Power Aware and Computationally Efficient Optical Network Design

George N. Rouskas

Department of Computer Science

North Carolina State University

Joint work with: Dr. Emre Yetginer (Tubitak, Turkey), Zeyu Liu (NCSU)

Outline

- Power-Aware Traffic Grooming
  - Power Consumption in Networks: Trends and Challenges
  - Optical Networks to the Rescue: Power-Aware Traffic Grooming
  - Results and Discussion
- Computationally Scalable Optical Network Design
  - Routing and Wavelength Assignment (RWA)
  - New Computationally Efficient ILP Formulations for Ring and Mesh
  - Numerical Results
- Conclusions and Future Research Directions

The Challenge of Power Consumption

- Power consumption a growing challenge for ICT industry:
  - high operating costs
  - high capital costs $\longrightarrow$ cooling equipment
- Significant environmental impact
  - industry responsible for $\approx$2-3% of man-made CO_2
  - growing at double-digit rates

Why Energy Efficiency For Networks

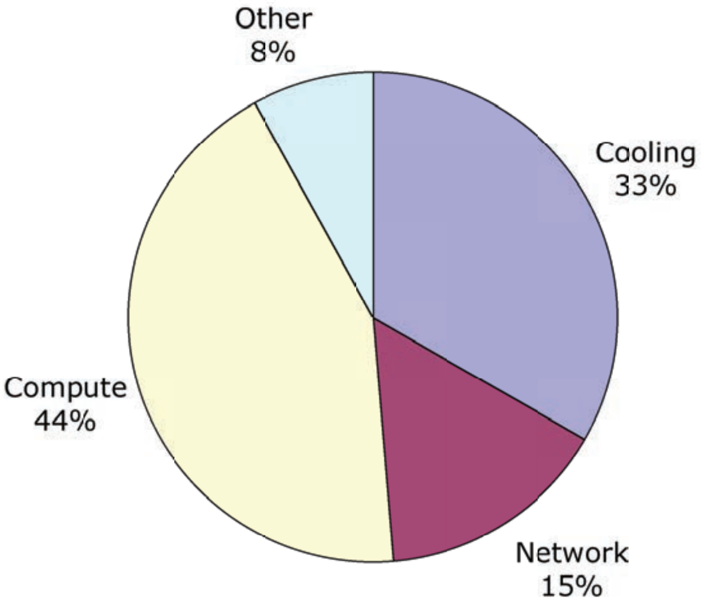

Why Energy Efficiency For Networks

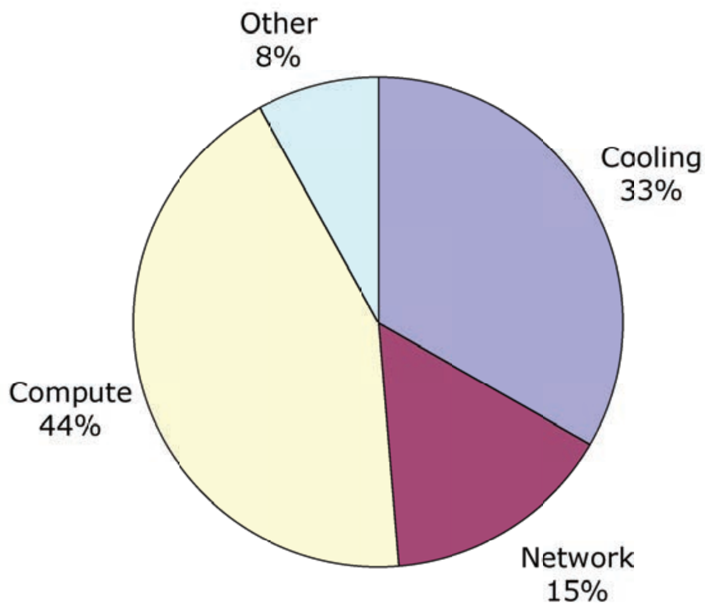

- So far, energy efficiency focus has been on servers and cooling
- Networks are shared resources $\longrightarrow$ always on
- In the US: 6 TWatts of power on networks

Addressing the Challenge

- Energy-efficient designs:
  1. low-power techniques in design of components
     - support low-power states in processors, memory, disks
     - disable clock signal to unused parts of processor
     - replace complex uniprocessors with multiple simple cores
  2. power management techniques across systems
     - intelligent policies to exploit low-power states
     - workload management

Addressing the Challenge

- Energy-efficient designs:
  1. low-power techniques in design of components
     - support low-power states in processors, memory, disks
     - disable clock signal to unused parts of processor
     - replace complex uniprocessors with multiple simple cores
  2. power management techniques across systems
     - intelligent policies to exploit low-power states
     - workload management
- Seek inexpensive energy sources

  $\longrightarrow$ build data/compute centers wherever energy is cheap

The Networking Infrastructure

- Forwarding table lookup $\longrightarrow$ routers operate at very high speeds
  - high energy consumption
  - low-power operation not feasible

The Networking Infrastructure

- Forwarding table lookup $\longrightarrow$ routers operate at very high speeds
  - high energy consumption
  - low-power operation not feasible
- New routing architecture?
  - partition Internet address space
  - multiple parallel networks of “virtual” routers
  - each network handles small address space $\longrightarrow$ energy-efficient routers

Trends: Traffic vs. Router Capacity Growth

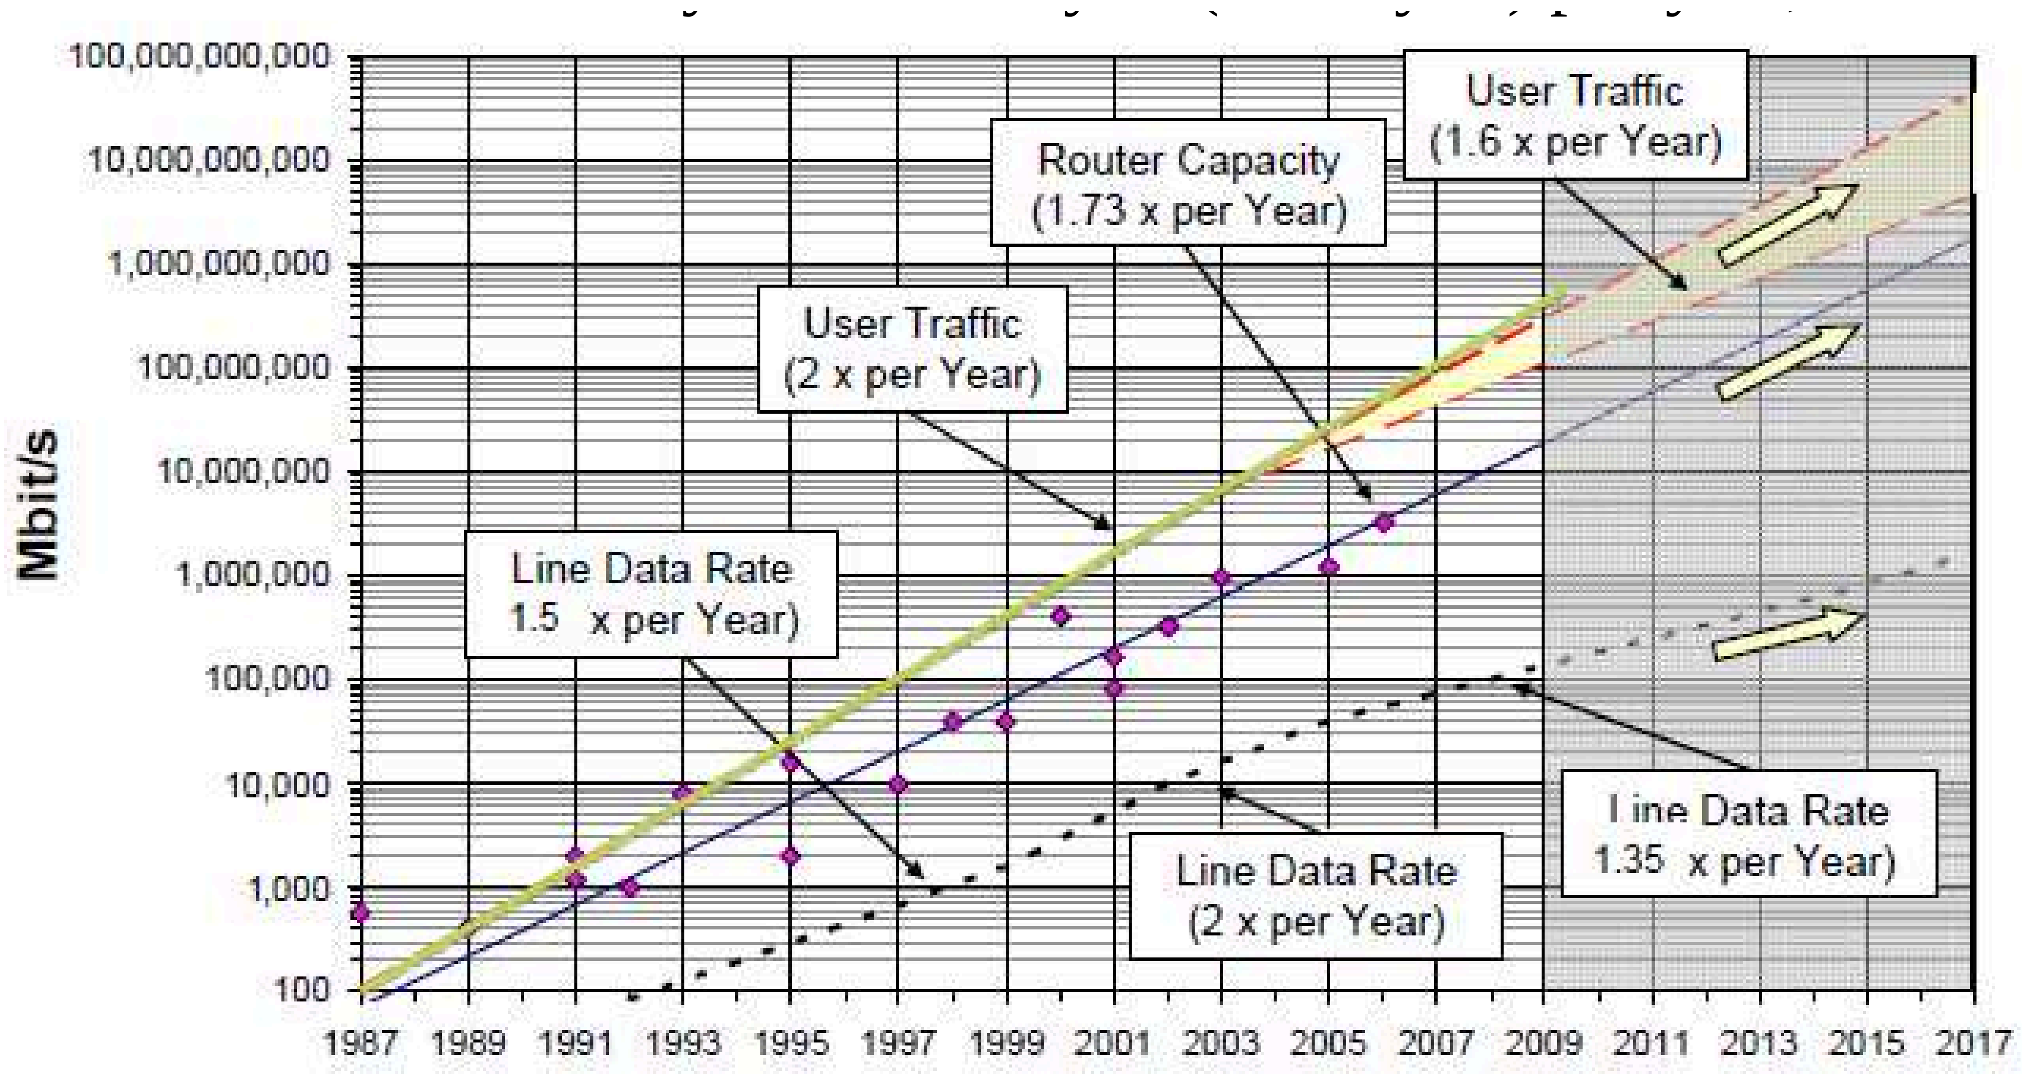

Trends: Router Power Consumption

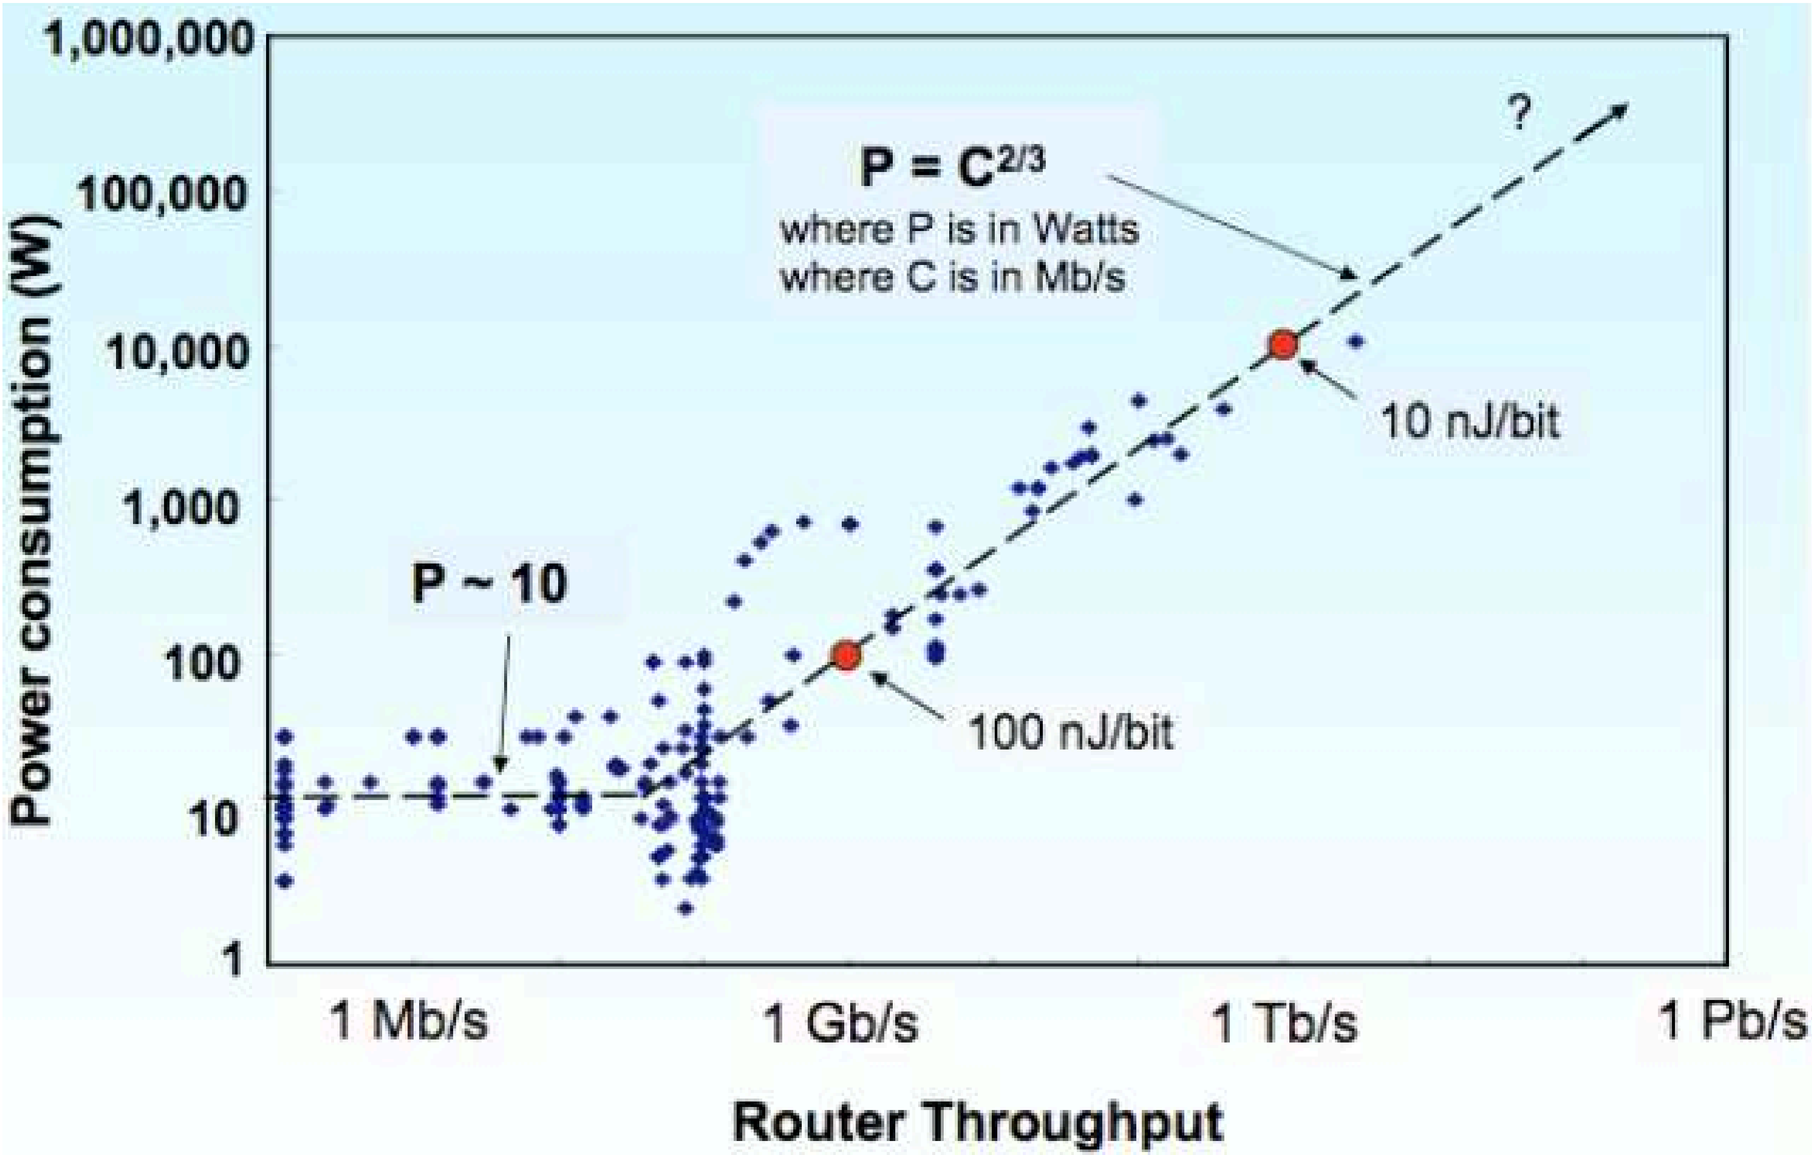

Trends: Energy Demand Will Exceed Supply

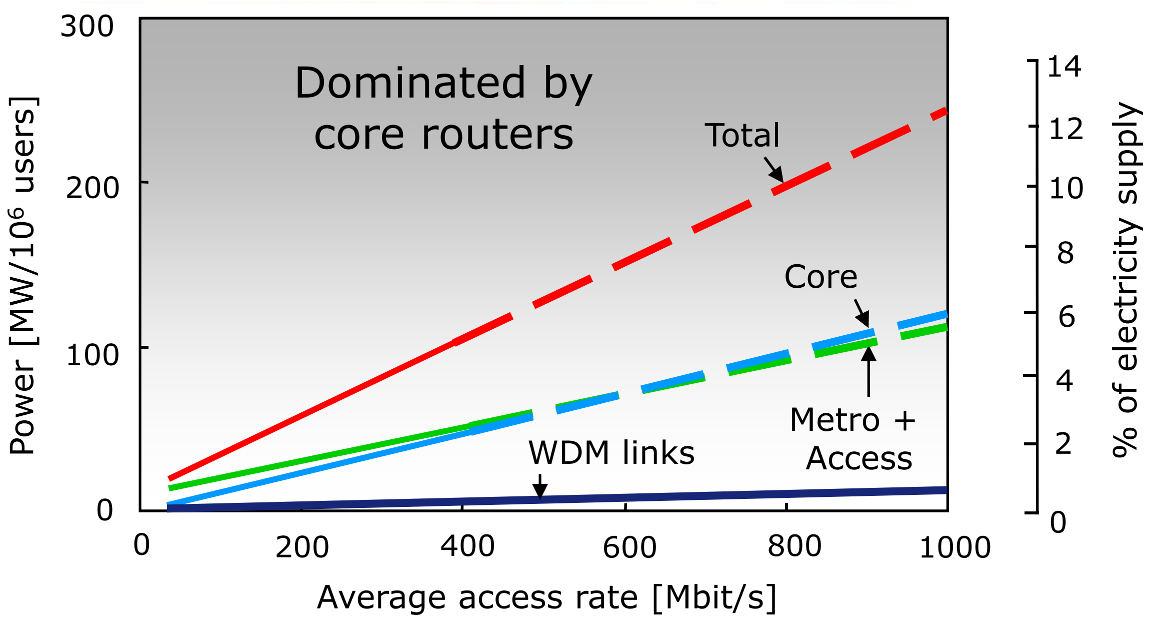

Trends: Energy Demand Will Exceed Supply

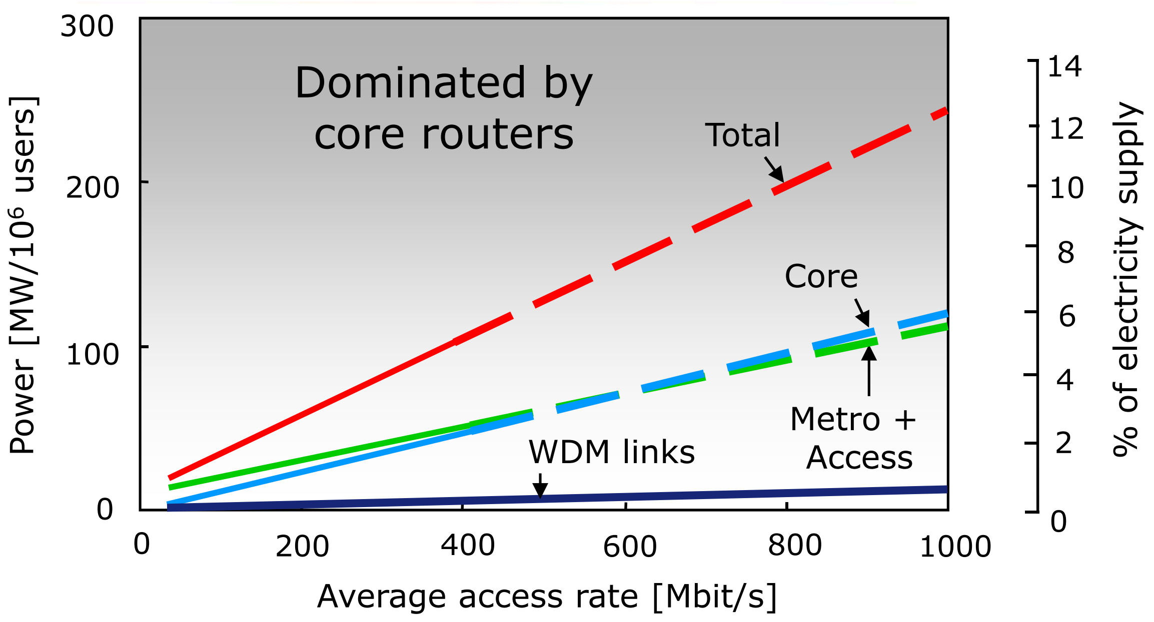

If 33% of the world's population were to obtain broadband access:

Access rate	1 Mbps	10 Mbps
Power consumption	100 GW	1 Tw
electricity supply	5%	50%

Optical Networks to the Rescue

- Optical networks:
  - energy efficient
    - many passive components
    - active components (e.g., repeaters) can be solar/wind-powered
  - low carbon footprint

Motivation: Router Power Consumption

Juniper Core Router T640

- 8 ports at 40 Gbps each
- Power consumption:
  - 4500 W overall, 550 W/port
- Cost (10c/kWh):
  - \$4000/year, \$500/port/year
- Add AC+UPS:
  - $\approx$ double power consumption $\longrightarrow$ \$1000/port/year
- Power consumption increases with line rate

Motivation: Optical Switch Power Consumption

Calient DiamondWave PXC 128

- 128×128 switch
- Power Consumption:
  - $<$ 750 W overall
  - $<$ 6 W/port
  - independent of line rate
- PXC consumes $\approx$ 1% of power per port consumed by the Juniper router

The Case for Optical Bypass

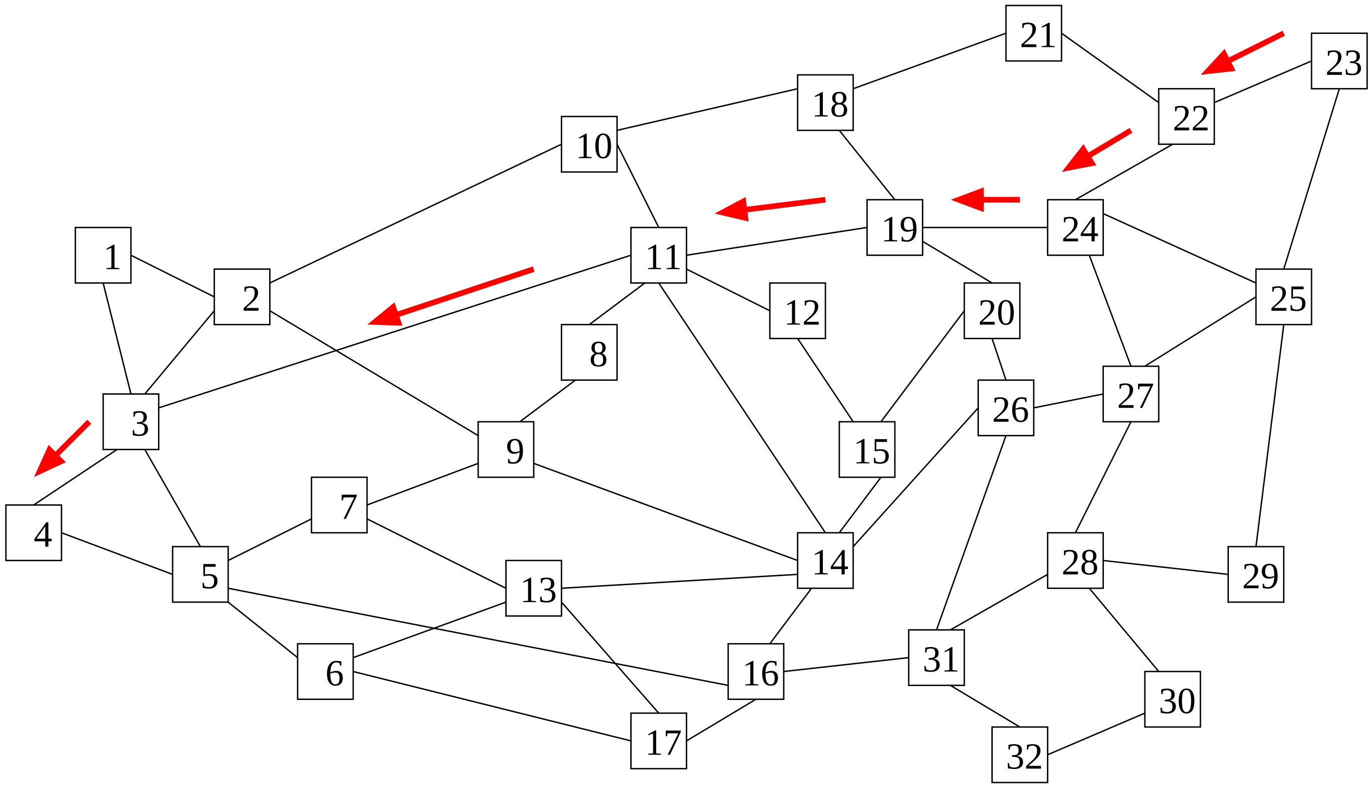

- Most ($\approx$ 80%) network links: $<$ 200 miles in length
- Most traffic demands ($\approx$ 80%): travel $>$ 200 miles

The Case for Optical Bypass

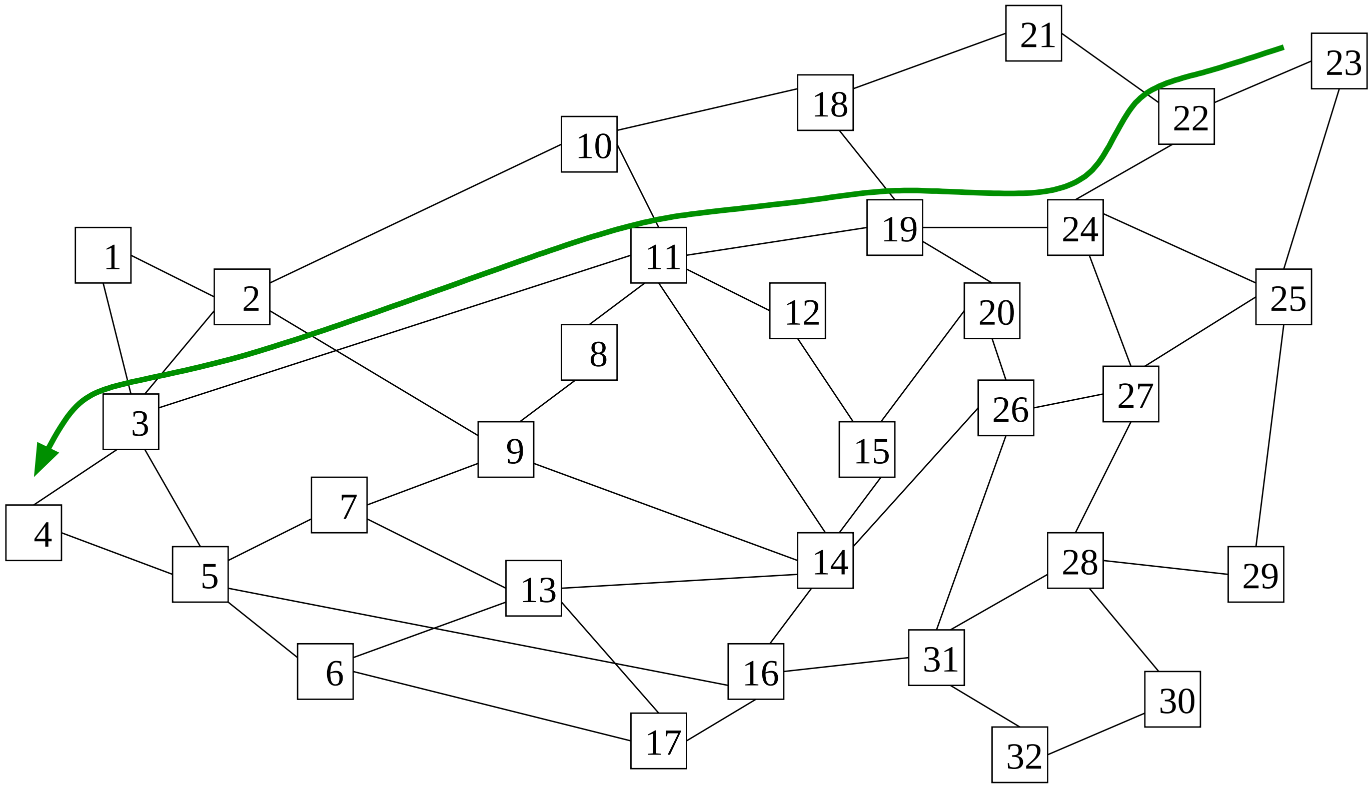

- Most (≈ 80%) network links: $<$ 200 miles in length
- Most traffic demands (≈ 80%): travel $>$ 200 miles

Grooming Networks

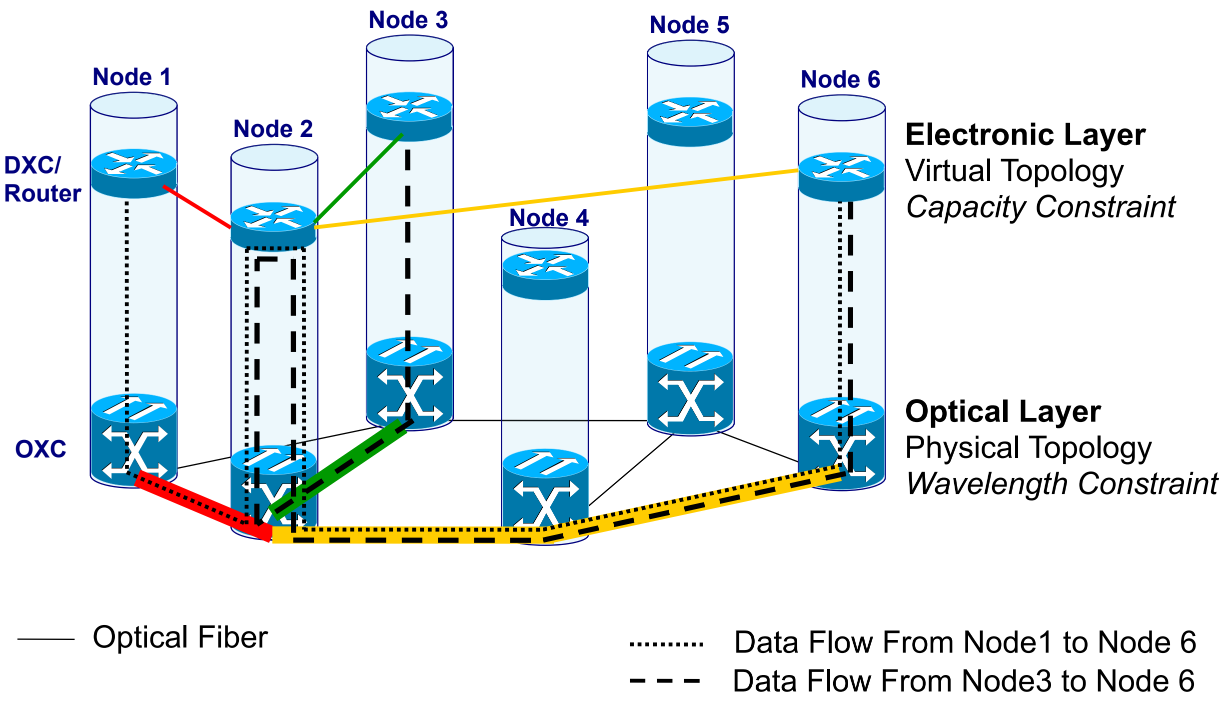

- What is traffic grooming?

  Efficiently set up lightpaths and groom (i.e., pack/unpack, switch, route, etc.) low-speed traffic onto high capacity wavelengths so as to minimize network resources

Traffic Grooming as Optimization Problem

- Inputs to the problem:
  - physical network topology (fiber layout)
  - traffic matrix $T = [t_{sd}] \longrightarrow$ int multiples of unit rate (e.g., OC-3)
- Output:
  - logical topology
  - lightpath routing and wavelength assignment (RWA)
  - traffic grooming on lightpaths

Traffic Grooming Subproblems

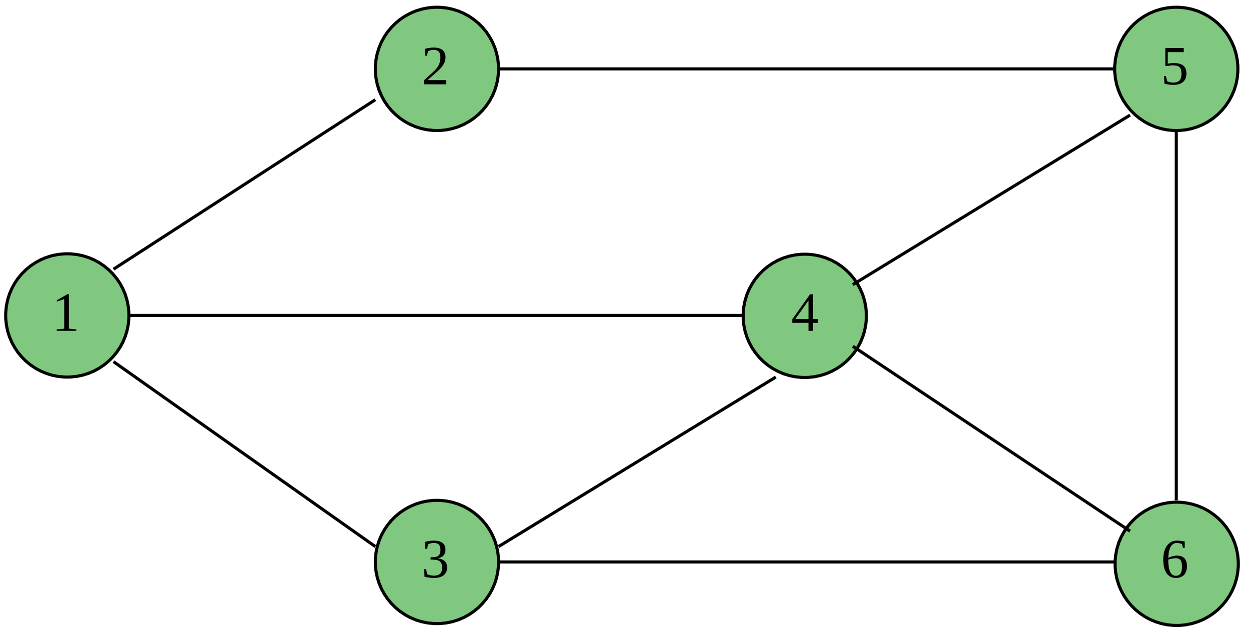

Traffic Grooming Subproblems

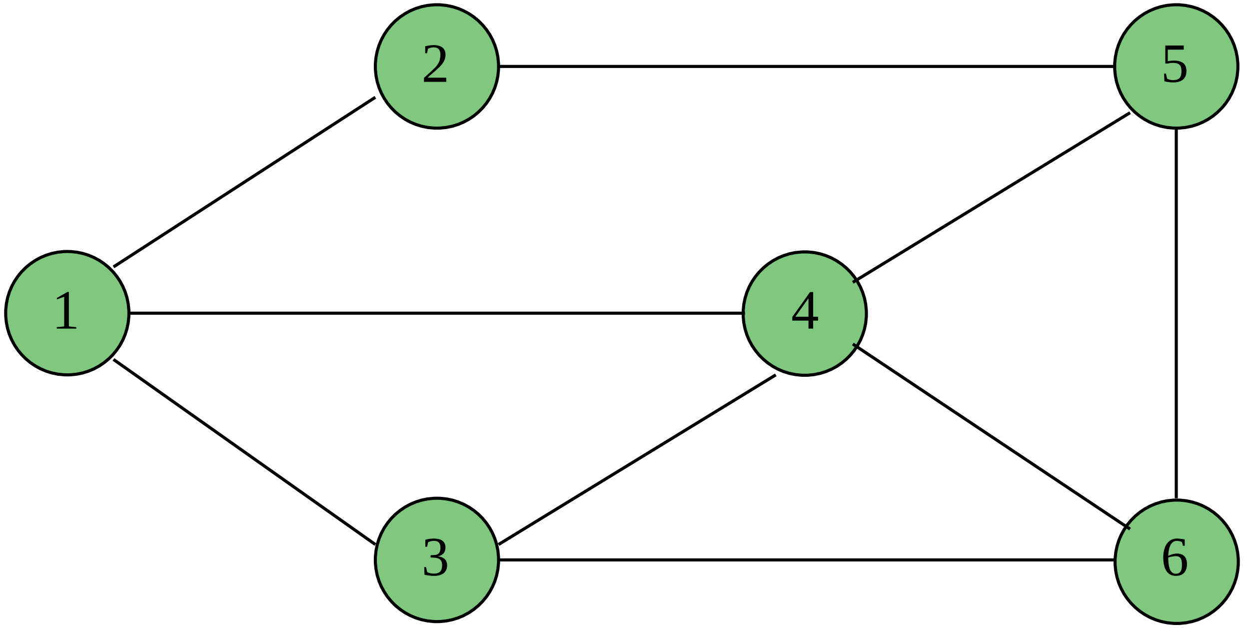

- Logical topology design → determine the lightpaths to be established

Traffic Grooming Subproblems

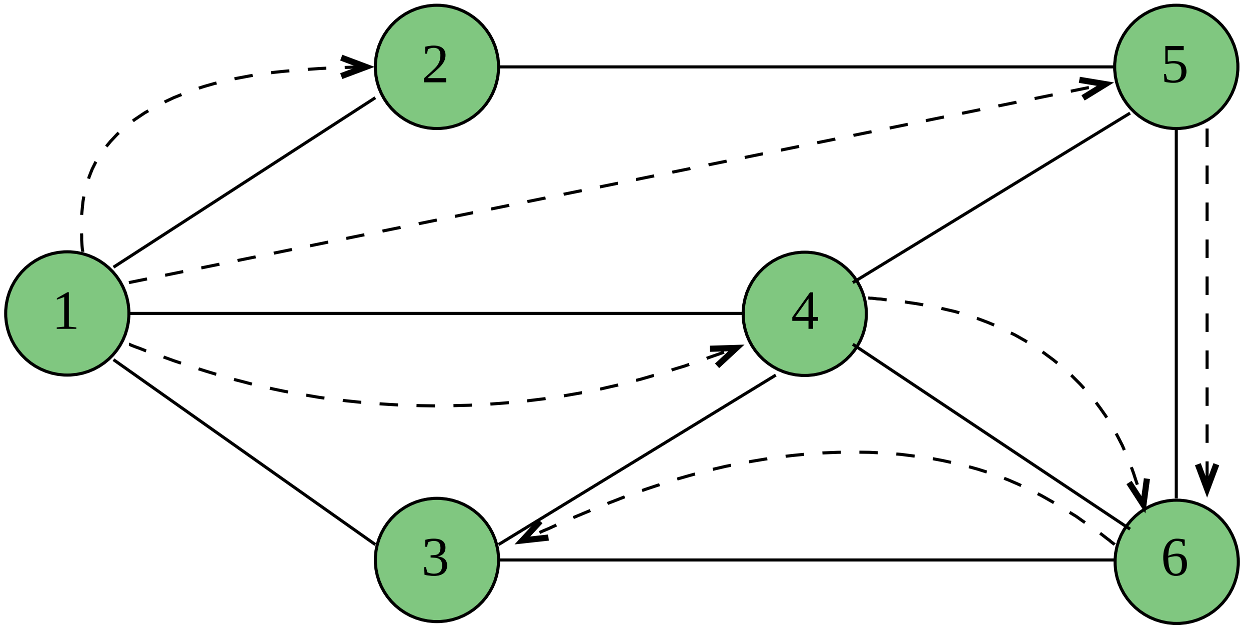

- Logical topology design → determine the lightpaths to be established
- Lightpath routing → route the lightpaths over the physical topology

