Dynamic Page Migration with Stochastic Requests

Marcin Bienkowski*
International Graduate School
of Dynamic Intelligent Systems
University of Paderborn, Germany
young@upb.de

ABSTRACT

The page migration problem is one of subproblems of data management in networks. It occurs in a distributed network of processors sharing one indivisible memory page of size D. During runtime, the processors access a unit of data from the page, and the system is allowed to migrate the page between the processors. The problem is to compute (online) a schedule of page movements to minimize the total communication cost.

The Dynamic Page Migration problem is an extension to the page migration. It attempts to model the network dynamics, occurring, for example, in mobile networks. However, the pace of changes is restricted, i.e. the distances between processors can change only by a constant per round. The movement of the nodes induce changes in the communication cost between each pair of nodes, which is proportional to the distance between them raised to some power α. This is typical for mobile wireless networks, where nodes can move with a constant speed, and the cost of communication is measured in terms of energy used for sending the data. Thus, by setting α equal to the propagation exponent of the medium, cost minimization becomes minimizing the total energy consumption in the system.

However, as proven in [7], if both network mobility and request sequence are created by an adversary, then the competitive ratio is polynomially large in D and in the number of the nodes. In our search for a reasonable, close-to-reality model, in this paper we consider a scenario in which the network mobility is adversarial, but the requests are generated randomly by a stochastic process. We design an algorithm MTFR for this scenario, and prove that it is $\mathcal{O}(1)$-competitive, on expectation and with high probability.

*Partially supported by DFG-Sonderforschungsbereich 376 "Massive Parallelität: Algorithmen Entwurfsmethoden Anwendungen", and by the Future and Emerging Technologies programme of the EU under EU Contract 001907 DELIS "Dynamically Evolving, Large Scale Information Systems".

SPAA'05, July 18–20, 2005, Las Vegas, Nevada, USA.

Categories and Subject Descriptors

E.1 [**Data Structures**]: Graphs and networks; F.2.2 [**Nonnumerical Algorithms and Problems**]: Computations on discrete structures; G.2.2 [**Graph Theory**]: Network problems; G.3 [**Probability and Statistics**]: Markov processes, Probabilistic algorithms, Stochastic processes; C.2.4 [**Distributed systems**]

General Terms

Algorithms, Theory

Keywords

online algorithms, data management, page migration, dynamic networks, mobile networks

1. INTRODUCTION

The page migration problem [1, 3, 9, 11, 14, 18] arises in a distributed network of processors which share some global memory. Shared variables or memory pages are stored at the local memory of these processors. In the page migration problem we restrict our attention to *one indivisible memory page* stored at some processor. If a processor wants to access (read or write) a single unit of data from the page, and the page is not stored in its local memory, it has to send a request to the processor holding the page copy and appropriate data is sent back. To avoid the problem of maintaining consistency among multiple copies of the page, the model allows only *one copy* of the page to be stored in the network. However, to reduce the communication costs, the system may migrate the page to a new processor. The problem is to decide, *online*, when and where to move the page to minimize the total cost of communication.

The processors are modelled as nodes with some positions in a metric space $(\mathcal{X}, d)$. This induces a complete graph of all nodes, with the distances given by the metric d. The single copy of the memory page is of size D and is stored at one of the processors. The cost of accessing a single unit of data from the page stored at a remote processor is a function of the distance between the accessing processor and the processor holding the page. On the other hand, migrating the page between two processors incurs a cost which is equal to the cost of sending one unit of data times D, the size of the page. The problem is considered in online setting where access requests are dictated by a request adversary and come up online. The performance of an (online) algorithm is measured by competitive analysis [17], i.e. by comparing its total cost to the total cost of the optimal offline algorithm on the same requests of inputs.

The *Dynamic Page Migration* (DPM) problem [5, 6, 7, 8] is an extension to this model where an additional network adversary can change the positions of the nodes during the runtime of the algorithm. This is typical in mobile wireless networks, and also tries to capture the dynamics of the wired ones. Obviously, the movement changes the distances and corresponding costs of sending data between the processors. The DPM model imposes a *basic restriction* on the network adversary, i.e. the speed of the network changes has to be bounded. In each round, a new position of each node has to be chosen within a ball of a constant diameter centered at its previous position.

Whereas the cost function in standard models of page migration is equal to the distance between two processors, in DPM the cost of accessing one unit of data is defined to be the distance, raised to an integer power α, between the requesting node and the node holding the page, plus a constant overhead for communication. First we explain why we introduced the overhead. In essence, it represents the fact that there are no zero-cost communications, even if nodes are very close. Moreover, without this overhead in the definition of the cost, no online algorithm is able achieve a finite competitive ratio. The explanation for α is the following. If we choose α to be the propagation exponent [16] (for example 2 for an ideally free space), then the communication cost between two nodes is proportional to the energy required to send a message. Thus, in this case the problem is to minimize the total energy consumption in the system, consisting of mobile stations moving with a constant speed.

It appears (see [5, 7]) that without further restrictions on the adversaries, the competitive ratio of the problem is polynomially large in size of the page and in number of the nodes. The results were obtained for the case $\alpha = 1$, but it can be shown that for larger α the competitive ratios are even higher. Therefore, another scenario, in which the network adversary was replaced by a stochastic process, was analyzed in [7] and further extended in [6]. It turns out that when the requests are given by the adversary, but the mobility of the nodes in some specific metric space is induced by a random walk of the nodes, then (with high probability) the competitive ratio of the problem can be substantially lowered.

In this paper we consider a DPM scenario, opposite to the one described above. In this scenario, the network adversary can dictate the changes of the network in any way, as long as it obeys the basic restriction, and we replace the request adversary by a random process. In each step the processor issuing the request is chosen independently at random, according to some fixed probability distribution. This is a natural situation if we know that the mobile nodes' accesses to the page appear periodically with some given frequencies.

1.1 Our model

Following [7], we define the DPM problem as follows. The network is modeled as a set of n nodes (processors) numbered from v_1 to v_n placed in some metric space $(\mathcal{X}, d)$. The distances between two nodes v_x and v_y in time step t are denoted by $d_t(v_x, v_y)$ and are determined by the metric d. However, we extend the notion of the distance between two nodes in the following way. If v_x and v_y are the same node, then we denote it by $v_x \equiv v_y$ and we denote the distance between them by 0_E. Note that this is different from the case when they just occupy the same point in the space, in which case we write $v_x = v_y$ and $d(v_x, v_y) = 0$.

We assume discrete time steps $t = 1, 2, \ldots$. An input consists of a *configuration sequence* $(\mathcal{C}_t)$ and a *request sequence* (σ_t). $\mathcal{C}_t$, called a *configuration in time step* t, is a tuple describing the positions of all nodes in time step t; σ_t denotes the single node that issues a request at time t. In this paper we consider a *stochastic requests scenario*, in which $(\mathcal{C}_t)$ is generated by a network adversary, and σ_t is generated by a random process. The basic restriction mentioned in the previous subsection is formalized as follows.

Definition 1. *A Δ-restricted network adversary is an adversary which is allowed to choose a new position of each node within a ball with radius Δ, centered at the previous position of this node.*

For the Δ-restricted network adversary and any node v_x, its positions x_t and x_{t+1} in two consecutive configurations $\mathcal{C}_t$ and $\mathcal{C}_{t+1}$ cannot be too far apart, i.e. $d(x_t, x_{t+1}) \le \Delta$.

Any two nodes are able to communicate directly with each other. The cost of sending a unit of data from node v_x to node v_y at time step t is defined by a cost function $c_t(v_x, v_y)$ as follows. If $v_x \equiv v_y$ then $c_t(v_x, v_y) = 0$. Otherwise,

$$c_t(v_x, v_y) = [d_t(v_x, v_y)]^\alpha + 1 \ , \qquad (1)$$

where $\alpha \ge 1$ is a fixed integer. We have one shared, indivisible memory page of size D, initially stored at the node v_1. The cost of moving the whole page from v_x to v_y in time step t is equal to $D \cdot c_t(v_x, v_y)$.

Formally, we construct the *stochastic requests scenario* of the DPM problem as follows. First, the whole configuration sequence $(\mathcal{C}_t)$, including the initial configuration C_1, is created by the Δ-restricted network adversary, for some constant Δ. The request adversary chooses the probability distribution over all indices of the nodes $\pi : \{1, \ldots, n\} \to (0, 1)$.[1] Then, in time step $t \ge 1$ the following happens.

