Packet Loss Signaling for Encrypted Protocols

draft-ferrieuxhamchaoui-tsvwg-lossbits

Alexandre Ferrieux – Orange Labs

Isabelle Hamchaoui – Orange Labs

Igor Lubashev – Akamai

Motivation: Loss Detection/Measurement Matters

Networks can look like a dumb pipes,
<u>only</u> if someone can find leaks and patch them quickly

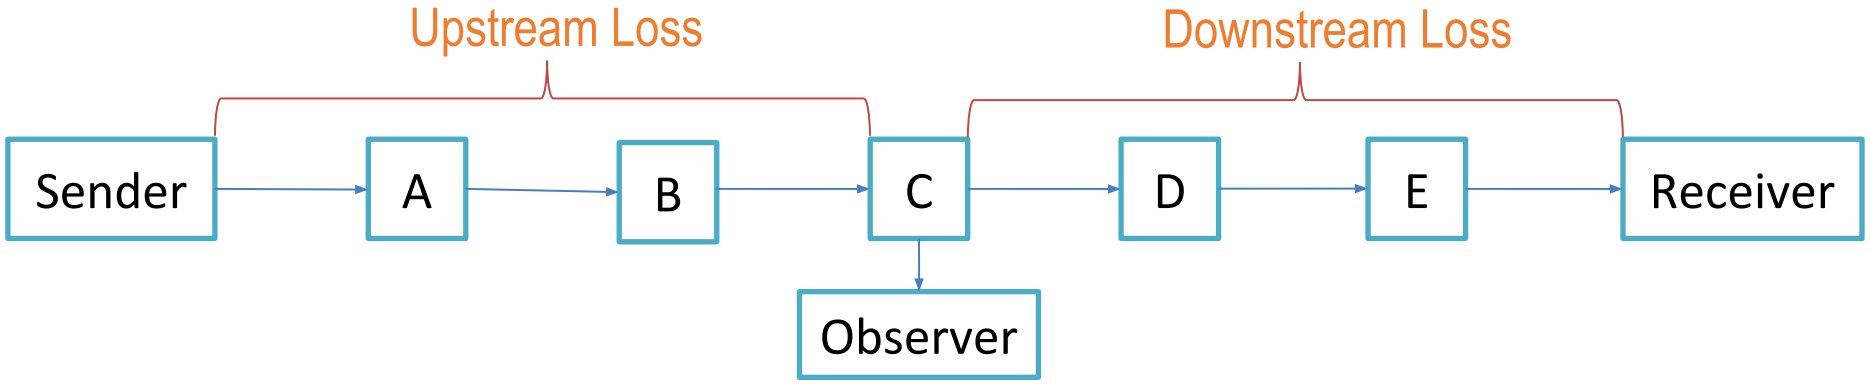

- TCP – observe seq# (and ack#/sack#s, if path is symmetric)

- Encrypted transport headers: ☹
  - QUIC has a “latency Spin bit”, so you may get an RTT estimate but not loss

Proposal: Two “Loss bits”

- **Q**: The “sQuare signal” bit is toggled every N outgoing packets (akin to *color* in RFC 8321)

- **L**: The “Loss event” bit is 1 when “Unreported Loss counter” (ULC) > 0
  - ULC is incremented for each packet deemed lost by the protocol
  - ULC is decremented for each packet sent with L=1

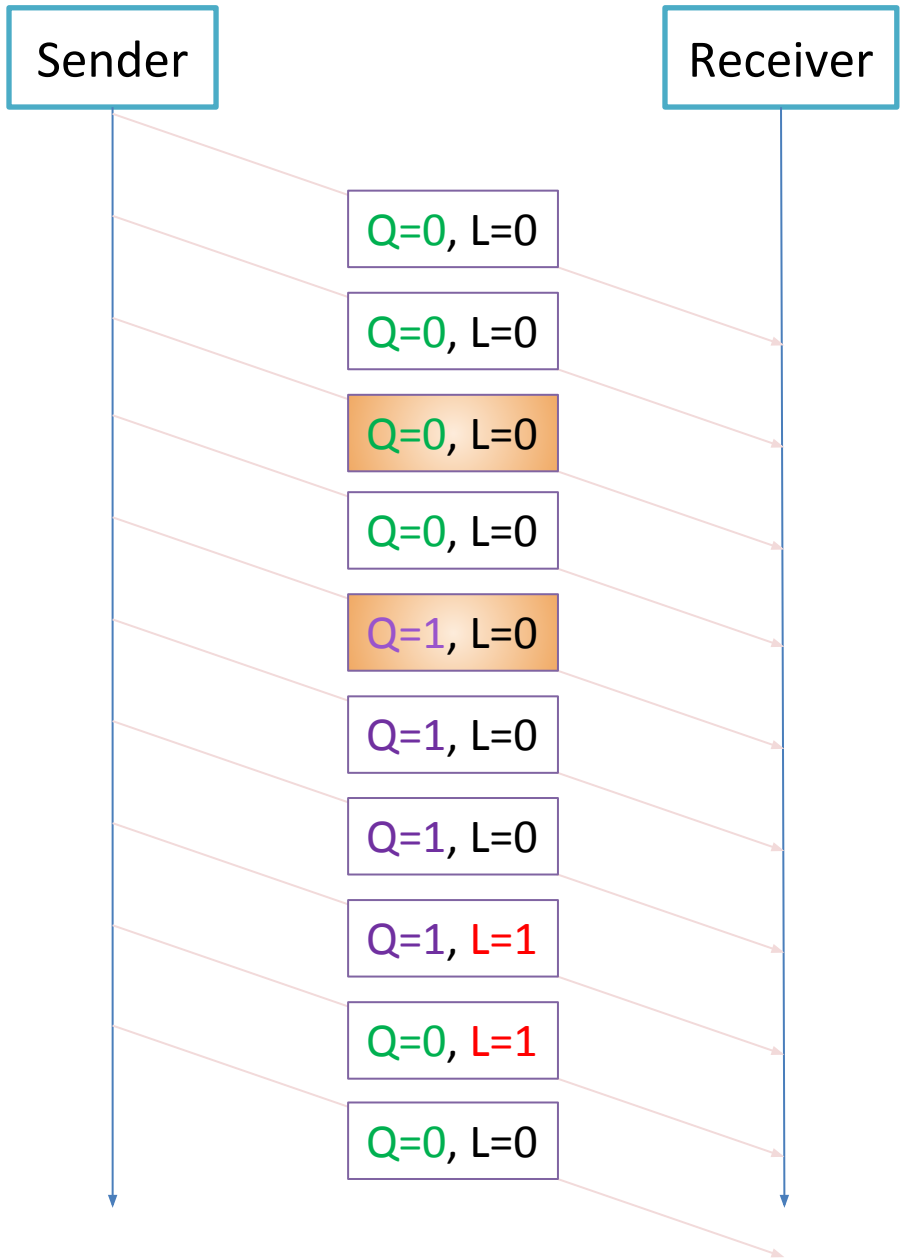

Loss Calculation

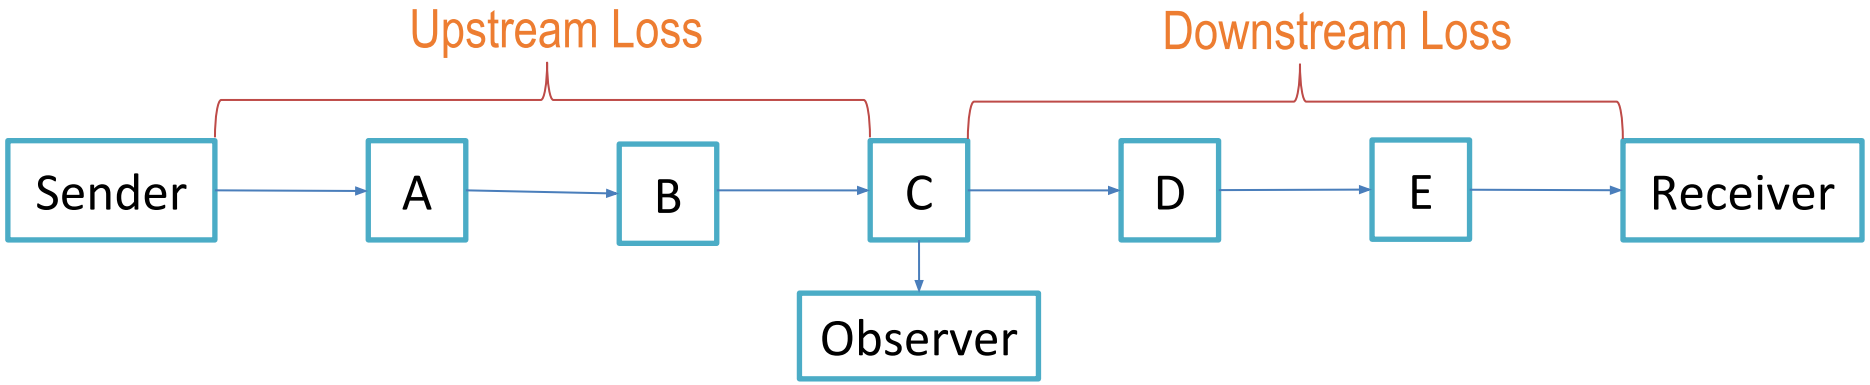

- End-to-End loss (e)

$$e = \text{fraction of packets with L=1}$$

- Upstream loss (u)

$$u = 1 - \frac{\text{average observed packets in a block (same Q)}}{\text{size of the block}}$$

- Downstream loss (d)

$$(1-u)(1-d) = 1-e \qquad d = \frac{e-u}{1-u} \approx \mathbf{e\text{-}u}$$

Experiment Setup

Implemented Q&L on QUIC traffic in Orange-Akamai deployments

- Orange networks (4 countries) with on-path observation points
- Akamai CDN servers with Q&L implementation
  - Q&L in 2 most significant `ip.ttl` and `ipv6.hoplimit` bits
- Unmodified *real clients*: no sim, no FUT :)

Experiment topology (x4)

