Noname manuscript No.
(will be inserted by the editor)

Concurrent Imitation Dynamics in Congestion Games

Heiner Ackermann · Petra Berenbrink · Simon Fischer · Martin Hoefer

Received: date / Accepted: date

Abstract Imitating successful behavior is a natural and frequently applied approach when facing complex decision problems. In this paper, we design protocols for distributed latency minimization in atomic congestion games based on imitation. We propose to study concurrent dynamics that emerge when each agent samples another agent and possibly imitates this agent's strategy if the anticipated latency gain is sufficiently large. Our focus is on convergence properties.

We show convergence in a monotonic fashion to stable states, in which none of the agents can improve their latency by imitating others. As our main result, we show rapid convergence to approximate equilibria, in which only a small fraction of agents sustains a latency significantly above or below average. Imitation dynamics behave like an FPTAS, and the convergence time depends only logarithmically on the number of agents.

An extended abstract of this work has been accepted for publication in the proceedings of the 28th Symposium on Principles of Distributed Computing (PODC 2009). This work was in part supported by DFG through Cluster of Excellence MMCI and UMIC Research Centre at RWTH Aachen University, and by an NSERC grant. Part of this work was done while the authors were at RWTH Aachen University.

Heiner Ackermann
Fraunhofer ITWM, Kaiserslautern, Germany.
E-mail: ackermann@itwm.fraunhofer.de

Petra Berenbrink
Simon Fraser University, Burnaby, Canada.
E-mail: petra@cs.sfu.ca

Simon Fischer
Rapid-I GmbH, Dortmund, Germany.
E-mail: fischer@rapid-i.com

Martin Hoefer
Max-Planck-Institut für Informatik and Saarland University, Saarbrücken, Germany.
E-mail: mhoefer@mpi-inf.mpg.de

Imitation processes cannot discover unused strategies, and strategies may become extinct with non-zero probability. For singleton games we show that the probability of this event occurring is negligible. Additionally, we prove that the social cost of a stable state reached by our dynamics is not much worse than an optimal state in singleton games with linear latency functions.

We concentrate on the case of symmetric network congestion games, but our results do not use the network structure and continue to hold accordingly for general symmetric games. They even apply to asymmetric games when agents sample within the set of agents with the same strategy space.

Finally, we discuss how the protocol can be extended such that, in the long run, dynamics converge to a pure Nash equilibrium.

1 Introduction

We study imitation dynamics that emerge if myopic agents concurrently imitate each other in order to improve on their own situation. In scenarios for which agents have little or no experience upon which they can base their decisions, or in which precise knowledge about the available options and their consequences is absent, it is a good strategy to *imitate* successful behavior. Thus, it is not surprising that such behavior can frequently be observed, and has been studied intensively in economics and game theory [25, 34].

In this paper we use the imitation paradigm to design protocols and study dynamics in the context of symmetric congestion games [30]. As an example of such a game consider a network congestion game in which agents strive to allocate paths with minimum latency between the same source-sink pair in a network. The latency of a path equals the sum of the latencies of the edges in that path and the latency of an edge depends on the number of agents sharing it.

Our main focus is the design and analysis of a simple imitation protocol, which can be used by agents in decentralized scenarios for latency minimization. Using the protocol agents strive to improve their individual latencies over time by imitating others in a concurrent and round-based fashion. Our Imitation Protocol has several appealing properties: it is simple, stateless, based on local information, and is compatible with the selfish incentives of the agents. Thus, it is a well-suited tool for implementing load-balacing and latency minimization in distributed systems with decentralized control and non-cooperative agents, such as e.g. in the channel allocation process in wireless networks.

The Imitation Protocol consists of a sampling and a migration step. First, each agent samples another agent uniformly at random. Then he considers the latency gain that he would have by adopting the strategy of the sampled agent, under the assumption that no-one else changes his strategy. If this latency gain is not too small our agent adopts the sampled strategy with a *migration probability* mainly depending on the anticipated latency gain. The major technical challenge in designing such a concurrent protocol is to avoid *overshooting effects*. Overshooting occurs if too many agents sample other agents currently using the same strategy, and if all of them migrate towards it. In this case their latency might be greater than before the migration. In order to avoid overshooting, the migration probabilities have to be defined appropriately without sacrificing the benefit of concurrency. We propose to scale the migration probabilities by the *elasticity* of the latency functions in order to avoid overshooting. The elasticity of a function at point x describes the proportional growth of the function value as a result of a proportional growth of its argument. Note that in case of polynomial latency functions with positive coefficients and maximum degree d the elasticity is upper bounded by d.

A natural solution concept in this scenario is imitation-stability. A state is *imitation-stable* if no more improvements are possible based on the Imitation Protocol. We analyse convergence properties with respect to this solution concept.

1.1 Our Results

As our first result we prove that the Imitation Protocol succeeds in avoiding overshooting effects and converges in a monotonic fashion (Section 3). More precisely, we show that a well-known potential function (Rosenthal [30]) decreases on expectation as long as the system is not yet at an imitation-stable state. Thus, the potential is a *super-martingale* and eventually reaches a local minimum, corresponding to an imitation-stable state. Hence, as a corollary, we see that an imitation-stable state is reached in pseudopolynomial time.

Our main result, presented in Section 3.3, however, is a much stronger bound on the time to reach approximate imitation-stable states. What is a natural definition of approximately stable states in our setting? By repeatedly sampling other agents, an agent gets to know the average latency of the system. It is approximately satisfied, if it does not sustain a latency much larger than the average. Hence, we say that a state is approximately stable if almost all agents are almost satisfied. More precisely, we consider states in which at most a δ-fraction of the agents deviates by more that an ϵ-fraction (in any direction) from the average latency. We show that the expected time to reach such a state is polynomial in the inverse of the approximation parameters δ and ϵ as well as in the maximum elasticity of the latency functions, and logarithmic in the ratio between maximum and minimum potential. Hence, if the maximum latency of a path is fixed, the time is only logarithmic in the number of agents and independent of the size of the strategy space and the number of resources.

We complement these results by various lower bounds. First, it is clear that pseudopolynomial time is required to reach exact imitation-stable states. This follows from the fact that there exist states in which all latency improvements are arbitrarily small, resulting in arbitrarily small migration probabilities. Hence, already a single step may take pseudopolynomially long. As a concept of approximate stable states one could have required *all* agents to be approximately satisfied, rather than only all but a δ-fraction. This, however, would require to wait a polynomial number of rounds for the last agent to become approximately satisfied, as opposed to our logarithmic bound.

In addition, we will consider sequential imitation processes in which only one agent may move at a time. In Section 3.2 we extend a construction from [2] to show that there exist instances in which the shortest sequence of imitations that leads to an imitation-stable state is exponentially long.

The Imitation Protocol has one drawback: It is not innovative in the following sense. It might happen with small but non-zero probability that all agents currently using the same strategy P migrate towards other strategies and no other agent migrates towards P. In this case, the knowledge about the existence of strategy P is lost and cannot be regained. For singleton games, i. e., games in which each strategy is a singleton set, in which empty links have latency zero, we show in Section 4 that the probability of this event occurring in a polynomial number of rounds is negligible. An important consequence of this result is that the cost of a state to which the Imitation Protocol converges is, on expectation, not much worse than the cost of a Nash equilibrium. More precisely, for the case of linear latency functions the expected cost of a state to which the Imitation Protocol converges is within a constant factor of the optimal solution. While we conjecture that this results holds in general for singleton games, we are able to prove it here only for games, in which the optimum solution has a significant number of agents on every link.

Alternatively, in cases, in which convergence to a Nash equilibrium is required, we can adjust the dynamics and occasionally let agents use a suitably defined Exploration Protocol. Using such a protocol, agents sample other strategies directly instead of sampling them by looking at other agents. In Section 5 we show that a suitable definition of such a protocol and a suitable combination with the Imitation Protocol guarantee convergence to Nash equilibria in the long run.

To the best of our knowledge, this is the first work that considers concurrent protocols for atomic congestion games that are not restricted to parallel links or linear latency functions. Results similar to the ones presented here have been obtained for the non-atomic Wardrop model in [19] where the analysis is significantly simplified by the fact that probabilistic effects do not have to be taken into account.

1.2 Related Work

Rosenthal [30] proves that every congestion game possesses a Nash equilibrium, and that better-response dynamics converge to Nash equilibria. In these dynamics agents have complete knowledge, and, in every round, only a single agent deviates to a better strategy than it currently uses. Fabrikant et al. [14], however, observe that, in general, from an appropriately chosen initial state it takes exponentially many steps until agents finally reach an equilibrium. This negative result still holds in games with ϵ-greedy agents, i. e., in games in which agents only deviate if their latency decreases by a relative factor of at least $1 + \epsilon$ [2,10,32]. Moreover, Fabrikant et al. [14] prove that, in general, computing a Nash equilibrium is PLS-complete. Their result still holds in the case of asymmetric network congestion games. In addition, Skopalik and Vöcking [32] prove that even computing an approximate Nash equilibrium is PLS-complete. On the positive side, best response dynamics converge quickly in singleton and matroid congestion games [2,26]. Additionally, Chien and Sinclair [10] consider the convergence time of best-response dynamics to approximate Nash equilibria in symmetric games. They prove fast convergence to approximate Nash equilibria provided that the latency of a resource increases by at most a factor for each additional user. Finally, Goldberg [23] considers a protocol applied to a scenario where n weighted users assign load to m parallel links and the latency equals the load of a resource. In this protocol, randomly selected agents move sequentially, and migrate to a randomly selected resource if this improves their latency. The expected time to reach a Nash equilibrium is pseudopolynomial. Results considering other protocols and links with latency functions are presented in [12].

The social cost of (approximate) Nash equilibria in congestion games has been subject to numerous studies. The most prominent concept has been the price of anarchy [29], which is the ratio of the worst cost of any Nash equilibrium over the cost of an optimal assignment. For atomic games and linear latencies, Awerbuch et al. [4] and Christodoulou and Koutsoupias [11] show a tight bound of 2.5, which was later translated into a unified argument for more general equilibrium concepts by Roughgarden [31]. The special case of (weighted) singleton games has been of particularly strong interest, and we refer the reader to [33] for an introduction to the numerous results. In terms of dynamics, Awerbuch et al. [5] consider the number of best-response steps required to reach a desirable state, which has a social cost only a constant factor larger than that of a social optimum. They show that even in congestion games with linear latencies there exist exponentially long best-response sequences for reaching such a desirable state. In contrast, Fanelli et al. [15] show that for linear latency functions not all such sequences are exponentially long. They describe a particular class of best response sequences that reach a desirable state after at most $\Theta(n \log \log n)$ steps. For more results on weighted congestion games or sequences with player movement patterns see, e.g., [16, 17].

Recently, concurrent protocols have been studied in various models and under various assumptions. Even-Dar and Mansour [13] consider concurrent protocols in a setting where the links have speeds. However, their protocols require global knowledge in the sense that the users must be able to determine the set of underloaded and overloaded links. Given this knowledge, the convergence time is doubly logarithmic in the number of agents. In [6] the authors consider a distributed protocol for the case that the latency equals the load that does not rely on this knowledge. Their bounds on the convergence time are also doubly logarithmic in the number of agents but polynomial in the number of links. In [7] the results are generalized to the case of weighted jobs. In this case, the convergence time is only pseudopolynomial, i. e., polynomial in the number of users, links, and in the maximum weight. As another generalization, Berenbrink et al. [8] generalize the scenario to load balancing over networks, in which agents are weighted jobs moving from machine to machine via an underlying graph. The convergence times here are logarithmic in the number of agents and polynomial in the number of machines and the maximum machine speed or maximum weight. For some improved bounds and further extensions of this scenario see [3].

Fotakis et al. [21] consider a scenario with latency functions for every resource. Their protocol involves local coordination among the agents sharing a resource. For the family of games in which the number of agents asymptotically equals the number of resources they prove fast convergence to almost Nash equilibria. Intuitively, an almost Nash equilibrium is a state in which there are not too many too expensive and too cheap resources. In [18,1], a load balancing scenario is considered in which no information about the target resource is available. The authors present efficient protocols in which the migration probability depends purely on the cost of the currently selected strategy.

Finally, Kleinberg et al. have analyzed the performance of multiplicative-weights algorithms for no-regret learning in load balancing [28] and congestion games [27]. However, these dynamics usually only converge in the history of play and only to a stable *distribution* over states, such as mixed Nash or correlated equilibria.

Most of these protocols involve direct sampling of resources or machines instead of imitation of agents. While this is unavoidable for reaching Nash equilibria (and we also incorporate this with our exploration protocol), it severely limits the convergence speed for large and complex strategy spaces, as finding cheap strategies can already be a non-trivial task. Instead, our main insight in this paper is that the imitation process allows to propagate good known strategies much more quickly. For example, while the protocol in [21] is similar to ours in terms of potential function and migration probabilities, the authors assume (1) parallel links and (2) a local coordination that restricts the number of migrating agents to at most one per link. This avoids problems with large strategy spaces and also allows to simplify the analysis. By sampling machines, the protocol yields a convergence time depending on the slope of the latency functions. Instead, our imitation process allows to use the relate the convergence time to the elasticity, which can be significantly smaller than the slope.

In this sense, our work is close to [19], where the authors consider congestion games in the non-atomic Wardrop model with an infinite population of agents carrying an infinitesimal amount of load each. They consider a protocol similar to ours and prove that with respect to approximate equilibria it behaves like an FPTAS, i. e., it reaches an approximate equilibrium in time polynomial in the approximation parameters and the representation length of the instance. In contrast to our work the analysis of the continuous model does not have to take into account probabilistic effects.

Our protocol is based on the notion of imitation, a concept frequently applied in evolutionary game theory. For an introduction to imitation dynamics, see, e g., [25, 34].

2 Congestion Games and Imitation Dynamics

In this section, we provide a formal description of our model. We define congestion games in terms of networks, that is, the strategy space of each agent corresponds to the set of paths connecting a particular source-sink pair in a network. We use this terminology only for convenience, and our main results presented in subsequent chapters are *independent of the relation to networks*. They continue to hold for *general symmetric congestion games*, in which strategy spaces of agents might not be efficiently searchable. Furthermore, we introduce the slope and the elasticity of latency functions, and give a precise definition of the IMITATION PROTOCOL.

2.1 Symmetric (Network) Congestion Games

A symmetric network congestion game is given by a tuple $(G, (s, t), \mathcal{N}, (\ell_e)_{e\in E})$, where $G = (V, E)$ denotes a network with vertices V and m directed edges E, and $s \in V$ and $t \in V$ denote a source and a sink vertex. $\mathcal{N}$ denotes the set of n *agents*. For measuring the delay or cost caused by usage there is $(\ell_e)_{e\in E}$, a family of non-decreasing and differentiable *latency functions* $\ell_e \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. We assume that for all $e \in E$, the latency functions satisfy $\ell_e(x) > 0$ for all $x > 0$. The strategy space of all agents equals the set of paths $\mathcal{P}$ connecting the source s with the sink t. If G consists of two nodes s and t only, which are connected by a set of parallel links, then we call the game a *singleton game*. Let p denote the number of paths, i. e., $p = |\mathcal{P}|$. A *state* x of the game is a vector $(x_P)_{P\in\mathcal{P}}$ where x_P denotes the number of agents utilizing path P in state x, and $x_e = \sum_{P\ni e} x_P$ is the *congestion* of edge $e \in E$ in state x. The latency of edge e in state x is given by $\ell_e(x_e)$, and the latency of path $P \in \mathcal{P}$ is $\ell_P(x) = \sum_{e\in P} \ell_e(x_e)$. The latency of an agent is the latency of the path it chooses.

