Welfare Maximization with Friends-of-Friends Network Externalities

Sayan Bhattacharya*[1], Wolfgang Dvořák[2], Monika Henzinger[2], and Martin Starnberger[2]

1 **Institute of Mathematical Sciences, Chennai, India.**
2 **University of Vienna, Faculty of Computer Science, Austria**

Abstract

Online social networks allow the collection of large amounts of data about the influence between users connected by a friendship-like relationship. When distributing items among agents forming a social network, this information allows us to exploit network externalities that each agent receives from his neighbors that get the same item. In this paper we consider Friends-of-Friends (2-hop) network externalities, i.e., externalities that not only depend on the neighbors that get the same item but also on neighbors of neighbors. For these externalities we study a setting where multiple different items are assigned to unit-demand agents. Specifically, we study the problem of welfare maximization under different types of externality functions. Let n be the number of agents and m be the number of items. Our contributions are the following: (1) We show that welfare maximization is APX-hard; we show that even for step functions with 2-hop (and also with 1-hop) externalities it is NP-hard to approximate social welfare better than $(1-1/e)$. (2) On the positive side we present (i) an $O(\sqrt{n})$-approximation algorithm for general concave externality functions, (ii) an $O(\log m)$-approximation algorithm for linear externality functions, and (iii) an $(1-1/e)\frac{1}{6}$-approximation algorithm for 2-hop step function externalities. We also improve the result from [6] for 1-hop step function externalities by giving a $(1-1/e)/2$-approximation algorithm.

1998 ACM Subject Classification F.2 Analysis of algorithms and problem complexity, G.1.2 Approximation

Keywords and phrases network externalities, welfare maximization, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.STACS.2015.90

1 Introduction

Assume you have to form a committee and need to decide whom to choose as a member. It seems like a good strategy to select members from your network that are well-connected to the whole field so that not only the knowledge of the actual members but also of their whole network can be called upon when needed. Along the same vein assume you want to play a multiplayer online game but you do not have enough friends who are willing to play with you. Then it is a good idea to ask these friends to contact their friends whether they are willing to play as well. Both these settings can be modeled by a social network graph and in both settings not the *direct* (or *1-hop*) neighbors alone, but instead the 1-hop neighbors in combination with the *neighbors of neighbors* (or *2-hop neighbors*) are the decisive factor. Note that the 2-hop neighborhoods cannot be modeled by 1-hop neighborhoods through the

* The work was done while the author was at the University of Vienna, Faculty of Computer Science.

32nd Symposium on Theoretical Aspects of Computer Science (STACS 2015).
Editors: Ernst W. Mayr and Nicolas Ollinger; pp. 90–102
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

insertion of an additional edge (to the neighbor of the neighbor) as we require that every *participating* neighbor of a neighbor is adjacent to a *participating* neighbor. In the above example, we can only get the opinion of a contact of a contact if we asked the contact before. In the same way, the participation of a friend of a friend will only be possible if there is a participating friend that invites him.

There has been a large body of work by social scientists and, in the last decade, also by computer scientists (see e.g., the influential paper by Kempe, Kleinberg, and Tardos [18] and its citations) to model and analyze the effect of 1-hop neighborhoods. The study of 2-hop neighborhoods has received much less attention (see e.g., [10, 17]). This is surprising as a recent study [14] of the Facebook network shows that the median Facebook user has 31k people as "friends of friends" and due to some users with very large friend lists, the average number of friends-of-friends reaches even 156k. Thus, even if each individual friend of a friend has only a small influence on a Facebook user, in aggregate the influence of the friends-of-friends might be large and should not be ignored.

We, therefore, initiate the study of the influence of 2-hop neighborhoods in the popular *assignment* setting, where items are assigned to users whose values for the item depend on who else in their neighborhood has the item. There is a large body of work on mechanisms and pricing strategies for this problem with a single [5, 15, 4, 1, 7, 19, 11, 3, 13] or multiple items [8, 2, 6, 12, 21, 22, 20, 16] when the valuation function of a user depends *solely* on the 1-hop neighborhood of a user and the user itself. All this work assumes that there is an infinite supply of items (of each type if there are different items) and the users have unit-demand, that is, they want to buy only *one* item. This is frequently the case, for example, if the items model competing products or if the user has to make a binary decision between participating or not participating. In the above examples, this requirement would model that each user can only be in one committee or play one game at a time.

Thus, we study the allocation of items to users in a setting with 2-hop network externalities, where the valuation that a user derives from the products depends on herself, her 1-hop, and her 2-hop neighborhood with the goal of maximizing the social welfare of the allocation. The prior work that is most closely related to our work is the work by Bhalgat et al. [6], where they study the multi-item setting with 1-hop externality functions and give approximation algorithms for different classes of externality functions. For linear externalities they give a 1/64-approximation algorithm and for step function externalities they get an approximation ratio of $(1 - 1/e)/16 \approx 0.04$. Additionally they present a $2^{O(d)}$-approximation algorithm for convex externalities that are bounded by polynomials of degree d and a polylogarithmic approximation algorithm for submodular externalities.

1.1 Our Results

The Model: Let $G = (V, E)$ be an undirected graph modeling the social network. Consider any agent $j \in V$ who receives item $i \in I$, and let $S_{ij} \subseteq V \setminus \{j\}$ denote the (2-hop) *support* of agent j for item i: this is the set of agents who contribute towards the valuation of j. Specifically, an agent $j' \in V \setminus \{j\}$ belongs to the set S_{ij} iff j' gets item i *and* the following condition holds: either j' is a neighbor of j (i.e., $(j, j') \in E$), or j and j' have a common neighbor j'' who also gets item i. The valuation received by agent j is equal to $\lambda_{ij} \cdot ext_{ij}(|S_{ij}|)$, where λ_{ij} is the agent's *intrinsic valuation* and $ext_{ij}(|S_{ij}|)$ is her 2-hop *externality* for item i. The goal is to compute an assignment of items to the agents that maximizes the social welfare, which is defined as the sum of the valuations obtained by the agents.

We study three types of 2-hop externality functions, namely concave, linear and step function externalities.

Step-function externalities: Consider a game requiring a minimal or fixed number of players (larger than two), e.g., Bridge or Canasta, then the externality is a step function. For step functions we show that it is NP-hard to approximate the social welfare within a factor of $(1 - 1/e)$. The result holds for 1-hop and 2-hop externalities. We also show that the problem remains APX-hard when the number of items is restricted to 2. Then we give an $(1 - 1/e)/6 \approx 0.1$-approximation algorithm for 2-hop step function externalities. Note that this is within a factor of 1/6 of the hardness bound. Our technique also leads to a *combinatorial* $(1-1/e)/2 \approx 0.3$-approximation algorithm for 1-hop step function externalities, improving the approximation ratio of the *LP-based* algorithm in [6].

Linear externalities: First we show that social welfare maximization for linear 2-hop externality functions is NP-hard.[1] Then we give an $O(\log n)$-approximation algorithm for linear 2-hop externalities. For these externality functions we can relax the unit-demand requirement. Specifically, we can handle the setting where each user j can buy up to c_j different items, where c_j is a parameter given in the input.[2]

Concave externalities: We give an $O(\sqrt{n})$ -approximation algorithm when the externality functions $ext_{ij}(.)$ are concave and monotone.

