


# Designing Energy-efficient Low-Diameter On-chip Networks with Equalized Interconnects


---

Ajay Joshi, Byungsub Kim and Vladimir Stojanović


HOTI 2009



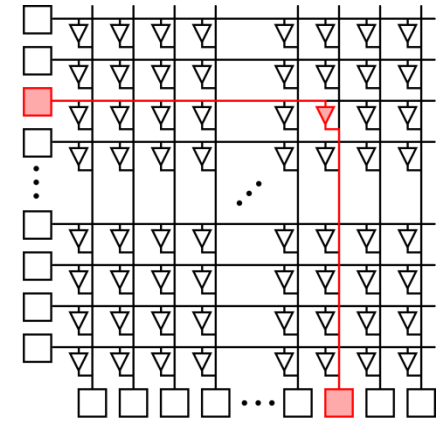

# Our target manycore system



# On-chip network topology spectrum



Mesh



Crossbar

Shorter wires

- Low area
- Low latency
- Low power dissipation


Difficult to map applications

Longer wires



- High area
- High latency
- High power dissipation

Easier to map applications

# Outline




Mesh



Crossbar

- ❑ Repeater-inserted interconnect vs Equalized interconnects
- ❑ Network design space exploration
- ❑ Summary

# Repeater-inserted pipelined interconnects




## □ Design constraints



- 32 nm technology
- 200 ps stage latency
- 5 Gbps throughput

## □ Design parameters

- Wire width
- Repeater size
- Repeater spacing



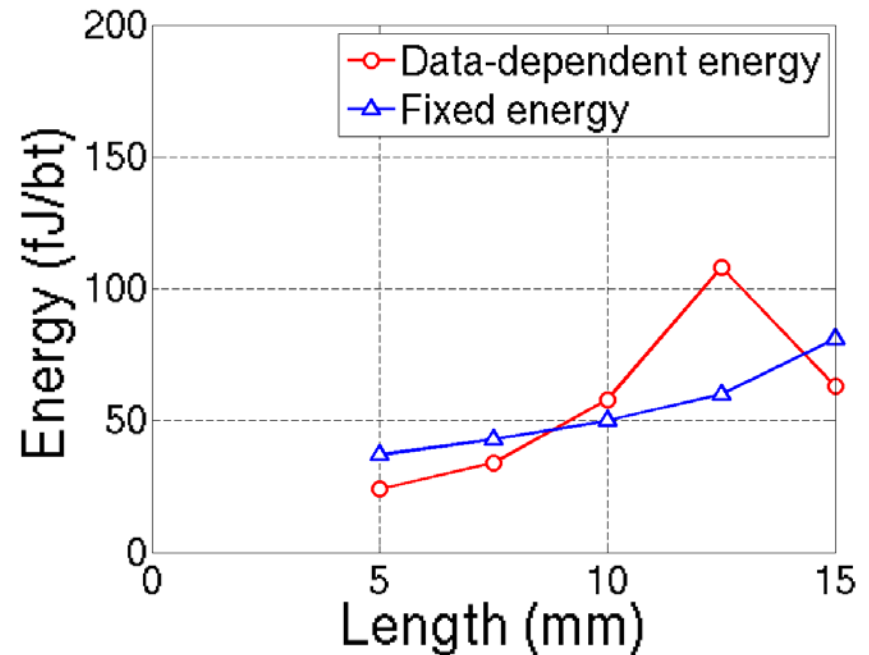
# Equalized interconnects



- ❑ FFE to shape transmitted pulse, DFE to cancel first trailing ISI tap
- ❑ Lower energy cost due to output voltage swing attenuation

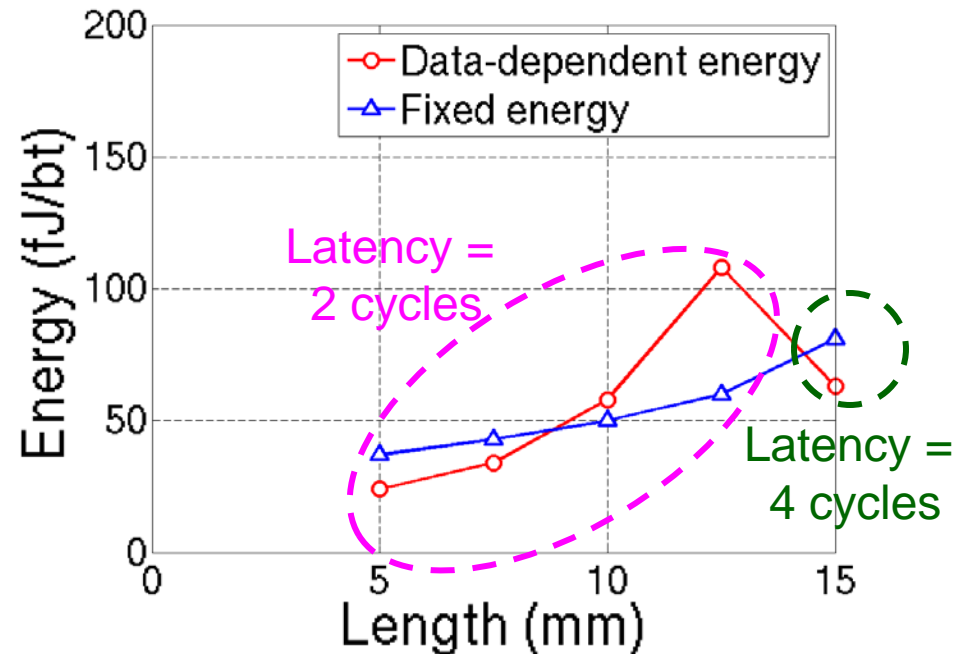
# Equalized interconnects

- ❑ Design constraints
  - 32 nm technology
  - 200 ps stage latency
  - 5 Gbps throughput
- ❑ Tx, Rx and wire dimensions jointly optimized to minimize power
- ❑ Long wires pipelined to minimize power

