

A per-Application Account of Bufferbloat: Causes and Impact on Users

Andrea Araldo

INFRES, Télécom ParisTech
LRI, Université Paris Sud
araldo@lri.fr

Dario Rossi

INFRES, Télécom ParisTech

dario.rossi@telecom-paristech.fr



*TRAC Workshop 2014
August 6th 2014, Nicosia - Cyprus*



Agenda

- Introduction on bufferbloat (high queueing delay)
- Passive methodology to infer the queueing delay in the Internet
 - *Bufferbloat Dissector*: Implementation in Tstat
 - (source code in <http://perso.telecom-paristech.fr/~araldo>)
 - Validation, calibration and analysis of overhead
- Results from a real ISP network
 - *Is bufferbloat on the user or on the ISP side?*
 - Bufferbloat vs ISP load correlation
 - *Impact on users? / Which applications suffer bufferbloat?*
 - We leverage Tstat DPI → Per application view and QoE
 - *Which application mix causes bufferbloat?*
 - We apply A-priori alg. → Root cause analysis

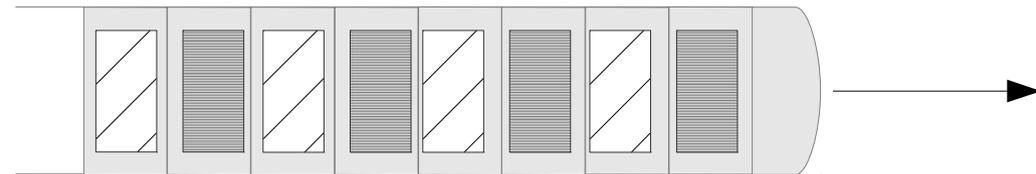
Agenda

- Introduction on bufferbloat (high queueing delay)
- Passive methodology to infer the queueing delay in the Internet
 - ~~Bufferbloat Dissector: Implementation in Tstat~~
 - (source code in <http://perso.telecom-paristech.fr/~araldo>)
 - ~~Validation, calibration and analysis of overhead~~
- Results from a real ISP network
 - *Is bufferbloat on the user or on the ISP side?*
 - Bufferbloat vs ISP load correlation
 - *Impact on users? / Which applications suffer bufferbloat?*
 - We leverage Tstat DPI → Per application view and QoE
 - *Which application mix causes bufferbloat?*
 - We apply A-priori alg. → Root cause analysis

What is bufferbloat? [GN11]

Memory is cheap

- ➔ Large buffers in net equipment
- ➔ Long queues in the uplink without TCP losses
- ➔ TCP sender cannot detect congestion and does not adjust rate
- ➔ Congestion is exacerbated



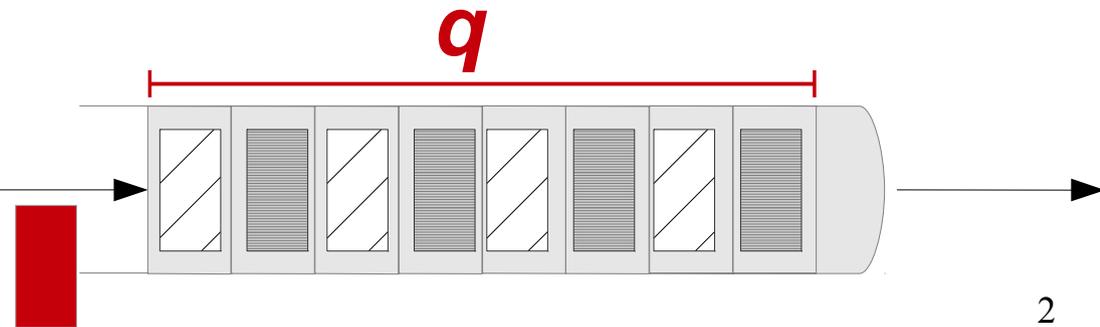
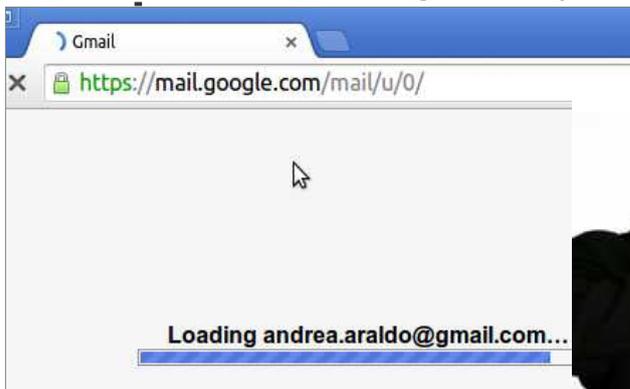
What is bufferbloat? [GN11]

Memory is cheap

- ➔ Large buffers in net equipment
- ➔ Long queues in the uplink without TCP losses
- ➔ TCP sender cannot detect congestion and does not adjust rate
- ➔ Congestion is exacerbated
- ➔ High queueing delay

Example of ADSL router [IMC 10]:

- **256 KB** buffer
- **512 Kbps** uplink
- ➔ Queueing delay up to **4 s**



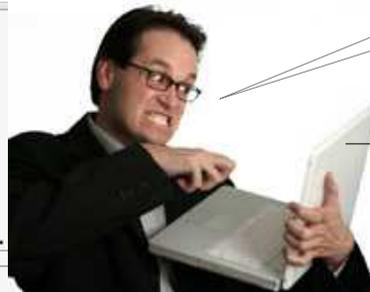
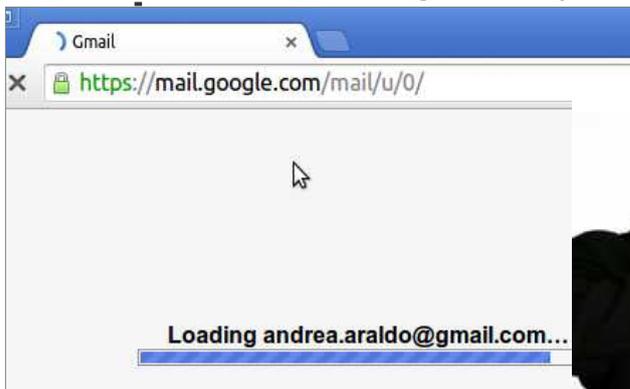
What is bufferbloat? [GN11]

Memory is cheap

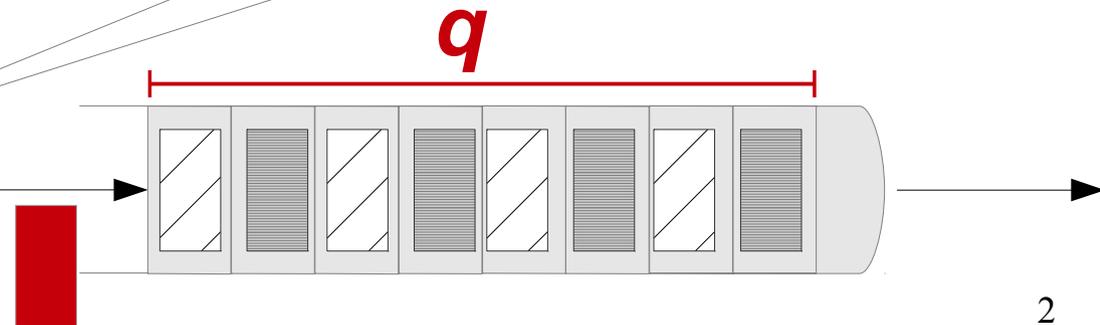
- ➔ Large buffers in net equipment
- ➔ Long queues in the uplink without TCP losses
- ➔ TCP sender cannot detect congestion and does not adjust rate
- ➔ Congestion is exacerbated
- ➔ High queueing delay

Example of ADSL router [IMC 10]:

- **256 KB** buffer
- **512 Kbps** uplink
- ➔ Queueing delay up to **4 s**



Grrrrrrrrrrrrrrrrrr!
Google is not working



Previous work

Active measures

[IMC 10], [IMC 12], [IMC 12 a],
[SIGCOMM 11], [SIGCOMM 11 a],
[BOB12]

Delay is measured under
controlled load
(stress tests)

- x *Limitation:*
they give **maximum** queueing
delay rather than the **typical**
one.

Passive measures

[PAM 13], [CCR 12 a], [PAM 13 a], [TMA 13]

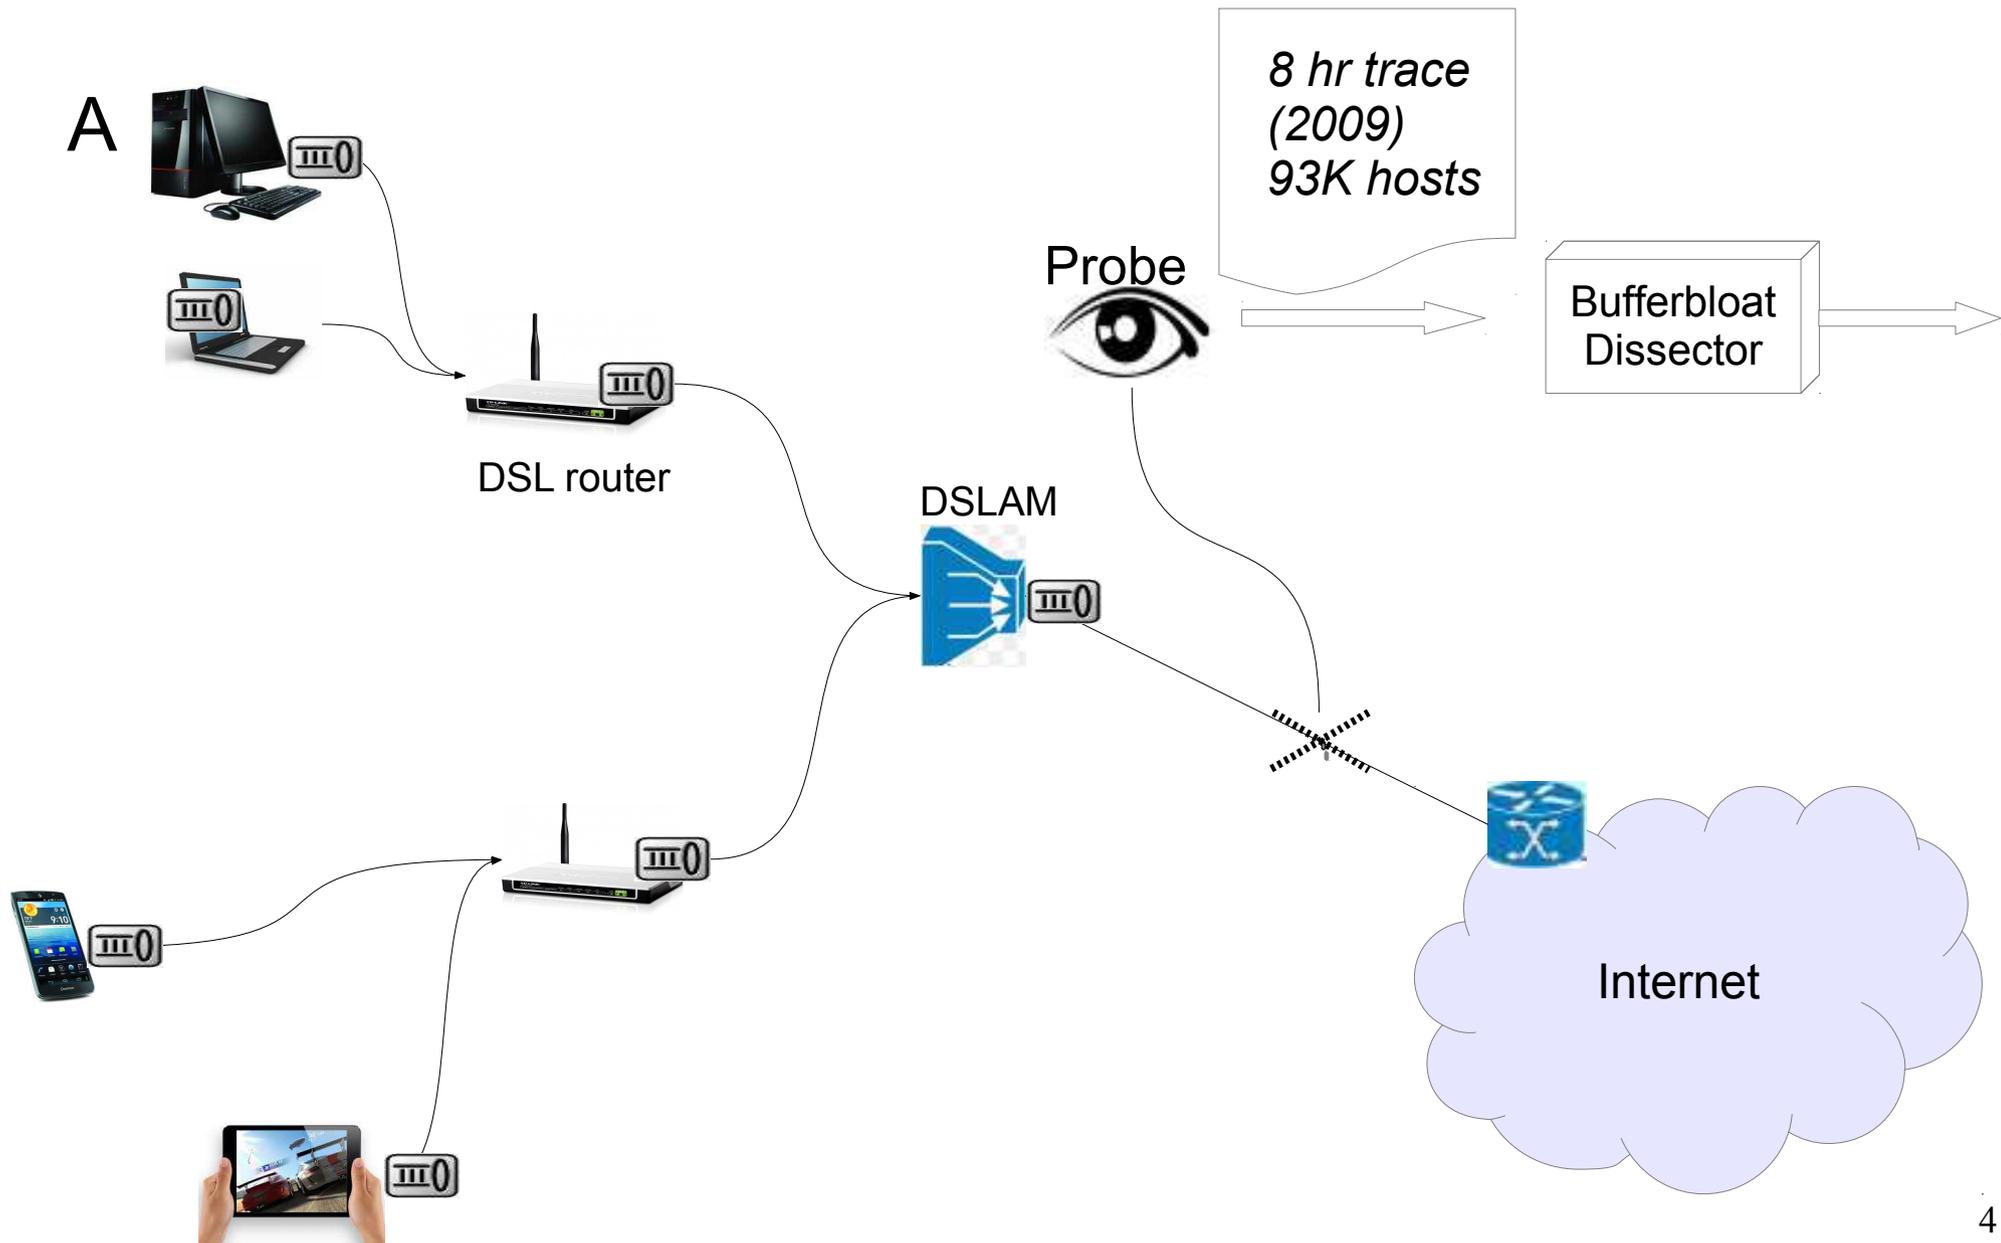
- x *Limitation:*
application-blind
 - High delay can be (in)tolerable
depending on the application
 - They say nothing about user
experience

[TMA 14]

