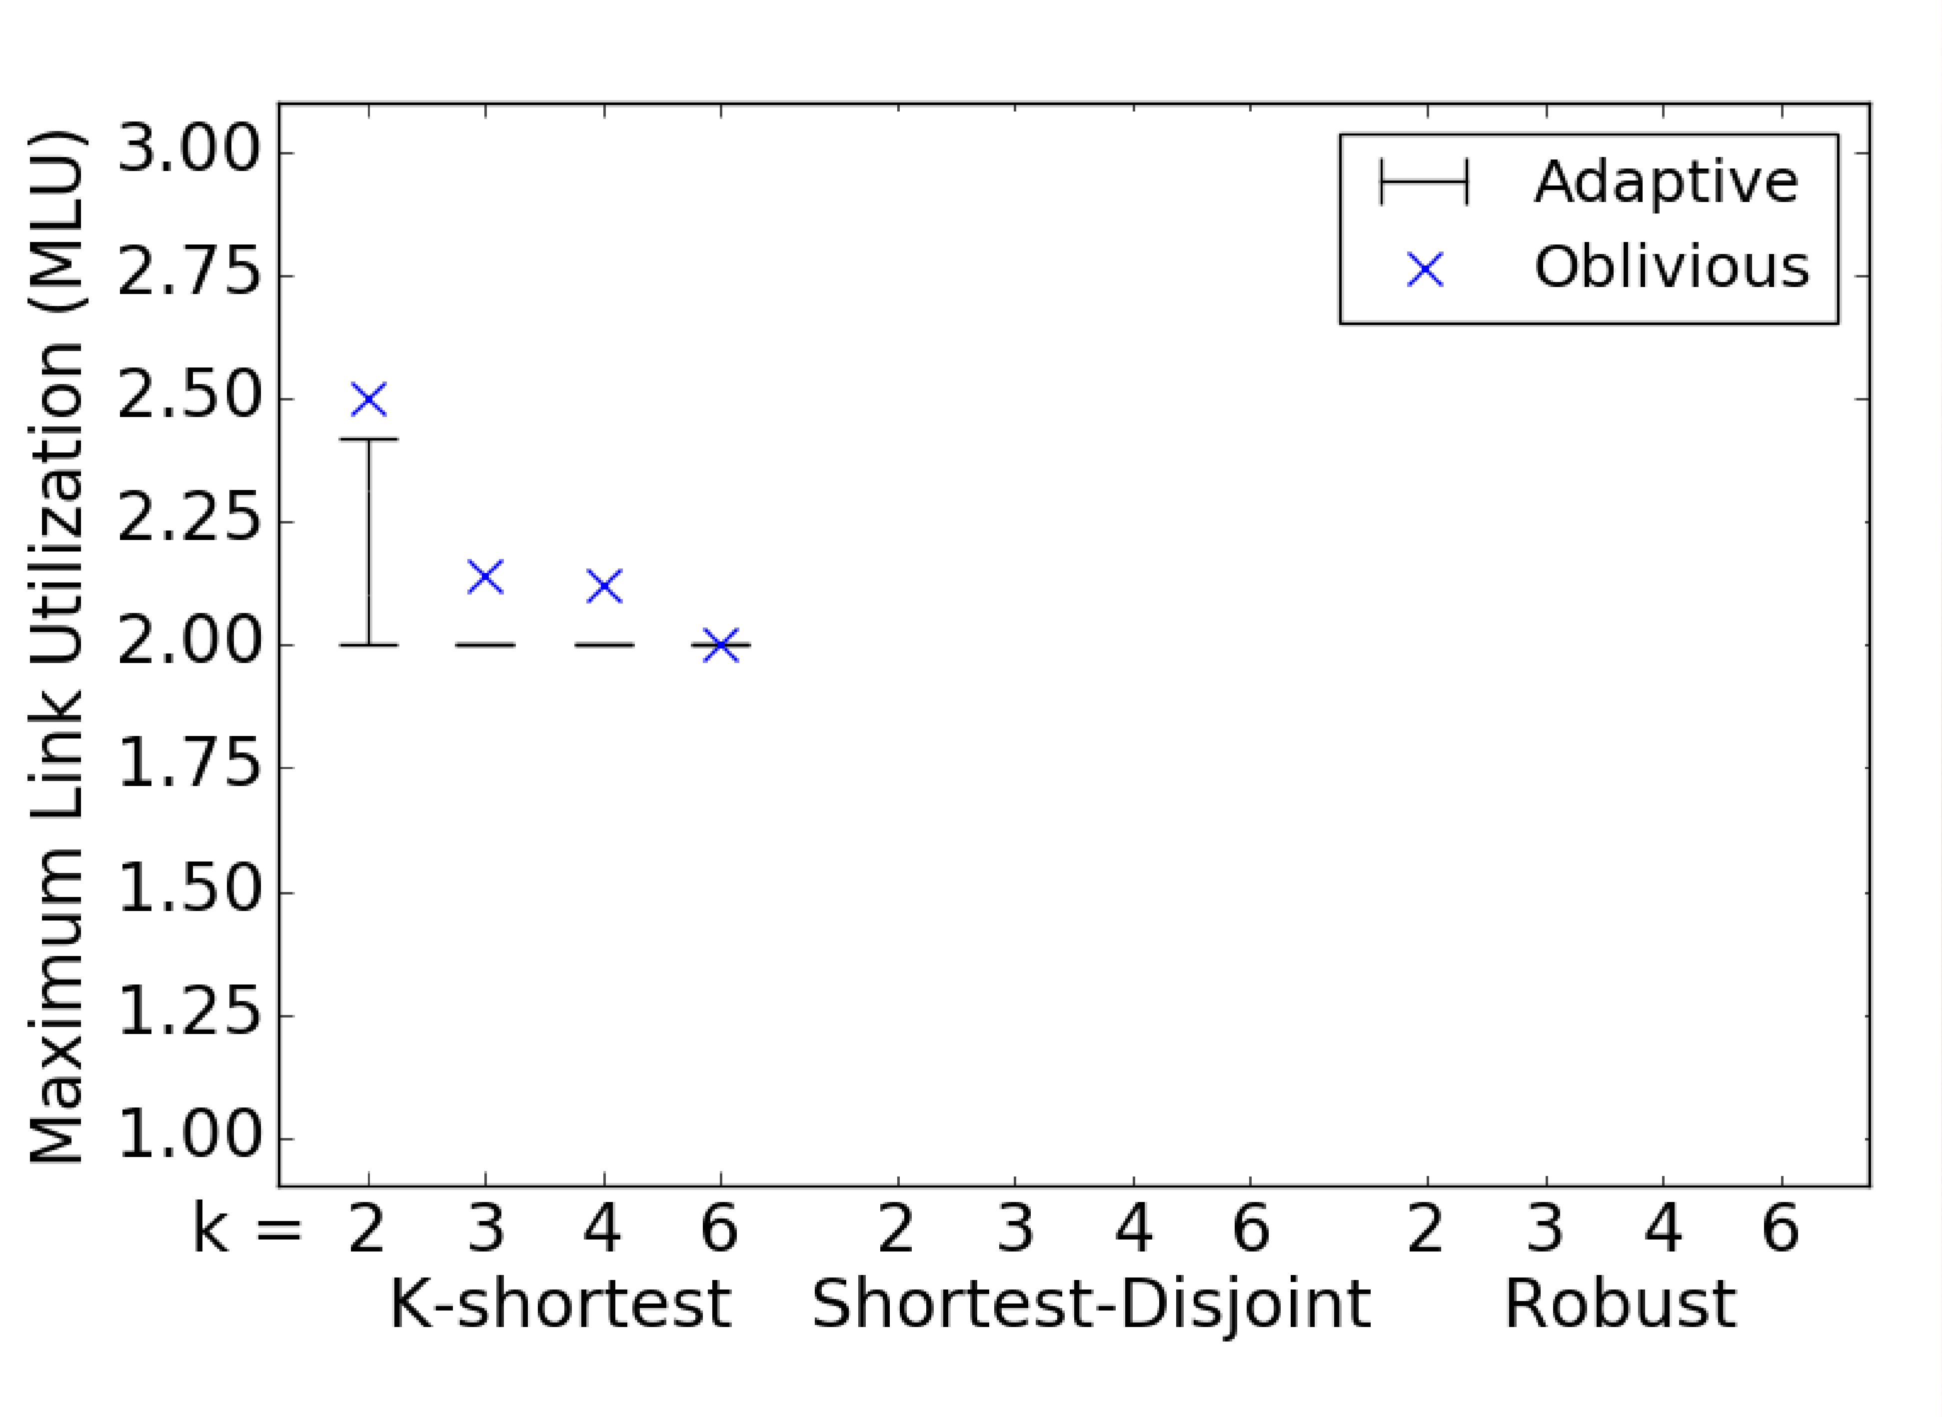

Bounds on utilization (Abilene)

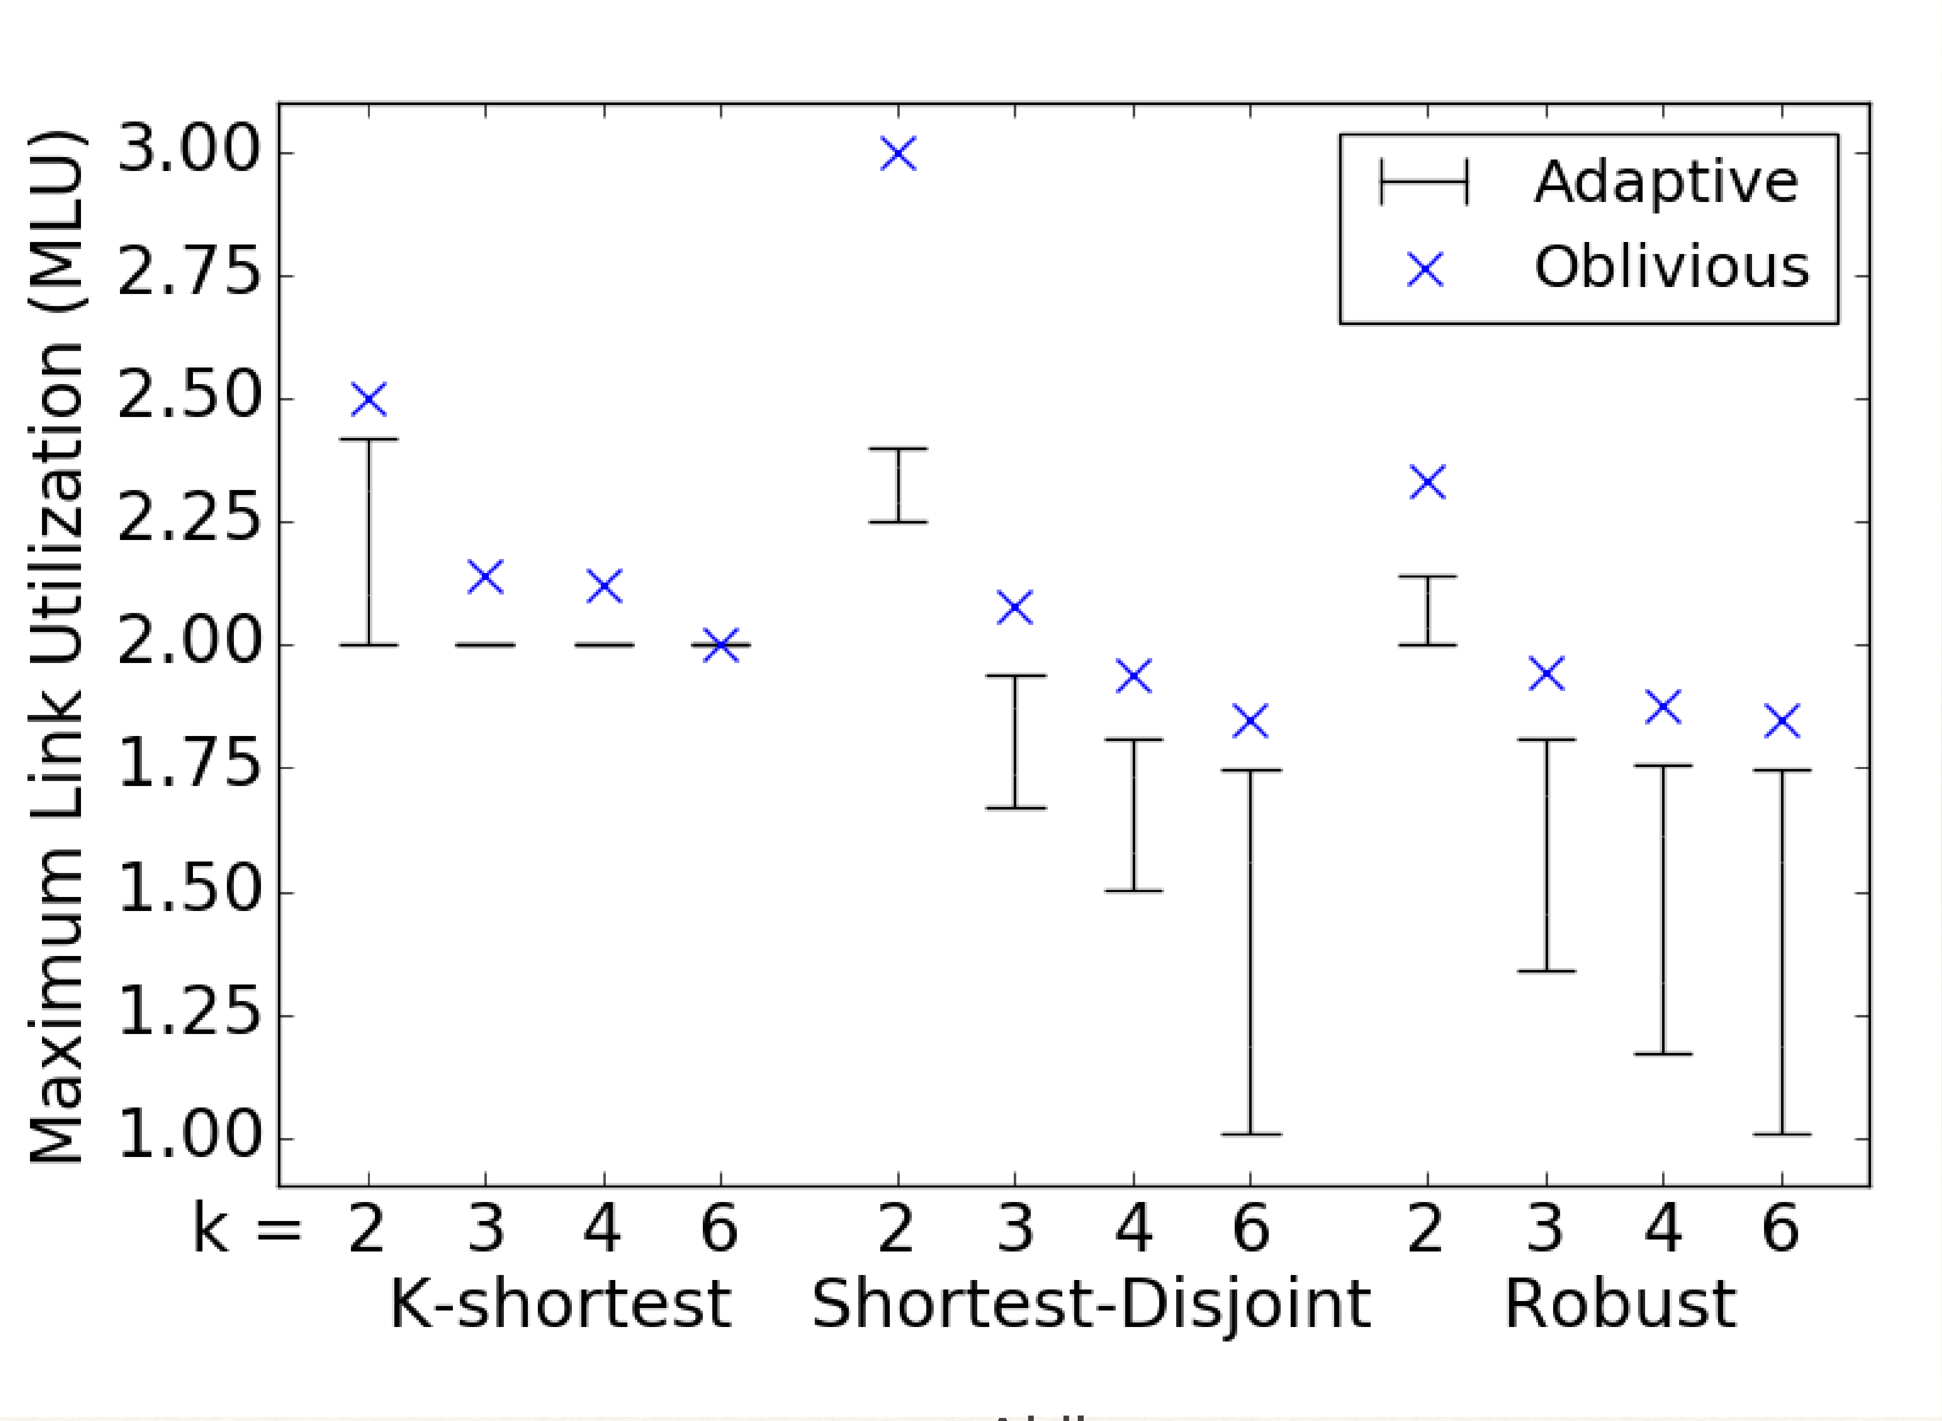

