IS THE PRICE OF ANARCHY THE RIGHT MEASURE FOR LOAD-BALANCING GAMES?

J. DONCEL, U. AYESTA, O. BRUN, AND B. J. PRABHU

ABSTRACT. Price of Anarchy is an oft-used worst-case measure of the inefficiency of non-cooperative decentralized architectures. For a non-cooperative load-balancing game with two classes of servers and for a finite or infinite number of dispatchers, we show that the Price of Anarchy is an overly pessimistic measure that does not reflect the performance obtained in most instances of the problem. We explicitly characterize the worst-case traffic conditions for the efficiency of non-cooperative load-balancing schemes, and show that, contrary to a common belief, the worst inefficiency is in general not achieved in heavy-traffic.

1. INTRODUCTION

Server farms are commonly used in a variety of applications, including cluster computing, web hosting, scientific simulation or even the rendering of 3D computer generated imagery. A central problem arising in the management of the distributed computing resources of a data center is that of balancing the load over the servers so that the overall performance is optimized¹. In a centralized architecture, a single dispatcher, or a routing agent, routes incoming jobs to a set of servers so as to optimize a certain performance objective, such as the mean processing time of jobs for instance. However, modern data centers commonly have thousands of processors and up, and it becomes difficult or even impossible to centrally implement a globally optimal load-balancing solution. For instance, Akamai Technologies revealed, in march 2012, that it operates 105,000 servers [12]. Similarly, it is estimated that Google has more than 900,000 servers, and the company recently revealed that container data center holds more than 45,000 servers in a single facility built in 2005 [11]. The ever growing size and complexity of modern server farms thus calls for decentralized control schemes.

In a decentralized routing architecture, several dispatchers are used with each one routing a certain portion of the traffic. There are several possible approaches for the implementation of decentralized routing mechanisms. Approaches based on distributed optimization techniques [3, 14], can be cumbersome to implement and can have significant synchronisation and communication overheads, thus reducing the scalability of the decentralized routing scheme.

An alternative approach is based on autonomous, self-interested agents [15]. Such routing schemes are also known as "selfish routing" since each dispatcher independently

(A1,A2,A3,A4) LAAS-CNRS, UNIVERSITÉ DE TOULOUSE, CNRS, TOULOUSE, FRANCE

E-mail addresses: jdoncel@laas.fr, urtzi@laas.fr, brun@laas.fr, bala@laas.fr.

Key words and phrases. Inefficiency, Price of Anarchy, Atomic Games.

¹We shall use the terms load-balancing and routing interchangeably.

seeks to optimize the performance perceived by the jobs it routes. This setting can be analysed within the framework of a non-cooperative routing game. The strategy that rational agents will choose under these circumstances is called a Nash Equilibrium and it is such that a unilateral deviation will not help any routing agent in improving the performance perceived by the traffic it routes. When the number of dispatchers grows to infinity (every incoming job is handled by a dispatcher and it takes its own routing decision) the corresponding equilibrium is given by the notion of Wardrop Equilibrium [18], where journey times are minimal and equal in all routes.

Apart from the obvious gain in scalability with respect to a centralized setting, there are wide-ranging advantages to non-cooperative routing schemes: ease of deployment, no need for coordination between the routing agents that just react to the observed performances of the servers, and robustness to failures and environmental disturbances. However, it is well-known that non-cooperative routing mechanisms are potentially inefficient. Indeed, in general, the Nash equilibrium resulting from the interactions of many self-interested routing agents with conflicting objectives does not correspond to an optimal routing solution; hence, the lack of regulation carries the cost of decreased overall performance.

A standard measure of the inefficiency of selfish routing is the Price of Anarchy (PoA) which was introduced by Koutsoupias and Papadimitriou [10]. It is defined as the ratio between the performance obtained by the worst Nash equilibrium and the global optimal solution. Thus the PoA measures the cost of having no central authority, irrespective of a specific data center architecture. A value of the PoA close to 1 indicates that, in the worst case, the gap between a Nash Equilibrium and the optimal routing solution is not significant, and thus that good performances can be achieved even without a centralized control. On the contrary, a high PoA value indicates that, under certain circumstances, the selfish behaviour of the dispatchers leads to a significant performance degradation.

Several recent works have shown that non-cooperative load-balancing can be very inefficient in the presence of non-linear delay functions, see, for example, [7], [2], [1], [17] and [4]. We just mention two of them here. First, Haviv and Roughgarden have considered in [7] the so-called non-atomic scenario where every arriving job can select the server in which it will be served. They have shown that in this scenario the PoA corresponds to the number of servers, implying that, in a server farm with S servers, the mean response time of jobs can be as high as S times the optimal one! Another important result on the PoA was proved by Ayesta *et al.* in [1]. They investigate the Price of Anarchy of a load balancing game with a finite number, say K , of dispatchers, the so-called atomic case, and with a price per unit time to be paid for processing a job, which depends on the server. They prove that for a system with two or more servers, the Price of Anarchy is of the order of $\sqrt{K}$, independently of the number of servers, implying that when the number of dispatchers grows large, the PoA grows unboundedly. The fact that the Nash equilibrium can be very inefficient has paved the way to a lot of research on mechanism design that aims at coming up with Nash equilibria that are efficient with respect to the centralized setting [8, 9, 16, 5, 15].

In this paper, we adopt the view that the worst-case analysis (PoA) of the inefficiency of selfish routing is overly pessimistic and that high PoAs are obtained in pathological instances that hardly occur in practice. For example, in [7], the worst-case architecture

has one server whose capacity is much larger (tending to infinity) compared to that of the other servers. It is doubtful that such asymmetries will occur in data-centers where processors are more than likely to have similar characteristics.

While the architecture of a data-center is more or less fixed, the incoming traffic volume can vary as a function of time. Thus, for applications such as data-centers, it seems more appropriate to compare the performance of selfish routing and the centralized setting for different traffic profiles and a *fixed data-center architecture* (number of servers and their capacities). For this reason, we define the *inefficiency* for a fixed architecture of a data-center as the performance ratio between the worst-case Nash equilibrium and the global optimal. The worst-case case is taken over all possible traffic profiles that the routing agents can be asked to route. As is true of the PoA, *inefficiency* can take values between 1 and ∞ . A higher value of *inefficiency* indicates a worse performance of selfish routing compared to centralized routing. As opposed to the PoA, the *inefficiency* depends on the parameters (the server speeds and the number of servers in our case) of the architecture. By calculating the worst possible *inefficiency*, one retrieves the PoA.

The main contributions in this work are the following:

- For an arbitrary architecture in the system, we characterize the traffic conditions (or load) associated with the *inefficiency*. Contrary to classical queueing theory, we show that the *inefficiency* is in general not achieved in heavy-traffic or close to saturation conditions. In fact, we show that the *inefficiency* is close to 1 in heavy-traffic. We also provide examples for which the *inefficiency* is obtained for fairly low values of the utilization rate.
- We observe, see beginning of section 4, that for each problem instance with arbitrary server capacities we can construct a scenario with two server speeds whose inefficiency is worse. We thus conjecture that the case of two classes of servers is the worst, and that our conclusions on the efficiency of non-cooperative load-balancing extend beyond this case.
- In the case of two server classes, we show that the *inefficiency* is obtained when selfish routing uses only one class of servers and is marginally using the second class of servers. This scenario was used in [7] and [1] to obtain a lower bound on the PoA for their models. We give a formal proof on why this is indeed the worst-case scenario for selfish routing. Further, we obtain a closed-form formula for the *inefficiency* which in particular depends only on the ratio of the number of servers in each class and on the ratio of the capacities of each class (but not on the total nor on their capacities). When the number of servers is large, we also show that the PoA is equal to $\frac{K}{2\sqrt{K}-1}$, where K is the number of dispatchers.
- We then show that the *inefficiency* is very close to 1 in most cases, and that it approaches the known upper bound (given by the PoA) only in a very specific setting, namely, when there is only one fast server that is infinitely faster than the slower ones.
- For an infinite number of dispatchers, i.e., the non-atomic case, and an arbitrary architecture of the system, we show that the performance of the decentralized and the centralized settings are not equal in the heavy-traffic regime, in contrast with the case of finite number of dispatchers. For the case of two servers classes, we give an expression of the *inefficiency* that depends only on the ratio of the

number of servers in each class and on the ratio of the capacities of each class. We show that the PoA equals the number of servers, which coincides with the result presented in [7] and that the PoA is achieved only when there is only one fast server which is infinitely faster than the slower ones. In all the other configurations, we observe that the *inefficiency* is very close to one.

We believe that our work opens a new avenue in the study of the PoA and we hope that future research will be done not only on the PoA, but also on the *inefficiency*.

The rest of the paper is organized as follows. In section 2 we describe the model. In section 3 we investigate the worst case traffic conditions. In section 4, we give more precise results for server farms with two classes of servers. We give the expression for the load which leads to *inefficiency*, and the corresponding value of the *inefficiency*. We study the *inefficiency* of a server farm with an infinite number of dispatchers in section 5. Finally, the main conclusions of this work are presented in section 6.

A conference version of this article appeared in [6].

2. PROBLEM FORMULATION

We consider a non-cooperative routing game with K dispatchers and S Processor-Sharing servers. Denote $\mathcal{C} = \{1, \dots, K\}$ to be the set of dispatchers and $\mathcal{S} = \{1, \dots, S\}$ to be the set of servers. Jobs received by dispatcher i are said to be jobs of stream i .

Server $j \in \mathcal{S}$ has capacity r_j . It is assumed that servers are numbered in the order of decreasing capacity, i.e., if $m \leq n$, then $r_m \geq r_n$. Let $\mathbf{r} = (r_j)_{j \in \mathcal{S}}$ denote the vector of server capacities and let $\bar{r} = \sum_{n \in \mathcal{S}} r_n$ denote the total capacity of the system.

