A Privacy-Preserving Social-Aware Incentive System for Word-of-Mouth Advertisement Dissemination on Smart Mobile Devices

Wei Peng*, Feng Li†, Xukai Zou* and Jie Wu‡
*Department of Computer and Information Science
†Department of Computer, Information and Leadership Technology
Indiana University-Purdue University, Indianapolis, Indianapolis, IN, U.S.A.
‡Department of Computer and Information Science
Temple University, Philadelphia, PA, U.S.A.

***Abstract*—The recent penetration of smart mobile devices into the consumer market sets a stage for novel network applications. In particular, we envision a paradigm shift in the commercial advertising model facilitated by the widespread uses of these devices: advertisements circulate in a *word-of-mouth* fashion among device users and reach potential customers based on the users' knowledge about their contacts. In this paper, we identify two major challenges baffling the deployment of such an application: users' selfishness and their privacy concerns. We address the selfishness issue by proposing an incentive scheme which aligns users' interest with that of advertisers in a way that the users are willing to *fully* explore their social knowledge for *effective* advertisement deliveries—the emphasis is not only on users' participation but also on the extent and effectiveness of their contributions. We address the privacy concerns by designing a privacy-preserving evidence-collection mechanism, on which the incentive scheme is based. In addition, our design is 1) appealing to advertisers by guaranteeing effectiveness and controllability of the incentive dispensing and 2) robust against users' misbehaviors. We perceive incentive and enforcement as the keys to unlock the power of users' collective intelligence for effective information dissemination.**

***Keywords*-incentive system; mobile application; privacy-preserving; social-aware; user intelligence; word-of-mouth**

I. Introduction

The recent proliferation of smart mobile devices (e.g., smart cell phones) in the consumer market opens up opportunities for novel applications. These applications, powered by diverse connectivity options and abundant computing resources on the devices, increasingly emphasize the role of users over underlying technology.

This shift of focus introduces a subtle but significant change in the inter-device communication pattern. Communications are no longer restricted by a technological infrastructure; users' mobility and social-connection patterns bring forth numerous ad-hoc communication opportunities. These ad-hoc communications are rich in semantics. For example, a short-range Bluetooth connection implies device users' co-location and, depending on the location, connotes social relationships among the users, e.g., coworkers. These semantics, imbued with knowledge about device users, are aptly captured in the term *user intelligence*.

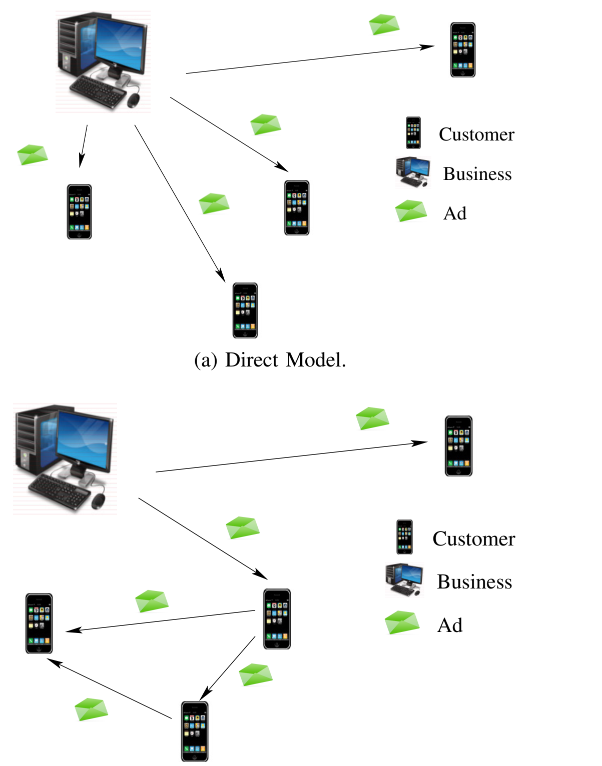

(a) Direct Model.

(b) Word-of-Mouth Model.

Fig. 1. Different advertising models. In the direct model (a), the advertisement is delivered directly from the business to the customers; customers play a *passive* role. In contrast, in the word-of-mouth model (b) facilitated by smart mobile devices, the advertisement is released to a (small) pool of initial customers and gets circulated in a customer-to-customer fashion; customers play an *active* role.

A promising application which benefits from the integration of user intelligence is *commercial advertisement dissemination*. In a conventional televised or roadside-billboard advertising campaign, advertisements are *indiscriminately* delivered to a *passive* audience. We call this approach the *direct model*. A recent U.S. patent [1] granted to Apple describes a location-based advertisement-delivery service in the direct model.

In contrast, with smart mobile devices, advertisements can be easily duplicated among the device users and reach the *target* audience based on the user intelligence; the users play an *active* part in delivering the advertisements. We call this approach the *word-of-mouth model*. The word-of-mouth model extends the direct model in terms of both reach and effectiveness. The contrast between these models is illustrated in Figure 1.

However, two challenges baffle the deployment of such an application. First, unlike the conventional advertising campaigns, in which activities are planned and executed by a single entity (the advertiser), advertisement dissemination on mobile devices largely depends on the cooperation of device users. Without the prospect of benefiting from participating, a user might refuse to join the program at the outset. This is called the *selfishness* problem.

The second challenge is users' *privacy concerns*. Because the advertisement propagation involves multiple users and their collective user intelligence, without protections in place, the users' private information might be divulged to a snoopy third party. An even more grave threat is the "prying eyes" among the users. Users' privacy needs to be protected from both outside and inside threats.

In this paper, we propose a *privacy-preserving social-aware incentive system* for word-of-mouth advertisement dissemination on smart mobile devices. We address the selfishness problem by proposing an incentive scheme which aligns users' interest with that of advertisers in a way that the users are willing to *fully* explore their user intelligence for *effective* advertisement deliveries—this is what we mean by a social-aware incentive system. We address the privacy concerns by designing a privacy-preserving evidence-collection mechanism in which users' participation is kept secret from prying parties during the advertisement propagation. The distinction of our incentive system is the emphasis on the extent and effectiveness of users' contributions in addition to their participation.

