

On the Social Cost of Distributed Selfish Content Replication

Gerasimos G. Pollatos¹ *, Orestis A. Telelis² **, and Vassilis Zissimopoulos¹

¹ Department of Informatics and Telecommunications, University of Athens, Greece
{gpol,vassilis}@di.uoa.gr

² Department of Computer Science, University of Aarhus, Denmark
telelis@daimi.au.dk

Abstract. We study distributed content replication networks formed voluntarily by selfish autonomous users, seeking access to information objects that originate from distant servers. Each user caters to minimization of its individual access cost by replicating locally (up to constrained storage capacity) a subset of objects, and accessing the rest from the nearest possible location. We show existence of stable networks by proving existence of pure strategy Nash equilibria for a game-theoretic formulation of this situation. Social (overall) cost of stable networks is measured by the average or by the maximum access cost experienced by any user. We study socially most and least expensive stable networks by means of tight bounds on the ratios of the Price of Anarchy and Stability respectively. Although in the worst case the ratios may coincide, we identify cases where they differ significantly. We comment on simulations exhibiting occurrence of cost-efficient stable networks on average.

1 Introduction

Distributed network storage constitutes a frequently exploited modern resource for the improvement of services offered to internet users. A network node can replicate locally content (files or services) that is accessed frequently by local users, so as to lessen bandwidth consumption incurred by frequent remote access. Non-locally replicated objects can be retrieved from the nearest location where they can be found replicated (neighbors/origin servers). *Web caching, content distribution networks* and, more recently, applications over P2P networks serve this scheme. Query protocols [1], succinct summaries [2], and distributed hash tables [3] can implement search of object replicas at remote nodes.

Such systems are frequently modeled as *Distributed Replication Groups* [4]: a distant server distributes content represented as a set of information objects to a network of nodes sharing proximity, each having constrained local storage capacity. Each node has an interest to a subset of information objects in the form

* Author supported by the 03ED694 research project (PENED), co-financed by E.U.-European Social Fund (75%) and the Greek Ministry of Development-GSRT (25%).

** Center for Algorithmic Game Theory, funded by the Carlsberg Foundation, Denmark

of a non-negative demand weight per object. Access cost incurred per non-locally replicated object is defined as the distance from which the object is retrieved times the demand weight. The related *Object (or Data) Placement* problem requires placement of object replicas in nodes, subject to capacity constraints, so that total access cost summed over all objects and nodes is minimized [4, 5].

We consider networks formed *voluntarily* by autonomous users for shared exploitation of distributed resources (storage capacity), such as file-sharing networks. We develop a game-theoretic analysis of distributed replication networks, in the light of which, users choose autonomously a replication strategy (a subset of objects to replicate locally), catering to minimization of their individual access cost. Through our study of the *Distributed Selfish Replication* game we address systematically the questions: “*Does the performance of the network ever stabilize?*” and “*What is the overall performance of a stable network, in absence of a central optimizing authority?*” Game theory predicts that users will eventually end up (converge) in a Nash equilibrium [6], i.e. a set of (randomized in general) replication strategies, under which no user can decrease his/her individual (in general, expectation of) access cost by unilaterally changing his/her strategy [7].

In view of equilibria as a notion of network stabilization, we are interested in deterministic strategies (pure equilibria), because randomized strategies do not naturally fit behavior of a rational user [7, 8], as he/she may observe inefficiency during network’s lifetime and change the replication strategy. Two natural ways of measuring the social (overall) network cost include the average cost experienced by network users (SUM of individual costs), and the maximum cost over all users (MAX). The *coordination ratio* of such systems, widely known as the *Price of Anarchy* (PoA) was introduced in the seminal work of [9] as the worst-case ratio of the social cost of the *most* expensive equilibrium over the optimum centrally enforceable social cost. The PoA measures worst-case performance loss due to the absence of a centralized coordinating authority. More recently in [10] it was proposed that the worst-case ratio of the social cost of the *least* expensive equilibrium to the optimum social cost is also of interest, since some kind of coordination may be provided to the users, so that they reach the most efficient stable state. This ratio, initially termed the *optimistic* PoA [10], was established in [11] as the *Price of Stability*, (PoS) i.e. the performance loss due to the requirement for stability.

