Topology Matters: Smoothed Competitiveness of Metrical Task Systems*

Guido Schäfer and Naveen Sivadasan

Max-Planck-Institut für Informatik, Saarbrücken, Germany,
{schaefer,ns}@mpi-sb.mpg.de

Abstract. We consider *metrical task systems*, a general framework to model online problems. Borodin, Linial and Saks [BLS92] presented a deterministic *work function algorithm* (WFA) for metrical task systems having a tight competitive ratio of $2n-1$. We present a *smoothed competitive analysis* of WFA. Given an adversarial task sequence, we smoothen the request costs by means of a symmetric additive smoothing model and analyze the competitive ratio of WFA on the smoothed task sequence. We prove upper and matching lower bounds on the smoothed competitive ratio of WFA. Our analysis reveals that the smoothed competitive ratio of WFA is much better than $O(n)$ and that it depends on several topological parameters of the underlying graph G, such as the maximum degree D and the diameter. For example, already for small perturbations the smoothed competitive ratio of WFA reduces to $O(\log n)$ on a clique or a complete binary tree and to $O(\sqrt{n})$ on a line. We also provide the first average case analysis of WFA showing that its expected competitive ratio is $O(\log(D))$ for various distributions.

1 Introduction

Borodin, Linial and Saks [BLS92] introduced a general framework to model online problems, called *metrical task systems*. Many important online problems can be formulated as metrical task systems; for example, the paging problem, the static list accessing problem and the k-server problem.

We are given an undirected and connected graph $G = (V, E)$, with node set V and edge set E, and a positive length function $\lambda : E \to \mathbb{R}^+$ on the edges of G. We extend λ to a metric δ on G. Let $\delta : V \times V \to \mathbb{R}^+$ be a distance function such that $\delta(u, v)$ denotes the shortest path distance (with respect to λ) between any two nodes u and v in G. A *task* τ is an n-vector $(r(v_1), \ldots, r(v_n))$ of *request costs*. The cost to process task τ in node v_i is $r(v_i) \in \mathbb{R}^+ \cup \{\infty\}$. The online algorithm starts from a given initial position $s_0 \in V$ and has to service a sequence $\mathcal{S} = \langle \tau_1, \ldots, \tau_r \rangle$ of tasks, arriving one at a time. If the online algorithm resides after task τ_{t-1} in node u, the cost to service task τ_t in node v is $\delta(u, v) + r_t(v)$; $\delta(u, v)$ is the *transition cost* and $r_t(v)$ is the *processing cost*. The objective is to minimize the total transition plus processing cost.

* Partially supported by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

Upper Bounds	
random tasks	$O\big(\frac{\sigma}{U_{\min}}\big(\frac{U_{\min}}{\sigma}+\log(D)\big)\big)$
arbitrary tasks	$O\big(\frac{Diam}{U_{\min}}\cdot\big(\frac{U_{\min}}{\sigma}+\log(D)\big)\big)$ and $O\left(\sqrt{n\cdot\frac{U_{\max}}{U_{\min}}\big(\frac{U_{\min}}{\sigma}+\log(D)\big)}\right)$
β-elementary tasks	$O\big(\beta\cdot\frac{U_{\max}}{U_{\min}}\big(\frac{U_{\min}}{\sigma}+\log(D)\big)\big)$

Table 1. Upper bounds on the competitive ratio of WFA.

Borodin, Linial and Saks [BLS92] gave a deterministic online algorithm, known as the *work function algorithm* (WFA), for metrical task systems. WFA has a competitive ratio of $2n-1$, which is optimal. However, the competitive ratio is often an over-pessimistic estimation of the true performance of an online algorithm.

Based on the idea underlying *smoothed analysis* [ST01], Becchetti et al. [BLMS$^+$03] recently proposed *smoothed competitive analysis* as an alternative to worst case competitive analysis of online algorithms. The idea is to randomly perturb, or *smoothen*, an adversarial input instance $\check{\mathcal{S}}$ and to analyze the performance of the algorithm on the perturbed instances. Let $\text{ALG}[\mathcal{S}]$ and $\text{OPT}[\mathcal{S}]$, respectively, be the cost of the online and the optimal offline algorithm on a smoothed instance $\mathcal{S}$ obtained from $\check{\mathcal{S}}$. The *smoothed competitive ratio* c of ALG with respect to a smoothing distribution f is defined as

$$c := \sup_{\check{\mathcal{S}}} \mathbf{E}_{\mathcal{S}\overset{f}{\leftarrow}\check{\mathcal{S}}}\left[\frac{\text{ALG}[\mathcal{S}]}{\text{OPT}[\mathcal{S}]}\right].$$

We use the notion of smoothed competitiveness to characterize the asymptotic performance of WFA. We smoothen the request costs of each task according to an additive symmetric smoothing model. Each cost entry is smoothed by adding a random number chosen from a probability distribution f, whose expectation coincides with the original cost entry. Our analysis holds for various probability distributions, including the uniform and the normal distribution. We use σ to refer to the standard deviation of f. Our analysis reveals that the smoothed competitive ratio of WFA is much better than its worst case competitive ratio suggests and that it depends on certain *topological parameters* of the underlying graph.

