

Limitations of using Real-World, Public Servers to Estimate Jitter Tolerance Of First Person Shooter Games

Grenville Armitage, Lawrence Stewart
{garmitage,lastewart}@swin.edu.au



Introduction

- Networked multiplayer games are an emerging market, making new demands on service quality
- ISPs and hosting companies must understand player tolerance to network characteristics
- Previous work has experimentally estimated network Latency tolerance for First Person Shooters
 - e.g. upper bound of 150-180ms for Quake3
- We attempted to estimate player tolerance to network Jitter, and learned lessons along the way

Our goal

- Watch people playing a public game server
 - Unknown, uncontrolled players 'out there on the internet'
 - Jitter caused by regular network congestion
- Track their:
 - Playing time and 'success'
 - Latency and Jitter
- Attempt to draw relationship between player satisfaction and network jitter



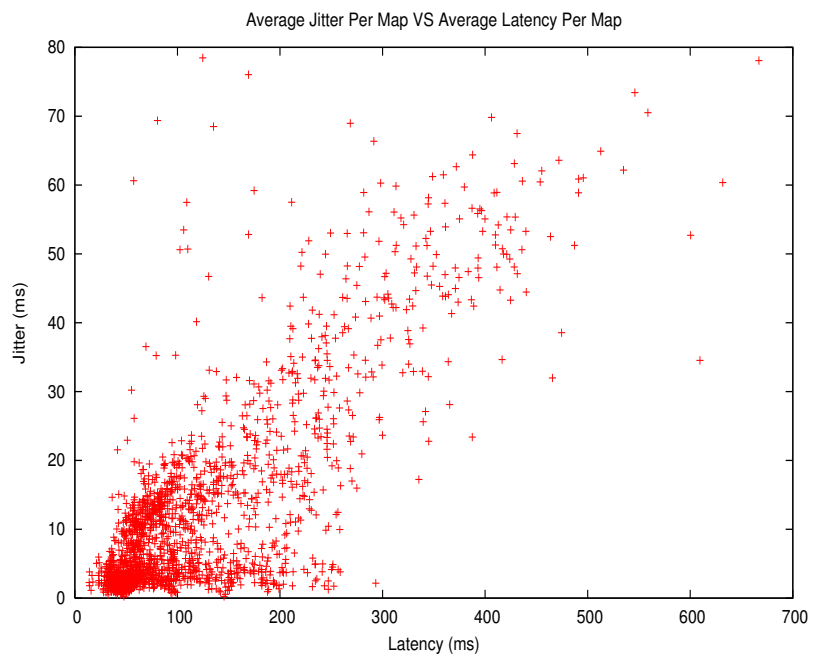
Summary of our results

- Ran an active, public Quake3 server for 3 months
 - 1837 unique players, 4931 games and 11138 maps
- Collected 'ping' samples ~20 times per second
- Discovered latency and jitter quite correlated
 - Jitter appears to be $\leq 20\%$ of latency
 - Latency is "too high" before jitter becomes significant
 - Difficult to isolate jitter's contribution
- Estimation of jitter tolerance should be done in controlled lab environments, not 'on the net'



