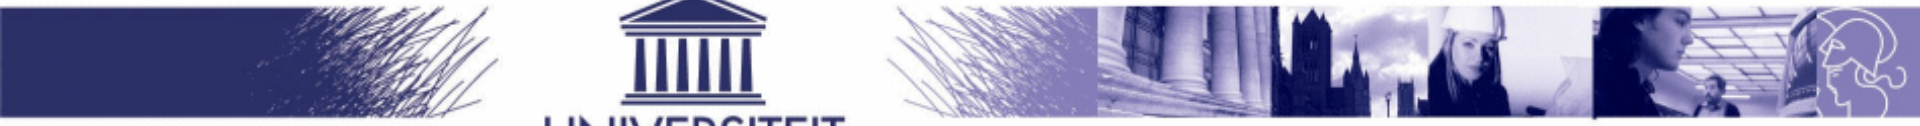

Synthetic Traffic Generation:
a Tool for Dynamic Interconnect Evaluation

W. Heirman, J. Dambre, J. Van Campenhout
ELIS Department, Ghent University, Belgium

Sponsored by IAP-V *PHOTON* & IAP-VI *photonics@be*,
Belgian Science Policy Office

Outline

- Introduction
- Synthetic traffic generation
- Results
- Conclusions

Distributed shared-memory architecture

Network is part of the memory hierarchy

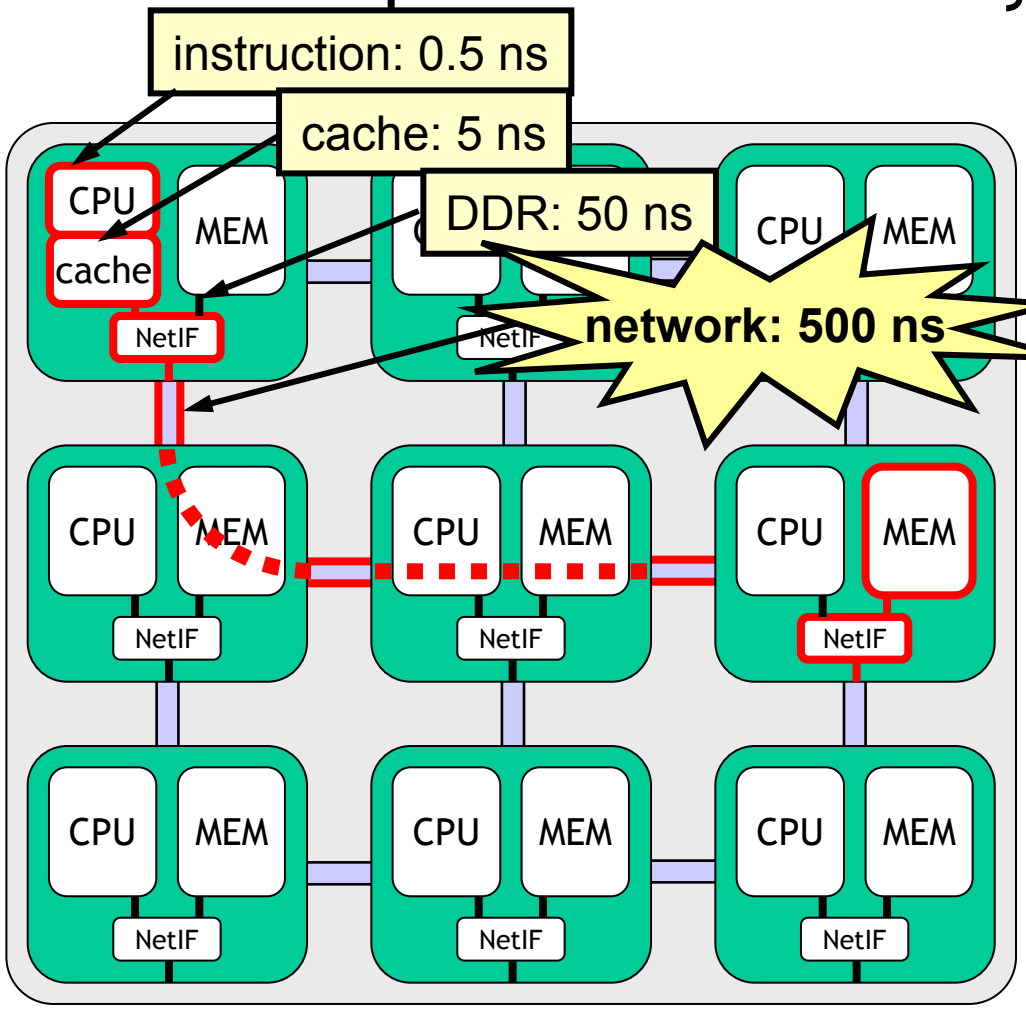

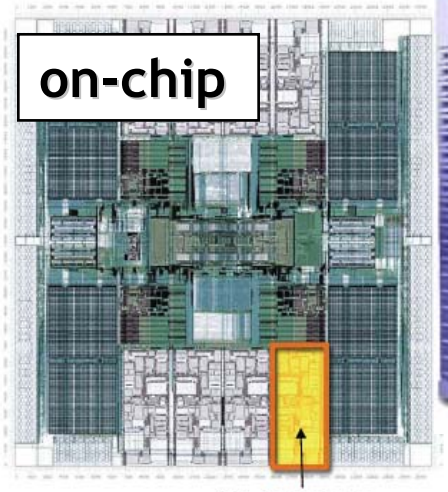

