Institut Eurécom
Department of Corporate Communications
2229, route des Crètes
B.P. 193
06904 Sophia-Antipolis
FRANCE

Research Report RR-08-213

A new approach to efficient peer-to-peer backup systems

Laszlo Toka and Pietro Michiardi

Tel : (+33) 4 93 00 81 45
Fax : (+33) 4 93 00 26 27
Contact at Eurecom : {Pietro.Michiardi}@eurecom.fr

[1]Institut Eurécom's research is partially supported by its industrial members: Bouygues Téécom, Orange, Hitachi Europe, SFR, Sharp, ST Microelectronics, Swisscom, Thales, BMW, CISCO.

Abstract

We present a realistic model of next generation peer-to-peer (P2P) backup/storage systems where users have the ability, besides making the usual decisions, to select the peers they want to cooperate with. Under the circumstances we can anticipate the final state where the system ends up, and interesting conclusions can be drawn: selfish users will form coalitions based on their qualities. The assumed extension, compared to the existing P2P systems, introduces direct incentives for users to contribute more to the system, moreover it increases user fairness, not to mention the fact that it results in easier peer selection and data management techniques to be implemented.

Keywords: peer-to-peer system, backup, storage, peer selection, game theory, user model, incentive, cooperative, coalitional game

1 Introduction

In a P2P backup/storage system users are expected to cooperate i.e. share their private resources (storage space on hard disks, bandwidth) with other participants to make the concept work. This cooperation needs to be a long-term one and dedicated to well-defined partners (not like in file sharing systems, where e.g. tit-for-tat cooperation operates and holds just for the transfer time, and the same resources are shared towards multiple partners). Especially in a backup system, users like to rely on the service on the long run, and the backuped data, exchanged between cooperating peer pairs, belongs only to its owner. In such an environment the need of incentives is evident.

Several works tackle economic modeling of P2P systems and the well-known user incentive problem has seen many approaches of possible solutions from different authors on every field of distributed systems, networks. However, in our opinion, an important aspect has not received enough attention so far: when game theoretic model is used to describe user strategy and utility, the possible actions towards various players should be considered differently, i.e. if information is at hand about players, then it needs to be taken into account. The information on user characteristics is not available in every case of distributed system or it is too limited to efficiently lie on (e.g. packet forwarding problem in ad-hoc wireless networks), but in a P2P backup/storage it is.

Thus, in this paper, we consider that the system enables users to have an effect on peer selection too. The technical implementation of this model is a double-overlay structure: one overlay is a DHT containing meta-data about peers to support the required reputation system and partner selection. The second overlay is the effective cooperation mesh of peers in coalitions. We also look into details of market solutions where every user can find the best buddy to cooperate with. We investigate a stock exchange-like approach where different quality unit storage spaces are offered and asked for various values received in also storage spaces. Necessary rules for simple, fast trades and robust system are then needed. We do not tackle the problem of optimal backup data amount or other user defined strategy profiles in details, but focus on peer selection which has an outstanding importance in user rationality, selfishness, system incentives and social welfare. We derive the expected outcome of such a system, and show that it is advantageous to integrate peer reputation into user strategy more substantially than it is done in previous works.

2 Related work

The study of incentives in P2P systems has a long history. We point the reader to references [10], [11] on economic modeling of backup/storage systems targeting particularly incentive issues and other existing literature cited in them.

2.1 Past contribution

In [6] the selfish user behavior is described by game theoretic strategies and related payoff functions. This latter is split in two parts: utility function which shows a given user's willingness to participate in the backup service in function of amount of data, and cost function to depict faced costs of backuping the partners' data at the user own resources. The built-up model handles user's online spent time, bandwidth issues, and introduces incentives for users to offer storage space to the system and to reduce churn. Two types of incentive management is analyzed and compared to reach the highest social welfare on every user set. In the model new data load at every peer (re)arrival is supposed - maybe too drastic.

2.2 Relevant works

Wuala [13], soon to be commercialized P2P storage/backup system has focused on symmetric exchanges between users, in terms of backuped data, based on a simple rule: each user has the right to backup the amount of data in the system that storage space she offers at her own resource reduced to her real availability on the network, i.e. offering 10 Gb but online only 50% of time, she can eat up only 5 Gb in the system (redundancy, replication excluded). Users must hold online availability above a defined level, backuped data existence and integrity is periodically checked by the data owner, and important peer parameters (offered/used storage space, availability, connection bandwidth, malicious behavior is observed, spread and stored in a distributed way). For the initial time, backuped data of each user is stored on central servers as well.

In [3] the authors present the performance evaluation of different strategies on peer selection from the churn-caused cost's point of view. In the sense that they argue on randomization's positive effects in the outcome of a stochastic model, their churn-decreasing efforts through the right peer selection fundamentally differs from our discourse (i.e. they want to decrease the effect of churn, in the meanwhile we target the churn itself), but already show that appropriate node selection is important in a distributed system.

2.3 Data management supporting tools

Data management is a crucial issue in a backup/storage system. A vast collection of literature is written on this domain, tackling data redundancy ([12], [9]), churn policy, i.e. on grace period at the end of which re-arranging the stored data happens after a peer departure, ([11], [1], [5], [2]) and reputation systems [4], to collect and spread meta-data.

