Distributed Selfish Load Balancing on Networks *

Petra Berenbrink[†] Martin Hoefer[‡] Thomas Sauerwald[§]

Abstract

We study distributed load balancing in networks with selfish agents. In the simplest model considered here, there are n identical machines represented by vertices in a network and $m \gg n$ selfish agents that unilaterally decide to move from one vetex to another if this improves their experienced load. We present several protocols for concurrent migration that satisfy desirable properties such as being based only on local information and computation and the absence of global coordination or cooperation of agents. Our main contribution is to show rapid convergence of the resulting migration process to states that satisfy different stability or balance criteria. In particular, the convergence time to a Nash equilibrium is only logarithmic in m and polynomial in n, where the polynomial depends on the graph structure. In addition, we show reduced convergence times to approximate Nash equilibria. Finally, we extend our results to networks of machines with different speeds or to agents that have different weights and show similar results for convergence to approximate and exact Nash equilibria.

1 Introduction

Load balancing is an essential requirement in large networks to ensure efficient utilization of resources and satisfactory performance of the system. In many large computer networks load balancing becomes a challenge because of the absence of global information and coordination. When there is only local information available about the load situation and even the existence of machines, a centralized approach to load balancing is inappropriate or even impossible. Instead, one then needs to develop protocols that respect the informational and computational restrictions of the scenario. In addition, the protocols should guarantee rapid convergence to balanced states.

Some distributed algorithmic approaches for load balancing have been proposed in algorithmic game theory, see, e.g., [5, 15, 2, 18]. In this context tasks are considered as selfish agents that act unilaterally and migrate concurrently between machines without global coordination. Such an approach has two main advantages over protocols that use more centralized optimization. Firstly, concurrent migration and unilateral decision making of multiple agents controlling the tasks reduce the coordination overhead and may still allow for rapid convergence (i.e., sublinear in the number of agents). Secondly, such "task agents" have an incentive to follow the protocol. This is an advantage in modern computer networks that are influenced by a variety of economic incentives and developments. In these networks centralized coordination is often absent and user actions are made in a selfish manner. While these properties make protocols for concurrent selfish load balancing desirable, their existence and convergence properties are not well-understood in many load balancing contexts.

*A preliminary version appeared at the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA'11) [8].

†Dept. of Computer Science, Simon Fraser University, Canada. Email: petra@cs.sfu.edu.

‡Supported by DFG grant Ho 3831/3-1. Dept. of Computer Science, RWTH Aachen University, Germany. Email: mhoefer@cs.rwth-aachen.de.

§Max-Planck Institute for Computer Science, Saarbrücken, Germany. Email: sauerwal@mpi-inf.mpg.de.

In this paper, we present protocols for selfish load balancing in a discrete network balancing model. There are n identical machines which represent vertices in an arbitrary, undirected graph $G = (V, E)$, and m tasks that are initially assigned arbitrarily to the machines. Our protocols proceed in a round-based fashion. In each round, every task picks a neighboring machine at random and decides probabilistically whether or not to migrate to that machine. Hence, the tasks only need local information, i.e., they have to know the load of machine they are currently assigned to and the load of the neighboring machines in the graph. The main challenge in the design of concurrent protocols such as ours is to carefully choose appropriate migration probabilities in order to guarantee rapid convergence while avoiding oscillation effects. Our scenario represents a significant extension of the existing literature on selfish load balancing, as concurrent protocols have been considered essentially only for complete graphs [4, 17, 2, 15, 5, 6]

In our model there are several concepts of a state in which the assignment of tasks is "stable" or "balanced". The standard notion of stability is the *Nash equilibrium (NE)*, in which no player can unilaterally improve by moving to a neighboring machine. Note that throughout this paper we restrict to *pure NE* and states that do not involve randomization. A weaker and more rapidly achievable condition is that of *ε-approximate Nash equilibrium (ε-apx. NE)*, where no player can decrease his personal load by less than a factor of $(1 - \varepsilon)$. However, in an NE the load difference between the least loaded and most loaded machines can still be in the order of the network diameter in our model. In this paper, we consider protocols and study the convergence times to approximate and exact NE. We extend our results to networks of machines with different speeds or to agents that have different weights and show similar results for convergence to approximate and exact NE.

1.1 Contribution and Techniques

We propose and analyze concurrent probabilistic protocols to obtain approximate and exact NE for identical machines or machines with speeds, and for weighted tasks on identical machines.

Identical machines. We present a protocol that reaches an exact NE after a number of rounds that depends only logarithmically on m. Let Δ be the maximum degree of the graph and μ_2 be the second smallest eigenvalue of the Laplace matrix of G. The dynamics reach a NE in expected time $\mathcal{O}\left(\Delta/\mu_2 \cdot (\ln m + \ln n) + |E| \cdot \Delta/\mu_2\right)$. For $m \geqslant \delta n^4$ and $\delta > 1$, the first part of this convergence time is the time needed to reach a $2/(1+\delta)$-apx. NE. The second part is the time needed so that no agent wants to unilaterally deviate to a neighboring machine. Therefore, the convergence time is only logarithmic in the number of agents m, but the dependency on n is polynomial and connected to the structure of the graph. Obviously, in general a polynomial dependence on n cannot be avoided (e.g., for paths).

Machines with speeds. In this case, each machine $i \in V$ has a speed $s_i \in \mathbb{N}$, $s_i \geqslant 1$, and thus processes tasks at a different rate. Let $S := \sum_{i \in V} s_i$ and $\delta > 1$, then using our protocol a set of $m \geqslant \delta \cdot 8 \cdot n^3 \cdot S$ agents converges to an $2/(1+\delta)$-apx. NE in expected time $\mathcal{O}(\ln(m) \cdot \text{poly}(n) \cdot \text{poly}(s_{\max}))$ on any graph, where $s_{\max}$ is the maximum speed. An exact NE is reached after additional time $O(\text{poly}(n) \cdot \text{poly}(s_{\max}))$, so in total also in time $\mathcal{O}(\ln(m) \cdot \text{poly}(n) \cdot \text{poly}(s_{\max}))$.

Weighted tasks. Finally, we also consider the case that each task ℓ has a weight $w_\ell \in \mathbb{N}$, $w_\ell \geqslant 1$. In this case, the expected convergence time to a NE is $\mathcal{O}(\Delta \cdot W^3 \cdot w_{\max})$, where W is the sum of all weights and $w_{\max}$ is the maximum weight.

Outline of the paper. Due to technical reasons we present our results in a different order than the one stated above. After presenting the necessary definitions and preliminaries in Section 2, we first derive in Section 3 the results on load balancing with speeds (Theorems 3.1 and 3.2), as this introduces the main technical framework. The results on networks with identical machines are discussed in Section 4. Then we enhance the general approach for machines with speeds to show convergence to apx. NE (Theorems 4.1 and 4.2). Finally, Section 5 contains the extension to weighted tasks (Theorem 5.1) and is mainly based on our previous approach.

1.2 Related Work

Most closely related to our paper is [5], where the case of identical machines in a complete graph is studied. There, a protocol that is equivalent to ours is shown to arrive at a NE in time $\mathcal{O}(\log\log m + \text{poly}(n))$. Note that for complete graphs the NE and optimal allocations are identical. An extension of this model to weighted tasks is studied in [6]. There the authors consider a similar protocol for the complete graph. They show that the expected time to reach an approximate NE is $O(mn \cdot w_{\max}^3)$. Here we extend the results of both papers significantly by studying dynamics on general graphs and machines that have speeds. This also requires to use different techniques that allow to capture the connections between convergence time and graph structure.

Our paper relates to a general stream of works for selfish load balancing on complete graphs. There is a variety of issues that have been considered, starting with seminal papers on algorithms and dynamics to reach NE [14, 16]. More directly related are concurrent protocols for selfish load balancing in different contexts that allow convergence results similar to ours. Whereas some papers consider protocols that use some form of global information [15] or coordinated migration [20], others consider infinitesimal or splittable tasks [19, 4] or work without rationality assumptions [17, 2]. The machine models in these cases range from identical, uniformly related (linear with speeds) to unrelated machines. The latter also contains the case when there are access restrictions of certain agents to certain machines. In contrast, in our model players migrate over a network and can access machines only *depending on their current location*. This is a fundamental difference to all previous related work in this area. For an overview of work on selfish load balancing see, e.g., [31].

A slightly different approach for discrete selfish load balancing on networks are finite congestion games [29], for which the convergence times of sequential best-response dynamics to exact and approximate NE have been extensively studied [10, 3, 30]. A general approach for a concurrent better-response protocol is [1], which is inspired by similar results for non-atomic congestion games [18]. In this protocol, agents pick strategies only by imitation of other agents, and there is rapid convergence even for general delay functions, but only to an approximate equilibrium concept where all agents experience a similar cost. While this represents a generally applicable approach, the obtained approximate equilibrium might be far from any (apx.) NE. A different line of research are no-regret and similar payoff-based learning dynamics [23, 24, 25, 26], but they usually converge (quickly) only in the history of play and/or to classes of mixed NE. In contrast, we present protocols that reach pure exact and approximate NE rapidly.

Our protocol is also related to a vast amount of literature on (non-selfish) load balancing over networks, where results usually concern the case of identical machines and unweighted tasks. Often there are additional restrictions on the graph structure such as regular graphs, expander graphs, tori, etc. A central measure of balance is the discrepancy, i.e., the difference between most and least loaded machine in the network. In expectation, our protocols mimic continuous diffusion, which has been studied initially in [12, 9] and later, e.g., in [27]. This work established the connection between convergence, discrepancy, and eigenvalues of graph matrices. Closer to our paper are discrete diffusion processes – prominently studied in [28], where the authors introduce a general technique to bound the load deviations between an idealized and the actual processes. Recently, randomized extensions of the algorithm in [28] have been considered, e.g., [13, 21]. However, either machines have to communicate with their neighbors to determine the number of tasks that should move [28, 21], or the tasks perform independent random walks [13]. In the first case, machines have a strong control over their tasks, while in the second case tasks may jump from an underloaded to an overloaded machine, which clearly is undesirable in a game-theoretic context.

2 Notation and Preliminaries

We consider an arbitrary, undirected and connected *graph* $G = (V, E)$ with $n = |V|$ vertices representing machines. The *degree* of a vertex $i \in V$ is $\deg(i)$. Δ denotes the *maximum degree* of any vertex in V. For two vertices i, j, $\deg(i, j) = \max\{\deg(i), \deg(j)\}$ is the maximum degree of i and j.

There are m *tasks* in the system, which are initially assigned arbitrarily to the n machines. We denote by x a *state* of the system, i.e., a fixed assignment of tasks to the machines. For any machine $i \in V$, we denote by x_i the set of tasks are that assigned to machine i in x. We consider a probabilistic migration process, in which the state is a random variable in each time step. In particular, let X^t be the state at (the end of) step t. Similarly, X_i^t is the subset of tasks assigned to machine $i \in V$, and it is a random variable for every $t \geqslant 1$ due to the probabilistic nature of our migration protocols. Each task ℓ has a weight $w_\ell \in \mathbb{N}$ and, unless stated otherwise, we assume *uniform tasks* with $w_\ell = 1$.

Each vertex $i \in V$ is a *machine* with a *speed* $s_i \in \mathbb{N}$, $s_i \geqslant 1$. We define $S = \sum_{i=1}^n s_i$. Note that we can also handle rational speeds by normalization to integers. By $s_{\max}$ and $s_{\min}$ we denote the *maximal and minimal speed* of a machine, respectively. If $s_{\max} = s_{\min}$, we have a network of *identical* machines and assume w.l.o.g. $s_i = 1$ for all $i \in V$.

In the case of uniform tasks the *load* of a machine is defined as the sum of tasks assigned to it, divided by the speed of the machine. In the case of weighted tasks the load is the sum of the weights of these jobs divided by the speed. In particular, by $W(x_i)$ we denote the *weight* on machine i in state x, i.e., the sum of weights of all tasks that are located on i. Similarly, $W(X_i^t)$ is the weight at the end of step t. Let $W := \sum_{i \in V} W(x_i)$. For a state x the load of vertex i is denoted by $L(x_i)$ and equals $L(x_i) = W(x_i)/s_i$. Each task on machine i experiences a disutility of $L(x_i)$ in state x. Naturally, for our process at time t, we obtain the random variable $L(X_i^t)$. A task is a *selfish agent* that strives to minimize the experienced load. The task is only aware of the load of the machine it is currently located at and is able to inspect the load on one of the neighboring machines. We consider a round-based process. In each iteration every task uses a protocol to decide to which of the neighbouring machines it potentially migrates.

The Laplace $\mathcal{L}$ matrix is a standard matrix associated to undirected graphs, which is based on adjacency and degree information (e.g., [11]). Formally, $\mathcal{L}$ is defined as the $n \times n$-matrix where $\mathcal{L}_{i,j}$ is equal to $\deg(i)$ if $j = i$, -1 if $\{i, j\} \in E(G)$ and 0 otherwise. The eigenvalues of the Laplace matrix are known to encode valuable structural information for dynamic load balancing processes, see, e.g. [22, 7].

