Fair Linking Mechanisms for Resource Allocation with Correlated Player Types

Agustín Santos[1], Antonio Fernández Anta[1(✉)], José A. Cuesta[2], and Luis López Fernández[3]

[1] IMDEA Networks Institute, Madrid, Spain
antonio.fernandez@imdea.org
[2] GISC, Mathematics Department, Universidad Carlos III, Madrid, Spain
[3] FUN-LAB, Universidad Rey Juan Carlos, Madrid, Spain

Abstract. Resource allocation is one of the most relevant problems in the area of Mechanism Design for computing systems. Devising algorithms capable of providing efficient and fair allocation is the objective of many previous research efforts. Usually, the mechanisms they propose use payments in order to deal with selfishness. Since using payments is undesirable in some contexts, a family of mechanisms without payments is proposed in this paper. These mechanisms extend the Linking Mechanism of Jackson and Sonnenschein introducing a generic concept of fairness with correlated preferences. We prove that these mechanisms have good incentive, fairness, and efficiency properties. To conclude, we provide an algorithm, based on the mechanisms, that could be used in practical computing environments.

Keywords: Linking mechanism · Fairness · Resource allocation

1 Introduction

The success of the Internet has made the problem of resource allocation to emerge in many versions like, for example, deciding which peer must receive bandwidth or disk in a file sharing P2P system [1], or deciding to which computational task some CPU is assigned in a collaborative distributed environment [2]. The problem may also appear with a negative formulation (i.e., instead of deciding who shall receive a resource, the problem is deciding who shall not receive it).

In all these scenarios, it is very important to conceive mechanisms that achieve efficient and fair resource allocation even when players present selfish or non-rational behavior. With that purpose, a number of interesting protocols and mechanisms based on Game Theory concepts [3,4] have been proposed.

This research was supported in part by the Comunidad de Madrid grant S2009TIC-1692, Spanish MICINN/MINECO grants FIS2011-22449 (PRODIEVO) and TEC2011-29688-C02-01 (E2Net), and National Natural Science Foundation of China grant 61020106002.

G. Noubir and M. Raynal (Eds.): NETYS 2014, LNCS 8593, pp. 70–83, 2014.
DOI: 10.1007/978-3-319-09581-3_6

In such works, it is often assumed that players can transfer their utilities (i.e., use payments). However, there are many systems in which this assumption is not realistic. Recently, some mechanisms without payments have been proposed, like those of Procaccia and Tennenholtz [5], or the seminal work of Jackson and Sonnenschein [6,7] in which a new type of mechanism (called *Linking Mechanism*) is proposed. A Linking Mechanism, instead of offering incentives or payments to players, limits the spectrum of players' responses to a probability distribution known by the game designer. The objective of this paper is to explore and extend Linking Mechanisms, introducing a wide spectrum of fairness concepts, while preserving all the original properties.

State of the Art. Mechanism Design has been gaining increasing popularity in distributed computing during the last few years (see, e.g., [8–10]). Even though the mechanisms proposed in these works are interesting, they are usually based on payment systems. Deploying such payment system in practice is often difficult. For this reason, mechanisms without payments have also been proposed. Related literature could be found in economics on cooperation [11,12] or similar problems in P2P systems such as reputation [13] and artificial currencies [14]. The work closest to our own, and in which we have based our proposal, is the *Linking Mechanism* proposed by Jackson and Sonnenschein [6,7]. Related to this work, Engelmann and Grimm [15] presents experimental research on linking mechanisms. An algorithm called QPQ (Quid Pro Quo) [16] has been proposed as an application of this kind of mechanisms to distribute task executions fairly among independent players.

QPQ reflects the main idea behind the concept of linking mechanism: when a game consists of multiple instances of the same basic decision problem (e.g., saying yes or no, choosing among a number of discrete options), it is possible to define selfishness-resistant algorithms by restricting the players' responses to a given distribution. Hence, in that case, the frequency with which a player declares a particular decision is established beforehand. Based on this, QPQ presents quite relevant features as the fact of not requiring payments, the flexibility on the definitions of the utility functions of the players, its applicability in iterative (i.e. repeated) games, the lack of central control authority, etc. While QPQ presents some very interesting properties, it only guarantees fairness and efficiency when users behave independently on each other. Nevertheless, this does not need to be the case in real environments, where users may have correlated preferences. The problem of fairness among players has been widely analyzed in the game theory literature and a wide range of fairness concept has been proposed, but, as far as we know, there is no fair linking mechanisms when players have correlated preferences. This motivates the research proposed in this paper.

Contributions. Our contributions are twofold. On the one hand, we have extended the idea of Linking Mechanism introducing fairness, while preserving desirable properties, like efficiency, truthful reporting, incentive compatibility, etc. On the other hand, we propose an algorithm based on these mechanisms that we expect to be used in practical scenarios.