Definition of Topological Parameters. In this paper, we assume that the underlying graph G has n nodes, minimum edge length $U_{\min}$, maximum edge length $U_{\max}$, and maximum degree D. Furthermore, we use *Diam* to refer to the *diameter* of G, i.e., the maximum length of a shortest path between any two nodes. Similarly, a graph has *edge* diameter *diam* if any two nodes are connected by a path of at most *diam* edges. Observe that $diam U_{\min} \leq Diam \leq diam U_{\max}$.

Lower Bounds	
arbitrary tasks	
– existential	$\Omega\left(\frac{Diam}{U_{\min}} \cdot \left(\frac{U_{\min}}{\sigma} + \log(D)\right)\right)$ and $\Omega\left(\sqrt{n \cdot \frac{U_{\max}}{U_{\min}} \cdot \left(\frac{U_{\min}}{\sigma} + \log(D)\right)}\right)$
– universal	$\Omega\left(\frac{U_{\min}}{\sigma} + \frac{U_{\min}}{U_{\max}} \log(D)\right)$ and $\Omega\left(\sqrt{diam \cdot \frac{U_{\min}}{U_{\max}} \cdot \left(\frac{U_{\min}}{\sigma} + 1\right)}\right)$
β-elementary tasks	$\Omega\left(\beta \cdot \left(\frac{U_{\min}}{\sigma} + 1\right)\right)$ (existential)

Table 2. Lower bounds on the competitive ratio of any deterministic online algorithm.

We emphasize that these topological parameters are defined with respect to G and its length function λ—not with respect to the resulting metric.

We prove several upper bounds; see also Table 1.

1. We show that if the request costs are chosen randomly from a distribution f, which is non-increasing in $[0, \infty)$, the expected competitive ratio of WFA is
$$O\left(1 + \frac{\sigma}{U_{\min}} \cdot \log(D)\right).$$
In particular, WFA has an expected competitive ratio of $O(\log(D))$ if $\sigma = \Theta(U_{\min})$. For example, we obtain a competitive ratio of $O(\log(n))$ on a clique and of $O(1)$ on a binary tree.
2. We prove two upper bounds on the smoothed competitive ratio of WFA:
$$O\left(\frac{Diam}{U_{\min}} \cdot \left(\frac{U_{\min}}{\sigma} + \log(D)\right)\right) \quad \text{and} \quad O\left(\sqrt{n \cdot \frac{U_{\max}}{U_{\min}} \left(\frac{U_{\min}}{\sigma} + \log(D)\right)}\right).$$
For example, if $\sigma = \Theta(U_{\min})$ and $U_{\max}/U_{\min} = \Theta(1)$, WFA has smoothed competitive ratio $O(\log(n))$ on any constant diameter graph and $O(\sqrt{n})$ on any constant degree graph. Note also that on a complete binary tree we obtain an $O(\log(n))$ upper bound.
3. We obtain a better upper bound on the smoothed competitive ratio of WFA if the adversarial task sequence only consists of *β-elementary tasks*. A task is β-elementary if it has at most β non-zero entries. We prove a smoothed competitive ratio of
$$O\left(\beta \cdot \frac{U_{\max}}{U_{\min}} \left(\frac{U_{\min}}{\sigma} + \log(D)\right)\right).$$
For example, if $\sigma = \Theta(U_{\min})$ and $U_{\max}/U_{\min} = \Theta(1)$, WFA has smoothed competitive ratio $O(\beta \log(D))$ for β-elementary tasks.

We also present lower bounds; see Table 2. All our lower bounds hold for *any* deterministic online algorithm and if the request costs are smoothed according to the additive symmetric smoothing model. We distinguish between *existential* and *universal* lower bounds. An existential lower bound, say $\Omega(f(n))$, means that there *exists* a class of graphs such that *every* deterministic algorithm has smoothed competitive ratio $\Omega(f(n))$ on these graphs. On the other hand, a universal lower bound $\Omega(f(n))$ states that for *any arbitrary* graph, *every* deterministic algorithm has smoothed competitive ratio $\Omega(f(n))$. Clearly, for metrical task systems, the best lower bound we can hope to obtain is $\Omega(n)$. Therefore, if we state a lower bound of $\Omega(f(n))$, we actually mean $\Omega(\min\{n, f(n)\})$.