These means are not studied in our paper, we just assume them to support the overall economic framework of the P2P backup/storage system model, i.e. we suppose that integrity check protocols exist to maintain "connection" between peers, reputation systems are in place to spread the information on system participants, DHT is built to store meta-data, etc. Observing a node's behavior by other peers and accessibility checks by data's owner on storing peers have outstanding importance.

3 Model

In this section we give the the base of our model that is supposed to describe users realistically in the economic point of view in a P2P backup/storage system. First we give the necessary assumptions, then we describe users, their actions and corresponding utility. At the end we outline the way users interact and the results this will eventually lead to.

3.1 Assumptions

We tackle the modeling and evaluation of a realistic P2P backup/storage system, where users store their own data at other members and vice versa in a symmetric way taking into account some rules. We import some concepts from existing models [6], i.e. exchange of abstract storage space units symmetrically between peers, linear user payoff function with disjoint valuation and cost terms, user strategy on backup amount, etc. User modeling by statistical parameters, including needs (in terms of data to backup), resources (exchangeable storage space, uplink/downlink bandwidth capacity, uptime) and the strategy set at hand with the expected utility, are all well described in past contributions. Just a few weaknesses need attention when building a more realistic model on backup/storage P2P systems: in this work we do not assume new data load at every peer re-arrival, we include bandwidth of partners in the user utility function of backuping, we manage grace period for departing nodes through advertised reputation, etc.

The main contribution of this work lies on the intention to take peer selection into account. In a fully distributed system, the peers are assumed to manage their own existence and actions, such as the selection of available peers that they are willing to exchange data for backup/storage with. Moreover we assume that peers act and decide selfishly, so a game theoretical model is suited to the analysis and performance evaluation of such a system.

3.2 Users & strategy

We define the realistic user model based on the above mentioned assumptions. A user profile parameter α, that suits with the given peer's availability to the system (probability to be found online), data possession behavior (results of data integrity checks) and accessibility (link bandwidth to reach the given peer), is introduced. This profile parameter is established for each peer by her partners in the system, so it can be produced and viewed differently from variant angles, i.e. a nearby peer may experience faster communication with a given peer, whereas a distant one might reduce the given peer's observed bandwidth parameter based on her observations. This "distributed information" is then handled by well-known means, like Bayesian techniques (e.g. [7]). The user set of αs affects each participant's payoff function when operating in the P2P backup/storage system. On one hand, backuping's utility depends on what type of partners store a given peer's data, on the other hand the occurring costs of offering storage space at own resources are partly due to the peers' profile set when e.g. considering churn generated heavy traffic (even if no new load is shipped, checks are desirable on re-arriving peers).

3.2.1 Profile parameter

α plays a crucial role in inter-peer relationships and so does in our model. Each and every peer wants to deal with the best peers ready for cooperation. But how to define peer i's α_i profile exactly? First of all, malicious behavior has to be dealt with by giving it a big importance in the parameter. In the case a given peer is observed to be dropping an other peer's stored data unilaterally, the malicious peer's reputation, provided by the data owner, has to fall drastically. Same action is needed when a peer is being unreachable for an unusually long period. Second most important factor is the online behavior: the probability that the given peer can be found online. This parameter is global in the sense that hopefully every observer will establish nearly the same value for a given peer. Thirdly, the connection bandwidth should be also one ingredient, but this latter might be relative from other peers' perspectives.

The previously described sequence of α_i's factors is important: generally a peer's reliability (in a backup/storage service) depends more on its tendency to behave maliciously than its availability, and the same relation stands for availability and the connection rate (as hopefully transfers of large data amount will occur rarely, especially in a backup system). Just for an initial hint, we suggest that $0 < \alpha_i \leq p_i$, where p_i is peer i's online availability, i.e. probability measure, hence $\alpha_i \in [0, 1]$.

We do assume that because of the public parameter α, the reputation system in place and the long-term characteristic of backup/storage service, the permanent disconnections happen far more rarely than temporal ones (e.g. transient node failures). That is why we consider definite peer departures (after *unusual* offline time) as malicious behavior.

3.2.2 Strategy

Strategy selection results in the way a given peer is willing to cooperate in the P2P system. Our game theoretical model assumes that users are selfish. Each one of them determines the capacity she wants to exchange (c_i) based on the observed system profile and her individual payoff function.

α_i might be also part of one's strategy as there are components she has effects on (e.g. could buy faster connection) in order to vary her profile.

3.3 User & utility: payoff function

Definition 1 *The assumed additive payoff function (P_i) holds several disjoint elements and can be described by the following form: $P_i = U_i - D_i - (O_i + T_i) - E_i$, where*

- U_i *stands for user i's* utility *regarding the backup service,*
- D_i *means the* service degradation *due to non-optimal collaborating peers,*
- O_i opportunity cost *of giving away private resources,*
- T_i *represents* transfer cost *related to the service,*
- E_i *describes the* effort cost *when peer i tries to ameliorate her profile α_i.*

As discussed in Subsection 3.2, user i can influence her payoff by moving on two independent axis (apart from the peers' quality selected for cooperation). The first dimension stands for the supplied capacity c_i of user i and the second one for her behavior described by α_i. O_i is due to the first element of user strategy (c_i) but every other term contains α_i: U_i depends on the demanded capacity $\alpha_i c_i$ (see exchange policy in 3.4), and as α_i is increased by peer i, the service degradation and the transfer cost decrease so the user's payoff grows, but the effort cost goes in the opposite direction. The user's best strategy is to find and apply the balance where these contrary effects result the maximum payoff. More details on the payoff's elements are presented hereby and deeper analysis is shown in Subsection 4.1.