In our model, fairness is a key element introduced to compensate current sacrifices in future iterations. Due to the large number of notions of fairness that could be defined, it is difficult to find a general model that encompasses any approach. Fairness is, in general, an elusive concept that can be seen from many different perspectives. In this work we have proposed a generic fairness definition, which we hope will serve as a reference to wider models. Hence, our contribution is clear: to the best of our knowledge, no other previous research work has offered a linking mechanism providing fair and efficient decisions.

In addition, from a theoretical perspective, we contribute to the progress of the state-of-the-art by proposing a mathematical framework suitable for proving all claimed algorithmic properties. This framework is inspired on previous work on theoretical economics but, as far as we know, it has never been adapted to the specific peculiarities of distributed computing (at least not to solve the resource allocation problem). This technique has proven to be extremely powerful for our specific problem, but it can be re-used in other scenarios with similar assumptions.

Based on the theoretical results, we propose a realization of the mechanism suitable for being implemented as an iterative game in real distributed environments. Unlike in the original linking mechanism, this algorithm does not need to know the probability distribution of the players' responses. We show that this realization does not require central entities and that its computational cost is affordable for current state of the art networks and devices. In addition, through simulations, we confirm the stability of the algorithm demonstrating that few iterations on a repeated game are enough for making the mechanism to converge to a fair equilibrium even when the players' distributions are strongly correlated.

To illustrate the application of this mechanism in a real environment, consider a P2P system used for computation to which requests arrive continuously. When a request arrives, the computational cost of processing it at a given node of the system will depend on the load of the node. Our mechanism could be used to enforce that all nodes process the same proportion of requests, while the total computational cost is minimized.

2 Model and Definitions

We start by presenting the usual mathematical framework for mechanism design and then we formally define the specific problem we face in this paper.

Mechanism Design Concepts. The following provides the usual theoretic framework that will be later applied to our problem. We assume that there are n players. The set of players is $N = \{1, 2, \dots, n\}$. Players are risk-neutral. The alternative or outcome set of the game played is D. In a general setting, D could be defined over $\Delta(N)$[1], but in this paper we define $D = N$ so that the outcome $d \in D$ is the player to whom the resource will be allocated.

[1] We denote by $\Delta(S)$ the set of all probability distribution over some set S.

Prior to making the collective choice in the game, each player privately observes her preferences over the alternatives in D. This is modeled by assuming that player i privately observes a parameter or signal θ_i that determines her preferences. (For instance, in resource allocation, θ_i could represent the value player i assigns to the resource.) For a given player i, we say that θ_i is the player type. The set of possible types of player i is Θ_i. We denote by $\theta = (\theta_1, \theta_2, \ldots, \theta_n)$ the vector of player types. The set of all possible vectors is $\Theta = \Theta_1 \times \Theta_2 \times \ldots \Theta_n$. We denote by θ_{-i} the vector obtained by removing θ_i from θ.

We denote by $\Pi = \Delta(\Theta)$ the set of all probability distributions over Θ. It is assumed that there is a common prior distribution $\pi \in \Pi$ that is shared by all the players. We denote by $\pi_i \in \Delta(\Theta_i)$ the marginal probability of θ_i. We define $\beta_i(\theta_{-i}|\theta_i)$ as the conditional probability distribution of θ_{-i} given θ_i. That is, for any possible type $\theta_i \in \Theta_i$, $\beta_i(\cdot|\theta_i)$ specifies a probability distribution over the set Θ_{-i} representing what player i would believe about the types of the other players if her own type were θ_i. Beliefs $(\beta_i)_{i \in N}$ are *consistent*, since individual belief functions β_i can all be derived from the common prior π. This implies that $\pi(\theta_{-i}|\theta_i) = \beta_i(\theta_{-i}|\theta_i)$.

Individual players have preferences over outcomes, which are represented by a utility function $u_i(d, \theta_i) \in \mathbb{R}$ defined over all $d \in D$ and $\theta_i \in \Theta_i$.

The set of outcomes D, the set of players N, the type sets in Θ, the common prior distribution $\pi \in \Pi$, and the payoff functions $u_i, i \in N$ are assumed to be *common knowledge* among all the players. The game rules defined by a specific mechanism are also common knowledge. However, the specific value θ_i observed by player i is *private information* of player i.

A strategy for the player i is any map $\sigma_i : \Theta_i \rightarrow \Delta(\Theta_i)$, where $\sigma_i(\hat{\theta}_i|\theta_i)$ is the conditional probability that the player reports $\hat{\theta}_i$ when her true type is θ_i. A reporting strategy σ_i is *truthful* if for every pair $(\hat{\theta}_i, \theta_i)$, $\sigma_i(\hat{\theta}_i|\theta_i) = 1$ if $\hat{\theta}_i = \theta_i$ and 0 otherwise. As usually done, we will use $\hat{\theta}_i$ to denote the reported type and θ_i the actual type.