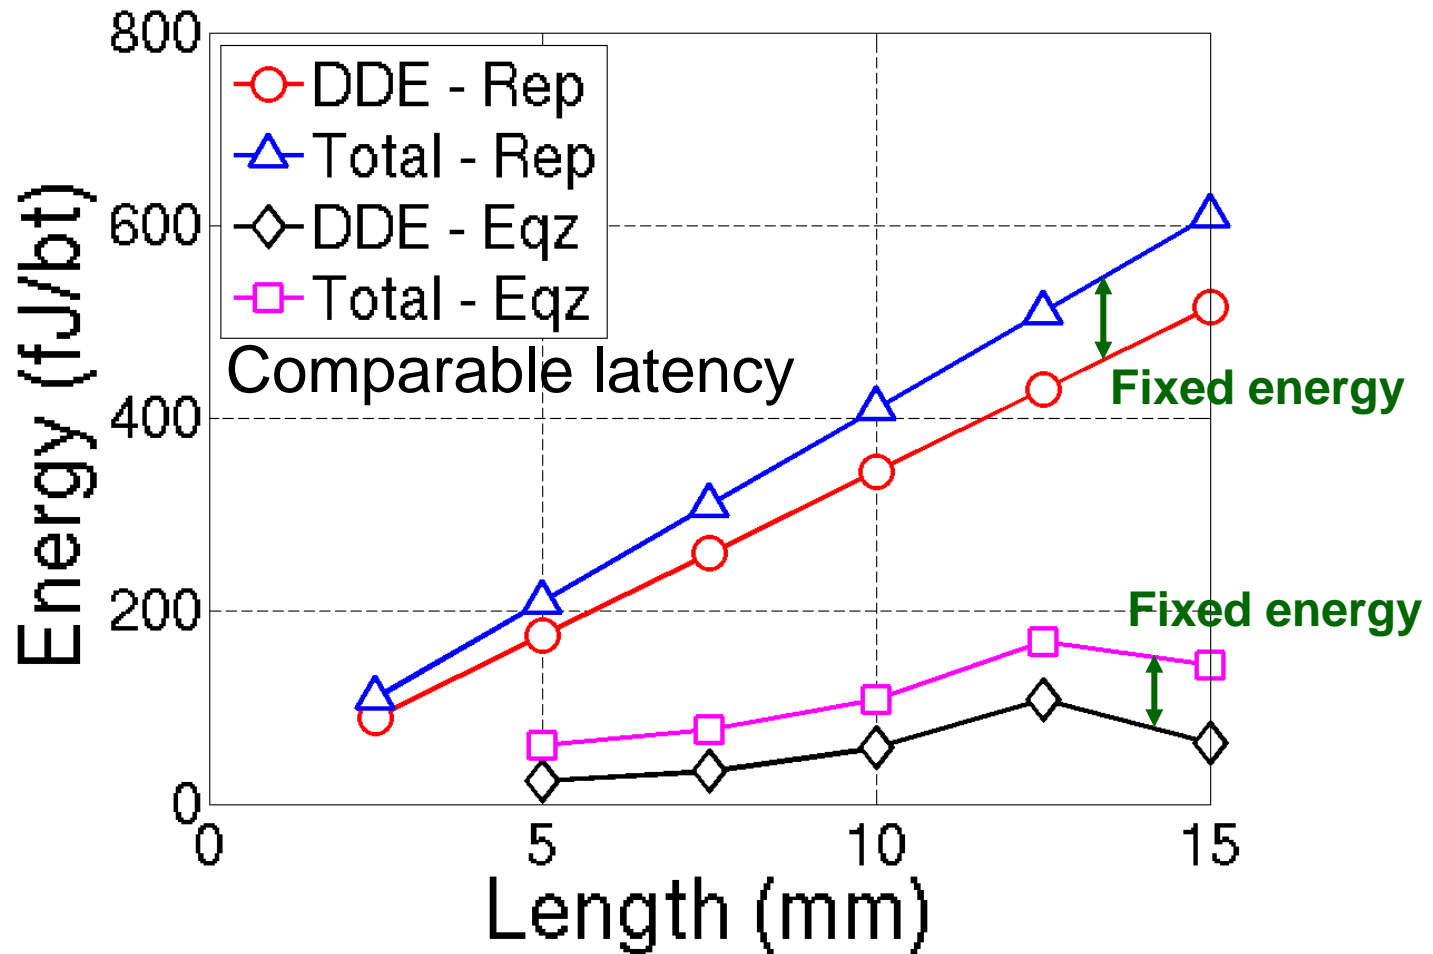


# Equalized interconnects

- Design constraints
  - 32 nm technology
  - 200 ps stage latency
  - 5 Gbps throughput
- Tx, Rx and wire dimensions jointly optimized to minimize power
- Long wires pipelined to minimize power




# Repeated interconnects vs Equalized interconnects




Data-dependent energy (DDE) is 4-10x lower for equalized interconnects, while fixed energy (FE) is comparable

# Outline




- ❑ Repeater-inserted interconnect vs Equalized interconnects
- ❑ **Network design space exploration**
- ❑ Summary

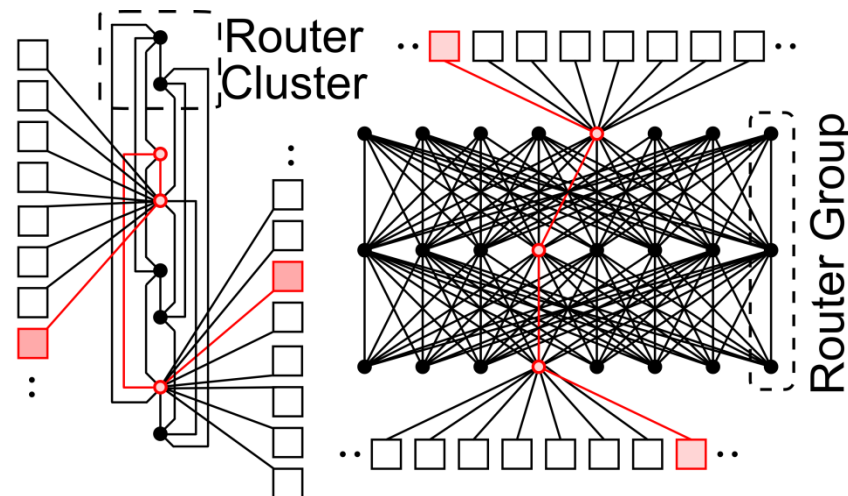
# Network topologies under consideration



Mesh



CMesh




Flattened Butterfly

Clos


- ❑ Channel widths chosen using bisection bandwidth criteria for uniform random traffic
- ❑ Cycle-accurate micro-architectural simulator
  - Accounts for pipeline latencies, router contention, flow control, etc
  - Capture events – channel utilization, switch traversal to calculate power

# Partition application model [Joshi'09]


- ❑ Tiles divided into logical partitions and communication is within partition
- ❑ Logical partitions mapped to physical tiles
  - Co-located tiles → Local traffic
  - Distributed tiles → Global traffic




Uniform random (UR)



2 tiles per partition that are distributed across the chip (P2D)

