Improved Recommendation Systems

Extended Abstract

Baruch Awerbuch*   Boaz Patt-Shamir†   David Peleg‡   Mark Tuttle§

July 19, 2004

Abstract

We consider a model of competitive recommendation systems proposed by Drineas *et al.* [DKR02]. In recommendation systems (such as book or movie recommendations), the system tracks which product each user chose in the past, and tries to deduce which other products an asking user is likely to be satisfied with. Obviously, recommendation systems can be effective only for users who share preferences with many other users. Such users are said to belong to a "dominant type." Current approaches to on-line recommendation systems involve using Singular Value Decomposition (SVD), which is computationally intensive and, more important, is applicable only under additional strong conditions. Specifically, correctness is guaranteed in [DKR02] only if users of different dominant types essentially do not share a product they like ("type separability"), and only if the number of users in non-dominant types is significantly smaller than the number of users in dominant types ("gap assumption"). The complexity of that algorithm is $O(mn)$, where m and n denote the number of users and products, respectively. Here, it is shown that in fact, very simple combinatorial algorithms can make good recommendations without using SVD. Our algorithms require neither the type separability nor the gap assumption, they are naturally amenable to distibuted computation, and their complexity is lower. In particular, the paper presents a centralized algorithm with $O(m + n)$ running time, and a distributed algorithm that can be implemented in a peer-to-peer model even in the presence of adaptively colluding malicious players, with only logarithmic overhead.

1 Introduction

Motivation. We consider an abstraction of the *recommendation problem* proposed by [DKR02, KRRT98], that can be described informally as follows. There are m users and n products. Given a product, a user can tell whether he considers it to be good. A recommendation algorithm tracks the

---

*Computer Science Department, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218. Research conducted while visiting CRL. Email: `baruch@acm.org`.

†Cambridge Research Laboratory, HP Labs, One Cambridge Center, Cambridge, MA 02142. On leave from Tel Aviv University. Email: `boaz@hp.com`.

‡Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot 76100, Israel. Research conducted while visiting CRL. Email: `david.peleg@weizmann.ac.il`.

§Cambridge Research Laboratory, HP Labs, One Cambridge Center, Cambridge, MA 02142. Email: `mark.tuttle@hp.com`.

choices of all users; in a basic step, the algorithm recommends a product to a user and asks what the user thought about the product. The task of the algorithm is to recommend a good product to most users, and its performance measure is the number of recommendations made to a user until a good product is recommended, as well as the number of users that are eventually satisfied. We believe that most readers have been exposed to commercial recommendation systems in one form or another (say, book recommendations in Amazon).

Obviously, recommendation systems (also known as *collaborative filtering* systems) may only be useful to a user who shares preferences with others: a user with esoteric preferences cannot rely on others to help him find a product he likes. In [DKR02], this property is captured as follows. On the outset, each user is assumed to have a grade for each product, so that each user is represented by his vector of grades. (The grades are revealed in an on-line fashion, but are assumed to exist always.) To model popular preferences, it is assumed that there exists a set of k "canonical" vectors (where k is a small constant), such that most user vectors are the result of adding random noise to one of the canonical vectors. Such a user is said to belong to a *dominant type*. The assumption that most users belong to a small number of dominant types appears to be a reasonable simplification; the algorithm aims at satisfying them.

The recommendation system proposed in [DKR02] is a centralized algorithm based on *matrix reconstruction* [AFK+01, SKKR00]. The algorithm of [DKR02], denoted MR henceforth, can be roughly described as follows. MR first chooses a small set of users and asks them to try *all* products and grade them. Based on these responses, each remaining user is asked to try a carefully chosen small set of products; the latter responses are used to associate users with types, which in turn are used to recommend products to users. It is shown that MR guarantees (with high probability) that each user from a dominant type gets a good recommendation. The heart of MR is a low-rank approximation of the full user-product grade matrix (the low rank is k, the number of dominant types). This approximation is based on Singular Value Decomposition (SVD). The algorithm works only when a few severe conditions are met. First is the *type separability* assumption, which requires that the canonical vectors be nearly orthogonal; this means that two users from different dominant types essentially do not share any product they both like. Second is the *gap* assumption, which says that any non-dominant type is far less popular than any dominant type. To complicate things further, MR must know the number of dominant types in order to produce correct results. Lastly, the computational complexity of the algorithm is $O(mn)$.

In this paper we show, surprisingly, that the algebraic approach can be abandoned for this particular problem without any loss in the quality of the results, and while dramatically simplifying the algorithms. Moreover, both type separability and gap assumptions turn out to be unnecessary. Most important, we show that a distributed peer-to-peer solution is possible, even in the presence of a very powerful adversary.

Broadly speaking, the basic idea behind our algorithms is the usage of "best-seller lists" via random sampling. Our distributed algorithm is based on techniques developed in the context of auction systems like eBay [APPT04].

In more detail, we present the following results. First, we give a centralized algorithm that runs in time $O(m + n)$ assuming that most users fall within one of $k = O(1)$ large types (but neither type

separability nor the existence of a gap are assumed). We note that $\Omega(n + m)$ is an obvious lower bound on the number of recommendations made by any recommendation algorithm even when all users are identical (i.e., $k = 1$): if only one of the products is good, then $\Omega(n)$ products must be recommended by the algorithm until that good product is discovered; and in addition, the algorithm must recommend to each of the m users at least one good product.

Our main result is a distributed solution to the recommendation problem. We show that under any asynchronous schedule, the total number of recommendations required to satisfy any set of users with shared interest is $O(n + m \log m)$ with high probability, coming within a logarithmic factor of the optimal. Our algorithm can tolerate malicious users in the form of an omniscient, adaptive, Byzantine adversary. (This is an important consideration for distributed peer-to-peer systems.) Another important feature of the algorithm is that it does not require users to be rigidly classified into types: users may have any preference vector whatsoever without changing the algorithm; only the performance will be affected, in the sense that users with esoteric preferences will have to work more.

We note that our algorithms imply that type separation is equivalent to full matrix reconstruction, but it is not required for competitive recommendations.

Related work. Most prior research on recommendation systems focused on a centralized, off-line version of the problem, where the algorithm is presented with a lot of historical preference data, and the task is to generate a single recommendation that maximizes the utility to the user. This is usually done by heuristically identifying clusters of users [RIS+94] (or products [SKKR01]) in the data set, and using past grades by users in a cluster to predict future grades by other users in the same cluster. SVD was shown also to be effective for the off-line problem [SKKR00]. Some of these systems enjoy industrial success, but they are known to perform poorly when prior data is less than plentiful [SPUP02], and they are extremely vulnerable even to mild attacks [LR04, OHS04]. Canny [Can02] gives a distributed secure and private SVD computation, but it is for the off-line version of the problem.