In this paper we prove existence of pure strategy Nash equilibria for the *Distributed Selfish Replication* game, (section 3) over a slightly more general network model than the one originally proposed in [4], by analyzing a converging *best response dynamics* algorithm. We prove a tight upper bound for the PoA and matching lower bounds for the PoS (section 4), for both the SUM and MAX social costs: our PoA upper bound is valid for more general network models than the one considered. However, the simple model that we study exhibits tight worst-case lower bounds for the PoS. We identify cases where the PoS can be significantly smaller than PoA, and comment on simulations that exhibit occurrence of cost-efficient equilibria on average (section 5). Our results generalize

and complement the work of [12], and answer a related open question posed in [8] with respect to considering constrained storage capacity in replication games.

Related Work. The study of voluntarily formed data replication networks by autonomous users with individual incentives was initiated in [8], where each replicated object could be payed for a fixed amount (as opposed to capacity constraints). It was shown that both PoA and PoS can be unbounded in general with the exception of special topologies. The authors considered only 0/1 demand weights in their analysis. They posed consideration of capacitated nodes along with demand weights as an open question. Let us note here that the notion of capacity should not be perceived under the narrow view of hard disk storage (which is quite affordable nowadays), but rather as a more general notion of *budget*, to be spent for maintenance and reliability purposes (particularly when replication concerns services). Existence of equilibria for the simple capacitated model of [4] was shown in [12], and for demand weights that constitute a distribution over the set of objects requested by each node. This work however does not concern social efficiency, but rather maximization of the gain of individuals when participating in the network as opposed to staying isolated. Our work is strongly related to the recent studies on *network formation* games initiated by the seminal work of [13], in an effort of understanding the efficiencies and defficiencies induced to a network created by autonomous users, that are motivated by colliding incentives. A wealth of recent results on such network creation models [11, 10, 14] have grown this direction of research into an exciting field [15] (chapter 19), that has provided useful insights with respect to the manipulation of the users' incentives towards creation of socially efficient networks.

2 Preliminaries

We consider a network involving k servers and a set N , $|N| = n$ of client-nodes (simply referred to as nodes), each node having at its disposal integral local storage capacity of $size_i$. A universe U of up-to-date unit-sized information objects originates from the k servers and each node i requests access to each object $o \in U$ by means of a demand weight $w_{io} \geq 0$. The request set of i is $R_i = \{o \in U | w_{io} > 0\}$. Node i may choose to replicate locally any at most $size_i$ -cardinality subset of R_i , and access every non-replicated object o from some node (or server) $j \neq i$ at cost $d_{ij}w_{io}$, where d_{ij} is the cost (distance) of accessing j per unit of demand.

If P_i is the subset of objects replicated by i , a *placement* over the network is denoted by $X = \{P_1, \dots, P_n\}$. Given a placement X , we use $d_i(o)$ to denote the distance from node i of the closest node (including origin servers: they can be considered as nodes with fixed placements) replicating object o . Formally: $d_i(o) = \min_{j \neq i: o \in P_j} d_{ij}$. The individual access cost of i under a placement X , is defined as:

$$c_i(X) = \sum_{o \in R_i \setminus P_i} w_{io} d_i(o) \tag{1}$$

Our study involves two versions of *social* (overall) cost: SUM and MAX, defined as $\sum_i c_i(X)$ and $\max_i c_i(X)$ respectively. We study the strategic *Distributed Selfish Replication* (DSR, and DSR(0/1) when demand weights are 0 or 1) game defined by the triple $\langle N, \{P_i\}, \{c_i\} \rangle$ in which every node i is a player with strategy space $\{P_i\}$ consisting of all $size_i$ -cardinality subsets of objects that i may replicate, and utility expressed by the access cost c_i , that i wishes to minimize. A placement X is then a strategy profile for the DSR game. In what follows we use the terms node and player interchangeably. Let $X_{-i} = \{P_1, \dots, P_{i-1}, P_{i+1}, \dots, P_n\}$ refer to the strategy profile of all players but i .

For the DSR game, it is easy to see that given a strategy profile X_{-i} , player i can determine optimally in polynomial time its *best response* P_i to the other players' strategies, simply by solving a special 0/1 *Knapsack* problem, that amounts to replicating $size_i$ objects with the greatest value $w_{io}d_i(o)$.

Network Model We consider a slightly more general network model than the one introduced by Leff, Wolf, and Yu in [4] and studied in [12], involving k origin servers, instead of 1. The distance of every node i from server l is d_l for $l = 0 \dots k - 1$, while two nodes i and j are at distance $d_{ij} = d_k$. We assume that distances form an ultra-metric such that $d_k < d_{k-1} < \dots < d_0$. We refer to this network model as **LWY**(k). The minimum and maximum distances appearing in the network are also referred to with $d_{\min} = d_k$ and $d_{\max} = d_0$.

