Implications of Rent's Rule for Network-on-Chip Design and Its Fault-Tolerance

Daniel Greenfield

UNIVERSITY OF CAMBRIDGE

Presented at 1st International Network on Chip Symposium
Princeton, May 2007

Introduction: NoC Evaluation

Design A

Design B

Introduction: NoC Evaluation

Introduction: NoC Evaluation

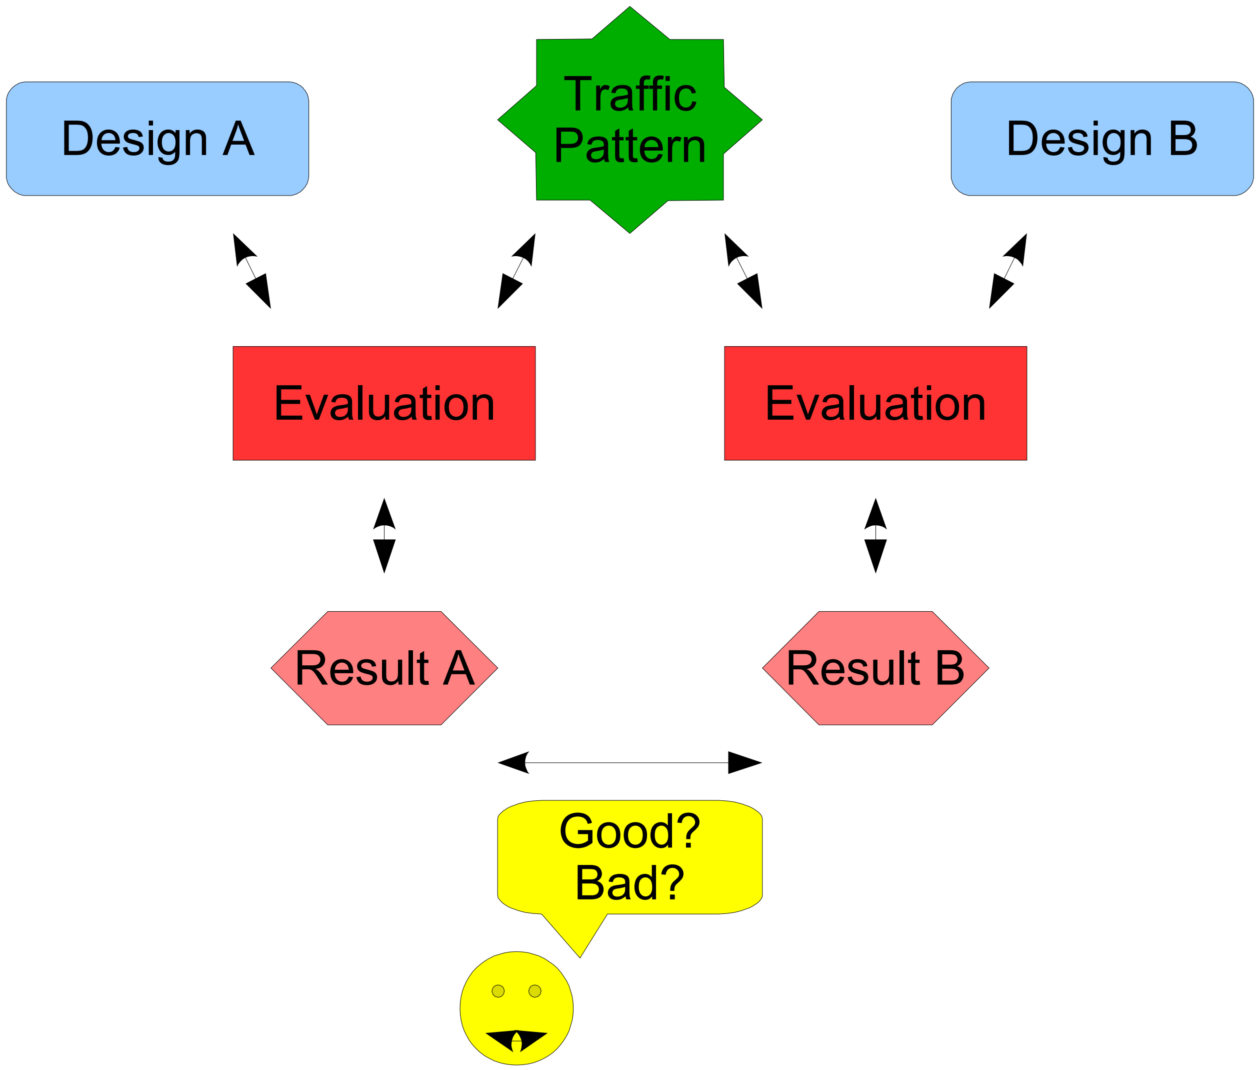

Introduction

- When we build a new Network-on-Chip
  - Is it better than existing networks on chip?
- Traffic models are used to evaluate and compare Networks-on-Chip
  - What does traffic look like?
  - Designing for the future, so what is its scaling behaviour?
  - Commonly use very simple non-local models. Is this wise?

Introduction

- When we build a new Network-on-Chip
  - Is it better than existing networks on chip?
- Traffic models are used to evaluate and compare Networks-on-Chip
  - What does traffic look like?
  - Designing for the future, so what is its scaling behaviour?
  - Commonly use very simple non-local models. Is this wise?
- How important is locality and how can we model it?

