

Reputation Systems for Self-Organized Networks: Lessons Learned

Sonja Buchegger
Deutsche Telekom
Laboratories
Ernst-Reuter-Platz 7,
D-10587 Berlin,
Germany
sonja@ieee.org

Jochen Mundinger
EPFL-IC-LCA
BC205, Station 14
CH-1015 Lausanne,
Switzerland
jochen.mundinger
@epfl.ch

Jean-Yves Le Boudec
EPFL-IC-LCA
BC203, Station 14
CH-1015 Lausanne,
Switzerland
jean-yves.leboudec
@epfl.ch

Abstract

Self-organized networks such as mobile ad-hoc, Internet-based peer-to-peer, wireless mesh and Fourth Generation (4G) Wireless networks depend on cooperation of nodes. Reputation systems help nodes decide with whom to cooperate and which nodes to avoid. They have been studied and applied almost separately in diverse disciplines such as economics, computer science and social science, resulting in effort duplication and inconsistent terminology. In this paper, we aim at bringing together these efforts by outlining features and fundamental questions common to reputation systems in general. We derive methodologies to address these questions and lessons for both reputation system design and research from our own experiences and evaluations by simulation and analytical modelling. We argue for using deviation tests, discounting, only passing on of first-hand information, secondary response, and stressing the importance of identity.

1 Reputation Systems

Self-organized communication systems such as mobile ad-hoc, Internet-based peer-to-peer and wireless mesh networks have received increasing attention, in terms of both deployment and research.

They are typically organized according to the peer-to-peer (P2P) organization principle. That is, participants in the system are equals in that they have equivalent capabilities and responsibilities – they are peers. Such P2P systems can also be found in a variety of other networks, for example social or biological networks. Thus it is not surprising that there is a wealth of problems that is also of interest in other disciplines.

In the novel Fourth Generation (4G) paradigm as well there are self-organized components such as ad-hoc connectivity among spatially close nodes.

One of the major issues in such self-organized communication systems is that of *cooperation*. Typically, users are concerned primarily about their own benefits and thus cooperation and fairness cannot be guaranteed. This selfish behavior is called *free-riding* and is a well-known phenomenon in economics. The *free-rider problem* is that as a result this service might not be provided at all or without sufficient quality of service [15]. Effects can be detrimental as shown, for example, in Internet-based P2P networks [6].

Although altruistic behavior has been observed, it is not clear to which extent this will help in communication systems. *Altruism* is the practice of being helpful to other people with little or no interest in being rewarded for one's efforts [1, 8, 14].

Incentive mechanisms (pricing mechanisms as well as rules) and artificial immune systems [16] have been proposed and investigated to address the issue of cooperation in communication systems.

Two other problems often incurred in P2P networks are malicious attacks and random failures. *Reputation systems* [5] address both these issues as well as incentive problems. Here, users keep track of their peers' behavior and exchange this information with others in order to compute a reputation value about their peers. Reputation values are then used to decide with whom to cooperate and which nodes to avoid, i.e. users with a good reputation are then favored [4].

Reputation systems have already proven useful and are popular in online auctioning systems such as eBay [17] or online book stores such as Amazon. However, unlike self-organized communication systems, those have a centralized component.

We next look at reputation systems in more detail before addressing fundamental features that a reputation system should have as well as fundamental questions that need to be addressed.

2 Terminology and Classification of Reputation Systems

Reputation systems have been studied and applied almost separately in diverse disciplines such as economics, computer science and social science, resulting in effort duplication and inconsistent terminology. Even within computer science research activities have not been very consistent, and have almost evolved separately in the artificial intelligence, Internet-based P2P and Mobile Ad-Hoc Networks communities. In fact, there is not even a consistent definition of *reputation* itself, nor, closely linked, of *trust*.

As a convention, following the Oxford English Dictionary, we shall adopt that reputation is an estimate about a person's actual quality. *Person* is the appropriate term for social networks. In the context of computer networks we shall replace it with *user* (of the system), *node* (in the network) or simply a *peer*. Similarly, *quality* refers to the behavior that is of interest in a given context. For 4G, it might refer to the packet forwarding behavior.