Jobs of stream $i \in \mathcal{C}$ arrive to the system according to a Poisson process and have generally distributed service-times. We do not specify the arrival rate and the characteristics of the service-time distribution due to the fact that in an $M/G/1 - PS$ queue the mean number of jobs depends on the arrival process and service-time distribution only through the traffic intensity, i.e., the product of the arrival rate and the mean service-time. Let λ_i be the traffic intensity of stream i . It is assumed that $\lambda_i \leq \lambda_j$ for $i \leq j$. Moreover, it will also be assumed that the vector $\boldsymbol{\lambda}$ of traffic intensities belongs to the following set: $\Lambda(\bar{\lambda}) = \{\boldsymbol{\lambda} \in \mathbb{R}^K : \sum_{i \in \mathcal{C}} \lambda_i = \bar{\lambda}\}$, where $\bar{\lambda}$ denotes the total incoming traffic intensity. It will be assumed throughout the paper that $\bar{\lambda} < \bar{r}$, which is the necessary and sufficient condition to guarantee the stability of the system. We denote by $\rho = \frac{\bar{\lambda}}{\bar{r}}$ the total traffic of the system.

We will sometimes be interested in what happens when $\bar{\lambda} \rightarrow \bar{r}$, a regime which we will refer to as heavy-traffic ($\rho \rightarrow 1$).

Let $\mathbf{x}_i = (x_{i,j})_{j \in \mathcal{S}}$ denote the routing strategy of dispatcher i , with $x_{i,j}$ being the amount of traffic it sends towards server j . Dispatcher i seeks to find a routing strategy that minimizes the mean sojourn times of its jobs, which, by Little's law, is equivalent to minimizing the mean number of jobs in the system as seen by this stream. This optimization problem can be formulated as follows:

$$\begin{aligned}
& \text{(ROUTE-}i\text{)} \quad \text{minimize } T_i(\mathbf{x}) = \sum_{j \in \mathcal{S}} \frac{x_{i,j}}{r_j - y_j} \\
(1) \quad & \text{subject to } \sum_{j \in \mathcal{S}} x_{i,j} = \lambda_i, \quad i = 1, \dots, K, \\
(2) \quad & \text{and } 0 \leq x_{i,j} \leq r_j, \quad \forall j \in \mathcal{S},
\end{aligned}$$

where $y_j = \sum_{k \in \mathcal{C}} x_{k,j}$ is the traffic offered to server j . Note that the optimization problem solved by dispatcher i depends on the routing decisions of the other dispatchers since $y_j = x_{i,j} + \sum_{k \neq i} x_{k,j}$. We let $\mathcal{X}_i$ denote the set of feasible routing strategies for dispatcher i , i.e., the set of routing strategies satisfying constraints (1)-(2). A vector $\mathbf{x} = (\mathbf{x}_i)_{i \in \mathcal{C}}$ belonging to the product strategy space $\mathcal{X} = \bigotimes_{i \in \mathcal{C}} \mathcal{X}_i$ is called a strategy profile.

A Nash equilibrium of the routing game is a strategy profile from which no dispatcher finds it beneficial to deviate unilaterally. Hence, $\mathbf{x} \in \mathcal{X}$ is a Nash Equilibrium Point (NEP) if $\mathbf{x}_i$ is an optimal solution of problem (ROUTE- i) for all dispatcher $i \in \mathcal{C}$.

Let $\mathbf{x}$ be a NEP for the system with K dispatchers. The global performance of the system can be assessed using the global cost

$$D_K(\boldsymbol{\lambda}, \mathbf{r}) = \sum_{i \in \mathcal{C}} T_i(\mathbf{x}) = \sum_{j \in \mathcal{S}} \frac{y_j}{r_j - y_j},$$

where the offered traffic y_j are those at the NEP. The above cost represents the mean number of jobs in the system. Note that when there is a single dispatcher, we have a single dispatcher with $\lambda_1 = \bar{\lambda}$. The global cost can therefore be written as $D_1(\bar{\lambda}, \mathbf{r})$ in this case.

We shall use the ratio between the performance obtained by the Nash equilibrium and the global optimal solution as a metric in order to assess the *inefficiency* of a decentralized scheme with K dispatchers and S servers. We define the *inefficiency* as the performance ratio under the worst possible traffic conditions, namely:

$$(3) \quad \text{inefficiency } I_K^S(\mathbf{r}) = \sup_{\boldsymbol{\lambda} \in \Lambda(\bar{\lambda}), \bar{\lambda} < \bar{r}} \frac{D_K(\boldsymbol{\lambda}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}.$$

The rationale for this definition is that in practice the system administrator controls neither the total incoming traffic nor how it is split between the dispatchers, whereas the number of servers and their capacities are fixed. Therefore it makes sense to consider the worst traffic conditions for the *inefficiency* of selfish routing, provided the system is stable.

The PoA for this system as defined in [1] can be retrieved by looking at the worst *inefficiency*, i.e.,

$$(4) \quad \text{PoA}(K, S) = \sup_{\mathbf{r}} I_K^S(\mathbf{r}).$$

3. WORST CASE TRAFFIC CONDITIONS

In this section, we show that for a sufficiently large load, the ratio $\frac{D_K(\boldsymbol{\lambda}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ decreases with $\bar{\lambda}$. This result implies that the *inefficiency* of a data-center is not achieved in the

heavy-traffic regime. Moreover, we also prove that when the system is in heavy-traffic, i.e., when $\bar{\lambda} \rightarrow \bar{r}$, the performance of both settings is the same.

The main difficulty in determining the behaviour of the *inefficiency* stems from the fact that for most cases there are no easy-to-compute explicit expressions for the NEP. A first simplification results from the following theorem which was proved in one of our previous works [1]. It states that, among all traffic vectors with total traffic intensity $\bar{\lambda}$, the global cost $D_K(\boldsymbol{\lambda}, \mathbf{r})$ achieves its maximum when all dispatchers control the same fraction of the total traffic. Formally,

[1]

$$D_K(\boldsymbol{\lambda}, \mathbf{r}) \leq D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right), \quad \forall \boldsymbol{\lambda} \in \Lambda(\bar{\lambda}).$$

where $\mathbf{e}$ is the all-ones vector.

Thus, we have identified the traffic vector in the set $\Lambda(\bar{\lambda})$ which has the worst-ratio of global cost at the NEP to the global optimal cost. It follows from the above result that

Corollary 3.1.

$$I_K^S(\mathbf{r}) = \sup_{\bar{\lambda} < \bar{r}} \frac{D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right)}{D_1(\bar{\lambda}, \mathbf{r})}.$$

Routing games in which players have exactly the same strategy set are known as *symmetric* games. These games belong to the class of *potential games* [13], that is, they have the property that there exists a function, called the *potential* such that the NEP can be obtained as the solution of an optimization problem with the *potential* as the objective. This property considerably simplifies the computation of the NEP. Another important consequence of the above results is that the *inefficiency* depends only the total traffic intensity and not on individual traffic flows to each of the dispatchers.

Another consequence of theorem 3 is that the inefficiency of decentralized routing increases with the number of dispatchers, that is,

Lemma 3.1.

$$I_K^S(\mathbf{r}) \leq I_{K+1}^S(\mathbf{r}), \quad \forall K \geq 1.$$

Proof. We have for all $\bar{\lambda} < \bar{r}$,

$$D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right) = D_{K+1}\left(\left(\frac{\bar{\lambda}}{K} \mathbf{e}, 0\right), \mathbf{r}\right) \leq D_{K+1}\left(\frac{\bar{\lambda}}{K+1} \mathbf{e}, \mathbf{r}\right),$$

where the last inequality follows from theorem 3. It yields

$$\sup_{\bar{\lambda} < \bar{r}} \frac{D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right)}{D_1(\bar{\lambda}, \mathbf{r})} \leq \sup_{\bar{\lambda} < \bar{r}} \frac{D_{K+1}\left(\frac{\bar{\lambda}}{K+1} \mathbf{e}, \mathbf{r}\right)}{D_1(\bar{\lambda}, \mathbf{r})},$$

i.e., $I_K(\mathbf{r}) \leq I_{K+1}(\mathbf{r})$. □

Before going further, let us take a look at the ratio $\frac{D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right)}{D_1(\bar{\lambda}, \mathbf{r})}$ as a function of the load $\rho = \bar{\lambda}/\bar{r}$, as is shown in figure 1 for two and five dispatchers. The data-center characteristics are the following: 200 servers of speed 6, 100 servers of speed 3, 300 servers of speed 2, and 200 servers of speed 1. It can be observed that as the load increases the ratio

FIGURE 1. Evolution of the ratio of social costs for $K = 2$ and $K = 5$ as the load in the system ranges from 0% to 100%.

goes through peaks and valleys, and finally it moves towards 1 as the load approaches 1. In the numerical experiments, we noted that the peaks corresponded to the total traffic intensity when selfish routing started to use one more class of servers. Moreover, just after these peaks the number of servers used by selfish routing and the centralized one was the same. A similar behaviour is observed on different sets of experiments.

In general, it is not easy to make formal the above observation, that is to say, there are no simple expressions for the value of loads which corresponds to the peaks and the valleys. However, in heavy-traffic, it helps to observe that both selfish and centralized routing will be using the same number of servers. Then, in order to show that heavy-traffic conditions are not inefficient, it is sufficient to show that the ratio decreases with load when both settings use the same number of servers.

Proposition 3.1. *If the total traffic intensity $\bar{\lambda}$ is such that centralized and the decentralized settings use the same number of servers (more than one), then the ratio of social costs $D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})/D_1(\bar{\lambda}, \mathbf{r})$ is decreasing with $\bar{\lambda}$.*

Proof. See appendix B.1. □

In the above result we exclude the case of one server so as to obtain a stronger result. If both the settings use just one server, then the ratio remains 1, which is non-increasing. For a sufficiently high load all the servers will be used by both settings in order to guarantee the stability of the system. It then follows that in a server farm with an arbitrary number of servers and with arbitrary server capacities, heavy-traffic regime is not inefficient. In fact, we can prove a stronger result which states that the *inefficiency* of the heavy-traffic regime is close to 1, that is, in heavy-traffic both settings have similar performances. Formally,

Theorem For a fixed $K < \infty$,

$$\lim_{\bar{\lambda} \rightarrow \bar{r}} \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = 1.$$

FIGURE 2. Evolution of the ratio of social costs when $K = 5$ and $S = 13$ as the load in the system ranges from 0% to 90% for a server farm with different values of the capacities.

Proof. See appendix B.2. □

It is important that the number of dispatchers be finite for the above result to hold. As we show in section 5, if the number of dispatchers is infinite, as in the case of non-atomic games, the above limit may be a value larger than 1.

