

Tube Streaming: Modelling Collaborative Media Streaming in Urban Railway Networks

Argyrios G. Tasiopoulos, Ioannis Psaras, Vasilis Sourlas and George Pavlou
Dept. of Electronic and Electrical Engineering, University College London, UK
Email: (argyrios.tasiopoulos, i.psaras, v.sourlas, g.pavlou)@ucl.ac.uk

Abstract—We propose a quality assessment framework for crowdsourced media streaming in urban railway networks. We assume that commuters either “tune in” to some TV/radio channel, or submit requests for content they desire to watch or listen to, which eventually forms a playlist of videos/podcasts/tunes. Given that connectivity is challenged by the movement of trains and the disconnection that this movement causes, users collaboratively download (through cellular and WiFi connections) and share content, in order to maintain uninterrupted playback. We model collaborative media streaming for the case of the London Underground train network. The proposed quality assessment framework comprises a utility function which characterises the Quality of Experience (QoE) that users (subjectively) perceive and takes into account all the necessary parameters that affect smooth playback. The framework can be used to assess the media streaming quality in any railway network, after adjusting the related parameters.

To the best of our knowledge, this is the first study to quantify the perceptual quality of collaborative media streaming in (underground) railway networks from a modelling perspective, as opposed to a systems perspective. Based on real commuter traces from the London Underground network, we evaluate whether audio and video can be streamed to commuters with acceptable QoE. Our results show that even with very high-speed Internet connection, users still experience disruptions, but a carefully designed collaborative mechanism can result in high levels of perceived QoE even in such disruptive scenarios.

Index Terms—Mobile Video/Audio Delivery, Collaborative Video/Audio Streaming, Crowdsourcing.

I. INTRODUCTION

We envision ubiquitous streaming of popular channels, *e.g.*, national broadcasters, in smart city environments. We argue that there is a lot of space for resource optimisation in mixed cellular and WiFi access environments complemented with local transmission of content between mobile devices. Arguably, the quality of Internet services deteriorates when several hundreds of users attempt to connect simultaneously through the same WiFi Access Point (AP) or cellular Base Station (BS). This situation is rather common in large metropolitan areas where hundreds of users commute (*e.g.*, in trains or buses), wait in stations or airports, or just move slowly towards their destination [1]. The case becomes even more challenging when connectivity is physically disrupted, *e.g.*, when (underground) trains travel between stations. In those cases, installing more bandwidth will not necessarily improve performance [2], simply because quality often suffers due to: *i*) frequent handovers, *ii*) the hundreds of sessions that an AP/BS

has to handle simultaneously, and *iii*) physical challenges, such as long disconnection periods. In urban railway networks, for instance, it is not uncommon the situation where in busy times more than 500 commuters are onboard a train (and might want to access online services), trains stay in stations (and therefore, connectivity is available) for as little as 20 seconds, followed by a disconnection period of 4 minutes (or more), when the train is in the tunnel. The situation is even worse for real-time streaming sessions, as opposed to downloading some static web-page content.

In the absence of any collaboration between users to download and stream content collectively, the users end up with degraded media quality and, more often than not, abandon their sessions.

Crowdsourcing-based approaches have been proposed recently in order to deal with the above challenges [3], [4], [5], [6]. Crowdsourced content retrieval is based on the premise that users share their storage, connectivity and energy resources given some direct or indirect incentives, *e.g.*, improved service quality [7], or some monetary reward [6], [8].

The implementation details of such a crowdsourced mobile video and audio on demand streaming service for commuters in urban railway networks have been studied in [3]. However, the question of *whether such a system can provide acceptable media experience (or QoE) to participating users has not been answered yet*. That said, in this study, we fill this gap by building a *quality assessment framework* that attempts to answer the above question.

Quality of Experience (QoE) is traditionally characterised by the subjective perception of users [9] and as such, it is difficult to quantify in terms of objective metrics [10]. Indeed, objective measures of user engagement, such as bitrate, throughput, startup delay and buffering, which by and large comprise the state of the art to date, are not adequate in a media-dominated Internet, let alone a mobile environment. The proposed framework, which is based on a utility function, extends our previous work in [5] and goes beyond the above primitive metrics to take into account the user’s tolerance to disruptions, the energy needed to download and share media content, as well as the cost of using the cellular network. The energy factor, for instance, can influence user engagement in collaborative streaming in fear of battery depletion. In other words, users will not spend energy (and cellular data) to participate in a collaborative streaming system unless quality compensates. As such, we argue that a quality assessment framework for media delivery is necessary to characterise the

perceptual and subjective quality as opposed to the network-dependent, objective quality.

An urban railway system exhibits a number of challenges, as well as a number of physical properties, which can be utilised to make design decisions. For example, the train route and timetable is fixed (as opposed to walking users or vehicles). On the other hand, connectivity is physically disrupted when trains move between stations. Furthermore, disconnection periods can be as long as eight times this of the connection period, during which there is no WiFi connectivity and cellular connection is rather poor. Finally, the users' commuting patterns are unknown, but the aggregate number of commuters can be approximated depending on factors such as the time of day (*i.e.*, peak or off-peak time) and the direction of travel (*i.e.*, bound towards the city centre or not).

The main contributions of this work are:

- We define a utility function as the key aspect of our quality assessment framework. This utility function takes into account playback disruption and energy factors and provides a quantitative measure of a media delivery system for railway networks.
- We analyse a 17-day sample of commuters' journey traces, provided by Transport for London (TfL, <http://www.tfl.gov.uk>). We identify and illustrate commuting patterns and approximate the number of users travelling in each train throughout the day, as well as the connection and disconnection periods they experience.
- We quantify the users' experience by applying the utility function to the commuters' traces. Interestingly, we find that although it is difficult to achieve uninterrupted multimedia playback, especially during the rush hours, a collaborative media streaming application can provide acceptable QoE.

The rest of the paper is organised as follows. In Section II, we discuss related works. Section III includes the analytical description of the problem under consideration, while in Section IV we give details of our commuter data trace and we evaluate the performance of the system. Finally, Section V concludes this study.

