Déjà Vu Switching for Multiplane NoCs

Ahmed Abousamra

Rami Melhem

Alex Jones

University of Pittsburgh

Motivation

- Power efficiency has become a primary concern in the design of CMPs.
- The NoC of Intel's TeraFLOPS processor consumes more than 28% of the chip's power.
- Network messages can be classified into ***critical*** and ***non-critical***
- It may be possible to send non-critical messages on a slower plane ***without*** hurting performance.

Setting

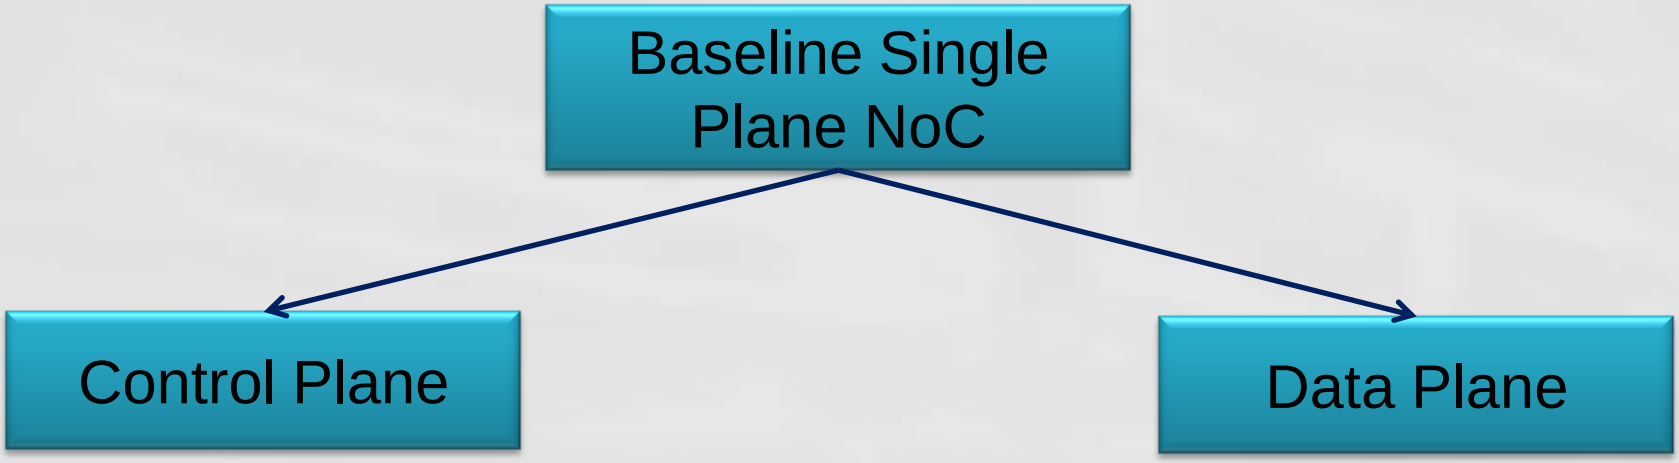

- Carries control and coherence messages: data requests, invalidates, acknowledgments, …
- Operates at baseline's voltage & frequency

- Carries data messages
- Operates at a lower voltage & frequency to save power

Outline

- Motivation & Related work
- **Importance of data messages to performance**
- The Optimization Problem
- Proposed Solution
  - Déjà Vu Switching
  - Analysis of the acceptable data plane speed
- Evaluation
- Summary

Importance of data messages to performance

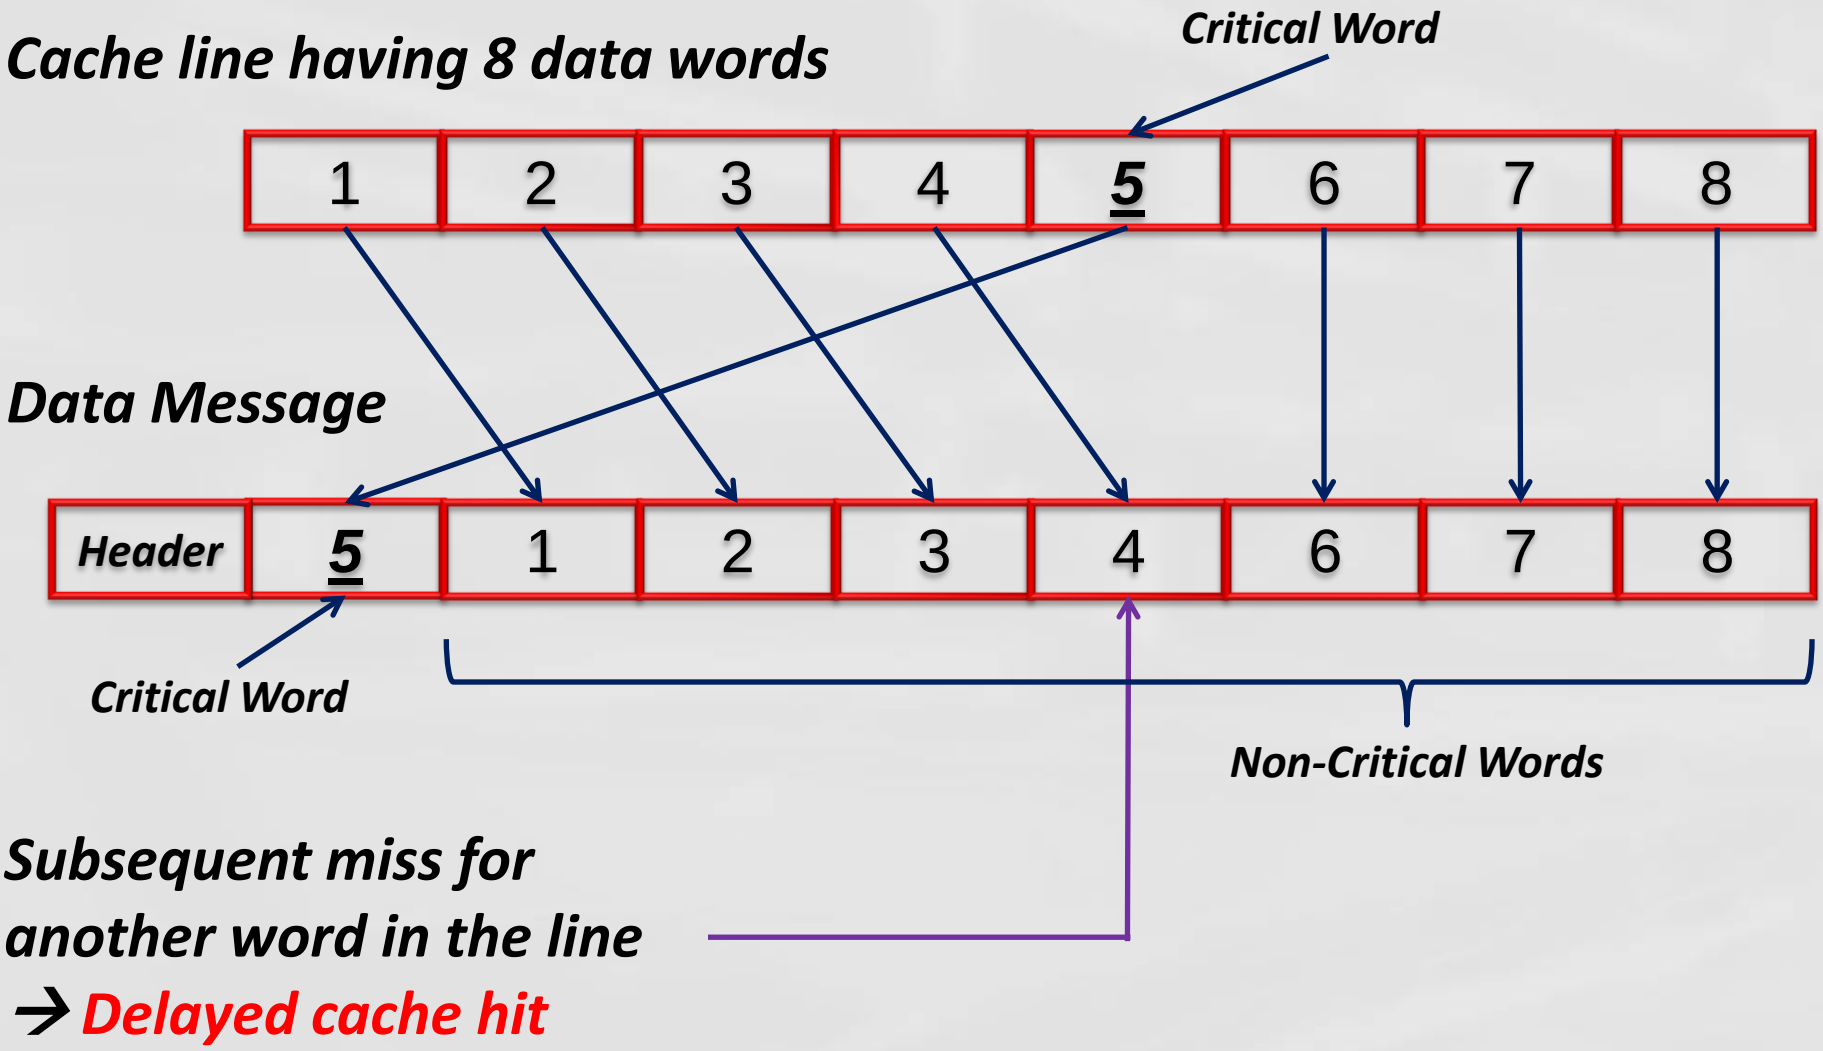

Importance of data messages to performance

- If delayed cache hits are overly delayed, performance can suffer.

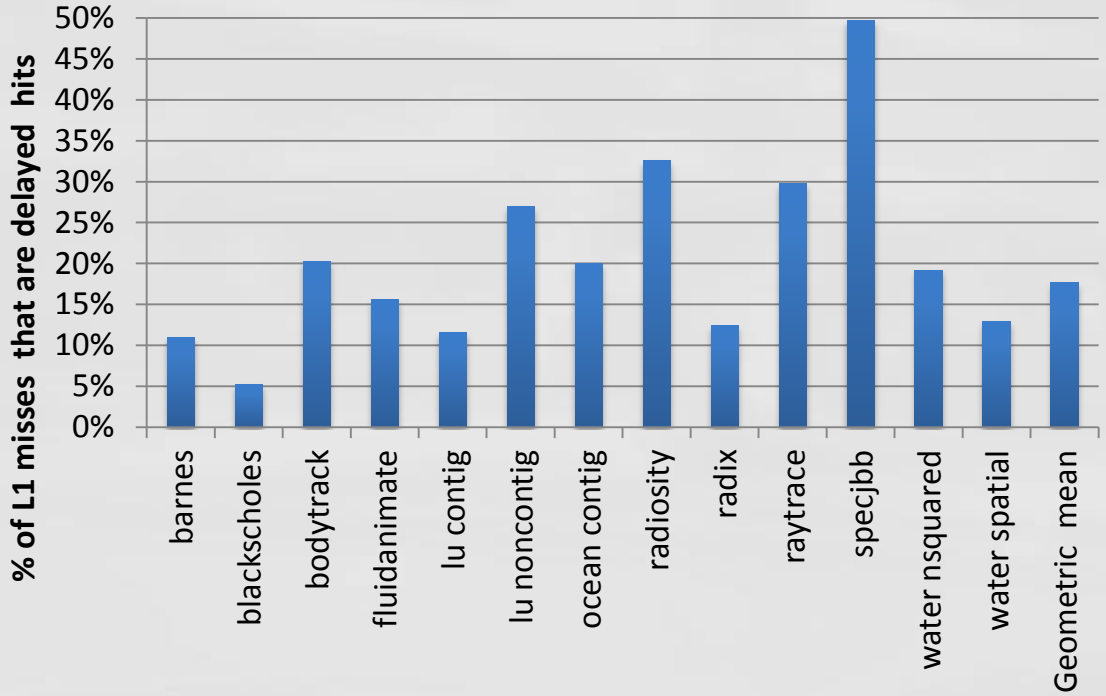

Percentage of L1 misses that are delayed cache hits

Outline

- Motivation & Related work
- Importance of data messages to performance
- **The Optimization Problem**
- **Proposed Solution**
  - Déjà Vu Switching
  - Analysis of the acceptable data plane speed
- Evaluation
- Summary

The Optimization Problem

- *How slow can we operate the data plane to maximize energy savings while not impacting performance?*

Proposed Solution

- Split NoC physically into 2 planes: control + data
- On data plane:
  - Use circuit-switching to speed-up communication.
  - Reduce voltage & frequency.
- Send a control message to establish the circuit once a cache hit is detected.
- Do not block circuit establishment message: *Déjà Vu Switching*
- Analyze acceptable slow down of the data plane to minimize energy while maintaining performance.

