UNIVERSITEIT GENT

Rent's Rule and Parallel Programs:
Characterizing Network Traffic Behavior

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

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

Outline

- Introduction
- Rent's rule & traffic locality
- Time-varying network traffic
- Conclusions

UNIVERSITEIT GENT

Evolution of Systems design

- VLSI systems get ever more complicated
- More software, processor IP blocks, hardware/software co-design
- Ad-hoc global wiring → Network-on-Chip (“communication IP block”), long wires → packets
- What with Rent’s rule?

UNIVERSITEIT GENT

Rent's rule: power law relation

components (G)
vs.
terminals (T)

$T = tG^p$ [1]

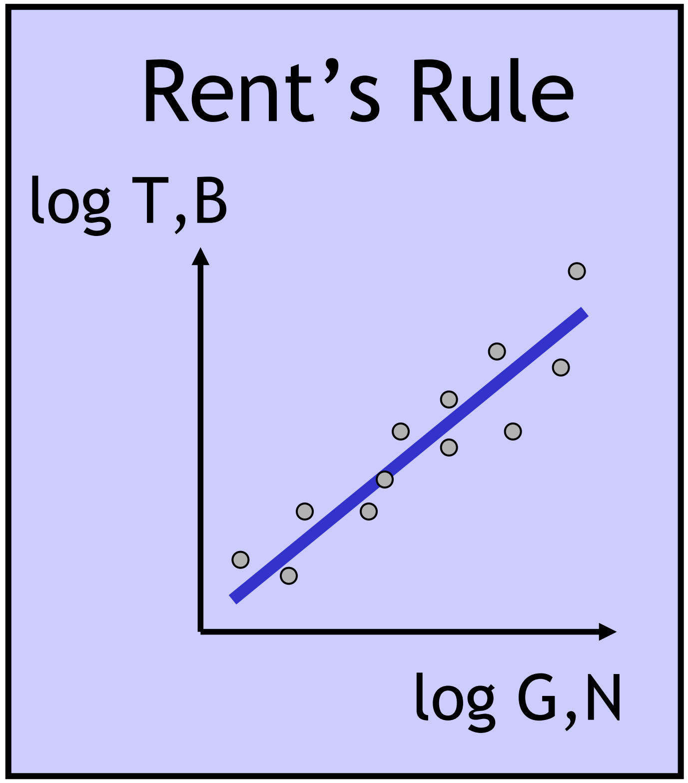

processors (N)
vs.
bandwidth (B)

$B = bN^p$ [2]

circuits
wires

processor cores
networks-on-chip

[1] Landman and Russo, IEEE Trans. on Computers, 1971
[2] D. Greenfield et. al, NOCS 2007.