This result is important because it is widely believed that the maximum inefficiency of the decentralized routing scheme is obtained in the heavy-traffic regime. Theorem 3 shows that this belief is false. As can be observed in figure 1, the worst case traffic conditions can occur at low or moderate utilization rates (in fact, the worst total traffic intensity can be arbitrary close to 0 if the server capacities are sufficiently close to each other). In heavy-traffic, even though the cost in both the settings will grow, the rate of growth is the same which results in a ratio close to 1.

The characterization of the exact traffic intensity which results in $I_K^S(\mathbf{r})$ proves to be a difficult task for arbitrary values of the capacities. In the following section we restrict ourselves to two server classes.

4. INEFFICIENCY FOR TWO-SERVER CLASSES

Before considering in detail the two-classes case, let us first observe how the inefficiency depends on the configuration of servers. Assume that there are 5 dispatchers and 13 servers. Figure 2 presents the evolution of the ratio $\frac{D_K(\frac{\lambda}{K}\mathbf{e}, \mathbf{r})}{D_1(\lambda, \mathbf{r})}$ according to the total load of the system for several vectors $\mathbf{r}$ of server capacities. We observe that the highest inefficiency is obtained in the case of two classes of servers with extreme capacity values. From extensive numerical experimentations, we conjecture that, given the number of servers, for each problem instance with arbitrary server capacities we can construct a scenario with the same number of servers and two server speeds whose inefficiency is worse. This is formally stated in conjecture 4.1.

Conjecture 4.1. *For a data-center of S servers with $r_1 > r_S$*

$$I_K^S(\mathbf{r}) \geq I_K^S(\mathbf{r}^*),$$

where $\mathbf{r} = (\overbrace{r_1, \dots, r_1}^m, \overbrace{r_S, \dots, r_S}^{S-m})$ and $\mathbf{r}^* = (\overbrace{r_1, \dots, r_1}^m, \overbrace{r_{m+1}, \dots, r_{S-1}, r_S}^{S-m})$, with $r_1 > r_{m+1} \geq \dots \geq r_{S-1} \geq r_S$ and $m \geq 1$.

A direct consequence of the above conjecture is that if for two classes of servers the inefficiency is very close to one except in some pathological cases, this should also be true for more than two classes.

In the following, we thus consider a server farm with two classes of servers. Let S_1 be the number of "fast" servers of capacity r_1 , and $S_2 = S - S_1$ be the number of "slow" servers, each one of capacity r_2 , with $r_1 > r_2$ ². The behaviour of the ratio of social costs is illustrated in figure 3 in the case of a server farm with $S_1 = 100$ fast servers of capacity $r_1 = 100$, and $S_2 = 300$ slow servers of capacity $r_2 = 10$. We plot the evolution of the ratio $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})/D_1(\bar{\lambda}, \mathbf{r})$ as the load on the system ranges from 0% to 1 for $K = 2$, $K = 5$.

It was observed that for low loads both the settings used the fast servers. The ratio in this regime was 1. After a certain point, the centralized setting started to use the slow servers as well, and the ratio increased with the load until the point when the decentralized setting also started to use the slow servers. From this point on, the ratio decreased with increase in the load.

We shall now characterize the point where the ratio starts to increase and where the peak occurs. Define

$$\bar{\lambda}^{OPT} = S_1 \sqrt{r_1} (\sqrt{r_1} - \sqrt{r_2}),$$

and

$$\bar{\lambda}^{NE} = S_1 r_1 \left(1 - \frac{2}{\sqrt{(K-1)^2 + 4K \frac{r_1}{r_2}} - (K-1)} \right).$$

The following lemma gives the conditions on $\bar{\lambda}$ under which the centralized setting and the decentralized one use only the fast class of servers, or both classes.

Lemma 4.1. *For $K \geq 2$, $\bar{\lambda}^{OPT} < \bar{\lambda}^{NE}$, and*

- (1) *if $\bar{\lambda} \leq \bar{\lambda}^{OPT}$, both settings use only the "fast" servers,*
- (2) *if $\bar{\lambda}^{OPT} \leq \bar{\lambda} \leq \bar{\lambda}^{NE}$, the decentralized setting uses only the "fast" servers, while the centralized one uses all servers,*
- (3) *if $\bar{\lambda} > \bar{\lambda}^{NE}$, both settings use all servers.*

Proof. See appendix C.1. □

Since $\bar{\lambda}^{OPT} < \bar{\lambda}^{NE}$, a consequence of lemma 4.1 is that the decentralized setting always uses a subset of the servers used by the centralized one. We immediately obtain expressions of the social cost in the centralized and decentralized settings, as given in corollary 4.1.

Corollary 4.1. *For the centralized setting, if $\bar{\lambda} < \bar{\lambda}^{OPT}$, then*

$$D_1(\bar{\lambda}, \mathbf{r}) = \bar{\lambda} / (r_1 - \frac{\bar{\lambda}}{S_1}),$$

²In the case $r_2 = r_1$, it is easy to see that the NEP is always an optimal routing solution, whatever the total traffic intensity.

FIGURE 3. Evolution of the ratio of social costs for $K = 2$ and $K = 5$ as the charge in the system ranges from 0% to 100%.

otherwise

$$D_1(\bar{\lambda}, \mathbf{r}) = \left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}} \right) \right] \frac{1}{r_1 - y_1},$$

where $y_1 = \sqrt{r_1} \frac{\bar{\lambda} - S_2 \sqrt{r_2} (\sqrt{r_2} - \sqrt{r_1})}{S_1 \sqrt{r_1} + S_2 \sqrt{r_2}}$, and $y_2 = (\bar{\lambda} - S_1 y_1) / S_2$ are the loads on each fast server and on each slow server in the case $\bar{\lambda} > \bar{\lambda}^{OPT}$, respectively. Similarly, if $\bar{\lambda} < \bar{\lambda}^{NE}$, then

$$D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right) = \bar{\lambda} / (r_1 - \frac{\bar{\lambda}}{S_1}),$$

and

$$D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right) = \frac{1}{2} \sum_{j=1}^2 S_j \left[\sqrt{(K-1)^2 + 4K r_j \gamma(K)} - (K+1) \right]$$

otherwise.

Proof. See appendix C.2. □

In lemma 4.1, we identified three intervals, namely, $[0, \bar{\lambda}^{OPT})$, $[\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$, $[\bar{\lambda}^{NE}, \bar{r})$, each one corresponding to a different set of servers used by the two settings. In proposition 4.1, we describe how the ratio of the social costs evolves in each of these three intervals.

Proposition 4.1. *The ratio $D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}) / D_1(\bar{\lambda}, \mathbf{r})$ is*

- (a) *equal to 1 for $0 \leq \bar{\lambda} \leq \bar{\lambda}^{OPT}$,*
- (b) *strictly increasing over the interval $(\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$,*
- (c) *and strictly decreasing over the interval $(\bar{\lambda}^{NE}, \bar{r})$.*

Proof. See appendix C.3. □

Moreover, the ratio of social costs has the following property.

Lemma 4.2. *The ratio $D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}) / D_1(\bar{\lambda}, \mathbf{r})$ is a continuous function of $\bar{\lambda}$ over the interval $[0, \bar{r})$.*

FIGURE 4. The evolution of the ratio of social costs for $K = 2$ and $K = 5$ with respect to ρ in a server farm with 3 server classes.

Proof. See appendix C.4. □

We can now state the main result of this section.

Theorem The inefficiency is worst when the total arriving traffic intensity equals $\bar{\lambda}^{NE}$, namely,

$$I_K^S(\mathbf{r}) = \frac{D_K(\frac{\bar{\lambda}^{NE}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}^{NE}, \mathbf{r})}.$$

Proof. It was shown in lemma 4.2 that $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})/D_1(\bar{\lambda}, \mathbf{r})$ is a continuous function of $\bar{\lambda}$ over the interval $[0, \bar{r})$. Proposition 4.1.(a) states that the ratio is minimum for $0 \leq \bar{\lambda} \leq \bar{\lambda}^{OPT}$. For $\bar{\lambda}$ in $(\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$, we know from proposition 4.1.(b) that this ratio is strictly increasing, which implies that $I_K^S(\mathbf{r}) \geq D_K(\frac{\bar{\lambda}^{NE}}{K}\mathbf{e}, \mathbf{r})/D_1(\bar{\lambda}^{NE}, \mathbf{r})$ by continuity. Since, according to proposition 4.1.(c), the ratio is decreasing over the interval $(\bar{\lambda}^{NE}, \bar{r})$, we can conclude that its maximum value is obtained for $\bar{\lambda} = \bar{\lambda}^{NE}$. □

Theorem 4 fully characterizes the worst case traffic conditions for a server farm with two classes of servers. It states that the worst inefficiency of the decentralized setting is achieved when (a) each dispatcher controls the same amount of traffic and (b) the total traffic intensity is such that the decentralized setting only starts using the slow servers. The behaviour described by proposition 4.1 can easily be observed in figure 3.

For more than two classes of servers, we were unfortunately not able to prove the above results concerning the worst traffic conditions. Nevertheless, we conjecture that a similar behaviour happens also in this case. As another illustration of this behaviour, in figure 4 we plot the ratio of social costs as a function of the load on the system, for a server farm with 3 server classes (and for $K = 2$, $K = 5$) with $S_1 = 100$ fast servers of capacity $r_1 = 30$, $S_2 = 200$ intermediate servers of capacity $r_2 = 20$ and $S_3 = 100$ slow servers of capacity $r_3 = 10$.

FIGURE 5. Evolution of the *inefficiency* as a function of α and β for $K = 5$ dispatchers and $S = 1000$ servers.

4.1. Inefficiency for a given architecture. We now give the expression for the *inefficiency* of selfish routing for data-centers with two classes of servers. Using theorem 4 we assume the worst traffic conditions for the inefficiency of selfish routing, i.e., the symmetric game obtained for $\bar{\lambda} = \bar{\lambda}^{NE}$.