- x *Limitation:*
application running on 55 hosts
 - Small scale

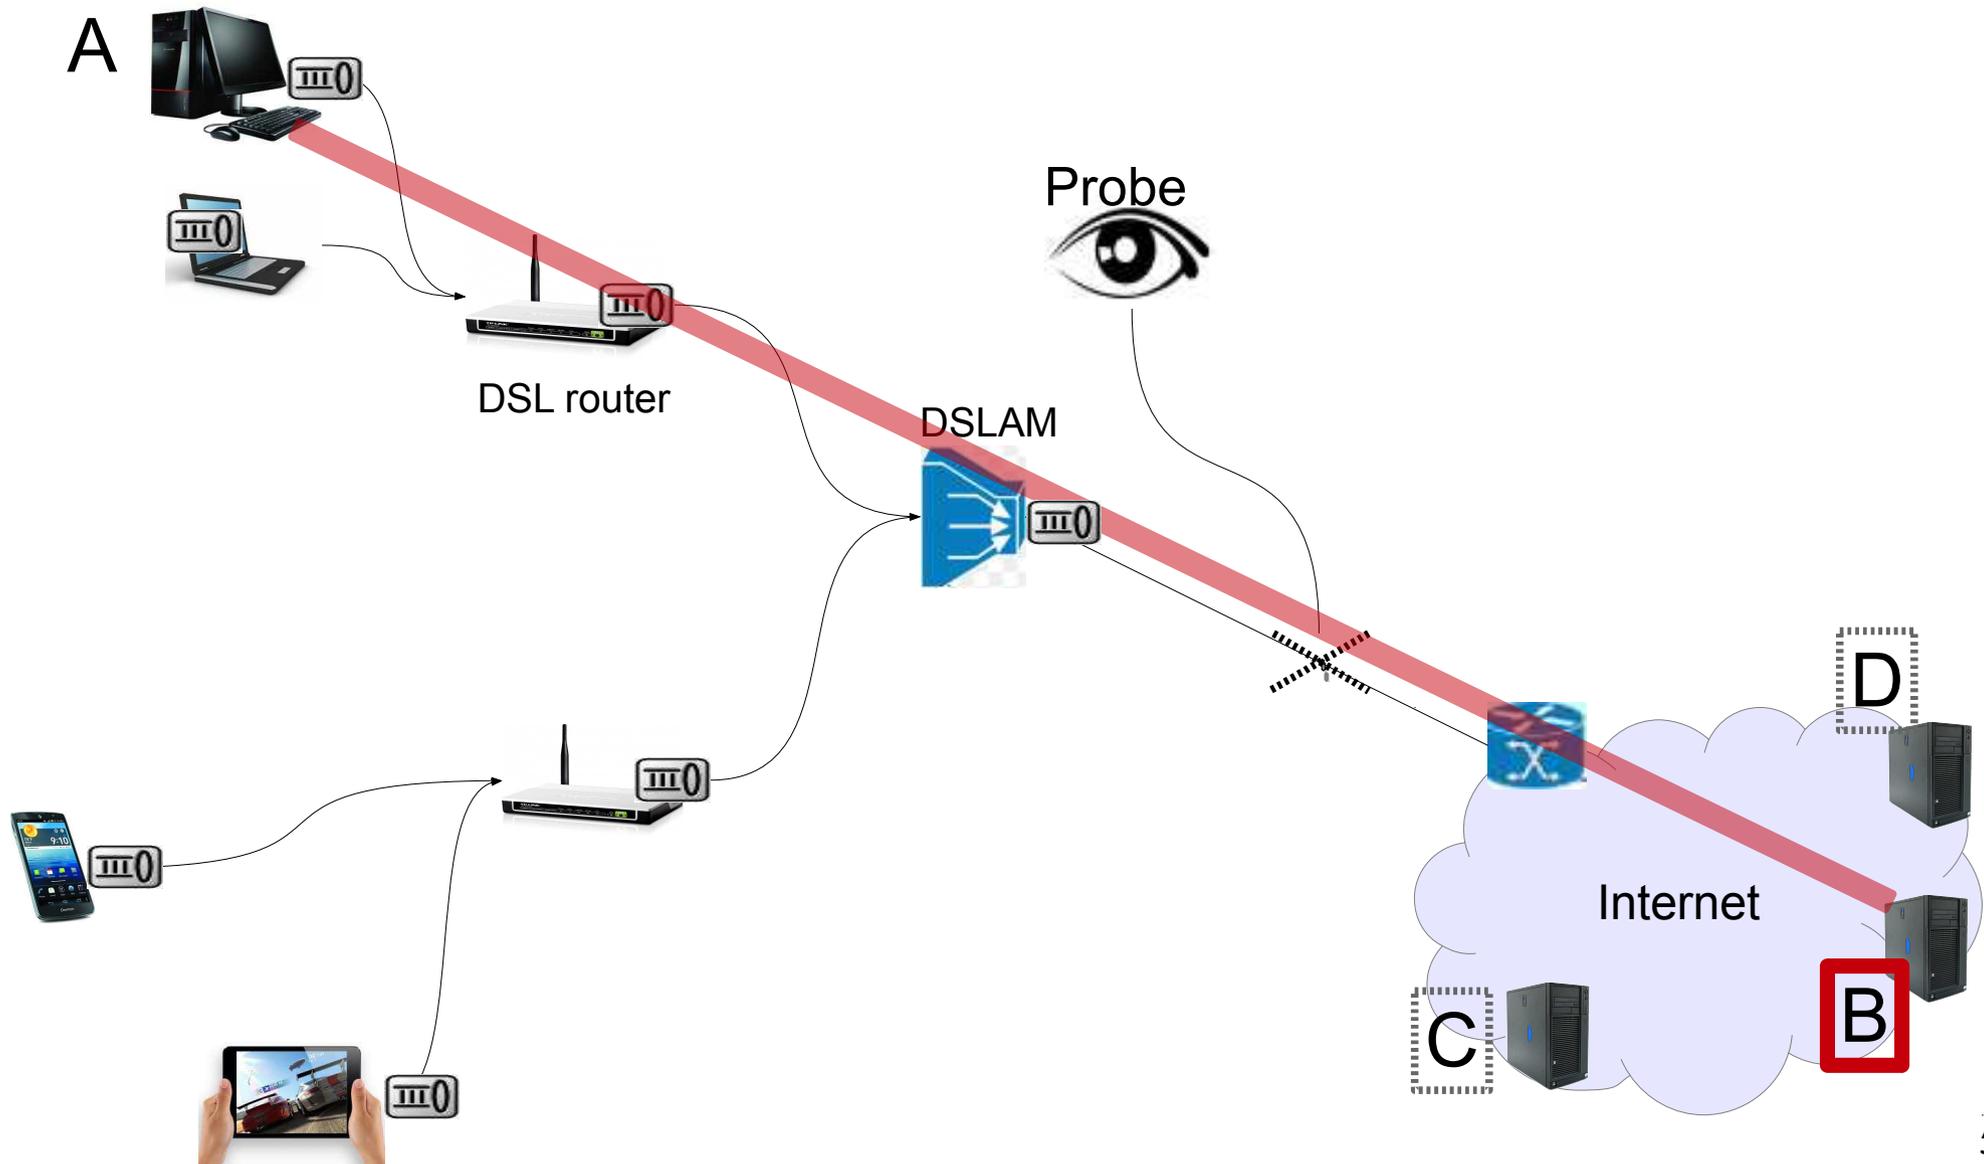
Methodology

Measurement scenario



Methodology

Measurement scenario



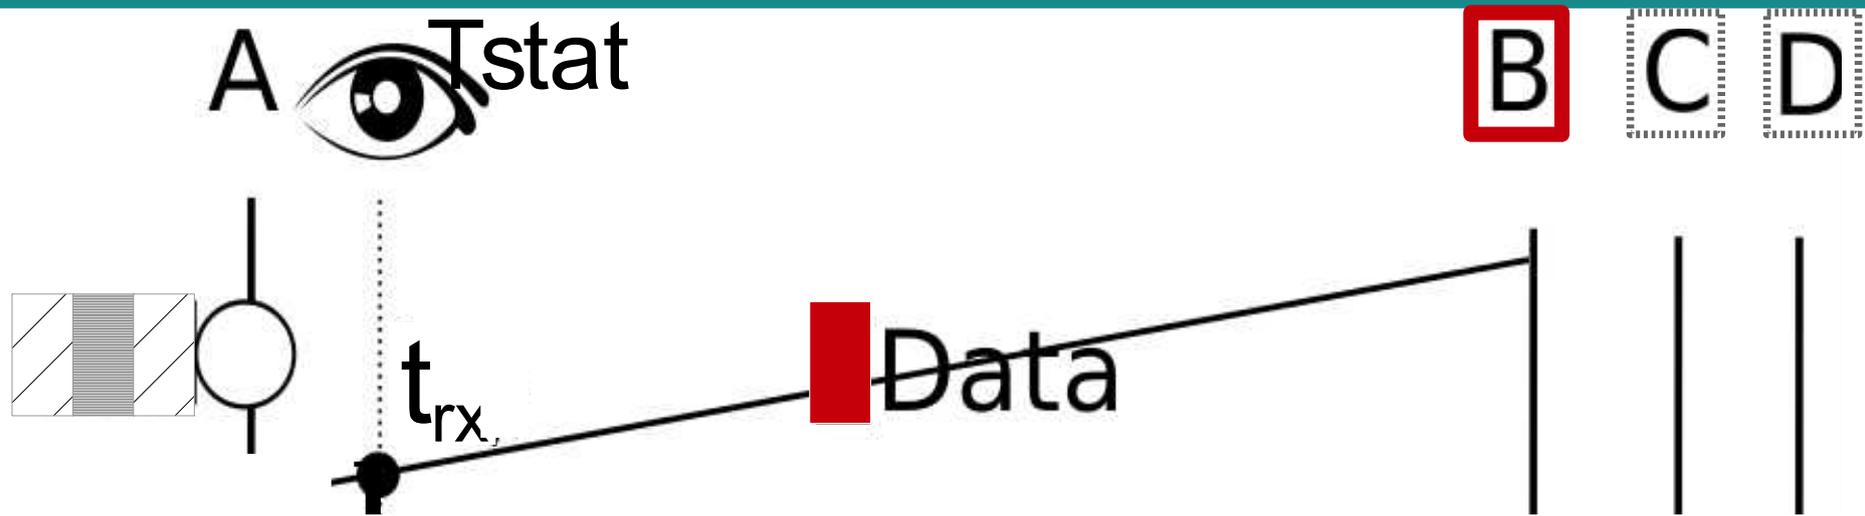
Methodology

Queueing delay inference



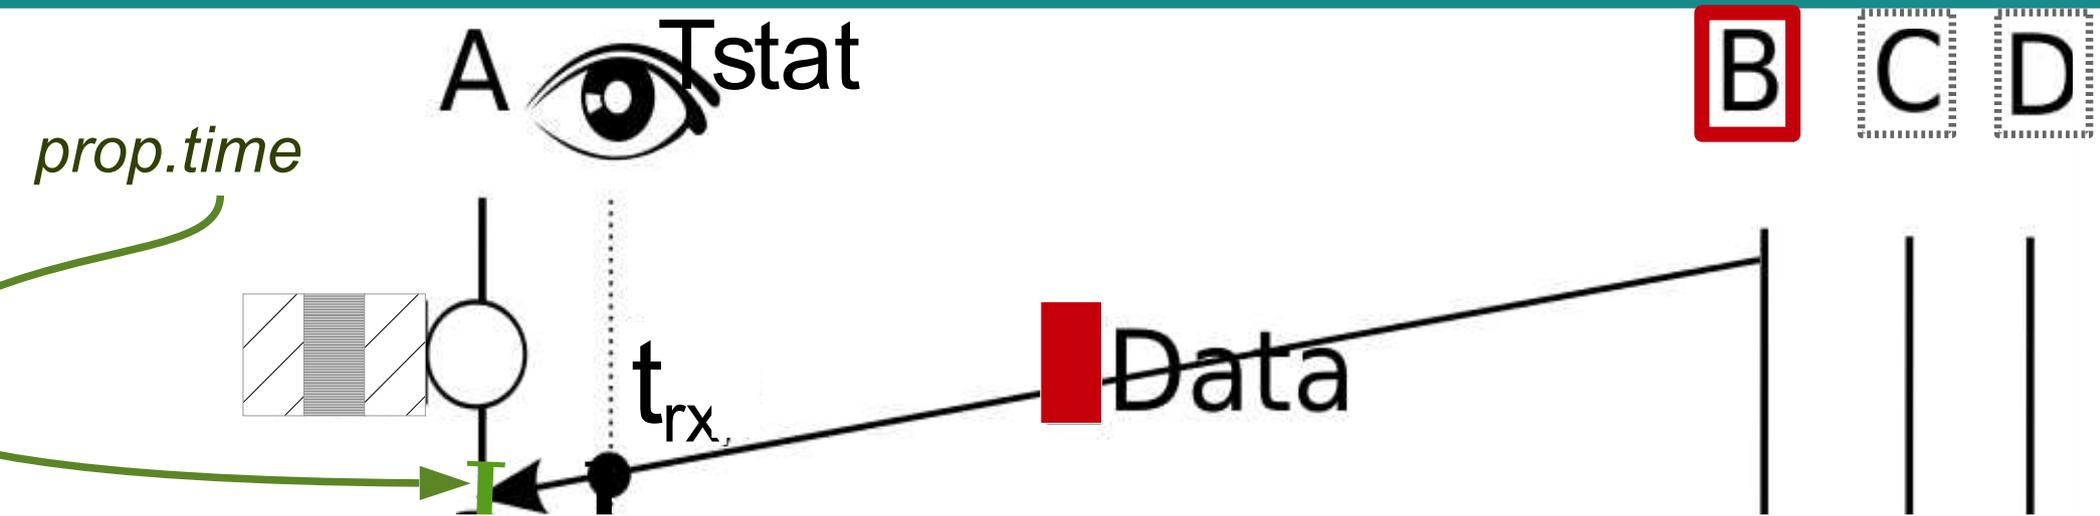
Methodology

Queueing delay inference



