European Commission — FP7

R/TE

REDUCING INTERNET TRANSPORT LATENCY

Project acronym: **RITE**

Project number: 317700

Work package: Dissemination & Exploitation

Deliverable number and name:

Deliverable 4.2: Dissemination and standardisation report

Title: Dissemination and standardisation report
Work Package: WP4

Version: 00B
Date: December 1, 2015
Pages: 32

Author:
Bob Briscoe

Co-Author(s):
Anna Brunström, Gorry Fairhurst, Andreas Petlund, Nicolas Kuhn, David Ros, Iffat Ahmed, Elisabeth Andersen, Koen De Schepper, Michael Welzl

To:
Jorge Carvalho
Project Officer

Status:
[ ] Draft
[ ] To be reviewed
[ ] Proposal
[X] Final / Released to CEC

Confidentiality:
[X] PU — Public
[ ] PP — Restricted to other programme participants
[ ] RE — Restricted to a group
[ ] CO — Confidential

Revision:
(Dates, Reviewers, Comments)

Contents:
This document reports dissemination activities over the 3 years of the RITE project. It includes research outputs; educational activites; coverage in the technical media; events organised by the project; standardisation activities; and software delivery.

Contents

Abbreviations

This section provides definitions of key terms and defines the abbreviations used in the remainder of the report.

3G 3rd Generation

3GPP 3rd Generation Partnership Project

ABE Alternative Backoff with ECN

AccECN Accurate ECN

ACM Association of Computing Machinery

AD Area Director (IETF)

API Application Programming Interface

AQM Active Queue Management

ARQ Automatic Repeat Request

ART Applications and and Real-Time (an IETF Area)

BCP Best Current Practice

BoF Birds-of-a-Feather

CAIA Centre for Advanced Internet Architectures (Swinburne University)

CB Circuit Breaker

CC Congestion Control

CDG CAIA Delay Gradient

CMT-SCTP Concurrent Multipath Transfer for SCTP

CoDel Controlled Delay

CSWS Capacity Sharing WorkShop

CUBIC Cubic function congestion control

CWV Congestion Window Validation

DAPS Delay-Aware Packet Scheduling

DASH Dynamic Adaptive Streaming over HTTP

DC Data Centre

DCLC Data Centre Latency Control (IRTF research group)

DCTCP Data Centre TCP

DualQ Dual Queue

E2E End-to-End

ECN Explicit Congestion Notification

eNB eNodeB, Evolved Node B (LTE Base-Station)

EuCNC European Conference on Networks and Communications

EXP Experimental (an IETF track)

FNMS Future Network & Mobile Summit

FQ Fair queuing

FQ-CoDel Flow queuing CoDel

HTTP HyperText Transfer Protocol

IAB Internet Architecture Board

ICC International Conference on Communications

ICCRG Internet Congestion Control RG (an IRTF RG)

IESG Internet Engineering Steering Group

IEEE Institution of Electrical and Electronic Engineers

IETF Internet Engineering Task Force

INF Informational (an IETF track)

IP Internet Protocol

IRTF Internet Research Task Force

ISCC IEEE Symposium on Computers and Communications

ISOC Internet SOCiety

ISP Internet Service Provider

ITU International Telecommunications Union

L4S Low Latency Low Loss Scalable throughput

LCN Local Communications Networks

LEDBAT Low Extra Delay BAckground Transport

LISA Linked Slow Start Algorithm

LTE Long Term Evolution

MAC Media Access Control

MMSYS MultiMedia SYStems

MPTCP Multipath TCP

MSS Maximum Segment Size

NEAT New, Evolutive API and Transport-layer

NOMS Network Operations and Management Symposium (IEEE)

NOSSDAV Network and Operating System Support for Digital Audio and Video

OWD One Way Delay

P2P Peer-to-Peer

PAM Passive and Active Measurement

PDCP Packet Data Convergence Protocol (3GPP)

PIE Proportional Integral controller Enhanced

QMR Quarterly Management Report

QoMEX Quality of Multimedia Experience workshop

QUIC Quick UDP Internet Connections

RED Random Early Detection

RFC Request For Comments

RG Research Group (IRTF)

RITE Reducing Internet Transport Latency End-to-End

RMCAT RTP Media Congestion Avoidance Techniques (IETF Working Group)

RT Real-Time

RTCWEB Real Time Collaboration on world wide WEB (IETF Working Group)

RDB Redundant Data Bundling

RTO Retransmission Time Out

RTOR RTO Restart

RTP Real-Time Protocol

RTT Round Trip Time

SBD Shared Bottleneck Detection

SCTP Stream Control Transmission Protocol

SDO Standards Development Organisation

SEMI Stack Evolution in a Middlebox Internet (IAB)

SNCNW Swedish National Computer Networking Workshop

SR ARQ Selective Repeat ARQ

STD Standards track (IETF)

TADA Tool for Automatic Detection of AQM

TAPS TrAnsPort Services (IETF WG)

TCP Transmission Control Protocol

TCPINC TCP INCreased security (IETF WG)

TCPM TCP Maintenance and minor modifications (IETF Working Group)

TLP Tail Loss Probe

TLPR TLP Restart

TSVAREA Transport Area (IETF)

TSVWG Transport Area WG

UDP User Datagram Protocol

UTRAN Universal Terrestrial Radio Access Network

VoLTE Voice over LTE

WiNMeE Wireless Network Measurements and Experimentations workshop

WiFi Wireless Fidelity

WoWMoM World of Wireless, Mobile and Multimedia networks symposium

WG Working Group

WP Work Package

Participant organisation name	**Participant Short Name**	**Country**
Simula Research Laboratory	SRL	Norway
BT	BT	UK
Alcatel-Lucent	ALU	Belgium
University of Oslo	UiO	Norway
Karlstad University	KaU	Sweden
Institut Mines-Télécom	IMT	France
University of Aberdeen	UoA	UK
Megapop	MEGA	Norway

1 Introduction

This report records all the different ways that the RITE partners have disseminated the outputs of their work during the three year project. On every occasion the funding of the European Community through the RITE project was acknowledged, thus reinforcing the leadership position of European research in the perception of audiences around the world.

The project Description of Work spells out the objectives of the dissemination work package as:

- *Reach system operators and software engineers that need lower latency for their applications to work or to be competitive*
- *Reach executives in the communications industry to raise low latency as an industry priority alongside bandwidth on product and marketing agendas*
- *Reach the technical media to bring low latency not just bandwidth onto the industry agenda*
- *Reach the research and scientific community through top-tier publication venues.*
- *Make the results available to the European public through operating systems implementation and standardisation*
- *Demonstrate the importance of low latency through lab events and showcases*

At the end of this report, a table is given to visualize how each dissemination activity of the project has addressed each of the above audiences.

2 Dissemination of Research

Title	Author(s)	Periodical or Journal	No./date	Pub.	Year	Del.
Enhancing TCP to Support Rate-Limited Traffic	Fairhurst, G. Secchi, R. Sathiaseelan, A.	Workshop on Capacity Sharing (CSWS'12)	2012/Dec	CSWS	2012	D1.1
Assessing LEDBAT's Delay Impact	Ros, D. Welzl, M.	IEEE Communication Letters	17(5)	IEEE	2013	D2.1
Impact of TCP Congestion Control on Bufferbloat in Cellular Networks	Alfredsson, S. Del Giudice, G. Garcia, J. Brunstrom, A. De Cicco, L. Mascolo, S.	The Fourteenth International Symposium on aWorld of Wireless, Mobile and Multimedia Networks	2013/Jun	IEEE	2013	D2.1
Observations of Bufferbloat in Swedish Cellular Networks	Alfredsson, S. Del Giudice, G. Garcia, J. Brunstrom, A. De Cicco, L. Mascolo, S.	The 9th Swedish National Computer Networking Workshop (SNCNW 2013)	2013/Jun	SNCNW	2013	Q2
Reducing Internet Transport Latency for Thin Streams and Short Flows	Brunstrom, A. Petlund, A. Rajiullah, M.	Future Network & Mobile Summit 2013	2013/Jul	FNMS	2013	Q2
A Survey of Latency Reducing Techniques and their Merits	Briscoe, B. Petlund, A. Brunstrom, A. Tsang, I-J. Ros, D. Gjessing, S. Hayes, D. Fairhurst, G.	Internet Society Workshop on Reducing Internet Latency	2013/Sep	ISOC	2013	D1.1, D2.1
On the Treatment of Application-Limited Streams	Petlund, A. Brunstrom, A. Markussen, J. Fuchs, M.	Internet Society Workshop on Reducing Internet Latency	2013/Sep	ISOC	2013	Q3
Delay-based Congestion Control for Low Latency	Ros, D. Hayes, D.	Internet Society Workshop on Reducing Internet Latency	2013/Sep	ISOC	2013	D2.1
ECN & Early ECN Marking	Welzl, M. Fairhurst, G. Khademi, N.	Internet Society Workshop on Reducing Internet Latency	2013/Sep	ISOC	2013	D2.1
Congestion-Aware Scalable Video Streaming	Fairhurst, G. Nuruddeen, I. Verdicchio, F.	Workshop on Capacity Sharing (CSWS'13)	2013/Oct	IEEE	2013	Q4

Table 2.1 — continued on next page

– continued from previous page

Title	Author(s)	Periodical or Journal	No./date	Pub.	Year	Del.
Evaluation of the Impact of Packet Drops due to AQM in Capacity Limited Paths	Fairhurst, G. Grigorescu, E. Kulatunga, C.	Workshop on Capacity Sharing (CSWS'13)	2013/Oct	IEEE	2013	D2.1
One Control to Rule Them All: Coupled Congestion Control for RTP Media	Islam, S. Welzl, M. Gjessing, S.	Packet Video Workshop 2013	2013/Dec		2013	Q4
Workshop report: Reducing internet latency, 2013	Ford, M.	Computer Communication Review	44(2)	ACM	2014	Q7
The New AQM Kids on the Block: An Experimental Evaluation of CoDel and PIE	Khademi, N. Ros, D. Welzl, M.	IEEE Global Internet Symposium	2014/Apr-May	IEEE	2014	D2.2
Performance Evaluation of Congestion Window Validation for DASH Transport	Nazir, S. Hossain, Z. Secchi, R. Broadbent, M. Petlund, A. Fairhurst, G.	NOSSDAV March 2014 Singapore	2014/Mar	NOSSDAV	2014	D1.2
A Measurement Based Study of TCP Protocol Efficiency in Cellular Networks	Garcia, J. Alfredsson, S. Brunstrom, A.	WiNMeE 2014	2014/May	IEEE	2014	Q6
DAPS: Intelligent Delay-Aware Packet Scheduling For Multipath Transport	Kuhn, N. Lochin, E. Mifdaoui, A. Sarwar, G. Mehani, O. Boreli, R.	IEEE International Conference on Communications	2014/Jun	IEEE	2014	D1.2
On the existence of optimal LEDBAT parameters	Kuhn, N. Trang, S. Lochin, E. Baudoin, C. Dubois, E. Gelard, P.	IEEE International Conference on Communications	2014/Jun	IEEE	2014	D2.3
Analysis of SR ARQ Delays Using Data-bundling over Markov Channels	Ahmed, I. Badia, L. Petlund, A. Griwodz, C. Halvorsen, P.	The Nineteenth IEEE Symposium on Computers and Communications (IEEE ISCC)	2014/Jun	IEEE	2014	Q7
Can SPDY Really Make the Web Faster?	Elkhatib, Y. Tyson, G. Welzl, M.	Networking 2014	2014/Jun	IEEE	2014	Q7
Revisiting Old Friends: Is CoDel Really Achieving What RED Cannot?	Kuhn, N. Lochin, E. Mehani, O.	Workshop on Capacity Sharing (CSWS´14)	2014/Jun	CSWS	2014	D2.3
Experimental Evaluation of Multipath TCP Schedulers	Paasch, C. Ferlin-Oliveira, S. Alay, Ö. Bonaventure, O.	Workshop on Capacity Sharing (CSWS´14)	2014/Aug	CSWS	2014	D1.2
Coupled congestion control for RTP media	Islam, S. Welzl, M. Gjessing, S. Khademi, N.	Workshop on Capacity Sharing (CSWS´14)	2014/Aug	CSWS	2014	D1.2
Practical Passive Shared Bottleneck Detection Using Shape Summary Statistics	Hayes, D. Ferlin-Oliveira, S. Welzl, M.	The 39th IEEE Conference on Local Computer Networks (LCN) 2014	2014/Sep	IEEE	2014	D1.2
Reducing Internet Latency: A Survey of Techniques and their Merits	Briscoe, B. Brunstrom, A. Petlund, A. Hayes, D. Ros, D. Tsang, I. Gjessing, S. Fairhurst, G. Griwodz, C. Welzl, M.	Communications Surveys & Tutorials, IEEE	PP(99)	IEEE	2014	D1.3*
Using Data Center TCP (DCTCP) in the Internet	Kühlewind, M. Wagner, D.P. Espinosa, J.M. Briscoe, B.	Globecom Workshop on Telecommunication Standards	2014/Dec	IEEE	2014	D2.2
An Evaluation of Tail Loss Recovery Mechanisms for TCP	Rajiullah, M. Hurtig, P. Brunstrom, A. Petlund, A. Welzl, M.	Computer Communication Review	45(1)	ACM	2015	D1.3
Ossification: a result of not even trying?	Welzl, M. Fairhurst, G. Ros, D.	Internet Architecture Board (IAB) Stack Evolution in a Middlebox Internet (SEMI) Workshop	2015/Jan	IAB	2015	Q9