Theoretical studies of recommendation systems usually take the latent variable model approach: a stochastic process is assumed to generate noisy observations, and the goal of an algorithm is to approximate some unknown parameters of the model. Kumar *et al.* [KRRT98] study the off-line problem for a model where preferences are identified with past choices (purchases). In this model there are clusters of products. Each user has a probability distribution over clusters; a user first chooses a cluster by his distribution, and then chooses a product uniformly at random from that cluster. The goal is to recommended a product from the user's most preferred cluster. Kleinberg and Sandler [KS03] generalized this model to the case where the choice within a cluster is governed by an arbitrary probability distribution, and also consider the mixture model, in which each cluster is a probability distribution over all products. Azar *et al.* [AFK+01] consider a model where there exists an unknown user-product preference matrix which can be approximated by a low-rank matrix. The system observes this matrix only after its entries were subjected to random additive noise and then to random omissions. They use SVD to reconstruct the original preferences.

In this work we use the on-line model of recommendation systems, as proposed in [DKR02], which is discussed extensively above, and even in more detail in Section 5. We note that measures and techniques used for off-line algorithms do not readily apply to the on-line scenario, since the off-line

case does not take into account the feedback effect an on-line algorithm has on its observed data, which skews the statistical properties of the observed data.

Paper organization. The remainder of this paper is organized as follows. We define the recommendation problem and our model in Section 2. In Sections 3 and 4 we give our centralized and distributed algorithms, respectively. In Section 5 we present a detailed discussion of our model in the context of previous models. Proofs are presented in the appendix.

2 Formal Model

The recommendation problem. Let us begin by defining the recommendation problem. Assume the existence of a set $\mathcal{U}$ of *users* and a set $\mathcal{P}$ of *products*. Denote $|\mathcal{U}| = m$ and $|\mathcal{P}| = n$. Each user has a numeric grade for each product. In most of this paper, we will assume for the sake of simplicity that the grades are binary, i.e., each user views some products as *good* and the rest as *bad* (more on that in Section 5). We use $good(u)$ to denote a function mapping each user $u \in \mathcal{U}$ to the set of products u views as good.

A *recommendation algorithm* has a single primitive operation, in which it atomically outputs a product name to a user ("making a recommendation"), asks the user for his rating of that product (good or bad), and gets the user's response as input. We often write "u probes p" in an algorithm description to mean that the algorithm recommends the product p to u and thus learns whether $p \in good(u)$. We define the set of products recommended to a user to be the algorithm's recommendation to that user. We want algorithms that recommend a small set of products to each user, and so that, with high probability, a large fraction λ of the users will find a good product in the set recommended to them.

The performance of an algorithm for the recommendation problem is evaluated in two ways. The traditional *time complexity* measures the worst-case number of RAM-model steps taken until the algorithm terminates. In addition, we measure the *recommendation complexity*, which is the total number of times users test recommended products. We further define the *individual recommendation complexity*, which is the worst-case number of products a single user tests. The recommendation complexity is obviously a lower bound on the time complexity, and our algorithms have essentially equal recommendation and time complexities.

The distributed model. We now turn to our model of a distributed system. Each user maintains a log of his probe results. An execution of the system consists of a sequence of user steps. In each atomic step, a user may send messages to other users asking them about their probe histories, get results back, and optionally probe a product (and keep the result in his local history record).[1] In contrast to the centralized model, where the algorithm can choose which user takes the next step, in the distributed model we assume that the order of user probes is under the control of an external *schedule*; the role of a distributed algorithm is limited to directing the user, as to which user to consult with and what product to probe during a step, when the user's turn comes up according to the schedule.

---

[1] We remark that in our distributed algorithm, only a single user is consulted at a step, and a valid response consists of at most one product, so communication complexity is kept at minimum.

We assume that some users are *honest* and some are *dishonest*. Honest users are required to follow the protocol, including, in particular, answering questions truthfully. Dishonest users are allowed to behave in an arbitrary (Byzantine) fashion, including giving false reports on their probes and colluding.

Formally, an execution of a recommendation algorithm is uniquely determined by the algorithm, the sequence of coins flips, and by three external entities:

1. the *user schedule* that determines the order in which users take steps,
2. the identity of *dishonest users* and of *good products*, and
3. the *adversary* that determines the behavior of the dishonest users.

The adversary we consider is the extremely powerful adaptive Byzantine adversary. Formally, an adversary is a function that takes the set of dishonest users, the set of goodness predicates, the schedule and a random binary sequence, and maps them to a sequence of actions for the dishonest users, telling them exactly what to do in each step. The random binary sequence represents the entire sequence of coin flip outcomes during the execution, including even future coin flips. Note that with this information, the adversary can reconstruct the entire state of the system at any point in time, and use any of this information to choose the next move for each dishonest user.

An *operating environment* consists of a user schedule, a set of dishonest users, a set of goodness predicates, and an adversary. The purpose of the operating environment is to factor out all of the nondeterministic choices made during an execution, leaving only the probabilistic choices to consider. When we deal with probabilities or compute expected values, we fix an operating environment and consider the distribution of executions with this operating environment induced by the coin flips. We note that the user schedule may schedule steps by a user long after the user has found a good product and halted, in which case the user just skips its step.

Each execution of the system yields a *probe sequence* consisting of the sequence of products probed during the execution. When we compute the recommendation complexity of a probe sequence, we count only the number of probes by the honest users, and ignore the probes by the dishonest users. Our goal is to minimize the recommendation complexity for the honest users.

3   Centralized algorithm

This section presents a centralized solution to the recommendation problem. The algorithm is a simplification of the algorithm in [DKR02], but it does not make the "gap" and "type separability" assumptions (see below). In this section we assume that all users are honest; dishonest users are considered in the next section.

Figure 1 gives the centralized algorithm $\mathcal{A}_{\mathrm{cntr}}(k, \gamma)$ that takes as parameters an integer $k \geq 0$ and a real number $0 < \gamma \leq 1$ and outputs recommendations for each user. The algorithm first selects a random sample of $k \ln(k/\gamma)$ users called the "committee," and recommends all n products to each committee member. Each committee member chooses the product he liked best, and these $k \ln(k/\gamma)$ are recommended to all remaining users.[2]

---

[2]The algorithm works as stated for arbitrary numeric grades, as well as for binary values.

Algorithm $\mathcal{A}_{\text{cntr}}(k, \gamma)$

Let $K = k \ln(k/\gamma)$.

Select K users a random. Call these users the *committee*.