1. The positions of the nodes in this time step are defined by $\mathcal{C}_t$.
2. A request is generated randomly, i.e. σ_t is chosen randomly according to the distribution π and node v_{σ_t} issues a request. For clarity, we sometimes abuse the notation and write *node* σ_t instead of *node* v_{σ_t}.
3. Let $P_{\mathrm{ALG}}(t)$ be the node at which algorithm Alg has its page. Alg has to pay $c_t(P_{\mathrm{ALG}}(t), \sigma_t)$ for serving the request.
4. Alg chooses, optionally, a new position $P'_{\mathrm{ALG}}(t)$ and moves its page to $P'_{\mathrm{ALG}}(t)$. Such a transaction incurs a cost $D \cdot c_t(P_{\mathrm{ALG}}(t), P'_{\mathrm{ALG}}(t))$.

Sometimes, we will abuse the notation by writing that an algorithm is at v_i or moves to v_j, meaning that this algorithm has its page at v_i or moves its page to v_j, respectively.

We focus on the online version of the problem, i.e. the algorithm has to base its decision what to do in step t solely on the initial parts of the two input sequences up to step t, i.e. on the sequence $\mathcal{C}_1, \sigma_1, \mathcal{C}_2, \sigma_2, \ldots, \mathcal{C}_t, \sigma_t$. We want to

[1] With a little technical difficulty, we could extend the proof of Theorem 2 to handle distributions $\pi : \{1, \ldots, n\} \to [0, 1]$.

emphasize that the network adversary is oblivious and creates its sequence knowing the deterministic algorithm but not the request sequence (σ_t).

In order to analyze the performance in the stochastic request scenarios we follow [7], and adapt classical competitive analysis [17, 10] for the model where the input sequence is created both by the adversary and the stochastic process. We say that the algorithm ALG achieves competitive ratio $\mathcal{R}$ (or is $\mathcal{R}$-competitive) with probability p, if there exists an integer $T_{\min}$, s.t. for all configuration sequences $(\mathcal{C}_t)$ of length $T \geq T_{\min}$ and all probability distributions π, it holds

$$\mathbf{Pr}_{(\sigma_t)}\left[C_{\mathrm{ALG}}(\mathcal{C}_t, \sigma_t) \leq \mathcal{R} \cdot C_{\mathrm{OPT}}(\mathcal{C}_t, \sigma_t)\right] \geq p \ ,$$

where $C_{\mathrm{ALG}}(\mathcal{C}_t, \sigma_t)$ and $C_{\mathrm{OPT}}(\mathcal{C}_t, \sigma_t)$ are costs of ALG and the optimal offline algorithm, respectively. The probability is taken over all possible random choices made for generating (σ_t) sequence. Similarly, we say that an algorithm achieves expected competitive ratio of $\mathcal{R}$ (or is $\mathcal{R}$-competitive on expectation), if there exists an integer $T_{\min}$, s.t. for all sequences $(\mathcal{C}_t)$ of length $T \geq T_{\min}$ and all π, it holds

$$\mathbf{E}_{(\sigma_t)}\left[\frac{C_{\mathrm{ALG}}(\mathcal{C}_t, \sigma_t)}{C_{\mathrm{OPT}}(\mathcal{C}_t, \sigma_t)}\right] \leq \mathcal{R} \ .$$

To be consistent with the previous definition, we assume that if $C_{\mathrm{ALG}}(\mathcal{C}_t, \sigma_t) = C_{\mathrm{OPT}}(\mathcal{C}_t, \sigma_t) = 0$, then the ratio is 1. However, we do not permit situations where $C_{\mathrm{OPT}}(\mathcal{C}_t, \sigma_t) = 0 \neq C_{\mathrm{ALG}}(\mathcal{C}_t, \sigma_t)$. Note that competitive ratio $\mathcal{R}$ achieved with high probability does not directly imply expected competitive ratio $\mathcal{O}(\mathcal{R})$, since the former notion does not explicitly exclude inputs, on which the competitive ratio is very high or infinite. We also note that this notion is usually stronger and gives more realistic estimates than the similar $\mathbf{E}[C_{\mathrm{ALG}}(\mathsf{input})]/\mathbf{E}[C_{\mathrm{OPT}}(\mathsf{input})]$ ratio introduced by Koutsoupias and Papadimitriou [13].

1.2 Contribution of this paper

In this paper we design and analyze a deterministic algorithm MOVE-TO-FIRST-REQUEST (MTFR) for the DPM problem. We prove that in the stochastic requests scenario, its competitiveness is constant[2] with high probability, for any constant-restricted network adversary. In this context, high probability means that for any constant γ we can achieve probability $1 - O(D^{-\gamma})$ by running the algorithm sufficiently long. Additionally, we prove that in the worst case (occurring with negligible probability) the competitive ratio of MTFR is finite. These results assure that the performance of MTFR stabilizes with time and additionally allows us to conclude that on expectation its competitive ratio is also constant.

1.3 Related work

For the page migration problem in a static network, Westbrook [18] gave the first randomized $\mathcal{O}(1)$-competitive algorithms against oblivious and adaptive adversaries. See [4] for a general discussion on the possible types of adversaries. This result was improved for some network topologies like trees and uniform graphs by Chrobak et al. [11] and Lund et al. [14]. The first constant competitive deterministic algorithm was given by Awerbuch, Bartal and Fiat in [1] and the competitive ratio was subsequently improved by Bartal, Charikar and Indyk [3].

Bienkowski, Korzeniowski and Meyer auf der Heide [7] defined the DPM model and proved that for $\alpha = 1$ in adversarial scenario (both request and configuration sequences chosen by an adversary) the competitive ratio is $\Theta(\min\{D, n \cdot \sqrt{D}\})$ for an adaptive-online adversary. In [5] this ratio was proven to hold also for deterministic algorithms, and the competitive ratio against an oblivious adversary was proven to be between $\Omega(\min\{\sqrt{D \cdot \log n}, D^{2/3}\})$ and $\mathcal{O}(\min\{\sqrt{D} \cdot \log n, D\})$.

Another contribution of [7] was defining a *Brownian motion scenario*, in which the network adversary is replaced by some specific stochastic process. Precisely, all nodes perform a random walk on a bounded area of diameter B. The preliminary result of [7] for this scenario was extended in [6], and the MAJ algorithm presented there achieves the competitive ratio of $\mathcal{O}(\min\{\sqrt[4]{D}, n\} \cdot \mathrm{polylog}(B, D, n))$ for the case $\alpha = 1$.

Our notion of network changes substantially differs from the one introduced by Awerbuch, Bartal and Fiat [2], in which nodes do not change their positions but can appear and disappear. In particular, their algorithms are inapplicable in our model.

2. THE ALGORITHM

In this section we present a deterministic algorithm MTFR, which, on sufficiently long input sequences, achieves constant competitive ratio on expectation and with high probability.

Let MTFR be a deterministic algorithm which divides time steps into phases of length $\ell := D^{\alpha+1}$. In the first step of each phase, after serving a request, MTFR moves the page to the node which issued a request. In this paper we prove the following theorem.

THEOREM 2. *MTFR is $\mathcal{O}(1)$-competitive, with high probability, and on expectation, in the stochastic requests scenario of the DPM.*

Let us introduce some notations. Let $\pi_{\min} = \min_i \{\pi(i)\}$. Let

$$T_\gamma = m_\gamma \cdot \ell \ , \tag{2}$$

for

$$m_\gamma = \left(c_T \cdot \gamma \cdot \frac{\ln D}{\pi_{\min}^2}\right)^{\alpha+1} \cdot \left(\frac{2 \cdot \ell}{\alpha}\right)^{\alpha} , \tag{3}$$

where c_T is a constant, which will be specified later. In the following, we prove that MTFR is $\mathcal{O}(1)$-competitive with probability $1 - \mathcal{O}\left(D^{-\gamma}\right)$ if run for $T \geq T_\gamma$ time steps. Additionally, we prove that on the inputs of length $T \geq T_{(\alpha+1)^2}$, MTFR achieves expected competitive ratio $\mathcal{O}(1)$.

We prove constant competitiveness against $\frac{1}{2 \cdot \alpha}$-restricted network adversary. The proof for any $\mathcal{O}(1)$-restricted network adversary follows from the Reduction Lemma proven in [5].

