Efficacy of techniques for responsiveness in a wide-area publish/subscribe system

Minkyong Kim	IBM Watson Research, NY
Kyriakos Karenos	Microsoft, UK (Postdoc at IBM)
Fan Ye	IBM Watson Research
Johnathan Reason	IBM Watson Research
Hui Lei	IBM Watson Research
Konstantin Shagin	IBM Haifa Research

R3 messaging in news

HOME PAGE | TODAY'S PAPER | VIDEO | MOST POPULAR | TIMES TOPICS

The New York Times
Thursday, December 2, 2010

Business Day
Markets

WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HEALTH | SPORTS | OPINION | ARTS

Search | Global | DealBook | Markets | Economy | Energy | Media

Boeing and IBM Research Pilot Experimental Air Traffic Management Initiatives

Published: December 1, 2010

ARMONK, N.Y., Dec. 1, 2010 /PRNewswire-FirstCall/ -- Boeing (NYSE: BA) and IBM (NYSE: IBM) announced today that they have successfully completed a pilot research project designed to show how officials from multiple organizations can have more timely, consistent and complete information to resolve fast-changing or unpredictable aviation events.

(Logo: http://photos.prnewswire.com/prnh/20090416/IBMLOGO)

The quicker availability of better information potentially can help authorities make more informed decisions to protect passenger and national security when airspace anomalies are detected, such as when an aircraft is found to be traveling off course. The project has demonstrated that advances in software can accelerate and orchestrate the flow of information from sensors and networks on a nationwide scale.

To better assess and manage these events, officials have been striving to implement "network centric operations" -- the quick exchange of information from a variety of electronic sources -- to gain better "situational awareness" of the many aspects of an aviation event. These situations require the sharing of updated and detailed information about temporary or ongoing airspace restrictions, flight plans, reports about weather and

ZDNet

News & Blogs | Reviews | Downloads | White Papers

US Edition | Companies | Hardware | Software | Mobile | Security

Between the Lines
Larry Dignan, Sam Diaz, Andrew Nusca

Mobile | RSS | Email Alerts

Comments | Share | Print | Facebook | Twitter | Recommend | Votes

Home / News & Blogs / Between the Lines

Boeing, IBM create real-time aviation analysis system

By Larry Dignan | November 30, 2010, 9:01pm PST

Summary

Boeing and IBM Research have cooked up a new technology dubbed R3 designed to better predict and manage air traffic volume as well as deliver more real-time information about an aviation event.

Boeing and IBM Research have cooked up a new technology dubbed R3 designed to better predict and manage air traffic volume as well as deliver more real-time information about an aviation event.

R3 stands for responsive, reliable and real-time and created at IBM's Watson research lab. The joint project was designed to coordinate more flights in the same airspace. The air traffic issue is a big one considering that flight traffic is expected to double or triple by 2025.

Boeing was involved with the pilot project for obvious reasons. IBM aims to take the R3 technology and deploy it in other areas ranging from emergency services to weather forecasting to power and water management.

CNNMoney.com
A Service of CNN, Fortune & Money

Symbol | Keyword

Home | Business News | Markets | Personal Finance | Retirement | Technology | Luxury | Small Busines

Boeing and IBM Research Pilot Experimental Air Traffic Management Initiatives

December 01, 2010: 12:01 AM ET

PR Newswire

Messaging middleware for wide-area

- Increased need to federate domains, which span across states, nations, or continents
- Messaging middleware to enable cross-domain, wide-area federation
- Applications requires certain Quality of Service (QoS) guarantees in message delivery

Air surveillance data distribution

- Goal: Track air crafts and alert important events
- Sensor data needs to be delivered to various agencies in timely manner