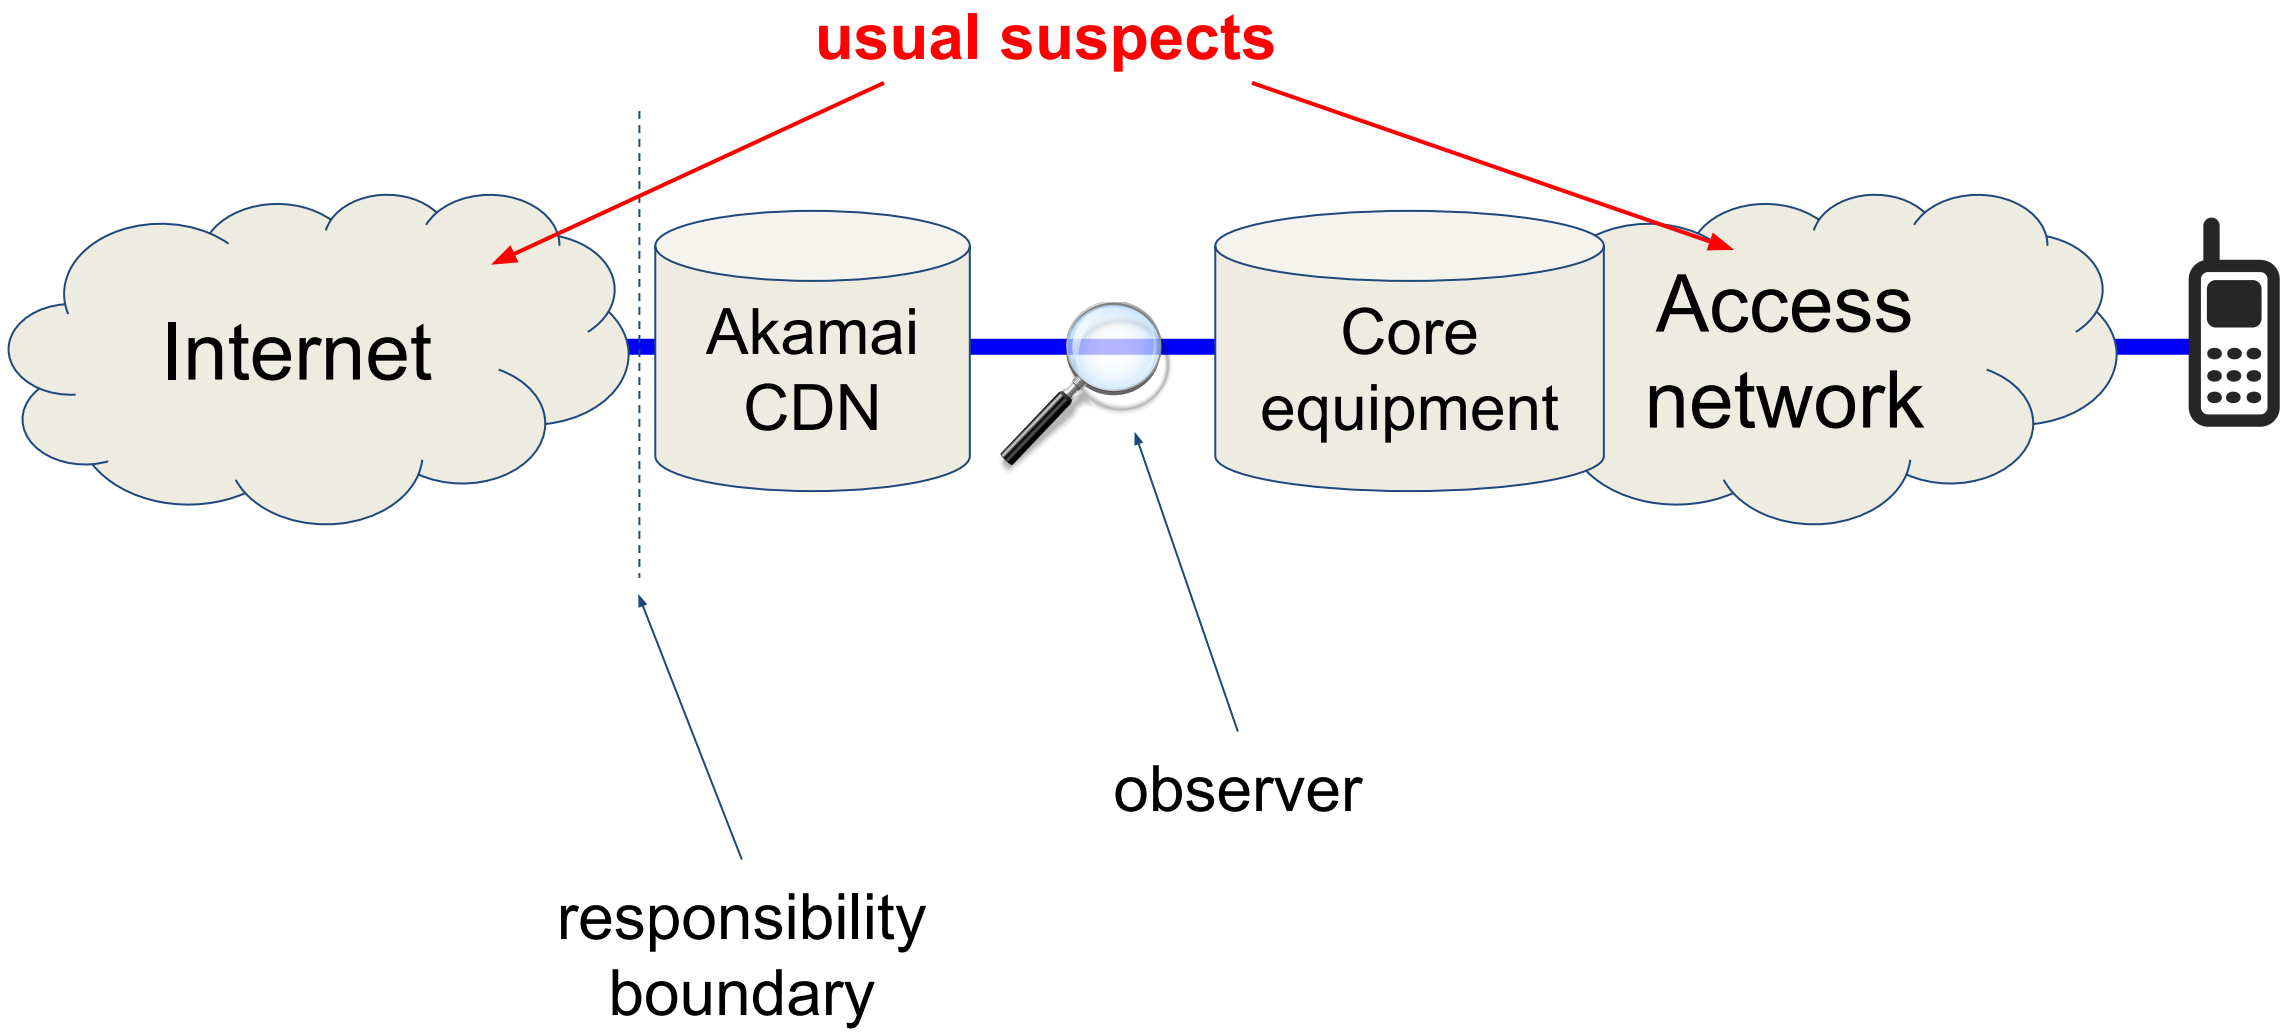

Implementation Details

- QUIC stack patched:
  - compute Q&L
  - insert them in TTL high bits
- Ugly details:
  - need to choose a period for Q: N=64 packets (half-period of Q)
  - constrain initial TTL to free high bits: <=63 (no big deal)

Observer Details

At the observation points:

- traffic capture, downlink only, truncated (full payload is useless)
- packet summaries extracted: timestamp, IP/ports, Q&L, payload size
- flow segmentation: IP/ports + inactivity timeout (60s)

- Capture into one compact text file per flow (~ QUIC connection)
  => many **post-processing** variants can be tried offline
  => both **statistical** and **unit** analysis can be performed

Flow Selection for Analysis

- Over 50k connections analyzed, spanning 2 weeks (4 countries: c1, c2, c3, c4)
- Flow size (# of packets) is critical due to Q granularity
- Select flows >= 5 Q-periods (640 packets) to limit last-chunk loss underestimation

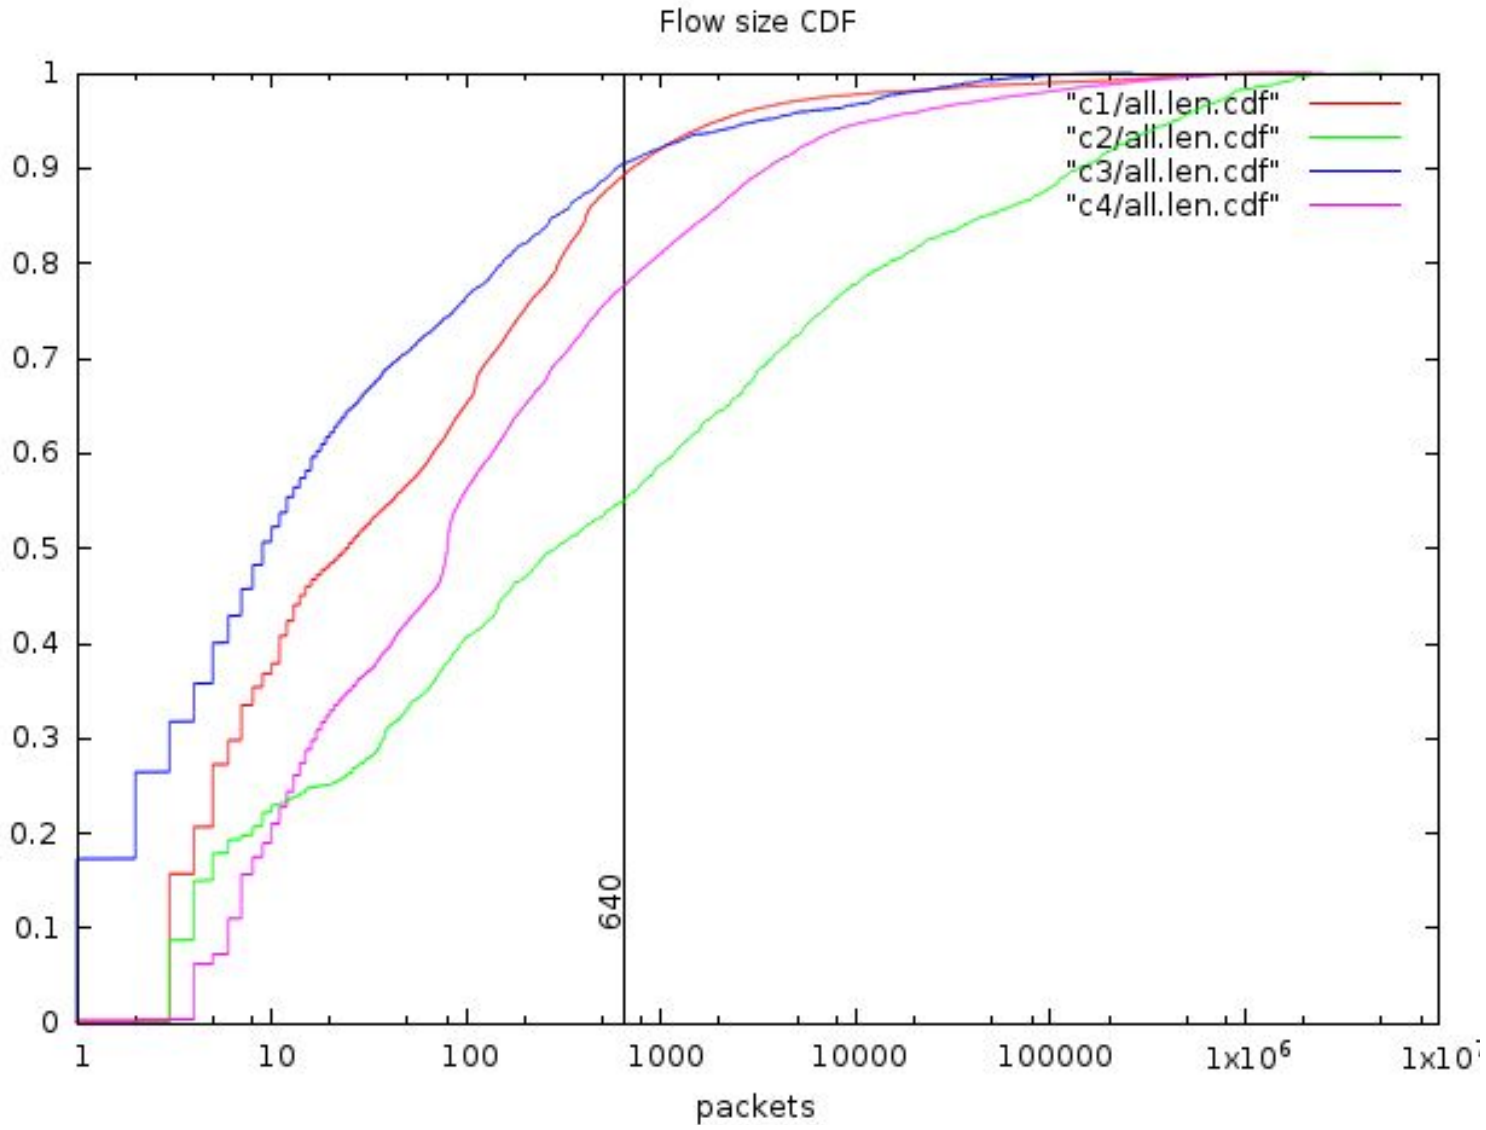

Result Views

- Statistical view = “Q&L scatterplots”

  = distribution of (Qloss,Lloss)
  (Q/Lloss = Q/L-derived loss rate)

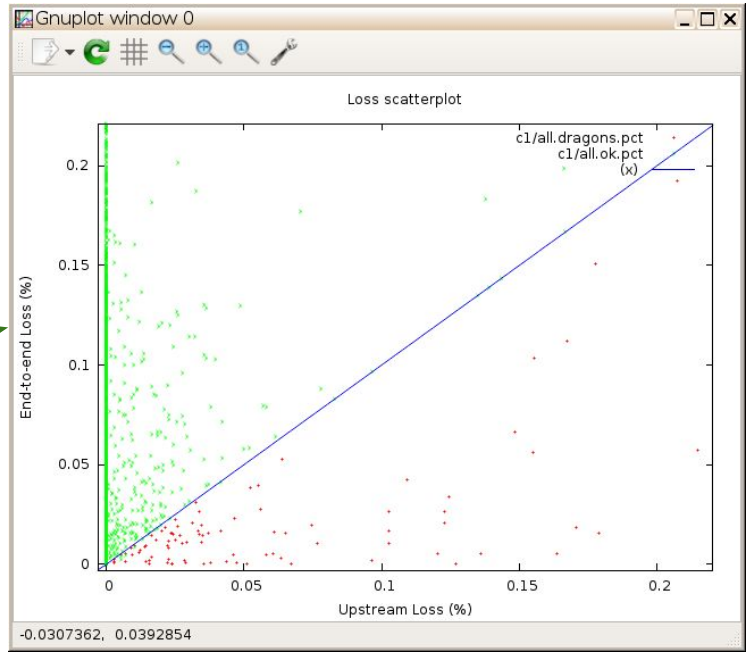

- Unit view = cumulative Q&L

  = sum of individual Q/L losses
  against time of a single connection

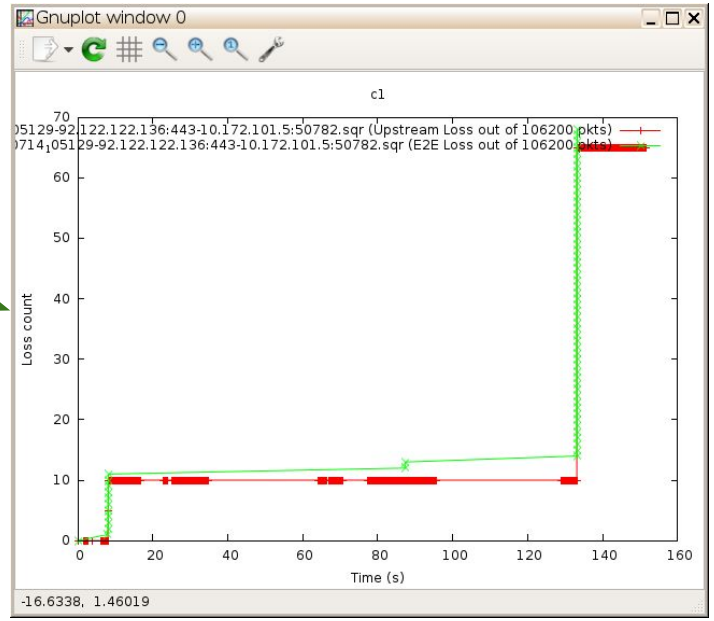

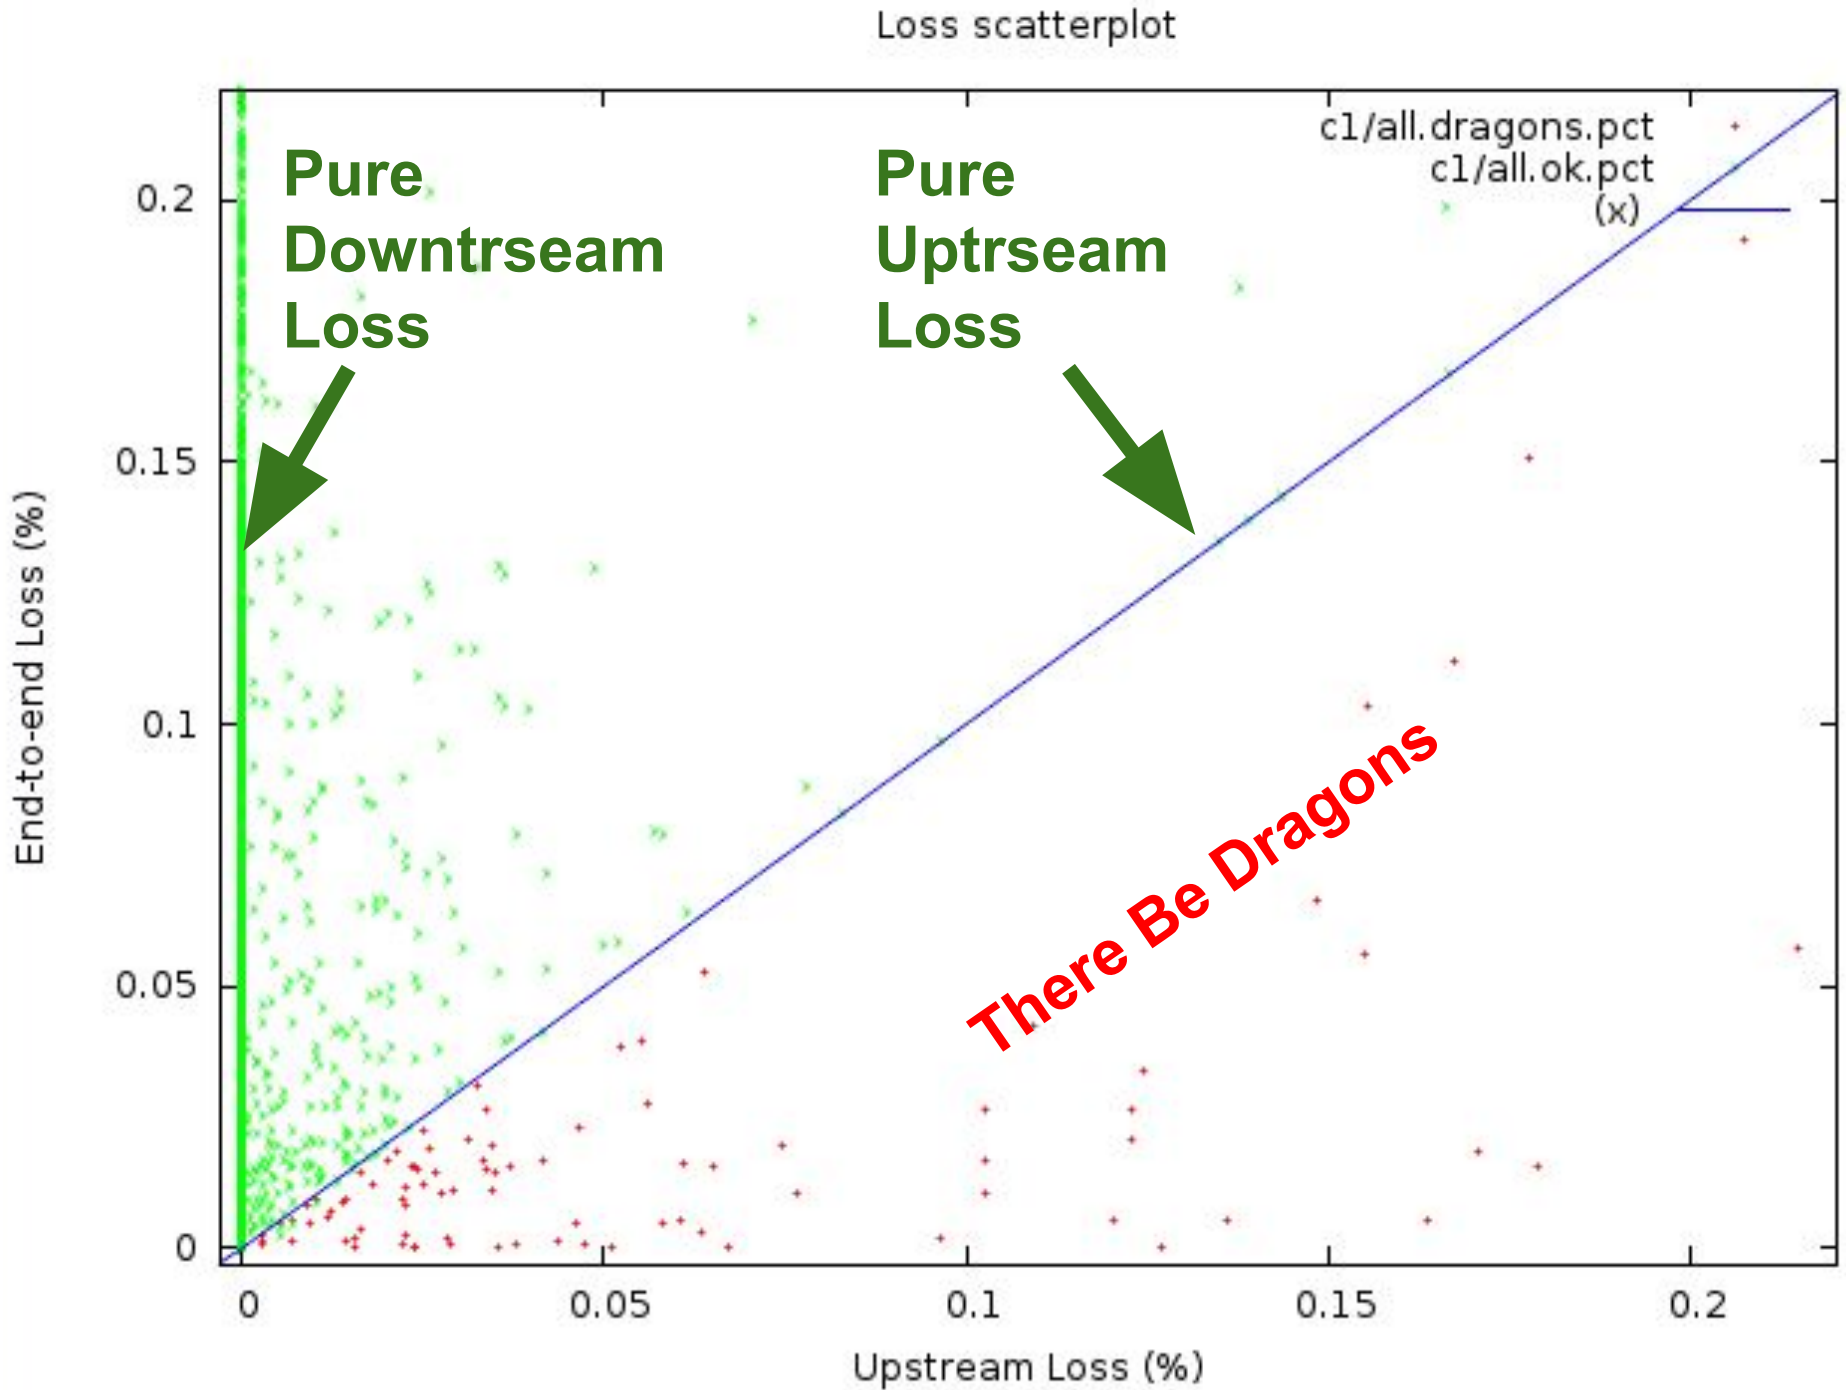
Loss scatterplot
c1/all.dragons.pct
c1/all.ok.pct
(x)
Pure Downtrseam Loss
Pure Uptrseam Loss
There Be Dragons
End-to-end Loss (%)
Upstream Loss (%)
0
0.05
0.1
0.15
0.2

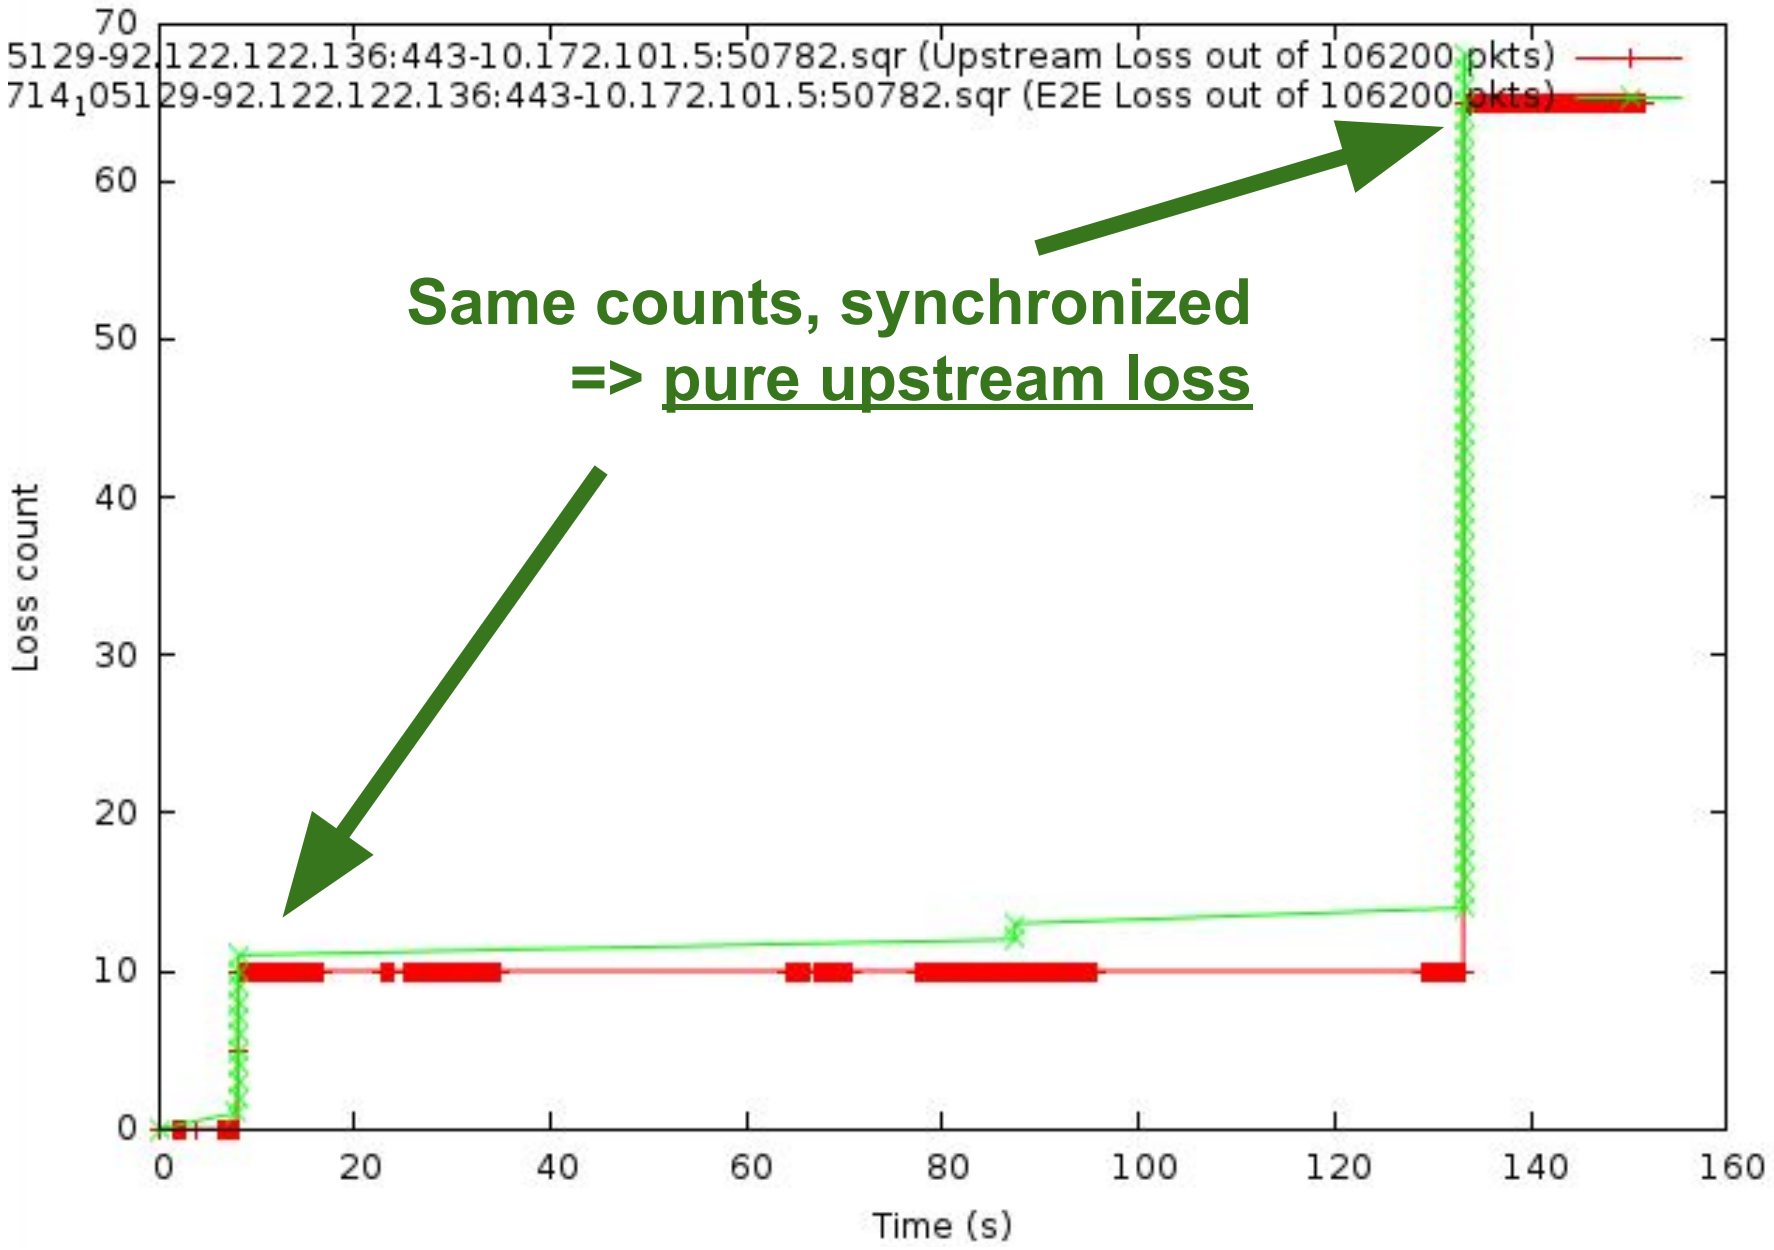
5129-92.122.122.136:443-10.172.101.5:50782.sqr (Upstream Loss out of 106200 pkts)
714 1 05129-92.122.122.136:443-10.172.101.5:50782.sqr (E2E Loss out of 106200 pkts)
Same counts, synchronized
=> pure upstream loss
Loss count
Time (s)
0
10
20
30
40
50
60
70
0
20
40
60
80
100
120
140
160

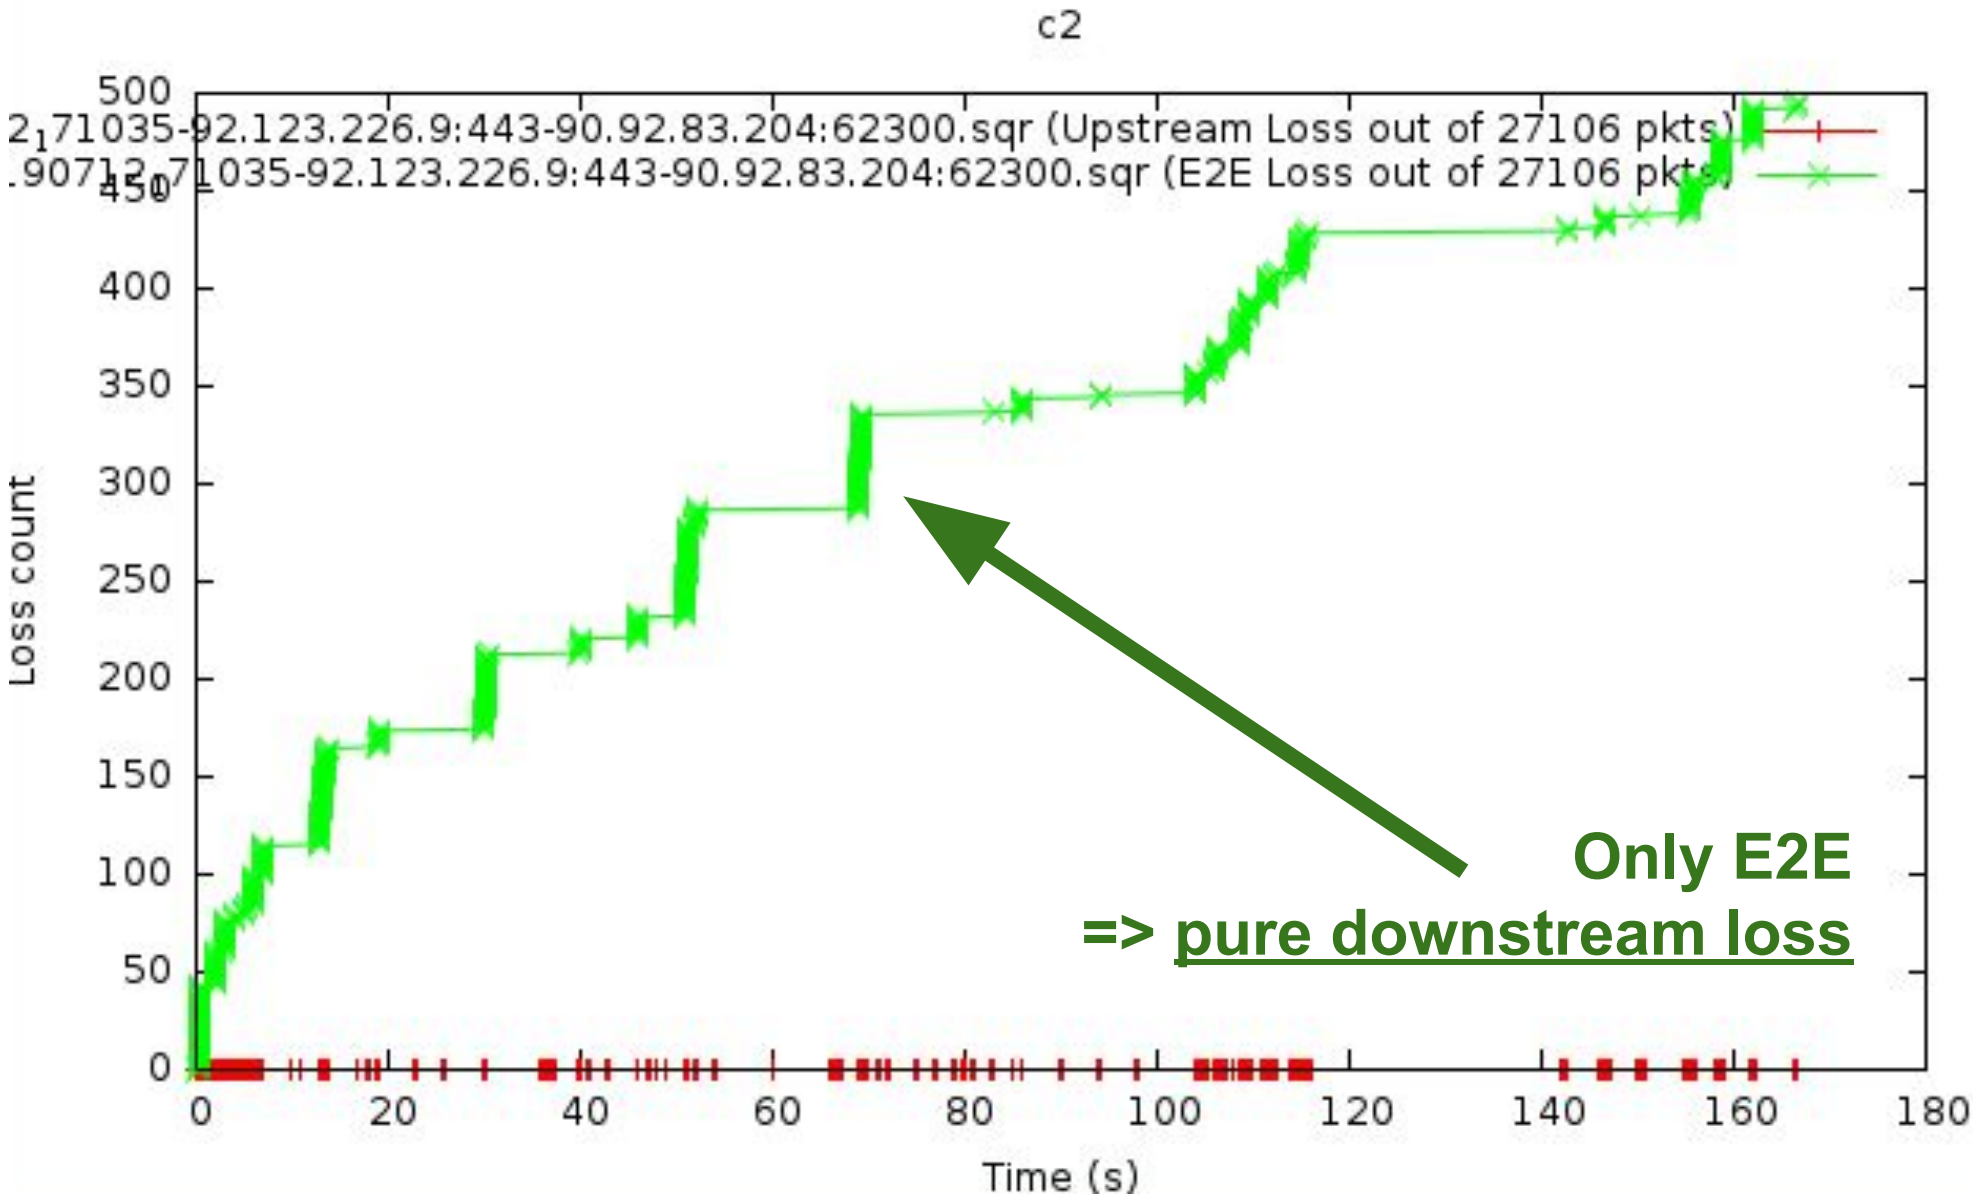
c2
2,71035-92.123.226.9:443-90.92.83.204:62300.sqr (Upstream Loss out of 27106 pkts)
.90712,71035-92.123.226.9:443-90.92.83.204:62300.sqr (E2E Loss out of 27106 pkts)
Loss count
Time (s)
0
50
100
150
200
250
300
350
400
450
500
0
20
40
60
80
100
120
140
160
180
Only E2E
=> pure downstream loss