Given that the prior distribution π is known, player i can not change it. Hence, we say that a player i has a *limited strategy space*, since her strategy can not change the beliefs of other players. Intuitively, player i has a limited strategy space if beliefs over reports are the same as actual beliefs.

For a given Bayesian mechanism $\langle \Theta, g \rangle$ we shall write $q_i(\cdot|\theta_i)$ for player i's interim probability density function on D conditional on player i's type being θ_i.

In this paper, we are looking for a mechanism $\langle \Theta, g \rangle$, where $g(\cdot)$ is the decision function, without utility transfers (payments) and that implements some social choice function f under some equilibrium when the induced game is Bayesian. In addition, we introduce fairness as a key tool to compensate or reward players. We call this kind of mechanisms as *Quid Pro Quo Mechanisms* (QPQ).

Fairness. In our model, we use fairness as a very abstract concept. For us, fairness is the property of balancing in expectation some game parameters (modelled with a real function) among all players. Our model was originally built with two examples in mind: *fairness in utility* ("players have same expected utility") and *fairness in assignment* ("same expected number of assignments"). But these two examples are just special cases of our model. Additionally, we have contemplated the possibility that some scenarios require allocations other than equiproportional; or than the game must be constrained to several fairness concepts at the same time. All of this is modelled introducing a set of functions $\eta_{i,l} : \theta \rightarrow \mathbb{R}$ and ratios $\delta_{i,l}$, all defined for each player $i \in N$ and for each fairness concept $l = 1, \cdots, m$ (m is the number of fairness concepts). The function $\eta_{i,l}$ represents a fairness concept. For instance, for *fairness in assignment* this function could be defined as $\eta_{i,l}(\theta) = 1$. Similarly, *fairness in utility* is applied when $\eta_{i,l}(\theta) = \theta_i$. On the other hand, $\delta_{i,l}$ is the ratio for player i when fairness l is applied. Typically, this ratio is $\delta_{i,l} = \frac{1}{n}$. Then, formally, our concept of fairness is defined as follows.

Definition 1 (Fairness). *Given functions $\eta_{i,l} : \Theta \rightarrow \mathbb{R}$, and values $\delta_{i,l}$, we say that a mechanism $\langle \Theta, g \rangle$ is fair (or η-fair) when, for all $i \in N$ and $l = 1, \cdots, m$,*

$$\int_{\Theta} \eta_{i,l}(\theta) q_i(\theta) \ d\pi(\theta) = \delta_{i,l} \sum_{j \in N} \int_{\Theta} \eta_{j,l}(\theta) q_j(\theta) \ d\pi(\theta) \tag{1}$$

In this paper, we deal mathematically with this general concept of fairness, but for the algorithm and simulations we used a particular concept of fairness, where players will have *equal number* of allocated resources (in expectation).

Resource Allocation Problem. We now formally define the problem we study in this work. Intuitively, the problem is like a repeated single-unit auction, where the mechanism that decides how to allocate the resource in each auction is a QPQ mechanism. Hence there are no payments and the allocation must satisfy a notion of fairness.

The problem of resource allocation is a tuple $\langle R, N, \Theta \rangle$ where, N and Θ are as defined above, and $R = \{r_1, r_2, \ldots\}$ is the ordered set of resources that have to be allocated by the system over time. Resources are received by the system in their order in R, they are independent among them, and the system must allocate resource r_k to a single player before receiving resource r_{k+1}. R is assumed to be infinite.

As was mentioned previously, in this problem the outcome set is $D = N$, where an outcome of $d \in D$ for resource r_k means that r_k is allocated to player d. In [16], we have proposed a QPQ algorithm that implements this function when the type of players follow mutually independent distributions. As in that work, we assume here that the type of each player is normalized using a *Probability Integral Transform* (PIT), so that it takes real values in the interval $[0, 1]$ and follows a uniform distribution within that support. Hence, we assume that $\Theta_i = [0, 1]$. Finally, as mentioned, we assume that players have a limited space strategy (i.e., π is known a priori and cannot be changed by the players).