II. RELATED WORK

The mobility properties of public transportation systems, as well as the limited connectivity during the commute makes communications in these environments a rather challenging problem. In the absence of any connectivity in underground train networks, authors in [11] build on the concept of *Familiar Strangers* [12] and analyse commuter traces to investigate whether collocation patterns exist. They find that such patterns indeed exist and build a content sharing and distribution network around content stored in commuters' devices. They identify users who possess content which is of interest to others and who will travel long enough within range of each other in order to successfully transfer the content in a *p2p* manner.

In [13] and [14] authors propose installation of a hub on public transportation vehicles, which commuters utilise in order to stay connected throughout their journeys. Although existence of a connectivity device on board the vehicle is desirable, the authors do not exploit aggregate connectivity

and storage opportunities offered by user devices. Peer-to-peer communications between mobile devices have been extensively studied in the context of delay-tolerant networks, taking also advantage of social links to drive connectivity decisions [15], [16], but also with a target to improve QoE [7], [17].

None of these works, however, has attempted to model the performance of real-time multimedia streaming in challenged, intermittent connectivity environments and as such we consider these works complementary to ours.

Closer to our work are studies that investigate co-operative download techniques for mobile users. Such studies are applicable to a broader scope of problems in the context of delay-/disruption-tolerant networks. For instance, co-operative techniques that exploit cellular and local WiFi connectivity have been studied to improve the capacity available to mobile users, or in other words, increase the aggregated downlink rate of each user, *e.g.*, [6], [18], [19]. In [20], a group of collocated train commuters, using several wireless Internet access links, jointly access a video stream and each user contributes to the group by sharing his downloaded content. In [4], the authors formulate a network utility maximisation problem where a single static group of commuters attempts to access a video stream. According to [4], users try to utilise Internet access links as well as the device-to-device capacity by taking into account packet loss and applying network coding. The model in [4] is complementary to any study that investigates co-operative download for mobile devices. Finally, the authors in [3] design, implement, and evaluate *Microcast*, a system that improves the streaming experience of a group of users, albeit in relatively static conditions.

This second group of studies focuses mainly on *the implementation principles* of a co-operative streaming service, ignoring to a large extent the performance that such systems would achieve in reality. *Quality of Service* (QoS) metrics cannot be used for this purpose as they focus on network-related aspects and ignore the human perception factor. There is a clear lack in the literature of studies that characterise QoE for mobile media delivery. *Mean Opinion Score* (MOS) is the most widely used ranking metric that takes into account the objective opinion of individual users [9], [10]. In this study, we fill this gap by building a utility function that characterises the user's QoE when streaming media in train networks. Given the connectivity challenges in this case, tolerance to disruptions is a central factor. Furthermore, given that in order to participate in collaborative media streaming, users spend battery resources, energy consumption also plays a central role in the QoE assessment.

III. QUALITY ASSESSMENT FRAMEWORK

A. System Model

We consider that users access the Internet through either WiFi or cellular links and can either individually *Pull*, or collectively *PULL and SHare* (PUSH) content. In this section we lay out the description of the model for each of the Internet access methods considered (*i.e.*, cellular and WiFi), as well as for each content retrieval approach (*i.e.*, *Pull* or *PUSH*). We define as an epoch i , Ep^i , a time interval of duration $|Ep^i|$,

which consists of a connection period C^i (of duration $\lfloor C^i \rfloor$) and a disconnection, or poor connection quality period \tilde{C}^i (of duration $\lfloor \tilde{C}^i \rfloor$). Clearly, $|Ep^i| = |C^i| + |\tilde{C}^i|$, which also implies that a new epoch starts at each station, where the previous one finishes. Connection period is the time that a train spends in a station and disconnection period is the time it takes for the train to arrive at the next station.

In the simple case, users *Pull* content individually from the Internet. Video or audio content is split in chunks, where every chunk contains y seconds of playback time at bit-rate b . Users can also form groups to *PULL and SHare (PUSH)* content with fellow-commuters in the vicinity. Initially, we present the basic framework which effectively quantifies the utility obtained by individual users (Sections III-B, III-C). We then extend the utility function to incorporate aspects such as energy consumption and cellular download data charges (Section III-D). Finally, we adjust the basic framework to each of the content retrieval approaches (*i.e.*, *Pull* and *PUSH*) and access method technologies in Sections III-E, III-F, III-G.

We assume that the functionality of the *PUSH* approach, is realised in the context of a mobile Backend as a Service (mBaaS) platform that runs in the cloud (similar in rationale to [3] and [6]) and is responsible for managing group synchronisation in terms of collaborative content retrieval. For the purposes of this study we ignore any potential implementation overhead (*e.g.*, chunk download scheduling). In particular, the mBaaS platform assigns to each member the content chunks that (s)he has to download and share with the rest of the group participants. All members of a group download and share equal number of chunks which implies *fairness* in terms of computation and communication effort. Our model notation is given on Table I.

B. Playback and Playback Disruption

The Internet access medium (*i.e.*, WiFi or cellular), as well as the content retrieval approach (*i.e.*, *Pull* or *PUSH*) affects the number of chunks that a user can receive over a predefined period of time. We denote as $X_{A,R}^i$ the number of chunks that a user receives over epoch i when (s)he accesses the Internet by technology A and retrieves a content by approach R .

Given reception of $X_{A,R}^i$ chunks, we calculate the *playback* and *playback disruption* periods that a user experiences over consecutive epochs. Firstly, we estimate the total number of chunks received by a user over epochs f to i as $X_{A,R}^{f \rightarrow i} = \sum_{j=f}^i X_{A,R}^j$, where f is the epoch during which the user enters the system.

Next, we express the number of chunks in terms of playback time. We assume that a chunk can be watched/listened only when it has been fully downloaded. Hence, the watching time worth of downloaded content from epoch f to epoch i for one user is:

$$L_{A,R}^{f \rightarrow i} = \lfloor X_{A,R}^{f \rightarrow i} \rfloor \times y, \quad (1)$$

where y is the playback duration of a content chunk.

