On bidder zones, cell phones and ballot stuffing

Wesley Brandi and Martin Olivier

Department of Computer Science, University of Pretoria, Pretoria

Abstract. Ballot stuffing is the process of illegitimately creating a good reputation and requires a number of colluding fraudsters. With the use of an anonymising network, a single user could pose as many separate users and possibly implement the ballot stuffing scheme without the complexities of managing multiple fraudsters. In this paper we propose a framework that deals with thwarting ballot stuffing within the context of anonymising networks.
The framework proposed depends on the ability to look up a single user's physical location with a fair degree of certainty. With this mechanism in place, the framework is able to compare the locations of all users involved in a transaction. If the coordinates are within close proximity to one another (this is unlikely in a global trading system) then the transaction is tagged accordingly.
The framework proposed does not necessarily apply to only thwarting ballot stuffing schemes, we discuss the fundamental problem of using anonymising networks to commit ballot stuffing and analyse another scheme (ticket scalping) that may be thwarted with this framework.

1 Introduction

As the Internet increasingly becomes more a part of our daily lives, it comes as no surprise that many people now depend on the Internet as a means of conducting trade. Where there was once a stall in a local flea market, there is now an online equivalent in a global trading community the likes of eBay. Since the normal social interactions traders depend upon is not possible in such an environment, Reputation Systems play an integral part of deciding whom traders should trust when conducting business and to what degree. Unfortunately, as much as fraud is a problem in real world trading environments, it threatens to topple the very foundation upon which online trading systems have been built.

One such threat is that of ballot stuffing. This involves fraudsters who collude with a network of buyers to provide them with positive feedback after conducting false business transactions. The feedback provided serves so as to increase the online reputation of the fraudster. In doing so, the fraudster can be rest assured that he will do better business than those with reputations that are not as competitive.

Ballot stuffing, although a simple scheme, requires the collaboration of a number of buyers in order to be successful. The management of these buyers (and payment thereof) has the potential to be a complex endeavour. For the

buyer looking to branch out on his own though, there is an alternative. By camouflaging himself in an anonymising network, a buyer may have the means to setup a false network of buyers and significantly reduce the overhead of the ballot stuffing scheme.

The intention of this paper is as follows: we propose a framework that promises to significantly reduce ballot stuffing through anonymising networks. The framework is dependent upon an inexpensive way of looking up the physical location of a potential buyer in an online sale. Cell phones are proposed as a potential means of achieving this. The physical location of online bidders is looked up and compared to every other bidder involved in the transaction. The entire transaction may be tagged as suspicious if any bidder falls within one of three physical zones of another bidder (either for the same transaction or even another transaction from the same seller).

This paper is structured as follows: In section 2 we briefly examine and discuss related work in the fields of Reputation Systems, ballot stuffing and anonymising networks. Where section 3 discusses what would be required to setup false buyers within the context of an anonymising network, section 4 proposes the framework that would be implemented to thwart such a scheme. Section 5 highlights the problem of masquerading through anonymising networks and discusses the application of the framework to thwart ticket scalping. This paper is then concluded in section 6

2 Background

2.1 Reputation Systems

Since the normal elements one associates with a face to face sale (holding or seeing the goods for example) are not present in an online environment, one must rely on other mechanisms to establish some degree of trust between the two parties involved before any transaction can occur. A Reputation System serves as a mechanism to establish this degree of trust. In the case of online trading systems, it will most likely be the case that a buyer has no idea who the seller is or in what condition the goods being sold are in (other than the condition that the seller says they are in).

With a Reputation System in place, the potential buyer can get an accurate idea as to how trustworthy the seller is based on his reputation. A seller's reputation is made up of the opinions of buyers the seller has dealt with over a period of time. Opinions are in the form of feedback after a sale. In the case of a binary feedback system[1], feedback can either be good (+1) or bad (-1) (neutral feedback (0) is also considered part of a binary feedback system).

A seller is said to have a better reputation than another seller if the ratio of positive to negative feedbacks is better. As an example, consider sellers Alice and Bob. Alice has conducted 500 transactions with only 1 negative feedback. Because of the ratio of negative to positive feedback Alice has a good reputation.

[1] eBay.com uses a binary feedback system.

Bob has also conducted 500 transactions but has 50 negative feedbacks. Although 90% of Bob's feedback is positive, and he is probably a trustworthy seller, he is not as trustworthy as Alice.