Extensions: Our algorithms for linear and concave externalities can be further generalized to allow a weighting of 2-hop neighbors so that 2-hop neighbors have a lower weight than 1-hop neighbors. This can be useful if it is important that the influence of 2-hop neighbors does not completely dominate the influence of the 1-hop neighbors.

Techniques: The main challenge in dealing with 2-hop externalities is as follows. Fix an agent j who gets an item i, and let $V_i \subseteq V$ denote the set of all agents who get item i. Recall that the agent j's externality is given by $ext_{ij}(|S_{ij}|)$, where the set S_{ij} is called the support of agent j. The problem is that $|S_{ij}|$, as a function of $V_i \setminus \{j\}$, is not submodular. This is in sharp contrast with the 1-hop setting, where the support for the agent's externality comes only from the set of her 1-hop neighbors who receive item i.

All the mechanisms in [6] use the same basic approach: First solve a suitable LP-relaxation and then round its values independently for each item i. In the 2-hop setting, however, the lack of submodularity of the support size (as described above) leads to many dependencies in the rounding step. Nevertheless, we show how to extend the technique in [6] to achieve the approximation algorithm for linear 2-hop externality functions, using a novel LP. We further give a simple combinatorial algorithm with an approximation guarantee of $O(\sqrt{n})$ for 2-hop concave externalities. For this, we show that either an $\Omega(1/\sqrt{n})$-fraction of the optimal social welfare comes from a single item, or we can reduce our problem to a setting with 1-hop step function externalities by losing an $(1 - \Omega(1/\sqrt{n}))$-fraction of the objective.

Our approach for 2-hop step functions is different. We use a novel decomposition of the graph into a maximal set of disjoint connected sets of size 3, 2, and 1. We say an assignment is *consistent* if it assigns all the nodes (i.e., users) in the same connected set the same item.

---

[1] Theorem 3.1 in [6] claims that the welfare maximization problem for linear 1-hop externality functions in complete graphs is MaxSNP-hard, which would imply our result, but, as we show in the full version, this claim is not true.

[2] This is also true for the results in [6]. In both results, the assumption is that the valuation functions are additive over the items.

We show first that restricting ourself to consistent assignments reduces the maximum welfare by at most a factor of 1/6. Finally, we show that finding the optimal consistent assignment is equal to maximizing social welfare in a scenario where agents are not unit demand, do not influence each other, and have valuation functions that are fractionally subadditive in the items they get assigned. For the latter we use the $(1-1/e)$-approximation algorithm by Feige [9].

2 Notations and Preliminaries

We are given a simple undirected graph $G=(V,E)$ with $|V|=n$ nodes. Each node $j \in V$ in this graph is an agent, and there is an edge $(j,j') \in E$ iff the agents j and j' are friends with each other. There is a set of m items $I=\{1,\dots,m\}$. Each item is available in unlimited supply, and each agent wants to get at most one item. An *assignment* $\mathcal{A}: V \to I$ specifies the item received by every agent, and under this assignment, $u_j(\mathcal{A},G)$ gives the *valuation* of an agent $j \in V$. Our goal is to find an assignment that maximizes the *social welfare* $\sum_{j\in V} u_j(\mathcal{A},G)$, i.e., the sum of the valuations of the agents.

Let $F_j^1(G)$ (resp. $F_j^2(G)$) be the 1-hop (resp. 2-hop) neighborhood of node j.

$$F_j^1(G)=\{j' \in V : (j,j') \in E\}, \quad F_j^2(G)=\bigcup_{j'\in F_j^1(G)} F_{j'}^1(G) \setminus (F_j^1(G) \cup \{j\}).$$

Define $V_i(\mathcal{A},G)=\{j \in V : \mathcal{A}(j)=i\}$ to be the set of agents who receive item $i \in I$ under the assignment $\mathcal{A}$. Let $N_j^1(i,\mathcal{A},G)=F_j^1(G)\cap V_i(\mathcal{A},G)$ denote the set of agents in $F_j^1(G)$ who receive item i under the assignment $\mathcal{A}$. Further, let $N_j^2(i,\mathcal{A},G)=F_j^2(G)\cap V_i(\mathcal{A},G)\cap \left(\bigcup_{j''\in N_j^1(i,\mathcal{A},G)} F_{j''}^1(G)\right)$ denote the set of agents in $F_j^2(G)$ who receive item i under the assignment $\mathcal{A}$ *and are adjacent to some node in* $N_j^1(i,\mathcal{A},G)$.

The *support* of an agent $j \in V$ for item $i \in I$ is defined as $S_{ij}(\mathcal{A},G)=N_j^1(i,\mathcal{A},G)\cup N_j^2(i,\mathcal{A},G)$. This is the set of agents contributing towards the valuation of j for item i. Let λ_{ij} be the *intrinsic valuation* of agent j for item i, and let $ext_{ij}(|S_{ij}(\mathcal{A},G)|)$ be the *externality* of the agent for the same item. The agent's valuation from the assignment $\mathcal{A}$ is given by the following equality.

$$u_j(\mathcal{A},G)=\lambda_{\mathcal{A}(j),j}\cdot ext_{\mathcal{A}(j),j}(|S_{\mathcal{A}(j),j}(\mathcal{A},G)|).$$

We consider three types of externalities in this paper.

▶ **Definition 1.** In *concave externality* it holds that $ext_{ij}(t)$ is a monotone and concave function of t, with $ext_{ij}(0)=0$, for every item $i \in I$ and agent $j \in V$.

▶ **Definition 2.** In *linear externality* it holds that for all $j \in V$, $i \in I$ and every nonnegative integer t, we have $ext_{ij}(t)=t$.

We extend the step function definition of [6] as follows to 2-hop neighborhoods.

▶ **Definition 3.** For integer $s \geq 1$, in *s-step function* externality it holds that for all $j \in V$, $i \in I$ and every nonnegative integer t, we have $ext_{ij}(t)$ is 1 if $t \geq s$ and 0 otherwise.

We omit the symbol G from these notations if the underlying graph is clear from the context. Some proofs are omitted due to space restrictions but are provided in a full version available at `http://eprints.cs.univie.ac.at/4240/1/paper-full.pdf`.

3 An $O(\sqrt{n})$-Approximation for Concave Externalities

For the rest of this section, we fix the underlying graph G, and assume that the agents have concave externalities as per Definition 1. We also fix the intrinsic valuations λ_{ij} and the externality functions $ext_{ij}(.)$.

- Let $\mathcal{A}^* \in \arg max_{\mathcal{A}} \left\{\sum_{j \in V} u_j(\mathcal{A})\right\}$ be an assignment that maximizes the social welfare, and let $\textsc{Opt} = \sum_{j \in V} u_j(\mathcal{A}^*)$ be the optimal social welfare.
- Let $X^* = \{j \in V : |S_{\mathcal{A}^*(j),j}(\mathcal{A}^*)| \geq \sqrt{n}\}$ be the set of agents with support size at least $\sqrt{n}$ under the assignment $\mathcal{A}^*$, and let $Y^* = V \setminus X^*$ be the remaining set of agents.