For brevity, for all $P \in \mathcal{P}$, let 1_P denote the p-dimensional unit vector with the one in position P. In state x an agent has an incentive to switch from path P to path Q if this would strictly decrease its latency, i. e., if $\ell_P(x) > \ell_Q(x + 1_Q - 1_P)$.

If no agent has an incentive to change its strategy, then x is at a *(pure) Nash equilibrium*.[1] It is well known [30], that the set of Nash equilibria corresponds to the set of states that minimize the potential function

$$\Phi(x) = \sum_{e\in E} \sum_{i=1}^{x_e} \ell_e(i) \ .$$

[1] We will restrict our attention to pure Nash equilibria throughout the paper.

In the following, let $\Phi^* = \min_x \Phi(x)$ be the minimum potential. Note that due to our definition of the latency functions $\Phi^* > 0$. For every path $P \in \mathcal{P}$ let

$$\ell_P^+(x) = \ell_P(x + 1_P) \ .$$

Note that for every path $Q \in \mathcal{P}$

$$\ell_P^+(x) \geq \ell_P(x + 1_P - 1_Q) \ .$$

Additionally, let

$$L_{\mathrm{av}}(x) = \sum_{P \in \mathcal{P}} \frac{x_P}{n} \ell_P(x)$$

denote the average latency of the paths in state x, and let

$$L_{\mathrm{av}}^+(x) = \sum_{P \in \mathcal{P}} \frac{x_P}{n} \ell_P(x + 1_P) \ .$$

Finally, let $\ell_{\max} = \max_x \max_{P \in \mathcal{P}} \ell_P(x)$ denote the maximum latency of any path. Throughout this paper, whenever we consider a fixed state x we simply drop the argument (x) from Φ, ℓ_P, ℓ_P^+, L_{av}, and L_{av}^+.

2.2 Elasticity and Slope

To bound the steepness of the latency functions and the effect that overshooting may have, we consider the elasticity of the latency functions. Let d denote an upper bound on the elasticity of the latency functions, i. e.,

$$d \geq \max_{e \in E} \sup_{x \in (0,n]} \left\{ \frac{\ell_e'(x) \cdot x}{\ell_e(x)} \right\} \ .$$

Now given a latency function with elasticity d, it holds that for any x and $\alpha \geq 1$, $\ell_e(\alpha\, x) \leq \ell_e(x) \cdot \alpha^d$ and for $0 \leq \alpha < 1$, $\ell_e(\alpha\, x) \geq \ell_e(x) \cdot \alpha^d$. As an example, the function $a\, x^d$ has elasticity d.

For almost empty resources, we will also need an upper bound on the slope of the latency functions. Let ν_e denote the maximum slope on almost empty edges, i. e., we define

$$\nu_e = \max_{x \in \{1,\ldots,d\}} \{\ell_e(x) - \ell_e(x-1)\} \ .$$

Finally, for $P \in \mathcal{P}$, let $\nu_P = \sum_{e \in P} \nu_e$ and choose ν such that $\nu \geq \max_{P \in \mathcal{P}} \nu_P$.

2.3 The Imitation Protocol

Our Imitation Protocol (Protocol 1) proceeds in two steps. First, an agent samples another agent uniformly at random. The agent then migrates with a certain probability from its old path P to the sampled path Q depending on the anticipated relative latency gain and on the elasticity of the latency functions. Our analysis concentrates on dynamics that result from the protocol being executed by the agents in parallel in a round-based fashion. These dynamics generate a sequence of states $x(0), x(1), \ldots$. The resulting dynamics converge to a state that is stable in the sense that imitation cannot produce further progress, i. e., $x(t+1) = x(t)$ with probability 1. Such a state is called an *imitation-stable state*. In other words, a state is imitation-stable if it is ϵ-Nash with $\epsilon = \nu$ with respect to the strategy space restricted to the current support. Here, ϵ-Nash means that no agent can improve its own payoff unilaterally by more than ϵ.

Protocol 1 Imitation Protocol, repeatedly executed by all agents.

Let P denote the path of the agent in state x.
Sample another agent uniformly at random. Let Q denote its path.
if $\ell_P - \ell_Q(x + 1_Q - 1_P) > \nu$ **then**
with probability

$$\mu_{PQ} = \frac{\lambda}{d} \cdot \frac{\ell_P(x) - \ell_Q(x + 1_Q - 1_P)}{\ell_P(x)}$$

migrate from path P to path Q.
end if

As discussed in the introduction, the main difficulty in the design of the protocol is to bound overshooting effects. To get an intuition of this problem, consider two parallel links of which the first has the constant latency function $\ell_1(x) = c$ and the second has the latency function $\ell_2(x) = x^d$. Recall that the elasticity of ℓ_2 is d. Furthermore, assume that only a small number of agents x_2 utilize link 2 whereas the majority of $n - x_2$ agents utilize link 1. Let $b = c - x_2^d > 0$ denote the latency difference between the two links. A simple calculation shows that using the protocol without the damping factor $1/d$, the expected latency increase on link 2 would be $\Omega(b \cdot d)$, overshooting the balanced state by a factor d. For this reason, we reduce the migration probability accordingly. The constant λ will be determined later.

Note that the arguments in the last paragraph hold for the *expected* load changes. Our protocol, however, has to take care of probabilistic effects, i. e., the realized migration vector may differ from its expectation. Typically, we can use the elasticity to bound the impact of this effect. However, if the congestion on an edge is very small, i. e., less than d, then the number of joining agents is not concentrated sharply enough around its expectation. In order to compensate for this, we add an additional requirement that agents only migrate if the anticipated latency gain is at least ν and use this to bound probabilistic effects if the congestion of the edge is less than d. Let us remark that we will see below (Theorem 4) that for a large class of singleton games it is very unlikely, that an edge will ever have a load of d or less, so the protocol will behave in the same way with high probability for a polynomial number of rounds even if this additional requirement is dropped.

3 General Strategy Spaces

In this section, we consider *imitation dynamics* that emerge if in each round agents concurrently apply the IMITATION PROTOCOL. At first, we observe that imitation dynamics converge to imitation stable states since in each round the potential $\Phi(x)$ decreases in expectation. From this result we derive a pseudopolynomial upper bound on the convergence time to imitation-stable states.

3.1 Convergence to Imitation-Stable States

Consider two states x and x' as well as a *migration vector* $\Delta x = (\Delta x_P)_{P \in \mathcal{P}}$ such that $x' = x + \Delta x$. We may imagine Δx as the result of one round of the IMITATION PROTOCOL although the following lemma is independent of how Δx is constructed. Furthermore, we consider Δx to be composed of a set of migrations of agents between pairs of paths. We use x_{PQ} to denote the number of agents who switch from path P to path Q. Then, Δx_P denotes the total increase or decrease of the number of agents utilizing path P, that is,

$$\Delta x_P = \sum_{Q \in \mathcal{P}} (x_{QP} - x_{PQ}) \ .$$

Also, let $\Delta x_e = \sum_{P \ni e} \Delta x_P$ denote the induced change of the number of agents utilizing edge $e \in E$. In order to prove convergence, we define the *virtual potential gain*

$$V_{PQ}(x, \Delta x) = x_{PQ} \cdot (\ell_Q(x + 1_Q - 1_P) - \ell_P(x))$$

which is the sum of the potential gains each agent migrating from path P to path Q would contribute to $\Delta\Phi$ if each of them was the only migrating agent. Note that if an agent improves the latency of his path, the potential gain is negative. The sum of all virtual potential gains is a very rough lower bound on the true potential gain $\Delta\Phi(x, \Delta x) = \Phi(x + \Delta x) - \Phi(x)$. In order to compensate for the fact that agents concurrently change their strategies, consider the *error term* on an edge $e \in E$:

$$F_e(x, \Delta x) = \begin{cases} \displaystyle\sum_{u=x_e+1}^{x_e+\Delta x_e} \ell_e(u) - \ell_e(x_e + 1) & \text{if } \Delta x_e > 0 \\ \displaystyle\sum_{u=x_e+\Delta x_e+1}^{x_e} \ell_e(x_e) - \ell_e(u) & \text{if } \Delta x_e < 0 \\ 0 & \text{if } \Delta x_e = 0 \end{cases}$$

Subsequently, we show that the sum of the virtual potential gains and the error terms is indeed an upper bound on the true potential gain $\Delta\Phi(x, \Delta x)$. A similar result is shown in [20] for a continuous model.

Lemma 1 *For any assignment x and migration vector Δx it holds that*

$$\Delta\Phi(x, \Delta x) \leq \sum_{P,Q\in\mathcal{P}} V_{PQ}(x, \Delta x) + \sum_{e\in E} F_e(x, \Delta x) \ .$$

Proof We first express the virtual potential gain in terms of latencies on the edges. Clearly,

$$\begin{aligned}
\sum_{P,Q\in\mathcal{P}} V_{PQ}(x, \Delta x) &= \sum_{P,Q\in\mathcal{P}} x_{PQ} \cdot (\ell_Q(x + 1_Q - 1_P) - \ell_P(x)) \\
&= \sum_{P,Q\in\mathcal{P}} x_{PQ} \cdot \left(\sum_{e\in Q\setminus P} \ell_e(x_e + 1) - \sum_{e\in P\setminus Q} \ell_e(x_e) \right) \\
&\geq \sum_{e:\Delta x_e>0} \Delta x_e \cdot \ell_e(x_e + 1) + \sum_{e:\Delta x_e<0} \Delta x_e \cdot \ell_e(x_e) \ . \quad (1)
\end{aligned}$$

The true potential gain, however, is

$$\begin{aligned}
\Delta\Phi(x, \Delta x) &= \sum_{e:\Delta x_e>0} \sum_{u=x_e+1}^{x_e+\Delta x_e} \ell_e(u) - \sum_{e:\Delta x_e<0} \sum_{u=x_e+\Delta x_e+1}^{x_e} \ell_e(u) \\
&= \sum_{e:\Delta x_e>0} \left(\Delta x_e \cdot \ell_e(x_e + 1) + \sum_{u=x_e+1}^{x_e+\Delta x_e} (\ell_e(u) - \ell_e(x_e + 1)) \right) \\
&\quad + \sum_{e:\Delta x_e<0} \left(\Delta x_e \cdot \ell_e(x_e) + \sum_{u=x_e+\Delta x_e+1}^{x_e} (\ell_e(x_e) - \ell_e(u)) \right) \ .
\end{aligned}$$

Substituting Equation (1) for the left term of each sum and the definition of F_e for the right term of each sum, we obtain the claim of the lemma. □

In the following, we consider Δx to be a migration vector generated by the Imitation Protocol rather than an arbitrary vector. In this case, we denote ΔX as a random variable and all probabilities and expectations are taken with respect to the Imitation Protocol. In order to prove that the potential decreases in expectation, we derive a bound on the size of the error terms. We show that the error terms alter the virtual potential gain by at most a factor of two. Put another way, we show that in expectation the absolute value of the true potential gain is at least half of the absolute value of the virtual potential gain.

Lemma 2 *Let x denote a state and let ΔX denote a random migration vector generated by the* Imitation Protocol*. Then,*

$$\mathbb{E}\left[\Delta\Phi(x, \Delta X)\right] \leq \frac{1}{2} \sum_{P,Q\in\mathcal{P}} \mathbb{E}\left[V_{PQ}(x, \Delta X)\right] \ .$$

where the expectation is taken over the randomness of the IMITATION PROTOCOL.

Proof For any given round, each term in V_{PQ}, $P, Q \in \mathcal{P}$ and F_e, $e \in E$ can be associated with an agent. Fix an agent i migrating from, say, P to Q. Its contribution to the $V_{PQ}(x, \Delta X)$ is $-\Delta\,\ell_{PQ}(x) = \ell_Q(x+1_Q-1_P)-\ell_P(x)$ (this is the same for all agents moving from P to Q). It may also contribute to F_e, $e \in P \cup Q$. The general idea of the proof is to split and allocate $V_{PQ}(x, \Delta X)$ and $\sum_e F_e(x, \Delta X)$ in a suitable manner to migrating agents that shows that the total error reduces the absolute value of the total virtual potential gain by no more than a half.

While agent i contributes $-\Delta\,\ell_{PQ}$ to $V_{PQ}(x, \Delta X)$, bounding its contribution to $\sum_e F_e(x, \Delta X)$ to at most $\Delta\,\ell_{PQ}/2$ is non-trivial. The contribution depends on ΔX_e and whether i migrates towards or away from e. Deriving suitable upper and lower bounds on these contributions depending on whether i migrates towards e or away from e is the central technical challenge in the proof.

For characterizing the error contribution we consider subsets $\mathcal{N}' \subset \mathcal{N}$ of the agents and assume that they are ordered with respect to ascending migration probabilities $\mu_{P_j Q_j}$, in which P_j and Q_j denote the origin and destination path of agent $j \in \mathcal{N}'$. Ties are broken arbitrarily. Note that all such orderings can be derived solely using the state x that yields the values $\mu_{P_j Q_j}$. They are independent of the random choices made by the agents.

Recall that we fixed an agent i that migrates from P to Q and strive to determine its expected contribution to the error term. Now fix an edge $e \in Q \setminus P$. We let $A^+(e)$ denote the random set of agents migrating to $e \in Q \setminus P$. Let $\Delta\tilde{X}_e$ denote the random number of agents in $A^+(e)$ which occur in our ordering with respect to μ_{PQ} before agent i. Agent i's contribution to $F_e(x, \Delta X)$, $e \in Q \setminus P$, is upper bounded by $\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)$ where we define the error function $\Delta\tilde{\ell}_e(\delta) = \ell_e(x_e+1+\delta)-\ell_e(x_e+1)$. In this case, we forgot about the positive effects agents departing from e might have. For an illustration, see Figure 1. For brevity, let us write $\ell_e = \ell_e(x_e)$ and $\ell_e^+ = \ell_e(x_e+1)$ as well as $\ell_P = \ell_P(x)$ and $\ell_Q^+ = \ell_Q(x_e + 1_Q - 1_P)$. For $e \in Q \setminus P$ we show that

$$\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right] \le \frac{1}{8} \cdot (\ell_P - \ell_Q^+) \cdot \left(\frac{\ell_e^+}{\ell_Q^+} + \frac{\nu_e}{\nu_Q}\right) \quad . \tag{2}$$

Now fix an edge $e \in P \setminus Q$. Let $A^-(e)$ denote the random set of agents migrating away from $e \in P \setminus Q$. Let $\Delta\tilde{X}_e$ denote the random number of agents in $A^-(e)$ which occur in our ordering with respect to μ_{PQ} before agent i. Agent i's contribution to $F_e(x, \Delta X)$, $e \in P \setminus Q$ is lower bounded by $\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)$ where $\Delta\tilde{\ell}_e(\delta)$ is defined as above. Hence, we forgot about the positive effects agents migrating towards e might have. For $e \in P \setminus Q$ we show that

$$\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right] \le \frac{1}{8} \cdot (\ell_P - \ell_Q^+) \cdot \left(\frac{\ell_e}{\ell_P} + \frac{\nu_e}{\nu_P}\right) \quad . \tag{3}$$

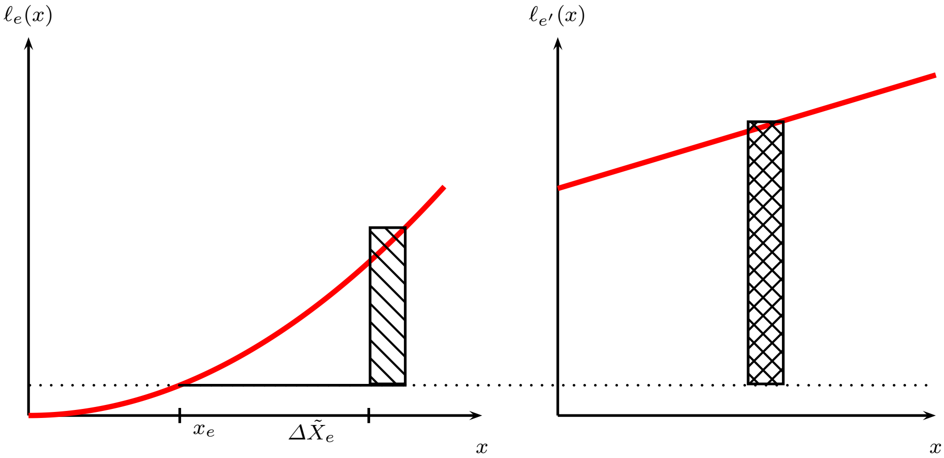

Fig. 1 Potential gain of an agent migrating from edge e' towards edge e. The hatched area is the agent's virtual potential gain. The shaded area on the left is this agents contribution to the error term, caused by the $\Delta\tilde{X}_e$ agents ranking before the agent under consideration (with respect to μ_{PQ}).

