CloudTracker: Accelerating Internet Content Distribution by Bridging Cloud Servers and Peer Swarms

Zhenhua Li (Ph.D. student), Yan Huang, Gang Liu, and Yafei Dai

Peking University and Tencent Research

Abstract

Large-scale content distribution on today's Internet primarily operates in two modes: 1) **cloud-based** content distribution, which relies on "cloud servers" (e.g., servers in data centers or web hosting farms) that are located closer to the "core" Internet; 2) **peer-to-peer** content distribution that relies on a large number of user controlled end hosts that are located on the "edge" of the Internet and collectively form a peer swarm. Both modes have inherent advantages and disadvantages in terms of manageability, costs, scalability, and so forth. To bridge the two modes, especially to leverage the inherent scalability of peer swarms while circumventing its key limitations, in this paper we propose and implement a **cloud tracking service** named "CloudTracker" to accelerate peer-to-peer content distribution. CloudTracker continuously tracks downloadable files in cloud servers all over the Internet, indexes these files, and schedules peer swarms to retrieve them, so as to accelerate the data transfer of peers as well as to facilitate load-balancing among cloud servers. By using only 53 commodity servers, CloudTracker schedules around 5 million peers to retrieve data in petabytes from over 1.5 million cloud servers per day, and the average download rate of a peer is enhanced from 57 KBps to 158 KBps.

Motivation

Content distribution on today's Internet operates in two modes

1. Cloud-based content distribution

- Rely on "cloud servers" such as data centers, conventional web servers, and so forth
- Locate closer to the "core" Internet
- Allow content providers better control
- Leverage the economy of scale
- Require significant upstart infrastructure costs
- Continuing operation costs: storage and bandwidth

2. P2P (Peer-to-Peer) content distribution

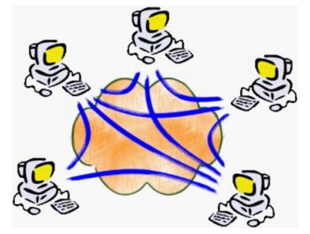

- Rely on a large number of end hosts such as home PCs, laptops, or hand-held mobile devices
- Locate on the "edge" of the Internet
- Collectively form a *peer data swarm*
- Little cost, highly scalable
- Limited capacity per node
- High dynamics (frequently join or leave)
- Working efficiency can be poor and unpredictable

Both modes have their inherent advantages and disadvantages

	Cloud	P2P
Manageability	Good	Poor
Costs	Expensive	Low
Scalability	Low	High
Stability	High	Poor

CloudTracker

CloudTracker: **bridge** the **cloud-based** and **P2P** content distribution

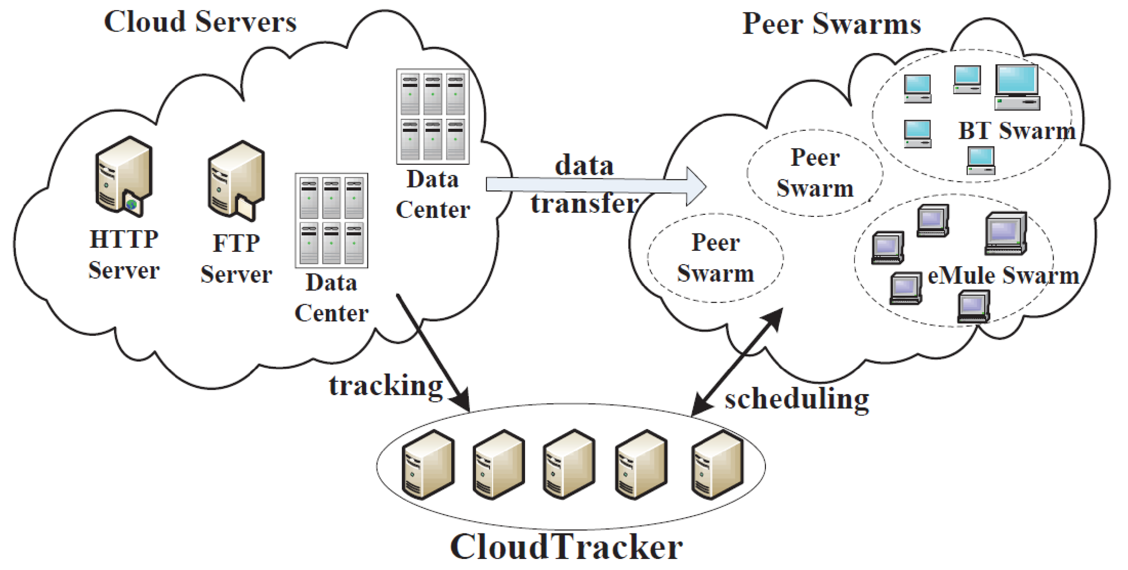

1. Continuously tracks and indexes *downloadable files in cloud servers*
- Allow end users to search and quickly find contents
- Enable peer nodes that are interested in the same content to find each other and form a peer swarm