Scaling of per-tile BW

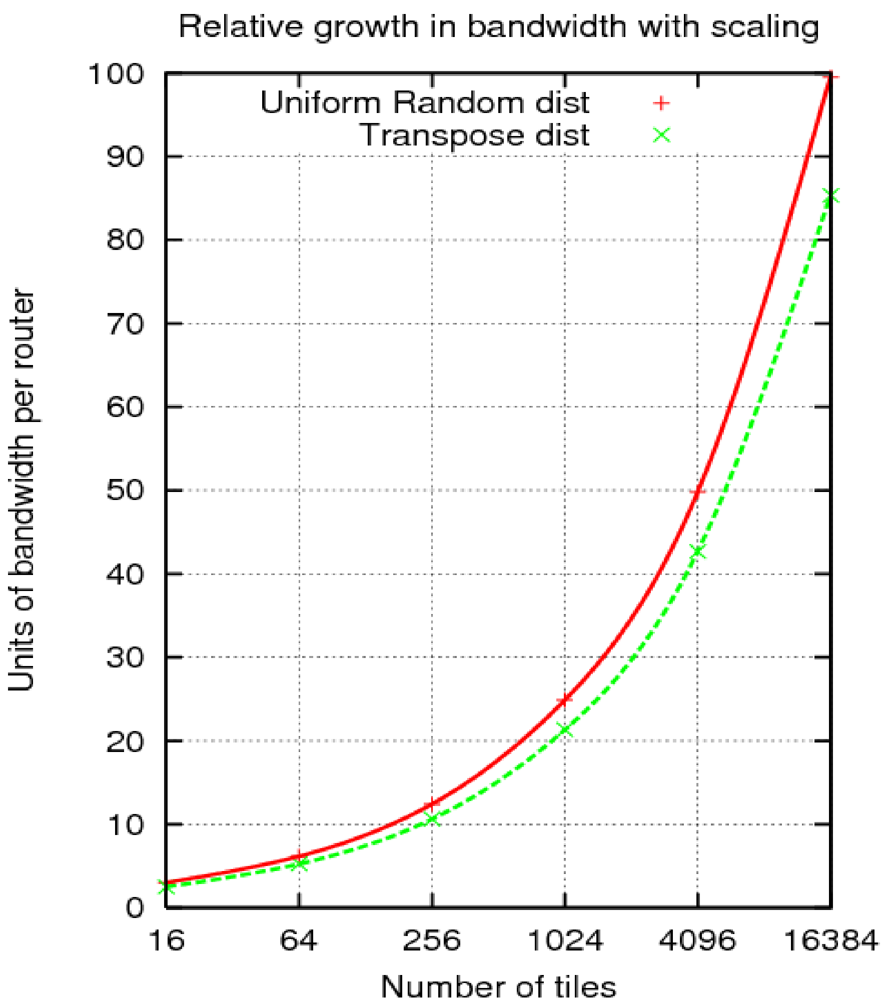

- Each tile generates one unit of BW
- BW through/routed by each tile grows exponentially with technology scaling
- Doubling per quadrupling of tiles
- # Metal layers, repeaters, power grows exponentially

Scaling of per-tile BW

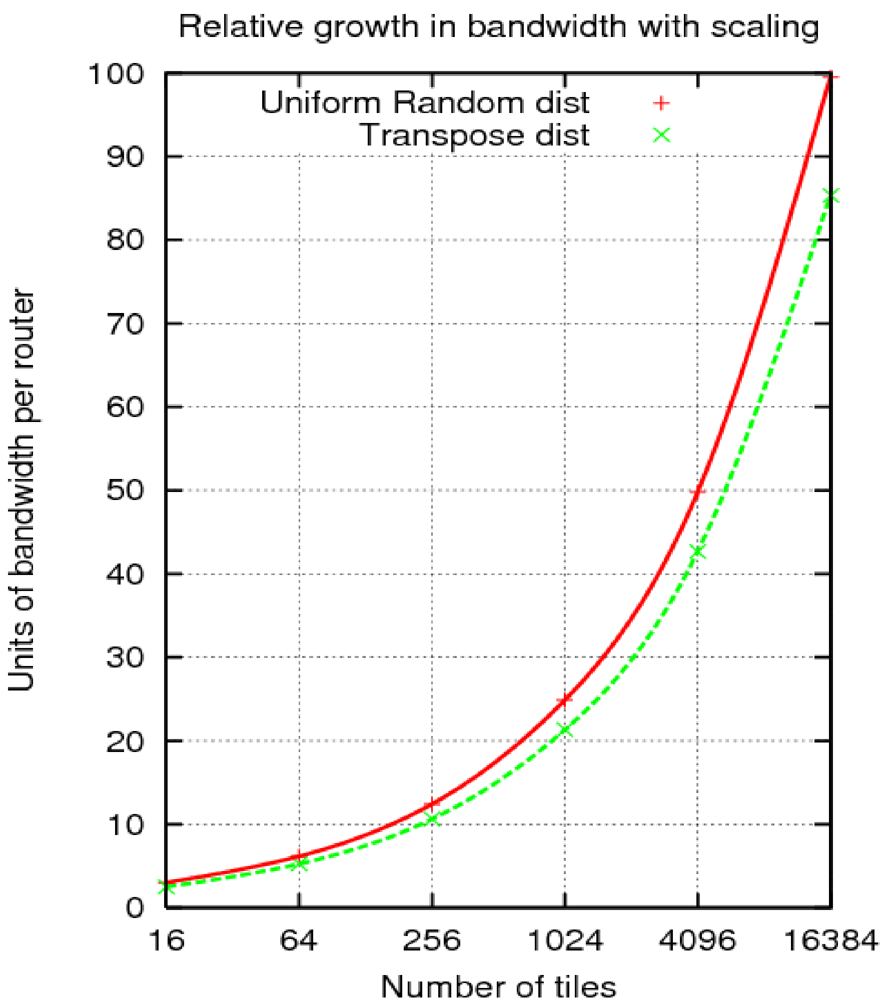

- Each tile generates one unit of BW
- BW through/routed by each tile grows exponentially with technology scaling
- Doubling per quadrupling of tiles
- # Metal layers, repeaters, power grows exponentially *for all on-chip topologies*

