A Additional figures and examples

A.1 Special cases of transductive regret.

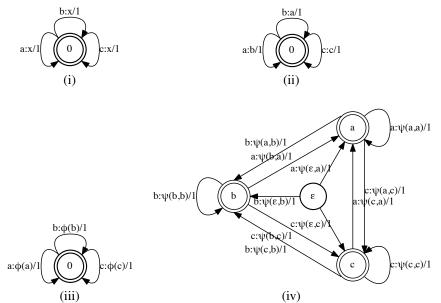


Figure 4: Several families of WFSTs for special cases of transductive regret for $\Sigma = \{a, b, c\}$. (i) External regret with parameter $x \in \Sigma$. (ii) Internal regret: family of transducers \mathcal{T}_{a_1,a_2} with $a_1 \neq a_2$, $a_1, a_2 \in \Sigma$; example shown for $\mathcal{T}_{a,b}$. (iii) Swap regret with parameter $\varphi \colon \Sigma \to \Sigma$. (iv) Bigram conditional swap regret with parameter $\psi \colon (\Sigma \cup \{\epsilon\}) \times \Sigma \to \Sigma$.

A.2 Example with a swapping subset.

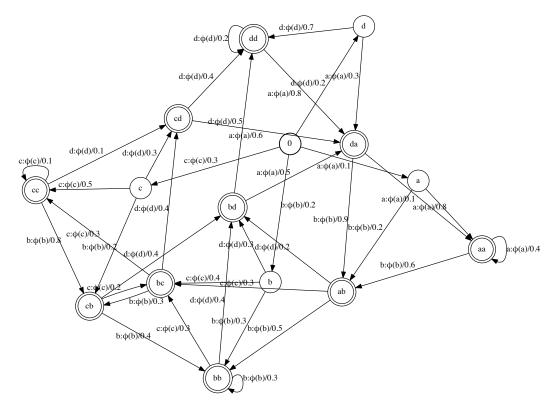


Figure 5: Example of a WFST with $\Sigma = \{a, b, c, d\}$ and where each state has a swapping subset.

B Pseudocode of FASTTRANSDUCE

Algorithm 3: FASTTRANSDUCE; $(\overline{A}_{u,i})_{u \in Q_{\mathcal{T}}, i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]}$ external regret minimization algorithms.

```
Algorithm: FASTTRANSDUCE(\mathcal{T}, (\mathcal{A}_{u,i})_{u \in Q_{\mathcal{T}}, i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]})u \leftarrow I_{\mathcal{T}}for t \leftarrow 1 to T dofor each i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]] doq_i \leftarrow \text{QUERY}(\mathcal{A}_{u,i})\mathbf{Q}^{t,u} \leftarrow [\mathsf{q}_1 \mathbf{1}_{1 \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \cdots \mathsf{q}_N \mathbf{1}_{N \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]}]^{\mathsf{T}}for each j \leftarrow 1 to N doc_j \leftarrow \min_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{j \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]}\alpha_t \leftarrow \|\mathbf{c}\|_1; \quad \tau_t \leftarrow \left\lceil \frac{\log(\frac{1}{\sqrt{t}})}{\log(1-\alpha_t)} \right\rceilif \tau_t < N thenp_t \leftarrow p_t^0 \leftarrow \frac{\mathsf{c}}{\alpha_t}for \tau \leftarrow 1 to \tau_t do(\mathsf{p}_t^{\tau})^{\mathsf{T}} \leftarrow (\mathsf{p}_t^{\tau})^{\mathsf{T}} (\mathbf{Q}^{t,u} - \vec{1}\mathbf{c}^{\mathsf{T}}); \mathsf{p}_t \leftarrow \mathsf{p}_t + \mathsf{p}_t^{\mathsf{T}}p_t \leftarrow \mathsf{sAMPLE}(\mathsf{p}_t); \quad \mathbf{l}_t \leftarrow \mathsf{RECEIVELOSS}(); \quad u \leftarrow \delta_{\mathcal{T}}(u, x_t)for each i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]] doATTRIBUTELOSS(\mathcal{A}_{u,i}, \mathsf{p}_t[i] \mathsf{l}_t)
```

C Pseudocode of FASTSLEEPTRANSDUCE

Algorithm 4: FASTSLEEPTRANSDUCE. $(\mathcal{A}_{u,i})$ sleeping regret minimization algorithms.

```
Algorithm: FASTSLEEPTRANSDUCE(\mathcal{T}, \{\mathcal{A}_{u,i}\}_{u \in Q_{\mathcal{T}}, i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]})
 u \leftarrow I_T
 for t \leftarrow 1 to T do
              A_t \leftarrow \text{AWAKESET}()
             for each i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}^{\smile}[u]] \cap A_t do
                         \mathbf{q}_i \leftarrow \mathrm{QUERY}(\mathcal{A}_{u,i}); \quad \mathbf{q}_i^{A_t} \leftarrow \frac{\mathbf{q}_i|_{A_t}}{\sum_{j \in A_t} \mathbf{q}_i}
            \mathbf{Q}^{t,u} \leftarrow [\mathsf{q}_1^{A_t} \mathbf{1}_{1 \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]] \cap A_t}; \dots; \mathsf{q}_N^{A_t} \mathbf{1}_{N \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]] \cap A_t}]for each j \leftarrow 1 to N do
                         c_j \leftarrow \min_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]] \cap A_t} \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{j \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]] \cap A_t}
             \alpha_t \leftarrow \|\mathbf{c}\|_1; \quad \tau_t \leftarrow \left[\frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha_t)}\right]
            if \tau_t < N then
                         \mathbf{p}_t \leftarrow \mathbf{p}_t^0 \leftarrow \frac{\mathbf{c}}{\alpha_t}
for \tau \leftarrow 1 to \tau_t do
                          \begin{array}{c} (\mathbf{p}_t^{\tau})^{\top} \leftarrow (\mathbf{p}_t^{\tau})^{\top} (\mathbf{Q}^{t,u} - [\mathbf{1}_{1 \in A_t}; \dots; \mathbf{1}_{|\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]| \in A_t}] \mathbf{c}^{\top}) \\ \mathbf{p}_t \leftarrow \mathbf{p}_t + \mathbf{p}_t^{\tau} \\ \mathbf{p}_t \leftarrow \frac{\mathbf{p}_t}{\|\mathbf{p}_t\|_1} \end{array} 
              else
                          \mathsf{p}_t^{\top} \leftarrow \text{FIXED-POINT}(\mathbf{Q}^{t,u})
             \mathbf{p}_{t}^{A_{t}} \leftarrow \frac{\mathbf{p}_{t}|A_{t}}{\sum_{j \in A_{t}} \mathbf{p}_{t,j}}; \quad x_{t} \leftarrow \mathsf{SAMPLE}(\mathbf{p}_{t}^{A_{t}}); \quad \mathbf{l}_{t} \leftarrow \mathsf{RECEIVELOSS}(); \quad u \leftarrow \delta_{\mathcal{T}}[u, x_{t}]
for each i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]] do
                          ATTRIBUTELOSS(\mathcal{A}_{u,i}, \mathsf{p}_t[i]\mathbf{l}_t)
```