3 Existence of Pure Strategy Nash Equilibria

We introduce a simple polynomial-time algorithm, that finds a feasible placement constituting a pure strategy Nash equilibrium for the DSR game on a **LWY**(k) networks. The algorithm constitutes a converging *iterative best response* dynamics that iterates for $l = 0 \dots k$ over the different distance values (in non-increasing order) appearing in **LWY**(k): for each distance value d_l , all distances at most equal to d_l are redefined to exactly d_l and a best response placement is computed for each node $i = 1 \dots n$ in turn. It is important to note that during computation of best response of a node i a tie for the values of two objects $w_{io}d_i(o)$ and $w_{io'}d_i(o')$ is resolved on the basis of index (for otherwise we cannot guarantee convergence to equilibrium). The order by which nodes compute their best response is fixed in all iterations. We refer to this algorithm (algorithm 1) as DSR-EQ. We show that DSR-EQ finds an equilibrium placement. In our analysis we use the iteration index l as superscript to the various quantities involved.

Proposition 1. *Consider the time right after computation of placements P_i^{l-1} and P_i^l for a fixed node i in iterations $l - 1$ and l of algorithm DSR-EQ. Then for every pair of objects o, o' with $o \in P_i^{l-1} \setminus P_i^l$ and $o' \in P_i^l \setminus P_i^{l-1}$ it is $d_i^l(o) = d_l$ and $d_i^l(o') > d_l$.*

The proof is by induction on indices $\langle i, l \rangle$ through the following claims:

Claim. (Basic Step) Proposition 1 holds for $l = 1$ (second iteration) and $i = 1$.

Algorithm 1: DSR-EQ

```
for all  $i, j$  set  $d_{ij}^{-1} \leftarrow d_{ij}$ ;  
for  $l=0 \dots k$  do  
  for all  $i, j$  set  $d_{ij}^l \leftarrow d_{ij}^{l-1}$ ;  
  for all  $i, j$ :  $d_{ij} \leq d_l$  set  $d_{ij}^l \leftarrow d_l$ ;  
  for  $i=1 \dots n$  do  
    Compute best response  $P_i^l$  of  $i$  with respect to  $[d_{ij}^l]$   
  end  
end
```

Proof. Right after computation of P_1^0 and P_1^1 the following hold with respect to objects o and o' :

$$d_1^0(o')w_{1o'} \leq d_1^0(o)w_{1o}, \quad \text{and} \quad d_1^1(o')w_{1o'} \geq d_1^1(o)w_{1o} \quad (2)$$

Equality in both cases cannot hold, because of the index rule for breaking ties. From the first inequality it must be that $w_{1o'} \leq w_{1o}$ because $d_1^0(o') = d_1^0(o) = d_0$. Furthermore, in iteration $l = 1$, by redefinition of distances less than d_0 in the original input to d_1 , it must be either that $d_1^1(o') = d_1^1(o) = d_0$ (both objects are found only at server distance d_0), or that $d_1^1(o') = d_1^1(o) = d_1$ (both objects are replicated somewhere at distance d_1), or that $d_1^1(o') \neq d_1^1(o)$. The first two imply $w_{1o'} \geq w_{1o}$ which contradicts $w_{1o'} \leq w_{1o}$, even in the case of equality (because of the tie-breaking index rule). For the third case, if $d_1^1(o') = d_1$ and $d_1^1(o) = d_0$, we obtain that $d_1^1(o')w_{1o'} < d_1^1(o)w_{1o}$, a contradiction to the second relation of (2). Thus it must be that $d_1^1(o') = d_0 > d_1$ and $d_1^1(o) = d_1$. \square

Claim. (Inductive Step) Proposition 1 holds for every iteration l .

