Fast Distributed Bandits for Online Recommendation Systems

Kanak Mahadik
Adobe Research
San Jose, USA
mahadik@adobe.com

Shuai Li
Cambridge University
Cambridge, United Kingdom
shuaili.sli@gmail.com

Qingyun Wu
University of Virginia
Charlottesville, USA
qw2ky@virginia.edu

Amit Sabne
Google Brain
Mountain View, USA
asabne@google.com

ABSTRACT

Contextual bandit algorithms are commonly used in recommender systems, where content popularity can change rapidly. These algorithms continuously learn latent mappings between users and items, based on contexts associated with them both. Recent recommendation algorithms that learn clustering or social structures between users have exhibited higher recommendation accuracy. However, as the number of users and items in the environment increases, the time required to generate recommendations deteriorates significantly. As a result, these cannot be deployed in practice. The state-of-the-art distributed bandit algorithm - DCCB - relies on a peer-to-peer network to share information among distributed workers. However, this approach does not scale well with the increasing number of users. Furthermore, it suffers from slow discovery of clusters, resulting in accuracy degradation.

To address the above issues, this paper proposes a novel distributed bandit-based algorithm called *DistCLUB*. This algorithm lazily creates clusters in a distributed manner, and dramatically reduces the network data sharing requirement, achieving high scalability. Additionally, *DistCLUB* finds clusters much faster, achieving better accuracy than the state-of-the-art algorithm. Evaluation over both real-world benchmarks and synthetic datasets shows that *DistCLUB* is on average 8.87x faster than DCCB, and achieves 14.5% higher normalized prediction performance.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent systems; Machine learning; Online learning settings; Distributed algorithms; Parallel algorithms.**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICS '20, June 29–July 2, 2020, Barcelona, Spain

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7983-0/20/06...\$15.00

<https://doi.org/10.1145/3392717.3392748>

KEYWORDS

bandit systems, machine learning, online learning, high-performance computing, recommendation systems, distributed systems

1 INTRODUCTION

Large-scale recommendation systems are deployed in businesses such as retail, video-on-demand, and music streaming. They are fundamental drivers for web services such as YouTube, Netflix, Amazon, and Google News. Recommendation systems provide the user with suggestions on what they should watch next, shop next, or do based on the time of the day. According to McKinsey, the recommendation feature contributed to 35% and 23.7% growth in revenue for Amazon and BestBuy respectively. The report also stated that 75% of the video-consumption and 60% views on web services Netflix and YouTube [28] respectively are due to recommendations. Hence, developing effective recommendation systems is extremely important for businesses such that appropriate content is suggested to users in a timely manner.

Recommendation systems are deployed in social environments where user preferences and the set of active users change with time. In addition, many web services comprise content that is not only dynamic, but also receives rapid updates. Traditional recommendation systems involve supervised learning algorithms that use collaborative and content-based filtering methods [13, 17]. The key issue faced by these supervised learning approaches is that they train *offline*. Real applications, however, face a phenomenon called *concept drift* [12, 34, 37], wherein user preferences and content could change within minutes. Concept drift makes results obtained by supervised learning sub-optimal. The large-scale at which these web services operate make it necessary to have algorithms that learn fast, in an *online* fashion. Bandit-based algorithms can adapt to these updates and therefore are prevalent in online settings [14, 24]. A *bandit learner* or *agent* picks an action from a set, i.e., an item to recommend to the user, and observes the user feedback to learn about her preferences for the set of items. If the feedback is positive, e.g. the user clicks on the item, the agent receives a reward. The agent is inclined to pick actions that provide higher rewards, thus reinforcing better user-item mappings.

As the number of users and items increases, these bandit-based algorithms face severe challenges in scaling up to satisfy the requirements of these modern applications that cater to millions of

users with billions of possible recommendable items. High-quality, speedy, and dynamic recommendations are critical to the success of these enterprises. At scale, the data sizes and compute volumes go well beyond the capacity of a single compute node. Therefore, the ability to compute recommendations in a distributed manner becomes paramount [18].

Unfortunately, most of the bandit-based recommendation algorithms are neither parallel nor distributed [14, 24, 27]. Bandit-based algorithms access and edit dynamic user and item information, necessitating data coherence across different cores/nodes. The resultant high synchronization costs limit the scalability of a distributed approach. The state-of-the-art distributed online recommendation algorithm, called Distributed Clustering of Confidence Ball (DCCB) [19] builds upon the methodology used in its sequential counterpart, named Cluster of Bandits (CLUB) [14]. CLUB operates by organizing users into different clusters and recommending items to users based on the cluster's cumulative properties. The process of cluster formation is dynamic, and therefore adaptable to changing user preferences. If a distributed algorithm, such as DCCB, had to dynamically create clusters and perform recommendations, it would need all-to-all communications after each interaction. DCCB overcomes this expensive need by a) updating clusters in a lazy, periodic fashion and b) it archives the interactions-generated user feedback for the period in a buffer. After every period, each node shares its buffer with a randomly chosen peer so as to propagate the newly learned user-item mappings. Bandit-based algorithms are resilient to such lazy updates [19].

In spite of the lazy distributed updates, DCCB runs into scalability issues owing to its reliance on peer-to-peer communications transferring large data buffers. To reduce communication penalties, DCCB transfers the data periodically. However, this results in slow discovery of clusters and a reduction in accuracy since the restricted peer-to-peer information exchange introduces bias and imperfections due to partial information visibility. The communication bottleneck and imperfect information sharing exacerbate in environments with large user counts. Because each machine must handle several users, the buffer space gets constrained, requiring more frequent communication.

In this work, we present *DistCLUB*¹, a novel distributed algorithm that overcomes these issues. *DistCLUB* eliminates the need for large buffer transfers by distributing execution across nodes based on user-level and cluster-level parallelism. It employs efficient reduction primitives, boosting the cluster discovery rate. *DistCLUB* adaptively alternates between recommending based on user and cluster information. It leverages heuristics to statistically estimate the volume of interaction data available for each user, and *intelligently* balances highly personalized and cluster-based recommendations. The reduced communication overheads allow *DistCLUB* to quickly discover clusters and achieve higher accuracy than DCCB. The paper presents a detailed regret analysis for *DistCLUB* which shows that it achieves the same asymptotic regret bound as DCCB and CLUB.

To the best of our knowledge, *DistCLUB* is the first large-scale, high-performance distributed system for online recommendations that can scale to 64 nodes (512 cores). *DistCLUB* is built using Spark [41], and carefully maps operations on Spark's distributed

primitives. We evaluate *DistCLUB* on several real-world applications, achieving an average execution speedup of 8.87x while obtaining 14.5% better normalized prediction performance than DCCB.

This paper makes the following contributions:

- (1) The paper presents a novel distributed algorithm named *DistCLUB* that provides online recommendations at scale.
- (2) The paper designs and implements a Spark-based distributed system to demonstrate the effectiveness of *DistCLUB*.
- (3) The paper experimentally demonstrates that *DistCLUB* can scale to millions of user interactions, while operating faster and achieving better accuracy than the state-of-the-art DCCB approach.

The rest of this paper is organized as follows. Section 2 describes the state-of-the-art online recommendation algorithm, CLUB and DCCB. Section 3 present the *DistCLUB* algorithm and its system design. Section 4 describes the theoretical regret analysis of *DistCLUB*. Section 5 presents experimental results. Section 6 overviews the related work, while section 7 concludes the paper.