Interconnect requirements

Non-uni…ork traffic in space and time

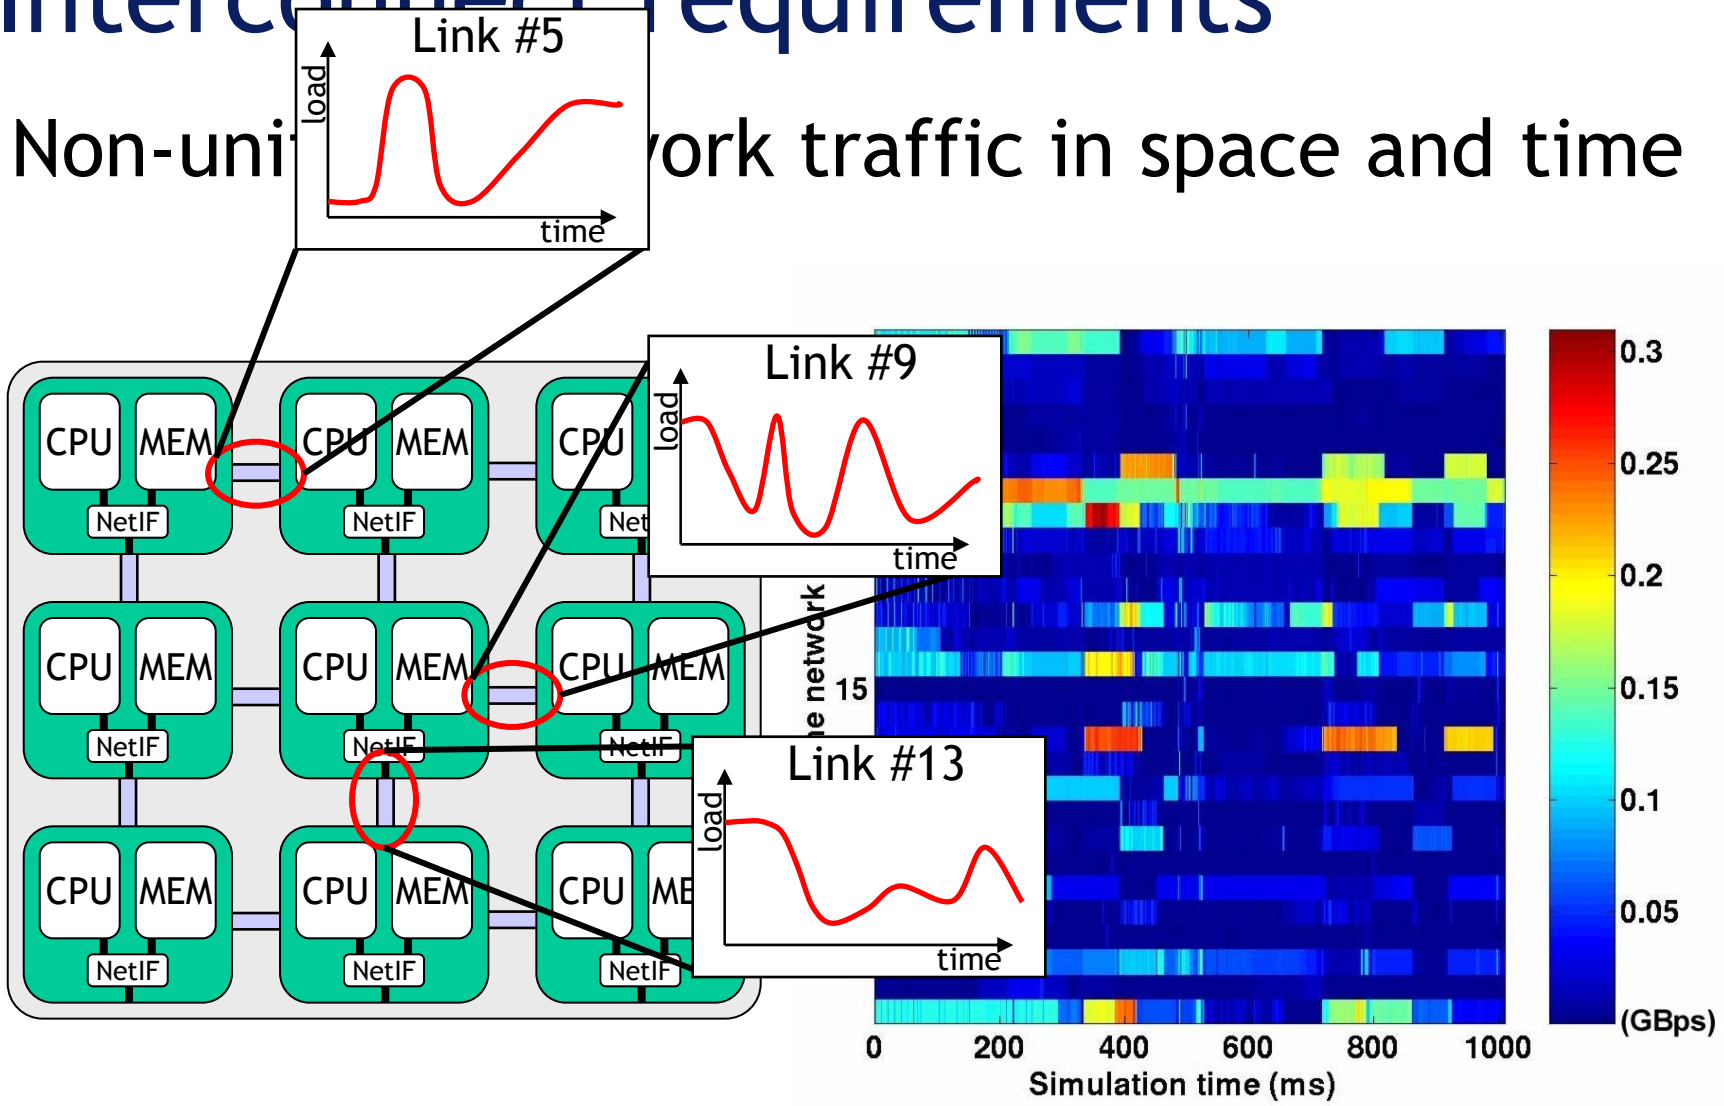

=> Reconfigurable network?

Reconfiguration implementation:
base network + extra reconfigurable links

other ‘dynamic networks’: e.g. per-link voltage scaling

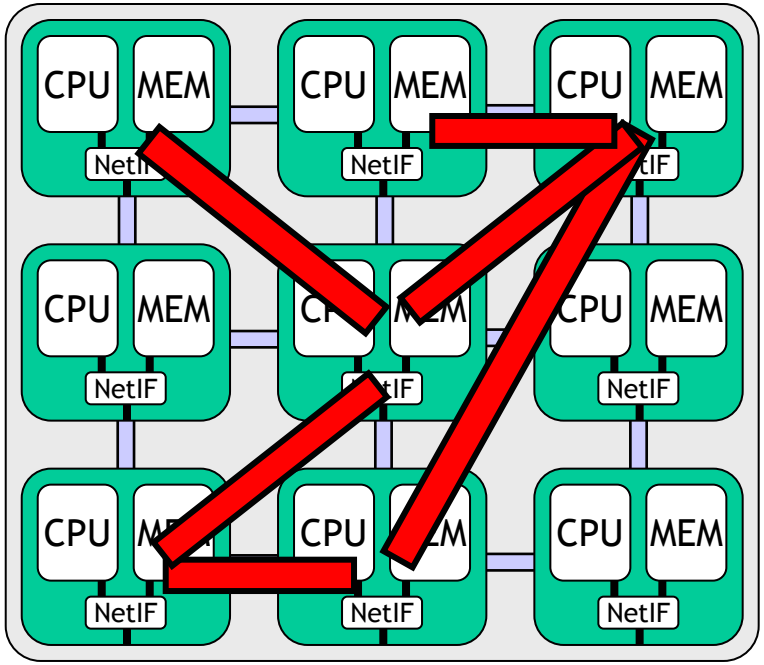

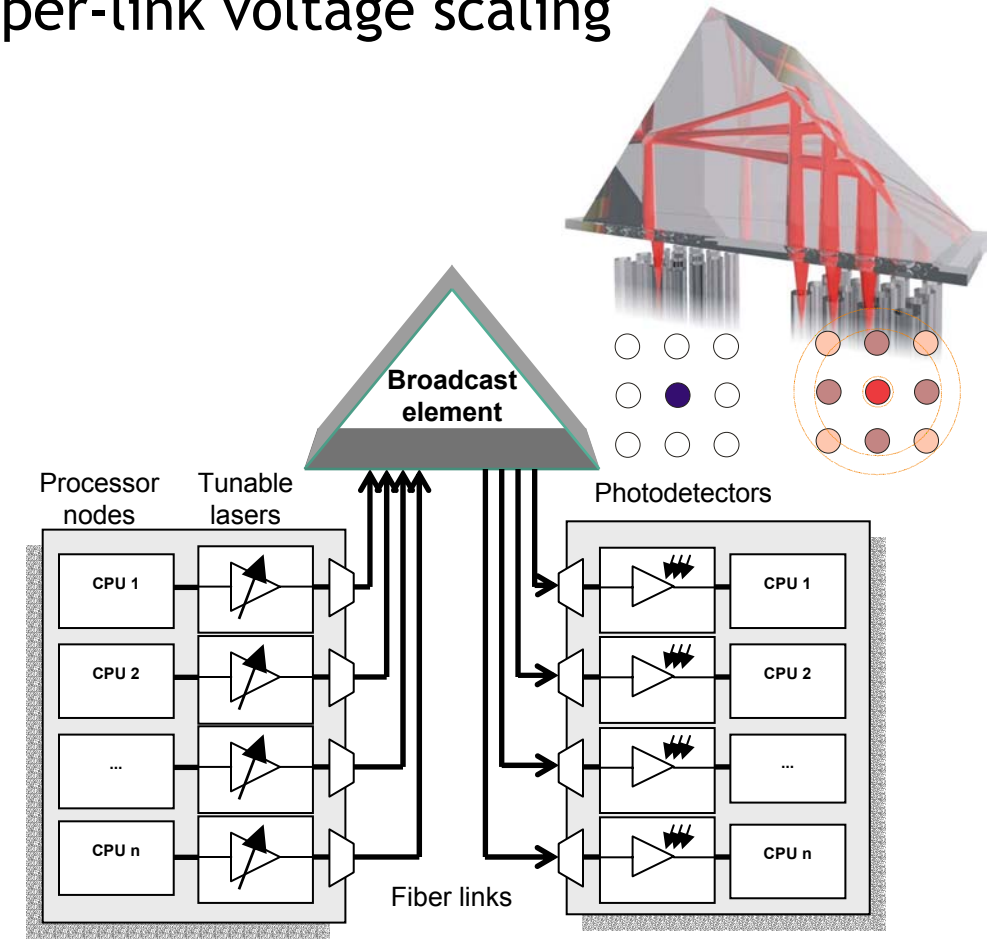

Evaluate networks with synthetic traffic

- Mimics the behavior of real traffic
- But without the computational cost of modeling application, OS, CPUs, caches, …

