Statistical Approach to NoC Design

Itamar Cohen, Ori Rottenstreich and Isaac Keslassy

Technion (Israel)

Problem

The traffic matrix in NoCs is often-changing and unpredictable

$\Rightarrow$ makes NoCs hard to design

Road Capacities

- **Goal:** Design road capacities between a city and its suburbs
- Let's model the traffic matrices…

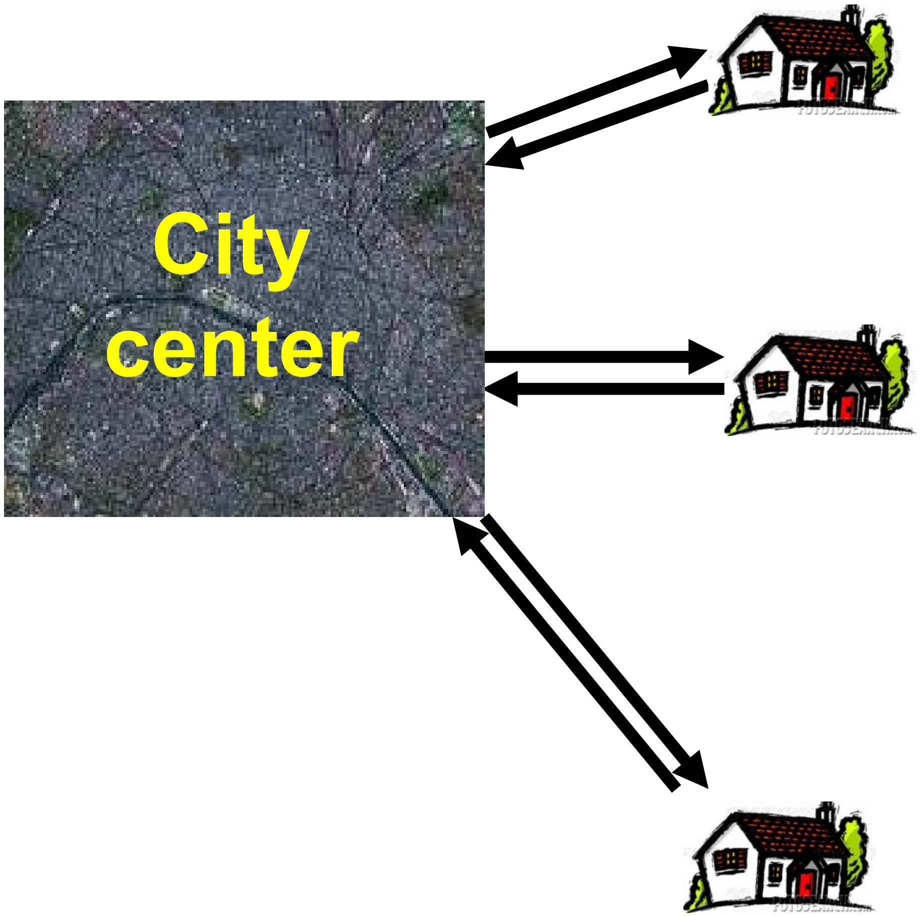

Road Capacities

- **Morning peak:** most traffic towards city

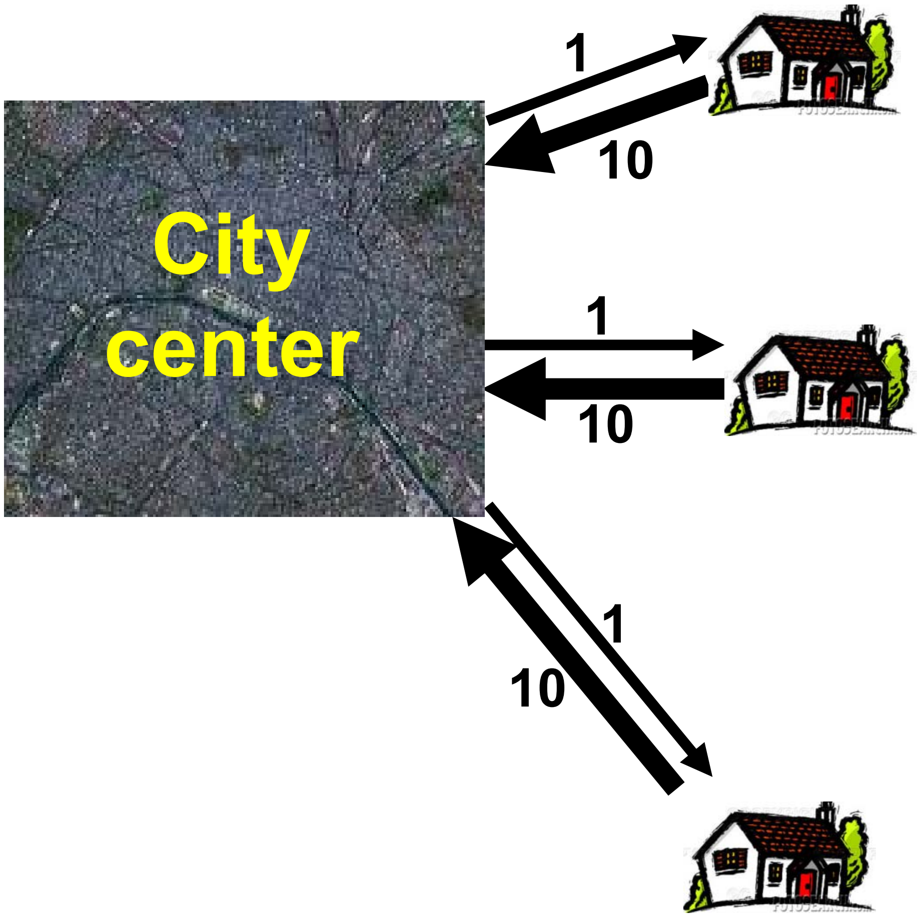

Road Capacities

- **Morning peak:** most traffic towards city
- **Afternoon peak:** most traffic leaving city

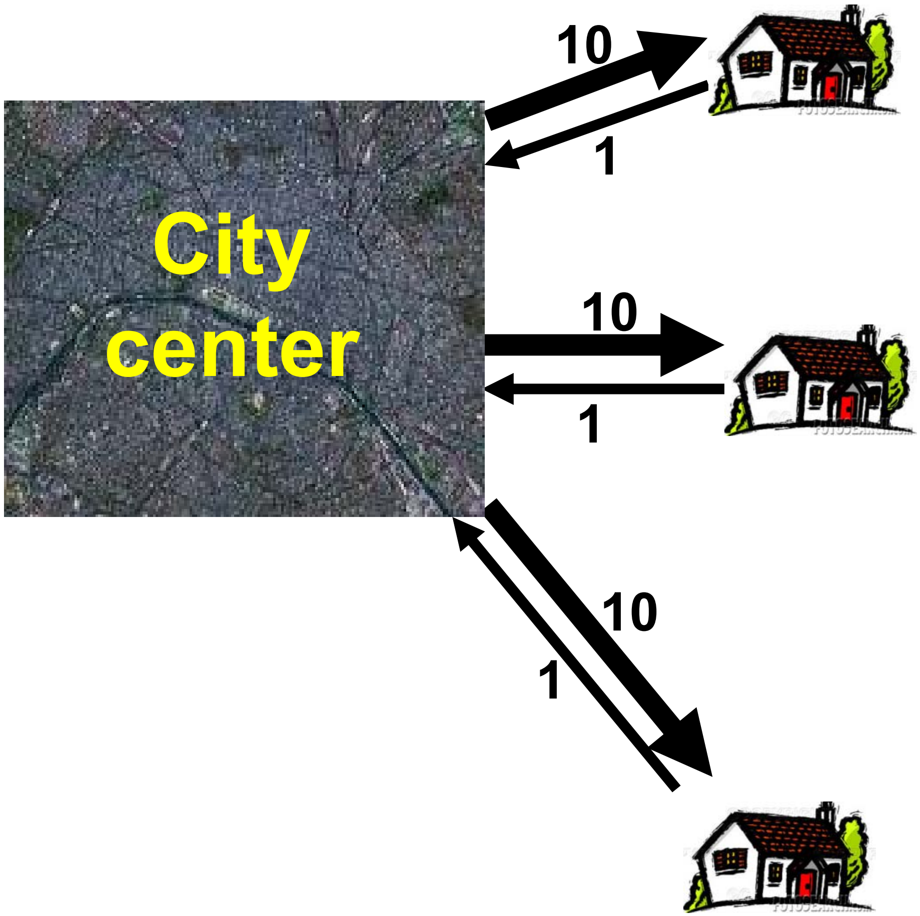

Solution (1): Average-Case

- Solution (1): plan for **average-case**
  - i.e. allocate capacity of ~5 for each link.
  - $\lambda < \mu$
- Problem: traffic jam during many hours, every day.

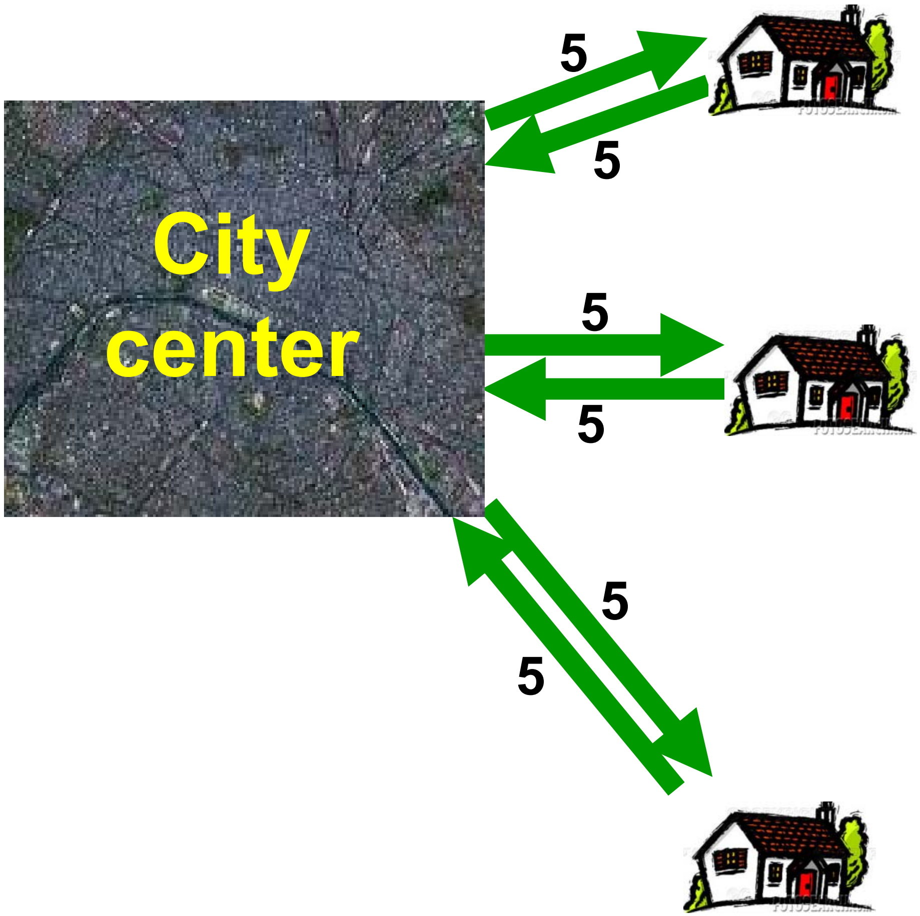

Solution (2): Worst-Case

- Solution (2): plan for **worst-case**
  - i.e. allocate capacity of 10 for each link.