LEMMA 3 (REDUCTION LEMMA). *Let ALG_A be a (possibly randomized) algorithm, which is k-competitive against an A-restricted network adversary. Then, ALG_A is also k-competitive against a B-restricted network adversary for*

[2]This constant depends on α, but we assume that α is a constant for all practical applications.

$B \leq A$. *Additionally, for any* $B \geq A$ *there exists an algorithm* ALG_B *which is* $\frac{B}{A} \cdot k$*-competitive against a* B*-restricted network adversary.*

The idea of the proof of Theorem 2 is as follows. We consider an input sequence $\mathcal{I}$ of length T. By $\mathcal{P}$ we denote a set of all finished phases, i.e., of length ℓ. At the end of $\mathcal{I}$ we might also have one unfinished phase; we denote it by p_{last}. This way $\mathcal{I} = (\biguplus_{p \in \mathcal{P}} p) \uplus p_{\text{last}}$.

For clarity, the proof was divided into subsections and the proofs of all technical lemmas were moved to the appendix. In this subsection we show some general relations concerning algorithm MTFR, and we define $\mathcal{K}_p$ as the average cost of communication in phase p. In Subsection 2.1 we show an $\Omega(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$ lower bound on the cost of OPT, which holds with probability $1 - \mathcal{O}(D^{-\gamma})$ on input sequences of length at least T_γ. In Subsection 2.2 we divide the input sequence into three disjoint parts and prove for each part separately that the cost it incurs on MTFR is, with high probability, at most $\mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$. Finally, in Subsection 2.3 we prove that even when the whole input sequence is generated by an adversary, the competitive ratio of MTFR is at most $\mathcal{O}(\ell^{\alpha+1})$. We conclude that MTFR has constant expected competitive ratio. In the proofs we are not aiming at minimizing constants, but rather at the simplicity.

Since the network adversary is $\frac{1}{2 \cdot \alpha}$-restricted, the distance between any two nodes can change only by at most $1/\alpha \leq 1$ per round. First, we present some general relations. Let K_t be the *average cost* of sending a unit of data between two nodes in time step t, i.e.,

$$K_t := \sum_{i=1}^{n} \sum_{j=1}^{n} \pi(i) \cdot \pi(j) \cdot c_t(v_i, v_j) \ . \tag{4}$$

We note that some of the terms contributing to the average cost K_t are always 0, because we counted also the nodes sending data to themselves. Therefore, we have

$$\begin{aligned} K_t &= \sum_i \sum_j \pi(i) \cdot \pi(j) \cdot c_t(v_i, v_j) \\ &= \sum_i \pi(i)^2 \cdot \underbrace{c_t(v_i, v_i)}_{=0} \\ &\quad + \sum_i \sum_{j \neq i} \pi(i) \cdot \pi(j) \cdot (1 + d_t^\alpha(v_i, v_j)) \ . \end{aligned} \tag{5}$$

Let $\beta = \sum_i \sum_{j \neq i} \pi(i)\pi(j)$. It follows that $K_t \geq \beta$ for any time step t. Let $k_t = \sqrt[\alpha]{K_t/\beta}$, i.e. $K_t = \beta \cdot k_t^\alpha$. Then $k_t \geq 1$ for all t, and it follows from the Jensen's Inequality [12] that

$$|k_{t+1} - k_t| \leq 1/\alpha \tag{6}$$

for any two consecutive steps t and $t+1$. The proof of Inequality 6 can be found in the appendix. As an immediate consequence we get

$$\frac{K_{t+1}}{K_t} \leq \frac{(k_t + \frac{1}{\alpha})^\alpha}{k_t^\alpha} \leq \left(1 + \frac{1}{\alpha}\right)^\alpha \leq e \ . \tag{7}$$

The same holds for the quotient $\left(\frac{c_{t+1}(v_a, v_b)}{c_t(v_a, v_b)}\right)$ for any two nodes v_a and v_b.

Let $c_t^{\max}$ be the maximum cost of communication between two nodes in time step t. We can also establish a relation between the average and the worst-case cost of communication in any round t, as stated in the lemma below.

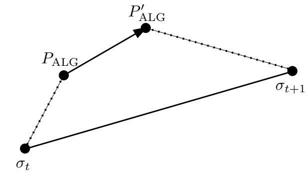

Figure 1: Illustration of Lemma 5

LEMMA 4. *For any time step* t, $K_t \geq b_1 \cdot \pi_{\min} \cdot c_t^{\max}$, *where* b_1 *is a constant.*

For any phase p we define the *average phase cost* $\mathcal{K}_p = \sum_{t \in p} K_t$. In the next subsection we relate the cost of MTFR and OPT to $\sum_{p \in \mathcal{P}} \mathcal{K}_p$.

2.1 Lower bound for the OPT's cost

In this subsection we show that for any phase $p \in \mathcal{P}$ the cost of the OPT in this phase is bounded from below by L_p, which is a random variable with expectation $\Omega(\mathcal{K}_p)$. Additionally, $0 \leq L_p \leq \sum_{t \in p} c_t^{\max}$ and all L_p are independent for different phases $p \in \mathcal{P}$.

Consider any phase $p \in \mathcal{P}$, and number all the time steps within this phase from 1 to ℓ. By $[\ell]_{\text{odd}}$ we denote a set of all odd numbers in the set $\{1, \ldots \ell\}$. For all $t \in [\ell]_{\text{odd}}$, we define

$$s_t = \frac{1}{3^\alpha} \cdot c_t(\sigma_t, \sigma_{t+1}) \ . \tag{8}$$

We note that for odd t, s_t are independent random variables. It follows from the fact that the adversary first chooses $c_t(\cdot, \cdot)$ functions, and then $(\sigma_t)_t$ is randomly picked.

We can show that sum of s_t constitutes a lower bound for the OPT's cost in phase p, i.e.

$$C_{\text{OPT}}(p) \ \geq \sum_{t \in [\ell]_{\text{odd}}} s_t =: L_p \tag{9}$$

To prove the inequality above, it is sufficient to prove the following lemma.

LEMMA 5. *For any algorithm* ALG *and any time step* t *the cost of serving requests* σ_t *and* σ_{t+1} *is at least* s_t.

PROOF. The situation in time step t is depicted in Figure 1. P_{ALG} and P'_{ALG} denote the nodes in which ALG has its page in step t and $t+1$, respectively. If $\sigma_t \equiv \sigma_{t+1}$, then $s_t = 0 \leq C_{\text{ALG}}(\sigma_t, \sigma_{t+1})$ and the lemma follows trivially. Otherwise, ALG has to pay at least $c_t(P_{\text{ALG}}, \sigma_t) + c_{t+1}(P'_{\text{ALG}}, \sigma_{t+1})$ for the requests in steps t and $t+1$ and $D \cdot c_t(P_{\text{ALG}}, P'_{\text{ALG}})$ for moving its page at the end of step t. Therefore,

$$\begin{aligned} & C_{\text{ALG}}(t, t+1) \\ &\quad = c_t(P_{\text{ALG}}, \sigma_t) + D \cdot c_t(P_{\text{ALG}}, P'_{\text{ALG}}) + c_{t+1}(P'_{\text{ALG}}, \sigma_{t+1}) \\ &\quad \geq c_t(\sigma_t, P_{\text{ALG}}) + c_t(P_{\text{ALG}}, P'_{\text{ALG}}) + \frac{1}{e} \cdot c_t(P'_{\text{ALG}}, \sigma_{t+1}) \end{aligned}$$

We apply Skewed Triangle Inequality (see appendix) to get

$$C_{\text{ALG}}(t, t+1) \geq \frac{1}{e} \cdot \frac{1}{3^{\alpha-1}} \cdot c_t(\sigma_t, \sigma_{t+1}) > \ s_t \ ,$$

which finishes the proof. $\square$

LEMMA 6. *For any phase p, $\mathbf{E}[L_p] \geq b_2 \cdot \mathcal{K}_p$, where b_2 is a constant.*