A state x is an *ε-approximate Nash equilibrium (ε-apx. NE)* for any $0 \leqslant \varepsilon \leqslant 1$ if no task can decrease the experience load by more than factor of $(1 - \varepsilon)$. In such a state we have for every machine i and every neighboring machine j

$$(1 - \varepsilon) \cdot \frac{W(x_i)}{s_i} \leqslant \frac{W(x_j) + 1}{s_j} \quad .$$

For $\varepsilon = 0$ we call such a state an *(exact) Nash equilibrium (NE).*

Our game can be expressed as an atomic congestion game, and thus the following function due to Rosenthal [29] is a potential function for our game with uniform tasks:

$$\Phi_1(x) = \sum_{i \in V} \sum_{k=1}^{W(x_i)} \frac{k}{s_i} = \sum_{i \in V} \frac{W(x_i) \cdot (W(x_i)+1)}{s_i} \tag{2.1}$$

Whenever a single player makes a unilateral assignment change, the change in the potential function equals the load change experienced by the player. Thus, the local optima of the potential function are exactly the NE of the game, and the potential function measures the progress to NE from a players point of view. We will also use the quadratic potential function that is standard in the load balancing literature.

$$\Phi_0(x) = \sum_{i \in V} \frac{(W(x_i))^2}{s_i}$$

It measures progress to a completely balanced state from a global point of view. This function will be helpful when we prove convergence to apx. NE. This leads to the following general definition.

Definition 2.1. *For $r \in \{0,1\}$, define*

$$\Phi_r(x) := \sum_{i \in V} \frac{W(x_i) \cdot (W(x_i)+r)}{s_i}.$$

For our migration process, we often consider the change of the potential, which is

$$\Delta\Phi_r(X^t) := \Phi_r(X^{t-1}) - \Phi_r(X^t) \ .$$

We will also use a normalized formulation of $\Phi_0(x)$ denoted by $\Psi_0(x)$.

Definition 2.2. *The function $\Psi_0(x)$ is defined as*

$$\Psi_0(x) = \Phi_0(x) - \frac{m^2}{S} \ .$$

The potential change of Ψ_0 is defined as

$$\Delta\Psi_0(X^t) := \Psi_0(X^{t-1}) - \Psi_0(X^t) = \Delta\Phi_0(X^t).$$

Note that for identical machines it holds

$$\Psi_0(x) = \Phi_0(x) - 2m \cdot \frac{m}{n} + n \cdot \frac{m^2}{n^2} = \sum_{i \in V} \left(W(x_i) - \frac{m}{n}\right)^2 \ .$$

As a standard convention, the term "with high probability" means with probability at least $1 - n^{-c}$ for some constant $c > 0$. Our main results are upper bounds on the expected time for our protocols to reach an (apx.) NE. We point out that one can get corresponding upper bounds which hold with high probability at the cost of a multiplicative increase by $\mathcal{O}(\log n)$.

---

for each task ℓ in parallel **do**
  Let $i = i(\ell)$ the current machine of task ℓ
  Choose a neighboring machine j u.a.r.
  if $L(X_i^{t-1}) - L(X_j^{t-1}) > 1/s_j$ **then**
    Move task ℓ from resource i to j with probability

$$\frac{\deg(i)}{\deg(i,j)} \cdot \frac{L(X_i^{t-1}) - L(X_j^{t-1})}{\alpha \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot W(X_i^{t-1})}$$

  end if
end for

Figure 1: Protocol I for uniform tasks and machines with speeds. We set $\alpha := 4s_{\max}$.

3 Uniform Tasks and Related Machines

In this section, we first present our results on general networks and machines with speeds, as our analysis of this case introduces our main approach. Protocol I in Figure 1 allows tasks to move to neighboring machines with a smaller load. In more detail, in each round every player randomly chooses a neighboring machine. If the anticipated load of the other machine is smaller, the player moves to it with a probability that depends on several parameters: the degree of the actual vertices, the speeds of both machines and their load difference. Note that these are all local parameters.

The factor $\deg(i)/\deg(i,j) = \deg(i)/(\max\{\deg(i), \deg(j)\})$ is crucial to prevent that too many tasks from low degree vertices move to neighbors with a much larger degree. Alternatively, we can use $1/\Delta$ in our analysis but this would require that all tasks know the maximum degree.

Our first result in this section shows that, for sufficiently large m, we reach an apx. NE in time logarithmic in m. The second result bounds the time to reach a NE.

Theorem 3.1. *Let $m \geqslant \delta \cdot (8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta \cdot S)$ for some $\delta > 1$ and $\varepsilon = 1/(1+\delta)$. Then Protocol I reaches a ε-approximate Nash equilibrium in expected time*

$$\mathcal{O}\left(\ln(m) \cdot (\mathrm{diam}(G))^2 \cdot \Delta \cdot S \cdot \frac{s_{\max}}{s_{\min}}\right),$$

which is $\mathcal{O}(\ln(m) \cdot \mathrm{poly}(n) \cdot \mathrm{poly}(s_{\max}))$ on any graph.

Theorem 3.2. *If $m \geqslant 8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta \cdot S$, then Protocol I reaches a Nash equilibrium in expected time*

$$\mathcal{O}\left(\ln(m) \cdot (\mathrm{diam}(G))^2 \cdot \Delta \cdot S \cdot \frac{s_{\max}}{s_{\min}} + (\mathrm{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{max}^3\right)$$

on any graph. Otherwise, Protocol I reaches a Nash equilibrium in expected time

$$\mathcal{O}(\mathrm{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{max}^3) \ .$$

In both cases, the time is in $\mathcal{O}(\ln(m) \cdot \mathrm{poly}(n) \cdot \mathrm{poly}(s_{\max}))$ on any graph.

In expectation, the protocol behaves like a a continuous diffusion process. To avoid oscillation we need $\alpha \geqslant 4s_{\max}$. This, however, implies that even though tasks can have a large incentive for migration (e.g., if they move from a full and slow to an empty and fast machine), they never migrate with a probability of more than $1/(4s_{\max})$. Thus, to reach an apx. NE where *all* players have a small incentive to migrate, it might take $\Omega(s_{\max})$ many rounds for the last players to move. Thus, a convergence time that polynomially depends on $s_{\max}$ is unavoidable, given the way our protocol is defined.

In Section 3.1 we prove some fundamental bounds on the potential change in one step of Protocol I. Using these insights we show Theorem 3.1 and Theorem 3.2 in Section 3.2.

3.1 Potential Function Analysis

First we introduce some additional definitions. Note that throughout the paper we first estimate the expected potential decrease occuring in a single round of the process. In particular, we usually condition on the event that X^{t-1} is some arbitrary but fixed state x.

Definition 3.3. *For any $i, j \in V$ with $\{i, j\} \in E$ and any given state x, the* expected flow *along this edge in a single round of our protocol starting from x is*

$$f_{i,j}(x) := \begin{cases} \frac{L(x_i)-L(x_j)}{\alpha\cdot\deg(i,j)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} & \text{if } L(x_i) - L(x_j) > \frac{1}{s_j} \\ 0 & \text{otherwise.} \end{cases}$$

Note that $f_{i,j}(x)$ is always non-negative. Further, let

$$\widetilde{E}(x) := \left\{ (i,j) \in E \colon L(x_i) - L(x_j) > \frac{1}{s_j} \right\} \quad ,$$

be the set of edges over which tasks have an incentive to move when the system is in state x. Finally, for any $r \in \{0, 1\}$ we define

$$\Lambda^r_{i,j}(x) := (2\alpha - 2) \cdot \deg(i,j) \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot f_{i,j}(x) + \frac{r}{s_i} - \frac{r}{s_j} \quad .$$

We usually use the shorthand $\widetilde{E}$ instead of $\widetilde{E}(x)$. Our aim is to prove that in one round the system makes progress towards a NE, i.e., to show that $\mathbf{E}\left[\Delta\Phi_r(X^t) \mid X^{t-1} = x\right] > 0$. Let us first consider the potential change when the number of tasks transferred over any edge is *exactly* its expected number. Hence, we define

$$\widetilde{\Delta}\Phi_r(X^t \mid X_i^{t-1} = x) :=$$
$$\sum_{i\in V} \frac{W(x)\cdot(W(x)+r)}{s_i} - \sum_{i\in V} \frac{\left(\mathbf{E}\left[W(X_i^t) \mid X^{t-1} = x\right]\right)^2}{s_i} - r\cdot\sum_{i\in V}\frac{\mathbf{E}\left[W(X_i^t) \mid X^{t-1} = x\right]}{s_i} \quad .$$

Note that the expected load at machine i at step t given the load vector x at step $t-1$ equals

$$\mathbf{E}\left[W(X_i^t) \mid X^{t-1} = x\right] = W(x_i) + \sum_{j\in N(i)} \left(f_{j,i}(x) - f_{i,j}(x)\right).$$

The following lemma generalizes [7, Lemma 2] to the setting with speeds.

Lemma 3.4. *For any round $t \in \mathbb{N}$ it holds that*

$$\widetilde{\Delta}\Phi_r(X^t \mid X^{t-1} = x) \geqslant \sum_{(i,j)\in\widetilde{E}(x)} f_{i,j}(x)\cdot\Lambda^r_{i,j}(x) \quad .$$

Proof. Let $X^{t-1} = x$ be fixed. For reasons of simplification we sometimes omit the conditioning in this proof. We order the edges in $\widetilde{E}(x)$ increasingly according to $|f_{i,j}(x)|$. Let $\widetilde{E}(x) = \{e_1, e_2, \ldots, e_{|\widetilde{E}|(x)}\}$ be the edges in that order. We activate the edges in $\widetilde{E}(x)$ sequentially and bound the potential change after the activation of each edge. For any $1 \leqslant k \leqslant |\widetilde{E}(x)|$, let $X_i^{(t,k)}$ be the set of tasks assigned to vertex i after the k-th activation of an edge in $\widetilde{E}(x)$ in round t. Moreover, for any $1 \leqslant k \leqslant |\widetilde{E}(x)|$,

$$\widetilde{\Delta}\Phi_r(X^{(t,k)}) = \Phi_r(X^{(t,k)}) - \Phi_r(X^{(t,k-1)})$$

and

$$\widetilde{\Delta}\Phi_r(X^t) = \sum_{k=1}^{|\widetilde{E}|} \widetilde{\Delta}\Phi_r(X^{(t,k)}).$$

Let us now fix k with $e_k = (i,j)$ and consider $\widetilde{\Delta}\Phi_r(X^{(t,k)})$. Note that $f_{i,j}(x) > 0$. We have

$$W(X_j^{(t,k-1)}) \leqslant W(x_j) + (\deg(j) - 1) \cdot f_{i,j}(x)$$

Similarly,

$$W(X_i^{(t,k-1)}) \geqslant W(x_i) - (\deg(i) - 1) \cdot f_{i,j}(x).$$

Consequently,

$$\begin{aligned}
\widetilde{\Delta}\Phi_r(X^{(t,k)}) &= \frac{W(X_i^{(t,k-1)}) \cdot (W_i^{(t,k-1)}) + r)}{s_i} + \frac{W(X_j^{(t,k-1)}) \cdot (W(X_j^{(t,k-1)}) + r)}{s_j} - \\
&\quad - \frac{W(X_i^{(t,k)}) \cdot (W(X_i^{(t,k)}) + r)}{s_i} - \frac{W(X_j^{(t,k)}) \cdot (W(X_j^{(t,k)}) + r)}{s_j} \\
&= \frac{W(X_i^{(t,k-1)}) \cdot (W(X_i^{(t,k-1)}) + r)}{s_i} + \frac{W(X_j^{(t,k-1)}) \cdot (W(X_j^{(t,k-1)})) + r)}{s_j} - \\
&\quad - \frac{(W(X_i^{(t,k-1)}) - f_{i,j}(x)) \cdot (W(X_i^{(t,k-1)}) - f_{i,j}(x) + r)}{s_i} - \\
&\quad - \frac{(W(X_j^{(t,k-1)}) + f_{i,j}(x)) \cdot (W(X_j^{(t,k-1)}) + f_{i,j}(x) + r)}{s_j} \\
&= f_{i,j}(x) \cdot \left(2 \cdot \left(\frac{W(X_i^{(t,k-1)})}{s_i} - \frac{W(X_j^{(t,k-1)})}{s_j}\right) - \frac{f_{i,j}(x)}{s_i} - \frac{f_{i,j}(x)}{s_j} + \frac{r}{s_i} - \frac{r}{s_j}\right) \\
&\geqslant f_{i,j}(x) \cdot 2 \cdot \left(\frac{W(x_i) - (\deg(i) - 1) \cdot f_{i,j}(x)}{s_i} - \frac{W(x_j) + (\deg(j) - 1) \cdot f_{i,j}(x)}{s_j}\right) \\
&\quad - f_{i,j}(x) \cdot \left(\frac{f_{i,j}(x)}{s_i} - \frac{f_{i,j}(x)}{s_j} + \frac{r}{s_i} - \frac{r}{s_j}\right) \\
&\geqslant f_{i,j}(x) \cdot \left(2 \cdot \left(\frac{W(x_i)}{s_i} - \frac{W(x_j)}{s_j}\right) - (2d(i,j) - 1)\, f_{i,j}(x) \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) + \frac{r}{s_i} - \frac{r}{s_j}\right),
\end{aligned} \tag{3.1}$$

where until this point, we have not used the particular definition of $f_{i,j}(x)$. In the next step we use the definition of $f_{i,j}(x)$ to obtain

$$\begin{aligned}&\geqslant f_{i,j}(x)\cdot\left(2\cdot\alpha d(i,j)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot f_{i,j}(x)-2d(i,j)\cdot f_{i,j}(x)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)+\frac{r}{s_i}-\frac{r}{s_j}\right)\\&=f_{i,j}(x)\cdot\left((2\alpha-2)\cdot d(i,j)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot f_{i,j}(x)+\frac{r}{s_i}-\frac{r}{s_j}\right),\end{aligned}$$

and summing up over all edges yields the claim. □

The previous lemma lower bounds the potential improvement given that the expectations are exactly attained. To relate this to the real potential change, we have to upper bound the variance of the loads.

