

Mobile Video Delivery via Human Movement

Gene Moo Lee, Swati Rallapalli, Wei Dong, Yi-Chao Chen, Lili Qiu, Yin Zhang
The University of Texas at Austin

Abstract—This paper proposes VideoFountain, a novel service that deploys kiosks at popular venues to store and transmit digital media to users’ personal devices using Wi-Fi access points, which may not have Internet connectivity. We leverage mobile users to deliver content to these kiosks. A key component in this design is an in-depth understanding of user mobility. We gather real mobility traces from two largest location-based social networks (Foursquare and Gowalla) and analyze both macroscopic and microscopic human mobility in different cities. Based on the insights we gain, we study several algorithms to determine the initial placement of content and design routing algorithms to optimize the content delivery. We further consider several practical issues, such as how to incentivize users to forward content, how to manage copyrights, how to ensure security, and how to achieve service discovery. We demonstrate the feasibility of VideoFountain using trace-driven simulations.

I. INTRODUCTION

People across the world spend significant time commuting. For example, Americans spend over 100 hours a year commuting [43]. The average trip is 38 minutes for the European Union [41]. One way to put the commute time into good use is for the passengers to download and watch interesting videos, such as popular movie clips, TV shows, and user generated videos, on-the-go. The current practice is to download videos on demand through cellular or Wi-Fi networks while watching them on the road. However, despite significant advances in cellular technologies, cellular providers have trouble keeping up with the rapidly increasing user demands and warn of crisis in mobile spectrum [6], [12], [27]. On the other hand, Wi-Fi with Internet connectivity is still very sparse due to its limited range and its coverage increases slowly due to high deployment cost.

Our approach: To support video watching on the go, we deploy kiosks equipped with Wi-Fi and a large storage at popular venues, such as restaurants, bus stops, railway stations and gas stations, where users can upload and download videos at a high speed and watch them on the go. To eliminate the need of significant effort of service personnel copying videos to every kiosk and to avoid the need for Internet deployment, we propose VideoFountain, which replicates videos through mobile users who travel between kiosks and can carry the content on their way. The kiosks will then serve the videos to interested customers. In order to make this system a reality, a number of significant challenges need to be addressed:

- An in-depth understanding of real user mobility is the necessary first step towards designing the system. While user mobility has been an extensively studied subject, existing research has been severely limited by the traces available. Most existing studies (*e.g.*, [19], [23], [30]) have conducted small-scale measurements to collect movement traces involving no more than 100 people over a few days. While insights gained from these studies are valuable,

it is not clear how well the observations apply to large populations for an extended period of time.

- When we have the flexibility of bootstrapping a small number of kiosks with new videos, the kiosk selection for the bootstrap has significant impact on the speed and delivery rate of the subsequent dissemination.
- We then need to deliver the videos from the bootstrapped kiosks to other interested ones. We propose to leverage users who travel between these locations anyway to carry the videos. This is a new kind of disruption tolerant network (DTN). Unlike most existing DTN research, which focuses on opportunistic communications between mobile nodes, the network in our context consists of two different types of nodes: mobile users and static venues. It calls for a new routing design. The static venues serve as useful rendezvous points and make the mobility prediction easier due to aggregation. How to effectively leverage such rendezvous points to facilitate routing and content distribution is an interesting issue.
- There are also many practical issues, such as how to incentivize users to forward videos that they may not be interested in viewing, how to protect copyright, how to preserve integrity of videos and prevent malware, and how to perform service discovery.

To address the first challenge, we analyze traces of Location-Based Social Networks (LBSNs), which have experienced an explosive growth in popularity (Section II). Foursquare [14], the most popular LBSN, has over 30 million users with 3 billion *check-ins* as of January 2013 [13]. Major social networking sites like Facebook, Twitter, and Google+ have also added location-based features into their services. LBSN data is particularly useful to study the potential of our system for the following reasons. First, all the check-ins are performed by mobile devices and they represent when and where people use their devices. These devices will be the targets for our video download. Second, users stay at these locations for a while (at least long enough to check-in manually), which is likely to be sufficient for downloading videos. Third, the check-in trace is massive: it contains mobility patterns of millions of users spanning across the world.

To address the second challenge, we study a range of algorithms to determine initial bootstrap kiosk locations (Section III-A). The algorithms use information such as venue popularity and the number of people traveling between different venues, which is readily available from the mobility traces.

To address the third challenge, we examine a variety of routing schemes for DTNs. We develop several new routing metrics to quantify the utility of the next hop and direct traffic towards the next hops that have high marginal utility

City	Users	Places	Check-ins
Austin, TX	13,718	836	51,638
London, UK	16,426	621	36,679
Manhattan, NY	115,605	5,824	481,976
Paris, France	11,197	504	25,341
Seoul, Korea	9,158	457	20,255

TABLE I

FOURSQUARE TRACES COLLECTED FROM JAN. 2 TO FEB. 6, 2012.

City	Users	Places	Check-ins
Austin, TX	25,515	32,086	1,422,987
Manhattan, NY	25,848	34,492	618,805

TABLE II

GOWALLA TRACES COLLECTED FROM JAN. 2009 TO DEC. 2011.

(Section III-B).

Then we introduce the methodology to evaluate the proposed system (Section IV) and measure its performance through trace-driven simulations (Section V).

Finally, we address a series of practical issues, such as incentivizing users by offering discounts, applying message authentication to ensure the integrity of the videos, leveraging digital right management tool to manage the copyright of the videos, and performing service discovery (Section VI).

II. ANALYZING USER MOBILITY

We first describe our data collection. Then we analyze venue popularity distribution and stability, the number of people traveling between venues, link capacity between venues, and the ability of predicting users' future check-ins. All of these metrics are important to the design of a mobile content distribution system.

A. Data Description

We collect and analyze data from Foursquare (4SQ) and Gowalla (GWL), two of the most popular LBSNs. Both services provide access to the data through open APIs ([14], [18]), which we use to collect the data. Traces we collected contain: (i) user information, (ii) detailed venue information (e.g., venue name, latitude and longitude, venue categories), (iii) check-in information (e.g., user identifier, venue identifier, time stamp). Tables I and II show the trace statistics. All the following results unless otherwise mentioned are based on the check-ins during one month (January 2, 2012 - February 6, 2012) in Foursquare and three years (January 2009 - December 2011) in Gowalla.