PROOF. From the definition of s_t we have

$$\mathbf{E}[s_t] = \frac{1}{3^\alpha} \cdot \sum_i \sum_j \pi(i) \cdot \pi(j) \cdot c_t(v_i, v_j) = \frac{1}{3^\alpha} \cdot K_t \ ,$$

and therefore $\mathbf{E}[L_p] = \mathbf{E}[\sum_{t \in [\ell]_{\text{odd}}} s_t] = \frac{1}{3^\alpha} \cdot \sum_{t \in [\ell]_{\text{odd}}} K_t$. It follows from Inequality 7 that two consecutive K_t can differ at most by a multiplicative factor of e. Therefore,

$$\begin{aligned} \mathbf{E}[L_p] &\geq \frac{1}{3^\alpha} \sum_{t \in [\ell]_{\text{odd}}} \frac{1}{2} \left(K_t + \frac{K_{t+1}}{e} \right) \\ &\geq \frac{1}{2 \cdot e \cdot 3^\alpha} \sum_{t \in [\ell]_{\text{odd}}} (K_t + K_{t+1}) \\ &= \frac{1}{2 \cdot e \cdot 3^\alpha} \cdot \mathcal{K}_p \ , \end{aligned}$$

which proves Lemma 6. □

Additionally, L_p are independent random variables fulfilling $0 \leq L_p = \sum_{t \in [\ell]_{\text{odd}}} s_t \leq \sum_{t \in p} c_t^{\max}$. Thus, using concentration bounds we are able to prove that the random variable $\sum_{p \in P} L_p$ is concentrated around its mean. We do not have any global upper bound on L_p variables, which renders classical Chernoff bound [15] unusable. However, we can relate $\mathcal{K}_p$ in any two consecutive phases, and use Hoeffding bound [15]. Then with the help of a combinatorial Lemma 7, we are able to guarantee concentration.

Let $\mathcal{A}_p := \sqrt[\alpha]{\mathcal{K}_p / (\ell \cdot \beta)}$, i.e. $\mathcal{K}_p = \ell \cdot \beta \cdot \mathcal{A}_p^\alpha$. $\mathcal{A}_p$ is approximately the average distance in phase p. Then we have an obvious relation $1 \leq \min_{t \in p}\{k_t\} \leq \mathcal{A}_p \leq \max_{t \in p}\{k_t\}$, and thus, by Inequality 6, for any two consecutive phases p_i and p_{i+1} holds

$$\mathcal{A}_{p_{i+1}} - \mathcal{A}_{p_i} \leq \max_{t \in p_{i+1}} \{k_t\} - \min_{t \in p_i} \{k_t\} \leq 2\ell \cdot (1/\alpha) \ . \tag{10}$$

From the symmetry, we have the same bound on the absolute difference, i.e., $|\mathcal{A}_{p_{i+1}} - \mathcal{A}_{p_i}| \leq \frac{2 \cdot \ell}{\alpha}$. We number all the phases in $\mathcal{P}$ from 1 to $m := |\mathcal{P}|$. Since all L_p are independent random variables, we can use the Hoeffding bound [15] to show that

$$\begin{aligned} &\mathbf{Pr}\left[\sum_{p \in P} L_p < \frac{1}{2} \cdot \sum_{p \in P} \mathbf{E}[L_p] \right] \\ &\leq \exp\Bigg(\underbrace{- 2 \cdot \Big(\frac{\sum_{p \in \mathcal{P}} \mathbf{E}[L_p]}{2} \Big)^2 \Big/ \sum_{p \in \mathcal{P}} \Big(\sum_{t \in p} c_t^{\max} \Big)^2}_{=: M} \Bigg) \end{aligned} \tag{11}$$

Using Lemma 4 and 6 we get the bound on M.

$$\begin{aligned} M &\geq \frac{b_1}{2} \cdot \pi_{\min}^2 \cdot \frac{\left(\sum_{p \in \mathcal{P}} b_2 \cdot \mathcal{K}_p \right)^2}{\sum_{p \in \mathcal{P}} (\mathcal{K}_p)^2} \\ &= \frac{b_1 \cdot b_2^2}{2} \cdot \pi_{\min}^2 \cdot \frac{\left(\sum_{p \in \mathcal{P}} \mathcal{A}_p^\alpha \right)^2}{\sum_{p \in \mathcal{P}} (\mathcal{A}_p^\alpha)^2} \end{aligned}$$

To bound the latter term we use the following combinatorial lemma.

LEMMA 7. *Let $\{A_i\}_{i=1}^m$ be the sequence of $m \geq 16$ real numbers, s.t. for any i, $A_i \geq 1$, and there exists δ, s.t. for any $i < m$, $|A_{i+1} - A_i| \leq \delta$. If $\delta \geq 1$ and $m \geq \delta^\alpha$, then*

$$\frac{\left(\sum_{i=1}^m A_i^\alpha \right)^2}{\sum_{i=1}^m (A_i^\alpha)^2} \geq b_3 \cdot \left(\frac{m}{\delta^\alpha} \right)^{\frac{1}{\alpha+1}} \ ,$$

where b_3 is a constant.

Finally, applying Lemma 7 with $\delta = 2\ell/\alpha$, results in

$$M \geq \frac{b_1 \cdot b_2^2 \cdot b_3}{2} \cdot \pi_{\min}^2 \cdot \left(\frac{m}{(2\ell/\alpha)^\alpha} \right)^{\frac{1}{\alpha+1}} \ . \tag{12}$$

If $T \geq T_\gamma$, then $m \geq m_\gamma$, and thus by setting $c_T := \frac{2}{b_1 \cdot b_2^2 \cdot b_3}$, we get $M \geq \gamma \cdot \ln D$. It follows from (11) that $C_{\text{OPT}}(\mathcal{P}) \geq \frac{1}{2} \cdot \sum_{p \in \mathcal{P}} \mathbf{E}[L_p] = \Omega(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$ with probability at least $1 - \exp(-M) \geq 1 - D^{-\gamma}$.

2.2 Upper bound for the MTFR's cost

In this subsection we analyze cost of MTFR on input sequence $\mathcal{I}$ of length $T \geq T_\gamma$ (for some integer γ). We divide $\mathcal{I}$ into three parts and we show for each part separately, that it incurs a cost $\mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$, with high probability.

$\mathcal{I}_1$ – a set of all time steps in all phases $p \in \mathcal{P}$, except for the first step of each phase,

$\mathcal{I}_2$ – the last phase p_{last} and time step 1,

$\mathcal{I}_3$ – the first steps of phases $p \in \mathcal{P}$, except for time step 1.

Part $\mathcal{I}_1$: We show that $C_{\text{MTFR}}(\mathcal{I}_1) = \mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$, with high probability. We consider any phase $p \in P$ and number all the steps within p from 1 to ℓ. By F_t we denote cost of serving requests by MTFR in step $t \geq 2$. We have

$$\begin{aligned} \mathbf{E}[F_t] &= \sum_{i,j} \mathbf{Pr}[\sigma_1 = i \wedge \sigma_t = j] \cdot c_t(v_i, v_j) \\ &= \sum_{i=1}^n \sum_{j=1}^n \mathbf{Pr}[\sigma_1 = i] \cdot \mathbf{Pr}[\sigma_t = j] \cdot c_t(v_i, v_j) \\ &= K_t \ . \end{aligned}$$

Let $(p)_{2 \ldots \ell}$ denote the set of all time steps in p from 2 to ℓ. From the linearity of expectation follows

$$\mathbf{E}[C_{\text{MTFR}}((p)_{2 \ldots \ell})] = \sum_{t=2}^{\ell} K_t \leq \mathcal{K}_p \ . \tag{13}$$

Since MTFR moves at the beginning of phase, random variables $C_{\text{MTFR}}((p)_{2 \ldots \ell})$ are independent. Similarly to the previous subsection's proof, $0 \leq C_{\text{MTFR}}((p)_{2 \ldots \ell}) \leq \sum_{t \in p} c_t^{\max}$, and thus we may apply the Hoeffding bound to prove that

$$\begin{aligned} &\mathbf{Pr}\left[\sum_{p \in \mathcal{P}} C_{\text{MTFR}}((p)_{2 \ldots \ell}) \geq 2 \cdot \sum_{p \in \mathcal{P}} \mathcal{K}_p \right] \\ &\leq \exp\left(-2 \cdot \Big(\sum_{p \in \mathcal{P}} \mathcal{K}_p \Big)^2 \Big/ \sum_{p \in \mathcal{P}} \Big(\sum_{t \in p} c_t^{\max} \Big)^2 \right) \\ &\leq e^{-M} \ . \end{aligned} \tag{14}$$

Therefore, $C_{\text{MTFR}}(\mathcal{I}_1) = \mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$ with the probability $1 - \mathcal{O}(D^{-\gamma})$.

