Scaling RPR with multiple rings: One control plane multiple transit paths

Fredrik Orava

Outline

- Definition--what is multiple rings
  - independently operating unidirectional transit paths
- Motivation--why multiple rings
  - scalability
  - protection
  - asymmetric capacity
- How to implement
  - simple MAC model
  - implications
- Proposal

Multiple rings

DYNARC

- Links:
  - multiple (>2)
  - independently operating
  - unidirectional
- Nodes:
  - connect to number of links
  - aggregate physical links into logical interfaces
    - several physical Mac:s into logical MAC

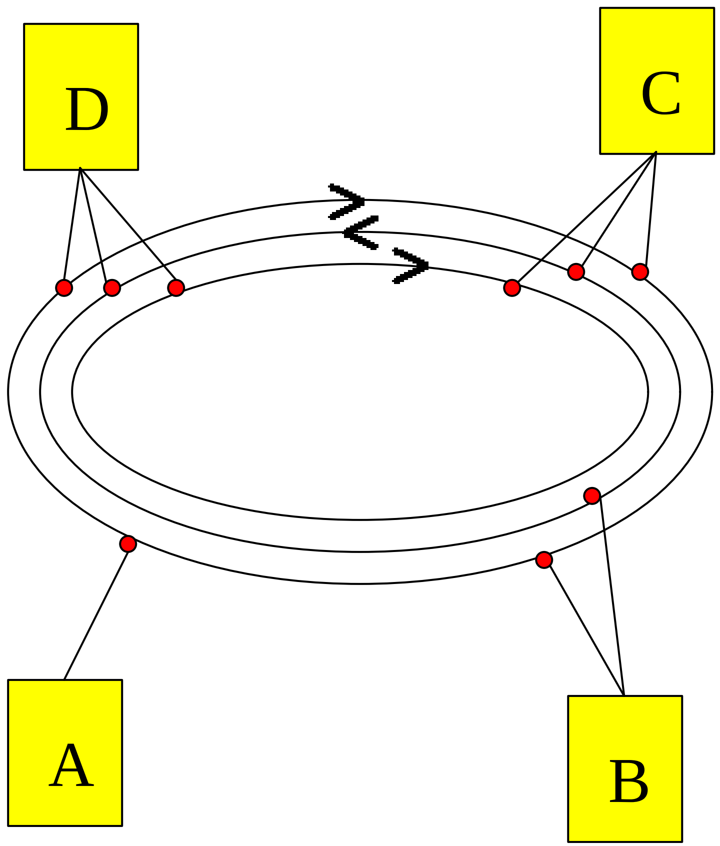

Scalability

- Problems:
  - some fibres has limited transmission capacity
    - cannot scale by increasing speed or use WDM
  - high speed optics expensive
    - costly to scale by increasing link speed
  - linear increase
    - no scaling by factor of four or increase in magnitude
- Conclusion:
  - scale by adding multiple rings
- Benefits
  - each additional ring increases the capacity
  - cheaper to add ring than to increase speed
  - individual rings can be operated at different speed
  - one logical Mac--several physical
    - the rings are managed as one aggregated link

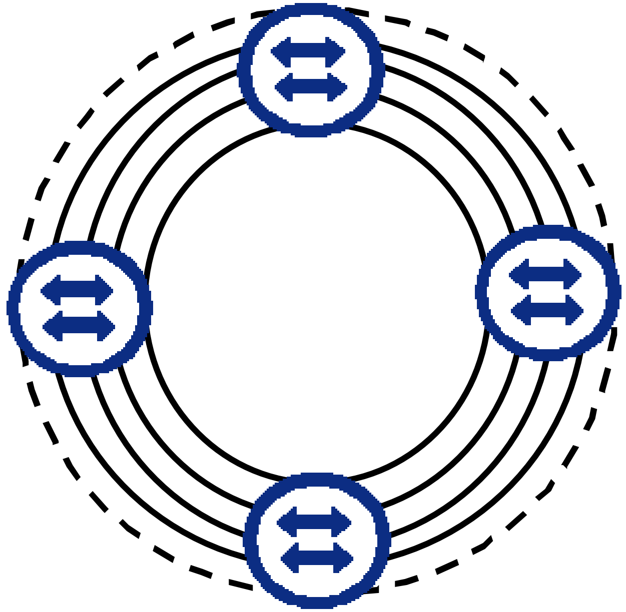

Protection

DYNARC

Goal: Should be efficient in terms of bandwidth and be able to handle multiple failures.