Since X^* and Y^* partition the set of agents V, there can be two possible cases. Half of the social welfare under $\mathcal{A}^*$ is coming (1) either from the agents in X^*, or (2) from the agents in Y^*. Lemma 4 shows that in the former case there is a *uniform assignment*, where every agent gets the same item, that retrieves $1/(2\sqrt{n})$-fraction of the optimal social welfare. We consider the latter case in Lemma 5, and reduce it to a problem with 1-hop externalities.

▶ **Lemma 4.** *If* $\sum_{j \in X^*} u_j(\mathcal{A}^*) \geq \textsc{Opt}/2$*, then there is an item* $i \in I$ *such that* $\sum_{j \in V} u_j(\mathcal{A}^i) \geq \textsc{Opt}/(2\sqrt{n})$*, where* $\mathcal{A}^i$ *is the assignment that gives item* i *to every agent in* V*, that is,* $\mathcal{A}^i(j) = i$ *for all* $j \in V$*.*

Proof. Define the set of items $I(X^*) = \bigcup_{j \in X^*} \{\mathcal{A}^*(j)\}$.

We claim that $|I(X^*)| \leq \sqrt{n}$. To see why the claim holds, let $V_i^* = \{j \in V : \mathcal{A}^*(j) = i\}$ be the set of agents who receive item i under $\mathcal{A}^*$. Now, fix any item $i \in I(X^*)$, and note that, by definition, there is an agent $j \in X^*$ with $\mathcal{A}^*(j) = i$. Thus, we have $|V_i^*| \geq |S_{ij}(\mathcal{A}^*)| \geq \sqrt{n}$. We conclude that $|V_i^*| \geq \sqrt{n}$ for every item $i \in I(X^*)$. Since $\sum_{i \in I(X^*)} |V_i^*| \leq |V| = n$, it follows that $|I(X^*)| \leq \sqrt{n}$.

To conclude the proof of the lemma, we now make the following observations.

$$\sum_{j \in X^*} u_j(\mathcal{A}^*) = \sum_{i \in I(X^*)} \sum_{j \in X^* : \mathcal{A}^*(j) = i} u_j(\mathcal{A}^*) \leq |I(X^*)| \cdot \max_{i \in I(X^*)} \left(\sum_{j \in X^* : \mathcal{A}^*(j) = i} u_j(\mathcal{A}^*) \right)$$

$$\leq \sqrt{n} \cdot \max_{i \in I(X^*)} \left(\sum_{j \in X^* : \mathcal{A}^*(j) = i} u_j(\mathcal{A}^i) \right) \leq \sqrt{n} \cdot \max_{i \in I(X^*)} \left(\sum_{j \in V} u_j(\mathcal{A}^i) \right)$$

The lemma holds since $\textsc{Opt}/(2\sqrt{n}) \leq \sum_{j \in X^*} u_j(\mathcal{A}^*)/\sqrt{n} \leq \max_{i \in I(X^*)} \left(\sum_{j \in V} u_j(\mathcal{A}^i) \right)$. ◀

For every item $i \in I$ and agent $j \in V$, we now define the externality function $\hat{ext}_{ij}(t)$ and the valuation function $\hat{u}_j(\mathcal{A})$.

$$\hat{ext}_{ij}(t) = \begin{cases} ext_{ij}(1) & \text{if } t \geq 1; \\ 0 & \text{otherwise.} \end{cases} \qquad \hat{u}_j(\mathcal{A}) = \lambda_{\mathcal{A}(j),j} \cdot \hat{ext}_{ij}(|N_j^1(i, \mathcal{A})|) \tag{1}$$

Clearly, for every assignment $\mathcal{A} : V \to I$, we have $0 \leq \sum_{j \in V} \hat{u}_j(\mathcal{A}) \leq \sum_{j \in V} u_j(\mathcal{A})$. Also note that the valuation function $\hat{u}_j(.)$ depends only on the 1-hop neighborhood of the agent j. Specifically, if an agent j gets an item i, then her valuation is $\lambda_{ij} \cdot ext_{ij}(1)$ if at least one of her 1-hop neighbors also gets the same item i, and zero otherwise. Bhalgat et al. [6] gave an LP-based $O(1)$-approximation for finding an assignment $\mathcal{A} : V \to I$ that maximizes the social welfare in this setting (also see Section 5 for a combinatorial algorithm). In the lemma below, we show that if the agents in Y^* contribute sufficiently towards $\textsc{Opt}$ under the assignment $\mathcal{A}^*$, then by losing an $O(\sqrt{n})$-factor in the objective, we can reduce our original problem to the one where the externalities are $\hat{ext}_{ij}(.)$ and the valuations are $\hat{u}_j(.)$.

▶ **Lemma 5.** *If* $\sum_{j \in Y^*} u_j(\mathcal{A}^*) \geq \textsc{Opt}/2$*, then* $\sum_{j \in V} \hat{u}_j(\mathcal{A}^*) \geq \textsc{Opt}/(2\sqrt{n})$*.*

Proof. Consider a node $j \in Y^*$ that makes nonzero contribution towards the objective (i.e., $u_j(\mathcal{A}^*) > 0$) and suppose that it gets items i (i.e., $\mathcal{A}^*(j) = i$). Since $u_j(A^*) > 0$, we have $S_{ij}(\mathcal{A}^*) = N_j^1(i, \mathcal{A}^*) \cup N_j^2(i, \mathcal{A}^*) \neq \emptyset$, which in turn implies that $N_j^1(i, \mathcal{A}^*) \neq \emptyset$. Thus, we have $\hat{u}_j(\mathcal{A}^*) = \lambda_{ij} \cdot ext_{ij}(1)$. Since $|S_{ij}(\mathcal{A}^*)| \leq \sqrt{n}$ and $ext_{ij}(.)$ is a concave function, we have $ext_{ij}(1) \geq ext_{ij}(|S_{ij}(\mathcal{A}^*)|)/|S_j(\mathcal{A}^*)| \geq ext_{ij}(|S_{ij}(\mathcal{A}^*)|)/\sqrt{n}$. Multiplying both sides of this inequality by λ_{ij}, we conclude that $\hat{u}_j(\mathcal{A}^*) \geq u_j(\mathcal{A}^*)/\sqrt{n}$ for all agents $j \in Y^*$ with $u_j(\mathcal{A}^*) > 0$. In contrast, if $u_j(\mathcal{A}^*) = 0$, then the inequality $\hat{u}_j(\mathcal{A}^*) \geq u_j(\mathcal{A}^*)/\sqrt{n}$ is trivially true. Thus, summing over all $j \in Y^*$, we infer that $\sum_{j \in Y^*} \hat{u}_j(\mathcal{A}^*, G) \geq \sum_{j \in Y^*} u_j(\mathcal{A}^*, G)/\sqrt{n} \geq \textsc{Opt}/(2\sqrt{n})$. The lemma now follows since $\sum_{j \in V} \hat{u}_j(\mathcal{A}^*, G) \geq \sum_{j \in Y^*} \hat{u}_j(\mathcal{A}^*, G)$. ◀

