

DISTRIBUTED SELFISH LOAD BALANCING*

PETRA BERENBRINK[†], TOM FRIEDETZKY[‡], LESLIE ANN GOLDBERG[§],
PAUL W. GOLDBERG[§], ZENGJIAN HU[†], AND RUSSELL MARTIN[§]

Abstract. Suppose that a set of m tasks are to be shared as equally as possible among a set of n resources. A game-theoretic mechanism to find a suitable allocation is to associate each task with a “selfish agent” and require each agent to select a resource, with the cost of a resource being the number of agents that select it. Agents would then be expected to migrate from overloaded to underloaded resources, until the allocation becomes balanced. Recent work has studied the question of how this can take place within a distributed setting in which agents migrate selfishly without any centralized control. In this paper we discuss a natural protocol for the agents which combines the following desirable features: It can be implemented in a strongly distributed setting, uses no central control, and has good convergence properties. For $m \gg n$, the system becomes approximately balanced (an ϵ -Nash equilibrium) in expected time $O(\log \log m)$. We show using a martingale technique that the process converges to a perfectly balanced allocation in expected time $O(\log \log m + n^4)$. We also give a lower bound of $\Omega(\max\{\log \log m, n\})$ for the convergence time.

Key words. load balancing, reallocation, equilibrium, convergence

AMS subject classifications. 68Q25, 68W20, 68W15, 91A80

DOI. 10.1137/060660345

1. Introduction. Suppose a consumer learns the price she would be charged by some domestic power supplier other than the one she is currently using. It is plausible that if the alternative price is lower than the price she is currently paying, then there is some possibility that she will switch to the new power supplier. Furthermore, she is more likely to switch if the ratio of the current price to the new price is large. If there is only a small savings, then it becomes unattractive to make the switch, since an influx of new business (hers and that of other consumers) may drive up the price of the new power supplier and make it no longer competitive.

We study a simple mathematical model of the above natural rule, in the context of a load balancing (or task allocation) scenario that has received a lot of recent attention. We assume the presence of many individual users who may assign their tasks to chosen resources. The users are *selfish* in the sense that they attempt to optimize their own situation, i.e., try to assign their tasks to minimally loaded resources, without trying to optimize the global situation. In general, a *Nash equilibrium* (NE) among a set of selfish users is a state in which no user has the incentive to change her current decision. In our setting, this corresponds to no user having an incentive to reallocate

*Received by the editors May 19, 2006; accepted for publication (in revised form) April 23, 2007; published electronically November 21, 2007. A preliminary version appeared in *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006, pp. 354–363. This work was partially supported by the EPSRC grants “Discontinuous Behaviour in the Complexity of Randomized Algorithms” and “Algorithmics of Network-sharing Games,” and by the Natural Sciences and Engineering Research Council of Canada (NSERC) discovery grant 250284-2002.

<http://www.siam.org/journals/sicomp/37-4/66034.html>

[†]School of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby V5A 1S6, BC, Canada (petra@cs.sfu.ca, zhu@cs.sfu.ca).

[‡]Department of Computer Science, Durham University, Science Labs, South Road, Durham, DH1 3LE, UK (tom.friedetzky@dur.ac.uk).

[§]Department of Computer Science, University of Liverpool, Ashton Building, Liverpool, L69 3BX, UK (l.a.goldberg@liverpool.ac.uk, pwg@csc.liv.ac.uk, r.martin@csc.liv.ac.uk).

her task to some other resource. An ϵ -Nash equilibrium (ϵ -NE) is a standard notion of an approximate NE, and is a state in which no user can reduce her cost by a multiplicative factor of less than $1 - \epsilon$ by changing action. Here we do not focus on the *quality* of equilibria but rather on the (perhaps more algorithmic) question of convergence time to such a state.

We assume a strongly distributed and concurrent setting; i.e., there is no centralized control mechanism whatsoever, and all users may choose to reallocate their tasks at the same time. Thus, we do not (and cannot) use the *elementary step system* [25] (discussed in more detail in the next section), where the assumption is that at most one user may reallocate her task at any given stage.

Throughout we let m denote the number of tasks (in the above discussion, customers) and n denote the number of resources (power suppliers). As hinted at in the above discussion, we assume that typically $m \gg n$. In a single time step (or round) each task does the following: Let i be the resource currently being used by the task. Select j uniformly at random from $\{1, \dots, n\}$ and find the load of resource j . Let X_i and X_j be the loads of resources i and j , respectively. If $X_j < X_i$, migrate from i to j with a probability of $1 - X_j/X_i$; the transition from round t to round $t + 1$ is given in Figure 1.1. Notice that if we had unconditional migrations, i.e., without an additional coin flip (move only with probability $1 - X_j(t)/X_i(t)$), then this might lead to an unstable system; consider, for example, the case $m = 2$ with initially most tasks assigned to one of the resources. The overload would oscillate between the two resources, with a load ratio tending towards 2:1. (This observation about the risk of oscillation has also been made in similar contexts in [12, 11], and we will not elaborate on it further.)

For each task b **do** in parallel
 Let i_b be the current resource of task b
 Choose resource j_b uniformly at random
 Let $X_{i_b}(t)$ be the current load of resource i
 Let $X_{j_b}(t)$ be the current load of resource j
If $X_{i_b}(t) > X_{j_b}(t)$ **then**
 Move task b from resource i_b to j_b with probability $1 - X_{j_b}(t)/X_{i_b}(t)$

FIG. 1.1. The protocol with “neutral moves” allowed.

It can easily be seen that if all tasks use the above policy, then the expected load of every resource at the next step is m/n .

OBSERVATION 1.1. *Regardless of the load distribution at time step t , the expected load of every resource at the next step is m/n .*

Proof. To see this, assume that the loads $X_i(t)$ are arranged in descending order so that $X_j(t) \geq X_{j+1}(t)$ and note that

$$\begin{aligned} \mathbb{E}[X_i(t+1)] &= X_i(t) + \sum_{\ell=1}^{i-1} \frac{1}{n} X_\ell(t) \left(1 - \frac{X_i(t)}{X_\ell(t)}\right) - \sum_{\ell=i+1}^n \frac{1}{n} X_i(t) \left(1 - \frac{X_\ell(t)}{X_i(t)}\right) \\ &= X_i(t) + \frac{1}{n} \sum_{\ell=1}^{i-1} (X_\ell(t) - X_i(t)) - \frac{1}{n} \sum_{\ell=i+1}^n (X_i(t) - X_\ell(t)) \\ &= X_i(t) + \frac{1}{n} \sum_{\ell=1}^n (X_\ell(t) - X_i(t)) = \frac{1}{n} \sum_{\ell=1}^n X_\ell(t) = \frac{m}{n}. \quad \square \end{aligned}$$

This provides a compelling motivation for the policy, which is that as a result, no task has an incentive to deviate unilaterally from this policy. This implies that in the terminology of [8] it is a *Nash rerouting policy*. It is also a simple regret-minimizing policy in the sense of [2] since the average cost of resources used by an agent is no higher than the best choice of a single resource to be used repeatedly. Although the above rule is very natural and has the nice properties described above, we show that it may take a long time to converge to a perfectly balanced allocation of tasks to resources. We address this problem as follows. Define a *neutral move* to be a task migration from a resource with load ℓ at time t to a resource with load $\ell - 1$ at time t (so, if no other task migrates, then the cost to the migrating task is unchanged). We consider a modification in which neutral moves are specifically disallowed (see Figure 2.1). That seemingly minor change ensures fast convergence from an almost balanced state to a perfectly balanced state. To summarize, here are the most important features of the modified protocol:

- We do not need any global information whatsoever (apart from the number of available resources); in particular, a task does not need to know the total number of tasks in the system. Also, it is strongly distributed and concurrent. If additional tasks were to enter the system, it would rapidly converge once again, with no outside intervention.
- A migrating task needs to query the load of only one other resource (thus, doing a constant amount of work in each round).
- When a task finds a resource with a significantly smaller load (that is, a load that is smaller by at least two), the migration policy is *exactly* the same as that used by the Nash rerouting policy of Figure 1.1, so the incentive is to use that probability.
- When a task finds a resource with a load that is smaller by exactly one unit, the migration policy is sufficiently close to the Nash rerouting policy that the difference in expected load is at most one, and there is little incentive to deviate.
- The protocol is simple (as well as provably efficient) enough to convince users to actually stick to it.

1.1. Related work. We are studying a simple kind of *congestion game*. In their general form, congestion games specify a set of agents, a set of resources, and, for each agent, a set of allowed strategies, where a strategy is the selection of a subset of the resources (in this paper, any singleton subset is allowed). The cost of a resource is a nondecreasing function of the number of agents using it, and the cost for an agent is the sum of the costs of resources it uses. A classical result due to Rosenthal [26] is that pure Nash equilibria (NEs) always exist for congestion games, and this is shown by exhibiting a potential function; they are a type of *potential game* [24]. The potential function also establishes that pure NEs can be found via sequences of “better-response” moves, in which agents repeatedly switch to lower-cost strategies. The potential function we use later in this paper is that of [26], modulo a linear rescaling.

These results do not show how to find NEs efficiently, the problem being that in the worst case, sequences of these self-improving moves may be exponentially long. The following questions arise: When can NEs be found by any efficient algorithm, and if so, can they be found via an algorithm that purports to be a realistic model of agents’ behavior? Regarding the first of these questions, the answer is no in the general setting (the problem is PLS-complete for general congestion games [9]; see

also [1, 3]). PLS-completeness (introduced in [17]) is a generally accepted criterion for intractability of computational problems in which we seek a local optimum of a given objective function.