Proposition 4.2. *Let $\beta = \frac{r_1}{r_2} > 1$ and $\alpha = \frac{S_1}{S_2} > 0$, then*

$$(5) \quad I_K^S(\mathbf{r}) = \frac{1}{2} \frac{\sqrt{(K-1)^2 + 4K\beta} - (K+1)}{\frac{(\frac{1}{\alpha} + \sqrt{\beta})^2}{\frac{1}{\alpha} + \sqrt{(K-1)^2 + 4K\beta - (K-1)}} - (\frac{1}{\alpha} + 1)}.$$

Proof. According to theorem 4, we have $I_K^S(\mathbf{r}) = D_K(\frac{\bar{\lambda}^{NE}}{K}\mathbf{e}, \mathbf{r})/D_1(\bar{\lambda}^{NE}, \mathbf{r})$. The proof is then obtained after some algebra by using the expressions for $D_K(\frac{\bar{\lambda}^{NE}}{K}\mathbf{e}, \mathbf{r})$ and $D_1(\bar{\lambda}^{NE}, \mathbf{r})$ given in corollary 4.1, and the expression for $\bar{\lambda}^{NE}$ given in lemma 4.1. $\square$

The *inefficiency* $I_K^S(\mathbf{r})$ does not depend on the total number of servers S , but only on the ratio of server capacities and on the ratio of the numbers of servers of each type. In figure 5, we plot the *inefficiency* $I_K(\mathbf{r})$ of the non-cooperative routing scheme with $K = 5$ dispatchers and $S = 1000$ servers as the parameters α and β change from $\frac{1}{S-1}$ to 2 and from 1 to 1000, respectively. It can be observed that even for unbalanced scenarios (α small and β large), the *inefficiency* is always fairly close to 1, indicating that, even in the worst case traffic conditions, the gap between the NEP and the optimal routing solution is not significant. With slight abuse of notation, let us denote the RHS of (5) by $I_K(\alpha, \beta)$.

Lemma 4.3. *The function $I_K(\alpha, \beta)$ is decreasing with α .*

Proof. See appendix C.5. $\square$

A consequence of the above result is that given the ratio of server speeds in a data-center, the *inefficiency* is largest when there is one fast server and all the other servers

are slow. Selfish routing has the tendency to use the fast servers more than the slow ones. When there is just one fast server, its performance tends to be the worst as compared to that of the centralized routing which reduces its cost by sending traffic to the slower ones as well. Thus, in decentralized routing architectures, it is best to avoid server configurations with this particular kind of asymmetry.

4.2. Price of Anarchy. The PoA is defined as the worst possible *inefficiency* when the server capacities are varied. Then, from (3), (4) and proposition 4.2, it follows that

$$PoA(K, S) = \sup_{\alpha, \beta} I_K(\alpha, \beta).$$

From lemma 4.3 and the fact that, for a fixed S , α can take values in $\{\frac{1}{S-1}, \frac{2}{S-2}, \dots, S-1\}$, we can deduce that

$$(6) \quad PoA(K, S) = \sup_{\beta} I_K\left(\frac{1}{S-1}, \beta\right).$$

We are able to prove that $I_K\left(\frac{1}{S-1}, \beta\right)$ is increasing with β . This means that the PoA of a server farm with S servers and K dispatchers is achieved when $\alpha = \frac{1}{S-1}$ and β infinity, i.e., when there is only one fast server and it is infinitely faster than the slower ones. While there is no simple expression for the PoA in terms of K and S , we can nonetheless derive a certain number of properties from the preceding set of results.

Proposition 4.3. *The Price of Anarchy has the following properties.*

- (1) *For fixed K , $PoA(K, S)$ is increasing in S ; and*
- (2) *for a fixed S , $PoA(K, S)$ is increasing in K .*

Proof. For fixed K and for every β , from lemma 4.3 and (6),

$$I_K\left(\frac{1}{S-1}, \beta\right) \leq I_K\left(\frac{1}{S}, \beta\right) \leq \sup_{\beta} I_K\left(\frac{1}{S}, \beta\right) = PoA(K, S+1),$$

where the last equality follows from (6). Taking the supremum over β in the above inequality, we obtain, for a fixed K ,

$$PoA(K, S) \leq PoA(K, S+1),$$

which proves the first property.

For a fixed S and β , from lemma 3.1,

$$I_K\left(\frac{1}{S-1}, \beta\right) \leq I_{K+1}\left(\frac{1}{S-1}, \beta\right) \leq \sup_{\beta} I_{K+1}\left(\frac{1}{S-1}, \beta\right) = PoA(K+1, S),$$

Again, taking the supremum over β in the above inequality, we obtain, for a fixed S ,

$$PoA(K, S) \leq PoA(K+1, S),$$

which proves the second property. □

In figure 6, the PoA is plotted as a function of S for different values of K . It is observed that this value remains modest even when the number of servers is 10,000.

We now give an upper bound the PoA. For this, we first need the following result.

FIGURE 6. The Price of Anarchy as a function of the number of servers for different values of the number of dispatcher

Lemma 4.4. *For a server farm with two server classes and K dispatchers,*

$$\lim_{S \rightarrow \infty} PoA(K, S) = \frac{K}{2\sqrt{K} - 1}.$$

Proof. See appendix C.6. □

Proposition 4.4. *For a server farm with two server classes and K dispatchers, and for all K and S ,*

$$PoA(K, S) \leq \min \left(\frac{K}{2\sqrt{K} - 1}, S \right).$$

Proof. From proposition 4.3, $PoA(K, S)$ is increasing with S . Combining this fact with lemma 4.4, we can conclude that

$$PoA(K, S) \leq \frac{K}{2\sqrt{K} - 1}.$$

Moreover, it was shown in [7] that, for the Wardrop case which is the limit of $K \rightarrow \infty$, $PoA(\infty, S) \leq S$. Thus,

$$PoA(K, S) \leq S.$$

We can deduce the desired result from the above two inequalities. □

In server farms with large number of servers, it follows from lemma 4.4 that the PoA will be $\frac{K}{2\sqrt{K}-1}$. In [1], it was shown that this value was a lower bound on the PoA. The model in that paper had server dependent holding cost per unit time. The lower bound was obtained in an extreme case with negligible (tending to 0) holding cost on the fast servers and the decentralized setting marginally using the slow servers. Our present results show that the lower bound is indeed tight. Moreover, even in a less asymmetrical setting of equal holding costs per unit time, one can construct examples in which the PoA is attained.

The PoA obtained in the non-atomic case in [7] comes into play when there are few servers and a relatively large number of dispatcher. However, for data-centers the configuration is reversed: there are a few dispatchers and a large number of servers. In this case it is more appropriate to use the upper bound given in lemma 4.4.

5. INEFFICIENCY WITH AN INFINITE NUMBER OF DISPATCHERS

We know from lemma 3.1 that the *inefficiency* increases with the number of dispatchers K . This motivates the analysis of the *inefficiency* when the number of dispatchers K grows to infinity. In this section, we show that in the heavy-traffic regime the *inefficiency* is not one, in contrast to the case of finite K . For the case of two classes of servers we give the expression of the *inefficiency*, we characterize the situation under which the PoA is achieved and show that it is equal to the number of servers.

We define *inefficiency* of a server farm with S servers and an infinite number of dispatchers as

$$(7) \quad I_{\infty}^S(\mathbf{r}) = \lim_{K \rightarrow \infty} I_K^S(\mathbf{r}).$$

5.1. Heavy-traffic analysis. We study the performance of a data center with an arbitrary number of servers S when the system is in the heavy-traffic regime. We give the value of $\lim_{K \rightarrow \infty} \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ in the following proposition:

Proposition 5.1. *For a data-center with S servers, we have*

$$\lim_{\bar{\lambda} \rightarrow \bar{r}} \lim_{K \rightarrow \infty} \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = \frac{S \bar{r}}{\left(\sum_{i=1}^S \sqrt{r_i}\right)^2}.$$

Proof. See appendix D.1. □

We observe that this result does not coincide with the one obtained for the Nash equilibrium for K large, as given in theorem 3. This implies that the limits of the number of dispatchers and heavy-traffic do not interchange, i.e.,

$$\lim_{K \rightarrow \infty} \lim_{\bar{\lambda} \rightarrow \bar{r}} \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} \neq \lim_{\bar{\lambda} \rightarrow \bar{r}} \lim_{K \rightarrow \infty} \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}.$$

5.2. The case of two-server classes. We focus on the case of a data center with servers of two different speeds, r_1 and r_2 respectively, where $r_1 > r_2$. For the case of $K = \infty$, we conjecture that the worst inefficiency is obtained for two server-classes, i.e., the conjecture 4.1 holds when the number of dispatchers grows to infinity.

Let S_1 be the number of "fast" servers and $S_2 = S - S_1$ be the number of "slow" servers. We give the expression of the *inefficiency* for a data center with two classes of servers in terms only on the parameters $\alpha = \frac{S_1}{S_2}$ and $\beta = \frac{r_1}{r_2} > 1$.

Corollary 5.1.

$$(8) \quad I_{\infty}(\alpha, \beta) = \frac{(\beta - 1)(1 + \frac{1}{\alpha})}{(\sqrt{\beta} + \frac{1}{\alpha})^2 - (\frac{1}{\alpha} + 1)^2}.$$

FIGURE 7. Evolution of the *inefficiency* as a function of α and β for $K = 10^6$ and $S = 1000$

Proof. The result follows from theorem 4 and (7), taking into account that $\sqrt{(K-1)^2 + 4Kx} - (K-1) = 2x$, when $K \rightarrow \infty$, $\forall x \geq 1$. $\square$

Using this expression, we can use the parameters $\alpha \in \{\frac{1}{S-1}, \frac{2}{S-2}, \dots, \frac{S-2}{2}, S-1\}$ and β to characterize the inefficiency for a data center with two-server classes and infinite number of servers. Lemma 5.1 states the main properties of the inefficiency for an infinite number of dispatchers.

Lemma 5.1. *We have*

- $I_\infty(\alpha, \beta)$ is decreasing with α , for all β ,
- $I_\infty(\alpha, \beta)$ is increasing with β , for all α .

Proof. The result follows from corollary 5.1. $\square$

A direct consequence of corollary 5.1 and lemma 5.1 is that the Price of Anarchy of a data center with two classes of servers and infinite number of dispatchers is equal to the number of servers S .