Traffic Grooming Subproblems

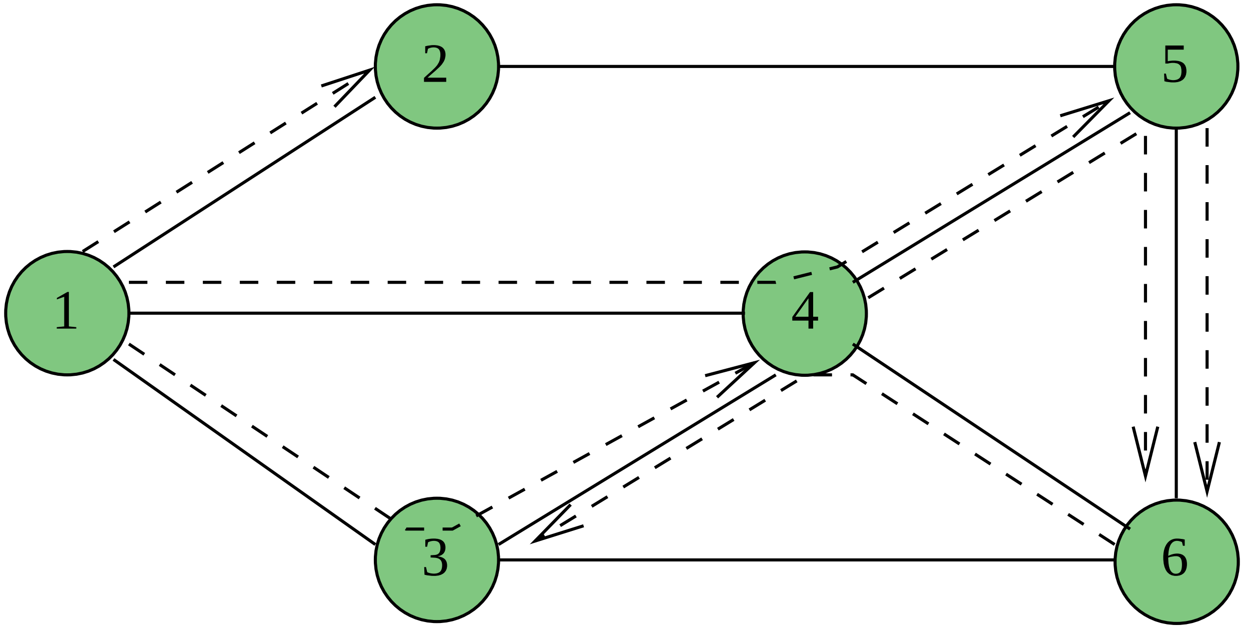

- Logical topology design → determine the lightpaths to be established
- Lightpath routing → route the lightpaths over the physical topology
- Wavelength assignment → assign wavelengths to lightpaths w/o clash

Traffic Grooming Subproblems

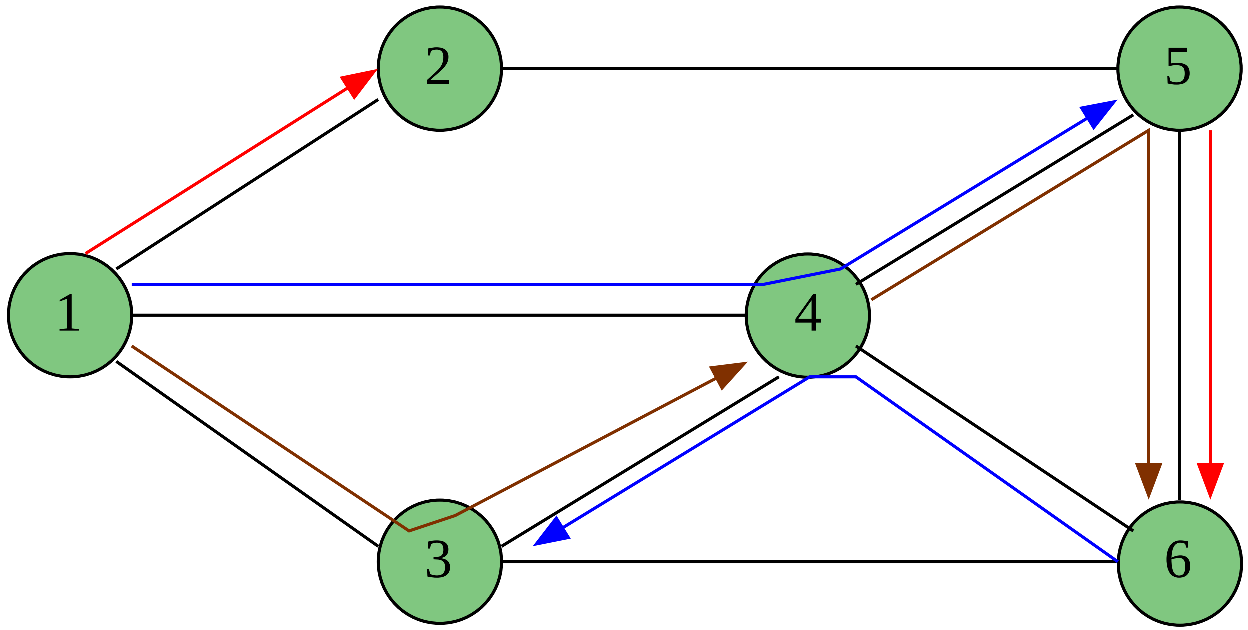

- Logical topology design $\longrightarrow$ determine the lightpaths to be established
- Lightpath routing $\longrightarrow$ route the lightpaths over the physical topology
- Wavelength assignment $\longrightarrow$ assign wavelengths to lightpaths w/o clash
- Traffic grooming $\longrightarrow$ route traffic on virtual topology

Grooming Objectives

- Minimize the number of lightpaths $\longrightarrow$ minL
  - equivalent to minimizing the number of electronic ports
  - minimizes initial deployment cost

Grooming Objectives

- Minimize the number of lightpaths $\longrightarrow$ minL
  - equivalent to minimizing the number of electronic ports
  - minimizes initial deployment cost
- Minimize the amount of electronically switched traffic $\longrightarrow$ minT
  - minimizes average processing delay
  - minimizes electronic switching capacity

Grooming Objectives

- Minimize the number of lightpaths $\longrightarrow$ minL
  - equivalent to minimizing the number of electronic ports
  - minimizes initial deployment cost
- Minimize the amount of electronically switched traffic $\longrightarrow$ minT
  - minimizes average processing delay
  - minimizes electronic switching capacity
- Minimize the amount of power consumption $\longrightarrow$ minP
  - maximizes power efficiency (in Watts/bit)
  - minimizes operational costs
  - most general objective

Power Consumption: Assumptions

- Optical power $\ll$ Electronic power
  $\longrightarrow$ energy consumed by optical ports is negligible
- Inactive ports and transceivers may be shut down
- Power consumption of each component (electronic input/output port, O/E and E/O converters) increases linearly with amount of traffic handled

Power Consumption: Router Port Model

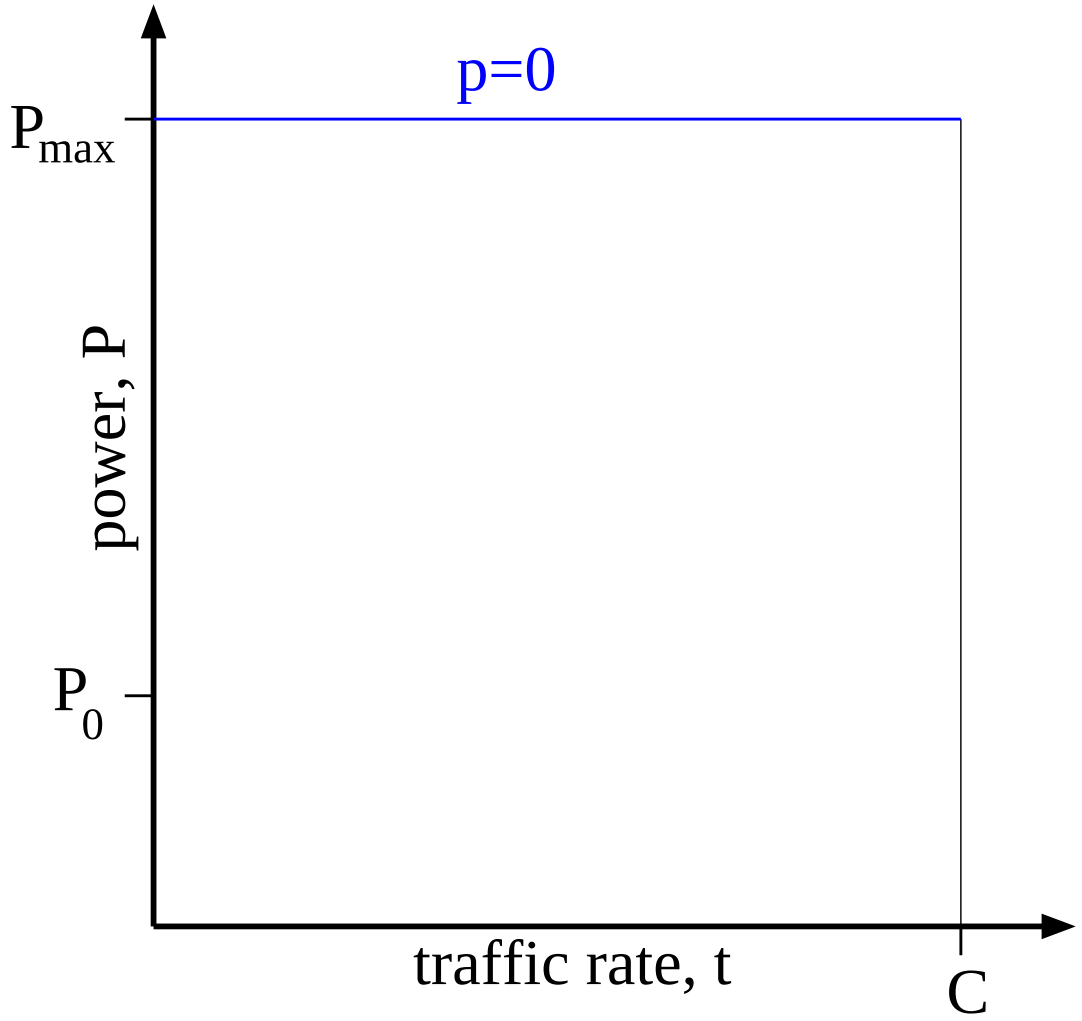

Equivalent to minimizing the number of lightpaths $\longrightarrow$ minL

Power Consumption: Router Port Model

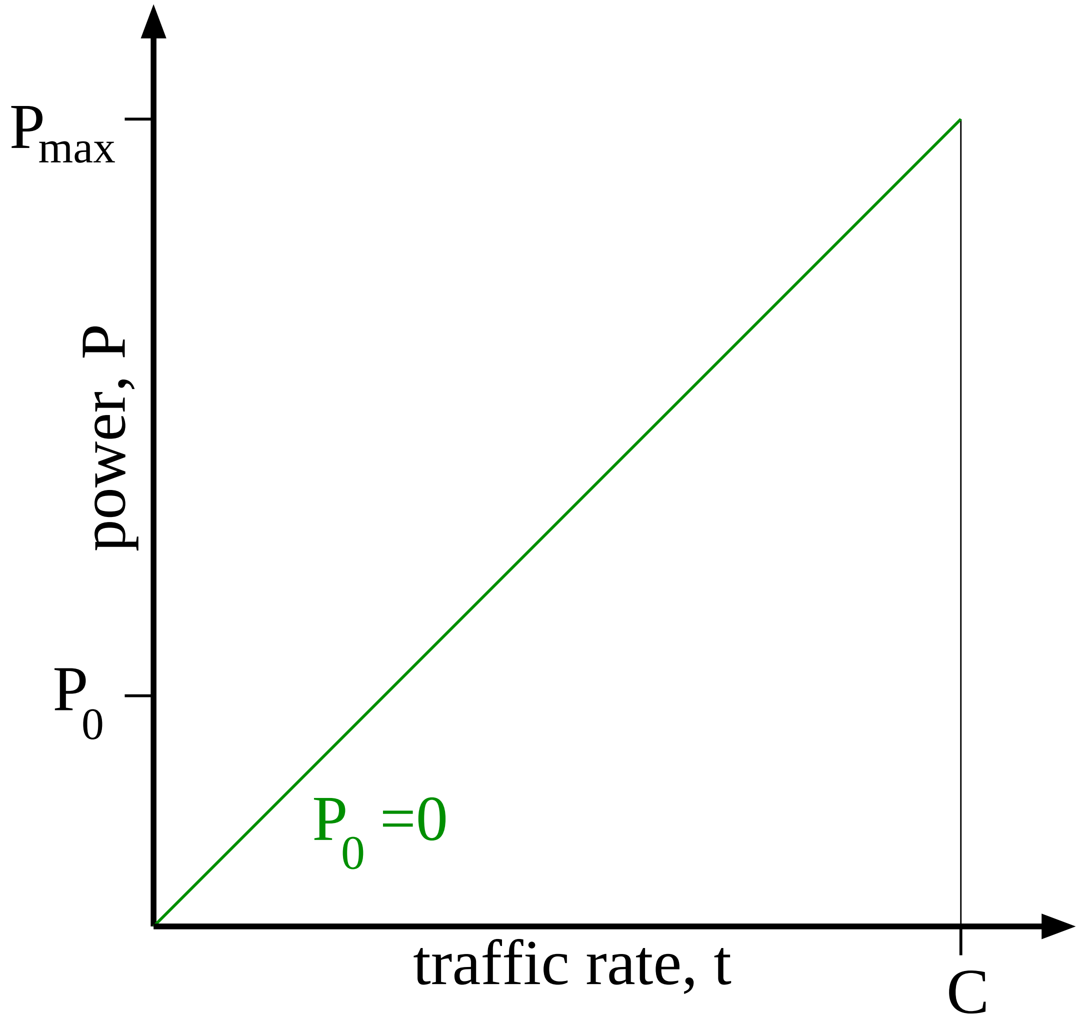

Equivalent to minimizing amount of electronically switched traffic $\longrightarrow$ minT

Power Consumption: Router Port Model

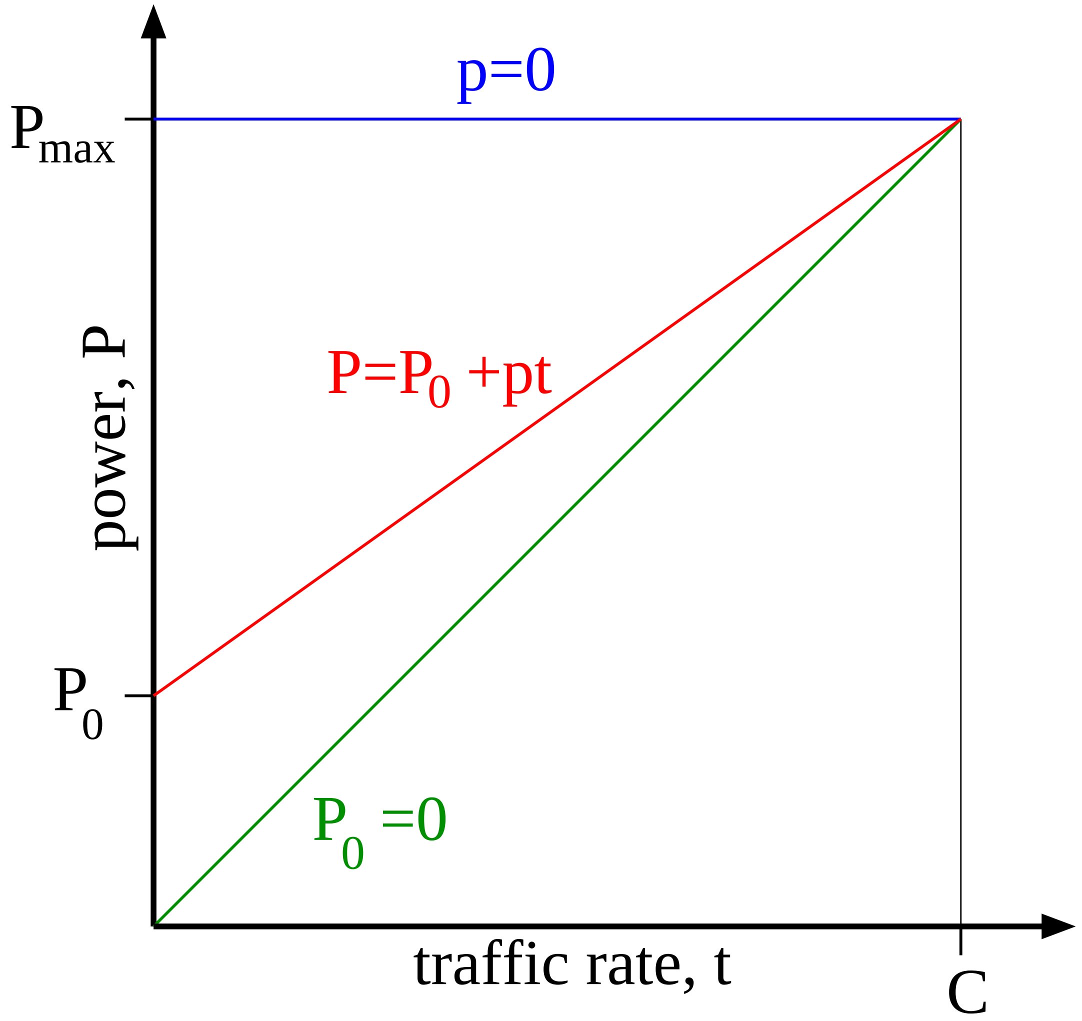

Most general model: minimize power consumption $\longrightarrow$ minP

ILP Formulation

- Objective:
  1. minL: min # of lightpaths
  2. minT: min amount of electronically switched traffic
  3. minP: min power consumption $\longrightarrow$ most general model
- subject to:
  - lightpath routing constraints
  - wavelength assignment constraints
  - traffic routing constraints

Performance Evaluation

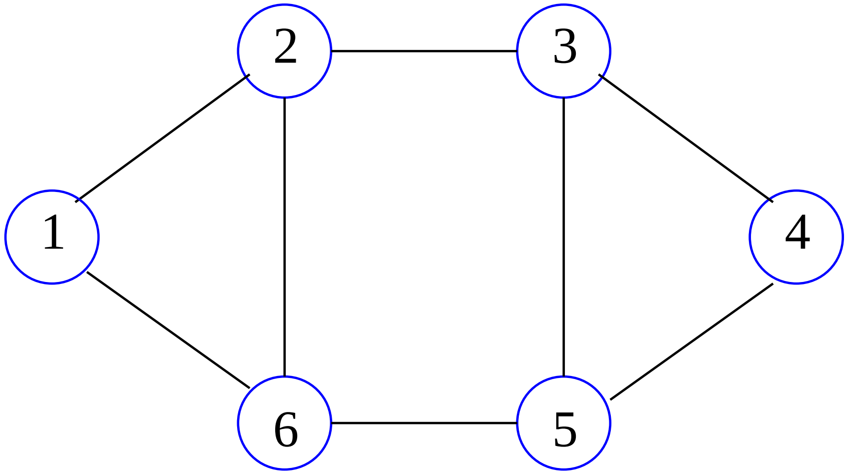

- $W = 3$ wavelengths
- $C = 48$ wavelength capacity
- Source-destination traffic $t_{sd} \leftarrow uniform[0, t_{max}]$
- $P_0 = 0.25$
- $P_{max} = 1$
- Each data point: average of 40 problem instances

