

## A Additional figures and examples

### A.1 Special cases of transductive regret.

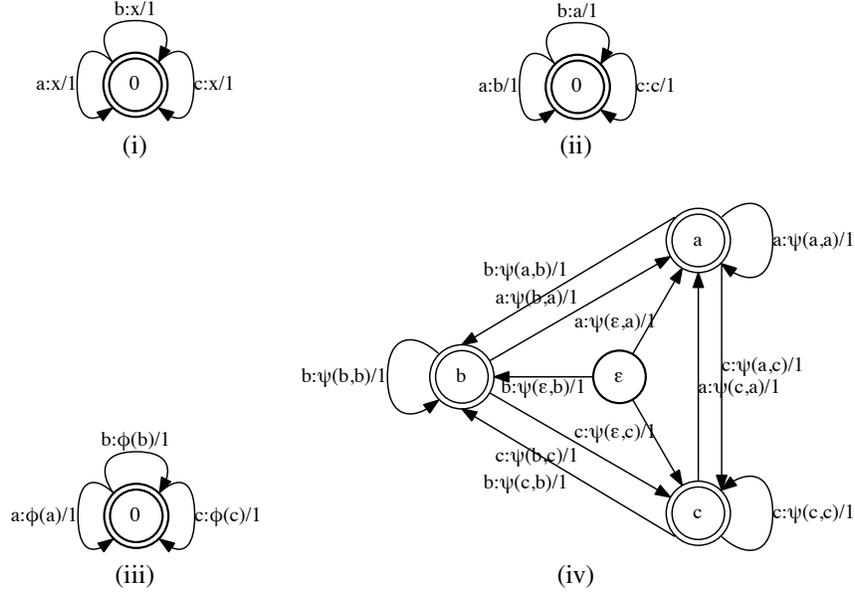


Figure 4: Several families of WFSTs for special cases of transductive regret for  $\Sigma = \{a, b, c\}$ . (i) External regret with parameter  $x \in \Sigma$ . (ii) Internal regret: family of transducers  $\mathcal{T}_{a_1, a_2}$  with  $a_1 \neq a_2$ ,  $a_1, a_2 \in \Sigma$ ; example shown for  $\mathcal{T}_{a,b}$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

## A.2 Example with a swapping subset.

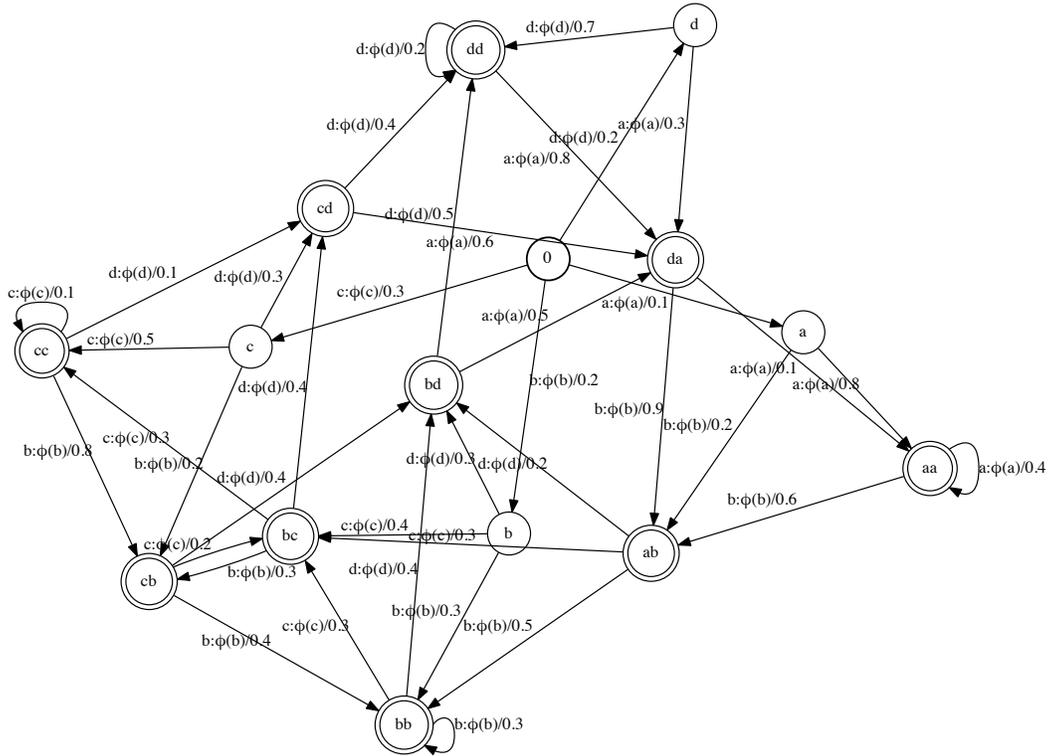


Figure 5: Example of a WFST with  $\Sigma = \{a, b, c, d\}$  and where each state has a swapping subset.

