

Incentive and Trust Issues in Assured Information Sharing

Ryan Layfield, Murat Kantarcioglu, and Bhavani Thuraisingham

Computer Science Department, University of Texas at Dallas
{layfield,muratk,bxt043000}@utdallas.edu

Abstract. Assured information sharing among different organizations in a coalitional environment is an important first step in accomplishing many critical tasks. For example, different security agencies may need to share intelligence information for detecting terrorist plots. At the same, each organization participating in the assured information sharing process may have different incentives. In this paper, we explore the effects of different incentives and potential trust issues among organizations on the assured information sharing process by developing an evolutionary game theoretic framework. In addition, we provide extensive simulation analysis that illustrates the impact of various different information sharing strategies.

1 Introduction

Many current challenges require different organizations to share critical information. Defending against the threats of international terrorism presents this scenario in an alarming manner. Countries that are concerned about an attack on native soil must frequently consider assailants that have collaborators which span the globe. A single country rarely has the necessary resources and jurisdiction to continuously investigate every possible suspect. Even if the resources are available, unless attacks are expected to have a considerable impact, it is not likely that the investment is cost-effective. Thus, a variety of agreements have been created in the course of history to attempt to unite multiple governing bodies under a common threat by exchanging information with their allies.

Unfortunately, in the scope of international politics, the only governing factor that ensures members of an alliance will always cooperate fully lies in their individual incentives. Unless information can be verified, there are no guarantees that the information supplied will be truthful. In addition, in some cases, there exists the potential for increasing gain from an information exchange by presenting false knowledge. In such cases, if the other party's knowledge is truthful, and the malicious party is not caught, a one-sided gain could occur. Thus, one of the biggest challenges in such endeavors is how to encourage honest assured information sharing.

The study of game theory deals directly with the motivations of participants, known as players, attempting to achieve some known goal and the choices they must make to do so. Out of all options available, game theory assumes that each participant wishes to maximize their own personal benefit in a rational manner. At any given

time in an information exchange, both participants have the option of telling the truth, providing false data or providing no data. While it may seem obvious that all parties would collectively benefit from the truth, each individual is often only concerned with their own gains [14]. If that gain comes at the expense of another participant with no threat of retribution, there is little encouragement to do otherwise.

However, when games are repeated, new constraints begin to emerge on a player's strategies. If a participant chooses to lie, they run the risk of being caught, leading to a potential net loss. When a central authority can observe actions and affect the payout a player receives, enforcement of the agreement often becomes a simple matter of finding an appropriate punishment. When considering agreements between multiple sovereign governing entities, there is not necessarily any central authority that can enforce such punishments. In such cases, the burden of ensuring an ideal situation is created is shifted to the collective actions of the group.

One option a player has within such a scenario is to simply refuse to participate. This can include all members, or just a selection of those that are not giving the desirable responses. If one player has information that is highly desirable to the rest, with little dependence on others, they can potentially influence the choices by the entire group. On a level playing field, where no player has information that is significantly more valuable, a single player which no longer communicates with the rest can be sacrificed with little trouble. Clearly, collective action must be taken by a significant number of participants to have any effect on group. Thus, several players must be willing to isolate the same player with undesirable behavior in the hopes that the malicious participant will change their ways.

Another more indirect method of enforcing behavior is to base punishment indirectly on the level of trust shared by the players. Normally, each player already has opinions of the rest, but they lack a broader view of the situation and must assume that how a player deals with them is how they deal with everyone else. Eigentrust [22] provides a means of collectively allowing each individual player to form accurate opinions of the group by providing ratings of their own experiences. Such opinions can be gathered in a distributed manner and provide the rough equivalent of an 'omniscient' view of the rest.

Several of these factors have already been explored in other works (e.g., [21]). However, a factor that has not always been considered is the cost of determining whether or not information is correct. While there is certainly data that can easily be verified, that nature of certain kinds of information requires a much more in-depth review. Thus, the cost of verification should always be considered in situations where the net gain of the interactions is paramount to the success or failure of an exchange.

