2016 IEEE 27th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications - (PIMRC): MAC & Cross-Layer Design

Joint Optimization for Social Content Dissemination in Wireless Networks

Xiangnan Weng, John S. Baras
Institute for Systems Research and Department of Electrical and Computer Engineering
University of Maryland, College Park, Maryland 20742
Email: {wengxn, baras}@umd.edu

***Abstract*—Over the last decade, the success of social networks has significantly reshaped how information is disseminated among human beings in terms of scale, intensity and speed. Much effort has been dedicated to learning the patterns of such dissemination. However, as users become dominantly mobile, little is done to consider the impact of the wireless environment, in particular, how to optimize the system in the presence of capacity shortage, to better disseminate contents. In this paper, we propose a novel method, utilizing: 1) a hybrid system to handle active dissemination requests; and 2) predictions of dissemination dynamics obtained from the social network applications. This method could mitigate the performance degradation for content dissemination caused by wireless delivery delay. Results indicate that utilizing predictions (even coarse ones) of the dissemination process would benefit the process itself. This is due to improved scheduling prioritization at the wireless layer, which in turn improves the accuracy of the predictions with the help of feedback.**

***Index Terms*—Social network, wireless network, diffusion, resource allocation.**

I. Introduction

Social networks, which enable information sharing among people, have long existed on the Internet in various forms. However, it is not until the past decade that we witnessed their huge commercial and social success. Facebook, Twitter, YouTube, along with innumerable social networks, greatly facilitate information exchange for their users all across the world. Their successes rely heavily, if not solely, on content consumptions and user engagements. Instead of burdening users to actively seek and/or disseminate contents themselves, modern social networks take control and 'push' contents to users, utilizing recommendations based on user profiles (including but not limited to social relationships, engagement history, etc.). On one hand, social network providers are collecting more data about users than ever ('big data') to deliver highly engaging contents; on the other hand, users are sharing more data with their trusted social network providers for better experience. Therefore, researchers in this area are working diligently to enhance user experience by delivering high quality predictions of how much reward the social networks obtained when a user consumes a specific content. These predictions are based on numerous features of both users and contents. Metrics of rewards are different for different systems. For example, for ads systems, one possible reward metric is the revenue earned from displaying ads to users; for video subscription systems, the time users spent on the video.

In the real world, such rewards are hardly static, even though the dynamics governing evolutions of rewards remain an open question. Intuitively, just like any contagious disease, the evolutions are 'viral', in that they require certain form of interactions between the users, be it 'share', 'like', rate or comment. Particularly, researchers have already confirmed the effectiveness of friend's recommendations compared to those generated by the algorithms [1]. In fact, as the commercial success of Facebook demonstrates, users are more susceptible to the influence of other users than that of machine suggestions. In other words, users are likely to consume contents (even ads) that are disseminated by actual human users, especially their friends. Correspondingly, modern social networks significantly reduce the degree of separation between users [2], as well as reducing the 'infection' time with attention-grabbing notifications. These developments make interactions between users more pervasive and more intense.

Most social networks are reporting that their users are dominantly mobile, i.e. majority of the users access contents mainly on their mobile devices. With insufficient radio spectrum resources, the mobile applications struggle to optimize user experience. Certain preliminary work has been done concerning the joint user experience optimization incorporating social and wireless networks. In [3], we investigated the performance gained from a joint optimization approach for a single base station. The result is encouraging: the joint optimization approach outperforms existing layered solutions, especially when the wireless resource is insufficient. We extended the scheduling framework to scenarios with multiple base stations in [4], by introducing scalable two-phase scheduling consisted of 1) each base station makes their own decisions; 2) the system eliminates redundancy.

Unlike [3] and [4], in this paper, we are taking into account the dynamics from social networks. The reward values are considered time-variant, based on the dynamics between users. An intuitive example of the importance of the dynamics is that certain contents become increasingly rewarding and popular if specific users (e.g. celebrities, public figures) spread them. According to [5], content dissemination (or 'cascades') in social networks generally follows a diffusion process. The cascades are predictable with high precision based on user profiles. Fast algorithms were developed in [6] to infer network diffusion parameters for a continuous time diffusion model. However, the models used at social network layer failed to

978-1-5090-3254-9/16/$31.00 ©2016 IEEE

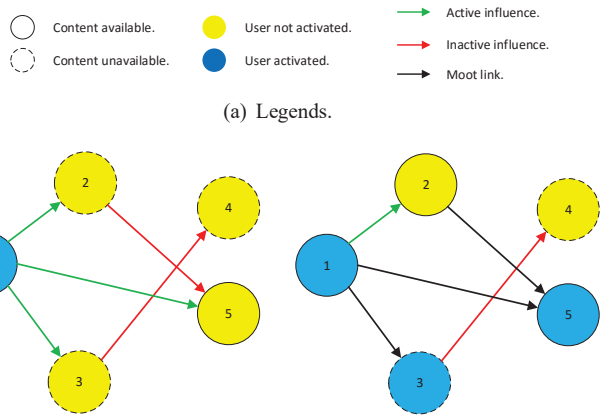

(a) Legends.

