

Convergence in Player-Specific Graphical Resource Allocation Games

Valentino Pacifici and György Dán

Abstract—As a model of distributed resource allocation in networked systems, we consider resource allocation games played over a influence graph. The influence graph models limited interaction between the players due to, e.g., the network topology: the payoff that an allocated resource yields to a player depends only on the resources allocated by her neighbors on the graph. We prove that pure strategy Nash equilibria (NE) always exist in graphical resource allocation games and we provide a linear time algorithm to compute equilibria. We show that these games do not admit a potential function: if there are closed paths in the influence graph then there can be best reply cycles. Nevertheless, we show that from any initial allocation of a resource allocation game it is possible to reach a NE by playing best replies and we provide a bound on the maximal number of update steps required. Furthermore we give sufficient conditions in terms of the influence graph topology and the utility structure under which best reply cycles do not exist. Finally we propose an efficient distributed algorithm to reach an equilibrium over an arbitrary graph and we illustrate its performance on different random graph topologies.

Index Terms—resource allocation, graphical games, player-specific congestion games, best reply dynamics.

I. INTRODUCTION

Resource allocation is a fundamental problem in communication systems. In general, resource allocation problems are characterized by a set of nodes competing for the same set of resources, and arise in a wide variety of contexts, from congestion control, through content management [1], [2] to dynamic channel allocation for opportunistic spectrum access [3], [4]. The nodes competing for the resources are often autonomous entities, such as end hosts or mobile nodes, and therefore the resource allocation problem has to be solved in a decentralized manner. Furthermore, the solution should be compatible with the nodes' own interests.

Congestion games are widely used as a game theoretic model of distributed resource allocation [5]. In congestion games a set of players allocate resources in order to maximize their own utility, but the utility provided by an individual resource to a player is a function of how many other players have allocated the resource. In the case of networked systems, the interactions between the

players are often limited by the network topology, which can be captured by introducing an influence graph in the model of congestion games [6], [7], [8]. Under certain conditions of symmetry, e.g., homogeneous or linear utility functions [6], [7], [8], [9], [10], [11], [12], congestion games allow a potential function [13], and thus the two most fundamental questions, the existence of pure Nash equilibria and the convergence of myopic learning rules to the equilibria are answered. For congestion games that do not admit a potential function, the answer to these two fundamental questions is, in general, not known [14], [15], [16].

In this work we consider a class of resource allocation games that gives rise to a graphical player-specific congestion game that does not admit a potential function. In the model we consider, the player-specific utility of a resource to a player is amortized if any of its neighbors allocates the same resource. This model captures problems arising in a number of fields: object placement in the context of CPU caching in computer architectures [17], the placement of contents in clean-slate information centric network architectures [1], the problem of cooperative caching between Internet service providers [2], the distributed allocation of radio spectrum to radio transmitters [3], [4], and the distributed scheduling of jobs [18].

Our contributions are the following. We show that Nash equilibria exist in graphical resource allocation games for arbitrary influence graphs and payoff structures, and give a bound on the complexity of finding equilibria. We show that nodes might cycle arbitrarily long before reaching equilibrium if the influence graph is non-complete, thus the game does not admit a potential function. We then provide sufficient conditions in terms of the graph topology and the payoff structure that guarantee that cycles do not exist if players play best replies. Furthermore, we give a sufficient condition under which a simple and efficient distributed algorithm can be used to reach an equilibrium, and illustrate the efficiency of the algorithm with numerical results.

The rest of the paper is structured as follows. We define graphical resource allocation games in Section II, and provide results for arbitrary graph topologies in Section III. We then analyze the convergence to equilibria for complete graphs in Section IV and for specific payoff structures in Section V. Section VI introduces the distributed algorithm for reaching equilibrium. In Section VII we discuss related work, and Section VIII concludes the paper.

The authors are with the Laboratory for Communication Networks, School of Electrical Engineering, and the ACCESS Linnaeus Centre, KTH Royal Institute of Technology. Osquidas väg 10, 10044 Stockholm. Tel.: +46 87904251. E-mail: {pacifici,gyuri}@kth.se. Manuscript received December 15, 2011; revised May 25, 2012; accepted July 15, 2012.

II. THE RESOURCE ALLOCATION GAME

In the following we formulate the problem of graphical resource allocation as a non-cooperative game played on a graph. We consider a set of nodes N and a set of resources \mathcal{R} . Every node is located at a vertex of an undirected graph $\mathcal{G} = (N, E)$, called the influence graph. We denote the set of neighbors of node i by $\mathcal{N}(i)$, i.e., $\mathcal{N}(i) = \{j | (i, j) \in E\}$. Each node i allocates $K_i \in \mathbb{N}_+$ different resources: we describe the set of resources allocated by node i with the $|\mathcal{R}|$ dimensional vector $a_i = (a_i^1, \dots, a_i^{|\mathcal{R}|})$, whose component $a_i^r \in \{0, 1\}$ is 1 if resource r is allocated by node i . Thus the set of feasible resource allocation vectors for node i is $\mathcal{A}_i = \{a_i | \sum_r a_i^r \leq K_i\} \subseteq \{0, 1\}^{|\mathcal{R}|}$. To ease notation we say that a resource $r \in \mathcal{R}$ is *i-busy* if it is allocated by at least one of node i 's neighbors and we define $\pi_i^r(a_{-i}^r) \triangleq \prod_{j \in \mathcal{N}(i)} (1 - a_j^r) = 0$, otherwise we say that resource $r \in \mathcal{R}$ is *i-free*.

We call c_{ir} the value of resource r for node i . The payoff $U_i^r(1, a_{-i}^r)$ that a node i gets from allocating a resource r is influenced by the resource allocation of its neighboring nodes $\mathcal{N}(i)$. A node i allocating resource r gets as payoff the value c_{ir} if resource r is *i-free*. If resource r is *i-busy*, node i gets payoff $\delta_i c_{ir}$, where δ_i is the cost of sharing for node i , $0 \leq \delta_i < 1$,

$$U_i^r(1, a_{-i}^r) = \begin{cases} c_{ir} & \text{if } \pi_i^r(a_{-i}^r) = 1 \\ \delta_i c_{ir} & \text{if } \pi_i^r(a_{-i}^r) = 0. \end{cases} \quad (1)$$

A node does not get any payoff from the resources it does not allocate, i.e., $U_i^r(0, a_{-i}^r) = 0$.

We model the resource allocation problem as a strategic game $\Gamma = \langle N, (\mathcal{A}_i)_{i \in N}, (U_i)_{i \in N} \rangle$, where the utility function of player i is the sum of the payoffs $U_i(a_i, a_{-i}) = \sum_r U_i^r(a_i^r, a_{-i}^r)$. Note that the influence graph influences the payoffs that player i gets from the allocated resources via the neighbor set $\mathcal{N}(i)$, i.e., the utility of player i is entirely specified by the actions of players $j \in \mathcal{N}(i)$. Under the assumption that every player i allocates always K_i resources, the graphical resource allocation game we consider is a graphical player-specific matroid congestion game.

In the following we outline two applications of graphical resource allocation games.

