# Experimentation with the YouTube Content Delivery Network (CDN)

**Siddharth Rao**
Department of Computer Science
Aalto University, Finland
siddharth.rao@aalto.fi

**Sami Karvonen**
Department of Computer Science
Aalto University, Finland
sami.karvonen@aalto.fi

**Abstract**
In this paper we analyze the behavior of YouTube Content Delivery Network from five different locations. The first part of the experiment deals with most famous videos on YouTube, where we identify the content servers IP address and analyze the criteria behind it. We have used Pytomo tool for fetching the data that we need for our analysis. In the second part of our experiment, we observe and analyze how a new content gets populated over the content delivery network.

***Keywords:*** Content Delivery Network, YouTube CDN analysis, Pytomo ,

## 1. Introduction

Internet is constantly evolving in terms of its topology and hence the way the content over internet that is being delivered to the users. Simultaneous unicast multimedia streaming to many customers have been enabled because of increasing access and bandwidth in fixed line as well as in mobile domains. Content Distribution Networks (CDNs) serves as the best choice at this point of Internet era to implement this along with many other large scale services.

Over multiple backbones CDN nodes are generally deployed in multiple locations. Benefits of using such a mechanism for content delivery are as follows:

- Reduction in bandwidth costs
- Improving page load times
- Increasing global availability of content

The number of nodes and servers making up a CDN varies, depending on the architecture, some reaching thousands of nodes with tens of thousands of servers on many remote Point of Presences (PoPs). Others build a global network and have a small number of geographical PoPs.

Requests for content are typically algorithmically directed to nodes that are optimal in one or the other way. When optimizing for performance, locations that are best for serving content to the user may be chosen. This may be measured by choosing locations that are the fewest hops, the least number of network seconds away from the requesting client, or the highest availability in terms of server performance (both current and historical), so as to optimize delivery across local networks. When optimizing for cost, locations that are least expensive may be chosen instead. In an optimal scenario, these two goals tend to align, as servers that are close to the end-user at the edge of the network may have an advantage in performance or cost.

## 2. Background

Users with high-speed connections often experience choppiness, loading lags and poor quality, especially when viewing live events or if they are located far from the hosting servers. CDNs minimize latency issues that cause image jitters, optimize delivery speeds and maximize available bandwidth for each viewer. Also search engines now penalize sites that load too slowly, and broadband penetration has created unprecedented traffic jams. Fast connection times do little good for viewers if the content delivers slowly. Customer attention spans become increasingly short because people spend good money on their connections and equipment, expecting the fastest page loads and highest quality videos and live-streaming events. CDN technology provides redundancy for fail-safe protection during partial Internet malfunctions. Duplication of content also protects against loss of data and image degradation.

### 2.1 Advantages of CDNs[2]

- Eliminate Pauses and Accommodate Heavy Traffic
- Minimize Packet Loss
- Faster Loading
- File Mirroring
- Optimize Live Delivery
- Enables Linear Networks
- Scalability

### 2.2 Disadvantages of CDNs

Disadvantages include additional costs, but large companies often save money by delivering their content along shorter routes.

- Sharing Resources - Delivery networks have many clients, and response times could vary due to the volume of website traffic of other CDN customers.
- Geographical Choice Considerations - Website owners need to research their clients and choose the CDNs that offer the most convenient server locations near where they get the most business.
- Content Management Problems - Some customers might need special widgets, content players or specific applications. Customer interactions could have different requirements in certain regions
- Lack of Direct Control **-** Changes to content must be made through CDN providers instead of directly, which could pose problems for editors and developers.
- New points of Failure: CDNs create new points of potential failure along the delivery chains.