2. Dynamically directs peer swarms to appropriate cloud servers
- "Hungry swarms" do not have adequate download bandwidth
- Directed to cloud servers to get extra data support
- Thus achieve more efficient content distribution
- Address key limitations of P2P (instability and unpredictability)

3. Facilitate load-balancing among cloud servers
- By directing peer swarms to appropriate cloud servers
- Avoid overloading a specific server

Light-weight Implementation
- CloudTracker does not upload any file to any peer by itself
- We have implemented CloudTracker using **only 53 commodity servers** and it serves millions of peers well every day

System Architecture
- File Index DB
  - Indexing and mapping contents
- Cloud Crawler
  - Discovering new contents
- Validation Servers
  - Validate contents/indexes
- Data Scheduler
  - Direct peers to cloud servers

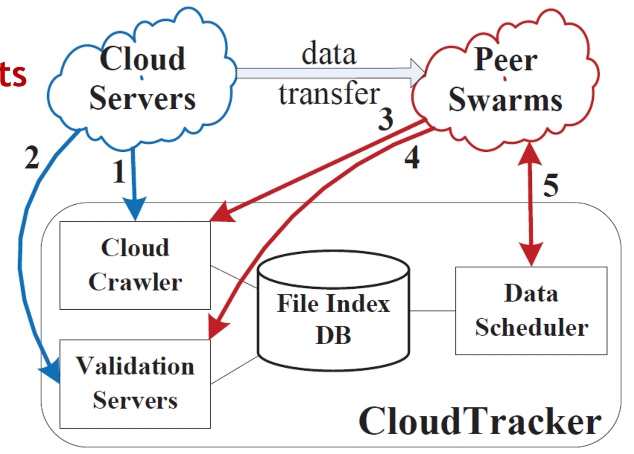

Building block	Number of servers	Memory	Storage	Bandwidth
File Index DB	15	2 GB / 8 GB / 16 GB	120 GB	1 Gbps
Cloud Crawler	12	2 GB / 8 GB	150 GB	1 Gbps
Validation Servers	12	2 GB / 8 GB	150 GB	1 Gbps
Data Scheduler	14	2 GB / 8 GB	150 GB	1 Gbps

Data Scheduling Strategy

Strives for *a proper tradeoff between the user experience and cloud bandwidth burden*

- *On one hand, we want to enhance* each peer's download rate to as high as possible by fully utilizing cloud bandwidth
  - so that each peer can have his best user experience
  - and CloudTracker can attract more and more users
- On the other hand, the extra bandwidth burden posed on a cloud server directed by CloudTracker should be within a certain limit
  - otherwise, the normal functionality of the cloud server may be interfered
- Moreover, cloud providers may be reluctant to support CloudTracker if they cannot obtain extra benefit from extra bandwidth expense

Data Scheduling Strategy utilizes three mechanisms:

1. *Discriminative acceleration of peer swarms*
- Every user hopes to achieve his best user experience
- We discover a user usually has his basic expectation for the download rate: d_{basic} = 30 KBps
- Provide discriminative acceleration of peer swarms
  - peer swarms are classified into three categories
  - according to download rates and data supply-demand condition
  - *a) hungry swarms, b) potentially hungry swarms, and c) high-demand swarms* (d_{high} = 100 KBps)
- *Different* peer swarms correspond to different acceleration methods → each user has a download rate at least above basic expectation

2. *Balanced utilization of cloud servers*
- Measurements indicate the original bandwidth utilization of a cloud server rarely exceeds 60%
- Thus, *extra bandwidth utilization (EBU)* directed by CloudTracker had better be controlled within 40% (= 1 – 60%).
- CloudTracker periodically collects the *EBU of each cloud server. If* the *EBU of a server S exceeds 40%, CloudTracker* will notify a part of the peers served by S to stop their data download from S.