Proof. Assume that the statement is true for all nodes $\{i+1, \dots, n, 1, \dots, i-1\}$ computing their best response right after computation of P_i^{l-1} and right before computation of P_i^l . For $o \in P_i^{l-1} \setminus P_i^l$ and $o' \in P_i^l \setminus P_i^{l-1}$ we can write:

$$d_i^{l-1}(o')w_{io'} \leq d_i^{l-1}(o)w_{io}, \quad \text{and} \quad d_i^l(o')w_{io'} \geq d_i^l(o)w_{io} \quad (3)$$

Equality in both cases cannot occur as before. Assume that $d_i^l(o) > d_l$. Then $d_i^l(o) \geq d_{l-1}$ and, if o was replicated at some node in iteration $l-1$ (i.e. $d_i^{l-1}(o) = d_{l-1}$), it must have remained so, by hypothesis. But then it should have been $d_i^l(o) = d_l$. Thus o was found only at some server in iteration $l-1$. Then $d_i^l(o) = d_i^{l-1}(o)$. For o' however, we have that $d_i^l(o') \leq d_i^{l-1}(o')$, because of the assumed behavior of the other nodes computing best response after P_i^{l-1} and before P_i^l : either o' reduces its distance from i due to redefinition of distances to d_l (in case some node has it replicated) or not. Then it is:

$$d_i^l(o')w_{io'} \leq d_i^{l-1}(o')w_{io'} \leq d_i^{l-1}(o)w_{io} = d_i^l(o)$$

which is a contradiction to (3). Thus it must be $d_i^l(o) = d_l$.

Now assume that $d_i^l(o') = d_l$. Since $d_i^l(o) = d_l$ also, we have that $w_{io'} \geq w_{io}$. Furthermore, in iteration $l-1$ it must have been $d_i^{l-1}(o') \geq d_{l-1}$ and $d_i^{l-1}(o) = d_{l-1}$. Then we obtain $d_i^{l-1}(o')w_{io'} \geq d_i^l(o')w_{io'} = d_l w_{io'} > d_l w_{io} = d_i^l(o)w_{io} = d_i^{l-1}(o)w_{io}$, which is a contradiction to (3). Thus it must be $d_i^l(o') > d_l$. \square

Using proposition 1 it is possible to show that:

Lemma 1. *In the end of iteration l of algorithm 1, $\{P_i^l | i = 1 \dots n\}$ is a pure strategy Nash equilibrium with respect to the current distance matrix $[d_{ij}^l]$.*

Proof. By proposition 1 we deduce that, in each iteration, a node's best response is not invalidated by other nodes' best responses, as follows. In the l -th iteration, for every node $j \neq i$ and an object $o \in P_j^{l-1} \setminus P_j^l$ the object remains at distance d_l from i (which was its distance when i was computing P_i^l). For an object $o' \in P_j^l \setminus P_j^{l-1}$, o' was at distance $> d_l$ and, therefore, not replicated in i either. Replication in j only decreases the value of o' for i (because its distance is reduced). \square

Theorem 1. *Pure strategy Nash equilibria exist for the DSR game on **LWY**(k) networks.*

Proof. By lemma 1 the placement produced in the last iteration of algorithm 1 is a pure strategy Nash equilibrium with respect to a distance matrix identical to the input one. \square

4 Worst-Case Equilibria

We study efficiency of pure equilibria for the DSR game, by studying the worst-case ratios of the PoA and PoS. We assume that all players (nodes) have a *marginal* interest for joining the network, expressed by $|R_i| \geq size_i$. If this does not hold, then equilibria may be unboundedly inefficient, because of unused storage space in certain nodes. In summary we show that the PoA can be at most $\frac{d_{\max}}{d_{\min}}$ for both the **SUM** and the **MAX** social cost functions, and the PoS can be at least as much in the worst-case. We compare a worst-case socially most expensive equilibrium placement X towards the socially optimum placement X^* , and use $d_i(\cdot)$ and $d_i^*(\cdot)$ for distances of objects in each placement. The analysis is performed for the individual cost of a single fixed node i , and therefore yields results with respect to the **SUM** and the **MAX** social cost functions.

Lemma 2. *If X and X^* are equilibrium and optimum placements respectively, for each node i we have:*

$$c_i(X) = \sum_{o \in R_i \setminus (P_i \cup P_i^*)} w_{io} d_i(o) + \sum_{o \in P_i^* \setminus P_i} w_{io} d_i(o) \quad (4)$$

$$c_i(X^*) = \sum_{o \in R_i \setminus (P_i^* \cup P_i)} w_{io} d_i^*(o) + \sum_{o \in P_i \setminus P_i^*} w_{io} d_i^*(o) \quad (5)$$

