

# Incentives in Peer-to-Peer Systems

---

Moshe Babaioff, John Chuang, and Michal Feldman

## Abstract

Peer-to-peer (p2p) systems support many diverse applications, ranging from file-sharing and distributed computation to overlay routing in support of anonymity, resiliency, and scalable multimedia streaming. Yet, they all share the same basic premise of voluntary resource contribution by the participating peers. Thus, the proper design of incentives is essential to induce cooperative behavior by the peers. With the increasing prevalence of p2p systems, we have not only concrete evidence of strategic behavior in large-scale distributed systems but also a live laboratory to validate potential solutions with real user populations. In this chapter we consider theoretical and practical incentive mechanisms, based on reputation, barter, and currency, to facilitate peer cooperation, as well as mechanisms based on contracts to overcome the problem of hidden actions.

## 23.1 Introduction

The public release of Napster in June 1999 and Gnutella in March 2000 introduced the world to the disruptive power of peer-to-peer (p2p) networking. Tens of millions of individuals spread across the world could now self-organize and collaborate in the dissemination and sharing of music and other content, legal or otherwise. Yet, within 6 months of its public release, and long before individual users are threatened by copyright infringement lawsuits, the Gnutella network saw two thirds of its users free-riding, i.e., downloading files from the network without uploading any in return.

Given the large-scale, high-turnover, and relative anonymity of the p2p file-sharing networks, most p2p transactions are one-shot interactions between strangers that will never meet again in the future. It is therefore unsurprising that cooperation is difficult to sustain in these networks. The problem is exacerbated by hidden action due to nondetectable defections, and by the ability of peers to create multiple identities at no cost. It quickly became clear to the p2p developers community that some form of incentives is needed to overcome this free-riding problem.

The subsequent generation of p2p file-sharing networks incorporated incentive mechanisms based on *currency* or *reputation*. For example, in Mojonation, peers earn *mojos* through contributions to others, and use the earned currency to redeem for service from others. In KaZaA, peers build up their reputation scores by uploading, and highly reputed peers receive preferential treatment in their downloads.

The BitTorrent file-sharing system went beyond currency and reputation, and adopted an incentive mechanism based on *barter*. By partitioning large files such as movies and software binaries into small chunks, file-sharing using the BitTorrent protocol necessitates repeat interactions among peers, allowing cooperation to flourish based on *direct reciprocity* rather than *indirect reciprocity*. From a system perspective, there is no need to keep long-term state information, in the form of either reputation or currency. This simplifies the design and improves its robustness against attacks. Empirical studies found much lower levels of free-riding in BitTorrent communities. Yet, theoretical analysis has demonstrated that the BitTorrent protocol can still be manipulated by selfish peers in their favor.

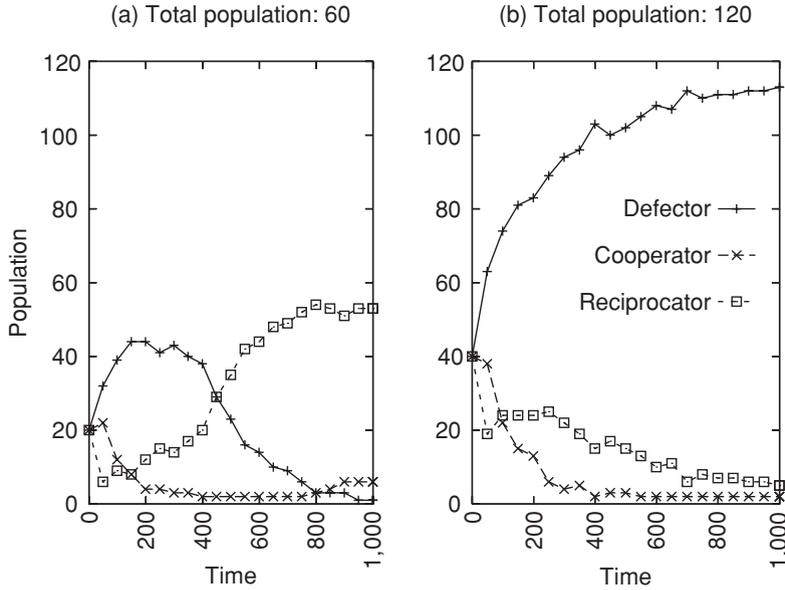
The issue of incentives in p2p systems goes far beyond free-riding in file-sharing networks. Grassroots contribution by autonomous peers are needed to sustain many networked systems, ranging from mobile ad hoc networks and community-based wireless mesh networks, to application layer overlay networks that support anonymous communications and live video streaming. Even interdomain routing over the Internet requires the cooperation of competing network operators.

The strategy space is also far richer than the binary choice of share/not-share in file-sharing networks. Peers make strategic decisions concerning the revelation of *private information*, such as local resource availability, workload, contribution cost, or willingness-to-pay. Peers decide on the amount of exerted effort, given the nonobservability of their *hidden actions*. Peers may adjust their spatial engagement with the network through *strategic network formation*, and temporal engagement through *strategic churning* (arrivals and departures). Finally, peers may choose to manage their own identities and treat the identities of others differently given the availability of cheap pseudonyms.

The increasing prevalence of p2p systems, coupled with the rich strategy space available to the peers, make the problem of *p2p mechanism design* a challenging and broadly relevant topic of study for algorithmic game theory. P2P systems offer a concrete example of strategic behavior in large-scale distributed systems, as well as a live laboratory to validate potential solutions with real user populations. In this chapter, we discuss some p2p incentive mechanisms based on reputation, barter, and currency, as well as mechanisms to overcome the problem of hidden actions. We refer readers to other chapters in this book on the related topics of distributed algorithmic mechanism design (Chapter 14), strategic network formation (Chapter 19), network pricing (Chapter 22), and reputation systems (Chapter 27).

## 23.2 The p2p File-Sharing Game

A p2p file-sharing system seeks to support efficient and scalable distribution of files by leveraging the upload bandwidth of the downloading peers. In a p2p file-sharing



**Figure 23.1.** The temporal evolution of strategy populations in a p2p file-sharing game. “Time” is the number of elapsed rounds. “Population” is the number of players using a strategy.

system, a peer plays one of two roles. For certain interactions, he is a client who wishes to download a file, and derives benefit from a successful download. For other interactions, he is a server who is requested to upload part or all of a file, and if he agrees he may bear some cost in the form of bandwidth and CPU usage. In such a one-shot game, “free-riding” is a dominant strategy – a player will download when he is a client, and refuse to upload when he is a server.

