Large-Scale Scanning of TCP's Initial Window

Jan Rüth, Christian Bormann, Oliver Hohlfeld

http://comsys.rwth-aachen.de/

London / IMC, November 2017

COM SYS | RWTH AACHEN UNIVERSITY

Why look at Initial Windows?

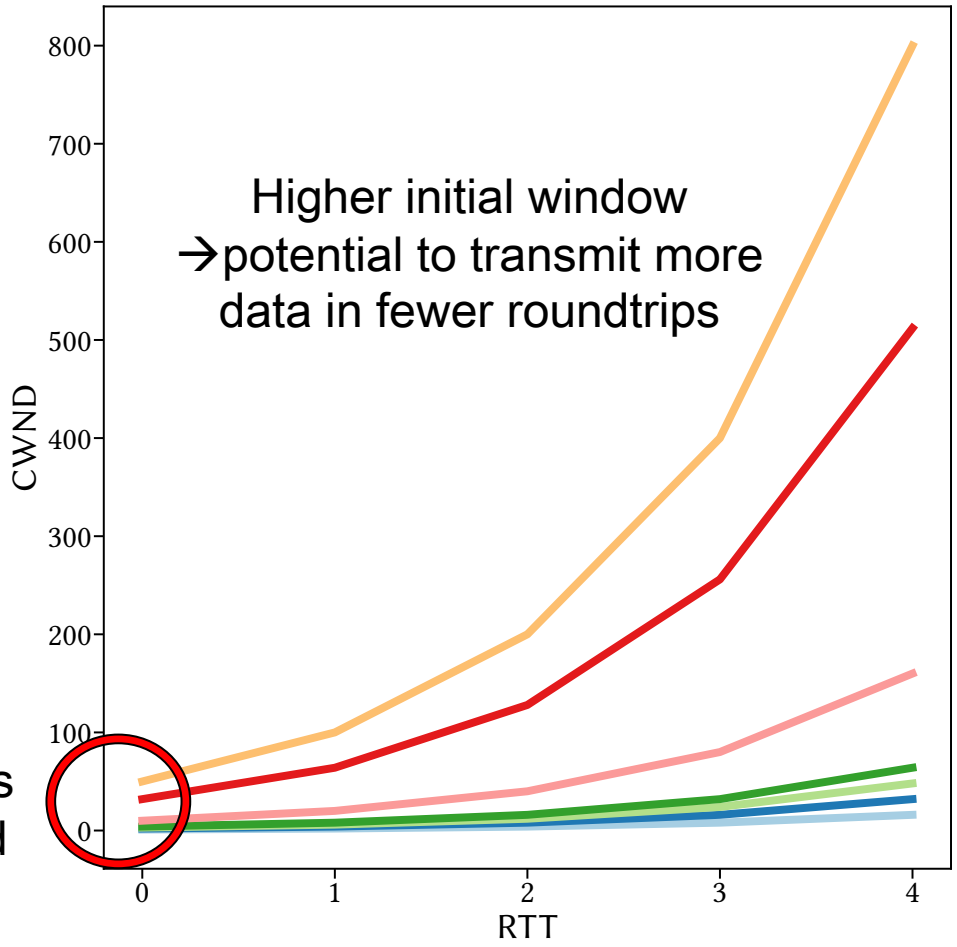

Initial Window
unacknowledged bytes
"in flight" in first round

typically as a
multiple of the MSS

Jan Rüth, Christian Bormann, Oliver Hohlfeld

Why look at Initial Windows?