The social choice function (scf) $g(\cdot)$ we are looking for is one that optimizes the social utility restricted by fairness conditions. The social choice function must be the solution to the following equation,

$$\begin{aligned}
&\max_{g} \ \{\textstyle\sum_{i\in N}\int_{\Theta} u_i(d(\theta),\theta_i)\ q_i(\theta)\ d\pi(\theta)\}\\
&\text{s.t.,}\\
&\textstyle\int_{\Theta}\eta_{i,l}(\theta)q_i(\theta)\ d\pi(\theta) = \delta_{i,l}\sum_{j\in N}\int_{\Theta}\eta_{j,l}(\theta)q_j(\theta)\ d\pi(\theta), l = 1,\cdots,m
\end{aligned} \tag{2}$$

As an example, we study the fairness concept where each player i will receive a *proportional number of resources* δ_i. Hence, we obtain that the scf is the solution of the following equation.

$$\begin{aligned}
&\max_{g}\{\sum_{i\in N}\int_{\Theta} u_i(d,\theta_i)\ q_j(\theta)\ d\pi(\theta)\}\\
&\text{s.t.}\\
&\int_{\Theta} q_j(\theta)\ d\pi(\theta) = \delta_i, \forall i \in N.
\end{aligned} \tag{3}$$

Another fairness concept that we study as an instance of this framework is *players with proportional utility.* Under this fairness concept every player will obtain a proportional expected utility. The equations are similar in this case.

$$\begin{aligned}
&\max_{g}\{\sum_{i\in N}\int_{\Theta} u_i(d,\theta_i)\ q_j(\theta)\ d\pi(\theta)\}\\
&\text{s.t.}\\
&\int_{\Theta} u_i(d,\theta_i)\ q_j(\theta)\ d\pi(\theta) = \delta_i\sum_{j\in N}\int_{\Theta}\ u_j(d,\theta_j)\ q_j(\theta)\ d\pi(\theta), \forall i \in N.
\end{aligned} \tag{4}$$

Without loss of generality, we can define the utility of a player i as follows,

$$u_i(d,\theta_i) = \begin{cases}\theta_i & \text{if } d = i,\\ 0 & \text{otherwise.}\end{cases} \tag{5}$$

In this paper, we are interested in dynamic mechanisms where truth-telling is a Bayesian equilibrium of the static QPQ mechanism. In that case we call the QPQ mechanism Bayesian incentive compatible. That means that a player obtains a higher utility when reporting truthfully.

3 The Fair Quid Pro Quo Mechanism

With the above definitions, we now derive QPQ Mechanisms that implement the social choice functions given by Eqs. 3 and 4 under equilibrium, as special cases of the solution to Eq. 2.

Theorem 1. *The QPQ Mechanism that implements the social function 2 with η-fairness is a set of functions $\psi = (\psi_1, \cdots \psi_n)$ that defines a line $y = \psi_i(x)$ for each player i with deterministic assignment $d = g_\psi(\theta) = \underset{i \in N}{\text{argmax}}\ (\psi_i(\theta))$ (except at some points where the decision is indifferent).*

Proof. The problem we aim to solve is to find the decision function g that maximizes

$$\int_\Theta \sum_{i \in N} \theta_i \ q_i(\theta) \ d\pi(\theta) \tag{6}$$

under the constraints given in Eq. 2. Using Lagrange multipliers, this is tantamount to maximizing the functional

$$\mathcal{F}[q] \equiv \int_\Theta \sum_{i \in N} \theta_i \ q_i(\theta) \ d\pi(\theta) + \sum_{k \in N} \sum_{l=1}^{m} \lambda_{k,l} \int_\Theta \left\{ \eta_{k,l}(\theta) q_k(\theta) - \delta_{k,l} \sum_{j \in N} \eta_{j,l}(\theta) \ q_j(\theta) \right\} d\pi(\theta), \tag{7}$$

which can be rewritten

$$\mathcal{F}[q] = \int_\Theta \sum_{i \in N} \psi_i(\theta) q_i(\theta) \ d\pi(\theta), \tag{8}$$

where

$$\psi_i(\theta) \equiv \theta_i + \sum_{l=1}^{m} \lambda_{i,l} \ \eta_{i,l}(\theta) - \sum_{k \in N} \sum_{l=1}^{m} \lambda_{k,l} \ \delta_{k,l} \ \eta_{i,l}(\theta). \tag{9}$$

Since $0 \le q_i(\theta) \le 1$ and $\sum_{i \in N} q_i(\theta) = 1$ for all $\theta \in \Theta$, then for each $\theta \in \Theta$,

$$\sum_{i \in N} \psi_i(\theta) q_i(\theta) \le \psi_j(\theta) \tag{10}$$

if $j \in N$ is such that $\psi_j(\theta) > \psi_k(\theta)$ for all $k \neq j$. The upper bound is reached if, and only if, for that value of θ we have $q_j(\theta) = 1$ and $q_k(\theta) = 0$ for all $k \neq j$.