The interaction between players in a p2p file-sharing system has many characteristics of the Prisoner’s Dilemma (PD) game. In the single-shot PD game, players have a dominant strategy to defect, which leads to a socially undesirable equilibrium outcome known as the “tragedy of the commons.” In the Iterated Prisoner’s Dilemma game, cooperation can be sustained through direct reciprocity (e.g., using the Tit-for-Tat or TFT strategy) since a defection in the current round can lead to retaliation by the other player in a future round. This “shadow of the future” can similarly sustain cooperation in the p2p file-sharing game, where a peer may decide to upload a file to another peer with the expectation that he may wish to download a file from the other peer sometime in the future.

Of course, there is no guarantee that two peers will engage in multiple transactions with each other in their lifetimes. Even if they do, there is no guarantee that they will do so with a proper reversal of client and server roles to facilitate reciprocity or retaliation. In a large dynamic population with random matching of players, the probability of repeat interactions between players may be too small to cast an effective “shadow of the future,” and free-riding might prevail.

Figure 23.1, taken from a simulation study of a p2p file-sharing game (Feldman et al., 2004), illustrates the inability of a reciprocative strategy to scale to large populations.

Starting with equal shares of players that (1) always defect, (2) always cooperate, and (3) play a reciprocative strategy (a generalization of TFT for interleaved interactions with multiple peers), the game proceeds in rounds where the size of the population that plays each strategy is proportional to its success in the previous round. We see in Figure 23.1(a) that with a relatively small population, the reciprocative strategy dominates the population after 1,000 rounds. However, the strategy does not scale to larger populations, as seen in Figure 23.1(b), since the interactions between pairs of players are not frequent enough to make the strategy effective against defectors.

This suggests that strategies based on the notion of direct reciprocity may not fit the environment of p2p systems with random matching and large populations. One way to overcome this is to enforce repeated interactions with a small number of peers, as is done in BitTorrent (discussed in further detail in Section 23.4). This design works well for the sharing of large and popular files, e.g., movies and software binaries, since there are large numbers of peers who are concurrently interested in a file, and are willing to engage in repeated interactions to exchange file segments with one another.

To support cooperation over multiple files and longer timescales, some form of information sharing among the peers may be needed. This marks a shift from direct reciprocity to indirect reciprocity. Reputation systems (discussed in Section 23.3) provide a means for a peer to condition his action against his opponent upon the opponent's past actions, not just against the peer himself, but against other peers in the system. This way, a peer may choose to serve a file to another peer on the grounds that the latter had cooperated with other peers in earlier interactions.

Because p2p systems are large, dynamic systems with high turnover rates, peers often interact with strangers with no prior history or reputation. It is therefore very important to think about how one deals with strangers. A tit-for-tat strategy that always cooperates with strangers may encourage newcomers to join the system, but it can be easily exploited by whitewashers who leave and rejoin the system with new identities. The problem arises because a whitewasher is indistinguishable from a legitimate newcomer. Always defecting against strangers is robust against whitewashers, but it discourages newcomers and may also initiate unfavorable cycles of defection. It has been shown that cooperating with strangers with a fixed probability  $0 < p < 1$  is not robust against whitewashers. On the other hand, adapting the probability of cooperation with strangers to the frequency of past cooperation by strangers appears to be effective against whitewashers, at least for a sufficiently small turnover rate.

In the next three sections, we will discuss incentive mechanisms for p2p systems based on reputation, barter, and currency.

### 23.3 Reputation

Reputation has an excellent track record at facilitating cooperation in very diverse settings, from evolutionary biology to online marketplaces like eBay. It is therefore unsurprising that many p2p systems have adopted some form of reputation scheme to reward good behavior and/or punish bad behavior by the peers.

In general, a p2p reputation scheme is coupled with a service differentiation scheme. Contributing peers possess good reputations and receive good service from other peers,

while noncontributing peers possess bad reputations and receive poor service from others. For example, peers in the KaZaA file-sharing network build up their reputation scores by uploading files to others, and are rewarded with higher priority when downloading files from others. Similar schemes have been proposed for p2p storage, p2p multicast, and mobile ad hoc networks.

Used in conjunction with other security techniques, a p2p reputation scheme can also be used to identify, isolate, and avoid malicious peers in a system. For example, the Eigentrust algorithm computes global trust values of peers by aggregating local trust values based on the notion of transitive trust, similar to the PageRank algorithm. Peers that introduce inauthentic files into the system receive a low global trust value and will be shunned by others. The Credence system extends the notion of reputation from peers to objects. Reputation scores are maintained for individual objects in the p2p system. These techniques can be used to defend against pollution and poisoning attacks in p2p file-sharing networks.

Reputation systems may be subject to a number of different attacks. Multiple colluding peers may boost one another's reputation scores by giving false praise, or punish a target peer by giving false accusations. The availability of *cheap pseudonyms* in p2p systems make reputation systems vulnerable to *Sybil attacks* and *whitewashing attacks*. In a Sybil attack, a single malicious peer generates multiple identities that collude with one another. In a whitewashing attack, a peer defects in every p2p transaction, but repeatedly leaves and rejoins the p2p system using newly created identities, so that it will never suffer the negative consequences of a bad reputation.

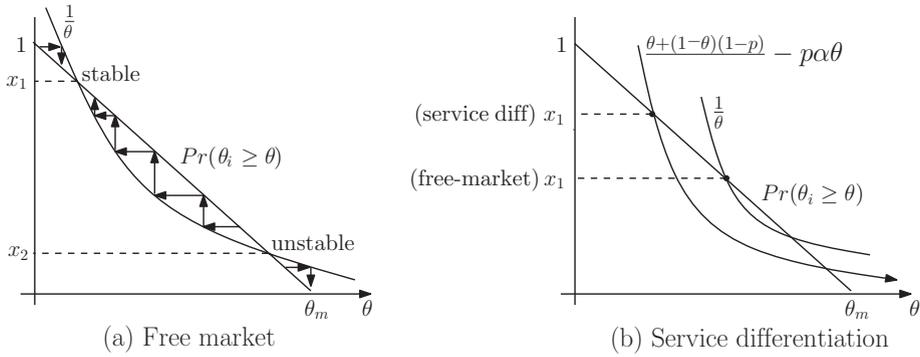
A comprehensive treatment of the design and implementation of reputation systems is provided in a separate chapter of this book. So we will focus our attention to the use of reputation and service differentiation schemes in establishing cooperation in p2p systems. In particular, we will construct a minimalistic model of a p2p system (in Section 23.3.1) to explore its dynamics and resulting equilibria in the absence of any reputation scheme, and see (in Section 23.3.2) how a reputation and service differentiation scheme can improve the performance of the system.