## 3. Experimentation setup

To understand the serving mechanism by the CDN servers( in the query -reply format), we tested the end to end process on our local machines and analyzed the TCP/UDP packet streaming with the Wireshark tool. The IP of the CDN server is analyzed based on the connection which has more traffic compared to all other proxy servers. This also indicates the the IP of the server which actually replies to the query from local machine using GET method. After this analysis, we tried to experiment in various locations for measuring. We crawled the same videos URLs about 7-8 times over all the locations that we chose for our experiment, which are carried out at different days and different time.

### 3.1 Measurement locations

From AWS we used medium sized Ubuntu instances in Singapore, California and Ireland and the other two locations were a home machine in Espoo and Kosh from Otaniemi. In all of these locations we ran the Pytomo tool with 23 static Youtube links[5] for the first part and the link to the uploaded video for the second part.

### 3.2 Tools

We used pytomo[4] as a measurement tool for both parts of the assignment. Pytomo measures all the information needed for both assignments but doesn't print it out in any good way so we made a little modification to it which prints out all the needed information to the log file. The modification is one output print which prints out timestamp, YouTube link, IP of the content server, ping to content server, file size and time to get first byte of the file. The print command is in the start_pytomo.py file at the end of compute_stats-method after all the necessary values are calculated and looks like this:

```
config_pytomo.LOG.info(‘ASSIGNMENTPRINTTAG
%s %s %s %s %s %s’ % (timestamp, url, ip_address,
ping_times[0], download_stats[3], download_stat[13]))
```

The static Youtube links can be added to the config_pytomo.py file to the STATIC_URL_LIST-variable. In the second assignment we were told to measure the time between the HTTP-request and response but we figured that getting the time between the HTTP-request and getting the first byte of the movie would measure the same thing and that time was already available at Pytomo so we decided to use it for both parts of the assignment.

### 3.3 Procedure (Commands)

After the modifications to Pytomo the experiment was very simple. We uploaded Pytomo to all the location. Then we ran the command:

**python start_crawl.py**

Then we used grep to get the information from the log-file with command:

**cat [log-file name] | grep -oP “ASSIGNMENTPRINTTAG\K.*”**

This prints out the needed information from every link crawled. After that we downloaded a video to Youtube changed the static links in Pytomo to the downloaded video and ran the two commands again in every location.

## 4. Analysis

Our data set for the experiment was 23 Youtube links which were selected by Pytomo randomly crawling weeks most popular links and one movie uploaded to Youtube by us for the second part of the assignment. The data for the 23 links were recorded with Pytomo from 5 locations which were Espoo [Table 1], Otaniemi [Table 2], Ireland [Table 3], California [Table 4] and Singapore [Table 5].

### 4.1 Content servers

In this experiment we located three Youtube content servers which are located in Sweden, Switzerland and California. The ip-addresses that start with 80.239. are from Switzerland, addresses starting with 109.105. are from Sweden and 173.194. or 74.125. are from California. There are more content servers around the world but these are the ones which were used from the five locations of this experiment.

All of the locations seem to have one or two favourite content server which they download the movies from but it isn't always the closest. This might indicate that Youtube CDN prefers to use the closest data centers but also redirects traffic to other data centers to balance the load between them. The other reason for this could be that the closest data center doesn't have the video but it's unlikely because these are most popular links and all the main data centers should have them already.

If you compare the content server used from Espoo and Otaniemi, you notice that from Espoo the videos come either from Switzerland or California and from Otaniemi they come from either Sweden or California. The videos that come from California are mostly the same videos in both cases so they are probably videos that are popular only in U.S. and the data centers in Europe doesn't have them. The reason why the other videos come either from Sweden or Switzerland could be the CDN balancing the load between the two data centers. In Ireland's case all the videos came from California which shows that CDN sometimes gets the videos from very far away even though there are many other data centers closer. This could be explained by the fact that the experiment was done during daytime in Europe which means that in California it was night so the content server there had a lot of spare capacity.

From both California and Singapore locations all the videos came from the content servers in California. For California it's is obvious that the videos come from the closest content server but in Singapore's case the reason is bit more unclear. It could be either load balancing or that the content servers closest to Singapore don't have the videos. Most probably it's the latter because the videos are popular videos in U.S. but might not be as popular in Asia.