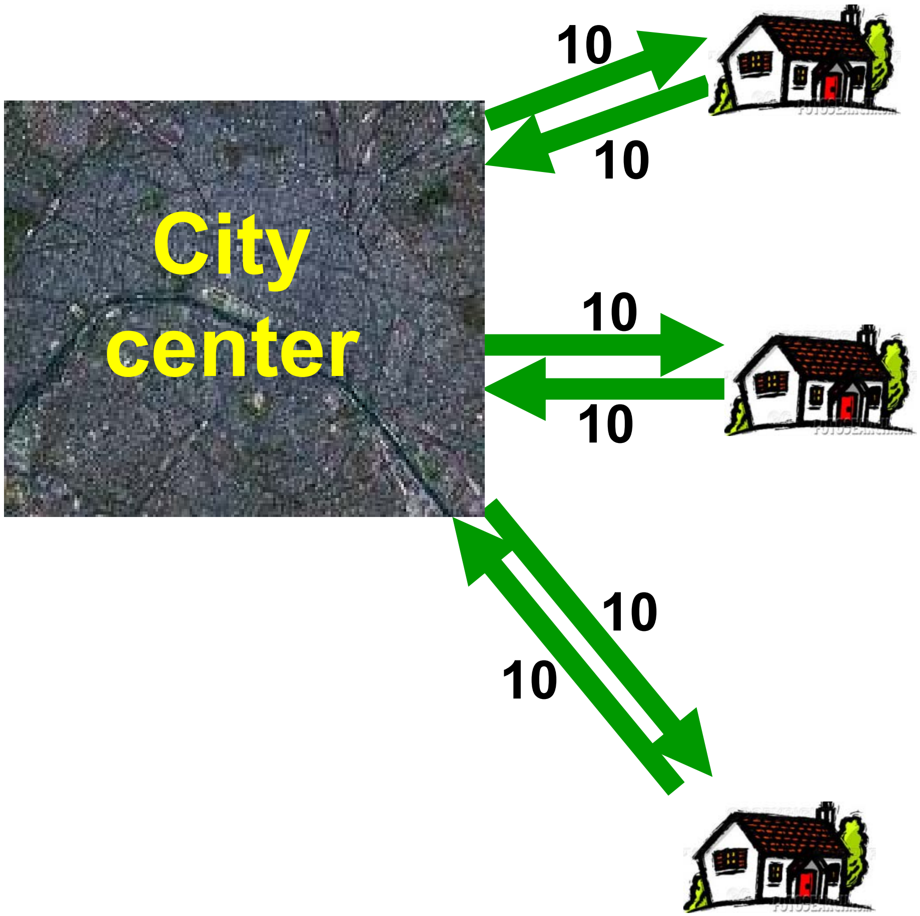

Holidays

- Problem: traffic burst in the holidays
  ⇒ excessive resources
- Solution (3): **statistical approach**
  - Enough capacity for 99% of the time
  - Allow for occasional congestion

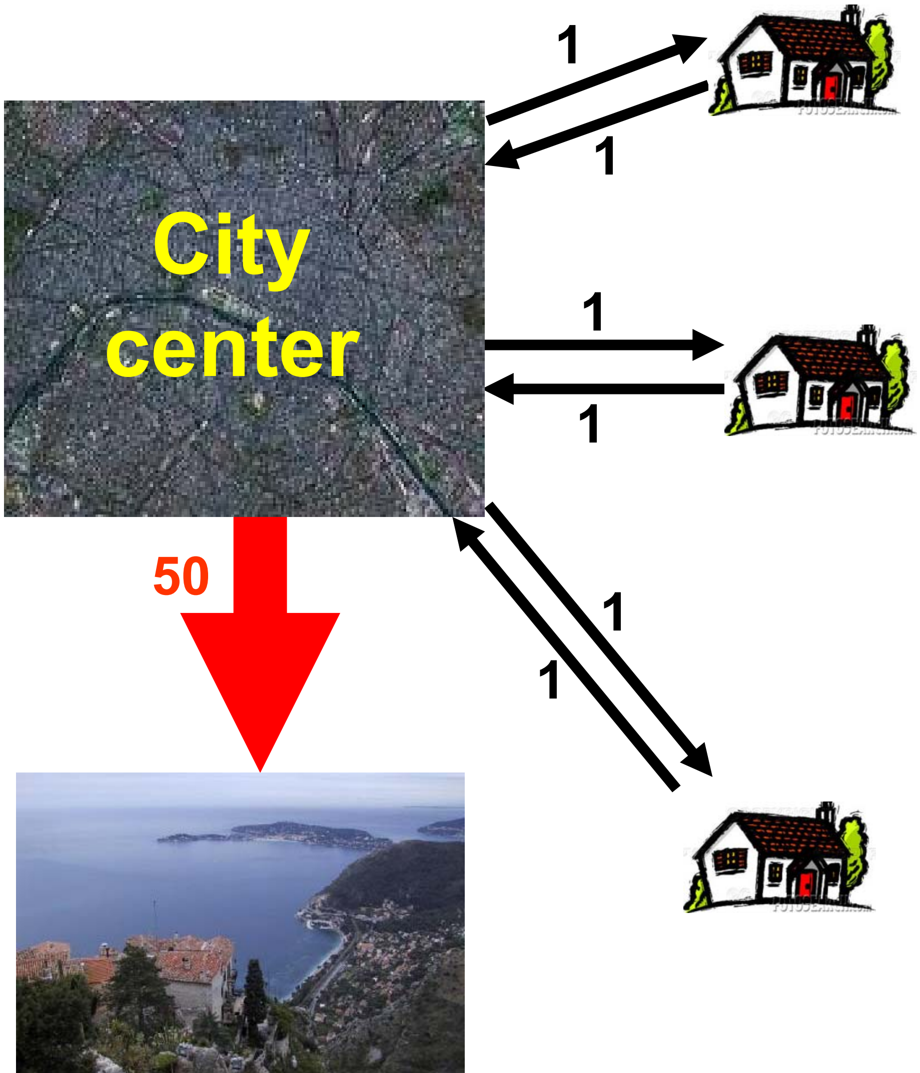

Back to the NoC world

- Similar problems in NoC design process
  - City → Shared cache
  - Suburbs → Cores
  - Many possible traffic matrices: writing, reading, etc.

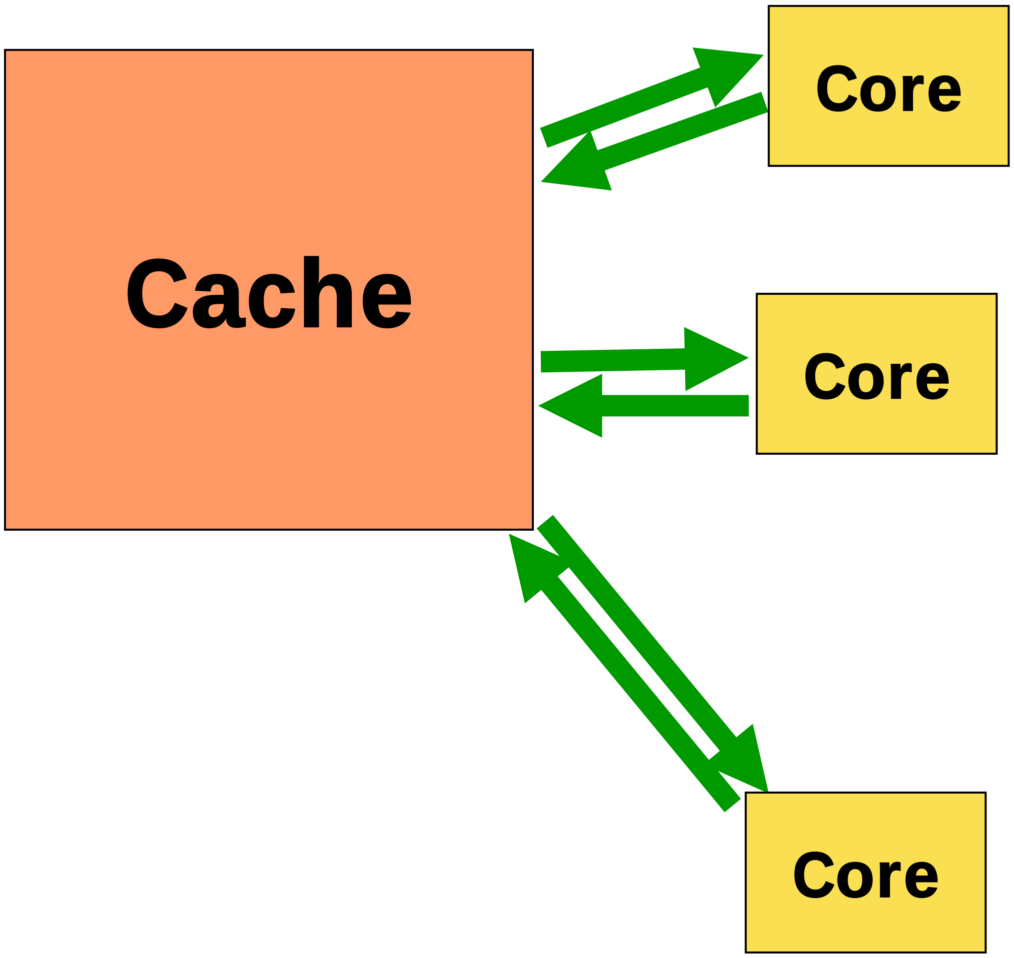

Motivation

Tradeoff:
more resources
vs. better guarantee

- **"Easy" solution: Worst-Case Guarantee**
  - E.g., no-congestion guarante[…]
    time

Better when no congestion is allowed

- **Objective: Statistical Guarantee**
  - E.g., no-congestion guarante[…]
    time

Better when power/area are most important

Statistical Approach to NoC Design

Given:

1
- Traffic matrix distribution
- Topology
- Routing
- Link capacities

2
Compute congestion guarantee

- “95% of traffic matrices will receive enough capacity”

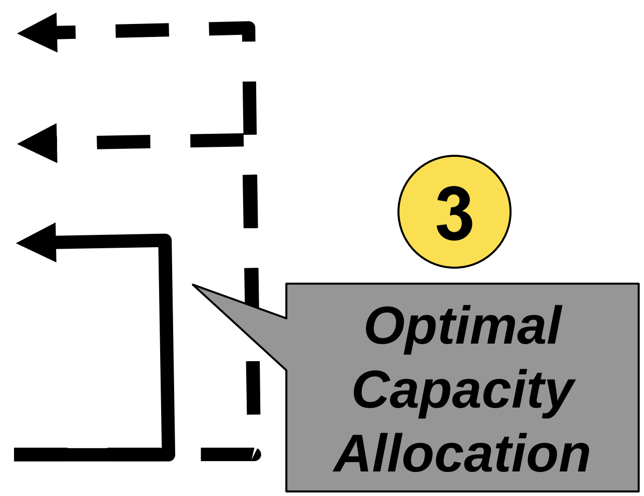

CMP (Chip Multi-Processor)

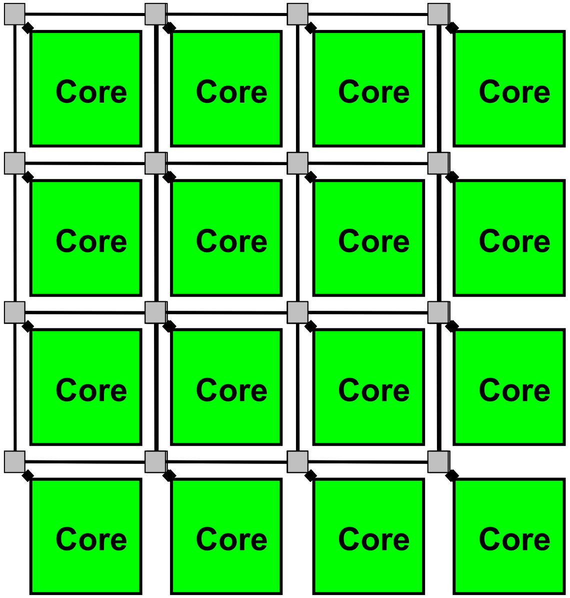

How can we model the future traffic matrix distribution?

Traffic Matrix Distribution

- **Measure** on real large CMP networks?
- Problems
  - Their future applications are unknown
  - Their future traffic types are unknown: Between processors? Cache accesses? Control traffic?
  - (Chicken-and-egg problem: traffic might depend on what architecture offers)

Traffic Matrix Distribution

- **General model:** any core can communicate with any core - but with a bounded core input/output rate.
  - Example: each core may send/receive up to 1 Gbps.
  - Used in CMPs [Murali et al. '07] – but also backbone networks [Dukkipati et al. '05], interconnection networks [Towles and Dally '02], VPN networks [Duffield et al. '99], routers [McKeown et al. '96]

Core 1 sends ⟶

$$T = \begin{pmatrix} 0 & 0.5 & 0.1 & 0.2 \\ 0 & 0 & 0 & 0.6 \\ 0.2 & 0.1 & 0 & 0 \\ 0.1 & 0.1 & 0.3 & 0 \end{pmatrix} \quad \leq 1$$

- Reminder: just one model among many…

Statistical Approach to NoC Design

Given:

1

- Traffic matrix distribution
- Topology
- Routing
- Link capacities

3

Optimal Capacity Allocation

2

Compute congestion guarantee

- "95% of traffic matrices will receive enough capacity"

Compute Congestion Guarantee

Show that "**95%** of all traffic matrices will receive enough capacity on link ℓ"

$\Leftrightarrow$

1. Compute ***Traffic-load distribution Plot*** *(or* ***T-Plot****)* **for link** ℓ
2. Show that the load on link ℓ is less than its capacity for **95%** of traffic matrices.