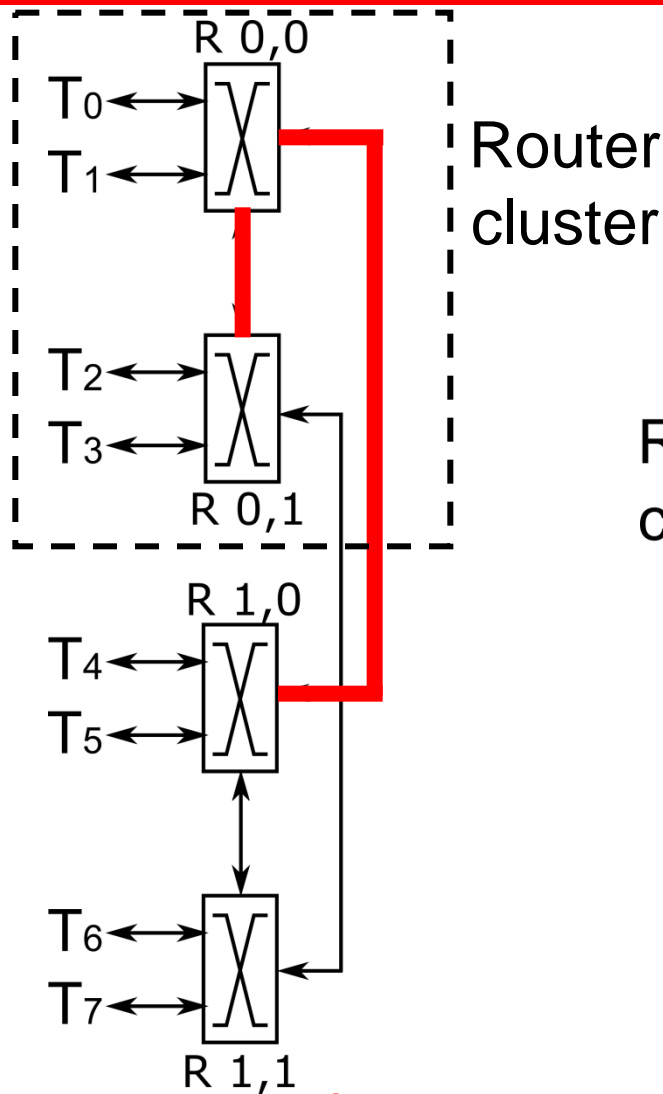


8 tiles per partition that are distributed across the chip (P8D)



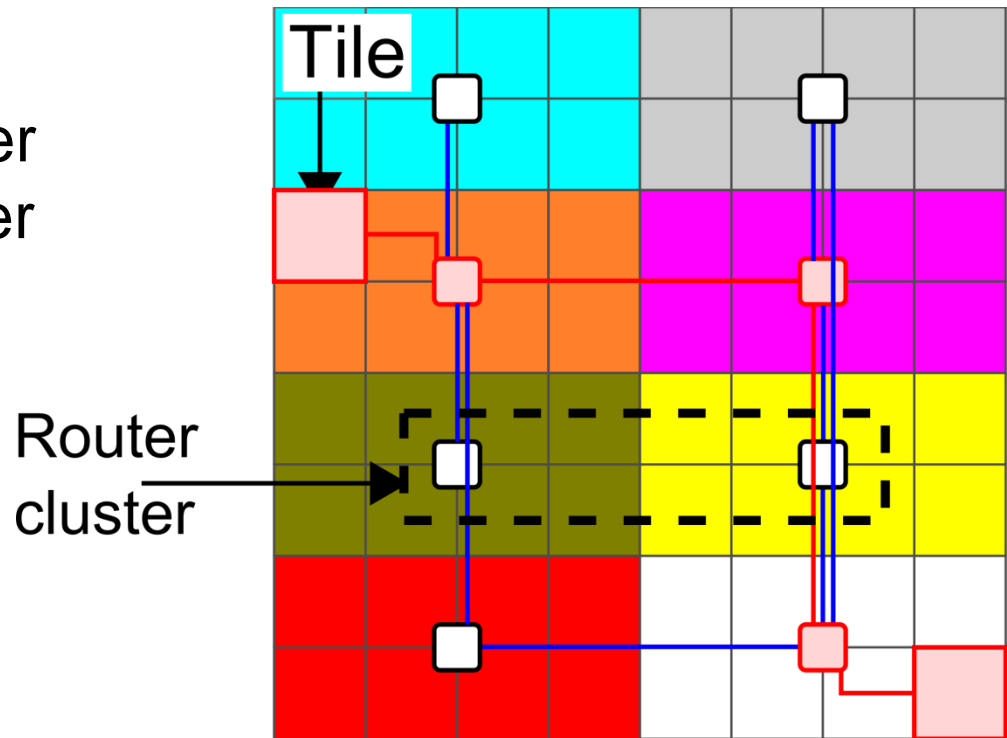
8 tiles per partition that are co-located (P8C)