Multiprocessor + Network architecture

Shared memory: network is part of memory hierarchy

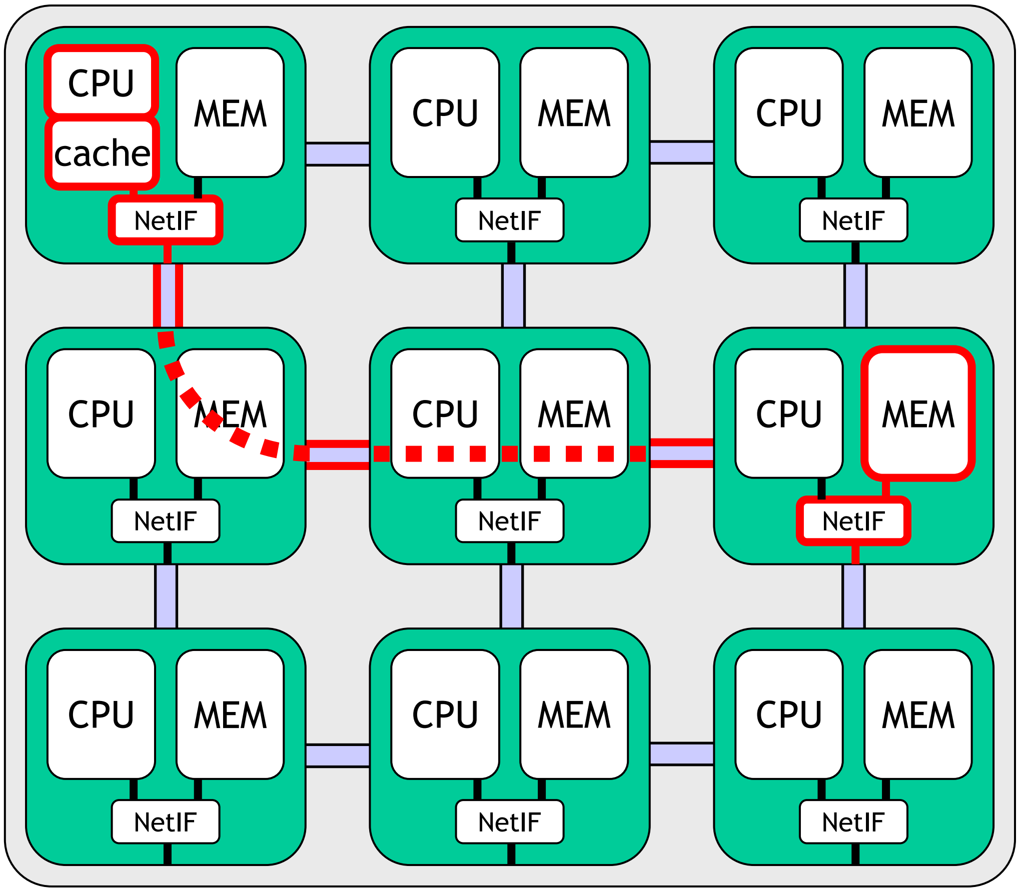

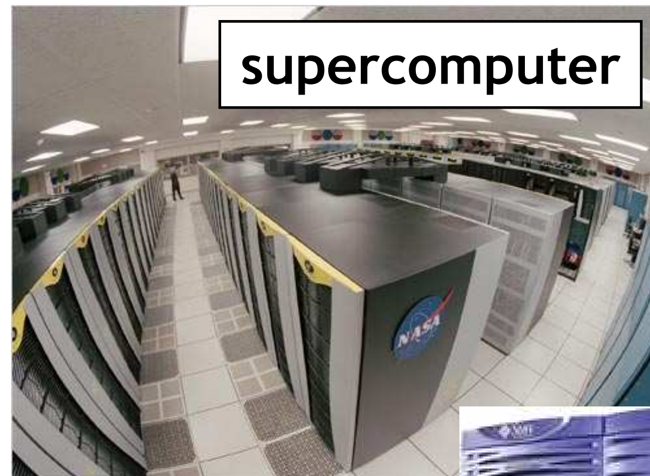

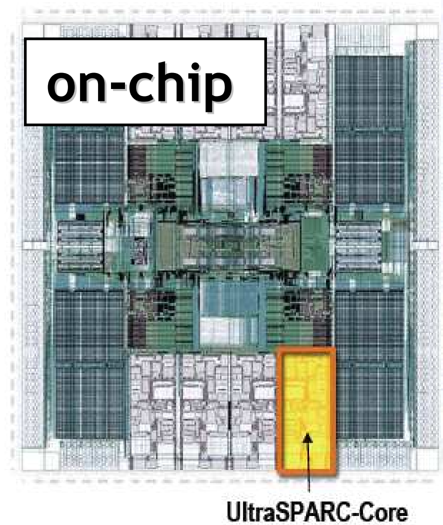

NoC design: problems and opportunities

- Simple traffic models: uniform, hot-spot, fixed bandwidth distribution
  - Ignores locality, time-variance in network traffic
  - Yields non-optimal NoC designs (uniform vs. non-uniform, static vs. reconfigurable)
- Opportunity: better traffic models, analytical tools vs. trial-and-error

Outline

- Introduction
- Rent's rule & traffic locality
- Time-varying network traffic
- Conclusions

Partitioning nodes by communication intensity

- Hierarchically partition nodes according to communication (hMETIS)
- Just as for wires, but:
  - Communication graph is usually fully connected
  - Weight on each connection
    = total communication between node pair
- Fit power law on
  (cluster size, bandwidth) distribution

UNIVERSITEIT GENT

Rent exponent

measured Rent exponent (dependent on application):

- 16 nodes: .55-.65
- 64 nodes: .66-.74

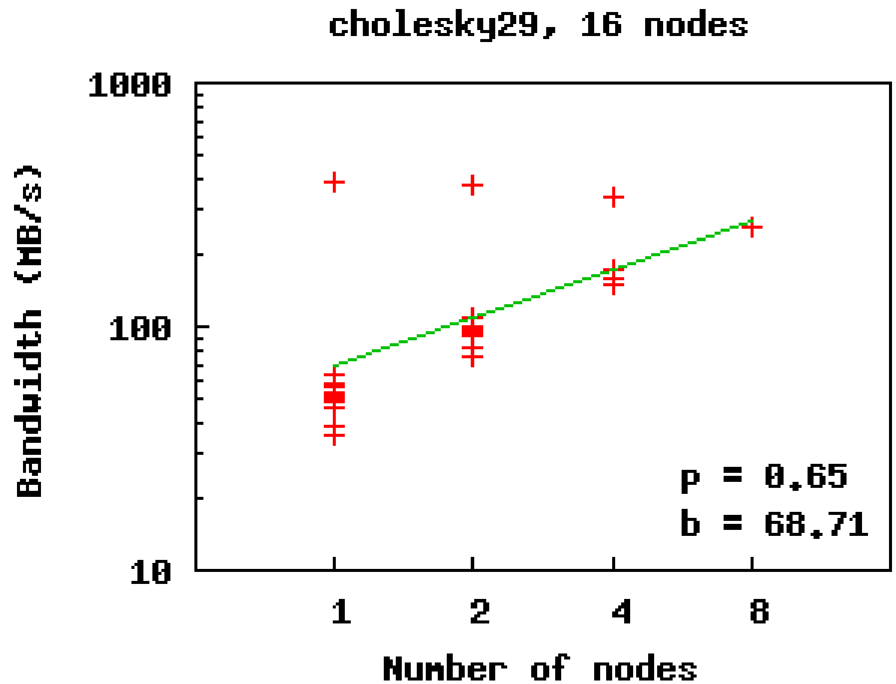

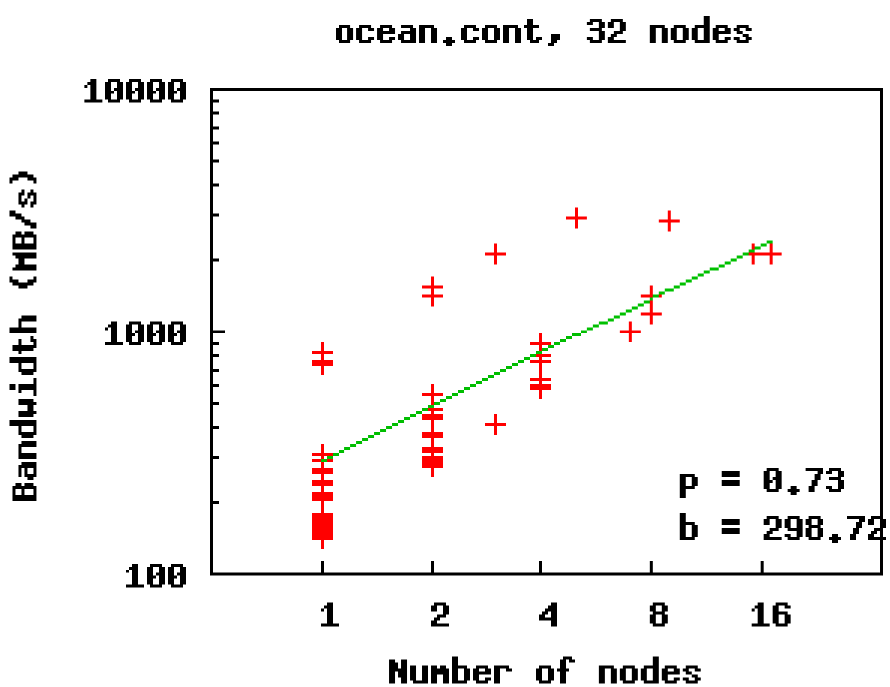

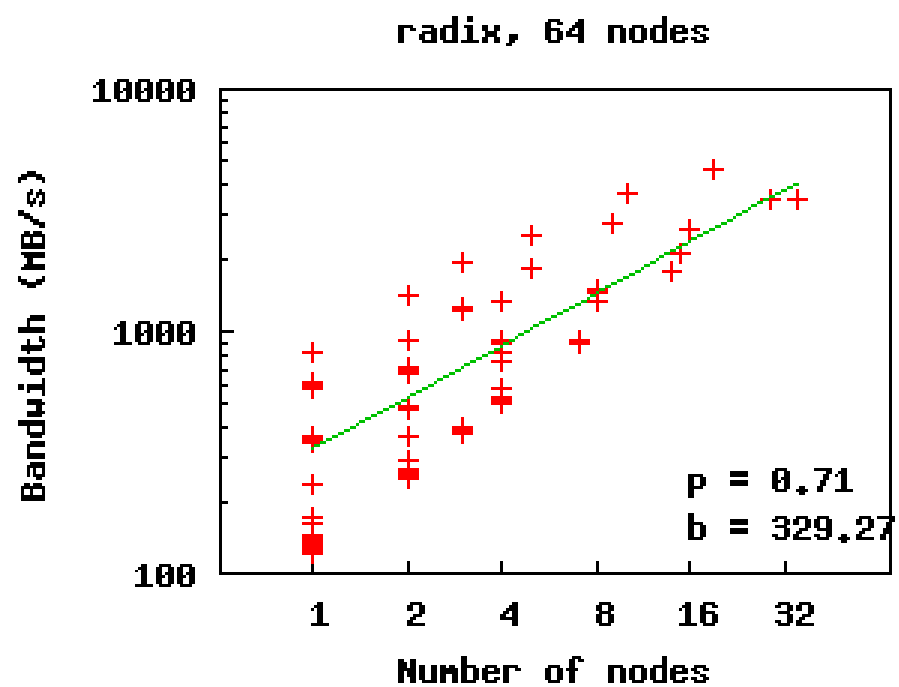

"Wire length" distribution

Distribution of communication vs. distance

distance(A, B) =
log2(size of smallest cluster containing both A and B)

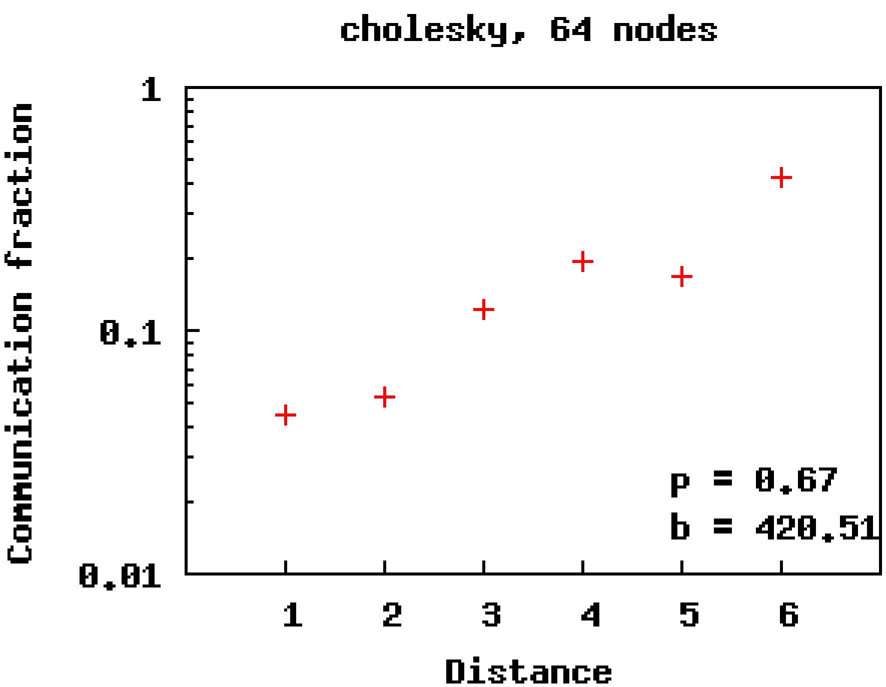

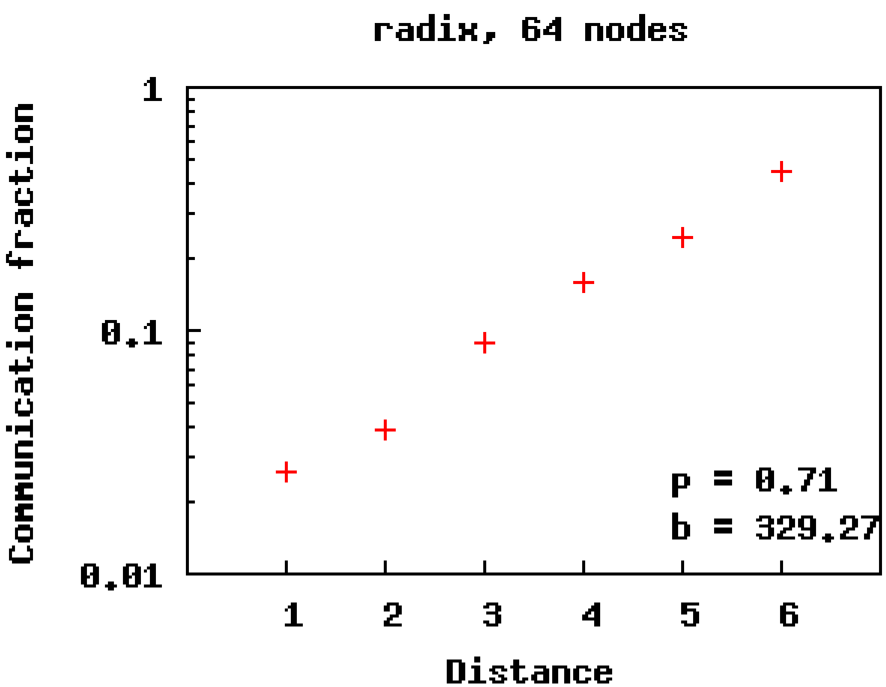

Outline

- Introduction
- Rent’s rule & traffic locality
- Time-varying network traffic
- Conclusions

Communication varies through time

- Hardware:
  - fixed function
  - traffic remains similar through time
- Software:
  - more complex, different phases (e.g. function call)
  - communication patterns can change trough time

Communication varies through time

- Repeat partitioning per interval of 100k clock cycles
- Periods of high and low communication alternate
- Rent exponent badly defined during periods of low communication

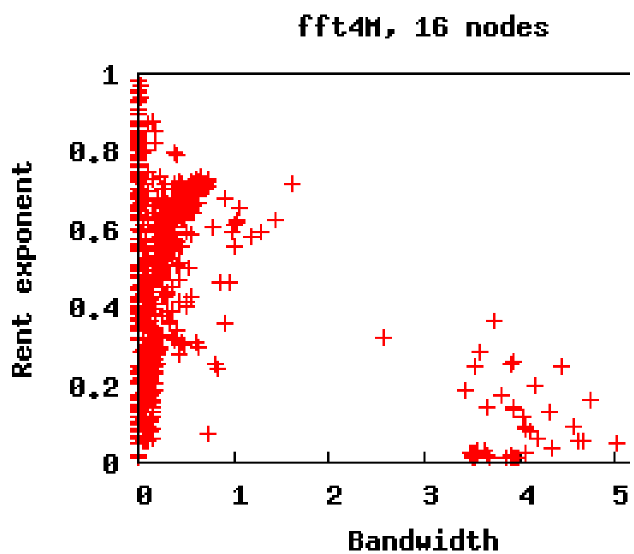

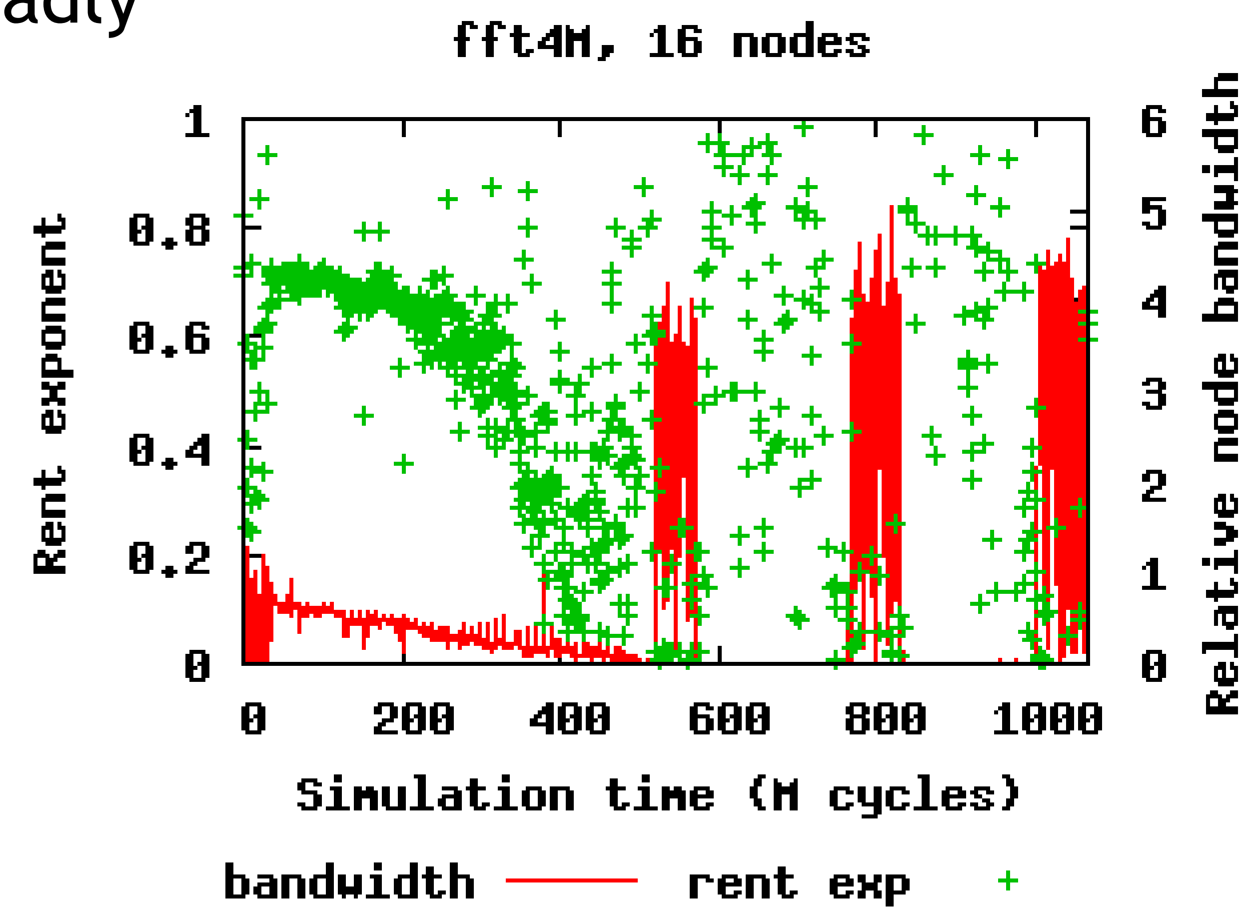

Communication varies through time

- Repeat partitioning per interval of 100k clock cycles
- Periods of high and low communication alternate
- Rent exponent badly defined during periods of low communication