D Proof of Theorem 1

Theorem 1. Let A_1, \ldots, A_N be external regret minimizing algorithms admitting data-dependent regret bounds of the form $O(\sqrt{L_T(A_i) \log N})$, where $L_T(A_i)$ is the cumulative loss of A_i after T rounds. Assume that, at each round, the sum of the minimal probabilities given to an expert by these algorithms is bounded below by some constant $\alpha > 0$. Then, FASTSWAP achieves a swap regret in $O(\sqrt{TN \log N})$ with a per-iteration complexity in $O(N^2 \min \{\frac{\log T}{\log(1/(1-\alpha))}, N\})$.

Proof. Let p_t be the distribution returned by FASTSWAP at round t. For any distribution $p_t^*, t \in [T]$, the following inequality holds:

$$\begin{split} \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} &= \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} \\ &- \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} \\ &\leq \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \|\mathsf{p}_{t} - \mathsf{p}_{t}^{*}\|_{1} \|l_{t}\|_{\infty} \mathbf{1}_{\tau_{t} < N} \\ &\leq \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \|\mathsf{p}_{t} - \mathsf{p}_{t}^{*}\|_{1} \|l_{t}\|_{\infty} \mathbf{1}_{\tau_{t} < N} \end{split}$$

Let p_t^* be the stationary distribution of the row stochastic matrix \mathbf{Q}^t , $\mathbf{p}_t^{*\top}\mathbf{Q}^t = \mathbf{p}_t^{*\top}$. Then, we can write

$$\begin{split} \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} &= \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{p}_{t,j}^{*} l_{t,j} \mathbf{1}_{\tau_{t} < N} \\ &= \sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{j=1}^{N} \mathsf{p}_{t,i}^{N} \mathbf{Q}_{i,j}^{t} l_{t,j} \mathbf{1}_{\tau_{t} < N} \\ &= \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t} \mathsf{p}_{t,i} l_{t,j} \mathbf{1}_{\tau_{t} < N} + \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t} (\mathsf{p}_{t,i}^{*} \\ &- \mathsf{p}_{t,i}) l_{t,j} \mathbf{1}_{\tau_{t} < N} \\ &\leq \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t} \mathsf{p}_{t,i} l_{t,j} \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \|\mathsf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N}. \end{split}$$

On the other hand, by design, if $\tau_t \ge N$, then $p_t = p_t^*$, so that

$$\sum_{t=1}^{T} \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] \mathbf{1}_{\tau_t \ge N} = \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^t \mathbf{p}_{t,i} l_{t,j} \mathbf{1}_{\tau_t \ge N}.$$

Thus, it follows that

$$\begin{split} \sum_{t=1}^{T} & \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}}[l_{t}(x_{t})] \leq \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t} \mathsf{p}_{t,i} l_{t,j} + 2 \sum_{t=1}^{T} \|\mathsf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} 1_{\tau_{t} < N} \\ & \leq \sum_{i=1}^{N} \left[\min_{j \in [N]} \sum_{t=1}^{T} \mathsf{p}_{t,i} l_{t,j} + \operatorname{Reg}_{T}(\mathcal{A}_{i}, \Phi_{\text{ext}}) \right] + 2 \sum_{t=1}^{T} \|\mathsf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} 1_{\tau_{t} < N} \\ & = \min_{\varphi \in \Phi_{\text{swap}}} \sum_{i=1}^{N} \left[\sum_{t=1}^{T} \mathsf{p}_{t,i} l_{t,\varphi(i)} + \operatorname{Reg}_{T}(\mathcal{A}_{i}, \Phi_{\text{ext}}) \right] + 2 \sum_{t=1}^{T} \|\mathsf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} 1_{\tau_{t} < N}. \end{split}$$

Now let $L_T(\mathcal{A}_i)$ denote the cumulative loss incurred by algorithm \mathcal{A}_i . Since the losses attributed to algorithm \mathcal{A}_i are scaled by $p_{t,i}$, at each round, the sum of the losses over all the algorithms is at most 1. Thus, by Jensen's inequality, the following inequalities hold:

$$\frac{1}{N} \sum_{i=1}^{N} \operatorname{Reg}_{T}(\mathcal{A}_{i}, \Phi_{\text{ext}}) = \frac{1}{N} \sum_{i=1}^{N} O\left(\sqrt{L_{T}(\mathcal{A}_{i}) \log N}\right)$$
$$\leq O\left(\sqrt{\frac{1}{N} \sum_{i=1}^{N} L_{T}(\mathcal{A}_{i}) \log N}\right) \leq O\left(\sqrt{\frac{T \log N}{N}}\right),$$

which implies $\sum_{i=1}^{N} \operatorname{Reg}_{T}(\mathcal{A}_{i}, \Phi_{ext}) \leq \sqrt{TN \log N}$.