The goal of our work is to explore the potential of punishment via isolation with regards to the introduction of trust computations. We wish to determine whether or not such methods are viable in large games with multiple players with different incentive structures, and consider various scenarios that such logic may face. Success of such a method would prove useful in a variety of decentralized assured information sharing platforms.

Our paper is organized as follows. Section 2 describes existing work on the subject, including our own. Section 3 describes our model. Section 4 discusses how we setup our experiments while section 5 details our results. Finally, we share our own conclusions on the subject in section 6 and consider future directions of the field.

2 Related Work

One closely related area to our current work is the research on incentive issues in peer-to-peer file sharing networks. Within these file sharing systems, independent players join and leave at their leisure, seeking to download a file or files with the help of other participants. Problems arise when a new participant joins the network and download a resource from other peers and never actually contribute to the group. This process, known as leeching, has been a large problem in piece-meal file sharing protocols such as the popular BitTorrent. The work of Gupta et al. [20] and Buragohain et al. [19] both deal with this behavior by creating a system of incentives for further contribution.

In our previous work [21], we provide a simpler model that excludes distributed trust from behavior influence and focuses on how individual agents indirectly can contribute towards achieving a common goal in a virtual game of assured information sharing. Agents were permitted to shift their behavioral choices to more effective means through analysis at a fixed interval. The LivingAgent (more on this later) was introduced as a competitive alternative to the Tit-for-Tat strategy, and the former worked well in the constraints provided. However, it still suffered from a need for ‘critical mass’ of the behavior’s presence before it was effective.

A great deal of research has dealt with trust in the realm of distributed systems. Information exchange methods have been particularly useful in the formation of ad-hoc networks. Seredynski et al. [1] used the concept of an evolving genetic algorithm to enhance security and trust in a wireless network among multiple nodes relaying data packets. Other works have taken a purely game-theoretic approach to trust. The work of Cascella [2] performed analysis of an infinitely repeated form of the prisoner’s dilemma with regard to persistent players randomly selected. They found the introduction of a reputation system allowed punishment systems based on discriminating between good and bad reputations succeeded as long as players were sufficiently patient to achieve the results. Other works have attempted to apply peer-to-peer trust in scenarios involving military joint cooperation in the field [3]. Given a military body as a dividable resource, they attempt to address the problem of resource allocation in situations where either a central authority is not robust enough or the resources span multiple international owners. The core issues they addressed were dealing with malicious reports attempting to sabotage trust ratings and attempting to give more control to an agent’s own rating.

3 Our Model

Our previous work on the subject of behavior enforcement in distributed information sharing focused solely on the feasibility of the verification and punishment process. Several experiments proved conclusively that the LivingAgent behavior outlined in [21] both helped to eliminate malicious agents and eventually become the dominate strategy in an adaptive game. We plan to build upon the success of that agent by integrating distributed trust metrics into the decision making process.

The scenario for our information exchange strategies is based on a loose alliance with no central authority to enforce behavior. Consider multiple nations that have

learned of an impending terrorist attack. They do not have conclusive data to suggest when or where the attack will occur, but each country has reason to believe it may occur on their own soil. Their objective is to attempt to thwart the current threat. However, given an indefinite time span in which the attack could occur and limited resources, they have each determined they must ask other countries for help. After discovering each of them had a common goal, they have formed an alliance in which they exchange information they have collected both at home and abroad. The nature of the game is one which occurs repeatedly for an indefinite amount of time.

Information is exchanged between members of the alliance individually at a regular interval. The transaction occurs between two countries in such a way that the data is swapped simultaneously; both countries must decide on their strategy before the transaction is complete. Each of these transactions occurs between all possible pairs of countries simultaneously, assuming each pair agrees to do business. The value of the information fluctuates within predictable boundaries, and no player has a considerable advantage over the rest.