A. Selfish Object Replication on Graphs

Consider a set N of nodes and a set \mathcal{R} of objects of unit size. The demand for object $r \in \mathcal{R}$ at node $i \in N$ is given by the rate $w_i^r \in \mathbb{R}_+$. Node i has integer storage capacity K_i which it uses to replicate objects locally. The marginal cost of serving requests for object r in node i is α_i if the object is replicated in node i , it is β_i if the object is replicated in a node $j \in \mathcal{N}(i)$ neighboring i , and it is γ_i otherwise. It is reasonable to consider that it is not more costly to access a resource replicated locally than one replicated at a neighbor, and it is less costly to access a resource replicated at a neighbor than retrieving it directly from the common set of resources. Formally $\alpha_i \leq \beta_i < \gamma_i$. The cost of node i due to object r is proportional to the

demand w_i^r , and is a function of a_i and the replication states a_{-i} of the neighboring nodes, and its total cost is

$$C_i(a_i, a_{-i}) = \sum_{r \in \mathcal{R}} C_i^r(a_i^r, a_{-i}^r) = \sum_{r \in \mathcal{R}} w_i^r \left[\alpha_i a_i^r + (1 - a_i^r) [\gamma_i \pi_i^r(a_{-i}^r) + \beta_i (1 - \pi_i^r(a_{-i}^r))] \right] \quad (2)$$

The goal of node i is to choose a replication strategy a_i^* that minimizes its total cost given the strategy profile a_{-i} of the other nodes. Observe that the cost of object r for node i can be expressed as

$$\begin{aligned} C_i^r(a_i^r, a_{-i}^r) &= C_i^r(0, a_{-i}^r) - (C_i^r(0, a_{-i}^r) - C_i^r(a_i^r, a_{-i}^r)) \\ &= C_i^r(0, a_{-i}^r) - CS_i^r(a_i^r, a_{-i}^r), \end{aligned}$$

where $CS_i^r(a_i^r, a_{-i}^r)$ is the cost saving that node i achieves through object r given the other nodes' replication strategies. Since the cost $C_i^r(0, a_{-i}^r)$ is independent of the action a_i^r of node i , finding the minimum cost is equivalent to maximizing the aggregated cost saving

$$\begin{aligned} \arg \min_{a_i} C_i(a_i, a_{-i}) &= \arg \min_{a_i} \sum_r C_i^r(a_i^r, a_{-i}^r) \\ &= \arg \min_{a_i} \left(\sum_r C_i^r(0, a_{-i}^r) - \sum_r CS_i^r(a_i^r, a_{-i}^r) \right) \\ &= \arg \max_{a_i} \sum_r CS_i^r(a_i^r, a_{-i}^r). \end{aligned}$$

We can express the cost saving $CS_i^r(a_i^r, a_{-i}^r)$ of node i by substituting (2)

$$\begin{aligned} CS_i^r(a_i^r, a_{-i}^r) &= w_i^r [\beta_i (1 - \pi_i^r(a_{-i}^r)) + \gamma_i \pi_i^r(a_{-i}^r)] \\ &\quad - w_i^r [\alpha_i a_i^r + (1 - a_i^r) [\beta_i (1 - \pi_i^r(a_{-i}^r)) + \gamma_i \pi_i^r(a_{-i}^r)]] \\ &= a_i^r w_i^r [\beta_i (1 - \pi_i^r(a_{-i}^r)) + \gamma_i \pi_i^r(a_{-i}^r) - \alpha_i]. \end{aligned}$$

Thus $CS_i^r(0, a_{-i}^r) = 0$. Then, by defining $\delta_i \triangleq \frac{\beta_i - \alpha_i}{\gamma_i - \alpha_i}$ and $w_i^r [\gamma_i - \alpha_i] = c_{ir}$, we obtain that $CS_i^r(1, a_{-i}^r) = U_i^r(1, a_{-i}^r)$ defined in (1). This model of object replication was used, for example, in [17] to model distributed cache allocation in a computer architecture. It was used to model cooperative caching of contents among Internet Service Providers (ISPs) in [2], where the ISPs are neighbors if they have a peering agreement, and the costs γ_i , β_i , and α_i correspond to the cost of downloading contents over transit, peering and local links, respectively. The model also captures the cost structure of cooperative caching for multiview video streaming systems considered in [19], in which case the costs correspond to the access of image partitions from a remote repository, remote servers and local servers, respectively.

B. Graph Multi-coloring, Distributed Radio Spectrum Allocation and Medium Access Control

Proper graph multi-coloring is a generalization of the proper graph coloring problem. Given a graph $\mathcal{G} = (N, E)$ and a set \mathcal{R} of colors the task is to assign K_i distinct colors to vertex $i \in N$ such that no adjacent vertices have the same color. In general, the goal is to use as few as possible

colors. In the case of an improper coloring the same color can be assigned to adjacent vertices, but this involves a penalty δ_i for vertex i .

Graph multi-coloring has a number of applications. In the case of scheduling the nodes are jobs, the colors are time units, and K_i is the time needed to finish job i . The minimum number of colors is then the makespan of all jobs [18]. In the case of medium access control there is a set N of nodes that contend for some radio channels or for time slots [3], [4]. The influence graph \mathcal{G} models the potential conflicts due to co-channel interference that can occur between nodes. The model of an undirected influence graph is appropriate for a system in which the pairs of nodes that communicate to each other are relatively close to each other, and fast fading is not a dominant fading component. In the case of channel assignment, each node $i \in N$ has K_i radio interfaces that it can use to transmit data. Alternatively, in the case of time slot allocation, each node $i \in N$ can use K_i time slots. The throughput that node i can achieve transmitting on resource $r \in \mathcal{R}$ in the absence of interference is c_{ir} . In the case of interference the expected throughput drops to $\delta_i c_{ir}$, where δ_i is the resource deterioration expected by node i under interference. In some applications a node can represent an entire network. In the case of coexisting Zigbee PRO networks, for example, the coordinator of each network $i \in N$ is in charge of selecting the channel for the entire network, so as to minimize the expected interference with other networks.

III. EXISTENCE OF EQUILIBRIA AND CONVERGENCE

We start with addressing two fundamental questions. First, we address the question whether in the graphical resource allocation problem there exist equilibrium resource allocations, from which no player would like to deviate unilaterally. Second, we address the question whether the players could reach an equilibrium state in a distributed manner.

A. Existence of equilibria

In order to formalize equilibrium allocations in the resource allocation game and to formulate our results, we start with the definition of three key terms used throughout the section.

Definition 1. A *best reply* of player i is a best strategy a_i^* of player i given the other players' strategies, that is, $U_i(a_i^*, a_{-i}) \geq U_i(a_i, a_{-i}), \forall a_i \in \mathcal{A}_i$.

A *best reply improvement path* is a sequence of strategy profiles, such that in every step t there is one player that *strictly increases* its utility by updating her strategy through performing a *best reply*. Note that a best reply improvement path implies that only one player at a time updates her strategy; this property is known in the literature as asynchronous updates. An improvement path terminates when no player can perform an improvement step. The resulting strategy profile is a pure strategy NE,

defined as a strategy profile a^* in which every player's strategy is a best reply to the other players' strategies

$$U_i(a_i^*, a_{-i}^*) \geq U_i(a_i, a_{-i}^*) \quad \forall a_i \in \mathcal{A}_i, \forall i \in N. \quad (3)$$

Example 1. Consider $|N| = 2$ players, $|\mathcal{R}| = 3$ resources and $K_i = 1$. Let $c_{11} > c_{12} > c_{13} > c_{11}\delta_1 > \dots$ and $c_{22} > c_{22}\delta_2 > c_{21} > c_{23} > \dots$. If the initial strategy profile is $(3, 1)$ then the following is a best reply improvement path where the unique deviator is marked at every step: $(3, 1) \xrightarrow{1} (2, 1) \xrightarrow{2} (2, 2) \xrightarrow{1} (1, 2)$. Observe that $(1, 2)$ is a NE.