Moreover, we shall adopt that *rater reputation* refers to the reporting behavior within the reputation system. This is different from our earlier papers [4] where we referred to it as trust. Although, again, the definition in the Oxford English Dictionary of trust fits our previous usage ("Confidence in or reliance on some quality or attribute of a person or thing, or the truth of a statement"), the term *trust* has been used in several different ways in the literature, also synonymously with *reputation*, and has therefore become too ambiguous for our purposes.

As opposed to reputation, rater reputation values are based on compatibility and thus indicate agreement.

As for classification, there have been various surveys in different computer science communities using partly overlapping criteria. However, most of them are relevant in all the other communities, too. For further information, the reader is referred to [10].

In this article, we focus on fundamental features and questions that concern reputation systems in general, independently of the application domain and only mention specifics for self-organized networks when they deviate from generally applicable aspects of reputation systems.

3 Features That a Reputation System Should Have

The basic premise of a reputation system is that one can predict future behavior by looking at past behavior. This does not hold for all cases, since there can be erratic behavior that is completely inconsistent with past behavior, as in the case of sudden failure. But the assumption is that such cases are the exception and not the norm and that past behavior can be used as a basis for the prediction of future behavior.

To provide this basis, the reputation system has to keep track of past behavior. This can be done in several ways. Here are some decision points to guide the design process of a reputation system.

What information is kept? About whom? Where? For how long? When is information added? How is information from others considered? How is it integrated? What does this information look like over time? What has to happen to change this information?

In summary, a reputation system needs a way of keeping information about the entity of interest, of updating it and of incorporating the information about that entity obtained from others. This provides the basis of decision making. Then the decision making itself has to take place to allow nodes to choose other nodes for cooperation.

Humans can look at graphical representations of reputation such as the number and color of stars on eBay, and glance over some qualitative information given in feedback comments. In self-organized networks, we want the reputation system to not only be able to present information about reputation to the user, but to make automatic and autonomous decisions. The reputation system therefore has to have a mechanism to make decision and classifications. In this paper, however, we concern ourselves mainly with the reputation information itself.

As time passes, the importance of parts of the reputation data collected can change. For instance, recent steady behavior is probably a better predictor of future behavior than behavior observed a long time ago. On the other hand, looking only at the most recent behavior can yield a distorted picture of past behavior as one instance observed is not enough to measure a trend. Reputation systems need to have a way of factoring in time in a reasonable manner that would either conform to the user's expectation or be proven to work well in the system environment.

Figure 3 shows the workings of an idealized reputation system. Several sources contribute to the generation of reputation values: direct (own) observations, indirect information from others, and time passing. Once the reputation value is determined, the subjects of interest can be classified and reactions according to these classifications can be triggered.

In the following we look at the features needed from a reputation system in more detail.

Keeping track of past behavior. Reputation is a function of past behavior and time, so a reputation system needs to collect data about past behavior. These data can be stored in a centralized or in a distributed way. For self-organized networks, a distributed storage of reputation data is needed, as there is no infrastructure in place to reliably ensure access to a centralized reputation authority.

The reputation system has to offer a way of collecting information about the entities of interest. In a self-organized network these entities of interest are neighboring nodes and nodes that are on any communication path to other nodes. Nodes can join and leave self-organized networks. There is a tradeoff between performance and overhead of reputation systems in terms of which entities to keep track of, as nodes may be short-lived in a network and every entry entails both storage and potential maintenance cost as computation.

In addition to the decision about which nodes to keep track of, and the feasibility of doing that in a particular instance of a self-organized network, there is the question of which data to




Figure 1: Reputation System Flow

collect about the behavior of an entity of interest. Assuming binary behavior, i.e. behavior is either good or bad, cooperative or defecting; the basic choice is between keeping track of the volume of good/bad behavior and the ratio of good to bad behavior, or both if the chosen algorithm for calculating a reputation value can take advantage of it.

A good reputation should be a reflection of good behavior. Just what constitutes good behavior needs to be clearly defined in order to determine how reputation should be calculated given data about behavior. Using the ratio of good to bad behavior, for instance, reflects the willingness to cooperate only in relation to a specific extent of demand or opportunity for cooperation. This extent (say the number of cooperation requests) is lost in the ratio and therefore unknown. The explanatory power of a cooperation ratio is thus limited. Conversely, if the absolute number of cooperation instances (or misbehavior, respectively) is taken as the basis for reputation calculation, it is not known what the number of opportunities where out of which the behavior was good or bad. A combination of ratio and volume calculations would capture willingness to cooperate in relation to opportunity.

