Department of Telecommunications and Media Informatics
Budapest University of Technology and Economics

Eurecom
Telecom Paris

Game Theoretic Analysis of Distributed Systems: Design and Incentives

László Toka

Thesis booklet

Advisors:

Dr. Attila Vidács

High Speed Networks Lab
Dept. of Telecommunications and Media Informatics
Budapest University of Technology and Economics

Dr. Pietro Michiardi

Dept. of Network and Security
Eurecom
Telecom Paris

2010.

1 Introduction

My research work aims at designing distributed systems with focus on fairness and incentives in resource provisioning. The thesis proposes game theoretic models of peer-to-peer (P2P, distributed network of functionally equal peers) backup and decentralized radio spectrum sharing solutions. The inherent selfish behavior of participants is mitigated by well-suited incentive schemes. Analytic and numeric results reflect the performance evaluation of the advised system design frameworks.

1.1 Background

Nowadays, more and more information technology services and applications employ the distributed paradigm in order to ensure scalability and robustness. Self-organizing and distributed systems, although different in most technical aspects, reflect similar incentive issues. Many distributed services currently rely upon altruistic behavior from their users. The phenomenon of selfish individuals who opt out of a voluntary contribution to the common welfare of the group has been widely studied, and is known as the free-rider problem. It is important to design mechanisms that encourage peers to contribute resources and reduce free-riding behavior in distributed systems.

Besides the extended research work on technical aspects, the related economic system characteristics have also been tackled in growing measures recently. Various incentive solutions are proposed for many distributed systems, such as network access sharing [25, 26], P2P file sharing [5, 2], network routing [10], packet forwarding in ad-hoc networks [4, 31, 1], spectrum allocation [20], P2P storage and backup [8, 29, 2, C9, J4], network content caching [22, 23], and network formation [9, 6, 21].

1.2 Motivation

I target distributed multi-user systems which cannot be operated socially optimally without adequate design. Specific incentive schemes must be deployed to enhance the quality of service for the system users. According to the decentralized nature of the systems, these incentive schemes must avoid relying largely on central entities.

The first part of the work aims at a relevant field of research: there is growing need for seamless, secure, reliable and easily accessible online backup as the daily used electronic devices are often connected to the Internet and increasing transmission rates make transfers of large data amounts possible. Since instead of public or central resources the users of the P2P backup system exploit those of one another, encouraging resource sharing requires a well-suited incentive scheme.

The second part of the work targets radio spectrum allocation, studying a distributed auction-based management scheme. The main motivation behind the investigated approach is that the sequence of central auctions to reallocate public resources should be transformed into a more scalable framework. Therefore, in the presented framework participants trade the acquired resources among themselves in a distributed design without the intervention of a central auctioneer.

2 Goals

I aim at building specially tailored design for two different types of distributed systems: P2P backup and distributed radio frequency allocation. In both cases, the non-cooperative selfish behavior of participants can jeopardize the operation. Based on user models, my goal is to design optimal economic incentive solutions that ensure desirable quality of service, and to evaluate their performance through both analytic and numeric investigations.

After recognizing the particularities of each application domain, I define user models accounting for their benefits in terms of system performance, their costs for resource sharing (if applicable), and their characteristics relevant from the investigated system perspective, *e.g.*, the heterogeneity of shared user resources or interference relations among users. I address resource allocation by building novel incentive frameworks for the studied systems. In order to foster cooperation among users in the service infrastructure barter-based and monetary payments, *i.e.*, pricing, are applied.

To evaluate the system models and the proposed incentive schemes, I utilize a broad class of analytic tools: matching theory models for the P2P backup system, and auction theory for the allocation of radio spectrum. I decompose the appeared optimization problems and develop distributed algorithms to solve them. I perform simulations of my theoretic models as proofs of concept to finally arrive at potential implementation of practical, feasible and scalable applications. These frameworks, with embedded incentive schemes, ensure the expected favorable outcomes in robust systems.

In Section 4.1, I present a P2P system design that adapts the resource contribution of users to a low level while ensuring a fair quality backup service. In Section 4.2, I analyze a model of peer selection in P2P backup systems in which users have the ability to selfishly select remote peers they want to exchange data with in a symmetric scheme. In Section 4.3, I investigate the possibility of allocating radio spectrum among multiple applicants dynamically in a distributed manner.

3 Methodology

I model distributed systems with the tool-set of *game theory* [28, 12, 27, 17], a means for modeling individual user preferences, strategies, costs and valuations. Analytical investigations tackle selfish user behavior, the existence of best-response strategies and equilibrium. I build incentive schemes [11, 3, 7] to align selfish participant behavior with the goals of the system design.

I also employ graph theory and matching theory to analyze the proposed incentive mechanisms. Matching theory [14, 18, 19], a field of combinatorial optimization, provides useful tools to investigate, among several other possible targets, *e.g.*, peer selection in P2P file sharing [24, 13] systems.

Finally, I perform numerical evaluations with simulations written in MATLAB.

4 New Results

4.1 The Design of a P2P Backup System

Thesis group 1: [C7, C8] I have proposed a P2P system design that adapts the resource contribution of users to a low level while ensuring a fair quality backup service.

I studied backup service in P2P systems, where users save their data on the unexploited storage devices of other users at different locations over the Internet, for free (Figure 1). As a main consequence, no scalability problems arise (higher number of users come with larger overall storage space), moreover the geographical as well as the ownership diversity of the storing hosts ensure great safety for the backed up data. The system must be designed to maintain high durability of the backup, and recoverability of local files at an event of data loss.

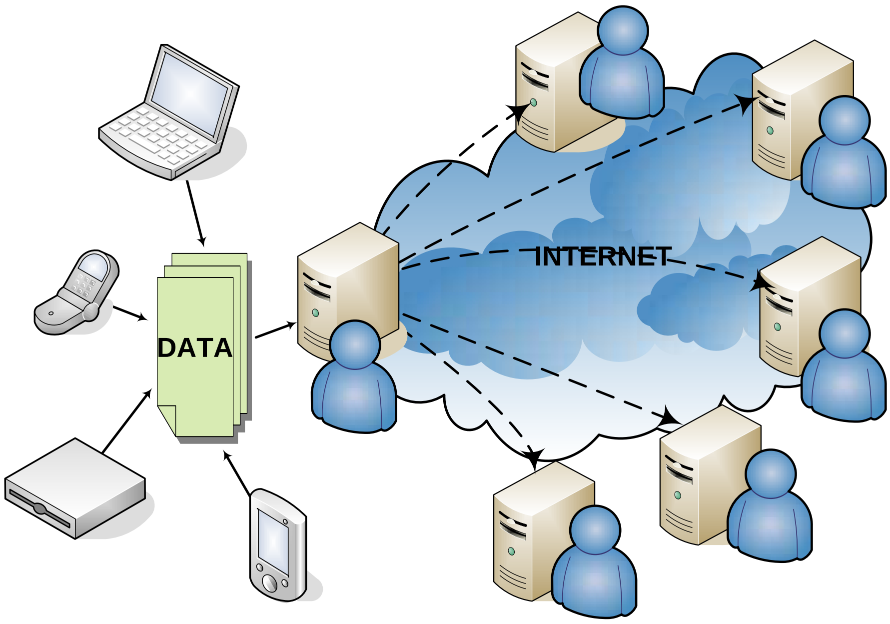

Figure 1: The backup application, running on the computer of the user, connected to the Internet, safely stores user backup by scattering the secure copies on other participating computers. After a loss of local data, the program recovers the required data from the storing participants.

Storing data on temporarily unavailable peers causes temporary data loss. To remedy this phenomenon, data must be stored in multiple copies, redundantly. The applied data redundancy scheme is erasure coding [30], in which backup data is divided into k fragments that are subsequently transformed to create $n > k$ redundant fragments. These latter are stored on remote peers, and any given subset of them, not less than k, are sufficient to restore the backup. The backup gets lost if less than k fragments can be retrieved from remote peers in case of need.

The management of data redundancy, and the network operations greatly affect the quality of service. The investigation of these aspects were the first subjects of my research.

In order to evaluate different system designs, first I define simple performance metrics to describe the quality of the backup service. The first metric reflects data durability by the probability of losing so many encoded fragments out of n, stored on remote peers, within a given time frame T, so that the remaining fragments are insufficient to restore the original data.

Definition 4.1. ***Data loss probability*** *(DLP) is $F_t^{n,k}(T)$, the value of the cumulative distribution function of the time duration t within which $n-k+1$ fragments are lost at input T. $F_t^{n,k}$ depends on the given set of the n storing remote peers,* i.e., *their crash rates, and k.*

The second set of metrics reflect the duration of data archiving and recovery processes.

Definition 4.2. ***Time-To-Backup (TTB) and Time-To-Retrieve (TTR)*** *TTB (resp. TTR) of a user is the time elapsed while the user uploads (resp. downloads) a number of encoded fragments to (resp. from) remote storage locations which meets the target redundancy (resp. is sufficient to restore the original data). TTR is definable only if the user has backed up at least k fragments before starting to retrieve a subset of them.*

In an ideal storage system with unlimited capacity and uninterrupted online time that backs up user data, TTB and TTR only depend on the backup data amount and on the bandwidth capacity and availability of the data owner. I defined $minTTB$ and $minTTR$ of a user as baselines for backup and restore operations which bound both its TTB and TTR. I used these reference values throughout the paper to compare the relative performance of the P2P application versus that of such an ideal system.