Moreover, our design has several other desirable features:

1) Rewards are dispensed only for *effective* advertisements (Section IV-A);
2) The incentive provider has control over the amount of resources allocated for the rewards (Section III-A);
3) The incentive system is robust against users' misbehaviors (Section III-C).

We organize the presentation in the following way. We formulate the problem into a network model and summarize the design objectives in Section II. A crucial mechanism is introduced in Section III, which paves the way to the incentive design discussed in Section IV. We evaluate our design with simulations in Section V. Related works are reviewed in Section VI. We conclude our discussion in Section VII.

II. Models

The incentive system is defined by three types of entities and the interactions among the entities.

The entity types are information sources, mobile device users, and a trusted third party which we call the *incentive authority*. Information sources could be the businesses which leverage the incentive system to disseminate advertisements for their products or services. Users are smart device owners who forward the advertisements. The incentive authority is trusted by both users and sources in the incentive process.

We propose a mechanism called *incentive ticket* to facilitate efficient advertisement propagation. The name is inspired by store-issued paper based coupons in the everyday life: a customer can redeem (and hence forfeit) a coupon in exchange for a price discount when paying for merchandise or services. However, these coupons are usually issued *directly* from the store to the customers, rather than exchanged between customers. In our design, on behalf of a information source, the incentive authority issues incentive tickets to users and encourages them to *duplicate* the incentive tickets to their friends. A user is willing to 1) *receive a incentive ticket* because he has the right to redeem it later for a price discount and 2) *duplicate the incentive ticket* to others because he might get rewarded by the shop for doing this.

To specify the design of our incentive scheme, we consider the following model. The *contact network* is a graph with nodes representing users and edges connecting each pair of users who share contact opportunities, e.g., they establish a short-range Bluetooth or Wi-Fi Direct connection between their devices. A user u is associated with a *redemption probability* p_u ($0 \leq p_u \leq 1$) representing the likelihood that u will redeem a incentive ticket he receives. We denote the degree of u in the network by k_u, i.e., u has k_u contacts.

In the contact network model, the user intelligence of a user u is characterized as the knowledge u has about his neighbors (who may be u's friends) $\mathcal{N}_u$: u knows p_v and k_v for every $v \in \mathcal{N}_u$. However, privacy concerns mandate that the incentive system should keep users' participation secret between pairs of users who do not share contact opportunities. This effectively confines the user intelligence to a *one-hop scope*: u *knows and only knows* p_v and k_v for $v \in \mathcal{N}_u$.

Because incentive ticket redemptions are associated with sales, the information source's expected profits from issuing a incentive ticket to a user u are estimated by u's redemption probability p_u: the larger p_u is, the higher the expected profits are. Therefore, the shop's interest can be interpreted as the extent to which the incentive ticket copies reach a population with high redemption probabilities.

We spell out the design objectives of the incentive system as follows:

- The rewards are dispensed in a way that users, equipped with the user intelligence and seeking to maximize their own interest, promote the shop's interest as well by duplicating the incentive ticket to a population *with high redemption probabilities*.
- Duplication activities between users are recorded on the incentive ticket (as the evidence for user to claim rewards) in a manner that users' participation is kept secret from snooping parties.
- The shop can impose a limit on the amount of resources allocated for the rewards.
- Users' misbehaviors should be detected and the offenders should be identified.

TABLE I
SYMBOLS AND NOTATIONS.

I	The incentive authority.
s	A shop.
u, v, w	Users.
p_u	User u's redemption probability.
k_u	The number of user u's contacts.
M	A text segment.
$M_1 \| M_2$	Concatenation of text segments.
C_n	Incentive ticket cached by n.
T_C	Front-page section of incentive ticket C.
W_C	Spray width of incentive ticket C.
L_C	Available slots in the authentication section of incentive ticket C.
K_n^+/K_n^-	n's public/private key.
$\{M\}_{K_n^-}$	n's digital signature on the *hash* of M.
$E_I(M)$	Encrypt M with I's public key.
x_n	A cryptographic nonce generated by n.
R_C	Reward amount for incentive ticket C.
$i_1, i_2, \cdots, i_l$	Identifiers in incentive ticket circulation chain.

III. INCENTIVE TICKET: THE MECHANISM

The incentive ticket mechanism, which includes a *format specification* and an *exchange protocol*, is the backbone of our design. It serves as the medium to bear advertisements and as the evidence to prove users' participation.

Table I lists all the symbols and notations used in the following discussions.

A. Format Specification

A incentive ticket consists of three sections.

- **Front-page** The content of this section is the advertisement along with an explanation of the reward program (e.g., product discount associated with a incentive ticket redemption). We denote it by the symbol T_C.
- **Restriction** This section contains two fields, W_C (spray width) and L_C (available slots), which together serve to regulate the incentive ticket C's duplication process. A user who receives a copy of C is allowed to duplicate it *at most* W_C times. L_C is the number of available slots in the *authentication* section.
- **Authentication** This section consists of slots used for authenticating incentive ticket duplications. The format of a slot is best explained with the incentive ticket exchange protocol (Section III-B).

B. Exchange Protocol

The incentive ticket exchange protocol coordinates communication and bookkeeping activities in three scenarios, namely, *issuance*, *duplication*, and *redemption*. Figure 2 shows the life cycle of a incentive ticket.

The protocol is built on top of a public key infrastructure (PKI). Each entity generates a public/private key pair. The incentive authority I's public key K_I^+ is known to all shops and users; so everyone can encrypt a message to I with K_I^+. We also assume that 1) I knows all parties' public keys and 2) a pair of contacting parties know each other's public key.

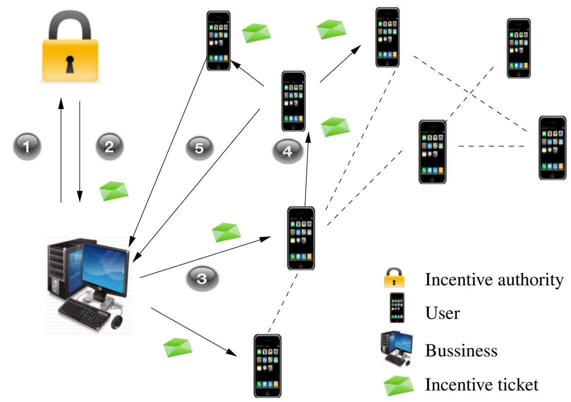

Fig. 2. Incentive ticket's life cycle. (1) The shop sends a incentive ticket request to the incentive authority; (2) The incentive authority issues a incentive ticket to the shop; (3) The shop offers the incentive ticket to its customers (users); (4) Users duplicate the incentive ticket; (5) A user redeems the incentive ticket at the shop.