Lemma 3.5. *For any step t and any state x,*

$$\sum_{i\in V}\frac{\mathbf{Var}\left[W(X_i^t)\mid X^{t-1}=x\right]}{s_i}=\sum_{(i,j)\in\widetilde{E}}f_{i,j}(x)\left(\frac{1}{s_i}+\frac{1}{s_j}\right).$$

Proof. Note that

$$\sum_{i\in V}\frac{\mathbf{Var}\left[W(X_i^t)\mid X^{t-1}=x\right]}{s_i}=\sum_{i\in V}\frac{\mathbf{Var}\left[C_i^t-A_i^t\right]}{s_i},$$

where C_i^t and A_i^t are the random variables (conditioned on $X^{t-1}=x$) defined as follows. C_i^t is the random variable that counts the number of tasks that come to link i from another link in x, and A_i^t is the random variable that counts the number of tasks that abandon link i in x. Note that linearity of the variance for independent variables and the fact that $\mathbf{Var}[-Y]=\mathbf{Var}[Y]$ yields

$$\frac{\mathbf{Var}\left[C_i^t-A_i^t\right]}{s_i}=\frac{\mathbf{Var}\left[C_i^t\right]}{s_i}+\frac{\mathbf{Var}\left[-A_i^t\right]}{s_i}=\frac{\mathbf{Var}\left[C_i^t\right]}{s_i}+\frac{\mathbf{Var}\left[A_i^t\right]}{s_i}.$$

Observe that $C_i^t=\sum_{j\colon (j,i)\in\widetilde{E}}Z_{j,i}^t$, where $Z_{j,i}^t\geqslant 0$ is the random variable that counts the number tasks that move from j to i from x. We have

$$Z_{j,i}^t\sim\mathsf{Bin}\left(W(x_j),\frac{1}{\alpha d(i,j)}\cdot\frac{L(x_j)-L(x_i)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W(x_j)}\right).$$

Similarly,

$$A_i^t\sim\mathsf{Bin}\left(W(x_i),\frac{1}{\alpha d(i,j)}\cdot\sum_{j\colon (i,j)\in\widetilde{E}}\frac{L(x_i)-L(x_j)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W(x_i)}\right).$$

Since $\left\{Z_{j,i}^t\colon j\in N(i)\right\}$ is a set of independent random variables, we have

$$\begin{aligned}\mathbf{Var}\left[C_i^t\right]&=\sum_{j\colon (j,i)\in\widetilde{E}}\mathbf{Var}\left[\mathsf{Bin}\left(W(x_j),\frac{1}{\alpha d(i,j)}\cdot\frac{L(x_j)-L(x_i)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W(x_j)}\right)\right]\\&\leqslant\sum_{j\colon (j,i)\in\widetilde{E}}\frac{1}{\alpha d(i,j)}\cdot\frac{L(x_j)-L(x_i)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)},\end{aligned}$$

where the last inequality follows from $\mathbf{Var}\left[\mathsf{Bin}(n,p)\right] = np(1-p) \leqslant np$. Similarly,

$$\mathbf{Var}\left[A_i^t\right] = \mathbf{Var}\left[\mathsf{Bin}\left(W(x_i), \frac{1}{\alpha d(i,j)} \cdot \sum_{j\colon (i,j)\in\widetilde{E}} \frac{L(x_i)-L(x_j)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W(x_i)}\right)\right]$$
$$\leqslant \sum_{j\colon (i,j)\in\widetilde{E}} \frac{1}{\alpha d(i,j)} \cdot \frac{L(x_i)-L(x_j)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} \quad .$$

Hence,

$$\sum_{i\in V} \frac{\mathbf{Var}\left[W(X_i^t) \mid X^{t-1}=x\right]}{s_i}$$
$$\leqslant \sum_{i\in V} \frac{1}{s_i}\cdot\left(\sum_{j\colon (j,i)\in\widetilde{E}} \frac{1}{\alpha d(i,j)}\cdot\frac{L(x_j)-L(x_i)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} + \sum_{j\colon (i,j)\in\widetilde{E}} \frac{1}{\alpha d(i,j)}\cdot\frac{L(x_i)-L(x_j)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\right) \quad .$$

Since every edge in $\widetilde{E}$ occurs exactly twice in above sum, once with weight $1/s_i$ and once with weight $1/s_j$, we obtain

$$\sum_{i\in V} \frac{\mathbf{Var}\left[W(X_i^t) \mid X^{t-1}=x\right]}{s_i} \leqslant \sum_{(i,j)\in\widetilde{E}} \left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot\frac{1}{\alpha d(i,j)}\cdot\frac{L(x_i)-L(x_j)}{\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}$$
$$= \sum_{(i,j)\in\widetilde{E}} \frac{L(x_i)-L(x_j)}{\alpha d(i,j)} = \sum_{(i,j)\in\widetilde{E}} f_{i,j}(x)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right) \quad .$$

□

With Lemma 3.5 at hand, we will now provide a bound on the real potential change, i.e., we include the deviation which occurs since the actual number of transferred tasks may differ from their expected values.

Lemma 3.6. *For any step t and any state x,*

$$\mathbf{E}\left[\Delta\Phi_r(X^t) \mid X^{t-1}=x\right] \quad \geqslant \quad \sum_{(i,j)\in\widetilde{E}(x)} f_{i,j}(x)\cdot\left(\Lambda_{i,j}^r(x) - \frac{1}{s_i} - \frac{1}{s_j}\right) \quad .$$

Proof. We obtain

$$\begin{aligned}
\mathbf{E}\left[\Delta\Phi_r(X^t) \mid X^{t-1}=x\right] &= \Phi_r(x) - \mathbf{E}\left[\Phi_r(X^t) \mid X^{t-1}=x\right] \\
&= \sum_{i\in V} \frac{W(x_i)\cdot(W(x_i)+r)}{s_i} - \sum_{i\in V} r\cdot\frac{\mathbf{E}\left[W(X_i^t) \mid X^{t-1}=x\right]}{s_i} - \\
&\quad - \sum_{i\in V} \frac{\mathbf{E}\left[(W(X_i^t))^2 \mid X^{t-1}=x\right]}{s_i} \\
&= \sum_{i\in V} \frac{W(x_i)\cdot(W(x_i)+r)}{s_i} - \sum_{i\in V} r\cdot\frac{\mathbf{E}\left[W(X_i^t) \mid X^{t-1}=x\right]}{s_i} - \\
&\quad - \sum_{i\in V} \frac{\mathbf{E}\left[W(X_i^t) \mid X^{t-1}=x\right]^2}{s_i} - \sum_{i\in V} \frac{\mathbf{Var}\left[W(X_i^t) \mid X^{t-1}=x\right]}{s_i}
\end{aligned}$$

$$
\begin{aligned}
&= \widetilde{\Delta}\Phi_r(X^t \mid X^{t-1} = x) - \sum_{i \in V} \frac{\mathbf{Var}\left[W(X_i^t) \mid X^{t-1} = x\right]}{s_i} \\
&\geqslant \sum_{(i,j) \in \widetilde{E}} f_{i,j}(x) \cdot \left(\Lambda_{i,j}^r(x) - \frac{1}{s_i} - \frac{1}{s_j} \right),
\end{aligned}
$$

where the last inequality follows by applying Lemma 3.4 and Lemma 3.5. □

The next technical lemma builds on Lemma 3.6 and shows that every edge $(i,j) \in \widetilde{E}(x)$ with a sufficiently large load difference contributes positively to $\mathbf{E}\left[\Delta\Phi_r(X^t) \mid X^{t-1} = x\right]$.

Lemma 3.7. *Set $\alpha := 4s_{\max} \geqslant 4$. Then the following two statements hold for any state x:*

(1) If $r = 1$, then for any $(i,j) \in \widetilde{E}(x)$ with

$$
L(x_i) - L(x_j) \geqslant \frac{1}{s_j} + \frac{1}{s_i \cdot s_j},
$$

we have

$$
\Lambda_{i,j}^1(x) - \frac{1}{s_i} - \frac{1}{s_j} \geqslant \frac{1}{2s_{\max}} \left(\frac{1}{s_i} + \frac{1}{s_j} \right) \quad .
$$

(2) If $r = 0$, then for any $(i,j) \in \widetilde{E}(x)$ with

$$
L(x_i) - L(x_j) \geqslant \frac{1}{s_j} + \frac{1}{s_i},
$$

we have

$$
\Lambda_{i,j}^0(x) - \frac{1}{s_i} - \frac{1}{s_j} \geqslant \frac{1}{2} \left(\frac{1}{s_i} + \frac{1}{s_j} \right) \quad .
$$

Proof. We first start with the claim for $r = 1$. We obtain that

$$
\begin{aligned}
\Lambda_{i,j}^1(x) &= (2\alpha - 2) \cdot d(i,j) \cdot \left(\frac{1}{s_i} + \frac{1}{s_j} \right) \cdot f_{i,j}(x) + \frac{r}{s_i} - \frac{r}{s_j} && \text{(by definition)} \\
&= (2\alpha - 2) \cdot d(i,j) \cdot \frac{L(x_i) - L(x_j)}{\alpha d(i,j)} + \frac{1}{s_i} - \frac{1}{s_j} \\
&\geqslant \left(2 - \frac{1}{2s_{\max}} \right) \cdot \left(\frac{1}{s_j} + \frac{1}{s_i s_j} \right) + \frac{1}{s_i} - \frac{1}{s_j} && \text{(by assumption)} \\
&= \frac{2}{s_j} - \frac{1}{2s_j s_{\max}} + \frac{2}{s_i s_j} - \frac{1}{2s_{\max} s_i s_j} + \frac{1}{s_i} - \frac{1}{s_j} \\
&\geqslant \frac{1}{s_i} + \frac{1}{s_j} + \frac{2}{s_i s_j} - \frac{1}{s_j s_{\max}} && \text{(since } s_i \geqslant 1\text{)} \\
&\geqslant \frac{1}{s_i} + \frac{1}{s_j} + \frac{1}{s_i s_j} = \left(1 + \frac{1}{s_i + s_j} \right) \cdot \left(\frac{1}{s_i} + \frac{1}{s_j} \right) \\
&\geqslant \left(1 + \frac{1}{2s_{\max}} \right) \cdot \left(\frac{1}{s_i} + \frac{1}{s_j} \right).
\end{aligned}
$$

Similarly, consider now the second claim where $r = 0$. Then

$$\begin{aligned}
\Lambda^0_{i,j}(x) &= (2\alpha - 2) \cdot d(i,j) \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot f_{i,j}(x) + \frac{r}{s_i} - \frac{r}{s_j} && \text{(by definition)} \\
&= (2\alpha - 2) \cdot d(i,j) \cdot \frac{L(x_i) - L(x_j)}{\alpha d(i,j)} \\
&\geqslant \left(2 - \frac{1}{2s_{\max}}\right) \cdot \left(\frac{1}{s_j} + \frac{1}{s_i}\right) && \text{(by assumption)} \\
&= \frac{2}{s_j} - \frac{1}{2s_j s_{\max}} + \frac{2}{s_i} - \frac{1}{2s_{\max} s_i} \\
&\geqslant \frac{3}{2s_i} + \frac{3}{2s_j}. && \text{(since } s_{\max} \geqslant 1)
\end{aligned}$$

□

3.2 Convergence to Approximate and Exact Nash Equilibria

In this section, we finally prove our main theorems. We use Lemma 3.4 for the migration in a single round and consider the potential drop w.r.t. an "ideal" potential value if exactly the expected loads are realized. Then, using Lemma 3.6, we bound the difference between ideal and realized potential values by analyzing the variance of the migration process. This yields bounds on the expected drop of the potential in one round and is the main ingredient to prove the theorems. Let us define for any given state x the *maximum load difference* as

$$L_\Delta(x) := \max_{i \in V} |L(x_i) - m/S| \quad .$$

In order to prove our first theorem, we use function Φ_0 and its normalized version Ψ_0. For any given state x, we observe a simple relation between the potential value $\Phi_0(x)$ and the maximum load difference $L_\Delta(x)$. Note that $\Phi_0(x) = \sum_{i \in V} (W(x_i))^2 / s_i$ is minimized when all $W(x_i) = m/S \cdot s_i$. This implies

$$\Phi_0(x) \geqslant \sum_{i \in V} (m^2 \cdot s_i)/S^2 = m^2/S.$$

The following observation considers the normalized version of $\Psi_0(x)$.

