Towards Understanding Mobile Video Ads

Maria Carpen Amarie[1], Ioannis Pefkianakis[2], Henrik Lundgren[2]
Université de Neuchâtel[1], Technicolor Research Paris[2]
maria.carpen-amarie@unine.ch, {ioannis.pefkianakis, henrik.lundgren}@technicolor.com

1. INTRODUCTION

Nowadays, online video streaming dominates the Internet traffic, corresponding to a stunning amount of 3 million years of videos crossing the Internet each month. Simultaneously, the use of Internet on mobile devices, such as smartphones and tablets, is steadily increasing. Today, more than half of this mobile traffic consists of video. Consequently, Internet advertisers also turn towards mobile video advertising. In the literature, although the online advertising has been extensively studied, only few efforts (e.g., [2, 1]) have focused on video ads.

In this work, we report on our initial efforts towards understanding mobile video ads. We perform a novel characterization of mobile video ads related to a well-defined subset of videos on YouTube, the largest online video service by far. Then, we observe the impact of a set of core features (category, length, display time and frequency) on ad lifetime and number of occurrences. Finally, we devise a caching algorithm based on the discovered results.

2. ANALYSIS

As it is not possible to crawl the millions of available YouTube videos, we focus on a well-defined subset with similar attributes (i.e., length, category type, size). For 3 months we crawl the www.top40-charts.com website, which features around 700 music videos hosted on YouTube. We analyze our resulting data set and make the following main observations:

- The distribution of observed ad categories is skewed: only 4 categories (out of 15) account for almost two-thirds of all *unique* ads
- 50% of the ads are shorter than 1 minute
- Roughly 50% of the ads are 5 MB or less in size (although the largest observed ad is almost 40 MB)
- Short length video ads tend to have higher number of occurrences.
- Short length video ads tend to live longer.

3. CACHING ALGORITHM

To this end, we design a simple caching system. It monitors the HTTP communication and when it identifies a YouTube video ad request, it checks if the ad is already cached; if not, a prediction module infers the number of occurrences of the ad and decides to cache it or not. The *local storage* (cache) keeps track of the predicted number of occurrences for each ad, decreases it when the ad is seen again and removes the ad when it reaches 0. To evaluate our caching system, we perform trace-driven simulations using the YouTube data set collected with our crawler. We assume that a user watches N videos per day (N varies from 10 to 90) and has a fixed cache size of 800 MB. We compare our prediction caching system with a random-caching system (flips a coin to decide to cache an ad which is not already cached) and a no-caching system (current policy on mobile devices, which do not cache video ads). Figure 1 shows that ad data consumption is reduced up to 52%, 32% when compared with no-caching, random-caching strategies, respectively. These savings increase with the number of the video ads to be watched by the user.

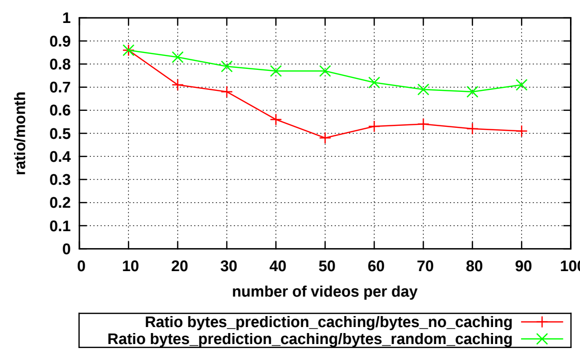

Figure 1: Performance evaluation

4. FUTURE WORK

This work is a first stepping stone to better understand mobile video ads. We plan to extend our work along the following three axes; First, we will study and derive the impact of these savings into exact battery savings and download delay. Second, we will refine and evaluate our caching system based on a wider set of the vast YouTube video sources. Finally, we will explore in further detail the impact of the user's profile on the displayed ads and extend the type of considered video types.

5. REFERENCES

[1] R. Kokku, R. Mahindra, S. Rangarajan, and H. Zhang. Opportunistic alignment of advertisement delivery with cellular basestation overloads. MobiSys '11.

[2] S. S. Krishnan and R. K. Sitaraman. Understanding the effectiveness of video ads: A measurement study.