1) Issuance: Since the incentive authority oversees the reward dispensing, the issuance of a incentive ticket originates from the shop s's request to the incentive authority I. s and I negotiate on the terms of the incentive ticket, which include the front-page section T_C and the duplication restrictions W_C and L_C. On agreement, I issues the following incentive ticket to s: $T_C, W_C|L_C, \emptyset$ ($\emptyset$ stands for an empty authentication section). Along with the incentive ticket is a signature $\{T_C|W_C|L_C|s\}_{K_I^-}$.

Upon receiving the incentive ticket, s checks that the negotiated terms have been correctly recorded by verifying the signature with I's public key K_I^+. Then, s generates a random number (a cryptographic nonce) x_s and caches the following incentive ticket

$$C_s = T_C, W_C|(L_C - 1), \quad E_I(\{T_C|W_C|L_C|s\}_{K_I^-}|x_s|I|s). \tag{1}$$

The decrement in L_C reflects the allocation of a slot for the encrypted content. The nonce x_s prevents the replay attack on the identifiers from a third party.

2) Duplication: Conceptually, a incentive ticket duplication happens in two different settings: 1) between the shop and a user and 2) between two users. Technically, as we will see below, there is only a *slight* difference between these settings: the shop *is not* subject to the spray width restriction W_C but a user *is*.

The shop determines how many copies of C_s it will offer to users based on the amount of resources allocated for rewards. Then, according to the plan, the shop sends incentive ticket offers to the users in the shop. A user u who accepts the offer will be given a copy of C_s along with a signature $\{C_s|u\}_{K_s^-}$.

First, u verifies the signature with s's public key K_s^+. Then, u generates a nonce x_u and caches the following incentive ticket

$$C_u = T_C, W_C|(L_C - 2), E_I(\{C_s|u\}_{K_s^-}|x_u|s|u) |E_I(\{T_C|W_C|L_C|s\}_{K_I^-}|x_s|I|s). \quad (2)$$

Again, L_C is decreased by one. Essentially, what u does is encrypting the received signature from s along with their identifiers and attaching the encrypted content to the incentive ticket.

Now, subject to the duplication restrictions W_C and L_C, u is allowed to duplicate his cached incentive ticket C_u to another user v in the same way as above (simply substitute s for u and u for v). A user who caches a incentive ticket with no available slot (i.e., $L_C = 0$) is forbidden to further duplicate the incentive ticket.

3) Redemption: A user can redeem his incentive ticket at the shop. The incentive ticket is subjected to a verification before being honored. Essentially, the incentive authority iteratively decrypts each slot in the authentication section and reconstructs the incentive ticket's *circulation chain* starting from the shop. Misbehaviors, if there are any, will be detected and reported.

As an illustration, let us consider how the incentive authority I reconstructs the circulation chain after u redeems C_u. I extracts and decrypts the first authentication slot in C_u and gets $\{C_s|u\}_{K_s^-}|x_s|s|u$. This indicates that u got the incentive ticket from s. With s's public key, I verifies the signature $\{C_s|u\}_{K_s^-}$ (I recovers C_s by removing the first authentication slot and adding one to the L_C field in C_u). In the same way, I recovers and verifies $I \to s$ from C_s. Then, I knows that s is the shop and reconstructs the circulation chain: $s \to u$.

C. Misbehavior Detection

It is instructive to consider how misbehaviors are thwarted by the exchange protocol.

1) ***Privacy Attack****:* A snooping party may find out who duplicate a incentive ticket by matching the encryption of an identifier pair (u and v) with those in the incentive ticket's authentication section. The nonce x introduces a linear factor into the complexity of this enumeration attack. If x can assume any of N values, the snooper will have to try all identifier pairs N times to find out u and v in the worst case. With a large nonce space, it is difficult to launch this attack.

2) ***Frame-up****:* A user is unable to frame either his (honest) predecessor or successor in the circulation chain for misbehaving. For example, let us consider a circulation chain $\cdots \to u \to v \to w \to \cdots$, in which u and w faithfully fulfill their duties but v does not.

v might remove a previous encrypted segment or replace his identifier in hope to frame u. But u's signature on $\{C_u|v\}_{K_u^-}$ protects u from being framed in these anomalies.

v might modify T_C or raise W_C in hope to get w into trouble. Indeed, w cannot detect v's malicious intent as long as v signs on the tampered incentive ticket. However, v's signature on the incentive ticket holds him responsible for the tampering. The innocent w will not be blamed.

3) ***Collusion****:* Suppose u sends C_u along with his signature $\{C_u|v\}_{K_u^-}$ to v. But instead of following the protocol, v tampers with the incentive ticket and finds another user w who is willing to sign the tampered incentive ticket for v. Now, v encrypts w's signature and their identifiers (to I) and attaches the encrypted content to the incentive ticket.

In the redemption verification, v will not be detected as a cheater. However, in the next iteration, I detects a mismatch in the current receiver and the sender in the previous iteration; an anomaly is reported. Because w is the one who signs off this anomalous incentive ticket and *no one is detected cheating before him*, w, rather than (the innocent) u, will be reported as a cheater and be punished accordingly. Since no rational user wants to be a scapegoat for others, w will refuse to collude with v in the first place.

The morale of the above analysis is that signatures hold users accountable and encryption keeps identifiers concealed. Abiding by the protocol is in each user's best interest. Therefore, the incentive authority can assure that the circulation chain reconstructed from a redeemed incentive ticket faithfully reflects the incentive ticket's circulation among users.