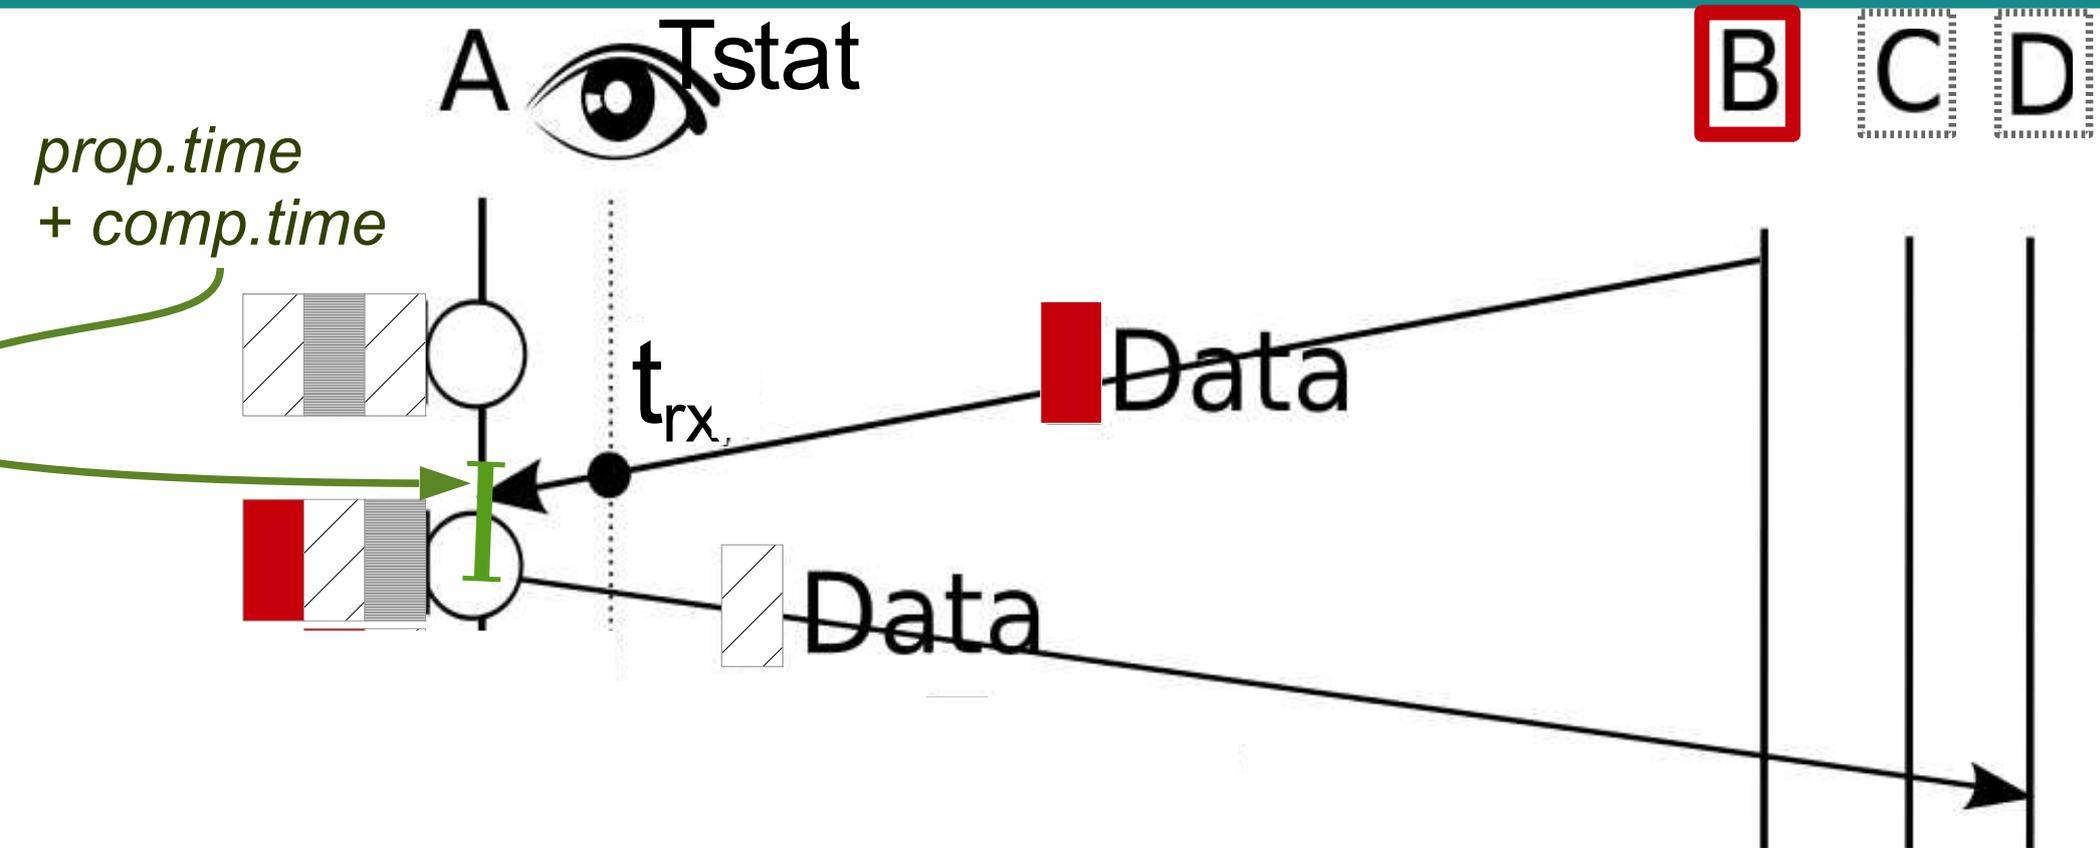
Methodology

Queueing delay inference



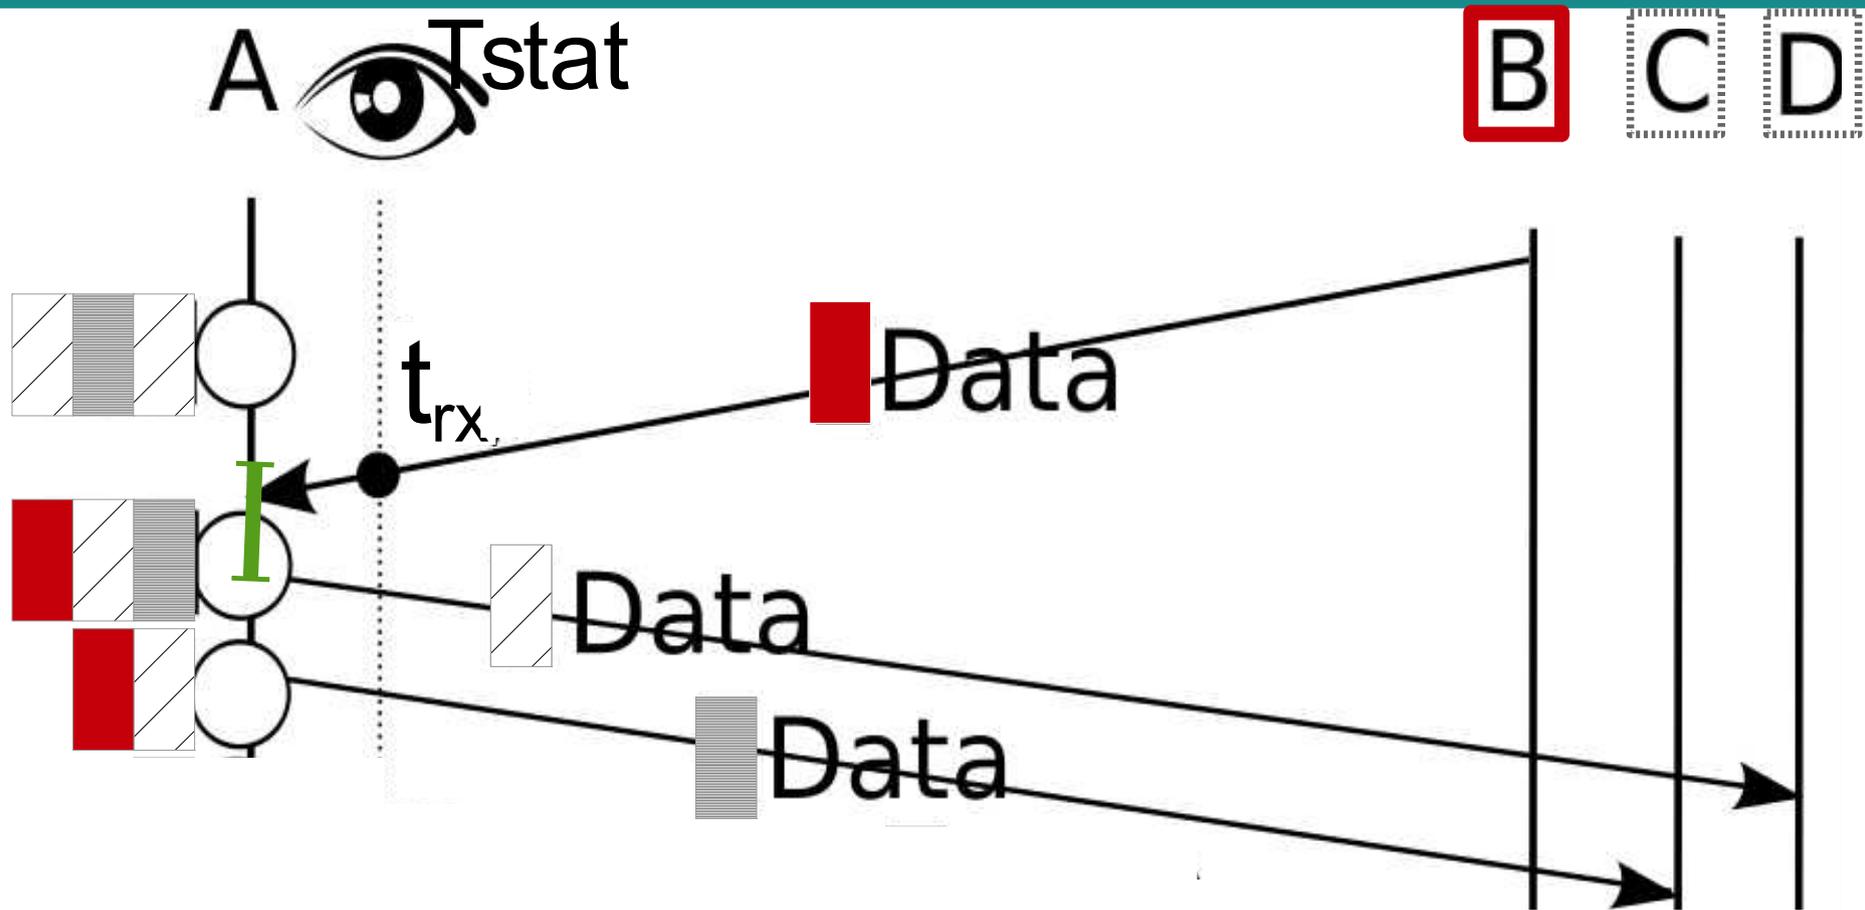
Methodology

Queueing delay inference



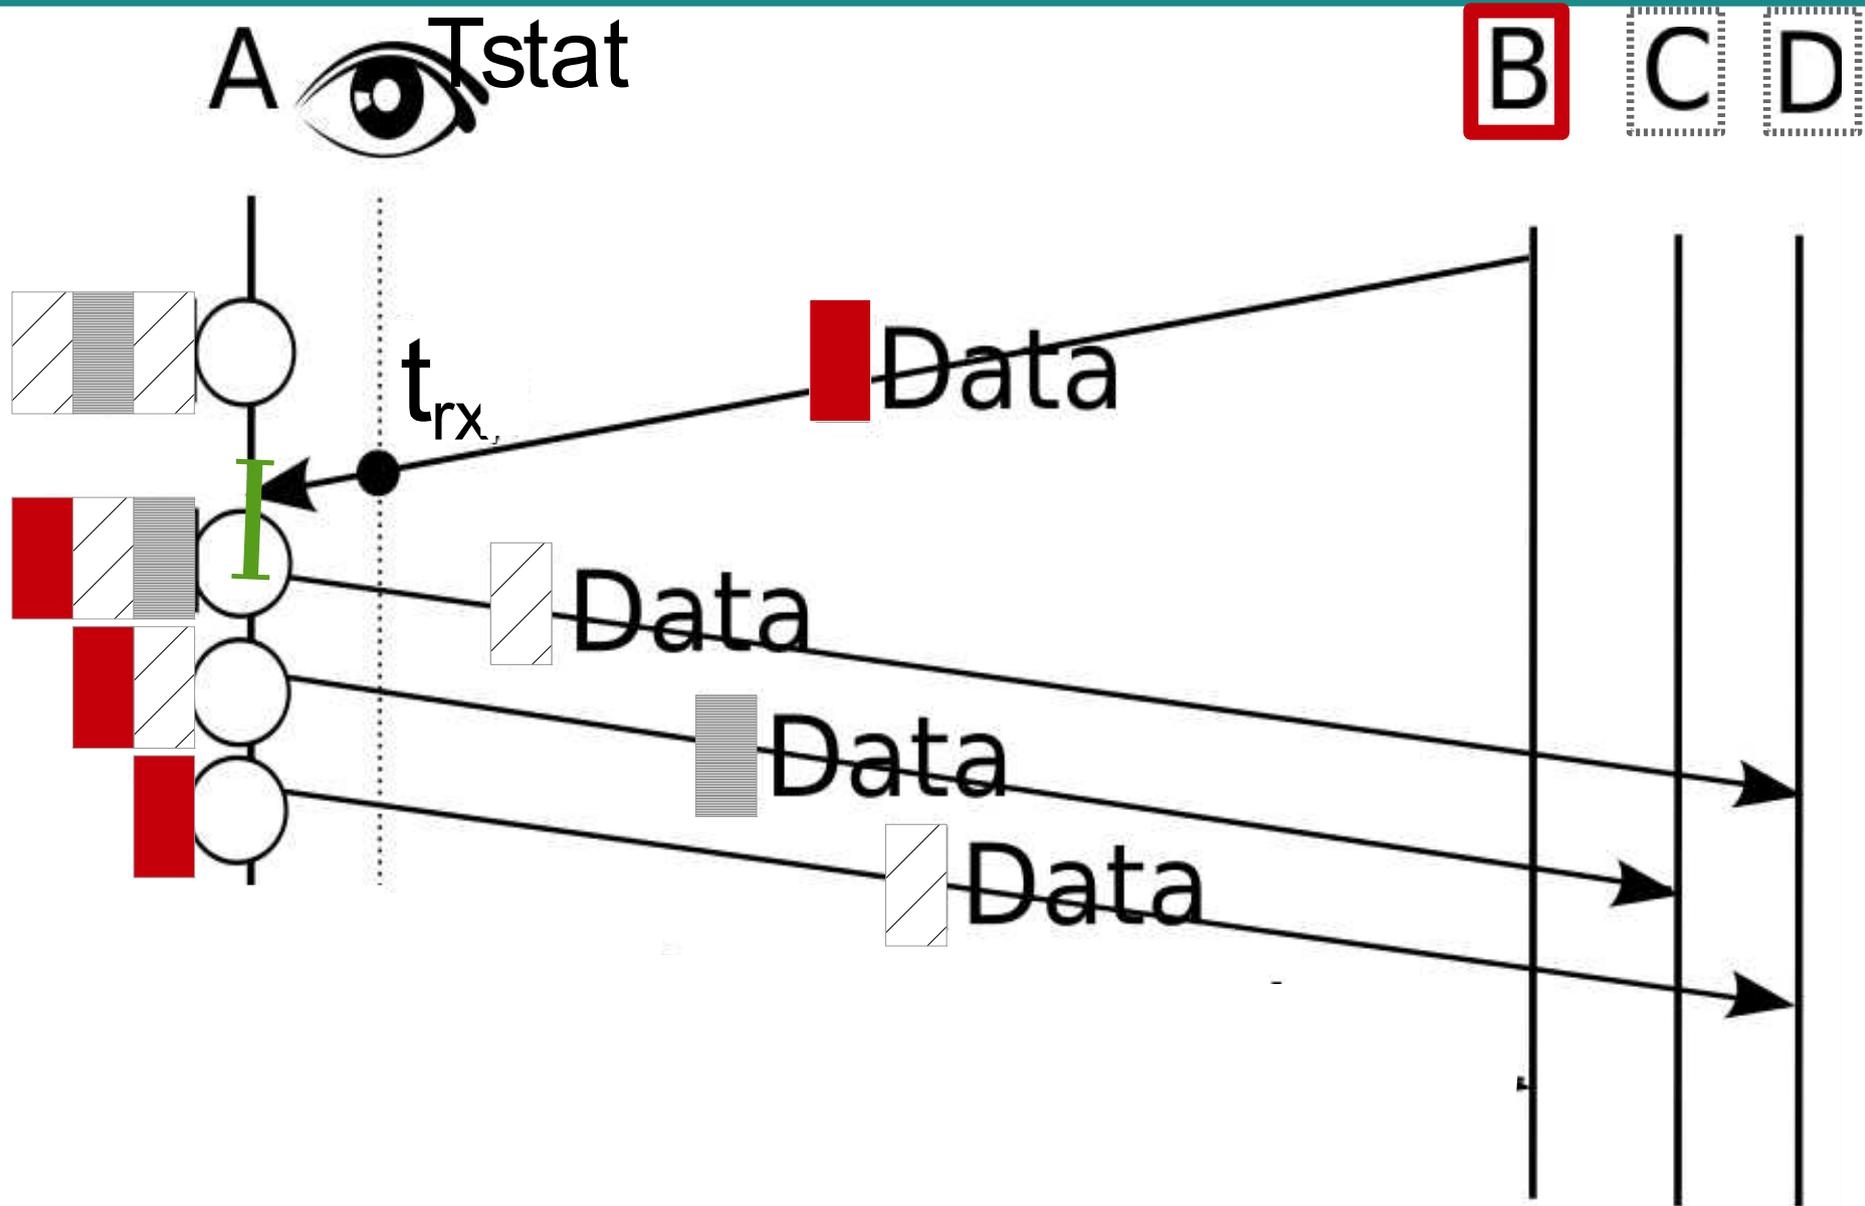
Methodology

Queueing delay inference



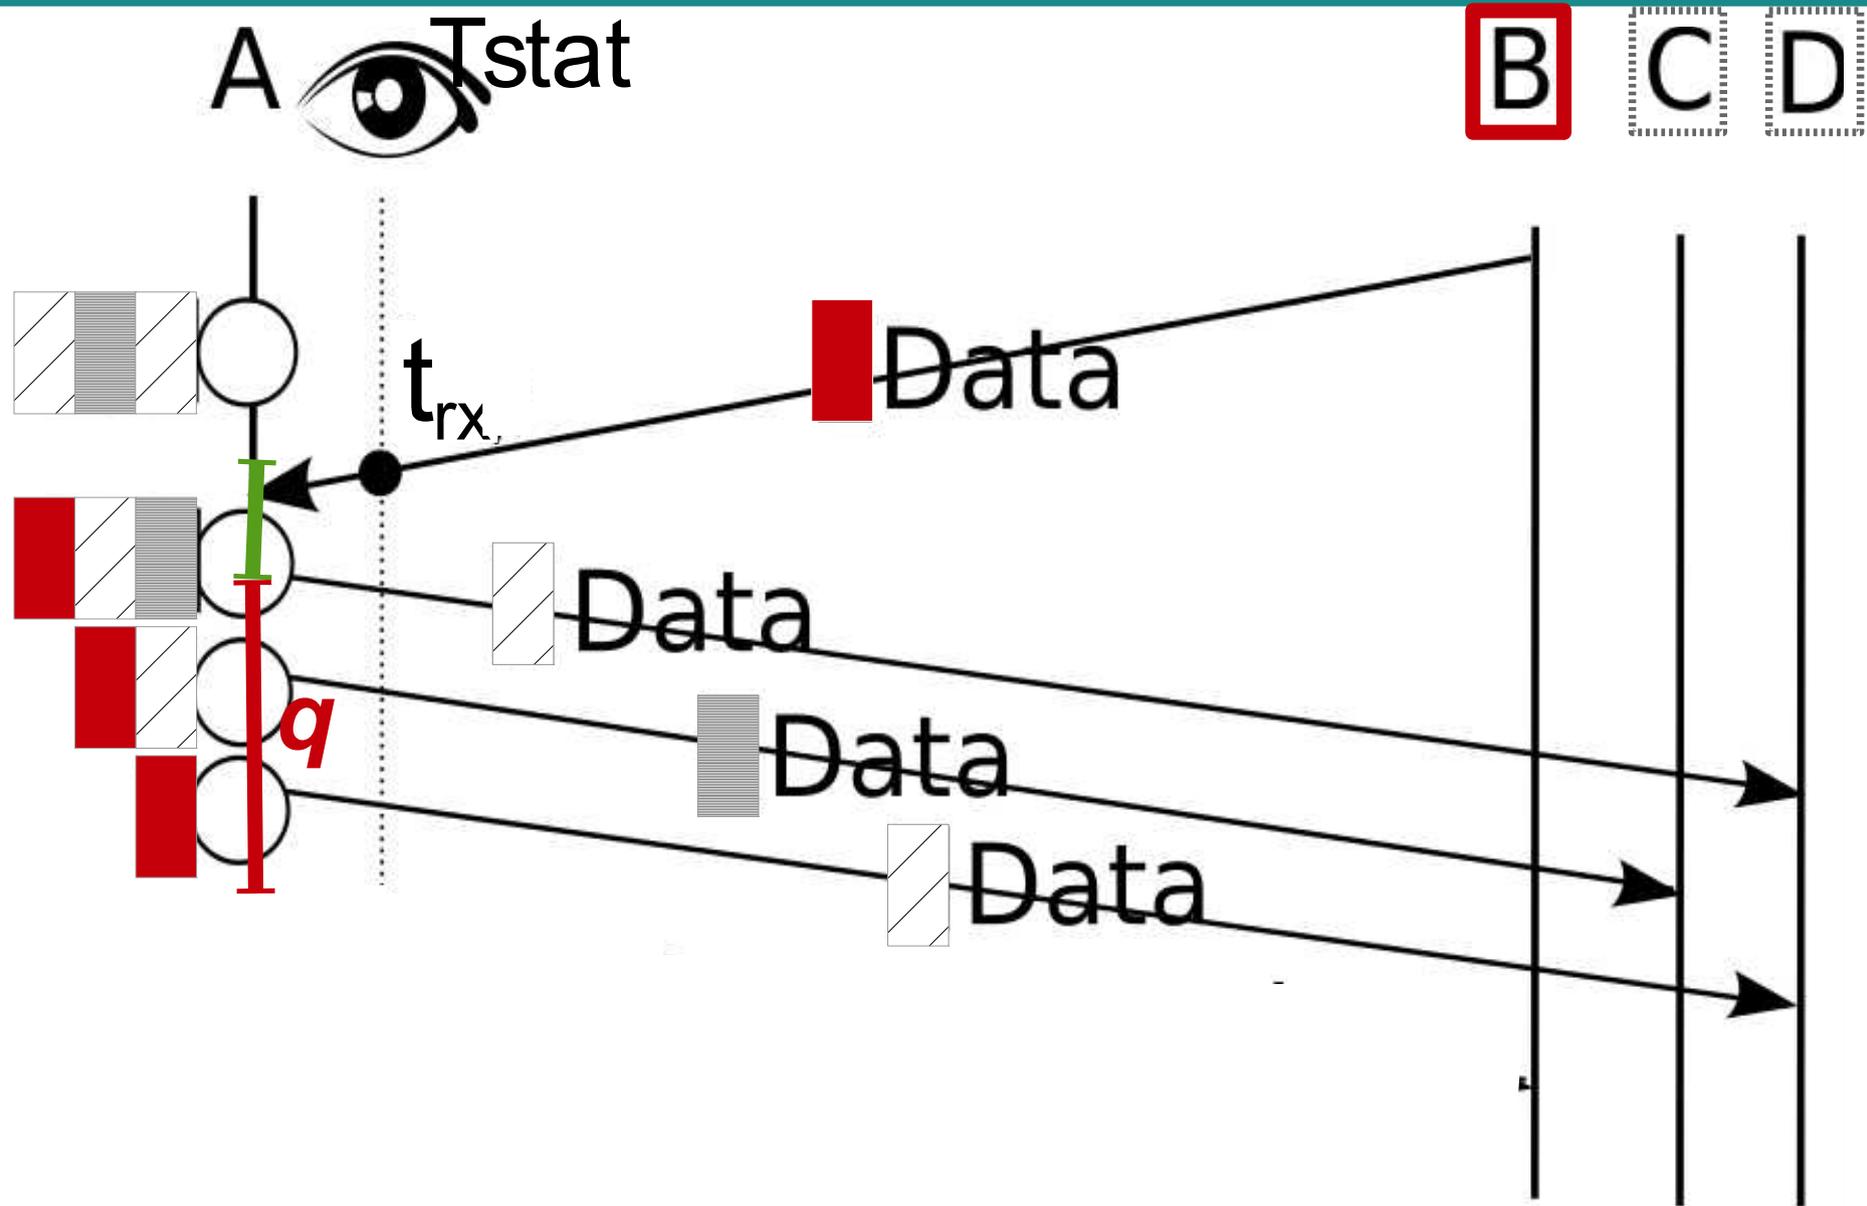
Methodology

Queueing delay inference