The algorithm for concave externalities. We run two procedures. Procedure (1) returns an assignment $\mathcal{A}' \in \arg\max_{i \in I} \left(\sum_{j \in V} u_j(\mathcal{A}^i)\right)$, where $\mathcal{A}^i(j) = i$ for all $i \in I$ and $j \in V$. Procedure (2) invokes the algorithm in [6] and returns an assignment $\mathcal{A}''$ such that $\sum_{j \in V} \hat{u}_j(\mathcal{A}'') \geq (1/\alpha) \cdot \max_{\mathcal{A}} \left(\sum_{j \in V} \hat{u}_j(\mathcal{A})\right)$ for some constant $\alpha \geq 1$, where the function $\hat{u}_j(.)$ is defined as in equation 1. Our algorithm now compares these two assignments $\mathcal{A}'$ and $\mathcal{A}''$ and returns the one that gives maximum social welfare, i.e, we output an assignment $\mathcal{A}''' \in \arg\max_{\mathcal{A} \in \{\mathcal{A}', \mathcal{A}''\}} \left(\sum_{j \in V} u_j(\mathcal{A})\right)$.

▶ **Theorem 6.** *The algorithm described above gives an $O(\sqrt{n})$-approximation for social welfare under 2-hop, concave externalities.*

Proof. Recall the notations introduced in the beginning of Section 3. Since the set of agents V is partitioned into $X^* \subseteq V$ and $Y^* = V \setminus X^*$, either $\sum_{j \in X^*} u_j(\mathcal{A}^*) \geq \textsc{Opt}/2$ or $\sum_{j \in Y^*} u_j(\mathcal{A}^*) \geq \textsc{Opt}/2$. In the former case, Lemma 4 guarantees that $\sum_{j \in \mathcal{A}'''} u_j(\mathcal{A}''') \geq \sum_{j \in \mathcal{A}'} u_j(\mathcal{A}') \geq \textsc{Opt}/(2\sqrt{n})$. In the latter case, by Lemma 5 we have $\sum_{j \in \mathcal{A}'''} u_j(\mathcal{A}''') \geq \sum_{j \in \mathcal{A}''} u_j(\mathcal{A}'') \geq \sum_{j \in \mathcal{A}''} \hat{u}_j(\mathcal{A}'') \geq \sum_{j \in \mathcal{A}^*} \hat{u}_j(\mathcal{A}^*)/\alpha \geq \textsc{Opt}/(2\alpha\sqrt{n})$. Since α is a constant, we conclude that the social welfare returned by our algorithm is always within an $O(\sqrt{n})$-factor of the optimal social welfare. ◀

4 An $O(\log m)$-Approximation for Linear Externalities

In this section, we assume that the input graph $G = (V, E)$ is of the following form. The set V is partitioned into three groups V_1, V_2 and V_3. Further, an edge in E either connects a node in V_1 with a node in V_2, or connects a node in V_2 with a node in V_3. Our goal is to assign the items to the agents in such a way as to maximize the social welfare from the set V_1. We refer to this problem as Restricted-Welfare.

▶ **Theorem 7.** *Any α-approximation algorithm for the* Restricted-Welfare *problem can be converted into an $O(\alpha)$-approximation algorithm for the welfare-maximization problem in general graphs with linear (or even concave) externalities.*

Consider the LP below. Here, the variable $\alpha(i, j, k)$ indicates if both the agents $j \in V_1$ and $k \in F_j^1$ received item $i \in I$. If this variable is set to one, then agent j gets one unit of externality from agent k. Similarly, the variable $\beta(i, j, l)$ indicates if both the agents $j \in V_1, l \in V_3 \cap F_j^2$ received item $i \in I$ *and* there is at least one agent $k \in F_j^1 \cap F_l^1$ who also received the same item. If this variable is set to one, then agent j gets one unit of externality from agent l. Clearly, the total valuation of agent j for item i is given by $\sum_{k \in V_2 \cap F_j^1} \lambda_{ij} \cdot \alpha(i, j, k) + \sum_{l \in V_3 \cap F_j^2} \lambda_{ij} \cdot \beta(i, j, l)$. Summing over all the items and all the agents in V_1, we see that the LP-objective encodes the social welfare of the set V_1.

$$\text{Maximize:} \quad \sum_{j\in V_1}\sum_{i\in I}\lambda_{ij}\cdot\left(\sum_{k\in V_2\cap F_j^1}\alpha(i,j,k)+\sum_{l\in V_3\cap F_j^2}\beta(i,j,l)\right) \tag{2}$$

$$\beta(i,j,l)\le \min\{w(i,l),y(i,j)\} \qquad \forall i\in I, j\in V_1, l\in V_3\cap F_j^2 \tag{3}$$

$$\beta(i,j,l)\le \textstyle\sum_{k\in F_j^1\cap F_l^1} z(i,k) \qquad \forall i\in I, j\in V_1, l\in V_3\cap F_j^2 \tag{4}$$

$$\alpha(i,j,k)\le \min\{y(i,j),z(i,k)\} \qquad \forall i\in I, j\in V_1, k\in V_2\cap F_j^1 \tag{5}$$

$$\textstyle\sum_i y(i,j)\le 1,\ \sum_i z(i,k)\le 1,\ \sum_i w(i,l)\le 1 \qquad \forall j,k,l \tag{6}$$

$$0\le z(i,k),y(i,j),w(i,l),\alpha(i,j,k),\beta(i,j,l) \qquad \forall i,j,k,l \tag{7}$$

The variables $y(i,j), z(i,k)$ and $w(i,l)$ respectively indicate if an agent $j \in V_1$, $k \in V_2$, $l \in V_3$ received item $i \in I$. Constraints 6 state that an agent can get at most one item. Constraint 5 says that if $\alpha(i,j,k) = 1$, then both $y(i,j)$ and $z(i,k)$ must also be equal to one. Constraint 3 states that if $\beta(i,j,l) = 1$, then both $y(i,j)$ and $w(i,l)$ must also be equal to one. Finally, note that if an agent $l \in V_3$ contributes one unit of externality to an agent $j \in V_1$ for an item $i \in I$, then there must be some agent $k \in F_j^1 \cap F_l^1$ in V_2 who received the same item. This condition is encoded in constraint 4. Thus, we have the following lemma.

▶ **Lemma 8.** *The LP is a valid relaxation of the* Restricted-Welfare *problem.*

Before proceeding towards the rounding scheme, we perform a preprocessing step as described in the next lemma.

▶ **Lemma 9.** *In polynomial time, we can get a feasible solution to the LP that gives an* $O(\log m)$ *approximation to the optimal objective, and ensures that each* $\alpha(i,j,k), \beta(i,j,l)$, $y(i,j), w(i,l) \in \{0,\gamma\}$ *for some real number* $\gamma \in [0,1]$, *and that each* $z(i,k) \le \gamma$.