- Any ring should be able to protect traffic carried on other rings
  - Can distribute protected traffic over several rings
  - More efficient than 1+1 protection (can utilise more of the total bandwidth)
  - 1+1 cannot handle more than one independent failure
  - 1+1 schema can never utilise more than 50% of total capacity
- Example:
  - Assume equal capacity on all rings, equal distribution of protection capacity
  - N: number of links
  - F: number of link failures (assume ring goes away on failure)
  - Protection capacity on each link (%): P = F/N
  - Available capacity on each link (%): A = 1-F/N
  - Assume four rings of capacity C, one link failure
    - 1+1: total capacity 2*C
    - Independent: total capacity 4*C*(1-1/4) = 12*C/4 = 3*C

Protection (cont'd)

Goal: Should be able to repair (replace) component without disturbing other traffic.

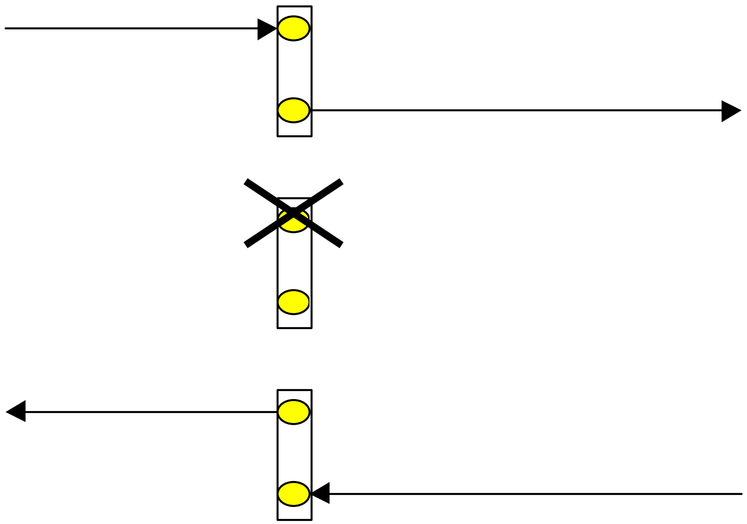

- Requirement:
  - No dependency between paths!
  - Rings should be operated independently!

Management of multiple rings

DYNARC

- Aggregate several physical MAC:s to one logical
  - Higher layers sees one interface towards the aggregated link.
  - MAC client routes traffic to logical MAC.
  - How to choose physical MAC?
    - map traffic aggregate to virtual output queue
    - assign VOQ to physical MAC
    - needs topology knowledge (topology/nodes attached)
    - needs resource knowledge (required/available)

MAC-model:

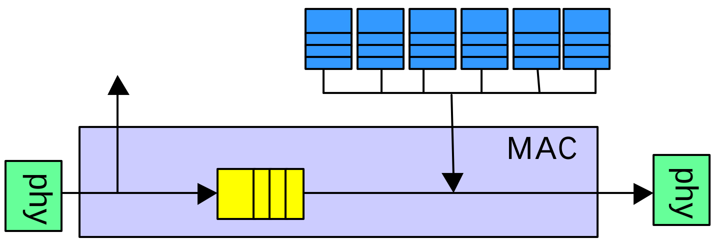

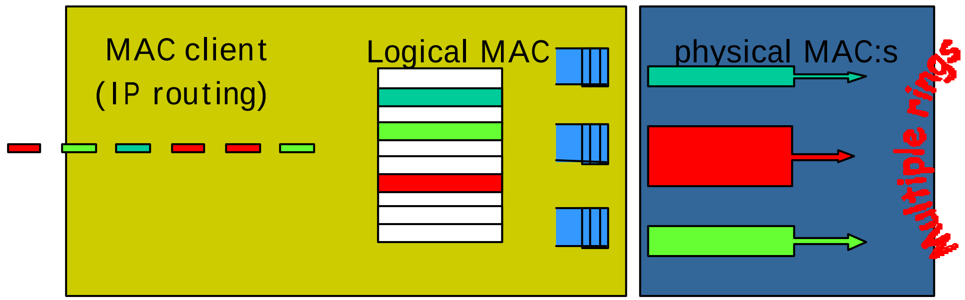

Use model from link aggregation in 802.3

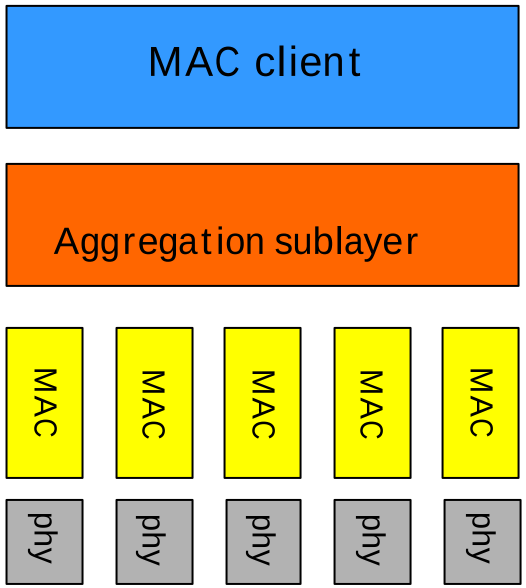

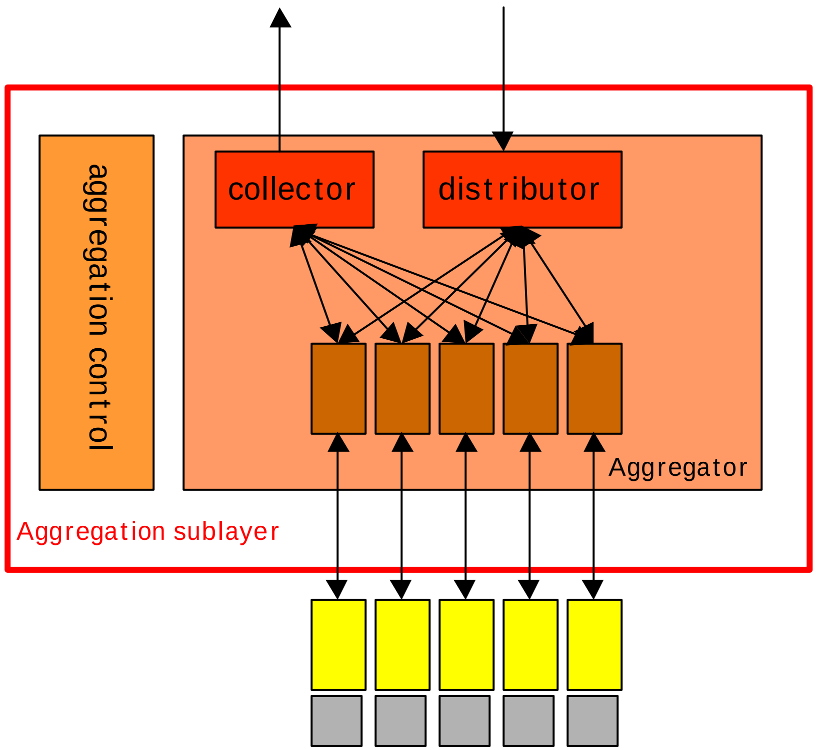

- Aggregator
  - allocates “conversations” to MAC:s
  - runs “marker protocol”
- Aggregation control
  - controls aggregation of links into multilink