### 23.3.1 A Minimalist p2p Model

Consider a population of rational peers with heterogeneous willingness to contribute resources to the system. Each peer  $i$  has a type  $\theta_i$ , reflecting his generosity or the maximum cost he is willing to incur in contribution. Each peer makes autonomous decisions whether to contribute or free-ride based on the relationship between the cost of contribution and her type. Since contributors have to carry the load of the system, the contribution cost can be modeled as inversely proportional to the fraction of contributors in the system. Thus, if at present a fraction  $x$  of the peers are contributing, the contribution cost is  $1/x$ , and therefore the decision of a rational peer with type  $\theta_i$  is:

Contribute, if  $\theta_i > 1/x$ ;  
Free-ride, otherwise.

Even within this simple framework we can already see some interesting implications. In this “free market” environment where no incentive mechanism is in place,



**Figure 23.2.** (a) The intersection points of the type distribution and cost curves represent two equilibria of the system. The curve  $1/\theta$  represents the contribution cost, and  $Pr(\theta_i \geq \theta)$  represents the generosity CDF, assuming  $\theta_i \sim U(0, \theta_m)$ . The higher equilibrium (contribution level  $x_1$ ) is stable. The point  $x = 0$  is an additional equilibrium of the system. (b) Under the service differentiation mechanism, the cost curve shifts from  $1/\theta$  to  $\frac{\theta+(1-\theta)(1-p)}{\theta} - p\alpha\theta$ . Consequently, the attractor ( $x_1$ ) shifts upward.

the contribution level  $x$  in equilibrium is determined as the intersection of the type distribution,  $x = Pr(\theta_i \geq \theta)$  with the curve  $x = 1/\theta$ .

Figure 23.2 shows the equilibria when the generosity type is uniformly distributed between 0 and some maximal value  $\theta_m$ . There are three equilibria in this system. The first two are the intersection points of the type distribution curve and the cost curve. The third equilibrium is  $x = 0$ , which always exists. Consider the natural fix-point dynamics of the system, i.e., starting at some initial  $x$ , peers arrive at individual decisions, their aggregate decisions define a new  $x$ , which leads to a new aggregate decision, and so on. When the system is out of equilibrium, the direction in which the system moves depends on the relative heights of the two curves. If the cost curve is above the type distribution curve, contribution cost is higher than the fraction of users who are willing to contribute at this cost, so the fraction of contributors decreases. For example, in Figure 23.2, this happens for  $x < x_2$  or  $x > x_1$ . Conversely, for  $x_1 < x < x_2$ , the contribution cost is lower than the willingness to contribute, so contribution level increases. Therefore,  $x = x_1$  and  $x = 0$  are the two attractors of the fixpoint dynamics. As long as the initial  $x$  lies above the lower intersection point ( $x_2$ ), the process converges to the upper one ( $x_1$ ). Otherwise, if the initial  $x$  is below the lower intersection point, or if there is no intersection; i.e., when there are too many selfish rascals around, then  $x$  converges to 0 and the system collapses.

The contribution level of the system,  $x$ , is derived by solving the fixpoint equation:  $x = Prob(\theta_i \geq 1/x)$ . If we consider the case in which the generosity of the peers is uniformly distributed between 0 and  $\theta_m$ , i.e.,  $\theta_i \sim U(0, \theta_m)$ , then  $Prob(\theta_i \geq 1/x) = 1 - \frac{1}{x\theta_m}$ , and the fixpoint equation is  $x = 1 - \frac{1}{x\theta_m}$ . The solutions are  $x_{1,2} = \frac{\theta_m \pm \sqrt{\theta_m^2 - 4\theta_m}}{2\theta_m}$ . The larger root  $x_1$  is a stable equilibrium while  $x_2$  is not.  $\theta_m$  denotes the maximal willingness to contribute resources, and reflects the overall generosity level of the system.

**Claim 23.1** *The stable nonzero equilibrium contribution level ( $x_1$ ) increases in  $\theta_m$  and converges to 1 as  $\theta_m$  goes to  $\infty$ , but falls to zero when  $\theta_m < 4$ .*

So far we have been interested only in costs. To understand system performance, we need to consider system benefits as well. We assume that the benefit a peer receives from participation in the system (whether or not she contributes) is proportional to the contribution level in the system, and thus a function of the form  $\alpha x$  for some constant  $\alpha \geq 1$ . We concentrate on cases where  $\alpha$  is large, in which  $x = 0$  is socially inefficient.

We define the performance of the system,  $W_S$ , as the difference between the total benefits received by all peers and the total contribution cost incurred by all peers (noting that free riders incur no costs). Normalizing network size to 1, for  $x > 0$  we have

$$W_S = \alpha x - (1/x)x = \alpha x - 1.$$

According to the definition of system performance and Claim 23.1, even if participation can provide high benefits to the peers, the system will still collapse if the maximal generosity is low, since the system performance is limited by the low contribution level. In the next section, we see how a reputation and service differentiation scheme can overcome this problem.

### 23.3.2 Reputation and Service Differentiation

Now let us introduce an incentive mechanism based upon reputation and service differentiation. Consider a reputation system that can catch free riders with probability  $p$ , and a service differentiation policy where identified free riders are excluded from the system. An alternate interpretation is a reputation system that can perfectly distinguish free riders and contributors, used in conjunction with a service differentiation policy where free riders are penalized with a reduced level of service of  $1 - p$  times that of a contributor.

Degrading the performance of the free riders has two effects, both of which lead to a higher contribution level. First, since free riders get only a fraction  $1 - p$  of the benefits, the load placed on the system decreases to  $x + (1 - x)(1 - p)$ . Therefore, contribution cost becomes  $\frac{x+(1-x)(1-p)}{x}$ . Second, the penalty introduces a threat, since peers who free ride know that they will receive reduced service or face the possibility of expulsion.

Let  $Q$ ,  $R$ , and  $T$  denote the individual benefit, reduced contribution cost, and threat, respectively. A contributor would realize a performance of  $Q - R = \alpha x - \frac{x+(1-x)(1-p)}{x}$  while a free rider would realize a performance of  $Q - T = \alpha x - p\alpha x$ . Then, the new equilibrium contribution level becomes  $x = \text{Prob}(\theta_i \geq R - T)$ , and is derived by solving the fixpoint equation:  $x = \text{Prob}(\theta_i \geq \frac{x+(1-x)(1-p)}{x} - p\alpha x)$ .

With the reputation and service differentiation mechanism in place, the system performance now becomes

$$W_S(p) = x(Q - R) + (1 - x)(Q - T) = (\alpha x - 1)(x + (1 - x)(1 - p))$$

Imposing a penalty on free riders, while increasing the contribution level, entails some social loss. The p2p system designer could set the value of  $p$  to achieve a target cooperation level. Note that if the penalty is set sufficiently high, the threat  $T$  will

exceed the contribution cost  $R$ , and peers will no longer have any reason to free ride. In this case, no penalty is actually imposed. With no free riders an optimal system performance of  $\alpha - 1$  will be achieved.