Given that we model an on-demand streaming service, the downloaded playback time will differ from the actually watched playback time over epochs f to i , $W_{A,R}^{f \rightarrow i}$, due to

$ C^i , \tilde{C}^i $	i -th connection/disconnection duration
$ Ep^i $	i -th epoch duration
A	Internet access technology (WiFi,cellular)
R	Content retrieval approach (<i>Pull,PUSH, Hybrid</i>)
f	Epoch user started downloading the content
y, b, S	Playback time, bit-rate, and size of a chunk
Y	Content playback time duration
$X_{A,R}^i$	Chunks received at epoch i for A and R
$X_{A,R}^{f \rightarrow i}$	Total chunks received until epoch i (incl.)
$L_{A,R}^{f \rightarrow i}$	Total playback time received until epoch i (incl.)
$W_{A,R}^{f \rightarrow i}$	Playback time watched until epoch i (incl.)
$D_{A,R}^{f \rightarrow i}$	Playback Disruption until epoch i (incl.)
$U_{A,R}^{f \rightarrow i}$	User utility until epoch i (incl.)
$\tilde{U}_{A,R}^{f \rightarrow i}$	User extended utility until epoch i (incl.)
$U_{A,R}^{C,f \rightarrow i}, U_{A,R}^{E,f \rightarrow i}$	User cellular and energy cost to epoch i
a_d, a_c, a_e	User disruption, cell, and energy sensitivity
X_{cell}, X_{WiFi}	Per second chunks delivery rate at cellular, WiFi, and sharing interface for a group g
X_{p2p}^g	# of users requesting access at moment t
$N(t)$	Total bandwidth assigned to a platform of a station
B_{WiFi}	User good and poor cell bandwidth
$B_{cell}, \tilde{B}_{cell}$	# of users of group g at moment t
$N_g(t)$	Bandwidth limit for local $p2p$ transfers
B_{p2p}	Maximum amount of content to be shared and downloaded over an epoch i for group g
$V_{max}^{i,g}(p2p)$	Identical chunks downloaded by all the members of group g during good and poor connectivity
$V_A^{i,g}, \tilde{V}_A^{i,g}$	

TABLE I: Model Notation

buffered content. We model this difference in a retrospective manner between epochs i and $i - 1$ and express the actually watched playback time $W_{A,R}^{f \rightarrow i}$ as the sum of: *i)* the watched playback time during the previous epochs $W_{A,R}^{f \rightarrow i-1}$ and *ii)* the difference between the total downloaded playback time between f and i , $L_{A,R}^{f \rightarrow i}$, and the actually watched playback time during the current epoch ($L_{A,R}^{f \rightarrow i} - W_{A,R}^{f \rightarrow i-1}$). Note that in case the user has buffered enough content to get through the current epoch, then $L_{A,R}^{f \rightarrow i} - W_{A,R}^{f \rightarrow i-1} = |Ep^i|$. Therefore, we have:

$$W_{A,R}^{f \rightarrow i} = \min[W_{A,R}^{f \rightarrow i-1} + \min(|Ep^i|, L_{A,R}^{f \rightarrow i} - W_{A,R}^{f \rightarrow i-1}), Y], \quad (2)$$

where Y is the total playback duration from the beginning of the journey, which apparently works as an upper bound of the watched time and $W_{A,R}^{f \rightarrow i-1} = 0$ for the first epoch.

Finally, the playback disruption time until epoch i , $D_{A,R}^{f \rightarrow i}$, is calculated as:

$$D_{A,R}^{f \rightarrow i} = \begin{cases} D_{A,R}^{f \rightarrow i-1}, & \text{if } W_{A,R}^{f \rightarrow i} = Y \\ \sum_{j=f}^i |Ep^j| - W_{A,R}^{f \rightarrow i}, & \text{otherwise,} \end{cases} \quad (3)$$

where $D_{A,R}^{f \rightarrow i-1} = 0$ for the first epoch.

C. Utility Function

The playback time, $W_{A,R}^{f \rightarrow i}$, as well as the playback disruption, $D_{A,R}^{f \rightarrow i}$, are the fundamental components that we use in order to describe the user's utility in terms of QoE. Clearly, *the ideal utility of a user until epoch i is equal to the sum of the epochs' duration, or the content duration Y , whichever is shorter:*

$$U_{ideal}^{f \rightarrow i} = \min(Y, \sum_{j=f}^i |Ep^j|). \quad (4)$$

We consider that the utility decreases due to the playback disruption $D_{A,R}^{f \rightarrow i}$ and express a *user's utility* in terms of undisrupted playback time according to the formula:

$$\text{Utility Function} : U_{A,R}^{f \rightarrow i} = W_{A,R}^{f \rightarrow i} - a_d \times D_{A,R}^{f \rightarrow i}, \quad (5)$$

where a_d is the user's tolerance to playback disruptions. The disruption tolerance factor decreases the user's utility by a_d . For instance, for a user who has watched 10 secs of undisrupted playback and therefore has built a utility function equal to 10, a disruption tolerance factor equal to 2 will decrease his utility to 8, after 1 second of disruption. We consider the disruption tolerance factor as a central component of the utility function for media delivery in connectivity-challenged mobile environments.

We express the “*efficiency*” of an Internet access technology A and content retrieval approach R for a commuter according to the following *utility ratio*:

$$\text{Utility Ratio / Efficiency} : Q_{A,R}^{f \rightarrow i} = \frac{T_d + U_{A,R}^{f \rightarrow i}}{T_d + U_{ideal}^{f \rightarrow i}}, \quad (6)$$

which is bounded by the interval $[0, 1]$. T_d is the “*Initial Tolerance Interval*”, which indicates the user's patience to start-up delay. In the rest of this work, we use Eq. 6 to quantify the QoE that users obtain during their journeys, which together with the *utility function* in Eq. 5 comprise the two main building blocks of the proposed quality assessment framework.

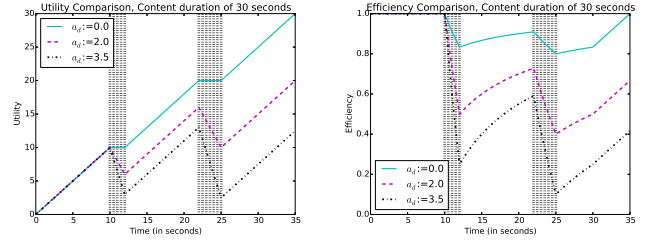
Fig. 1 illustrates the *utility* (Eq. 5) and *efficiency* (Eq. 6) fluctuation for 3 users with disruption tolerance (*i.e.*, a_d) 0, 2, and 3.5, respectively; the users experience a playback disruption of 2 seconds after 10 seconds of playback time and another disruption of 3 seconds after playback time of 20 seconds; the total playback duration is 30 seconds, which also means that the ideal utility is equal to 30. Apparently, delay sensitivity equal to 0 leads to efficiency equal to 1 (Fig. 1b), while for delay sensitivity equal to 2 and 3.5 in this setting the produced efficiency is 0.666 and 0.416, respectively. This result demonstrates that *efficiency* is subjective when it comes to QoE, since it is subject to the users' temporal utility/satisfaction and their personal tolerance to disruptions.