Incorporating data from different sources. There is a tradeoff between speed and accuracy: the more second hand information is used, the faster an estimate of some subject's behavior can be obtained, however, the more vulnerable it is to liars. In order to be useful, reputation values need to be accurate, at least to some degree. This has to be assessed.

Whereas direct observations should always be accepted, second hand information should be accepted only if considered as likely, i.e. only if it does not differ by more than is acceptable (e.g. as measured by a threshold Δ or other suitable metrics) from the user's current reputation value. This behavior is comparable to the concept of *confirmation bias* in the social sciences, and can also be motivated by observations in everyday life. Even if accepted, one might want to weight them by a factor ω_{weight} .

Confirmation bias is an example of human faulty reasoning, discarding information (facts) that do not fit a theory and favoring confirming information. The use of such a fallacy might seem counterintuitive. Besides the motivation of excluding spurious information rather than protecting an already formed opinion, there are differences in how we employ confirmation bias that allow us to use it as a rational tool. First, the own information of a node, i.e. its own observations are not subject to the confirmation bias, all direct information is incorpo-

rated as-is. Second, the observed behavior is binary (cooperate/defect) and made with high certainty, hence the trust in the node's own judgment is justified. These two points amount to what corresponds to *undeniable facts* in human confirmation bias. When facts are undeniable, they are included despite the bias toward only accepting facts that confirm one's belief. Third, the confirmation bias is applied only to decide whether to accept third-party indirect observations and thereby is a measure of compatibility with a view on reality with a high probability of accuracy. This rests on the assumption of behavior being constant over different interaction partners. This can be compared to collaborative filtering.

Forgetting reputation over time. Taking into account the passing of time in a reputation system allows for two features: emphasizing the importance of behavior at one time over another, e.g. of recent behavior over behavior observed long ago, and providing the possibility of revising the action toward a node that was triggered by a particular reputation value, i.e. redemption of a node after it has been repaired. We suggest to include a discounting with a factor ρ , so that old observations gradually become less important. This is a form of forgetting.

Secondary Response. To avoid that such forgetting backfires, a mechanism such as secondary response can be introduced which provides increased sensitivity to misbehavior by nodes that have been deemed misbehaving in the past. By increased sensitivity, appropriate action can be triggered faster than in a regular case.

4 Questions for Reputation Systems

There are some fundamental questions regarding effectiveness and robustness that need to be addressed for reputation systems in all communities.

- What is the impact potential liars could expect to achieve on the reputation value about the subject in question?
- Which kind of information should be passed on to other nodes to achieve an accurate reputation system? The comparison of two different scenarios as regards to the second hand information. In the first one, *Reputation* – based on all previous observations including indirect ones – is passed on as second hand information. In the second one, only *Direct Observations* are passed on as second hand information. What difference does this make?
- What strategies can an attacking node employ to distort the reputation system, in addition to lying?
- How can the reputation system recover from false positives or negatives?
- What is the impact of incomplete information? Especially in distributed reputation systems, such as those for self-organized networks, nodes only have a partial view of the environment, only a subset of all peers and only a subset of the behavior of these peers is known.
- What is the impact of wrong observations? These can happen due to the inherent lack of unambiguous observability such as the difficulty to distinguish between deliberate packet dropping and congestion, mobility, loss of connectivity in wireless networks.
- Why should nodes participate in a reputation system? Is there an incentive to cooperate and contribute to a reputation system, and to do so honestly?
- How accurate and fair is the reputation system? How well do the reputation records represent past behavior? The goal of reputation systems is usually twofold: to enable nodes to find

good peers and to give an incentive for cooperation. It is not straightforward to do this in an accurate and fair way, taking into account not only the actual instances of cooperation, but both the opportunity and the willingness to cooperate. What is the metric for accuracy?

5 Methodology to Answer Fundamental Questions

To answer fundamental questions about general reputation systems independently of implementation details, we suggest the consideration of an abstract model supported by simulations and measurements.

For the modelling part, we are not concerned with the detection and response components of a system, but focus on the actual formation of reputation. The detection component depends on the application scenario and we merely assume that misbehavior can be told apart from good behavior. Moreover, we assume that if reputation values can be computed accurately, then there exists a response mechanism using them to obtain the desired effects. Typically, this might mean exclusion of the misbehaving user from benefits.

