BREAKING MULTICAST DEADLOCK BY VIRTUAL CHANNEL ADDRESS/DATA FIFO DECOUPLING

Ka-Ming Keung, Akhilesh Tyagi

Iowa State University

On-Chip System with On-Chip Network

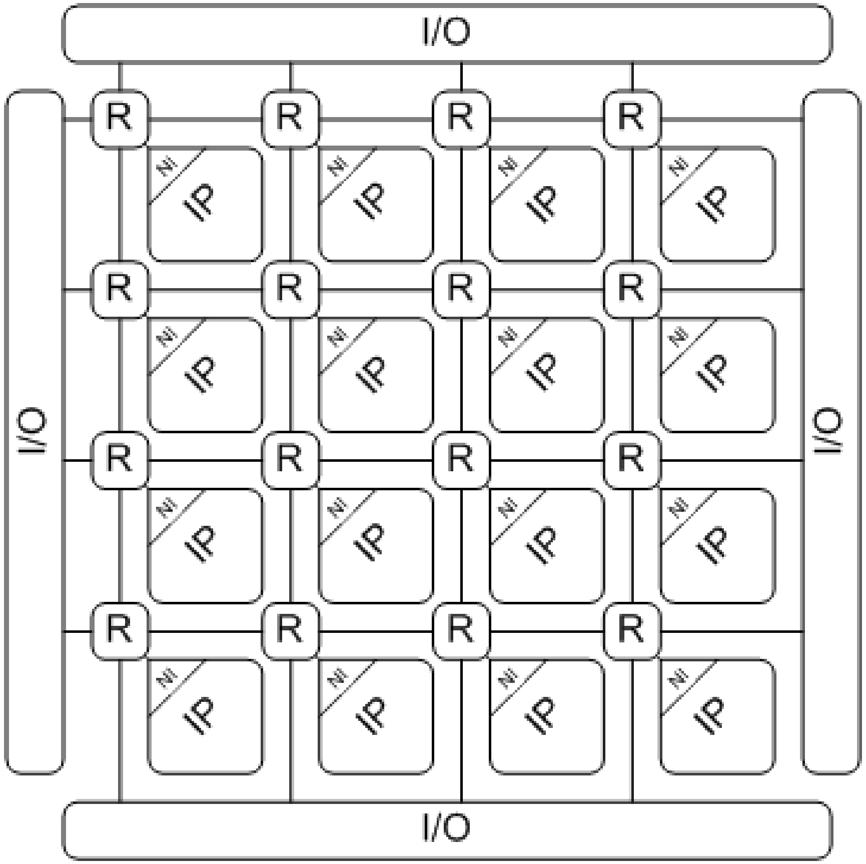

- Many tiles on a chip
- Communication among Tiles is supported by 2D Mesh Network

Adaptive Routing

- Allows packets being router through less congested channel.

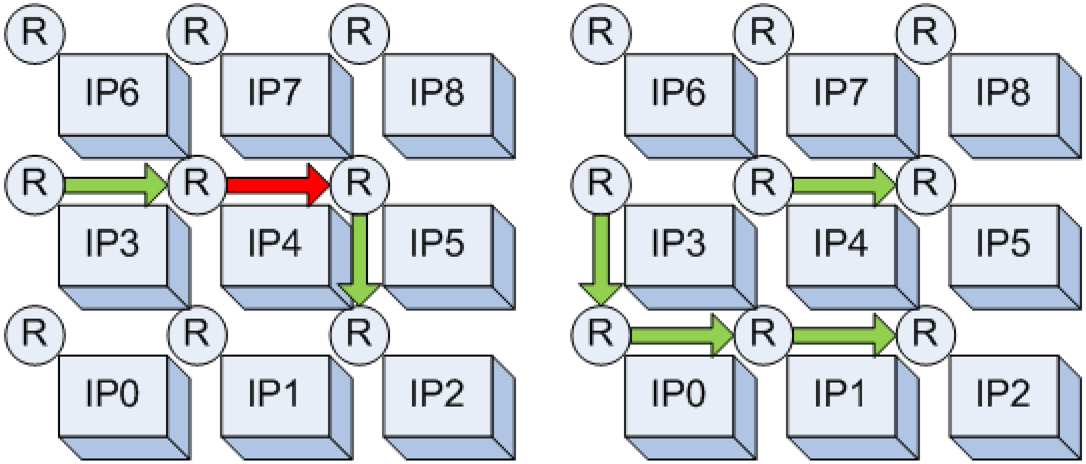

(a) XY Routing

(b) Adaptive Routing

Native Multicast Support

- Avoid redundant unicast packets
  - → Decrease Network Load
  - → Reduce Packet Latency

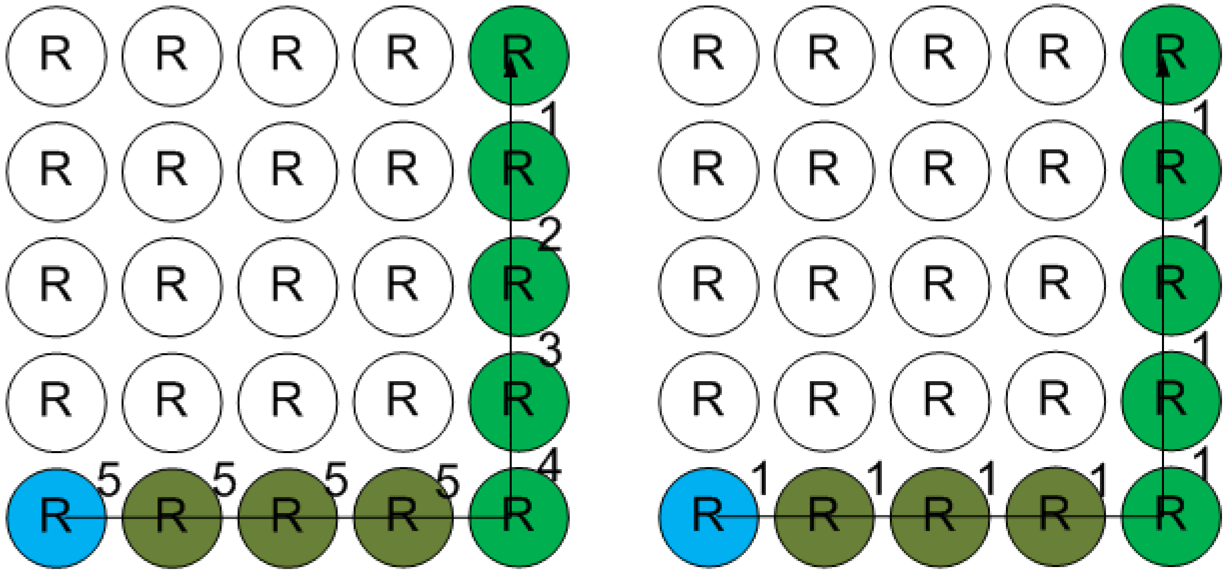

(a) XY Unicast (b) XY Multicast

BufferWrite: 30 V.S. 8

Adaptive Routing + Native Multicast Support

- Allow dynamic multicast packet divergent points

  → Decrease Network Load

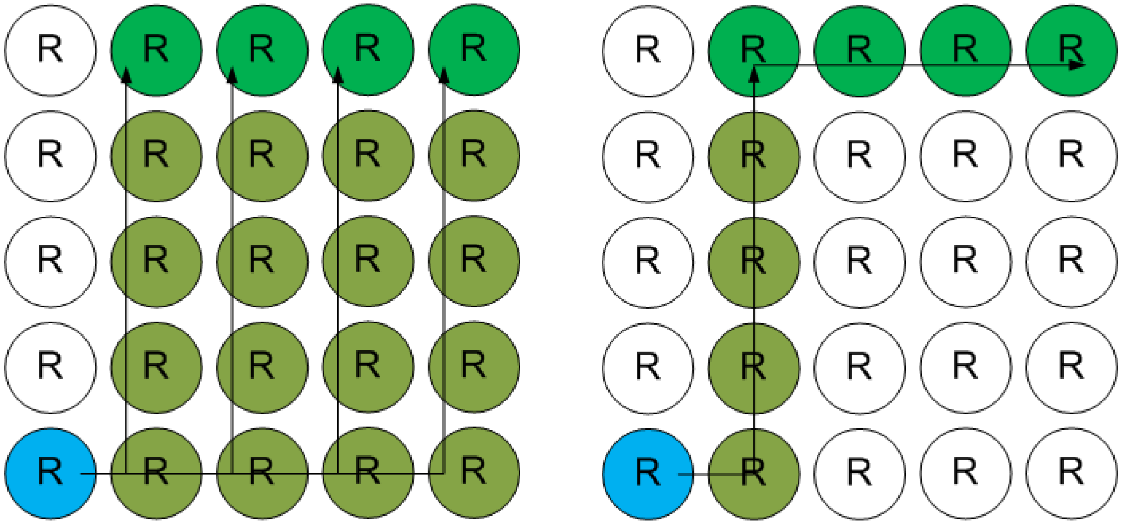

(a) XY Multicast Tree (b) Adaptive Multicast Tree

BufferWrite: 20 V.S. 8

Path Based Adaptive Routing

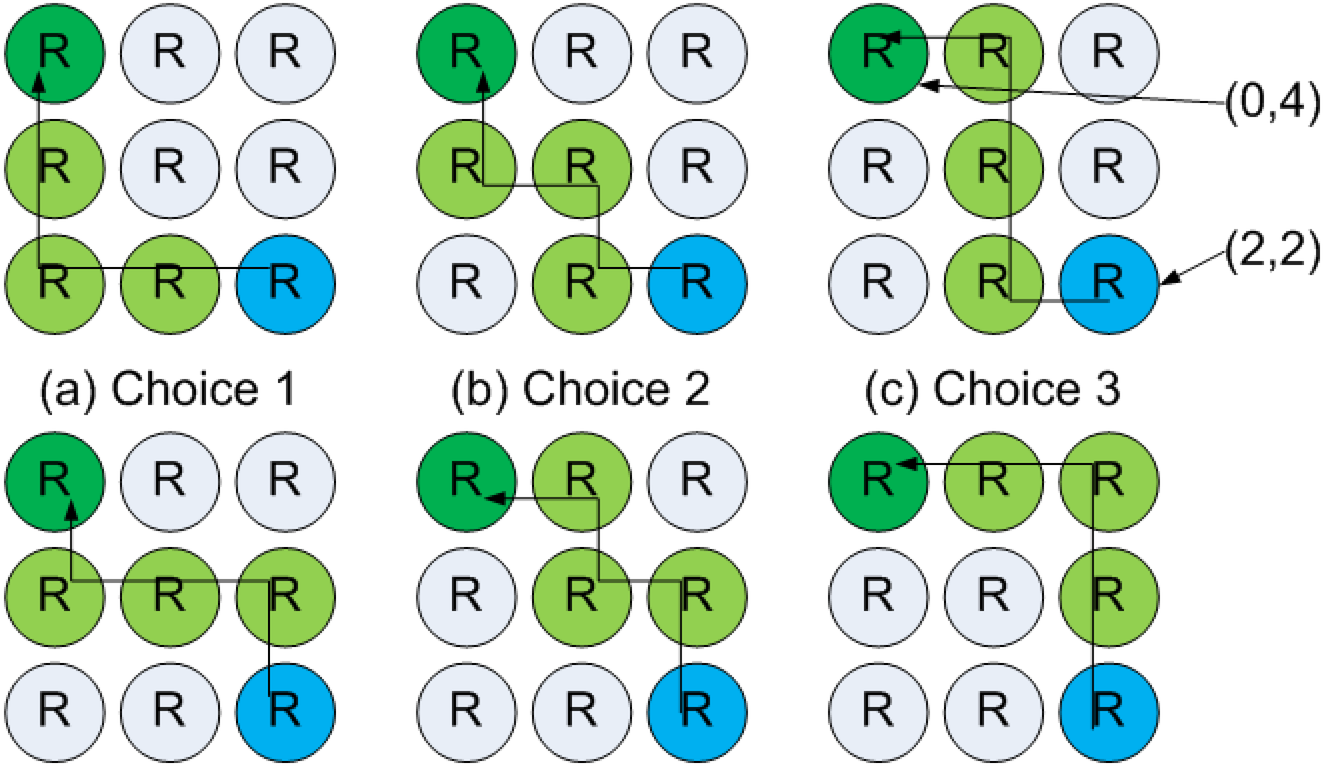

(a) Choice 1 (b) Choice 2 (c) Choice 3

(d) Choice 4 (e) Choice 5 (f) Choice 6

Valid Path

Choice	Route	Stop 0	Stop 1	Stop 2	Stop 3	OE Viol	Valid
1	WWNN	(2,2)	(1,2)	(0,2)	(0,3)	Free	Yes
2	WNWN	(2,2)	(1,2)	(1,3)	(0,3)	Even	No
3	WNNW	(2,2)	(1,2)	(1,3)	(1,4)	Even	No
4	NWWN	(2,2)	(2,3)	(1,3)	(0,3)	Odd	Yes
5	NWNW	(2,2)	(2,3)	(1,3)	(1,4)	Both	No
6	NNWW	(2,2)	(2,3)	(2,4)	(1,4)	Odd	Yes

Odd-Even Turn Model(Chiu et. al.) to ensure the network is deadlock free.
Only route 1,4,6 are valid.
Route 2,3,5 violate the odd-even routing rule