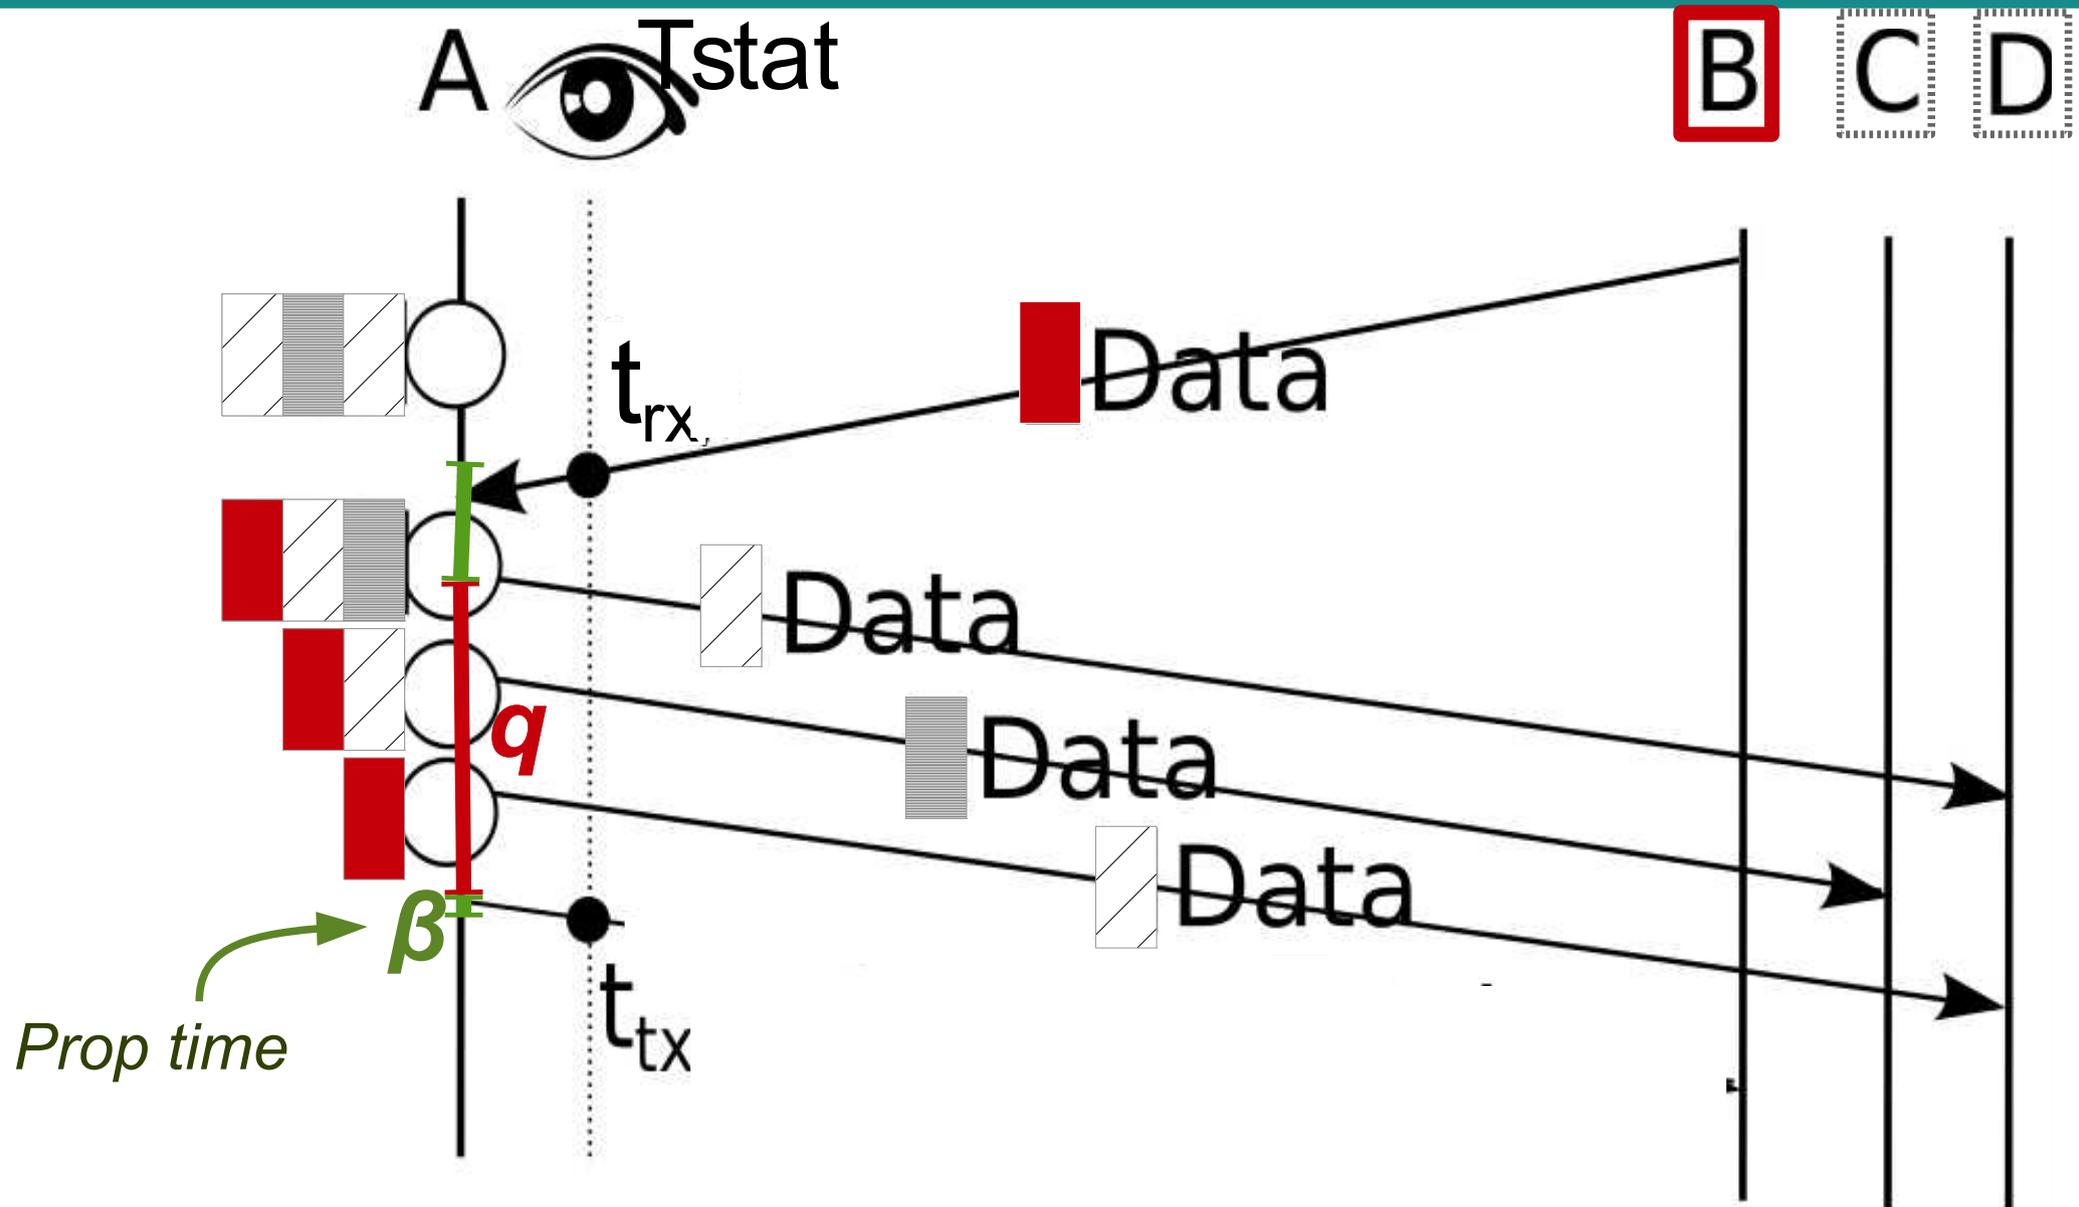
Methodology

Queueing delay inference



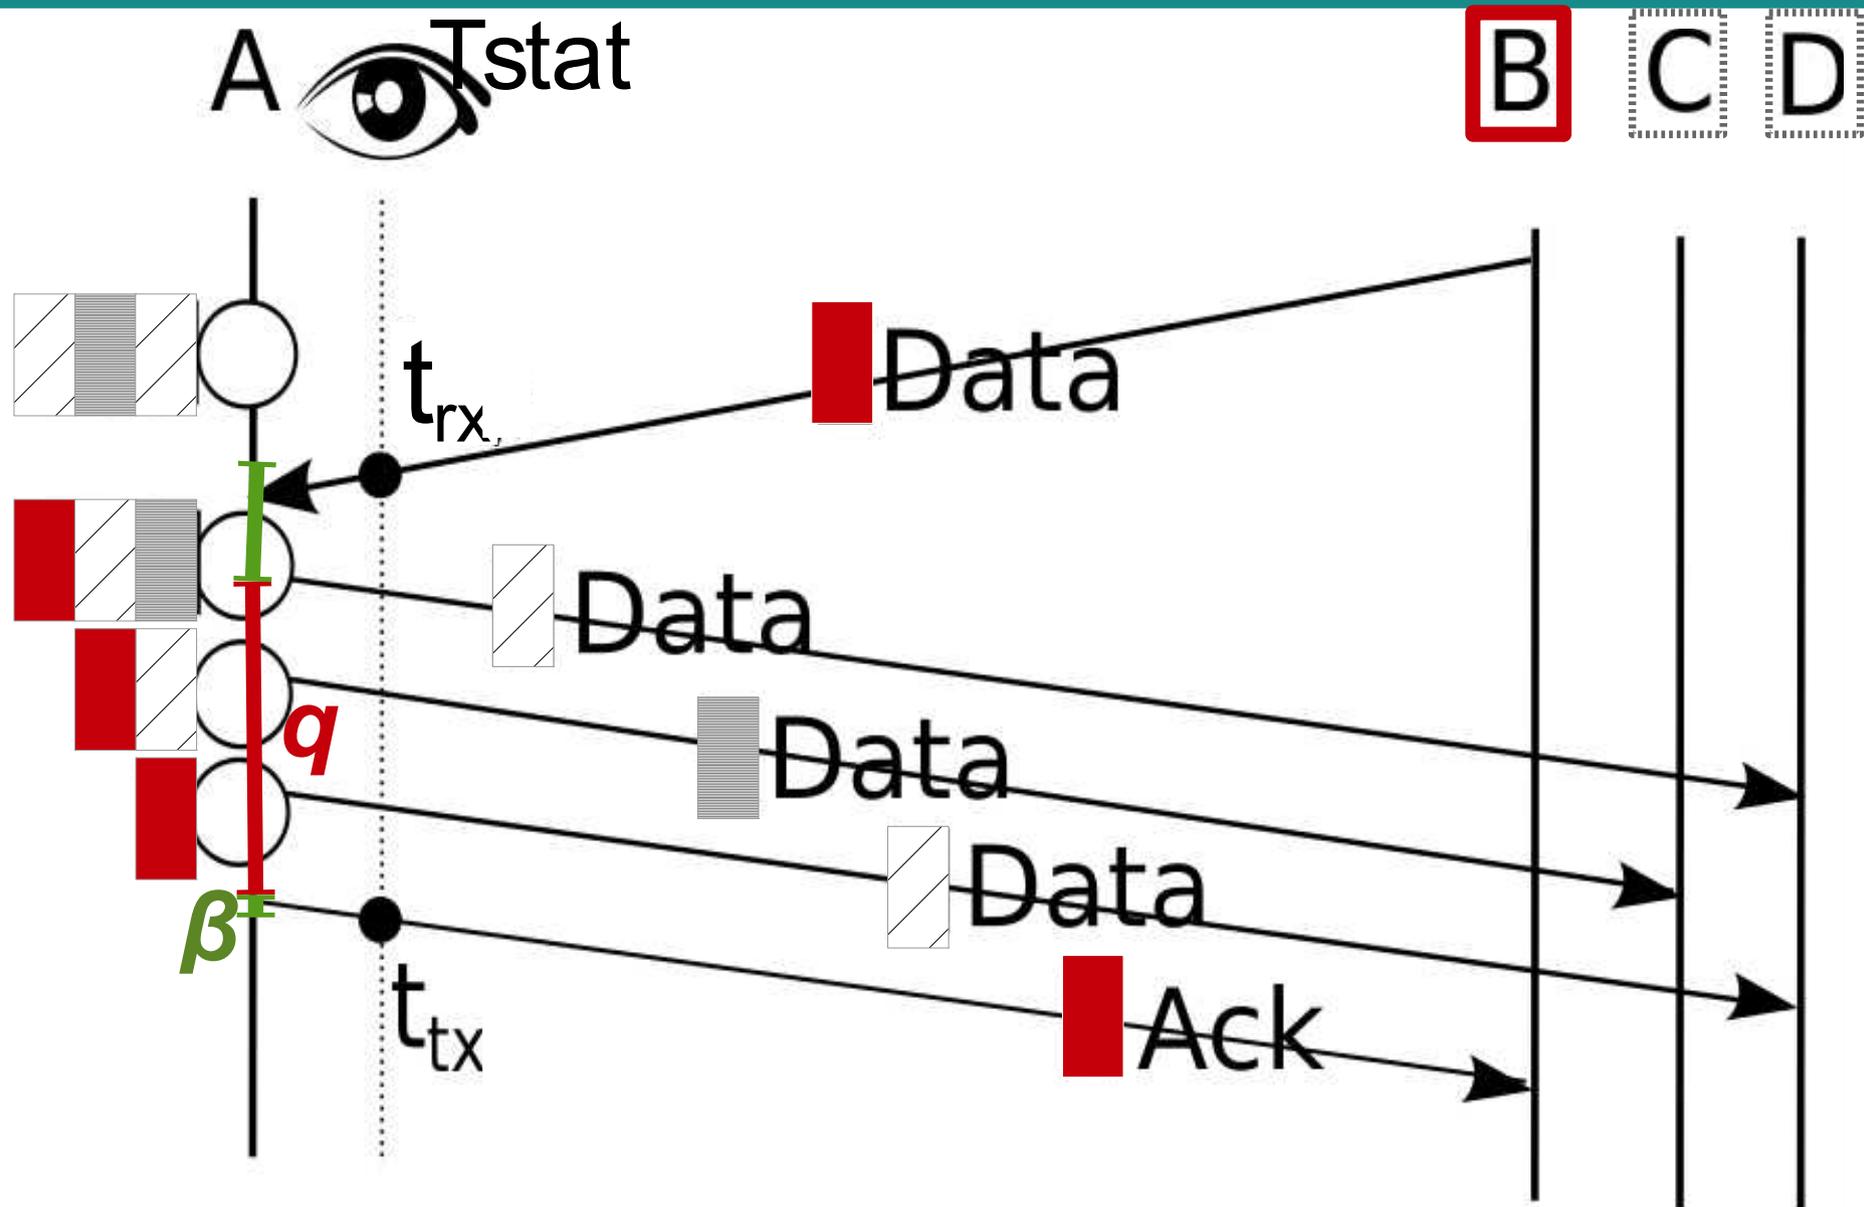
Methodology

Queueing delay inference



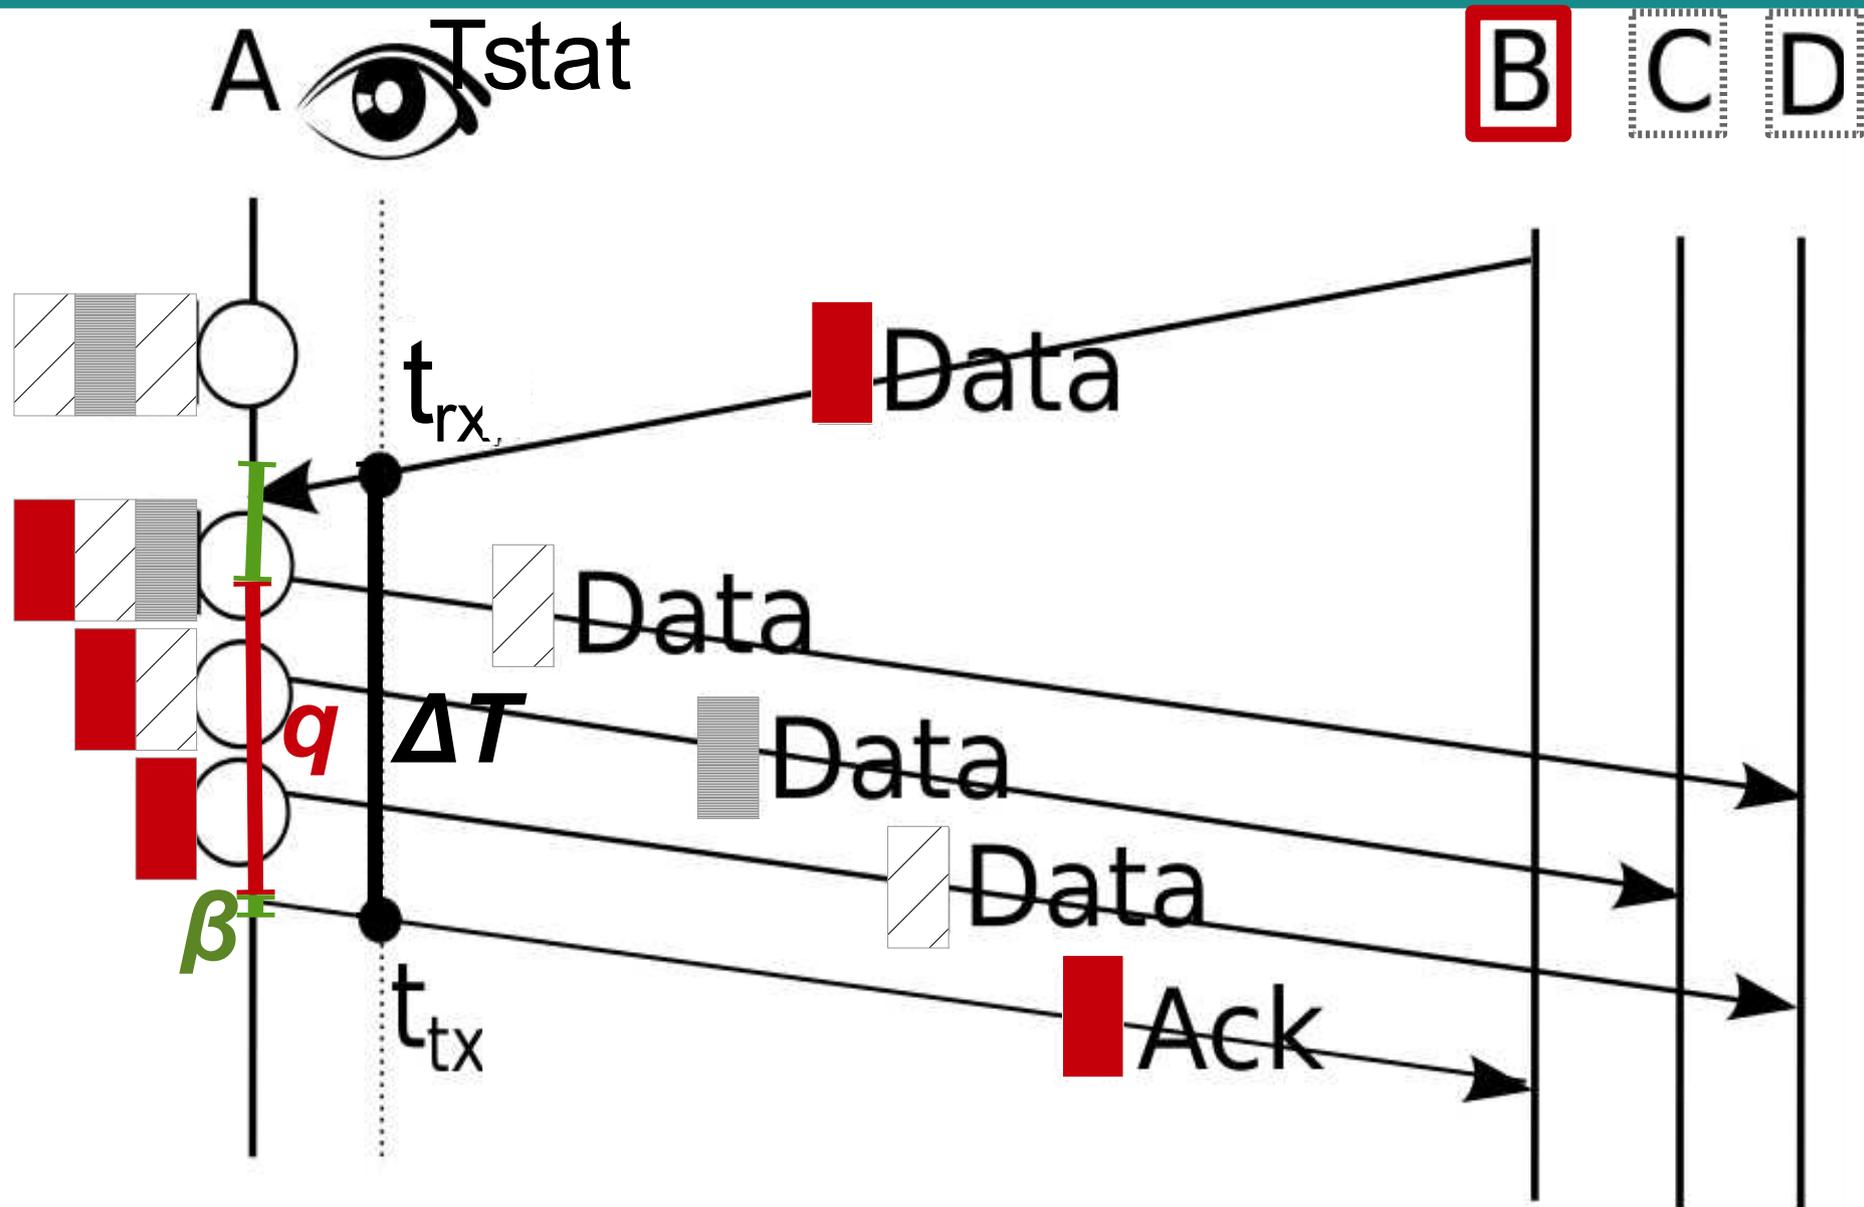
Methodology

Queueing delay inference



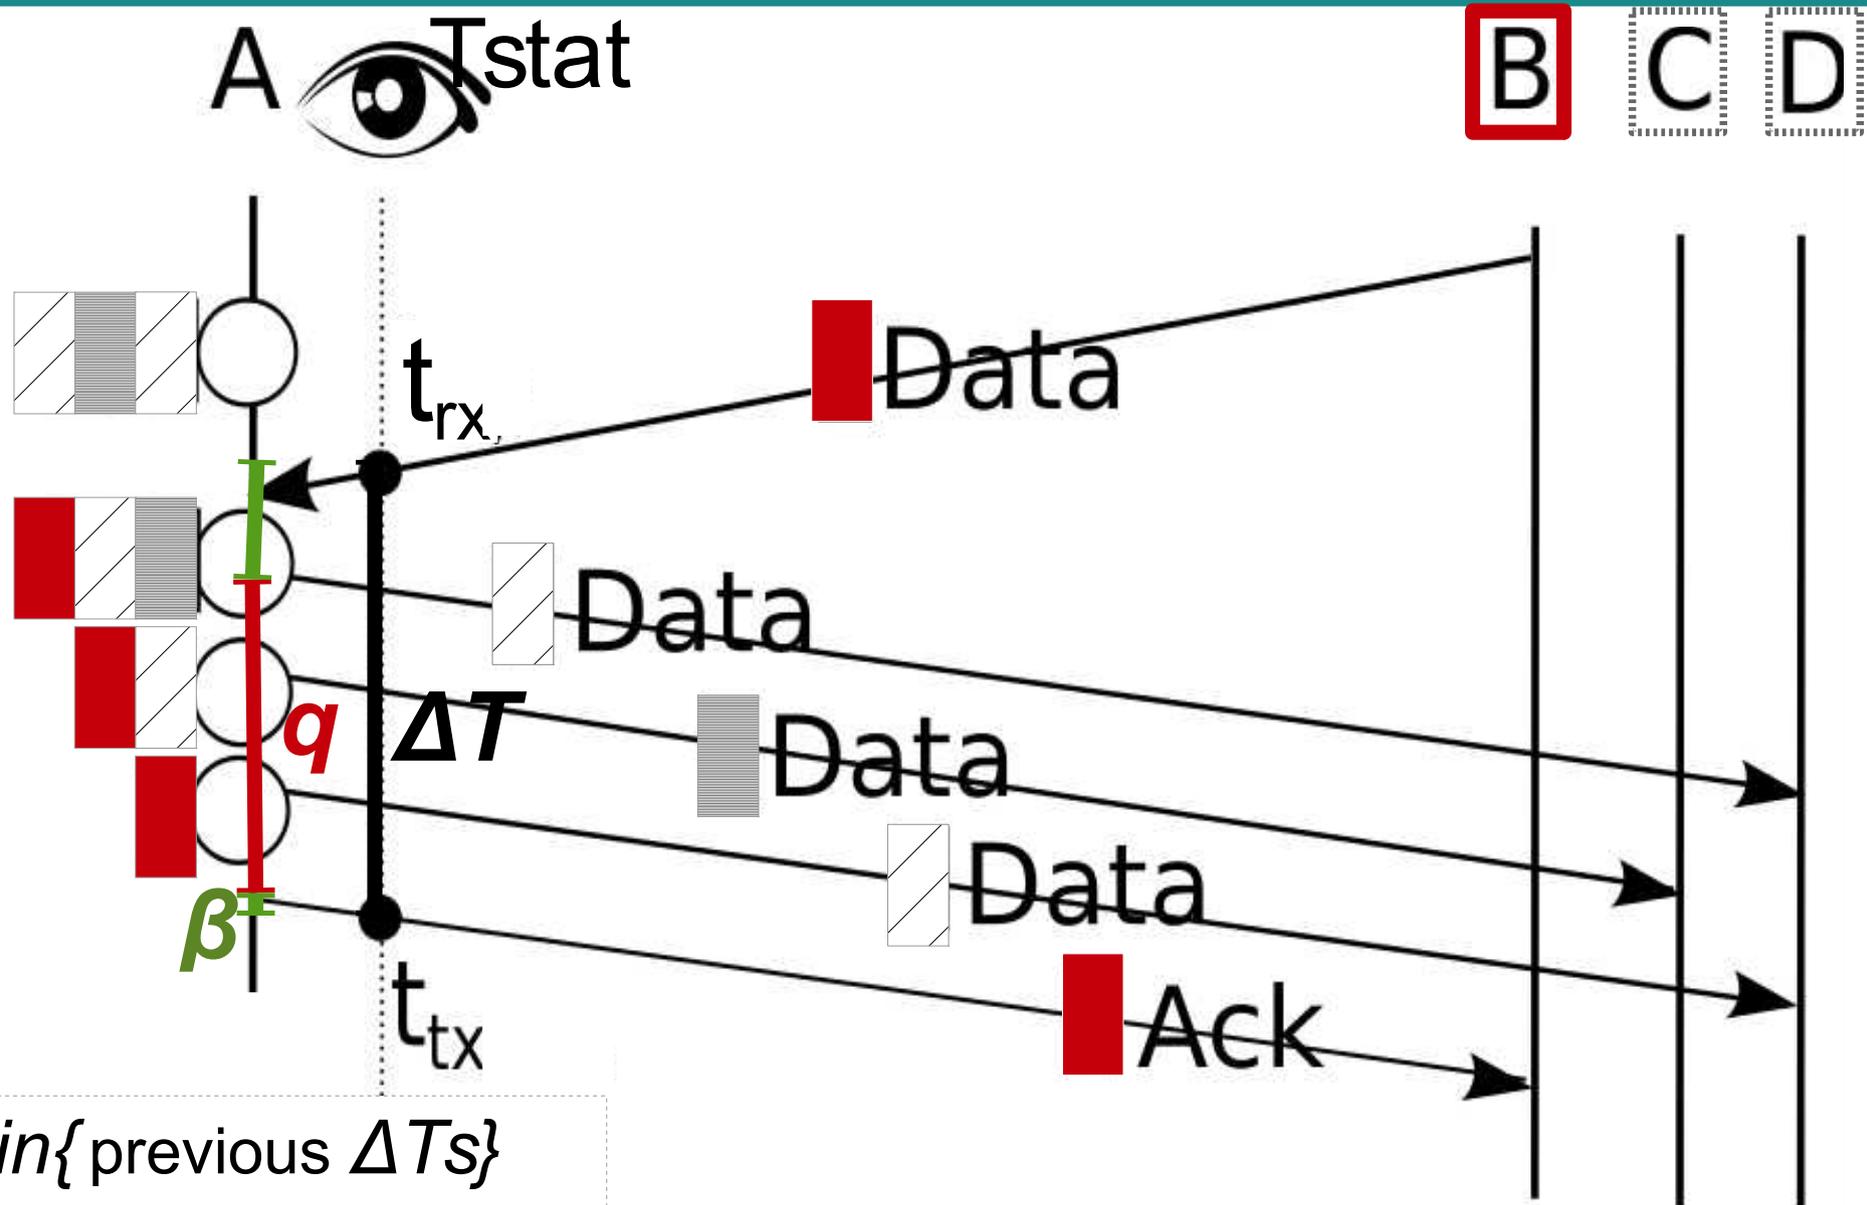
Methodology

Queueing delay inference



Methodology