3.3.1 Service degradation

The issue of backup service's quality should appear in the valuation function of a user's payoff. Authors in [11] model the utility as a function of the backuped data amount and the uploading/retrieving process is supposed to be optimal. However depending on the cooperating partner profiles, these steps could be less appropriate if a specific peer does not gather right peers to cooperate with (e.g. bad connectivity of a peer makes much longer upload time when backuping on that specific target and not on a better connected one).

Thus this term reflects the service's degradation compared to an optimal situation, mainly due to "light" peers, so it could be viewed as a non-monetary cost function taking one's cooperating partners' αs as inputs. The optimal situation means the visionary state where a user's actions are not limited by other peers (by their constraints or wills). A service degradation a given user might face is e.g. being able to upload data for backup to partners by a lower connection rate than she would support. Checking on partner peers and the backuped data's integrity are also costly procedures, especially with lower α partners, not to mention eventual data re-settlements in the case when a partner becomes un-trusted/lost.

Definition 2 *The cost term that service degradation causes is user specific, but in each case takes the collaborating peers' αs, weighted by the backuped data amount at each one, as inputs. $D_i := f_i(c_i, (\alpha_j)) \to \mathbb{R}$, where (α_j) means the α vector of user i's partner peers.*

More will be discussed in 4.1.

3.3.2 Opportunity storage cost

As we intend to drop [11]'s "new data load at every re-arrival" policy, the cost term describing losses due to space offering on own disks needs to be reshaped. The opportunity cost remains the same extended with the little note that it is not proportional with online time, as peers intend to show reliable characteristics and thus keep backuped data persistently. Even in lower classes, peers should endeavor to respect the backup data stored on their disks otherwise their αs drop to low values which results in very costly backuping.

Definition 3 *The opportunity cost term is also user specific, and is a function of the given user i's c_i: $O_i := f_i(c_i) \to \mathbb{R}$.*

See 3.4 and 4.1 for further explanation.

3.3.3 Transfer cost

The transfer cost reflects the model's data management policy: main backup load transfers happen only when new partnerships are born. Looking at a long term evaluation (excluding moments where a given peer changes her strategy regarding the exchanged capacity) new partnerships are born at the same rate as existing partnerships break. A partner peer with a lower α is more susceptible to vanish from the system or more desirable to be swapped for a "better" peer to make a new deal of storage exchange. Both events require the backup data load to be moved (from storing peers to owner in order to regenerate the raw data from stored pieces and then from owner peer to storing peers again, in the case the data management is not delegated, i.e. the owner needs to handle cutting into pieces, erasure coding, encryption, etc.). It is why a given peer's transfer cost highly depends on the quality of its partners, and then the quality of their partners, etc. In general we can state that in this interdependent mesh of cooperating peers each and every peer experiences a transfer cost depending on its neighbors (in the sense of the collaborating overlay).

Besides the load transfers, when storing other users' data, the peer has to face with periodical integrity checks on itself (likewise the peer in question also supervises its partners) resulting in additional computational and bandwidth burden which also imply non-monetary cost. This term now depends on the given user itself, namely on its α_i, based on which partners will check it more or less often (a more reliable node requires less frequent checks) and these checks will hurt differently (e.g. a well connected peer will not lose large portion of its bandwidth during the procedure).

Definition 4 *The transfer cost term is also user specific, and takes in function user i's c_i, α_i and (α_j), i.e. the α vector of user i's partner peers: $T_i := f_i(c_i, \alpha_i, (\alpha_j)) \to \mathbb{R}$.*

3.3.4 Effort cost

Given an initial profile, a user may want to "upgrade" herself in the hope that she could cooperate with better peers to achieve a less costly quality service (seen that cost terms are due to one's partners weak αs). We call the cost of this step an effort cost. It may mean e.g. leaving a storing device online for longer periods, sacrificing more connection bandwidth to the service, etc., anything that eventually may cause the increase of the α parameter observed by partners. The cost is a function of one's initial (effortless) profile, the targeted level and of course a user specific parameter which describes her external bonds to overcome.

Definition 5 *The effort cost term is also user specific and depends on user i's initial effortless α'_i and α_i, the desired α strategy: $E_i := f_i(\alpha'_i, \alpha_i) \to \mathbb{R}$.*

3.4 The game

As we outlined it before the control level of data exchange also draws our attention not just the upper level management (consisting strategies and incentives) when we suppose that the distributed system consists of selfish, strategic users. Basically by control level we mean the peer selection process that determines the directions of partnerships and flows among users. Exchanges between storage-lenders need first selecting appropriate peers to cooperate with. This selection is user driven (as we are facing a distributed system), and as users are selfish, typical behavior in the way they treat this issue may be observed. We suppose that each peer seeks partners with higher α profiles to reduce own costs. Less cost is resulted by less required redundancy (thus less traffic of data shipping) when cooperating with a peer that is online most of the time; less traffic when checking less frequently a more reliable peer; faster up/downloads to a better connected peer, all of them come together with a higher α. The higher club of partners the user belongs to, the less replicas she needs to store: redundancy decreases, less overall storage, less traffic, less cost of checks.

3.4.1 Exchanged goods

Most of the existing works (e.g. [13]) consider a specific exchange rule in order to address the data availability management issue. Briefly this means that if a user offers c_i storage capacity for other users, she gets the right to store $\hat{c_i} <$ discount(c_i,α_i) in the system. The storage discount assures the continuous availability of backuped data despite the churn. A proper discounting function is subject of ongoing and further research.

