Rotating Priority Queues

Rotating Priority Queues

- **Idea:**
  - Approximate a sorted scheduler queue (e.g., EDF, WFQ) by FIFO queues
  - FIFO queues are periodically rearranged (“rotated”)

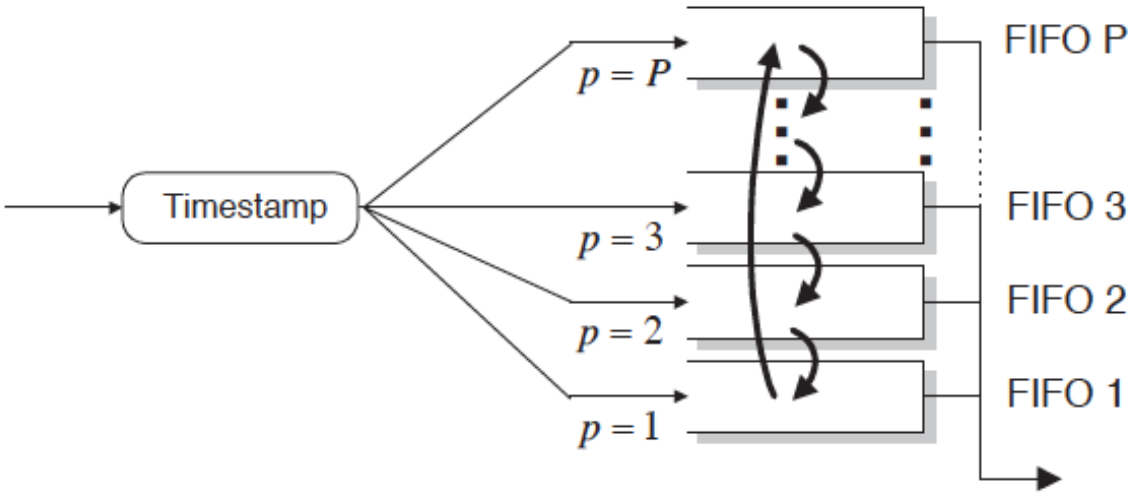

  - Here: Approximate Earliest-Deadline-First with FIFO queue

Rotating-Priority-Queues (RPQ)

Design Principles:

- P priority sets.
- $P + 1$ FIFO queues with labels.
- Relabel queues every Δ time units.
- One delay bound for each priority set: $d_p = p \cdot \Delta$.