2 BANDIT-BASED ONLINE RECOMMENDATION ALGORITHMS

Figure 1 shows a typical bandit-based recommendation system. For each interaction t , the bandit agent picks an action, i.e. an item I_k to recommend to the user. Context C_t represents user and items features. For example, in a movie recommendation system, user features could be geographical information, historical behaviour while item features could be genre, rating, etc. The action of selecting I_k is based on context C_t for this interaction, and feedback obtained in the previous interactions [3]. The problem of sequentially selecting an item among multiple items with each having an unknown reward distribution is an instance of *multi-armed bandit* problem². The user responds to this recommendation by either clicking on it or not, which is considered as the reward R_t to the agent. If the user clicks the item the agent receives a non-zero reward for the item. Based on the user feedback, the agent learns about user preferences for the set of items. Note that the agent only observes the reward for the recommended item.

This section first describes the sequential Cluster of Bandits (CLUB) [14], followed by the state-of-the-art distributed online recommendation algorithm named Distributed Clustering Confidence Ball (DCCB) [19].

2.1 Cluster of Bandits (CLUB)

Listing 1 shows the CLUB [14] algorithm. It sequentially clusters users via linear contextual bandits.

CLUB begins with a strongly connected graph. Through the incoming user interactions, it partitions the users into m clusters such that $m \ll n$.

Let $G(V, E)$ be the undirected graph consisting of vertices V and edges $E \subseteq V \times V$. Each vertex $\{u_1, u_2, \dots, u_n\} \in V$ represents a user, and an edge between two vertices indicates that the users have similar preferences.

¹Our software is available at <https://bit.ly/distclub>

²The term comes from a casino setting wherein a player is facing multiple slot machines (*arms*) at once, and needs to repeatedly choose where to insert the next coin. "Bandit" is an American slang for slot machine.

Figure 1: Bandit-based Recommendation System

Features of items in the recommendation system are encoded using a d -dimensional vector. The bandit agent maintains a d -dimensional user-vector for each user which encapsulates the user's preference for these features of the items. In other words, these item vectors, denoted as $\{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_n\}$, are the agent's estimation of the user's preferences. In the linear contextual bandit setting [1], $\vec{v}_i$ can be obtained by multiplying the inverse of correlation matrix M_i and the bias vector $\vec{b}_i$ [14]. Cluster vectors are calculated by aggregating user-vectors in the same cluster. For m clusters, we denote the cluster vectors as $\{\vec{w}_1, \vec{w}_2, \dots, \vec{w}_m\}$.

A user interaction entails a bandit learner recommending an item to the user, and observing their reaction to it, obtained as a *reward*. During each interaction, the learner recommends an item from an item-list, i.e., context, as illustrated in Figure 1.

Upper Confidence Bound (UCB) bandit algorithm [6, 7, 21] estimates an upper confidence bound for each arm as the sum of the sample mean reward and the confidence interval (exploration bonus). It then picks the arm with the highest calculated upper confidence bound on the reward for a user on each interaction. Note that UCB associates higher confidence to users that were previously interacted with, i.e. confidence intervals for an arm decrease after its true reward is observed. Importantly, estimated cluster vectors are used to make recommendations and not user-vectors. After observing the obtained reward, CLUB updates the user-vector according to the contextual choice using a standard linear-bandit algorithm [10].

Clusters are recomputed according to the updated user vectors. The key principle exploited by the algorithm for clustering is that similar users have similar ground truth vectors. If the difference between two vectors is greater than a threshold, the corresponding users are considered dissimilar and are disconnected from each other. In practice, recomputing clusters after each interaction becomes expensive, and hence is typically done once in a few interactions.

2.2 Distributed Clustering Confidence Ball (DCCB)

Listing 2 shows the DCCB [19] algorithm. DCCB is a decentralised, distributed variant of the CLUB algorithm.

DCCB allows catering to different user interactions in parallel. Initially, similar to CLUB, all users are connected to each other. For each user, DCCB maintains a fixed-sized buffer of interactions-ordered correlation matrices, and b -vectors. All elements in the

Algorithm 1: Cluster of Bandit (CLUB) - UCB-based clustering; a sequential approach to online recommendations

Input: α : UCB hyperparameter
Input: δ : Network update delay
Input: γ : cluster-partitioning hyperparameter
Input: $G(V, E)$: Graph where vertices represent users. Edges represent correlation between users. Initially it is a complete graph.
Input: occ : Array of counters representing interaction occurrences for the users, initialized to 0
Input: context: table with features for all recommendable items
Initial state: $Mu[users] = I \quad \forall users \in [0, numUsers]$
 $\vec{b}u[users] = 0 \quad \forall users \in [0, numUsers]$

```
function UCB( $\vec{w}, occ, context, Minv$ )
for each item  $k$  in context do
    estimate  $\leftarrow k \times \vec{w}$ 
    bonus  $\leftarrow \alpha \times \sqrt{k^T \cdot Minv \cdot k} \times \sqrt{\log(1 + occ)}$ 
    UCB-VAL[ $k$ ] = estimate + bonus
end
return argmax(UCB-VAL)
end function
```

```
function UPDATENETWORK()
for each user pair  $(u_1, u_2)$  do
    if exists  $E[u_1, u_2]$  then
         $\vec{v}_1 \leftarrow Mu[u_1]^{-1} \times \vec{b}u[u_1]$ 
         $\vec{v}_2 \leftarrow Mu[u_2]^{-1} \times \vec{b}u[u_2]$ 
        if  $|\vec{v}_1 - \vec{v}_2| \geq \gamma \times getThreshold(\vec{v}_1, \vec{v}_2)$  then
            remove  $E[u_1, u_2]$ 
        end
    end
end
end function
```

```
for each user interaction  $i$  do
    userID  $\leftarrow$  getUserIdFromInteraction( $i$ )
    clusterID  $\leftarrow$  clusters[userID]
     $M_c \leftarrow I, \vec{b}_c \leftarrow 0$ 
    for each user  $n \in clusterID$  do
         $M_c \leftarrow M_c + Mu[n]$ 
         $\vec{b}_c \leftarrow \vec{b}_c + \vec{b}u[n]$ 
    end
     $\vec{w} \leftarrow M_c^{-1} \times \vec{b}_c$ 
    choice  $\leftarrow$  UCB( $\vec{w}, occ[userID], getContext(i), M_c^{-1}$ )
    reward  $\leftarrow$  getReward(choice)
    // standard linear bandit algorithm
     $Mu[userID] \leftarrow context[choice] \cdot context[choice]^T$ 
     $\vec{b}u[userID] \leftarrow reward \times context[choice]$ 
    if  $i \% \delta == 0$  then
        UPDATENETWORK()
        clusters  $\leftarrow$  recomputeClusters( $G$ )
    end
     $occ[userID] \leftarrow occ[userID] + 1$ 
end
```

buffers are initialized to zero. Furthermore, each user has an active and a current correlation matrix and a bias-vector associated with it. For each user i , current objects are used to compute the user-vector,

Algorithm 2: Distributed Clustering Confidence Ball (DCCB) : performs clustering in a distributed manner.