Definition 4.3. minTTB *and* minTTR *are the TTB and TTR of a user respectively in an ideal system. A peer i with upload and download bandwidth u_i and d_i starting the backup of data with size o at time t completes its backup at time t', after having spent $\frac{o}{u_i}$ time online. Analogously, i restores a backup object with the same size at t'' after having spent $\frac{o}{d_i}$ time online. I defined $minTTB(i,t) = t' - t$ and $minTTR(i,t) = t'' - t$.*

I evaluated various options of data redundancy and its maintenance in numerical simulations based on these metrics. Furthermore, I advocated an *adaptive* redundancy scheme that is based on the estimations of DLP and TTR, and determines the redundancy rate, *i.e.*, $\frac{n}{k}$, accordingly. The scheme handles the unavailability of backup data flexibly, and focuses on reaching relatively low TTR values, while keeping the estimated DLP low.

Thesis 1.1. *[C7] I have proposed to determine the data redundancy rate $\frac{n}{k}$ with a method that is based on TTR and DLP estimations. The scheme guarantees high service quality for backup purposes, and keeps the storage and bandwidth requirements imposed on users low.*

The method works as follows. After generating encoded fragments, each peer uploads an increasing subset of those until its DLP *and* TTR estimations fall below predefined thresholds. Otherwise, more fragments are stored on remote peers, and the *adaptive* redundancy rate is continuously increased. I propose a heuristic estimations for the DLP and TTR metrics, since the stored fragment losses and the length of TTR after a local crash at a given moment in time is hard to predict due to the unreliable nature of remote peers.

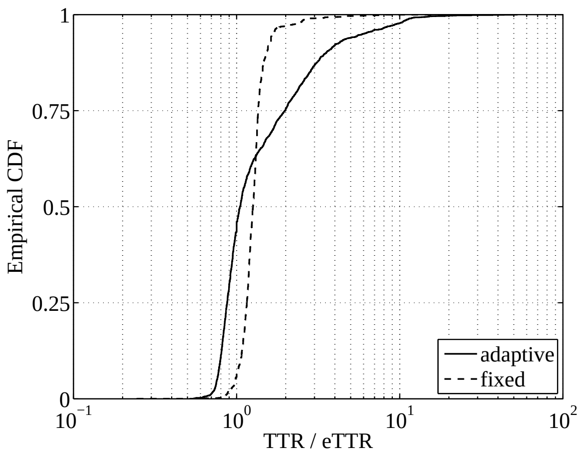

Figure 2: Measured vs. estimated TTR

The estimated TTR (eTTR) is calculated as the time a peer would need to download k fragments from the storage peer that has the kth highest product of long-term time average of the probability to be found online, called as availability (a) and average upload bandwidth (u). Formally, $a = \frac{1}{t_c - t_0} \int_{t_0}^{t_c} \mathbb{P}^t(\text{peer is online})dt$, where t_0 denotes the point in time where the peer started to use the service, and t_c is the current time. My heuristic eTTR writes as (1) where j is the kth "fastest" storage peer. The downlink of the retrieving peer might be saturated with maximum p^d parallel downloads, and the lowest TTR is given by $minTTR = \frac{o}{d}$, which is attained only if no useless unfinished fragment parts are downloaded with the download capacity d. eTTR can be computed for peers that have already uploaded at least k fragments, the number of fragments into which the backup data is originally divided into.

$$eTTR = \frac{o}{\min\left(a_j u_j p^d, d\right)} \tag{1}$$

Upon a crash, a peer with n fragments placed on remote peers can lose its backed up data if more than $n-k$ of them crash as well before k fragments are completely retrieved. Considering a delay that can pass between the crash event and the beginning of the retrieval phase, I compute the estimated DLP (eDLP) within a total delay of $t = delay + eTTR$ by considering peer crashes to be memoryless events, with constant probability for any peer and at any time. Therefore the times before crash are exponentially distributed stochastic variables with a parametric average $\overline{t}$: a peer crashes by time t with probability $1 - e^{-t/\overline{t}}$:

$$eDLP = \sum_{i=n-k+1}^{n} \binom{n}{i} \left(1 - e^{-t/\overline{t}}\right)^{i} \left(e^{-t/\overline{t}}\right)^{n-i}. \tag{2}$$

Simulation results of fixed-rate (that aims to ensure that at least k fragments of any backup are online with 99% probability, based on the average peer availability) and my adaptive-rate redundancy schemes are shown in Figure 3. eTTR values are calculated when crashes occur, and compared to the subsequently measured TTR: eTTR gives a fairly good estimate (Figure 2). The significantly lower adaptive-rates result in decreased TTB values, in return of prolongated TTR results, offering remedy for every user against long TTB, and punishing *only* those with longer TTR who lose their local copy.

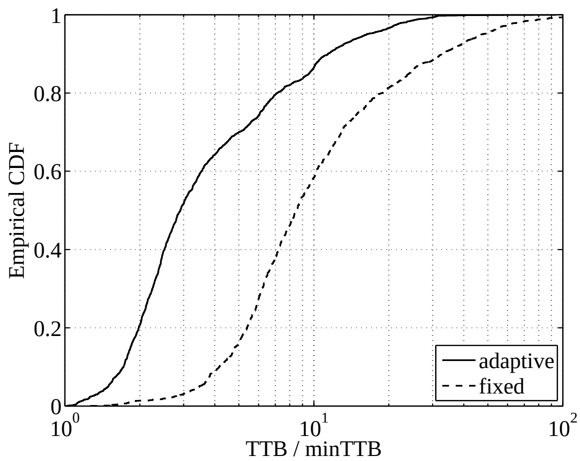

(a) TTB

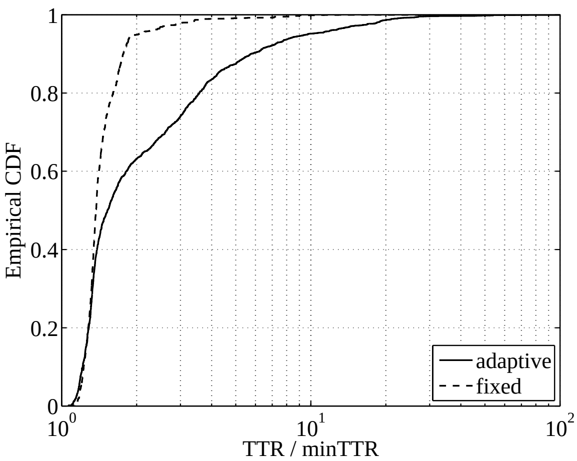

(b) TTR

Figure 3: Analysis of the adaptive redundancy scheme

In order to maintain a P2P backup system, participating users must share three types of resources: disk space, bandwidth and time spent connected to the network. When peer selection is random, peers with high availability and good connectivity receive excessive storage burden: since they are found online more often, more fragments are uploaded to them than to peers with lower reachability. Therefore, building on my novel data redundancy scheme, I have focused on fairness: I suggested a barter-based approach to encourage resource sharing in return for high service quality.

Thesis 1.2. *[C7, C8] I have proposed a method that ensures fairness in terms of quality of service and resource contribution. This scheme groups peers based on their availability and bandwidth characteristics and impose symmetric fragment exchanges within the groups.*

The adaptive redundancy scheme considers estimations based on the product of availability (a) and average upload bandwidth (u) of remote peers. I denote this value as $g_i = a_i u_i$, and call it the grade of peer i. I have shown that grouping peers based on their grades, and imposing symmetric fragment exchanges within the groups lead to less storage and traffic burden on peers that contribute high online availability and bandwidth (Figure 4(b)) than in a random peer selection framework (Figure 4(a)).