Observation 3.8. *For any state x of the system it holds that*

$$(L_\Delta(x))^2 \leqslant \Psi_0(x) = \Phi_0(x) - \frac{m^2}{S} \leqslant S \cdot (L_\Delta(x))^2.$$

Proof. Clearly,

$$\begin{aligned}
\Phi_0(x) - \frac{m^2}{S} &= \sum_{i \in V} \frac{(m/S \cdot s_i + (W(x_i) - m/S \cdot s_i))^2}{s_i} - \frac{m^2}{S} \\
&= \sum_{i \in V} \frac{(m/S \cdot s_i)^2}{s_i} + 2 \sum_{i \in V} \frac{m/S \cdot s_i \cdot (W(x_i) - m/S \cdot s_i)}{s_i} + \sum_{i \in V} \frac{(W(x_i) - m/S \cdot s_i)^2}{s_i} - \frac{m^2}{S} \\
&= 2 \sum_{i \in V} \frac{m/S \cdot s_i \cdot W(x_i) - (m/S)^2 \cdot s_i^2}{s_i} + \sum_{i \in V} \frac{(W(x_i) - m/S \cdot s_i)^2}{s_i} \\
&= \sum_{i \in V} \frac{(W(x_i) - m/S \cdot s_i)^2}{s_i}
\end{aligned}$$

and

$$(L_\Delta(x))^2 \leqslant \left(\max_{i\in V}\left|\frac{W(x_i)}{s_i} - m/S\right|\right)^2 = \max_{i\in V}\frac{(W(x_i) - m/S\cdot s_i)^2}{s_i^2} \leqslant \sum_{i\in V}\frac{(W(x_i) - m/S\cdot s_i)^2}{s_i}.$$

Moreover,

$$\sum_{i\in V}\frac{(W(x_i) - m/S\cdot s_i)^2}{s_i} = \sum_{i\in V} s_i\cdot (W(x_i)/s_i - m/S)^2 \leqslant S\cdot (L_\Delta(x))^2.$$

□

Lemma 3.9. *Let* $\gamma := 32\,S\cdot(\mathrm{diam}(G))^2\cdot s_{\max}/s_{\min}\cdot\Delta$. *Let* x *be a state such that* $L_\Delta(x)\geqslant 8\cdot\mathrm{diam}(G)\cdot n\cdot\Delta$. *Then*

$$\mathbf{E}\left[\Psi_0(X^t)\mid X^{t-1}=x\right]\leqslant\left(1-\frac{1}{\gamma}\right)\cdot\Psi_0(x)$$

Proof. In this proof, we consider the progress of $\Phi_0(X^t)$ for an arbitrary but fixed assignment $X^{t-1}=x$. Consider now the protocol before the execution of the next step t. Let $l\in V$ be a vertex with $|L(x_l)-m/S| = L_\Delta(x)$, and assume w.l.o.g. that $L(x) > m/S$. Then there must be another vertex $k\in V$ with $L(x_k)\leqslant m/S$. This implies that there is an edge $\{p,q\}\in E$ on a path from l to k such that

$$L(x_p) - m/S \geqslant L(x_q) - m/S + L_\Delta(x)/\,\mathrm{diam}(G)\ \ ,$$

which is equivalent to

$$L(x_p) - L(x_q) \geqslant L_\Delta(x)/\,\mathrm{diam}(G)\ \ .$$

By Lemma 3.6,

$$\mathbf{E}\left[\Delta\Phi_0(X^t)\mid X^{t-1}=x\right]\geqslant\sum_{(i,j)\in\widetilde{E}} f_{i,j}(x)\cdot\left(\Lambda_{i,j}(x)-\frac{1}{s_i}-\frac{1}{s_j}\right)\quad .$$

Let us now define for any edge $(i,j)\in\widetilde{E}$,

$$\Delta_{i,j}\Phi_0(X^t\,|\,X^{t-1}=x) := f_{i,j}(x)\cdot\left(\Lambda^0_{i,j}(x)-\frac{1}{s_i}-\frac{1}{s_j}\right)\quad ,$$

so that

$$\mathbf{E}\left[\Delta\Phi_0(X^t)\mid X^{t-1}=x\right]\geqslant\sum_{(i,j)\in\widetilde{E}}\Delta_{i,j}\Phi(X^t\mid X^{t-1}=x).$$

Let us now group $\widetilde{E}$ into three disjoint groups (where we omit the argument x throughout):

$$\widetilde{E}_1 := \left\{(i,j)\in\widetilde{E}\colon L(x_i)-L(x_j)\geqslant\frac{1}{s_i}+\frac{1}{s_j}\right\}\Big\backslash (p,q)$$
$$\widetilde{E}_2 := \left\{(i,j)\in\widetilde{E}\colon L(x_i)-L(x_j)<\frac{1}{s_i}+\frac{1}{s_j}\right\}\Big\backslash (p,q)$$
$$\widetilde{E}_3 := (p,q)$$

The intuition behind this grouping is as follows. $\widetilde{E}_1$ contains all edges with a load difference large enough so that, in expectation, the edge yields a decrease of the potential. $\widetilde{E}_2$ contains all edges with a small load difference. Due to the variance, these edges could increase the potential. However, since the load difference is small, their total impact on the potential is limited. Finally, $\widetilde{E}_3$ consists of the edge (p,q) which has a large load difference and thereby decreases the potential significantly (in expectation).

Group 1: Lemma 3.7 with $r=0$ shows that

$$\sum_{(i,j)\in\widetilde{E}_1} \Delta_{i,j}\Phi_0(X^t \mid X^{t-1}=x) = \sum_{(i,j)\in\widetilde{E}_1} f_{i,j}(x)\cdot\left(\Lambda^0_{i,j}(x)-\frac{1}{s_i}-\frac{1}{s_j}\right) \geqslant 0 \ .$$

Group 2: In this case we consider the value of $\sum_{(i,j)\in\widetilde{E}_2}\Delta_{i,j}\Phi_0(X^t \mid X^{t-1}=x)$. Let (i,j) be an edge in $\widetilde{E}_2$. Then, plugging in the definition of $\Lambda^0_{i,j}(x)$ and $f_{i,j}(x)$ yields

$$\begin{aligned}
&\Delta_{i,j}\Phi_0(X^t \mid X^{t-1}=x)\\
&= f_{i,j}(x)\cdot\Lambda^0_{i,j}(x) - f_{i,j}(x)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\\
&= \frac{L(x_i)-L(x_j)}{\alpha\cdot d(i,j)\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\left((2\alpha-2)\cdot d(i,j)\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot f_{i,j}(x)\right) - \frac{L(x_i)-L(x_j)}{\alpha\cdot d(i,j)}\\
&= \frac{L(x_i)-L(x_j)}{\alpha\cdot d(i,j)\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\left(\left(2-\frac{2}{\alpha}\right)\cdot\underbrace{(L(x_i)-L(x_j))}_{\geqslant 0}\right) - \frac{L(x_i)-L(x_j)}{\alpha\cdot d(i,j)}\\
&\geqslant -\frac{L(x_i)-L(x_j)}{\alpha\cdot d(i,j)} > -\frac{\frac{1}{s_j}+\frac{1}{s_i}}{\alpha} \geqslant -\frac{2}{\alpha}\ ,
\end{aligned}$$

since we assumed that all speeds are larger than one. Summing up over all edges in $\widetilde{E}_2$, we obtain

$$\sum_{(i,j)\in\widetilde{E}_2} \Delta_{i,j}\Phi_0(X^t \mid X^{t-1}=x) \geqslant -\frac{2|E|}{\alpha}\ .$$

Group 3: Consider now the edge (p,q), i.e., the set $\widetilde{E}_3$:

$$\begin{aligned}
\Delta_{p,q}\Phi_0(X^t \mid X^{t-1}=x) &= f_{p,q}(x)\cdot\left(\Lambda^0_{p,q}(x)-\frac{1}{s_p}-\frac{1}{s_q}\right)\\
&\geqslant \frac{L(x_p)-L(x_q)}{\alpha\cdot d(p,q)\left(\frac{1}{s_p}+\frac{1}{s_q}\right)}\cdot\left(\left(2-\frac{2}{\alpha}\right)\cdot(L(x_p)-L(x_q))-2\right)\\
&\geqslant \frac{L(x_p)-L(x_q)}{\alpha\cdot d(p,q)\left(\frac{1}{s_p}+\frac{1}{s_q}\right)}\cdot\left(\frac{L_\Delta(x)}{\operatorname{diam}(G)}-2\right)\\
&\geqslant \frac{L_\Delta(x)/\operatorname{diam}(G)}{\alpha\cdot d(p,q)\left(\frac{1}{s_p}+\frac{1}{s_q}\right)}\cdot\frac{L_\Delta(x)}{2\operatorname{diam}(G)}\\
&= \frac{(L_\Delta(x))^2}{2(\operatorname{diam}(G))^2\cdot\alpha\cdot d(p,q)\left(\frac{1}{s_p}+\frac{1}{s_q}\right)}\ ,
\end{aligned}$$

where the second last inequality holds since $L(x_p) - L(x_q) \geqslant L_\Delta(x)/\operatorname{diam}(G)$ and $\alpha \geqslant 2$. The last inequality holds since by assumption, $L_\Delta(x) \geqslant 8 \cdot \operatorname{diam}(G) \cdot n \cdot \Delta$. Combining the contribution of all three groups yields

$$\begin{aligned}
&\mathbf{E}\left[\Delta\Phi_0(X^t) \mid X^{t-1} = x\right] \\
&\geqslant \sum_{\{i,j\} \in \widetilde{E}_1} \Delta_{i,j}\Phi_0(X^t \mid X^{t-1} = x) + \sum_{\{i,j\} \in \widetilde{E}_2} \Delta_{i,j}\Phi_0(X^t \mid X^{t-1} = x) + \Delta_{p,q}\Phi_0(X^t \mid X^{t-1} = x) \\
&\geqslant 0 - \frac{2|E|}{\alpha} + \frac{(L_\Delta(x))^2}{2\,(\operatorname{diam}(G))^2 \cdot \alpha \cdot d(p,q) \cdot \left(\frac{1}{s_p} + \frac{1}{s_q}\right)} \\
&\geqslant \frac{(L_\Delta(x))^2}{4\,(\operatorname{diam}(G))^2 \cdot \alpha \cdot d(p,q) \cdot \left(\frac{1}{s_p} + \frac{1}{s_q}\right)},
\end{aligned}$$

where the last inequality uses $L_\Delta(x) \geqslant 8 \cdot \operatorname{diam}(G) \cdot n \cdot \Delta$.