3. *Benefit collaboration with cloud providers*
- Two most important benefit metrics of a cloud provider:
  a) *page view (PV)* and b) *paid-to-click (PTC)*
- *PV*: if a cloud provider uploads a file to a peer swarm by the direction of CloudTracker, its *PV will be increased* in the web search engine "Soso" (the 4th biggest web search engine in China)
- *PTC*: If a cloud provider has contributed a lot to CloudTracker, it will be rewarded by a nontrivial share of our *PTC revenue*

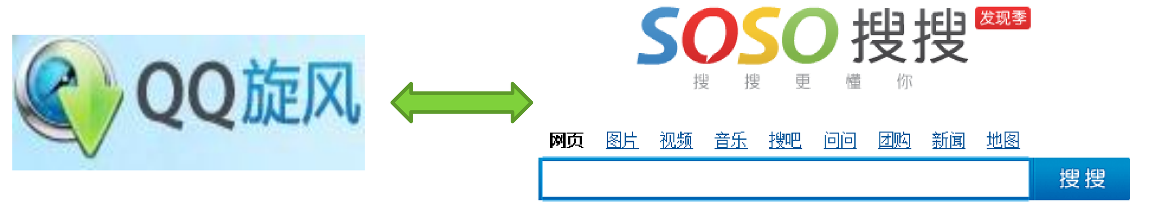

http://xf.qq.com http://www.soso.com

Performance Evaluation

System Performance Summary
- CloudTracker has gained over 120 million accumulated users
- Schedules 5 million peers to retrieve data in petabytes from over 1.5 million cloud servers per day
- Peers' average download rate: 57 KBps → 158 KBps (177% increment)
- 45% of the download rate is obtained from cloud servers
- Extra bandwidth utilization of cloud servers controlled within 40%

Metrics
- *Acceleration effect on peer swarms*
- *Bandwidth contribution of cloud servers*
- *Extra bandwidth utilization of cloud servers*

