

Selfish Content Replication on Graphs

Valentino Pacifici and György Dán
ACCESS Linnaeus Center, School of Electrical Engineering
KTH, Royal Institute of Technology, Stockholm, Sweden
E-mail: {pacifici,gyuri}@ee.kth.se

Abstract—Replication games are a model of the problem of content placement in computer and communication systems, when the participating nodes make their decisions such as to maximize their individual utilities. In this paper we consider replication games played over arbitrary social graphs; the social graph models limited interaction between the players due to, e.g., the network topology. We show that in replication games there is an equilibrium object placement for arbitrary social graphs. Nevertheless, if all nodes follow a myopic strategy to update their object placements then they might cycle arbitrarily long before reaching an equilibrium if the social graph is non-complete. We give sufficient conditions under which such cycles do not exist, and propose an efficient distributed algorithm to reach an equilibrium over a non-complete social graph.

I. INTRODUCTION

Replicating content close to its users has long been used in computer architectures and in the Internet to improve system performance. In computer architectures CPU caches have been used to decrease memory access latencies [1]. In the Internet caches are employed to provide faster access to content for local customers and at the same time to decrease the amount of network traffic [2], [3], [4]. Content replication is also at the core of clean-slate information centric network architectures [5].

The problem of content replication is often modeled by a distributed replication group [6], [7], [8], [9]. A replication group consists of nodes and users located near the nodes. The nodes can replicate objects, which are accessed by local or remote users. The cost incurred by a user accessing an object depends on whether the object is replicated locally, in a remote node, or is not replicated in any node. A commonly studied problem of replication is how to minimize the total cost of users accessing the objects through either optimal node placement [6] or through the optimal allocation of objects to nodes [7].

Often there is no central authority that would be able to enforce the optimal solution on the nodes. This can be the case when nodes are autonomous, for example, in the case of Internet Web and peer-to-peer content replication [3], [4]. In lack of a central authority nodes would not implement the optimal solution if they can individually fare better by deviating from it [8], [9], [10]. Instead, nodes would replicate the objects that minimize their own costs, and would update the set of replicated objects as a response to the decisions made by other nodes. An

important question in this case is whether there exists an equilibrium state in terms of the objects replicated by the nodes from which no node has an interest to deviate. A similarly important question is whether in a distributed system the nodes would be able to reach an equilibrium state if each node follows a myopic strategy to minimize its own cost. Answering these questions is a key to the design of efficient distributed replication algorithms.

In this paper we model the problem of selfish replication as a non-cooperative graphical game. The model of a graphical game allows us to capture the limited interactions between nodes imposed by an underlying communication network [10], modeled by a so called *social graph*. The social graph defines neighborhood relationships between the nodes, which influence the cost of accessing replicas. We show that equilibrium states exist for arbitrary social graphs, but nodes might cycle arbitrarily long before reaching an equilibrium state if the social graph is non-complete. Based on our results 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 structure of the paper is as follows. In Section II we define the problem of replication and graphical replication games. In Section III we discuss the existence of equilibrium states, and in Section IV we address the question whether nodes that follow simple learning rules would reach an equilibrium state. In Section V we show that an efficient distributed algorithm can be used to reach an equilibrium state under certain conditions. In Section VI we discuss related work, and Section VII concludes the paper.

II. SYSTEM MODEL

In the following we describe the system model, and formulate the problem of replication as a non-cooperative game.

A. The replication problem

We consider a set of nodes N and a set of objects \mathcal{O} . The demand for object $o \in \mathcal{O}$ at node $i \in N$ is given by the rate $w_i^o \in \mathbb{R}_+$. We consider that every object $o \in \mathcal{O}$ has unit size, $S^o = 1$, which is a reasonable simplification if objects are divisible into unit-sized chunks. Node i has integer storage capacity $K_i \in \mathbb{N}_+$, which it uses to replicate objects locally. We describe the set of objects replicated at

node i with the $|\mathcal{O}|$ dimensional vector $r_i = (r_i^1, \dots, r_i^{|\mathcal{O}|})$, whose component $r_i^o \in \{0, 1\}$ is 1 if object o is replicated in node i . Due to the limited storage capacity $\sum_o r_i^o S^o \leq K_i$, thus the set of feasible replication vectors for node i is $\mathcal{R}_i = \{r_i \mid \sum_o r_i^o S^o \leq K_i\} \subseteq \{0, 1\}^{|\mathcal{O}|}$.

Every node is located at a vertex of an undirected graph $\Gamma(N, E)$, called the social graph. We denote the set of neighbors of node i by $\mathcal{N}(i)$, i.e., $\mathcal{N}(i) = \{j \mid (i, j) \in E\}$. The social graph allows us to consider a generalized version of the cost model described in [1], whose variations were used in [8], [9], [10]. In our model the marginal cost of accessing object o in node i is α_i if the object is replicated in node i , it is β_i if it is replicated in a node $j \in \mathcal{N}(i)$ neighboring i , and it is γ_i otherwise. We consider the practically relevant case when it is not more costly to access an object replicated locally than one replicated at a neighbor, and it is less costly to access an object replicated at a neighbor than retrieving it directly from the common set of objects. Formally $\alpha_i \leq \beta_i < \gamma_i$, or equivalently

$$0 \leq \delta_i \triangleq \frac{\beta_i - \alpha_i}{\gamma_i - \alpha_i} < 1. \quad (1)$$

To ease notation we say that an object $o \in \mathcal{O}$ is *i-available* if it is replicated by at least one of node i 's neighbors in which case

$$\pi_i^o \triangleq \prod_{j \in \mathcal{N}(i)} (1 - r_j^o) = 0,$$

otherwise we say that object $o \in \mathcal{O}$ is *not i-available*.

