


# Power Efficient Cache Coherence


---

Craig Saldanha  
Mikko Lipasti



# Motivation


- Power consumption becoming a serious design constraint.
- Market demand for faster and more complex servers.
- Complex coherence protocol and interconnect.
- $f_{clock} \uparrow + Complexity \uparrow = P_{interconnect} \uparrow$



# Motivation

---


- Traditional power saving methodologies ineffective.
- Minimize number of transaction packets.
- At the end-points: *Jetty*.
- At the source: *Serial Snoop*.



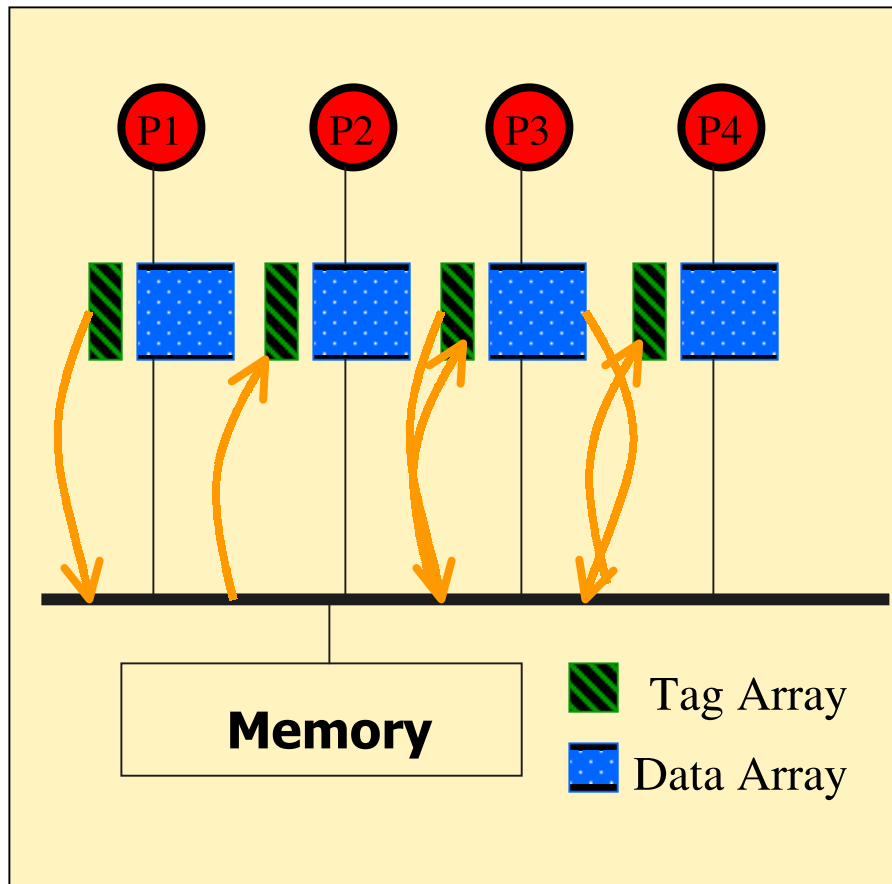
# OUTLINE

---

- Overview of Snoop based protocols and opportunity for power savings
- Latency and Power consumption of parallel snooping techniques
- Serial Snooping
- Results
- Conclusions
- Future Work