## B Pseudocode of FASTTRANSDUCE

---

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

---

**Algorithm:** FASTTRANSDUCE( $\mathcal{T}, (\mathcal{A}_{u,i})_{u \in Q_{\mathcal{T}}, i \in \text{ilab}[E_{\mathcal{T}}[u]]}$ )

```

 $u \leftarrow I_{\mathcal{T}}$ 
for  $t \leftarrow 1$  to  $T$  do
  for each  $i \in \text{ilab}[E_{\mathcal{T}}[u]]$  do
     $q_i \leftarrow \text{QUERY}(\mathcal{A}_{u,i})$ 
   $\mathbf{Q}^{t,u} \leftarrow [q_1 \mathbf{1}_{1 \in \text{ilab}[E_{\mathcal{T}}[u]]} \cdots q_N \mathbf{1}_{N \in \text{ilab}[E_{\mathcal{T}}[u]]}]^{\top}$ 
  for each  $j \leftarrow 1$  to  $N$  do
     $c_j \leftarrow \min_{i \in \text{ilab}[E_{\mathcal{T}}[u]]} \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{j \in \text{ilab}[E_{\mathcal{T}}[u]]}$ 
   $\alpha_t \leftarrow \|\mathbf{c}\|_1; \quad \tau_t \leftarrow \left\lceil \frac{\log\left(\frac{1}{\alpha_t}\right)}{\log(1-\alpha_t)} \right\rceil$ 
  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
       $(\mathbf{p}_t^{\tau})^{\top} \leftarrow (\mathbf{p}_t^{\tau-1})^{\top} (\mathbf{Q}^{t,u} - \vec{\mathbf{1}} \mathbf{c}^{\top}); \mathbf{p}_t \leftarrow \mathbf{p}_t + \mathbf{p}_t^{\tau}$ 
       $\mathbf{p}_t \leftarrow \frac{\mathbf{p}_t}{\|\mathbf{p}_t\|_1}$ 
  else
     $\mathbf{p}_t^{\top} = \text{FIXED-POINT}(\mathbf{Q}^{t,u})$ 
   $x_t \leftarrow \text{SAMPLE}(\mathbf{p}_t); \quad \mathbf{l}_t \leftarrow \text{RECEIVELOSS}(); \quad u \leftarrow \delta_{\mathcal{T}}(u, x_t)$ 
  for each  $i \in \text{ilab}[E_{\mathcal{T}}[u]]$  do
     $\text{ATTRIBUTELOSS}(\mathcal{A}_{u,i}, \mathbf{p}_t[i] \mathbf{l}_t)$ 

```

---

## C Pseudocode of FASTSLEEPTRANSDUCE

---

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

---

**Algorithm:** FASTSLEEPTRANSDUCE( $\mathcal{T}, \{\mathcal{A}_{u,i}\}_{u \in Q_{\mathcal{T}}, i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]}$ )

$u \leftarrow I_{\mathcal{T}}$

**for**  $t \leftarrow 1$  **to**  $T$  **do**

$A_t \leftarrow \text{AWAKESET}()$

**for each**  $i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]] \cap A_t$  **do**

$\mathbf{q}_i \leftarrow \text{QUERY}(\mathcal{A}_{u,i}); \quad \mathbf{q}_i^{A_t} \leftarrow \frac{\mathbf{q}_i|_{A_t}}{\sum_{j \in A_t} \mathbf{q}_j}$

$\mathbf{Q}^{t,u} \leftarrow [\mathbf{q}_1^{A_t} \mathbf{1}_{1 \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]] \cap A_t}; \dots; \mathbf{q}_N^{A_t} \mathbf{1}_{N \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]] \cap A_t}]$