The cost of node i due to object o is proportional to the demand w_i^o , and is a function of r_i and the replication states r_{-i}^o of the neighboring nodes

$$C_i^o(r_i^o, r_{-i}^o) = w_i^o (\alpha_i r_i^o + (1 - r_i^o) [\gamma_i \pi_i^o + \beta_i (1 - \pi_i^o)]) \quad (2)$$

B. The graphical replication game

We consider a system in which the goal of every node is to minimize its own total cost. We model this problem of selfish replication as a multiplayer non-cooperative game played on a graph, called a graphical game. The players are the nodes, the set of actions of player i is \mathcal{R}_i , and the cost function of player i is given by $C_i(r_i, r_{-i}) = \sum_o C_i^o(r_i^o, r_{-i}^o)$. The social graph influences the cost function via the neighbor set $\mathcal{N}(i)$, i.e., the cost function of player i is entirely specified by the actions of players $j \in \mathcal{N}(i)$.

The goal of player i is to choose a replication strategy r_i that minimizes its total cost given the strategy profile r_{-i} of the other players

$$\arg \min_{r_i} C_i(r_i, r_{-i}) = \arg \min_{r_i} \sum_o C_i^o(r_i^o, r_{-i}^o). \quad (3)$$

Observe that the cost of object o for player i can be expressed as

$$\begin{aligned} C_i^o(r_i^o, r_{-i}^o) &= C_i^o(0, r_{-i}^o) - (C_i^o(0, r_{-i}^o) - C_i^o(r_i^o, r_{-i}^o)) \\ &= C_i^o(0, r_{-i}^o) - U_i^o(r_i^o, r_{-i}^o), \end{aligned}$$

where the utility $U_i^o(r_i^o, r_{-i}^o)$ is the cost saving that player i achieves through object o given the other players' replication strategies. The utility function of player i is the sum of the cost savings $U_i(r_i, r_{-i}) = \sum_o U_i^o(r_i^o, r_{-i}^o)$.

Since the cost $C_i^o(0, r_{-i}^o)$ is independent of the action r_i^o of player i , finding the minimum cost is equivalent to finding the maximum utility

$$\begin{aligned} \arg \min_{r_i} C_i(r_i, r_{-i}) &= \arg \min_{r_i} \sum_o C_i^o(r_i^o, r_{-i}^o) \\ &= \arg \min_{r_i} \left(\sum_o C_i^o(0, r_{-i}^o) - \sum_o U_i^o(r_i^o, r_{-i}^o) \right) \\ &= \arg \max_{r_i} \sum_o U_i^o(r_i^o, r_{-i}^o). \end{aligned}$$

Consequently, the problem of replication can be modeled by the strategic game $\mathcal{G} = \langle N, (\mathcal{R}_i), (U_i) \rangle$.

We can express the utility $U_i^o(r_i^o, r_{-i}^o)$ of player i by substituting (2) into the definition of the utility

$$\begin{aligned} U_i^o(r_i^o, r_{-i}^o) &= w_i^o [\beta_i (1 - \pi_i^o) + \gamma_i \pi_i^o] - \\ &- w_i^o (\alpha_i r_i^o + (1 - r_i^o) [\beta_i (1 - \pi_i^o) + \gamma_i \pi_i^o]) \quad (4) \\ &= r_i^o w_i^o [\beta_i (1 - \pi_i^o) + \gamma_i \pi_i^o - \alpha_i]. \quad (5) \end{aligned}$$

We note a property of the utility defined in (5): the utility of player i due to object o is independent of the other players' strategies if she does not replicate the object, i.e., $r_i^o = 0 \Rightarrow U_i^o(0, r_{-i}^o) = 0$. If player i replicates object o then the cost saving is

$$U_i^o(1, r_{-i}^o) = \begin{cases} w_i^o [\gamma_i - \alpha_i] = c_{io} & \text{if } \pi_i^o = 1 \\ w_i^o [\beta_i - \alpha_i] = \delta_i c_{io} & \text{if } \pi_i^o = 0 \end{cases} \quad (6)$$

Note that replicating a *not i-available* object o yields to player i a cost saving of c_{io} , while replicating an *i-available* object o yields $\delta_i c_{io} < c_{io}$.

III. EXISTENCE OF EQUILIBRIA

The first question we address is whether in a system of selfish nodes there is a state from which no node has an interest to deviate unilaterally. Such an equilibrium state is in fact a pure strategy *Nash equilibrium* (NE) of the graphical replication game. Hence to answer the question we have to find whether pure strategy NE exist in graphical replication games. It is known that for a complete social graph pure strategy NE exist in a replication game [9], but it is not known whether pure strategy NE exist for non-complete social graphs. In what follows we show that pure strategy NE exist for arbitrary social graphs.