Let us now turn to the question whether all graphical resource allocation games have at least one pure strategy Nash equilibrium. Consider the strategy profile $a(0)$ that consists of the best replies that the players would play on an edgeless social graph. Let us consider now a best reply path starting from the strategy profile $a(0)$. For $t \leq n$ each player $i \in N$ has a chance to play her first best reply at $t = i$. For $t > n$ they play in an arbitrary order. We can make two important observations about the players' best replies.

Lemma 1. Player $i \in N$ allocates only i -free resources when she first updates her strategy at $t = i$.

Proof: Player i first updates her strategy at $t = i$. Define the evicted set as $E_i(t) = \{r | a_i^r(0) = 1 \wedge a_i^r(t) = 0\}$ and the inserted set as $I_i(t) = \{r | a_i^r(0) = 0 \wedge a_i^r(t) = 1\}$. Consider now two resources $r \in E_i(t = i)$ and $r' \in I_i(t = i)$. By definition $a_i^r(0) = 1$ and $a_i^{r'}(0) = 0$, thus by the definition of best reply and because of the edgeless social graph at $t = 0$

$$c_{ir} \geq c_{ir'} \quad (4)$$

Since r' is allocated and r is evicted, at the improvement step $t = i$ it must hold that

$$U_i^{r'}(1, a_{-i}^{r'}(i)) > U_i^r(1, a_{-i}^r(i)). \quad (5)$$

Consider now the definition of the payoff (1) and (4). If r is i -free then $U_i^r(1, a_{-i}^r(i)) = c_{ir}$ and consequently (5) cannot hold. Similarly, if r' is i -busy then $U_i^{r'}(1, a_{-i}^{r'}(i)) = \delta_i c_{ir'}$ and consequently (5) cannot hold. Thus the only possibility to satisfy (5) is that r' is i -free and r is i -busy. ■

Lemma 2. Consider a sequence of best reply steps and the best reply $a_i(t)$ played by player i at step $t > 0$. A necessary condition for $a_i(t)$ not being a best reply for i at step $t' > t$ is that at least one of the following holds

- (i) An i -free resource r allocated by i ($a_i^r(t) = 1, \pi_i^r(t) = 1$) becomes i -busy by step t' ,
- (ii) An i -busy resource r not allocated by i ($a_i^r(t) = 0, \pi_i^r(t) = 0$) becomes i -free by step t' .

The proof of Lemma 2 is in the Appendix. Lemma 1 implies that a resource allocated by player i cannot change from i -free to i -busy during the first round of best reply steps ($1 \leq t \leq n$). Consequently, condition (i) in Lemma 2 cannot hold, and the only reason why player i would update her strategy a second time (at some $t > n$) is that

condition (ii) holds, and she would start allocating an *i-free* resource. Thus, by induction, condition (i) will never hold, and in every step t player $i \in N$ starts allocating only *i-free* resources. Since no player ever starts allocating an *i-busy* resource, according to the expression of the payoff in (1) the utilities of the players cannot decrease for $t > 0$. Nevertheless, every time a player updates her strategy her utility must strictly increase. Since the players' utilities cannot increase indefinitely, the best reply path must end in a NE. Hence we can state the following.

Theorem 1. *Every graphical resource allocation game possesses a pure strategy Nash equilibrium.*

For a complete influence graph every player makes at most one best-reply improvement step [20], but this does not hold even for a simple non-complete influence graph: on a ring of 4 players with $c_{ir} = c_{jr} (\forall r \in \mathcal{R})$ at least one player updates her strategy twice. We can however bound the number of required improvement steps. We know that from $a(0)$ every player can only start allocating *i-free* resources. In the worst case each player i inserts exactly one *i-free* resource r at every best reply step. According to condition (ii) in Lemma 2, r becomes *i-free* as an effect of a best reply of a player $j \in \mathcal{N}(i)$. The number of resources that can change from *i-busy* to *i-free* for an arbitrary player i are at most $\sum_{j \in \mathcal{N}(i)} K_j$, thus we obtain the following result

Theorem 2. *It is possible to compute a Nash equilibrium of a graphical resource allocation game in at most $\sum_{i \in N} \sum_{j \in \mathcal{N}(i)} K_j$ steps.*

Due to space limitations we do not address issues such as the price of anarchy and price of stability, and the occurrence of phenomena such as Bélády's anomaly [21]. Instead, in the rest of the paper our focus is on how to reach Nash equilibria in an efficient way, with the aim of providing guidelines for designing distributed algorithms that would be able to reach an equilibrium allocation with little overhead.

B. Convergence to Nash Equilibria

The existence of equilibrium states is important, but in a distributed system it is equally important to have efficient distributed algorithms that the nodes can use to reach an equilibrium state. The algorithm in Theorem 1 can be used to reach an equilibrium state if the values c_{ir} of the resources in the nodes never change, so once a NE is reached, the nodes will not deviate from it. Nevertheless, the algorithm would be inefficient if the values c_{ir} or the influence graph can change over time, as the equilibrium states for different values and different influence graphs are, in general, different. Hence, an important question is whether the players will reach a NE given an arbitrary initial strategy profile, e.g., a NE for past c_{ir} or a NE for a past influence graph, and given the myopic decisions the players make to update their strategies.

A straightforward distributed algorithm would be to let every player play her best reply at the same time (syn-

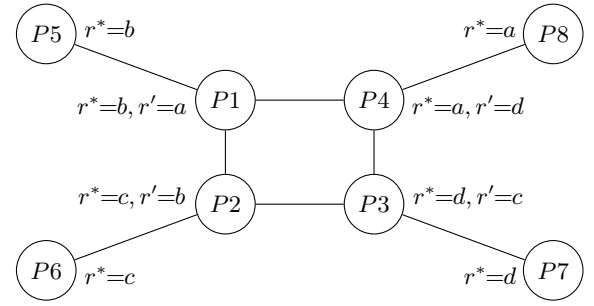


Fig. 1: Influence graph and players' preferences that allow a cycle in best replies.

chronously) until an equilibrium is reached. The algorithm only requires synchronization between neighboring players, thus it is relatively easy to implement. Unfortunately, it is easy to find resource allocation games where players cannot find an equilibrium this way.

Example 2. *Consider two players and $|\mathcal{R}| \geq 2$. Let $c_{i1} > c_{i2} > c_{i1}\delta_i > c_{i2}\delta_i$ and $K_i = 1$. If the initial allocation strategies are $a_i(0) = (1, 0)$ then we have $(1, 0) \xrightarrow[1,2]{} (0, 1) \xrightarrow[1,2]{} (1, 0)$, etc.*

As an alternative, consider an algorithm that only allows one player to play a best reply at a time. This is the case of the asynchronous updates used in Section III-A, which generate best reply improvement paths. The algorithm requires global synchronization to ensure that no players perform an update simultaneously. In what follows we show that even best reply improvement paths can be arbitrarily long.

