

Distributed Clustering of Linear Bandits in Peer to Peer Networks

Nathan Korda

MLRG, University of Oxford

NATHAN@ROBOTS.OX.AC.UK

Balázs Szörényi

EE, Technion & MTA-SZTE Research Group on Artificial Intelligence

SZORENYI.BALAZS@GMAIL.COM

Shuai Li

DiSTA, University of Insubria

SHUAILI.SLI@GMAIL.COM

Abstract

We provide two distributed confidence ball algorithms for solving linear bandit problems in peer to peer networks with limited communication capabilities. For the first, we assume that all the peers are solving the same linear bandit problem, and prove that our algorithm achieves the optimal asymptotic regret rate of any centralised algorithm that can instantly communicate information between the peers. For the second, we assume that there are clusters of peers solving the same bandit problem within each cluster, and we prove that our algorithm discovers these clusters, while achieving the optimal asymptotic regret rate within each one. Through experiments on several real-world datasets, we demonstrate the performance of proposed algorithms compared to the state-of-the-art.

1. Introduction

Bandits are a class of classic optimisation problems that are fundamental to several important application areas. The most prominent of these is recommendation systems, and they can also arise more generally in networks (see, e.g., (Li et al., 2013; Hao et al., 2015)).

We consider settings where a network of agents are trying to solve collaborative linear bandit problems. Sharing experience can improve the performance of both the whole network and each agent simultaneously, while also increasing robustness. However, we want to avoid putting too much strain on communication channels. Communicating every piece of information would just overload these chan-

nels. The solution we propose is a gossip-based information sharing protocol which allows information to diffuse across the network at a small cost, while also providing robustness.

Such a set-up would benefit, for example, a small start-up that provides some recommendation system service but has limited resources. Using an architecture that enables the agents (the client’s devices) to exchange data between each other directly and to do all the corresponding computations themselves could significantly decrease the infrastructural costs for the company. At the same time, without a central server, communicating all information instantly between agents would demand a lot of bandwidth.

Multi-Agent Linear Bandits In the simplest setting we consider, all the agents are trying to solve the same underlying linear bandit problem. In particular, we have a set of *nodes* V , indexed by i , and representing a finite set of *agents*. At each time, t :

- a set of *actions* (equivalently, the *contexts*) arrives for each agent i , $\mathcal{D}_t^i \subset \mathcal{D}$ and we assume the set \mathcal{D} is a subset of the unit ball in \mathbb{R}^d ;
- each agent, i , chooses an action (context) $x_t^i \in \mathcal{D}_t^i$, and receives a *reward*

$$r_t^i = (x_t^i)^\top \theta + \xi_t^i,$$

where θ is some unknown *coefficient vector*, and ξ_t^i is some zero mean, R -subGaussian noise;

- last, the agents can share information according to some protocol across a communication channel.

We define the *instantaneous regret* at each node i , and, respectively, the *cumulative regret* over the whole network to be:

$$\rho_t^i := \left(x_t^{i,*}\right)^\top \theta - \mathbb{E}r_t^i, \quad \text{and} \quad \mathcal{R}_t := \sum_{k=1}^t \sum_{i=1}^{|V|} \rho_k^i,$$

where $x_t^{i,*} := \arg \max_{x \in \mathcal{D}_t^i} x^\top \theta$. The aim of the agents is to minimise the rate of increase of cumulative regret. We also wish them to use a sharing protocol that does not impose much strain on the information-sharing communication channel.

Gossip protocol In a gossip protocol (see, e.g., (Kempe et al., 2003; Xiao et al., 2007; Jelasity et al., 2005; 2007)), in each round, an overlay protocol assigns to every agent another agent, with which it can share information. After sharing, the agents aggregate the information and, based on that, they make their corresponding decisions in the next round. In many areas of distributed learning and computation gossip protocols have offered a good compromise between low-communication costs and algorithm performance. Using such a protocol in the multi-agent bandit setting, one faces two major challenges.

First, information sharing is not perfect, since each agent acquires information from only one other (randomly chosen) agent per round. This introduces a bias through the unavoidable doubling of data points. The solution is to mitigate this by using a delay (typically of $O(\log t)$) on the time at which information gathered is used. After this delay, the information is sufficiently mixed among the agents, and the bias vanishes.

Second, in order to realize this delay, it is necessary to store information in a buffer and only use it to make decisions after the delay has been passed. In (Szörényi et al., 2013) this was achieved by introducing an epoch structure into their algorithm, and emptying the buffers at the end of each epoch.

The Distributed Confidence Ball Algorithm (DCB) We use a gossip-based information sharing protocol to produce a distributed variant of the generic Confidence Ball (CB) algorithm, (Abbasi-Yadkori et al., 2011; Dani et al., 2008; Li et al., 2010). Our approach is similar to (Szörényi et al., 2013) where the authors produced a distributed ϵ -greedy algorithm for the simpler multi-armed bandit problem. However their results do not generalise easily, and thus significant new analysis is needed. One reason is that the linear setting introduces serious complications in the analysis of the delay effect mentioned in the previous paragraphs. Additionally, their algorithm is epoch-based, whereas we are using a more natural and simpler algorithmic structure. The downside is that the size of the buffers of our algorithm grow with time. However, our analyses easily transfer to the epoch approach too. As the rate of growth is logarithmic, our algorithm is still efficient over a very long time-scale.

The simplifying assumption so far is that all agents are solving the same underlying bandit problem, i.e. finding the same unknown θ -vector. This, however, is often unre-

alistic, and so we relax it in our next setup. While it may have uses in special cases, DCB and its analysis can be considered as a base for providing an algorithm in this more realistic setup, where some variation in θ is allowed across the network.

Clustered Linear Bandits Proposed in (Gentile et al., 2014; Li et al., 2016a;b), this has recently proved to be a very successful model for recommendation problems with massive numbers of users. It comprises a multi-agent linear bandit model agents' θ -vectors are allowed to vary across a clustering. This clustering presents an additional challenge to find the groups of agents sharing the same underlying bandit problem before information sharing can accelerate the learning process. Formally, let $\{U^k\}_{k=1,\dots,M}$ be a clustering of V , assume some coefficient vector θ^k for each k , and let for agent $i \in U^k$ the reward of action x_t^i be given by

$$r_t^i = (x_t^i)^\top \theta^k + \xi_t^i.$$

Both clusters and coefficient vectors are assumed to be initially unknown, and so need to be learnt on the fly.

The Distributed Clustering Confidence Ball Algorithm (DCCB) The paper (Gentile et al., 2014) proposes the initial centralised approach to the problem of clustering linear bandits. Their approach is to begin with a single cluster, and then incrementally prune edges when the available information suggests that two agents belong to different clusters. We show how to use a gossip-based protocol to give a distributed variant of this algorithm, which we call DCCB.

Our main contributions In Theorems 1 and 6 we show our algorithms DCB and DCCB achieve, in the multi-agent and clustered setting, respectively, near-optimal improvements in the regret rates. In particular, they are of order almost $\sqrt{|V|}$ better than applying CB without information sharing, while still keeping communication cost low. And our findings are demonstrated by experiments on real-world benchmark data.

2. Linear Bandits and the DCB Algorithm

The generic Confidence Ball (CB) algorithm is designed for a single agent linear bandit problem (i.e. $|V| = 1$). The algorithm maintains a confidence ball $C_t \subset \mathbb{R}^d$ within which it believes the true parameter θ lies with high probability. This confidence ball is computed from the observation pairs, $(x_k, r_k)_{k=1,\dots,t}$ (for the sake of simplicity, we dropped the agent index, i). Typically, the covariance matrix $A_t = \sum_{k=1}^t x_k x_k^\top$ and b -vector, $b_t = \sum_{k=1}^t r_k x_k$, are sufficient statistics to characterise this confidence ball. Then, given its current action set, \mathcal{D}_t , the agent selects the optimistic action, assuming that the true parameter sits in C_t , i.e. $(x_t, \sim) = \arg \max_{(x,\theta') \in \mathcal{D}_t \times C_t} \{x^\top \theta'\}$. Pseudocode for CB is given in the Appendix A.1.

Gossip Sharing Protocol for DCB We assume that the agents are sharing across a peer to peer network, i.e. every agent can share information with every other agent, but that every agent can communicate with only one other agent per round. In our algorithms, each agent, i , needs to maintain

- (1) a *buffer* (an ordered set) \mathcal{A}_t^i of covariance matrices and an *active* covariance matrix \tilde{A}_t^i ,
- (2) a *buffer* \mathcal{B}_t^i of b-vectors and an *active* b-vector \tilde{b}_t^i ,

Initially, we set, for all $i \in V$, $\tilde{A}_0^i = I$, $\tilde{b}_0^i = 0$. These *active* objects are used by the algorithm as sufficient statistics from which to calculate confidence balls, and summarise only information gathered before or during time $\tau(t)$, where τ is an arbitrary monotonically increasing function satisfying $\tau(t) < t$. The buffers are initially set to $\mathcal{A}_0^i = \emptyset$, and $\mathcal{B}_0^i = \emptyset$. For each $t > 1$, each agent, i , shares and updates its buffers as follows:

- (1) a random permutation, σ , of the numbers $1, \dots, |V|$ is chosen uniformly at random in a decentralised manner among the agents,¹
- (2) the buffers of i are then updated by averaging its buffers with those of $\sigma(i)$, and then extending them using their current observations²

$$\mathcal{A}_{t+1}^i = \left(\left(\frac{1}{2}(\mathcal{A}_t^i + \mathcal{A}_t^{\sigma(i)}) \right) \circ \left(x_{t+1}^i (x_{t+1}^i)^\top \right) \right),$$

$$\mathcal{B}_{t+1}^i = \left(\left(\frac{1}{2}(\mathcal{B}_t^i + \mathcal{B}_t^{\sigma(i)}) \right) \circ \left(r_{t+1}^i x_{t+1}^i \right) \right),$$

$$\tilde{A}_{t+1}^i = \tilde{A}_t^i + \tilde{A}_t^{\sigma(i)}, \text{ and } \tilde{b}_{t+1}^i = \tilde{b}_t^i + \tilde{b}_t^{\sigma(i)}.$$

- (3) if the length $|\mathcal{A}_{t+1}^i|$ exceeds $t - \tau(t)$, the first element of \mathcal{A}_{t+1}^i is added to \tilde{A}_{t+1}^i and deleted from \mathcal{A}_{t+1}^i . \mathcal{B}_{t+1}^i and \tilde{b}_{t+1}^i are treated similarly.