Table 2.1 — continued on next page

– continued from previous page

Title	Author(s)	Periodical or Journal	No./date	Pub.	Year	Del.
Tunneling Through Inner Space	Briscoe, B.	Internet Architecture Board (IAB) Stack Evolution in a Middlebox Internet (SEMI) Workshop	2015/Jan	IAB	2015	D1.3
How much delay is there really in current games?	Raaen, K. Petlund, A.	MMSys 2015	2015/Mar	MMSYS	2015	Q10
Enabling Internet-Wide Deployment of Explicit Congestion Notification	Trammell, B. Kühlewind, M. Boppart, D. Learmonth, I. Fairhurst, G. Scheffenegger, R.	Passive and Active Measurement Conference 2015	2015/Mar	PAM	2015	D2.3
Analysing User Satisfaction in Next Generation Networks for Multimedia Multicast Transmission	Ahmed, I. Petlund, A.	QoMEX 2015	2015/May	IEEE	2015	Q10
Analyzing the impact of bufferbloat on latency-sensitive applications	Iya, N. Kuhn, N. Verdicchio, F. Fairhurst, G.	IEEE ICC 2015	2015/Jun	IEEE	2015	D1.3
Coupled Congestion Control for WebRTC		EUCNC 2015	2015/Jun	EUCNC	2015	D1.3
Internet Latency: Causes, Solutions and Trade-offs		EUCNC 2015	2015/Jun	EUCNC	2015	Q11
Evaluation of Priority Scheduling and Flow Starvation for Thin Streams with FQ-CoDel	Grigorescu, E. Kulatunga, C. Kuhn, N. Fairhurst, G.	EUCNC 2015	2015/Jun	EUCNC	2015	D2.3
Reducing Transport Latency using Multi-path Protocols	Hurtig, P. Alay, O. Ferlin, S. Yedugundla, K.	EUCNC 2015	2015/Jun	EUCNC	2015	Q11
Tackling Bufferbloat in Capacity-limited Networks	Kuhn, N. Kulatunge, C. Fairhurst, G. Ros, D.	EUCNC 2015	2015/Jun	EUCNC	2015	D2.3
Evaluating the Performance of Next Generation Web Access via Satellite	R. Secchi, A. Mohideen, and G. Fairhurst	EIA conf on Wireless and Satellite Systems (WISATS)	2015/Jul	EIA	2015	D1.3
Evaluating the Performance of Next Generation Web Access via Satellite	R. Secchi, A. Mohideen, and G. Fairhurst	Wireless and Satellite Systems in the series Lecture Notes of the Inst. for Comp Sci, Social Informatics and Telecoms Engineering	154:163–176	Springer	2015	D1.3
IETF Recommendations Regarding Active Queue Management	F. Baker, G. Fairhurst	Request for Comments	RFC 7567, BCP 197	RFC Edit'r	2015	D2.3
Problem Statement and Requirements for Increased Accuracy in Explicit Congestion Notification (ECN) Feedback	Kühlewind, M., Scheffenegger, R. & Briscoe, B.	Request for Comments	RFC 7560	RFC Edit'r	2015	D2.3
Updating TCP to support Rate-Limited Traffic	G. Fairhurst, A. Sathiaseelan, R. Secchi	Request for Comments	RFC 7661	RFC Edit'r	2015	D1.3
Latency and Fairness Trade-Off for Thin Streams using Redundant Data Bundling in TCP	Opstad B. R., Markussen J., Ahmed I., Petlund A., Griwodz C., and Halvorsen P.	40th IEEE LCN	2015/Oct	IEEE	2015	D1.3
TCP and SCTP RTO Restart	Hurtig, P. Brunstrom, Petlund, A. Welzl, M.	Request for Comments	RFC ?**	RFC Edit'r	2015	D1.2
The Benefits of using Explicit Congestion Notification (ECN)	M. Welzl and G. Fairhurst	Request for Comments	RFC ?**	RFC Edit'r	2015	D2.2
TADA: An Active Measurement Tool for Automatic Detection of AQM	Bideh, M.K. Petlund, A. Griwodz, C. Ahmed, I. Halvorsen, P. Behjati, B. Brunstrom, A. Alfredsson, S.	Int'l Conf on Performance Evaluation Methodologies and Tools (ValueTools 2015)	2015/Dec	EAI	2015	Q12

Table 2.1 — continued on next page

– continued from previous page

Title	Author(s)	Periodical or Journal	No./date	Pub.	Year	Del.
Managing Real-Time Media Flows through a Flow State Exchange	Islam, S. Welzl, M. Hayes, D. Gjessing, S.	NOMS 2016 (to appear)	2016/Apr	IEEE	2016	
Improving PIE's performance over high-delay paths	Kuhn, N., Ros, D.	IEEE ICC 2016	2016/May	IEEE	2016	D2.3
Operating ranges, tunability and performance of CoDel and PIE	Kuhn, N., Ros, D., Bagayoko, A., Culatunga, C., Fairhurst, G., Khademi, N.	Computer Communications (under submission)				D2.3
Revisiting Congestion Control for Multipath TCP with Shared Bottleneck Detection	Ferlin, S., Alay, O., Hayes, D., Dreibholz, T. and Welzl, M.	Infocom 2016 (to appear)	2016/Apr	IEEE	2016	
Is Multi-Path Transport Suitable for Latency Sensitive Traffic?	Yedugundla, K., Ferlin, S., Dreibholz, T., Alay, O., Kuhn, N., Hurtig, P. and Brunstrom, A.	Computer Networks (under submission)				D1.3

* Also in Supplementary Report "RITE - Methodology and Choices", Jan 2014

** Approved pending formalities

Table 2.1: Peer-reviewed Publications (or submitted where stated)

Title	Author	Institution or Publication	Type	Number	Date	Del.
New AQM Kids on the Block	Khademi, N. Ros, D. Welzl, M.	Uni of Oslo Dept of Informatics	Technical Report	434	23/10/13	D2.3
RTO Restart Experimental Results	M. Rajiullah, P. Hurtig, A. Brunstrom, A. Petlund, and M. Welzl	RITE Project	Internal Report		2014/Jun	D1.2
Validation of CDG implementation in ns-2	Kuhn, N.	Institut Mines Telecom / Telecom Bretagne	Technical Report		2014	D1.2
Scaling TCP's Congestion Window for Small Round Trip Times	Briscoe, B. De Schepper, K.	BT	Technical Report	TR-TUB8-2015-002	10/06/15	
Latency considerations for thin-stream Congestion Control	B.R. Opstad, J. Markussen	RITE project	Technical Report (Work in Progress)		17/07/15	D3.4
Data Center to the Home	De Schepper, K., Bondarenko, O., Tsang, I. & Briscoe, B.	RITE project	Work in Progress		28/07/15	D2.3
Insights from Curvy Random Early Detection (RED)	Briscoe, B.	BT	Technical Report	TR-TUB8-2015-003	29/07/15	D2.3
On the coexistence of AQM and LBE	N. Kuhn, X. Corbillon, and E. Lochin	RITE Project	(under submission)			D2.3
Up to Speed with Queue View	Briscoe, B., and Hurtig, P.	RITE Project	Technical Report		09/09/15	D2.3
Ultra-Low Delay for All	Briscoe, B.	IETF Journal	Invited Paper		06/11/15	Q12
What Use is Top Speed without Acceleration?	B. Briscoe, M. Rajiullah, A. Brunstrom, and A. Petlund	RITE Project	Work in Progress		2015/Nov	D2.3

Table 2.2: RITE non-peer-reviewed papers

Table 2.1 lists RITE papers accepted in peer-reviewed publication venues. Given pressure on space, some of the columns are quite curt. The "No./date" column contains either the number for a journal article (in standard "volume(issue):pages" format), or the date for a conference. "Pub." gives the publisher and the "Del." column gives the RITE deliverable where each paper was reported. In this column, "Q" refers to the Quarterly Management Report in

which a paper was reported, for those papers that were not featured in full in the main technical deliverables.

Table 2.2 lists papers that have not been peer-reviewed for various reasons. Most are technical reports, published to support other published work. There is also an invited paper. A few are reports of work in progress that were appended to a RITE deliverable.

RITE has produced papers on a wide range of topics related to latency reduction in a wide range of publication fora. Here we briefly highlight some published (or to appear) in venues generally recognised as of higher quality:

- The paper "Reducing Internet Latency: A Survey of Techniques and their Merits", which some of us have called our "magnum opus", discusses all aspects of Internet latency on a total of 53 pages, categorising and summarising 322 references. It is without a doubt the most comprehensive survey on the topic ever written. It has been published in IEEE Communications Surveys & Tutorials magazine. According to the Thomson-Reuters 2014 Journal Citation Reports (released in Jun 2015), this journal is considered:
  - first by Impact Factor of all journals in Telecommunications
  - first by Impact Factor of all journals in Computer Science—Information Systems
  - first by Article Influence score in Telecommunications.

  RITE partners also prepared a 2-page extended abstract of the survey for the Internet Society Workshop on Reducing Internet Latency in London, 2013, as well as a digest for the EU Conference on Network Communications, 2015.
- RITE partners have also published papers in a number of conferences ranked 'A' by the Computing Research and Education Association of Australasia (CORE)[1], which is widely respected in computing circles.[2] These include:
  - The observations on bufferbloat in Swedish cellular networks in IEEE WoWMoM (2013)
  - The performance evaluation of Congestion Window Validation for Dynamic Adaptive Streaming over HTTP (DASH) in NOSSDAV (2014)
  - The performance evaluation of Google's SPDY protocol in IEEE Networking (2014)
  - The paper on shared bottleneck detection in IEEE LCN (2014)
  - The paper on Redundant Data Bundling, also in IEEE LCN (2015)
- RITE researchers have been well-represented in a number of high-quality workshops. Of particular note was the paper on coupled congestion control for real-time media that won best paper award at the Capacity Sharing Workshop co-located with ACM SIGCOMM (2014).
- It will be noticed that five IETF RFCs (requests for comments) have been included in our list of publications. Although these are featured in §6 on Standards, they are also included as peer-reviewed publications because the process of authoring a sufficiently robust document to get it approved as an RFC is an order of magnitude more rigorous than the typical peer-review that a top-flight journal or conference paper is subjected to.[3] In §6, we give three possible reasons for this extra rigour: the materially significant consequences of errors; the depth and breadth of expertise of many IETF practitioners; and the fact that RFCs can only be replaced, not amended, once published.
- Fifteen further Internet Drafts are not included, even though five have received sufficient peer-review to have been adopted onto the IETF's standardisation agenda. They are not included in the non-peer-reviewed publications either, purely on the grounds that §6 is the more appropriate section.
- It may also be noted that the list includes a few papers associated with industry workshops, particularly those held by the Internet Architecture Board (IAB). These are different from academic workshops in that attendance is often gated by acceptance of a position paper, and acceptance is determined partly by the quality of the paper but also partly by the reputation of the author(s).

[1] http://www.core.edu.au/
[2] CORE ranks conferences A*, A, B, C, D etc. where 'C' represents a conference of the minimum standard.
[3] The RFC Series also qualifies on the grounds that it is a stable archival series with an ISSN and a DOI.

3 Education

3.1 Curriculum Development

Integration of core concepts, research results and knowledge gained through RITE into the undergraduate and graduate curricula is an important method of dissemination. This helps prepare the next generation of Internet engineers with the appropriate skills on latency-related problems and possible countermeasures and also transfers the knowledge into the industrial environments in which they start their careers. All academic partners have capitalized on this opportunity.

- KaU has integrated understanding of latency issues and solutions gained from the RITE project into the "Topics in Computer Networking" masters course starting from Spring 2013 and the course now has a focus on latency-related issues. Several student projects have also been run in relation to the topics studied in RITE.
- UiO now discuss the "bufferbloat" problem as a part of its basic computer networks education. The course material hints at potential solutions, pointing to some of the RITE mechanisms. UiO/SRL has also had a number RITE-related master theses.
- SRL and UiO is planning a new master/PhD level course on Computer Systems Analysis where part of the curriculum will focus on network latency measurement and analysis with use-cases from the core topics of the RITE project.
- UoA has added topics on bufferbloat and latency to its undergraduate network course.
- IMT has integrated RITE-related topics in its Master's level courses starting from the 2013 Fall semester. Specific courses on TCP now feature topics on bufferbloat, AQM, and latency. IMT has also run a number of student projects on RITE topics such as MPTCP.

3.2 Doctoral Research

No less than nine PhD students (2 KAU, 2 UiO, 2 SRL, 3 UoA) have also been working actively in RITE, with three having completed their doctorates:

- Z. Hossain, "A Study of Mechanisms to Support Variable-rate Internet Applications over a Multi-service Satellite Platform," Ph.D. dissertation, School of Engineering, University of Aberdeen, Aberdeen, 2015.
- N. Khademi, "Reducing Latency in Internet Access Links with Mechanisms in Endpoints and within the Network," Ph.D. dissertation, University of Oslo, Oslo, May 2015.
- M. Rajiullah, "Towards a Low Latency Internet: Understanding and Solutions," Ph.D. dissertation, Karlstad University, Karlstad, Nov. 2015

Educating PhD students with expert knowledge on latency is very important as they will bring their knowledge and experiences from RITE along to their new positions within academia and industry, being part of the new generation of experts of Internet transport.

3.3 Educational Material

While dissemination by the RITE partners within their curricula is important, an even more efficient approach is to produce materials designed to be convenient for educators to use, thus passing on our message virally. A carefully produced animation or presentation that saves a lecturer time explaining difficult concepts (e.g. the cause of bufferbloat) can influence many hundreds of thousands of students (academic learning or certification training).

In October 2014 we released a highly acclaimed, 5-minute animated video[4] explaining why more bandwidth is often not a solution to slow Internet performance. It is targeted at the general public aiming to stimulate understanding and

[4] http://youtu.be/F1a-eMF9xdY

therefore demand for latency improvements from the industry. The video have already received thousands of views on YouTube.