Proof. Taking expression (1) for c_i at X and X^* and noticing that:

$$R_i \setminus P_i = (R_i \setminus (P_i \cup P_i^*)) \cup (P_i^* \setminus P_i) \text{ and } (R_i \setminus (P_i \cup P_i^*)) \cap (P_i^* \setminus P_i) = \emptyset$$

$$R_i \setminus P_i^* = (R_i \setminus (P_i^* \cup P_i)) \cup (P_i \setminus P_i^*) \text{ and } (R_i \setminus (P_i^* \cup P_i)) \cap (P_i \setminus P_i^*) = \emptyset$$

yields the result. \square

Lemma 3. *If X and X^* are equilibrium and optimum placements respectively, for every player i we have that $|P_i \setminus P_i^*| = |P_i^* \setminus P_i|$ and for every pair of objects $o \in P_i \setminus P_i^*$, $o' \in P_i^* \setminus P_i$ it is $w_{io}d_i(o) \geq w_{io'}d_i(o')$.*

Proof. Since for every player i $size_i < |R_i|$, we obtain $|P_i \setminus P_i^*| = |P_i^* \setminus P_i|$. For every pair of objects $o \in P_i \setminus P_i^*$, $o' \in P_i^* \setminus P_i$ it must be $w_{io}d_i(o) > w_{io'}d_i(o')$, for otherwise, player i would have an incentive to unilaterally deviate in placement X , by substituting o for o' , thus decreasing its individual access cost. \square

Theorem 2. *The price of anarchy for the DSR game is upper bounded by $\frac{d_{max}}{d_{min}}$, with respect to either of the SUM or the MAX social cost functions.*

Proof. Starting from expression (5) and using lemma 3 we have:

$$\begin{aligned} c_i(X^*) &= \sum_{o \in R_i \setminus (P_i^* \cup P_i)} w_{io}d_i^*(o) + \sum_{o \in P_i \setminus P_i^*} w_{io}d_i^*(o) \\ &\geq \sum_{o \in R_i \setminus (P_i^* \cup P_i)} w_{io}d_i^*(o) \frac{d_i(o)}{d_i(o)} + \sum_{o' \in P_i^* \setminus P_i} w_{io'} \frac{d_i(o')}{d_i(o)} d_i^*(o) \\ &\geq \sum_{o \in R_i \setminus (P_i^* \cup P_i)} w_{io}d_i(o) \frac{d_{min}}{d_{max}} + \sum_{o' \in P_i^* \setminus P_i} w_{io'}d_i(o') \frac{d_{min}}{d_{max}} \end{aligned}$$

By expression (4) $c_i(X^*) \geq \frac{d_{min}}{d_{max}} c_i(X)$. Summing over all i yields the result for the SUM. For the MAX notice that for some j : $\max_i c_i(X) = c_j(X) \leq \frac{d_{max}}{d_{min}} c_j(X^*) \leq \frac{d_{max}}{d_{min}} \max_i c_i(X^*)$. \square

Note that we did not make use of any particular network topology. The bound is valid on any network topology, provided an equilibrium exists. Interestingly, matching lower bounds for the PoS can be shown on a **LWY**(1) network:

Proposition 2. *The Price of Stability for the DSR game has a lower bound arbitrarily close to d_{max}/d_{min} in the worst-case, with respect to SUM and MAX social cost functions, even for 1 server and 0/1 demand weights.*

Proof. Take d_{min} inter-node distance and d_{max} the distance of the single server.

For SUM: we show that for every fixed integer I the PoS is lower bounded by $\frac{I}{I+1} \frac{d_{max}}{d_{min}}$ asymptotically with n . Take $size_i = 1$ and set: $R_i = \{a\}$ for $i =$

$1 \dots n - I$ and $R_i = \{b_j | j = 1 \dots n - 1\}$, for $i = n - I + 1 \dots n$. In every equilibrium of this instance the I last nodes replicate collectively at most I objects from their (common) request set, while the rest $n - I$ replicate object a . The social cost is $I(I - 1)d_{\min} + I(n - 1 - I)d_{\max}$. In the socially optimum placement we let nodes $i = 1 \dots n - I - 1$ replicate $n - I - 1$ objects that are not replicated in the I last nodes. But, we save a node for replicating a . The social cost becomes $I(n - 2)d_{\min} + (n - 1 - I)d_{\min} = [(I + 1)n - 3I - 1]d_{\min}$. Since I is fixed, the PoS becomes $\frac{I}{I+1} \frac{d_{\max}}{d_{\min}}$ as n grows.