Each player faces that challenge that they do not know of the kind of behavior the members will engage in. While all countries involved are assumed to have a common goal, they may also see an opportunity to advance other political statements. For example, one country may wish to keep what they know a secret from the rest, in the hopes of learning more at no real cost.

The strategies chosen by each country is determined by the overall behavior they have chosen. Each country wishes to find the optimal strategy to reduce the impact of defense on their national budget. We assume thus that countries are willing to adapt by altering the behavior to reflect the one they believe has performed the best. At the same time, as behaviors shift, the payouts of strategy choices may shift as well, leading to a dynamic balance of power. For example, a behavior to always lie may perform well when other countries are not verifying, but as others learn of the benefits of the behavior, others may follow suit. This would result in several liars always lying to each other and never gaining any information.

Determining whether or not the data is received is legitimate is the responsibility of the country itself. Since the data is primarily intelligence, verifying it has a substantial cost due to the resources, manpower, and time required. In our scenario, verification is always less than the value of the information itself, which means consistently doing so will still result in a net gain. However, it is not necessarily the most efficient.

The use of adaptation to improved behaviors within a game raises an interesting point about the duration of punishment. One option is to punish a deviating agent indefinitely. When this is done, any future benefit from that agent is simply not possible, potentially allowing more forgiving agents to flourish. Instead, punishment in our game is done in such a way that the other player simply loses a significant amount of their own potential earnings, reducing their net gain from the game. This indirectly makes the ideal behaviors much more likely to be chosen. Eventually, if this is practiced widely enough, overall behavioral choices yield an ideal environment where everyone can benefit. Likewise, when we have a fixed interval when agents may take the opportunity to change behaviors, we would potentially discourage what may otherwise be an excellent source of profit; thus, forgiveness must be performed by all agents during this round. An example of this would be when a government agency has a new leader or a business comes to the end of a fiscal quarter.

Determining reputations within a distributed network can be a difficult endeavor. Since it is possible for a malicious participant to deal honestly with some players and dishonestly with others, a trust value must extend beyond a local perspective. This necessitates querying others for their opinions on opponents within the game, which introduces the possibility of the same malicious agents simply telling others they have an outstanding rating while their peers have just the opposite. This introduces the additional possibility that different players will come to separate conclusions, based on the ‘noise’ introduced by the subversion. Sepandar [22] et al. devised the EigenTrust algorithm as an answer to these problems.

The algorithm itself is relatively straightforward. Each player queries every other player for their opinion on the rest. This forms a matrix of relative trust, based on a score built from history among individual agents. From here, a normalized matrix is constructed, then evaluated with the Eigenvalue Decomposition technique. When all players perform this properly, they will all come up with the same left-principle eigenvector. This vector represents the EigenTrust rating of each player. The algorithm has been well received as a foundation for more robust distributed systems, though trust itself needs further refinement to be properly defined [7]. In real distributed system deployments, EigenTrust would be done in a distributed fashion [22].

Table 1. Utilities for various actions

		Play (agent j)		Do Not Play
		Truth	Lie	
Play (Agent i)	Truth	$\left(\frac{\delta_{\min} + \delta_{\max}}{2}\right) - C_v(P_i + t_j)\sigma_i$	$-C_v(P_i + t_j)\sigma_i$	0
	Lie	$\left(\frac{\delta_{\min} + \delta_{\max}}{2}\right) - C_v(P_j + t_i)\sigma_j$	$\left(\frac{\delta_{\min} + \delta_{\max}}{2}\right) - C_v(P_j + t_i)\sigma_j$	0
Do Not Play		0	0	0