Thus, the expected sum of the error terms of an agent migrating from P to Q is at most

$$\frac{\ell_P - \ell_Q^+}{8}\left(\sum_{e\in P\setminus Q}\left(\frac{\ell_e}{\ell_P}+\frac{\nu_e}{\nu_P}\right)+\left(\sum_{e\in Q\setminus P}\frac{\ell_e^+}{\ell_Q^+}+\frac{\nu_e}{\nu_Q}\right)\right)\leq\frac{1}{2}(\ell_P-\ell_Q^+)\ ,$$

i. e., half of its virtual potential gain, which proves the lemma. We now proceed to prove Inequality (2). The case of Inequality (3) is very similar.

Consider $e \in Q \setminus P$ where Q denotes the destination path of agent i. For brevity, let us write $I_{PQ} = (\ell_P - \ell_Q^+)/\ell_P$ for the incentive to migrate from P to Q. Then, due to our ordering of the agents,

$$\mathbb{E}\left[\Delta\tilde{X}_e\right] \leq n\cdot\frac{x_e}{n}\cdot\mu_{PQ}\leq\frac{\lambda\cdot x_e\cdot I_{PQ}}{d}\ , \tag{4}$$

implying

$$x_e \geq \frac{\mathbb{E}\left[\Delta\tilde{X}_e\right]\cdot d}{\lambda\cdot I_{PQ}}\ . \tag{5}$$

Furthermore, due to the elasticity of ℓ_e, and using $(1+1/x)^x \leq \exp(1)$, we obtain

$$\begin{aligned}\Delta\tilde{\ell}_e(\delta) &\leq \ell_e^+ \cdot \left(\frac{x_e+1+\delta}{x_e+1}\right)^d - \ell_e^+ \\ &\leq \ell_e^+ \cdot \left(1+\frac{\delta}{x_e}\right)^d - \ell_e^+ \\ &\leq \ell_e^+ \cdot \left(\mathrm{e}^{\frac{d\,\delta}{x_e}} - 1\right) \ . \end{aligned} \tag{6}$$

Subsequently, we consider two cases.

Case 1: $\mathbb{E}\left[\Delta\tilde{X}_e\right] \geq \frac{1}{64}$. Substituting Inequality (5) into Inequality (6), we obtain for every $\kappa \in \mathbb{R}_{\geq 0}$

$$\Delta\tilde{\ell}_e\left(\kappa\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right) \leq \ell_e^+ \cdot \left(\mathrm{e}^{\kappa\,\lambda\,I_{PQ}} - 1\right) \ .$$

Now, note that for every $k \in \mathbb{N}$ and $\kappa \in [k, k+1]$

$$\begin{aligned}\mathbb{P}\left[\Delta\tilde{X}_e \geq \kappa\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right] &\leq \mathbb{P}\left[\Delta\tilde{X}_e \geq k\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right] \quad \text{and} \\ \Delta\tilde{\ell}_e(\kappa\,\mathbb{E}\left[\Delta\tilde{X}_e\right]) &\leq \Delta\tilde{\ell}_e((k+1)\,\mathbb{E}\left[\Delta\tilde{X}_e\right])\end{aligned}$$

hold. Applying a Chernoff bound (Fact 7 in the appendix), we obtain an upper bound for the expectation of $\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right]$ as follows.

$$\begin{aligned}&\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right] \\ &\leq \sum_{k=1}^{\infty} \mathbb{P}\left[\Delta\tilde{X}_e \geq k\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right] \cdot \Delta\tilde{\ell}_e((k+1)\,\mathbb{E}\left[\Delta\tilde{X}_e\right]) \\ &\leq \Delta\tilde{\ell}_e\left(5\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right) + \sum_{k=5}^{\infty} \mathbb{P}\left[\Delta\tilde{X}_e \geq k\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right] \cdot \Delta\tilde{\ell}_e((k+1)\,\mathbb{E}\left[\Delta\tilde{X}_e\right]) \\ &\leq \ell_e^+ \cdot \left(\mathrm{e}^{5\,\lambda\,I_{PQ}} - 1\right) + \sum_{k=5}^{\infty} \mathrm{e}^{-\frac{1}{4}\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\,k\,\ln k} \cdot \ell_e^+ \cdot \left(\mathrm{e}^{(k+1)\,\lambda\,I_{PQ}} - 1\right) \\ &\leq \ell_e^+ \cdot \left(\mathrm{e}^{5\,\lambda\,I_{PQ}} - 1\right) + \sum_{k=5}^{\infty} \mathrm{e}^{-\frac{1}{4}\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\,k} \cdot \ell_e^+ \cdot \left(\mathrm{e}^{2\,k\,\lambda\,I_{PQ}} - 1\right) \\ &\leq \ell_e^+ \cdot \left(\mathrm{e}^{5\,\lambda\,I_{PQ}} - 1\right) + \ell_e^+ \cdot \sum_{k=5}^{\infty} \mathrm{e}^{k(2\,\lambda\,I_{PQ} - \frac{1}{4}\,\mathbb{E}\left[\Delta\tilde{X}_e\right])} \\ &\leq \ell_e^+ \cdot \left(\mathrm{e}^{5\,\lambda\,I_{PQ}} - 1\right) + \ell_e^+ \cdot \int_4^{\infty} \mathrm{e}^{u(2\,\lambda\,I_{PQ} - \frac{1}{4}\,\mathbb{E}\left[\Delta\tilde{X}_e\right])} du \\ &= \ell_e^+ \cdot \left(\mathrm{e}^{5\lambda I_{PQ}} - 1 + \mathrm{e}^{-\mathbb{E}\left[\Delta\tilde{X}_e\right]} \frac{\mathrm{e}^{8\,\lambda\,I_{PQ}}}{\frac{1}{4}\mathbb{E}\left[\Delta\tilde{X}_e\right] - 2\,\lambda\,I_{PQ}}\right) \ .\end{aligned}$$

where the last inequality is true for $\lambda < 1/512$, because due to our assumption that $\mathbb{E}\left[\Delta\tilde{X}_e\right] \geq 1/64$ the sum term is monotonically decreasing. Now, due to Fact 8 (with $r = 1$, see appendix) and our assumption that $\mathbb{E}\left[\Delta\tilde{X}_e\right] \geq 1/64$, we obtain

$$\begin{aligned}
\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right] &\leq \lambda \cdot \ell_e^+ \cdot I_{PQ} \cdot \left(5\,(\mathrm{e}-1) + \frac{8\,(\mathrm{e}-1)}{\frac{1}{4 \cdot 64} - 2\,\lambda}\right) \\
&\leq c \cdot \lambda \cdot \ell_e^+ \cdot \frac{\ell_P - \ell_Q^+}{\ell_P} \\
&\leq c \cdot \lambda \cdot \ell_e^+ \cdot \frac{\ell_P - \ell_Q^+}{\ell_Q^+}
\end{aligned}$$

for some constant c. The first inequality holds for $\lambda < 1/512$, proving Equation (2) if λ is chosen small enough.

Case 2: $\mathbb{E}\left[\Delta\tilde{X}_e\right] < \frac{1}{64}$. Again, in this case we can apply a Chernoff bound (Fact 7 in the appendix) to upper bound $\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right]$.

$$\begin{aligned}
\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right] &\leq \sum_{k=1}^{n} \mathbb{P}\left[\Delta\tilde{X}_e = k\right] \cdot \Delta\tilde{\ell}_e(k) \\
&\leq \sum_{k=1}^{n} \mathbb{P}\left[\Delta\tilde{X}_e \geq \frac{k}{\mathbb{E}\left[\Delta\tilde{X}_e\right]}\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\right] \cdot \Delta\tilde{\ell}_e(k) \\
&\leq \sum_{k=1}^{n} \mathrm{e}^{-k\,(\ln(k/\mathbb{E}[\Delta\tilde{X}_e])-1)} \cdot \Delta\tilde{\ell}_e(k)
\end{aligned}$$

There are two sub-cases:

Case 2a: $x_e > d$. In order to bound the expected latency increase, we apply the elasticity bound on ℓ_e:

$$\begin{aligned}
&\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right] \\
&\leq \sum_{k=1}^{n} \mathrm{e}^{-k\,(\ln(k/\mathbb{E}[\Delta\tilde{X}_e])-1)} \cdot \ell_e^+ \cdot \left(\mathrm{e}^{\frac{k\,d}{x_e}} - 1\right) \\
&\leq \ell_e^+ \cdot \sum_{k=1}^{n} \mathrm{e}^{-k\,(\ln(k)-\ln(\mathbb{E}[\Delta\tilde{X}_e])-1)} \cdot \left(\mathrm{e}^{\frac{k\,d}{x_e}} - 1\right) \\
&\leq \ell_e^+ \cdot \sum_{k=1}^{n} \left(\mathbb{E}\left[\Delta\tilde{X}_e\right]\,(\mathrm{e}^k\,\mathbb{E}\left[\Delta\tilde{X}_e\right]^{k-1})\right)\,\mathrm{e}^{-k\,(\ln k)} \cdot \left(\mathrm{e}^{\frac{k\,d}{x_e}} - 1\right) \\
&\leq \ell_e^+ \cdot \mathbb{E}\left[\Delta\tilde{X}_e\right] \cdot \sum_{k=1}^{n} \mathrm{e}^{-k\,(\ln k)} \cdot \left(\mathrm{e}^{\frac{k\,d}{x_e}} - 1\right) \quad .
\end{aligned}$$

Now, splitting up the sum, we define

$$\begin{aligned}
L_1 &= \mathbb{E}\left[\Delta\tilde{X}_e\right] \sum_{k=1}^{\left\lfloor \frac{8\,x_e}{d} \right\rfloor} \mathrm{e}^{-k\,(\ln k)} \cdot \left(\mathrm{e}^{\frac{k\,d}{x_e}} - 1\right) \\
&\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \frac{(\mathrm{e}^8 - 1)\,d}{8\,x_e} \sum_{k=1}^{\left\lfloor \frac{8\,x_e}{d} \right\rfloor} \mathrm{e}^{-k\,(\ln k)} \cdot k \\
&\le \frac{\mathrm{e}^8}{4} \cdot \mathbb{E}\left[\Delta\tilde{X}_e\right] \frac{d}{x_e} \\
&\le \frac{\mathrm{e}^8}{4} \cdot \lambda\, I_{PQ} \quad ,
\end{aligned}$$

where the first inequality uses the observation that $\mathrm{e}^{\frac{k\,d}{x_e}} \le \mathrm{e}^8$ since $k \le \lfloor 8x_e/d \rfloor$, and Fact 8 in the appendix (with $r = 8$). Additionally, the second inequality uses the observation that $\sum_{k=1}^{\infty} \mathrm{e}^{-k\,(\ln k)} \cdot k \le 2$ (see Fact 9 in the appendix), and finally the last inequality uses Inequality (4).

For the second part of the sum, let

$$\begin{aligned}
L_2 &= \mathbb{E}\left[\Delta\tilde{X}_e\right] \sum_{k=\left\lceil \frac{8\,x_e}{d} \right\rceil}^{\infty} \mathrm{e}^{-k\,(\ln k)} \cdot \left(\mathrm{e}^{\frac{k\,d}{x_e}} - 1\right) \\
&\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \sum_{k=\left\lceil \frac{8\,x_e}{d} \right\rceil}^{\infty} \mathrm{e}^{-k\,(\ln k) + \frac{k\,d}{x_e}} \\
&\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \sum_{k=\left\lceil \frac{8\,x_e}{d} \right\rceil}^{\infty} \mathrm{e}^{-k\,(\ln k - 1)} && \text{(since } x_e > d\text{)} \\
&\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \sum_{k=\left\lceil \frac{8\,x_e}{d} \right\rceil}^{\infty} \mathrm{e}^{-\frac{1}{2}\,k\,\ln k} && \text{(since } k \ge \left\lceil \tfrac{8\,x_e}{d} \right\rceil \ge 8\text{)} \\
&\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \sum_{k=\left\lceil \frac{8\,x_e}{d} \right\rceil}^{\infty} \left(\frac{d}{8\,x_e}\right)^{\frac{1}{2}k} \quad .
\end{aligned}$$

Due to Fact 11 in the appendix and since $x_e > d$

$$\begin{aligned}
L_2 &\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \frac{\left(\frac{d}{8\,x_e}\right)^{\frac{8}{2}}}{1 - \sqrt{\frac{d}{8\,x_e}}} \\
&\le \mathbb{E}\left[\Delta\tilde{X}_e\right] \frac{d}{x_e} \\
&\le \lambda\, I_{PQ} \quad .
\end{aligned}$$

Reassembling the sum, we obtain

$$\begin{aligned}\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right] &\leq \ell_e^+ \cdot (L_1 + L_2) \\ &\leq \ell_e^+ \cdot \left(\frac{\mathrm{e}^8}{4} + 1\right) \lambda\, I_{PQ} \ .\end{aligned}$$

Again, by the same arguments as at the end of Case 1 this proves Equation (2) if λ is less than $1/(2\mathrm{e}^8 + 8)$.

