

# Experiments with ECN

Stuart Cheshire & Dave Täht

TSVAREA WG, 89<sup>th</sup> IETF, London, England, March 2014

# Disclaimer

- Opinions of a relative newcomer to ECN (Stuart)
- Generously assisted by Dave Täht

# Apparent State of ECN

- Implemented in Mac OS X, Windows, Linux, etc.
- Turned off — present but not actually used
- Turning on has no effect
  - Routers on Internet not doing ECN marking
- Reducing lost packets is nice
  - But who cares?
  - Loss rates are modest and TCP handles it fine

# TCP for Streaming Video

- TCP handles packet loss extremely well
- But... causes irregular data delivery to client
- No problem for file transfer (e.g. sending an email)
- Big problem for streaming video over TCP
  - Netflix
  - YouTube
  - etc.

# Test Setup

- Linux web server
  - 3.8.0-19-generic #30-Ubuntu (TCP CUBIC)
- CeroWRT gateway
  - 3.10.18-1
- Mac OS X client using “curl”
  - 10.8.5

# Test: 10Mb/s downstream

(< 1 ms intrinsic delay, plus queueing delay)

256kB FIFO queue

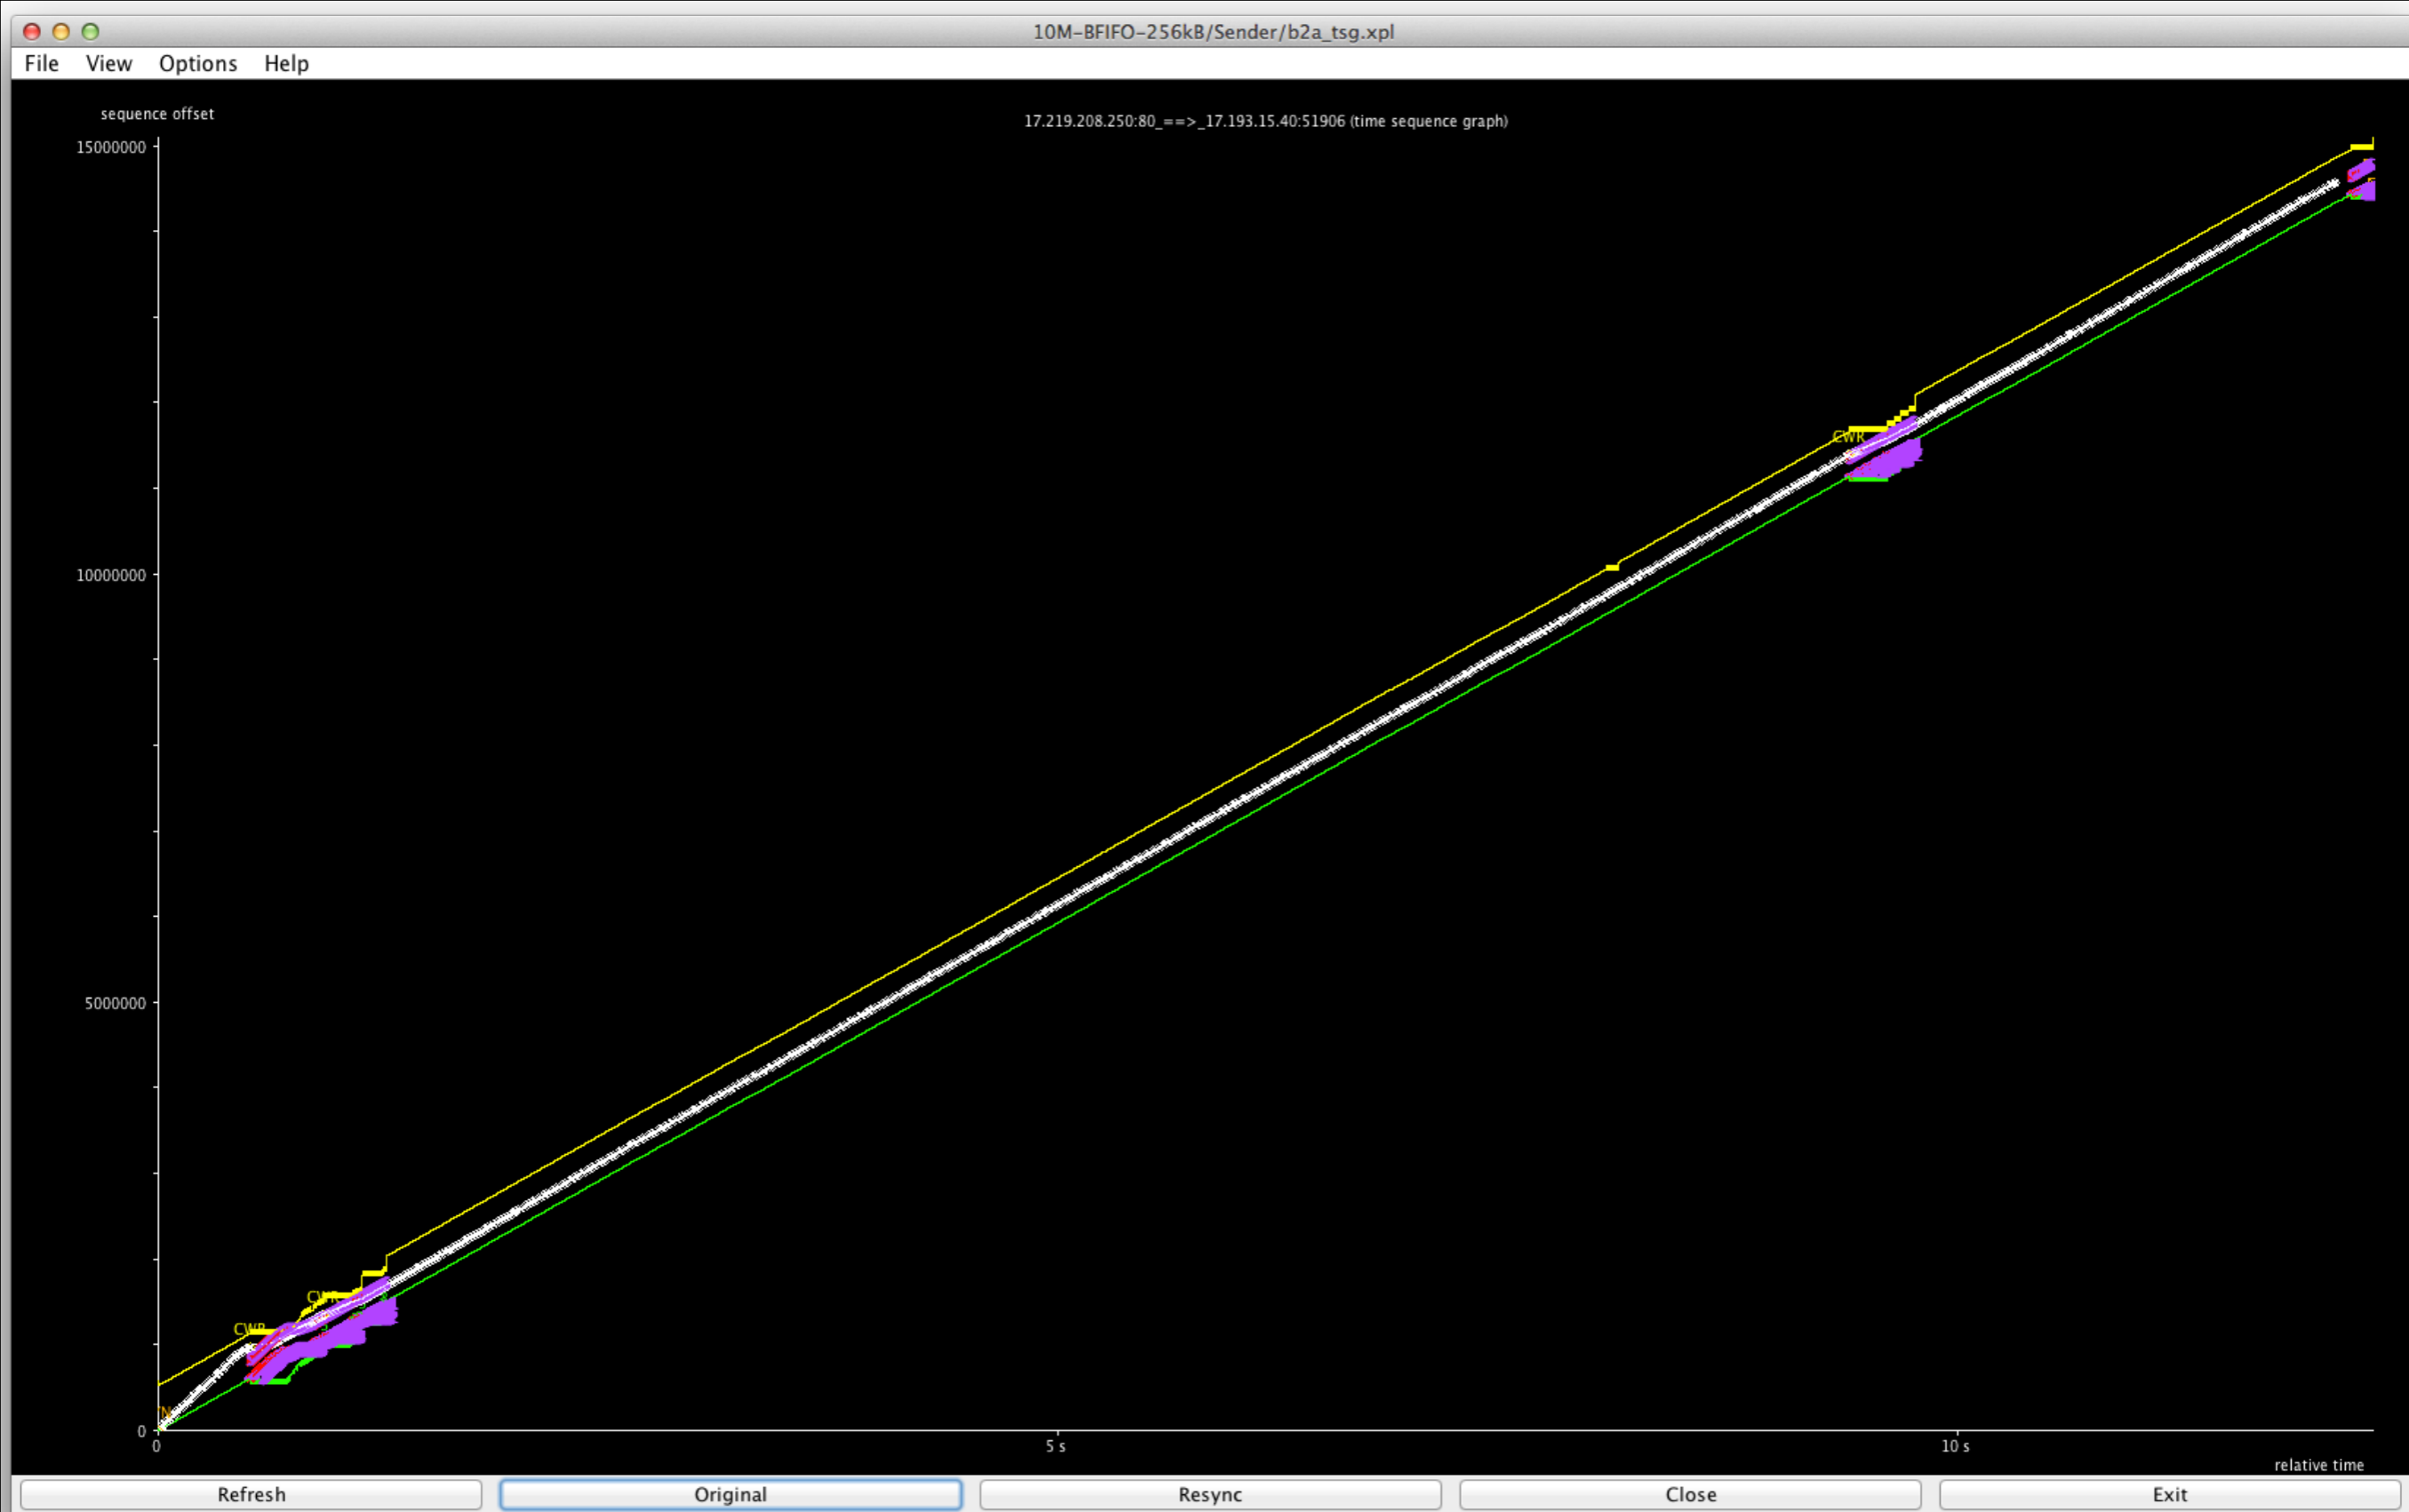
vs.

