26th International Teletraffic Congress (ITC 26)
September 9th-11th, 2014, Karlskrona, Sweden

A Socially-aware ISP-friendly Mechanism for Efficient Content Delivery

George Darzanos, Ioanna Papafili, George D. Stamoulis
Athens University of Economics and Business
Network Economic and Services Laboratory (http://nes.aueb.gr)

Context

- **Video sharing** has been an increasingly popular application in Online Social Networks (OSNs).
- Most of the content shared via OSNs, including videos, is **User Generated Content** (UGC).
- UGC popularity is **long-tailed.**
- Drawbacks of main caching approaches (e.g. CDN, Web caches):
  - Do not address long-tailed content
  - Face scalability problems – possible QoE degradation

Motivation and Objectives

- Valuable **information from OSNs** can be extracted and used for effective content placement strategies.

- **We aim** to design a **scalable Content Distribution system** in order to:
  - **Improve QoE** of OSNs users by reducing latency and eliminating stalling events.
  - **Reduce expenses of**
    - ✓ **ISPs**: reduce Inter-domain traffic, and thus inter-connection (transit) charges too
    - ✓ **OSNs and CDNs**: reduce the workload of the video server, and thus the related costs (bandwidth, storage, energy) too.

Use Cases Addressed (I)

Video Viewing Cases: video hosted in Facebook Video Server

1. Upload video to Facebook Video Server
2. Download video from Facebook Video Server

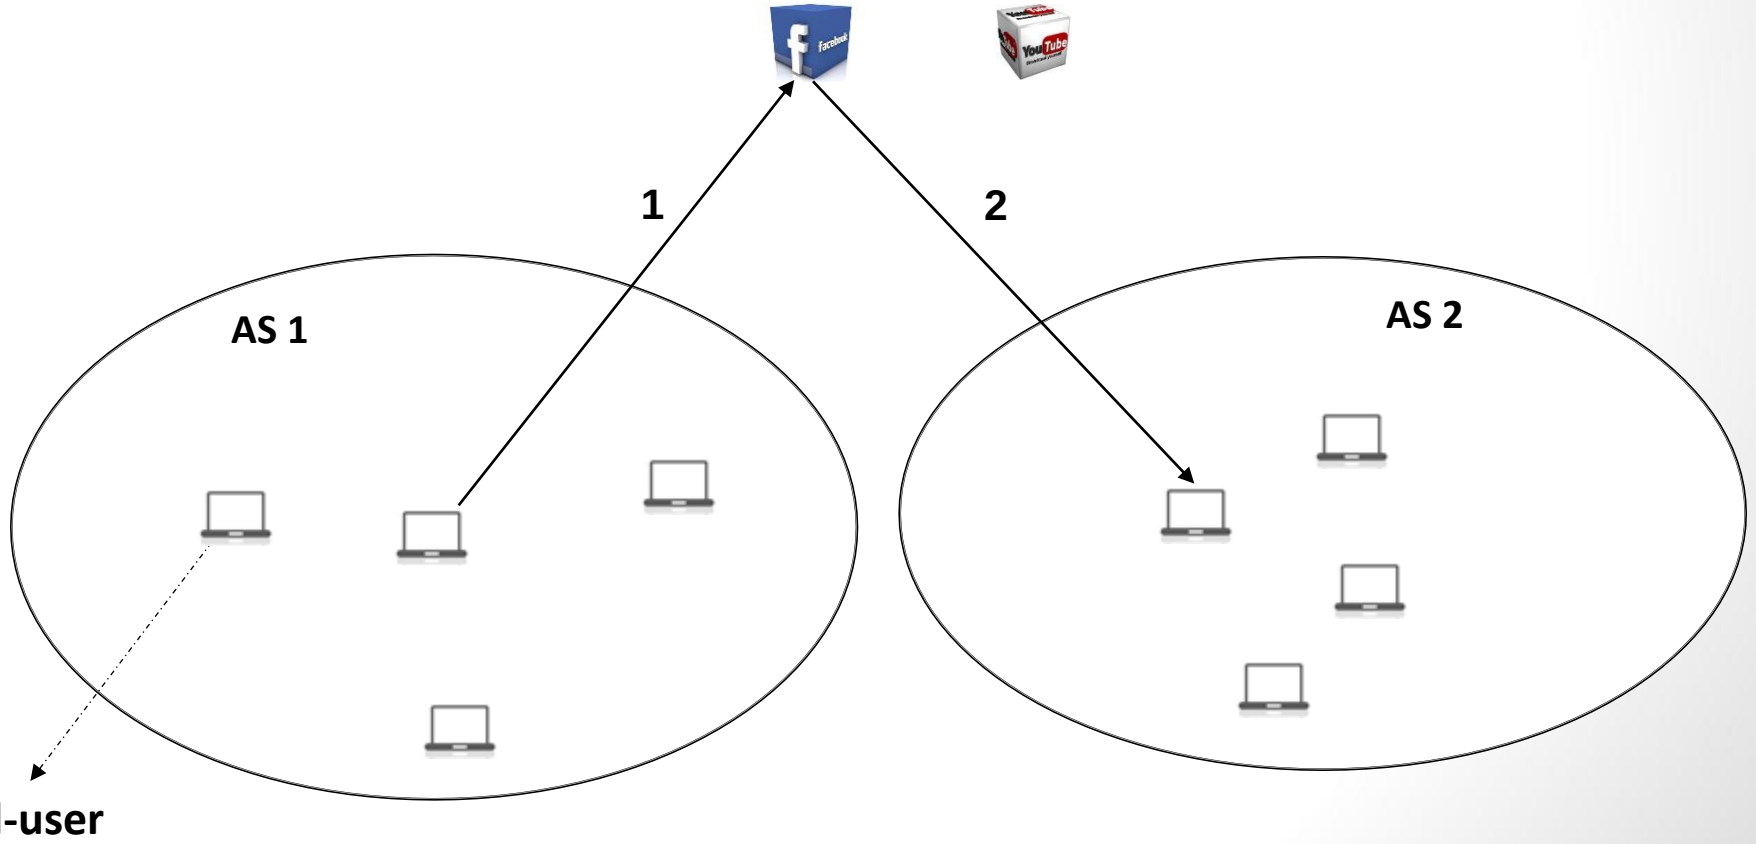

Use Cases Addressed (II)

Video Viewing Cases: video hosted in YouTube video server

1. Copy link from YouTube website
2. Share link on Facebook Wall
3. Click on link of YouTube video on Facebook Wall
4. Redirection to YouTube Video Server
5. Download video from YouTube Video Server