Input: $G(V, E)$: User Graph, All users in the same cluster initially.
Input: MBuff, bBuff - Buffers of length bufferSize, Mw, $\vec{bw}$ - current copies, Mw_{local} , $\vec{bw}_{local}$ - local working copies
Initial state: $Mw[users] = I \quad \forall users \in [0, numUsers]$
 $\vec{bw}[users] = 0 \quad \forall users \in [0, numUsers]$
while user interactions exist **do**
 while no user buffer gets bufferSize new elements **do**
 // loop iterations can execute in parallel
 $j \leftarrow \text{getUserIdFromInteraction}()$
 $\vec{w} = Mw[j]^{-1} \times \vec{bw}[j]$
 choice $\leftarrow \text{UCB}(\vec{w}, \text{occ}[j], \text{getContext}(j), Mw[j]^{-1})$
 reward $\leftarrow \text{getReward}(\text{choice})$
 update $\leftarrow \text{context}[\text{choice}] \cdot \text{context}[\text{choice}]^\top$
 $Mw_{local}[j] += \text{update}$
 $\vec{bw}_{local}[j] += \text{reward} \times \text{context}[\text{choice}]$
 $Mw[j] += \text{MBuff}[j].\text{pop}()$
 $\vec{bw}[j] += \text{bBuff}[j].\text{pop}()$
 MBuff[j].push(update)
 bBuff[j].push(reward $\times$ context[choice])
 occ[j] $\leftarrow \text{occ}[j] + 1$
 end
 shuffledUserOrder $\leftarrow \text{getRandomizedOrder}(numUsers)$
 for each user i **do**
 peer $\leftarrow \text{getRandomPeer}(\text{shuffledUserOrder}[i])$
 $\vec{w} \leftarrow Mw_{local}[i]^{-1} \times \vec{bw}_{local}[i]$
 $\vec{v} \leftarrow Mw_{local}[\text{peer}]^{-1} \times \vec{bw}_{local}[\text{peer}]$
 if $|\vec{w} - \vec{v}| \geq \gamma \times \text{getThreshold}(\text{peer}, i)$ **then**
 remove $E[i, \text{peer}]$
 Reset all entries of MBuff[i], MBuff[peer], bBuff[i], bBuff[peer]
 Reset $Mw[i]$, $\vec{bw}[i]$, $Mw[\text{peer}]$, $\vec{bw}[\text{peer}]$
 end
 else if peer and i share same neighbors **then**
 $\text{MBuff}[i] \leftarrow (\text{MBuff}[i] + \text{MBuff}[\text{peer}]) / 2$
 $\text{bBuff}[i] \leftarrow (\text{bBuff}[i] + \text{bBuff}[\text{peer}]) / 2$
 $Mw[i] \leftarrow (Mw[i] + Mw[\text{peer}]) / 2$
 $\text{bw}[i] \leftarrow (\text{bw}[i] + \text{bw}[\text{peer}]) / 2$
 end
 end
end

i. e., $\vec{w}[i] = Mw[i] \times \vec{bw}[i]$. Each interaction updates the active and the current correlation matrix and the b -vector by dequeuing the corresponding entries at the head of the queue. Correspondingly, a new entry of both the correlation matrix and the b -vector are inserted into the respective buffers.

The above parallel user interaction process continues until one of the users has observed an interaction count of *bufferSize*. This stage marks the beginning of network updates. Each user obtains the buffers from a randomly selected user it is connected with. If the user vectors of the two differ by more than a threshold, then the edge is removed, potentially placing the two users into different clusters. Otherwise, if the neighbor sets of the two users match, then the buffers of the user are updated with average values of the peer's buffers, and its own. After this cluster formation stage is over, user interactions are addressed in parallel again. Note that unlike CLUB,

DCCB does not actively recommend on the basis of a cluster's information.

Unlike CLUB, which updates network per user interaction, DCCB updates the network lazily. It therefore needs to archive the history via buffers. DCCB can be considered as having two repeating stages. The first stage is a massively parallel stage that handles user interactions, while the second stage performs cluster updates. The cluster formation stage requires peer-to-peer communication, although each user only communicates with a single peer.

3 DESIGN OF *DISTCLUB*: A HIGHLY DISTRIBUTED ONLINE RECOMMENDATION SYSTEM

This section describes limitation of DCCB. It presents the *DistCLUB* algorithm, and explains how the algorithm overcomes the limitation faced by DCCB. The section also describes the Spark-based *DistCLUB* system design.

3.1 Limitations of DCCB

DCCB, the state-of-the-art distributed online recommendation algorithm achieves accuracy similar to that of CLUB, but faces the following key challenges prohibiting it from achieving high performance.

- (1) Cluster formation requires a global understanding of the graph. DCCB's lazy update-based approach where a node communicates with a single peer to share updates fails to gather the global view quickly. This slows down the cluster discovery, lowering accuracy.
- (2) DCCB approach archives user-interaction information history in buffers, and communicates the entire history with peers. This results in communication becoming the bottleneck.

The *DistCLUB* algorithm overcomes these issues.

3.2 *DistCLUB* Algorithm

We continue with the notation introduced in Section 2. Listing 3 shows the *DistCLUB* algorithm. *DistCLUB* employs three main strategies a) Efficient parallelism exploitation : When recommendations do not use clustering information, and are based on the user's past interactions alone, user-level interaction parallelism can be exploited explicitly. However, good recommendations require clustering information too. As clusters are discovered, updates to a cluster's information must occur sequentially in order to maintain consistency. A key principle of *DistCLUB* algorithm is to interleave stages of user-based and cluster-based interactions. In the latter, although interactions of users in a given cluster cannot be parallelized, interactions *across* clusters can occur in parallel. b) Limiting communication bottleneck : *DistCLUB* uses no buffers. Instead, it only communicates each user's correlation matrix and user-vector, each time clustering is updated. c) Adaptive user-based and cluster-based round limits : *DistCLUB* dynamically switches between using cluster-specific information, when information for a user is sparse, and highly personalized user-specific information, when information for the user is abundant. This allows *DistCLUB* to achieve high accuracy.

Figure 2: Comparison of clustering in DCCB and *DistCLUB*. Graph G is made of $|V| = 5$ vertices $\{1, 2, 3, 4, 5\}$ representing users.. Clustering in DCCB proceeds lazily - hence fewer edges are removed at a given time, as shown in Figure 2(a), where edges between users $(1, 4)$, $(2, 4)$, and $(1, 5)$ have not yet been broken. However, in *DistCLUB* Figure 2(b), clustering is performed on the entire graph G , forming the clusters more quickly. This allows parallel interaction processing on clusters C_1 and C_2 .

To realize the above strategies we design *DistCLUB* algorithm to have four major repeating stages described as follows:

3.2.1 Stage 1 : Recommending via user information alone. Initially, user correlation matrices and bias vectors contain very little information. However, in subsequent iterations of the algorithm, the reward-based feedback mechanism updates the user correlation matrix and the user-vector. User-vector updates facilitate two outcomes: i) Improved recommendations as more user information is gathered, and ii) Clusters discovery. The recommendations are generated using the linear contextual bandit setting, similar to CLUB. Since users share no data elements, interactions in this stage are embarrassingly parallel. Note that interactions of a given user are processed in order. Cluster updates are deferred to the third stage of the algorithm.