4. For a large range of values for *Diam* and D, we present existential lower bounds that are asymptotically tight to the upper bounds stated in 2.
5. We also prove two universal lower bounds on the smoothed competitive ratio:

$$\Omega\left(\tfrac{U_{\min}}{\sigma} + \tfrac{U_{\min}}{U_{\max}} \log(D)\right) \quad \text{and} \quad \Omega\left(\min\left\{ diam, \sqrt{diam \cdot \tfrac{U_{\min}}{U_{\max}} \cdot \left(\tfrac{U_{\min}}{\sigma} + 1\right)}\right\}\right).$$

   Assume that $U_{\max}/U_{\min} = \Theta(1)$. Then, the first bound matches the first upper bound stated in 2 if the edge diameter *diam* is constant, e.g., for a clique. The second bound matches the second upper bound in 2 if $diam = \Omega(n)$ and the maximum degree D is constant, e.g., for a line.
6. For β-elementary tasks, we prove an existential lower bound of

$$\Omega\left(\beta \cdot \left(\tfrac{U_{\min}}{\sigma} + 1\right)\right).$$

   This implies that the bound in 3 is tight up to a factor of $(U_{\max}/U_{\min}) \log(D)$.

Constrained Balls into Bins Game. Our analysis crucially relies on a lower bound on the cost of an optimal offline algorithm. We therefore study the growth of the work function values on a sequence of random requests. It turns out that the increase in the work function values can be modeled by a version of a balls into bins game with dependencies between the heights of the bins, which are specified by a constraint graph. We call this game the *constrained balls into bins game*. We believe that this game is also interesting independently of the context of this paper.

Due to lack of space, we omit the lower bounds and some upper bound proofs from this extended abstract. We refer the reader to [SS03] for a complete version of this paper.

2 Work Function Algorithm

Let $\mathcal{S} = \langle \tau_1, \ldots, \tau_\ell \rangle$ be a request sequence, and let $s_0 \in V$ denote the initial position. Let $\mathcal{S}_t$ denote the subsequence of the first t tasks of $\mathcal{S}$. For each t, $0 \le t \le \ell$, we define a function $w_t : V \to \mathbb{R}$ such that for each node $u \in V$, $w_t(u)$ is the minimum offline cost to process $\mathcal{S}_t$ starting in s_0 and ending in u. The function w_t is called the *work function* at time t with respect to $\mathcal{S}$ and s_0.

Let OPT denote an optimal offline algorithm. Clearly, the optimal offline cost OPT$[\mathcal{S}]$ on $\mathcal{S}$ is equal to the minimum work function value at time ℓ, i.e., OPT$[\mathcal{S}] = \min_{u \in V}\{w_\ell(u)\}$. We can compute $w_t(u)$ for each $u \in V$ by dynamic programming:

$$w_0(u) := \delta(s_0, u), \quad \text{and} \quad w_t(u) := \min_{v \in V}\{w_{t-1}(v) + r_t(v) + \delta(v, u)\}. \qquad (1)$$

We next describe the online work function algorithm; see also [BLS92,BEY98]. Intuitively, a good strategy for an online algorithm to process task τ_t is to move to a node where OPT would reside if τ_t would be the final task. However, the competitive ratio of an algorithm that solely sticks to this policy can become arbitrarily bad. A slight modification gives a $2n-1$ competitive algorithm: Instead of blindly (no matter at what cost) traveling to the node of minimum work function value, we additionally take the transition cost into account. Essentially, this is the idea underlying the work function algorithm.

Work Function Algorithm (WFA): Let $s_0, \dots, s_{t-1}$ denote the sequence of nodes visited by WFA to process $\mathcal{S}_{t-1}$. Then, to process task τ_t, WFA moves to a node s_t that minimizes $w_t(v) + \delta(s_{t-1}, v)$ for all $v \in V$. There is always a choice for s_t such that in addition $w_t(s_t) = w_{t-1}(s_t) + r_t(s_t)$. More formally,

$$s_t := \arg\min_{v \in V}\{w_t(v) + \delta(s_{t-1}, v)\} \quad \text{such that} \quad w_t(s_t) = w_{t-1}(s_t) + r_t(s_t). \quad (2)$$

In the sequel, we use WFA and OPT, respectively, to denote the work function and the optimal offline algorithm. For a given sequence $\mathcal{S} = \langle \tau_1, \dots, \tau_\ell \rangle$ of tasks, WFA$[\mathcal{S}]$ and OPT$[\mathcal{S}]$ refer to the cost incurred by WFA and OPT on $\mathcal{S}$, respectively. By $s_0, \dots, s_\ell$ we denote the sequence of nodes visited by WFA.

We state the following facts without proof.

