

Multi-Winner Contests for Strategic Diffusion in Social Networks

Wen Shen, Yang Feng & Cristina V. Lopes

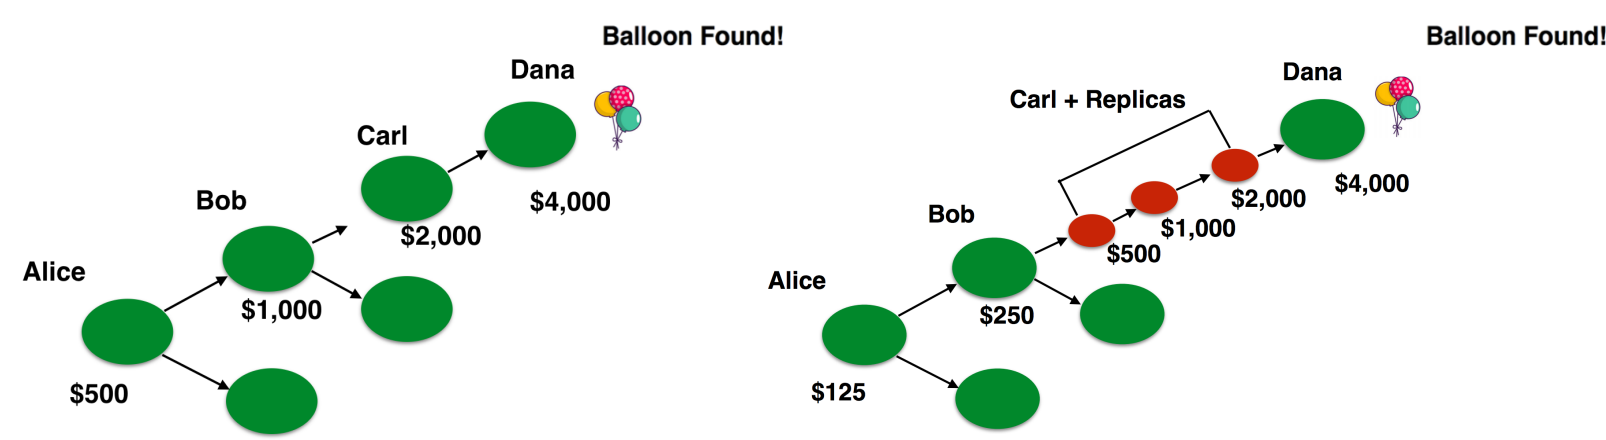
Department of Informatics, UC Irvine

{wen.shen, yang.feng, lopes}@uci.edu



Strategic Diffusion in Social Networks

Setting: Consider a principal aims to solicit as many efforts (e.g., making purchases, answering questions, reporting software bugs) as possible from the users of a social network. To encourage more players to participate, the principal offers rewards for both direct and indirect referrals. This form of referral mechanisms are usually called *incentive tree mechanisms* (or *geometric mechanisms*) [2].



(a) Honest diffusion (b) Diffusion with manipulations

Figure 1: An illustration of strategic network diffusion.

Challenges: Three challenges arise in mechanism design for strategic diffusion:

- **False-name attacks:** A strategic player may create multiple fake accounts (i.e., replicas) or identities on his behalf with one referring another to increase his rewards.




Figure 2: False-name attacks at node 2: nodes $2_1, 2_2, \dots, 2_m$ are replicas of node 2.

- **Incentives for successful referrals:** Under many mechanisms, some of the less influential players in the referral networks will receive no rewards for successful indirect referrals [2].
- **Scalability:** Most Sybil-proof mechanisms assume a tree structure. It is unclear how to extend them to large graphs that consists of thousands of nodes.

Question: How should a principal determine the rewards for each participant such that the participants cannot make profits by creating replicas (a.k.a. Sybils), and such that the players would receive positive rewards for each of their successful referrals in large-scale social networks?