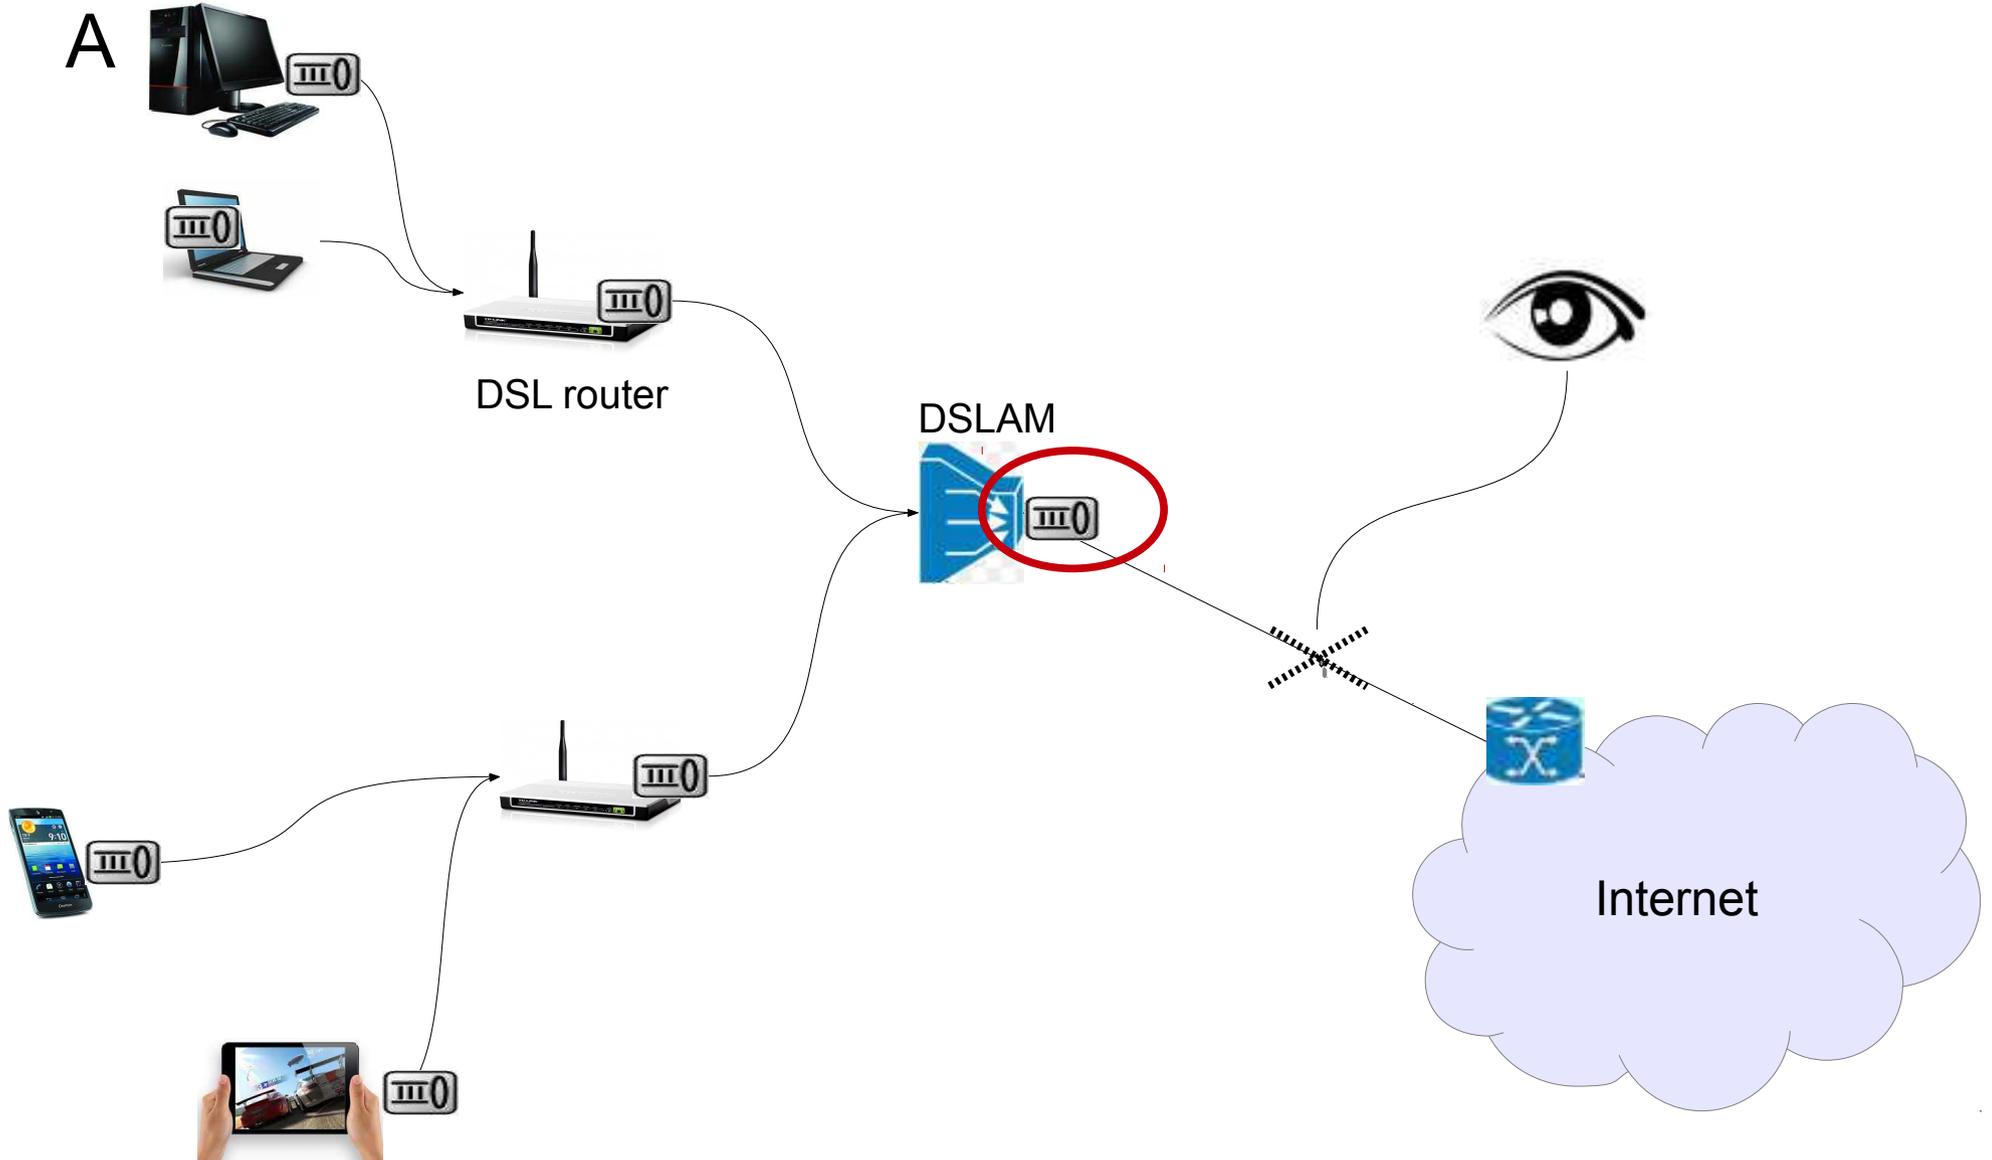
Queueing delay inference



$\beta = \min\{\text{previous } \Delta T\}$
 $q = \Delta T - \beta$

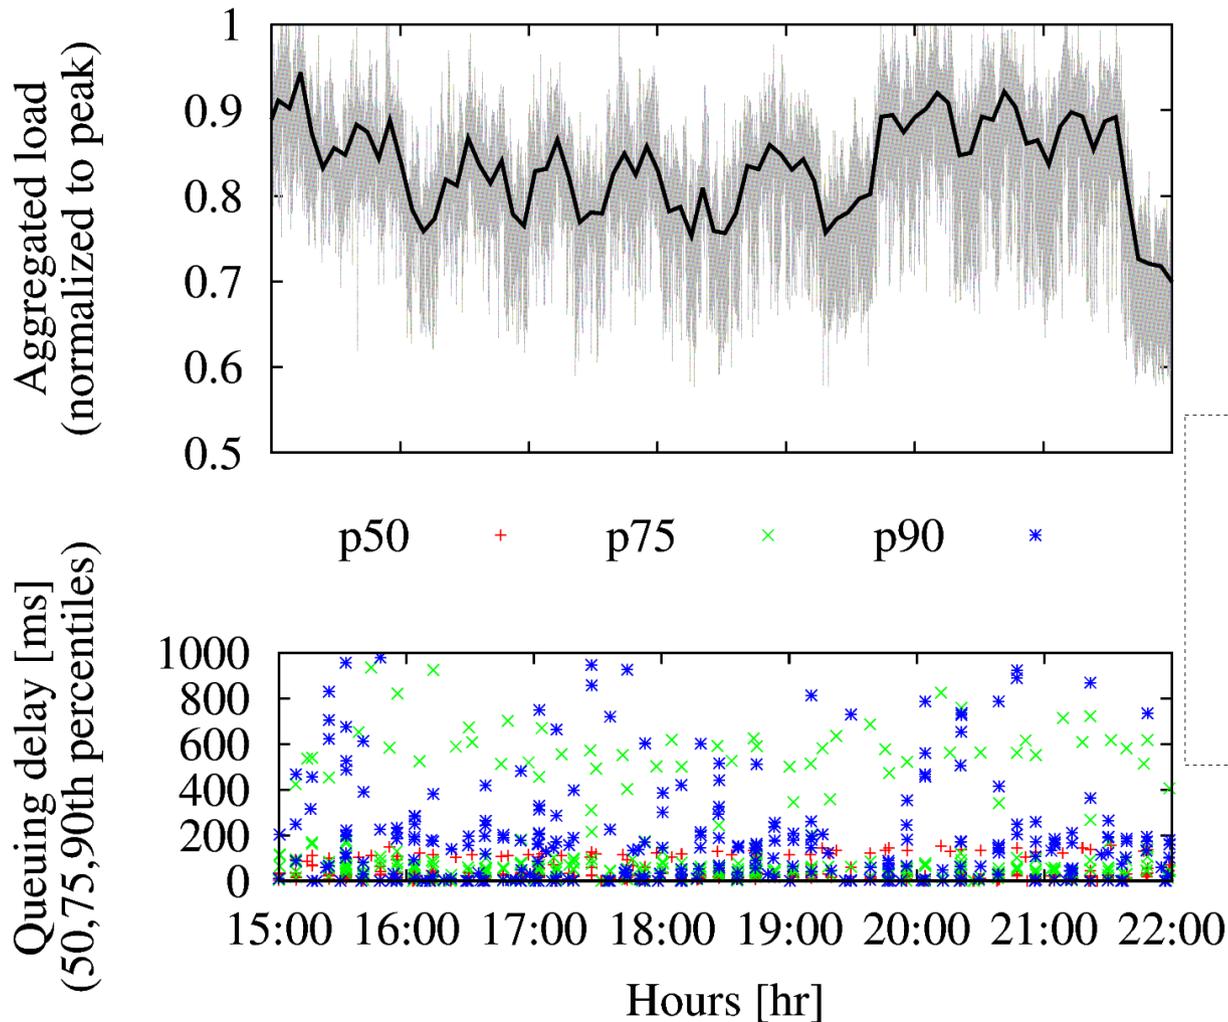
Bufferbloat vs ISP load

Where does bufferbloat happen?



Bufferbloat vs ISP load

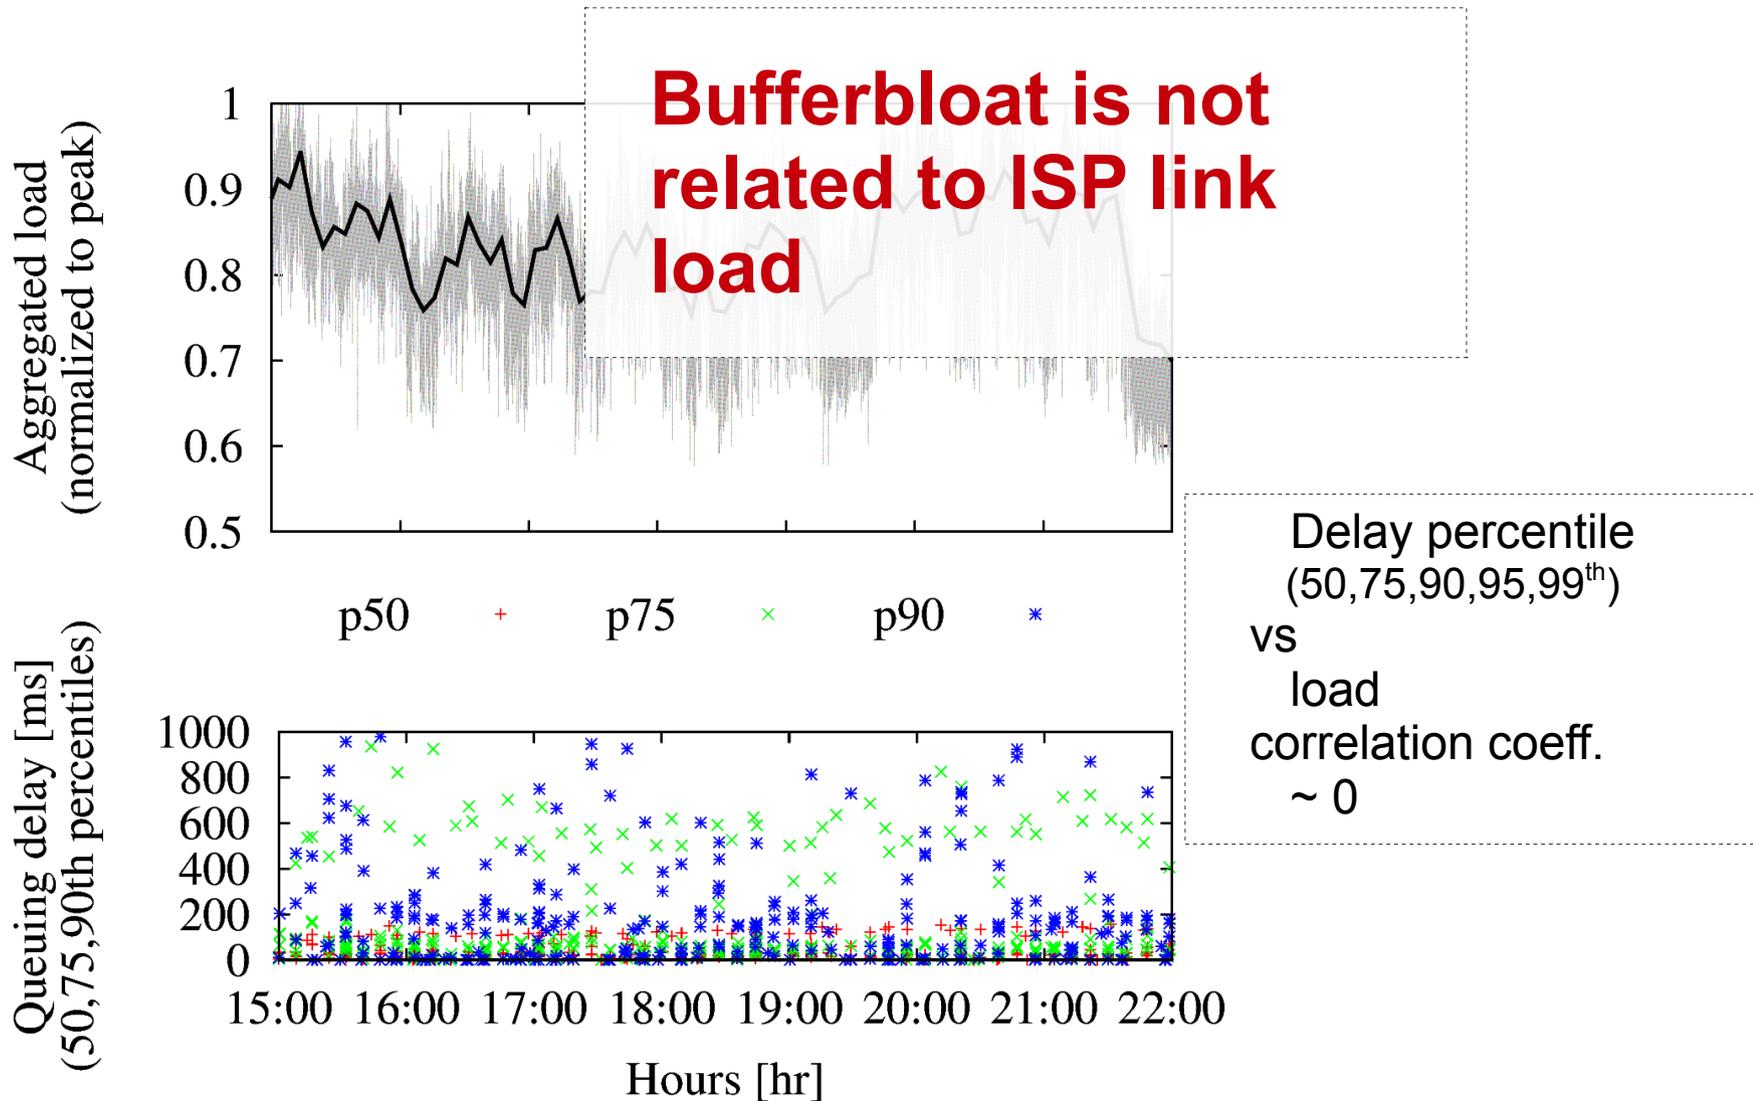
Where does bufferbloat happen?