**for each**  $j \leftarrow 1$  **to**  $N$  **do**

$c_j \leftarrow \min_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]] \cap A_t} \mathbf{Q}_{i,j}^{t,u} \mathbf{1}_{j \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]] \cap A_t}$

$\alpha_t \leftarrow \|\mathbf{c}\|_1; \quad \tau_t \leftarrow \left\lceil \frac{\log\left(\frac{1}{\alpha_t}\right)}{\log(1-\alpha_t)} \right\rceil$

**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**

$(\mathbf{p}_t^\tau)^\top \leftarrow (\mathbf{p}_t^\tau)^\top (\mathbf{Q}^{t,u} - [\mathbf{1}_{1 \in A_t}; \dots; \mathbf{1}_{|\text{ilab}[\mathbb{E}_{\mathcal{T}}[q]|] \in A_t}] \mathbf{c}^\top)$

$\mathbf{p}_t \leftarrow \mathbf{p}_t + \mathbf{p}_t^\tau$

$\mathbf{p}_t \leftarrow \frac{\mathbf{p}_t}{\|\mathbf{p}_t\|_1}$

**else**

$\mathbf{p}_t^\top \leftarrow \text{FIXED-POINT}(\mathbf{Q}^{t,u})$

$\mathbf{p}_t^{A_t} \leftarrow \frac{\mathbf{p}_t|_{A_t}}{\sum_{j \in A_t} \mathbf{p}_t^{t,j}}; \quad x_t \leftarrow \text{SAMPLE}(\mathbf{p}_t^{A_t}); \quad \mathbf{l}_t \leftarrow \text{RECEIVELOSS}(); \quad u \leftarrow \delta_{\mathcal{T}}[u, x_t]$

**for each**  $i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]$  **do**

$\text{ATTRIBUTELOSS}(\mathcal{A}_{u,i}, \mathbf{p}_t[i] \mathbf{l}_t)$

---

## D Proof of Theorem 1

**Theorem 1.** Let  $\mathcal{A}_1, \dots, \mathcal{A}_N$  be external regret minimizing algorithms admitting data-dependent regret bounds of the form  $O(\sqrt{L_T(\mathcal{A}_i)} \log N)$ , where  $L_T(\mathcal{A}_i)$  is the cumulative loss of  $\mathcal{A}_i$  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  $\mathbf{p}_t$  be the distribution returned by FASTSWAP at round  $t$ . For any distribution  $\mathbf{p}_t^*$ ,  $t \in [T]$ , the following inequality holds:

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] \mathbf{1}_{\tau_t < N} &= \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] \mathbf{1}_{\tau_t < N} + \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] \mathbf{1}_{\tau_t < N} \\ &\quad - \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] \mathbf{1}_{\tau_t < N} \\ &\leq \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] \mathbf{1}_{\tau_t < N} + \sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \|l_t\|_\infty \mathbf{1}_{\tau_t < N} \\ &\leq \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] \mathbf{1}_{\tau_t < N} + \sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \mathbf{1}_{\tau_t < N}. \end{aligned}$$

Let  $\mathbf{p}_t^*$  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

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

On the other hand, by design, if  $\tau_t \geq N$ , then  $\mathbf{p}_t = \mathbf{p}_t^*$ , so that

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

Thus, it follows that

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] &\leq \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^N \mathbf{Q}_{i,j}^t \mathbf{p}_{t,i} l_{t,j} + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 \mathbf{1}_{\tau_t < N} \\ &\leq \sum_{i=1}^N \left[ \min_{j \in [N]} \sum_{t=1}^T \mathbf{p}_{t,i} l_{t,j} + \text{Reg}_T(\mathcal{A}_i, \Phi_{\text{ext}}) \right] + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 \mathbf{1}_{\tau_t < N} \\ &= \min_{\varphi \in \Phi_{\text{swap}}} \sum_{i=1}^N \left[ \sum_{t=1}^T \mathbf{p}_{t,i} l_{t,\varphi(i)} + \text{Reg}_T(\mathcal{A}_i, \Phi_{\text{ext}}) \right] + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 \mathbf{1}_{\tau_t < N}. \end{aligned}$$

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:

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \text{Reg}_T(\mathcal{A}_i, \Phi_{\text{ext}}) &= \frac{1}{N} \sum_{i=1}^N O(\sqrt{L_T(\mathcal{A}_i) \log N}) \\ &\leq O\left(\sqrt{\frac{1}{N} \sum_{i=1}^N L_T(\mathcal{A}_i) \log N}\right) \leq O\left(\sqrt{\frac{T \log N}{N}}\right), \end{aligned}$$

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

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

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

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

□

## E Proof of Theorem 2

**Theorem 2.** Let  $(\mathcal{A}_{u,i})_{u \in Q, i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]}$  be external regret minimizing algorithms admitting data-dependent 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  $O(\sqrt{T} |\mathbb{E}_{\mathcal{T}}| \log N)$  with a per-iteration complexity in  $O\left(N^2 \min\left\{\frac{\log T}{\log(1/(1-\alpha))}, N\right\}\right)$ .