Part $\mathcal{I}_2$: To prove $C_{\text{MTFR}}(\mathcal{I}_2) = \mathcal{O}(\sum_{p \in P} \mathcal{K}_p)$, we number all the steps within phase p_{last} from 1 to $|p_{\text{last}}| < \ell$.

$$\begin{aligned} C_{\text{MTFR}}(p_{\text{last}}) &= D \cdot c_1^{\max} + \sum_{t \in p_{\text{last}}} c_t^{\max} \\ &\le \frac{1}{b_1 \cdot \pi_{\min}} \cdot \Big(D \cdot K_1 + \sum_{t \in p_{\text{last}}} K_t \Big). \end{aligned}$$

Let Q be the value of K_t in the last time step of the last phase from $\mathcal{P}$. We may relate Q to the sum of K_t before and after that step, i.e. to $\sum_{p \in \mathcal{P}} \sum_{t \in p} K_t$ and to $\sum_{t \in p_{\text{last}}} K_t$ by the following lemma.

LEMMA 8. *Fix any sequence $\{K_t\}$ of length s and choose any step t_0. Let $Q := K_{t_0}$. Then,*

1. $\sum_{t=1}^{s} K_t = \mathcal{O}(\beta \cdot s^{\alpha+1} + s \cdot Q)$, *and*
2. $\sum_{t=1}^{s} K_t = \Omega(\beta \cdot s + s^{\frac{1}{\alpha+1}} \cdot Q)$.

By Lemma 8, $\sum_{t \in p_{\text{last}}} K_t = \mathcal{O}(\beta \cdot \ell^{\alpha+1} + \ell \cdot Q)$. By Inequality 6, we have $D \cdot K_1 \le D \cdot Q \cdot e = \mathcal{O}(\ell \cdot Q)$. Thus, the total cost of MTFR in the last phase is

$$C_{\text{MTFR}}(p_{\text{last}}) = \frac{1}{\pi_{\min}} \cdot \mathcal{O}(\beta \cdot \ell^{\alpha+1} + \ell \cdot Q) \quad . \tag{15}$$

On the other hand, since $T \ge T_\gamma \ge T_1$, it follows from Lemma 8 that

$$\begin{aligned} \sum_{p \in \mathcal{P}} \mathcal{K}_p &= \Omega(\beta \cdot T_1 + T_1^{\frac{1}{\alpha+1}} \cdot Q) \\ &= \Omega\left(\beta \cdot \left(\frac{\ell}{\pi_{\min}^2} \right)^{\alpha+1} + \frac{\ell}{\pi_{\min}^2} \cdot Q \right) \quad . \end{aligned} \tag{16}$$

We may combine (15) with (16) and get $C_{\text{MTFR}}(p_{\text{last}}) = \mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$. Note that this bound holds in the worst case, not only with high probability.

The cost incurred in time step 1 can be bounded similarly, and therefore $C_{\text{MTFR}}(\mathcal{I}_2) = \mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$.

Part $\mathcal{I}_3$: Third, we prove that with high probability $C_{\text{MTFR}}(\mathcal{I}_3) = \mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$. For any phase p, let $(p)_1$ be its first step. We have $0 \le C_{\text{MTFR}}((p)_1) \le (1 + D) \cdot c_{(p)_1}^{\max}$ and $\mathbf{E}[C_{\text{MTFR}}((p)_1)] = (1 + D) \cdot K_{(p)_1}$. Moreover, if we consider the random variables $C_{\text{MTFR}}((p)_1)$ in odd phases only, then all these variables are independent. Analogously to the proof for part $\mathcal{I}_1$, we can prove that

$$\mathbf{Pr}[\sum_{\substack{p \in \mathcal{P} \\ p \text{ odd}}} C_{\text{MTFR}}((p)_1) = \mathcal{O}(D \cdot \sum_{\substack{p \in \mathcal{P} \\ p \text{ odd}}} K_{(p)_1})] = 1 - \mathcal{O}(D^{-\gamma}).$$

The same bound holds for even phases. Since $\ell = D^{\alpha+1}$, it follows from Lemma 8 that $D \cdot K_{(p)_1} = \mathcal{O}(\sum_{t \in p} K_t)$, and thus $C_{\text{MTFR}}(\mathcal{I}_3) = \mathcal{O}(\sum_{p \in \mathcal{P}} \mathcal{K}_p)$ with probability $1 - \mathcal{O}(D^{-\gamma})$.

2.3 Expected competitive ratio

When we combine the two results of the previous subsections we get that on a sequence $\mathcal{I}$ of length $T \ge T_\gamma$, the competitive ratio of MTFR is constant with probability $1 - \mathcal{O}(D^{-\gamma})$. In this subsection we prove that it is also constant on expectation. To prove it we need the following bound.

LEMMA 9. *For any input sequence (even when both request and configuration sequences are chosen by the adversary) the competitive ratio of* MTFR *is at most $O(\ell^{\alpha+1})$.*

In the proof of this lemma we will often use the following bound which follows from Hölder's Inequality and is proven in the appendix.

LEMMA 10. *For any sequence of k non-negative numbers $a_1, a_2, \ldots, a_k$ and any integer $s \ge 1$ holds*

$$(\sum_i a_i)^s \le (\sum_i a_i^s) \cdot k^{s-1} \quad .$$

A inequality which follows from this lemma we denote, by $\overset{(\text{H})}{\le}$

PROOF OF LEMMA 9. Fix any input sequence $\mathcal{I}$ of length T. We divide this input into phases. In the following we concentrate on any phase p of length $\ell_0 \le \ell = D^{\alpha+1}$. We number all time steps within p from 1 to ℓ_0.

Let $f = 2 \cdot 3^\alpha$. We define a potential at the beginning of p as $\Phi_B(p) := f \cdot D \cdot c_0(P_{\text{MTFR}}(1), P_{\text{OPT}}(1))$, where c_0 is a cost function at the very beginning of the first time step of p, i.e. *before* the network adversary moves the nodes.[3] We define a potential at the end of p as $\Phi_F(p) := f \cdot D \cdot c_{\ell_0}(P'_{\text{MTFR}}(\ell_0), P'_{\text{OPT}}(\ell_0))$. At the beginning of first phase p_0, $\Phi_B(p_0) = 0$, because MTFR and OPT have their pages at the same node, v_1. For any two consecutive phases p_i and p_{i+1} holds $\Phi_F(p_i) = \Phi_B(p_{i+1})$. Thus, for proving that MTFR is $\mathcal{O}(\ell^{\alpha+1})$-competitive, it is sufficient to show that for any phase p

$$C_{\text{MTFR}}(p) + \Phi_F(p) - \Phi_B(p) \le \mathcal{O}(\ell^{\alpha+1}) \cdot C_{\text{OPT}}(p) \quad . \tag{17}$$

Let $F := d_0(P_{\text{MTFR}}(1), P_{\text{OPT}}(1))$ be the distance between the pages of MTFR and OPT in the very beginning of the phase. Let X_t be the distance between σ_t and $P_{\text{OPT}}(t)$. Finally, for any time step $1 \le t \le \ell_0$, let Y_t be the distance across which OPT moves its page in time t (0_E if it remains in one node). We consider two cases.

If $X_t = Y_t = 0_E$ for all t, then OPT remains at $P_{\text{OPT}}(1)$ for the whole phase p, paying 0. On the other hand, in the first step MTFR has to pay for serving request and moving the page to $\sigma_1 \equiv P_{\text{OPT}}(1)$, i.e.,

$$\begin{aligned} C_{\text{MTFR}}(1) &= (1 + D) \cdot c_1(P_{\text{MTFR}}(1), P_{\text{OPT}}(1)) \\ &\le 2 \cdot D \cdot e \cdot c_0(P_{\text{MTFR}}(1), P_{\text{OPT}}(1)) \\ &\le \Phi_B(p) \ , \end{aligned}$$

where the first inequality follows from (7). After this step the algorithm remains at the same node as OPT, paying 0. Thus, $\Phi_F(p) = 0$, and the lemma follows in this case.