# Snoop Based Coherence




## ■ Local Node

- Tag Lookup
- Bus Arbitration
- Snoop Transmission

## ■ Remote Node

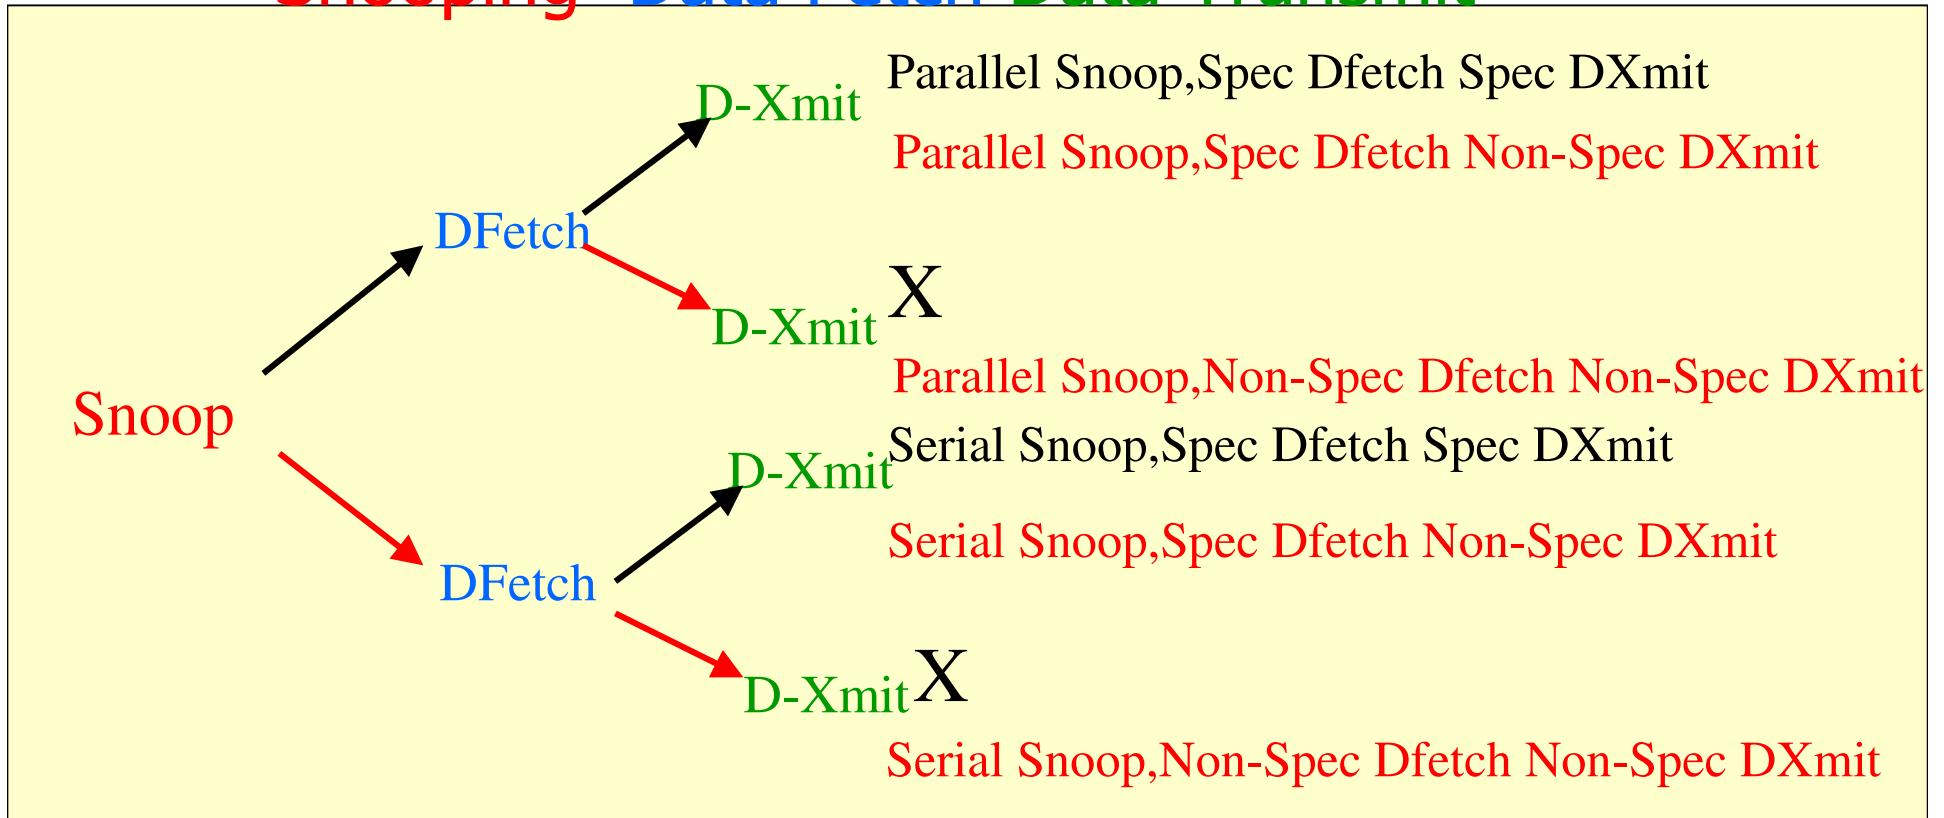
- Tag Lookup
- Snoop Response
- Combination of Responses
- Data Fetch
- Data Transmit