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

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] 1_{\tau_t < N} &= \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] 1_{\tau_t < N} + \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] 1_{\tau_t < N} \\ &\quad - \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] 1_{\tau_t < N} \\ &\leq \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] 1_{\tau_t < N} + \sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \|l_t\|_{\infty} 1_{\tau_t < N} \\ &\leq \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t^*} [l_t(x_t)] 1_{\tau_t < N} + \sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 1_{\tau_t < N}. \end{aligned}$$

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 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:

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

On the other hand, by design, if  $\tau_t \geq N$ , then  $\mathbf{p}_t = \mathbf{p}_t^*$ , so that

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

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

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] &\leq \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^N \sum_{u \in Q_{\mathcal{T}}} \mathbf{Q}_{i,j}^{t,u} 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i} l_{t,j} + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 1_{\tau_t < N} \\ &= \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]} \sum_{t=1}^T \sum_{j=1}^N \mathbf{Q}_{i,j}^{t,u} 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i} l_{t,j} \end{aligned}$$

$$\begin{aligned}
& + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 1_{\tau_t < N} \\
& \leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]} \min_{i^* \in \text{olab}[\mathbb{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_t]]} \sum_{t=1}^T 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i} l_{t,i^*} \\
& + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 1_{\tau_t < N} + \sum_{i=1}^N \sum_{u \in Q_{\mathcal{T}}} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{ext}}) \\
& \leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[q]]} \sum_{e \in \mathbb{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_t]} \sum_{t=1}^T 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i} w[e] l_t(\text{olab}[e]) \\
& + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 1_{\tau_t < N} + \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[q]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{ext}}) \\
& = \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} \left[ \sum_{e \in \mathbb{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_t]} w[e] l_t(\text{olab}[e]) \right] + 2 \sum_{t=1}^T \|\mathbf{p}_t^* - \mathbf{p}_t\|_1 1_{\tau_t < N} \\
& + \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[q]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{ext}}).
\end{aligned}$$

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  $1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i}$ , 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

$$\begin{aligned}
& \frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{ext}}) \\
& = \frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]} \sqrt{L_T(\mathcal{A}_{u,i}) \log(N)} \\
& \leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]} L_T(\mathcal{A}_{u,i}) \log(N)} \\
& \leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathbb{E}_{\mathcal{T}}[u]]|} T \log(N)},
\end{aligned}$$

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

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

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

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

□

### F Proof of Theorem 3

**Theorem 3.** Let  $(\mathcal{A}_{I,u,i})_{I \in \mathcal{I}, u \in \mathcal{Q}_T, i \in \text{lab}[\mathbb{E}_{\mathcal{T}}[q]]}$  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, FASTTIMSELECTTRANSDUCE achieves a time-selection transductive regret with respect to the time-selection family  $\mathcal{I}$  and WFST family  $\mathcal{T}$  that is in  $O\left(\sqrt{T(\log(|\mathcal{I}|) + |\mathbb{E}_{\mathcal{T}}| \log N)}\right)$  with a per-iteration complexity in  $O\left(N^2\left(\min\left\{\frac{\log(T)}{\log((1-\alpha)^{-1})}, N\right\} + |\mathcal{I}|\right)\right)$ .