When exchanging a unit storage space, the two participating partners understand that the unit means "useful" amount, which is the discounted value of the real reserved capacity. So when going to the market, a given peer i with α_i has to reserve c_i capacity at her own disk to be able to offer (and in the meantime ask for) a unit liability (discount($\alpha_i c_i$)). In the simplest approach the derivation to be made could be $\alpha_i c_i$. In this case the question of a proper discount function is shifted to the construction of the α parameter.

3.4.2 Exchange market

In our vision the peer selection process in the distributed backup/storage system would be similar to an open market, like a stock exchange. Let us suppose that peers introduce unit storage space fractions to the market: a given peer (with α_i) would like to exchange, i.e. offers an abstract capacity unit on her own disk and wants to store the same amount of data on an other peer's disk, holding *at least* α_j. In this case she introduces her ask on the market of α_i quality unit spaces as of α_j, and in the meantime she places bids on the markets of higher quality than α_i with her true value α_i. If an exchange becomes possible (the given peer's ask or one of her bids find a corresponding counter party), a deal (and the neighborhood relation) is born, and the relevant bids and asks are cleared from the markets.

In an extreme setting on the stock exchange example no peer would find any partners if α_i was strictly lower than α_j for every peer and no one undertook any slight loss in order to a make a deal (and thus to gain after all). In reality, as these peer profile parameters can not be determined with precise accuracy, we assume that individual markets are established on coarse α segments, and not on exact α values. We assume that a peer belonging to a specific α segment would also place a bid in the same segment's market, supposing that the expected service quality would be at least as good as she provides. Inside a market the quality difference between peers stating asks is then negligible, so whenever two ask-bid pairs appear from two different peers (belonging to the same market, i.e. with similar α) they clear right away. One of them will surely lose on the deal, because there are no two exact same peers in the system, but this loss is negligible based on the mentions above. On the contrary, if there is only one "local" ask-bid pair in a segment's market, the peer in question could decide if she cooperates with a lower bid peer. The decision should be the function of the possible partner's α and the relevant costs, and the expected waiting time for an other "local" peer in relation to the urgency of the exchange.

3.4.3 Clubs

The question of profile segments' granularity then arises: what kind of metric will determine if an α_i profiled peer is possibly able to trade on a given segment's market. If segments become too coarse, the higher profile peers face losses in deals among members that are not necessarily negligible any more. On the contrary if a segment were too fine (i.e. no large differences are allowed in peers' profiles), the number of eligible peers could drop to a critical level where dispersing data for backup is not possible. We will show later, that there is no need for determining the segments' granularity, the distributed user purposes will shape it.

The system will see a "clustering" process where users do not cooperate with much "worse" peers in terms of α. This global user profile parameter gives birth to "elite clubs" with members having outstanding attributes, and also direct incentives to less valuable users to try harder to get into one, since backuping at similarly "bad" peers costs a lot more. Our study targets the price of efforts a peer faces getting into and staying inside a "club", and the mutual benefits this kind of segmentation creates. As non-cooperative selfish peers take into account only their own payoffs, the question rises whether a given user is better off getting along the normal way, or it is beneficial to bring sacrifices for a higher quality of service inside a club. We consider the problem with the tool set of cooperative game theory as the above mentioned clubs may be seen as coalitions where selfish users gather to increase their payoffs, and the rest of the peer-set outside the coalition is not relevant. User rationality and incentive compatibility issues are targeted (e.g. incentive to increase α_i so to reduce churn).

4 Analysis

In this section we first present a deeper analysis of the user model, namely the effect of user's strategy on the terms of the user's payoff function. We also sketch the hints of the maximization problem a user faces. Afterwards we turn to the evaluation of the overall system when such users compose it, and give forecast on the run of the game (describtion in Subsection 3.4) how selfish users will team up in coalitions and ruin the social welfare.

4.1 Micro level: user strategies

Players will only join a coalition if they expect to gain from it. A user with an initial profile parameter may decide to pay on extra effort costs to save up on other costs by joining a higher class, but this is an issue of individual strategy-making. Participating in the system is the question of rationality, i.e. positive gain from the service.

As mentioned in Subsection 3.2 a given user i may take her choice on two variables: c_i and α_i. Each of them has effects on her payoff function of course.

- c_i determines the following value:
  - $O_i(c_i)$ which is generally an increasing quasi-convex function (like the general opportunity cost in literature) of the storage capacity that is offered for partners.
- c_i *and* α_i determine the following values:
  - $U_i(c_i, \alpha_i)$ which is generally an increasing quasi-concave function (like the general utility in literature) of the backuped data amount (determined by c_i and α_i).
  - $D_i(c_i, \alpha_i, n_i)$ which is the degradation due to partners with similar profile as peer i. n_i represents the number of peers available for cooperation inside a club. If this value drops drastically, the service degradation increases, so we assume that D is increasing and convex in n. This also implies a maximal limit of exchanged capacity between two given peers in order to avoid the birth of storage "centers". The approximation (i.e. similar profile partners are supposed) of partners' profiles is based on our beliefs about the form the game is going to achieve (see Subsection 4.2).
  - $T_i(c_i, \alpha_i)$ which is the transfer cost due to partners with similar profile to peer i's. The approximation again is assumed to be accurate because of the same reasoning as mentioned above at the service degradation cost.
- α_i determines the following value:

– $E_i(\alpha_i)|_{\alpha'_i}$ where α'_i is user i's initial (effortless) profile. E_i is assumed to be positive, quasi-convex increasing function for $\alpha_i > \alpha'_i$.