### 4.2 Our Video

In the second assignment we downloaded our own video to Youtube and used Pytomo again to figure out how it'll propagate in the CDN. From most of the locations our video was downloaded from the same content server as the other videos so it seems our video propagated quite fast. Only exception was when downloading from home machine in Espoo which used either Switzerland or California server for the 23 random links but changed to Swedish server for our video. This indicates that Switzerland server didn't have the video and redirected to the Swedish server. Also latency and the time to get the first byte of the movies is quite a lot higher than in the other videos downloaded from this location.

From the data we gathered we can conclude that our video propagated to the California server and the Swedish server pretty fast but it didn't go to the Switzerland server and this is because the California server is the main server and the Swedish server is closest to the upload location. From Ireland and Singapore the video was downloaded from California which unfortunately don't tell us much because these locations downloaded all of their videos from California.

## Results :

| Sl.No | Date | Time Stamp | IP | Ping to Content Server | File Size | Time to get first byte |
|---|---|---|---|---|---|---|
| 1 | 2013-12-01 | 13:21:41,45 | 80.239.229.238 | 2,02 | 12286640 | 0,010454178 |
| 2 | 2013-12-01 | 13:22:42,41 | 173.194.48.24 | 9,152 | 10058880 | 0,071233988 |
| 3 | 2013-12-01 | 13:23:43,96 | 173.194.48.41 | 9,364 | 542134 | 0,053607941 |
| 4 | 2013-12-01 | 13:24:03,48 | 80.239.229.236 | 1,527 | 11295343 | 0,037833929 |
| 5 | 2013-12-01 | 13:24:53,63 | 80.239.229.239 | 1,827 | 22390541 | 0,010363817 |
| 6 | 2013-12-01 | 13:25:44,28 | 80.239.229.241 | 1,837 | 5582937 | 0,048713923 |
| 7 | 2013-12-01 | 13:26:15,02 | 80.239.229.239 | 1,833 | 26941828 | 0,00950098 |
| 8 | 2013-12-01 | 13:27:05,33 | 80.239.229.205 | 2,143 | 20734550 | 0,00880003 |
| 9 | 2013-12-01 | 13:28:04,87 | 74.125.163.70 | 9,031 | 22472544 | 0,061602831 |
| 10 | 2013-12-01 | 13:29:06,41 | 173.194.48.70 | 9,099 | 15855766 | 0,047533989 |
| 11 | 2013-12-01 | 13:29:56,07 | 80.239.229.209 | 1,739 | 14697109 | 0,009244204 |
| 12 | 2013-12-01 | 13:30:55,06 | 74.125.163.42 | 9,178 | 31437962 | 0,077466011 |
| 13 | 2013-12-01 | 13:31:54,10 | 74.125.163.54 | 9,293 | 23823489 | 0,054216146 |
| 14 | 2013-12-01 | 13:32:53,19 | 74.125.163.73 | 9,104 | 8174529 | 0,060096025 |
| 15 | 2013-12-01 | 13:33:52,92 | 74.125.163.9 | 9,169 | 399742 | 0,055113792 |
| 16 | 2013-12-01 | 13:34:22,56 | 74.125.163.115 | 9,075 | 3343766 | 0,043750048 |
| 17 | 2013-12-01 | 13:34:56,54 | 74.125.163.115 | 9,241 | 5671998 | 0,068995953 |
| 18 | 2013-12-01 | 13:35:34,59 | 80.239.229.205 | 1,745 | 3630300 | 0,011542082 |
| 19 | 2013-12-01 | 13:36:00,65 | 80.239.229.208 | 1,584 | 2344045 | 0,009463787 |
| 20 | 2013-12-01 | 13:36:29,91 | 74.125.163.113 | 9,022 | 19682496 | 0,052381992 |
| 21 | 2013-12-01 | 13:37:20,36 | 80.239.229.205 | 1,79 | 4731056 | 0,044341087 |
| 22 | 2013-12-01 | 13:38:11,28 | 74.125.218.53 | 35,128 | 5662350 | 0,11991787 |
| 23 | 2013-12-01 | 13:38:51,34 | 80.239.229.206 | 1,81 | 63069979 | 0,009178162 |