We formulate a stochastic process formulation, based on which we derive an ordinary differential equation by averaging the dynamics and passing to a *fast-time scaling* limit. That is, we scale time so that events occur more frequently, i.e. users make observations at a higher rate, but at the same time the impact of each observation is reduced by the same factor. We then derive the solutions of the differential equation and study their fixed points. Thus, our approach can be called a *mean-field* approach [9]. Moreover, we use simulation and direct computation to confirm the analytical results.

Reputation or recommender systems can collect opinions about the quality of objects such as films. Here, we are concerned with reputation systems for self-organized networks where nodes give reputation ratings about other nodes, i.e. active subjects and peers in the network.

There are no benchmarks available for the simulation of reputation systems in self-organized networks. It is customary to simulate reputation systems in self-organized networks by augmenting the simulation of regular network behavior by a reputation system component and simulating specific scenarios of node misbehavior.

The parameters for the regular network behavior typically include the number of nodes, at least an initial topology (in the case of model Internet-based peer-to-peer networks potentially also topology control of the overlay network), and a routing protocol. Simulations of wireless networks additionally include a mobility model, assumptions about physical characteristics of devices and terrains.

While the simulation of self-organized networks without reputation systems already offer many potential pitfalls in the choice of parameters that make the obtained results difficult to reproduce and achieve statistical significance and generalizability, additional care has to be exercised when simulating reputation systems on top of regular network behavior. The model for node behavior (and misbehavior) of both the network (e.g. forwarding) and the reputation itself (e.g. lying) strongly influences the results and determines the scope of the conclusions that can be drawn from these results. Using a wide range of scenarios and node behaviors (threat models, attacker/failure models) can help to expose vulnerabilities of a reputation system by simulation.

6 Lessons Learnt

Motivated originally by observations in everyday life as well as by research in the Social Sciences, and supported by our analytical modelling as well as simulation results, we have learnt the following for the design of a reputation system.

Lesson 1: Deviation tests mitigate spurious ratings. Using a test to evaluate each piece of information with respect to how it conforms to a node’s own view, i.e. quantifying the congruence of views necessary for confirmation bias, turns out to be a more fine-grained and adaptive approach than only considering the rater reputation of the node providing the reputation information. The deviation test works as follows. Every time a node receives reputation information from another node, it has to decide whether and how to consider this information. It compares the received information with its own prior knowledge and only accepts it if the deviation is less than a specified acceptable deviation, otherwise the received information is discarded. This produces a non-linear effect.

More precisely, we find by analysis [12, 11, 13], that in order to have an impact, the number of liars N_l in the network needs to exceed a certain threshold. That is, there is a phase transition behavior. The phase transition behavior can be phrased in terms of the parameter Δ (the deviation threshold) rather than in terms of N_l . If Δ is below a certain threshold, that can be computed, the liars have no impact.

Lesson 2: Discounting adds resilience. Giving more weight to recent behavior and discounting past behavior as time passes achieves two objectives: better correlation to future behavior and allowing for node redemption: When past behavior is discounted, nodes cannot capitalize on previous good behavior but have to consistently behave well to maintain a good reputation. Information about nodes has to be constantly reinforced to stay current. Node redemption allows for a node to regain at least a neutral reputation after a specified time period (determined by the discount rate) without bad behavior. This is crucial for example for dealing with formerly faulty nodes that have been repaired and useful in general to adapt to behavior changes of nodes regardless of the reason.

Lesson 3: Passing on first-hand information only (*direct observations*) improves accuracy.

Passing on information received from others, as opposed to direct observation, i.e. rumor spreading turns out to not only offer no gain in reputation accuracy or speed, but also to introduce vulnerabilities by creating a spiral of self-reinforcing information [3].

This is confirmed by a theoretical analysis. We find that the performance of *Direct Observations Reputation* coincide on some range (namely, if and only if the $\theta > 2\Delta$, where θ is the probability that a well behaving node is indeed observed as behaving well), but that otherwise *Direct Observations* is more robust against liars. For *Reputation*, second hand information does not improve accuracy, whereas for *Direct Observations* it does and overall *Direct Observations* is better.