We have also provided supporting materials to complement the video:

- A quiz made in Kahoot!;
- An activity sheet with experiments to perform using any computer;
- A fact sheet

The latency video and associated material is being used in the education by the academic partners (e.g. UiO introductory network courses INF1060 & INF3190) to help explain the latency problem and the sources that contribute to the problem. We have also been notified that the National Technical University of Norway is using the RITE Latency video and complementary material in their introductory ICT courses "TDT4105" & "TDT4110" for over 2000 students a year.

To help spread the video, it was also promoted in two separate blogposts on the Internet Society's influential "Tech Matters" blog: A blog-post by RITE SAB member Mat Ford[5], and in a guest blog-post on the long-term challenge of Internet latency by Andreas Petlund[6] (see also § 4.1).

Based on the content in the RITE latency video, KAU also released a video on latency in the Swedish language[7]. The video and associated press material received coverage on national radio[8], as well as on local radio and newspapers.

A second RITE animated video has also been produced. The audience for this video is industry decision-makers. The video illustrates an ISP that uses a consultant to discover the reason for a competing ISP's success—a shift from focus on bandwidth to latency, in particular use of AQM. Then to continue the same theme into a third competing ISP that uses ECN as well. In order to get the most publicity around the finalization of the project, the launch of this video has been re-scheduled to coincide with the project-end press release.

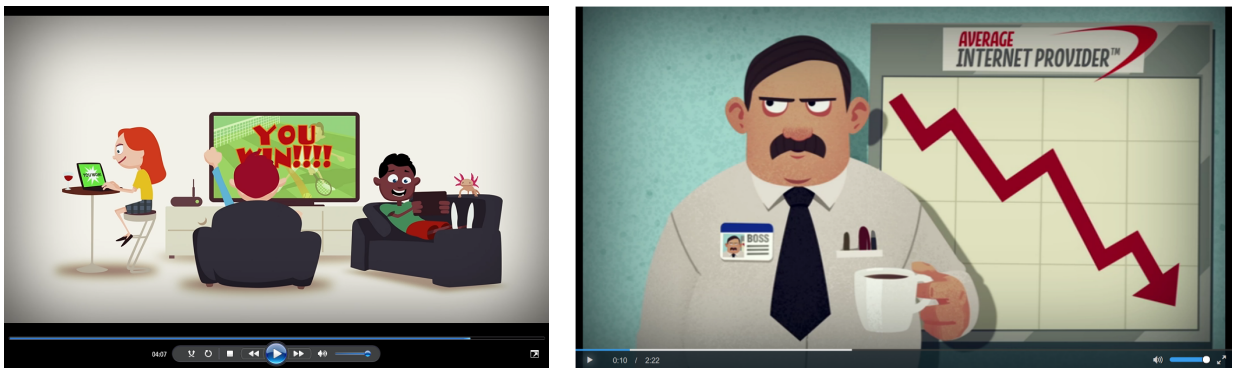

Figure 3.1: RITE animated videos aimed at young people in education (left) and industry decision-makers (right)

4 Technical Media & RITE Web SIte

4.1 Technical Media and Blogs

Table 4.1 lists articles (including blog posting) that we have found about the RITE project in the Technical Media, both offline and online.

Some articles appeared spontaneously, but most were triggered by our press releases:

[5] http://www.internetsociety.org/blog/tech-matters/2014/10/why-%E2%80%98megafast%E2%80%99-internet-often-isn%E2%80%99t-video

[6] http://www.internetsociety.org/blog/tech-matters/2014/11/reducing-internet-latency-long-term-challenge-making-internet-faster

[7] https://youtu.be/6mYHom7ijiE

[8] http://sverigesradio.se/sida/artikel.aspx?programid=1646&artikel=6278594

Publication	**Title**	Date
ALU News Item	RITE: Working Smarter not Harder will make the Internet Faster	12-Dec-12
UiO News Item	RITE: Working Smarter not Harder will make the Internet Faster	07-Jan-13
SRL News Item	RITE: Working Smarter not Harder will make the Internet Faster	10-Jan-13
KAU News Item	RITE: Snabbare Internet med smartare mjukvara	10-Jan-13
Värmlands Folkblad	Karlstadsforskare snabbar upp nätet	19-Jan-13
RITE project-summary flyer	RITE: Working Smarter not Harder will make the Internet Faster	Feb-13
UoA News Item	An end to suffering buffering? Researchers to speed up Internet	13-Feb-13
Computer Weekly	European Commission funds research to eliminate latency	13-Feb-13
Scottish regional radio	Interview	13-Feb-13
Aberdeen radio (Northsound)	Interview	13-Feb-13
v3.co.uk	Aberdeen University aims to boost internet performance by cutting network latency	13-Feb-13
USA Today	"Aberdeen University aims to boost internet performance by cutting network latency"	13-Feb-13
Evening Express (Aberdeen)	"Researchers in Bid to Increase Internet Speed"	13-Feb-13
Giacom Research	Latency is a pain, and these researchers say they have the cure	14-Feb-13
Advanced Television	Researchers seek end to Internet video buffering	14-Feb-13
Techtings	Latency is a pain, and these researchers say they have the cure	14-Feb-13
The Register	"British and Belgian boffins battle buffering bandwidth bogeyman"	14-Feb-13
ISP Review	UK Researchers Working to End Internet Video Buffering on ISP Connections	14-Feb-13
Computerworld, New Zealand	Researchers get EU funding to solve jittery internet	14-Feb-13
Computerworld, UK	Researchers get EU funding to solve jittery internet	14-Feb-13
Computerworld, Norge	"Researchers get EU funding to solve jittery internet"	14-Feb-13
CIO Asia	Researchers get EU funding to solve jittery internet	14-Feb-13
Telecompaper	"EU research project looks to speed up the internet"	14-Feb-13
Clickonline.com	Project team tasked with eliminating lag	14-Feb-13
Star Plus Ireland	"Project team tasked with eliminating lag"	14-Feb-13
Digital Gadget Magazine (Blog)	RITE Project aims to conquer internet lag, eliminate excuses for game noobs	14-Feb-13
Original 106fm (Web)	"Aberdeen Experts Bid To Give Internet A Boost"	14-Feb-13
MIS Asia (Web)	Researchers get EU funding to solve jittery internet	14-Feb-13
Individual.com (Web)	"Researchers get EU funding to solve jittery internet"	14-Feb-13
Zecco (Web)	"EU research project looks to speed up the internet"	14-Feb-13
Easier (Web)	"An end to suffering buffering? Researchers to speed up Internet"	14-Feb-13
Technet	European Commission funds RITE Project to defeat Internet lag	14-Feb-13
Parity News	EU Funded RITE Aims to Speed Up Internet by Reducing Latency	14-Feb-13
VPWsystemsUK	European Commission Backs Latency Elimination Research	14-Feb-13
übergizmo	Europe Spends 3.5M Euros To Fight Network Latency	14-Feb-13
elmundo (Web)	Proponen eliminar retrasos en la velocidad de Internet	19-Feb-13
TU.no	Raskere nett uten bredere bånd	20-Feb-13
The European Sting	The European Commission to stop Buffering	26-Feb-13
ISOC Blog	Get Happy! – Reducing Internet Latency	20-Jun-13
Caliphjingo (Blog, Uganda)	Get Happy! – Reducing Internet Latency	15-Jul-13
ISOC Blog	Speeding Up the Internet - Reducing Latency	05-Dec-13
NorNet	NorNet discusses multipath transport with RITE	26-Feb-14
ISOC Blog (incl. video)	Why 'Megafast' Internet Often Isn't	23-Oct-14
P5 (Oslo city radio station)	Radio interview (in Norwegian)	Oct-14
ACM SIGMM Records 6(3):13	Slow internet?: more bandwidth is not the answer	Oct-14
Megapop News Item	Megapop supports RITE, the EU-project for Reducing Internet Transport Latency	23-Oct-14
Lexians, Télécom Bretagne	A simple explanation of the origin of latency in the internet	24-Oct-14
ISOC Blog	Reducing Internet Latency: The Long-term Challenge of Making the Internet Faster	06-Nov-14
lwn.net	Delay-gradient congestion control	20-May-15

Table 4.1: A Selection of the Coverage of RITE Activities in the Technical Media

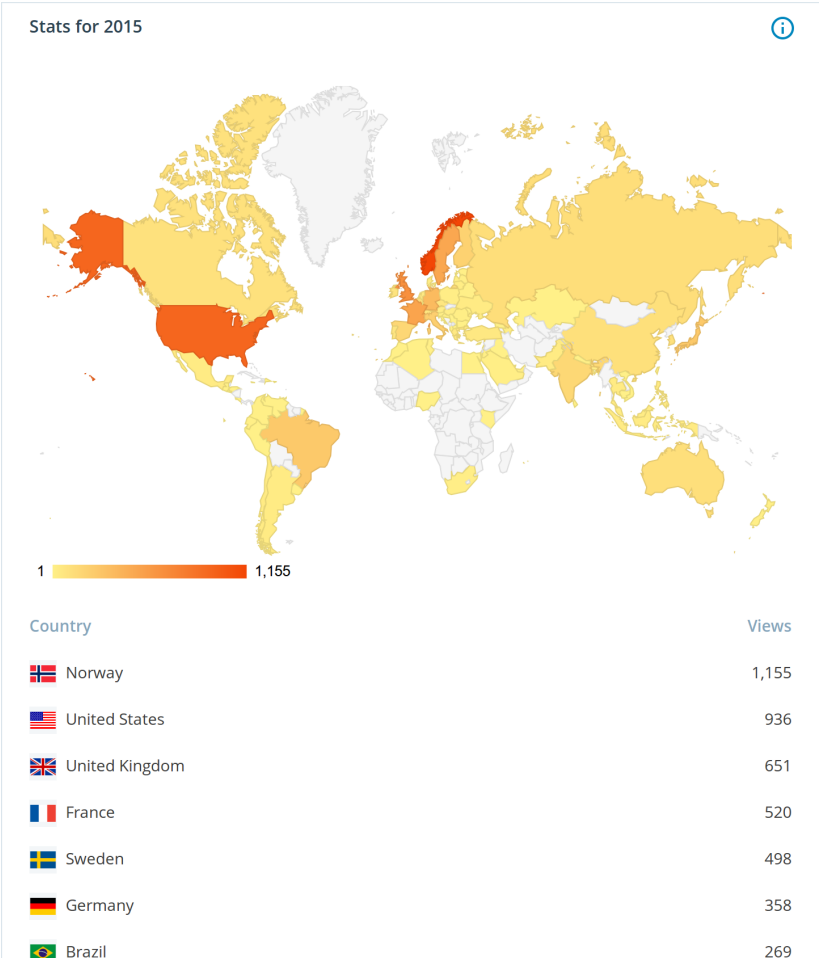

Figure 4.1: Distribution of Hits on RITE Web Site in 2015, by country (stats collected prior to press release at close of project)

- 10-Jan-13: RITE: Working Smarter not Harder will make the Internet Faster notifying the launch of the RITE project;
- 23-Oct-2014: Slow Internet? -more bandwidth is not the answer notifying the launch of the first RITE video.

One further press release is being prepared to signal the end of the project pointing out the project's main achievements. This and the response it receives will be reported at the final review.

We also prepared an article for the IETF Journal about the RITE Bits-N-Bytes exhibition stand at the Prague IETF in Jul-2015 (see Table 2.2 in § 2).

Figure 4.1 shows the global distribution of hits on the RITE Web site by country in 2015.

4.2 RITE Web Site

The RITE Web site <www.riteproject.eu> was launched in the first month of the project and has carried regular news items throughout the project (about once a month on average); and material about the project and the partners; as well as acting as a repository for all the project's outputs as soon as they have been released, as would be expected:

- Publications;
- Standards outputs;

- Downloadable resources, source code, etc.;

The Web site also carries a number of pages focused on specific activities within the project, such as:

- general material on Internet latency;
- the TCP evaluation suite;
- the New-CWV (Congestion Window Validation) activity;
- the RTO Restart (RTOR) activity;
- the Coupled Congestion Control activity;
- the Redundant Data Bundling (RDB) activity;
- the "Ultra-Low Q Delay for All" activity;
- and the Rite Integrated Linux Kernel.

5 Events

Promoting the importance of latency has been a major target for RITE and a number of events have been organized in order to draw attention to the problem and its possible solutions. Indeed we made specific commitments to organise such events in the project Description of Work:

> *Task 4.3: Arrange Dissemination Events*
>
> *RITE will arrange at least two specific dissemination events or activities. We would like to hold more than two events, depending on the opportunities that arise. As well as raising awareness of RITE outputs, the objective will also be to gather valuable feedback for the project's research process.*
>
> *Our choice of events will target audiences with the greatest potential impact on future delivered systems. The plans for these meetings will be reviewed at the RITE Project meetings, because our choice has to be tempered by practical issues such as availability of venues, opportunities for co-location and our degree of influence over the organisers.*

Three main events were organized by the project, each engaging a majority of all partners: A joint ISOC/RITE workshop on reducing Internet latency, a special session on Reducing Internet Latency at EuCNC 2015 and a Bits-n-Bites stand and demo at the 93rd IETF in Prague. These events are further detailed below along with a summary of additional events organized by or participated in by a smaller number of RITE partners.

5.1 Workshop on Reducing Internet Latency

The Workshop on Reducing Internet Latency took place on 25-26 September 2013 in London[9]. The RITE-project jointly organised the workshop with the Internet Society (ISOC). This workshop gathered many of the world's foremost experts on Internet architecture and latency. Godred Fairhurst (UoA), Anna Brunstrom (KaU), David Ros (IMT), Bob Briscoe (BT) and David Hayes (UiO) represented RITE in the workshop. The workshop was organised and co-chaired by Bob Briscoe and Mat Ford of ISOC, and Godred Fairhurst also served on the programme committee. The workshop brought together different communities working on latency, particularly researchers, developers, operational practitioners and those working on standardisation. Delegates from industry included Google (x2), Cisco, CableLabs, Comcast, Orange, BT, NSN, Al-Lu etc.