Link aggregation in RPR

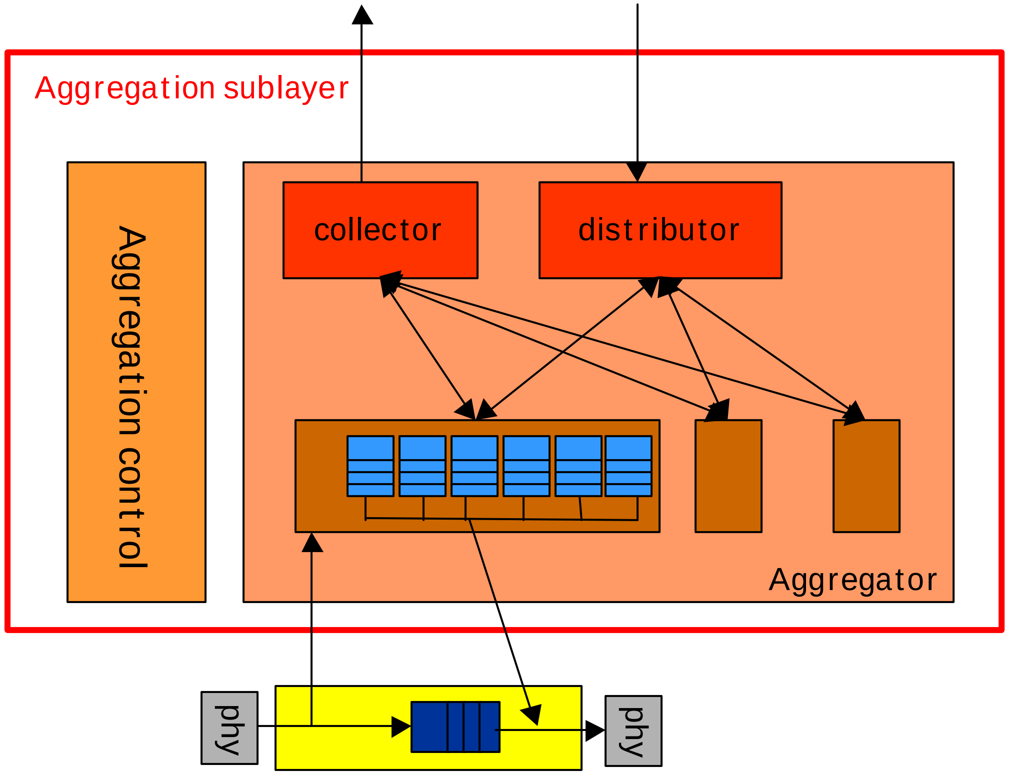

- Distributor--selects MAC and VOQ based on:
  - destination
  - resources needed by traffic aggregate ("conversation")
  - resources available on ringlet
- Main differences:
  - 802.3
    - only point to point links
    - links connect same systems
  - RPR
    - add-drop links
    - links may connect different systems

Asymmetric capacity of segments

- Use rings to which not all nodes are connected
  - Speed on individual rings is the same over all segments
  - Aggregated capacity of segments could differ
  - nodes can communicate (w/o router/ bridge hop) if they share ringlet
- Implications:
  - need boot process to detect topology and nodes of each ring
    - se Frederic Thepot's presentation
  - need mechanism to route traffic to suitable ring
    - as described earlier

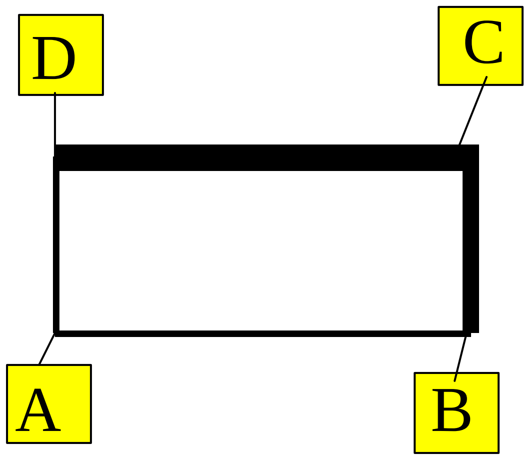

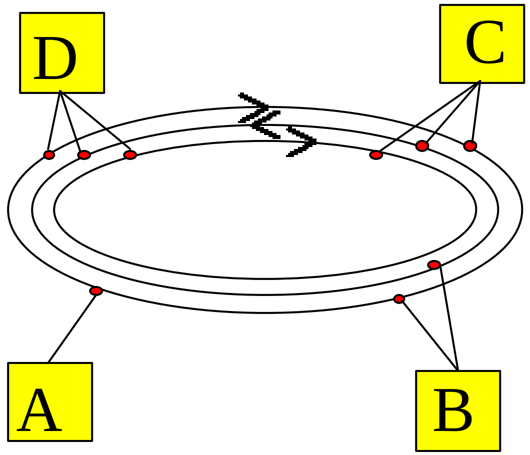

Proposal:

PRP MAC should not preclude aggregation of multiple unidirectional independently operating multiple rings.