


Measuring ECN

Geoff Huston

Chief Scientist, APNIC

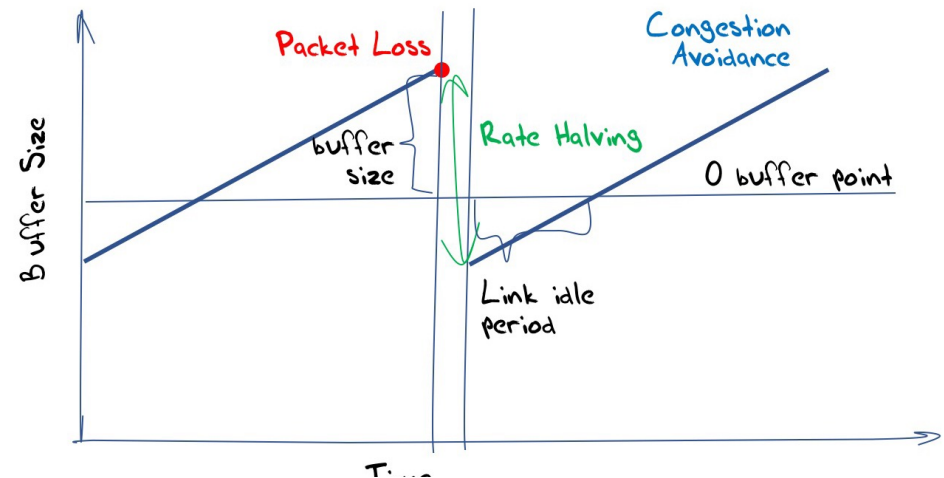
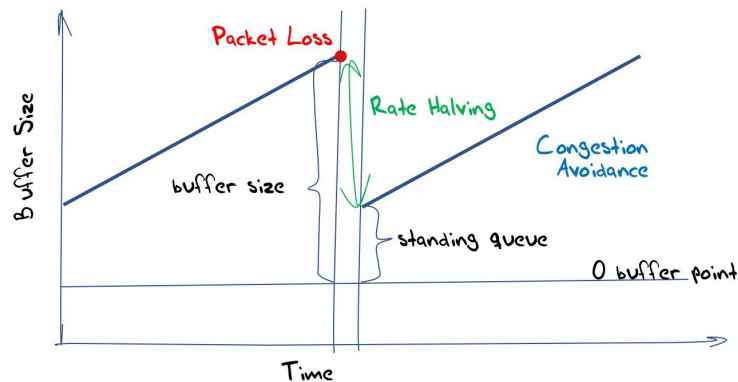
Loss-Based Congestion Control

- Most of today's network traffic uses an extremely simple algorithm that avoids networks collapsing under extreme overload:
 - Continually increase the amount of data being pushed into the network
 - Until a network queue overloads and the sender detects the consequent packet loss
 - Halve the sending rate and do it again!




Refinements

- Change the congestion avoidance inflation algorithm
- Try to detect the difference between isolated damage packet loss and queue overload loss
- Better understand the relationship between network buffers and protocol performance