(b) Initial state: user 1 and user 5 have the contents available; user 1 is activated. Therefore, all the influence links originating from user 1 are active, but other links are inactive because their influencers are not activated.

(c) Possible evolution: user 5 is activated; user 2 has the content delivered. We are unable to claim the reward from user 3 yet because the content is not available. User 4 is not under any active influence because its influencer (user 3) has not yet acquired the content. Any link with both nodes activated is moot.

Fig. 1. Illustration of Content Dissemination and Delivery

consider the fact that mobile delivery delay can be a bottleneck and stopping point of the dissemination process, as illustrated in Fig.1.

In this paper, we investigate the content dissemination problem with capacity constraints in a system with centralized wireless infrastructure. The scheduling framework decides what contents to deliver to which users and how to transmit them. We investigate how severely the capacity shortage degrades the dissemination performance and propose a solution to mitigate such degradation. We show by simulations that both wireless and social networks perform well utilizing the predictions of social dynamics.

We introduce our system model and evaluation framework in Section II. We propose social-lookahead scheduling based on predictions of social dynamics in Section III. The performance is analyzed in Section IV and we summarize our conclusions in Section V.

II. Problem Formulation

A. General System Model

In this paper, we focus on a slotted (slot length T) single base station scenario with K multicast modes. We complete transmitting any content within one scheduling slot to avoid the management complexity of multicast groups. Channel information of all users are known to the system. We assume the wired bandwidth between base station and content server(s) is sufficient such that the base station could access contents as if they are stored locally. Additionally, the storage on user devices is sufficiently large to precache all the contents scheduled for delivery.

The objective of the system is to maximize overall user rewards delivered subject to capacity constraints. The contents

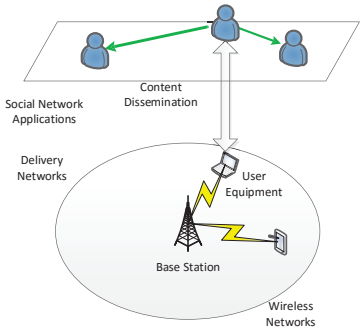

Fig. 2. System Model: The disseminating contents are propagated from influencer(s) to influencee(s) based on influence graphs (the boldness of the edge demonstrate its 'strength'), while the wireless networks are responsible to deliver them to user devices.

might be disseminating during the scheduling horizon, as illustrated in Fig.2, therefore, the reward value for transmitting content j to user i is time-variant $f_{ij}^{(t)} \in [0,1]$ (unlike [3] [4]).

The optimization problem could be formulated as in (1). Only contents with strictly-positive reward values could be scheduled for transmission.

$$\begin{aligned}
\max \quad & \sum_{t=1}^{T_H} \sum_{k} \sum_{i,j} \alpha_{ij}^{(t),k} f_{ij}^{(t)} \\
\text{s.t.} \quad & \alpha_{ij}^{(t),k} \in \{0,1\} && \forall i,j,k,t \\
& \sum_{t=1}^{T_H} \sum_{k} \alpha_{ij}^{(t),k} \leq 1 && \forall i,j \\
& s_j^{(t)} \cdot R_k \geq W_j && \forall \alpha_{ij}^{(t),k} = 1 \\
& \mathrm{SINR}_i^{(t)} \geq \mathrm{SINR}_k^{th} && \forall \alpha_{ij}^{(t),k} = 1 \\
& \sum_{j} s_j^{(t)} \leq B^{(t)} T && \forall t \\
& s_j^{(t)} \geq 0 && \forall j,t
\end{aligned} \tag{1}$$

The time-variance of rewards usually comes from the following two aspects:

1) social dynamics: users become more (or less) interested in a content than before, due to other users (either their own behaviors or interactions); and/or
2) information revelation: after gathering sufficient information regarding users (or their peer), the social networks may adjust previous predictions.

In this paper, we model the time-variance based on only social dynamics. Specifically, we focus on the content dissemination with the help of influence graphs.

B. Content Dissemination Without Delivery Delay

At the social layer, without delivery delay, we model the continuous-time reward change for a given content as activation process: reward for user i to consume the content (notation j omitted for simplicity) at time slot t is dependent on intrinsic reward value f_i and the user's binary activation state $\gamma_i(t)$.

$$\tilde{f}_i^{(t)} = f_i \cdot \tilde{\gamma}_i(t) \tag{2}$$

The activation process is based on the directed social graph between users $G = (V, E)$. The binary activation state of user u is dependent on the previous activation states of the influencer set $I_u^{\leftarrow}$.

$$I_u^{\leftarrow} = \{v : (v, u) \in E\} \tag{3}$$

Conversely, the influencee set is $I_u^{\rightarrow}$.