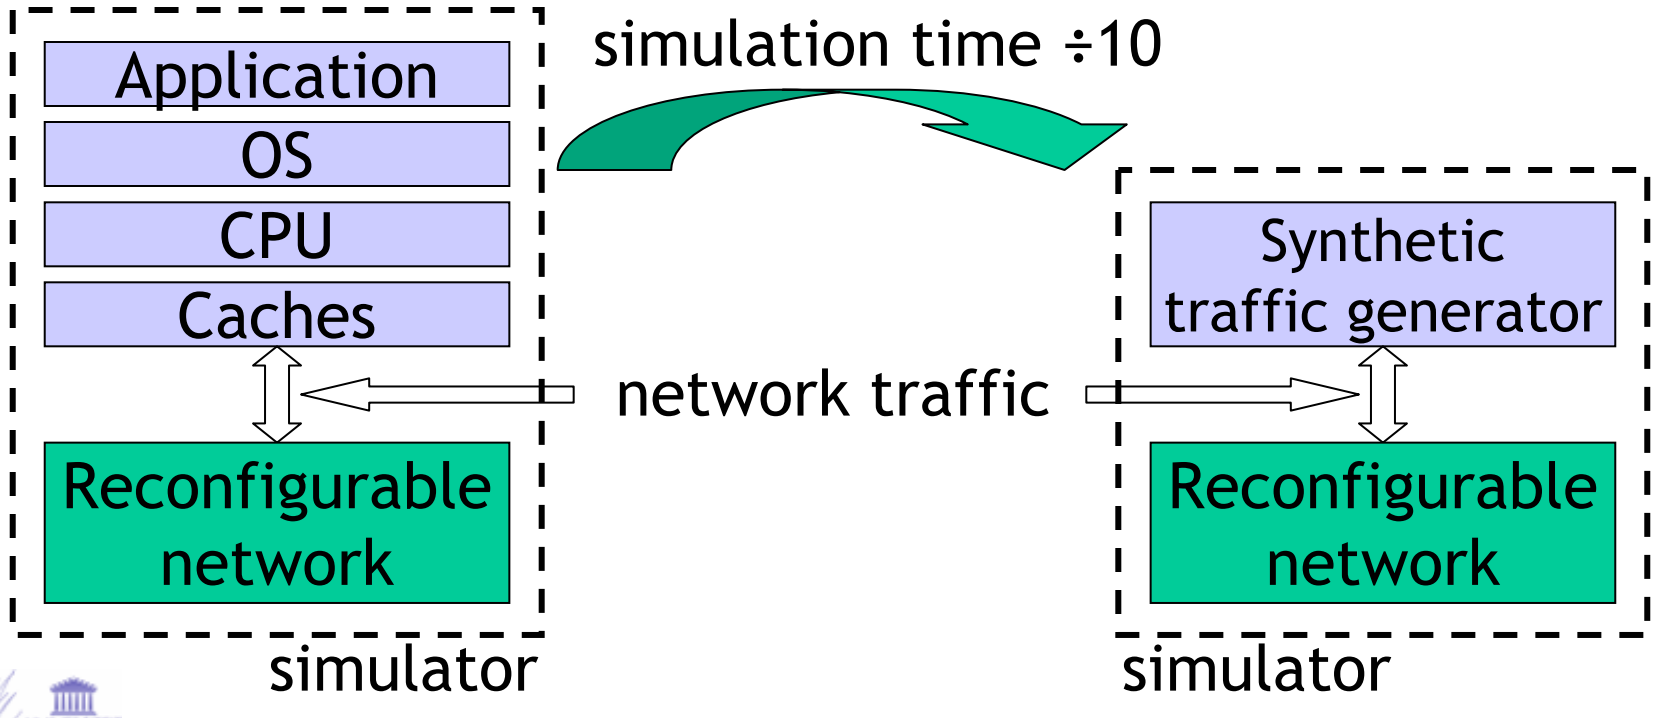

We need better synthetic traffic

Reconfiguration exploits low-frequency dynamics in the network traffic

- Trace-driven simulation using static traffic patterns (uniform, hotspot, shuffle, …) won't do!
- Full execution-driven simulation (traffic is driven by application: FFT, weather forecast, database) is too slow!

