Distributed and Autonomic Architecture for Real-Time Traffic Analysis

Cristian Morariu

Communication Systems Group, CSG
Department of Informatics IFI, University of Zurich
Switzerland

Outline

- Introduction and Motivation
- Distributed Architecture for Real-Time Traffic Analysis
- Applications
- Conclusions

Introduction and Motivation

- Link speeds double every year
- DRAM speed increases 7-9% every year
- → Main problem with IP traffic measurement: scalability

- Solution: packet sampling and flow sampling
  - Problem: measurement accuracy
    - lower sampling rate → lower acuracy, few memory
- *IP Flow*: unidirectional stream of data between two endpoints
- *Flow Keys*: IP header fields that define an IP flow
- *FlowID*: result of a hash function applied on the flow keys of an IP packet
- *FlowID space*: whole range of possible *FlowIDs*

Use Case

- Traditional approach:
  - Traffic increase → replace traffic analysis hardware

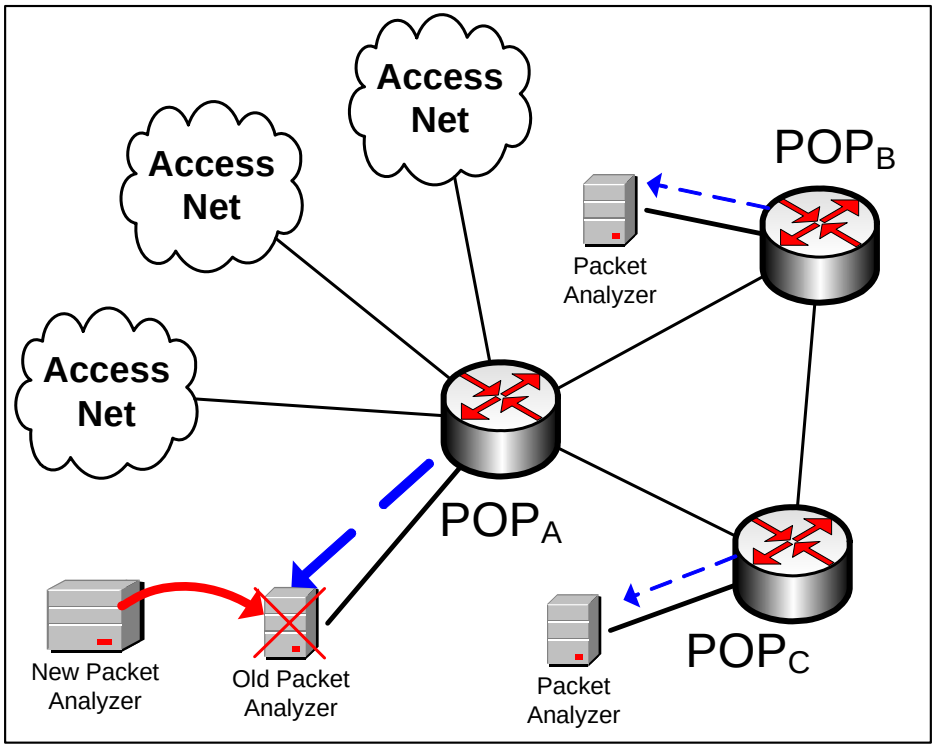

- Distributed approach:
  - Traffic increase → make use of available resources

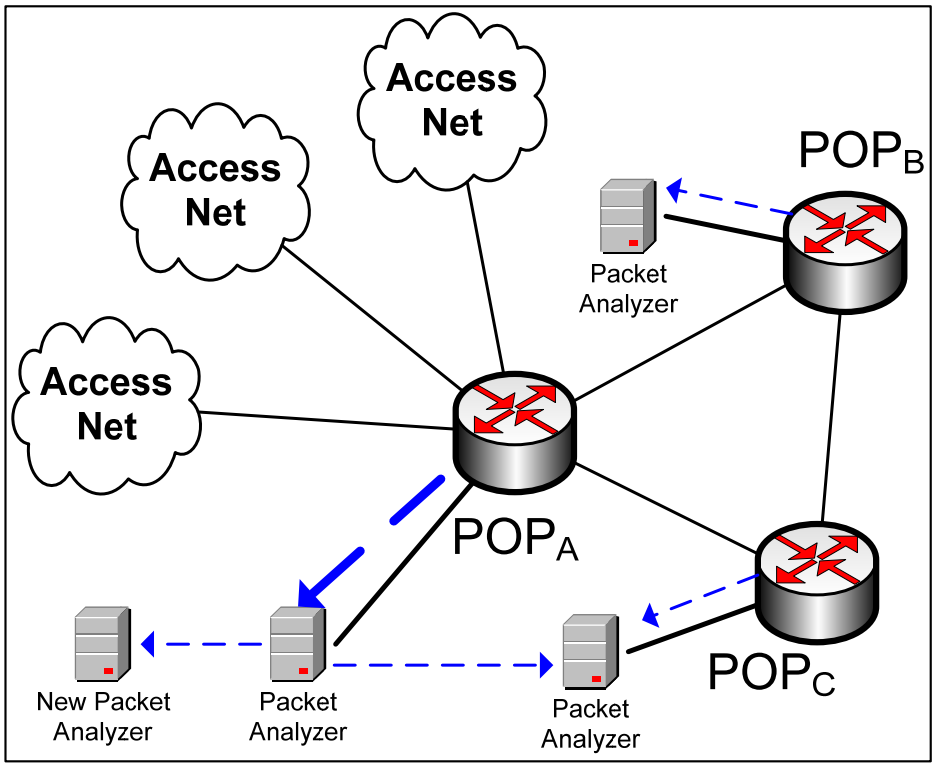

Approach

- Distribute all the tasks of traffic analysis using P2P mechanisms
  - Packet capturing
  - Packet analysis
  - Data storage