Consider a resource allocation game played over the influence graph shown in Figure 1, where $\mathcal{R} = \{a, b, c, d\}$, and $K_i = 1$. Each player $1 \leq i \leq 4$ has a resource $r^* \in \mathcal{R}$ such that $c_{ir^*}\delta_i > c_{ir}\delta_i \forall r \neq r^*$, and at least one resource $r' \in \mathcal{R}$ such that $c_{ir'} > c_{ir^*}\delta_i$. For players $5 \leq i \leq 8$ there is a resource $r^* \in \mathcal{R}$ such that $c_{ir^*}\delta_i > c_{ir} \forall r \neq r^*$. In the following we show an asynchronous best reply dynamic that cycles, we omit the strategies of players $5 \leq i \leq 8$ since they always allocate the resource that has the highest value at their respective neighboring player $i - 4$. The *i-busy* resources are in bold.

$$\begin{aligned} (\mathbf{a}, b, \mathbf{d}, \mathbf{a}) &\xrightarrow[3]{} (\mathbf{a}, b, c, \mathbf{a}) \xrightarrow[1]{} (\mathbf{b}, \mathbf{b}, c, \mathbf{a}) \xrightarrow[4]{} (\mathbf{b}, \mathbf{b}, c, \mathbf{d}) \\ &\xrightarrow[2]{} (\mathbf{b}, \mathbf{c}, \mathbf{c}, \mathbf{d}) \xrightarrow[1]{} (\mathbf{a}, \mathbf{c}, \mathbf{c}, \mathbf{d}) \xrightarrow[3]{} (\mathbf{a}, \mathbf{c}, \mathbf{d}, \mathbf{d}) \\ &\xrightarrow[2]{} (\mathbf{a}, b, \mathbf{d}, \mathbf{d}) \xrightarrow[4]{} (\mathbf{a}, b, \mathbf{d}, \mathbf{a}) \end{aligned}$$

The existence of a cycle in the best reply paths implies that resource allocation games played over an arbitrary graph do not admit a potential function [13]. It also raises the question whether there could be an initial strategy from which *every* best reply improvement path is infinite. Starting a best reply improvement path from such an initial strategy profile, the players would never reach a NE using the asynchronous algorithm.

We show that fortunately this is not the case; from every

strategy profile there exists at least one finite best reply path that leads to a NE. This property is called weak acyclicity [14], [22].

Definition 2. A game is *weakly acyclic under best replies* if from every strategy profile a , there is a best reply improvement path starting from a and ending in a pure Nash equilibrium.

To show that resource allocation games are weakly acyclic, let us consider a best reply $a_i(t)$ of a player i at step t . We can define four not mutually exclusive properties of $a_i(t)$, depending on whether the involved resources are *i-busy* as shown in Table I.

1	$\exists r \in E_i(t), r \text{ i-busy}$	$\exists r' \in I_i(t), r' \text{ i-busy}$
2	$\exists r \in E_i(t), r \text{ i-busy}$	$\exists r' \in I_i(t), r' \text{ i-free}$
3	$\exists r \in E_i(t), r \text{ i-free}$	$\exists r' \in I_i(t), r' \text{ i-busy}$
4	$\exists r \in E_i(t), r \text{ i-free}$	$\exists r' \in I_i(t), r' \text{ i-free}$

TABLE I: Four properties of a best reply move

Note that a best reply necessarily has at least one of the listed properties. The following two lemmas state that the best replies in a cycle cannot be arbitrary.

Lemma 3. *In every best reply cycle there exists at least one strategy profile $a(t)$ in which at least one player $i \in N$ performs a best reply that has property 1). Furthermore, in a best reply cycle it is not possible to perform a best reply that has property 3).*

Lemma 4. *Assume that a player i performs a best reply with property 1) at step t in a best reply cycle such that $r' \in I_i(t)$ and r' is *i-busy*. Then every resource $r'' \in \mathcal{R}$ for which $c_{ir''} > c_{ir'}$ is allocated by player i at step t ($a_i^{r''}(t)=1$).*

The proofs of the lemmas are given in the Appendix. Let us consider a strategy profile $a(t)$ in a best reply cycle in which at least one player can perform a best reply that has property 1). Such a strategy profile exists according to Lemma 3. Starting from $a(t)$, let us perform a sequence of best replies with property 1). From Lemma 4, it follows that player i cannot evict the resources that she inserted as *i-busy* through these best replies, thus every player i can perform at most K_i best replies with property 1). So after at most $\sum_{i \in N} K_i$ best replies with property 1) we reach a strategy profile $a(t')$ in which there is no player that can perform a best reply that has property 1). If in $a(t')$ no player can make a best reply then $a(t')$ is a NE. Otherwise, if a player i can perform a best reply in $a(t')$ then it will have neither property 1) nor 3). As a consequence of a best reply that has property 2) or 4) only, condition (i) of Lemma 2 cannot hold for any player, so the only new reason why a player j would perform a best reply is that condition (ii) in Lemma 2 is satisfied, and she would start allocating *j-free* resources. Thus, by induction, condition (i) of Lemma 2 will never hold, and in every step $t'' > t'$ players only perform best replies that have property 2) or 4) but not 1) or 3). A player's

utility strictly increases when it performs a best reply, and a best reply with property 2) or 4) does not decrease any other player's utility. Since the players' utilities cannot increase indefinitely, this path must end in a NE after a finite number of steps. Hence we can state the following.

Theorem 3. *Every graphical resource allocation game is weakly acyclic under best replies.*

Weak acyclicity has two important consequences for system design. First, if the nodes that compete for the resources update their allocations one at a time in a myopic way and the order of updates is fixed then the nodes might cycle forever without reaching an equilibrium. Therefore, it is advisable that the next node to perform the update step is chosen at random by the system. By doing so the system will reach an equilibrium after a finite number of update steps *on average*, even though the number of update steps can be arbitrarily high. Second, weak acyclicity in best replies implies that various complex learning rules can be used to reach an equilibrium with high probability. One example is adaptive play, described in [22], when nodes select the next allocation simultaneously based on a random sample of a finite history of the allocations. Such a learning rule resembles a system in which the nodes update their allocations based on a sliding window estimator of the other players' allocations. Another example is regret based learning [23], in which case every node updates its allocation simultaneously so as to minimize its loss of utility in retrospect given the history of all allocations.

In the discussion above we showed that every best reply path that starts from a strategy profile $a(t')$ in which no player can perform a best reply that has property 1) ends in a NE. It follows that in every strategy profile of a best reply cycle there is at least one player i that can perform a best reply with property 1). Let us consider now the number of steps needed to reach a NE starting from an arbitrary strategy profile $a(t)$. We have already shown that after performing at most $\sum_{i \in N} K_i$ best replies with property 1) we reach a strategy profile $a(t')$ in which there is no player that can perform a best reply that has property 1). Given that there is no player that can perform a best reply with property 1), we can use the same reasoning that we formulated to show Theorem 2 to bound the maximum number of best reply steps from $a(t')$ to a NE to $\sum_{i \in N} \sum_{j \in \mathcal{N}(i)} K_j$. Summing it to the maximum number of steps required from $a(t)$ to $a(t')$ we obtain the following upper bound.

Corollary 1. *From an arbitrary strategy profile there exists a best reply path that reaches a NE in at most $\sum_{i \in N} K_i + \sum_{i \in N} \sum_{j \in \mathcal{N}(i)} K_j$ steps.*

As an example, consider a graphical resource allocation game with $K_i = K \forall i \in N$ and a regular influence graph of degree d . In this case the maximum length of a best reply path would be $|N| \cdot K + |N| \cdot d \cdot K = |N|K(1 + d)$.

IV. THE CASE OF COMPLETE INFLUENCE GRAPH

In this section we consider the case that the influence graph is *complete* and use it to illustrate the influence of the graph topology on the convergence to equilibria. We make use of the notion of the finite best reply property [14].

Definition 3. A game possesses the *finite best reply property* (FBRP) if every best reply improvement path is finite.