$$I_u^{\rightarrow} = \{v : (u, v) \in E\} \tag{4}$$

For generality, we do not restrict the influencer/influencee set to include only the explicit neighbors of the user, but note that due to system complexity, the influence propagation is usually calculated based on a small set of users.

$$\tilde{\gamma}_u(t) = \mathcal{A}_u \left(\{\tilde{\gamma}_v(\tau)\}_{\substack{v \in I_u^{\leftarrow} \\ \tau < t}} \right) \tag{5}$$

For simplicity, we assume that once users are activated, they will stay activated, or

$$\tilde{\gamma}_u(t_1) \leq \tilde{\gamma}_u(t_2), \forall t_1 < t_2 \tag{6}$$

We denote the activation time $\tilde{\mathcal{T}}_u$ as

$$\tilde{\gamma}_u(t) = \begin{cases} 0 & \forall t < \tilde{\mathcal{T}}_u \\ 1 & \forall t \geq \tilde{\mathcal{T}}_u \end{cases} \tag{7}$$

Based on the assumptions of independence and time-shift-invariance as in [6], we rewrite the state transition as

$$\tilde{\gamma}_u(t) = \max_{v \in I_u^{\leftarrow}} \mathcal{A}_{vu} \left(t - \tilde{\mathcal{T}}_v \right) \tag{8}$$

The mappings $\{\mathcal{A}(\cdot)\}$ are obtained from social network applications based on user profiles, and they are subject to different models.

The mappings satisfy causality

$$\mathcal{A}_{vu}(t) = 0, \forall t < 0 \tag{9}$$

In general cases, the mappings $\mathcal{A}(\cdot)$ are defined on a random space. In this paper, we assume that the pairwise activation time τ_{vu} for user v to 'infect'/activate user u is independent and follows exponential distribution $\tau_{vu} \sim \text{Exp}(\lambda_{vu})$. The diffusion parameter λ_{vu} reflects the influence rate of the influence link $v \rightarrow u$; the larger, the more quickly user u is activated due to user v, and vice versa. Trivially, user v has no influence on user u iff $\lambda_{uv} = 0$.

Then we could obtain the activation time of user u:

$$\tilde{\mathcal{T}}_u = \min_{v \in I_u^{\leftarrow}} \left(\tilde{\mathcal{T}}_v + \tau_{vu} \right) \tag{10}$$

Denote the social dissemination graph for a specific content as $G(V, E)$, with users as its nodes V, directed influence links as its edges $E = \{(u, v) : \lambda_{uv} > 0\}$. For time t, the binary user activation state vector $\tilde{\boldsymbol{\gamma}}_G(t) \in \{0, 1\}^{|V|}$ in the graph is:

$$\tilde{\boldsymbol{\gamma}}_G(t) = (\tilde{\gamma}_u(t))_{u \in V} \tag{11}$$

Then, the state transition process is essentially a Markov Chain with transition rate:

$$\lambda_{\tilde{\boldsymbol{\gamma}}_G} = \sum_{(u,v) \in E} \tilde{\gamma}_u \cdot (1 - \tilde{\gamma}_v) \cdot \lambda_{uv} \tag{12}$$

and transition probability

$$\mathbb{P}\left[\tilde{\boldsymbol{\gamma}}_G + \boldsymbol{e}_v \mid \tilde{\boldsymbol{\gamma}}_G\right] = \frac{1}{\lambda_{\tilde{\boldsymbol{\gamma}}_G}} \left(\sum_{u \in I_v^{\leftarrow}} \mathbf{1}(\tilde{\gamma}_u) \cdot [1 - \mathbf{1}(\tilde{\gamma}_v)] \, \lambda_{uv} \right) \tag{13}$$

Note that modern social network applications significantly increase the influence rate λ's (in a selective way), resulting in much shorter activation time for certain social connections. This is due to the facts that notifications on mobile devices are more visible to the users, and at the same time users are generally more attentive when they use mobile devices.

Table I summarizes definitions of parameters.

TABLE I
SUMMARY OF VARIABLES

Notations	Definition
M	Number of users.
N	Number of contents.
f_{ij}	Reward for delivering content j to user i.
$\gamma_{ij}^{(t)}$	Binary activation state for user i to consume content j at time slot t.
$\phi_{ij}^{(t)}$	Binary delivery state whether content j has been transmitted to user i before time slot t.
$I_u^{\leftarrow}$	Influencer set for user u.
$I_u^{\rightarrow}$	Influencee set for user u.
τ_{vu}	Influence time for user v to activate user u.
λ_{vu}	Diffusion parameter of (exponential) influence time for user v to activate user u.
$\alpha_{ij}^{(t)}$	Binary decision variable whether to transmit content j to user i at time slot t.
$B^{(t)}$	Total available bandwidth at time slot t.
$s_j^{(t)}$	Wireless resource allocated for content j at time slot t.
$\text{SINR}_i^{(t)}$	Signal-to-Interference-Noise ratio of user i at time slot t.
SINR_k^{th}	SINR threshold for multicast mode k.
W_j	Size of content j in bits.