RPQ Scheduler

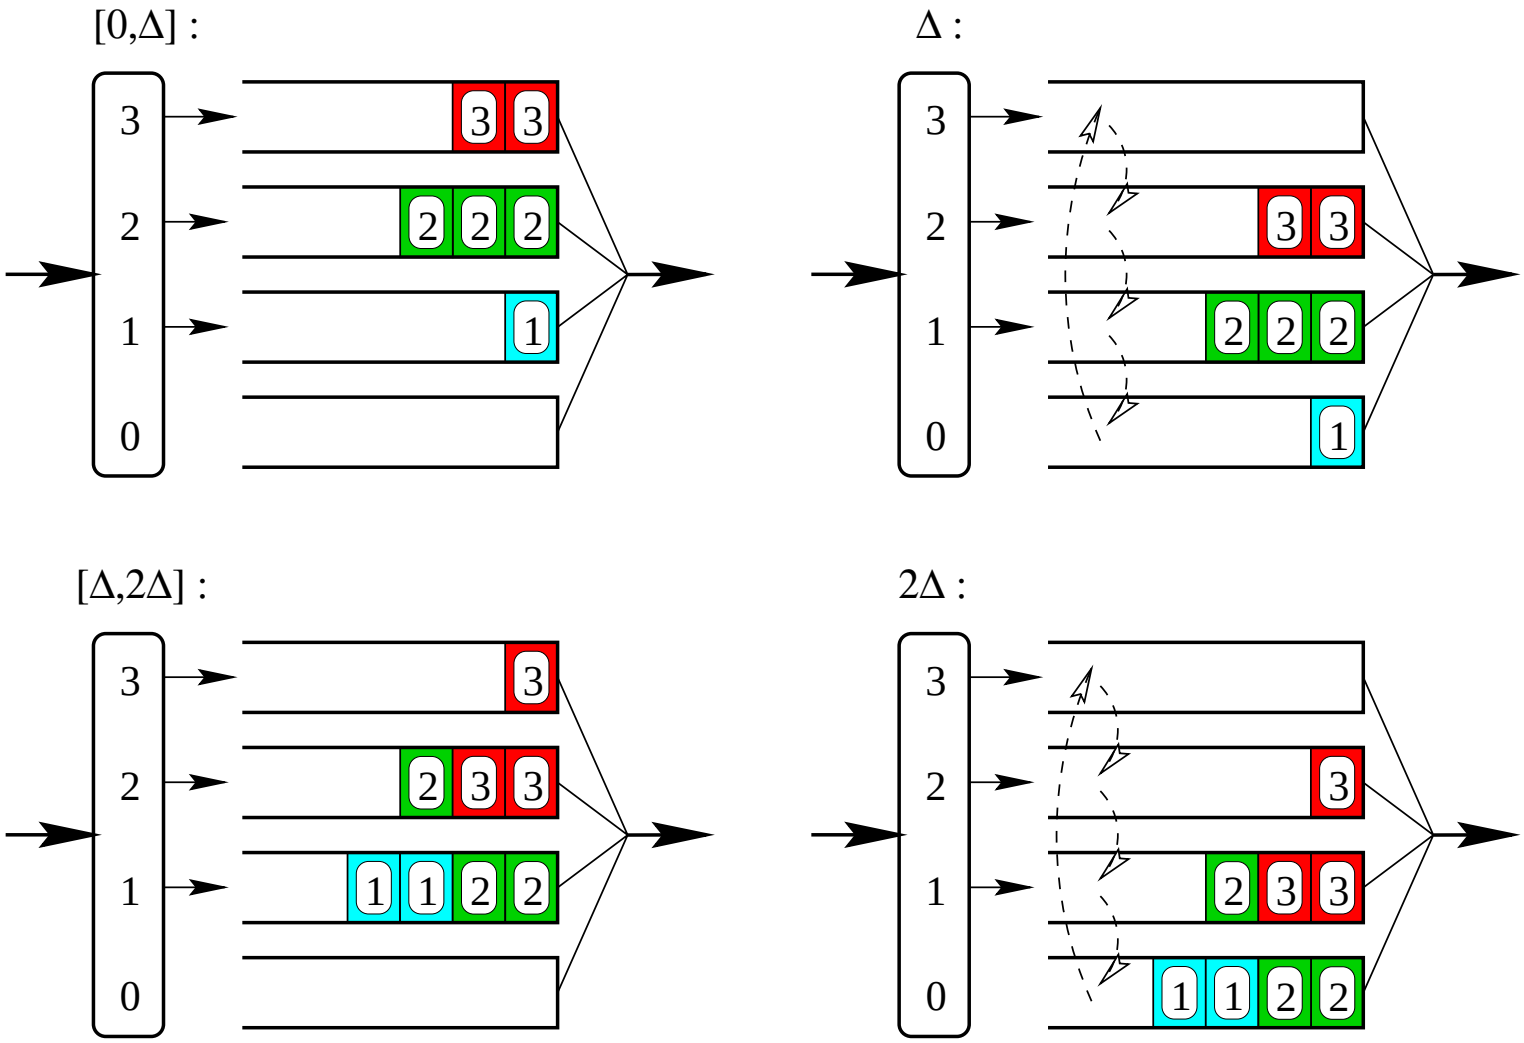

Experimental Setup

- Single 155 Mbps switch.
- Three connection groups *Low, Medium, High Delay*.