For a complete influence graph the resource allocation game is a special case of the player-specific matroid congestion games investigated in [16], and it is known that if $|\mathcal{A}_i| = 2 \forall i \in N$ or if $|N| = 2$, then the game has the FBRP. It is not known whether the game possesses the FBRP in general.

Let us denote by $F_i(t)$ the set of *i-free* resources and by $B_i(t)$ the set of *i-busy* resources allocated by i at step t . Note that $|F_i(t)| + |B_i(t)| = K_i$. We can prove the following lemmas.

Lemma 5. *On a complete influence graph, a best reply step $a_i(t)$ performed by player i in a best reply cycle cannot affect the number of *i-busy* resources allocated by i , that is, $|B_i(t)| = |B_i(t-1)|$.*

The proof of Lemma 5 is given in the Appendix.

Lemma 6. *On a complete influence graph, for a best reply step $a_i(t)$ performed by player i in a best reply cycle it holds that $c_{ir'} < c_{ir}$, $\forall r' \in E_i(t), r \in I_i(t)$.*

Proof: Player i solves a knapsack problem to construct her best reply. Recall that according to Lemma 5 we have $|B_i(t-1)| = |B_i(t)|$ and consequently $|F_i(t-1)| = |F_i(t)|$. Hence we can construct the best reply of player i by dividing the knapsack problem into two similar subproblems: we can solve the knapsack problem for all the *i-busy* resources and populate the set $B_i(t)$, and do the same with the set $F_i(t)$ using the *i-free* ones. Suppose that k resources are evicted from one set, consequently k are inserted, and in order for the result to be the solution of the knapsack problem, the cost saving yielded by each inserted resource must be higher than the cost saving yielded by each evicted resource. That is, for every $r, r' \in \mathcal{R}$ such that r was inserted and r' was evicted from $B_i(t)$, we have $c_{ir}\delta_i > c_{ir'}\delta_i \Rightarrow c_{ir} > c_{ir'}$. Similarly, for every $r, r' \in \mathcal{R}$ such that r was inserted and r' was evicted from $F_i(t)$, we have $c_{ir} > c_{ir'}$. This proves the lemma. ■

Assume now that there is a best reply cycle. According to Lemma 6 each best reply step can only move towards resources with higher value. The number of resources $|\mathcal{R}|$ is finite, hence every player can only perform a finite number of best replies and the best reply path terminates after a finite number of steps. Thus the following result holds.

Theorem 4. *Every resource allocation game played over a complete influence graph has the FBRP.*

Theorem 4 has important practical implications. Recall that based on the results for a general graph topology

shown in Section III-B it was advisable to ensure that nodes would update their allocations one at a time and in a random order so as to ensure that an equilibrium allocation can be reached; the number of update steps was unbounded in general but finite on average. By Theorem 4, if the influence graph is complete and the nodes that compete for the resources update their allocations one at a time by using some synchronization protocol, then they will reach an equilibrium allocation after a *bounded* number of update steps starting from an arbitrary initial allocation independent of the order of the updates. Since the order of updates can be arbitrary (e.g., fixed), the system would require less coordination between the nodes.

V. THE IMPORTANCE OF THE UTILITY STRUCTURE

In this section we allow an arbitrary influence graph but we introduce a constraint on the payoff structure: we consider the case when $\delta_i = 0$. This case was used in the context of replication to model the problem of cooperative caching between peering ISPs in [2], and in the context of radio spectrum allocation $\delta_i = 0$ corresponds to the case when interference leads to zero expected throughput (i.e., proper multi-coloring). Throughout the section we consider a notion of improvement paths that we refer to as lazy.

Definition 4. A step $a_i(t+1)$ of player i is an *improvement step* if $U_i(a_i(t+1), a_{-i}(t)) > U_i(a_i(t), a_{-i}(t))$. A *lazy improvement step* of player i is an improvement step such that the payoff of every inserted resource exceeds that of every evicted resource. That is, for every $r \in E_i(t+1)$ and $r' \in I_i(t+1)$ we have $U_i^r(1, a_{-i}^r(t)) < U_i^{r'}(1, a_{-i}^{r'}(t))$.

Clearly, every best reply improvement step is a lazy improvement step. The following result shows that all lazy improvement paths are finite for arbitrary graph topologies if $\delta_i = 0$.

Proposition 5. *In a graphical resource allocation game with $\delta_i = 0 \forall i \in N$ every lazy improvement path is finite, i.e., the game possesses the finite lazy improvement property.*

Proof: We prove the proposition by showing that under the condition $\delta_i = 0$ the resource allocation game has a generalized ordinal potential function [13] for lazy improvement steps. A function $\Psi : \times_i(\mathcal{A}_i) \rightarrow \mathbb{R}$ is a generalized ordinal potential function for the game if the change of Ψ is strictly positive if an arbitrary player i increases its utility by changing her strategy from a_i to a'_i . Formally, $U_i(a_i, a_{-i}) - U_i(a'_i, a_{-i}) > 0 \Rightarrow \Psi(a_i, a_{-i}) - \Psi(a'_i, a_{-i}) > 0$. In the following we show that the function

$$\Psi(\mathbf{a}) = \sum_i U_i(a_i, a_{-i}), \quad (6)$$

where the utility function was defined in (1), is a generalized ordinal potential for the game. Substituting $\delta_i = 0$ into (1), one notes that player i benefits only from allocating *i-free* resources. Furthermore, the utility of player i does not depend on resources that she does

not allocate herself. Given a strategy profile $\mathbf{a} = (a_i, a_{-i})$ player i can improve its utility by combining three kinds of lazy improvement steps.

First, if player i has free storage capacity (that is, $\sum_r a_i^r < K_i$) then she has to allocate a resource r for which $a_i^r = 0$ but $U_i^r(1, a_{-i}^r) > 0$. By (1) we know that resource r is *i-free* and hence the utility of her neighbors will not be affected if she allocates resource r . Consequently, if we denote the new strategy of player i by $a_i' = (a_i^1, \dots, a_i^{r-1}, 1, a_i^{r+1}, \dots, a_i^{|\mathcal{R}|})$ then

$$\begin{aligned} \Psi(a_i', a_{-i}) - \Psi(a_i, a_{-i}) &= U_i(a_i', a_{-i}) - U_i(a_i, a_{-i}) \\ &= U_i^r(1, a_{-i}^r) > 0. \end{aligned} \quad (7)$$

Second, if player i stops allocating a resource r for which $U_i^r(1, a_{-i}^r) = 0$ and starts allocating a resource r' for which $U_i^{r'}(1, a_{-i}^{r'}) > 0$. By (1) we know that resource r is *i-busy*, but resource r' is *i-free*. Let us denote the strategy of player i after the change by a_i' . We first observe that the utility of the neighboring players cannot decrease when player i stops allocating resource r (it can potentially increase). At the same time the utility of the neighboring players does not change when player i starts allocating resource r' . Hence we have that

$$\Psi(a_i', a_{-i}) - \Psi(a_i, a_{-i}) \geq U_i^{r'}(1, a_{-i}^{r'}) > 0. \quad (8)$$

Third, if player i stops allocating a resource r for which $U_i^r(1, a_{-i}^r) > 0$ and starts allocating a resource r' for which $U_i^{r'}(1, a_{-i}^{r'}) > U_i^r(1, a_{-i}^r)$. By (1) we know that both resource r and r' are *i-free*. Hence the utility of the neighboring players is not affected by the change. The utility of player i increases, however. Let us denote the strategy of player i after the change by a_i' , then

$$\Psi(a_i', a_{-i}) - \Psi(a_i, a_{-i}) = U_i^{r'}(1, a_{-i}^{r'}) - U_i^r(1, a_{-i}^r) > 0. \quad (9)$$