Fact 1. *For any two nodes u and v and any time t, $|w_t(u) - w_t(v)| \leq \delta(u, v)$.*

Fact 2. *At any time t, $w_t(s_t) = w_t(s_{t-1}) - \delta(s_{t-1}, s_t)$.*

Fact 3. *At any time t, $r_t(s_t) + \delta(s_{t-1}, s_t) = w_t(s_{t-1}) - w_{t-1}(s_t)$.*

3 Smoothing Model

Let the *adversarial task sequence* be given by $\check{\mathcal{S}} := \langle \check{\tau}_1, \dots, \check{\tau}_r \rangle$. We smoothen each task vector $\check{\tau}_t := (\check{r}_t(v_1), \dots, \check{r}_t(v_n))$ by perturbing each *original cost* entry $\check{r}_t(v_j)$ according to some probability distribution f as follows

$$r_t(v_j) := \max\{0, \check{r}_t(v_j) + \epsilon(v_j)\}, \quad \text{where } \epsilon(v_j) \leftarrow f.$$

That is, to each original cost entry we add a random number which is chosen from f. The obtained *smoothed* task is denoted by $\tau_t := (r_t(v_1), \dots, r_t(v_n))$. We use μ and σ, respectively, to denote the expectation and the standard deviation of f. We assume that f is symmetric around $\mu := 0$. We take the maximum of zero and the smoothing outcome in order to assure that the smoothed costs are non-negative. Thus, the probability for an original zero cost entry to remain zero is amplified to $\frac{1}{2}$.

A major criticism to the additive model is that zero entries are destroyed. However, one can easily verify that the lower bound proof of $2n-1$ [BLS92,MMS88,BEY98] on the competitive ratio of any deterministic algorithm for metrical task systems goes through for any smoothing model that does not destroy zeros.

Our analysis holds for a large class of probability distributions, which we call *permissible*. We say f is permissible if (i) f is symmetric around $\mu = 0$ and (ii) f is non-increasing in $[0, \infty)$. For example, the uniform and the normal distribution are permissible. Since the stated upper bounds on the competitive ratio of WFA do not further improve by choosing σ much larger than $U_{\min}$, we assume that $\sigma \leq 2U_{\min}$. Moreover, we use c_f to denote a constant depending on f such that for a random ϵ chosen from f, $\mathbf{P}[\epsilon \geq \sigma/c_f] \geq \frac{1}{4}$.

All our results hold against an *adaptive adversary*. An adaptive adversary reveals the task sequence over time, thereby taking decisions made by the online algorithm in the past into account.

4 A Lower Bound on the Optimal Offline Cost

In this section, we establish a lower bound on the cost incurred by an optimal offline algorithm OPT when run on smoothed task sequences. For the purpose of proving the lower bound, we first investigate an interesting version of a balls into bins experiment, which we call the *constrained balls into bins game*.

4.1 Constrained Balls into Bins Game

We are given n bins $B_1, \ldots, B_n$. In each round, we place a ball independently in each bin B_i with probability p; with probability $1-p$ no ball is placed in B_i. We define the *height* $h_t(i)$ of a bin B_i as the number of balls in B_i after round t. We have dependencies between the heights of different bins that are specified by an (undirected) *constraint graph* $G_c := (V_c, E_c)$. The node set V_c of G_c contains n nodes $u_1, \ldots, u_n$, where each node u_i corresponds to a bin B_i. All edges in E_c have uniform edge lengths equal to Q. Let D be the maximum degree of a vertex in G_c. Throughout the experiment, we maintain the following invariant.

Invariant: The difference in height between two bins B_i and B_j is at most the shortest path distance between u_i and u_j in G_c.

If the placement of a ball into a bin B_i would violate this invariant, the ball is *rejected*; otherwise we say that the ball is *accepted*. Observe that if two bins B_i and B_j do not violate the invariant in round t, then, in round $t+1$, B_i and B_j might cause a violation only if one bin, say B_i, receives a ball, and the other, B_j, does not receive a ball; if both receive a ball, or both do not receive a ball, the invariant remains true.

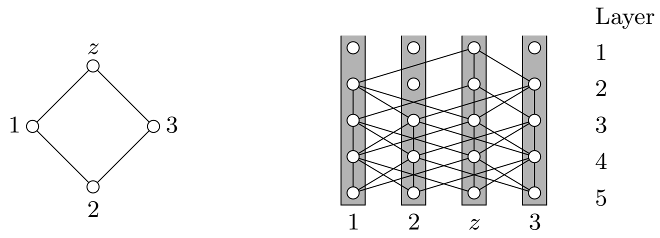

Fig. 1. Illustration of the "unfolding" for $Q = 1$ and $h = 5$. Left: constraint graph G_c. Right: layered dependency graph $\mathcal{D}_h$.

Theorem 1. *Fix any bin B_z. Let R_z be the number of rounds needed until the height of B_z becomes $h \geq \log(n)$. Then,* $\mathbf{P}[R_z > c_3 h\,(1 + \log(D)/Q)] \leq 1/n^4$.

We remark that there are instances, where the above bound is indeed tight.

We next describe how one can model the growth of the height of B_z by an alternative, but essentially equivalent, game on a *layered dependency graph*. A layered dependency graph $\mathcal{D}_h$ consists of h layers, $V_1, \ldots, V_h$, and edges are present only between adjacent layers. The idea is to "unfold" the constraint graph G_c into a layered dependency graph $\mathcal{D}_h$.