Classical Rent's Rule

- A description of locality

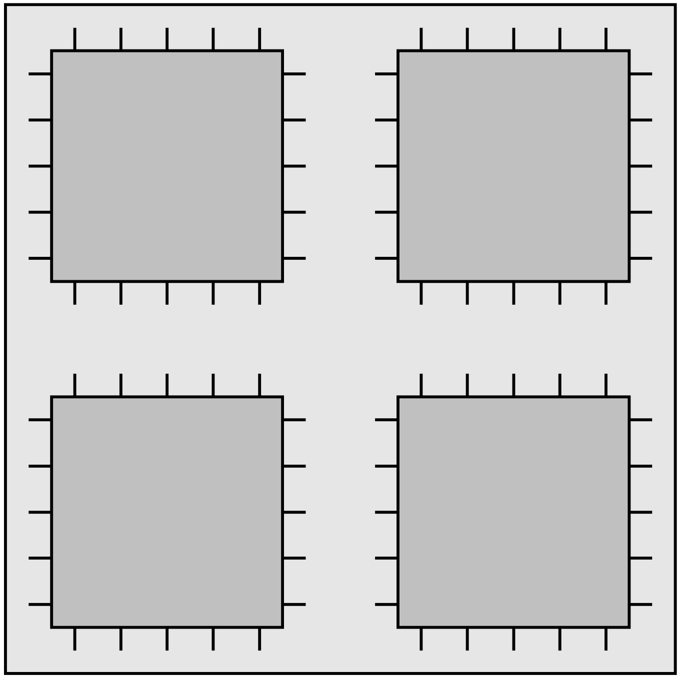

Classical Rent's Rule

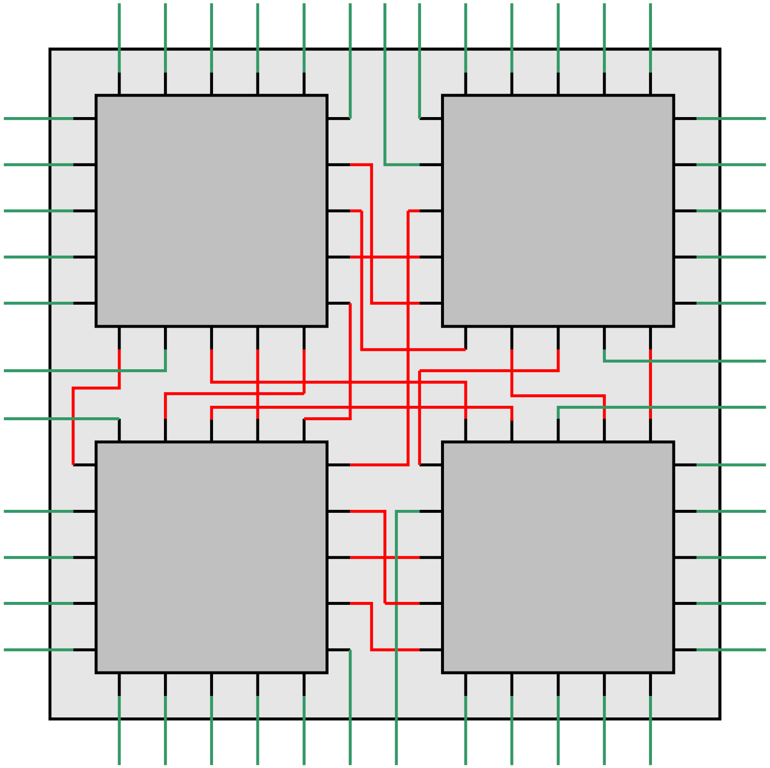

- A description of locality
- Rent's exponent characterises the split between internal and external nets

Reasons for Placement Locality

Domain to minimise	Wires	NoC
Delay	Wire delay	NoC latency (& congestion)
Congestion	Wire-density	Cross-sectional bandwidth
Power	Wire buffering & length	Hop-length & router utilisation

Bandwidth Version

- Bandwidth locality

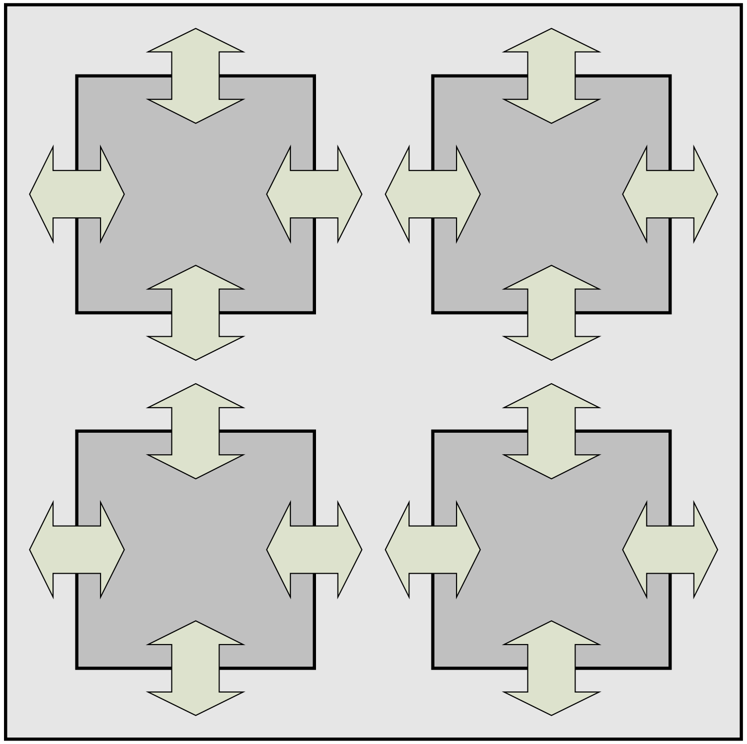

Bandwidth Version

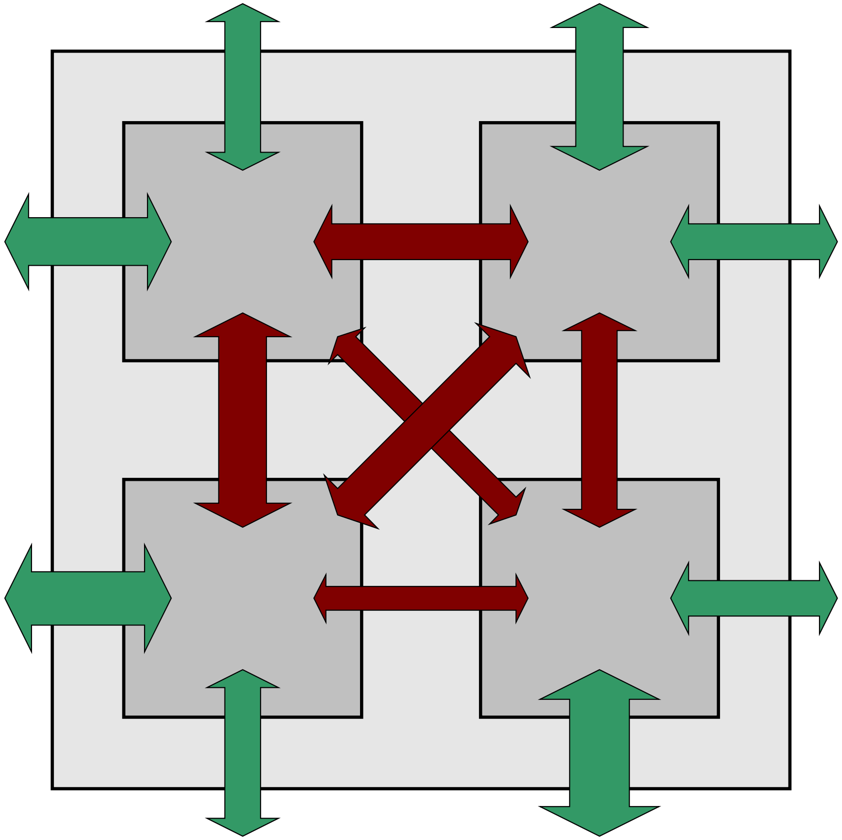

- Bandwidth locality
- Rent's bandwidth-exponent characterises split between internal and external traffic