Delay percentile
(50,75,90,95,99th)
vs
load
correlation coeff.
~ 0

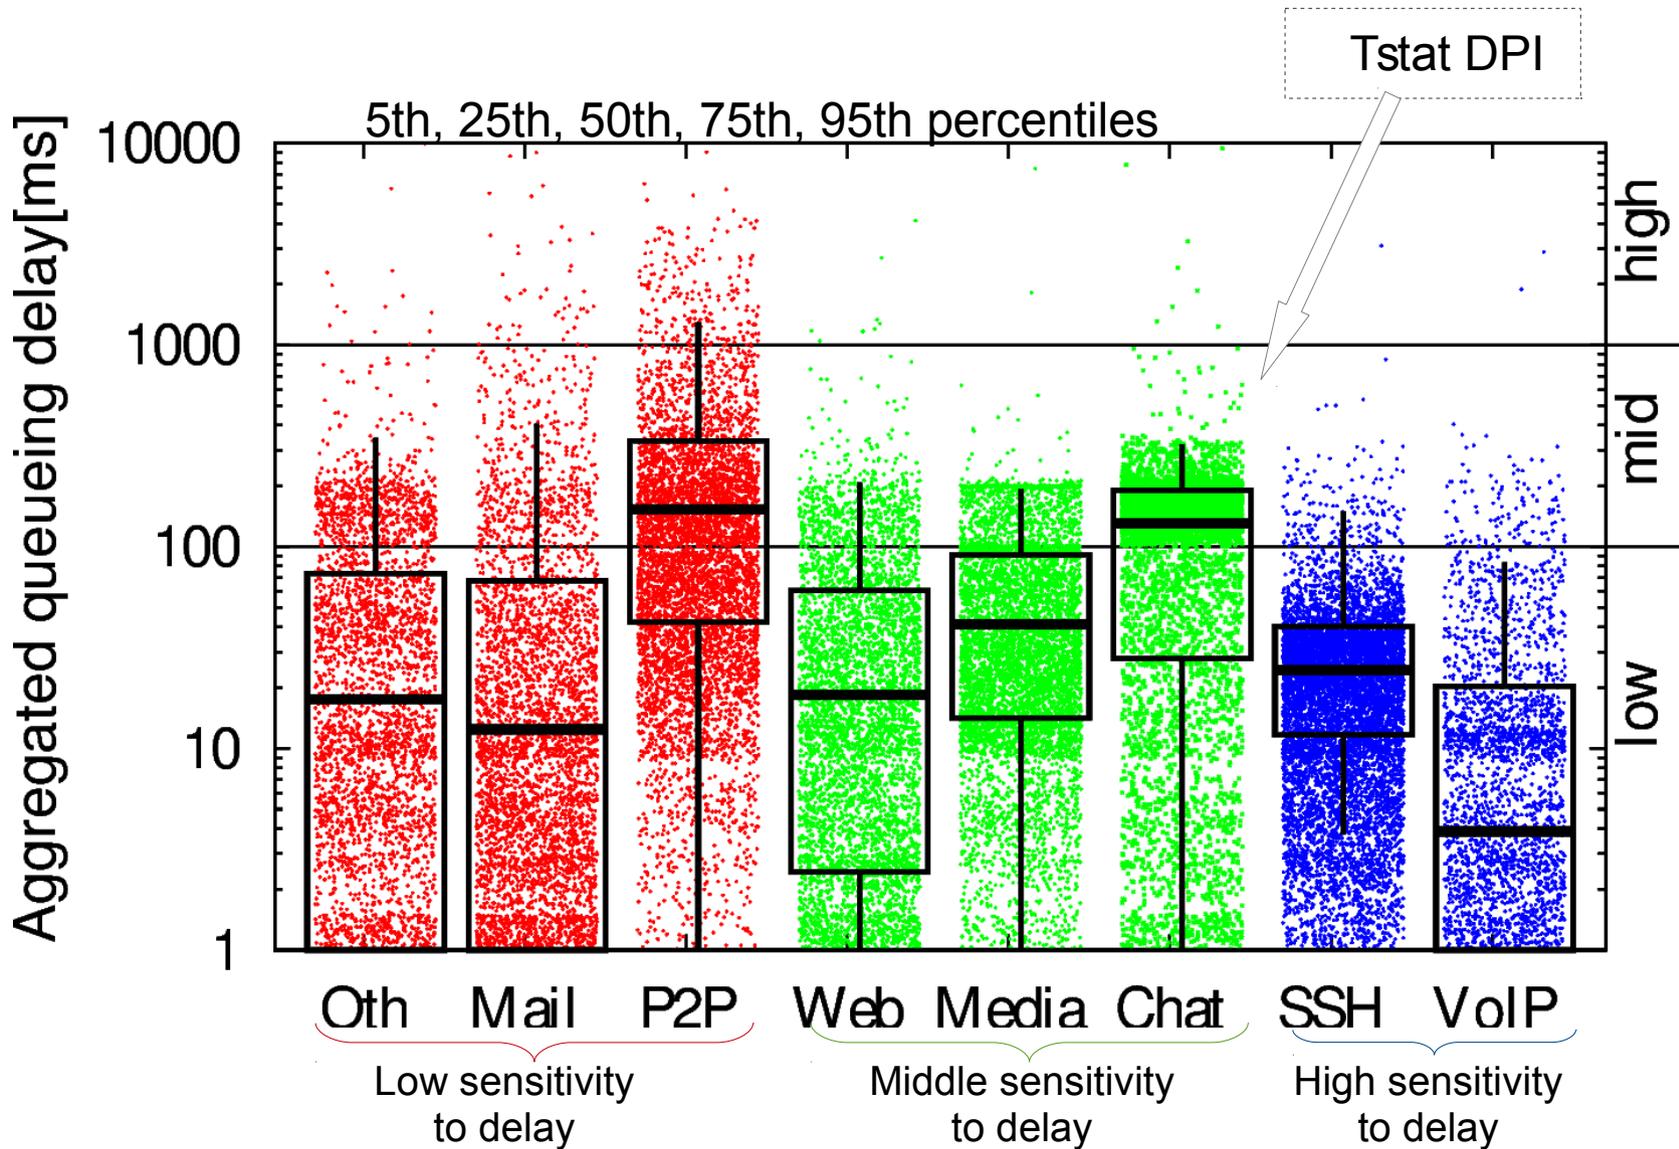
Bufferbloat vs ISP load

Where does bufferbloat happen?



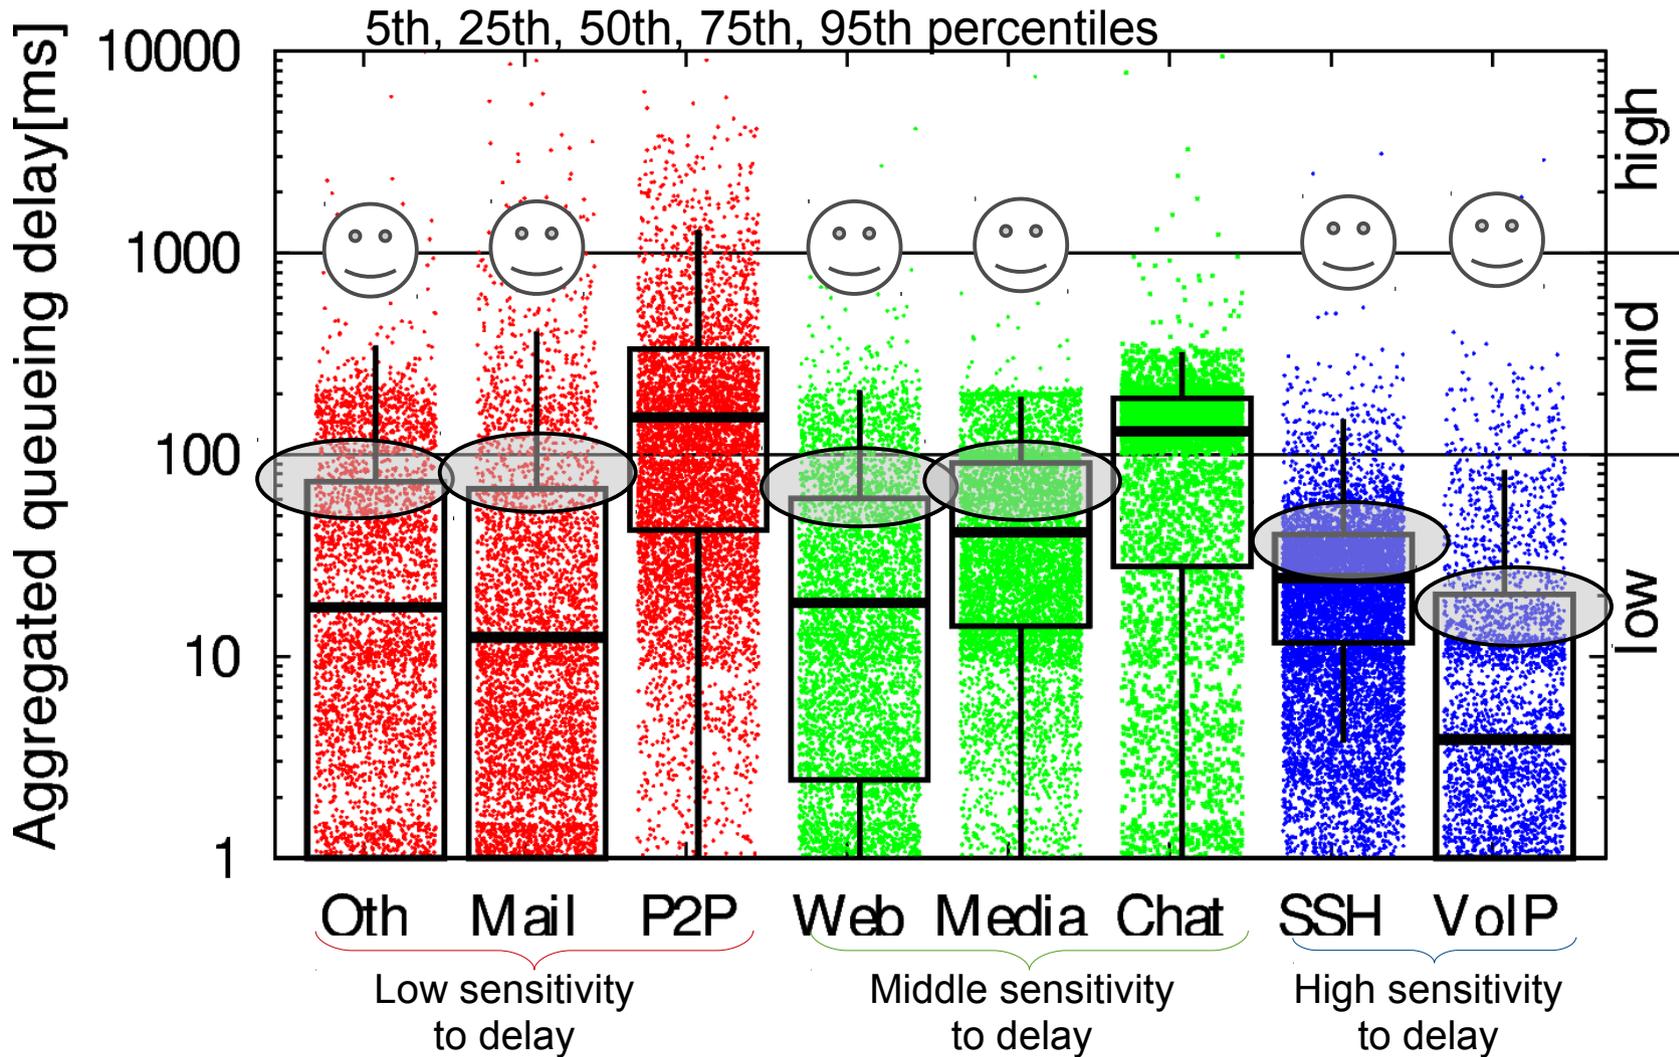
Per application view and QoE

Which applications suffer bufferbloat?



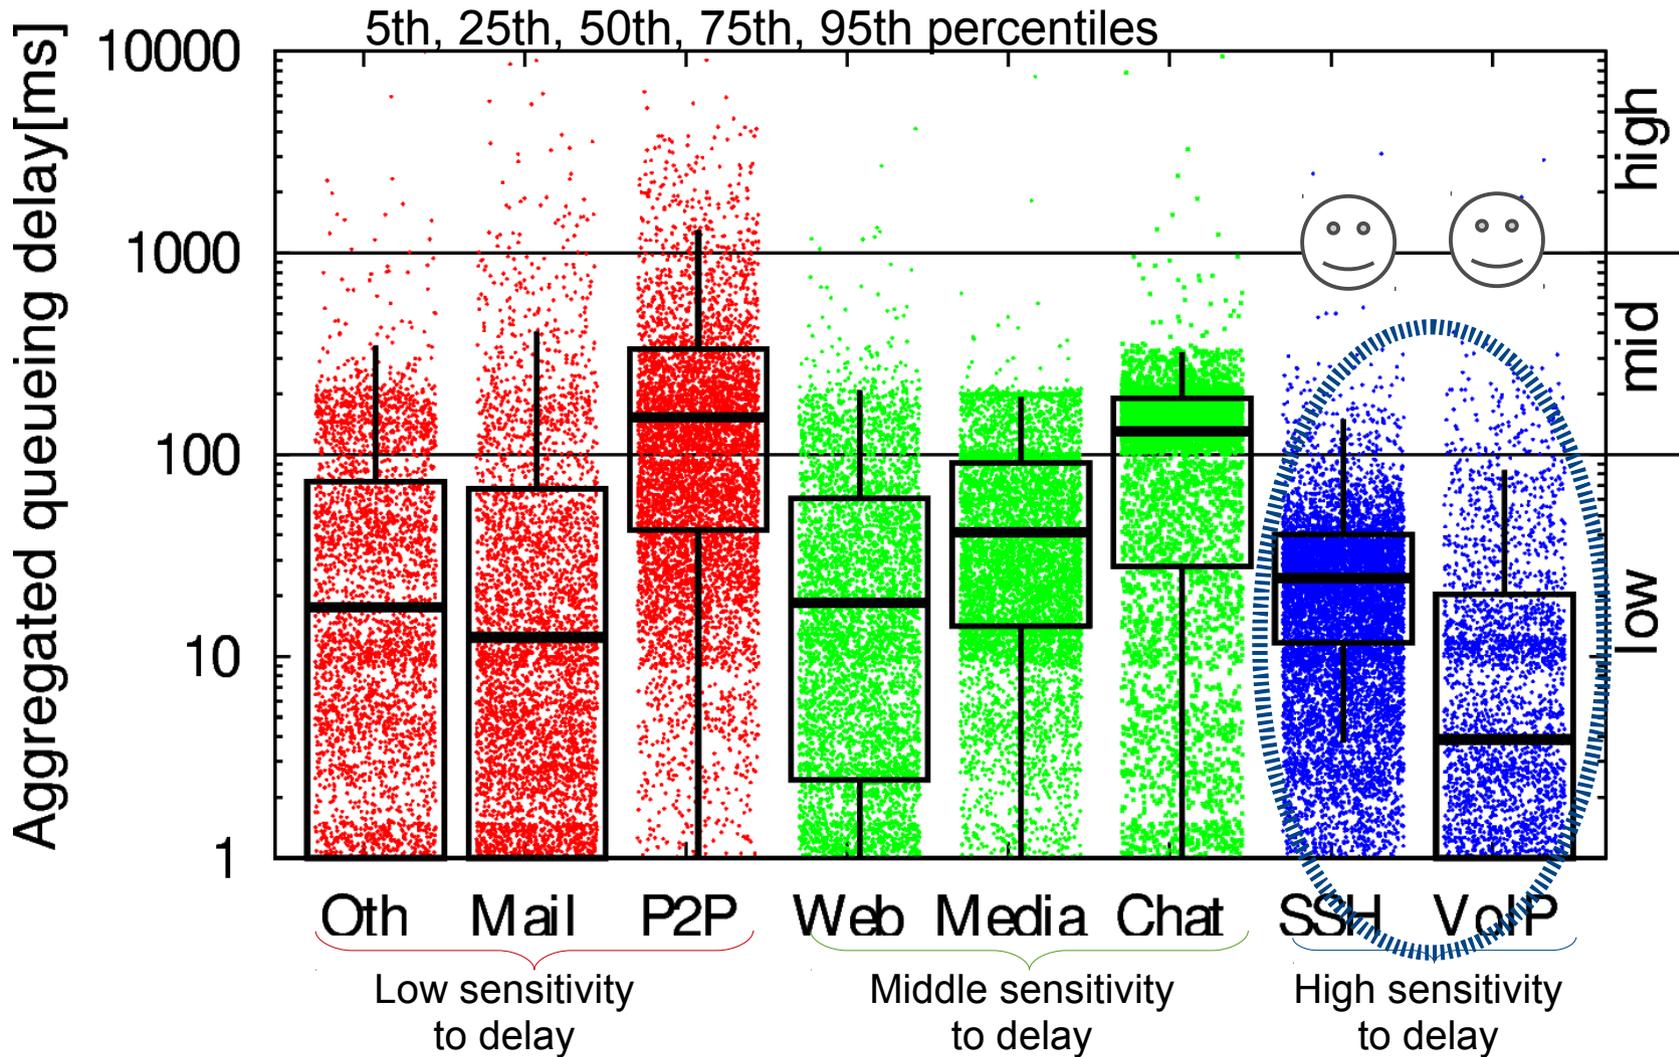
Per application view and QoE

Which applications suffer bufferbloat?