	Index j	Delay Bound d_j	Burst Size B_j	Rate r_j
Low	1	12 ms	4,000 cells	10-155 Mbps
Medium	2	24 ms	2,000 cells	10-155 Mbps
High	3	36 ms	4,000 cells	10-155 Mbps

Evaluation

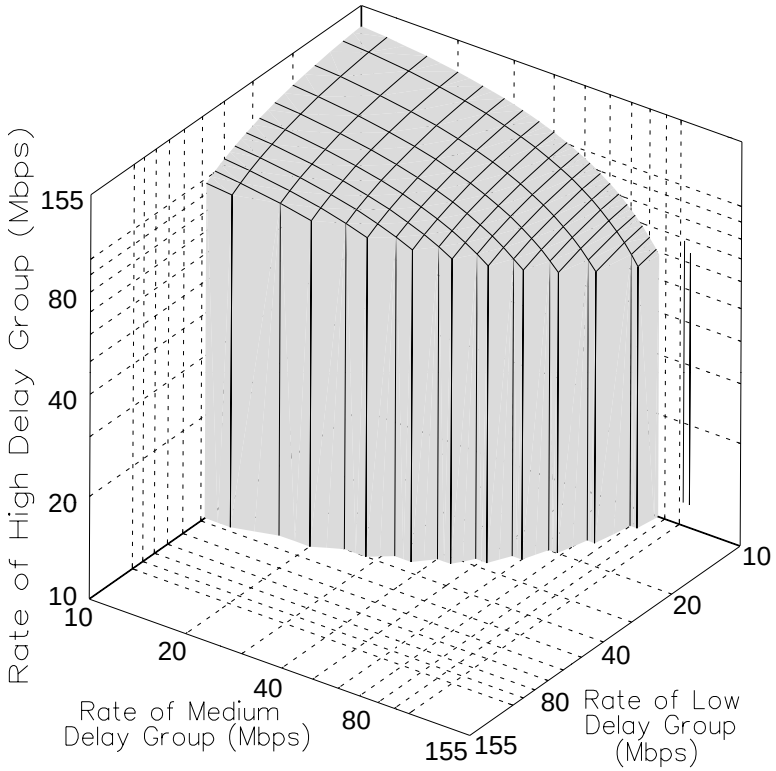

EDF

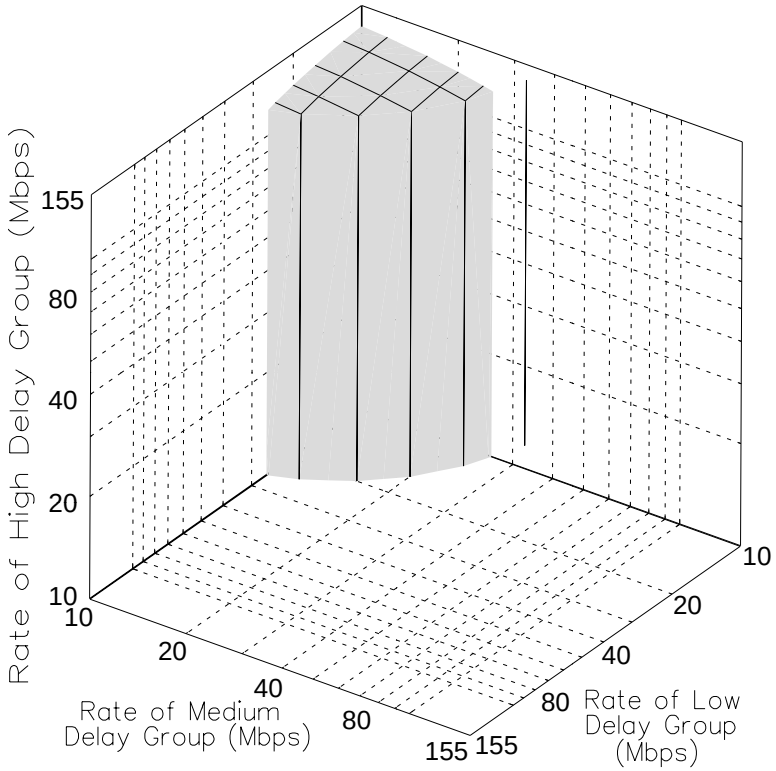

SP

Evaluation of RPQ

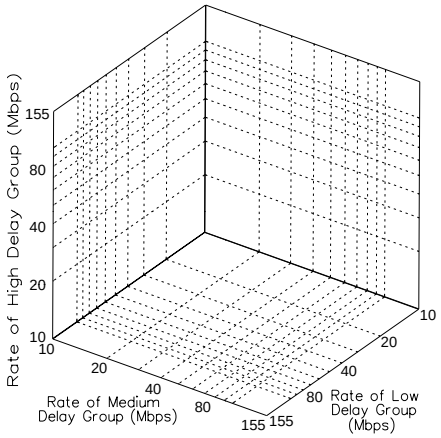

RPQ ($\Delta = 12ms$; 6 FIFOs)