However, due to the basic fact of [26, 24] that pure NEs are sure to result from a sufficiently long better-response sequence, many algorithms for finding them are based on such sequences. An important subclass is the *elementary step system* (ESS), proposed in Orda, Rom, and Shimkin [25], which consists of best-response moves (where a migrating agent switches not to any improved choice but to one that is optimal at the time of migration). For matroid games (a class of congestion games that includes the ones we consider here), Ackermann, Röglin, and Vöcking [1] show that best-response sequences must have length polynomial in the number of players, resources, and maximal rank of the matroids. In this paper we consider the special case of singleton congestion games (where players' strategies are always single resources, and thus the ranks of the matroids is 1). For these games, Jeong et al. [16] give polynomial bounds for best-response and better-response sequences. Chien and Sinclair [3] study a version of the ESS in the context of approximate NEs, and show that in some cases the ϵ -Nash dynamics may find an ϵ -NE where finding an exact NE is PLS-complete. Mirrokni and Vetta [22] study the convergence rate of the ESS to solutions, and the quality of the approximation after limited iterations.

While best- and better-response dynamics are a plausible model of selfish behavior, the associated algorithms typically require that migrations be done one-by-one, and another common assumption is that best- (not better-) responses are always selected. This means that to some extent, agents are being assumed to be governed by a centralized algorithm that finds an NE, raising the question of what sort of *distributed* algorithms can do so, especially if agents have limited information about the state of the system (and so may not be able to find best responses). That issue is of central importance to us in this paper. Goldberg [14] studied situations where simple better-response approaches can be realized as weakly distributed algorithms (where each agent looks for moves independently of the others, but it is assumed that moves take place consecutively, not simultaneously). In a strongly distributed setting (as we study here), where moves may occur simultaneously, we need to address the possibility that a change of strategy may increase an agent's cost. It may happen that after a best response has been identified, it is not optimal at the time it is executed. Even-Dar and Mansour [8] consider concurrent, independent rerouting decisions where tasks are allowed to migrate from overloaded to underloaded resources. Their rerouting process terminates in expected $O(\log \log m + \log n)$ rounds when the system reaches an NE. Note that their convergence rate as a function of the number n of resources is faster than the one we obtain in this paper. The reason is that it requires agents to have a certain amount of global knowledge. A task is required to know whether its resource is overloaded (having above-average load), and tasks on underloaded resources do not migrate at all. Our rerouting policy does not require that agents know anything other than their current resource load and the load of a randomly chosen alternative. Even-Dar and Mansour also present a general framework that can be used to show a logarithmic convergence rate for a wide class of rerouting strategies. Our protocol does not fall into that class, since we do not require migrations to occur only from overloaded resources. Note that our lower bound is linear in n (thus, more than logarithmic).

Distributed algorithms have been studied in the Wardrop setting (the limit of infinitely many agents), for which recent work has also extensively studied the coordination ratio [28, 27]. Fischer, Räcke, and Vöcking [11] investigate convergence to

Wardrop equilibria for games where agents select paths through a shared network to route their traffic. (Singleton games correspond to a network of parallel links.) Their rerouting strategies are slightly different to ours—they assume that in each round, an agent queries a path with probability proportional to the traffic on that path. Here we assume that paths (individual elements of a set of parallel links) are queried uniformly at random, so that agents can be assumed to have minimal knowledge. As in this paper, the probability of switching to a better path depends on the latency difference, and care has to be taken to avoid oscillation. Also in the Wardrop setting, Blum, Even-Dar, and Ligett [2] show that approximate NE is the outcome of *regret-minimizing* rerouting strategies, in which an agent’s cost, averaged over time, should approximate the cost of the best individual link available to that agent.

Certain generalizations of singleton games have also been considered. These generalizations are not strictly congestion games according to the standard definition we gave above, but many ideas carry over. One version introduced by Koutsoupias and Papadimitriou [18] has been studied extensively in different contexts (for example, [20, 6, 13, 4, 28]). In this generalization, each task may have a numerical *weight* (sometimes called traffic, or demand), and each resource has a *speed* (or capacity). The cost of using a resource is the total weight of tasks using it, divided by its speed. Even-Dar, Kesselman, and Mansour [7] give a generalized version of the potential function of [26] that applies to these games and which was subsequently used in [14]. For these games, however, it seems harder to find polynomial-length best-response sequences. Feldman et al. [10] show how a sequence of steps may lead to NEs, under the weaker condition that the maximal cost experienced by agents must not increase, but individual steps need not necessarily be “selfish.” They also note that poorly chosen better-response moves may lead to an exponential convergence rate. Another generalization of singleton games is *player-specific* cost functions [21], which allow different agents to have different cost functions for the same resource. In this setting there is no potential function, and better-response dynamics may cycle, although it remains the case that pure NEs always exist.

Our rerouting strategy is also related to reallocation processes for balls-into-bins games. The goal of a balls-into-bins game is to allocate m balls as evenly as possible into n bins. It is well known that a fairly even distribution can be achieved if every ball is allowed to randomly choose d bins and then the ball is allocated to the least loaded among the chosen bins (see [23] for an overview). Czumaj, Riley, and Scheideler [5] consider such an allocation where each ball initially chooses two bins. They show that, in a polynomial number of steps, the reallocation process ends up in a state with maximum load at most $\lceil m/n \rceil + 1$. Sanders, Egner, and Korst [29] show that a maximum load of $\lceil m/n \rceil + 1$ is optimal if every ball is restricted to two random choices.

In conclusion, this paper sits at one end of a spectrum in which we study a very simple load-balancing game, but we seek solutions in a very adverse setting in which agents have, at any point in time, a minimal amount of information about the state of their environment and carry out actions simultaneously in a strongly distributed sense.

1.2. Overview of our results. Section 3 deals with upper bounds on convergence time. The main result, Theorem 3.1, is that the protocol of Figure 2.1 converges to an NE within expected time $O(\log \log m + n^4)$.

The proof of Theorem 3.1 shows that the system becomes *approximately* balanced very rapidly. Specifically, Corollary 3.11 shows that if $n \leq m^{1/3}$, then for all ϵ , either

version of the distributed protocol (with or without neutral moves allowed) attains an ϵ -NE (where all load ratios are within $[1 - \epsilon, 1 + \epsilon]$; we use ϵ to denote a multiplicative factor as in [3]) in expected $O(\log \log m)$ rounds. The rest of section 3 analyzes the protocol of Figure 2.1. It is shown that within an additional $O(n^4)$ rounds the system becomes optimally balanced.

In section 4, we provide two lower bound results. The first one, Theorem 4.1, shows that the first protocol (of Figure 1.1, including moves that do not necessarily yield a strict improvement for an individual task but allow for simply “neutral” moves as well) results in exponential (in n) expected convergence time. Finally, in Theorem 4.2 we provide a general lower bound (regardless of which of the two protocols is being used) on the expected convergence time of $\Omega(\log \log m)$. This lower bound matches the upper bound as a function of m .

2. Notation. There are m tasks and n resources. An assignment of tasks to resources is represented as a vector (x_1, \dots, x_n) in which x_i denotes the number of tasks that are assigned to resource i . In the remainder of this paper, $[n]$ denotes $\{1, \dots, n\}$. The assignment is an NE if for all $i \in [n]$ and $j \in [n]$, $|x_i - x_j| \leq 1$. We study a distributed process for constructing an NE. The states of the process, $X(0), X(1), \dots$, are assignments. The transition from state $X(t) = (X_1(t), \dots, X_n(t))$ to state $X(t+1)$ is given by the greedy distributed protocol in Figure 2.1.

For each task b **do** in parallel
 Let i_b be the current resource of task b
 Choose resource j_b uniformly at random
 Let $X_{i_b}(t)$ be the current load of resource i
 Let $X_{j_b}(t)$ be the current load of resource j
If $X_{i_b}(t) > X_{j_b}(t) + 1$ **then**
 Move task b from resource i_b to j_b with probability $1 - X_{j_b}(t)/X_{i_b}(t)$

FIG. 2.1. The modified protocol, with “neutral moves” disallowed.

Note that if $X(t)$ is an NE, then $X(t+1) = X(t)$ so the assignment stops changing. Here is a formal description of the transition from a state $X(t) = x$. Independently, for every $i \in [n]$, let $(Y_{i,1}(x), \dots, Y_{i,n}(x))$ be a random variable drawn from a multinomial distribution with the constraint $\sum_{j=1}^n Y_{i,j}(x) = x_i$. (Y_{ij} represents the number of migrations from i to j in a round.) The corresponding probabilities $(p_{i,1}(x), \dots, p_{i,n}(x))$ are given by

$$p_{i,j}(x) = \begin{cases} \frac{1}{n} \left(1 - \frac{x_j}{x_i}\right) & \text{if } x_i > x_j + 1, \\ 0 & \text{if } i \neq j \text{ but } x_i \leq x_j + 1, \\ 1 - \sum_{j \neq i} p_{i,j}(x) & \text{if } i = j. \end{cases}$$

Then $X_i(t+1) = \sum_{\ell=1}^n Y_{\ell,i}(x)$.