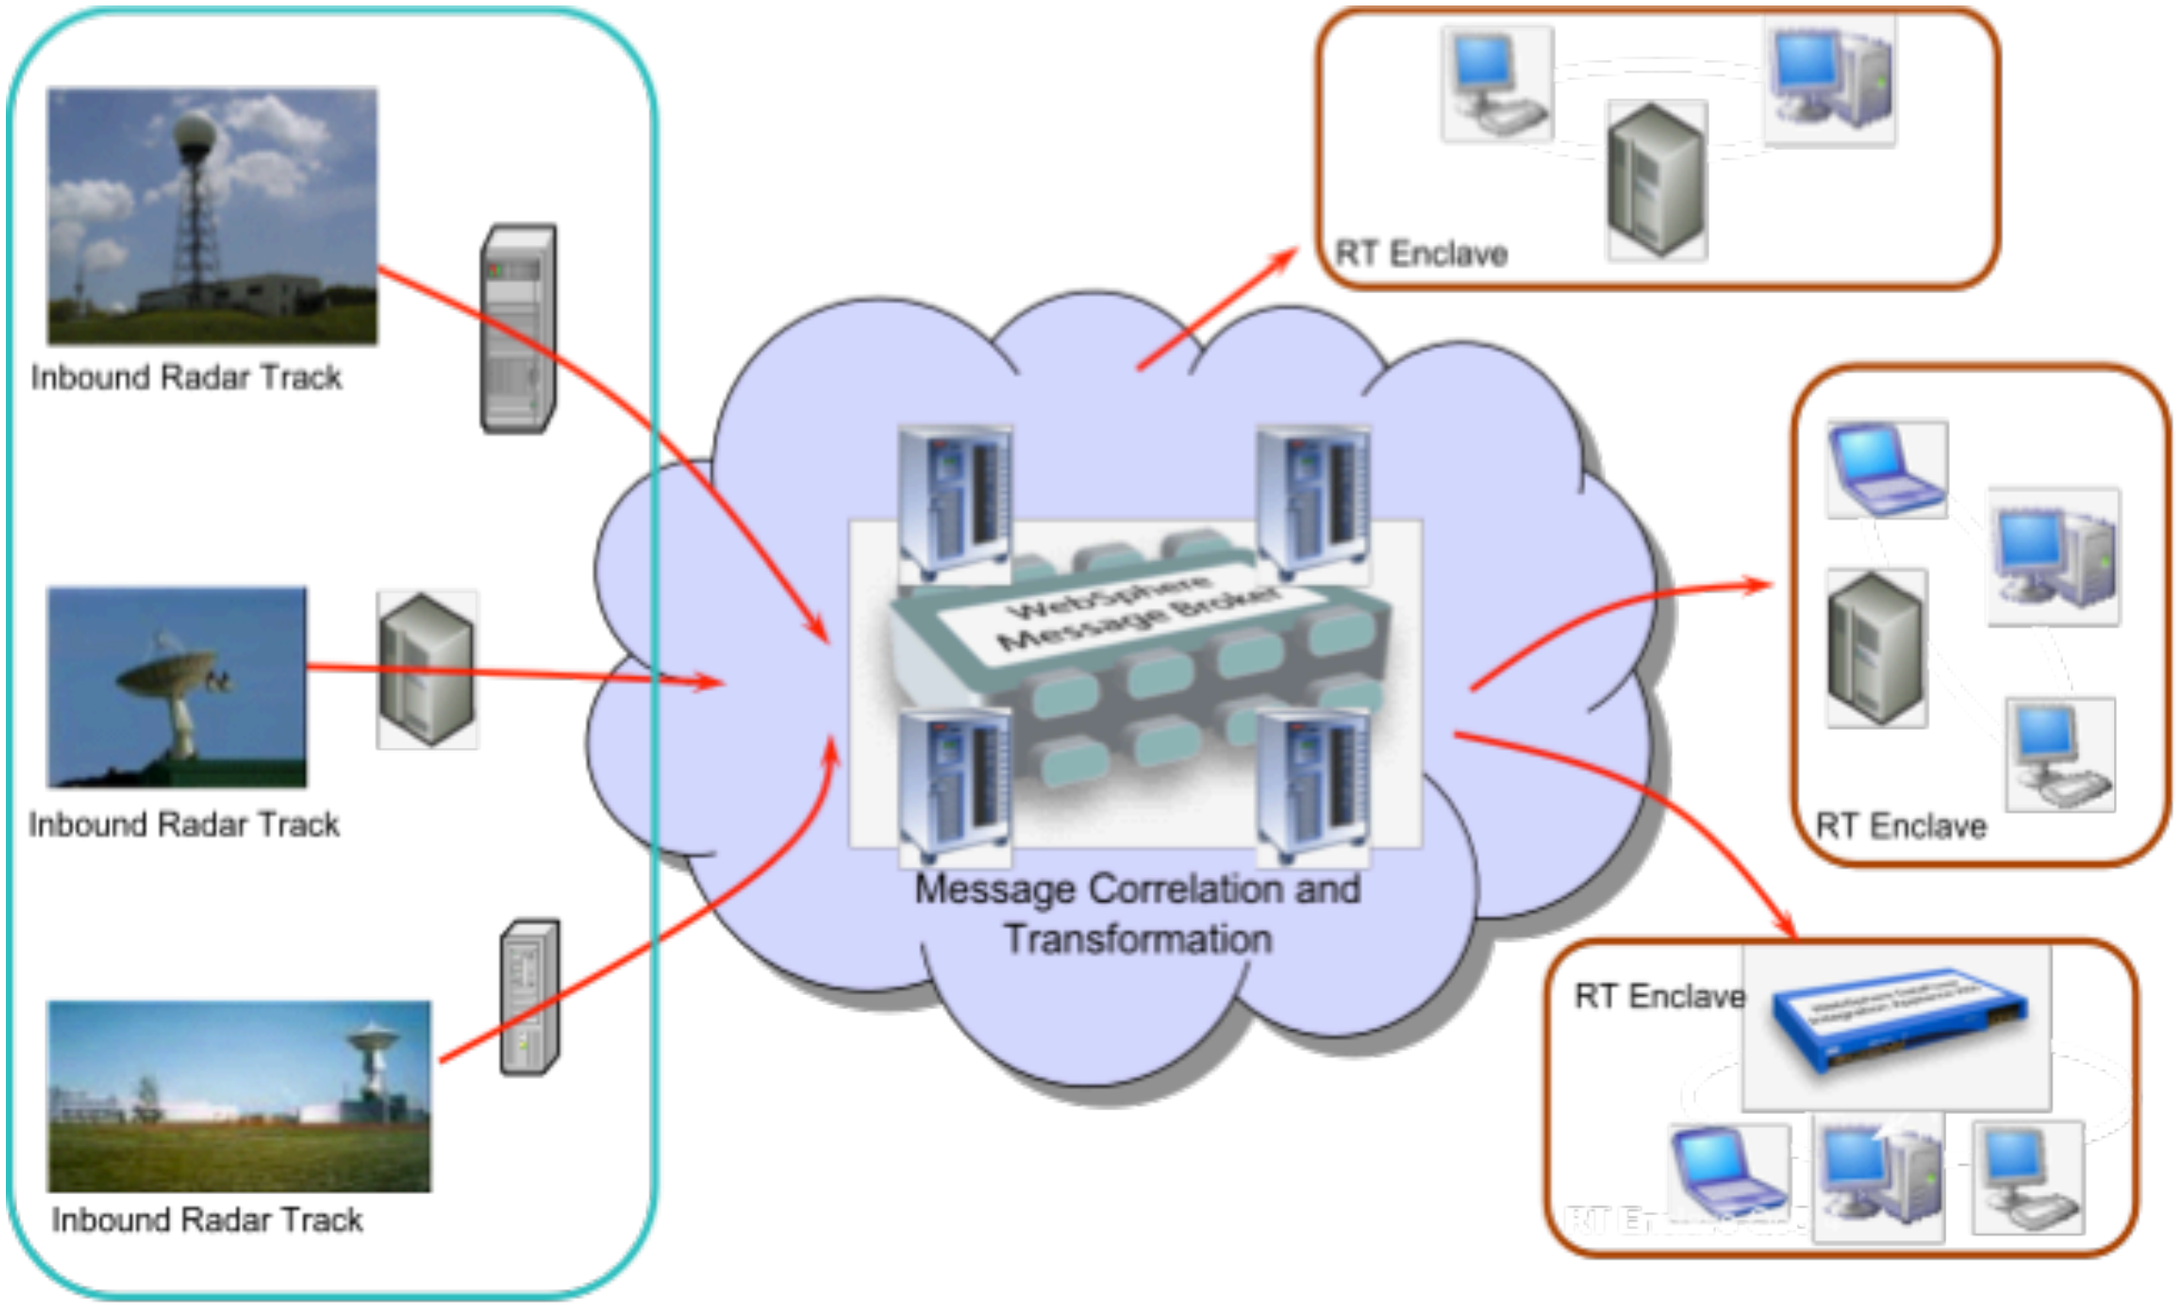

Smart electric grid communication

- Goal: Better manage electricity by providing pricing incentives
- Demand signal flows from individual houses to utilities to providers
- Pricing signal flows in the reverse direction