The *Utility Ratio or Efficiency* can get negative values due to extended disruptions. In this case we assume that a rational user would quit the attempt to watch (listen) this video (music playlist), which would set the Efficiency's value to zero, or more formally:

$$Q_{A,R}^{f \rightarrow i,t} = \max\left(\frac{T_d + U_{A,R}^{f \rightarrow i,t}}{T_d + U_{ideal}^{f \rightarrow i,t}}, 0\right). \quad (7)$$

D. Cellular and Energy Cost

To further extend our model with a realistic representation of a working system we incorporate two more factors in the utility function of users. The first relates to the medium used to download content and is associated with the cost of using the cellular network, denoted as *cellular cost*. The second factor relates to the energy consumed in order to use this medium, denoted as *energy cost*. Given the description and structure of the utility function in Eq. 5 the cellular and energy costs need to be converted to actual playback time. That said, the cellular cost is the equivalent of the playback time downloaded



(a) Utility Comparison

(b) Efficiency Comparison

Fig. 1: Comparison of user utilities and efficiencies over time for a disruption of 2 and 3 seconds after 10 and 20 seconds of playback time.

through the cellular network, denoted as $U_{A,R}^{C,f \rightarrow i}$ for epochs f to i .

As regards the energy cost, an interface supporting radio technologies works in different *power states*, each one related to a different workload as well as power consumption. These states include the *Idle* state, when no Internet connection is required, and consumes the least possible energy. Starting from this state, an interface can be promoted to one or more productive states by spending a fixed amount of energy and time, where productive means that an interface becomes able to receive or transmit data. Finally, an interface experiences a tail power phenomenon where it stays in a high power state, in anticipation of more data exchange, before returning to its initial *Idle* state.

We denote as $E_{A,R}^{prom}$ and $E_{A,R}^{tail}$ the fixed energy cost of all involved interfaces, of technology A and approach R , when being promoted to a productive state and when experiencing the tailing phenomenon, respectively. Thus, the energy spent until epoch i , $E_{A,R}^{f \rightarrow i}$, is:

$$E_{A,R}^{f \rightarrow i} = E_{A,R}^{prom} + E_{A,R}^{prod,i} + E_{A,R}^{tail} + E_{A,R}^{f \rightarrow i-1}, \quad (8)$$

where $E_{A,R}^{prod,i}$ is the energy spent on the productive state of all interfaces for receiving and sharing content during epoch i . Therefore, by identifying the technology, A^* , and approach, R^* , which spend the smallest possible amount of energy until epoch i , $E_{A^*,R^*}^{f \rightarrow i}$, we can express the energy cost of a candidate technology, A , and approach, R , in terms of playback time, $U_{A,R}^{E,f \rightarrow i}$, by:

$$U_{A,R}^{E,f \rightarrow i} = \left(\frac{E_{A,R}^{f \rightarrow i}}{E_{A^*,R^*}^{f \rightarrow i}} - 1\right) \times L_{A^*,R^*}^{f \rightarrow i}. \quad (9)$$

The value of $U_{A,R}^{E,f \rightarrow i}$ in Eq. 9 is effectively the additional content (in terms of playback time) that a user would download if (s)he used technology, A^* , and approach, R^* .

Building on Eq. 5, the extended utility function that integrates the energy and cellular costs is:

$$\tilde{U}_{A,R}^{f \rightarrow i} = U_{A,R}^{f \rightarrow i} - a_c \times U_{A,R}^{C,f \rightarrow i} - a_e \times U_{A,R}^{E,f \rightarrow i}, \quad (10)$$

where a_c and a_e are the corresponding weights of cell and energy cost, called *cell* and *energy sensitivity*.

In the following, we first discuss the general chunk reception rates over WiFi and cellular connectivity, which are independent of the content retrieval approach (Section III-E). Then we adjust those reception rates to the *Pull* and *PUSH* cases, in

Sections III-F and III-G, respectively. The target is to define the total chunks received until epoch i , that is, the $[X_{A,R}^{f \rightarrow i}]$ component of Eq. 1, which then feeds Eq. 2 and eventually the Utility Function defined in Eq. 5 and extended in Eq. 10.

E. Chunk Reception Rate over WiFi and Cellular

The size of a single chunk, S , is $S = b \times y$, where b is the bitrate and y is the chunk's playback duration.¹ Then, assuming interference-free cellular downlinks, the cellular chunk delivery rate (per second) will be $X_{cell}(t) = B_{cell}/S$, where B_{cell} is the bandwidth allocated to the user by the cellular network provider. For simplicity, we consider that the cellular bandwidth is stable, irrespectively of the number of active users.² We assume that B_{cell} is the bandwidth availability when the train is stopped at some station and that the connection quality deteriorates when the train is moving (*i.e.*, $\ddot{B}_{cell} < B_{cell}$).

Without loss of generality, we assume that the available WiFi access bandwidth per platform of each station is equal to B_{WiFi} , which is equally shared among the users/commuters at each platform. Therefore, the chunk reception rate per second and per user at time t , received through the WiFi AP of a platform is, $X_{WiFi}(t) = \frac{B_{WiFi}}{N(t) \cdot S}$, where $N(t)$ is the number of users requesting Internet access at moment $t \in [t_i, t_{i+1})$ at the given platform. Please note that there is no WiFi connectivity during disconnection periods (*i.e.*, when trains are in-between stations), which means that the chunk reception rate is: $\ddot{X}_{WiFi} = 0$.

