Towards QoS-friendly Content Recommendations: An Experimental Framework

Savvas Kastanakis[1], Mateus Nogueira[2], Felipe Assis[2], Pavlos Sermpezis[3], Daniel Menasché[2], Thrasyvoulos Spyropoulos[4]

[1] *University of Crete, Greece,* [2] *UFRJ, Brazil,* [3] *FORTH,Greece,* [4] *EURECOM, France*

1. Background and Motivation

- Internet mobile traffic, especially for online video services (e.g., YouTube, Netflix), increases exponentially
- Mobile networks struggle to attain high QoE while serving *all* content requests, e.g., due to QoS impairments
- A traditional solution: caching → cached contents delivered in high QoS, but small gains due to small caches
- A recent solution: QoS-friendly recommendations, e.g., bias recommendations towards cached videos [1, 2, 3]

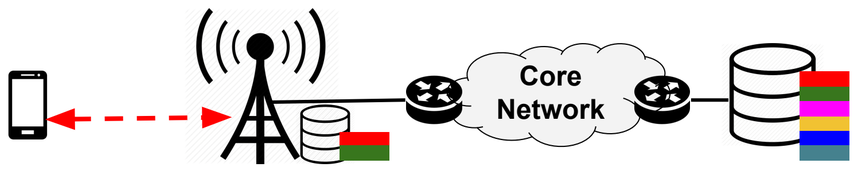

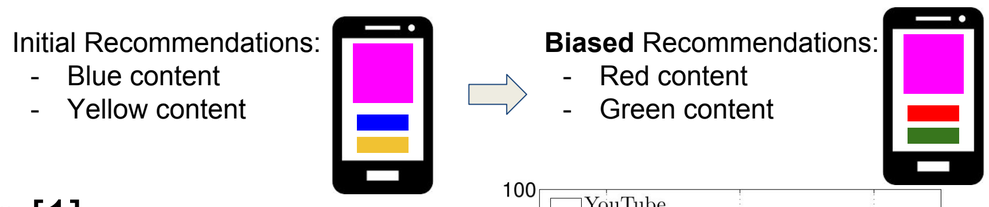

  - Promising gains, e.g., *10x* increase in the cache hit ratio [1]
  - But, no experiments with real users (e.g., simulation-based evaluation[1,2,3])

- **Goal**: experimentally assess the role and implications of recommendation systems on QoE with real users
- **Contribution**: platform to conduct experiments with real users

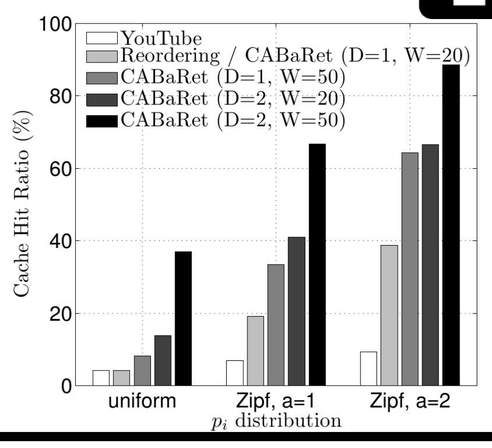

2. Controlled Experiments

http://www.experimentoqos.com/en

Try it yourself

Goal: quantify interplay between (i) QoS and (ii) user interests, on users preferences and QoE.
→ ultimately, build better recommendation systems!

Experiment: users watch pairs of videos
Video A: low QoS, but hypothesized as interesting
Video B: high QoS, but hypothesized as less interesting
→ "*Which video do you prefer?*"

Results:

(a) Word clouds reflect user experience:

QoS-insensitive users
focus on content

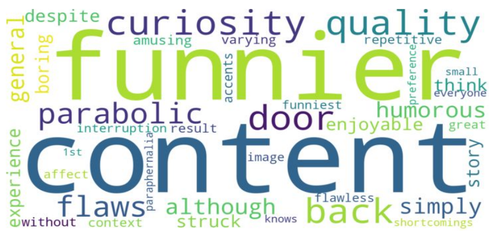

QoS-sensitive users
focus on quality

(b) Which video do you prefer to watch?

User interests and QoS together impact choices:
QoS role more prominent for users with elastic interest.