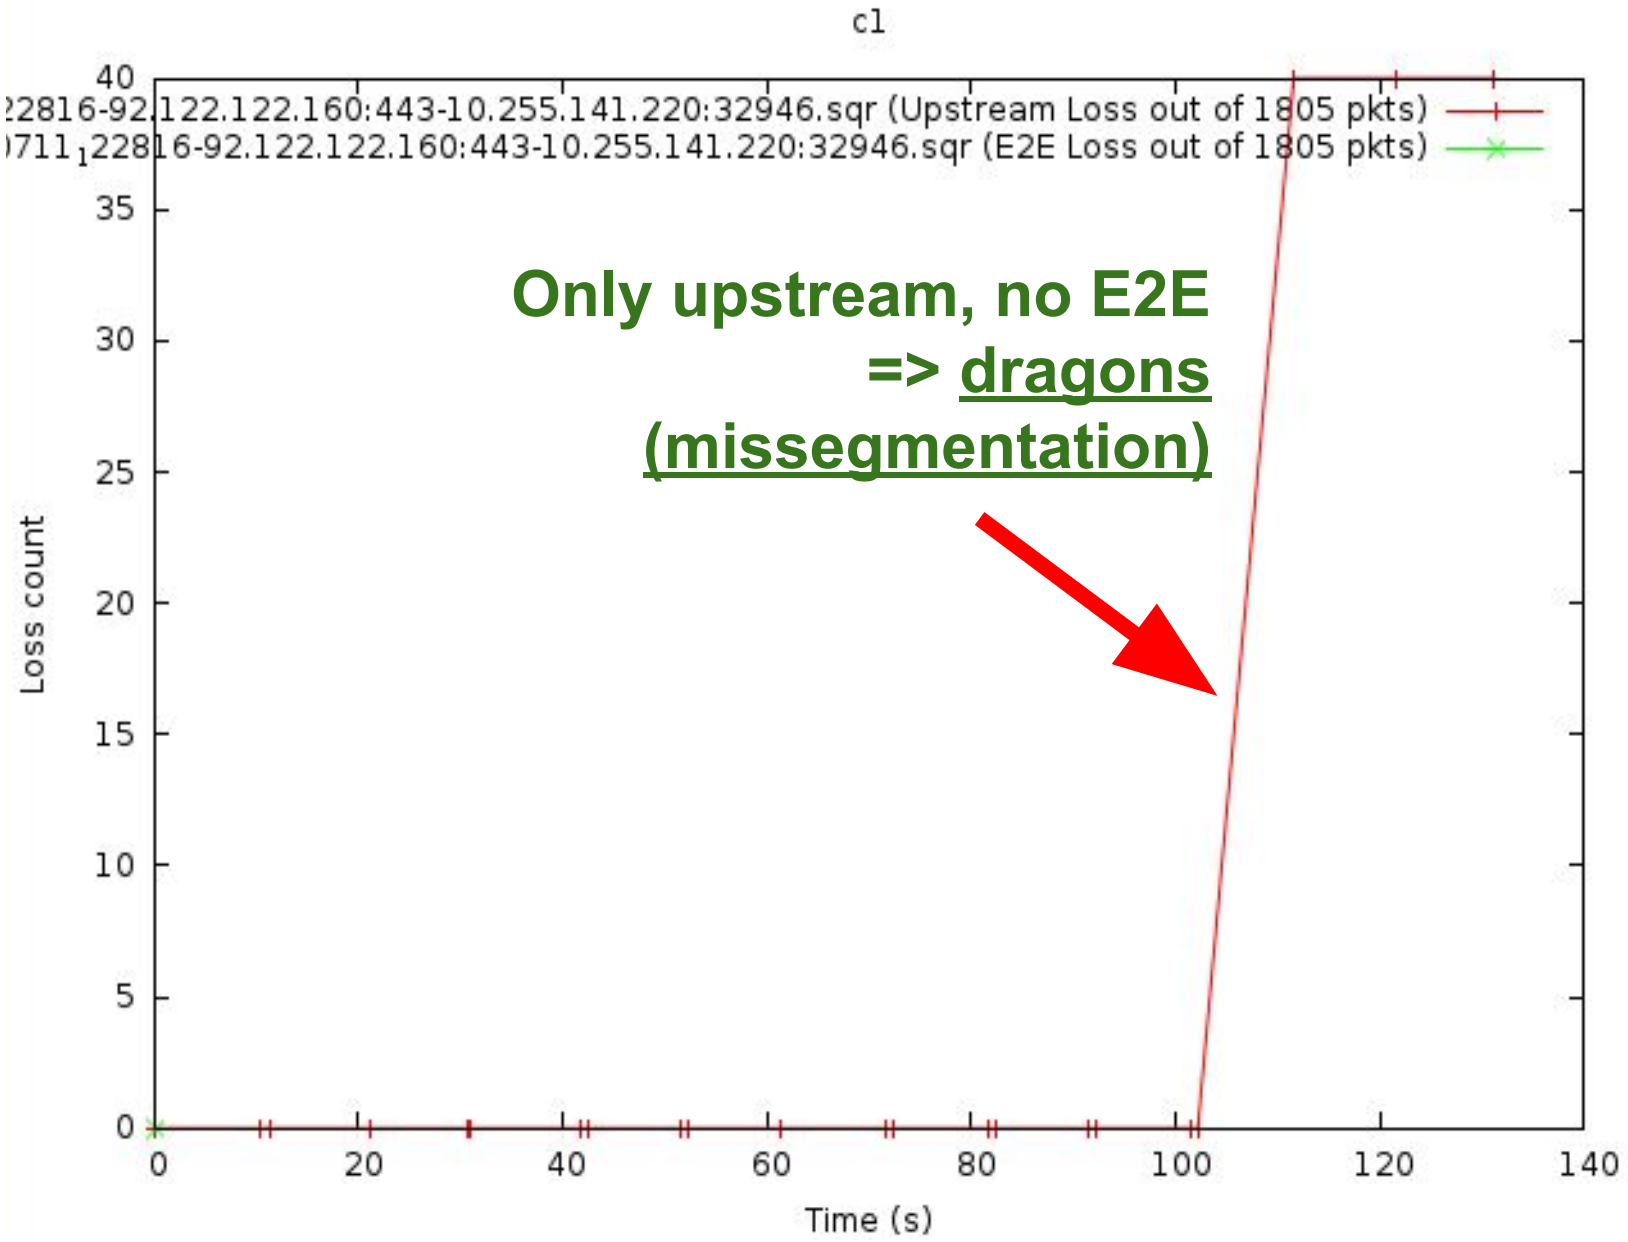
c1
22816-92.122.122.160:443-10.255.141.220:32946.sqr (Upstream Loss out of 1805 pkts)
0711_122816-92.122.122.160:443-10.255.141.220:32946.sqr (E2E Loss out of 1805 pkts)
Only upstream, no E2E
=> dragons
(missegmentation)
Loss count
Time (s)
0
5
10
15
20
25
30
35
40
0
20
40
60
80
100
120
140

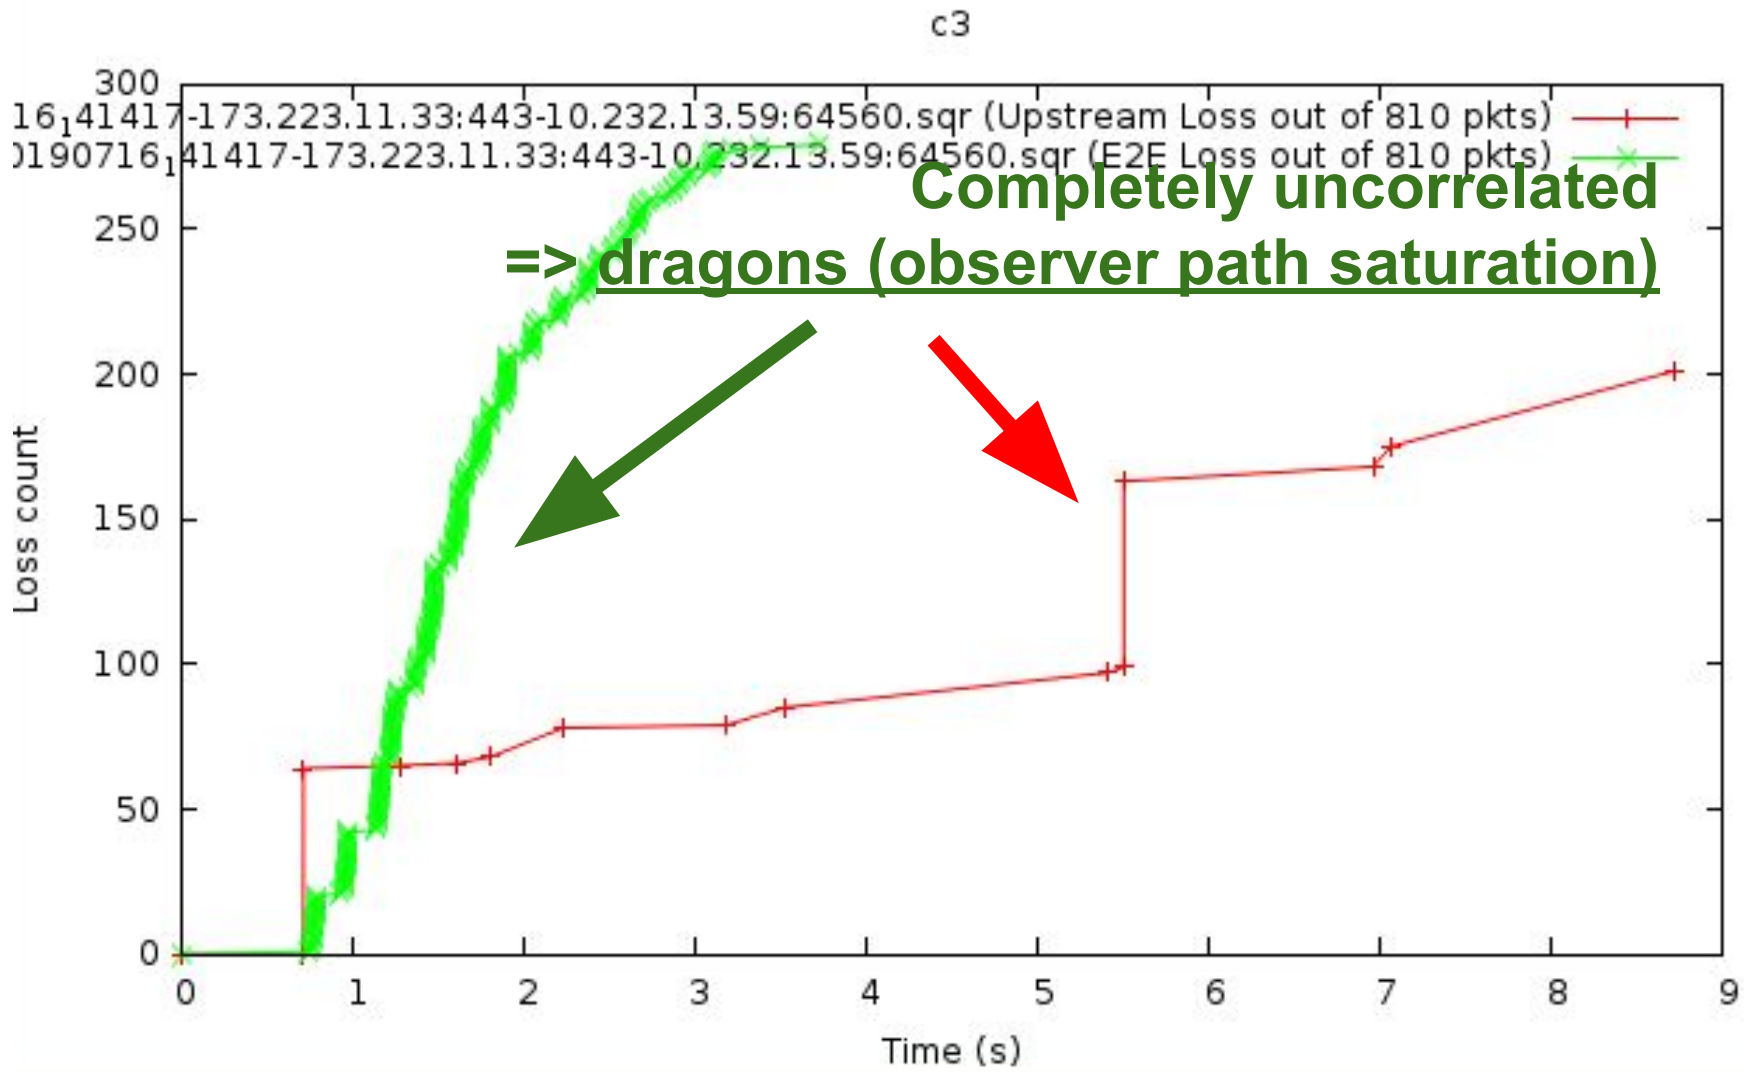
c3
16_141417-173.223.11.33:443-10.232.13.59:64560.sqr (Upstream Loss out of 810 pkts)
0190716_141417-173.223.11.33:443-10.232.13.59:64560.sqr (E2E Loss out of 810 pkts)
Completely uncorrelated
=> dragons (observer path saturation)
Loss count
Time (s)
0
50
100
150
200
250
300
0
1
2
3
4
5
6
7
8
9