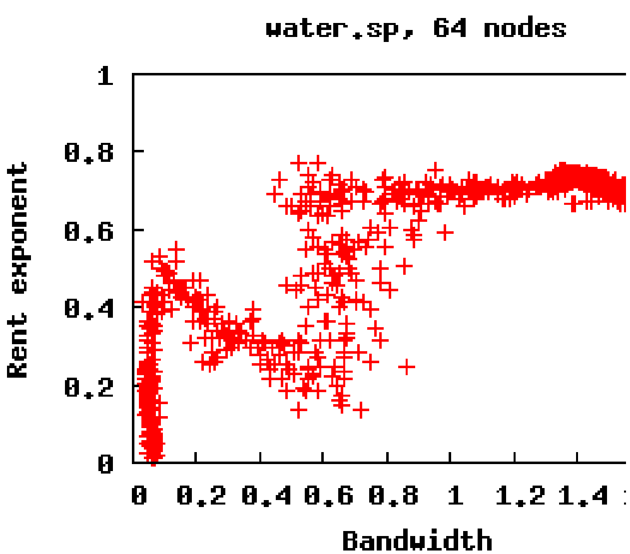

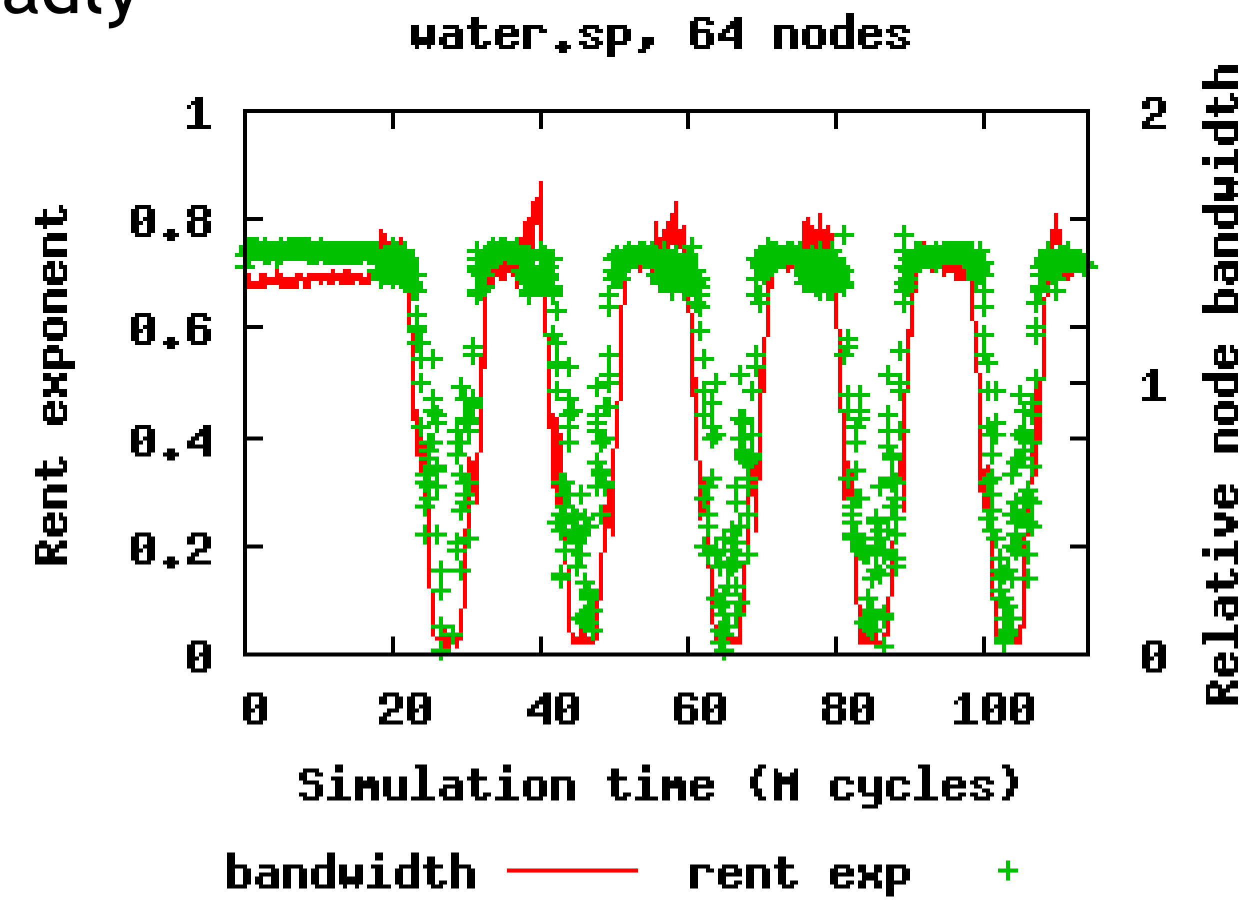

Node placement vs. variable traffic

- Node partitioning can lead to optimal node placement (minimal communication distances)
- But: varying traffic → varying optimal placement?
- Compute interval similarity, based on partitionings
- Account for traffic intensity (moving non-communicating nodes has no effect)