Now we continue with the normalized version of $\Phi_0(x)$, i.e., with $\Psi_0(x)$. By Observation 3.8, $\Psi_0(x) \leqslant S \cdot (L_\Delta(x))^2$. Since $\Delta\Psi_0(X^t) = \Delta\Phi_0(X^t)$ and $\alpha = 4s_{\max}$ we obtain (as long as $L_\Delta(x) > 8 \cdot \operatorname{diam}(G) \cdot n \cdot \Delta$) that

$$\begin{aligned}
\mathbf{E}\left[\Delta\Psi_0(X^t) \mid X^{t-1} = x\right] &\geqslant \frac{(L_\Delta(x))^2}{4\,(\operatorname{diam}(G))^2 \cdot \alpha \cdot d(p,q) \cdot \left(\frac{1}{s_p} + \frac{1}{s_q}\right)} \\
&\geqslant \frac{\Psi_0(x)}{4\,S \cdot (\operatorname{diam}(G))^2 \cdot \alpha \cdot \Delta \cdot \left(\frac{1}{s_p} + \frac{1}{s_q}\right)} \\
&\geqslant \frac{\Psi_0(x)}{32\,S \cdot (\operatorname{diam}(G))^2 \cdot \frac{s_{\max}}{s_{\min}} \cdot \Delta} \enspace .
\end{aligned}$$

With $\gamma := 32\,S \cdot (\operatorname{diam}(G))^2 \cdot s_{\max}/s_{\min} \cdot \Delta$ we have $\mathbf{E}\left[\Delta\Psi_0(X^t) \mid X^{t-1} = x\right] \geqslant \frac{1}{\gamma} \cdot \Psi_0(x)$, or equivalently

$$\mathbf{E}\left[\Psi_0(X^t) \mid X^{t-1} = x\right] \leqslant \left(1 - \frac{1}{\gamma}\right) \cdot \Psi_0(x).$$

□

Lemma 3.10. *Let $\gamma := 32\,S \cdot (\operatorname{diam}(G))^2 \cdot s_{\max}/s_{\min} \cdot \Delta$ and $\tau := \gamma \cdot (2\ln n + \ln(\Psi_0(X^0)))$. Then, with probability at least $1 - n^{-1}$, there exists a round $t \in [1, \tau]$ such that*

$$L_\Delta(X^t) < 8\operatorname{diam}(G) \cdot n \cdot \Delta.$$

Proof. Let us define an auxiliary random variable $\widetilde{\Psi}_0$ by $\widetilde{\Psi}_0(X^0) := \Psi_0(X^0)$, and for any round $t \geqslant 1$,

$$\widetilde{\Psi}_0(X^t) = \begin{cases} \Psi_0(X^t) & \text{if } [L_\Delta(X^{t-1}) \geqslant 8\operatorname{diam}(G) \cdot n \cdot \Delta] \wedge [\widetilde{\Psi}_0(X^{t-1}) > 0] \\ 0 & \text{otherwise.} \end{cases}$$

Then, for any $t \geqslant 1$, it follows by Lemma 3.9,

$$\mathbf{E}\left[\widetilde{\Psi}_0(X^t) \mid X^{t-1} = x\right] \leqslant \left(1 - \frac{1}{\gamma}\right) \cdot \widetilde{\Psi}_0(X^{t-1}).$$

We have for $\tau = \gamma \cdot (2\ln n + \ln \Psi_0(X^0))$,

$$\begin{aligned}\mathbf{E}\left[\widetilde{\Psi}_0(X^t)\right] &= \sum_x \mathbf{E}\left[\widetilde{\Psi}_0(X^t) \mid X^{t-1} = x\right] \cdot P[X^{t-1} = x] \\ &\leqslant \sum_x \left(1 - \frac{1}{\gamma}\right) \cdot \widetilde{\Psi}_0(X^{t-1}) \cdot P[X^{t-1} = x] \\ &\leqslant \left(1 - \frac{1}{\gamma}\right)^{\tau} \cdot \widetilde{\Psi}_0(X^0) \leqslant n^{-2}.\end{aligned}$$

Hence by Markov's inequality,

$$\mathbf{Pr}\left[\widetilde{\Psi}_0(X^\tau) \geqslant n^{-1}\right] \leqslant n^{-1}.$$

We consider two cases.

Case 1: For all time-steps $t \in [0, \ldots, \tau]$, $\widetilde{\Psi}_0(X^t) = \Psi_0(X^t)$. Then by Observation 3.8,

$$L_\Delta(X^\tau) \leqslant \sqrt{\Psi_0(X^t)} = \sqrt{\widetilde{\Psi}_0(X^t)} \leqslant n^{-1/2}.$$

Case 2: There exists a step $t \in [1, \ldots, \tau]$ such that $\widetilde{\Psi}_0(X^t) \neq \Psi_0(X^t)$. Let t be the smallest time step with that property. Hence, $\widetilde{\Psi}_0(X^t) \neq \Psi_0(X^t)$, but $\widetilde{\Psi}_0(X^{t-1}) = \Psi_0(X^{t-1})$. If $\widetilde{\Psi}_0(X^{t-1}) = 0$, then

$$L_\Delta(X^{t-1}) \leqslant \sqrt{\Psi_0(X^{t-1})} = \sqrt{\widetilde{\Psi}_0(X^{t-1})} = 0.$$

If $\widetilde{\Psi}_0(X^{t-1}) \neq 0$, then by definition of $\widetilde{\Psi}_0(X^t)$,

$$\left(\widetilde{\Psi}_0(X^t) \neq \Psi_0(X^t)\right) \bigwedge \left(\widetilde{\Psi}_0(X^{t-1}) \neq 0\right) \Rightarrow L_\Delta(X^t) < 8\,\mathrm{diam}(G) \cdot n \cdot \Delta.$$

In all cases we have shown that there exists a step $\in [0, \tau]$ so that $L_\Delta(X^t) < 8\,\mathrm{diam}(G) \cdot n \cdot \Delta$. This completes the proof of the lemma. □

3.2.1 Proof of Theorem 3.1

First note that for any state x we have $\Psi_0(x) \leqslant m^2/s_{\min} \leqslant m^2$. We now show that for $m \geqslant 8\delta \cdot n^3$, any state x with $L_\Delta(x) \leqslant 8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta \leqslant 8 \cdot n^3$ is indeed a $1/(1+\delta)$-apx. NE. For every $i \in V$ we have

$$|W(x_i)/s_i - m/S| \leqslant L_\Delta(x) \leqslant 8 \cdot n^3 \quad .$$

Consider now any pair i, j with $\{i, j\} \in E$. Then $W(x_i)/s_i \leqslant 8 \cdot n^3 + m/S$ and similarly

$$(W(x_j) + 1)/s_j \geqslant \max\left\{1/s_j, (m/S) - 8n^3 + (1/s_j)\right\} \quad .$$

We are looking for the smallest possible $\varepsilon \in [0, 1)$ such that $(1 - \varepsilon) \cdot W(X_i^T)/s_i \leqslant (W(X_j^T) + 1)/s_j$. Plugging in our bounds from above, a simple calculation yields the result $m \geqslant 8\delta \cdot n^3$.

From Lemma 3.10, it follows now that with probability at least $1 - n^{-4}$ after

$$\begin{aligned}\tau &= (32S \cdot (\mathrm{diam}(G))^2 \cdot s_{\max}/s_{\min} \cdot \Delta) \cdot (4\ln n + \ln \Psi_0(X^0)) \\ &= \mathcal{O}\left(\ln(m) \cdot (\mathrm{diam}(G))^2 \cdot \Delta \cdot S \cdot \frac{s_{\max}}{s_{\min}}\right)\end{aligned}$$

steps, we reach a $1/(1+\delta)$-apx. NE in the time-interval $[0, 1, \ldots, \tau]$. Hence after expected $\mathcal{O}\left(\ln(m) \cdot (\mathrm{diam}(G))^2 \cdot \Delta \cdot S \cdot \frac{s_{\max}}{s_{\min}}\right)$ rounds, we have reached a $1/(1+\delta)$-apx. NE. This proves Theorem 3.1.

3.2.2 Proof of Theorem 3.2

Now we show that there is still a small potential drop even if L_Δ is already relatively small. First we show some results that will be needed in our later proofs.

Observation 3.11. *Let x be a state that is not a Nash equilibrium. Then there exist two neighbouring vertices i, j such that $L(x_i) - L(x_j) > \frac{1}{s_j}$. If all speeds are integers, then we also have*

$$L(x_i) - L(x_j) \geqslant \frac{1}{s_j} + \frac{1}{s_i s_j} \quad .$$

Proof. The fact that two neighbouring vertices i, j exist with $L(x_i) - L(x_j) > \frac{1}{s_j}$ follows directly from the definition of a NE. Also,

$$\frac{W(x_i)}{s_i} - \frac{W(x_j)}{s_j} > \frac{1}{s_j} \Leftrightarrow W(x_i) > \frac{s_i}{s_j} \cdot (W(x_j) + 1). \tag{3.2}$$

Let us now fix any s_i, s_j and sum of weighted tasks $W(x_j)$. Our goal is to lower bound the smallest possible $W(x_i)$ such that the strict inequality above is fulfilled. We proceed by a case distinction. First we assume $s_i/s_j \cdot (W(x_j) + 1)$ is an integer. Then,

$$W(x_i) \geqslant \frac{s_i}{s_j} \cdot (W(x_j) + 1) + 1,$$

and therefore,

$$L(x_i) - L(x_j) = \frac{W(x_i)}{s_i} - \frac{W(x_j)}{s_j} \geqslant \frac{1}{s_j} + \frac{1}{s_i} \geqslant \frac{1}{s_j} + \frac{1}{s_i s_j} \quad .$$

Now we assume that $\frac{s_i}{s_j} \cdot (W(x_j) + 1)$ is not an integer. Hence,

$$W(x_i) \geqslant \left\lceil \frac{s_i}{s_j} \cdot (W(x_j) + 1) \right\rceil .$$

We now use the fact that $\lceil \frac{a}{b} \rceil - \frac{a}{b} \geqslant \frac{1}{b}$ for any pair of integers a, b (where a is not a multiple of b) to conclude that

$$W(x_i) \geqslant \frac{s_i}{s_j} \cdot (W(x_j) + 1) + \frac{1}{s_j}$$

and

$$L(x_i) - L(x_j) = \frac{W(x_i)}{s_i} - \frac{W(x_j)}{s_j} \geqslant \frac{1}{s_j} + \frac{1}{s_i \cdot s_j}.$$

□

Again, we define a normalized version of Φ_1. For every state x, we define

$$\Psi_1(x) = \Phi_1(x) - \left(\frac{m^2}{S} + n \cdot \frac{m}{S} - \frac{n}{4 s_{\min}} \right) .$$

The next Lemma shows a relation between L_Δ and Ψ_1.

Lemma 3.12. *Let x be an arbitrary state. With the definition above we have*

$$0 \leqslant \Psi_1(x) \leqslant S \cdot (L_\Delta(x))^2 + n \cdot L_\Delta(x) + n \quad .$$

Proof. We first show the left inequality, i.e., we show a lower bound on $\Phi_1(x)$. In order to bound the minimum, we characterize the NE with minimum potential. The loads are equilibrated to m/S on all machines when the weight is $W(x_i) = m \cdot s_i/S$ for every machine i. Thus, in every state x of the game we have a weight $W(x_i) = m \cdot s_i/S + \delta_i$, where $\delta_i \in \left[-\frac{ms_i}{S}, m - \frac{ms_i}{S}\right]$ and $\sum_{i \in V} \delta_i = 0$. Using this insight, we bound

$$\begin{aligned}
\Phi_1(x) &= \sum_{i \in V} \frac{\left(\frac{ms_i}{S} + \delta_i\right)^2}{s_i} + \frac{\frac{ms_i}{S} + \delta_i}{s_i} \\
&= \sum_{i \in V} \frac{m^2 s_i}{S^2} + \frac{2m\delta_i}{S} + \frac{\delta_i^2}{s_i} + \frac{m}{S} + \frac{\delta_i}{s_i} \\
&= \frac{m^2}{S} + \frac{nm}{S} + \frac{2m}{S}\left(\sum_{i \in V} \delta_i\right) + \sum_{i \in V} \frac{\delta_i^2}{s_i} + \frac{\delta_i}{s_i} \\
&= \frac{m^2}{S} + \frac{nm}{S} + \sum_{i \in V} \frac{\delta_i(\delta_i + 1)}{s_i} \\
&\geqslant \frac{m^2}{S} + \frac{nm}{S} - \frac{n}{4s_{\min}} \ ,
\end{aligned}$$

because the function $y(y+1) \geqslant -1/4$ for any y. This proves the lower bound in the lemma.

To show the upper bound, we note that $L(x_i) - m/S \leqslant L_\Delta(x)$ and calculate

$$\begin{aligned}
& \Phi_1(x) - \frac{m^2}{S} - \frac{nm}{S} + \frac{n}{4s_{\min}} \\
=\ & \Phi_0(x) + \left(\sum_{i \in V} L(x_i)\right) - \frac{m^2}{S} - \frac{nm}{S} + \frac{n}{4s_{\min}} \\
\leqslant\ & S \cdot (L_\Delta(x))^2 + \left(\sum_{i \in V} L(x_i) - \frac{m}{S}\right) + \frac{n}{4s_{\min}} \\
\leqslant\ & S \cdot (L_\Delta(x))^2 + n \cdot L_\Delta(x) + n \ ,
\end{aligned}$$

where the first inequality follows from Observation 3.8. This finishes the proof of the lemma. □

Lemma 3.13. *Assume that at step t the system is not in a Nash equilibrium. We have*

$$\mathbf{E}\left[\Psi_1(X^{t+1}) \,|\, X^t = x\right] \leqslant \Psi_1(x) - \frac{1}{8\Delta s_{\max}^3} \ .$$

Proof. Note that Observation 3.11 provides the minimum decrease required to apply Lemma 3.7 for $\Phi_1(x)$. Plugging in this bound for $\Lambda^1_{i,j}$ in Lemma 3.6 yields an expected (additive) decrease of at least $1/(8\Delta s_{\max}^3)$ per round, as long as x is not a NE. □

Lemma 3.14. *Assume that at step τ we have $L_\Delta(X^\tau) \leqslant 8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta$. Let T be the additional number of steps such that $X^{\tau+T}$ is a Nash equilibrium. Then*

$$\mathbf{E}\left[T\right] \leqslant 560 \cdot (\mathrm{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{\max}^3 \ .$$

Proof. For reasons of simplification let us assume $\tau = 0$. Lemma 3.12 implies that for every state x with $L_\Delta(x) \leqslant 8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta$

$$\begin{aligned}
\Psi_1(x) &\leqslant S \cdot (8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta)^2 + n \cdot 8 \cdot \mathrm{diam}(G) \cdot n \cdot \Delta + n \\
&\leqslant S \cdot 70 \cdot (\mathrm{diam}(G))^2 \cdot n^2 \cdot \Delta^2 \ .
\end{aligned}$$

Hence we can assume that $\mathbf{E}\left[\Psi_1(X^0)\right] \leqslant 70 \cdot S \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^2$.