A report on the workshop on latency has been agreed by all the workshop participants and placed on-line at the workshop site. It included contributions from BT, UoA, KAU, IMT, UiO. BT also produced a variant of the report for the ACM Computer Communications Review Journal, which appeared in 2014 (see § 2).

[9] http://www.internetsociety.org/latency2013

5.2 EuCNC Special Session

A special session on Reducing Internet Latency was arranged by RITE at the European Conference on Networks and Communications (EuCNC 2015) in Paris. EuCNC is a major European event in the strategic domain of communication networks and systems, bringing together the worldwide research and innovation community and attracting over 500 key actors from all over Europe and beyond. The special session was held on Tuesday June 30 and featured four paper presentations:

- Internet Latency: Causes, Solutions and Trade-offs
- Reducing Transport Latency using Multi-path Protocols
- Evaluation of Priority Scheduling and Flow Starvation for Thin Streams with FQ-CoDel
- Coupled Congestion Control for WebRTC

All papers were authored by RITE researchers, with contributions from in total seven RITE partners. The session was well received and spurred a number of interesting discussions.

5.3 IETF Bits-n-Bites Exhibition Stand and Demo

Figure 5.1: RITE exhibition stand and demonstration at IETF 93 Bits-N-Bites

The RITE project was present with a stand at the Bits-n-Bites session at IETF-93 in Prague. The Bits-n-Bites session features a number of exhibition stands and demos along with plenty of food and drinks. It attracts a large crowd with around a thousand people registered for the Prague Bits-n-Bites.

The technical focus of the RITE stand was to show the latency benefit of Explicit Congestion Notification (ECN) and to demonstrate the Dual Queue AQM solution developed in the project. This solution enables coexistence between Data Centre TCP-like congestion control and traditional TCP congestion control over a bottleneck. The demonstration was a great success and generated a queue of interested people throughout the session—significantly surpassing any

other exhibition stand. The technology gathered interest from several major equipment vendors, ISPs and operating system developers.

ALU's presentation of the same technology to the IETF AQM WG (about 400 people) the previous day was part of the reason for the interest. A video of this is available in the IETF's regular proceedings. It shows the sub-millisecond broadband queuing latency being demonstrated live on ALU's video testbed by remote access. We also advertised the Bit-n-Bites exhibition in a number of working groups relevant to real-time applications.

The RITE animated video was also on display during Bits-n-Bites. As discussed in Sec. 3 the video explains the fundamental questions that the RITE project tackles, providing people with a brief introduction to the challenges of Internet latency. Quizzes allowed the crowd to compete for RITE give-aways.

A professional video of the event has been produced by SRL and ALU, using a professional producer who shot the footage. It combines introductory head-and shoulders material with material taken during the Bits-N-Bites session, which shows the crowds around the RITE stand as well as a full demonstration of ALU's Panoramic Interactive Video application and an explanation of the technology.

5.4 Executive Briefings

RITE partners have all given frequent presentations about the project and its outputs to management and executives within their own organisations. The list below gives a flavour of the types of audience who have been influenced by RITE activities, with a particular emphasis on those likely to make product decisions:

- Presentations to various relevant ALU business-line managers, followed-up with ongoing technical R&D meetings to investigate feasibility for development;
- Presentation to the president of Bell Labs and CTO of ALU.
  As documented in the recent book from Bell Labs "The Future X Network: A Bell Labs Perspective", which outlines how Bell Labs sees the future unfolding and the key technological breakthroughs needed at both the architectural and systems levels, the future cloud-integrated network will be driven by 2 key performance drivers: ultra-high access capacity, and ultra-low cloud-connectivity latency (Ref: The Future X Network: A Bell Labs Perspective, Marcus K. Weldon, October 22, 2015);
- Presentations to customers of ALU; three Telecom Operators to show future network capabilities with regard to interactive services requiring high throughput and low latency. For one of these, a European Operator, we demonstrated the importance of high throughput and low latency to its CEO and its the board of directors;
- Presentation in BT "Lunchtime Byte" series, to about 80 BT product line executives and managers: "Quick as a Flash" – using RITE's first animated video and other demonstrations to illustrate the importance of latency not just bandwidth;
- Upcoming follow-up presentation in the BT "Lunchtime Byte" series on "Quality of Service";
- Knowledge Share on "Cutting End-to-End Latency" for the CTO community and Technical Sales community of BT's Global Banking and Finance Markets division;
- Presentations to various of BT's enterprise customers in the banking sector, on reducing end-to-end latency, e.g. Thomson-Reuters, Morgan Stanley, Credit Suisse, etc.
- SRL presentation to the Norwegian Defence Research Establishment ICT seminar on Reducing Internet Transport Latency (scheduled for December 16th 2015).
- Presentation by IMT on RITE to ISAE (French Institut Supérieur d'Aéronautique et de l'Espace)
- Scheduled IMT presentations to CNES (Centre National d'Etudes Spatiales), TAS (Thales Alenia Space), IRIT (Institut de Recherche en Informatique de Toulouse), etc.

5.5 Other Events

RITE partners attended five events to facilitate coordination between EU projects ('concertation'). Given that these were more coordination than dissemination, they are reported in the final management report.

SRL and KAU presented[10,11] some of the RITE prototypes that have been implemented in the Linux kernel at Linux

[10] https://events.linuxfoundation.org/sites/events/files/slides/RITE-LinuxCon2015-v3.pdf
[11] https://lceeu2015.sched.org/event/1d959dedb3951725907e8af78262acae

Conference Europe in October 2015, to promote uptake of the solutions. LinuxCon Europe 2015 had 1484 registered attendees and gathers core competence from the open source community, both developers, businesses and academics (see the photo in Figure 5.2). The RITE talk was one of the most popular in the regular technical (non-keynote) track. In RITE, one of the goals is to develop networking technology that enables lower latency transport available to the public and industry, and Linux is at the centre of this focus. Prototypes implemented by RITE in the Linux kernel that were presented at LinuxCon included methods for keeping the congestion window appropriately open for bursty traffic (new-CWV), Faster retransmissions for application limited flows (RTO restart and TLP restart), Redundant bundling to avoid retransmissions for thin streams and bringing hybrid delay-based congestion control to the Linux kernel for less queue build up in bottlenecks.

Figure 5.2: LINUXCON Europe, 2015, including the talk on RITE prototypes implemented in the Linux kernel

UiO organised an invited session on low-latency interactive video, focused on RITE topics and RMCAT, at the Packet Video Workshop, December 2013. The session featured three papers and was chaired by Stein Gjessing of UiO.

The first workshop on Large Scale Distributed Virtual Environments on Clouds and P2P was a RITE-featured workshop that focused on the RITE use-case of online gaming. Andreas Petlund from SRL co-chaired the workshop, and several of the accepted papers addressed network latency for online games. The workshop was held in conjunction with the Euro-Par conference in Aachen, August 2013.

RITE was present at MMSYS 2013 in Oslo, Norway. The RITE FP7 factsheet was handed out to the MMSYS participants. RITE, as represented by SRL, also had a poster explaining the key topics for the project as well as some preliminary results.

A latency tutorial based on RITE material was held for high-school and college teachers of Norway at the annual teachers' theme day at the University of Oslo in October 2014.

6 Standards

6.1 Standardisation Highlights

IETF Direction Setting : RITE partners have exerted an unprecedented level of influence on the Transport Areas of the IETF, IRTF and IAB, particularly impressive given the small size of the project. RITE has been instrumental in refocusing the agenda of the IETF towards latency, evidenced by being highly influential in the creation of two completely new IETF working groups during the project, and a third in the lead-up to the project, as well as providing three chairs of other groups.

Demonstrating European Leadership: RITE emerged as a new highly visible force demonstrating the leadership of European research and industry across the transport performance areas of the IETF. The project demonstrated greater depth and breadth than the major US companies (Google, Facebook, Netflix, Microsoft, Apple). that have been recruiting heavily in these areas. Each has only been visible in the group that covers their particular interest, whereas RITE has been influential across them all.

IETF specifications delivered: RITE partners have achieved an impressive ten new Internet Drafts adopted onto the standardisation agenda of the IETF, of which five have already been approved as RFCs. Partners have produced a further ten Internet Drafts being considered for adoption.

Influencing Other Standards Bodies : It is well-known that other standards bodies tend to follow the direction set by the IETF, so RITE influence over the IETF has been designed to percolate through to the rest of the industry. Nonetheless, RITE partners have also had a small but tangible explicit influence on bodies such as the 3GPP through informal and formal standards liaison processes.

The above general highlights have been evidenced by delivery of the following specific achievements:

- Influential in the launch of 2 new IETF working groups (WGs), on Active Queue Management (AQM) and Transport Services (TAPS);
- 3 RFCs published: as outputs of the AQM (1) and the TCP maintenance (2) WGs of the IETF;
- An additional 2 specifications approved for publication as IETF RFCs, now in the RFC Editor formalities stage (also outputs of the TCPM and AQM WGs);
- A further 5 specifications adopted onto the IETF's standardisation agenda, and progressing through the WG phase, on their way to publication as IETF RFCs. One in the AQM WG is in the last call for comments. Another in the Transport Area WG is waiting for liaison with the 3GPP to complete (1 Dec 2015) before the WG phase con conclude. The other two are progressing through the RTP Media Congestion Avoidance Techniques (RMCAT) WG;
- A call to adopt a 6th draft onto the IETF's standardisation agenda is in progress in the TCP maintenance WG. Support during the Nov'15 meeting was unanimous, but mailing list ratification is also necessary;
- A further 10 drafts resulting from work later in the project are being pursued for adoption onto the IETF's standardisation agenda. Other sources of funding are being sought to continue with these 10, as well as the 7 adopted drafts that still require further processing;
- Most popular exhibition stand at the July'15 IETF, educating on latency via RITE animated video and quiz, and demonstrating "Ultra Low Latency for All" using an interactive football match video over a portable variant of ALU's testbed;
- Thought Leadership: RITE partners have led a shift in architectural thinking concerning the latency potential of Explicit Congestion Notification (ECN);
- Arranged liaison between the IETF and the IEEE & 3GPP concerning the design of protocols in these other standards bodies that interact with Explicit Congestion Notification (ECN) in IP (v4 & v6).

6.2 IETF, IRTF & IAB

This section outlines the standardisation activities of RITE project through the Internet Engineering Task Force (IETF) and its related organisations, the Internet Research Task Force (IRTF) and the Internet Architecture Board (IAB).

Open standards document are the key foundation of the Internet, where the Internet Engineering Task Force (IETF) is the principal Internet standards organisation, overseen by the IAB and supported by the IRTF. The IETF meets three times a year and produces the RFC-series of documents. A Request for Comments (RFC) is a formal document from the Internet Engineering Task Force (IETF) that is the result of working group process, peer review by the working group, expert review and a final general review process by interested parties. Some RFCs are informational in nature. Of those that are intended to become Internet standards (including Best Current Practice, BCP), the final version of the RFC becomes the standard and no further comments or changes are permitted, unless the document is formally replaced. This publication method therefore demands much more stringent review than commonly found for initial standards published by other organisations. It also dictates a longer standards process than some other standards development organisations.

6.2.1 IETF/IRTF/IAB Direction Setting

New IETF working groups (WGs) RITE partners were influential in setting up whole new IETF WGs to standardise work on latency. We helped create one WG before starting the project and two during the lifetime of the project:

RTP Media Congestion Avoidance Techniques (RMCAT) WG: Before RITE launched, the IETF had started an ambitious programme of work to standardise real-time communication between Web browsers. It became apparent that congestion avoidance would block progress if the necessary expertise was not gathered together. Initially, the pre-existing IRTF Internet congestion control research group (ICCRG) was used to oversee this work. However, in the run-up to the start of the RITE project, a number of RITE partners were involved in the Internet Architecture Board (IAB) workshop and ensuing discussions that led to the formation the RMCAT working group to address this question. Nonetheless, RMCAT was chartered about 9 months earlier than the launch of RITE, so its creation is not strictly an output of the RITE project.

Active Queue Management (AQM) WG: The same IAB workshop and discussions led to the formation of this important new WG. The bufferbloat project had helped people understand that AQM would be needed to complement RMCAT to address the wider problem of queuing delay. RITE partners were heavily involved in all the activity required to charter this WG both behind the scenes and visible. The AQM WG held its first meeting in Nov'13—the first week of the RITE project—perfectly timed as a dissemination channel for the project.

Transport Services (TAPS) WG: RITE partners are widely acknowledged as the leading forces behind the creation of this WG. The process to approve the TAPS WG charter was tortuous, but it was eventually agreed in Sep'14, and the TAPS WG held its first meeting in Jul 2014, and it is still active and progressing. The scope of TAPS concerns the API to transport services. So, although expected latency benefits were presented, it was clear that latency would be just one aspect of this wider problem, which is therefore being addressed in another EU project (NEAT).

New IETF Initiatives and Thought Leadership :

Unlocking the Latency Benefits of Explicit Congestion Notification (ECN): During the life of the RITE project, partners have changed the way IETF delegates think about ECN. Previously, ECN had generally been considered solely as a replacement for loss. Delegates in the transport area now understand and accept that ECN has much greater potential as a way to reduce the delay taken to signal available capacity, and that this will require the meanings of ECN and loss signals to be separated. This subtle but important architectural shift has been achieved through numerous presentations and demonstrations of research, particularly in IETF Transport Area plenaries, and through the process of publishing an RFC on the Benefits of ECN. RITE's second animated video seeks to share this insight more widely with a simple message for industry decision-makers.