Link T-Plot

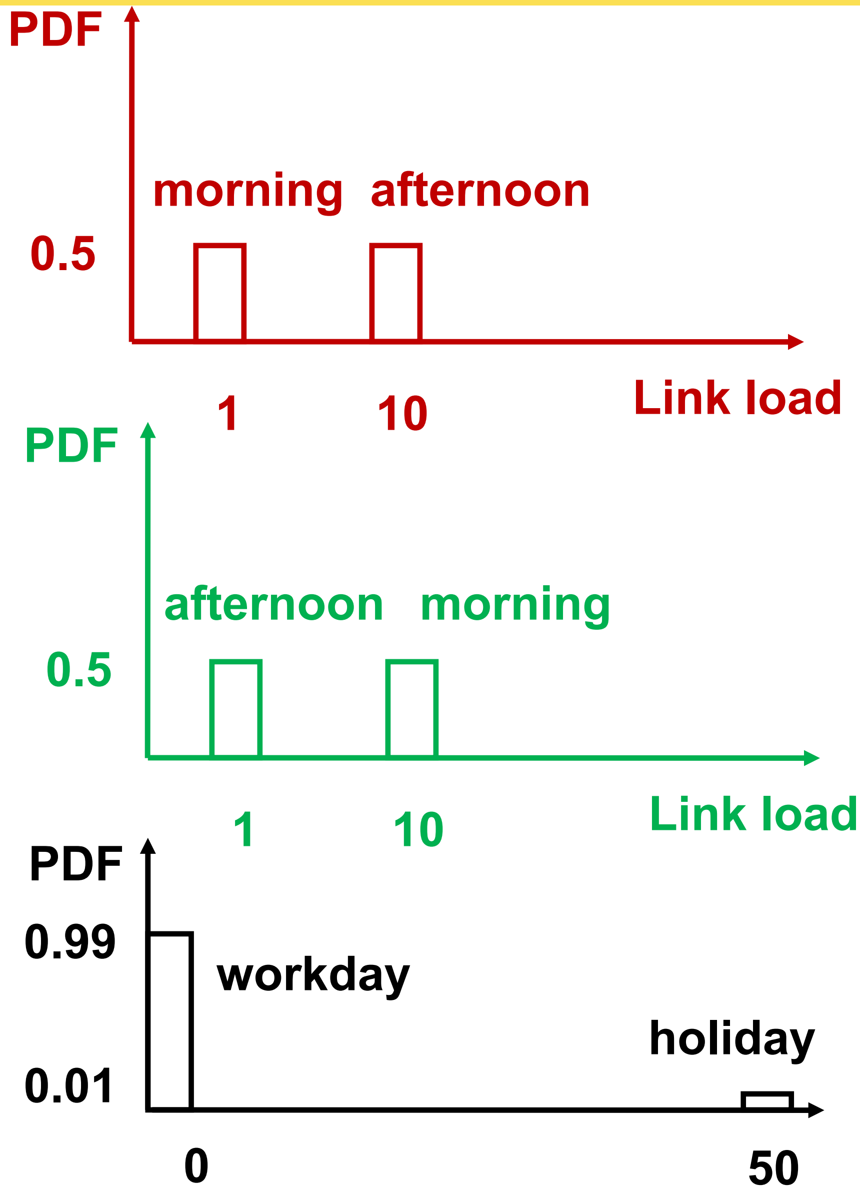
PDF
morning
afternoon
0.5
1
10
Link load
PDF
afternoon
morning
0.5
1
10
Link load
PDF
0.99
workday
holiday
0.01
0
50

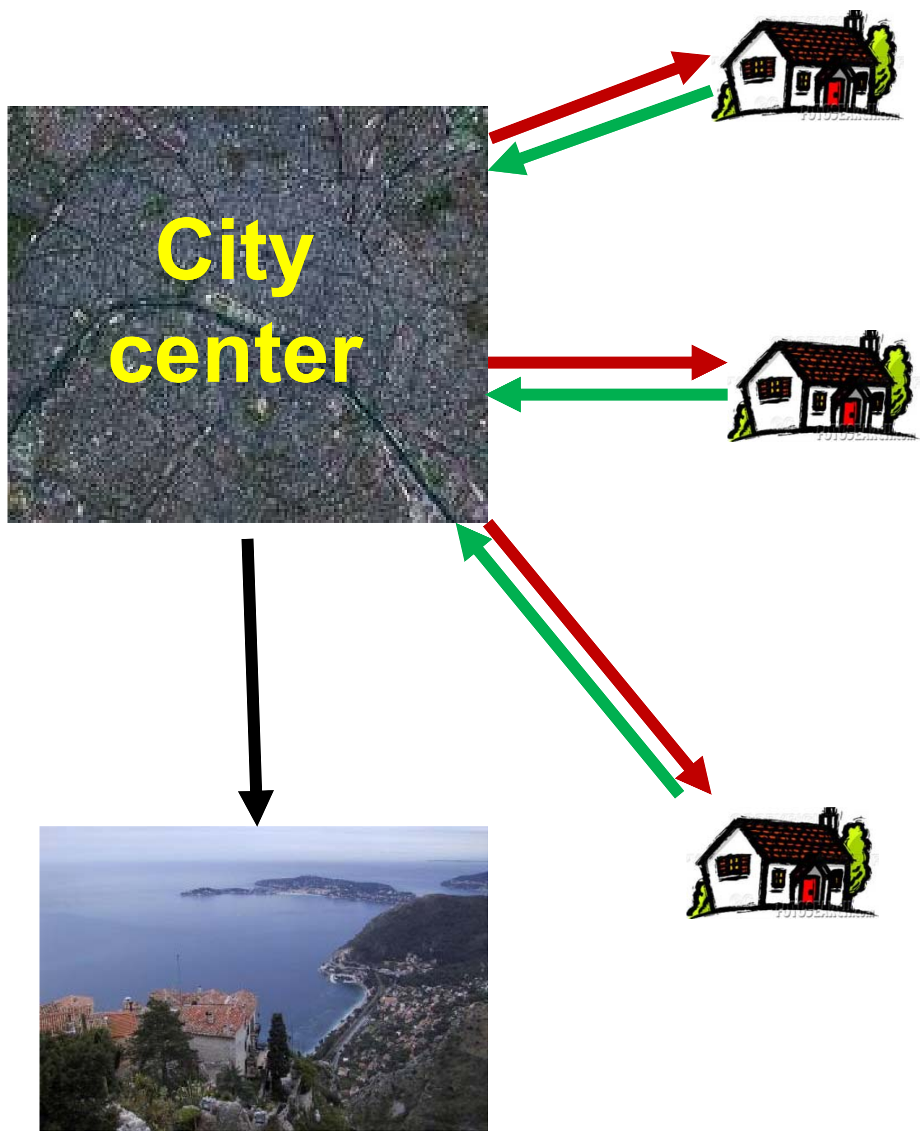
City
center

Global T-Plot

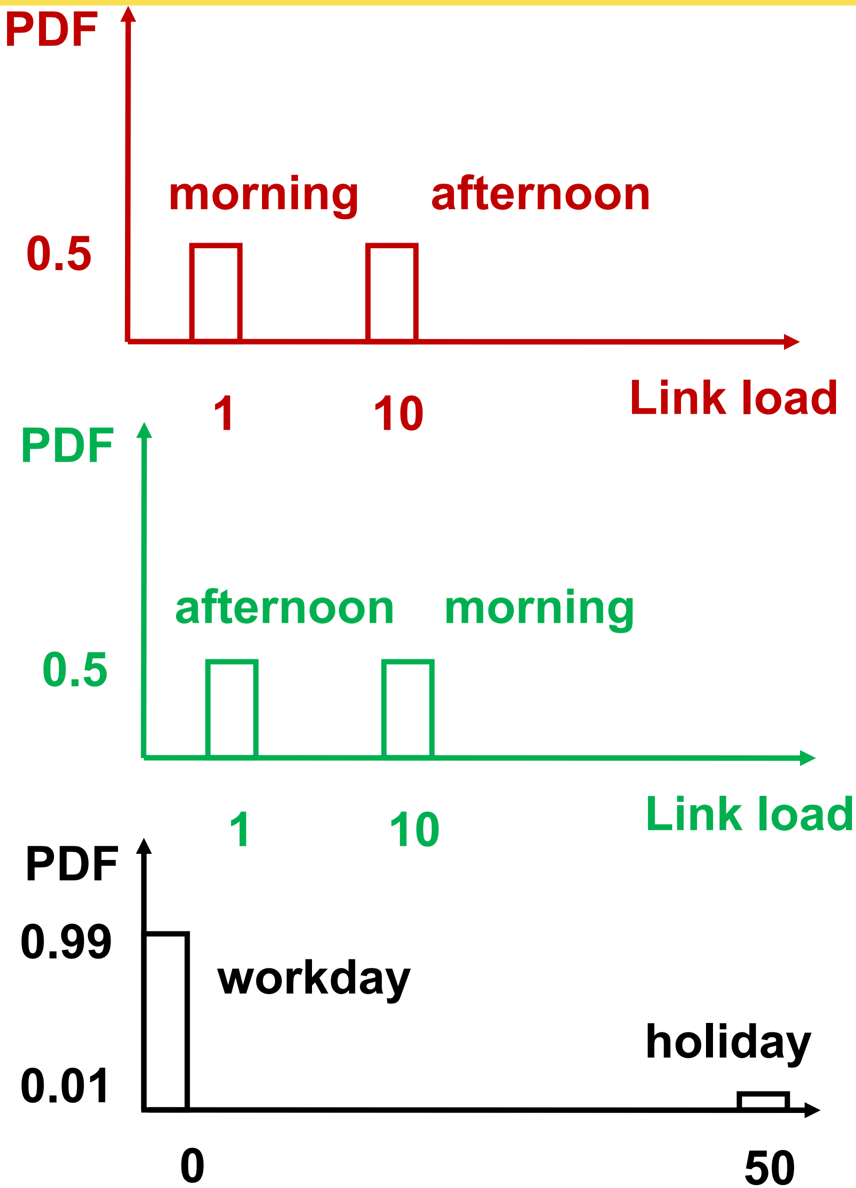

- We want to provide statistical guarantees on the whole network
  - i.e. on **all** links
- **Global T-Plot:** for each traffic matrix ***T***, measure maximum load among **all** links