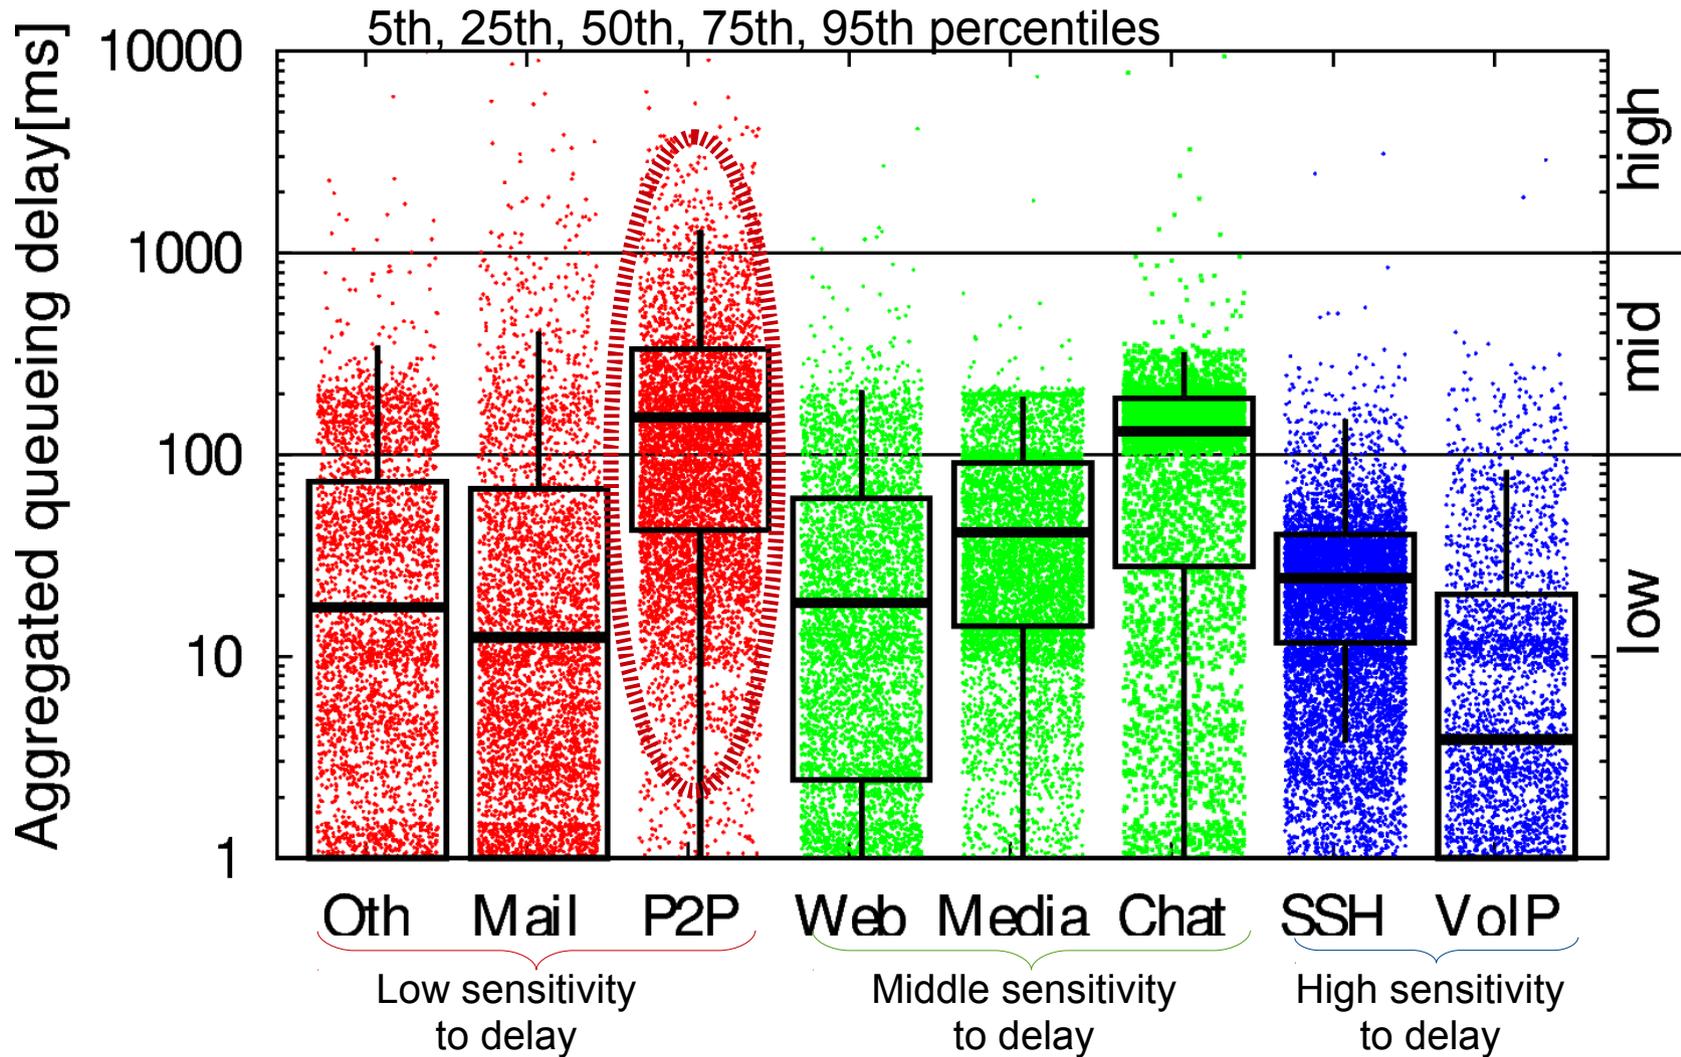
Per application view and QoE

Which applications suffer bufferbloat?



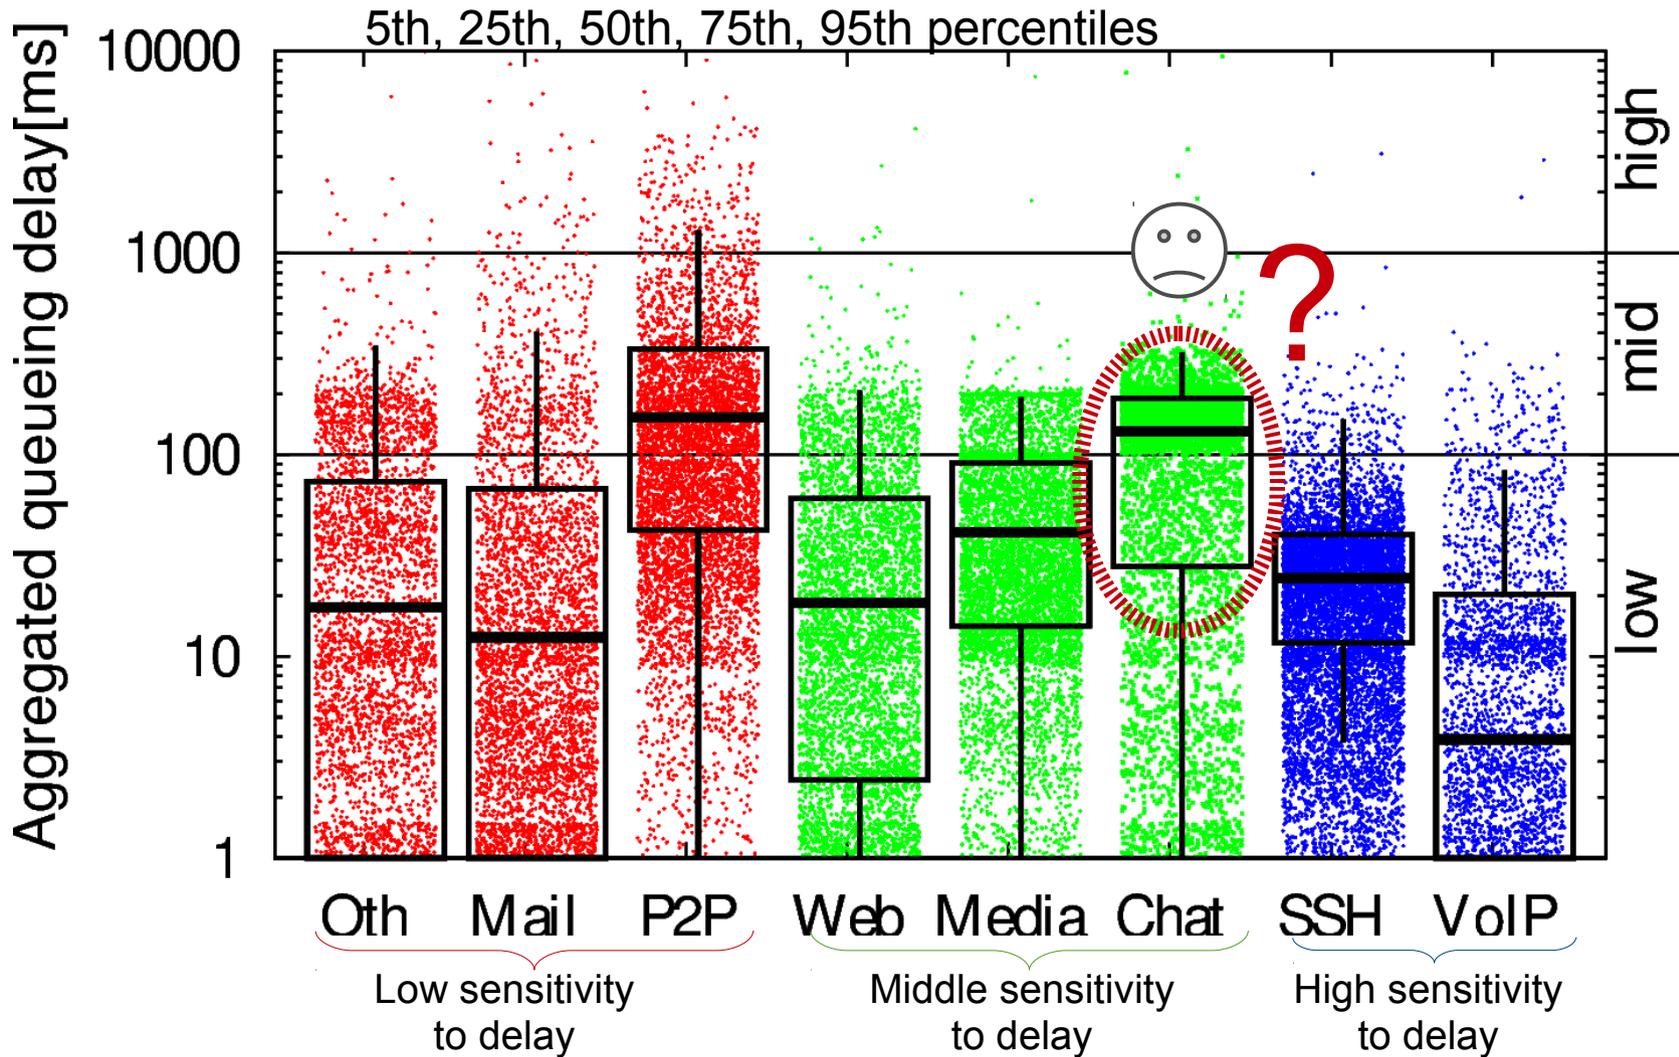
Per application view and QoE

Which applications suffer bufferbloat?



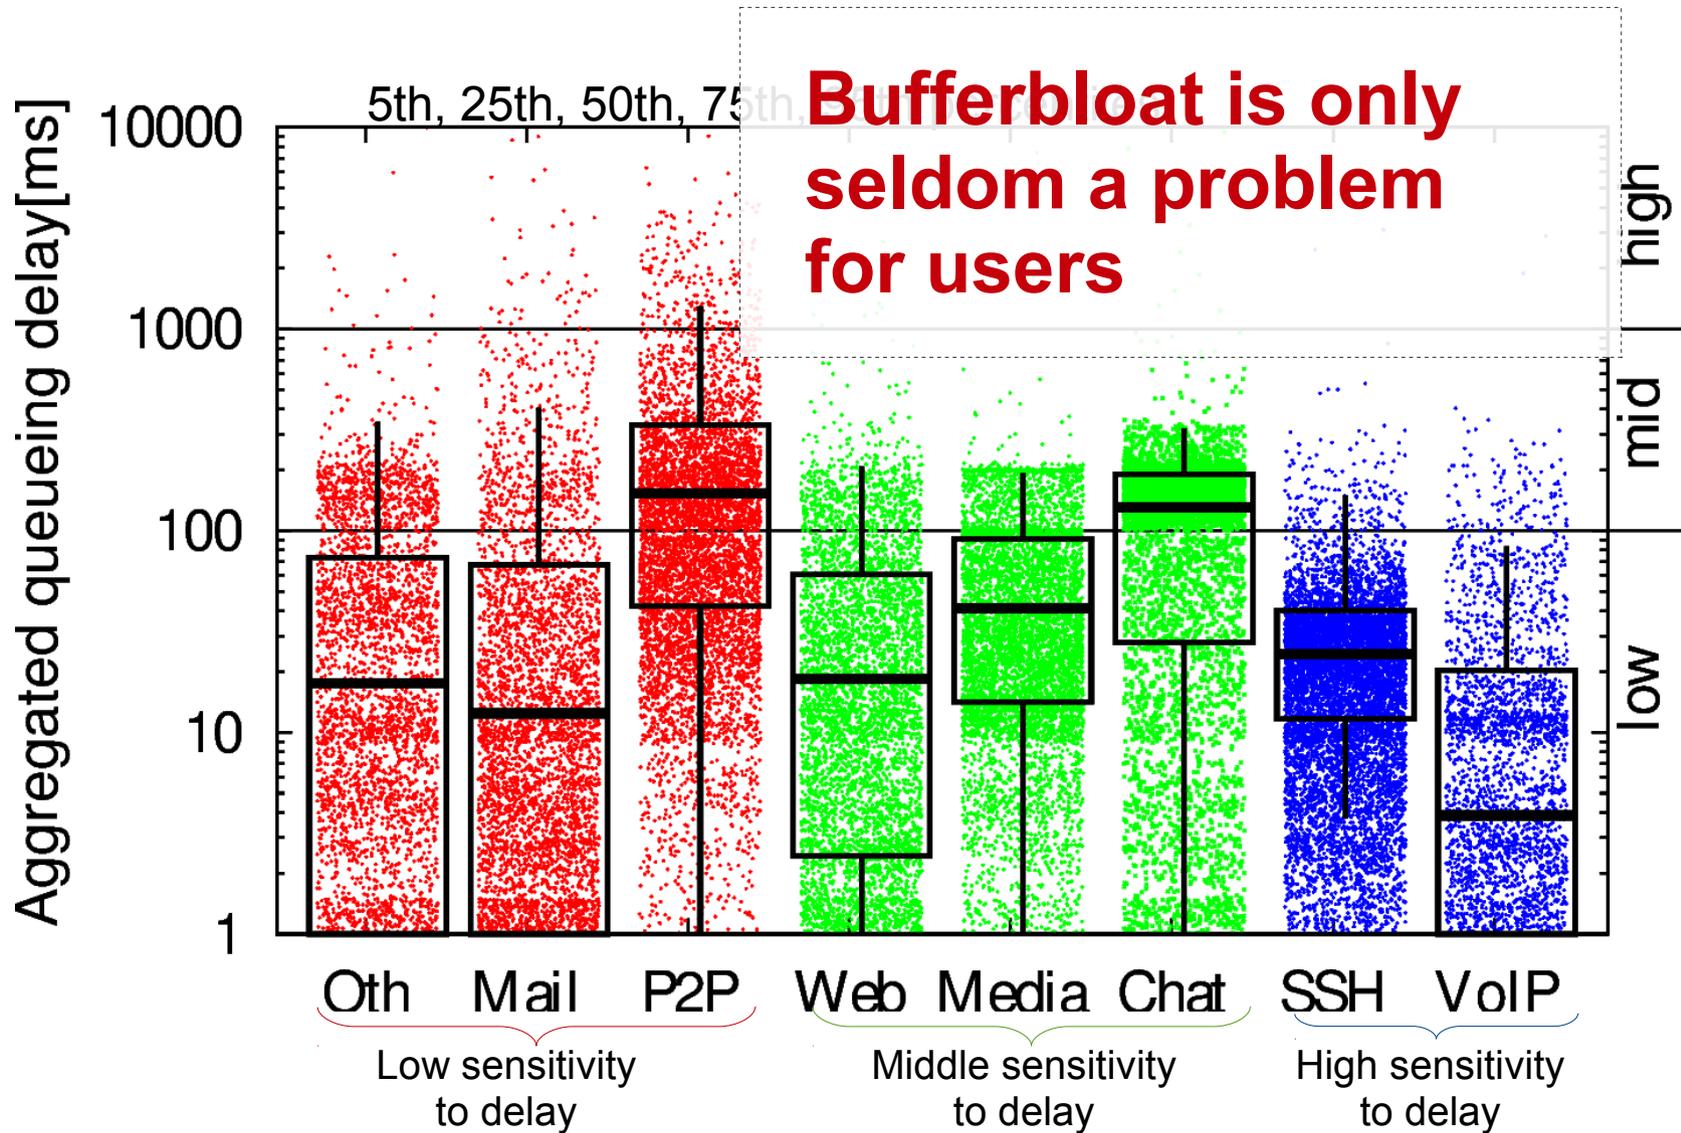
Per application view and QoE

Which applications suffer bufferbloat?



Per application view and QoE

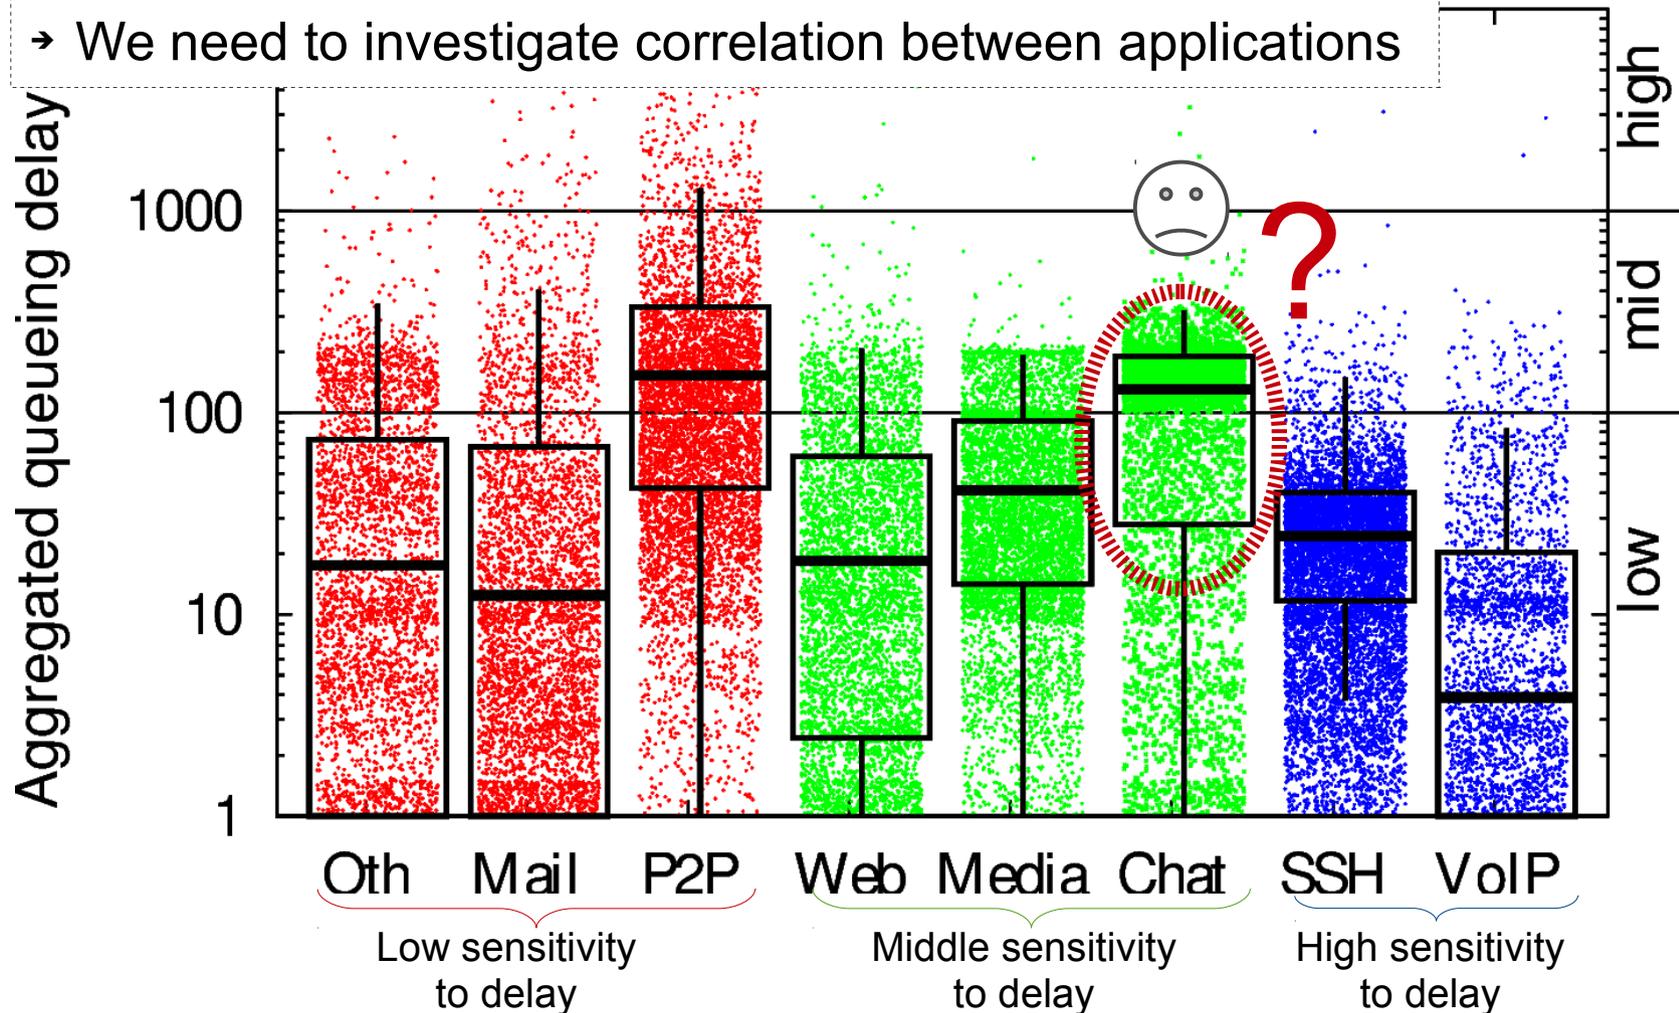
Which applications suffer bufferbloat?