It is then straightforward to see that the optimal user strategy (c_i^*,α_i^*) is defined by solving the following equations:

$$c_i^*|_{\alpha_i} : \frac{\partial(U_i - O_i - D_i - T_i)}{\partial c_i}|_{\alpha_i} = 0, \tag{1}$$

$$\alpha_i^*|_{c_i} : \frac{\partial(U_i - D_i - T_i - E_i)}{\partial \alpha_i}|_{c_i} = 0. \tag{2}$$

For further study, more exact assumptions or real-life measurements are to be made, which is subject of our future work.

4.2 Macro level - cooperation - incentive compatibility

When carrying out analysis of the system on the macro level, one would be interested by the peers' partner selection strategies. As the basic assumption in our model is about the additional (compared to actual existing P2P systems) right of users to select their own partners to cooperate with, we are about to distinguish the two main types of decision a peer can make: either a given peer is selfish and tries to make bonds with the best peers she can, or the peer just does not care about her partners' quality. We show that the first choice is more advantageous for the user (it is in the best interest of every user not to deviate from the policy), especially if she does not have any information about the others' decisions and supposes that they lean to the first one. These effects result in a well-describable stable state to which the system approaches. To arrive to this conclusion, we deal with cooperative game theory when considering the segmentation of users.

4.2.1 User behavior

Definition 6 *A dominant user strategy of a strategic game $\langle N, (A_i), (\succeq_i)\rangle$, where N is the set of layers, (A_i) is the set of possible actions of users, and $(\succeq_i)$ represents the set of user preferences over the action set, is $a_i^* \in A_i$ of actions with the property that for player $i \in N$ we have $(a_{-i}, a_i^*) \succeq_i (a_{-i}, a_i)$ for all $a \in A$.*

When a user chooses to follow policy, she enters a club based on her strategy regarding her profile and possible efforts. Given a final profile, she would not cooperate needlessly with peers less "valuable" than she is, furthermore higher profile policy follower peers will not cooperate with her because of the same reasoning.

Definition 7 *A policy-follower peer's P_i is a function of $(\widetilde{\overline{\alpha_i}}, \alpha_i)$, where $\widetilde{\overline{\alpha_i}}$ is the average profile of defector peers having at least α_i, and α_i represents the same profiled partners inside the club.*

A user may decide not to follow the clear incentives of appropriate peer selection seen in 3.4.3. In this case we call the peer a *defector*, and suppose that her payoff function will reflect the service degradation and transfer costs due to average α partners, since a defector peer is assumed to pick cooperating partners randomly who are open for collaboration with her, i.e. other users who defect and policy-follower users with at most the same type of profile as the user in question.

Definition 8 *A defector peer's P_i then has $(\widetilde{\alpha}, \underline{\widetilde{\alpha_i}})$ as inputs, meaning that her costs will be defined by defector (with the average profile $\widetilde{\alpha}$) and (not strictly) lower profiled $(\underline{\alpha_i})$ policy follower partners.*

Proposition 1 *Under the assumption that both type of player (i.e. defector/follower) exist in large number in the system (i.e. finding and picking a partner of each type happens with similar probability) and one's stored data load is spread uniformly on peers independently of their behavior strategy, we state that the policy-following strategy is dominating.*

Proof: Based on the assumptions, the payoff function can be viewed as a function of selected partner qualities (partners from the policy defecting and following groups respectively) while the other user parameters (form of subsidiary functions, like U_i, D_i, etc.) and strategies (C_i) are compacted into the lower index of P, i.e. P_i. After the previous sections on reasoning, collaborating with a less-profiled peer brings more cost to a given user than cooperating with a better peer (if nothing else is changed) that is to say that $P(\widetilde{\overline{\alpha_i}}, \alpha_i) > P(\widetilde{\alpha}, \alpha_i)$, and similarly $P(\widetilde{\alpha}, \alpha_i) > P(\widetilde{\alpha}, \underline{\widetilde{\alpha_i}})$. So following the policy is always the best strategy (assures the highest payoff regardless the counter-strategy set). $\square$

4.2.2 Transition from grand coalition to clubs

The previous proposition assumes that every possible type of peer is present in the system to cooperate with. Starting with a peer set where nobody follows the policy (second input of Ps does not exist), the motivation level to carefully select buddies lies on the first input, i.e. possible defector partners' qualities, thus high profile peers have clear incentives. First the best peers will collude in order to reduce their own costs. As the highest quality peers change their strategies, and quit the camp of deviator peers, this latter will see a decreasing $\widetilde{\alpha}$, which re-generates the motivation of the best remaining peers to do the same, i.e. build coalitions. At the end hopefully every peer will sign up for a coalition but the worst ones. Thus the game can be characterized as a special coalitional game where the worth of a coalition is not divisible arbitrarily among its members. The suitable definition is from [8].