Let each committee member u probe all n products, and let p_u be the product u views
as best (break ties arbitrarily).

Output $\{p_1, \ldots, p_K\}$ as the set of recommended products.

Figure 1: The centralized algorithm $\mathcal{A}_{\text{cntr}}(k, \gamma)$ for integer $k \geq 0$ and error $\gamma > 0$.

Informally, to analyze the algorithm we assume that users can be partitioned into equivalence classes called *types*, where each equivalence class represents a set of users with similar preferences. The success of the algorithm depends on the abundance of large types. Formally, we have the following definitions.

Definition 3.1 *A type T is a set of users $U(T)$ and a nonempty set of products $P(T)$ satisfying the condition that the set of good products for a user in $U(T)$ is exactly $P(T)$, namely,*

$$U(T) \subseteq \{u \in \mathcal{U} \mid P(T) = good(u)\}.$$

Given a collection $\mathcal{T}$ of types, we write $U(\mathcal{T}) \overset{\text{def}}{=} \bigcup_{T \in \mathcal{T}} U(T)$.

Definition 3.2 *Let $0 < \lambda \leq 1$ be a real number, and let $k > 0$ be an integer. A collection $\mathcal{T}$ of types is a (λ, k)-type-cover if $|U(\mathcal{T})| \geq \lambda m$ and $|U(T)| \geq m/k$ for each $T \in \mathcal{T}$.*

Let us highlight a few properties of our definition of type cover. First, the product sets of different types in a type cover may have arbitrary intersection. This is in contrast with the "type separability" assumption [DKR02], which requires the product sets to be nearly disjoint. This is quite a severe restriction: for example, if types were required to be separable, then a common user whose preferences intersect—but not exactly identify—with others, will remain "typeless." Second, types in a cover must have a minimum number of users, but nothing is required of types not in the cover. This in contrast to the "gap" property required by [DKR02], that stipulates that the number of users of each type in the cover is much larger than the number of users in each type not in the cover. Finally, we note that the definition of a type implies that if T and T' are types with different product sets, then $U(T) \cap U(T') = \emptyset$, i.e., they have disjoint user sets. This restriction will be lifted in the next section.

We now characterize the performance of $\mathcal{A}_{\text{cntr}}(k, \gamma)$ in the context of a (λ, k)-type-cover:

Theorem 3.1 *Let $\gamma > 0$. If there is a (λ, k)-type-cover for some $\lambda > 0$ and $k > 0$, then with probability at least $1 - \gamma$ the algorithm $\mathcal{A}_{cntr}(k, \gamma)$ generates a set of recommendations that satisfies λm users with recommendation and time complexity $O(k(m + n) \log \frac{k}{\gamma})$.*

The proof is given in the appendix; intuitively, each user in the type cover has, with high probability, a "representative" user from his type in the committee, whose job is to discover the product his "constituent" users like; this representative reports his findings to the benefit of all others. The reader may note that in some sense, committee members play a role similar to that of critics in human society.

6

Before we close this section, let us take another look at the algorithm. Theorem 3.1 bounds the probability that a user finds a good product after testing all $K = k \ln(k/\gamma)$ recommendations. It may be interesting also to understand how many products the user should test *on average* until it finds that good product. We assume that the order in which non-committee users try recommendations is a random permutation.

Theorem 3.2 *Let $\gamma > 0$. If there is a (λ, k)-type-cover $\mathcal{T}$ for some $\lambda > 0$ and $k > 0$, then the expected individual recommendation complexity for users in $\mathcal{T}$ is $O(k(1 + \frac{n \log(k/\gamma)}{m}))$.*

A detailed comparison of these results with those of [DKR02] appears in Section 5.

Critique of Committee Algorithms. Algorithm $\mathcal{A}_{\mathrm{cntr}}$ was proven to work correctly, but it has some serious disadvantages due to its basic committee-based approach. First, the individual recommendation complexity of $\mathcal{A}_{\mathrm{cntr}}$ is $\Omega(n)$, since all committee members must test all products before others can use their recommendations. This is not merely a formal objection, but rather a symptom of the real difficulty in implementing a committee-based algorithm (in [DKR02] it was proposed to compensate committee members for their efforts).

The second point is manifest in a distributed implementation with malicious users. In this case, even choosing a committee is a non-trivial task: how will membership be determined? If users decide on their own whether they are committee members, many dishonest users might "volunteer" into the committee and effectively hide the honest members. One possible workaround for this problem is a random beacon [Rab83] that generates public coin flips. But even if we had a random beacon at our disposal, any committee-based algorithm is vulnerable to *adaptive* Byzantine attacks that target committee members. This is a very practical threat: one real-world adaptive Byzantine adversary is the one who thwarts the algorithm by bribing committee members.

Subsequently, in the following section we present a distributed algorithm that does not use a committee and does not suffer from these problems.

4 Distributed Algorithm

This section presents a distributed solution to the recommendation problem that is resilient to arbitrarily malicious behavior from any fraction of the users. This algorithm shows that it is possible to do away with the committee altogether in a solution to the recommendation problem. Even more interestingly, the correctness of the algorithm does not depend on the existence of any kind of cover. The analysis of the algorithm's recommendation complexity does depend on covers, but it uses a significantly more relaxed notion of a cover by "special interest groups," and not type covers.

4.1 Algorithm

Our distributed algorithm $\mathcal{A}_{\mathrm{dist}}$, presented formally in Figure 2, is extremely simple: each honest user repeatedly either chooses a random product and probes it, or chooses a random user and probes the product that user recommends. This is done by sending a message to that user, to which the consulted user responds with the identity of the best product he's probed so far (ties broken arbitrarily). When

a user finds a good product, the user is said to be *satisfied* and he stops running the algorithm.

Algorithm $\mathcal{A}_{\text{dist}}$ followed by each user
  repeat
    Flip a coin.
    If the result is "heads," select a product uniformly at random and probe it.
    If the result is "tails," select a user uniformly at random and probe the product
      that user recommends.
  until a good product is found.
  Recommend the good product and halt.

Figure 2: The distributed algorithm $\mathcal{A}_{\text{dist}}$.

4.2 Analysis of recommendation complexity

The basic property of $\mathcal{A}_{\text{dist}}$ is stated in the following theorem, which bounds the total recommendation complexity of any set of users with a shared good product under any asynchronous schedule and any adversary.