The rest of the proof is done by a standard martingale argument. Let T be the end of the first time step after the system reaches a NE. Let $t \wedge T$ be the minimum of t and T and let $V = 1/(8\Delta s_{\max}^3)$. We define $Z_t = \Psi_1(X^t) + t \cdot V$. Observe that T is a stopping time for $(Z_t)_{t \geqslant 0}$. $\{Z_t\}_{t \wedge T}$ is a supermartingale since by Lemma 3.13 with $X^t = x$,

$$\begin{aligned}
\mathbf{E}\left[Z_{t+1} \,|\, Z_t = z\right] &= \mathbf{E}\left[\Psi_1(X^{t+1}) + V \cdot (t+1) \mid \Psi_1(x) + t \cdot V = z\right] \\
&\leqslant \mathbf{E}\left[\Psi_1(X^{t+1}) \mid \Psi_1(x) = z - t \cdot V\right] + (t+1) \cdot V \\
&\leqslant (z - tV - V) + (t+1)V = z.
\end{aligned}$$

Hence $\mathbf{E}\left[Z_{t+1}\right] = \sum_z \mathbf{E}\left[Z_{t+1} \,|\, Z_t = z\right] \cdot \mathbf{Pr}\left[Z_t = z\right] \leqslant \sum_z z \cdot \mathbf{Pr}\left[Z_t = z\right] = \mathbf{E}\left[Z_t\right]$. We obtain

$$\begin{aligned}
V \cdot E[T] &\leqslant E[\Psi(X^T)] + V \cdot E[T] \\
&= E[Z_T] \leqslant \ldots \leqslant E[Z_0] \\
&\leqslant 70 \cdot S \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^2 \ .
\end{aligned}$$

Therefore,

$$\begin{aligned}
E[T] &\leqslant 70 \cdot S \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^2 \cdot (1/V) \\
&\leqslant 70 \cdot S \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^2 \cdot (8\Delta s_{\max}^3) \\
&= 560 \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{\max}^3 \ .
\end{aligned}$$

□

Now we are ready to show the theorem. From Lemma 3.10 we know that there is an integer

$$T_1 \leqslant \gamma \cdot (2 \ln n + \ln \Psi_0(X^0))$$

such that

$$\mathbf{Pr}\left[\neg \exists t \leqslant T_1 : L_\Delta(X^t) \leqslant 8 \cdot \text{diam}(G) \cdot n \cdot \Delta\right] \leqslant 1/4 \ .$$

Let T_2 be the additional number of steps until X^{t+T_2} is a NE. From Lemma 3.14 we get

$$\mathbf{E}\left[T_2\right] \leqslant 560 \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{max}^3 \ ,$$

which implies

$$\mathbf{Pr}\left[T_2 \geqslant 4 \cdot 560 \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{max}^3\right] \leqslant 1/4 \ .$$

Hence, for

$$T = T_1 + 2240 \cdot (\text{diam}(G))^2 \cdot n^2 \cdot \Delta^3 \cdot S \cdot s_{max}^3 \ ,$$

the probability that X^T is not a NE is at most $2/4$. Now divide the time into intervals of length T. The expected number of these time intervals is at most $1 + 1/2 + (1/2)^2 + (1/2)^3 \ldots \leqslant 2$. In addition, after $c \cdot \log_2 n$ intervals the state X^t is a NE with probability at least $1 - 1/n^c$. This proves Theorem 3.2.

4 Uniform Tasks and Identical Machines

In this section we consider the case of identical machines. We use Protocol I with $s_{\max} = 1$ and $\alpha = 3$. For uniform tasks and identical machines, a task assigned to i prefers machine j if and only if $W(X_i^{t-1}) - W(X_j^{t-1}) > 1$. With μ_2 being the second smallest eigenvalue of the Laplace matrix, we show the following result.

Theorem 4.1. *Let $m \geqslant \delta n^5$ for some $\delta > 1$. With $\alpha = 3$, Protocol I reaches a $2/(1+\delta)$-approximate Nash equilibrium in expected time*

$$\mathcal{O}((\Delta/\mu_2) \cdot \ln m) \ .$$

Theorem 4.2. *With $\alpha = 3$, Protocol I reaches a Nash equilibrium in expected time*

$$\mathcal{O}\left((\Delta/\mu_2) \cdot (\ln m + \ln n) + (|E| \cdot \Delta)/\mu_2\right) \ .$$

4.1 Proof of Theorem 4.1

For a state x, we define $f_{i,j}(x)$ as in Section 3, i.e., as the expected load that is sent from i to j in one round of Protocol I. The next lemma bounds the potential change for the edges with a load difference of at most two.

Lemma 4.3. *Let μ_2 be the second smallest eigenvalue of the Laplace matrix and let $\alpha = 3$. Assume $\Psi_0(x) \geqslant 2|E|/\mu_2$, then*

$$\mathbf{E}\left[\Delta\Psi_0(X^t) \,|\, X^{t-1} = x\right] \geqslant \left(1 - \frac{\mu_2}{36\Delta}\right) \cdot \Psi_0(x) \ .$$

Proof. Consider Rosenthal's potential function, which simplifies to $\Phi_1(x) = \sum_{i\in V} W(x_i)\cdot(W(x_i)+1)$ since all $s_i = 1$. We obtain

$$\begin{aligned}\Phi_1(x) &= \sum_{i\in V} W(x_i) \cdot (W(x_i) + 1) = \sum_{i\in V} (W(x_i))^2 + \sum_{i\in V} W(x_i)\\ &= \Phi_0(x) + m\\ &= \Psi_0(x) + \frac{m^2}{n} + m.\end{aligned}$$

Hence, for any given state x functions $\Phi_1(x)$, $\Phi_0(x)$ and $\Psi_0(x)$ differ only by additive terms and $\Delta\Phi_1(X^t \,|\, X^{t-1} = x) = \Delta\Psi_0(X^t \,|\, X^{t-1} = x)$. Note that for uniform tasks and with $\alpha = 3$ we have

$$f_{i,j}(x) = \begin{cases} \frac{W(x_i) - W(x_j)}{6\cdot\deg(i,j)} & \text{if } W(x_i) - W(x_j) > 1\\ 0 & \text{otherwise.}\end{cases}$$

We also have

$$\Lambda^0_{i,j}(x) = 8 \cdot \deg(i,j) \cdot f_{i,j}(x) \ ,$$

and

$$\widetilde{E}(x) := \{(i,j) \in E : \quad W(x_i) - W(x_j) > 1\} \ ,$$

where we omit the argument x whenever possible. Applying Lemma 3.6 we get

$$\begin{aligned}\mathbf{E}\left[\Delta\Psi_0(X^t)\,|\,X^{t-1}=x\right] &\geqslant \sum_{(i,j)\in\widetilde{E}} f_{i,j}(x)\cdot\left(\Lambda^0_{i,j}(x)-2\right)\\ &= \sum_{\{i,j\}\in\widetilde{E}} \frac{W(x_i)-W(x_j)}{6\cdot d(i,j)}\cdot\left(8\cdot d(i,j)\cdot\frac{W(x_i)-W(x_j)}{6\cdot d(i,j)}\right)\\ &\quad -\sum_{\{i,j\}\in\widetilde{E}}\frac{W(x_i)-W(x_j)}{3\cdot d(i,j)}\\ &= \sum_{\{i,j\}\in\widetilde{E}}\frac{2\,(W(x_i)-W(x_j))^2}{9\cdot d(i,j)}-\sum_{\{i,j\}\in\widetilde{E}}\frac{W(x_i)-W(x_j)}{3\cdot d(i,j)}\end{aligned}$$

Now we apply $x\leqslant x^2/2$ for $x\geqslant 2$ to obtain

$$\begin{aligned}\mathbf{E}\left[\Delta\Psi_0(X^t)\,|\,X^{t-1}=x\right] &\geqslant \sum_{\{i,j\}\in\widetilde{E}}\frac{2\,(W(x_i)-W(x_j))^2}{9\cdot d(i,j)}-\sum_{\{i,j\}\in\widetilde{E}}\frac{\frac{1}{2}\cdot(W(x_i)-W(x_j))^2}{3\cdot d(i,j)}\\ &= \sum_{\{i,j\}\in\widetilde{E}}\frac{(W(x_i)-W(x_j))^2}{18\cdot d(i,j)}.\end{aligned}$$

We conclude from Lemma 4.3 that

$$\mathbf{E}\left[\Delta\Psi_0(X^t)\,|\,X^{t-1}=x\right]\geqslant\sum_{\{i,j\}\in\widetilde{E}}\frac{(W(x_i)-W(x_j))^2}{18\cdot d(i,j)}\geqslant\sum_{\{i,j\}\in\widetilde{E}}\frac{(W(x_i)-W(x_j))^2}{18\cdot\Delta}\ .$$

On the other hand, a classic result from standard (continuous) diffusion [7, Theorem 4] (see also [22, Proof of Theorem 1]) shows that

$$\sum_{\{i,j\}\in E}\frac{(W(x_i)-W(x_j))^2}{18\Delta}\geqslant\frac{\mu_2}{18\cdot\Delta}\cdot\Psi_0(x)\ ,$$

where we recall that μ_2 is the second smallest eigenvalue of the Laplace matrix. By noting that $(W(x_i)-W(x_j))^2\leqslant 1$ for all $e\in E\setminus\widetilde{E}$, we derive

$$\sum_{\{i,j\}\in\widetilde{E}}\frac{(W(x_i)-W(x_j))^2}{18\cdot\Delta}\geqslant\frac{\mu_2}{18\cdot\Delta}\cdot\Psi_0(x)-\frac{|E|}{18\cdot\Delta}\ .$$

Hence if $\mu_2\cdot\Psi_0(x)\geqslant 2|E|$, we get

$$\mathbf{E}\left[\Delta\Psi_0(X^t)\,|\,X^{t-1}=x\right]\geqslant\frac{\mu_2}{36\cdot\Delta}\cdot\Psi_0(x)\ .$$

This implies that as long as $\Psi_0(x)\geqslant 2|E|/\mu_2$,

$$\mathbf{E}\left[\Psi_0(X^t)\,|\,X^{t-1}=x\right]\leqslant\Psi_0(x)-\mathbf{E}\left[\Delta\Psi_0(X^t)\,|\,X^{t-1}=x\right]=\left(1-\frac{\mu_2}{36\Delta}\right)\cdot\Psi_0(x)\ .$$

This completes the proof. □

As in Lemma 3.10, we prove that:

Lemma 4.4. *Let* $\gamma := \frac{36\Delta}{\mu_2}$ *and* $\tau := \gamma \cdot (2 \ln n + \ln(\Phi_0(X^0)))$. *Then with probability at least* $1 - n^{-1}$, *there exists a round* $t \in [1, \tau]$ *such that*

$$\Phi_0(X^t) < 2|E|/\mu_2.$$

Proof. The proof is the same as Lemma 3.10, except that the conditions on $L_\Delta(X^t)$ are replaced by the conditions on $\Psi_0(X^t)$. □

Now we are ready to show Theorem 4.1.

Proof of Theorem 4.1. First we show that for $m \geqslant \delta n^5$, any state x with $\Psi_0(x) \leqslant 2|E|/\mu_2 \leqslant n^5$ is a $2/(1+\delta)$-apx. NE. Note that $2|E|/\mu_2 \leqslant 4|E|^2 \operatorname{diam}(G) \leqslant n^5$, using the fact that for any graph G, $\mu_2 \geqslant \frac{1}{\operatorname{diam}(G) \cdot 2|E|}$ ([11, Lemma 1.9]).

For the approximation ratio we note that, given $\Psi_0(x) \leqslant n^5$, the maximum load must satisfy

$$W_{\max}(x) = \max_{i \in V} W(x_i) \leqslant m/n + n^3.$$

The minimum load must satisfy

$$W_{\min}(x) = \min_{i \in V} W(x_i) \geqslant m/n - n^3.$$

Recall that in an ε-apx. NE, for all neighboring vertices i, j the following inequality must hold

$$(1 - \varepsilon) W(x_i) \leqslant W(x_j) + 1.$$

Thus, as we have $W(x_i) \leqslant m/n + n^3$ and $W(x_j) \geqslant \max\{0, m/n - n^3\} + 1$ for all edges $(i, j) \in E$ along which players want to move, it suffices for ε to satisfy $(1 - \varepsilon)(m/n + n^3) \leqslant m/n - n^3$. By solving this for ε and using $m \geqslant \delta n^5$, we get $\varepsilon \leqslant 1 - (\delta - 1)/(\delta + 1) = 2/(\delta + 1)$ as desired.