Results: Number of Lightpaths

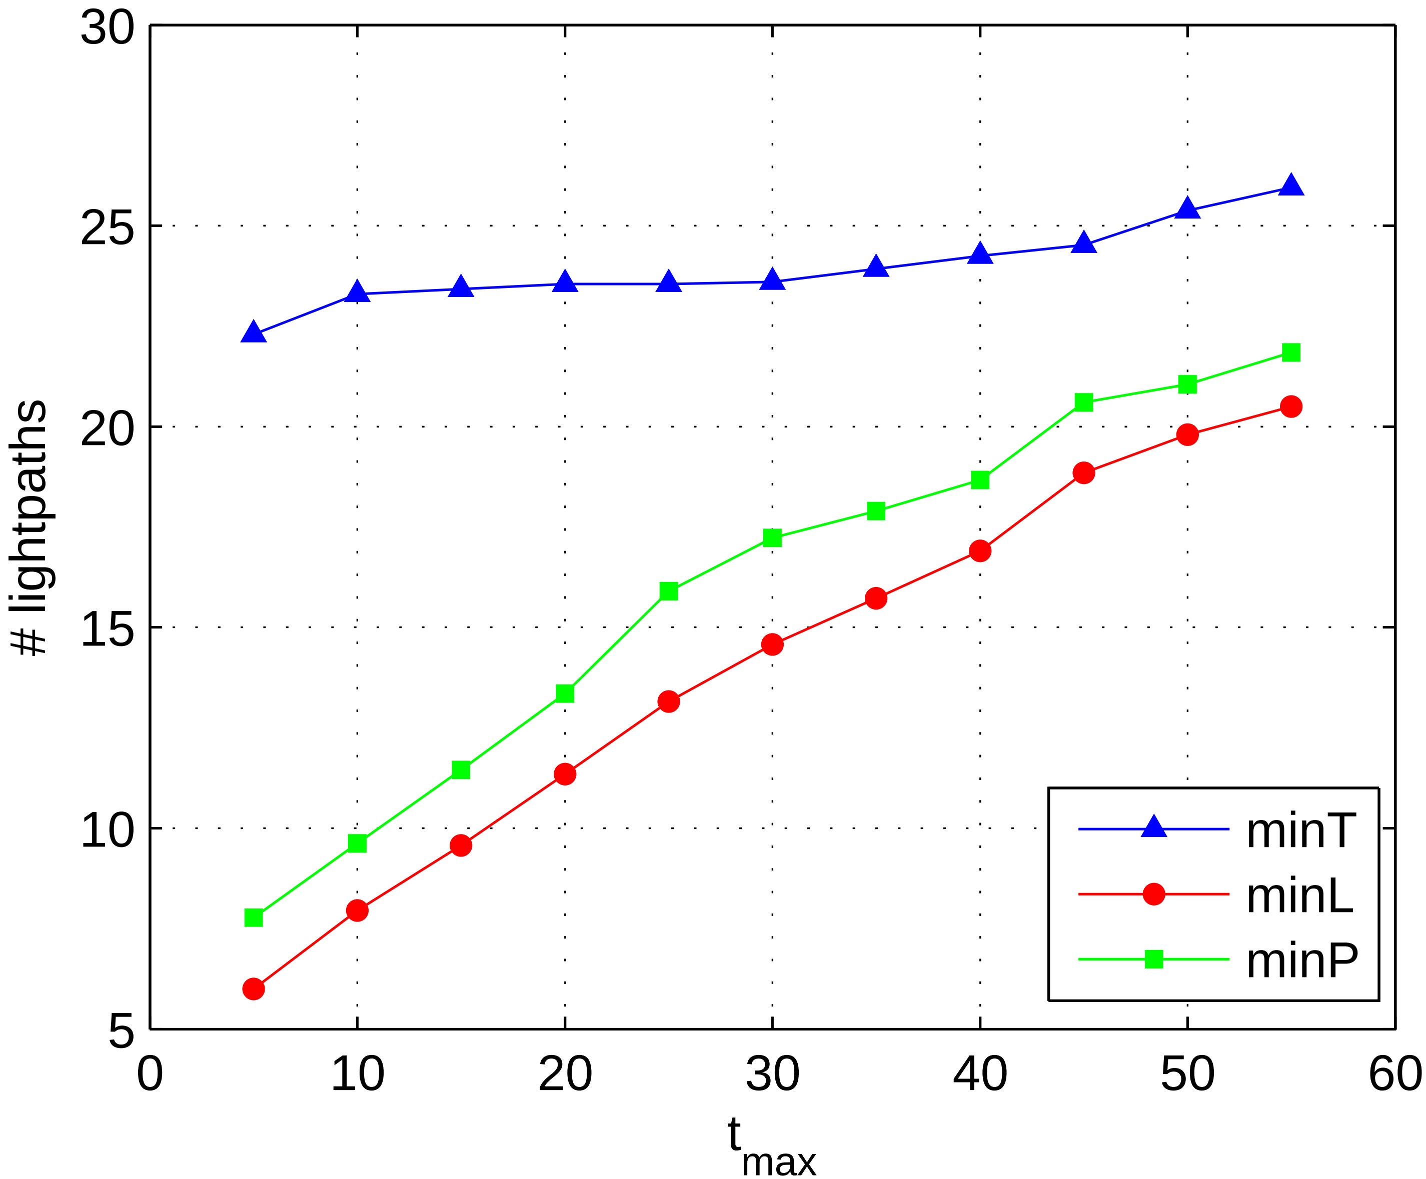

Results: Amount of Electronically Switched Traffic

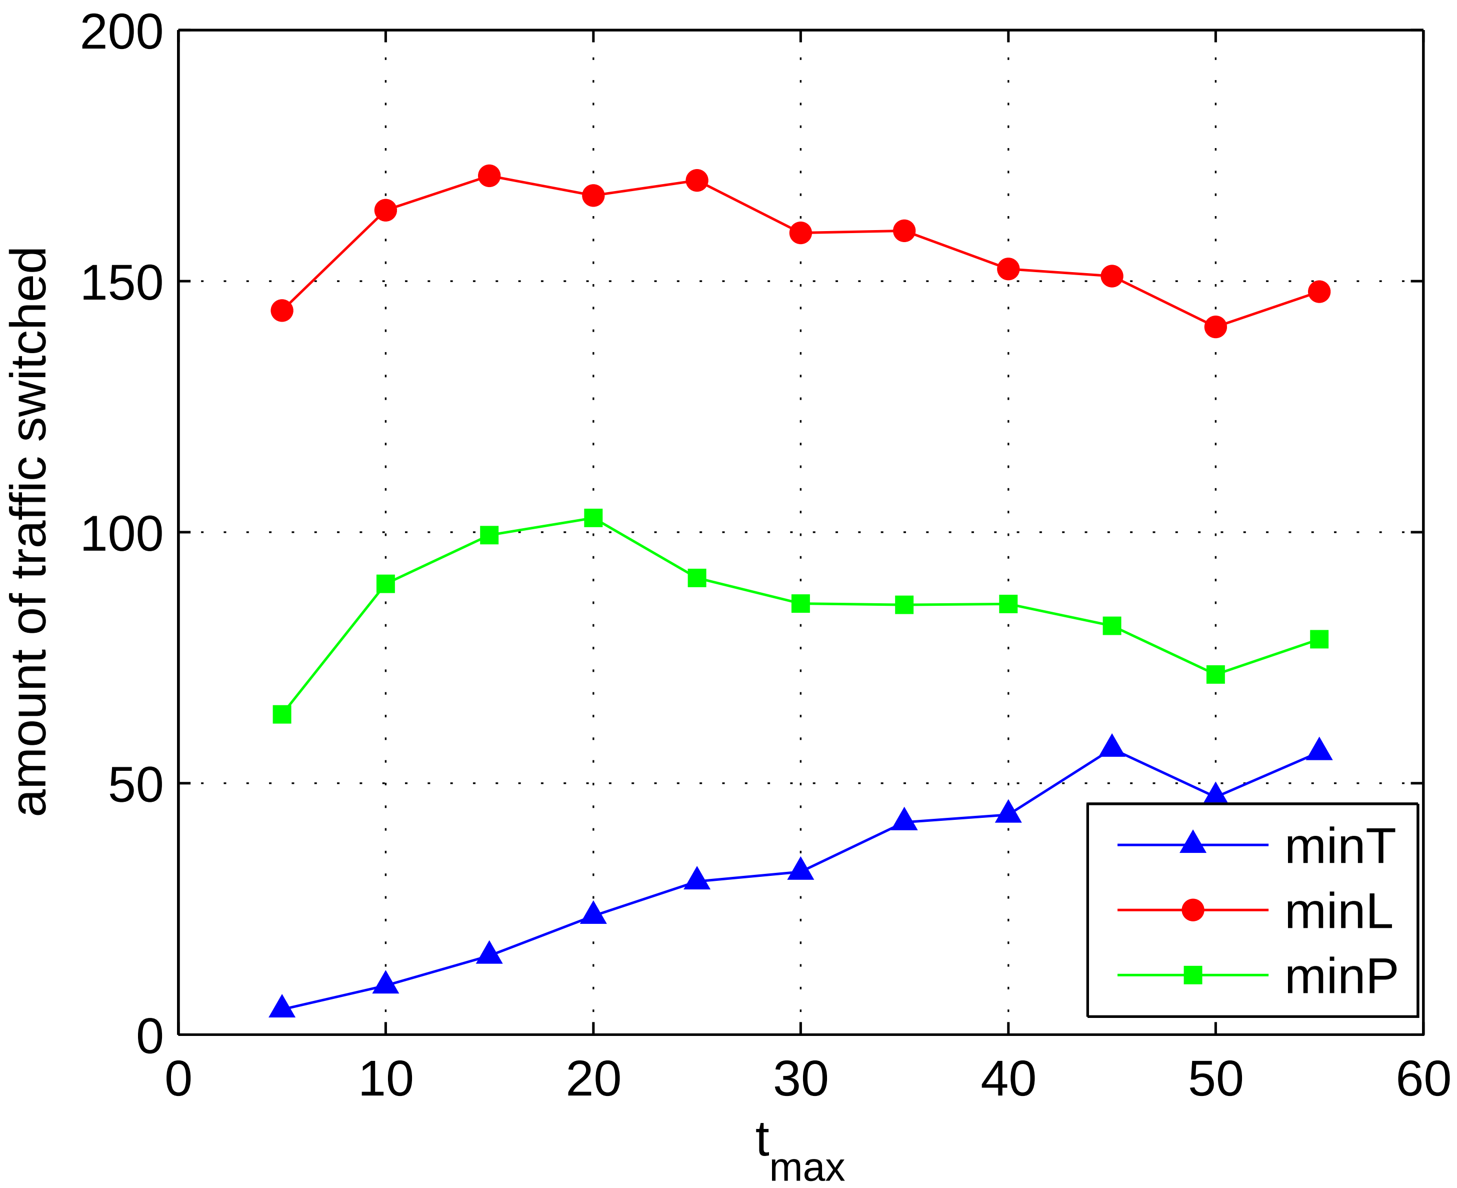

Results: Relative Power Consumption

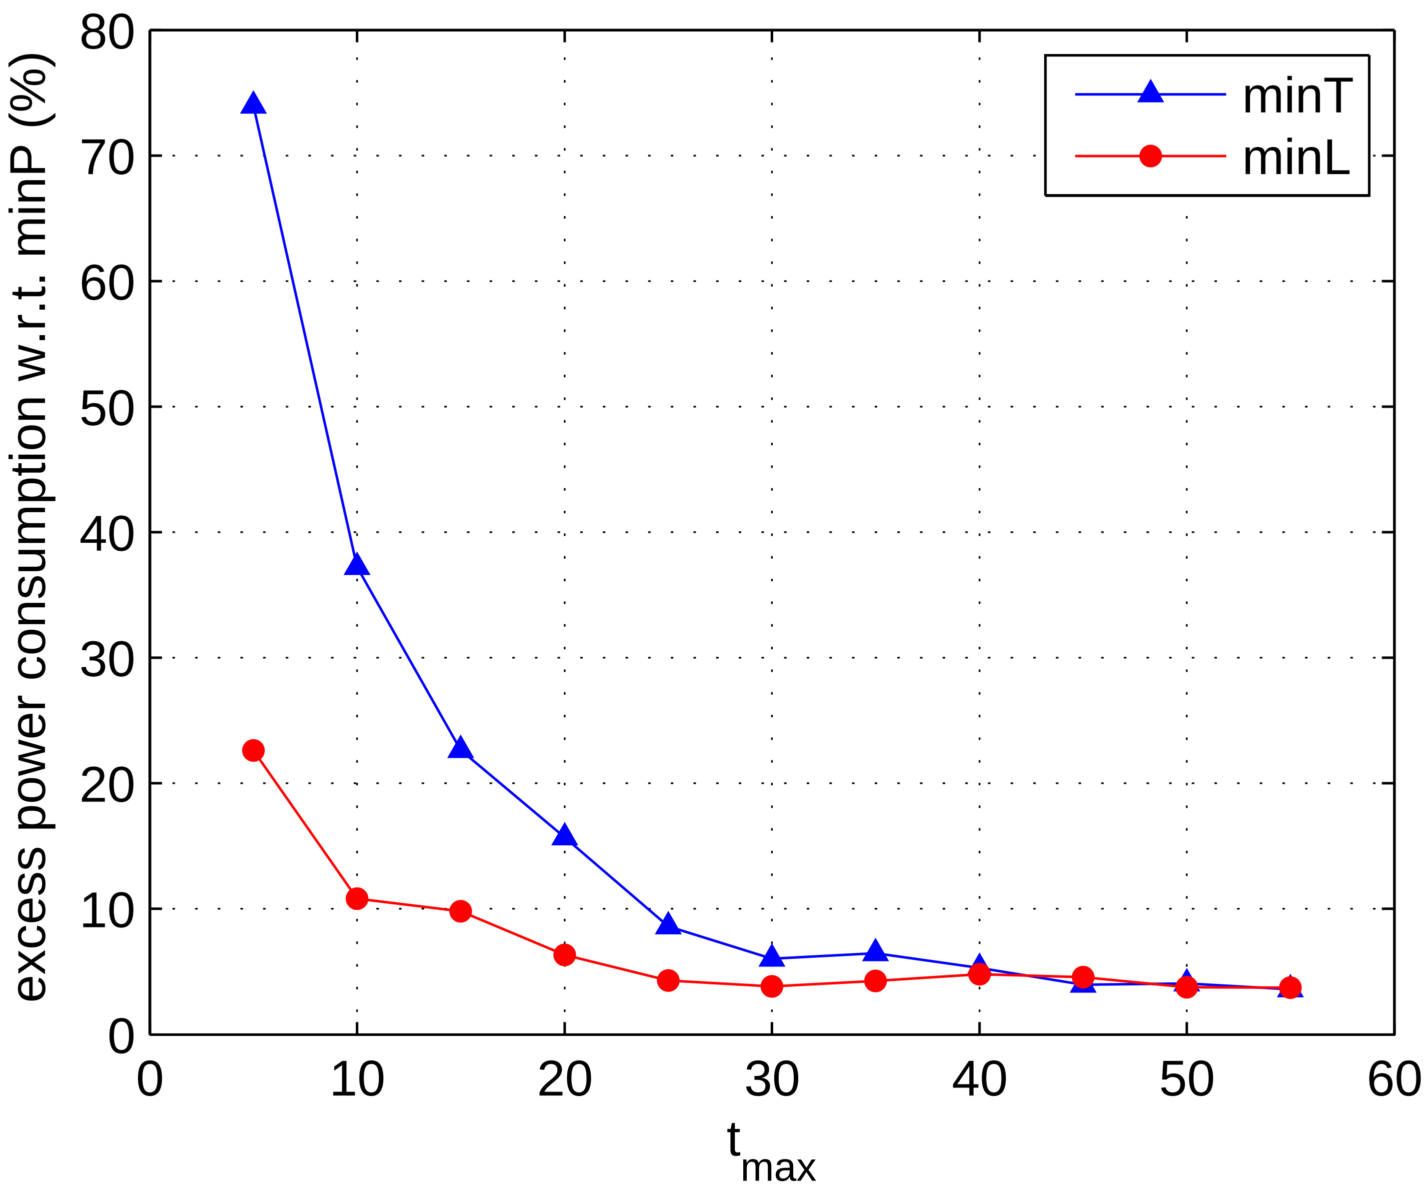

Summary and Discussion

- Power-aware design may lead to significant energy savings even for small networks
- The benefits are expected to increase with the network size
- Challenges:
  - existing ILPs do not scale to realistic networks
  - performance of heuristics difficult to characterize

Routing and Wavelength Assignment (RWA)

- Fundamental control problem in optical networks
- Objective: for each connection request determine a lightpath, i.e.,
  - a path through the network, and
  - a wavelength
- Two variants:
  1. online RWA: connection requests arrive/depart dynamically
  2. static RWA: a set of traffic demands to be established simultaneously

Static RWA

- Input:
  - network topology graph $G = (V, E)$
  - traffic demand matrix $T = [t_{sd}]$
- Objective:
  - minRWA: establish all demands with the minimum # of λs
  - maxRWA: maximize established demands for a given # of λs
- Constraints:
  - wavelength continuity: each lightpath uses the same λ along path
  - distinct wavelength: lightpaths using the same link assigned distinct λs
- NP-hard problem (both variants)