TCP Prague ad hoc "Birds of a Feather" (BoF) meeting: At the July'15 IETF in Prague, ALU demonstrated the RITE DualQ Coupled AQM live by remote login to the testbed in Antwerp, for the benefit of about 400 delegates in the AQM WG. The audience was so enthused by the near-removal of all queuing delay, that we decided to arrange an ad hoc BoF the next evening, to bring together developers working on evolution of DCTCP. Despite the short notice and everyone's hectic schedules, all the main developers were present, from Apple (iOS/MacOS/FreeBSD), Google (Android/Linux), Microsoft (Windows), NetApp (FreeBSD), etc. and the room was packed with about thirty people, all leaders in the evolution of TCP-related congestion control. An agenda of work was agreed and prioritised and the work is progressing.

RITE Bits-N-Bites exhibition stand: (see also Sec. 5) The IETF sets aside an evening for exhibition stands to demonstrate technology in a social setting. The RITE project stand in Jul'15 was by far the most popular, also bearing in mind attendance at this IETF in Prague was record-breaking and about a thousand of those delegates registered for Bits-N-Bites. Our stand was located outside the main hall, beside the entrance to the main foyer, so the crowds started at midday once we set up, and did not abate until late evening. The stand took a two -pronged approach:

- The first RITE animated video "Latency not just Bandwidth" ran continuously with an on-line quiz. This allowed those less involved in the transport area to understand the importance of latency reducing techniques.
- A portable version of the ALU testbed demonstrating "Ultra-Low Latency for All", which show-cased the stunning responsiveness of finger-gesture interactions with a panoramic video of a football match. Two laptops demonstrated the same application, comparing the poor responsiveness of a legacy TCP stack with that of a scalable TCP (Data Centre TCP). For those interested in how it worked, on demand we ran a wide-screen presentation explaining the interaction between our DualQ Coupled AQM and DCTCP, and how the AQM enables incremental deployment of scalable TCP traffic on the public Internet.

Internet Architecture Board (IAB) workshop on Stack Evolution in a Middlebox Internet (SEMI): Three RITE partners were invited to this workshop based on position-paper submissions, and between us we presented on tunnelling new protocol designs through middleboxes (joint work with the Trilogy 2 project) and on TAPS (after this, TAPS work continued outside RITE).

IETF/IRTF Community Service

Chairing: Three RITE participants co-chaired relevant WGs/RGs throughout the duration of the RITE project:

- Michael Welzl (UiO) and David Ros (SRL) co-chair the IRTF Internet Congestion Control Research Group (ICCRG);
- Gorry Fairhurst (UoA) co-chairs the IETF Transport Area Working Group (TSVWG)

Chairing involves a continuous heavy workload, particularly in the case of the TSVWG, which handles standardisation of all transport-related technologies that do not fit into one of the other WGs (which includes maintenance of technologies that used to have their own WG). The chair has to understand and track all the drafts progressing through the WG, as well as assessing the relevance of each of the continual stream of drafts proposing new areas of work (some latency-related, some not). The chair also has to liaise with the Area Director and with other standards bodies.

Technical Review: During the lifetime of RITE, project participants produced numerous technical reviews of the work of others, ensuring that all technology progressing through the IETF meets high standards. Reviews range in size and detail, but most are far more thorough and detailed than a typical peer review of a research paper for a journal or conference. This is perhaps not surprising given that the implications of an error would be far more serious. Unlike delegates to other standards bodies such as the 3GPP or ITU, IETF delegates include the engineers and implementers who are developing the technology and often the researchers who created the ideas as well. The serious consequences of failure and the nature of the delegates as well as the fact that an RFC cannot be amended once published are all factors that explain the IETF's extremely demanding reviewing environment. Producing an RFC that has survived this hostile review process is an extremely demanding exercise. But producing a detailed review is in itself highly demanding (and often poorly recognised) work.

6.2.2 IETF Specifications

Before the IETF publishes a specification as an RFC, it has to pass through multiple stages. Below, the specifications submitted to the IETF by the RITE project are grouped by the stage they reached by the date of this report. In each group, first the citations of the documents are listed, along with their current state and intended status. The abbreviations used for the status of an IETF RFC are: STD: Standards track, EXP: Experimental track, BCP: Best Current Practice or INF: Informational. Then a brief couple of sentences describes each specification, using the same order.

3 RITE specifications published as IETF RFCs:

- F. Baker, G. Fairhurst, "IETF Recommendations Regarding Active Queue Management," <RFC 7567> Status: Best Current Practice (BCP 197) (Jul 2015)

- M. Kühlewind, R. Scheffenegger, and B. Briscoe, "Problem Statement and Requirements for Increased Accuracy in Explicit Congestion Notification (ECN) Feedback"
  <RFC 7560> Status: Informational (INF) (Aug 2015)
- G. Fairhurst, A. Sathiaseelan, R. Secchi, "Updating TCP to support Rate-Limited Traffic"
  <RFC 7661> Status: Experimental (EXP) (Oct 2015)

Very brief summaries of each are given below. But for fuller details the final report from WP2 on host-network interaction (D2.3) explains the first two, and the final report from WP1 on end-to-end techniques (D1.3) introduces the third. These deliverables also include the full text of each in an appendix.

AQM Recommendations was the first and primary milestone of the IETF's new AQM WG. It specifies best current practice in AQM design, setting the framework for standardisation of specific AQM mechanisms.

More accurate ECN feedback is needed because the existing ECN feedback in TCP only indicates the existence of congestion in each round trip, not the extent. The richer information has been shown to nearly eliminate queuing delay and make TCP much more responsive to changes in available capacity. Publishing Requirements for Accurate ECN Feedback in TCP was a pre-requisite for adoption of candidate proposal(s).

The change to TCP to support Rate-Limited Traffic (termed New-CWV) changes the evolution of the TCP congestion window following an idle period, given the previous congestion window validation (CWV) procedure was too conservative.

2 RITE specifications approved for publication as IETF RFCs:

These have been approved by the Internet Engineering Steering Group (IESG), so they now only have to pass through the editorial formalities of the final RFC Editor phase:

- P. Hurtig, A. Brunstrom, A. Petlund, M. Welzl, "TCP and SCTP RTO Restart,"
  <draft-ietf-tcpm-rtorestart-10> Intended RFC status: "EXP". Current state: Approved for publication as RFC (2015-11-19)
- M. Welzl and G. Fairhurst, "The Benefits of using Explicit Congestion Notification (ECN)"
  <draft-ietf-aqm-ecn-benefits-07.txt> Intended RFC status: "INF". Current state: "Approved-announcement to be sent" (2015-11-01)

Very brief summaries of each are given below. But for fuller details the final report from WP1 (D1.3) explains the first, and the final report from WP2 (D2.3) introduces the second. These deliverables also include the full text of each in an appendix.

RTO Restart provides faster loss recovery when there is a small amount of outstanding data, which gives more predictable low latency for short or application-limited flows.

The Benefits of using ECN explains why it is useful to deploy ECN and how not to obstruct others who want to deploy it. It also looks forward to potential ECN latency benefits through future changes to standards.

5 RITE drafts adopted by an IETF working group:

A huge number of Internet Drafts are submitted as proposals to the IETF. Such a draft has no status unless sufficient IETF participants agree that it should be adopted to address a chartered milestone of a WG. The IETF then allows the string 'ietf' to be included in the file-name, as is the case for the 5 drafts below. It is then highly likely, but not inevitable, that the IETF will eventually publish the draft as an RFC. The first three of the drafts below are in the final stages of WG consideration. Indeed the Circuit Breakers draft has progressed beyond the WG stage on its way to approval by the Internet Engineering Steering Group (IESG).

- N.Kuhn, P. Natarajan, N. Khademi and D. Ros. "AQM Characterization Guidelines."
  <draft-ietf-aqm-eval-guidelines-09.txt> (2015-11-27). Intended RFC status: "INF". Current state: In AQM Working Group last call, Sep 2015.
- G. Fairhurst, "Network Transport Circuit Breakers"
  <draft-ietf-tsvwg-circuit-breaker-10> (2015-11-26) Intended RFC status: "BCP". Current state: Area Director Evaluation, (2015-10-05)

- B. Briscoe, J. Kaippallimalil and P. Thaler "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP",
  <draft-ietf-tsvwg-ecn-encap-guidelines-04> (2015-10-08). Intended RFC status: "BCP". Current state: TSVWG WG document; Liaison Correspondence in Progress with 3GPP (promised 1 Dec 2015) and completed with IEEE.
- M. Welzl, S. Islam, S. Gjessing, "Coupled Congestion Control for RTP Media," <draft-ietf-rmcat-coupled-cc-00.txt> (2015-09-14). Intended RFC status: "EXP". Current state: RMCAT WG document;
- D. Hayes, S. Ferlin, M. Welzl, "Shared Bottleneck Detection for Coupled Congestion Control for RTP Media," <draft-ietf-rmcat-sbd-03.txt> (2015-10-19). Intended RFC status: "EXP". Current state: RMCAT WG document;

Very brief summaries of each are given below. But for fuller details the final report from WP2 (D2.3) explains the first three, and the final report from WP1 (D1.3) introduces the last two. These deliverables also include the full text of each in an appendix.

The Guidelines on AQM Characterization specify the types of experiments that should be conducted to establish the merit of a particular AQM, particularly to determine its suitability for standardization.

The Network Transport Circuit Breakers draft explains what is meant by the term "Network transport Circuit Breaker"(CB). It describes how circuit breakers can be used to avoid the implications (particularly queuing delay) of severe long-term congestion and defines requirements for building a circuit breaker for use in the Internet.

The Guidelines for adding Congestion Notification to protocols layered below IP is intended to help designers of lower layer protocols to interface to the ECN facility in IP. It is written for protocol designers in other IETF WGs, and/or other standards bodies, who might be experts in protocol layering and tunnelling, but not necessarily experts in congestion control.

Coupled Congestion Control addresses the common scenario where a real-time application is controlling multiple related streams, to take advantage of combining the congestion information from each stream. Examples of how to add coupling to existing congestion controls are given, including Google's CC and Cisco's NADA.

Shared Bottleneck Detection complements Coupled Congestion Control above. Because streams between the same real-time application endpoints might take different paths through the network, it provides a mechanism to detect if the congestion information from each stream is likely to be due to the same bottleneck.

10 'Individual' RITE drafts have been submitted to be considered for adoption onto the IETF's standardisation agenda:

'Individual' drafts have no formal status in the IETF until they are adopted. So, many individual drafts are submitted that there is no guarantee the proposers will even be allowed to present their ideas. Even if an individual draft is adopted it may have had to be presented, modified, and re-presented multiple times, each time having to compete with the continual flow of other new drafts. Even if the authors of a new draft are already well-known and respected, it might stlll typically require a year of presentations (3 IETF meeting cycles) before it would even be considered for adoption as WG business. That is an average for those that are eventually successful, but it can take much longer. Therefore, for those that are not successful, you can go on trying for years without knowing whether you ought to give up.

- B. Briscoe, M. Kühlewind and R. Scheffenegger, "More Accurate ECN Feedback in TCP"
  <draft-kuehlewind-tcpm-accurate-ecn-05> (2015-10-19). Intended RFC status: "EXP". Current state: TCPM WG adoption call in progress;
- K. De Schepper, B. Briscoe (Ed), O. Bondarenko and I. Tsang, "DualQ Coupled AQM for Low Latency, Low Loss and Scalable Throughput"
  <draft-briscoe-aqm-dualq-coupled-00> (2015-08-07). Intended RFC status: "EXP". Current state: Active Individual Internet Draft. Presented once to the AQM WG and once to the DISPATCH WG (as potential 'customers' of the work).
- K. De Schepper, B. Briscoe (Ed) and I. Tsang, "Identifying Modified Explicit Congestion Notification (ECN) Semantics for Ultra-Low Queuing Delay" <draft-briscoe-tsvwg-ecn-l4s-id-00> (19-Oct-2015) Intended RFC status: "EXP". Current state: Active Individual Internet Draft. Not yet presented.
- N. Khademi, M. Welzl, G. Armitage, G. Fairhurst "TCP Alternative Backoff with ECN (ABE)"

<draft-khademi-alternativebackoff-ecn-01> (2015-09-21) Intended RFC status: "EXP". Current state: Active Individual Internet Draft. Presented twice to the TCPM WG.

- R. Barik, S. Ferlin, M. Welzl. "A Linked Slow-Start Algorithm for MPTCP," <draft-barik-mptcp-lisa-00>, (2015-10-19). Intended RFC status: "STD". Current state: Active Individual Internet Draft. Presented once to the MPTCP WG.
- A. Zimmermann and R. Scheffenegger and B. Briscoe "The TCP Echo and TCP Echo Reply Options," <draft-zimmermann-tcpm-echo-option-00> (2015-06-30) Intended RFC status: "STD". Current state: Active Individual Internet-Draft. Presented once to the TCPM WG.
- B. Briscoe , "Inner Space for TCP Options," <draft-briscoe-tcpm-inner-space-01> (2015-04-30) Intended RFC status: "EXP". Current state: Expired Individual Internet-Draft. Presented once to the TCPM WG and once to the TCPINC WG.
- B. Briscoe , "Inner Space for All TCP Options," <draft-briscoe-tcpm-inspace-mode-tcpbis-00> (2015-09-10) Intended RFC status: "EXP". Current state: Expired Individual Internet-Draft. Presented once to the TCPM WG
- T. Moncaster, B. Briscoe and A. Jacquet "A TCP Test to Allow Senders to Identify Receiver Non-Compliance," <draft-moncaster-tcpm-rcv-cheat-03> (2014-07-03) Intended RFC status: "EXP". Current state: Expired Individual Internet Draft. Presented 3 times to the IETF TCPM WG.
- D. Hayes, D. Ros, L. Andrew and S. Floyd, "Common TCP Evaluation Suite," <draft-irtf-iccrg-tcpeval-01> (2014-07-04) Intended RFC status: "INF'. Current state: Expired Individual Internet-Draft.