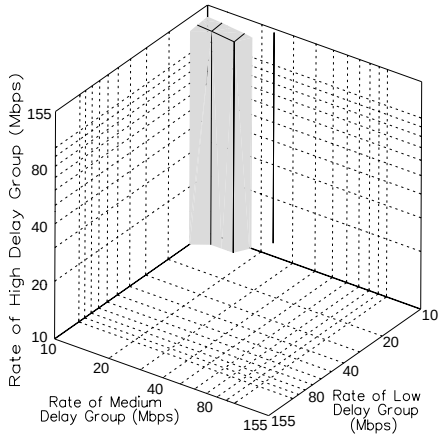

RPQ ($\Delta = 6ms$; 12 FIFOs)

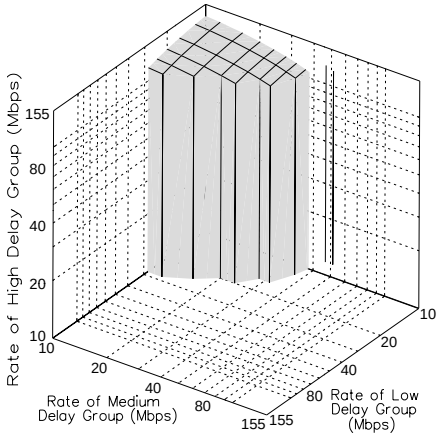

RPQ ($\Delta = 4ms$; 18 FIFOs)

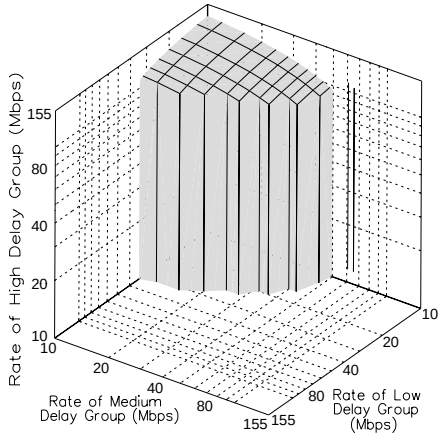

RPQ ($\Delta = 3ms$; 24 FIFOs)

Rotating Anomaly

- For $\Delta \rightarrow 0$: RPQ = EDF
- For $\Delta \rightarrow \infty$: RPQ = SP
- For Δ small: RPQ is worse than SP

- Why?

Rotation Anomaly

Arrival scenario:

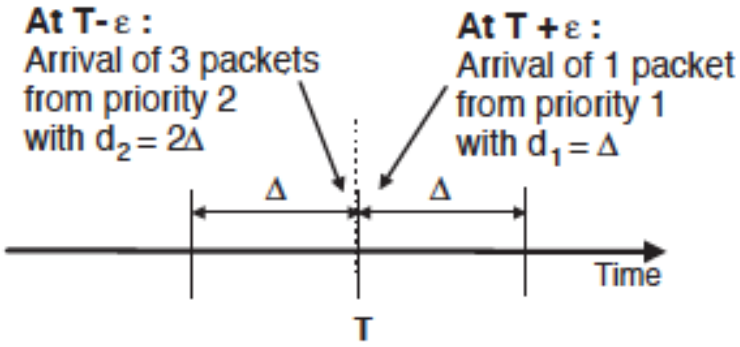

delay bounds:
priority 1: Δ
priority 2: 2Δ

at T-ε:

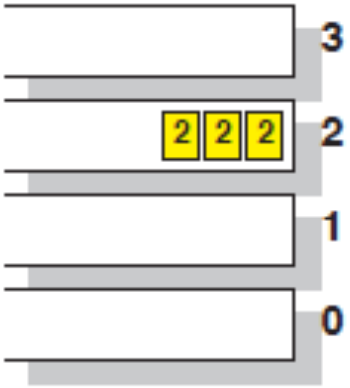

at T+ε:

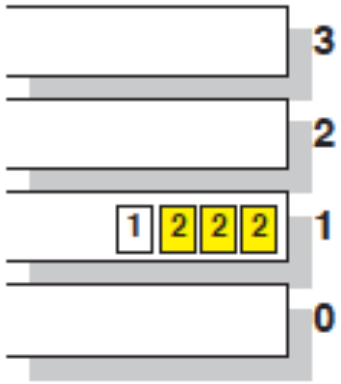

- Observation: Rotation can put packets with a later deadline ahead (“deadline inversion)
- Can rotation anomaly be avoided?