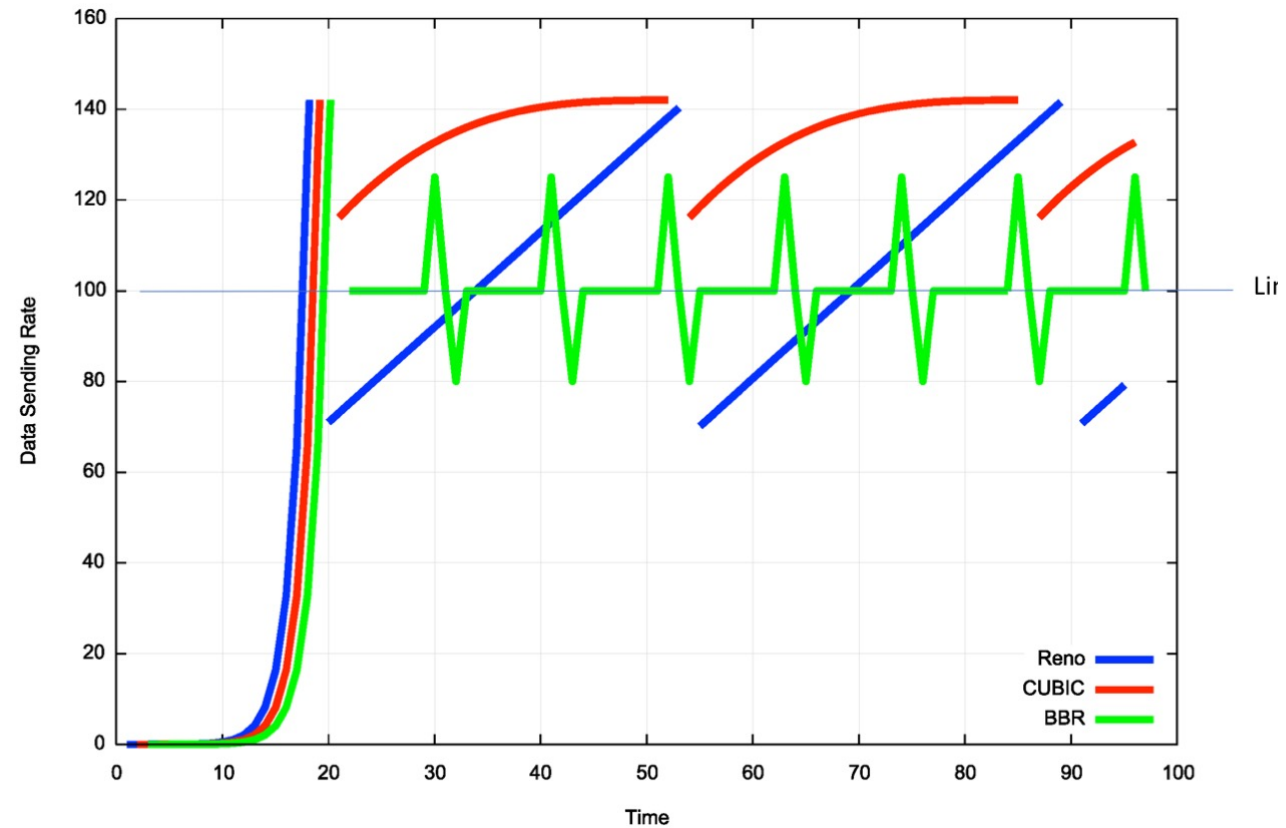
Is there a "better" way?

Trigger the congestion response at the onset of queue formation rather than at the point of catastrophic queue collapse



BBR

- Detect queue formation through pulsed testing and delay sensitivity



Or get the router to help!

May 1990!

The screenshot shows the ACM Digital Library interface. At the top, there's a navigation bar with links like 'Journals', 'Magazines', 'Proceedings', 'Books', 'SIGs', 'Conferences', and 'People'. Below this is a search bar and a secondary navigation bar with links like 'Journal Home', 'Just Accepted', 'Latest Issue', 'Archive', 'Authors', 'Editors', 'Reviewers', 'About', and 'Contact Us'. The main content area displays the title 'A binary feedback scheme for congestion avoidance in computer networks'. Below the title, it lists the editor 'Anita K. Jones' and authors 'K. K. Ramakrishnan' and 'Raj Jain'. It also shows the journal 'ACM Transactions on Computer Systems (TOCS), Volume 8, Issue 2' and the publication date 'Published: 01 May 1990'. A blue arrow points from the text 'May 1990!' to the publication date. The abstract is visible below the title, starting with 'We propose a scheme for congestion avoidance in networks using a connectionless protocol at the network layer.' The page also features a sidebar with a table of contents and a right sidebar with a search bar and a 'PDF' button.