We first define a *best reply* of player i as a best strategy r_i^* of player i given the other players' strategies

$$U_i(r_i^*, r_{-i}) \geq U_i(r_i, r_{-i}) \quad \forall r_i \in \mathcal{R}_i. \quad (7)$$

The NE is a strategy profile r^* in which every player's strategy is a best reply to the other players' strategies

$$U_i(r_i^*, r_{-i}^*) \geq U_i(r_i, r_{-i}^*) \quad \forall r_i \in \mathcal{R}_i, \quad \forall i \in N. \quad (8)$$

Finally, we define a *best reply path* as a sequence of strategy profiles, such that in every step t there is one player that strictly increases its utility by updating her strategy to a best reply $r_i(t)$ with respect to the other players' most recent strategies $r_{-i}(t-1)$. A best reply path terminates when no player can increase its utility, in which case a NE is reached. Hence, to show the existence of NE it is enough to show that there is a particular strategy profile starting from which there is at least one finite best reply path.

Consider the strategy profile $r(0)$ that consists of the best replies that the players would play on an edgeless social graph. In this strategy profile every player i replicates the K_i objects with highest demands w_i^o . Let us consider now a best reply path starting from the strategy profile $r(0)$. 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 on the players' best replies.

Lemma 1. *Player $i \in N$ inserts only not i -available objects when she first updates her strategy at $t = i$.*

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

- (i) A not i -available object o replicated by i ($r_i^o(t) = 1$, $\pi_i^o(t) = 1$) becomes i -available by step t' ,
- (ii) An i -available object o not replicated by i ($r_i^o(t) = 0$, $\pi_i^o(t) = 0$) becomes not i -available by step t' .

The proofs of the Lemmas are in the Appendix. A consequence of Lemma 1 is that an object replicated by player i cannot change from *not i -available* to *i -available* 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 replicating a *not i -available* object. Thus, by induction, condition (i) will never hold, and in every step t player $i \in N$ only inserts *not i -available* objects. Since no player ever inserts an *i -available* object, according to the definition of cost saving in (6) 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 Nash equilibrium. Hence we can state the following.

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

For a complete social graph every player makes at most one improvement step [9], but this does not hold even for a simple non-complete social graph: on a ring of 4 players with $w_i^o = w_j^o$ ($\forall o \in \mathcal{O}$) at least one player updates her strategy twice.

The existence of equilibrium states is important, but in a distributed system it is equally important that the nodes would be able to reach an equilibrium state using a distributed algorithm. One algorithm that the nodes could use to reach an equilibrium state is the algorithm used to prove Theorem 1. This algorithm can be adequate if the demands for the objects in the nodes w_i^o never change, so once a NE is reached, the nodes will not deviate from it. Nevertheless, the algorithm would be inefficient if the demands can change over time, as the equilibrium states for different demands 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 demands, and given the myopic decisions they make to update their strategies.

Consider, for example, that in every time step t every player i simultaneously updates her strategy to her best reply with respect to $r_{-i}(t-1)$. Such a synchronous update rule would require little coordination in a distributed system, but unfortunately it would be difficult to reach an equilibrium this way.

Example 1. *Consider two nodes and $|\mathcal{O}| \geq 2$. Let $c_{i1} > c_{i2} > c_{i1}\delta_i$ and $K_i = 1$. If the initial replication strategies are $r_i(0) = (1, 0)$ then after one step both players will have $r_i(1) = (0, 1)$. After the second step both players will have $r_i(2) = (1, 0)$, etc.*

As an alternative, consider a sequence of best reply steps as defined in Section III, the result of an asynchronous best reply dynamic. A natural question is whether all best reply paths are finite irrespective of the initial strategy profile $r(0)$ and the social graph. The following example shows that, unfortunately, the answer is no in general:

Example 2. *Consider a replication game played over the social graph shown in Figure 1, where $\mathcal{O} = \{A, B, C, D\}$ and $K_i = 1 \forall i \in N$. Each player $1 \leq i \leq 4$ has an object $o^* \in \mathcal{O}$ such that $c_{io^*} > c_{io} \forall o \neq o^*$, and at least one object $o' \in \mathcal{O}$ such that $c_{io'} > c_{io^*}\delta_i$. For players $5 \leq i \leq 8$ there is an object $o^* \in \mathcal{O}$ such that $c_{io^*}\delta_i > c_{io} \forall o \neq o^*$. An asynchronous best reply dynamic that cycles is shown in Table I. Observe that players $5 \leq i \leq 8$ always replicate the object that has the highest demand at their respective neighboring node $i - 4$.*

These negative examples raise the question under what conditions nodes would reach an equilibrium state in a finite number of steps. In the following we distinguish between complete and non-complete social graphs. We first prove that on a complete social graph the asynchronous best reply dynamic always results in a finite best reply path, and then we show that under certain conditions cycles do not exist even on non-complete social graphs.