Table 1. Local Machine in Espoo

| Sl.No | Date | Time Stamp | IP | ping to Content Sever | File Size | Time to get first byte |
|---|---|---|---|---|---|---|
| 1 | 2013-12-01 | 14:07:50,07 | 109.105.109.206 | 7,403 | 12286640 | 0,041049004 |
| 2 | 2013-12-01 | 14:08:21,03 | 74.125.163.24 | 34,952 | 10058880 | 0,13194418 |
| 3 | 2013-12-01 | 14:08:54,19 | 173.194.48.41 | 30,758 | 542134 | 0,087753057 |
| 4 | 2013-12-01 | 14:09:27,56 | 74.125.163.77 | 34,836 | 11295343 | 0,074902058 |
| 5 | 2013-12-01 | 14:09:48,76 | 109.105.109.207 | 7,415 | 22390541 | 0,053025961 |
| 6 | 2013-12-01 | 14:10:18,00 | 173.194.48.74 | 31,018 | 5582937 | 0,101792812 |
| 7 | 2013-12-01 | 14:10:38,96 | 109.105.109.207 | 7,406 | 26941828 | 0,018375158 |
| 8 | 2013-12-01 | 14:10:58,82 | 109.105.109.205 | 7,219 | 20734550 | 0,050659895 |
| 9 | 2013-12-01 | 14:11:27,91 | 74.125.163.70 | 30,644 | 22472544 | 0,164355993 |
| 10 | 2013-12-01 | 14:11:49,60 | 109.105.109.204 | 7,402 | 15855766 | 0,589711905 |
| 11 | 2013-12-01 | 14:12:11,13 | 109.105.109.205 | 7,231 | 14697109 | 0,061920881 |
| 12 | 2013-12-01 | 14:12:40,16 | 173.194.48.42 | 30,513 | 31437962 | 0,066213131 |
| 13 | 2013-12-01 | 14:13:10,93 | 74.125.163.54 | 35,158 | 23823489 | 0,096378088 |
| 14 | 2013-12-01 | 14:13:41,55 | 74.125.163.73 | 30,739 | 8174529 | 0,099963903 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 15 | 2013-12-01 | 14:14:11,78 | 74.125.163.9 | 30,765 | 399742 | 0,08831811 |
| 16 | 2013-12-01 | 14:14:41,86 | 74.125.163.115 | 34,873 | 3343766 | 0,111600876 |
| 17 | 2013-12-01 | 14:15:11,93 | 74.125.163.115 | 34,973 | 5671998 | 0,100214958 |
| 18 | 2013-12-01 | 14:15:32,78 | 109.105.109.205 | 7,221 | 3630300 | 0,01853013 |
| 19 | 2013-12-01 | 14:15:52,21 | 109.105.109.204 | 7,404 | 2344045 | 0,054746151 |
| 20 | 2013-12-01 | 14:16:21,36 | 74.125.163.113 | 34,717 | 19682496 | 0,095088959 |
| 21 | 2013-12-01 | 14:16:43,96 | 109.105.109.205 | 7,23 | 4731056 | 0,102582932 |
| 22 | 2013-12-01 | 14:17:13,02 | 173.194.48.10 | 34,946 | 5662350 | 1,054959059 |
| 23 | 2013-12-01 | 14:17:35,17 | 109.105.109.206 | 7,408 | 63069979 | 0,04776907 |