By summing (7)-(9) we can see that the function Ψ is a generalized ordinal potential function for an arbitrary combination of lazy improvement steps. Finite games that allow a generalized ordinal potential function were shown to have the finite improvement property in [13]. Following the arguments of (Lemma 2.3, [13]) it follows that the resource allocation game has the finite lazy improvement property. We note that if non-lazy improvement steps are allowed, then a player can allocate an *i-busy* resource as part of an improvement step and thus the proof would not hold. ■

The finite lazy improvement property shown in Proposition 5 allows more freedom for system design than the results in Theorems 3 and 4 that are valid for best reply updates. Based on Proposition 5 an update step can be as simple as evicting the worst allocated resource and inserting the best non-allocated resource, and the system is guaranteed to reach an equilibrium for an arbitrary graph topology in a bounded number of steps.

VI. FAST CONVERGENCE BASED ON INDEPENDENT SETS

Until now we considered asynchronous updates. Unfortunately, the implementation of the asynchronous update rule in a distributed system would require global synchronization, which can be impractical in large distributed systems. Hence, an important question is whether the previous results of convergence to a NE would be preserved even if some players would update their strategies simultaneously. In the following we show that relaxing the requirement of asynchronicity is indeed possible: both Theorem 3 and Proposition 5 hold if the players that simultaneously make an update at each step form an independent set of the influence graph \mathcal{G} , that is, two players $i, j \in N$ can make an update simultaneously only if $j \notin \mathcal{N}(i)$. We refer to this dynamic as the *plesiochronous* dynamic.

Proposition 6. *Every resource allocation game is weakly acyclic under plesiochronous best replies.*

The proof of the proposition is analogous to the proof of Theorem 3.

Proposition 7. *In a graphical resource allocation game with $\delta_i = 0 \forall i \in N$ every plesiochronous lazy improvement path is finite.*

Proof: Consider a sequence of the subsets of the players $N^*(t) \subseteq N$ ($t = 0, \dots$) such that for $i, j \in N^*(t)$ we have $j \notin \mathcal{N}(i)$. Observe that, by definition, each set $N^*(t)$ is an independent set of the influence graph \mathcal{G} . The players $i \in N^*(t)$ make an improvement step at step t simultaneously from $a_i(t-1)$ to $a_i(t)$. Each player can combine the three kinds of lazy improvement steps discussed in the proof of Proposition 5 to increase its utility.

Since we require that none of the players that update their strategies are neighbors of each other, then their updates do not affect each others' utilities. Formally, for $i \in N^*(t)$ we have $U_i(a_i(t), a_{-i}(t)) = U_i(a_i(t), a_{-i}(t-1))$. Consequently, we can use the same arguments as in the proof of Proposition 5 to show for every $i \in N^*(t)$ that Ψ defined in (6) satisfies $U_i(a_i(t), a_{-i}(t)) - U_i(a_i'(t-1), a_{-i}(t-1)) > 0 \Rightarrow \Psi(a(t)) - \Psi(a(t-1)) > 0$.

The rest of the proof is similar to that of Proposition 5. ■

In order to maximize the convergence speed of the plesiochronous dynamic we need to find a minimum vertex coloring of \mathcal{G} , i.e., we have to find the chromatic number $\chi(\mathcal{G})$ of graph \mathcal{G} . Finding the chromatic number is NP-hard in general, but efficient distributed graph coloring algorithms exist [24], and can be used to find a coloring in a distributed system. The chromatic number can be bounded based on the largest eigenvalue $\lambda_{max}(\mathcal{G})$ of the graph's adjacency matrix [25],

$$\chi(\mathcal{G}) \leq 1 + \lambda_{max}(\mathcal{G}) \leq \max_{i \in N} |\mathcal{N}(i)|$$

where the second inequality follows from the Perron-Frobenius theorem ([26], Lemma 2.8). Given a coloring,

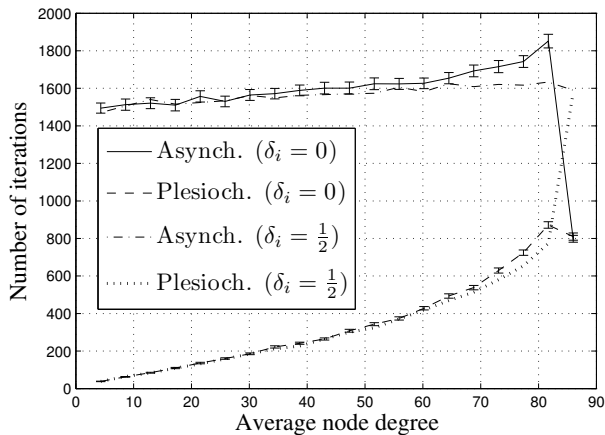


Fig. 2: Average number of improvement steps needed to reach a NE for asynchronous and plesiochronous dynamics vs. the average node degree of the ER random graphs.

the number of steps needed to reach equilibrium can be significantly smaller than for the corresponding asynchronous dynamic for sparse graphs.

We illustrate the convergence speedup of the plesiochronous dynamic compared to the asynchronous dynamic in Figures 2 and 3. For the plesiochronous dynamic we used the Welsh-Powell algorithm to find a coloring [27] of the influence graphs and we denote the number of colors by $\xi(\mathcal{G})$. Each player can allocate $K_i = 5$ resources and we considered two scenarios: $\delta_i = 0$ and $\delta_i = 0.5$. Figure 2 shows the average number of lazy improvement steps needed to reach equilibrium as a function of the average node degree in Erdős-Rényi random graphs with 87 vertices. Each data point is the average of the results obtained on 160 random graphs with the same average node degree. The figure shows the 95% confidence intervals for the case $\delta_i = 0$. We omitted the confidence intervals for $\delta_i = 0.5$ to improve readability. The results confirm that the plesiochronous dynamic converges significantly faster than the asynchronous dynamic, especially over sparse influence graphs. The figure also confirms that the convergence properties are different on a complete influence graph than on a sparse graph, as the number of steps necessary for the asynchronous dynamic to reach a NE drops for average node degree equal to 86.

Figure 3 shows the speedup (the ratio of the number of steps needed to reach equilibrium under the asynchronous dynamic and the plesiochronous dynamic) as a function of the average node degree for Erdős-Rényi (ER) and for Barabási-Albert (BA) random graphs with 87 vertices for $\delta_i = \frac{1}{2}$. Each data point is the average of the results obtained on 160 random graphs with the same average node degree. The results show that the speedup for BA random graphs is slightly smaller than for ER random graphs but shows a similar trend. A comparison of the speedup with the average size of the independent sets after the coloring of the random graphs ($\frac{|N|}{\xi(\mathcal{G})}$) indicates that the speedup of the plesiochronous dynamic exceeds the aver-

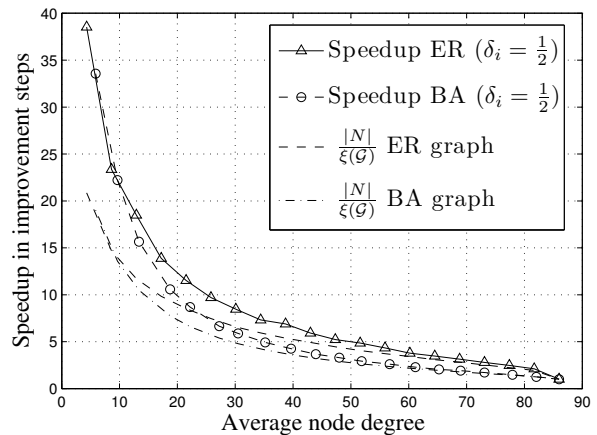


Fig. 3: Speedup in terms of the number of improvement steps needed to reach a NE vs. the average node degree of the ER and the BA random graphs.

age number of players that can perform an improvement step simultaneously.