Furthermore, the quality of service that each participant may receive is limited by its contribution: low grade peer cannot exploit the abundant resources of high grade peers. As a consequence, the number of data loss events changes by switching from random to the grouped peer selection scheme (Figure 5), especially due to long backup phases in the low grade class. I categorize data losses in the following way: if the crashed peer

1. has not spent enough time online to upload k fragments until its crash the data loss is unavoidable: no online backup system could have saved the data, because data loss is determined by the limits in resources of the data owner;

2. has spent enough time online to upload at least k fragments until its crash, but it has not succeeded: in this case, the limited resources of remote peers are the cause of data loss, since backing up on an always-on data center with plenty of bandwidth would have succeeded;

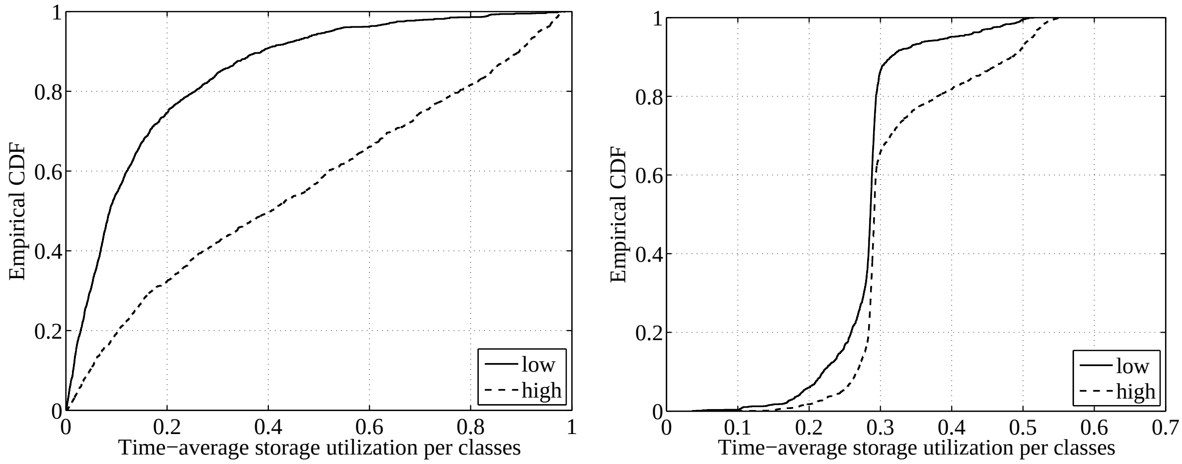

(a) Storage load on peers (random)

(b) Storage load on peers (grouped)

Figure 4: Fairness in grouped peer selection

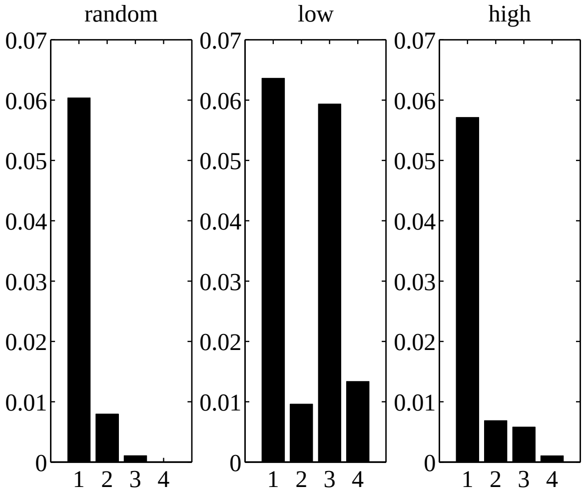

Figure 5: Fatal fraction of peer crashes per group in grouped peer selection

3. has uploaded at least k fragments, but its backup phase has never completed, then the loss is due to elongated TTB and to the fact that remote peers can crash;

4. has completed its backup phase, but it fails to retrieve at least k fragments after the crash, before its storing peers crash in a fatal number (prolongated TTR).

While more data losses happen in both groups than in a system applying random peer selection, the high grade group members suffer from a significantly lower number of data losses than low grade peers. Indeed, the fairness introduced by the grouped peer selection does not affect every user in the same extent: fast fragment transfers between the data owner and storer, both high grade peers, results in short backup phases, hence lower probability for fatal crashes than among low grade users.

A P2P system might not be able to guarantee the appropriate quality of service with the low number of users and/or based only on the shared user resources, *e.g.*, insufficient storage capacity. Therefore, I examined the effects of introducing a central storage server in order to avoid such (provisional) situations: I have shown the cost implications of per-

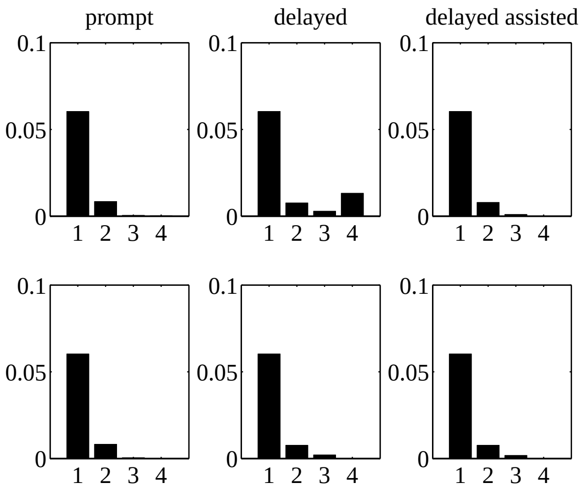

Figure 6: Fatal fraction of peer crashes with adaptive-rates (top) and fixed-rates (bottom)

sistent quality guarantees. In such a hybrid system the central, reliable (highly available) server might be used to store data in exchange for the reimbursement of costs. I expressed the relatively low occurring costs of such hybrid systems.

Thesis 1.3. *[C7] I have proposed a hybrid system architecture in which a data center complements the P2P resources in order to improve the quality of service by assisting to the backup phases and to the repairs of lost fragments for a reasonable cost.*

The data center can quickly retrieve and store fragments in cases where the backup is jeopardized because many fragments are lost and the data owner is temporarily offline, precluding any redundancy maintenance performed by it. Figure 6 represents the number of the observed data losses, labeled based on the above categories. The first row of plots are the results when using the adaptive redundancy scheme, the second row are the outcomes of the fixed-rate scheme defined above. The first column represents the case where the crashed peer appears online right after its crash, and remote peers are notified without delay; the second column shows a scenario when crashed peer remain offline for a week on average, and others are notified only after a week; the third column stands for the simulations where maintenance is similarly delayed, but assisted by a data center.

As shown in the third column of Figure 6, the number of data losses is lowered back to the case of prompt retrieves and repairs, in return for upload traffic (and intermittent storage load) of the data center, plotted in Figure 7(a). The role of the data center is more important when low adaptive-rates are applied, hence it performs more intensive outbound traffic which generates higher cost. On the other hand, the crash-related repair traffic, performed by peers, is proportional to the overall redundant data amount. Therefore it is significantly lower when applying adaptive-rates (Figure 7(b)), constituting an advantage of the adaptive scheme.

With assisted backups, peers may upload fragments to a data center during their backup phases if remote peers are not available temporarily. In my scheme (labeled as “opportunistic”) uploads to the data center are performed only with excess upload bandwidth. An alternative gives priority to uploads toward the data center and backing up to remote peers is started only after having uploaded all original fragments to the data center (thus called “pessimistic”). The pessimistic data placement policy ensures that

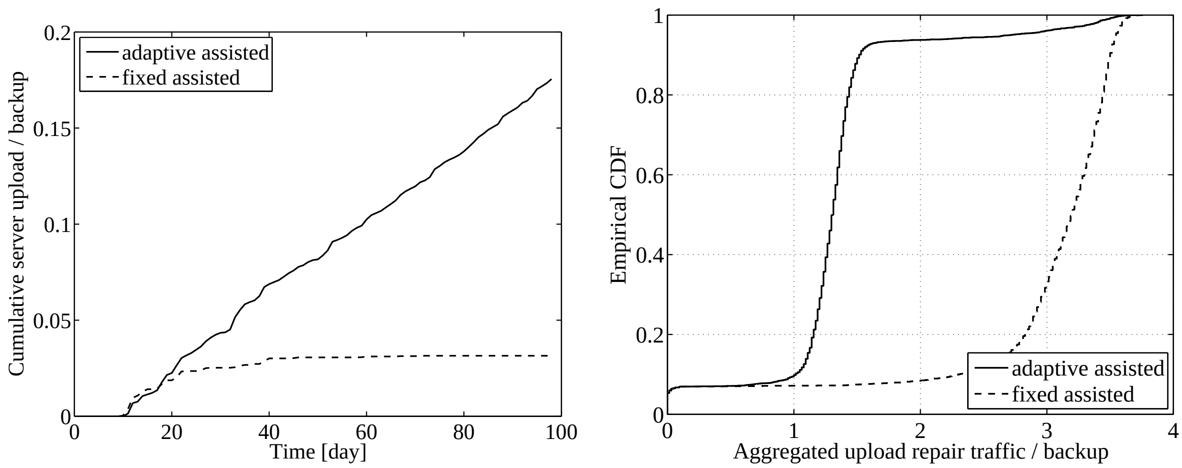

(a) Data center (b) Peers

Figure 7: Repair traffic with assisted repairs

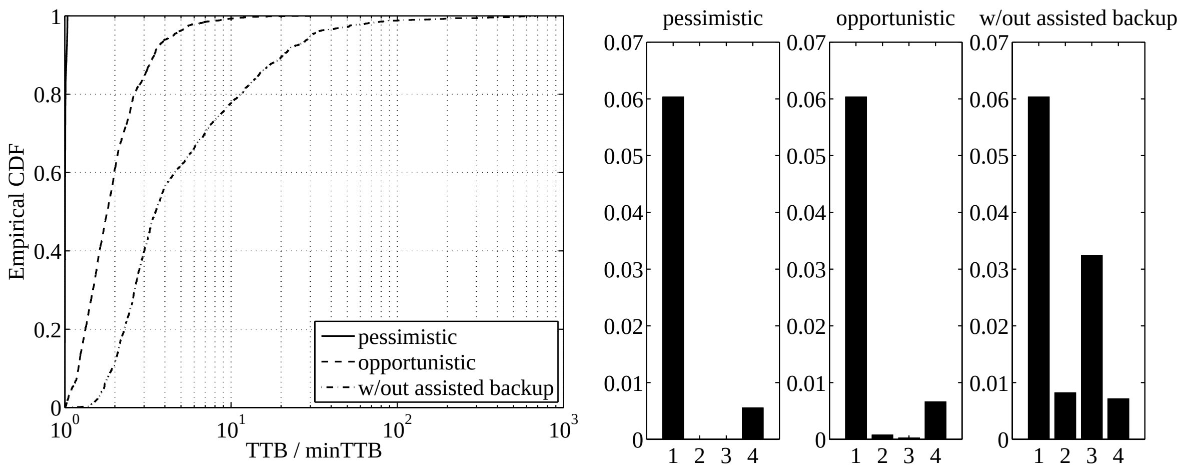

(a) TTB (b) Fatal fraction of peer crashes

Figure 8: Benefits of assisted backup

the TTB of every peer is effectively decreased to its $minTTB$. In this "safer" scheme the backup data is first entirely uploaded to data center to build reliable backup as soon as possible, then the centrally stored data is continuously deleted to save on the storage costs as the backup amount, successfully transferred to remote peers, is growing.