Table 2. Machine in Otaniemi

| Sl.No | Date | Time Stamp | IP | ping to Content Sever | File Size | Time to get first byte |
|---|---|---|---|---|---|---|
| 1 | 2013-12-01 | 17:22:36,08 | 74.125.12.74 | 69,861 | 12286640 | 0,165047169 |
| 2 | 2013-12-01 | 17:22:57,58 | 74.125.96.152 | 1,605 | 10058880 | 0,041352987 |
| 3 | 2013-12-01 | 17:23:21,08 | 74.125.96.137 | 1,7 | 542134 | 0,02485919 |
| 4 | 2013-12-01 | 17:23:40,61 | 173.194.49.45 | 2,816 | 11295343 | 0,036267996 |
| 5 | 2013-12-01 | 17:24:04,01 | 173.194.49.44 | 2,095 | 22390541 | 0,008991957 |
| 6 | 2013-12-01 | 17:24:25,46 | 173.194.49.10 | 2,328 | 5582937 | 0,015633821 |
| 7 | 2013-12-01 | 17:24:45,63 | 173.194.59.217 | 69,701 | 26941828 | 0,143646002 |
| 8 | 2013-12-01 | 17:25:10,61 | 74.125.96.138 | 1,692 | 20734550 | 0,008189917 |
| 9 | 2013-12-01 | 17:25:30,76 | 173.194.59.198 | 68,989 | 22472544 | 0,174438 |
| 10 | 2013-12-01 | 17:25:52,71 | 74.125.12.70 | 68,757 | 15855766 | 0,177608967 |
| 11 | 2013-12-01 | 17:26:13,89 | 74.125.96.170 | 1,722 | 14697109 | 0,00952816 |
| 12 | 2013-12-01 | 17:26:33,45 | 74.125.96.138 | 1,778 | 31437962 | 0,065854073 |
| 13 | 2013-12-01 | 17:26:53,70 | 74.125.96.182 | 1,685 | 23823489 | 0,070862055 |
| 14 | 2013-12-01 | 17:27:14,18 | 173.194.59.201 | 69,787 | 8174529 | 0,18502593 |
| 15 | 2013-12-01 | 17:27:37,41 | 74.125.96.169 | 1,85 | 399742 | 0,02572608 |
| 16 | 2013-12-01 | 17:27:57,04 | 173.194.59.243 | 70,665 | 3343766 | 0,164361954 |
| 17 | 2013-12-01 | 17:28:27,88 | 74.125.209.208 | 303,384 | 5671998 | 0,985977888 |
| 18 | 2013-12-01 | 17:28:57,02 | 74.125.96.170 | 1,792 | 3630300 | 0,009833813 |
| 19 | 2013-12-01 | 17:29:16,58 | 74.125.96.134 | 2,163 | 2344045 | 0,042320013 |
| 20 | 2013-12-01 | 17:29:37,15 | 173.194.59.241 | 70,111 | 19682496 | 0,168532133 |
| 21 | 2013-12-01 | 17:29:58,58 | 173.194.22.42 | 11,08 | 4731056 | 0,057880878 |
| 22 | 2013-12-01 | 17:30:24,43 | 74.125.96.138 | 1,896 | 5662350 | 3,286892891 |
| 23 | 2013-12-01 | 17:30:50,42 | 74.125.96.168 | 2,003 | 63069979 | 0,010391951 |