Very brief summaries of each are given below. But for fuller details the final report from WP2 (D2.3) explains the first four, and the final report from WP1 (D1.3) explains the rest. These deliverables also include the full text of each in an appendix (in the case of LISA, the appendix contains the paper on which the Internet draft was based).

Accurate ECN (AccECN) Feedback is a candidate proposal to satisfy the requirements recently published as an RFC (see earlier). The IETF puts up a high bar against changes to the main TCP protocol, so it has been an achievement to get the requirements adopted then standardised and now get a candidate solution adopted all within the three years of the RITE project.

The DualQ Coupled AQM documents the AQM developed in the RITE project that enables ultra-low latency, low loss and scalable throughput (L4S) service. It isolates any legacy TCP traffic in a separate queue from modern L4S traffic, which keeps its queuing delay ultra-low by using modern scalable congestion controls and a new type of ECN, which also removes congestion losses.

Identifying ECN Semantics for Ultra-Low Queuing Delay specifies the identifier for new type of ECN mentioned above. This and the previous two items (AccECN and the DualQ Coupled AQM) all fit together as the three main parts of the system that provides the new L4S service.

Alternative Backoff with ECN (ABE) proposes a sender-side only change to TCP which addresses the under-utilization problem for long round trip time (RTT) flows when passing through a modern AQM with a shallow target queuing delay. It is solely proposed for ECN-marked flows, on the basis that ECN marking can only emanate from an AQM.

Inner Space for TCP Options specifies a way to embed framing headers within the TCP payload stream. This allows new extensions to TCP to traverse middleboxes like proxies, but it is also designed so that the two endpoints can determine if they both support the new framing approach and any new extension it enables, without any extra rounds of protocol negotiation delay.

The Inner Space protocol can be switched into different modes, typically for the duration of a connection. Inner Space for *All* TCP Options switches TCP into a mode that can deliver unordered datagrams. In this specification, unordered delivery is solely used for control options. Nonetheless, it demonstrates how TCP could be modified to deliver multiple streams within one TCP flow that are ordered within themselves, but each stream can be delivered to the application in any order. This removes the head-of-line blocking delay, providing the same facility as Google's QUIC protocol. However QUIC is built over UDP which is blocked on about 15% of Internet paths, whereas Inner Space is designed to traverse nearly all middleboxes.

The Linked Slow start Algorithm (LISA) addresses the problem where parallel slow-starts in multiple sub-flows of multipath TCP (MPTCP) can overwhelm a bottleneck. It couples the initial window of each sub-flow to remove the adverse effect.

6.2.3 IETF/IRTF/IAB Presentations

As explained above, presentation of a new draft at the IETF is in itself a small victory; to be heard above the continual stream of other new drafts competing for everyone's attention. It then requires persistence to keep demonstrating that there is interest in your work and it is evolving to become suitable for adoption. In the unlikely event that the work is adopted, it is necessary to continue to demonstrate the level of activity around the work, whether implementation, testing, evaluation or reviewing. This also demands that you have to compete for a presentation slot at each IETF.

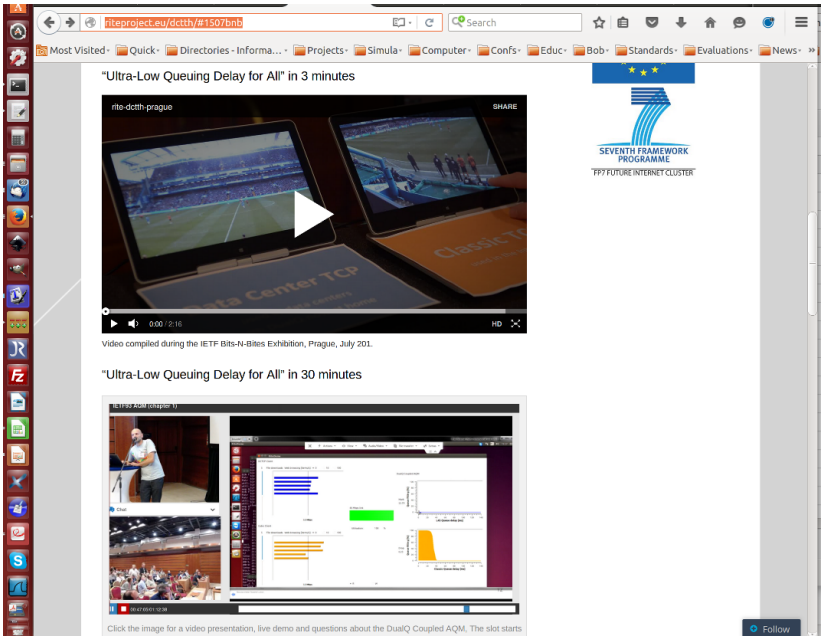

Figure 6.1: Screen shot of a selected page from the RITE Web site
This page provides video material to support ALU's demo and presentation to the IETF AQM WG

Therefore, the sheer number of presentations (91) that WG chairs have invited RITE partners to present at the IETF demonstrates the high regard the project partners have earned, and the impact the project has made on the community. They are listed below:

- R. Barik, S. Ferlin and M.Welzl, "LISA: A Linked Slow-Start Algorithm for MPTCP", MPTCP WG, IETF-94, Yokohama.
- S. Islam, "Coupled congestion control with RTP", RMCAT WG, IETF-94, Yokohama.
- D. Hayes, "Shared Bottleneck Detection for Coupled Congestion Control for RTP Media", RMCAT WG, IETF-94, Yokohama.
- N. Khademi, "TCP Alternative Backoff with ECN" TCPM WG, IETF-94, Yokohama.
- B. Briscoe, M. Kühlewind (presenter), R. Scheffenegger, "More Accurate ECN Feedback in TCP", TCPM WG, IETF-94, Yokohama.
- B. Briscoe, J. Kaippallimalil and P. Thaler "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP", TSVWG WG, IETF-94, Yokohama.
- G. Fairhurst, "Network Transport Circuit Breakers" Transport Area (tsvwg) Working Group, TSVWG WG, IETF-94, Yokohama.
- K. De Schepper, I. Tsang, O. Bondarenko and B. Briscoe "Ultra Low Queuing Delay for ALL Applications", DISPATCH WG, IETF-94, Yokohama.

- RITE Project, "Reducing Internet Transport Latency," Bits-N-Bites exhibition stand, IETF-93, Prague.
- N. Kuhn, "AQM Characterization Guidelines", AQM WG, IETF-93, Prague.
- G. Fairhurst and M. Welzl, "The Benefits of using Explicit Congestion Notification (ECN)", AQM WG, IETF-93, Prague.
- B. Briscoe, "Questioning a Fixed Delay Target", AQM WG, IETF-93, Prague.

- K. De Schepper, "DualQ Coupled AQM"', AQM WG, IETF-93, Prague.
- B. Briscoe, "Review: PIE AQM" (In IETF AQM WG, IETF-93, Prague.
- S. Islam, "Coupled congestion control with RTP", RMCAT WG, IETF-93, Prague.
- D. Hayes, "Shared Bottleneck Detection for Coupled Congestion Control for RTP Media", RMCAT WG, IETF-93, Prague.
- N. Khademi, "TCP Alternative Backoff with ECN" TCPM WG, IETF-93, Prague.
- G. Fairhurst, "Network Transport Circuit Breakers" Transport Area (tsvwg) Working Group, TSVWG WG, IETF-93, Prague.
- A. Zimmermann, R. Scheffenegger and B. Briscoe, "The TCP Echo and TCP Echo Reply Options", TCPM WG, IETF-93, Prague.
- B. Briscoe, J. Kaippallimalil (presenter) and P. Thaler "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP", TSVWG WG, IETF-93, Prague.
- B. Briscoe and K. De Schepper, "Scaling TCP's Congestion Window for Small Round Trip Times", IRTF ICCRG, IETF-93, Prague.

- N. Kuhn, N. Khademi, P. Natarajan, D. Ros, "AQM Evaluation Guidelines", 92nd IETF, AQM WG, Dallas, TX, US, Mar 2015
- G. Fairhurst & M. Welzl, "The Benefits and Pitfalls of using ECN" draft-ietf-aqm-ecn-benefits-01, 92nd IETF, AQM WG, Dallas, TX, US, Mar 2015
- M. Welzl, "Coupled Congestion Control for RTP Media (draft-welzl-rmcat-coupled-cc-04)", RMCAT, IETF-92, Dallas, USA.
- D. Hayes, S. Ferlin & M. Welzl, "Shared Bottleneck Detection for Coupled Congestion Control for RTP Media", RMCAT, IETF-92, Dallas, USA.
- G. Fairhurst & R. Secchi (UoA), "Updating TCP to support Rate-Limited Traffic", 92nd IETF, TCPM WG, Dallas, TX, US, Mar 2015
- P. Hurtig, A. Brunstrom (presenter), A. Petlund, M. Welzl, "TCP and SCTP RTO Restart", 92nd IETF, TCPM WG, Dallas, TX, US, Mar 2015
- B. Briscoe, J. Kaippallimalil, P. Thaler, "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP," Transport Area (tsvwg) Working Group, IETF-92, Dallas, TX, US, Mar 2015
- G. Fairhurst, "Network Transport Circuit Breakers" Transport Area (tsvwg) Working Group, IETF-92, Dallas, TX, US, Mar 2015.
- P. Hurtig, A. Brunstrom (presenter), A. Petlund, M. Welzl, "TCP and SCTP RTO Restart", 92nd IETF, TSVWG, Dallas, TX, US, Mar 2015
- K. de Schepper, I. Tsang, O. Bondarenko and B. Briscoe, "Data Centre to the Home," Internet Research Task Force (IRTF) Internet congestion control research group (ICCRG), Dallas, TX, US, Mar 2015.
- B. Trammell, M. Kühlewind, D. Boppart, I. Learmonth, G. Fairhurst, and R. Scheffenegger, "Enabling Internet-Wide Deployment of Explicit Congestion Notification", Internet Research Task Force (IRTF) Internet congestion control research group (ICCRG), Dallas, TX, US, Mar 2015.

- B. Briscoe, "Tunneling through Inner Space," Internet Architecture Board (IAB) workshop on Stack Evolution in a Middlebox Internet, Zürich, Jan 2015
- M. Welzl, G. Fairhurst, "Ossification: a result of not even trying?" Internet Architecture Board (IAB) workshop on Stack Evolution in a Middlebox Internet, Zürich, Jan 2015

- N. Khademi, "AQM Evaluation Guidelines" (draft-kuhn-aqm-eval-guidelines)", AQM WG, IETF-91, Honolulu.
- M. Welzl, "The Benefits and Pitfalls of using Explicit Congestion Notification" (ECN) (draft-ietf-aqm-ecn-benefits-00)", AQM WG, IETF-91, Honolulu.
- D. Hayes, "Shared Bottleneck Detection for Coupled Congestion Control for RTP Media"(draft-hayes-rmcat-sbd-01), RMCAT WG, IETF-91, Honolulu.
- S. Islam, "Coupled Congestion Control for RTP Media" (draft-welzl-rmcat-coupled-cc-04), RMCAT WG, IETF-91, Honolulu.
- B. Briscoe, "Echo Cookie TCP Option," TCPM WG, IETF-91, Honolulu, US
- B. Briscoe, "Inner Space for TCP increased security (tcpinc)," TCP Increased Security (tcpinc) WG, IETF-91, Honolulu, US, Nov 2014
- B. Briscoe, "Inner Space," TCPM WG, IETF-91, Honolulu, US, Nov 2014

- M. Welzl, "TCP and SCTP RTO Restart" (draft-ietf-tcpm-rtorestart-04)", TCPM, IETF-91, Honolulu.

- N. Kuhn, P. Natarajan, D. Ros, N. Khademi, "AQM Evaluation", AQM WG, IETF-90, Toronto.
- F. Baker, G. Fairhurst, "AQM Recommendation (2309bis)", AQM WG, IETF-90, Toronto.
- M. Welzl, G. Fairhurst. "The Benefits to Applications of using Explicit Congestion Notification (ECN)", AQM WG,IETF-90, Toronto.
- M. Welzl, S. Islam, S. Gjessing, "Practical Shared Bottleneck Detection for Coupled Congestion Control", RMCAT WG, IETF-90, Toronto.
- B. Trammell, "TAPS BOF", TAPS BoF, IETF-90, Toronto.
- G. Fairhurst, R. Secchi, A. Sathiaseelaan, "Updating TCP to support Rate-Limited Traffic", TCPM WG, IETF-90, Toronto.
- P. Hurtig, A. Brunstrom, A. Petlund, M. Welzl, "TCP and SCTP RTO Restart", TCPM WG, IETF-90, Toronto.
- B. Briscoe, R. Scheffenegger, M. Kühlewind, "Accurate ECN", TCPM WG, IETF-90, Toronto.
- T. Moncaster, B. Briscoe, A. Jacquet, "A Test To Allow TCP Senders to Identify Receiver Non-Compliance", TCPM WG, IETF-90, Toronto.
- G. Fairhurst, "Circuit Breakers", TSVWG WG, IETF-90, Toronto.
- M. Welzl, G. Fairhurst, "The Benefits to Applications of using ECN", TSVWG WG, IETF-90, Toronto.
- B. Briscoe, R. Scheffenegger, M. Kühlewind, "More Accurate ECN Feedback Reflector", IRTF Data Centre Latency RG, IETF-90, Toronto.
- B. Briscoe, A. Brunstrom, A. Petlund, D. Hayes, D. Ros, I-J. Tsang, S. Gjessing, G. Fairhurst, C. Griwodz, M. Welzl, "Reducing Internet Latency", IRTF Data Centre Latency RG, IETF-90, Toronto.