We now present the rounding scheme for LP (see Algorithm 1). Here, the set W_i denotes the set of agents that have not yet been assigned any item when the rounding scheme enters the For loop for item i (see Step 2). Note that the sets T_i might overlap, but these conflicts are resolved in Line 9 by intersecting T_i with W_i, which is disjoint with all previous $T_j, j < i$.

Algorithm 1 Rounding Scheme for LP

1. In accordance with Lemma 9, compute a feasible solution to the LP.
   Set $T_0 \leftarrow \emptyset$, and $W_0 \leftarrow V = V_1 \cup V_2 \cup V_3$.
2. For all items $i \in I = \{1,\ldots,m\}$:
3. Set $W_i \leftarrow W_{i-1} \setminus T_{i-1}$, and $T_i \leftarrow \emptyset$.
4. Pick a value η_i uniformly at random from $[0,1]$.
5. If $\eta_i \le \gamma$:
6. For all nodes $j \in V_1$:
   If $y(i,j) = \gamma$, then with probability $1/4$, set $T_i \leftarrow T_i \cup \{j\}$.
7. For all nodes $l \in V_3$:
   If $w(i,l) = \gamma$, then with probability $1/4$, set $T_i \leftarrow T_i \cup \{l\}$.
8. For all nodes $k \in V_2$:
   With probability $z(i,k)/(4\gamma)$, set $T_i \leftarrow T_i \cup \{k\}$.
9. Assign item i to all nodes in $W_i \cap T_i$, i.e., set $\mathcal{A}(t) \leftarrow i$ for all $t \in W_i \cap T_i$.
10. Return the (random) assignment $\mathcal{A}$.

▶ **Lemma 10.** *For all* $t \in V$ *and all* $i \in I$, *we have* $\mathbf{P}[t \in W_i] \ge 3/4$. *Thus,* $\mathbf{P}[\{t_1,t_2,t_3\} \subseteq W_i] \ge 1/4$ *for all* $t_1, t_2, t_3 \in V$.

Proof. We will prove the lemma for a node in V_1, the argument extends to $V_2 \cup V_3$.

Fix any node $j \in V_1$ and any item $i \in I$, and consider an indicator random variable $\Gamma_{i'j}$ that is set to one iff $j \in T_{i'}$. It is easy to check that $\mathbf{E}[\Gamma_{i'j}] = y(i', j)/4$ for all items $i' \in I$. By constraint 6 and linearity of expectation, we thus have: $\mathbf{E}[\sum_{i'<i} \Gamma_{i'j}] = \sum_{i'<i} y(i', j)/4 \leq 1/4$. Applying Markov's inequality, we get $\mathbf{P}[\sum_{i'<i} \Gamma_{i'j} = 0] \geq 3/4$. In other words, with probability at least $3/4$, we have that $j \notin T_{i'}$ for all $i' < i$. Under this event, we must have $j \in W_i$.

We have $\mathbf{P}[t \notin W_i] \leq 1/4$ for all $t \in \{t_1, t_2, t_3\}$. $\mathbf{P}[\{t_1, t_2, t_3\} \subseteq W_i] \geq 1/4$ now follows from applying union-bound over these three events. ◀

In the first step, when we find a feasible solution to the LP in accordance with Lemma 9, we lose a factor of $O(\log m)$ in the objective. Below, we will show that the remaining steps in the rounding scheme result in a loss of at most a constant factor in the approximation ratio.

For all items $i \in I$, nodes $j \in V_1$, and nodes $k \in F_j^1$, $l \in F_j^2$, we define the random variables $X(i, j, k)$ and $Y(i, j, l)$. Their values are determined by the outcome $\mathcal{A}$ of our randomized rounding. To be more specific, we have that $X(i, j, k) = 1$ if both j and k receive item i, and $X(i, j, k) = 0$ otherwise. Further, $Y(i, j, l) = 1$ if both j and l receive item i *and* there is some node in $F_j^1 \cap F_l^1$ that also received item i, and $Y(i, j, l) = 0$ otherwise. Now, the valuation of any agent $j \in V_1$ from the (random) assignment $\mathcal{A}$ is:

$$u_j(\mathcal{A}) = \sum_{i \in I} \left(\sum_{k \in F_j^1} \lambda_{ij} \cdot X(i, j, k) + \sum_{l \in F_j^2} \lambda_{ij} \cdot Y(i, j, l) \right) \tag{8}$$

We will analyze the expected contribution of the rounding scheme to each term in the LP-objective. Towards this end, we prove the following lemmas.

▶ **Lemma 11.** *For all $i \in I, j \in V_1, k \in F_j^1$, we have $\mathbf{E}_{\mathcal{A}}[X(i, j, k)] \geq \delta \cdot \alpha(i, j, k)$, where $\delta > 0$ is a sufficiently small constant.*

▶ **Lemma 12.** *For all $i \in I, j \in V_1, l \in F_j^2$, we have $\mathbf{E}_{\mathcal{A}}[Y(i, j, l)] \geq \delta \cdot \beta(i, j, l)$, where δ is a sufficiently small constant.*

Proof. Fix an item $i \in I$, a node $j \in V_1$ and a node $l \in F_j^2$. If $\beta(i, j, l) = 0$ the lemma is trivially true. Otherwise suppose for the rest of the proof that $\beta(i, j, l) = y(i, j) = w(i, l) = \gamma$.

Let $\mathcal{E}_i$ be the event that $\eta_i \leq \gamma$ (see Step 4 in Algorithm 1). Let $Z(i, k)$ be an indicator random variable that is set to one iff node $k \in V_2$ is included in the set T_i by our rounding scheme (see Step 8 in Algorithm 1). We have:

$$\mathbf{P}[\mathcal{E}_i] = \gamma, \text{ and } \mathbf{P}[Z(i, k) = 1 \,|\, \mathcal{E}_i] = z(i, k)/4\gamma \text{ for all } k \in V_2 \tag{9}$$

Thus, conditioned on the event $\mathcal{E}_i$, the expected number of common neighbors of j and l who are included in the set T_i is given by

$$\mu_i := \mathbf{E}\left[\sum_{k \in F_j^1 \cap F_l^1} Z(i, k) \,\middle|\, \mathcal{E}_i \right] = \sum_{k \in F_j^1 \cap F_l^1} z(i, k)/4\gamma \geq 1/4 \tag{10}$$

Note that conditioned on the event $\mathcal{E}_i$, the random variables $Z(i, k)$ are mutually independent. Thus, applying Chernoff bound on Equation 10, we infer that with constant probability, at least one common neighbor of j and l will be included in the set T_i. To be more precise, define $T_{i,j,l} = T_i \cap F_j^1 \cap F_l^1$. For some sufficiently small constant δ_1, we have:

$$\mathbf{P}\left[T_{i,j,l} \neq \emptyset \,\middle|\, \mathcal{E}_i \right] = \mathbf{P}\left[\sum_{k \in F_j^1 \cap F_l^1} Z(i, k) > 0 \,\middle|\, \mathcal{E}_i \right] \geq 1 - e^{-1/8} = \delta_1 \tag{11}$$