Mechanism Design for Strategic Network Diffusion

In strategic diffusion, the principal is interested in a *reward mechanism* π that determines the rewards for each player that has exerted efforts.

Definition 1. A reward mechanism π is a tuple of payments for each player $v \in V$, where $G = (V, E)$. That is, $\pi = (\pi(v))_{v \in V}$, where $\pi(v) : \Theta \rightarrow \mathbb{R}$.

Design Constraints: It is desirable if a reward mechanism π satisfies the following properties:

- **False-Name-Proofness:** A reward mechanism π is false-name-proof if for all $v \in G$: $\pi(v) \geq \sum_{r \in R} \pi(r)$, where R is the set of replicas due to a false-name attack at node v .
- **Individual Rationality:** A reward mechanism π is individually rational if for each player $v \in V$ with diffusion contributions $d_v > 0$, we have v 's utility $U(v) > 0$, and the diffusion reward v receives due to referring $u \pi_d(v, u) > 0$. Here, u is a successor of v . That is, $u \in \kappa_v^+$. The utility $U(v)$ is determined by

$$U(v) = \pi(v) - \delta_v \cdot t_v. \quad (1)$$
- Here, $\pi(v)$ is the sum of the rewards for both the task effort $\pi_t(v)$ and the diffusion effort $\pi_d(v)$, $\delta_v > 0$ is a private coefficient that determines the player's marginal cost for exerting extra unit effort, and t_v is the task effort that v has exerted.
- **Budget Constraint:** A reward mechanism π is budget constrained if: $\sum_{v \in G} \pi(v) \leq \vartheta \cdot \sum_{v \in G} t_v$, where ϑ is a positive constant.
- **Monotonicity:** A reward mechanism π is monotonic if v_2 is a successor of v_1 , adding a direct successor v_i to v_2 increases v_1 's diffusion rewards $\pi_d(v_1)$ as least as much as the diffusion rewards $\pi_d(v_1')$ by adding a direct successor v_j to a successor of v_2 , where $t_{v_i} = t_{v_j}$, i.e., $\pi_d(v_1) \geq \pi_d(v_1')$.
- **Subgraph Constraint:** A reward mechanism π is subgraph-constrained if $\pi(v)$ only depends on the rooted subgraph G_v .
- **Scalability (informal):** The reward mechanism should be capable of handling large graphs that consist of thousands of nodes.

The Multi-Winner Contests Mechanism

The MWC mechanism has two key ingredients: it first calculates the virtual credits for the diffusion contributions of each player with successful referrals; it then determines the diffusion rewards by holding a contest [3] among players that are in his rooted subgraph. The MWC mechanism is computationally efficient and satisfies several desirable properties, including false-name-proofness, individual rationality, budget constraint, monotonicity, and subgraph constraint.

Virtual Credits for Diffusion Contributions: For each newly joined player v that has exerted task efforts t_v , the MWC mechanism pays $\pi_t(v) = \mu \cdot t_v$ ($\mu > 0$) for his task contributions and allocates virtual credits $\eta \cdot (t_v)^2$ for his diffusion contributions. If v has either directly or indirectly referred player $u \in \kappa_v^+$ to participate ($t_u > 0$), his virtual credits b_v are computed by:

$$b_v = \eta \cdot (t_v)^2 + t_v \cdot \sum_{u \in \kappa_v^+} \sum_{p \in P_{vu}} t_u \cdot \omega(p) \cdot \lambda^{|p|}, \quad (2)$$

where $0 < \lambda < 1$, $\eta \geq \lambda/2$, and P_{vu} is the set of paths from v to u .