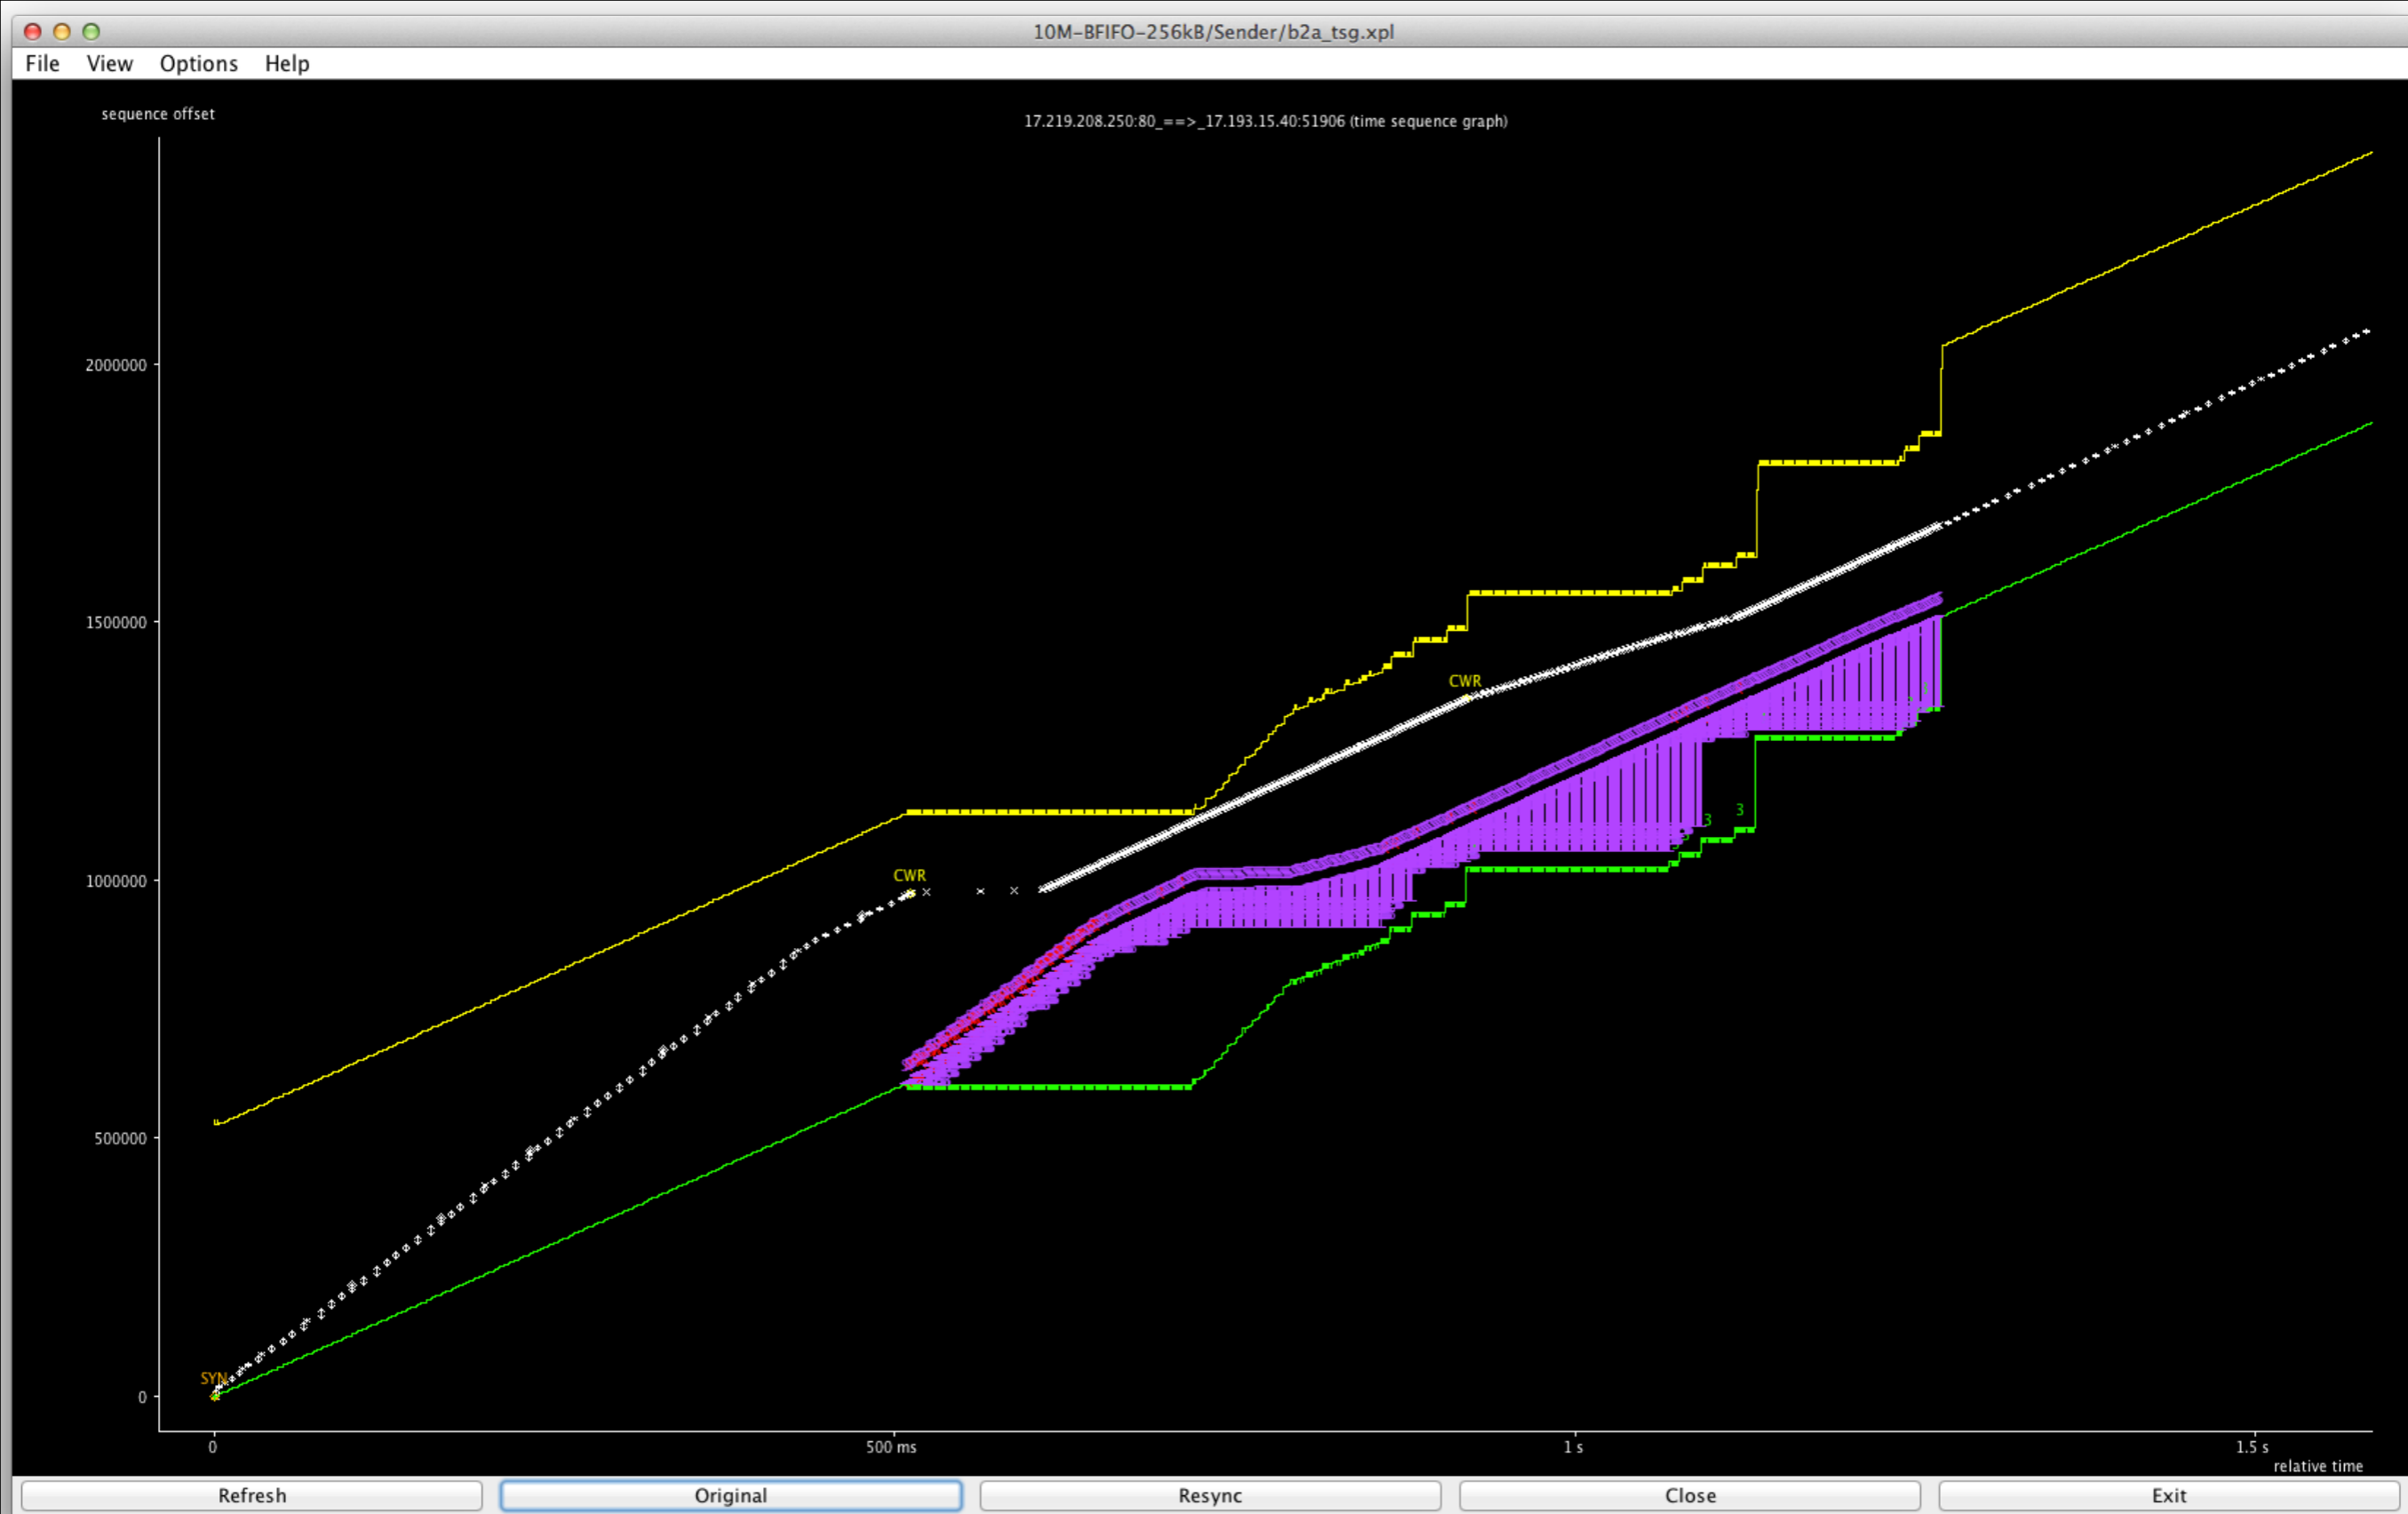
FQ\_CODEL with ECN