Case 2b: $x_e \leq d$. In this case we separate the upper bound on $\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right]$ into the section up to d and above d. For the first section we use the fact that each additional agent on resource e causes a latency increase of at most ν_e as long as the load is at most d. We define the contribution to the expected latency increase by the events that up to $d - x_e$ join resource e, i. e., afterwards the congestion is still at most d. In this case, we may use ν_e to bound the contribution of each agent:

$$\begin{aligned}L_1 &\leq \sum_{k=1}^{d-x_e} \mathrm{e}^{-k\left(\ln\left(\frac{k}{\mathbb{E}[\Delta\tilde{X}_e]}\right)-1\right)} \cdot k\,\nu_e \\ &\leq \mathrm{e}\,\nu_e\,\mathbb{E}\left[\Delta\tilde{X}_e\right] + \nu_e\,\mathbb{E}\left[\Delta\tilde{X}_e\right]^2 \sum_{k=2}^{d-x_e} \mathrm{e}^{-k\,(\ln(k)-1)} \cdot k \\ &\leq \mathrm{e}\,\nu_e\,\mathbb{E}\left[\Delta\tilde{X}_e\right] \cdot \left(1 + \frac{8\,\mathbb{E}\left[\Delta\tilde{X}_e\right]}{\mathrm{e}}\right) \\ &\leq 3\,\nu_e\,\mathbb{E}\left[\Delta\tilde{X}_e\right] \ ,\end{aligned}$$

where the third inequality holds since $\sum_{k=2}^{d-x_e} \mathrm{e}^{-k\,(\ln(k)-1)} \cdot k \leq 8$ (see Fact 10 in the appendix), and where the last inequality holds since $\mathbb{E}\left[\Delta\tilde{X}_e\right] < 1/64$.

For the contribution of the agents increasing the load on resource e to above d we use the elasticity constraint again. This time, we do not consider the latency increase with respect to $\ell_e^+(x_e)$ but with respect to $\ell_e(d)$:

$$L_2 = \sum_{k=d-x_e+1}^{n} \mathrm{e}^{-k\cdot\left(\ln\left(\frac{k}{\mathbb{E}[\Delta\tilde{X}_e]}\right)-1\right)} \cdot \ell_e(d) \cdot \left(\mathrm{e}^{\frac{d\,(k-(d-x_e))}{d}} - 1\right) \ .$$

As in case (2a),

$$\begin{aligned} L_2 &\le \ell_e(d) \cdot \mathbb{E}\left[\Delta \tilde{X}_e\right] \cdot \sum_{k=d-x_e+1}^{\infty} \mathrm{e}^{-k\,\ln k + k - (d-x_e)} \\ &= \ell_e(d) \cdot \mathbb{E}\left[\Delta \tilde{X}_e\right] \cdot \sum_{k=1}^{\infty} \mathrm{e}^{-(k+(d-x_e))\,\ln(k+(d-x_e))+k} \\ &= \ell_e(d) \cdot \mathbb{E}\left[\Delta \tilde{X}_e\right] \cdot \mathrm{e}^{-(d-x_e)} \cdot \sum_{k=1}^{\infty} \mathrm{e}^{-(k+(d-x_e))\,\ln(k+(d-x_e))+k+d-x_e} \quad . \end{aligned}$$

Consider the series in the above expression as a function of $u = (d - x_e)$ and denote it by $S(u)$. Note that $S(u)$ converges for every $u \ge 0$ and $S(u) \to 0$ as $u \to \infty$. In particular, $S(u) = \sum_{k=1}^{\infty} (\mathrm{e}/(k+u))^{k+u}$ is monotonically decreasing in u, as for $k \ge 1$ each summand is monotonically decreasing in u. For $u = 0$, the series is upper bounded by the one in Fact 10 (see the appendix), and thus $S(u) \le 8$ for any $u \ge 0$. Therefore,

$$\begin{aligned} L_2 &\le 8\,\ell_e(d) \cdot \mathbb{E}\left[\Delta \tilde{X}_e\right] \cdot \mathrm{e}^{-(d-x_e)} \\ &\le 8\,(\ell_e(x_e) + (d-x_e)\,\nu_e) \cdot \mathbb{E}\left[\Delta \tilde{X}_e\right] \cdot \mathrm{e}^{-(d-x_e)} \quad . \end{aligned}$$

Since $(d - x_e) \cdot \mathrm{e}^{-(d-x_e)} < 1/2$,

$$L_2 \le 4\,(\ell_e(x_e) + \nu_e) \cdot \mathbb{E}\left[\Delta \tilde{X}_e\right] \quad .$$

Altogether,

$$\begin{aligned} \mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta \tilde{X}_e)\right] &\le L_1 + L_2 \\ &\le 7\,\nu_e\,\mathbb{E}\left[\Delta \tilde{X}_e\right] + 4\,\ell_e(x_e)\,\mathbb{E}\left[\Delta \tilde{X}_e\right] \\ &\le 7\,\nu_e\,\mathbb{E}\left[\Delta \tilde{X}_e\right] + 4\,\frac{\lambda\,x_e\,I_{PQ}}{d} \cdot \ell_e(x_e) \\ &\le \frac{7}{64}\,\nu\,\frac{\nu_e}{\nu_Q} + \frac{4\,\lambda\,x_e\,I_{PQ}}{d} \cdot \ell_e(x_e) \end{aligned}$$

where we have used Equation (4) for the third inequality, and the inequalities $\mathbb{E}\left[\Delta \tilde{X}_e\right] < 1/64$ and $\nu \ge \nu_Q$ for the last step. Since $x_e \le d$ and $\ell_P - \ell_Q^+ \ge \nu$,

$$\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta \tilde{X}_e)\right] \le \frac{1}{8}\,(\ell_P - \ell_Q^+)\,\frac{\nu_e}{\nu_Q} + \frac{4\,\lambda\,(\ell_P - \ell_Q^+)}{\ell_P} \cdot \ell_e(x_e)$$

again proving Equation (2) if $\lambda \le 1/32$.

Finally, the case $e \in P$ is very similar. $\square$

Note that all migrating agents add a negative contribution to the virtual potential gain since they migrate only from paths with currently higher latency to paths with lower latency. Hence, together with Lemma 2, we can derive the next corollary.

Corollary 1 *Consider a symmetric network congestion game Γ and let x and X' denote states of Γ such that X' is a random state generated after one round of executing the* Imitation Protocol. *Then,* $\mathbb{E}[\Phi(X')] \leq \Phi(x)$ *with strict inequality as long as x is not imitation-stable. Thus, Φ is a super-martingale.*

It is obvious that the sequence of states generated by the Imitation Protocol terminates at an imitation-stable state. From Lemma 2 we can immediately derive an upper bound on *the time to reach* such a state. However, since for arbitrary latency functions the minimum possible latency gain may be very small, this bound can clearly be only pseudo-polynomial. To see this, consider a state in which only one agent can make an improvement. Then, the expected time until the agent moves is inversely proportional to its latency gain.

Theorem 1 *Consider a symmetric network congestion game in which all agents use the* Imitation Protocol. *Let x denote the initial state of the dynamics. Then the dynamics converge to an imitation-stable state in expected time*

$$\mathcal{O}\left(\frac{d\,n\,\ell_{\max}\,\Phi(x)}{\nu^2}\right) .$$

Proof By definition of the Imitation Protocol, the expected virtual potential gain in any state x' which is not yet imitation-stable is bounded by

$$\mathbb{E}\left[\sum_{P,Q\in\mathcal{P}} V_{PQ}(x', \Delta X')\right] \leq -\nu \cdot \frac{\lambda}{d\,n} \cdot \frac{\nu}{\ell_{\max}} .$$

Hence, also the expected potential gain $\mathbb{E}[\Delta\Phi(X')]$ in every intermediate state x' of the dynamics is bounded from above by at least half of the above value. From this, it follows that the expected time until the potential drops from at most $\Phi(x)$ to the minimum potential $\Phi^* > 0$ is at most

$$\mathcal{O}\left(\frac{d\,n\,\ell_{\max}(\Phi(x) - \Phi^*)}{\lambda\,\nu^2}\right) .$$

Formally, this is a consequence of Lemma 7 in the Appendix. $\square$

It is obvious that this result cannot be significantly improved since we can easily construct an instance and a state such that the only possible improvement that can be made is ν. Hence, already a single step takes pseudopolynomially long. In case of polynomial latency functions Theorem 1 reads as follows.

Corollary 2 *Consider a symmetric network congestion game with polynomial latency functions with maximum degree d and minimum and maximum coefficients a_{min} and a_{max}, respectively. Then the dynamics converge to an imitation-stable state in expected time*

$$\mathcal{O}\left(d^3 m^2 n^{2d+2} \cdot \left(\frac{a_{max}}{a_{min}}\right)^2\right) .$$

In case of a singleton congestion game with monomial latency functions $\ell_e(x) = a \cdot x^{d_e}$ we can improve the corollary as follows.

Corollary 3 *Consider a symmetric singleton congestion game with monomial latency functions with maximum degree d. Then the dynamics converge to an imitation-stable state in expected time $\mathcal{O}\left(d^3 n^{2d+2}\right)$.*

Let us remark that all proofs in this section do not rely on the assumption that the underlying congestion game is symmetric or the underlying structure is a network. In fact, the results hold accordingly for general asymmetric congestion games in which each agent samples only among agents that have the same strategy space.

3.2 Sequential Imitation Dynamics and a Lower Bound

In the previous subsection we proved that agents applying the Imitation Protocol reach an imitation-stable state after a pseudopolynomial number of rounds. Recall that in this case each agent decreases its latency by at least ν if it were the only agent to change its strategy. In this section, we consider sequential imitation dynamics such that in each round a single agent is permitted to imitate someone else. Furthermore, we assume that each agent changes its path regardless of the anticipated latency gain. For the discussion in this section, we call states to be imitation-stable when no agent can improve by imitating the strategy of any other agent. Now, it is obvious that sequential imitation dynamics converge towards imitation-stable states as the potential Φ strictly decreases after every strategy change. Hence, we focus on the convergence time of such dynamics.

For such sequential imitation dynamics we prove an exponential lower bound on the number of rounds to reach an imitation-stable state. To be precise, we present a family of symmetric network congestion games with corresponding initial states such that every sequence of imitation leading to an imitation-stable state is exponentially long. To some extent, this result complements Theorem 1 as it presents an exponential lower bound in a slightly different model. However, in this lower bound ν is arbitrarily large and almost every state is imitation-stable with respect to the Imitation Protocol.

Theorem 2 *For every $n \in \mathbb{N}$, there exists a symmetric network congestion game with n agents, initial state S^{init}, polynomially bounded network size, and linear latency functions such that* every *sequential imitation dynamics that start in S^{init} is exponentially long.*

We do not give a complete proof of the theorem but we discuss how to adapt a series of constructions as presented in [2] which show that there exists a family of symmetric network congestion games with the same properties as stated in the above theorem such that *every best response dynamics* starting in S^{init} is exponentially long. To be precise, they prove that in every intermediate state of the best response dynamics *exactly* one agent can improve its latency. Recall that in best response dynamics agents know the entire strategy space and that in each round one agent is permitted to switch to the best available path.

In the following, we summarize the constructions presented in [2]. At first, a PLS-reduction from the local search variant of MaxCut to threshold games is presented. In a threshold game, each agent either allocates a single resource on its own or shares a bunch of resources with other agents. Hence, in a threshold game each agent chooses between two strategies only. The precise definition of these games is given below. Then, a PLS-reduction from threshold games to asymmetric network congestion games is presented. Finally, the authors of [2] show how to transform an asymmetric network congestion game into a symmetric one such that the desired properties of best response dynamics are preserved. All PLS-reductions are embedding reductions, and there exists a family of instances of MaxCut with corresponding initial configurations such that in every intermediate configuration generated by a local search algorithm exactly one node can be moved to the other side of the cut. Therefore, there exists a family of symmetric network congestion games with the properties as stated above.

A naive approach to prove a lower bound on the convergence time of imitation dynamics in symmetric network congestion games is as follows. Building upon the lower bound of the convergence time of best response dynamics, for every path an agent is added to the game. Then the latency functions are adopted accordingly. However, in this case we would introduce an exponential number of additional agents. In threshold games, however, the agents' strategy spaces only have size two. Hence, in the following we apply this approach to threshold games. It is then straightforward to verify that the PLS-reductions mentioned above can be reworked in order to prove Theorem 2. However, note that this does not imply that computing a imitation-stable state is PLS-complete since one can always assign all agents to the same strategy which obviously constitutes an imitation-stable state.

Threshold games are a special class of congestion games in which the set of resources $\mathcal{R}$ can be divided into two disjoint sets $\mathcal{R}_{\text{in}}$ and $\mathcal{R}_{\text{out}}$. The set $\mathcal{R}_{\text{out}}$ contains exactly one resource r_i for every agent $i \in \mathcal{N}$. This resource has a fixed latency T_i called the *threshold* of agent i. Each agent i has only two strategies, namely a strategy $S_i^{\text{out}} = \{r_i\}$ with $r_i \in \mathcal{R}_{\text{out}}$, and a strategy $S_i^{\text{in}} \subseteq \mathcal{R}_{\text{in}}$. The preferences of agent i can be described in a simple and intuitive way: Agent i prefers strategy S_i^{in} to strategy S_i^{out} if the latency of S_i^{in} is smaller than the threshold T_i. *Quadratic threshold games* are a subclass of threshold games in which the set $\mathcal{R}_{\text{in}}$ contains exactly one resource r_{ij} for every unordered pair of agents $\{i, j\} \subseteq \mathcal{N}$. Additionally, for every agent $i \in \mathcal{N}$ of a quadratic threshold game, $S_i^{\text{in}} = \{r_{ij} \mid j \in \mathcal{N}, j \neq i\}$. Moreover, for every resource $r_{ij} \in \mathcal{R}_{\text{in}}$: $\ell_{r_{ij}}(x) = a_{i,j} \cdot x$ with $a_{ij} \in \mathbb{N}$, and for every resource r_i: $\ell_{r_i}(x) = 1/2 \sum_{j \neq i} a_{ij} \cdot x$ to r_i.

Let Γ be a quadratic threshold game that has an initial state S^{init}, such that every best response dynamics which starts in S^{init} is exponentially long, and every intermediate state has a unique agent that can improve its latency. Suppose now that we replace every agent i in Γ by three agents i_1, i_2 and i_3 which all have the same strategy spaces as agent i has. Additionally, suppose that we choose new latency functions ℓ' for every resource r_i as follows: $\ell'_{r_i}(x) = 1/2 \sum_{j \neq i} a_{ij} \cdot x + 3/2 \sum_{j \neq i} a_{ij}$. Hence, we add an additional offset of $3/2 \sum_{j \neq i} a_{ij}$.

Suppose now that we assign every agent i_1 to S_i^{out}, and every agent i_2 to S_i^{in}. For every possible strategy that the i_3 agents can use, their latency increases by $2 \sum_{j \neq i} a_{ij}$, compared to the equivalent state in the original game, in which every agent i chooses the same strategy as agent i_3 does. Hence, if we assign every agent i_3 to the strategy chosen by agent i in S^{init} and if the agents i_1 and i_2 were not permitted to change their strategies, then we would obtain the desired lower bound on the convergence time of imitation dynamics in threshold games. However, since also i_1 and i_2 are permitted to imitate, it remains to show that whenever agent i_3 has changed its strategy, then both i_1 and i_2 do not want to change their strategies anymore.