**Claim 23.2** *Under the penalty mechanism, if  $p \geq 1/\alpha$ , then there exists an equilibrium in which  $x = 1$ .*

This means that if the benefits of participating in the p2p system are high ( $\alpha$  is large), either a service differentiation policy that imposes a small performance penalty on free riders or a mechanism that can catch and exclude free riders with a small probability is sufficient to induce a high level of cooperation (with any maximal generosity level). Otherwise, a more severe penalty or a finer sieve for catching free riders would be necessary.

### 23.4 A Barter-Based System: BitTorrent

BitTorrent is a popular p2p file-sharing system with incentives as an integral part of its design. It departs from earlier p2p file-sharing systems in that its incentive mechanism is based loosely on direct reciprocity rather than indirect reciprocity.

In BitTorrent, a seeding peer divides a large file into small fixed size pieces, and provides different pieces to different peers, who in turn exchange pieces with one another. A peer can reconstruct the file once it has obtained all the pieces. This technique is known as swarming download or parallel download. To induce peers to upload their pieces, a peer's download rate is influenced by his upload rate through a direct reciprocity or barter scheme.

BitTorrent attempts to alleviate the problem of random matching in large populations (Figure 23.1(b) in Section 23.2) by enforcing repeated transactions among peers. When a peer initiates a file download, it is matched with a small set of around 40 peers who are also downloading or uploading pieces of the same file. The peer selects four or five peers out of the set to connect to as neighbors, and periodically updates the list of neighbors with those peers that provide the best download rates. Through an *opportunistic unchoking* mechanism, a peer occasionally selects a random peer from the set to upload to, with the hope of finding new peers that can provide better download rates than the current neighbors.

With this design, BitTorrent peers engage in multiple interactions with a small number of peers for the duration of a file download period. For the exchange of large files such as movies and software binaries, the number of repeated interactions can be quite large, allowing cooperation to take hold through direct reciprocity. However, the BitTorrent barter scheme does not address cooperation beyond the file download period. As a result, peers have no incentive to serve as a *seeder*, i.e., to continue uploading after their own download is complete. To overcome this problem, a number of BitTorrent communities employ some form of reputation scheme on top of the existing barter scheme, and exclude peers with low contribution levels.

BitTorrent represents the state of the art in p2p file-sharing, and appears to be able to establish cooperative communities in practice. However, several theoretical and ex-

perimental studies have revealed flaws associated with its incentive scheme. Through the formalization of *specification faithfulness*, Shneidman et al. (2004) demonstrate that the BitTorrent protocol is vulnerable to a number of rational manipulations by a selfish peer, including (1) pretending to have a lower upload bandwidth while retaining relative order with respect to the upload rate of other peers, so as to reduce its upload rate without compromising its download rate; (2) pretending to be split into multiple nodes (Sybil attack) to increase its chance of being randomly selected for download; (3) replacing identities when it is beneficial to do so (whitewashing attack); and (4) uploading garbage data to boost its upload rate. Therefore, it remains an open question if and how BitTorrent (or any other p2p barter scheme) can be made robust against all forms of rational manipulations.

The Fair, Optimal eXchange (FOX) protocol offers a different, theoretical approach to solving the free-riding problem in p2p file swarming systems. Assuming that all peers are homogeneous with a capacity to serve  $k$  requests in parallel, and seeks to minimize its download completion time, FOX runs a distributed, synchronized protocol based on a static structured  $k$ -ary tree to schedule the exchange of file blocks between peers. Optimal download completion times can be achieved by all peers if all peers comply with the protocol.

FOX employs a “grim trigger” strategy to enforce compliance. When a peer finds out that its neighbor deviates from the protocol, it can trigger a “meltdown” of the entire system. This threat results in an equilibrium where all rational nodes execute the protocol as specified, since any deviation will lead to an infinite download completion time. However, the equilibrium is not a subgame perfect equilibrium, and the threat is not credible. The protocol has limited practicality since the system is vulnerable to meltdown caused by a single malicious or faulty node.

## 23.5 Currency

A p2p system can also employ a currency scheme to facilitate resource contributions by rational peers. Generally, peers would earn currency by contributing resources to the system, and spend the currency to obtain resources from the system. MojoNation and Karma are two examples of currency-based p2p systems.

Golle et al. (2001) provide the first equilibrium analysis of a p2p payment system. In the model, each peer makes an independent decision regarding his download and upload amounts. If each peer is charged an amount proportional to the gap between his downloads and uploads, then a unique strict Nash equilibrium exists where all peers would maximize their upload and download amounts.

A more recent work by Friedman et al. (2006) looks at the efficiency of a currency-based p2p system. First, it establishes the existence, for each fixed amount of money supply in the system, a nontrivial Nash equilibrium where all peers play a threshold strategy, given a large enough discount rate. When playing a threshold strategy, a peer will satisfy a request (and earn some money) if his current balance is less than some threshold value, and refuse to satisfy a request if his current balance is above the threshold. By comparing the efficiency of equilibria at different money supply levels, it is possible to determine the money supply level that maximizes efficiency for a system

of a given size. It is interesting to note that the effective money supply level can be controlled either via the explicit injection or removal of currency or via changing the price of servicing a request. This means that inflation can be used as a tool to maintain the efficiency of the system as it grows in size.

Robustness against Sybil and whitewashing attacks is still an important requirement for currency-based p2p system design in general. For example, a currency system can still be vulnerable to the whitewashing attack if newcomers are endowed with a positive opening balance, or if the balance is allowed to become negative, even temporarily.

### 23.6 Hidden Actions in p2p Systems

As we mentioned in the Introduction, strategic behavior in p2p systems goes far beyond free-riding in file-sharing networks. Peers may make strategic decisions on the timing of their arrivals and departures from the network, in selecting which peers to connect to, on whether to truthfully report to the system private information such as costs and valuations, or engage in other ways of manipulating the system protocol or mechanism. In this section, we will consider the issue of hidden action in p2p systems – how peers may behave strategically when their actions are hidden from the rest of the network, and how currency-based incentive mechanisms could be devised to overcome this problem.

Consider the case of p2p file-sharing. In addition to sharing files, the peers in file-sharing networks such as Gnutella and KaZaA are also expected to forward protocol messages to and from their neighbors. For example, when a peer receives a query message from one of its neighbors, it is expected to forward the message to its other neighbors, in addition to responding to the query if it is able to. However, the peer could strategically choose to drop the message or forward the message probabilistically, so as to reduce its message forwarding costs. In many systems, such an action is not easily observable, nor can a defecting node be readily identified, since messages are forwarded on a best-effort basis and the topology is continually changing as peers enter and leave the network. Clearly, such a system would cease to function if all peers strategically decide not to forward any messages. How can the querying node provide incentives for the other nodes to perform the message forwarding task?