Assisted backups mitigate the negative effects of long data transfers toward peers with poor availability and connectivity (Figure 8). While uploading only to remote peers may be slow due their unavailability, in assisted backup schemes uploading to data center is only constrained by the availability and uplink capacity of the peer, therefore the DLP and TTR targets are reached earlier, hence smaller TTB values. Due to much longer TTBs, the data loss results (categorized in regard to the cause of the losses, as in Figure 6) during the backup phase are significantly worse without assistance. Once the backup is considered complete on remote peers, the achieved DLP and TTR targets ensure the same quality of service in all schemes, *i.e.*, the rates of data loss due to long TTR are similar.

The costly central storage (Figure 9) grows rapidly in the beginning of the backup

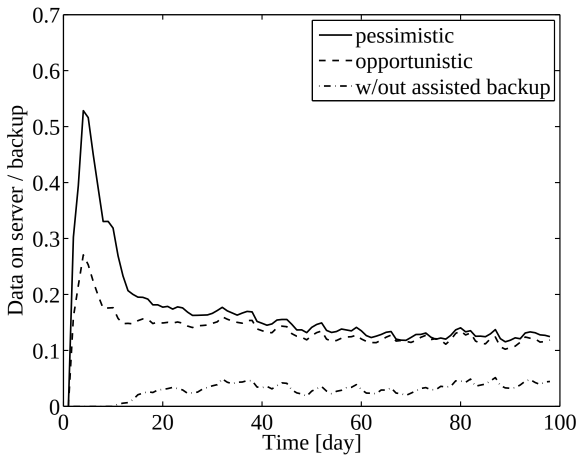

Figure 9: Storage costs of assisted backup

phases in the assisted backup cases due to the slow transfers toward peers; then decreases as more fragments are stored on remote peers. After a peer reaches its targets of eDLP and eTTR, further fragment transfers toward remote peers are carried out for cost-minimizing: storing more fragments on peers for free, and deleting fragments from the data center in exchange decreases the cost paid for the backup service at the data center. In the backup assisted systems more fragments are stored on the data center than in the systems where only the repairs are assisted: besides the transitional storage of repairs, peers, having difficulties with outsourcing their fragments to sufficiently reliable remote peers in the necessary number, keep some storage on the data center. The "price of fairness" of the grouped symmetric peer selection scheme therefore implies higher, although still reasonable, costs in these systems.

In the system design, network data transfers are programmed randomly towards remote peers. I studied this scheduling policy by analyzing theoretic models, and I have validated my analytic estimations and bonds on the length of fragment transfer periods by simulating practical settings.

Thesis 1.4. *I have proposed to use the efficient maximum flow algorithm to compute the optimal scheduling solution for hypothetical cases where future peer uptimes are known. I have used the optimal results to evaluate the applied scheduling policy, and I have proposed practical settings in which the performance of random decisions is close to optimal.*

During the backup and retrieval phases, a peer schedules upload and download transfers. The objective is to find the schedule that minimizes the latest completion time among the fragment transfers that satisfy the goals of the backup or the retrieval phase. In order to find the schedule that minimizes the required time to transfer N fragments (called as *mintime* problem), I determined the maximal number of transferable fragments in T time-slots (denoted as $maxfrag$ problem). Then, the solution for the original problem is the smallest T within which N fragments can be transferred.

Definition 4.4. ***The*** *mintime* ***and*** $maxfrag$ ***problems*** *are defined as follows: s stands for any scheduling, $t(s)$ gives the duration, and $n(s)$ provides the transferred fragments of a given schedule,*

$$mintime(N) = \min\{T \mid \exists s \ : \ t(s) = T \wedge n(s) \geq N\}; \tag{3}$$

$$maxfrag(T) = \max\{N \mid \exists s \ : \ t(s) \leq T \wedge n(s) = N\}. \tag{4}$$

The two problems are intertwined in the following way:

Proposition 4.5. *$mintime(N) = \min\{T \mid maxfrag(T) \geq N\}$.*

The following Integer Linear Programming (ILP) problem formulation is analogous to $maxfrag(T)$ of (4).

Definition 4.6. ***The scheduling problem with full knowledge*** *maximizes the number of transmitted fragments within a given duration T. x_i^t is a variable that encodes scheduling decisions: the number of fragments scheduled with remote peer i in time-slot t. The availability of remote peers is given as constraints: $a_i^t = 1$ if remote peer i is available in time-slot t, 0 otherwise. Another constraint is the maximal number of fragments m that can be placed on each remote peer. The solution of $maxfrag(T)$ can be found by solving the following ILP problem:*

$$\begin{array}{lll}
\max \sum_{t=0}^{T} \sum_{i=1}^{I} x_i^t & & \textit{maximize number of transmitted fragments} \\
\textit{s.t.} & x_i^t = [0, \min(u, d_i)] & \textit{transferable fragments in a time-slot to a peer} \\
& x_i^t \leq m a_i^t & \textit{transfer to online remote peers} \\
& \sum_{t=0}^{T} x_i^t \leq m & \textit{no more than } m \textit{ fragments on a remote peer} \\
& \sum_{i=1}^{I} x_i^t \leq u & \textit{no more than } u \textit{ transfers within a time-slot.}
\end{array}$$

In order to translate the above backup scheduling problem to the retrieval case, I write d instead of u, and m is replaced by the vector of stored fragment number on each remote peer.

I transformed the ILP problem into a *maximum flow* formalization. The same underlying problem thus becomes solvable in polynomial time. I present the maximum flow formulation of the $maxfrag(T)$ problem in Figure 10. Nodes *ts i* with $i = 1, 2, \ldots, T$ represent the time slots up to T. *peer i* with $i = 1, 2, \ldots, I$ depict the remote peers. *ts i* is connected to *peer j* if and only if j is online in time-slot i. Edges with capacity u give the constraints on the uplink capacity of the peer, edges with capacity m describe the maximum number of fragments to be stored on each remote peer. The maximum flow from the *source* to the *target* yields the largest number of fragments that can be uploaded within time T. Similarly to the ILP formulation, the retrieval problem formulation is attained if u is replaced by d, and m is replaced by the stored fragment number on each remote peer respectively.

I can now iteratively compute $maxfrag(T)$ with growing values of T; Proposition 4.5 guarantees that the first value T that satisfies $maxfrag(T) \geq N$ will be the desired result for our initial scheduling problem. The original problem, *i.e.*, finding an optimal schedule that minimizes the time to transfer N fragments, can be solved by performing $O(\log T)$ max-flow computations. For a flow network with V nodes and E edges, the maximum flow can be computed with time complexity $O\left(VE\log\left(\frac{V^2}{E}\right)\right)$ [16]. In our case, when we have I nodes and an optimal solution of T time-slots, V is $O(I+T)$ and E is $O(IT)$.