Finally, during the rounds in which $1_{\tau_t < N}$, p_t is an RPM approximation of p_t^* using τ_t iterations. Thus, by Equation 3.7 in [Nesterov and Nemirovski, 2015] the following inequality holds: $||\mathbf{p}_t - \mathbf{p}_t^*||_1 \le (1 - \alpha_t)^{\tau_t}$. Since τ_t is chosen so that the inequality $(1 - \alpha_t)^{\tau_t} \le 1/\sqrt{t}$ holds, it follows that $\sum_{t=1}^T ||\mathbf{p}_t - \mathbf{p}_t^*|| 1_{\tau_t < N} \le \sum_{t=1}^T 1/\sqrt{t} \le \sqrt{T}$, which proves the regret bound $\operatorname{Reg}_T(\mathcal{A}, \Phi_{swap}) \le O(\sqrt{TN \log N})$.

Furthermore, the computational cost of the *t*-th iteration of the algorithm is dominated by τ_t matrix multiplications or the solution of the linear system. τ_t can be bounded as follows: $\tau_t = \left\lceil \frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha_t)} \right\rceil \le \frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha)} + 1$. Thus, the computational cost of the *t*-th iteration is in

$$O\left(N^2 \min\left\{\frac{\log t}{\log(1/(1-\alpha_t))}, N\right\}\right) \le O\left(N^2 \min\left\{\frac{\log T}{\log(1/(1-\alpha))}, N\right\}\right).$$

E Proof of Theorem 2

Theorem 2. Let $(\mathcal{A}_{u,i})_{u \in Q, i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]}$ be external regret minimizing algorithms admitting datadependent regret bounds of the form $O(\sqrt{L_T(\mathcal{A}_{u,i})\log N})$, where $L_T(\mathcal{A}_{u,i})$ is the cumulative loss of $\mathcal{A}_{u,i}$ after T rounds. Assume that, at each round, the sum of the minimal probabilities given to an expert by these algorithms is bounded below by some constant $\alpha > 0$. Then, FASTTRANSDUCE achieves a transductive regret against \mathcal{T} that is in $O(\sqrt{T|\mathsf{E}_{\mathcal{T}}|_{in}\log N})$ with a per-iteration complexity in $O\left(N^2\min\left\{\frac{\log T}{\log(1/(1-\alpha))},N\right\}\right)$.

Proof. Let p_t be the distribution output by FASTTRANSDUCE at round t. For any distribution p_t^* , $t \in [T]$, the following inequalities hold:

$$\sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} = \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N}$$
$$- \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N}$$
$$\leq \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \|\mathsf{p}_{t} - \mathsf{p}_{t}^{*}\|_{1} \|l_{t}\|_{\infty} \mathbf{1}_{\tau_{t} < N}$$
$$\leq \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \|\mathsf{p}_{t} - \mathsf{p}_{t}^{*}\|_{1} \|l_{t}\|_{\infty} \mathbf{1}_{\tau_{t} < N}$$

Let u_t be the state that the algorithm is in at time t as a result of its past actions. Consider the matrix \mathbf{Q}^{t,u_t} defined in the algorithm. The restriction of the matrix \mathbf{Q}^{t,u_t} to its non-zero rows and columns is a row stochastic matrix. Let \mathbf{p}_t^* be its stationary distribution, and by augmenting it with zeros in the zero rows of \mathbf{Q}^{t,u_t} , we may take $\mathbf{p}_t^* \in \Delta_N$ as a fixed point of \mathbf{Q}^{t,u_t} . Then, we can write:

$$\begin{split} \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathbf{p}_{t}^{*}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} &= \sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{j=1}^{N} \mathbf{p}_{t,i}^{N} \mathbf{Q}_{i,j}^{t,u_{t}} l_{t,j} \mathbf{1}_{\tau_{t} < N} \\ &= \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t,u_{t}} \mathbf{p}_{t,i} l_{t,j} \mathbf{1}_{\tau_{t} < N} \\ &+ \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t,u_{t}} (\mathbf{p}_{t,i}^{*} - \mathbf{p}_{t,i}) l_{t,j} \mathbf{1}_{\tau_{t} < N} \\ &\leq \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t,u_{t}} \mathbf{p}_{t,i} l_{t,j} \mathbf{1}_{\tau_{t} < N} + \sum_{t=1}^{T} \|\mathbf{p}_{t}^{*} - \mathbf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N}. \end{split}$$

On the other hand, by design, if $\tau_t \ge N$, then $p_t = p_t^*$, so that

$$\sum_{t=1}^{T} \mathbb{E}_{x_t \sim \mathsf{p}_t}[l_t(x_t)] \mathbf{1}_{\tau_t \geq N} = \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t,u_t} \mathsf{p}_{t,i} l_{t,j} \mathbf{1}_{\tau_t \geq N}.$$

Thus, it follows that for any WFST $\mathfrak{T} \in \mathcal{T}$,

$$\begin{split} \sum_{t=1}^{T} \mathbb{E}_{x_{t} \sim \mathsf{p}_{t}}[l_{t}(x_{t})] &\leq \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \sum_{u \in Q_{\mathcal{T}}} \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1})=u} \mathsf{p}_{t,i} l_{t,j} + 2 \sum_{t=1}^{T} \|\mathsf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N} \\ &= \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1})=u} \mathsf{p}_{t,i} l_{t,j} \end{split}$$

$$\begin{split} &+ 2\sum_{t=1}^{T} \|\mathbf{p}_{t}^{*} - \mathbf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N} \\ &\leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \min_{i^{*} \in \mathsf{olab}[\mathsf{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_{t}]]} \sum_{t=1}^{T} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}) = u} \mathsf{p}_{t,i} l_{t,i^{*}} \\ &+ 2\sum_{t=1}^{T} \|\mathbf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N} + \sum_{i=1}^{N} \sum_{u \in Q_{\mathcal{T}}} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{ext}}) \\ &\leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]} \sum_{e \in \mathsf{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_{t}]} \sum_{t=1}^{T} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}) = u} \mathsf{p}_{t,i} w[e] l_{t}(\mathsf{olab}[e]) \\ &+ 2\sum_{t=1}^{T} \|\mathbf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N} + \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{ext}}) \\ &= \sum_{t=1}^{T} \sum_{x_{t} \sim \mathsf{p}_{t}} \left[\sum_{e \in \mathsf{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_{t}]} w[e] l_{t}(\mathsf{olab}[e]) \right] + 2\sum_{t=1}^{T} \|\mathbf{p}_{t}^{*} - \mathsf{p}_{t}\|_{1} \mathbf{1}_{\tau_{t} < N} \\ &+ \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{ext}}). \end{split}$$