3.2.2 Stage 2: Network updates and cluster formation. This stage is logically equivalent to the network update stage in CLUB. It first updates the network by removing unnecessary edges. Next, it generates clusters and computes their correlation matrices and cluster-vectors. This stage can exploit tree-reduction style parallelism, as will be described in Sec. 3.4.

3.2.3 Stage 3: Cluster-based recommendations. Figure 2 presents a comparison of how clustering occurs in *DistCLUB* and DCCB. The faster cluster formation in *DistCLUB* creates an opportunity for generating cluster-specific recommendations.

This stage is similar to stage 1, except that instead of using the user's information for recommendations, cluster's information is used. This helps normalize the recommendations across all users in the cluster. A distinguishing aspect of this stage is that for users with significantly higher interaction counts than other users in the cluster, the recommendation is still provided on the basis of the user's information, making the recommendation more personalized. This is crucial because in spite of the clustering, users with an exceptionally long history are more likely to choose an item based on their own past activity. To identify such users, *DistCLUB* computes the number of interactions processed for the user, and determines if it is β times more than the average interactions for users in the cluster. If so, then the individual user's information is utilized for recommendations,

else cluster information is used. Interactions where the catered users belong to different clusters can be processed in parallel in this stage.

3.2.4 Stage 4: Balancing stage 1 and 2 sizes. For users who have observed more interactions than others in a given cluster, personalized recommendations will be more accurate. On the contrary, for users with few interactions, recommending based off cluster's information offers better rewards. The number of user interactions to be processed according to cluster information and user information are dynamically chosen in *DistCLUB*. Initially, stage 1 and 2 operate on $uRounds$ and $cRounds$ arrays respectively, which are initialized to the same values for all users. This stage factors in information regarding interaction volume for each user, and updates user and cluster round limits accordingly, by a factor δ , as shown in Algorithm 3. This mechanism automatically steers newly coming users towards using cluster-based recommendations. This approach makes *DistCLUB* a hybrid of CLUB and a fully personalized linear bandit approaches, with the decision to switch taken in an adaptive fashion.

3.3 Discussion

While both DCCB and *DistCLUB* are distributed online recommendation algorithms where nodes share information, there are important differences in their design. Figure 3 compares the data communication sizes of DCCB and *DistCLUB*. In DCCB, a user shares the archived buffer with a random peer. This partial information sharing results in biasing, and only gets mitigated when several communication rounds have occurred. Design alternatives to reduce DCCB data transfers are i) Reducing buffer lengths - if buffer lengths were to be reduced, frequency of communication will increase, amplifying the associated communication overheads, while keeping the overall communication sizes the same. ii) Reducing the number of users that share information in communication round - This approach will maintain the frequency of communication and buffer lengths, but will lower the overall communication volume. Unfortunately, this results in accuracy reduction. Both these alternatives are thus rendered ineffective. The communication bottleneck exacerbates in environments with large user counts because each machine must handle several users, constraining data transfer limits. On the other hand,

Figure 3: Comparing communication requirements: This figure shows a scenario where each node caters to a single user. N is the total number of nodes, D is the feature count. Figure 2(a)- In DCCB each node maintains a buffer of length L which holds an archive of user matrices for each interaction, along with the active matrices. The buffers result in large communications. On the other hand, data exchange in *DistCLUB* is much lighter (Figure 2(b)). Table 4 shows cumulative data transfer sizes for DCCB and *DistCLUB* required in practice, where each node caters to several users.

DistCLUB aggregates and processes information from all nodes during a communication stage removing bias in information sharing.

3.4 Implementation and Optimizations

Figure 4: *DistCLUB* System Implementation in Spark. Figure depicts RDDs in *DistCLUB*'s Spark Jobs.

Choice of Framework We chose to use the Spark Framework [42] to implement *DistCLUB*, owing to its support for large-scale distributed processing, fault-tolerance, in-memory data persistence, and automatic data locality exploitation. Furthermore, Spark automatically exploits both inter-node and intra-node parallelism. Other traditional frameworks such as MPI do not provide automatic in-memory data persistence and out-of-the-box data locality exploitation like in Spark. Spark is highly popular, easy to use, and well-integrated with other big-data processing frameworks. It also supports parallel graph processing libraries. Dask [11] could be a handy framework to parallelize such applications and is lighter weight than Spark, however it is not yet adapted in the Big Data enterprise world.

Spark uses a master-worker paradigm where the Spark Executors are the workers and they execute Map tasks in parallel. Spark jobs have an Application Master that controls launches for all tasks.

Figure 4 depicts the mapping of *DistCLUB* algorithm on Spark. We represent the user correlation matrix M , the vector b , the user-graph G , the cluster correlation matrix M and the user vectors as RDDs (Resilient distributed datasets) [41]. RDD is a Spark primitive, representing a fault-tolerant collection of elements that can be operated on in parallel. Additionally, we make Spark persist these RDDs in the node's memory after each update to them. This way, a subsequent interaction for the same user is fast. This is primarily because of Spark's scheduling mechanism that tries to allocate next such task to the same node. Spark replicates RDDs across multiple nodes in the cluster for fault-tolerance thus making *DistCLUB* robust. Spark Map transformation applies a function to each element in the RDD in parallel and returns a new RDD. We leverage the Map transformation to process the user interactions with the key-value pair being $\langle userID, user-interaction \rangle$ for stage 1 and with a slight modification in stage 3 of *DistCLUB*.

For both these stages, if there are multiple interactions with the same user, their execution is serialized. Spark automatically schedules and executes these Map tasks such that the data which they process is available in the local disk, and preferably in the memory of the node. This typically results in tasks for a given user being executed on the same set of machines. The updated user correlation matrix and b-vector are passed back to the Spark Master.

For Stage 2 of *DistCLUB*, to process network updates efficiently, we perform the edge updates for each user in parallel Map tasks. This merely involves communicating the user-vector information, which is small in size. Clusters are generated based on the updated edges; by performing computation of *connected components* in a distributed and parallel manner in Spark [40]. Updated user vectors are computed using the updated correlation matrix and b-vector. Cluster vectors are calculated for the newly formed clusters. This is done by summing up the constituent user vectors in each cluster. To prevent this step from becoming a bottleneck, aggregation of

Algorithm 3: *DistCLUB* : Performs clustering lazily, and uses both user-driven and cluster-driven recommendations.