The problem of hidden action in message forwarding can be readily generalized to other *peer-to-peer* settings. For example, devices in mobile ad hoc networks (MANETs) strategically drop packets to conserve their constrained energy resources. Internet Service Providers (ISPs) commonly practise *hot potato routing* to avoid the cost of transporting packets over their own networks. Indeed, the problem of hidden action is hardly unique to networks, and has long been studied by economists as the problem of *moral hazard* in contexts ranging from insurance to labor contracts. In the next section, we will apply the principal-agent framework to analyze the efficiency loss due to hidden action, and the design of optimal contracts to induce effort by the agents.

#### 23.6.1 The Principal-Agent Model

A principal employs a set of  $n$  agents,  $N$ . Each agent  $i \in N$  has a set of possible actions  $A_i = \{0, 1\}$ , and a cost (effort)  $c(a_i) \geq 0$  for each possible action  $a_i \in A_i$ . The cost of

low effort is zero while the cost of high effort is  $c > 0$ , i.e.,  $c(0) = 0$  and  $c(1) = c$ . The actions of the agents collectively and probabilistically determine a “contractible” outcome,  $o \in \{0, 1\}$ , where the outcomes 0 and 1 denote project failure and success, respectively. The principal’s valuation of a successful project is given by a scalar  $v > 0$ , while he gains no value from a project failure. The outcome is determined according to the project *technology*, or a success function  $t : A_1 \times \cdots \times A_n \rightarrow [0, 1]$ , where  $t(a_1, \dots, a_n)$  denotes the probability of project success when agents adopt the action profile  $a = (a_1, \dots, a_n) \in A_1 \times \cdots \times A_n = A$ .

We identify a subclass of technologies that can be represented by *read-once networks*. Read-once networks are given by a graph with two special nodes, a *source* and a *sink*, and each agent  $i$  controls a single edge. If an agent exerts low effort, he succeeds with probability  $\gamma_i$ , and if he exerts high effort, the success probability increases to  $\delta_i > \gamma_i$ . The project succeeds if there is a successful source-sink path, where the technology maps the individual successes and failures of agents (denoted by  $x_i = 1$  and  $x_i = 0$  respectively) into the probability of project success. Two natural examples are the “AND” and the “OR” technologies. We consider the case in which the technology is anonymous (symmetric in the agents) and is further determined by a single parameter  $\gamma \in (0, 1/2)$  that satisfies  $1 - \delta_i = \gamma_i = \gamma$  for all  $i$ .

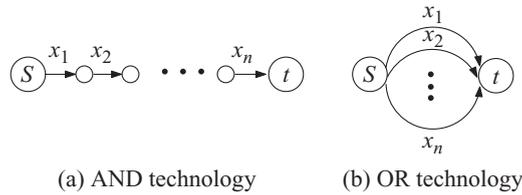
**The “AND” technology**  $f(x_1, \dots, x_n)$  is the logical conjunction of  $x_i$  ( $f(x) = \bigwedge_{i \in N} x_i$ ). Thus the project succeeds if and only if all agents succeed in their tasks (shown graphically in Figure 23.3(a)). If  $m$  agents exert effort ( $\sum_i a_i = m$ ), then  $t(a) = \gamma^{n-m}(1 - \gamma)^m$ .

For example, packet forwarding in a mobile ad hoc network can be represented by the AND technology. Each edge on the path is controlled by a single agent who succeeds in forwarding the packet with probability  $\gamma \in (0, \frac{1}{2})$  if he exerts low effort ( $a_i = 0$ ), and with probability  $1 - \gamma \in (\frac{1}{2}, 1)$  if he exerts high effort ( $a_i = 1$ ). The message is delivered to the final destination if and only if all the individual agents have succeeded in their single-hop deliveries. The sender can only observe whether the message has reached the destination.

**The “OR” technology**  $f(x_1, \dots, x_n)$  is the logical disjunction of  $x_i$  ( $f(x) = \bigvee_{i \in N} x_i$ ). Thus the project succeeds if and only if at least one of the agents succeed in their tasks (shown graphically in Figure 23.3(b)). If  $m$  agents exert effort ( $\sum_i a_i = m$ ), then  $t(a) = 1 - \gamma^m(1 - \gamma)^{n-m}$ .

For example, the practice of *multipath routing* (Ganesan et al., 2001; Xu and Rexford, 2006), where a message is duplicated and sent over multiple paths to a single destination, can be represented by the OR technology if each path is represented by a single agent.<sup>1</sup> Each agent succeeds in forwarding the message with probability  $\gamma \in (0, \frac{1}{2})$  if he exerts low effort ( $a_i = 0$ ), and with probability  $1 - \gamma \in (\frac{1}{2}, 1)$  if he exerts high effort ( $a_i = 1$ ). The project is considered a success if at least one of the messages is successfully delivered to the destination.

<sup>1</sup> Query message forwarding in p2p file-sharing networks may be modeled by OR-of-AND technology since the messages may be forwarded multiple hops along multiple paths.



**Figure 23.3.** Graphical representations of (a) *AND* and (b) *OR* technologies. The project succeeds if there is a successful path from  $s$  to  $t$ . Each agent controls an edge and succeeds with probability  $\gamma$  with no effort, and with probability  $1 - \gamma$  with effort.

The principal may design enforceable contracts based on the observable outcome.<sup>2</sup> We impose the *limited liability* constraint, thus negative payments to the agents (or fines paid by agents to the principal) are disallowed. A contract is thus a commitment to pay agent  $i$  an amount  $p_i \geq 0$  upon project success, and nothing upon project failure.

Given this setting, the agents have been placed in a game, where the utility of agent  $i$  under the profile of actions  $a = (a_1, \dots, a_n)$  is given by  $u_i(a) = p_i \cdot t(a) - c(a_i)$ . Following convention, we denote by  $a_{-i} \in A_{-i}$  the vector of the actions of all agents excluding agent  $i$ , i.e.,  $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ . The principal's problem is that of designing the contracts  $p_i$  for each agent  $i$ , so as to maximize his own expected utility  $u(a, v) = t(a) \cdot (v - \sum_{i \in N} p_i)$ , where the actions  $a_1, \dots, a_n$  are at Nash equilibrium. In the case of multiple Nash equilibria, the principal can choose a desired one and "suggest" it to the agents. While this is a standard assumption, in our setting it is further justified by the fact that the best Nash equilibrium is also a strong equilibrium (i.e., equilibrium in which no subgroup of agents can coordinate a joint deviation such that every member of the subgroup strictly improves his utility), and the unique strong equilibrium in many scenarios.

As we wish to concentrate on motivating agents, rather than on the coordination between agents, we assume that more effort by an agent always leads to a higher probability of success. Formally,

$$\forall i \in N, \forall a_{-i} \in A_{-i} \quad t(1, a_{-i}) > t(0, a_{-i})$$

In addition, we assume that  $t(a) > 0$  for any  $a \in A$ .