Déjà Vu Switching – Circuit Switching

Request Packet  **Reservation Packet**  **Data Packet** 

Control Plane

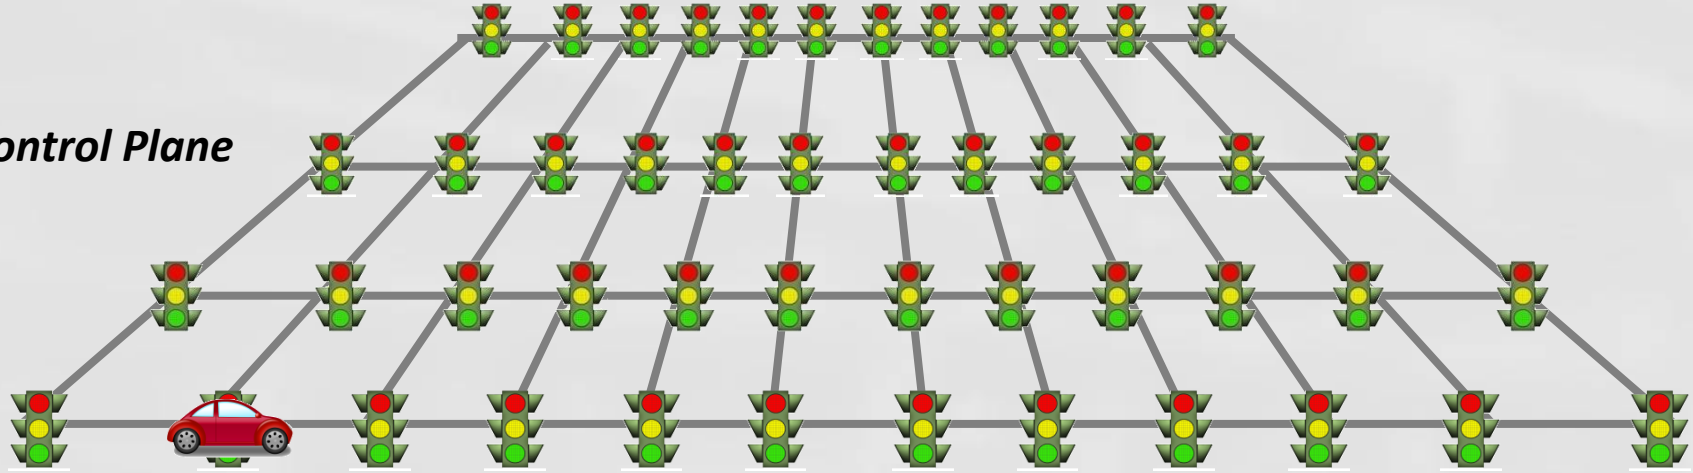

1) Upon a cache miss, a data request is sent

Data Plane

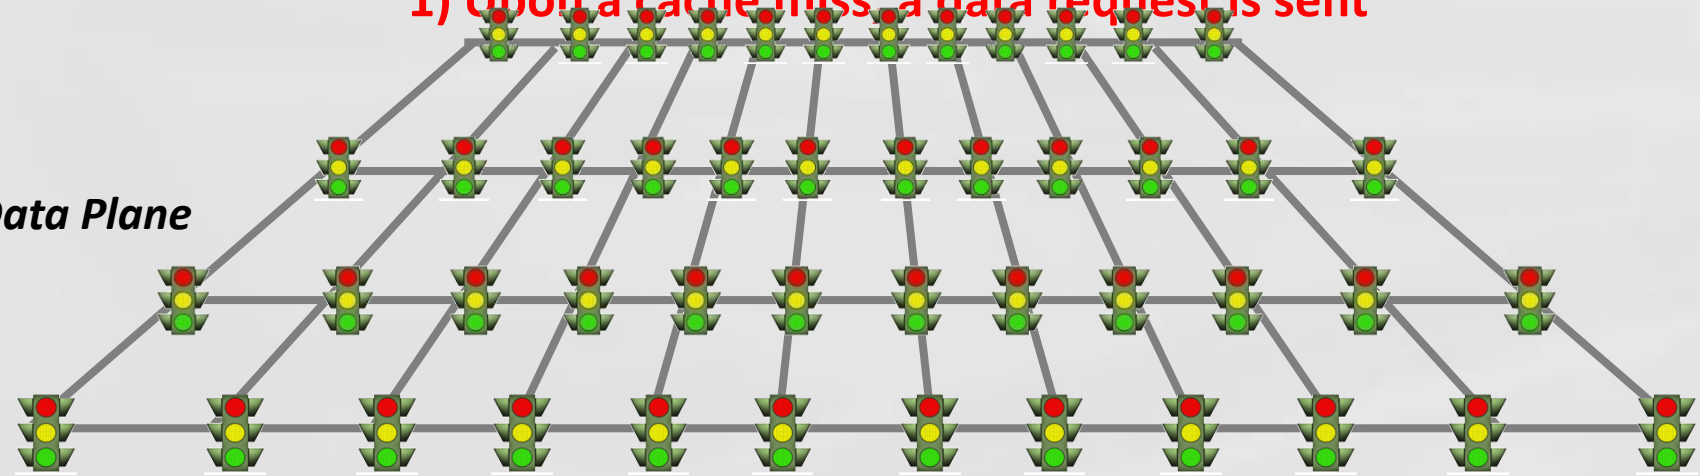

Déjà Vu Switching – Circuit Switching

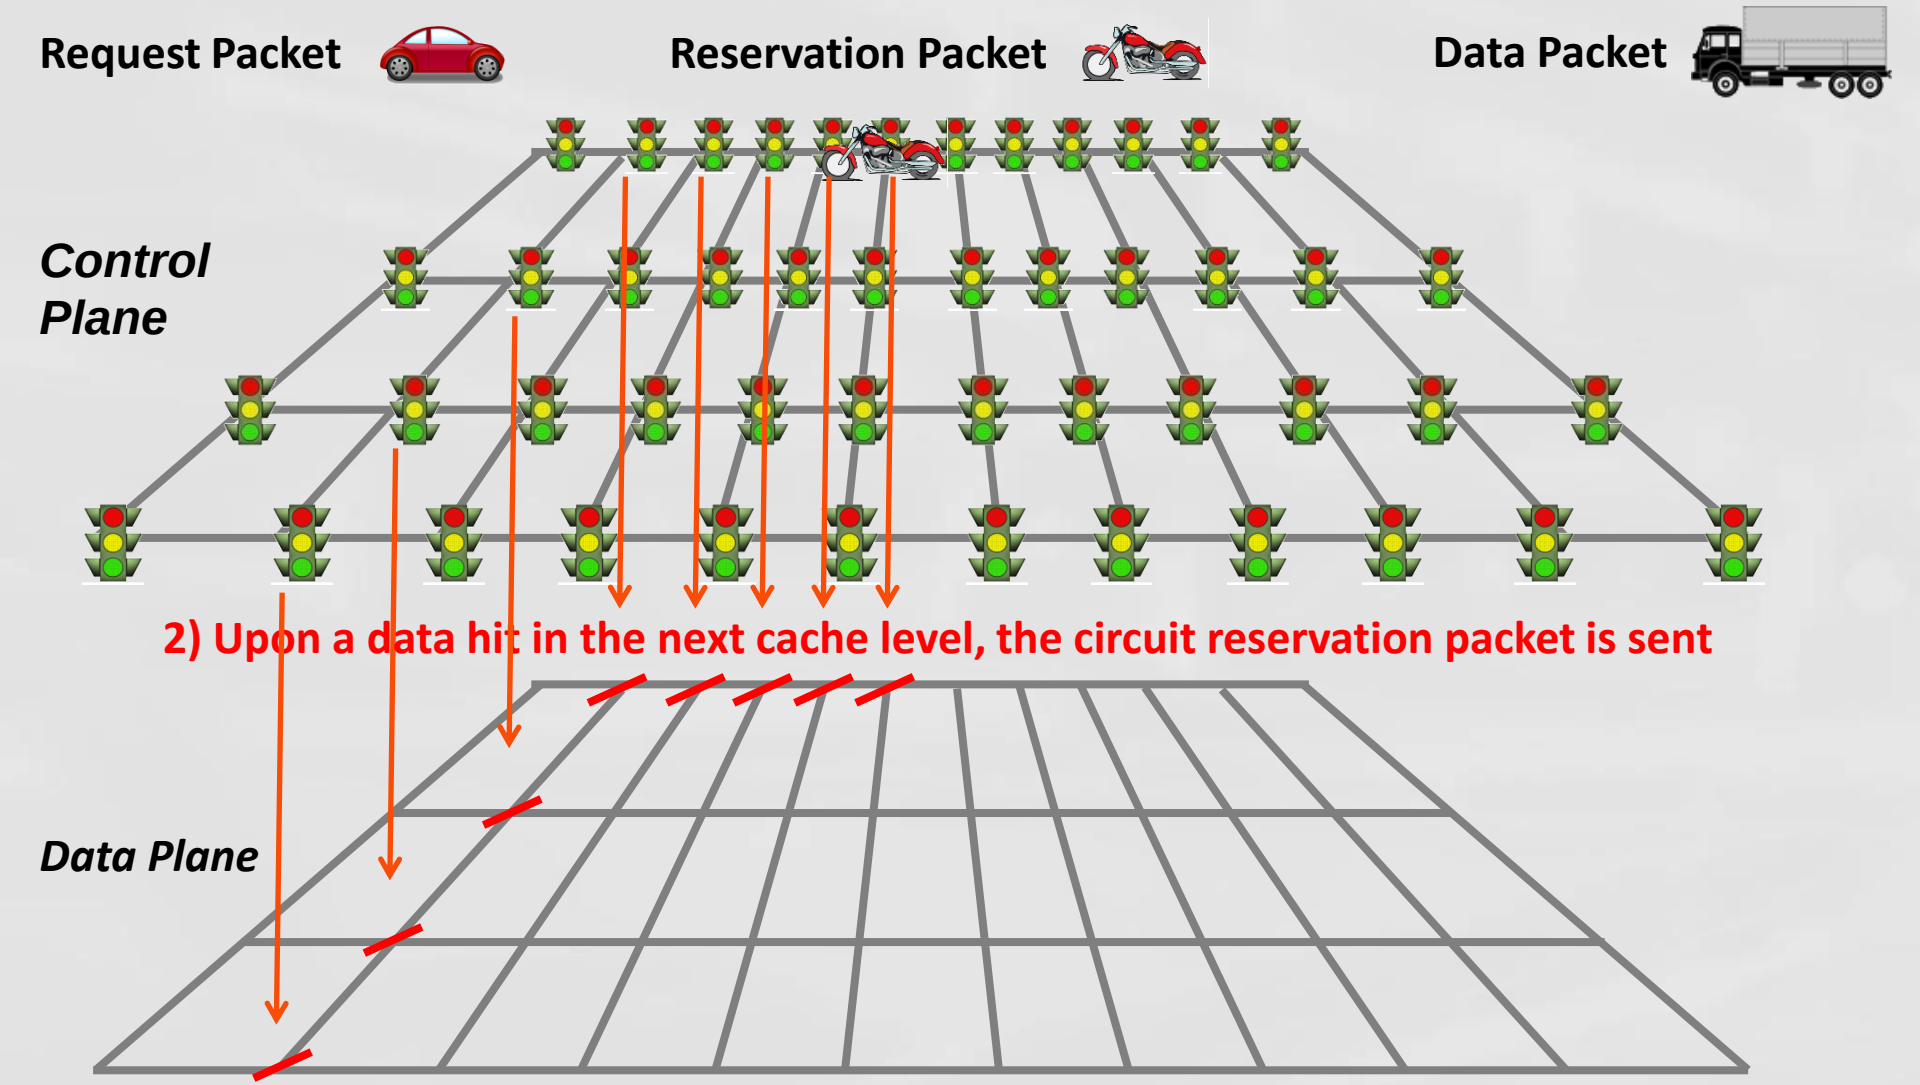

Déjà Vu Switching – Circuit Switching