In the *PUSH* approach, apart from the chunks received through the Internet, users receive chunks from group members too. The chunk reception rate between group members in the *PUSH* approach is proportional to the group size. Chunks are shared over bandwidth B_{p2p} and the number of chunks that can be shared among a group g of $N_g(t)$ members at a moment t , $X_{p2p}^g(t)$, is defined as:

$$X_{p2p}^g(t) = \frac{B_{p2p}}{N(t) \cdot S} \times N_g(t), \quad (11)$$

Eq. 11 implies an underlying pseudo-broadcast sharing mechanism (similar to [3]) where a single peer at a time unicasts its content to another peer, while the rest of the members of the group overhear the transmission. Note that the mBaaS platform (mentioned earlier in Section III-A) is also responsible for organising the “sharing turns”, given that only one user can unicast at a time.

F. Pull Reception Rate

In the simple *Pull* approach users *pull* content individually. For cellular Internet access, the total number of chunks received during epoch i is equal to:

¹Note that we do not consider dynamic rate adaptation (*e.g.*, DASH) for simplicity of modelling. We evaluate the proposed model under the lowest possible rate, hence, any higher bandwidth availability will only increase the performance we observe.

²In reality, even the cellular bandwidth assigned to each user can be influenced by the level of contention (that is, number of users), but this happens for larger number of users, possibly in the order of thousands, which is not the case of a train (station).

$$X_{cell,Pull}^i = |C^i| \times X_{cell} + |\tilde{C}^i| \times \ddot{X}_{cell}. \quad (12)$$

In the WiFi AP case, where the medium is shared between users, the total number of chunks that a user receives through a WiFi AP over epoch i is:

$$X_{WiFi,Pull}^i = \int_{t=t_i}^{t_i+|C^i|} X_{WiFi}(t) dt, \quad (13)$$

where t_i is the starting time of epoch i .

Finally, in case users utilise both interfaces for Internet access, in a *Hybrid* way, the total number of chunks they receive during epoch i is:

$$X_{Hybrid,Pull}^i = X_{cell,Pull}^i + X_{WiFi,Pull}^i. \quad (14)$$

G. Pull and Share (PUSH) Reception Rate

In the *PUSH* approach, users who belong to a group g of $N_g(t)$ members, at moment t , share their downloaded content. The chunk reception rate from fellow group members is given in Eq. 11. Our model does not include multi-hop transmissions which implies that all members of a group have to be within transmission range of each other. For that purpose, we consider WiFi Direct as the technology of choice in order to transmit in long distances with high rates [21]. We also assume that devices can make simultaneous use of two separate half-duplex WiFi interfaces, one for downloading through the WiFi AP and another one for local sharing. Our approach is also applicable in the simple scenario where only one WiFi interface is available per device, but in that case downloading and sharing should take place sequentially, that is, the users would first download the required chunks and then share them with the rest of the group.

The total volume of content that can be shared between a group (according to Eq. 11) is X_{p2p}^g , or more formally, *the theoretical maximum amount of content that can be shared over an epoch i for a group g , $V_{max(p2p)}^{i,g}$, is:*

$$V_{max(p2p)}^{i,g} = \int_{t=t_i}^{t_i+|Ep^i|} X_{p2p}^g(t) dt. \quad (15)$$

On the other hand, *the maximum volume of content that can be downloaded collaboratively over the connection period of an epoch by access approach A , $V_{A,C}^{i,g}$, is:*

$$V_{A,C}^{i,g} = \int_{t=t_i}^{t_i+|C^i|} X_A(t) \times N_g(t) dt, \quad (16)$$

where $X_A(t)$ is the number of downloaded chunks per second per group member, during the connection period.

Finally, the corresponding maximum downloaded content under poor connectivity (*i.e.*, using cellular connection when the train is in-between stations), $V_{A,\tilde{C}}^{i,g}$, is:

$$V_{A,\tilde{C}}^{i,g} = \int_{t=t_i+|C^i|}^{t_i+|Ep^i|} \ddot{X}_A(t) \times N_g(t) dt, \quad (17)$$

where $\ddot{X}_A(t)$ is the number of downloaded chunks per second per group member, during poor connection (or disconnection)




Fig. 2: Average journey duration and number of commuters per end-to-end journey of the line.

period. It follows that *the total volume of content that can be downloaded over an entire epoch i , $X_{A,PUSH}^i$, is:*

$$X_{A,PUSH}^i = V_{A,C}^{i,g} + V_{A,\tilde{C}}^{i,g}. \quad (18)$$

Eq. 18 effectively expresses the total number of chunks that the group can download *and* share within epoch i . In case $X_{A,PUSH}^i > V_{max(p2p)}^{i,g}$ the group has downloaded exactly as much content as it can share during epoch i . This would happen when the Internet access bandwidth is higher than the $p2p$ one and therefore, users can download faster than they can share. In this case, users continue downloading individually without sharing.

IV. EVALUATION

A. Commuter Journey Traces Dataset

We analyse the anonymised commuter trace dataset from London’s Oyster Radio Frequency Identification cards. The dataset is a 17-day trace of all journeys made in all of the 11 lines of the London Underground network; the total number of journeys in our dataset is in the order of one million per day. Each journey is identified by an entry and exit point in the network, which is recorded at the granularity of one minute. The traces do not include information regarding the specific lines used by commuters, hence, in order to get an insight of how many passengers are on board a train at some given point throughout the day, we build a custom simulator which takes into account both the entry/exit points of commuters, but also the specific topology of the network. That is, we also consider the distance between stations, as well as the intersections of lines at each station and the routes that each line follows.

Given that a journey’s route might cross more than one lines, and that commuters follow the shortest path with the fewest interchanges between their entry and exit points, we process the trace and assign commuters to specific lines and then to individual trains.

Due to space limitations and for better visualisation of our results, we choose one line and present results for this line only. However, we report that evaluations with most of the lines of the network present similar results.

B. Group Formation Insights/Potential

Fig. 2 depicts the average number of commuters that our chosen line serves per one end-to-end journey. Obviously, not all commuters travel from one end to the other, but this plot includes all commuters that at some point in their journey use this line of the network. The number of commuters is averaged over all trains of the day in per hour time slots. We observe




Fig. 3: Number of commuters per train (single direction) for a station in the beginning, middle, and end of the line.

that there are two clear peaks during the morning and the afternoon rush hours. During those times, one train can serve up to 1400 commuters in its end-to-end journey. Depending on their entry and exit points, these commuters can form groups to *PULL* and *SHARE* content from the network.