Outline

- Introduction
- Synthetic traffic generation
- Results
- Conclusions

Realistic synthetic traffic generation

- One execution-driven simulation
- Resulting traffic profile re-used many times

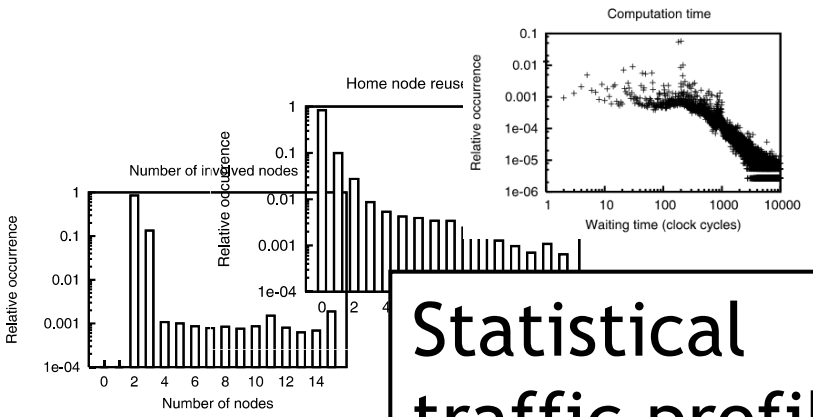

Statistical traffic profile

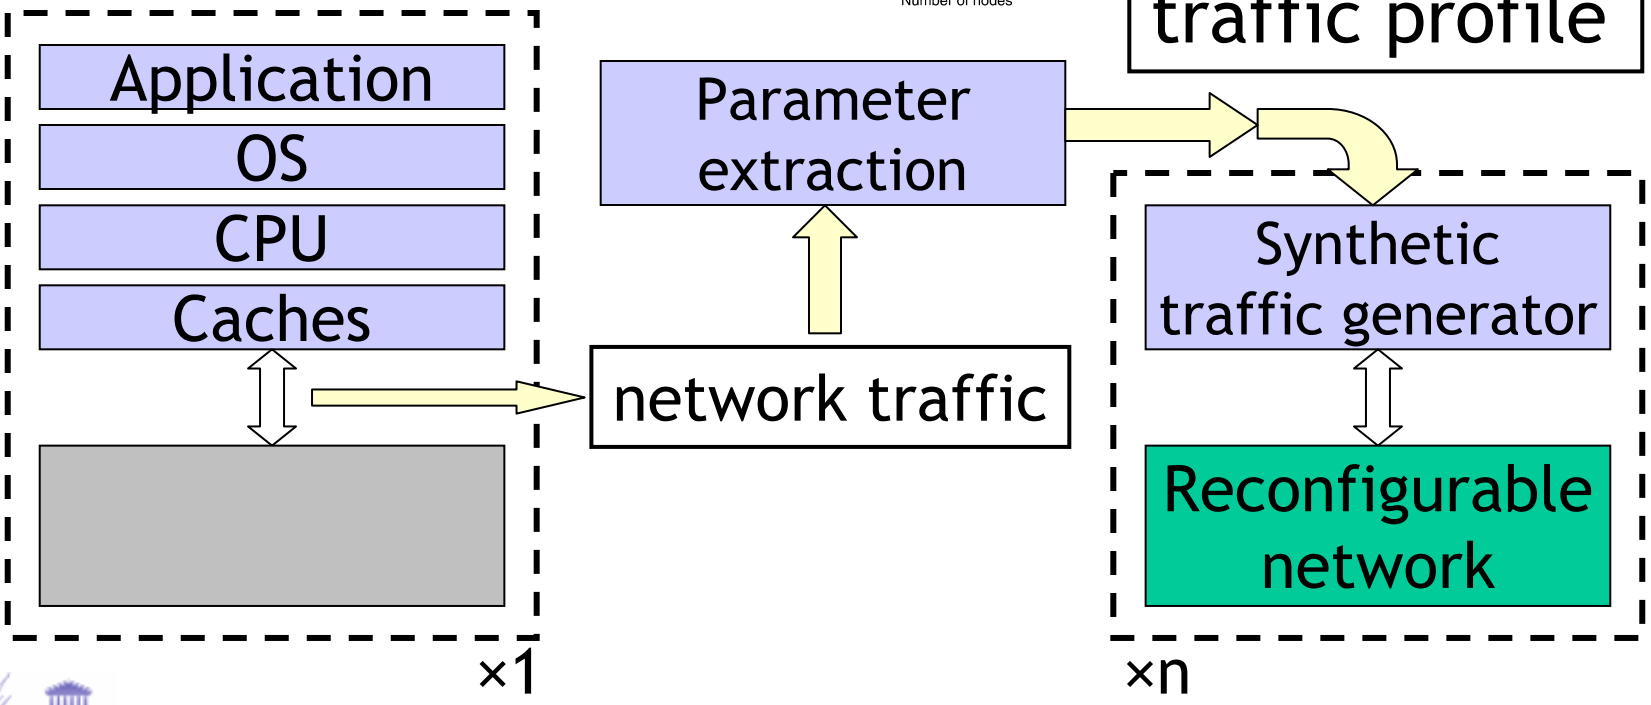

Preserve packet-interdependencies by using packet groups

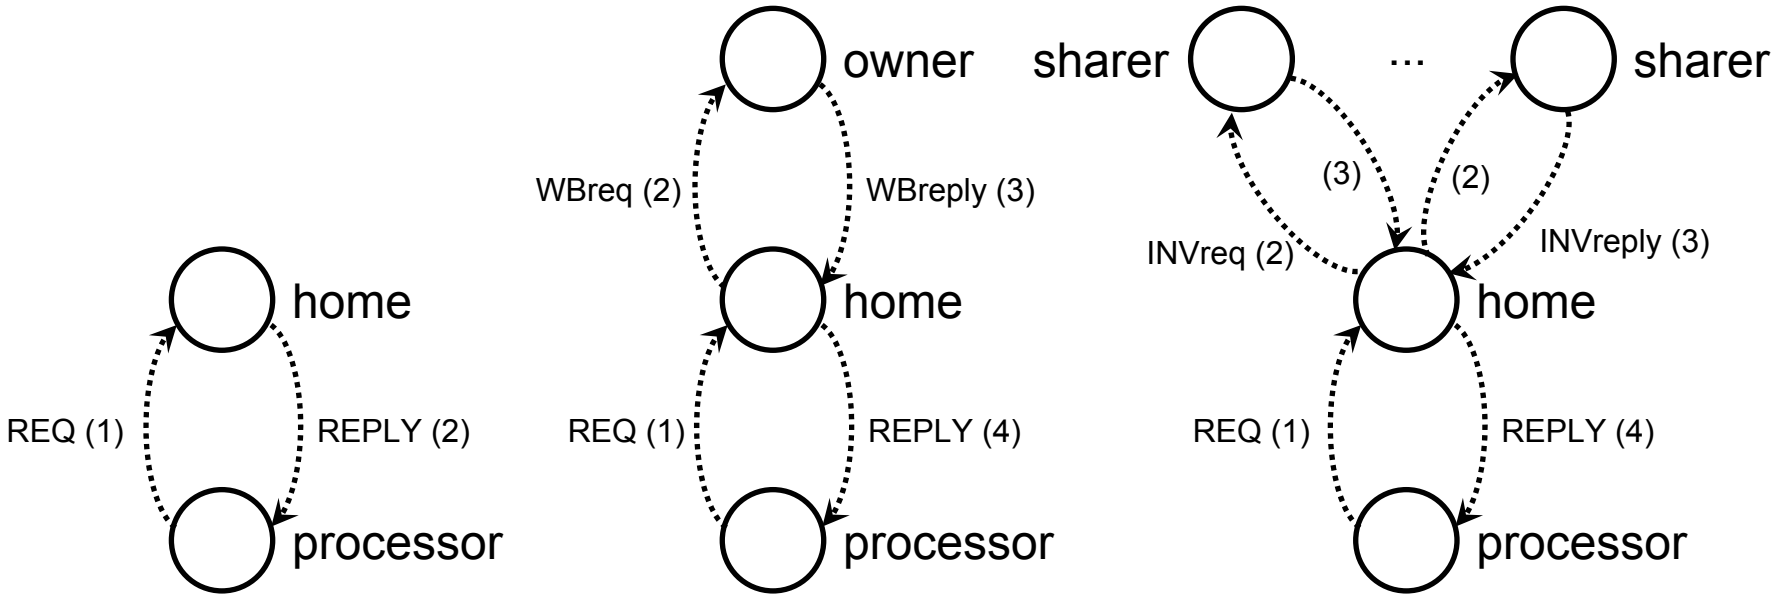

Packets are processed/generated in groups, corresponding to one memory operation each