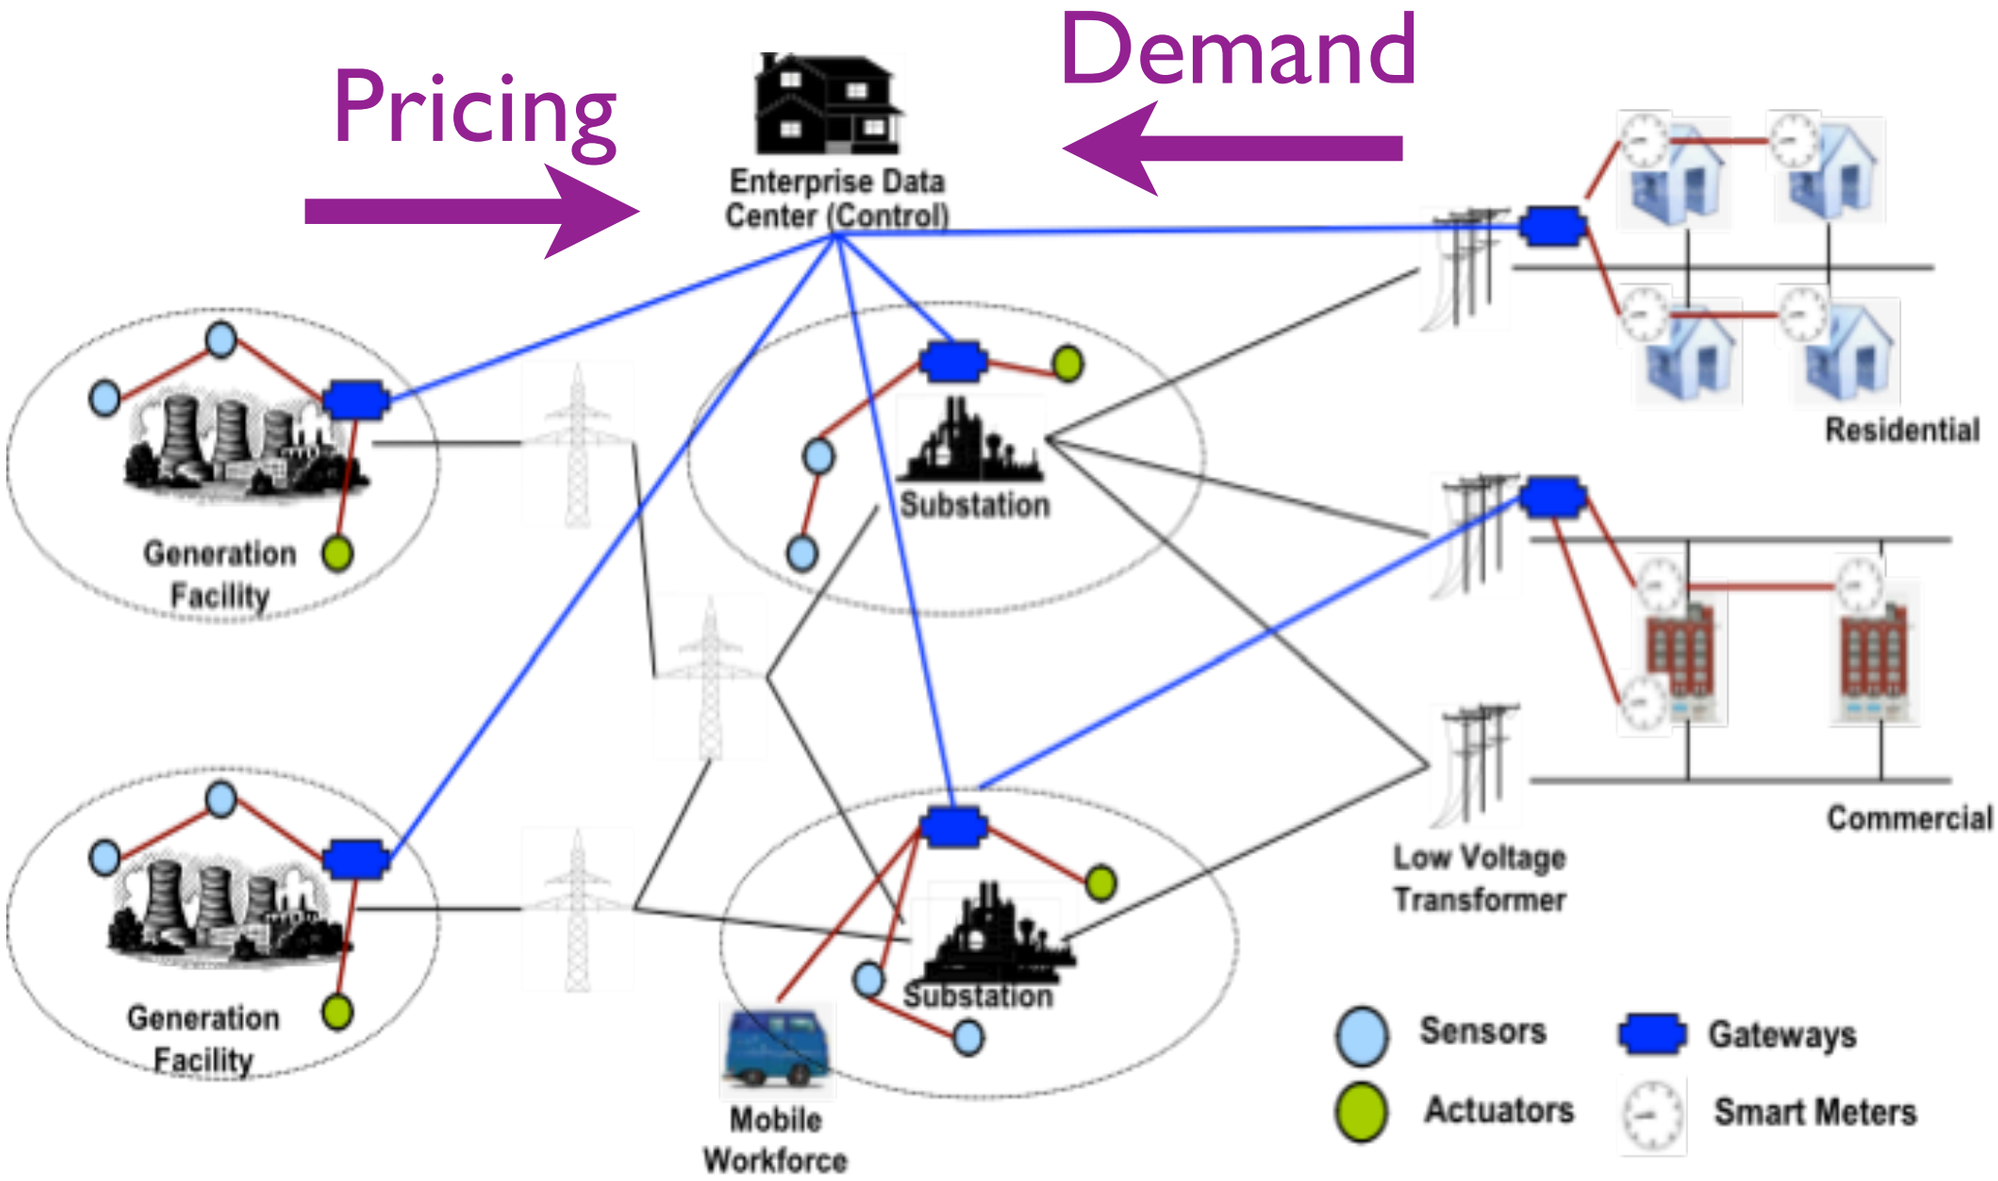

Harmony messaging middleware

- Harmony: Qos-aware topic-based messaging system for federation of domains in wide-area network environment
- Allow subscribers to specify end-to-end QoS requirement such as latency for a topic
- Paths are built on overlay network to connect publishers and subsribers
- Each hop uses pub/sub for sending/receiving messages
- *Tempore* (transport layer) alerts QoS violation on individual hop
- Harmony manages QoS over end-to-end path

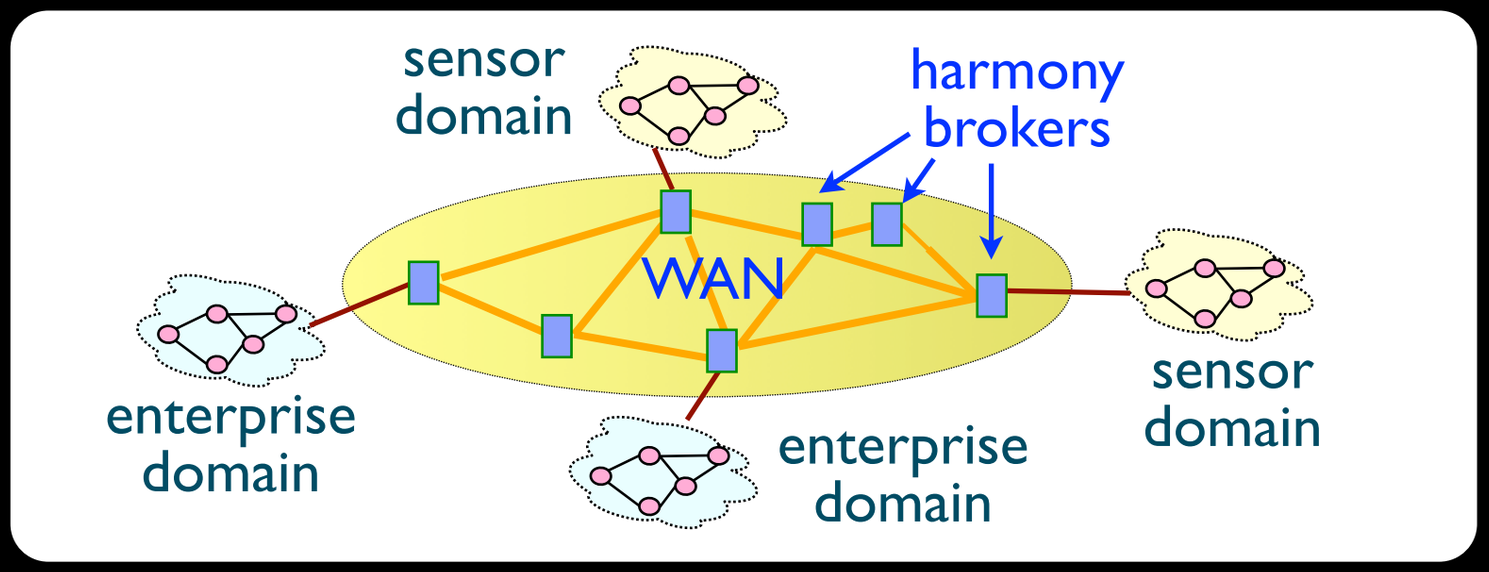

QoS-aware messaging techniques

- Proactive best-path routing (BP): Periodically search for the lowest latency path. If the improvement would be greater than threshold, switch to a new path.
- Reactive QoS-aware routing (QoS): Upon QoS violation based on budgeting mechanism, search for a better path and switch.
- Multipath (MP): Keep multiple paths and send a copy of message through all paths in parallel

Evaluation

- Goal: Which techniques are effective or necessary under various scenarios?
- Testbed: 5 nodes with 3 publishers and 6 subscribers
  - Linux /sbin/tc to emulate link delay and failures
- Compare 4 techniques
  - DP: Direct Path, serving as the baseline
  - 3 Harmony variations: BP, BP+QoS, BP+QoS+MP
- Metric: Latency between each publisher and subscriber pair, loss, overhead

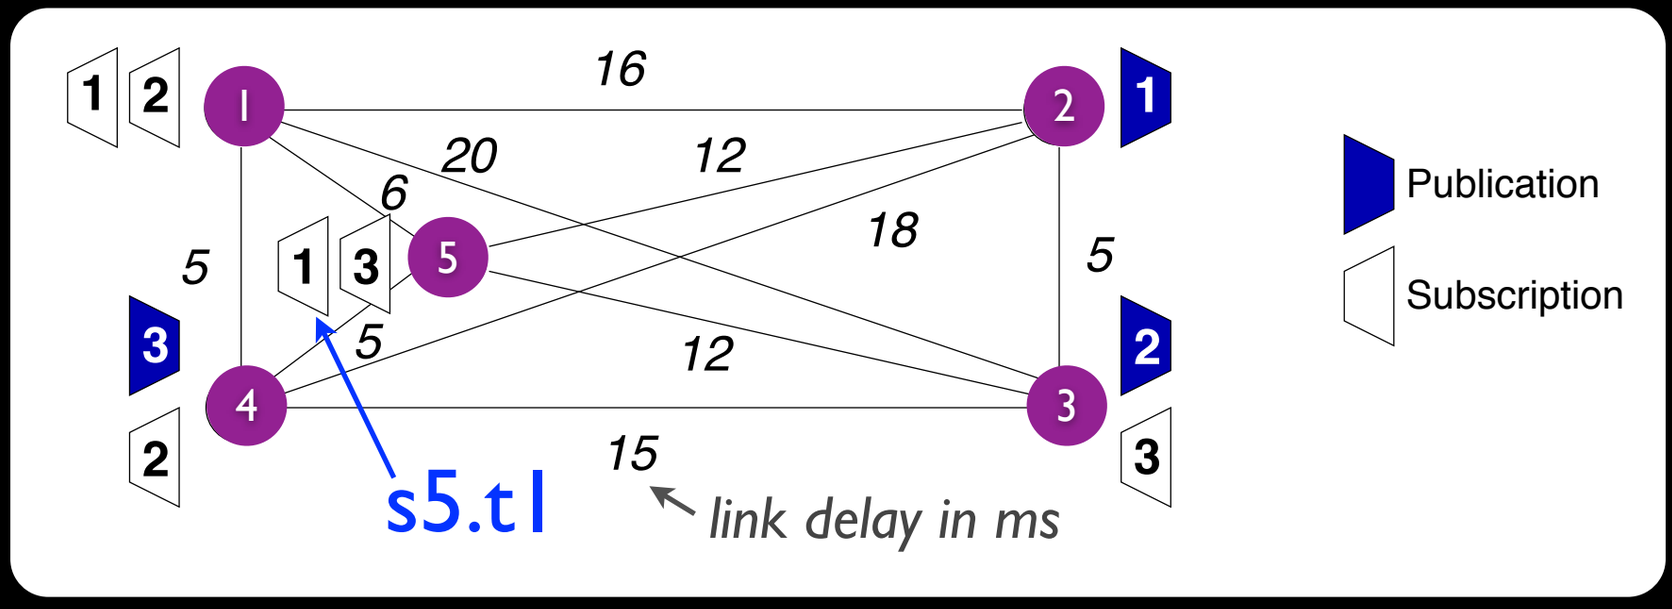

Dynamic path delay

- Scenario: Every 3 minutes, link delay increased to 100 ms with prob of 0.12
- Graph: CDF of latency for one subscriber, s5.t1