In order to get a closer view of the potential to form groups, in Fig. 2, we also depict the average journey time per commuter, as well as the standard deviation of the per commuter journey time. According to this plot, commuters use this line on average for approximately 9 minutes with a standard deviation of around 5 minutes. Although the proportion of time that commuters physically share journeys depends on how much the chosen routes overlap (*i.e.*, entry and exit points), this time is enough to watch or listen to the headline stories of some TV/radio channel.

Clearly, from Fig. 2 we see that a line serves different volumes of commuters at different times of day, but their distribution over the stations of the end-to-end line needs further investigation. A more detailed analysis though, shows that commuters’ volumes differ hugely at each station depending on the train direction and the time of day. We note that the lines normally start and terminate at the outskirts of the city centre, but always cross the city centre itself. In Fig. 3 we present the average number of commuters per train at three different stations over the duration of a day. The three stations chosen are towards the beginning of the line (Station 1), the middle of the line (a central location, Station 6) and the end of the line (Station 12). Interestingly, and somewhat expectedly, we see that the station at the beginning of the line serves most users in the morning rush hours (*i.e.*, users move towards the city centre), whereas the station towards the other end of the line serves most users during the afternoon rush hours (*i.e.*, on the way back). The station in the city centre presents two peaks one in the morning and one in the afternoon rush hour.

This result makes clear the need for provisioning of Internet streaming services according to journey patterns, station locations and the time of day. That is, given that connectivity is available only when trains are within stations, where they normally stay for about 20-30 seconds and that disconnection periods last 1-4 minutes, the amount of content that needs to be buffered in order to avoid playback disruptions differs hugely. Given also that different stations are busy at different times of day depending on the area and the direction of the train, the system has to be carefully modelled to include these factors in order to guarantee uninterrupted playback.

Standard Setting Variable	Value
Total bandwidth per station, B_{WiFi}	0.5 Gbps
Cell rate - moving (\tilde{B}_{cell})/static (B_{cell}) train	150-550 Kbps
Sharing bandwidth, B_{p2p}	54 Mbps
Connection period duration for each station i , $ C^i $	20"
Participation probability, p_{list}	50%
Music/Video bit-rate	160/419 Kbps
Chunk playback time, y	5"
Playlist duration, Y	15'
Zipf's exponent, a	1
Playlists generated per station, Z	5-95
Delay (a_d), Energy (a_e), Cell (a_c) Sensitivity	3, 0, 0

TABLE II: Standard Evaluation Setting

C. Simulation and Evaluation Setup

We assume that users have access to both WiFi connectivity when trains are in stations and cellular connectivity throughout their journeys. Although in the case of the London Underground cellular connectivity is not available, here we include this access option for completeness. We also assume that users can use both their cellular and their WiFi interfaces simultaneously to receive content (*Hybrid* approach).

Although the current WiFi access deployment at the London Underground network provides download speeds of up to 100Mbps, here, we consider a future, overprovisioned network where each station is connected to the Internet with 500Mbps links. This bandwidth is split between all platforms and among all lines passing through each station.

For the cellular case, we assume that users get between 150Kbps and 550Kbps - closer to the lower value when the train is moving and to the highest one when the train is static in the station. Finally, we set the users' sensitivity to delay (a_d in Eq. 5) equal to 3, which means that for every second of disruption the users' utility decreases by three, whereas their utility increases by one for every second of undisturbed playback. This value is specifically chosen as an extreme scenario, where users are not tolerant to disruptions.

In addition to the different access methods, we evaluate the users' QoE (*i.e.*, Eq. 6) when two different types of content are available, that is, music content (at 160Kbps) and video content (at 419Kbps). The chosen bit rates are the lowest possible for acceptable quality streaming. In cases where the system can achieve minimum disruptions, an adaptive increase of the corresponding bit rates can be applied (*e.g.*, through DASH), but this is out of the scope of this paper. Unless mentioned otherwise, the full list of settings of our evaluation setup is given on Table II.

Due to the limitations of the environment under investigation, we do not assume that users individually choose to stream content and only if requests match, then users form groups. This would clearly result in very few and very small groups, effectively reflecting the *Pull* case. Instead, we evaluate two specific scenarios. In the first one, we assume that users create music or video "playlists" according to genre preferences, *e.g.*, sports clips, or jazz music. Users then join a group and add their own preferences to the list. We assume that between 5 and 95 playlists are generated in each epoch, resulting in more than 500 playlists available throughout the train's end-to-end journey. The playlists' duration is set to 15 minutes to reflect the average journey time plus standard deviation. In our second scenario, users tune in to radio or TV channels [22], [23]




Fig. 4: 15-min Playlists scenario *Pull* efficiency for WiFi, Cell, and Hybrid access methods.

and we assume that 50 channels are supported by the system, which would cover the main broadcasters of a country. The difference between the *playlist* and *channel* scenario is that in case of the latter, users form larger groups. In both cases, users choose which playlist or channel to subscribe to following a Zipf distribution [24].

In both scenarios, and after the users have created new playlists at each station, these playlists are suggested to the rest of the commuters who are not participating in any other group. A commuter is interested in this list and participates with a probability p_{list} while choosing the r -th element of the list according to a Zipf distribution with probability $f(r; a, Z)$ where a is the Zipf exponent and Z is the number of elements in the list.

We measure the users' perceptual QoE as expressed through Eq. 6. The initial tolerance interval (T_d in Eq. 6) is set to 5 seconds, after the expiration of which each user abandons the attempt to receive content. This setting is based on studies that assess the patience of users to load web content (*e.g.*, [25]).

D. Performance Evaluation

1) *Pull Approach - Playlists scenario*: Fig. 4 illustrates the average (over all trains) efficiency of WiFi, Cell, and Hybrid access methods of the *Pull* approach for music and video streaming over one operational day of the tube network. Given that with the *Cell* approach each user gets between 150Kbps and 550Kbps and that the bitrate for streaming music rate is 160Kbps, it is straightforward that the network can support such a service. Hence, the *Cell* approach (and stemming from that, obviously, the *Hybrid* approach) can provide good QoE. On the other hand, the *WiFi* access method experiences disruptions which increase as the volume of commuters increases (*i.e.*, during rush hours - see Fig. 2 and Fig. 3). In the case of video streaming all three access approaches, *WiFi*, *Cell* and *Hybrid* perform worse than music streaming. In the following we exclude the music bitrate for the cell access approach, as it clearly can be supported when cell access is available (not for the case of the London Underground network though).