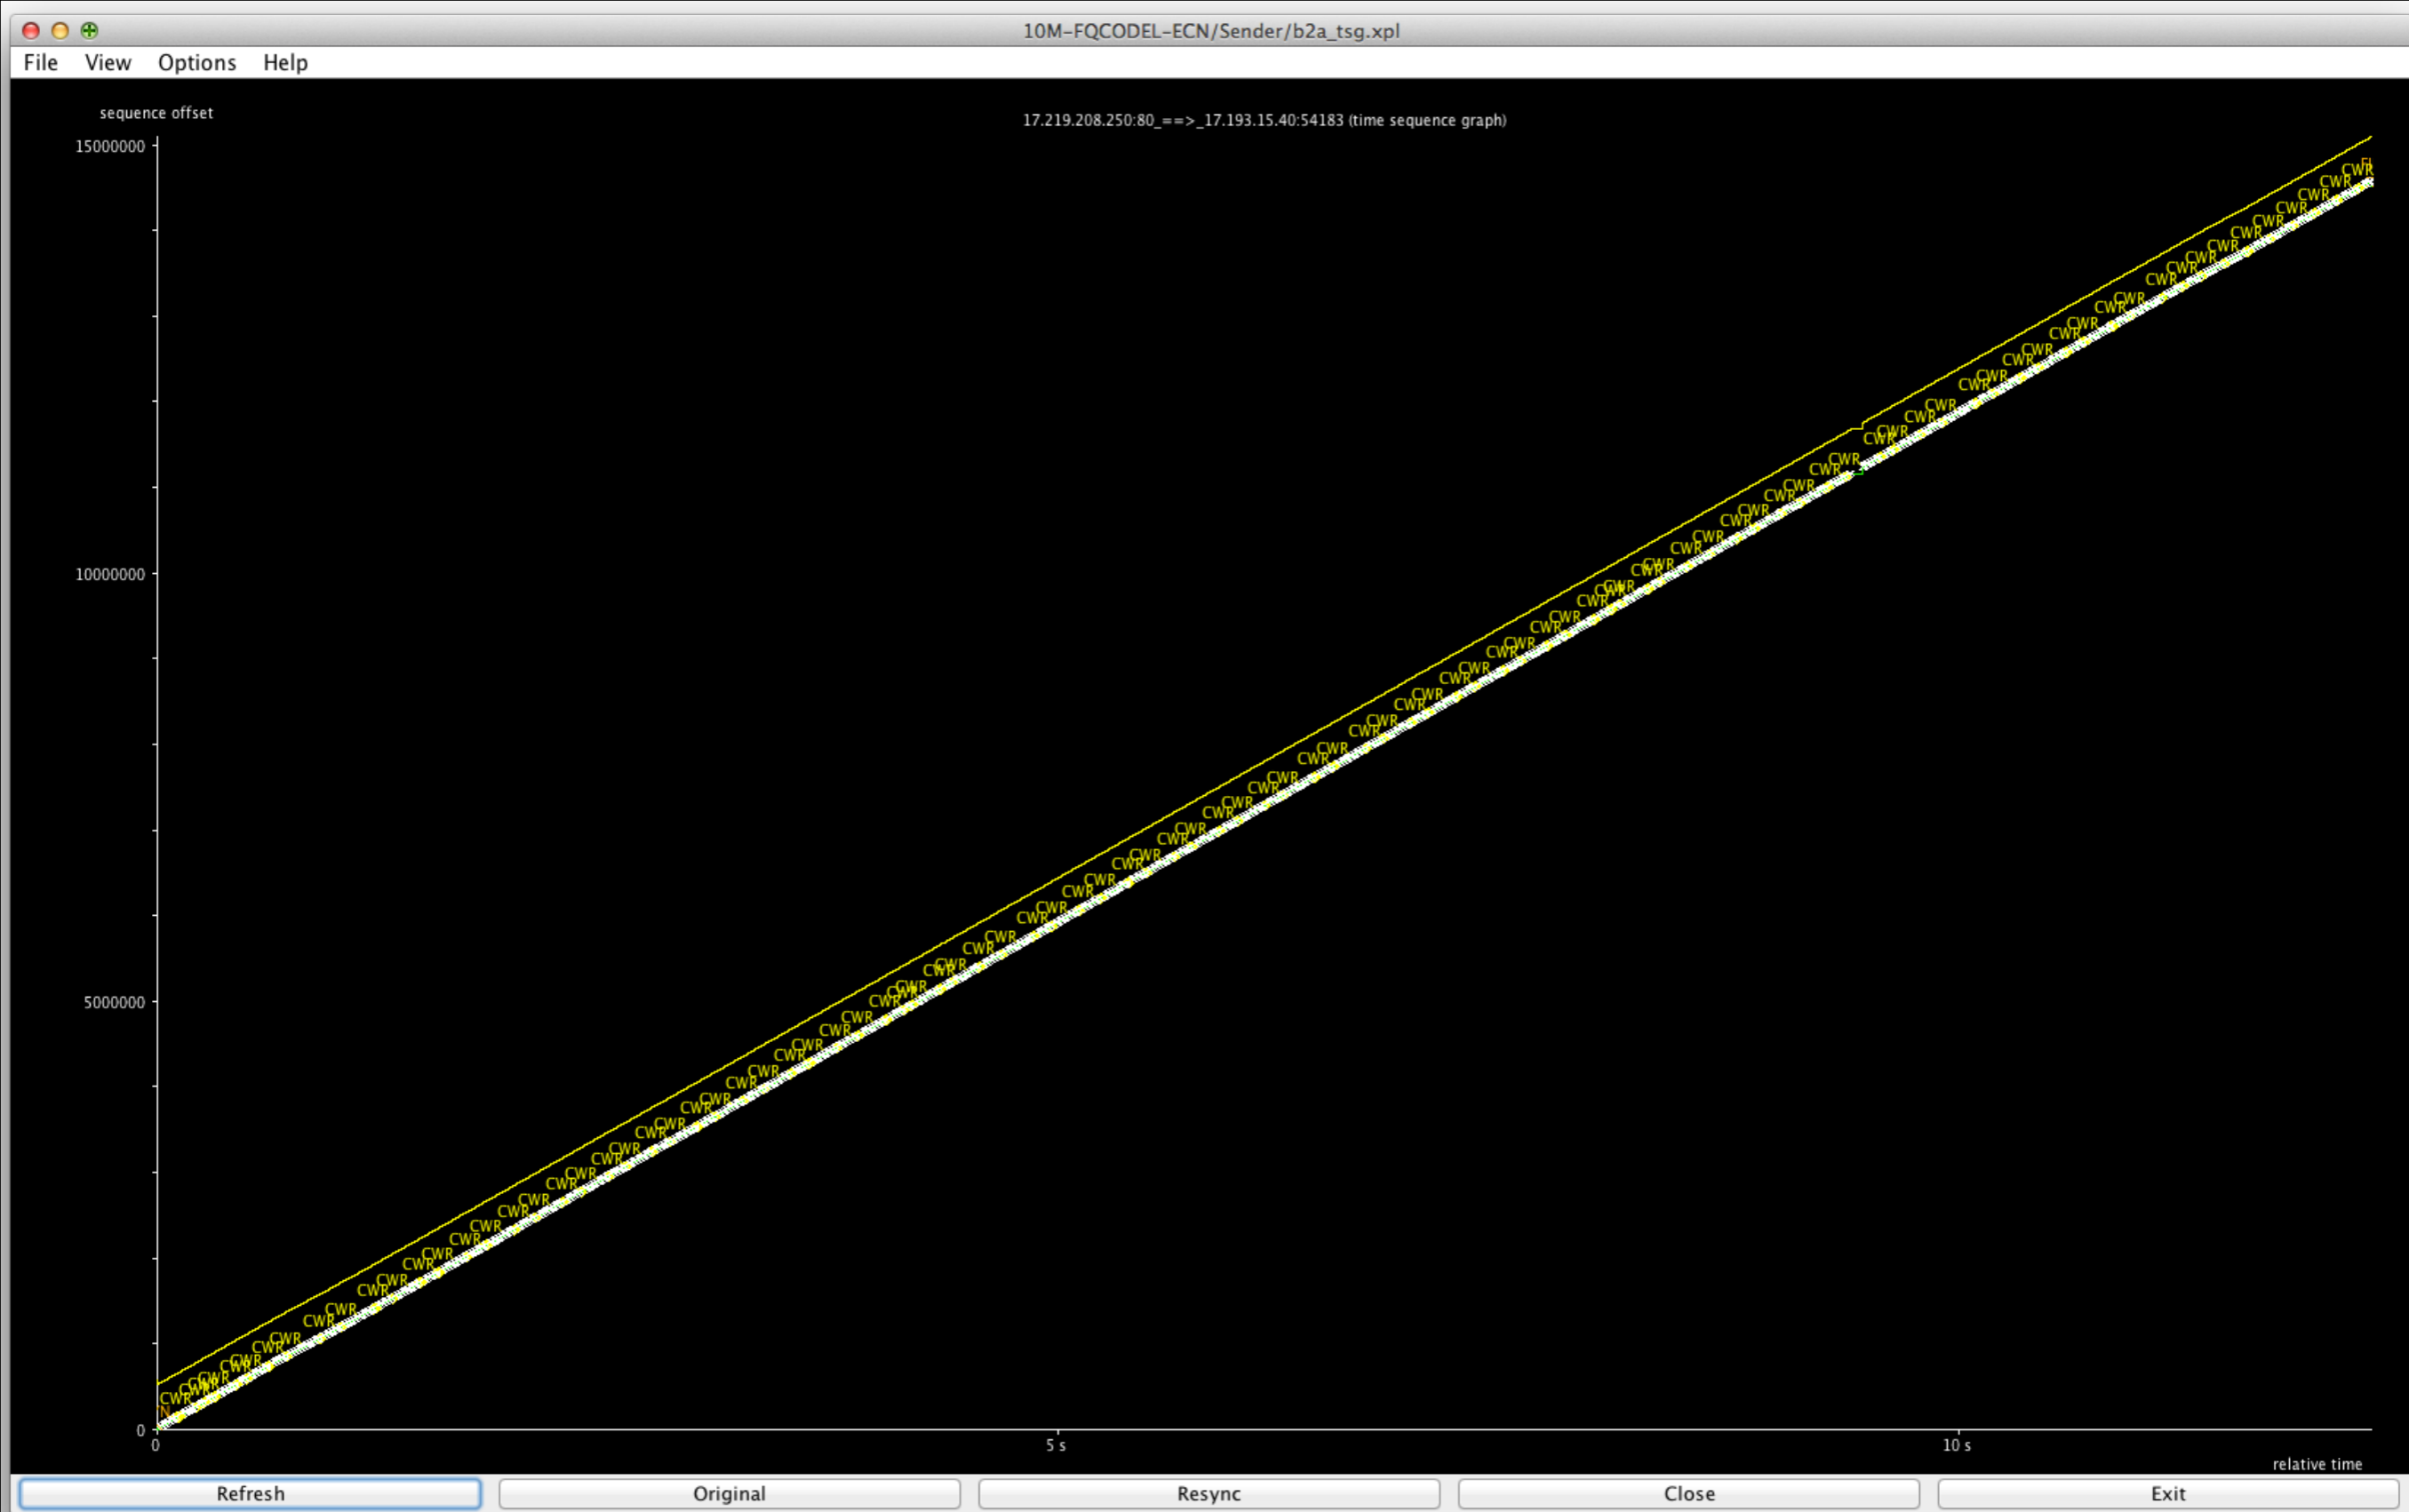
(Results for FQ\_CODEL non-ECN, PIE, etc., also available)