Input: β cluster penalizing hyperparameter
Input: σ initial round splitting hyperparameter
Input: $G(V, E)$: Graph where vertices represent users. Edges represent correlation between users. Initially it is a complete graph.
Input: occ : Array of counters for number of interactions with each user, initialized to 0
Input: context: table containing features for all items
Input: $uRounds, cRounds$: arrays representing interaction round limits per user for user-based and cluster-based recommendations
Initial state: $Mu[user] = I \quad \forall \text{ users } \in [0, \text{numUsers}]$
 $\vec{bu}[user] = 0 \quad \forall \text{ users } \in [0, \text{numUsers}]$
 $uRounds[users] = cRounds[users] = \sigma \quad \forall \text{ users } \in [0, \text{numUsers}]$
while user interactions exist **do**
 while user interactions $\leq uRounds[user] \quad \forall \text{ users } \text{ do}
 // explicitly parallel
 $userID \leftarrow \text{getUserInteraction}()$
 $\vec{u} \leftarrow Mu[userID]^{-1} \cdot \vec{bu}[userID]$
 $\text{choice} \leftarrow \text{UCB}(\vec{u}, occ, \text{context}(t), Mu[userID]^{-1})$
 $\text{reward} \leftarrow \text{getReward}(\text{choice})$
 $Mu[userID] += \text{context}[\text{choice}] \cdot \text{context}[\text{choice}]^T$
 $\vec{bu}[userID] += \text{reward} \times \text{context}[\text{choice}]$
 $occ[userID] += 1$
 end
 $\text{UPDATENETWORK}()$
 $\text{clusters} \leftarrow \text{recomputeClusters}(G)$
 for each cluster c in clusters **do**
 for each user n in c **do**
 $Mu[c] += Mu[n]$
 $\vec{bc}[c] += \vec{bu}[n]$
 end
 end
 while user interactions $\leq cRounds[user] \quad \forall \text{ users } \text{ do}
 // interactions across clusters can execute in parallel
 $userID \leftarrow \text{getUserInteraction}()$
 $\text{clusterID} \leftarrow \text{clusters}[userID]$
 if $occ[userID] \geq \beta \times \frac{\text{numClusterInteractionsSeen}}{\text{numUsersInCluster}}$ **then** $\vec{\theta} \leftarrow Mu[userID]^{-1} \cdot \vec{bu}[userID]$;
 else $\vec{\theta} \leftarrow Mc[\text{clusterID}]^{-1} \cdot \vec{bc}[\text{clusterID}]$;
 $\text{choice} \leftarrow \text{UCB}(\vec{\theta}, occ, \text{context}(userID), Mc[\text{clusterID}]^{-1})$
 $\text{reward} \leftarrow \text{getReward}(\text{choice})$
 $Mu[userID] += \text{context}[\text{choice}] \cdot \text{context}[\text{choice}]^T$
 $\vec{bu}[userID] += \text{reward} \times \text{context}[\text{choice}]$
 $occ[userID] += 1$
 $\text{numClusterInteractionsSeen}[\text{clusterID}] += 1$
 end
 for each cluster c in clusters **do**
 $\text{meanOccInCluster} \leftarrow \frac{\text{numClusterInteractionsSeen}[c]}{\text{numUsersInCluster}}$
 for each user n in c **do**
 $\delta \leftarrow \frac{occ[n] - \text{meanOccInCluster}}{\beta}$
 $uRounds[n] += \delta$
 $cRounds[n] -= \delta$
 end
 end
end$$

user-vectors into the cluster vectors is performed in parallel, for each cluster. To this end, we make use of Spark's tree aggregation pattern, and utilize treeReduce to collect the correlation matrix and b-vector of users in a cluster.

Stage 4 of the algorithm is trivially parallelizable using Map tasks in Spark. In this manner, all four stages of *DistCLUB* are *optimized* and mapped into *efficient* distributed and parallel primitives in Spark.

4 ANALYSIS OF *DISTCLUB*

The regret of a bandit learner relative to a learning policy is defined to be the difference between the total expected reward using that policy for T recommendation rounds and the total expected reward collected by the learner over T rounds. Therefore, regret analysis is used to compare the learning rate of *DistCLUB* with alternative approaches, such as DCCB and CLUB. This section shows that regret bounds for *DistCLUB* are similar to those of DCCB and CLUB. To derive the regret bounds, we first establish our learning model.

4.1 Learning model for *DistCLUB*

Let $V = \{1, \dots, n\}$ represent the set of n users. User behaviour similarity induces a partitioning of V into c sets $V_1, V_2, \dots, V_c$ with $c \ll n$. Users belonging to the same set or cluster share similar behavior while users from different clusters have different behaviors. Note that, some sets in these partitioning can be singleton sets, such that the user does not share similar behaviour with other users. No partitioning information of V or user behavior are known to the learner in the beginning. In each user interaction t the learner receives a user i_t , together with a set of context vectors $C_{i_t} \subset R^d$ from which the learner chooses an action $A(t)$. d represents the number of features for each action in C_{i_t} . Subsequently, the learner observes reward $W(t)$, based on the chosen action. We assume i_t is selected uniformly at random from V .

4.2 Regret Analysis

Denote T as the total number of iterations and T_i as the total number of iterations when user i is selected up to time T . So we have $T = \sum_{i \in V} T_i$. Standard contextual bandit analysis shows that, for a particular user, as T_i grows large, the cumulative regret can be bounded with high probability [1, 14] as

$$R_{\text{IND}}(T, i) = O(\sqrt{dT_i}) \quad (1)$$

Where $r_{t,i}$ is the regret from interaction processed at time t with user i , and $R_{\text{IND}}(T, i)$ is the aggregate regret over T interactions for user i by running an independent linear contextual bandit model. The total aggregate regret over all the users, denoted as $R_{\text{IND}}(T)$ can be written as,

$$R_{\text{IND}}(T) = \sum_{i \in V} R(T, i) = O\left(\sum_{i \in V} \sqrt{dT_i}\right) \quad (2)$$

If users are clustered, which means that observations are shared across users, the regret bound of a cluster of user V_j and over all regret bound can be written as,

$$R_{\text{CLU}}(T, V_j) = O\left(\sqrt{d \sum_{i \in V_j} T_i}\right) \quad (3)$$

$$R_{\text{CLU}}(T) = \sum_{j=1}^c R_{\text{CLU}}(T, V_j) = O\left(\sum_{j=1}^c \sqrt{d \sum_{i \in V_j} T_i}\right) \quad (4)$$

Since users i_t are chosen randomly, we can consider $T_i = \frac{T}{n}$, the above regret can be simplified as,

$$R_{\text{IND}}(T) = O(\sqrt{dT} \sum_{i \in V} \sqrt{1/n}) \quad (5)$$

$$R_{\text{CLU}}(T) = O(\sqrt{dT} \sum_{j=1}^c \sqrt{|V_j|/n}) \quad (6)$$

If we consider regret in the clustering case, i.e. $R_{\text{CLU}}(T)$ as the reference regret, regret is only affected when user interactions are processed in *DistCLUB*. This happens in Stage 1 and Stage 3 of the algorithm. In the other stages, regret remains unchanged. The regret bound depends on the size of clusters V_j . In Stage 1, the user correlation matrix and the user-vector are considered only for choosing the action $A(t)$, and no information is shared across users. In this case, regret is the same as running independent linear contextual bandit models. Hence the regret bound in this stage, denoted as R_{S1} , can be written as,

$$R_{S1}(T) = R_{\text{IND}}(T) = O(\sqrt{ndT}) \quad (7)$$

After clustering of users in Stage 2, assume that m clusters are formed such that they contain n_1 users overall, and let there be n_2 users ($n_2 \ll n, m \ll n$) whose rewards are generated without considering their cluster information. Therefore, $n_1 + n_2 = n$ also $m + n_2 = c$. In the worst case, i.e. make the summation term significant, let m clusters have the same size n_1/m . Hence substituting this information back in equation (5) and equation (6), we have,

$$\begin{aligned} R_{S2}(T) &= O(\sqrt{dT}(m \times \sqrt{(n_1/m)/n} + n_2 \times \sqrt{1/n})) \\ &= O(\sqrt{dT}(\sqrt{n_1 m/n} + \sqrt{n_2^2/n})) \end{aligned} \quad (8)$$

Since $n_2 \ll n, n_1/n = 1$ and n_2^2/n can be dropped because of the O notation. Hence

$$R_{S2}(T) = O(\sqrt{dTm}) \quad (9)$$

In the degenerate case, the regret will fall back to equation (7). The regret thus has the same bounds as DCCB and CLUB, and as the number of user interactions processed by *DistCLUB* tends to ∞ , the final regret is asymptotically equivalent to CLUB and DCCB [14, 19].