Using Lemma 4.4 along with the fact that $\Psi_0(X^0) \leqslant m^2$, it follows that with probability at least $1 - n^{-1}$ that there is a step $t \in [1, \tau]$ with $\Psi_0(X^t) \leqslant 2|E|/\mu_2 \leqslant n^5$ and $\tau \in \mathcal{O}((\Delta/\mu_2) \cdot (\ln n + \ln m))$ as in Lemma 4.4. Since this holds for an arbitray initial state at step 0 with m tasks, it follows that the expected time to reach such a state is at most $\mathcal{O}((\Delta/\mu_2) \cdot (\ln n + \ln m))$. With $m \in \Omega(n)$, the theorem follows. □

4.2 Proof of Theorem 4.2

The proof of Theorem 4.2 is very similar to the proof of Theorem 3.2. Using Lemma 4.4, we can bound the number of steps to reach a state τ with $\Phi_0(X^\tau) < n^5$. Therefore, it remains to consider the number of additional steps to reach a NE from such a state.

Lemma 4.5. *Assume that at step* t *the system is not in a Nash equilibrium. Then we have*

$$E[\Psi_0(X^{t+1}) \mid X^t = x] = \Psi_0(x) - \frac{1}{18\Delta} \ .$$

Proof. If in a step t we are not in a NE, there is an edge $\{i, j\} \in E$ with $|W(X_i^t) - W(X_j^t)| > 1$. Then Equation 4.2 of Lemma 4.3 yields

$$E[\Psi_0(X^{t+1}) \mid X^t = x] \leqslant \Psi_0(x) - \frac{1}{18\Delta}.$$

□

---

for each task ℓ in parallel **do**
  Let $i = i(\ell)$ the current machine of task ℓ and w_ℓ the weight of task ℓ
  Choose neighbor machine j u.a.r.
  if $W(X_i^{t-1}) - W(X_j^{t-1}) > w_\ell$ **then**
    Move task ℓ from machine i to j with probability

$$\frac{\deg(i)}{\deg(i,j)} \cdot \frac{W(X_i^{t-1}) - W(X_j^{t-1})}{2\alpha \cdot W_i^{t-1}}$$

  end if
end for

Figure 2: Protocol II for weighted tasks and uniform speeds.

Lemma 4.6. *Assume that at step τ we have $\Phi_0(X^\tau) \leqslant 2|E|/\mu_2$. Let T be the additional number of steps until $X^{\tau+T}$ is a Nash equilibrium. Then*

$$\mathbf{E}[T] \leqslant 36|E|/\mu_2 \cdot \Delta \ .$$

Similarly to Theorem 3.2, we are now able to prove Theorem 4.2 by combining Lemma 4.4 (to reach a state τ with $\Phi_0(X^\tau) \leqslant n^5$) and Lemma 4.6 (to reach from there a NE). By Markov's inequality and the union bound, with probability at least 1/2, the number of total steps needed for this is at most

$$\frac{36}{\mu_2} \cdot (2\ln n + \ln(\Phi_0(X^0))) + 54\frac{|E| \cdot \Delta}{\mu_2}.$$

Since for any state x, $\Phi_0(x) \leqslant m^2$, it follows that the total number of expected steps for Protocol I to reach a NE is at most

$$\mathcal{O}\left(\frac{1}{\mu_2} \cdot (\ln n + \ln m) + \frac{|E| \cdot \Delta}{\mu_2}\right).$$

□

5 Weighted Tasks and Identical Machines

In this section we assume that every task $\ell \in [m]$ has an integral weight $w_\ell \geqslant 1$.

5.1 Protocol and Results

We denote the maximal and minimal weight of any task by $w_{\max}$ and $w_{\min}$, respectively. A task ℓ with weight $w_\ell \in \mathbb{N}$ that is located at machine i in state x prefers machine j over i if and only if $W(x_i) > W(x_j) + w_\ell$, which is equivalent to $W(x_i) - W(x_j) > w_x$. Our protocol for weighted tasks is given in Figure 2.

Theorem 5.1. *With $\alpha = 4w_{\max}$, Protocol II reaches a Nash equilibrium in expected time $\mathcal{O}(\Delta \cdot W^3 \cdot w_{\max})$.*

5.2 Proof of Theorem 5.1

Let x be an arbitrary state. We define for every task ℓ an indicator variable $H^{\ell}_{i,j}(x) \in \{0,1\}$ (H for happy) which is one if task ℓ prefers machine j over its current location i, and zero otherwise. $H_{i,j}(x)$ is defined as the fraction of the total weight of jobs on machine i that prefer machine j over i, i.e.,

$$H_{i,j}(x) := \sum_{\ell \in x_i :\, H^{\ell}_{i,j}(x)=1} \frac{w_\ell}{W(x_i)}.$$

Note that for uniform weights, $H_{i,j}(x) \in \{0,1\}$, i.e., either all tasks or no task on i prefer j over i. For non-uniform weights however, $H_{i,j}(x)$ may potentially be as small as $w_{\min}/W$ (apart from being 0). For any $j \in N(i)$ at least one of the variables $H_{i,j}(x)$ or $H_{j,i}(x)$ must be zero. With these definitions we get

$$\begin{aligned}
\mathbf{E}\left[W(X_i^t) \,|\, X^{t-1} = x\right] &= W(x_i) - \sum_{j \in N(i): W(x_i) > W(x_j)} \; \sum_{\ell \in x_i :\, H^{\ell}_{i,j}(x)=1} w_\ell \cdot \mathbf{Pr}\left[r \text{ moves to } j\right] \\
&\quad + \sum_{j \in N(i): W(x_i) < W(X_i)} \; \sum_{\ell \in x_j :\, H^{\ell}_{j,i}(x)=1} w_\ell \cdot \mathbf{Pr}\left[\ell \text{ moves to } i\right] \\
&= W(x_i) - \sum_{j \in N(i): W(x_i) > W(x_j)} H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{2\alpha \cdot d(i,j)} \\
&\quad + \sum_{j \in N(i): W(x_i) < W(x_j)} H_{j,i}(x) \cdot \frac{W(x_j) - W(x_i)}{2\alpha \cdot d(i,j)}.
\end{aligned}$$

Again, we denote by $\widetilde{E}(x)$ the set of edges (i,j) with $H_{i,j}(x) > 0$ (and omit x whenever possible). Then

$$\mathbf{E}\left[W_i(X^t) \,|\, X^{t-1} = x\right] = W(x_i) - \sum_{j:(i,j) \in \widetilde{E}} H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{2\alpha \cdot d(i,j)} + \sum_{j:(j,i) \in \widetilde{E}} H_{j,i}(x) \cdot \frac{W(x_j) - W(x_i)}{2\alpha \cdot d(i,j)}.$$

So the expected flow from i to j, $\{i,j\} \in E$, in a single round of the protocol starting from state x is

$$f_{i,j}(x) := H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{2\alpha \cdot d(i,j)} \quad .$$

Note that at least one of the flows, $f_{i,j}(x)$ or $f_{j,i}(x)$ has to be zero. Further, observe that $(i,j) \in \tilde{E}(x)$ if and only if $f_{i,j}(x) > 0$.

Let us first consider the potential change $\widetilde{\Delta}\Phi(X^t \,|\, X^{t-1} = x)$ (recall, this is the change under the assumption that the flow on every edge is exactly its expected value).

$$\widetilde{\Delta}\Phi_0(X^t \,|\, X^{t-1} = x) := \sum_{i \in V} (W(x_i))^2 - \sum_{i \in V} (\mathbf{E}\left[W(X_i^t)\right])^2.$$

The following lemma generalizes [7, Lemma 2] to the setting with weights. The statement and the proof is similar to Lemma 3.4.

Lemma 5.2. *For any step $t \in \mathbb{N}$ and any state x it holds that*

$$\widetilde{\Delta}\Phi_0(X^t \,|\, X^{t-1} = x) \geqslant \sum_{(i,j) \in \widetilde{E}} f_{i,j}(x) \cdot \left(\left(\frac{\alpha d(i,j)}{H_{i,j}(x)} - d(i,j) \right) \cdot 4 \cdot f_{i,j}(x) \right).$$

Proof. We follow the same proof as for Lemma 3.4. Recall that in the proof of Lemma 3.4, we did not use the particular definition of $f_{i,j}(x)$. Hence we can use the same notations and definitions as in the proof of Lemma 3.4 to obtain a simplified version of equation (3.1) for weighted tasks but with uniform speeds, i.e.,

$$\widetilde{\Delta}\Phi_0(X^{(t,k)}) \geqslant f_{i,j}(x) \cdot (2 \cdot (W(x_i) - W(x_j)) - 2 \cdot (2d(i,j) - 1)\, f_{i,j}(x))\,,$$

where $e_k = (i,j)$ represents the k-th activation of an edge in $\tilde{E}(x)$. Plugging in the definition of $f_{i,j}(x)$, we get

$$\begin{aligned} \widetilde{\Delta}\Phi_0(X^{(t,k)}) &\geqslant f_{i,j}(x) \cdot \left(\frac{4\alpha \cdot d(i,j)}{H_{i,j}(x)} \cdot f_{i,j}(x) - 2 \cdot (2d(i,j) - 1) \cdot f_{i,j}(x) \right) \\ &\geqslant f_{i,j}(x) \cdot \left(\left(\frac{\alpha \cdot d(i,j)}{H_{i,j}(x)} - d(i,j) \right) \cdot 4 \cdot f_{i,j}(x) \right). \end{aligned}$$

Summing over all $(i,j) \in \widetilde{E}$ yields the result. □

Now we bound the sum of the variances of the random variables $W(X_i^t)$ conditioned on $X^{t-1} = x$ for any x. We define $w_{i,j}(x)$ as the maximum weight of a task that prefers (in state x) either j over i or i over j. We define $w_{i,j}(x) = 0$ if there is no such task.

Lemma 5.3. *For any step $t \in \mathbb{N}$ and state x it holds that,*

$$\sum_{i \in V} \mathbf{Var}\left[W(X_i^t \,|\, X^{t-1} = x)\right] = \sum_{(i,j) \in \widetilde{E}} w_{i,j}(x) \cdot H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha \cdot d(i,j)}$$

Proof. We have

$$\sum_{i \in V} \mathbf{Var}\left[W(X^t) \,|\, X^{t-1} = x\right] = \sum_{i \in V} \mathbf{Var}\left[C_i^t - A_i^t\right],$$

where the random variable C_i^t counts the number of tasks that come to link i from another link in step t, and the variable A_i^t counts the number of tasks that abandon link i in step t. Again we have

$$\mathbf{Var}\left[C_i^t - A_i^t\right] = \mathbf{Var}\left[C_i^t\right] + \mathbf{Var}\left[A_i^t\right].$$

Observe that

$$C_i^t = \sum_{j \in N(i)} \sum_{\ell \in x_j :\, H_{j,i}^{\ell}(x) = 1} w_\ell \cdot Y_{\ell,j}^t,$$

where $Y_{\ell,j}^t$ is the Bernoulli random variable indicating whether task ℓ jumps to j at step t. Note that

$$Y_{\ell,j}^t \sim \mathsf{Ber}\left(\frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_j) - W(x_i)}{2 \cdot W(x_j)} \right).$$

Similarly,

$$A_i^t = \sum_{\ell \in x_i} w_\ell \cdot Z_{\ell,i}^t,$$

where $Z^t_{\ell,i}$ is the Bernoulli random variable indicating whether a task ℓ located at i before step t moves to another vertex. Again, we have

$$Z^t_{\ell,i} \sim \mathsf{Ber}\left(\frac{1}{\alpha \cdot d(i,j)} \cdot \sum_{j \in N(i): H^\ell_{i,j}(x)=1} \frac{W(x_i) - W(x_j)}{2 \cdot W(x_i)} \right).$$

By independence,

$$\begin{aligned}
\mathbf{Var}\left[C^t_i\right] &= \sum_{j \in N(i)} \sum_{\ell \in x_j:\, H^\ell_{j,i}(x)=1} w_\ell^2 \cdot \mathbf{Var}\left[\mathsf{Ber}\left(\frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_j) - W(x_i)}{2 \cdot W(x_j)} \right)\right] \\
&\leqslant \sum_{j \in N(i)} w_{i,j}(x) \cdot \sum_{\ell \in x_j:\, H^\ell_{j,i}(x)=1} w_\ell \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_j) - W(x_i)}{2 \cdot W(x_j)},
\end{aligned}$$

Similarly,

$$\begin{aligned}
\mathbf{Var}\left[A^t_i\right] &= \sum_{\ell \in x_i} w_\ell^2 \cdot \mathbf{Var}\left[\mathsf{Ber}\left(\frac{1}{\alpha \cdot (i,j)} \cdot \sum_{j \in N(i): H^\ell_{i,j}(x)=1} \frac{W(x_i) - W(x_j)}{2 \cdot W(x_i)} \right)\right] \\
&\leqslant \sum_{\ell \in x_i} \sum_{j \in N(i): H^\ell_{i,j}(x)=1} w_\ell^2 \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_i) - W(x_j)}{2 \cdot W(x_i)} \\
&= \sum_{j \in N(i)} \sum_{\ell \in x_i: H^\ell_{i,j}(x)=1} w_\ell^2 \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_i) - W(x_j)}{2 \cdot W(x_i)} \\
&\leqslant \sum_{j \in N(i)} w_{i,j}(x) \cdot \sum_{\ell \in x_i: H^\ell_{i,j}(x)=1} w_\ell \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_i) - W(x_j)}{2 \cdot W(x_i)}.
\end{aligned}$$

This gives

$$\begin{aligned}
\sum_{i \in V} \mathbf{Var}\left[W(X^t \,|\, X^{t-1} = x)\right] &= \sum_{(i,j) \in \widetilde{E}} w_{i,j}(x) \cdot 2 \cdot \sum_{\ell \in x_i:\, H^\ell_{i,j}(x)=1} w_\ell \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_i) - W(x_j)}{2 \cdot W(x_i)} \\
&= \sum_{(i,j) \in \widetilde{E}} w_{i,j}(x) \sum_{\ell \in x_i:\, H^\ell_{i,j}(x)=1} w_\ell \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \frac{W(x_i) - W(x_j)}{W(x_i)} \\
&= \sum_{(i,j) \in \widetilde{E}} w_{i,j}(x) \cdot H_{i,j}(x) \cdot \frac{1}{\alpha \cdot d(i,j)} \cdot \left(W(x_i) - W(x_j)\right).
\end{aligned}$$

□

The following lemma is similar to Lemma 3.6 (where we have different speeds but uniform weights).