2) *PUSH Approach - Playlists scenario*: From Fig. 5 we observe that streaming music over the WiFi network is challenged by disruptions/disconnections, achieving performance close to 0.8 at best, while during rush hours this can go as low as 0.4. In the case of Video playlists, which is even more demanding (*i.e.*, higher bit rates), achieves lower performance ranging between 0.7-0.8 when utilising both the WiFi and the cellular connection. Note that the performance




Fig. 5: 15-min Playlists scenario *PUSH* efficiency for WiFi, Cell, and Hybrid access methods.




Fig. 6: Streaming Channels - *PUSH* efficiency for WiFi, Cell, and Hybrid access methods.

of the *PUSH* approach is acceptable during off-peak time, especially considering the extreme delay sensitivity assumed here. Overall, based on the parameters used here (especially for the delay sensitivity) collaborative download can provide acceptable service quality, especially when it is combined with cellular connectivity. In the following we experiment with different parameters to investigate whether quality can improve further.

3) *PUSH Approach - Channels scenario*: Efficiency performance follows similar trends in the case of the second scenario (*i.e.*, TV/radio channel streaming - see Fig. 6), where, however, we observe a small increase in the QoE perceived by the users. This is because, in the “channel” case, groups form for longer time periods [26] and also more users join in as trains move towards their destination. That said, groups are bigger in size compared to the playlists scenario, giving the opportunity to move more content locally. Of course, during the rush hours the available WiFi bandwidth per commuter is significantly small and despite the size of the formed groups the efficiency is still quite low, since each user can download and share a very small portion of the desired content. As we show next, it requires a large amount of available WiFi bandwidth at each station in order to increase the performance of the system. Furthermore, another important factor that affects performance is the initial tolerance interval as we show later in this section.

4) *PUSH Approach - WiFi Bandwidth Factor*: In general more bandwidth available at each platform/station improves performance, but here we examine whether investing more on bandwidth would pay off in terms of users’ QoE. In Fig. 7 we present the case for video streaming to groups of commuters (*i.e.*, *PUSH* approach) for the WiFi and the Hybrid access methods.

Interestingly, the efficiency achieved by the WiFi access method during peak times overtakes the one achieved during off-peak times when the available bandwidth at each station is significantly large ($\geq 1\text{Gbps}$). In those cases, larger number

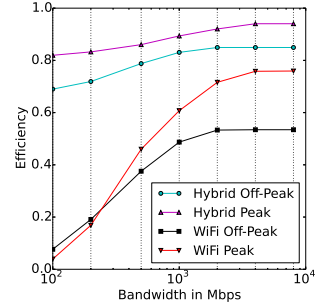


Fig. 7: 15-min Video Playlists scenario *PUSH* Efficiency for WiFi and Hybrid access methods for increasing bandwidth (Peak time: 8am, Off-Peak: 11am).

of commuters on board the trains result in larger groups. In turn, each group gets a larger stake of the available bandwidth. Overall, we see that even when 1Gbps is readily available at each of the hundreds of tube stations throughout the network, it is still difficult to achieve uninterrupted video streaming relying on the WiFi connectivity only, even at off-peak times. On the other hand, when collaborative streaming is combined with cellular connectivity, we observe performance close to 0.8, even for smaller bandwidth values and rather disruption-sensitive users (*i.e.*, $a_d = 3$).

5) *PUSH Approach - Cell and Energy Sensitivity*: The cell sensitivity integrates the cost factor, that is, the monetary cost to download through the cellular network (*i.e.*, a_c in Eq. 10), as opposed to the WiFi access. In general the efficiency decreases linearly in all the retrieval approaches (*i.e.*, *Pull* and *PUSH*) with respect to the cell sensitivity. In more details in the *Pull* case it declines with exactly the same rate during both peak and off-peak times, since the amount of data that each commuter downloads from the cellular interface in each case remains the same. On the other hand, for the *PUSH* approach (Fig. 8) the efficiency decline rate is less steep and the difference between the peak and off-peak time is proportional to the average group size. In this approach users also exchange a significant amount of data minimizing the data to be directly downloaded from the cellular network.

We use the power state machine presented in [27] to evaluate the energy sensitivity of the cellular and WiFi interfaces of smartphone devices for each one of the power state (*i.e.*, Promotion, Productive and Tail). Our findings show that despite the fact that *PUSH* uses an additional interface for sharing data, the low promotion and tail energy required by the sharing interface decreases the overall performance as we increase the energy sensitivity. Finally, we notice that increasing the sharing energy transmission coefficient, a_{tr} , causes only slight performance decline, proportional to the coefficient’s actual value (Fig. 9). This is partly because local transfers (through the sharing interface) completes much faster, therefore, spending little time in “transmission mode”.

6) *PUSH Approach - Tolerance Interval*: The “Initial Tolerance Interval” of users indicates the patience of users to startup delay. Throughout our evaluation this interval was set to 5 seconds, according to related studies on users’ tolerance to delays [25]. In this experiment, we investigate the effect of

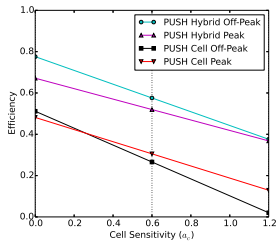


Fig. 8: Cell Sensitivity Increase - Scenario 1: Video Playlist

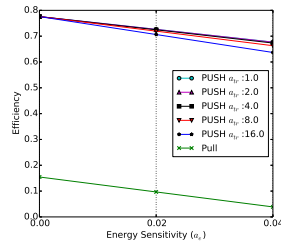


Fig. 9: Energy Sensitivity Increase - Scenario 1 Video Playlist, Hybrid

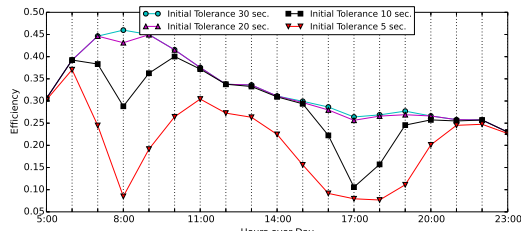


Fig. 10: 15-min Playlists scenario PUSH- WiFi Efficiency for different Initial Tolerance Intervals.

the tolerance interval on users' QoE (in Fig. 10); we assume a 15-min playlist when users are pulling and sharing content over WiFi and the tolerance interval is set to 10, 20 and 30 seconds. As the users' tolerance increases, we observe that the QoE increases too. This is a straightforward result, given that the groups formed in this case are larger and can therefore, get larger share of the available WiFi bandwidth. However, we also observe that after the 20-second threshold the performance does not improve further.