Table 3. Machine in Ireland

| Sl.No | Date | Time Stamp | IP | ping to Content Sever | File Size | Time to get first byte |
|---|---|---|---|---|---|---|
| 1 | 2013-12-01 | 17:22:36,08 | 74.125.12.74 | 69,861 | 12286640 | 0,165047169 |
| 2 | 2013-12-01 | 17:22:57,58 | 74.125.96.152 | 1,605 | 10058880 | 0,041352987 |
| 3 | 2013-12-01 | 17:23:21,08 | 74.125.96.137 | 1,7 | 542134 | 0,02485919 |
| 4 | 2013-12-01 | 17:23:40,61 | 173.194.49.45 | 2,816 | 11295343 | 0,036267996 |
| 5 | 2013-12-01 | 17:24:04,01 | 173.194.49.44 | 2,095 | 22390541 | 0,008991957 |
| 6 | 2013-12-01 | 17:24:25,46 | 173.194.49.10 | 2,328 | 5582937 | 0,015633821 |
| 7 | 2013-12-01 | 17:24:45,63 | 173.194.59.217 | 69,701 | 26941828 | 0,143646002 |
| 8 | 2013-12-01 | 17:25:10,61 | 74.125.96.138 | 1,692 | 20734550 | 0,008189917 |
| 9 | 2013-12-01 | 17:25:30,76 | 173.194.59.198 | 68,989 | 22472544 | 0,174438 |
| 10 | 2013-12-01 | 17:25:52,71 | 74.125.12.70 | 68,757 | 15855766 | 0,177608967 |
| 11 | 2013-12-01 | 17:26:13,89 | 74.125.96.170 | 1,722 | 14697109 | 0,00952816 |
| 12 | 2013-12-01 | 17:26:33,45 | 74.125.96.138 | 1,778 | 31437962 | 0,065854073 |
| 13 | 2013-12-01 | 17:26:53,70 | 74.125.96.182 | 1,685 | 23823489 | 0,070862055 |
| 14 | 2013-12-01 | 17:27:14,18 | 173.194.59.201 | 69,787 | 8174529 | 0,18502593 |
| 15 | 2013-12-01 | 17:27:37,41 | 74.125.96.169 | 1,85 | 399742 | 0,02572608 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 16 | 2013-12-01 | 17:27:57,04 | 173.194.59.243 | 70,665 | 3343766 | 0,164361954 |
| 17 | 2013-12-01 | 17:28:27,88 | 74.125.209.208 | 303,384 | 5671998 | 0,985977888 |
| 18 | 2013-12-01 | 17:28:57,02 | 74.125.96.170 | 1,792 | 3630300 | 0,009833813 |
| 19 | 2013-12-01 | 17:29:16,58 | 74.125.96.134 | 2,163 | 2344045 | 0,042320013 |
| 20 | 2013-12-01 | 17:29:37,15 | 173.194.59.241 | 70,111 | 19682496 | 0,168532133 |
| 21 | 2013-12-01 | 17:29:58,58 | 173.194.22.42 | 11,08 | 4731056 | 0,057880878 |
| 22 | 2013-12-01 | 17:30:24,43 | 74.125.96.138 | 1,896 | 5662350 | 3,286892891 |
| 23 | 2013-12-01 | 17:30:50,42 | 74.125.96.168 | 2,003 | 63069979 | 0,010391951 |