Reasons for using LBSN traces: These traces are large scale and contain millions of users over an extended period of time. In addition to their sheer sizes, they have several important characteristics that make them appropriate for understanding the potential of VideoFountain: (i) The LBSN traces only consist of smartphone users, who are the potential customers of VideoFountain. (ii) Check-ins are voluntary and thus users only check-in when they are 'idle' to use their smartphones, similar to what VideoFountain targets. (iii) The traces naturally capture different properties of different venues. In some venues, users tend to be busy and they are less likely to check-in, while in other venues users have more free time to check-in. When

City (Dataset)	Top 1%	Top 2%	Top 5%
Austin (4SQ)	67%	80%	67%
Manhattan (4SQ)	65%	69%	59%
Austin (GWL)	100%	100%	67%
Manhattan (GWL)	100%	100%	100%

TABLE III

OVERLAP OF TOP $K\%$ VENUES ACROSS SNAPSHOTS.

VideoFountain is successful, users may have more incentives to use their smartphones, which will help to further improve the performance of video distribution.

Our traces also have limitations: (i) check-in information does not come with the *check-out* time. So we assume that the dwell time follows an exponential distribution according to the observations from ([46], [16]). (ii) As check-ins are voluntary, we may miss opportunities of video transfer. But we avoid over-estimation of the network connectivity.

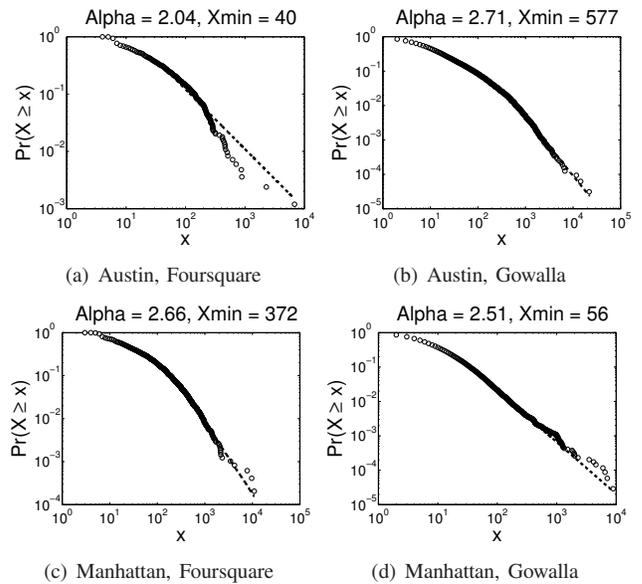


Fig. 1. Fitting of number of check-ins per venue to a Zipf-like distribution.

B. Venue Popularity

We first look at the popularity of venues, since VideoFountain will target the popular venues as the candidate locations of kiosks.

Zipf-like distribution: We plot the complementary CDF of the number of check-ins per venue in Figure 1, where y-axis is the probability of having more than x check-ins at the venue. We fit the curve using the maximum likelihood fitting method developed in [9]. All the points starting from X_{min} are used for fitting. $X_{min} = 1$ indicates that the fit is good even for the initial portion of the curve. We see the curves fit closely with a line on the log-log scale, indicating a Zipf-like distribution. The exponent α varies between 2.04 and 2.71. This means the popular venues get much higher number of visits and will be especially useful for video distribution.

Temporal stability: We compare the top $K\%$ venues in one snapshot against those in the next snapshot, where each snapshot lasts two weeks. Table III summarizes the results and shows the top few venues are relatively stable over time.

Spatial distribution: Figure 2 shows the CDF of pair-wise distances between top 20 and 100 popular venues. Only 6%-14% of top 20 venues are within 1 km from each other and only 3%-11% of top 100 venues are within 1 km distance. Hence the popular venues are spread across the city.

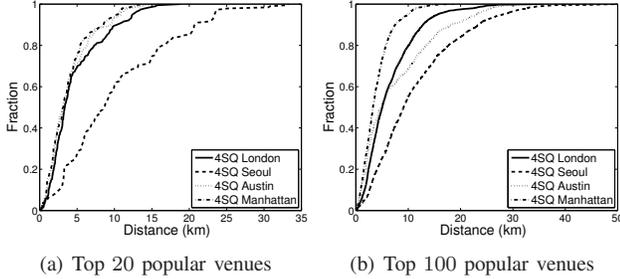


Fig. 2. CDF of inter venue distances.

C. Degree of Separation Between Venues

We further compute the degree of separation between venues, *i.e.*, the smallest number of hops a video needs to traverse between any two venues. More specifically, we construct a spatio-temporal graph according to the mobility traces as follows. Suppose there are N venues and K snapshots. We create $N \times K$ nodes, where each node (v, t) represents a venue v at snapshot t . There are two ways of creating edges. To be optimistic, we connect (v, t) with (v', t) if there is at least one user moving from v to v' during t . This is optimistic because in reality user may reach v' around the end of t , hence movements out of v' during t may not be able to relay the content brought to v' via this edge. We also create a conservative spatio-temporal graph by connecting two nodes (v, t) and $(v', t+1)$ if there is at least one user moving from v to v' during t . This conservatively assumes all users reaching v' at the end of snapshot. It does not utilize the users leaving in t to carry data arriving in t even though sometimes this is possible. We then compute the shortest path distance between any pairs of venues.

First, the graph is constructed using the first two week data for venues with at least 25 check-ins in Foursquare. We use daily snapshots (*i.e.*, t is a day). Figure 3(a) and (b) show the degree of separation in the optimistic and conservative cases, respectively. The percentage number in parenthesis indicates the fraction of reachable pairs. As expected, in the optimistic case the degree of separation is lower and connectivity is higher (85%-100% connected pairs in optimistic case vs. 61%-90% in conservative case). In the optimistic case, 28%-63% pairs are within two hops and 77%-93% pairs are within three hops. Even in conservative case 22%-62% pairs are within two hops and 62%-84% pairs are within three hops. Since most venue pairs are within 2-3 hops, routing in our system is relatively easy.