First, suppose that agent i_3 switches from the strategy of agent i_2 to the strategy of agent i_1. Obviously, agent i_1 does not want to change its strategy as otherwise i_3 would not have imitated i_1. Suppose now that i_2, whose strategy is dropped by i_3, also wants to imitate i_1. In this case all three agents would allocate S_i^{out}, and hence have latency $3 \sum_{r \in j \neq i} a_{ij}$. However, if agent i_2 would stay with strategy S^{in} then its latency is upper bounded by $2 \sum_{r \in S_i^{\text{in}}} a_{ij}$. Hence, agents i_1, i_2, i_3 will never select S^{out} at the same time.

Second, suppose that agent i_3 switches from the strategy of agent i_1 to the strategy of agent i_2. Now, agent i_2 does not want to change its strategy as otherwise i_3 would not have imitated i_2. Suppose now that i_1, whose strategy is dropped by i_3, also wants to imitate i_3. In this case, the latency would increase to at least $3 \sum_{r \in j \neq i} a_{ij}$, whereas agent i_1 would have latency $2 \sum_{r \in j \neq i} a_{ij}$ if it would stay with strategy S^{out}. Hence, agents i_1, i_2, i_3 will never select S^{in} at the same time.

By applying the argument that all three agents never allocate the same strategy at the same point in time we can conclude our claim and Theorem 2 follows.

3.3 Convergence to Approximate Equilibria

Theorem 1 guarantees convergence of concurrent imitation dynamics generated by the IMITATION PROTOCOL to an imitation-stable state in the long run. However, it does not give a reasonable bound on the time due to the small progress that can be made. Hence, as our main result, we present bounds on the time to reach an *approximate equilibrium.* Here we relax the definition of a state that is imitation-stable (with respect to the Imitation Protocol) in two aspects: We allow only a small minority of agents to deviate by more than a small amount from the average latency. Our notion of an approximate equilibrium is similar to the notion used in [9,**?**,**?**]. It is motivated by the following observation. When sampling other agents each agent gets to know its latency if it would adopt that agent's strategy. Hence to some extent each agent can compute the average latency L^+_{av} and determine if its own latency is above or below that average.

Definition 1 ((δ,ϵ,ν)-equilibrium) Given a state x, let the set of *expensive* paths be $\mathcal{P}^+_{\epsilon,\nu} = \{P \in \mathcal{P} : \ell_P(x) > (1+\epsilon)\, L^+_{\mathrm{av}} + \nu\}$ and let the set of *cheap* paths be $\mathcal{P}^-_{\epsilon,\nu} = \{P \in \mathcal{P} : \ell_P(x) < (1-\epsilon)\, L_{\mathrm{av}} - \nu\}$. Let $\mathcal{P}_{\epsilon,\nu} = \mathcal{P}^+_{\epsilon,\nu} \cup \mathcal{P}^-_{\epsilon,\nu}$. A configuration x is at a *(δ,ϵ,ν)-equilibrium* iff it holds that $\sum_{P\in\mathcal{P}_{\epsilon,\nu}} x_P \le \delta \cdot n$.

Intuitively, a state at (δ,ϵ,ν)-equilibrium is a state in which almost all agents are almost satisfied when comparing their own situation with the situation of other agents. One may hope that it is possible to reach a state in which *all* agents are almost satisfied quickly. This would be a relaxation of the concept of Nash equilibrium. We will argue below, however, that there is no rapid convergence to such states.

Theorem 3 *For an arbitrary initial assignment $x(0)$, let τ denote the first round in which the* Imitation Protocol *reaches a (δ,ϵ,ν)-equilibrium. Then,*

$$\mathbb{E}\left[\tau\right] = \mathcal{O}\left(\frac{d}{\epsilon^2\,\delta} \cdot \log\left(\frac{\Phi(x(0))}{\Phi^*}\right)\right) \quad .$$

Proof We consider a state $x(t)$ that is not at a (δ,ϵ,ν)-equilibrium and derive a lower bound on the expected potential gain. There are two cases. Either at least half of the agents utilizing paths in $\mathcal{P}_{\epsilon,\nu}$ utilize paths in $\mathcal{P}^+_{\epsilon,\nu}$ or at least half of them utilize paths in $\mathcal{P}^-_{\epsilon,\nu}$.

Case 1: Many agents use expensive paths, i. e., $\sum_{P\in\mathcal{P}^+_{\epsilon,\nu}} x_P \ge \delta\, n/2$. Let us define the volume T and the average ex-post latency C of potential destination paths, i. e., paths with ex-post latency at most $(1+\epsilon)L^+_{\mathrm{av}}$, by

$$T = \sum_{Q:\ell^+_Q \le (1+\epsilon)L^+_{\mathrm{av}}} \frac{x_Q}{n} \quad \text{and} \quad C = \frac{1}{T} \sum_{Q:\ell^+_Q \le (1+\epsilon)L^+_{\mathrm{av}}} \frac{x_Q}{n}\ell^+_Q \ .$$

Clearly,

$$L^+_{\mathrm{av}} = \sum_P \frac{x_P}{n}\ell^+_P \ge T \cdot C + (1-T)\cdot(1+\epsilon)\, L^+_{\mathrm{av}} \ ,$$

and solving for T yields

$$T \ge \frac{\epsilon\, L^+_{\mathrm{av}}}{(1+\epsilon)\, L^+_{\mathrm{av}} - C} \quad . \tag{7}$$

We now give an upper bound on (i.e., a lower bound on the absolute value of) the expected virtual potential gain given that the current state is not at a (δ,ϵ,ν)-equilibrium. We consider only the contribution of agents utilizing paths in $\mathcal{P}^+_{\epsilon,\nu}$ and sampling paths with ex-post latency below $(1+\epsilon)\,L^+_{\mathrm{av}}$. Then,

$$\begin{aligned}\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] &\leq -\frac{\lambda}{d} \sum_{P \in \mathcal{P}^+_{\epsilon,\nu}} x_P \sum_{Q:\ell^+ \leq (1+\epsilon) L^+_{\mathrm{av}}} \frac{x_Q}{n} \cdot \frac{(\ell_P - \ell_Q(x + 1_Q - 1_P))^2}{\ell_P} \\ &= -\frac{\lambda}{d} \sum_{P \in \mathcal{P}^+_{\epsilon,\nu}} x_P \ell_P \sum_{Q:\ell^+ \leq (1+\epsilon) L^+_{\mathrm{av}}} \frac{x_Q}{n} \cdot \left(\frac{\ell_P - \ell^+_Q}{\ell_P}\right)^2 \quad .\end{aligned}$$

Using Jensen's inequality (Fact 12) and substituting $\ell_P \geq L^+_{\mathrm{av}}$ yields

$$\begin{aligned}\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] \leq & -\frac{\lambda}{d} L^+_{\mathrm{av}} \sum_{P \in \mathcal{P}^+_{\epsilon,\nu}} x_P \left(\sum_{Q:\ell^+ \leq (1+\epsilon) L^+_{\mathrm{av}}} \frac{x_Q}{n} \cdot \frac{\ell_P - \ell^+_Q}{\ell_P}\right)^2 \\ & \cdot \frac{1}{\sum_{Q:\ell^+_Q \leq (1+\epsilon) L^+_{\mathrm{av}}} \frac{x_Q}{n}} \quad .\end{aligned}$$

Now we substitute $\ell_P \geq (1+\epsilon)\,L^+_{\mathrm{av}}$ and use the fact that the squared expression is monotone in ℓ_P. Furthermore, we substitute the definition of T and C to obtain

$$\begin{aligned}&\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] \\ &\leq -\frac{\lambda}{d} L^+_{\mathrm{av}} \sum_{P \in \mathcal{P}^+_{\epsilon,\nu}} x_P \left(\frac{T\,(1+\epsilon) L^+_{\mathrm{av}} - \sum_{Q:\ell^+ \leq (1+\epsilon) L^+_{\mathrm{av}}} \frac{x_Q\,\ell^+_Q}{n}}{(1+\epsilon) L^+_{\mathrm{av}}}\right)^2 \cdot \frac{1}{T} \\ &\leq -\frac{\lambda}{d} L^+_{\mathrm{av}} \sum_{P \in \mathcal{P}^+_{\epsilon,\nu}} x_P \left(\frac{T\,(1+\epsilon) L^+_{\mathrm{av}} - T\,C}{(1+\epsilon) L^+_{\mathrm{av}}}\right)^2 \cdot \frac{1}{T} \\ &= -\frac{\lambda}{d} L^+_{\mathrm{av}} \cdot \left(\frac{(1+\epsilon) L^+_{\mathrm{av}} - C}{(1+\epsilon) L^+_{\mathrm{av}}}\right)^2 \cdot T \cdot \sum_{P \in \mathcal{P}^+_{\epsilon,\nu}} x_P \quad .\end{aligned}$$

We can now use the tradeoff shown in Equation (7), $C \le L_{\mathrm{av}}^+$, and $\sum_{P \in \mathcal{P}_{\epsilon,\nu}^+} x_P > \delta\, n/2$ to obtain

$$\begin{aligned}\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] &\le -\frac{\lambda}{d} \cdot L_{\mathrm{av}}^+ \cdot \frac{(1+\epsilon)L_{\mathrm{av}}^+ - C}{((1+\epsilon)L_{\mathrm{av}}^+)^2} \cdot \epsilon\, L_{\mathrm{av}}^+ \cdot \sum_{P \in \mathcal{P}_{\epsilon,\nu}^+} x_P \\ &\le -\frac{\lambda}{d} \cdot \epsilon \cdot \frac{\epsilon\, L_{\mathrm{av}}^+}{(1+\epsilon)^2} \cdot \frac{\delta\, n}{2} \\ &\le -\Omega\left(\frac{\epsilon^2 \cdot \delta}{d} \cdot n\, L_{\mathrm{av}}^+\right) \quad .\end{aligned}$$

Since $nL_{\mathrm{av}}^+ \ge \Phi$, we have by Lemma 2

$$\mathbb{E}\left[\Phi(x(t+1))\right] \le \Phi(x(t)) - \frac{1}{2}\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] \le \Phi(x(t))\left(1 - \Omega\left(\frac{\epsilon^2 \cdot \delta}{d}\right)\right) \quad .$$

Case 2: Many agents use cheap paths, i. e., $\sum_{P \in \mathcal{P}_{\epsilon,\nu}^-} x_P \ge \delta\, n/2$. This time, we define the volume T and average latency C of paths which are potential origins of agents migrating towards $\mathcal{P}_{\epsilon,\nu}^-$.

$$T = \sum_{Q : \ell_Q \ge (1-\epsilon)L_{\mathrm{av}}} \frac{x_Q}{n} \quad \text{and} \quad C = \frac{1}{T} \sum_{Q : \ell_Q \ge (1-\epsilon)L_{\mathrm{av}}} \frac{x_Q}{n} \ell_Q \quad .$$

This time,

$$L_{\mathrm{av}} \le T \cdot C + (1-T) \cdot (1-\epsilon)\, L_{\mathrm{av}}$$

implying

$$T \ge \frac{\epsilon\, L_{\mathrm{av}}}{C - (1-\epsilon)\, L_{\mathrm{av}}} \quad . \tag{8}$$

Similarly as in Case 1 we now give a lower bound on the contribution to the absolute value of the virtual potential gain caused by agents with latency at least $(1-\epsilon)L_{\mathrm{av}}$ sampling agents in $\mathcal{P}_{\epsilon,\nu}^-$.

$$\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] \le -\frac{\lambda}{d} \sum_{Q : \ell_Q \ge (1-\epsilon)L_{\mathrm{av}}} x_Q\, \ell_Q \sum_{P \in \mathcal{P}_{\epsilon,\nu}^-} \frac{x_P}{n} \cdot \left(\frac{\ell_Q - \ell_P^+}{\ell_Q}\right)^2 \quad .$$

we rearrange the sum, apply Jensen's inequality (Fact 12) to obtain

$$\begin{aligned}
&\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] \\
&\leq -\frac{\lambda}{d} \sum_{P \in \mathcal{P}^-_{\epsilon,\nu}} x_P \sum_{Q : \ell_Q \geq (1-\epsilon) L_{\mathrm{av}}} \frac{x_Q\, \ell_Q}{n} \cdot \left(\frac{\ell_Q - \ell^+_P}{\ell_Q} \right)^2 \\
&\leq -\frac{\lambda}{d} \sum_{P \in \mathcal{P}^-_{\epsilon,\nu}} x_P \left(\sum_{Q : \ell_Q \geq (1-\epsilon) L_{\mathrm{av}}} \frac{x_Q\, \ell_Q}{n} \cdot \frac{\ell_Q - \ell^+_P}{\ell_Q} \right)^2 \cdot \frac{1}{\sum_{Q : \ell_Q \geq (1-\epsilon) L_{\mathrm{av}}} \frac{x_Q\, \ell_Q}{n}} \\
&= -\frac{\lambda}{d} \sum_{P \in \mathcal{P}^-_{\epsilon,\nu}} x_P \left(\sum_{Q : \ell_Q \geq (1-\epsilon) L_{\mathrm{av}}} \frac{x_Q}{n} \cdot (\ell_Q - \ell^+_P) \right)^2 \cdot \frac{1}{C\,T} \\
&= -\frac{\lambda}{d} \sum_{P \in \mathcal{P}^-_{\epsilon,\nu}} x_P \left(T \cdot (C - \ell^+_P) \right)^2 \cdot \frac{1}{C\,T} \\
&\leq -\frac{\lambda}{d} \left(T \cdot (C - (1-\epsilon)\, L_{\mathrm{av}}) \right)^2 \cdot \frac{1}{C\,T} \cdot \sum_{P \in \mathcal{P}^-_{\epsilon,\nu}} x_P \quad .
\end{aligned}$$

Finally, using Equation (8) and $C\,T \leq L_{\mathrm{av}}$,

$$\begin{aligned}
\mathbb{E}\left[\sum_{P,Q} V_{PQ}\right] &\leq -\frac{\lambda}{d} \left(\epsilon\, L_{\mathrm{av}} \right)^2 \cdot \frac{1}{C\,T} \cdot \sum_{P \in \mathcal{P}^-_{\epsilon,\nu}} x_P \\
&\leq -\frac{\lambda\, \epsilon^2\, L_{\mathrm{av}}}{d} \delta n \\
&\leq -\Omega\left(\frac{\delta\, \epsilon^2\, \Phi}{d} \right) \quad .
\end{aligned}$$

In both cases, the potential decreases by at least a factor of $(1 - \Omega(\epsilon^2\, \delta/d))$ in expectation, which, by Lemma 7, implies that the expected time to reach a state with $\Phi(x(t)) \leq \Phi^*$ is at most the time stated in the theorem. □

From Theorem 3 we can immediately derive the next corollary.

Corollary 4 *Consider a symmetric network congestion game with polynomial latency functions with maximum degree d and minimum and maximum coefficients $a_{\min}$ and $a_{\max}$, respectively. Then the dynamics converge to an (δ,ϵ,ν)-equilibrium in expected time*

$$\mathcal{O}\left(\frac{d^2}{\epsilon^2\, \delta} \cdot \log\left(n\, m \frac{a_{\max}}{a_{\min}} \right) \right) \quad .$$

Let us remark that (δ,ϵ,ν)-equilibria are transient. They can be left again once they are reached, e.g., if the average latency decreases or if agents migrate towards low-latency paths. However, our proofs actually do not only bound the time until a (δ,ϵ,ν)-equilibrium is reached for the first time, but rather the expected total number of rounds in which the system is not at a (δ,ϵ,ν)-equilibrium.