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

Player	P1	P2	P3	P4	P5	P6	P7	P8
r(0)	A	B	D	A	B	C	D	A
r(1)	↓A	B	↓C	A	B	C	D	A
r(2)	↓B	B	C	↓A	B	C	D	A
r(3)	B	↓B	C	↓D	B	C	D	A
r(4)	↓B	↓C	C	D	B	C	D	A
r(5)	↓A	C	↓C	D	B	C	D	A
r(6)	A	↓C	↓D	D	B	C	D	A
r(7)	A	↓B	D	↓D	B	C	D	A
r(8)	A	B	D	↓A	B	C	D	A

TABLE I: Cycle in best replies for the game played over the social graph in Figure 1. An arrow shows the best reply played at each step.

A. Best reply dynamic on a complete social graph

In the following we show that if the social graph is complete then every best reply path is finite, that is, it cannot contain a cycle.

Theorem 2. *Every best reply path in a replication game played over a complete social graph is finite.*

Proof: We prove the theorem by contradiction. We assume that a best reply cycle exists in the dynamic. We then show that a strategy profile, in order to be in such a cycle, can only contain a well-defined subset of objects and that a best reply of a player i cannot change the number of i -available objects she is replicating. Following these constraints we show that such a cycle cannot exist.

Let us denote by $F_i(t)$ the set of *not* i -available objects and by $B_i(t)$ the set of i -available objects replicated by i at step t . Note that $|F_i(t)| + |B_i(t)| = K_i$. For a best reply step performed by player i at step t let us define the evicted set as $E_i(t) = \{o | r_i^o(t-1) = 1 \wedge r_i^o(t) = 0\}$ and the inserted set as $I_i(t) = \{o | r_i^o(t-1) = 0 \wedge r_i^o(t) = 1\}$.

Lemma 3. *A best reply step $r_i(t)$ performed by player i in a best reply cycle cannot affect the number of i -available objects replicated by i , that is,*

$$|B_i(t)| = |B_i(t-1)|.$$

Lemma 4. *For a best reply step $r_i(t)$ performed by player i in a best reply cycle it holds that*

$$w_i^{o'} < w_i^o \quad \forall o' \in E_i(t), o \in I_i(t).$$

The proofs of the lemmas are given in the Appendix. Consider now a best reply cycle. According to Lemma 4

each best reply step can only move towards objects with higher demand. The number of objects $|\mathcal{O}|$ 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. ■

By Theorem 2 we know that if the social graph is complete and nodes update their replication strategies to their best replies one at a time then they reach an equilibrium state without going through any state twice. By Example 2 we know, however, that if the social graph is non-complete then this is not necessarily the case.

B. The case of non-complete social graph

Given that cycles might exist in the best reply paths for non-complete social graphs, an important question is whether the players might cycle forever without being able to reach an equilibrium state. In the following we show that if we consider replication games with $K_i = 1$, then from every strategy profile there exists at least one finite best reply path that leads to a NE.

Consider a best reply of a player i at step t such that $E_i(t) = \{o\}$ and $I_i(t) = \{o'\}$. We can distinguish between four types of best replies depending on whether the involved objects are i -available as shown in Table II together with the cost savings of the objects.

Type	Evicted o	Inserted o'
1	i -available	$c_{io}\delta_i$
2	i -available	$c_{io'}\delta_i$
3	<i>not</i> i -available	c_{io}
4	<i>not</i> i -available	$c_{io'}$

TABLE II: Four possible types of a best reply move

The following two lemmas state that the strategy profiles in a best reply cycle cannot be arbitrary.

Lemma 5. *In a best reply cycle, if a player i inserts an i -available object o then*

- (i) o is the object with highest demand for player i .
- (ii) player i cannot insert an i -available object in her next best reply.

Lemma 6. *In every best reply cycle there exists at least one strategy profile $r(t)$ in which at least one player $i \in N$ can perform a best reply of type 1). Furthermore, no best reply cycle contains a best reply of type 3).*

The proofs of the lemmas are given in the Appendix. Let us consider a strategy profile $r(t)$ in a best reply cycle in which at least one player $i \in N$ can perform a best reply of type 1). Such a strategy profile exists according to Lemma 6. Starting from $r(t)$, let us perform a sequence of best replies of type 1). According to Lemma 5 (ii) every player can perform only one best reply of type 1) before performing a best reply of another type. There is a finite number of players, so after a finite number of best replies of type 1) we reach a strategy profile $r(t')$ in which there is no player that can perform a best reply of type 1).

If in $r(t')$ no player can make a best reply then $r(t')$ is a NE. Otherwise, if a player i can play a best reply in $r(t')$ then it must be of type 2) or of type 4). Due to a best reply of type 2) or of type 4) 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 replicating a *not j-available* object. Thus, by induction, condition (i) of Lemma 2 will never hold, and in every step $t'' > t'$ players only perform best replies of type 2) and of type 4). A player's utility strictly increases when it performs a best reply, and a best reply of type 2) and of type 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. *From every strategy profile of a graphical replication game with $K_i = 1 \forall i \in N$ there exists a best reply path that leads to a NE in a finite number of steps.*

To summarize, if the social graph is non-complete and $K_i = 1$, then there exists at least one finite best reply path to an equilibrium from every strategy profile, but not all best reply paths are finite because cycles can exist. That is, the game is weakly acyclic in best replies. In the following we show that if we introduce a small restriction on the marginal costs of accessing objects then cycles cannot exist even if the social graph is non-complete.

We define an *improvement step* of player i at step t as an update of her strategy $r_i(t)$ to $r_i(t+1)$, such that its utility increases

$$U_i(r_i(t+1), r_{-i}(t)) > U_i(r_i(t), r_{-i}(t)). \quad (9)$$

The *best reply step* defined in (7) is a special case of an improvement step. A sequence of improvement steps is called an *improvement path*. An improvement path terminates when no player can perform an improvement step, i.e., in an equilibrium.

In order to avoid cycles, we consider a subset of the set of improvement paths. We define a *lazy improvement step* of player i as an improvement step with minimal number of changes among all improvement steps that lead to the same utility. Formally, $r_i(t+1)$ is a lazy improvement step if there is no $r'_i(t+1) \neq r_i(t+1)$ for which

$$\begin{aligned} U_i(r_i(t+1), r_{-i}(t)) &= U_i(r'_i(t+1), r_{-i}(t)) \text{ and} \\ |I_i(t+1)| &> |I'_i(t+1)|. \end{aligned}$$

If players only make lazy improvement steps then $\beta_i = \alpha_i$ is a sufficient condition under which all improvement paths are finite. The case $\beta_i = \alpha_i$ was considered as a model of cooperative caching between peering ISPs in [10].

Proposition 4. *In a graphical replication game with $\beta_i = \alpha_i \forall i \in N$ every lazy improvement path is finite.*

Proof: We prove the proposition by showing that under the condition $\beta_i = \alpha_i$ the replication game has

a generalized ordinal potential function [11] for lazy improvement steps. A function $\Psi : \times_i(\mathcal{R}_i) \rightarrow \mathbb{R}$ 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 r_i to r'_i . Formally,

$$U_i(r_i, r_{-i}) - U_i(r'_i, r_{-i}) > 0 \Rightarrow \Psi(r_i, r_{-i}) - \Psi(r'_i, r_{-i}) > 0.$$

In the following we show that the function

$$\Psi(\mathbf{r}) = \sum_i U_i(r_i, r_{-i}), \quad (10)$$

where the utility function was defined in (5), is a generalized ordinal potential for the game. We substitute $\beta_i = \alpha_i$ into (5) to obtain the utility function of player i

$$U_i(r_i, r_{-i}) = \sum_o U_i^o(r_i^o, r_{-i}^o) = r_i^o w_i^o [\pi_i^o(\gamma_i - \alpha_i)]. \quad (11)$$

Note that player i benefits only from replicating *not i-available* objects. Furthermore, the utility of player i does not depend on objects that she does not replicate herself. Given a strategy profile $\mathbf{r} = (r_i, r_{-i})$ player i can improve its utility by combining three kinds of lazy improvement steps.

First, if player i has free storage capacity (that is, $\sum_o r_i^o < K_i$) then she has to choose an object o for replication for which $r_i^o = 0$ but $U_i^o(1, r_{-i}) > 0$. By (11) we know that object o is *not i-available* and hence the utility of her neighbors will not be affected if she replicates object o . Consequently, if we denote the new strategy of player i by $r'_i = (r_i^1, \dots, r_i^{o-1}, 1, r_i^{o+1}, \dots, r_i^{|\mathcal{O}|})$ then

$$\begin{aligned} \Psi(r'_i, r_{-i}) - \Psi(r_i, r_{-i}) &= U_i(r'_i, r_{-i}) - U_i(r_i, r_{-i}) \\ &= U_i^o(1, r_{-i}^o) > 0. \end{aligned} \quad (12)$$

Second, if player i stops replicating an object o for which $U_i^o(1, r_{-i}^o) = 0$ and starts replicating an object o' for which $U_i^{o'}(1, r_{-i}^{o'}) > 0$. By (11) we know that object o is *i-available*, but object o' is *not i-available*. Let us denote the strategy of player i after the change by r'_i . We first observe that the utility of the neighboring players cannot decrease when player i stops replicating object o (it can potentially increase). At the same time the utility of the neighboring players does not change when player i starts replicating object o' . Hence we have that

$$\Psi(r'_i, r_{-i}) - \Psi(r_i, r_{-i}) \geq U_i^{o'}(r'_i, r_{-i}) > 0. \quad (13)$$

Third, if player i stops replicating an object o for which $U_i^o(1, r_{-i}^o) > 0$ and starts replicating an object o' for which $U_i^{o'}(1, r_{-i}^{o'}) > U_i^o(1, r_{-i}^o)$. By (11) we know that neither object o nor o' are *i-available*. 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 r'_i , then

$$\Psi(r'_i, r_{-i}) - \Psi(r_i, r_{-i}) = U_i^{o'}(1, r_{-i}) - U_i^o(1, r_{-i}) > 0. \quad (14)$$

The function Ψ satisfies (10) for the three kinds of lazy improvement steps, and by summing (12)-(14) we can see that it does so for the combinations of the improvement steps too. Hence, Ψ is a generalized ordinal potential for the game under lazy improvement steps.

Finite games that allow a generalized ordinal potential function were shown to have the finite improvement property in [11]. Following the arguments of (Lemma 2.3, [11]) it follows that the replication game has the finite lazy improvement property. ■

It is easy to see that the above proof does not hold if non-lazy improvement steps are allowed.