III. SOCIAL-LOOKAHEAD SCHEDULING

Obviously, if we could leverage the information from the social network on how contents disseminate, the system might be able to precache smartly, utilize wireless resource efficiently, and provide better user experience.

Unfortunately, for general graphs, calculating the exact probability of activation before social scheduling horizon $\mathscr{T}_H$ is not quite easy. For a directed forest graph (in which each node has at most one parent, i.e. in-degree of each node is at most one $|I_u^{\leftarrow}| \leq 1, \forall u$), it yields to analytical form [7], provided that a node u's root ancestor is activated at time 0:

$$\mathbb{P}[\gamma_u(\mathscr{T}_H) = 1] = \sum_{\omega \in \theta_u} \left[\left(1 - e^{-\lambda_\omega \mathscr{T}_H}\right) \cdot \prod_{\substack{\omega' \in \theta_u \\ \omega' \neq \omega}} \frac{\lambda_{\omega'}}{\lambda_{\omega'} - \lambda_\omega} \right] \tag{14}$$

where θ_u is the path (set of directed edges) from node u's root ancestor to itself.

However, if the graph is moderately complicated, e.g. the graph has rings or alternative paths, it is computationally expensive to calculate the exact probability in its analytic form.

Require: User activation state $\boldsymbol{\gamma}$, directed and weighted diffusion graph $G = (V, E)$ with diffusion parameter λ_{uv} as the weight of the edge from u to v.

procedure ESTIMATEUSERACTIVATION($\mathscr{T}_H$)
  $H \leftarrow \emptyset$ ▷ H is a min-heap for tuple $(i, \hat{\mathcal{T}}_i)$ ordered on ascending estimated activation time $\hat{\mathcal{T}}_i$.
  $\hat{\boldsymbol{\gamma}} \leftarrow \boldsymbol{\gamma}$
  for $i \leftarrow \{i : \gamma_i = 1\}$ **do**
    H.insert $((i, 0))$
  while $H \neq \emptyset$ **do**
    $u, \hat{\mathcal{T}}_u \leftarrow H$.poll() ▷ Next activated node.
    $\hat{\gamma}_u \leftarrow 1$
    for $v \leftarrow \{v : v \in I_u^{\rightarrow}, \gamma'_v = 0\}$ **do**
      $\hat{\tau}_{uv} \leftarrow exp(\lambda_{uv})$
      $\hat{\mathcal{T}}'_v = \hat{\mathcal{T}}_u + \hat{\tau}_{uv}$
      if $\hat{\mathcal{T}}'_v \leq \mathscr{T}_H$ **then** ▷ Only update when the activation is within scheduling horizon.
        if $v \notin H$ **then**
          H.insert $\left((v, \hat{\mathcal{T}}'_v)\right)$
        else
          H.update $\left(v, \min(\hat{\mathcal{T}}_v, \hat{\mathcal{T}}'_v)\right)$
  return $\{\hat{\boldsymbol{\gamma}}\}$

Fig. 3. Monte Carlo Estimation of User Reward

To the best of our knowledge, there are no reported results regarding analytical forms of the solution.

To resolve the disadvantage, we employ the Monte Carlo method to approximate the activation state. The basic idea is to simulate a large number of instances and use instance average to substitute expectation. For each instance, we scan the dissemination path with breadth first search, using min-heap H (ordered by activation time) to store the nodes to be activated before scheduling horizon $\mathscr{T}_H$ and their respective activation times. The complexity of each heap operation is $\mathcal{O}(\log |H|)$, hence the overall complexity for running a simulation instance is $\mathcal{O}(|E| + |V| \log |V|)$. The algorithm is stated in Fig.3. Note that if the content has been delivered before (denoted by binary user delivery state ϕ_u, 0 not delivered and 1 delivered), we do not attribute the reward to the current transmission, but rather previous transmissions. Given the fact that the activation graph is sparse, we could run each instance relatively fast and different instances could run in parallel.

Obviously, we need to run sufficient number M_s of instances and use the instance result to predict activation states for each disseminating content given its current user activation states.

$$\hat{\boldsymbol{\gamma}} = \boldsymbol{\Psi}\left(\hat{\boldsymbol{\gamma}}^1, \ldots, \hat{\boldsymbol{\gamma}}^{M_s}\right) \tag{15}$$

where $\boldsymbol{\Psi}$ is a chosen detector function.

Trivially, social predictions always respect the current activations:

$$\hat{\gamma}_u \geq \gamma_u \tag{16}$$

A. Estimation Runtime

The runtime for each simulation instance depends on the number of edges evaluated. We present the statistics of runtimes for 100,000 simulation instances running on a machine with 2.2 GHz CPU in Table II, for graph parameters used in Section IV. Apparently, the simulations are fast and hence could be run real-time in actual systems.