I simulated optimal and random scheduling of some $I-n$ settings (1000 runs, moreover the random scheduling scenario is performed 1000 times on each input). I plot the median of the optimal and random solutions for each $I-n$ case in Figure 11. As the number of fragments to transfer grows, the optimal solution gets closer to the theoretic lower

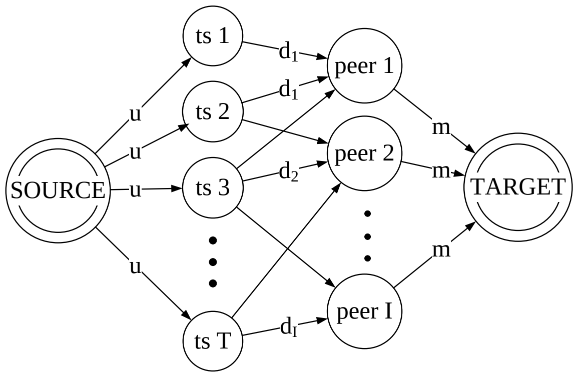

Figure 10: Maximum flow problem formulation of data transfer scheduling

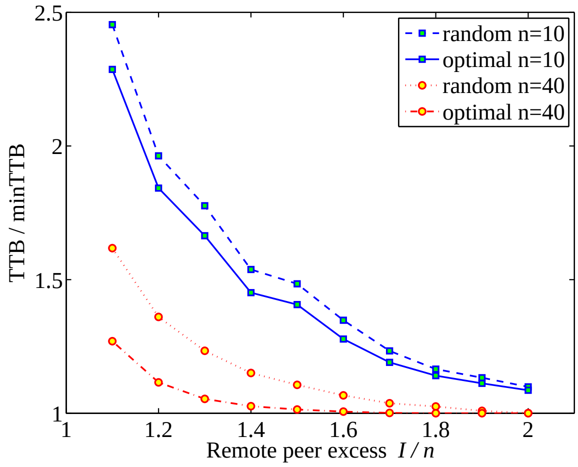

(a) Relative number of peers

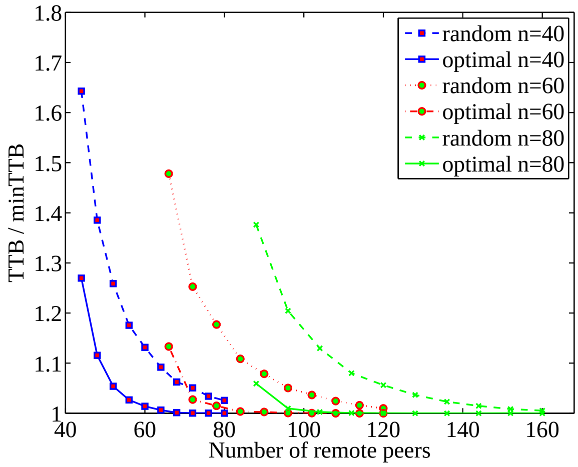

(b) Absolute number of peers

Figure 11: Backup and retrieval inefficiency with realistic peer online phases

bound, *i.e.*, $minTTB$. Moreover, with larger remote peer set, the performance of random scheduling approaches to that of the optimal, irrespective to the number of fragments to transfer. As heuristic thresholds, the values of $n = 60$ and $I = 90$ are sufficient to complete backup within a tolerable (around 10%) deviation from $minTTB$ with random scheduling.

4.2 User-Driven Peer Selection in a P2P Backup System

Thesis group 2: [J1, C1, C2, C3, C4] I have analyzed a model of peer selection in P2P backup systems in which users have the ability to selfishly select remote peers they want to exchange data with in a symmetric scheme.

In the symmetric system model users "selfishly" choose the peers, on which they store data. Storing partners are selected based on their characteristics (on-line availability and dedicated upload bandwidth) that are reflected through a single parameter, termed peer grades. In case of mutual willingness, two users offer the same amount of storage space to each other.

Definition 4.7. ***Peer selection model***

- *Let $\mathcal{I}$ denote the set of peers taking part in the system; $I = |\mathcal{I}| > 1$.*
- *Every peer $i \in \mathcal{I}$ splits each of its backup objects into k original fragments, out of which it creates $\hat{c}_i$ redundant fragments of the same size, to store on separate remote peers.*
- *Then every peer establishes a set of links, denoted by n_i for peer i: $n_i = \{j \in \{\mathcal{I} \backslash i\}\}$, where peer i stores a fragment on peer j. By policy, more than one fragment from a backup object cannot be stored on the same remote peer.*
- *The symmetric fragment exchange design dictates that $i \in n_j$ if $j \in n_i$,* i.e., *peers i and j are storage partners. Therefore peer i must share a local storage capacity equivalent to $|n_i|$ fragments, $|n_i| \leq \hat{c}_i$.*
- *$g_i = a_i u_i$ characterizes peer $i \in \mathcal{I}$, and it is called as the grade of i. $\hat{g}_i$ represents the* effortless *grade of i,* i.e., *the attributes composing it do not require additional strain from i other than its normal behavior.*

The peer selection model is in line with the presented P2P backup system design: symmetric fragment exchanges are performed with remote peers selected by their grades. The quality of service is analogous to the attributes of the partners: their number and their grades. If either of these values increases, the service quality improves too (up to a certain limit). The cost of the service consists of the grade improvement of a peer, if it chooses to increase its effortless resource contribution.

Thesis 2.1. *[J1, C1, C2, C3, C4] I have proposed a realistic, but analytically tractable game theoretic [27] payoff model for user "selfishness" (Definition 4.8).*

Definition 4.8. ***The payoff*** *that every peer maximizes when establishing links to remote peers, is composed by two terms, the value of the service and the effort cost:*

$$p_i = \min\left(|n_i|\underline{g}_i, 1\right) - (g_i - \hat{g}_i),$$

where $|n_i|$ denotes the number of partners of i and $\underline{g}_i = \min_{j \in n_i}(g_j)$ is the lowest grade among them.

Furthermore, I have defined a non-cooperative game, termed as exchange game, built on the payoff function to reflect the selfish context of user-driven peer selection.

Definition 4.9. ***The exchange game*** *is defined by the collection of player strategy sets $\{\mathcal{S}_i \ \forall i \in \mathcal{I}\}$, and the payoff function p that yields user payoffs $\{p_i \ \forall i \in \mathcal{I}\}$ on the combination of the strategy sets ($p : \mathcal{S}_1 \times \cdots \times \mathcal{S}_I \to \mathbb{R}^I$). A strategy of player $i \in \mathcal{I}$ consists of a grade $g_i \in (0, 1]$* and *a set of links n_i.*

The exchange game is in equilibrium, if every peer makes a strategy that yields the highest payoff, given the strategies of other peers. To determine these best response strategies, I dissect the joint optimization process that the strategic peers face: I analyze the grade and peer selection algorithmic problems separately. By decomposing the optimization problem of each player in the game to find its best response strategy, I assume that player decisions regarding the selection of their own grades and their remote peers are interleaved. For sake of tractability, first, I cast the peer selection as a stable fixtures problem [19].

Thesis 2.2. *[J1, C1, C2, C3, C4] I have proposed to reduce the exchange game with fixed grades to a matching problem. Based on the peer selection preferences induced by the payoff, I have shown that matches are created between peers with similar grades.*

The stable fixtures algorithm [19] solves every peer selection problem instance that may occur in the interleaved algorithm which finds the exchange game equilibrium. This is, in fact, a direct consequence of their particular attributes: player preferences are directly determined by their payoffs, indirectly by the grades of other peers. Therefore, during the stable fixtures algorithm match proposals from high grade peers are always reciprocated by lower grade peers.

The stratification phenomenon, *i.e.*, peers are linked respective to their grade order, leads to a matching where a higher grade peer has no poorer quality partners than a lower grade peer. As a result, low grade peers might realize less matches than their capacity. After formulating the algorithmic problem of peer selection, I investigated the grade selection of peers in the exchange game. I have proved that the stratification creates incentives for users with low grades to improve their contribution to the system. When selfish users are encouraged to raise their devoted contribution to the system, in terms of online availability and dedicated bandwidth, it essentially improves the overall quality of service offered by the system.

Thesis 2.3. *I have proven the existence of equilibrium in the exchange game, and I have given the best-response user strategies in respect to grade and remote peer selection.*

Proposition 4.10. *If $\min_{i \in \mathcal{I}} \hat{g}_i \geq \frac{1}{I-1}$, then a possible equilibrium is the effortless grand clique: everyone is linked to every other player. Generally, the best response grade strategy is to join a clique according to the effortless grade order rank. This might necessitate improvement on the effortless grade.*

Players unite in groups, driven by their payoffs based on the size and the worst grade member of the clique. If the group size becomes large, *i.e.*, many peers group together, the players might be better off excluding the ones with the worst grade. If the grade-segment becomes too narrow, *i.e.*, low grades are not eligible, the quality of service declines due to the critical number of member peers. This duality is due to the payoff: low grade clique members decrease the value at high grade members in the group, excluding them causes a drop in the value at the remaining players because of the reduced clique size. In equilibrium the two opposing effects are balanced inside disjoint groups.

In contrast to heterogeneous effortless grade settings, if every player has the same initial grade, they have no incentive to improve: if $\hat{g}_i = \hat{g}\ \forall i \in \mathcal{I}$, the best response strategy is $g_i = \hat{g}\ \forall i \in \mathcal{I}$.

4.3 Distributed Dynamic Spectrum Allocation

Thesis group 3: [J2, C5, C6] I have investigated the possibility of allocating radio spectrum among multiple applicants dynamically in a distributed manner.

The actual spectrum allocation policies, *i.e.*, governmental licenses for frequency bands sold for long-terms, are not efficient because peak traffic planning causes temporal underutilization in less busy periods, furthermore, the spatial and spectral restrictions on frequency re-usage due to rigid interference handling policies exclude many potential frequency exploitation opportunities. The emergence of novel radio technologies enables allocating spectrum bands for licensees with various spectral, spatial and temporal parameters, thus possibly improving spectrum utilization.