For any assignment $x = (x_1, \dots, x_n)$, let $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$. We define the potential function $\Phi(x) = \sum_{i=1}^n (x_i - \bar{x})^2$. Note that $\Phi(x) = \sum_{i=1}^n x_i^2 - n\bar{x}^2$ and that a single selfish move reduces the potential.

3. Upper bound on convergence time. Our main result is the following.

THEOREM 3.1. *Let T be the number of rounds taken by the protocol of Figure 2.1 to reach an NE for the first time. Then $\mathbb{E}[T] = O(\log \log m + n^4)$.*

The proof of this theorem proceeds as follows. First (Lemma 3.6) we give an upper bound on $\mathbb{E}[\Phi(X(t))]$ which implies (Corollary 3.10) that there is a $\tau = O(\log \log m)$ such that, with high probability, $\Phi(X(\tau)) = O(n)$. We also show (Observation 3.5 and Corollary 3.14) that $\Phi(X(t))$ is a supermartingale and (Lemma 3.15) that it has enough variance. Using these facts, we obtain the upper bound on the convergence time.

DEFINITION. Let $S_i(x) = \{j \mid x_j < x_i - 1\}$. $S_i(x)$ is the set of resources that are significantly smaller than resource i in state x (in the sense that their loads are at least two tasks smaller than the load of resource i). Similarly, let $L_i(x) = \{j \mid x_j > x_i + 1\}$ and let $d_i(x) = \frac{1}{n} \sum_{j:|x_i-x_j|\leq 1} (x_i - x_j)$.

OBSERVATION 3.2. $\mathbb{E}[X_i(t+1) \mid X(t) = x] = \bar{x} + d_i(x)$.

Proof.

$$\begin{aligned} \mathbb{E}[X_i(t+1) \mid X(t) = x] &= \sum_{\ell=1}^n \mathbb{E}[Y_{\ell,i}(x)] = \sum_{\ell=1}^n x_\ell p_{\ell,i}(x) \\ &= \sum_{\ell \in L_i(x)} x_\ell \frac{1}{n} \left(1 - \frac{x_i}{x_\ell}\right) + x_i \left(1 - \sum_{j \in S_i(x)} \frac{1}{n} \left(1 - \frac{x_j}{x_i}\right)\right) \\ &= x_i + \frac{1}{n} \left(\sum_{\ell \in L_i(x)} (x_\ell - x_i) - \sum_{j \in S_i(x)} (x_i - x_j) \right) \\ &= x_i + \frac{1}{n} \sum_{\ell \in L_i(x) \cup S_i(x)} (x_\ell - x_i) \\ &= x_i + \frac{1}{n} \sum_{\ell=1}^n (x_\ell - x_i) - \frac{1}{n} \sum_{\ell \notin L_i(x) \cup S_i(x)} (x_\ell - x_i) \\ &= \bar{x} - \frac{1}{n} \sum_{\ell \notin L_i(x) \cup S_i(x)} (x_\ell - x_i) \\ &= \bar{x} + \frac{1}{n} \sum_{\ell \notin L_i(x) \cup S_i(x)} (x_i - x_\ell). \quad \square \end{aligned}$$

OBSERVATION 3.3. $\sum_{i=1}^n (\mathbb{E}[X_i(t+1) \mid X(t) = x])^2 = n\bar{x}^2 + \sum_{i=1}^n d_i(x)^2$.

Proof. Using Observation 3.2,

$$\sum_{i=1}^n (\mathbb{E}[X_i(t+1) \mid X(t) = x])^2 = \sum_{i=1}^n (\bar{x} + d_i(x))^2 = n\bar{x}^2 + 2\bar{x} \sum_{i=1}^n d_i(x) + \sum_{i=1}^n d_i(x)^2,$$

and the second term is zero since $d_i(x) = \mathbb{E}[X_i(t+1) \mid X(t) = x] - \bar{x}$. \square

OBSERVATION 3.4. $\text{var}[X_i(t+1) \mid X(t) = x] \leq \frac{1}{n} \sum_{\ell \in L_i(x)} (x_\ell - x_i) + \frac{1}{n} \sum_{j \in S_i(x)} (x_i - x_j)$.

Proof.

$$\begin{aligned} \text{var}(X_i(t+1) \mid X(t) = x) &= \sum_{\ell=1}^n \text{var}(Y_{\ell,i}(x)) = \sum_{\ell=1}^n x_\ell p_{\ell,i}(x) (1 - p_{\ell,i}(x)) \\ &= \sum_{\ell \in L_i(x)} x_\ell \frac{1}{n} \left(1 - \frac{x_i}{x_\ell}\right) (1 - p_{\ell,i}(x)) \end{aligned}$$

$$\begin{aligned}
 &+ x_i p_{i,i}(x) \left(\sum_{j \in S_i(x)} \frac{1}{n} \left(1 - \frac{x_j}{x_i} \right) \right) \\
 &= \frac{1}{n} \sum_{\ell \in L_i(x)} (x_\ell - x_i)(1 - p_{\ell,i}(x)) + p_{i,i}(x) \frac{1}{n} \sum_{j \in S_i(x)} (x_i - x_j) \\
 &\leq \frac{1}{n} \sum_{\ell \in L_i(x)} (x_\ell - x_i) + \frac{1}{n} \sum_{j \in S_i(x)} (x_i - x_j). \quad \square
 \end{aligned}$$

DEFINITION. For any assignment x , let $s_i(x) = |\{j \mid x_j = x_i - 1\}|$ and $l_i(x) = |\{j \mid x_j = x_i + 1\}|$. Let $u_1(x) = \sum_{i=1}^n \sum_{j \in [n]: |x_i - x_j| > 1} |x_i - x_j|$ and $u_2(x) = \sum_{i=1}^n (s_i(x) - l_i(x))^2$. Let $u(x) = u_1(x)/n + u_2(x)/n^2$. We will show that $u(x)$ is an upper bound on the expected potential after one step, starting from state x . The quantity $u_1(x)$ corresponds to the contribution arising from the sum of the variances of the individual loads, and $u_2(x)$ corresponds to the rest.

OBSERVATION 3.5. $\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] \leq u(x)$.

Proof.

$$\begin{aligned}
 \mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] + n\bar{x}^2 &= \sum_{i=1}^n \mathbb{E}[X_i(t+1)^2 \mid X(t) = x] \\
 &= \sum_{i=1}^n (\mathbb{E}[X_i(t+1) \mid X(t) = x])^2 \\
 &\quad + \sum_{i=1}^n \text{var}(X_i(t+1) \mid X(t) = x).
 \end{aligned}$$

Using Observations 3.3 and 3.4, this is at most $n\bar{x}^2 + \sum_{i=1}^n d_i(x)^2 + u_1(x)/n$. But

$$d_i(x) = \frac{1}{n} \sum_{j: |x_i - x_j| \leq 1} (x_i - x_j) = \frac{1}{n} (s_i(x) - l_i(x)),$$

so the result follows. \square

LEMMA 3.6. $\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] \leq n + 2n^{1/2}\Phi(x)^{1/2}$.

Proof. In the proof of Observation 3.5, we established that $\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] \leq \sum_{i=1}^n d_i(x)^2 + u_1(x)/n$. Upper-bounding $u_1(x)$ and using $d_i(x) \leq 1$, we have

$$\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] \leq n + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n |x_i - x_j|,$$

and since $|x_i - x_j| \leq |x_i - \bar{x}| + |x_j - \bar{x}|$, this is at most $n + 2 \sum_{i=1}^n |x_i - \bar{x}|$. By Cauchy-Schwarz, $(\sum_i |x_i - \bar{x}| \cdot 1)^2 \leq \sum_i |x_i - \bar{x}|^2 \sum_i 1$; thus

$$\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] \leq n + 2 \left(n \sum_{i=1}^n |x_i - \bar{x}|^2 \right)^{1/2}. \quad \square$$

COROLLARY 3.7. $\mathbb{E}[\Phi(X(t+1))] \leq n + 2n^{1/2}(\mathbb{E}[\Phi(X(t))])^{1/2}$.

Proof. Using Lemma 3.6, $\mathbb{E}[\Phi(X(t+1))] \leq n + 2n^{1/2}\mathbb{E}[f^{1/2}]$, where f denotes the random variable $\Phi(X(t))$. By Jensen's inequality $\mathbb{E}[f^{1/2}] \leq (\mathbb{E}[f])^{1/2}$ since the square-root function is concave, so we get $\mathbb{E}[\Phi(X(t+1))] \leq n + 2n^{1/2}(\mathbb{E}[f])^{1/2}$. \square

LEMMA 3.8. *Either there is a $t' < t$ such that $\mathbb{E}[\Phi(X(t'))] \leq 18n$ or $\mathbb{E}[\Phi(X(t))] \leq 9^{1-2^{-t}} n^{1-2^{-t}} \Phi(X(0))^{2^{-t}}$.*

Proof. The proof is by induction on t . The base case is $t = 0$. For the inductive step, note that $1 - 2^{-t} = \sum_{k=1}^t 2^{-k}$. Suppose that for all $t' < t$, $\mathbb{E}[\Phi(X(t'))] > 18n$ (otherwise we are finished). Then by Corollary 3.7,

$$\mathbb{E}[\Phi(X(t))] = n + 2n^{1/2}(\mathbb{E}[\Phi(X(t-1))])^{1/2} \leq 3n^{1/2}(\mathbb{E}[\Phi(X(t-1))])^{1/2}.$$