We describe the construction for $Q = 1$; the details for $Q > 1$ can be found in [SS03]. Each layer of $\mathcal{D}_h$ corresponds to a subset of nodes in G_c. Layer 1 consists of z only, the node corresponding to bin B_z. Assume we have constructed layers $V_1, \ldots, V_i$, $i < h$. Then, V_{i+1} is constructed from V_i by adding all nodes, $\Gamma_{G_c}(V_i)$, that are adjacent to V_i in G_c, i.e., $V_{i+1} := V_i \cup \Gamma_{G_c}(V_i)$. For every pair $(u, v) \in V_i \times V_{i+1}$, we add an edge (u, v) to $\mathcal{D}_h$ if $(u, v) \in E_c$, or $u = v$. See Figure 1 for an example.

Now, the following game on $\mathcal{D}_h$ is equivalent to the balls and bins game. Each node in $\mathcal{D}_h$ is in one of three states, namely UNFINISHED, READY or FINISHED. Initially, all nodes in layer h are READY and all other nodes are UNFINISHED. In each round, all READY nodes toss a coin; each coin independently turns up *head* with probability p and *tail* with probability $1 - p$. A READY node changes its state to FINISHED if the outcome of its coin toss is *head*. At the end of each round, an UNFINISHED node in layer j changes its state to READY, if all its neighbors in layer $j + 1$ are FINISHED.

Note that the nodes in layer V_j are FINISHED if and only if the corresponding bins B_i, $i \in V_j$, have height at least j. Consequently, the number of rounds needed until the root node z in $\mathcal{D}_h$ becomes FINISHED is equal to the number of rounds needed for the height of B_z to become h.

Proof (Theorem 1). Let $\mathcal{D}_h$ be a layered dependency graph constructed from G_c as described above. As argued above, the event $(R_z \leq t)$ is equivalent to the event that the root node becomes FINISHED within t rounds in $\mathcal{D}_h$. Consider the event that the root node z does not become FINISHED after t rounds. Then, there exists a *bad* path $P := (v_1, \ldots, v_h)$ from $z = v_1$ to some node v_h in the bottom layer h such that no node v_i of P was delayed by nodes other than $v_{i+1}, \ldots, v_h$. Put differently, P was delayed independently of any other path. Consider the outcome of the coin flips only for the nodes along P. If P is bad then the number of coin flips, denoted by X, that turned up *head* within t rounds is at most $h-1$. Let $\alpha(t)$ denote the probability that P is bad, i.e., $\alpha(t) := \mathbf{P}[X \le h-1]$. Clearly, $\mathbf{E}[X] = tp$.

Observe that in $\mathcal{D}_h$ any node has at most $D+1$ neighbors in the next larger layer. That is, the number of possible paths from z to any node v in layer h is bounded by $(D+1)^h$.

Thus, $\mathbf{P}[R_z > t] \le \alpha(t)(D+1)^h$. We want to choose t such that this probability is at most $1/n^4$. If we choose $t \ge (32/p)(h + h\log(D))$ and use Chernoff's bound [MR95] on X, we obtain for $h \ge \log(n)$

$$\alpha(t) = \mathbf{P}[X \le h-1] \le \mathbf{P}[X \le pt/2] \le e^{-pt/8} \le \frac{1}{n^4(D+1)^h}.$$

□

4.2 Lower Bound

We are now in a position to prove the following lemma.

Lemma 1. *Let $\check{S}$ be an adversarial sequence of $\ell := \lceil c_2 n\gamma(U_{\min}/\sigma + \log(D))\rceil$ tasks, for a fixed constant c_2 and some $\gamma \ge 1$. Then, $\mathbf{P}[\textsc{opt}[\mathcal{S}] < n\gamma U_{\min}] \le 1/n^3$.*

We relate the growth of the work function values to the balls and bins game as follows. For each node v_i of G we have a corresponding bin B_i. We obtain the constraint graph G_c from G by setting all edge lengths to $Q := \lfloor U_{\min}/\Delta \rfloor$, where $\Delta := \min\{U_{\min}, \sigma/c_f\}$. Since for any v_i and any time t, $\mathbf{P}[r_t(v_i) \ge \sigma/c_f] \ge \frac{1}{4}$, we place a ball into B_i with probability $\frac{1}{4}$. The following lemma establishes a relation between the work function value of v_i and the height $h_t(i)$ of B_i.