# Degrees of Speculation

- 3 Degrees of Freedom to Speculate

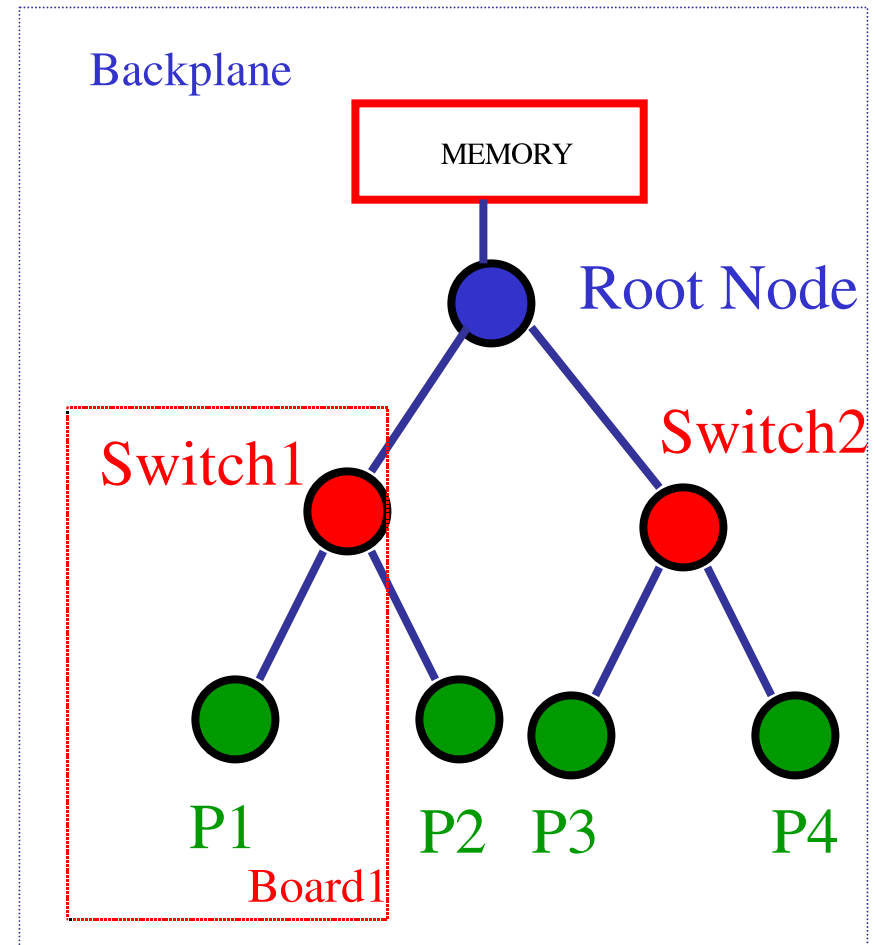
Snooping Data Fetch Data Transmit






# Latency and Power assumptions

- Consider only load misses
- Tree of point-point connections.
- Latency to traverse a link: 1Bus cycle (7ns)
- Tag Look up :1 bus cycle (7ns)
- D-Fetch: 2 bus cycles (14ns)
- DRAM access: 10 bus cycles (70ns)





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Snoop Broadcast

*Latency*




*Power*

*Plink*







# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Snoop Broadcast


*Latency*



*Power*

Plink+Pswitch





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Snoop Broadcast


*Latency*



*Power*

$2P_{link} + P_{switch}$





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Snoop Broadcast


*Latency*



*Power*

2Plink+2Pswitch






# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Snoop Broadcast


*Latency*



*Power*

5Plink+2Pswitch





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)

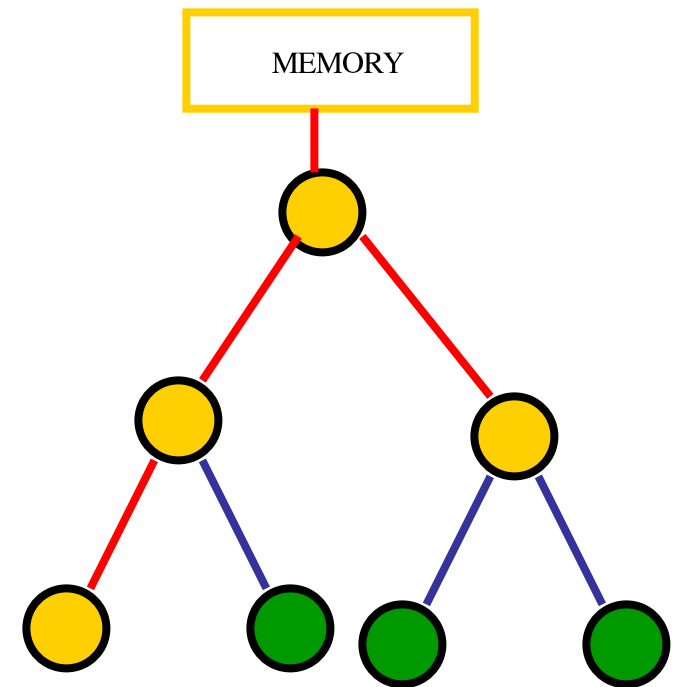
## Snoop Broadcast


*Latency*



*Power*

5Plink+4Pswitch






# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Snoop Broadcast


*Latency*



*Power*

8Plink+4Pswitch





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Memory Access :Data Fetch


*Latency*



*Power*

$P_{mem}$






# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Memory Access : Data Transmit

