SWINBURNE

CENTRE FOR ADVANCED INTERNET ARCHITECTURES

An Experimental Estimation of Latency Sensitivity in Multiplayer Quake 3

Grenville Armitage

garmitage@swin.edu.au

Introduction

- This talk reports on an estimation of latency sensitivity using real-world Quake 3 servers in 2001
- Why?
  - Game designer: What do I need from the network?
  - Game hosting company: Where (realistically) will my regular customers be located?
  - ISP: What are my performance objectives?
  - “Latency is bad” just isn't enough of a guideline

Background

- Hypothesis
  - Server usage patterns will reflect *topological locality* of players (and relate to latency tolerance)
- Methodology
  - Establish two <u>non-colocated</u> QuakeIII servers that appear identical to client-side selection process
  - Log players, their IP addresses, and in-game 'ping' samples over period of months
  - Assess topological locality of players, and distribution of observed ping values.

Test Environment

- Well connected servers
  - *Californian server:* 600MHz Celeron, 128MB, FreeBSD4.2, T1 link to PAIX (hosted in Palo Alto)
  - *London server:* 900MHz Athlon, 128MB, Linux kernel 2.4.2, 10Mb link to UK net (hosted at University College London)
- <u>Both</u> servers advertised their location as "Palo Alto, California" *(to GameSpy3d and other master-server game selection clients)*

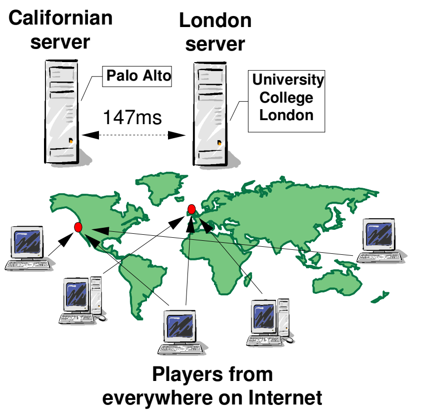

Test conditions

Duration of Trials:

- **Californian server:**
  May 17 to Aug 18, 2001
  5290 unique clients
  338 clients played >= 2hrs each
  164 'days' aggregate played time
- **London server:**
  May 29 to Sep 12, 2001
  4232 unique clients
  131 clients played >= 2hrs each
  77 'days' aggregate played time

Common server details:

- Quake III version 1.17 (linux binary)
- Same 6 maps, fixed cycle sequence
- 20 minutes per map
- Up to 6 remote players
- 2 permanent 'bots' to attract players
- Identical registration with master server (clients see latency as only difference)
- Server-side 'ping' sampled everytime player runs over an object, dies, or kills another player

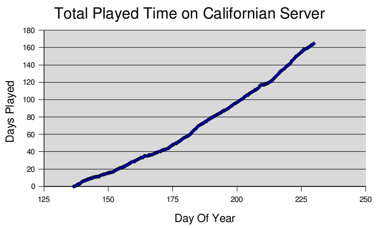

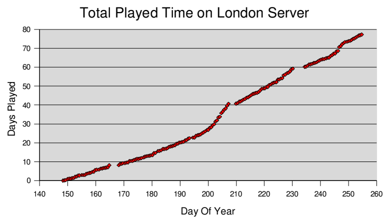

Median Latency results

- Each player's 'ping' sampled > 10 times per game
- Median values <u>per player per game</u>
- Cumulative plot reflects most frequently appearing median ping values
- California and London curves similar

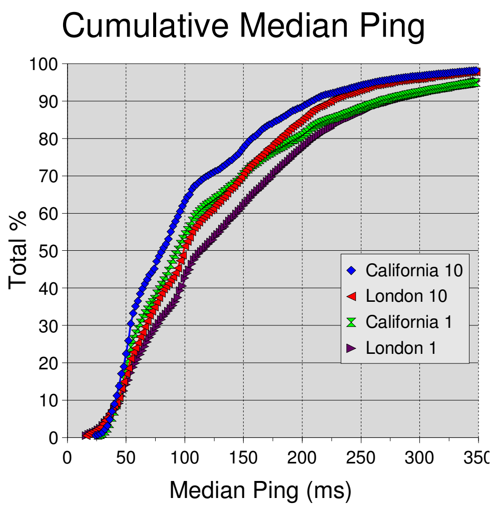

Reading the curve...

- Players who picked up at least 1 item per minute (minimal activity)
  - California 1: 80% of player.games < ~196ms
  - London 1: 80% of player.games < ~210ms
- Players who picked up at least 10 items per minute (reasonably active)
  - California 10: 80% < ~158ms
  - London 10: 80% < ~182ms

Don't forget why they play...