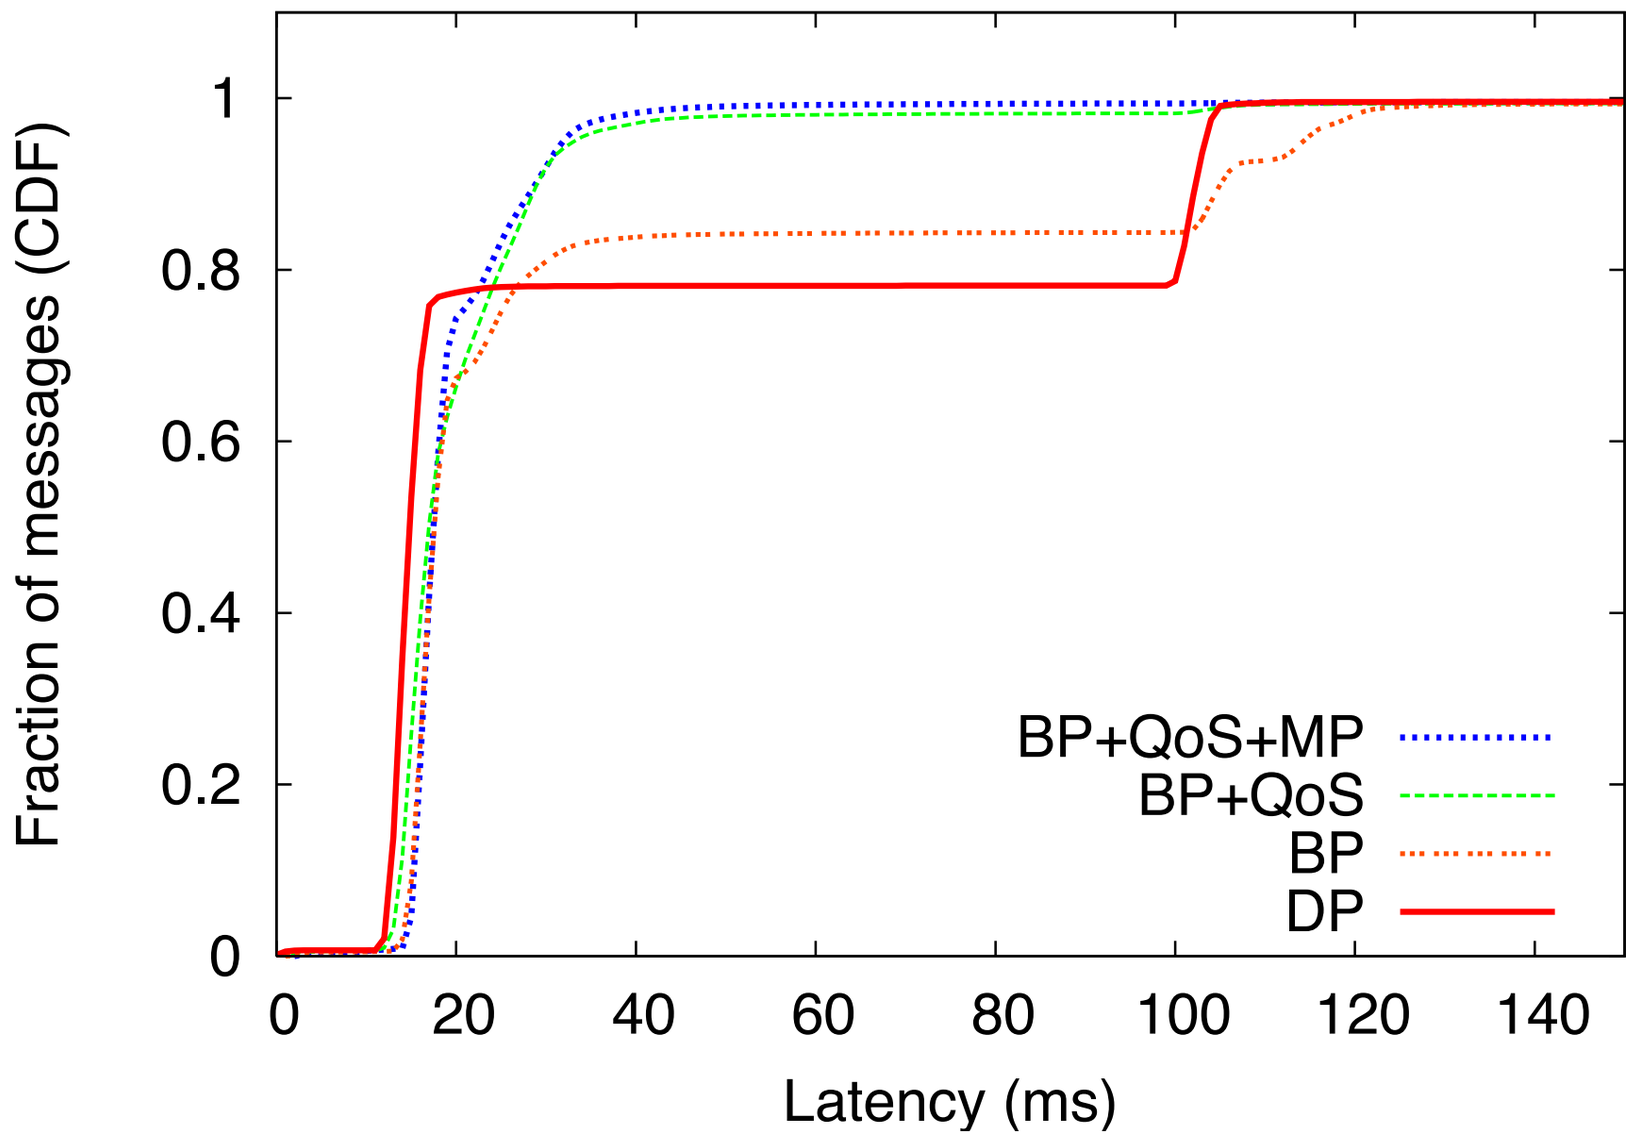

Dynamic path delay

- Scenario: Every 3 minutes, link delay increased to 100 ms with prob of 0.12
- Graph: CDF of latency for one pair, Topic 1 subscriber at node 5