Reputation Systems can be divided into two types of architectures: Centralised Reputation Systems and Distributed Reputation Systems. Centralised Reputation Systems have a central authority responsible for the construction of a reputation for a user (eBay is an example of such a system). A Distributed Reputation System is one in which agents participating in the system are responsible for the construction of a users reputation.

2.2 Ballot Stuffing

One of the aims in the experiment conducted by Resnick et al [9] was to determine if a user with a high reputation would be more successful (as a seller) than a user with a low reputation. The results showed that in a centralised, binary Reputation System the likes of eBay, a user with a high reputation is likely to conduct more business transactions than a user with a low reputation. Increasing one's reputation is therefore of interest to any serious user of a trading system. Unfortunately, users keen on increasing their reputation by illegitimate means may be able to do so via a process referred to as ballot stuffing [3, 4].

The process of ballot stuffing is simple: a user wishing to increase his reputation would collude with willing buyers to buy false items from him. Upon sale of the items (usually for a miniscule amount of money), each of the colluding buyers would give the seller a +1 rating for the transaction. The more colluding buyers involved, the higher the reputation of the seller after the transactions have been completed. Bhattacharjee et al [1] point out a similar process where buyers collude to provide negative feedback to a seller (possibly to make it easier for rival sellers to compete), this process is referred to as bad mouthing. In their paper, they show that a Reputation System can be made to resist ballot stuffing and bad mouthing as long as transaction costs are incurred for the seller (as is the case with eBay). Within this field of research, Dellarocas [3] proposes several mechanisms to thwart ballot stuffing using controlled anonymity and cluster filtering.

Douceur [6] discusses an attack of a similar nature i.e. a user assuming multiple identities in a peer-to-peer environment. In this paper the term "Sybil attack" is coined when referring to this type of attack. The author concludes that in the absence of a central authority, "Sybil attacks" will always be possible.

2.3 Anonymising Networks

There are a number of techniques that can be used to obtain varying degree's of anonymity on the Internet. A relatively cost-effective and easily configurable solution lies in the usage of anonymous proxies [2]. One is effectively anonymous to the end server being visited when traffic is passed through an anonymising

[2] http://anonymizer.com

proxy. Unfortunately, although some degree of anonymity is offered from the end server, this is not the case from the proxy itself.

To circumvent this problem there are a number of privacy enhancing technologies that offer anonymity from an end server as well as from the entities used to obtain the anonymity (in most cases, this is a network of machines). Most of these technologies are based, in one way or another, on the mix proposed by Chaum [2]. Chaum effectively employs the use of mix machines to delay the delivery of encrypted messages so as to hide the source a message. Wright et al [11] provide an overview on a host of different mix technologies as well as on the current state of research within the field of anonymity.

Tor[3] is an example of an anonymising network that is based on second generation onion routing [5, 8] (which in turn is based on the mix). Developed by the Naval Research Laboratory and the Free Haven Project, Tor effectively provides a high degree of unlinkability between the sender and receiver of a message even in the case of compromised mixes. With approximately 450 server nodes participating in the Tor network [7], Tor makes for a good candidate through which to commit the ballot stuffing scheme.

The JAP[4] network is similar to Tor in that it is also based on the mix. What makes it slightly different though, is that it only offers a limited selection of static mix machines (mix cascades) through which requests can be routed. As a result, many users of the JAP network are likely to share a single static IP address. This characteristic makes the JAP communication network an unlikely candidate in the ballot stuffing scheme.

3 A Network of False Buyers

Within the field of Reputation Systems, the following entities in a ballot stuffing scheme are involved :

- A user wishing to increase his reputation. This would typically be done by increasing the number of positive feedbacks he has had from sales.
- A number of buyers that are willing to collude with the user so as to increase his reputation.

The relationship between these two entities is said to be a one to many relationship i.e. there are a number of colluding buyers for a single seller. With the increasing usage and availability of anonymising networks though, we argue in this paper that the relationship between a seller and buyers need not be a one to many relationship. An anonymising network could allow for a fraudulent buyer to setup a separate network of agents (acting as buyers) that fall within his control. Though there are still a number of buyers, there is effectively only one person behind the buyers.

But what does it take to be a buyer? In general, setting up an account on an online trading system requires the following from a user:

[3] An anonymous Internet communication system - http://tor.eff.org/
[4] Anonymity & Privacy - http://anon.inf.tu-dresden.de/index_en.html

1. Personal details the likes of his full name and address. This can be faked.
2. Credit card details. These can be setup legitimately or even stolen. In some trading systems, one can use the same credit card for each of the accounts, possibly due to legislation prohibiting the storage of credit card details in plain text.
3. A unique IP address. This is not always a prerequisite, but is sure to raise alarms if multiple accounts are registering/bidding using the same IP address.