- Have self-configuration mechanisms at all layers

General Architecture

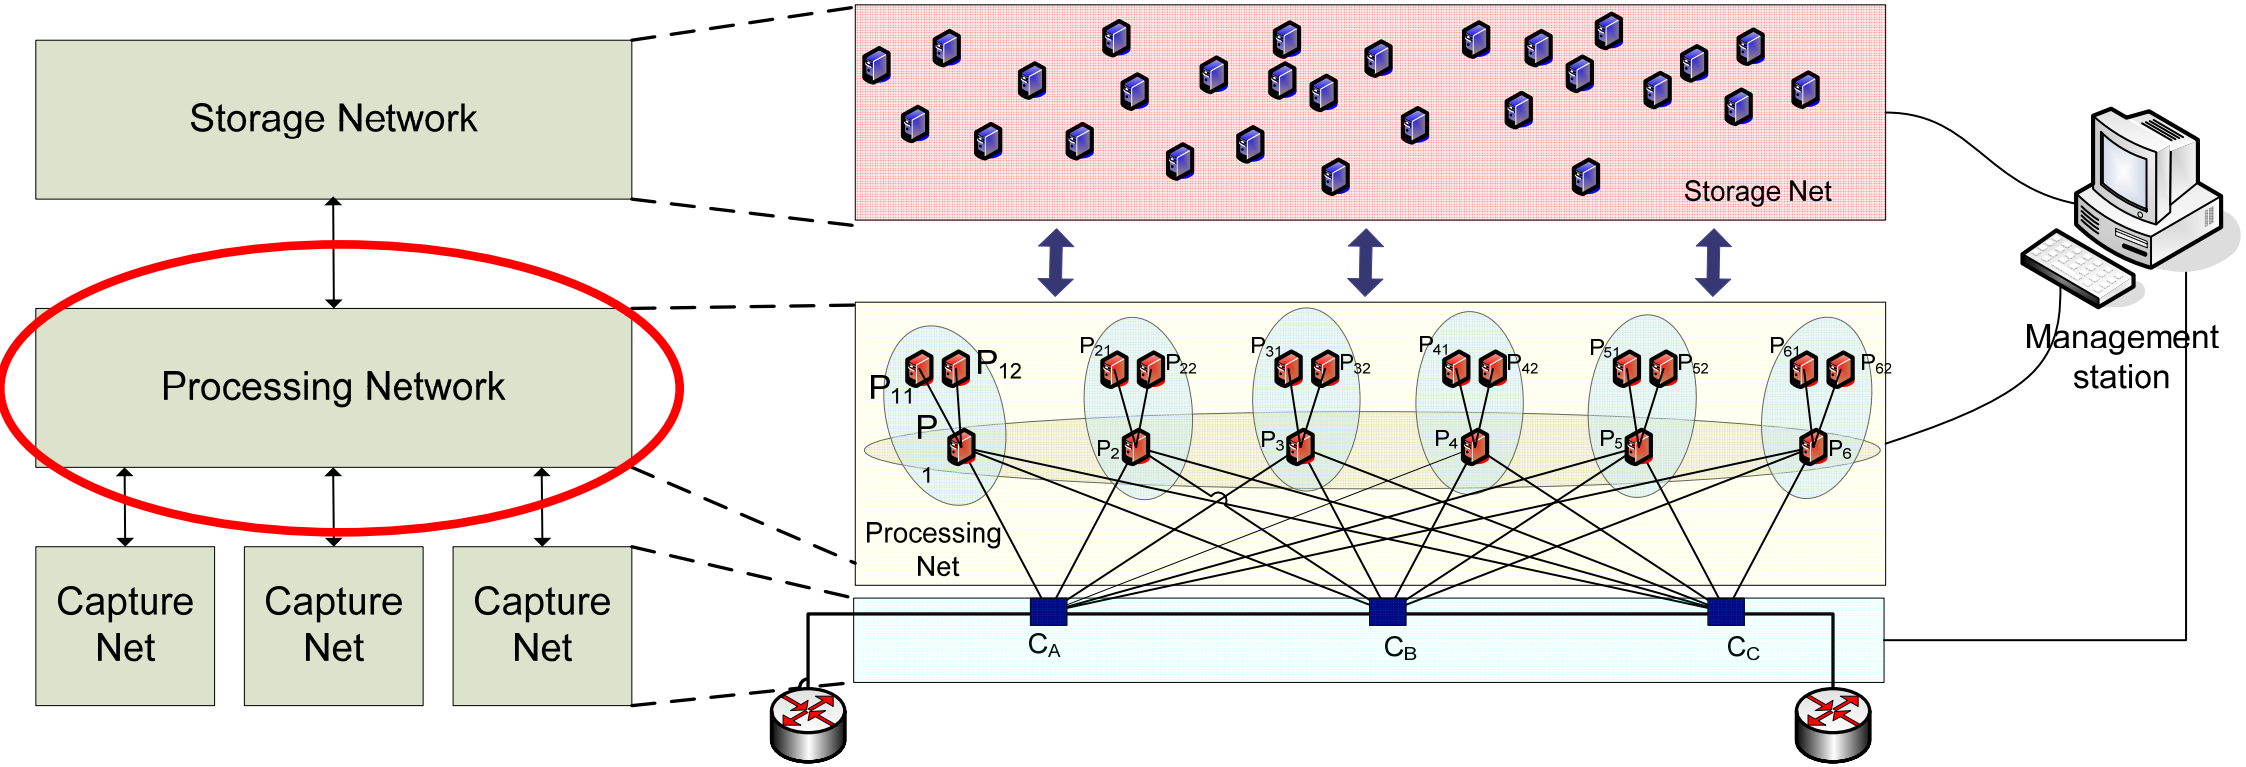

C_A, C_B, C_C: Capturing Nodes
P_X, P_{XY}: Analyzer nodes

Processing Layer

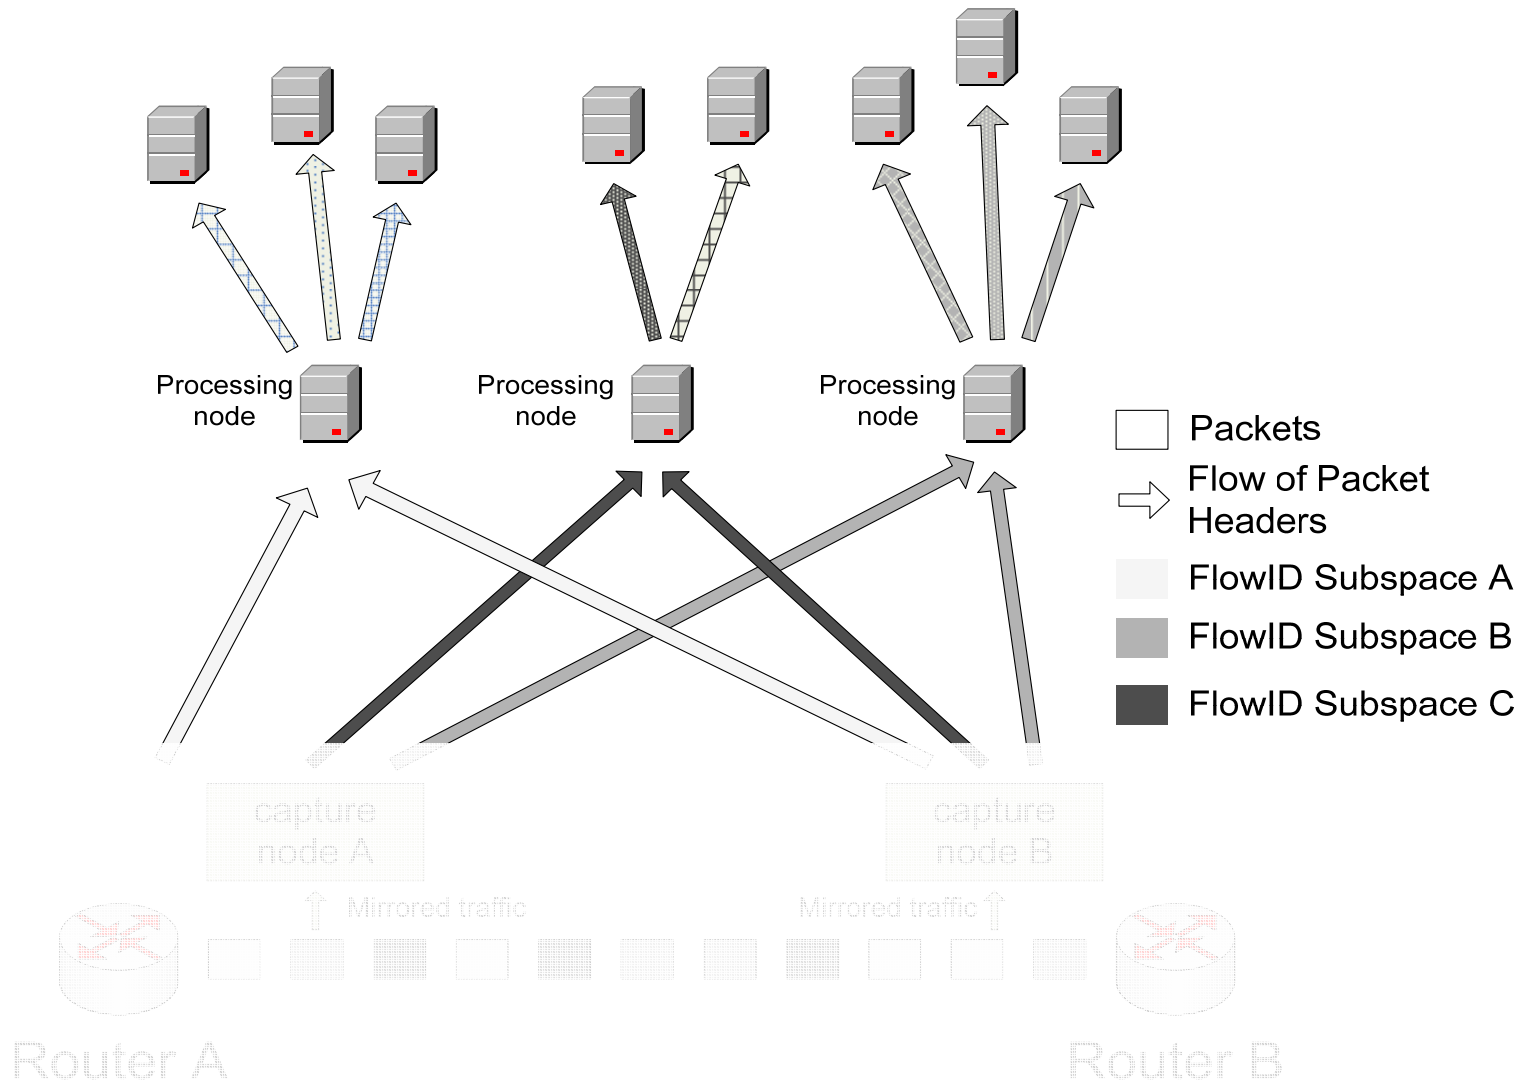

Concepts

- Split the analysis responsibility based on flow ID
- Maintain the same sampling rate on the whole Flow ID space
- Permanent monitoring of processing performance achieved in all nodes
- Workload balance

Processing Network

- Organized as a P2P overlay network
- Logical hierarchical organization
  - Different organization strategies shall be analyzed
- Every subtree represents a continuous flow ID subspace
- Each processing node is responsible with a part of the flow ID space