- N. Kuhn, P. Natarjan, D. Ros, N, Khademi. "AQM Evaluation", AQM WG, IETF-89, London.
- G. Fairhurst, F. Baker. "AQM Recommendation", AQM WG, IETF-89, London.
- M. Welzl, G. Fairhurst. "The Benefits to Applications of using Explicit Congestion Notification (ECN)", AQM WG, IETF-89, London.
- M. Welzl, S. Islam, S. Gjessing. "Coupled Congestion Control for RTP media". RMCAT WG, IETF-89, London.
- D. Hayes, S. Ferlin & M. Welzl, "Shared Bottleneck Detection," RMCAT WG, IETF-89, London.
- M.J. Montpetit, I. Zhovnirovsky, B. Reuther. "How could Transport Services be provided?" TAPS BoF, IETF-89 London.
- G. Fairhurst, R. Secchi, Z. Hossain. "Updating TCP to support Rate-Limited Traffic". TCPM WG, IETF-89 London.
- P. Hurtig, A. Brunstrom, A. Petlund, M. Welzl. "TCP and SCTP RTO Restart". TCPM WG, IETF-89, London.
- M. Kühlewind, R. Scheffenegger, B. Briscoe, "Problem Statement and Requirements for a More Accurate ECN Feedback", TCPM WG, IETF-89, London.
- B. Briscoe, J. Kaippallimalil, P. Thaler. "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP". TSVWG, IETF-89, London.
- B. Briscoe, M. Kühlewind, D. Wagner, K. De Schepper. "ECN; the Identifier of a New Service Model". TSVWG, IETF-89, London.
- G. Fairhurst, F. Baker. "AQM Recommendation," AQM WG, IETF-88, Vancouver.
- N. Khademi, A. Bagayoko, G. Fairhurst, C. Kulatunga, D. Ros, M. Welzl. "AQM Evaluation". AQM WG, IETF-88, Vancouver.
- M. Welzl, S. Islam, S. Gjessing. "Coupled Congestion Control for RTP media". RMCAT WG, IETF-88, Vancouver.
- G. Fairhurst, A. Sathiaseelan, R. Secchi. "TCP new-CWV". TCPM WG, IETF-88, Vancouver.
- P.Hurtig, A. Brunstrom, A. Petlund, M. Welzl. "TCP and SCTP RTO Restart", TCPM WG, IETF-88, Vancouver.
- B. Briscoe, J. Kaippallimalil, P. Thaler. "ECN Encaps Guidelines", TSVWG, IETF-88, Vancouver.
- T. Moncaster, M. Welzl, D. Ros, "Transport Services BOF plan". TSVWG, IETF-88, Vancouver.
- T. Moncaster, M. Welzl, D. Ros. "Transport Services BOF plan". ICCRG, IETF-88, Vancouver.
- M.Khademi, D. Ros, M. Welzl, "Evaluating CoDel, FQ_CoDel and PIE: how good are they really?" ICCRG, IETF-88, Vancouver.
- M. Welzl, S. Islam, S. Gjessing. "Coupled Congestion Control for RTP media". RMCAT WG, IETF-88, Vancouver.
- B. Briscoe, M. Kühlewind, D. Wagner, M. Espinosa. "Immediate ECN". TSVWG, IETF-88, Vancouver.

- M. Ford, “Latency Workshop Report,” TSVAREA, IETF-88, Vancouver.

- F. Baker, “AQM Recommendation”. AQM BoF, IETF-87, Berlin.
- N. Kuhn, G. Sarwar, E. Lochin, R. Boreli, A. Mifdaoui. “MPTCP Scheduling”. ICCRG, IETF-87, Berlin.
- M. Welzl, S. Islam, S. Gjessing. “Coupled Congestion Control for RTP media”. RMCAT WG, IETF-87, Berlin.
- G. Fairhurst, A. Sathiaseelan, R. Secchi. “TCP new-CWV”. TCPM WG, IETF-87, Berlin.
- A. Brunstrom. “TCP RTO Restart”. TCPM WG, IETF-87, Berlin.

- G. Fairhurst, A. Sathiaseelan, R. Secchi. “TCP new-CWV”, TCPM WG, IETF-86, Orlando.
- B. Briscoe, J. Kaippallimalil, P. Thaler. “ECN Encaps Guidelines”. TSVWG, IETF-86, Orlando.
- M. Welzl. “Coupled Congestion Control for RTP media”. RMCAT WG, IETF-86, Orlando.
- G. Fairhurst, B. Briscoe. “Advice on network buffering”. ICCRG, IETF-86, Orlando.

- G. Fairhurst, A. Sathiaseelan, “Enhancing TCP to support Rate-Limited Traffic”, ICCRG, IETF-85, Atlanta
- B. Briscoe. “ECN Encaps Guidelines”. TSVWG, IETF-85, Atlanta.
- G. Fairhurst, A. Sathiaseelan. “TCP new-CWV”. ICCRG, IETF-85, Atlanta.
- M. Welzl, Per Hurtig, A Brunstrom, A. Petlund, “TCP RTO Restart”. TCPM WG, IETF-85, Atlanta.

6.2.4 IETF/IRTF Summary

The RITE project has had an unprecedented impact on the Transport Areas of the IETF, IRTF and IAB. In recent years, the major US service providers have recruited a significant proportion of the world’s transport performance expertise and they are each attempting to dominate their particular open-source community: Google now employs many of the influential Linux stack developers, while Netflix and NetApp have the top FreeBSD developers, and of course Microsoft and Apple represent their interests.

RITE has given the European companies and institutions a common identity that is more consistently active across all the transport performance related groups of the IETF than any of the major US companies, even though RITE may not be as dominant in certain specific areas (e.g. not as dominant as Google in the tcpm wg). The RITE partners are also perceived as more independent of a commercial position, whereas the opinions of the delegates from the US companies tend to be assumed to be tinged with self-interest, even if the particular individual is behaving perfectly objectively. It will be interesting to see how the distribution of influence changes now that the major mobile operators and equipment vendors are starting to recognise how critical transport protocol behaviour is to end-to-end performance.

Figure 6.2 attempts to provide a visualization of RITE’s impact. The IETF is divided into six areas, the Transport Area being the one relevant to RITE. To visualize the impact on the whole of the transport area, all the transport working groups that were active at the start of the project are shown. It can be seen that RITE has contributed to half of them. It can also be seen that RITE has dominated much of the activity of 5 working groups: aqm, rmcat, taps, tcpm and tsvwg, as well as the iccrg in the IRTF. Note that only the creation phase of the TAPS WG is shown—as already explained, the creation of TAPS was very much led by RITE partners, but the scope of TAPS is wider than the RITE project, so a separate EU project has been created to support activity in TAPS, and TAPS presentations and drafts are being reported there. The expansion of these abbreviations is given below:

wg working group
rg research group
aqm active queue management
rmcat RTP media congestion avoidance (where RTP = Real-Time Protocol)
taps transport services
tcpinc TCP INCrease security
tcpm TCP maintenance & minor mods
tsvwg transport area wg
dclcrg data centre latency control rg
iccrg Internet congestion control rg

The diagram also aims to show how much work and persuasion is necessary to get an IETF RFC published. The

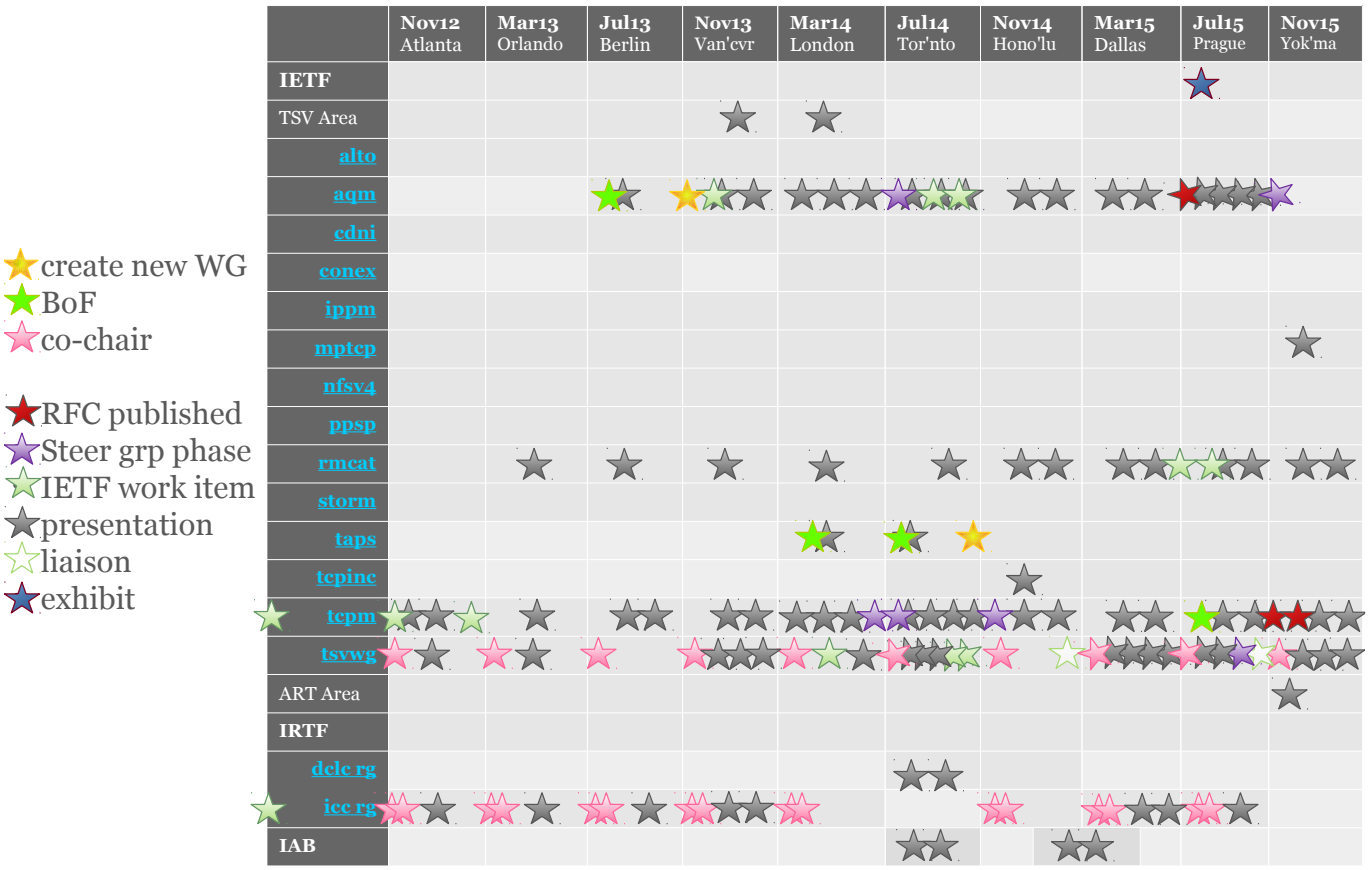

Figure 6.2: The Impact of the RITE Project on the IETF, IRTF and IAB

average time from when the first draft is submitted to RFC publication is perhaps 3 years (the duration of the RITE project). And it should be noted that only a small proportion of initial drafts survive the whole process through to RFC status.

It can be seen that RITE partners managed to squeeze the whole process of starting a completely new working group (aqm) and getting its first RFC published within two years, plus of course all the preparatory work before the first Birds of a Feather (BoF) meeting to socialize the idea of creating it.[12] It can be seen how many times it is necessary to keep presenting work to keep it progressing through the process: first to get it adopted by a WG; then to respond to all the reviews in the WG stage so that it can be approved to go to the steering group stage; then responding to all the Steering Group reviews and the RFC Editor's formalities, before it can finally be published as an RFC.

The TCP maintenance WG is probably the hardest to progress through quickly, of all WGs in the transport area, unsurprisingly given that changes to TCP need to be heavily tested and assessed before releasing on the public Internet. It can be seen that the RITE project focused on getting three drafts adopted in TCPM early in the project, based on work done prior to the formal project launch. All three moved to the IESG (Internet Engineering Steering Group) in the latter half of 2014, and we managed to get two published as RFCs before the end of the project. One of these was a requirements specification that the TCPM WG requested as a prerequisite to standardising any actual technology. So we developed the technology in parallel, and now that the requirements are published, the technology itself is in the process of being adopted for standardisation (expected before the end of the RITE project).

Finally, it will be seen that the RITE project used the IETF to disseminate its outputs more widely than just the transport area. At the top of the figure, our highly popular stand at the IETF Bits-N-Bites exhibition can be seen. And at the bottom, it can be seen that RITE partners were invited to contribute to Internet Architecture Board (IAB) activities. We also demonstration the impressive DualQ Coupled AQM to the DISPATCH WG of the Apps

[12] The RITE project was one of a number of organisations that cooperated in its formation.

& Real-Time (ART) Area, as "customers" of the transport area. The DISPATCH WG is chartered to dispatch new work on real-time application protocols to the appropriate ART WG, but we persuaded the chairs to allow us to give them a "heads-up".

6.3 Other Standards Bodies

The IETF is the sole forum for standardisation of technologies at the transport layer. However, RITE partners interacted with other standards bodies where it was necessary to coordinate or prevent unfortunate conflicts. RITE did not budget for attendance at other standards bodies, therefore these interactions were typically achieved via informal associations with colleagues.