Distribution of # involved nodes

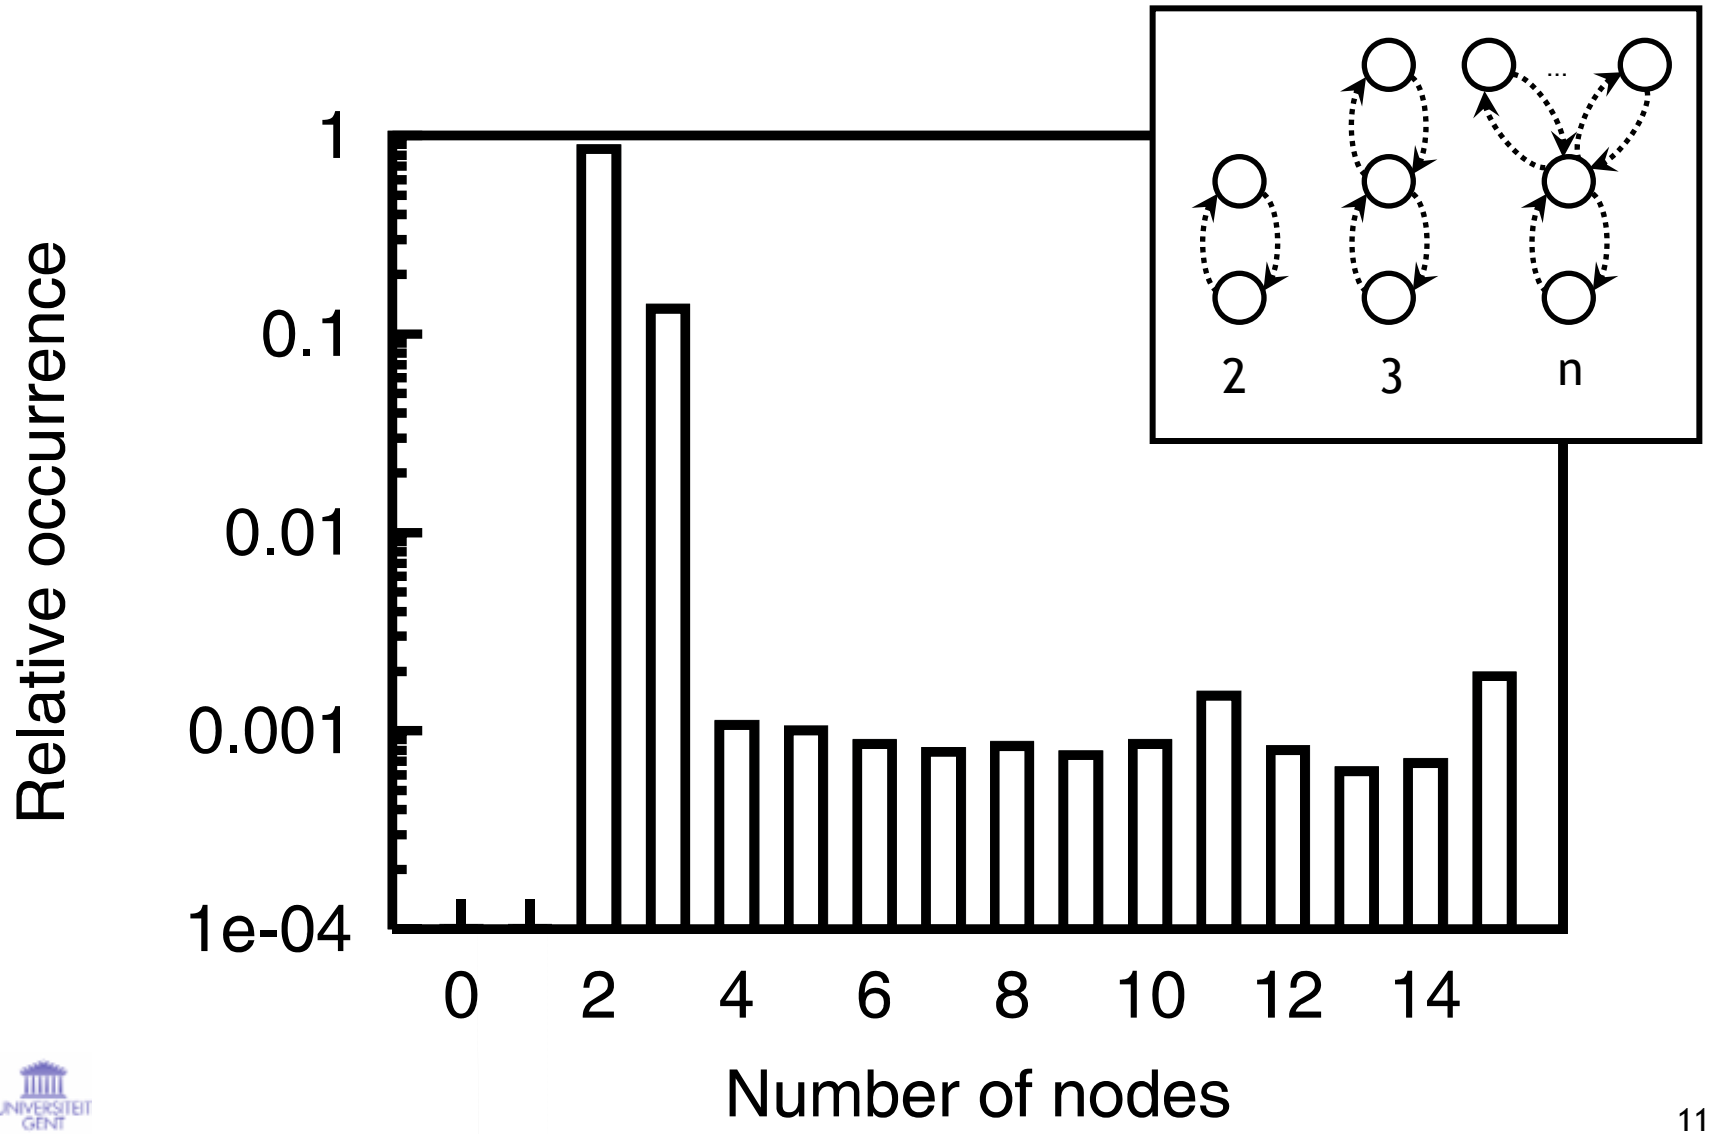

Reuse distance of home nodes: introduce locality

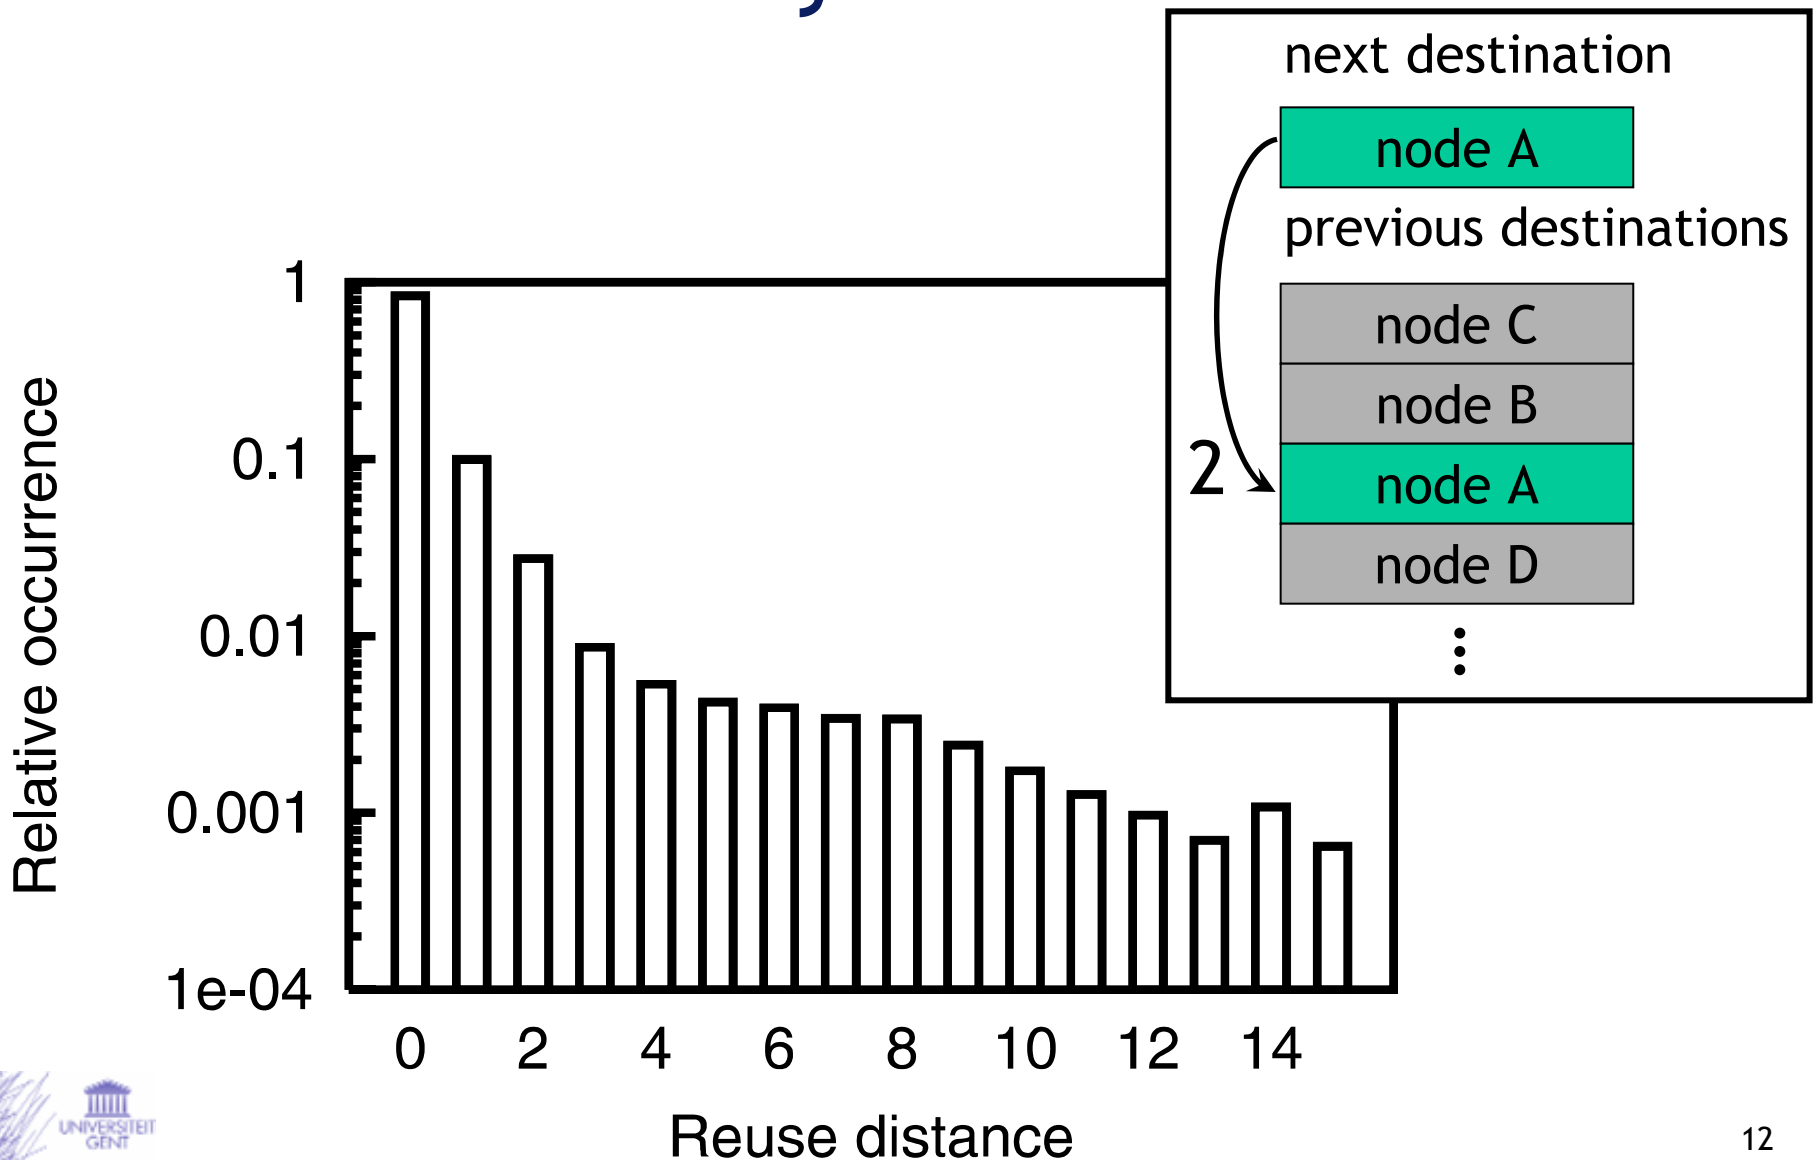

Computation or ‘think’ time

Models time delay between subsequent requests

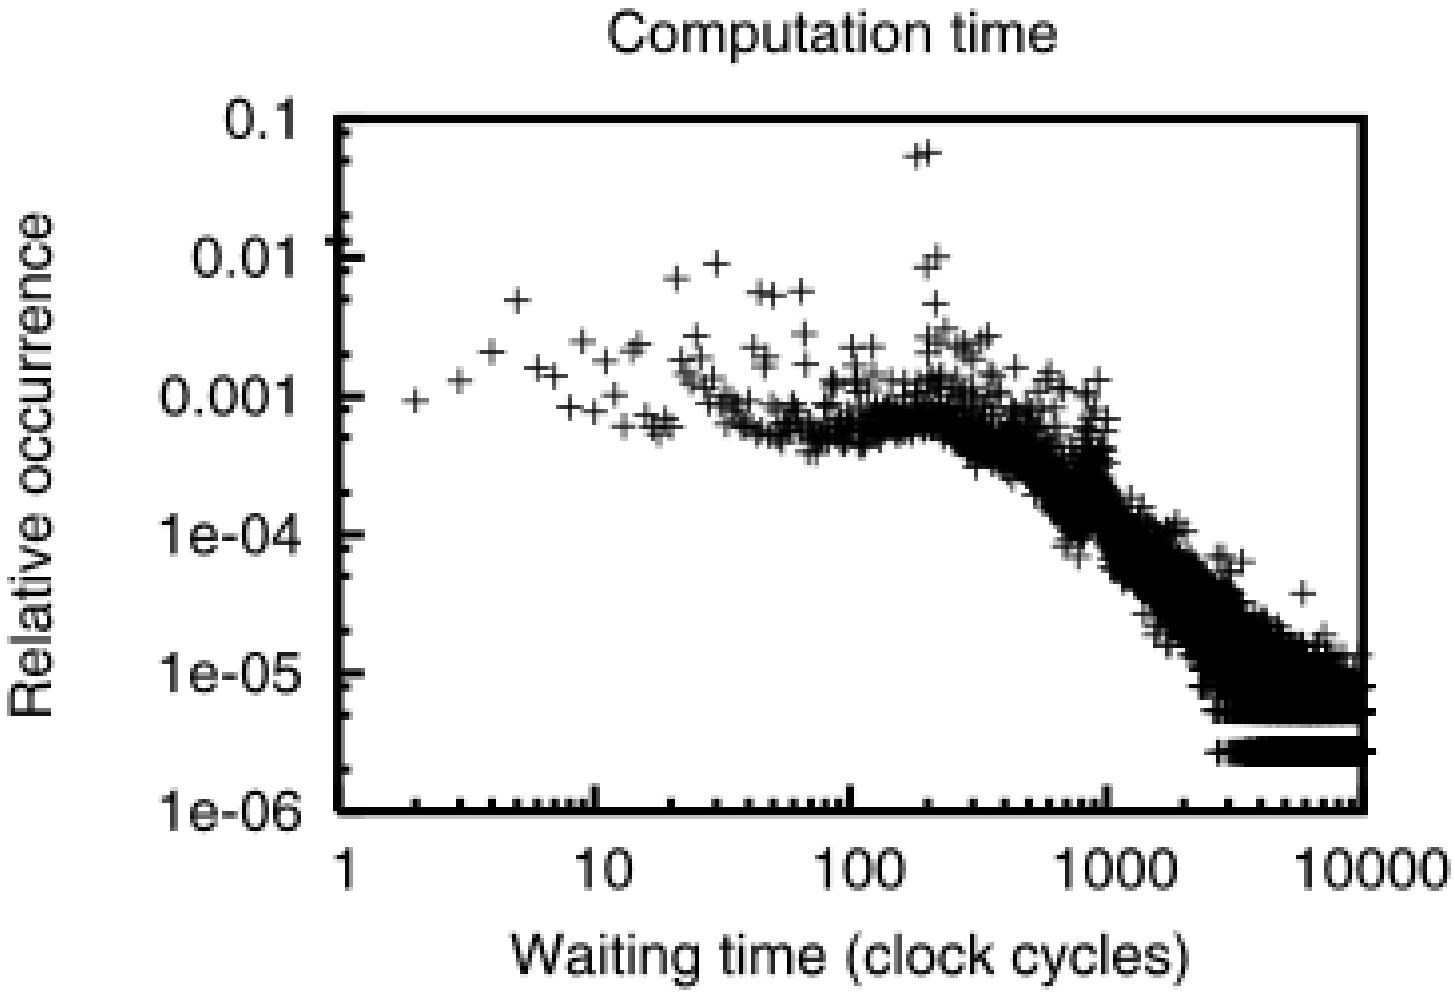

Outline

- Introduction
- Synthetic traffic generation
- Results
- Conclusions

Simulations

- Simulation platform: Simics, providing *functional* multiprocessor simulation
  - 16 UltraSPARC III processors
  - SPLASH-2 parallel benchmarks
- Timing model:
  - Computes the latency for each memory access
  - Models caches, interconnection network