ACM DIGITAL LIBRARY Association for Computing Machinery

Browse About Sign in Register

Journals Magazines Proceedings Books SIGs Conferences People

Search ACM Digital Library

Journal Home Just Accepted Latest Issue Archive Authors Editors Reviewers About Contact Us

A binary feedback scheme for congestion avoidance in computer networks

Editor: Anita K. Jones Authors: K. K. Ramakrishnan Raj Jain [Authors Info & Claims](#)

ACM Transactions on Computer Systems (TOCS), Volume 8, Issue 2 • Pages 158 - 181 • <https://doi.org/10.1145/78952.78955>

Published: 01 May 1990 [Publication History](#) [Check for updates](#)

321 1,446

PDF eReader

Abstract

We propose a scheme for *congestion avoidance* in networks using a connectionless protocol at the network layer. The scheme uses a minimal amount of feedback from the network to the users, who adjust the amount of traffic allowed into the network. The routers in the network detect congestion and set a *congestion-indication* bit on packets flowing in the forward direction. The congestion indication is communicated back to the users through the transport-level acknowledgment. The scheme is distributed, adapts to the dynamic state of the network, converges to the optimal operating point, is quite simple to implement, and has low overhead. The scheme maintains *fairness* in service provided to multiple sources. This paper presents the scheme and the analysis that went into the choice of the various decision mechanisms. We also address the performance of the scheme under transient changes in the network and pathological overload conditions.

Or get the router to help!

Network Working Group
Request for Comments: 2481
Category: Experimental

K. Ramakrishnan
AT&T Labs Research
S. Floyd
LBNL
January 1999

A Proposal to add Explicit Congestion Notification (ECN) to IP

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

This note describes a proposed addition of ECN (Explicit Congestion Notification) to IP. TCP is currently the dominant transport protocol used in the Internet. We begin by describing TCP's use of packet drops as an indication of congestion. Next we argue that with the addition of active queue management (e.g., RED) to the Internet infrastructure, where routers detect congestion before the queue overflows, routers are no longer limited to packet drops as an indication of congestion. Routers could instead set a Congestion Experienced (CE) bit in the packet header of packets from ECN-capable transport protocols. We describe when the CE bit would be set in the routers, and describe what modifications would be needed to TCP to make it ECN-capable. Modifications to other transport protocols (e.g., unreliable unicast or multicast, reliable multicast, other reliable unicast transport protocols) could be considered as those protocols are developed and advance through the standards process.

January 1999!

Or get the router to help!

Network Working Group
Request for Comments: 2481
Category: Experimental

K. Ramakrishnan
AT&T Labs Research
S. Floyd
LBNL
January 1999

A Proposal to add Explicit Congestion Notification (ECN) to IP

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

This note describes a proposed addition of ECN (Explicit Congestion Notification) to IP. TCP is currently the dominant transport protocol used in the Internet. We begin by describing TCP's use of packet drops as an indication of congestion. Next we argue that with the addition of active queue management (e.g., RED) to the Internet infrastructure, where routers detect congestion before the queue overflows, routers are no longer limited to packet drops as an indication of congestion. Routers could instead set a Congestion Experienced (CE) bit in the packet header of packets from ECN-capable transport protocols. We describe when the CE bit would be set in the routers, and describe what modifications would be needed to TCP to make it ECN-capable. Modifications to other transport protocols (e.g., unreliable unicast or multicast, reliable multicast, other reliable unicast transport protocols) could be considered as those protocols are developed and advance through the standards process.

January 1999!

Or get the router to help!

Network Working Group
Request for Comments: 3168
Updates: ~~2474, 2401, 793~~
Obsoletes: 2481
Category: Standards Track

K. Ramakrishnan
TeraOptic Networks
S. Floyd
ACIRI
D. Black
EMC
September 2001

The Addition of Explicit Congestion Notification (ECN) to IP

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This memo specifies the incorporation of ECN (Explicit Congestion Notification) to TCP and IP, including ECN's use of two bits in the IP header.