Request Packet  **Reservation Packet**  **Data Packet** 

Control Plane

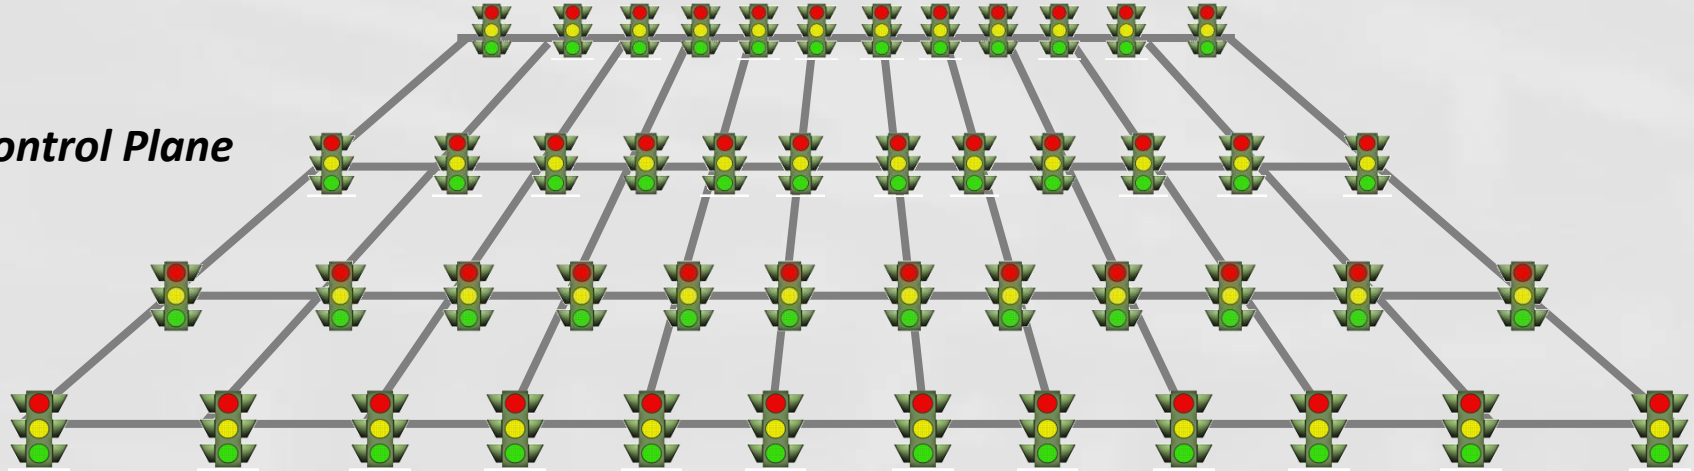

3) Finally, the data packet is sent to the requester on the reserved circuit

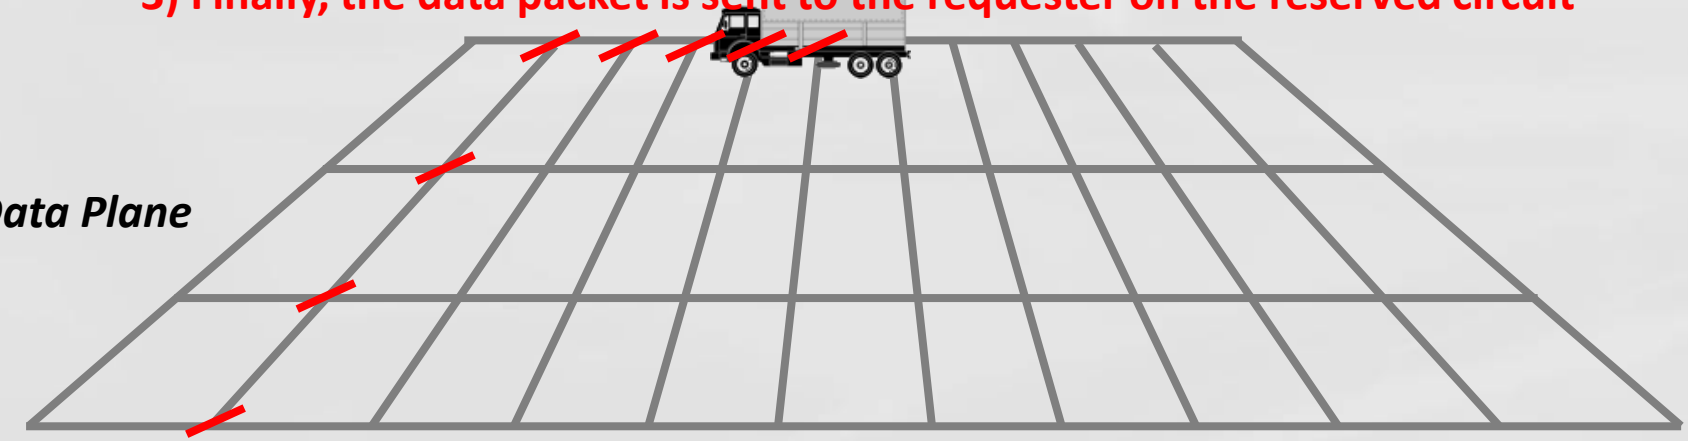

Data Plane

Déjà Vu Switching – Conflicting Reservations

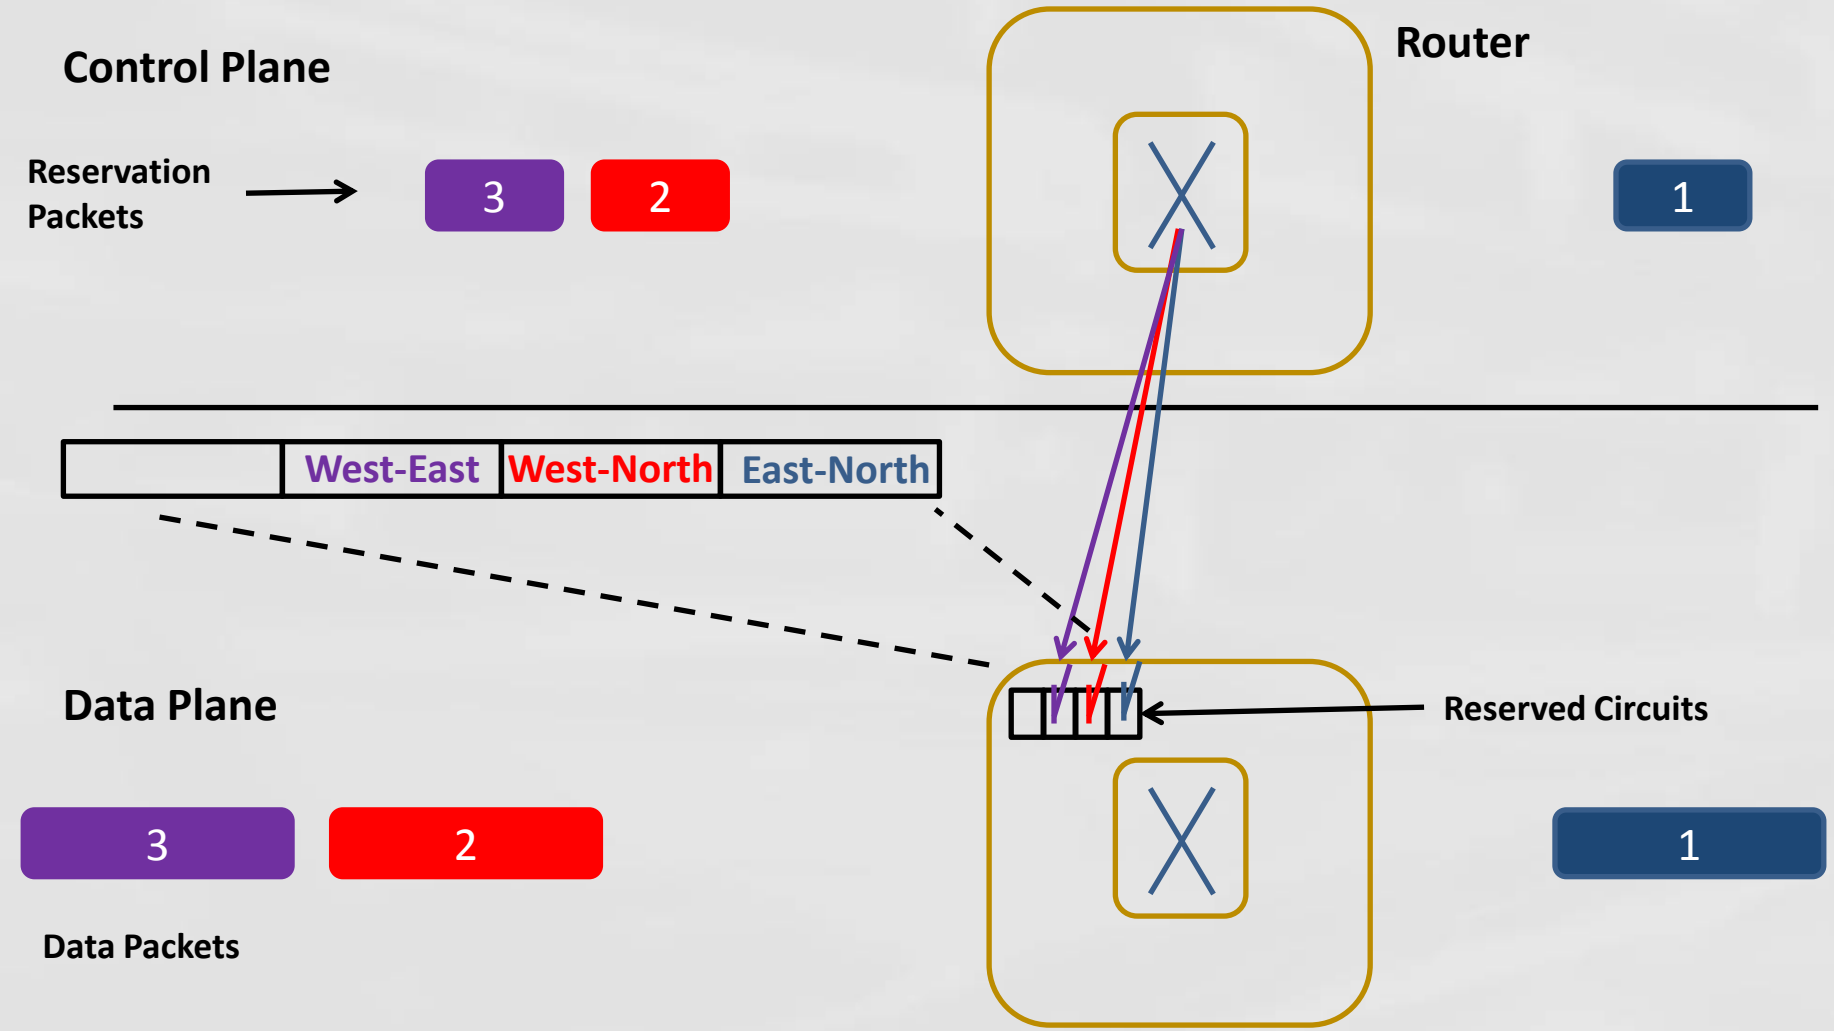

Déjà Vu Switching - Reservation Scheme

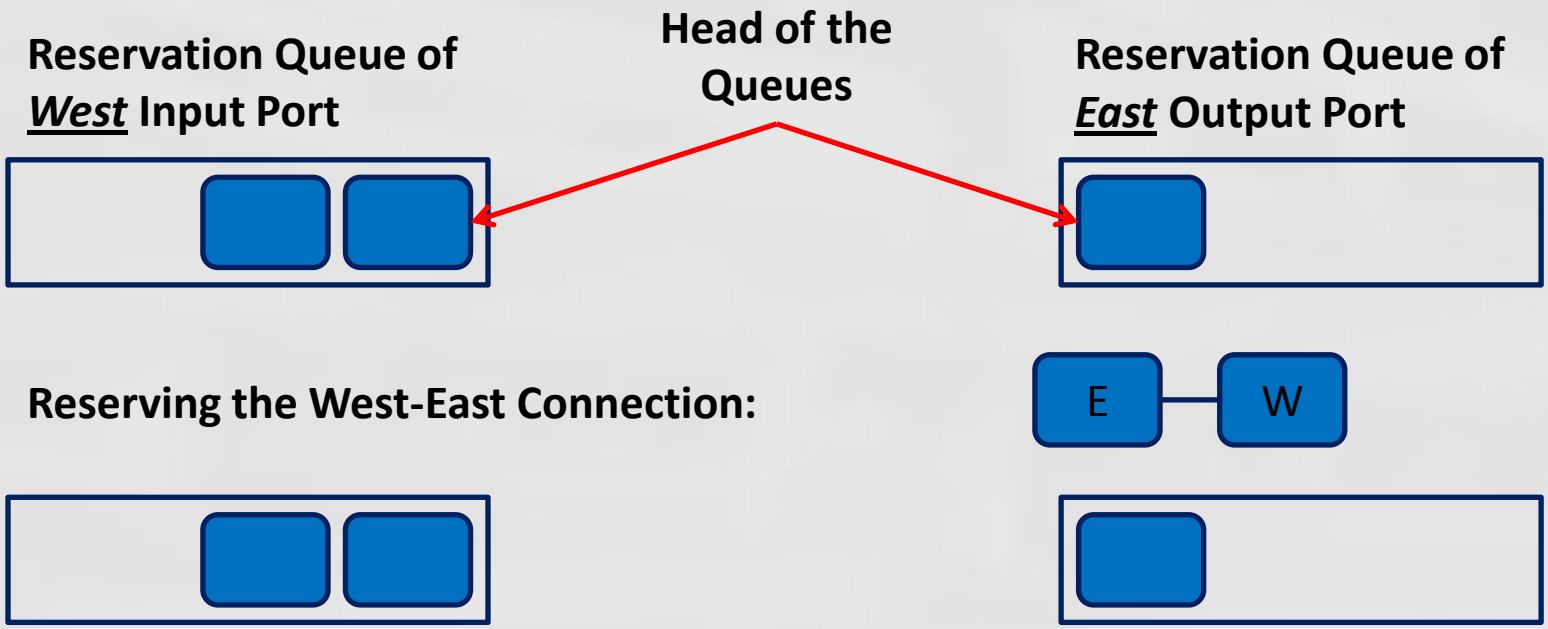