Theorem 4.1 *Let U be a nonempty set of honest users and P be a nonempty set of products such that $P \subseteq good(u)$ for every $u \in U$. For every operating environment, the expected recommendation complexity for users in U in an execution of Algorithm $\mathcal{A}_{dist}$ is at most $2\left(n/|P| + m\ln|U|\right)$.*

The proof is given in the appendix. Intuitively, the analysis divides the execution into two stages: in the first stage, the users search for a good object until one of them finds it due to the "random product" probes; in the second stage, the identity of that product is spread among the users by the epidemic-style "random recommendation" probes. It may be interesting to note that in a "steady state," when many users have halted, the algorithm gives a concrete distributed implementation of best-seller lists by randomly sampling the users.

Theorem 4.1 holds unconditionally, but gives a good bound only for a set of users with a common interest. A good bound for most users can be obtained under an assumption similar to that made in Theorem 3.2, namely, the existence of (λ, k)-type-cover. However, Theorem 4.1 indicates that we can use a much weaker notion, referred to as a *special interest group* or *SIG*. Intuitively, while users of the same type must have exactly the same preferences, users in the same SIG need only have a non-empty intersection of preferences. Formally, we have the following definition.

Definition 4.1 *A special interest group (SIG) S is a set of honest users $U(S)$ and a nonempty set of products $P(S)$ satisfying the condition that every product in $P(S)$ is good for every user in $U(S)$, namely,*

$$U(S) \subseteq \{u \in \mathcal{U} \mid P(S) \subseteq good(u)\}.$$

Given a collection $\mathcal{S}$ of special interest groups we write $U(\mathcal{S}) \overset{\text{def}}{=} \bigcup_{S \in \mathcal{S}} U(S)$.

A special interest group S, therefore, is a set of users $U(S)$ that would be satisfied by any product in $P(S)$ (and possibly by other products as well). The set of products $P(S)$ represents a common or special interest of the users in $U(S)$.

Using the concept of SIGs, we generalize Theorem 4.1 as follows.

Corollary 4.2 *Let $\mathcal{S} = \{S_1, \ldots, S_\ell\}$ be any collection of SIGs, and denote*

$$\tilde{m}(\mathcal{S}) \;=\; \max\{|U(S_1)|, \ldots, |U(S_\ell)|\} \;, \qquad \tilde{n}(\mathcal{S}) \;=\; \min\{|P(S_1)|, \ldots, |P(S_\ell)|\} \;.$$

The expected recommendation complexity for the users in $U(\mathcal{S})$ is at most

$$2\ell \left(\frac{n}{\tilde{n}(\mathcal{S})} + m \ln \tilde{m}(\mathcal{S}) \right) \;\leq\; 2\ell(n + m \ln m).$$

Now we can conveniently compare $\mathcal{A}_{\mathrm{dist}}$ to $\mathcal{A}_{\mathrm{cntr}}$: suppose, in the spirit of Theorem 3.2, that there exists a small collection of SIGs that cover most users (i.e., $\ell = O(1)$ in Corollary 4.2, and $\frac{|U(\mathcal{S})|}{m} \geq 1 - \lambda$ for a small λ). Then the expected total recommendation complexity over most users is just $O(n + m \ln m)$ by Corollary 4.2, as opposed to $O(m + n)$ in Theorem 3.2. However, unlike Theorem 3.2, we need only assume the existence of SIGs rather than types, and the assumption is not required for the algorithm to work correctly.

4.3 Individual recommendation complexity in the synchronous model

Corollary 4.2 gives the expected recommendation complexity for the algorithm $\mathcal{A}_{\mathrm{dist}}$, but what can we say about the individual recommendation complexity? Clearly the number of products a user must probe to find a good one depends how many other users from the same SIG are probing. If the schedule is such that a user is running alone, it will have to carry out the entire search on its own. If a user is running concurrently with many others from its SIG, it can expect the search to be distributed over all of these users, reducing the number of products it has to probe itself. To get a handle on individual costs, we study individual recommendation complexity in a *synchronous model* where an execution proceeds in a sequence of rounds. In each round, users first send messages to other users, then get replies, and finally probe a product. This model is different from the one obtained by restricting the asynchronous model to round robin schedules: round robin schedules consist of a sequence of atomic send-receive-probe steps, whereas the synchronous rounds entail concurrent actions, and therefore require a little more effort to analyze.

Given any special interest group S, we analyze the individual recommendation complexity of users in S as follows. We first make the following trivial observation about the round complexity of $\mathcal{A}_{\mathrm{dist}}$, that follows directly from the "random object" probes.

Lemma 4.3 *Let S be any special interest group. With probability at least $1 - m^{-\Omega(1)}$, all the users in $U(S)$ have found a good product by time $O\left(\log m \cdot \frac{n}{|P(S)|}\right)$.*

Our main result for this model is stated below.

Theorem 4.4 *Let S be any special interest group. If $|U(S)| \geq \Omega(\log m)$, then with probability at least $1 - m^{-\Omega(1)}$, at least half of the users in $U(S)$ have found a good product by time $O\left(\frac{\log m}{|U(S)|} \cdot \left(\frac{n}{|P(S)|} + m\right)\right)$.*

9

The proof is given in the appendix; it is based on showing that the number of satisfied users grows in an exponential rate once some user from $U(S)$ finds a product in $P(S)$. Note that the theorem does not guarantee that all users in the SIG will find a good product quickly: this is because some of the users may be satisfied with products not in $P(S)$. The following lemma says that if sufficiently many users of $U(S)$ actually recommend products from $P(S)$, then *all* users in $U(S)$ will be done quickly.

Lemma 4.5 *Let S be any special interest group, and suppose that at some time t, there are at least $|U(S)|/2$ recommendations for products in $P(S)$. Then with probability at least $1 - m^{-\Omega(1)}$, all users in $U(S)$ will find a good product by time $t + O(\frac{m \log m}{|U(S)|})$.*

Theorem 4.4 and Lemma 4.5 yield the following implication for the special case of types.

Corollary 4.6 *Let T be any type. If $|U(T)| \geq \Omega(\log m)$, then with probability at least $1 - m^{-\Omega(1)}$, all users in $U(T)$ have found a good product by time $O\left(\frac{\log m}{|U(S)|} \cdot \left(\frac{n}{|P(S)|} + m\right)\right)$.*

We note that one can use either Theorem 4.4 or Corollary 4.6 to bound the individual recommendation complexity of users. It is clear that we can always partition the set $\mathcal{U}$ of users into equivalence classes where two users u and v are in the same equivalence class iff $good(u) = good(v)$. Each equivalence class naturally gives rise to a type T where $U(T)$ is the set of users in the class and $P(T)$ is the set of products they consider good. In this sense, we can always find a type T containing u, and use T and Corollary 4.6 to bound the recommendation complexity for u. However, $|U(T)|$ might be small. If there is a large SIG containing u, then Theorem 4.4 allows us to get a tighter bound on the individual recommendation complexity for at least half of the members of that SIG.