Lemma 2. *Consider any node v_i and its corresponding bin B_i. Let $h_t(i)$ denote the number of balls in bin B_i after t rounds. Then, for any $t \ge 0$, $w_t(v_i) \ge h_t(i)\,\Delta$.*

Put differently, the number of rounds needed until a bin B_i has height h stochastically dominates the time t needed until $w_t(v_i) \ge h\Delta$. Applying Theorem 1, we obtain that after $\ell := \lceil c_2 n\gamma(U_{\min}/\sigma + \log(D))\rceil$ rounds, for an appropriate constant c_2, the probability that there exists a bin of height less than $2n\gamma Q$ is at most $1/n^3$. That is, with probability at least $1 - 1/n^3$, all v_i satisfy $w_\ell(v_i) \ge 2n\gamma Q\Delta \ge n\gamma U_{\min}$. Since $\textsc{opt}[\mathcal{S}] = \min_{u \in V}\{w_\ell(u)\}$, the theorem follows.

We will use the Lemma 1 several times as follows.

Corollary 1. *Let $\check{S}$ be an adversarial sequence of $\ell := \lceil c_2 n\gamma(U_{\min}/\sigma + \log(D))\rceil$ tasks, for a fixed constant c_2 and an some $\gamma \ge 1$. Then, the smoothed competitive ratio of* WFA *is at most $\mathbf{E}[\textsc{wfa}[\mathcal{S}]]/(n\gamma U_{\min}) + o(1)$.*

Proof. Let $\mathcal{S}$ be a random variable denoting a smoothed sequence obtained from $\check{\mathcal{S}}$. We define $\mathcal{E}$ as the event that OPT incurs a cost of at least $n\gamma U_{\min}$ on $\mathcal{S}$. By Lemma 1, we have $\mathbf{P}[\neg\mathcal{E}] \leq 1/n^3$. Thus,

$$\begin{aligned}\mathbf{E}\left[\frac{\text{WFA}[\mathcal{S}]}{\text{OPT}[\mathcal{S}]}\right] &= \mathbf{E}\left[\frac{\text{WFA}[\mathcal{S}]}{\text{OPT}[\mathcal{S}]}\,\middle|\,\mathcal{E}\right]\mathbf{P}[\mathcal{E}] + \mathbf{E}\left[\frac{\text{WFA}[\mathcal{S}]}{\text{OPT}[\mathcal{S}]}\,\middle|\,\neg\mathcal{E}\right]\mathbf{P}[\neg\mathcal{E}] \\ &\leq \frac{\mathbf{E}[\text{WFA}[\mathcal{S}]\,|\,\mathcal{E}]\,\mathbf{P}[\mathcal{E}]}{n\gamma U_{\min}} + \frac{2n-1}{n^3} \leq \frac{\mathbf{E}[\text{WFA}[\mathcal{S}]]}{n\gamma U_{\min}} + o(1),\end{aligned}$$

where the second inequality follows from the definition of $\mathcal{E}$ and the fact that the (worst case) competitive ratio of WFA is $2n-1$. □

5 Upper Bounds

5.1 First Upper Bound

We derive the first upper bound on the smoothed competitive ratio of WFA. The idea is as follows. We derive two upper bounds on the smoothed competitive ratio of WFA. The first one is a deterministic bound, and the second one uses the probabilistic lower bound on OPT. We combine these two bounds using the following fact to obtain the theorem stated below.

Fact 4. *Let A, B, and X_i, $1 \leq i \leq m$, be positive quantities. We have*

$$\min\left\{\frac{A\sum_{i=1}^m X_i}{\sum_{i=1}^m X_i^2}, \frac{B\sum_{i=1}^m X_i}{m}\right\} \leq \sqrt{A\,B}.$$

Consider any deterministic input sequence $\mathcal{K}$ of length ℓ. Let $s_0, s_1, \ldots, s_\ell$ denote the sequence of nodes visited by WFA. Define $C(t) := r_t(s_t) + \delta(s_{t-1}, s_t)$ as the service cost plus the transition cost incurred by WFA in round t. With respect to $\mathcal{K}$ we define T as the set of rounds, where the increase of the work function value of s_{t-1} is at least one half of the transition cost, i.e., $t \in T$ if and only if $w_t(s_{t-1}) - w_{t-1}(s_{t-1}) \geq \delta(s_{t-1}, s_t)/2$. We use $\bar{T}$ to refer to the complement of T. Due to Fact 2 we have $w_t(s_{t-1}) = w_t(s_t) + \delta(s_{t-1}, s_t)$. Therefore, the above definition is equivalent to

$$T := \left\{t : w_t(s_t) - w_{t-1}(s_{t-1}) \geq -\tfrac{1}{2}\delta(s_{t-1}, s_t)\right\}. \quad (3)$$

We first prove that the total cost of WFA on $\mathcal{K}$ is bounded by a constant times the contribution of rounds in T.

Lemma 3. *Let $\mathcal{K}$ be a sufficiently long sequence such that* WFA$[\mathcal{K}] \geq 6\,Diam$. *Then,* WFA$[\mathcal{K}] \leq 8\sum_{t\in T} C(t)$.

We partition T into T^1 and T^2, where $T^1 := \{t \in T : w_t(s_t) - w_{t-1}(s_t) \leq 4U_{\max}\,diam\}$, and $T^2 := T \setminus T_1$.