Proposition 5.2.

$$PoA(\infty, S) = S.$$

Proof. According to lemma 5.1, the worst inefficiency is obtained when $\alpha = \frac{1}{S-1}$ and $\beta \rightarrow \infty$. Using the formula of corollary 5.1 for this values of α and β , we get the desired result. $\square$

We observe that this result coincides with the result given by [7].

We illustrate in figure 7 the evolution of the inefficiency of a server farm of $S = 1000$ servers and $K = 10^6$ dispatchers when α changes from $\frac{1}{S-1}$ to 0.15 and β from 1 to 10^8 . We observe that the inefficiency equals the PoA when $\alpha = \frac{1}{S-1}$ and $\beta = 10^8$, i.e., when there is only one fast server and it is 10^8 times faster than the slower ones. We also see in figure 7 that the inefficiency stays very close to one in most cases, even if the worst

inefficiency is 1000. Thus, we can conclude that although the inefficiency can be as bad as the number of servers when the number of dispatchers is infinity, the decentralized setting is almost always as efficient as the centralized setting.

5.3. The Price of Anarchy. We have seen that the PoA equals the number of servers in case of an infinity number of dispatchers, while for K finite the upper-bound is given by the minimum of $\frac{K}{2\sqrt{K}-1}$ and S .

We observe that, given the number of servers S , there exist a threshold on K^* such that

- if $K \leq K^*$, then $PoA \leq \frac{K}{2\sqrt{K}-1}$,
- if $K \geq K^*$, then $PoA \leq S$.

For a sufficiently large S , we can say that

$$S = \frac{K^*}{2\sqrt{K^*}-1} \approx 0.5\sqrt{K^*}.$$

Thus, we claim that for a sufficiently large S , then $K^* \approx 4S^2$ and this means that if the number of dispatchers is smaller than $4S^2$ the upper-bound on the PoA is given by $\frac{K}{2\sqrt{K}-1}$, and by the number of servers S otherwise.

6. CONCLUSIONS

Price of Anarchy is an oft-used worst-case measure of the inefficiency of noncooperative decentralized architectures. In spite of its popularity, we have observed that the Price of Anarchy is an overly pessimistic measure that does not reflect the performance obtained in most instances of the load-balancing game. For an arbitrary architecture in the system, we have seen that, contrary to a common belief, the inefficiency is in general not achieved in the heavy-traffic regime. Surprisingly, we have shown that inefficiency might be achieved at arbitrarily low load. For the case of two classes of servers we give an explicit expression of the inefficiency and we have shown that noncooperative load-balancing has close-to-optimal performances in most cases. We also show that the worst-case performances given by the Price of Anarchy occur only in a very specific setting, namely, when there is only one fast server and it is infinitely faster than the slower ones. We conjecture that our conclusions will also be true for more than two classes of servers.

We believe that our study opens up a new complementary point of view on the PoA and we hope that in the future researchers will not only investigate the PoA, but also the inefficiency. As our work suggests, even if the PoA is very bad, the inefficiency might be low in most instances of the problem. We believe that this issue should be investigated for other models.

REFERENCES

- [1] U. Ayesta, O. Brun, and B. J. Prabhu. Price of anarchy in non-cooperative load-balancing games. *Performance Evaluation*, 68:1312–1332, 2011.
- [2] C. H. Bell and S. Stidham. Individual versus social optimization in the allocation of customers to alternative servers. *Management Science*, 29:831–839, 1983.
- [3] D.P. Bertsekas and J.N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989.

- [4] H. L. Chen, J. Marden, and A. Wierman. The effect of local scheduling in load balancing designs. In *Proceedings of IEEE Infocom*, 2009.
- [5] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, 2005.
- [6] J. Doncel, U. Ayesta, O. Brun, and B. Prabhu. On the efficiency of non-cooperative load balancing. In *Proceedings of IFIP Networking*, 2013.
- [7] M. Haviv and T. Roughgarden. The price of anarchy in an exponential multi-server. *Operations Research Letters*, 35:421–426, 2007.
- [8] Y. A. Korilis, A. A. Lazar, and A. Orda. Achieving network optima using stackelberg routing strategies. *IEEE/ACM TRANSACTIONS ON NETWORKING*, 5(1), February 1997.
- [9] Y.A. Korilis, A.A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE J.Sel. A. Commun.*, 13(7), September 2006.
- [10] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS 1999*, 1999.
- [11] R. Miller. Google unveils its container data center, 2009.
- [12] R. Miller. Who has the most web servers?, 2009.
- [13] D. Monderer and L. S. Shapley. Potential games. *Games and Econ. Behavior*, 14:124–143, 1996.
- [14] D. Mosk-aoyama, T. Roughgarden, and D. Shah. Fully distributed algorithms for convex optimization problems. *SIAM Journal on Optimization*, 2010.
- [15] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, May 2005.
- [16] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, 2009.
- [17] S. Suri, C.D. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. In *Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures*, 2004.
- [18] J.G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers*, 1:325–378, 1952.

APPENDIX A. SOME KNOWN RESULTS

The results in this section are taken from [1]. Since they are cited several times in the present work, we choose to present them here for its easy perusal. Let $W(K, z) = \sum_{j \in \mathcal{S}} W_j(K, z)$, we define the function

$$(9) \quad W_j(K, z) = \mathbb{1}_{\{z \in [\frac{1}{r_j}, \frac{1}{r_{j+1}})\}} \cdot \left(\sum_{s=1}^j \frac{2r_s}{\sqrt{(K-1)^2 + 4Kr_s z} - (K-1)} - \sum_{s=1}^j r_s + \bar{\lambda} \right).$$

The following proposition gives the solution of the symmetric game.

Proposition A.1. *The subset of servers that are used at the NEP is $\mathcal{S}^*(K) = \{1, 2, \dots, j^*(K)\}$, where $j^*(K)$ is the greatest value of j such that $W(K, 1/r_{j+1}) \leq 0 < W(K, 1/r_j)$. The equilibrium flows are $x_{i,j}(K) = y_j(K)/K$, $i \in \mathcal{C}, j \in \mathcal{S}^*(K)$, where the offered traffic of server j is given by*

$$(10) \quad y_j(K) = r_j \frac{\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K+1)}{\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K-1)},$$

with $\gamma(K)$ the unique root of $W(K, z) = 0$ in $[\frac{1}{r_1}, \infty)$.

APPENDIX B. PROOFS OF THE RESULTS IN SECTION 3

B.1. Proof of proposition 3.1. Before proving proposition 3.1, we establish closed-form expressions for the value of the social costs in the centralized and decentralised

settings. Recall that we assume a server farm with S servers with decreasing values of the capacities, i.e, $r_i \leq r_j$, if $i > j$.

Lemma B.1. *Let n be the number of servers that the centralized setting uses, for $n = 1, \dots, S$, then*

$$D_1(\bar{\lambda}, \mathbf{r}) = \frac{\left(\sum_{j=1}^n \sqrt{r_j}\right)^2}{\sum_{j=1}^n r_j - \bar{\lambda}} - n.$$

Similarly, if the decentralized setting uses n servers, we have

$$D_K\left(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}\right) = \frac{1}{2} \sum_{j=1}^n \left[\sqrt{(K-1)^2 + 4Kr_j\gamma(K)} - (K+1) \right].$$

Proof. We first prove the results for the centralized setting. When the centralized setting uses n servers proposition A.1 states that in this case,

$$y_j(1) = r_j \left(1 - \frac{1}{\sqrt{\gamma(1)}\sqrt{r_j}} \right), \quad j = 1, \dots, n,$$

that yields

$$\frac{y_j(1)}{r_j - y_j(1)} = \sqrt{\gamma(1)}\sqrt{r_j} - 1, \quad j = 1, \dots, n.$$

We thus obtain

$$D_1(\bar{\lambda}, \mathbf{r}) = \sum_{j=1}^n \frac{y_j(1)}{r_j - y_j(1)} = \sqrt{\gamma(1)} \sum_{j=1}^n \sqrt{r_j} - n.$$

Since $\gamma(1)$ is the unique root of $W(K, \gamma(1)) = 0$ as defined in appendix A and, according to proposition A.1, then $\gamma(1)$ is the solution of

$$(11) \quad \frac{1}{\sqrt{\gamma(1)}} \sum_{j=1}^n \sqrt{r_j} = \sum_{j=1}^n r_j - \bar{\lambda}.$$

Thus, it follows that $\sqrt{\gamma(1)} = \sum_{j=1}^n \sqrt{r_j} / (\sum_{j=1}^n r_j - \bar{\lambda})$ and hence

$$D_1(\bar{\lambda}, \mathbf{r}) = \frac{\left(\sum_{j=1}^n \sqrt{r_j}\right)^2}{\sum_{j=1}^n r_j - \bar{\lambda}} - n.$$

Let us now consider the decentralized setting. If the number of servers used by the decentralized setting is n , then (10) gives that for $j = 1, \dots, n$

$$(12) \quad 1 - \frac{y_j(K)}{r_j} = \frac{2}{\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K-1)}.$$

From (10) and (12), it yields the desired result

$$D_K\left(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}\right) = \sum_{j=1}^n \frac{y_j(K)}{r_j} \left(1 - \frac{y_j(K)}{r_j} \right)^{-1} = \frac{1}{2} \sum_{j=1}^n \left[\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K+1) \right].$$

□

We show in the following lemma an important property to prove proposition 3.1.

Lemma B.2. *Let $a_k = \sqrt{(K-1)^2 + 4K\gamma(K)r_k} + (K-1)$, then for all $i > j$, $\frac{a_j}{a_i}$ is increasing with $\bar{\lambda}$.*

Proof. First, we define $b_j = \sqrt{(K-1)^2 + 4K\gamma(K)r_j}$ and we see that $\frac{b_j}{b_i}$ is increasing with $\bar{\lambda}$ if $\frac{(K-1)^2 + 4K\gamma(K)r_j}{(K-1)^2 + 4K\gamma(K)r_i}$ is increasing with $\bar{\lambda}$ because $\frac{b_j}{b_i}$ is positive and thus:

$$\frac{\partial}{\partial \bar{\lambda}} \left(\frac{(K-1)^2 + 4K\gamma(K)r_j}{(K-1)^2 + 4K\gamma(K)r_i} \right) \geq 0 \iff 4K \frac{\partial \gamma(K)}{\partial \bar{\lambda}} (K-1)^2 (r_j - r_i) \geq 0$$

that it is true due to $r_j \geq r_i$ if $i > j$. Hence, we have proved that $\frac{(K-1)^2 + 4K\gamma(K)r_j}{(K-1)^2 + 4K\gamma(K)r_i}$ is increasing with $\bar{\lambda}$ and this implies that $\frac{b_j}{b_i}$ is increasing with $\bar{\lambda}$.