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

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

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  $\mathbf{Q}^{t,u_t}$ . Then, by expanding the definition of  $\tilde{\mathbf{I}}^t$ , we can rewrite the expression on the left-hand side as

$$\begin{aligned} \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t \tilde{l}_I^t \mathbf{1}_{\tau_t < N} &= \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^\top \mathbf{M}^{t,u_t,I} \mathbf{1}_t - \mathbf{p}_t^\top \mathbf{1}_t) \mathbf{1}_{\tau_t < N} \\ &= \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) \mathbf{p}_t^\top \mathbf{M}^{t,u_t,I} \mathbf{1}_t \mathbf{1}_{\tau_t < N} - \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) \mathbf{p}_t^\top \mathbf{1}_t \mathbf{1}_{\tau_t < N} \\ &\geq \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{M}^{t,u_t,I} \mathbf{1}_t \mathbf{1}_{\tau_t < N} - \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{1}_t \mathbf{1}_{\tau_t < N} \\ &\quad - \sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \mathbf{1}_{\tau_t < N}. \end{aligned}$$

On the other hand, by design, if  $\tau_t \geq N$ , then  $\mathbf{p}_t = \mathbf{p}_t^*$ , so that

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

Thus, it follows that

$$\sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t \tilde{l}_I^t \geq \sum_{t=1}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{M}^{t,u_t,I} \mathbf{1}_t - \sum_{t=0}^T \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{1}_t - \sum_{t=1}^T \|\mathbf{p}_t - \mathbf{p}_t^*\|_1 \mathbf{1}_{\tau_t < N}.$$

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

$$\sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{M}^{t,u_t,I} \mathbf{1}_t = \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{1}_t.$$

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

$$\sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{M}^{t,u_t,I} \mathbf{1}_t = \sum_{I \in \mathcal{I}} \tilde{\mathbf{q}}_I^t I(t) (\mathbf{p}_t^*)^\top \mathbf{1}_t = 0,$$

which implies that

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

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

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

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

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

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

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

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

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

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

and as in Theorem 2, we can bound the  $l_1$  approximation error of  $\mathbf{p}_t$  for  $\mathbf{p}_t^*$  by

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

by the algorithm's choice of  $\tau_t$ . 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  $O\left(\sqrt{T\left(\log(|\mathcal{I}|) + \sum_{q \in \mathcal{Q}_{\Phi}} |\text{lab}[\mathbf{E}_{\mathcal{I}}[q]]| \log(N)\right)}\right)$ .

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

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

□

## 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  $O\left(\sqrt{\sum_{t=1}^T u^*(A_t) \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)] \log(N)}\right)$ . 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  $\text{Reg}_T(\text{FASTSLEEPTRANSDUCE}, \mathcal{T}, A_1^T)$  of FASTSLEEPTRANSDUCE is upper bounded by  $O\left(\sqrt{\sum_{t=1}^T u(A_t) |\mathbb{E}_{\mathcal{T}} \log(N)}\right)$ . 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  $\mathbf{p}_t^{A_t}$  be the distribution output by FASTSLEEPTRANSDUCE at round  $t$ . For any distribution  $\mathbf{p}_t^*$ ,  $t \in [T]$ , the following inequalities hold:

$$\begin{aligned} u(A_t) \mathbb{E}_{x_t \sim \mathbf{p}_t^{A_t}} [l_t(x_t)] \mathbf{1}_{\tau_t < N} &= u(A_t) \left( \mathbb{E}_{x_t \sim \mathbf{p}_t^{A_t, *}} [l_t(x_t)] + \mathbb{E}_{x_t \sim \mathbf{p}_t^{A_t}} [l_t(x_t)] - \mathbb{E}_{x_t \sim \mathbf{p}_t^{A_t, *}} [l_t(x_t)] \right) \mathbf{1}_{\tau_t < N} \\ &\leq u(A_t) \left( \mathbb{E}_{x_t \sim \mathbf{p}_t^{A_t, *}} [l_t(x_t)] + \|\mathbf{p}_t^{A_t} - \mathbf{p}_t^{A_t, *}\|_1 \|l_t\|_\infty \right) \mathbf{1}_{\tau_t < N} \\ &\leq u(A_t) \left( \mathbb{E}_{x_t \sim \mathbf{p}_t^{A_t, *}} [l_t(x_t)] + \|\mathbf{p}_t^{A_t} - \mathbf{p}_t^{A_t, *}\|_1 \right) \mathbf{1}_{\tau_t < N}. \end{aligned}$$

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:

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

On the other hand, by design, if  $\tau_t \geq N$ , then  $\mathbf{p}_t = \mathbf{p}_t^*$ , so that