Path Based Adaptive Routing

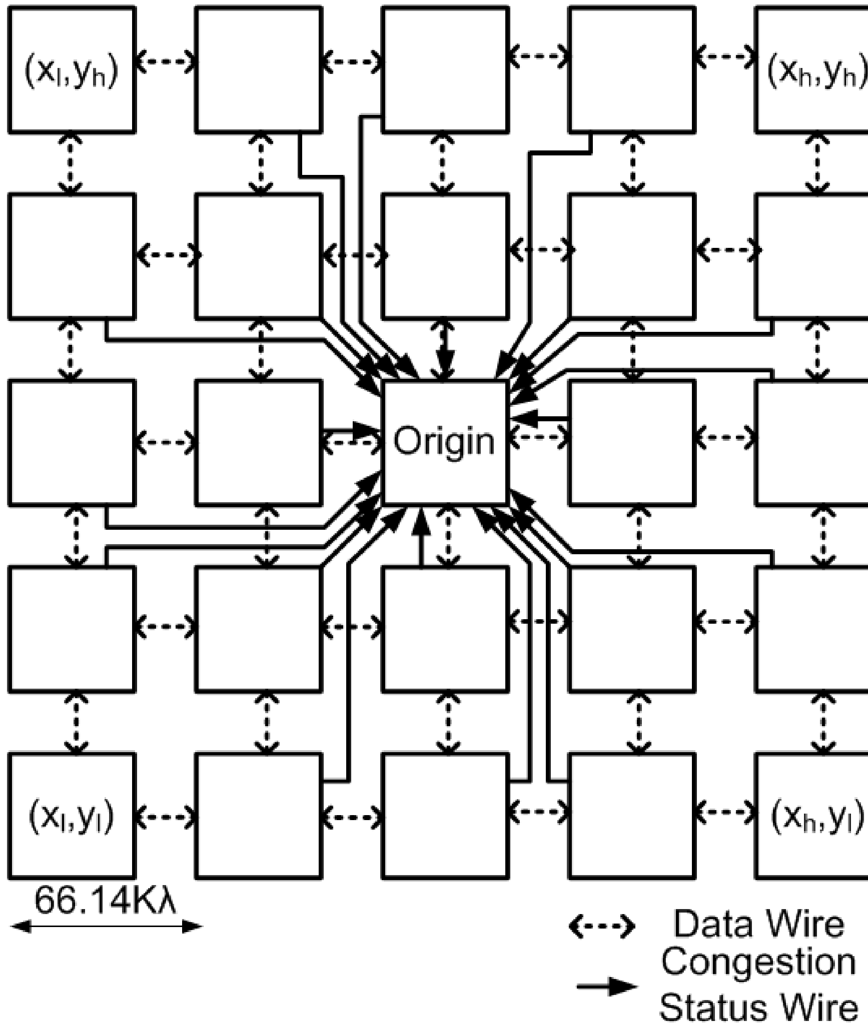

Path Selection

- Channel Congestion($CC_{x,y,j}$) is measured by the total Channel Demand($CD_{x,y,i,j}$) by all router input buffers:

$$CC_{x,y,j}=CD_{x,y,north,j}+CD_{x,y,east,j}+CD_{x,y,west,j}+CD_{x,y,south,j}+CD_{x,y,local,j}$$

- Path Congestion (PC_i) is the sum of the channel congestion along the path.

$$PC_i = \sum CC_{x,y,j}$$

- Pick the valid path i with the lowest PC_i

Observation Range

- Intuition: Bigger observation range leads to better network performance.
- Bigger observation range requires
  - More congestion status wires from the remote router
  - Longer cost computation path

    → Potentially affects router clock frequency
  - More adders and comparators

    → Higher Area Cost

Observation Range

Uniform Traffic Test:

- Low-Load Latency
  - Stay the same
- Throughput
  - 5x5 is 29% higher than 3x3
  - 7x7 is 6% higher than 5x5
  - 9x9 is 5.5% higher than 7x7
- RC Path
  - 5x5 is 219ps longer than 3x3
  - 7x7 is 439ps longer than 5x5
  - 9x9 is 453ps longer than 7x7

- We pick 5x5 to avoid RC stage becomes the critical stage

Range	Throughput(flits)		Latency(cycles)	
	vc9	vc14	vc9	vc14
3×3	76444	103828	105.43	102.89
5×5	97541	133848	104.99	102.84
7×7	104081	142282	104.99	102.87
9×9	106842	150109	104.89	102.83
XY	95529	118201	103.84	101.69

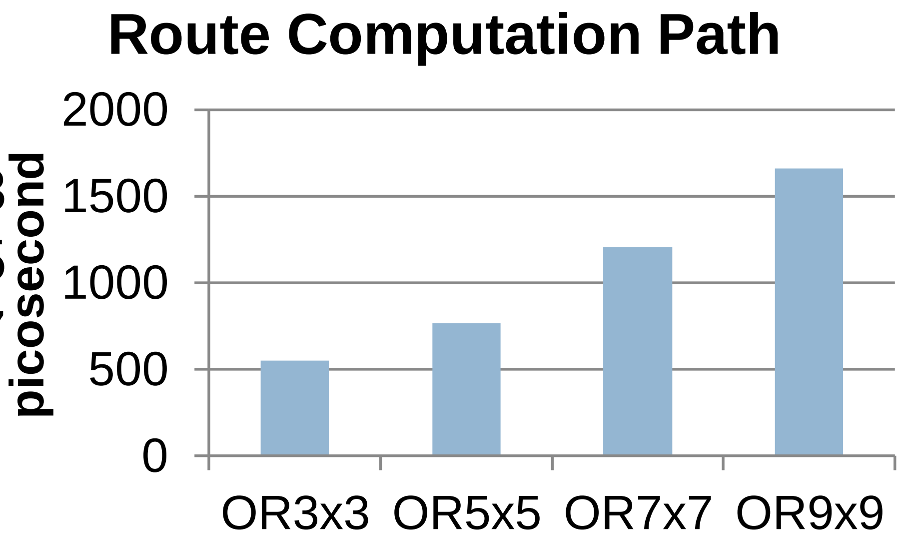

Virtual Destinations

- Not all destination lies within the observation range
- For those destinations, we assume they lie on the observation range boundary

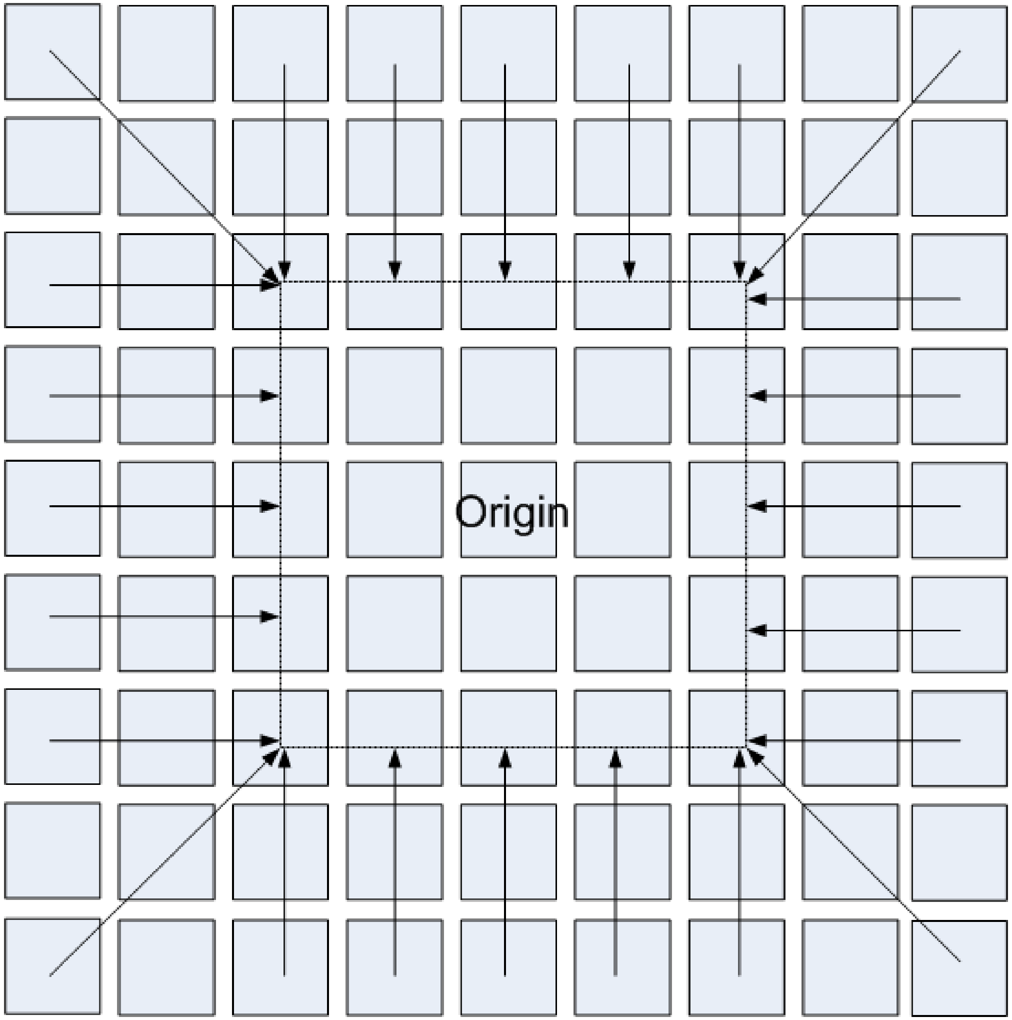

Multicast Adaptive Routing

- Objective:
  - Reduce the number of buffer write by diverging the packet as late as possible

Multicast Adaptive Routing

Rule 1(XY Destinations):

- If the packet has directions in North, East, West and South, packet will be routed to the corresponding direction.

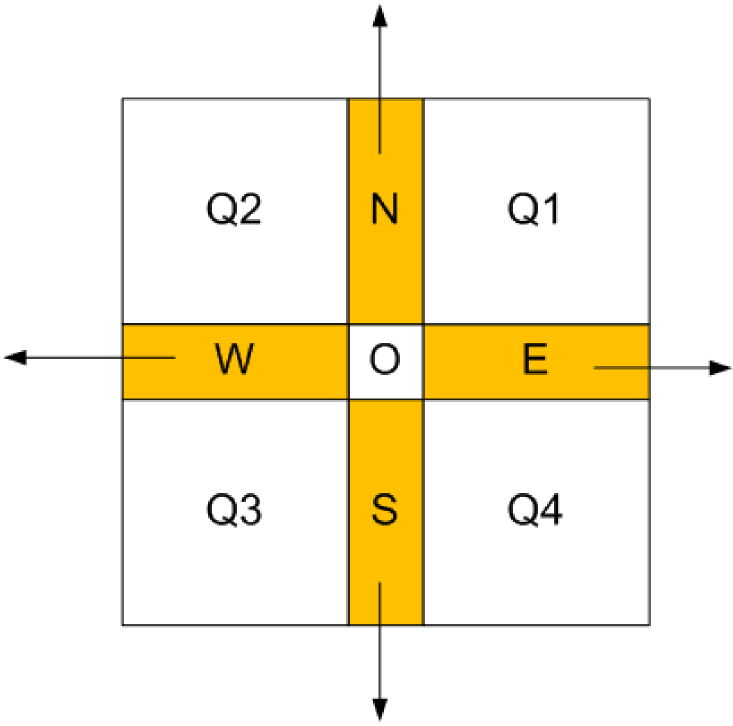

Multicast Adaptive Routing

Rule 2 (Quadrant Destinations) :

- In minimal routing, destinations at the quadrants can be routed horizontally(D_h) or vertically(D_v). If the packet has destination on either D_h or D_v, quadrant destinations will be routed to that direction.

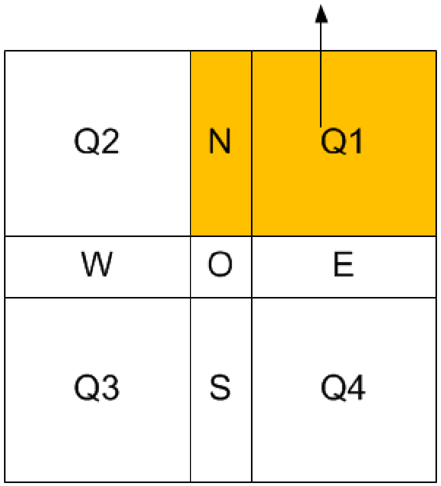

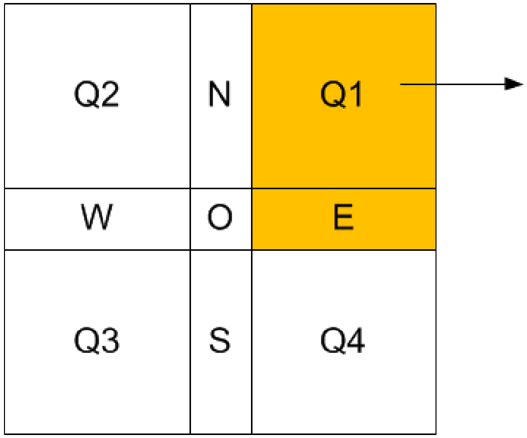

Multicast Adaptive Routing

Rule 3 (Quadrant Destinations):

- Group the destinations which can't be routed by Rule 2 to a single routing direction.

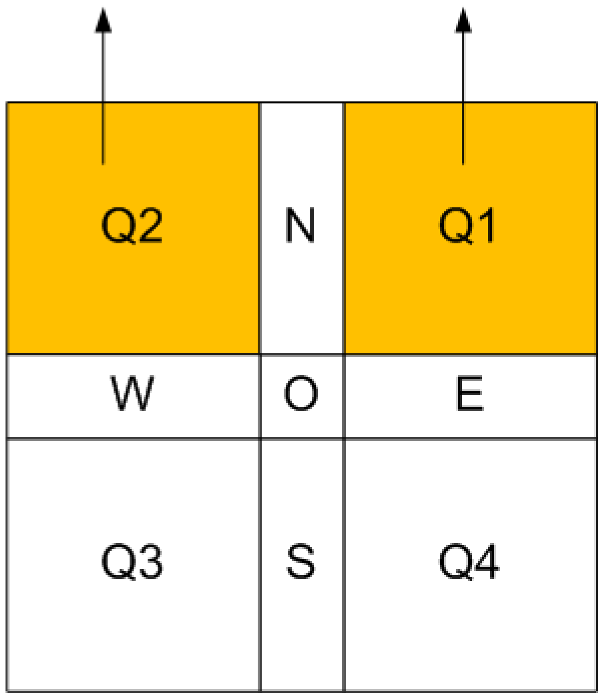

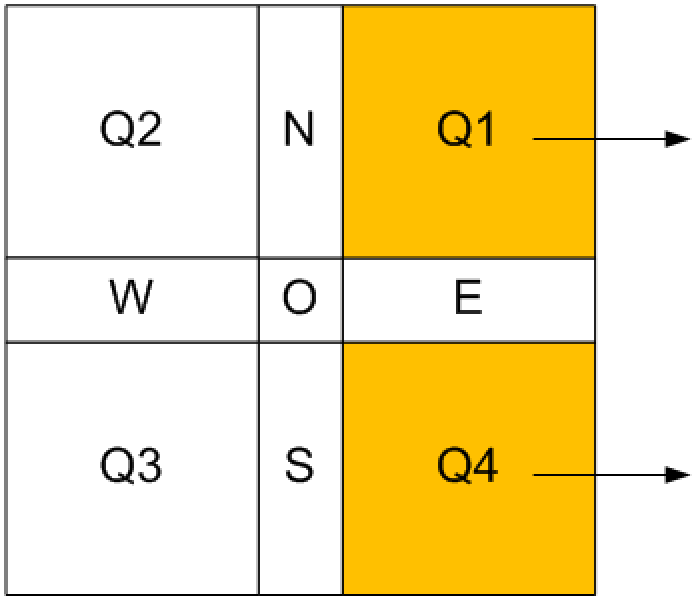

Multicast Adaptive Routing

Rule 4 (Quadrant Destinations):

- Destinations which can't be routed by Rule 3 are routed using unicast adaptive routing to the virtual destination at the corner of the observation range.