Applying the inductive hypothesis,

$$\mathbb{E}[\Phi(X(t))] \leq 3n^{1/2}(3^{2(1-2^{-(t-1)})}n^{1-2^{-(t-1)}}\Phi(X(0))^{2^{-(t-1)}})^{1/2}. \quad \square$$

COROLLARY 3.9. *There is a $\tau \leq \lceil \lg \lg \Phi(X(0)) \rceil$ such that $\mathbb{E}[\Phi(X(\tau))] \leq 18n$.*

Proof. Take $t = \lceil \lg \lg \Phi(X(0)) \rceil$. Either there is a $\tau < t$ with $\mathbb{E}[\Phi(X(\tau))] \leq 18n$ or, by the lemma,

$$\mathbb{E}[\Phi(X(t))] \leq 9n\Phi(X(0))^{2^{-t}} \leq 18n. \quad \square$$

COROLLARY 3.10. *There is a $\tau \leq \lceil \lg \lg \Phi(X(0)) \rceil$ such that $\Pr(\Phi(X(\tau)) > 720n) \leq 1/40$.*

Proof. Consider the (nonnegative) random variable $Y = \Phi(X(\tau))$, where τ is the quantity from Corollary 3.9. Markov's inequality says that for any $a > 0$, $\Pr(Y \geq a) \leq \mathbb{E}[Y]/a$. Now use Corollary 3.9 with $a = 720n$. \square

COROLLARY 3.11. *For all $\epsilon > 0$, provided that $n < m^{1/3}$, the expected time to reach an ϵ -NE is $O(\log \log m)$.*

Proof. Since the bound is asymptotic as a function of m for fixed ϵ , we can assume without loss of generality that $m > (60/\epsilon)^2$ and that $\epsilon m/(2n)$ is an integer. We show that for any starting assignment $X(0)$, there exists $\tau \leq \log \log(m^2)$ such that $\Pr(X(\tau) \text{ is } \epsilon\text{-Nash}) > \frac{39}{40}$. This implies the statement of the result since the number of blocks of τ steps needed to reach an ϵ -NE is at most

$$1 + \left(\frac{1}{40}\right) + \left(\frac{1}{40}\right)^2 + \dots = \frac{40}{39} < 2.$$

Suppose assignment x is not ϵ -Nash. If $X(t) = x$, there exist resources i, j with $X_i(t) - X_j(t) > \epsilon m/n$. We use the following notation. Let $\Delta = \epsilon m/(2n)$. Let $\beta = X_i(t) - X_j(t) - 2\Delta$. Note $\beta > 0$. If $X(t+1)$ is obtained from $X(t)$ by transferring Δ tasks from i to j , then

$$\begin{aligned} & \Phi(X(t)) - \Phi(X(t+1)) \\ &= X_i(t)^2 + X_j(t)^2 - X_i(t+1)^2 - X_j(t+1)^2 \\ &= (2\Delta + \beta + X_j(t))^2 + X_j(t)^2 - (\Delta + \beta + X_j(t))^2 - (\Delta + X_j(t))^2 \\ &= 2\Delta(\Delta + \beta + X_j(t)) + \Delta^2 - (2\Delta X_j(t) + \Delta^2) \\ &= 2\Delta(\Delta + \beta) \geq \Delta^2 = (\epsilon m/2n)^2. \end{aligned}$$

It follows that $\Phi(X(t)) \geq (\epsilon m/2n)^2$. From Corollary 3.10, $\Pr(\Phi(X(\tau)) < 720n) > \frac{39}{40}$, for $\tau = \log \log(\Phi(0)) = O(\log \log m)$.

An assignment $X(\tau)$ with $\Phi(X(\tau)) \leq 720n$ must be ϵ -Nash if $(\epsilon m/2n)^2 > 720n$. Note that $m > n^3$ and $m > (60/\epsilon)^2$. Hence, from $\epsilon^2(60/\epsilon)^2 n^3 > 4.720.n^3$, we can deduce $\epsilon^2 m^2 > 4.720.n^3$; hence $(\epsilon m/2n)^2 > 720n$. \square

Corollary 3.10 tells us that $\Phi(X(\tau))$ is likely to be $O(n)$. We want to show that $\Phi(X(t))$ quickly gets even smaller (all the way to an NE) and to this end, we show that $\Phi(X(t))$ is a supermartingale. By Observation 3.5, it suffices to show $u(x) \leq \Phi(x)$, and we proceed with this. In the following, we shall consider the cases $|x_i - \bar{x}| < 2.5$ for all $i \in [n]$ (Lemma 3.12) and $\exists i \in [n] : |x_i - \bar{x}| \geq 2.5$ (Lemma 3.13) separately.

LEMMA 3.12. *Suppose that assignment $x = (x_1, \dots, x_n)$ satisfies $|x_i - \bar{x}| < 2.5$ for all $i \in [n]$. Then $u(x) \leq \Phi(x)$.*

Proof. For all $i \in [n]$ and $j \in [n]$ we have $|x_i - x_j| \leq |x_i - \bar{x}| + |x_j - \bar{x}| < 5$. Let $z = \min_i x_i$ so that every $x_i \in \{z, \dots, z + 4\}$. Let $n_i = |\{j \mid x_j = z + i\}|$. Then

$$n^2\Phi(x) = n^2 \sum_{i=1}^n x_i^2 - n \left(\sum_{i=1}^n x_i \right)^2 = n^2 \left(\sum_{j=0}^4 n_j(z+j)^2 \right) - \left(\sum_{j=0}^4 n_j(z+j) \right)^2.$$

Also, $n^2u(x) = nu_1(x) + u_2(x)$, where

$$u_1(x) = n_0(2n_2 + 3n_3 + 4n_4) + n_1(2n_3 + 3n_4) + n_2(2n_0 + 2n_4) + n_3(3n_0 + 2n_1) + n_4(4n_0 + 3n_1 + 2n_2)$$

and

$$u_2(x) = n_0n_1^2 + n_1(n_0 - n_2)^2 + n_2(n_1 - n_3)^2 + n_3(n_2 - n_4)^2 + n_4n_3^2.$$

Plugging in these expressions and simplifying, we get

$$\begin{aligned} n^2\Phi(x) - n^2u(x) &= 4n_0n_1n_2 + 3n_0^2n_3 + 4n_0n_1n_3 + 4n_0n_2n_3 + 4n_1n_2n_3 + 3n_0n_3^2 \\ &\quad + 8n_0^2n_4 + 12n_0n_1n_4 + 3n_1^2n_4 + 8n_0n_2n_4 + 4n_1n_2n_4 + 12n_0n_3n_4 \\ &\quad + 4n_1n_3n_4 + 4n_2n_3n_4 + 8n_0n_4^2 + 3n_1n_4^2, \end{aligned}$$

which is clearly nonnegative since all coefficients are positive. \square

LEMMA 3.13. *Suppose that assignment $x = (x_1, \dots, x_n)$ satisfies $|x_n - \bar{x}| \geq 2.5$ and, for all $i \in [n]$, $|x_i - \bar{x}| \leq |x_n - \bar{x}|$. Let $w = (w_1, \dots, w_{n-1})$ be the assignment with $w_i = x_i$ for $i \in [n - 1]$. Then $\Phi(x) - u(x) \geq \Phi(w) - u(w)$; that is, the lower bound on the potential drop for x is at least as big as that for w .*

Proof. Let $k = |x_n - \bar{x}|$. We will show that

- (1) $\Phi(x) - \Phi(w) \geq k^2$ and
- (2) $u(x) - u(w) \leq 2k + 1$.

Then

$$\Phi(x) - u(x) - (\Phi(w) - u(w)) \geq k^2 - (2k + 1),$$

which is nonnegative since $k \geq 2.5 \geq 1 + \sqrt{2}$.

First, we prove (1). Let $f(z) = \sum_{i=1}^{n-1} (x_i - z)^2$. Note that the derivative of $f(z)$ is

$$f'(z) = 2(n - 1)z - 2 \sum_{i=1}^{n-1} x_i = 2(n - 1)z - 2(n - 1)\bar{w}.$$

Furthermore, the second derivative is $f''(z) = 2(n - 1) \geq 0$. Thus, $f(z)$ is minimized at $z = \bar{w}$. Now note that

$$\Phi(x) - \Phi(w) = k^2 + \sum_{i=1}^{n-1} (x_i - \bar{x})^2 - \sum_{i=1}^{n-1} (x_i - \bar{w})^2 \geq k^2.$$

Now we finish the proof by proving (2). Assume first that $x_n = \bar{x} + k$. Then

$$u_1(x) - u_1(w) = 2 \sum_{i \in [n]: |x_i - x_n| > 1} |x_i - x_n| \leq 2 \sum_{i=1}^n |x_i - x_n| = 2 \sum_{i=1}^n (x_n - x_i) = 2nk.$$

Let $z_j = |\{\ell \mid x_\ell = j\}|$. Clearly $z_j = 0$ for $j > x_n$. Let $\xi = \lceil x_n - 2k \rceil$. For $\ell \in [n]$ we have $x_\ell \geq \bar{x} - k = x_n - 2k$, so $z_j = 0$ for $j < \xi$. Now $u_2(x) = \sum_{j=\xi}^{x_n} z_j (z_{j-1} - z_{j+1})^2$. The representation of w in terms of z_j s is the same as the representation of x except that z_{x_n} is reduced by one. Therefore,

$$\begin{aligned} u_2(x) - u_2(w) &= z_{x_n-1} \left((z_{x_n-2} - z_{x_n})^2 - (z_{x_n-2} - z_{x_n} + 1)^2 \right) + (z_{x_n-1} - z_{x_n+1})^2 \\ &= z_{x_n-1} (-2z_{x_n-2} + 2z_{x_n} + z_{x_n-1} - 1) \leq z_{x_n-1} (2z_{x_n} + z_{x_n-1}). \end{aligned}$$