In the definition of (δ,ϵ,ν)-equilibria we require the majority of agents to deviate by no more than a small amount from L_{av}^{+}. This is because the expected latency of a path sampled by an agent is L_{av}, but the latency of the destination path becomes larger if the agent migrates. We use L_{av}^{+} as an upper bound in our proof, although we could use a slightly smaller quantity in cases where the origin Q and the destination P intersect, namely $\ell_P(x + 1_P - 1_Q)$. Using an average over P and Q of this quantity rather than L_{av}^{+} results in a slightly stronger definition of (δ,ϵ,ν)-equilibria. Here we used the weaker definition for the sake of clarity.

Finally we outline fundamental limitations of fast convergence. One could hope to show fast convergence towards a state in which $\delta = 0$ and *all* agents are approximately satisfied. Any protocol that proceeds by sampling either a strategy or an agent and then possibly migrates, takes at least expected time $\Omega(n)$ to reach a state in which all agents sustain a latency that is within a constant factor of L_{av}^{+}. To see this, consider an instance with $n = 2\,m$ agents and identical linear latency functions. Now, let $x_1 = 3$, $x_2 = 1$ and $x_i = 2$ for $3 \le i \le n$. Then, the probability that one of the agents currently using resource 1 samples resource 2 is at most $\mathcal{O}(1/m) = \mathcal{O}(1/n)$. Since this is the only possible improvement step, this yields the desired bound.

4 Singleton Games and the Price of Imitation

4.1 The Lack of Innovation

In this section, we improve on our previous results and consider the special case of singleton congestion games. In a singleton game every strategy of every agent is a single element $e \in E$. A major drawback of the Imitation Protocol is that agents who rely on this protocol cannot access any edges which are unused in the starting state of the dynamics. Even worse, although an edge has been used initially, it can become unused in later states. It is clear, however, that when starting from a random initial distribution of agents among the edges, the probability of emptying an edge becomes negligible *as the number of agents increases.*

Subsequently, we formalize this statement in the following sense. Consider a family of singleton congestion games over the *same* set of edges with latency functions without offsets. If the strategies of agents are initialized uniformly at random, the probability that an edge becomes unused is exponentially small in the number of agents. To this end, consider a vector of continuous latency functions $\mathcal{L} = (\ell_e)_{e\in E}$ with $\ell_e : [0, 1] \to \mathbb{R}_{\ge 0}$. To use these functions for games with a finite number of agents, we have to normalize them appropriately. For any such function $\ell \in \mathcal{L}$, let ℓ^n with $\ell^n(x) = \ell(x/n)$ denotes the respective scaled function. We may think of this as having n agents with weight $1/n$ each. Note that this transformation leaves the elasticity unchanged, whereas the step size ν decreases as n increases. For a vector of latency functions $\mathcal{L} = (\ell_e)_{e\in E}$, let $\mathcal{L}^n = (\ell_e^n)_{e\in E}$.

For the proof of the theorems we need the following definitions. For a singleton game with n agents, a set E of resources, and linear latency functions $\ell_e(x) = a_e\, x$ for every $e \in E$, the *optimal fractional solution* $\tilde{x}$ is the solution $\tilde{x} \in [0, n]^m$ such that $\sum_{e\in E} \tilde{x}_e = n$ and $\tilde{x} = \arg\min SC(x)$, where the *social cost* $SC(x) = \sum_{e\in E}(x_e/n) \cdot \ell_e(x_e)$ is the average latency of the agent's total demand.

Theorem 4 *Fix a vector of latency functions $\mathcal{L}$ with $\ell_e(0) = 0$ for all $e \in E$. For the singleton congestion game over $\mathcal{L}^n$ with n agents, the probability that the* Imitation Protocol *with random initialization generates a state with $x_e = 0$ for some $e \in E$ within* $\mathrm{poly}(n)$ *rounds is bounded by $2^{-\Omega(n)}$.*

Proof Let d denote an upper bound on the elasticity of the functions in $\mathcal{L}$, and let $\mathrm{opt}_{\mathcal{L}} = \min_y\{L_{\mathrm{av}}(y)\}$ where the minimum is taken over all $y \in \{y' \in \mathbb{R}^m_{\geq 0} \mid \sum_e y'_e = 1\}$. In other words, $\mathrm{opt}_{\mathcal{L}}$ corresponds to the minimum average latency achievable in a fractional solution, i.e., to $SC(\tilde{x})$ for the current set of latency functions $\mathcal{L}$. For any $e \in E$, by continuity and monotonicity, there exists an $y_e > 0$ such that $\ell_e(y_e) < \mathrm{opt}_{\mathcal{L}}/4^d$ and $y_e < 1/m$.

Consider the congestion game with n agents and fix an arbitrary edge $e \in E$. In the following, we upper bound the probability that the congestion on edge e falls below $n\, y_e/2$. First, consider the random initialization in which each resource receives an expected number of n/m agents. The probability that $x_e < n\, y_e/2 \leq n/(2\, m)$ is at most $2^{-\Omega(n\, y_e)}$. Now, consider any assignment x with $x_j > n\, y_j/2$ for all $e \in E$. There are two cases.

Case 1: $x_e > y_e\, n$. Since in expectation, our policy removes at most a λ/d fraction of the agents from edge e, the expected load in the subsequent round is at least $(1 - \lambda/d)\, x_e$. Since for sufficiently small λ it holds that $1 - \lambda/d \geq 3/4$, we can apply a Chernoff bound (Fact 7) in order to obtain an upper bound of $2^{-\Omega(x_e)}$ for the probability that the congestion on e decreases to below $x_e/2 \geq y_e\, n/2$.

Case 2: $y_e\, n/2 < x_e \leq y_e\, n$. Hence, $\ell_e^n(x_e) \leq \mathrm{opt}_{\mathcal{L}}/4^d$. In the following, let n^- denote the number of agents on edges r with $\ell_r^n(x_r + 1) < \ell_e^n(x_e)$, and let n^+ denote the number of agents utilizing edges with latency above $\mathrm{opt}_{\mathcal{L}}$. There are two subcases:

Case 2a: $n^- = 0$. Then, the probability that an agent leaves edge e is 0.

Case 2b: $n^- \geq 1$. We first show that $n^+ \geq 4 \max\{n^-, x_e\}$. For the sake of contradiction, assume that $n^+ < 4\, n^-$. Now, consider an assignment where all of these agents are shifted to edges r with latency $\ell_r^n(x_r) < \ell_e^n(x_e) \leq \mathrm{opt}_{\mathcal{L}}/4^d$, where edge r receives $n^+ \cdot x_r/n^-$ (fractional) agents. In this assignment, the congestion on all edges is increased by no more

than a factor of $n^+/n^- < 4$. Hence, due to the limited elasticity, this increases the latency by strictly less than a factor of 4^d. Then, all edges have a latency of less than $\mathrm{opt}_{\mathcal{L}}/4 \cdot 4 = \mathrm{opt}_{\mathcal{L}}$ and some have latency strictly less than $\mathrm{opt}_{\mathcal{L}}$, a contradiction. The same argument also holds if we consider only resource e rather than all resources r considered above. Hence, also $n^+ \geq 4\,x_e$.
Now, consider the number of agents leaving edge e. Clearly,

$$\mathbb{E}\left[\Delta X_e^-\right] \leq x_e \cdot \frac{\lambda}{d} \sum_{r:\ell_r^n(x_r+1)<\ell_e^n(x_e)} \frac{x_r}{n} = x_e \cdot \frac{\lambda\, n^-}{d\, n} \quad .$$

All agents with current latency at least $\mathrm{opt}_{\mathcal{L}}$ can migrate to resource e since the anticipated latency gain is larger than ν. Hence, the number of agents migrating towards e, is at least

$$\begin{aligned}
\mathbb{E}\left[\Delta X_e^+\right] &\geq \sum_{r:\ell_r^n(x_r)\geq \mathrm{opt}_{\mathcal{L}}} x_r \cdot \frac{\lambda\, x_e \cdot (\ell_r^n(x_r) - \ell_e^n(x_e+1))}{n\, d\, \ell_r^n(x_r)} \\
&\geq \frac{\lambda\, x_e}{n\, d} \cdot \sum_{r:\ell_r^n(x_r)\geq \mathrm{opt}_{\mathcal{L}}} x_r \cdot \frac{\ell_r^n(x_r) - 2^d \cdot \ell_e^n(x_e)}{\ell_r^n(x_r)} \\
&\geq \frac{\lambda\, x_e}{n\, d} \cdot (1 - \frac{1}{2^d}) \cdot n^+ \\
&\geq 2 \cdot x_e \cdot \frac{\lambda}{d\, n} \max\{n^-, x_e\} \quad .
\end{aligned}$$

The third inequality holds since $\ell_r^n \geq \mathrm{opt}_{\mathcal{L}}$ and $\ell_e^n \leq \mathrm{opt}_{\mathcal{L}}/4^d$ and the last inequality holds since $d \geq 1$. For any $T \geq 0$ it holds that

$$\begin{aligned}
\mathbb{P}\left[\Delta X_e \geq 0\right] &\geq \mathbb{P}\left[(\Delta X_e^+ \geq T) \wedge (\Delta X_e^- \leq T)\right] \\
&\geq \left(1 - \mathbb{P}\left[\Delta X_e^+ < T\right]\right) \cdot \left(1 - \mathbb{P}\left[\Delta X_e^- > T\right]\right) \quad .
\end{aligned}$$

Due to our lower bounds on $\mathbb{E}\left[\Delta X_e^+\right]$ and $\mathbb{E}\left[\Delta X_e^-\right]$ we can apply a Chernoff bound (Fact 7) on these probabilities. We set

$$T = 1.5\,\lambda\, \max\{x_e, n^-\}\, x_e/(d\, n)$$

which upper bounds $\mathbb{E}\left[\Delta X_e^-\right]$ and lower bounds on $\mathbb{E}\left[\Delta X_e^+\right]$, so

$$\begin{aligned}
&\mathbb{P}\left[\Delta X_e^+ < T\right] \leq 2^{-\Omega(T)} \leq 2^{-\Omega(\lambda\, x_e^2/(d\, n))} \quad \text{and} \\
&\mathbb{P}\left[\Delta X_e^- > T\right] \leq 2^{-\Omega(T)} \leq 2^{-\Omega(\lambda\, x_e^2/(d\, n))} \quad .
\end{aligned}$$

Altogether,

$$\begin{aligned}
\mathbb{P}\left[\Delta X_e \geq 0\right] &\geq \left(1 - 2^{-\Omega\left(\frac{\lambda\, x_e^2}{d\, n}\right)}\right) \cdot \left(1 - 2^{-\Omega\left(\frac{\lambda\, x_e^2}{d\, n}\right)}\right) \\
&= 1 - 2^{-\Omega\left(\frac{\lambda\, x_e^2}{d\, n}\right)} \quad .
\end{aligned}$$

Finally, since $x_e \geq n\, y_e/2$, $\mathbb{P}\left[\Delta X_e < 0\right] \leq 2^{-\Omega(\lambda\, n\, y_e^2/d)} = 2^{-\Omega(x_e)}$.

In all cases, the probability that the edge becomes unused is bounded by $2^{-\Omega(x_e)} = 2^{-\Omega(n)}$. Hence, the same holds also for $m = \text{poly}(n)$ edges and $\text{poly}(n)$ rounds. □

The proof does not only show that edges do not become empty with high probability, but also that the congestion does not fall below any constant congestion value. This is the only place where our analysis relies on the parameter ν. Thus, if the number of agents is large, we can remove ν from the protocol, and the dynamics converge to an exact Nash equilibrium with high probability.

4.2 The Price of Imitation

In the preceding section we have seen that it is unlikely for resources to become unused when the granularity of an agent decreases. If the instance, i. e., the latency functions and the number of users, is fixed, it is an interesting question, how much the performance can suffer from the fact that the Imitation Protocol is not innovative. We measure this degradation of performance by introducing the *Price of Imitation* which is defined as the ratio between the expected social cost of the state to which the Imitation Protocol converges, denoted I_Γ, and the optimum social cost. The expectation is taken over the random choices of the Imitation Protocol. Let us point out that throughout this section we again use the assumption that strategies are initialized uniformly at random for each agent. When we consider expectations over random choices of the Imitation Protocol, random choices during the initialization are naturally included.

We bound the performance degradation for the case of linear latency functions of the form $\ell_e(x) = a_e\, x$. Then, $d = 1$ is an upper bound on the elasticity and $\nu = a_{\max} = \max_{e \in E}\{a_e\}$. Choosing the average latency $SC(x) = \sum_{e \in E}(x_e/n) \cdot \ell_e(x_e)$ as the social cost measure, we show in Theorem 5 that the Price of Imitation is bounded by a constant. It is, however, obvious that the same also holds if we consider the maximum latency as social cost function.

The performance of the dynamics can be artificially degraded by introducing an extremely slow edge. Thus, $a_{\max}$ can be chosen extremely large such that any state is imitation-stable. However, such a resource can be removed from the instance without harming the optimal solution at all since it would not be used anyhow. We will call such resources *useless* and make this notion precise as follows. For a set of resources $M \subseteq E$, let $A_M = \sum_{e \in M} \frac{1}{a_e}$ and let $A_\Gamma = A_E$. Note that the optimal fractional solution $\tilde{x}_e$ can be computed as $\tilde{x}_e = n/(A_\Gamma\, a_e)$. For this solution, the latency of all resources is $a_e \cdot \tilde{x}_e = n/A_\Gamma$. A resource $e \in E$ is *useless* if $\tilde{x}_e < 1$.

Let us give a rough outline of the proof. We do not compare the outcome of the Imitation Protocol to the optimum solution, but rather to a lower bound, namely the *optimal fractional solution*. In particular, we assume that there are no useless resources. Then, we can show that the social cost at an imitation-stable state in which all resources are used, does not differ by more than a small constant from the optimal social cost and that the Price of Imitation is small. In fact, whereas we have $\tilde{x}_e \geq 1$ for the former statement, for bounding the Price of Imitation we need a slightly stronger assumption, namely that $x_e = \Omega(\log n)$.

Theorem 5 *Suppose that in a singleton congestion game with linear latency functions the optimal fractional solution satisfies $\tilde{x}_e = \Omega(\log n)$. Then the Price of Imitation with random initialization is at most $(3 + o(1))$. In particular, for $\delta > 0$, and any $n \geq n_0(\delta)$ for a large enough value $n_0(\delta)$ (which is independent of the instance), $I_\Gamma \leq (3 + \delta) \cdot \frac{n}{A_\Gamma}$.*

We start by proving two lemmas.

Lemma 3 *Let x be a state in which no agent can gain more than $a_{\max}$. Then,*

$$\frac{n}{A_\Gamma} \leq SC(x) \leq 3\frac{n}{A_\Gamma} \quad .$$

Proof The lower bound has been proven above since n/A_Γ is the social cost of an optimal fractional solution. Also note that, since there are no useless resources, $\tilde{x}_e \geq 1$ and hence $n/A_\Gamma \geq a_{\max}$.