Table 4. Machine in California

| **Sl.No** | **Date** | **Time Stamp** | **IP** | **ping to Content Sever** | **File Size** | **Time to get first byte** |
|---|---|---|---|---|---|---|
| 1 | 2013-12-01 | 17:22:36,08 | 74.125.12.74 | 69,861 | 12286640 | 0,165047169 |
| 2 | 2013-12-01 | 17:22:57,58 | 74.125.96.152 | 1,605 | 10058880 | 0,041352987 |
| 3 | 2013-12-01 | 17:23:21,08 | 74.125.96.137 | 1,7 | 542134 | 0,02485919 |
| 4 | 2013-12-01 | 17:23:40,61 | 173.194.49.45 | 2,816 | 11295343 | 0,036267996 |
| 5 | 2013-12-01 | 17:24:04,01 | 173.194.49.44 | 2,095 | 22390541 | 0,008991957 |
| 6 | 2013-12-01 | 17:24:25,46 | 173.194.49.10 | 2,328 | 5582937 | 0,015633821 |
| 7 | 2013-12-01 | 17:24:45,63 | 173.194.59.217 | 69,701 | 26941828 | 0,143646002 |
| 8 | 2013-12-01 | 17:25:10,61 | 74.125.96.138 | 1,692 | 20734550 | 0,008189917 |
| 9 | 2013-12-01 | 17:25:30,76 | 173.194.59.198 | 68,989 | 22472544 | 0,174438 |
| 10 | 2013-12-01 | 17:25:52,71 | 74.125.12.70 | 68,757 | 15855766 | 0,177608967 |
| 11 | 2013-12-01 | 17:26:13,89 | 74.125.96.170 | 1,722 | 14697109 | 0,00952816 |
| 12 | 2013-12-01 | 17:26:33,45 | 74.125.96.138 | 1,778 | 31437962 | 0,065854073 |
| 13 | 2013-12-01 | 17:26:53,70 | 74.125.96.182 | 1,685 | 23823489 | 0,070862055 |
| 14 | 2013-12-01 | 17:27:14,18 | 173.194.59.201 | 69,787 | 8174529 | 0,18502593 |
| 15 | 2013-12-01 | 17:27:37,41 | 74.125.96.169 | 1,85 | 399742 | 0,02572608 |
| 16 | 2013-12-01 | 17:27:57,04 | 173.194.59.243 | 70,665 | 3343766 | 0,164361954 |
| 17 | 2013-12-01 | 17:28:27,88 | 74.125.209.208 | 303,384 | 5671998 | 0,985977888 |
| 18 | 2013-12-01 | 17:28:57,02 | 74.125.96.170 | 1,792 | 3630300 | 0,009833813 |
| 19 | 2013-12-01 | 17:29:16,58 | 74.125.96.134 | 2,163 | 2344045 | 0,042320013 |
| 20 | 2013-12-01 | 17:29:37,15 | 173.194.59.241 | 70,111 | 19682496 | 0,168532133 |
| 21 | 2013-12-01 | 17:29:58,58 | 173.194.22.42 | 11,08 | 4731056 | 0,057880878 |
| 22 | 2013-12-01 | 17:30:24,43 | 74.125.96.138 | 1,896 | 5662350 | 3,286892891 |
| 23 | 2013-12-01 | 17:30:50,42 | 74.125.96.168 | 2,003 | 63069979 | 0,010391951 |

Table 5. Machine in Singapore

## Conclusion

As a conclusion we can say that Youtube CDN uses many different strategies to redirect traffic between the content servers and even though some locations seem to favor some content servers it doesn’t always mean that this content server is the closest one. A good example of this is the home machine in Espoo and Kosh in Otaniemi which are really close together and still seem to use different content servers in Youtube.

While experimenting at the same time on machines which are located close together like Kosh and Espoo home machine we noticed that Youtube CDN can redirect them to different content servers even though both are closer to one than the other. This indicates that load balancing is also used in the Youtube CDN to balance the traffic between content servers.

**References:**

[1] Vijay Kumar Adhikari, Sourabh Jain, and Zhi-Li Zhang. Youtube traffic dynamics and its interplay with a tier-1 isp: an isp perspective. In Proceedings of the 10th ACM SIGCOMM conference on Internet measurement, IMC '10, pages 431{443, New York, NY, USA, 2010.ACM.

[2] Ruben Torres, Alessandro Finamore, Jin Ryong Kim, Marco Mellia, Maurizio M. Munaf, and Sanjay G. Rao. Dissecting video server selection strategies in the youtube cdn. In ICDCS, pages 248{257. IEEE Computer Society, 2011

[3] Article about CDN in whatismyipadress.com

[4] Parikshit Juluri, Louis Plissonneau, and Deep Medhi. Pytomo: A tool for analyzing playback quality of youtube videos. In ke Arvidsson, Gustavo de Veciana, Steven H. Low, Charles R.Kalmanek, and Deep Medhi, editors, International Teletrac Congress, pages 304{305. IEEE, 2011.

[5] List of  Most viewed Youtube links