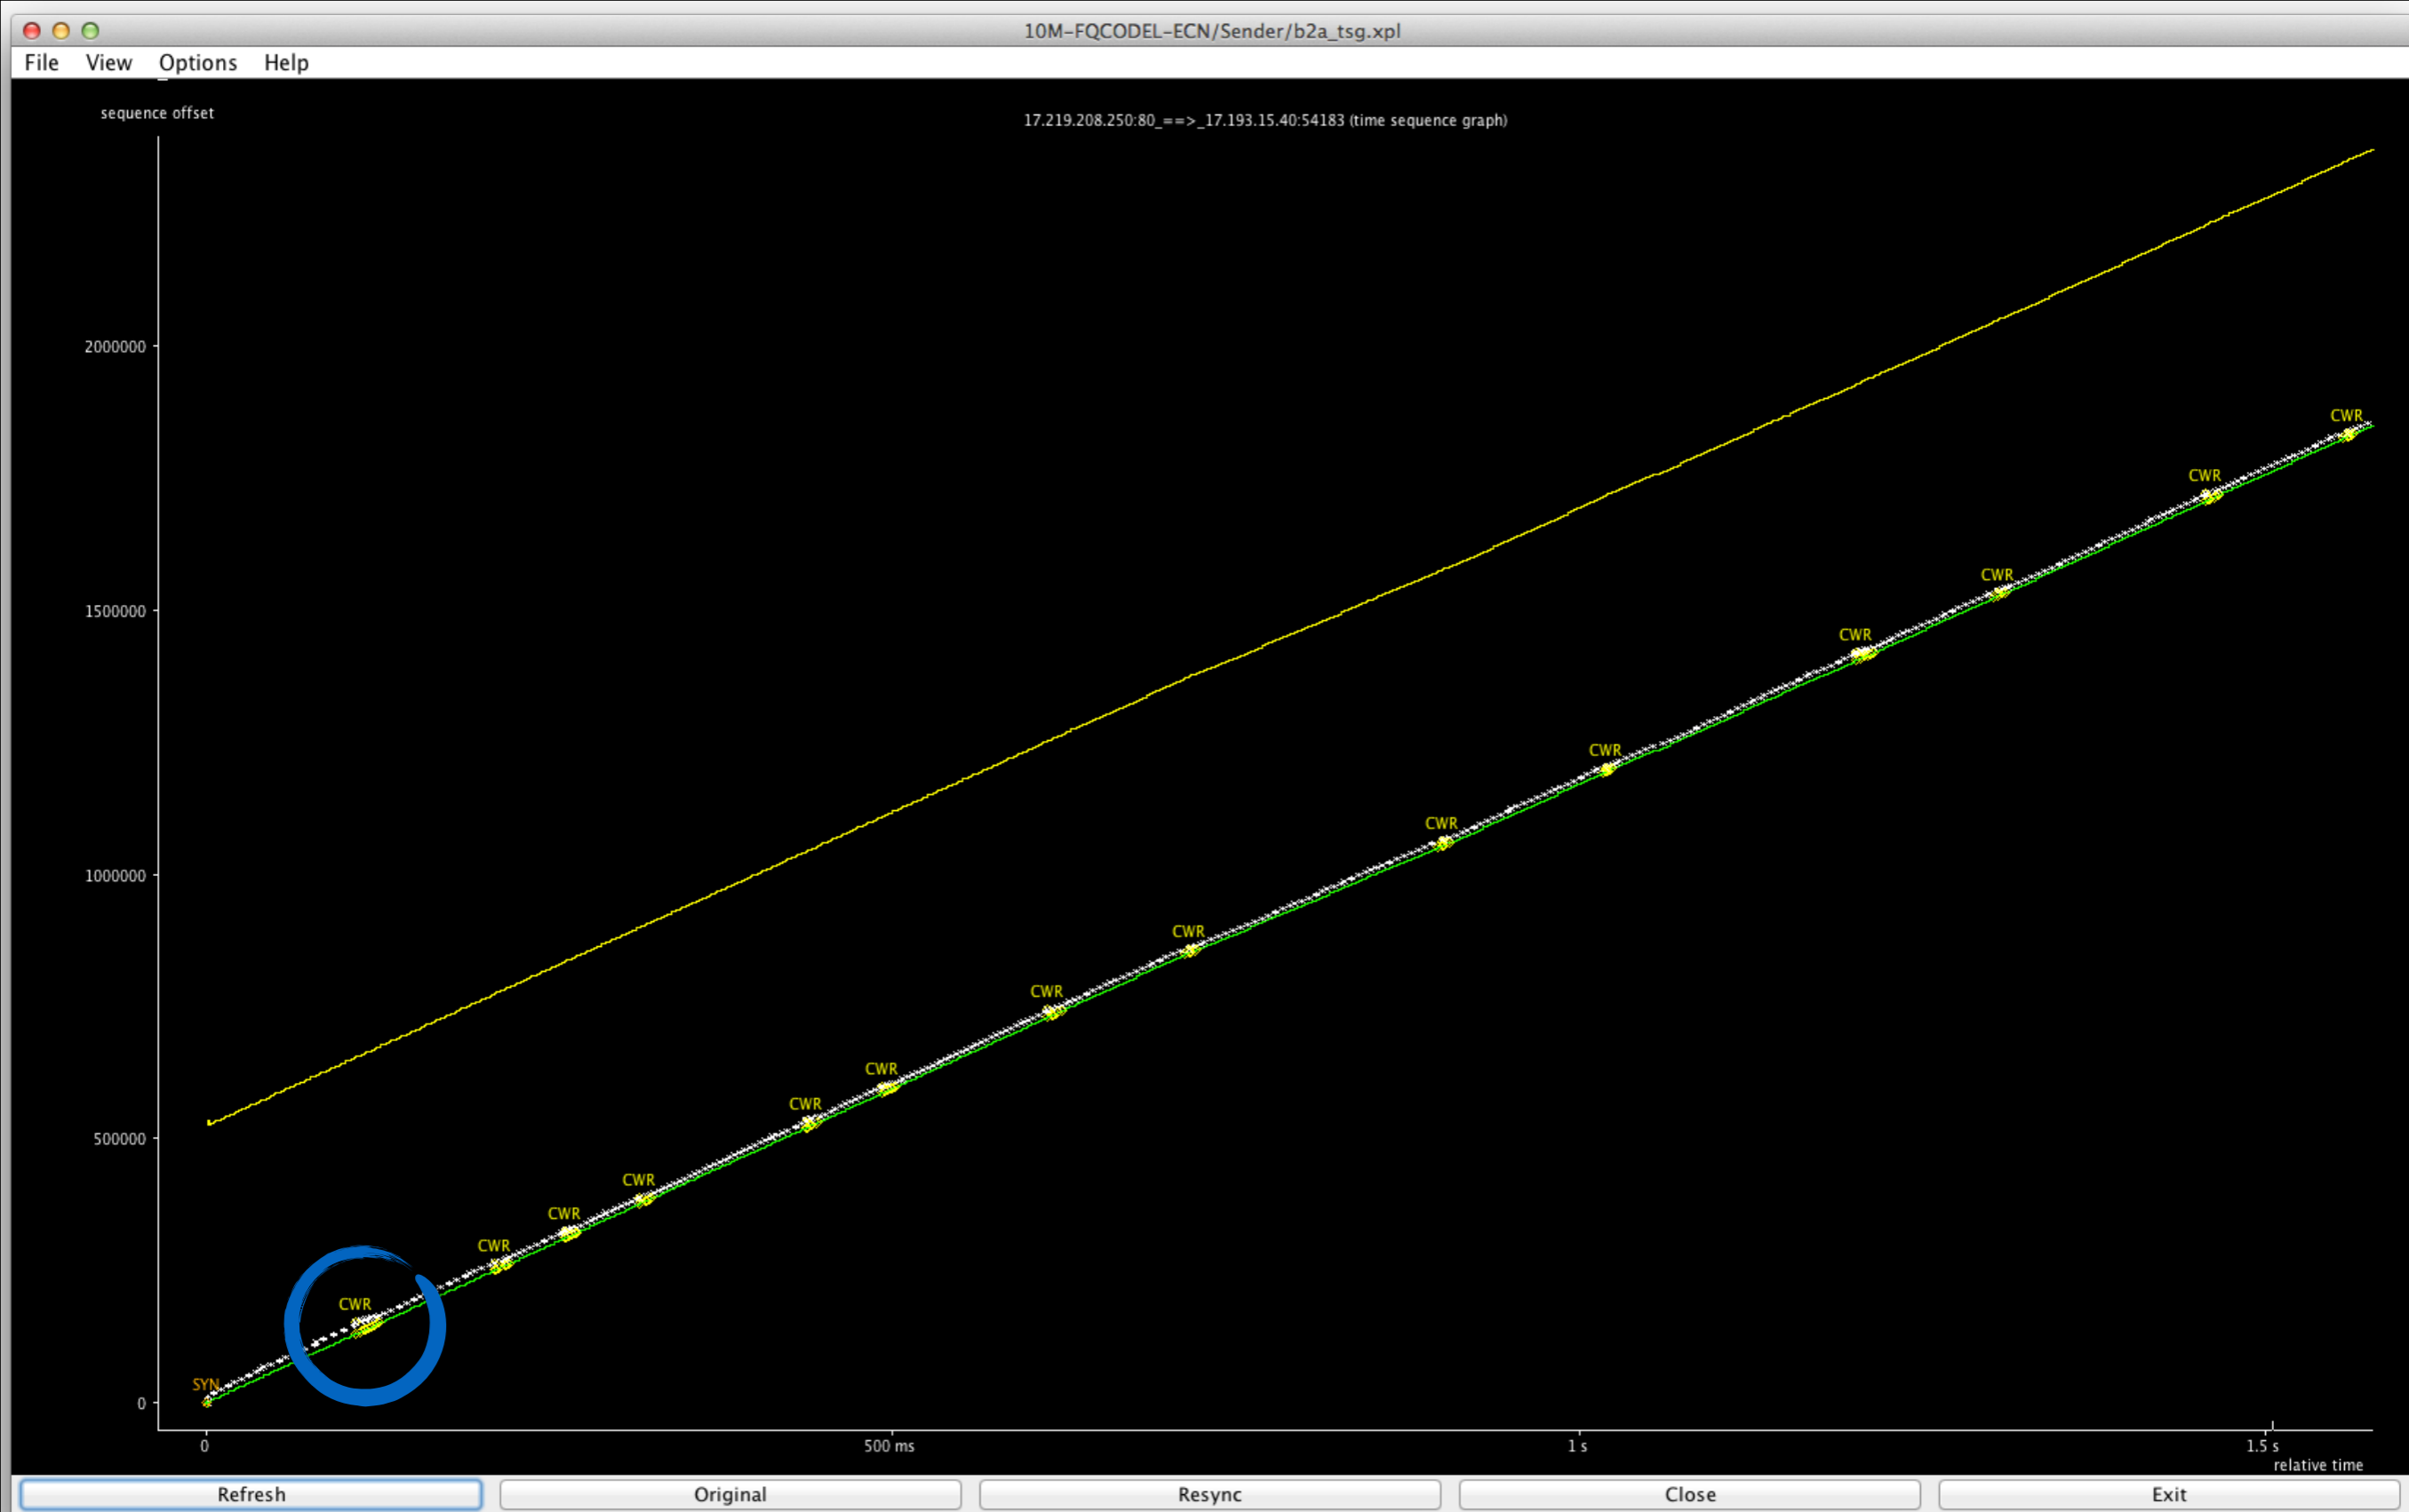
# tcptrace

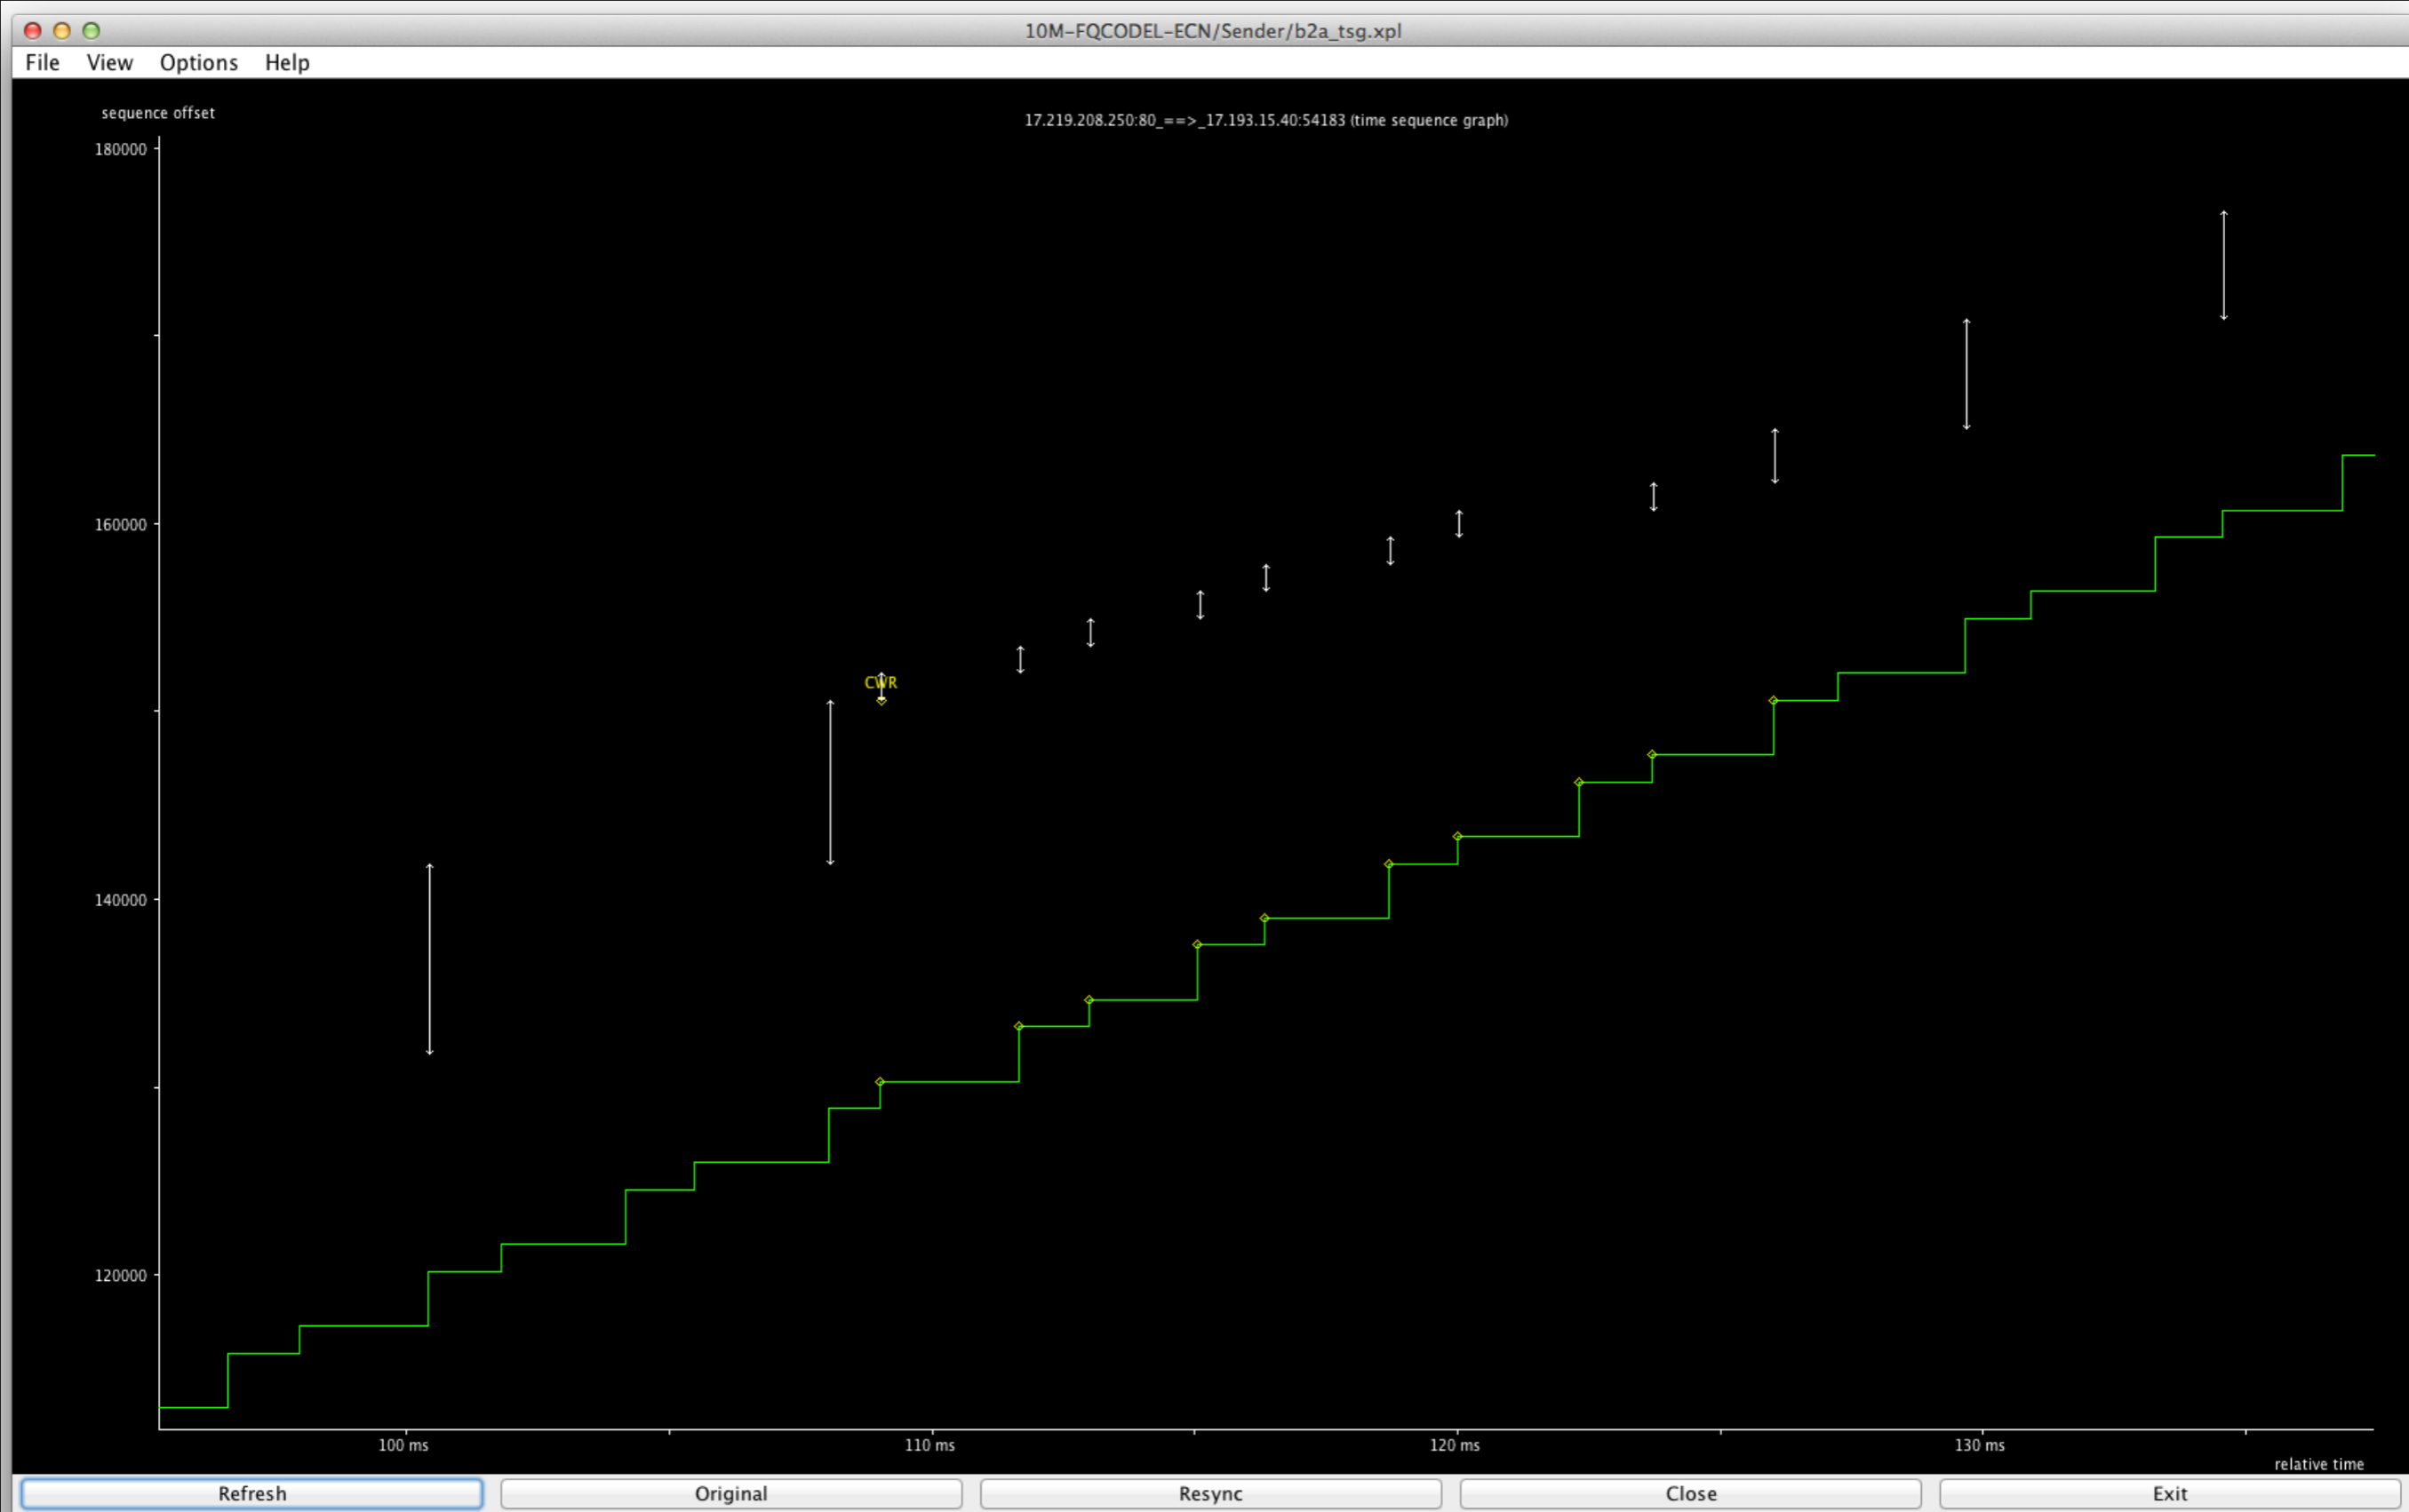
- White lines are packets
- Green line is ack line
- Yellow line is window ceiling
- Red lines are retransmissions
- Purple lines are SACK indications