TABLE II
MONTE CARLO INSTANCE RUNTIME (IN μs)

Min	P90	P95	P99	Max	Mean
2.8	58.8	77.4	137.6	900.9	32.6

B. Estimation Performance

The major metrics describing the performance of the estimations are false prediction ratio P_{F} and missed prediction ratio P_{M}.

$$P_{\mathrm{F}} = \frac{|\{i : \gamma_i = 0 \wedge \hat{\gamma}_i = 1\}|}{|\{i : \hat{\gamma}_i = 1\}|} \tag{17}$$

$$P_{\mathrm{M}} = \frac{|\{i : \gamma_i = 1 \wedge \hat{\gamma}_i = 0\}|}{|\{i : \gamma_i = 1\}|} \tag{18}$$

In this paper, we choose a simple hard-threshold detector with relative threshold χ in (19) for all users.

$$\Psi_i(\{\hat{\boldsymbol{\gamma}}^q\}) = \mathbf{1}\left(\sum_q \hat{\gamma}_i^q > \chi \cdot M_s\right) \tag{19}$$

We choose the number of estimation instances to be $M_s = 50$ and plot the performance for different relative threshold ratio χ in Fig.4 with different parameters (number of initial activated users $\|\boldsymbol{\gamma}_0\|_1$ and prediction horizon $\mathcal{T}_H$) in each sub graph.

We could conclude that:

1) the performance of Monte Carlo Estimation based on a simple detector is somewhat acceptable;
2) the more initial activated users ($\|\boldsymbol{\gamma}_0\|_1$) and/or the longer the prediction horizon ($\mathcal{T}_H$), the more precise Monte Carlo Estimation is;
3) trivially, when the threshold χ is higher, missed prediction is higher but false prediction is lower.

The lookahead reward for delivering content j to user i at slot t based on social prediction is thus:

$$\hat{f}_{ij}(t) = \hat{\gamma}_{ij}^{(t)} \cdot \left(1 - \phi_{ij}^{(t)}\right) \cdot f_{ij} \tag{20}$$

even though the reward might not be immediately claimed at the same slot of delivery.

Note that we only need to run these Monte Carlo simulations once new user activations occur.

Given the computation complexity and in light of [3], we obtain the scheduling decisions for each time slot t by solving mixed integer programming problem (21) for the optimal. We

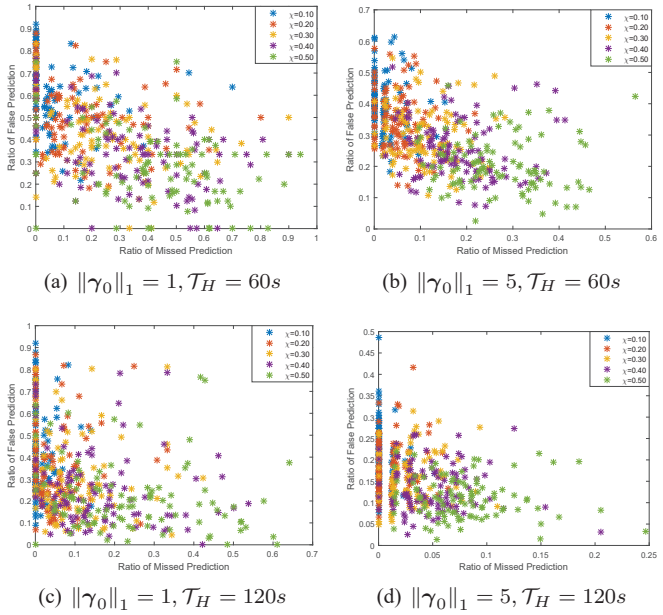

(a) $\|\gamma_0\|_1 = 1, \mathcal{T}_H = 60s$ (b) $\|\gamma_0\|_1 = 5, \mathcal{T}_H = 60s$

(c) $\|\gamma_0\|_1 = 1, \mathcal{T}_H = 120s$ (d) $\|\gamma_0\|_1 = 5, \mathcal{T}_H = 120s$

Fig. 4. Performance for Monte Carlo Estimation

could easily confirm that the scheduling solution satisfies all constraints in (1).

$$\begin{aligned} \max \quad & \sum_k \sum_{i,j} \bar{\alpha}_{jk}^{(t)} \cdot \bar{f}_{jk}^{(t)} \\ \text{s.t.} \quad & \bar{\alpha}_{jk}^{(t)} \in \{0,1\} && \forall j,k,t \\ & \sum_k \bar{\alpha}_{jk}^{(t)} \le 1 && \forall j,k \\ & s_j^{(t)} \cdot R_k \ge \alpha_{jk}^{(t)} \cdot W_j && \forall j,k,t \\ & \sum_j s_j^{(t)} \le B^{(t)} T && \forall t \\ & s_j^{(t)} \ge 0 && \forall j,t \end{aligned} \tag{21}$$

with lookahead reward $\bar{f}_{jk}^{(t)}$ for content j transmitted at mode k based on predicted activations in (22).