5 Discussion of results

In this section we discuss the differences between the models used by the *MR* algorithm in [DKR02] and our centralized model as describes in Section 2. Since our model is a streamlined version of the *MR* model, some (straightforward) analysis is required to show that up to constant factors, our model is no weaker than the *MR* model.

5.1 User preferences

The *MR* model uses real values to model user preferences for products, whereas we use binary values. We argue that using binary values is simpler and does not restrict generality.

Specifically, the *MR* model assumes that there is a *preference vector* A_u for each user u, a vector of nonnegative real values such that the pth entry $A_{u,p}$ in the vector is the *preference value* of product p in the eyes of u. For each user u, the model defines a threshold θ_u in terms of A_u, and defines a product p to be *good* for u if $A_{u,p} > \theta_u$. The model defines a *recommendation* to be a small set of products, and defines a recommendation to be *good* for u if it contains a product that is good for u.

In contrast, we use binary preference values to products, saying that u has preference 1 for p if $p \in good(u)$ and 0 otherwise. This is justified simply because our algorithms work for arbitrary numeric values just as well. (As an aside, it appears that binary data is often seems easier to collect, say by

interpreting a clicking on a product as endorsing it.) It may be interesting to consider translating the *MR* model to the binary model by replacing real preferences with binary preferences by defining *good*(u) to be the products with real preferences above θ_u. This approach gives rise to the more important question: can a user u decide on-line if $p \in good(u)$? The implication of the existence of an on-line test (say, if the threshold value θ_u is known to u), is that users can stop as soon as they find a good product. If no on-line test is available, then the best one can hope for from an algorithm is to prove that if the users test enough of the recommendations made by the algorithm, they will be satisfied. Our distributed algorithm is written assuming that an on-line test is available (thus allowing early stopping), but can easily be adapted to the case where no such test is available.

5.2   Covers and effectiveness

The analysis of the *MR* algorithm uses a concept called (λ, k)-effectiveness. A set of users is said to be (λ, k)-*effective* for some real $0 \leq \lambda \leq 1$ and integer $k \geq 0$ if there exists a subset of at least λm users that belong to at most k types. In the analysis of Algorithm $\mathcal{A}_{\text{cntr}}$, we use a dual concept of (λ, k)-type-covers. The following lemma shows that the concepts of user effectiveness and covers are equivalent up to constant factors.

Lemma 5.1 *Let $\mathcal{U}$ be a set of users and good(u) be the good products for a each user $u \in \mathcal{U}$.*

 (1) *If there exists a (λ, k)-type-cover for $\mathcal{U}$, then $\mathcal{U}$ is (λ', k')-effective for $\lambda' = \lambda$ and $k' = \lceil \lambda k \rceil$.*

 (2) *If $\mathcal{U}$ is (λ, k)-effective, then for any $0 < \epsilon \leq 1$, $\mathcal{U}$ has a (λ', k')-type-cover for $\lambda' = (1 - \epsilon)\lambda$ and $k' = k/(\lambda\epsilon)$.*

Casting the results of Theorems 3.1 and 3.2 in the terminology of [DKR02] and applying Lemma 5.1, we get the following:

Corollary 5.2 *If the users are (λ', k')-effective for some $\lambda' > 0$ and $k' > 0$, then for all $\epsilon > 0$ Theorems 3.1 and 3.2 hold with $\lambda = (1 - \epsilon)\lambda'$ and $k = k'/(\lambda'\epsilon)$.*

5.3   Type deviation

An important aspect of the *MR* model that we abstract is type deviation. Specifically, in the *MR* model, it is assumed that the preference vector of each user is a random variable, obtained by adding to the user's type vector an "error vector." The error vector is assumed to consist of n independent random variables with mean $\mu = 0$ and variance $\sigma^2 = \frac{\epsilon^2}{m+n}$ for some small $0 < \epsilon < 1$.

Assuming such a small magnitude of variance of the errors, it is easy to justify our abstraction: This assumption implies that for each user and any $p > 0$, with probability $1 - p$, all preferences are at most $\delta \stackrel{\text{def}}{=} \epsilon/\sqrt{p}$ away from the type vector. This is because Chebychev's Inequality says that the probability that a user's preference for any specific product differs from the canonical preference by more than δ is less than

$$\frac{\sigma^2}{\delta^2} = \frac{\epsilon^2}{(m+n)\delta^2} \leq \frac{\epsilon^2}{n\delta^2} = \frac{\epsilon^2 p}{n\epsilon^2} = \frac{p}{n}.$$

It follows that the probability that *all* the preferences of a user differ from the type vector by *at most δ* is at least $\left(1 - \frac{p}{n}\right)^n \geq 1 - p$. Hence, for any given $p > 0$ we can define a goodness predicate using a

threshold θ so that all but a fraction of p of the users have preferences that agree with all other users of their type.

As an aside, we note that SIGs may be viewed as another (deterministic) kind of type deviation.

Acknowledgment

References

[AFK+01] Yossi Azar, Amos Fiat, Anna Karlin, Frank McSherry, and Jared Saia. Spectral analysis of data. In *Proc. 33rd ACM Symp. on Theory of Computing (STOC)*, pages 619–626, 2001.

[APPT04] Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Mark Tuttle. Collaboration of untrusting peers with changing interests. In *Proc. 5th ACM Conf. on Electronic Commerce (EC)*, pages 112–119, May 2004.

[Can02] John F. Canny. Collaborative filtering with privacy. In *IEEE Symposium on Security and Privacy*, pages 45–57, 2002.

[DKR02] Petros Drineas, Iordanis Kerenidis, and Prabhakar Raghavan. Competitive recommendation systems. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 82–90, 2002.

[KRRT98] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. Recommendation systems: A probabilistic analysis. In *Proc. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 664–673, 1998.

[KS03] Jon Kleinberg and Mark Sandler. Convergent algorithms for collaborative filtering. In *Proc. 4th ACM Conf. on Electronic Commerce (EC)*, pages 1–10, 2003.

[LR04] Shyong K. Lam and John Riedl. Shilling recommender systems for fun and profit. In *13th conference on World Wide Web*, pages 393–402. ACM Press, 2004.

[OHS04] Michael P. OMahony, Neil J Hurley, and Guénolé C. M. Silvestre. Utility-based neighbourhood formation for efficient and robust collaborative filtering. In *5th ACM conference on Electronic commerce*, pages 260–261, 2004.

[Rab83] Michael O. Rabin. Transaction protection by beacons. *Journal of Computer and System Sciences*, 27(2):256–267, 1983.