But since $z_{x_n} \leq n - z_{x_n-1}$, the upper bound on the right-hand side is at most

$$z_{x_n-1} (2n - 2z_{x_n-1} + z_{x_n-1}) = 2z_{x_n-1} (n - z_{x_n-1}/2),$$

which is at most n^2 since the right-hand side is maximized at $z_{x_n-1} = n$. To finish the proof of (2), use the definition of u to deduce that

$$u(x) - u(w) \leq \frac{u_1(x) - u_1(w)}{n} + \frac{u_2(x) - u_2(w)}{n^2}.$$

The proof of (2) when $x_n = \bar{x} - k$ is similar. □

COROLLARY 3.14. *For any assignment $x = (x_1, \dots, x_n)$, $\Phi(x) - u(x) \geq 0$.*

Proof. The proof is by induction on n . The base case, $n = 1$, follows from Lemma 3.12. Suppose $n > 1$. Neither $\Phi(x)$ nor $u(x)$ depends upon the order of the components in x , so assume without loss of generality that $|x_i - \bar{x}| \leq |x_n - \bar{x}|$ for all i . If $|x_n - \bar{x}| < 2.5$, then apply Lemma 3.12. Otherwise, use Lemma 3.13 to find an assignment $w = (w_1, \dots, w_{n-1})$ such that $\Phi(x) - u(x) \geq \Phi(w) - u(w)$. By the inductive hypothesis, $\Phi(w) - u(w) \geq 0$. □

Together, Observation 3.5 and Corollary 3.14 tell us that $\mathbb{E}[\Phi(X(t+1)) \mid X(t) = x] \leq \Phi(x)$. The next lemma will be used to give a lower bound on the variance of the process. Let $V = 0.4n^{-2}$.

LEMMA 3.15. *Suppose that $X(t) = x$ and that x is not an NE. Then*

$$\Pr(\Phi(X(t+1)) \neq \Phi(x) \mid X(t) = x) \geq V.$$

Proof. Choose s and ℓ such that for all $i \in [n]$, $x_s \leq x_i \leq x_\ell$. Since x is not an NE, $x_\ell > x_s + 1$. Assuming $X(t) = x$, consider the following experiment for choosing $X(t+1)$.

The intuition behind the experiment is as follows. We wish to show that the transition from $X(t)$ to $X(t+1)$ has some variance in the sense that $\Phi(X(t+1))$ is sufficiently likely to differ from $\Phi(X(t))$. To do this, we single out a “least loaded” resource s and a “most loaded” resource ℓ as above. In the transition from $X(t)$ to $X(t+1)$ we make transitions from resources other than resource ℓ in the usual way. We pay special attention to transitions from resource ℓ (and particular attention to transitions from resource ℓ which could either go to resource s or stay at resource ℓ). It helps to be very precise about how the random decisions involving tasks that start at resource ℓ are made. In particular, for each task b that starts at resource ℓ , we first make a decision about whether b would *accept* the transition from resource ℓ to

resource s if b happened to choose resource s . Then we make the decision about which resource task b should choose. Of course, we cannot cheat and we have to sample from the original required distribution. Here are the details.

Independently, for every $i \neq \ell$, choose $(Y_{i,1}(x), \dots, Y_{i,n}(x))$ from the multinomial distribution described in section 2. (In the informal description above, this corresponds to making transitions from resources other than resource ℓ in the usual way.) Now, for every task $b \in x_\ell$, let $z_b = 1$ with probability $1 - x_s/x_\ell$ and $z_b = 0$ otherwise. (In the informal description above, this corresponds to deciding whether b would *accept* the transition to s if resource s were (later) chosen.) Let x_ℓ^+ be the number of tasks b with $z_b = 1$ and let x_ℓ^- be the number of tasks b with $z_b = 0$. Choose $(Y_{\ell,1}^+(x), \dots, Y_{\ell,n}^+(x))$ from a multinomial distribution with the constraint $\sum_{j=1}^n Y_{\ell,j}^+(x) = x_\ell^+$ and probabilities given by

$$p_{\ell,j}^+(x) = \begin{cases} \frac{1}{n} & \text{if } j = s, \\ \frac{1}{n} \left(1 - \frac{x_j}{x_\ell}\right) & \text{if } j \neq s \text{ and } x_\ell > x_j + 1, \\ 0 & \text{if } \ell \neq j \text{ but } x_\ell \leq x_j + 1, \\ 1 - \sum_{j \neq \ell} p_{\ell,j}(x) & \text{if } \ell = j. \end{cases}$$

Similarly, choose $(Y_{\ell,1}^-(x), \dots, Y_{\ell,n}^-(x))$ from a multinomial distribution with the constraint $\sum_{j=1}^n Y_{\ell,j}^-(x) = x_\ell^-$ and probabilities given by

$$p_{\ell,j}^-(x) = \begin{cases} 0 & \text{if } j = s, \\ \frac{1}{n} \left(1 - \frac{x_j}{x_\ell}\right) & \text{if } j \neq s \text{ and } x_\ell > x_j + 1, \\ 0 & \text{if } \ell \neq j \text{ but } x_\ell \leq x_j + 1, \\ 1 - \sum_{j \neq \ell} p_{\ell,j}(x) & \text{if } \ell = j. \end{cases}$$

For all j , let $Y_{\ell,j}(x) = Y_{\ell,j}^+(x) + Y_{\ell,j}^-(x)$. Informally, the $p_{\ell,j}^+$ transition probabilities are set up so that packets which decided that they would accept a transition to s behave appropriately, and the $p_{\ell,j}^-$ transition probabilities are set up so that packets which decided that they would *not* accept a transition to s behave appropriately. By combining the probabilities, we see that $X(t+1)$ is chosen from the correct distribution in this way.

Now, consider the transition from x to $X(t+1)$. Condition on the choice for $(Y_{i,1}(x), \dots, Y_{i,n}(x))$ for all $i \neq \ell$. Suppose $x_\ell^+ > 2$. Condition on the choice for $(Y_{\ell,1}^-(x), \dots, Y_{\ell,n}^-(x))$. Flip a coin for each of the first $x_\ell^+ - 2$ tasks with $z_b = 1$ to determine which of $Y_{\ell,1}^+(x), \dots, Y_{\ell,n}^+(x)$ the task contributes to. Condition on these choices. Consider the following options:

- (1) Let x_1 be the resulting value of $X(t+1)$ when we add both of the last two tasks to $Y_{\ell,\ell}^+(x)$.
- (2) Let x_2 be the resulting value of $X(t+1)$ when we add one of the last two tasks to $Y_{\ell,\ell}^+(x)$ and the other to $Y_{\ell,s}^+(x)$.
- (3) Let x_3 be the resulting value of $X(t+1)$ when we add both of the last two tasks to $Y_{s,s}^+(x)$.

Note that, given the conditioning, each of these choices occurs with probability at least n^{-2} . Also, $\Phi(x_1)$, $\Phi(x_2)$, and $\Phi(x_3)$ are not all the same. Thus, $\Pr(\Phi(X(t+1)) \neq \Phi(x) \mid X(t) = x, x_\ell^+ > 2) \geq n^{-2}$. Also,

$$\Pr(x_\ell^+ > 2) = 1 - \left(\frac{x_s}{x_\ell}\right)^{x_\ell} - x_\ell \left(1 - \frac{x_s}{x_\ell}\right) \left(\frac{x_s}{x_\ell}\right)^{x_\ell-1}.$$

Since the derivative with respect to x_s is negative, this is minimized by taking x_s as large as possible, namely $x_\ell - 2$; thus $\Pr(x_\ell^+ > 2) \geq 1 - 7e^{-2} \geq 0.4$, and the result follows. \square

In order to finish our proof of convergence, we need the following observation about $\Phi(x)$.

OBSERVATION 3.16. *For any assignment x , $\Phi(x) \leq m^2$. Let $r = m \bmod n$. Then $\Phi(x) \geq r(1 - r/n)$, with equality if and only if x is an NE.*

Proof. Suppose that in assignment x there are resources i and j such that $x_i - x_j \geq 2$. Let x' be the assignment constructed from x by transferring a task from resource i to resource j . Then

$$\begin{aligned} \Phi(x) - \Phi(x') &= x_i^2 - x_i'^2 + x_j^2 - x_j'^2 = x_i^2 - (x_i^2 - 2x_i + 1) + x_j^2 - (x_j^2 + 2x_j + 1) \\ &= 2x_i - 2x_j - 2 = 2(x_i - x_j) - 2 > 0. \end{aligned}$$

Now suppose that, in some assignment x' , resources i and j satisfy $x'_i \geq x'_j > 0$. Let x be the assignment constructed from x' by transferring a task from resource j to resource i . Since $(x'_i + 1) - (x'_j - 1) \geq 2$, the above argument gives $\Phi(x) > \Phi(x')$. We conclude that an assignment x with maximum $\Phi(x)$ must have all of the tasks in the same resource, with $\Phi(x) = m^2$.