5 EVALUATION

In this section we present an experimental comparison of *DistCLUB* with DCCB and CLUB. The experiments compare execution times and prediction performance of *DistCLUB*, CLUB, DCCB for multiple datasets. The section also compares scalability of *DistCLUB* and DCCB.

5.1 Experimental Setup

We used Amazon EC2 to run our experiments. Each compute node is a c5.4x [5] instance with 32 GB of RAM. Each node contains eight cores. The c5 instance has an Intel Xeon Platinum 8000 series processor with CPU clock speed of up to 3.6 GHz. The machines in this cluster have a network bandwidth of up to 10 Gbps.

Dataset	Number of User Interactions	Number Of Users	Number of Item Features
MovieLens	80,000	943	19
LastFM	10,000	1,888	25
Delicious	10,000	1,816	25
Yahoo	50,000	5,045	1
Synthetic	4,000,000	20,000	25

Table 1: Datasets used in the evaluation

We used Spark version 2.4.3 and EMR version 5.26.0 in the experiments. The EMR cluster was run in cluster mode, wherein the Application Master and Spark driver together with the Spark Executors occupy all cores in the system.

Table 1 shows the datasets we used for the experiments. The values used for all configurable parameters are shown in Table 2. We followed the dataset preparation described in Li et al. [14]. These datasets are used to evaluate *DistCLUB* in dynamic environments where the set of active users and user preferences change rapidly, and have been used in past for evaluating CLUB and DCCB algorithms [14, 19]. The evaluation is performed over four real-world benchmark datasets and one synthetic dataset, each having different user, interaction, and item feature counts. The MovieLens dataset [22] contains personal ratings and tags for multiple users. Delicious is a social bookmarking webservice system, while LastFM is a social music streaming system. These datasets were obtained from [15]. The Yahoo dataset is obtained from the ICML 2012 Exploration and Exploitation Challenge for news article recommendation and was preprocessed following the directions in [14]. The synthetic dataset was created with twenty thousand users and four million user interactions. We created the synthetic dataset to stress *DistCLUB* with much larger number of users, clusters, and user-interactions than the above datasets. The synthetic dataset models today's recommendation systems that have high volume of users and are expected to process interactions at a high rate. Table 2 shows the various configurations/hyperparameter values used in the evaluation.

Parameter	Value
α : Exploration Parameter in UCB	0.03
β : Cluster penalizing parameter in <i>DistCLUB</i>	2
γ : Cluster Exploration Parameter in CLUB	0.7
δ : networkUpdate Delay for CLUB	2000
<i>bufferSize</i> : Buffer length for DCCB	5000
σ : Initial split sizes in <i>DistCLUB</i>	2500

Table 2: Hyper-parameter values used in the experiments

DCCB Implementation: Since no open-source system implementation for DCCB is available, we implemented the DCCB algorithm (Listing 2) in Spark, by closely following the description in their paper [19]. In DCCB, interactions of different users are processed in parallel. We used Spark Map transformation to implement this stage, where the key-value pair is <userID, user-interaction to process>. After this stage, information sharing happens in a peer-to-peer manner. Note that big data processing systems such as Hadoop or Spark do not directly allow processes to share information outside

of their reduction primitives. To facilitate this information sharing to occur in a parallel fashion, we created a Spark Map transformation where the key-value pair is $\langle \text{receiving-userID}, (\text{sender-ID}, \text{senderM}, \text{senderb}, \text{senderMBuffer}, \text{senderbBuffer}) \rangle$. Inside this transformation, the receiving user's buffers are active objects and are updated according to the algorithm.

5.2 Execution Time Performance Comparison

Figure 5: Speedup of *DistCLUB* over DCCB on different datasets. The geometric mean speedup is 8.87x.

Table 3 shows the execution times in seconds for CLUB, DCCB, and *DistCLUB*. CLUB cannot exploit cross interaction parallelism, while DCCB and *DistCLUB* exploit interaction-level parallelism executing on 8 nodes, where each node has 8 cores. CLUB can exploit intra-interaction parallelism, such as in the matrix multiplication routines. Figure 5 plots the speedup achieved by *DistCLUB* over DCCB. We could not finish running the synthetic dataset in CLUB because it did not finish within three hours. We executed each run thrice, and calculated the average execution times. *DistCLUB* outperforms both DCCB and CLUB. The geo-mean speedup achieved by *DistCLUB* over DCCB and CLUB is 8.87x and 29.13x respectively. For the synthetic dataset, *DistCLUB* runs 17.65 times faster than DCCB because communication overheads become severe with large user and interaction counts.

Execution Time in seconds			
Dataset	CLUB	DCCB	<i>DistCLUB</i>
MovieLens	2395	452.5	68
LastFM	1459	294	49
Delicious	2074	249	31
Yahoo	1785	1693	174
Synthetic	-	7730	438

Table 3: Execution times of *DistCLUB*, CLUB, and DCCB in seconds. DCCB and *DistCLUB* use 64 cores (8 nodes). CLUB does not exploit interaction-level parallelism.

5.3 Scalability

To evaluate the scalability of *DistCLUB* and DCCB we ran them on a real dataset (Yahoo) and a synthetic dataset. The synthetic dataset has greater number of interactions and users than any other dataset.

Figure 6: Scalability comparison of *DistCLUB* and DCCB : DCCB does not scale beyond 128 cores owing to communication overheads.

We increased the number of cores from 64 to 512 (going from 8 to 64 nodes) and measured the speedup achieved, compared to the 64 core variant. *DistCLUB* has an average scaling efficiency [29] of 81.2%. This shows that *DistCLUB* can effectively leverage all cores in the system and scale to larger number of nodes.

For large user and interaction counts, each node in the distributed setup ends up catering to several users. In DCCB, this results in the reduction of effective memory available to each user. For large user and interaction counts, buffer sizes shrink significantly, resulting in more frequent communication rounds.

5.4 Size of information passed around in *DistCLUB* and DCCB

Table 4 shows the cumulative data shared by DCCB and *DistCLUB* during the execution of various datasets over the total interactions. DCCB shares significantly more data than *DistCLUB*. The data transferred in DCCB during each communication stage comprises both the buffer and the active objects, and the overall data size depends on the number of interactions and buffer length (see Figure 3). On the other hand, in *DistCLUB*, data transfer is driven by the interaction count alone.