The results formulated in Propositions 6 and 7 allow one to simplify the design of large systems significantly. The convergence results of Sections III, IV and V required that only one node at a time would update its allocation. In order to ensure this, there needs to be a coordination protocol for the entire system that ensures that there be only one node at a time that updates its allocation. In contrast, Propositions 6 and 7 show that local coordination is sufficient to ensure that the system would reach an equilibrium allocation, and the numerical results show that convergence is in fact very fast. In practice the information needed for the coordination can be piggybacked with the information about the updated allocations, and thus an equilibrium allocation can be reached with very little communication overhead.

VII. RELATED WORK

There is a large body of works on congestion games [5] on complete influence graphs. Most works consider congestion games that allow a potential function [13], and analyze the number of steps needed to reach equilibria [9], the price of anarchy and the price of stability [10], [11], or the complexity of calculating equilibria [12].

Player-specific congestion games on a complete influence graph that do not admit a potential function were considered in [14], [20], [15], [16]. In [14] it was shown that for non-weighted player-specific congestion games with *singleton action sets* the best reply paths are finite. In [20] the authors showed the existence of equilibria for a game of replication. In [15], the authors considered congestion games with player-specific constants, which correspond to all players having the same cost of sharing ($\delta_i = \delta$) in our model, and showed that improvement paths are finite in the unweighted version of the game. In [16] the authors showed that player-specific congestion games with matroid action sets are weakly acyclic in better replies, and provided bounds on the length of the shortest improvement

paths. They also showed that games with 2 players, or with 2 actions per player are acyclic in best replies, similar to [14].

A few recent works considered graphical congestion games that allow potential functions [6], [7], [8]. [6] gave results on the price of anarchy and stability for games with linear payoff functions, and showed that cycles can exist if the influence graph is directed and contains cycles. [7] addressed similar questions for weighted graphical congestion games with linear payoff functions. In [8] the authors analyzed the number of steps needed to reach equilibrium in graphical congestion games with homogeneous resources and singleton action sets.

The resource allocation game we consider combines the concept of graphical congestion games with player-specific payoffs on non-singleton action sets. Our results on equilibrium existence, and the results on convergence rely on non-standard techniques and could be of interest for the analysis of congestion games that do not admit a potential function. Furthermore, the proposed plesiochronous update dynamic based on independent sets is a promising candidate for implementation in large distributed systems.

VIII. CONCLUSION

In this work we considered a player-specific graphical resource allocation game, a resource allocation game played over an influence graph. The game models a system in which every node can choose a subset of resources, and the value of a selected resource to a node is decreased if any of its neighbors chooses the same resource. We showed that pure strategy Nash equilibria exist in the game for arbitrary influence graph topologies even though the game does not admit a potential function, and gave a bound on the complexity of finding equilibria. We then considered the problem of reaching an equilibrium when nodes update their allocations one at a time to their best allocation, that is, every node performs a best reply. We showed that for non-complete influence graphs there might be cycles in best replies but the system would reach an equilibrium eventually if updates are done in a random order. For complete influence graphs there are no cycles, and hence an equilibrium is reached after a finite number of steps. We also showed that if neighboring nodes try to allocate disjoint sets of resources then there are no cycles, even if the nodes' update steps are improvements instead of best replies. Finally, we showed that the convergence properties hold even if improvement steps are performed simultaneously by nodes that form an independent set of the influence graph, and proposed an efficient algorithm to reach a NE over sparse influence graphs that only requires local coordination between neighboring nodes.

REFERENCES

- [1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proc. of ACM CoNEXT*, 2009.
- [2] G. Dán, "Cache-to-cache: Could ISPs cooperate to decrease peer-to-peer content distribution costs?" *IEEE Trans. Parallel Distrib. Syst.*, vol. 22, no. 9, pp. 2469–2482, 2011.
- [3] L. Narayanan, "Channel assignment and graph multicoloring," in *Handbook of wireless networks and mobile computing*, 2002, pp. 71–94.
- [4] K. I. Aardal, V. Hoesel, P. M. Stan, A. M. C. A. Koster, C. Mannino, and A. Sassano, "Models and solution techniques for frequency assignment problems," *Annals of Operations Research*, vol. 153, no. 1, pp. 79–129, 2007.
- [5] R. W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.
- [6] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli, "Graphical Congestion Games," in *Proc. of WINE*, vol. 5385, 2008, pp. 70–81.
- [7] D. Fotakis, V. Gkatzelis, A. C. Kaporis, and P. G. Spirakis, "The Impact of Social Ignorance on Weighted Congestion Games," in *Proc. of WINE*, 2009, pp. 316–327.
- [8] R. Southwell and J. Huang, "Convergence Dynamics of Resource-Homogeneous Congestion Games," in *Proc. of GameNets*, 2011.
- [9] S. Chien and A. Sinclair, "Convergence to approximate Nash equilibria in congestion games," in *Proc. of the 18th annual ACM-SIAM Symposium on Discrete Algorithms*, 2007, pp. 169–178.
- [10] G. Christodoulou and E. Koutsoupias, "The price of anarchy of finite congestion games," in *Proc. of ACM Symposium on Theory of Computing*, 2005, pp. 67–73.
- [11] B. Awerbuch, Y. Azar, and A. Epstein, "The price of routing unsplittable flow," in *Proc. of ACM Symposium on Theory of Computing*, 2005, pp. 57–66.
- [12] A. Fabrikant, C. Papadimitriou, and K. Talwar, "The complexity of pure Nash equilibria," in *Proc. of ACM Symposium on Theory of Computing*, 2004, pp. 604–612.
- [13] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, pp. 124–143, 1996.
- [14] I. Milchtaich, "Congestion games with player-specific payoff functions," *Games and Economic Behavior*, vol. 13, no. 1, pp. 111–124, 1996.
- [15] M. Mavronicolas, I. Milchtaich, B. Monien, and K. Tiemann, "Congestion games with player-specific constants," *Mathematical Foundations of Computer Science 2007*, pp. 633–644, 2007.
- [16] H. Ackermann, H. Röglin, and B. Vöcking, "Pure Nash equilibria in player-specific and weighted congestion games," *Theor. Computer Science*, vol. 410, no. 17, pp. 1552–1563, 2009.
- [17] A. Leff, J. L. Wolf, and P. S. Yu, "Replication Algorithms in a Remote Caching Architecture." *IEEE Trans. Parallel Distrib. Syst.*, vol. 4, no. 11, pp. 1185–1204, 1993.
- [18] M. M. Halldórsson and G. Kortsarz, "Tools for Multicoloring with Applications to Planar Graphs and Partial k-Trees," *Journal of Algorithms*, vol. 42, no. 2, pp. 334–366, 2002.
- [19] H. Huang, B. Zhang, G. Chan, G. Cheung, and P. Frossard, "Coding and Replication Co-Design for Interactive Multiview Video Streaming," in *Proc. of mini-conference in IEEE INFOCOM*, 2012, pp. 1–5.
- [20] N. Laoutaris, O. Telelis, V. Zissimopoulos, and I. Stavrakakis, "Distributed selfish replication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 12, pp. 1401–1413, 2006.
- [21] L. A. Belady, R. A. Nelson, and G. S. Shedler, "An Anomaly in Space-Time Characteristics of Certain Programs Running in a Paging Machine," in *Communications of the ACM*, vol. 12, no. 6, 1969, pp. 349–353.
- [22] H. P. Young, "The evolution of conventions," *Econometrica: Journal of the Econometric Society*, vol. 61, no. 1, pp. 57–84, 1993.
- [23] J. R. Marden, G. Arslan, and J. S. Shamma, "Regret Based Dynamics: Convergence in Weakly Acyclic Games," in *Proc. of AAMAS*, 2007, pp. 194–201.
- [24] F. Kuhn, "Weak Graph Colorings: Distributed Algorithms and Applications," in *ACM SPAA*, 2009.
- [25] H. S. Wilf, "The eigenvalues of a graph and its chromatic number," *J. London Math. Soc.*, vol. 42, pp. 330–332, 1967.
- [26] R. S. Varga, *Matrix Iterative Analysis*, 2000.
- [27] D. J. A. Welsh and M. B. Powell, "An upper bound for the chromatic number of a graph and its application to timetabling problems," *The Computer Journal*, vol. 10, no. 1, pp. 85–86, 1967.