*Latency*




*Power*

$$P_{\text{mem}} + 3P_{\text{link}} + 2P_{\text{switch}}$$








# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Remote Node: Tag Lookup


*Latency*



*Power*

$3P_{tag}$





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Remote Node :Snoop Response


*Latency*



*Power*

$$3P_{tag} + 3*(P_{switch} + 2P_{link}) + P_{link} + 2P_{switch} + 2P_{link}$$





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Remote Node: Data Fetch


*Latency*



*Power*

3Pcache





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Remote Node :Data Transmit


*Latency*



*Power*

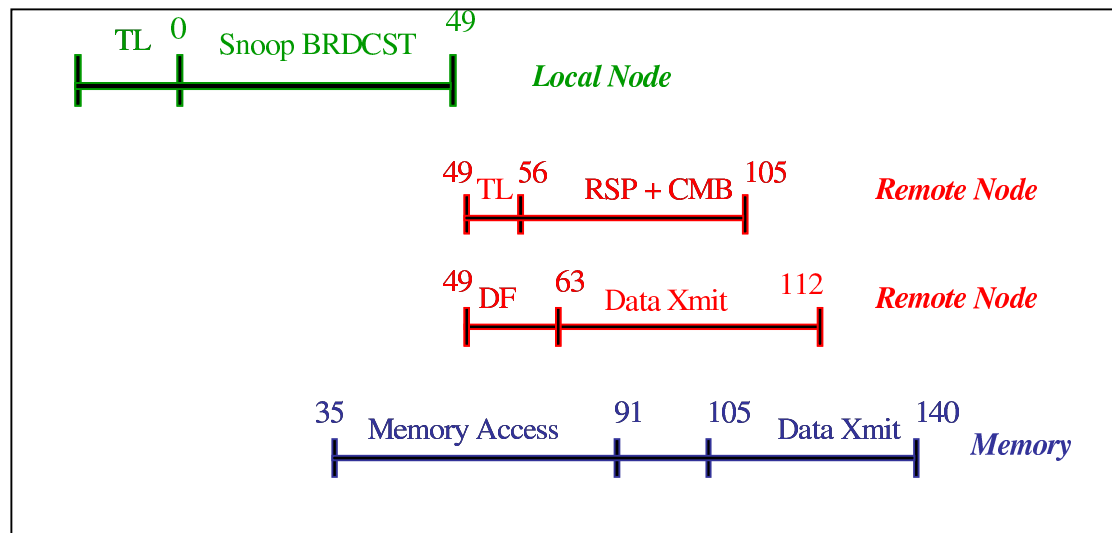
$$3P_{tag} + 3*(4P_{link} + 3P_{switch})$$





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)


## Latency



## Power

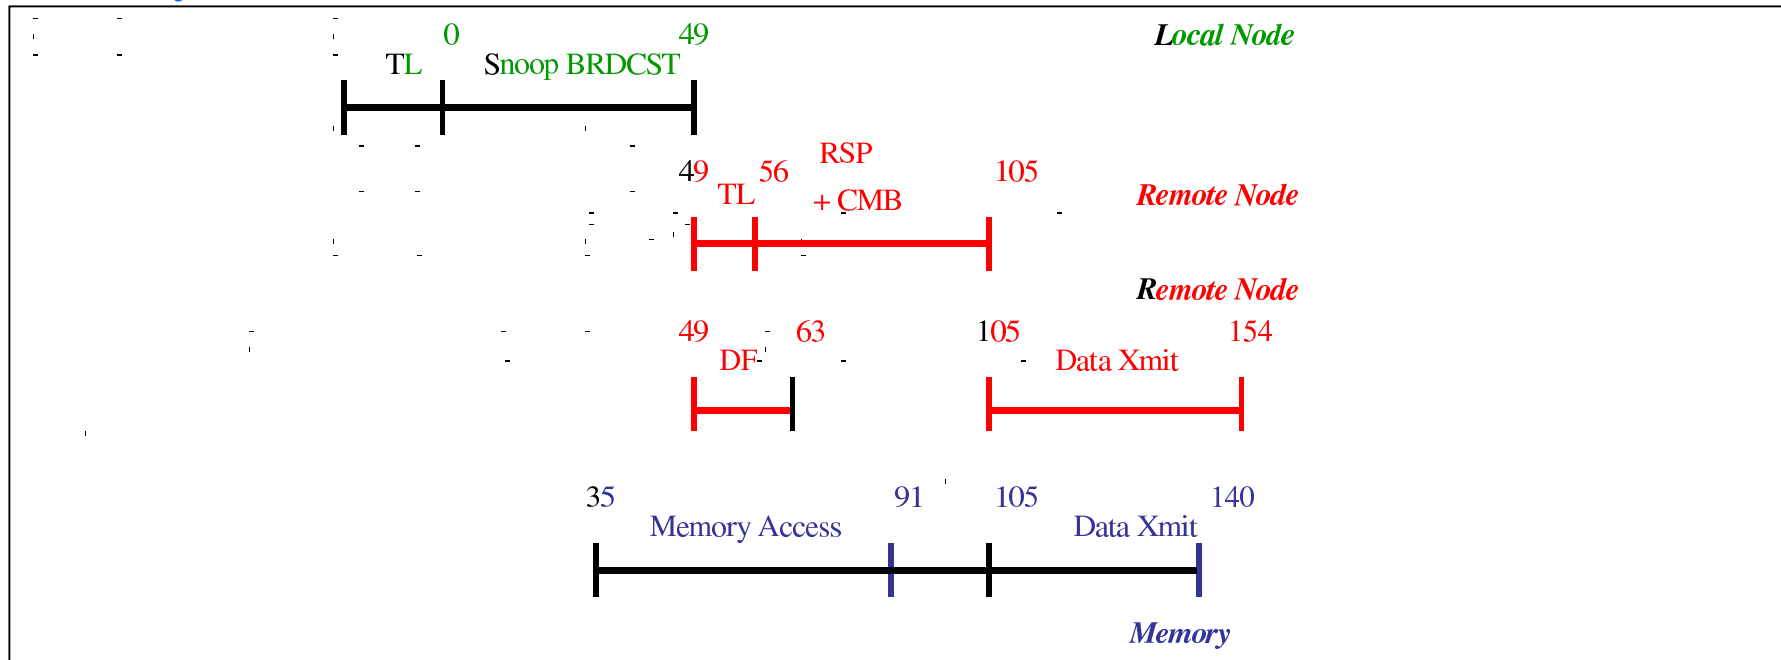
Remote Node supplies the data  $29P_{link} + 18P_{switch} + 3P_{tag} + 3P_{cache} + P_{mem}$