Given these ingredients, the basics of each round of our game can be described with an immediate snapshot of the game matrix. There are essentially three choices every player can choose: lie, tell the truth, or stop playing with the other player. The potential benefit from the truth is an average of δ_{\min} and δ_{\max} , the upper and lower bounds of what the information is worth. A lie of course carries no value, but checking as to whether a piece is legitimate or not does carry a cost. The expected cost of verification is determined by the cost of verification C_v times the probability that verification will occur coupled. The probability of verification is determined by three factors: (i) the type of the agent i (σ_i) (ii) the minimal probability of verification for agent i (P_i) (iii) the EigenTrust value t_j assigned to the opponent j . However, certain

behaviors never consider verification as a possibility, regardless of trust, and as such the verification has no effect on the result. For example, a player may want to accept the provided information without any verification. In that case, we can set the σ_i as 0. The table 1 summarizes the payoff matrix for each information exchange transaction.

3.1 Behaviors

In our analysis, we considered various different types of agent behaviors. The Honest behavior takes a naïve approach to other players. Truth is the only strategy ever chosen, and it never verifies the strategies of other players. It has the advantage of never incurring the cost of verification, and it always maximizes the potential gains with other players by never severing the links. An example may be a country that wishes to set an example, or perhaps is simply under significant amounts of scrutiny. While this may prevail against behaviors that perform even the slightest verification, they will always lose in a competition with the Dishonest behavior. Essentially serving as the opposite of the Honest behavior, this behavior simply chooses the Lie strategy regardless of the outcome.

Not every player may believe that a predictable behavior is optimal. The Random behavior picks either Truth or Lie with equal probability. No punishment or verification is ever performed. Countries that wish to avoid being anticipated may choose this strategy. It carries the same benefits as the Dishonest behavior, but potentially only gains at most half the benefit.

In our prior work [21], we encountered a unique yet simple approach to dealing with undesirable behavior known as the Tit-for-Tat behavior. Devised by Anatol Rapoport [13, 14], it follows a simple strategy selection process. Initially, it selects the desirable strategy (Truth). From that point on, it simply selects the same strategy as its' opponent within the game. This was proven quite effective against all but the most sophisticated collaborative opponents [23]. In our simulation, however, being able to mimic the opponent's actions requires constant verification.

Our devised behavior from our prior work is the LivingAgent. Initially, like Tit-for-Tat, the Truth strategy is chosen. During each transaction, there is a probability P_i that the player will verify whether or not the other player told the truth. If a lie was told, the other player is punished by severing the link for R_S rounds. This is a sacrifice in the sense that, if the other player is telling the truth for at least part of the time, further opportunity will be lost. The goal with this behavior is to place a high price on deviation. A country which behaves in this manner may be attempting to send a message to the rest of the participants, or may simply be unwilling to waste time with participants whom are equally unwilling to share valid information.

A variant on this behavior is the SubtleLiar, which obeys the same principle but has a threshold in which it will automatically choose the Lie strategy. The net effect is a behavior which can take advantage of a low P_i and net a slightly larger gain in information. An example of this behavior may be countries that believe their fellow members trust them enough to be taken advantage of.

Finally, we have the Liar. This behavior is almost identical to that of LivingAgent, with one notable exception. Is called Liar simply because it essentially tries to pass itself off as a honest participant while consistently trying to take advantage of the right situations. Assuming that the value of the information about to be received is

known in advance by both, the Liar will always lie if the received value is within a certain threshold of maximum value. Thus, this agent appears to capitalize on advanced information by attempting to only take a risk when the gain appears to be sufficient. This threshold is determined by the constant δ_{valuable} .

4. Experimental Setup

Our experimental setup involves creating an alliance of 100 virtual countries. They begin with equal levels of trust, and hold all of their peers in the same regard. Each experiment begins by distributing initial behaviors based on a configuration file. The behaviors are assigned to each player based on the distribution specified within the file. During each round of the simulation, a transaction is executed between all possible combinations of players through a virtual link. After this round, the trust metrics are updated, history files are recorded, and agents receive their payoff based on the results. This payoff is used to directly determine the performance of the agent itself. The value of the information varies between $\delta_{\text{min}} = 3$ and $\delta_{\text{max}} = 7$. If the information provided is false, it has no value. In the event verification is performed, it comes at a cost C_v , set to a value of 2. Thus, even if all information is verified, a net gain is still possible. When verification occurs, and no lie was told, it is noted as waste. The overall score a player receives is simply the total value of all information sans any cost incurred. The threshold for Liar to lie is at 6.9.