Otherwise, for some time step t either $X_t \ne 0_E$ or $Y_t \ne 0_E$, and therefore $C_{\text{OPT}}(p) \ge 1$. Additionally, the cost of OPT in step t is at least $C_{\text{OPT}}(t) \ge X_t^\alpha + D \cdot Y_t^\alpha$. In the first step, the distance between MTFR and σ_t is at most $F + 1 + X_1$ (the distance from MTFR to OPT is at most $F + 1$). Thus, MTFR's cost in the first step is at most

$$\begin{aligned} C_{\text{MTFR}}(1) &\le (1 + D) \cdot [F + 1 + X_1]^\alpha \\ &\overset{(\text{H})}{\le} 2D + 2D \cdot 3^{\alpha-1} \cdot (F^\alpha + 1 + X_1^\alpha) \\ &\le \mathcal{O}(D) + \Phi_B(p) + \mathcal{O}(D) \cdot X_1^\alpha \\ &\le \mathcal{O}(D) \cdot C_{\text{OPT}}(p) + \Phi_B(p) \ . \end{aligned}$$

[3] If p is the first phase then $c_0 \equiv c_1$.

At the end of first step, MTFR moves to σ_1 and the distance between $P_{\text{MTFR}}(1)$ and $P_{\text{OPT}}(1)$ becomes at most $X_1 + Y_1$. In each step $t > 1$, the distance between $P_{\text{MTFR}}(t) \equiv P_{\text{MTFR}}(1)$ and $P_{\text{OPT}}(t)$ can increase by at most 1 due to the changes in the network made by the network adversary, and additionally at most Y_t due to the movement of the OPT page. Thus, for $t \in p$, $d_t(P_{\text{MTFR}}(t), P_{\text{OPT}}(t)) \leq X_1 + \ell_0 + \sum_{t=1}^{\ell_0} Y_t =: U$. We have

$$U^\alpha \overset{\text{(H)}}{=} O(1) \cdot \left(X_1^\alpha + \ell_0^\alpha + \left(\sum_{t=1}^{\ell_0} Y_t \right)^\alpha \right)$$
$$\overset{\text{(H)}}{=} O(1) \cdot \left(X_1^\alpha + \ell_0^\alpha + \ell_0^{\alpha-1} \cdot \sum_{t=1}^{\ell_0} Y_t^\alpha \right).$$

Therefore, for all $1 < t \leq \ell_0$ holds $C_{\text{MTFR}}(t) \leq 1 + (U + X_t)^\alpha \overset{\text{(H)}}{=} 1 + \mathcal{O}(1) \cdot [U^\alpha + X_t^\alpha]$. Since $\Phi_F(p) \leq f \cdot D \cdot [1 + U^\alpha]$, we have

$$\begin{aligned}
&\sum_{t=2}^{\ell_0} C_{\text{MTFR}}(t) + \Phi_F(p) \\
&\leq \mathcal{O}(\ell_0 + D) + \mathcal{O}(\ell_0 + D) \cdot U^\alpha + \mathcal{O}(1) \cdot \sum_{t=2}^{\ell_0} X_t^\alpha \\
&\leq \mathcal{O}(\ell) + \mathcal{O}(\ell) \cdot \left(X_1^\alpha + \ell_0^\alpha + \ell_0^{\alpha-1} \cdot \sum_{t=1}^{\ell_0} Y_t^\alpha \right) \\
&\quad + \mathcal{O}(1) \cdot \sum_{t=2}^{\ell_0} X_t^\alpha \ . \\
&\leq \mathcal{O}(\ell^{\alpha+1}) \cdot C_{\text{OPT}}(p)
\end{aligned}$$

which finishes the proof of $\mathcal{O}(\ell^{\alpha+1})$-competitiveness. $\square$

PROOF OF THEOREM 2. Now we can apply this result to compute the expected competitive ratio of MTFR on any input sequence $(\mathcal{C}_t)$ of length $T \geq T_{(\alpha+1)^2}$.

$$\begin{aligned}
\mathbf{E}_{(\sigma_t)} \left[\frac{C_{\text{ALG}}(\mathcal{C}_t, \sigma_t)}{C_{\text{OPT}}(\mathcal{C}_t, \sigma_t)} \right] &= \left(1 - \mathcal{O}\left(\frac{1}{D^{2 \cdot (\alpha+1)}} \right) \right) \cdot \mathcal{O}(1) \\
&\quad + \mathcal{O}\left(\frac{1}{D^{(\alpha+1)^2}} \right) \cdot \mathcal{O}(\ell^{\alpha+1}) \\
&= \mathcal{O}(1) \ ,
\end{aligned}$$

which finishes the proof of Theorem 2. $\square$

3. CONCLUDING REMARKS

Our result shows that the scenario considered in this paper is much more favorable than the one where both configuration and request sequences are created by an adversarial entity. This supports the claim that in fact the ratio in adversarial scenario is so high, because the network and sequence adversaries may combine and *synchronize* their efforts. The constant competitive ratio of stochastic scenario presented here follows from the fact that it is extremely unlikely that the random requests sequence contains *hard subsequences.* By hard subsequences we, informally, mean sequences similar to the ones used in construction of lower bounds against oblivious adversaries [7, 5].

On the other hand, it might be interesting to consider another scenario where π is not equal for each time step, but depends, for example on π in the last step $t-1$. In particular, the case in which $\pi_t = A \cdot \pi_{t-1}$, for a fixed stochastic matrix A might be interesting. Such a modes captures the locality of accesses (e.g. if a processor v_a accesses the page, then either processor v_b or v_c will access it in the next step).

The complexity of such *Markovian scenario* remains unknown. However, we conjecture that if A does not contain entries which are equal to 1, then it is possible to construct an algorithm, which achieves a reasonably low competitive ratio.

Another open question is whether it is possible to construct an algorithm for the stochastic scenario, which will be better than $\mathcal{O}\left(D^{(\alpha+1)^2}\right)$-competitive in the worst case, still assuring constant ratio on expectation. Combining MTFR with algorithms for adversarial scenario [7, 5] might be challenging and yield interesting results.

4. REFERENCES

[1] B. Awerbuch, Y. Bartal, and A. Fiat. Competitive distributed file allocation. In *Proc. of the 25th ACM Symp. on Theory of Computing (STOC)*, pages 164–173, 1993.

[2] B. Awerbuch, Y. Bartal, and A. Fiat. Distributed paging for general networks. *Journal of Algorithms*, 28(1):67–104, 1998. Also appeared in *Proc. of the 7th SODA*, pages 574–583, 1996.

[3] Y. Bartal, M. Charikar, and P. Indyk. On page migration and other relaxed task systems. *Theoretical Computer Science*, 268(1):43–66, 2001. Also appeared in *Proc. of the 8th SODA*, pages 43–52, 1997.

[4] S. Ben-David, A. Borodin, R. M. Karp, G. Tardos, and A. Wigderson. On the power of randomization in online algorithms. In *Proc. of the 22nd ACM Symp. on Theory of Computing (STOC)*, pages 379–386, 1990.

[5] M. Bienkowski, M. Dynia, and M. Korzeniowski. Improved algorithms for dynamic page migration. In *Proc. of the 22nd Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 365–376, 2005.

[6] M. Bienkowski and M. Korzeniowski. Dynamic page migration under brownian motion. In *Proc. of the European Conf. in Parallel Processing (Euro-Par)*, 2005. To appear.

[7] M. Bienkowski, M. Korzeniowski, and F. Meyer auf der Heide. Fighting against two adversaries: Page migration in dynamic networks. In *Proc. of the 16th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pages 64–73, 2004.

[8] M. Bienkowski and F. Meyer auf der Heide. Page migration in dynamic networks. In *Proc. of the 30th Int. Symp. on Mathematical Foundations of Computer Science (MFCS)*, 2005. To appear.

[9] D. L. Black and D. D. Sleator. Competitive algorithms for replication and migration problems. Technical Report CMU-CS-89-201, Department of Computer Science, Carnegie-Mellon University, 1989.

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

[11] M. Chrobak, L. L. Larmore, N. Reingold, and J. Westbrook. Page migration algorithms using work functions. In *Proc. of the 4th Int. Symp. on Algorithms and Computation (ISAAC)*, pages 406–415, 1993.

[12] I. S. Gradshteyn, I. M. Ryzhik, A. Jeffrey, and D. Zwillinger. *Table of Integrals, Series, and Products.* San Diego, CA: Academic Press, 6th edition, 2000.