In this way, each buffer remains of size at most $t - \tau(t)$, and contains only information gathered after time $\tau(t)$. The result is that, after t rounds of sharing, the current covariance matrices and b-vectors used by the algorithm to make decisions have the form:

$$\tilde{A}_t^i := I + \sum_{t'=1}^{\tau(t)} \sum_{i'=1}^{|V|} w_{i,t}^{i',t'} x_{t'}^{i'} x_{t'}^{i'\top},$$

$$\text{and } \tilde{b}_t^i := \sum_{t'=1}^{\tau(t)} \sum_{i'=1}^{|V|} w_{i,t}^{i',t'} r_{t'}^{i'} x_{t'}^{i'}.$$

where the weights $w_{i,t}^{i',t'}$ are random variables which are unknown to the algorithm. Importantly for our analysis, as

¹This can be achieved in a variety of ways.

²The \circ symbol denotes the concatenation operation on two ordered sets: if $x = (a, b, c)$ and $y = (d, e, f)$, then $x \circ y = (a, b, c, d, e, f)$, and $y \circ x = (d, e, f, a, b, c)$.

a result of the overlay protocol's uniformly random choice of σ , they are identically distributed (*i.d.*) for each fixed pair (t, t') , and $\sum_{i' \in V} w_{i,t}^{i',t'} = |V|$. If information sharing was perfect at each time step, then the current covariance matrix could be computed using all the information gathered by all the agents, and would be:

$$A_t := I + \sum_{i'=1}^{|V|} \sum_{t'=1}^t x_{t'}^{i'} (x_{t'}^{i'})^\top. \quad (1)$$

DCB algorithm The OFUL algorithm (Abbasi-Yadkori et al., 2011) is an improvement of the confidence ball algorithm from (Dani et al., 2008), which assumes that the confidence balls C_t can be characterised by A_t and b_t . In the DCB algorithm, each agent $i \in V$ maintains a confidence ball C_t^i for the unknown parameter θ as in the OFUL algorithm, but calculated from \tilde{A}_t^i and \tilde{b}_t^i . It then chooses its action, x_t^i , to satisfy $(x_t^i, \theta_t^i) = \arg \max_{(x, \theta) \in \mathcal{D}_t^i \times C_t^i} x^\top \theta$, and receives a reward r_t^i . Finally, it shares its information buffer according to the sharing protocol above. Pseudocode for DCB is given in Appendix A.1, and in Algorithm 1.

2.1. Results for DCB

Theorem 1. *Let $\tau(\cdot) : t \rightarrow 4 \log(|V|^{\frac{3}{2}} t)$. Then, with probability $1 - \delta$, the regret of DCB is bounded by*

$$\mathcal{R}_t \leq (N(\delta)|V| + \nu(|V|, d, t)) \|\theta\|_2 + 4e^2 (\beta(t) + 4R) \sqrt{|V|t \ln \left((1 + |V|t/d)^d \right)},$$

where $\nu(|V|, d, t) := (d+1)d^2(4|V| \ln(|V|^{\frac{3}{2}} t))^3$, $N(\delta) := \sqrt{3}/((1 - 2^{-\frac{1}{4}})\sqrt{\delta})$, and

$$\beta(t) := R \sqrt{\ln \left(\frac{(1 + |V|t/d)^d}{\delta} \right)} + \|\theta\|_2. \quad (2)$$

The term $\nu(t, |V|, d)$ describes the loss compared to the centralised algorithm due to the delay in using information, while $N(\delta)|V|$ describes the loss due to the incomplete mixing of the data across the network.

If the agents implement CB independently and do not share any information, which we call *CB-NoSharing*, then it follows from the results in (Abbasi-Yadkori et al., 2011), the equivalent regret bound would be

$$\mathcal{R}_t \leq |V|\beta(t) \sqrt{t \ln((1 + t/d)^d)} \quad (3)$$

Comparing Theorem 1 with (3) tells us that, after an initial ‘‘burn in’’ period, the gain in regret performance of DCB over *CB-NoSharing* is of order almost $\sqrt{|V|}$.

Corollary 2. *We can recover a bound in expectation from Theorem 1, by using the value $\delta = 1/\sqrt{|V|t}$:*

$$\begin{aligned} \mathbb{E}[\mathcal{R}_t] &\leq O(t^{\frac{1}{4}}) + \sqrt{|V|t}\|\theta\|_2 \\ &+ 4e^2 \left(R\sqrt{\ln\left((1+|V|t/d)^d\sqrt{|V|t}\right)} + \|\theta\|_2 + 4R \right) \\ &\quad \times \sqrt{|V|t \ln\left((1+|V|t/d)^d\right)}. \end{aligned}$$

This shows that DCB exhibits asymptotically optimal regret performance, up to log factors, in comparison with any algorithm that can share its information perfectly between agents at each round.

COMMUNICATION COMPLEXITY

If the agents communicate their information to each other at each round without a central server, then every agent would need to communicate their chosen action and reward to every other agent at each round, giving a communication cost of order $d|V|^2$ per-round. We call such an algorithm *CB-InstSharing*. Under the gossip protocol we propose each agent requires at most $O(\log_2(|V|t)d^2|V|)$ bits to be communicated per round. Therefore, a significant communication cost reduction is gained when $\log(|V|t)d \ll |V|$.

Using an epoch-based approach, as in (Szörényi et al., 2013), the per-round communication cost of the gossip protocol becomes $O(d^2|V|)$. This improves efficiency over any horizon, requiring only that $d \ll |V|$, and the proofs of the regret performance are simple modifications of those for DCB. However, in comparison with growing buffers this is only an issue after $O(\exp(|V|))$ number of rounds, and typically $|V|$ is large.

While the DCB has a clear communication advantage over *CB-InstSharing*, there are other potential approaches to this problem. For example, instead of randomised neighbour sharing one can use a deterministic protocol such as *Round-Robin* (RR), which can have the same low communication costs as DCB. However, the regret bound for RR suffers from a naturally larger delay in the network than DCB. Moreover, attempting to track potential doubling of data points when using a gossip protocol, instead of employing a delay, leads back to a communication cost of order $|V|^2$ per round. More detail is included in Appendix A.2.

PROOF OF THEOREM 1

In the analysis we show that the bias introduced by imperfect information sharing is mitigated by delaying the inclusion of the data in the estimation of the parameter θ . The proof builds on the analysis in (Abbasi-Yadkori et al., 2011). The emphasis here is to show how to handle the extra difficulty stemming from imperfect information sharing, which results in the influence of the various rewards

at the various peers being unbalanced and appearing with a random delay. Proofs of the Lemmas 3 and 4, and of Proposition 1 are crucial, but technical, and are deferred to Appendix A.3.

Step 1: Define modified confidence ellipsoids. First we need a version of the confidence ellipsoid theorem given in (Abbasi-Yadkori et al., 2011) that incorporates the bias introduced by the random weights:

Proposition 1. *Let $\delta > 0$, $\tilde{\theta}_t^i := (\tilde{A}_t^i)^{-1}\tilde{b}_t^i$, $W(\tau) := \max\{w_{i,t}^{i',t'} : t, t' \leq \tau, i, i' \in V\}$, and let*

$$\begin{aligned} C_t^i &:= \left\{ x \in \mathbb{R}^d : \|\tilde{\theta}_t^i - x\|_{\tilde{A}_t^i} \leq \|\theta\|_2 \right. \\ &\quad \left. + W(\tau(t))R\sqrt{2 \log\left(\det(\tilde{A}_t^i)^{\frac{1}{2}}/\delta\right)} \right\}. \end{aligned} \quad (4)$$

Then with probability $1 - \delta$, $\theta \in C_t^i$.

In the rest of the proof we assume that $\theta \in C_t^i$.

Step 2: Instantaneous regret decomposition. Denote by $(x_t^i, \theta_t^i) = \arg \max_{x \in D_t^i, y \in C_t^i} x^\top y$. Then we can decompose the instantaneous regret, following a classic argument (see the proof of Theorem 3 in (Abbasi-Yadkori et al., 2011)):

$$\begin{aligned} \rho_t^i &= (x_t^{i,*})^\top \theta - (x_t^i)^\top \theta \leq (x_t^i)^\top \theta_t^i - (x_t^i)^\top \theta \\ &= (x_t^i)^\top \left[(\theta_t^i - \tilde{\theta}_t^i) + (\tilde{\theta}_t^i - \theta) \right] \\ &\leq \|x_t^i\|_{(\tilde{A}_t^i)^{-1}} \left[\|\theta_t^i - \tilde{\theta}_t^i\|_{\tilde{A}_t^i} + \|\tilde{\theta}_t^i - \theta\|_{\tilde{A}_t^i} \right] \end{aligned} \quad (5)$$

Step 3: Control the bias. The norm differences inside the square brackets of the regret decomposition are bounded through (4) in terms of the matrices \tilde{A}_t^i . We would like, instead, to have the regret decomposition in terms of the matrix A_t (which is defined in (1)). To this end, we give some lemmas showing that using the matrices \tilde{A}_t^i is almost the same as using A_t . These lemmas involve elementary matrix analysis, but are crucial for understanding the impact of imperfect information sharing on the final regret bounds.

Step 3a: Control the bias coming from the weight imbalance.

Lemma 3 (Bound on the influence of general weights). *For all $i \in V$ and $t > 0$,*

$$\begin{aligned} \|x_t^i\|_{(\tilde{A}_t^i)^{-1}}^2 &\leq e^{\sum_{t'=1}^{\tau(t)} \sum_{i'=1}^{|V|} |w_{i,t}^{i',t'} - 1|} \|x_t^i\|_{(A_{\tau(t)})}^2, \\ \text{and } \det(\tilde{A}_t^i) &\leq e^{\sum_{t'=1}^{\tau(t)} \sum_{i'=1}^{|V|} |w_{i,t}^{i',t'} - 1|} \det(A_{\tau(t)}). \end{aligned}$$

Using Lemma 4 in (Szörényi et al., 2013), by exploiting the random weights are identically distributed (*i.d.*) for each

fixed pair (t, t') , and $\sum_{i' \in V} w_{i,t}^{i',t'} = |V|$ under our gossip protocol, we can control the random exponential constant in Lemma 3, and the upper bound $W(T)$ using the Chernoff-Hoeffding bound:

Lemma 4 (Bound on the influence of weights under our sharing protocol). *Fix some constants $0 < \delta_{t'} < 1$. Then with probability $1 - \sum_{t'=1}^{\tau(t)} \delta_{t'}$*

$$\sum_{i'=1}^{|V|} \sum_{t'=1}^{\tau(t)} \left| w_{i,t}^{i',t'} - 1 \right| \leq |V|^{\frac{3}{2}} \sum_{t'=1}^{\tau(t)} \left(2^{(t-t')} \delta_{t'} \right)^{-\frac{1}{2}},$$

and $W(T) \leq 1 + \max_{1 \leq t' \leq \tau(t)} \left\{ |V|^{\frac{3}{2}} \left(2^{(t-t')} \delta_{t'} \right)^{-\frac{1}{2}} \right\}$.

In particular, for any $\delta \in (0, 1)$, choosing $\delta_{t'} = \delta 2^{\frac{t'-t}{2}}$, with probability $1 - \delta/(|V|^3 t^2 (1 - 2^{-1/2}))$ we have

$$\sum_{i'=1}^{|V|} \sum_{t'=1}^{\tau(t)} \left| w_{i,t}^{i',t'} - 1 \right| \leq \frac{1}{(1 - 2^{-\frac{1}{4}}) t \sqrt{\delta}},$$

and $W(\tau(t)) \leq 1 + \frac{|V|^{\frac{3}{2}}}{t \sqrt{\delta}}$. (6)

Thus Lemma 3 and 4 give us control over the bias introduced by the imperfect information sharing. Combining them with Equations (4) and (5) we find that with probability $1 - \delta/(|V|^3 t^2 (1 - 2^{-1/2}))$:

$$\rho_t^i \leq 2e^{C(t)} \|x_t^i\|_{(A_{\tau(t)}^i)^{-1}} (1 + C(t)) \times \left[R \sqrt{2 \log \left(e^{C(t)} \det(A_{\tau(t)})^{\frac{1}{2}} \delta^{-1} \right)} + \|\theta\| \right] \quad (7)$$

where $C(t) := 1/(1 - 2^{-1/4}) t \sqrt{\delta}$

Step 3b: Control the bias coming from the delay. Next, we need to control the bias introduced from leaving out the last $4 \log(|V|^3 t^2)$ time steps from the confidence ball estimation calculation:

Proposition 2. *There can be at most*

$$\nu(k) := (4|V| \log(|V|^3 t^2 k))^3 (d+1) d(\text{tr}(A_0) + 1) \quad (8)$$

pairs $(i, k) \in 1, \dots, |V| \times \{1, \dots, t\}$ for which one of

$$\|x_k^i\|_{A_{\tau(k)}^{-1}}^2 \geq e \|x_k^i\|_{(A_{k-1} + \sum_{j=1}^{i-1} x_k^j (x_k^j)^\top)^{-1}},$$

or $\det(A_{\tau(k)}) \geq e \det \left(A_{k-1} + \sum_{j=1}^{i-1} x_k^j (x_k^j)^\top \right)$ holds.

Step 4: Choose constants and sum the simple regret.

Defining a constant

$$N(\delta) := \frac{1}{(1 - 2^{-\frac{1}{4}}) \sqrt{\delta}},$$

we have, for all $k \geq N(\delta)$, $C(k) \leq 1$, and so, by (7) with probability $1 - (|V|k)^{-2} \delta / (1 - 2^{-1/2})$

$$\rho_k^i \leq 2e \|x_k^i\|_{A_{\tau(k)}^{-1}} \times \left[2R \sqrt{2 \log \left(\frac{e \det(A_{\tau(k)})^{\frac{1}{2}}}{\delta} \right)} + \|\theta\|_2 \right]. \quad (9)$$

Now, first applying Cauchy-Schwarz, then step 3b from above together with (9), and finally Lemma 11 from (Abbasi-Yadkori et al., 2011) yields that, with probability $1 - (1 + \sum_{t=1}^{\infty} (|V|t)^{-2} / (1 - 2^{-1/2})) \delta \geq 1 - 3\delta$,

$$\begin{aligned} \mathcal{R}_t &\leq N(\delta) |V| \|\theta\|_2 + \left[|V| t \sum_{t'=N(\delta)}^t \sum_{i=1}^{|V|} (\rho_{t'}^i)^2 \right]^{\frac{1}{2}} \\ &\leq (N(\delta) |V| + \nu(|V|, d, t)) \|\theta\|_2 \\ &\quad + 4e^2 (\beta(t) + 2R) \left[|V| t \sum_{t'=1}^t \sum_{i=1}^M \|x_{t'}^i\|_{(A_{t'})^{-1}}^2 \right]^{\frac{1}{2}} \\ &\leq (N(\delta) |V| + \nu(|V|, d, t)) \|\theta\|_2 \\ &\quad + 4e^2 (\beta(t) + 2R) \sqrt{|V| t (2 \log(\det(A_t)))}, \end{aligned}$$

where $\beta(\cdot)$ is as defined in (2). Replacing δ with $\delta/3$ finishes the proof.

PROOF OF PROPOSITION 2

This proof forms the major innovation in the proof of Theorem 1. Let $(y_k)_{k \geq 1}$ be any sequence of vectors such that $\|y_k\|_2 \leq 1$ for all k , and let $B_n := B_0 + \sum_{k=1}^n y_k y_k^\top$, where B_0 is some positive definite matrix.

Lemma 5. *For all $t > 0$, and for any $c \in (0, 1)$, we have*

$$\left| \left\{ k \in \{1, 2, \dots\} : \|y_k\|_{B_{k-1}^{-1}}^2 > c \right\} \right| \leq (d+c) d (\text{tr}(B_0^{-1}) - c) / c^2,$$

Proof. We begin by showing that, for any $c \in (0, 1)$

$$\|y_k\|_{B_{k-1}^{-1}}^2 > c \quad (10)$$

can be true for only $2dc^{-3}$ different k .

Indeed, let us suppose that (10) is true for some k . Let $(e_i^{(k-1)})_{1 \leq i \leq d}$ be the orthonormal eigenbasis for B_{k-1} , and, therefore, also for B_{k-1}^{-1} , and write $y_k = \sum_{i=1}^d \alpha_i e_i$. Let, also, $(\lambda_i^{(k-1)})$ be the eigenvalues for B_{k-1} . Then,

$$\begin{aligned} c &< y_k^\top B_{k-1}^{-1} y_k = \sum_{i=1}^d \frac{\alpha_i^2}{\lambda_i^{(k-1)}} \leq \text{tr}(B_{k-1}^{-1}), \\ \implies \exists j \in \{1, \dots, d\} : \frac{\alpha_j^2}{\lambda_j^{(k-1)}}, \frac{1}{\lambda_j^{(k-1)}} &> \frac{c}{d}, \end{aligned}$$

where we have used that $\alpha_j^2 < 1$ for all i , since $\|y_k\|_2 < 1$. Now,

$$\begin{aligned} & \text{tr}(B_{k-1}^{-1}) - \text{tr}(B_k^{-1}) \\ &= \text{tr}(B_{k-1}^{-1}) - \text{tr}((B_{k-1} + y_k y_k^\top)^{-1}) \\ &> \text{tr}(B_{k-1}^{-1}) - \text{tr}((B_{k-1} + \alpha_j^2 e_j e_j^\top)^{-1}) \\ &= \frac{1}{\lambda_j^{(k-1)}} - \frac{1}{\lambda_j^{(k-1)} + \alpha_j^2} = \frac{\alpha_j^2}{\lambda_j^{(k-1)}(\lambda_j^{(k-1)} + \alpha_j^2)} \\ &> (d^2 c^{-2} + d c^{-1})^{-1} > \frac{c^2}{d(d+c)} \end{aligned}$$

So we have shown that (10) implies that

$$\text{tr}(B_{k-1}^{-1}) > c \text{ and } \text{tr}(B_{k-1}^{-1}) - \text{tr}(B_k^{-1}) > \frac{c^2}{d(d+c)}.$$

Since $\text{tr}(B_0^{-1}) \geq \text{tr}(B_{k-1}^{-1}) \geq \text{tr}(B_k^{-1}) \geq 0$ for all k , it follows that (10) can be true for at most $(d+c)d(\text{tr}(B_0^{-1}) - c)c^{-2}$ different k . \square

Now, using an argument similar to the proof of Lemma 3, for all $k < t$

$$\begin{aligned} \|y_{k+1}\|_{B_{\tau(k)}^{-1}} &\leq e^{\sum_{s=\tau(k)+1}^k \|y_{s+1}\|_{B_s^{-1}}} \|y_{k+1}\|_{B_k^{-1}}, \\ \text{and } \det(B_{\tau(t)}) &\leq e^{\sum_{k=\tau(t)+1}^t \|y_k\|_{B_k^{-1}}^2} \det(B_t). \end{aligned}$$

Therefore,

$$\begin{aligned} \|y_{k+1}\|_{B_{\tau(k)}^{-1}} &\geq c \|y_{k+1}\|_{B_k^{-1}} \text{ or } \det(B_{\tau(k)}) \geq c \det(B_k) \\ \implies \sum_{s=\tau(k)}^{k-1} \|y_{s+1}\|_{B_s^{-1}} &\geq \ln(c) \end{aligned}$$

However, according to Lemma 5, there can be at most

$$\nu(t) := \left(d + \frac{\ln(c)}{\Delta(t)}\right) d \left(\text{tr}(B_0^{-1}) - \frac{\ln(c)}{\Delta(t)}\right) \left(\frac{\Delta(t)}{\ln(c)}\right)^2$$

times $s \in \{1, \dots, t\}$, such that $\|y_{s+1}\|_{B_s^{-1}} \geq \ln(c)/\Delta(t)$, where $\Delta(t) := \max_{1 \leq k \leq t} \{k - \tau(k)\}$. Hence $\sum_{s=\tau(j)+1}^k \|y_{s+1}\|_{B_s^{-1}} \geq \ln(c)$ is true for at most $\Delta(t)\nu(|V|, d, t)$ indices $k \in \{1, \dots, t\}$.

Finally, we finish by setting $(y_k)_{k \geq 1} = \circ_{t \geq 1} (x_t^i)_{i=1}^{|V|}$.

3. Clustering and the DCCB Algorithm

We now incorporate distributed clustering into the DCB algorithm. The analysis of DCB forms the backbone of the analysis of DCCB.

DCCB Pruning Protocol In order to run DCCB, each agent i must maintain some local information buffers in addition to those used for DCB. These are:

Algorithm 1 Distributed Clustering Confidence Ball

Input: Size of network $|V|$, $\tau : t \rightarrow t - 4 \log_2 t$, α , λ
Initialization: $\forall i \in V$, set $\tilde{A}_0^i = I_d$, $\tilde{b}_0^i = \mathbf{0}$, $\mathcal{A}_0^i = \mathcal{B}_0^i = \emptyset$, and $V_0^i = V$.

for $t = 0, \dots, \infty$ **do**

Draw a random permutation σ of $\{1, \dots, V\}$ respecting the current local clusters

for $i = 1, \dots, |V|$ **do**

Receive action set \mathcal{D}_t^i and construct the confidence ball C_t^i using \tilde{A}_t^i and \tilde{b}_t^i

Choose action and receive reward:

Find $(x_{t+1}^i, *) = \arg \max_{(x, \bar{\theta}) \in \mathcal{D}_t^i \times C_t^i} x^\top \bar{\theta}$, and get reward r_{t+1}^i from context x_{t+1}^i .

Share and update information buffers:

if $\|\hat{\theta}_{local}^i - \hat{\theta}_{local}^j\| > c_\lambda^{thresh}(t)$

Update local cluster: $V_{t+1}^i = V_t^i \setminus \{\sigma(i)\}$, $V_{t+1}^{\sigma(i)} = V_t^{\sigma(i)} \setminus \{i\}$, and reset according to (13)

elseif $V_t^i = V_t^{\sigma(i)}$

Set $\mathcal{A}_{t+1}^i = \left(\frac{1}{2}(\mathcal{A}_t^i + \mathcal{A}_t^{\sigma(i)})\right) \circ (x_{t+1}^i (x_{t+1}^i)^\top)$

and $\mathcal{B}_{t+1}^i = \left(\frac{1}{2}(\mathcal{B}_t^i + \mathcal{B}_t^{\sigma(i)})\right) \circ (r_{t+1}^i x_{t+1}^i)$

else Update: Set $\mathcal{A}_{t+1}^i = \mathcal{A}_t^i \circ (x_{t+1}^i (x_{t+1}^i)^\top)$ and $\mathcal{B}_{t+1}^i = \mathcal{B}_t^i \circ (r_{t+1}^i x_{t+1}^i)$

endif

Update local estimator: $A_{local,t+1}^i = A_{local,t}^i + x_{t+1}^i (x_{t+1}^i)^\top$, $b_{local,t+1}^i = b_{local,t}^i + r_{t+1}^i x_{t+1}^i$, and $\hat{\theta}_{local,t+1}^i = \left(A_{local,t+1}^i\right)^{-1} b_{local,t+1}^i$

if $|\mathcal{A}_{t+1}^i| > t - \tau(t)$ set $\tilde{A}_{t+1}^i = \tilde{A}_t^i + \mathcal{A}_{t+1}^i(1)$, $\mathcal{A}_{t+1}^i = \mathcal{A}_{t+1}^i \setminus \mathcal{A}_{t+1}^i(1)$. Similarly for \mathcal{B}_{t+1}^i .

end for

end for

- (1) a local covariance matrix $A_{local}^i = A_{local,t}^i$, a local b-vector $b_{local}^i = b_{local,t}^i$,
- (2) and a local neighbour set V_t^i .

The local covariance matrix and b-vector are updated as if the agent was applying the generic (single agent) confidence ball algorithm: $A_{local,0}^i = A_0$, $b_{local,0}^i = 0$,

$$\begin{aligned} A_{local,t}^i &= x_t^i (x_t^i)^\top + A_{local,t-1}^i, \\ \text{and } b_{local,t}^i &= r_t^i x_t^i + b_{local,t-1}^i. \end{aligned}$$

DCCB Algorithm Each agent's local neighbour set V_t^i is initially set to V . At each time step t , agent i contacts one other agent, j , at random from V_t^i , and both decide whether they do or do not belong to the same cluster. To do this

they share local estimates, $\hat{\theta}_t^i = A_{local,t}^i{}^{-1}b_{local,t}^i$ and $\hat{\theta}_t^j = A_{local,t}^j{}^{-1}b_{local,t}^j$, of the unknown parameter of the bandit problem they are solving, and see if they are further apart than a threshold function $c = c_\lambda^{thresh}(t)$, so that if

$$\|\hat{\theta}_t^i - \hat{\theta}_t^j\|_2 \geq c_\lambda^{thresh}(t), \quad (11)$$

then $V_{t+1}^i = V_t^i \setminus \{j\}$ and $V_{t+1}^j = V_t^j \setminus \{i\}$. Here λ is a parameter of an extra assumption that is needed, as in (Gentile et al., 2014), about the process generating the context sets \mathcal{D}_t^i :

- (A) Each context set $\mathcal{D}_t^i = \{x_k\}_k$ is finite and contains *i.i.d.* random vectors such that for all, k , $\|x_k\| \leq 1$ and $\mathbb{E}(x_k x_k^\top)$ is full rank, with minimal eigenvalue $\lambda > 0$.

We define $c_\lambda^{thresh}(t)$, as in (Gentile et al., 2014), by

$$c_\lambda^{thresh}(t) := \frac{R\sqrt{2d \log(t) + 2 \log(2/\delta)} + 1}{\sqrt{1 + \max\{A_\lambda(t, \delta/(4d)), 0\}}} \quad (12)$$

where $A_\lambda(t, \delta) := \frac{\lambda t}{\delta} - 8 \log \frac{t+3}{\delta} - 2\sqrt{t \log \frac{t+3}{\delta}}$.

The DCCB algorithm is pretty much the same as the DCB algorithm, except that it also applies the pruning protocol described. In particular, each agent, i , when sharing its information with another, j , has three possible actions:

- (1) if (11) is not satisfied and $V_t^i = V_t^j$, then the agents share simply as in the DCB algorithm;
- (2) if (11) is not satisfied but $V_t^i \neq V_t^j$, then no sharing or pruning occurs.
- (3) if (11) is satisfied, then both agents remove each other from their neighbour sets and reset their buffers and active matrices so that

$$\begin{aligned} \mathcal{A}^i &= (0, 0, \dots, A_{local}^i), \mathcal{B}^i = (0, 0, \dots, b_{local}^i), \\ \text{and } \tilde{\mathcal{A}}^i &= A_{local}^i, \tilde{b}^i = b_{local}^i, \end{aligned} \quad (13)$$

and similarly for agent j .

It is proved in the theorem below, that under this sharing and pruning mechanism, in high probability after some finite time each agent i finds its true cluster, i.e. $V_t^i = U^k$. Moreover, since the algorithm resets to its local information each time a pruning occurs, once the true clusters have been identified, each cluster shares only information gathered within that cluster, thus avoiding introducing a bias by sharing information gathered from outside the cluster before the clustering has been identified. Full pseudo-code for the DCCB algorithm is given in Algorithm 1, and the differences with the DCB algorithm are highlighted in blue.




Figure 1. Here we plot the performance of DCCB in comparison to CLUB, CB-NoSharing and CB-InstSharing. The plots show the ratio of cumulative rewards achieved by the algorithms to the cumulative rewards achieved by the random algorithm.

3.1. Results for DCCB

Theorem 6. Assume that (A) holds, and let γ denote the smallest distance between the bandit parameters θ^k . Then there exists a constant $C = C(\gamma, |V|, \lambda, \delta)$, such that with probability $1 - \delta$ the total cumulative regret of cluster k when the agents employ DCCB is bounded by

$$\mathcal{R}_t \leq \left[\max \left\{ \sqrt{2}N(\delta), C + 4 \log_2(|V|^{\frac{3}{2}}C) \right\} |U^k| \right. \\ \left. + \nu(|U^k|, d, t) \right] \|\theta\|_2 \\ + 4e(\beta(t) + 3R) \sqrt{|U^k|t \ln \left((1 + |U^k|t/d)^d \right)},$$

where N and ν are as defined in Theorem 1, and $\beta(t) := R\sqrt{2 \ln \left((1 + |U^k|t/d)^d \right)} + \|\theta\|_2$.

The constant $C(\gamma, |V|, \lambda, \delta)$ is the time that you have to wait for the true clustering to have been identified,

The analysis follows the following scheme: When the true clusters have been correctly identified by all nodes, within each cluster the algorithm, and thus the analysis, reduces to the case of Section 2.1. We adapt results from (Gentile et al., 2014) to show how long it will be before the true clusters are identified, in high probability. The proof is deferred to Appendices A.4 and A.5.

4. Experiments and Discussion

Experiments We closely implemented the experimental setting and dataset construction principles used in (Li et al., 2016a;b), and for a detailed description of this we refer the reader to (Li et al., 2016a). We evaluated DCCB on three real-world datasets against its centralised counterpart CLUB, and against the benchmarks used therein, *CB-NoSharing*, and *CB-InstSharing*. The LastFM dataset comprises of 91 users, each of which appear at least 95 times. The Delicious dataset has 87 users, each of which appear at least 95 times. The MovieLens dataset contains 100 users, each of which appears at least 250 times. The performance was measured using the ratio of cumulative reward of each algorithm to that of the predictor which chooses a random action at each time step. This is plotted in in Figure 1. From the experimental results it is clear that DCCB performs comparably to CLUB in practice, and both outperform *CB-NoSharing*, and *CB-InstSharing*.

Relationship to existing literature There are several strands of research that are relevant and complimentary to this work. First, there is a large literature on single agent linear bandits, and other more, or less complicated bandit problem settings. There is already work on distributed approaches to multi-agent, multi-armed bandits, not least

(Szörényi et al., 2013) which examines ϵ -greedy strategies over a peer to peer network, and provided an initial inspiration for this current work. The paper (Kalathil et al., 2014) examines the extreme case when there is no communication channel across which the agents can communicate, and all communication must be performed through observation of action choices alone. Another approach to the multi-armed bandit case, (Nayyar et al., 2015), directly incorporates the communication cost into the regret.

Second, there are several recent advances regarding the state-of-the-art methods for clustering of bandits. The work (Li et al., 2016a) is a faster variant of (Gentile et al., 2014) which adopt the strategy of boosted training stage. In (Li et al., 2016b) the authors not only cluster the users, but also cluster the items under collaborative filtering case with a sharp regret analysis.

Finally, the paper (Tekin & van der Schaar, 2013) treats a setting similar to ours in which agents attempt to solve contextual bandit problems in a distributed setting. They present two algorithms, one of which is a distributed version of the approach taken in (Slivkins, 2014), and show that they achieve at least as good asymptotic regret performance in the distributed approach as the centralised algorithm achieves. However, rather than sharing information across a limited communication channel, they allow each agent only to ask another agent to choose their action for them. This difference in our settings is reflected worse regret bounds, which are of order $\Omega(T^{2/3})$ at best.

Discussion Our analysis is tailored to adapt proofs from (Abbasi-Yadkori et al., 2011) about generic confidence ball algorithms to a distributed setting. However many of the elements of these proofs, including Propositions 1 and 2 could be reused to provide similar asymptotic regret guarantees for the distributed versions of other bandit algorithms, e.g., the Thompson sampling algorithms, (Agrawal & Goyal, 2013; Kaufmann et al., 2012; Russo & Van Roy, 2014).

Both DCB and DCCB are synchronous algorithms. The work on distributed computation through gossip algorithms in (Boyd et al., 2006) could alleviate this issue. The current pruning algorithm for DCCB guarantees that techniques from (Szörényi et al., 2013) can be applied to our algorithms. However the results in (Boyd et al., 2006) are more powerful, and could be used even when the agents only identify a sub-network of the true clustering.

Furthermore, there are other existing interesting algorithms for performing clustering of bandits for recommender systems, such as COFIBA in (Li et al., 2016b). It would be interesting to understand how general the techniques applied here to CLUB are.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments. We would also like to thank Gergely Neu for very useful discussions. NK thanks the support from EPSRC Autonomous Intelligent Systems project EP/I011587. SL thanks the support from MIUR, QCRI-HBKU, Amazon Research Grant and Tsinghua University. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 306638.

References

- Abbasi-Yadkori, Yasin, Pál, Dávid, and Szepesvári, Csaba. Improved algorithms for linear stochastic bandits. In *NIPS*, pp. 2312–2320, 2011.
- Agrawal, Shipra and Goyal, Navin. Thompson sampling for contextual bandits with linear payoffs. In *ICML*, 2013.
- Boyd, Stephen, Ghosh, Arpita, Prabhakar, Balaji, and Shah, Devavrat. Randomized gossip algorithms. *IEEE/ACM Transactions on Networking (TON)*, 14(SI): 2508–2530, 2006.
- Dani, Varsha, Hayes, Thomas P, and Kakade, Sham M. Stochastic linear optimization under bandit feedback. In *COLT*, pp. 355–366, 2008.
- Gentile, Claudio, Li, Shuai, and Zappella, Giovanni. Online clustering of bandits. In *ICML*, 2014.
- Hao, Fei, Li, Shuai, Min, Geyong, Kim, Hee-Cheol, Yau, Stephen S, and Yang, Laurence T. An efficient approach to generating location-sensitive recommendations in ad-hoc social network environments. *IEEE Transactions on Services Computing*, 2015.
- Jelasity, M., Montresor, A., and Babaoglu, O. Gossip-based aggregation in large dynamic networks. *ACM Trans. on Computer Systems*, 23(3):219–252, August 2005.
- Jelasity, M., Voulgaris, S., Guerraoui, R., Kermarrec, A.-M., and van Steen, M. Gossip-based peer sampling. *ACM Transactions on Computer Systems*, 25(3):8, 2007.
- Kalathil, Dileep, Nayyar, Naumaan, and Jain, Rahul. Decentralized learning for multiplayer multiarmed bandits. *IEEE Transactions on Information Theory*, 60(4):2331–2345, 2014.
- Kaufmann, Emilie, Korda, Nathaniel, and Munos, Rémi. Thompson sampling: An asymptotically optimal finite-time analysis. In *Algorithmic Learning Theory*, pp. 199–213. Springer, 2012.
- Kempe, D., Dobra, A., and Gehrke, J. Gossip-based computation of aggregate information. In *Proc. 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS’03)*, pp. 482–491. IEEE Computer Society, 2003.
- Li, Lihong, Chu, Wei, Langford, John, and Schapire, Robert E. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pp. 661–670. ACM, 2010.
- Li, Shuai, Hao, Fei, Li, Mei, and Kim, Hee-Cheol. Medicine rating prediction and recommendation in mobile social networks. In *Proceedings of the International Conference on Grid and Pervasive Computing*, 2013.
- Li, Shuai, Gentile, Claudio, and Karatzoglou, Alexandros. Graph clustering bandits for recommendation. *CoRR:1605.00596*, 2016a.
- Li, Shuai, Karatzoglou, Alexandros, and Gentile, Claudio. Collaborative filtering bandits. In *The 39th SIGIR*, 2016b.
- Nayyar, Naumaan, Kalathil, Dileep, and Jain, Rahul. On regret-optimal learning in decentralized multi-player multi-armed bandits. *CoRR:1505.00553*, 2015.
- Russo, Daniel and Van Roy, Benjamin. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Slivkins, Aleksandrs. Contextual bandits with similarity information. *JMLR*, 2014.
- Szörényi, Balázs, Busa-Fekete, Róbert, Hegedűs, István, Ormándi, Róbert, Jelasity, Márk, and Kégl, Balázs. Gossip-based distributed stochastic bandit algorithms. In *ICML*, pp. 19–27, 2013.
- Tekin, Cem and van der Schaar, Mihaela. Distributed online learning via cooperative contextual bandits. *IEEE Trans. Signal Processing*, 2013.
- Xiao, L., Boyd, S., and Kim, S.-J. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing*, 67(1):33–46, January 2007.

A. Supplementary Material

A.1. Pseudocode for the generic CB algorithm and the DCB algorithm

Algorithm 2 Confidence Ball

Initialization: Set $A_0 = I$ and $b_0 = 0$.
for $t = 0, \dots, \infty$ **do**
 Receive action set \mathcal{D}_t
 Construct the confidence ball C_t using A_t and b_t
Choose action and receive reward:
 Find $(x_t, *) = \arg \max_{(x, \tilde{\theta}) \in \mathcal{D}_t \times C_t} x^\top \tilde{\theta}$
 Get reward r_t^i from context x_t^i
 Update $A_{t+1} = A_t + x_t x_t^\top$ and $b_{t+1} = b_t + r_t x_t$
end for

Algorithm 3 Distributed Confidence Ball

Input: Network V of agents, the function $\tau : t \rightarrow t - 4 \log_2(|V|^{\frac{3}{2}} t)$.
Initialization: For each i , set $\tilde{A}_0^i = I_d$ and $\tilde{b}_0^i = \mathbf{0}$, and the buffers $\mathcal{A}_0^i = \emptyset$ and $\mathcal{B}_0^i = \emptyset$.
for $t = 0, \dots, \infty$ **do**
 Draw a random permutation σ of $\{1, \dots, |V|\}$
for each agent $i \in V$ **do**
 Receive action set \mathcal{D}_t^i and construct the confidence ball C_t^i using \tilde{A}_t^i and \tilde{b}_t^i
Choose action and receive reward:
 Find $(x_{t+1}^i, *) = \arg \max_{(x, \tilde{\theta}) \in \mathcal{D}_t^i \times C_t^i} x^\top \tilde{\theta}$
 Get reward r_{t+1}^i from context x_{t+1}^i .
Share and update information buffers:
 Set $\mathcal{A}_{t+1}^i = \left(\frac{1}{2}(\mathcal{A}_t^i + \mathcal{A}_t^{\sigma(i)})\right) \circ (x_{t+1}^i (x_{t+1}^i)^\top)$ and $\mathcal{B}_{t+1}^i = \left(\frac{1}{2}(\mathcal{B}_t^i + \mathcal{B}_t^{\sigma(i)})\right) \circ (r_{t+1}^i x_{t+1}^i)$
if $|\mathcal{A}_{t+1}^i| > t - \tau(t)$ **set** $\tilde{A}_{t+1}^i = \tilde{A}_t^i + \mathcal{A}_{t+1}^i(1)$ and $\mathcal{A}_{t+1}^i = \mathcal{A}_{t+1}^i \setminus \mathcal{A}_{t+1}^i(1)$. Similary for \mathcal{B}_{t+1}^i .
end for
end for

A.2. More on Communication Complexity

First, recall that if the agents want to communicate their information to each other at each round without a central server, then every agent would need to communicate their chosen action and reward to every other agent at each round, giving a communication cost of $O(d|V|^2)$ bits per-round. Under DCB each agent requires at most $O(\log_2(|V|t)d^2|V|)$ bits to be communicated per round. Therefore, a significant communication cost reduction is gained when $\log(|V|t)d \ll |V|$.

Recall also that using an epoch-based approach, as in (Szörényi et al., 2013), we reduce the per-round communication cost of the gossip-based approach to $O(d^2|V|)$. This makes the algorithm more efficient over any time horizon, requiring only that $d \ll |V|$, and the proofs of the regret performance are simple modifications of the proofs for DCB. In comparison with growing buffers this is only an issue after $O(\exp(|V|))$ number of rounds, and typically $|V|$ is large. This is why we choose to exhibit the growing-buffer approach in this current work.

Instead of relying on the combination of the diffusion and a delay to handle the potential doubling of data points under the randomised gossip protocol, we could attempt to keep track which observations have been shared with which agents, and thus simply stop the doubling from occurring. However, the per-round communication complexity of this is at least quadratic in $|V|$, whereas our approach is linear. The reason for the former is that in order to be efficient, any agent j , when sending information to an agent i , needs to know for each k which are the latest observations gathered by agent k that agent i already knows about. The communication cost of this is of order $|V|$. Since every agent shares information with somebody in each round, this gives per round communication complexity of order $|V|^2$ in the network.

A simple, alternative approach to the gossip protocol is a *Round-Robin* (RR) protocol, in which each agent passes the information it has gathered in previous rounds to the next agent in a pre-defined permutation. Implementing a RR protocol

leads to the agents performing a distributed version of the *CB-InstSharing* algorithm, but with a delay that is of size at least linear in $|V|$, rather than the logarithmic dependence on this quantity that a gossip protocol achieves. Indeed, at any time, each agent will be lacking $|V|(|V| - 1)/2$ observations. Using this observation, a cumulative regret bound can be achieved using Proposition 2 which arrives at the same asymptotic dependence on $|V|$ as our gossip protocol, but with an additive constant that is worse by a multiplicative factor of $|V|$. This makes a difference to the performance of the network when $|V|$ is very large. Moreover, RR protocols do not offer the simple generalisability and robustness that gossip protocols offer.

Note that the pruning protocol for DCCB only requires sharing the estimated θ -vectors between agents, and adds at most $O(d|V|)$ to the communication cost of the algorithm. Hence the per-round communication cost of DCCB remains $O(\log_2(|V|t)d^2|V|)$.

Algorithm	Regret Bound	Per-Round Communication Complexity
CB- <i>NoSharing</i>	$O(V \sqrt{t})$	0
CB- <i>InstSharing</i>	$O(\sqrt{ V t})$	$O(d V ^2)$
DCB	$O(\sqrt{ V t})$	$O(\log_2(V t)d^2 V)$
DCCB	$O(\sqrt{ U^k t})$	$O(\log_2(V t)d^2 V)$

Figure 2. This table gives a summary of theoretical results for the multi-agent linear bandit problem. Note that CB with no sharing cannot benefit from the fact that all the agents are solving the same bandit problem, while CB with instant sharing has a large communication-cost dependency on the size of the network. DCB successfully achieves near-optimal regret performance, while simultaneously reducing communication complexity by an order of magnitude in the size of the network. Moreover, DCCB generalises this regret performance at not extra cost in the order of the communication complexity.

A.3. Proofs of Intermediary Results for DCB

Proof of Proposition 1. This follows the proof of Theorem 2 in (Abbasi-Yadkori et al., 2011), substituting appropriately weighted quantities.

For ease of presentation, we define the shorthand

$$\tilde{X} := (\sqrt{w_1}y_1, \dots, \sqrt{w_n}y_n) \text{ and } \tilde{\eta} = (\sqrt{w_1}\eta_1, \dots, \sqrt{w_n}\eta_n)^\top,$$

where the y_i are vectors with norm less than 1, the η_i are R -subgaussian, zero mean, random variables, and the w_i are positive real numbers. Then, given samples $(\sqrt{w_i}y_i, \sqrt{w_i}(\theta y_i + \eta_i))$, the maximum likelihood estimate of θ is

$$\begin{aligned} \tilde{\theta} &:= (\tilde{X}\tilde{X}^\top + I)^{-1}\tilde{X}(\tilde{X}^\top\theta + \tilde{\eta}) \\ &= (\tilde{X}\tilde{X}^\top + I)^{-1}\tilde{X}\tilde{\eta} + (\tilde{X}\tilde{X}^\top + I)^{-1}(\tilde{X}\tilde{X}^\top + I)\theta - (\tilde{X}\tilde{X}^\top + I)^{-1}\theta \\ &= (\tilde{X}\tilde{X}^\top + I)^{-1}\tilde{X}\tilde{\eta} + \theta - (\tilde{X}\tilde{X}^\top + I)^{-1}\theta \end{aligned}$$

So by Cauchy-Schwarz, we have, for any vector x ,

$$x^\top(\tilde{\theta} - \theta) = \langle x, \tilde{X}\tilde{\eta} \rangle_{(\tilde{X}\tilde{X}^\top + I)^{-1}} - \langle x, \theta \rangle_{(\tilde{X}\tilde{X}^\top + I)^{-1}} \quad (14)$$

$$\leq \|x\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}} \left(\|\tilde{X}\tilde{\eta}\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}} + \|\theta\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}} \right) \quad (15)$$

Now from Theorem 1 of (Abbasi-Yadkori et al., 2011), we know that with probability $1 - \delta$

$$\|\tilde{X}\tilde{\eta}\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}}^2 \leq W^2 R^2 2 \log \sqrt{\frac{\det(\tilde{X}\tilde{X}^\top + I)}{\delta^2}}.$$

where $W = \max_{i=1, \dots, n} w_i$. So, setting $x = (\tilde{X}\tilde{X}^\top + I)^{-1}(\tilde{\theta} - \theta)$, we obtain that with probability $1 - \delta$

$$\|\tilde{\theta} - \theta\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}} \leq WR \left(2 \log \sqrt{\frac{\det(\tilde{X}\tilde{X}^\top + I)}{\delta^2}} \right)^{\frac{1}{2}} + \|\theta\|_2$$

since ³

$$\begin{aligned} \|x\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}} \|\theta\|_{(\tilde{X}\tilde{X}^\top + I)^{-1}} &\leq \|x\|_2 \lambda_{\min}^{-1}(\tilde{X}\tilde{X}^\top + I) \|\theta\|_2 \lambda_{\min}^{-1}(\tilde{X}\tilde{X}^\top + I) \\ &\leq \|x\|_2 \|\theta\|_2. \end{aligned}$$

Conditioned on the values of the weights, the statement of Proposition 1 now follows by substituting appropriate quantities above, and taking the probability over the distribution of the subGaussian random rewards. However, since this statement holds uniformly for any values of the weights, it holds also when the probability is taken over the distribution of the weights. \square

Proof of Lemma 3. Recall that \tilde{A}_t^i is constructed from the contexts chosen from the first $\tau(t)$ rounds, across all the agents. Let i' and t' be arbitrary indices in V and $\{1, \dots, \tau(t)\}$, respectively.