# Conclusions

- CODEL (or similar AQM) helps
- ECN helps
- AQM+ECN really helps a lot

# Why don't we have ECN?

- Clients
  - Routers aren't doing marking (so, no reward)
  - Some routers might drop the packets (so, risk)
  - All risk; no reward
- Routers
  - What clients are requesting ECN? (none)
  - Enabling ECN might expose code bugs (risk)
  - All risk; no reward

# Where can we turn on ECN?

- Staged Deployment
  1. Data Centre
  2. Residential Cable/DSL ISP network
  3. Entire Internet
- Not advocating for more walled gardens, but...
  - ISP's want ECN for their own video services
  - CMTS and DSLAM start to support ECN marking
  - ECN-capable equipment becomes widely available
  - ECN deployment in public Internet becomes attractive

# Enable ECN on Mac OS X

```
sudo sysctl -w net.inet.tcp.ecn_initiate_out=1
```

```
sudo sysctl -w net.inet.tcp.ecn_negotiate_in=1
```

To make the settings persistent,  
put following lines in `/etc/sysctl.conf`:

```
net.inet.tcp.ecn_initiate_out=1
```

```
net.inet.tcp.ecn_negotiate_in=1
```

# Enable ECN on Windows

```
netsh interface tcp set global ecncapability=enabled
```

# Enable ECN on Linux

```
sudo sysctl -w net.ipv4.tcp_ecn=1
```

Like Mac OS X, to make setting persistent,  
add line in `/etc/sysctl.conf`

[http://en.wikipedia.org/wiki/Explicit\\_Congestion\\_Notification](http://en.wikipedia.org/wiki/Explicit_Congestion_Notification)

# ECN Queueing on Linux

To enable an ECN-enabled AQM in Linux:

```
tc add dev your_device root fq_codel ecn
```

(Linux 3.13 makes swapping the default qdisc a sysctl also)

# Fix the Bufferbloat please!

(It's not just for Skype any more)