Furthermore, an assignment x with minimum $\Phi(x)$ must have $|x_i - x_j| \leq 1$ for all i, j . In this case there must be r resources with loads of $q + 1$ and $n - r$ resources with loads of q , where $m = qn + r$. So

$$\Phi(x) = r(q + 1 - \bar{x})^2 + (n - r)(q - \bar{x})^2 = r\left(1 - \frac{r}{n}\right)^2 + (n - r)\left(\frac{r}{n}\right)^2 = r\left(1 - \frac{r}{n}\right).$$

Note that x is a Nash assignment if and only if $|x_i - x_j| \leq 1$ for all i and j . \square

Combining Observation 3.16 and Corollary 3.10, we find that there is a $\tau \leq \lceil \lg \lg m^2 \rceil$ such that $\Pr(\Phi(X(\tau)) > 720n) \leq 1/40$. Let $B = 7200n + \lceil \frac{m^2}{n} \rceil - \frac{m^2}{n}$. Let $t' = \tau + \lceil 10B^2/V \rceil$.

LEMMA 3.17. *Given any starting state $X(0) = x$, the probability that $X(t')$ is an NE is at least $3/4$.*

Proof. The proof is based on a standard martingale argument; see [19]. Suppose that $\Phi(X(\tau)) \leq 720n$. Let $W_t = \Phi(X(t + \tau)) - r(1 - r/n)$ and let $D_t = \min(W_t, B)$. Note that $D_0 \leq 720n$. Together, Observation 3.5 and Corollary 3.14 tell us that W_t is a supermartingale. This implies that D_t is also a supermartingale since

$$\mathbb{E}[D_{t+1} \mid D_t = x < B] \leq \mathbb{E}[W_{t+1} \mid W_t = x < B] \leq W_t = D_t,$$

and

$$\mathbb{E}[D_{t+1} \mid D_t = B] \leq B = D_t.$$

Together, Lemma 3.15 and Observation 3.16 tell us that if $x > 0$, $\Pr(W_{t+1} \neq W_t \mid W_t = x) \geq V$. Thus, if $0 < x < B$,

$$\begin{aligned} \Pr(D_{t+1} \neq D_t \mid D_t = x) &= \Pr(\min(W_{t+1}, B) \neq W_t \mid W_t = x) \\ &\geq \Pr(W_{t+1} \neq W_t \wedge B \neq W_t \mid W_t = x) \\ &= \Pr(W_{t+1} \neq W_t \mid W_t = x) \geq V. \end{aligned}$$

Since $D_{t+1} - D_t$ is an integer, $\mathbb{E}[(D_{t+1} - D_t)^2 \mid 0 < D_t < B] \geq V$. Let T be the first time at which either (a) $D_t = 0$ (i.e., $X(t + \tau)$ is an NE), or (b) $D_t = B$.

Note that T is a stopping time. Define $Z_t = (B - D_t)^2 - Vt$, and observe that $Z_{t \wedge T}$ is a submartingale, where $t \wedge T$ denotes the minimum of t and T . Let p be the probability that (a) occurs. By the optional stopping theorem $\mathbb{E}[D_T] \leq D_0$; thus $(1 - p)B = \mathbb{E}[D_T] \leq D_0$ and $p \geq 1 - D_0/B \geq \frac{9}{10}$. Also, by the optional stopping theorem

$$pB^2 - V\mathbb{E}[T] = \mathbb{E}[(B - D_T)^2] - V\mathbb{E}[T] = \mathbb{E}[Z_T] \geq Z_0 = (B - D_0)^2 > 0,$$

and thus $\mathbb{E}[T] \leq pB^2/V$. Conditioning on the occurrence of (a), it follows that $\mathbb{E}[T \mid D_T = 0] \leq B^2/V$. Hence $\Pr(T > 10B^2/V \mid D_T = 0) \leq \frac{1}{10}$. So, if we now run for $10B^2/V$ steps, then the probability that we do not reach an NE is at most $\frac{1}{40} + 2 \cdot \frac{1}{10} < 1/4$. \square

Now we can give the proof of Theorem 3.1.

Proof. Subdivide time into intervals of t' steps. The probability that the process has not reached an NE before the $(j + 1)$ st interval is at most $(1/4)^{-j}$. \square

4. Lower bounds. In this section we prove the lower bound results stated in the introduction. We will use the following Chernoff bound which can be found, for example, in [15]. Let $N \geq 1$ and let $p_i \in [0, 1]$ for $i = 1, \dots, N$. Let X_1, X_2, \dots, X_N be independent Bernoulli random variables with $\Pr(X_i = 1) = p_i$ for $i = 1, \dots, N$ and let $X = X_1 + \dots + X_N$. Then we have $\mathbb{E}[X] = \sum_{i=1}^N p_i$ and for $0 \leq \epsilon \leq 1$,

$$(4.1) \quad \Pr(X \leq (1 - \epsilon) \cdot \mathbb{E}[X]) \leq \exp\left(-\frac{\epsilon^2 \cdot \mathbb{E}[X]}{3}\right).$$

The following theorem gives an exponential lower bound for the expected convergence time of the process in Figure 1.1.

THEOREM 4.1. *Let $X(t)$ be the process in Figure 1.1 with $m = n$. Let $X(0)$ be the assignment given by $X(0) = (n, 0, \dots, 0)$. Let T be the first time at which $X(t)$ is an NE. Then $\mathbb{E}[T] = \exp(\Theta(\sqrt{n}))$.*

Proof. For an assignment x , let $n_0(x)$ denote the number of resources i with $x_i = 0$. Thus, $n_0(X(0)) = n - 1$. The (unique) NE x assigns one task to each resource; thus $n_0(x) = 0$. Let $k = \lfloor \sqrt{n} \rfloor$. We will show that for any assignment x with $n_0(x) \geq k$,

$$\Pr(n_0(X(t)) < k \mid X(t - 1) = x) \leq \exp(-\Theta(\sqrt{n})).$$

This implies the result.

Suppose $X(t - 1) = x$ with $n_0(x) \geq k$. For convenience, let n_0 denote $n_0(x)$. Let x' denote $X(t)$, and let n'_0 denote $n_0(x')$. We will show that, with probability at least $1 - \exp(-\Theta(\sqrt{n}))$, $n'_0 \geq k$. During the course of the proof, we will assume, where necessary, that n is sufficiently large. This is without loss of generality given the Θ notation in the statement of the result.

Case 1. $n_0 > 8k$.

Consider the protocol in Figure 1.1. Let $U = \{b \mid x_{j_b} = 0\}$. $\mathbb{E}[|U|] = n_0$, so by the Chernoff bound (4.1), $\Pr(|U| \leq \lceil \frac{n_0}{2} \rceil + \lceil \frac{3n_0}{8} \rceil) \leq \Pr(|U| \leq \frac{8}{9}n_0) = \exp(-\Theta(\sqrt{n}))$. Thus, $|U| \geq \lceil n_0/2 \rceil + \lceil 3n_0/8 \rceil$ with probability at least $1 - \exp(-\Theta(\sqrt{n}))$. Suppose this is the case. Partition U into U_1 and U_2 with $|U_1| = \lceil n_0/2 \rceil$. Let $W = \cup_{b \in U_1} \{j_b\}$. First, suppose $|W| \leq \frac{3}{8}n_0$. In that case

$$|\{j \mid x'_j > 0\}| \leq n - |U_1| + \frac{3}{8}n_0 = n - \lceil n_0/2 \rceil + \frac{3}{8}n_0 \leq n - k,$$

so $n'_0 \geq k$. Otherwise, let $U' = \{b \in U_2 \mid j_b \in W\}$. Since

$$\mathbb{E}[|U'|] = |U_2| \frac{|W|}{n_0} \geq \frac{9}{64} n_0 > \frac{9}{8} k,$$

by the Chernoff bound (4.1), $\Pr(|U'| \leq k) = \Pr(|U'| \leq (1 - \frac{1}{9})\mathbb{E}[|U'|]) = \exp(-\Theta(\sqrt{n}))$, recalling that $k = \lfloor \sqrt{n} \rfloor$. Thus $|U'| \geq k$ with probability at least $1 - \exp(-\Theta(\sqrt{n}))$, which implies $n'_0 \geq k$.

Case 2. $k \leq n_0 \leq 8k$.

Consider the protocol in Figure 1.1. Let L be the set of “loners” defined by $L = \{i \mid x_i = 1\}$ and let $\ell = |L|$. The number of resources i with $x_i > 1$ is $n - n_0 - \ell$, and this is at most half as many as the number of tasks assigned to such resources (which is $n - \ell$), so $\ell \geq n - 2n_0$. Let $U = \{b \mid i_b \in L \text{ and } x_{j_b} = 0\}$. $\mathbb{E}[|U|] = \frac{\ell n_0}{n} \geq \frac{(n - 2n_0)n_0}{n} = \Theta(\sqrt{n})$, so by the Chernoff bound (4.1), $\Pr(|U| \leq 2\lceil \frac{1}{4}\ell \frac{n_0}{n} \rceil) \leq \Pr(|U| \leq \frac{2}{3}\mathbb{E}[|U|]) \leq \exp(-\Theta(\sqrt{n}))$. Thus, $|U| \geq 2\lceil \frac{1}{4}\ell \frac{n_0}{n} \rceil$ with probability at least $1 - \exp(-\Theta(\sqrt{n}))$. Suppose this is the case. Let U_1 and U_2 be disjoint subsets of U of size $\lceil \frac{1}{4}\ell \frac{n_0}{n} \rceil$. Order tasks in U arbitrarily and let $S = \{b \in U \mid \text{for some } b' \in U \text{ with } b' < b, j_{b'} = j_b\}$. (Note that $|S|$ does not depend on the ordering.) Let $W = \cup_{b \in U_1} \{j_b\}$.