Scaling of per-tile BW (comparison)

- Rent's-type locality allows graceful increase in per-tile bandwidth

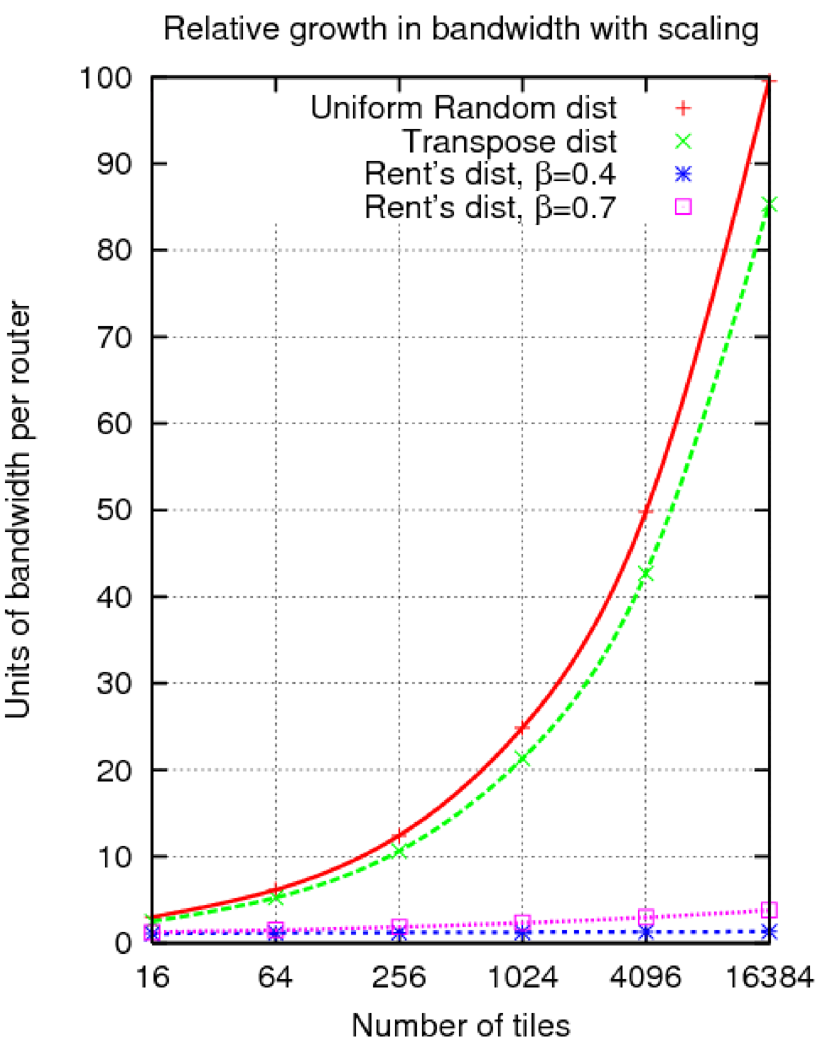

Scaling of per-tile BW (Rent's)

- Many-fold increase with scaling for high versus low Rent's bandwidth-exponent
- This is in line with ordinary VLSI expectations

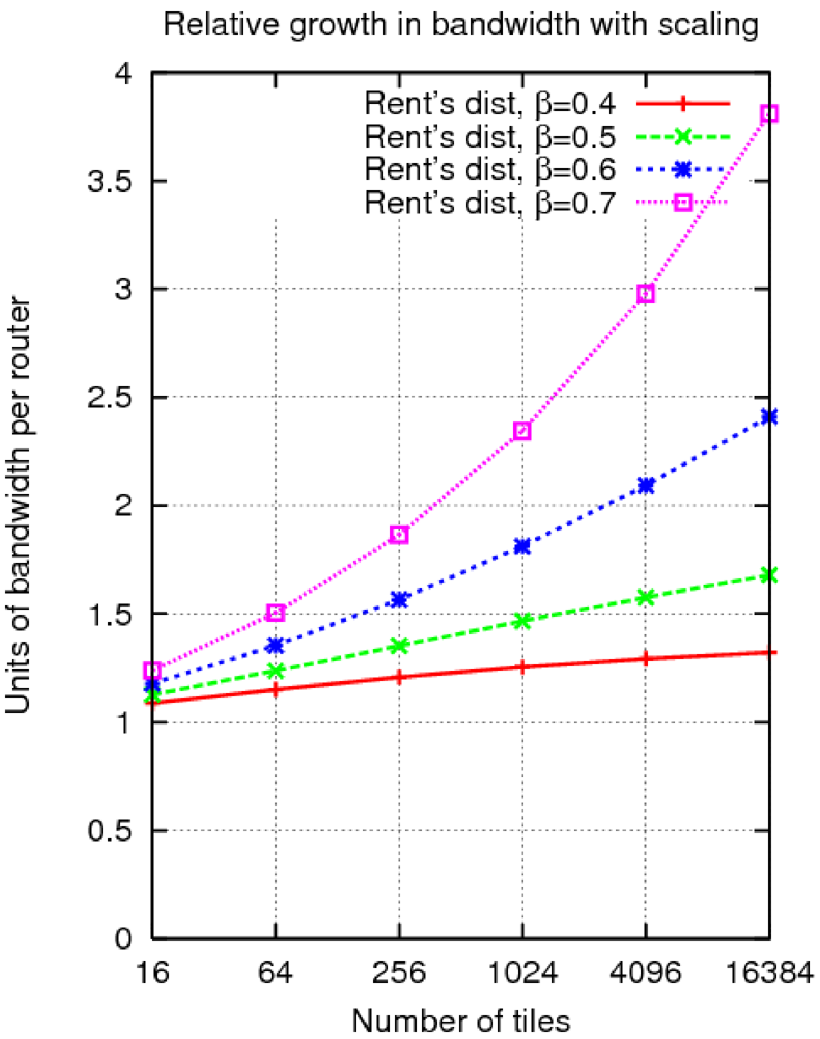

Example Hop-length Distribution (8x8)