IV. Incentive Design

The incentive ticket mechanism testifies on *participating* users' faithfulness; it alone, however, is insufficient to guarantee that users will be *interested* in participating. The incentive ticket's benefits (e.g., discount) encourage *redeeming* rather than *duplicating* the incentive ticket. The whole point of exploring the user intelligence will be lost if users do not duplicate the incentive ticket to others. Therefore, *extra* rewards should be provided by the shop to the users who duplicate a incentive ticket. Moreover, the reward dispensing should be designed in such a way that a user, seeking to maximize his own interest, promotes the shop's interest as well.

This leads to three questions: 1) *Where* are the rewards from? 2) *Who* should be rewarded? 3) *How* should the rewards be dispensed?

A. Reward Origination

In our design, the rewards come from the shop's sales profits associated with incentive ticket redemptions. More precisely, prior to a incentive ticket's issuance, the shop negotiates with the incentive authority I on the amount of rewards allocated for each valid redemption. The amount depends on the profits and the shop's intended incentive strength. I only needs to know the reward amount R_C. Note that R_C is only an *upper limit*: I is free to dispense an amount of rewards less than R_C.

Two implications follow from this design choice.

First, shops are exempted from the risks associated with advertisement's effectiveness: rewards are dispensed only if users' effort leads to incentive ticket redemptions (and profits for the shops). This is appealing to the shops: they do not waste money on ineffective advertisements.

Second, incentive ticket circulation should be regulated. The shop must honor every legitimate copy of a incentive ticket. Meanwhile, duplications are not centrally coordinated.

Without regulation, a popular incentive ticket's circulation might go beyond control. This is undesirable for those shops which exert tight control over their advertising campaign's scale, evidenced by the "duplication not valid" notice on many real-world incentive tickets. Therefore, our design takes this issue into account and enforces the regulation by 1) introducing the restriction section into the incentive ticket mechanism (Section III-A); and 2) detecting and punishing over-limit duplications (Section IV-F).

B. Reward Target

Associated with each redeemed incentive ticket is a circulation chain $i_1 \to i_2 \to \cdots \to i_l$ (with $l \leq L_C$). As mandated by the incentive ticket exchange protocol, i_1 is the shop and i_l is the user who benefits from redeeming the incentive ticket. Therefore, only the intermediate users $(i_2, \cdots, i_{l-1})$ are rewarded for their effort in duplicating this incentive ticket[1].

C. Reward Dispensing

In this section, we evaluate a few reward dispensing strategies, which culminate in the social-aware scheme.

1) Uniform Dispensing: A straightforward strategy is the *uniform* reward dispensing, in which $i_2, \cdots, i_{l-1}$ equally share the rewards R_C. Each user receives $\frac{R_C}{l-2}$.

However, this simple strategy leads to an undesirable situation of *diminishing incentive ticket attractiveness to receivers*. For example, if $l = 3$, i_2 receives the full amount of rewards R_C. However, if $l = 4$, each of i_2 and i_3 receives $\frac{R_C}{2}$. In this case, i_3 is less eager to receive the incentive ticket than i_2 because his expected gain is smaller. The duplication process becomes weaker as the chain grows longer.

A variation which apparently addresses this issue was proposed in a few previous works [2**?**], in which each user receives an equal and *fixed* reward. In this case, no matter how long the circulation chain becomes, each user receives the same amount of reward as in a shorter chain. Then, the incentive ticket's attractiveness to receivers remains the same. However, in order to observe the reward upper limit R_C imposed by the shop, the reward amount is *at most* $\frac{R_C}{L_C-2}$. For a small R_C and a large L_C, this is *unattractive to everyone*—we trade a sagging process for a consistently weak one.

Uniform dispensing yields another undesirable situation. A pair of users might increase their expected gain by passing the incentive ticket back and forth before duplicating it to others: their identifiers occur multiple times in the chain and they claim more rewards. This is called the "looping" strategy. Looping causes a smaller user coverage and diminished expected profits for the shop.

One solution to the looping problem is to eliminate duplicated identifiers by taking the first occurrence of an identifier as its position in the circulation chain. Social-aware incentive scheme solves it by making looping unprofitable.

[1]A provision is made in the case where $l = 2$, i.e., the user who receives the incentive ticket from the shop redeems it. In this case, no reward is dispensed.

2) Geometric Dispensing: Intuitively, the users' contributions are uneven; the closer a user is to i_l in the circulation chain, the greater his contribution is.

Based on this observation, we consider a strategy in which the reward is dispensed among $i_2, \cdots, i_{l-1}$ in a ratio of $p : 1$ $(0 < p < 1)$ between consecutive users. If the reward for i_{l-1} is 1, the reward for $i_{l-2}, i_{l-3}, \cdots$ will be $p, p^2, \cdots$. We call this strategy *geometric dispensing* because the reward distribution is a *geometric* series with a constant ratio p.

In fact, the uniform dispensing can be viewed as a special case of the geometric dispensing in which $p \approx 1$. The discussions in Section IV-C1 apply to a geometric dispensing when p is close to one.

It is also instructive to consider $p \approx 0$. In this case, the geometric dispensing effectively becomes a single-level reward scheme. i_{l-1}, whose duplication directly leads to i_l's incentive ticket redemption, is rewarded with (almost) the whole amount R_C; in contrast, $i_2, \cdots, i_{l-2}$ get (almost) nothing. Under this scheme, if a user duplicates his incentive ticket to n neighbors with redemption probabilities $p_1, \ldots, p_n$, his expected gain is $\sum_{i=1}^{n} p_i$. Therefore, a rational incentive ticket holder will duplicate his incentive ticket to those neighbors with the *largest* redemption probabilities. Even he knows a neighbor v's degree k_v, he has no motive to assimilate k_v into his strategy because his expected gain $\sum_{i=1}^{n} p_i$ is independent of k_v. This may lead to premature demises of the incentive ticket duplication process in cases where the users with high redemption probabilities happen to be the reclusive ones.

3) Social-Aware Dispensing: The parameter which characterizes the above dispensing strategies is the *reward level*. Uniform dispensing is an *all-level* reward scheme: $i_2, \cdots, i_{l-1}$ are all rewarded. In contrast, as noted above, geometric dispensing with $p \approx 0$ is a *single-level* reward scheme: only i_{l-1} is rewarded.