$$\bar{f}_{jk}^{(t)} = \sum_{i:\text{SINR}_i^{(t)} \ge \text{SINR}_k^{th}} \hat{f}_{ij}^{(t)} \tag{22}$$

Then we have the reverse mapping and state transitioning

$$\alpha_{ij}^{(t)} = \sum_{\substack{i,j:\hat{f}_{ij}^{(t)} > 0 \\ k:\text{SINR}_k^{th} \le \text{SINR}_i^{(t)}}} \bar{\alpha}_{jk}^{(t)} \tag{23}$$

$$\phi_{ij}^{(t+1)} = \max\left(\phi_{ij}^{(t)}, \alpha_{ij}^{(t)}\right), \forall i,j,t \tag{24}$$

C. Hybrid System and Predictions

To mitigate the performance degradation for disseminating contents due to wireless delivery delay, as in [3], we reuse the joint optimization scheduling framework that handles the hybrid requests ('push' for system recommendation and 'pull' for active user requests) to properly prioritize different types of requests. However, unlike [3], we impose no deadlines, but rather introduce unbounded additional rewards for active user requests.

Specifically, at time slot t, if user i is activated for content j, then there is an additional reward growing with time, controlled by weight factor ρ. The longer such requests are not delivered, the larger the (unbounded) additional reward.

$$f_{ij}^{(t)} = \gamma_{ij}^{(t)} \cdot \left(1 - \phi_{ij}^{(t)}\right) \cdot \left[f_{ij} + \rho\left(t - \lceil \mathcal{T}_{ij} \rceil\right)\right] \tag{25}$$

Of course, the system reduces to a 'push'-only system when $\rho = 0$.

Employing social predictions, for scheduling purposes, we rewrite (25) to properly prioritize:

$$\hat{f}_{ij}^{(t)} = \hat{\gamma}_{ij}^{(t)} \cdot \left(1 - \phi_{ij}^{(t)}\right) \cdot \left[f_{ij} + \rho \cdot \gamma_{ij}^{(t)} \cdot \left(t - \lceil \mathcal{T}_{ij} \rceil\right)\right] \tag{26}$$

It is easy to verify that

1) (20) and (26) are equivalent when $\rho = 0$.
2) Only unserved actual user requests would accumulate additional reward as time progresses.

IV. Simulations and Results

In this section, we conduct simulations to demonstrate the relationship between delay, total rewards obtained for content cascade and overall system reward delivered.

We simulate the performance based on synthetic Kronecker graphs [8], generated by SNAP [9]. In the simulations, we use the default parameters (with seed matrix $[0.9, 0.5; 0.5, 1]$) to generate Kronecker synthetic graph instances of size 256 (or 2^8) nodes. The set of initial activated users are selected as follows: 1) choose the number of activated users M_A uniformly from $[1, 10]$; 2) choose a total of M_A users randomly with equal probability from the top 20 influencers (i.e. users with largest out-degree $|I_i^{\rightarrow}|$). For simplicity, we assume that the reward for an activated user to consume a disseminating content is 1.

System level parameters are shown in Table III [10].

TABLE III
System Level Simulation Parameters

Simulation Parameter	Value
UE distribution	UEs dropped with uniform density within the macro coverage area.
Carrier frequency	2.0 GHz
Channel model	Typical Urban (TU)
Inter-site distance	1500 m
Noise power spectral density	-174 dBm/Hz
Macro BS transmit power	40 W (46 dBm)
Macrocell path loss model	128.1 + 37.6 $\log_{10} R$ (R in km)
Macrocell shadowing model	Log normal fading with std. 10 dB
Macro BS antenna gain	15 dBi

A. Scheduling Without Social Lookahead

In the simulations, we choose the following parameters:

1) the number of users $M = 300$, the number of contents $N = 600$;