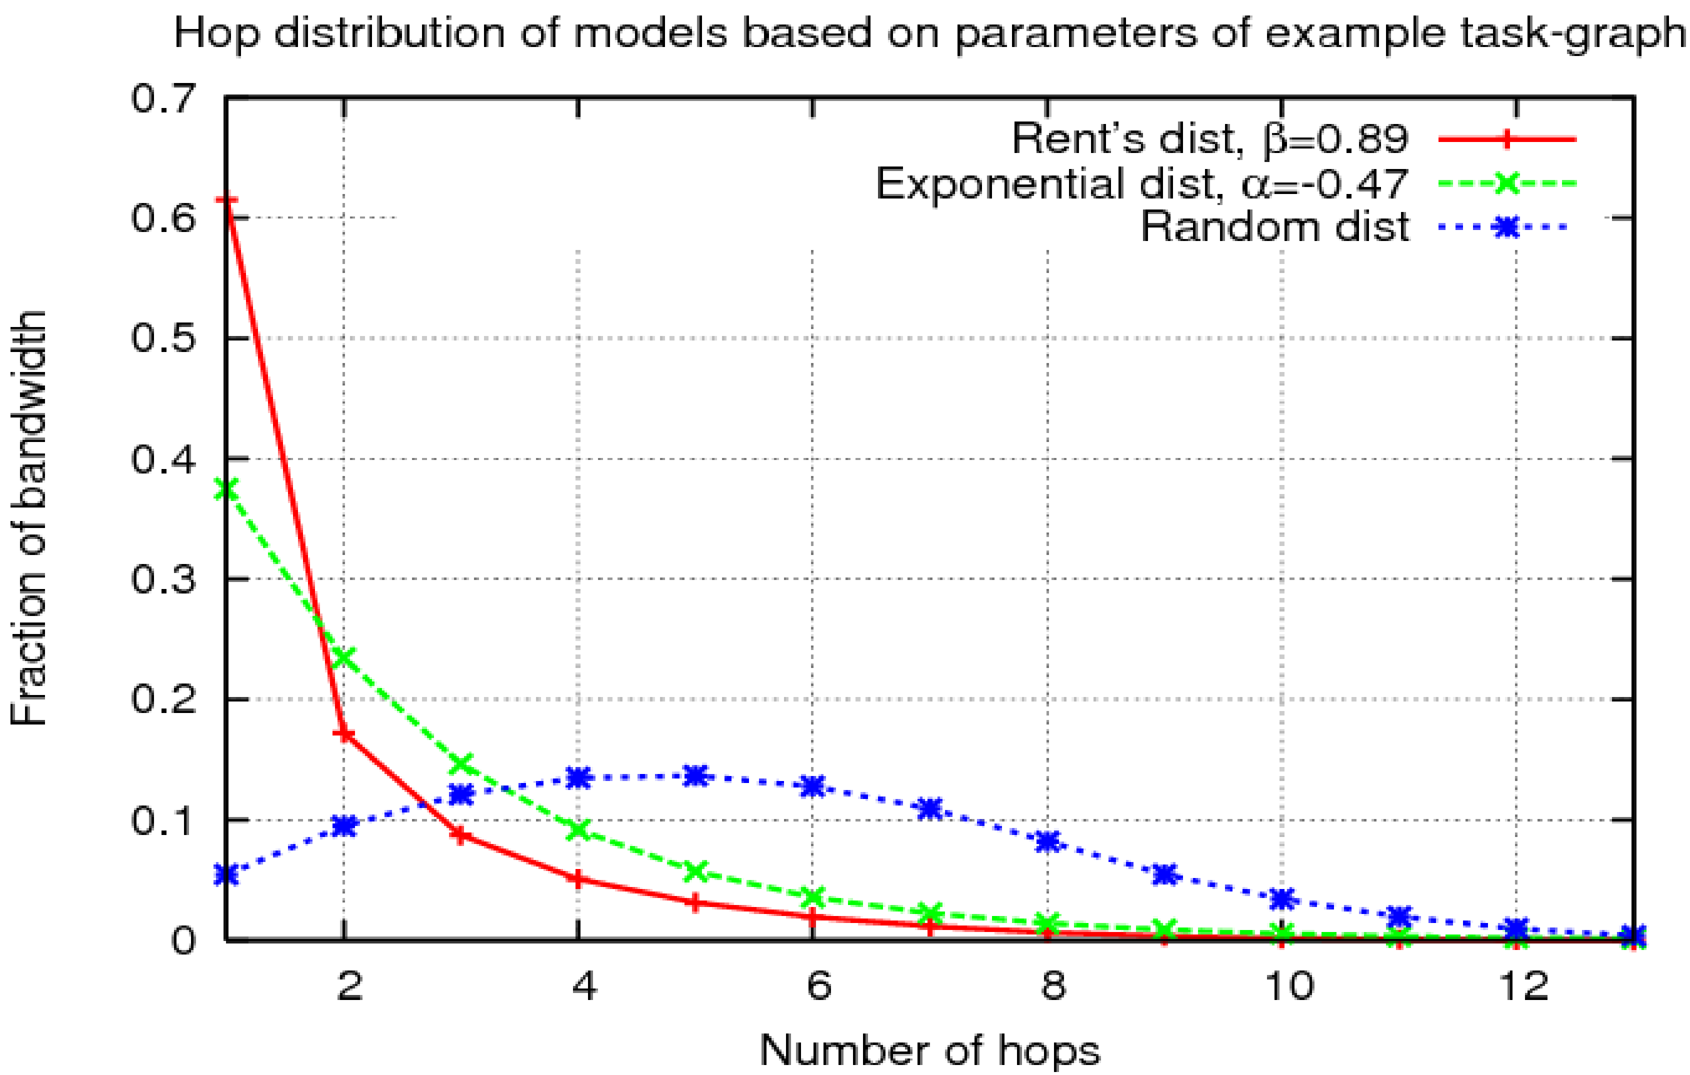

Traffic Types

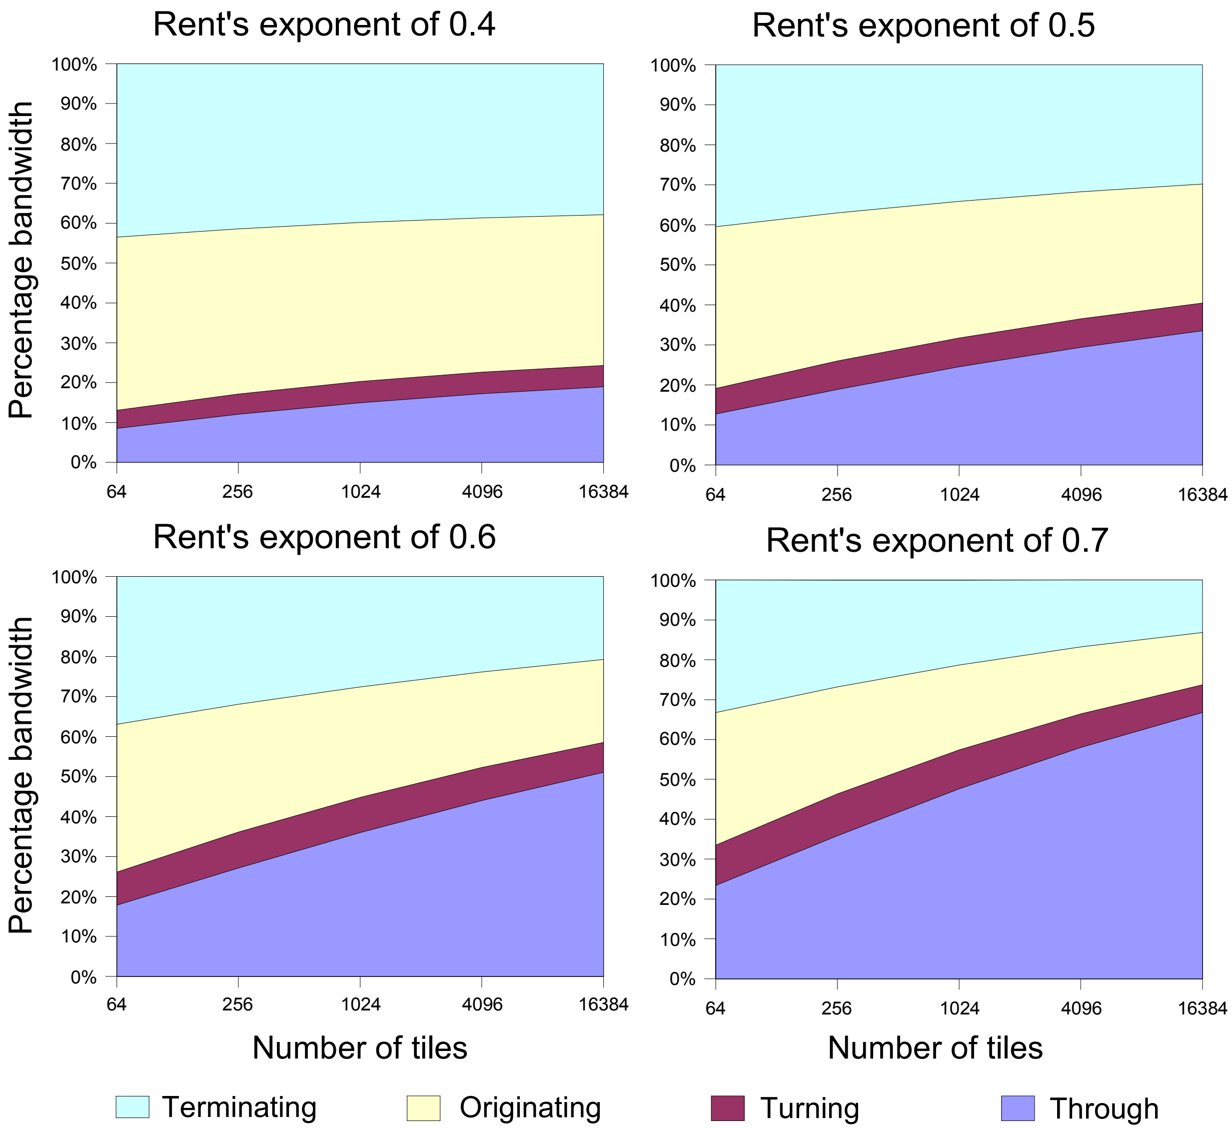

Case-study: A Fault Tolerance Approach

- Add bypass of faulty router logic
- Effectively a pipelined interconnect
- Implicit credit-scheme preserves flow-control