- Depending on the node's position it may delegate some work to other nodes

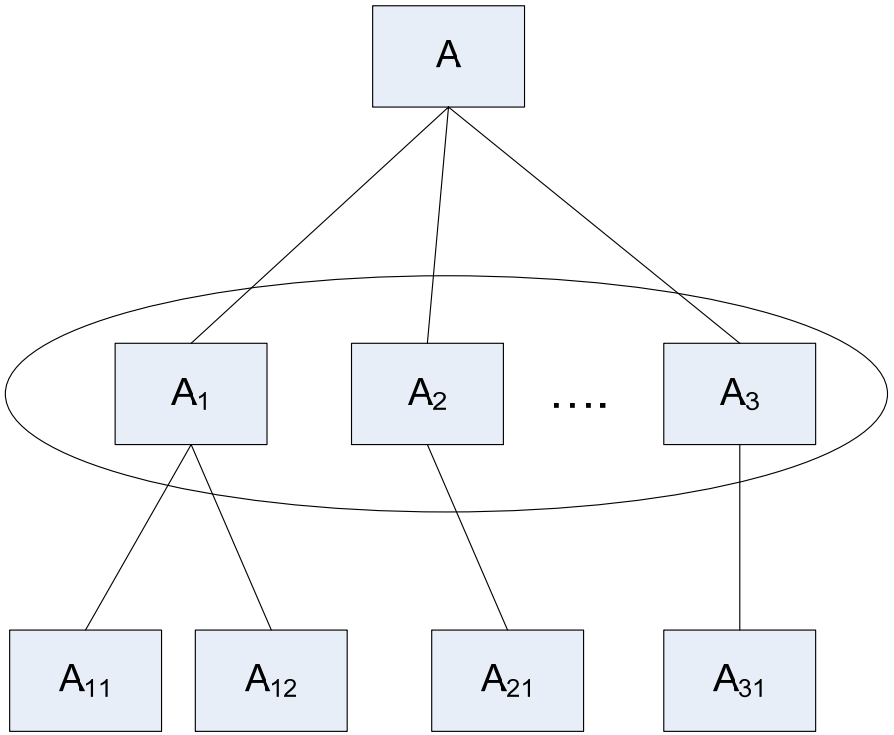

Load Balancing

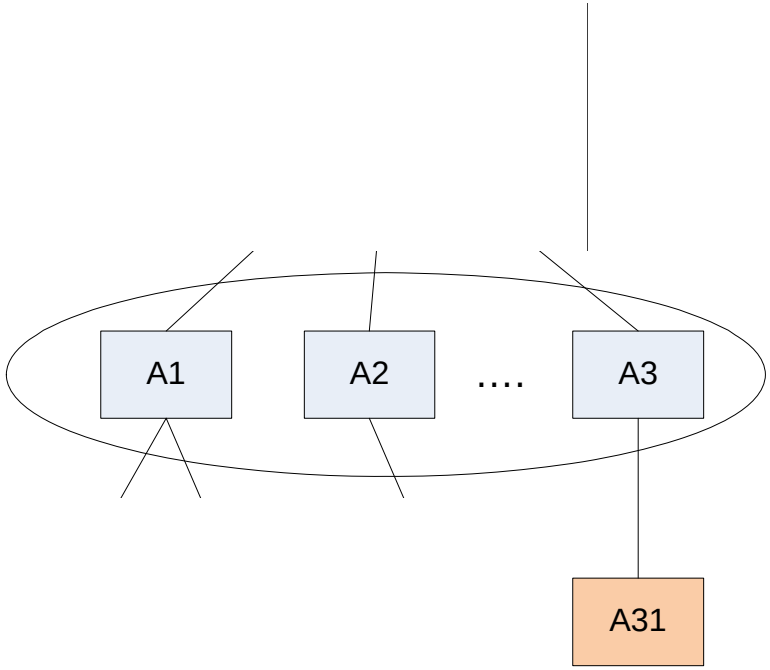

Horizontal Load Balance

- Shifting of flow ID subspace

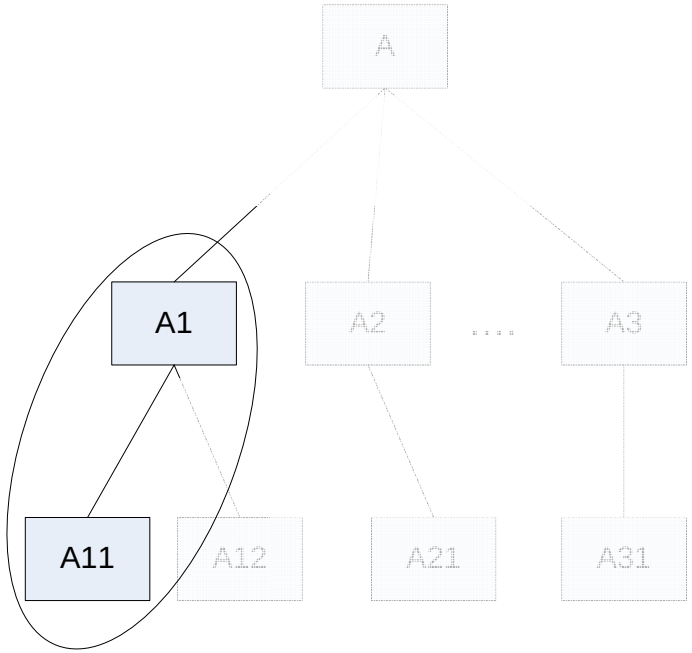

Vertical Load Balance

- Changing places

Applications

- Distributed storage of IP flow records (DIPStorage)
- IP flow accounting
- Multi-point delay measurements
- Asymmetric route detection

Distributed Storage of IP Flow Records

- Motivation:
  - Scalable storage for IP flow records
  - Fast query response
- Idea: Store the flow records within the processing network.
- Flow records are routed to the appropriate node based on their *FlowID*.
- Tank: A set of nodes that have the same set of routing rules.