**Definition 23.3** The marginal contribution of agent  $i$ , given  $a_{-i} \in A_{-i}$  is

$$\Delta_i(a_{-i}) = t(1, a_{-i}) - t(0, a_{-i})$$

$\Delta_i(a_{-i})$  is the increase in success probability due to agent  $i$  moving from no effort to effort, given the effort of the others. The best strategy of agent  $i$  can be easily determined as a function of the other agents' effort levels,  $a_{-i} \in A_{-i}$ , and his contract  $p_i$ .

<sup>2</sup> An alternate approach is to maintain a trusted clearinghouse to whom agents report intermediate outcomes, and the challenge is to induce the agents to report truthfully (Zhong et al., 2003).

**Claim 23.4** Given a profile of actions  $a_{-i}$ , agent  $i$ 's best strategy is  $a_i = 1$  if  $p_i \geq \frac{c}{\Delta_i(a_{-i})}$ , and is  $a_i = 0$  if  $p_i \leq \frac{c}{\Delta_i(a_{-i})}$ . (In the case of equality the agent is indifferent between the two alternatives.)

As  $p_i \geq \frac{c}{\Delta_i(a_{-i})}$  if and only if  $u_i(1, a_{-i}) = p_i \cdot t(1, a_{-i}) - c \geq p_i \cdot t(0, a_{-i}) = u_i(0, a_{-i})$ , agent  $i$ 's best strategy in this case is to choose  $a_i = 1$ . This allows us to specify the principal's optimal contracts for inducing a given equilibrium.

**Claim 23.5** The best contracts for the principal that induce  $a \in A$  as an equilibrium are  $p_i = 0$  for agent  $i$  who exerts no effort ( $a_i = 0$ ), and  $p_i = \frac{c}{\Delta_i(a_{-i})}$  for agent  $i$  who exerts effort ( $a_i = 1$ ).

In this case, the expected utility of agent  $i$  who exerts effort is  $c \cdot (\frac{t(1, a_{-i})}{\Delta_i(a_{-i})} - 1)$ , and 0 for an agent who shirk. The principal's expected utility is given by  $u(a, v) = (v - \sum_{i|a_i=1} \frac{c}{\Delta_i(a_{-i})}) \cdot t(a)$ .

If  $a_i = 1$  in the induced equilibrium  $a$ , we say that the principal contracts with agent  $i$ . Note that the utility of the principal is lower than in the observable-actions case, as the payment to each agent is higher than the agent cost. In economic terms, the principal can only obtain the "second best" but not the "first best" solution under hidden-actions.<sup>3</sup>

The principal's goal is to determine the profile of actions  $a^* \in A$ , which gives the highest utility  $u(a, v)$  in equilibrium, given his valuation  $v$ . Choosing  $a \in A$  corresponds to choosing a set  $S$  of agents that exert effort ( $S = \{i|a_i = 1\}$ ). The set of agents  $S^*(v)$  that the principal contracts with in  $a^*$  ( $S^*(v) = \{i|a_i^* = 1\}$ ) is an *optimal contract* for the principal at value  $v$ . We will abuse notation and denote  $t(S)$  instead of  $t(a)$ , when  $S$  is exactly the set of agents that exert effort in  $a \in A$ .

A natural yardstick by which to measure this decision is the observable-actions case. When the principal can observe the individual actions of each agent, it can induce effort with a payment  $p_i = c_i$  to each agent  $i$ . In this case the principal's utility is exactly the social welfare, and so the principal will simply choose the profile  $a \in A$  that optimizes the social welfare or global efficiency,  $t(a) \cdot v - \sum_{i|a_i=1} c$ . The worst case ratio between the optimal principal's utility in this observable-actions case and his optimal utility in the hidden-actions case can be termed the *price of unaccountability*.

Given a technology  $t$ , recall that  $S^*(v)$  denote the optimal contract in the hidden-actions case and let  $S_{oa}^*(v)$  denote an optimal contract in the observable-actions case, when the principal's valuation is  $v$ .

**Definition 23.6** The *price of unaccountability*  $POU(t)$  of a technology  $t$  is defined as the worst ratio (over  $v$ ) between the principal's utility in the observable-

<sup>3</sup> In the case of "AND" technology where  $\gamma_i = 0 \forall i$ , it is shown in Feldman et al. (2005) that the principal can obtain the first best. While it is shown for the case in which agents take sequential actions, the same qualitative results also apply to the case of simultaneous actions (as  $\Delta_i(a_{-i}) = t(1, a_{-i})$  the expected utility of each agent is 0). It is also shown that the principal achieves the first best either through direct contracts (i.e., the principal contracts with each agent directly) or through recursive contracts (i.e., each agent contracts with its subsequent agent).

actions case and the hidden-actions case:

$$POU(t) = \text{Sup}_{v>0} \frac{t(S_{oa}^*(v)) \cdot v - \sum_{i \in S_{oa}^*(v)} c}{t(S^*(v)) \left( v - \sum_{i \in S^*(v)} \frac{c}{t(S^*(v)) - t(S^*(v) \setminus \{i\})} \right)}$$

For example, in the packet forwarding example, the POU measures the worst multiplicative loss incurred by the sender due to his inability to monitor the individual actions taken by the intermediate nodes.

### 23.6.2 Results

We wish to understand how the optimal set of contracted agents should be selected as a function of the principal's valuation of project success. A basic observation is that the optimal contract weakly "improves" with an increase in the valuation  $v$ .

**Lemma 23.7 (Monotonicity lemma)** *For any technology  $t$ , in both the hidden- actions and the observable-actions cases, the expected utility of the principal at the optimal contracts, the success probability of the optimal contracts, and the expected payment of the optimal contract, are all monotonically nondecreasing with the valuation  $v$ .*

For technologies in which the success probability depends only on the number of agents that exert effort (e.g., anonymous *AND* and *OR*), the above implies that the number of contracted agents is a monotonically non-decreasing function of the valuation. We find that the *AND* and *OR* technologies have very different structures on the optimal contracts: *AND* has just a single transition, from 0 agents to  $n$  agents, while *OR* has all transitions.

