A Novel Utility and Game-Theoretic Based Security Mechanism for Mobile P2P Systems

Brent Lagesse and Mohan Kumar
Department of Computer Science Engineering
University of Texas at Arlington
{Brent.Lagesse, mkumar}@uta.edu

Abstract

Research on security in Peer-to-Peer (P2P) systems is dominated by reputation-based solutions. These solutions propagate opinions about other peers in order to help identify the best set of peers to utilize. In this paper, we model peers with utility functions and use those functions to examine the case in which an individual peer participates in an unfamiliar and untrusted system, similar to one in which a mobile peer can enter when moving into a new location. We additionally introduce a novel security mechanism for P2P systems called resource exploration in order to mitigate the problems inherent in reputation-based systems and analyze its effect on a 2-player game (between an attacker and the benign peer).

1. Introduction

There have been many efforts dealing with economics in P2P computing applications [1, 3, 4]; however, these methods primarily deal with forms of freeloading. Instead, this paper will discuss the economics of end-to-end transactions in a P2P system as perceived through the utility functions of individual peers.

The goal of this research work is to analyze security mechanisms in P2P systems in order to better compare them, define new mechanisms to improve security, and better understand the economic cause of vulnerabilities. One of the insufficiencies of reputation-based security in P2P systems is that it requires prior experience in order to make decisions. As a result, a malicious peer must have previously attacked another peer in order to be recognized as malicious. In a foreign system with no known trusted peers, an entering peer is vulnerable to attack as it has no means to determine the trustworthiness of any other peers in the system. This fact can be exploited by an individual malicious peer or by a set of collaborating peers. Furthermore, a peer can initially behave benignly, be recognized as such, and then act maliciously (either intentionally or due to being compromised). These attacks are described as risks that require deeper investigation in [5].

In this paper, we have identified utility functions that define the behavior of peers in a P2P system. These utility functions take into consideration the costs and benefits as perceived by each peer by being connected to the P2P system and particular events that occur within the P2P system. Our modeling approach using utility functions allows a better understanding of the mechanisms for providing protection from malicious peers. We also introduce a resource exploration (Res-Exp) mechanism aimed at mitigating the vulnerabilities in reputation-based security that are listed above.

The rest of this paper is organized as follows. Section 2 is a brief overview of related work. Section 3 describes the general utility model we use for the remainder of the paper. Section 4 contains the analysis of a P2P system based on the utility model and introduction of a new security mechanism for P2P systems. Section 5 contains future work, and finally we conclude in Section 6.

2. Related Work

Trust, in the form of reputation management[9, 6, 5, 10], and incentives[4] have largely been a focus in P2P security from an end to end perspective. Reputation systems focus on accumulating reputations and propagating them through the network, so other peers can interpret the reputations to make decisions on who they should trust. Incentive solutions provide some form of payment to peers to encourage good behavior. The problem with reputation systems is that they require prior knowledge to work. In other words, peers are vulnerable to attack if they do not have knowledge or correct knowledge of other peers in a reputation system. As stated in the introduction, the vulnerability is most evident when a peer first enters a system or a peer previously recog-

nized as benign chooses to betray trust (or is compromised). Since that peer would have a good reputation up until that point, a reputation system would give no reason not to trust that peer. Incentive systems are vulnerable because they do not prevent attack, they just give more reason to cooperate in the system, but the vulnerability is still there if the malicious peer prefers acting maliciously enough.

Research in economics, particularly utility functions and game theory, has had a large influence in computer science. While much of the research is focused on auctions, some similar concepts that are discussed in this paper are being researched [7]. In particular, economic-based approaches have permeated both security [4, 8] and P2P computing [1, 3, 4]. These solutions do very little to address general malicious behavior in P2P systems. Instead, those related to P2P systems are largely focused on incentives to prevent freeloading.

In the remainder of this paper, we borrow techniques from utility and game theory in order to model and analyze peers in a P2P system.

3. Utility Model

We define a set of relationships between benefits and costs that are intended to capture the potential sources of benefit and cost that would drive a generic peer. By generic, we mean that we do not require a peer to be purely malicious or purely benign. Instead, a peer's actions will be evaluated based on its utility function. The following terms are used in the relationships.