September 2001!

Then ...

Crickets!

Until around 2022...

L4S!

Internet Engineering Task Force (IETF)
Request for Comments: 9330
Category: Informational
ISSN: 2070-1721

B. Briscoe, Ed.
Independent
K. De Schepper
Nokia Bell Labs
M. Bagnulo
Universidad Carlos III de Madrid
G. White
CableLabs
January 2023


Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service:
Architecture

Abstract

This document describes the L4S architecture, which enables Internet applications to achieve low queuing latency, low congestion loss, and scalable throughput control. L4S is based on the insight that the root cause of queuing delay is in the capacity-seeking congestion controllers of senders, not in the queue itself. With the L4S architecture, all Internet applications ~~could (but do not have to) transition away from~~ congestion control algorithms that cause substantial queuing delay and instead adopt a new class of congestion controls that can seek capacity with very little queuing. These are aided by a modified form of Explicit Congestion Notification (ECN) from the network. With this new architecture, applications can have both low latency and high throughput.


The architecture primarily concerns incremental deployment. It defines mechanisms that allow the new class of L4S congestion controls to coexist with 'Classic' congestion controls in a shared network. The aim is for L4S latency and throughput to be usually much better (and rarely worse) while typically not impacting Classic performance.

ECN Control Loop




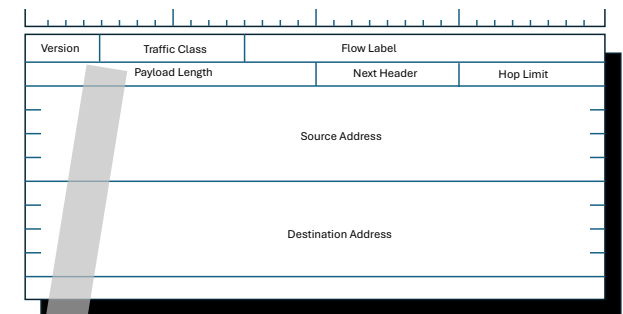
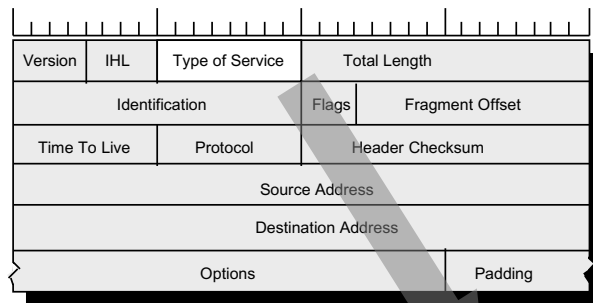
- A router “marks” IP packets at the onset of queue formation with a bit signal
- The Receiver echoes this bit up into the transport protocol reverse flow
- The sender reduces its sending window size (and notifies the receiver that it was performed this window reduction)

ECN Control Loop



- A router “marks” IP packets at the onset of queue formation with a bit signal
- The Receiver echoes this bit up into the transport protocol reverse flow
- The sender reduces its sending window size (and notifies the receiver that it was performed this window reduction)