RWA Example

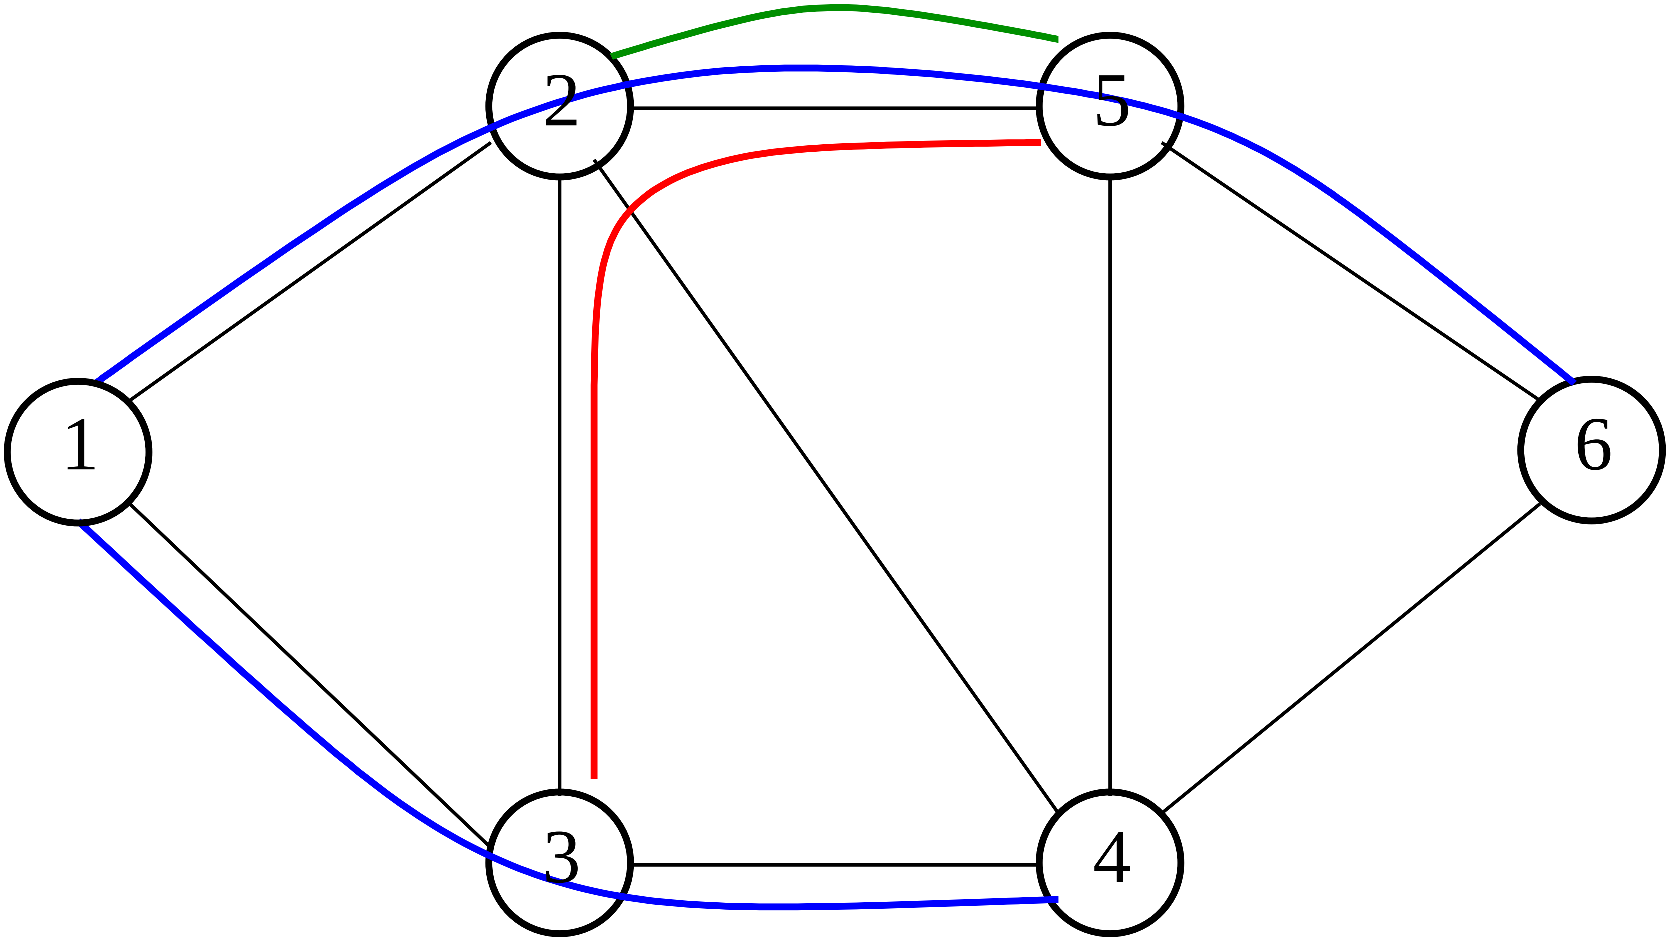

RWA: Symmetry

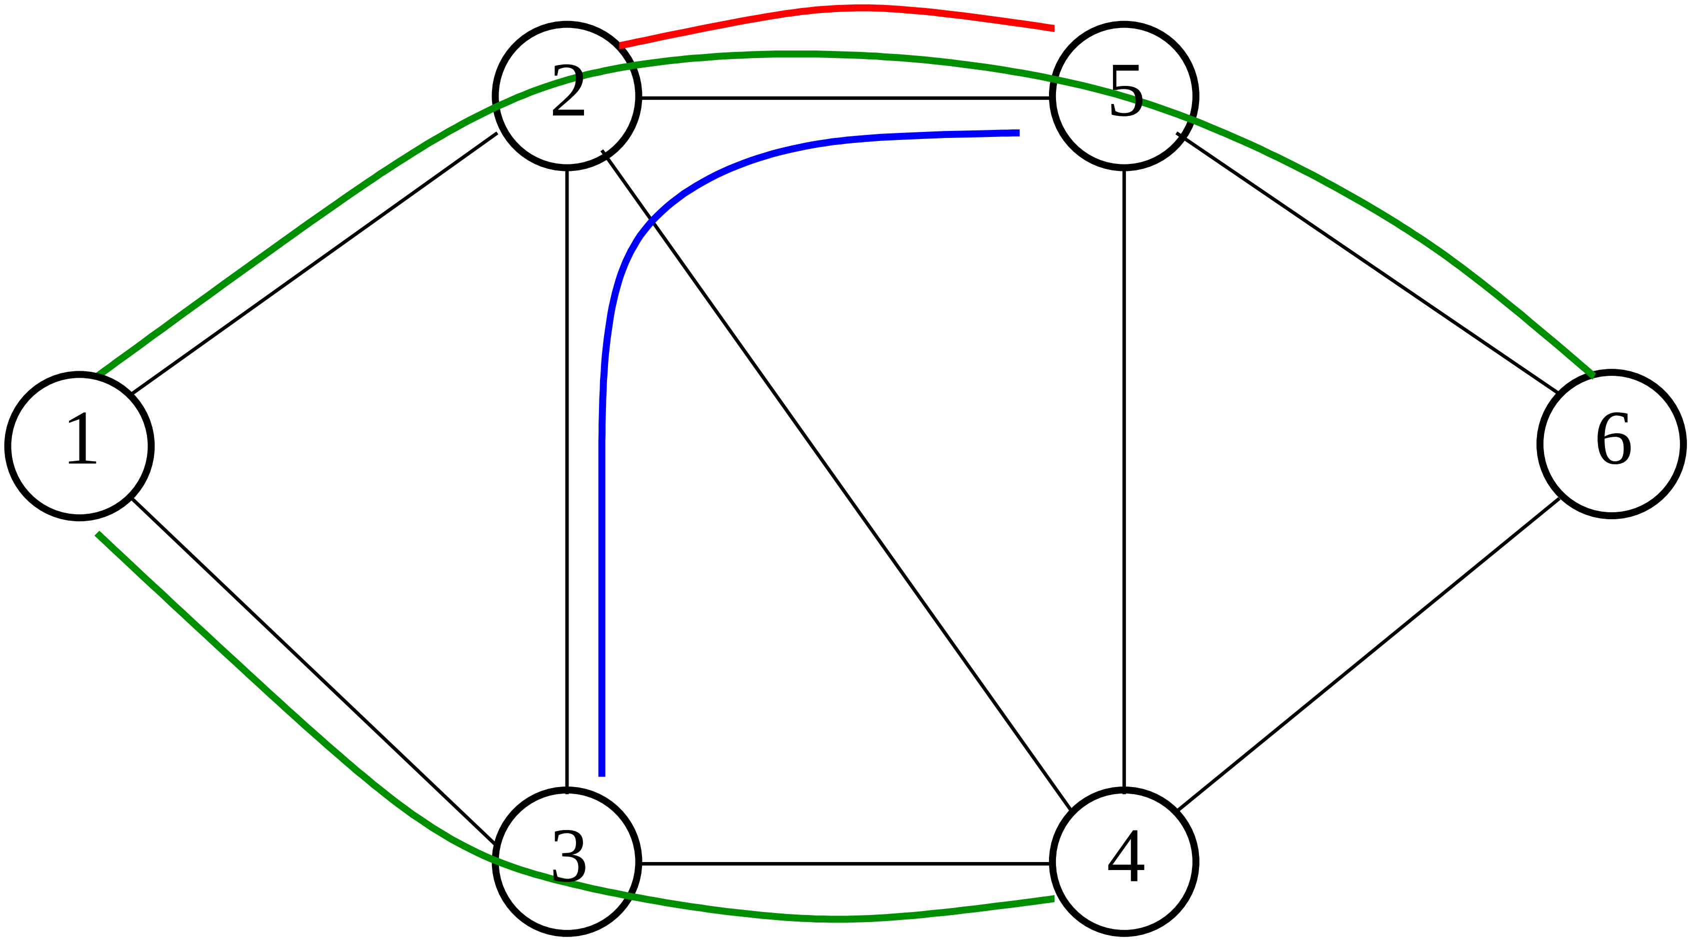

Link ILP Formulation

- Nodes/links are entities of interest
- Focus on traffic demand to and from nodes, on links

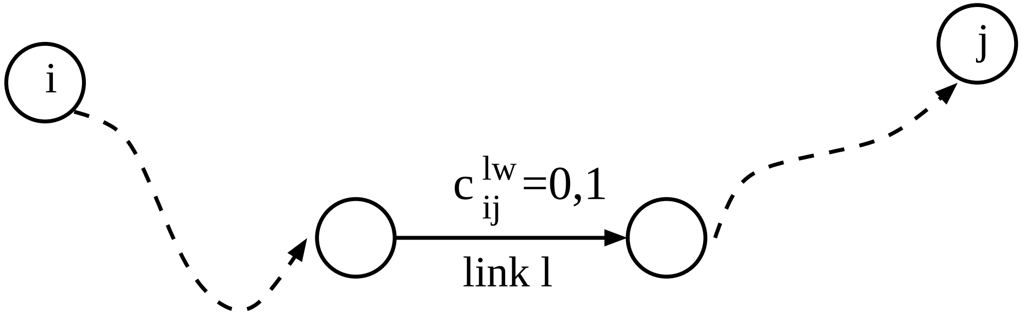

- Bridging variable: demand between nodes on links

Path ILP Formulation

- Nodes/paths are entities of interest
- Demand is still between nodes
- For each given demand node pair, list all paths
  $\longrightarrow$ typically, a subset of all paths

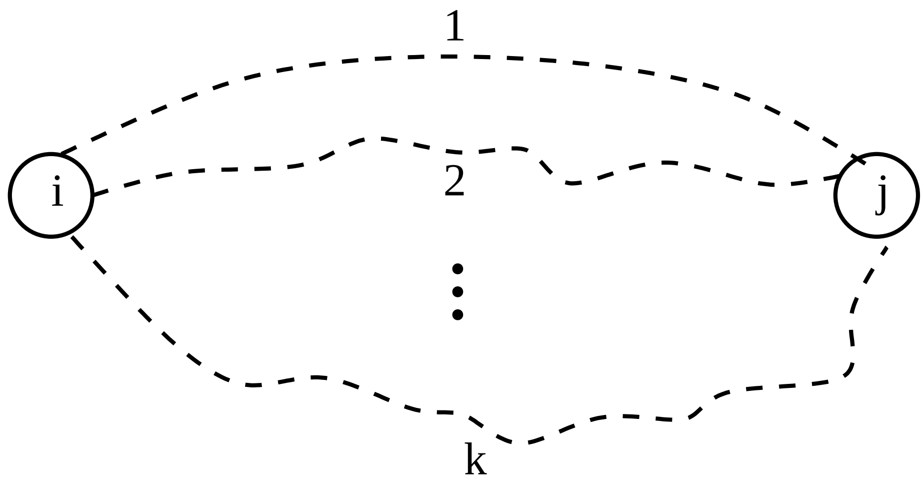

- assign variable to path traffic flow $\longrightarrow$ implicitly identifies demand
- for each link, sum up path flow variables
  $\longrightarrow$ constrain with capacities

RWA As Graph Coloring

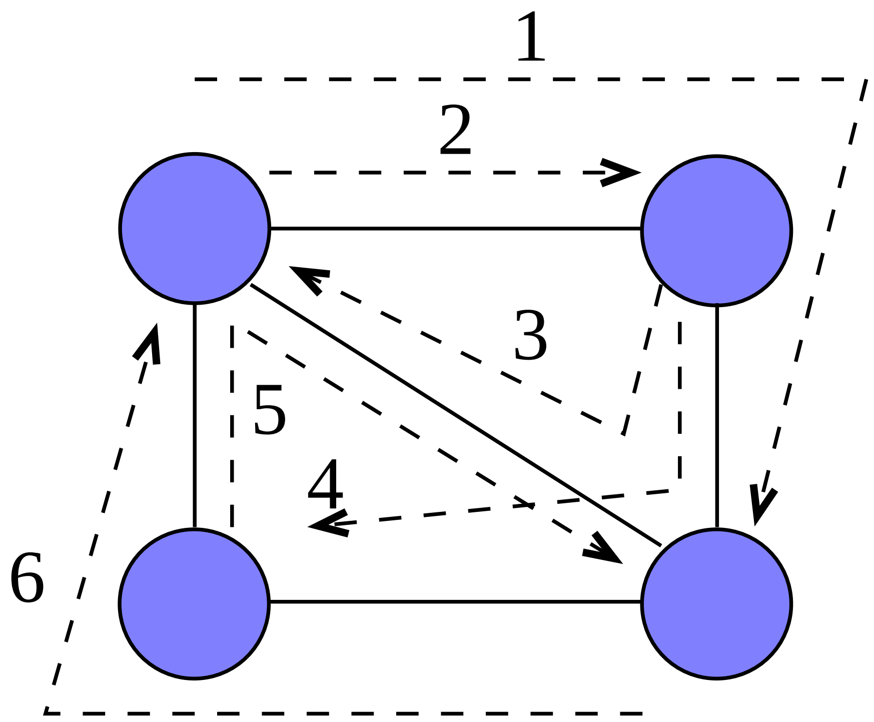

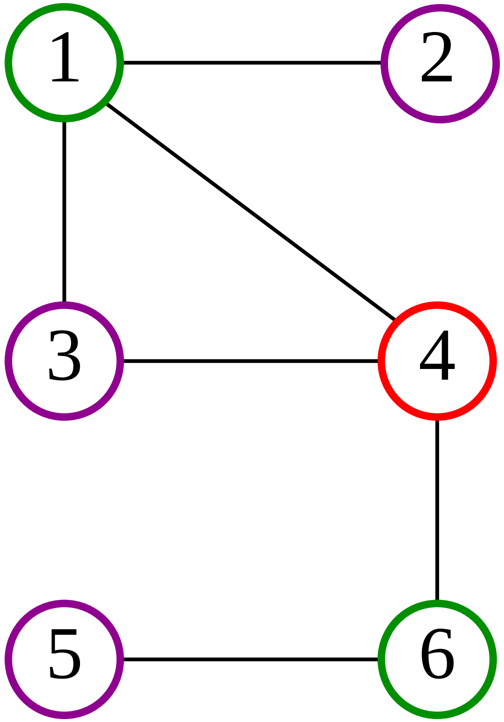

Maximal Independent Sets

- Independent set: a set of vertices in a graph no two of which are adjacent
- Maximal independent set: not a subset of any other independent set

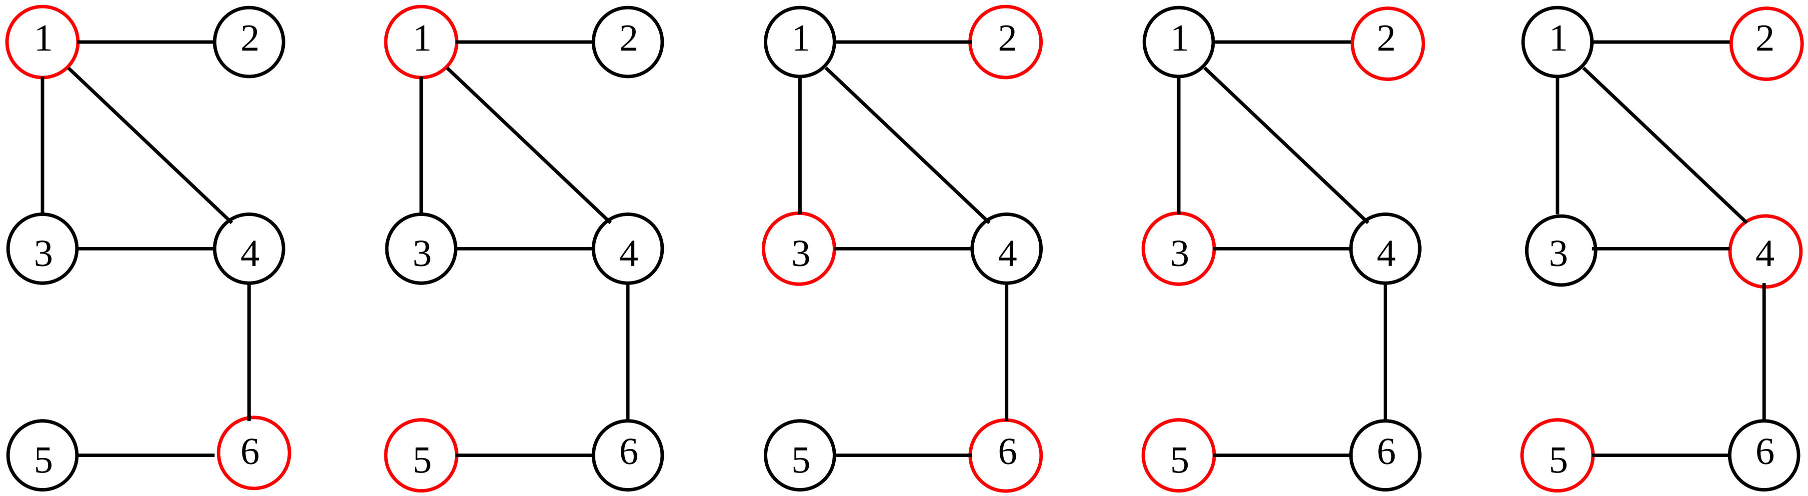

MIS ILP Formulation

- Precompute k paths for each source-destination pair
- Create the path graph G_p:
  - each node in G_p corresponds to a path in the original network
  - two nodes connected in G_p if corresponding paths share a link
- Enumerate the MISs of G_p
- Set up ILP to assign wavelengths to each MIS

Comparison

- Link ILP formulation
  - $O(N^4W)$ variables
  - $O(N^3W)$ constraints
  - symmetry with respect to λ permutations
- Path ILP formulation
  - $O(N^2W)$ variables
  - $O(N^2W)$ constraints
  - symmetry with respect to λ permutations
- MIS ILP formulation
  - $O(3^{N^2/3})$ variables
  - $O(N^2)$ constraints
  - no symmetry
  - size independent of W $\longrightarrow$ future-proof

RWA in Ring Networks

- Vast parts of network infrastructure based on SONET/SDH rings
- AT&T operates $\approx$ 6,700 rings in North America

  $\longrightarrow$ optimal solutions for rings important for foreseeable future
- Max size of SONET ring: 16 nodes
- Operators have started transition to mesh networks $\longrightarrow$ next ...