Amount of data bursted in an unprobed network

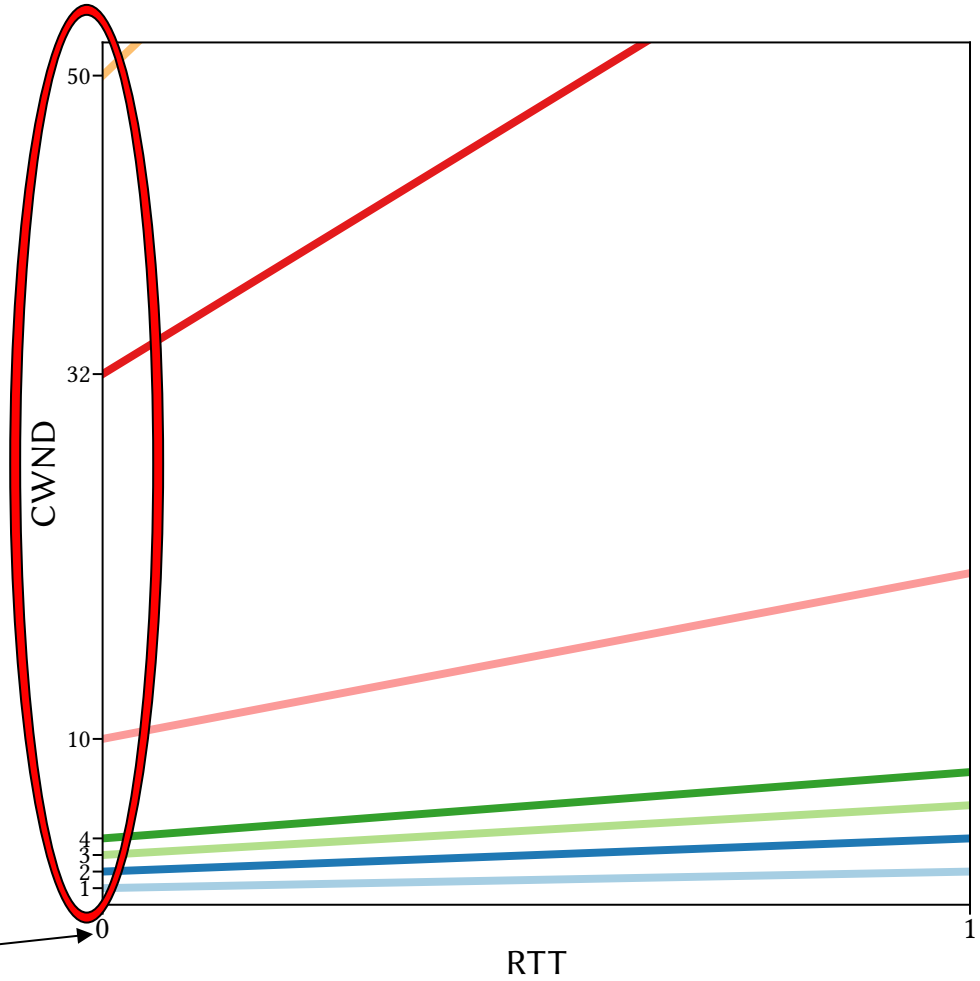

At the start, we don't know the bottleneck capacity

Why now?

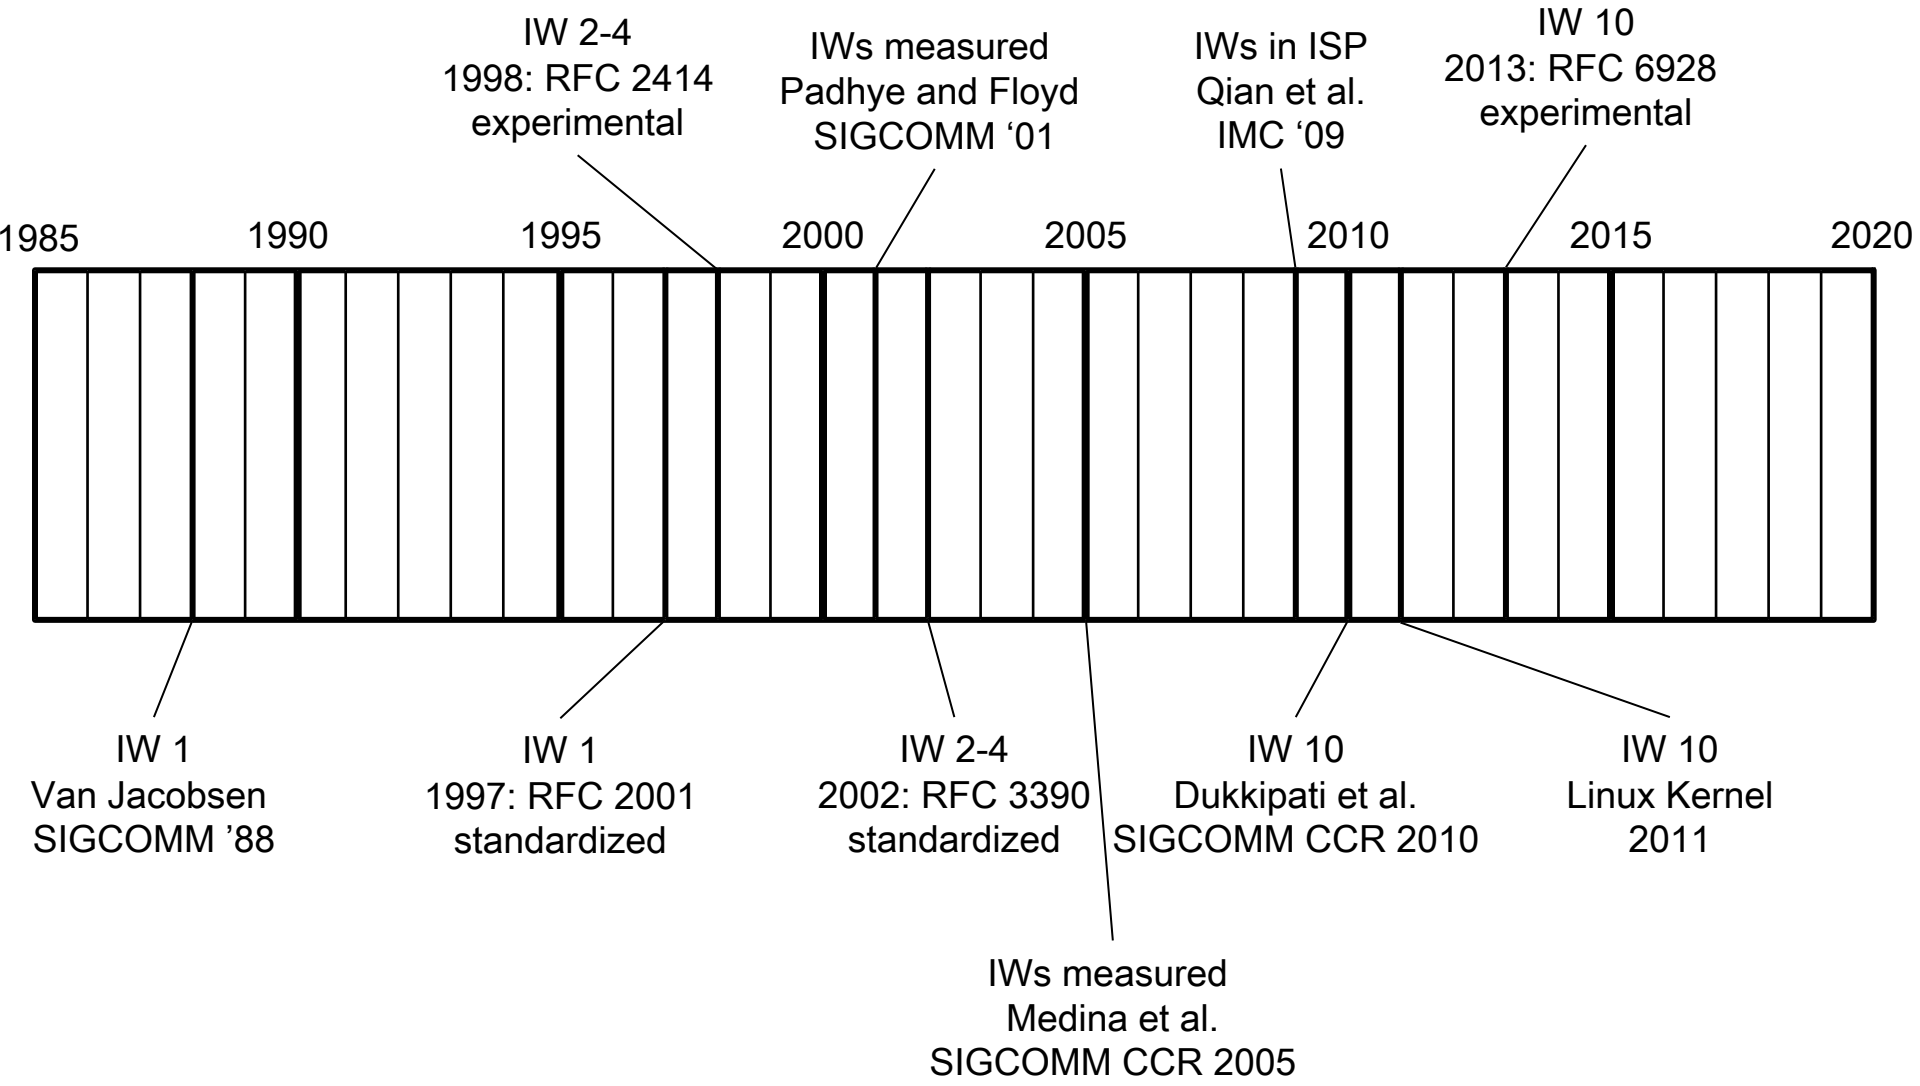
IW 2-4
1998: RFC 2414
experimental
IWs measured
Padhye and Floyd
SIGCOMM '01
IWs in ISP
Qian et al.
IMC '09
IW 10
2013: RFC 6928
experimental
1985
1990
1995
2000
2005
2010
2015
2020
IW 1
Van Jacobsen
SIGCOMM '88
IW 1
1997: RFC 2001
standardized
IW 2-4
2002: RFC 3390
standardized
IWs measured
Medina et al.
SIGCOMM CCR 2005
IW 10
Dukkipati et al.
SIGCOMM CCR 2010
IW 10
Linux Kernel
2011