Where an appropriate level of attention was in order, formal liaison channels between the IETF and other standards bodies were used. This approach was taken in one case where we felt a high level of awareness of the guidelines the project had produced was necessary. The Guidelines were for adding congestion notification to protocols layered beneath IP (v4 & v6). RITE partners arranged for the following two liaison statements to be agreed then sent through the formal liaison channels to the IEEE and to the 3GPP:

- IETF Liaison to IEEE 802 "Explicit Congestion Notification (ECN) and IEEE Protocols" (2014-12-02); https://datatracker.ietf.org/liaison/1364/
- IETF Liaison to 3GPP "Explicit Congestion Notification for Lower Layer Protocols" (2015-07-21); Current state: Action needed. https://datatracker.ietf.org/liaison/1424/

The IEEE liaison resulted in the IEEE 802 committee confirming that it has no current relevant work on congestion notification.

The 3GPP liaison unearthed a case where the 3GPP specification of the Universal Terrestrial Radio Access Network (UTRAN) MAC recommends that a base-station (eNodeB) ECN-marks IP packets, giving the example of the Voice over LTE (VoLTE) service. However, an eNodeB is normally not aware of the IP layer, because it acts at the lower PDCP layer.

Also, during the liaison, two 3GPP specifications were found that recommended ECN marking that would be incompatible with the IETF RFC they cited.

The 3GPP is likely to be able to find simple solutions to these issues. It is currently going through the formal process of responding to the IETF's liaison.

7 Dissemination of Open-Source Software

The RITE project had a stated goal to make tools and mechanisms available to a wider audience through Open Source code. This section describes our open source contributions, their main audience, and how to access each of them. Table 7.1 lists the open source software shared through the project. We have divided the table into two sections: **Mechanisms** and **Tools**.

7.1 Mechanisms

The mechanisms that were developed in the project and prototyped in the Linux kernel are shown in the first part of table 7.1. For mechanisms already accepted into the mainline Linux kernel, the kernel version first integrating the mechanism has been listed. For other mechanisms, links to where the patch can be found are shown. The table also points to which RITE deliverable describes the Mechanism in detail.

Getting code accepted into the Linux kernel network stack is a process starting with a functioning prototype of a mechanism developed as a patch or module that is compatible with a certain version of the Linux kernel. For the RITE mechanisms, we identified the 3.18.5 kernel as a candidate for our patches. Once we have a running prototype, extensive experiments needs to be run to verify that the mechanism has the wanted effect, that it does not have any

RITE Mechanisms

Mechanism	Details	How shared
New-CWV patch	D1.3	https://github.com/rsecchi/newcwv/
RTO Restart patch	D1.3	https://github.com/perhurt/rtorestart
TLP restart patch	D1.3	https://github.com/perhurt/rtorestart
ABE patch	D1.3, D2.3, D3.3	http://heim.ifi.uio.no/naeemk/research/ABE/
RDB patch	D1.3	https://bitbucket.org/mpg_code/rdb
CAIA Delay Gradient	D1.3	Mainline Linux kernel (from 4.2)
Dual Queue AQM	D2.3	http://riteproject.eu/dctth/ when approval process completed

RITE Tools

Tool	Details	How shared
TCP Evaluation Suite	QMR4+	https://bitbucket.org/hayesd/tcp-evaluation-suite-public http://trac.tools.ietf.org/group/irtf/trac/wiki/ICCRG_tcpeval
pcap latency analysis tool	D3.1	https://bitbucket.org/mpg_code/tstools_analysetcp
Active measurement AQM detection tool (TADA)	D2.3	https://bitbucket.org/mkargar/tada-tool_automatic_detection_aqm
RITE integrated kernel	D1.3	http://riteproject.eu/resources/rite-integrated-linux-kernel/

Table 7.1: Open-source software shared from the RITE project.

adverse effects and that it does not break other kernel functionality. To upstream the patch to the mainline kernel, the code needs to follow a strict set of formatting and code style requirements. The patch is then submitted to a mailing list, for network mechanisms, it is the *netdev* list, where it is commented and criticised by other developers in the community and tested by other parties than the developers. If the group reaches a positive consensus, the maintainer may consider taking the patch into the official development tree. This is beneficial for several reasons, mainly 1) it saves the developer/project from having to update the patch as the kernel develops, 2) it makes the mechanism available to a large user base.

Currently the RITE project has gotten one of its target mechanisms, the Linux implementation of CAIA Delay Gradients, that has been improved in several ways from the FreeBSD version, into the mainline Linux kernel[13]. There are several patches that have been refined to a point where kernel mailing list submission is viable, most notably RTO Restart, RDB[14] and New CWV.

New CWV

Lead partner: UoA

The New CWV mechanism attempts to choose a good size for the TCP congestion window after idle periods when the window has not been recently validated. The mechanism is useful for applications with bursty traffic patterns. It has been thoroughly tested and the algorithm has been standardised in the IETF as RFC 7661 (experimental). The mechanism and patch was presented to the Linux community at Linux Conference Europe 2015. A patch is being

[13] https://lwn.net/Articles/644440/
[14] https://lkml.org/lkml/2015/10/23/720

refined for kernel mailing list submission.

RTO Restart

Lead partners: KaU, UiO, SRL

Both TCP and SCTP use a retransmission time-out (RTO) timer as a last resort for data loss recovery. The standardised way in which this timer is restarted often unnecessarily extends the loss recovery time by at least one round-trip time (RTT). The RTO Restart (RTOR) mechanism can be used to reduce the retransmission delay when recovery is done by RTO. The mechanism reduces the recovery time for tail loss and is applicable for short flows like most web traffic. It has been thoroughly tested and the the mechanism has come far in the IETF process of becoming an experimental RFC. The RTOR mechanism and patch was presented to the Linux community at Linux Conference Europe 2015. A patch has been developed and prepared for mailing list submission.

TLP Restart

Lead partner: KaU, UiO, SRL

Linux makes use of a non-standardised retransmission scheme called Tail Loss Probe (TLP) that is designed to tackle the slow loss recovery in the end of flows. Like the Retransmission time-out, TLP also makes use of a timer that has the same restart problem. To mitigate this performance problem, RITE has developed a restart strategy for TLP, which we call TLP Restart (TLPR), that is similar to RTOR. The mechanism reduces the recovery time for tail loss and is applicable for short flows like most web traffic. The TLPR mechanism and patch was presented to the Linux community at Linux Conference Europe 2015. A patch has been developed and prepared for mailing list submission.

Alternative Backoff with ECN (ABE)

Lead partners: UiO, UoA

ABE is a simple sender side TCP modification that allows end-hosts to gain performance benefits from Explicit Congestion Notification (ECN) signals from the network. This provides incentives for latency lowering Active Queue Management (AQM) schemes with ECN to be deployed in the network. ABE was developed by RITE in collaboration with the Centre for Advanced Internet Architectures (CAIA), Swinburne University of Technology (Melbourne, Australia). An experimental patch has been made available for both FreeBSD and Linux.

Redundant Data Bundling

Lead partner: SRL

Redundant data Bundling (RDB) is a sender-side only mechanism that tries to preempt the experience of packet loss by redundantly bundling all unacknowledged data as long as the resulting segment does not surpass 1 Maximum Segment Size (MSS). This allows for effectively avoiding high retransmission delays without sending any more packets into the network. The mechanism will be effective for flows that transmit only a small amount of data with relatively large inter-packet interval times or *thin streams*. The RDB mechanism and patch was presented to the Linux community at Linux Conference Europe 2015. A patch has been submitted to the *netdev* mailing list and are currently being discussed[15].

CAIA Delay Gradient Linux implementation (CDG)

Lead partners: SRL, UiO CAIA Delay Gradient is a delay-based congestion control that takes its backoff decisions from estimates of whether the queue is growing or shrinking based on the gradient of delay measurements within RTT periods. The mechanism includes a heuristic that switches the congestion control to ”New Reno” when it detects that

[15] https://lkml.org/lkml/2015/10/23/720

it is competing with other loss-based congestion control mechanisms. The aim of this is to get the queueing delay benefits when alone on the bottleneck or competing with delay-based congestion controls, and still be able to achieve a fair throughput when competing with loss-based congestion control. When turning off the loss competition heuristics, CDG can be deployed as a less-than-best-effort congestion control. Kenneth Klette Jonassen developed the Linux version in collaboration with David Hayes and Andreas Petlund. As of version 4.2, CDG is a part of the mainline Linux kernel.

Dual Queue AQM

Lead partners: ALU, SRL, BT The Dual Queue Coupled AQM is a novel AQM that allows DCTCP (or other scalable / L4S) flows to coexist with Classic TCP flows. It assures that the Classic TCP flows have an approximately equal throughput (TCP-Fair) compared to the DCTCP flows, while the DCTCP flows maintain their ultra low latency capability. Throughput compatibility is achieved by applying a dropping/marking probability to Classic TCP packets that is the square of the marking probability applied to the L4S flows. Low latency for the L4S packets is achieved by scheduling the L4S packets with priority. At the moment of writing, the approval process within ALU for the release of the code is still ongoing. The latest status can be checked at the DCttH web-page http://riteproject.eu/dctth/. When approval is received, also the patches will be published.

7.2 Tools

The bottom section of table 7.1 lists tools developed and used in the RITE project. This is code that can be used for testing, analysis or measurement related to the study of Internet latency. The list also contains a patch set and build script that allows for building an integrated kernel with the key RITE mechanisms available as Linux patches. The table points to deliverables where more details about the tool can be found.

TCP evaluation suite

Lead partner: UiO

This suite of realistic tests implemented in the ns2 simulator facilitates testing of TCP modifications in a number of standard scenarios. The software has built upon previous work by Gang Wang and Yong Xia (NEC China) and David Hayes at Swinburne University of Technology (Melbourne, Australia). The public bitbucket repository contains test scripts along with a tailored ns2 2.35 simulator. The traffic traces are available from the ICCRG tools page. Elements of this tool have been used in the evaluation of a number of RITE proposals.

Pcap latency analysis tool

Lead partners: SRL, KaU

RITE has developed a tool for analysis of PCAP formatted files with a special focus on latency. The tool can be used to get separate statistics on retransmission latency as well as aggregated latency statistic over a group of different connections. Functionality for calculating one-way delay above the base one-way delay (OWD) and compensate for clock drift between sender and receiver computer is implemented. It also has support for analysing the byte-wise latency for a TCP stream when the sender bundles old data with new using RDB or if bundles are made upon retransmissions, which occur in the Linux kernel. The tool was developed by Andreas Petlund, Kristian Evensen, Bendik Rønning Opstad and Jonas Markussen. It is available through a public Bitbucket repository.

Active measurement AQM detection tool (TADA)

Lead partners: SRL, KaU

The last years have seen a lot of effort to show the benefits of AQMs over simple tail-drop queuing and to encourage deployment. Yet it is still unknown to what extent AQMs are deployed in the Internet. RITE has developed an end-to-

end active measurement method to detect AQMs on the path bottleneck. We have developed an active measurement tool, *TADA*, and evaluated our measurement methodology on a controlled experimental testbed. Experimental results show that the proposed approach provides the basis to identify whether an AQM is deployed on the bottleneck. The latest stable version of the tool is available through a public Bitbucket repository.

RITE Integrated Kernel

Lead partners: SRL, KaU, MEGA

The RITE project has prepared a patch set for the 3.18.5 Linux kernel containing the key end-host mechanisms developed in the project. The patch set also includes a build script that will download the correct kernel source, apply the patches, build the kernel copying the running kernel configuration and install the kernel for testing. The mechanisms provided with the kernel are New CWV, RTO Restart, TLP Restart, Redundant Data Bundling and a tuning option for experimenting with different sizes for the TCP initial window upon restart. The integrated kernel allows for interested parties to have a low-threshold way of experimenting with and build upon RITE mechanisms. It is shared from the Resources page on the RITE website.

8 Conclusions

This report has catalogued an impressive array of dissemination activities undertaken by the RITE partners. The promised focus on IETF standardisation has reaped the highest rewards with the creation of two new whole working groups and adoption of ten specifications onto the IETF's standardisation agenda, five of which have already been approved for publication as RFCs.

To achieve such widespread visibility of European research in the IETF, it was necessary to sink a huge amount of effort into standardisation. Nonetheless, partners still managed to devote considerable resources to production of high quality research and code, delivery of high quality educational courses and materials and to organisation of high impact events.

For instance, partners have had papers accepted in high quality conferences, such as IEEE LCN, Networking & WoWMoM as well as NOSSDAV, and all partners collaborated in our highly acclaimed and hugely comprehensive survey of latency reducing techniques that has been published in one of the highest impact journals in communications: IEEE Communications Surveys and Tutorials. Our workshop on Reducing Internet Latency, jointly organised with ISOC, focused the subsequent actions of many of the main thought-leaders and practitioners in Internet performance, from both academia and industry. And our prominent exhibition stand at the July'15 IETF challenged the preconceptions of many hundreds of delegates about Internet quality of service and performance.

Each of the dissemination activities reported here cut across the types of audience we had committed to target in the Description of Work, as quoted in the Introduction. Table 8.1 summarises the type of audience targeted by each activity:

	System operators & software engineers	Industry executives	Technical media	Research & scientific community	European public	Lab events & showcases
Research						
Papers				✓		✓
Education						
Curriculum				✓	✓	
Doctoral				✓		
Materials	✓	✓	✓	✓	✓	✓
Media						
Tech media & blogs	✓	✓	✓	✓	✓	
RITE Web site	✓	✓	✓	✓		✓
Events						
Latency workshop	✓	✓	✓	✓		
EUCNC workshop				✓		
IETF Bits-n-Bites	✓	✓	✓	✓	✓	
Other	✓	✓	✓	✓	✓	✓
Standards						
IETF	✓	✓		✓	✓	
Software						
Mechanisms	✓		✓	✓	✓	✓
Tools	✓			✓		

Table 8.1: Mapping of Each Dissemination Activity in this Report to its Audience