Diffusion Rewards: The MWC mechanism utilizes the probability of winning to determine the proportion of the total diffusion rewards for each player. The diffusion rewards of v is determined as:

$$\pi_d(v) = \frac{(b_v)^\sigma}{\sum_{u \in V'} (b_u)^\sigma} \cdot \phi \cdot \sum_{u \in V'} t_u \cdot \frac{deg_{G_v}^-(u)}{deg^-(u)}, \quad (3)$$

where $0 < \sigma \leq 1$ is the noise factor, $\phi > 0$, and $deg_{G_v}^-(u)$ denotes the number of direct predecessors of u in $G_v = (V', E')$. Here, $deg_{G_v}^-(u) \leq deg^-(u)$.

The MWC Mechanism: The MWC mechanism $\pi = (\pi(v))_{v \in G}$ uses a post-price mechanism $\pi_t(v) = \mu \cdot t_v$ to reward a player that exerts t_v task efforts. The total reward for player v includes both the task rewards $\pi_t(v)$ and the diffusion rewards $\pi_d(v)$. Thus, we have the total reward for v :

$$\pi(v) = \begin{cases} \mu \cdot t_v + \pi_d(v) & \text{if } b_v \geq \eta \cdot (t_v)^2 \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

where $\mu \in \mathbb{R}_{>0}$ is the parameter that characterizes to what extent the principal values agents' efforts.

Experiments

We used four public datasets: Twitter, Flickr, Flixster, and Digg. For each dataset, we first estimated the influence diffusion probabilities for each node using the learning algorithms by [4] with the Bernoulli distribution under the static model. We then simulated the influence diffusion process with the general threshold model [5] based on the estimated diffusion probabilities. After preprocessing, each dataset produced a largest weakly connected component.

Dataset	#Nodes	#Edges	#Seeds	M.D.	A.D.
Twitter	323,185	2,148,717	1,715	8,822	52
Flickr	145,305	2,149,882	768	6,731	34
Flixster	95,969	484,865	502	3,109	27
Digg	17,817	128,587	107	1,375	20

Table 1: Dataset configuration: M.D. – maximum degree, A.D. – average degree.

large social networks with hundreds of thousands of nodes.