Let $\mathcal{E}_{i,j,l}$ be the event that the following two conditions hold simultaneously: (a) $T_{i,j,l} \neq \emptyset$, AND (b) j, l, and an arbitrary node from $T_{i,j,l}$—each of these three nodes is included in W_i. Now, Equation 11 and Lemma 10 imply that $\mathbf{P}[\mathcal{E}_{i,j,l} \,|\, \mathcal{E}_i] \geq \delta_2$ for $\delta_2 = \delta_1/4$. Putting all these observations together, we obtain that $\mathbf{P}[Y(i,j,l) = 1] = \mathbf{P}[\mathcal{E}_i] \cdot \mathbf{P}[\mathcal{E}_{i,j,l} \,|\, \mathcal{E}_i] \cdot \mathbf{P}[j, l \in T_i \,|\, \mathcal{E}_{i,j,l} \cap \mathcal{E}_i] = \gamma \cdot \delta_2 \cdot (1/4) \cdot (1/4) = \delta \cdot \gamma = \delta \cdot \beta(i,j,l)$ for $\delta = \delta_2/16$. ◀

▶ **Theorem 13.** *The rounding scheme in Algorithm 1 gives an $O(\log m)$-approximation to the* Restricted-Welfare *problem.*

Proof. In the first step, when we find a feasible solution to the LP in accordance with Lemma 9, we lose a factor of $O(\log m)$ in the objective. At the end of the remaining steps, the expected valuation of an agent $j \in V_1$ is given by:

$$\mathbf{E}_{\mathcal{A}}[u_j(\mathcal{A})] = \sum_{i \in I} \lambda_{ij} \cdot \Bigg(\sum_{k \in F_j^1} \mathbf{E}_{\mathcal{A}}[X(i,j,k)] + \sum_{l \in F_l^1} \mathbf{E}_{\mathcal{A}}[Y(i,j,l)] \Bigg)$$

$$= \Theta\Bigg(\sum_{i \in I} \lambda_{ij} \cdot \big(\sum_{k \in F_j^1} \alpha(i,j,k) + \sum_{l \in F_l^1} \beta(i,j,l) \big) \Bigg)$$

The first equality follows from linearity of expectation, while the second equality follows from Lemma 11 and Lemma 12. Thus, the expected valuation of any agent in V_1 is within a constant factor of the fractional valuation of the same agent under the feasible solution to the LP obtained at the end of Step 1 (see Algorithm 1). Summing over all the agents in V_1, we get the theorem. ◀

We can generalize the above approach to the following setting: Each user j is given an integer c_j and can be assigned up to c_j different items (each at most once). For this we replace for each item i and node j the constraint $\sum_i y(i,j) \leq 1$ by the two constraints $\sum_i y(i,j) \leq c_j$ and $y(i,j) \leq 1$ and adapt the proof of Lemma 10.

Finally, we state NP-hardness for linear externalities, not only in the 2-hop setting but also for 1-hop. [3]

▶ **Theorem 14.** *Maximizing social welfare under linear externalities is* NP*-hard.*

5 Constant Approximation for Step Function Externalities

In this section, our goal is to maximize the social welfare when the agents have general step function externalities, i.e., to receive externality an agent needs a certain number of 1- and 2-hop neighbors having the same product. We will show that no constant factor approximation is possible unless a bound on the number of neighbors an agent needs to receive externality is given. Thus we consider the case of 2-step function externalities, where only two neighbors are needed (see Definition 3) and give a $\frac{1}{6} \cdot (1 - 1/e)$-approximation algorithm for this problem. Notice that if we consider step functions that just require one neighbor the problem reduces to the 1-hop step function scenario in [6]. However, our algorithm gives a $\frac{1}{2} \cdot (1 - 1/e)$-approximation for this scenario improving the result in [6].

In the following we assume 2-step function externalities. Let $G_{V'}$ denote the subgraph induced by $V' \subseteq V$. For the rest of this section, the term "triple" will refer to any (unordered)

[3] Theorem 3.1 in [6] claims that the welfare maximization problem for linear 1-hop externality functions in complete graphs is MaxSNP-hard, which would imply our result, but this claim is not true.

set of three nodes $T = \{j_1, j_2, j_3\}$ such that G_T is connected. Similarly, the term "pair" will refer to any (unordered) set of two nodes $\{j_1, j_2\}$ that are connected by an edge in E.

We first compute a maximal collection of mutually disjoint triples in the graph G. We denote this collection by $\mathcal{T}$, and let $V(\mathcal{T}) = \bigcup_{T \in \mathcal{T}} T \subseteq V$. The graph $G_{V \setminus V(\mathcal{T})}$, by definition, consists of a mutually disjoint collection of pairs (say $\mathcal{P}$) and a set of isolated nodes (say B). We thus have the following lemma.

▶ **Lemma 15.** *In $G = (V, E)$, there is no edge that connects a node $j \in B$ with another node in B or with a node belonging to a pair in $\mathcal{P}$. Furthermore, there is no edge that connects two nodes j, j' belonging to two different pairs $P, P' \in \mathcal{P}$.*

▶ **Definition 16.** An assignment $\mathcal{A}$ is *consistent* iff two agents get the same item whenever they belong to the same triple or the same pair. To be more specific, for all $j, j' \in V$, we have that $\mathcal{A}(j) = \mathcal{A}(j')$ if either (a) $j, j' \in T$ for some triple $T \in \mathcal{T}$ or (b) $\{j, j'\} \in \mathcal{P}$.

The next lemma shows that by losing a factor of 6 in the approximation ratio, we can focus on maximizing the social welfare via a consistent assignment.

▶ **Lemma 17.** *The social welfare from the optimal consistent assignment is at least $(1/6) \cdot \textsc{Opt}$, where $\textsc{Opt}$ is the maximum social welfare over all assignments.*

Proof. Let $\mathcal{A}^*$ be an assignment (not necessarily consistent) that gives maximum social welfare. We convert it into a (random) consistent assignment $\mathcal{A}$ as follows. For each triple $\{j_1, j_2, j_3\} \in \mathcal{T}$, we pick one of the items $\mathcal{A}^*(j_1), \mathcal{A}^*(j_2), \mathcal{A}^*(j_3)$ uniformly at random, and assign that item to all the three agents j_1, j_2, j_3. Similarly, for each pair $\{j_1, j_2\} \in \mathcal{P}$, we pick one of the items $\mathcal{A}^*(j_1), \mathcal{A}^*(j_2)$ uniformly at random, and assign that item to both the agents j_1, j_2. The events corresponding to different triples and pairs are mutually independent. Finally, the remaining agents (those who are in B) get the same items as in $\mathcal{A}^*$. It is easy to see that the resulting assignment $\mathcal{A}$ is consistent. We claim that $\mathbf{E}[u_j(\mathcal{A})] \geq (1/6) \cdot u_j(\mathcal{A}^*)$ for all $j \in V$. To prove this claim, we consider three cases.