Now let $L_T(\mathcal{A}_{u,i})$ denote the cumulative loss incurred by algorithm $\mathcal{A}_{u,i}$. Since the losses attributed to algorithm $\mathcal{A}_{u,i}$ are scaled by $1_{\delta_T(I_T, x_{1:t-1})=u} p_{t,i}$, it follows that at each round, the sum of the losses over all the algorithms is at most 1. Thus, by Jensen's inequality, it follows that

$$\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{ext}}) \\
= \frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \sqrt{L_{T}(\mathcal{A}_{u,i}) \log(N)} \\
\leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|}} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} L_{T}(\mathcal{A}_{u,i}) \log(N) \\
\leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|}} T \log(N),$$

so that $\sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{ext}}) \leq \sqrt{T \sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]| \log(N)}.$

Finally, during the rounds in which $1_{\tau_t < N}$, p_t is an RPM approximation of p_t^* using τ_t iterations. Thus, it follows from Equation 3.7 in [Nesterov and Nemirovski, 2015] that $\|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \le (1 - \alpha_t)^{\tau_t}$. By the algorithm's choice of τ_t , $\|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \le \frac{1}{\sqrt{t}}$. Thus, it follows that $\sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 1_{\tau_t < N} \le \sqrt{T}$, so that $\operatorname{Reg}_T(\mathcal{A}, \mathcal{T}) \le O(\sqrt{T \sum_{u \in Q_{\mathcal{T}}} ||\operatorname{abl}[\mathsf{E}_{\mathcal{T}}[q]]| \log(N)})$.

Moreover, the computational cost of the *t*-th iteration of the algorithm is dominated by τ_t matrix multiplications or the solution of the linear system. τ_t can be bounded as follows: $\tau_t = \left\lceil \frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha_t)} \right\rceil \le \frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha)} + 1$. Thus, the computational cost of the *t*-th iteration is in

$$O\left(N^2 \min\left\{\frac{\log t}{\log(1/(1-\alpha_t))}, N\right\}\right) \le O\left(N^2 \min\left\{\frac{\log T}{\log(1/(1-\alpha))}, N\right\}\right).$$

F Proof of Theorem 3

Theorem 3. Let $(\mathcal{A}_{I,u,i})_{I \in \mathcal{I}, u \in Q_{\mathcal{T}}, i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]}$ be external regret minimizing algorithms admitting data-dependent regret bounds of the form $O(\sqrt{L_T(\mathcal{A}_{I,u,i})\log N})$, where $L_T(\mathcal{A}_{I,u,i})$ is the cumulative loss of $\mathcal{A}_{I,u,i}$ after T rounds. Let $\mathcal{A}_{\mathcal{I}}$ be an external regret minimizing algorithm over \mathcal{I} that admits a regret in $O(\sqrt{T\log(|\mathcal{I}|)})$ after T rounds. Assume further that at each round, the sum of the minimal probabilities given to an expert by these algorithms is bounded below by some constant $\alpha > 0$. Then, FASTTIMESELECTTRANSDUCE achieves a time-selection transductive regret with respect to the time-selection family \mathcal{I} and WFST family \mathcal{T} that is in $O(\sqrt{T(\log(|\mathcal{I}|) + |\mathsf{E}_{\mathcal{T}}|_{in}\log N)})$ with a per-iteration complexity in $O(N^2(\min\{\frac{\log(T)}{\log((1-\alpha)^{-1})}, N\} + |\mathcal{I}|))$.

Proof. We first note that since $\mathcal{A}_{\mathcal{I}}$ is designed to minimize external regret against the losses $(\tilde{I}^t)_{t=1}^T$, it follows that for any $I^* \in \mathcal{I}$,

$$\sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} \tilde{l}_{I}^{t} \leq \sum_{t=1}^{T} \tilde{l}_{I^{*}}^{t} + \operatorname{Reg}_{T}(\mathcal{A}_{\mathcal{I}}).$$

Let u_t be the state that the algorithm is in at time t as a result of its past actions. Consider the matrix \mathbf{Q}^{t,u_t} defined in the algorithm. The restriction of the matrix \mathbf{Q}^{t,u_t} to its non-zero rows and columns is a row stochastic matrix. Let \mathbf{p}_t^* be its stationary distribution, and by augmenting it with zeros in the zero rows of \mathbf{Q}^{t,u_t} , we may take $\mathbf{p}_t^* \in \Delta_N$ as a fixed point of of \mathbf{Q}^{t,u_t} . Then, by expanding the definition of $\tilde{\mathbf{l}}^t$, we can rewrite the expression on the left-hand side as

$$\begin{split} \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_{I}^{t} \tilde{l}_{I}^{t} \mathbf{1}_{\tau_{t} < N} &= \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_{I}^{t} I(t) \left(\mathbf{p}_{t}^{\top} \mathbf{M}^{t, u_{t}, I} \mathbf{l}_{t} - \mathbf{p}_{t}^{\top} \mathbf{l}_{t} \right) \mathbf{1}_{\tau_{t} < N} \\ &= \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_{I}^{t} I(t) \mathbf{p}_{t}^{\top} \mathbf{M}^{t, u_{t}, I} \mathbf{l}_{t} \mathbf{1}_{\tau_{t} < N} - \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_{I}^{t} I(t) \mathbf{p}_{t}^{\top} \mathbf{l}_{t} \mathbf{1}_{\tau_{t} < N} \\ &\geq \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_{I}^{t} I(t) (\mathbf{p}_{t}^{*})^{\top} \mathbf{M}^{t, u_{t}, I} \mathbf{l}_{t} \mathbf{1}_{\tau_{t} < N} - \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_{I}^{t} I(t) (\mathbf{p}_{t}^{*})^{\top} \mathbf{l}_{t} \mathbf{1}_{\tau_{t} < N} \\ &- \sum_{t=1}^{T} \|\mathbf{p}_{t} - \mathbf{p}_{t}^{*}\|_{1} \mathbf{1}_{\tau_{t} < N}. \end{split}$$