D. Human Traffic Matrices Between Venues

Definition: To understand aggregated user movements at a macroscopic level, we define a human matrix $T(a, b, t)$ as the number of users moving from a venue a to another venue b

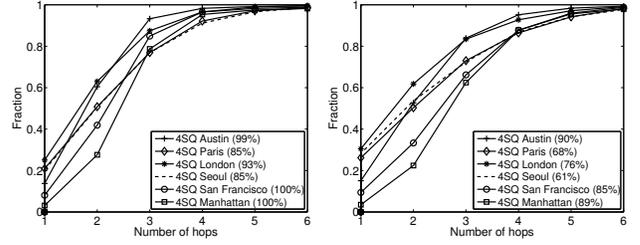


Fig. 3. Degree of separation among venues.

during a day t . If a user has visited A then B , we consider it as a one unit movement from A to B . Human TMs are derived by aggregating all user movements in a given day.

Zipf-like distribution: We plot the complementary CDF of TM elements in log-log scale in Figure 4, where y-axis is the probability of having more than x users moving from one location to another. We further fit the curve using the maximum likelihood fitting as done in Section II-B.

User movements closely follow a line in the log-log scale with the slope α ranging from 2.59 to 3.07. Such a close fit suggests the inter-venue traffic closely follows a Zipf-like distribution (*i.e.*, the volume of users moving between venues decreases exponentially with the exponent of α), which shows that popular venues have much more traffic.

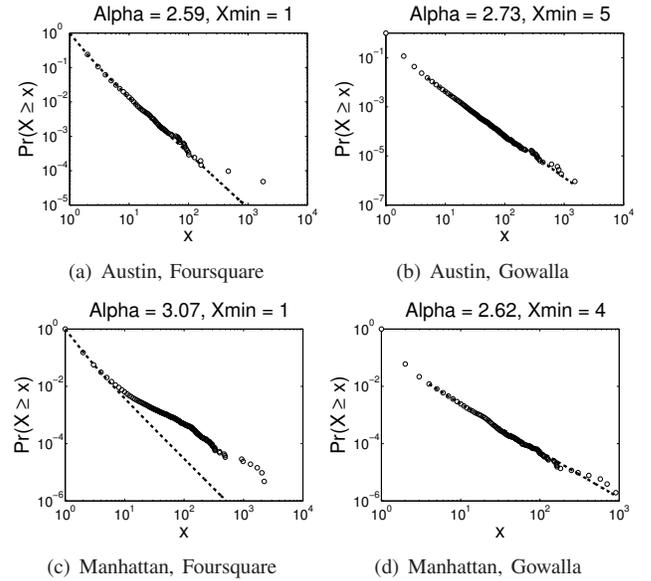


Fig. 4. Fitting of traffic matrices to a Zipf-like distribution.

Temporal stability: Next we analyze the temporal stability of TMs, *i.e.*, how much human traffic changes over time. We quantify the temporal stability using the following metric: Normalized Change in Traffic (NCT), defined as follows:

$$NCT(a, b, t) = \frac{|(T(a, b, t) - T(a, b, t-1))|}{\text{mean}_t T(a, b, t)} \quad (1)$$

where $T(a, b, t)$ is the TM element from a to b during time t . This metric is used to capture the change in the number of people moving between two venues. A similar metric has been used in [34].

For both traces, we only consider the TM elements between active venues with more than 25 check-ins during the first two weeks of the analysis period. We observe that NCT is 0 for 87.72%-92.9% cases in the daily TMs for the five cities in Foursquare traces. For Gowalla, we analyze TM snapshots from May 1, 2010 to June 4, 2010. We observe NCT is 0 for 92.6% and 89.7% cases in Austin and Manhattan, respectively. These numbers indicate that the TMs exhibit high temporal stability. Thus we can exploit such stability for predicting venue-to-venue traffic. This is critical for routing design.

E. Link Capacity Between Venues

To understand how much capacity we can achieve between different venues with human mobility, we perform the following analysis. We assume users can download content cached at the kiosks at Wi-Fi bandwidth L . If a user u stays at a venue v_i during time t_i and goes to another venue v_j and stays there for t_j , then he can carry $\min(t_i, t_j) \times L/2$ amount of data from v_i to v_j . We only consider half the time as we assume half the time is used for upload and the other half for download. The latency of this data transfer T is his inter-check-in time between v_i and v_j . The capacity of this transfer is equal to the total data carried divided by T . Assuming $L = 50$ Mbps (we can achieve even higher bandwidths using 802.11n), the average capacity of all inter-venue movements in Austin, London, Paris, Seoul is 0.65 Mbps in Foursquare. Even if the kiosks are connected via the Internet, this is a reasonable amount of bandwidth to complement Internet capacity as the average capacity is around 1.5 Mbps at many public hotspots [45].

F. Individual User Mobility Prediction

At a microscopic level, predicting an individual user's next check-in is essential to the design of content distribution, because content should be downloaded to users who are likely to visit the destination venue or at least make progress towards the destination. Below we quantify the prediction accuracy.

Methodology: Suppose a user u_i visits a venue v_j and we want to predict the probability of visiting a venue v_k right after current venue v_j . We consider three cases: (i) the user has previously visited v_j and then v_k right after that, (ii) the user has previous check-in history (at v_k and other venues) but visits v_j for the first time, and (iii) the user does not have any check-in history. For the case (i), we can naturally estimate his next check-in probability at v_k as the ratio between the number of times that the user visits v_k right after visiting v_j and the total number of times that the user visits v_j . For the case (ii), we estimate his next check-in probability at v_k as the ratio between the number of times that the user visits v_k and the total number of his check-ins at all venues. For case (iii), since we do not have per user information, we estimate user's next check-in using all other users' check-ins. The probability of next check-in at v_k is the ratio between the total number of times any user visits v_k right after visiting v_j and the total number of times any user visits v_j .

City (Dataset)	Accuracy (largest)	Accuracy (weighted)
Austin (4SQ)	24.17%	22.67%
London (4SQ)	40.70%	38.97%
Manhattan (4SQ)	21.91%	18.38%
Paris (4SQ)	55.46%	53.59%
Seoul (4SQ)	47.79%	46.56%
Austin (GWL)	26.39%	22.42%
Manhattan (GWL)	41.40%	38.98%

TABLE IV
ACCURACY OF NEXT CHECK-IN PREDICTION.