Download rate distribution

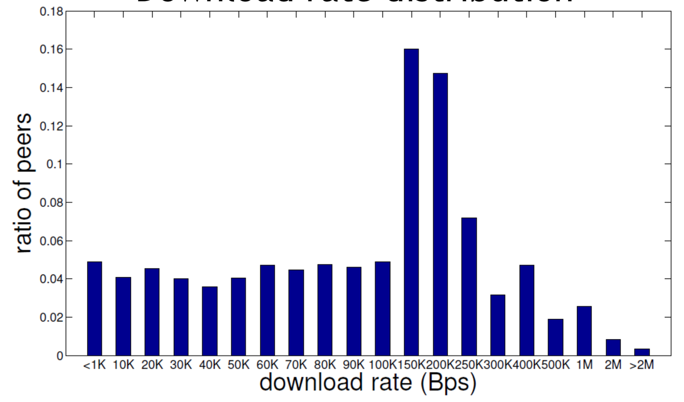

Acceleration effect brought by CloudTracker

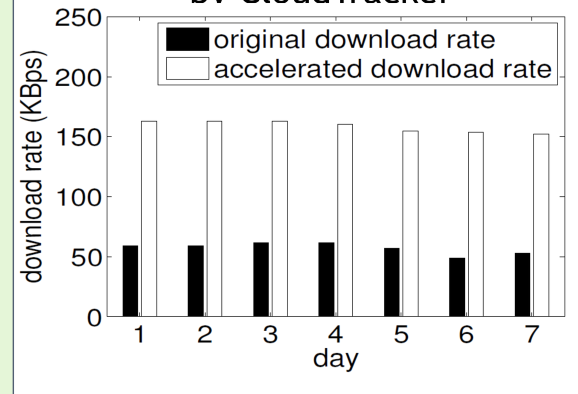

Acceleration effect as the system scales

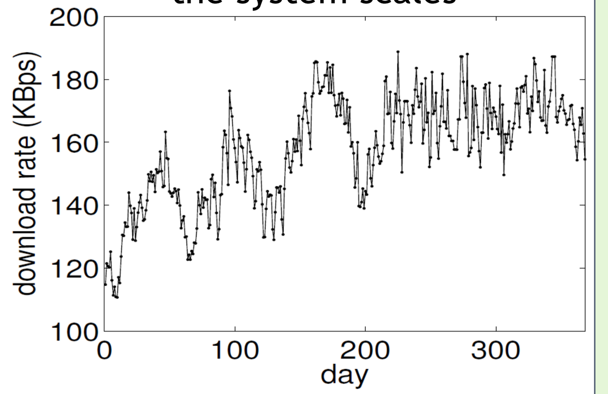

Bandwidth Contribution of Cloud Servers

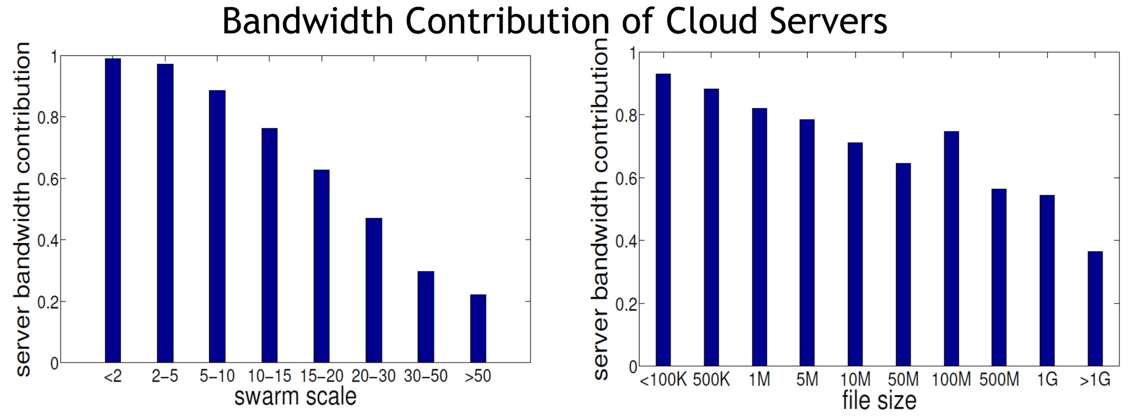

Extra Bandwidth Utilization of Cloud Servers

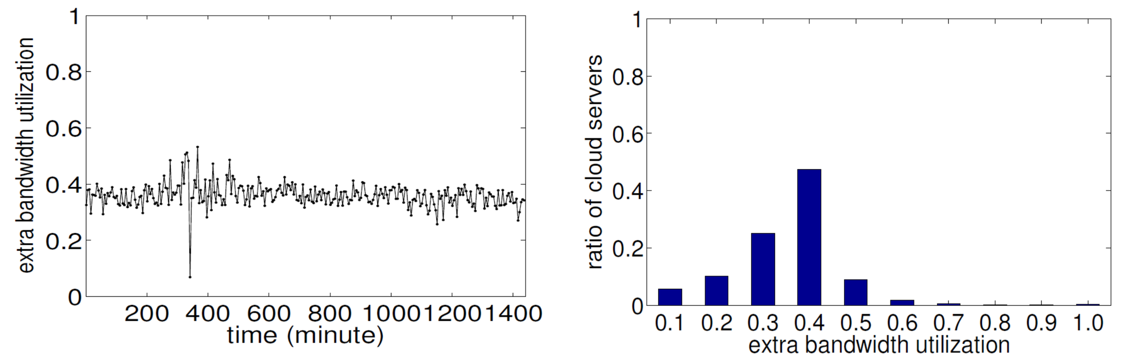

Future Work

Dynamically tuning system parameters
- Current system parameters are statically configured based on comprehensive measurements, although work fine, may not well adapt to changes of network environments.
- How to design proper mechanisms to automatically collect/analyze measurements and thus dynamically tune system parameters to match the new situations.

More specific and powerful crawler
- Filtering web search engine results at present
- Slow, weak, and incomplete
- Specific crawler focusing on cloud tracking

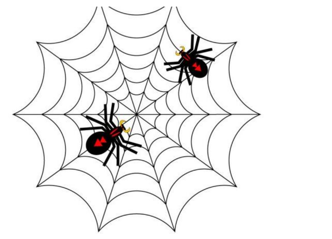

Related Work

BitTorrent Trackers
- Only tracks the peers in one swarm
- Does not care about the data transfer status of peers
- Only supports a single protocol: bittorrent.

P2SP Systems

- C. Wu, B. Li, and S. Zhao. "On Dynamic Server Provisioning in Multichannel P2P Live Streaming," To appear in TON.
- Y. Sun, F. Liu, et al. "Fs2you: Peer-assisted semi-persistent online storage at a large scale," In IEEE INFOCOM, 2009.
- F. Liu, S. Shen, B. Li, B. Li, H. Yin, and S. Li. "Novasky: Cinematic-Quality VoD in a P2P Storage Cloud," In IEEE INFOCOM, 2011.
- H. Yin, X. Liu, T. Zhan, V. Sekar, F. Qiu, C. Lin, H. Zhang, and B. Li. "Design and deployment of a hybrid CDN-P2P system for live video streaming: experiences with LiveSky," In ACM Multimedia, 2009.

Thunder (larger and older): http://www.xunlei.com.

Web browsers using cloud tracking

Contact

Peking University (Beijing, Shenzhen) is one of the oldest and best comprehensive universities in China, which is founded in 1898.

Tencent (Shenzhen, Shanghai, Beijing, Chengdu) is one of the biggest Internet companies in China, whose business involves IM, Web portal, (Micro)Blog, Video Streaming, etc.

Zhenhua Li, *Peking University* lzh@net.pku.edu.cn
Yan Huang, *Tencent Research* galehuang@tencent.com
Gang Liu, *Tencent Research* sinbadliu@tencent.com
Yafei Dai, *Peking University* dyf@net.pku.edu.cn