Why now?

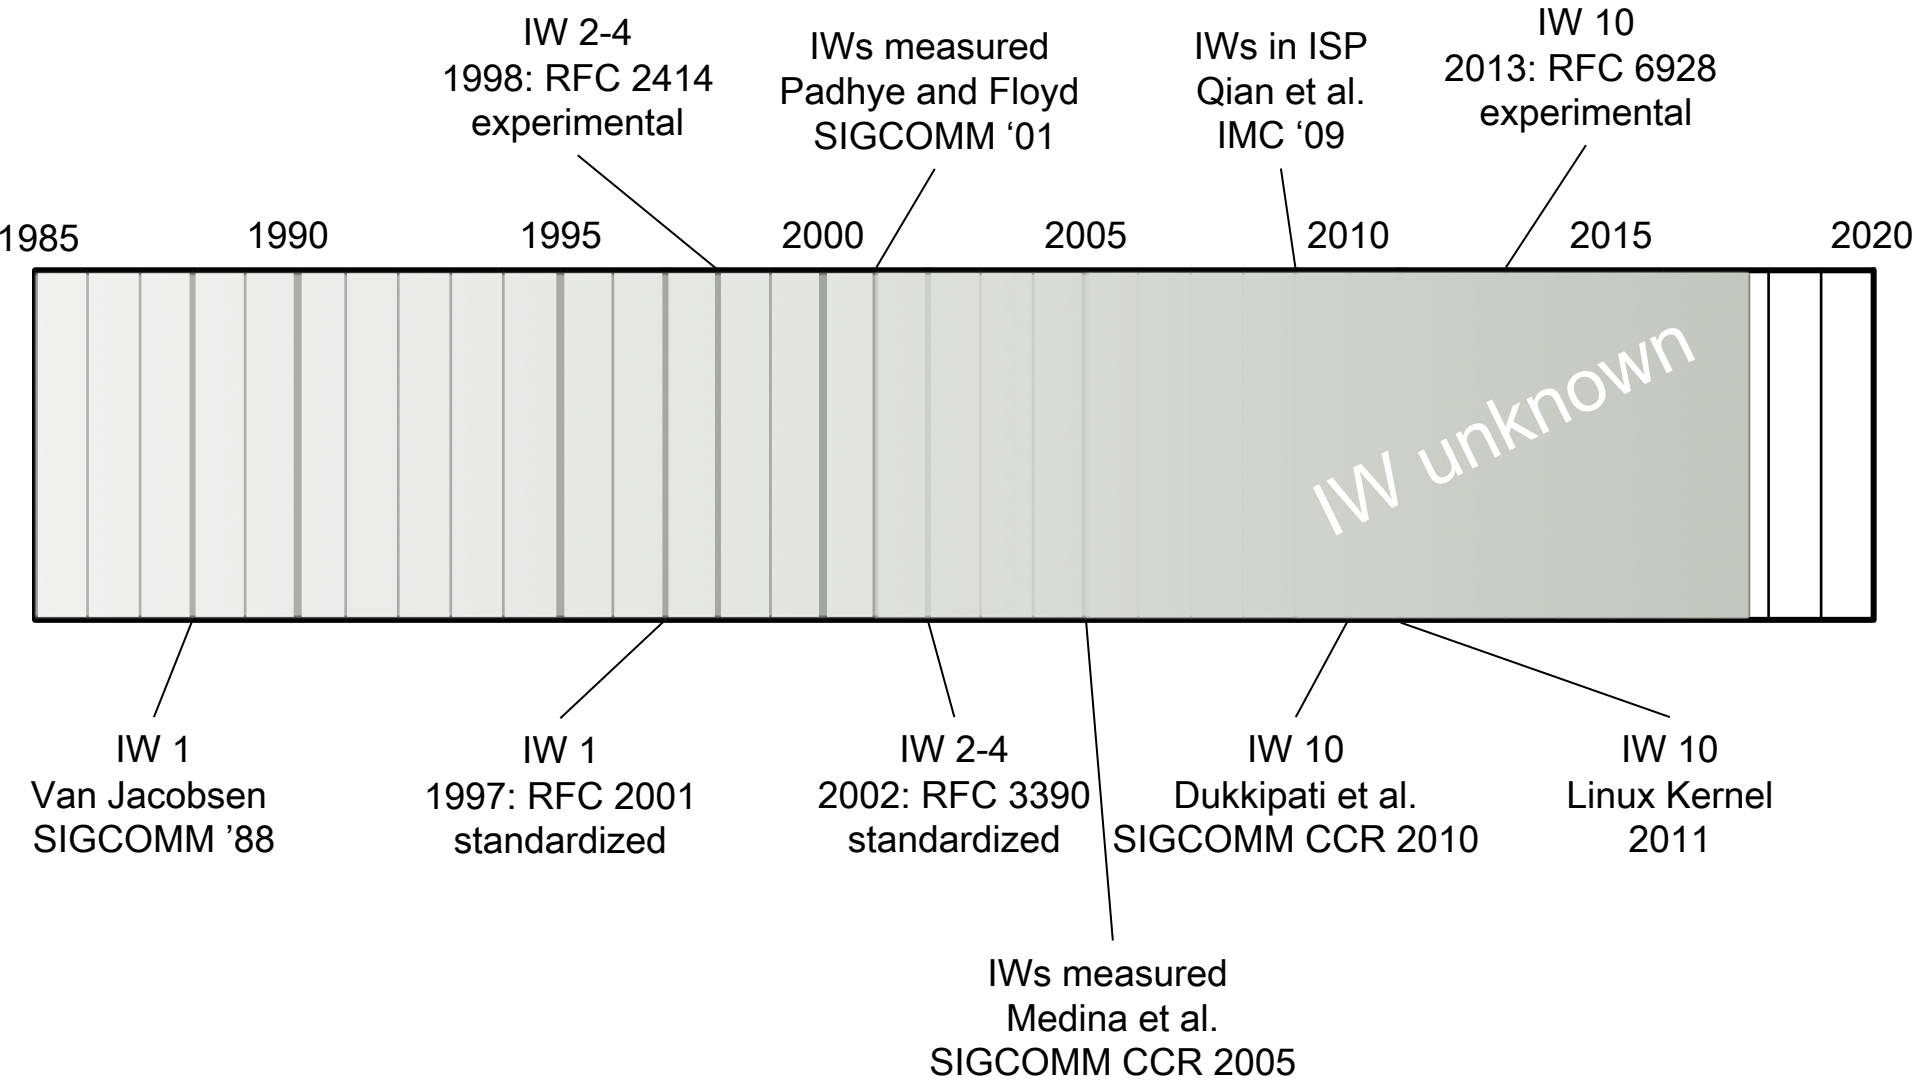
IW 2-4
1998: RFC 2414
experimental
IWs measured
Padhye and Floyd
SIGCOMM '01
IWs in ISP
Qian et al.
IMC '09
IW 10
2013: RFC 6928
experimental
1985
1990
1995
2000
2005
2010
2015
2020
IW unknown
IW 1
Van Jacobsen
SIGCOMM '88
IW 1
1997: RFC 2001
standardized
IW 2-4
2002: RFC 3390
standardized
IWs measured
Medina et al.
SIGCOMM CCR 2005
IW 10
Dukkipati et al.
SIGCOMM CCR 2010
IW 10
Linux Kernel
2011

Why now?