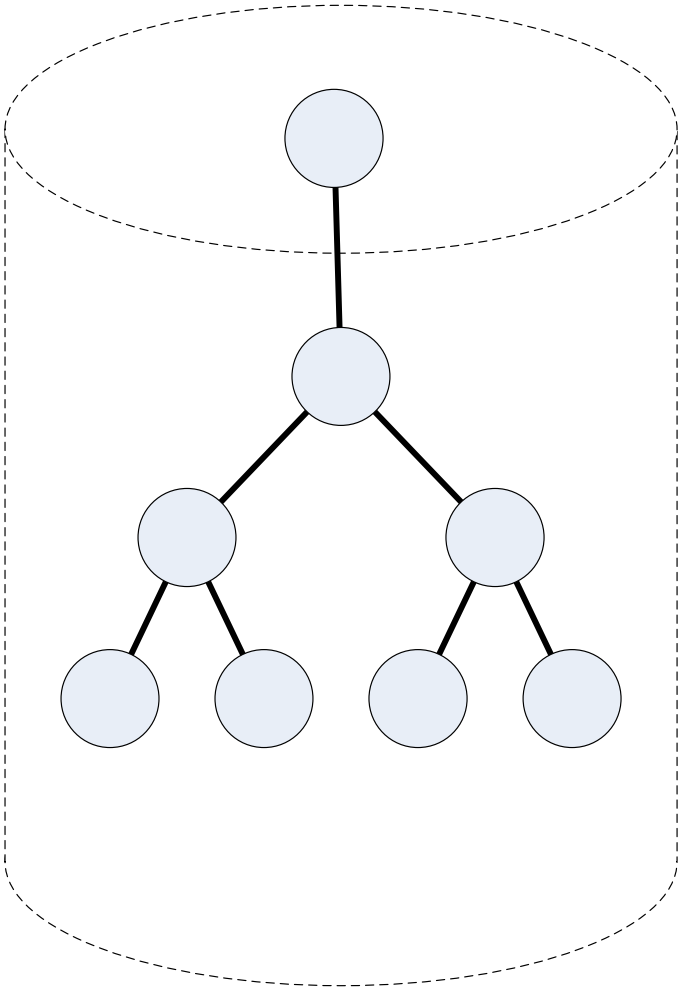

Tank

DIPStorage Architecture

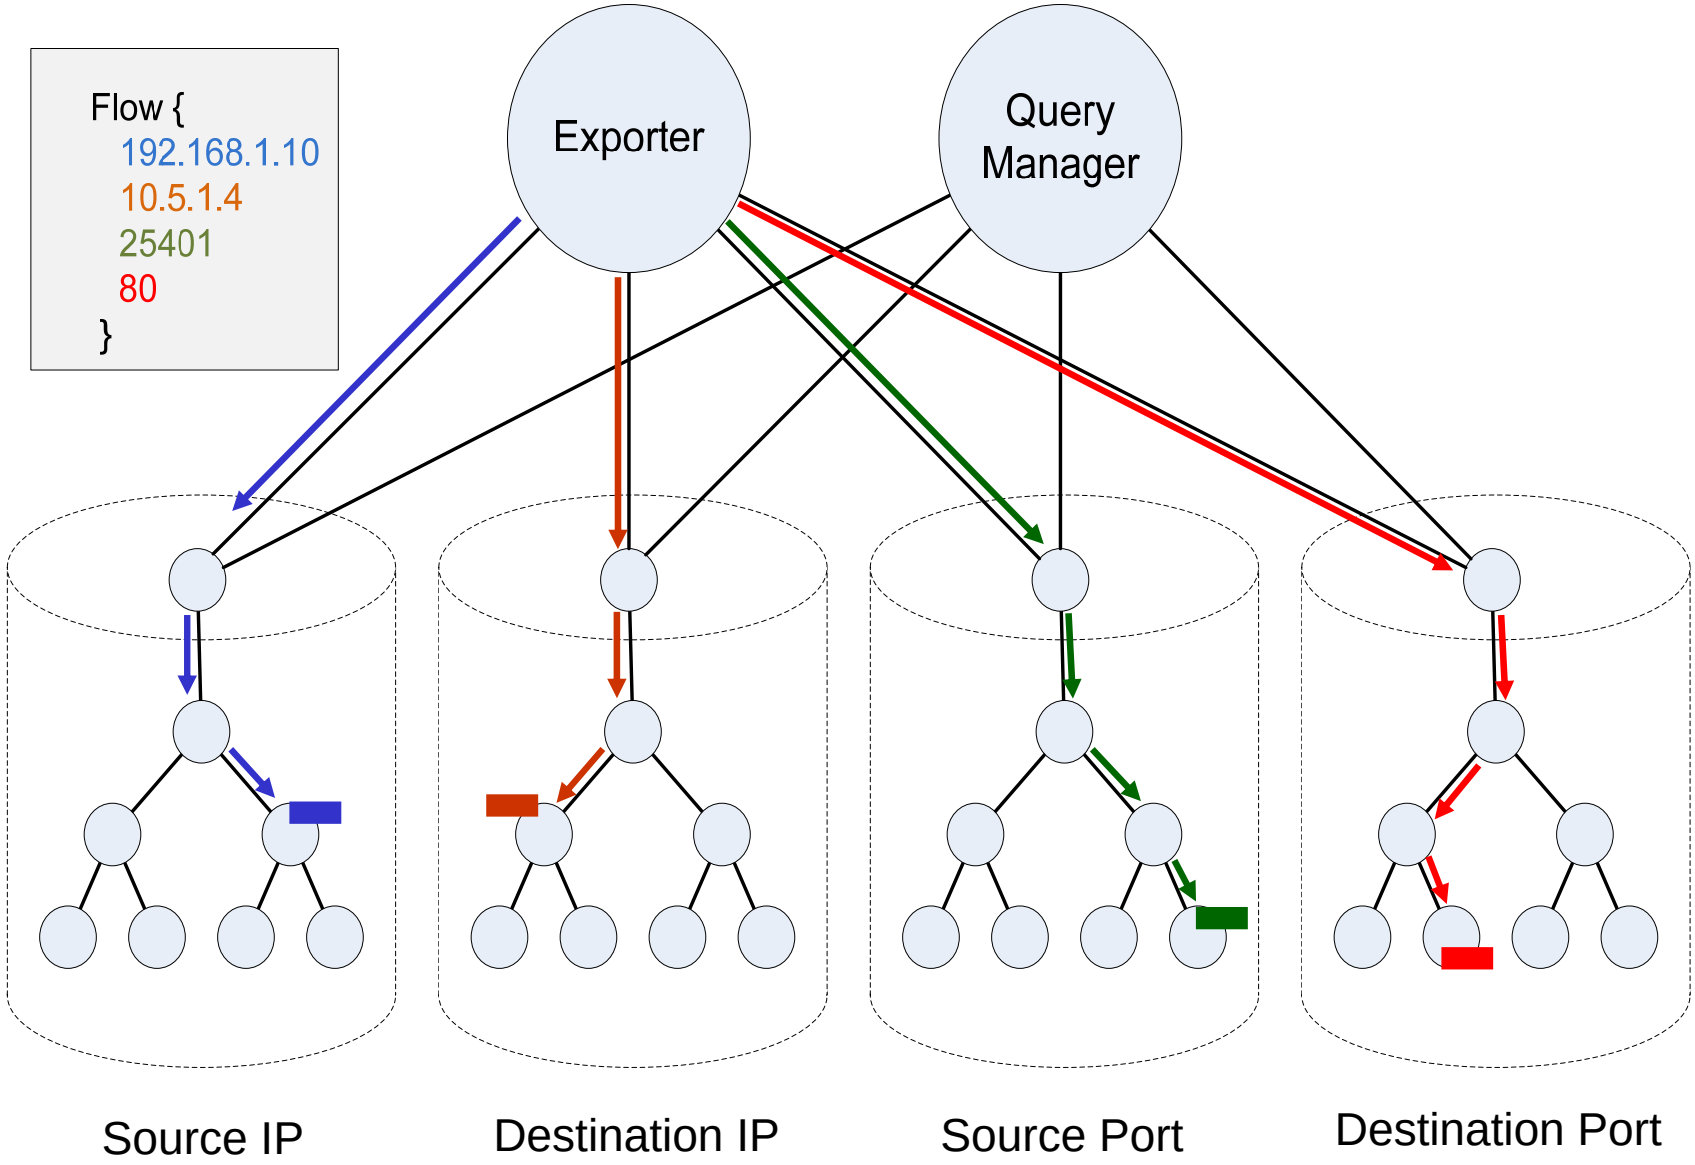

DipStorage Evaluation

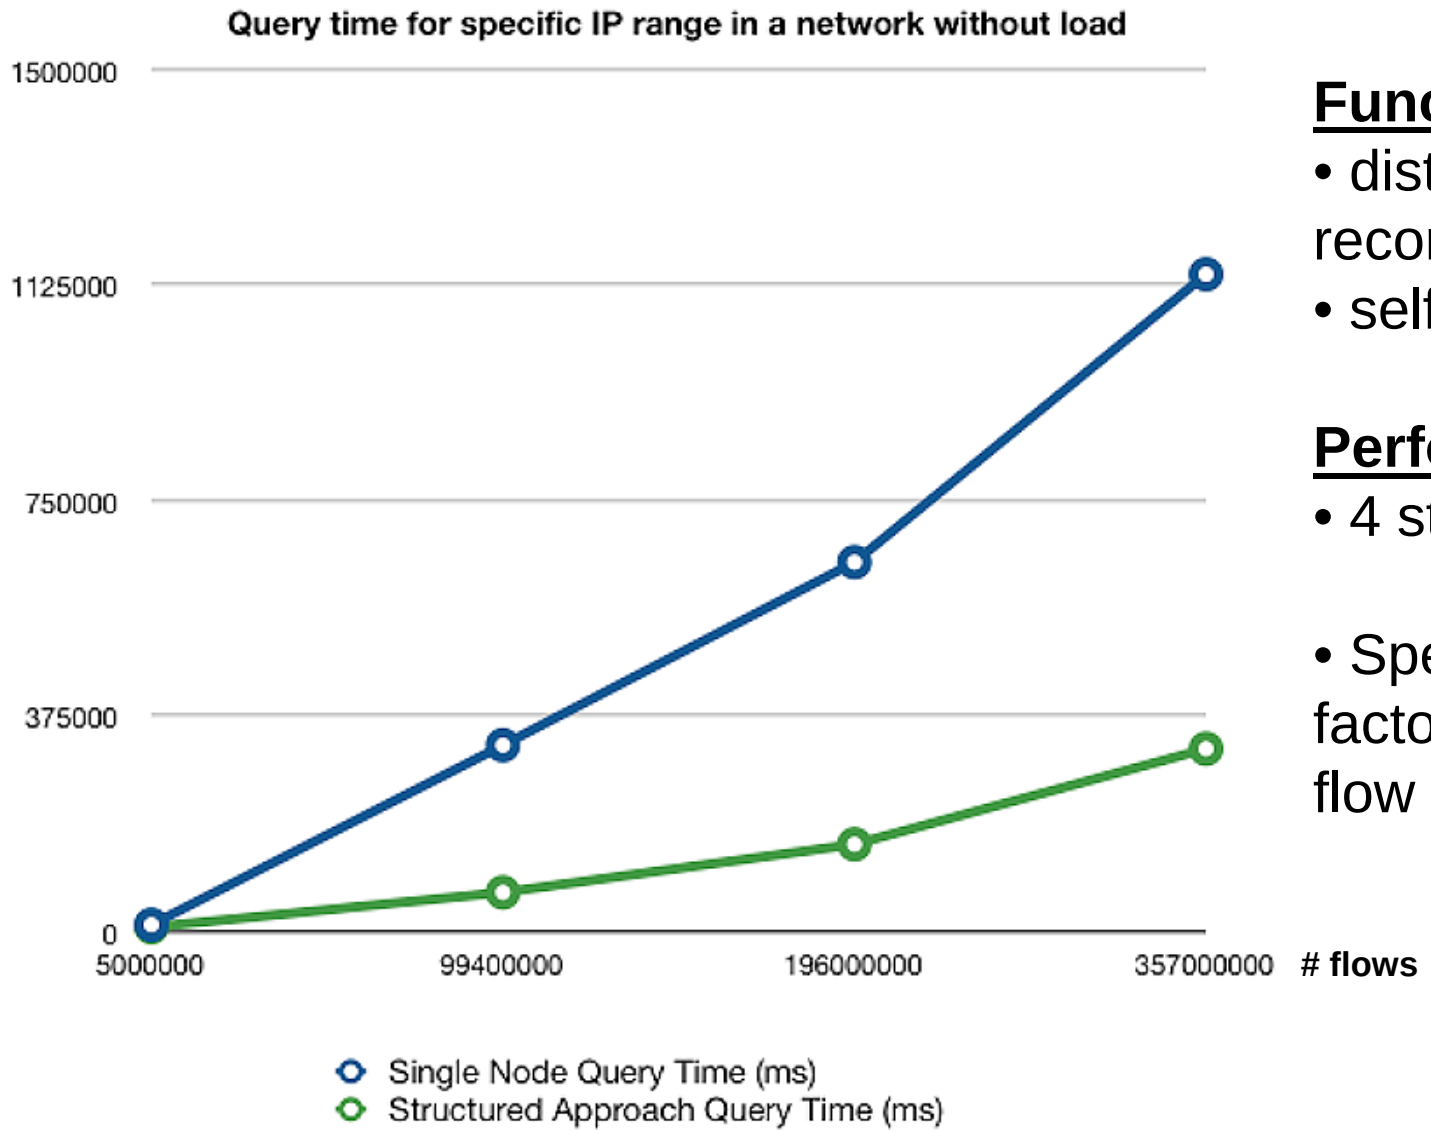

Functional Evaluation

- distribution of IP flow records
- self organizing

Performance evaluation

- 4 storage nodes used

- Speedup increase by a factor of 4 at high number of flow records

Applications: IP Flow Accounting

- Motivation: IP flow accounting requires a stateful process for every active flow
  - Expensive overhead
  - For each incoming packet header:
    - Lookup flow record
    - Update counters
  - If distributed:
    - each node has fewer active flows → faster lookup
    - Multiple lookups may be done in parallel