Definition 4.11. ***Dynamic spectrum allocation framework (presented in [20])***

- *The radio spectrum, denoted by F, is divided into small predefined sized, non-overlapping, homogeneous frequency slots, denoted by f.*
- *The possible frequency leasers ($\mathcal{I}$),* i.e., *the nodes, are distinguishable entities exploiting radio spectrum at fixed, and/or confined geographic locations,* i.e., *base stations of wireless service providers, private radio systems.*
- *Nodes require a given size frequency band ($q_i\ \forall i \in \mathcal{I}$ expressed in frequency slot units) and they are willing to pay a given amount of money (called utility, $q_i u_i\ \forall i \in \mathcal{I}$) for it.*
- *F_i denotes the frequency band allocated by node $i\ \forall i \in \mathcal{I}$; the size of a band F_i is $|F_i|$, and $\mathcal{F}^f$ is the set of nodes that allocate frequency slot f within their bands,* i.e., $\mathcal{F}^f = \{i : f \in F_i, \forall i \in \mathcal{I}\}$.
- *Interference may occur if the same frequency is used by other nodes than i, and it is defined as the* maximal *measured Signal to Interference-plus-Noise Ratio (SINR) value on the operating area of i and denoted and approximated by $\sum_{j\in\mathcal{I}} \omega_{ji}^f$, ω_{ji}^f being the interference caused by node j at node i on f.*
- *The maximal aggregate interference level node i can bear from the other nodes $j \neq i$ on frequency slot f by α_i^f.*

The interference level depends on many aspects: geographic distance, transmission power, applied technologies, coding, type of radio transmitters, etc. I model issues related to radio interference by the frequency-dependent relations ω and thresholds α. Nodes allocate frequency bands dynamically by paying fee to each other and to the authority if these frequency aspects make it necessary, *i.e.*, the co-existence of nodes on a given frequency slot is not feasible.

Definition 4.12. ***The allocation and pricing rules*** *ensure that sequentially arriving nodes may exclude others, if necessary due to unbearable interference constraints. At any exclusion the two interfering parties,* i.e., *the newcomer and the actual frequency leaser of the given frequency band, execute a second-price auction: both issue their bids, then the higher wins and pays the second bid. The possible outcomes are the following.*

- *Successful buy-out: if the bid of the newcomer node is higher, it pays the lower bid to the authority and the frequency leaser node is excluded.*
- *Successful defense: if the bid of the frequency leaser is higher, the authority does not impose any fee on it, and the attempting node is* implicitly *excluded.*

When node i attempts to exclude node j on some frequency slot f with bid b_i, if $b_i > b_j$, where b_j is the defense bid *of j, then j is excluded, and i pays b_j to the authority. Furthermore, $b_j \leq u_j - c_j^f$, where c_j^f stands for the sum of the expenses of j paid for prior exclusions of other, interfering nodes on frequency slot f. By paying b_j to the authority, c_i^f increases,* i.e., *$c_i^f := c_i^f + b_j$, which lowers the budget of i for further exclusion bids. An exclusion attempt is unsuccessful if $b_i < b_j$: neither i nor j pays the authority and i is implicitly excluded in this case.*

The distributed allocation and pricing make the system flexible in terms of possible allocation of spectrum at any time without a centrally announced or periodical auction. The pricing rule provides efficiency of spectrum utilization, *i.e.*, buyers with higher bids get the right to use the frequency spectrum. Furthermore, it ensures fairness despite the fact that exclusions are only unidirectional: those nodes that cause high interference are targeted with exclusion attempts, but nodes with high utility may get interference-free frequency while they cause high interference for other nodes.

Definition 4.13. ***Interference-friendliness*** *means high interference tolerance and low interference that is caused to other nodes. i is more interference-friendly than k, if $\alpha_i^f > \alpha_k^f$ and $\omega_{ij}^f < \omega_{kj}^f$ and $\omega_{ji}^f < \omega_{jk}^f$ $\forall j \in \mathcal{I} \setminus i, k, f \in F$.*

Thesis 3.1. *[J2, C5] I have shown that the allocation and pricing rules of Definition 4.12 ensure rationality (the payoff cannot be negative), incentive compatibility (nodes bid with their true utility u), and less interference-friendly nodes pay relatively more for the spectrum.*

Excluded nodes receive compensation from the authority or from the nodes that excluded them, therefore the payoff $p_i^f = u_i^f - c_i^f$ (c_i^f is the sum of second-prices paid in prior exclusions) of any node i on any frequency slot f cannot be negative. Furthermore, the well-known truthfulness property of second-price auctions makes bidders bid their true utilities as dominant strategies.

At a possible buyout, the actual owner node i and the interested parti(es) (let j be the one with the highest utility) are going to play a second-price auction. Let c_i^f denote the price that i paid for the frequency unit f, subject to the auction. None of the utilities among the nodes are known, however the utility of the actual owner has a lower bound:

its former winning price c_i^f. If j therefore decides to bid beyond its u_j with $u_j' > u_j$, and if $u_j' > u_i > u_j$, then the fee to pay to the authority results in a negative payoff.

Let two nodes i and k holding the same utility allocate a given frequency unit f at the same time, at the same geographic location. Let $\alpha_i^f > \alpha_k^f$ and $\sum_{j\in\mathcal{I}} \omega_{ij}^f < \sum_{j\in\mathcal{I}} \omega_{kj}^f$. If k gets f, and then i can also get f at a lower price: k is only disturbed by the interference of other nodes up to α_k^f, thus the quality of f would be sufficient for i as well. Since k paid the possible exclusions of other nodes in order to decrease the interference level below the threshold from its utility u_k, this could be also done by i, since $u_i = u_k$. Also, the interference that k causes to others is higher then i would cause, therefore i could resist to even less intensive re-allocation attempts of other nodes than k does.

The sequence of node "arrivals" has critical importance, but node characteristics partly predetermine the success of its spectrum allocation. However, selecting a frequency band to allocate and then the disturbing nodes to exclude is not straightforward.

Definition 4.14. ***Node roles*** *on frequency slot f are categorized into the following sets:*

- disturbing nodes*: the group of nodes whose exclusion by newcomer node i is necessary in order to assure that the cumulative interference on f is kept below α_i^f: $\mathcal{D}_i^f \subseteq \mathcal{F}^f$ such that $\sum_{j\in\mathcal{F}^f\setminus\mathcal{D}_i^f} \omega_{ji}^f \leq \alpha_i^f$;*

- excluding nodes*: a group of nodes from which exclusion attempts are expected because their perceived interference, increased by newcomer node i, is higher than their tolerance levels: $\mathcal{E}_i^f \subseteq \mathcal{F}^f \setminus \mathcal{D}_i^f$ such that for $\mathcal{C}_i^f = \mathcal{F}^f \setminus \mathcal{D}_i^f \setminus \mathcal{E}_i^f$ (called as the set of* coexistent nodes*) $\forall k \in \mathcal{C}_i^f \sum_{j\in\mathcal{C}_i^f} \omega_{jk}^f \leq \alpha_k^f$ and $\forall k \in \mathcal{E}_i^f \sum_{j\in\mathcal{C}_i^f} \omega_{jk}^f > \alpha_k^f$.*

Thesis 3.2. *[C6] I have shown that optimizing the set of disturbing nodes to be excluded is an NP-complete problem. I have suggested heuristic node exclusion strategies.*

The problem is formulated as follows. For given $\{\omega_{ji}^f, u_j\}$ $\forall j \in \mathcal{I}$ and for α_i^f, u, is there any $\mathcal{D}_i^f \subseteq \mathcal{I}$ which satisfies $\sum_{j\in\mathcal{D}_i^f} \omega_{ji}^f \geq \sum_{j\in\mathcal{I}} \omega_{ji}^f - \alpha_i^f$ and $\sum_{j\in\mathcal{D}_i^f} u_j \leq u$? This is equivalent to the knapsack problem, which is known to be NP-complete [15].

Hindered by the difficulty of the problem, exclusion of nodes is carried out in the increasing order of their "interference prices", *i.e.*, for newcomer node i: $\forall j \in \mathcal{D}_i^f$ and $\forall k \notin \mathcal{D}_i^f$ $\frac{u_j - c_j^f}{\omega_{ji}^f} \leq \frac{u_k - c_k^f}{\omega_{ki}^f}$. Furthermore, exclusion attempts, targeting the newcomer, should arrive from nodes in the increasing order of their interference thresholds, *i.e.*, $\forall j \in \mathcal{E}_i^f$ and $\forall k \notin \mathcal{E}_i^f$ $\alpha_j \leq \alpha_k$. In case of equality, by policy, those nodes are included in the sets that cause higher interference or arrived sooner respectively.

Newcomer selfish nodes strive to allocate the required size frequency band by spending the minimum on occurring costs for the maximal expected lifetime. Each node i may perform exclusions up to its budget of u_i on each of its frequency slots, but one should maintain sufficient defense bid against exclusion attempts. At exclusion, in the worst case node i needs to set its bid b_i at the target node j to u_j. This is the case when node j has not excluded other nodes earlier. As the number of nodes that should be excluded in order to fulfill the interference requirements of i grows, its competitiveness worsens, as $u_i - c_i$ falls.