All players are assumed to be willing to change their behavior to a more effective one, based on the performance of their neighbors. To simulate this, every 5,000 rounds, each agent is assigned a new behavior based on a weighted probability assigned based on the total gain achieved. For example, if the Honest agent has an total net payoff of 10,000, while Dishonest has a total net payoff of 20,000, each agent has a 33% chance of choosing Honest and a 67% chance of choosing Dishonest. Thus, the new distribution reflects the relative performance of all agents. Note that while this method does not necessarily guarantee an ineffective behavior will be eliminated, it will ensure that any effective choice will be much more likely to emerge victorious. The simulation ends when either 100,000 rounds have passed or all players have chosen the same behavior, the latter considered a 'win' by the behavior adopted by the rest.

There are several verification rates possible for a LivingAgent-derived behavior. To ensure a larger search space is explored during the experiment, a small mutation rate is introduced on P_i for each player behavior that is copied. This allows players to adapt over time and consider reducing the potential waste as the system approaches an equilibrium in which all participants tell the truth.

5 Experimental Results

The LivingAgent performed admirably against Tit-for-Tat, even when the latter behavior began in the experiment with twice as many players adopting it. Out of all experiments, neither the Dishonest nor the Tit-for-Tat behavior ever successfully became the dominant behavior. However, the LivingAgent behavior only won the

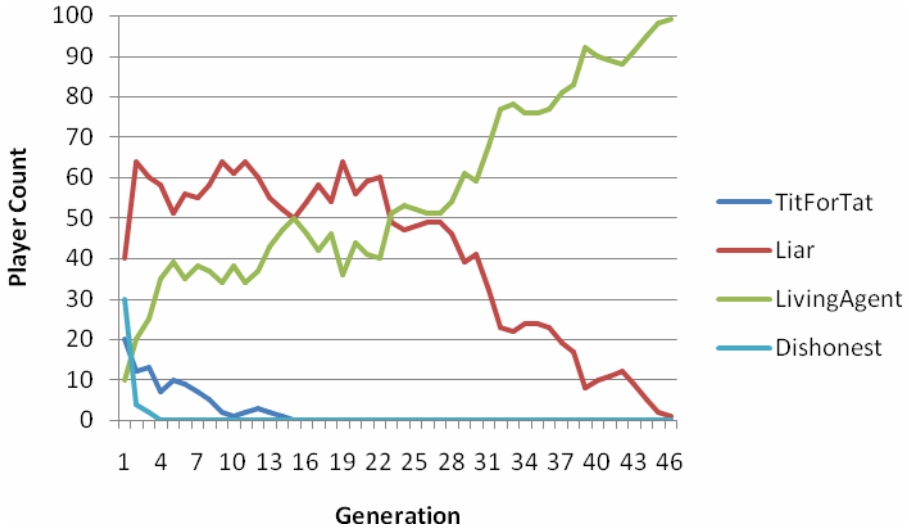


Fig. 1. Success of various strategies against LivingAgent

game 87% of the time, while Liar won the rest. The average verification rate was 14.8% in the final behavior tally, while the average standard deviation was only 2%, taking an average of 15 generations to declare a winner.

In figure 1, we see that LivingAgent successfully achieved the majority behavior despite beginning as a much smaller population. At first, Liar benefits from the fact that it and the rest of the behaviors engage in punishment. Those that leave the Dishonest behavior go towards both LivingAgent and Liar, with more towards Liar. However, Dishonest is no longer being used, Tit-for-Tat begins to help LivingAgent by working slightly against Liar. A trend in the results is that once LivingAgent achieves half of the population, the Liar behavior rapidly loses ground to the point of complete loss. This appears to be the critical mass for the behavior in such a situation.