Definition 9 *A coalitional game without transferable payoff consists of*

- *a finite set N (the set of players)*
- *a set X (the set of consequences)*
- *a function V that assigns to every nonempty subset S of N (a coalition) a set $V(S) \subseteq X$*
- *for each player $i \in N$ a preference relation $\succeq_i$ on X.*

Coalitions will appear on segments of α in such way that they will contain enough peer, that satisfy the entry criteria, for a reliable service. So, in fact, segment granularity will be defined by the game itself. To have the ability deciding whether a bit lower level peer is welcome to a club or not, users have to hold a utility degradation factor due to the fact that fewer partners are available in a club than should be. The peers, as players of the game, will have the preference relation (based on their payoffs, function of c, α, partners' αs and the coalition's characteristics they belong to) on the set of possible consequences. Several conditions have to be satisfied before coalitions are formed based on α segments:

- The considered coalitional game is without transferable cost so the set of consequences for possible coalitions needs to be established in order to obtain the game's core.
- A newcomer's marginal contribution to a club highly depends on the latter's size. This issue is tackled by segment granularity, in order to have enough peers in each club.
- The variation of members' payoff caused by a quiting peer is the function of the leaving peer's profile (relative to the other members') and also the size of the club.

Definition 10 *The core of a coalitional game $\langle N, V, X, (\succeq_i)_{i \in N} \rangle$ is the set of all $x \in V(N)$ for which there is no coalition S and $y \in V(S)$ for which $y \succeq_i x$ for all $i \in S$.*

Proposition 2 *The core of the cooperative game is not empty: in equilibrium each user in the set will follow the peer selection policy and cooperate with similar profile peers inside a club that suits her the best.*

Proof:

A user payoff depends on her club's size and its segment size as well. If a club's size becomes large, i.e. many peers collude together, the highest part of the segment would be better off excluding the others. Given a coalition large enough with mixed profile peers, there will always be a subset coalition of the similarly high level peers who each have a better consequence if they secede from the grand coalition and leave the worse peers behind. But this is to be done in a way to reserve enough peers in the smaller segment club to be able to maintain a backup system (e.g. in an extreme setting, two very similar profile peers can not guarantee backup service for each other). In fact the same process goes along as in the initial period when small coalitions were created out of the "grand" coalition. The proposition states that this transformation stops at a point and the system finds itself in a stable equilibrium.

To describe the balance between the two opposing effects, a given user's payoff should be examined: at similarly high level α peers, low value club members will cause the same increase in service degradation and transfer cost, but excluding these peers may cause higher service degradation effect on each. Again, we assume that cooperation is spread among many peers inside a club, so every high level member is affected by low level members. Keeping them inside the club is just reasoned by the mandatory minimal critical size of a club to maintain the backup service.

Let us consider the ordered (by profile parameter) set of peers $(1, 2, \ldots, n)$. We suggest that at equilibrium the peers form clubs "continuously" on this ordering, i.e. a given club contains $(i, i+1, \ldots, j-1, j)$ peers and no one else. We prove that no group would deviate from this state. $i, i+1, \ldots, j$ are members of the same club and the lowest profiled peer is j. No peer can be excluded from the group, since the payoff of the remaining peers would drop because of the too low size of the group (it is straightforward to see that if they represented a sufficiently sized population, by re-organizing the collaboration bonds among themselves, the high level peers would gain more if they gave up the cooperation with the "weak link" j). No peer can be exchanged for an other one: peers indexed higher than j would cause more cost than any actual member, on the other hand peers with lower index than i would lose by cooperating with the club in question instead of their club in the stable set. For the same reasoning, no additional peer should join the given club: higher indexed peers are not welcome, and lower indexed ones would not want to.

At equilibrium $\frac{\partial(D_i+T_i)}{\partial\alpha_i} = -\frac{\partial D_i}{\partial n_i}$ for $\forall i$ inside a club when trying to exclude the lowest profile or to receive a new member, where the partial differentiation in α_i means the change in partners' average profile. $\square$

4.2.3 Cost or gain of anarchy

By anarchy one would think of the grand coalition scheme where each user picks partners randomly without taking into account their profile levels. However under some assumptions, that would be the setting for the optimal social welfare, and the core consisting smaller clubs (in Proposition 2) is equivalent to the "anarchy" state, where every user acts selfishly.

Assumption 1 *We suppose that the joint value of disjoint coalitions is no less than the sum of their values (i.e. we assume super additivity, like most of the coalitional games). The requirement for our super additive game is the following:*

$$\frac{\partial(D_i + T_i)}{\partial\alpha} \geq \frac{\partial(D_j + T_j)}{\partial\alpha}$$

for $\forall i, j$ such that $\alpha_i \geq \alpha_j$ when the two peers of type α_i and α_j start to cooperate (better peer's payoff's input (α) decreases and the cost functions grow, worse peer's payoff's input increases so costs drop).

To derive the optimal pairing of peers in the P2P backup system under Assumption 1 is straightforward. This means that when a new peer joins a group of backup data exchanging peers, ones' payoffs may drop, but others gain in the way that the sum of all peers' payoff increases. As our case is a game with non-transferable payoff, the growth in social welfare can not be divided among the users arbitrarily, thus some of them will face loss, so user selfishness opposes social welfare.

Proposition 3 *In the super additive case the grand coalition has the highest overall social welfare, which means that in fact choosing randomly the collaborating peers (on the whole peer set) is the optimal solution in the usual terms.*

However, we emphasize again, this scheme is not aligned with user selfishness since more contributing peers will suffer in the cooperation with lower profiled peers; and also lacks incentives towards the peers to maintain certain level of profile.