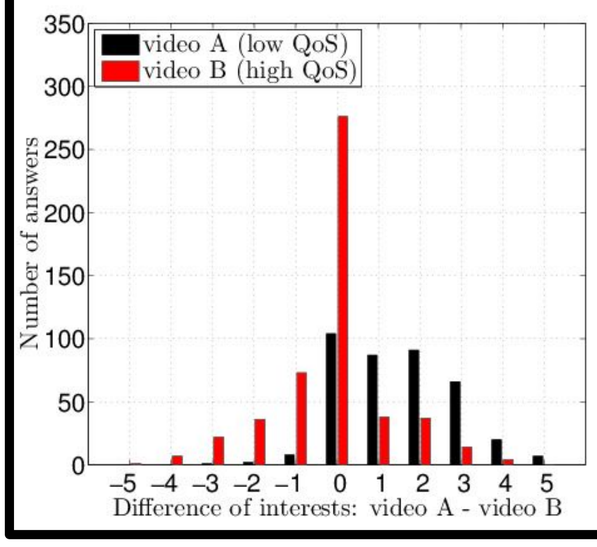

Key finding: 50% of the users with more interest in the video A (x-axis:1,2) enjoy watching more the high QoS video B
→ *take QoS into account in recommender systems!*

3. Experiments in the Wild

http://cabaret.ddns.net

Try it yourself

Platform: CABaRet [1] + *real YouTube experience*

- YouTube API (retrieve lists of related contents)
- Assume top trending/popular videos as cached
- Play non-cached videos with interruptions

Experiment:

- the user selects and watches YouTube videos; for each:
- we recommend 5 videos (cached & related, CABaRet)
- the user rates her QoS, QoE, and QoR

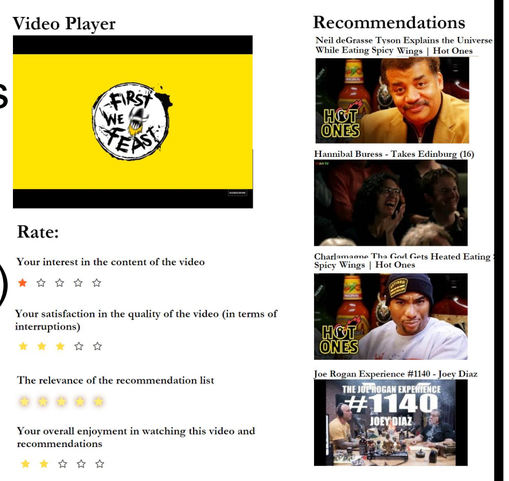

Goals:

- Test CABaRet performance with real users
- Quantify interplay of QoS, QoR, QoE

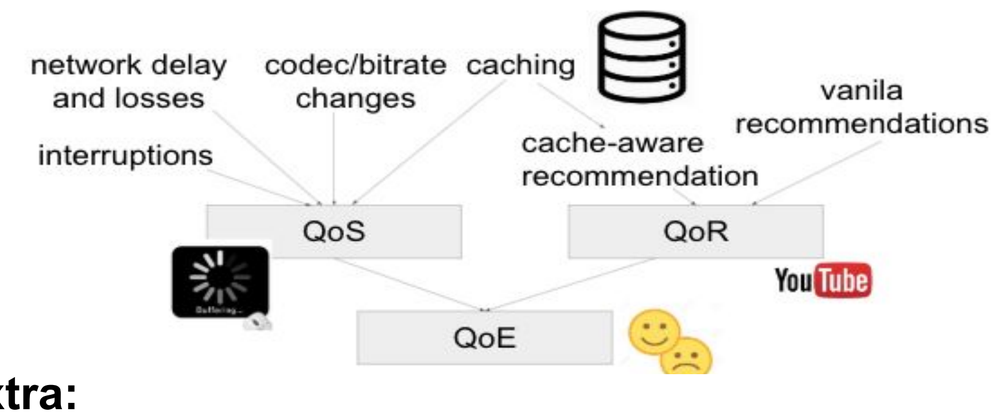

Extra:
→Build your own experiment on the CABaRet platform!

[1] S. Kastanakis et al., "CABaRet: Leveraging Recommendation Systems for Mobile Edge Caching", in ACM MECOMM (SIGCOMM workshop) 2018.
[2] P. Sermpezis et al., "Soft Cache Hits: Improving Performance through Recommendation and Delivery of Related Content", in IEEE JSAC, 2018.
[3] D. Munaro et al., "Content recommendation and service costs in swarming systems", in Proc. IEEE ICC, 2015.

Contact us at: kastan@csd.uoc.gr, mateussnogs@gmail.com, felipe.link1996@poli.ufrj.br, sermpezis@ics.forth.gr, sadoc@dcc.ufrj.br, spyropou@eurecom.fr