Dataset	DCCB	<i>DistCLUB</i>
MovieLens	114.6 GB	29.9 MB
LastFM	46 GB	11.1 MB
Delicious	42.2 GB	10.4MB
Yahoo	2.1 GB	423 KB
Synthetic	204 TB	43.6 GB

Table 4: Comparison of overall data transfer sizes in DCCB and *DistCLUB* : Cumulatively, DCCB shares significantly more information than *DistCLUB*

5.5 Cluster Discovery Speed

Figure 7 compares the cluster formation rate of DCCB and *DistCLUB* on the MovieLens dataset. The figure shows that *DistCLUB* finds clusters faster than DCCB. Other datasets show similar patterns

Dataset	<i>DistCLUB</i> Reward units	DCCB Reward units	CLUB Reward units	<i>DistCLUB</i> to DCCB reward
MovieLens	5,620	5,321	7,014	1.06
LastFM	1,242	896	1,494	1.39
Delicious	484	456	428	1.06
Yahoo	2,279	2,138	2105	1.07

Table 5: Comparison of normalized cumulative rewards units obtained by *DistCLUB* in comparison to DCCB and CLUB. The last column highlights *DistCLUB*'s prediction performance compared to DCCB for each of the dataset. *DistCLUB*'s prediction performance is 14.5% higher than DCCB normalized over all datasets

Figure 7: Cluster discovery rate in DCCB and *DistCLUB*. As interaction count grows, *DistCLUB* discovers more clusters than DCCB. Faster cluster discovery results in better recommendations.

in cluster growth rate. Information sharing in DCCB is restricted to a single peer node, while in *DistCLUB*, information from all nodes is globally aggregated. As a result, information sharing is much slower in DCCB, and this slows down the formation of clusters as shown in Fig 7.

5.6 Cumulative Reward Comparison

We measure the ratio of cumulative reward obtained for *DistCLUB*, CLUB, and DCCB to that obtained by a predictor which chooses an item within a context fully at random as a recommendation. Figure 8 shows results of this comparison, where each run was executed thrice. The algorithm behaviour is erratic near the zero horizon. This is because the exploration in the beginning interaction rounds is noisy. The objective is to achieve largest possible cumulative reward over all user interactions [23]. We divide the cumulative reward obtained by the algorithm with the cumulative reward obtained by a predictor which chooses a random item, and then obtain average prediction performance across datasets. *DistCLUB* achieves 14.5% higher prediction performance than DCCB normalized over all the datasets. Because reward value is 0 or 1 in all datasets, achieved rewards can be averaged and compared across datasets. Table 5 shows the performance on each individual dataset.

On LastFM and Yahoo datasets, *DistCLUB* performs much better than DCCB by efficiently finding more clusters, indicating that social information plays higher role in achieving better predictions on these datasets. CLUB performs the best, owing to the fact that it serially processes individual interactions, finding better clustering.

5.7 Cumulative Regret Comparison

Figure 9 compares the ratio of the cumulative regret for *DistCLUB*, CLUB and DCCB to a random predictor RAN which randomly chooses an item. The cumulative regret is the difference between the algorithm's accumulated reward and the maximum the algorithm could have obtained if all parameters are known, so lower regret is better. Regret is compared in the same setting as the rewards. For all datasets, overall regret for *DistCLUB* is lower than DCCB. The experiment reaffirms the regret analysis presented in Sec. 4.2 which showed that asymptotic regret bounds for *DistCLUB* are same as that of DCCB and CLUB.

6 RELATED WORK

Offline recommendation algorithms Recommendation systems' algorithms can be classified into two major methods - Collaborative Filtering (CF) and Content-based methods. CF techniques calculate similarity based on interactions. CF techniques use algorithms like k-nearest neighbors [32] and matrix factorization [20] to predict the utility of items to users. SGD [45] or ALS [33] are used in training such offline algorithms. Content-based recommendation systems recommend items to users based on user's interest profile and item descriptions [31]. Techniques such as Decision Trees [9], Naive Bayes [43], nearest neighbor methods, and linear classifiers [2] are used to learn a model for the user. DNN-based recommenders have been used to combine CF (through autoencoders) and content-based (training on item attributes) approaches [8, 38, 39]. Distributed CF methods have been proposed to solve scalability issues in centralized CF schemes [4, 16]. However, all these methods are trained offline, and suffer in performance when the item-set or user-preferences change dynamically. In most social systems of today, it is common to see an item popularity change and evolution in user interests. It therefore becomes essential to learn continuously, requiring online algorithms.

Bandit-based recommendation systems: Many bandit-based approaches for online recommendations exist. Stochastic multi-armed bandits with single and multiple instances of linear bandits were outlined in [24]. Specifically, LinUCB models expected reward through linear regression on context vectors. CLUB [14] introduces bandit-based clustering techniques. DynUCB [30] also propose dynamic clustering using contextual multi-armed bandits. They propose a flat clustering structure, and allow users to switch clusters. Our approach can leverage their clustering technique to allow such user movement between clusters. While CLUB, DCCB, and *DistCLUB* only identify user-based clustering, it is possible to simultaneously identify user and item clusters collaboratively such as in [25].

Among distributed and decentralized bandits, DCCB [19] proposes decentralised clustering algorithms to solve linear bandit problems in peer-to-peer networks by sharing information among collaborating workers. However, DCCB suffers scalability issues owing

Figure 8: Cumulative reward compared to random action (RAN) for *DistCLUB* (red), DCCB (green), and CLUB (blue) various datasets: Near the 0 horizon the algorithm behaviour is erratic since the initial exploration is noisy. Prediction performance (cumulative reward) is higher for *DistCLUB* than DCCB for all datasets.

to large transfer volumes. Tekin et al. propose a decentralized recommendation system [36], where a group of agents can recommend items to users, and privacy and autonomy are important. Agents are prohibited from learning from other agents, and can only use their own experience. This is a fundamentally different constraint than other online systems, such as *DistCLUB*, DCCB, and CLUB.

RL-based online recommendation algorithms: Reinforcement learning (RL) techniques are an alternative to bandit-based recommendation systems. Among those, model-based techniques, such as [35] are unsuited for online recommendations owing to their high complexity [26]. Deep RL-based techniques are suitable for online recommendations. Zhao et al. [44] use a deep Q-network trained with embeddings of users’ historical clicked/ordered items (state) and a recommended item (action) as input. Training this state-space generates recommendations. However, such approaches need to evaluate the Q-values of all the actions under a specific state, which is inefficient for large interaction counts. In general, RL-based approaches involve significant complexity owing to the state space they preserve. For instance, DDR [26] stores state transitions in buffers.

Such requirements prohibit RL-based techniques from being decentralized, leaving bandit-based approaches alone to be amenable to distribution.

7 CONCLUSION

This paper has presented *DistCLUB*, a distributed algorithm and system applicable to contextual bandit algorithms that are prevalent in online recommendation systems where content and user preferences change dynamically. Although DCCB, the state-of-the-art online recommendation system, distributes the computation, it fails to scale effectively. The key reason for this behaviour is the communication bottleneck. *DistCLUB* algorithm overcomes this issue by intelligently mixing cluster-based and user-based recommendation schemes. This mixing allows *DistCLUB* to generate clusters faster, thereby obtaining better accuracy.

DistCLUB is implemented using Spark, and efficiently employs distributed primitives in the Spark framework. The paper showed that *DistCLUB* can scale to 512 cores with high efficiency. Experimental results over both real-world and synthetic datasets show that *DistCLUB* operates on average 8.87x faster than DCCB, and

Figure 9: Regret values obtained for various algorithms : *DistCLUB* (red), *DCCB* (green), and *CLUB* (blue) for MovieLens, LastFM, Delicious, and Yahoo datasets. *DCCB* regret is higher than *DistCLUB* for all datasets (lower is better).

achieves 14.5% better normalized prediction performance. These results indicate that *DistCLUB* can be readily deployed in practice.