Port Names: North West South East Local

Déjà Vu Switching – Realizing Connections

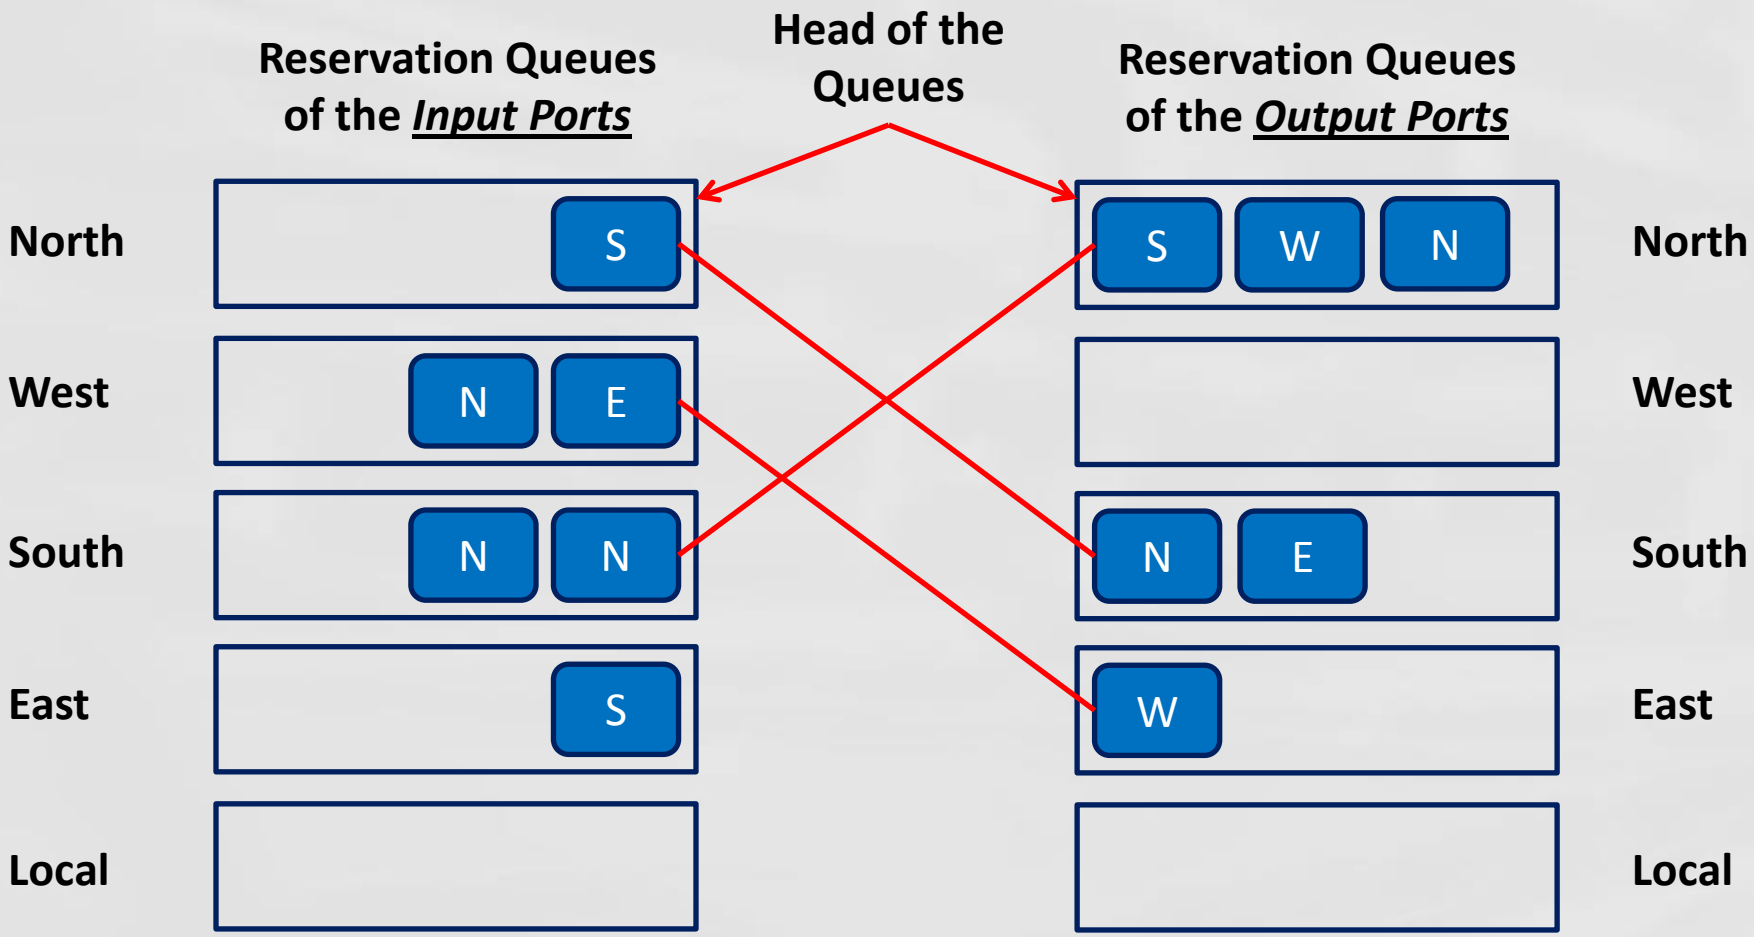

Port Names: North West South East Local

Déjà Vu Switching – Realizing Connections

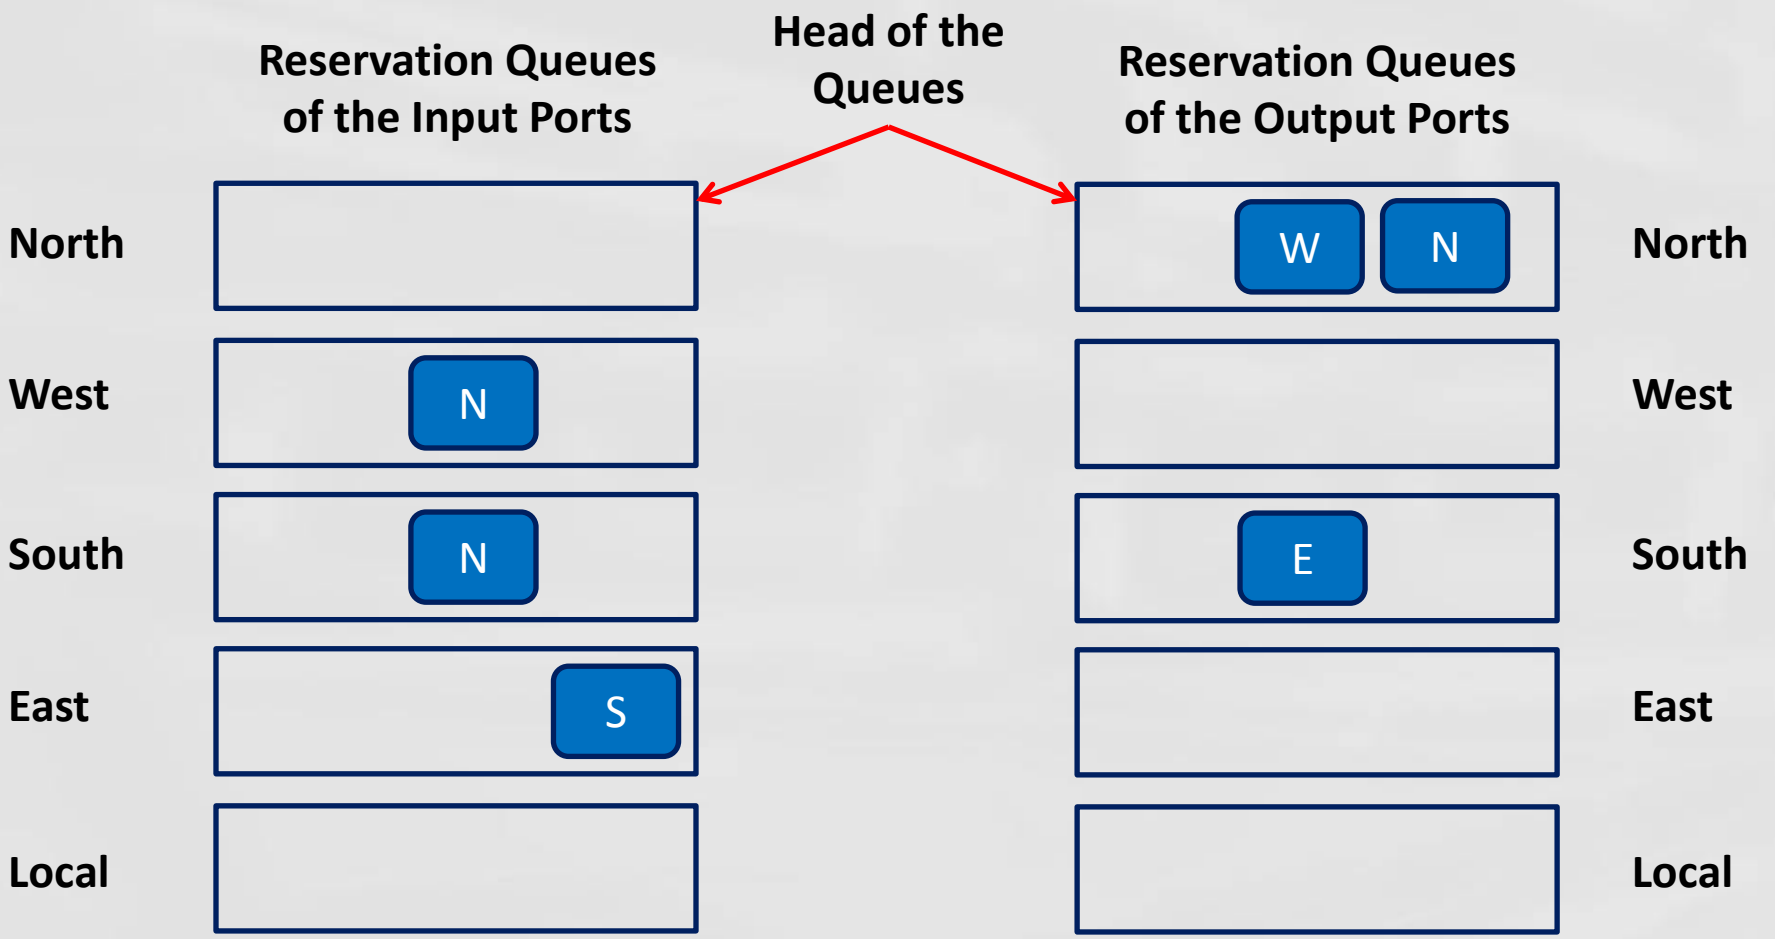

Port Names: North West South East Local

Déjà Vu Switching – Realizing Connections

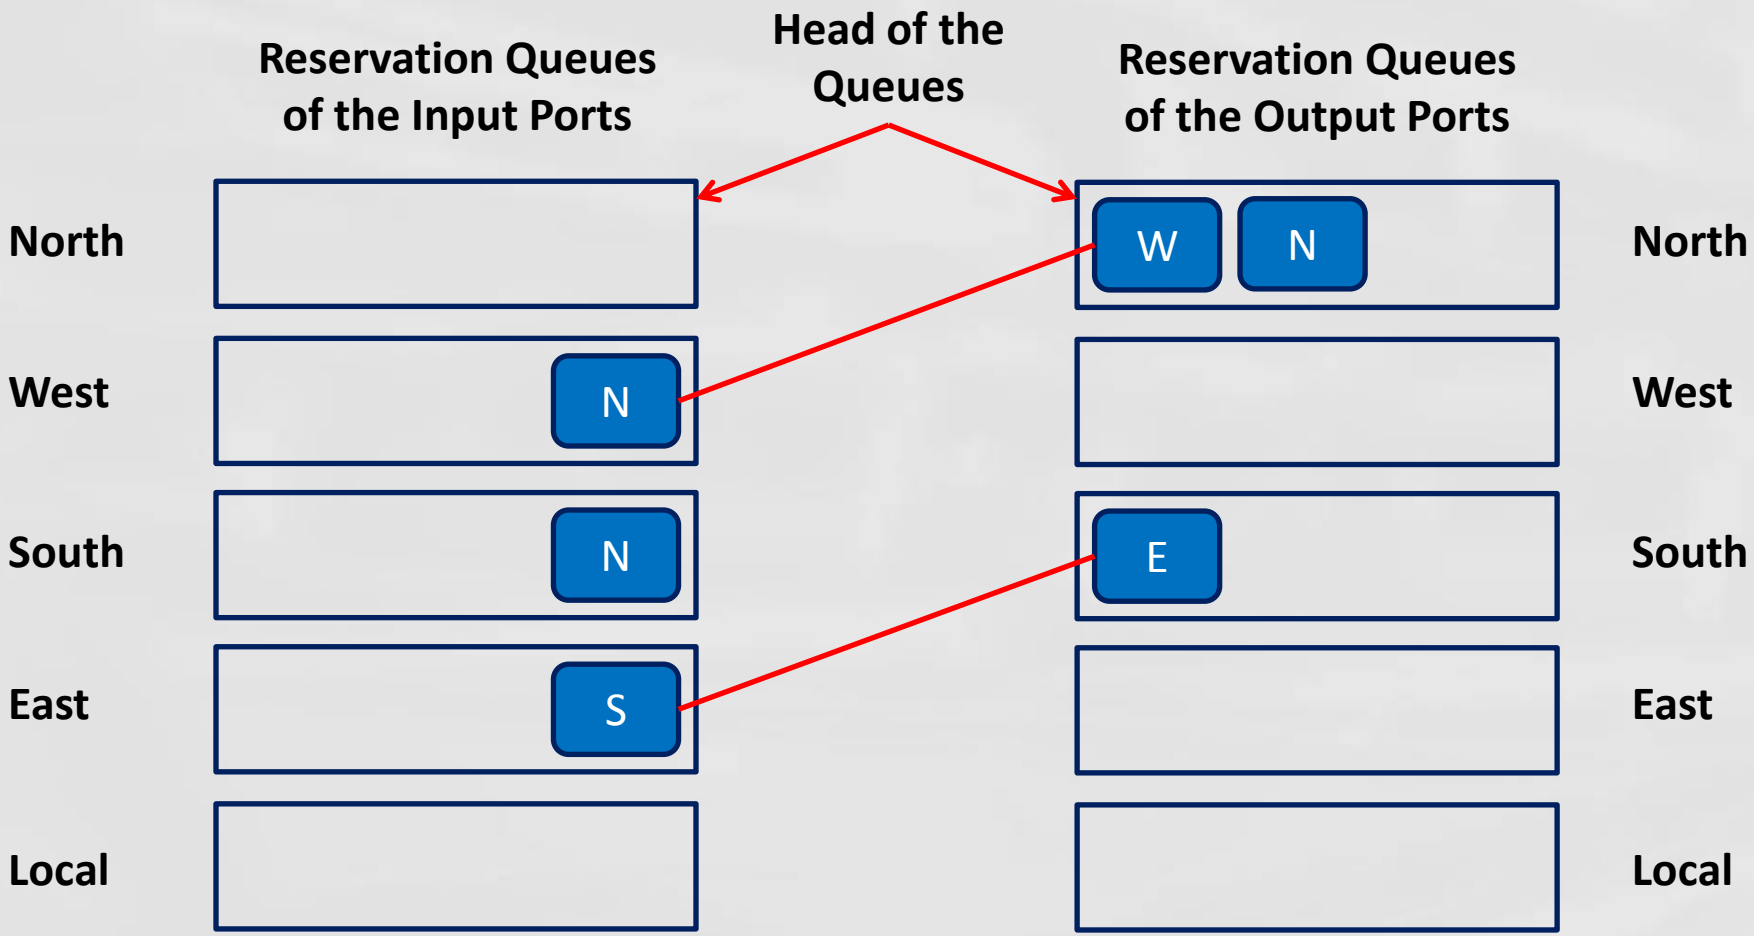

Port Names: North West South East Local

Déjà Vu Switching – Ensuring Correct Routing

- Each input and output port must independently track the reserved circuits it is part of.
- Any two reservation packets that share part of their paths, must traverse all the shared links in the same order
- Data packets must be injected onto the data plane in the same order their reservation packets are injected onto the control plane.