# Physical layout for Flattened Butterfly



**2-ary 2-fly**

**flattened butterfly**

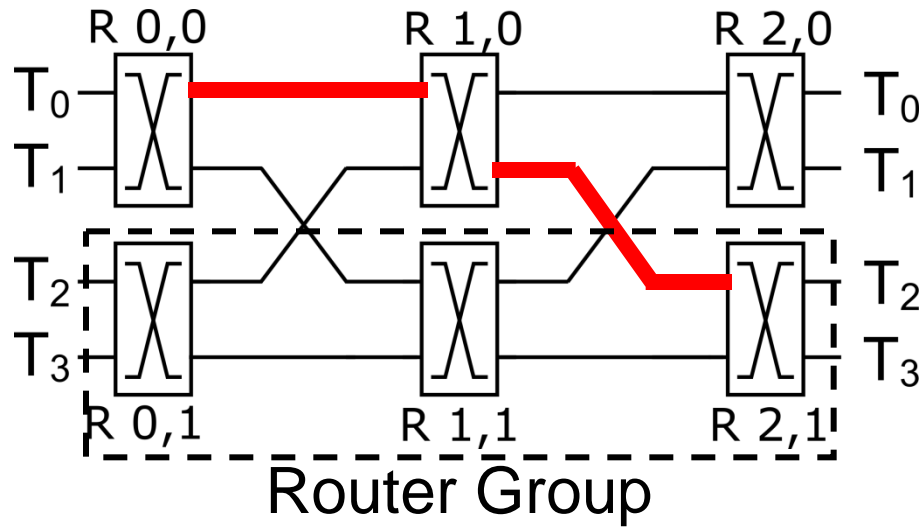


□ One 12 x 12 Router

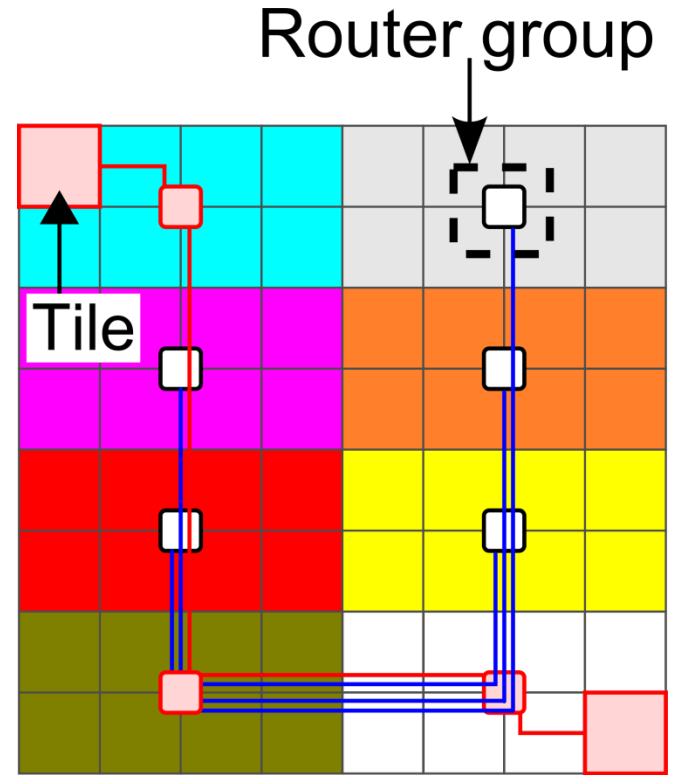
**8-ary 2-fly flattened butterfly**

- One router per group of 8 tiles
- Dimension-ordered routing → Intra-cluster followed by Inter-cluster

# Physical layout for Clos



**2-ary 3-stage clos**




Three 8 x 8 Routers


**8-ary 3-stage clos**

- 3 routers per router group
- Randomized oblivious routing

# Physical layouts for Clos



**2-ary 3-stage clos**




- Two 8 x 8 Routers
- Eight 8 x 8 Routers

**8-ary 3-stage clos**

- Intermediate routers located at the center
- Randomized oblivious routing

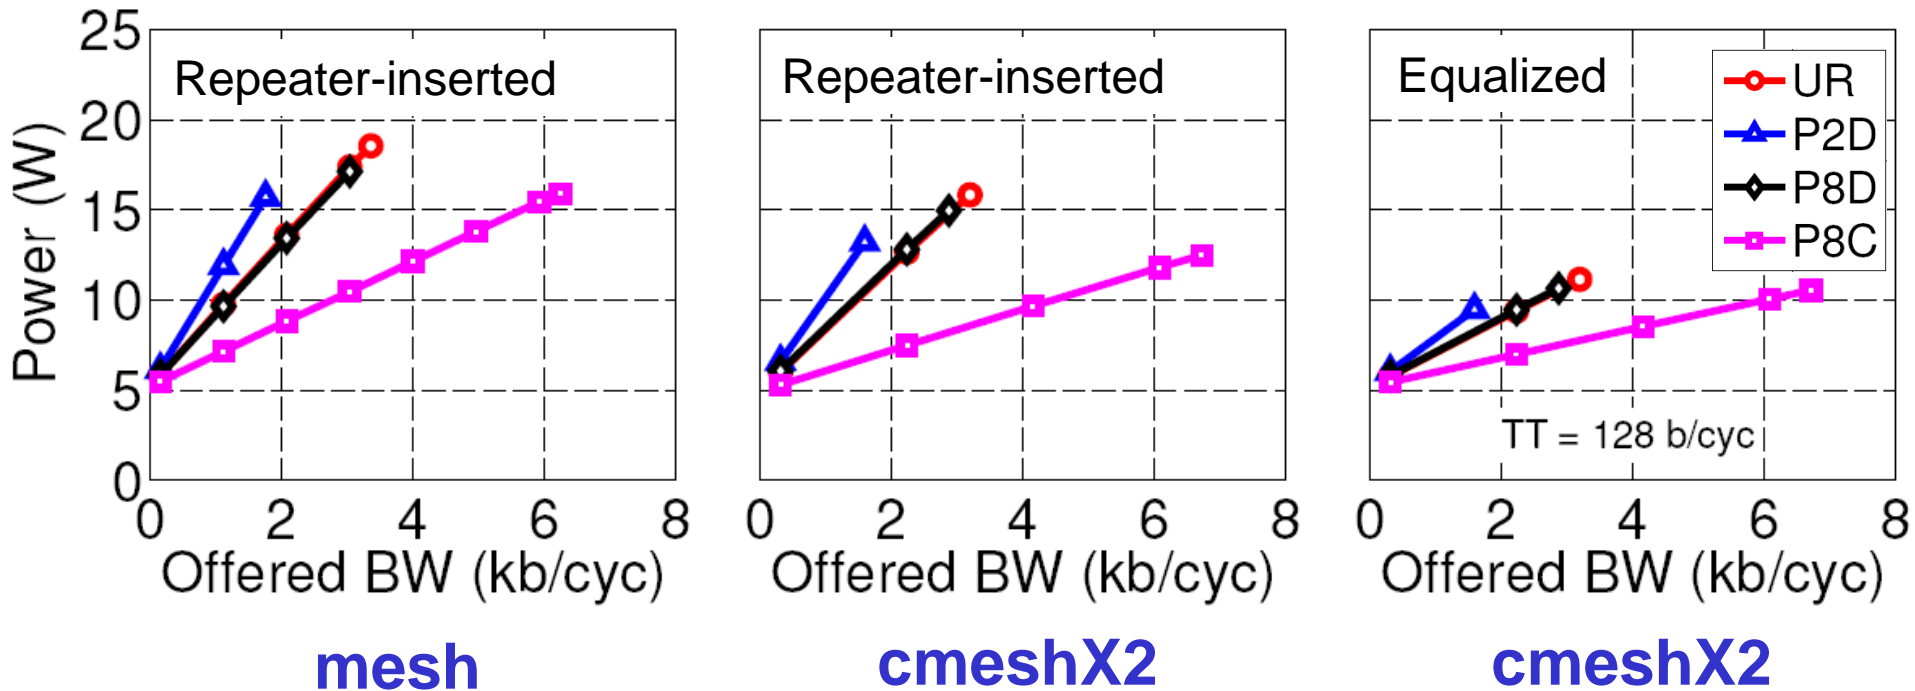
# Latency vs BW



**Ideal throughput = 8 kb/cyc for UR**

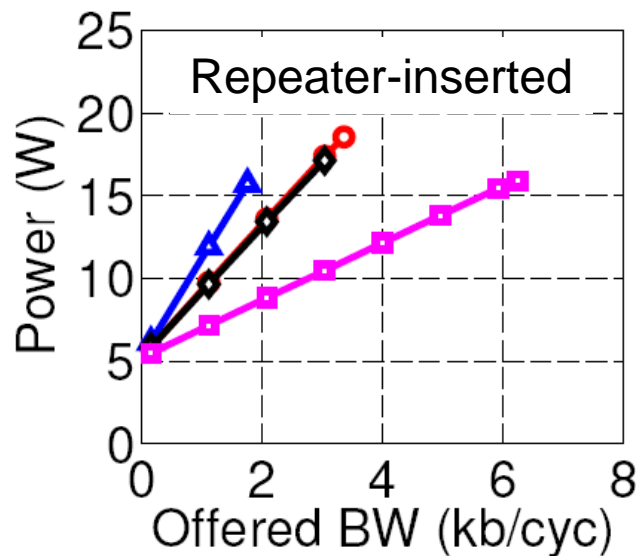
- ❑ flatFlyX2 vs mesh/cmeshX2
  - Saturation BW → comparable (UR, P8D and P2D), higher in flatFlyX2 for P8C
  - Latency → flatFlyX2 has lower latency
- ❑ clos vs mesh/cmeshX2/flatFlyX2
  - Saturation BW → uniform across all traffic patterns, comparable to UR of mesh
  - Latency → uniform across all traffic patterns, comparable to UR of mesh