Memory supplies the data  $20P_{link} + 11P_{switch} + 3P_{tag} + 3P_{cache} + P_{mem}$



# Parallel Snoop Speculative Fetch Non-Speculative Transmit (PS/SF/NT)

## Latency




## Power

Remote Node supplies the data

$$21P_{link} + 12P_{switch} + 3P_{tag} + P_{cache} + P_{mem}$$

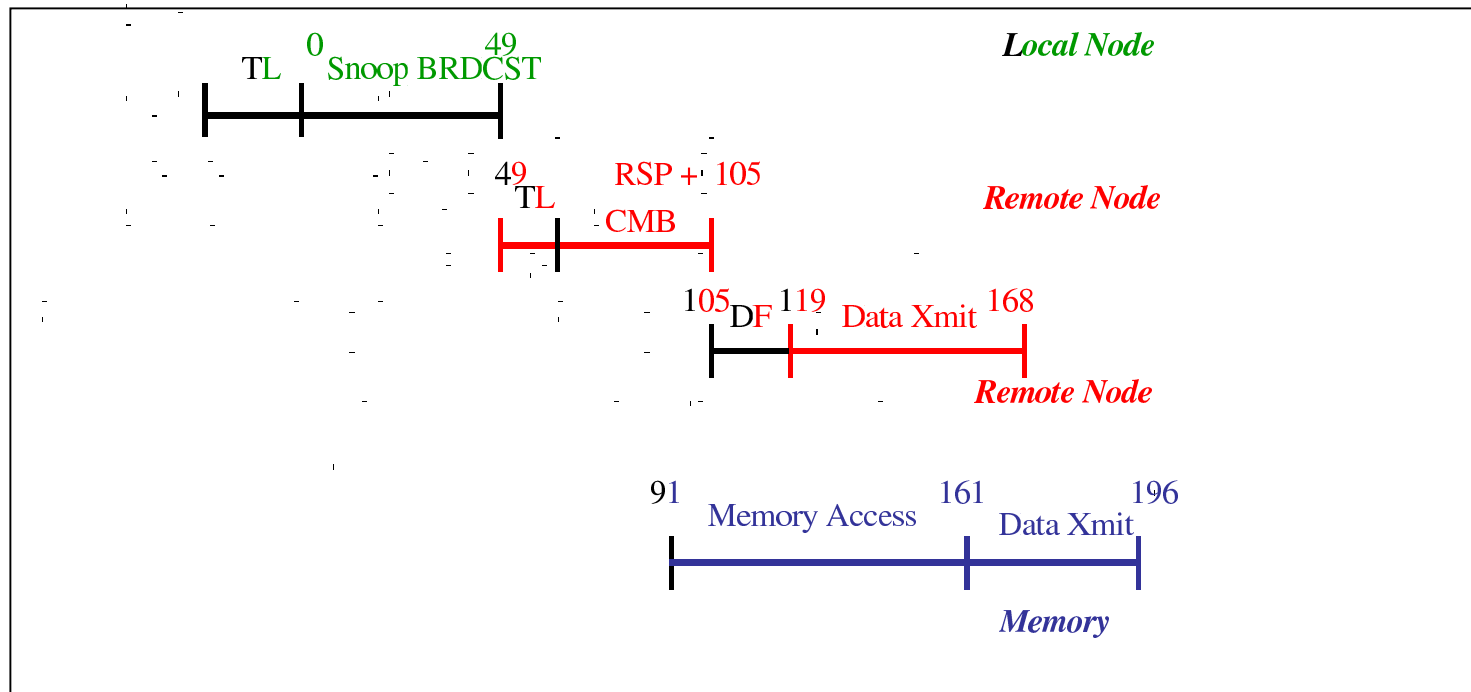
Memory supplies the data

$$20P_{link} + 11P_{switch} + 3P_{tag} + 3P_{cache} + P_{mem}$$



# Parallel Snoop Non-Speculative Fetch Non-Speculative Transmit (PS/NF/NT)

## Latency




## Power

Remote Node supplies the data

$$21P_{link} + 12P_{switch} + 3P_{tag} + P_{cache}$$


Memory supplies the data

$$20P_{link} + 11P_{switch} + 3P_{tag} + P_{mem}$$




## Serial Snooping

- Avoids Speculative transmission of Snoop packets.









## Serial Snooping


- Avoids Speculative transmission of Snoop packets.