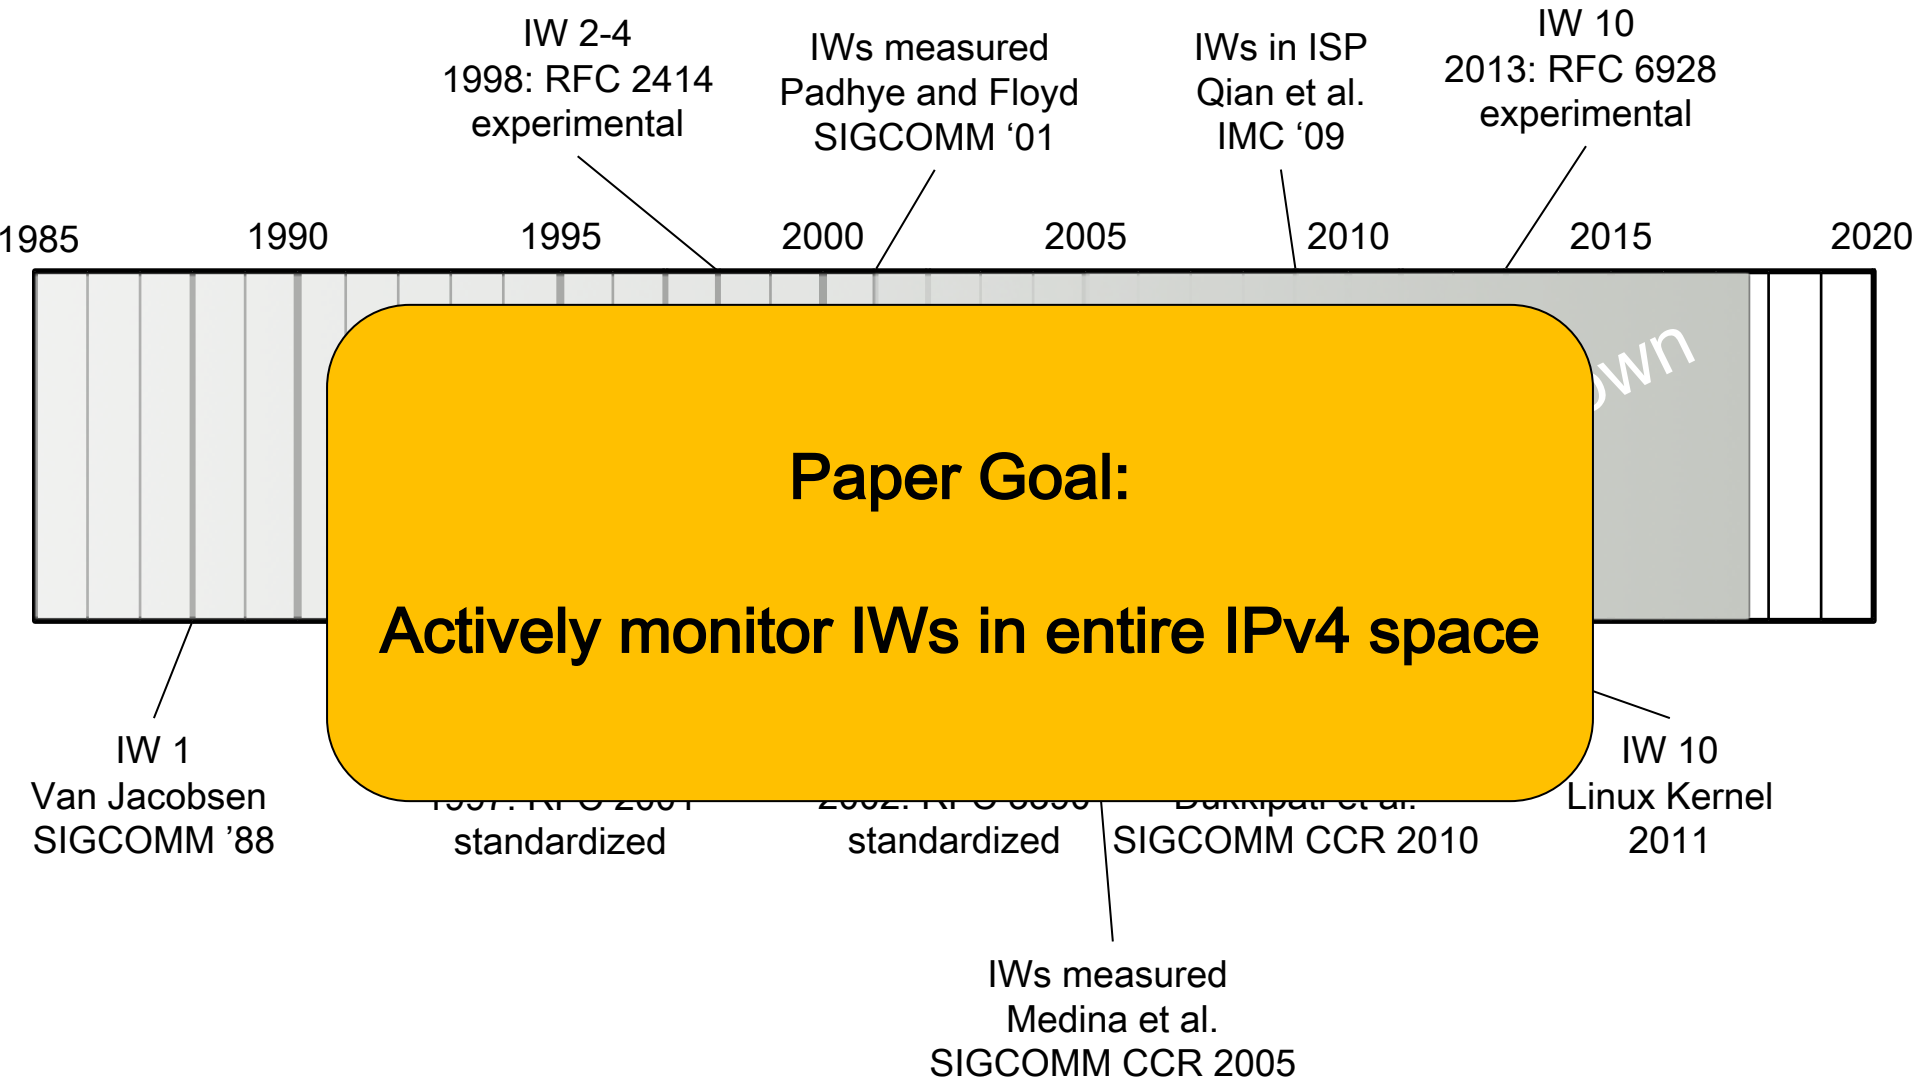
IW 2-4
1998: RFC 2414
experimental
IWs measured
Padhye and Floyd
SIGCOMM ‘01
IWs in ISP
Qian et al.
IMC ‘09
IW 10
2013: RFC 6928
experimental
1985
1990
1995
2000
2005
2010
2015
2020
Paper Goal:
Actively monitor IWs in entire IPv4 space
IW 1
Van Jacobsen
SIGCOMM ’88
standardized
standardized
SIGCOMM CCR 2010
IW 10
Linux Kernel
2011
IWs measured
Medina et al.
SIGCOMM CCR 2005