For the upper bound, consider a state x in which no agent can gain more than $a_{\max}$. For the sake of contradiction assume that there exists a resource $e \in E$ with $\ell_e(x_e) > 3\,n/A_\Gamma$. Since $x \neq \tilde{x}$ there exists a resource $f \neq e$ with $x_f < \tilde{x}_f$. In particular, $\ell_f(x_f + 1) < n/A_\Gamma + a_{\max} \leq 2\,n/A_\Gamma \leq \ell_e(x_e) - a_{\max}$. The last inequality holds due to our assumption on $\ell_e(x_e)$ and since $n/A_\Gamma \geq a_{\max}$. Hence, any agent on resource e can improve by $a_{\max}$ by migrating to f, a contradiction. □

Lemma 4 *The* Imitation Protocol *converges towards an imitation-stable state in time $\mathcal{O}\left(n^4 \log n\right)$.*

Proof Consider a state $x(t)$ in which there is at least one agent who can make an improvement of $a_{\max}$. Since its current latency is at most $n \cdot a_{\max}$ and the probability to sample the correct resource is at least $1/n$, the probability to do so is at least $\lambda \cdot (1/n) \cdot (a_{\max}/(n\, a_{\max})) = \lambda/n^2$ and the virtual potential gain of such a step is $a_{\max} \geq \Phi/n^2$. Hence, the absolute value of the expected virtual potential gain in state $x(t)$ is at least $\lambda\,\Phi(x(t))/n^4$. Hence, by Lemma 2,

$$\mathbb{E}\left[\Phi(x(t+1))\right] \leq \Phi(x(t)) \cdot \left(1 - \frac{\lambda}{2\,n^4}\right) \quad .$$

Note that $\Phi^* \geq n\, a_{\min}$ and $a_{\max} \leq n\, a_{\min}$ by the assumption that no resource is useless. Also, $\Phi(x(0)) \leq n^2\, a_{\max}$. Now, the lemma follows by an application of Lemma 8 in the Appendix. □

Based upon the proof of Theorem 4 we can now bound the probability that a resource becomes empty for the case of linear latency functions more specifically.

Lemma 5 *The probability that all resources of the subset $M \subseteq E$ become empty in one round simultaneously is bounded from above by*

$$\prod_{e \in M} 2^{-\Omega\left(\frac{n}{A_\Gamma\, a_e}\right)} \ .$$

Proof Recall the bounds on the probability that a resource $e \in E$ becomes empty in the proof of Theorem 4. Since we now consider linear latency functions, we may explicitly compute the value of $y_e = 1/(A_\Gamma\, a_e)$. Recall the two cases and the failure probability in the initialization:

Initialization: Here, the error probability was at most $2^{-\Omega(n\, y_e)} = 2^{-\Omega\left(\frac{n}{A_\Gamma\, a_e}\right)}$.
Case 1: $x_e > y_e\, n$. Here, the error probability was at most $2^{-\Omega(x_e)} = 2^{-\Omega\left(\frac{n}{A_\Gamma\, a_e}\right)}$.
Case 2: $y_e\, n/2 < x_e \leq y_e\, n$. Here, the error probability was at most $2^{-\Omega(x_e^2/n)} = 2^{-\Omega\left(\frac{n}{(A_\Gamma\, a_e)^2}\right)}$.

In all cases, the probability that resource i becomes empty is at most $2^{-\Omega\left(\frac{n}{A_\Gamma\, a_e}\right)}$.

Furthermore, consider resources e and e' and let B and B' denote the events that e and e' become empty, respectively. It holds that, $\mathbb{P}[B' \mid B] \leq \mathbb{P}[B']$. Therefore, $\mathbb{P}[B \cap B'] = \mathbb{P}[B] \cdot \mathbb{P}[B' \mid B] \leq \mathbb{P}[B] \cdot \mathbb{P}[B']$. Extending this argument to several resources yields the statement of the lemma. □

Using the above two lemmas, we can now prove the main theorem of this section.

Proof (Proof of Theorem 5) The proof is by induction on the number of resources m. Clearly, the statement holds for $m = 1$, in which case there is only one assignment. In the following we divide the sequence of states generated by the Imitation Protocol into *phases* consisting of several rounds. The phase is terminated by one of the following events, whatever happens first:

1. A subset of resources M becomes empty.
2. The Imitation Protocol reaches an imitation-stable state.
3. The protocol enters round $\Theta(n^5 \log n)$.

For $M \subseteq E$ let $\Gamma \setminus M$ denote the instance obtained from Γ by removing all resources in M. If a phase ends because Event 1 occurs, we start a new phase for the instance $\Gamma \setminus M$. If it ends because of Event 3, we start a new phase for the original instance.

The probability for Event 1 is bounded by Lemma 5. Note that the probability is also bounded for up to poly(n) many rounds. If a phase ends with Event 2 we have $I_\Gamma \leq 3\,\frac{n}{A_\Gamma}$ (Lemma 3). We bound the probability of this event by 1, which is trivially true. Event 3 happens with a probability at most $\mathcal{O}(1/n)$. This can be shown using Lemma 4 and Markov's inequality. Note that the expected social cost is still at most I_Γ. Summing up over all three events, we obtain the following recurrence:

$$I_\Gamma \leq \sum_{M \subset E} \prod_{e \in M} 2^{-\Omega\left(\frac{n}{A_\Gamma\, a_e}\right)} \cdot I_{\Gamma \setminus M} + 3 \cdot \frac{n}{A_\Gamma} + \mathcal{O}\left(\frac{1}{n}\right) \cdot I_\Gamma$$

implying

$$I_\Gamma \cdot \left(1 - \mathcal{O}\left(\frac{1}{n}\right)\right) \le 3 \cdot \frac{n}{A_\Gamma} + \sum_{M \subset E} \prod_{e \in M} 2^{-\Omega\left(\frac{n}{A_\Gamma a_e}\right)} \cdot I_{\Gamma \setminus M} \ .$$

Substituting the induction hypothesis for $I_{\Gamma\setminus M}$, and introducing a constant c for the constant in the $\Omega()$,

$$\begin{aligned} I_\Gamma \cdot \left(1 - \mathcal{O}\left(\frac{1}{n}\right)\right) &\le 3 \cdot \frac{n}{A_\Gamma} + \sum_{M \subset E} \prod_{e \in M} 2^{-\frac{c\,n}{A_\Gamma a_e}} \cdot 4\,\frac{n}{A_{\Gamma\setminus M}} \\ &= 3 \cdot \frac{n}{A_\Gamma} + 4\,\frac{n}{A_\Gamma} \sum_{M \subset E} 2^{-\frac{c\,n\,A_M}{A_\Gamma}} \cdot \frac{A_\Gamma}{A_{\Gamma\setminus M}} \ . \end{aligned}$$

Now, by our assumption that for all $e \in M$, $\tilde{x}_e = n/(A_\Gamma \cdot a_e) \ge \Omega(\log n)$, we know that for all e, $1/a_e \ge c'\,A_\Gamma \cdot \log n/n$ for a constant c' which we may choose appropriately. In particular, $A_M \ge |M| c'\,A_\Gamma \cdot \log n/n$ and $A_{\Gamma\setminus M} \ge c'\,A_\Gamma \cdot \log n/n$. Altogether,

$$\begin{aligned} I_\Gamma \cdot \left(1 - \mathcal{O}\left(\frac{1}{n}\right)\right) &\le \frac{n}{A_\Gamma}\left(3 + 4 \sum_{M \subset E} 2^{-c\,c'\,|M|\log n} \cdot \frac{n}{c'\,\log n}\right) \\ &= \frac{n}{A_\Gamma}\left(3 + 4 \sum_{k=1}^{m-1} \binom{m}{k} 2^{-c\,c'\,k\log n} \cdot \frac{n}{c'\,\log n}\right) \\ &\le \frac{n}{A_\Gamma}\left(3 + 4 \sum_{k=1}^{m-1} n^k \cdot 2^{-c\,c'\,k\log n} \cdot \frac{n}{c'\,\log n}\right) \\ &\le \frac{n}{A_\Gamma}\left(3 + 4 \sum_{k=1}^{m-1} 2^{-(c\,c'-1)\,k\log n} \cdot \frac{n}{c'\,\log n}\right) \\ &\le \frac{n}{A_\Gamma}\left(3 + 4 \sum_{k=1}^{m-1} \frac{n^{-(c\,c'-1)\,k+1}}{c'\,\log n}\right) \\ &\le (3 + o(1))\,\frac{n}{A_\Gamma} \ , \end{aligned}$$

since the last sum is bounded by $o(n)$. This implies our claim. □

5 Exploring New Strategies

In Section 3 we have seen that in the long run the dynamics resulting from the Imitation Protocol converges to an imitation-stable state in pseudopolynomial time. The Imitation Protocol and the concept of an imitation-stable state have the drawback that the dynamics can stabilize in a very disadvantageous state, e.g. when all agents play the same expensive strategy. This results from the strategy space being restricted to the current strategy choices of the agents. Strategies that might be attractive and offer a large latency gain are "lost" once no agent uses them anymore.

A stronger result would be convergence towards a Nash equilibrium. In the literature on congestion games, several other protocols are discussed. For all of the protocols we are aware of, the probability to migrate from one strategy to another depends in some continuous, non-decreasing fashion on the anticipated latency gain, and it becomes zero for zero gain. Hence, in a setting with arbitrary latency functions which we consider here, there always exist simple instances and states that are not at equilibrium and in which only one improvement step is possible which has an arbitrarily small latency gain. Hence, it takes pseudopolynomially long until an exact Nash equilibrium is reached. Still, even without efficient convergence time it might be desirable to design a protocol which reaches a Nash equilibrium in the long run. There are several ways to achieve this goal. We will discuss three of them here.

Theorem 4 states the following for a particular class of singleton congestion games. With an increasing number of agents it becomes increasingly unlikely that useful strategies are lost. This allows to omit the parameter ν from the protocol. If no strategies are lost for a long period of time, the dynamics will converge to an exact Nash equilibrium. Hence, when the setting corresponds to an instance from this class of congestion games, convergence to a Nash equilibrium can be achieved simply by omitting the parameter ν.

Second, we may add an additional "virtual agent" to every strategy, such that the probability to sample a strategy never becomes zero. This has two implications on our analysis. On the one hand, there is a certain base load on all resources, denoted by x_e^0. We then need to have an upper bound on the elasticity of $\ell_e(x - x_e^0)$ which may be larger than the elasticity of $\ell_e(x)$ itself. Furthermore, we have to add $|\mathcal{P}|$ virtual agents, which leaves the analysis of the time of convergence unchanged only if $n = \Omega(|\mathcal{P}|)$.

As a third alternative, we can add an exploration component to the protocol. With a probability of $1/2$, the agents can sample another path uniformly at random rather than another agent. In this case, however, the elasticity d cannot be used as a damping factor anymore, since the expected increase of congestion may be much larger than the current load. Rather, we have to reduce the migration probability by a factor $\min\left\{1, \frac{|\mathcal{P}|\,\ell_{\min}}{\beta\, n}\right\}$ where β is an upper bound on the maximum slope and $\ell_{\min} = \min_{e\in E} \ell_e(1)$ is the minimum latency of an empty resource.

Lemma 6 *Let x denote a state and let ΔX denote a random migration vector generated by the* Exploration Protocol. *Then,*

$$\mathbb{E}\left[\Delta\Phi(x, \Delta X)\right] \leq \frac{1}{2} \sum_{P,Q\in\mathcal{P}} \mathbb{E}\left[V_{PQ}(x, \Delta X)\right] \quad .$$

Proof We use a similar approach as in Lemma 2. We first separate the potential gain into virtual potential gain and error term. Recall that Lemma 1 states

Protocol 2 EXPLORATION PROTOCOL, repeatedly executed by all agents.

Let P denote the path of the agent in state x.
Sample another path $Q \in \mathcal{P}$ uniformly at random.
if $\ell_P(x) > \ell_Q(x + 1_Q - 1_P)$ **then**
  with probability

$$\mu_{PQ} = \min\left\{1, \lambda \cdot \frac{|\mathcal{P}|\,\ell_{\min}}{\beta\, n} \cdot \frac{\ell_P(x) - \ell_Q(x + 1_Q - 1_P)}{\ell_P(x)}\right\}$$

  migrate from path P to path Q.
end if

the following for every state x and every migration vector Δx

$$\Delta\Phi(x, \Delta x) \le \sum_{P,Q\in\mathcal{P}} V_{PQ}(x, \Delta x) + \sum_{e\in E} F_e(x, \Delta x) \ .$$

Recall that the main difficulty in Lemma 2 was to prove the upper bound on $\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right]$ in Equations (2) and (3). Once this upper bound was established, it was straightforward to argue that the sum of error terms $\sum_{einE} F_e(x, \Delta x)$ is at most half of the absolute value of the virtual potential gain. Thus, the error alters the potential gain by at most a factor of 2 and the lemma was proven.

Here we use exactly the same approach. The only part we adjust is the upper bound on $\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right]$, because for the EXPLORATION PROTOCOL deriving this bound turns out to be extremely simple. First consider the case of an edge $e \in Q\backslash P$. Then due to the linearity of expectation,

$$\begin{aligned}\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right] &\le \beta\,\mathbb{E}\left[\Delta\tilde{X}_e\right]\\ &\le \beta\, n \cdot \lambda \cdot \frac{\ell_{\min}\,|\mathcal{P}|}{\beta\, n} \cdot \frac{1}{|\mathcal{P}|} \cdot \frac{\ell_P - \ell_Q^+}{\ell_P}\\ &\le \lambda \cdot \frac{\ell_e^+}{\ell_Q^+} \cdot (\ell_P - \ell_Q^+) \ ,\end{aligned}$$

where we have substituted the migration probability of the protocol and the fact that there are at most n agents that may sample a path containing e. When λ is chosen small enough, this proves Equation (2), i.e.,

$$\mathbb{E}\left[\Delta\tilde{\ell}_e\left(\Delta\tilde{X}_e\right)\right] \le \frac{1}{8} \cdot (\ell_P - \ell_Q^+) \cdot \left(\frac{\ell_e^+}{\ell_Q^+} + \frac{\nu_e}{\nu_Q}\right) \ .$$

A similar argument can be used to prove the statement of Equation (3) for the EXPLORATION PROTOCOL in the case $e \in P$. □

Note that we have omitted the parameter ν from the protocol. Thus, in principle, agents can make arbitrarily small improvements. However, in order to give an upper bound on the convergence time, we need some lower bound on the minimum improvement that is possible when the system is not yet at an imitation-stable state. Formally, let

$$\kappa = \min_x \min_{\substack{P,Q \in \mathcal{P} \\ \ell_p(x) > \ell_Q(x+1_Q-1_P)}} \{\ell_P(x) - \ell_Q(x+1_Q-1_P)\} \ .$$

Theorem 6 *Consider a symmetric network congestion game in which all agents use the* Exploration Protocol*. Let x denote the initial state of the dynamics. Then the dynamics converge to a Nash equilibrium in expected time*

$$\mathcal{O}\left(\frac{\Phi(x)\,\beta\,n\,\ell_{\max}}{\ell_{\min}\,\kappa^2}\right) \ .$$

Proof In every state which is not a Nash equilibrium there exists an agent currently utilizing path $P \in \mathcal{P}$ and a path $Q \in \mathcal{P}$ such that $\ell_Q \leq \ell_P - \kappa$. Hence, the (absolute value of the) expected virtual potential gain is

$$\mathbb{E}\left[V_{PQ}\right] \leq -\frac{1}{|\mathcal{P}|} \cdot \frac{\lambda\,|\mathcal{P}|\,\ell_{\min}}{\beta\,n} \cdot \frac{\kappa}{\ell_P} \cdot \kappa \leq -\frac{\lambda\,\ell_{\min}}{\beta\,n} \cdot \frac{\kappa^2}{\ell_{\max}} \ ,$$

and the true potential gain differs from this only by a factor of at most 1/2. Again, Lemma 7 yields the expected time until the potential decreases from at most Φ to $\Phi^* \geq 0$ and proves the theorem. □

It is obvious that an analogue of Lemmata 2 and 6 also holds for any protocol that is a combination of the Imitation Protocol and the Exploration Protocol, e. g., a protocol in which in every round every agent executes the one or the other with probability one half. Then, in order to bound the value of $\mathbb{E}\left[\Delta\tilde{\ell}_e(\Delta\tilde{X}_e)\right]$, we must make a case differentiation based on whether proportional or uniform sampling dominates the probability that other agents migrate towards resource e. Such a protocol combines the advantages of the Imitation Protocol and the Exploration Protocol: In the long run, it converges to a Nash equilibrium, and it reaches an approximate equilibrium as quickly as stated by Theorem 3 (up to a factor of 2).