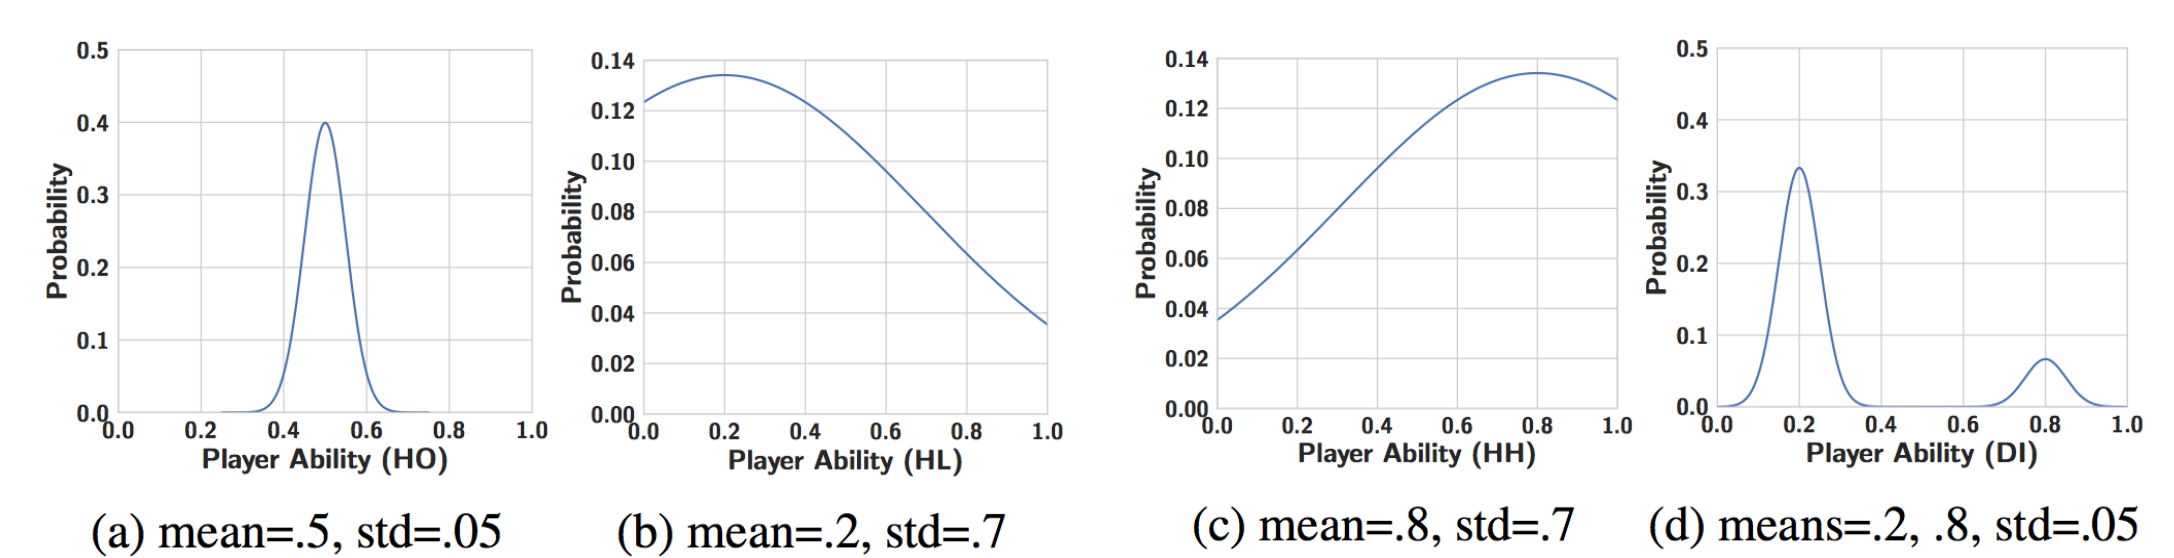


Figure 3: Probability density function of the means of players' abilities by four different groups: *homogeneous (HO)*-players with similar levels of abilities; *heterogeneously low (HL)*-players with different level of abilities, and the average abilities are low; *heterogeneously high (HH)*-players with different level of abilities, and the average abilities are high; and *distinct (DI)*-a portion of players with low average abilities, the other portion with high average abilities.