Measuring IWs

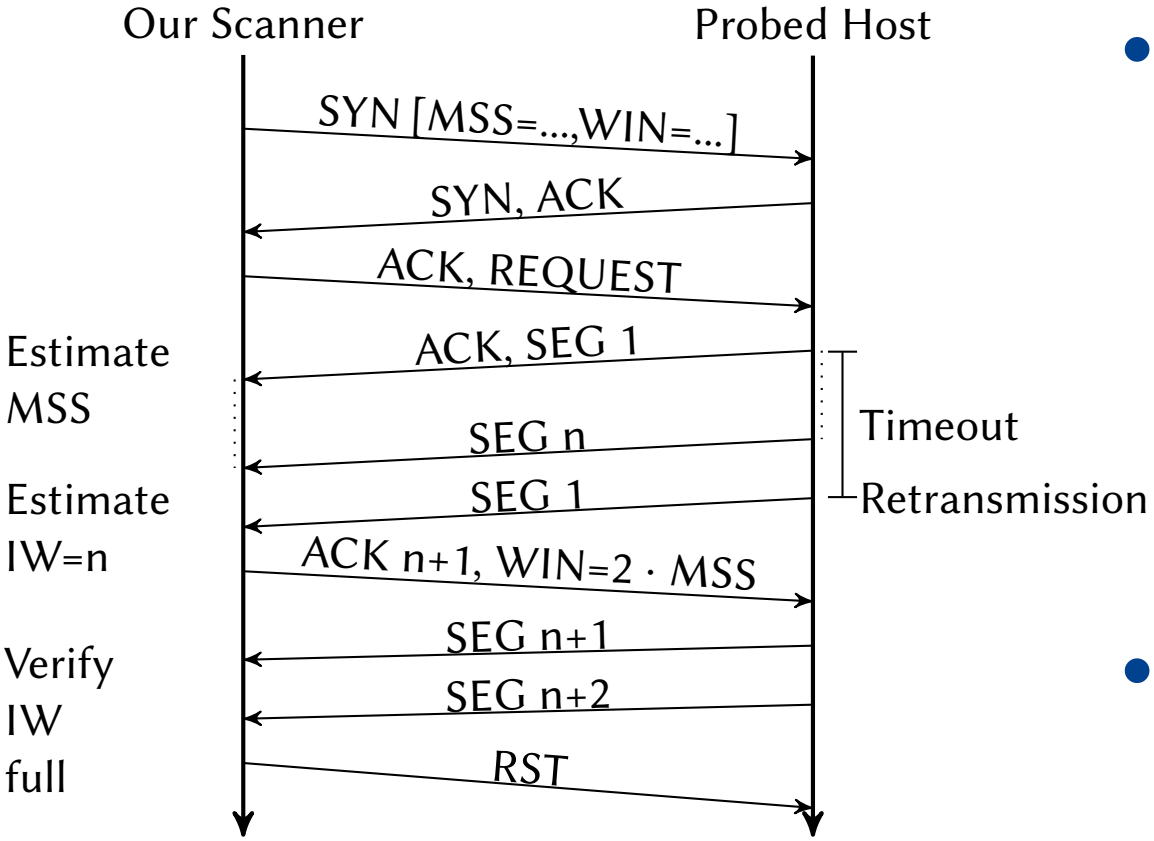

- **Loss is a problem**
  - Actually tail-loss
  - **Do multiple scans**
  - Scan early in the morning

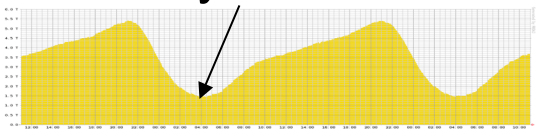

  - Disable tail-loss probes
    - Do not enable SACK
- **Trigger big response**
  - HTTP and TLS

- **Announce small MSS and large receive window**
- **Use ACK to test for more data**
  - Was the host out of data or was the IW actually full?

Jan Rüth, Christian Bormann, Oliver Hohlfeld

Measuring IWs – Probe without prior knowledge

- **Send a client hello as the request**
- **Server hello contain certificate chains**
  - We further request options enlarging the reply (e.g., cert stapling)