When Tit-for-Tat was able to sustain itself for at least 10 generations, LivingAgent often benefited from this indirectly. The Liar players, even when they initially surged ahead, would usually observe that Tit-for-Tat was a better choice. As Tit-for-Tat increased, LivingAgent took some losses, but the efforts ended up working in concert to reduce the threat of a lying behavior. Ultimately, once Tit-for-Tat was no longer in play, LivingAgent needed only to compete with a smaller pool of malicious players.

Another trend that made itself apparent is that the verification rates of the LivingAgent did not correlate directly with the number of malicious agents present. In many instances, as the number of Liars increased, the average rate continued to drop. In these cases, it appears that players simply cannot afford to refuse doing business with other players that ran the risk of lying. Essentially, the punishment method ended up only punishing the enforcer.

In order to demonstrate the effectiveness of the LivingAgent approach, we ran experiments involving all behaviors not derived from it (Figures 2 and 3). The behaviors here were only Dishonest, Honest, Tit-for-Tat, and Random. The first thing we noticed is that there was automatically a large increase in the number of iterations




Fig. 2. Evolution of various strategies without LivingAgents: Honest behavior winning case




Fig. 3. Evolution of various strategies without Livingagents: Random agent winning case

necessary. The winner was not always clear, and it appeared that the fluctuation in the payoffs alone caused some agents to benefit more than others. The biggest competitor to Honest was strangely Dishonest, even though both used the polar extremes of selecting a strategy, and neither made any efforts to check information validity.

Another issue that arose was how quickly Tit-For-Tat was eliminated at times within the first generations. This, however, was no surprise; the verification costs no doubt came at a high toll to the payouts. It contributed to the game by removing malicious behaviors consistently, but this often only resulted in a surge of Honest behavior adoptions. Once Tit-for-Tat was eliminated, malicious behaviors again rose substantially in numbers.

Figure 3 shows an experiment where we find that the Random behavior has won. Again, Tit-for-Tat helps Honest surge ahead briefly, but the verification costs cause it to fail to function after only 7 generations. From that point on, Honest appears to be the certain victor, eliminating Dishonest. At generation 92, Random begins to succeed. Essentially, because Random does not discriminate against which players it lies to, it runs the risk of dropping off at just shy of 50% of the player market. This happens twice during the simulation, but due to fluctuations in information value, it eventually achieves victory. Note that this particular equilibrium took 197 generations to achieve, and it was primarily based on the delta in the value gained.

There were only three winning behaviors out of all the experiments. The Honest behavior only achieved the majority 26.1% of the time. Out of those instances, only 84% of them actually resulted in an equilibrium behavior. The Random behavior faired equally well, with only a slightly smaller number of wins at 23.3%. However, the Dishonest behavior beat both of them twice as often at 53.3% as either of them.

In reality, players within these games can vary widely in their ulterior motives, beliefs, and decisions. To observe this, we wanted to observe all of the devised behaviors in action (Figure 4). Unsurprisingly, Honest won 99% of all games played. The abundance of malicious agents, coupled with a roughly 3:1 starting ratio to the Honest behavior, allows LivingAgent to flourish briefly. However, as malicious