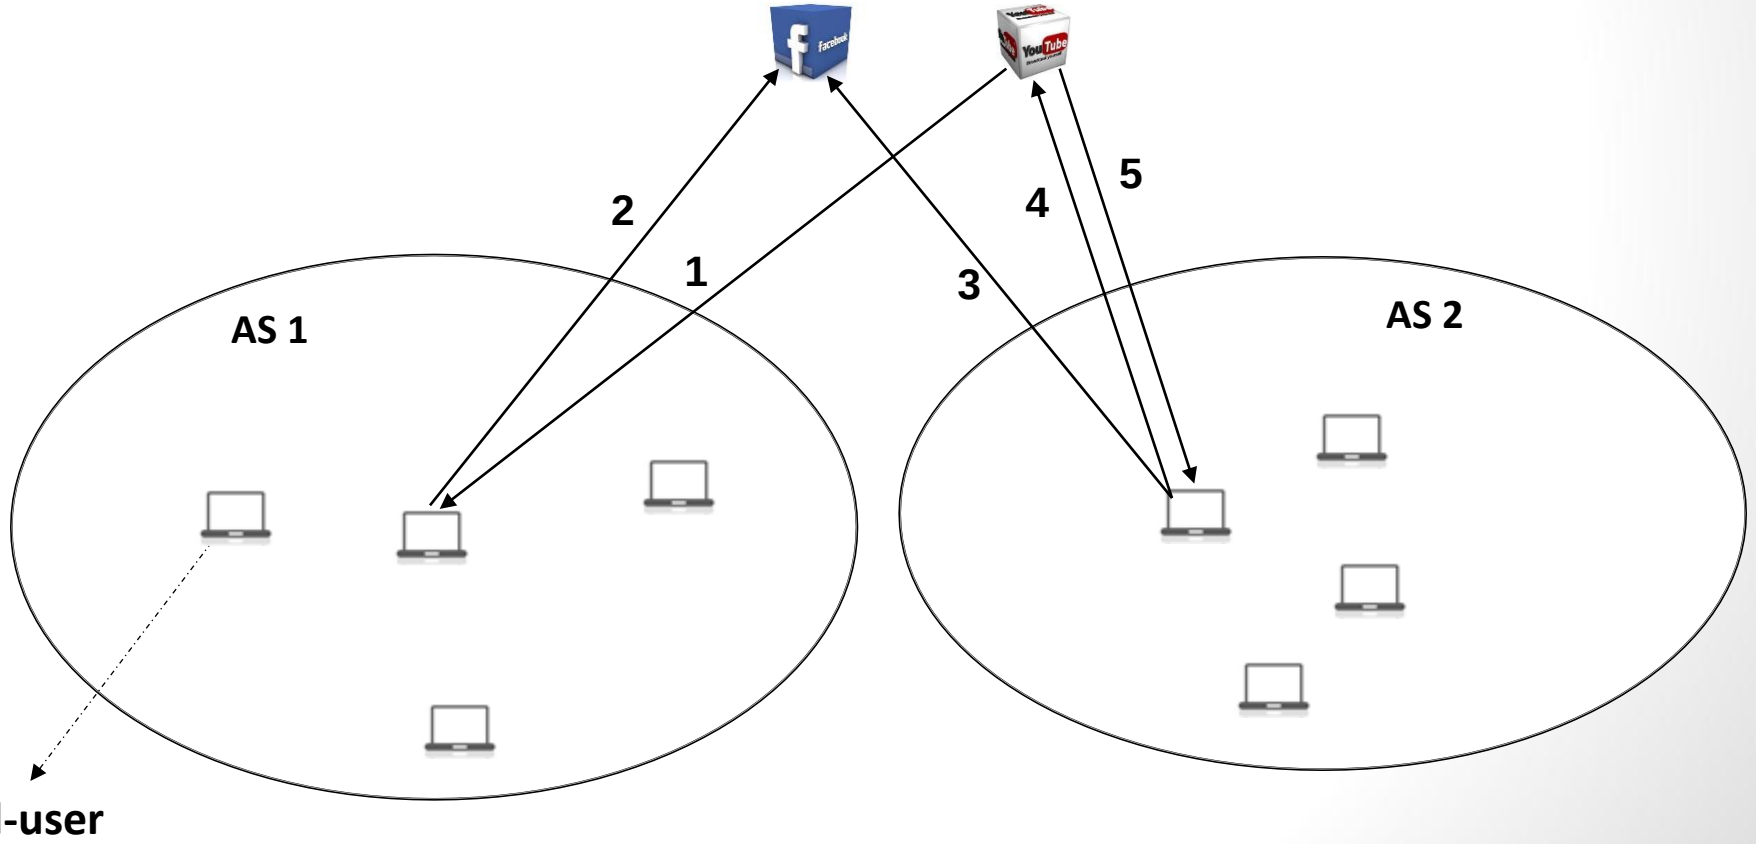

SEConD Overview

- A **S**ocially Aware ISP-friendly Mechanism for **E**fficient **Con**tent **D**elivery exploiting:
  - Social relationships
  - Similarities of interests with respect to content category
  - Locality of demand for OSN-published content

- Constituent elements
  - 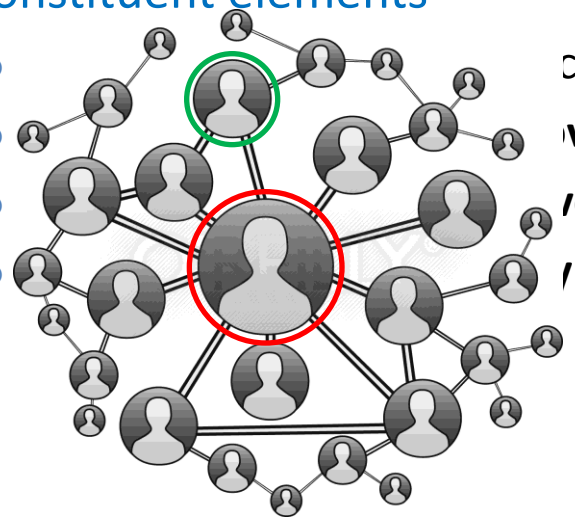
  - cated in every AS: SEC
  - **verlays:** trigger pull-b
  - **erlays**: enables intra-
  - **:** end-users and SPS

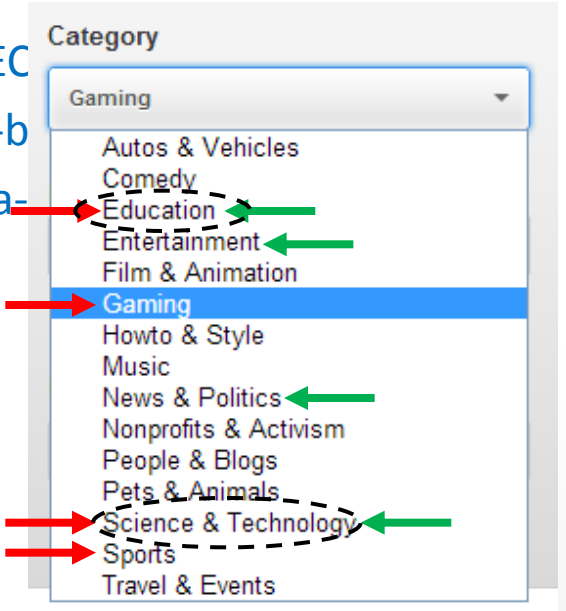

Social Proxy Server Functionality

- **Monitors** user interactions related to videos.
- **Forms/updates** the messaging overlays, based on monitoring information.
- **Pushes video prefixes** (first chunk) to users, in order to server requests produced my messaging overlays.
- **Adds users** requesting a video to the local content-specific P2P swarm **(P2P Tracker).**
- **Assists swarms** with inadequate upload bandwidth, acting as **super-peer.**
- **Caches** video prefixes and videos following requests.