[13] E. Koutsoupias and C. H. Papadimitriou. Beyond competitive analysis. 30(1):300–317, 2000.

[14] C. Lund, N. Reingold, J. Westbrook, and D. C. K. Yan. Competitive on-line algorithms for distributed data management. *SIAM Journal on Computing*, 28(3):1086–1111, 1999. Also appeared as On-Line Distributed Data Management in *Proc. of the 2nd ESA*, pages 202–214, 1994.

[15] S. Rajesekaran, P. M. Pardalos, J. H. Reif, and J. Rolim. *Handbook of Randomized Computing*, volume II. Kluwer Academic Publishers, 2001.

[16] T. S. Rappaport. *Wireless Communications: Principles and Practices.* Prentice Hall, 1996.

[17] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.

[18] J. Westbrook. Randomized algorithms for multiprocessor page migration. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 7:135–150, 1992.

APPENDIX

A. PROOFS OF TECHNICAL LEMMAS

Proof of Inequality 6. We define k'_t as $\sqrt[\alpha]{K_t/\beta - 1}$. Equivalently $K_t = \beta \cdot (1 + (k'_t)^\alpha)$ and $k_t^\alpha = 1 + (k'_t)^\alpha$. Let $\delta = 1/\alpha$. As an intermediate step we will prove that $k'_{t+1} - k'_t \leq \delta$; later we will use it to prove $k_{t+1} - k_t \leq \delta$.

From the definition of k'_t we have $\beta \cdot (k'_t)^\alpha = \sum_i \sum_{j \neq i} \pi(i) \cdot \pi(j) \cdot d_t^\alpha(v_i, v_j)$. For succinctness of the proof we denote $\frac{1}{\beta} \cdot \pi(i) \cdot \pi(j)$ by $p_{i,j}$. Then

$$(k'_t)^\alpha = \sum_{i \neq j} p_{i,j} \cdot (d_t(v_i, v_j))^\alpha \ ,$$

and $\sum_{i \neq j} p_{i,j} = 1$. We will prove that $k'_{t+1} \leq k'_t + \delta$.

We fix any constant non-negative integer $s < \alpha$, and consider variables $y_t(i,j) := (d_t(v_i, v_j))^{\alpha - s}$ for all $i \neq j$ where $i, j \in [n]$. Since the function $f(y) = y^{\frac{\alpha}{\alpha - s}}$ is convex and $\sum_{i \neq j} p_{i,j} = 1$, we can apply Jensen's Inequality [12] for the $y_t(i,j)$ variables and get that

$$\left[\sum_{i \neq j} p_{i,j} \cdot y_t(i,j) \right]^{\frac{\alpha}{\alpha - s}} \leq \sum_{i \neq j} p_{i,j} \cdot [\, y_t(i,j) \,]^{\frac{\alpha}{\alpha - s}} \ ,$$

and thus raising both sides to power $\frac{\alpha - s}{\alpha}$ we obtain

$$\begin{aligned} \sum_{i \neq j} p_{i,j} \cdot [\, d_t(v_i, v_j) \,]^{\alpha - s} &= \left[\sum_{i \neq j} p_{i,j} \cdot [\, d_t(v_i, v_j) \,]^\alpha \right]^{\frac{\alpha - s}{\alpha}} \\ &= (k'_t)^{\alpha - s} \ . \end{aligned}$$

Note that the inequality above holds trivially also for the case $s = \alpha$. Hence, we can multiply both sides of the above inequality by $\binom{\alpha}{s} \cdot \delta^s$ and sum them over all $s \in \{0, \ldots, \alpha\}$.

$$\sum_{s=0}^{\alpha} \sum_{i \neq j} p_{i,j} \cdot \binom{\alpha}{s} \cdot [\, d_t(v_i, v_j) \,]^{\alpha - s} \cdot \delta^s \leq \sum_{s=0}^{\alpha} \binom{\alpha}{s} \cdot (k'_t)^{\alpha - s} \cdot \delta^s$$

By folding the binomial formula, we get

$$\sum_{i \neq j} p_{i,j} \cdot [\, d_t(v_i, v_j) + \delta \,]^\alpha \leq \left(k'_t + \delta\right)^\alpha \ .$$

But from the definition of k'_{t+1} we have

$$(k'_{t+1})^\alpha = \sum_{i \neq j} p_{i,j} \cdot [\, d_{t+1}(v_i, v_j) \,]^\alpha \leq \sum_{i \neq j} p_{i,j} \cdot [\, d_t(v_i, v_j) + \delta \,]^\alpha$$

Combining these two inequalities above, and taking α-th root from both sides, we finally get $k'_{t+1} \leq k'_t + \delta$.

We have $k_{t+1} - k_t = \sqrt[\alpha]{(k'_{t+1})^\alpha + 1} - \sqrt[\alpha]{(k'_t)^\alpha + 1} \leq \sqrt[\alpha]{(k'_t + \delta)^\alpha + 1} - \sqrt[\alpha]{(k'_t)^\alpha + 1} =: g(k'_t)$. Function g is monotonically increasing, as its first derivative is greater than 0 for $k'_t, \delta \geq 0$. Additionally, $\lim_{k'_t \to \infty} g(k'_t) = \delta$, and therefore $g(k_t) < \delta$ for all k'_t.

The proof of $k_t - k_{t+1} \leq \delta$ is analogous, and thus Inequality 6 follows. $\square$

Proof of Lemma 4. Fix any time step t. Without loss of generality we can assume that the largest communication cost occurs between the nodes v_1 and v_2, i.e. $c_t(v_1, v_2) = c_t^{\max}$. Then we have

$$\begin{aligned} K_t &= \sum_i \sum_j \pi(i) \cdot \pi(j) \cdot c_t(v_i, v_j) \\ &\geq \pi(1) \cdot \pi(2) \cdot c_t(v_1, v_2) + \pi(2) \cdot \pi(1) \cdot c_t(v_2, v_1) \\ &\quad + \sum_{i=3}^{n} [\pi(i) \cdot \pi(1) \cdot c_t(v_1, v_i) + \pi(i) \cdot \pi(2) \cdot c_t(v_i, v_2)] \\ &\geq \pi(1) \cdot \pi_{\min} \cdot c_t(v_1, v_2) + \pi(2) \cdot \pi_{\min} \cdot c_t(v_1, v_2) \\ &\quad + \sum_{i=3}^{n} \pi(i) \cdot \pi_{\min} \cdot [c_t(v_i, v_1) + c_t(v_i, v_2)] \end{aligned}$$

By an argument similar to the one used in Lemma 11 we get $c_t(v_1, v_i) + c_t(v_i, v_2) \geq \frac{1}{2^{\alpha - 1}} \cdot c_t(v_1, v_2)$, and thus

$$\begin{aligned} K_t &\geq \frac{1}{2^{\alpha - 1}} \cdot \pi_{\min} \cdot c_t(v_1, v_2) \cdot \sum_{i=1}^{n} \pi(i) \\ &= \frac{1}{2^{\alpha - 1}} \cdot \pi_{\min} \cdot c_t^{\max} \ , \end{aligned}$$

which finishes the proof. $\square$

Proof of Lemma 7. First, we scale down all elements of sequence $\{A_i\}$ by dividing them by $\min\{A_i\}$. Note that the ratio $R := \frac{\left(\sum_{i=1}^m A_i^\alpha\right)^2}{\sum_{i=1}^m (A_i^\alpha)^2}$ remains invariant, the property $|A_{i+1} - A_i| \leq \delta$ holds, and after scaling we have $\min\{A_i\} = 1$. Let $k = \max\{A_i\}$ and let S be the smallest possible (in terms of the number of elements) subset of $\{A_i\}$ with the following properties.

- $1, k \in S$.
- Let $\{S_i\}$ be the sequence of all elements from set S, sorted in non-descending order. Then for a pair of consecutive elements S_i and S_{i+1} holds $|S_{i+1} - S_i| \leq \delta$.

For any integer j, let I_j denote interval $(\delta(j-1), \delta j]$. and let $\kappa := \lceil k/\delta \rceil$. It is straightforward that all elements of the sequence $\{S_i\}$ belong to the $\biguplus_{j=1}^{\kappa} I_j$. Furthermore, each interval I_j from this union contains at least one element of

$\{S_i\}$ (from the second property of set S) and at most two elements of $\{S_i\}$ (from the minimality of S).