Running Time Results, $W = 120$

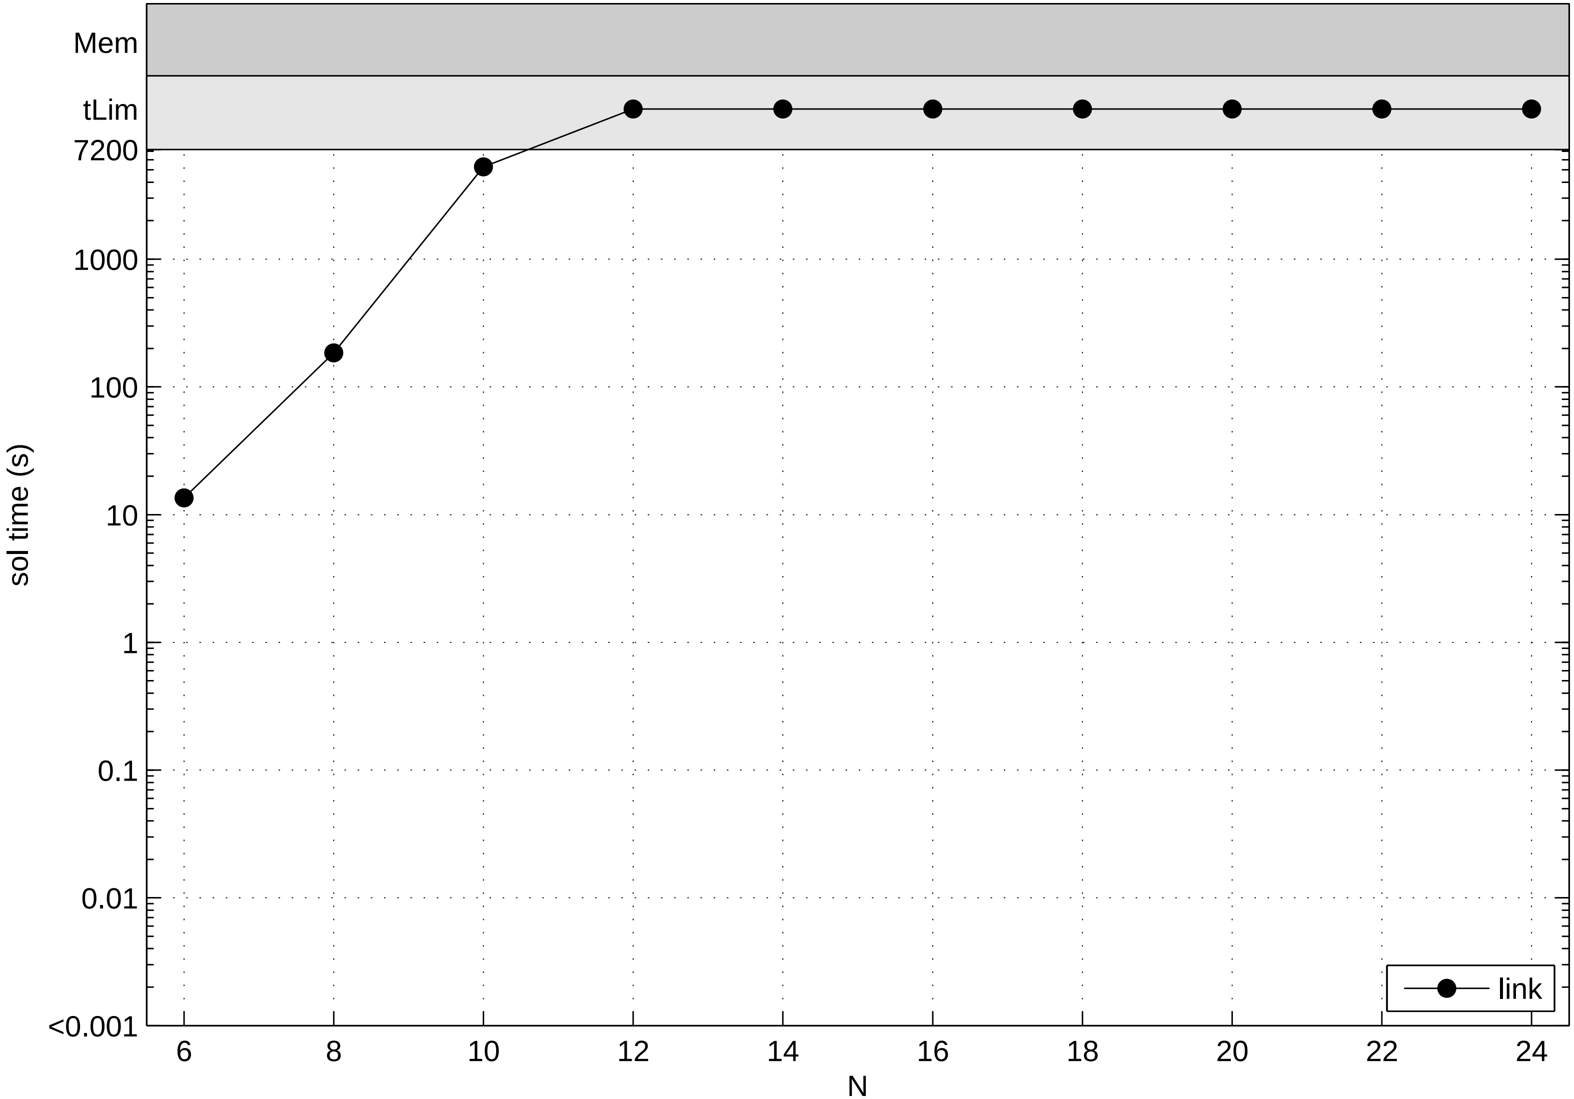

Running Time Results, $W = 120$

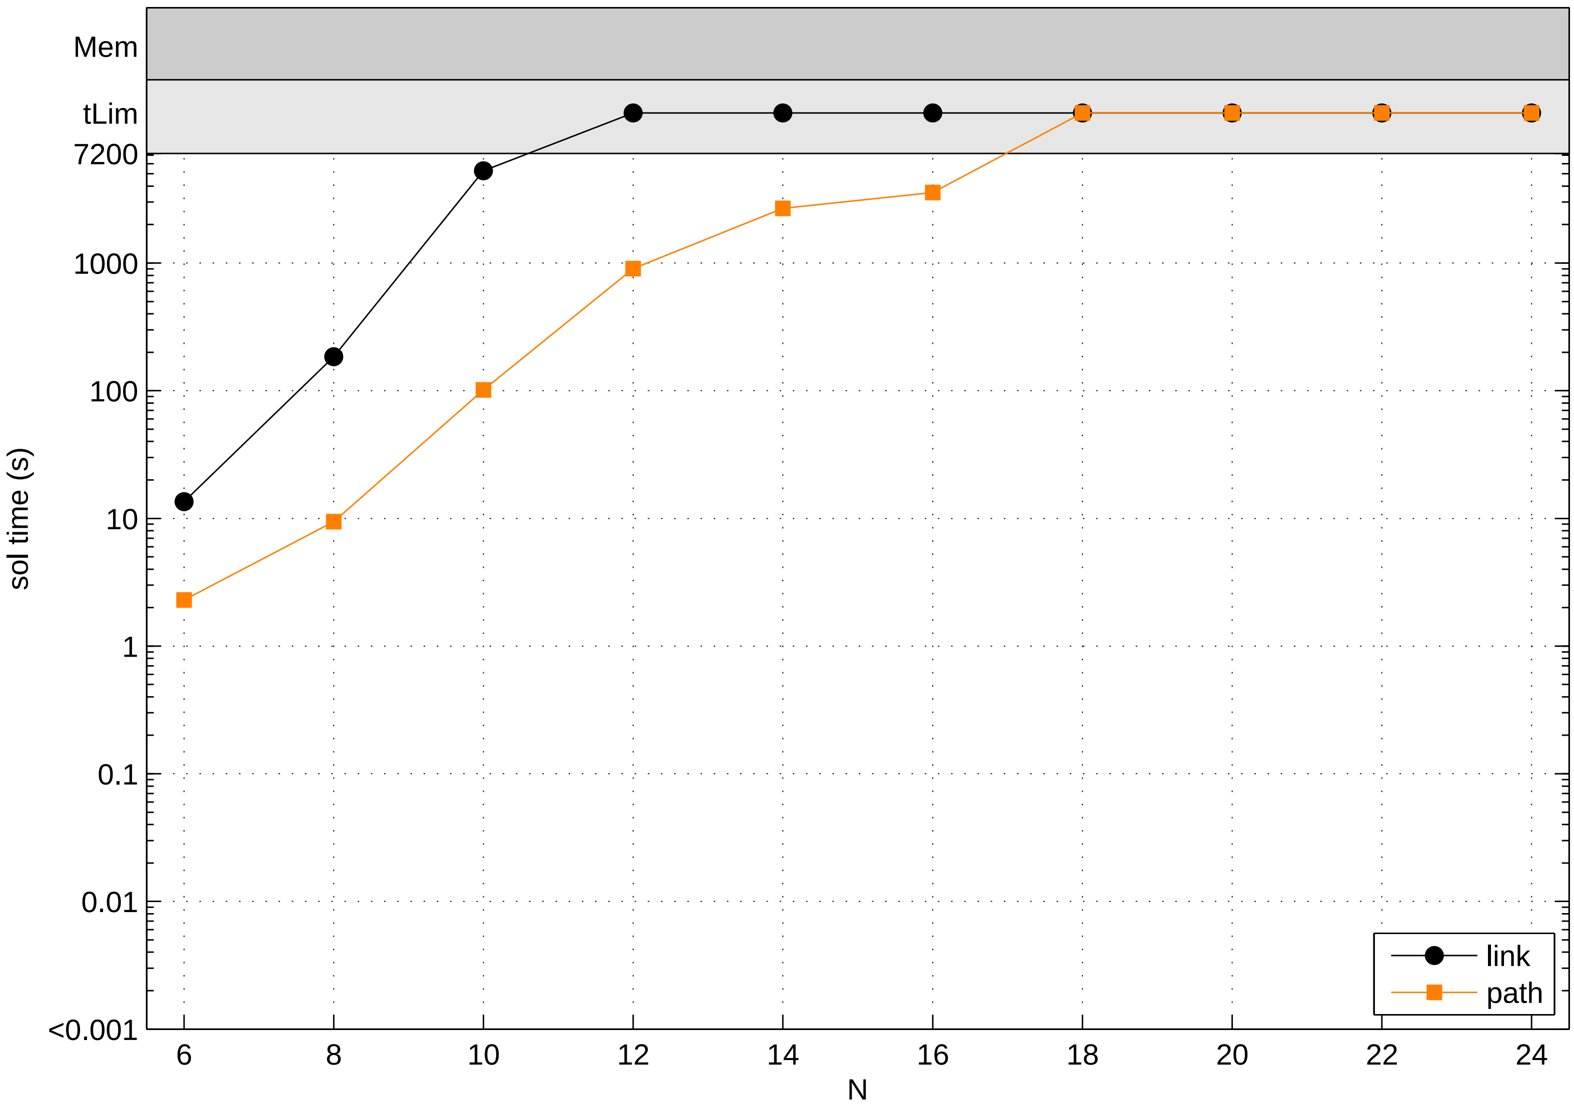

Running Time Results, $W = 120$

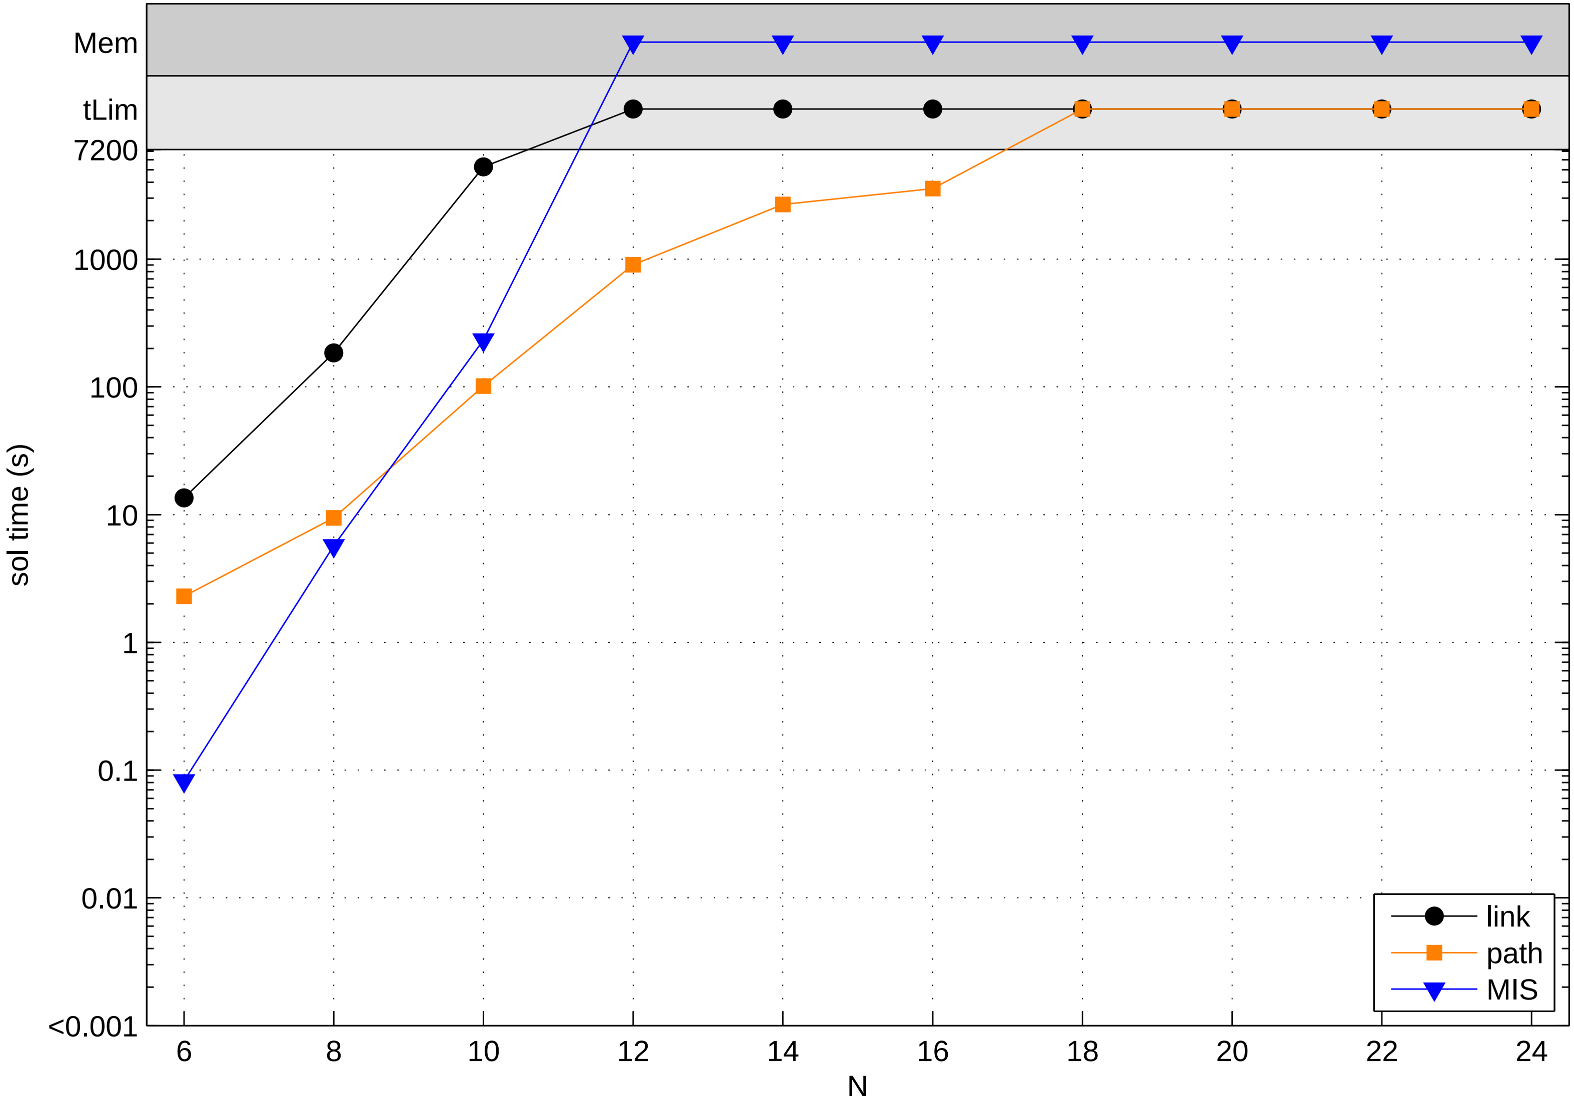

MIS Decomposition for Rings: MISD-2

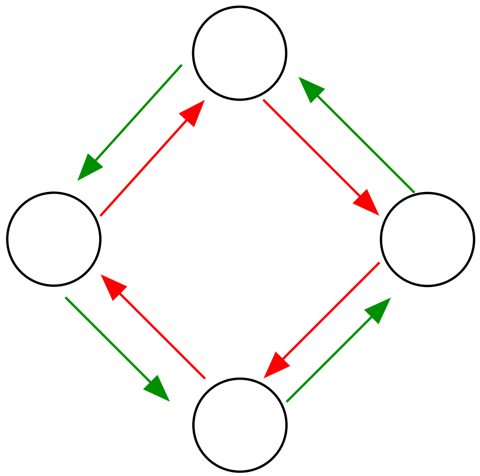

- Clockwise paths do not intersect with counter-clockwise paths:

$$G_p = G_p^{cw} \cup G_p^{ccw}$$

- M, M^{cw}, M^{ccw} : # of MISs of G_p, G_p^{cw}, G_p^{ccw}:

$$M^{cw} = M^{ccw} = \sqrt{M}$$

$\longrightarrow$ orders of magnitude decrease in # of variables/size of formulation

- Slight modifications to formulation

Further Decomposition: MISD-4

- Consider clockwise direction only
  $\longrightarrow$ similar steps for counter-clockwise
- Partition ring in two parts such that:

$$G_p^{cw} = G_p^{cw,0} \cup G_p^{cw,1} \cup G_p^{cw,core}$$

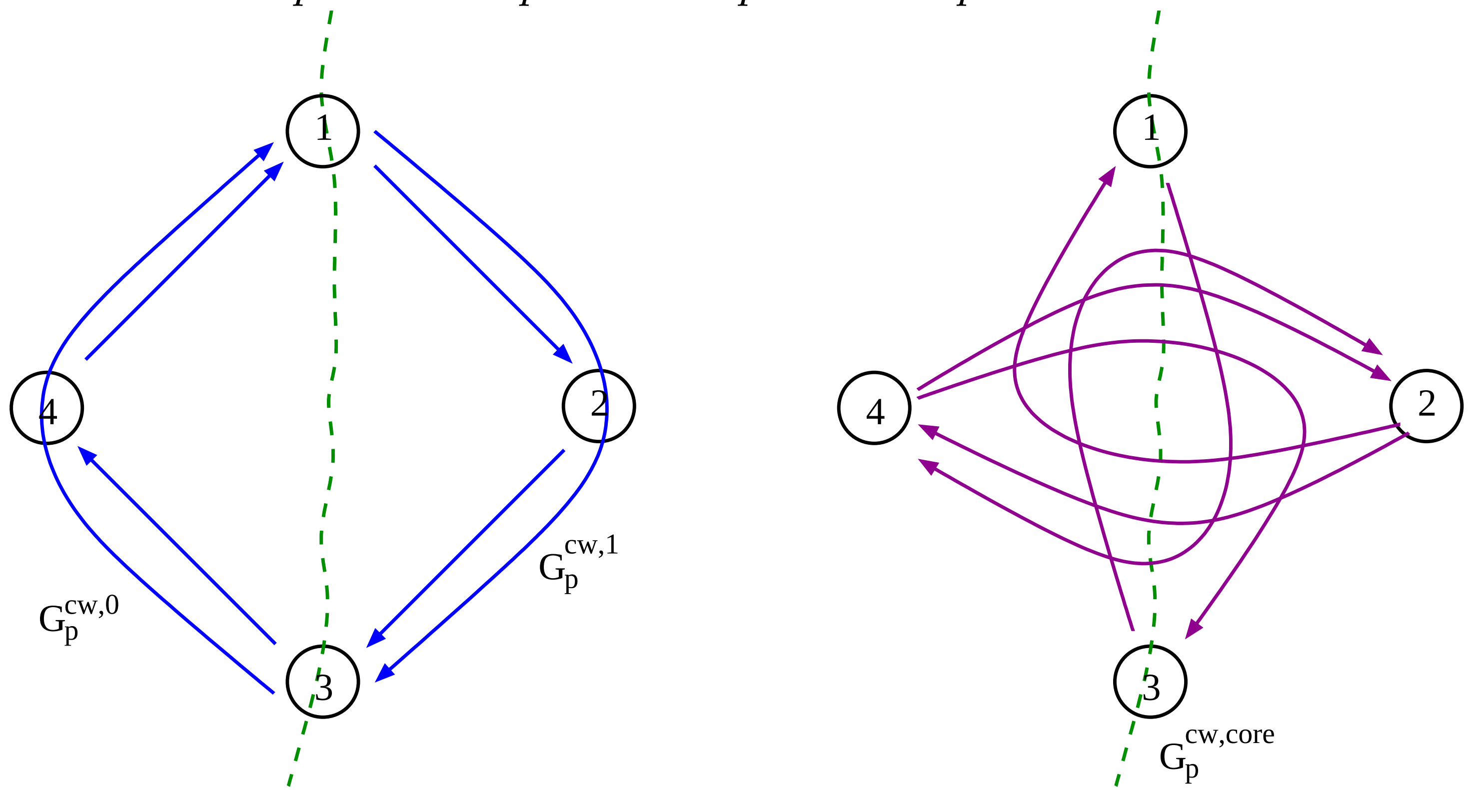

MISD-4 (cont'd)

- Express each MIS m of G_p^{cw} as:

$$m = m^0 \cup m^1 \cup q$$

- Modify the formulation appropriately
  - # MIS variables $\downarrow$
  - # constraints $\uparrow$
- Recursively partition the two ring parts to effect higher-order decompositions (MISD-8, MISD-16, . . .)