  - Base network: 4x4 torus
  - Extra links: configurable number, fan-out, reconfiguration interval

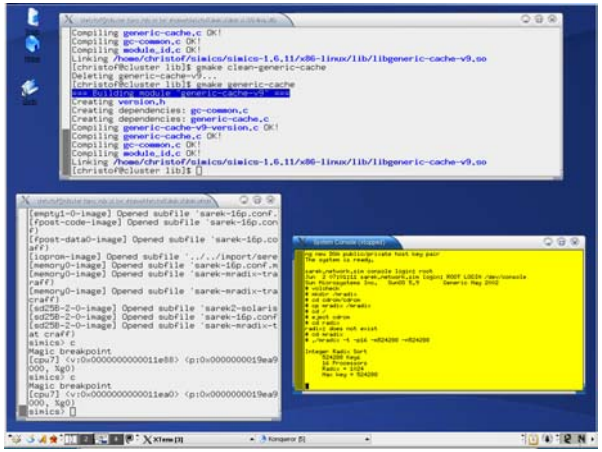

Simulations

Once per benchmark:

Simulate execution of the benchmark, base network only, measuring traffic profile (1)

For each set of extra link parameters:

- Execution-driven simulation with reconfigurable network (2)
  ➔ *"correct" result*
- Trace-driven simulation using (simplified) traffic from (2)
  ➔ *tracing error*
- Trace-driven simulation using (simplified) traffic from (1)
  ➔ *traffic-dependence on network*
- Trace-driven simulation using synthetic traffic
  ➔ *total error*