The utility function is normalized to unit-less (typically non-negative) values. Since most of the low-level components that make up the utility relationships are preferences, such as an aversion to being subjected to a denial of service attack, we do not provide any formal method for determining the values of those costs and benefits, though in many cases these benefits and costs could be described financially.

The $VictimCost$ is a relation that captures the negative effect on a peer when it becomes the victim of an attack. It allows us to describe a peer's aversion to being attacked and plays a large role in determining how much effort should go into avoiding attacks or whether to participate in a system at all.

$$VictimCost = SpyVictim + DenyVictim + MisinformVictim \quad (1)$$

Benign benefit captures the benefit gained by legitimate participation in a P2P system. It consists of the benefit a peer perceives from accessing resources and any benefit that is derived from mechanisms in the system (for example, incentives for sharing useful resources).

$$BenignBenefit = AccessToResourceBenefit + MechanismBenefit \quad (2)$$

Malicious benefit captures the benefit gained from acting maliciously. This is described by the actions of spying on a peer, denying access to a peer, and providing faulty information to a peer.

$$MaliciousBenefit = Spy + Deny + Misinform \quad (3)$$

Benign cost is the cost of participating in the system. This is the overhead cost of staying in the system (as derived and normalized from energy, memory, bandwidth, etc.) in addition to the costs incurred from providing resources and any costs from mechanisms incorporated in the system (such as punishments to prevent freeloading).

$$BenignCost = TimeConnected + ProvidingResources + MechanismCost \quad (4)$$

Malicious costs are costs associated with malicious actions, and include bandwidth costs or processing costs. While these are likely to be relatively small for a malicious peer, they do exist and are incorporated into the relationships.

$$MaliciousCost = Spy + Deny + Misinform \quad (5)$$

All of these benefits are tied together and then related to each other to provide the overall utility as described in Equation 8. Also included in the cost is $DiscoveryCost$ which is the cost of an attacker being discovered (which may take the form of having to exit and re-enter the system or even just a decrease in available peers to attack).

$$Benefit = BenignBenefit + MaliciousBenefit \quad (6)$$

$$Cost = BenignCost + MaliciousCost + VictimCost + DiscoveryCost \quad (7)$$

$$Utility = Benefit - Cost \quad (8)$$

For the remainder of this paper we will simplify some of the details of specific attacks by dealing only with $BenignBenefit$, $MalciousBenefit$, $BenignCost$, $MaliciousCost$, $DiscoveryCost$ and $VictimCost$. We feel that handling the relationships at this level of details provides sufficient information without sacrificing simplicity and generality. Furthermore, this approach allows us to

speak generally about attacks rather than identifying specific attacks during the analysis.

For simplicity in our analysis, we examine purely benign versus purely malicious peers, though we could model hybrid peers which would represent a peer that uses the system both as it is intended and maliciously. Malicious peers only gain utility from successfully attacking other peers and can be modeled by using the components $MaliciousBenefit$, $BenignCost$, $DiscoveryCost$ and $MaliciousCost$. Purely benign peers only gain utility from successful transactions and can be modeled with the components $BenignBenefit$, $BenignCost$, and $VictimCost$.

4. Resource Exploration

4.1. Theory

From a utility perspective, security mechanisms can be introduced in order to manipulate the utility perceived by a peer without having to change the peer's preferences. The goal of such a security mechanism is to reduce (or ideally eliminate) the effectiveness of an attacker while not being a significant burden on benign peers. A mechanism should simultaneously require an attacker to change to a strategy that involves less malicious activity in order to maximize its utility and at least not decrease the utility of benign peers. In other words, the benign peer should have its expected utility improved by suffering less $VictimCost$, and that reduction in expected $VictimCost$ should at least balance the costs inflicted by implementing the mechanism.

We introduce a novel mechanism called resource exploration (Res-Exp) in this section. The main idea of the Res-Exp mechanism is to send out exploratory requests in addition to real requests. These exploratory messages are designed to reveal the nature of the peers. Peers will incur a cost by sending the exploratory messages ($ResExpCost$), but the messages reduce the likelihood of being attacked through increasing the costs incurred by the attacker if discovered as a malicious peer ($DiscoveryCost$).