For **MAX**: all nodes have $size_i = 1$ and let $R_1 = \{a, b\}$, whereas for $i > 1$ set $R_i = \{c\}$. Take $n \geq 3$ nodes. In equilibrium node 1 pays d_{\max} for either a or b . In the social optimum some $i > 1$ replicates one of a, b that is not replicated by node 1, and pays d_{\min} for c , while 1 also pays d_{\min} . \square

5 Less Expensive Equilibria

Proposition 2 states that the least expensive and the most expensive equilibria can be of the same cost in the worst case. In this section we investigate existence of less expensive equilibria than the ones indicated by the worst-case PoA and PoS. In particular, we elaborate on conditions under which the PoS can be at most $\frac{d_{\max}}{2d_{\min}}$ for the DSR(0/1) game, and present some indicative experimental results on random instances of the DSR game, showing that social efficiency of equilibria does not degrade much as the ratio d_{\max}/d_{\min} grows. For the rest of this section's theoretical analysis we turn our attention to the DSR(0/1) game and restrict our discussion to the **SUM** social cost function.

In proving proposition 2 we used “over-demanding” nodes, each requesting $O(n)$ times more objects, than the storage capacity they offer to the group. In what follows we will assume *modestly demanding* nodes. We define such nodes formally through the *demand ratio* $q = \max_i \frac{|R_i| - size_i}{size_i}$ as being $q = O(1)$. We will also use the assumption that the absolute difference of any pair of distinct distance values in the network is at least d_{\min} . This is e.g. the case when the network is modeled as a graph with inter-node distance being one hop, while the nearest server being at least 2 hops away.

We will compare the least expensive equilibrium $X = \{P_i\}$ against the socially optimum placement $X^* = \{P_i^*\}$, so as to obtain the maximum possible difference $c(X) - c(X^*)$. We consider objects $a^* \in P_i^* \setminus P_i$, that we call *insertions* and study an imaginary procedure that transforms X into X^* , by performing such insertions. An insertion is termed *significant* if it contributes to $c(X) - c(X^*)$.

Lemma 4. *If r significant insertions occur in comparing a least expensive equilibrium X to a socially optimum placement X^* , and at most n_0 nodes benefit by any insertion, then: $c(X) - c(X^*) \leq rn_0(d_{\max} - d_{\min}) - rd_{\min}$.*

Proof. Take $a^* \in P_i^* \setminus P_i$ for some node i . There is an object $a \in P_i \setminus P_i^*$ evicted from i due to insertion of a^* . We use $d_i(\cdot)$ to denote the current distance of an object from i . Our target is to show that every significant insertion contributes a

social cost increase at least d_{\min} while potentially offering at most $(d_{\max} - d_{\min})$ cost decrease to some nodes. At first suppose that $w_{ia^*} = 0$: i suffers an access cost increase of at least d_{\min} , because a has to be retrieved remotely at minimum distance d_{\min} . Every node $j \neq i$ with $w_{ja^*} = 1$ benefits at most $(d_{\max} - d_{\min})$ though. This is a significant insertion. Now consider the case $w_{ia^*} = 1$. Then $d_i(a) \geq d_i(a^*)$, because X is equilibrium. We examine the following cases:

1. $d_i(a) = d_i(a^*) = d_{\min}$: then i does not benefit or loose, and the insertion is indifferent to every other node $j \neq i$ with $w_{ja^*} = 1$ (they also do not benefit or loose, because a^* is already replicated somewhere in the network). Such an insertion is not significant.
2. $d_i(a) > d_i(a^*) = d_{\min}$: the cost of i increases by at least $2d_{\min}$ (due to eviction of a and $d_i(a) - d_i(a^*) \geq d_{\min}$) and decreases its cost by at most d_{\min} (due to insertion of a^*). Such an insertion is potentially significant.
3. $d_i(a) = d_i(a^*) > d_{\min}$: i does not benefit or loose, but any node $j \neq i$ with $w_{ja} = 1$ and $w_{ja^*} = 0$ loses at least $d_i(a^*) - d_{\min} > d_{\min}$ (as long as a was stored in i it was retrieved by j from distance d_{\min}). If no such node exists, then we have a new equilibrium, potentially less expensive than X , a contradiction. Such an insertion can be significant.
4. $d_i(a) > d_i(a^*) > d_{\min}$: eviction of a from i is caused at distance $d_i(a)$ which is at least by d_{\min} greater than $d_i(a^*)$. Thus i loses at least d_{\min} in this case. This insertion can be significant.