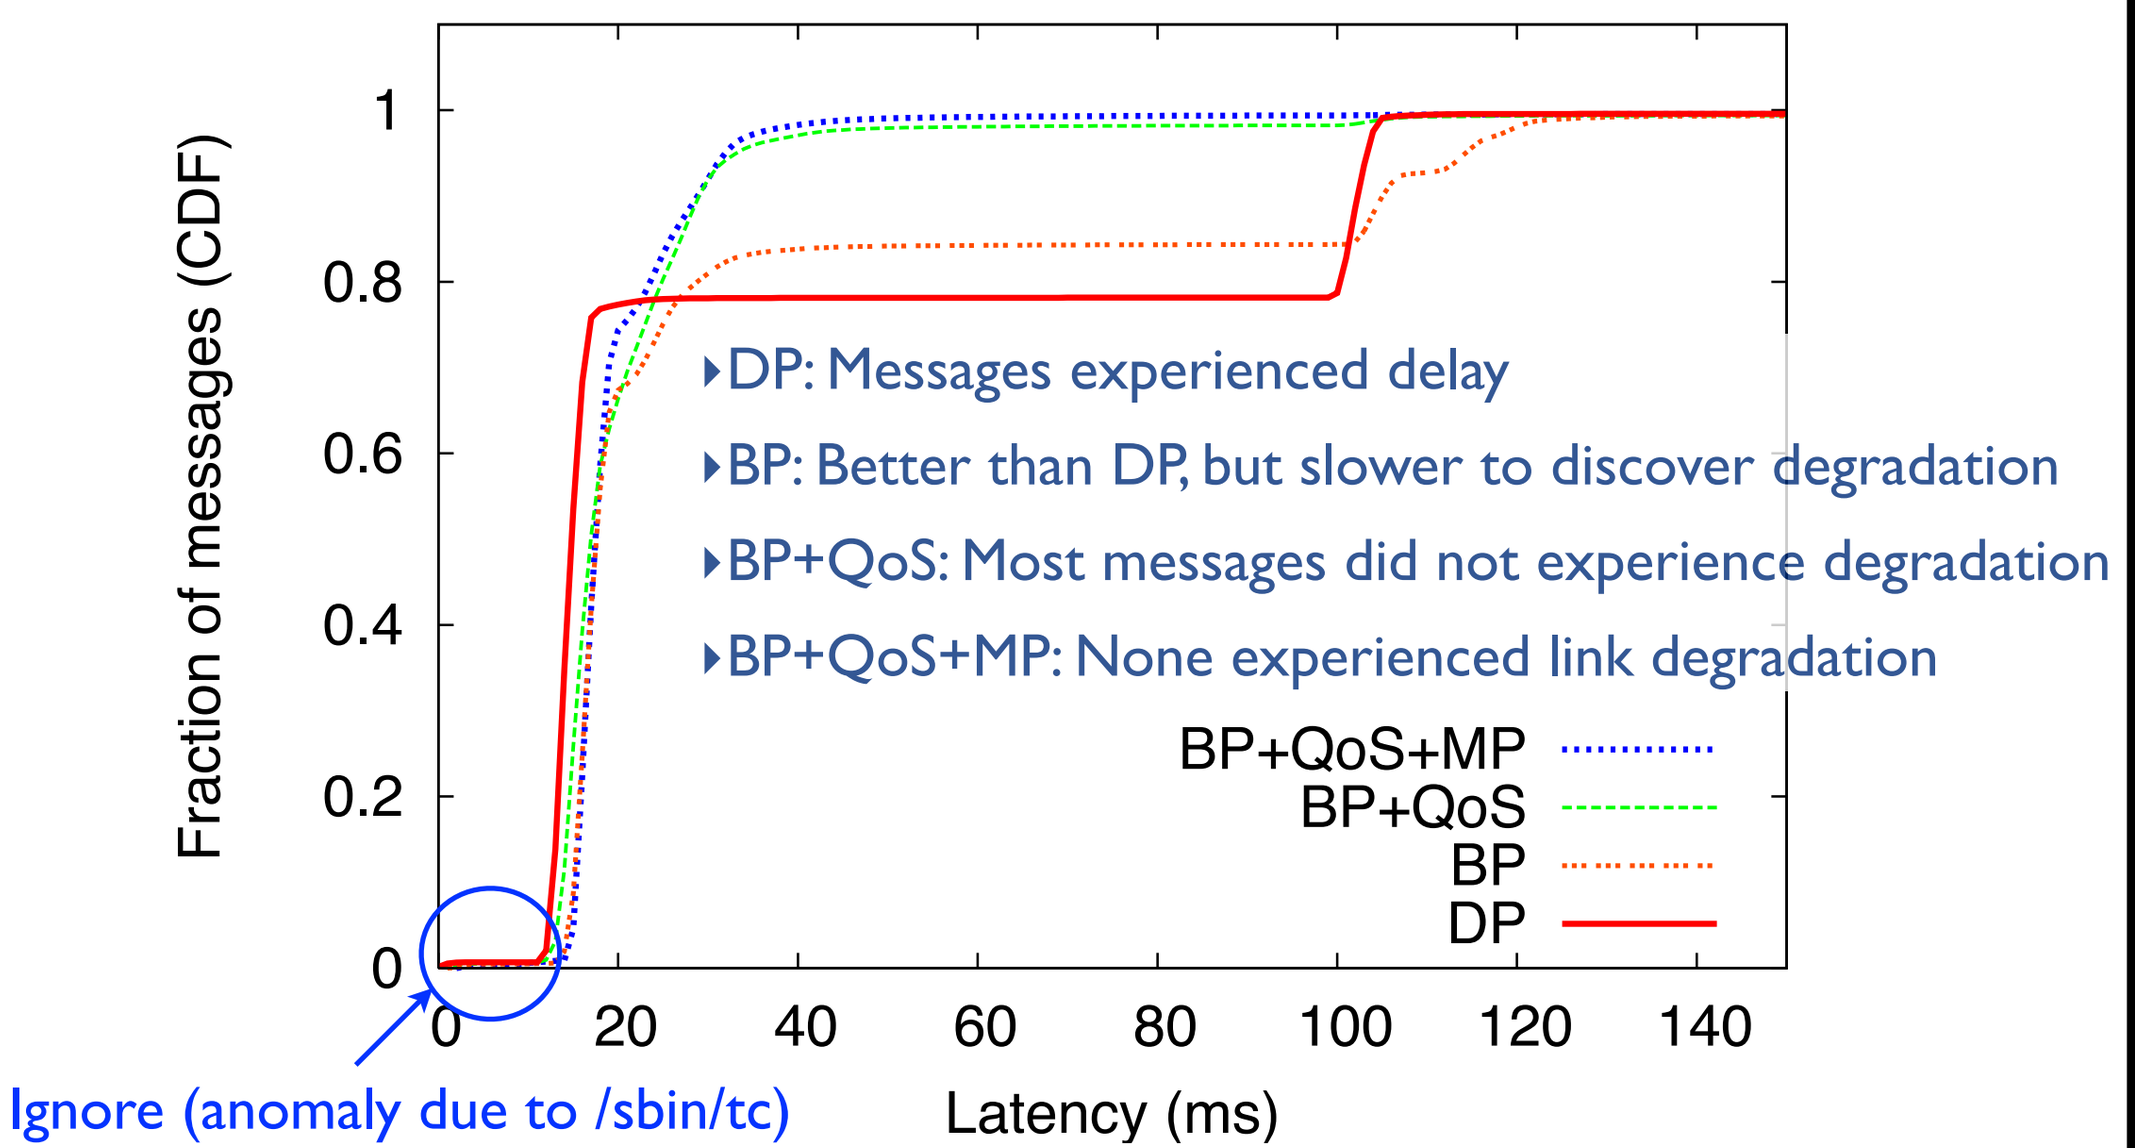

Path failures

- Scenario: Every 3 minutes, a link gets failed with probability of 0.12
- Graph: Loss rate for each subscriber; y-axis in log scale