Messaging overlay construction

Each messaging overlay, comprises the **source user** and his **potential viewers** for videos of a **specific category of interest**.

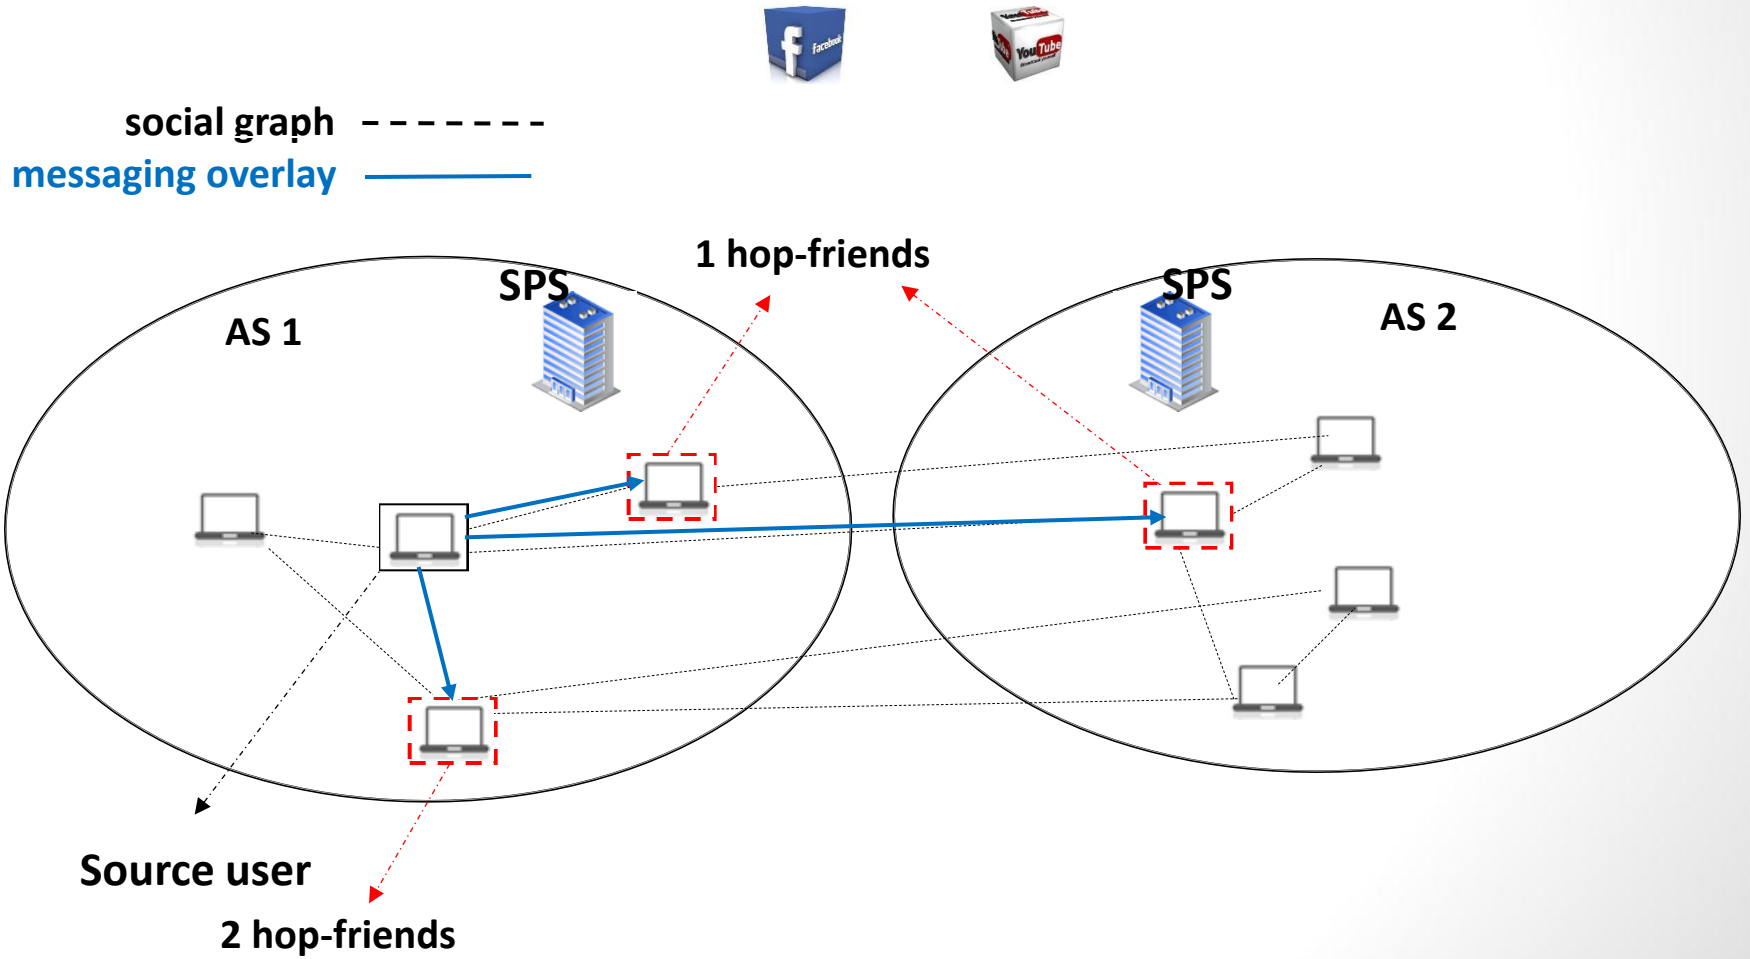

Prefetching Approach

- SPS downloads/stores the video and pushes the **video prefix** from local SPS
- Potential viewers store the prefix in their **cache**

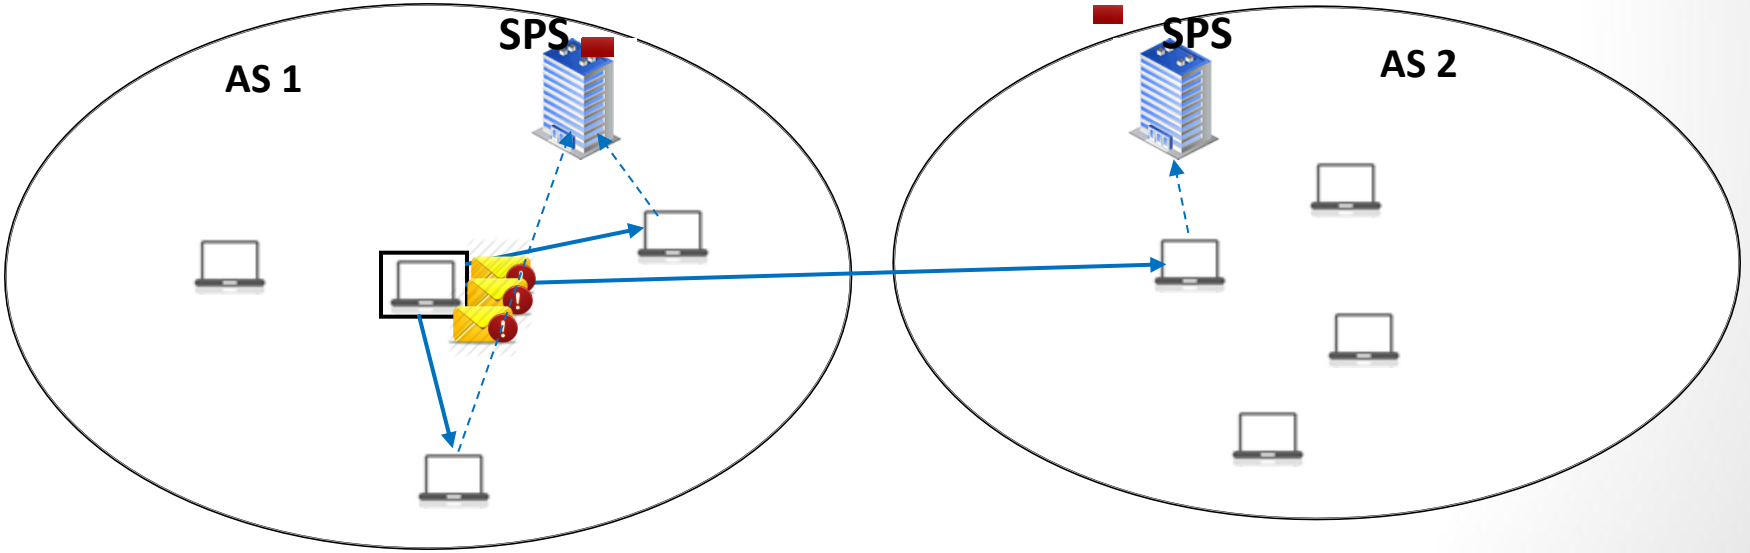

Local P2P and Caching

- **Local Content-based P2P overlays (SPS as Tracker):** perform local video sharing
  - End-users **added by SPS** in content-specific P2P swarms:
    - ✓ as **seeders**, for videos stored in their cache
    - ✓ as **leechers**, when they are watching a video
  - SPS participates as **super-peer,** only when the total upload bandwidth in the swarm is lower than the video-bit-rate,
    - ✓ thus **eliminating stalling events**
- **Caching strategy**
  - **Two level caching** of videos and prefixes: in **SPS** and **Users' Equipment**
  - **Caching Policies** applied when the cache is full

	SPS	User
Prefixes	Replace the oldest in cache	Replace the oldest in cache
Videos	Replace one of the two oldest in cache depending on number of prefix requests pending	Replace the oldest in cache

Evaluation Framework (I)

- **Developed based n observations in the literature about**
  - **Viewers' characteristics**: mainly within 1-2 social hops, viewers categories (watch %), …
  - **Timing of users' activities:** 66% DAU, 20min online daily, …
  - **Videos and related interactions:** interest categories, popularity, …

- **Supply model: Content generation**
  - **19 interest categories**: based on YouTube categories
  - **Pool of videos**: popularity (Power Law), interest category (weighted random)
  - Modeling of daily **video uploading and sharing**
  - Each user **pushes alert messages** for the videos he uploads.