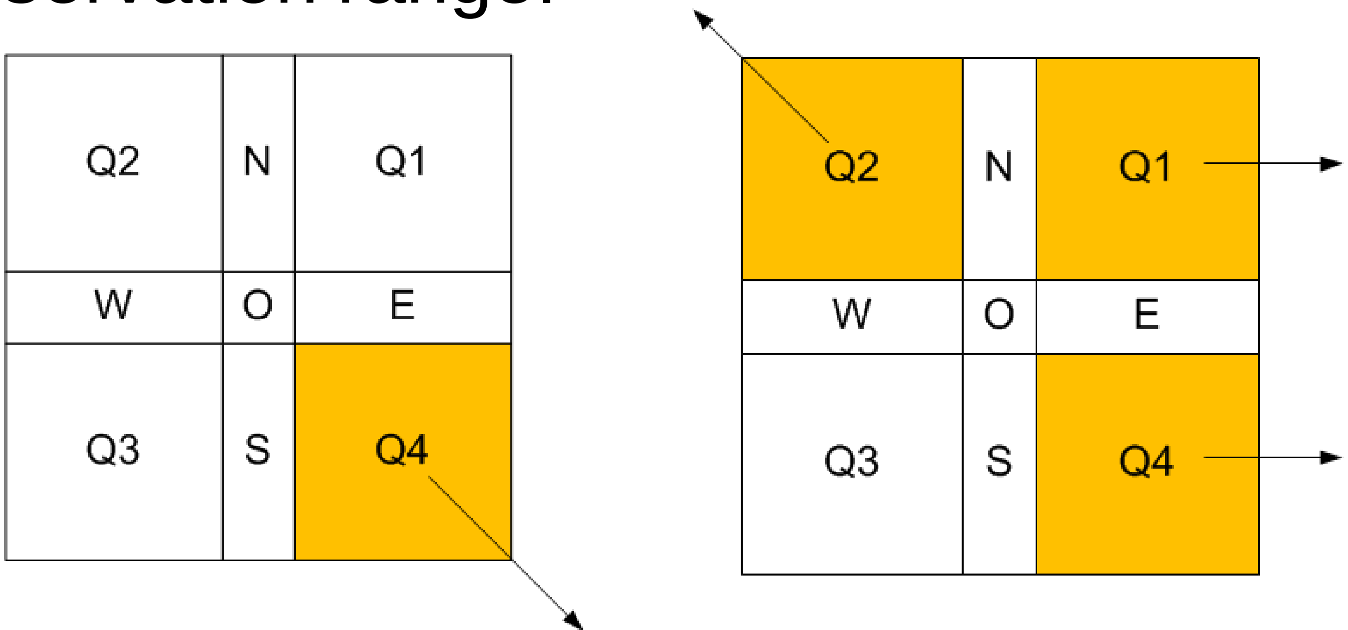

Unicast Deadlock

- Lock because of channel dependence
- XY-Routing is free from Unicast Deadlock
- Previous Solutions:

1. Ordered nodes and virtual channel (Dally et al.)
2. West-First, North-Last and Negative-first (Glass et al.)
3. Odd-Even Routing (Chiu)

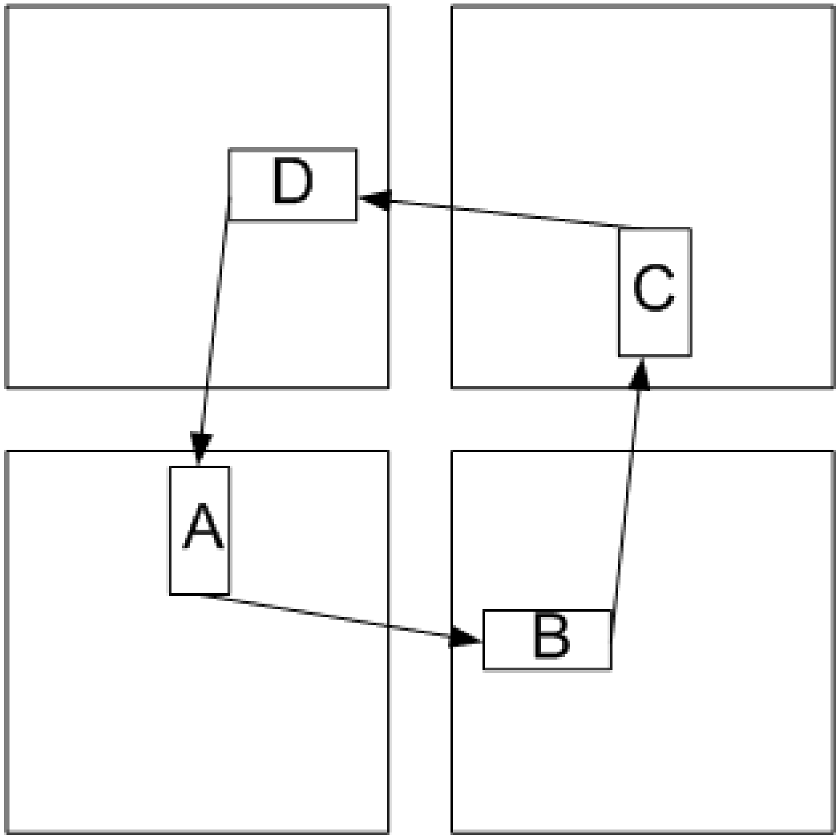

Multicast Deadlock

- Lock because of channel dependence
- Even XY-Routing could suffer from Multicast Deadlock
- Example:
  - Tile(1,1) sends multicast packet 55 to Tile (0,1), (3,1)
  - Tile(2,1) sends multicast packet 77 to Tile(0,1), (3,1)
  - Packet 55 does not release (0,1) E until it gets (3,1) W
  - Packet 77 does not release (3,1) W until it gets (0,1) E

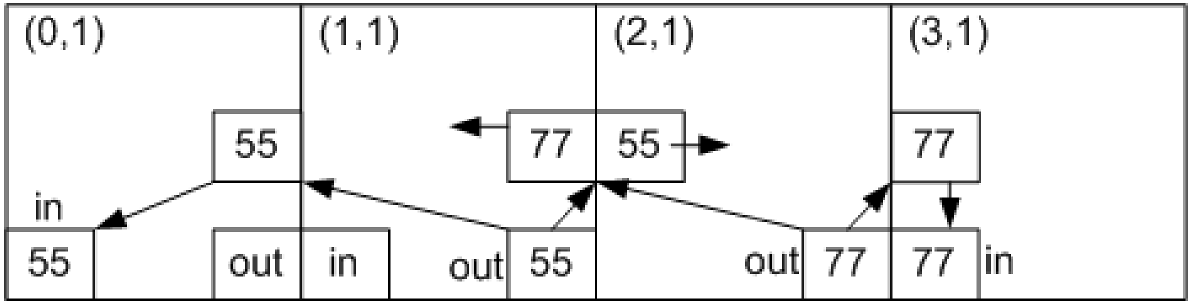

Multicast Deadlock

Previous Solution 1:

- Send four packets to regions (X+,Y+),(X+,Y-), (X-,Y+) and (X-,Y-) separately. (Lin et al.)

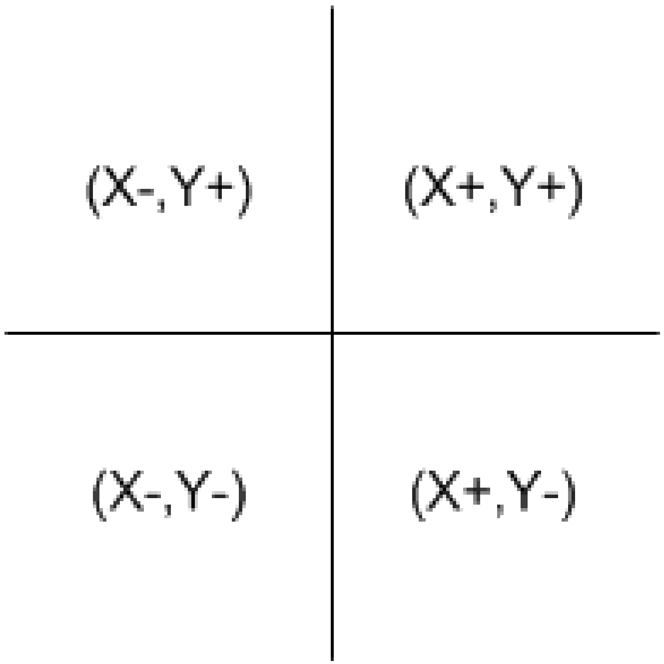

Multicast Deadlock

Previous Solution 2:

- Hamiltonian Path

Pre-compute deadlock free path and store it in the packet header.

Routers route the packet following the stored path (Lin et al.)

Multicast Deadlock

Previous Solution 3:

- Planar Network(Chien et al.)

Use two subnet networks X+ and X-.

X+ sub-network for packet with increasing X co-ordinate.

X- sub-network for packet with non-increasing X co-ordinate.

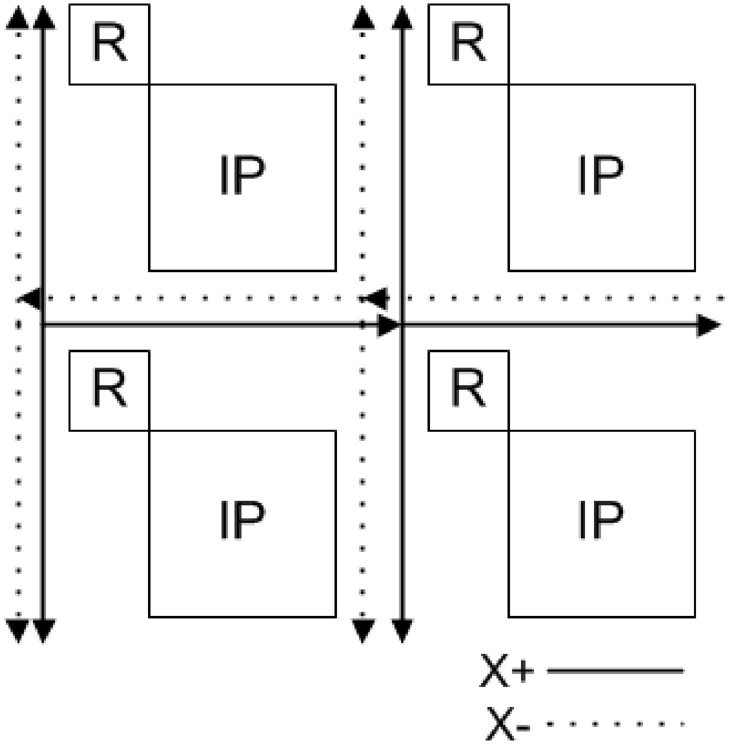

Multicast Deadlock

Simple Solution:

- Use Virtual Cut-through routing instead of wormhole routing.
- Router (0,1) East and (2,1) West can store the whole packet 55

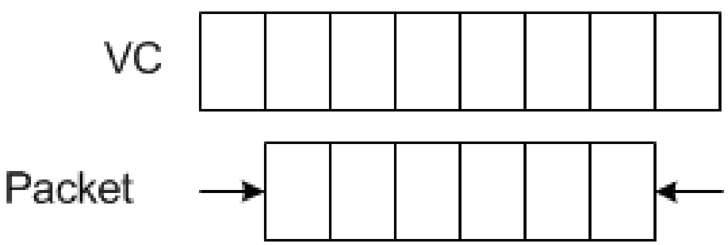

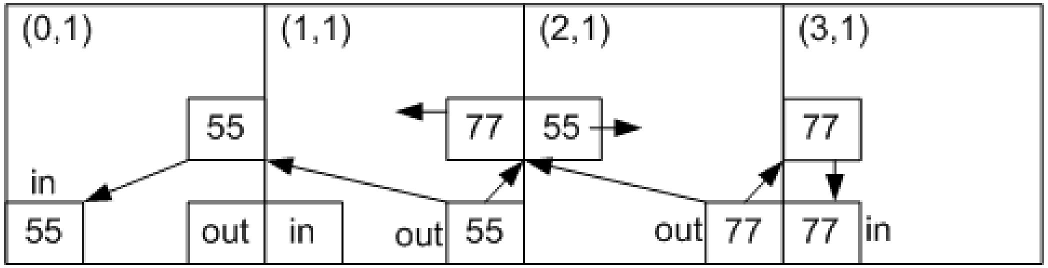

Multicast Deadlock

- (0,1) East channel and (3,1) West channel are empty when the deadlock occurs.
- (1,1) out has no new flit for (0,1) East (Packet 55)
- (2,1) out has no new flit for (3,1) West (Packet 77)
- Deadlock is broken if packet 55 releases (0,1) East Channel and packet 77 releases (3,1) West channel