A key insight behind the *social-aware dispensing* is that the reward level should be *fixed* and *as few as full user-intelligence utilization allows*. Fixed-level rewarding ensures consistent incentive ticket attractiveness to receivers; the "as few as..." clause prevents trivializing the rewards on the premise of making full use of user intelligence.

In the contact network model, each user knows his neighbors' redemption probabilities and degrees. By restricting spray width W_C and associating rewards with redemptions, the shop has made (partial) use of the user intelligence: a user will duplicate his incentive ticket to the neighbors with high redemption probabilities in order to maximize his expected gain. The neighbor degree information will influence a user's duplication decision *if* he can benefit from his neighbors' duplication activities. Therefore, the reward level should be *at least* two to utilize the degree information—this is where single-level rewarding falls short.

However, privacy concerns mandate the identifiers being encrypted and hence kept secret from users who do not share contact opportunity. In effect, as briefly noted in Section II, this confines the user intelligence to a local scope: a user *knows and only knows* his neighbors' redemption probabilities and degrees. Thus, a two-level reward dispensing strategy *suffices* for full user intelligence utilization.

Based on these observations, we propose the social-aware dispensing strategy as follows. Consider again the circulation chain $i_1 \to i_2 \to \cdots \to i_l$ ($l \geq 2$), in which i_1 is the shop and i_l is the user who redeems the incentive ticket. If $l \geq 4$, only the user i_{l-1} (whose duplication to i_l leads to redemption) and his predecessor i_{l-2} are rewarded; for generality, i_{l-1} and i_{l-2} share the full reward amount R_C in a ratio of $1 : \alpha$: i_{l-1} gets $\frac{1}{1+\alpha}R_C$ and i_{l-2} gets $\frac{\alpha}{1+\alpha}R_C$. If $l = 3$, only $i_{l-1} = i_2$ is rewarded; he gets the same amount $\frac{1}{1+\alpha}R_C$ as mandated by consistent incentive ticket attractiveness to receivers. If $l = 2$, no reward is dispensed.

We call the parameter α the *social weight*, for it determines the relative importance of the social information (neighbors' degrees) in a user's incentive ticket duplication decision.

D. User Strategy under Social-Aware Dispensing

Suppose a user u receives a incentive ticket from another user v. Let us consider the incentive ticket duplication strategy which maximizes u's expected gain under the social-aware reward dispensing scheme. This is the strategy u will adopt.

First, looping is unprofitable in two-level reward dispensing. Hence, u will not pass the incentive ticket back to v.

Suppose u has N neighbors beside v. Let a neighbor w's redemption probability and degree be p_w and k_w respectively. If u (somehow) knows that w will pass a incentive ticket to N' users with redemption probabilities $p'_1, p'_2, \ldots, p'_{N'}$. Then, u's expected gain for duplicating his incentive ticket to w is

$$\frac{R_C}{1+\alpha}p_w + \frac{\alpha R_C}{1+\alpha}\sum_{i=1}^{N'} p'_i. \qquad (3)$$

The optimal strategy for u is ordering his neighbors from high to low by Equation 3 and duplicating his incentive ticket to the first $\min\{k_u - 1, W_C\}$ neighbors.

However, the catch is that u does not actually *know* $p'_1, p'_2, \ldots, p'_{N'}$. Nevertheless, he can *estimate* them. For instance, a sociable tech-savvy user may have many friends who have high redemption probability for a incentive ticket of the latest gadget in the market.

For quantitative insights, we assume the existence of a (continuous) redemption probability distribution $g(p) : [0, 1] \to [0, +\infty]$. We have 1) $\int_0^1 g(p)\mathrm{d}p = 1$; and 2) $100 \cdot \int_a^b g(p)\mathrm{d}p$ is the percentage of population having a redemption probability in $[a, b]$.

If u assumes that his neighbor w will choose the W_C neighbors of w with largest redemption probabilities (excluding u) after receiving a incentive ticket from u, u estimates his gain from duplicating the incentive ticket to w in the following way. Define a function $G : [0, 1] \to [0, 1], a \mapsto \int_a^1 g(p)\mathrm{d}p$. Since $G(a)$ is a continuous decreasing function with $G(0) = 1$ and $G(1) = 0$, there exists[2] a largest $a^* \in [0, 1]$ satisfying $G(a^*) = \min\{1, W_C/(k_w - 1)\}$. Then, u's estimated gain from duplicating his incentive ticket to w is

$$\frac{p_w + \alpha \min\{k_w - 1, W_C\} \int_{a^*}^1 pg(p)\mathrm{d}p}{1+\alpha} R_C. \qquad (4)$$

Thus, the optimal strategy for u is ordering his neighbors from high to low by Equation 4 and duplicating his incentive ticket to the first $\min\{k_u - 1, W_C\}$ neighbors.

E. Choosing System Parameters

In this section, we derive a few guidelines for choosing the system parameters based on the social-aware reward dispensing scheme (Section IV-C3),

R_C has been discussed in Section IV-A; L_C (along with W_C) is used by the shop to control the maximal number of valid incentive ticket copies. Both R_C and L_C are inherent in the shop's marketing decision.

In contrast, W_C and α are more interesting because they have subtle interplays with users.

Let us look at the spray width W_C first. We propose choosing a *small* W_C. Our suggestion is based on 1) scarcity of duplications makes them appear (psychologically) more valuable; 2) a large W_C imposes a greater burden on an individual user and induces spamming. Besides, from the incentive authority's perspective, if W_C is large enough to be close to or even exceed a user's degree in the contact network, the user will become *insensitive* to the incentives: he will duplicate the incentive ticket to almost all his neighbors under *any* incentive scheme.

We now turn to the social weight parameter α. Suppose the shop shares with users the same redemption probability distribution estimation $g(p)$ (Section IV-D). As a first approximation, the shop's goal is that the incentive ticket copies reach the users with largest redemption probabilities within $L_C - 2$ hops from *the user who receives the incentive ticket from the shop*, i.e., the *initial spreader*.