Global T-Plot

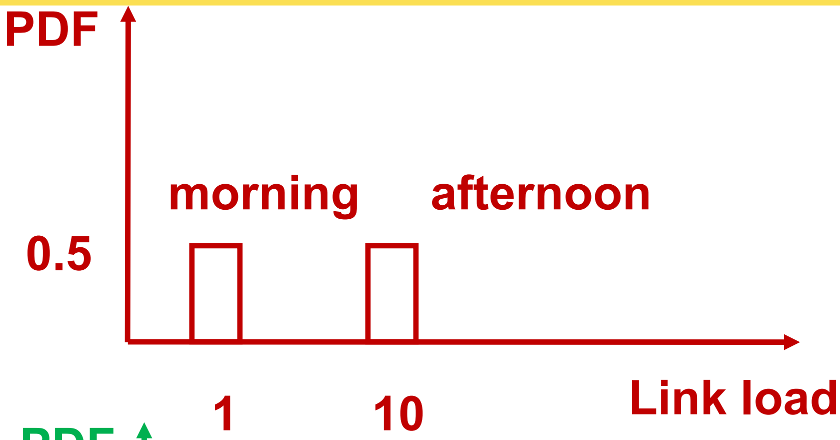

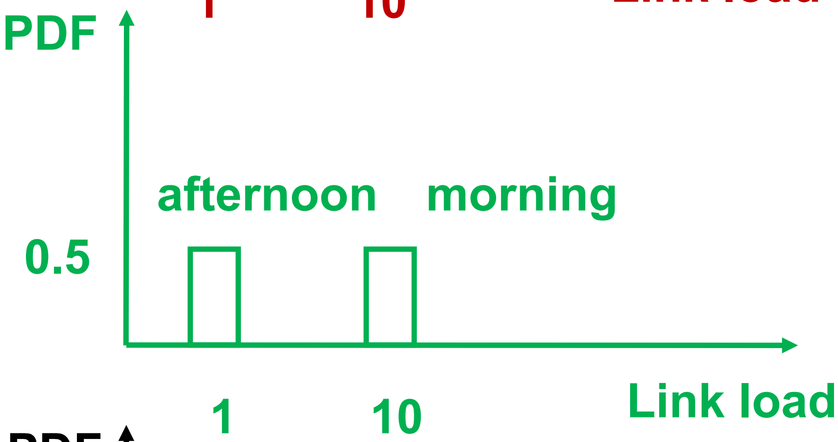

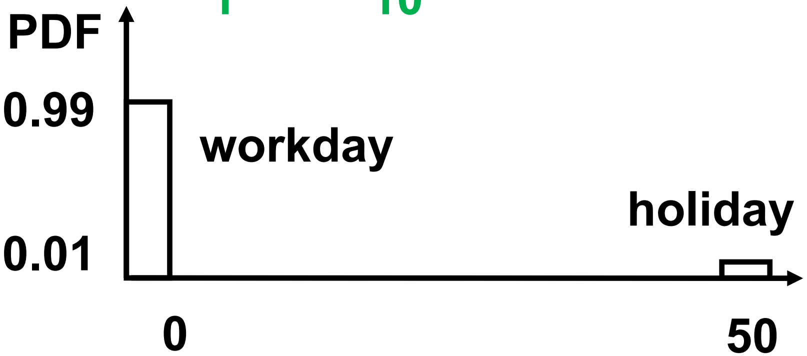

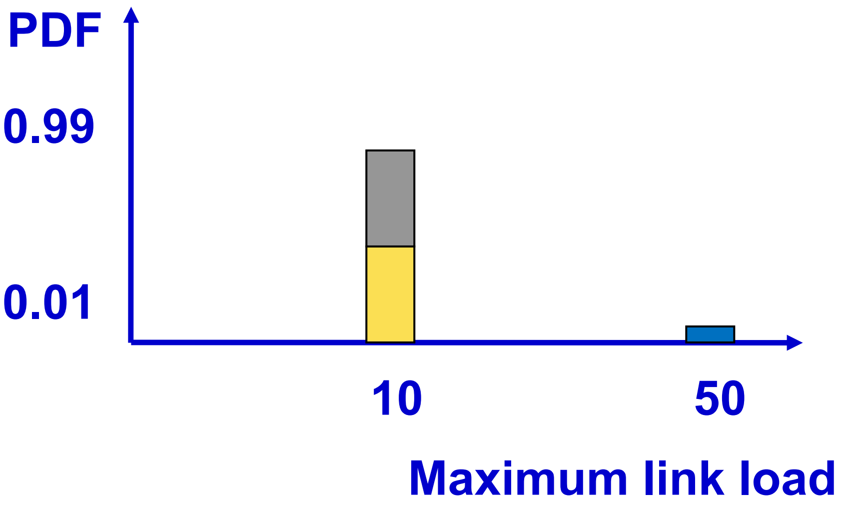

Local T-Plots in NoCs

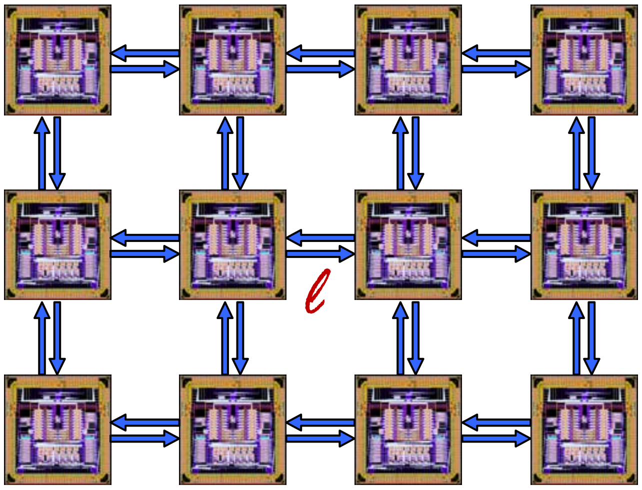

- Given:
  - Traffic matrices *T*∈***S***
  - 3x4 mesh topology
  - XY routing
  - Link ℓ
- Find T-Plot on ℓ

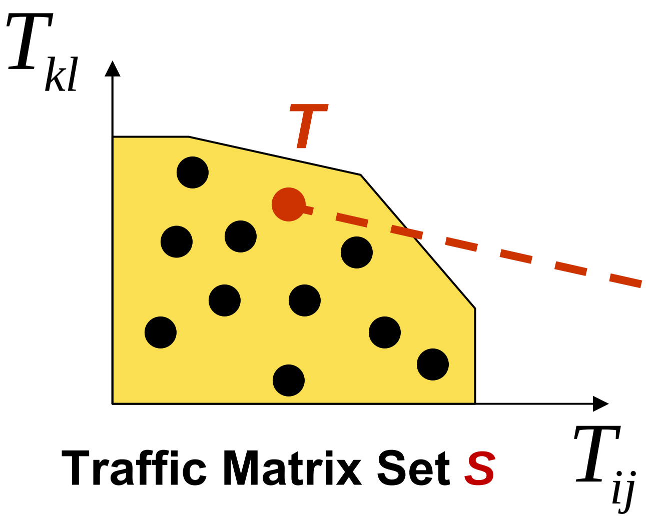

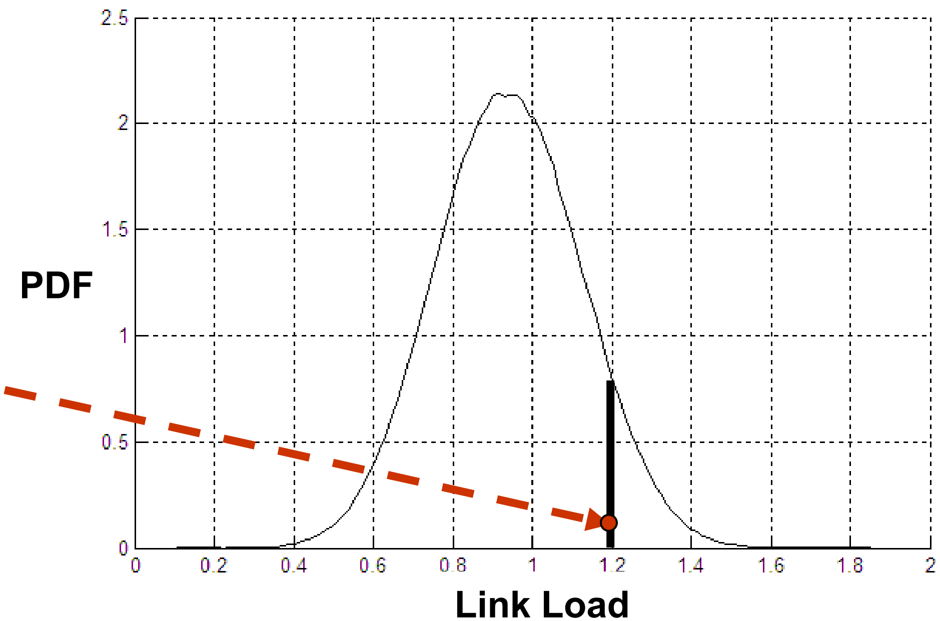

T-Plot

- Close-up view:

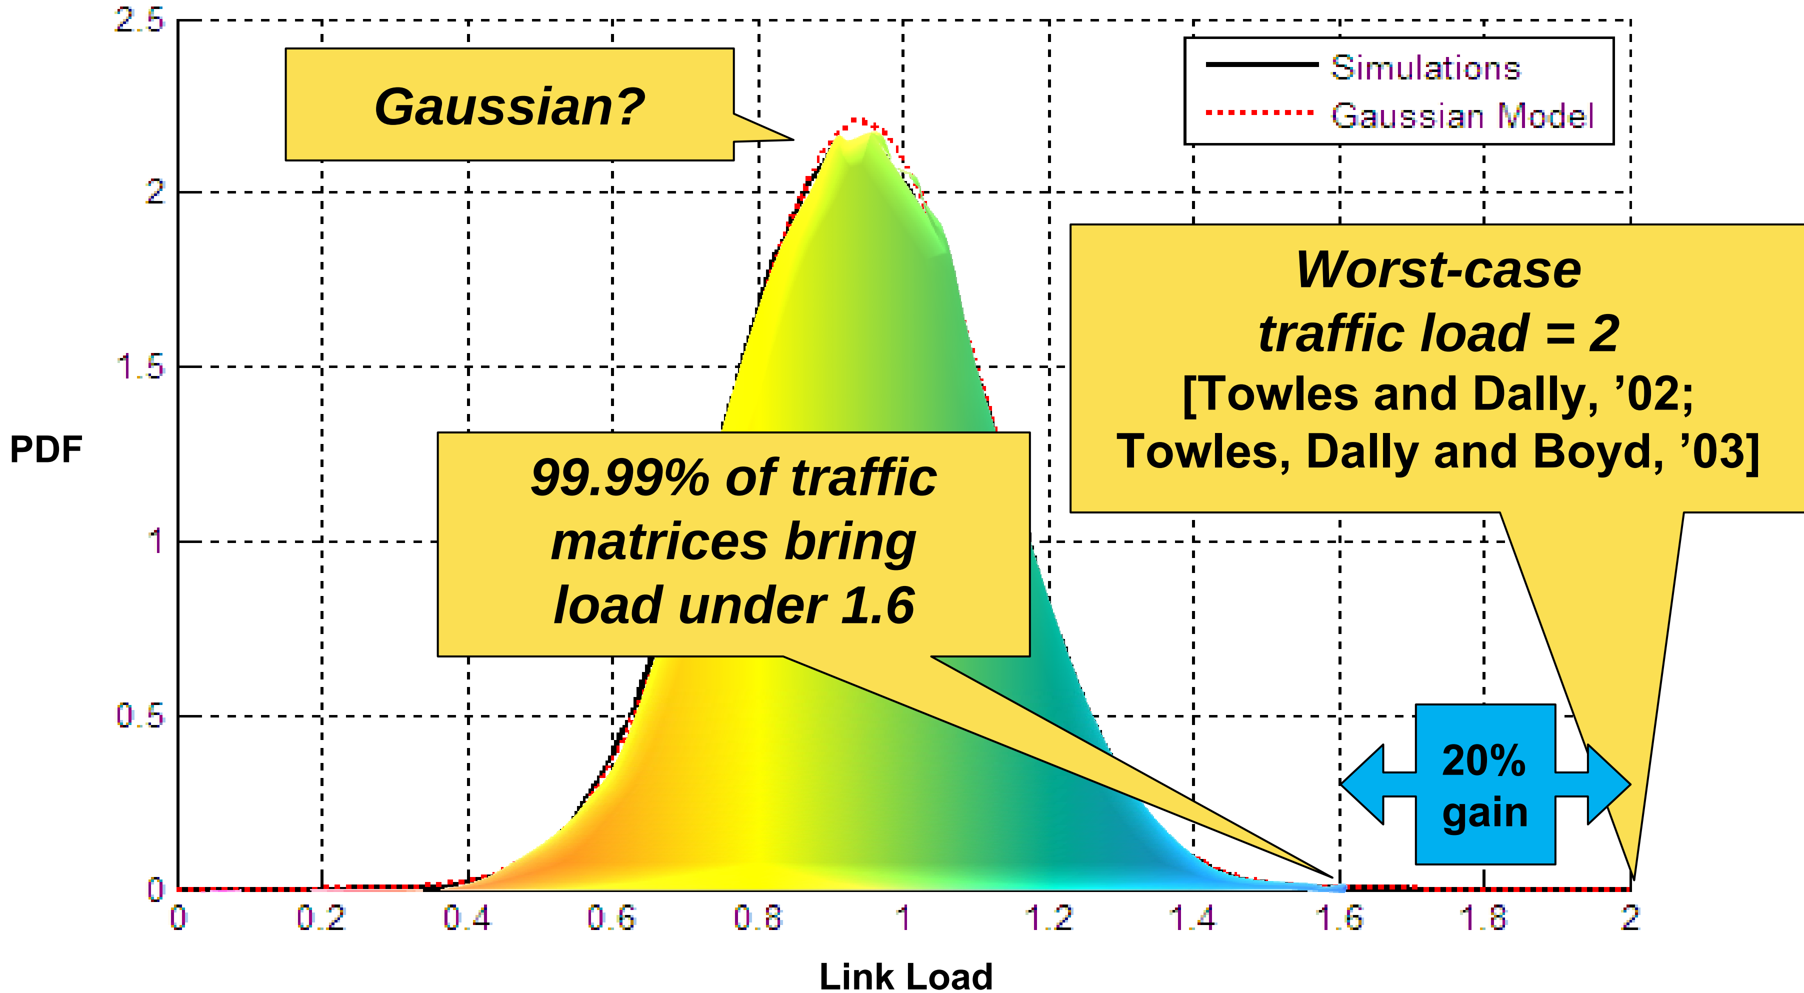

Are T-Plots Gaussian?