Lesson 4: Secondary responses accelerate classification. To offset a potential vulnerability of granting redemption (by means of discounting as described above, to a node previously classified as having a reputation rating too low for cooperation) and thereby providing a chance for misbehavior, the mechanism of secondary response has turned out to be useful. Secondary response is inspired by the human immune system and means, in the context of reputation systems, that the tolerance of misbehavior is reduced for nodes that have been deemed misbehaving previously.

Lesson 5: Identity is an issue. If the identity of a node cannot be established at all or only for a short while, the accuracy and robustness of reputation systems suffer. While in some environments the need for reputation information might be limited to short periods of time

(e.g. establishing a path in a network) and thus can deal with short-lived identities, in general a reputation system needs to have longer lived identities to make use of its features. In general, identities have to persist longer than the detection time of a misbehaving node. A fundamental requirement for identity in reputation systems is that one can be sure that an observed behavior has actually been exhibited by the observed node. Identity spoofing, be it impersonation or the creation of false identities (e.g. Sybil attacks [7]), and identity persistence over time are crucial for reputation systems effectiveness. The assumption of accurate and stable identities, generally used for reputation systems, is a strong one and difficult to realize in self-organized environments. 4G networks potentially can provide a solution to this problem by offering a mix of self-organized networks with access to infrastructure that would enable central authorities.

7 Issues For Reputation Research

We now indicate conclusions that might be of interest for future research on reputation systems.

Issue 1: Coherent terminology. It has become apparent that both definition and representation of reputation vary widely even within computer science. While it is debatable whether or not the same definition and representation should be used for all applications, a more coherent terminology would certainly be desirable. Moreover, it would be useful to have a coherent classification of reputation systems.

Issue 2: Coherent classification. Based on a more coherent terminology, it would also be desirable to bring together the different strands of research, within computer science, but also between disciplines. This would avoid lots of effort duplication that can be observed at the moment. The reputations research network¹ went only some way towards this. In this article, we have attempted to at least point out work on reputation in the different communities. A number of interesting examples of the successful combination of different disciplines (physics, economics, social sciences) can be found in [2].

Issue 3: Fundamental questions as well as specific implementations. Clearly, specific applications for distributed reputation systems are of crucial importance and many papers address various scenarios. However, it is also important not to get lost in the details. There are fundamental questions that are important in all these scenarios that should be addressed on a suitable level of abstraction. We have provided a list of them in Section 4.

Issue 4: Models as well as simulation and measurements. Finally, computer science research is often based on simulations, measurements or implementation and testing of prototypes. For example, prototype protocols are typically evaluated using a network simulator such as ns-2² or GloMoSim³. Apart from game and graph theoretic investigations, there are comparatively few analytical studies although they often provide insight that is hard to obtain otherwise. An example of this is the stochastic process formulation of the model referred to in Section 5. Even though results are typically proven to be valid under clearly defined assumptions, some of which might be unrealistic, it is often the case that results are valid at least qualitatively even if the assumptions are violated. In the social sciences context in particular, such approaches are rare although it would be desirable to enhance predictive capabilities of social networks. Moreover, there might yet be other approaches (than game theoretic, graph theoretic and stochastic models) that have not been considered so far.

¹<http://databases.si.umich.edu/reputations/index.html>

²<http://www.isi.edu/nsnam/ns/>

³<http://pcl.cs.ucla.edu/projects/glomosim/>

8 Reputation Systems in 4G

4G wireless networks provide new opportunities and challenges for reputation systems but share the same fundamental questions mentioned in this paper. The additional questions that are specific to 4G networks arise from their combination of self-organization and access to infrastructure, e.g. could 4G address identity better? 4G networks would enable centralized authorities at most times, but also have distributed, self-organized components. More centralization means more data can be kept, identities could be more easily verified by e.g. public key infrastructures. The reputation system should ideally take advantage of the added benefit provided by mostly-on central authorities but function in a self-organized way. When in disconnected mode or switching back and forth between centralized and decentralized modes, updates need to be consistent: there is a challenge to combine the two worlds.