V. FAST CONVERGENCE BASED ON GRAPH COLORING

In the previous section we showed that, under certain conditions, the players always reach a Nash equilibrium if they update their strategies asynchronously. 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 players would always reach a Nash equilibrium even if some players would update their strategies simultaneously. In the following we show that relaxing the requirement of asynchronicity is indeed possible.

Proposition 5. *Consider a graphical replication game with $\beta_i = \alpha_i \forall i \in N$. If player i makes an improvement step at time t only if no neighboring player $j \in \mathcal{N}(i)$ makes an improvement step at time t , then every 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 $j \notin \mathcal{N}(i)$. The players $i \in N^*(t)$ make an improvement step at step t simultaneously from $r_i(t-1)$ to $r_i(t)$. Each player can combine the three kinds of lazy improvement steps discussed in the proof of Proposition 4 to increase its utility.

Recall that for lazy improvement steps Ψ defined in (10) is a generalized ordinal potential function for the game if $\beta_i = \alpha_i$. 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

$$U_i(r_i(t), r_{-i}(t)) = U_i(r_i(t), r_{-i}(t-1)). \quad (15)$$

Consequently, we can use the same arguments as in the proof of Proposition 4 to show for every $i \in N^*(t)$ that

$$\begin{aligned} U_i(r_i(t), r_{-i}(t)) - U_i(r_i'(t-1), r_{-i}(t-1)) &> 0 \Rightarrow \\ \Psi(r(t)) - \Psi(r(t-1)) &> 0. \end{aligned}$$

Combining these yields that Ψ satisfies

$$\begin{aligned} U_i^o(r_i(t), r_{-i}(t)) - U_i^o(r_i'(t-1), r_{-i}(t-1)) &> 0 \quad \forall i \in N^*(t) \\ \Rightarrow \Psi(r(t)) - \Psi(r(t-1)) &> 0. \end{aligned}$$

At every improvement step the players in $N^*(t)$ increase their utilities, and hence the function Ψ increases. The function Ψ is bounded, hence it has to attain its maximum after a finite number of lazy improvement steps. ■

We refer to this dynamic as the *plesiochronous better reply dynamic* (PBRD), as opposed to the asynchronous dynamic (ABRD) considered in Proposition 4. In order to maximize the convergence speed of PBRD we need to find a minimum vertex coloring of Γ , i.e., we have to find the chromatic number $\chi(\Gamma)$ of graph Γ . Finding the chromatic number is NP-hard in general, but efficient distributed graph coloring algorithms exist [12], which can be used to find a coloring in a distributed system. Given a coloring, the number of steps required to reach equilibrium can be significantly smaller than for ABRD for sparse graphs.

The speedup of PBRD is bounded by the graph's chromatic number. In general, the chromatic number can be bounded based on the largest eigenvalue $\lambda_{max}(\Gamma)$ of the graph's adjacency matrix [13],

$$\chi(\Gamma) \leq 1 + \lambda_{max}(\Gamma), \quad (16)$$

and the largest eigenvalue can be bounded by the maximum node degree based on the Perron-Frobenius theorem ([14], Lemma 2.8)

$$\lambda_{max}(\Gamma) \leq \max_{i \in N} |\mathcal{N}(i)|. \quad (17)$$

In (17) equality holds for regular graphs ($|\mathcal{N}(i)| = |\mathcal{N}(j)|$ for $i \neq j$), and strict inequality holds otherwise. A stronger result can be obtained for random graphs. For almost every Erdős-Rényi random graph with $|N|$ vertices and edge probability p the chromatic number is [15]

$$\chi(\Gamma_{|N|,p}) = \left\lceil \frac{1}{2} + o(1) \right\rceil \log \frac{1}{1-p} \frac{|N|}{\log |N|}. \quad (18)$$

We illustrate the convergence speedup of PBRD compared to ABRD in Figure 2. The figure shows the average number of steps needed to reach equilibrium as a function of the edge probability in Erdős-Rényi random graphs with 87 vertices. For the PBRD we used the Welsh-Powell algorithm to find a coloring [16]. Each player had storage capacity $K_i = 5$ and we used $\gamma_i = 10$. We considered two scenarios: $\alpha_i = \beta_i$ and $\alpha_i \neq \beta_i$. Each data point is the average of the results obtained on 160 random graphs with the same edge probability. The figure shows the 95% confidence intervals for the case $\alpha_i = \beta_i$. We omitted the confidence intervals for $\alpha_i \neq \beta_i$ to improve readability.

The results confirm that PBRD converges significantly faster compared to ABRD, especially over sparse social graphs. The existence of cycles when $\alpha_i \neq \beta_i$ does not make a significant impact on the results. The figure also confirms that the convergence properties are different on a complete social graph than on a sparse graph, as the number of steps necessary for ABRD to reach a NE drops for $p = 1$. This observation is in accordance with the difference in computational complexity of finding the optimal object replication strategy [1]: the problem




Fig. 2: Average number of improvement steps needed to reach a NE for ABRD and PBRD vs. the edge probability of the Erdős-Rényi random graphs used as social graph.

is NP-complete on a non-complete social graph, but is polynomial in the number of players if the social graph is complete. A thorough analysis of the relationship between the complexity of finding the optimal solution and the convergence properties of improvement paths as a function of the graph topology is subject of our future work.