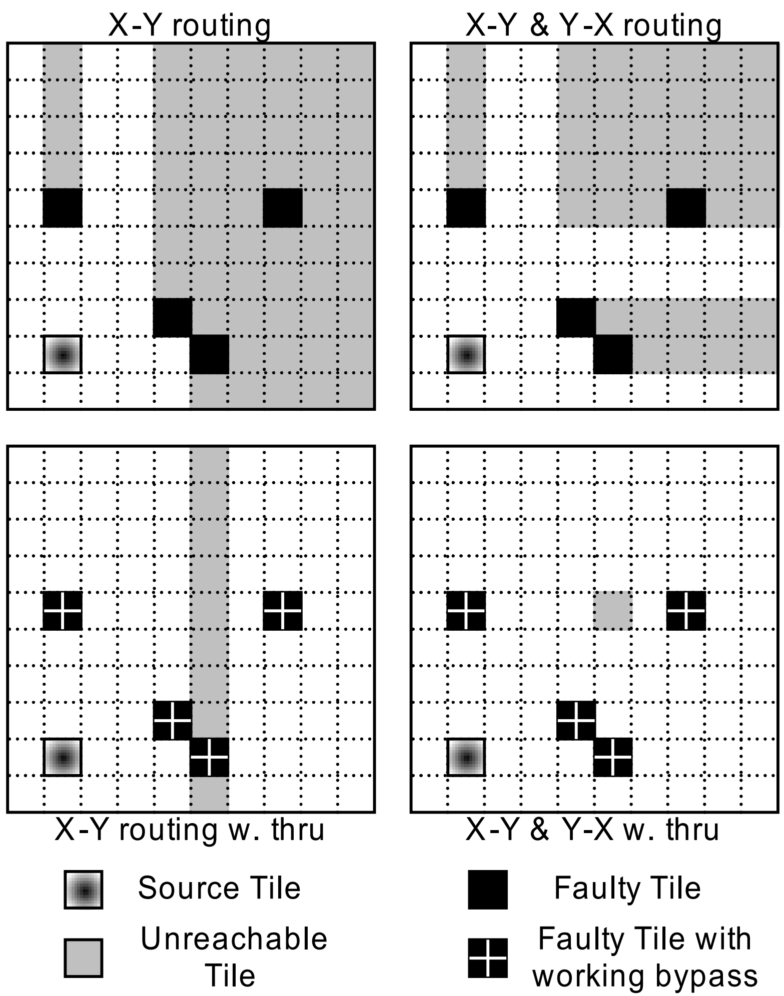

Reachability Analysis

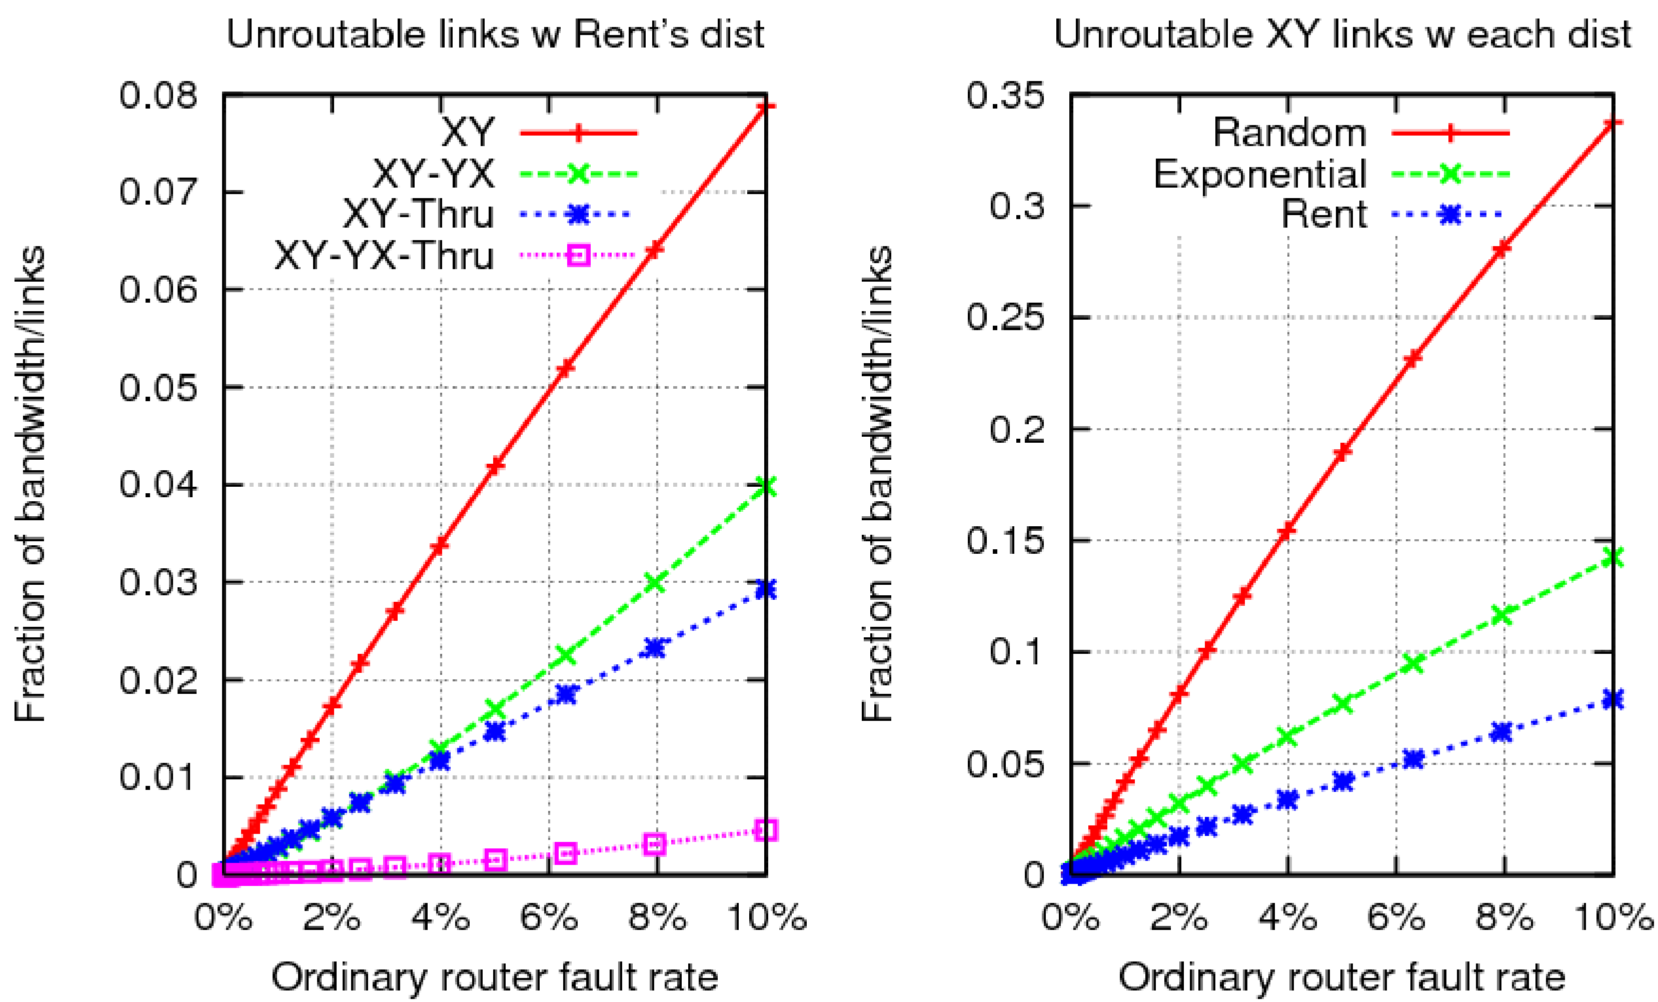

Congestion Surrounding Faulty Tile

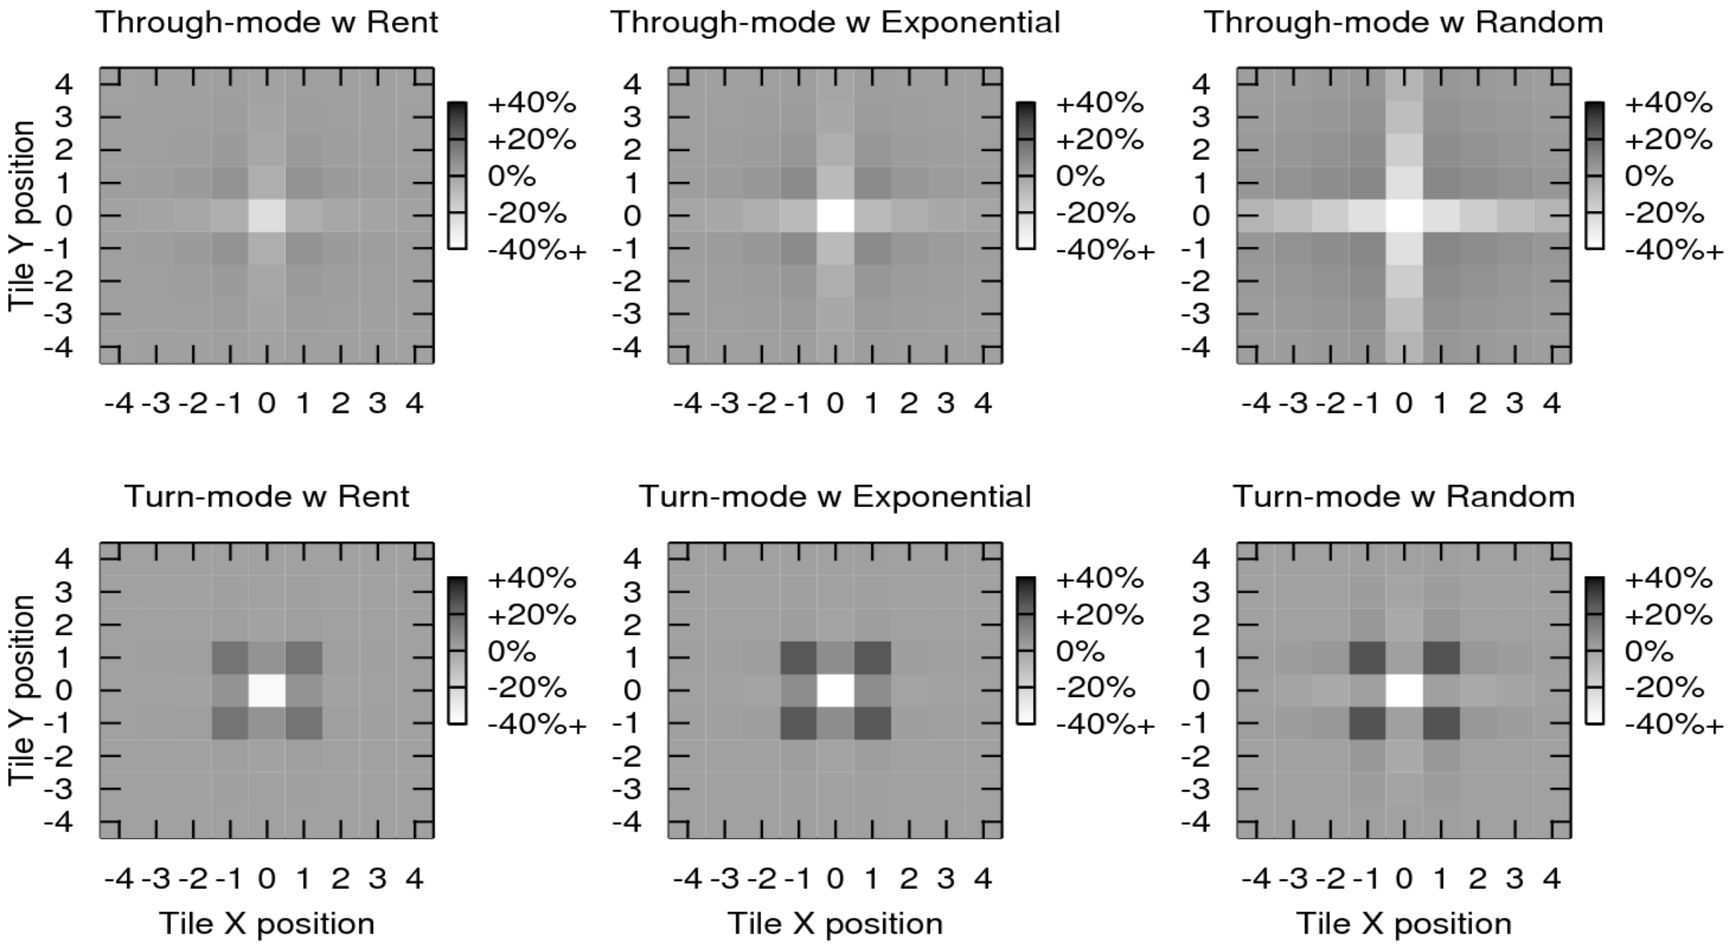

Conclusions

- Locality is essential. Non-local traffic models lead to unacceptable scaling behaviour
- Reasons for locality in software-circuits are analogous to ones in hardware-circuits
- A bandwidth version of Rent's rule for NoC can be derived to describe locality in software circuits
  - Distinct from classical Rent's rule
  - Hop-length distribution versus wire-length distribution
- Can analyse router utilisation according to traffic type
- In fault-tolerance analysis:
  - The traffic model used is important
  - Different traffic models lead to different results

Contact Information

Daniel.Greenfield@cl.cam.ac.uk
Computer Laboratory
Cambridge University, UK
http://www.cl.cam.ac.uk/research/comparch