Rotating-Priority-Queues$^+$ (RPQ$^+$)

Design Principles:

- P priority sets.
- $2P$ FIFO queues with labels.
- Relabel queues every Δ time units.
- One delay bound for each priority set: $d_p = p \cdot \Delta$.

RPQ^{+} Scheduler

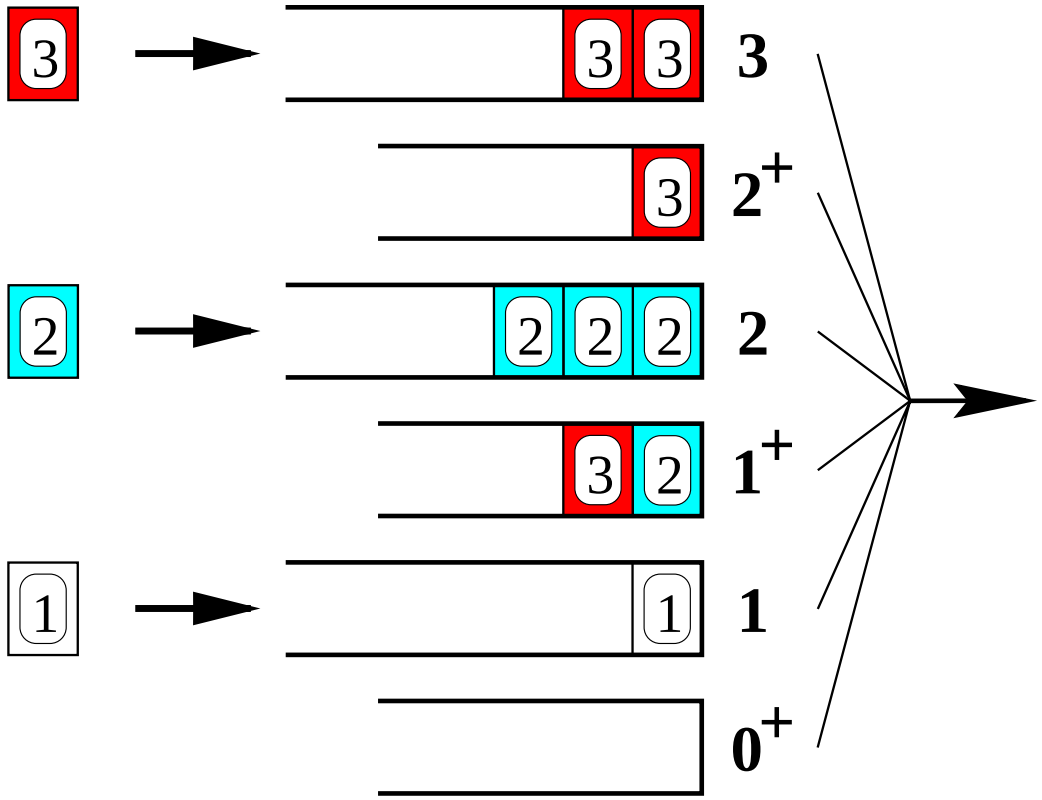

RPQ$^+$ Queue Rotation

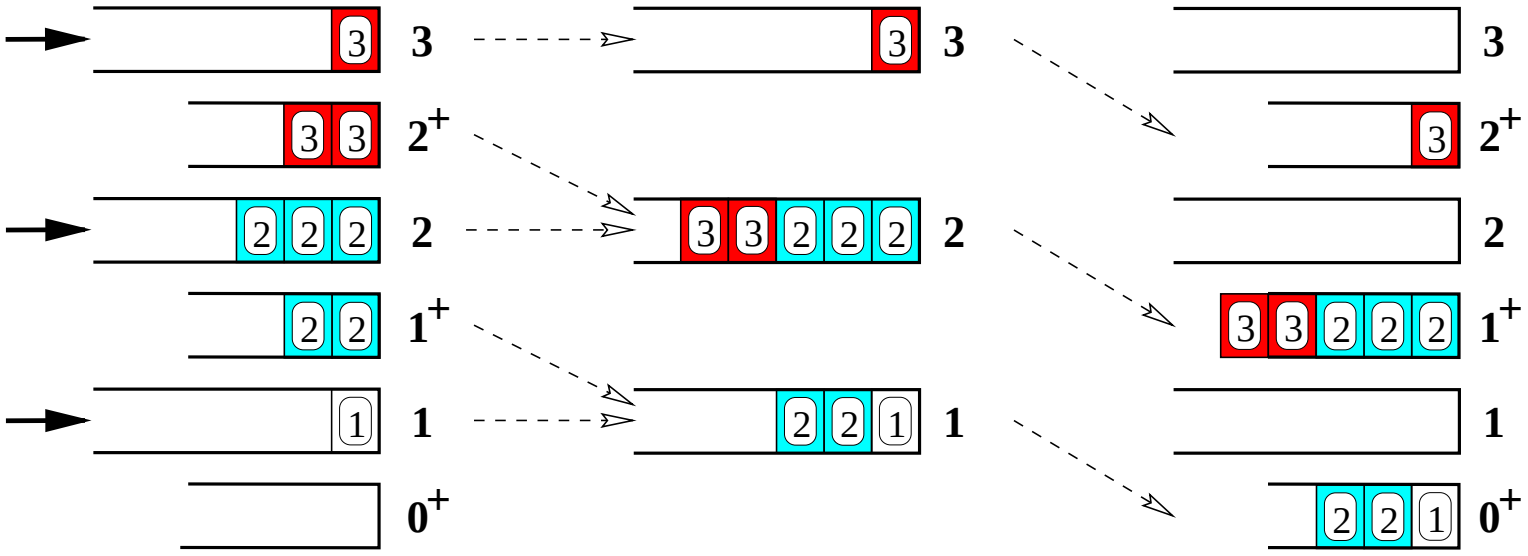

Implementating RPQ$^+$ in Shared Memory

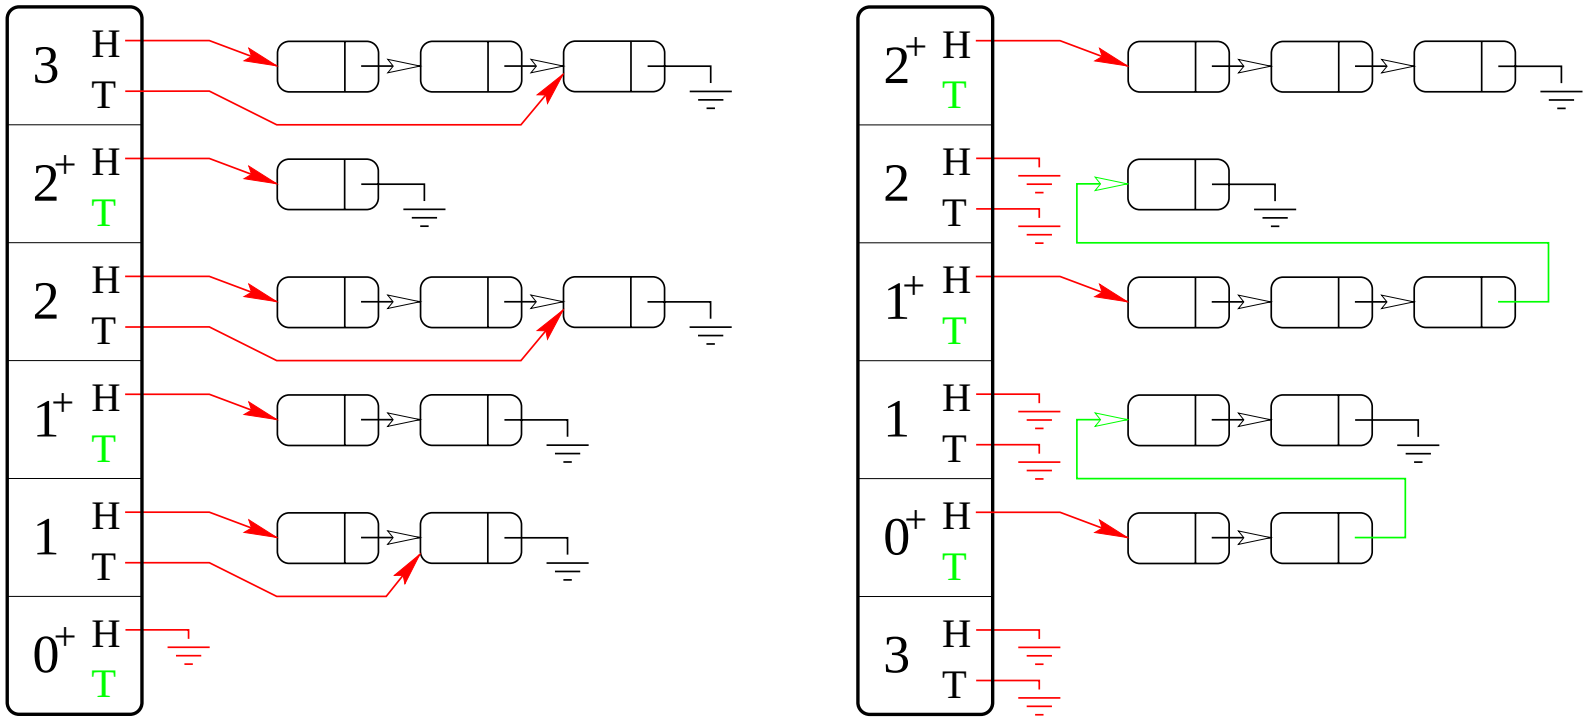

- No movement of packets.
- Operations independent of queued packets.