We also observe that $\frac{\partial b_j}{\partial \bar{\lambda}} \geq \frac{\partial b_i}{\partial \bar{\lambda}}$, if $i > j$:

$$\frac{\partial b_j}{\partial \bar{\lambda}} \geq \frac{\partial b_i}{\partial \bar{\lambda}} \iff \frac{2K \frac{\partial \gamma(K)}{\partial \bar{\lambda}} r_j}{b_j} \geq \frac{2K \frac{\partial \gamma(K)}{\partial \bar{\lambda}} r_i}{b_i} \iff \frac{1}{\sqrt{\frac{(K-1)^2}{r_j^2} + \frac{4K\gamma(K)}{r_j}}} \geq \frac{1}{\sqrt{\frac{(K-1)^2}{r_i^2} + \frac{4K\gamma(K)}{r_i}}}$$

and this inequality holds since $r_j \geq r_i$ when $i > j$.

As $\frac{\partial b_j}{\partial \bar{\lambda}} = \frac{\partial a_j}{\partial \bar{\lambda}}$ and $a_j = b_j + (K-1)$, for $j = 1, \dots, n$, we are able to state that if $\frac{b_j}{b_i}$ is increasing, then $\frac{a_j}{a_i}$ is increasing with $\bar{\lambda}$:

$$\frac{\partial}{\partial \bar{\lambda}} \left(\frac{a_j}{a_i} \right) > 0 \iff \frac{\frac{\partial b_j}{\partial \bar{\lambda}} a_i - \frac{\partial b_i}{\partial \bar{\lambda}} a_j}{a_i^2} > 0 \iff \frac{\partial b_j}{\partial \bar{\lambda}} b_i - \frac{\partial b_i}{\partial \bar{\lambda}} b_j + (K-1) \left(\frac{\partial b_j}{\partial \bar{\lambda}} - \frac{\partial b_i}{\partial \bar{\lambda}} \right) > 0$$

and we know the inequality is satisfied because $\frac{\partial \left(\frac{b_j}{b_i} \right)}{\partial \bar{\lambda}} > 0$ and $\frac{\partial b_j}{\partial \bar{\lambda}} > \frac{\partial b_i}{\partial \bar{\lambda}}$. □

We can now prove proposition 3.1.

Proof. We show that when both settings use n servers ($n = 1, \dots, S$), then the ratio $\frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is decreasing with $\bar{\lambda}$. We use the expressions of $D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})$ and $D_1(\bar{\lambda}, \mathbf{r})$ given in lemma B.1 and we modify the ratio $\frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ as follows:

$$\begin{aligned} \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} &= \frac{\frac{1}{2} \sum_{j=1}^n \left[\sqrt{(K-1)^2 + 4K r_j \gamma(K)} - (K+1) \right]}{-n + \sqrt{\gamma(1)} \sum_{j=1}^n \sqrt{r_j}} \\ &= \frac{-n + \frac{1}{2} \sum_{j=1}^n \left[\sqrt{(K-1)^2 + 4K r_j \gamma(K)} - (K-1) \right]}{-n + \sqrt{\gamma(1)} \sum_{j=1}^n \sqrt{r_j}} \\ &= \frac{f_1 + f_2}{f_1 + g_2} \end{aligned}$$

where we define $f_1 = \frac{-n}{\sqrt{\gamma(1)}}$, $g_2 = \sum_{j=1}^n \sqrt{r_j}$ and $f_2 = \frac{1}{2\sqrt{\gamma(1)}} \sum_{j=1}^n \left[\sqrt{(K-1)^2 + 4K r_j \gamma(K)} - (K-1) \right]$.

We want to prove that the derivative of $\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ with respect to $\bar{\lambda}$ is negative and we see that

$$\frac{\partial}{\partial \bar{\lambda}} \left(\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} \right) < 0 \iff \frac{\partial f_1}{\partial \bar{\lambda}} (g_2 - f_2) + \frac{\partial f_2}{\partial \bar{\lambda}} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{\sqrt{\gamma(1)}} < 0.$$

We observe that f_1 is increasing with $\bar{\lambda}$, because $\gamma(1)$ increases with $\bar{\lambda}$, and $D_1(\bar{\lambda}, \mathbf{r}) \leq D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})$ implies that $g_2 \leq f_2$. Therefore, if we show that f_2 is decreasing with $\bar{\lambda}$ and we can conclude that $\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is decreasing with $\bar{\lambda}$. From (11) and (9), if both settings use n servers then

$$\frac{1}{\sqrt{\gamma(1)}} = \frac{\sum_{j=1}^n r_j - \bar{\lambda}}{\sum_{j=1}^n \sqrt{r_j}} = \frac{1}{\sum_{j=1}^n \sqrt{r_j}} \sum_{s=1}^n \frac{2r_s}{\bar{a}_s},$$

where $\bar{a}_s = \sqrt{(K-1)^2 + 4K\gamma(K)r_s} - (K-1)$.

We rewrite f_2 as follows:

$$f_2 = \frac{1}{\sum_{j=1}^n \sqrt{r_j}} \sum_{j=1}^n \bar{a}_j \sum_{s=1}^n \frac{r_s}{\bar{a}_s} = \frac{1}{\sum_{j=1}^n \sqrt{r_j}} \left(\sum_{j=1}^n r_j + \sum_{j=1}^n \sum_{i>j} \left[r_j \frac{\bar{a}_i}{\bar{a}_j} + r_i \frac{\bar{a}_j}{\bar{a}_i} \right] \right).$$

We define $a_s = \sqrt{(K-1)^2 + 4K\gamma(K)r_s} + (K-1)$ and we notice that if we multiply and divide $\bar{a}_s$ by a_s it yields $\bar{a}_s = \frac{4K\gamma(K)r_s}{a_s}$. So f_2 gets modified as follows with this property:

$$f_2 = \frac{1}{\sum_{j=1}^n \sqrt{r_j}} \left(\sum_{j=1}^n r_j + \sum_{j=1}^n \sum_{i>j} \left[r_j \frac{a_i}{a_j} + r_i \frac{a_j}{a_i} \right] \right).$$

Now, we show that $r_j/a_j^2 > r_i/a_i^2$ for all $i > j$ since $\frac{r_k}{a_k^2}$ is decreasing with k because we can write it as

$$\frac{r_k}{a_k^2} = \left[\left(\sqrt{\frac{(K-1)^2}{r_k} + 4K\gamma(K)} + \frac{K-1}{\sqrt{r_k}} \right)^{-1} \right]^2$$

and r_k decreases with k .

Finally, we see that f_2 is decreasing with $\bar{\lambda}$:

$$\frac{\partial f_2}{\partial \bar{\lambda}} = \frac{1}{\sum_{j=1}^n \sqrt{r_j}} \sum_{j=1}^n \sum_{i>j} \left[\left(\frac{\partial a_j}{\partial \bar{\lambda}} a_i - \frac{\partial a_i}{\partial \bar{\lambda}} a_j \right) \left(\frac{r_i}{a_i^2} - \frac{r_j}{a_j^2} \right) \right] < 0$$

and we conclude that this is true because from lemma B.2 $\frac{a_j}{a_i}$ is increasing with $\bar{\lambda}$ if $i > j$ (so that $\frac{\partial a_j}{\partial \bar{\lambda}} a_i - \frac{\partial a_i}{\partial \bar{\lambda}} a_j > 0$) and we have observed that $r_j/a_j^2 > r_i/a_i^2$ when $i > j$. $\square$

B.2. Proof of theorem 3.

Proof. First, we know that in heavy-traffic all the servers are used, so we consider that S servers are used in both settings. We also observe that in heavy-traffic $\gamma(K)$, as defined in proposition A.1, tends to ∞ , and thus the following approximation is satisfied when $\bar{\lambda} \rightarrow \bar{r}$:

$$(13) \quad \sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K-1) \approx 2\sqrt{K\gamma(K)r_j}$$

From (13) and the definition of $\gamma(K)$, we obtain

$$\sqrt{K\gamma(K)} \approx \frac{\sum_{j=1}^S \sqrt{r_j}}{\sum_{j=1}^S r_j - \bar{\lambda}}.$$

Now, using this expression, (13) and lemma B.1, we show that $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) \approx D_1(\bar{\lambda}, \mathbf{r})$ in heavy-traffic:

$$\begin{aligned} D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) &= \frac{1}{2} \sum_{j=1}^S \left[\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K+1) \right] \\ &= -S + \frac{1}{2} \sum_{j=1}^S \left[\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K-1) \right] \\ &\approx -S + \sqrt{K\gamma(K)} \sum_{j=1}^S \sqrt{r_j} \\ &= -S + \frac{(\sum_{j=1}^S \sqrt{r_j})^2}{\sum_{j=1}^S r_j - \bar{\lambda}} \\ &= D_1(\bar{\lambda}, \mathbf{r}). \end{aligned}$$

□

APPENDIX C. PROOFS OF THE RESULTS IN SECTION 4**C.1. Proof of lemma 4.1.**

Proof. Let us first prove that $\bar{\lambda}^{OPT} < \bar{\lambda}^{NE}$:

$$\begin{aligned} \bar{\lambda}^{OPT} < \bar{\lambda}^{NE} &\iff \sqrt{r_1 r_2} > \frac{2r_1}{\sqrt{(K-1)^2 + 4Kr_1/r_2} - (K-1)} \\ &\iff \sqrt{r_1 r_2} \sqrt{(K-1)^2 + 4Kr_1/r_2} > \sqrt{r_1} [2\sqrt{r_1} + (K-1)\sqrt{r_2}] \\ &\iff 4Kr_1 > 4r_1 + 4(K-1)\sqrt{r_1 r_2} \\ &\iff r_1 > r_2. \end{aligned}$$

We now turn to the second part of the proof. According to proposition A.1, the centralized setting uses only the fast servers (S_1 servers of capacity r_1) for all values of $\bar{\lambda}$ such that $W_2(1, \frac{1}{r_2}) \leq 0$. It yields

$$\bar{\lambda} \leq (S_1 r_1 + S_2 r_2) - \sqrt{r_2}(S_1 \sqrt{r_1} - S_2 \sqrt{r_2})$$

which is equivalent to $\bar{\lambda} \leq \bar{\lambda}^{OPT}$.