The presence of large matrix-vector products make this approach highly suitable to be executed on GPUs. We leave GPU porting of *DistCLUB* to the future work.

REFERENCES

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. 2011. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*. 2312–2320.
- [2] David Adedayo Adeniyi, Zhaoqiang Wei, and Y Yongquan. 2016. Automated web usage data mining and recommendation system using K-Nearest Neighbor (KNN) classification method. *Applied Computing and Informatics* 12, 1 (2016), 90–108.
- [3] Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. 2014. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*. 1638–1646.
- [4] Kamal Ali and Wijnand Van Stam. 2004. TiVo: making show recommendations using a distributed collaborative filtering architecture. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. 394–401.
- [5] Amazon. 2017. Amazon EC2 C5 Instances. (2017). <https://aws.amazon.com/ec2/instance-types/c5/>
- [6] Peter Auer. 2002. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research* 3, Nov (2002), 397–422.
- [7] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [8] Pierre Baldi. 2012. Autoencoders, unsupervised learning, and deep architectures. In *Proceedings of ICML workshop on unsupervised and transfer learning*. 37–49.
- [9] Yoon Ho Cho, Jae Kyeong Kim, and Soung Hie Kim. 2002. A personalized recommender system based on web usage mining and decision tree induction. *Expert systems with Applications* 23, 3 (2002), 329–342.
- [10] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. 2011. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. 208–214.
- [11] Dask Development Team. 2016. *Dask: Library for dynamic task scheduling*. <https://dask.org>
- [12] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM computing surveys (CSUR)* 46, 4 (2014), 1–37.
- [13] Florent Garcin, Boi Faltings, Olivier Donatsch, Ayar Alazzawi, Christophe Bruttin, and Amr Huber. 2014. Offline and online evaluation of news recommender systems at swissinfo. ch. In *Proceedings of the 8th ACM Conference on Recommender systems*. 169–176.
- [14] Claudio Gentile, Shuai Li, and Giovanni Zappella. 2014. Online clustering of bandits. In *Proceedings of the 31st International Conference on Machine Learning, Journal of machine learning research workshop and conference proceedings*. 757–765.
- [15] GroupLens. 2011. *GroupLens Datasets*. www.grouplens.org/node/462
- [16] Peng Han, Bo Xie, Fan Yang, and Ruimin Shen. 2004. A scalable P2P recommender system based on distributed collaborative filtering. *Expert systems with applications* 27, 2 (2004), 203–210.
- [17] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 549–558.

- [18] Eshcar Hillel, Zohar S Karnin, Tomer Koren, Ronny Lempel, and Oren Somekh. 2013. Distributed exploration in multi-armed bandits. In *Advances in Neural Information Processing Systems*. 854–862.
- [19] Nathan Korda, Balázs Szörényi, and Shuai Li. 2016. Distributed clustering of linear bandits in peer to peer networks. In *Proceedings of the 33rd International Conference on Machine Learning, Journal of machine learning research workshop and conference proceedings*, Vol. 48. 1301–1309.
- [20] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 8 (2009), 30–37.
- [21] Tze Leung Lai et al. 1987. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics* 15, 3 (1987), 1091–1114.
- [22] Shyong Lam and Jon Herlocker. 2006. MovieLens data sets. *Department of Computer Science and Engineering at the University of Minnesota* (2006).
- [23] Tor Lattimore and Csaba Szepesvári. 2018. Bandit algorithms. *Cambridge University Press* (2018).
- [24] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*. ACM, 661–670.
- [25] Shuai Li, Alexandros Karatzoglou, and Claudio Gentile. 2016. Collaborative filtering bandits. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 539–548.
- [26] Feng Liu, Ruiming Tang, Xutao Li, Yunming Ye, Haokun Chen, Huifeng Guo, and Yuzhou Zhang. 2018. Deep Reinforcement Learning based Recommendation with Explicit User-Item Interactions Modeling. *CoRR* abs/1810.12027 (2018). arXiv:1810.12027 <http://arxiv.org/abs/1810.12027>
- [27] James McInerney, Benjamin Lacker, Samantha Hansen, Karl Higley, Hugues Bouchard, Alois Gruson, and Rishabh Mehrotra. 2018. Explore, exploit, and explain: personalizing explainable recommendations with bandits. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 31–39.
- [28] McKinsey. 2013. *How retailers can keep up with consumers*. <https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers>
- [29] Mediawiki. 2016. *Measuring Scaling Performance*. http://www.sharcnet.ca/help/index.php/Measuring_Parallel_Scaling_Performance#Calculating_Strong_Scaling_Efficiency Accessed: 2016-03-01.
- [30] Trong T Nguyen and Hady W Lauw. 2014. Dynamic clustering of contextual multi-armed bandits. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 1959–1962.
- [31] Michael J Pazzani and Daniel Billsus. 2007. Content-based recommendation systems. In *The adaptive web*. Springer, 325–341.
- [32] Leif E Peterson. 2009. K-nearest neighbor. *Scholarpedia* 4, 2 (2009), 1883.
- [33] István Pilászy, Dávid Zibriczky, and Dávid Tikk. 2010. Fast als-based matrix factorization for explicit and implicit feedback datasets. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 71–78.
- [34] Jeffrey C Schlimmer and Richard H Granger. 1986. Beyond Incremental Processing: Tracking Concept Drift. In *AAAI*. 502–507.
- [35] Guy Shani, David Heckerman, and Ronen I. Brafman. 2005. An MDP-Based Recommender System. *J. Mach. Learn. Res.* 6 (Dec. 2005), 1265–1295. <http://dl.acm.org/citation.cfm?id=1046920.1088715>
- [36] Cem Tekin, Simpson Zhang, and Mihaela van der Schaar. 2014. Distributed online learning in social recommender systems. *IEEE Journal of Selected Topics in Signal Processing* 8, 4 (2014), 638–652.
- [37] Alexey Tsymbal. 2004. The problem of concept drift: definitions and related work. *Computer Science Department, Trinity College Dublin* 106, 2 (2004), 58.
- [38] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative deep learning for recommender systems. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 1235–1244.
- [39] Xinxi Wang and Ye Wang. 2014. Improving content-based and hybrid music recommendation using deep learning. In *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 627–636.
- [40] Reynold S Xin, Joseph E Gonzalez, Michael J Franklin, and Ion Stoica. 2013. Graphx: A resilient distributed graph system on spark. In *First International Workshop on Graph Data Management Experiences and Systems*. ACM, 2.
- [41] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: Cluster computing with working sets. *HotCloud* 10, 10-10 (2010), 95.
- [42] Matei Zaharia, Reynold S Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J Franklin, et al. 2016. Apache spark: a unified engine for big data processing. *Commun. ACM* 59, 11 (2016), 56–65.
- [43] Tong Zhang and Vijay S Iyengar. 2002. Recommender systems using linear classifiers. *Journal of machine learning research* 2, Feb (2002), 313–334.
- [44] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep reinforcement learning for page-wise recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, 95–103.
- [45] Yong Zhuang, Wei-Sheng Chin, Yu-Chin Juan, and Chih-Jen Lin. 2013. A fast parallel SGD for matrix factorization in shared memory systems. In *Proceedings of the 7th ACM conference on Recommender systems*. ACM, 249–256.