At this point, we must state some assumptions about the Res-Exp approach. First, we have to assume that the peer responding to the request cannot differentiate between exploratory and regular requests. We justify this assumption by noting that a peer can reuse previously obtained results, self-generated results, or pre-programmed results depending on the specific application. Second, we have to assume that the peer sending the Res-Exp requests can verify whether or not an attack has occurred. We justify this assumption by noting that the assumption is common, though often not explicitly stated, in P2P reputation research [2, 5, 12] since in order to provide opinions, peers must know that they were attacked. In many cases, this is manually determined by human users, but due to the cost of human intervention, our work is most useful if the attacks can be automatically determined (for instance, comparing the known checksum of a file to a generated checksum of the file sent by another peer).

Before a peer can decide to utilize the above Res-Exp mechanism, it needs to determine at what rate to send out exploratory messages. We will show two similar methods for determining this rate. The first is to show the mixed-strategy Nash equilibrium. This approach is difficult since it requires knowledge of the attacker's preference. In order to overcome this limitation, we provide a second approach that allows a peer to determine bounds that it is willing to operate within.

4.2. Nash Equilibrium

In order to determine a Nash equilibrium we have to set up the parameters of the game as shown in Figure 1. The setup is a two player game comprising of a requesting peer and a serving peer. The requesting peer can either send an exploratory message or a request message. The serving peer can either respond with an attack or with a legitimate response. In all cases, each peer will incur $BenignCost$. For the sake of simplicity, we will also assume that the cost of serving maliciously is the same as serving benignly and that cost of sending an exploratory request is the same as the cost of sending a legitimate request. This does not have to be true for this approach to work though, it just requires us to work at a lower level and with more terms in the utility relationships. Given this, we can reduce the effects (not captured in $BenignCost$) of an action down to the cost of being discovered when the serving peer attacks an exploratory message, the benefit from acting malicious and the cost of being a victim when a serving peer attacks a legitimate request, and the benefit a peer receives from being served during a legitimate request. By setting the expected value of each action a peer could take equal to the alternative action, we can determine the mixed-strategy equilibrium. As a result, the serving peer should attack with a probability defined by Equation 10 and the requesting peer should use exploratory messages with a probability defined by Equation 9.

$$P_{exp} = \frac{MaliciousBenefit}{DiscoveryCost + MaliciousBenefit} \quad (9)$$

$$P_{attack} = \frac{BenignBenefit}{VictimCost + BenignBenefit} \quad (10)$$

An obvious downside to this approach is that it requires a knowledge of the opponent's preferences. It also does not assure that the transaction is beneficial (not playing the game could result in a higher expected utility than playing). For instance, if there is no $DiscoveryCost$, then P_{exp}

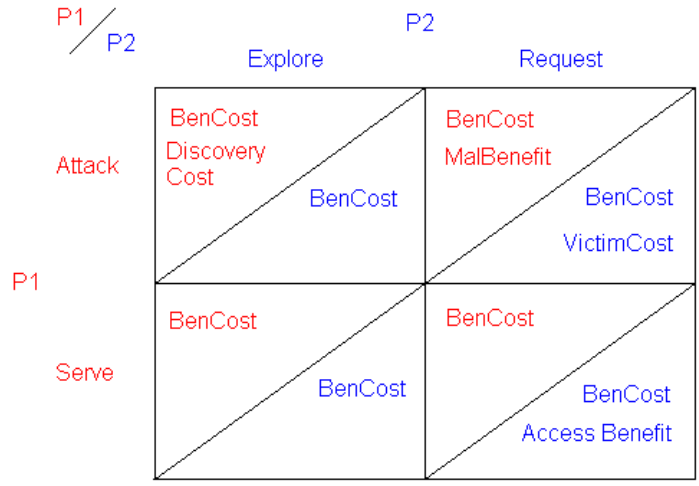

Figure 1. Payoff Matrix for a Benign Peer and a Malicious Peer

equals one, which will never result in a resource access. In order to overcome these two shortcomings, we also look at the Res-Exp mechanism from the perspective of selecting a rate based on maximizing our own utility and producing a preference bound on what attackers will benefit from attacking.