Lemma 5.4. *For any step t and any state x it holds that*

$$\mathbf{E}\left[\Delta\Phi_0(X^t) \,|\, X^{t-1} = x\right] \geqslant \sum_{(i,j) \in \widetilde{E}} \left(1 - \frac{(w_{i,j}(x) + 1) \cdot H_{i,j}(x)}{\alpha}\right) \cdot H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha \cdot d(i,j)}.$$

Proof. From the definition of Φ_0 we get

$$\begin{aligned}\mathbf{E}\left[\Phi_0(X^t) \mid X^{t-1} = x\right] &= \sum_{i \in V} \mathbf{E}\left[(W(X_i^t))^2 \mid X^{t-1} = x\right] \\ &= \sum_{i \in V} \mathbf{Var}\left[W(X_i^t) \mid X^{t-1} = x\right] + \left(\mathbf{E}\left[W(X_i^t) \mid X^{t-1} = x\right]\right)^2.\end{aligned}$$

Therefore,

$$\begin{aligned}&\mathbf{E}\left[\Delta\Phi_0(X^t) \mid X^{t-1} = x\right] \\ &= \sum_{i \in V}(W(x_i))^2 - \sum_{i \in V}\left(\mathbf{E}\left[W(X_i^t) \mid X^{t-1} = x\right]\right)^2 - \sum_{i \in V} \mathbf{Var}\left[W(X_i^t) \mid X^{t-1} = x\right].\end{aligned}$$

We first lower bound the first two sums. Using Lemma 5.2 we get

$$\begin{aligned}&\sum_{i \in V}(W(x_i))^2 - \sum_{i \in V}\left(\mathbf{E}\left[W(X_i^t) \mid X^{t-1} = x\right]\right)^2 \\ &\geqslant \sum_{(i,j) \in \widetilde{E}} f_{i,j}(x) \cdot \left(\left(\frac{\alpha \cdot d(i,j)}{H_{i,j}(x)} - d(i,j)\right) \cdot 4 \cdot f_{i,j}(x)\right) \\ &= \sum_{(i,j) \in \widetilde{E}} H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{2\alpha \cdot d(i,j)} \cdot \left(\left(\frac{\alpha \cdot d(i,j)}{H_{i,j}(x)} - d(i,j)\right) \cdot 4 \cdot f_{i,j}(x)\right) \\ &= \sum_{(i,j) \in \widetilde{E}} H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha} \cdot \left(\left(\frac{\alpha}{H_{i,j}(x)} - 1\right) \cdot 2 \cdot f_{i,j}(x)\right)\end{aligned}$$

Now we define

$$A(x) := \left(\frac{\alpha}{H_{i,j}(x)} - 1\right) \cdot 2 \cdot f_{i,j}(x) \cdot d(i,j) \quad .$$

Then by the above and using Lemma 5.3 to bound the sum of variances, we arrive at

$$\begin{aligned}\mathbf{E}\left[\Delta\Phi_0(X^t) \mid X^{t-1} = x\right] &\geqslant \sum_{(i,j) \in \widetilde{E}} H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha} \cdot \left(\frac{\alpha}{H_{i,j}(x)} - 1\right) \cdot 2 \cdot f_{i,j}(x) \\ &\quad - \sum_{(i,j) \in \widetilde{E}} w_{i,j}(x) \cdot H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha \cdot d(i,j)} \\ &= \sum_{(i,j) \in \widetilde{E}} (A(x) - w_{i,j}(x)) \cdot H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha \cdot d(i,j)}.\end{aligned}$$

We lower bound $A(x)$ as follows:

$$\begin{aligned}A(x) &= \left(\frac{\alpha}{H_{i,j}(x)} - 1\right) \cdot 2 \cdot f_{i,j}(x) \cdot d(i,j) \\ &= \left(\frac{\alpha}{H_{i,j}(x)} - 1\right) \cdot 2 \cdot H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{2\alpha} \\ &\geqslant \left(\frac{1}{H_{i,j}(x)} - \frac{1}{\alpha}\right) \cdot H_{i,j}(x) \cdot (w_{i,j}(x) + 1) \\ &= \left(1 - \frac{H_{i,j}(x)}{\alpha}\right) \cdot (w_{i,j}(x) + 1) \\ &= w_{i,j}(x) + 1 - \frac{(w_{i,j}(x) + 1) \cdot H_{i,j}(x)}{\alpha} \quad .\end{aligned}$$

Hence

$$\mathbf{E}\left[\Delta\Phi_0(X^t) \mid X^{t-1} = x\right] \geqslant \sum_{\{i,j\} \in \widetilde{E}} \left(1 - \frac{(w_{i,j}(x)+1) \cdot H_{i,j}(x)}{\alpha}\right) \cdot H_{i,j}(x) \cdot \frac{W(x_i) - W(x_j)}{\alpha \cdot d(i,j)}.$$

□

With this lemma at hand, it is easy to conclude the proof of Theorem 5.1.

Proof of Theorem 5.1. Assume there are tasks on machine i that prefer machine j over i. Then, since $H_{i,j}(x) \leqslant 1$, $w_{i,j}(x) \leqslant w_{\max}$ and $\alpha \geqslant 4w_{\max}$, we have

$$\frac{(w_{i,j}(x)+1) \cdot H_{i,j}(x)}{\alpha} \leqslant \frac{w_{\max}+1}{4w_{\max}} \leqslant \frac{1}{2} \ .$$

We also have $H_{i,j}(x) \geqslant w_{\min}/W$ if $H_{i,j}(x) > 0$. Hence as long as we have not reached a NE, Lemma 5.4 implies that

$$\mathbf{E}\left[\Delta\Phi_0(X^t) \mid X^{t-1} = x\right] \geqslant \frac{1}{2} \cdot \frac{w_{\min}}{W} \cdot \frac{w_{\min}}{4w_{\max} \cdot \Delta} = \frac{1}{8} \cdot \frac{(w_{\min})^2}{\Delta \cdot W \cdot w_{\max}} \ .$$

Next observe that for any state x, we have $\Phi_0(x) \leqslant W^2$, where W is the sum of the weights of all tasks. Hence as in the proof of Theorem 3.2, we obtain the following upper bound on the expected time to reach a NE:

$$W^2 \cdot 8 \cdot \frac{\Delta \cdot W \cdot w_{\max}}{(w_{\min})^2} \leqslant 8 \cdot \Delta \cdot W^3 \cdot w_{\max},$$

where the last inequality holds, since $w_{\min} \geqslant 1$. This completes the proof of Theorem 5.1. □

6 Conclusions

In this paper we initiate the study of concurrent protocols for selfish load balancing on general networks. Our protocols rely only on local information and computation and yield rapid convergence times to approximate and exact NE for systems with many agents. We show a number of generalisations e.g., to networks with uniformly related machines of different speeds or agents with weights.

There are many open problems that arise from our work, such as finding improved convergence times or general lower bounds for concurrent protocols. On a more technical side, the generalization of our techniques is an interesting open problem, e.g., to more general potential functions that work for networks with both speeds and weights [14]. Another interesting problem is whether approaches that do not use potential functions (e.g., [4]) can be applied here.

References

[1] Heiner Ackermann, Petra Berenbrink, Simon Fischer, and Martin Hoefer. Concurrent imitation dynamics in congestion games. In *Proc. 28th Symp. Principles of Distributed Computing (PODC)*, pages 63–72, 2009.

[2] Heiner Ackermann, Simon Fischer, Martin Hoefer, and Marcel Schöngens. Distributed algorithms for QoS load balancing. *Distributed Computing*, 23(5–6):321–330, 2011.

[3] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6), 2008.

[4] Baruch Awerbuch, Yossi Azar, and Rohit Khandekar. Fast load balancing via bounded best response. In *Proc. 19th Symp. Discrete Algorithms (SODA)*, pages 314–322, 2008.

[5] 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.

[6] 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.

[7] Petra Berenbrink, Tom Friedetzky, and Zengjian Hu. A new analytical method for parallel, diffusion-type load balancing. *J. Parallel and Distributed Comput.*, 69:54–61, 2009.

[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] Jacques Boillat. Load balancing and poisson equation in a graph. *Concurrency: Pract. Exper.*, 2(4):289–313, 1990.

[10] Steve Chien and Alistair Sinclair. Convergence to approximate Nash equilibria in congestion games. *Games Econom. Behav.*, 71(2):315–327, 2011.

[11] Fan Chung. *Spectral Graph Theory.* Number 92 in CBMS Regional Conference Series in Mathematics. Amer. Math. Society, 1997.

[12] George Cybenko. Dynamic load balancing for distributed memory multiprocessors. *J. Parallel and Distributed Comput.*, 7:279–301, 1989.

[13] Robert Elsässer, Burkhard Monien, and Stefan Schamberger. Distributing unit size workload packages in heterogeneous networks. *J. Graph Alg. Appl.*, 10(1):51–68, 2006.

[14] Eyal Even-Dar, Alexander Kesselman, and Yishay Mansour. Convergence time to Nash equilibria in load balancing. *ACM Trans. Algorithms*, 3(3), 2007.

[15] Eyal Even-Dar and Yishay Mansour. Fast convergence of selfish rerouting. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 772–781, 2005.

[16] Rainer Feldmann, Martin Gairing, Thomas Lücking, Burkhard Monien, and Manuel Rode. Nashification and the coordination ratio for a selfish routing game. In *Proc. 30th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 514–526, 2003.

[17] 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.

[18] Simon Fischer, Harald Räcke, and Berthold Vöcking. Fast convergence to Wardrop equilibria by adaptive sampling mehtods. *SIAM J. Comput.*, 39(8):3700–3735, 2010.

[19] Simon Fischer and Berthold Vöcking. Adaptive routing with stale information. *Theoret. Comput. Sci.*, 410(36):3357–3371, 2008.

[20] Dimitris Fotakis, Alexis Kaporis, and Paul Spirakis. Atomic congestion games: Fast, myopic and concurrent. *Theory Comput. Syst.*, 47(1):38–49, 2010.

[21] Tobias Friedrich and Thomas Sauerwald. Near-perfect load balancing by randomized rounding. In *Proc. 41st Symp. Theory of Computing (STOC)*, pages 121–130, 2009.

[22] Bhaskar Ghosh and S. Muthukrishnan. Dynamic Load Balancing by Random Matchings. *J. Comput. Syst. Sci.*, 53(3):357–370, 1996.

[23] 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.

[24] Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24(1):21–29, 2011.

[25] Jason Marden, Gürdal Arslan, and Jeff Shamma. Regret based dynamics: Convergence in weakly acyclic games. In *Proc. 6th Conf. Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2007.

[26] Jason Marden, Peyton Young, Gürdal Arslan, and Jeff Shamma. Payoff-based dynamics for multi-player weakly acyclic games. *SIAM J. Control Opt.*, 48(1):373–396, 2009.

[27] S. Muthukrishnan, Bhaskar Ghosh, and Martin Schultz. First- and second-order diffusive methods for rapid, coarse, distributed load balancing. *Theory Comput. Syst.*, 31(4):331–354, 1998.

[28] Yuval Rabani, Alistair Sinclair, and Rolf Wanka. Local divergence of Markov chains and the analysis of iterative load balancing schemes. In *Proc. 39th Symp. Foundations of Computer Science (FOCS)*, pages 694–705, 1998.

[29] Robert Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory*, 2:65–67, 1973.

[30] Alexander Skopalik and Berthold Vöcking. Inapproximability of pure Nash equilibria. In *Proc. 40th Symp. Theory of Computing (STOC)*, pages 355–364, 2008.

[31] 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.