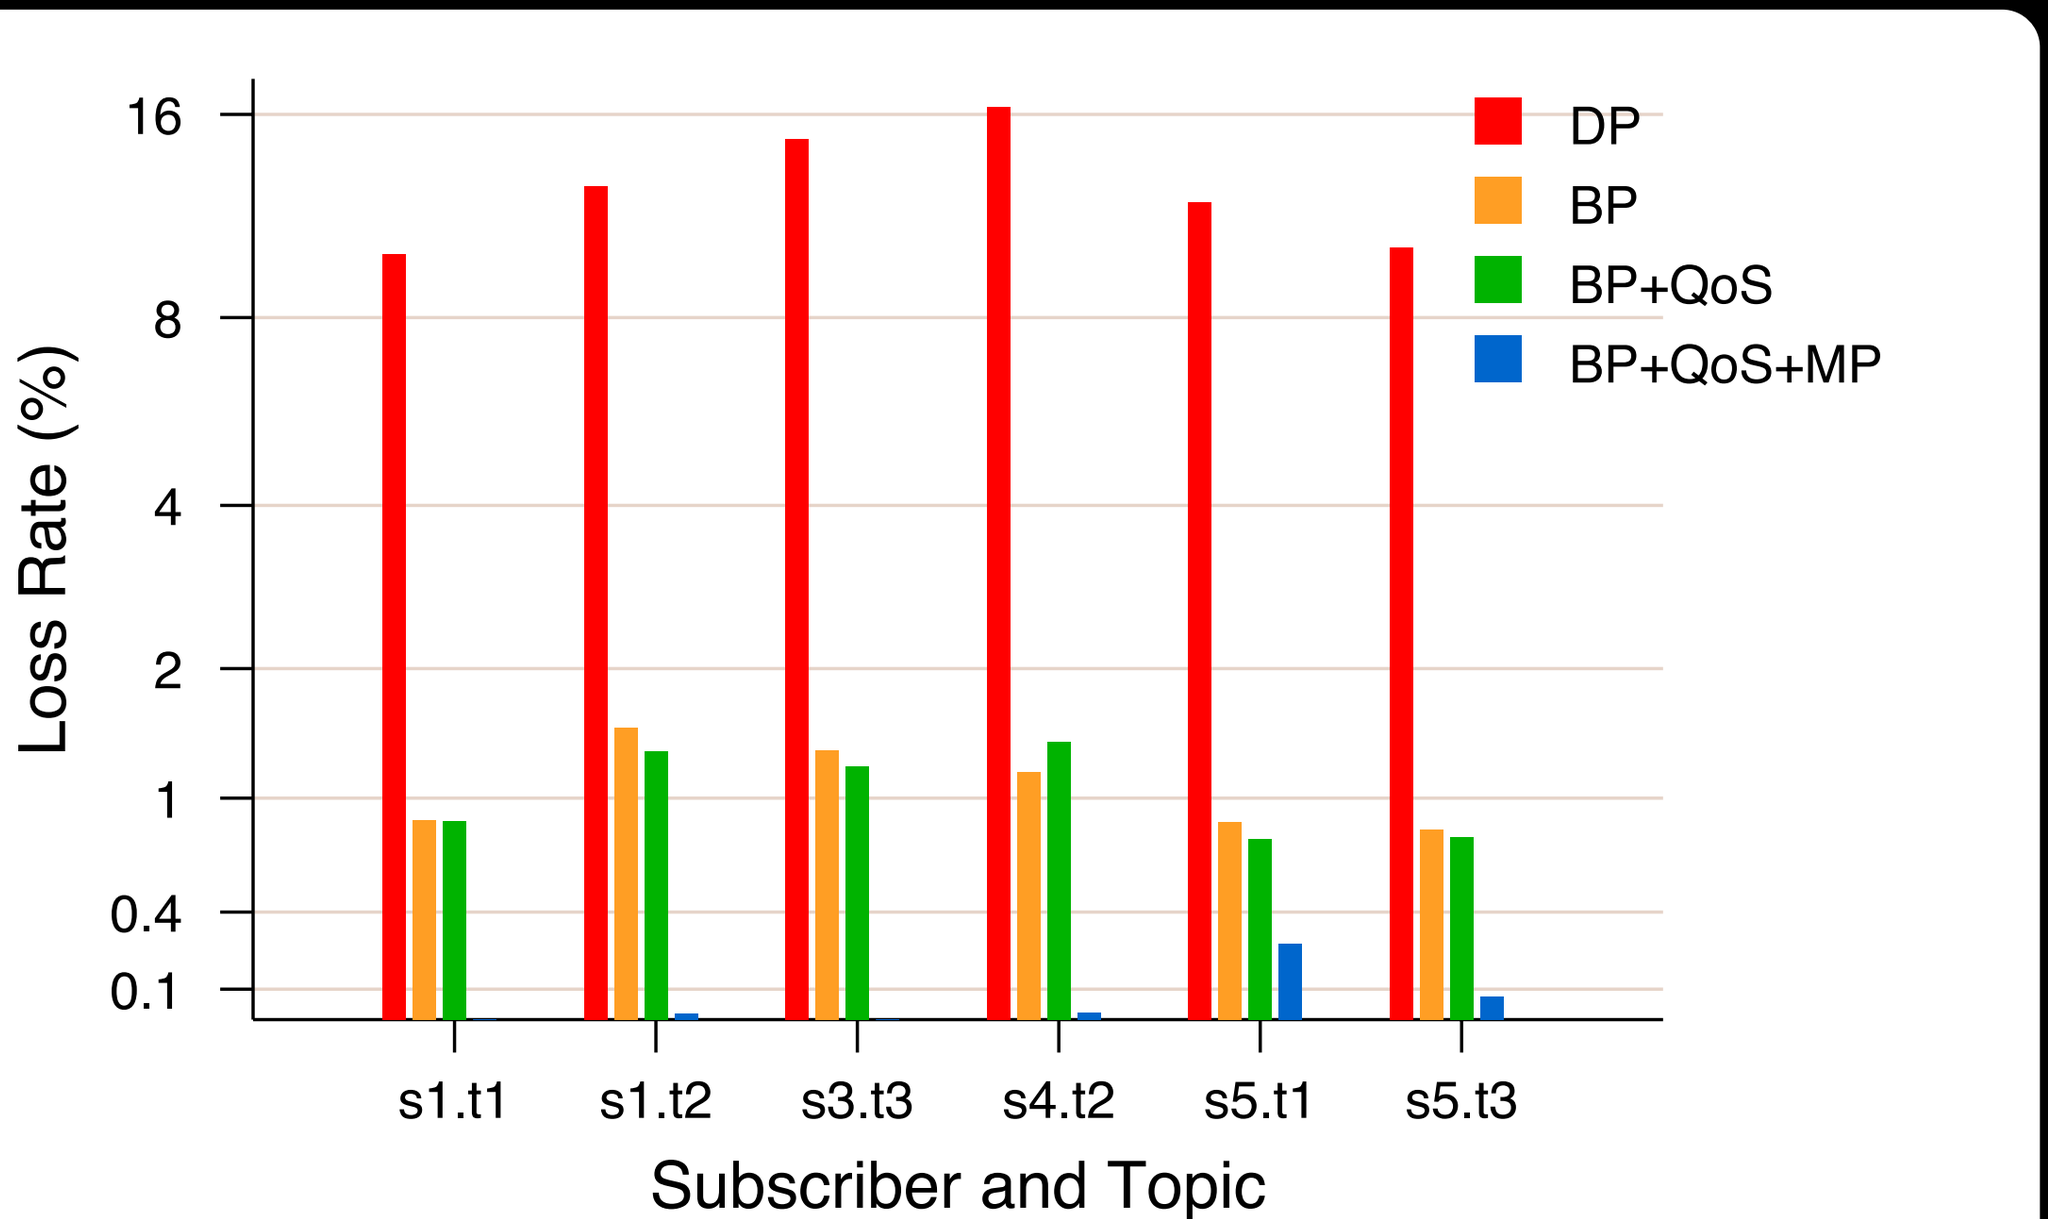

Path failures

- Scenario: Every 3 minutes, a link gets failed with probability of 0.12
- Graph: Loss rate for each subscriber; y-axis in log scale

