Experimental Determination of a Lower Bound for Ta

(STUN check interval)

Or, how fast is too fast?

Let's Gather Data

1. Choose random sample of Chrome dev users (about 400,000 over 7 days)
2. Choose 3 STUN servers
3. Choose a frequency (5-100ms)
4. Send 10 STUN requests to each server with the chosen frequency (ABCABCABC)
5. Count how many responses come back
6. Compare the success rate between different frequencies
7. Bonus: See if the user is behind a NAT, and if it's symmetric or not

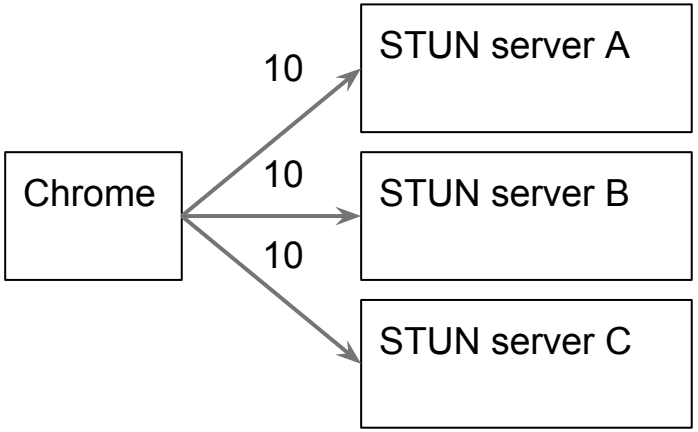

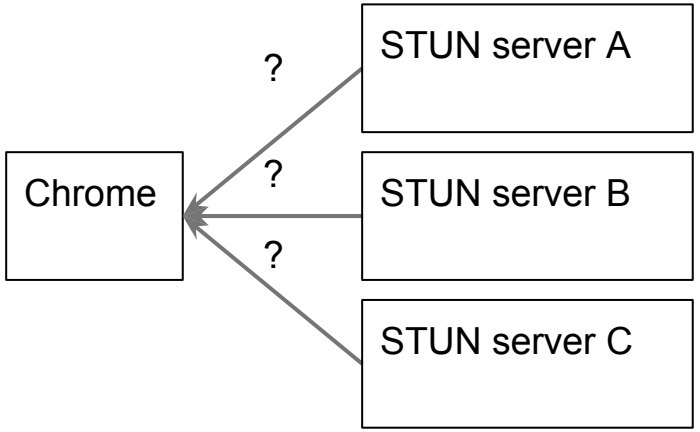

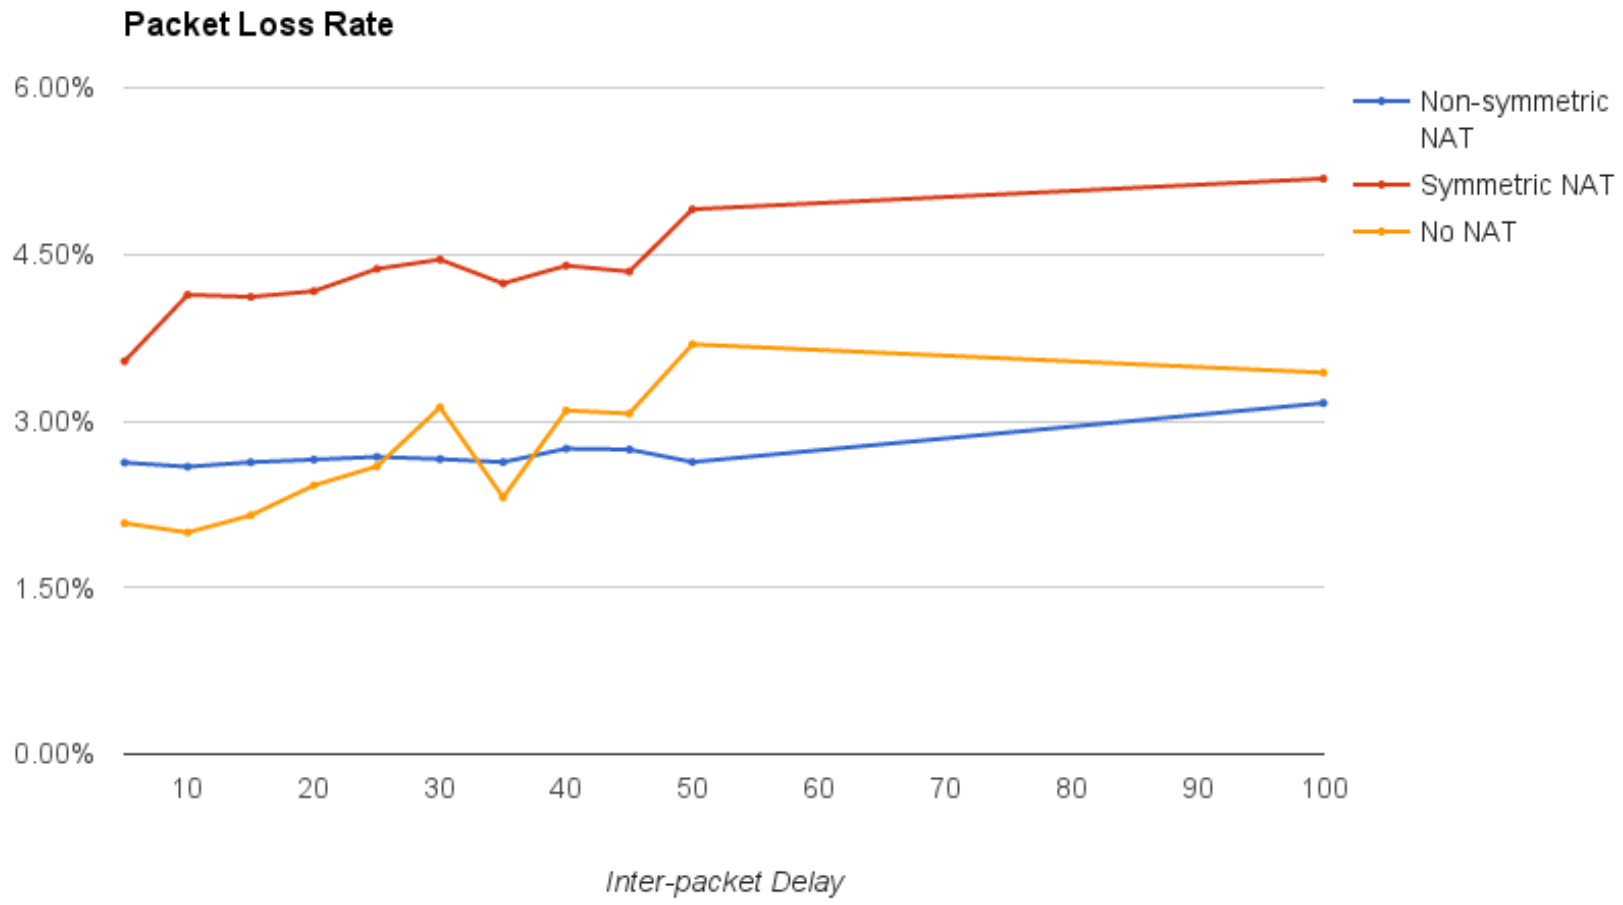
Packet Loss Rate
6.00%
4.50%
3.00%
1.50%
0.00%
10
20
30
40
50
60
70
80
90
100
Inter-packet Delay
Non-symmetric NAT
Symmetric NAT
No NAT