[RIS+94] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. Grouplens: an open architecture for collaborative filtering of netnews. In *1994 ACM conference on Computer supported cooperative work*, pages 175–186. ACM Press, 1994.

[SKKR00] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Analysis of recommendation algorithms for e-commerce. In *2nd ACM conference on Electronic commerce*, pages 158–167. ACM Press, 2000.

[SKKR01] Badrul Sarwar, George Karypis, Joseph Konstan, and John Reidl. Item-based collaborative filtering recommendation algorithms. In *10th international conference on World Wide Web*, pages 285–295. ACM Press, 2001.

[SPUP02]   Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. Methods and metrics for cold-start recommendations. In *25th annual international ACM SIGIR conference on Research and development in information retrieval (SIGIR '02)*, pages 253–260, 2002.

Appendix: Proofs

Proof of Theorem 3.1: Let $\mathcal{T}$ be a (λ, k)-type-cover. For a type $T \in \mathcal{T}$, the users of $U(T)$ will find a good product in the set recommended by Algorithm $\mathcal{A}_{\mathrm{cntr}}$ if some user of $U(T)$ has been chosen to the committee. Each user of $U(T)$ is chosen to the committee with probability $|U(T)|/m \geq 1/k$, since $T \in \mathcal{T}$ and $\mathcal{T}$ is a (λ, k)-cover. Hence the probability that no user of $U(T)$ has been chosen to the committee is at most

$$\left(1 - \frac{1}{k}\right)^K = \left(1 - \frac{1}{k}\right)^{k \ln(k/\gamma)} \leq e^{-\ln k/\gamma} = \gamma/k \ .$$

Consequently, the probability that the users of *some* type $T \in \mathcal{T}$ will fail to find a good product in the set recommended by Algorithm $\mathcal{A}_{\mathrm{cntr}}$ is at most $k(\gamma/k) = \gamma$. The theorem now follows from the fact that $|U(\mathcal{T})| \geq \lambda m$. ∎

Proof of Theorem 3.2: A random user u is chosen for the committee with probability $\frac{K}{m} = O(\frac{k \log(k/\gamma)}{m})$, in which case it tests all n products, accounting for the second term in the bound. Now suppose that $u \in U(\mathcal{T})$ is a non-committee member. In this case u only tests products until it finds a product it likes, or until all K recommendations are exhausted. Consider the ith recommendation it tests. Since committee members are picked at random, and since u picked a random recommendation as its ith product, we have that u follows the recommendation of a random user. Let T_u be the type of u. Since $T_u \in \mathcal{T}$, the probability that a random user is of type T_u is $|U(T_u)|/m \geq 1/k$. Since the ith recommendation is good for u if the ith committee member is of type T_u, we get that the probability that the first good recommendation is the ith one is at least $(1 - 1/k)^{i-1} \cdot (1/k)$. The probability that u does not find a good recommendation after trying all K of them is $(1 - 1/k)^K$. Therefore, the expected number φ_u of products u tests is bounded by

$$\mathbf{E}\left[\varphi_u\right] \leq \sum_{i=1}^{K} i \frac{1}{k}\left(1 - \frac{1}{k}\right)^{i-1} + K\left(1 - \frac{1}{k}\right)^K \leq k + K\left(1 - \frac{1}{k}\right)^{k \ln(k/\gamma)}$$

$$\leq k + K e^{-\ln(k/\gamma)} = k + \frac{k \ln(k/\gamma)}{k/\gamma} = O(k) \ . \qquad ∎$$

Proof of Theorem 4.1: To prove that the bound on expected recommendation complexity holds for every operating environment, let us fix an arbitrary environment $\mathcal{E}$. Let σ be the random variable whose value is the sequence of probes by users in U during an execution of $\mathcal{A}_{\mathrm{dist}}$ in the context of $\mathcal{E}$. Let $\sigma = \sigma_0 \sigma_1 \cdots \sigma_{|U|}$ where σ_ℓ is the random variable whose value is the subsequence of probes preceded by *exactly* ℓ probes by users in U of products in P. In other words, exactly ℓ users in U have found a product in P at the start of σ_ℓ, and the final probe of σ_ℓ is by the $(\ell + 1)$st user in U finding a product in P. Notice that some users in U may be satisfied by products outside of P, meaning that only a subset of users in U actually probe a product in P, in which case some suffix $\sigma_j \sigma_{j+1} \cdots \sigma_{|U|}$ of σ may actually be empty. Since users halt once they find a product in P, we can bound the recommendation complexity by counting the number of probes by users in U to products not in P. Clearly $count(\sigma) = \sum_{\ell=0}^{|U|} count(\sigma_\ell)$ is an upper bound on the total number of products tested by users in U as they search for a good product.

Consider the cost of σ_0. With probability $1/2$ the algorithm chooses a random product to sample. Since P is nonempty and there are n products, each probe in σ_0 probes a product in P with probability at least $|P|/2n$, and hence $\mathbf{E}\left[count(\sigma_0)\right] \leq 2n/|P|$. Now consider the cost of σ_ℓ for $\ell > 0$. With probability $1/2$ the algorithm chooses a user and samples its recommendation. Since exactly ℓ users in U have found a product in P when σ_ℓ begins, each probe in σ_ℓ finds a product in P with probability at least $\ell/2m$, so the expected number of probes in σ_ℓ is $2m/\ell$. Thus,

$$
\mathbf{E}\left[count(\sigma)\right] = \mathbf{E}\left[count(\sigma_0)\right] + \sum_{\ell=1}^{|U|} \mathbf{E}\left[count(\sigma_\ell)\right] \leq \frac{2n}{|P|} + \sum_{\ell=1}^{|U|} \frac{2m}{\ell} = 2\left(\frac{n}{|P|} + m\ln|U|\right). \quad \blacksquare
$$

Proof of Lemma 4.3: Fix a constant $c \geq 2$. Consider some user in $U(S)$. While the user is unsatisfied, it probes in every round a product in $P(S)$ with probability at least $\frac{1}{2}\frac{|P(S)|}{n} = \frac{\beta_S}{2}$. Thus the probability that the user fails to probe a product of $P(S)$ in the first $r_0 = \lceil 2(c+1)\log m/\beta_S \rceil$ rounds is at most

$$
\left(1 - \frac{\beta_S}{2}\right)^{r_0} \leq \left(1 - \frac{\beta_S}{2}\right)^{2(c+1)\log m/\beta_S} \leq \exp(-(c+1)\log m) = m^{-c-1} .
$$

It follows that the probability that *some* user of $U(S)$ fails to probe a product in $P(S)$ in any of the first r_0 rounds is at most m^{-c}, implying the theorem. $\quad \blacksquare$

Proof of Theorem 4.4: Fix a constant $c \geq 2$, and assume that $\alpha_S > 32c\log m/m$. Denote by $\rho(t)$ the number of recommendations by users in $U(S)$ for products in $P(S)$ after t rounds. Any execution of the algorithm can be viewed as having $q = \lceil \log \frac{\alpha_S m}{4c\log m} \rceil$ phases, defined as follows. Let $f_1 = 1$ and $f_k = 2^k c\log m$ for $2 \leq k \leq q - 1$.