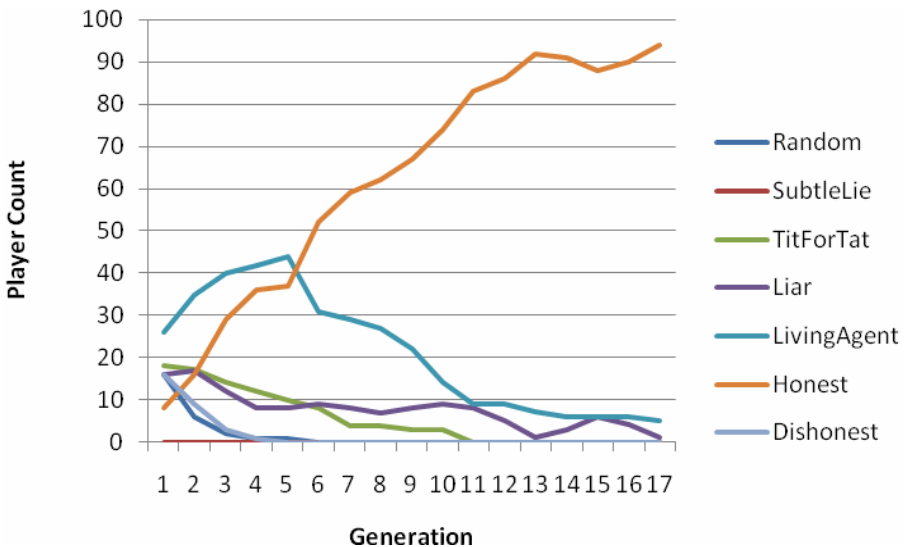


Fig. 4. Evolution of all possible strategies together

agents begin to disappear due to lack of relative performance, the appeal of the Honest approach eventually coerces a majority of agents away from it. The loss of Dishonest and Random predictably caused a 25% drop in.

On the surface, it may appear that our own LivingAgent is a failure under these circumstances. If our goal was simply to find the 'perfect' behavior, this would indeed be true. However, our constructed behavior helps to create an environment in which the Honest behavior can flourish. Thus, the end result is still the ideal, truth-telling environment. Since previous results demonstrated that Honest would normally fail against the malicious competitors, the introduction of our behavior has acted as an indirect policing force within the system. The Honest behavior achieved the majority roughly 97% of the time within our initial experiments, over three times what it achieved under similar conditions on its' own.

Additional experiments were performed to observe the minimum behavioral mix necessary to ensure Honest would succeed, performed in the form of ratios between LivingAgent and Honest. When equal parts of both behaviors were present, Honest won 86% of the time. Increasing the ratio of LivingAgent to Honest to 2:1, it increased to 92% of the time. At a ratio of 3:1, an effectiveness of nearly 100% was achieved. Thus, although LivingAgent benefited the rest of the group in achieving a truth-telling majority, significant numbers of LivingAgent were needed to guarantee it.

6. Conclusions

The overall experiment was a relative success. When enough players choose a behavior that reflects our approach to punishment, the malicious behaviors were successfully eliminated from consideration. The underlying nature of LivingAgent allowed it to defeat even variants of its' own behavior involving light amounts of deviation. However, the same nature of the persistent verification meant that the behavior did not succeed against the Honest behavior, which performed no verification whatsoever despite the circumstances. Even in these circumstances, the ideal situation still arose, allowing players to conclude that honesty is indeed the best choice.

We hope to increase the robustness of our punishment method in scenarios closer to reality. In real life, information cannot always be verified with 100% accuracy, nor do even the best intelligence agencies guarantee that information provided will be completely true. Such inadvertent mistakes would result in a system which potentially punish otherwise trustworthy players. This can be addressed with a mixture of a higher tolerance for lies and a slightly more relaxed punishment. Our future work will explore this problem in depth.

The growing size of networks such as the internet and the increasing use of distributed systems suggest that centralized authority approaches will be insufficient. Insuring behavioral choices by members of peer-to-peer networks requires an approach which can scale as much as the system itself. We believe our work offers a solution to the problem of encouraging behavior when players become responsible for their own outcome.