- Check the nearest neighbor
- Data supplied with minimum latency and power






# Serial Snooping


- Forward snoop to next level






# Serial Snooping


- Forward snoop to next level






# Serial Snooping


- Search other half of tree






# Serial Snooping


- Search other half of tree
- Search leaf nodes serially






## Serial Snooping

- Search other half of tree
- Search leaf nodes serially






## Serial Snooping : Features

---


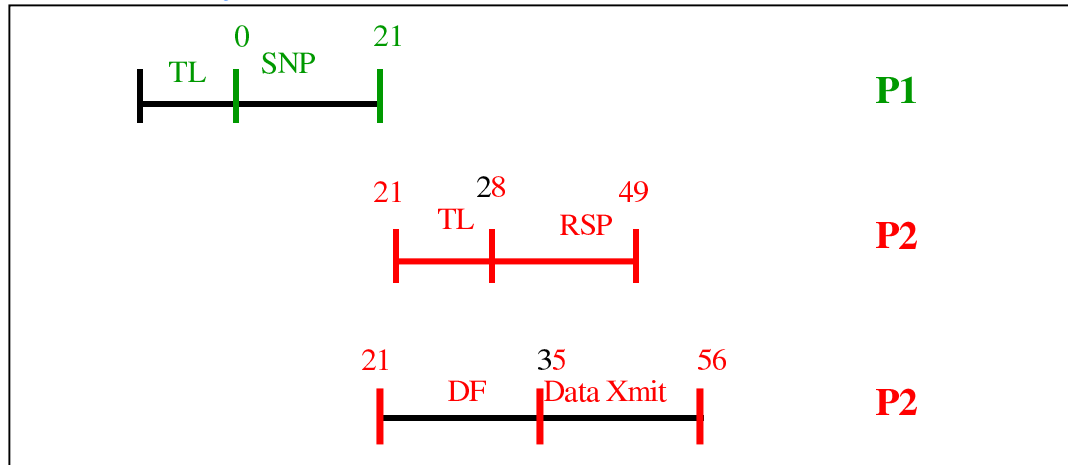
- Latency to satisfy a request dependent on distance from requestor.
- Data resident at the nearest neighbor supplied with the lowest latency and power.
- Requests visible to memory controller only at root node.
- Latency is adversely affected when requested data present at the farthest node
- Worst case power consumption is still less than the parallel snooping.