In the trivial case, $\kappa = 1$, all elements A_i are between 1 and δ. Therefore,

$$R = \frac{\left(\sum_{i=1}^m A_i^\alpha\right)^2}{\sum_{i=1}^m (A_i^\alpha)^2} \geq \frac{\left(\sum_{i=1}^m A_i^\alpha\right)^2}{\max_i\{A_i^\alpha\} \cdot \sum_{i=1}^m A_i^\alpha} \geq \frac{m}{\delta^\alpha} \geq \left(\frac{m}{\delta^\alpha}\right)^{\frac{1}{\alpha+1}} .$$

In the general case, $\kappa > 1$ and we have

$$\sum_{i=1}^{|S|} S_i^\alpha \geq \sum_{j=1}^{\kappa} [\delta(j-1)]^\alpha \geq a_1 \cdot \delta^\alpha \cdot \kappa^{\alpha+1} ,$$

for some constant a_1, which depends only on α. On the other hand,

$$\sum_{i=1}^{|S|} S_i^{2\alpha} \leq \sum_{j=1}^{\kappa} 2\,(\delta \cdot j)^{2\alpha} \leq a_2 \cdot \delta^{2\alpha} \cdot \kappa^{2\alpha+1} ,$$

for some constant a_2. Thus,

$$R = \frac{\left(\sum_{i\in S} A_i^\alpha + \sum_{i\notin S} A_i^\alpha\right)^2}{\sum_{i\in S} A_i^{2\alpha} + \sum_{i\notin S} A_i^{2\alpha}} \geq \frac{\left(a_1 \cdot \delta^\alpha \cdot \kappa^{\alpha+1} + \sum_{i\notin S} A_i^\alpha\right)^2}{a_2 \cdot \delta^{2\alpha} \cdot \kappa^{2\alpha+1} + \sum_{i\notin S} A_i^{2\alpha}} .$$

Since $\sum_{i\notin S} A_i^{2\alpha} \leq \max\{A_i^\alpha\} \cdot \sum_{i\notin S} A_i^\alpha \leq k^\alpha \cdot \sum_{i\notin S} A_i^\alpha \leq (2\delta \cdot \kappa)^\alpha \cdot \sum_{i\notin S} A_i^\alpha$, we obtain

$$R \geq \frac{a_1^2 \delta^{2\alpha} \kappa^{2\alpha+2} + 2a_1 \delta^\alpha \kappa^{\alpha+1} \sum_{i\notin S} A_i^\alpha + \left(\sum_{i\notin S} A_i^\alpha\right)^2}{a_2 \delta^{2\alpha} \kappa^{2\alpha+1} + 2^\alpha \delta^\alpha \kappa^\alpha \sum_{i\notin S} A_i^\alpha} .$$

By omitting either the first of the third term from the numerator above, we get $R \geq a_3 \cdot \kappa$ and $R \geq a_4 \cdot \frac{\sum_{i\notin S} A_i^\alpha}{\delta^\alpha \cdot \kappa^\alpha}$, for some constants a_3, a_4. Thus if $\kappa \geq (m/\delta^\alpha)^{\frac{1}{\alpha+1}}$, then the lemma follows immediately. Otherwise, S contains at most $2 \cdot \kappa \leq 2 \cdot m^{\frac{1}{\alpha+1}}$ elements, and thus $\sum_{i\notin S} A_i^\alpha \geq 1 \cdot (m - 2 \cdot m^{\frac{1}{\alpha+1}}) \geq \frac{m}{2}$. Therefore, in this case

$$R \geq \frac{a_4}{2} \cdot \frac{m}{\delta^\alpha \cdot \kappa^\alpha} \geq \frac{a_4}{2} \cdot (m/\delta^\alpha)^{\frac{1}{\alpha+1}} ,$$

and the lemma holds. □

Proof of Lemma 8. We have $Q = \beta \cdot k_{t_0}^\alpha$. Additionally by Inequality 6 follows that for any t, $|k_{t+1} - k_t| \leq 1$.

First, we prove $\sum_{t=1}^{s} K_t = \mathcal{O}(\beta \cdot s^{\alpha+1} + s \cdot Q)$. We proceed with case analysis.

1. If $k_{t_0} \leq s$, then for any t, $k_t \leq s + s = 2 \cdot s$, and thus $\sum_{t=1}^{s} K_t = \beta \cdot s \cdot (2 \cdot s)^\alpha = \mathcal{O}(\beta \cdot s^{\alpha+1})$.
2. If $k_{t_0} > s$, then $Q > \beta \cdot s^\alpha$ and then for any t, $k_t \leq (Q/\beta)^{\frac{1}{\alpha}} + s \leq 2 \cdot (Q/\beta)^{\frac{1}{\alpha}}$. Hence, all values of K_t are at most $2^\alpha \cdot Q$, and thus $\sum_{t=1}^{s} K_t = \mathcal{O}(s \cdot Q)$.

Second, we prove $\sum_{t=1}^{s} K_t = \Omega(\beta \cdot s + s^{\frac{1}{\alpha+1}} \cdot Q)$. Since all $k_t \geq 1$, $\sum_{t=1}^{s} K_t = \Omega(\beta \cdot s)$. Thus, it is sufficient to prove that $\sum_{t=1}^{s} K_t = \Omega(s^{\frac{1}{\alpha+1}} \cdot Q)$ for $Q \geq \beta \cdot s^{\frac{\alpha}{\alpha+1}}$. In this case $k_{t_0} = s^{\frac{1}{\alpha+1}}$, and thus there exist at least $\frac{1}{2} \cdot s^{\frac{1}{\alpha+1}}$ time steps t in which $k_t \geq \frac{1}{2} \cdot s^{\frac{1}{\alpha+1}}$, or alternatively, the value of K_t is at least $\frac{1}{2^\alpha} \cdot Q$. Thus, $\sum_{t=1}^{s} K_t \geq \frac{1}{2} \cdot s^{\frac{1}{\alpha+1}} \cdot \frac{1}{2^\alpha} \cdot Q = \Omega(s^{\frac{1}{\alpha+1}} \cdot Q)$. This finishes the proof of the second part of the lemma. □

Proof of Lemma 10. For $s = 1$ the lemma follows trivially. Hölder's Inequality [12] states that for any $p, q > 1$ s.t. $\frac{1}{p} + \frac{1}{q} = 1$ and for any non-negative sequences $(a_i)_{i=1}^k$ and $(b_i)_{i=1}^k$ holds $\sum_i (a_i \cdot b_i) \leq (\sum_i a_i^p)^{1/p} \cdot (\sum_i b_i^q)^{1/q}$. By setting $p = s$, $q = \frac{s}{s-1}$ and $b_i = 1$ for all i we obtain

$$\sum_{i=1}^{k} a_i \leq \left(\sum_{i=1}^{k} a_i^s\right)^{1/s} \left(\sum_{i=1}^{k} 1\right)^{(s-1)/s} .$$

By raising both sides to the s-th power we get the lemma. □

Lemma 11 (Skewed Triangle Inequality). *Choose any four nodes $A, B, C, D \in \mathcal{X}$ and any time step t. Then it holds $c_t(A,B) + c_t(B,C) + c_t(C,D) \geq \frac{1}{3^{\alpha-1}} \cdot c_t(A,D)$.*

Proof. If $A \equiv D$ then $c_t(A,D) = 0$ and the inequality follows trivially. Otherwise, at least one of inequalities $A \not\equiv B$, $B \not\equiv C$, $C \not\equiv D$ holds. In either case

$$\begin{aligned} c_t(A,B) &+ c_t(B,C) + c_t(C,D) \\ &\geq d_t^\alpha(A,B) + d_t^\alpha(B,C) + d_t^\alpha(C,D) + 1 . \end{aligned}$$

Applying Lemma 10 we get

$$\begin{aligned} c_t(A,B) &+ c_t(B,C) + c_t(C,D) \\ &\geq \frac{1}{3^{\alpha-1}} \cdot (d_t(A,B) + d_t(B,C) + d_t(C,D))^\alpha + 1 \\ &\geq \frac{d_t^\alpha(A,D)}{3^{\alpha-1}} + 1 \\ &> \frac{1}{3^{\alpha-1}} \cdot c_t(A,D) . \end{aligned}$$

This concludes the proof. □