4.3. Utility Driven

The utility of a benign peer and that of a malicious peer are modeled by Equations 11 and 12 respectively where P_{exp} is the probability of sending an exploratory message.

$$Utility_{benign} = Access - \frac{BenignCost}{1 - P_{exp}} - (1 - P_{exp}) \times VictimCost \quad (11)$$

$$Utility_{attacker} = (1 - P_{exp}) \times MaliciousBenefit - \frac{MaliciousCost}{1 - P_{exp}} - BenignCost \quad (12)$$

Input: Peer Preferences
Output: Resource Access Results
while *Resource Not Accessed* **do**
  Calculate P_{exp};
  Generate Request (P_{exp} are exploratory);
  Send Request;
  if *Request Is Exploratory* **then**
    if *Attacked* **then**
      Blacklist Attacker;
    end
  end
end

Algorithm 1: Algorithm for Sending Exploratory Requests

Algorithm 1 utilized the relationships listed above in order to implement a Res-Exp mechanism. By examining these relationships we can show the effects of exploratory messages as in Figure 2. Maximum $VictimCost$ is the plot that shows what a peer's maximum victim cost can be to still expect positive utility from an interaction with a given exploratory message rate. Maximum Attacker $MaliciousBenefit$ is the plot that shows the maximum $MaliciousBenefit$ that the given exploratory message rate would return a negative expected value if an attacker was to act maliciously.

We produce an example application of this mechanism in Figure 2. As we cannot generically speak for all peers that may possibly enter a system, we make some reasonable

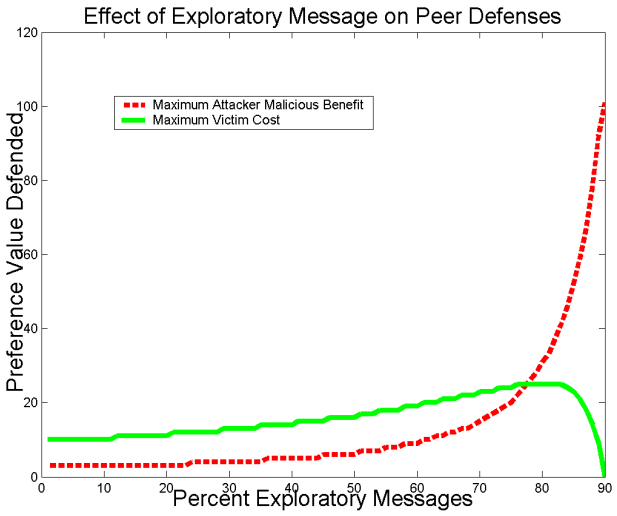

Figure 2. Effect of Exploration Message on a Benign and an Attacking Peer

assumptions in order to provide this example. The figure is produced assuming simple values of $BenignCost$ and $MaliciousCost$ as 1 and a $BenignBenefit$ of 10 in order to provide an example. In other words, the cost of participating in the system is the same for both the attacker and the benign peer (for example, the each peer consumes one unit of energy). The benefit that a benign peer receives from a successful resource access is 10 times the cost of participation (for example, accessing the resource remotely saves the benign peer 10 units of energy).

In the case of this example, the benign peer will maximize its own expected value, given no knowledge about its potential attackers by exploring at a rate of P_{exp} equal to 82%. Similarly, the relationships show that a rational attacker would not find it beneficial to attack the peer if its $MaliciousBenefit$ for the attack was less than 40. While not all sets of preferences will produce the same results, the form of the graph will be similar and still provide a set of bounds and guidelines for a peer to use to determine at what rate to send exploratory messages, and what protective bounds exist. A benign peer playing the utility-based strategy at its highest point is guaranteed the minimum vulnerability to attacks while still acheiving positive utility.

Since a malicious peer will be discovered by an exploratory message, if it is not aware that the peer it is attacking is utilizing the Res-Exp mechanism, it will reveal itself on the first exploratory request, thus mitigating attacks on peers recently entering the system and allowing such peers to use the system confidently than when using a reputation mechanism. Furthermore, the peer will continue to use the Res-Exp mechanism, so any one-time or compromised peer attacks will also be mitigated because of the high likelihood of being discovered. If the malicious peer knows that the Res-Exp mechanism is being used, then it will have to probabilistically increase its benign service in order to not be discovered, which increases resource availability in the system and benefiting all benign peers.