**Theorem 23.8** *For any anonymous *AND* technology with  $n$  agents and with  $\gamma = \gamma_i = 1 - \delta_i \in (0, \frac{1}{2})$  for all  $i$ :*

- *there exists a valuation<sup>4</sup>  $v_* < \infty$  such that for any  $v < v_*$  it is optimal to contract with no agent, for  $v > v_*$  it is optimal to contract with all  $n$  agents, and for  $v = v_*$ , both contracts (0 and  $n$ ) are optimal.*
- *the price of unaccountability is obtained at the transition point of the hidden-actions case, and is  $POU = (\frac{1}{\gamma} - 1)^{n-1} + (1 - \frac{\gamma}{1-\gamma})$*

Notice that the POU is not bounded across the *AND* family of technologies (for various  $n, \gamma$ ) as  $POU \rightarrow \infty$  either if  $\gamma \rightarrow 0$  (for any given  $n \geq 2$ ) or  $n \rightarrow \infty$  (for any fixed  $\gamma \in (0, \frac{1}{2})$ ).

This means that in the message forwarding example, the sender will induce either all or none of the agents to exert effort in forwarding a message. Moreover, the loss incurred by the sender due to his inability to monitor the individual actions may be

<sup>4</sup>  $v_*$  is a function of  $n, \gamma, c$ .

very large. This suggests a possible role for a network monitoring system, even if it is costly to implement.

Next we consider the OR technology.

**Theorem 23.9** *For any anonymous OR technology with  $n$  agents and with  $\gamma = \gamma_i = 1 - \delta_i \in (0, \frac{1}{2})$  for all  $i$ :*

- *there exist finite positive values  $v_1 < v_2 < \dots < v_n$  such that for any  $v$  where  $v_k < v < v_{k+1}$ , contracting with exactly  $k$  agents is optimal. (For  $v < v_1$ , no agent is contracted, for  $v > v_n$ , all  $n$  agents are contracted, and for  $v = v_k$ , the principal is indifferent between contracting with  $k - 1$  or  $k$  agents.)*
- *the POU for OR technology with any  $n, c$  and  $\gamma \in (0, \frac{1}{2})$  is upper bounded by  $5/2$ .*

This means that in the multipath routing example, the sender may induce any number of paths to exert effort in forwarding the message, depending on his valuation of successful message delivery. Moreover, the loss incurred by the sender due to his inability to monitor individual actions is always bounded by a factor of  $5/2$ .

For general read-once networks, it is not sufficient to determine the number of contracted agents, but the actual *set* of contracted agents. It turns out that computing the optimal contract for any read-once network, is at least as hard as computing the success probability  $t(E)$  (the network reliability), which is known to be #P-hard (Provan and Ball, 1983).

**Theorem 23.10** *The Optimal Contract Problem for Read-Once Networks is #P-hard (under Turing reductions).*

**PROOF SKETCH** We will show that an algorithm for this problem can be used to solve the network reliability problem. Given an instance of a network reliability problem  $\langle G, \{\zeta_e\}_{e \in E} \rangle$  (where  $\zeta_e$  denotes  $e$ 's probability of success), we define an instance of the optimal contract problem as follows: first define a new graph  $G'$ , which is obtained by "And"ing  $G$  with a new player  $x$ , with  $\gamma_x$  very close to  $\frac{1}{2}$  and  $\delta_x = 1 - \gamma_x$ . For the other edges, we let  $\delta_e = \zeta_e$  and  $\gamma_e = \zeta_e/2$ . By choosing  $\gamma_x$  close enough to  $\frac{1}{2}$ , we can make sure that player  $x$  will enter the optimal contract only for very large values of  $v$ , after all other agents are contracted. The critical value of  $v$ , where player  $x$  enters the optimal contract of  $G'$ , can be found using the algorithm that supposedly finds the optimal contract. At this critical value, the principal is indifferent between the set  $E$  and  $E \cup \{x\}$ . Now, from the expression for this indifference (in terms of  $t(E)$  and  $\Delta'_t(E)$ ), the value of  $t(E)$  is derived.  $\square$

A natural research problem is to characterize families of technologies whose optimal contracts can be computed in polynomial time. In addition, while there exists fully polynomial time approximation schemes (FPTAS) to various versions of the network reliability problem (Karger, 1995), it remains an open question how well one can approximate the optimal contract problem.

## 23.7 Conclusion

The fundamental premise of peer-to-peer systems is that of voluntary contribution of resources by individual users. However, there is an inherent tension between individual rationality and collective welfare. Therefore, the design of p2p incentives is of both theoretical and practical interest. In this chapter, we have reviewed different classes of p2p incentive mechanisms based on reputation, barter, and currency. We saw that cooperation can be sustained through barter if the p2p system can enforce repeat transactions among peers. Otherwise, incentive mechanisms based on reputation or currency may be necessary to overcome the free-riding problem. We also discussed the problem of hidden actions in p2p systems, and illustrated the use of contracts to induce the desired behavior by the peers.

Many challenges and open problems remain in the design and evaluation of p2p incentives, of which we highlight two. First, what is the range of possible rational manipulations against a p2p system that are either specific to, or independent of, the type of incentive mechanism in use? For example, we have seen that robustness against Sybil and whitewashing attacks are important design requirements for reputation-, barter-, and currency-based incentive mechanisms. Given a design, can we test its robustness against a comprehensive catalog of rational manipulations? Second, how should we relax the rationality assumption in the analysis and design of p2p systems, to account for heterogeneous populations of peers that may be perfectly rational, bounded rational, altruistic, malicious, and/or faulty? What would be the appropriate solution concepts for p2p systems, and for distributed systems more generally? This appears to call for cross-fertilization with both behavioral economics and computer security.

The ease of deploying p2p systems has led to their flowering in a short period of time. Today, we have a large number of p2p systems of varying scales running real applications of great value to real users. This offers us a unique opportunity to validate, using empirical data taken from real users, different designs and theories on p2p incentives. With hope, this will advance the theory and practice of incentive design for both online and offline systems.

## 23.8 Bibliographic Notes

Adar and Huberman's (2000) empirical evidence of prevalent free-riding in the Gnutella file-sharing network inspired a concerted study of incentives in p2p systems, leading to the incorporation of incentives in many p2p systems, including those discussed in this chapter: Eigentrust (Kamvar et al., 2003), Credence (Walsh and Sirer, 2005), BitTorrent (Cohen, 2003), FOX (Levin et al., 2006), Mojonation (Wilcox-O'Hearn, 2002), and Karma (Vishnumurthy et al., 2003).

Sybil and whitewashing attacks are introduced in Douceur (2002) and Friedman and Resnick (1998) and further studied in Cheng and Friedman (2005) and Feldman et al. (2006). Strategies for dealing with strangers are explored in Feldman et al. (2004) and Feldman and Chuang (2005). The minimalist p2p model in Section 23.3 is due to Feldman et al. (2006).

The study of hidden actions in p2p systems is initiated in Feldman et al. (2005), Babaioff et al. (2006a, 2006b), and the model in Section 23.6 is due to Babaioff et al. (2006a).

Strategic network formation in p2p systems is studied in Fabrikant et al. (2003), Chun et al. (2004), Albers et al. (2006), and Andelman et al. (2007). while strategic churning is discussed in (Christin and Chuang, 2005).