However, as each user's knowledge is confined to his one-hop neighbors, the best the shop can hope for is that a user u will duplicate his incentive ticket to a set of users $\mathcal{N}_u$ which has a maximal

$$\sum_{w \in \mathcal{N}_u} (p_w + \min\{k_w - 1, W_C\} \int_{a^*}^1 pg(p)\mathrm{d}p). \qquad (5)$$

The similarity in form between Equation 5 and the numerator of Equation 4 prompts us to look for a deeper connection between them. We notice that if $\alpha = 1$, Equation 4 becomes

$$\frac{R_C}{2}(p_w + \min\{k_w - 1, W_C\} \int_{a^*}^1 pg(p)\mathrm{d}p). \qquad (6)$$

The optimal strategy for user u under $\alpha = 1$ is ordering his neighbors from high to low by Equation 6 and duplicating his incentive ticket to the first $\min\{k_u - 1, W_C\}$ neighbors—*this is exactly the set of neighbors $\mathcal{N}_u$ which maximizes Equation 5.*

This leads to an insight: *if users and the shop share the same estimation about redemption probability distribution in the population, a social weight $\alpha = 1$ will lead to a desirable situation in which a user, acting on his own interest, serves*

[2]The existence of a^* is based on the intermediate value theorem of a continuous function.

the shop's interest in the best way. This bears a resemblance to Adam Smith's invisible hand metaphor [3].

Based on the above insight, we suggest that the shop publish their estimation on the redemption probability distribution and adopt a social-aware reward dispensing scheme with social weight $\alpha = 1$.

F. Over-limit Duplication Detection and Punishment

In this section, we deliver the promise at the end of Section IV-A by demonstrating how over-limit incentive ticket duplication is detected and punished. This also justifies the reconstruction of the whole circulation chain while at most two hops in the chain are rewarded under the social-aware incentive scheme.

The incentive authority I maintains a *prefix count table* for each distinct incentive ticket (duplications are considered the same); the index is a circulation-chain prefix and the value is the count of the prefix's occurrences in all redeemed copies of the incentive ticket. For instance, if a incentive ticket with circulation chain $i_1 \to i_2 \to \cdots \to i_l$ is redeemed, I increases the count of each prefix (i.e., i_1, $i_1 i_2$, ..., and $i_1 i_2 \cdots i_l$) by one. If a prefix count exceeds the spray-width limit W_C (set by the shop before the issuance of the incentive ticket), the last *user*[3] in the prefix is convicted of over-limit duplication.

Let us illustrate this with an example. Suppose a user u duplicates his incentive ticket (with circulation chain $i_1 \to \ldots \to u$) to $W_C + 1$ users. Because the duplications happen between users, the incentive authority I cannot immediately detect this foul. However, if all of the $W_C + 1$ users redeem their incentive tickets, the count of the prefix $i_1 \ldots u$ will increase by $W_C + 1$ and will exceed the W_C limit. u, as the last user in the prefix $i_1 \ldots u$, will be convicted of over-limit duplication. Hence, u's foul is detected.

The punishment for over-limit duplication can range from a mild warning for a first-time delinquent up to account freezing for a habitual offender. We leave it open to the implementation.

V. Evaluation

To evaluate the proposed incentive schemes, we conducted an array of simulations over a real-world social-network dataset GondolaGen[4] and a synthesized contact network shown in Figure 3. Statistical profiles of these datasets are collected in Table II.

We implemented the following user duplication strategies: 1) under social-aware incentives with uniform redemption probability estimation ($g(p) = 1$); 2) under social-aware incentives with linear redemption probability estimation ($g(p) = 2 - 2p$); 3) under single-level incentives; and 4) the user duplicates the incentive ticket to a random set of his neighbors. To cover more users, a user only duplicated the incentive ticket to the neighbors who had not received it yet. Therefore, each user received *at most* one incentive ticket.

[3]As mentioned in Section III-B2, shops are not subject to the spray width restriction W_C; only users are checked for over-limit duplication.

[4]The dataset and its visualization are publicly available at http://www.infovis-wiki.net/index.php/Social_Network_Generation.

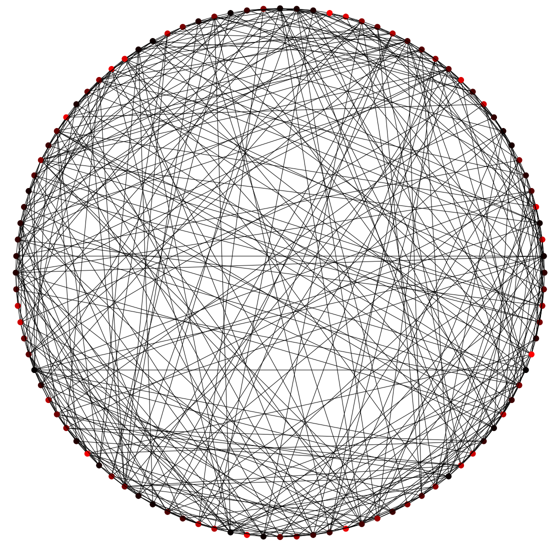

Fig. 3. The synthesized contact network.

TABLE II
Dataset profiles.

Dataset	Vertices	Edges	Max Deg.	Min Deg.
GondolaGen	242	255	14	1
Synthesized	100	483	18	2

When the duplication process finished, we computed the sum of redemption probabilities of the users who received the incentive ticket, which we called the *aggregated covered redemption probability* (ACRP hereafter), as the performance metric. A larger ACRP corresponds to higher expected profits for the shops, and hence is of better performance.

Under a given setting (i.e., dataset, W_C, L_C, and the first user who receives the incentive ticket—which we call the *initial spreader*), to account for the bias introduced by random factors, we 1) ran the simulation for random user strategy for 50 times and took the average ACRP as its performance metric; and 2) generated 50 random redemption probability settings for each dataset and took the average ACRP as the final performance results. By doing this, we mitigated the influence on performance evaluation of a particular combination of network topology and redemption probability distribution.