If, on the other hand, $j_1, \ldots, j_r$ are such that $\psi_{j_1}(\theta) = \cdots = \psi_{j_r}(\theta) > \psi_k(\theta)$ for all $k \neq j_1, \ldots, j_r$, then the upper bound is $\psi_{j_1}(\theta)$, but this time is reached for any choice of the functions $q_i(\theta)$ such that $q_{j_1}(\theta) + \cdots + q_{j_r}(\theta) = 1$ and $q_k(\theta) = 0$ for all $k \neq j_1, \ldots, j_r$. □

For convenience, we build the decision function of our mechanism introducing a *transformation function* $\psi : \Theta \to \mathbb{R}^n$ that returns a vector of n real values. The decision function is then obtained as $d = g(\theta) = g_\psi = \underset{i \in N}{\text{argmax}}\ (\psi_i(\theta))$. We say that ψ determines the "decision rule" or "decision function". Our main theorem give us insight into what can we expect about the set of functions ψ. Given our definition of $\psi_i(\theta)$ we can derive some intuition about the decision function. The theorem tells us that we can restrict our attention to deterministic solutions except when $\psi_i(\theta) = \psi_j(\theta), i, j \in N$. At these points, the decision is indifferent. The above theorem also gives us an optimality result.

Corollary 2. *Assume that all players are honest, mechanism M defined using the decision function $d = \underset{i \in N}{\operatorname{argmax}} \, (\psi_i(\theta))$ maximizes the utility of the system subject to fairness constraints.*

Finally, when fairness is symmetric in the sense that each player has the same fairness function, then each ψ_i depends only on the player's profile θ_i and therefore $\psi_i(\theta_i, \theta_{-i})$ could be reduced to $\psi_i(\theta_i)$. This last aspect allows us to state the following corollary.

Corollary 3. *When fairness is symmetric in the sense of $\eta_i(\theta) = \eta(\theta_i)$ $\forall i \in N$, and players have limited space strategy, then the probability q_i depends only on the player's value, that is $q_i(\theta) = q_i(\theta_i)$.*

Proof. The proof follows from the definition of $\psi_i(\theta)$ and therefore the decision function could be reduced to $d = \underset{i \in N}{\operatorname{argmax}} \, (\psi_i(\theta_i))$. As beliefs can not be changed by the strategy of others players, the probability $q_i(\theta)$ is only defined as a function of θ_i. □

Revisiting our particular cases of fairness defined as equal-number of resources (Eq. 3) and equal number of utility (Eq. 4) we can check that the solutions for ψ are in both cases straight lines. When fairness is defined as equal-number of resources (Eq. 3), $\psi_i(\theta)$ becomes

$$\psi_i(\theta) \equiv \theta_i + \lambda_i - \textstyle\sum_{k=1}^{n} \lambda_k \delta_k, \tag{11}$$

and therefore $\psi(\theta_i) = \theta_i + \lambda_i - \sum_{k=1}^{n} \lambda_k \delta_k$.

This solution has a very nice property that was already observed in our original work (QPQ with independent players). The mechanism designer could aggregate players when studying a single player. The mechanism designer can see the game as player i against the system formed by all other players ($j \in N, j \neq i$). In this case, player i has to compute just two values for λ, her own value λ_i and the aggregate value $\lambda_j = \sum_{k=1}^{n} \lambda_k \delta_k$. That is: $\psi(\theta_i) = \theta_i + \lambda_i - \lambda_j$, or even simpler: $\psi(\theta_i) = \theta_i + \lambda$. if we redefine λ as a new single real parameter that represents $\lambda_i - \sum_{k=1}^{n} \lambda_k \delta_k$.

This confirms that the decision function is a straight line where the parameter λ determines the point at which the line crosses the y-axis. And this is true for all players.

On the other hand, when fairness is defined as a function of utility (Eq. 4), our ψ function could be defined using

$$\psi_i(\theta) \equiv \theta_i \ (1 + \lambda_i - \textstyle\sum_{k=1}^{n} \lambda_k \delta_k), \tag{12}$$

and therefore $\psi(\theta_i) = \theta_i \ (1 + \lambda_i - \sum_{k=1}^{n} \lambda_k \delta_k)$.

Again, the decision function is a straight line where λ determines the slope. Aggregating players, the above solution could be reduced to $\psi(\theta_i) = \ (1 + \lambda_i - \frac{1}{n}\lambda_j)$, or $\psi(\theta_i) = \lambda \ \theta_i$.

Properties. The Fair QPQ Mechanism with Correlated players (M_{fair}) has the following properties:

1. M_{fair} is (ex-ante) **individual-rational**. This means that the expected utility of a player is at least its expected outside utility.
2. M_{fair} is not **allocative-efficient**, but assign tasks efficiently subject to equal number of tasks for each player. This property is a clear conclusion from Corollary 2.
3. There is no incentive for any of the players to lie about or hide their private information from the other players. Players will report truthfully in a Bayesian equilibrium. We said that M_{fair} is Bayesian incentive compatible. This property prevents selfish players from obtaining a benefit by misbehaving.