Evaluation: Our evaluation focuses on the check-ins at the popular venues, *i.e.*, venues that have at least 25 check-ins in the first two weeks in the analysis period. We use a one month trace for each dataset: January 2012 for Foursquare and May 2010 for Gowalla. Table IV shows the prediction accuracy. We find that predicting the next check-ins as the ones with the highest probability of visit is more accurate than predicting based on the weighted average of their visit probabilities (*e.g.*, if a user visits venue v_k and we predict the probability of visiting v_k as 0.8, then the prediction error is 0.2). So we predict based on the highest probability of visit in the remaining evaluation.

G. Summary of Analyses

To summarize, our analyses reveal the following important characteristics of LBSN traces:

- The popularity of venues exhibits Zipf-like distribution with a sharp decay. This indicates significant value of placing content at popular venues.
- The venue popularity is stable over time, which is useful for a practical system to leverage.
- Degree of separation between venues is relatively small (*e.g.*, 2-3 hops) between popular venues. The smaller degree of separation significantly eases the routing design.
- Analysis of aggregated human traffic matrices indicates Zipf-like distribution and strong temporal stability.
- Link capacity we can achieve with human mobility is significant with respect to typical Internet bandwidth and therefore we can leverage it to transfer content between these venues.
- We predict individual users' next check-ins, which is very useful for content distribution.

III. CONTENT DISTRIBUTION VIA HUMAN MOBILITY

In this section, we describe algorithms to determine the initial placement of content and routing schemes to distribute the content from the initial locations to the desired destinations through human mobility.

A. Initial Content Placement

Certain popular content can be placed at a small subset of venues through traditional channels (*e.g.*, the Internet and postal services) to facilitate further delivery. When that is possible, the selection of these seed venues has significant

impact on the performance of subsequent distribution. We use the following heuristics for placing the seeds. Note that our system does not depend on seed selection and can work without it (e.g., in case of user generated videos). We choose to do it for popular content as it improves performance when we can strategically bootstrap content to a few venues.

- Random placement (*Random*): It randomly selects the required number of seeds from all participating venues.
- Popularity based placement (*Popularity*): It selects the venues that are most popular (i.e., with the largest numbers of check-ins). The intuition of this placement is that popular venues have more visits and can replicate the content to more users, which increases the likelihood of successful delivery.
- Utility based placement (*Utility*): We place the seeds to maximize/minimize a certain utility function (e.g., the average distance to the closest seed, the average number of hops to the closest seed, the total human traffic to the destination venues). We use a simple greedy approach to compute the utility based placement: given a utility function, we pick the first seed that gives the highest utility; until the required number of seeds is reached, we iteratively pick the next seed such that in conjunction with the seeds selected so far it gives the best utility. Note that the performance of utility based placement depends on which utility function is considered, and thus we should use the utility function that best captures the performance of content distribution. We will discuss and identify a few important metrics in Section III-B.

B. Routing in VideoFountain

In VideoFountain, we leverage user mobility to distribute the content from the seeds to the requesting venues. It is similar to routing in DTNs, where mobile nodes serve as relays and communication happens during the contacts. Different from traditional DTNs focused on routing between mobile users, routing in VideoFountain involves two types of nodes: mobile users and static venues, where venues serve as rendezvous points. Therefore an effective routing design should leverage the following unique characteristics of the venues: (i) stable popularity with Zipf-like distribution, (ii) stable human traffic between venues, which decays rapidly with rankings, and (iii) the predictability of users' next check-in venues.

Problem formulation: There is a set of videos to be distributed. Let i denote the index of the video. Our goal is to maximize the number of videos going from the initial seeds S_i to the destination venues D_i . So this is routing for multi-source to multi-destination multi-commodity flows.

Existing routing schemes: There have been many routing schemes proposed in various DTN scenarios [25], [39], [44], [3]. Under lightly loaded cases, *Flooding* can be used to achieve high delivery rate under the excessive resource usage. *Epidemic routing* [44] has been proposed to exchange content between mobile nodes while restricting the hop count and the

number of replications. The notion of utility based replication was first proposed by *RAPID* [3]. The concept is intuitive: every packet is associated with a utility and a node replicates packets from its buffer in a descending order of marginal utility. The effectiveness of such a routing scheme depends on the utility function.

Utility based replication: We apply the utility-based routing framework in VideoFountain. The framework works as follows: given a contact between a venue and a user, we first identify all potential downloads and uploads. We then compute the marginal utility of each potential replication, which is the improvement in the utility function if we complete this replication. The utility value of a flow i is computed based on the existing replicas of i at all venues and the utility of the current user or venue.

We first consider single destination cases and then generalize to multi-destination cases. The marginal utility of uploading a flow i to a venue v is defined as the amount of improvement in utility achieved with this replication, over the utility of the existing replicas of i (i.e., $U(\{E, v\}) - U(E)$), where $U(\{E, v\})$ is the utility when we replicate to the new node v while still keeping the set of existing replicas E and $U(E)$ is the utility of the existing replicas. Similarly, the marginal utility of downloading a flow i to a user u is defined as $U(\{E, next(u)\}) - U(E)$, where $next(u)$ is the predicted next check-in venue of u from the results of Section II-F.

In multi-destination case, for each destination there is a utility value, thus we maintain a utility vector for each flow, where one vector element corresponds to one destination in the flow. We consider the marginal utility for this flow as the sum of the (non-negative) marginal utilities of all destinations and make the routing decisions in the same way.

Once we have the marginal utility value of all potential uploads and downloads, we sort them based on the marginal utility divided by the flow size and start replicating from the ones with the highest marginal per-packet utility.

Utility functions: We consider the following utility functions:

- Minimize expected delay (*Delay*): The expected delay for the packet to arrive at its destination from the current location. We use a training phase to learn the expected delay between the venues from the traces.
- Minimize geographic distance (*Geo*): The distance between the packet's current location and its destination.
- Maximize human traffic (*Single-hop Traffic*): The number of users who directly moved from the current venue to the destination within a period.
- Maximize multi-hop traffic decayed with hop counts (*Multi-hop Traffic*): The above human traffic captures only the traffic on one-hop links. But multihop connectivity is useful and should not be ignored. However, given the same volume of traffic, a shorter hop is preferred. To capture this notion, we use the following metric: $\sum_{l=1}^K \beta^l T(l)$, where l is the hop count, $T(l)$ is the total volume of traffic with the hop count of l , and β is a decay factor to penalize longer paths.