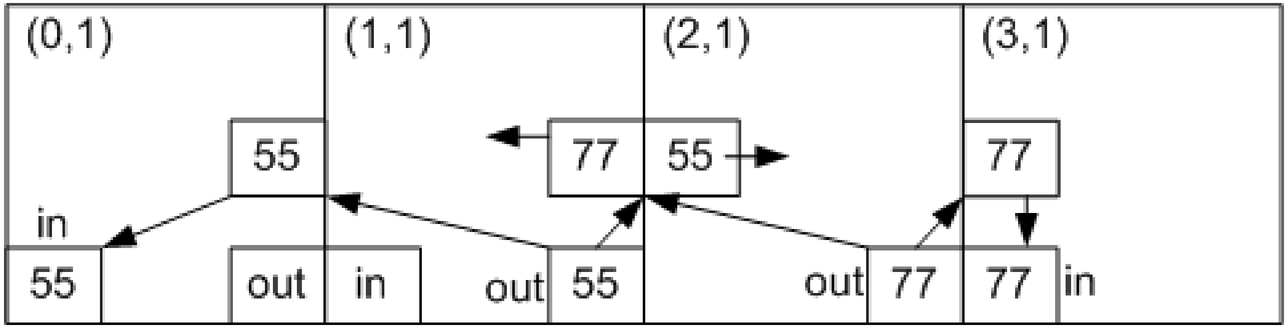

Address-Data FIFO Decoupling

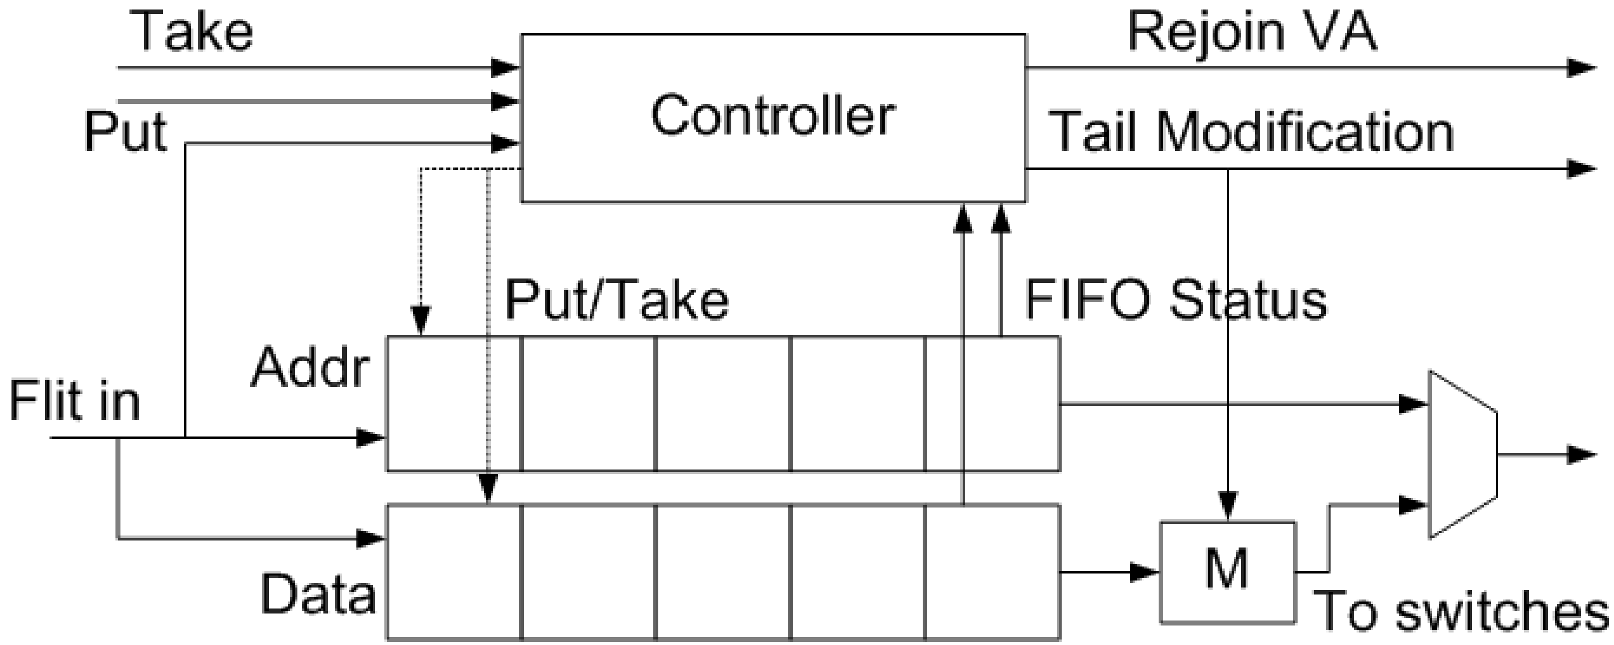

Example: Each Virtual Channel can store 2 addr flits + 2 data flits

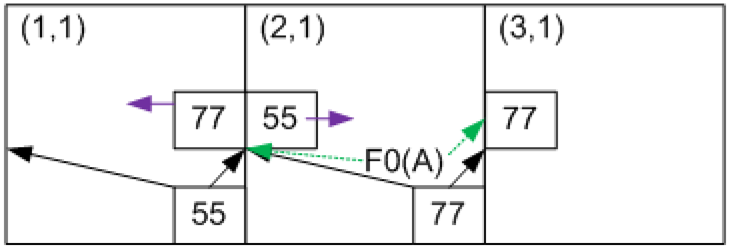

Step 1

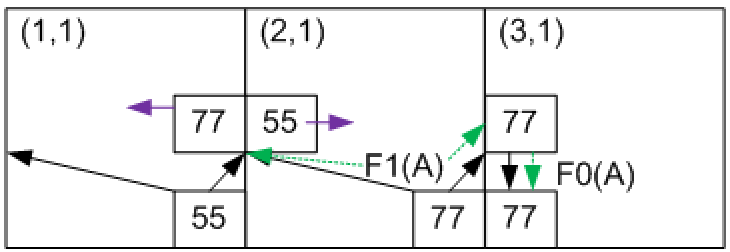

Step 2

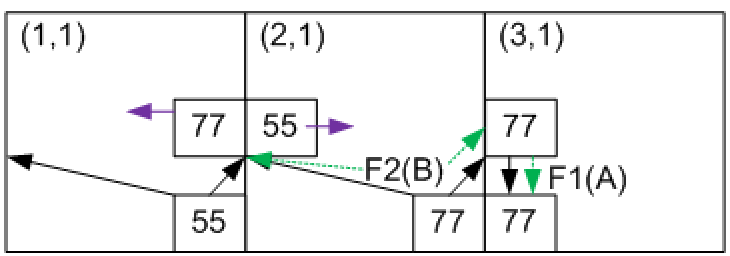

Step 3

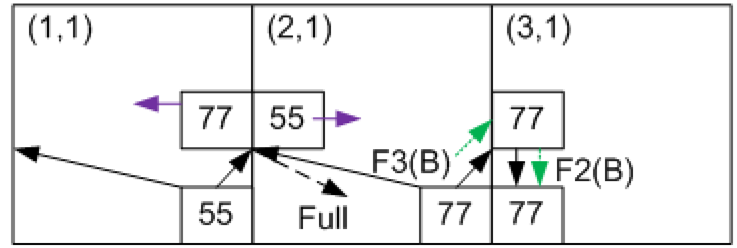

Step 4: (1,1) East is Full

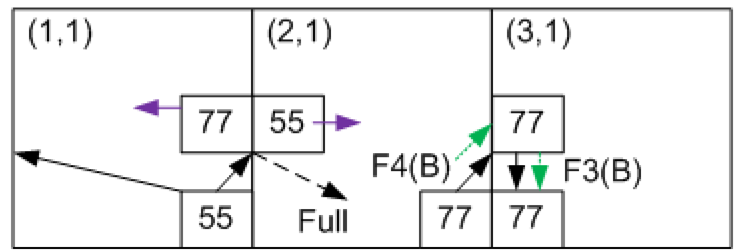

Step 5

Packet 77

F0(A)	F1(A)	F2(B)	F3(B)	F4(B)	F5(B)	F6(B)	F7(T)

A:Address B:Body T:Tail

Example: Virtual Channel can store 2 addr flits + 2 data flits

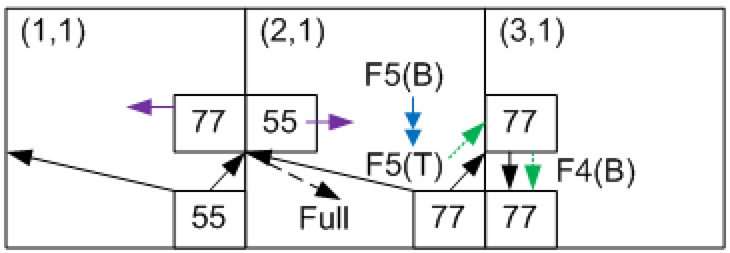

Step 6: (2,1)out does not have F6
→ Body to Tail Modification

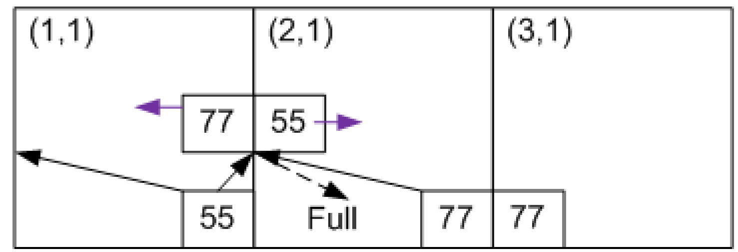

Step 8: (3,1) West is Free

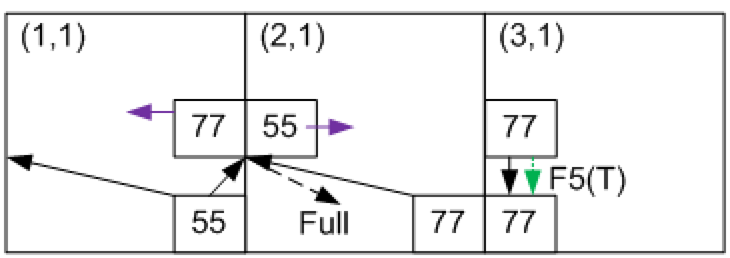

Step 7 Packet 77 Releases (3,1) West

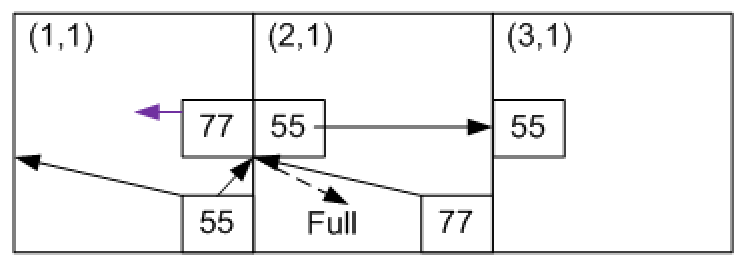

Step 9: Packet 55 obtains (3,1) West

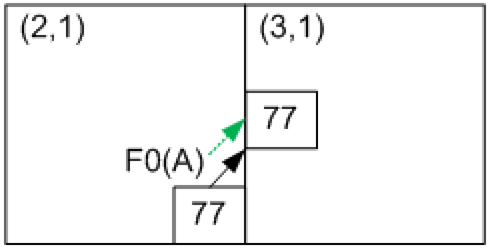

Step 10: Resend Address Flit F0

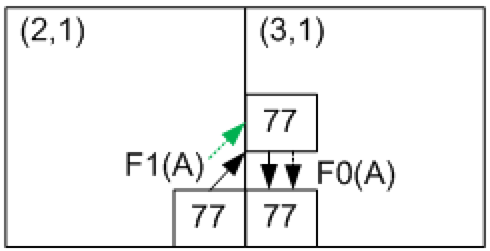

Step 11: Resend Address Flit F1

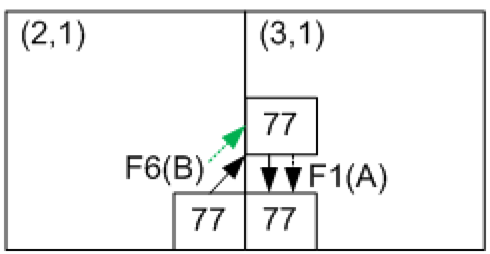

Step 12: Send Data Flit F6

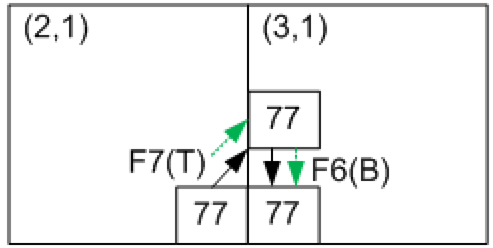

Step 13

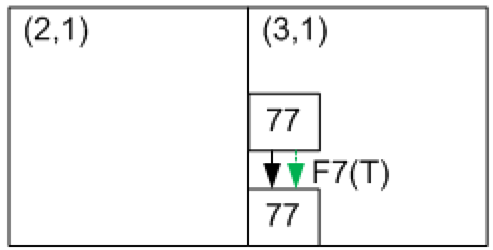

Step 14

Packet 77 Received

F0(A)	F1(A)	F2(B)	F3(B)	F4(B)	F5(T)
F0(A)	F1(A)	F6(B)	F7(T)		

A:Address B:Body T:Tail

Synthetic Traffic

- Four Types of synthetic traffics:
  - Uniform Traffic
  - Transpose Traffic (x,y) → (N-1-y,N-1-x)
  - Transpose2 Traffic (x,y) → (y,x)
  - Tornado Traffic
- Multicast Group Size: 10
- Multicast Probability: 5%

Experimental Setup

- Mesh Size: 20x20
- Flit Size: 128-bit
- Simulation Cycle: 30000
- Packet Length: 10 flits
- #Virtual Channel:
  - 3 unicast channels (unicast router)
  - 2 unicast + 1 multicast channel (multicast router)
- Virtual Channel Depth:
  - 14 (Virtual Cut-Through)
  - 9 (Address-Data FIFO decoupling)