Similarity of communication between intervals

- For time intervals X and Y, each with traffic pattern *traffic* and optimal partitioning *part*
- part[X] cuts minimal fraction of traffic[X]
- assume we use part[X] in interval Y, what fraction of traffic[Y] is cut? → cut[X,Y]
- always more than part[Y] would = cut[Y,Y]
- similarity of partitionings, accounting for traffic intensity:

$$sim[X,Y] = \frac{cut[X,X] + cut[Y,Y]}{cut[Y,X] + cut[X,Y]}$$

UNIVERSITEIT GENT

Similarity measure properties

$$sim[X,Y] = \frac{cut[X,X] + cut[Y,Y]}{cut[Y,X] + cut[X,Y]}$$

- cut[X,X] < cut[Y,X] and cut[Y,Y] < cut[X,Y]
  → 0 ≤ sim[X,Y] ≤ 1
- sim[X,X] = 1
- when traffic[X] >> traffic[Y]:
  cut[*,Y] ~ 0 → sim[X,Y] ~ cut[X,X]/cut[Y,X]
  (only dependent on traffic[X])

UNIVERSITEIT GENT

Similarity matrix: FFT

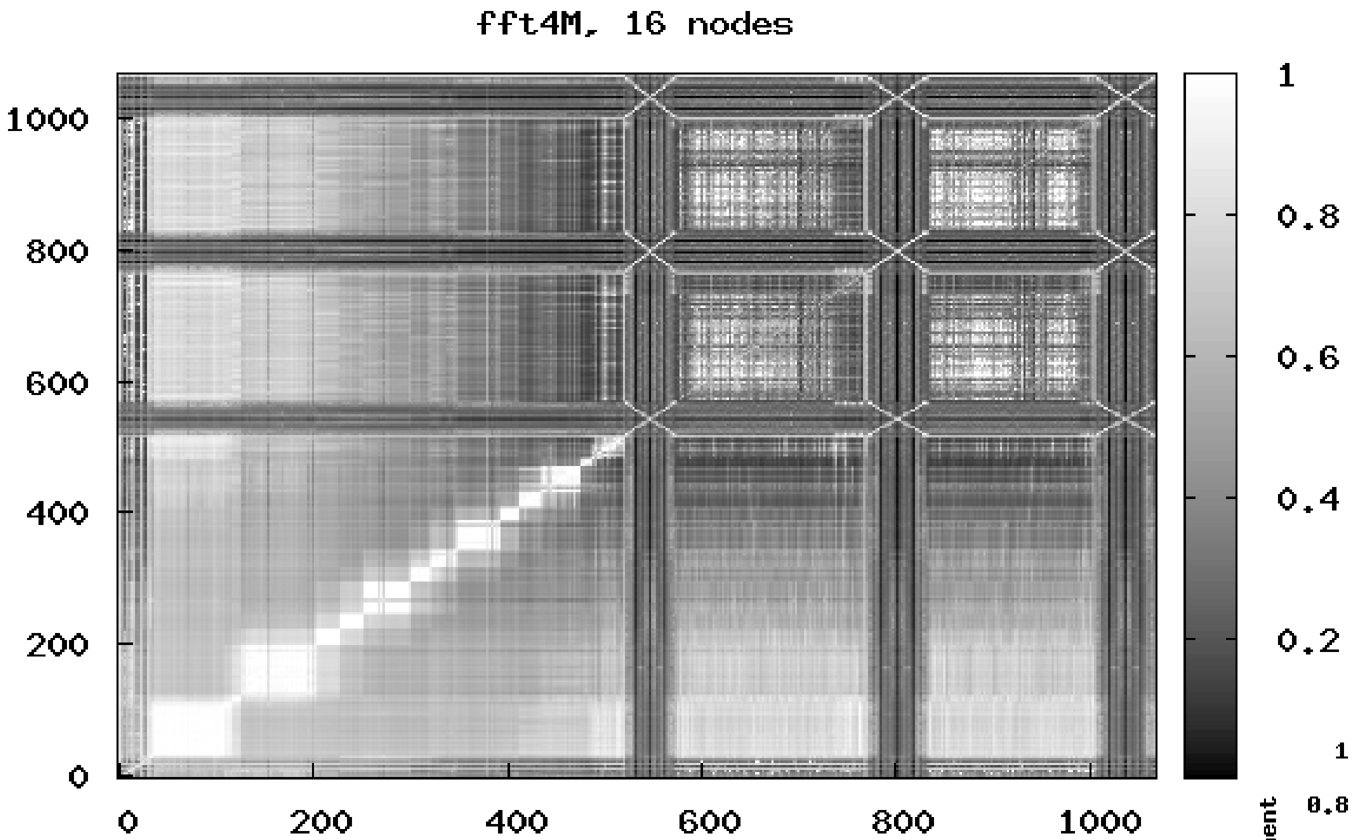

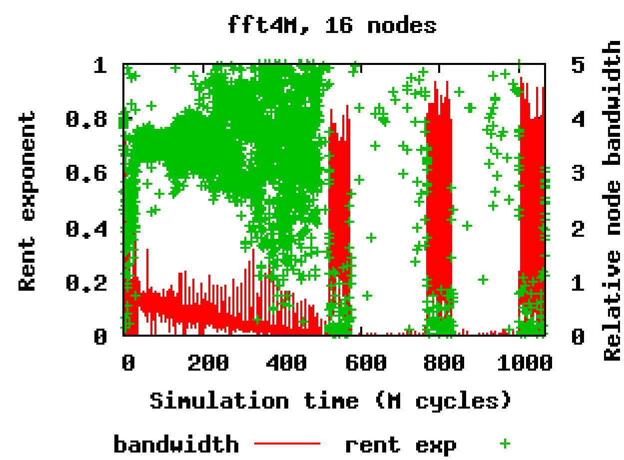

Similarity matrix: Water

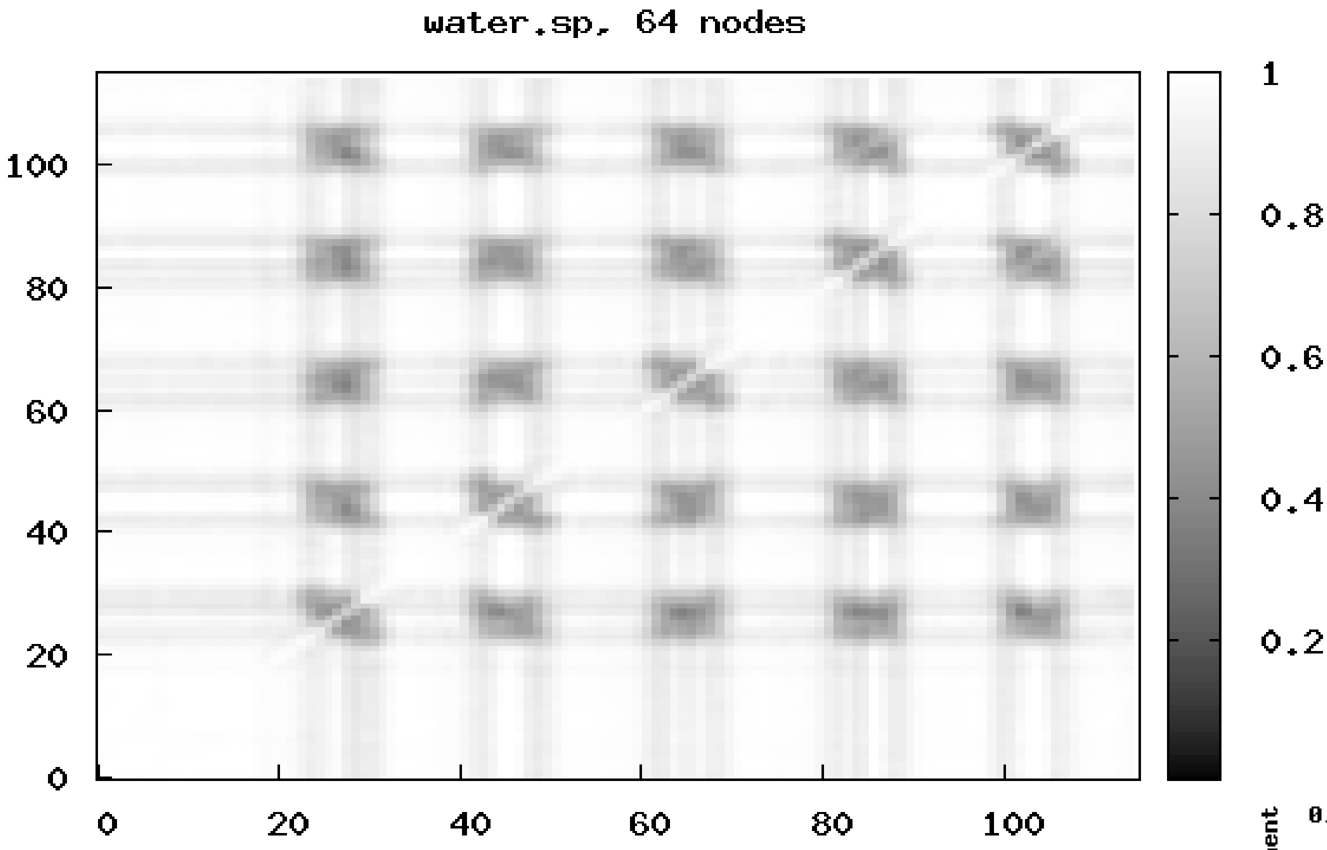

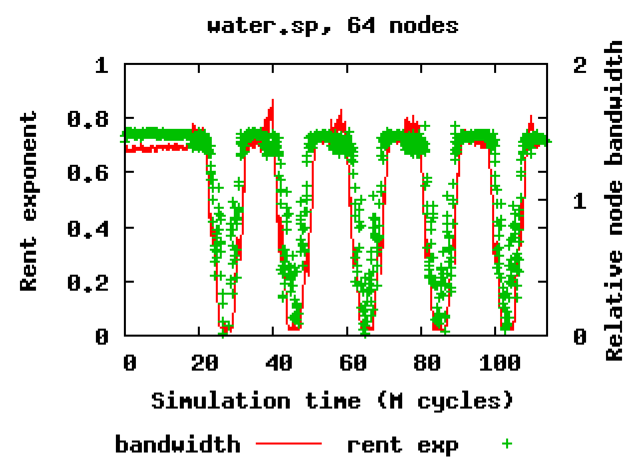

Suitability of a single placement

- Static network → one single placement
- How suitable is this placement through time?
- Suitability measure: based on partitionings (as are placements)
- Optimal partitioning for traffic[X]: part[X], cutting a bandwidth cut[X,X].
- Suitability of partitioning P: cut[X,X] / cut[P,X]