On the other hand, by design, if $\tau_t \ge N$, then $p_t = p_t^*$, so that

$$\sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} \tilde{l}_{I}^{t} \mathbf{1}_{\tau_{t} \geq N} = \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} I(t) (\mathsf{p}_{t}^{*})^{\top} \mathbf{M}^{t, u_{t}, I} \mathbf{l}_{t} \mathbf{1}_{\tau_{t} \geq N} - \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} I(t) (\mathsf{p}_{t}^{*})^{\top} \mathbf{l}_{t} \mathbf{1}_{\tau_{t} \geq N}.$$

Thus, it follows that

$$\sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} \tilde{l}_{I}^{t} \geq \sum_{t=1}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} I(t) (\mathsf{p}_{t}^{*})^{\top} \mathbf{M}^{t, u_{t}, I} \mathbf{l}_{t} - \sum_{t=0}^{T} \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_{I}^{t} I(t) (\mathsf{p}_{t}^{*})^{\top} \mathbf{l}_{t} - \sum_{t=1}^{T} \|\mathsf{p}_{t} - \mathsf{p}_{t}^{*}\|_{1} \mathbf{1}_{\tau_{t} < N}.$$

If $\sum_{I \in \mathcal{I}} I(t) \tilde{\mathsf{q}}_I^t \neq 0$, then the fact that p_t^* is a stationary distribution of $\mathbf{Q}^t = \frac{\sum_{I \in \mathcal{I}} I(t) \tilde{\mathsf{q}}_I^t \mathbf{M}^{t,u_t,I}}{\sum_{I \in \mathcal{I}} I(t) \tilde{\mathsf{q}}_I^t}$ implies that

$$\sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_I^t I(t)(\mathsf{p}_t^*)^\top \mathbf{M}^{t,u_t,I} \mathbf{l}_t = \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_I^t I(t)(\mathsf{p}_t^*)^\top \mathbf{l}_t$$

On the other hand, if $\sum_{I \in \mathcal{I}} I(t) \tilde{q}_I^t = 0$, then by non-negativity, it must be the case that $I(t) \tilde{q}_I^t = 0$ for every $I \in \mathcal{I}$. Thus, it follows that

$$\sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_I^t I(t)(\mathsf{p}_t^*)^\top \mathbf{M}^{t,u_t,I} \mathbf{l}_t = \sum_{I \in \mathcal{I}} \tilde{\mathsf{q}}_I^t I(t)(\mathsf{p}_t^*)^\top \mathbf{l}_t = 0,$$

which implies that

$$\sum_{t=1}^{T} -\tilde{l}_{I^*}^t \leq \sum_{t=1}^{T} \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \mathbf{1}_{\tau_t < N} + \operatorname{Reg}_T(\mathcal{A}_{\mathcal{I}}).$$

By expanding the definition of $\tilde{l}^t_{I^*},$ we can write

$$\sum_{t=1}^{T} -\tilde{l}_{I^*}^t = \sum_{t=1}^{T} -I^*(t) \left(\mathsf{p}_t^\top \mathbf{M}^{t,u_t,I^*} \mathbf{l}_t - \mathsf{p}_t^\top \mathbf{l}_t \right) = \sum_{t=1}^{T} I^*(t) \mathsf{p}_t^\top \mathbf{l}_t - I^*(t) \mathsf{p}_t^\top \mathbf{M}^{t,u_t,I^*} \mathbf{l}_t.$$

Moreover, for any $\mathfrak{T} \in \mathcal{T}$, we can bound the second term in the following way:

$$\begin{split} \sum_{t=1}^{T} I^{*}(t) \mathbf{p}_{t}^{\top} \mathbf{M}^{t,u_{t},I^{*}} \mathbf{l}_{t} &= \sum_{t=1}^{T} I^{*}(t) \sum_{i=1}^{N} \mathbf{p}_{t,i} \sum_{j=1}^{N} \mathbf{M}_{i,j}^{t,u_{t},I^{*}} l_{t,j} \\ &= \sum_{u \in Q_{\mathcal{T}}} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{M}_{i,j}^{t,u_{t},I^{*}} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1}) = u} I^{*}(t) \mathbf{p}_{t,i} l_{t,j} \\ &= \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{M}_{i,j}^{t,u_{t},I^{*}} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1}) = u} I^{*}(t) \mathbf{p}_{t,i} l_{t,j} \\ &\leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \min_{i^{*} \in \text{olab}[\mathsf{E}_{\mathcal{T}}[u]]} \sum_{t=1}^{T} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1}) = u} I^{*}(t) \mathbf{p}_{t,i} l_{t,i^{*}} \\ &+ \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{I,u,i}, \Phi_{\text{ext}}) \\ &\leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \sum_{e \in \mathsf{E}_{\mathcal{T}}[u]} w[e] \sum_{t=1}^{T} \mathbf{1}_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1}) = u} I^{*}(t) \mathbf{p}_{t,i} l_{t,olab}[e] \\ &+ \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{I,u,i}, \Phi_{\text{ext}}) \\ &= \sum_{u \in Q_{\mathcal{T}}} I^{*}(t) \sum_{x_{t} \sim \mathsf{P}_{t}} \left[\sum_{e \in \mathsf{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1}),x_{t}]} w[e] l_{t}(\mathsf{olab}[e]) \right] \\ &+ \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{I,u,i}, \Phi_{\text{ext}}), \end{split}$$

using the fact that algorithm $\mathcal{A}_{I,u,i}$ minimizes external regret against the surrogate losses $I(t)1_{\delta_{\mathcal{T}}(I_{\Phi},x_{1:t-1})=u}\mathsf{p}_{t,i}\mathbf{l}_{t}$.