# Mesh Vs CMeshX2



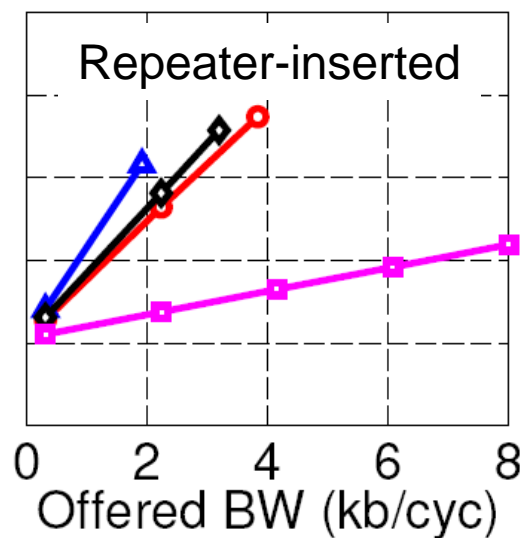
- Repeater-inserted interconnects
  - cmeshX2 consumes lower power than mesh at comparable throughput
- Equalized interconnects
  - cmeshX2 has further 1.5x reduction in power

# Power vs BW plots –repeater inserted pipelined vs equalized

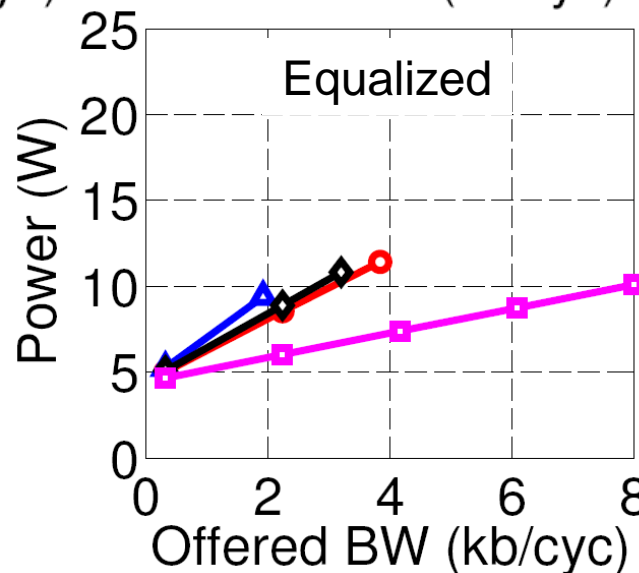



mesh


1.5-2x lower power obtained using equalized interconnects at comparable throughput