- Skill and response time influence a player's ability to frag (kill) others in the game
  - Response time has human and network components
- Average frag rate vs median ping hints at the negative impact of high latency
  - 45ms ping averages 1 frag/min better than 200ms ping

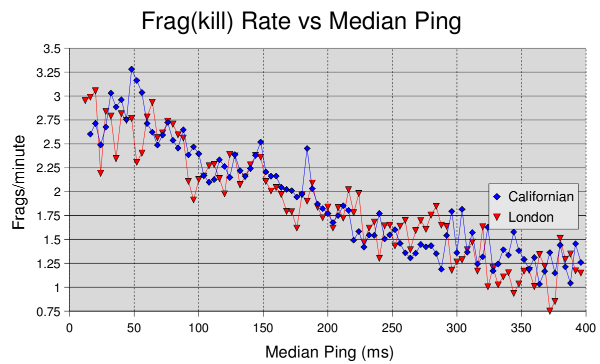

But what does this prove?

- Perhaps nothing!
  - Maybe “the Internet” is only 250ms wide?
- Unless there's evidence of regional/topological locality in the usage of each server....

Weekly pattern

- Pattern consistent with daily curve
  - Geographic locality ~ topological locality

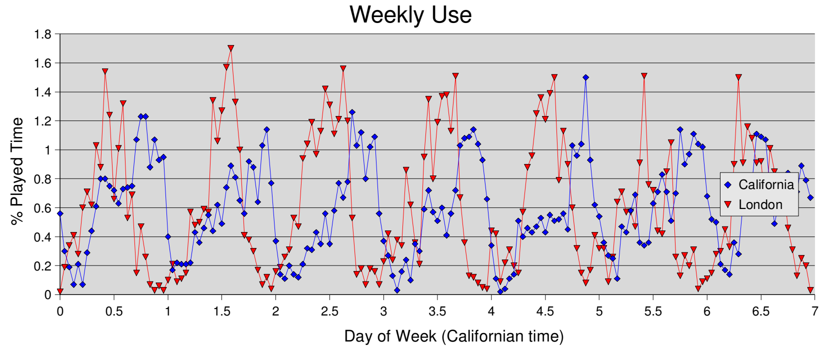

Regional locality - Daily pattern

- Usage peaks around afternoon/evening in their respective time zones
  - (London 8 hours ahead of Palo Alto)
- Servers attract regional players
  - Supports hypothesis that clients prefer 'closer' server, other things being equal

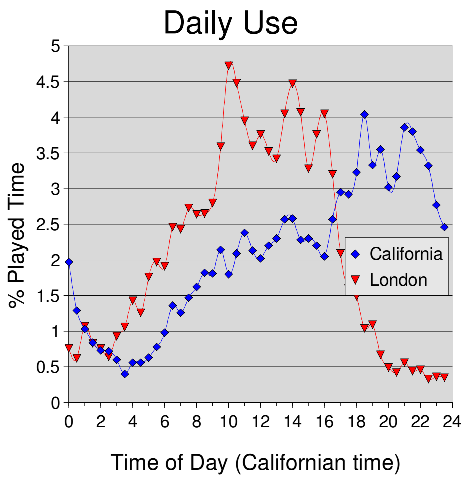

Locality based on IP addresses

- Reverse lookups:
  - Californian server: mostly *North America*
  - London server: mostly *Europe and US East Coast*
- Given otherwise identical servers, latency seems plausible as the *client-observable* metric on which a player chooses their server

Using active players who picked up at least 10 items per minute during each game:

Rank	*Calforinia Games/Time(min)*	*Calforinia Origin*	*London Games/Time*	*London Origin*
1	323 / 3005	.ed.shawcable.net	108 / 1027	.pit.adelphia.net
2	192 / 2072	.cruzio.com	73 / 690	.Uni-Mainz.DE
3	124 / 1383	(RogersEAST/@Home)	75 / 679	.upc-d.chello.nl
4	119 / 1246	.018.popsite.net	50 / 606	(telnordia.se)
5	118 / 1221	.tx.home.com	53 / 604	.dyn.optonline.net
6	150 / 1200	.mediaone.net	44 / 565	(Rogers EAST/@Home)
7	132 / 1178	.pit.adelphia.net	35 / 463	.dyn.optonline.net
8	115 / 1151	.socal.rr.com	53 / 448	.dialup.tiscalinet.it
9	87 / 980	.pa.home.com	34 / 430	.pa.home.com
10	93 / 938	.sfba.home.com	20 / 288	.tx.home.com
11	69 / 799	.hsia.telus.net	24 / 273	.btinternet.com

() bracketed origins involved looking up 'whois' database after .in-addr.arpa failed.

Table above shows origins of top 11 players on each server. Outside the top 11, the Californian server also saw dedicated players from ".jp" while the London server saw dedicated ".nl" and ".uk" players. There is also cross-over by players equidistant from either server.

Conclusions

- Players self-selected based on topological locality of servers, even though servers lied about their actual location
  - Latency was the visible metric by which this selection occurred (typical players unlikely/unable to check IP address)
- Thus the median-ping per player-game stats seem reasonably likely to reflect player preference
- Open question: High latency ~= high hop count paths, thus could be correlated with high jitter and loss rates.... research ongoing

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES