Appendices

A. Proof of Theorem 1

Proof. Theorem 1 can be proved based on the definitions of monotonicity and submodularity. Note that from Assumption 1, for any seed set $\mathcal{S} \in \mathcal{C}$, any seed node $u \in \mathcal{S}$, and any target node $v \in \mathcal{V}$, we have $F(\{u\}, v) \leq F(\mathcal{S}, v)$, which implies that

$$f(\mathcal{S}, v, p^*) = \max_{u \in \mathcal{S}} F(\{u\}, v) \leq F(\mathcal{S}, v),$$

hence

$$f(\mathcal{S}, p^*) = \sum_{v \in \mathcal{V}} f(\mathcal{S}, v, p^*) \leq \sum_{v \in \mathcal{V}} F(\mathcal{S}, v) = F(\mathcal{S}).$$

This proves the first part of Theorem 1.

We now prove the second part of the theorem. First, note that from the first part, we have

$$f(\widetilde{\mathcal{S}}, p^*) \leq F(\widetilde{\mathcal{S}}) \leq F(\mathcal{S}^*),$$

where the first inequality follows from the first part of this theorem, and the second inequality follows from the definition of $\mathcal{S}^*$. Thus, we have $\rho \leq 1$. To prove that $\rho \geq 1/K$, we assume that $\mathcal{S} = \{u_1, u_2, \ldots, u_K\}$, and define $\mathcal{S}_k = \{u_1, u_2, \ldots, u_k\}$ for $k = 1, 2, \ldots, K$. Thus, for any $\mathcal{S} \subseteq \mathcal{V}$ with $|\mathcal{S}| = K$, we have

$$\begin{aligned} F(\mathcal{S}) &= F(\mathcal{S}_1) + \sum_{k=1}^{K-1} [F(\mathcal{S}_{k+1}) - F(\mathcal{S}_k)] \\ &\leq \sum_{k=1}^{K} F(\{u_k\}) = \sum_{k=1}^{K} \sum_{v \in \mathcal{V}} F(\{u_k\}, v) \\ &\leq \sum_{v \in \mathcal{V}} K \max_{u \in \mathcal{S}} F(\{u\}, v) = K \sum_{v \in \mathcal{V}} f(\mathcal{S}, v, p^*) = K f(\mathcal{S}, p^*), \end{aligned}$$

where the first inequality follows from the submodularity of $F(\cdot)$. Thus we have

$$F(\mathcal{S}^*) \leq K f(\mathcal{S}^*, p^*) \leq K f(\widetilde{\mathcal{S}}, p^*),$$

where the second inequality follows from the definition of $\widetilde{\mathcal{S}}$. This implies that $\rho \geq 1/K$. □

B. Proof of Theorem 2

We start by defining some useful notations. We use $\mathcal{H}_t$ to denote the "history" by the end of time t. For any node pair $(u, v) \in \mathcal{V} \times \mathcal{V}$ and any time t, we define the upper confidence bound (UCB) $U_t(u, v)$ and the lower confidence bound (LCB) $L_t(u, v)$ respectively as

$$\begin{aligned} U_t(u, v) &= \mathrm{Proj}_{[0,1]} \left(\langle \widehat{\boldsymbol{\theta}}_{u,t-1}, \boldsymbol{x}_v \rangle + c \sqrt{\boldsymbol{x}_v^T \Sigma_{u,t-1}^{-1} \boldsymbol{x}_v} \right) \\ L_t(u, v) &= \mathrm{Proj}_{[0,1]} \left(\langle \widehat{\boldsymbol{\theta}}_{u,t-1}, \boldsymbol{x}_v \rangle - c \sqrt{\boldsymbol{x}_v^T \Sigma_{u,t-1}^{-1} \boldsymbol{x}_v} \right) \end{aligned} \tag{8}$$

Notice that U_t is the same as the UCB estimate $\overline{p}$ defined in Algorithm 1. Moreover, we define the "good event" $\mathcal{F}$ as

$$\mathcal{F} = \left\{ |x_v^T (\widehat{\boldsymbol{\theta}}_{u,t-1} - \boldsymbol{\theta}_u^*)| \leq c \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}, \, \forall u, v \in \mathcal{V}, \, \forall t \leq T \right\}, \tag{9}$$

and the "bad event" $\overline{\mathcal{F}}$ as the complement of $\mathcal{F}$.

B.1. Regret Decomposition

Recall that the realized scaled regret at time t is $R_t^{\rho\alpha} = F(S^*) - \frac{1}{\rho\alpha}F(\mathcal{S}_t)$, thus we have

$$R_t^{\rho\alpha} = F(S^*) - \frac{1}{\rho\alpha}F(\mathcal{S}_t) \overset{(a)}{=} \frac{1}{\rho}f(\widetilde{\mathcal{S}}, p^*) - \frac{1}{\rho\alpha}F(\mathcal{S}_t) \overset{(b)}{\leq} \frac{1}{\rho}f(\widetilde{\mathcal{S}}, p^*) - \frac{1}{\rho\alpha}f(\mathcal{S}_t, p^*), \tag{10}$$

where equality (a) follows from the definition of ρ (i.e. ρ is defined as $\rho = f(\widetilde{\mathcal{S}}, p^*)/F(S^*)$), and inequality (b) follows from $f(\mathcal{S}_t, p^*) \leq F(\mathcal{S}_t)$ (see Theorem 1). Thus, we have

$$\begin{aligned}
R^{\rho\alpha}(T) &= \mathbb{E}\left[\sum_{t=1}^{T} R_t^{\rho\alpha}\right] \\
&\leq \frac{1}{\rho}\mathbb{E}\left\{\sum_{t=1}^{T}\left[f(\widetilde{\mathcal{S}}, p^*) - f(\mathcal{S}_t, p^*)/\alpha\right]\right\} \\
&= \frac{P(\mathcal{F})}{\rho}\mathbb{E}\left\{\sum_{t=1}^{T}\left[f(\widetilde{\mathcal{S}}, p^*) - f(\mathcal{S}_t, p^*)/\alpha\right]\middle|\mathcal{F}\right\} + \frac{P(\overline{\mathcal{F}})}{\rho}\mathbb{E}\left\{\sum_{t=1}^{T}\left[f(\widetilde{\mathcal{S}}, p^*) - f(\mathcal{S}_t, p^*)/\alpha\right]\middle|\overline{\mathcal{F}}\right\} \\
&\leq \frac{1}{\rho}\mathbb{E}\left\{\sum_{t=1}^{T}\left[f(\widetilde{\mathcal{S}}, p^*) - f(\mathcal{S}_t, p^*)/\alpha\right]\middle|\mathcal{F}\right\} + \frac{P(\overline{\mathcal{F}})}{\rho}nT,
\end{aligned} \tag{11}$$

where the last inequality follows from the naive bounds $P(\mathcal{F}) \leq 1$ and $f(\widetilde{\mathcal{S}}, p^*) - f(\mathcal{S}_t, p^*)/\alpha \leq n$. Notice that under "good" event $\mathcal{F}$, we have

$$L_t(u, v) \leq p_{uv}^* = x_v^T \theta_u^* \leq U_t(u, v) \tag{12}$$

for all node pair (u, v) and for all time $t \leq T$. Thus, we have $f(\mathcal{S}, L_t) \leq f(\mathcal{S}, p^*) \leq f(\mathcal{S}, U_t)$ for all $\mathcal{S}$ and $t \leq T$ under event $\mathcal{F}$. So under event $\mathcal{F}$, we have

$$f(\mathcal{S}_t, L_t) \overset{(a)}{\leq} f(\mathcal{S}_t, p^*) \overset{(b)}{\leq} f(\widetilde{\mathcal{S}}, p^*) \overset{(c)}{\leq} f(\widetilde{\mathcal{S}}, U_t) \leq \max_{\mathcal{S} \in \mathcal{C}} f(\mathcal{S}, U_t) \overset{(d)}{\leq} \frac{1}{\alpha}f(\mathcal{S}_t, U_t)$$

for all $t \leq T$, where inequalities (a) and (c) follow from (12), inequality (b) follows from $\widetilde{\mathcal{S}} \in \arg\max_{\mathcal{S} \in \mathcal{C}} f(\mathcal{S}, p^*)$, and inequality (d) follows from the fact that `ORACLE` is an α-approximation algorithm. Specifically, the fact that `ORACLE` is an α-approximation algorithm implies that $f(\mathcal{S}_t, U_t) \geq \alpha \max_{\mathcal{S} \in \mathcal{C}} f(\mathcal{S}, U_t)$.

Consequently, under event $\mathcal{F}$, we have

$$\begin{aligned}
f(\widetilde{\mathcal{S}}, p^*) - \frac{1}{\alpha}f(\mathcal{S}_t, p^*) &\leq \frac{1}{\alpha}f(\mathcal{S}_t, U_t) - \frac{1}{\alpha}f(\mathcal{S}_t, L_t) \\
&= \frac{1}{\alpha}\sum_{v \in \mathcal{V}}\left[\max_{u \in \mathcal{S}_t} U_t(u, v) - \max_{u \in \mathcal{S}_t} L_t(u, v)\right] \\
&\leq \frac{1}{\alpha}\sum_{v \in \mathcal{V}}\sum_{u \in \mathcal{S}_t}\left[U_t(u, v) - L_t(u, v)\right] \\
&\leq \sum_{v \in \mathcal{V}}\sum_{u \in \mathcal{S}_t}\frac{2c}{\alpha}\sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}.
\end{aligned} \tag{13}$$

So we have

$$R^{\rho\alpha}(T) \leq \frac{2c}{\rho\alpha}\mathbb{E}\left\{\sum_{t=1}^{T}\sum_{u \in \mathcal{S}_t}\sum_{v \in \mathcal{V}}\sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}\middle|\mathcal{F}\right\} + \frac{P(\overline{\mathcal{F}})}{\rho}nT. \tag{14}$$

In the remainder of this section, we will provide a worst-case bound on $\sum_{t=1}^{T}\sum_{u \in \mathcal{S}_t}\sum_{v \in \mathcal{V}}\sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}$ (see Appendix B.2) and a bound on the probability of "bad event" $P(\overline{\mathcal{F}})$ (see Appendix B.3).

B.2. Worst-Case Bound on $\sum_{t=1}^{T} \sum_{u \in \mathcal{S}_t} \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}$

Notice that

$$\sum_{t=1}^{T} \sum_{u \in \mathcal{S}_t} \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v} = \sum_{u \in \mathcal{V}} \sum_{t=1}^{T} \mathbf{1}\left[u \in \mathcal{S}_t\right] \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}$$

For each $u \in \mathcal{V}$, we define $K_u = \sum_{t=1}^{T} \mathbf{1}\left[u \in \mathcal{S}_t\right]$ as the number of times at which u is chosen as a source node, then we have the following lemma:

Lemma 1. *For all $u \in \mathcal{V}$, we have*

$$\sum_{t=1}^{T} \mathbf{1}\left[u \in \mathcal{S}_t\right] \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v} \leq \sqrt{n K_u} \sqrt{\frac{dn \log\left(1 + \frac{nK_u}{d\lambda\sigma^2}\right)}{\lambda \log\left(1 + \frac{1}{\lambda\sigma^2}\right)}}.$$

Moreover, when $X = I$, we have

$$\sum_{t=1}^{T} \mathbf{1}\left[u \in \mathcal{S}_t\right] \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v} \leq \sqrt{n K_u} \sqrt{\frac{n \log\left(1 + \frac{K_u}{\lambda\sigma^2}\right)}{\lambda \log\left(1 + \frac{1}{\lambda\sigma^2}\right)}}.$$

Proof. To simplify the exposition, we use Σ_t to denote $\Sigma_{u,t}$, and define $z_{t,v} = \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}$ for all $t \leq T$ and all $v \in \mathcal{V}$. Recall that

$$\Sigma_t = \Sigma_{t-1} + \frac{\mathbf{1}\left[u \in \mathcal{S}_t\right]}{\sigma^2} X X^T = \Sigma_{t-1} + \frac{\mathbf{1}\left[u \in \mathcal{S}_t\right]}{\sigma^2} \sum_{v \in \mathcal{V}} x_v x_v^T.$$

Note that if $u \notin \mathcal{S}_t$, $\Sigma_t = \Sigma_{t-1}$. If $u \in \mathcal{S}_t$, then for any $v \in \mathcal{V}$, we have

$$\begin{aligned}
\det\left[\Sigma_t\right] &\geq \det\left[\Sigma_{t-1} + \frac{1}{\sigma^2} x_v x_v^T\right] \\
&= \det\left[\Sigma_{t-1}^{\frac{1}{2}} \left(I + \frac{1}{\sigma^2} \Sigma_{t-1}^{-\frac{1}{2}} x_v x_v^T \Sigma_{t-1}^{-\frac{1}{2}}\right) \Sigma_{t-1}^{\frac{1}{2}}\right] \\
&= \det\left[\Sigma_{t-1}\right] \det\left[I + \frac{1}{\sigma^2} \Sigma_{t-1}^{-\frac{1}{2}} x_v x_v^T \Sigma_{t-1}^{-\frac{1}{2}}\right] \\
&= \det\left[\Sigma_{t-1}\right] \left(1 + \frac{1}{\sigma^2} x_v^T \Sigma_{t-1}^{-1} x_v\right) = \det\left[\Sigma_{t-1}\right] \left(1 + \frac{z_{t-1,v}^2}{\sigma^2}\right).
\end{aligned}$$

Hence, we have

$$\det\left[\Sigma_t\right]^n \geq \det\left[\Sigma_{t-1}\right]^n \prod_{v \in \mathcal{V}} \left(1 + \frac{z_{t-1,v}^2}{\sigma^2}\right). \tag{15}$$

Note that the above inequality holds for any X. However, if $X = I$, then all Σ_t's are diagonal and we have

$$\det\left[\Sigma_t\right] = \det\left[\Sigma_{t-1}\right] \prod_{v \in \mathcal{V}} \left(1 + \frac{z_{t-1,v}^2}{\sigma^2}\right). \tag{16}$$

As we will show later, this leads to a tighter regret bound in the tabular ($X = I$) case.

Let's continue our analysis for general X. The above results imply that

$$n \log\left(\det\left[\Sigma_t\right]\right) \geq n \log\left(\det\left[\Sigma_{t-1}\right]\right) + \mathbf{1}\left(u \in \mathcal{S}_t\right) \sum_{v \in \mathcal{V}} \log\left(1 + \frac{z_{t-1,v}^2}{\sigma^2}\right)$$

and hence

$$\begin{aligned} n\log\left(\det\left[\Sigma_T\right]\right) &\geq n\log\left(\det\left[\Sigma_0\right]\right) + \sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}\log\left(1+\frac{z_{t-1,v}^2}{\sigma^2}\right) \\ &= nd\log(\lambda) + \sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}\log\left(1+\frac{z_{t-1,v}^2}{\sigma^2}\right). \end{aligned} \tag{17}$$

On the other hand, we have that

$$\begin{aligned} \mathrm{Tr}\left[\Sigma_T\right] &= \mathrm{Tr}\left[\Sigma_0 + \sum_{t=1}^{T}\frac{\mathbf{1}\left[u\in\mathcal{S}_t\right]}{\sigma^2}\sum_{v\in\mathcal{V}}x_v x_v^T\right] \\ &= \mathrm{Tr}\left[\Sigma_0\right] + \sum_{t=1}^{T}\frac{\mathbf{1}\left[u\in\mathcal{S}_t\right]}{\sigma^2}\sum_{v\in\mathcal{V}}\mathrm{Tr}\left[x_v x_v^T\right] \\ &= \lambda d + \sum_{t=1}^{T}\frac{\mathbf{1}\left[u\in\mathcal{S}_t\right]}{\sigma^2}\sum_{v\in\mathcal{V}}\|x_v\|^2 \leq \lambda d + \frac{nK_u}{\sigma^2}, \end{aligned} \tag{18}$$

where the last inequality follows from the assumption that $\|x_v\| \leq 1$ and the definition of K_u. From the trace-determinant inequality, we have $\frac{1}{d}\mathrm{Tr}\left[\Sigma_T\right] \geq \det\left[\Sigma_T\right]^{\frac{1}{d}}$. Thus, we have

$$dn\log\left(\lambda+\frac{nK_u}{d\sigma^2}\right) \geq dn\log\left(\frac{1}{d}\mathrm{Tr}\left[\Sigma_T\right]\right) \geq n\log\left(\det\left[\Sigma_T\right]\right) \geq dn\log(\lambda) + \sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}\log\left(1+\frac{z_{t-1,v}^2}{\sigma^2}\right).$$

That is

$$\sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}\log\left(1+\frac{z_{t-1,v}^2}{\sigma^2}\right) \leq dn\log\left(1+\frac{nK_u}{d\lambda\sigma^2}\right)$$

Notice that $z_{t-1,v}^2 = x_v^T\Sigma_{t-1}^{-1}x_v \leq x_v^T\Sigma_0^{-1}x_v = \frac{\|x_v\|^2}{\lambda} \leq \frac{1}{\lambda}$. Moreover, for all $y\in[0,1/\lambda]$, we have $\log\left(1+\frac{y}{\sigma^2}\right) \geq \lambda\log\left(1+\frac{1}{\lambda\sigma^2}\right)y$ based on the concavity of $\log(\cdot)$. Thus, we have

$$\lambda\log\left(1+\frac{1}{\lambda\sigma^2}\right)\sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}z_{t-1,v}^2 \leq dn\log\left(1+\frac{nK_u}{d\lambda\sigma^2}\right).$$

Finally, from Cauchy-Schwarz inequality, we have that

$$\sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}z_{t-1,v} \leq \sqrt{nK_u}\sqrt{\sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}z_{t-1,v}^2}.$$

Combining the above results, we have

$$\sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}z_{t-1,v} \leq \sqrt{nK_u}\sqrt{\frac{dn\log\left(1+\frac{nK_u}{d\lambda\sigma^2}\right)}{\lambda\log\left(1+\frac{1}{\lambda\sigma^2}\right)}}. \tag{19}$$

This concludes the proof for general X. Based on (16), the analysis for the tabular ($X = I$) case is similar, and we omit the detailed analysis. In the tabular case, we have

$$\sum_{t=1}^{T}\mathbf{1}\left(u\in\mathcal{S}_t\right)\sum_{v\in\mathcal{V}}z_{t-1,v} \leq \sqrt{nK_u}\sqrt{\frac{n\log\left(1+\frac{K_u}{\lambda\sigma^2}\right)}{\lambda\log\left(1+\frac{1}{\lambda\sigma^2}\right)}}. \tag{20}$$

□

We now develop a worst-case bound. Notice that for general X, we have

$$\begin{aligned}\sum_{u \in \mathcal{V}} \sum_{t=1}^{T} \mathbf{1}\left[u \in \mathcal{S}_t\right] \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v} &\leq \sum_{u \in \mathcal{V}} \sqrt{n K_u} \sqrt{\frac{dn \log\left(1+\frac{nK_u}{d\lambda\sigma^2}\right)}{\lambda \log\left(1+\frac{1}{\lambda\sigma^2}\right)}} \\ &\overset{(a)}{\leq} n \sqrt{\frac{d \log\left(1+\frac{nT}{d\lambda\sigma^2}\right)}{\lambda \log\left(1+\frac{1}{\lambda\sigma^2}\right)}} \sum_{u \in \mathcal{V}} \sqrt{K_u} \\ &\overset{(b)}{\leq} n \sqrt{\frac{d \log\left(1+\frac{nT}{d\lambda\sigma^2}\right)}{\lambda \log\left(1+\frac{1}{\lambda\sigma^2}\right)}} \sqrt{n} \sqrt{\sum_{u \in \mathcal{V}} K_u} \\ &\overset{(c)}{=} n^{\frac{3}{2}} \sqrt{\frac{dKT \log\left(1+\frac{nT}{d\lambda\sigma^2}\right)}{\lambda \log\left(1+\frac{1}{\lambda\sigma^2}\right)}},\end{aligned} \tag{21}$$

where inequality (a) follows from the naive bound $K_u \leq T$, inequality (b) follows from Cauchy-Schwarz inequality, and equality (c) follows from $\sum_{u \in \mathcal{V}} K_u = KT$. Similarly, for the special case with $X = I$, we have

$$\sum_{u \in \mathcal{V}} \sum_{t=1}^{T} \mathbf{1}\left[u \in \mathcal{S}_t\right] \sum_{v \in \mathcal{V}} \sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v} \leq \sum_{u \in \mathcal{V}} \sqrt{n K_u} \sqrt{\frac{n \log\left(1+\frac{K_u}{\lambda\sigma^2}\right)}{\lambda \log\left(1+\frac{1}{\lambda\sigma^2}\right)}} \leq n^{\frac{3}{2}} \sqrt{\frac{KT \log\left(1+\frac{T}{\lambda\sigma^2}\right)}{\lambda \log\left(1+\frac{1}{\lambda\sigma^2}\right)}}. \tag{22}$$

This concludes the derivation of a worst-case bound.

B.3. Bound on $P\left(\overline{\mathcal{F}}\right)$

We now derive a bound on $P\left(\overline{\mathcal{F}}\right)$ based on the "Self-Normalized Bound for Matrix-Valued Martingales" developed in Theorem 3 (see Theorem 3). Before proceeding, we define $\mathcal{F}_u$ for all $u \in \mathcal{V}$ as

$$\mathcal{F}_u = \left\{ |x_v^T(\widehat{\boldsymbol{\theta}}_{u,t-1} - \boldsymbol{\theta}_u^*)| \leq c\sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}, \; \forall v \in \mathcal{V}, \; \forall t \leq T \right\}, \tag{23}$$

and the $\overline{\mathcal{F}}_u$ as the complement of $\mathcal{F}_u$. Note that by definition, $\overline{\mathcal{F}} = \bigcup_{u \in \mathcal{V}} \overline{\mathcal{F}}_u$. Hence, we first develop a bound on $P\left(\overline{\mathcal{F}}_u\right)$, then we develop a bound on $P\left(\overline{\mathcal{F}}\right)$ based on union bound.

Lemma 2. *For all $u \in \mathcal{V}$, all $\sigma, \lambda > 0$, all $\delta \in (0,1)$, and all*

$$c \geq \frac{1}{\sigma} \sqrt{dn \log\left(1 + \frac{nT}{\sigma^2 \lambda d}\right) + 2 \log\left(\frac{1}{\delta}\right)} + \sqrt{\lambda} \|\boldsymbol{\theta}_u^*\|_2$$

we have $P\left(\overline{\mathcal{F}}_u\right) \leq \delta$.

Proof. To simplify the expositions, we omit the subscript u in this proof. For instance, we use θ^*, Σ_t, $\mathbf{y}_t$ and $\mathbf{b}_t$ to respectively denote θ_u^*, $\Sigma_{u,t}$, $\mathbf{y}_{u,t}$ and $\mathbf{b}_{u,t}$. We also use $\mathcal{H}_t$ to denote the "history" by the end of time t, and hence $\{\mathcal{H}_t\}_{t=0}^{\infty}$ is a filtration. Notice that U_t is $\mathcal{H}_{t-1}$-adaptive, and hence $\mathcal{S}_t$ and $\mathbf{1}\left[u \in \mathcal{S}_t\right]$ are also $\mathcal{H}_{t-1}$-adaptive. We define

$$\eta_t = \begin{cases} \mathbf{y}_t - X^T \theta^* & \text{if } u \in \mathcal{S}_t \\ 0 & \text{otherwise} \end{cases} \in \Re^n \quad \text{and} \quad X_t = \begin{cases} X & \text{if } u \in \mathcal{S}_t \\ 0 & \text{otherwise} \end{cases} \in \Re^{d \times n} \tag{24}$$

Note that X_t is $\mathcal{H}_{t-1}$-adaptive, and η_t is $\mathcal{H}_t$-adaptive. Moreover, $\|\eta_t\|_\infty \leq 1$ always holds, and $\mathbb{E}\left[\eta_t | \mathcal{H}_{t-1}\right] = 0$. To simplify the expositions, we further define $\mathbf{y}_t = 0$ for all t s.t. $u \notin \mathcal{S}_t$. Note that with this definition, we have

$\eta_t = \mathbf{y}_t - X_t^T \theta^*$ for all t. We further define

$$\begin{aligned}\overline{V}_t &= n\sigma^2 \Sigma_t = n\sigma^2 \lambda I + n \sum_{s=1}^t X_s X_s^T \\ \overline{S}_t &= \sum_{s=1}^t X_s \eta_s = \sum_{s=1}^t X_s \left[\mathbf{y}_s - X_s^T \theta^*\right] = \mathbf{b}_t - \sigma^2 \left[\Sigma_t - \lambda I\right] \theta^* \end{aligned} \tag{25}$$

Thus, we have $\Sigma_t \widehat{\boldsymbol{\theta}}_t = \sigma^{-2} \mathbf{b}_t = \sigma^{-2} \overline{S}_t + \left[\Sigma_t - \lambda I\right] \boldsymbol{\theta}^*$, which implies

$$\widehat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^* = \Sigma_t^{-1} \left[\sigma^{-2} \overline{S}_t - \lambda \boldsymbol{\theta}^*\right]. \tag{26}$$

Consequently, for any $v \in \mathcal{V}$, we have

$$\begin{aligned}\left| x_v^T \left(\widehat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*\right)\right| &= \left| x_v^T \Sigma_t^{-1} \left[\sigma^{-2} \overline{S}_t - \lambda \boldsymbol{\theta}^*\right]\right| \leq \sqrt{x_v^T \Sigma_t^{-1} x_v} \|\sigma^{-2} \overline{S}_t - \lambda \boldsymbol{\theta}^*\|_{\Sigma_t^{-1}} \\ &\leq \sqrt{x_v^T \Sigma_t^{-1} x_v} \left[\|\sigma^{-2} \overline{S}_t\|_{\Sigma_t^{-1}} + \|\lambda \boldsymbol{\theta}^*\|_{\Sigma_t^{-1}}\right], \end{aligned} \tag{27}$$

where the first inequality follows from Cauchy-Schwarz inequality and the second inequality follows from triangular inequality. Note that $\|\lambda \boldsymbol{\theta}^*\|_{\Sigma_t^{-1}} = \lambda \|\boldsymbol{\theta}^*\|_{\Sigma_t^{-1}} \leq \lambda \|\boldsymbol{\theta}^*\|_{\Sigma_0^{-1}} = \sqrt{\lambda} \|\boldsymbol{\theta}^*\|_2$. On the other hand, since $\Sigma_t^{-1} = n\sigma^2 \overline{V}_t^{-1}$, we have $\|\sigma^{-2} \overline{S}_t\|_{\Sigma_t^{-1}} = \frac{\sqrt{n}}{\sigma} \|\overline{S}_t\|_{\overline{V}_t^{-1}}$. Thus, we have

$$\left| x_v^T \left(\widehat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}^*\right)\right| \leq \sqrt{x_v^T \Sigma_t^{-1} x_v} \left[\frac{\sqrt{n}}{\sigma} \|\overline{S}_t\|_{\overline{V}_t^{-1}} + \sqrt{\lambda} \|\boldsymbol{\theta}^*\|_2\right]. \tag{28}$$

From Theorem 3, we know with probability at least $1 - \delta$, for all $t \leq T$, we have

$$\|S_t\|_{\overline{V}_t^{-1}}^2 \leq 2 \log \left(\frac{\det\left(\overline{V}_t\right)^{1/2} \det\left(V\right)^{-1/2}}{\delta}\right) \leq 2 \log \left(\frac{\det\left(\overline{V}_T\right)^{1/2} \det\left(V\right)^{-1/2}}{\delta}\right),$$

where $V = n\sigma^2 \lambda I$. Note that from the trace-determinant inequality, we have

$$\det\left[\overline{V}_T\right]^{\frac{1}{d}} \leq \frac{\operatorname{Tr}\left[\overline{V}_T\right]}{d} \leq \frac{n\sigma^2 \lambda d + n^2 T}{d},$$

where the last inequality follows from $\operatorname{Tr}\left[X_t X_t^T\right] \leq n$ for all t. Note that $\det\left[V\right] = \left[n\sigma^2 \lambda\right]^d$, with a little bit algebra, we have

$$\|S_t\|_{\overline{V}_t^{-1}} \leq \sqrt{d \log \left(1 + \frac{nT}{\sigma^2 \lambda d}\right) + 2 \log \left(\frac{1}{\delta}\right)} \quad \forall t \leq T$$

with probability at least $1 - \delta$. Thus, if

$$c \geq \frac{1}{\sigma} \sqrt{dn \log \left(1 + \frac{nT}{\sigma^2 \lambda d}\right) + 2 \log \left(\frac{1}{\delta}\right)} + \sqrt{\lambda} \|\boldsymbol{\theta}^*\|_2,$$

then $\mathcal{F}_u$ holds with probability at least $1 - \delta$. This concludes the proof of this lemma. □

Hence, from the union bound, we have the following lemma:

Lemma 3. *For all $\sigma, \lambda > 0$, all $\delta \in (0, 1)$, and all*

$$c \geq \frac{1}{\sigma} \sqrt{dn \log \left(1 + \frac{nT}{\sigma^2 \lambda d}\right) + 2 \log \left(\frac{n}{\delta}\right)} + \sqrt{\lambda} \max_{u \in \mathcal{V}} \|\boldsymbol{\theta}_u^*\|_2 \tag{29}$$

we have $P\left(\overline{\mathcal{F}}\right) \leq \delta$.

Proof. This lemma follows directly from the union bound. Note that for all c satisfying Equation 29, we have $P\left(\overline{\mathcal{F}}_u\right) \leq \frac{\delta}{n}$ for all $u \in \mathcal{V}$, which implies $P\left(\overline{\mathcal{F}}\right) = P\left(\bigcup_{u \in \mathcal{V}} \overline{\mathcal{F}}_u\right) \leq \sum_{u \in \mathcal{V}} P\left(\overline{\mathcal{F}}_u\right) \leq \delta$. □

B.4. Conclude the Proof

Note that if we choose

$$c \geq \frac{1}{\sigma}\sqrt{dn \log\left(1 + \frac{nT}{\sigma^2 \lambda d}\right) + 2\log\left(n^2 T\right)} + \sqrt{\lambda} \max_{u \in \mathcal{V}} \|\boldsymbol{\theta}_u^*\|_2, \tag{30}$$

we have $P\left(\overline{\mathcal{F}}\right) \leq \frac{1}{nT}$. Hence for general X, we have

$$\begin{aligned} R^{\rho\alpha}(T) &\leq \frac{2c}{\rho\alpha}\mathbb{E}\left\{\sum_{t=1}^{T}\sum_{u \in \mathcal{S}_t}\sum_{v \in \mathcal{V}}\sqrt{x_v^T \Sigma_{u,t-1}^{-1} x_v}\middle|\mathcal{F}\right\} + \frac{1}{\rho} \\ &\leq \frac{2c}{\rho\alpha} n^{\frac{3}{2}} \sqrt{\frac{dKT \log\left(1 + \frac{nT}{d\lambda\sigma^2}\right)}{\lambda \log\left(1 + \frac{1}{\lambda\sigma^2}\right)}} + \frac{1}{\rho}. \end{aligned} \tag{31}$$

Note that with $c = \frac{1}{\sigma}\sqrt{dn \log\left(1 + \frac{nT}{\sigma^2 \lambda d}\right) + 2\log\left(n^2 T\right)} + \sqrt{\lambda} \max_{u \in \mathcal{V}} \|\boldsymbol{\theta}_u^*\|_2$, this regret bound is $\widetilde{O}\left(\frac{n^2 d\sqrt{KT}}{\rho\alpha}\right)$. Similarly, for the special case $X = I$, we have

$$R^{\rho\alpha}(T) \leq \frac{2c}{\rho\alpha} n^{\frac{3}{2}} \sqrt{\frac{KT \log\left(1 + \frac{T}{\lambda\sigma^2}\right)}{\lambda \log\left(1 + \frac{1}{\lambda\sigma^2}\right)}} + \frac{1}{\rho}. \tag{32}$$

Note that with $c = \frac{n}{\sigma}\sqrt{\log\left(1 + \frac{T}{\sigma^2 \lambda}\right) + 2\log\left(n^2 T\right)} + \sqrt{\lambda} \max_{u \in \mathcal{V}} \|\boldsymbol{\theta}_u^*\|_2 \leq \frac{n}{\sigma}\sqrt{\log\left(1 + \frac{T}{\sigma^2 \lambda}\right) + 2\log\left(n^2 T\right)} + \sqrt{\lambda n}$, this regret bound is $\widetilde{O}\left(\frac{n^{\frac{5}{2}}\sqrt{KT}}{\rho\alpha}\right)$.

C. Self-Normalized Bound for Matrix-Valued Martingales

In this section, we derive a "self-normalized bound" for matrix-valued Martingales. This result is a natural generalization of Theorem 1 in Abbasi-Yadkori et al. (2011).

Theorem 3. *(Self-Normalized Bound for Matrix-Valued Martingales) Let* $\{\mathcal{H}_t\}_{t=0}^{\infty}$ *be a filtration, and* $\{\eta_t\}_{t=1}^{\infty}$ *be a* $\Re^K$*-valued Martingale difference sequence with respect to* $\{\mathcal{H}_t\}_{t=0}^{\infty}$*. Specifically, for all* t, η_t *is* $\mathcal{H}_t$*-measurable and satisfies (1)* $\mathbb{E}\left[\eta_t | \mathcal{H}_{t-1}\right] = 0$ *and (2)* $\|\eta_t\|_\infty \leq 1$ *with probability 1 conditioning on* $\mathcal{H}_{t-1}$*. Let* $\{X_t\}_{t=1}^{\infty}$ *be a* $\Re^{d \times K}$*-valued stochastic process such that* X_t *is* $\mathcal{H}_{t-1}$ *measurable. Assume that* $V \in \Re^{d \times d}$ *is a positive-definite matrix. For any* $t \geq 0$*, define*

$$\overline{V}_t = V + K\sum_{s=1}^{t} X_s X_s^T \qquad S_t = \sum_{s=1}^{t} X_s \eta_s. \tag{33}$$

Then, for any $\delta > 0$*, with probability at least* $1 - \delta$*, we have*

$$\|S_t\|_{\overline{V}_t^{-1}}^2 \leq 2\log\left(\frac{\det\left(\overline{V}_t\right)^{1/2} \det\left(V\right)^{-1/2}}{\delta}\right) \quad \forall t \geq 0. \tag{34}$$

We first define some useful notations. Similarly as Abbasi-Yadkori et al. (2011), for any $\lambda \in \Re^d$ and any t, we define D_t^λ as

$$D_t^\lambda = \exp\left(\lambda^T X_t \eta_t - \frac{K}{2}\|X_t^T \lambda\|_2^2\right), \tag{35}$$

and $M_t^\lambda = \prod_{s=1}^{t} D_s^\lambda$ with convention $M_0^\lambda = 1$. Note that both D_t^λ and M_t^λ are $\mathcal{H}_t$-measurable, and $\left\{M_t^\lambda\right\}_{t=0}^{\infty}$ is a supermartingale with respect to the filtration $\{\mathcal{H}_t\}_{t=0}^{\infty}$. To see it, notice that conditioning on $\mathcal{H}_{t-1}$, we have

$$\lambda^T X_t \eta_t = (X_t^T \lambda)^T \eta_t \leq \|X_t^T \lambda\|_1 \|\eta_t\|_\infty \leq \|X_t^T \lambda\|_1 \leq \sqrt{K} \|X_t^T \lambda\|_2$$

with probability 1. This implies that $\lambda^T X_t \eta_t$ is conditionally $\sqrt{K}\|X_t^T \lambda\|_2$-subGaussian. Thus, we have

$$\mathbb{E}\left[D_t^\lambda \middle| \mathcal{H}_{t-1}\right] = \mathbb{E}\left[\exp\left(\lambda^T X_t \eta_t\right) \middle| \mathcal{H}_{t-1}\right] \exp\left(-\frac{K}{2}\|X_t^T \lambda\|_2^2\right) \leq \exp\left(\frac{K}{2}\|X_t^T \lambda\|_2^2 - \frac{K}{2}\|X_t^T \lambda\|_2^2\right) = 1.$$

Thus,

$$\mathbb{E}\left[M_t^\lambda \middle| \mathcal{H}_{t-1}\right] = M_{t-1}^\lambda \mathbb{E}\left[D_t^\lambda \middle| \mathcal{H}_{t-1}\right] \leq M_{t-1}^\lambda.$$

So $\left\{M_t^\lambda\right\}_{t=0}^{\infty}$ is a supermartingale with respect to the filtration $\{\mathcal{H}_t\}_{t=0}^{\infty}$. Then, following Lemma 8 of Abbasi-Yadkori et al. (2011), we have the following lemma:

Lemma 4. *Let τ be a stopping time with respect to the filtration $\{\mathcal{H}_t\}_{t=0}^{\infty}$. Then for any $\lambda \in \Re^d$, M_τ^λ is almost surely well-defined and $\mathbb{E}\left[M_\tau^\lambda\right] \leq 1$.*

Proof. First, we argue that M_τ^λ is almost surely well-defined. By Doob's convergence theorem for nonnegative supermartingales, $M_\infty^\lambda = \lim_{t\to\infty} M_t^\lambda$ is almost surely well-defined. Hence M_τ^λ is indeed well-defined independent of $\tau < \infty$ or not. Next, we show that $\mathbb{E}\left[M_\tau^\lambda\right] \leq 1$. Let $Q_t^\lambda = M_{\min\{\tau,t\}}^\lambda$ be a stopped version of $\left\{M_t^\lambda\right\}_{t=1}^{\infty}$. By Fatou's Lemma, we have $\mathbb{E}\left[M_\tau^\lambda\right] = \mathbb{E}\left[\liminf_{t\to\infty} Q_t^\lambda\right] \leq \liminf_{t\to\infty} \mathbb{E}\left[Q_t^\lambda\right] \leq 1$. □

The following results follow from Lemma 9 of Abbasi-Yadkori et al. (2011), which uses the "method of mixtures" technique. Let Λ be a Gaussian random vector in $\Re^d$ with mean 0 and covariance matrix V^{-1}, and independent of all the other random variables. Let $\mathcal{H}_\infty$ be the tail σ-algebra of the filtration, i.e. the σ-algebra generated by the union of all events in the filtration. We further define $M_t = \mathbb{E}\left[M_t^\Lambda \middle| \mathcal{H}_\infty\right]$ for all $t = 0, 1, \ldots$ and $t = \infty$. Note that M_∞ is almost surely well-defined since M_∞^λ is almost surely well-defined.

Let τ be a stopping time with respect to the filtration $\{\mathcal{H}_t\}_{t=0}^{\infty}$. Note that M_τ is almost surely well-defined since M_∞ is almost surely well-defined. Since $\mathbb{E}\left[M_\tau^\lambda\right] \leq 1$ from Lemma 4, we have

$$\mathbb{E}\left[M_\tau\right] = \mathbb{E}\left[M_\tau^\Lambda\right] = \mathbb{E}\left[\mathbb{E}\left[M_\tau^\Lambda \middle| \Lambda\right]\right] \leq 1.$$

The following lemma follows directly from the proof for Lemma 9 of Abbasi-Yadkori et al. (2011), which can be derived by algebra. The proof is omitted here.

Lemma 5. *For all finite $t = 0, 1, \ldots$, we have*

$$M_t = \left(\frac{\det(V)}{\det(\overline{V}_t)}\right)^{1/2} \exp\left(\frac{1}{2}\|S_t\|_{\overline{V}_t^{-1}}\right). \tag{36}$$

Note that Lemma 5 implies that for finite t, $\|S_t\|^2_{\overline{V}_t^{-1}} > 2\log\left(\frac{\det(\overline{V}_t)^{1/2}\det(V)^{-1/2}}{\delta}\right)$ and $M_t > \frac{1}{\delta}$ are equivalent. Consequently, for any stopping time τ, the event

$$\left\{\tau < \infty,\, \|S_\tau\|^2_{\overline{V}_\tau^{-1}} > 2\log\left(\frac{\det\left(\overline{V}_\tau\right)^{1/2}\det\left(V\right)^{-1/2}}{\delta}\right)\right\}$$

is equivalent to $\left\{\tau < \infty,\, M_\tau > \frac{1}{\delta}\right\}$. Finally, we prove Theorem 3:

Proof. We define the "bad event" at time $t = 0, 1, \ldots$ as:

$$B_t(\delta) = \left\{\|S_t\|^2_{\overline{V}_t^{-1}} > 2\log\left(\frac{\det\left(\overline{V}_t\right)^{1/2}\det\left(V\right)^{-1/2}}{\delta}\right)\right\}.$$

We are interested in bounding the probability of the "bad event" $\bigcup_{t=1}^{\infty} B_t(\delta)$. Let Ω denote the sample space, for any outcome $\omega \in \Omega$, we define $\tau(\omega) = \min\{t \geq 0 : \omega \in B_t(\delta)\}$, with the convention that $\min \emptyset = +\infty$. Thus, τ is a stopping time. Notice that $\bigcup_{t=1}^{\infty} B_t(\delta) = \{\tau < \infty\}$. Moreover, if $\tau < \infty$, then by definition of τ, we have $\|S_\tau\|^2_{\overline{V}_\tau^{-1}} > 2\log\left(\frac{\det(\overline{V}_\tau)^{1/2}\det(V)^{-1/2}}{\delta}\right)$, which is equivalent to $M_\tau > \frac{1}{\delta}$ as discussed above. Thus we have

$$
\begin{aligned}
P\left(\bigcup_{t=1}^{\infty} B_t(\delta)\right) &\overset{(a)}{=} P(\tau < \infty) \\
&\overset{(b)}{=} P\left(\|S_\tau\|^2_{\overline{V}_\tau^{-1}} > 2\log\left(\frac{\det\left(\overline{V}_\tau\right)^{1/2}\det(V)^{-1/2}}{\delta}\right), \tau < \infty\right) \\
&\overset{(c)}{=} P\left(M_\tau > 1/\delta, \tau < \infty\right) \\
&\leq P\left(M_\tau > 1/\delta\right) \\
&\overset{(d)}{\leq} \delta,
\end{aligned}
$$

where equalities (a) and (b) follow from the definition of τ, equality (c) follows from Lemma 5, and inequality (d) follows from Markov's inequality. This concludes the proof for Theorem 3. □

We conclude this section by briefly discussing a special case. If for any t, the elements of η_t are statistically independent conditioning on $\mathcal{H}_{t-1}$, then we can prove a variant of Theorem 3: with $\overline{V}_t = V + \sum_{s=1}^{t} X_s X_s^T$ and $S_t = \sum_{s=1}^{t} X_s \eta_s$, Equation 34 holds with probability at least $1 - \delta$. To see it, notice that in this case

$$
\begin{aligned}
\mathbb{E}\left[\exp\left(\lambda^T X_t \eta_t\right) \middle| \mathcal{H}_{t-1}\right] &= \mathbb{E}\left[\prod_{k=1}^{K} \exp\left((X_t^T \lambda)(k)\eta_t(k)\right) \middle| \mathcal{H}_{t-1}\right] \\
&\overset{(a)}{=} \prod_{k=1}^{K} \mathbb{E}\left[\exp\left((X_t^T \lambda)(k)\eta_t(k)\right) \middle| \mathcal{H}_{t-1}\right] \\
&\overset{(b)}{\leq} \prod_{k=1}^{K} \exp\left(\frac{(X_t^T \lambda)(k)^2}{2}\right) = \exp\left(\frac{\left\|X_t^T \lambda\right\|^2}{2}\right),
\end{aligned}
\tag{37}
$$

where (k) denote the k-th element of the vector. Note that the equality (a) follows from the conditional independence of the elements in η_t, and inequality (b) follows from $|\eta_t(k)| \leq 1$ for all t and k. Thus, if we redefine $D_t^\lambda = \exp\left(\lambda^T X_t \eta_t - \frac{1}{2}\|X_t^T \lambda\|_2^2\right)$, and $M_t^\lambda = \prod_{s=1}^{t} D_s^\lambda$, we can prove that $\{M_t^\lambda\}_t$ is a supermartingale. Consequently, using similar analysis techniques, we can prove the variant of Theorem 3 discussed in this paragraph.

D. Laplacian Regularization

As explained in section 7, enforcing Laplacian regularization leads to the following optimization problem:

$$
\widehat{\boldsymbol{\theta}}_t = \arg\min_{\boldsymbol{\theta}} [\sum_{j=1}^{t} \sum_{u \in \mathcal{S}_t} (y_{u,j} - \boldsymbol{\theta}_u X)^2 + \lambda_2 \sum_{(u_1,u_2) \in \mathcal{E}} ||\boldsymbol{\theta}_{u_1} - \boldsymbol{\theta}_{u_2}||_2^2]
$$

Here, the first term is the data fitting term, whereas the second term is the Laplacian regularization terms which enforces smoothness in the source node estimates. This can optimization problem can be re-written as follows:

$$
\widehat{\boldsymbol{\theta}}_t = \arg\min_{\boldsymbol{\theta}} \left[\sum_{j=1}^{t} \sum_{u \in \mathcal{S}_t} (y_{u,j} - \boldsymbol{\theta}_u X)^2 + \lambda_2 \boldsymbol{\theta}^T (L \otimes I_d) \boldsymbol{\theta}\right]
$$

Here, $\boldsymbol{\theta} \in \Re^{dn}$ is the concatenation of the n d-dimensional $\boldsymbol{\theta}_u$ vectors and $A \otimes B$ refers to the Kronecker product of matrices A and B. Setting the gradient of equation 38 to zero results in solving the following linear system:

$$[XX^T \otimes I_n + \lambda_2 L \otimes I_d]\widehat{\boldsymbol{\theta}}_t = b_t \tag{38}$$

Here b_t corresponds to the concatenation of the n d-dimensional vectors $b_{u,t}$. This is the Sylvester equation and there exist sophisticated methods of solving it. For simplicity, we focus on the special case when the features are derived from the Laplacian eigenvectors (Section 7).

Let β_t be a diagonal matrix such that $\beta_t u, u$ refers to the number of times node u has been selected as the source. Since the Laplacian eigenvectors are orthogonal, when using Laplacian features, $XX^T \otimes I_n = \beta \otimes I_d$. We thus obtain the following system:

$$[(\beta + \lambda_2 L) \otimes I_d]\widehat{\boldsymbol{\theta}}_t = b_t \tag{39}$$

Note that the matrix $(\beta + \lambda_2 L)$ and thus $(\beta + \lambda_2 L) \otimes I_d$ is positive semi-definite and can be solved using conjugate gradient (Hestenes & Stiefel, 1952).

For conjugate gradient, the most expensive operation is the matrix-vector multiplication $(\beta + \lambda_2 L) \otimes I_d]\mathbf{v}$ for an arbitrary vector $\mathbf{v}$. Let vec be an operation that takes a $d \times n$ matrix and stacks it column-wise converting it into a dn-dimensional vector. Let V refer to the $d \times n$ matrix obtained by partitioning the vector $\mathbf{v}$ into columns of V. Given this notation, we use the property that $(B^T \otimes A)\mathbf{v} = vec(AVB)$. This implies that the matrix-vector multiplication can then be rewritten as follows:

$$(\beta + \lambda_2 L) \otimes I_d \mathbf{v} = \mathrm{vec}(V\left(\beta + \lambda_2 L^T\right)) \tag{40}$$

Since β is a diagonal matrix, $V\beta$ is an $O(dn)$ operation, whereas VL^T is an $O(dm)$ operation since there are only m non-zeros (corresponding to edges) in the Laplacian matrix. Hence the complexity of computing the mean $\widehat{\boldsymbol{\theta}}_t$ is an order $O((d(m+n))\kappa)$ where κ is the number of conjugate gradient iterations. In our experiments, similar to (Vaswani et al., 2017), we warm-start with the solution at the previous round and find that $\kappa = 5$ is enough for convergence.

Unlike independent estimation where we update the UCB estimates for only the selected nodes, when using Laplacian regularization, the upper confidence values for each reachability probability need to be recomputed in each round. Once we have an estimate of $\boldsymbol{\theta}$, calculating the mean estimates for the reachabilities for all u, v requires $O(dn^2)$ computation. This is the most expensive step when using Laplacian regularization.

We now describe how to compute the confidence intervals. For this, let $\boldsymbol{D}$ denote the diagonal of $(\beta + \lambda_2 L)^{-1}$. The UCB value $z_{u,v,t}$ can then be computed as:

$$z_{u,v,t} = \sqrt{\boldsymbol{D}_u}||x_v||_2 \tag{41}$$

The ℓ_2 norms for all the target nodes v can be pre-computed. If we maintain the $\boldsymbol{D}$ vector, the confidence intervals for all pairs can be computed in $O(n^2)$ time.

Note that $\boldsymbol{D}_t$ requires $O(n)$ storage and can be updated across rounds in $O(K)$ time using the Sherman Morrison formula. Specifically, if $\boldsymbol{D}_{u,t}$ refers to the u^{th} element in the vector $\boldsymbol{D}_t$, then

$$\boldsymbol{D}_{u,t+1} = \begin{cases} \dfrac{\boldsymbol{D}_{u,t}}{(1 + \boldsymbol{D}_{u,t})}, & \text{if} u \in \mathcal{S}_t \\ \boldsymbol{D}_{u,t}, & \text{otherwise} \end{cases}$$

Hence, the total complexity of implementing Laplacian regularization is $O(dn^2)$. We need to store the $\boldsymbol{\theta}$ vector, the Laplacian and the diagonal vectors β and $\boldsymbol{D}$. Hence, the total memory requirement is $O(dn + m)$.