The two first properties are quite evident. The last property follows from Theorem 4.

Theorem 4. *When players have limited space strategy, and fairness is symmetric in the sense that $\eta_i(\theta) = \eta(\theta_i)$ $\forall i \in N$, then M_{fair} is Bayesian incentive compatible.*

Proof. For the sake of contradiction, let us suppose this proposition is false. Hence, there is some set of assignments for which, if i is not honest, she will obtain more utility in expectation.

From Corollary 3, this holds for any strategy of the aggregate player j, and in particular when all its players are honest. Hence, we can consider in the rest of the proof that the rest of $n-1$ players behave honestly.

Additionally, using the same corollary, we know that every player, $j \neq i \in N$, will obtain the same expected utility (independently whether i lies or not),

$$\int_{\Theta} u_j(d, \theta_j)\ q_j(\theta)\ d\pi(\theta) = \int_{\Theta} u_j(d, \theta_j)\ \hat{q}_j(\theta)\ d\hat{\pi}(\theta)$$

Now we can define a new mechanism M that assigns a task to player i (when i is honest and declares θ_i) with the same probability as the original QPQ assigns the task to the player i when she declares a false value $\hat{\theta}_i$. Then, $q_i(\theta_i) = \hat{q}_i(\hat{\theta})$. Note this new mechanism conserves the same fairness constraints as the original one. However, if the above were true, QPQ would not be optimal, since a mechanism that reproduces the same decisions under i lying (in presence of honest players) would different (lower) utility. Clearly, this is in contradiction of optimality of QPQ. Therefore, the best strategy for a player (the one optimizing her normalized utility) is to be honest. □

4 Practical QPQ Algorithm

After describing the different ingredients of our solution, we are able to propose an application of our mechanism. Due to space restrictions, we will only discuss an algorithm for a particular case. We propose an algorithm where the resource

Algorithm 1. QPQ Correlated mechanism (code for node i)

Estimate the preference θ_i
Publish the normalized value $\bar{\theta}_i = PIT(\theta_i)$
Wait to receive the normalized values $\bar{\theta}_j$ from the other players
for all $j \in N$ **do**
    if not $GoF\_Test(\bar{\theta}_j, Historic)$ **then**
        $\bar{\theta}_j \leftarrow Random(\bar{\theta}_{-j}, Historic)$
    end if
end for
$Historic \leftarrow Historic \cup \{\bar{\theta}\}$
Let $d = \underset{j \in N}{\mathrm{argmax}} \{\psi_j(\bar{\theta}_j)\}$
if $d = i$ **then**
    Resource is assigned to node i
end if
Update $\lambda_j, \forall j \in N$: $\lambda_{k+1,j} = \lambda_{k,j} + \epsilon_k(T_{k,j} - 1/n)$.

allocation achieves fairness in the number of resources allocated to each player. This algorithm could be extended to other fairness concepts. The details can be observed in Algorithm 1.

In the algorithm, $T_{k,j}$ denotes the percentage of decisions assigned to player j, computed at round k. As it can be observed, for each round, each player estimates her own value and publishes it. Publication means broadcasting a message with the value to all players (although any other means of distribution, like shared memory, can be used). By assumption, a player sends its value before it receives any of the others (concurrency, which implies that they do not depend from each other), and all of the values are correctly received at each player (reliability). Then, the algorithm assigns the resource to the player that publishes the highest value modified by a particular ψ_k.

Acceptance Test. We are assuming that players are reporting values using a uniform distribution. If their original distribution is not the uniform, we apply here the same normalization transformation proposed in [16] based on the *Probability Integral Transform* (PIT). Given the properties of the PIT, the idea is that any player applying correctly the PIT on her real type distribution, must generate a uniform distribution on the unit interval on her published normalized values. Hence, from the point of view of the mechanism designer, the problem consists on determining whether these published values follow or not that uniform distribution. There are a wide range of tests that allow checking that. These tests are called Goodness-of-Fit or GoF tests.

Continuing with this argument, we propose to implement the acceptance test of our algorithm by using some GoF test on the declared transformed sequence of values published by the player. Whenever a player is honest and she declares the values by applying the PIT transformation on her own distribution, these values will be uniformly distributed in the unit interval. In that case (with high probability) the GoF tests will accept the samples. More importantly, this process has an error which tends to zero when the number of samples (rounds) increases for any reasonable value of the threshold. For the study of our analytical results, we assume that GoF tests are perfect and this error is zero.