Oracle routing: To compare our schemes with what can be the best achievable in VideoFountain, we try to compute the optimal throughput. To do so, we again construct a spatio-temporal graph as in Section II-C. Except that now the edges are added with the capacity that each user can carry. Suppose a user u stays at a venue v for time t_1 and at the next venue v' for time t_2 then the amount of traffic that u can carry from v to v' is $\min(t_1, t_2) * L/2$, where L is the wireless link capacity between the venue and the user. We divide by 2 since the user is assumed to upload for half the time and download for the other half. We construct optimistic and conservative graphs as described in Section II-C. We then use the Linear Programming technique from [47] to optimize the throughput achievable in our network. Note that this is still an approximation to the absolute optimal since the best division of time for upload and download may not be half and half. Furthermore there may be multiple users uploading and downloading simultaneously, so the link capacity may be shared. Nevertheless, this gives a good estimate to the optimal performance and we find that our utility based schemes perform reasonably in comparison with the oracle routing.

IV. EVALUATION METHODOLOGY

Trace-driven simulation: We evaluate the performance of VideoFountain with trace-driven simulations. We first get the check-in time stamps from the traces. Check-out time stamps are generated from an exponential distribution with 60 minutes as the mean, which is the average Wi-Fi session time in real traces [46], [16]. Each pair of check-in and check-out corresponds to one contact between a user and a venue.

Control channel: We assume a control channel through which all venues share the distribution status such as the utility values of different flows, which venues have what content, and whether the destinations already have the content so that they can be removed from the utility vectors. Similar assumption has been used in previous work [3]. We can realize this assumption using a thin control channel (*e.g.*, by leveraging mobile users' cellular connectivity).

Flow generation: Sources are generated based on the placement strategies in Section III-A. For each flow, destinations are randomly chosen among the venues that are not seeds. We generate the number of destinations for each flow using a uniform distribution, which can be up to 10% of the venues. We inject all the flows at the beginning of the simulation.

We then simulate the content distribution by playing the trace. For the utility-based schemes, upon every check-in at a venue, we update the potential uploads and downloads at the venue and compute their marginal utilities. Upon every check-out of a user, we compute what content is downloaded or uploaded by the user and update the utility values for all the involved flows. For *Flooding*, we always replicate the packets sequentially. *Epidemic* is similar to *Flooding* with limited hop count (two hops in our implementation).

We run the evaluation on the Foursquare trace from January 17 to January 30, 2012. For the utility functions that need

training data, we use data from January 2 to January 16, 2012 as the learning period. We compare the delivery rate (of complete videos) of different routing schemes. Note that while the delivery ratios with respect to the venues are sometimes low, the system is still useful since users can search for the venues where the deliveries are successful and only download videos from those venues.

We assume that a user device has 10 GB available space (which high-end smartphones already have and other phones can achieve via external microSD cards), and a venue has 1 TB. Each video consumes 1 GB space. Each venue has 50 Mbps wireless capacity. For users' next check-ins, we use the individual mobility prediction from Section II-F. We generate 50 flows in a run and report the average of three experiments for each configuration. All results presented are from Foursquare, as Gowalla results show similar trend.

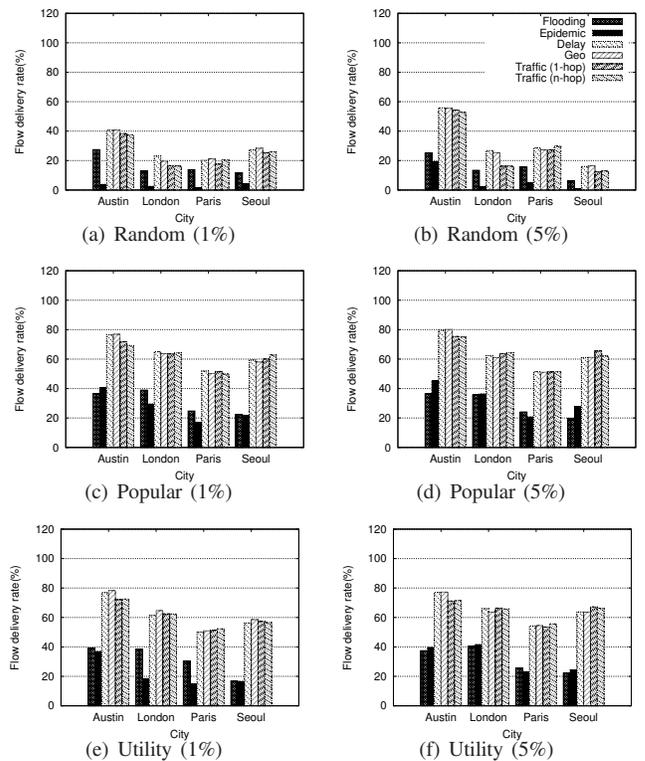


Fig. 5. Impact of seed placement.

V. PERFORMANCE RESULTS

Impact of seed placement: We first study how different placement schemes affect the routing performance. We compare random placement (*Random*), popularity based placement (*Popularity*), and utility based placement (*Utility*). We vary the number of seeds from 1% to 5% of all the venues. For *Utility* placement, we use *Multi-hop Traffic* as the utility function, *i.e.*, we maximize the *Multi-hop Traffic* from the seeds to the venues. We also tried other utility functions (*e.g.*, *Delay*, *Geo*, *Single-hop Traffic*) and found *Multi-hop Traffic* to be the best.

Figure 5 shows the evaluation result. We can see that placement is very important for all routing schemes. *Random* placement results in much lower delivery rate than *Popularity* placement and *Utility* placement. In London with 1% seeds,

Random placement achieves up to 23% delivery rate while Popularity placement and Utility placement achieve 77% and 64%, respectively. Using 1% seeds, we find the performance of Popularity placement is slightly better than Utility placement, (e.g., within 20%). However, Utility placement outperforms Popularity based scheme within 10% using 5% seeds. Comparing 1% seeds and 5% seeds, we find that with same placement scheme, having more seeds consistently improves the delivery rate as we would expect. For example, with Utility placement, 67% delivery rate is achieved in Seoul when 5% venues are seeds, while the number is 56% when only 1% seeds are used.