(i) We have

$$\begin{aligned} \det(\tilde{A}_t^i) &= \det\left(\tilde{A}_t^i - (w_{i,t}^{i',t'} - 1) x_{t'}^{i'} (x_{t'}^{i'})^\top + (w_{i,t}^{i',t'} - 1) x_{t'}^{i'} (x_{t'}^{i'})^\top\right) \\ &= \det\left(\tilde{A}_t^i - (w_{i,t}^{i',t'} - 1) x_{t'}^{i'} (x_{t'}^{i'})^\top\right) \\ &\quad \cdot \left(1 + (w_{i,t}^{i',t'} - 1) \|x_{t'}^{i'}\|_{(\tilde{A}_t^i - (w_{i,t}^{i',t'} - 1) x_{t'}^{i'} (x_{t'}^{i'})^\top)^{-1}}\right) \end{aligned}$$

The second equality follows using the identity $\det(I + cB^{1/2}xx^\top B^{1/2}) = (1 + c\|x\|_B)$, for any matrix B , vector x , and scalar c . Now, we repeat this process for all $i' \in V$ and $t' \in \{1, \dots, \tau(t)\}$ as follows. Let $(t_1, i_1), \dots, (t_{|V|\tau(t)}, i_{|V|\tau(t)})$ be an arbitrary enumeration of $V \times \{1, \dots, \tau(t)\}$, let $B_0 = \tilde{A}_t^i$, and $B_s = B_{s-1} - (w_{i,t}^{i_s, t_s} - 1) x_{t_s}^{i_s} (x_{t_s}^{i_s})^\top$ for $s = 1, \dots, |V|\tau(t)$. Then $B_{|V|\tau(t)} = A_{\tau(t)}$, and by the calculation above we have

$$\begin{aligned} \det(\tilde{A}_t^i) &= \det(A_{\tau(t)}) \prod_{s=1}^{|V|\tau(t)} \left(1 + (w_{i,t}^{i_s, t_s} - 1) \|x_{t_s}^{i_s}\|_{(B_s)^{-1}}\right) \\ &\leq \det(A_{\tau(t)}) \exp\left(\sum_{s=1}^{|V|\tau(t)} (w_{i,t}^{i_s, t_s} - 1) \|x_{t_s}^{i_s}\|_{(B_s)^{-1}}\right) \\ &\leq \exp\left(\sum_{t'=1}^{\tau(t)} \sum_{i'=1}^{|V|} |w_{i,t}^{i', t'} - 1|\right) \det(A_T) \end{aligned}$$

(ii) Note that for vectors x, y and a matrix B , by the Sherman-Morrison Lemma, and Cauchy-Schwarz inequality we have that:

$$\begin{aligned} x^\top (B + yy^\top)^{-1} x &= x^\top B^{-1} x - \frac{x^\top B^{-1} y y^\top B^{-1} x}{1 + y^\top B^{-1} y} \geq x^\top B^{-1} x - \frac{x^\top B^{-1} x y^\top B^{-1} y}{1 + y^\top B^{-1} y} \\ &= x^\top B^{-1} x (1 + y^\top B^{-1} y)^{-1} \end{aligned} \quad (16)$$

Taking

$$B = \left(\tilde{A}_t^i - (w_{i,t}^{i', t'} - 1) x_{t'}^{i'} (x_{t'}^{i'})^\top\right) \text{ and } y = \sqrt{w_{i,t}^{i', t'} - 1} x_{t'}^{i'},$$

and using that $y^\top B^{-1} y \leq \lambda_{\min}(B)^{-1} y^\top y$, by construction, we have that, for any $t' \in \{1, \dots, \tau(t)\}$ and $i' \in V$,

$$x^\top (\tilde{A}_t^i)^{-1} x \geq x^\top \left(\tilde{A}_t^i - (w_{i,t}^{i', t'} - 1) x_{t'}^{i'} (x_{t'}^{i'})^\top\right)^{-1} x (1 + |w_{i,t}^{i', t'} - 1|)^{-1}.$$

Performing this for each $i' \in V$ and $t' \in \{1, \dots, \tau(t)\}$, taking the exponential of the logarithm and using that $\log(1 + a) \leq a$ like in the first part finishes the proof. \square

³ $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue of its argument.

A.4. Proof of Theorem 6

Throughout the proof let i denote the index of some arbitrary but fixed agent, and let k denote the index of its cluster.

Step 1: Show the true clustering is obtained in finite time. First we prove that with probability $1 - \delta$, the number of times agents in different clusters share information is bounded. Consider the statements

$$\forall i, i' \in V, \forall t, \left(\|\hat{\theta}_{local,t}^i - \hat{\theta}_{local,t}^{i'}\| > c_\lambda^{thresh}(t) \right) \implies i' \notin U^k \quad (17)$$

and,

$$\forall t \geq C(\gamma, \lambda, \delta) = c_\lambda^{thresh}^{-1} \left(\frac{\gamma}{2} \right), i' \notin U^k, \|\hat{\theta}_{local,t}^i - \hat{\theta}_{local,t}^{i'}\| > c_\lambda^{thresh}(t). \quad (18)$$

where c_λ^{thresh} and A_λ are as defined in the main paper. Lemma 4 from (Gentile et al., 2014) proves that these two statements hold under the assumptions of the theorem with probability $1 - \delta/2$.

Let i be an agent in cluster U^k . Suppose that (17) and (18) hold. Then we know that at time $t = \lceil C(\gamma, \lambda, \delta) \rceil$, $U^k \subset V_t^i$. Moreover, since the sharing protocol chooses an agent uniformly at random from V_t^i independently from the history before time t , it follows that the time until $V_t^i = U^k$ can be upper bounded by a constant $C = C(|V|, \delta)$ with probability $1 - \delta/2$. So it follows that there exists a constant $C = C(|V|, \gamma, \lambda, \delta)$ such that the event

$$E := \{(17) \text{ and } (18) \text{ hold, and } (t \geq C(|V|, \gamma, \lambda, \delta) \implies V_t^i = U^k)\}$$

holds with probability $1 - \delta$.

Step 2: Consider the properties of the weights after clustering. On the event E , we know that each cluster will be performing the algorithm DCB within its own cluster for all $t > C(\gamma, |V|)$. Therefore, we would like to directly apply the analysis from the proof of Theorem 1 from this point. In order to do this we need to show that the weights, $w_{i,t}^{i',t'}$, have the same properties after time $C = C(\gamma, |V|, \lambda, \delta)$ that are required for the proof of Theorem 1.

Lemma 7. *Suppose that agent i is in cluster U^k . Then, on the event E ,*

- (i) for all $t > C(|V|, \gamma, \lambda, \delta)$ and $i' \in V \setminus U^k$, $w_{i,t}^{i',t'} = 0$;
- (ii) for all $t' \geq C(|V|, \gamma, \lambda, \delta)$ and $i' \in U^k$, $\sum_{i \in U^k} w_{i,C(|V|,\gamma)}^{i',t'} = |U^k|$;
- (iii) for all $t \geq t' \geq C(|V|, \gamma, \lambda, \delta)$ and $i' \in U^k$, the weights $w_{i,t}^{i',t'}$, $i \in U^k$, are i.i.d..

Proof. See Appendix A.5. □

We must deal also with what happens to the information gathered before the cluster has completely discovered itself. To this end, note that we can write, supposing that $\tau(t) \geq C(|V|, \gamma, \lambda, \delta)$,

$$\tilde{A}_t^i := \sum_{i' \in U^k} \frac{w_{i,t}^{i',C}}{|U^k|} \tilde{A}_C^{i'} + \sum_{t'=C+1}^{\tau(t)} \sum_{i' \in U^k} w_{i,t}^{i',t'} x_{t'}^{i'} \left(x_{t'}^{i'} \right)^\top. \quad (19)$$

Armed with this observation we show that the fact that sharing within the appropriate cluster only begins properly after time $C = C(|V|, \gamma, \lambda, \delta)$ the influence of the bias is unchanged:

Lemma 8 (Bound on the influence of general weights). *On the event E , for all $i \in V$ and t such that $T(t) \geq C(|V|, \gamma, \lambda, \delta)$,*

- (i) $\det(\tilde{A}_t^i) \leq \exp \left(\sum_{t'=C}^{\tau(t)} \sum_{i' \in U^k} |w_{i,t}^{i',t'} - 1| \right) \det(A_{\tau(t)}^k)$,
- (ii) and $\|x_t^i\|_{(\tilde{A}_t^i)^{-1}}^2 \leq \exp \left(\sum_{t'=C}^{\tau(t)} \sum_{i' \in U^k} |w_{i,t}^{i',t'} - 1| \right) \|x_t^i\|_{(A_{\tau(t)}^k)^{-1}}^2$.