- Some T-Plots may be well approximated as Gaussian.

- **Intuition (Central Limit Theorem):** when ***N*** grows, the sum of ***N*** i.i.d. (independent and identically distributed) random variables converges to a Gaussian distribution.

Gaussian T-Plot Example

- Example: *n* x *n* mesh with XY routing.
- Assume i.i.d traffic from *i* to *j*
- ***Theorem***: As *n* grows, T-Plot for *ℓ* converges to Gaussian.

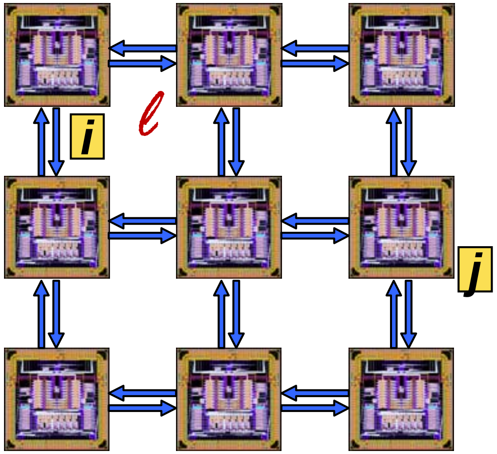

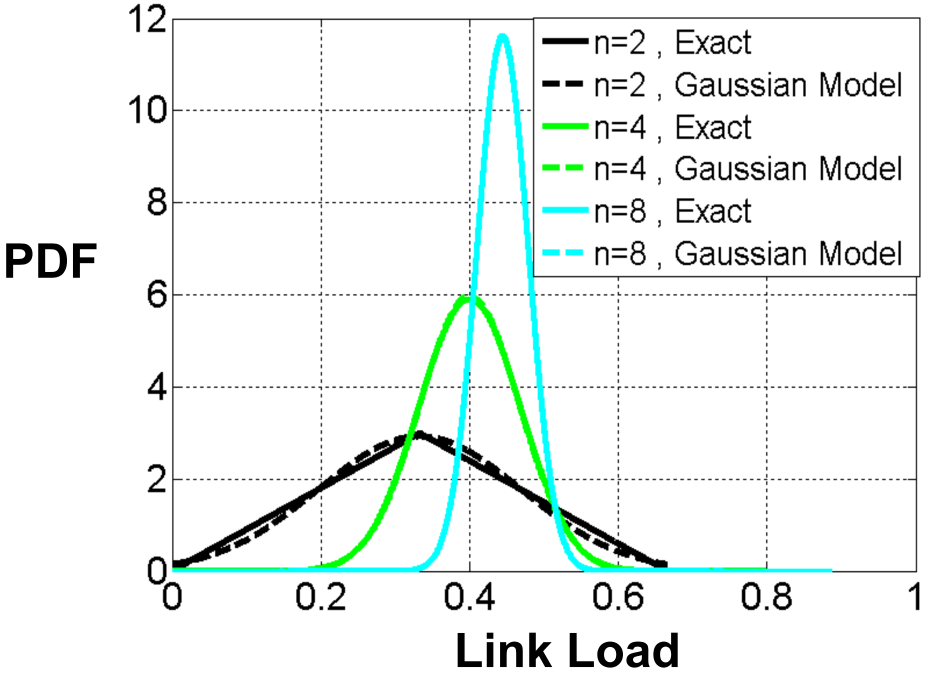

Computing the T-Plot

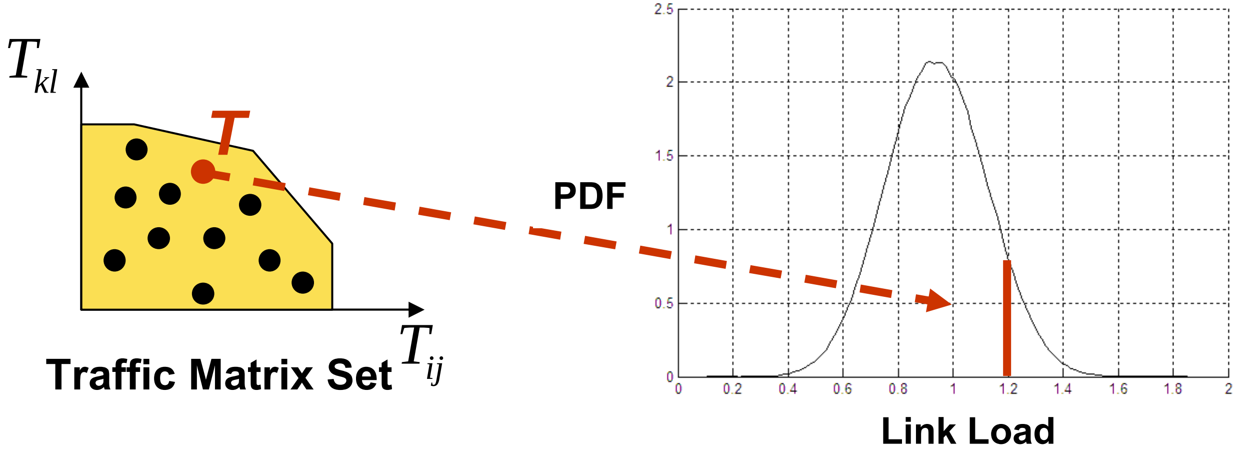

- ***Theorem***: for an arbitrary graph and routing, computing the *T-Plot* is ***#P-complete***.
- #P-complete problems are at least as hard as NP-complete problems.
  - NP: "Is there a solution?"
  - #P: "How many solutions?"

Computing the T-Plot

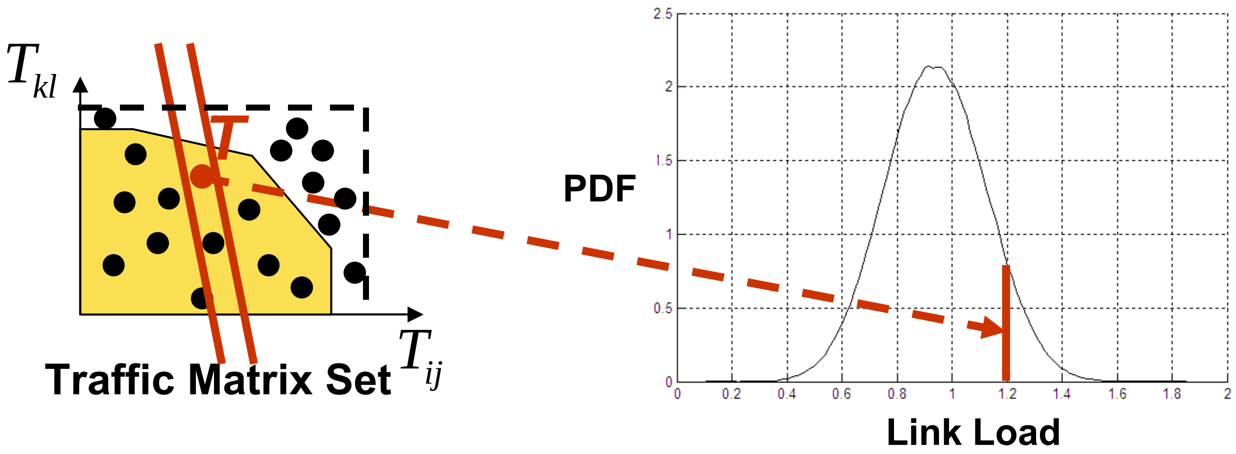

- Idea: use Monte Carlo simulations to approximate T-Plots
  - Plot many points, and count number of points to approximate density.

Link T-Plot Vs Global T-Plot

- **Global T-Plot:** for each traffic matrix ***T***, measure maximum load among **all** links

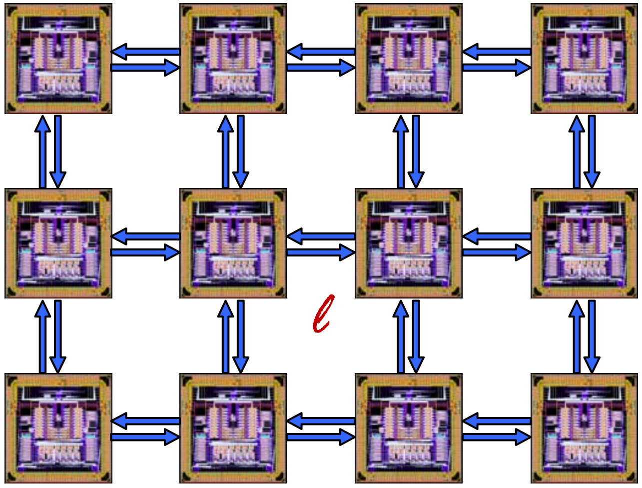

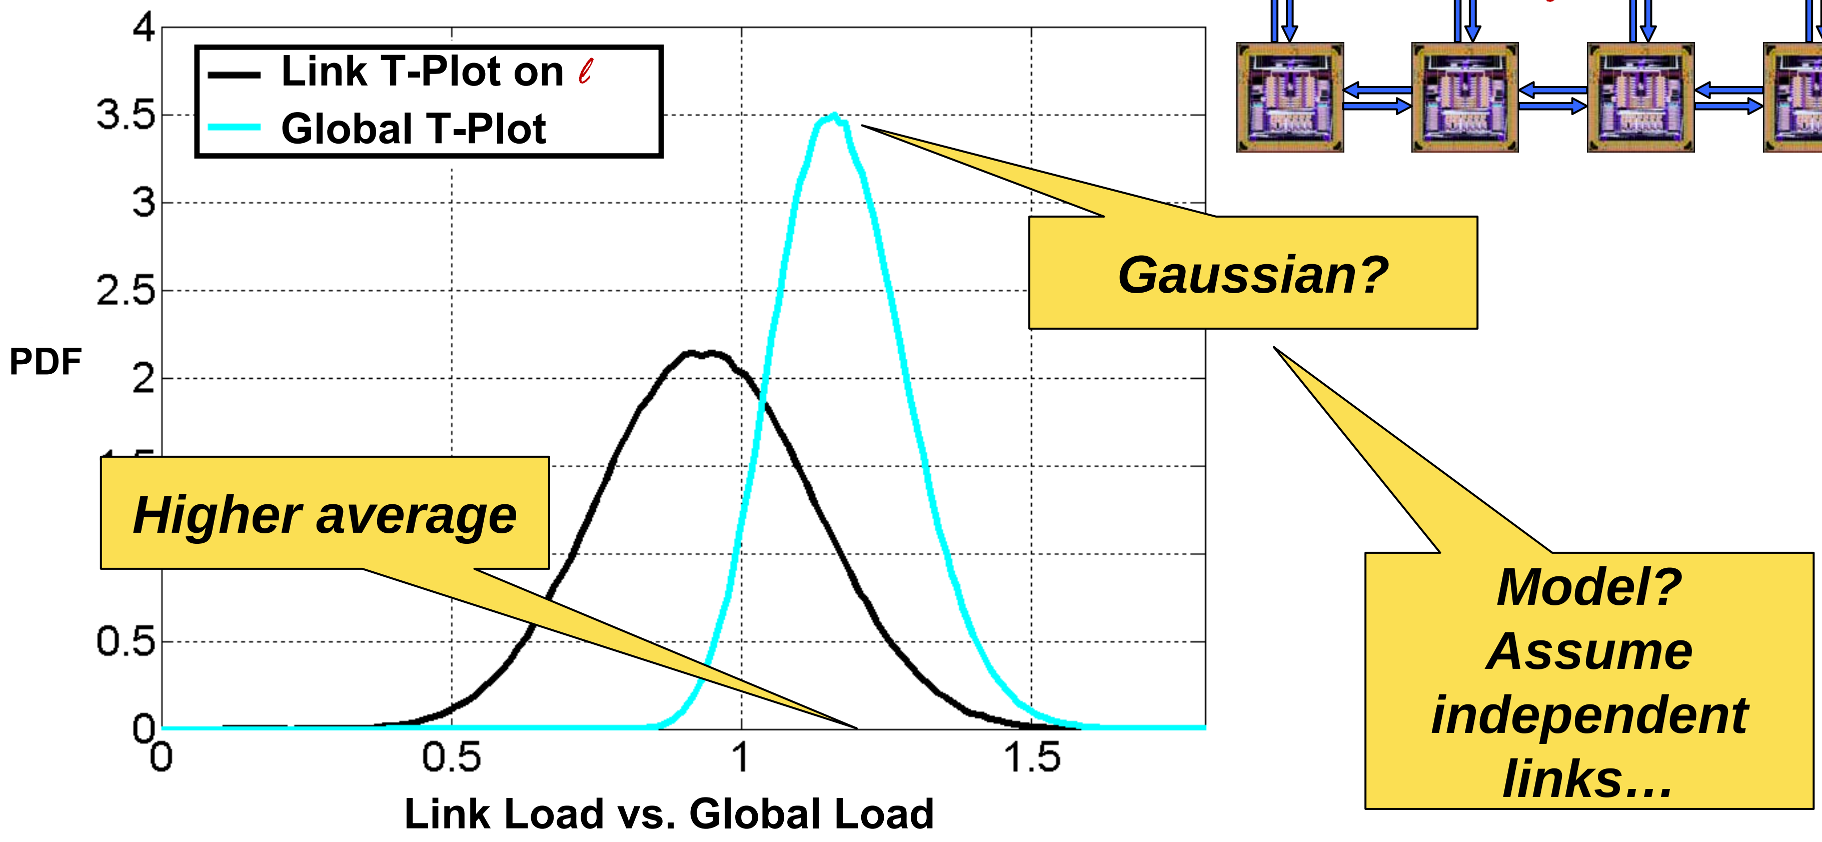

Link-Independent Model: Simple Example

- Assume:

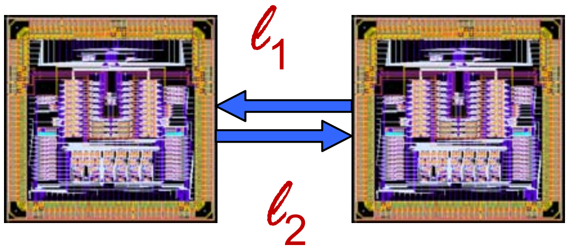

$$\ell_1: \begin{cases} 0 & w.p. \frac{1}{2} \\ 1 & w.p. \frac{1}{2} \end{cases}, \quad \ell_2: \begin{cases} 0 & w.p. \frac{1}{2} \\ 1 & w.p. \frac{1}{2} \end{cases}$$

- If ℓ_1 and ℓ_2 independent:

$$\text{Global Load: } \begin{cases} 0 & w.p. \frac{1}{4} \\ 1 & w.p. \frac{3}{4} \end{cases}$$

Independent-Gaussian Model

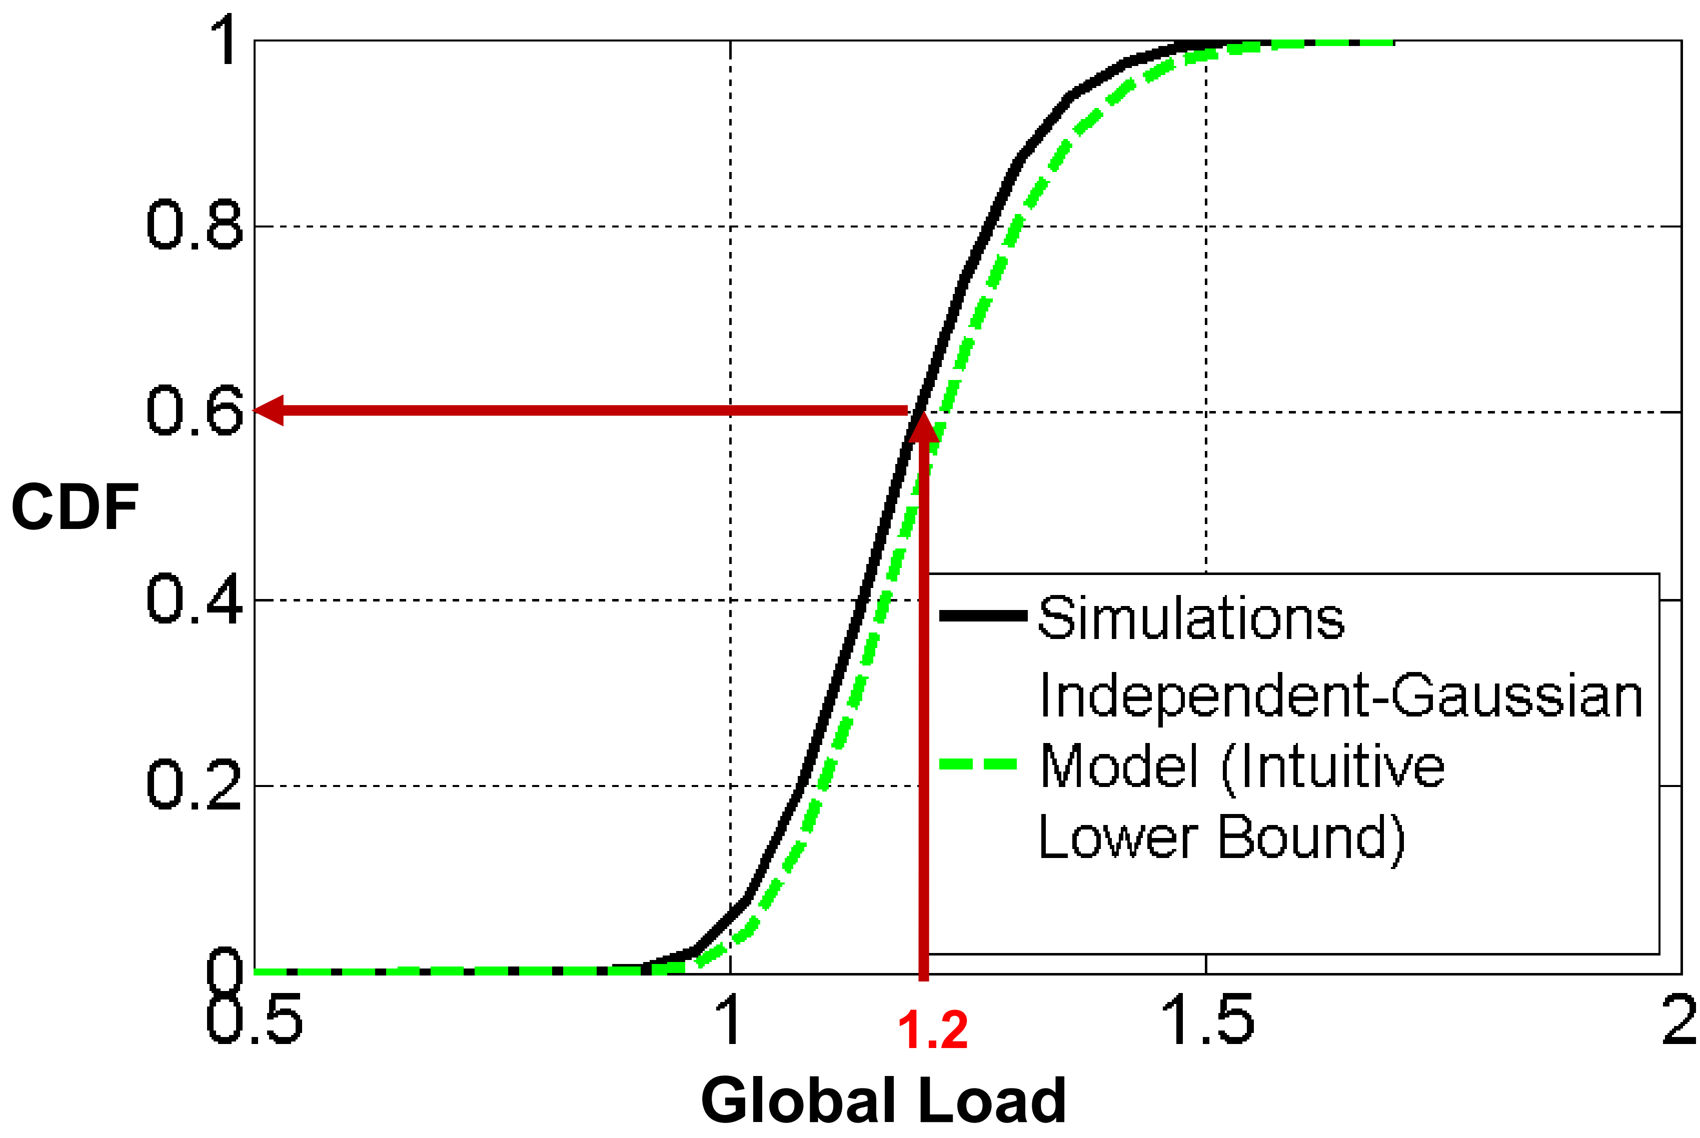

Statistical Approach to NoC Design

Given:

1

- Traffic matrix distribution
- Topology
- Routing
- Link capacities

2

Compute congestion guarantee

- “95% of traffic matrices will receive enough capacity”

3

Optimal Capacity Allocation

Capacity Allocation Intuition

- **Idea:** given a total capacity, distribute it so that the overflow probability on each link is **the same**.

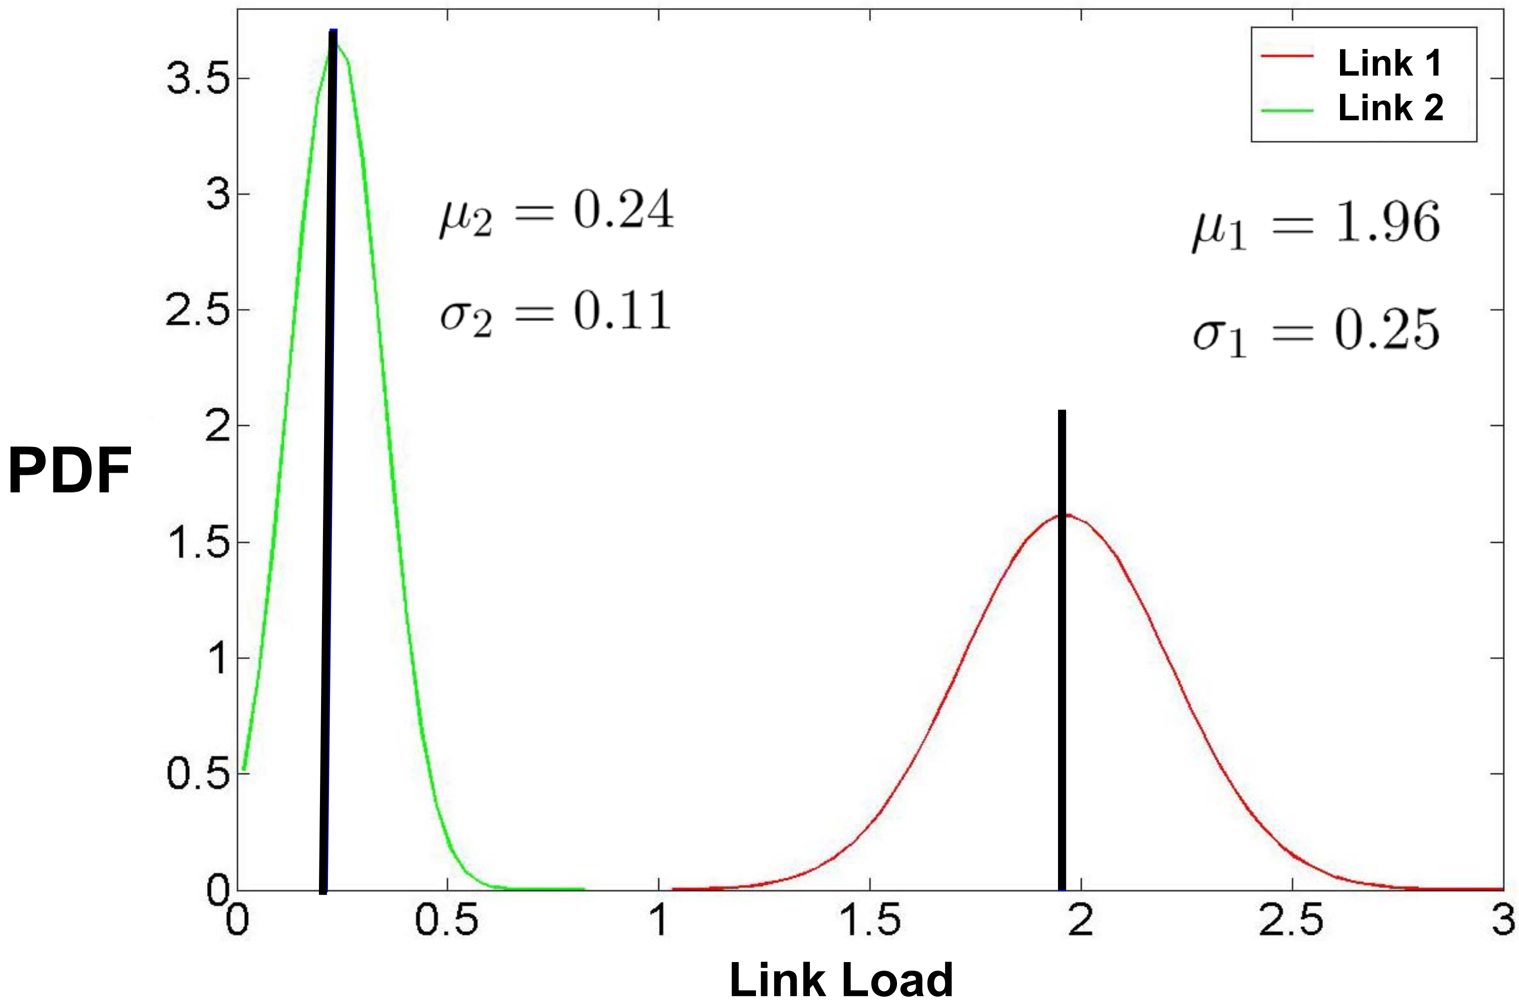

Capacity Allocation Intuition

- **Idea:** given a total capacity, distribute it so that the overflow probability on each link is **the same**.