6 Conclusion

We have proposed and analyzed a natural protocol based on imitating profitable strategies for distributed selfish agents in symmetric congestion games. If agents use our Imitation Protocol, the resulting dynamics converge rapidly to approximate equilibria, in which only a small fraction of agents have latency significantly above or below the average. In addition, in finite time the dynamics converge to an imitation-stable state, in which no agent can improve its latency by more than ν by imitating a different agent. The Imitation Protocol and the concept of an imitation-stable state have the drawback that dynamics can stabilize in a quite disadvantegous situation, e.g. when all agents play the same expensive strategy. This is due to the fact that the strategy space is essentially restricted to the current strategy choices of the agents. Strategies that might be attractive and offer large latency gain are "lost" once no agent uses them anymore. For singleton congestion games we showed that this event becomes unlikely to occur as the number of agents increases. Then, by removing parameter ν from the protocol, the dynamics become likely to converge to Nash equilibria. Another approach to avoid losing strategies is to include exploration of the strategy space. To this end, we can use an Exploration Protocol, in which agents sample from the strategy space directly and then migrate with a certain probability. If every agent uses a suitably designed Exploration Protocol (or any random combination of Exploration Protocol and Imitation Protocol), then the dynamics are always guaranteed to converge to a Nash equilibrium, and they still reach approximate equilibria rapidly.

References

1. Heiner Ackermann, Simon Fischer, Martin Hoefer, and Marcel Schöngens. Distributed algorithms for QoS load balancing. *Distributed Computing*, 23(5–6):321–330, 2011.
2. Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6), 2008.
3. Clemens Adolphs and Petra Berenbrink. Distributed selfish load balancing with weights and speeds. In *Proc. 31st Symp. Principles of Distrib. Comput. (PODC)*, pages 135–144, 2012.
4. Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. *SIAM J. Comput.*, 42(1):160–177, 2013.
5. Baruch Awerbuch, Yossi Azar, Amir Epstein, Vahab Mirrokni, and Alexander Skopalik. Fast convergence to nearly optimal solutions in potential games. In *Proc. 9th Conf. Electronic Commerce (EC)*, pages 264–273, 2008.
6. Petra Berenbrink, Tom Friedetzky, Leslie Ann Goldberg, Paul Goldberg, Zengjian Hu, and Russel Martin. Distributed selfish load balancing. *SIAM J. Comput.*, 37(4):1163–1181, 2007.
7. Petra Berenbrink, Tom Friedetzky, Iman Hajirasouliha, and Zengjian Hu. Convergence to equilibria in distributed, selfish reallocation processes with weighted tasks. *Algorithmica*, 62(3–4):767–786, 2012.
8. Petra Berenbrink, Martin Hoefer, and Thomas Sauerwald. Distributed selfish load balancing on networks. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1487–1497, 2011.
9. Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. *Theory of Computing*, 6(1):179–199, 2010.
10. Steve Chien and Alistair Sinclair. Convergence to approximate Nash equilibria in congestion games. *Games Econom. Behav.*, 71(2):315–327, 2011.
11. George Christodoulou and Elias Koutsoupias. The price of anarchy in finite congestion games. In *Proc. 37th Symp. Theory of Computing (STOC)*, pages 67–73, 2005.
12. Eyal Even-Dar, Alexander Kesselman, and Yishay Mansour. Convergence time to Nash equilibria in load balancing. *ACM Trans. Algorithms*, 3(3), 2007.
13. Eyal Even-Dar and Yishay Mansour. Fast convergence of selfish rerouting. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 772–781, 2005.
14. Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proc. 36th Symp. Theory of Computing (STOC)*, pages 604–612, 2004.

15. Angelo Fanelli, Michele Flammini, and Luca Moscardelli. The speed of convergence in congestion games under best-response dynamics. *ACM Trans. Algorithms*, 8(3):25, 2012.
16. Angelo Fanelli and Luca Moscardelli. On best response dynamics in weighted congestion games with polynomial delays. *Distributed Computing*, 24(5):245–254, 2011.
17. Angelo Fanelli, Luca Moscardelli, and Alexander Skopalik. On the impact of fair best response dynamics. In *Proc. 37th Intl. Symp. Math. Foundations of Computer Science (MFCS)*, pages 360–371, 2012.
18. Simon Fischer, Petri Mähönen, Marcel Schöngens, and Berthold Vöcking. Load balancing for dynamic spectrum assignment with local information for secondary users. In *Proc. Symp. Dynamic Spectrum Access Networks (DySPAN)*, 2008.
19. Simon Fischer, Harald Räcke, and Berthold Vöcking. Fast convergence to Wardrop equilibria by adaptive sampling mehtods. *SIAM J. Comput.*, 39(8):3700–3735, 2010.
20. Simon Fischer and Berthold Vöcking. Adaptive routing with stale information. *Theoret. Comput. Sci.*, 410(36):3357–3371, 2008.
21. Dimitris Fotakis, Alexis Kaporis, and Paul Spirakis. Atomic congestion games: Fast, myopic and concurrent. *Theory Comput. Syst.*, 47(1):38–49, 2010.
22. Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. Atomic congestion games among coalitions. *ACM Trans. Algorithms*, 4(4), 2008.
23. Paul Goldberg. Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game. In *Proc. 23rd Symp. Principles of Distrib. Comput. (PODC)*, pages 131–140, 2004.
24. Torben Hagerup and Christine Rüb. A guided tour of Chernoff bounds. *Inf. Process. Lett.*, 33:305–308, 1990.
25. Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
26. Samuel Ieong, Robert McGrew, Eugene Nudelman, Yoav Shoham, and Qixiang Sun. Fast and compact: A simple class of congestion games. In *Proc. 20th Conf. Artificial Intelligence (AAAI)*, pages 489–494, 2005.
27. Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proc. 41st Symp. Theory of Computing (STOC)*, pages 533–542, 2009.
28. Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24(1):21–29, 2011.
29. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
30. Robert Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory*, 2:65–67, 1973.
31. Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. 41st Symp. Theory of Computing (STOC)*, pages 513–522, 2009.
32. Alexander Skopalik and Berthold Vöcking. Inapproximability of pure Nash equilibria. In *Proc. 40th Symp. Theory of Computing (STOC)*, pages 355–364, 2008.
33. Berthold Vöcking. Selfish load balancing. In Noam Nisan, Éva Tardos, Tim Roughgarden, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 20. Cambridge University Press, 2007.
34. Jörgen Weibull. *Evolutionary Game Theory*. MIT Press, 1995.

A Appendix

Throughout the technical part of this paper, we apply the following two Chernoff bounds.

Fact 7 (Chernoff, see [24])**.** *Let X be a sum of Bernoulli variables. Then,* $\mathbb{P}\left[X \geq k \cdot \mathbb{E}\left[X\right]\right] \leq \mathrm{e}^{-\mathbb{E}[X]\,k\cdot(\ln k-1)}$*, and, for* $k \geq 4 > \mathrm{e}^{4/3}$*,* $\mathbb{P}\left[X \geq k \cdot \mathbb{E}\left[X\right]\right] \leq \mathrm{e}^{-\frac{1}{4}\,\mathbb{E}[X]\,k\,\ln k}$*. Equivalently, for* $k \geq 4\,\mathbb{E}\left[X\right]$*,* $\mathbb{P}\left[X \geq k\right] \leq \mathrm{e}^{-\frac{1}{4}\,k\,\ln(k/\mathbb{E}[X])}$*.*

The following fact yields a linear approximation of the exponential function.

Fact 8. *For any $r > 0$ and $x \in [0, r]$, it holds that* $(\mathrm{e}^x - 1) \leq x \cdot \frac{\mathrm{e}^r-1}{r}$*.*

Proof The function $\exp(x) - 1$ is convex and it goes through the points $(0, 0)$ and $(r, \mathrm{e}^r - 1)$, as does the function $x \cdot \frac{\mathrm{e}^r-1}{r}$. □

Fact 9. *It holds that*

$$\sum_{k=1}^{\infty} \mathrm{e}^{-k(\ln k)} \cdot k < 2 \ .$$

Proof We have

$$\sum_{k=1}^{\infty} \mathrm{e}^{-k(\ln k)} \cdot k = \sum_{k=1}^{\infty} \frac{1}{k^{k-1}} \quad = \quad 1 + \sum_{k=2}^{\infty} \frac{1}{k^{k-1}} \quad < \quad 1 + \sum_{k=1}^{\infty} \frac{1}{2^k} \quad \leq \quad 2 \ .$$

□

Fact 10. *It holds that*

$$\sum_{k=2}^{\infty} \mathrm{e}^{-k(\ln(k)-1)} \cdot k < 8 \ .$$

Proof We have

$$\sum_{k=2}^{\infty} \mathrm{e}^{-k(\ln(k)-1)} \cdot k = \sum_{k=1}^{\infty} \mathrm{e} \cdot \left(\frac{\mathrm{e}}{k+1}\right)^k \quad = \quad \sum_{k=1}^{4} \mathrm{e} \cdot \left(\frac{\mathrm{e}}{k+1}\right)^k + \sum_{k=5}^{\infty} \mathrm{e} \cdot \left(\frac{\mathrm{e}}{k+1}\right)^k$$
$$< 7.1 + \mathrm{e} \cdot \sum_{k=5}^{\infty} \frac{1}{2^k} \quad < \quad 8 \ .$$

□

Fact 11. *For every $c \in]0, 1[$ it holds*

$$\sum_{k=0}^{\infty} c^k = \frac{c}{1-c}$$
$$\sum_{k=l}^{\infty} c^k = \frac{c^l}{1-c}$$

Fact 12 (Jensen's Inequality)**.** *Let $f\colon \mathbb{R} \to \mathbb{R}$ be a convex function, and let $a_1, \ldots, a_k, x_1, \ldots, x_k \in \mathbb{R}$. Then*

$$f\left(\frac{\sum_{i=1}^k a_i x_i}{\sum_{i=1}^k a_i}\right) \le \frac{\sum_{i=1}^k a_i f(x_i)}{\sum_{i=1}^k a_i} \quad .$$

If $f(x) = x^2$, then

$$\left(\frac{\sum_{i=1}^k a_i x_i}{\sum_{i=1}^k a_i}\right)^2 \le \frac{\sum_{i=1}^k a_i (x_i)^2}{\sum_{i=1}^k a_i}$$

$$\Leftrightarrow \frac{1}{\sum_{i=1}^k a_i} \cdot \left(\sum_{i=1}^k a_i x_i\right)^2 \le \sum_{i=1}^k a_i f(x_i) \quad .$$

Lemma 7 *Let $X_0, X_1, \ldots$ denote a sequence of non-negative random variables and assume that for all $i \geq 0$*

$$\mathbb{E}\left[X_i \mid X_{i-1} = x_{i-1}\right] \le x_{i-1} - 1$$

and let τ denote the first time t such that $X_t = 0$. Then,

$$\mathbb{E}\left[\tau \mid X_0 = x_0\right] \le x_0 \quad .$$

The proof follows, e.g., from standard martingale arguments in combination with the optional stopping theorem and is omitted here.

Lemma 8 *Let $X_0, X_1, \ldots$ denote a sequence of non-negative random variables and assume that for all $i \geq 0$ $\mathbb{E}\left[X_i \mid X_{i-1} = x_{i-1}\right] \le x_{i-1} \cdot \alpha$ for some constant $\alpha \in (0,1)$. Furthermore, fix some constant $x^* \in (0, x_0]$ and let τ be the random variable that describes the smallest t such that $X_t \le x^*$. Then,*

$$\mathbb{E}\left[\tau \mid X_0 = x_0\right] \le \frac{4}{1-\alpha} \cdot \ln\left(\frac{2x_0}{x^*}\right) \quad .$$

Proof Let us define $\gamma = \frac{1}{1-\alpha}$ and an auxiliary random variable Y^t by $Y^0 := X^0$, and for any round $t \geq 1$,

$$Y^t = \begin{cases} X^t & \text{if } X^t > x^* \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any $t \geq 1$, it follows

$$\mathbb{E}\left[Y^t \mid X^{t-1} = x\right] \le \alpha x.$$

We have for $\kappa = \gamma \cdot (\ln(x^0) - \ln(x^*/2))$,

$$\begin{aligned} \mathbb{E}\left[Y^t\right] &= \sum_x \mathbb{E}\left[Y^t \mid X^{t-1} = x\right] \cdot \mathbb{P}\left[X^{t-1} = x\right] \\ &\le \sum_x \alpha \cdot Y^{t-1} \cdot \mathbb{P}\left[X^{t-1} = x\right] \\ &\le \alpha^\tau \cdot Y^0 \le x^*/2. \end{aligned}$$

Hence by Markov's inequality,

$$\mathbb{P}\left[Y^\kappa \geq x^*\right] \le \frac{1}{2}. \tag{9}$$

We consider two cases.

Case 1: For all time steps $t \in [0, \ldots, \kappa]$, $Y^t = X^t$. Then, as seen above $X^\kappa \leq x^*$ with probability at least 1/2.

Case 2: There exists a step $t \in [1, \ldots, \kappa]$ such that $Y^t \neq X^t$. Let t be the smallest time step with that property. Hence, $Y^t \neq X^t$, but $Y^{t-1} = X^{t-1}$. If $Y^{t-1} = 0$, then $X^{t-1} = 0$. If $Y^{t-1} \neq 0$, then by definition of Y^t,

$$\left(Y^t \neq X^t\right) \bigwedge \left(Y^{t-1} \neq 0\right) \Rightarrow X^t \leq x^*.$$

In all cases we have shown that with probability at least 1/2, there exists a step $t \in [0, \kappa]$ so that $X^t \leq x^*$. If such a step does not exist, we simply repeat the analysis and consider the next κ steps. The probability that we do not observe a step as desired decreases exponentially in the number of restarts. In expectation, we need only $\sum_{k=1}^{\infty} k/2^{k-1} = 4$ phases of κ steps to observe a step as desired. Thus, the expected number of steps is at most $\tau = 4\kappa$. This completes the proof of the lemma. $\square$