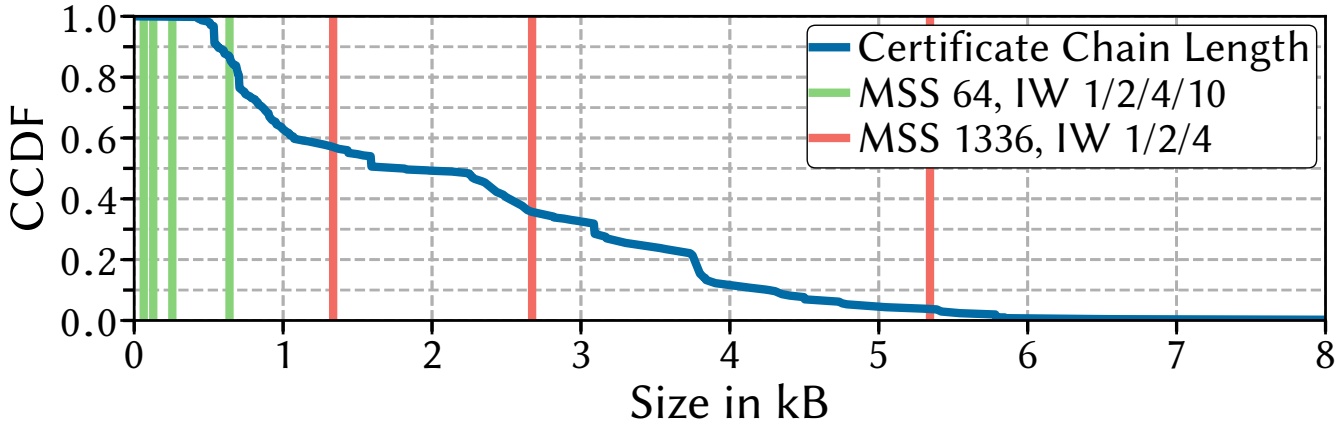

- **Fails when SNI is enforced**

Scanner implementation

- **We want to probe all reachable IPv4 HTTP/TLS hosts**

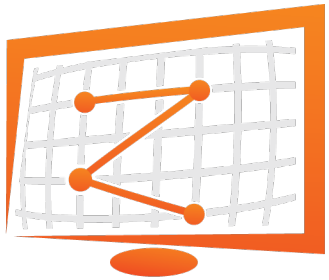

- **We implement the methodology in Zmap**
  - Bypasses the kernel stack
  - Typically only used for enumeration
  - We enable Zmap to send multiple packets
  - We can manually craft connections and manipulate them

- **Modified Zmap, HTTP/TLS scanners available on Github**
  - https://github.com/COMSYS/zmap

Results – IPv4 HTTP/TLS

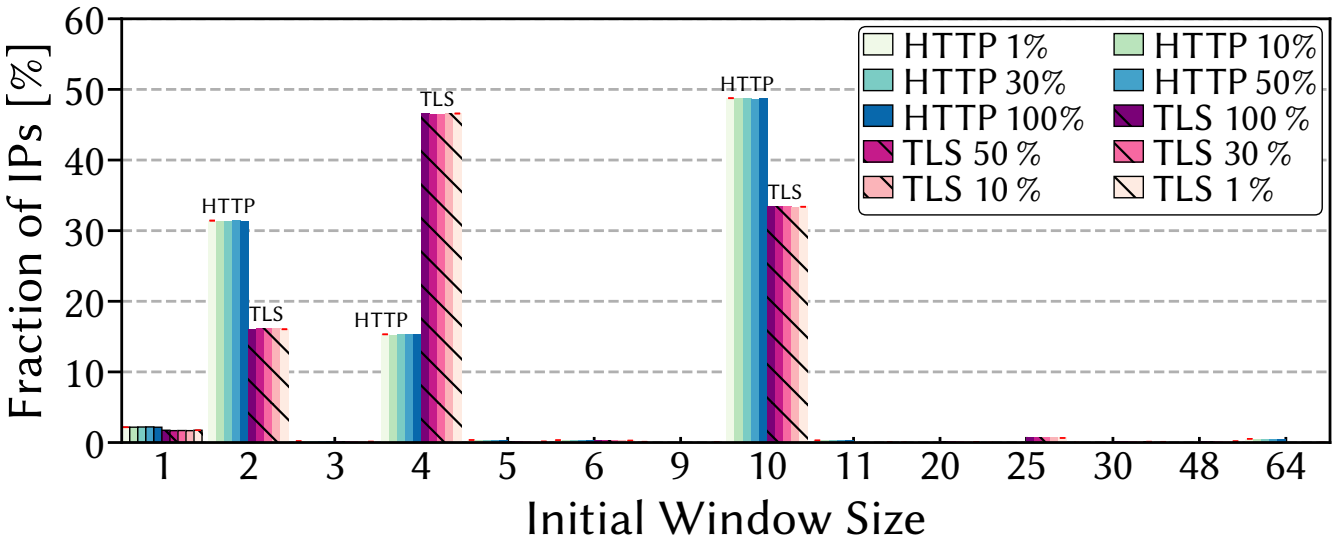

- **TLS and HTTP do not agree**
  - Many TLS hosts still use IW 4
- **HTTP scan triggers many abuse mails**
  - In contrast to TLS, this appears in access logs
- **How much scanning is enough?**