Similarly, we know from proposition A.1, that the decentralized setting starts using the second group of servers if and only if

$$\begin{aligned}\bar{\lambda} &\geq S_1 r_1 + S_2 r_2 - \sum_{s=1}^{S_1+S_2} \frac{2r_s}{\sqrt{(K-1)^2 + 4Kr_s/r_2} - (K-1)} \\ &= S_1 r_1 + S_2 r_2 - \sum_{s=1}^2 S_s \frac{2r_s}{\sqrt{(K-1)^2 + 4Kr_s/r_2} - (K-1)} \\ &= S_1 r_1 - S_1 \frac{2r_1}{\sqrt{(K-1)^2 + 4Kr_1/r_2} - (K-1)},\end{aligned}$$

which is equivalent to $\bar{\lambda} \geq \bar{\lambda}^{NE}$, as claimed. $\square$

C.2. Proof of corollary 4.1.

Proof. We first prove the results for the centralized setting. For $\bar{\lambda} < \bar{\lambda}^{OPT}$, we know from lemma 4.1 that the centralized setting uses only the first group of servers, and we thus have

$$D_1(\bar{\lambda}, \mathbf{r}) = \sum_{j=1}^{S_1} \frac{\frac{\bar{\lambda}}{S_1}}{r_1 - \frac{\bar{\lambda}}{S_1}} = \frac{\bar{\lambda}}{r_1 - \frac{\bar{\lambda}}{S_1}}.$$

For $\bar{\lambda} \geq \bar{\lambda}^{OPT}$, we know from lemma 4.1 that the centralized setting uses all the servers. According to the KKT conditions,

$$(14) \quad \frac{r_1}{(r_1 - y_1)^2} = \frac{r_2}{(r_2 - y_2)^2}$$

where $\bar{\lambda} = S_1 y_1 + S_2 y_2$. After some algebra, it yields

$$(15) \quad y_1 = \sqrt{r_1} \frac{\bar{\lambda} - S_2 \sqrt{r_2} (\sqrt{r_2} - \sqrt{r_1})}{S_1 \sqrt{r_1} + S_2 \sqrt{r_2}}.$$

With the constraint of $S_1 y_1 + S_2 y_2 = \bar{\lambda}$ and (14), we obtain

$$\begin{aligned}D_1(\bar{\lambda}, \mathbf{r}) &= S_1 \frac{y_1}{r_1 - y_1} + S_2 \frac{y_2}{r_2 - y_2} \\ &= S_1 \frac{y_1}{r_1 - y_1} + (\bar{\lambda} - S_1 y_1) \frac{1}{r_2 - y_2} \\ &= S_1 \frac{y_1}{r_1 - y_1} + (\bar{\lambda} - S_1 y_1) \sqrt{\frac{r_1}{r_2}} \frac{1}{r_1 - y_1} \\ &= \left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}} \right) \right] \frac{1}{r_1 - y_1},\end{aligned}$$

where y_1 is given in (15).

Let us now consider the decentralized setting. For $\bar{\lambda} < \bar{\lambda}^{NE}$, we know from lemma 4.1 that the decentralized setting uses only the servers of the first class, and we thus have $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) = \frac{\bar{\lambda}}{r_1 - \frac{\bar{\lambda}}{S_1}}$. For $\bar{\lambda} \geq \bar{\lambda}^{NE}$, all the servers are used and using (10) we obtain

$$D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) = \sum_{j=1}^{S_1+S_2} \frac{y_j(K)}{r_j} \left(1 - \frac{y_j(K)}{r_j}\right)^{-1} = \frac{1}{2} \sum_{j=1}^{S_1+S_2} \left[\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K+1) \right]$$

and thus

$$D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) = \frac{1}{2} \sum_{j=1}^2 S_j \left[\sqrt{(K-1)^2 + 4K\gamma(K)r_j} - (K+1) \right].$$

□

C.3. Proof of proposition 4.1. We shall break-down the proof of the proposition in three parts according to the three intervals which define the behaviour of the ratio of the social costs. In the following two lemmata, we present that the ratio is increasing in one of the intervals and decreasing in the other.

Lemma C.1. *The ratio $\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is strictly increasing over the interval $(\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$.*

Proof. In order to prove that $\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is an increasing function of $\bar{\lambda}$ for $\bar{\lambda} \in (\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$ we shall prove that

$$\frac{\frac{\partial}{\partial \bar{\lambda}} \left(D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) \right)}{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})} > \frac{\frac{\partial}{\partial \bar{\lambda}} (D_1(\bar{\lambda}, \mathbf{r}))}{D_1(\bar{\lambda}, \mathbf{r})}.$$

Since $\bar{\lambda} < \bar{\lambda}^{NE}$, we have $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) = \frac{\bar{\lambda}}{r_1 - \frac{\bar{\lambda}}{S_1}}$ and

$$\frac{\partial D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{\partial \bar{\lambda}} = \frac{r_1}{(r_1 - \frac{\bar{\lambda}}{S_1})^2} = \frac{r_1}{\bar{\lambda}(r_1 - \frac{\bar{\lambda}}{S_1})} D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}),$$

from which we deduce that

$$\frac{\frac{\partial}{\partial \bar{\lambda}} D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})} = \frac{r_1}{\bar{\lambda}(r_1 - \frac{\bar{\lambda}}{S_1})} = \frac{r_1}{\bar{\lambda}^2} D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}).$$

For the centralized setting, $\bar{\lambda} \geq \bar{\lambda}^{OPT}$ means that

$$D_1(\bar{\lambda}, \mathbf{r}) = S_1 \frac{y_1}{r_1 - y_1} + S_2 \frac{y_2}{r_2 - y_2}.$$

According to (14), the derivative of $D_1(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})$ gets modified as follows:

$$\frac{\partial D_1(\bar{\lambda}, \mathbf{r})}{\partial \bar{\lambda}} = S_1 \frac{r_1 y_1'}{(r_1 - y_1)^2} + S_2 \frac{r_2 y_2'}{(r_2 - y_2)^2} = (S_1 y_1' + S_2 y_2') \frac{r_1}{(r_1 - y_1)^2}.$$

The constraint $S_1 y_1 + S_2 y_2 = \bar{\lambda}$ implies that $S_1 y_1' + S_2 y_2' = 1$ and hence

$$\frac{\partial D_1(\bar{\lambda}, \mathbf{r})}{\partial \bar{\lambda}} = \frac{r_1}{(r_1 - y_1)^2}.$$

It yields

$$\frac{\frac{\partial}{\partial \bar{\lambda}} D_1(\bar{\lambda}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = \frac{r_1}{(r_1 - y_1)^2} \frac{1}{[D_1(\bar{\lambda}, \mathbf{r})]^2} D_1(\bar{\lambda}, \mathbf{r}) = \frac{r_1}{\left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}}\right)\right]^2} D_1(\bar{\lambda}, \mathbf{r}).$$

As a result, for $\bar{\lambda} \in (\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$, $\frac{\frac{\partial}{\partial \bar{\lambda}} D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})} > \frac{\frac{\partial}{\partial \bar{\lambda}} D_1(\bar{\lambda}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is equivalent to

$$\frac{r_1}{\bar{\lambda}^2} D_K\left(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r}\right) > \frac{r_1}{\left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}}\right)\right]^2} D_1(\bar{\lambda}, \mathbf{r})$$

and

$$\frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} > \frac{\bar{\lambda}^2}{\left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}}\right)\right]^2}.$$

Now, we assume that there exist $\bar{\lambda} \in (\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$ such that $\frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is not increasing with $\bar{\lambda}$.

Since $\frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} \geq 1$, it results

$$1 \leq \frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} \leq \frac{\bar{\lambda}^2}{\left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}}\right)\right]^2}.$$

Since we have that

$$\frac{\bar{\lambda}^2}{\left[\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}}\right)\right]^2} \geq 1,$$

then

$$\frac{\bar{\lambda}}{\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 y_1 \left(1 - \sqrt{\frac{r_1}{r_2}}\right)} \geq 1$$

because the numerator and the denominator are both positive. Thus, with (15), the above yields

$$1 \leq \frac{\bar{\lambda}}{\bar{\lambda} \sqrt{\frac{r_1}{r_2}} + S_1 \sqrt{r_1} \frac{\bar{\lambda} - S_2 \sqrt{r_2} (\sqrt{r_2} - \sqrt{r_1})}{S_1 \sqrt{r_1} + S_2 \sqrt{r_2}} \left(1 - \sqrt{\frac{r_1}{r_2}}\right)} \leq \frac{\bar{\lambda} (S_1 \sqrt{r_1} + S_2 \sqrt{r_2})}{\bar{\lambda} \sqrt{r_1} (S_1 + S_2) - S_1 S_2 \sqrt{r_1} (\sqrt{r_2} - \sqrt{r_1})^2}.$$

This is equivalent to

$$\bar{\lambda} \sqrt{r_1} (S_1 + S_2) - S_1 S_2 \sqrt{r_1} (\sqrt{r_2} - \sqrt{r_1})^2 \leq \bar{\lambda} (S_1 \sqrt{r_1} + S_2 \sqrt{r_2})$$

and after rearranging both sides of the expression, we arrive at the following condition

$$\bar{\lambda} \leq S_1 \sqrt{r_1} (\sqrt{r_1} - \sqrt{r_2}) = \bar{\lambda}^{OPT},$$

that it is a contradiction since $\bar{\lambda} \in (\bar{\lambda}^{OPT}, \bar{\lambda}^{NEP})$.