5 Illustrative example

For a better understanding, let us present a little case study. On the one hand this short example is made of rules and functions which are extremely simplified and could be not in line with all the assumptions previously discussed, but on the other hand our main target, the effects of selfish peer selection and the rise of a coalitional game are presented clearly. When comparing Table 3 to Tables 1 and 2, we can notice that the best peers have higher payoffs when gathering into an elite club than ever before, so if they have the possibility to act selfishly (in actual P2P system peer selection is not the exclusive right of peers), coalitions will be made on the peer-set.

We consider a user set with 10 peers, and establish the sub-functions of their payoffs as follows:

- each user wants to store 9 units of backup data and has the utility of service of 100.
- user parameters (αs) are 1, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2 and 0.1.
- storage exchanges are based on $\alpha_i c_i$, i.e. to store 9 units in the system, a peer with $\alpha_i = 0.3$ has to offer 27 units of capacity for others, so as to introduce redundancy for efficient and successful data management.
- opportunity cost is quadratic and normalized, i.e. 27 offered units yield $\frac{27^2}{100}$ cost.
- degradation cost is the derivation of exchanged capacity and partners' profiles: $\sum_j \frac{\alpha_i c_i}{n-1}(1-\alpha_j)$, where n is the user i's club's size, and α_js are the profiles of its partners.
- a club needs to contain at least $\left\lceil \frac{1}{\min_{i \in S} \alpha_i} \right\rceil + 1$ members[1] otherwise degradation becomes ∞. Surely this is not a bullet-proof data management policy, but as we describe just a toy example, we would like keep it simple and attract the reader's focus on the main issues.
- transfer cost is the derivation of exchanged capacity, partners' profile and user's profile: $\sum_j \frac{\alpha_i c_i}{n-1}(1-\alpha_j)\alpha_i$.
- effort cost is supposed to be equal to $100(\alpha - \alpha_i)$, i.e. the profile hop multiplied by 100.

In this example setting, the simple grand coalition yields the following payoffs summarized in Table 1. Again, we have to emphasize that this is a simplified example just to show the effects of coalition forming. User utility on backup service is the same for every peer, and we do not consider individual strategy making on selecting the optimal capacity to exchange (exchanged abstract data amount is the same for every user) either. Furthermore we suppose, which is realistic in fact, that every user spreads its data uniformly on the other members of its club, i.e. 1 abstract data unit on each partner in the grand coalition.

If we suppose that peers try to maximize their payoffs by making efforts to improve their profiles in order to decrease the mandatory offered capacity and thus the opportunity cost, the

[1]Supposing a club sized n where members spread their data-to-backup evenly on partners, the user with profile α_j will store $\frac{c_i}{n-1}\frac{1}{\alpha_j}$ for member i which should not exceed c_i. This constraint gives the simple rule on minimal club size.

	User payoff function					
	U_i	O_i	D_i	T_i	E_i	P_i
1.	100	0.81	4.5	0	0	94.69
2.	100	1	4.5	0.45	0	94.05
3.	100	1.27	4.5	0.9	0	93.33
4.	100	1.65	4.5	1.35	0	92.5
5.	100	2.25	4.5	1.8	0	91.45
6.	100	3.24	4.5	2.25	0	90.01
7.	100	5.06	4.5	2.7	0	87.74
8.	100	9	4.5	3.15	0	83.35
9.	100	20.25	4.5	3.6	0	71.65
10.	100	81	4.5	4.05	0	10.45

Table 1: **Grand coalition:** No peer takes effort to increase its profile.

payoffs become the following values in Table 2. After our simple rules (i.e. effort cost is assumed to be linear on profile growth), the worst two peers decided to enhance their profiles (from 0.2 and 0.1 up to 0.3, by not going into more fine-tuned profile improvements) and thus to minimize their costs. This transformation changed the general degradation and transfer costs over the whole peer set too.

	User payoff function					
	U_i	O_i	D_i	T_i	E_i	P_i
1.	100	0.81	4.2	0	0	94.99
2.	100	1	4.2	0.42	0	94.38
3.	100	1.27	4.2	0.84	0	93.69
4.	100	1.65	4.2	1.26	0	92.89
5.	100	2.25	4.2	1.68	0	91.87
6.	100	3.24	4.2	2.1	0	90.46
7.	100	5.06	4.2	2.52	0	88.22
8.	100	9	4.2	2.94	0	83.86
9.	100	9	4.2	2.94	10	73.86
10.	100	9	4.2	2.94	20	63.86

Table 2: **Grand coalition:** Last two peers make effort to increase their profiles.

Based on our Proposition 2, the steady state of the cooperative game will be realized if users act selfishly and select their buddies in that view. Again, we suppose that data load is spread evenly among club members, i.e. 4.5 units on each partner in a club of 3. The first 3 peers will create an elite club and the second group of 3 an other one, mainly due to our simple degradation rule on the minimal size of a club. The 4 worst peers in the set are left alone, and effort is required of 3 of them to be able to create a coalition of 4 (following the simple minimal-size policy) otherwise they would not be capable of participating to the system. We do not consider here the possibility that they would raise their profiles such that they could kick out some peers of a club, because with the too-simplified applied effort cost function it would lead us to unrealistic results. However we shall notice that the drastic drop of payoffs at the worst peers may have a strong incentive effect on the peers ahead to improve their profiles. The worst peers of clubs may also think about improving themselves, but here we do not consider these phenomena either.