With the first two prerequisites addressed, a fraudulent buyer can use one of several anonymising networks on the Web when accessing the trading system so that each account looks as though it is using a unique IP address. Figure 1 depicts the simple process of a user (with a single IP) accessing a trading system via the Web. Figure 2 then moves on to depict the same user making use of an anonymising network to access the trading system via the Web. Note that although the user still only has one IP address, his usage of the anonymising network makes it seem to the trading system that it is being accessed from a number of different IP addresses.

Fig. 1. A single user making use of the trading system.

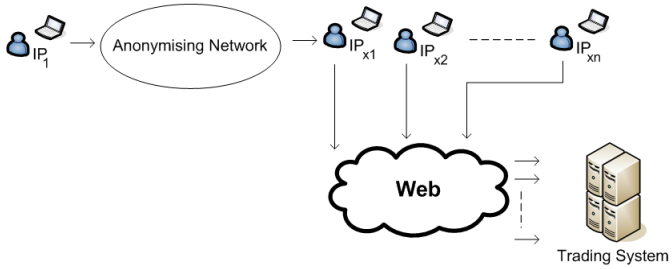

Fig. 2. A single user employing the services of an anonymising network to use the trading system.

Within the context of ballot stuffing, a single buyer controlling a number of agents on the trading system is sure to experience less complexities than a buyer

colluding with a number of other buyers. This single buyer, with a network of buyers under his control, may also find it easier to sell the idea of ballot stuffing to potential clients. The idea behind ballot stuffing is simple. With a network of agents behind the single buyer, implementing ballot stuffing may potentially be simple as well.

4 Proposed Framework

The intention of this paper is to propose a framework that will contribute to thwarting ballot stuffing from users who are employing the usage of a false network of buyers i.e. using an anonymising network. This framework is not necessarily specific to Reputation Systems, it may prove useful in several areas, for example, in social networking communities where users vote on the popularity of movies, news items or various Web sites.

The framework is based on the unlikelihood of users involved in a transaction being within the same physical area or zone. With trading systems the size of eBay, this is a reasonable assumption. It is highly unlikely that a business transaction will involve a seller and several buyers that are within the same physical zone (say, within 500 meters of each other or even the same city).

Before users of the system complete a bid or cast a vote, it will be the responsibility of the Reputation System to determine in what zone the current user is in. If the user is in the same zone as other users involved in the transaction, the transaction is flagged as being in one of three states (depending on the range the user is from the other users):

Highly suspicious - applicable to users that are within 500 meters of one another.
Suspicious - applicable to users within the same suburb.
Possibly suspicious - applicable to users within the same city.

4.1 Determining Zones

What is imperative to the framework is a means of determining the location of a user i.e. what geographical zone a user is in. This can be achieved by each user carrying around or being within close proximity to a device of some kind. In determining the location of the device, the framework must be relatively certain that the user is indeed within close proximity to the device.

The initial design of the framework had the users employing tags similar to those used in vehicle tracking systems[5]. This type of tracking system uses the Global Positioning System to pinpoint the position of an object. Since this is an inexpensive means of tracking ones location, this technology would suit our framework quite well. Unfortunately, if the framework employed this technology it would be susceptible to a very simple trick: a fraudulent user could purchase

[5] GPS Vehicle Tracking - http://www.networkcar.com/

multiple tags and scatter them at various intervals during a trip to an international conference. As a result, the Reputation System would register each of the tags in different zones and would not raise any alarms.

What was needed was another inexpensive mechanism that was not susceptible to this trick and yet still offered an easy way to look up one's location. The mechanism we propose is that of cell phones.

4.2 Cell Phones

Cell phones not only provide the ability to look up one's location but also the ability to directly communicate with the user in a manner that does not depend on a central authority. There is a plethora of cell phone providers that offer cheap access to multiple networks world wide.

By communicating with a user via the cell phone, the system can be relatively certain that the user is in the same location as the cell phone. To take advantage of this feature, we therefore propose that the bidding process consist of two stages:

1. The normal bidding process.
2. The authentication process. This serves as a means to authenticate the validity of the bid.

In step 2, the Reputation System would send a random string to the bidder's cell phone, the bidder would then have to enter this string online in order to complete the bid. This process is similar to the *two factor authentication* system[6] employed by several banks that offer online services[7]. The process employed in the framework proposed in this paper takes the authentication system further by looking up the location at which the cell phone is currently situated.