As in Theorem 2, the scaling assumption on the external regret minimizing algorithms and Jensen's inequality imply that

$$\sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{I,u,i}, \Phi_{\mathsf{ext}}) \leq O\left(\sqrt{T \sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]| \log(N)}\right)$$

Thus, we can write for any $I^* \in \mathcal{I}$ that

$$\sum_{t=1}^{T} I^*(t) \mathbf{p}_t^\top \mathbf{l}_t - I^*(t) \mathbf{p}_t^\top \mathbf{M}^{t, u_t, I^*} \mathbf{l}_t - \sum_{t=1}^{T} I^*(t) \mathop{\mathbb{E}}_{x_t \sim \mathbf{p}_t} \left[\sum_{e \in \mathsf{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}, x_{1:t-1}}), x_t]} w[e] l_t(\mathsf{olab}[e]) \right]$$
$$\leq \operatorname{Reg}_T(\mathcal{A}_{\mathcal{I}}) + O\left(\sqrt{T \sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]| \log(N)} \right) + \sum_{t=1}^{T} \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \mathbf{1}_{\tau_t < N},$$

and as in Theorem 2, we can bound the l_1 approximation error of p_t for p_t^* by

$$\|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \le (1 - \alpha_t)^{\tau_t} \le \frac{1}{\sqrt{t}},$$

by the algorithm's choice of τ_t . Thus, by applying regret guarantee of algorithm $\mathcal{A}_{\mathcal{I}}$ together with the above calculations, the time-selection transductive regret of FASTTIMESELECTTRANSDUCE is

in
$$O\left(\sqrt{T\left(\log(|\mathcal{I}|) + \sum_{q \in Q_{\Phi}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]|\log(N)\right)}\right)$$

Moreover, at each round t, the computational cost of the algorithm is dominated by two quantities: the update of $|\mathcal{I}|N$ external regret minimizing algorithms over the N experts, which is in $O(|\mathcal{I}|N^2)$, and the fixed-point approximation or solution of the linear system, which is in

$$O\left(N^2 \min\left\{\frac{\log(t)}{\log\left((1-\alpha_t)^{-1}\right)}, N\right\}\right) \le O\left(N^2 \min\left\{\frac{\log(T)}{\log\left((1-\alpha)^{-1}\right)}, N\right\}\right).$$

G Proof of Theorem 4

Theorem 4. Assume that the sleeping regret minimizing algorithms used as inputs of FASTSLEEPTRANSDUCE achieve data-dependent regret bounds such that, if the algorithm selects the distributions $(\mathbf{p}_t)_{t=1}^T$ and observes losses $(\mathbf{l}_t)_{t=1}^T$ with awake sets $(A_t)_{t=1}^T$, then the regret of \mathcal{A}_i^q is at most $O\left(\sqrt{\sum_{t=1}^T \mathbf{u}^*(A_t) \mathbb{E}_{x_t \sim \mathbf{p}_t}[l_t(x_t)] \log(N)}\right)$. Assume further that at each round, the sum of the minimal probabilities given to an expert by these algorithms is bounded below by some constant $\alpha > 0$. Then, the sleeping regret $\operatorname{Reg}_T(\operatorname{FASTSLEEPTRANSDUCE}, \mathcal{T}, \mathcal{A}_1^T)$ of FASTSLEEPTRANSDUCE is upper bounded by $O\left(\sqrt{\sum_{t=1}^T \mathbf{u}(A_t) |\mathbf{E}_T|_{in} \log(N)}\right)$. Moreover, FASTSLEEPTRANSDUCE admits a per-iteration complexity in $O\left(N^2 \min\left\{\frac{\log T}{\log(1/(1-\alpha))}, N\right\}\right)$.

Proof. Let $u \in \Delta_N$, and let $p_t^{A_t}$ be the distribution output by FASTSLEEPTRANSDUCE at round t. For any distribution p_t^* , $t \in [T]$, the following inequalities hold:

$$\begin{aligned} \mathsf{u}(A_{t}) & \underset{x_{t} \sim \mathsf{p}_{t}^{A_{t}}}{\mathbb{E}} [l_{t}(x_{t})] \mathbf{1}_{\tau_{t} < N} = \mathsf{u}(A_{t}) \left(\underset{x_{t} \sim \mathsf{p}_{t}^{A_{t},*}}{\mathbb{E}} [l_{t}(x_{t})] + \underset{x_{t} \sim \mathsf{p}_{t}^{A_{t}}}{\mathbb{E}} [l_{t}(x_{t})] - \underset{x_{t} \sim \mathsf{p}_{t}^{A_{t},*}}{\mathbb{E}} [l_{t}(x_{t})] \right) \mathbf{1}_{\tau_{t} < N} \\ & \leq \mathsf{u}(A_{t}) \left(\underset{x_{t} \sim \mathsf{p}_{t}^{A_{t},*}}{\mathbb{E}} [l_{t}(x_{t})] + \|\mathsf{p}_{t}^{A_{t}} - \mathsf{p}_{t}^{A_{t},*}\|_{1} \|l_{t}\|_{\infty} \right) \mathbf{1}_{\tau_{t} < N} \\ & \leq \mathsf{u}(A_{t}) \left(\underset{x_{t} \sim \mathsf{p}_{t}^{A_{t},*}}{\mathbb{E}} [l_{t}(x_{t})] + \|\mathsf{p}_{t}^{A_{t}} - \mathsf{p}_{t}^{A_{t},*}\|_{1} \right) \mathbf{1}_{\tau_{t} < N}. \end{aligned}$$