Online route finding would be an example application where the Res-Exp mechanism would apply well. For instance, Alice is visiting a new campus and wants to find the the room where she is supposed to be giving a presentation from her current location, the University Center. She uses her PDA to contact a route service in the University Center, but is uncertain if the service is reliable (or even if it is malicious). She has several known routes on campus already stored in her PDA, so her PDA employees a resources exploration mechanism in order to test the reliability of the service and reduce her risk of receiving bad results. In this case, the cost of making several requests is minimal relative to the cost of being attacked (going to the wrong building or room and missing her presentation), which is the type of application in which Res-Exp is most appropriate.

5. Future Work and Conclusions

There are many exciting areas of future works that can stem from this project. As mentioned earlier, the work with Nash equilibria is only useful if we know the preferences of the opponent. In auctions and multi-agent systems, there is already work being done in preference elicitation [11],

but similar work could be expanded in P2P systems to find Nash equilibria for security. In particular we have already begun work in this area. We have designed an architecture that implements the Res-Exp mechanism in a P2P system. In this architecture, we begin by sending exploratory message at the rate that will minimize attacks, but still leave the peer with positive utility or potentially least loss of utility. Our further work in this area will involve approximating the Nash equilibrium described in Section 4.2. In particular, we will be analyzing techniques for learning the Nash equilibrium and showing the resulting performance against many different types of attackers.

Using utility functions and game theoretical approaches discussed in this paper, we can analyze the security mechanisms in existing systems to discover exploits in these systems and to design new mechanisms to prevent those exploits. In particular, we have begun to develop a framework for quantitatively analyzing and comparing reputation mechanisms through our proposed utility relationships.

We also have begun analyzing the use of individual mechanisms, such as resource exploration, in combination with cooperative mechanisms such as reputation mechanisms. This research will provide insight into the appropriate methods for integrating the two approaches and providing even greater security benefits in P2P systems.

In this paper we have introduced a utility model for evaluating security in P2P systems. We have also introduced a security mechanism, Res-Exp, for individuals in untrusted environments and analyzed its usefulness based on peers' utility functions. Res-Exp is a mechanism that can be implemented individually without the cooperation of other peers in a system. As a result, it mitigates the problems inherent and often unaddressed by reputation mechanisms. The proposed unique approach to P2P security is applicable to mobile systems where peers roam into uncertain and unfamiliar environments. It is particularly suited to applications in which the cost of making exploratory requests is significantly less than that of being attacked. Res-Exp mechanism can provide benefit to peers and mitigate vulnerabilities in reputation mechanisms.

References

[1] Figueirdo, Shapiro, Towsley. Incentives to Promote Availability in Peer-to-Peer Anonymity Systems. Proceedings of the 13TH IEEE International Conference on Network Protocols, 2005.

[2] Kamvar, Schlosser, Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. In Proceedings of the Twelfth International World Wide Web Conference, 2003.

[3] Chuang. Economics of Peer to Peer Systems. Academia Sinica 2004 Summer Institute on P2P Computing, 2004.

[4] Feldman, Lai, Stoica, Chuang. Robust Incentive Techniques for Peer-to-Peer Networks. In the Proceedings of EC, 2004.

[5] Xiong, Liu. PeerTrust: Supporting Reputation Based Trust for Peer-to-Peer Electronic Communities. IEEE Transactions on Knowledge and Data Engineering, Vol 16, No. 7, 2004.

[6] Srivatsa, Xiong, Liu. TrustGuard: Countering Vulnerabilities in Reputation Management for Decentralized Overlay Networks. In Proceedings of WWW, 2005.

[7] Nisan and Ronen. Algorithmic Mechanism Design. In The Thirty First Annual ACM symposium on Theory of Computing, pages 129-140, May 1999.

[8] Anderson and Moore. The Economics of Information Security. Science 314, pp 610-613, October 27, 2006.

[9] Song, Hwang, Zhou, Kwok. Trusted P2P Transactions with Fuzzy Reputation Aggregation. IEEE Internet Computing Magazine, Nov/Dec 2005.

[10] Suryanarayana, Erenkrantz, and Taylor. An Architectural Approach to Decentralized Trust Management. IEEE Internet Computing Magazine, Nov/Dec 2005.

[11] Chen and Pu. Survey of Preference Elicitation Methods. A Technical Report. 2004.

[12] Kevin Walsh, Emin Gn Sirer. Experience With A Distributed Object Reputation System for Peer-to-Peer Filesharing. In Proceedings of the Symposium on Networked System Design and Implementation (NSDI), San Jose, California, May 2006.