Figure 3 illustrates the concept. Step 1 denotes a bidder making a regular bid for an item in an online sale. The process is extended by steps 2 (a) and (b). In step 2 (a) the trading system looks up the location of the phone associated with the bidder. If it is not within one of the zones that any other entity involved in the sale is in, then the transaction is not marked as suspicious. To confirm that the bidder is near his phone (and therefore within the zone the system thinks he is in) a message is sent to the phone. In step 2 (b) the bidder receives the message and inputs the token online to validate and complete the bid.

Of course, the authentication process of step 2 could be foiled by writing an application on a cell phone that simply forwards the random string to the user in question. This problem may be circumvented by introducing a slightly different process. Instead of the user typing in the string online, he may have to use the phone itself to input the string by calling through to an automated digital helpdesk. Whilst this process ensures that there is someone behind the phone, it does not ensure that the person making the bid is with the phone. The

[6] http://en.wikipedia.org/wiki/Two-factor_authentication

[7] http://www.smh.com.au/news/Breaking/NZ-bank-adds-security-online/2004/11/08/1099781306318.html

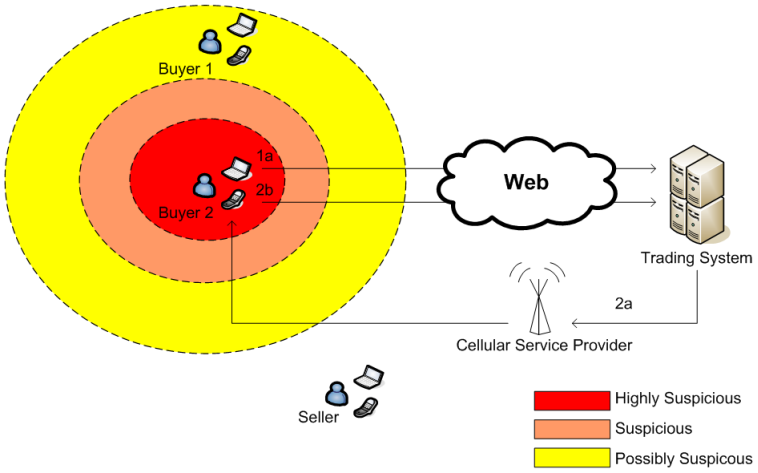

Fig. 3. The proposed bidding process.

fraudulent bidder could have given several phones to a number of people around the country that are participating in the scheme. This approach however brings with it more complexity (since there are more people involved) and defeats the original intention of a one man ballot stuffing scheme.

From a privacy perspective, the framework may have to be subjected to a number trade-offs. For maximum user privacy the framework would be prohibited from storing the coordinates of a user each time his location is looked up. Although this is beneficial from a privacy point of view, it gives the framework a far more limited scope from which to make decisions. To balance the effects of privacy vs functionality, one might employ a time limit on the amount of time that a user's location could be stored.

Of course, it could be argued that the framework does not need to store or know the locations of the users at all. This could be placed in the hands of a trusted third party. The framework would use the service to ask whether the phone X is within a specified range of phone Y. The trusted third party would simply provide a yes or no response.

4.3 Thwarting Anonymous Networks

As noted earlier, the authentication process described in this paper is not a new idea. This process has been employed by banks in a much simpler form for a number of years. Unfortunately, the process employed by the banks alone will not suffice in thwarting users of an anonymising network. Simply using the cell phone as a means to communicate with the user (by sending him a random string) only ensures that the user making the bid is the owner of the phone associated with the user. As mentioned in the previous section, an illegitimate

user wishing to commit ballot stuffing through an anonymising network would foil this process by purchasing a large number of cell phones.

Since the process employed by our framework includes the act of looking up the zone in which each user's cell phone currently resides, the illegitimate user can not simply purchase a number of cell phones. This additional step makes it far more complicated for a user looking for a simpler way in which to commit ballot stuffing; the alternative being the complex management (and payment) of a number of users keen on committing ballot stuffing.

5 Masquerading Through Anonymous Networks

Although the framework proposed in this paper addresses the problem of a one man ballot stuffing scheme, we argue that the underlying solution may be used to thwart a number of similar schemes. The ballot stuffing scheme discussed in this paper essentially depends on a single user that is able to present himself as many users by masquerading through an anonymising network. By looking up the location of the user in addition to comparing it to the locations of other users involved in the transaction, the trading system is able to make a fairly accurate decision as to whether or not the transaction is in some way fraudulent.