	User payoff function					
	U_i	O_i	D_i	T_i	E_i	P_i
1.	100	0.81	1.35	0	0	97.84
2.	100	1	0.9	0.09	0	98.01
3.	100	1.27	0.45	0.09	0	98.19
4.	100	1.65	4.05	1.22	0	93.08
5.	100	2.25	3.6	1.44	0	92.71
6.	100	3.24	3.15	1.58	0	92.03
7.	100	5.06	6	3.6	0	85.34
8.	100	7.44	5.2	3.47	3.33	80.56
9.	100	7.44	5.2	3.47	13.33	70.56
10.	100	7.44	5.2	3.47	20.33	60.56

Table 3: **Little coalitions:** with only unavoidable efforts.

6 Discussion

6.1 Conclusion

In the paper we present a realistic game theoretical user model for P2P backup/storage systems and we draw the conclusion that introducing some specific system management policy (namely user evaluation and data redundancy both based on α) will lead to a cooperative game where selfish users get together to found well describable clusters. This work depicts the possible outcomes of handing over peer selection into user's jurisdiction, which may be the novel feature in the new generation of P2P systems. External user-forcing incentives then become needless, and fairness reaches higher levels then ever from the user's point of view, while the P2P system turns into a real distributed system where commonly accepted policies drive users' actions, but nothing else. The approach we define realistic user characteristics and strategies is novel in the domain of P2P system analysis, and results in useful findings for the design of distributed systems, whatever their use is, if selfish autonomous users are considered.

6.2 Future work

The presented work is however not completed: many questions are left open. These are all subjects of our future research.

6.2.1 Simulation

We are planning on building a sophisticated simulation for modeling a real life implementation of a P2P backup/storage system. User profiles are to be modeled based on measurements, distribution describing e.g. user connectivity. The numerical simulator also needs the exact definition of realistic payoff function forms, thus user preferences and faced costs are also subject of further research. The future outcomes of the simulation will give us the opportunity for deeper analysis.

6.2.2 Profile parameter α

As it is supposed to be clear by now, the α parameter plays a crucial role in the system. Thus establishing α needs more attention, and has to take into account the following requirements and thoughts:

- to avoid system design enforcements, the background's content of α (i.e. the importance preference for malicious behavior, online availability, connectivity, etc.) must be the same for every peer in the system, so it must reflect real considerations,
- it must assure sufficient data redundancy since the capacity exchange rule is based on the abstract αc policy,

- since α determines peer selection, the constraint on minimal club size, instead of $\min n_i$, should also be connected to it through the best-fit assumption, e.g. $var(\alpha) \geq k$ or $\sum_{i \in S} \alpha_i \geq l$ in order to successfully model fundamental user preferences.

References

[1] Ranjita Bhagwan, Kiran Tati, Yu-Chung Cheng, Stefan Savage, and Geoffrey M. Voelker. Total recall: System support for automated availability management. In *Symposum on Networked Systems Design and Implementation (NSDI)*, 2004.

[2] Byung-Gon Chun, Frank Dabek, Andreas Haeberlen, Emil Sit, Hakim Weatherspoon, M. Frans Kaashoek, John Kubiatowicz, and Robert Morris. Efficient replica maintenance for distributed storage systems. In *Symposum on Networked Systems Design and Implementation (NSDI)*, 2006.

[3] P. Brighten Godfrey, Scott Shenker, and Ion Stoica. Minimizing churn in distributed systems. *SIGCOMM Comput. Commun. Rev.*, 36(4):147–158, 2006.

[4] Dominik Grolimund, Luzius Meisser, Stefan Schmid, and Roger Wattenhofer. Havelaar: A Robust and Efficient Reputation System for Active Peer-to-Peer Systems. In *1st Workshop on the Economics of Networked Systems (NetEcon), University of Michigan, Ann Arbor, Michigan, USA*, June 2006.

[5] A. Haeberlen, A. Mislove, and P. Druschel. Glacier: Highly durable, decentralized storage despite massive correlated failures. In *Symposum on Networked Systems Design and Implementation (NSDI)*, 2005.

[6] Patrick Maille and Laszlo Toka. Managing a peer-to-peer data storage system in a selfish society. *IEEE JSAC Special issue on "Game Theory in Communication Systems"*, 2008. Submitted.

[7] Petteri Nurmi. A bayesian framework for online reputation systems. In *AICT-ICIW '06: Proceedings of the Advanced Int'l Conference on Telecommunications and Int'l Conference on Internet and Web Applications and Services*, page 121, Washington, DC, USA, 2006. IEEE Computer Society.

[8] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory.* The MIT Press, July 1994.

[9] R. Rodrigues and B. Liskov. High availability in dhts: Erasure coding vs.replication. In *International Workshop on Peer-to-Peer Systems (IPTPS)*, 2005.

[10] Laszlo Toka. Peer-to-peer sharing of hars disks: study of incentive mechanisms. Master's thesis, Eurecom Institute, http://www.eurecom.fr/ toka/pub/toka.pdf, 2007.

[11] László Toka and Patrick Maillé. Managing a peer-to-peer backup system: Does imposed fairness socially outperform a revenue-driven monopoly? In *GECON*, pages 150–163, 2007.

[12] Hakim Weatherspoon and John D. Kubiatowicz. Erasure coding vs. replication: A quantitative comparison. In *Proceedings of IPTPS'02*, Cambridge, MA, 2002.

[13] Wuala. Peer-to-peer storage system. http://wua.la/en/home.html.