# Serial Snooping :

## Request satisfied by Nearest Node

### Latency



### Power


Xmit Snoop:  $2P_{link} + P_{switch}$

P2 Tag access and snoop response:  $P_{tag} + 2P_{link} + P_{switch}$

P2 Data Fetch and Xmit:  $P_{cache} + 2P_{link} + P_{switch}$

$P_{total} = 6P_{link} + 3P_{switch} + P_{tag} + P_{cache}$


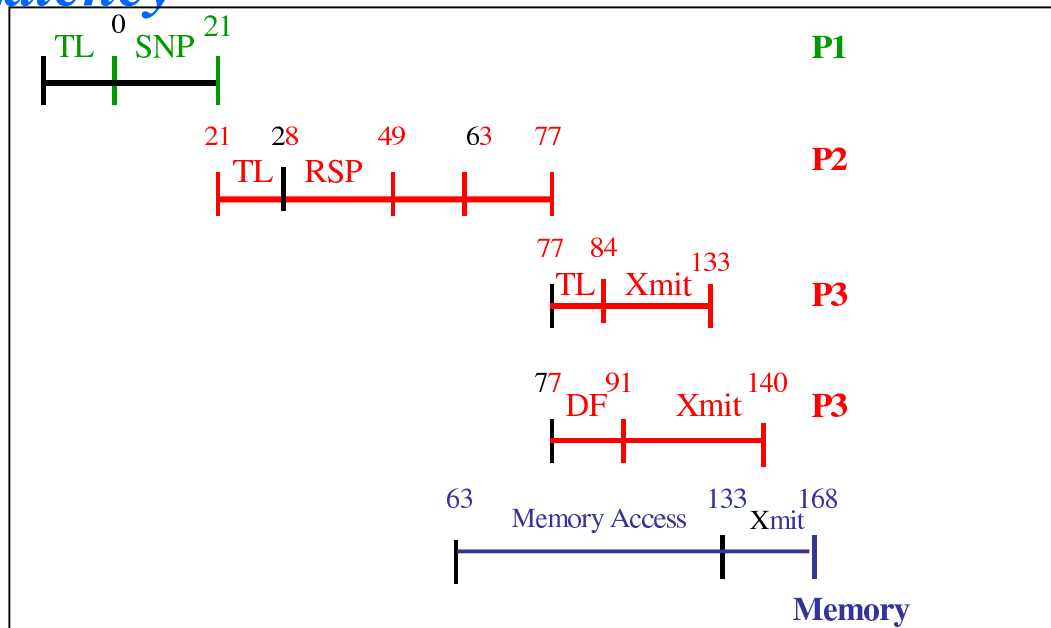




# Serial Snooping :


## Request satisfied by Next-Nearest Neighbor

### Latency



### Power


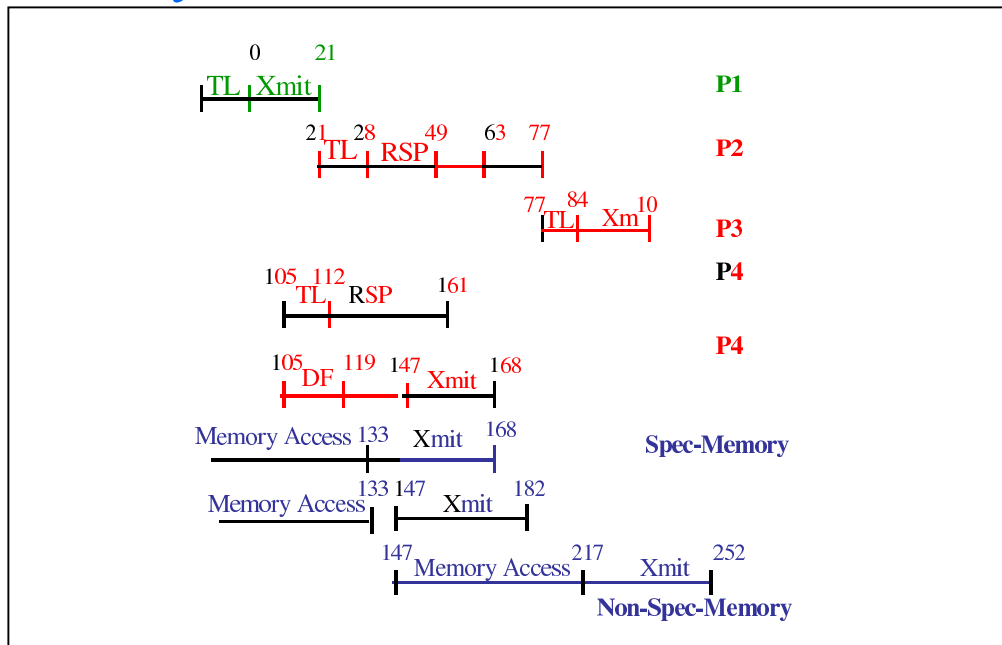
$$16P_{link} + 10P_{switch} + 2P_{tag} + P_{cache}$$