Figure 4 shows the simulation results under two typical settings: 1) $W_C = 2$, $L_C = 3$; and 2) $W_C = 4$, $L_C = 4$. W_C and L_C are values of the spray width and available slot fields on the incentive ticket issued by the shop to the initial spreader. Therefore, there are *at most* $\sum_{i=1}^{L_C} W_C^{i-1} = (W_C^{L_C} - 1)/(W_C - 1)$ valid copies of the incentive ticket. In particular, there are at most $2^3 - 1 = 7$ valid copies in Figure 4a and 4c and $(4^4 - 1)/(4 - 1) = 85$ valid copies in Figure 4b and 4d.

In all instances, the ACRP for the incentive-driven duplication strategies (social-aware or single-level) is significantly better than that of the random duplication strategy. In the case of Figure 4c, the average difference in ACRP between incentive-driven and random duplication strategies is around 2, which is very significant considering that the maximal ACRP for this setting is 7. This justifies our design decision

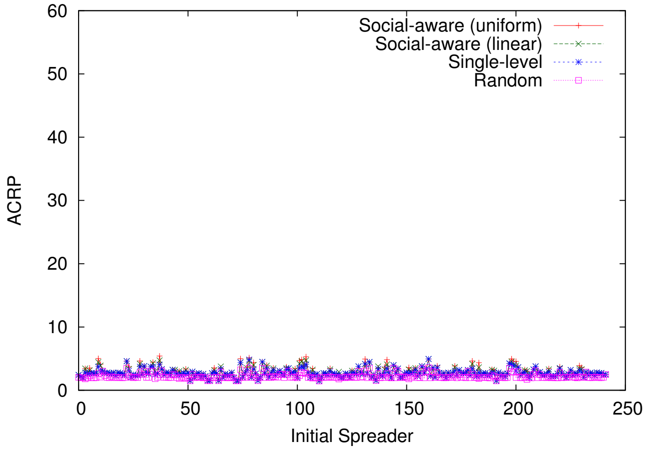

(a) GondolaGen ($W_C = 2, L_C = 3$)

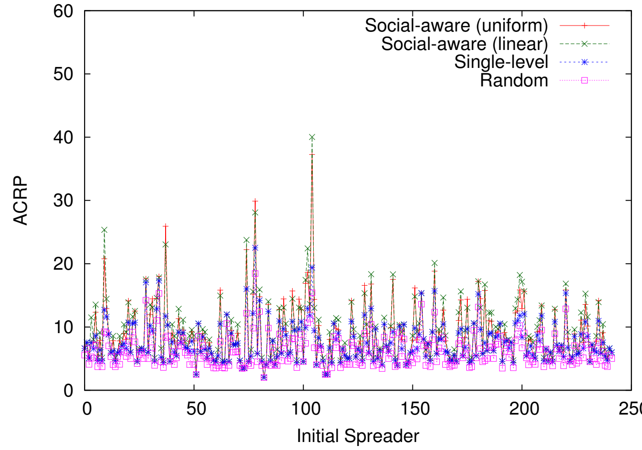

(b) GondolaGen ($W_C = 4, L_C = 4$)

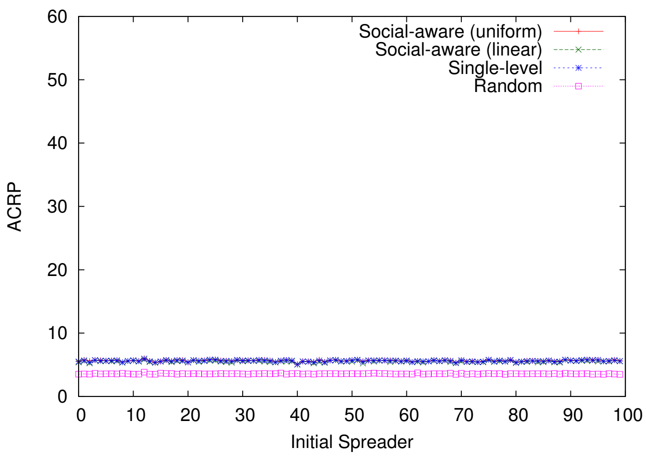

(c) Synthesized ($W_C = 2, L_C = 3$)

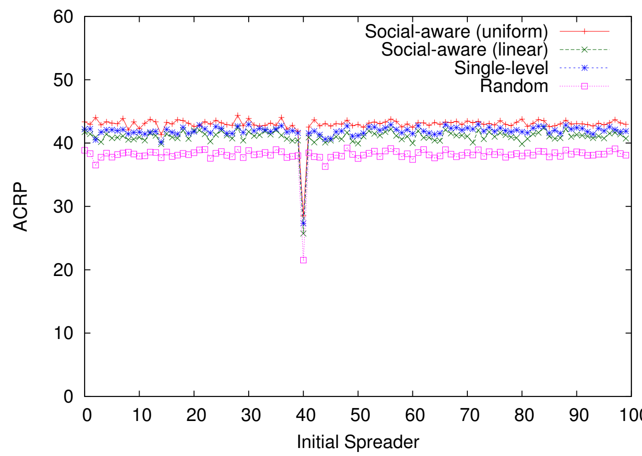

(d) Synthesized ($W_C = 4, L_C = 4$)

Fig. 4. Simulation results under two typical settings. W_C and L_C are the values of the spray width and available slot fields on the incentive ticket issued by the shop to the initial spreader (the first user who received the incentive ticket).

to associate rewards with incentive ticket *redemption* rather than *duplication alone*: a user will consider his neighbors' redemption probabilities rather than indiscriminately duplicate his incentive ticket.

The reader might notice the variation in ACRP over different initial spreaders is smaller in the synthesized network than in the GondolaGen dataset. This is due to the different underlying topological structures. A comparison of network visualizations shows that the synthesized network is much more homogeneous in terms of vertex degree than the GondolaGen dataset. While the synthesized network forms a single well-connected community, the GondolaGen dataset is much like a hierarchical society: most members are connected to one of a few well-connected hubs; the hubs form a community, which serves to connect the whole network together.

While the ACRP under the social-aware-incentive-driven strategies is very close to that under the single-level-incentive-driven strategy in Figure 4c and 4d, the former shows significant advantage over the latter for a considerable number of initial spreaders in Figure 4a and 4b. A closer examination indicates that *these initial spreaders are the social hubs* in the GondolaGen dataset. This illuminates the strength of the social-aware incentive scheme: while a single-level incentive induces the hub to favor a non-hub with high redemption probability, the social-aware incentive induces a preference for another well-connected hub. In other words, the social-aware incentive scheme encourages a duplication preference for social celebrities, who have better chances to find high-redemption-probability individuals.