Evaluation Framework (II)

- **Demand model: Video viewing**
  - **Split viewers of each uploader into categories**
    - ✓ **Followers**
    - ✓ **Non-followers**
    - ✓ **Other viewers**
  - Each viewer watches **1-5 videos** randomly selected from his 1-social hop friends
  - Each viewer has **4 interests** and watches videos only on these
  - Videos belonging to user's **top interest or popular** videos are more likely to watch

- **Performance Metrics**
  - **Inter/intra AS traffic**
  - **Contribution of server hosting the video**
  - **Prefetching accuracy (QoE proxy)**
  - **Caching accuracy**
  - **Useless and redundant prefetching**

Evaluation Framework (III)

- **Setup**
  - **Social graph:** imported from crawled dataset
  - Zipf distribution of **3963 users** into **4 ASes**
  - **Users cache size** is fixed to **300 MB**
  - **SPS cache size** proportional to the number of bound users, namely **33 MB per user**
  - **Pool of 9000 videos** with Power law **popularity** distribution.
  - Fixed **video-bit-rate to 330 Kbps**
  - 4 minutes **video length**
  - 30 **simulation cycles** corresponding to 30 days.
  - Slotted system, **with 20-minute slot**

Evaluation Results

Performance of **SEConD** and **SocialTube** [1] compared to **traditional Client-server**

- **~80% prefetching accuracy** for both mechanisms, thus improving users' QoE by avoiding initial stalling events.

- **~87% reduction by SEConD of the total inter-AS traffic** (including alerts and prefetching).
  - SocialTube achieves a reduction of ~18%

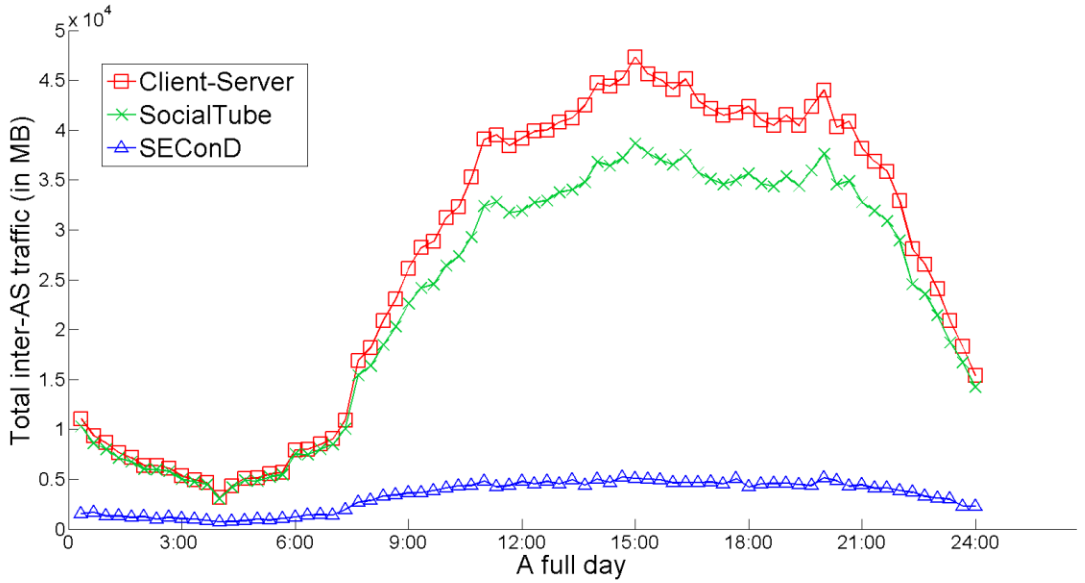

- **~88% reduction by SEConD to the contribution of the origin server** hosting the video
  - SocialTube achieves a reduction of ~45%

[1] H. Li, Y. Liu, G. Huang, H. Liu, C. Wu & C. Wu (2012), 'SocialTube: P2P-assisted video sharing in online social networks.', *in* Albert G. Greenberg & Kazem Sohraby, ed., 'INFOCOM' , IEEE, , pp. 2886-2890 .