Let u_t be the state that the algorithm is in at time t as a result of its past actions. Consider the matrix \mathbf{Q}^{t,u_t} defined in the algorithm. The restriction of \mathbf{Q}^{t,u_t} to its non-zero rows and columns is a row stochastic matrix. Let $\mathbf{p}_t^{A_t,*}$ be its stationary distribution, and by augmenting it with zeros in the zero rows of \mathbf{Q}^{t,u_t} , we may take $\mathbf{p}_t^{A_t,*} \in \Delta_N$ as a fixed point of \mathbf{Q}^{t,u_t} . Then, we can write:

$$\begin{split} &\sum_{t=1}^{T} \mathsf{u}(A_t) \underbrace{\mathbb{E}}_{x_t \sim \mathsf{p}_t^{A_t,*}} [l_t(x_t)] \mathbf{1}_{\tau_t < N} \\ &= \sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{j=1}^{N} \mathsf{u}(A_t) \mathsf{p}_{t,i}^{A_t,*} \mathbf{Q}_{i,j}^{t,u_t} l_{t,j} \mathbf{1}_{\tau_t < N} \\ &= \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{u}(A_t) \mathbf{Q}_{i,j}^{t,u_t} \mathsf{p}_{t,i}^{A_t} l_{t,j} \mathbf{1}_{\tau_t < N} + \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{u}(A_t) \mathbf{Q}_{i,j}^{t,u_t} (\mathsf{p}_{t,i}^{A_t,*} - \mathsf{p}_{t,i}^{A_t}) l_{t,j} \mathbf{1}_{\tau_t < N} \\ &\leq \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{u}(A_t) \mathbf{Q}_{i,j}^{t,u_t} \mathsf{p}_{t,i}^{A_t} l_{t,j} \mathbf{1}_{\tau_t < N} + \sum_{t=1}^{T} \mathsf{u}(A_t) \| \mathsf{p}_t^{A_t,*} - \mathsf{p}_t^{A_t} \|_1 \mathbf{1}_{\tau_t < N}. \end{split}$$

On the other hand, by design, if $\tau_t \ge N$, then $p_t = p_t^*$, so that

$$\sum_{t=1}^{T} \mathsf{u}(A_t) \underset{x_t \sim \mathsf{p}_t^{A_t}}{\mathbb{E}} [l_t(x_t)] \mathbf{1}_{\tau_t \ge N} \le \sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{u}(A_t) \mathbf{Q}_{i,j}^{t,u_t} \mathsf{p}_{t,i}^{A_t} l_{t,j} \mathbf{1}_{\tau_t \ge N}.$$

Thus, it follows that for any WFST $\mathfrak{T} \in \mathcal{T}$,

 σ

$$\begin{split} &\sum_{t=1}^{I} \mathsf{u}(A_t) \sum_{x_t \sim \mathsf{p}_t^{A_t}} [l_t(x_t)] \\ &\leq \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{u}(A_t) \mathbf{Q}_{i,j}^{t,u_t} \mathsf{p}_{t,i}^{A_t} l_{t,j} + 2 \sum_{t=1}^{T} \mathsf{u}(A_t) \| \mathsf{p}_t^{A_t,*} - \mathsf{p}_t^{A_t} \|_1 \mathbf{1}_{\tau_t < N} \end{split}$$

$$\begin{split} &= \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{j=1}^{N} \sum_{u \in Q_{T}} \mathsf{u}(A_{t}) \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{\delta_{T}(I_{T},x_{1:t-1})=u} \mathsf{p}_{t,i}^{A_{t}} l_{t,j} + 2 \sum_{t=1}^{T} \mathsf{u}(A_{t}) \| \mathsf{p}_{t}^{A_{t,*}} - \mathsf{p}_{t}^{A_{t}} \|_{1} \mathbf{1}_{\tau_{t} < N} \\ &= \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[u]]} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathsf{u}(A_{t}) \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{\delta_{T}(I_{T},x_{1:t-1})=u} \mathsf{p}_{t,i}^{A_{t}} l_{t,j} \\ &+ 2 \sum_{t=1}^{T} \mathsf{u}(A_{t}) \| \mathsf{p}_{t}^{A_{t,*}} - \mathsf{p}_{t}^{A_{t}} \|_{1} \mathbf{1}_{\tau_{t} < N} \\ &\leq \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[u]]} \sum_{j \in A_{t}} \sum_{u_{j}^{q,i}=u(A_{t})} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{1}_{\delta_{T}(I_{T},x_{1:t-1})=u} \mathsf{u}_{j}^{q,i} \mathbf{1}_{j \in A_{t}} \mathsf{p}_{t,i}^{A_{t}} l_{t,j} \\ &+ 2 \sum_{t=1}^{T} \mathsf{u}(A_{t}) \| \mathsf{p}_{t}^{A_{t,*}} - \mathsf{p}_{t}^{A_{t}} \|_{1} \mathbf{1}_{\tau_{t} < N} + \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{sleep}}) \\ &\leq \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[u]]} \sum_{e \in \mathsf{E}_{T}[q]} \sum_{t=1}^{T} \sum_{j=1}^{N} \mathbf{1}_{\delta_{T}(I_{T},x_{1:t-1})=u} \mathsf{u}_{j} \mathbf{1}_{j \in A_{t}} w[e] \mathsf{p}_{t,i}^{A_{t}} l_{t,j} \\ &+ 2 \sum_{t=1}^{T} \mathsf{u}(A_{t}) \| \mathsf{p}_{t}^{A_{t,*}} - \mathsf{p}_{t}^{A_{t}} \|_{1} \mathbf{1}_{\tau_{t} < N} + \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[q]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{sleep}}) \\ &\leq \sum_{u \in Q_{T}} \sum_{t=1}^{T} \mathsf{u}(A_{t}) \| \mathsf{p}_{t}^{A_{t,*}} - \mathsf{p}_{t}^{A_{t}} \|_{1} \mathbf{1}_{\tau_{t} < N} + \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[q]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{sleep}}) \\ &= \sum_{t=1}^{T} \sum_{u \in \mathcal{A}_{T}} \mathsf{w}_{u}(A_{t}) \| \mathsf{p}_{t}^{A_{t,*}} - \mathsf{p}_{t}^{A_{t}} \|_{1} \mathbf{1}_{\tau_{t} < N} + \sum_{u \in Q_{T}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{T}[q]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{sleep}}). \end{aligned}$$

For any distribution $u^* \in \Delta_N$ and awake sequence A_1^T , Let $L_T^{\mathbf{u},A_1^T} = \sum_{t=1}^T \mathbf{u}^*(A_t) \mathbb{E}_{x_t \sim \mathbf{p}_t}[l_t(x_t)]$, Thus, algorithm $\mathcal{A}_{u,i}$ achieves a regret in $O(\sqrt{L_T^{\mathbf{u}_i^{q,*},A_1^T}\log(N)})$, where $\mathbf{u}_i^{q,*}$ is a maximizer of algorithm $\mathcal{A}_{u,i}$'s sleeping regret.