Figure 5 also shows that utility based replication performs the best. For example, we achieve 54% delivery rate in Paris with the utility based replication, while Flooding only delivers 27% and Epidemic only delivers 20%. We will use 5% seeds with Utility placement as the default setting in all following evaluation as it gives good performance with a reasonably small number of seed nodes.

Comparison with Oracle routing: We also compare practical routing schemes against Oracle routing strategy described in Section III-B. We use 50 flows with Utility placement and 5% seeds. We found that the utility based routing gives reasonable performance compared to Oracle routing. For example in Seoul, Oracle routing performs between 86% and 90% in terms of traffic delivery rate including partial deliveries, where Multi-hop Traffic achieves 69%. For Paris, Oracle routing achieves 81% - 87%, where the number for Multi-hop Traffic is 60%. For London, Oracle routing performs 84% - 85%, where Multi-hop Traffic achieves 69%.

Impact of number of flows: The number of flows determines how congested the network is and thus affects the performance. Figure 6 shows the performance as we vary the number of flows from 10 to 200. We find that as the number of flows increases, the delivery rate reduces for all the schemes. Flooding suffers the most from network congestion, with the average delivery rate going from 87% to 21% as the number of flows increases from 10 to 100. In comparison, the delivery rate of Multi-hop Traffic only drops from 88% to 54% with the same increase in the system load.

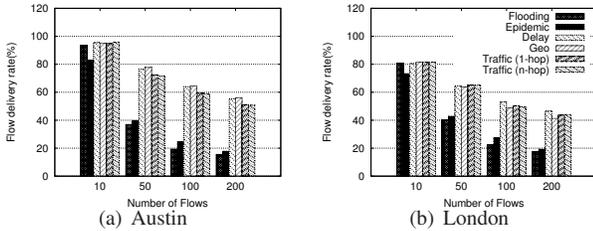


Fig. 6. Impact of number of flows.

Impact of video size: Video size also impacts the system performance. Figure 7 shows the delivery rate as we increase the video size from 100 MB to 5 GB. As expected, the delivery rate decreases as the file size increases. The results show that 500 MB videos can be delivered with good delivery rates (77% - 88%) using utility based replication. The delivery rate is around 70% even with larger 1 GB videos. This means that most mobile videos can be supported in our system.

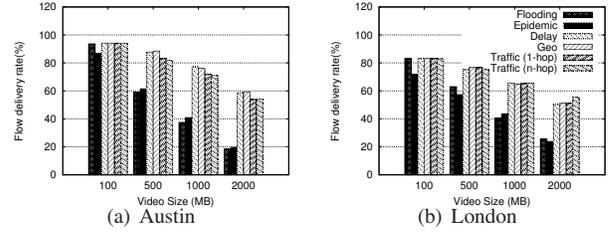


Fig. 7. Impact of video sizes.

Impact of contact time: We then measure the impact of the contact time in Figure 8 by varying the mean contact time from 10 to 90 minutes, while using exponential distribution. With 30-minute contact time, we can achieve 65% - 70% delivery. Increasing the contact time to 60 minutes (our default), the delivery rate increases to 67% - 78%.

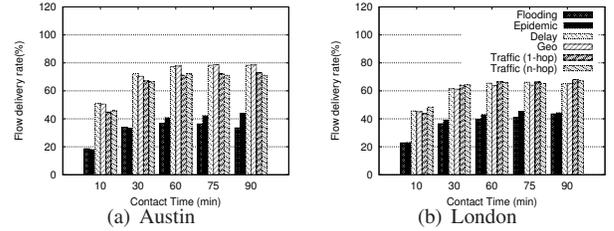


Fig. 8. Impact of contact time.

Impact of wireless capacity: Figure 9 shows the impact of the wireless capacity, while varying the capacity from 10 to 200 Mbps. We find that an increase in wireless capacity significantly improves the delivery rate. For example, delivery rate with 200 Mbps is 2.5 times the delivery rate when the capacity is 10 Mbps. This indicates that with wider deployment of the latest wireless technology, such as 802.11n in the future, the performance of VideoFountain can greatly improve. The utility based schemes consistently outperform other schemes under different wireless capacities. Specifically, the utility based schemes achieve about 28% higher delivery rate than other schemes when the wireless capacity is 100 Mbps and 23% higher when the capacity is 10 Mbps.

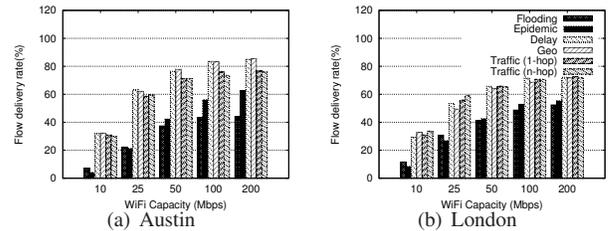


Fig. 9. Impact of wireless capacity.

Impact of user storage: We then study the impact of the user storage as we vary the numbers from 1 GB to 10 TB while assuming the venue has 50 Mbps capacity (Figure 10). We find that when the user storage is small (e.g. 1 GB), the delivery rate can be low. For example, we get only 30% traffic delivered with 1 GB and 76% with 20 GB when we use Multi-hop Traffic in Austin. However once the user storage is as large as 10 GB, further increasing the storage size does not significantly improve the delivery rate. That is because in that case the wireless link becomes the bottleneck, i.e., although we still have available space, the amount of data that can be

downloaded or uploaded is limited by the contact time and the wireless capacity.

We also compare different content replacement strategies in user storage. And we find that utility based replacement strategy always performs better than FIFO and random strategies. We omit the results for brevity.

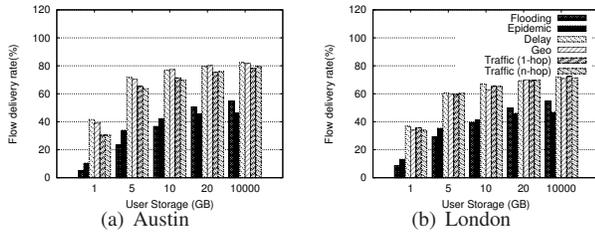


Fig. 10. Impact of user storage.