In order to find the cheapest frequency band, newcomer i tries to position its F_i on the spectrum, so that, on the one hand, the cost of exclusion of other nodes ($\mathcal{D}_i^f$) would be minimal; on the other hand the cost of defense against excluding nodes ($\mathcal{E}_i^f$) would cost the least possible on the average of all the frequency slots of F_i.

Thesis 3.3. *[C6] I have given a necessary condition for successful allocation. Node i is able to allocate an adequate spectrum band if $\exists F_i : |F_i| = q_i$ and $\forall f \in F_i$:*

$$u_i \geq \sum_{j \in \mathcal{D}_i^{f*}} \left(u_j - c_j^f\right) + \max_{j \in \mathcal{E}_i^{f*}} \left(u_j - c_j^f\right), \textit{ where}$$

$$\mathcal{D}_i^{f*} = \arg\min_{\mathcal{D}_i^f} \sum_{j \in \mathcal{D}_i^f} \left(u_j - c_j^f\right) \textit{ and } \mathcal{E}_i^{f*} = \arg\min_{\mathcal{E}_i^f} \max_{j \in \mathcal{E}_i^f} \left(u_j - c_j^f\right),$$

provided that the node exclusion policy requires the newcomer to exclude first the interfering nodes, then the old nodes may make attempts to exclude the newcomer if needed.

I have suggested a heuristic algorithm (in Algorithm 1) with various frequency band selection strategies:

- *cost-minimizing: finds the lowest expected cost based on my node selection heuristics;*
- *interference-aware: takes into account* only *the* actual *aggregate interference on each frequency band;*
- *deliberate: considers both the aggregate interference* and *the remaining budget of the other nodes.*

Algorithm 1 Implemented heuristic algorithm for spectrum allocation

newcomer node i: q_i, u_i, α_i^f
for all $F_i \in F$ such that $|F_i| = q_i$ **do**
    for all $f \in F_i$ **do**
        $\mathcal{D}_i^{f*} = \{j : \sum_{k \notin \mathcal{D}_i^{f*}} \omega_{ki}^f \leq \alpha_i^f, \frac{u_j - c_j^f}{\omega_{ji}^f} \leq \frac{u_k - c_k^f}{\omega_{ki}^f} \forall k \notin \mathcal{D}_i^{f*}\}$
        $\mathcal{E}_i^{f*} = \{j : \sum_{k \notin \mathcal{E}_i^{f*}} \omega_{kj}^f > \alpha_j^f, \alpha_j^f \leq \alpha_k^f \forall k \notin \mathcal{E}_i^{f*}\}$
    end for
end for
if cost-minimizing **then**
    $F_i^* = \arg\min_{F_i} \sum_{f \in F_i} \left(\sum_{j \in \mathcal{D}_i^{f*}} \left(u_j - c_j^f\right) + \max_{j \in \mathcal{E}_i^{f*}} \left(u_j - c_j^f\right)\right)$
else if interference-aware **then**
    $F_i^* = \arg\min_{F_i} \sum_{f \in F_i} \sum_{j \in \mathcal{F}^f} \omega_{ji}^f$
else if deliberate **then**
    $F_i^* = \arg\min_{F_i} \sum_{f \in F_i} \sum_{j \in \mathcal{F}^f} (\omega_{ji}^f - (u_j - c_j))$
end if
if $\forall f \in F_i^*$ $u_i \geq \sum_{j \in \mathcal{D}_i^{f*}} \left(u_j - c_j^f\right) + \max_{j \in \mathcal{E}_i^{f*}} \left(u_j - c_j^f\right)$ **then**
    successful allocation: i buys out nodes in $\mathcal{D}_i^{f*}$, and implicitly excludes nodes in $\mathcal{E}_i^{f*}$ $\forall f \in F_i^*$
else
    unsuccessful allocation: i buys out subset of nodes in $\mathcal{D}_i^{f*}$, and implicitly excludes subset of nodes in $\mathcal{E}_i^{f*}$ $\forall f \in F_i^*$
end if

I evaluated numerically the proposed framework, and compared the outcomes of applying these heuristics in a toy example (Figure 12). The income of the authority is less and less if nodes switch to interference-aware and deliberate frequency band selection strategies from the cost-minimizing selection. This is due to the failure of the allocation attempts of type-4 (DVB-T-like) nodes. Interestingly, while with interference-aware strategy type-1 nodes (GSM-like) may expect longer lifetime (Figure 12(d)), the difference is less observable in Figure 12(f) despite the fact that type-4 nodes are even more punished in the latter case. Since tpye-2 (UMTS-like) and type-3 (UWB-like) nodes experience similar lifetimes irrespective to the applied frequency band selection, the interference-aware strategy proposes a fair trade-off among the income and the allocation success of type-1 and type-3 nodes in this setting.

Based on the simplistic simulation scenarios, the distributed dynamic spectrum allocation framework proved to be a suitable approach to efficient and flexible spectrum utilization.

5 Application of the results

The findings of my research on P2P backup systems consist of both theoretic and practical results. The presented combination of matching-, and game theoretic models provides new insights on existing problems in distributed systems. The defined performance metrics and the novel data redundancy calculation approach may affect other P2P storage system designs, not necessarily only for backup purposes. In addition, my implemented prototype can provide a basis for developing practical backup applications. These latter can be commercialized and deployed shortly on end-user devices, on set-top-boxes of subscribers of Internet service providers, or on corporate networks of companies.

The analysis of spectrum allocation gives insights about the potential future radio frequency management schemes. I have shown a distributed, incentive-compatible model with goals of ensuring fairness and efficiency on spectrum utilization, although I highlighted the algorithmic complexity of exclusions and band selection. The allocation and pricing scheme is a well-engineered environment for selfish participants, provides temporal flexibility and capability for fast, local response to any variation in the frequency demand.

Acknowledgments

The work presented in Sections 4.1 and 4.2 was done at Eurecom. I wish to thank Pietro Michiardi and Matteo Dell'Amico for their help. The research demonstrated in Section 4.3 was carried out in High Speed Network Laboratory at Budapest University of Technology and Economics. I am thankful for the guidance provided by Attila Vidács.