Abilene

Bounds on utilization (ANS)

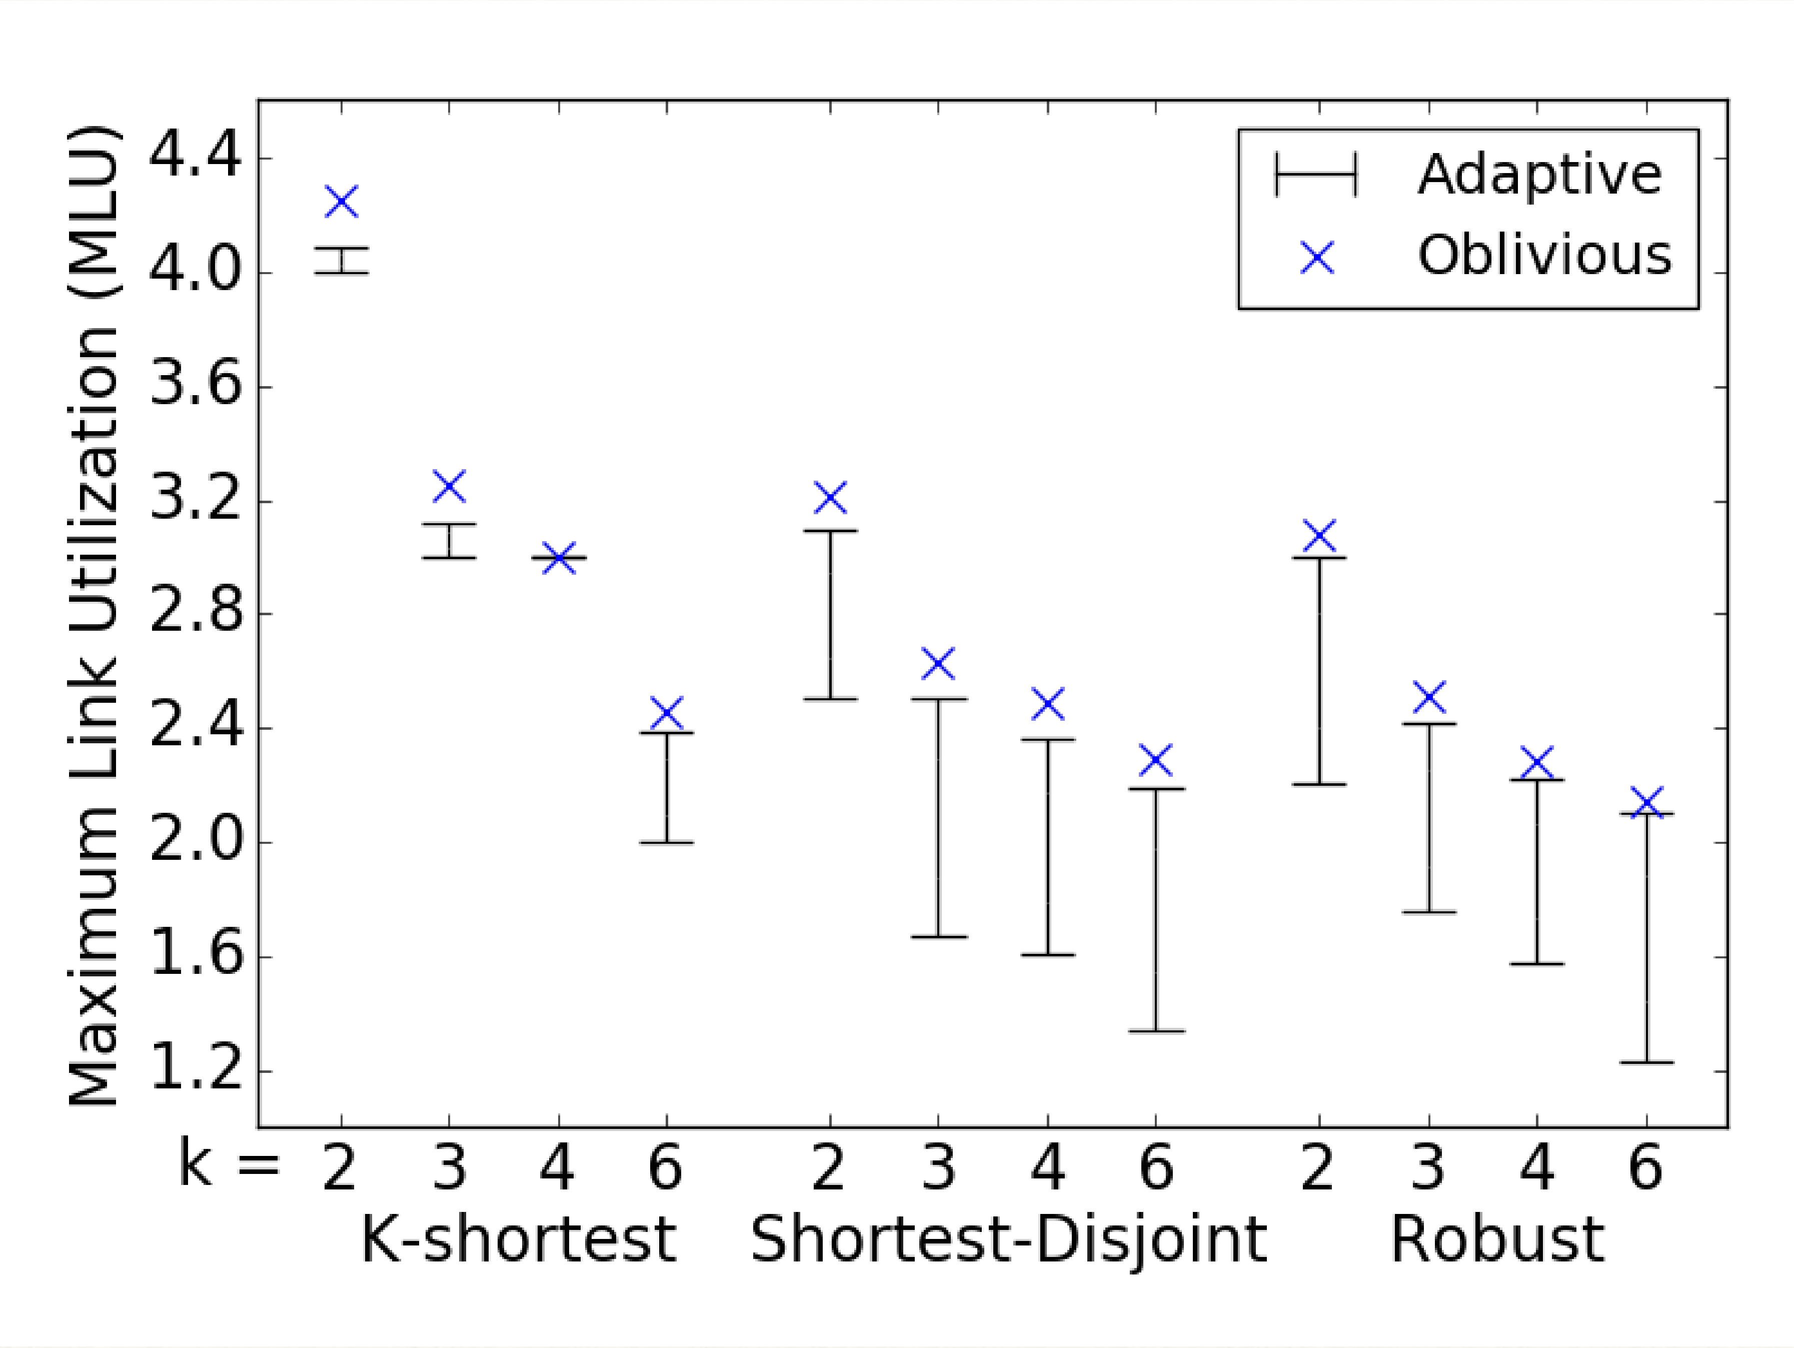