Per-country loss scatterplots

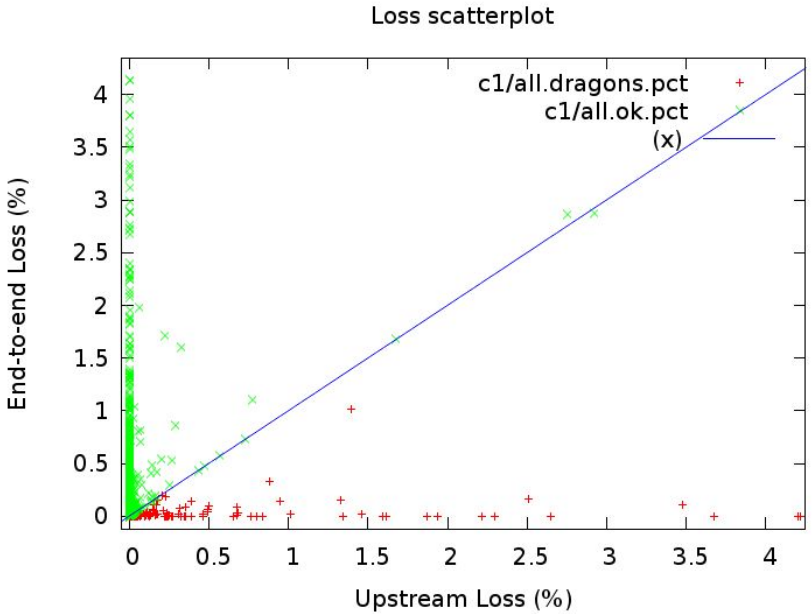

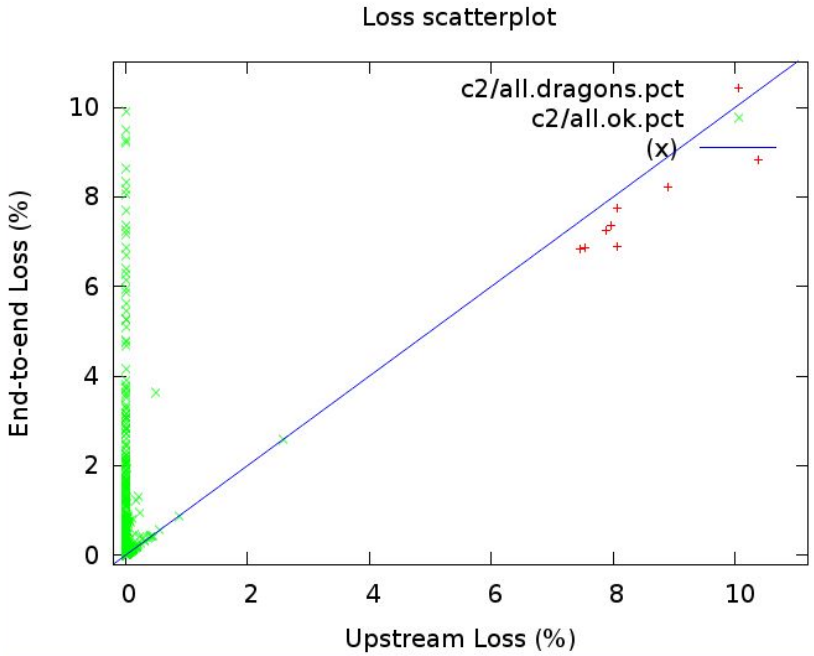

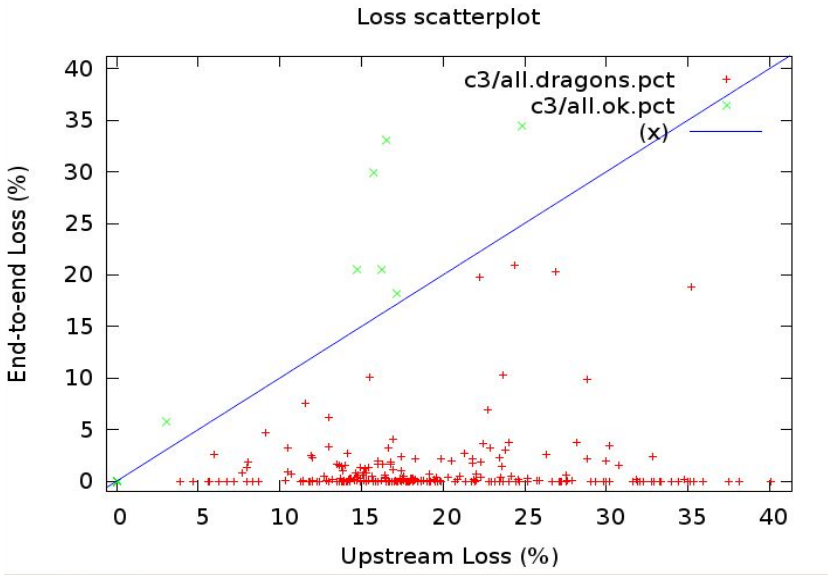

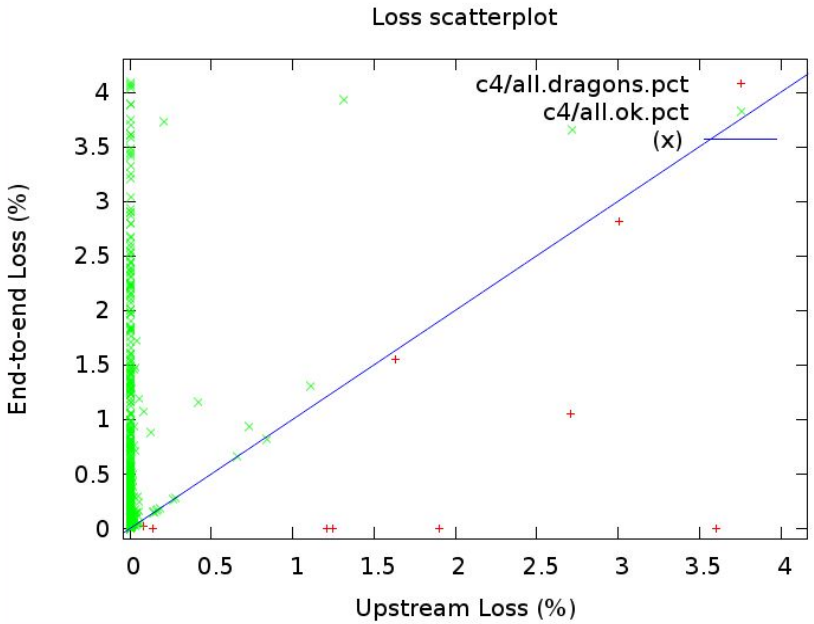

Results: Country 1

- mainly pure-downstream loss
- some noise from reordering

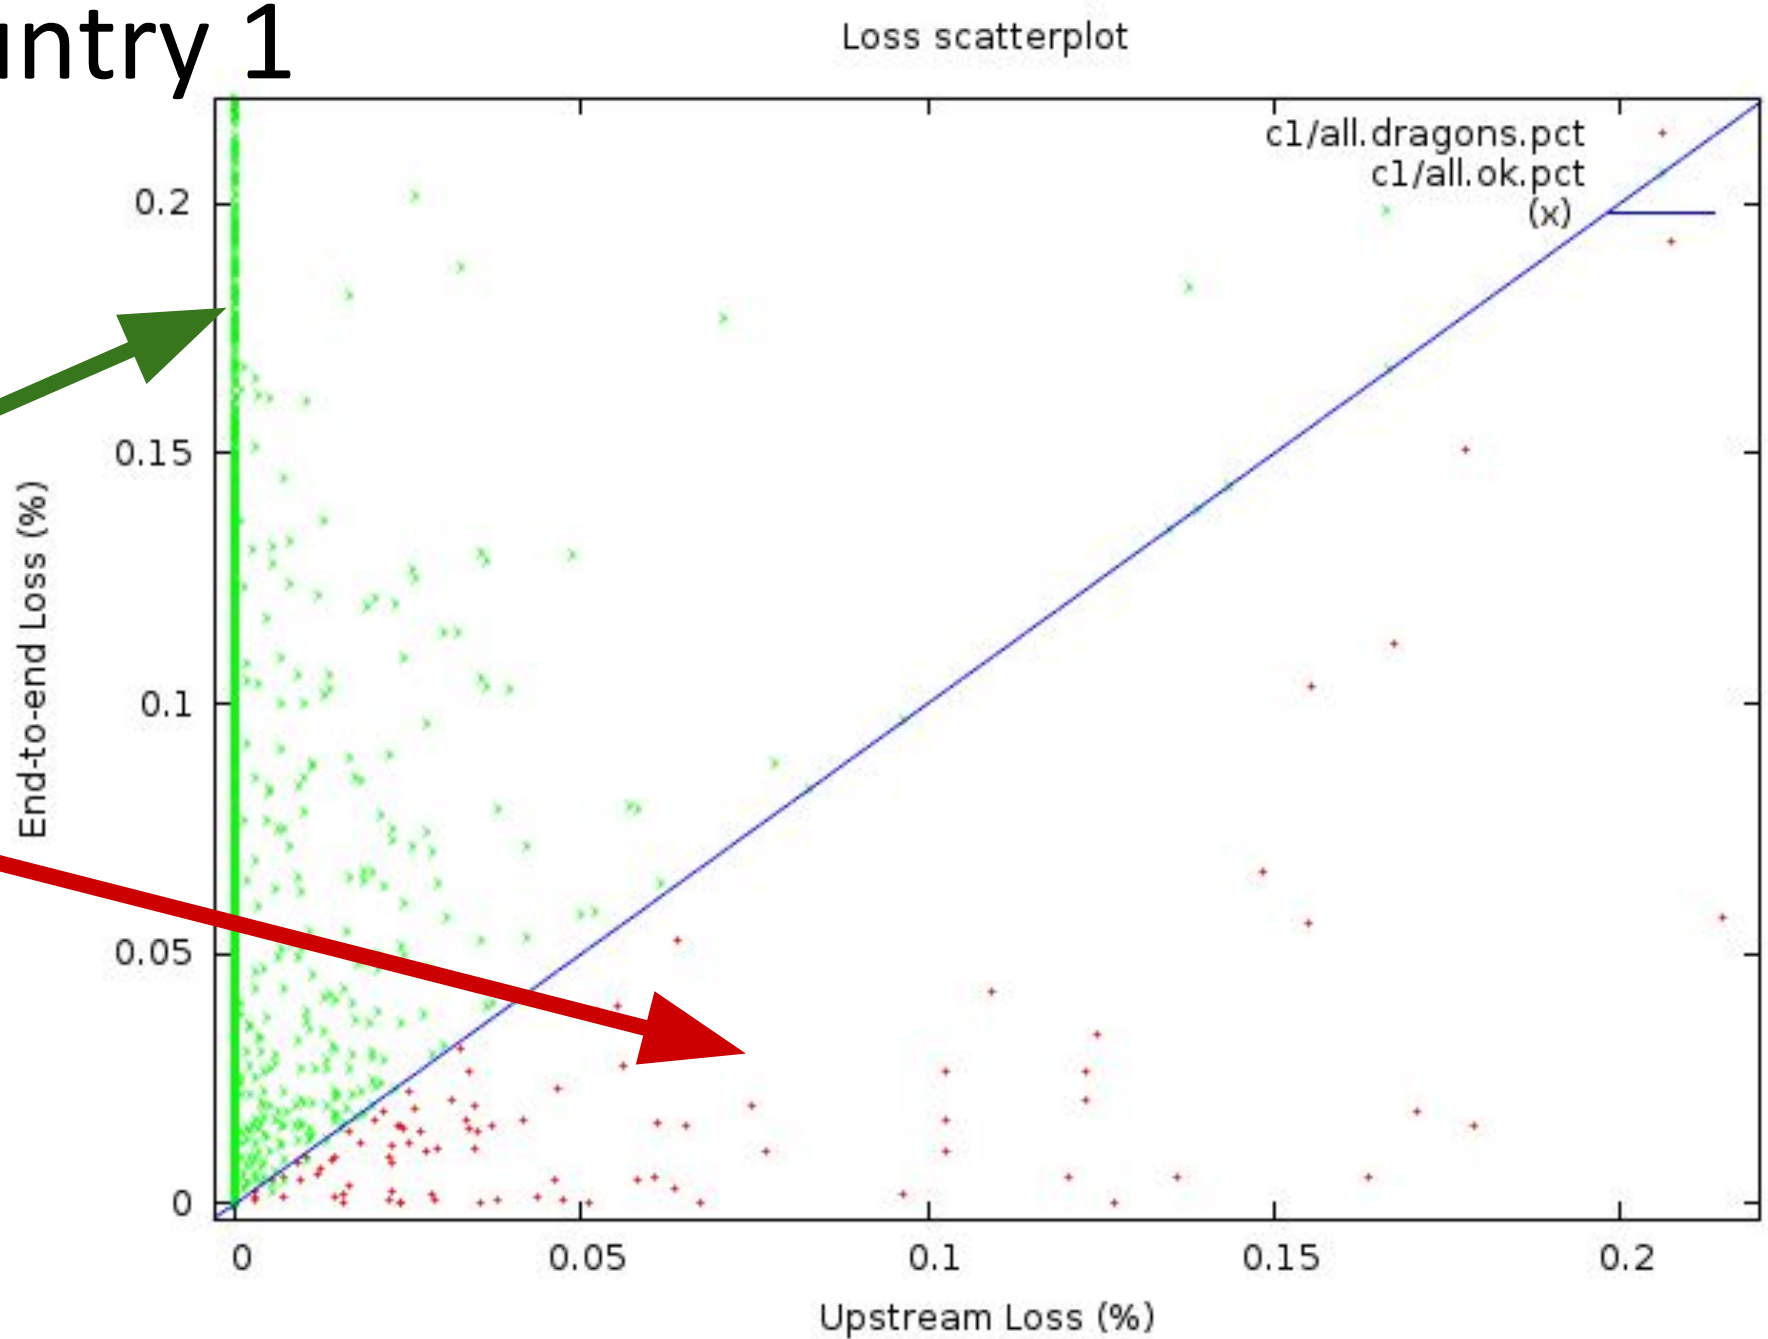

Results: Country 2

- a few strong upstream loss events
- no noise from reordering

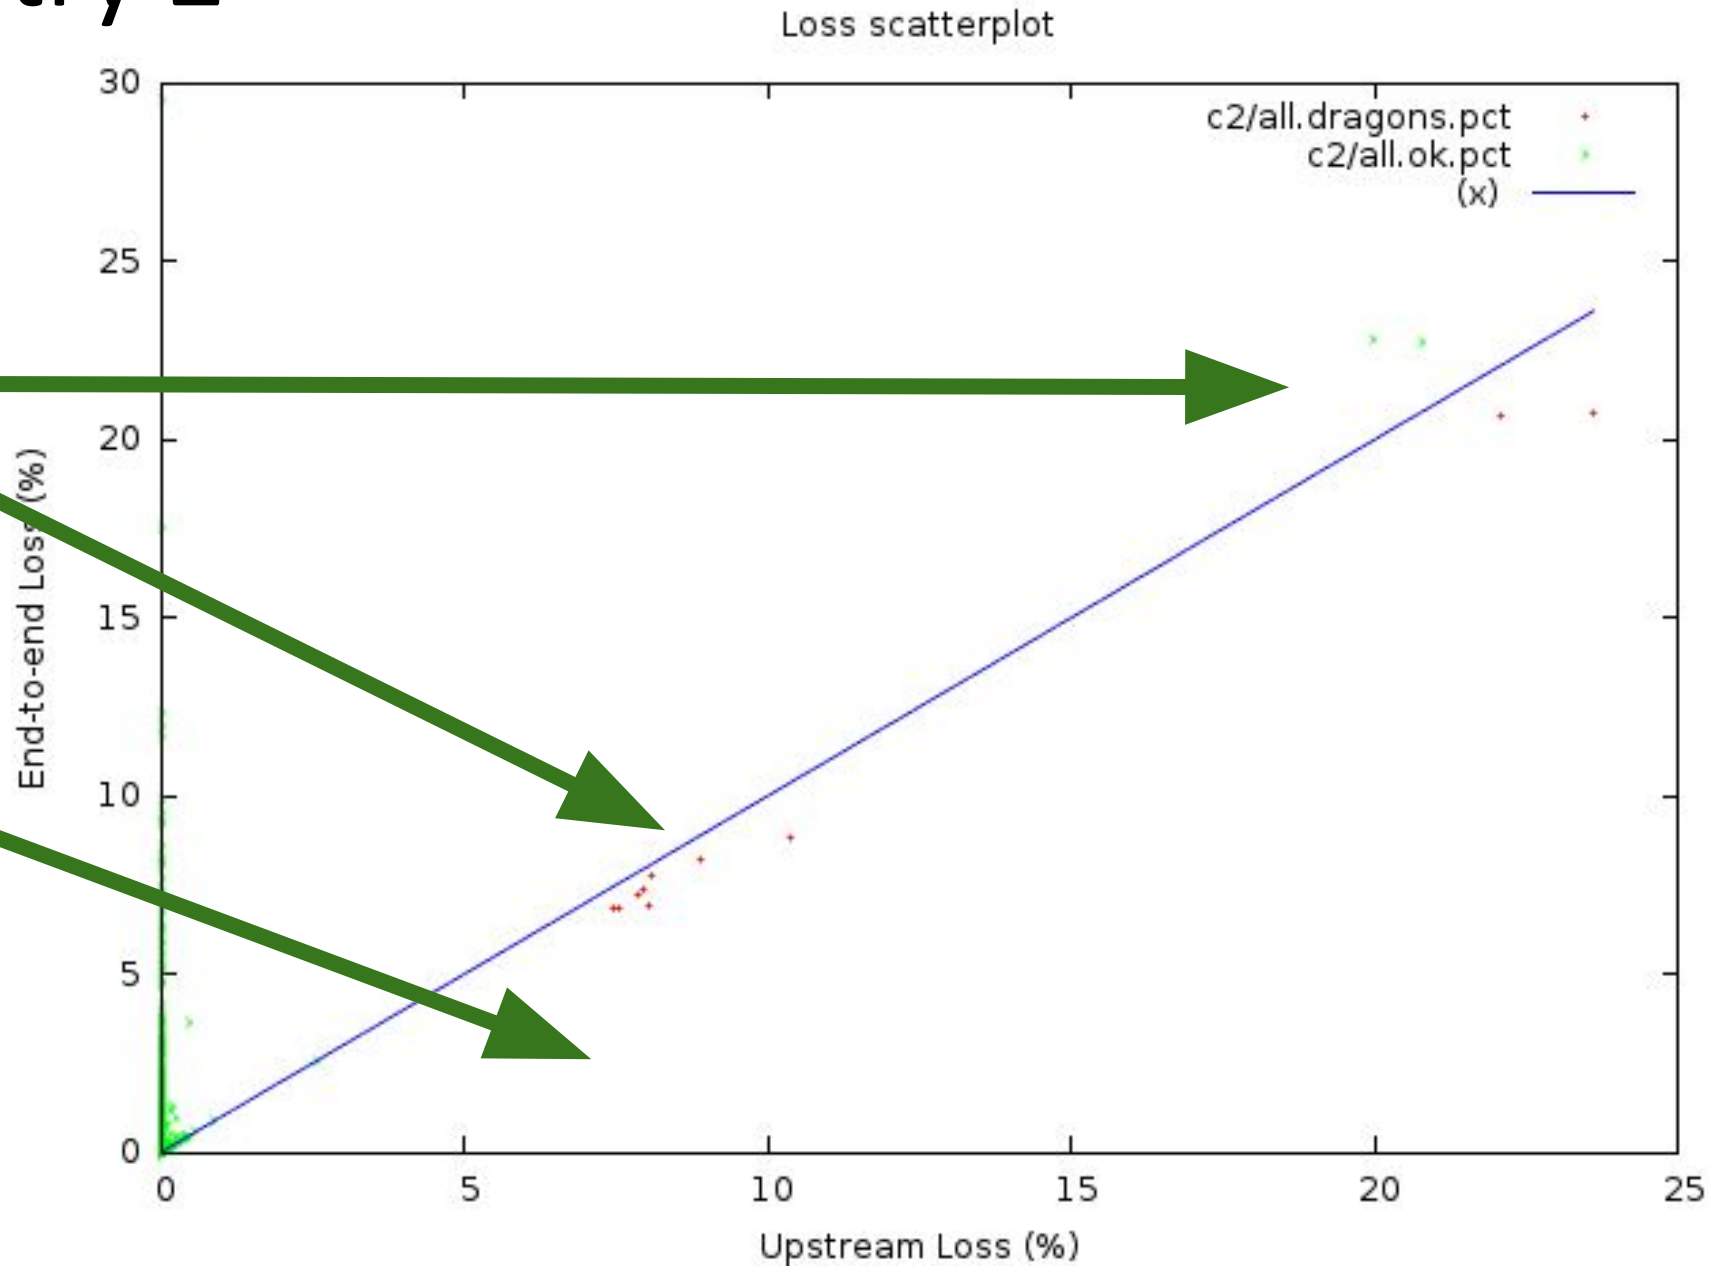

Results: Country 3

- strong noise from saturated capture chain

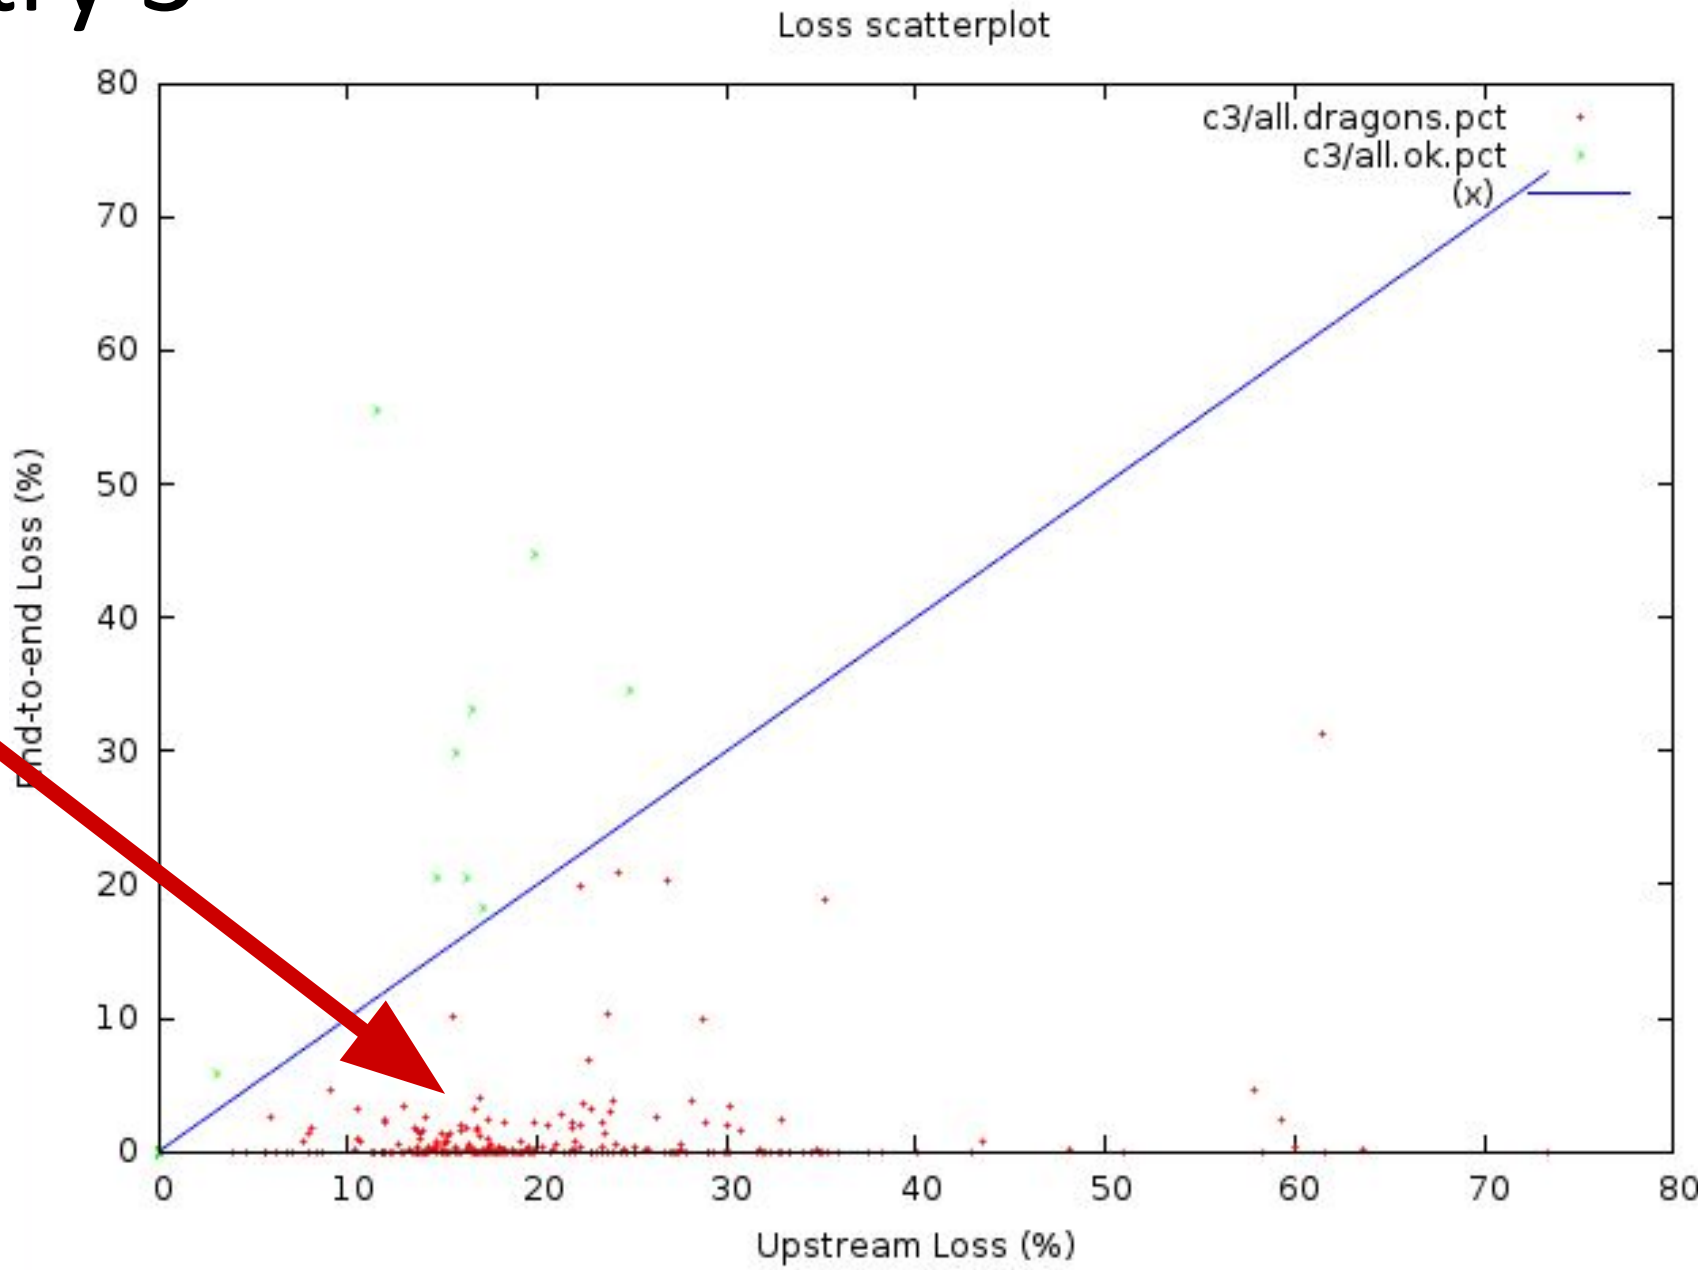

Results: Country 4

- mostly downstream loss
- marginal noise from missegmentation

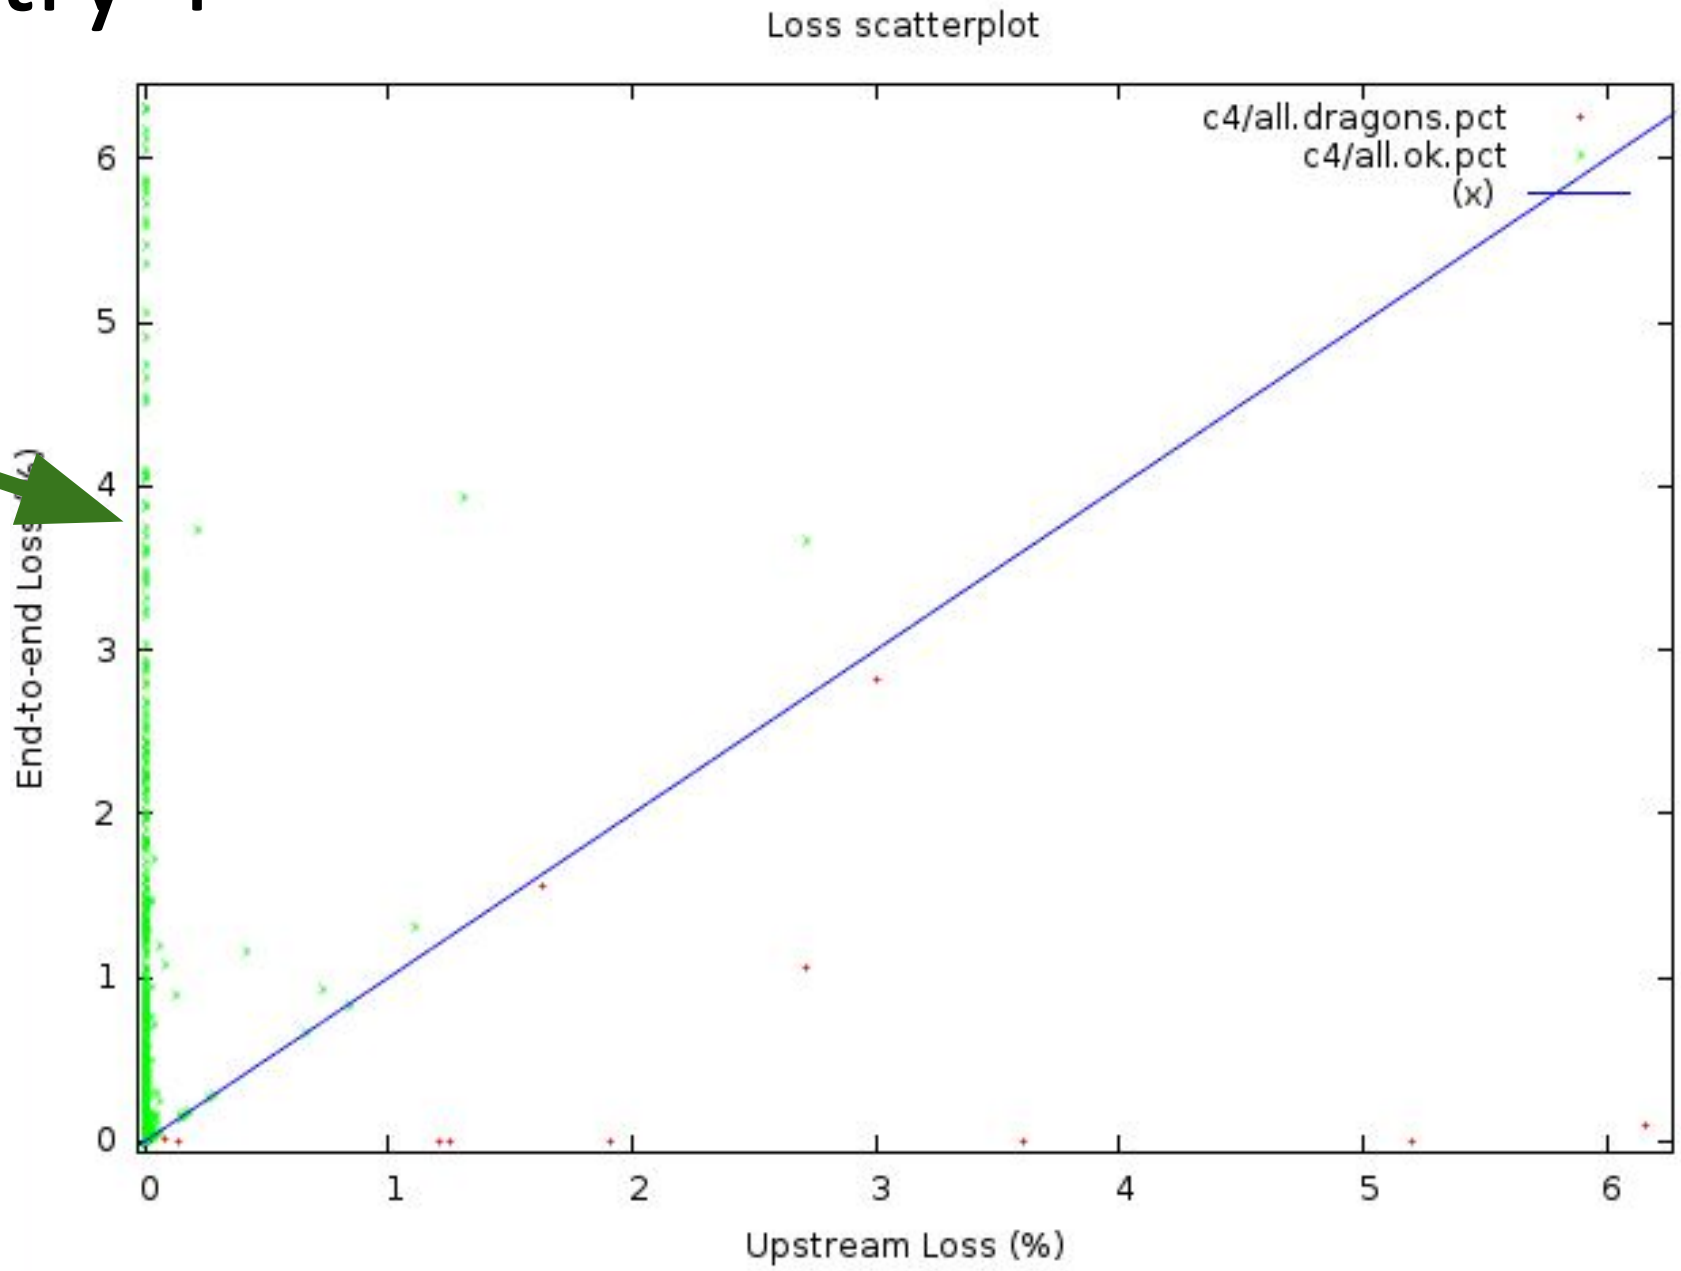

Results: CDF of downstream loss ratio