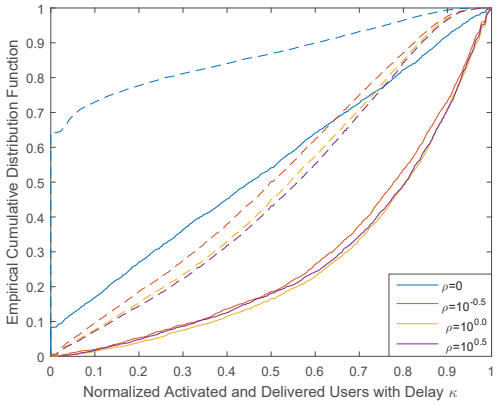

Fig. 5. Social Performance for Disseminating Contents

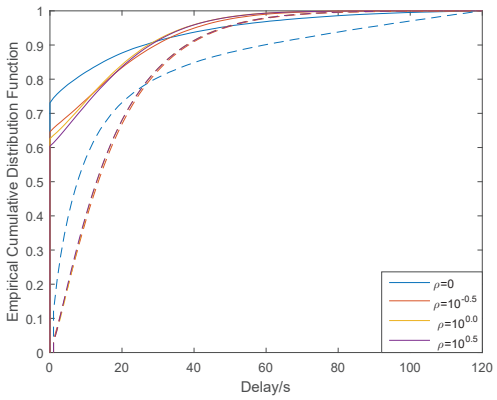

Fig. 6. Transmission Delay for Disseminating Contents

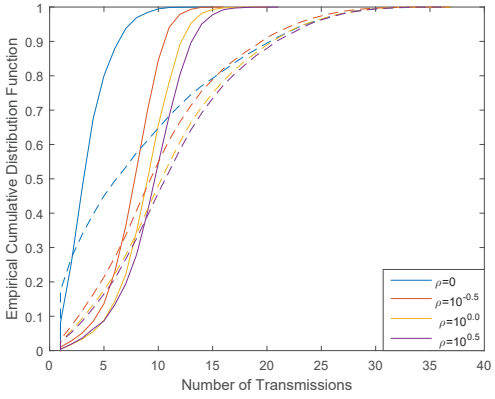

Fig. 7. Total Transmissions per Disseminating Content

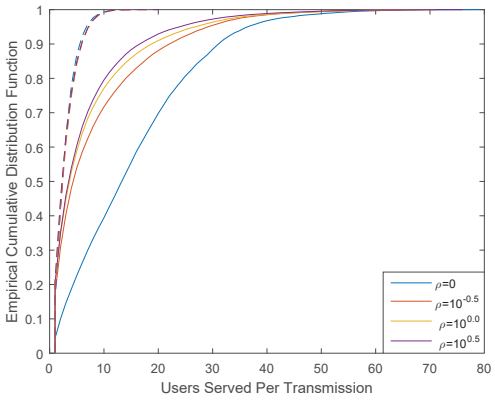

Fig. 8. Users per Transmission of Disseminating Contents

2) the number of disseminating contents $N_d = 0.05N = 30$;
3) scheduling horizon $\mathcal{T}_H = 120s$, slot length $T = 1s$;
4) system bandwidth $B = 20$ MHz.

We use ML-1M from MovieLens data [11] as the user reward values for the non-disseminating contents.

We map the nodes V_j of social dissemination graph G_j for content j randomly to the users.

All the curves in IV-A (without predictions from social networks) are plotted in dash lines, compared with those with predictions (IV-B) in solid lines.

The performance for disseminating contents is evaluated using the following metrics:

1) normalized activated and delivered users are plotted in Fig.5 (ideal situation yields $u(t-1)$);
2) delay in Fig.6 (ideal situation yields $u(t)$);
3) number of total transmissions per content in Fig.7;
4) number of users served per transmission in Fig.8.

The overall reward for the system is shown in Fig.9.

Without any prioritization ($\rho = 0$), the disseminating contents suffer from significant performance degradation.

B. Lookahead Scheduling

As described in Section III, we employ Monte Carlo Estimation to predict the activation states $\left\{\hat{\gamma}_{ij}^{(t)}\right\}$ based on the social dissemination graph without delivery delay. We use threshold value of $\xi = 0.3$ for estimation.

All the curves in IV-B (with predictions from social networks) are plotted in solid lines, compared with those without predictions (IV-A) in dash lines.

Compared with on-demand/no-lookahead scheduling (illustrated in dash lines), we note that with predictions (illustrated in solid lines):