1. For $1 \leq k \leq q - 1$, Phase i ends when either half the users in $U(S)$ are satisfied or $\rho(t) \geq f_k$, whichever is earlier.

2. Phase q ends when at least half the users in $U(S)$ are satisfied.

(Note that the two events defining the end of a phase are not necessarily ordered, since generally, users in $U(S)$ can be satisfied by products other than those in $P(S)$ and similarly, products in $P(S)$ can be recommended by users other than those in $U(S)$.)

For $1 \leq k \leq q$, let t_k be a random variable counting the number of rounds of Phase k. Let $r_1 = \lceil 4c\frac{\ln m}{\alpha_S \beta_S m} \rceil$, $r_2 = \lceil \frac{64c\log m}{\alpha_S} \rceil$, $r_k = \lceil \frac{16}{\alpha_S} \rceil$ for $3 \leq k \leq q - 1$ and $r_q = \lceil \frac{16(c+1)\log m}{\alpha_S} \rceil$. Let $\mathcal{E}_k$ denote the event that $t_k > r_k$.

Lemma .3 $\mathbf{P}\left[\mathcal{E}_k\right] \leq m^{-c}$ *for every* $1 \leq k \leq q$.

Proof: We prove the lemma for each phase separately.

Phase 1: In every round of Phase 1, each unsatisfied user probes a product in $P(S)$ with probability at least $\frac{1}{2}\frac{|P(S)|}{n} = \frac{\beta_S}{2}$. By definition of Phase 1, at least $\alpha_S m/2$ users in $U(S)$ are unsatisfied throughout the phase. Hence $\mathbf{P}\left[\mathcal{E}_1\right]$ is bounded from above by the probability that none of these users

probes a product in $P(S)$ in any of the first r_1 rounds, i.e.,

$$\mathbf{P}\left[\mathcal{E}_1\right] \;\leq\; \left(1-\frac{\beta_S}{2}\right)^{\alpha_S m r_1/2} \;\leq\; \exp(-\alpha_S \beta_S m r_1/4) \;=\; \exp\left(-\frac{\alpha_S \beta_S m}{4}\cdot\frac{4c\ln m}{\alpha_S \beta_S m}\right) \;=\; m^{-c}\,,$$

implying the lemma for Phase 1.

Phase 2: Clearly, if Phase 1 ends when half the users in $U(S)$ are satisfied then the ends of Phase 1 and 2 coincide and there is nothing to prove. So the interesting case is when Phase 2 begins with at least one recommendation for a product in $P(S)$, but fewer than half the users in $U(S)$ are satisfied. In this case, each unsatisfied user u samples a product of $P(S)$ with probability at least $\frac{1}{2m}$ in every round of the phase. Hence u finds a product of $P(S)$ in the first r_k rounds of Phase 2 with probability at least $1-p$ for $p = (1-\frac{1}{2m})^{r_2}$. Note that

$$p \;\leq\; \left(1-\frac{1}{2m}\right)^{\frac{64c\log m}{\alpha_S}} \;\leq\; \exp\left(-\frac{32c\log m}{\alpha_S m}\right) \;\leq\; 1-\frac{16c\log m}{\alpha_S m}\,,$$

where the last inequality relies on the fact that $\exp(-x)\leq 1-x/2$ for $0\leq x\leq 1$ and $\alpha_S > 32c\log m/m$. By definition of the phase, there are at least $\alpha_S m/2$ unsatisfied users in $U(S)$ throughout the phase. Let H_k denote the random variable counting the number of users of $U(S)$ that find a product of $P(S)$ during the first r_k rounds of Phase k. It follows that the expected value of H_2 satisfies

$$\mathbf{E}\left[H_2\right] \;\geq\; \frac{\alpha_S m}{2}(1-p) \;\geq\; \frac{\alpha_S m}{2}\cdot\frac{16c\log m}{\alpha_S m} \;=\; 8c\log m\,. \tag{1}$$

By Chernoff's bound,

$$\mathbf{P}\left[\mathcal{E}_2\right] \;\leq\; \mathbf{P}\left[H_2 < 4c\log m\right] \;\leq\; \exp\left(-\frac{1}{2}\cdot\frac{1}{4}\cdot 8c\log m\right) \;=\; m^{-c}\,,$$

implying the lemma for Phase 2.

Phase k for $3\leq k\leq q-1$: Again, if Phase $k-1$ ends when half the users in $U(S)$ are satisfied then there is nothing to prove, so suppose Phase k begins with at least f_{k-1} recommendation for a product in $P(S)$, but fewer than half the users in $U(S)$ are satisfied.

Assuming Phase k begins with at least f_{k-1} recommendations for products in $P(S)$, each unsatisfied user u samples a product of $P(S)$ with probability at least $\frac{f_{k-1}}{2m}$ in every round, hence u finds a product of $P(S)$ in the first r_k rounds of Phase k with probability at least $1-p$ for $p = \left(1-\frac{f_{k-1}}{2m}\right)^{r_k}$. Note that

$$p \;\leq\; \left(1-\frac{f_{k-1}}{2m}\right)^{16/\alpha_S} \;\leq\; \exp\left(-\frac{8f_{k-1}}{\alpha_S m}\right) \;\leq\; 1-\frac{4f_{k-1}}{\alpha_S m}\,,$$

where the last inequality relies on the fact that $\exp(-x)\leq 1-x/2$ for $0\leq x\leq 1$ and $f_{k-1}\leq \alpha_S m/8$ for $k\leq q-1$. By definition of the phase, there are at least $\alpha_S m/2$ unsatisfied users throughout it. It follows that the expected value of H_k satisfies

$$\mathbf{E}\left[H_k\right] \;\geq\; \frac{\alpha_S m}{2}(1-p) \;\geq\; \frac{\alpha_S m}{2}\cdot\frac{4f_{k-1}}{\alpha_S m} \;=\; 2f_{k-1}\,.$$

16

By the Chernoff bound,

$$\mathbf{P}\left[H_k < f_{k-1}\right] \;\leq\; \exp\left(-\frac{1}{2}\cdot\frac{1}{4}\cdot 2f_{k-1}\right) \;\leq\; \exp\left(-\frac{1}{8}\cdot 8c\log m\right) \;=\; m^{-c}\;.$$

As Phase k begins with at least f_{k-1} recommendations for products in $P(S)$, we have that if $H_k \geq f_{k-1}$ then after r_k rounds of Phase k, the number of such recommendations is at least $f_{k-1}+H_k \geq 2f_{k-1} = f_k$, implying $\neg\mathcal{E}_k$. Hence $\mathbf{P}\left[\mathcal{E}_k\right] \;\leq\; m^{-c}$.