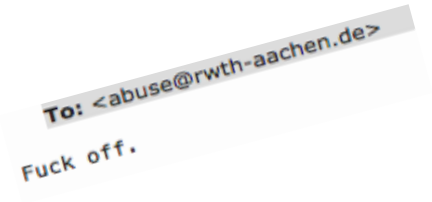

Jan Rüth, Christian Bormann, Oliver Hohlfeld

Results – IPv4 HTTP/TLS

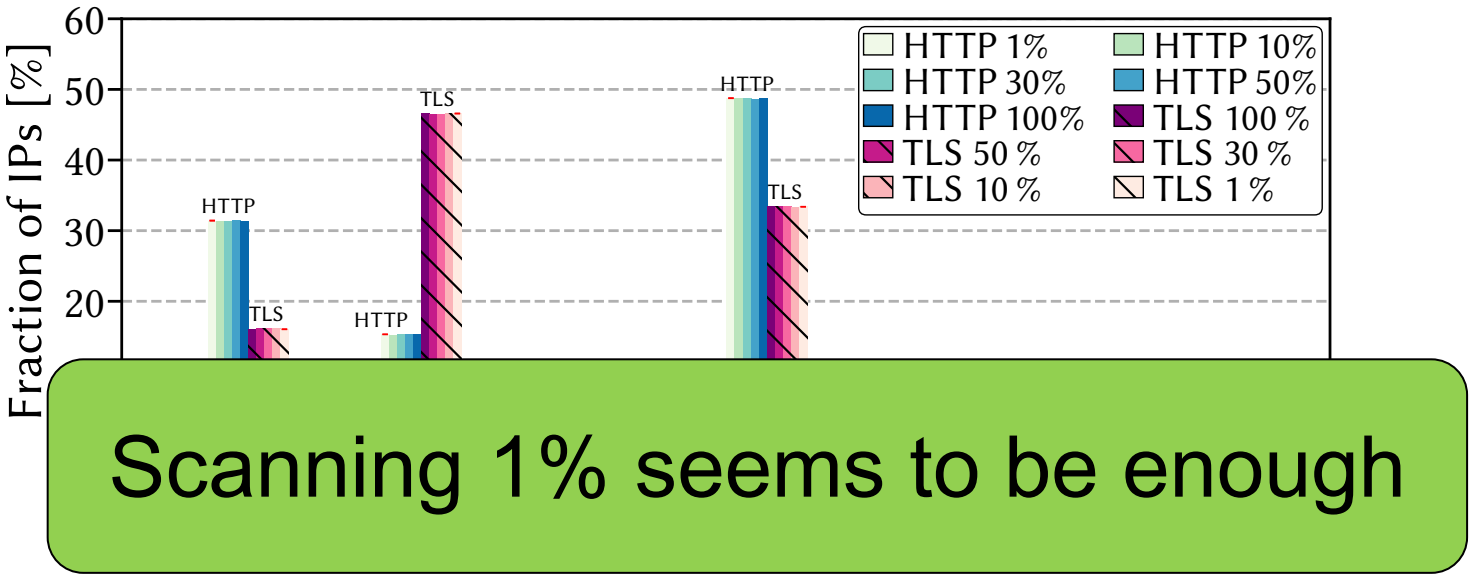

Scanning 1% seems to be enough

- **TLS and HTTP do not agree**
  - Many TLS hosts still use IW 4
- **HTTP scan triggers many abuse mails**
  - In contrast to TLS, this appears in access logs
- **How much scanning is enough?**

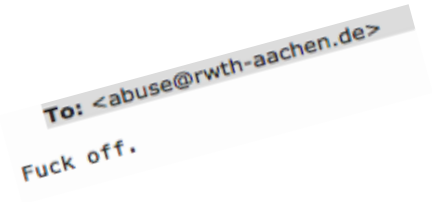

Jan Rüth, Christian Bormann, Oliver Hohlfeld

Results – Who uses which IW?

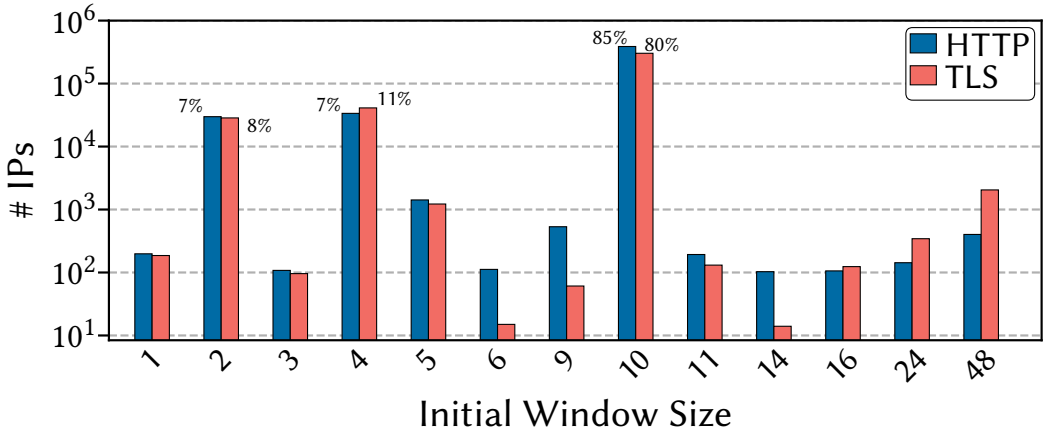

Service	HTTP IW1	HTTP IW2	HTTP IW4	HTTP IW10	TLS IW1	TLS IW2	TLS IW4	TLS IW10
Akamai	–	–	–	–	0.0	0.0	100.0	0.0
EC2	0.0	1.8	3.4	94.7	0.2	1.3	2.6	95.8
Cloudflare	0.0	0.0	0.0	100.0	0.0	0.0	0.0	100.0
Azure	0.0	7.8	54.9	37.1	0.1	4.1	73.3	21.9
Access NW	3.5	50.2	20.8	21.7	4.5	17.6	67.1	10.4

- **Most people in the Alexa list follow current RFCs**
  - Here: similar distribution for HTTP and TLS
- **Generally, we see older IWs in Access Networks**
- **CDNs and Cloud seem to be more up to date**

Jan Rüth, Christian Bormann, Oliver Hohlfeld

Conclusion

- **Distributions dominated by RFC-recommended values**
  - Still a lot of IW 2 and IW 4
  - Heavily used infrastructure and popular hosts seem to be on IW 10
- **We also find some customization**
  - Some hosts have very large IWs

- **Periodic 1% scans are available at https://iw.comsys.rwth-aachen.de**
- **Source code available at https://github.com/COMSYS/zmap**