Bounds on utilization (GEANT)

Memory requirements high for RLT relaxation (not done)
Standard decomposition techniques could be employed to reduce requirements

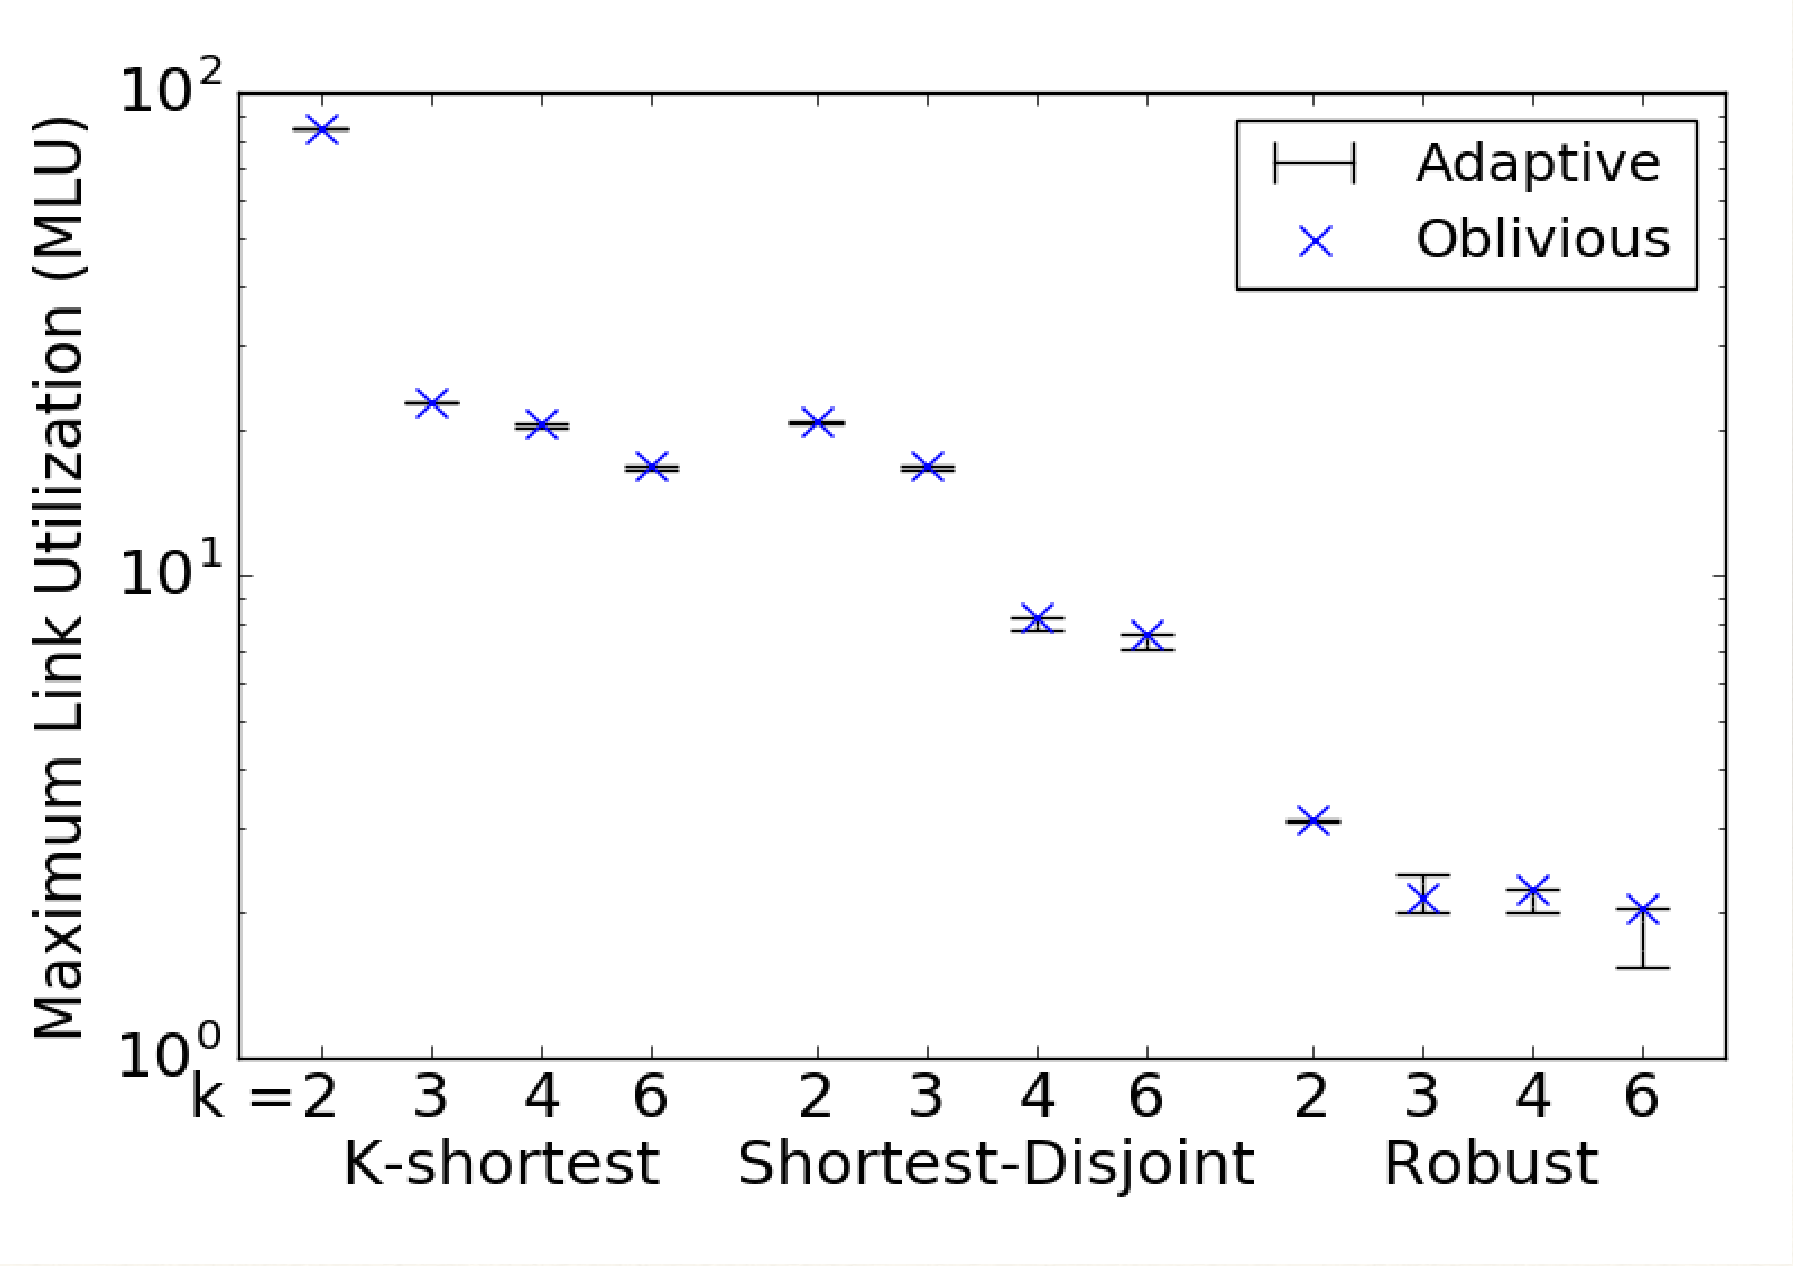

Conclusions

- Generic optimization framework to verify bounds on network link utilizations across failures/traffic demands
- RLT relaxations provide tighter bounds than oblivious
  - Oblivious relaxations still valuable
- Open questions for theoretical researchers:
  - Limits and opportunities with RLT hierarchies
  - Robust optimization: relating degree of adaptivity to level in RLT hierarchy