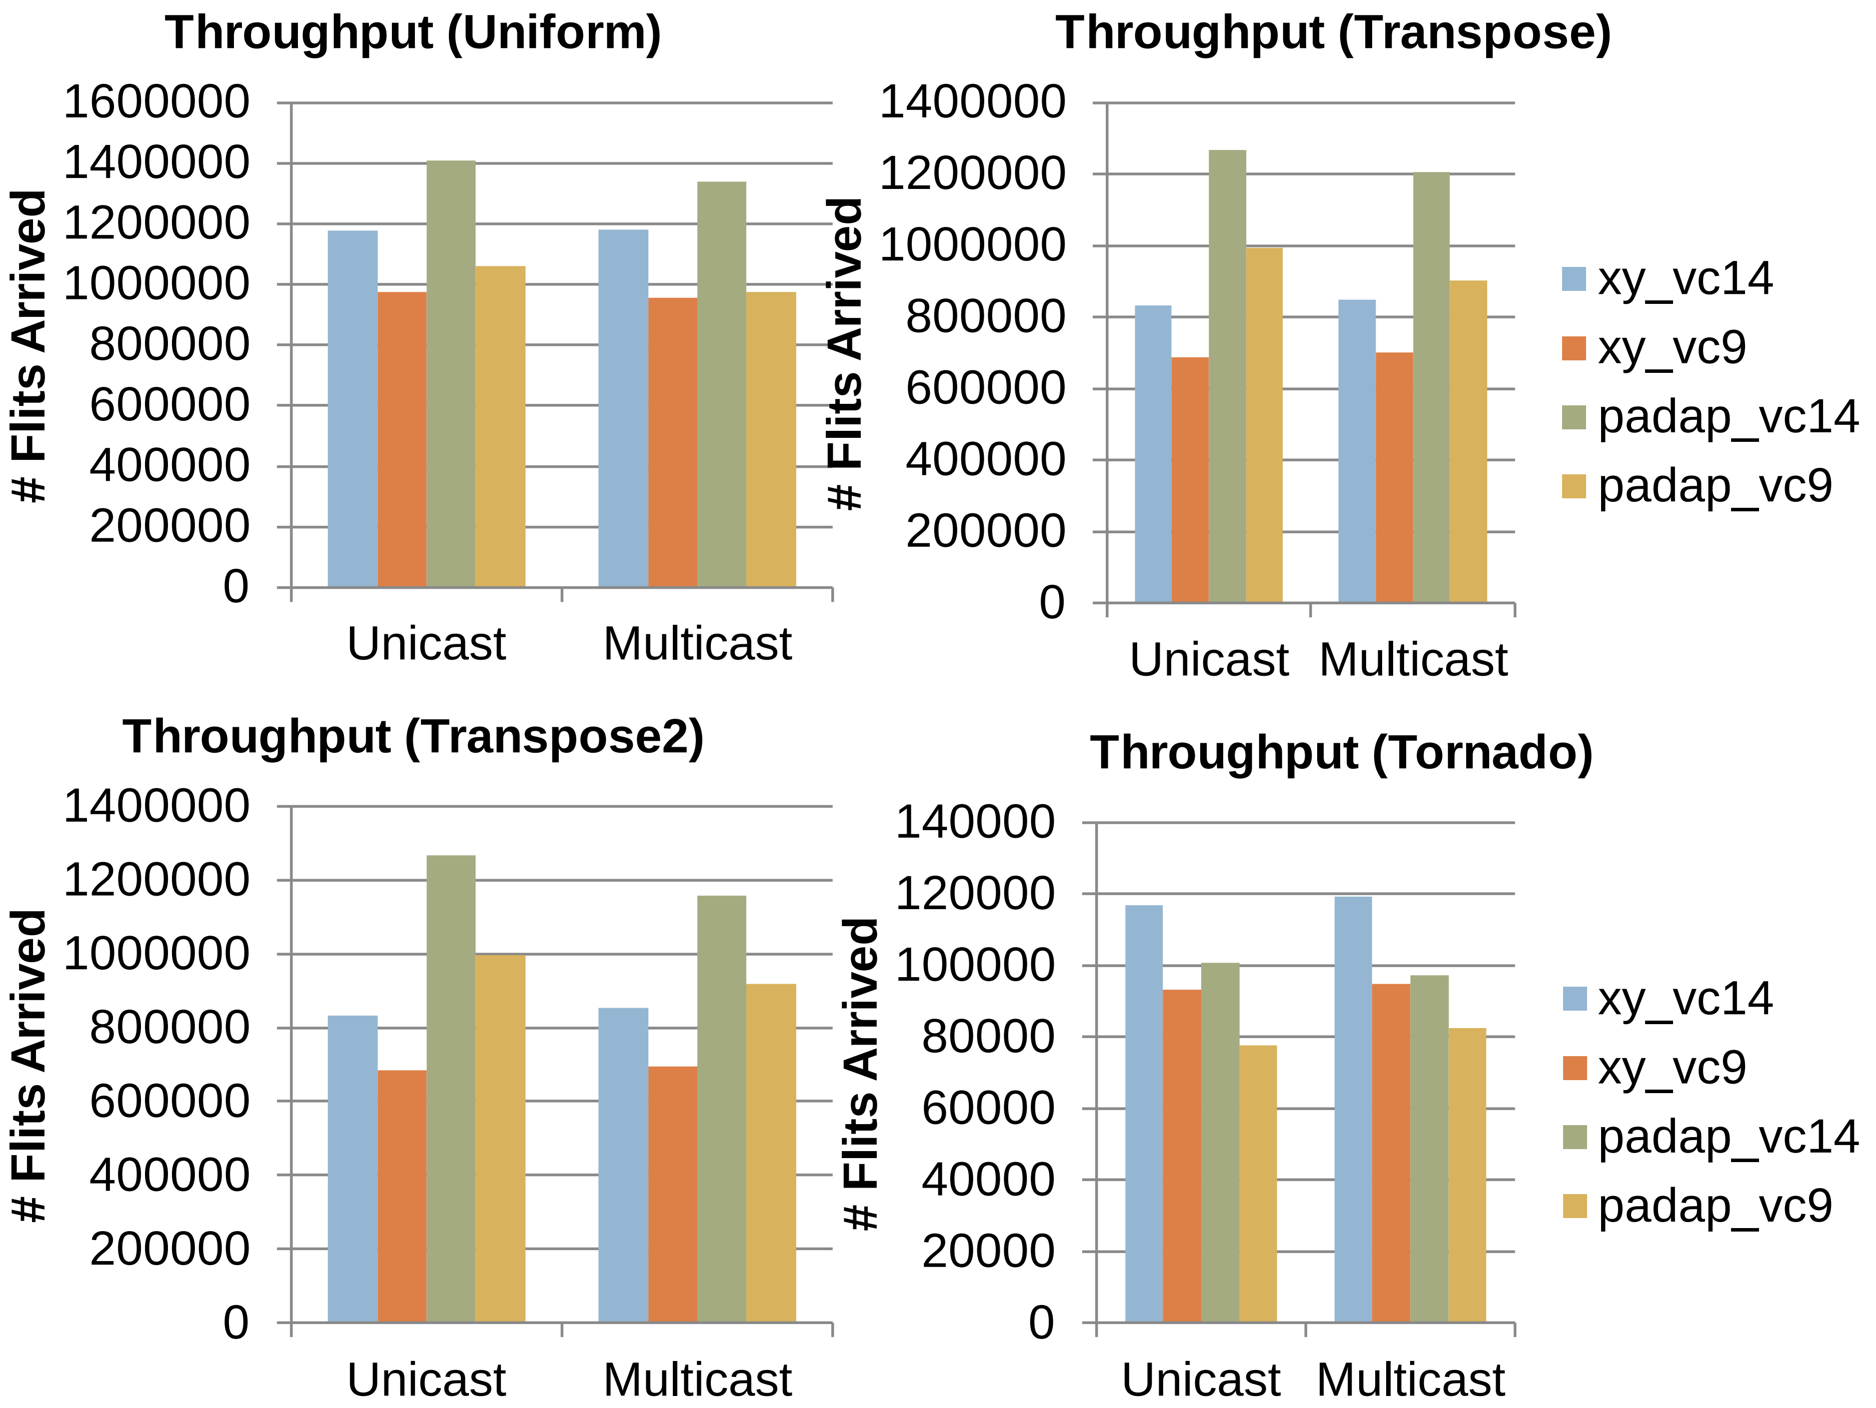
Throughput (Uniform)
Flits Arrived
1600000
1400000
1200000
1000000
800000
600000
400000
200000
0
Unicast
Multicast
Throughput (Transpose)
Flits Arrived
1400000
1200000
1000000
800000
600000
400000
200000
0
Unicast
Multicast
xy_vc14
xy_vc9
padap_vc14
padap_vc9
Throughput (Transpose2)
Flits Arrived
1400000
1200000
1000000
800000
600000
400000
200000
0
Unicast
Multicast
Throughput (Tornado)
Flits Arrived
140000
120000
100000
80000
60000
40000
20000
0
Unicast
Multicast
xy_vc14
xy_vc9
padap_vc14
padap_vc9

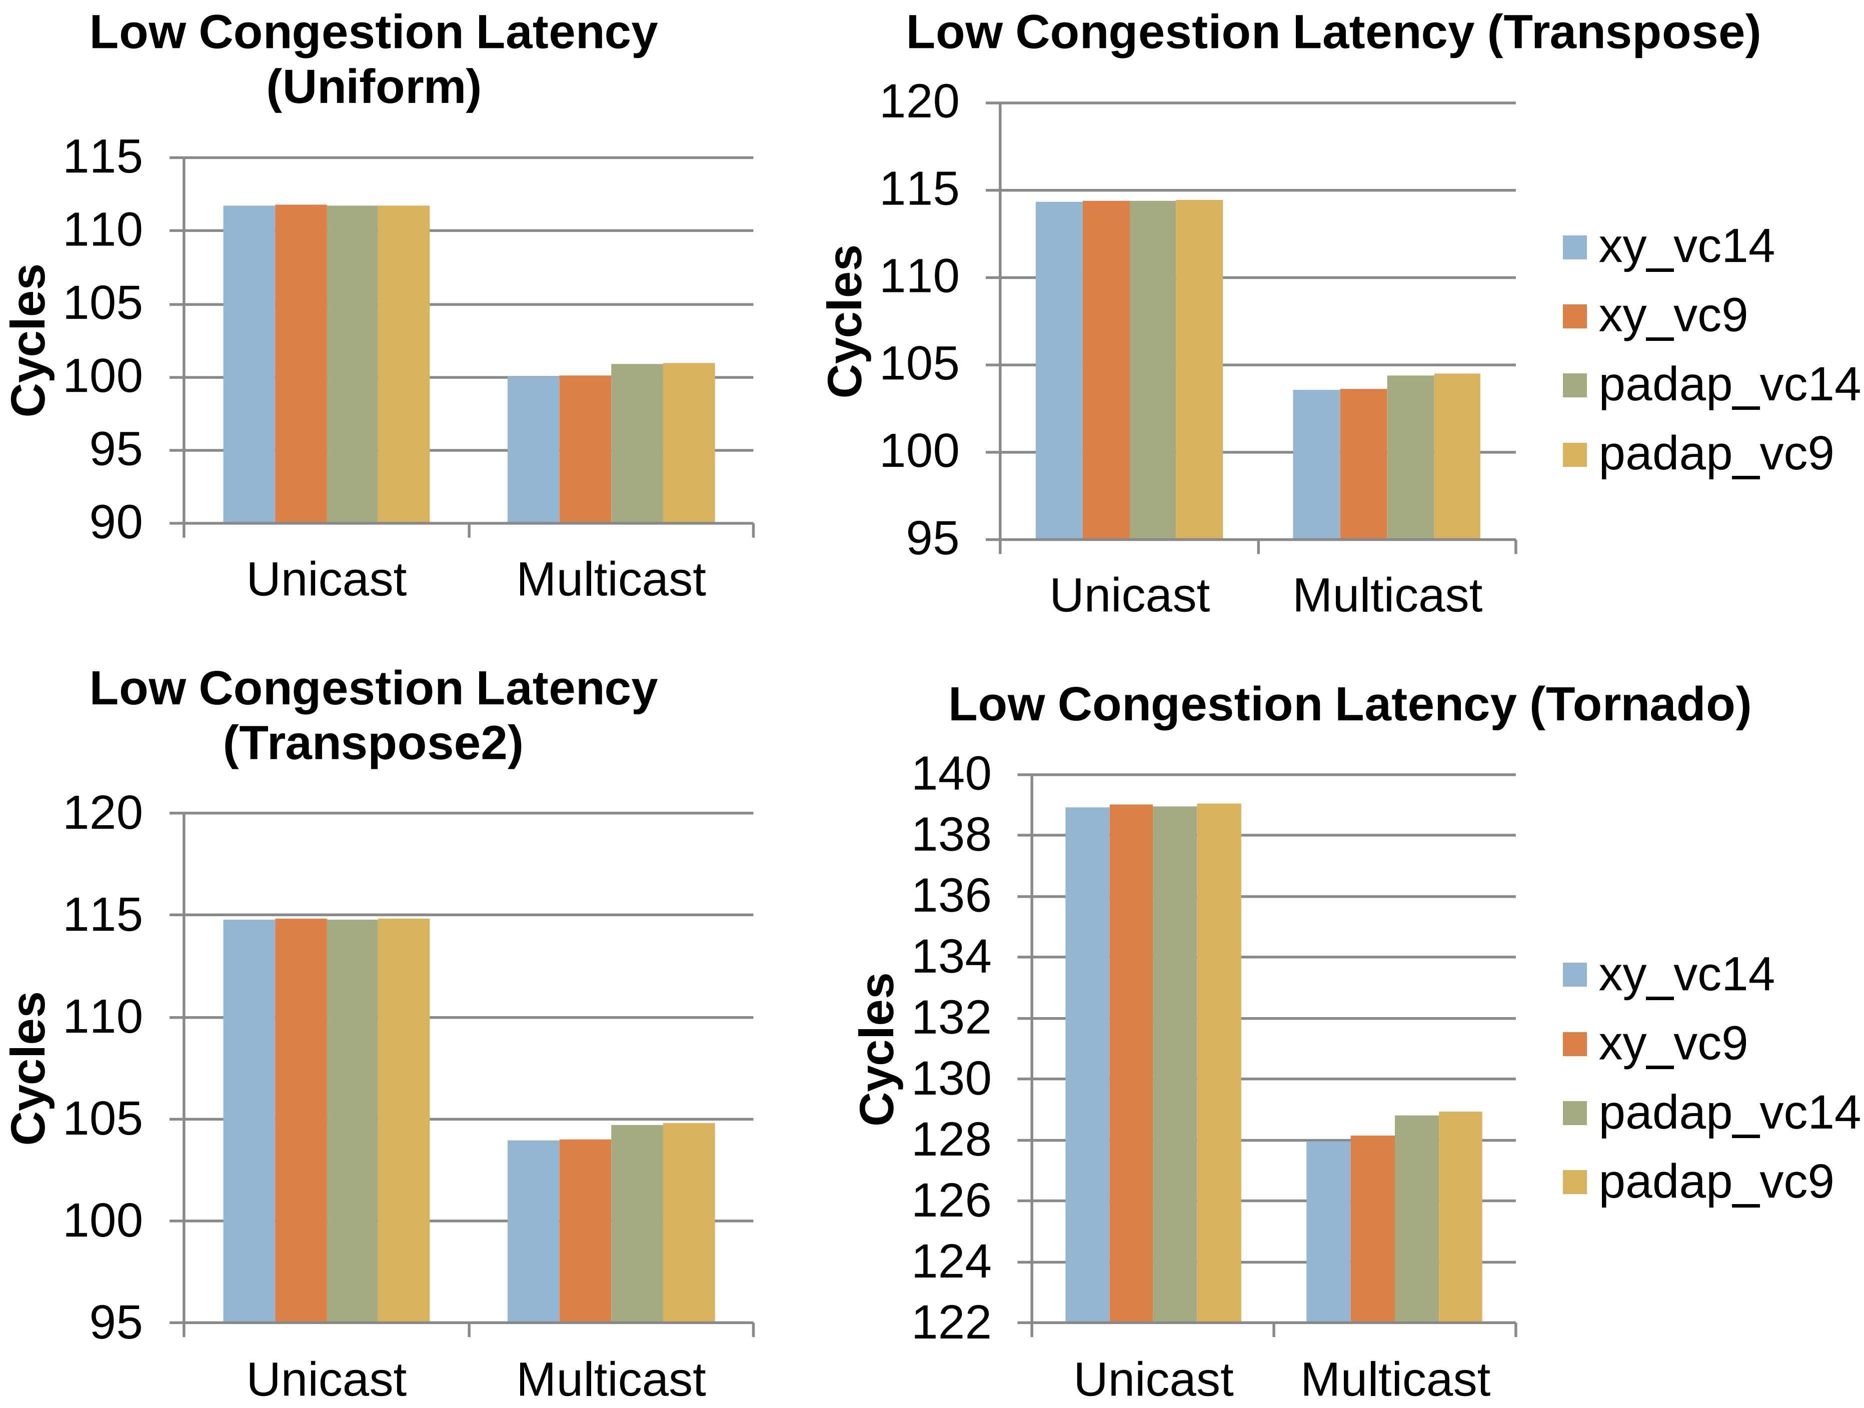
Low Congestion Latency (Uniform)
Cycles
115
110
105
100
95
90
Unicast
Multicast
Low Congestion Latency (Transpose)
Cycles
120
115
110
105
100
95
Unicast
Multicast
xy_vc14
xy_vc9
padap_vc14
padap_vc9
Low Congestion Latency (Transpose2)
Cycles
120
115
110
105
100
95
Unicast
Multicast
Low Congestion Latency (Tornado)
Cycles
140
138
136
134
132
130
128
126
124
122
Unicast
Multicast
xy_vc14
xy_vc9
padap_vc14
padap_vc9