flatFlyX2

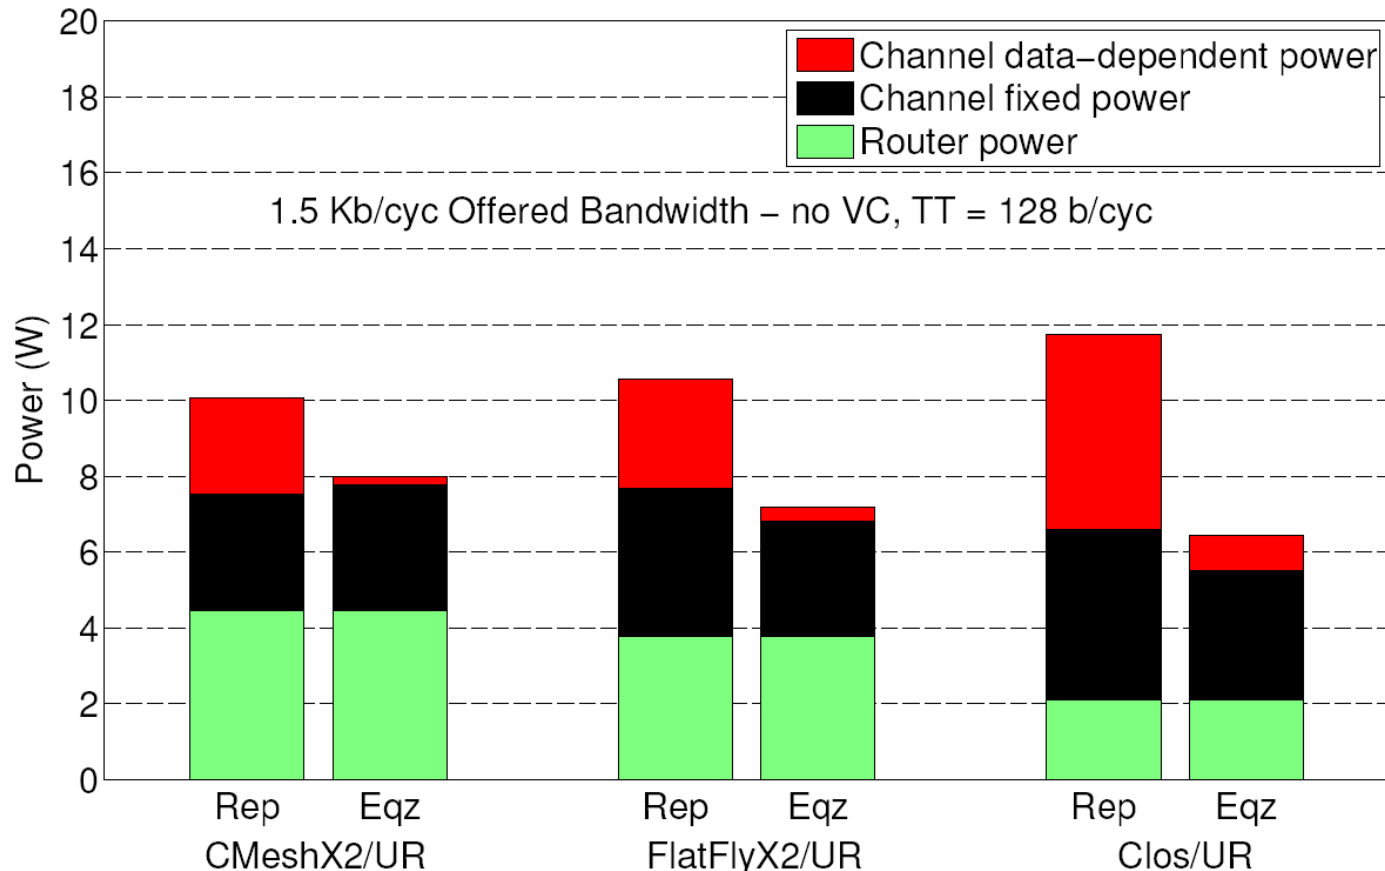


flatFlyX2



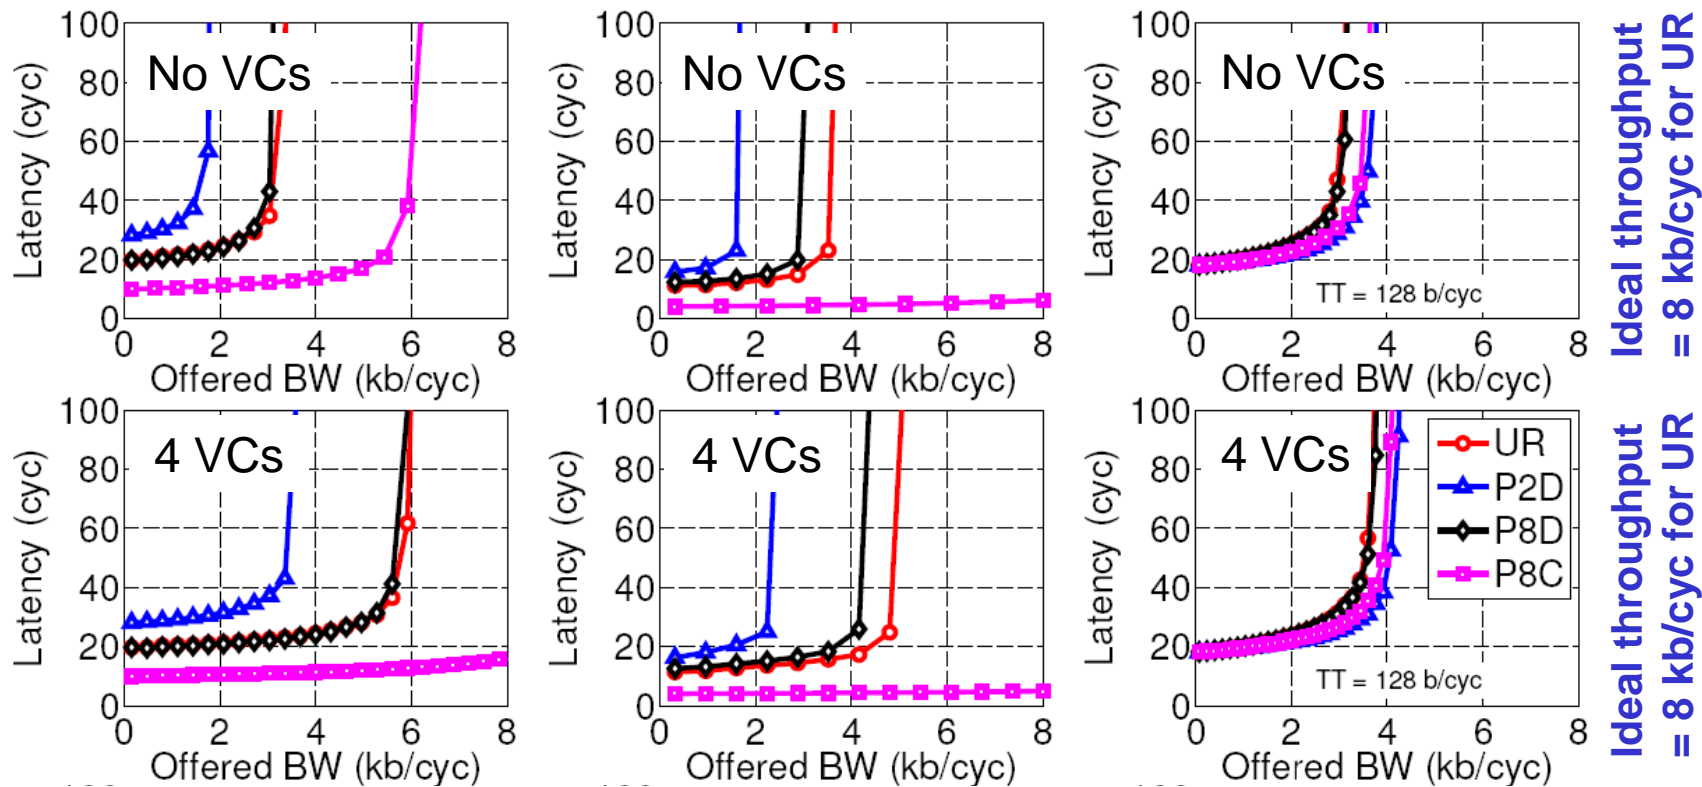
clos

# Power split



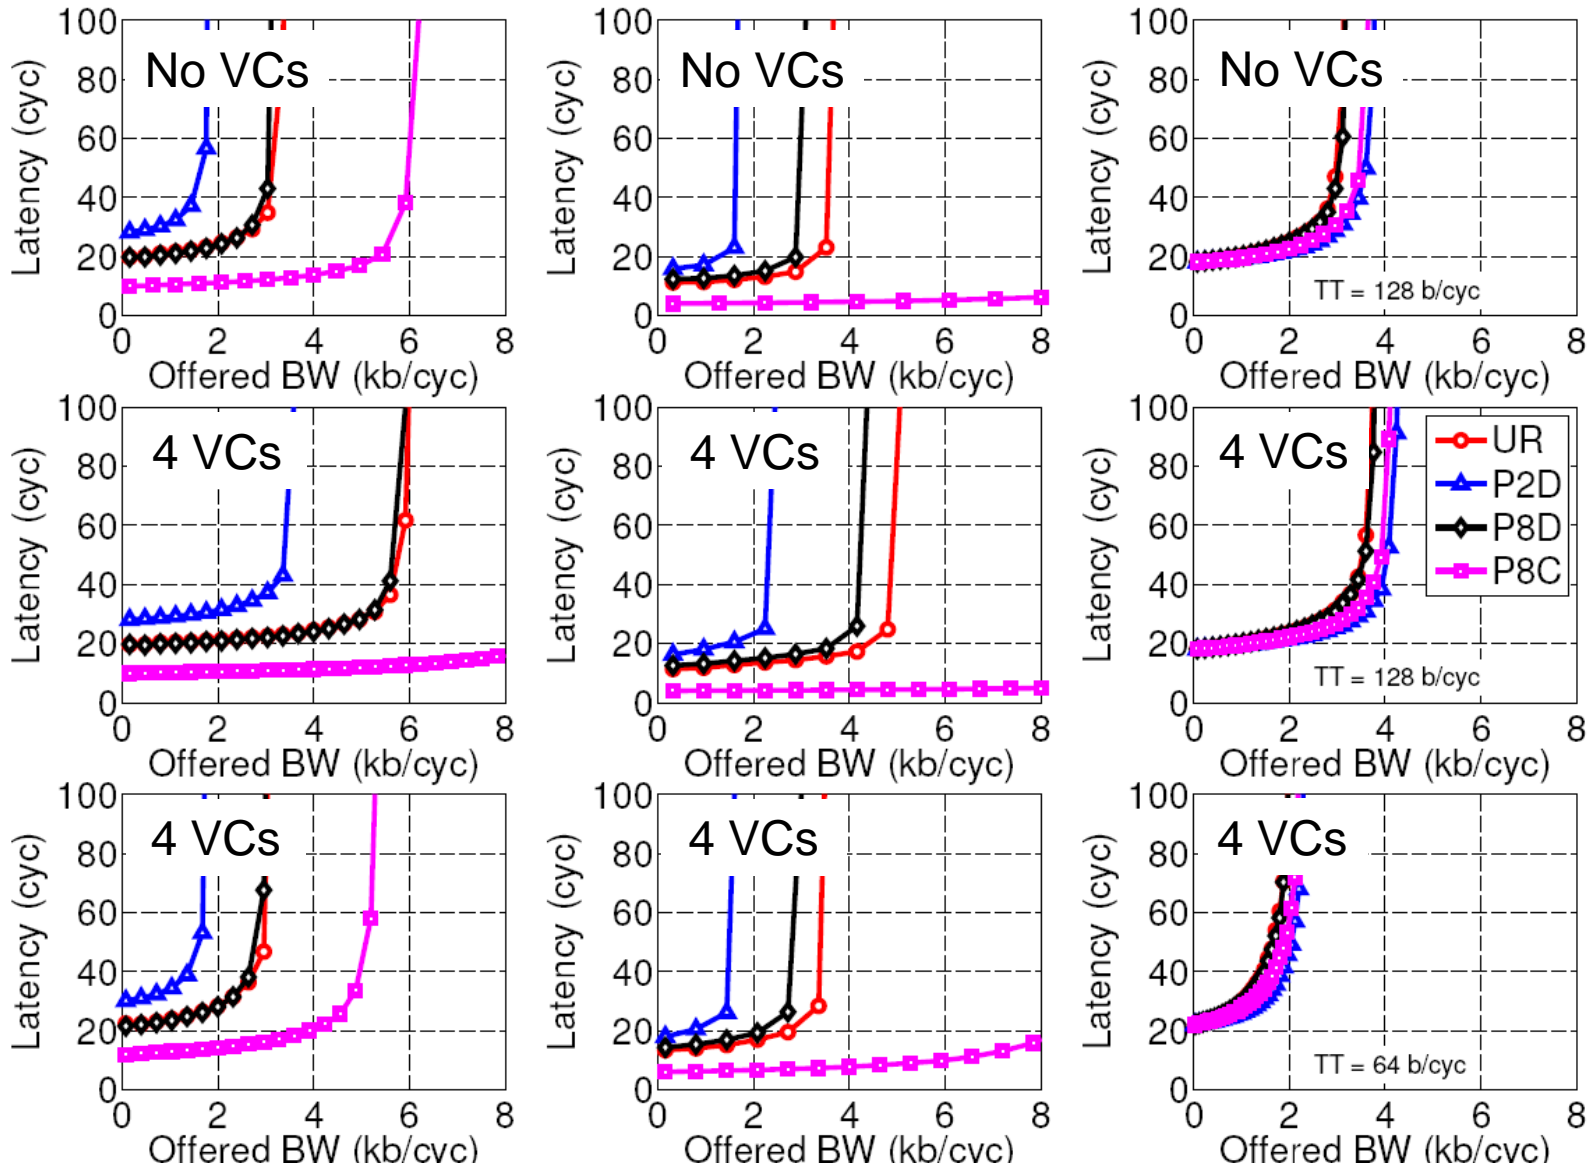
- ❑ Channel DDE reduces by 4-10x using equalized interconnects
- ❑ Channel fixed power and router power need to be tackled

# Latency vs BW – no VC vs 4 VCs



Saturation throughput improves using VC,  
but no change in power at comparable throughput