Case 1 ($j \in B$): Let $\mathcal{A}^*(j) = i$. Since $j \in B$, it always gets the same item under $\mathcal{A}$, i.e., $\mathcal{A}(j) = i$. Now, if $u_j(\mathcal{A}^*) = 0$, then the claim is trivially true. Otherwise it must be the case that $\mathcal{A}^*(j') = i$ for some neighbor j' of j. Since $j \in B$, this neighbor j' must be part of some triple $T \in \mathcal{T}$ (see Lemma 15). With probability at least 1/3 all the three nodes in T are assigned item i under $\mathcal{A}$ and at least two nodes of T are in the 2-hop neighborhood of j. In that event j gets the same valuation as in $\mathcal{A}^*$, and we have that $\mathbf{E}[u_j(\mathcal{A})] \geq (1/3) \cdot u_j(\mathcal{A}^*)$.

Case 2 (j belongs to a pair in $\mathcal{P}$): Consider the pair $P = \{j, j'\} \in \mathcal{P}$, which has j and another node (say j') as its members. Let $\mathcal{A}^*(j) = i$. As in Case 1, if $u_j(\mathcal{A}^*) = 0$, then the claim is trivially true. Otherwise it must be the case that there exists a node j'' with $\mathcal{A}^*(j'') = i$ such that j'' is either a neighbor of j or a neighbor of j'. Since $\{j, j'\} \in \mathcal{P}$, this agent j'' must be part of some triple $T \in \mathcal{T}$ (see Lemma 15). Let $\mathcal{E}_1$ be the event that all the three nodes in T are assigned item i under $\mathcal{A}$. Similarly, let $\mathcal{E}_2$ be the event that both the nodes $j, j' \in P$ get the same item i under $\mathcal{A}$. Since these two events are mutually independent, we have that $\mathbf{P}[\mathcal{E}_1 \cap \mathcal{E}_2] \geq (1/3) \cdot (1/2) = 1/6$, and in the event $\mathcal{E}_1 \cap \mathcal{E}_2$, we have $u_j(\mathcal{A}) = u_j(\mathcal{A}^*)$. It follows that $\mathbf{E}[u_j(\mathcal{A})] \geq (1/6) \cdot u_j(\mathcal{A}^*)$.

Case 3 (j belongs to a triple in $\mathcal{T}$): Consider the triple $T = \{j, j', j''\} \in \mathcal{T}$ which has, besides j, two other nodes (say j' and j'') as its members. With probability at least 1/3, all these three nodes are assigned item $\mathcal{A}^*(j)$ under $\mathcal{A}$, and in this event we have $u_j(\mathcal{A}) \geq u_j(\mathcal{A}^*)$. It follows that $\mathbf{E}[u_j(\mathcal{A})] \geq (1/3) \cdot u_j(\mathcal{A}^*)$.

Now, we take a sum of the inequalities $\mathbf{E}[u_j(\mathcal{A})] \geq (1/6) \cdot u_j(\mathcal{A}^*)$ over all agents $j \in V$, and by linearity of expectation infer that the expected social welfare under the consistent assignment $\mathcal{A}$ is within a factor of 6 of the optimal social welfare. This concludes the proof of the lemma. ◀

Next, we will give an $(1-1/e)$-approximation algorithm for finding a consistent assignment of items that maximizes the social welfare. Along with Lemma 17, this will imply the main result of this section (see Theorem 20).

We use the term "resource" to refer to either a pair $P \in \mathcal{P}$ or an agent $j \in B$. Let $\mathcal{R} = \mathcal{P} \cup B$ denote the set of all resources. We say that a resource $r \in \mathcal{R}$ *neighbors* a triple $T \in \mathcal{T}$ iff in the graph $G = (V, E)$ either (a) $r = \{j, j'\} \in \mathcal{P}$ and some node in $\{j, j'\}$ is adjacent to some node in T, or (b) $r = j \in B$ and j is adjacent to some node in T. We slightly abuse the notation (see Section 2) and let $N(T) \subseteq \mathcal{R}$ denote the set of resources that are neighbors of $T \in \mathcal{T}$.

By definition, every consistent assignment ensures that if two agents belong to the same triple in $\mathcal{T}$ (resp. the same pair in $\mathcal{P}$), then both of them get the same item. We say that *the item is assigned to a triple (resp. resource)*. Note that the triples do not need externality from outside. To be more specific, the contribution of a triple $T \in \mathcal{T}$ to the social welfare is always equal to $\sum_{j \in T} \lambda_{i,j}$, where i is the item assigned to T. Resources, however, do need outside externality, which by Lemma 15 can come only from a triple in $\mathcal{T}$.

▶ **Lemma 18.** *In a consistent assignment, if a resource $r \in \mathcal{R}$ makes a positive contribution to the social welfare, then it neighbors some triple $T_r \in \mathcal{T}$, and both the resource r and the triple T_r receive the same item.*

Proof. If a resource contributes a nonzero amount to the social welfare, then it must receive nonzero externality from the assignment. By Lemma 15, such externality can come only from a triple in $\mathcal{T}$. The lemma follows. ◀

Thus, given a consistent assignment $\mathcal{A}$ consider the following mapping $T_{\mathcal{A}}(r)$ of a resource $r \in \mathcal{R}$ to triples in $\mathcal{T}$ in accordance with Lemma 18: If the resource r makes zero contribution towards the social welfare (a case not covered by the lemma), then we let $T_{\mathcal{A}}(r)$ be any arbitrary triple from $\mathcal{T}$. Otherwise $T_{\mathcal{A}}(r)$ denotes an (arbitrary) neighboring triple of $\mathcal{T}$ that receives the same item as r. We say that the triple $T_{\mathcal{A}}(r)$ *claims* the resource r.

For ease of exposition, let $\lambda_{i,r}(T)$ be the valuation of the resource r when both the resource and the triple T that claims it get item $i \in I$, i.e.,

$$\lambda_{i,r}(T) = \begin{cases} \lambda_{i,j} + \lambda_{i,j'} & \text{if } r = \{j, j'\} \in \mathcal{P} \text{ and } r \in N(T); \\ \lambda_{i,j} & \text{if } r = j \in B \text{ and } r \in N(T); \\ 0 & \text{if } r \notin N(T). \end{cases}$$

Now, any consistent assignment $\mathcal{A}$ can be interpreted as follows. Under such an assignment, every triple $T \in \mathcal{T}$ claims the subset of the resources $S_T = \{r \in \mathcal{R} : T_{\mathcal{A}}(r) = T\}$; the subsets corresponding to different triples being mutually exclusive. A triple T and the resources in S_T all get the same item (say $i \in I$). The valuation obtained from them is $u_T(S_T, i) = \sum_{j \in T} \lambda_{i,j} + \sum_{r \in S_T} \lambda_{i,r}(T)$.

If our goal is to maximize the social welfare, then, naturally, for every triple T, we will pick the item that maximizes $u_T(S_T, i)$, thereby extracting a valuation of $u_T(S_T) = \max_i u_T(S_T, i)$. The next lemma shows that this function is fractionally subadditive.