Results: # of MIS Variables

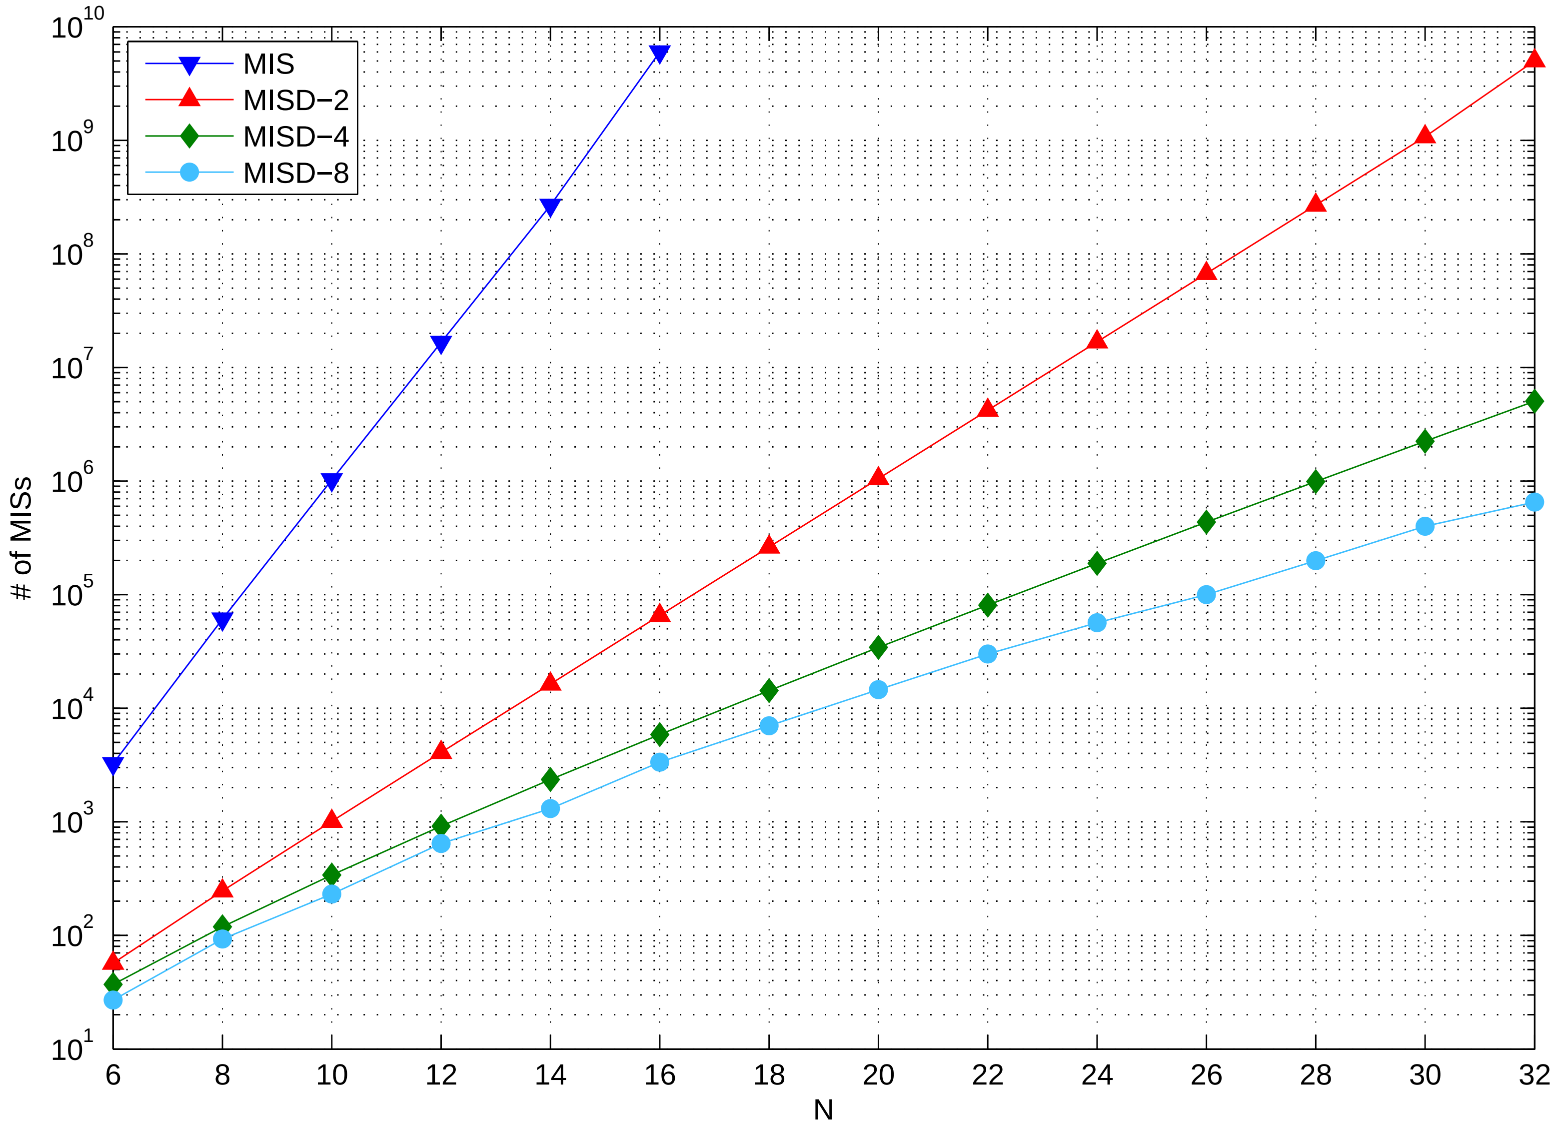

Running Time Results, $W = 120$

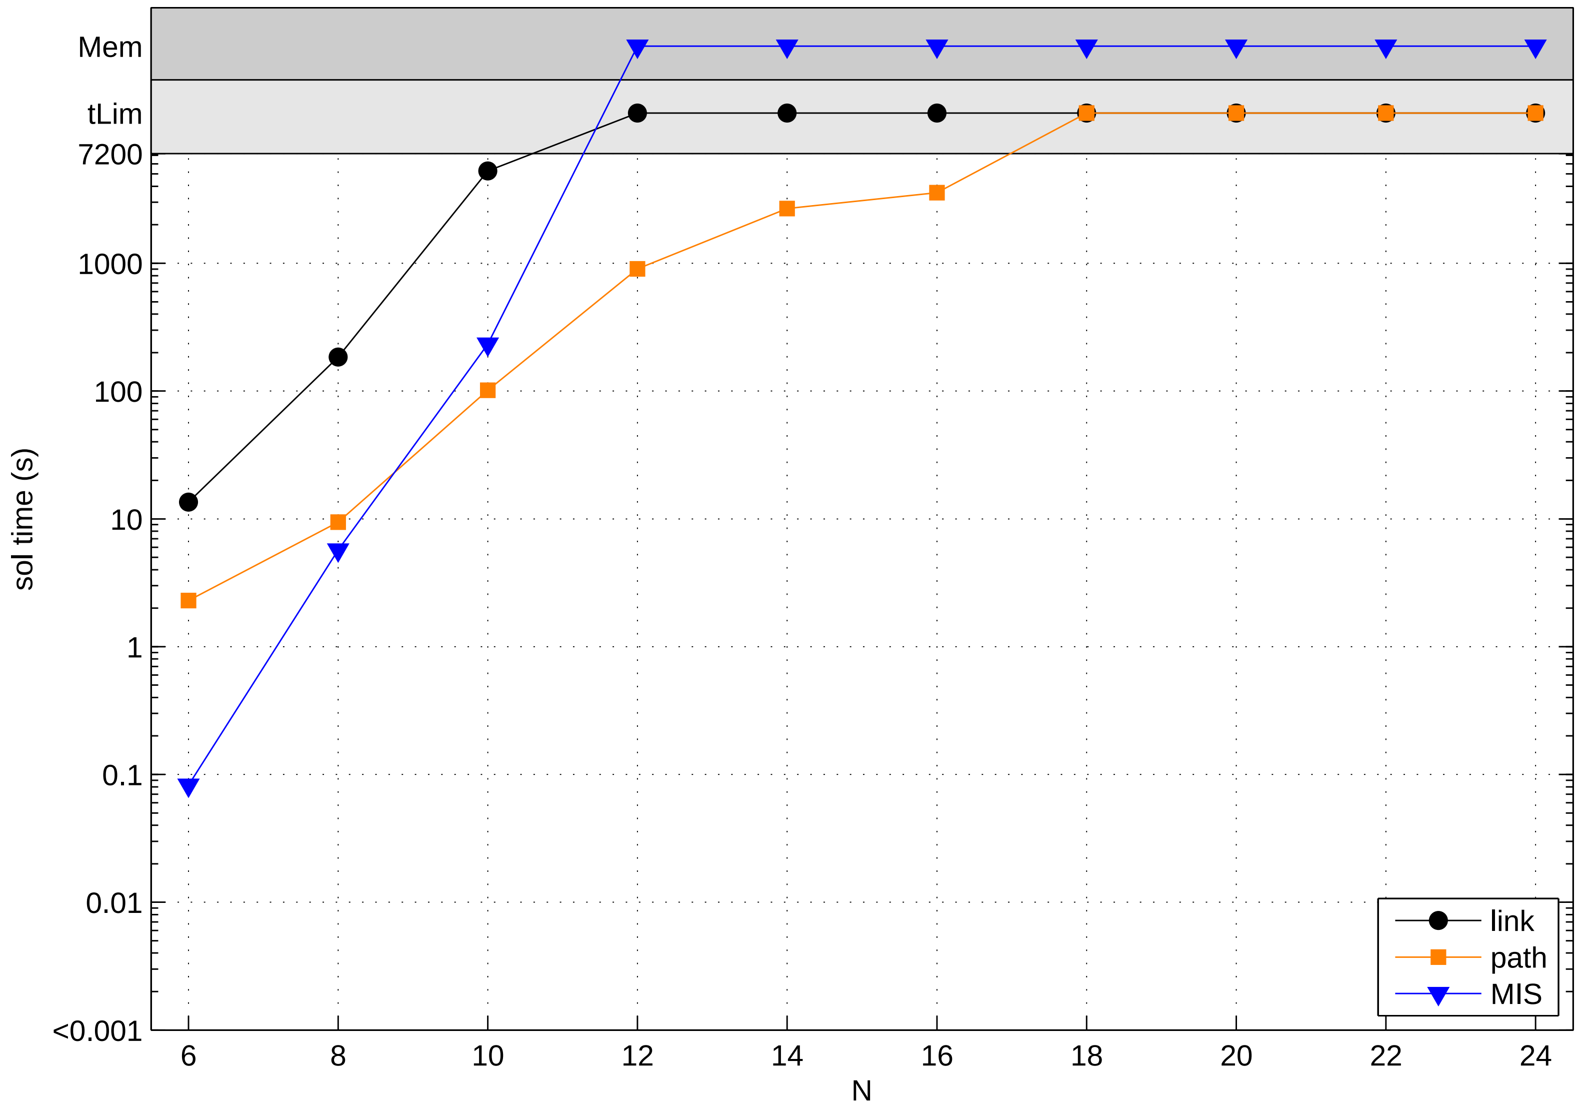

Running Time Results, $W = 120$

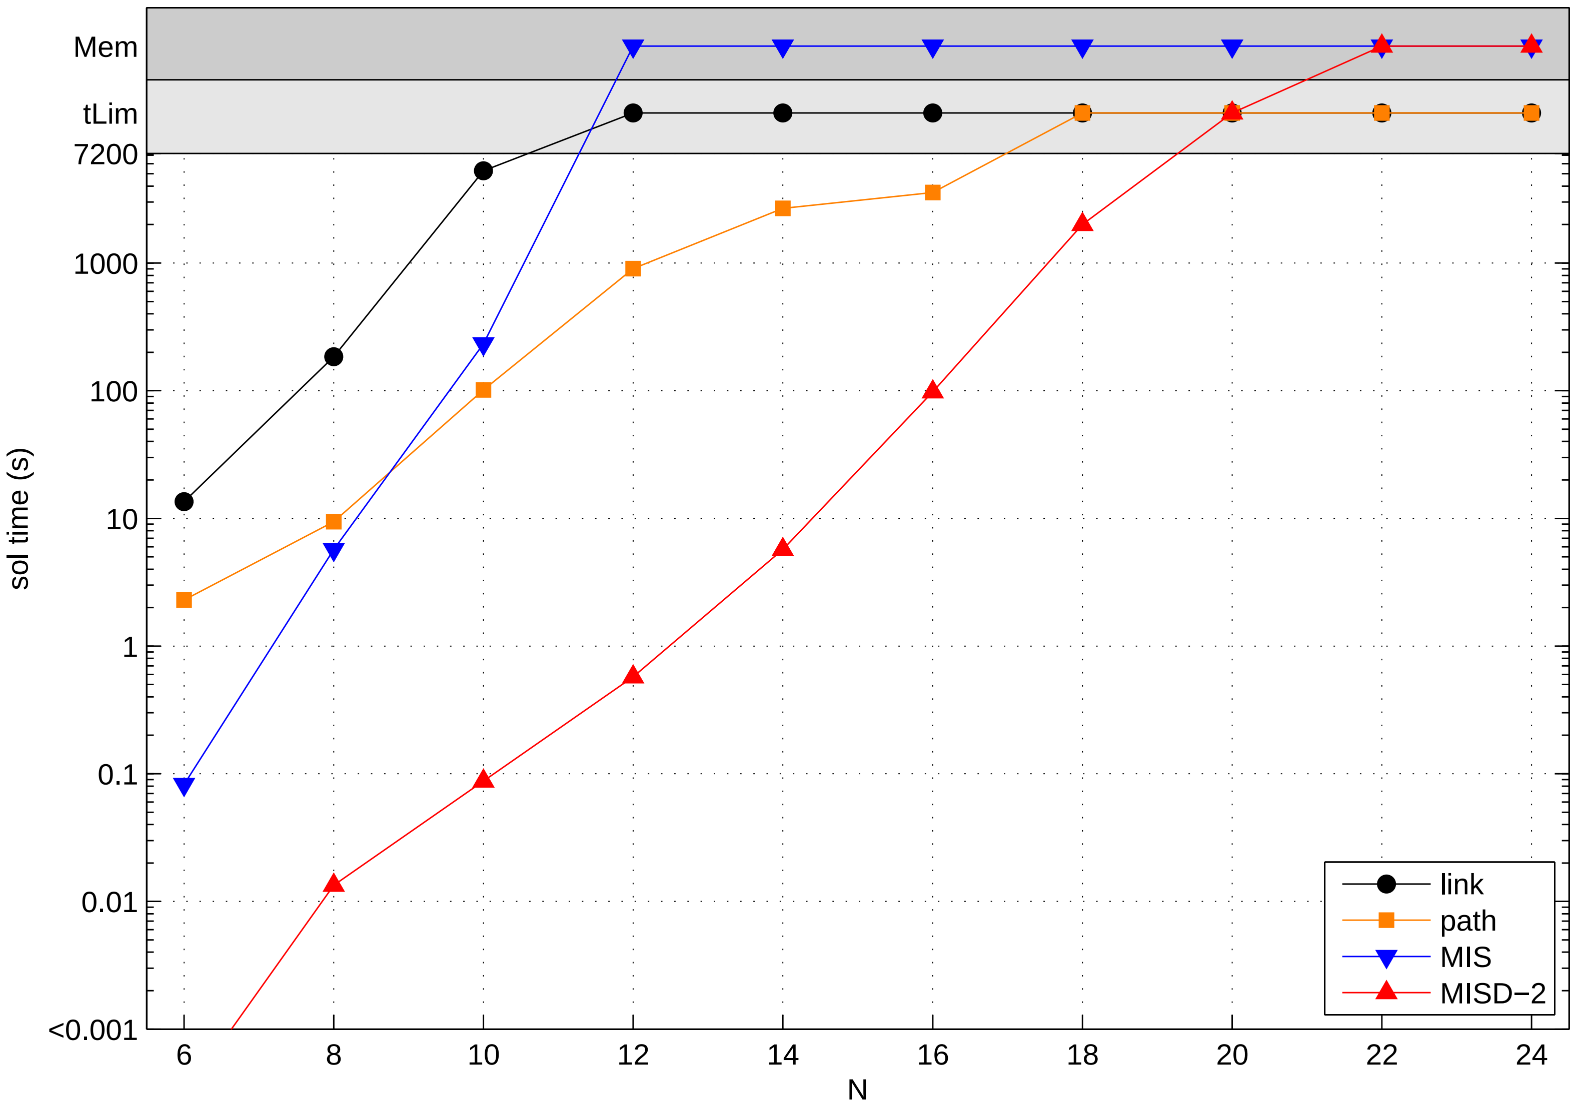

Running Time Results, $W = 120$

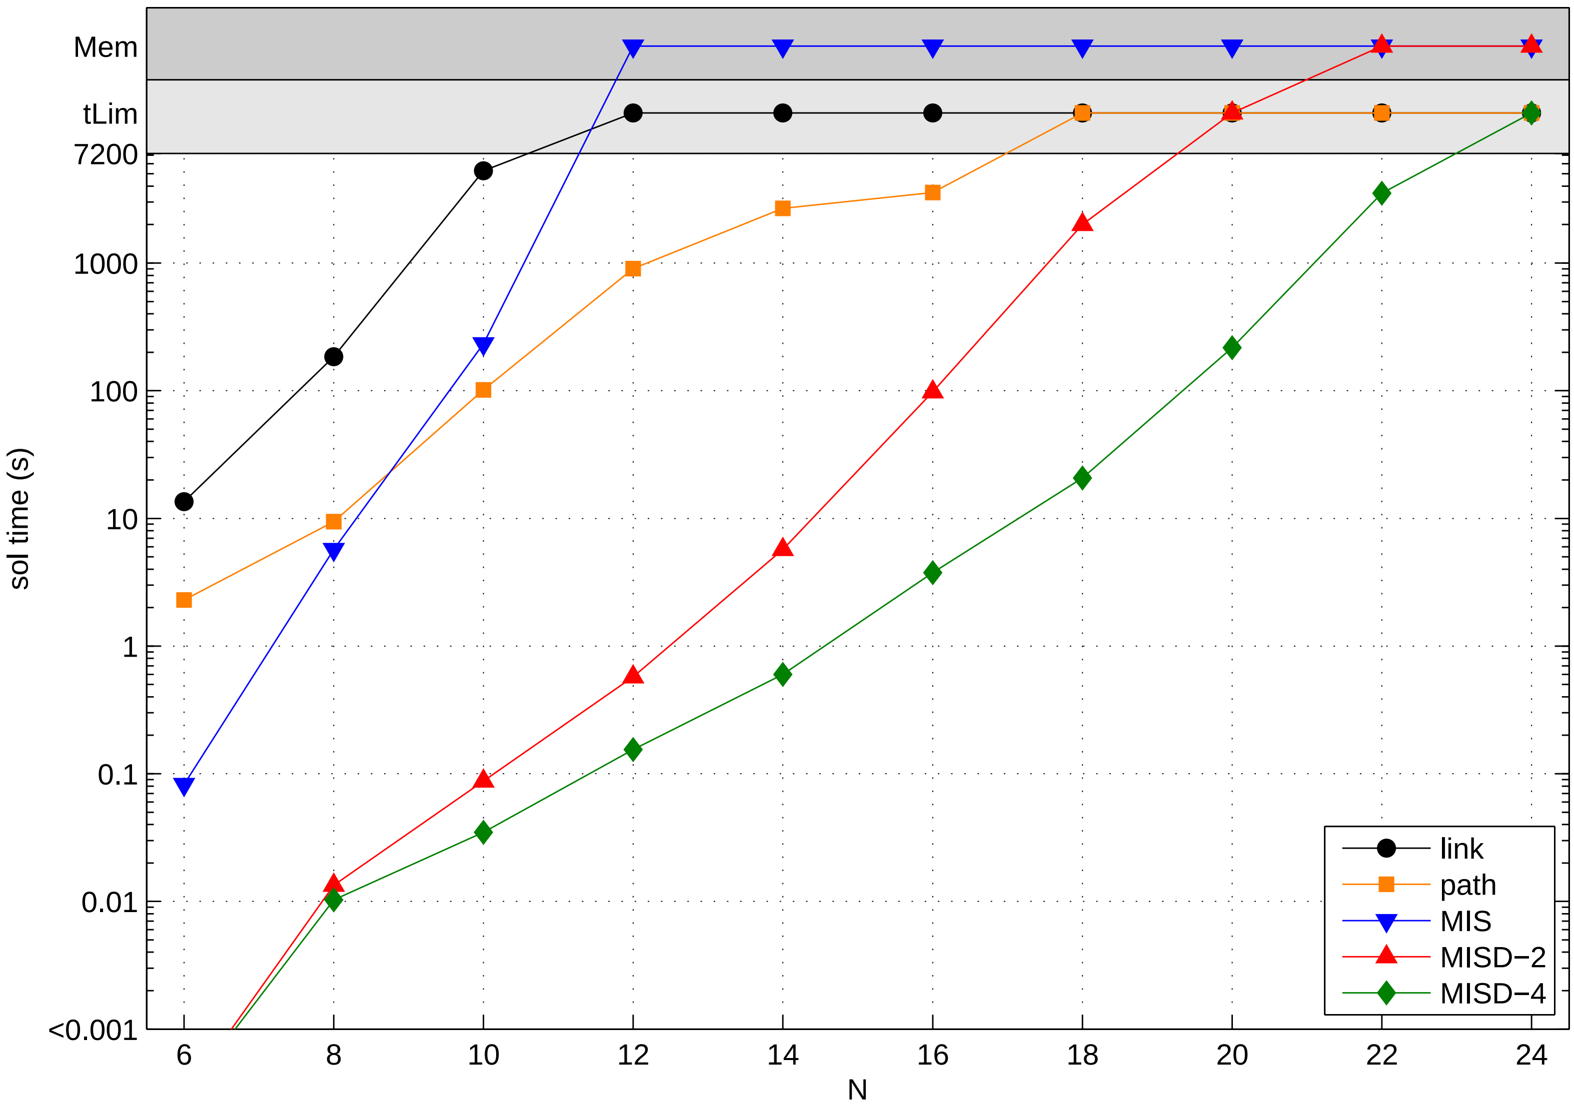

Running Time Results, $W = 120$

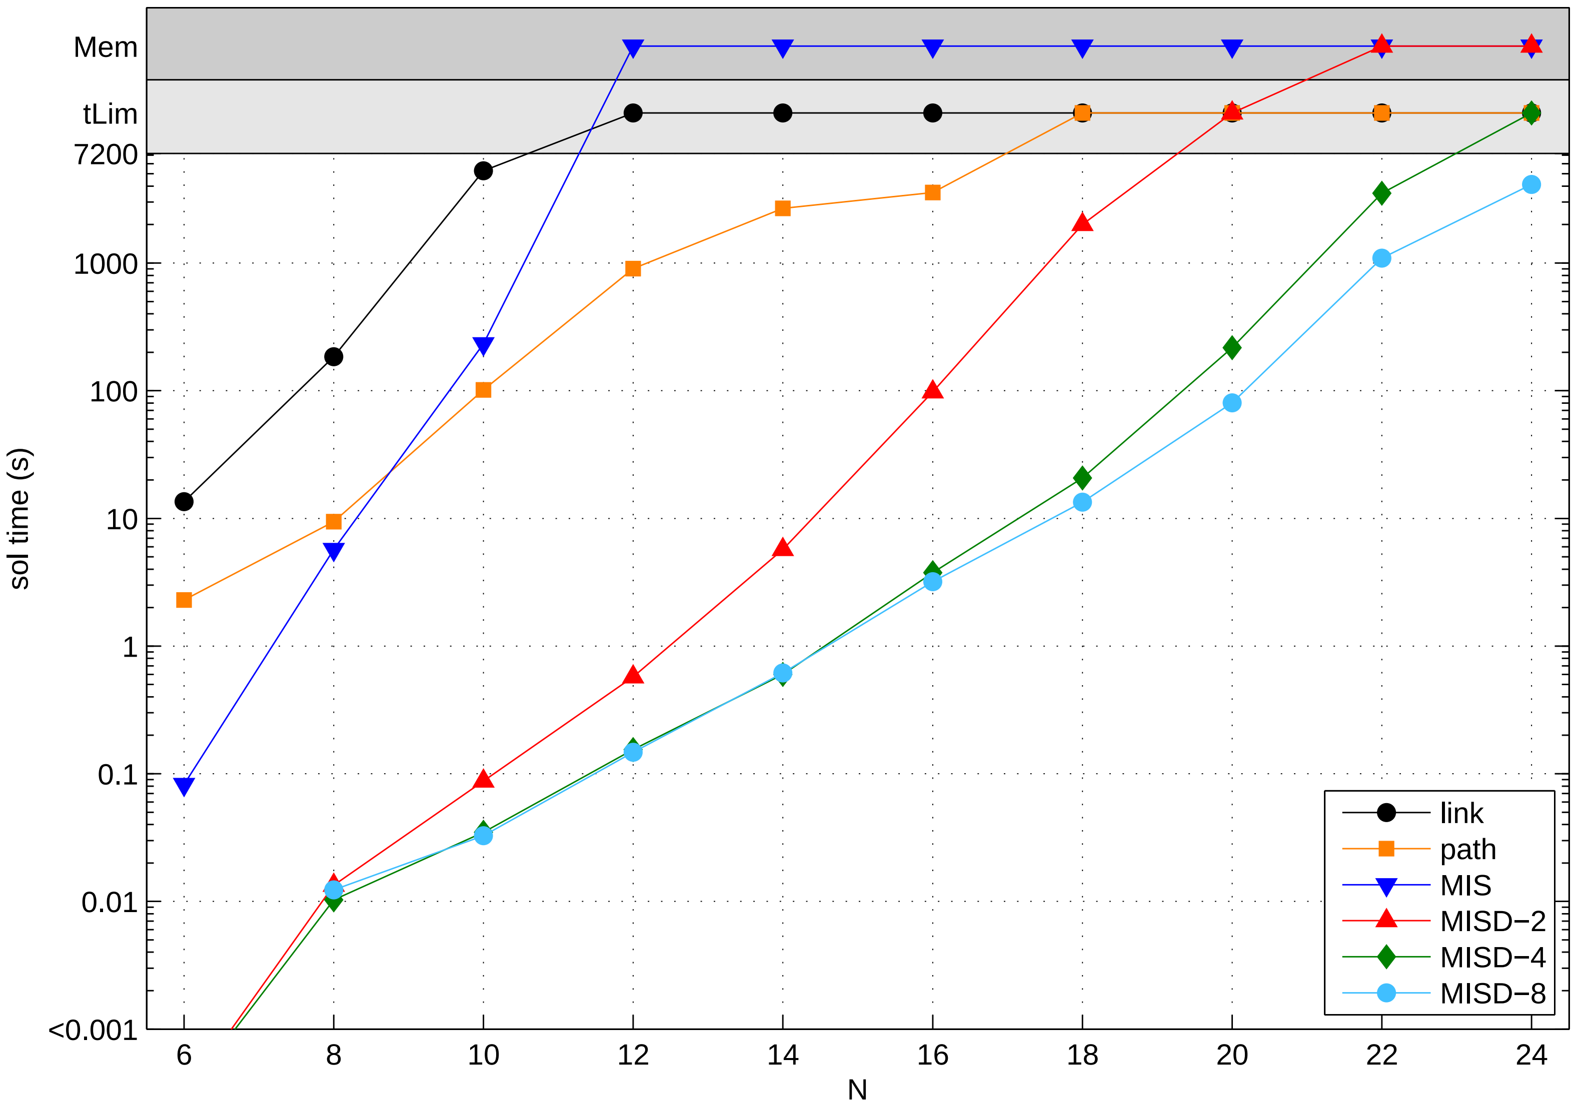

RWA in Mesh Networks

- MIS decomposition does not work
- Devised new exact decompositions for path formulation
- May solve efficiently 40-node networks