References

1. Seredynski, M. Bouvry, P. Klopotek, M.A.: Modelling the Evolution of Cooperative Behavior in Ad Hoc Networks using a Game Based Model. *Computation Intelligence and Games*, 96–103 (2007)
2. Cascella, R.G.: The “Value” of Reputation in Peer-to-Peer Networks. In: *Consumer Communications and Networking Conference*, 516–520. (2008)
3. Oh, J.C., Gemelli, N., Wright, R.: A Rationality-based Modeling for Coalition Support. *Intelligent Systems*, 172–177 (2004)
4. Li, X.: A Grassroots Approach in P2P Reputation Studies. In: *Distributed Computing Systems Workshops*, p. 234 (2008)
5. Yi, X., Han, J., Yu, P.: Truth Discover with Multiple Conflicting Information Providers on the Web. *IEEE Transactions on Knowledge and Data Engineering*, 796–808 (2007)
6. Eckel, C., Wilson, R.: The Human Face of Game Theory: Trust and Reciprocity in Sequential Games. In: *Trust Working Group Meeting* (1999)
7. Morselli, R., Katz, J., Bhattacharjee, B.: A Game-Theoretic Framework for Analyzing Trust Inference Protocols. In: *Workshop on Economics of Peer-to-Peer Systems* (2004)
8. Cascella, R., Battiti, R.: Social Networking and Game Theory to Foster Cooperation, http://www.enisa.europa.eu/doc/pdf/Workshop/June2007/Papers/reputation/REP_UniversityTrento.pdf
9. De Paola, A., Tamburo, A.: Reputation Management in Distributed Systems. In: *2008 IEEE Symposium on Game theory, evolutionary approach, distributed systems* (March 2008)
10. Song, S., Hwang, K., Zhou, R., Kwok, Y.: Trusted P2P Transactions with Fuzzy Reputation Aggregation. *IEEE Internet Computer*, 24–34 (November 2005)
11. Agarwal, N.: Equilibrium Game Theory Under the Conditions of Repeatability. In: *Proceedings of the 10th International Conference on Extending Database Technology*, pp. 240–256 (2006)
12. Andrade, N., Mowbray, M., Lima, A., Wagne, G., Ripeanu, M.: Influences on cooperation in BitTorrent communities. In: *Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pp. 111–115 (2005)
13. Axelrod, R.: *The Evolution of Cooperation*. Basic Books, New York (1984)
14. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press, Cambridge (1991)
15. Mahajan, R., Rodrig, M., Wetherall, D., Zahorjan, J.: Experiences applying game theory to system design. In: *Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems*, Portland, Oregon, pp. 183–190 (2004)
16. Monderer, D., Tennenholtz, M.: Distributed games: from mechanisms to protocols. In: *Proceedings of the sixteenth national conference on Artificial intelligence and the eleventh Innovative applications of artificial intelligence conference innovative applications of artificial intelligence*, Orlando, Florida, pp. 32–37 (1999)
17. Myerson, R.: *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge (1991)
18. Riolo, R., Worzel, B.: *Genetic Programming Theory and Practice*. Kluwer Academic, Boston (2003)
19. Buragohain, C., Agrawal, D., Suri, S.: A Game Theoretic Framework for Incentives in P2P Systems. *Peer-to-Peer Computing*, 48–56 (2003)
20. Gupta, M., Judge, P., Ammar, M.: A Reputation System for Peer-to-Peer Networks. In: *Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, Monterey, CA, pp. 144–152 (2003)

21. Layfield, R., Kantarcioglu, M., Thuraisingham, B.: Enforcing Honesty in Assured Information Sharing Within a Distributed System. In: Barker, S., Ahn, G.-J. (eds.) *Data and Applications Security 2007*. LNCS, vol. 4602, pp. 113–128. Springer, Heidelberg (2007)
22. Kamvar, S., Schlosser, M., Garcia-Molina, H.: The EigenTrust algorithm for reputation management in P2P networks. In: *Proceedings of the 12th international conference on World Wide Web*, Budapest, Hungary, pp. 640–651 (2003)
23. Grossman, W.: New Tack Wins Prisoner's Dilemma. *Wired Magazine*, <http://www.wired.com/culture/lifestyle/news/2004/10/65317>