Evaluation Results

AS Size vs Origin Server, SPS and P2P Contribution

	AS1	AS2	AS 3	AS4
Number of Users	1925	928	634	476
Proxy Contribution	48%	61%	75%	79%
Proxy Cache Size	63GB	30GB	21GB	15.7GB
Proxy Cache Hit Accuracy (Videos/prefixes)	90/94%	78/86%	68/80%	60/74%
External Server Contribution	5%	14%	24%	31%

decreasing → increasing

decreasing → decreasing

- As the size of the **AS increases**, the **contribution of SPS** decreases.
- **P2P contributes more** in larger ASes.
- The contribution of the **origin server** decreases with the hit accuracy of the **SPS cache**
- **Relatively lower** SPS caching capacity is needed in large ASes than in small ones

Conclusions

- **SEConD**:
  - **Improves users' QoE,** by achieving overall high **prefetching accuracy** and **availability of bandwidth** within swarms.
  - Achieves high **reduction in inter-domain** traffic.
  - **Reduces the contribution** of OSN and CDN **video server.**
  - **Eliminates redundant prefixes,** leading to reduction of traffic congestion within the AS.

- **SEConD is deployable by ISPs, CDNs, OSNs**
  - **ISPs and CDNs** need to either derive social information (e.g. by crawling) or establish agreement with OSN
  - Incentive compatibility of the **OSN**:
    - ✓ performance improvement of its users (indirect)
    - ✓ lower server contribution (direct)

Future work

- Assess the impact of **direct caching** in local SPS, in case of OSN hosted video.
- Develop a **monitoring component** of SPS for users' interactions and refine the algorithm for users' categorization based on the extra information.
- Try and evaluate more **caching policies** and different sizes of Proxy cache
- Consider the potential extension of SEConD, for efficient traffic management in the **backhaul network** too.

Thank you for your attention!

Questions?

Backup Slides

Viewers categorization

- Most **viewers** of an uploader are within two social hops [1]

- On the average, for the viewers of an uploader, the following applies: [1]
  - **25%** of viewers **watched all** videos (mostly 1-hop friends)
  - **33%** of viewers watched **80%** of videos
  - **all** viewers watched at least **20%** of videos

- **We adopt a categorization of viewers**
  - **Followers**: watch almost all the videos of the user (over 80%)
  - **Non-followers**: watch videos less than a high threshold (80%) and more than a low threshold (30%) of the user
  - **Other-viewers**: watch videos less than the low threshold (30%)

[1] Li, Z.; Shen, H.; Wang, H.; Liu, G. & Li, J. (2012), SocialTube: P2P-assisted video sharing in online social networks., *in* Albert G. Greenberg & Kazem Sohraby, ed., 'INFOCOM' , IEEE, , pp. 2886-2890 .

Video viewing in SEConD – Success scenario

- **User** requests a video from SPS
- **SPS** adds user to this video's local P2P swarm
  - If swarm does not exist, creates one
  - Downloads and stores video, if not already