Evaluation of RPQ$^+$

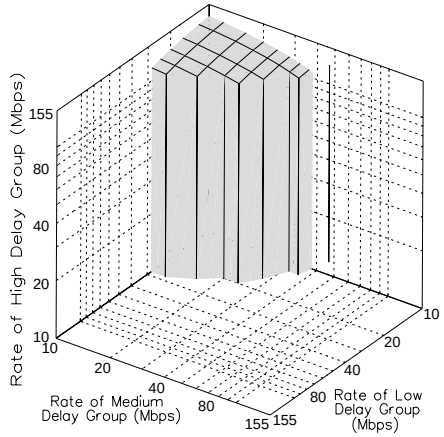

RPQ$^+$ ($\Delta = 12ms$; 6 FIFOs)

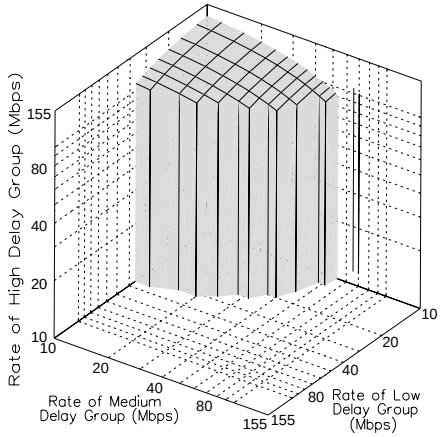

RPQ$^+$ ($\Delta = 6ms$; 12 FIFOs)

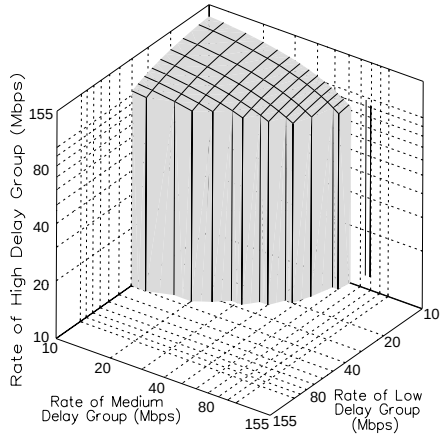

RPQ$^+$ ($\Delta = 4ms$; 18 FIFOs)