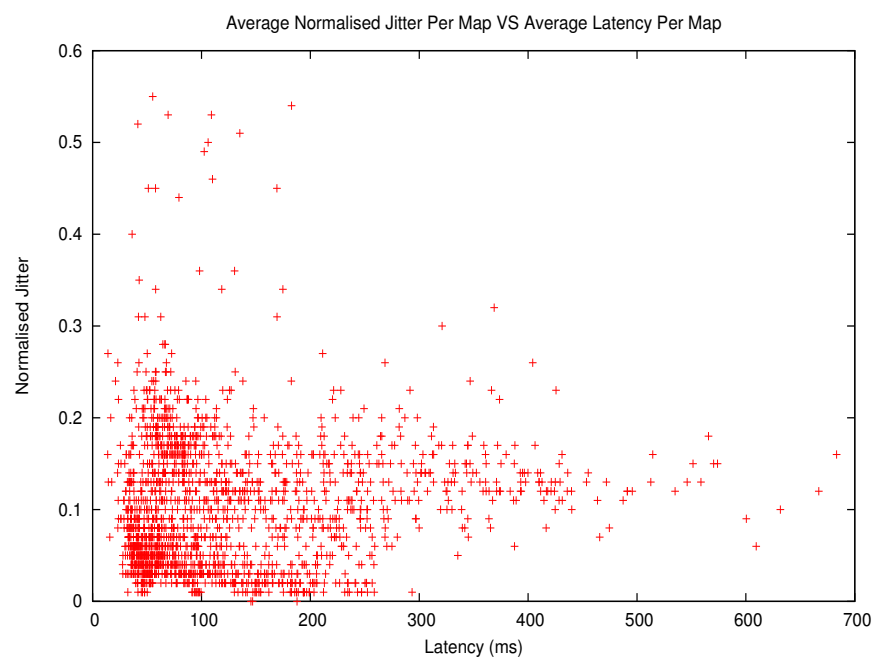
Jitter vs Latency

- Scatter plot per map
- Two regions
- Many short hops
 - High latency, high jitter
- Few long hops
 - High latency, modest jitter



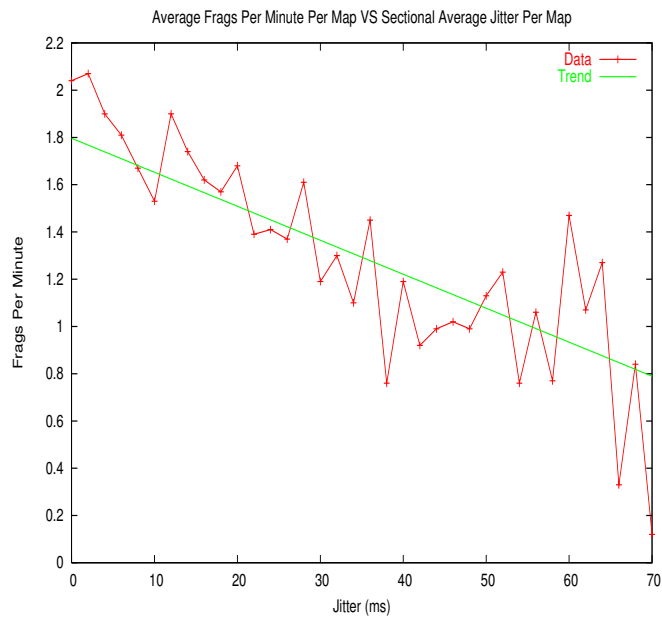
Normalised Jitter vs Latency

- Jitter/Latency vs Latency
- Most jitter is $\leq 20\%$ of latency
- (Also, under 250ms we see mix of longhop and multi-shorthop paths. Over 250ms, mostly multi-shorthop paths)



Player effectiveness?

- One measure is “frag rate”
 - “kills per minute”
- And yes, seems to be a relationship
 - But latency is a hidden influence



Consequences of Correlation

- Using uncontrolled players from around the Internet limits the observable latency-jitter permutations
 - We see high latency/high jitter, low latency/low jitter
 - Rare to see high latency/low jitter or low latency/high jitter
- Does jitter really matter?
 - By the time jitter reaches 30ms, the latency is likely to be up around 150ms, where FPS players noticeably degrade anyway
 - Or is previous work flawed because 30ms jitter PLUS 150ms latency is what kills the players?

Conclusion

- This paper is a cautionary tale to experimental researchers
- The Internet's own characteristics limit the spread of data you'll accumulate
 - User experience trials should be backed up with controlled lab-based experiments to test jitter/latency combinations not seen 'in the wild'
- But then again, the data suggests that current ISP engineering approaches are fine – keep the latency down, and the jitter will stay tolerable