□

Lemma C.2. *The ratio $\frac{D_K(\frac{\bar{\lambda}}{K} \mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is strictly decreasing over the interval $(\bar{\lambda}^{NE}, \bar{r})$.*

Proof. In the interval $(\bar{\lambda}^{NE}, \bar{r})$ we know that all servers are used. Thus, according to proposition 3.1, the ratio $\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is decreasing as a function of $\bar{\lambda}$. $\square$

The proof of proposition 4.1 results from corollary 4.1, and lemmata C.1 and C.2.

C.4. Proof of lemma 4.2.

Proof. First, we state that this function is continuous in $[0, \bar{\lambda}^{OPT})$, $(\bar{\lambda}^{OPT}, \bar{\lambda}^{NE})$ and $(\bar{\lambda}^{NE}, \bar{r})$ which follows from the definitions of $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})$ and $D_1(\bar{\lambda}, \mathbf{r})$ in corollary 4.1.

Now, we show that $\frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})}$ is continuous in $\bar{\lambda}^{OPT}$ and $\bar{\lambda}^{NE}$:

$$\begin{aligned} \bullet \lim_{\bar{\lambda} \rightarrow \bar{\lambda}^{OPT}-} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} &= \lim_{\bar{\lambda} \rightarrow \bar{\lambda}^{OPT}+} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = \frac{D_K(\frac{\bar{\lambda}^{OPT}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}^{OPT}, \mathbf{r})} = 1, \\ \bullet \lim_{\bar{\lambda} \rightarrow \bar{\lambda}^{NE}-} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} &= \lim_{\bar{\lambda} \rightarrow \bar{\lambda}^{NE}+} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = \frac{D_K(\frac{\bar{\lambda}^{NE}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}^{NE}, \mathbf{r})} = \frac{\frac{1}{2}S_1[\sqrt{(K-1)^2+4Kr_1/r_2}-(K+1)]}{\frac{(S_1\sqrt{r_1}+S_2\sqrt{r_2})^2}{2S_1r_1}-(S_1+S_2)} \\ &\quad \frac{1}{\sqrt{(K-1)^2+4Kr_1/r_2}-(K-1)} + S_2r_2 \end{aligned}$$

$\square$

C.5. Proof of lemma 4.3.

Proof. First, we modify $I_K(\alpha, \beta)$ and we obtain

$$I_K(\alpha, \beta) = \frac{1}{2} \frac{(x-2)(\frac{1}{\alpha} + \frac{2\beta}{x})}{\frac{1}{\alpha}(2\sqrt{\beta}-1-\frac{2\beta}{x}) + \beta(1-\frac{2}{x})},$$

where $x = \sqrt{(K-1)^2 + 4K\beta} - (K-1)$.

We now show that the derivative of $I_K(\alpha, \beta)$ with respect to α is negative using that the derivative of $\frac{1}{\alpha}$ with respect to α is negative:

$$\frac{\partial I_K(\alpha, \beta)}{\partial \alpha} = \frac{\frac{1}{2}(\frac{1}{\alpha})' \left[\beta(1-\frac{2}{x}) - \frac{2\beta}{x}(2\sqrt{\beta}-1-\frac{2\beta}{x}) \right]}{\left[\frac{1}{\alpha}(2\sqrt{\beta}-1-\frac{2\beta}{x}) + \beta(1-\frac{2}{x}) \right]^2} = \frac{\frac{1}{2}(\frac{1}{\alpha})' \left(1 - \frac{2\beta}{x} \right)^2}{\left[\frac{1}{\alpha}(2\sqrt{\beta}-1-\frac{2\beta}{x}) + \beta(1-\frac{2}{x}) \right]^2} < 0.$$

$\square$

C.6. Proof of lemma 4.4.

Proof. From (6), we observe that $\lim_{S \rightarrow \infty} PoA(K, S) = \sup_{\beta} \lim_{S \rightarrow \infty} I_K\left(\frac{1}{S-1}, \beta\right)$. In order to compute the limit of the PoA, we shall first compute the limit of I_K and then we shall compute the supremum. Let $x = \sqrt{(K-1)^2 + 4K\beta} - (K-1)$. We rewrite (5) as

$$I_K\left(\frac{1}{S-1}, \beta\right) = \frac{1}{2} \frac{x-2}{\frac{(S-1+\sqrt{\beta})^2}{S-1+\frac{2\beta}{x}} - S}.$$

For large S , we show that

$$\frac{(S-1+\sqrt{\beta})^2}{S-1+\frac{2\beta}{y}} - S \approx 2\sqrt{\beta} - 1 - \frac{2\beta}{x}$$

as follows:

$$\begin{aligned}
\frac{(S-1+\sqrt{\beta})^2}{S-1+\frac{2\beta}{x}} &\approx \frac{(S-1)^2+2\sqrt{\beta}(S-1)+\beta}{(S-1)\left(1+\frac{2\beta}{(S-1)x}\right)} \\
&\approx \left(S-1+2\sqrt{\beta}+\frac{\beta}{(S-1)}\right) \cdot \left(1-\frac{2\beta}{(S-1)x}\right) \\
&\approx S-1+2\sqrt{\beta}-\frac{2\beta}{x}.
\end{aligned}$$

Now that we have computed the limit of I_K as $S \rightarrow \infty$, we shall compute the supremum with respect to β . In order to do this, we shall show that the limit computed previously is an increasing function of β .

Denote $F_K(\beta) = \lim_{S \rightarrow \infty} I_K(1/(S-1), \beta)$. We shall show that it is an increasing function of β . Let $y = \sqrt{(K-1)^2 + 4K\beta}$. We write

$$F_K(\beta) = \frac{1}{2} \frac{y - (K+1)}{2\sqrt{\beta} - 1 - \frac{y+K-1}{2K}}.$$

Using that $\beta > 1$ we show that the derivative of $F_K(\beta)$ with respect to β is positive:

$$\frac{\partial F_K(\beta)}{\partial \beta} > 0 \iff y' \left[2\sqrt{\beta} - 1 - \frac{y+K-1}{2K} \right] - (y - (K+1)) \left(\frac{1}{\sqrt{\beta}} - \frac{1}{2K} y' \right) > 0.$$

This expression can be simplified as follows:

$$2y'(\sqrt{\beta} - 1) - (y - (K+1)) \frac{1}{\sqrt{\beta}} > 0.$$

Since y' is the derivative of $y = \sqrt{(K-1)^2 + 4K\beta}$ with respect to β , we derive it and substituting in the above expression it yields:

$$\frac{4K(\sqrt{\beta} - 1)}{y} - (y - (K+1)) \frac{1}{\sqrt{\beta}} > 0.$$

Multiplying by $b\sqrt{\beta}$ and using that $y = \sqrt{(K-1)^2 + 4K\beta}$ we obtain

$$\begin{aligned}
4K\sqrt{\beta}(\sqrt{\beta} - 1) - [(K-1)^2 + 4K\beta - (K+1)y] &> 0 \\
\iff 4K\beta - 4K\sqrt{\beta} - [(K+1)^2 + 4K(\beta - 1) - (K+1)y] &> 0.
\end{aligned}$$

This expression can be simplified and it results:

$$-4K(\sqrt{\beta} - 1) + (K+1)(y - (K+1)) > 0.$$

Now, we notice that $(y - (K+1)) = \frac{4K(\beta-1)}{y+K+1}$ and thus

$$-4K(\sqrt{\beta} - 1) + (K+1) \frac{4K(\beta-1)}{y+K+1} > 0.$$

Since $\beta - 1 = (\sqrt{\beta} - 1)(\sqrt{\beta} + 1)$, we get

$$-1 + (K+1) \frac{(\sqrt{\beta} + 1)}{y+K+1} > 0.$$

If we multiply this expression by $y + K + 1$, it results

$$\sqrt{\beta}(K + 1) - y > 0.$$

Now, we divide and multiply this expression by $\sqrt{\beta}(K + 1) + y$ and we obtain

$$\frac{\beta(K + 1)^2 - [(K + 1)^2 + 4K(\beta - 1)]}{\sqrt{\beta}(K + 1) + y} > 0.$$

This expression can be simplified as follows:

$$(\beta - 1)(K + 1)^2 - 4K(\beta - 1) > 0 \iff (K - 1)^2(\beta - 1) > 0,$$

and thus we have proved that $F_K(\beta)$ is a decreasing function of β . □

APPENDIX D. PROOFS OF THE RESULTS IN SECTION 5

D.1. Proof of proposition 5.1.

Proof. Let n be the number of servers used by the decentralized setting for an infinity number of dispatchers, with $n = 1, \dots, S$. We observe that

$$\sqrt{(K - 1)^2 + 4K\gamma(K)r_j} \rightarrow 2\gamma(\infty)r_j, \text{ when } K \rightarrow \infty.$$

We use this property in (9) to show that $\gamma(\infty) = \frac{n}{\sum_{j=1}^n r_j - \bar{\lambda}}$ and for the expression of $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})$ given in lemma B.1, we obtain that

$$\lim_{K \rightarrow \infty} D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r}) = \frac{1}{2} \sum_{j=1}^n [2r_j\gamma(\infty) - 2] = \frac{n \sum_{j=1}^n r_j}{\sum_{j=1}^n r_j - \bar{\lambda}} - n.$$

Now, we evaluate $D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})/D_1(\bar{\lambda}, \mathbf{r})$ when $K \rightarrow \infty$ and it yields to

$$\lim_{K \rightarrow \infty} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = \frac{\frac{n \sum_{j=1}^n r_j}{\sum_{j=1}^n r_j - \bar{\lambda}} - n}{\frac{(\sum_{j=1}^n \sqrt{r_j})^2}{\sum_{j=1}^n r_j - n} - n} = \frac{\bar{\lambda}}{\bar{\lambda} - \sum_{j=1}^n r_j + \frac{1}{n} \left(\sum_{j=1}^n \sqrt{r_j} \right)^2},$$

and we know that in heavy-traffic all the servers are used, i.e., when $\bar{\lambda} \rightarrow \bar{r}$, we have $n = S$. Hence, it results

$$\lim_{\rho \rightarrow 1} \lim_{K \rightarrow \infty} \frac{D_K(\frac{\bar{\lambda}}{K}\mathbf{e}, \mathbf{r})}{D_1(\bar{\lambda}, \mathbf{r})} = \frac{S \bar{r}}{\left(\sum_{i=1}^S \sqrt{\bar{r}_i} \right)^2}.$$

□