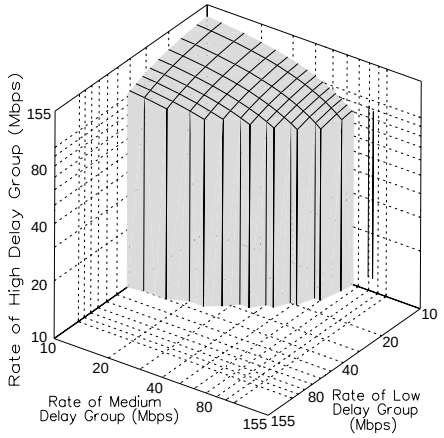

RPQ$^+$ ($\Delta = 3ms$; 24 FIFOs)

Summary of Evaluation

- Compare volume of the schedulable regions:
  $\frac{V^{\Sigma}(\Delta)}{V^{\infty}} \cdot 100\%$

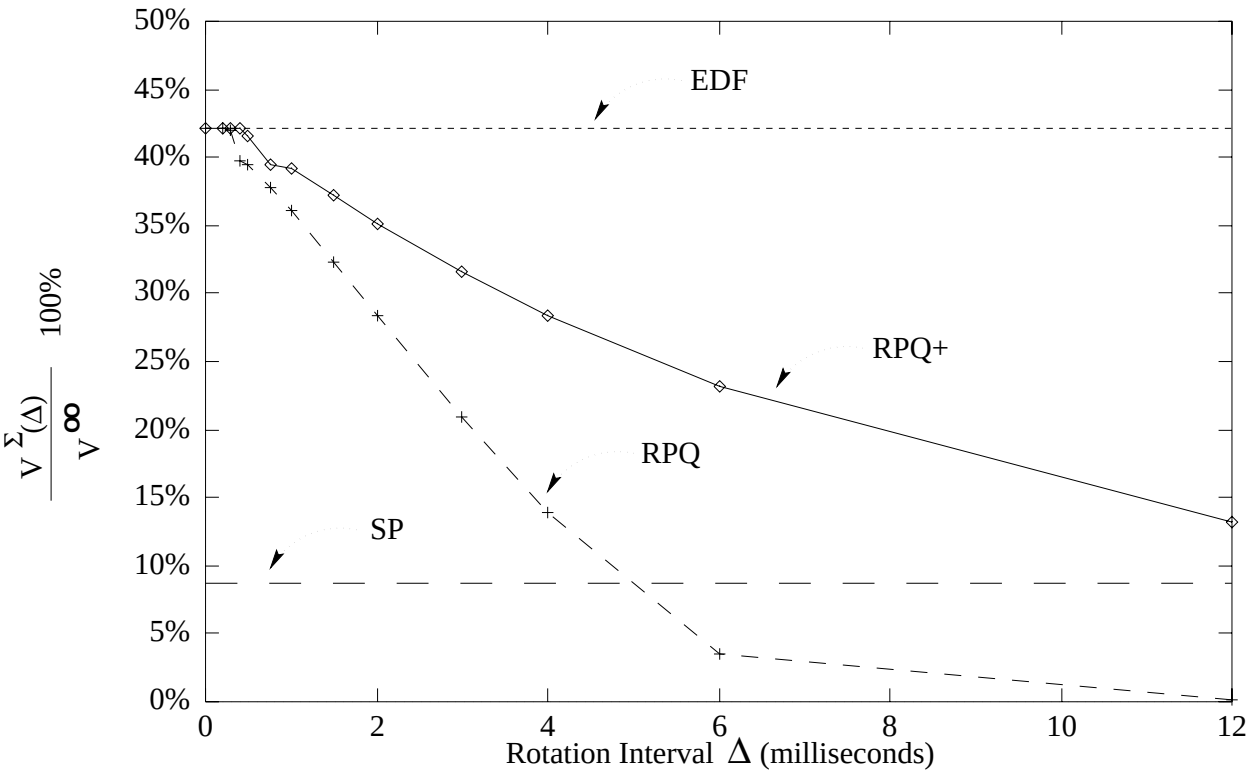

Summary of RPQ$^+$

- For $\Delta \to 0$: RPQ$^+$ = EDF
- For $\Delta \to \infty$: RPQ$^+$ = SP
- For Δ small: RPQ$^+$ always better than SP

- No rotation anomaly.