With 4G, nodes can be in more and different networks at the same time. Multihoming and the availability of ubiquitous network access via several technologies in the same physical space enables more decision space for network selection, trading off bandwidth, power, cost, and convenience. Selecting a network for each transaction can be facilitated by using reputation systems that capture the properties of the available networks. Besides direct interaction with the 4G infrastructure by the network provider, nodes can decide to cooperate directly to take advantage of low-cost high-bandwidth local connections to share content (e.g. obtained from low-bandwidth provider links) in real-time or pass on messages in delay-tolerant networks.

Reputation systems are useful in cases of cooperation and when making choices in a more informed way. Although the applications for 4G themselves are yet to clearly emerge, it is clear that 4G networks provide such cases of cooperation and choice more than traditional networks and we submit that reputation systems are a valuable tool in the box to make 4G networks work well.

9 Conclusions

Reputation systems across different domains share some fundamental features and questions that need to be addressed concerning the accuracy of reputation values and the robustness against spurious information. In addition, 4G networks call for a combination of decentralized and infrastructure, which poses a set of new challenges. We have learned some lessons from the development and analysis of reputation systems in self-organized networks and have developed methodologies to address some of the fundamental questions of reputation systems. For the remaining open questions, we think that using a coherent terminology and classification across disciplines and also addressing fundamental questions independent from their environment by means of both analytic modelling and simulation will help finding solutions.

References

- [1] J. Andreoni and J. H. Miller. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70:737–753, 2002.
- [2] P. Ball. *Critical mass: how one thing leads to another*. Farrar, Straus and Giroux, 2004.
- [3] S. Buchegger and J.-Y. L. Boudec. The effect of rumor spreading in reputation systems in mobile ad-hoc networks. Wiopt’03, Sofia-Antipolis, March 2003.
- [4] S. Buchegger and J.-Y. L. Boudec. A robust reputation system for peer-to-peer and mobile ad-hoc networks. P2PEcon, Harvard University, Cambridge, MA, USA, June 2004.

- [5] S. Buchegger and J.-Y. L. Boudec. Self-policing mobile ad-hoc networks by reputation systems. *IEEE Communications Magazine*, pages 101–107, July 2005.
- [6] J. Chu, K. Labonte, and B. N. Levine. Availability and locality measurements of peer-to-peer file systems. In *ITCom: Scalability and Traffic Control in IP Networks, Proceedings of SPIE*, volume 4868, 2002.
- [7] J. R. Douceur. The sybil attack. In Proc. of the IPTPS02 Workshop, Cambridge, MA (USA), March 2002.
- [8] U. Ebert and O. von dem Hagen. Altruism, redistribution and social insurance. *Review of Economic Design*, 5:365–385, 2000.
- [9] J.-Y. Le Boudec, D. McDonald, and J. Munding. A Generic Mean Field Convergence Result for Systems of Interacting Objects. In *QEST'07*, 2007. Keynote Speaker.
- [10] J. Munding and J.-Y. L. Boudec. Reputation in self-organized communication systems and beyond. In *Interperf '06: Proceedings from the 2006 workshop on Interdisciplinary systems approach in performance evaluation and design of computer & communications systems*, page 3, New York, NY, USA, 2006. ACM Press.
- [11] J. Munding and J.-Y. Le Boudec. Analysis of a reputation system for mobile ad-hoc networks with liars. In *Proceedings of WiOpt 2005: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, pages 41–46, 2005.
- [12] J. Munding and J.-Y. Le Boudec. Analysis of a robust reputation system for self-organised networks. *European Transactions on Telecommunications, Special Issue on Self-Organisation in Mobile Networking*, 16(5):375–384, October 2005.
- [13] J. Munding and J.-Y. Le Boudec. The impact of liars on reputation in social networks. In *Proceedings of Social Network Analysis: Advances and Empirical Applications Forum*, Oxford, UK, July 2005.
- [14] D. C. North. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, 1990.
- [15] P. A. Samuelson. The pure theory of public expenditure. *Review of Economics and Statistics*, 36(4):387–389, 1954.
- [16] S. Sarafijanovic and J.-Y. Le Boudec. An Artificial Immune System Approach with Secondary Response for Misbehavior Detection in Mobile Ad-Hoc Networks. *IEEE Transactions on Neural Networks, Special Issue on Adaptive Learning Systems in Communication Networks*, 16(5):1076 – 1087, 2005.
- [17] The Economist. Special report eBay. The Economist, June 11th, 2005, 2005.