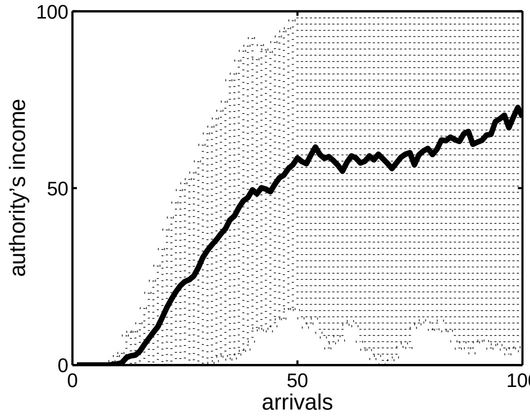

(a) Cost minimizing frequency band selection

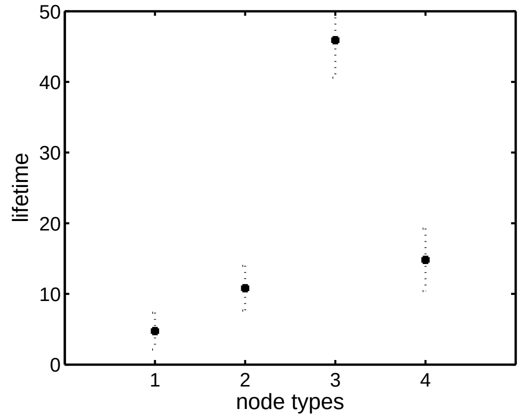

(b) Cost minimizing frequency band selection

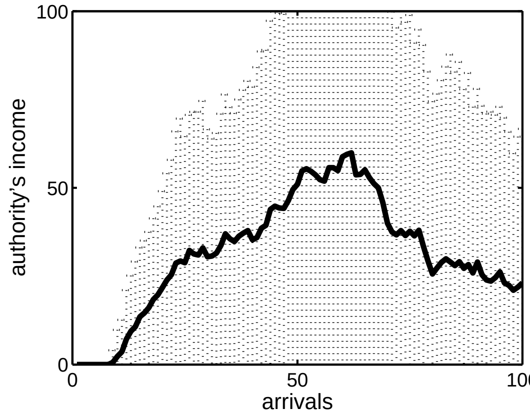

(c) Interference-aware frequency band selection

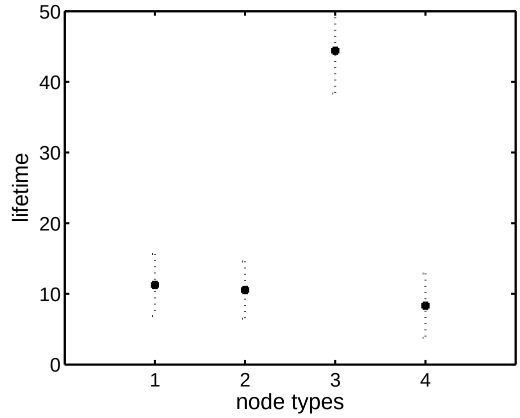

(d) Interference-aware frequency band selection

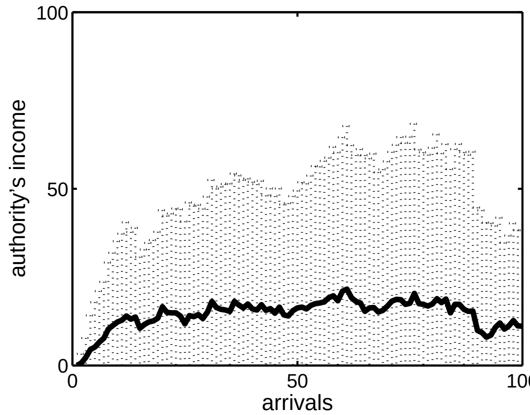

(e) Deliberate frequency band selection

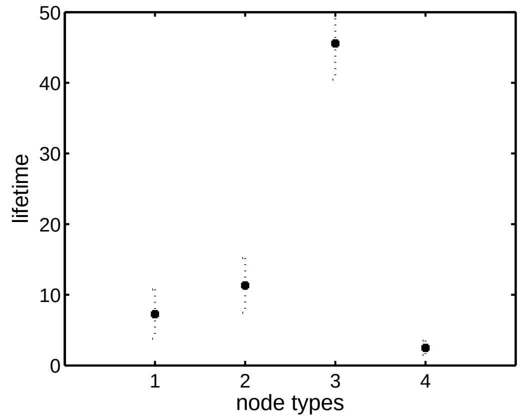

(f) Deliberate frequency band selection

Figure 12: Authority income and node lifetimes with different frequency band selection strategies

References

[1] E. Altman, A. A. Kherani, P. Michiardi, and R. Molva. Non-cooperative forwarding in ad-hoc networks. In *IFIP NETWORKING*, 2005.

[2] K. Anagnostakis and M. Greenwald. Exchange-based incentive mechanisms for peer-to-peer file sharing. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2004.

[3] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Computer Networks*, 46(1):133–146, 2004.

[4] L. Buttyán and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *Mobile Networks and Applications*, 8(5):579–592, 2003.

[5] B. Cohen. Incentives build robustness in bittorrent. In *Workshop on the Economics of Peer-to-Peer Systems (NetEcon)*, 2003.

[6] J. Corbo and D. Parkes. The price of selfish behavior in bilateral network formation. In *ACM Symposium on Principles of Distributed Computing (PODC)*, 2005.

[7] C. Courcoubetis and R. Weber. Incentives for large peer-to-peer systems. *IEEE Journal on Selected Areas in Communications*, 24(5):1034 – 1050, may. 2006.

[8] L. Cox and B. Noble. Samsara: Honor among thieves in peer-to-peer storage. In *ACM Symposium on Operating Systems Principles (SOSP)*, 2003.

[9] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *ACM Symposium on Principles of Distributed Computing (PODC)*, 2003.

[10] J. Feigenbaum, V. Ramachandran, and M. Schapira. Incentive-compatible interdomain routing. In *ACM Conference on Electronic Commerce (EC)*, 2006.

[11] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *ACM Conference on Electronic Commerce (EC)*, 2004.

[12] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[13] A. Gai, F. Mathieu, F. de Montgolfier, and J. Reynier. Stratification in p2p networks: Application to BitTorrent. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2007.

[14] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[15] M. R. Garey and D. S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1990.

[16] A. V. Goldberg and R. E. Tarjan. A new approach to the maximum flow problem. In *ACM Symposium on Theory of Computing (STOC)*, 1986.

[17] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.

[18] R. W. Irving. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms*, 6(4):577–595, 1985.

[19] R. W. Irving and S. Scott. The stable fixtures problem - a many-to-many extension of stable roommates. *Discrete Applied Mathematics*, 155(17):2118–2129, 2007.

[20] L. Kovács and A. Vidács. Interference-tolerant spatio-temporal dynamic spectrum allocation. In *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2007.

[21] N. Laoutaris, G. Smaragdakis, A. Bestavros, and J. W. Byers. Implications of selfish neighbor selection in overlay networks. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2007.

[22] N. Laoutaris, G. Smaragdakis, K. Oikonomou, I. Stavrakakis, and A. Bestavros. Distributed placement of service facilities in large-scale networks. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2007.

[23] N. Laoutaris, G. Zervas, A. Bestavros, and G. Kollios. The cache inference problem and its application to content and request routing. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2007.

[24] D. Lebedev, F. Mathieu, L. Viennot, A.-T. Gai, J. Reynier, and F. de Montgolfier. On using matching theory to understand p2p network design. In *International Network Optimization Conference*, 2007.

[25] P. Maillé and B. Tuffin. Pricing the internet with multibid auctions. *IEEE/ACM Transactions on Networking*, 14(5):992–1004, 2006.

[26] P. Maillé and B. Tuffin. Analysis of price competition in a slotted resource allocation game. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2008.

[27] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, July 1994.

[28] J. M. Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

[29] V. Vishnumurthy, S. Chandrakumar, and E. Sirer. KARMA: A secure economic framework for peer-to-peer resource sharing. In *Workshop on the Economics of Peer-to-Peer Systems (NetEcon)*, 2003.

[30] H. Weatherspoon and J. D. Kubiatowicz. Erasure coding vs. replication: A quantitative comparison. In *International Workshop on Peer-To-Peer Systems (IPTPS)*, 2002.

[31] S. Zhong, J. Chen, and Y. Yang. Sprite: a simple, cheat-proof, credit-based system for mobile ad-hoc networks. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2003.

Own publications related to the thesis

Journal papers

[J1] **L. Toka**, P. Michiardi. Analysis of User-driven Peer Selection in Peer-to-Peer Backup and Storage Systems. *Telecommunication Systems, Special Issue dedicated to GameComm'08, appeared online*, 2010.

[J2] **L. Toka**, A. Vidacs. General distributed economic framework for dynamic spectrum allocation. *Computer Communications, Special Issue dedicated to Spectrum Sharing Systems, 39(18):1955–1964*, 2009.

Conference papers

[C1] **L. Toka**, P. Michiardi. Brief Announcement: A dynamic exchange game. In *ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, 2008.

[C2] **L. Toka**, P. Michiardi. Analysis of user-driven peer selection in peer-to-peer backup and storage systems. In *International Workshop on Game Theory in Communication Networks*, 2008.

[C3] **L. Toka**, P. Michiardi. Uncoordinated peer selection in P2P backup and storage applications. In *IEEE Global Internet Symposium*, 2009.

[C4] P. Michiardi, **L. Toka**. Selfish Neighbor Selection in Peer-to-Peer Backup and Storage Applications. In *International Conference on Parallel and Distributed Computing (Euro-Par)*, 2009.

[C5] **L. Toka**, A. Vidacs. Distributed Dynamic Spectrum Management. In *International Conference on Telecommunications and Signal Processing*, 2009.

[C6] **L. Toka**, A. Korosi, A. Vidacs. On Distributed Dynamic Spectrum Allocation for Sequential Arrivals. In *IEEE Symposia on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2010.

[C7] **L. Toka**, M. Dell'Amico, P. Michiardi. Online Data Backup: a Peer-Assisted Approach. In *IEEE International Conference on Peer-to-Peer Computing (P2P)*, 2010.

[C8] A. Csoma, **L. Toka**, A. Vidacs. A Peer-to-Peer Backup System with Incentives. In *International Conference on Telecommunications and Signal Processing*, 2010.

Other own publications

Journal papers

[J3] I. Ráncsik, **L. Toka**, A. Vidács. Auction-Based Pricing Mechanisms in IP Networks. *HÍRADÁSTECHNIKA*, 60(5):27–34, 2005.

[J4] P. Maillé, **L. Toka**. Managing a peer-to-peer data storage system in a selfish society. *IEEE Journal on Selected Areas in Communications*, 26(7):1295 –1301, 2008.

[J5] **L. Toka**, A. Vidács. Peer-to-peer adattároló rendszer menedzselése önző társadalomban. *HÍRADÁSTECHNIKA*, 63(2):31–36, 2008.

[J6] **L. Toka**, A. Vidács. Managing Users in a Peer-to-Peer Storage System. *HÍRADÁSTECHNIKA*, 63(7):4–8, 2008.

[J7] G. Biczók, **L. Toka**, A. Gulyás, T. A. Trinh, A. Vidács. Incentivizing the Global Wireless Village. *COMPUTER NETWORKS*, elfogadva, 2010.

[J8] **L. Toka**, L. Kovács, A. Vidács. General distributed economic framework for dynamic spectrum allocation. *HÍRADÁSTECHNIKA*, 65(1):20–24, 2010.

Conference papers

[C9] **L. Toka**, P. Maillé. Managing a peer-to-peer backup system: Does imposed fairness socially outperforma revenue-driven monopoly? In *International Workshop on Grid Economics and Business Models*, 2007.

[C10] G. Biczók, **L. Toka**, A. Vidács, T. A. Trinh. On Incentives in Global Wireless Communities. In *ACM Workshop on User-Provided Networking*, 2009.