Several parameters can be measured

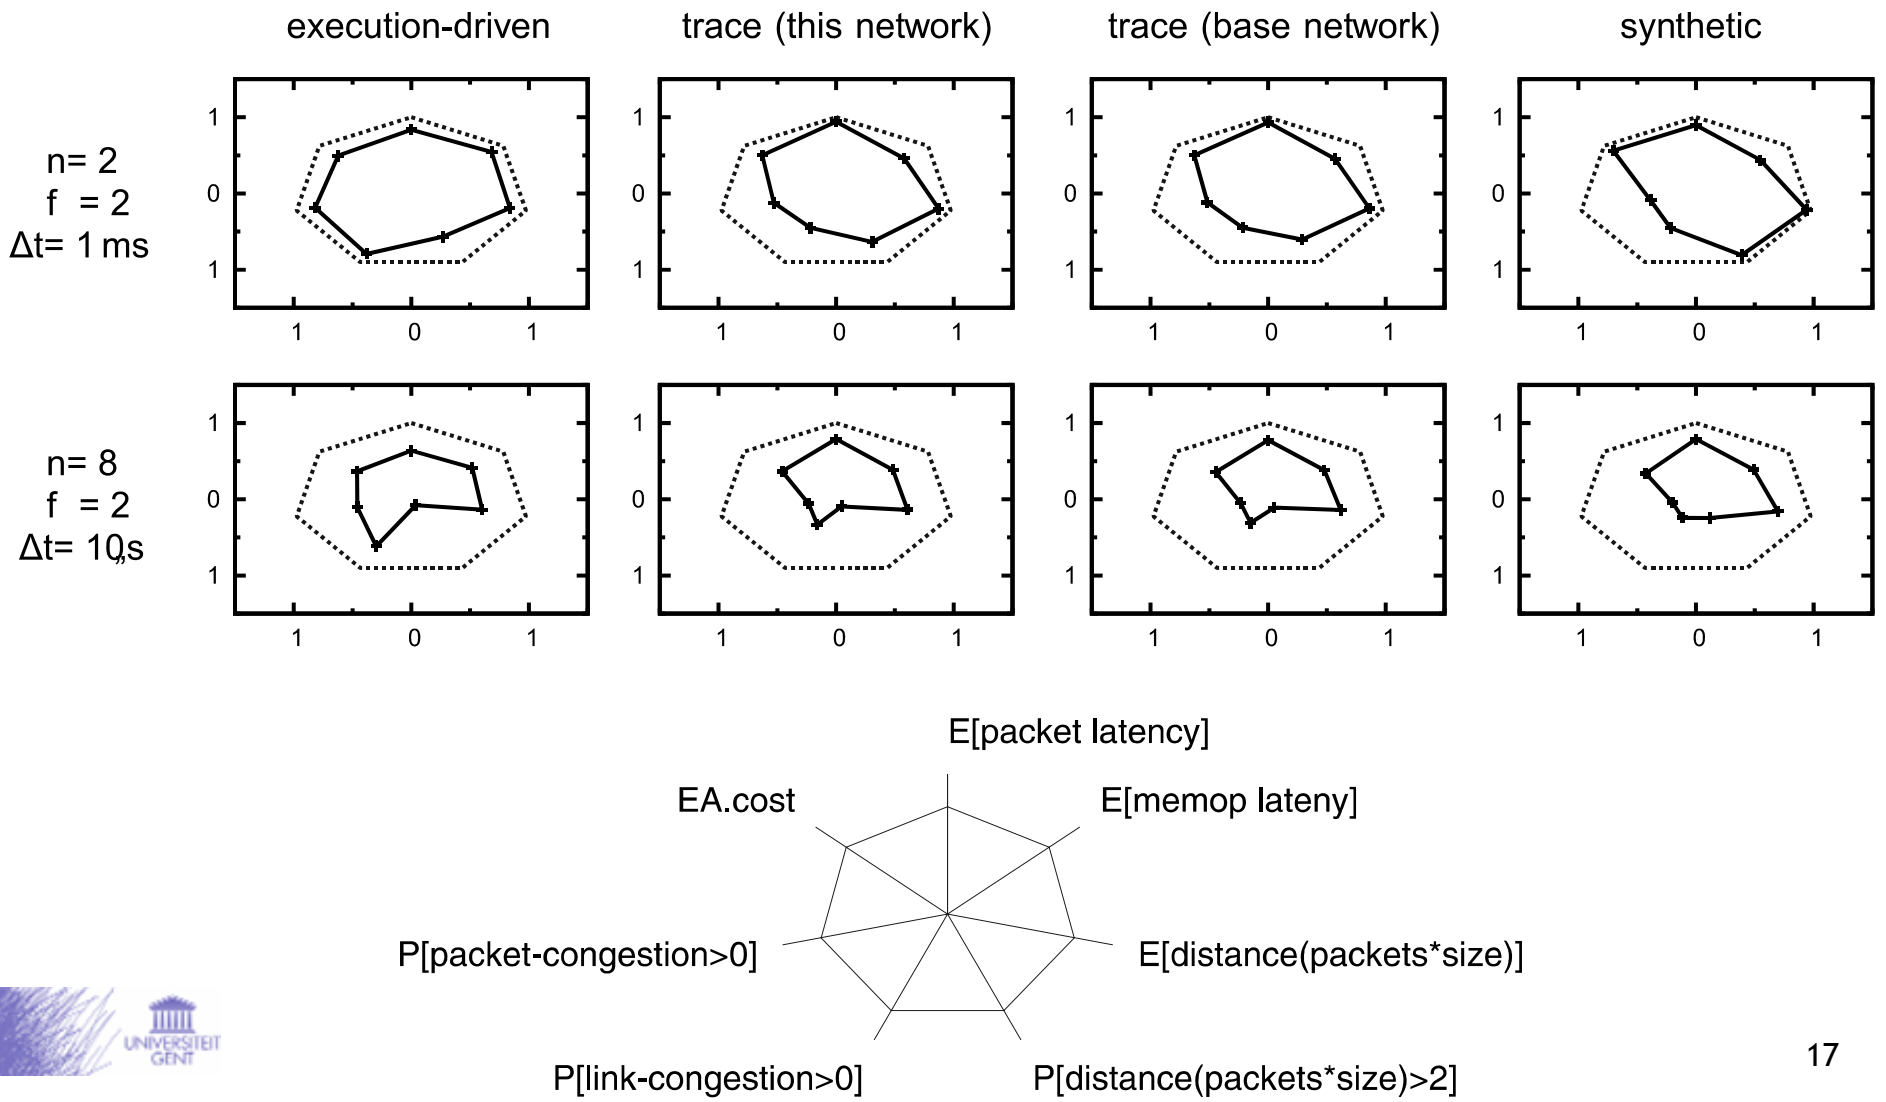

Detailed view of "average packet latency"