## Bibliography

- E. Adar and B.A. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.
- S. Albers, S. Elits, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *ACM-SIAM Symp. on Discrete Algorithms*, pp. 89–98, 2006.
- N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *ACM-SIAM Symp. on Discrete Algorithms*, 2007.
- M. Babaioff, M. Feldman, and N. Nisan. Combinatorial agency. In *ACM Conf. on E-Commerce*, pp. 18–28, 2006.
- M. Babaioff, M. Feldman, and N. Nisan. Mixed strategies in combinatorial agency. In *2nd Intl. Workshop on Internet and Network Economics*, 2006.
- A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *ACM SIGCOMM Workshop on the Economics of Peer-to-Peer Systems (P2PECON'05)*, 2005.
- N. Christin and J. Chuang. A cost-based analysis of overlay routing geometries. In *IEEE INFOCOM*, 2005.
- B.-G. Chun, R. Fonseca, I. Stoica, and J. Kubiatowicz. Characterizing selfishly constructed overlay routing networks. In *INFOCOM*, 2004.
- B. Cohen. Incentives build robustness in bittorrent. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.
- J.R. Douceur. The Sybil attack. In *Electronic Proc. Intl. Workshop on Peer-to-Peer Systems*, 2002.
- A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *ACM Symp. Princ. of Distributed Computing*, 2003.
- M. Feldman and J. Chuang. The evolution of cooperation under cheap pseudonyms. In *Proc. 7th Intl. IEEE Conf. on E-Commerce Technology*, 2005.
- M. Feldman, J. Chuang, I. Stoica, and S. Shenker. Hidden-action in multi-hop routing. In *ACM Conf. on Electronic Commerce (EC'05)*, pp. 117–126, 2005.
- M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *ACM Conf. on Electronic Commerce*, 2004.
- M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. *IEEE J. Selected Areas in Commun., Special Issue on Price-Based Access Control and Economics of Networking*, 24(5), 2006.
- E.J. Friedman, J.Y. Halpern, and I. Kash. Efficiency and Nash equilibria in a scrip system for P2P networks. In *ACM Conf. Electronic Commerce, (EC'06)*, June 2006.
- E. Friedman and P. Resnick. The social cost of cheap pseudonyms. *J. Econ. Management Strategy*, 10(2):173–199, 1998.
- D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 5(4):11–25, 2001.
- P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. In *Proc. 3rd ACM Conf. on Electronic Commerce*, 2001.

- S.D. Kamvar, M.T. Schlosser, and H. Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. In *Proc. 12th Intl. World Wide Web Conference*, May 2003.
- D.R. Karger. A randomized fully polynomial time approximation scheme for the all terminal network reliability problem. In *Symp. on Theory of Computing*, pp. 11–17, 1995.
- D. Levin, R. Sherwood, and B. Bhattacharjee. Fair file swarming with FOX. In *5th Intl. Workshop on Peer-to-Peer Systems (IPTPS)*, 2006.
- J.S. Provan and M.O. Ball. The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM J. Comput.*, 12(4):777–788, 1983.
- J. Shneidman, D. Parkes, and L. Massoulie. Faithfulness in Internet Algorithms. In *Proc. SIGCOMM Workshop on Practice and Theory of Incentives and Game Theory in Networked Systems*, 2004.
- V. Vishnumurthy, S. Chandrakumar, and E.G. Sirer. KARMA: A Secure Economic Framework for P2P Resource Sharing. In *Workshop on Economics of Peer-to-Peer Networks*, 2003.
- K. Walsh and E.G. Sirer. Fighting Peer-to-Peer SPAM and Decoys with Object Reputation. In *Proc. Third Workshop on the Economics of Peer-to-Peer Systems (P2PECON)*, 2005.
- B. Wilcox-O’Hearn. Experiences Deploying A Large-Scale Emergent Network. In *Proc. of the Intl. Workshop on Peer-to-Peer Systems*, 2002.
- W. Xu and J. Rexford. Miro: Multi-path interdomain routing. In *ACM SIGCOMM*, 2006.
- S. Zhong, J. Chen, and Y.R. Yang. Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks. In *22nd Annual Joint Conf. IEEE Comp. Commun. Soc.*, 2003.

---

## Exercises

---

- 23.1** Consider the p2p model in Section 23.3.1. The generosity of the peers is now distributed as follows: a fraction  $\phi$  of the peers have their type  $\theta_i$  uniformly distributed between 0 and  $\theta_m$ , a fraction  $(1 - \phi)/2$  are of type  $\theta_i = 0$ , and the remaining  $(1 - \phi)/2$  are of type  $\theta_i = \theta_m$ . How would the resulting equilibrium be different from that of Claim 23.1?
- 23.2** In the p2p model of Section 23.3.1, suppose that the system designer has full information on each peer’s type (i.e., generosity level), and could exclude peers based on their types (rather than based on their behavior, as suggested in Section 23.3.2). Let  $z$  denote the fraction of peers who are excluded from the system. Provide an explicit expression, as a function of  $\theta_m$  and  $z$ , for the stable equilibrium in the system under such an exclusion mechanism. Would it always (for any value of  $\theta_m$ ) be beneficial to exclude some nonzero fraction of the population? Explain.
- 23.3** Provide a proof for Theorem 23.8. Hint: First show that at  $v_*$  the principal’s utility when contracting with  $n$  agents is greater than that when contracting with  $1 \leq i < n$  agents. Then, use the monotonicity lemma to show that there must be a single transition for any AND technology. Finally, compute the price of unaccountability.
- 23.4** Provide a proof for Part 1 of Theorem 23.9, showing that for any OR technology there are  $n$  transitions. Hint: Let  $v_{i,i+1}$  ( $i \in \{0, \dots, n - 1\}$ ) be the value of  $v$  for which the principal has the same utility from contracting with  $i$  agents and with  $i + 1$  agents. First show that  $v_{i,i+1} < v_{i+1,i+2}$  for any  $i \in \{0, \dots, n - 2\}$ . Then, show that the above is sufficient to prove the theorem.
- 23.5** Prove or provide a counterexample to the following claim: For any technology, the number of transitions in the hidden-actions case is equal to the number of transitions in the observable-actions case.

- 23.6** A strategy profile  $a \in A$  is a *strong equilibrium (SE)* if there does not exist any coalition  $\Gamma \subseteq N$  and a strategy profile  $a'_\Gamma \in \times_{i \in \Gamma} A_i$  such that for any  $i \in \Gamma$ ,  $u_i(a'_{-\Gamma}, a_\Gamma) > u_i(a)$ . Prove that under the optimal payments that induce the optimal contract  $S^*$  in Section 23.6.1,  $S^*$  is a strong equilibrium.