VI. RELATED WORK

Content replication has been considered in a number of contexts with the aim to improve system performance [1], [2], [3], [4], [17]. Most of these works considered centralized algorithms to maximize the overall system performance, or compared the performance of a centralized algorithm to that of a decentralized one [1], [17].

A few recent works provided a game theoretic analysis of content replication by selfish agents [8], [9], [10]. [8] considered the case with unit storage capacity and an infinite number of objects, showed the existence of equilibria and analyzed the price of anarchy for some special cases. The problem considered there is equivalent to the facility location problem, which is known to be a potential game [11]. [9] considered replication on a complete social graph and homogeneous access costs, and showed the existence of equilibria. [10] considered a version of the game where players can replicate a fraction of objects, and showed the existence of equilibria.

Our work extends recent game theoretical works on replication in several aspects. We considered the impact of the social graph on the existence of and on the convergence to equilibria, and provided sufficient conditions for convergence on arbitrary graph topologies. Finally, we considered a non-standard game-theoretic model of learning, which leverages from the social graph topology, and enables to design efficient distributed replication algorithms.

VII. CONCLUSION

In this paper we considered the problem of replication of contents by a set of selfish nodes, which replicate content to minimize their own costs. We modeled the problem as a

graphical replication game, a replication game played over a social graph. We showed that independently of the social graph there always exists an equilibrium state from which no node wants to deviate, but the social graph affects the ease of reaching such an equilibrium state. Over a complete social graph the nodes can follow a simple myopic strategy and would always reach an equilibrium in a finite number of steps, but over a non-complete social graph they could cycle arbitrarily long before reaching an equilibrium. Finally, we provided a condition under which cycles do not exist, and based on this result we proposed an efficient algorithm to reach an equilibrium state over sparse social graphs.

REFERENCES

- [1] A. Leff, J. Wolf, and P. Yu, "Replication algorithms in a remote caching architecture," *IEEE Trans. Parallel Distrib. Syst.*, vol. 4, no. 11, pp. 1185–1204, 1993.
- [2] S. Dykes and K. Robbins, "A viability analysis of cooperative proxy caching," in *Proc. of IEEE INFOCOM*, Apr. 2001.
- [3] A. Wolman, M. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. M. Levy, "On the scale and performance of cooperative web proxy caching," in *Proc. of ACM Symposium on Operating Systems Principles (SOSP)*, 1999, pp. 16–31.
- [4] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should internet service providers fear peer-assisted content distribution?" in *Proc. of Internet Measurement Conference*, 2005, pp. 63–76.
- [5] 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*, Dec. 2009.
- [6] L. Qiu, V. Padmanabhan, and G. Voelker, "On the placement of web server replicas," in *Proc. of IEEE Infocom*, Apr. 2001.
- [7] I. D. Baev and R. Rajaraman, "Approximation algorithms for data placement in arbitrary networks," in *Proc. of ACM Symposium on Discrete Algorithms (SODA)*, Jan. 2001.
- [8] B. Chun, K. Chaudhuri, H. Wee, M. Barreno, C. Papadimitriou, and J. Kubiatowicz, "Selfish caching in distributed systems: a game-theoretic approach," in *Proc. of ACM Symposium on Principles of Distributed Computing (PODC)*, July 2004.
- [9] 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.
- [10] G. Dán, "Cache-to-cache: Could ISPs cooperate to decrease peer-to-peer content distribution costs?" *IEEE Trans. Parallel Distrib. Syst.*, to appear.
- [11] D. Monderer and L. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, pp. 124–143, 1996.
- [12] F. Kuhn, "Weak graph colorings: Distributed algorithms and applications," in *ACM SPAA*, Aug. 2009.
- [13] H. Wilf, "The eigenvalues of a graph and its chromatic number," *J. London Math. Soc.*, vol. 42, pp. 330–332, 1967.
- [14] R. Varga, *Matrix Iterative Analysis*. Springer Verlag, 2000.
- [15] B. Bollobás, "The chromatic number of random graphs," *Combinatorica*, vol. 8, no. 1, pp. 49–55, 1988.
- [16] D. Welsh and M. 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.
- [17] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proc. of IEEE Infocom*, 2010, pp. 1478–1486.

APPENDIX

Proof of Lemma 1:

Player i first updates her strategy at $t = i$. Define the evicted set as $E_i(t) = \{o | r_i^o(0) = 1 \wedge r_i^o(t) = 0\}$ and the inserted set as $I_i(t) = \{o | r_i^o(0) = 0 \wedge r_i^o(t) = 1\}$. Consider now two objects $o \in E_i(t = i)$ and $o' \in I_i(t = i)$. By definition $r_i^o(0) = 1$ and $r_i^{o'}(0) = 0$, thus by the definition

of best reply and because of the edgeless social graph at $t = 0$

$$c_{io} \geq c_{io'} \quad (19)$$

Since o' is inserted in place of the evicted o , at the improvement step $t = i$ it must hold that

$$U_i^{o'}(1, r_{-i}(i)) > U_i^o(1, r_{-i}(i)) \quad (20)$$

Considering (6), we can enumerate all the possible relations between o and o' according to (19):