▶ **Lemma 19.** *The function $u_T(S_T)$ is fractionally subadditive in S_T.*

The preceding discussion shows that the problem of computing a consistent assignment for welfare maximization is equivalent to the following setting. We have a collection of triples $\mathcal{T}$, and a set of resources $\mathcal{R}$. We will distribute these resources amongst the triples, i.e., every triple T will get a subset $S_T \subseteq \mathcal{R}$, and these subsets will be mutually exclusive. The goal is to maximize the sum $\sum_{T \in \mathcal{T}} u_T(S_T)$, where the functions $u_T(\cdot)$'s are fractionally subadditive. By a celebrated result of Feige [9], we can get an $(1 - 1/e)$-approximation algorithm for this problem if we can implement the following subroutine (called *demand oracle*) in polynomial time: Each resource r is given a "cost" $p(r)$ and we need to determine for each triple T a set of resources S_T^* that maximizes $u_T(S_T) - \sum_{r \in S_T} p(r)$ over all sets S_T. Such a demand oracle can be implemented in polynomial time using a simple greedy algorithm for each T and each item i: Add a resource r to S_T^* iff $\lambda_{i,r}(T) > p(r)$. The result of the approximation algorithm assigns each triple T a subset S_T and we then pick the item i that maximizes $u_T(S_T, i)$ over all items i. Together with Lemma 15, this implies the theorem stated below.

▶ **Theorem 20.** *We can get a polynomial-time $\frac{1}{6} \cdot (1 - 1/e)$-approximation algorithm for the problem of maximizing social welfare under 2-step function externalities.*

The algorithm can be easily adapted for 1-hop step function externalities. The difference is that instead of computing a maximal collection $\mathcal{T}$ of mutually disjoint triples, one computes a maximal collection of mutually disjoint pairs.

▶ **Theorem 21.** *We can get a polynomial-time $\frac{1}{2} \cdot (1 - 1/e)$-approximation algorithm for maximizing social welfare under 1-step function externalities.*

Finally, we present our hardness results for step functions. By a reduction from Max Independent Set we can show that, for unbounded s, there is no constant factor approximation. The main idea is that we modify the graph such that we replace each edge by a path of length three and each of the original nodes j wants a different item, while j can only get positive externalities when having a support of $2\delta_j$ (δ_j the node degree of j). The valuations of the newly introduced nodes are set to 0. That is, nodes that are adjacent in the original graph have two common neighbors in the constructed graph, want different items, need all their neighbors as support, and thus only one of them can have positive valuation.

▶ **Theorem 22.** *For any $\varepsilon > 0$ the problem of maximizing social welfare under arbitrary s-step function externalities is not approximable within $O(n^{1/4-\varepsilon})$ unless $\mathsf{NP} = \mathsf{P}$, and not approximable within $O(n^{1/2-\varepsilon})$ unless $\mathsf{NP} = \mathsf{ZPP}$.*

Second, we show that maximizing social welfare under 2-step function externalities is APX-hard and thus no PTAS can exists. This is by a reduction from Max Coverage. The APX-hardness for two items is by a reduction from SAT.

▶ **Theorem 23.** *The problem of maximizing social welfare under step function externalities is APX-hard, in particular, there is no polynomial time $1 - \frac{1}{e} + \epsilon$-approximation algorithm (unless $\mathsf{P} = \mathsf{NP}$). Furthermore, the problem remains APX-hard (although with a larger constant) even if there are only two items.*

Acknowledgments The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 317532 and from the Vienna Science and Technology Fund (WWTF) through project ICT10-002.

References

1 Hessameddin Akhlaghpour, Mohammad Ghodsi, Nima Haghpanah, Vahab S. Mirrokni, Hamid Mahini, and Afshin Nikzad. Optimal iterative pricing over social networks. In *6th WINE*, pages 415–423, 2010.

2 Noga Alon, Michal Feldman, Ariel D. Procaccia, and Moshe Tennenholtz. A note on competitive diffusion through social networks. *Inf. Process. Lett.*, 110(6):221–225, 2010.

3 Nima Anari, Shayan Ehsani, Mohammad Ghodsi, Nima Haghpanah, Nicole Immorlica, Hamid Mahini, and Vahab S. Mirrokni. Equilibrium pricing with positive externalities. *Theor. Comput. Sci.*, 476:1–15, 2013.

4 David Arthur, Rajeev Motwani, Aneesh Sharma, and Ying Xu. Pricing strategies for viral marketing on social networks. In *5th WINE*, pages 101–112, 2009.

5 Bernard Bensaid and Jean-Philippe Lesne. Dynamic monopoly pricing with network externalities. *Int. J. of Industrial Organization*, 14(6):837–855, 1996.

6 Anand Bhalgat, Sreenivas Gollapudi, and Kamesh Munagala. Mechanisms and allocations with positive network externalities. In *13th EC*, pages 179–196, 2012.

7 Sayan Bhattacharya, Dmytro Korzhyk, and Vincent Conitzer. Computing a profit-maximizing sequence of offers to agents in a social network. In *8th WINE*, pages 482–488, 2012.

8 Pradeep Dubey, Rahul Garg, and Bernard De Meyer. Competing for customers in a social network: The quasi-linear case. In *2nd WINE*, pages 162–173, 2006.

9 Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.

10 Scott L. Feld. Why your friends have more friends than you do. *American J. of Sociology*, 96(6):1464–1477, 1991.

11 Dimitris Fotakis and Paris Siminelakis. On the efficiency of influence-and-exploit strategies for revenue maximization under positive externalities. In *8th WINE*, pages 270–283, 2012.

12 Sanjeev Goyal and Michael Kearns. Competitive contagion in networks. In *44th STOC*, pages 759–774, 2012.

13 Nima Haghpanah, Nicole Immorlica, Vahab S. Mirrokni, and Kamesh Munagala. Optimal auctions with positive network externalities. *ACM Trans. Economics and Comput.*, 1(2):13:1–13:24, 2013.

14 Keith N. Hampton, Lauren Sessions Goulet, Cameron Marlow, and Lee Rainie. Why most facebook users get more than they give. *Pew Internet & American Life Project*, 2012.

15 Jason Hartline, Vahab S. Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *17th WWW*, pages 189–198, 2008.

16 Xinran He and David Kempe. Price of anarchy for the n-player competitive cascade game with submodular activation functions. In *9th WINE*, pages 232–248, 2013.

17 Matthew O. Jackson and Brian W. Rogers. Meeting strangers and friends of friends: How random are social networks? *American Economic Review*, 97(3):890–915, 2007.

18 David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *9th KDD*, pages 137–146, 2003.

19 Vahab S. Mirrokni, Sebastien Roch, and Mukund Sundararajan. On fixed-price marketing for goods with positive network externalities. In *8th WINE*, pages 532–538, 2012.

20 Sunil Simon and Krzysztof R. Apt. Choosing products in social networks. In *8th WINE*, pages 100–113, 2012.

21 Reiko Takehara, Masahiro Hachimori, and Maiko Shigeno. A comment on pure-strategy nash equilibria in competitive diffusion games. *Inf. Process. Lett.*, 112(3):59–60, 2012.

22 Vasileios Tzoumas, Christos Amanatidis, and Evangelos Markakis. A game-theoretic analysis of a competitive diffusion process over social networks. In *8th WINE*, pages 1–14, 2012.