A tremendous amount of GoF tests have been proposed in the scientific literature. We propose to use the Kolmogorov-Smirnov (KS) test [17,18] test as the GoF test of QPQ. In contrast to our previous work with independent players [16], in this case it is necessary to add a second test. The goal for this new test is to check if a player is trying to modify the joint distribution. For our work, we have used the "Copula" R-Cran package. We note that no approach is always the best.

Punishment. In the case that a dishonest player tried to lie, one possible strategy is to generate increasing $\hat{\theta}$ values, so that the PIT transformed values are close to the unit. However, this type of behavior is quickly detected by the test. In that case, the question is how to establish a punishment. Inspired on previous works on linking mechanisms, the proposal is to reject the value declared by the player and generate a new random value according to the join prior distribution.

Practical computation of λ. The above solution reduces the problem to finding the value of λ that adjusts the tasks performed by players. In principle, we can ask the players to declare the joint distribution and calculate that parameter accordingly. But in general, we should not expect to find an analytical equation. That is, it is possible that π does not have an analytical expression, or even if it exists, players must estimate it empirically. There are multiple methods for π estimation, both parametric and nonparametric. The major difficulty with these systems is the convergence speed making it necessary a large number of samples. There is a relationship between the dimension of the feature and the number of samples needed. In our case, the dimension would be given by the number of players. Fortunately, each player can compute the QPQ mechanism using just only two dimensions (itself and the aggregate system).

However, players do not need to know the joint density function π, they only need to know the function $T(\cdot)$ that indicates the number of tasks performed given a parameter λ. We denote by $T(\lambda)$ the number of tasks that run the player when the decision value ψ is determined by the parameter λ. Again, we can not expect an analytic form for T. But under the right assumptions, we can approximate λ using stochastic approximation methods. Due to the characteristics of the transformation function and noting how it influences the number of tasks, we can expect that the function $T(\lambda)$ is continuous and decreasing (or increasing in the direction of λ). That is, there is always a value of λ for each percentage of desired tasks. Our proposal is to approximate λ by a sequence $\lambda_0, \lambda_1, \lambda_2, \lambda_3, \cdots \to \lambda$ constructed using a stochastic approximation method. The best known method is perhaps the Robbins-Monro method [19] although not the only one. Then, our algorithm must compute, for each iteration k,

$$\lambda_{k+1} = \lambda_k + \epsilon_k(T_k - 1/n). \tag{13}$$

Where T_k is an estimation of the average number of tasks performed by the player and where ϵ_k is a sequence of values that satisfies $\epsilon_k > 0$, $\epsilon_k \rightarrow 0, \sum_k \epsilon_k = \infty$. Note that, in order to estimate T_k we don't need to store previous samples and memory consumption is low.

Simulations. By performing simulations, we have checked various aspects of our proposal. Mainly, we wondered how Robbins-Monro algorithm performs in time. We have simulated several alternatives for the generation of the sequence of values ϵ_k. In our simulation we have used two methods: $\epsilon_k = 1/k$ and $\epsilon_k = \frac{1}{log(k)+k}$. The first one is the original proposal of Robbins-Monro's work. With this sequence, our experiments produce some oscillations in the λ estimation and the speed of convergence was far from ideal. We found better results with the second approach. Figure 1 presents an experiment with three players, the first two are correlated and the third one is independent. Without our algorithm, the independent player will obtain less utility than the two other players. On the other hand, with our proposal, fairness is achieved and every player will have a proportional number of assignments.

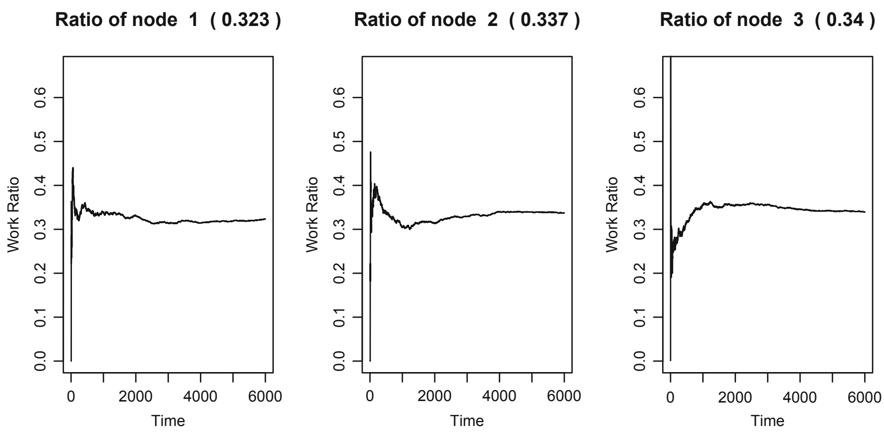

Fig. 1. Evolution of work ratio (number of tasks) with QPQ.