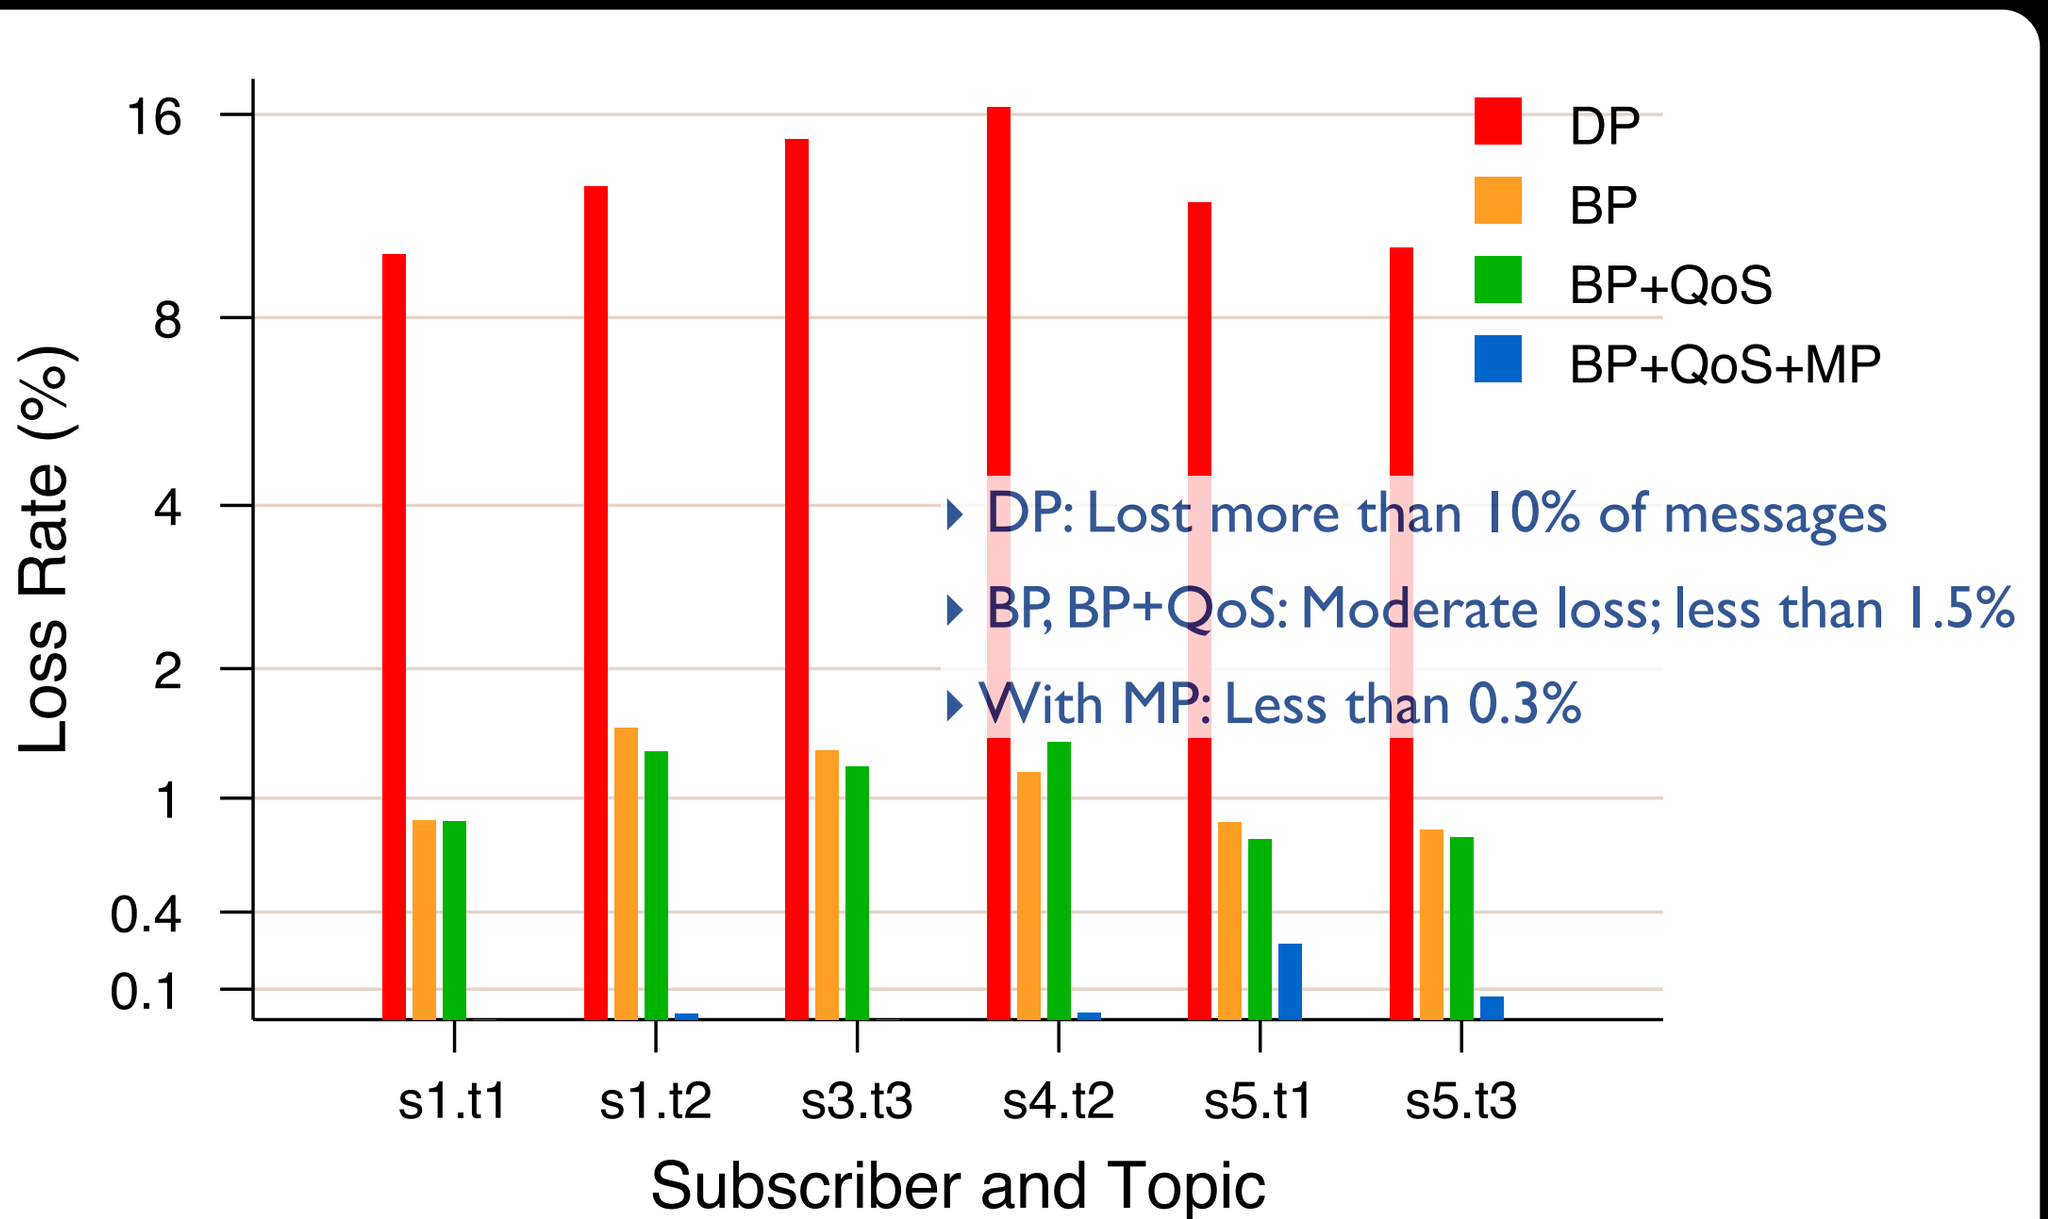

Overhead

- Scenario: Default link delay, repeated for three different message sizes (payload)
- Graph: Total traffic for different message sizes

  - BP+Qos over DP: Small control overhead for keeping multi-hop paths
  - BP+QoS+MP over BP+QoS: Control overhead 3-4 time for maintaining multipath; Data traffic (content) is about 2-3 times due to multipath

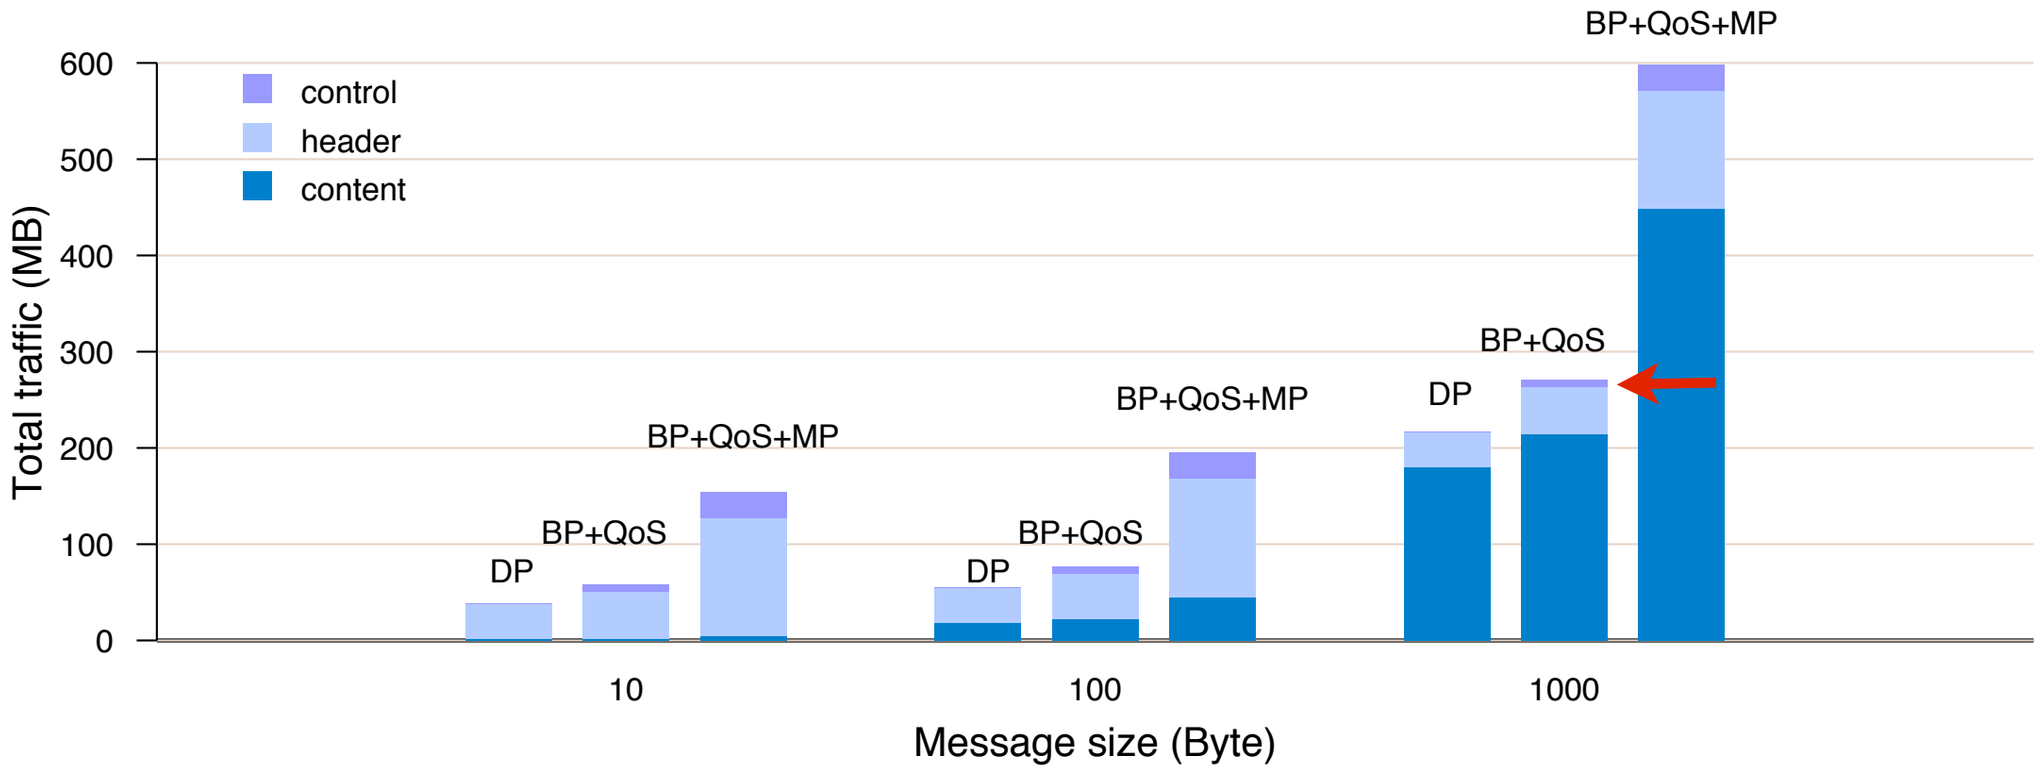

Overhead

- Scenario: Default link delay, repeated for three different message sizes (payload)
- Graph: Total traffic for different message sizes

  - BP+Qos over DP: Small control overhead for keeping multi-hop paths
  - BP+QoS+MP over BP+QoS: Control overhead 3-4 time for maintaining multipath; Data traffic (content) is about 2-3 times due to multipath