Note that if $|W| \leq \frac{1}{5}\ell \frac{n_0}{n}$, then $|S| \geq \frac{1}{20}\ell \frac{n_0}{n} > \frac{n_0}{40} (\frac{\ell}{n})^2$. Otherwise, let $U' = \{b \in U_2 \mid j_b \in W\}$. Since

$$\mathbb{E}[|U'|] = |U_2| \frac{|W|}{n_0} \geq \frac{n_0}{20} \left(\frac{\ell}{n}\right)^2,$$

by the Chernoff bound (4.1), $\Pr(|U'| \leq \frac{1}{2}\frac{n_0}{20} (\frac{\ell}{n})^2) \leq \exp(-\Theta(\sqrt{n}))$ (recall that $n_0 (\frac{\ell}{n})^2 \geq n_0 (\frac{n - 2n_0}{n})^2 \geq k (\frac{n - 16k}{n})^2 = \Theta(\sqrt{n})$), and thus $|U'| \geq \frac{n_0}{40} (\frac{\ell}{n})^2$ with probability at least $1 - \exp(-\Theta(\sqrt{n}))$; hence $|S| \geq \frac{n_0}{40} (\frac{\ell}{n})^2$.

Suppose then that $|S| \geq \frac{n_0}{40} (\frac{\ell}{n})^2$. Assuming that n is sufficiently large, $|S| \geq k/41$. Let $B_0 = \cup_{b \in U} \{j_b\}$ and $B_1 = \cup_{b \in L - U} \{i_b\}$. Note that every resource in $B_0 \cup B_1$ is used in x' for some task $b \in L$. Thus, $|B_0 \cup B_1| \leq \ell - |S|$. Let $R = \{i \mid x_i = 0\} \cup L - B_0 - B_1$. Then $|R| \geq n_0 + \ell - (\ell - |S|) \geq n_0 + |S| \geq (1 + \frac{1}{41})k$.

Let $T = \{b \mid i_b \notin L, j_b \in R\}$. $\mathbb{E}[T] = (n - \ell) \frac{|R|}{n}$, and

$$\Pr\left(T \geq \frac{|R|}{100}\right) \leq \binom{n - \ell}{\frac{|R|}{100}} \left(\frac{|R|}{n}\right)^{|R|/100} \leq \left(\frac{2n_0 e 100}{n}\right)^{|R|/100};$$

thus with probability at least $1 - \exp(-\Theta(\sqrt{n}))$, $T < |R|/100$. In that case, $n'_0 \geq |R|(1 - \frac{1}{100}) \geq k$. \square

The following theorem provides a lower bound on the expected convergence time regardless of which of the two protocols is being used.

THEOREM 4.2. *Suppose that m is even. Let $X(t)$ be the process in Figure 2.1 with $n = 2$. Let $X(0)$ be the assignment given by $X(0) = (m, 0)$. Let T be the first time at which $X(t)$ is an NE. Then $\mathbb{E}[T] = \Omega(\log \log m)$. The same result holds for the process in Figure 1.1.*

Proof. Note that both protocols have the same behavior since m is even, and, therefore, the situation $x_1 = x_2 + 1$ cannot arise. For concreteness, focus on the protocol in Figure 2.1.

Let $y(x) = \max_i x_i - m/2$ and let $y_t = y(X(t))$; thus $y_0 = m/2$ and, for an NE x , $y(x) = 0$. We will show that for any assignment x , $\Pr(y_{t+1} > y(x)^{1/10} \mid X(t) = x) \geq 1 - y_t^{-1/4}$. (There is nothing very special about the exact value “1/10”—this value is being used as part of an explicit “lack of concentration” inequality in the proof, noting that for a lower bound we essentially want to lower-bound the variances of the load distributions. This seems to require a somewhat ad hoc approach, in contrast with the usage of concentration inequalities.)

Suppose $X(t) = x$ is an assignment with $x_1 \geq x_2$. As we have seen in section 2, $Y_{1,2}(x)$ (the number of migrations from resource 1 to resource 2 in the round) is a binomial random variable

$$B\left(x_1, \frac{1}{2} \left(1 - \frac{x_2}{x_1}\right)\right) = B\left(\frac{m}{2} + y_t, \frac{2y_t}{m + 2y_t}\right).$$

In general, let T_t be the number of migrations from the most-loaded resource in $X(t)$ to the least-loaded resource and note that the distribution of T_t is $B\left(\frac{m}{2} + y_t, \frac{2y_t}{m + 2y_t}\right)$ with mean y_t . If $T_t = y_t + \ell$ or $T_t = y_t - \ell$, then $y_{t+1} = \ell$. Thus $\Pr(y_{t+1} > y_t^{1/10}) = \Pr(|T_t - \mathbb{E}[T_t]| > y_t^{1/10})$. We continue by showing that this binomial distribution is sufficiently “spread out” in the region of its mode that we can find an upper bound on $\Pr(y_{t+1} \leq y_t^{1/10})$. This will lead to our lower bound on the expected time for $(y_t)_t$ to decrease below some constant (we use the constant 16):

$$\begin{aligned} \Pr(T_t = y_t) &= \binom{\frac{1}{2}m + y_t}{y_t} \left(\frac{2y_t}{m + 2y_t}\right)^{y_t} \left(\frac{m}{m + 2y_t}\right)^{\frac{1}{2}m}, \\ \Pr(T_t = y_t + j) &= \binom{\frac{1}{2}m + y_t}{y_t + j} \left(\frac{2y_t}{m + 2y_t}\right)^{y_t + j} \left(\frac{m}{m + 2y_t}\right)^{\frac{1}{2}m - j}. \end{aligned}$$

Suppose $j > 0$. Then

$$\begin{aligned} \frac{\Pr(T_t = y_t + j)}{\Pr(T_t = y_t)} &= \left(\frac{2y_t}{m + 2y_t}\right)^j \left(\frac{m}{m + 2y_t}\right)^{-j} \left(\frac{y_t! (\frac{1}{2}m)!}{(y_t + j)! (\frac{1}{2}m + y_t - (y_t + j))!}\right) \\ &= \left(\frac{2y_t}{m}\right)^j \left(\prod_{\ell=1}^j \frac{\frac{1}{2}m + 1 - \ell}{y_t + \ell}\right) = \left(\frac{2y_t}{m}\right)^j \left(\prod_{\ell=1}^j \frac{m + 2 - 2\ell}{2y_t + 2\ell}\right) \\ &> \left(\frac{2y_t}{m}\right)^j \left(\prod_{\ell=1}^j \frac{m - 2j}{2y_t + 2j}\right) = \left[\left(\frac{2y_t}{m}\right) \left(\frac{m - 2j}{2y_t + 2j}\right)\right]^j. \end{aligned}$$

Similarly, for $j < 0$,

$$\begin{aligned} \frac{\Pr(T_t = y_t + j)}{\Pr(T_t = y_t)} &= \left(\frac{2y_t}{m}\right)^j \left(\prod_{\ell=1}^{|j|} \frac{y_t + 1 - \ell}{\frac{1}{2}m + \ell}\right) = \left(\frac{m}{2y_t}\right)^{|j|} \left(\prod_{\ell=1}^{|j|} \frac{2y_t + 2 - 2\ell}{m + 2\ell}\right) \\ &> \left(\frac{m}{2y_t}\right)^{|j|} \left(\frac{2y_t - 2|j|}{m + 2|j|}\right)^{|j|} = \left[\left(\frac{m}{2y_t}\right) \left(\frac{2y_t - 2|j|}{m + 2|j|}\right)\right]^{|j|} \\ &= \left[\left(\frac{2y_t}{m}\right) \left(\frac{m - 2j}{2y_t + 2j}\right)\right]^j. \end{aligned}$$

Thus for all j ,

$$\frac{\Pr(T_t = y_t + j)}{\Pr(T_t = y_t)} > \left[\left(\frac{2y_t}{m}\right) \left(\frac{m - 2j}{2y_t + 2j}\right)\right]^j = \left[\left(\frac{y_t}{y_t + j}\right) \left(\frac{m - 2j}{m}\right)\right]^j.$$

Thus for all j with $|j| \leq y_t^{1/4}$, where $y_t^{1/4}$ is the positive fourth root of y_t , this is at least

$$\begin{aligned} & \left(\frac{y_t}{y_t + y_t^{1/4}}\right)^{y_t^{1/4}} \left(\frac{m - 2y_t^{1/4}}{m}\right)^{y_t^{1/4}} \\ & \geq \left(\frac{y_t}{y_t + y_t^{1/4}}\right)^{y_t^{1/4}} \left(\frac{2y_t - 2y_t^{1/4}}{2y_t}\right)^{y_t^{1/4}} = \left(\frac{y_t - y_t^{1/4}}{y_t + y_t^{1/4}}\right)^{y_t^{1/4}} = \left(\frac{y_t + y_t^{1/4} - 2y_t^{1/4}}{y_t + y_t^{1/4}}\right)^{y_t^{1/4}} \\ & = \left(1 - \frac{2y_t^{1/4}}{y_t + y_t^{1/4}}\right)^{y_t^{1/4}} \geq \left(1 - \frac{2y_t^{1/4}}{y_t}\right)^{y_t^{1/4}} = \left(1 - 2y_t^{-3/4}\right)^{y_t^{1/4}} \\ & \geq 1 - 2y_t^{-3/4}y_t^{1/4} = 1 - 2y_t^{-1/2} \geq \frac{1}{2}, \end{aligned}$$

where the last inequality just requires $y_t \geq 16$.