Applications: Multi-Point Delay Calculation

Problem: Calculate delays for each packet in blue and orange flow

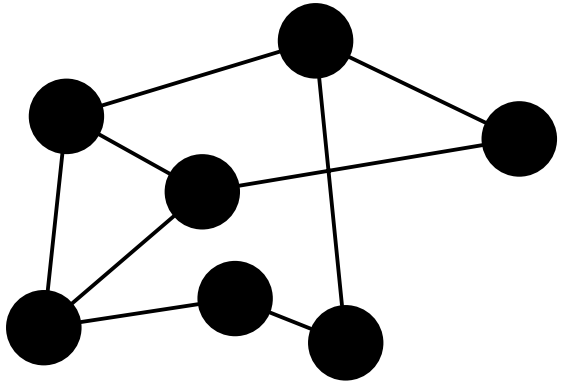

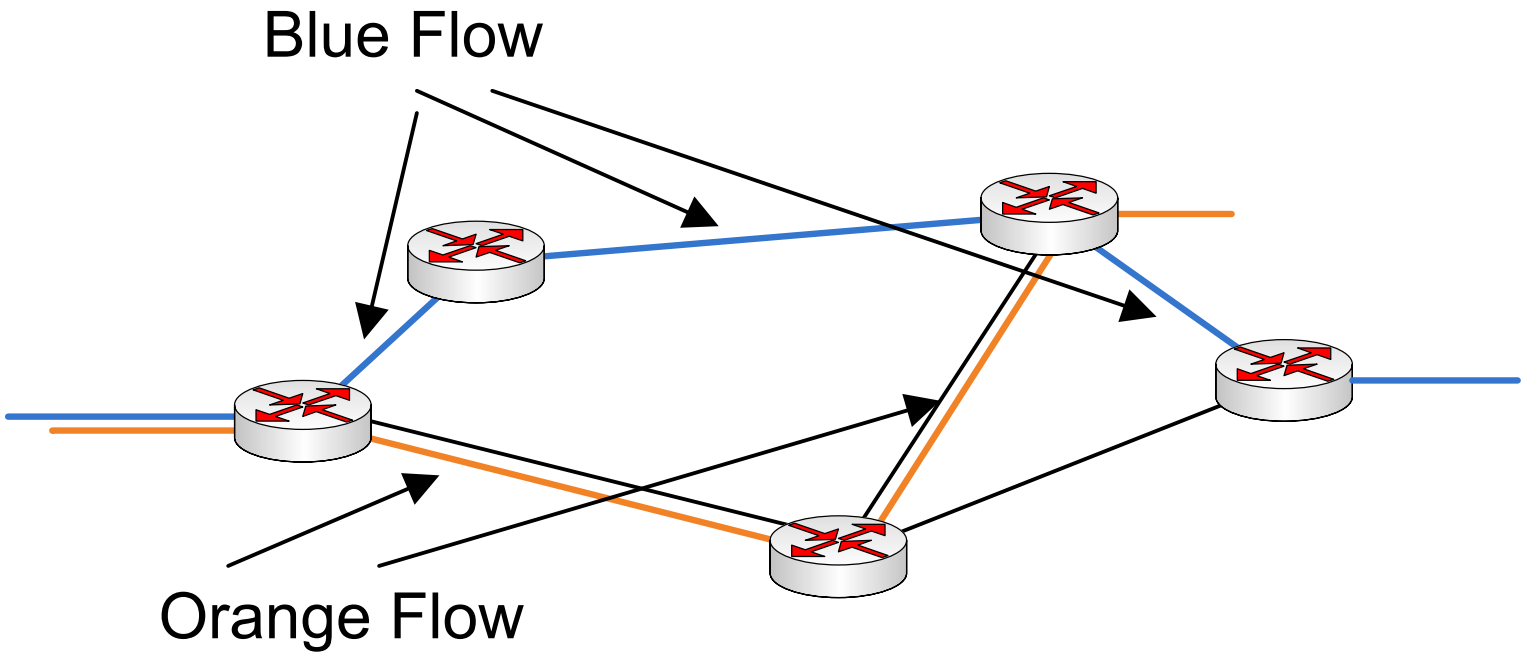

Applications: Multi-Point Delay Calculation (1)

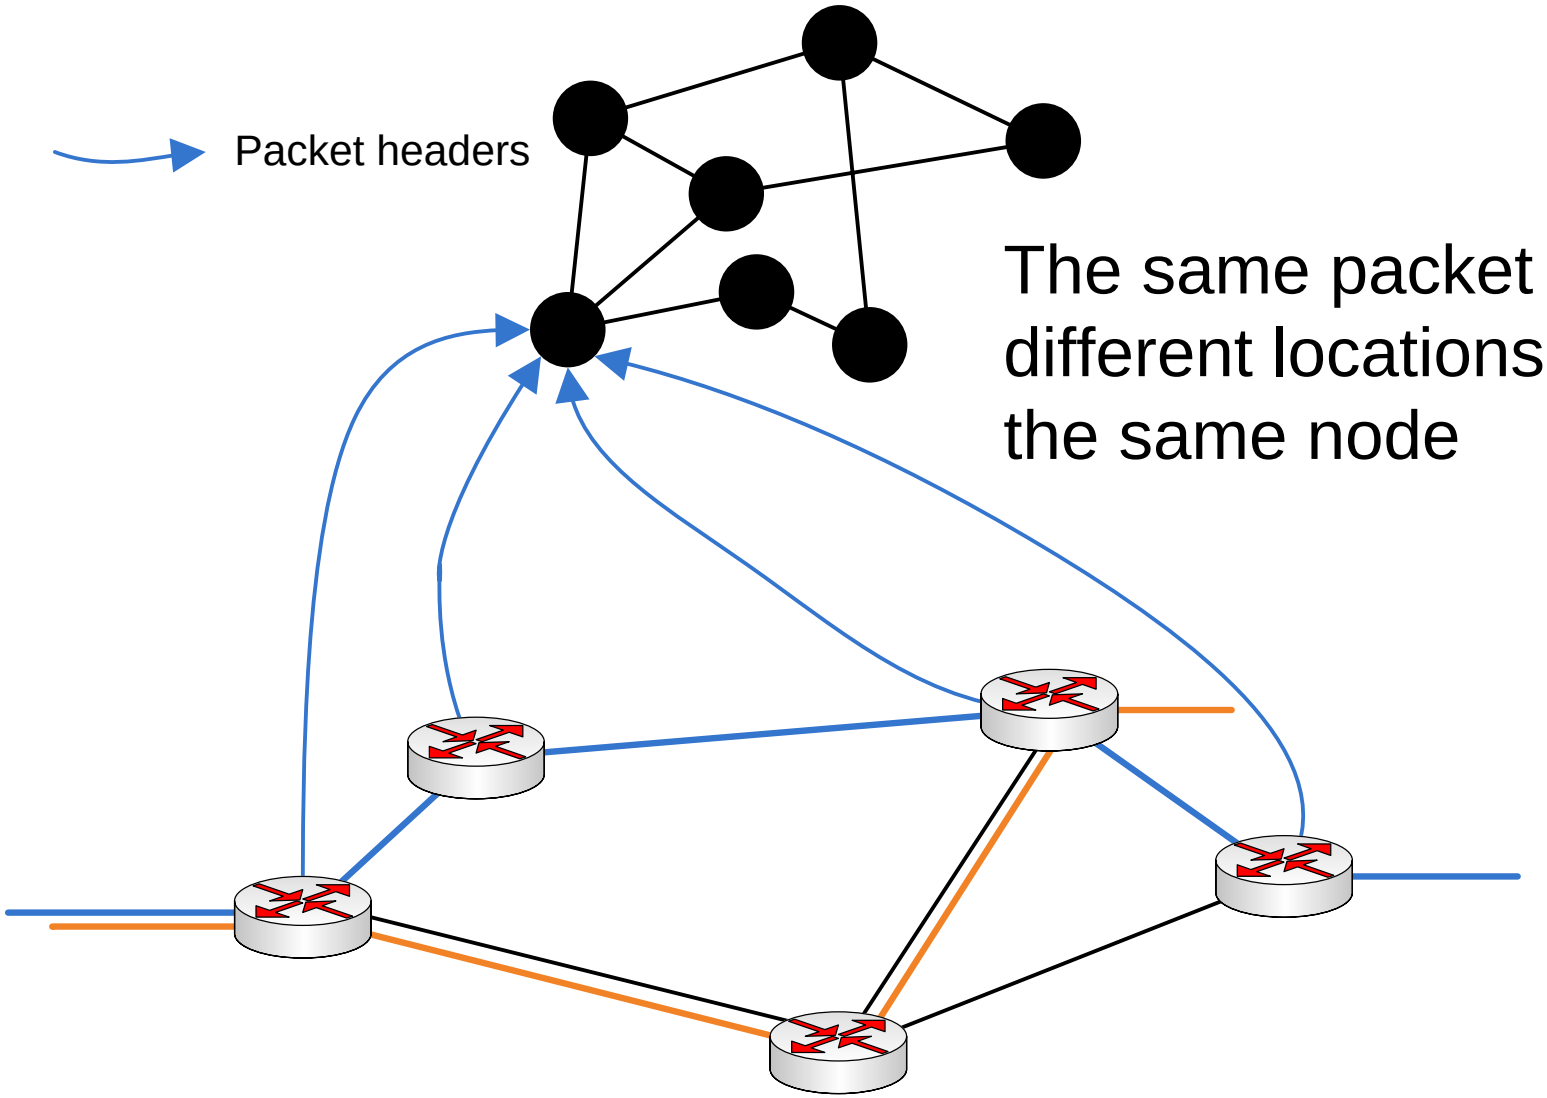

The same packet captured at different locations is routed to the same node

Applications: Multi-Point Delay Calculation (2)