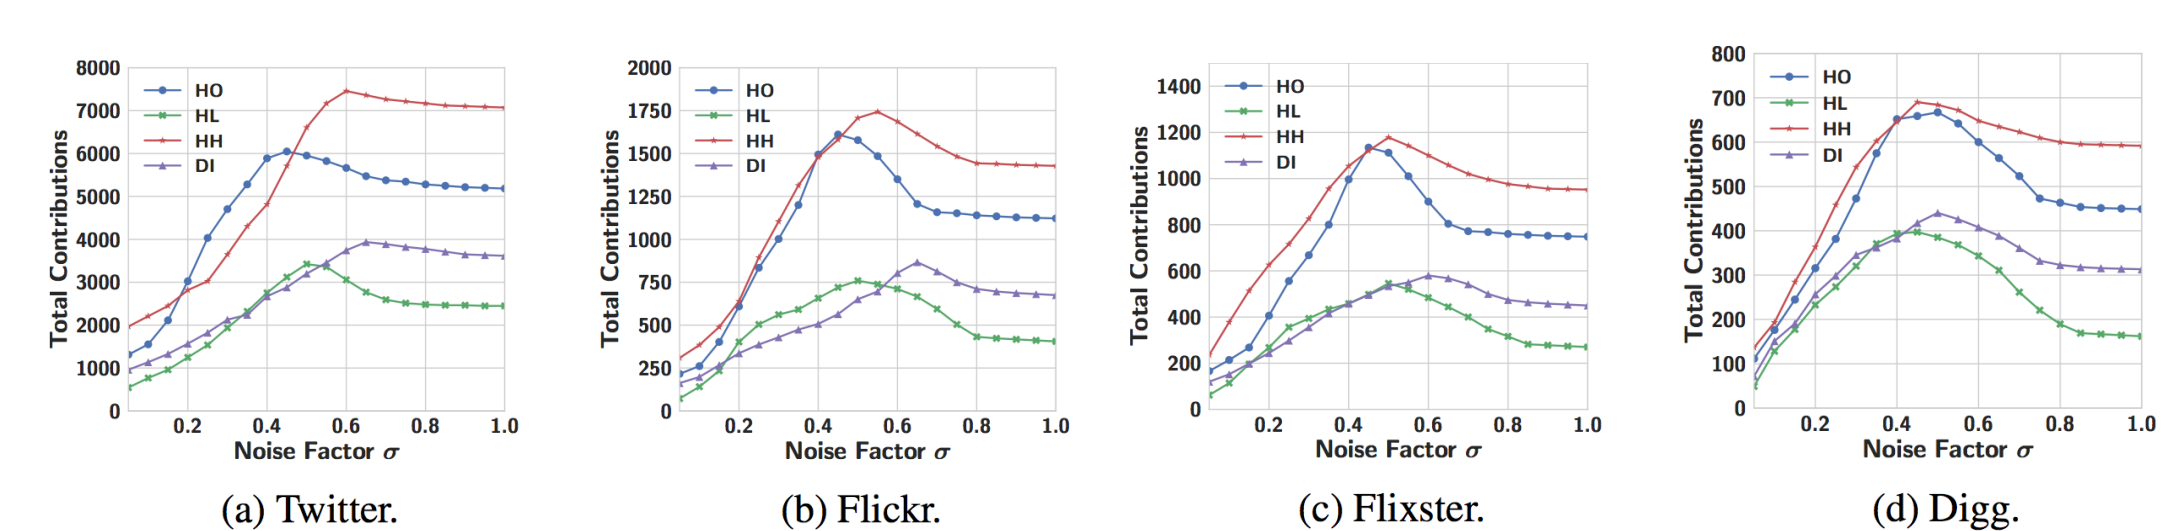


Figure 4: A comparison of total contributions with different bias factors: all the optimal noise factors fell within the range between 0.45 and 0.65. The results suggest that medium noise factors were typically superior than the others; there were "sweet spots" for stakeholders to maximize the total efforts.

Conclusions

- We introduce a novel multi-winner contests mechanism for strategic diffusion in social networks. The mechanism is false-name-proof and individual rational for players with successful referrals. It is computationally efficient, budget-constrained, monotonic, and subgraph-constrained.
- Experiments on four real-world social networking datasets show that stakeholders can boost the performance of players by selecting proper noise factors. They further indicate that the MWC mechanism is both false-name-proof and scalable. The results demonstrate the promising prospects of bringing contests to mechanism design.
- Our work opens several exciting avenues for future research: (1) agents with discounting values; (2) automated parameter selection; and (3) truthful mechanisms with contests.

References

- [1] A. Burnap, Y. Ren, P.Y. Papalambros, R. Gonzalez, and R. Gerth. A simulation based estimation of crowd ability and its influence on crowdsourced evaluation of design concepts. In *Proc. ASME 2013 IDETC*, 2013.
- [2] F. A. Drucker and L. K. Fleischer. Simpler sybil-proof mechanisms for multi-level marketing. In *Proc. of EC 2012*, pages 441–458, 2012.
- [3] Q. Fu and J. Lu. The beauty of "bigness": On optimal design of multi-winner contests. *Games and Economic Behavior*, 66(1):146–161, 2009.
- [4] A. Goyal, F. Bonchi, and L.V.S. Lakshmanan. Learning influence probabilities in social networks. In *Proc. of WSDM 2010*, pages 241–250, 2010.
- [5] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *Proc. of KDD 2003*, pages 137–146, 2003.