Suitability of a single placement

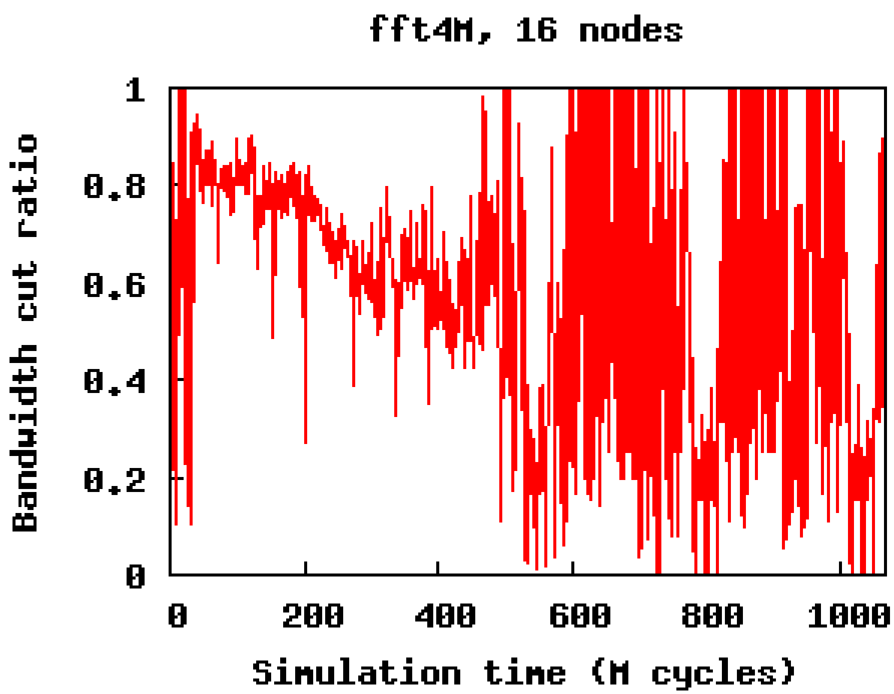

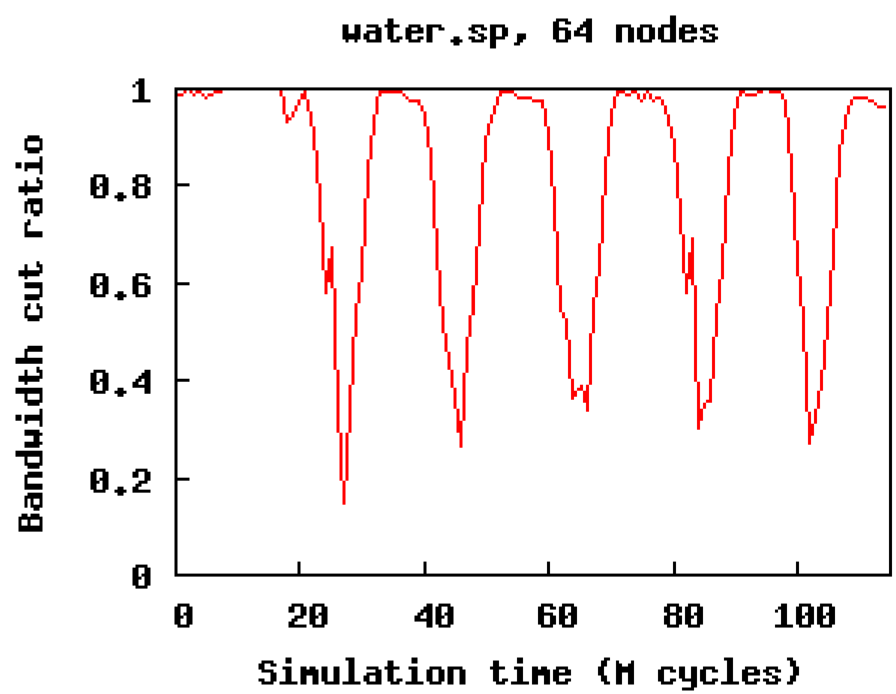

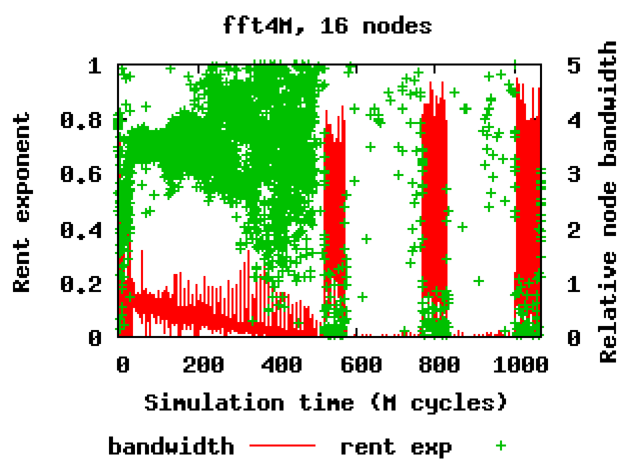

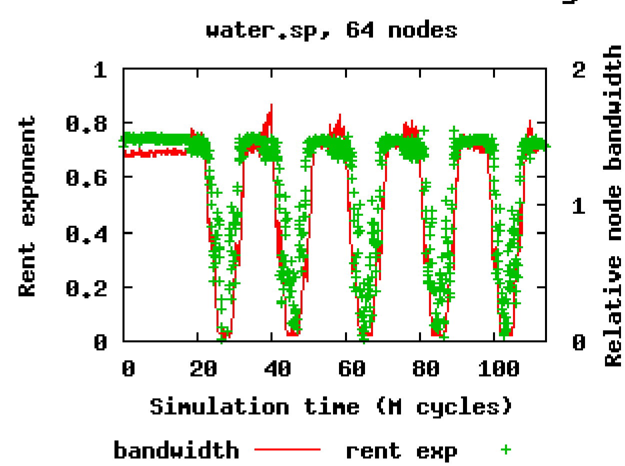

Outline

- Introduction
- Rent’s rule & traffic locality
- Time-varying network traffic
- Conclusions

UNIVERSITEIT GENT

Conclusions

- Measuring Rent exponents:
  - small number of nodes: difficult to measure, lots of noise
  - shared-memory: implicit communication, lots of non-essential communication → better/other results with message-passing?
- Still, difference in locality is visible, can be traced back to the benchmark's algorithm
- Time-variant communication!
- Rent's Rule (partitioning) is helpful to study communication behavior

UNIVERSITEIT GENT

Thank you!