- **Theorem:** if Gaussian-independent model holds, with same standard-deviation σ on all links, then **this scheme is optimal.**

Capacity Allocation Performances

- Total needed ca... us Capacity Allocation (CA)

Worst-case capacity saving with different link capacities

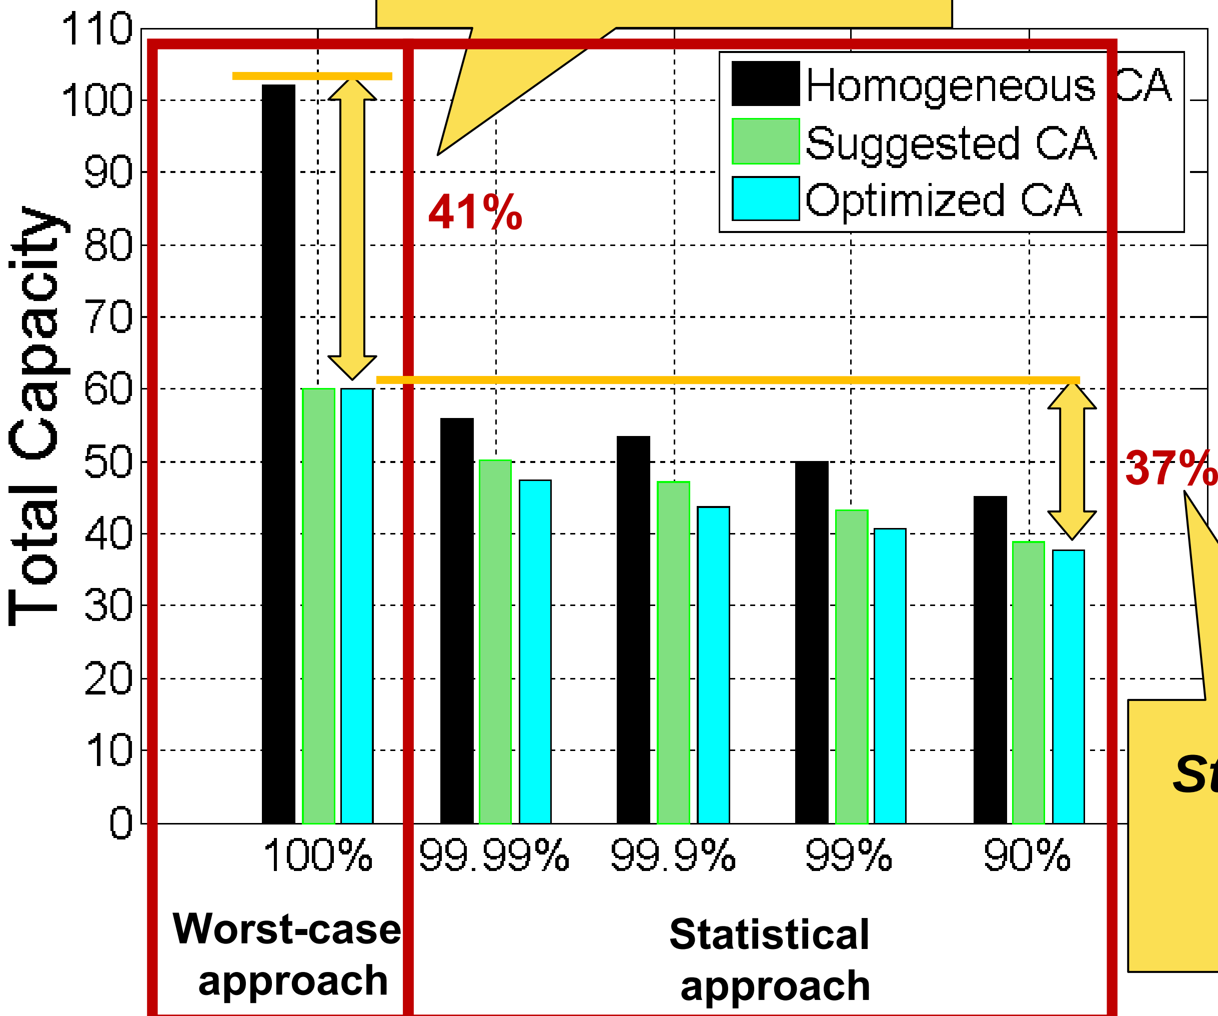

Statistical gain with 90% guarantee

Average Flow Delay Model

- Assume M/M/1 delay model
- Average flow delay distribution over all traffic matrices, for two different Capacity Allocation (CA) schemes:

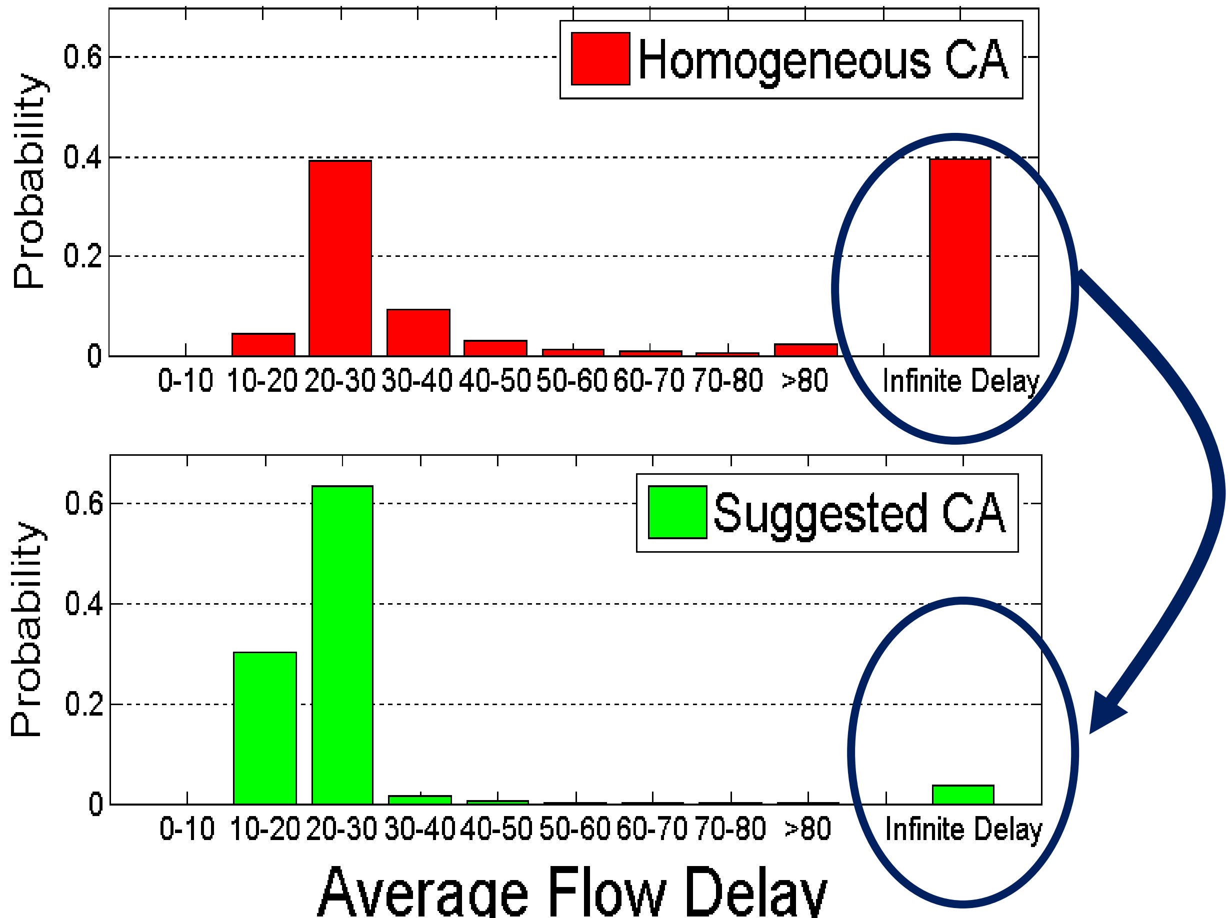

Statistical Approach to NoC Design

Given:

1
- Traffic matrix distribution
- Topology
- **Routing**
- Link capacities

2

Compute congestion guarantee

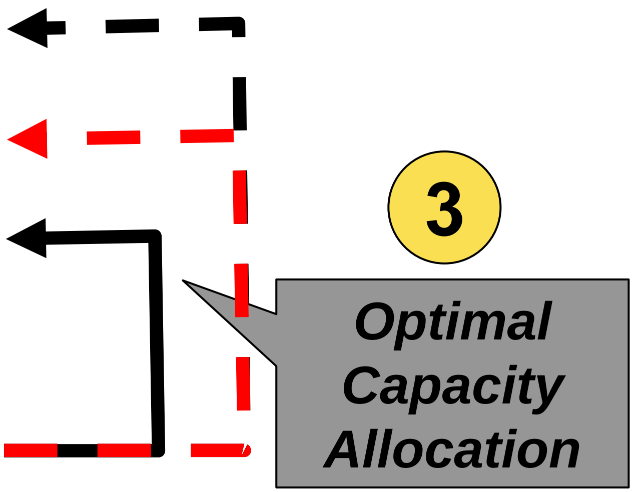

- "95% of traffic matrices will receive enough capacity"

Comparing routing algorithms

- DOR (XY) and O1TURN (½ XY, ½ YX) on 3x4 mesh:

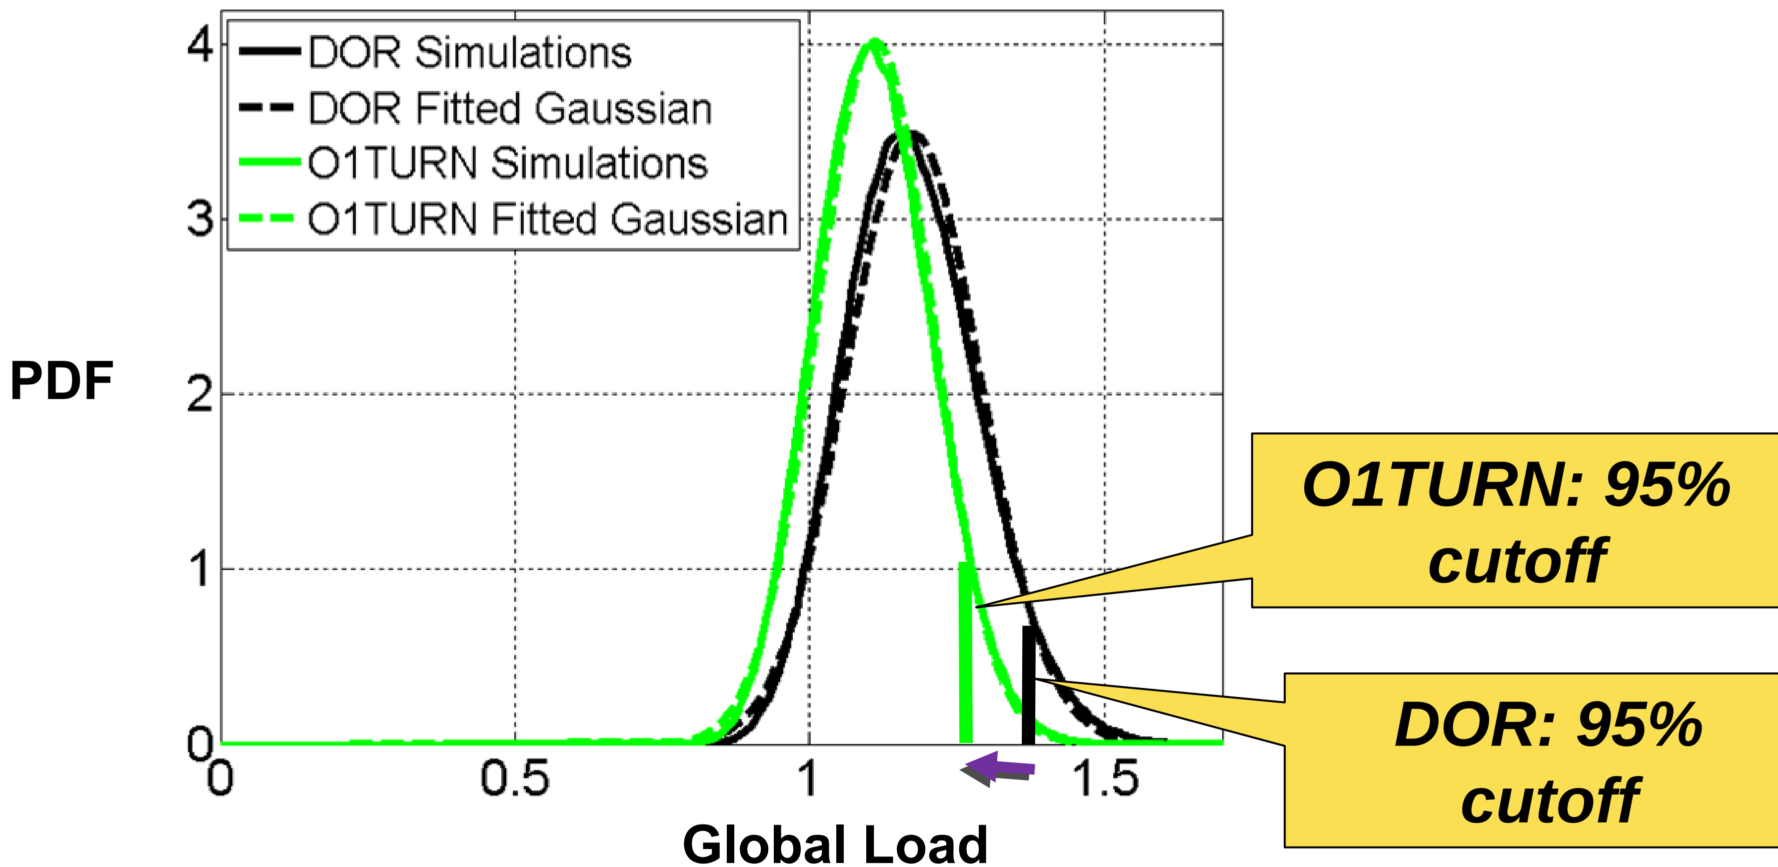

NUCA network

- More complex CMP architectures show similar results
  - **NUCA** (Non-Uniform Cache Architecture) with sharing degree 4.
  - Traffic model: each core (cache) may only send/receive traffic to/from caches (cores) in its sub-network.

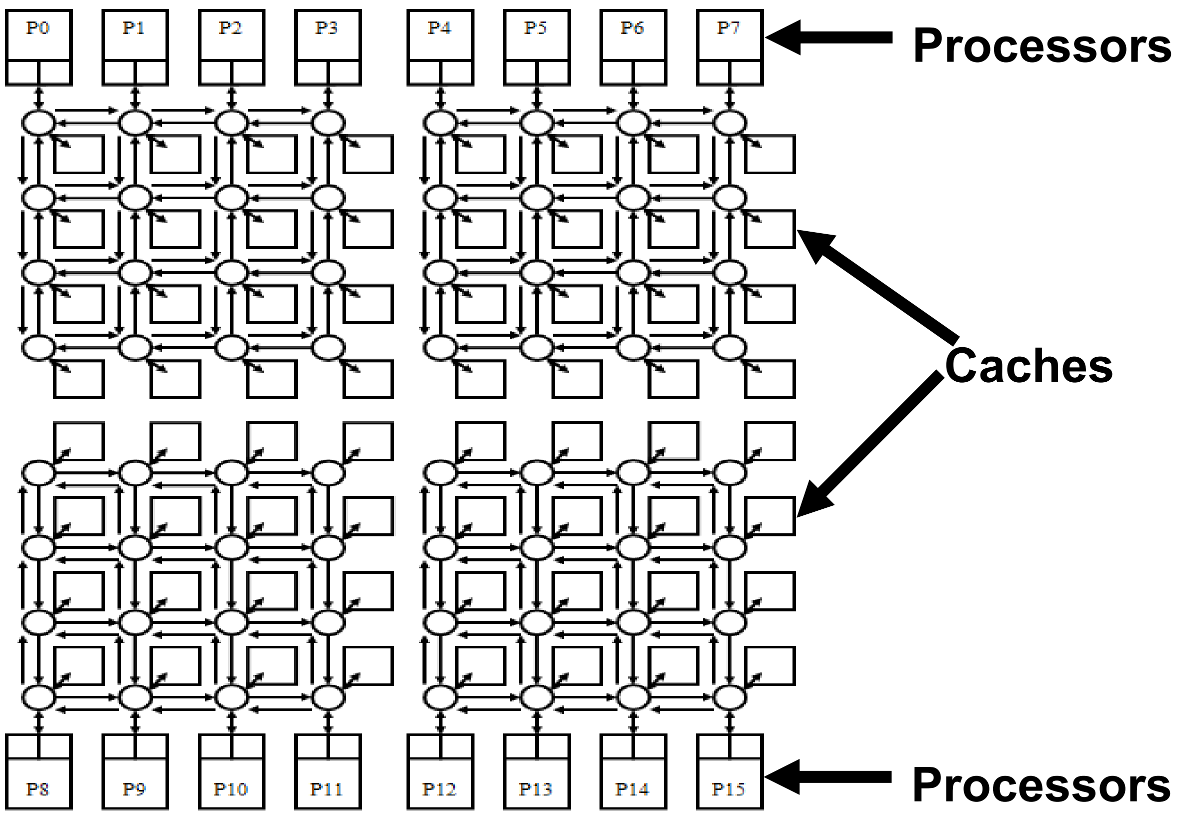

NUCA network – Total capacity

- Total capacity required for various Capacity Allocation (CA) targets.

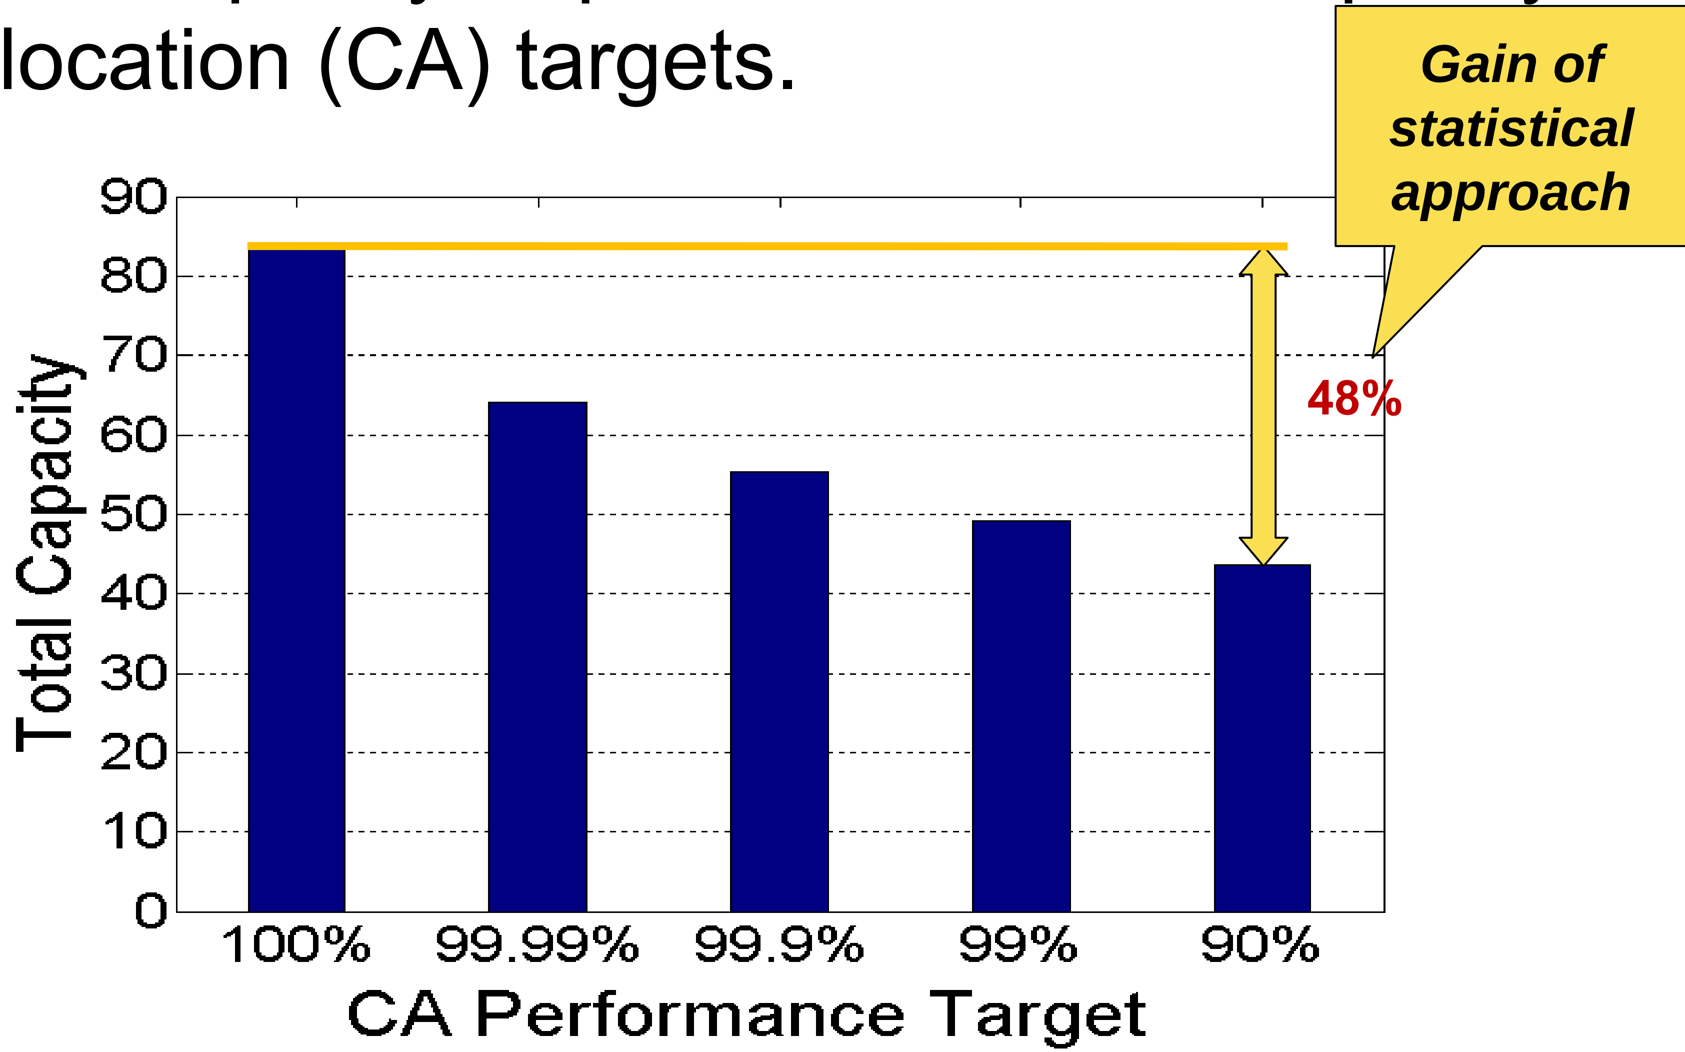

Summary

- Statistical approach to NoC design
  - Deals with several traffic matrices
  - Can optimize capacity allocation and routing

- Can be applied to any
  - Traffic matrix distribution
  - Topology
  - Oblivious routing
  - Link capacities

Thank you.