User mobility prediction: Lastly, we compare the mobility prediction scheme against the *Oracle* knowledge of next check-in in Figure 11. For all the utility based schemes, the performance difference between *Oracle* and our prediction scheme is between 18% and 38%. Given the fact that we use a very simple prediction scheme, the results are still promising.

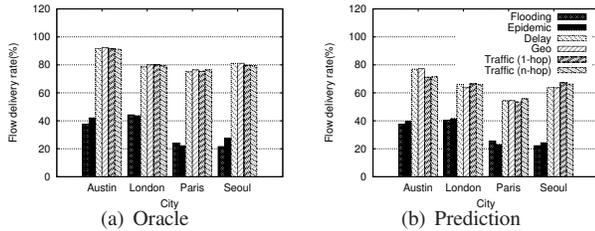


Fig. 11. Impact of user mobility prediction.

VI. PRACTICAL ISSUES

Providing incentive: Carrying and transferring content consumes storage and power of users' mobile devices. Therefore we should incentivize users to participate in forwarding. In VideoFountain, users are rewarded after every successful upload. Based on the amount of the uploaded traffic and the importance of the venue that the content is uploaded to, certain amount of reward points are issued to the user. The reward points can be used to consume video content in the future, or get discounts in participating businesses. To further reduce users' concern about the energy consumption, we can have charging stations co-located with the venues.

Content verification: Content uploaded by users may be corrupted or even modified maliciously. Thus it needs to be verified, even if it is only part of a video. Content verification also ensures only valid upload will be rewarded.

For partial content verification, we use segmented downloading methods [37]. Each video is divided into segments and a secure hash value is computed for each segment and signed by the content server. A segment is the smallest unit in data transfer. The integrity of a segment is verified by the hash value and authenticated by verifying the signature.

Copyright management: Some videos may have copyright. In this case, we need digital copyright management to ensure users can only view the content up to a limited number of times or a given amount of duration (*e.g.*, 24 hours) and

cannot share copies with others. There are many existing DRM (Digital Rights Management) solutions [10] that solve similar problems (*e.g.*, protecting the copyright of songs, movies, and e-books). They either use software solutions to provide DRM on general computing platforms (*e.g.*, used by Digiboo [11]), or tie the key to the identity of the hardware to prevent it from transferring to other devices (*e.g.*, used by Amazon Kindle). We can leverage the existing works to manage copyright.

Service discovery: If users are interested in the content themselves, then they can walk to a kiosk to select the videos to download. But if they are merely carrying the video, they may not have the motivation to voluntarily look for video kiosk to start the upload. We leverage existing service discovery protocols *e.g.*, UPnP [42], Apple Bonjour [2], to discover when to start the upload.

VII. RELATED WORK

Mobility analysis: There has been significant amount of work on mobility analysis using various location datasets, ranging from coarse grained locations from cellular networks [1], [32], [17] to finer grained locations from Wi-Fi networks [30], [38]. However, locations in these studies are approximated based on the locations of cell towers or Wi-Fi hotspots and therefore have large estimation errors.

Other studies use the exact GPS tracking [19], [21], [33], [20]. But most of these traces involve a smaller number of user movements and only a handful of studies examine movements of over 1000 users [17], [48].

Recently LBSN data has received attention as a good source for mobility analysis because of its scale. Interesting properties of human mobility are discovered such as periodic patterns [8], [7] and close relations between mobility and social interactions [36], [8].

Our work complements them as we go beyond the individual mobility and analyze aggregated movements between venues. Moreover, to the best of our knowledge, this paper is the first attempt to apply the properties learned from LBSN data to the human mobility based content dissemination domain.

Routing in DTNs: DTN routing has received wide attention (*e.g.*, [3], [35], [20], [4], [22], [39]). We leverage the insights from these works. Many of them are already discussed in Section III-B and here we review other works. In [22], the routing across a time varying connectivity graph is considered, where the dynamics may be known in advance. Authors in [15] also define spatio-temporal graphs to capture the time varying nature of the links.

In the vehicular DTN context, authors in [35] cluster taxis into concentration points (CPs) and show that the locations of CPs and flows of taxis between them are stable over time, which they then leverage for routing. VTube [26] proposes a vehicular content distribution using inter-vehicle contacts and roadside buffers which serve as rendezvous points. Our work uses a similar approach to use venues as the content aggregation points. But our work is different from them since human movement is quite different from vehicular movement.

Mobility based content distribution: Opportunistic networks leverage human mobility to distribute various content [31], [24]. Instead of relying on the infrastructure, this approach emphasizes communications among mobile devices. In this context, various properties of mobile nodes have been studied such as node degree [28], contact frequency [33], contact time [20], inter-contact time [5], temporal-spatial connectivity [40], social community [28], [29], to improve the content dissemination in user-to-user communications.

VIII. CONCLUSION

In this paper, we propose VideoFountain. We study the feasibility of such a system by first analyzing large-scale location-based social network traces. These traces show that the popularity of venues exhibits Zipf-like distribution and is stable over time; the degree of separation between the venues is small; the human traffic between venues is Zipf-like distributed and temporally stable; inter-venue link capacity is comparable to Internet bandwidth; we can predict users' next check-ins with a simple method. Based on the insights from the traces, we develop schemes to place initial content and routing algorithms to distribute the content from the initial seed venues to other destination venues. Our trace-driven simulation shows that the initial placement is important to the success of the system; with an appropriate placement, a simple utility-based routing scheme can perform very well due to the small degree of separation between the venues. These results suggest VideoFountain is a promising way of disseminating mobile videos in the near future.

Acknowledgements

This work is supported in part by NSF Grants CNS-0916106, CCF-1117009, and CCF-0916309. We are thankful to the anonymous reviewers for their constructive feedbacks.