# Latency vs BW – no VC vs 4 VCs

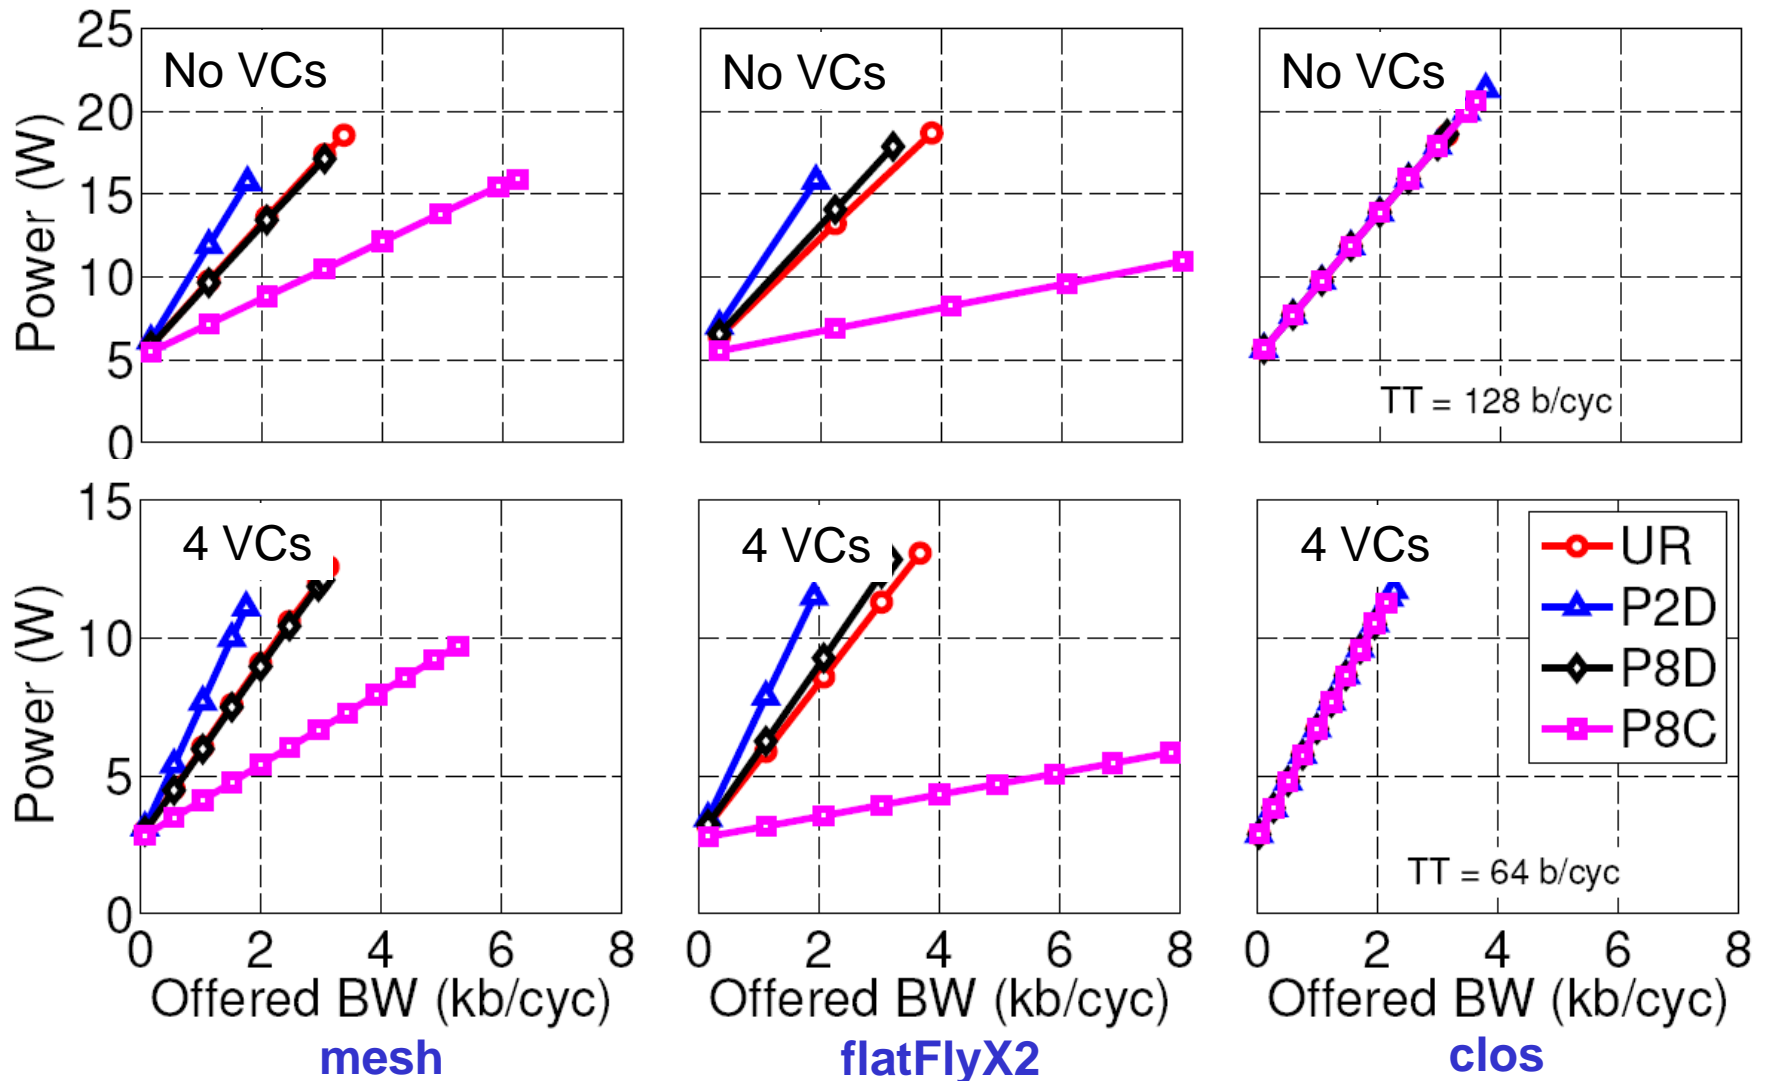


Ideal throughput = 8 kb/cyc for UR

Ideal throughput = 8 kb/cyc for UR

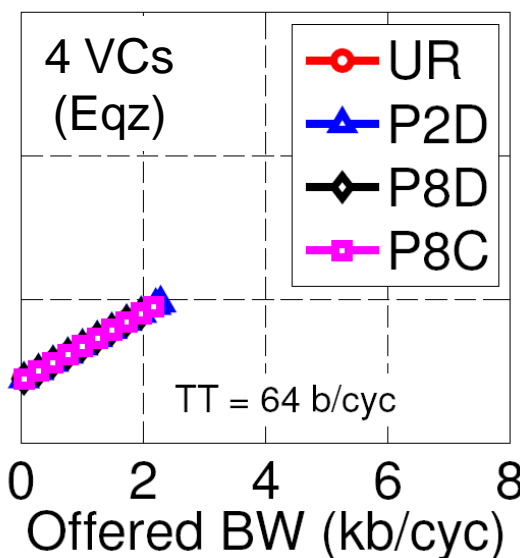
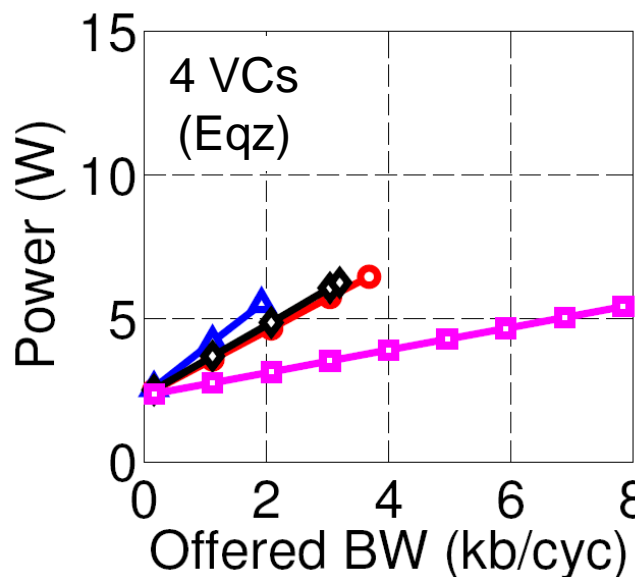



Ideal throughput = 4 kb/cyc for UR

# Power vs BW – no VC vs 4 VCs, repeater inserted pipelined



25-50% lower power using VCs at comparable throughput

# Power vs BW— no VC case, repeater inserted pipelined vs 4 VCs, equalized




2-3x lower power obtained using equalized interconnects and VCs at comparable throughput

mesh



flatFlyX2

clos

# Outline



Mesh



Crossbar

- ❑ Repeater-inserted interconnect vs Equalized interconnects
- ❑ Network design space exploration
- ❑ **Summary**

# Summary

---

- ❑ Repeater-inserted vs Equalized interconnects
  - Equalized interconnects have 4-10x lower data-dependent energy
  - Comparable fixed energy
- ❑ Equalized interconnects with no VCs gives 1.5-2x power savings at comparable throughput in low-diameter networks
- ❑ Using VCs provides further 25-50% reduction in power and 2x savings in area
- ❑ Router design needs to be improved