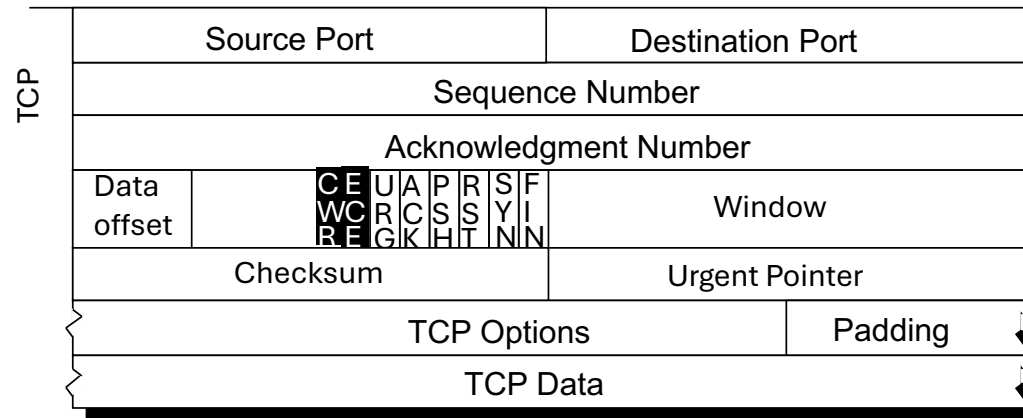
IP Header



ECN Bits

- 0 0 – Non-ECN Capable Transport
- 0 1 – ECN Capable Transport
- 1 0 – ECN Capable Transport
- 1 1 – Congestion Experienced

TCP



ECE – receiver back to sender – CE received

CWR – sender to receiver – Congestion Window Reduced

SYN+ECE+CWR – ECN capable on session start

SYN+ACK+ECE – ECN capable response

ECN Measures

Packet Count (by remote IP addresses) for Feb/March 2025

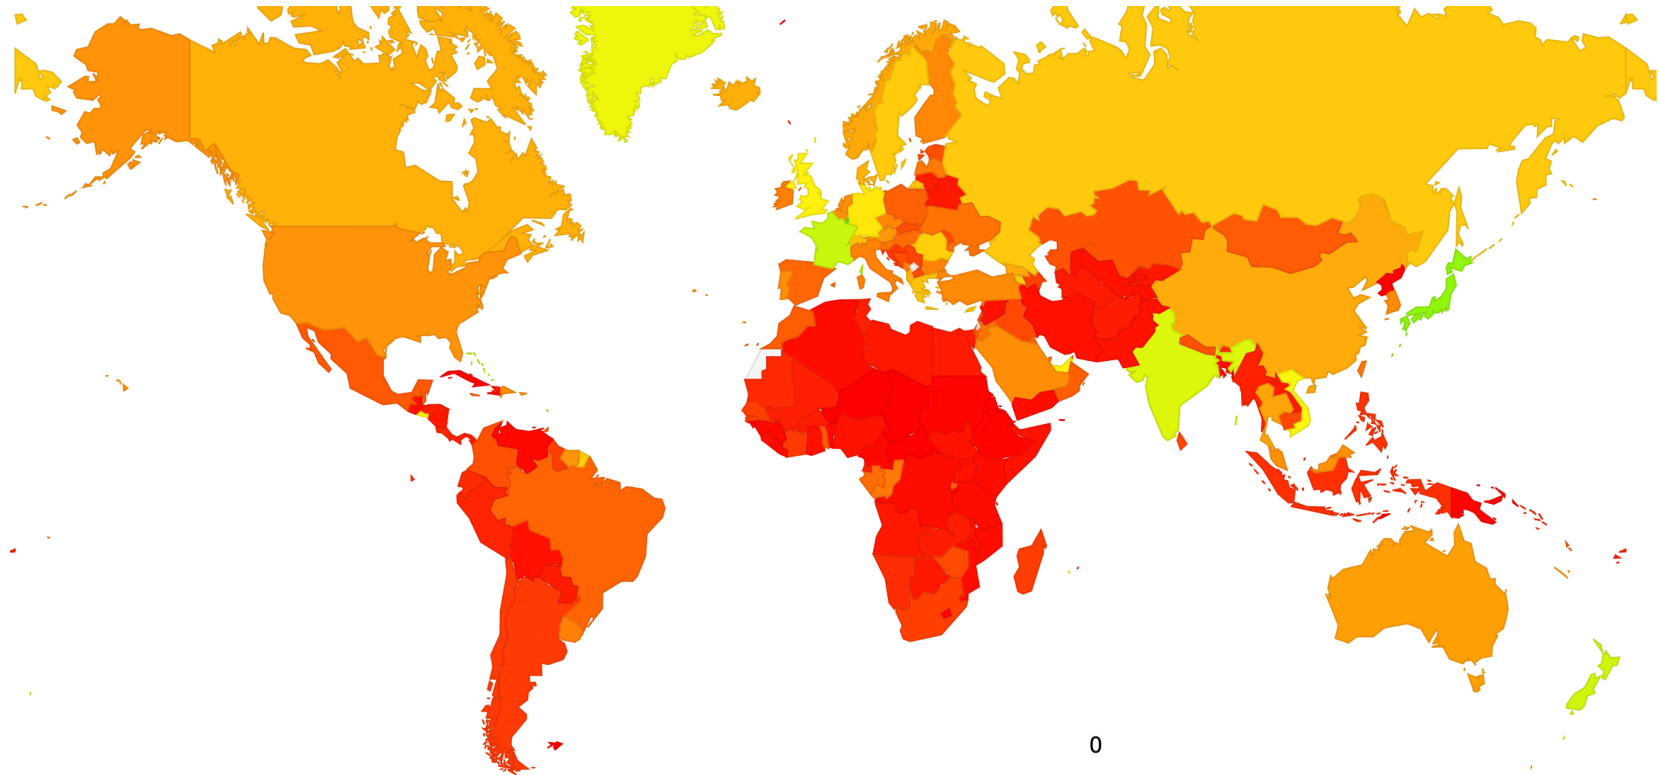
- | | |
|---------------|-----------------------------------|
| 0. IP Sources | 303,545,388 |
| 1. IP ECT | 6,815,753 (2.45% of sources) |
| 2. IP CE | 1,098,965 (16.12% of ECT sources) |