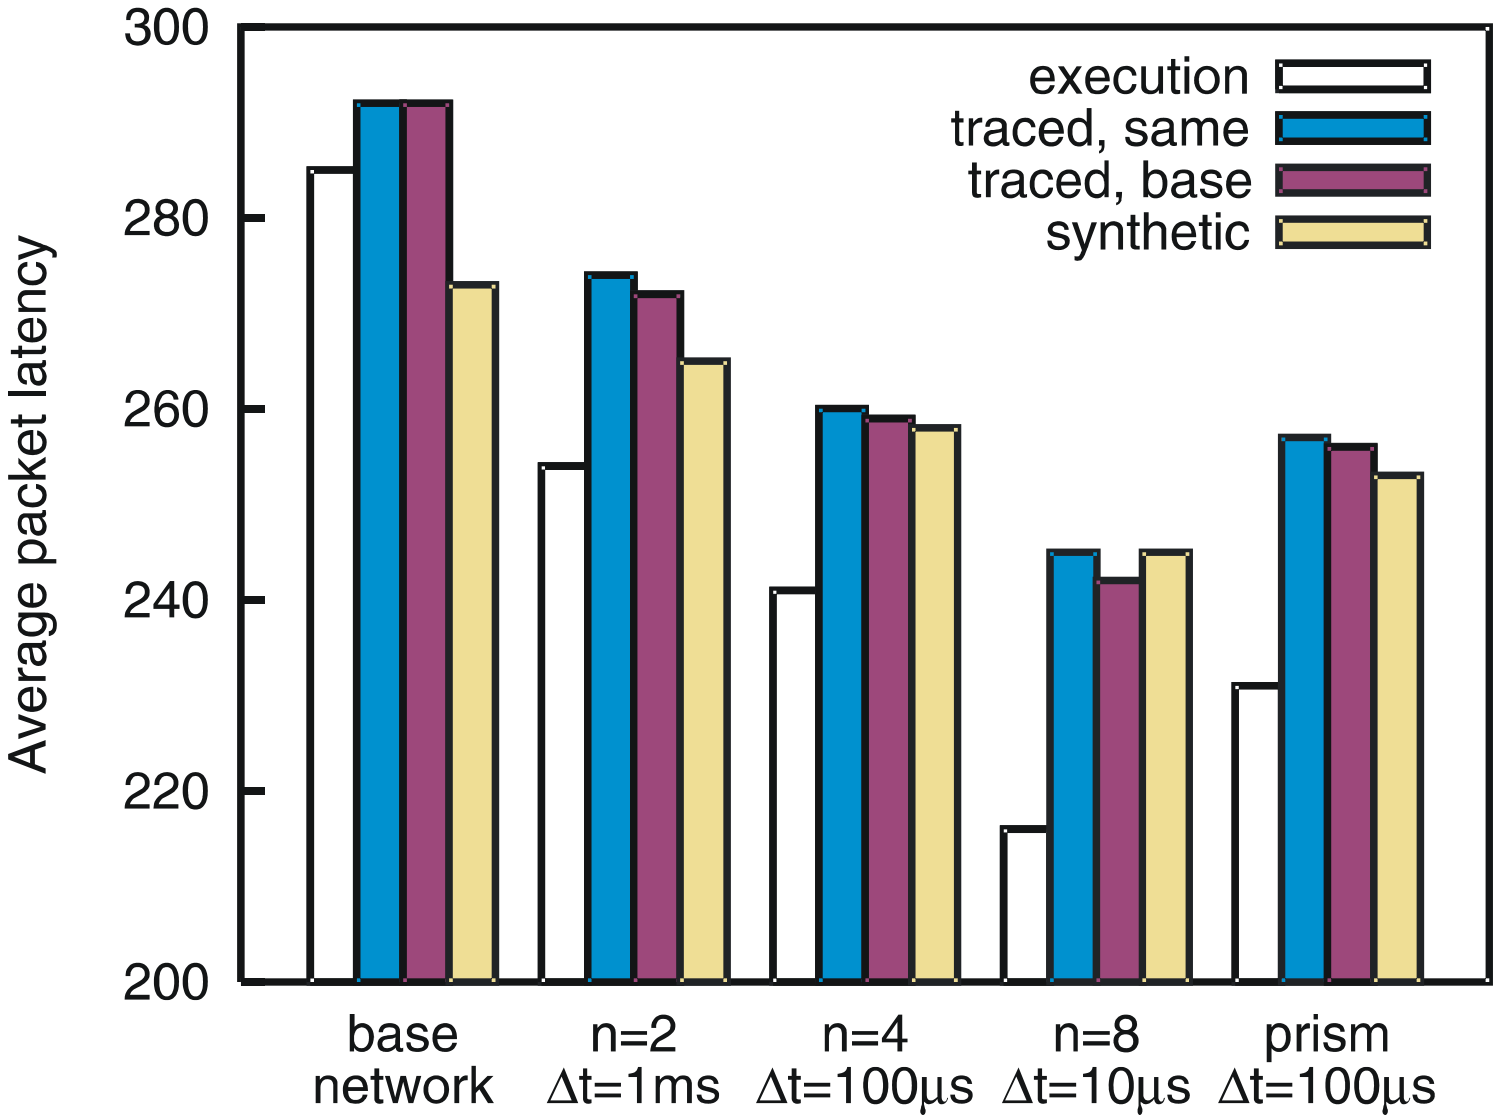

Variability for shorter traces

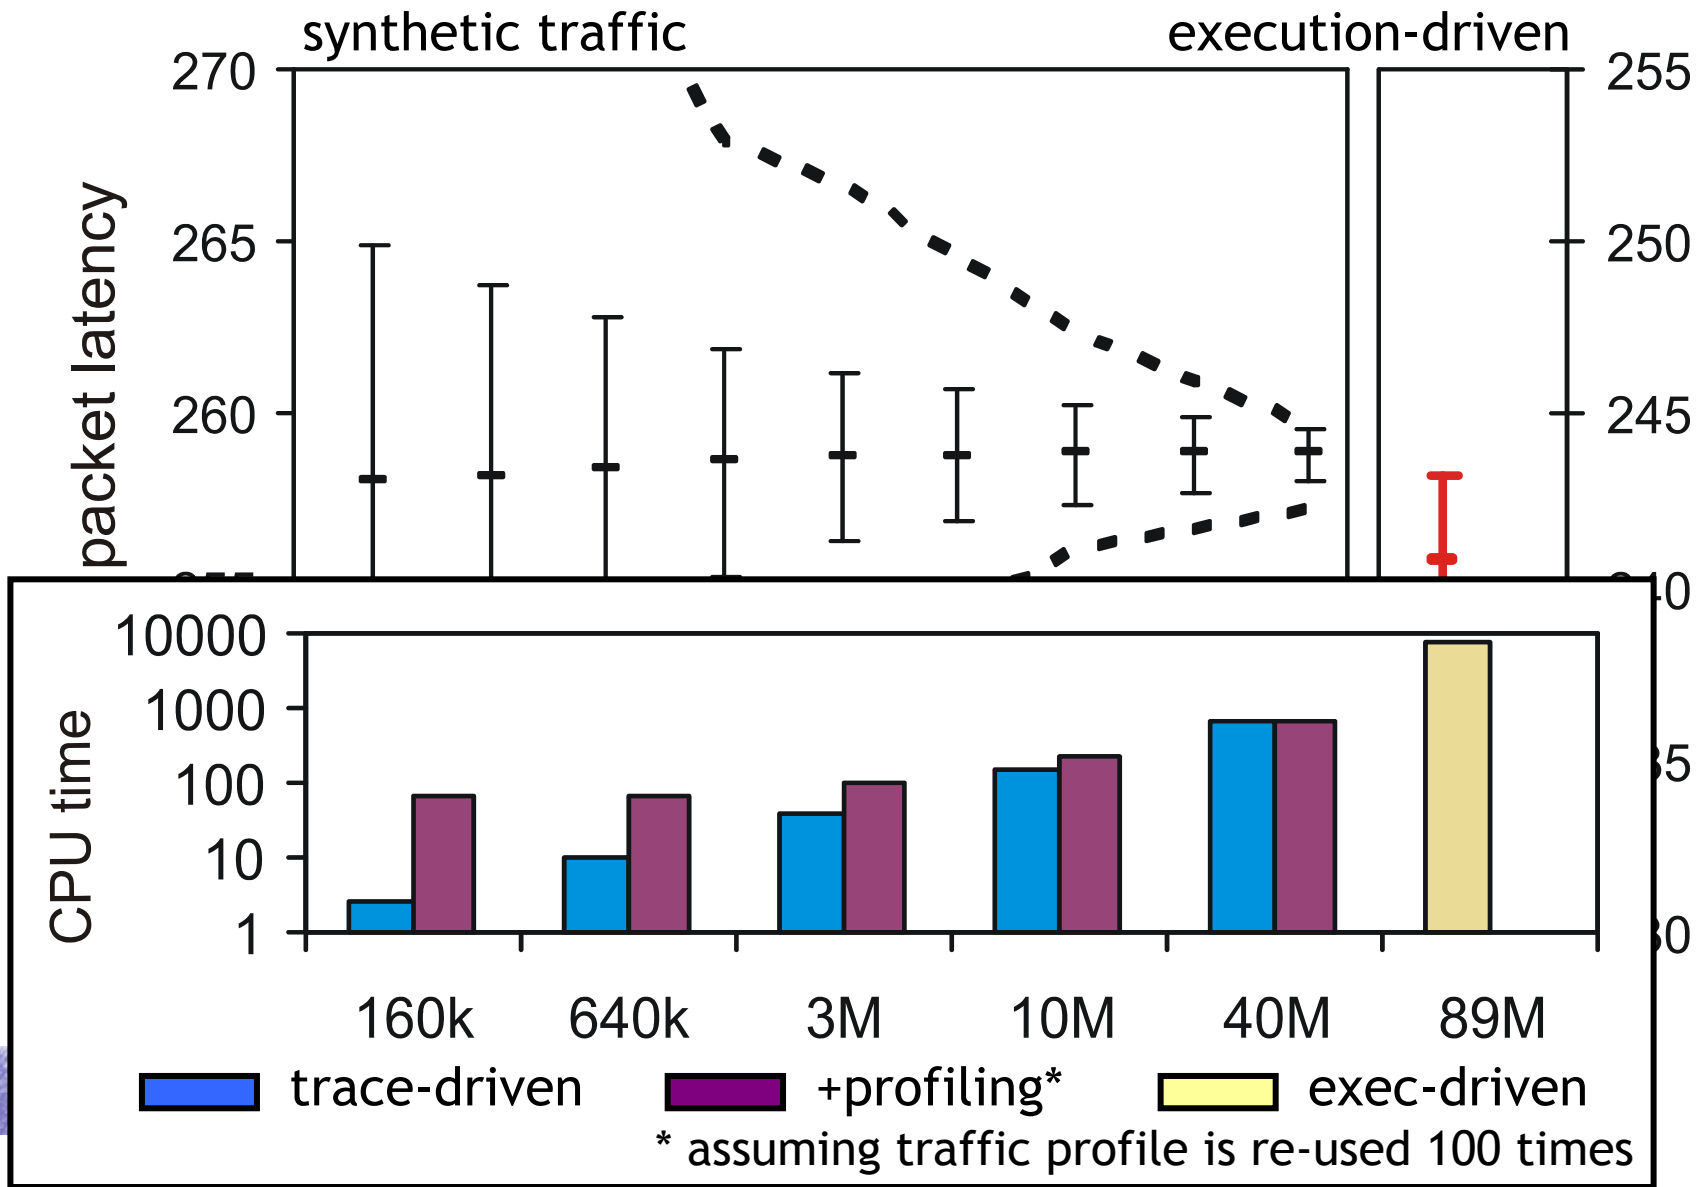

Outline

- Introduction
- Synthetic traffic generation
- Results
- Conclusions

Conclusions

- Synthetic traffic generation was extended to
  - shared-memory cache-coherence protocols,
  - reconfigurable networks
- Good relative accuracy for different network topologies
- Much less computationally expensive (x10), even more so for shorter traces (x100)
- Reproducibility equal to or better than execution-driven simulations