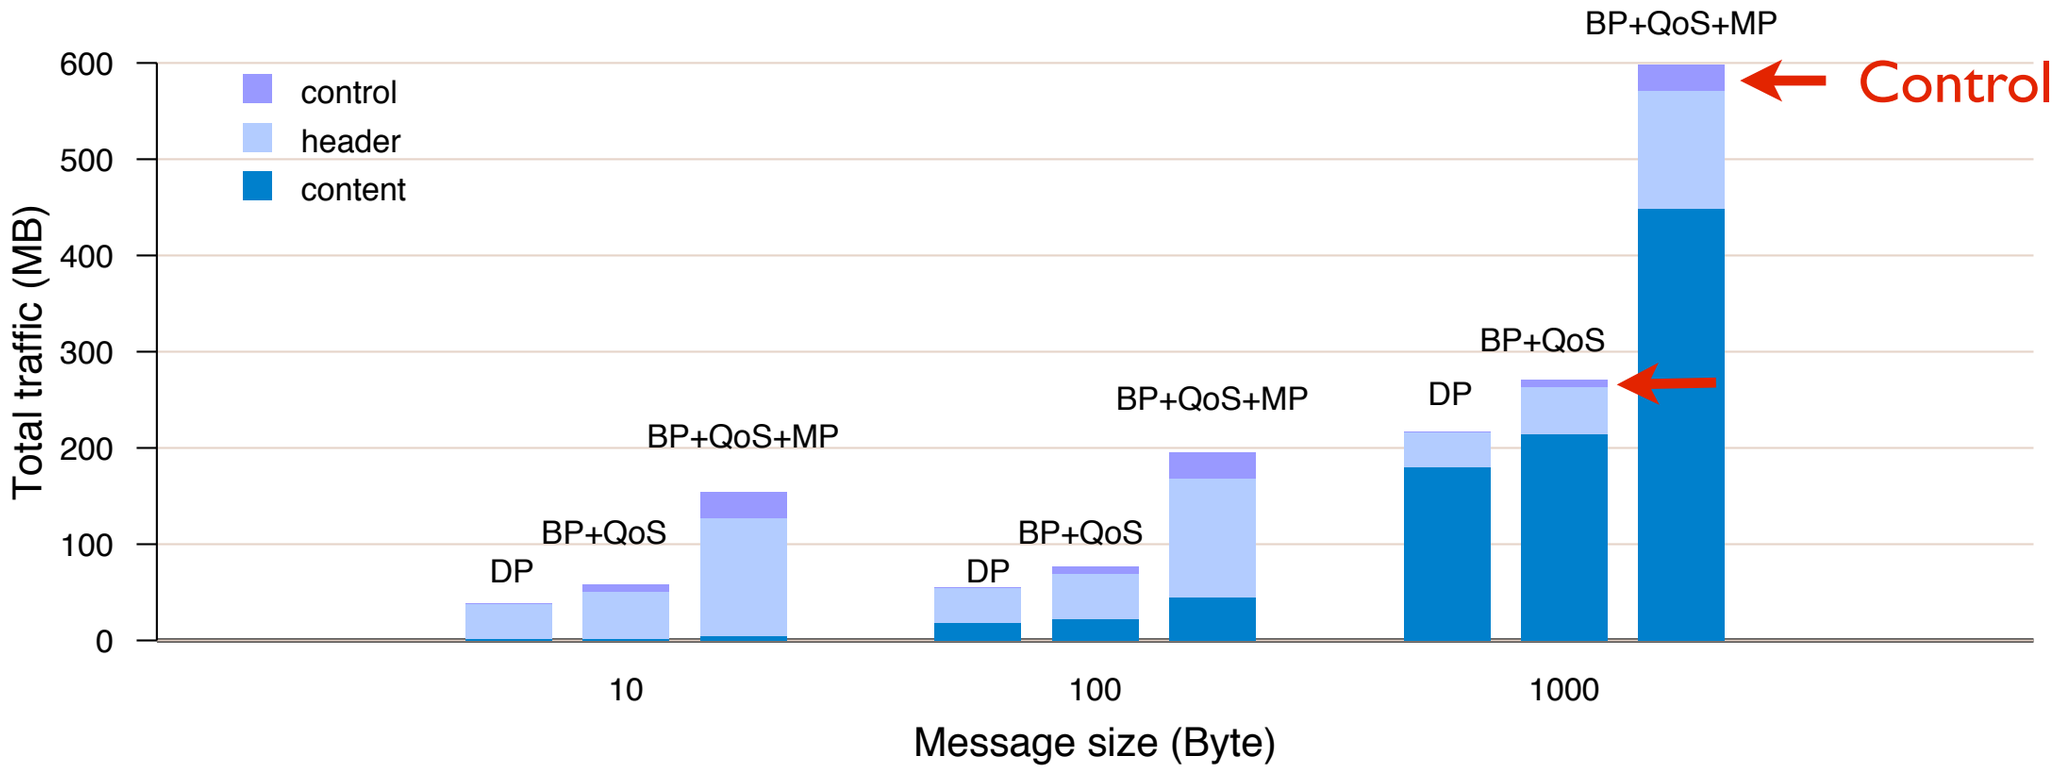

Overhead

- Scenario: Default link delay, repeated for three different message sizes (payload)
- Graph: Total traffic for different message sizes

  - BP+Qos over DP: Small control overhead for keeping multi-hop paths
  - BP+QoS+MP over BP+QoS: Control overhead 3-4 time for maintaining multipath; Data traffic (content) is about 2-3 times due to multipath

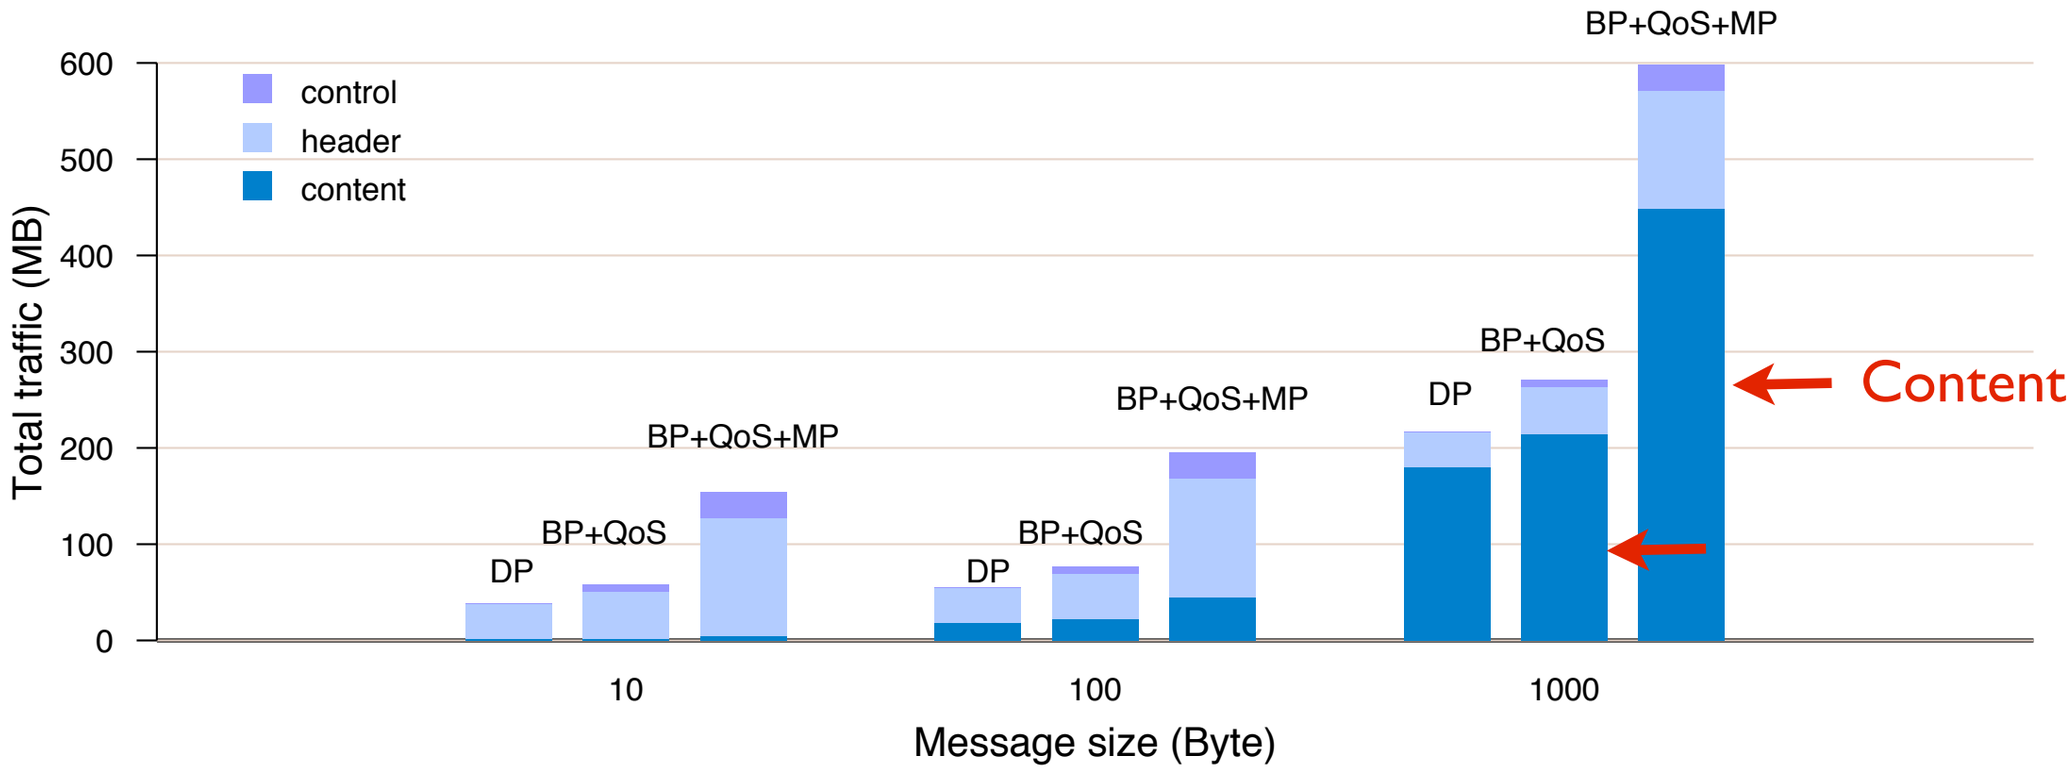

Related work

- JECho [Zhou01]: Opportunistic overlay topology to minimize end-to-end latency; no mechanism for specifying latency requirement
- IndiQos [Carvalho05]: Resource reservation approach for *stable* network (latency and bandwidth remain satisfied for all active flows)
- Reliable message delivery [Bhola02, Kazemzadeh09]: Guarantee delivery but are not latency-aware; intended for stable environment, such as data centers, and not for dynamic WAN

Conclusion

- Harmony provides QoS aware pub/sub messaging in wide-area networks
- Summary
  - Dynamic delay changes: BP+QoS
  - Packet loss is prevalent: BP+QoS+MP
  - Overhead (control traffic) is small
- Future work
  - Adapt coding scheme for multipath to reduce redundant data traffic

Authors

Minkyong Kim

IBM Watson Research, NY

Kyriakos Karenos

Microsoft, UK

Fan Ye

IBM Watson Research, NY

Johnathan Reason

IBM Watson Research, NY

Hui Lei

IBM Watson Research, NY

Konstantin Shagin

IBM Haifa Research