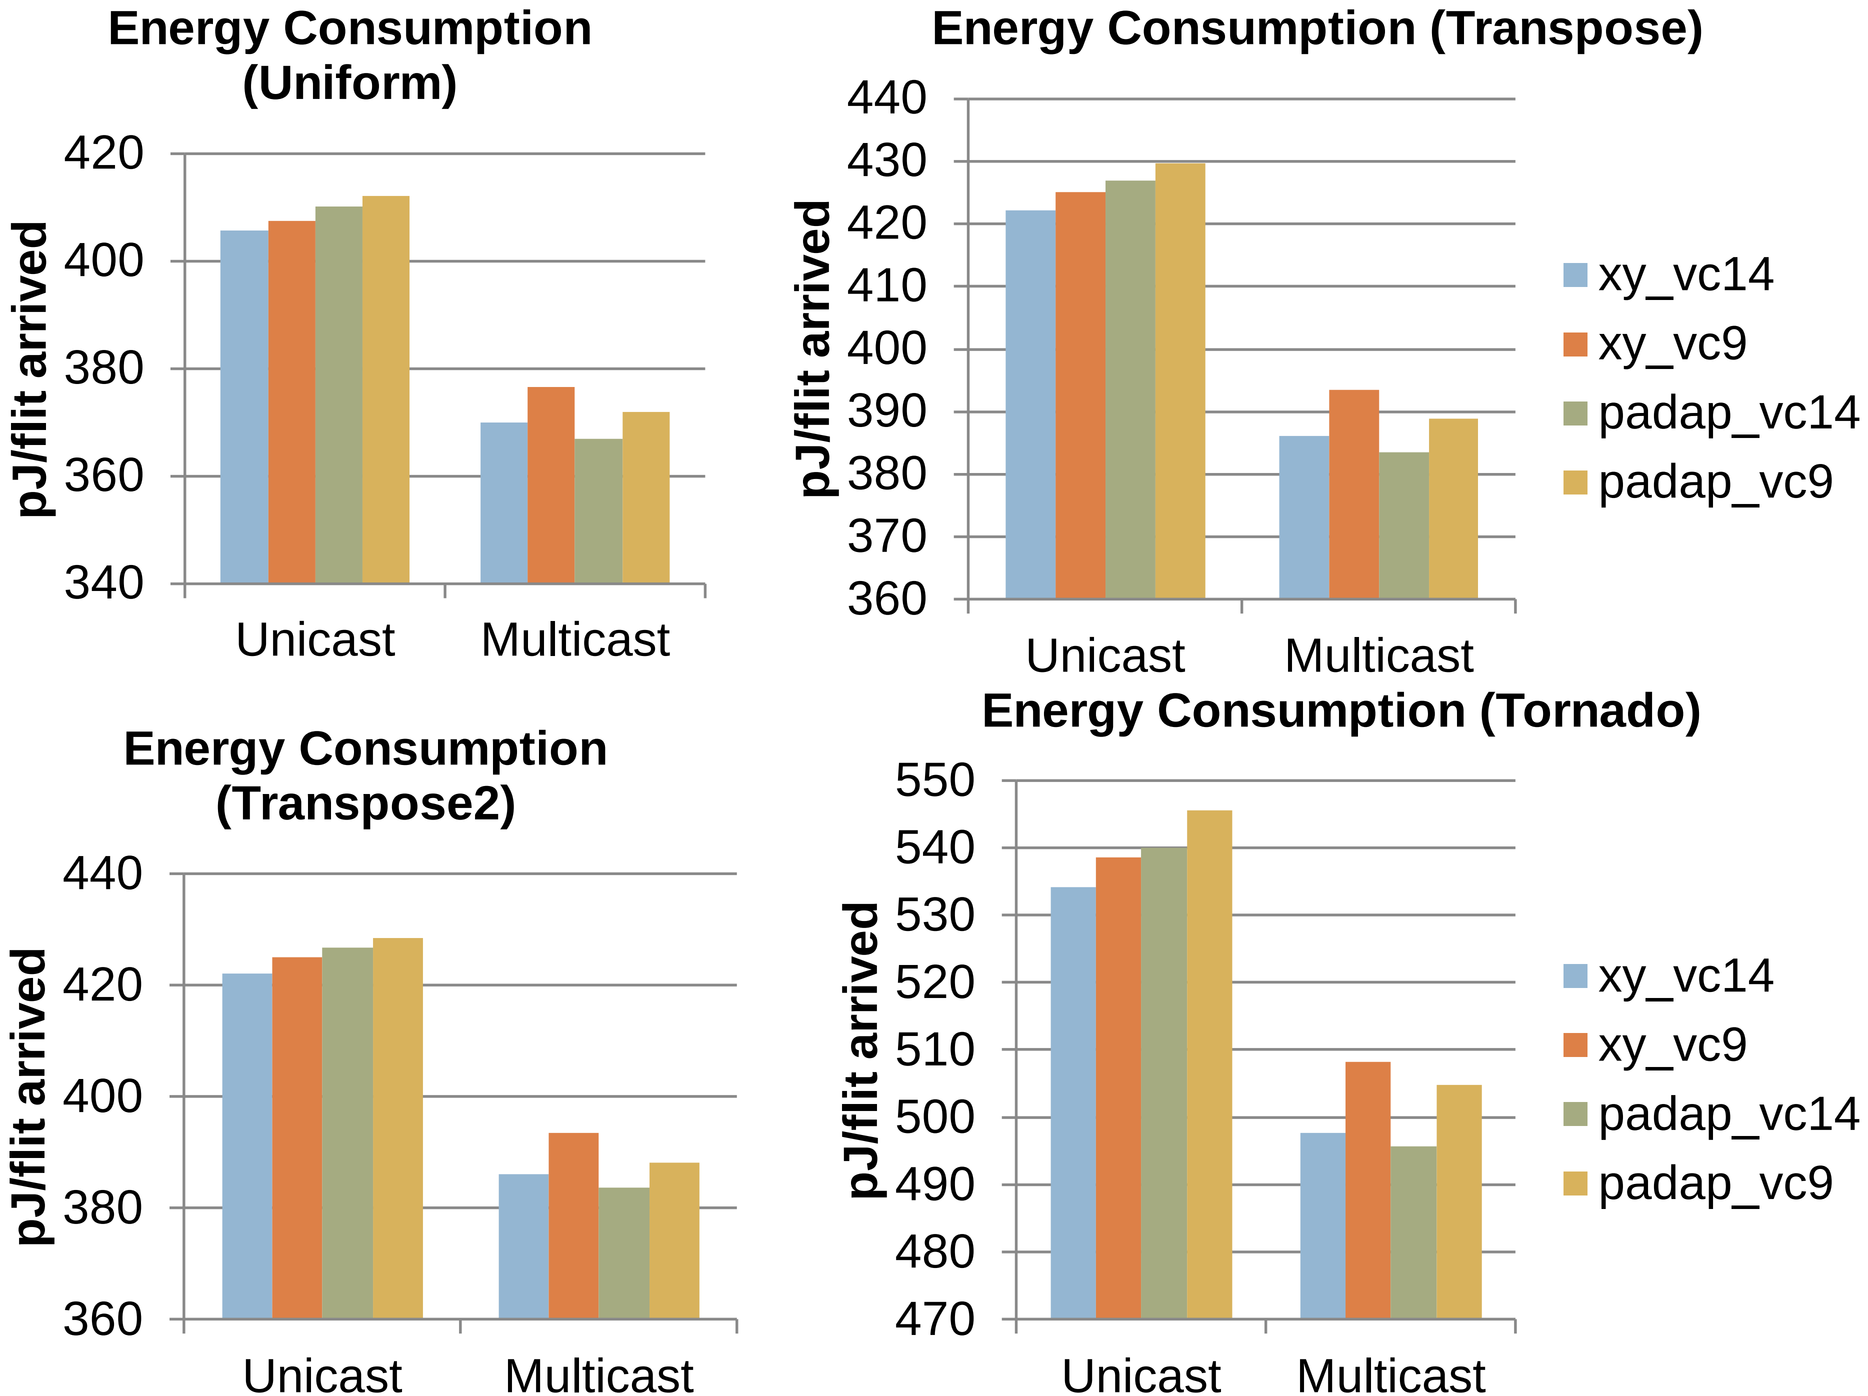
Energy Consumption (Uniform)
pJ/flit arrived
420
400
380
360
340
Unicast
Multicast
Energy Consumption (Transpose)
pJ/flit arrived
440
430
420
410
400
390
380
370
360
Unicast
Multicast
xy_vc14
xy_vc9
padap_vc14
padap_vc9
Energy Consumption (Transpose2)
pJ/flit arrived
440
420
400
380
360
Unicast
Multicast
Energy Consumption (Tornado)
pJ/flit arrived
550
540
530
520
510
500
490
480
470
Unicast
Multicast
xy_vc14
xy_vc9
padap_vc14
padap_vc9