ECN Measures

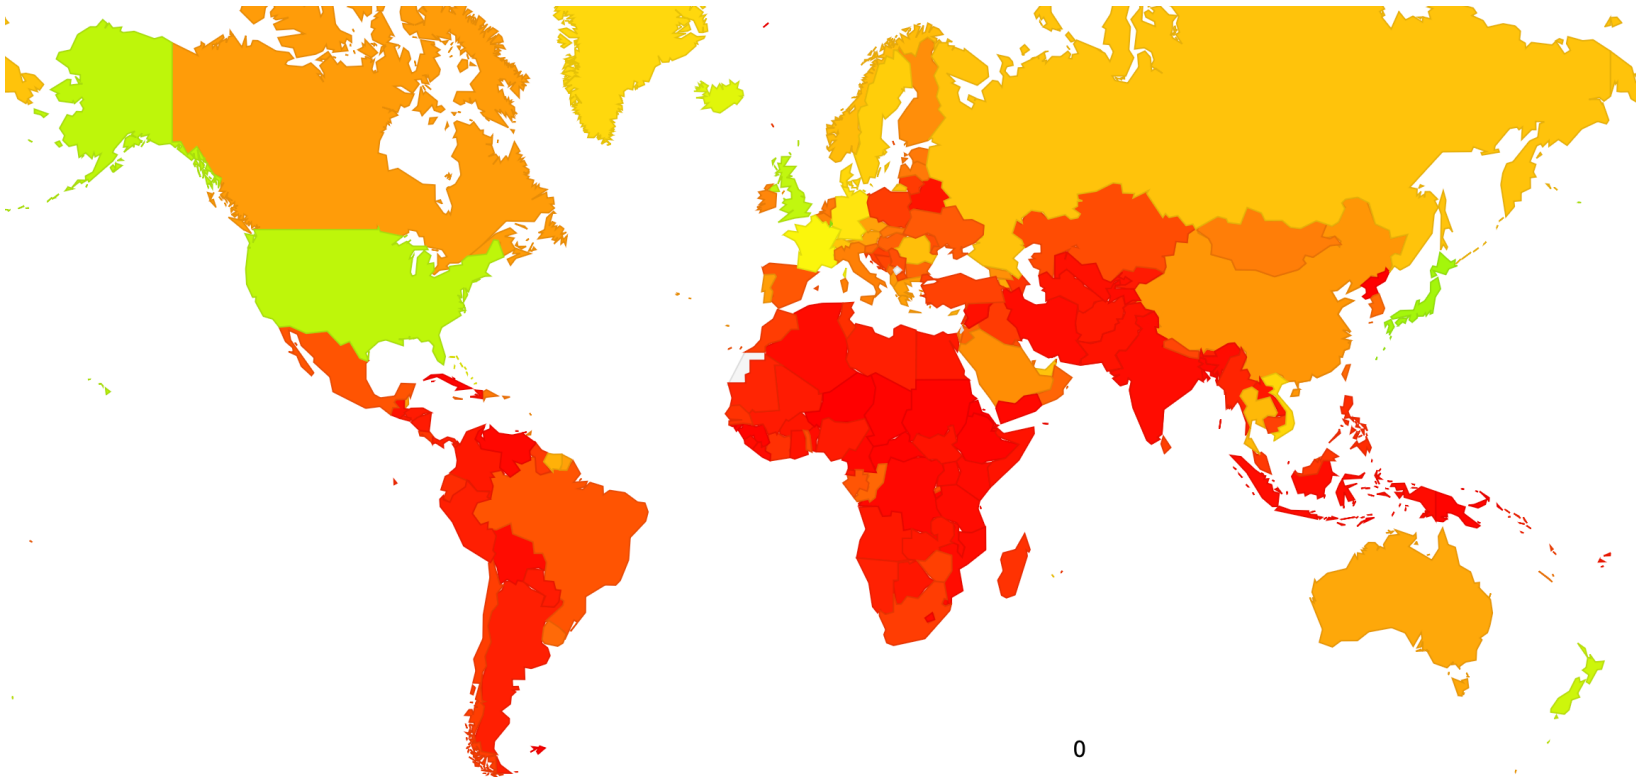
Packet Count (by remote IP addresses)