Wait.... what?

Success rate seems to *improve* as we increase the frequency.

We don't know exactly why yet. We do have some hypotheses...

Hypothesis #1: slow computers

- If slower computers are correlated with worse networks
- And if slower computers are sampled more with lower frequencies (because when we send checks, the actual delay is longer than the delay we asked for)
- Then worse networks would be correlated with lower frequencies

But we tracked how often we got a longer delay than we asked for, and adding that data in or taking that data out did not affect the results at all.

So this hypothesis seems bogus.

Hypothesis #2: Bursty loss

- If packet loss is bursty
- And we send all the packets in a shorter period of time
- Then we'll be less likely to hit a packet loss period
- But more likely to have all the packets fail
- But we ignore the data when we get 100% failure (since we can't differentiate it from UDP being blocked)

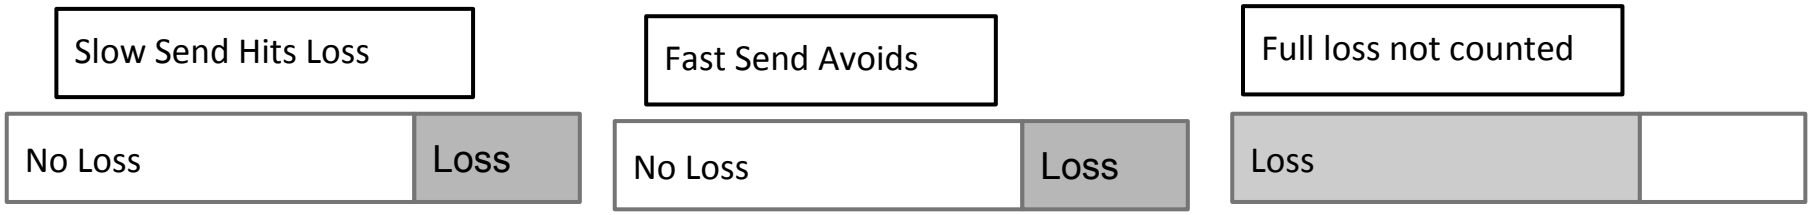

But when we look more closely at the high-loss data, such as 70-90% packet loss, the trend still supports the conclusion that higher frequency leads to lower loss, and that the number of samples would be very small. So this theory doesn't seem to match the data either.

Hypothesis #3: Experimental error

- The more flat curve (Symmetric NAT) has the most samples
- The least flat curves (No NAT and Symmetric NAT) have the fewest samples

We could go get some more samples.

But we have already have a lot of samples, even for No NAT and Symmetric NAT

Hypothesis #4: Dark Energy

- We only observe non-flat curves with no NAT and symmetric NAT
- Enterprise firewalls or CGNs may be doing something that's affecting the curves

This is the best hypotheis we have so far.

Conclusions

- **A 20ms value for Ta is too conservative**
- The real bottleneck for sending checks becomes bandwidth
  - Typical check size = 140 bytes
  - 140 bytes * 50hz (20ms delay) = **56kbps** (+responses!)
  - 140 bytes * 100hz (10ms delay) = **112kbps**
- This suggests that we need to think more deeply about exactly how we do connectivity checks
  - Can we reduce check size? Current minimum is 116 bytes; possibly, this could be reduced to 80 bytes (@ 100hz = **64kbps**)
  - Possibly send 'most likely' checks first at a fast rate, then fall back to slower rate

How you can collect the data yourself

Here's the code:

https://chromium.googlesource.com/external/webrtc/+/master/webrtc/p2p/stunprober/

Please try to replicate these results.

Can I have the data?

We need to check with our data protection folks.

NAT Types