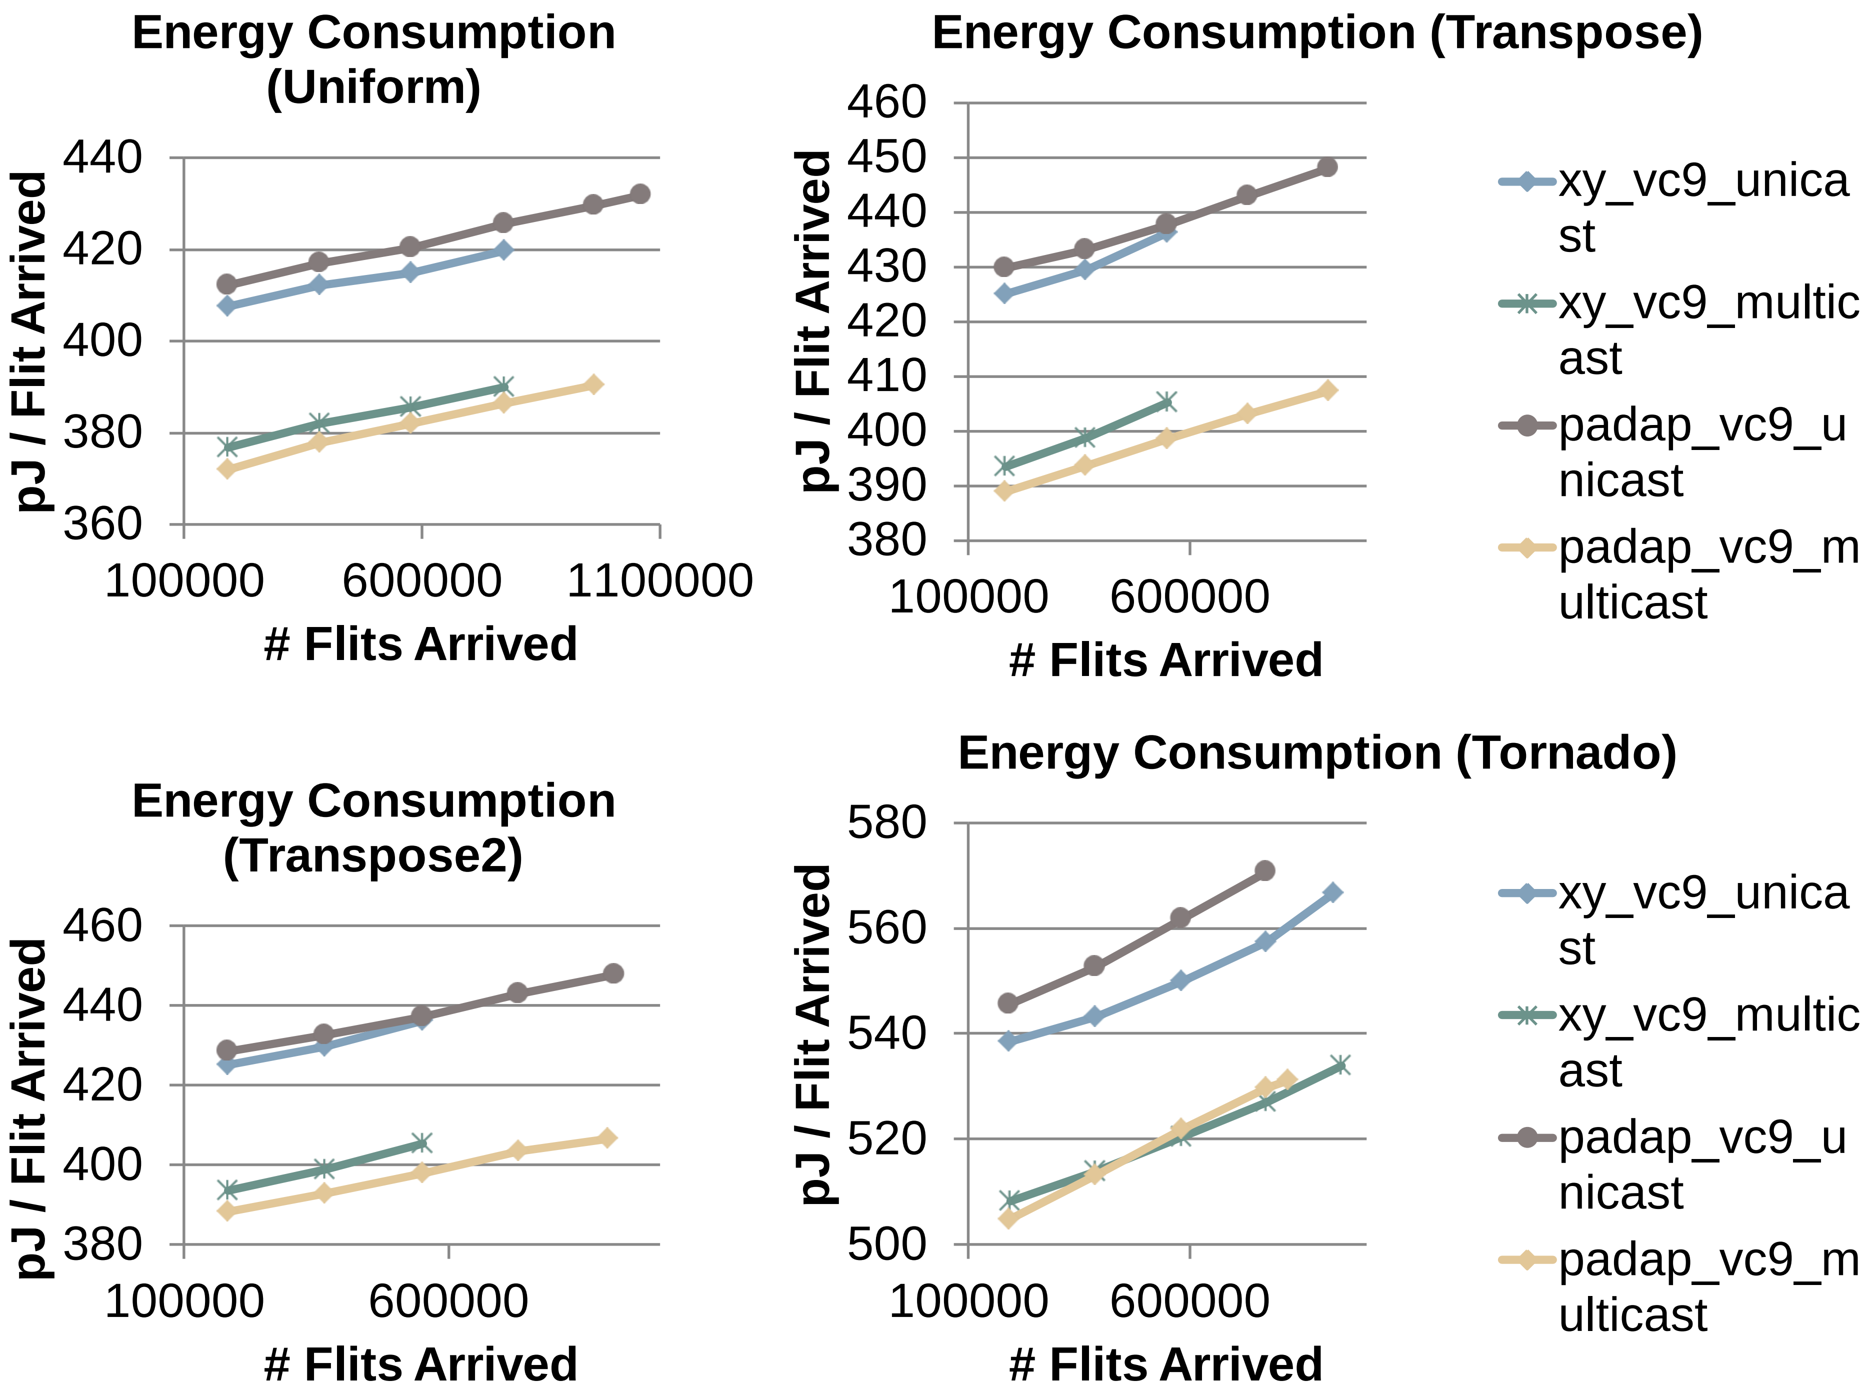
Energy Consumption (Uniform)
pJ / Flit Arrived
440
420
400
380
360
100000
600000
1100000
Flits Arrived
Energy Consumption (Transpose)
pJ / Flit Arrived
460
450
440
430
420
410
400
390
380
100000
600000
Flits Arrived
xy_vc9_unicast
xy_vc9_multicast
padap_vc9_unicast
padap_vc9_multicast
Energy Consumption (Transpose2)
pJ / Flit Arrived
460
440
420
400
380
100000
600000
Flits Arrived
Energy Consumption (Tornado)
pJ / Flit Arrived
580
560
540
520
500
100000
600000
Flits Arrived
xy_vc9_unicast
xy_vc9_multicast
padap_vc9_unicast
padap_vc9_multicast

FPGA Traffic

- CPU controls the application jobs scheduling and placement
- Each tile contains its own configuration bitstream controller

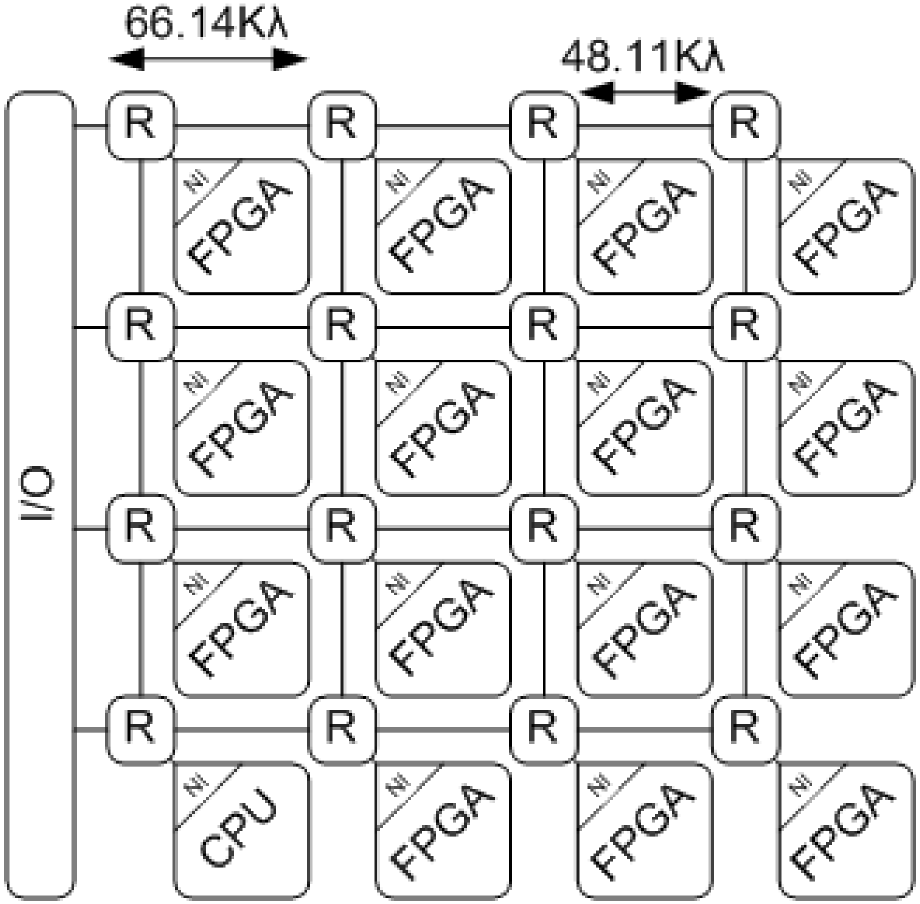

Applications

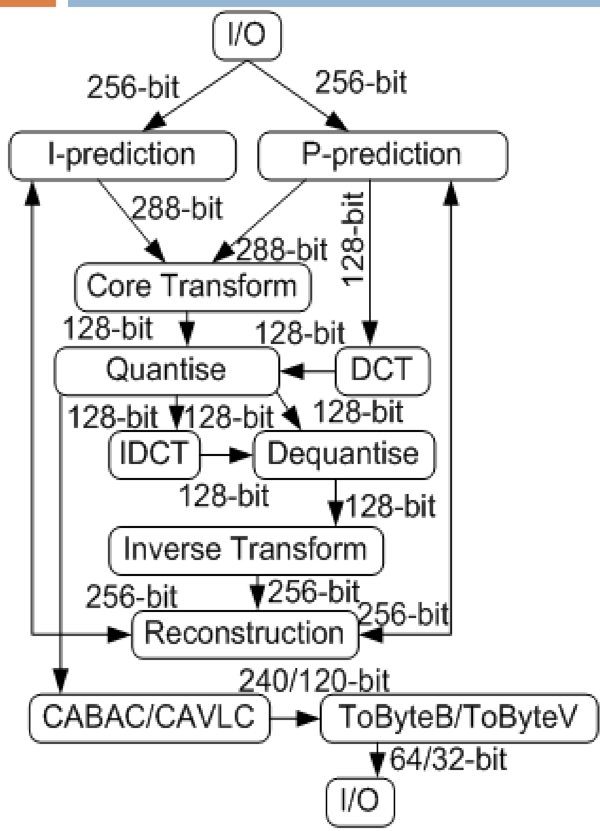

(a) MPEG4 Encoder

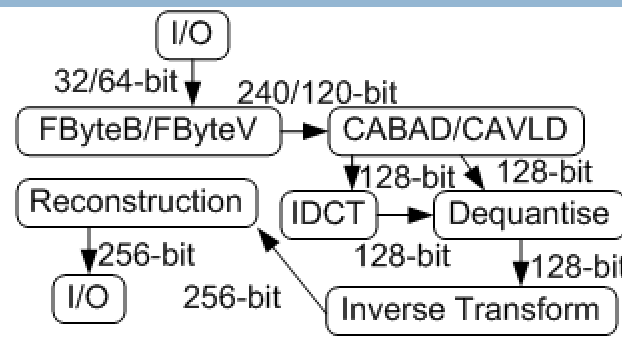

(b) MPEG4 Decoder

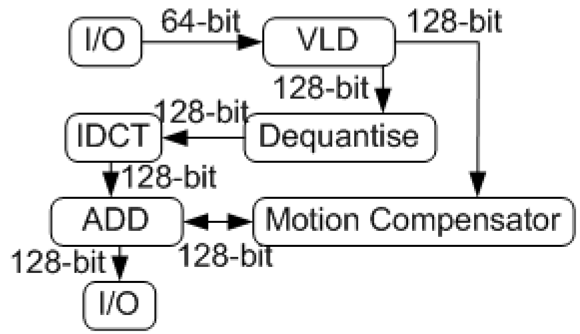

(c) MPEG2 Decoder

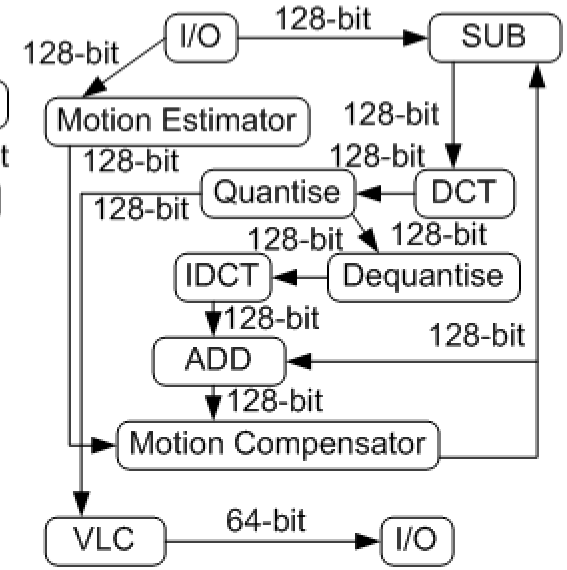

(d) MPEG2 Encoder

Experimental Setup

- Mesh Size: 20x20
- Flit Size: 128-bit
- Simulation Cycle: 200,000,000
- Virtual Channel Depth: 14
- Max Packet Length:
  - 10 (Virtual Cut-Through)
  - 20 (Address-Data FIFO decoupling)
- #Virtual Channel:
  - 3 unicast channels (unicast router)
  - 2 unicast + 1 multicast channel (multicast router)