1) The overall system reward is significantly improved

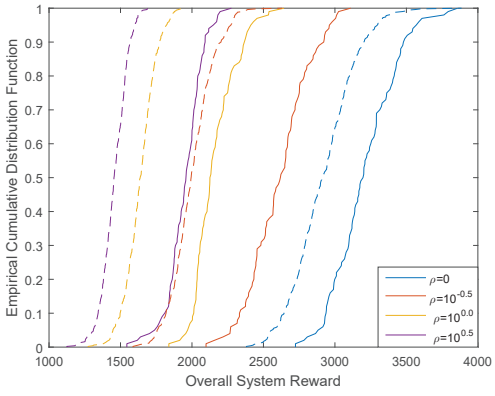

Fig. 9. Overall System Reward for Hybrid System

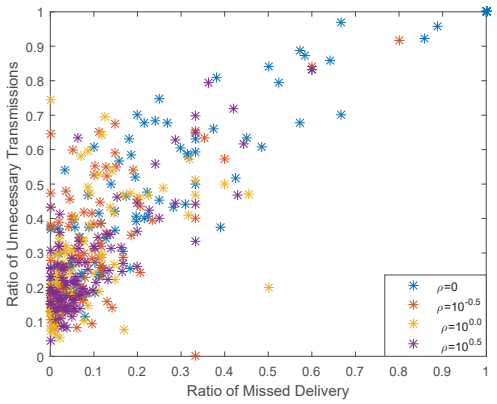

Fig. 10. Prediction Precision for Hybrid System

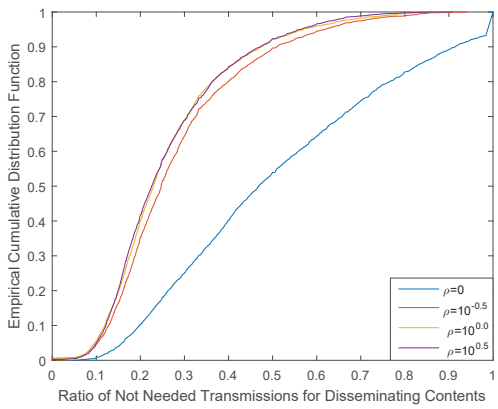

Fig. 11. Ratio of Unnecessary Transmissions for Hybrid System

(Fig.9).

2) The performance for disseminating contents is greatly improved (Fig.5). More specifically, the performance for 'push'-only system ($\rho = 0$) is comparable with hybrid system without predictions.
3) Delay performance for the disseminating contents is greatly improved and now delay could be 0 due to the precaching (Fig.6).
4) Redundant transmissions are greatly reduced (Fig.7) and the users served per transmission are increased (Fig.8). This is financially rewarding for the wireless system operators, since more efficient resource utilization usually brings about lower costs and thus higher profits.

We also plot the ratio of missed deliveries (in which users are activated but do not receive the contents) versus ratio of unnecessary transmissions (users are not activated but are delivered the contents) in Fig.10, from 100 randomly selected disseminating contents up until the scheduling horizon $\mathcal{T}_H$.

$$P'_{\mathrm{F}} = \frac{|\{i : \gamma_i = 0 \wedge \phi_i = 1\}|}{|\{i : \phi_i = 1\}|} \tag{27}$$

$$P'_{\mathrm{M}} = \frac{|\{i : \gamma_i = 1 \wedge \phi_i = 0\}|}{|\{i : \gamma_i = 1\}|} \tag{28}$$

In fact, the definitions (27)(28) are somewhat identical to the concept in social networks (17)(18), as plotted in Fig.4. The empirical distributions of the ratio of unnecessary transmissions are plotted in Fig.11. We observe that the hybrid scheduling framework greatly improves the prediction precision. This is most probably due to the fact that the hybrid system drives towards actual user requests, increasing the number of activations in general and triggering more re-evaluations.

V. Conclusions

In this paper, we investigate how to facilitate content dissemination with the wireless capacity constraints. We leverage the predictions of social dynamics to transmit the contents before the users actually request them. Results indicate that with the help of predictions, we could mitigate the performance degradation for disseminating contents and improve the overall system performance. Results also indicate that the hybrid system proposed in [3] greatly improve the performance of content dissemination, because they introduce proper prioritization for active user requests and thus prioritize social content dissemination.

Future work includes extending the analysis with more accurate predictions and extending the design to scenarios with multiple base stations.

Acknowledgment

Research partially supported by grants US AFOSR MURI FA9550-09-1-0538, AFOSR MURI FA-9550-10-1-0573, NSF CNS-1035655, NIST 70NANB11H148, and DARPA contract FA8750-14-C-0019 Phase 2.

References

[1] R. R. Sinha and K. Swearingen, "Comparing recommendations made by online systems and friends.," in *DELOS workshop: personalisation and recommender systems in digital libraries*, vol. 1, 2001.
[2] S. Edunov, C. Diuk, I. O. Filiz, S. Bhagat, and M. Burke, "Three and a half degrees of separation." https://research.facebook.com/blog/three-and-a-half-degrees-of-separation/, Feb. 2016.
[3] X. Weng and J. Baras, "Joint optimization for social content delivery in wireless networks," in *IEEE ICC 2016 - Communication QoS, Reliability and Modeling Symposium (ICC'16 CQRM)*, May 2016.
[4] X. Weng and J. Baras, "Joint Optimization for Social Content Delivery in Heterogeneous Wireless Networks," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt), 2016 14th International Symposium on*, May 2016.
[5] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, and J. Leskovec, "Can cascades be predicted?," in *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, pp. 925–936, 2014.
[6] M. Gomez Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1019–1028, ACM, 2010.
[7] M. Akkouchi, "On the convolution of exponential distributions," *J. Chungcheong Math. Soc*, vol. 21, no. 4, pp. 501–510, 2008.
[8] J. Leskovec and C. Faloutsos, "Scalable modeling of real graphs using kronecker multiplication," in *Proceedings of the 24th international conference on Machine learning*, pp. 497–504, ACM, 2007.
[9] J. Leskovec and R. Sosič, "SNAP: A general purpose network analysis and graph mining library in C++." http://snap.stanford.edu/snap, June 2014.
[10] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Frequency (RF) requirements for LTE Pico Node B," TR 36.931, 3rd Generation Partnership Project (3GPP), 9 2014.
[11] "MovieLens dataset." http://www.grouplens.org/data/, 2003.