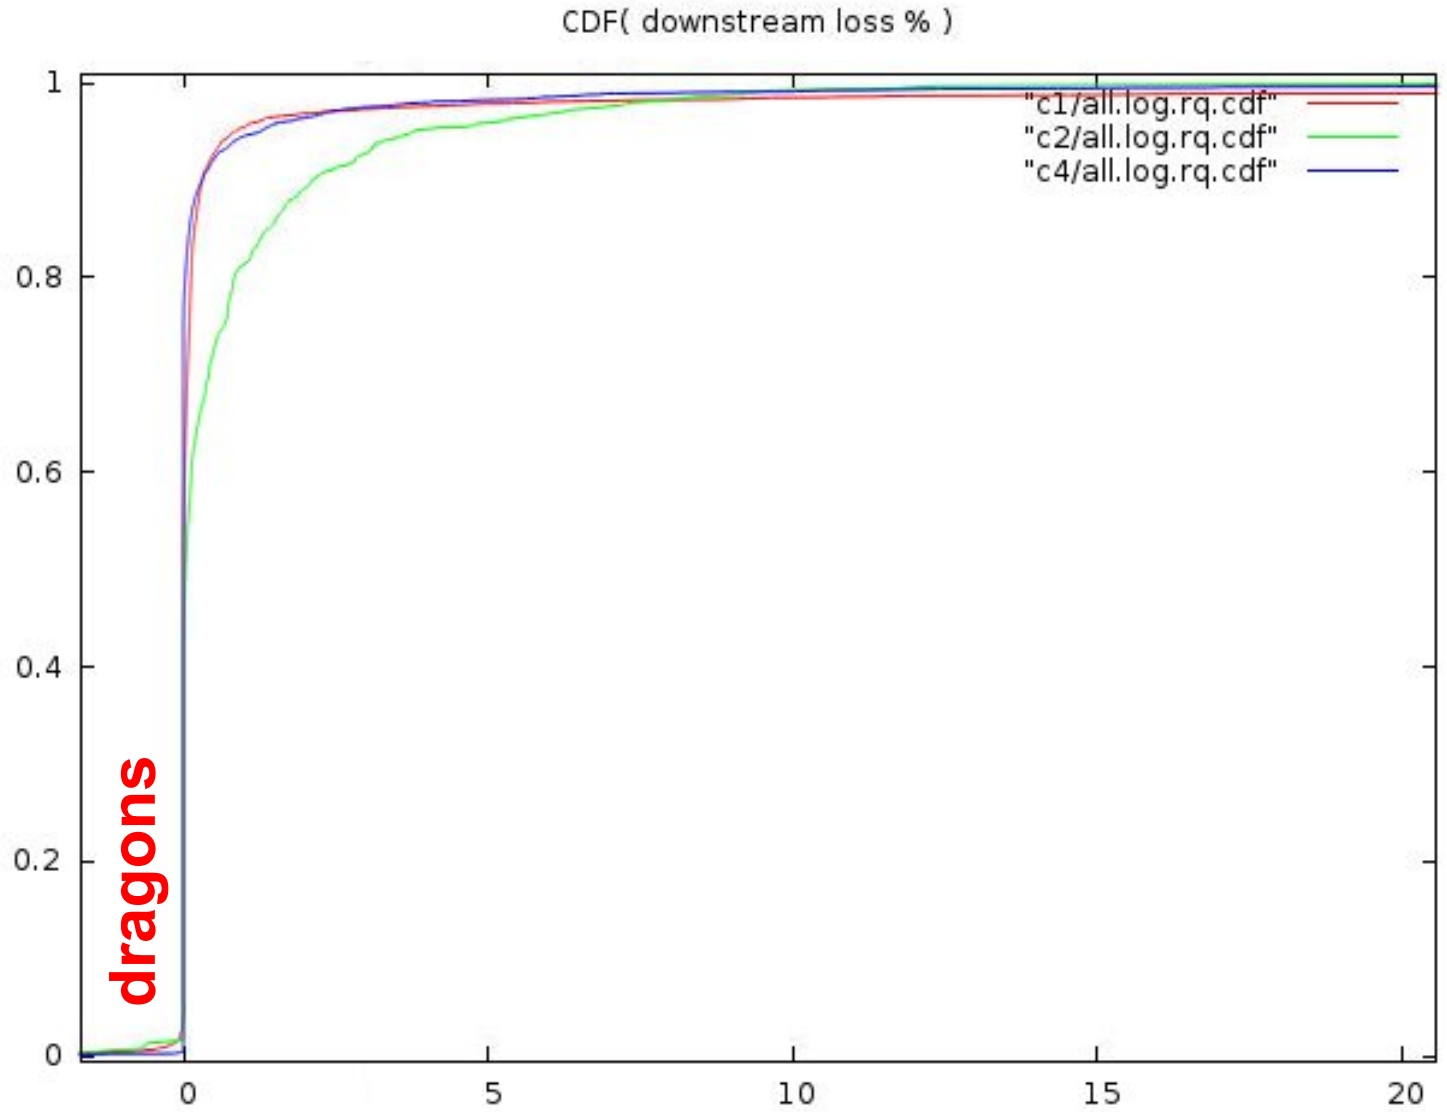

“Dragon Hunt”

- some noise is due to imperfections of denoising heuristics with medium reordering
- some is due to mis-segmentation (port reuse)
- some is due to observer loss (on the capture path only ; should not happen)

All these noise sources can be disambiguated by unit analysis
=> heuristics can be improved

Wrap Up

- Unilateral deployment (server-side) is key to **large-scale** experimentation: Q&L are nice in this respect.
- Q&L signals restore **TCP-like ability** to locate and quantify packet loss
- Short of direct (e.g.) QUIC support, another vehicle is needed.
  TTL>>6 is just an example. (Discuss other options)
- The mechanism can be applied to **any protocol** with **sender-side evaluation** of loss.

we need this in order to keep maintaining networks !

Alternate method: in-band flow summaries

- Multiplex special “flow summary” packets into the same tuple (=> new protocol)
- Meant to travel end-to-end, contrary to IOAM (=> needs to be ignored by receiver stack. Easy with QUIC. Doable with TCP). <=> special case of PBT-M, non-marking mode, with collector = on-path observers.
- Contents (strawman):
  - some magic number
  - sequence number (++ per summary)
  - Q counter (= egress packet counter on that flow)
  - L counter (= sum of L bits)
- Frequency:
  - once every N egress packets + once at flow end
  - may take number of L into decision (report only in the presence of end-to-end loss)