REFERENCES

- [1] I. Akyildiz and W. Wang. The predictive user mobility profile framework for wireless multimedia networks. *IEEE/ACM Trans. on Networking*, 2004.
- [2] Apple Bonjour. <https://developer.apple.com/bonjour/>.
- [3] A. Balasubramanian, B. Levine, and A. Venkataramani. DTN routing as a resource allocation problem. In *Proc. of ACM SIGCOMM*, 2007.
- [4] J. Burgess, B. Gallagher, D. Jensen, and B. Levine. MaxProp: Routing for vehicle-based disruption-tolerant networks. In *Proc. of IEEE INFOCOM*, 2006.
- [5] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Trans. on Mobile Computing*, 2007.
- [6] B. X. Chen. Carriers warn of crisis in mobile spectrum. *New York Times*, Apr. 2012. <http://www.nytimes.com/2012/04/18/technology/mobile-carriers-warn-of-spectrum-crisis-others-see-hyperbole.html>.
- [7] Z. Cheng, J. Caverlee, K. Lee, and D. Z. Sui. Exploring millions of footprints in location sharing services. In *Proc. of AAAI ICWSM*, 2011.
- [8] E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proc. of KDD*, 2011.
- [9] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, Nov. 2009.
- [10] K. Coyle. The technology of rights: Digital rights management. 2003.
- [11] Digiboo. <http://www.digiboo.com/>.
- [12] L. Downes. Averting a spectrum disaster. *CNET*, Feb. 2012.
- [13] Foursquare by the numbers. <https://foursquare.com/about/>.
- [14] Foursquare API. <http://developer.foursquare.com/docs/>.
- [15] B. George, S. Kim, and S. Shekhar. Spatio-temporal network databases and routing algorithms: A summary of results. *Network*, 2007.
- [16] A. Ghosh, R. Jana, V. Ramaswami, J. Rowland, and N. K. Shankaranarayanan. Modeling and characterization of large-scale Wi-Fi traffic in public hot-spots. In *Proc. of IEEE INFOCOM*, 2011.
- [17] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *Nature*, June 2008.
- [18] Gowalla API. <http://gowalla.com/api/docs>.
- [19] Huggle project. <http://www.huggleproject.org>.
- [20] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and human mobility in conference environments. In *Proc. of ACM WDTN*, 2005.
- [21] P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: social-based forwarding in delay tolerant networks. In *Proc. of ACM MobiHoc*, 2008.
- [22] E. P. C. Jones, L. Li, J. K. Schmidtke, and P. A. S. Ward. Practical routing in delay-tolerant networks. *IEEE Trans. on Mobile Computing*, 2007.
- [23] Human mobility traces. <http://crawdad.cs.dartmouth.edu/ncsu/mobilitymodels>.
- [24] V. Lenders, M. May, G. Karlsson, and C. Wacha. Wireless ad hoc podcasting. *MC2R*, 2008.
- [25] A. Lindgren, A. Doria, and S. Olov. Probabilistic routing in intermittently connected networks. *MC2R*, 2003.
- [26] T. Luan, L. Cai, J. Chen, X. Shen, and F. Bai. VTube: Towards the media rich city life with autonomous vehicular content distribution. In *Proc. of IEEE SECON*, 2011.
- [27] R. Marwah. Spectrum crisis: What's there for carriers and customers? *The Motley Fool Blog Network*, May 2012.
- [28] A. Miklas, K. Gollu, K. Chan, S. Saroiu, K. Gummadi, and E. De Lara. Exploiting social interactions in mobile systems. In *Proc. of ACM UbiComp*, 2007.
- [29] A. Mtibaa, A. Chaintreau, J. LeBrun, E. Oliver, A. Pietilainen, and C. Diot. Are you moved by your social network application? In *Proc. of ACM WOSN*, 2008.
- [30] A. J. Nicholson and B. D. Noble. Breadcrumbs: forecasting mobile connectivity. In *Proc. of ACM MobiCom*, 2008.
- [31] J. Ott, E. Hyttiä, P. Lassila, T. Vaegs, and J. Kangasharju. Floating content: Information sharing in urban areas. In *Proc. of IEEE PerCom*, 2011.
- [32] P. Pathirana, A. Savkin, and S. Jha. Mobility modelling and trajectory prediction for cellular networks with mobile base stations. In *Proc. of ACM MobiHoc*, 2003.
- [33] A.-K. Pietiläinen and C. Diot. Dissemination in opportunistic social networks: the role of temporal communities. In *Proc. of ACM MobiHoc*, 2012.
- [34] S. Rallapalli, L. Qiu, Y. Zhang, and Y.-C. Chen. Exploiting temporal stability and low-rank structure for localization in mobile networks. In *Proc. of ACM MobiCom*, 2010.
- [35] N. Sarafijanovic-djukic, M. Piwski, and M. Grossglauser. Island hopping: Efficient mobility-assisted forwarding in partitioned networks. In *Proc. of IEEE SECON*, 2006.
- [36] S. Scellato, A. Noulas, R. Lambiotte, and C. Mascolo. Socio-spatial properties of online location-based social networks. In *Proc. of AAAI ICWSM*, 2011.
- [37] Segmented downloading. http://en.wikipedia.org/wiki/Segmented_downloading/.
- [38] L. Song, U. Deshpande, U. C. Kozat, D. Kotz, and R. Jain. Predictability of WLAN mobility and its effects on bandwidth provisioning. In *Proc. of IEEE INFOCOM*, 2006.
- [39] T. Spyropoulos, K. Psounis, and C. Raghavendra. Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In *Proc. of ACM WDTN*, 2005.
- [40] J. Tang, M. Musolesi, C. Mascolo, and V. Latora. Characterising temporal distance and reachability in mobile and online social networks. *SIGCOMM CCR*, 2010.
- [41] UK commute. http://news.bbc.co.uk/2/hi/uk_news/3085647.stm.
- [42] UPnP. <http://upnp.org/about/what-is-upnp/>.
- [43] US commute. <http://usgovinfo.about.com/od/censusandstatistics/a/commutetimes.htm>.
- [44] A. Vahdat and D. Becker. Epidemic routing for partially-connected ad hoc networks. *Duke Technical Report*, 2000.
- [45] Wi-Fi capacity. <http://www.popularmechanics.com/technology/how-to/tips/how-does-starbuck%ks-free-wi-fi-compare-to-other-hotspots>.
- [46] Wi-Fi traces. <http://crawdad.cs.dartmouth.edu/meta.php?name=ilesansfil/wifidog>.
- [47] J. Yuan, Z. Li, W. Yu, and B. Li. A cross-layer optimization framework for multicast in multi-hop wireless networks. In *Proc. of WICON*, 2005.
- [48] G. Zyba, G. Voelker, S. Ioannidis, and C. Diot. Dissemination in opportunistic mobile ad-hoc networks: The power of the crowd. In *Proc. of IEEE INFOCOM*, 2011.