V. CONCLUSIONS

We have designed a model that characterises the QoE of users in an urban railway network, when they attempt to stream real-time media content. Our model applies to both simple *Pull* cases, and to collaborative download, *PULL* and *SHARE* (*PUSH*) cases and assumes WiFi connectivity in train stations, as well as cellular connection throughout the journey. The QoE is expressed in terms of the efficiency that users enjoy and takes into account connection and disconnection periods, sensitivity to disruptions and to the energy spent to download content, as well as the cost to use the cellular network. We have analysed commuters' traces and have applied our model to these mobility patterns. We found that it is difficult to maintain uninterrupted playback, especially in case of high bit rates, *i.e.*, video content, but at the same time, well thought-out collaborative mechanisms can increase the perceived QoE even under such challenged conditions. When cellular connectivity is available, performance improves considerably, given that users can utilise both interfaces (the WiFi and the cellular one) simultaneously.

ACKNOWLEDGMENTS

V. Sourlas is supported by the EC through the FP7-PEOPLE-IEF INTENT project, (GA no. 628360). The rest

of the authors are supported by the EU-Japan initiative under EC FP7 GreenICN project (GA no. 608518, NICT no. 167), the EC H2020 UMOBILE project (GA no. 645124) and the EPSRC INSP Fellowship (no. EP/M003787/1).

REFERENCES

- [1] G. Lyons and K. Chatterjee, "A human perspective on the daily commute: Costs, benefits and trade-offs," *Transport Reviews*, 2008.
- [2] L. Li, K. Xu, D. Wang, C. Peng, Q. Xiao, and R. Mijumbi, "A measurement study on tcp behaviors in HSPA+ networks on high-speed rails," in *IEEE INFOCOM 2015*.
- [3] L. Keller and et al., "Microcast: Cooperative video streaming on smartphones," in *MobiSys '12*.
- [4] H. Seferoglu, L. Keller, B. Cici, A. Le, and A. Markopoulou, "Cooperative video streaming on smartphones," in *Allerton Conference*, 2011.
- [5] A. G. Tasiopoulos, I. Psaras, and G. Pavlou, "Mind the gap: modelling video delivery under expected periods of disconnection," in *ACM CHANTS '14*.
- [6] D. Sryvelis, G. Iosifidis, D. Delimpasis, K. Chounos, A. Korakis, and T. Leandros, "Bits and coins: Supporting collaborative consumption of mobile internet," in *IEEE INFOCOM '15*.
- [7] C.-L. Tsao and R. Sivakumar, "On effectively exploiting multiple wireless interfaces in mobile hosts," in *ACM CoNEXT '09*.
- [8] W. Wu, R. Ma, and J. Lui, "Distributed caching via rewarding: An incentive scheme design in P2P-VoD systems," *Parallel and Distributed Systems, IEEE Transactions on*, 2014.
- [9] M. Venkataraman and M. Chatterjee, "Inferring video qoe in real time," *Network, IEEE*, 2011.
- [10] A. Balachandran, V. Sekar, A. Akella, S. Seshan, I. Stoica, and H. Zhang, "A quest for an internet video quality-of-experience metric," in *ACM HotNets-XI*, 2012.
- [11] L. McNamara, C. Mascolo, and L. Capra, "Media sharing based on colocation prediction in urban transport," in *ACM MobiCom '08*.
- [12] E. Paulos and E. Goodman, "The familiar stranger: anxiety, comfort, and play in public places," in *ACM SIGCHI '04*.
- [13] J. LeBrun and C.-N. Chuah, "Bluetooth content distribution stations on public transit," in *Workshop on Decentralized resource sharing in mobile computing and networking*, 2006.
- [14] F. P. Tso and et al., "Dragonnet: A robust mobile internet service system for long-distance trains," *Mobile Computing, IEEE Transactions on*, 2013.
- [15] C. Boldrini, M. Conti, and A. Passarella, "Exploiting users' social relations to forward data in opportunistic networks: The HiBOP solution," *Pervasive and Mobile Computing*, 2008.
- [16] P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: Social-based forwarding in delay-tolerant networks," *Mobile Computing, IEEE Transactions on*, 2011.
- [17] G. Ananthanarayanan, V. N. Padmanabhan, L. Ravindranath, and C. A. Thekkath, "Combine: leveraging the power of wireless peers through collaborative downloading," in *ACM MobiSys '07*.
- [18] S. Ioannidis, A. Chaintreau, and L. Massoulié, "Optimal and scalable distribution of content updates over a mobile social network," in *IEEE INFOCOM 2009*.
- [19] J. Whitbeck, M. Amorim, Y. Lopez, J. Leguay, and V. Conan, "Relieving the wireless infrastructure: When opportunistic networks meet guaranteed delays," in *IEEE WoWMoM '11*.
- [20] M. Stiemerling and S. Kiesel, "A system for peer-to-peer video streaming in resource constrained mobile environments," in *ACM U-NET Workshop '09*.
- [21] "WiFi alliance. Wi-Fi direct: <http://www.wi-fi.org>."
- [22] "PPLive: <http://www.pplive.com>."
- [23] "TVUnetworks: <http://www.tvunetworks.com>."
- [24] "Cisco. cisco visual networking index: Global mobile data traffic forecast update, 2013-2018. white paper, [online] <http://goo.gl/177haj>, 2014."
- [25] D. F. Galletta, R. Henry, S. McCoy, and P. Polak, "Web site delays: How tolerant are users?" *Journal of the Association for Information Systems*, 2004.
- [26] F. Wang, J. Liu, and Y. Xiong, "Stable peers: Existence, importance, and application in peer-to-peer live video streaming," in *IEEE INFOCOM '08*.
- [27] N. Ding, D. Wagner, X. Chen, Y. C. Hu, and A. Rice, "Characterizing and modeling the impact of wireless signal strength on smartphone battery drain," in *ACM SIGMETRICS '13*.