Proof. See Appendix A.5. \square

The final property of the weights required to prove Theorem 1 is that their variance is diminishing geometrically with each iteration. For the analysis of DCB this is provided by Lemma 4 of (Szörényi et al., 2013), and, using Lemma 7, we can prove the same result for the weights after time $C = C(|V|, \gamma, \lambda, \delta)$:

Lemma 9. *Suppose that agent i is in cluster U^k . Then, on the event E , for all $t \geq C = C(|V|, \gamma, \lambda, \delta)$ and $t' < t$, we have*

$$\mathbb{E} \left((w_{i,t}^{j,t'} - 1)^2 \right) \leq \frac{|U^k|}{2^{t-\max\{t', C\}}}.$$

Proof. Given the properties proved in Lemma 7, the proof is identical to the proof of Lemma 4 of (Szörényi et al., 2013). \square

Step 3: Apply the results from the analysis of DCB. We can now apply the same argument as in Theorem 1 to bound the regret after time $C = C(\gamma, |V|, \lambda, \delta)$. The regret before this time we simply upper bound by $|U^k|C(|V|, \gamma, \lambda, \delta)\|\theta\|$. We include the modified sections below as needed.

Using Lemma 9, we can control the random exponential constant in Lemma 8, and the upper bound $W(T)$:

Lemma 10 (Bound in the influence of weights under our sharing protocol). *Assume that $t \geq C(\gamma, |V|, \lambda, \delta)$. Then on the event E , for some constants $0 < \delta_{t'} < 1$, with probability $1 - \sum_{t'=1}^{\tau(t)} \delta_{t'}$*

$$\sum_{t'=C}^{\tau(t)} \sum_{i' \in U^k} |w_{i,t}^{i',t'} - 1| \leq |U^k|^{\frac{3}{2}} \sum_{t'=C}^{\tau(t)} \sqrt{\frac{2^{-(t-\max\{t', C\})}}{\delta_{t'}}},$$

and $W(\tau(t)) \leq 1 + \max_{C \leq t' \leq \tau(t)} \left\{ |U^k|^{\frac{3}{2}} \sqrt{\frac{2^{-(t-\max\{t', C\})}}{\delta_{t'}}} \right\}.$

In particular, for any $1 > \delta > 0$, choosing $\delta_{t'} = \delta 2^{-(t-\max\{t', C\})/2}$, and $\tau(t) = t - c_1 \log_2 c_2 t$ we conclude that with probability $1 - (c_2 t)^{-c_1/2} \delta / (1 - 2^{-1/2})$, for any $t > C + c_1 \log_2(c_2 C)$,

$$\sum_{i' \in U^k} \sum_{t'=C}^{\tau(t)} |w_{i,t}^{i',t'} - 1| \leq \frac{|U^k|^{\frac{3}{2}} (c_2 t)^{-\frac{c_1}{4}}}{(1 - 2^{-\frac{1}{4}}) \sqrt{\delta}}, \text{ and } W(\tau(t)) \leq 1 + \frac{|U^k|^{\frac{3}{2}} (c_2 t)^{-\frac{c_1}{4}}}{\sqrt{\delta}}. \quad (20)$$

Thus lemmas 8 and 10 give us control over the bias introduced by the imperfect information sharing. Applying lemmas 8 and 10, we find that with probability $1 - (c_2 t)^{-c_1/2} \delta / (1 - 2^{-1/2})$:

$$\rho_t^i \leq 2 \exp \left(\frac{|U^k|^{\frac{3}{2}}}{(1 - 2^{-\frac{1}{4}}) c_2^{\frac{c_1}{4}} t^{\frac{c_1}{4}} \sqrt{\delta}} \right) \|x_t^i\| (A_{\tau(t)}^i)^{-1} \cdot \left[\left(1 + \frac{|U^k|^{\frac{3}{2}}}{(1 - 2^{-\frac{1}{4}}) c_2^{\frac{c_1}{4}} t^{\frac{c_1}{4}} \sqrt{\delta}} \right) \left[R \sqrt{2 \log \left(\exp \left(\frac{|U^k|^{\frac{3}{2}}}{(1 - 2^{-\frac{1}{4}}) c_2^{\frac{c_1}{4}} t^{\frac{c_1}{4}} \sqrt{\delta}} \right) \frac{\det(A_{\tau(t)}^i)^{\frac{1}{2}}}{\delta} \right) + \|\theta\|} \right] \right]. \quad (21)$$

Step 4: Choose constants and sum the simple regret. Choosing again $c_1 = 4$, $c_2 = |V|^{\frac{3}{2}}$, and setting $N_\delta = \frac{1}{(1-2^{-\frac{1}{4}})\sqrt{\delta}}$, we have on the event E , for all $t \geq \max\{N_\delta, C + 4 \log_2(|V|^{\frac{3}{2}} C)\}$, with probability $1 - (|V|t)^{-2} \delta / (1 - 2^{-1/2})$

$$\rho_t^i \leq 4e \|x_t^i\| (A_{t-1}^k + \sum_{i'=1}^{i-1} x_{i'}^i (x_{i'}^i)^\top)^{-1} (\beta(t) + R\sqrt{2}),$$

where $\beta(\cdot)$ is as defined in the theorem statement. Now applying Cauchy-Schwarz, and Lemma 11 from (Abbasi-Yadkori et al., 2011) yields that on the event E , with probability $1 - (1 + \sum_{t=1}^{\infty} (|V|t)^{-2} / (1 - 2^{-1/2})) \delta \geq 1 - 3\delta$,

$$\mathcal{R}_t \leq \left(\max\{N_\delta, C + 4 \log_2(|V|^{\frac{3}{2}} C)\} + 2(4|V|d \log(|V|t))^3 \right) \|\theta\|_2$$

$$+ 4e \left(\beta(t) + R\sqrt{2} \right) \sqrt{|U^k|t \left(2 \log \left(\det \left(A_t^k \right) \right) \right)}.$$

Replacing δ with $\delta/6$, and combining this result with Step 1 finishes the proof.

A.5. Proofs of Intermediary Results for DCCB

Proof of Lemma 7. Recall that whenever the pruning procedure cuts an edge, both agents reset their buffers to their local information, scaled by the size of their current neighbour sets. (It does not make a difference practically whether or not they scale their buffers, as this effect is washed out in the computation of the confidence bounds and the local estimates. However, it is convenient to assume that they do so for the analysis.) Furthermore, according to the pruning procedure, no agent will share information with another agent that does not have the same local neighbour set.

On the event E , there is a time for each agent, i , before time $C = C(\gamma, |V|, \lambda\delta)$ when the agent resets its information to their local information, and their local neighbour set becomes their local cluster, i.e. $V_t^i = U^k$. After this time, this agent will only share information with other agents that have also set their local neighbour set to their local cluster. This proves the statement of part (i).

Furthermore, since on event E , after agent i has identified its local neighbour set, i.e. when $V_t^i = U^k$, the agent only shares with members of U^k , the statements of parts (ii) and (iii) hold by construction of the sharing protocol. \square

Proof of Lemma 8. The result follows the proof of Lemma 3. For the the iterations until time $C = C(\gamma, |V|, \lambda\delta)$ is reached, we apply the argument there. For the final step we require two further inequalities.

First, to finish the proof of part (i) we note that,

$$\begin{aligned} \det \left((A_T^k - A_C^k) + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C(\gamma,|V|)}}{|U^k|} \tilde{A}_C^{i'} \right) &= \det \left(A_T^k + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C} - 1}{|U^k|} \tilde{A}_C^{i'} \right) \\ &= \det(A_T^k) \det \left(I + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C} - 1}{|U^k|} A_T^{k-\frac{1}{2}} \tilde{A}_C^{i'} A_T^{k-\frac{1}{2}} \right) \\ &\leq \det(A_T^k) \det \left(I + \left[\sum_{i' \in U^k} |w_{i,t}^{i',C} - 1| \right] A_T^{k-\frac{1}{2}} \sum_{i' \in U^k} \frac{\tilde{A}_C^{i'}}{|U^k|} A_T^{k-\frac{1}{2}} \right) \leq \det(A_T^k) \left(1 + \sum_{i' \in U^k} |w_{i,t}^{i',C} - 1| \right). \end{aligned}$$

For the first equality we have used that $|U^k|A_C^k = \sum_{i' \in U^k} \tilde{A}_C^{i'}$; for the first inequality we have used a property of positive definite matrices; for the second inequality we have used that 1 upper bounds the eigenvalues of $A_T^{k-1/2} A_C^k A_T^{k-1/2}$.

Second, to finish the proof of part (ii), we note that, for any vector x ,

$$\begin{aligned} x^\top \left(A_{\tau(t)}^k + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C} - 1}{|U^k|} \tilde{A}_C^{i'} \right)^{-1} x \\ &= \left(A_{\tau(t)}^k \right)^{-\frac{1}{2}} x^\top \left(I + \sum_{i' \in U^k} \frac{w_{i,t}^{i',C} - 1}{|U^k|} A_{\tau(t)}^{k-\frac{1}{2}} \tilde{A}_C^{i'} A_{\tau(t)}^{k-\frac{1}{2}} \right)^{-1} \left(A_{\tau(t)}^k \right)^{-\frac{1}{2}} x \\ &\geq \left(A_{\tau(t)}^k \right)^{-\frac{1}{2}} x^\top \left(I + \sum_{i' \in U^k} \frac{|w_{i,t}^{i',C} - 1|}{|U^k|} A_{\tau(t)}^{k-\frac{1}{2}} \tilde{A}_C^{i'} A_{\tau(t)}^{k-\frac{1}{2}} \right)^{-1} \left(A_{\tau(t)}^k \right)^{-\frac{1}{2}} x \\ &\geq \left(1 + \sum_{i' \in U^k} |w_{i,t}^{i',C} - 1| \right)^{-1} x^\top A_{\tau(t)}^k^{-1} x. \end{aligned}$$

The first inequality here follows from a property of positive definite matrices, and the other steps follow similarly to those in the inequality that finished part (i) of the proof. \square