Lemma 4. *Let $\mathcal{K}$ be a sufficiently long sequence such that* $\text{OPT}[\mathcal{K}] \geq 2Diam$*. There exists a constant b such that*

$$\text{OPT}[\mathcal{K}] \geq \frac{1}{bn}\left(\frac{1}{U_{\max}}\sum_{t\in T^1} C(t)^2 + n\sum_{t\in T^2} C(t)\right).$$

Theorem 2. *The smoothed competitive ratio of* WFA *is* $O(\sqrt{n\cdot(U_{\max}/U_{\min})(U_{\min}/\sigma+\log(D))})$.

Proof. Let $\check{\mathcal{S}}$ be an adversarial task sequence of length $\ell := \lceil c_2 n\gamma(U_{\min}/\sigma + \log(D))\rceil$, and let $\mathcal{S}$ be a random variable denoting a smoothed sequence obtained from $\check{\mathcal{S}}$. Due to the proof of Corrollary 1 it suffices to bound $\mathbf{E}[\text{WFA}[\mathcal{S}]/\text{OPT}[\mathcal{S}]\,|\,\mathcal{E}]$, where $\mathcal{E}$ is the event $(\text{OPT}[\mathcal{S}] \geq n\gamma U_{\min})$. Consider a smoothing outcome $\mathcal{S}$ such that the event $\mathcal{E}$ holds. We fix γ sufficiently large such that $\text{OPT}[\mathcal{S}] \geq 6Diam$. Observe that $\text{WFA}[\mathcal{S}] \geq \text{OPT}[\mathcal{S}] \geq 6Diam$.

First, assume $\sum_{t\in T^1} C(t) < \sum_{t\in T^2} C(t)$. Due to Lemma 3 and Lemma 4,

$$\text{WFA}[\mathcal{S}] \leq 16\sum_{t\in T^2} C(t) \quad \text{and} \quad \text{OPT}[\mathcal{S}] \geq \frac{1}{b}\sum_{t\in T^2} C(t).$$

Hence, $\mathbf{E}[\text{WFA}[\mathcal{S}]/\text{OPT}[\mathcal{S}]\,|\,\mathcal{E}] = O(1)$.

Next, assume $\sum_{t\in T^1} C(t) \geq \sum_{t\in T^2} C(t)$. By Lemma 3 and Lemma 4 we have

$$\text{WFA}[\mathcal{S}] \leq 16\sum_{t\in T^1} C(t) \quad \text{and} \quad \text{OPT}[\mathcal{S}] \geq \frac{1}{bn}\left(\frac{1}{U_{\max}}\sum_{t\in T^1} C(t)^2\right). \tag{4}$$

Thus,

$$\frac{\text{WFA}[\mathcal{S}]}{\text{OPT}[\mathcal{S}]} \leq 16bnU_{\max}\left(\frac{\sum_{t\in T^1} C(t)}{\sum_{t\in T^1} C(t)^2}\right). \tag{5}$$

Since $\mathcal{E}$ holds, we also have

$$\frac{\text{WFA}[\mathcal{S}]}{\text{OPT}[\mathcal{S}]} \leq \frac{16\ell\sum_{t\in T^1} C(t)}{\ell n\gamma U_{\min}} \leq \frac{c}{U_{\min}}\left(\frac{U_{\min}}{\sigma} + \log(D)\right)\left(\frac{\sum_{t\in T^1} C(t)}{|T^1|}\right), \tag{6}$$

where the last inequality holds for an appropriate constant c and since $\ell \geq |T^1|$. Observe that (6) is well-defined since $\sum_{t\in T^1} C(t) \geq \frac{1}{16}\text{WFA}[\mathcal{S}]$ (by (4)) and $\text{WFA}[\mathcal{S}] \geq 6Diam$ imply that $|T^1| \geq 1$.

Applying Fact 4 to (5) and (6), these two bounds are combined to the one stated in the theorem. □

5.2 Second Upper Bound

Our second upper bound easily follows from the proof of Corrollary 1 and the following deterministic relation between WFA and OPT.

Lemma 5. *Let $\mathcal{K}$ be any request sequence of length ℓ. Then,* $\text{WFA}[\mathcal{K}] \leq \text{OPT}[\mathcal{K}] + Diam\cdot\ell$.

Theorem 3. *The smoothed competitive ratio of* WFA *is* $O((Diam/U_{\min})\cdot(U_{\min}/\sigma+\log(D)))$.

5.3 Potential Function

The next lemma can be proved using a potential function argument. Intuitively, it states that the expected cost of WFA is bounded by the expected cost of a simple greedy online algorithm.