Outline

- Motivation & Related work
- Importance of data messages to performance
- The Optimization Problem
- Proposed Solution
  - Déjà Vu Switching
  - **Analysis of the acceptable data plane speed**
- Evaluation
- Summary

Reservations should always be ahead of data packets

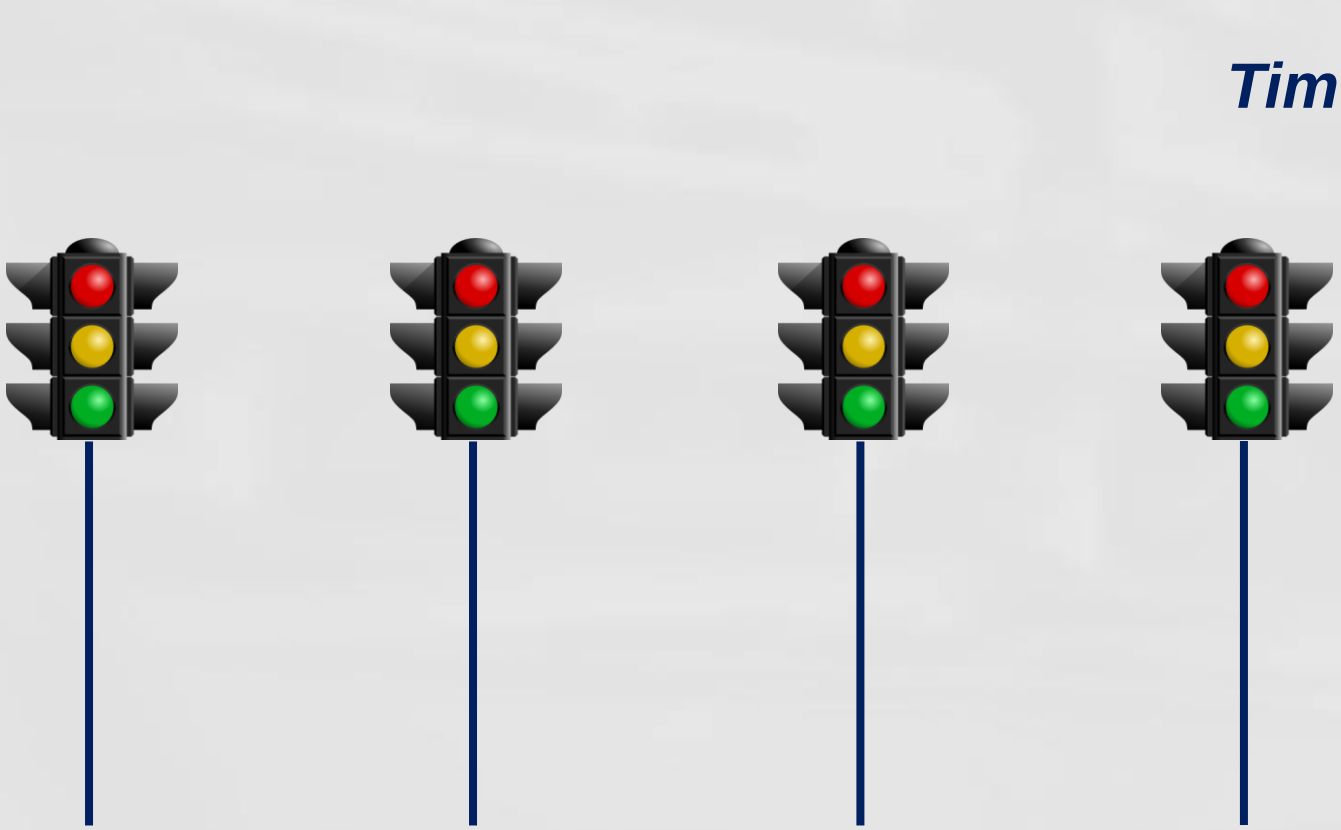

Reservation Packet injected early at cycle 1

Data Packet injected at cycle 3

Reservations should always be ahead of data packets

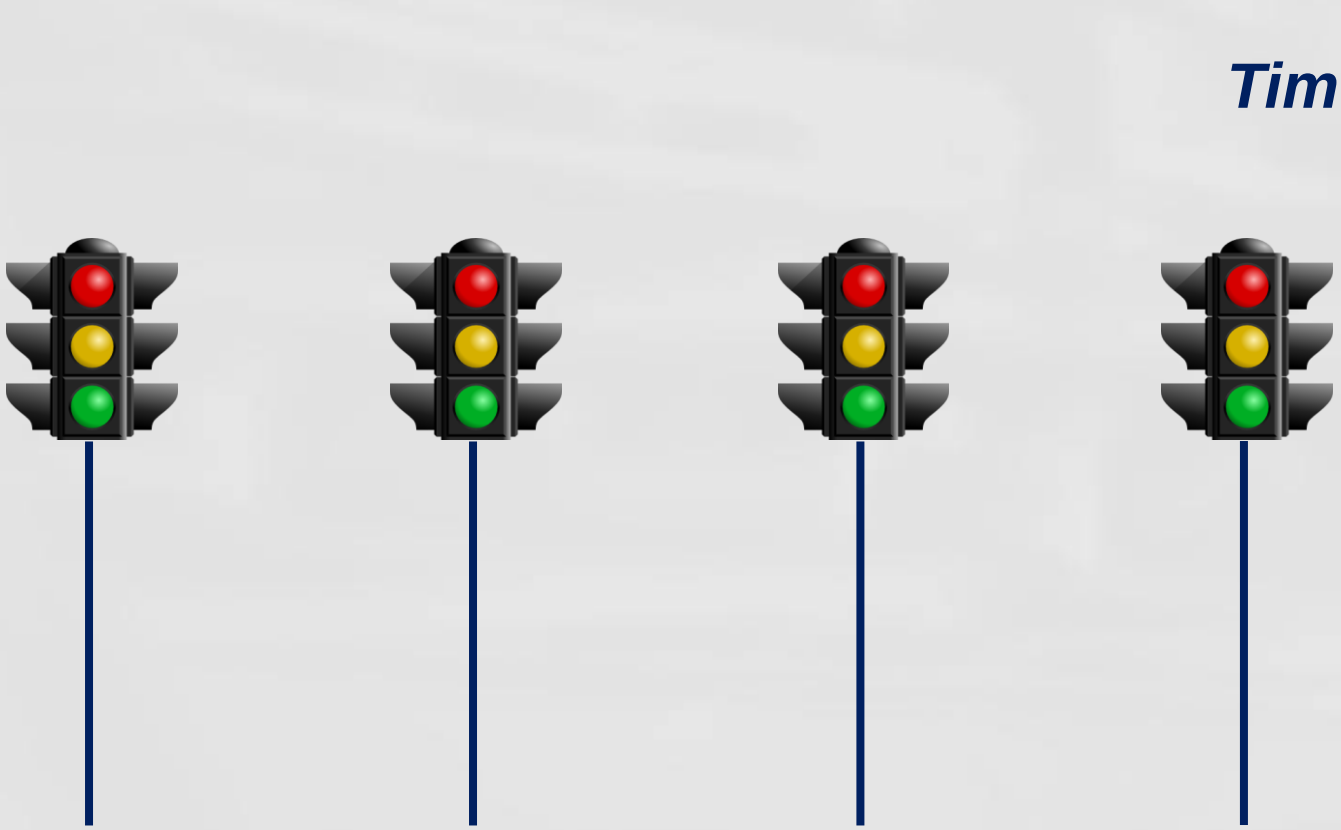

Reservation Packet injected early at cycle 1

Data Packet injected at cycle 3

Acceptable data plane slowdown - Message Latency

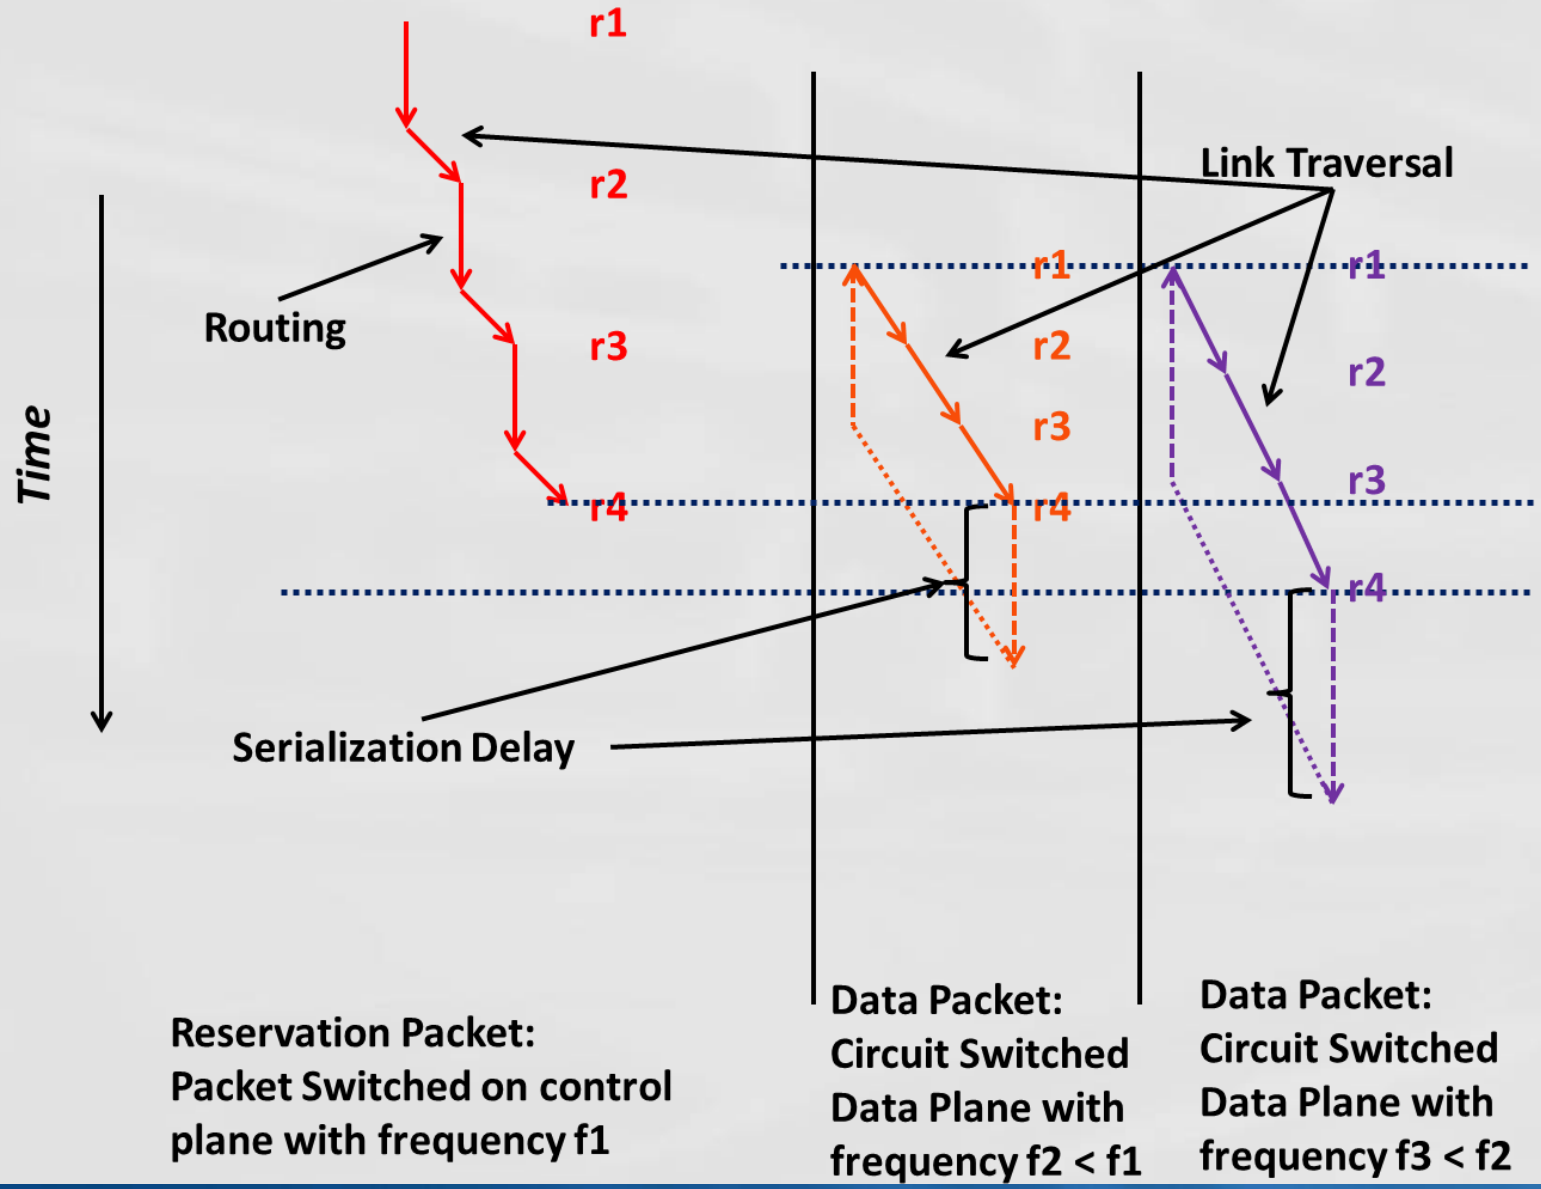

Outline

- Motivation & Related work
- Importance of data messages to performance
- The Optimization Problem
- Proposed Solution
  - Déjà Vu Switching
  - Analysis of the acceptable data plane speed
- **Evaluation**
- Summary