Running Time Results: Torus

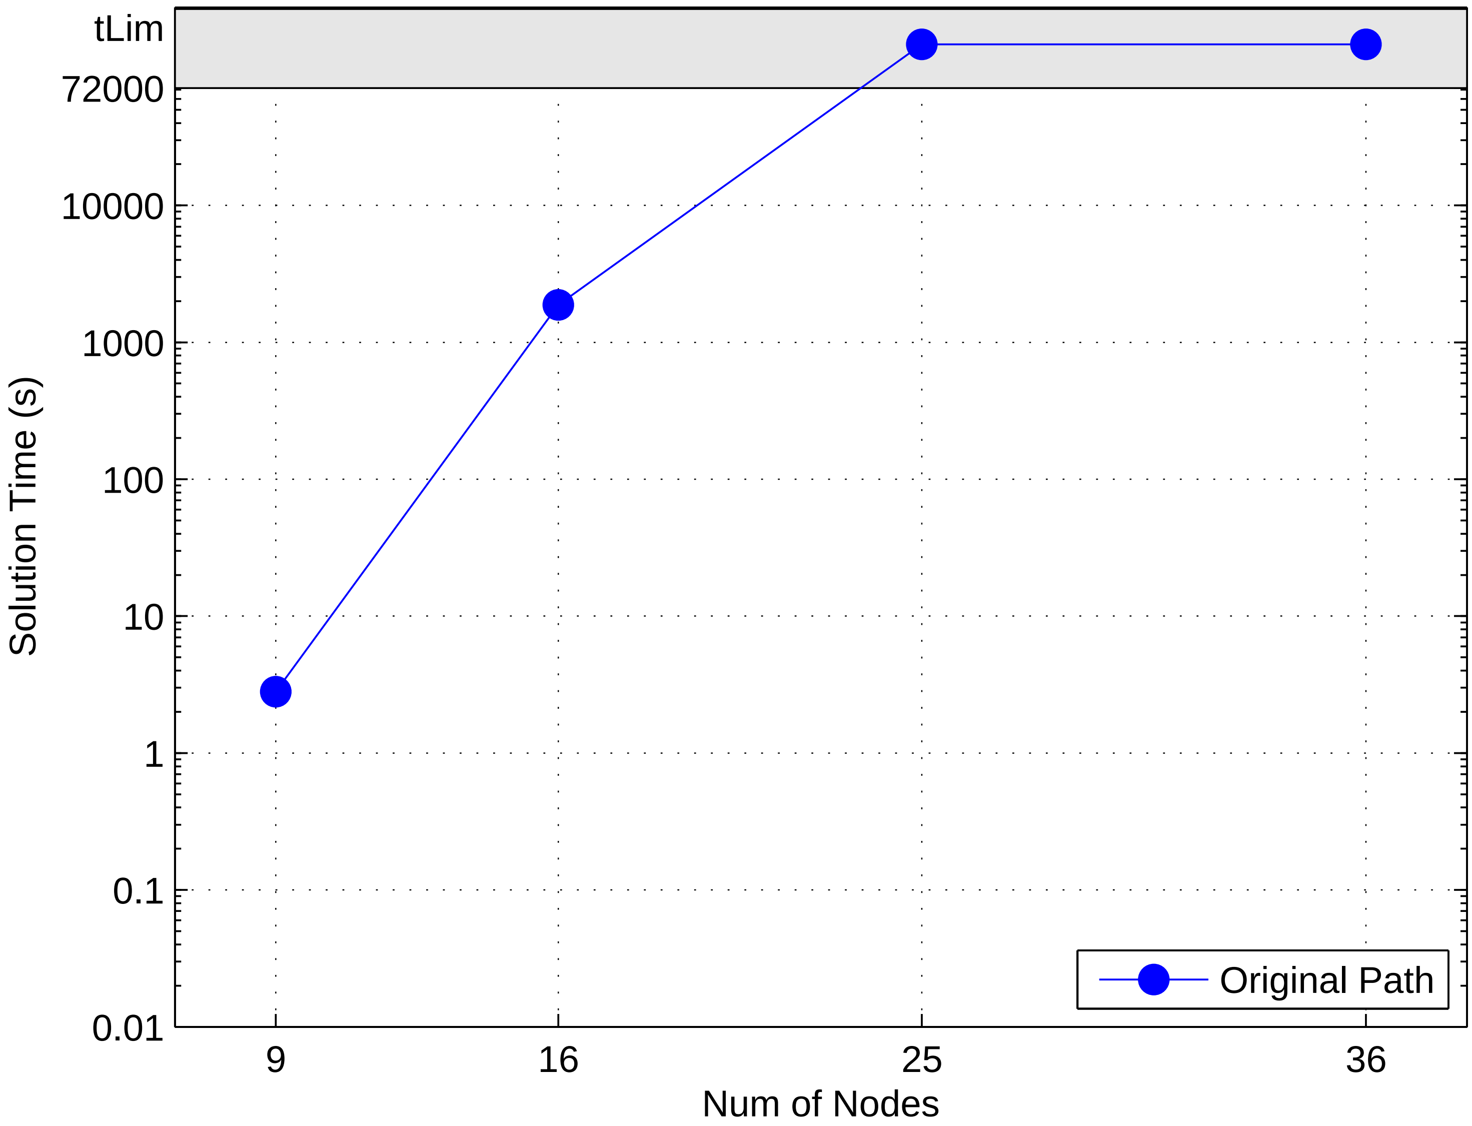

Running Time Results: Torus

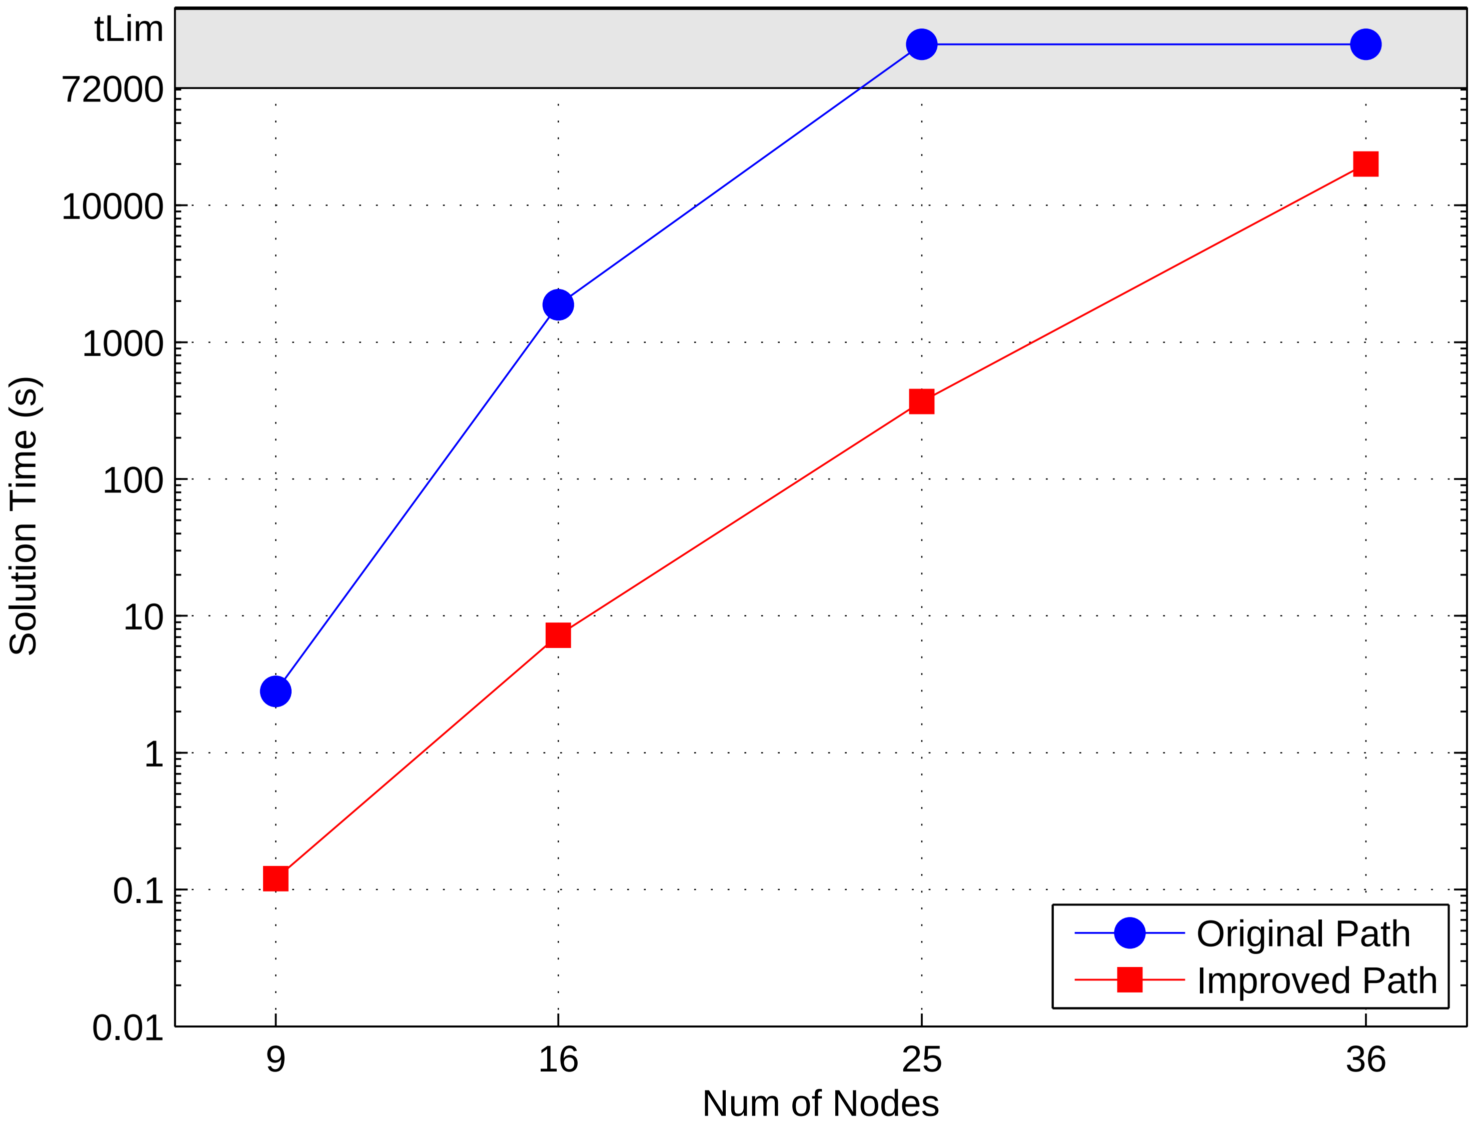

Running Time Results: Asymmetric Topologies

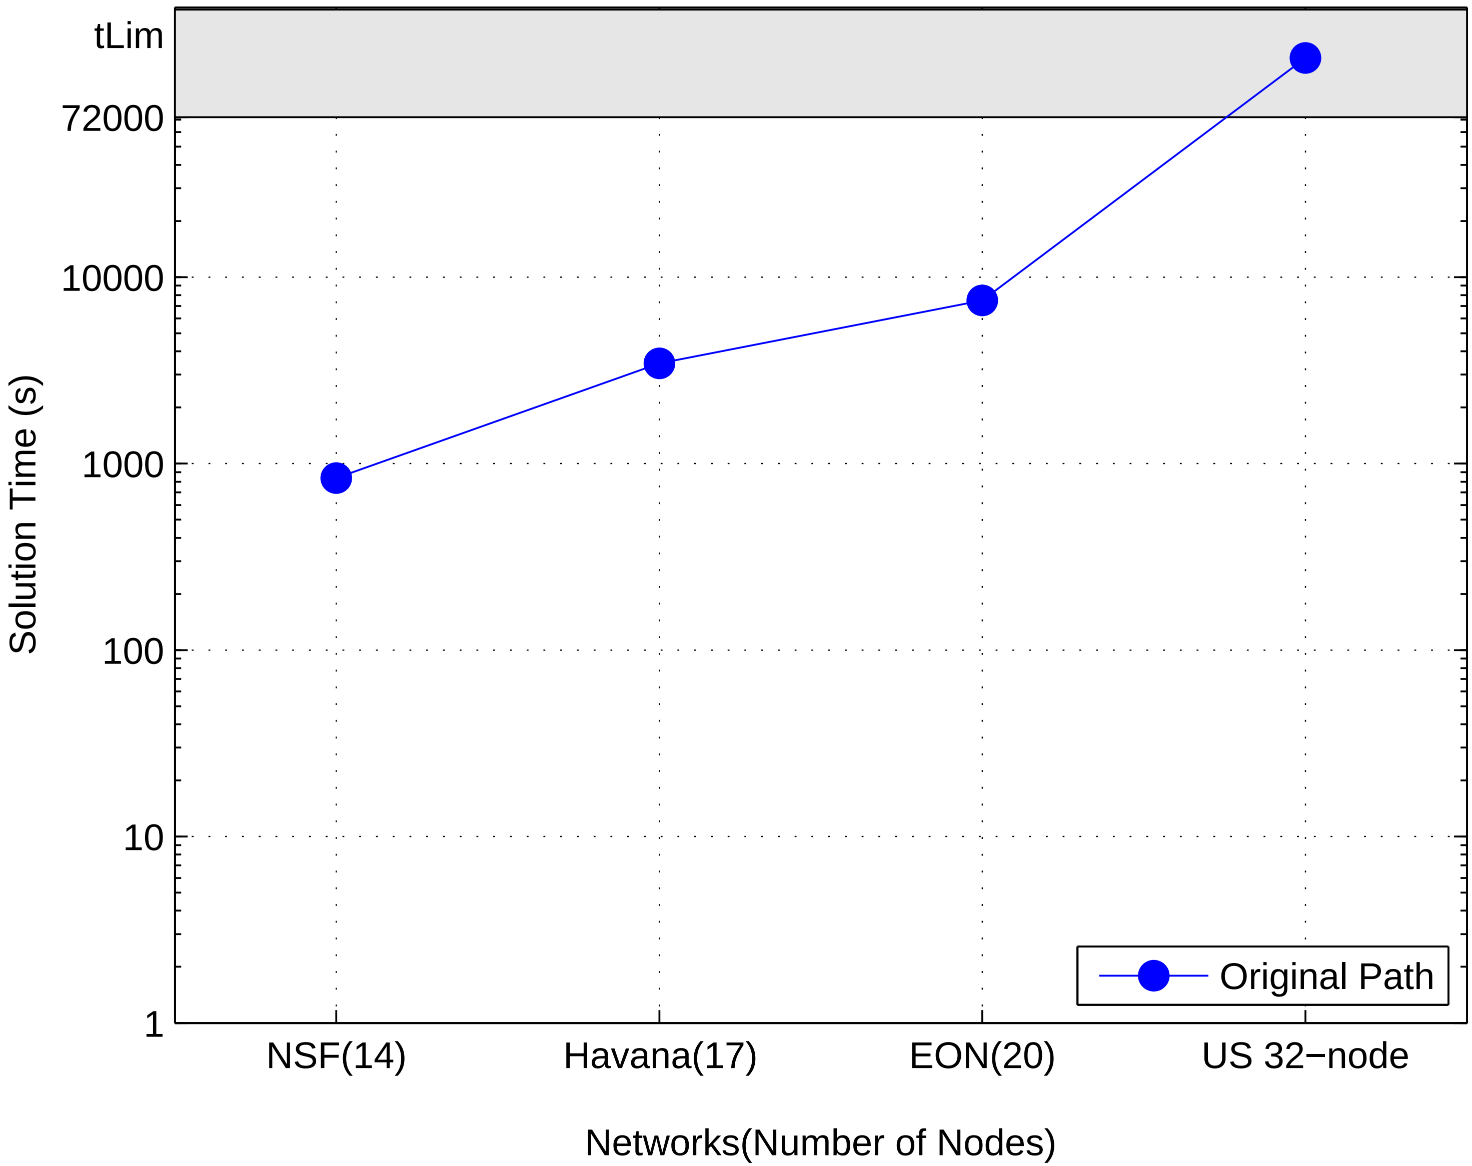

Running Time Results: Asymmetric Topologies

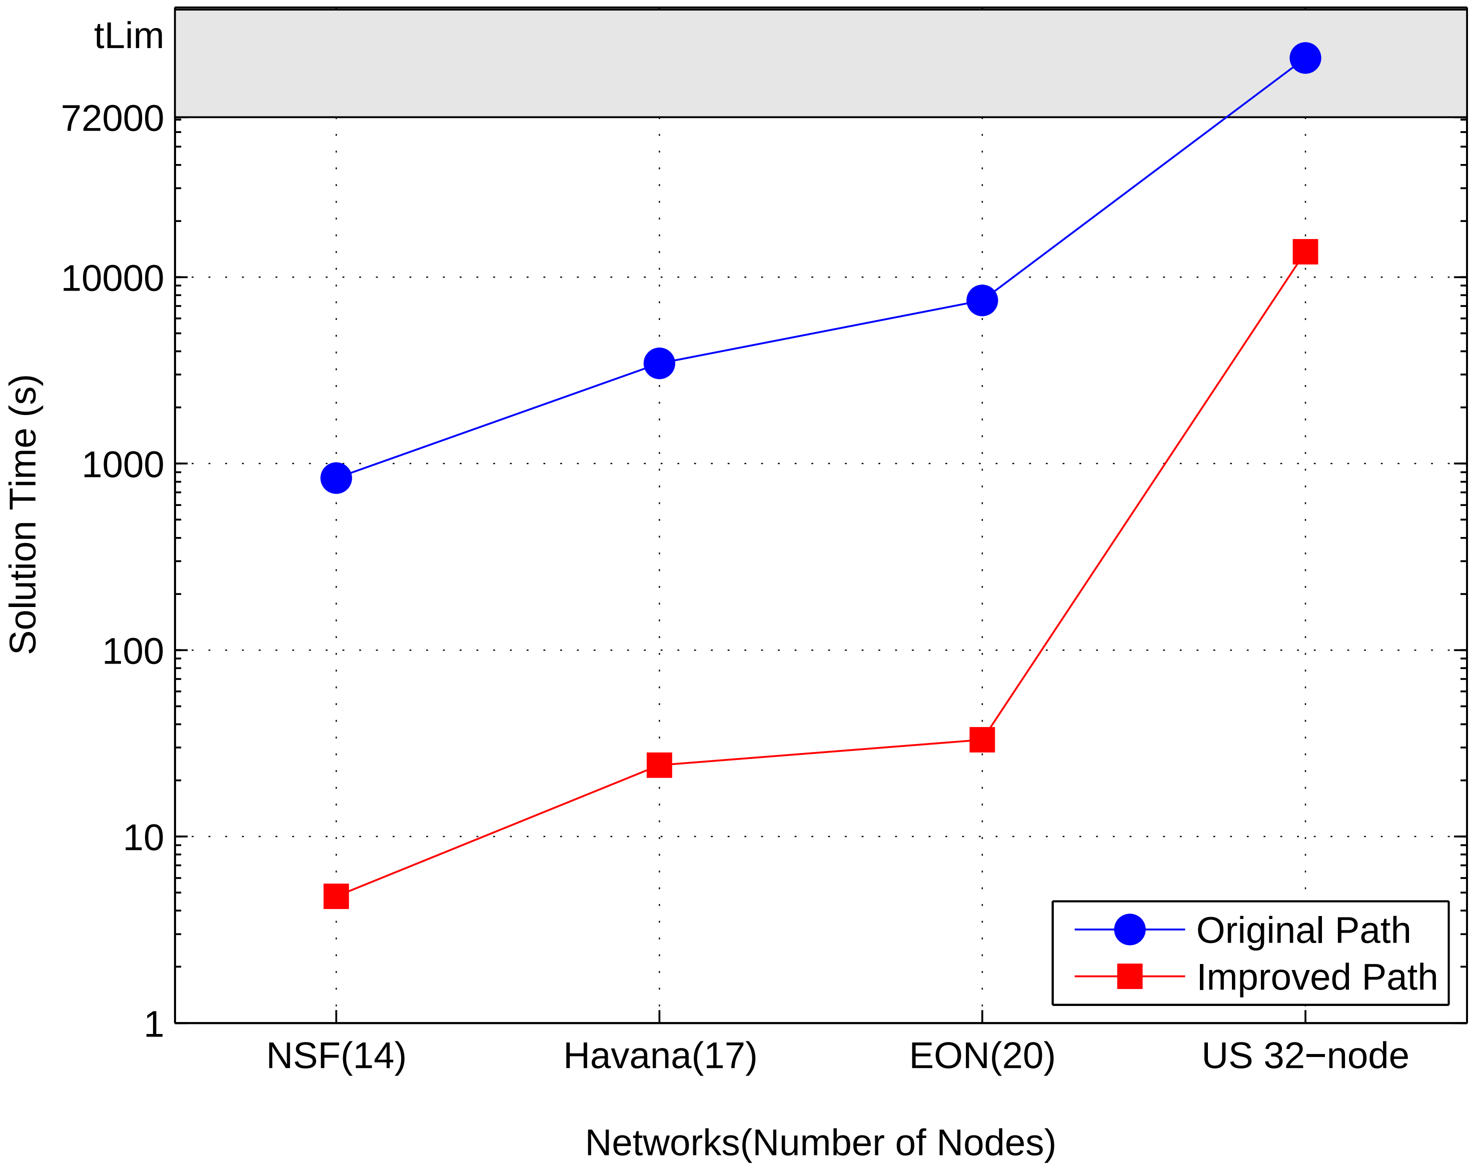

Conclusion & Ongoing Research

- Traffic grooming is ideal candidate for encompassing energy concerns
- Power-aware network design may lead to significant energy savings
- RWA subproblem can be solved efficiently
- Current research focuses on:
  - more accurate power consumption models for traffic grooming
  - computationally efficient formulations for optical network design problems
    - traffic grooming
    - impairment-aware RWA
    - multicast RWA and grooming