Lemma 6. *Let $\mathcal{S}$ be a smoothed sequence of ℓ tasks. For each t, $1 \leq t \leq \ell$, and a given node s, define a random variable $\Delta_t(s) := \min_{u \in V}\{r_t(u) + \delta(u, s)\}$. Let $\kappa > 0$. If $\mathbf{E}[\Delta_t(s)] \leq \kappa$ for each $s \in V$ and for each t, $1 \leq t \leq \ell$, then $\mathbf{E}[\text{WFA}[\mathcal{S}]] \leq 4\kappa\ell + Diam$.*

5.4 Random Tasks

We derive an upper bound on the expected competitive ratio of WFA if each request cost is chosen independently from a probability distribution f which is non-increasing in $[0, \infty)$. We need the following fact.

Fact 5. *Let f be a continuous, non-increasing distribution over $[0, \infty)$ with mean μ and standard deviation σ. Then, $\mu \leq \sqrt{12}\sigma$.*

Theorem 4. *If each request cost is chosen independently from a non-increasing probability distribution f over $[0, \infty)$ with standard deviation σ then the expected competitive ratio of* WFA *is $O(1 + (\sigma/U_{\min}) \cdot \log(D))$.*

Proof. Let $\mathcal{S}$ be a random task sequence of length $\ell := \lceil c_2 n\gamma(U_{\min}/\sigma) + \log(D))\rceil$, for an appropriate $\gamma \geq U_{\max}$, generated from f. Observe that since $\gamma \geq U_{\max}$, we have $\ell \geq Diam$. For any t and any node s, we have $\Delta_t(s) = \min_{u \in V}\{r_t(u) + \delta(u, s)\} \leq r_t(s)$. Since $r_t(s)$ is chosen from f, Fact 5 implies that $\mathbf{E}[\Delta_t(s)] \leq \kappa := \sqrt{12}\sigma$. Thus, by Lemma 6, we have $\mathbf{E}[\text{WFA}[\mathcal{S}]] = 4\sqrt{12}\sigma\ell + Diam = O(\sigma\ell)$.

Note that we can use the lower bound established in Section 4 to bound the cost of OPT: The generation of $\mathcal{S}$ is equivalent to smoothing (according to f) an adversarial task sequence consisting of all-zero request vectors only. Here, we do not need that the distribution f is symmetric around its mean. The theorem now follows from Corrollary 1. □

5.5 β-Elementary Tasks

We can strengthen the upper bound on the smoothed competitive ratio of WFA if the adversarial task sequence only consists of β-elementary tasks. Recall that in a β-elementary task the number of non-zero request costs is at most β.

Theorem 5. *If the adversarial task sequence only consists of β-elementary tasks then the smoothed competitive ratio of* WFA *is $O(\beta(U_{\max}/U_{\min})(U_{\min}/\sigma + \log(D)))$.*

The proof follows easily from the following lemma, Lemma 6 and Corrollary 1.

Lemma 7. *Let τ_t be a task obtained by smoothing a β-elementary task, where $\beta < n$. Then, $\mathbf{E}[\Delta_t(s)] \leq \sigma + \beta U_{\max}$ for each node $s \in V$.*

6 Conclusion

In this paper we focused on the asymptotic behaviour of WFA if the request costs of an adversarial task sequence are perturbed by means of a symmetric additive smoothing model. We showed that the smoothed competitive ratio of WFA is much better than its worst case competitive ratio suggests and that it depends on topological parameters of the underlying graph. Moreover, all our bounds, except the one for β-elementary tasks, are tight up to constant factors. We believe that our analysis gives a strong indication that the performance of WFA in practice is much better than $2n-1$.

An open problem would be to strengthen the universal lower bounds. Moreover, it would be interesting to obtain exact (and not only asymptotic) bounds on the smoothed competitive ratio of WFA.

Acknowledgements

We thank Alberto Marchetti-Spaccamela for his comments on a first draft of this paper. Moreover, we thank Kurt Mehlhorn for helpful suggestions and discussions; due to one of his suggestions we were able to significantly shorten the proof of Theorem 1.

References

[BEY98] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[BLMS+03] L. Becchetti, S. Leonardi, A. Marchetti-Spaccamela, G. Schäfer, and T. Vredeveld. Average case and smoothed competitive analysis of the multi-level feedback algorithm. In *Proceedings of the Forty-Fourth Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 462–471, 2003.

[BLS92] A. Borodin, N. Linial, and M. Saks. An optimal online algorithm for metrical task systems. *Journal of the ACM*, 39:745–763, 1992.

[MMS88] M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive algorithms for on-line problems. In *ACM Symposium on Theory of Computing*, pages 322–333, 1988.

[MR95] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, 1 edition, 1995.

[SS03] G. Schäfer and N. Sivadasan. Topology matters: Smoothed competitiveness of metrical task systems. Technical Report MPI-I-2003-1-016, Max-Planck-Institut für Informatik, Saarbrücken, Germany, 2003.

[ST01] D. Spielman and S. H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. In *ACM Symposium on Theory of Computing*, 2001.