Root cause analysis

Concurrent applications cause queueing delay

→ We need to investigate correlation between applications

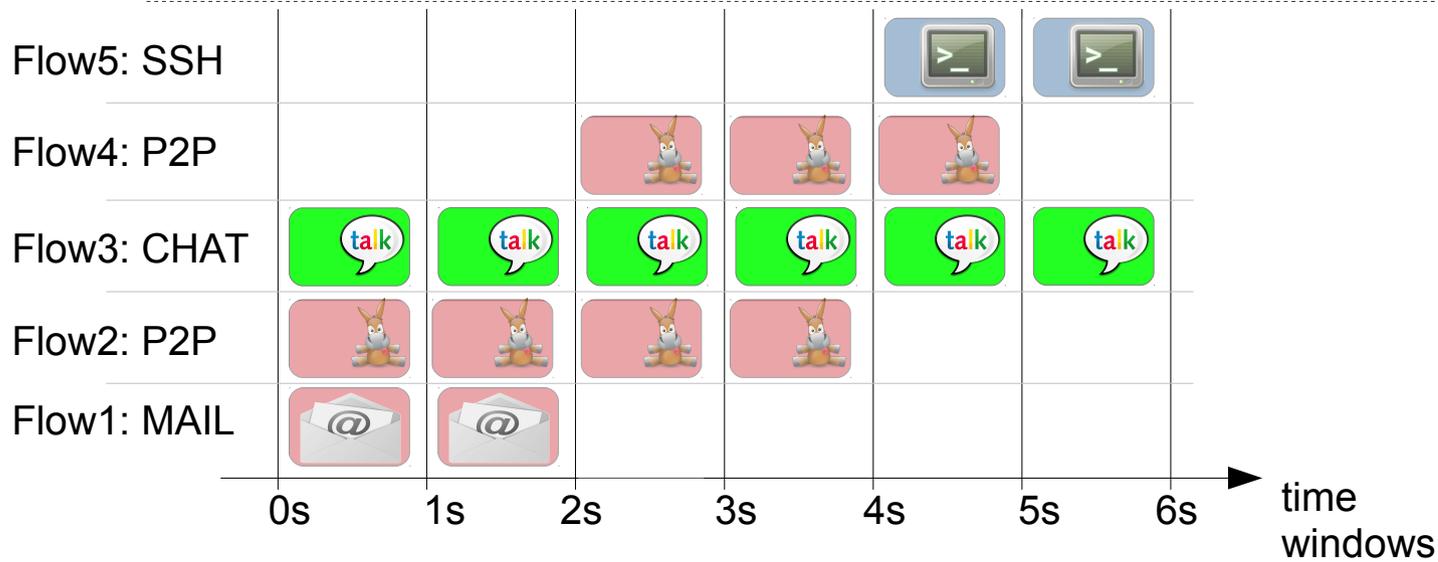


Root cause analysis

Which application mix causes bufferbloat?

Concurrent applications cause queueing delay

→ We need to investigate correlation between applications

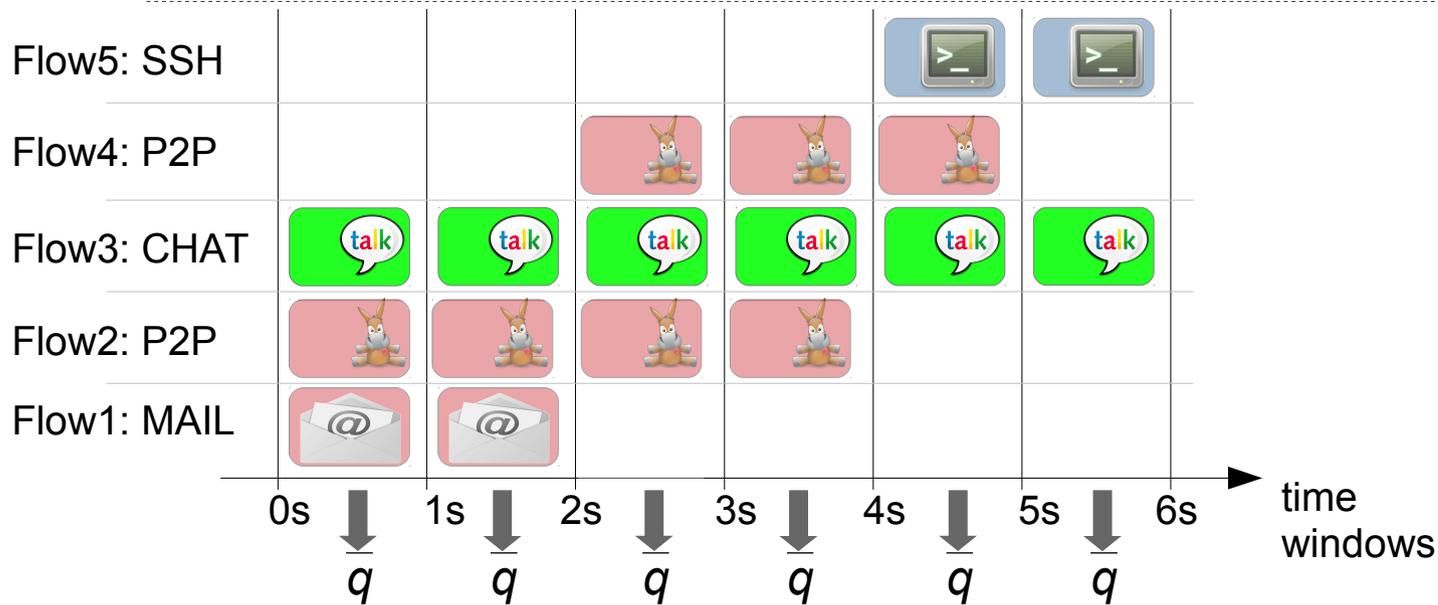


Root cause analysis

Which application mix causes bufferbloat?

Concurrent applications cause queueing delay

→ We need to investigate correlation between applications

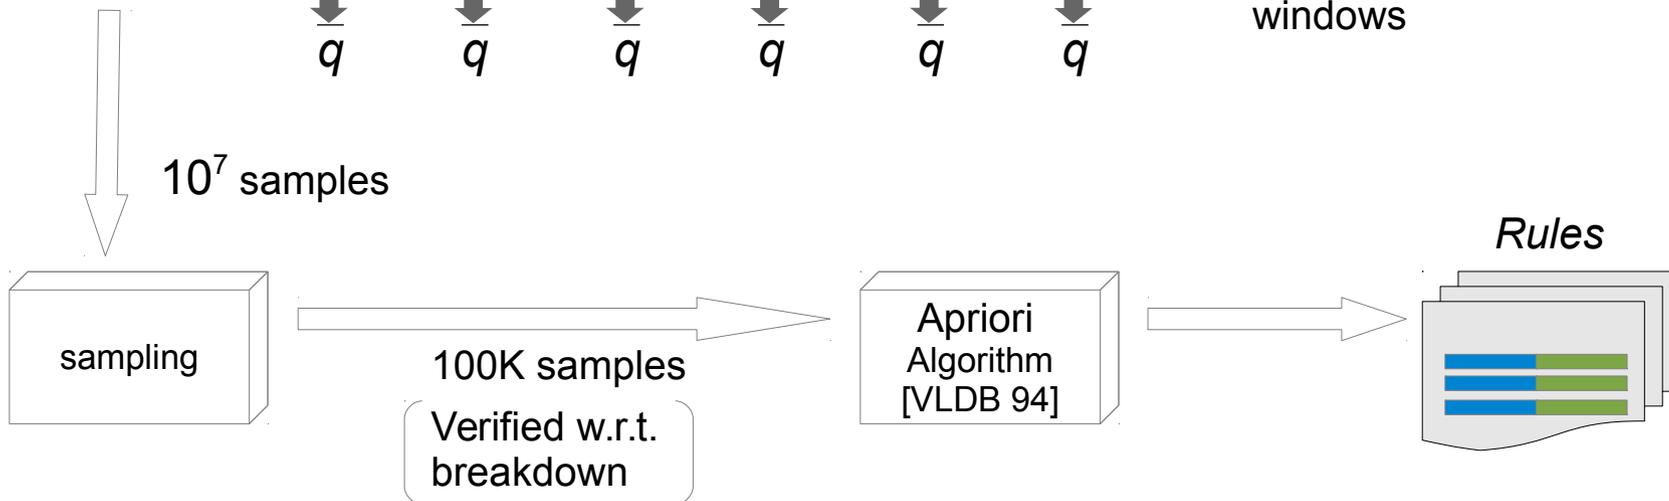
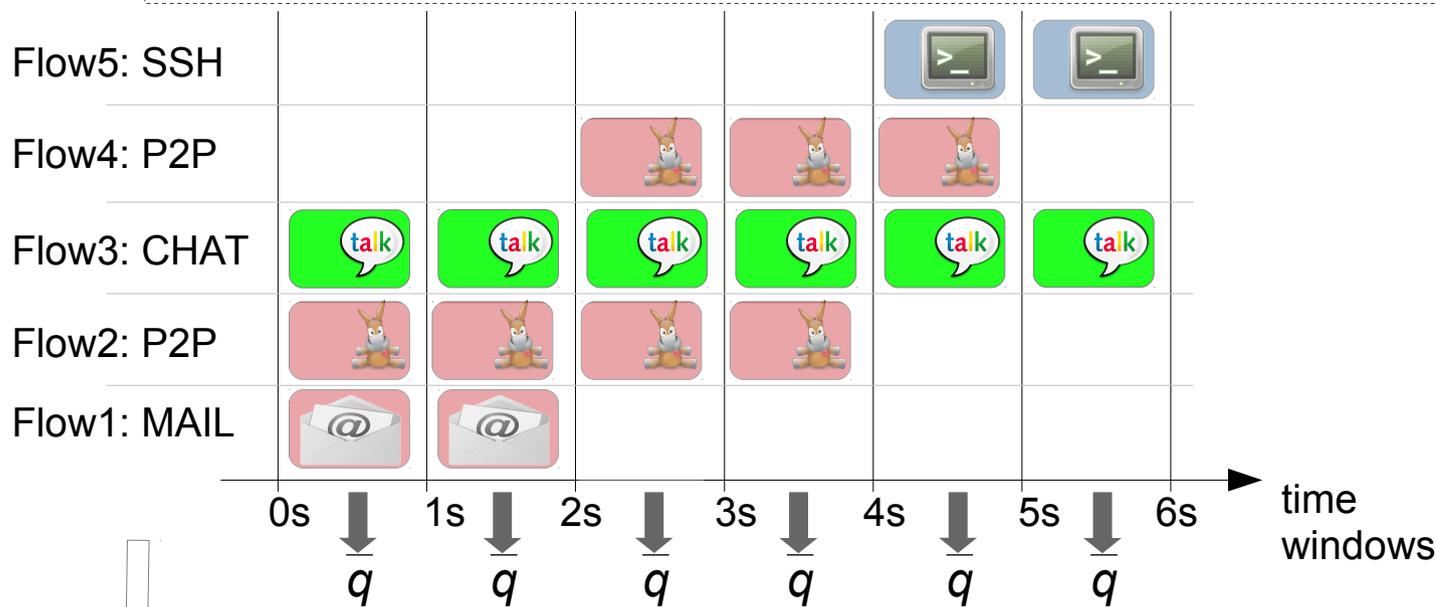


Root cause analysis

Which application mix causes bufferbloat?

Concurrent applications cause queueing delay

→ We need to investigate correlation between applications



Root cause analysis

Which application mix causes bufferbloat?

high

Applications						Supp.	Conf.
Num	Other	Mail	P2P	Web	Chat		

(ref)

1		✓				11	99	(a)
1			++			10	96	(b)
1				✓		10	64	(c)
2			++	✓		20	98	(d)
2	✓		++			28	87	(e)
2	✓			✓		23	80	(f)
2	++		++			14	61	(g)
2			++	++		19	54	(h)
3	✓		++	✓		13	99	(i)

mid

Applications						Supp.	Conf.
Num	Other	Mail	P2P	Web	Chat		

(ref)

1			✓			35	76	(j)
1				++		44	53	(k)
2			✓		✓	10	85	(l)
2	++		✓			11	88	(m)
2			✓	✓		19	85	(n)
2			✓	++		10	77	(o)
2	✓		✓			15	76	(p)
2	++			++		16	72	(q)
2				✓	✓	10	51	(r)
2	✓			++		20	50	(s)

Root cause analysis

Which application mix causes bufferbloat?

high

Applications						Supp.	Conf.	(ref)
Num	Other	Mail	P2P	Web	Chat			

1		✓				11	99	(a)
1			++			10	96	(b)
1				✓		10	64	(c)
2			++	✓		20	98	(d)
2	✓		++			28	87	(e)
2	✓			✓		23	80	(f)
2	++		++			14	61	(g)
2			++	++		19	54	(h)
3	✓		++	✓		13	99	(i)

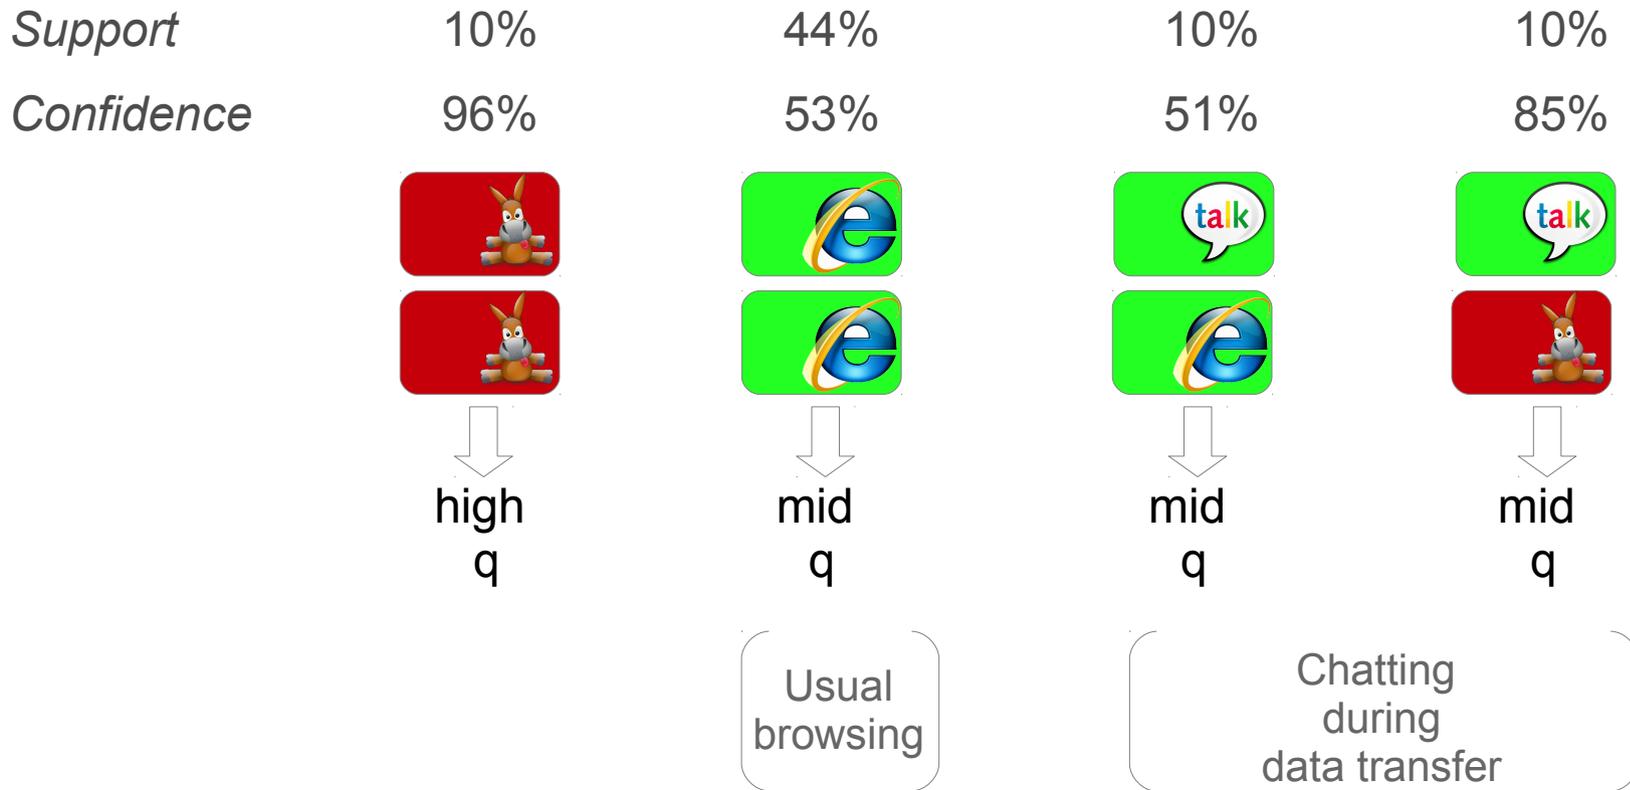
mid

Applications						Supp.	Conf.	(ref)
Num	Other	Mail	P2P	Web	Chat			

1			✓			35	76	(j)
1				++		44	53	(k)
2			✓		✓	10	85	(l)
2	++		✓			11	88	(m)
2			✓	✓		19	85	(n)
2			✓	++		10	77	(o)
2	✓		✓			15	76	(p)
2	++			++		16	72	(q)
2				✓	✓	10	51	(r)
2	✓			++		20	50	(s)

Root cause analysis

Which application mix causes bufferbloat?



Conclusion

- Summary
 - We propose and validate a methodology to infer the queueing delay
 - We implement the methodology in an open source tool
 - We quantify queueing delay per application
 - We apply data mining techniques to infer root causes
- Insights
 - Bufferbloat is not related to ISP link load
 - Bufferbloat is seldom a problem for users
 - Mutual impact of concurrent applications must be considered

References

- [IMC 10] B. Nechaev V. Paxson C. Kreibich, N. Weaver, *Netalyzr: Illuminating the edge network*, Internet Measurement Conference (IMC), 2010.
- [GN11] Jim Gettys and Kathleen Nichols, *Bufferbloat: Dark buffers in the internet*, Queue, 9(11): 40, 2011.
- [IMC 12] R. Teixeira C. Kreibich M. Allman N. Weaver M. Dhawan, J. Samuel and V. Paxson. Fathom. a browser-based network measurement platform. ACM IMC, 2012.
- [IMC 12 a] K. Lee H. Jiang, Y. Wang and I. Rhee. Tackling bufferbloat in 3g/4g networks. ACM IMC, 2012
- [SIGCOMM 11] Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescapè. Broadband internet performance: A view from the gateway. Proc. ACM SIGCOMM 2011
- [SIGCOMM 11 a] Zachary S. Bischof, John S. Otto, Mario A. Sánchez, John P. Rula, David R. Choffnes, and Fabián E. Bustamante. Crowdsourcing isp characterization to the network edge. In Proceedings of the first ACM SIGCOMM workshop on Measurements up the stack, USA, 2011.
- [CCR 12] Zachary S. Bischof, John S. Otto, and Fabián E. Bustamante. Up, down and around the stack: Isp characterization from network intensive applications. SIGCOMM Comput. Commun. Rev., 42(4)
- [PAM 13] Sriharsha Gangam, Jaideep Chandrashekar, Ítalo Cunha, and Jim Kurose. Estimating tcp latency approximately with passive measurements. In Passive and Active Measurement, pages 83–93. Springer, 2013.
- [CCR 12 a] Mark Allman. Comments on bufferbloat. ACM SIGCOMM Comput. Communi. Rev., 43(1): 30–37, 2012.
- [PAM 13 a] C Chirichella, D Rossi, C Testa, T Friedman, and Antonio Pescapè. Remotely gauging upstream bufferbloat delays. In Passive and Active Measurement, pages 250–252. Springer, 2013.
- [TMA 13] Chiara Chirichella and Dario Rossi. To the moon and back: are internet bufferbloat delays really that large? In IEEE INFOCOM Workshop on Traffic Measurement and Analysis, 2013
- [TMA 14] Stéphane Wustner, Renata Texeira and Jaideep Chandrashekar. Characterizing Bufferbloat and its Impact at End-hosts, Traffic Measurement and Analysis, 2014
- [VLDB 94] R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In VLDB, 1994.

End