Notice that the cost increase effect of cases (2), (3) and (4) may be compensated by a re-insertion of the evicted object a in some node $j \neq i$. By the previous arguments, such a re-insertion causes eviction of some object a' from j and therefore, induces a new cost increase of at least d_{\min} , or a new equilibrium contradiction. Such re-insertions are insignificant, as they may form chains of “movements” of objects around the network’s nodes, always carrying a cost increase, but no decrease. If n_0 nodes benefit from any of the significant insertions, then assuming that all of them benefit from all the significant insertions only enlarges the difference $c(X) - c(X^*)$, and the result follows. \square

Lemma 5. *The cost of the socially least expensive equilibrium placement X for the DSR(0/1) game on a $LWY(k)$ network in the worst case is at least:*

$$c(X) \geq n_0(n_0 - 1)Sd_{\min} + n_0rd_{\max} \quad (6)$$

where r is the number of significant insertions, S is the minimum storage capacity, and n_0 is the number of nodes that benefit from all insertions.

Proof. For every node j that benefits from an insertion a^* it is $w_{ja^*} = 1$, and for maximum benefit $(d_{\max} - d_{\min})$, this object is payed d_{\max} in the best equilibrium. In the worst case all n_0 nodes benefit from all r significant insertions, as argued in lemma 4 (this assumption enlarges the difference $c(X) - c(X^*)$). The second term of inequality (6) is thus justified.

For the first term, we notice that since all n_0 nodes benefit from an insertion a^* , then in the least expensive equilibrium X every object stored in these nodes

is of interest to at least n_0 nodes of the network. Otherwise, we show that there may be an equilibrium of less cost: take for example an instance where every one of the n_0 nodes is either interested to replicas stored in the rest $n_0 - 1$ of them or not, and no other nodes are interested in replicas stored in these n_0 nodes. If there is a node j not interested to a replica a stored in some node i , then substitution of a for a^* in i yields a less expensive equilibrium. We assume that all n_0 nodes benefiting from the r insertions are also interested in objects replicated in them in equilibrium. This yields a lower bound of $n_0(n_0 - 1)Sd_{\min}$ cost. \square

Now we can prove the following:

Theorem 3. *The Price of Stability of the DSR(0/1) game on **LWY**(k) networks with modestly demanding nodes and distance differences at least d_{\min} is upper bounded by $\frac{d_{\max}}{2d_{\min}}$.*

Proof. By lemma 4:

$$PoS \leq \frac{c(X)}{c(X) + rd_{\min} - r(d_{\max} - d_{\min})}$$

This bound is a monotonically decreasing function of $c(X)$. Therefore we use for $c(X)$ the lower bound given by (6) for worst-case PoS , and obtain:

$$PoS \leq \frac{n_0(n_0 - 1)Sd_{\min} + n_0rd_{\max}}{(r + rn_0 + n_0(n_0 - 1)S)d_{\min}}$$

By the analysis of lemma 5, the number of insertions occurring is upper bounded as $r \leq (q - (n_0 - 1))S$. Taking the maximum value $r = (q - (n_0 - 1))S$ increases the ratio for every value of n_0 :

$$PoS \leq \frac{n_0(n_0 - 1)Sd_{\min} + n_0(q - (n_0 - 1))Sd_{\max}}{[(q - (n_0 - 1))S + n_0(q - (n_0 - 1))S + n_0(n_0 - 1)S]d_{\min}} \Rightarrow$$

$$PoS \leq \frac{n_0 - 1}{q + \frac{q - (n_0 - 1)}{n_0}} + \frac{q - (n_0 - 1)}{q + \frac{q - (n_0 - 1)}{n_0}} \cdot \frac{d_{\max}}{d_{\min}}$$

Now notice that, since every node of the n_0 accesses replicas stored in all other $n_0 - 1$ nodes, it must be $n_0 - 1 \leq q$. Thus the first term can be at most 1, while for $q = O(1)$, the second term is maximized to $\frac{d_{\max}}{2d_{\min}}$ for $n_0 = 1$. \square

Simulation on Random Instances We performed simulations on random instances of the general DSR game, that provide evidence of existence of efficient equilibria on average. Fig. 1 depicts the evolution of empirically extracted values for PoS, PoA, and an average value of coordination ratio, as d_{\max}/d_{\min} increases. We used an **LWY**(10) network of 64 nodes, with $size_i = 5$, and $|R_i| = 10$. For each node i the request set R_i was constructed by uniform random sampling




Fig. 1: Evolution of PoS, PoA, and average coordination ratio (AVG) with d_{\max}/d_{\min} .

of 10 objects out of a universe U of 400 objects. Demand weights were taken randomly in $(0, 1]$ for each node. We took $d_{\min} = 1$ and distributed 10 servers, each hosting 40 objects, uniformly within the range $[d_{\min}, d_{\max}]$ for each different value of d_{\max} .