Since the losses attributed to algorithm $\mathcal{A}_{u,i}$ are scaled by $1_{\delta_{\mathcal{T}}(I_{\mathcal{T}},x_{1:t-1})=u}\mathsf{p}_{t,i}^{A_t}$, it follows that at each round, the sum of the losses over all the algorithms is at most 1. Thus, by Jensen's inequality, it follows that

$$\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{sleep}}) \\
= \frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \sqrt{L_{T}^{\mathsf{u}_{i}^{q,*}, A_{1}^{T}}(\mathcal{A}_{u,i}) \log(N)} \\
\leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} L_{T}^{\mathsf{u}, A_{1}^{T}}(\mathcal{A}_{u,i}) \log(N)} \\
\leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]|} \sum_{t=1}^{T} \mathsf{u}(A_{t}) \log(N)},$$

so that $\sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \operatorname{Reg}_{T}(\mathcal{A}_{u,i}, \Phi_{\mathsf{sleep}}) \leq \sqrt{\sum_{t=1}^{T} \mathsf{u}(A_{t}) \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[u]]} \log(N)}.$

Finally, during the rounds in which $1_{\tau_t < N}$, p_t is an RPM approximation of p_t^* using τ_t iterations. Thus, by Equation 3.7 in [Nesterov and Nemirovski, 2015] the following inequality holds: $||p_t - p_t^*||_1 \le (1 - \alpha_t)^{\tau_t}$. Since τ_t is chosen so that the inequality $(1 - \alpha_t)^{\tau_t} \le 1/\sqrt{t}$ holds, it follows that $\sum_{t=1}^{T} u(A_t) \| \mathbf{p}_t^{A_t} - \mathbf{p}_t^{A_t,*} \|_1 \leq \sqrt{T}$, which proves the regret bound

$$\begin{split} & \sum_{t=1}^{T} \mathsf{u}(A_t) \mathop{\mathbb{E}}_{x_t \sim \mathsf{p}_t^{A_t}}[l_t(x_t)] - \sum_{t=1}^{T} \mathop{\mathbb{E}}_{x_t \sim \mathsf{p}_t^{A_t}} \left[\sum_{e \in \mathsf{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_t]} (\mathsf{u}|_{A_t})_{\mathsf{olab}[e]} w[e] l_t(\mathsf{olab}[e]) \right] \\ & \leq O\left(\sqrt{\sum_{t=1}^{T} \mathsf{u}(A_t)} \sum_{q \in Q_{\Phi}} \sum_{i \in \mathsf{ilab}[\mathsf{E}_{\mathcal{T}}[q]]} \log(N) \right). \end{split}$$

Furthermore, the computational cost of the *t*-th iteration of the algorithm is dominated by τ_t matrix multiplications or the solution of the linear system. τ_t can be bounded as follows: $\tau_t = \left\lceil \frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha_t)} \right\rceil \le \frac{\log\left(\frac{1}{\sqrt{t}}\right)}{\log(1-\alpha)} + 1$. Thus, the computational cost of the *t*-th iteration is in $O\left(N^2 \min\left\{\frac{\log t}{\log(1-\alpha_t)}, N\right\}\right) < O\left(N^2 \min\left\{\frac{\log T}{\log(1-\alpha_t)}, N\right\}\right).$

$$O\left(N^2 \min\left\{\frac{\log t}{\log(1/(1-\alpha_t))}, N\right\}\right) \le O\left(N^2 \min\left\{\frac{\log T}{\log(1/(1-\alpha))}, N\right\}\right).$$

H Connections with game-theoretic equilibria

There is an elegant connection between regret minimization in online learning and convergence to game-theoretic equilibria in repeated games [Nisan et al., 2007]. As an example, remarkably, if all players in a repeated game follow a swap regret minimization algorithm, then the empirical distribution of their play converges to a correlated equilibrium (see for example [Blum and Mansour, 2007]). Similarly, if all players follow a conditional swap regret minimization algorithm, then the empirical distribution of their play converges to a conditional swap regret minimization algorithm, then the empirical distribution of their play converges to a conditional correlated equilibrium [Mohri and Yang, 2014]. Hazan and Kale [2008] showed a result generalizing this property to the case of a Φ -regret and Φ -equilibrium. Moreover, the authors showed that the existence of an efficient Φ -regret minimizing algorithm is equivalent to the possibility of efficiently computing a fixed point associated to Φ -regret. However, their characterization of efficiency is a per iteration time complexity of $O(|\Phi|)$, which may be very large, in fact exponential in the number of experts, as in the case of the examples discussed in this paper. Here, we proved the existence of a large class of Φ -equilibria, *transductive equilibria*, i.e. those induced by a WFST, that are realizable in time that is polynomial in the number of experts.

I Lower bound

Auer [2017] proved a lower bound of $\Omega(\sqrt{TN})$ for swap regret. Since swap regret is a special case of transductive regret, that lower bound applies to the setting of transductive regret as well. This is further detailed in an extended version of this paper.

J Bandit setting

Blum and Mansour [2007] and Mohri and Yang [2014] respectively showed that swap and conditional swap regret-minimizing algorithms can be extended to the bandit setting. Similarly, our more general transductive regret-minimizing can be extended to the bandit setting, as shown and detailed in the extended version of this paper.