Each node in the analysis network calculates delays for a set of flows

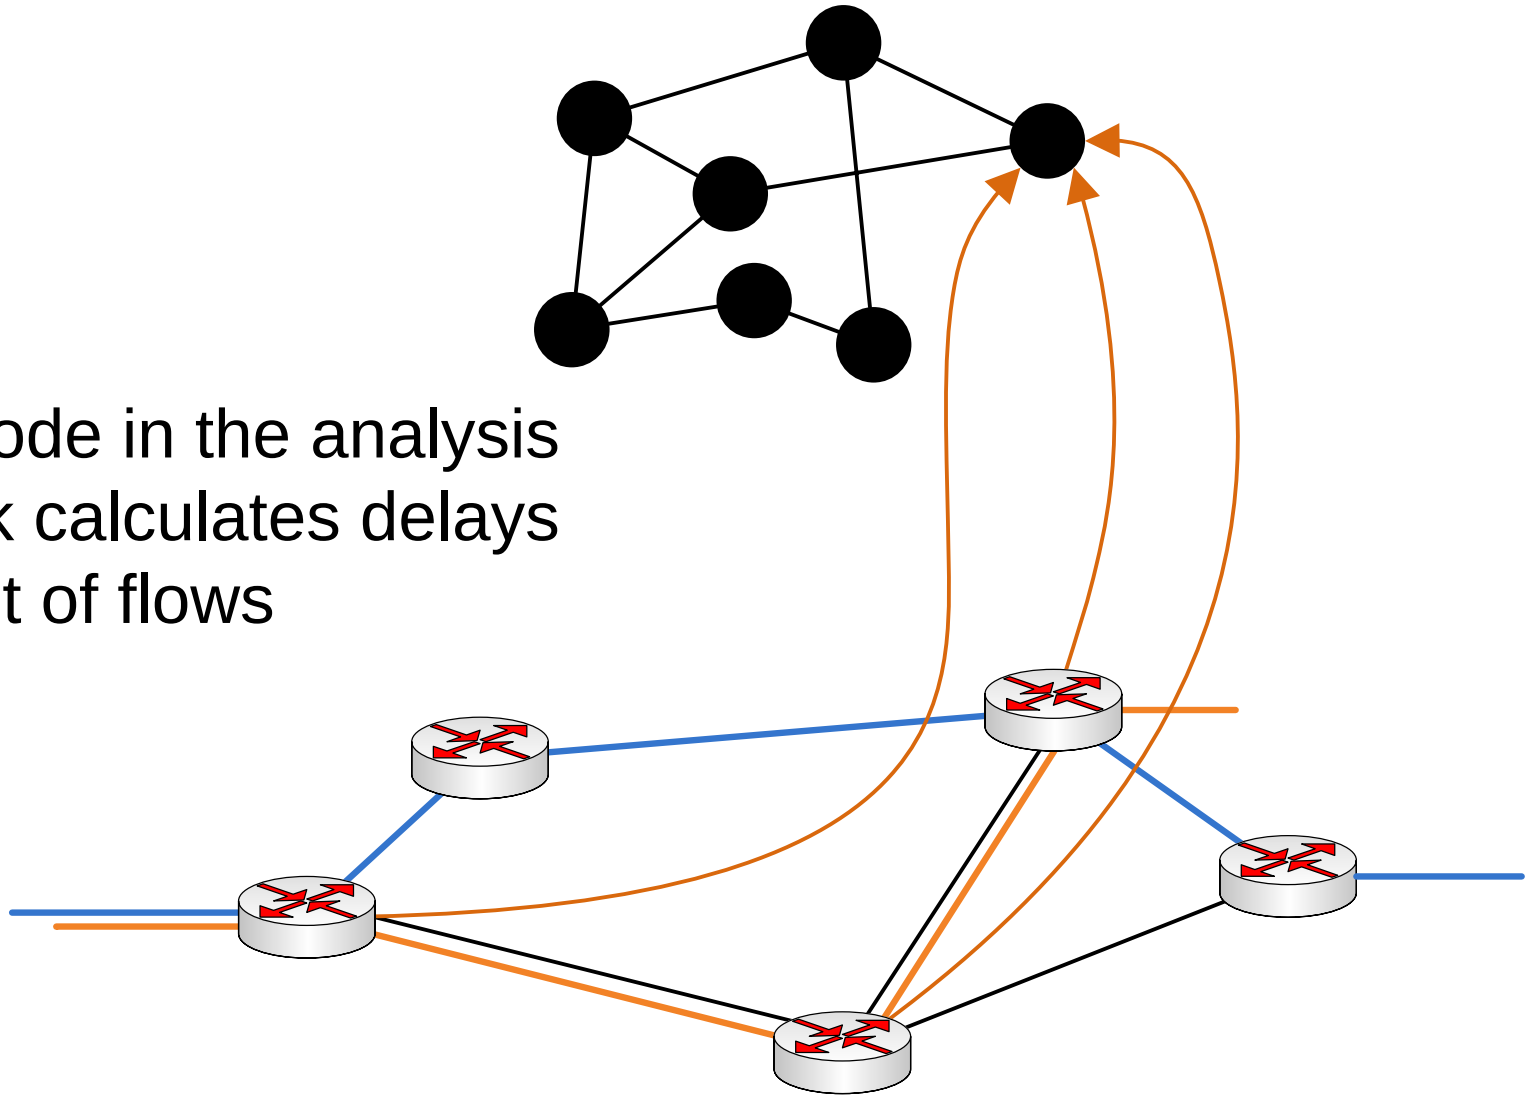

Applications: Asymmetric Route Detection

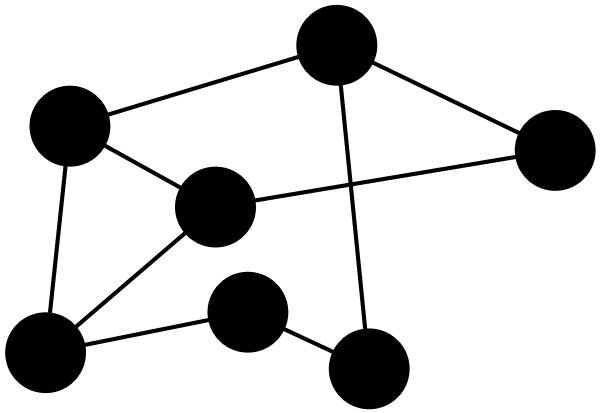

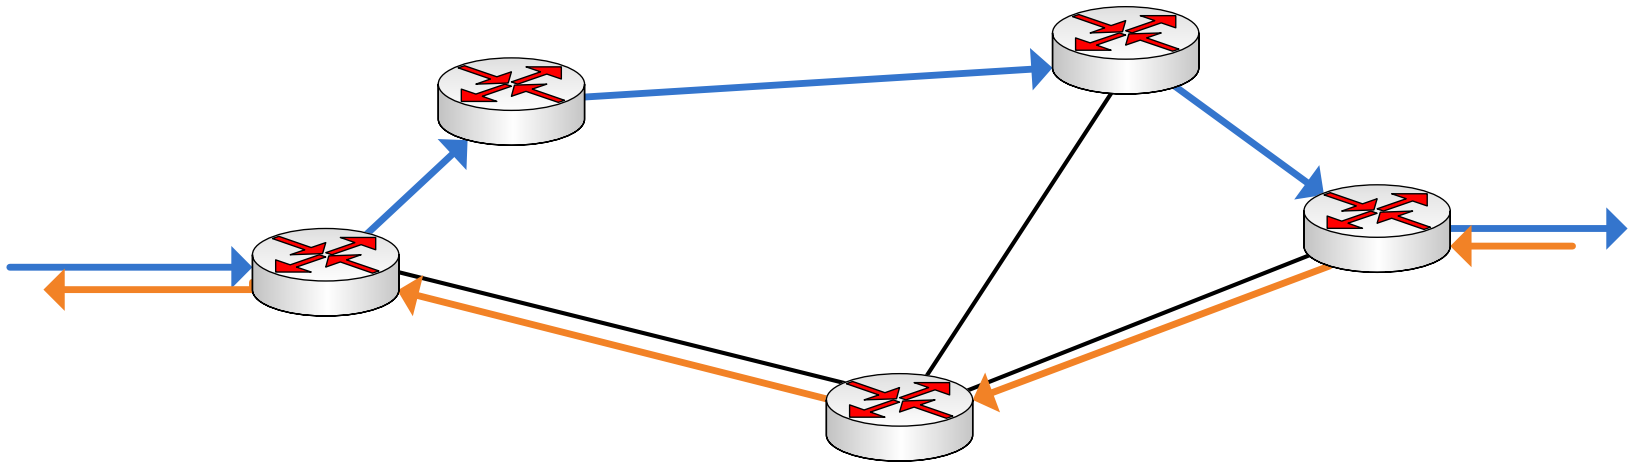

Applications: Asymmetric Route Detection

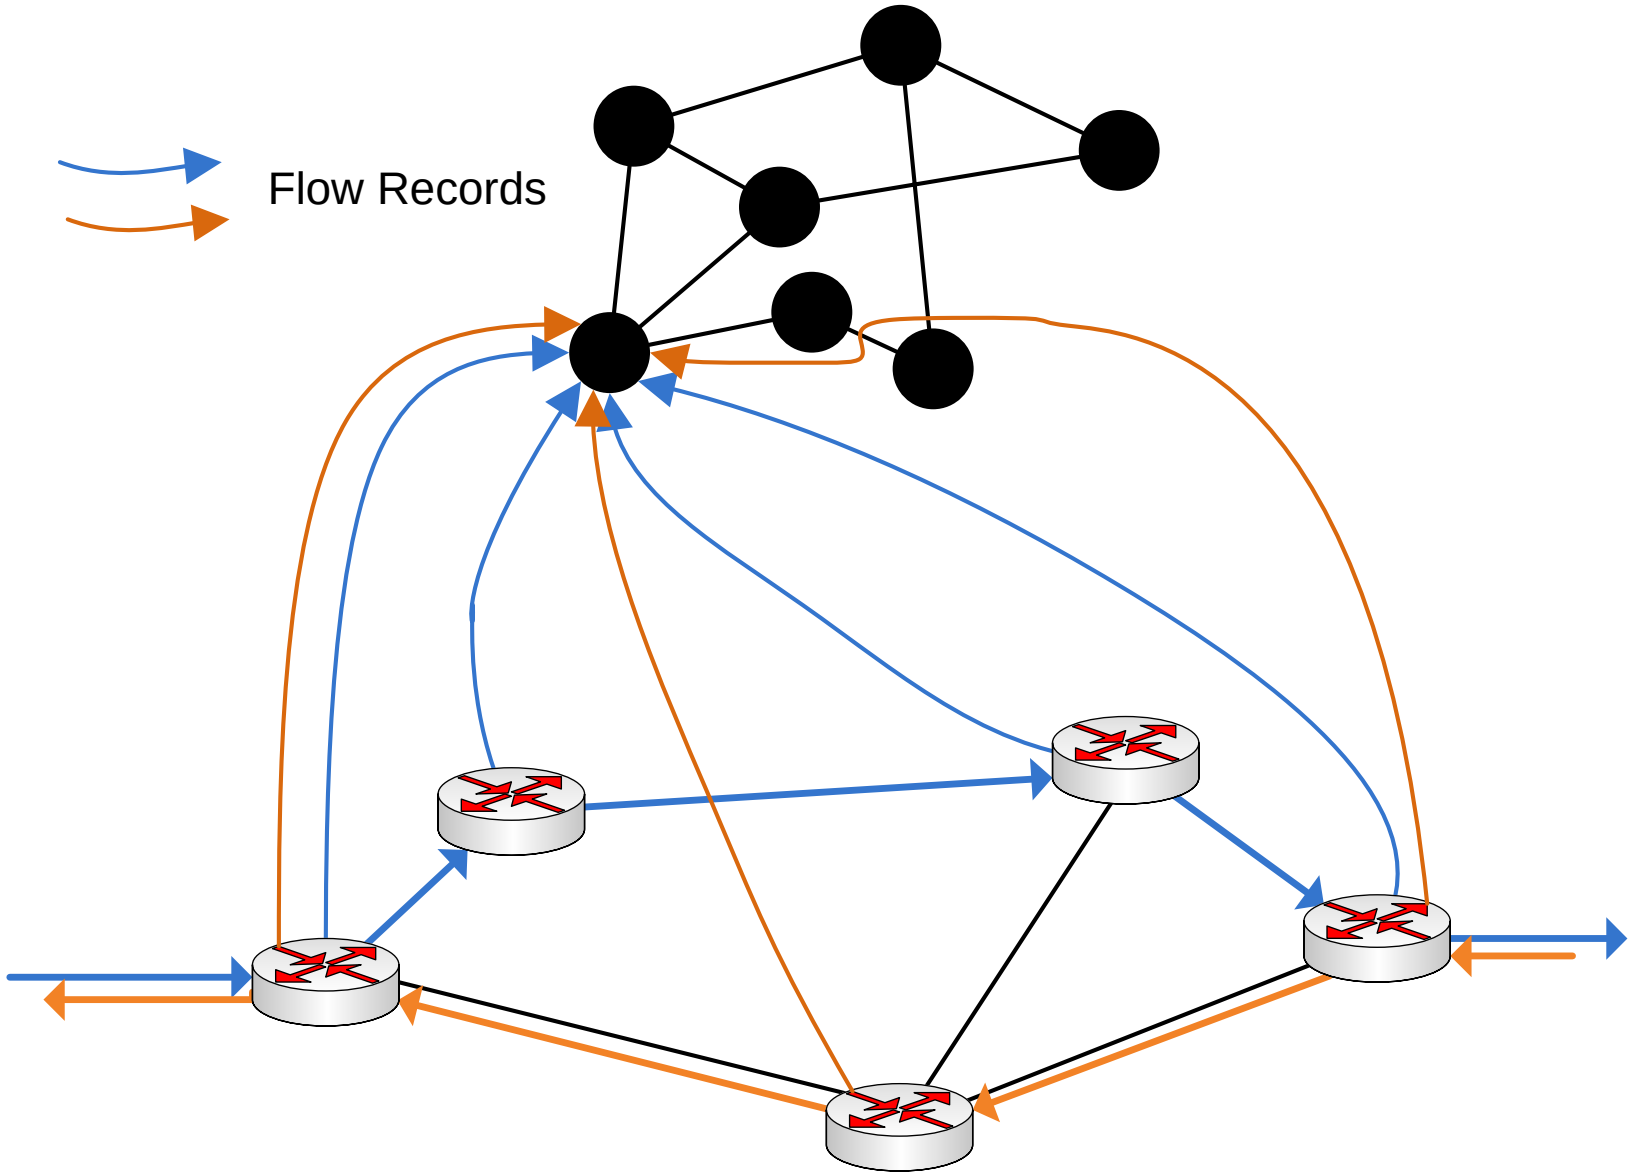

Conclusion Remarks

- Existing distributed approaches are based on “fix” configurations
- A P2P based approach was not yet investigated
- Such an approach allows:
  - Scalability
  - More accurate results by processing more data
  - Increased storage space for flow records
  - Faster query response for IP flow records repositories
  - Support for different analysis applications

Thank you

for your attention

Questions ?