1	$c_{io'}$	\leq	c_{io}
2	$c_{io'}$		$\delta_i c_{io}$
3	$\delta_i c_{io'}$	$<$	c_{io}
4	$\delta_i c_{io'}$	\leq	$\delta_i c_{io}$

Case 2 represents the only possibility to satisfy (20) and it represents the case when o' is *not* i -available and o is i -available. ■

Proof of Lemma 2:

According to the structure of the utility function $U_i(r_i, r_{-i}) = \sum_{\{o|r_i^o=1\}} U_i^o(1, r_{-i})$, a best reply $r_i(t)$ can stop to be such in two situations:

- (i) The cost savings of one or more replicated objects $\{o \in \mathcal{O} | r_i^o(t) = 1\}$ decrease;
- (ii) The cost savings of one or more not replicated objects $\{o \in \mathcal{O} | r_i^o(t) = 0\}$ increase;

According to the definition of cost saving in (6), case (i) can happen only if some *not* i -available objects replicated by i become i -available. This requires that some player $j \in \mathcal{N}(i)$ starts replicating some j -available objects. Similarly, case (ii) can happen only if a neighbor $j \in \mathcal{N}(i)$ replicating an object o evicts it, making object o *not* i -available. ■

Proof of Lemma 3:

Part A: First we show that $|B_i(t-1)| \geq |B_i(t)|$. Player i can only increase the number of i -available replicated objects if she evicts at least one *not* i -available object o' from $r_i(t-1)$ and inserts an i -available object o at step t . Thus by (6) we have

$$c_{io'} < c_{io} \delta_i \quad (21)$$

Since we are in a best reply cycle, at some step $t' > t$ the strategy profile $r_i(t-1)$ must become a best reply for player i , i.e. $r_i(t') = r_i(t-1)$. This requires either $c_{io'} > c_{io}$ or $c_{io'} > c_{io} \delta_i$, and both contradict (21).

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 $A(t)$ the set of the active objects, the objects replicated by at least one player in $r(t)$, $A(t) = \{o | r_j^o(t) = 1 \text{ for some } j \in N\} \subseteq \mathcal{O}$. Similarly, we denote by $A(t)_{-i}$ the set of objects replicated by the players not including i , $A(t)_{-i} = \{o | r_j^o(t) = 1 \text{ for some } j \in N \setminus \{i\}\}$. It is easy to see that the sets $|F_i(t)|$ and $|A(t)|$ are related

$$|A(t)| = |A(t)_{-i}| + |F_i(t)| \quad (22)$$

On one hand, a best reply for which $|F_i(t-1)| = |F_i(t)|$ does not affect $|A(t)|$ since $|A(t-1)_{-i}| = |A(t)_{-i}|$. On the

other hand, a best reply for which $|F_i(t-1)| > |F_i(t)|$ decreases the size of set A

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

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 A would increase indefinitely. Hence in a cycle $|F_i(t-1)| = |F_i(t)|$, which proves the lemma. ■

Proof of Lemma 4:

Recall that according to Lemma 3 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 global knapsack problem into two similar subproblems: we can solve the knapsack problem for all the i -available objects and populate the set $B_i(t)$ and do the same with the set $F_i(t)$ using the *not* i -available ones. Suppose that k objects 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 object must be higher than the cost saving yielded by each evicted object. That is, for every $o, o' \in \mathcal{O}$, given that o was inserted and o' was evicted from $B_i(t)$

$$c_{io} \delta_i > c_{io'} \delta_i \Rightarrow w_i^o > w_i^{o'}. \quad (23)$$

Similarly, for every $o, o' \in \mathcal{O}$, given that o was inserted and o' was evicted from $F_i(t)$

$$c_{io} > c_{io'} \Rightarrow w_i^o > w_i^{o'}. \quad (24)$$

This proves the lemma. ■

Proof of Lemma 5: Assume that o is not the object with highest demand for player i . Then $\exists o' \in \mathcal{O}$ such that $c_{io'} > c_{io} \Rightarrow c_{io'} \delta_i > c_{io} \delta_i$. Object o' would yield to player i a higher cost saving than $c_{io} \delta_i$. This contradicts o being a best reply and proves (i).

We continue by proving (ii). Assume that player i inserts the i -available object o' in her next best reply. From (6) it follows that either $c_{io'} \delta_i > c_{io} \delta_i$ or $c_{io'} \delta_i > c_{io}$. None of these cases is possible because according to (i) o is the object with highest demand for player i , i.e., $c_{io'} < c_{io}$. ■

Proof of Lemma 6: We prove the lemma by showing that a move of type 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 (6), in order for the utility $U_j(r(t))$ of a player j to decrease, some neighbor $i \in \mathcal{N}(j)$ needs to start replicating some i -available object replicated by j . It follows that, from Table II, there has to be at least one best reply of type 1) or of type 3) in a cycle. A best reply of type 3) would imply $c_{io} < c_{io'} \delta_i$, but this contradicts that object o was a best reply for player i earlier, because its cost saving (6) is less than the cost saving of o' . Since there cannot be a best reply of type 3) there must be at least one best reply of type 1). This proves the lemma. ■