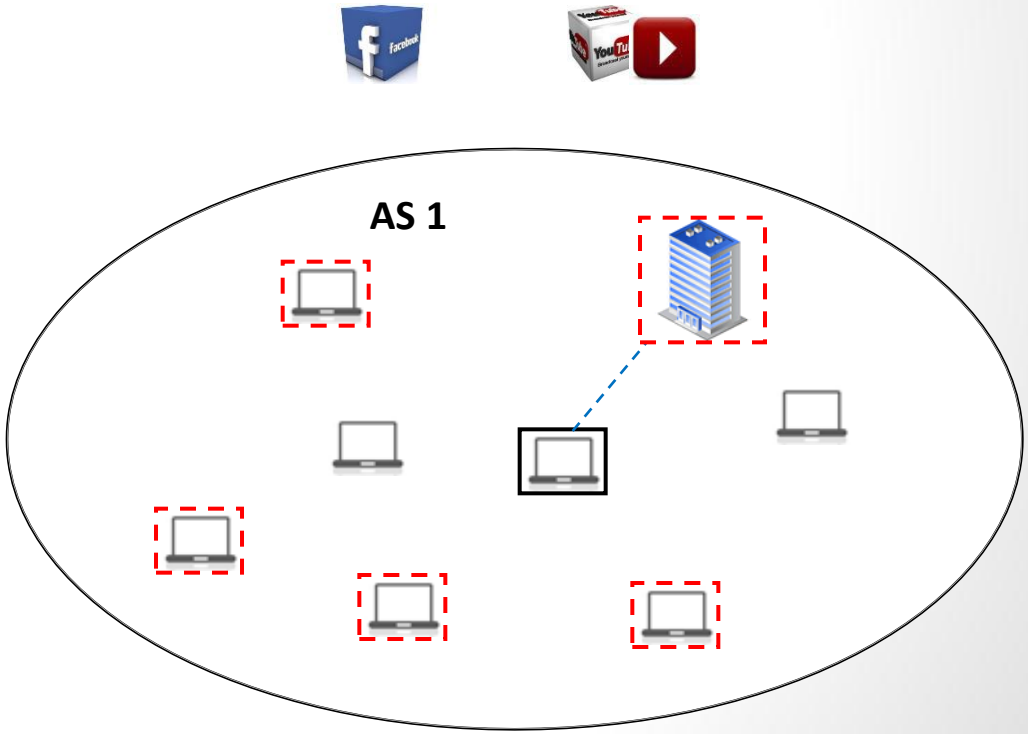

Evaluation Results

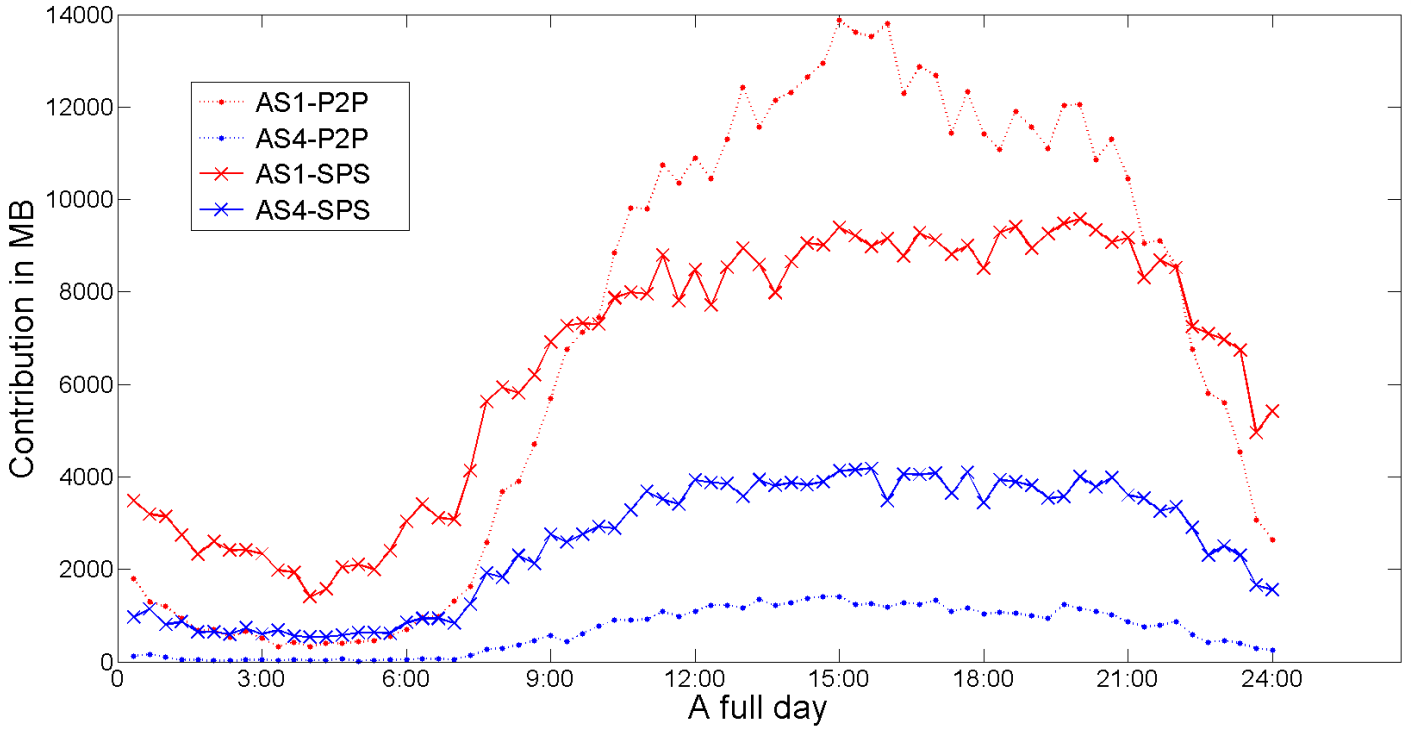

Evaluation Framework (I)

- **Built based upon observations in literature**
  - **Viewers characteristics**: mainly 1-2 social hops, driven by social relationships and interests, viewers categories (watch %), audience are 61% of friends …..
  - **Timing of users' activities**: 66% DAU, 20min online daily, intensity of interactions varies during the day, 140min spent on internet daily….
  - **Videos and related interactions:** interest categories, popularity, 4min average length, 86% of videos are external links….

- **Supply model: Content generation**
  - **19 interest categories**: based on YouTube categories
  - **Pool of videos**: popularity (Power Law), interest category (weighted random)
  - **Video uploading and sharing:**
    - 1 video for every ~20 users uploaded daily
    - Only this day's active users can upload video(s)
    - A user can upload video only the 20min is active
    - 11,8% of videos is re-shared from friends while 88,2% of videos are new uploaded
  - Each user **push alert messages** for the videos he uploads.