Average Tile Configuration Time

- Adaptive routing can reduce the configuration time by at most 10%
- With address-data decoupling, configuration time can be reduced by at most 25%
- Multicast support reduces configuration time by at most 40%

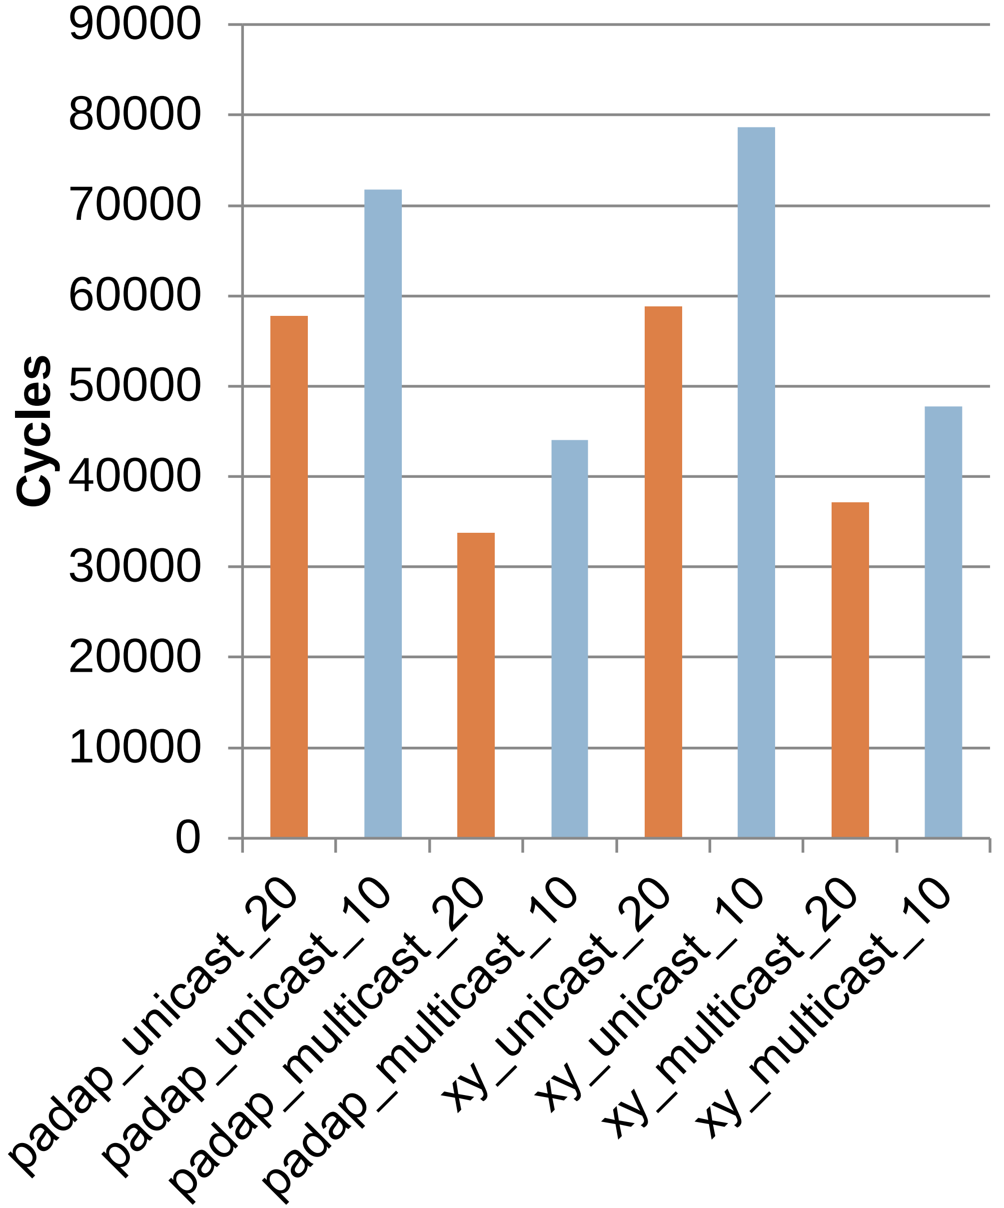

- Adaptive routing can reduce the application runtime by at most 6%
- With address-data decoupling, application runtime can be reduced by at most 10%
- Multicast support reduces application runtime by at most 4%

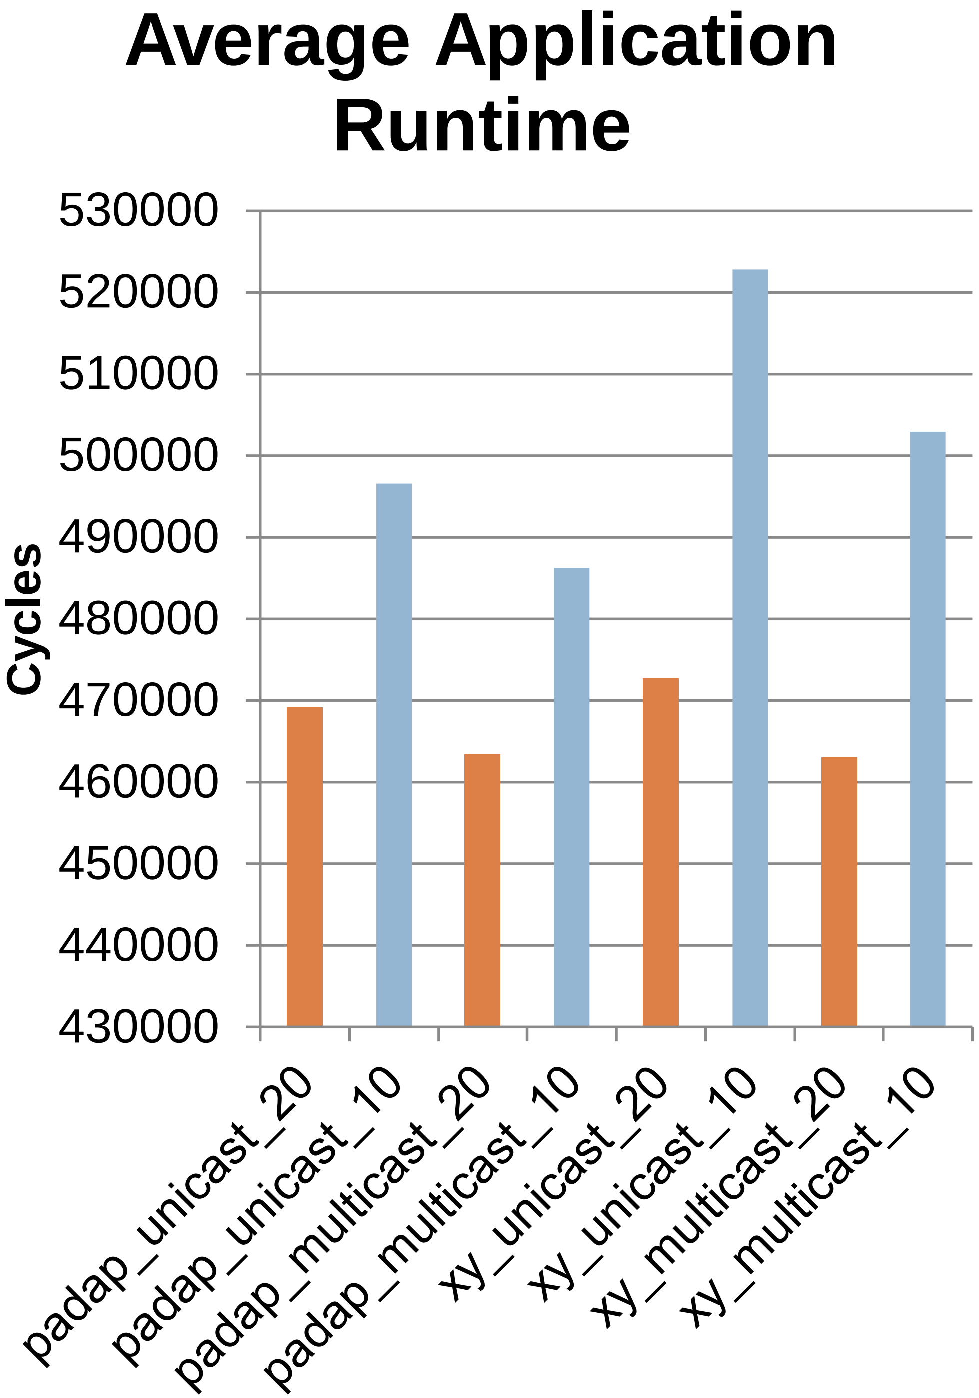

- Adaptive routing can reduce the configuration energy by at most 4%
- With address-data decoupling, configuration energy can be reduced by at most 5%
- Multicast support reduces energy consumption by at most 20%

Network Energy Consumption per Tile Reconfiguration

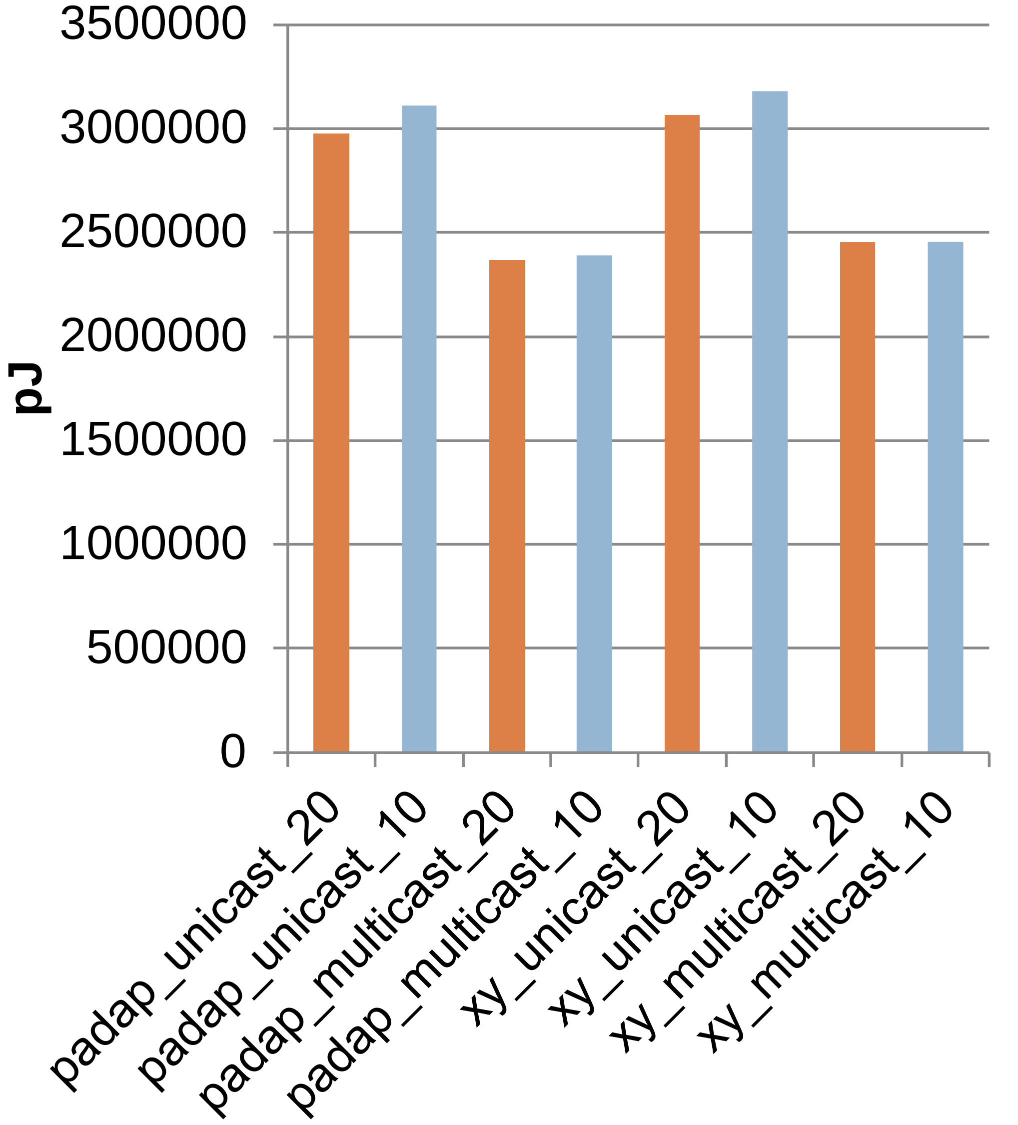

Hardware Cost

Router type	XY/ Adaptive	Unicast/ Multicast	VC Depth	Area(Mλ^2)	RC(ps)	SA(ps)
Wormhole	XY	Unicast	9	208	375	1151
Decouple	XY	Multicast	9	308	528	1552
VCT	XY	Multicast	14	396	515	1161
Wormhole	Adaptive	Unicast	9	216	843	1125
Decouple	Adaptive	Multicast	9	328	1242	1616
VCT	Adaptive	Multicast	14	417	1197	1134

- With address-data decoupling, router can support packet with length 20 using 23% lesser area
- Address-data decoupling increases critical path by 35%
- Multicast support increases area by 50% in 2U+1M

This is the end of the presentation.

Thank You.