Social optimum for SUM and MAX was obtained by solution of the *Object Placement Problem* ILP formulation described in [5], using GLPK [16]. Each point of the diagrams was extracted using the maximum (for PoA), minimum (for PoS) and average (for AVG) social cost obtained over 1000 executions of algorithm DSR-EQ with randomly chosen node and object orderings per execution.

Two main observations can be extracted by these experiments. First, that the rate of growth of all coordination ratios degrades significantly as d_{\max}/d_{\min} increases, thus giving the impression of a stabilization of the efficiency of the system. Second, that the range of different stable social cost values is quite narrow. These observations motivate a probabilistic analysis of the efficiency of equilibria, that constitutes an aspect of future work.

6 Conclusions and Open Questions

We studied the effects of selfish behavior in a distributed replication group with restricted storage capacity per node. We proved existence of pure strategy Nash equilibria for the related strategic game for networks with multiple origin servers, thus generalizing the result of [12]. Although worst-case maximum and minimum social cost of equilibria were shown to be $\frac{d_{\max}}{d_{\min}}$ times the optimum social cost in the worst case, for 0/1 demand weights and certain topologies a better upper bound of $\frac{d_{\max}}{2d_{\min}}$ was shown for the Price of Stability. Simulations on random instances have shown that even the Price of Anarchy may grow much less rapidly than the ratio d_{\max}/d_{\min} .

We were not able to find a single example without any pure equilibria when arbitrary-sized objects are involved, or disprove existence of pure equilibria for more general topologies. We consider investigation of these situations as an aspect of future work, along with enrichment of the game model with other realistic features, such as bounded node degrees.

References

1. Wessels, D., Claffy, K.: ICP and the Squid web cache. *IEEE Journal on Selected Areas in Communications* **16**(3) (1998)
2. Fan, L., Cao, P., Almeida, J., Broder, A.Z.: Summary cache: a scalable wide-area web cache sharing protocol. *IEEE/ACM Transactions on Networking* **8**(3) (2000) 281–293
3. Stoica, I., Morris, R., Liben-Nowell, D., Karger, D.R., Kaashoek, M.F., Dabek, F., Balakrishnan, H.: Chord: A scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Transactions on Networking* **11**(1) (2003) 17–32
4. Leff, A., Wolf, J., Yu, P.S.: Replication Algorithms in a Remote Caching Architecture. *IEEE Transactions on Parallel and Distributed Systems* **4**(11) (1993) 1185–1204
5. Baev, I., Rajaraman, R.: Approximation algorithms for data placement in arbitrary networks. In: *Proceedings of the ACM-SIAM Annual Symposium on Discrete Algorithms (SODA)*. (2001) 661–670
6. Nash, J.F.: Non-Cooperative games. *Annals of Mathematics* **54** (1951) 286–295
7. Osborne, M.J., Rubinstein, A.: *A course in game theory*. MIT Press (1994)
8. Chun, B., Chaudhuri, K., Wee, H., Barreno, M., Papadimitriou, C.H., Kubiawicz, J.: Selfish caching in distributed systems: a game-theoretic analysis. In: *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*. (2004)
9. Koutsoupias, E., Papadimitriou, C.H.: Worst-case Equilibria. In: *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS)*. (1999) 404–413
10. Anshelevich, E., Dasgupta, A., Tardos, E., Wexler, T.: Near-optimal network design with selfish agents. In: *Proceedings of the ACM Annual Symposium on Theory of Computing (STOC)*. (2003) 511–520
11. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, E., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: *Proceedings of IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. (2004) 295–304
12. Laoutaris, N., Telelis, O.A., Zissimopoulos, V., Stavrakakis, I.: Distributed Selfish Replication. *IEEE Transactions on Parallel and Distributed Systems* **17**(12) (2006) 1401–1413
13. Fabrikant, A., Luthra, A., Maneva, E.N., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: *Proceedings of the ACM Symposium on Principle of Distributed Computing (PODC)*. (2003) 347–351
14. Chen, H., Roughgarden, T., Valiant, G.: Designing networks with good equilibria. In: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)* (to appear). (2008)
15. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V., (eds.): *Algorithmic Game Theory*. Cambridge University Press (2007)
16. GNU Linear Programming Kit. (<http://www.gnu.org/software/glpk/>)