0. IP Sources	303,545,388
1. IP ECT	6,815,753 (2.45% of sources)
2. IP CE	1,098,965 (16.12% of ECT sources)
3. TCP ECN Opt	7,478,207 (2.46% of sources)
4. TCP ECE (Rec'd)	20,862 (0.27% of TCP sources)
5. TCP CWR (Rec'd)	335,209 (4.48% of TCP sources)

IP ECT by Country



TCP ECN Option by Country



Country Table

CC	Country	IP ECT Rate	TCP ECN Opt Rate	Samples
LI	Liechtenstein, Western Europe, Europe	7.72%	7.72%	661
JE	Jersey, Northern Europe, Europe	7.23%	7.27%	2,490
LU	Luxembourg, Western Europe, Europe	6.41%	6.63%	23,518
JP	Japan, Eastern Asia, Asia	5.30%	6.37%	3,423,773
MO	Macao Special Administrative Region of China, Eastern Asia, Asia	4.62%	5.26%	73,602
BS	Bahamas, Caribbean, Americas	4.53%	4.81%	14,423
FR	France, Western Europe, Europe	4.39%	4.58%	1,949,995
NZ	New Zealand, Australia and New Zealand, Oceania	4.29%	5.53%	95,892
IN	India, Southern Asia, Asia	4.06%	0.25%	17,880,174
GL	Greenland, Northern America, Americas	3.78%	3.96%	2,649
VN	Vietnam, South-Eastern Asia, Asia	3.60%	3.92%	1,777,402
SV	El Salvador, Central America, Americas	3.53%	0.45%	109,548
AE	United Arab Emirates, Western Asia, Asia	3.43%	3.57%	397,903
GB	United Kingdom of Great Britain and Northern Ireland, Northern Europe, Europe	3.41%	5.78%	1,788,972
AD	Andorra, Southern Europe, Europe	3.38%	4.40%	2,157
DE	Germany, Western Europe, Europe	3.28%	4.18%	1,955,293

Thanks!