VI. Related Works

The autonomous (i.e., individuals make their own decisions) and selfish (i.e., individuals contribute only if they are rewarded) assumptions in networked systems have evoked numerous researches on incorporating incentives into such systems. Metaphors from economics are often borrowed to motivate designs [4, 5].

Three schools of incentive design emerge from studies: bartering [6], credit-based [7], and reputation-based [8]. Zhang et al. [9] unified credit and reputation in their distributed incentive scheme for P2P applications. In the context of P2P applications over an underlying social network, Liu et al. [10] suggested the sweet spot for incentive system design lying between the extremes of bartering and global credit .

Lee et al. [11] discussed advertisement dissemination in the context of vehicular ad hoc networks. They proposed an incentive scheme using receipts as the evidence for redeeming rewards. However, the incentive provider has no control over the number of valid receipts and hence the resources allocated for incentives; this scheme falls prey to an uncontrollable budget for the incentive provider, whereas ours explicitly addresses this issue.

Alternatively, in a series of works [12, 13, 14], Garyfalos and Almeroth proposed the use of incentive ticket as an incentive to facilitate information dissemination. Their focus was on improving bandwidth utilization efficiency in terms of redundant transmission, whereas ours is on utilizing user intelligence for effective coverage. They did not discuss users' response to different reward dispensing strategies; the reward was simply dispensed evenly among all contributors, i.e., a uniform dispensing.

In regard to reward dispensing, Douceur and Moscibroda [15] focused on the bootstrap phase of a networked system and proposed a lottery tree scheme to probabilistically reward users for recruiting and contributing. Li and Wu [16] considered the impact of users' conservativeness on the lottery-based incentive's effectiveness and introduced a sweepstake component to address this issue. The incentive dispensing schemes in these works are based on a pre-specified temporal interval; user contributions are aggregated over the interval. In contrast, our application scenario dictates the incentive being dispensed at incentive ticket redemptions—moments which can not be specified by the system.

VII. Concluding Remarks

In this paper, we propose a promising application based on smart mobile devices of exploring user intelligence for effective commercial advertisement dissemination in a word-of-mouth fashion and identify the unique challenges and requirements for such an application. We design a incentive ticket mechanism to address the privacy concerns and present a social-aware incentive scheme to encourage users to fully contribute their user intelligence. In addition, our incentive system is cost-effective to and controllable by the incentive provider and is robust against user misbehaviors.

Commercial advertisement dissemination is just a natural and convenient scenario to describe the application of concepts and methods presented in this paper. The essence is to make *full* use of user intelligence for *effective* information deliveries. We perceive incentive and enforcement as the keys to unlock the power of users' collective intelligence for effective information dissemination.

References

[1] S. Herz, "Location specific content," U.S. Patent 20 100 120 450, May, 2010.

[2] J. Kangasharju and A. Heinemann, "Incentives for electronic coupon systems," in *Proc. Intl. Wrkshp. on Decentralized Resource Sharing in Mobile Computing and Networking*. ACM, 2006.

[3] F. Hahn, "Reflections on the invisible hand," *The Warwick Economics Research Paper Series (TWERPS)*, 1981.

[4] V. Srinivasan, P. Nuggehalli, C. Chiasserini, and R. Rao, "Cooperation in wireless ad hoc networks," in *Proc. Intl. Conf. on Computer Communications (INFOCOM)*. IEEE, 2003.

[5] M. Lelarge and J. Bolot, "Economic incentives to increase security in the internet: The case for insurance," in *Proc. Intl. Conf. on Computer Communications (INFOCOM)*. IEEE, 2009.

[6] B. Cohen, "Incentives build robustness in bittorrent," in *Wrkshp. on Economics of Peer-to-Peer systems*. IEEE, 2003.

[7] P. Marbach and Y. Qiu, "Cooperation in wireless ad hoc networks: a market-based approach," *IEEE/ACM Transactions on Networking (TON)*, vol. 13, no. 6, p. 1338, 2005.

[8] S. Zhong, J. Chen, and Y. Yang, "Sprite: a simple, cheat-proof, credit-based system for mobile ad-hoc networks," in *Proc. Intl. Conf. on Computer Communications (INFOCOM)*. IEEE, 2003.

[9] Z. Zhang, S. Chen, and M. Yoon, "March: a distributed incentive scheme for peer-to-peer networks," in *Proc. Intl. Conf. on Computer Communications (INFOCOM)*. IEEE, 2007.

[10] Z. Liu, H. Hu, Y. Liu, K. Ross, Y. Wang, and M. Mobius, "P2P trading in social networks: the value of staying connected," in *Proc. Intl. Conf. on Computer Communications (INFOCOM)*. IEEE, 2010.

[11] S. Lee, G. Pan, J. Park, M. Gerla, and S. Lu, "Secure incentives for commercial ad dissemination in vehicular networks," in *Proc. Intl. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*. ACM, 2007.

[12] A. Garyfalos and K. Almeroth, "Coupons: a multilevel incentive scheme for information dissemination in mobile networks," *Transactions on Mobile Computing (ToMC)*, vol. 7, no. 6, p. 792, 2008.

[13] K. Almeroth and A. Garyfalos, "Coupons: wide scale information distribution for wireless ad hoc networks," in *Proc. Global Telecommunications Conference (Globecom)*. IEEE, 2004.

[14] A. Garyfalos and K. Almeroth, "Coupon based incentive systems and the implications of equilibrium theory," in *Proc. Intl. Conf. on E-commerce Technology*. IEEE, 2004.

[15] J. Douceur and T. Moscibroda, "Lottery trees: motivational deployment of networked systems," *SIGCOMM Computer Communication Review (CCR)*, vol. 37, no. 4, p. 132, 2007.

[16] F. Li and J. Wu, "Frame: An innovative incentive scheme in vehicular networks," in *Proc. Intl. Conf. on Communications (ICC)*. IEEE, 2009.