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

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

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

$$\begin{aligned}
&= \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^N \sum_{u \in Q_{\mathcal{T}}} u(A_t) \mathbf{Q}_{i,j}^{t,u} 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i}^{A_t} l_{t,j} + 2 \sum_{t=1}^T u(A_t) \|\mathbf{p}_t^{A_t,*} - \mathbf{p}_t^{A_t}\|_1 1_{\tau_t < N} \\
&= \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} \sum_{t=1}^T \sum_{j=1}^N u(A_t) \mathbf{Q}_{i,j}^{t,u} 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i}^{A_t} l_{t,j} \\
&\quad + 2 \sum_{t=1}^T u(A_t) \|\mathbf{p}_t^{A_t,*} - \mathbf{p}_t^{A_t}\|_1 1_{\tau_t < N} \\
&\leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} \min_{\substack{u^{q,i} \in \Delta_N \\ \sum_{j \in A_t} u_j^{q,i} = u(A_t)}} \sum_{t=1}^T \sum_{j=1}^N 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} u_j^{q,i} 1_{j \in A_t} \mathbf{p}_{t,i}^{A_t} l_{t,j} \\
&\quad + 2 \sum_{t=1}^T u(A_t) \|\mathbf{p}_t^{A_t,*} - \mathbf{p}_t^{A_t}\|_1 1_{\tau_t < N} + \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{sleep}}) \\
&\leq \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} \sum_{e \in \mathcal{E}_{\mathcal{T}}[q]} \sum_{t=1}^T \sum_{j=1}^N 1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} u_j 1_{j \in A_t} w[e] \mathbf{p}_{t,i}^{A_t} l_{t,j} \\
&\quad + 2 \sum_{t=1}^T u(A_t) \|\mathbf{p}_t^{A_t,*} - \mathbf{p}_t^{A_t}\|_1 1_{\tau_t < N} + \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[q]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{sleep}}) \\
&= \sum_{t=1}^T \mathbb{E}_{x_t \sim \mathbf{p}_t} \left[ \sum_{e \in \mathcal{E}_{\mathcal{T}}[\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1}), x_t]} (u|_{A_t})_{\text{olab}[e]} w[e] \mathbf{p}_{t,i}^{A_t} l_{t,j} (\text{olab}[e]) \right] \\
&\quad + 2 \sum_{t=1}^T u(A_t) \|\mathbf{p}_t^{A_t,*} - \mathbf{p}_t^{A_t}\|_1 1_{\tau_t < N} + \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[q]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{sleep}}).
\end{aligned}$$

For any distribution  $\mathbf{u}^* \in \Delta_N$  and awake sequence  $A_1^T$ , Let  $L_T^{u, A_1^T} = \sum_{t=1}^T u^*(A_t) \mathbb{E}_{x_t \sim \mathbf{p}_t} [l_t(x_t)]$ , Thus, algorithm  $\mathcal{A}_{u,i}$  achieves a regret in  $O(\sqrt{L_T^{u_i^{q,*}, A_1^T} \log(N)})$ , where  $u_i^{q,*}$  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  $1_{\delta_{\mathcal{T}}(I_{\mathcal{T}}, x_{1:t-1})=u} \mathbf{p}_{t,i}^{A_t}$ , 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

$$\begin{aligned}
&\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} \text{Reg}_T(\mathcal{A}_{u,i}, \Phi_{\text{sleep}}) \\
&= \frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} \sqrt{L_T^{u_i^{q,*}, A_1^T}(\mathcal{A}_{u,i}) \log(N)} \\
&\leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]|} \sum_{u \in Q_{\mathcal{T}}} \sum_{i \in \text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]} L_T^{u, A_1^T}(\mathcal{A}_{u,i}) \log(N)} \\
&\leq \sqrt{\frac{1}{\sum_{u \in Q_{\mathcal{T}}} |\text{ilab}[\mathcal{E}_{\mathcal{T}}[u]]|} \sum_{t=1}^T u(A_t) \log(N)},
\end{aligned}$$

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

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

$\sum_{t=1}^T \mathbf{u}(A_t) \|\mathbf{p}_t^{A_t} - \mathbf{p}_t^{A_t,*}\|_1 \leq \sqrt{T}$ , which proves the regret bound

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

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

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

□

## 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 correlated equilibrium [Mohri and Yang, 2014]. Hazan and Kale [2008] showed a result generalizing this property to the case of a  $\Phi$ -regret and  $\Phi$ -equilibrium. Moreover, the authors showed that the existence of an efficient  $\Phi$ -regret minimizing algorithm is equivalent to the possibility of efficiently computing a fixed point associated to  $\Phi$ -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  $\Phi$ -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.