Note that the mode of a binomial distribution is one or both of the integers closest to the expectation, and the distribution is monotonically decreasing as one moves away from the mode. But, for $|j| \leq y_t^{1/4}$, $\Pr(T_t = y_t + j) \geq \frac{1}{2} \Pr(T_t = y_t)$; hence $\Pr(T_t = y_t) \leq 2/(1 + 2y_t^{1/4})$. Since $\Pr(T_t = y_t + j) \leq \Pr(T_t = y_t)$, it follows that

$$\Pr(T_t \in [y_t - y_t^{1/10}, y_t + y_t^{1/10}]) \leq (2y_t^{1/10} + 1) \Pr(T_t = y_t) < 3y_t^{-3/20}.$$

We say that the transition from y_t to y_{t+1} is a “fast round” if $y_{t+1} \leq y_t^{1/10}$ (equivalently, it is a fast round if $T_t \in [y_t - y_t^{1/10}, y_t + y_t^{1/10}]$). Otherwise it is a slow round. Recall that $y_0 = m/2$. Let

$$r = \left\lceil \log_{10} \left(\frac{\log(y_0)}{\log(12^{20/3})} \right) \right\rceil.$$

If the first j rounds are slow, then $y_j \geq y_0^{10^{-j}}$. If $j \leq r$, then $y_0^{10^{-j}} \geq 12^{20/3}$; thus the probability that the transition from y_j to y_{j+1} is the first fast round is at most $3(y_0^{10^{-j}})^{-3/20} \leq 1/4$.

Also, if $j < r$, then these probabilities increase geometrically so that the ratio of the probability that the transition to y_{j+1} is the first fast round and the probability that the transition to y_j is the first fast round is

$$\frac{3(y_0^{10^{-(j+1)}})^{-3/20}}{3(y_0^{10^{-j}})^{-3/20}} = (y_0^{10^{-j} - 10^{-(j+1)}})^{3/20} \geq (y_0^{10^{-(j+1)}})^{3/20} \geq 12 \geq 2;$$

thus $\sum_{j=0}^{r-1} \Pr(\text{transition from } y_j \text{ to } y_{j+1} \text{ is the first fast round}) \leq 2 \cdot 1/4 = \frac{1}{2}$. Therefore, with probability at least $1/2$, all of the first r rounds are slow. In this case, $\arg \min_t (y_t \leq 16) = \Omega(\log \log(m))$, which proves the theorem. \square

We also have the following observation.

OBSERVATION 4.3. *Let $X(t)$ be the process in Figure 2.1 with $m = n$. Let $X(0)$ be the assignment given by $X(0) = (2, 0, 1, \dots, 1)$. Let T be the first time at which $X(t)$ is an NE. Then $\mathbb{E}[T] = \Omega(n)$.*

The observation follows from the fact that the state does not change until one of the two tasks assigned to the first resource chooses the second resource.

5. Summary. We have analyzed a very simple, strongly distributed rerouting protocol for m tasks on n resources. We have proved an upper bound of $(\log \log m + n^4)$ on the expected convergence time (convergence to an NE), and for $m > n^3$ an upper bound of $O(\log \log m)$ on the time to reach an approximate NE. Our lower bound of $\Omega(\log \log m + n)$ matches the upper bound as a function of m . We have also shown an exponential lower bound on the convergence time for a related protocol that allows “neutral moves.”

REFERENCES

- [1] H. ACKERMANN, H. RÖGLIN, AND B. VÖCKING, *On the impact of combinatorial structure on congestion games*, in Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2006, pp. 613–622.
- [2] A. BLUM, E. EVEN-DAR, AND K. LIGETT, *Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games*, in Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing, 2006, pp. 45–52.
- [3] S. CHIEN AND A. SINCLAIR, *Convergence to approximate Nash equilibria in congestion games*, in Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), New Orleans, LA, 2007, pp. 169–178.
- [4] A. CZUMAJ, P. KRISTA, AND B. VÖCKING, *Selfish traffic allocation for server farms*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC), Montreal, 2002, pp. 287–296.
- [5] A. CZUMAJ, C. RILEY, AND C. SCHEIDELER, *Perfectly balanced allocation*, in Proceedings of the 7th Annual International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM-APPROX), Princeton, NJ, Lecture Notes in Comput. Sci. 2764, Springer-Verlag, Berlin, 2003, pp. 240–251.
- [6] A. CZUMAJ AND B. VÖCKING, *Tight bounds for worst-case equilibria*, in Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, San Francisco, CA, 2002, pp. 413–420.
- [7] E. EVEN-DAR, A. KESSELMAN, AND Y. MANSOUR, *Convergence time to Nash equilibria*, in Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP), Eindhoven, The Netherlands, 2003, pp. 502–513.
- [8] E. EVEN-DAR AND Y. MANSOUR, *Fast convergence of selfish rerouting*, in Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Vancouver, BC, 2005, pp. 772–781.
- [9] A. FABRIKANT, C. H. PAPADIMITRIOU, AND K. TALWAR, *The complexity of pure Nash equilibria*, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC), 2004, pp. 604–612.
- [10] R. FELDMAN, M. GAIRING, T. LÜCKING, B. MONIEN, AND M. RODE, *Nashification and the coordination ratio for a selfish routing game*, in Proceedings of the 30th Annual International Colloquium on Automata, Languages and Programming (ICALP 30), Lecture Notes in Comput. Sci. 2719, Springer-Verlag, Berlin, 2003, pp. 514–526.
- [11] S. FISCHER, H. RÄCKE, AND B. VÖCKING, *Fast convergence to Wardrop equilibria by adaptive sampling methods*, in Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), Seattle, WA, 2006, pp. 53–662.
- [12] S. FISCHER AND B. VÖCKING, *Adaptive routing with stale information*, in Proceedings of the 24th Annual ACM SIGACT/SIGOPS Symposium on Principles of Distributed Computing (PODC), 2005, pp. 276–283.
- [13] D. FOTAKIS, S. KONTOGIANNIS, E. KOUTSOPIAS, M. MAVRONICOLAS, AND P. SPIRAKIS, *The structure and complexity of Nash equilibria for a selfish routing game*, in Proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP), Malaga, Spain, 2002, pp. 123–134.
- [14] P. GOLDBERG, *Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game*, in Proceedings of the 23rd Annual ACM SIGACT/SIGOPS Symposium on Principles of Distributed Computing (PODC), St. John’s, Newfoundland, 2004, pp. 131–140.
- [15] T. HAGERUP AND C. RÜB, *A guided tour of Chernoff bounds*, Inform. Process. Lett., 33 (1989), pp. 305–308.
- [16] S. IEONG, R. MCGREW, E. NUDELMAN, Y. SHOHAM, AND Q. SUN, *Fast and compact: A simple class of congestion games*, in Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI), 2005, pp. 489–494.

- [17] D. S. JOHNSON, C. H. PAPADIMITRIOU, AND M. YANNAKAKIS, *How easy is local search?*, J. Comput. System Sci., 37 (1988), pp. 79–100.
- [18] E. KOUTSOPIAS AND C. H. PAPADIMITRIOU, *Worst-case equilibria*, in Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Trier, Germany, 1999, pp. 404–413.
- [19] M. LUBY, D. RANDALL, AND A. SINCLAIR, *Markov chain algorithms for planar lattice structures*, SIAM J. Comput., 31 (2001), pp. 167–192.
- [20] M. MAVRONICOLAS AND P. SPIRAKIS, *The price of selfish routing*, in Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), Crete, Greece, 2001, pp. 510–519.
- [21] I. MILCHTAICH, *Congestion games with player-specific payoff functions*, Games Econom. Behav., 13 (1996), pp. 111–124.
- [22] V. S. MIRROKNI AND A. VETTA, *Convergence issues in competitive games*, in APPROX-RANDOM, Lecture Notes in Comput. Sci. 3122, Springer-Verlag, Berlin, 2004, pp. 183–194.
- [23] M. MITZENMACHER, A. RICHA, AND R. SITARAMAN, *The power of two random choices: A survey of techniques and results*, in Handbook of Randomized Computing, P. Pardalos, S. Rajasekaran, and J. Rolim, eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, 2000, pp. 255–312.
- [24] D. MONDERER AND L. S. SHAPLEY, *Potential games*, Games Econom. Behav., 14 (1996), pp. 124–143.
- [25] A. ORDA, R. ROM, AND N. SHIMKIN, *Competitive routing in multi-user communication networks*, IEEE/ACM Trans. Networking, 1 (1993), pp. 510–521.
- [26] R. W. ROSENTHAL, *A class of games possessing pure-strategy Nash equilibria*, Internat. J. Game Theory, 2 (1973), pp. 65–67.
- [27] T. ROUGHGARDEN, *Selfish Routing and the Price of Anarchy*, MIT Press, Cambridge, MA, 2005.
- [28] T. ROUGHGARDEN AND É. TARDOS, *How bad is selfish routing?*, J. ACM, 49 (2002), pp. 236–259.
- [29] P. SANDERS, S. EGNER, AND J. KORST, *Fast concurrent access to parallel disks*, in Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), San Francisco, CA, 2000, pp. 849–858.