Evaluation Environment

- We use the functional simulator, Simics, to simulate cache coherent CMPs of 16 and 64 cores.
- We use Orion2 to get power numbers for the interconnect routers
- We evaluate with
  - Synthetic traces: allows varying the network load
  - Execution driven simulation of parallel benchmarks from the SPLASH-2 and PARSEC suits
- Interconnect Topology: Mesh

Evaluation Environment

- Baseline NoC:
  - Single plane, 16 byte links, packet switched with 3 cycles router pipeline, clocked at 4 GHz
- Evaluated NoC:
  - Control plane: 6 byte links, packet switched, 4GHz
  - Data plane: 10 byte links, circuit switched.
- Control and Coherence Packets: 1-flit
- Data Packets:
  - Baseline NoC: 5 flits
  - Data Plane: 7 flits

Evaluation – Synthetic Traces

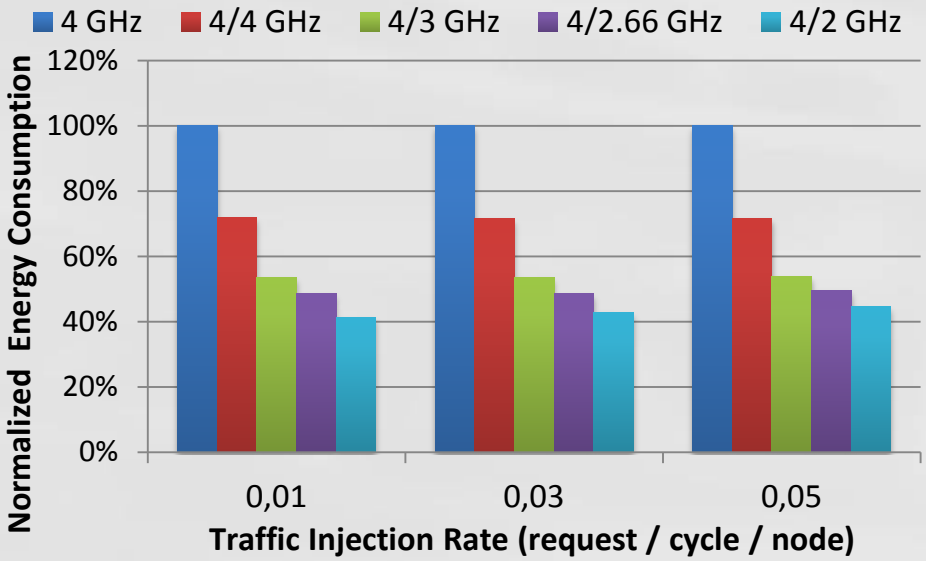

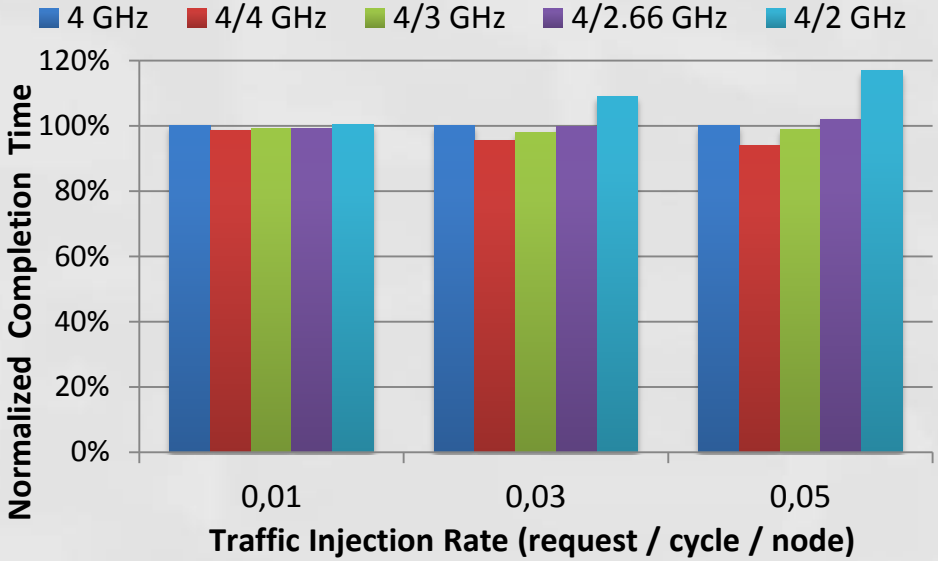

Normalized NoC energy and completion time of *synthetic traces* on a 64-core CMP

*Each request receives a reply data packet

Evaluation – Execution Driven Simulation

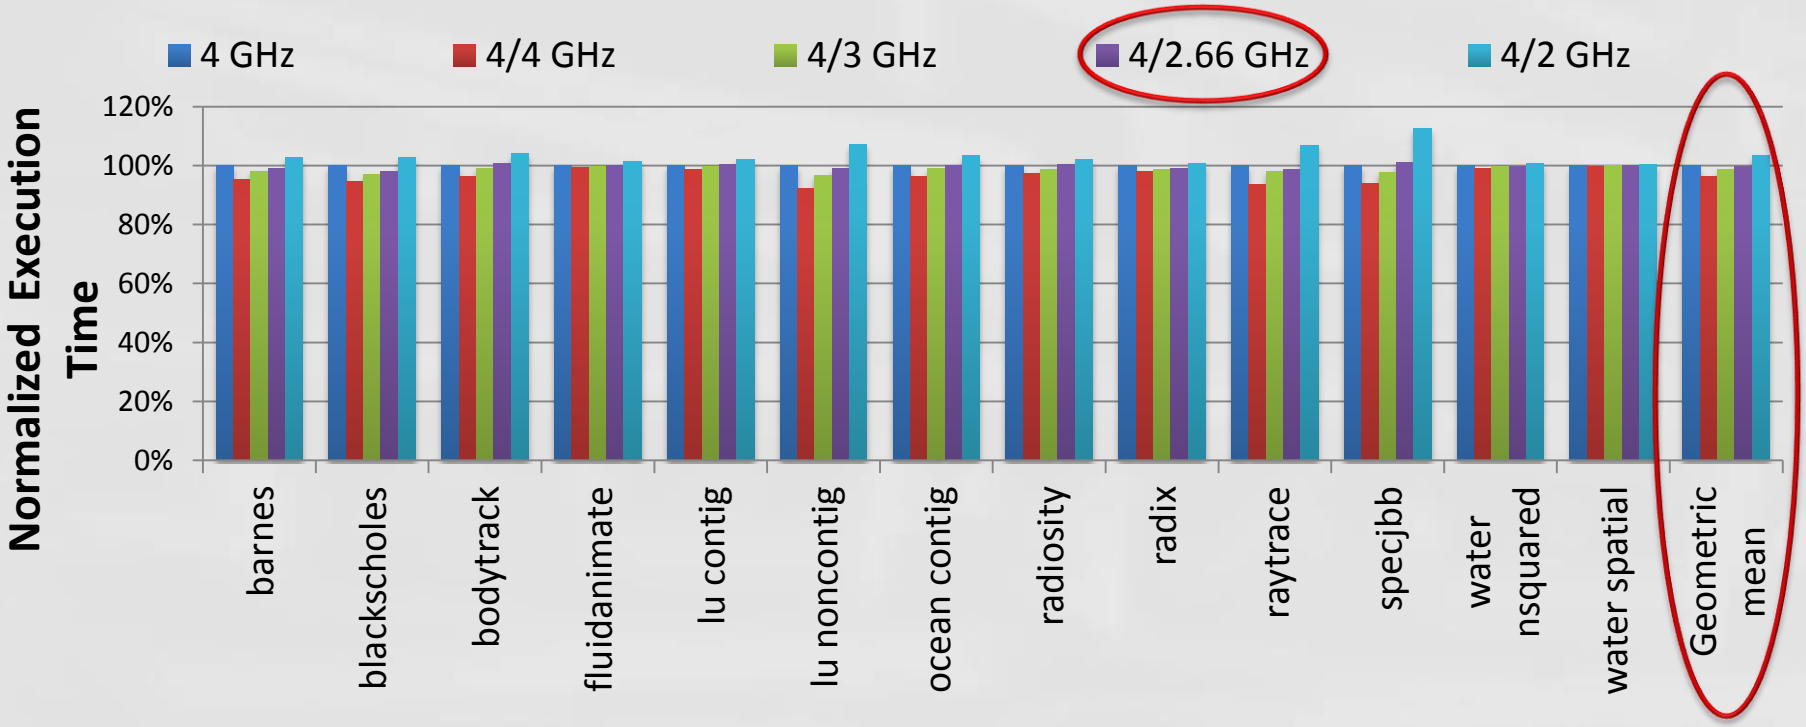

Normalized execution time on a 16-core CMP

Evaluation – Execution Driven Simulation

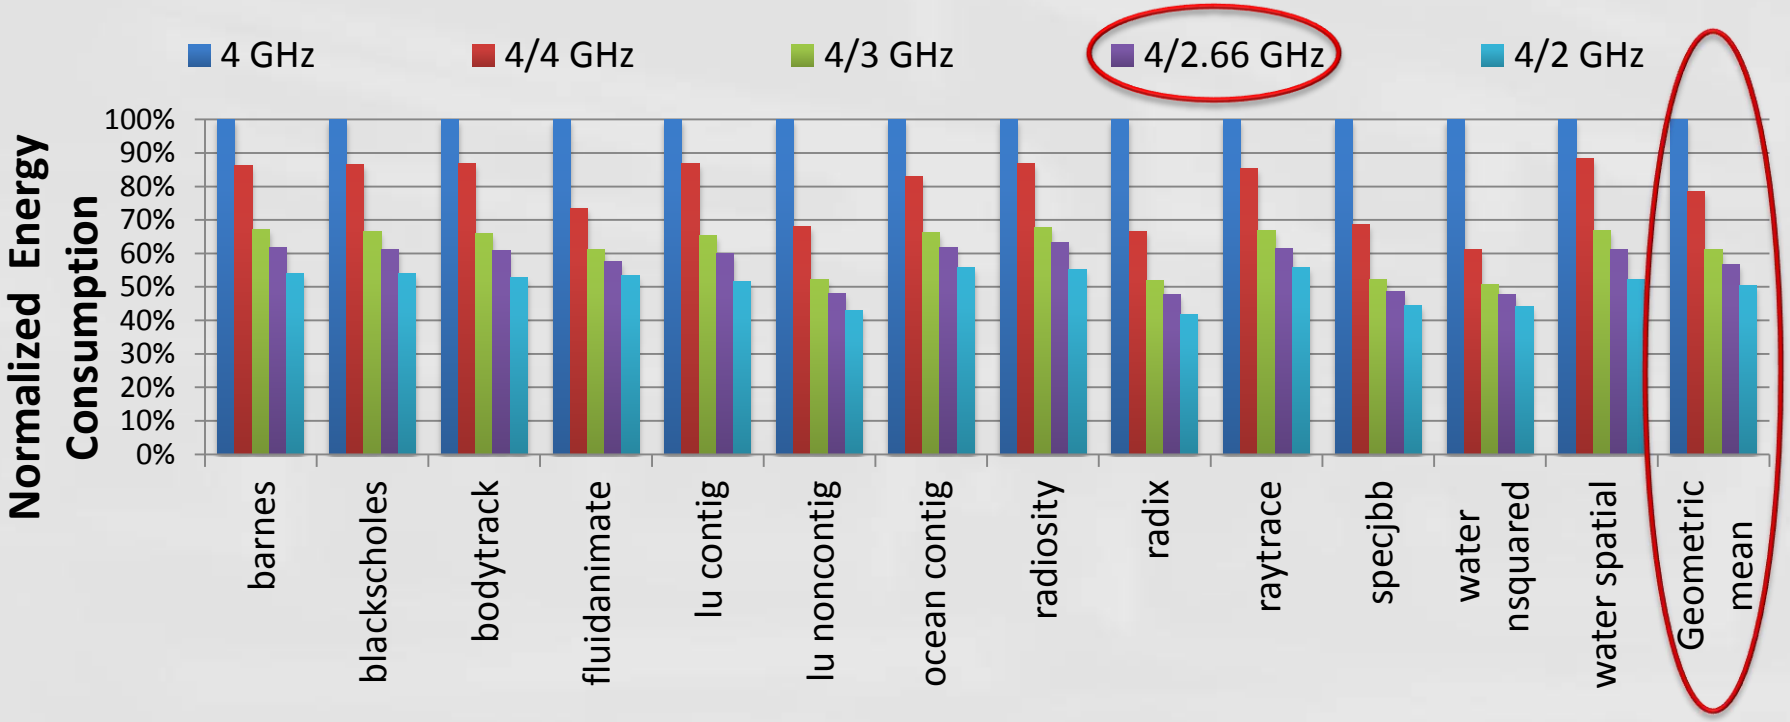

Normalized NoC energy on a 16-core CMP