5 Conclusions and Future Work

In this paper we have created a novel scheme capable of providing efficient resource allocation in distributed systems even in the presence of selfish correlated players. We have shown that, for a general notion of fairness, the mechanism can be proved to perform efficiently and to maintain the incentive of players to participate. In addition, we have proposed a specific realization of the mechanism as an algorithm implementable in real distributed environments with affordable computational and communication costs. This algorithm is susceptible of being used in repeated task allocations given that our simulations demonstrate its rapid convergence, which open new horizons for systems based on open systems for distributed collaborative tasks execution.

Despite this, the authors consider necessary to extend the current research in several directions. First, the model requires knowledge on the number of players that participate. We may find scenarios where this is not reasonable, e.g., scenarios in which several players "hide" and play the game with a single identity, which may resulting on the mechanism not achieving fairness. Second, it would be important to analyze the problem when more flexible space strategies are possible. One of our main assumptions has been to consider that correlations are fixed and that players are not able to alter them through their strategies. This assumption is reasonable when information is private and the mechanism is designed in such a way that players cannot make their declared (true or false) values on an iteration dependent on the values of others at the same iteration. However, there are many real-live scenarios where players may be able to share their values making more complex interdependent strategies possible. This would break the properties of our proposed algorithm.

References

1. Jun, S., Ahamad, M.: Incentives in BitTorrent induce free riding. In: Proceedings of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems. P2PECON '05, pp. 116–121. ACM, New York (2005)
2. Ahmad, I., Ranka, S., Khan, S.: Using game theory for scheduling tasks on multi-core processors for simultaneous optimization of performance and energy. In: IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2008, pp. 1–6. IEEE (2008)
3. Bell, M.G.: A game theory approach to measuring the performance reliability of transport networks. Transp. Res. Part B: Methodol. **34**, 533–545 (2000)
4. Srivastava, V., Neel, J., Mackenzie, A.B., Menon, R., Dasilva, L.A., Hicks, J.E., Reed, J.H., Gilles, R.P.: Using game theory to analyze wireless ad hoc networks. IEEE Commun. Surv. Tutor. **7**, 46–56 (2005)
5. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: Proceedings of the 10th ACM Conference on Electronic Commerce, pp. 177–186. ACM (2009)
6. Jackson, M.O., Sonnenschein, H.F.: The linking of collective decisions and efficiency. Microeconomics 0303007, EconWPA (2003)
7. Jackson, M.O., Sonnenschein, H.F.: Overcoming incentive constraints by linking decisions. Econometrica **75**, 241–257 (2007)
8. Papadimitriou, C.H.: Games, algorithms, and the internet. In: Srinivasan, S., Ramamritham, K., Kumar, A., Ravindra, M.P., Bertino, E., Kumar, R., (eds.) WWW, pp. 5–6. ACM (2011)
9. Jackson, M.: Mechanism theory (2003). http://web.stanford.edu/~jacksonm/mechtheo.pdf
10. Jackson, M.O.: A crash course in implementation theory. Working papers 1076, Division of the Humanities and Social Sciences, California Institute of Technology (1999)

11. Kandori, M.: Social norms and community enforcement. Rev. Econ. Stud. **59**, 63–80 (1992)
12. Ellison, G.: Cooperation in the prisoner's dilemma with anonymous random matching. Rev. Econ. Stud. **61**, 567–588 (1994)
13. Kamvar, S.D., Schlosser, M.T., Garcia-Molina, H.: The Eigentrust algorithm for reputation management in P2P networks. In: Proceedings of the 12th International Conference on World Wide Web, pp. 640–651. ACM (2003)
14. Friedman, E.J., Halpern, J.Y., Kash, I.: Efficiency and nash equilibria in a scrip system for P2P networks. In: Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 140–149. ACM (2006)
15. Engelmann, D., Grimm, V.: Mechanisms for efficient voting with private information about preferences. Econ. J. **122**, 1010–1041 (2012)
16. Santos Méndez, A., Fernández Anta, A., López Fernández, L.: Quid pro quo: a mechanism for fair collaboration in networked systems. PLOS ONE (2013), See CoRR abs/1207.6045 (2012, in press)
17. Kolmogorov, A.N.: Sulla determinazione empirica di una legge di distribuzione. Giornale dellIstituto Italiano degli Attuari **4**, 83–91 (1933)
18. Smirnov, N.V.: On the estimation of the discrepancy between empirical curves of distribution for two independent samples. Bull. Moscow Univ. **2**, 3–16 (1939)
19. Robbins, H., Monro, S.: A stochastic approximation method. Ann. Math. Stat. **22**, 400–407 (1951)