# Serial Snooping :

## Request satisfied by farthest node

### Latency



### Power


Remote Node supplies the data

$$18P_{link} + 11P_{switch} + 3P_{tag} + P_{cache}$$


If Memory supplies the data

$$17P_{link} + 10P_{switch} + 3P_{tag} + P_{mem}$$

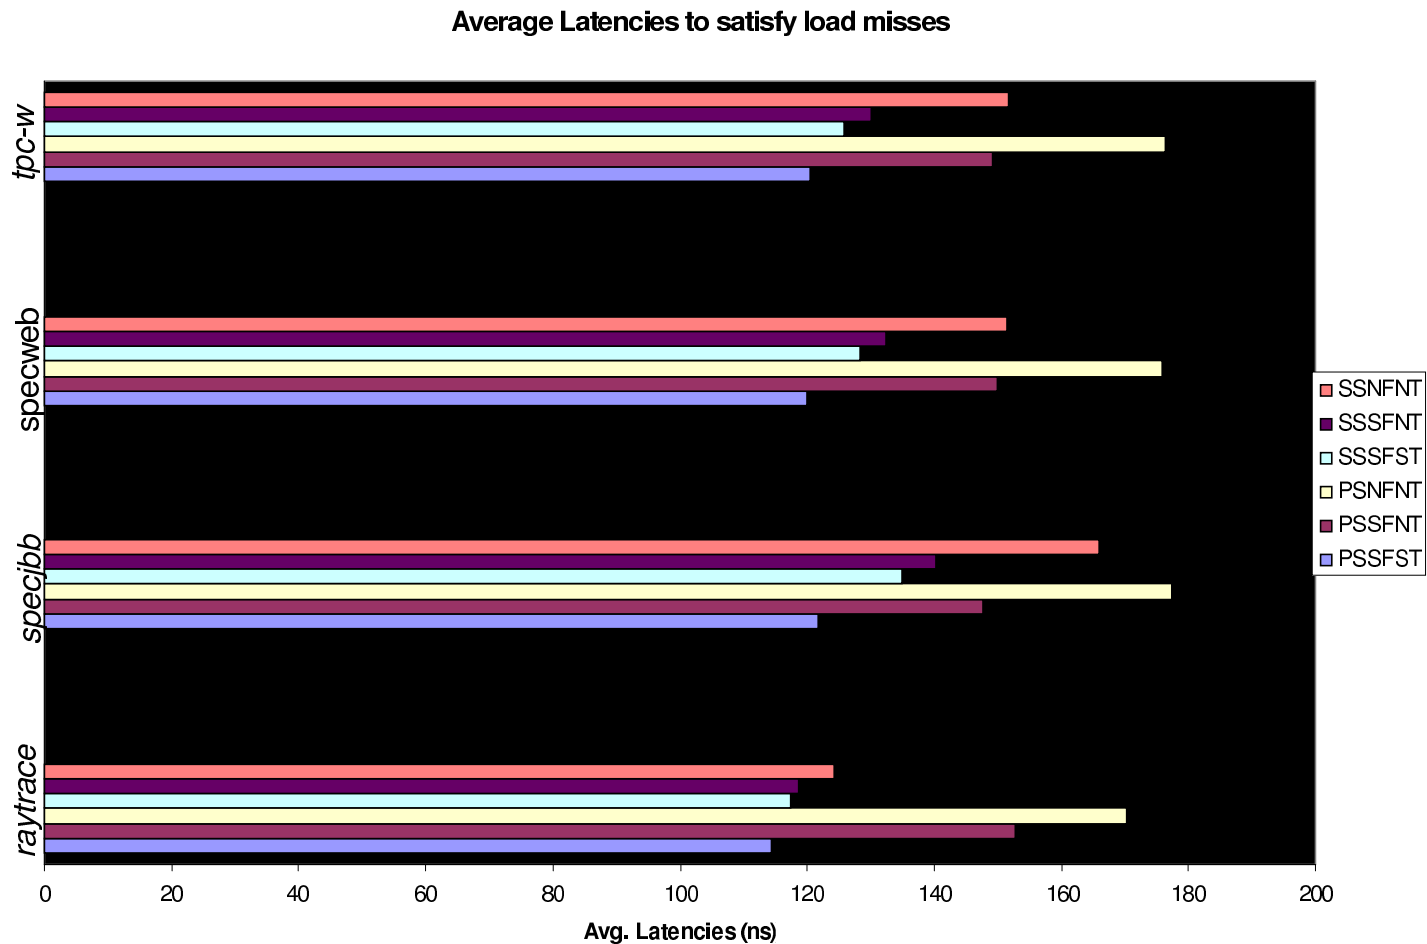

# RESULTS : Load Miss Distributions




Load Miss Distribution




# RESULTS: Average Latencies to satisfy load misses




# RESULTS: Relative Power Savings



Factors contributing to power consumption






# CONCLUSIONS

---


- Reducing degree of speculation has potential for significant power savings
- Performance degradation is minimal for the set of benchmarks studied.
- Serial Snooping with speculative memory fetch provides optimal latency and power consumption.



# Future Work

---


- Develop detailed execution-driven Power Model
- Explore different interconnect topologies.
- Examine the viability of adaptive mechanisms for protocol policy.



# Serial Snooping


## Nearest Neighbor

*Latency*




*Power*

$2P_{link} + P_{switch}$









# Questions

---





# Parallel Snoop Speculative Fetch Speculative Transmit (PS/SF/ST)