Evaluation – Execution Driven Simulation

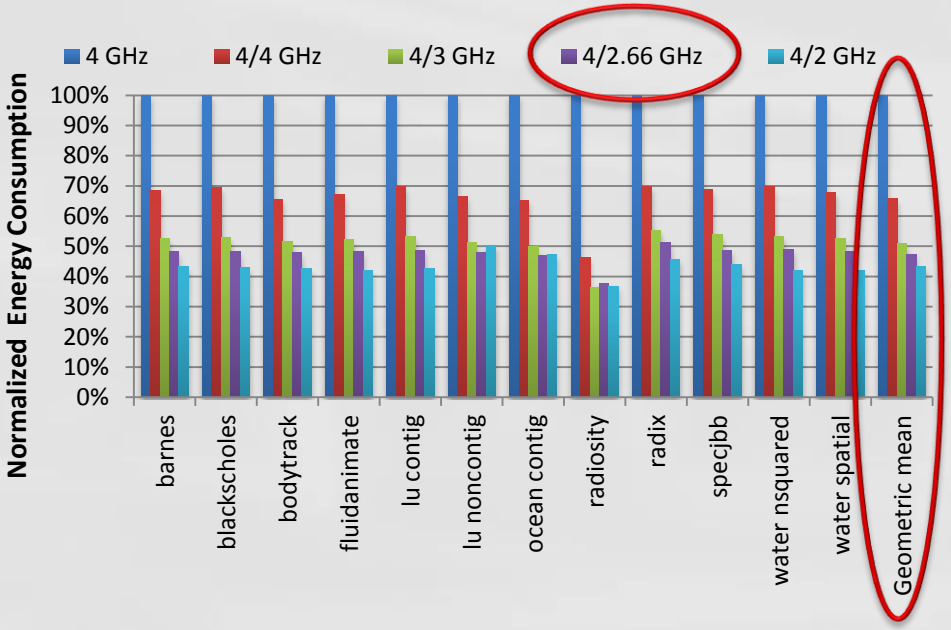

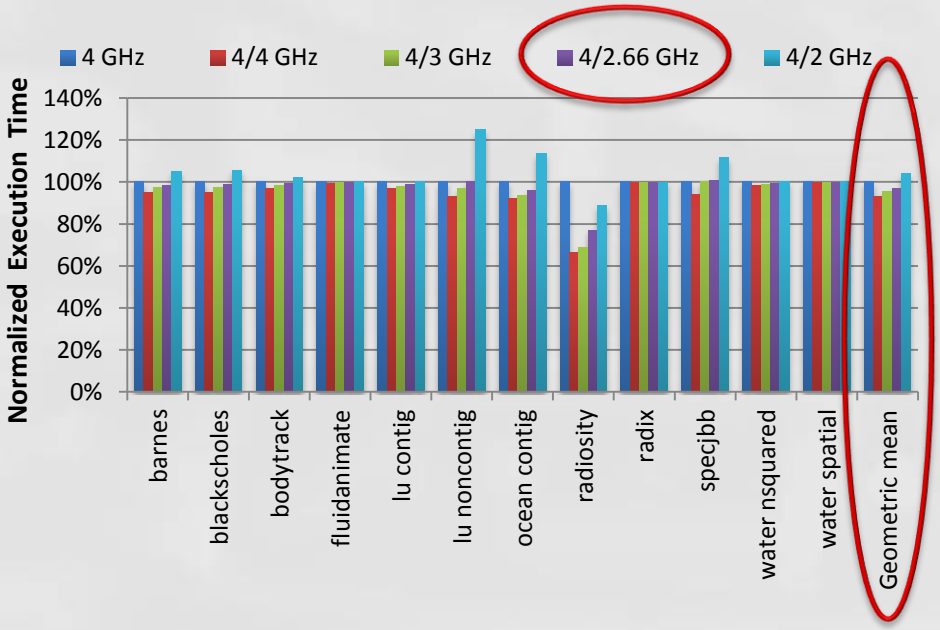

Normalized NoC energy and execution time on a 64-core CMP

Evaluation – Isolating Effect of Déjà Vu Switching

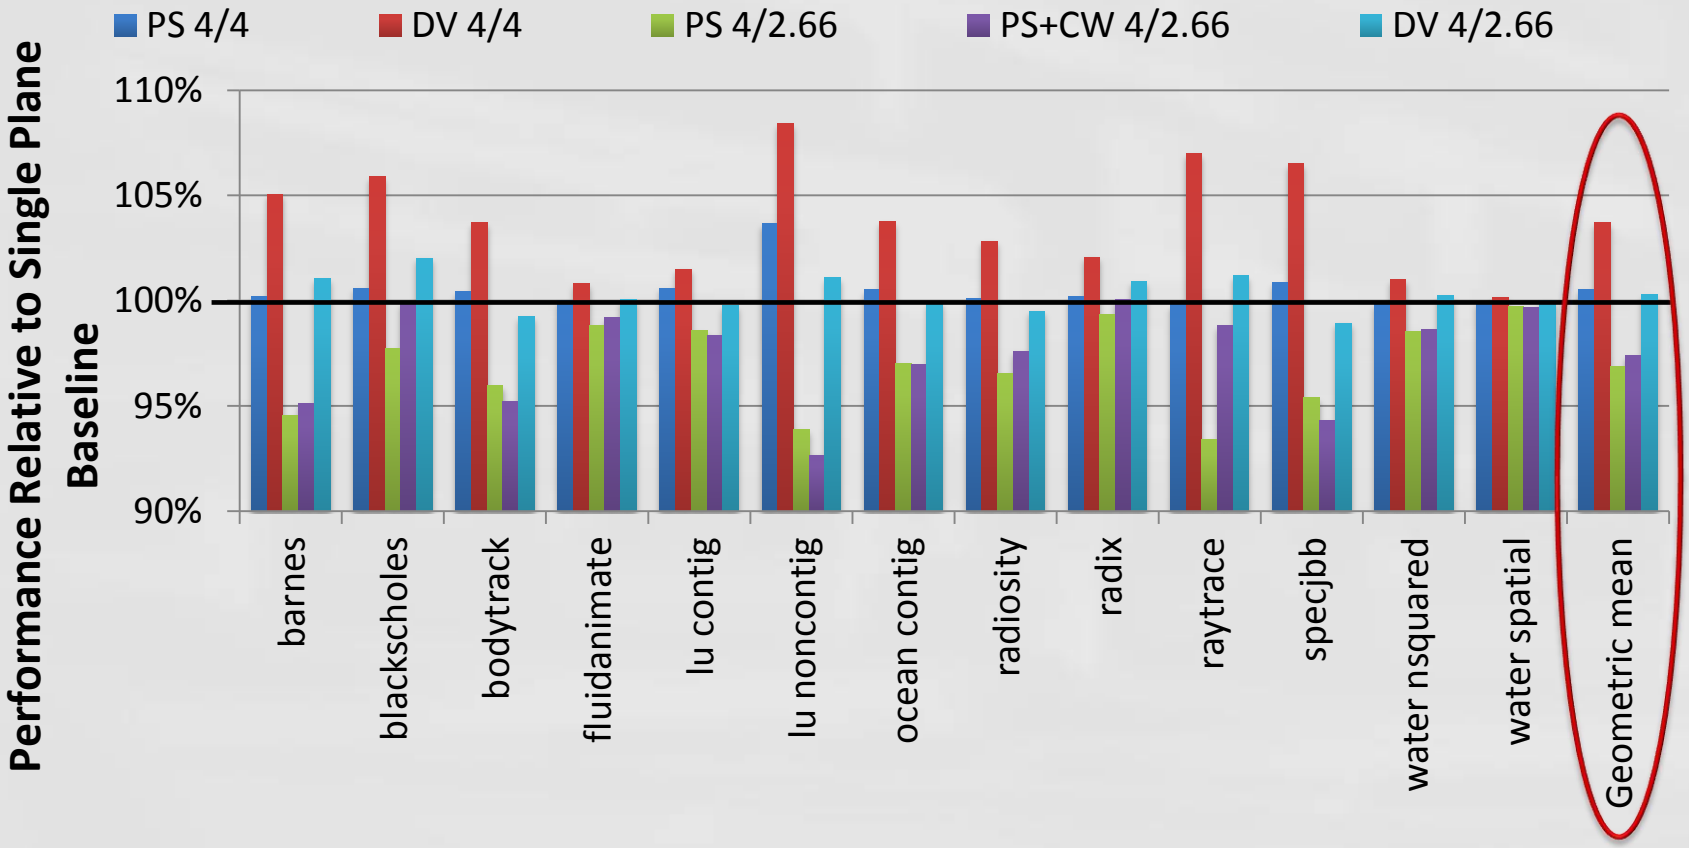

Performance relative to CMP with a single plane NoC on a 16-core CMP

Related Work

- L. Cheng et al. (ISCA'06) and A. Flores et al. (IEEE Trans. Computers'10): Heterogeneous NoC using wires of different latency and power characteristics to improve performance and reduce NoC energy.
- Proposal requires wide links (75 bytes), but performance degrades with narrow links.
- Our work differs in:
  - Tying latency of messages to performance
  - Using Déjà Vu Switching to Compensate for slower data plane

Summary

- Problem: Saving power in the NoC by reducing the data plane's power consumption without impacting performance.
- Delayed cache hits are important to performance.
- Operating data plane in circuit-switched mode allows it to operate at reduced frequency.
- Déjà Vu Switching allows reservation to proceed when resources are not currently available.
- The constraints governing the speed of the data plane can be estimated analytically.

Thank you!