Phase q: Again, if Phase $q-1$ ends when half the users in $U(S)$ are satisfied then we are done, so consider the case when Phase q begins with at least f_{q-1} recommendation for a product in $P(S)$, but fewer than half the users in $U(S)$ are satisfied.

Assuming Phase q begins with at least

$$f_{q-1} \;=\; 2^{q-1}c\log m \;\geq\; c\log m\cdot\frac{\alpha_S m}{8c\log m} \;=\; \frac{\alpha_S m}{8}$$

recommendations for products in $P(S)$, each unsatisfied user u samples a product of $P(S)$ with probability at least $\frac{f_{q-1}}{8m}$ in every round, hence u fails to find a product of $P(S)$ in the first r_q rounds of Phase q with probability at most

$$\left(1-\frac{f_{q-1}}{8m}\right)^{r_q} \;\leq\; \exp\left(-\frac{f_{q-1}\cdot r_q}{8m}\right) \;\leq\; \exp\left(-\frac{\frac{\alpha_S m}{8}\cdot\frac{16(c+1)\log m}{\alpha_S}}{2m}\right) \;=\; \exp(-(c+1)\log m) = m^{-c-1}\;.$$

Thus the probability that not all users of $U(S)$ will be satisfied after r_q rounds is at most m^{-c}. ∎

The theorem now follows by combining the claims of the lemma for the separate phases and concluding that the probability that any of the events $\mathcal{E}_i$ has occurred, for $1\leq i\leq q+1$, is at most $(q+1)m^{-c}\leq m^{1-c}$. Hence with probability at least $1-m^{1-c}$, the total number of rounds until at least half the users of $U(S)$ were satisfied was bounded by

$$\sum_{k=1}^{q} r_k \;\leq\; \left\lceil 4c\frac{\ln m}{\alpha_S\beta_S m}\right\rceil + \left\lceil\frac{64c\log m}{\alpha_S}\right\rceil + (q-3)\cdot\left(\frac{16}{\alpha_S}+1\right) + \left\lceil\frac{16(c+1)\log m}{\alpha_S}\right\rceil$$

$$=\; O\left(\log m\left(\frac{1}{\alpha_S\beta_S m}+\frac{1}{\alpha_S}\right)\right)\;,$$

as $q=O(\log m)$. ∎

Proof of Lemma 4.5: Let $c>4$. We show that with probability at least $1-m^{1-c/4}$, all users in $U(S)$ will find a good product by time $t+\frac{c\ln m}{\alpha_S}$. Consider any user $u\in U(S)$ that has not found a good product by time t. Let p_u denote the probability that u does not find a good product within $\frac{c\ln m}{\alpha_S}$ rounds after time t. In each round after time t, u follows the recommendation of a random user with probability $1/2$. By assumption, this user recommends a product from $P(S)$ with probability at least $\frac{\alpha_S m/2}{m}=\alpha_S/2$. Hence in each round after time t, u will try a product from $P(S)$ with probability at least $\frac{1}{2}\cdot\frac{\alpha_S}{2}=\frac{\alpha_S}{4}$. It follows that

$$p_u \leq \left(1-\frac{\alpha_S}{4}\right)^{(c\ln m)/\alpha_S} \leq e^{-(c\ln m)/4} = m^{-c/4}\;.$$

17

The result follows, since the probability that not all users in $U(S)$ have finished by time $t + \frac{c \ln m}{\alpha_S}$ is at most $p_u |U(S)| \leq p_u m$. $\blacksquare$

Proof of Corollary 4.6: By Theorem 4.4, with probability at least $1 - m^{-\Omega(1)}$ there exists some time $t_1 = O(\log m(\frac{1}{\alpha_T \beta_T m} + \frac{1}{\alpha_T}))$ by which at least half of the users of $U(T)$ have found a good product. Since T is a type, the good products all these users found are in $P(T)$. The result hence follows from Lemma 4.5. $\blacksquare$

Proof of Lemma 5.1: Suppose $\mathcal{U}$ has a (λ, k)-type-cover $\mathcal{T}'$. To show that $\mathcal{U}$ is $(\lambda, \lceil \lambda k \rceil)$-effective, consider any subset $\mathcal{T}$ of $\mathcal{T}'$ such that $|\mathcal{T}| = \lceil \lambda k \rceil$. It is sufficient to show that $|U(\mathcal{T})| \geq \lambda m$. This is true because by definition, $|U(T)| \geq \frac{m}{k}$ for each $T \in \mathcal{T}'$, and $U(T) \cap U(T') = \emptyset$ for any two distinct $T, T' \in \mathcal{T}'$. Therefore $|U(\mathcal{T})| \geq \lceil \lambda k \rceil \frac{m}{k} \geq \lambda m$.

Conversely, suppose $\mathcal{U}$ is (λ, k)-effective, and fix $0 \leq \epsilon < 1$. Let $\mathcal{T} = \{T_1, \ldots, T_k\}$ be the type collection indicated by the (λ, k)-effectiveness assumption. Then $|U(\mathcal{T})| \geq \lambda m$. Define $\mathcal{T}_\epsilon = \{T \in \mathcal{T} \mid |U(T)| \geq \epsilon \lambda m / k\}$. We claim that $\mathcal{T}_\epsilon$ is a $((1 - \epsilon)\lambda, k/(\lambda \epsilon))$-type-cover. By definition, each type in $\mathcal{T}_\epsilon$ has at least $\lambda \epsilon m / k$ users, hence the second cover requirement follows. To see that also $|U(\mathcal{T}_\epsilon)|$ is as required, define $\mathcal{T}'_\epsilon = \mathcal{T} \setminus \mathcal{T}_\epsilon$. By definition, $|U(T)| < \epsilon \lambda m / k$ for each $T \in \mathcal{T}'_\epsilon$. This yields the required size bound, since

$$|U(\mathcal{T}_\epsilon)| \;=\; |U(\mathcal{T} \setminus \mathcal{T}'_\epsilon)| \;\geq\; \lambda m - k \cdot \frac{\epsilon \lambda m}{k} \;=\; (1 - \epsilon)\lambda m . \qquad \blacksquare$$