The framework in this paper effectively proposes a technique to recognise a single user despite the fact that he may be using an anonymising network. As a result, schemes that are dependent on single users masquerading as many users may be circumvented with the framework proposed.

An example of such a scheme is that of ticket scalping[8]. The online version of this scheme is as simple as ballot stuffing and involves the increase of ticket prices (to a ball game or show of some kind) by buying as many of the tickets as possible and selling them privately. With only a few people holding a large number of tickets, demand for the remaining tickets increases and so does the price. In the real world, there is a predefined limit as to how many tickets can be purchased by a single person for a single event. The scheme therefore requires a number of colluding fraudsters in order for it to be successful.

Unfortunately, in the online world, ticket scalping is well within the reach of the one man operation. This is partly due to the ease with which one can switch identities (with help from anonymising networks) in addition to being able to automate the entire procedure (from buying the tickets to auctioning them).

With the use of the framework proposed, circumventing online ticket scalping would only require a few minor adjustments:

1. The parameters of each zone must be reexamined. It is more than likely that there will be a number of people from the same city buying tickets to a football game within the city. Zones must therefore be more fine grained. The same idea applies i.e. it is unlikely that a thousand tickets will be purchased by a thousand people that are within 150 meters of one another.

[8] http://www.oag.state.ny.us/press/reports/scalping/full_text.html

2. The authorisation process of the framework would then be included in the online ticket purchase. Before confirmation of the sale of a ticket, the system confirms the zone of the buyer. If it is within very close proximity to a predefined threshold of other buyers then the transaction is tagged accordingly.

6 Conclusion

In this paper we have discussed the problem of ballot stuffing and suggested that it may be implemented by a single user with the aid of an anonymising network. This type of network allows a user to pose as many individual users, and in doing so, makes it easier to commit schemes the likes of ballot stuffing without a network of colluding fraudsters.

The framework proposed assumes that it is highly unlikely for users of a transaction in a global system to be within close proximity to one another. It then attempts to deal with the problem of ballot stuffing by looking up the physical location of a user in a transaction and comparing it to the location of other users involved in the same transaction. The transaction may then be tagged according to the proximity of each of the users to each other.

Using cell phones to lookup the location of each user was shown to be a better alternative because it ensures that there is someone next to the device. Since the focus of this paper is that of a single user using an anonymising network it is most likely that the person behind the phone will be the fraudulent user in question.

Although the framework proposed has been presented as a means to thwart ballot stuffing, the underlying problem is that of using anonymising networks to present a single user as many users. Since the framework is able to identify the location of a user regardless of the means used to access the system, it may be the case that the framework is applicable to a number of scenarios. At the very least, it is applicable to the ticket scalping scheme.

References

1. R. Bhattacharjee and A. Goel. Avoiding Ballot Stuffing in eBay-like Reputation Systems. Economics of peer-to-peer systems, 2005.
2. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. Commun. ACM 24, 2 (Feb.), 84-88, 1981.
3. C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. Proceedings of the second ACM Conference on Electronic Commerce, 2000.
4. C. Dellarocas. Mechanisms for coping with unfair ratings and discriminatory behavior in online reputation reporting systems. International conference on Information Systems, 2000.
5. R. Dingledine, N. Mathewson and P. Syverson. Tor: The Second-Generation Onion Router. Proceedings of the 13th USENIX Security Symposium, 2004.

6. J. Douceur. The Sybil Attack. Proceedings of the IPTPS02 Workshop, Cambridge, MA (USA), 2002.
7. L. Overlier and P. Syverson. Locating Hidden Servers. Proceedings of the 2006 IEEE Symposium on Security and Privacy, 2006.
8. M. Reed, P. Syverson and D. Goldschlag. Anonymous connections and onion routing. IEEE J. Selected Areas Commun., 1998.
9. P. Resnick, R. Zeckhauser, J. Swanson, and K. Lockwood. The Value of Reputation on eBay: a Controlled Experiment. Working paper, University of Michigan, 2002.
10. M. K. Reiter and A. D. Rubin. Crowds: Anonymity for Web Transactions. Communications of the ACM, Vol. 42, No. 2, pp. 32-48, 1999.
11. J. Wright, S. Stepney, J. Clark and J. Jacob. Designing anonymity - a formal basis for identity hiding. Internal yellow report, York University, York, UK, 2004.