APPENDIX

Proof of Lemma 2: According to the structure of the utility function $U_i(a_i, a_{-i}) = \sum_{\{r|a_i^r=1\}} U_i^r(1, a_{-i})$, a best reply $a_i(t)$ can stop to be such in two situations:

- (i) The payoffs of one or more allocated resources $\{r \in \mathcal{R}|a_i^r(t) = 1\}$ decrease;
- (ii) The payoffs of one or more not allocated resources $\{r \in \mathcal{R}|a_i^r(t) = 0\}$ increase;

According to the definition of cost saving in (1), case (i) can happen only if some i -free resources allocated by i become i -busy. This requires that some player $j \in \mathcal{N}(i)$ starts allocating some j -busy resources. Similarly, case (ii) can happen only if a neighbor $j \in \mathcal{N}(i)$ allocating an resource r evicts it, making resource r i -free. ■

Proof of Lemma 3: We prove the lemma by showing that a best reply satisfying condition 1) must occur in a best reply cycle in order for the cycle to exist. The utility of at least one player must decrease at least once in a best reply cycle. According to the definition of cost saving in (1), in order for the utility $U_j(a(t))$ of a player j to decrease, some neighbor $i \in \mathcal{N}(j)$ needs to start allocating some i -busy resource allocated by j . It follows that, from Table I, in a best reply cycle there has to be at least one best reply satisfying condition 1) or 3). Let $r \in E_i(t)$ and $r' \in I_i(t)$ be two resources in the evicted and inserted sets, respectively, during a best reply of player i at step t in a best reply cycle. It follows from the definition of inserted and evicted set that $a_i^r(t-1) = 1$ and $a_i^{r'}(t-1) = 0$. Assume now that this best reply satisfies 3). This implies $c_{ir} < c_{ir'}\delta_i$. Note that this inequality also implies that player i always prefers r' over r (r' always provides a higher cost saving) and thus if $a_i^r(t-1) = 1$, then $a_i^{r'}(t-1) = 1$. Since there cannot be a best reply satisfying 3), there must be at least one best reply satisfying 1). This proves the lemma. ■

Proof of Lemma 4: Since r' is i -busy it yields to player i the payoff $c_{ir'}\delta_i$. Consider an object r'' for which $c_{ir''} > c_{ir'}$. Since player i is performing a best reply, if r'' is i -free then its payoff is $c_{ir''} > c_{ir'}\delta_i$ and consequently $a_i^{r''}(t) = 1$. Similarly, if r'' is i -busy then its payoff is $c_{ir''}\delta_i > c_{ir'}\delta_i$, and consequently $a_i^{r''}(t) = 1$. ■

Proof of Lemma 5: Part A: First we show that $|B_i(t-1)| \geq |B_i(t)|$. Player i can only increase the number of i -busy allocated resources if she evicts at least one i -free resource r' from $a_i(t-1)$ and inserts an i -busy resource r at step t . Thus by (1) we have

$$c_{ir'} < c_{ir}\delta_i \quad (10)$$

Since we are in a best reply cycle, at some step $t' > t$ the strategy $a_i(t-1)$ must become a best reply for player i ,

i.e. $a_i(t') = a_i(t-1)$. This requires either $c_{ir'} > c_{ir}$ or $c_{ir'} > c_{ir}\delta_i$, and both contradict (10). *Part B:* Second we show that for every step in a best reply cycle $|B_i(t-1)| = |B_i(t)|$ must hold. We denote by $C(t)$ the set of the chosen resources, the resources allocated by at least one player in $a(t)$, $C(t) = \{r|a_j^r(t) = 1 \text{ for some } j \in N\} \subseteq \mathcal{R}$.

Similarly, we denote by $C(t)_{-i}$ the set of resources allocated by the players not including i , $C(t)_{-i} = \{r|a_j^r(t) = 1 \text{ for some } j \in N \setminus \{i\}\}$. It is easy to see that the sets $|F_i(t)|$ and $|C(t)|$ are related and on a complete influence graph it holds that $|C(t)| = |C(t)_{-i}| + |F_i(t)|$. On one hand, a best reply for which $|F_i(t-1)| = |F_i(t)|$ does not affect $|C(t)|$ since $|C(t-1)_{-i}| = |C(t)_{-i}|$. On the other hand, a best reply for which $|F_i(t-1)| > |F_i(t)|$ decreases the size of set C

$$\begin{aligned} |C(t)| &= |C(t)_{-i}| + |F_i(t)| = |C(t-1)_{-i}| + |F_i(t)| \\ &< |C(t-1)_{-i}| + |F_i(t-1)| = |C(t-1)| \end{aligned}$$

Since best replies for which $|F_i(t-1)| > |F_i(t)|$ do not exist in a cycle (*Part A*), best replies for which $|F_i(t-1)| < |F_i(t)|$ cannot exist either, as otherwise the size of set C would increase indefinitely. Hence $|F_i(t-1)| = |F_i(t)|$, which proves the lemma. ■



Valentino Pacifici studied computer engineering at Politecnico di Milano in Milan, Italy. In October 2010, he completed a joint M.Sc. degree in computer engineering, between Politecnico di Milano and KTH, Royal Institute of Technology in Stockholm, Sweden. Now he is a researcher at the Laboratory for Communication Networks in KTH, Royal Institute of Technology and he is pursuing his Ph.D. His research interests include the modeling, design and analysis of content management systems using game theoretical tools.



György Dán received the M.Sc. degree in computer engineering from the Budapest University of Technology and Economics, Hungary in 1999 and the M.Sc. degree in business administration from the Corvinus University of Budapest, Hungary in 2003. He worked as a consultant in the field of access networks, streaming media and videoconferencing 1999-2001. He received his Ph.D. in Telecommunications in 2006 from KTH, Royal Institute of Technology, Stockholm, Sweden, where he currently works as an assistant professor. He was visiting researcher at the Swedish Institute of Computer Science in 2008. His research interests include the design and analysis of distributed and peer-to-peer systems and cyber-physical system security.